Title
A Search for Dark Matter and Investigations in Discrimination Between Neutrons and Gamma-rays

Author
Johnson, Jyothisraj

Publication Date
2022

Peer reviewed|Thesis/dissertation

A Search for Dark Matter and Investigations in Discrimination Between Neutrons and Gamma-rays

By

JYOTHISRAJ JOHNSON

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Physics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

S. Mani Tripathi, Chair

Michael Mulhearn

Eric Prebys

Committee in Charge

2022

Abstract

Several independent studies have provided evidence supporting a cold, non-baryonic, non-luminous component of the universe, which is colloquially referred to as dark matter. Weakly interacting massive particles (WIMPs) are a well-motivated candidate to describe the particle nature of this component of the universe. There are several approaches to searching for evidence of these particles. This dissertation focuses on LUX-ZEPLIN (LZ), a direct dark matter search using a dual-phase xenon time projection chamber (TPC). Standard searches for WIMP recoils off xenon nuclei assume a spin-independent or spin-dependent interaction in the zero-momentum limit. This work focuses on a more general interaction basis derived from a non-relativistic effective field theory (EFT) approach to WIMP-nucleus interactions. More specifically, we set limits on the coupling strengths for the fourteen operators that form the basis of inelastic WIMP-nucleus interactions in the non-relativistic regime. The limit-setting analysis was conducted using an extended, unbinned profile likelihood ratio (PLR) method. World-leading limits are achieved for the lower mass-splitting values considered in the analysis.

The second half of this dissertation expands on a key requirement in such dark matter searches: discrimination between nuclear (NR) and electronic recoil (ER) events. One of many methods to achieve this capability is pulse shape discrimination (PSD). Although xenon is not a suitable detection medium to take advantage of this technique, plastic/organic

scintillators exhibit this capability. This subclass of scintillators are widely used in the radiation detection and nuclear security fields. This dissertation presents the design of and results from a custom SiPM + Scintillator test bed. The setup was used to study the PSD capability of commercially available plastic/organic scintillators when optically coupled to silicon photomultipliers (SiPMs). Also discussed are the results of a feasibility study, conducted using data taken with the test bed, on the design of a prototype fully custom application specific integrated circuit (ASIC) developed to read out signals from plastic/organic scintillators using SiPMs and provide real-time PSD via a custom analog circuit design on the ASIC. The end goal for this ASIC is integration into a compact, portable, and segmented neutron scatter camera design. The design, simulation, and initial testing results for the chip are presented. Condensed versions of the data sheets are provided as appendices for both versions of the chip that have currently been fabricated.

To my parents,

who immigrated to the United States 21 years ago with five-year-old me and my ten-month-old brother in tow and successfully made a new life in a completely foreign country while raising young kids. Thank you for encouraging me to pursue my interests and supporting me in paving my own path in life.

Contents

List of Figures

List of Tables

Acknowledgments

There are a lot of people who have supported and encouraged me to get to this point. I know that if I try to name everyone individually, I will end up forgetting to mention someone. So, I'm going to keep this short and sweet and I apologize if I didn't mention you by name.

I want to start by thanking Mani Tripathi, my advisor, for encouraging me to pursue graduate school in the first place. I have to admit, there were several times throughout my five years as a graduate student, especially during my first couple of years, where I wondered if I had made the right choice. Looking back now though, I can honestly say that it was all worth it in the end. You were always supportive of the fact that I didn't want to continue in academia and even set up non-traditional (for physics) research opportunities for me, such as two summer internships with the IC Design Group at Berkeley Lab. Thank you for all the opportunities and for always having my back.

Speaking of the Berkeley Lab IC Design Group, I want to thank Carl Grace, who has been a second advisor to me over the past two years. I really appreciate that you were always willing to answer all my beginner questions and that you always had the perfect resource to point me to whenever I asked. Thank you for teaching me the basics of IC design. I've learned a lot and I'm looking forward to learning a lot more in my new role.

Next up, I want to give a big shout out to the many members of the HEP postdoc mafia who I've worked with. You all were great. I want to especially thank Jon Balajthy for taking me under his wing and teaching me the ins and outs of conducting a physics analysis

and the many crash courses on dark matter, the LZ detector, xenon physics, etc over the years. I want to also say thanks for letting me be your deputy for the entire LZ analog electronics QA/QC testing procedure at UC Davis. That experience has given me a solid foundation for when I went on to design my own IC. I want to thank Ben Godfrey, the newest addition to the HEP postdoc mafia, who as a graduate student was always my go-to person for any issues I was up against. Thanks for always being willing to tackle and debug the most random problems that I've come across throughout graduate school. I also want to thank Billy Boxer for taking it upon himself to set me up with a last minute SR1 analysis when it became obvious I needed to accelerate my graduate school timeline because of a job opportunity. I know the 8-hour time difference has made work-life balance very difficult for you, so I really appreciate all the quick replies regardless of the time. Thanks for always being willing to provide support whenever I needed.

I can't forget to mention the UC Davis Boxing Club and the many friends I've made being a member over the last five years. You've all played a big part in keeping me sane and letting my life not completely revolve around physics. On top of the workouts, it's been an amazing journey going from a boxing newbie to "assistant" coach.

Now, moving outside of Davis, I gotta acknowledge the Lindo Squad. I'm not gonna list you all out by name but every visit back home to NYC wouldn't be the same without hanging out with all you crazy people. Props to y'all who actually made it out to Cali to visit me. Here's to the future where we are all going to be kicking back at one of our future summer vacation homes. Can't forget about the "Nerd Squad" either. I know we all getting

advanced degrees in various fields but graduate school is still the same no matter where or what we are doing, as we've all found out. Here's to everyone being out soon enough. We also going to be chilling in those future vacation homes together.

In addition to my parents, I want to thank my brother, Karun, for always checking in on me during graduate school. I'm glad you managed to visit me out here in the boondocks with Chris and Christina as well. On that note, I want to say thanks to all my cousins who have all been supportive of me. I know we don't see eachother as often as when we were kids. Back then, life was a seemingly never-ending series of sleep-overs, pool parties and family functions. I'm glad that, at very least, we can all catch up in the never ending series of cousin weddings.

Last, but definitely not least, I want to say thank you to Urvi. Thanks for being an amazing pillar of support over the last four years. I don't think I would have survived graduate school without you always being there for me. Our road trips and random adventures have been put on hold for a while but its time to get back to it. The world awaits.

Chapter 1

Introduction to Dark Matter

1.1 Missing Matter in the Universe

Several independent studies have provided evidence in support of a cold, non-baryonic, non-luminous component to the universe. Such evidence comes from the cosmic microwave background (CMB) [1][2], ultra diffuse galaxies [3], gravitational lensing [4], large scale structure of the universe [5], rotational curves of galaxies [6], and interacting clusters [7]. Of the above mentioned, ultra diffuse galaxies, the CMB, interacting clusters and rotational curves provide some of the strongest evidence to date for this dark matter (DM).

In the ensuing, we shall describe in detail three of the most compelling observations in order to provide an introduction to missing (dark) matter in the universe. Further, we will discuss weakly interacting massive particles (WIMPs) as a possible candidate to solve this puzzle. From these studies, we can summarize that there is good evidence to suggest that

DM is produced cold (non-relativistic), is not composed of Standard Model (SM) particles, probably does not clump on "small" scales (i.e. planet to stellar scales), and it played a vital role in the early universe and structure formation.

1.1.1 Galactic Rotation Curves

Rotational curves are defined as distributions of the variation in orbital velocity of stars and gas clouds as a function of radius from the galactic center. The profile of radial variations for a galaxy directly reflects the mass distribution within it. For most galaxies, which have visible mass concentrated towards the center, we should expect to observe the velocity decrease from its peak as a function of $\sqrt{\mathrm{r}}$. With the discovery of the 21 cm hydrogen line in the mid-twentieth century, we can actually compare predicted galactic rotation velocities with those measured through observations of this line [8]. For several galaxies, the observations showed a peak in the velocity curve followed by a relatively flat section past its optical edge. A flat rotation curve over a specific subsection in radii indicates that mass is increasing. Thus, applying this logic to the observed rotation curves of galaxies, the matter distribution must be increasing past the visible edge of the galaxy. Figure 1.1 shows a distribution of observed orbital velocities for galaxy NGC 3198, which clearly shows an example of this behavior [6]. Also shown is the assumed mass contribution to the galaxy from a dark matter halo surrounding the galaxy that would be needed to explain the observed rotation curve.

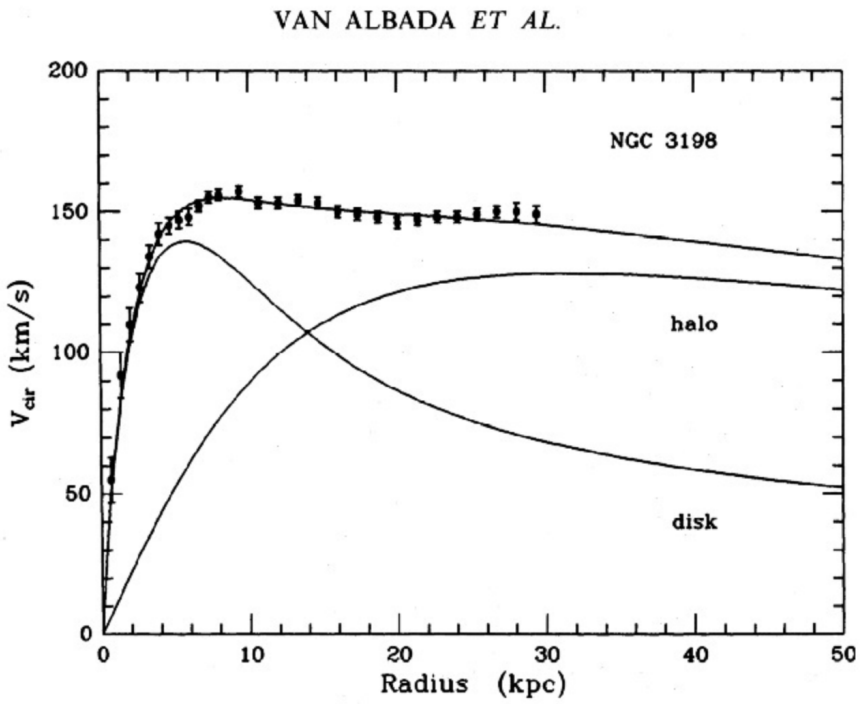

Figure 1.1: An example of a fitted rotational curve (for NGC 3198) plotting observed rotational velocity as a function of radius from galactic center. The disk and (required) halo contributions are also plotted. Figure from [6].

1.1.2 Interacting Galaxy Clusters

Another powerful evidence for dark matter is derived from observing interactions of colliding galaxy clusters. The most famous example is colloquially known as the Bullet Cluster (see Figure 1.2). It is a specific set of two colliding galaxy clusters observed in 2006. During the collision, we would expect the stars and the hot gas portions of the respective clusters to pass by one another at varying speeds. The individual stars are sparsely distributed, and interact only via gravity. Thus, they are not slowed considerably as the two clusters pass

by each other. The hot gas, however, which represents most of the baryonic matter in the clusters interact electromagnetically, causing them to slow down much more than that of the stars during collision.

Now, if that was the whole story, then gravitational lensing studies of the Bullet Cluster should provide clear evidence that lensing is strongest near the gas as the individual stars only form a small fraction of total mass within a cluster. Gravitational lensing is a phenomenon resulting from a massive foreground galaxy/galaxy cluster having the critical mass, and consequently, the required gravitational potential to bend the path of light from a source (i.e. a galaxy, quasar, etc) located directly behind it along the line of sight. Famously, this can produce multiple images of the source as seen from the observer's point of view [4]. Lensing effects are directly correlated to the mass of the object that acts as the lens. Returning to the Bullet Cluster, it was actually observed that lensing was strongest in the two separated regions near the visible portion of the two clusters. Thus, we can reasonably conclude that most of the mass in the clusters, constrained in two regions, must be composed of dark matter that only interacts via a weak force (and gravitation but not the electromagnetic force) during the collision [7].

1.1.3 The Cosmic Microwave Background

The cosmic microwave background (CMB) is an almost uniform distribution of microwave radiation that permeates the current universe. Most of this radiation was initially released

Figure 1.2: A reconstructed image of the Bullet Cluster. Shown in pink is the x-ray distribution of hot gas. Shown in blue is the derived matter distribution profile of the two galaxy clusters. Figure from the Chandra X-ray Observatory.

in the visible/UV portion of the electromagnetic (EM) spectrum but was red-shifted to the microwave band by the expansion of the universe. This radiation background was formed approximately 380,000 years after the beginning of the universe. Previous to this moment in time, the universe consisted of hot, ionized gas (plasma). This gas was almost entirely uniform, except for slight $O(10^{-5})$ deviations due to quantum fluctuations [9]. The pockets of higher density, via gravity, pulled in more surrounding matter. As mass accumulated in these pockets, it would heat up, and create outward radiation pressure. These competing processes evolve with time, and are formally described as baryonic acoustic oscillations (BAO) [10]. Eventually, the universe, as it continued to expand, had cooled enough that neutral atoms

could form. At this point, radiation emitted was no longer continually being absorbed and re-emitted and these photons were free to escape. Therefore, they contain the last imprint of these density fluctuations within the gas, which are presently seen as hot and cold spots in the temperature map of the CMB as shown in Figure 1.3. Mapping of the density fluctuations have been performed by various space missions over the last 30 years including, most recently, the Planck mission [1][2].

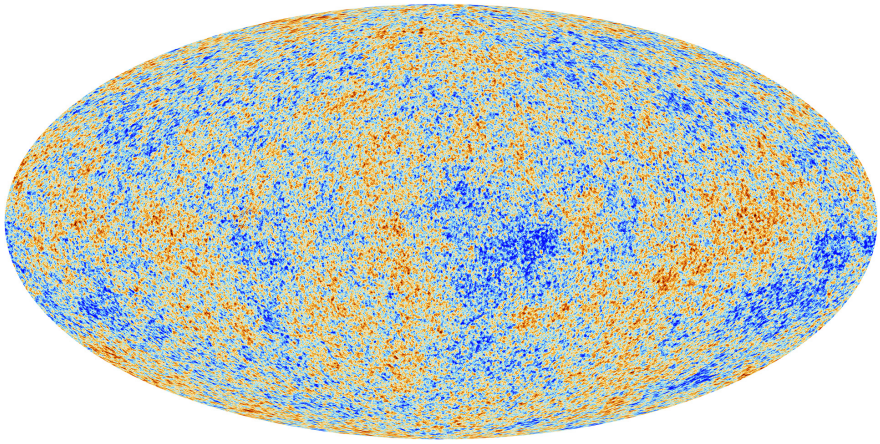

Figure 1.3: A reconstructed image of the cosmic microwave background from the Planck survey. Warmer (redder) colors show hot spots and cooler (bluer) colors show cold spots in the background.

The evidence for a dark matter component to the universe is derived from the extracted multi-pole power spectrum of the CMB distribution, as shown in Figure 1.4. There are three prominent peaks observed. The first peak in the spectrum gives evidence for a flat geometry to the universe. The second peak can be associated with the baryonic density of the universe and the third, to a combination of baryonic and dark matter density [2]. Thus, the ratios

between the second and third peaks provide a direct measurement for dark matter density composition of the universe. Simulations of the evolution of the universe cannot explain the ratio of the second and third peaks in the spectrum without including dark matter [11]. Planck's 2018 results found that approx. 31.5% of the energy-matter density of the universe is dark matter and only 4.9% is baryonic matter. The rest is dark energy [1], which is in itself a perplexing phenomenon, but is not the topic of study in this dissertation.

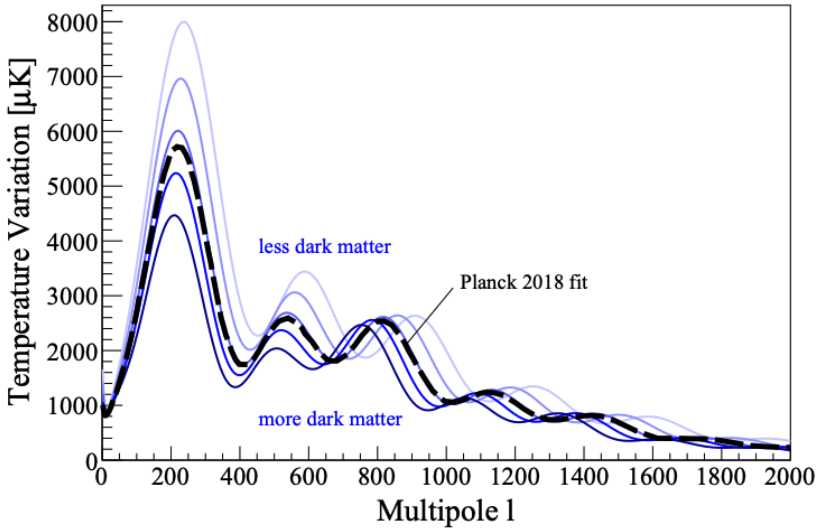

Figure 1.4: The calculated power spectrum from Planck is shown in the dashed black line. Simulation results varying dark matter density and the observed shift in peaks is shown. Lighter purple shows less dark matter and darker purple shows more dark matter. Figure from [11].

1.2 The WIMP Model for Missing/Dark Matter

1.2.1 Background on WIMPs

All of this overwhelming evidence for a dark component to the universe has prompted an enormous set of theoretical predictions for possible particle candidates that constitute dark matter. Due to our ignorance as to the mass of DM particles, it can range anywhere from 10^{-22} to 10^{28} eV. The multitude of theories for dark matter cover various regions of this parameter space [12]. However, this is an incredibly large range for any single experiment to explore. Thus, we need to settle on a well-motivated theory that can be probed with a reasonable experimental setup.

One of the most popular conjectures for dark matter (and what we will focus on) is that of weakly interacting massive particles (WIMPs). The preferred mass range for these candidate particles are between approx. 10 - 1000 GeV [12]. The theory postulates that WIMPs would have been produced in thermal equilibrium in the early universe and exist today as thermal relics.

In equilibrium, we expect production and annihilation rates of WIMPs in particle-antiparticle collisions would be equal because for T $>>$ m_χ (mass of WIMP), the kinetic energies are much higher than any mass production threshold. As the temperature of the universe decreased with time, this condition would at some point become invalid and the number of WIMPs produced would have decreased as $e^{-m_\chi/T}$ (the Boltzmann factor) [13]. During that phase, only the Standard Model (SM) particles, with much smaller masses, with

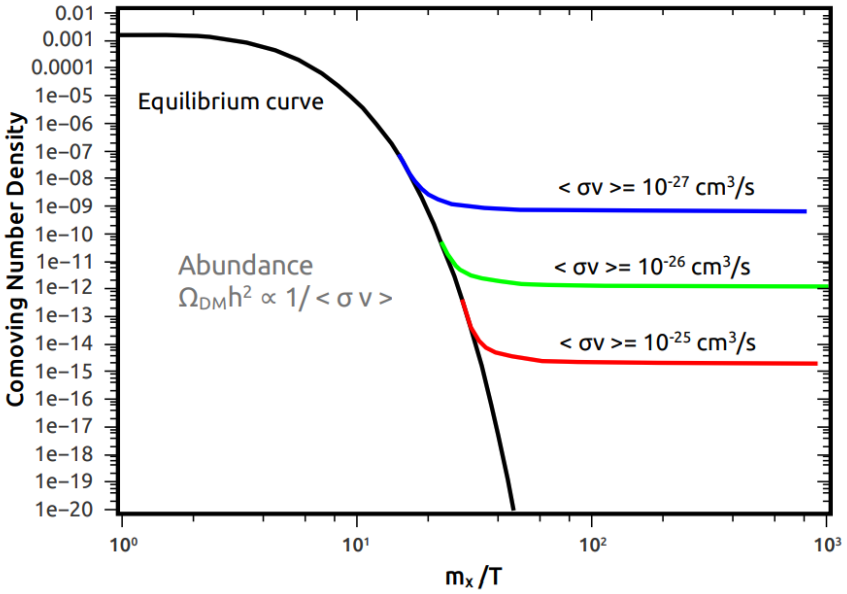

Figure 1.5: The co-moving number density is plotted as a function of time (average temperature) of the universe. Different thermal relic densities for WIMPs are shown for various assumed thermally averaged interaction cross sections with the green curve closest to the calculated relic density. Figure from [13].

energies in the tail of the distribution had enough kinetic energy to produce WIMP pairs in particle-antiparticle collisions. At the same time, as the universe expanded, the number density of all particles also decreased. Thus, we expect that the annihilation rates for WIMPs also decreased with time. When this rate became less than H, the Hubble constant, or equivalently stated, the mean-free path for WIMP-producing collisions became longer than the Hubble radius, WIMP production ceased and we were left with a constant thermal relic density of WIMPs in the universe. This is shown in Figure 1.5 with the various colored curves representing different possible thermally averaged WIMP interaction cross sections, $< \sigma\nu >$, and the corresponding respective final relic density. The actual relic density has

been measured as $\Omega h^2 = (3 * 10^{-22} cm^3 s^{-1}) / < \sigma\nu >$, which is $\approx 0.1198 \pm 0.0015$ [14]. To match this, the required thermally averaged cross section directly implies an interaction cross section consistent with the Weak interaction.

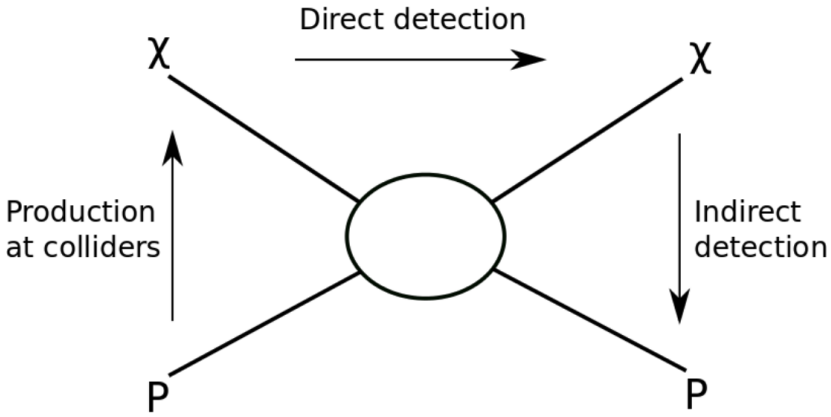

Figure 1.6: The various possible detection channels for dark matter,from [15].

Potential detection of WIMPs involves exploiting possible interactions with the SM. There are three possible interaction mechanisms, all derived from an assumed interaction, as shown in Figure 1.6. The arrows indicate the three rotations of initial state particles, and are labeled accordingly. Each will be described in the proceeding sections below. Indirect and collider searches for WIMPs will only be touched on briefly.

1.2.2 Experimental Searches for WIMPs

1.2.2.1 Indirect Detection Searches

Indirect detection searches for WIMPs are conducted by looking for evidence of SM products from dark matter pair annihilation or decay. Because of the small measured dark matter density in the galaxy ($\rho_0 \approx 0.2 - 0.6$ GeV/cm^3 [16]), such searches are constrained to observations of regions which have been measured to have very high concentration of DM. These include dwarf spheroidal satellite galaxies of the Milky Way and the cores of large scale structures in the universe (i.e. galaxies and galaxy clusters). We expect, in general, an electromagnetic (EM) spectrum from such sources. Thus, indirect experimental searches look for anomalous excesses over the expected EM background. There are large uncertainties associated both with the background and signal models in these searches. Choices in the dark matter halo parametrization and, as a consequence, velocity distribution of dark matter in the halo along with assumptions on its thermally averaged cross section all play a large role in the expectation of the anomalous excess. In general, these experiments are conducted either with ground-based telescopes (IceCube; see [17], VERITAS; see [18], MAGIC; see [19], HESS; see [20]) or satellite-mounted space observatories (FERMI; see [21]). A very thorough overview of background and the history of experimental searches in this realm is given in [22].

1.2.2.2 Collider Searches

Searches for direct production of WIMPs have to be conducted at particle colliders. Given the small expected Weak-scale WIMP production cross-sections, the high energies involved in the collisions and high luminosities make it possible to have significant dark matter production rates. However, an important note is that collider searches cannot definitively provide proof of WIMPs, because they are detected via signatures consisting of imbalanced transverse momenta. There are many other phenomena that can explain such signatures. Thus, any hint of WIMP production will need to be cross-checked and verified against direct detection and/or indirect detection results. Currently, there is a substantial portfolio of dark matter candidates being probed (including WIMPs) at the Large Hadron Collider (LHC). See [23] for an example.

The searches for missing energy are conducted in association with recoiling visible particles, which help with triggering on such events. Besides WIMPs, the LHC experiments are also engaged in searching for dark photons [24], which are conjectured to have mass and are the mediator of an additional U(1) symmetry. Thus far, all the attempts have achieved null results, however, vast regions of the parameter space have been excluded [25].

1.2.2.3 Direct Detection Searches

The main focus of this work is direct detection searches for dark matter. For this type of search, experiments are not directly trying to observe dark matter, despite the name.

Rather, the goal is to observe WIMPs (that are gravitationally bound to the local Milky Way galactic halo) scatter off target nuclei within a detector's active volume. Given that the escape velocity, v_{esc}, is usually taken to be 544 km/s for our galaxy [26], WIMPs are treated as non-relativistic in these searches. For most of the history of these searches, WIMP-nucleus elastic scattering was the main focus. However, recently, inelastic scattering models are becoming more popular [27]. We will discuss the latter in more detail in Chapter 3. In this section, we shall focus on a simplified theory for WIMP-nucleus elastic scattering. More specifically, we will examine spin-independent and spin-dependent elastic scattering in the zero-momentum limit (motivated by the expected non-relativistic velocities for WIMPs).

Generally, both electrons and nuclei can participate in the WIMP scattering interactions. However, WIMP-electron scattering is kinematically suppressed. On the other hand, WIMP-nucleus scattering is well matched for most commonly used detection target nuclei. From the experimental point of view, the main parameter of interest is the number of scattering events we expect to see and the expected energy spectrum of the produced recoils within the active volume. In other words, we need to find an equation for the expected differential rate as a function of recoil energy. We will derive this here following the steps given in [28].

Let us start with the kinematic equation for the angular dependence of the nuclear recoil energy deposited for a given scattering event:

$$E_R = \frac{1}{2} m_\chi v^2 \frac{4 m_\chi m_A}{(m_\chi + m_A)} \frac{1 + \cos\theta}{2}, \tag{1.1}$$

with E_R the recoil energy, m_χ the mass of WIMP, v the non-relativistic velocity of the WIMP, m_A the mass of the target nucleus (A) and θ the scattering angle. Maximum recoil energy occurs when θ equals 0 (head on collision) and m_A equals m_χ. Assuming 200 km/s for v and 100 GeV for m_χ, we get that the max energy transfer in such a scenario is 50 keV. We can now motivate why the region of interest (ROI) for direct detection experiments is usually in the range of O(1) - O(10) keV_{nr} energy depositions.

Now, to find an equation for the WIMP-nucleus differential rate, we can start with its standard relationship to the differential cross section, written as

$$\frac{dR}{dE_R} = nvN_T\frac{d\sigma}{dE_R}, \tag{1.2}$$

where R is the event rate, N_T the number of target nucleons, n the number density for WIMPs and v the WIMP velocity. However, it is important to take into account that WIMP velocities are usually assumed to follow a Maxwellian distribution. Thus, Eqn. 1.2 should be rewritten as an integral over an allowed velocity range as

$$\frac{dR}{dE_R} = nN_T\int_{v_{\min}} vf(\vec{v})\frac{d\sigma}{dE_R}d\vec{v}, \tag{1.3}$$

with $f(\vec{v})$ standing in for the range of possible WIMP velocities. The minimum velocity, $v_{\min}$, in the integral is defined as the minimum velocity needed for a WIMP with mass m_χ to produce a recoil energy E_R. Finally, converting from number of targets to a detector mass and substituting appropriately for n allows us to write the differential rate equation as

$$\frac{dR}{dE_R} = \frac{\rho_0}{m_\chi m_A} \int_{v_{\min}} v f(\vec{v}) \frac{d\sigma}{dE_R} d\vec{v} = \frac{2\rho_0}{m_\chi} \int_{v_{\min}} v f(\vec{v}) \frac{d\sigma}{d|q|^2} d\vec{v}, \tag{1.4}$$

with ρ_0 the WIMP local density. The second form for the differential rate is arrived at by substituting $q = \sqrt{2m_A E_R}$, where q is defined as the momentum transfer in the scattering interaction. The final piece of the puzzle is to define $\frac{d\sigma}{dE_R}$ or, equivalently, $\frac{d\sigma}{d|q|^2}$. In general, we expect that the differential cross section depends on interactions with the nucleus, the type of coupling and the detailed nuclear structure of the atom.

To start, we can relate the differential cross section to $|\mathcal{M}|$, the scattering amplitude, as

$$\frac{d\sigma}{d|q|^2} = \frac{1}{\pi v^2} |\mathcal{M}|^2. \tag{1.5}$$

We can then plug in the scattering amplitude relationship to a nuclear form factor, $F^2(q)$, in Eqn. 1.5 and get

$$\frac{d\sigma}{d|q|^2} = \frac{CG^2}{v^2} F^2(q), \tag{1.6}$$

where G is the standard Fermi constant and C a dimensionless number that carries particle physics model information. Now, due to the non-relativistic velocities involved, in this simplified framework for WIMP-nucleus interactions, only non-vanishing terms in the zero-momentum limit are kept. Thus, we can find an equation for the zero-momentum cross section by integrating Eqn. 1.6 in the $q \to 0$ limit. This can then be substituted back in to rewrite Eqn. 1.6 in terms of this zero-momentum cross-section as

$$\frac{d\sigma}{d|q|^2} = \frac{\sigma_{0,A}}{4\mu_{\chi,A}^2 v^2} F^2(q), \tag{1.7}$$

with $\sigma_{0,A}$ the zero-momentum WIMP-nucleus cross-section defined as

$$\sigma_{0,A} = 4G^2\mu_{\chi,A}^2 C \tag{1.8}$$

and $\mu_{\chi,A}^2$ the reduced WIMP-nucleus mass. Now, the next step is to replace the cross section with explicit forms for the spin-independent and spin-dependent interactions (which are the two interactions that survive in the zero-momentum limit). For this, it is easier to switch to a WIMP-nucleon framework rather than a WIMP-nucleus one. From [28], we can write the spin-independent WIMP-nucleon zero-momentum cross section as

$$\sigma_{0,N}^{SI} = \frac{4\mu_N^2}{\pi} \left[f_p Z + f_n (A - Z)\right]^2, \tag{1.9}$$

with μ_N the WIMP-nucleon reduced mass, A and Z the atomic mass and number, and f_n and f_p coefficients for n,p respectively. Assuming the scale of such coefficients are roughly the same for n and p, we can simplify to

$$\sigma_{0,N}^{SI} = \frac{4\mu_N^2}{\pi} f_N^2 A^2, \tag{1.10}$$

with f_N representing a general nucleon coefficient. The A^2 term in Eqn. 1.10 is known as the coherence enhancement in the zero-momentum limit spin-independent rate, and it

motivates the choice of a heavier target nuclei medium for a direct dark matter search. This can be related back to the WIMP-nucleus zero-momentum cross section as

$$\sigma_{0,N}^{SI} = \frac{\mu_{\chi,A}^2}{\mu_N^2} A^2 \sigma_{0,A}^{SI} \tag{1.11}$$

Using the above relationship, we can substitute Eqn. 1.10 into Eqn. 1.7 to get a final equation for spin-independent differential scattering. We can further substitute this version of Eqn. 1.7 into Eqn. 1.4 to give us a final spin-independent WIMP-nucleon differential rate as:

$$\left(\frac{dR}{dE_R}\right)_{SI} = \frac{2\rho_0}{m_\chi} \int_{v_{\min}} v f(\vec{v}) \left(\frac{A^2 \sigma_{0,N}}{4\mu_N^2 v^2} F^2(q)\right) d\vec{v}. \tag{1.12}$$

Now, let's tackle the spin-dependent case. We can go through a similar set of steps as above, following [28], and arrive at a zero-momentum WIMP-nucleon cross-section for this scenario as

$$\sigma_{0,N} = \frac{32}{\pi} G^2 \mu_N^2 \frac{J+1}{J} \left[a_p \langle S_p \rangle + a_n \langle S_n \rangle\right]^2 \tag{1.13}$$

with $\langle S_p \rangle$ and $\langle S_n \rangle$ the proton and neutron spin contents of the nucleons and J the total angular momentum. To arrive at this equation, two assumptions were made: we only allow spin 1/2 Majorana WIMPs and the total spin content, $\langle S_N \rangle$, of a nucleus can be attributed to the unpaired nucleon in a nucleus. Similar to the spin-independent case, we can then define a relationship between the WIMP-nucleus and WIMP-nucleon cross sections as

$$\sigma_{0,A} = \frac{4}{3}\frac{J+1}{J}\frac{\mu_{\chi,A}^2}{\mu_N^2}\langle S_N\rangle^2 \sigma_{0,N} \tag{1.14}$$

and the final differential rate can be written after substituting this into Eqn. 1.4 as

$$\left(\frac{dR}{dE_R}\right)_{SD} = \frac{2\rho_0}{m_\chi}\int_{v_{\min}} v f(\vec{v}) \left(\frac{1}{3\mu_N^2 v^2}\frac{J+1}{J}\langle S_N\rangle^2 \sigma_{0,N} F^2(q)\right) d\vec{v}. \tag{1.15}$$

The $F^2(q)$ in both differential rate equations represents the nuclear form factor, which encodes the underlying physics of how the WIMP resolves the nucleus in scattering interactions. We will discuss this in greater detail in Chapter 3. However, by convention, this is defined as the Helm form factor, in which the nucleus is treated as a solid sphere with uniform nucleon density [29], for the simplified zero-momentum limit framework.

It is important to note that although we have defined equations to model the differential rate (and cross-sections) of WIMP-nucleon scattering (in the zero-momentum limit), the validity of such assumptions is up for debate. Furthermore, ρ_0, m_χ, and $f(\vec{v})$ are either completely unknown quantities or quantities that have to be modeled resulting in large associated uncertainties. By convention, we assume a Maxwellian velocity distribution for WIMPs and a value for its density ($\rho_0 = 0.3$ GeV/c^3) within the Milky Way galactic halo to probe $\sigma_{\chi,\mathrm{N}}$-m_χ parameter space, which allows for easy comparisons of results between various direct detection experiments.

One other important note to make about direct detection experiments is that, at their core, they are rare event searches and as such, require careful mitigation and understanding

of expected backgrounds. In order to suppress cosmic ray induced backgrounds, most experiments are located deep underground and even then, contain other shielding measures such as water tanks and liquid scintillator veto volumes. This results in the reduction of the cosmic ray backgrounds to manageable levels. Conducting these experiments at the surface level is infeasible due to rates seen at the surface. Background radiation from detector materials and the surrounding underground cavern walls also needs to be mitigated and modelled. Water shields are primarily designed for the latter and furthermore, neutron veto systems are also often used. Detector materials, on the other hand, are an inescapable source of background that cannot be easily mitigated against. As a result, most experiments spend a considerable amount of effort in screening and selecting very low radioactivity materials for detector construction. Once the detector components are obtained, they are subjected to extensive assays in order to build a robust background model.

1.2.3 Other Dark Matter Candidates

While WIMPs are a leading candidate for explaining dark matter, there are several other theories to explain the question of missing matter in the universe. In order to not ignore them completely, we shall very briefly touch on two of the popular alternatives.

1.2.3.1 Axions

The axion is a pseudoscalar particle proposed primarily as a solution to the strong CP problem. One of the predictions of quantum chromodynamics (QCD) is that CP can (and

should) be violated in the Strong sector. However, to date, no experimental evidence of this exists. A key prediction is that the neutron should have an electric dipole moment and experimental probes have resulted in a much smaller value for this than predicted. As with most open questions in physics, there are no lack of theories to explain the source of the discrepancy. The standard explanation, the Peccei-Quinn theory [30], introduces the axion to explain the lack of CP violation in the Strong sector. The theory proposes that the axion is created through spontaneous symmetry breaking, which cancels out the CP violating term in the SM Lagrangian. Due to its very low predicted mass, it is fairly unlikely that axions can account for the full missing matter of the universe. A much more detailed overview of the theory is given in [31]. In addition, [32] presents an overview of one such experiment looking for these particles.

1.2.3.2 Sterile Neutrinos

In general, neutrinos are described by the Standard Model as massless, left-handed particles with three distinct flavours that are independently conserved. Evidence for both neutrino oscillation and mass implies that the Standard Model is incomplete in its formulation of neutrinos. One explanation for neutrino mass involves introducing a sterile neutrino. These massive, right-handed neutrinos, according to the description, have no interaction with other SM particles (no electric, strong or weak interactions) except through gravity. It only interacts with the standard three neutrino flavours through a proposed seesaw mechanism. A full theoretical description of the particle and overview of current searches is given in [33].

Chapter 2

The LUX-ZEPLIN Experiment

In this work, we shall be focusing on one particular direct detection experiment using the LUX-ZEPLIN (LZ) detector. By the end of the chapter, we will have covered in sufficient detail the design of this experiment. Further details will also be given in Chapter 4, while discussing analysis steps and presenting results using data from the so-called science run one (SR1). LZ is currently the world's largest xenon-based direct detection experiment looking for evidence of dark matter-nucleon interactions. It is housed nearly one mile underground at the Sanford Underground Research Facility (SURF) within the Black Hills of South Dakota. The earthen overburden provides approximately 4300 meters of water equivalent (m.w.e) shielding against cosmic rays for the experiment. However, accessing the experiment cavern is quite an ordeal. It takes approximately 10 minutes of straight vertical descent in an elevator (the "cage"), which is powered by a giant winch on the surface.

Before talking specifically about the LZ detector, it is important to have a brief discussion

on why xenon in particular was chosen as a target medium for the experiment. In doing so, we shall briefly cover the details of how liquid xenon responds to radiation-induced interactions, and the formation of resultant signals. Next, an overview of dual-phase time projection chamber (TPC) experiments will be provided, as this is the detection technique employed by LZ. Finally, the chapter will conclude with an overview of the main subsystems that come together to form the LZ detector. Particular focus will be given to the analog electronics subsystem as the group at UC Davis was primarily responsible for its design, fabrication, testing, integration into the overall detector, and commissioning.

2.1 Dual-Phase Xenon Time Project Chambers

2.1.1 Signal Formation in Xenon

Xenon is a noble element that has very few long-lived radio-isotopes. Considering the strict constraint on radioactive background in rare event searches, it is a perfect choice for detector medium. Furthermore, it has a high atomic mass, making it an excellent target to take advantage of the A^2 enhancement in the zero-momentum spin-independent differential rate. The high density of liquid xenon (3.057 g/cm^3) results in an excellent self-shielding property as well. Due to the relatively short attenuation lengths of gamma-rays in xenon, a few cm of self-shielding is sufficient to reduce contamination from background radioactivity (both detector materials and surrounding natural radioactivity of the experiment cavern) to levels

required for performing a rare-event search. Figure 2.1 demonstrates this self-shielding capability. The inner volume of liquid xenon is usually referred to as the fiducial volume and only energy depositions within it are considered for dark matter searches. Further, liquid xenon is transparent to the scintillation light that it emits, thus making it possible for such photons to exit the liquid volume. While this is a considerable list of advantages, an additional useful property is that natural xenon is composed of a variety of different isotopes, each with differing total angular momentum and unpaired proton and neutron spin compositions. This proves to be a boon when looking at the zero-momentum spin-dependent differential rate. A summary of key properties of xenon are explicitly given in Table 2.1.

Property	Symbol	Value	Unit
Atomic Number	Z	54	num. protons
Density	ρ_{Xe}	3.057	g/cm^3
Peak scintillation wavelength	λ	177.6	nm
Isotopes (Atomic mass and natural abundance)	A	^{124}Xe(0.09), ^{126}Xe(0.09), ^{128}Xe(1.92), ^{129}Xe(26.44), ^{130}Xe(4.08), ^{131}Xe(21.18), ^{132}Xe(26.89), ^{134}Xe(10.44), ^{136}Xe(8.87)	%

Table 2.1: This table summarizes some key properties for xenon as listed in [34].

Energy depositions and subsequent signal formation in liquid xenon is a very complicated topic. We will not attempt to provide a full overview in this text. An excellent source for those interested can be found in the dissertation of Eric Dahl [35].

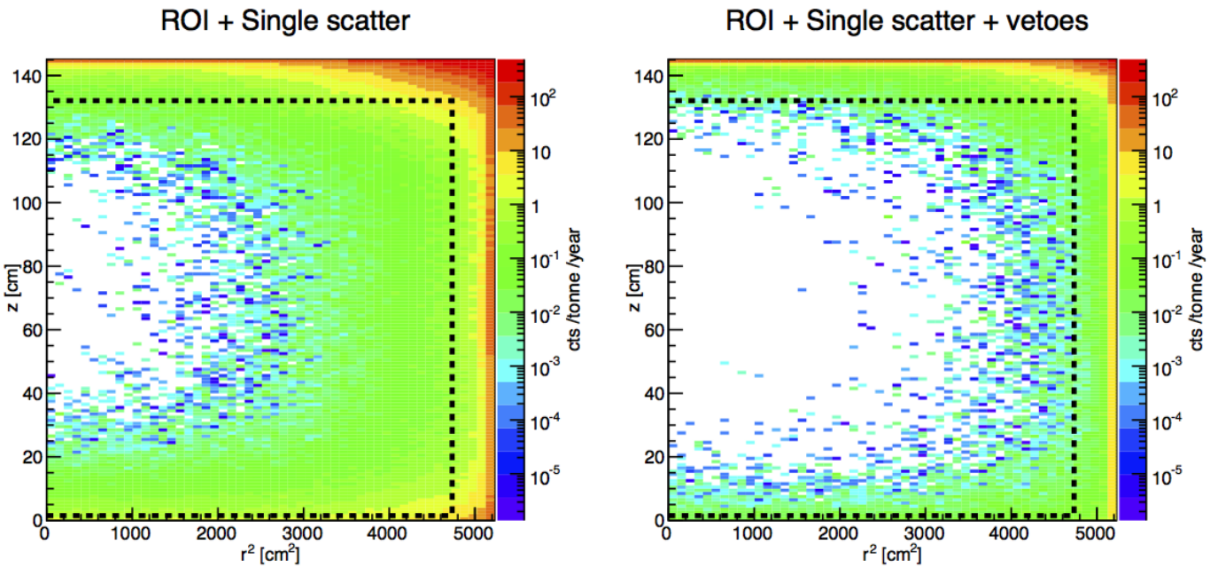

Figure 2.1: From [36], the event rate of single scatters in the ROI without (left) and with (right) applied vetoes. In both cases, the rate is lowest in the innermost region of the detector, due to self-shielding.

A typical interaction within liquid xenon will generate primary scintillation, ionization and heat. We cannot directly detect heat within a dual-phase TPC experiment and thus, energy loss to heat must be modelled using a Lindhard quenching factor. The initial distribution between all three channels will depend on whether the incident particle produced a nuclear (through interactions with nuclei; NR) or electronic recoil (through interactions with bound electrons; ER). Figure 2.2 shows a flowchart for the partition of energy losses, and the subsequent processes that result in the formation of signals. The scintillation light yield is labeled S1, while S2 refers to the ionized electrons that survive recombination.

The yield of light and charge depends on the type of recoil. For NR, the deposited

charge density is much higher than in the case of ER, which results in differences in the recombination rates. Figures 2.3 and 2.4 show the relative scale of yields in the three channels for both types of recoil events. As we can see, the ratio of primary scintillation and ionization for ER and NR origin events is different. This provides the basis for how xenon dual-phase TPC experiments are able to discriminate between background (gammas, betas, etc) from potential signal events (dark matter scattering). This will be discussed in more detail in Chapter 4.

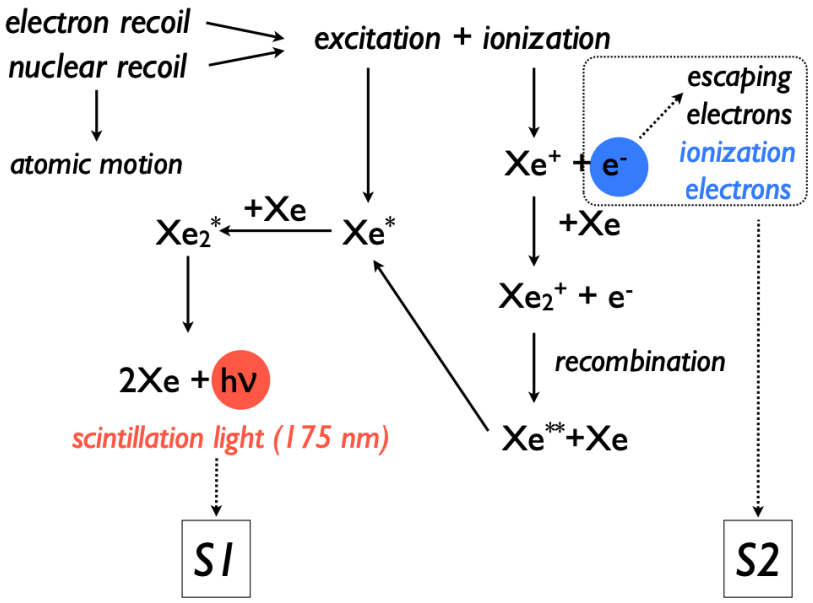

Figure 2.2: Signal formation in xenon for both electronic and nuclear recoils, from [37].

There is one added complication that we need to mention. We are not able to directly observe the initial distribution of primary scintillation and ionization. Recombination pre-

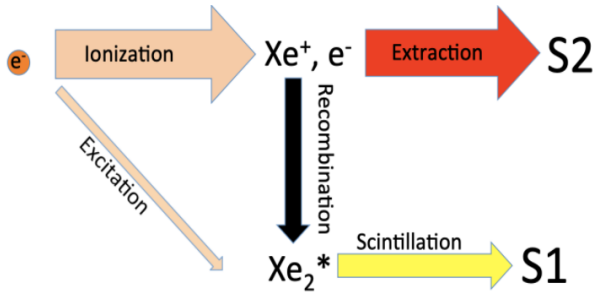

Figure 2.3: A schematic diagram, which illustrates the scale of each signal formation channel in xenon for electronic recoils using width of corresponding arrows. Figure courtesy of the LZ collaboration.

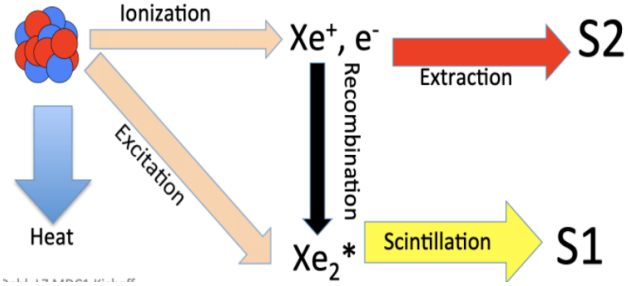

Figure 2.4: A similar diagram for nuclear recoils. Figure courtesy of the LZ collaboration.

vents this possibility. In fact, we can write the following set of equations to describe this effect:

$$\begin{aligned} n_y &= N_{ex} + rN_i \\ n_e &= (1-r)N_i \end{aligned} \tag{2.1}$$

with N_{ex} and N_i representing the initial numbers of produced quanta (electrons and photons) via ionization and scintillation, r the recombination fraction and n_e and n_y the

final numbers of remaining ionization and scintillation quanta, respectively. It is important to note that we do not change the total number of quanta produced, only change the numbers of net scintillation and electron-ion pairs. Thus, we can write

$$N_{ex} + N_i = n_e + n_\gamma. \tag{2.2}$$

The recoil energy can then be measured as

$$E_R = W\left(n_e + n_\gamma\right) \tag{2.3}$$

for electronic recoil events and as

$$E_R = \left(W/\mathcal{L}\right)/\left(n_e + n_\gamma\right) \tag{2.4}$$

for nuclear recoil events. $\mathcal{L}$ represents the Lindhard quenching factor mentioned previously, and W the average energy required to produce a photon or electron-ion pair. It is important to state here that we are not able to detect every produced electron-ion pair or scintillation photon for a given interaction. Detector impurities and light collection efficiencies, among other factors, degrade the signals. As a result, these efficiencies need to be accounted for and calculated for every detector. We shall see how this is done for TPCs in the next section.

We will now return to the processes involved in the scintillation mechanism in liquid xenon, as was shown in Figure 2.2 in a flowchart format. Through either direct excitation

or through recombination, interactions produce excited dimers (excimers) of Xe molecules, which then decay to produce the VUV scintillation light (approx. 178 nm). Because the scintillation light is produced by the decay of the dimer, there are no matching molecules in the xenon volume that can re-absorb these photons, thus leading to the transparency of liquid xenon to its own scintillation light. The process for scintillation proceeds as

$$\begin{aligned} \mathrm{Xe}^* + \mathrm{Xe} &\rightarrow \mathrm{Xe}_2^* \\ \mathrm{Xe}_2^* &\rightarrow 2\mathrm{Xe} + h\nu(\mathrm{VUV}) \end{aligned} \tag{2.5}$$

and, likewise, the process for recombination can be described as

$$\begin{aligned} \mathrm{Xe}^+ + \mathrm{Xe} &\rightarrow \mathrm{Xe}_2^+ \\ \mathrm{Xe}_2^+ + \mathrm{e} &\rightarrow \mathrm{Xe}^{**} + \mathrm{Xe} \\ \mathrm{Xe}^{**} &\rightarrow \mathrm{Xe}^* + \text{ heat} \\ \mathrm{Xe}^* + \mathrm{Xe} &\rightarrow \mathrm{Xe}_2^* \\ \mathrm{Xe}_2^* &\rightarrow 2\mathrm{Xe} + h\nu(\mathrm{VUV}). \end{aligned} \tag{2.6}$$

It is also important to note that the short-lived excimer that is formed in either scenario can be in a singlet or triplet state and although the de-excitation process is different at a base level between the two, they both produce the same scintillation light. There is a small difference in the decay times between the singlet (2.2ns) and the triplet (27ns) states [34]. However, in a detector with a large volume, these differences are eclipsed in the process of light collection, which involves several reflections from the detector walls. Thus, the

technique of pulse shape discrimination (discussed in more detail in Chapter 5) for ER/NR classification is not applicable for liquid xenon detectors. In contrast, liquid argon, another noble element that can be used for direct detection experiments, has a separation of O(1μs) between the singlet and triplet states. For such experiments, this type of discrimination plays an important role in effectively separating ER and NR events [38].

2.1.2 Dual-Phase Time Projection Chambers

The history of time projection chambers (TPCs) and their use in physics experiments spans over a half-century. To speed things along a bit, we shall fast forward a few decades to the introduction of dual-phase TPCs and their usage in dark matter direct detection experiments. As one might expect, dual-phase in the name refers to the fact that two phases of xenon are present in the detector volume. While the bulk of the detector xenon is liquid, a small region at the top of the detector is in the gas phase. A diagram of a dual-phase TPC and expected signal readout is given in Figure 2.5.

Before explaining the signal readout, we must introduce one more key concept. As we have discussed, interactions in liquid xenon produce both primary scintillation (S1) and electron-ion pairs from ionization. There are two methods to read out the ionization signal: a) a direct collection of charge on an anode, or b) devise a method for generating scintillation light that is proportional to the number of electrons surviving recombination. The latter is the primary method used in current dark matter dual-phase TPC experiments.

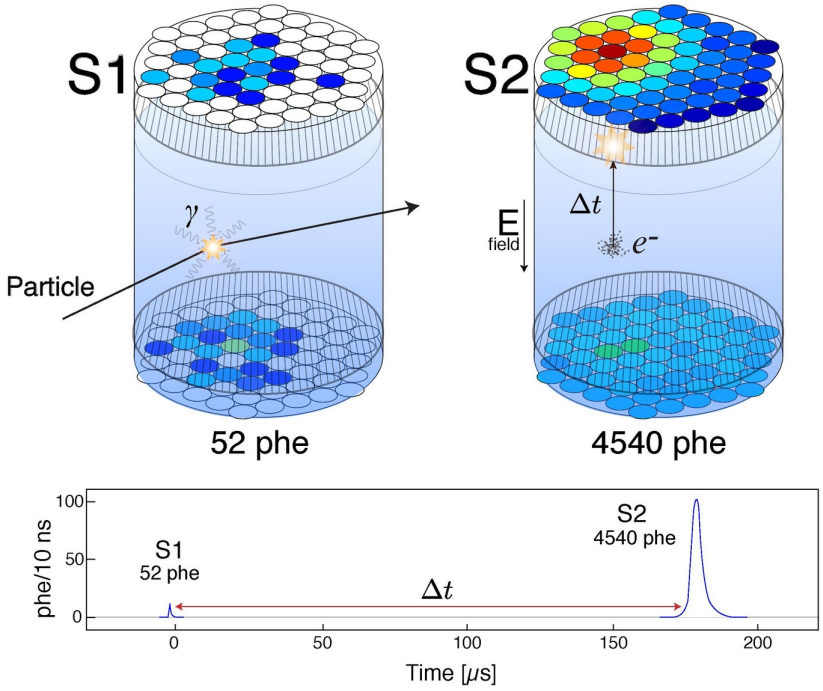

Figure 2.5: An interaction event within a dual-phase LXe TPC and the resulting S1 + S2 waveform. Figure courtesy of the LUX collaboration.

The basic set up in these detectors is to use three primary grids within the cylindrical TPC volume (cathode, gate and anode) to produce a vertical electric field between the cathode (negative voltage) and gate (positive voltage relative to cathode). Appropriately applied voltages to create drift fields of O(100V/cm) then drift electrons from the interaction site and up the detector volume towards the gate. A second, much stronger extraction electric field (several kV/cm) is set up between the gate and the anode at top of the detector to extract drifted electrons across the xenon liquid-gas interface. During this process, the electrons produce proportional electroluminescence as they are drifted through the gas phase.

The electrons gain energy from the electric field, and in turn radiate it away in collisions. Typical light yields can vary from O(10) - O(100) ph/e^- depending on the extraction voltage (directly related to electron extraction efficiency) and the height of the extraction region. The delayed scintillation light (from electroluminescence) produced is referred to as secondary scintillation (S2) light.

Now, the time-projection aspect of a TPC is that it is possible to extract z-position information based on the observed delay between the S1 and S2 scintillation signals. Longer Δt implies that the electrons had to drift farther in the volume before being extracted across the liquid-gas phase change. Refer to Figure 2.5 for a visual representation.

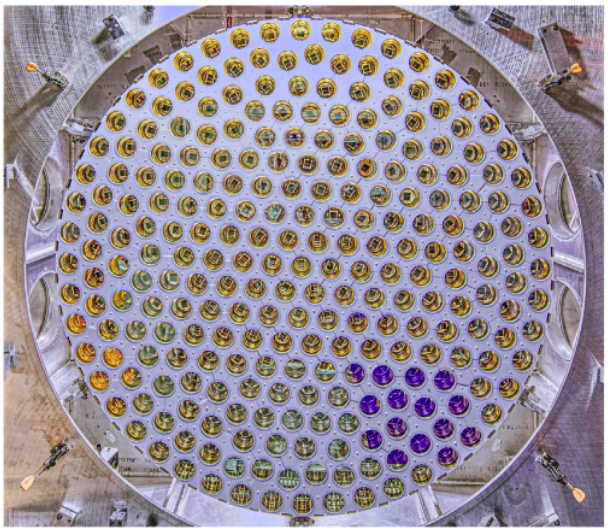

Figure 2.6: From [39], a photograph of the top PMT array for LZ, consisting of 253 PMTs.

The actual detection of both S1 and S2 light is done through the use of two arrays of photomultiplier tubes (PMTs) at either end of the liquid xenon volume. Figure 2.6 shows

a photograph of the top array deployed in LZ; the bottom array is similar. PMTs are photodetectors that work by taking advantage of the photoelectric effect. The window of a PMT is called the photocathode. It is designed to have large efficiency in generating electrons in response to incident light through the photoelectric effect. This efficiency is referred to as the quantum efficiency of a PMT. Gains of $O(10^6)$ in these photodetectors are achieved using multiple subsequent dynode stages that are held at successively higher voltages. The dynodes accelerate the electrons produced at previous stages to then produce secondary, tertiary, etc showers at the successive dynode stages. The final macroscopic current pulse is then read out by the last electrode (the anode). PMTs are sensitive to magnetic fields and also require their internals to be in vacuum. Furthermore, they typically require operating voltages of O(1 kV). Optimal performance is obtained by adjusting the voltage distribution in the dynode chain.

Using the hit pattern in the top array (amount of light seen at each of the PMTs in the array) for a given event, it is also possible to reconstruct the (x,y) position of the interaction using maximum likelihood techniques. Resolution/error depends directly on the total light collected.

We can conclude this section by revisiting Eqns 2.3 and 2.4. Having now introduced S1 and S2 light, these equations can, respectively, be rewritten as

$$E_R = W\left(\frac{\mathrm{S1}}{g_1} + \frac{\mathrm{S2}}{g_2}\right) \tag{2.7}$$

and

$$E_R = \frac{W}{\mathcal{L}} \left(\frac{\mathrm{S1}}{g_1} + \frac{\mathrm{S2}}{g_2} \right), \tag{2.8}$$

where we have introduced two new quantities, g_1 and g_2, to describe the average number of photons detected per primary scintillation photon (a direct efficiency) and per e^- (a combination of electroluminescence yield, and other detector efficiencies). This is how energy reconstruction of events is performed in dual-phase TPC experiments. The detector's $g1$ and $g2$ values are obtained via extensive calibrations using radioactive sources and some short-lived isotopes introduced into the liquid xenon volume.

2.2 The LZ Detector

A large scale overview of the LZ detector is provided in Figure 2.7. A very thorough breakdown of the detector can be found in the technical design report [36]. A more manageable overview of the detector is given in [39]. Presented in this work will be a concise introduction to the key details of the LZ detector.

The detector is a dual-phase xenon TPC with two complementary veto systems. The inner volume of the TPC (approx. 1.5 m in height and diameter) contains 7 active metric tonnes of xenon (approx. 5.6 tonnes in the defined fiducial volume). The TPC contains a total of 494 3-inch Hamamatsu R11410-22 PMTs, 253 in the top PMT array and 241 PMTs in the bottom array. The gate is positioned 5 mm below the top of the liquid level, and

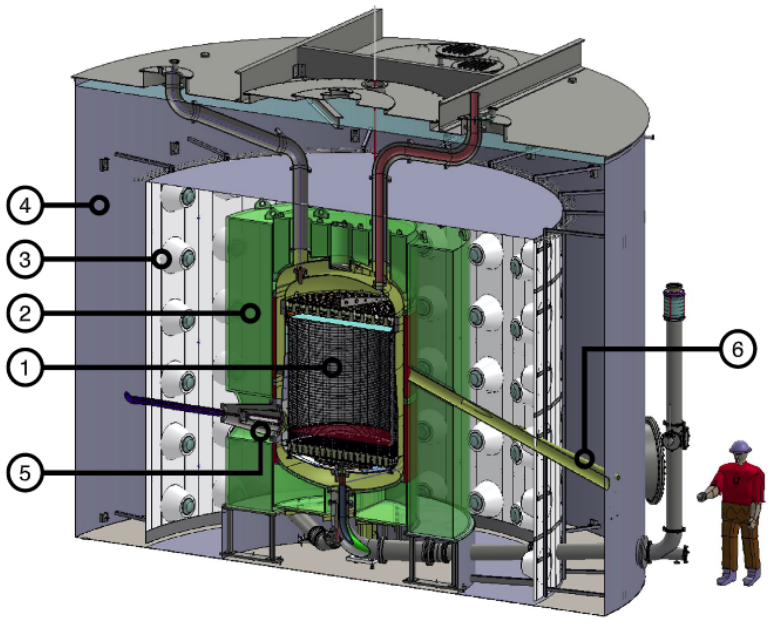

Figure 2.7: From [39], the main sub-components of the LZ detector. (1) is the liquid xenon TPC, with a top and bottom array of PMTs. The TPC is contained in a titanium cryostat and surrounded on all sides by a GdLS Outer Detector (2). The cathode high voltage connection is made at the lower left (5). The GdLS is monitored by 8" PMTs (3) in the water tank (4) which provides shielding for the detector. The conduit on the right (6) allows for neutron calibration sources a direct path to the inner detector.

the anode 8 mm above it in the gaseous Xe region. The drift region between the cathode and gate is 145.6 cm long. To prevent strong negative voltages near the bottom PMT array, a grounded bottom grid is placed below the cathode to shield the PMTS. However, this results in the creation of the so-called reverse field region (RFR) below the cathode, in which electrons drift towards the bottom grid.

The first of two complementary veto systems consists of the region between the TPC instrumentation and the cryostat walls, which contains approximately 2 metric tonnes of

liquid xenon. Instrumented with a total of 131 PMTs (93 1-inch Hamamatsu R8520-406 PMTs and 38 2-inch Hamamatsu R8778 PMTs), this skin detector is intended primarily as a scintillation only veto for gammas scattering into or out of the inner TPC volume. It also serves as a secondary point of comparison for any particle interactions or detector pathologies that can cause unexpected signals in the TPC.

The second system is an outer detector (OD) containing 17 metric tonnes of gadolinium doped liquid scintillator (GdLs) that surrounds both the inner TPC volume and the skin detector. All of this is enclosed within a water tank. Both the water tank and the OD are monitored by 128 Hamamatsu R5912 8-inch PMTs that line the inner surface area of the outer water tank. The OD serves as a highly efficient neutron veto system ($> 95\%$ efficiency) due to the enormous thermal capture cross section of Gd. This is important because neutrons present a nuclear recoil background that cannot be separated from potential dark matter nuclear recoil interactions. Sources of neutrons are expected to primarily come from the experiment cavern walls due to natural radioactivity. Thus, the OD neutron-tagging system identifies any neutrons entering the detector volume, and more importantly, neutrons that exit the TPC volume after an interaction. Because the water tank is also instrumented, we can detect Cerenkov radiation from muons or muon-induced cascades.

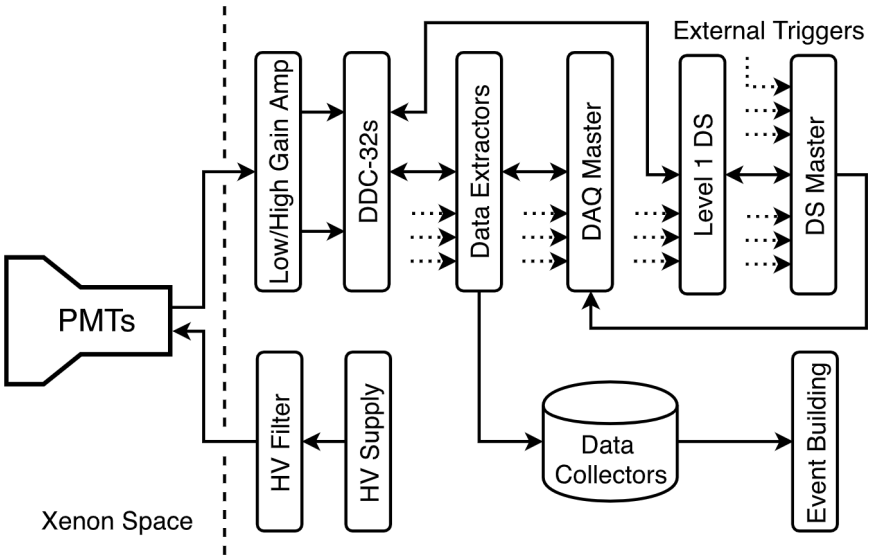

Figure 2.8: From [39], a diagram of the full signal processing chain for LZ. The analog electronics chain involves the HV supplies, filters and the amplifier + slow control boards (not shown).

2.2.1 Analog Electronics Chain

The front-end electronics provide a low noise amplification system for the PMT signals brought out of the TPC via co-axial cables to vacuum feed-through flanges. It is the most critical component in the entire readout system, because its performance in detecting single photons establishes the lowest recoil energy that can be detected by LZ. It has been quite a success considering that the specifications called for a $> 90\%$ detection efficiency of single photoelectrons (SPHEs), but we actually achieved $> 99\%$ detection efficiency. Figure 2.9 shows the distribution of S1 filter peak outputs for SPHEs. The S1 filter threshold is represented by the green dotted line. At this threshold, the SPHE detection efficiency is

99.8%.

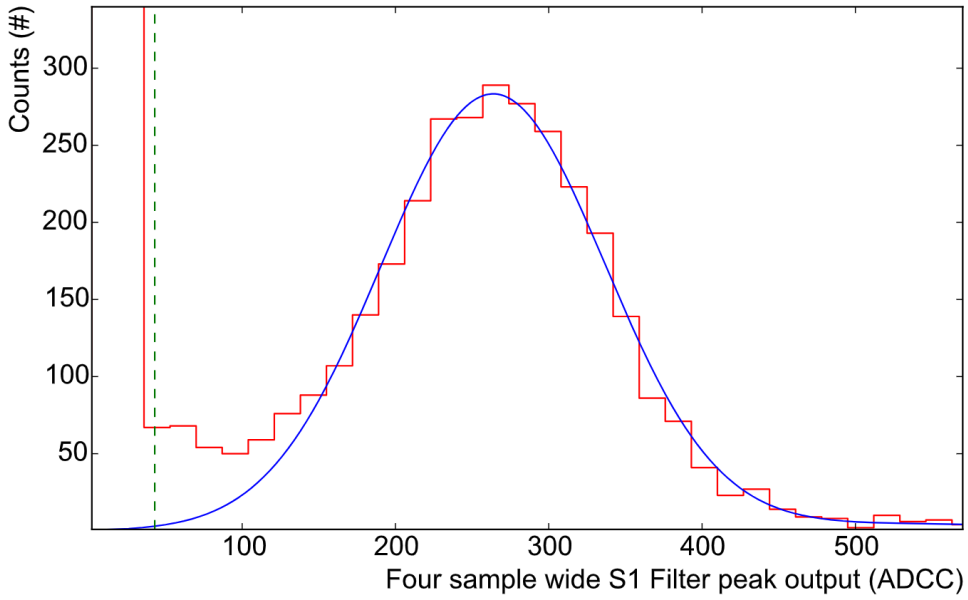

Figure 2.9: A histogram of S1 filter peak outputs. The filter threshold is shown as a green dotted line. The calculated SPHE detection efficiency with this threshold is 99.8%. Figure from [39].

The analog pulses emerging from the front-end are digitized at a sampling rate of 100MS/s and with 14-bit resolution using custom digitizer boards (DDC-32s). The main components of the LZ electronics chain are shown in Figure 2.8. The analog section consists of the signal chain from the PMTs to the input of the DDC-32s along with the HV filters and supplies that provide the optimal bias voltages for each PMT independently.

We, at UC Davis, were responsible for the design, simulation, manufacturing, testing and quality assurance (QA) of the amplifier and control boards as well as the crates that house them. The control boards facilitate communication with and control of the amplifiers

boards. We were also responsible for testing the PMT HV supply modules, the design and manufacturing of the HV filter boards for those modules and the crates that house them. Furthermore, we also performed verification and QA/QC (quality control) of all the custom HV cables that connect the power supply modules to the filter crates.

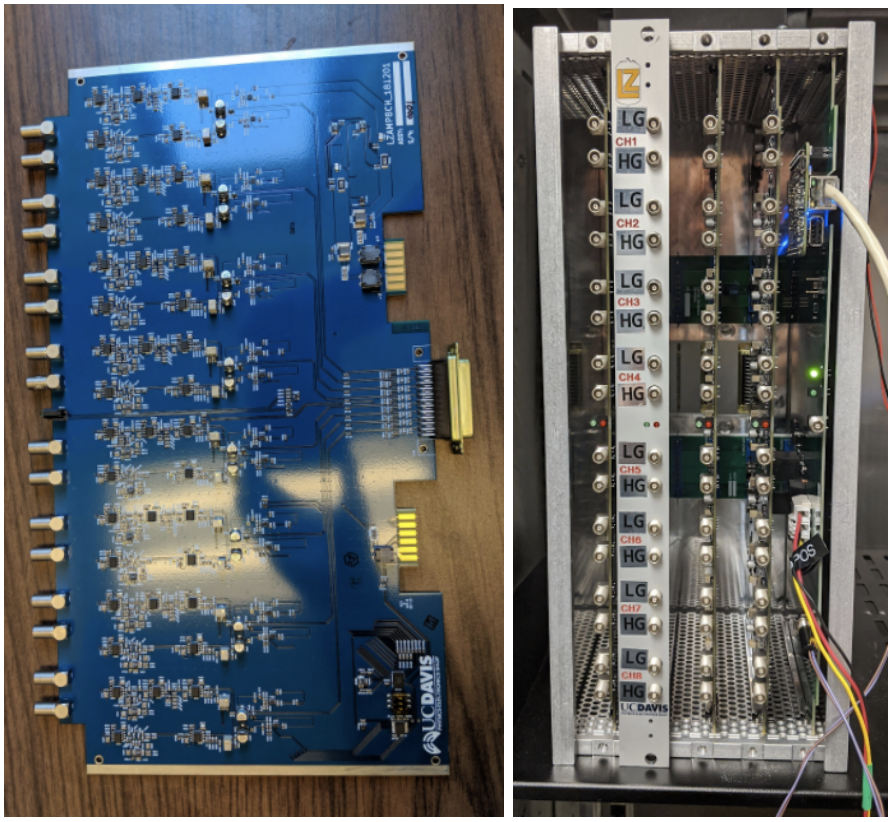

Figure 2.10: Photographs of one of the 100 LZ amplifier boards used for initial signal amplification and shaping (Left) and of one of the 25 LZ amplifier crates housing the amplifiers. All of these modules were tested for QA/QC at UC Davis.

Each amplifier crate, shown in Figure 2.10 (Right), houses four amplifier boards, as shown in Figure 2.10 (Left). The latter picture is of one of the 100 LZ amplifier boards used for initial signal amplification and shaping. Figure 2.11 shows the control board (one per

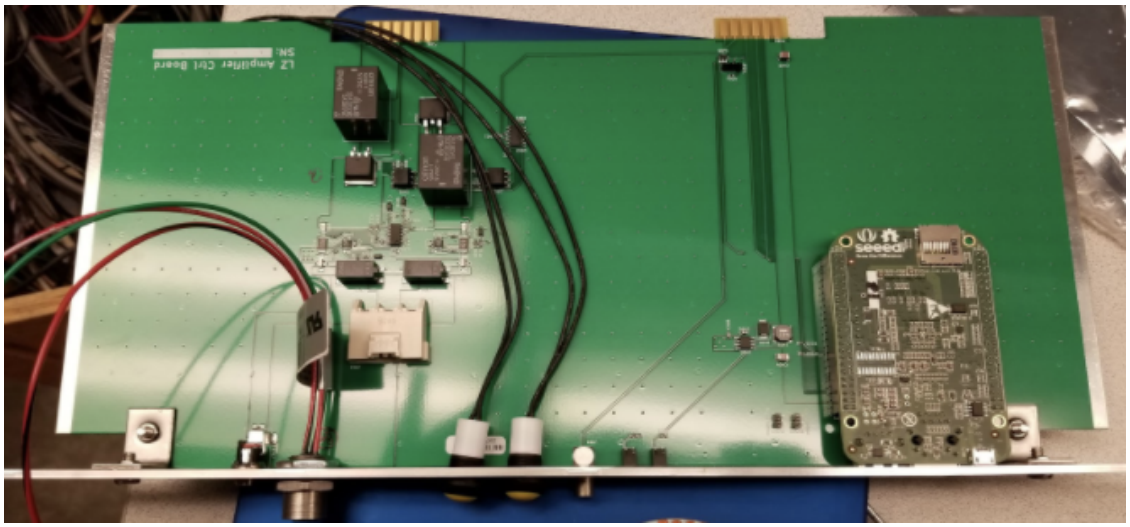

Figure 2.11: A picture of one of 25 slow control boards used to control the amplifiers and facilitate communication between the amplifiers and Ignition, the primary slow control interface.

crate) that enables communication and control of the 4 amplifiers in the crate using I2C (a standard serial digital communication protocol allowing for interfacing to multiple chips on a single line using addressing) on custom bus-lines at the back of the crate. A BeagleBone (a commercially available microprocessor board with various GPIO pins) on the control board serves as the master. The slaves consist of various monitoring sensors located both on the control board and each amplifier board. The quantities monitored consist of airflow, temperature, voltage, current, and DC offset. There are a total of 25 such crates and each of these crate's control boards are connected via ethernet to a client called Ignition using the MODBUS communication protocol. This protocol was initially developed in the 1970s to provide a standard for serial communication for programmable logic controllers. It is still an

industry standard for communication between electronic devices. Ignition, which serves as a master slow control system for the LZ experiment, can query any of the slow control boards and ask for updated sensor values, DC channel offsets, current enable/disable status for any channel/amplifier board. These are continually updated and shown in real-time graphs within the GUI interface.

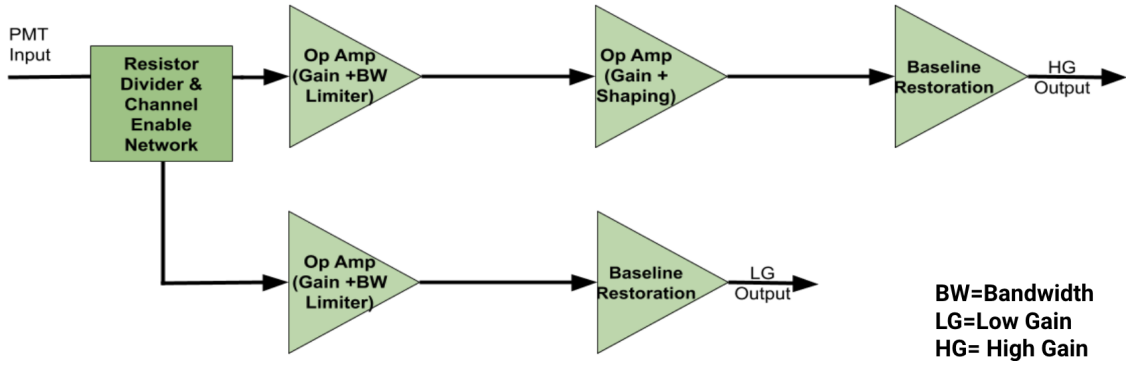

Figure 2.12: A system level overview of the main stages of signal progression in a single input channel of the LZ amplifier board.

For the overall success of the experiment, the primary parameter of interest is the measurement of the area for any given pulse. At the amplifier board level, there are 8 input channels for each board and 16 output channels corresponding to a HG (high gain) and LG (low gain) output for each input (refer back to Figure 2.10 (Left) and to Figure 2.13 for an test pulse input example). This was done in order to ensure maximum dynamic range for signals that the amplifiers see coming from the PMTs; the S1 signals are much smaller than S2, and hence it is best to treat them separately. As a result, the boards have sufficient dynamic range to amplify signals (without saturating) for energies between 0 - 2.5 MeV

in the TPC and up to about 9 MeV in the OD before saturation of the amplifiers occur. The LG output is meant for large S2 signals and is designed for about 4x area gain and a shaping time (full width tenth maximum; FWTM) of 30 ns. Figure 2.12 shows a system level overview of a single channel's main circuit blocks. Likewise, the HG output is meant for amplification of single photo-electrons contained in small S1 signals, and is designed for about 40x area gain and a shaping time factor of 60 ns. The larger shaping time allows for approximately 10 samples of the single photo-electron pulse, thus providing an excellent area measurement. For the TPC and OD PMTs, both HG and LG channels are digitized. For the skin PMTs, only the HG waveforms are digitized. A digitized single scatter event (1 S1 + 1 S2 pulse) within the TPC fiducial volume is shown in Figure 2.14.

While performing QA/QC tests, the primary goal was to ensure signal uniformity over all 1600 (HG and LG) amplifier channels. Thus, for all 100 amplifier boards, we measured and adjusted output baseline offset (using trim potentiometers on the board) to within 0 ± 10mV to ensure maximum dynamic range, characterized output noise levels, measured the area gain, peak-peak delay, FWHM ($\propto$ shaping time) of a test pulse and performed a Bode plot analysis to confirm the gain versus frequency response. We rejected all boards with values outside 3σ of the average across all amplifier boards for all the tests above. This ensured minimal error (and spread) in reconstruction of the deposition energy (area of the pulse) in offline software.

Similarly, for all 25 control boards, we ensured successful I2C communication to associated amplifier boards. We tested the ability to trigger all 32 channels (across the 4 amplifier

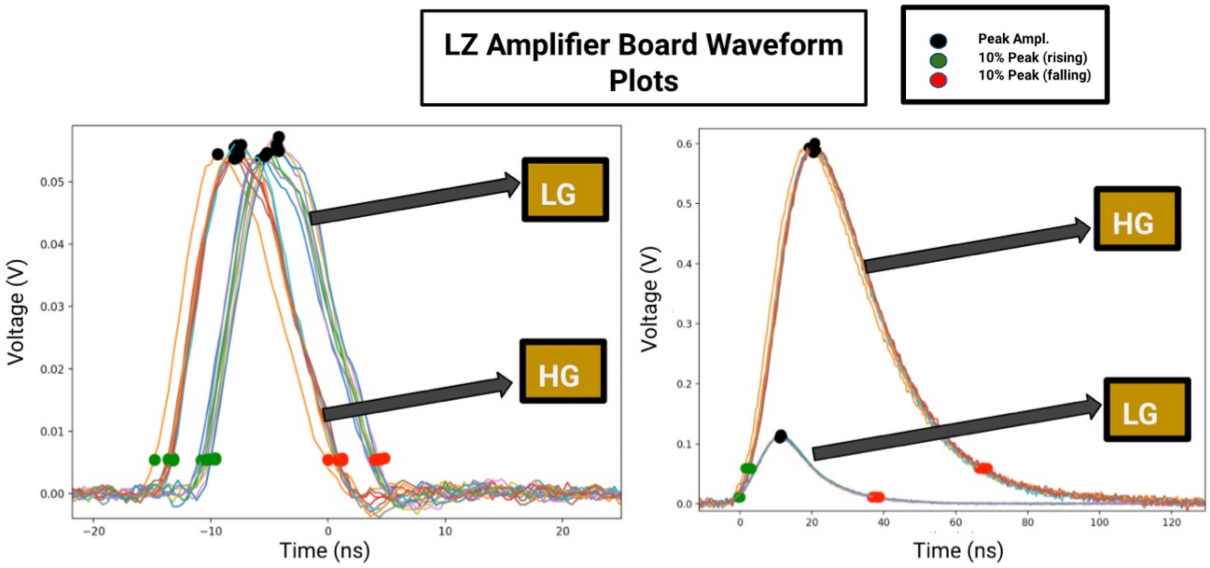

Figure 2.13: Digitized outputs of an input test waveform going through the HG and LG channels of one of the amplifier boards from the QA/QC procedure are plotted. The peak amplitude and 10% of peak (rising/falling) locations for the various channels are marked. It is easy to see signal shape and propagation uniformity within the channels.

boards) in a crate with a test pulse, the channel enable/disable functionality, including capability to hold the present state of all channels in a crate through a power cycle, and the ability to read airflow, temperatures, voltage, current, and all channel DC offsets on amplifier boards when queried within a reasonable amount of time. We also took care to ensure that no ground shorts were present at any point in the signal chain for any board and/or crate.

The entire analog front-end, including crates, amplifiers, power supplies, and communication system was installed in LZ in 2020 and has performed remarkably well to date.

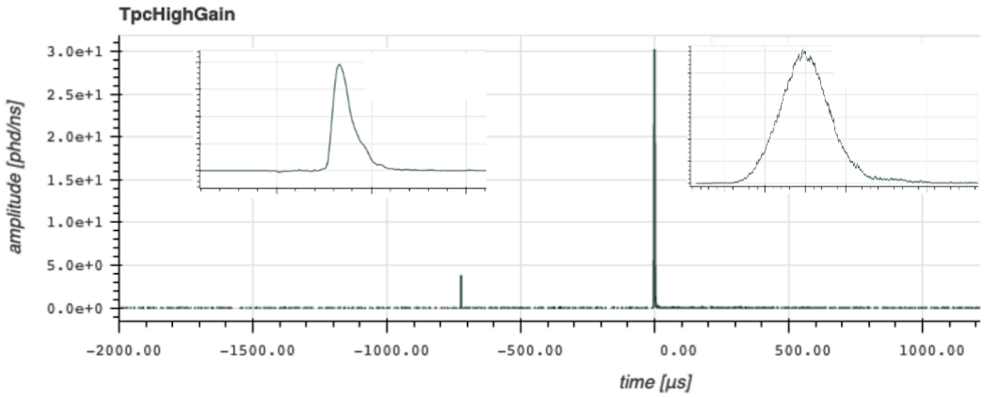

Figure 2.14: A typical single scatter (1 S1 + 1 S2) event within the LZ TPC volume from the Event Viewer. Left inlay zooms into the S1 pulse and the right inlay does the same for the S2 pulse. Event taken from Rn220 calibration data.

Chapter 3

A Search for Dark Matter based on an Effective Field Theory

3.1 Limitations of the SI and SD WIMP Search

Because WIMP velocities are expected to be non-relativistic, we work under the framework of non-relativistic scattering as explained in Chapter 1. In that chapter, we focused only on interactions that are non-vanishing in the zero-momentum limit. In other words, we focused on zero-momentum WIMP-nucleon cross sections. When we did this, only the spin-independent (SI) and spin-dependent (SD) interactions survive. However, as pointed out by Fitzpatrick et al. [40], such simplified models for WIMP-nucleon interactions are limited in scope. In reality, we know almost nothing about possible interactions (outside of gravitational) between dark matter (DM) and the Standard Model (SM). Reducing the

full particle zoo of possible candidates for dark matter to overly simplified models leaves us blind to the possibility that dark matter can, in theory, be as complex as the Standard Model itself. If it is, it is not hard to argue that momentum-dependent interactions can exist even for non-relativistic scattering. In fact, momentum-independent interactions can become vanishing and sub-dominant if dark matter is composite [40].

We are currently engaged in the second generation of dual-phase, noble gas, TPC-based direct dark matter searches. LZ has just recently published the most stringent limits on standard spin-independent WIMP-nucleon cross sections with data taken over nearly 4 calendar months (corresponding to approximately two months of live time) [41]. Now, with more than two decades of ever improving null results in the search for evidence of DM-nucleon interactions, it is time to take a more serious look at expanded possibilities for dark matter interactions. In the next section, we approach the topic in a first principles and model independent manner closely following the derivation given in [40].

3.2 A General Effective Field Theory Model for Dark Matter-Nucleon Interactions

It is important to note, before we begin, that only the main steps of the derivation for an effective field theory (EFT) approach to elastic/inelastic dark matter-nucleus interactions are presented in this dissertation. This is because we do not need to directly implement the

final differential rates presented in [40], since the authors have put together a Mathematica package that already does so, as detailed in [42]. A modified package was then put together by the authors of [27] to include the inelastic extension. These packages were used in the previous generation of (xenon-based) direct dark matter experiments to produce the differential rate spectra as a function of nuclear recoil energy for all of the elastic and inelastic operators in liquid xenon (i.e. the input signal models). Thus, they are also used to do the same for the limit-setting analysis in Chapter 4. The main goal of this chapter is to introduce the theory behind these software packages, which are used to generate the differential rate spectra for the upper limit-setting analysis presented in the next chapter.

Instead of trying to account for all possible models in a more general dark matter-nucleon interaction framework, an effective field theory provides a thorough and expedient description of the problem. This approach allows us to build up a basis of operators involved in describing interactions between dark matter particles and nucleons in the non-relativistic limit without relying on specific models for dark matter [40]. These operators cover a full range of possible interactions and include (in addition to the usual SI and SD interactions) momentum dependent and momentum plus spin dependent interactions. An important limitation is that the operators presented in this section will only be complete for exchanges of spin-0 and spin-1 mediators. This is quite reasonable in light of the fact that all known propagators in the SM are bosons. In this framework, the nuclear properties of different target nuclei play a significant role in the overall strength of couplings beyond the basic A^2 enhancement for SI interactions in the simplified scenario derived in Chapter 1. For details

beyond what is shown in this section, it is advised to consult [40] as the full derivation is quite long and complex.

In any effective field theory, it is important to set a cut off scale, Λ_{UV}. For DM, this cut off is determined by the scale of momentum transfers probed in direct detection experiments. As discussed in Chapter 1, these experiments measure the recoil energy, E_R, of the atomic nuclei from the momentum transfer. For a target mass, m_N, momentum transfer scales as

$$q = \sqrt{2 m_N E_R}. \tag{3.1}$$

The above also sets the minimum velocity needed for an incoming dark matter particle to elastically scatter and produce a measured nuclear recoil energy deposition as

$$v_{min} = \frac{q}{\mu_N}, \tag{3.2}$$

where μ_N is the dark matter-nucleus reduced mass. Taking into account the escape velocity of the galactic dark matter halo, v_{esc}, of approx. 544 km/s, i.e. O(10^{-3} c), and typical target masses of O(100GeV), maximum momentum transfers in such experiments will generally be no more than O(100MeV). Thus, Λ_{UV} should be several times this maximum momentum transfer in order to be applicable in describing direct detection experiments.

Chapter 4 will present limits on inelastic scattering of DM with nuclei in LZ's fiducial volume using SR1 (science run one) data. In that regard, in this chapter we will build a theory that will be applicable to the measurements. So, our end goal in this chapter is to

present the basis of operators that govern inelastic DM-nucleus scattering and an associated differential rate equation. As derived in [27], this can be done with a simple modification of the operators in the elastic scattering EFT framework. Thus, we will start with a summary of the detailed elastic scattering derivation presented by [40] and [43].

It is easier to begin by discussing DM-nucleon interactions and then make the transformation to DM-nucleus interactions. For elastic scattering of dark matter particles off nucleons, we are interested in all effective operators corresponding to interaction Lagrangians of the form

$$\mathcal{L}_{\text{int}} = \sum_{N=n,p} \sum_{i} c_i^{(N)} \mathcal{O}_i \chi^+ \chi^- N^+ N^-, \tag{3.3}$$

with $\mathcal{O}_i$ defined as the effective operators and c_i defined as coefficients included to account for effects of higher order physics. This Lagrangian is a sum over all possible operators and both types of nucleons: protons and neutrons. The goal is to find the form of these operators. However, in attempting this, we need to keep in mind that there are several symmetries that restrict the form of allowed interactions and hence, the operators.

The first symmetry is Galilean invariance, which involves translation of velocities. To see the effect of this symmetry, we can first define momentum transfer in dark matter scattering off nucleons as

$$q = p' - p = k - k', \tag{3.4}$$

where p', p and k', k are the initial and final momenta of the WIMP and nucleon, respectively. Following [40], we can state that $\vec{q}$, the momentum transfer, is Galilean invariant. The incoming dark matter particle velocity in the nucleon rest frame (in general, the relative incoming velocity) is also Galilean invariant. This can be defined explicitly as

$$\vec{v} \equiv \vec{v}_{\chi,\,\mathrm{in}} - \vec{v}_{N,\,\mathrm{in}}\,. \tag{3.5}$$

For elastic scatters, energy conservation is an additional constraint to be obeyed. Setting final and initial energies of the dark-matter-nucleon system as equal, we end up with the following equation to satisfy

$$\vec{v}\cdot\vec{q} = -\frac{q^2}{2\mu_N}. \tag{3.6}$$

Next, we need to satisfy Hermiticity of the scattering interaction. This is equivalent to satisfying crossing symmetry. Because $\vec{q}$ is anti-Hermitian, we need to work with the quantity, $i\vec{q}$, which is Hermitian. One more hiccup to deal with in satisfying Hermiticity is that $\vec{v}$, as defined in Eqn. 3.5, is not Hermitian either. Instead, following [40], we must replace it with a perpendicular velocity defined as,

$$\vec{v}^{\perp} \equiv \vec{v} + \frac{\vec{q}}{2\mu_N}, \tag{3.7}$$

which is a Hermitian invariant velocity.

Having now introduced a Galilean and Hermitian invariant velocity, we now need to include particle spins. Because we are working in the realm of non-relativistic scattering, we can write the dark matter and nuclear spins down directly as $\vec{S}_\chi$ and $\vec{S}_N$. With this, we now have formed a basis of four Galilean and Hermitian invariant quantities to use in generating the complete set of operators (to leading order in velocity) for the non-relativistic interaction Lagrangian (Eqn. 3.3): $i\vec{q}, \quad \vec{v}^\perp, \quad \vec{S}_\chi, \quad \vec{S}_N$. The operators that can be formed from these quantities are given as

$$\begin{aligned}
&\mathcal{O}_1 = \mathbf{1}_\chi \mathbf{1}_N, \quad \mathcal{O}_2 = \left(v^\perp\right)^2, \quad \mathcal{O}_3 = i\vec{S}_N \cdot \left(\frac{\vec{q}}{m_N} \times \vec{v}^\perp\right), \\
&\mathcal{O}_4 = \vec{S}_\chi \cdot \vec{S}_N, \quad \mathcal{O}_5 = i\vec{S}_\chi \cdot \left(\frac{\vec{q}}{m_N} \times \vec{v}^\perp\right), \quad \mathcal{O}_6 = \left(\vec{S}_\chi \cdot \frac{\vec{q}}{m_N}\right)\left(\vec{S}_N \cdot \frac{\vec{q}}{m_N}\right), \\
&\mathcal{O}_7 = \vec{S}_N \cdot \vec{v}^\perp, \quad \mathcal{O}_8 = \vec{S}_\chi \cdot \vec{v}^\perp, \quad \mathcal{O}_9 = i\vec{S}_\chi \cdot \left(\vec{S}_N \times \frac{\vec{q}}{m_N}\right), \\
&\mathcal{O}_{10} = i\vec{S}_N \cdot \frac{\vec{q}}{m_N}, \quad \mathcal{O}_{11} = i\vec{S}_\chi \cdot \frac{\vec{q}}{m_N}, \quad \mathcal{O}_{12} = \vec{S}_\chi \cdot \left(\vec{S}_N \times \vec{v}^\perp\right), \\
&\mathcal{O}_{13} = i\left(\vec{S}_\chi \cdot \vec{v}^\perp\right)\left(\vec{S}_N \cdot \frac{\vec{q}}{m_N}\right), \quad \mathcal{O}_{14} = i\left(\vec{S}_\chi \cdot \frac{\vec{q}}{m_N}\right)\left(\vec{S}_N \cdot \vec{v}^\perp\right), \\
&\mathcal{O}_{15} = -\left(\vec{S}_\chi \cdot \frac{\vec{q}}{m_N}\right)\left(\left(\vec{S}_N \times \vec{v}^\perp\right) \cdot \frac{\vec{q}}{m_N}\right).
\end{aligned} \tag{3.8}$$

As mentioned, the endgame is to calculate differential rates associated with each of these operators for a given target in a direct dark matter search. In order to do so, we now translate the interaction Lagrangian formed by these operators from a DM-nucleon framework to a DM-nucleus framework. In other words, we calculate matrix elements for these nucleon operators inside a given target nucleus, because the specific nuclear form factors for different target nuclei are expected to play a much more important role in this EFT framework.

Initially, we can define an atomic nucleus as a many-body bound state of some X number of nucleons. Next, we need to think of $\vec{v}^{\perp}$ from Eqn. 3.7 as containing two parts. One part acts on the center of mass of a nucleus and the other on relative velocities of the nucleons within the nucleus. Explicitly, we can write

$$\vec{v}^{\perp} \equiv \vec{v} + \frac{\vec{q}}{2\mu_N} = \vec{v}_T^{\perp} + \vec{v}_N^{\perp} \tag{3.9}$$

where $\vec{v}_T^{\perp}$ (component acting on nucleus center of mass) and $\vec{v}_N^{\perp}$ (component acting on relative velocities of individual nucleons) are defined as the following

$$\vec{v}_T^{\perp} = \frac{1}{2}\left(\vec{v}_{\chi,\,\text{in}} + \vec{v}_{\chi,\,\text{out}} - \vec{v}_{T,\,\text{in}} - \vec{v}_{T,\,\text{out}}\right) = \vec{v}_T + \frac{\vec{q}}{2\mu_T} \tag{3.10}$$

and

$$\vec{v}_N^{\perp} = -\frac{1}{2}\left(\vec{v}_{N,\,\text{in}} + \vec{v}_{N,\,\text{out}}\right), \tag{3.11}$$

with $\vec{v}_T$ the incoming dark matter velocity in the lab frame.

Keeping in mind that each $\vec{v}^{\perp}$ term in the Lagrangian (and thus in the presented operator basis) can be separated into these two parts, we can write out the interaction Lagrangian in the center of mass frame of the nucleus, following [40], as:

$$
\begin{aligned}
\mathcal{L}_{\text{int}} &= a_1 \\
&+ ia_3\vec{S}_N \cdot \left(\vec{q} \times \vec{v}^{\perp}\right) \\
&+ a_4\vec{S}_\chi \cdot \vec{S}_N \\
&+ ia_5\vec{S}_\chi \cdot \left(\vec{q} \times \vec{v}^{\perp}\right) \\
&+ a_6\left(\vec{S}_\chi \cdot \vec{q}\right)\left(\vec{S}_N \cdot \vec{q}\right) \\
&+ a_7\vec{S}_N \cdot \vec{v}^{\perp} \\
&+ a_8\vec{S}_\chi \cdot \vec{v}^{\perp} \\
&+ ia_9\vec{S}_\chi \cdot \left(\vec{S}_N \times \vec{q}\right) \\
&+ ia_{10}\vec{S}_N \cdot \vec{q} \\
&+ ia_{11}\vec{S}_\chi \cdot \vec{q} \\
&+ a_{12}\vec{S}_\chi \cdot \left(\vec{S}_N \times \vec{v}^{\perp}\right) \\
&+ ia_{13}\left(\vec{S}_\chi \cdot \vec{v}^{\perp}\right)\left(\vec{S}_N \cdot \vec{q}\right) \\
&+ ia_{14}\left(\vec{S}_\chi \cdot \vec{q}\right)\left(\vec{S}_N \cdot \vec{v}^{\perp}\right) \\
&\quad - a_{15}\left(\vec{S}_\chi \cdot \vec{q}\right)\left(\left(\vec{S}_N \times \vec{v}^{\perp}\right) \cdot \vec{q}\right).
\end{aligned}
\tag{3.12}
$$

In this Lagrangian, [40] defines three nuclear "charges" as

$$
1, \vec{v}_N^{\perp} \cdot \vec{v}_N^{\perp}, \text{ and } \vec{S}_N \cdot \vec{v}_N^{\perp} \tag{3.13}
$$

and three nuclear "currents" as

$$\vec{v}_N^{\perp}, \vec{S}_N, \text{ and } \vec{S}_N \times \vec{v}_N^{\perp}. \tag{3.14}$$

We are primarily concerned with leading order interactions in this derivation, so we can safely ignore the $\vec{v}_N^{\perp} \cdot \vec{v}_N^{\perp}$ term. Thus, we can state that there are a total of five unique nuclear currents and charges that can couple with dark matter in scattering interactions. We can re-group the surviving fourteen terms of this Lagrangian into five nuclear current and charge terms with the coefficients being the combinations of coefficients, a_i, and non-nuclear operator terms that contain each of the five terms respectively. At the same time, we can expand a_i in the isospin basis following [40] as

$$a_i \rightarrow \left(a_i^0 + a_i^1 \tau_3(i)\right). \tag{3.15}$$

Performing both of these steps allows the Lagrangian to be re-written more succinctly as

$$\mathcal{L} = l_0 1 + l_0^A \left[-2\vec{v}_N^{\perp} \cdot \vec{S}_N\right] + \vec{l}_5 \cdot \left[2\vec{S}_N\right] + \vec{l}_M \cdot \left[-\vec{v}_N^{\perp}\right] + \vec{l}_E \cdot \left[2i\vec{v}_N^{\perp} \times \vec{S}_N\right] \tag{3.16}$$

with the coefficients for each of the charges and currents given explicit names corresponding to densities: l_0 (charge), l_0^A (axial charge), $\vec{l}_5$ (axial vector), $\vec{l}_M$ (vector magnetic), and $\vec{l}_E$ (vector electric). As [40] explains, in a Galilean invariant theory, all contributions proportional to $\vec{v}_N^{\perp}$ reside in the l_0, $\vec{l}_5$, $\vec{l}_M$, and $\vec{l}_E$ terms. Expanded, these four densities have the following definitions:

$$
\begin{aligned}
l_0 &= \left(a_1^0 + a_1^1\tau_3(i)\right) - i\left(\vec{q}\times\vec{S}_\chi\right)\cdot\vec{v}_T^{\perp}\left(a_5^0 + a_5^1\tau_3(i)\right) + \vec{S}_\chi\cdot\vec{v}_T^{\perp}\left(a_8^0 + a_8^1\tau_3(i)\right) \\
&\quad + i\vec{q}\cdot\vec{S}_\chi\left(a_{11}^0 + a_{11}^1\tau_3(i)\right) \\
\vec{l}_5 &= \frac{1}{2}\Big[i\vec{q}\times\vec{v}_T^{\perp}\left(a_3^0 + a_3^1\tau_3(i)\right) + \vec{S}_\chi\left(a_4^0 + a_4^1\tau_3(i)\right) + \vec{S}_\chi\cdot\vec{q}\vec{q}\left(a_6^0 + a_6^1\tau_3(i)\right) + \\
&\quad + \vec{v}_T^{\perp}\left(a_7^0 + a_7^1\tau_3(i)\right) + i\vec{q}\times\vec{S}_\chi\left(a_9^0 + a_9^1\tau_3(i)\right) + i\vec{q}\left(a_{10}^0 + a_{10}^1\tau_3(i)\right)\Big] \\
\vec{l}_M &= i\vec{q}\times\vec{S}_\chi\left(a_5^0 + a_5\tau_3(i)\right) - \vec{S}_\chi\left(a_8^0 + a_8^1\right)\tau_3(i) \\
\vec{l}_E &= \frac{1}{2}\vec{q}\left(a_3^0 + a_3^1\tau_3(i)\right).
\end{aligned}
\tag{3.17}
$$

In preparation for translating the Lagrangian in Eqn. 3.16 into a Hamiltonian density, we can rewrite it with the following substitutions,

$$
\begin{aligned}
-\vec{v}_{N_i}^{\perp} &= \frac{\vec{p}_{N_i} + \vec{p}_{N_f}}{2m_N} \\
&\text{and} \\
\vec{S}_{N_i} &= \vec{\sigma}(i),
\end{aligned}
\tag{3.18}
$$

as

$$
\mathcal{L} = l_0 1 + l_0^A\left(\frac{\vec{p}_i + \vec{p}_f}{2m_N}\right)\cdot\vec{\sigma} + \vec{l}_5\cdot\vec{\sigma} + \vec{l}_M\cdot\left(\frac{\vec{p}_i + \vec{p}_f}{2m_N}\right) + \vec{l}_E\cdot\left(-i\frac{\vec{p}_i + \vec{p}_f}{2m_N}\times\vec{\sigma}\right). \tag{3.19}
$$

The index, i, in $\vec{v}_{N_i}^{\perp}$ and $\vec{S}_{N_i}$ specifies the ith nucleon within a target nucleus and $\vec{p}_i$, $\vec{p}_f$ represent the initial and final momenta of that nucleon. $\vec{\sigma}(i)$ is the Pauli spin matrix for that ith nucleon, $\vec{r}_i$ gives the position for the ith nucleon and, finally, $\vec{r}$ represents the position of the dark matter particle. We are almost in a position to write the Hamiltonian density. In

order to do so, we now need to convert from momentum to coordinate space. The following identities are used to do this:

$$\frac{\vec{p}_{N_i} + \vec{p}_{N_f}}{2m_N} \to \frac{1}{2m_N} i \left(\overleftarrow{\nabla} \delta \left(\vec{r} - \vec{r}_i \right) - \delta \left(\vec{r} - \vec{r}_i \right) \vec{\nabla} \right) \tag{3.20}$$

and

$$1 \to \delta \left(\vec{r} - \vec{r}_i \right). \tag{3.21}$$

After doing so, we can then (skipping a few steps) write down the Hamiltonian density, following [40], as

$$\begin{aligned} \mathcal{H}(\vec{x}) = & \sum_{i=1}^{A} l_0(i) \delta \left(\vec{x} - \vec{x}_i \right) + \sum_{i=1}^{A} l_0^A(i) \frac{1}{2m_N} \left[-\frac{1}{i} \overleftarrow{\nabla}_i \cdot \vec{\sigma}(i) \delta \left(\vec{x} - \vec{x}_i \right) + \delta \left(\vec{x} - \vec{x}_i \right) \vec{\sigma}(i) \cdot \frac{1}{i} \vec{\nabla}_i \right] \\ & + \sum_{i=1}^{A} \vec{l}_5(i) \cdot \vec{\sigma}(i) \delta \left(\vec{x} - \vec{x}_i \right) + \sum_{i=1}^{A} \vec{l}_M(i) \cdot \frac{1}{2m_N} \left[-\frac{1}{i} \overleftarrow{\nabla}_i \delta \left(\vec{x} - \vec{x}_i \right) + \delta \left(\vec{x} - \vec{x}_i \right) \frac{1}{i} \vec{\nabla}_i \right] \\ & + \sum_{i=1}^{A} \vec{l}_E(i) \cdot \frac{1}{2m_N} \left[\overleftarrow{\nabla}_i \times \vec{\sigma}(i) \delta \left(\vec{x} - \vec{x}_i \right) + \delta \left(\vec{x} - \vec{x}_i \right) \vec{\sigma}(i) \times \vec{\nabla}_i \right]. \end{aligned} \tag{3.22}$$

We can extract the Hamiltonian from this Hamiltonian density equation by integrating over the positions of all nucleons in the nucleus and the position of the DM particle. Then, with the Hamiltonian in hand, and averaging over initial nuclear spins and summing over final state spins, we write down the dark matter scattering probability as

$$
\begin{aligned}
&\frac{1}{2J_i+1}\sum_{M_i,M_f}|\langle J_iM_f|H|J_iM_i\rangle|^2=\frac{4\pi}{2J_i+1}\left[\sum_{J=1,3,\ldots}^{\infty}\left|\left\langle J_i\left\|\vec{l}_5\cdot\hat{q}\Sigma''_J(q)\right\|J_i\right\rangle\right|^2\right.\\
&+\sum_{J=0,2,\ldots}^{\infty}\left\{|\langle J_i\,\|l_0M_J(q)\|\,J_i\rangle|^2+\left|\left\langle J_i\left\|\vec{l}_E\cdot\hat{q}\frac{q}{m_N}\Phi''(q)\right\|J_i\right\rangle\right|^2\right.\\
&\left.+2\,\mathrm{Re}\left[\left\langle J_i\left\|\vec{l}_E\cdot\hat{q}\frac{q}{m_N}\Phi''(q)\right\|J_i\right\rangle\langle J_i\,\|l_0M_J(q)\|\,J_i\rangle^*\right]\right\}\\
&+\frac{q^2}{2m_N^2}\sum_{J=2,4,\ldots}^{\infty}\left(\left\langle J_i\left\|\vec{l}_E\tilde{\Phi}'_J(q)\right\|J_i\right\rangle\cdot\left\langle J_i\left\|\vec{l}_E\tilde{\Phi}'_J(q)\right\|J_i\right\rangle^*-\left|\left\langle J_i\left\|\vec{l}_E\cdot\hat{q}\tilde{\Phi}'_J(q)\right\|J_i\right\rangle\right|^2\right)\\
&+\sum_{J=1,3,\ldots}^{\infty}\left\{\frac{q^2}{2m_N^2}\left(\left\langle J_i\left\|\vec{l}_M\Delta_J(q)\right\|J_i\right\rangle\cdot\left\langle J_i\left\|\vec{l}_M\Delta_J(q)\right\|J_i\right\rangle^*-\left|\left\langle J_i\left\|\vec{l}_M\cdot\hat{q}\Delta_J(q)\right\|J_i\right\rangle\right|^2\right)\right.\\
&+\frac{1}{2}\left(\left\langle J_i\left\|\vec{l}_5\Sigma'_J(q)\right\|J_i\right\rangle\cdot\left\langle J_i\left\|\vec{l}_5\Sigma'_J(q)\right\|J_i\right\rangle^*-\left|\left\langle J_i\left\|\vec{l}_5\cdot\hat{q}\Sigma'_J(q)\right\|J_i\right\rangle\right|^2\right)\\
&\left.\left.+\,2\,\mathrm{Re}\left[i\hat{q}\cdot\left\langle J_i\left\|\vec{l}_M\frac{q}{m_N}\Delta_J(q)\right\|J_i\right\rangle\times\left\langle J_i\left\|\vec{l}_5\Sigma'_J(q)\right\|J_i\right\rangle^*\right]\right\}\right].
\end{aligned}
\tag{3.23}
$$

In this equation for the scattering probability, [40] has introduced six nuclear responses and constructed them in terms of vector and spherical harmonics. Rewritten here, they are defined explicitly in this basis as:

$$
\begin{aligned}
\Delta_{JM}(q\vec{x}) &\equiv \vec{M}^{M}_{JJ}(q\vec{x}) \cdot \frac{1}{q}\vec{\nabla} \\
\Sigma'_{JM}(q\vec{x}) &\equiv -i\left\{\frac{1}{q}\vec{\nabla} \times \vec{M}^{M}_{JJ}(q\vec{x})\right\} \cdot \vec{\sigma} = [J]^{-1}\left\{-\sqrt{J}\vec{M}^{M}_{JJ+1}(q\vec{x}) + \sqrt{J+1}\vec{M}^{M}_{JJ-1}(q\vec{x})\right\} \cdot \vec{\sigma} \\
\Sigma''_{JM}(q\vec{x}) &\equiv \left\{\frac{1}{q}\vec{\nabla} M_{JM}(q\vec{x})\right\} \cdot \vec{\sigma} = [J]^{-1}\left\{\sqrt{J+1}\vec{M}^{M}_{JJ+1}(q\vec{x}) + \sqrt{J}\vec{M}^{M}_{JJ-1}(q\vec{x})\right\} \cdot \vec{\sigma} \\
\tilde{\Phi}'_{JM}(q\vec{x}) &\equiv \left(\frac{1}{q}\vec{\nabla} \times \vec{M}^{M}_{JJ}(q\vec{x})\right) \cdot \left(\vec{\sigma} \times \frac{1}{q}\vec{\nabla}\right) + \frac{1}{2}\vec{M}^{M}_{JJ}(q\vec{x}) \cdot \vec{\sigma} \\
\Phi''_{JM}(q\vec{x}) &\equiv i\left(\frac{1}{q}\vec{\nabla} M_{JM}(q\vec{x})\right) \cdot \left(\vec{\sigma} \times \frac{1}{q}\vec{\nabla}\right),
\end{aligned}
\tag{3.24}
$$

where $M_{JM}(q\vec{x})$ and $\vec{M}^{M}_{JL}(q\vec{x})$ are Bessel and vector spherical harmonics, $\vec{\sigma}$ are Pauli spin matrices and $[J] = \sqrt{2J+1}$. For further details on these nuclear responses, including a comparison of these responses for different target nuclei, we can refer to the original derivation. However, we will move forwards and, as done in [43], relate the nuclear responses $\left(k = M, \Sigma'', \Sigma', \Delta, \Phi'', \tilde{\Phi}'\right)$ to operator form factors $\left(F^{(N,N')}_{i,j}\right)$ through linear combinations of the nuclear form factors $\left(F^{(N,N')}_{k}\right)$. This can be written as:

$$
F^{(N,N')}_{ij}\left(q^2, v^2\right) = \sum_{k=M,\Sigma'',\Sigma',\Delta,\Phi'',\tilde{\Phi}'} a_{ijk}\left(j_\chi, v^2, q^2\right) F^{(N,N')}_{k}\left(q^2\right). \tag{3.25}
$$

In this equation, all properties of the WIMP, including mass, spin and relative velocity, are folded into the coefficient term. There are actually two scenarios for providing definitions of the nuclear form factors, as stated in [43]. The first is if there are no interference terms involved:

$$F_k^{(N,N')}\left(q^2\right) \equiv \frac{4\pi}{2j+1}\sum_{J=0}^{2j+1}\left\langle j\left\|k_J^{(N)}\right\|j\right\rangle\left\langle j\left\|k_J^{(N')}\right\|j\right\rangle. \tag{3.26}$$

The second is in the case of interference between Φ'' and M or Σ' and Δ (the two non-zero combinations). In this scenario, the nuclear form factor equation can be written as:

$$F_{k_1,k_2}^{(N,N')}\left(q^2\right) \equiv \frac{4\pi}{2j+1}\sum_{J=0}^{2j+1}\left\langle j\left\|k_{1J}^{(N)}\right\|j\right\rangle\left\langle j\left\|k_{2J}^{(N')}\right\|j\right\rangle. \tag{3.27}$$

We now, at last, can proceed to the final task left: relating this back to a differential rate. To start, we can rewrite Eqn. 1.2 from Chapter 1 here as:

$$\frac{dR}{dE_R} = nvN_T\frac{d\sigma}{dE_R}. \tag{3.28}$$

This is the relationship of the differential rate to the differential scattering rate. In general, we can find the latter using Fermi's Golden Rule and can write it as

$$\frac{d\sigma}{dE_R} = \frac{1}{32\pi v^2 m_\chi^2 m_A}\frac{1}{(2j_A+1)}\frac{1}{(2j_\chi+1)}\sum_{\text{spins}}\left|\mathcal{M}\right|^2\left(q^2,v^2\right)dv \tag{3.29}$$

with the usual protocol of summing over final states and averaging over initial states. This specific equation written here does contain an additional normalization to account for the non-relativistic operators being considered [40]. The scattering amplitude can be written as:

$$\frac{1}{(2j_A+1)}\frac{1}{(2_\chi j+1)}\sum_{\text{spins}}|\mathcal{M}|^2 \equiv \frac{m_A^2}{m_N^2}\sum_{i,j=1}^{15}\sum_{N,N'=p,n} c_i^{(N)}c_j^{(N')}F_{ij}^{(N,N')}\left(q^2,v^2\right). \tag{3.30}$$

Now, $F_{ij}^{(N,N')}\left(q^2,v^2\right)$ can be substituted by the definition given in Eqn. 3.25 and then plugging Eqn. 3.29 back into Eqn. 3.28 gives a final expression for the dark matter-nucleus differential rate:

$$\frac{dR}{dE_R} = \frac{\rho_0}{32\pi m_x^3 m_N^2}\int_{v>v_{\min}}\frac{f(\vec{v})}{v}\sum_{i,j}\sum_{N,N'=n,p} c_i^N c_j^{N'}\sum_{k=\mathrm{M},\Sigma'',\Sigma',\Delta,\Phi'',\tilde{\Phi}'} a_{ijk}F_k^{(N,N')}\left(v^2,q^2\right)dv. \tag{3.31}$$

This equation, which was derived using the full basis of operators that survive in a general effective field theory of dark matter-nucleus elastic scattering interactions, can be compared to the simplified WIMP-nucleon elastic scattering differential rates derived in Chapter 1 for the zero-momentum limit spin-independent and spin-dependent cases given by Eqns. 1.12 and 1.15. More details on coefficient terms are found in [40].

3.2.1 The Transformation From Elastic to Inelastic Scattering

In the previous section, we arrived at an equation for the elastic DM-nucleus scattering differential rate that can be compared to the one derived for the much simpler zero-momentum limit differential rates. However, in Chapter 4, we will present limits using SR1 data for a basis of operators formed from a general effective field theory, which describes inelastic

scattering of dark matter with target nuclei. Therefore, we still have some work to do in this chapter.

Fortunately, we do not need to re-derive everything. As stated previously, the transition from elastic to inelastic dark matter-nucleus scattering mainly requires a correction to $\vec{v}^{\perp}$, the perpendicular component of velocity that was derived as a Galilean invariant and Hermitian quantity [27]. We will proceed in this task by sketching out the main steps of the transition from elastic to inelastic scattering as presented by Barello, et al. in [27], similar to what we did in the previous derivation of the differential rate for the elastic scattering scenario. For inelastic scattering, we are concerned with scatterings of the type,

$$\chi_1(\vec{p})N(\vec{k}) \rightarrow \chi_2\,(\vec{p}\,')\,N(\vec{k}'), \tag{3.32}$$

where $\chi_1(\vec{p})$ and $\chi_2\,(\vec{p}\,')$ represent the initial and final state dark matter particle and N represent a nucleon within a target nucleus. Figure 3.1 shows a Feynman diagram of this process.

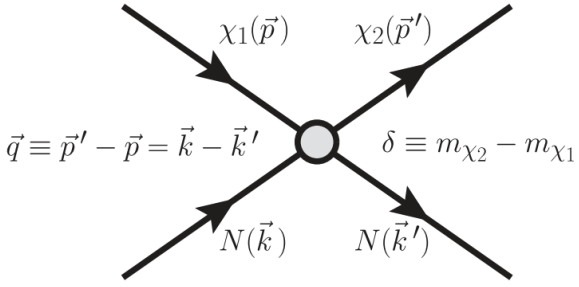

Figure 3.1: The general Feynman diagram for inelastic dark matter scattering off a nucleon within a target nucleus [27].

We now define a mass-splitting between the final and initial dark matter particles as:

$$\delta = m_{\chi_2} - m_{\chi_1}. \tag{3.33}$$

This term represents the additional energy required for such a transition of dark matter to occur. This value is not, at a model-independent level, constrained to be either positive or negative. In fact, [27] states that positive values indicate a preference for dark matter scattering off heavier target nuclei, which in turn will show up as an increase in the annual modulation fraction in such experiments. Negative values, on the other hand, imply exothermic transitions. The scale of mass splitting must be $\mathrm{O}(v^2)$ in order to be able to perform a consistent velocity expansion and this implies that the range of mass splitting values is:

$$\delta \sim 100keV \left(\frac{m_\chi}{100GeV}\right). \tag{3.34}$$

We now proceed to include this mass-splitting term in the kinematics of the scattering. In the previous derivation, we had identified two Galilean invariant quantities (Eqns. 3.4 and 3.5), however, in the case of inelastic scattering, $\vec{q}$ is no longer strictly a Galilean invariant due to δ. It is now only invariant to leading order in velocity expansion. Thus, we shall observe this constraint as we proceed. Earlier, we had imposed energy conservation as a constraint, but this is not precisely applicable in the inelastic case. We must include the mass splitting term in equating final and initial energies involved in this non-relativistic scattering process. Expanding to second order in velocities, following [27], we get the following expressions for

initial and final energies:

$$E_{\text{in}} \approx m_{\chi_1} + m_N + \frac{1}{2}\mu_N v^2 \tag{3.35}$$

and

$$\begin{aligned} E_{\text{out}} &= m_{\chi_2} + m_N + \frac{1}{2m_{\chi_2}}|\vec{p}+\vec{q}|^2 + \frac{1}{2m_N}|\vec{k}-\vec{q}|^2 \\ &\approx E_{\text{in}} + \delta + \vec{v}\cdot\vec{q} + \frac{|\vec{q}|^2}{2\mu_N}. \end{aligned} \tag{3.36}$$

In the latter equation, we treat all momenta as $O(v)$, and δ as $O(v^2)$. Now, we can state that, in analog to Eqn. 3.6, we end up with

$$\delta + \vec{v}\cdot\vec{q} + \frac{|\vec{q}|^2}{2\mu_N} = 0 \tag{3.37}$$

as a constraint from energy conservation in inelastic scattering. Here, let us pause for a moment to note that from Eqn. 3.37, we can get an equation for minimum velocity required for inelastic scattering with some mass splitting, δ. We see that we must satisfy

$$|\vec{v}| \geq \frac{1}{|\vec{q}|}\left|\frac{|\vec{q}|^2}{2\mu_N} + \delta\right|. \tag{3.38}$$

Remembering that q can be defined as $\sqrt{2m_N E_R}$, we can rewrite this constraint as

$$v_{\text{min}} = \frac{1}{\sqrt{2m_N E_R}}\left|\frac{m_N E_R}{\mu_N} + \delta\right|. \tag{3.39}$$

We can now move back to the main discussion. Following the derivation for elastic scattering, the next constraint we had to satisfy was Hermiticity. We had stated that $\vec{v}$ was not a Hermitian invariant quantity, and it had to be replaced with $\vec{v}^{\perp}$ (Eqn 3.7). For the inelastic case, this term must be replaced with an analog $\vec{v}^{\perp}_{\rm inel}$ defined as

$$\vec{v}^{\perp}_{\rm inel} \equiv \vec{v} + \frac{\vec{q}}{2\mu_N} + \frac{\delta}{|\vec{q}|^2}\vec{q} = \vec{v}^{\perp}_{\rm el} + \frac{\delta}{|\vec{q}|^2}\vec{q}, \tag{3.40}$$

where $\vec{v}^{\perp}_{\rm el}$ is exactly $\vec{v}^{\perp}$ from Eqn 3.7.

With the inclusion of $\vec{v}^{\perp}_{\rm inel}$, we now once again have four Galilean and Hermitian invariant quantities for inelastic scattering: $i\vec{q}$, $\vec{v}^{\perp}_{\rm inel}$, $\vec{S}_{\chi}$, $\vec{S}_N$. To leading order in velocity, all that has changed in transitioning to inelastic dark matter scattering is a re-definition of $\vec{v}^{\perp}$. Thus, Eqn. 3.8, which lists all allowed operators formed from the invariant quantities in the elastic case, is still valid. The only change required is the replacement of $\vec{v}^{\perp}$ with $\vec{v}^{\perp}_{\rm inel}$. Written again here with this substitution, these inelastic operators explicitly are:

$$
\begin{aligned}
&\mathcal{O}_1 = \mathbf{1}_\chi \mathbf{1}_N, \quad \mathcal{O}_2 = \left(v_{\rm inel}^\perp\right)^2, \quad \mathcal{O}_3 = i\vec{S}_N \cdot \left(\frac{\vec{q}}{m_N} \times \vec{v}_{\rm inel}^\perp\right) \\
&\mathcal{O}_4 = \vec{S}_\chi \cdot \vec{S}_N, \quad \mathcal{O}_5 = i\vec{S}_\chi \cdot \left(\frac{\vec{q}}{m_N} \times \vec{v}_{\rm inel}^\perp\right), \quad \mathcal{O}_6 = \left(\vec{S}_\chi \cdot \frac{\vec{q}}{m_N}\right)\left(\vec{S}_N \cdot \frac{\vec{q}}{m_N}\right) \\
&\mathcal{O}_7 = \vec{S}_N \cdot \vec{v}_{\rm inel}^\perp\,, \quad \mathcal{O}_8 = \vec{S}_\chi \cdot \vec{v}_{\rm inel}^\perp\,, \quad \mathcal{O}_9 = i\vec{S}_\chi \cdot \left(\vec{S}_N \times \frac{\vec{q}}{m_N}\right) \\
&\mathcal{O}_{10} = i\vec{S}_N \cdot \frac{\vec{q}}{m_N}, \quad \mathcal{O}_{11} = i\vec{S}_\chi \cdot \frac{\vec{q}}{m_N}, \quad \mathcal{O}_{12} = \vec{S}_\chi \cdot \left(\vec{S}_N \times \vec{v}_{\rm inel}^\perp\right) \\
&\mathcal{O}_{13} = i\left(\vec{S}_\chi \cdot \vec{v}_{\rm inel}^\perp\right)\left(\vec{S}_N \cdot \frac{\vec{q}}{m_N}\right), \quad \mathcal{O}_{14} = i\left(\vec{S}_\chi \cdot \frac{\vec{q}}{m_N}\right)\left(\vec{S}_N \cdot \vec{v}_{\rm inel}^\perp\right) \\
&\mathcal{O}_{15} = -\left(\vec{S}_\chi \cdot \frac{\vec{q}}{m_N}\right)\left(\left(\vec{S}_N \times \vec{v}_{\rm inel}^\perp\right) \cdot \frac{\vec{q}}{m_N}\right).
\end{aligned}
\tag{3.41}
$$

One implicit change that does occur is that coefficients for these operators may be complex as long as the constraint that they appear in complex conjugate pairs in the Hamiltonian is satisfied. More details are given in Appendix C of [27]. The final step is to understand the consequence of the change to $\vec{v}_{\rm inel}^\perp$ for the nuclear responses derived previously. In deriving the nuclear responses, we had previously split $\vec{v}^\perp$ into two components (Eqns. 3.9 to 3.11). We can do the same for the inelastic case:

$$
\vec{v}_{\rm inel}^\perp = \vec{v}_{{\rm inel}\,T}^\perp + \vec{v}_{\rm N} \tag{3.42}
$$

with $\vec{v}_{{\rm inel}\,T}^\perp$ defined explicitly as:

$$
\vec{v}_{{\rm inel}\,T}^\perp = \frac{1}{2}\left(\vec{v}_{\chi_1} + \vec{v}_{\chi_2} - \vec{v}_{T_{\rm in}} - \vec{v}_{T_{\rm out}}\right) + \frac{\delta}{|\vec{q}|^2}\vec{q}. \tag{3.43}
$$

It is important to remember that the nuclear form factors, as derived previously, only depended on interactions with the nucleons that form the target nuclei. As a result, the $\vec{v}_{\rm N}$ that is used as an operator to derive the nuclear response functions is identical to the elastic case. The mass splitting term only shows up in $\vec{v}^{\perp}_{\text{inel } T}$. In fact, we can follow the derivation of the differential rate from Eqn. 3.12 to Eqn. 3.31 with just replacing $\vec{v}^{\perp}_{T}$ with $\vec{v}^{\perp}_{\text{inel } T}$.

3.3 Current Experimental Limits

There are two dual-phase Xenon TPC direct dark matter experimental searches thus far that have looked at setting limits on the coupling strengths of the individual elastic and inelastic operators. These are the LUX and XENON100 experiments. In presenting these results here, many of the details of the analyses will be deferred to Chapter 4. While the size of the TPCs used have increased, detector and physics modeling have improved, and the statistical framework has become more robust, the core components of the analyses remain the same. After all, in the continuing null result scenario, each experiment sets ever more stringent limits on the coupling strengths of each operator. Since the XENON100 results were published prior to the LUX results, it shall be presented first in this section.

In looking at Eqn. 3.31, we see that regardless of whether we have converted from individual operator form factors to nuclear response form factors in writing the differential rate equation, we have a very large parameter space. In the (n, p) basis, we have 28 operator coupling parameters and a mass splitting term, δ, for the inelastic extension. It is imprac-

tical to be able to carefully consider such a large parameter space, especially if considering possible interference between the various operators. As a result, XENON100 has adopted a convention of only treating one operator as active at a time and LUX has maintained it. We shall follow this in Chapter 4 as well. Furthermore, they have also decided to work in the isoscalar basis, where operator couplings to protons and neutrons are equal. As mentioned, XENON100 (and LUX) has used the Mathematica package by Anand et al. [42] to produce the signal models (expected differential rates) for each operator. For the inelastic case, they considered mass splittings between 0 - 300 keV. In both cases, nuclear recoil energy depositions up to 240 keV_{nr} were used in the analysis. By convention, their results scale the operator coefficients by M^2_{weak}(246.2 GeV^2), which is the vacuum expectation value in the SM. The coefficients, in general, contain units of mass^{-2} and [42] picks this scaling factor to convert the coefficients to dimensionless quantities. This convention has been adopted by XENON100, LUX and LZ. Figure 3.2 shows the final elastic operator coupling strength limits for a range of dark matter masses from 10 GeV to 1 TeV. In the inelastic case, a single mass of 1 TeV was selected and limits were presented per operator in Figure 3.3, for the range of mass splitting values as stated above.

LUX published initial results on EFT operator limits in 2020, where they only considered the elastic operators [45]. In 2021, the collaboration presented updated EFT limits based on data collected in an extended running period [46], containing results for both inelastic and elastic EFT operators. Here, we only consider these latest limits from LUX. Unlike XENON100, LUX presented elastic limits in the (n, p) basis to maintain consistency with

conventions used in their initial results. Their nuclear recoil energy range considered was between 0 - 180 keV$_{nr}$. WIMP-n limits are shown in Figure 3.4 and WIMP-p limits are shown in Figure 3.5. For the inelastic case, only mass splitting terms between 50 - 200 keV were considered. LUX, however, did use the isoscalar basis for the inelastic limits for direct comparison to XENON100 results. This is shown in Figure 3.6. In the next chapter, the inelastic limits that will be shown will be compared against the XENON100 and LUX limits.

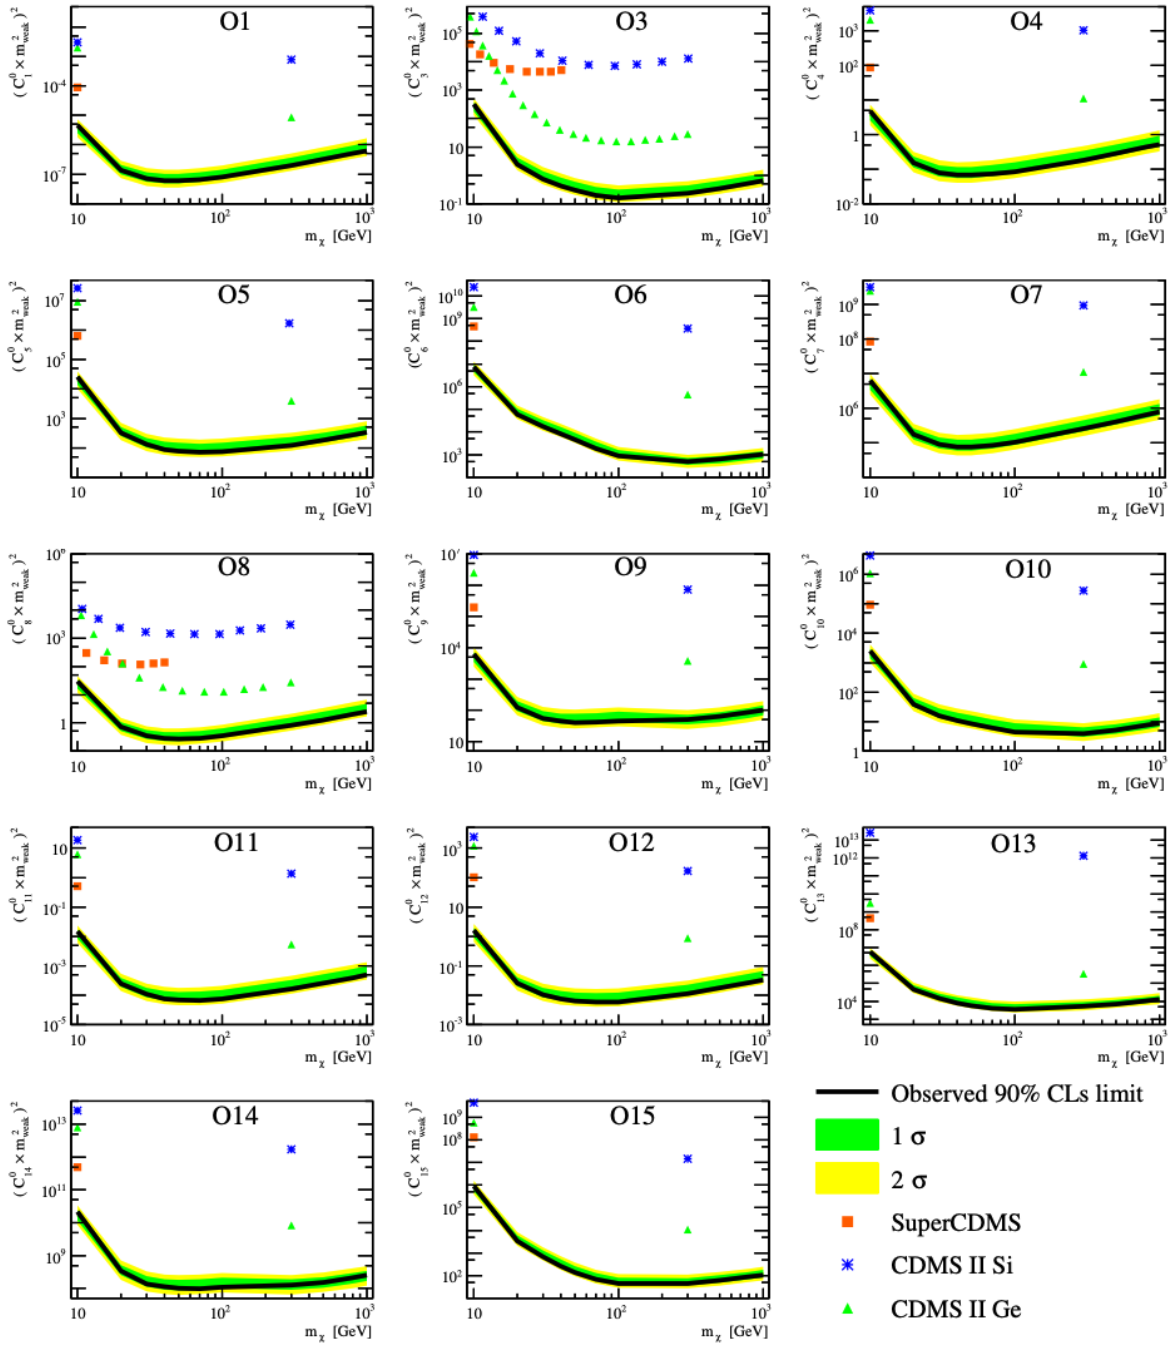

Figure 3.2: Results from XENON100's elastic EFT analysis [44]. Limits on coupling strength for each operator are shown.

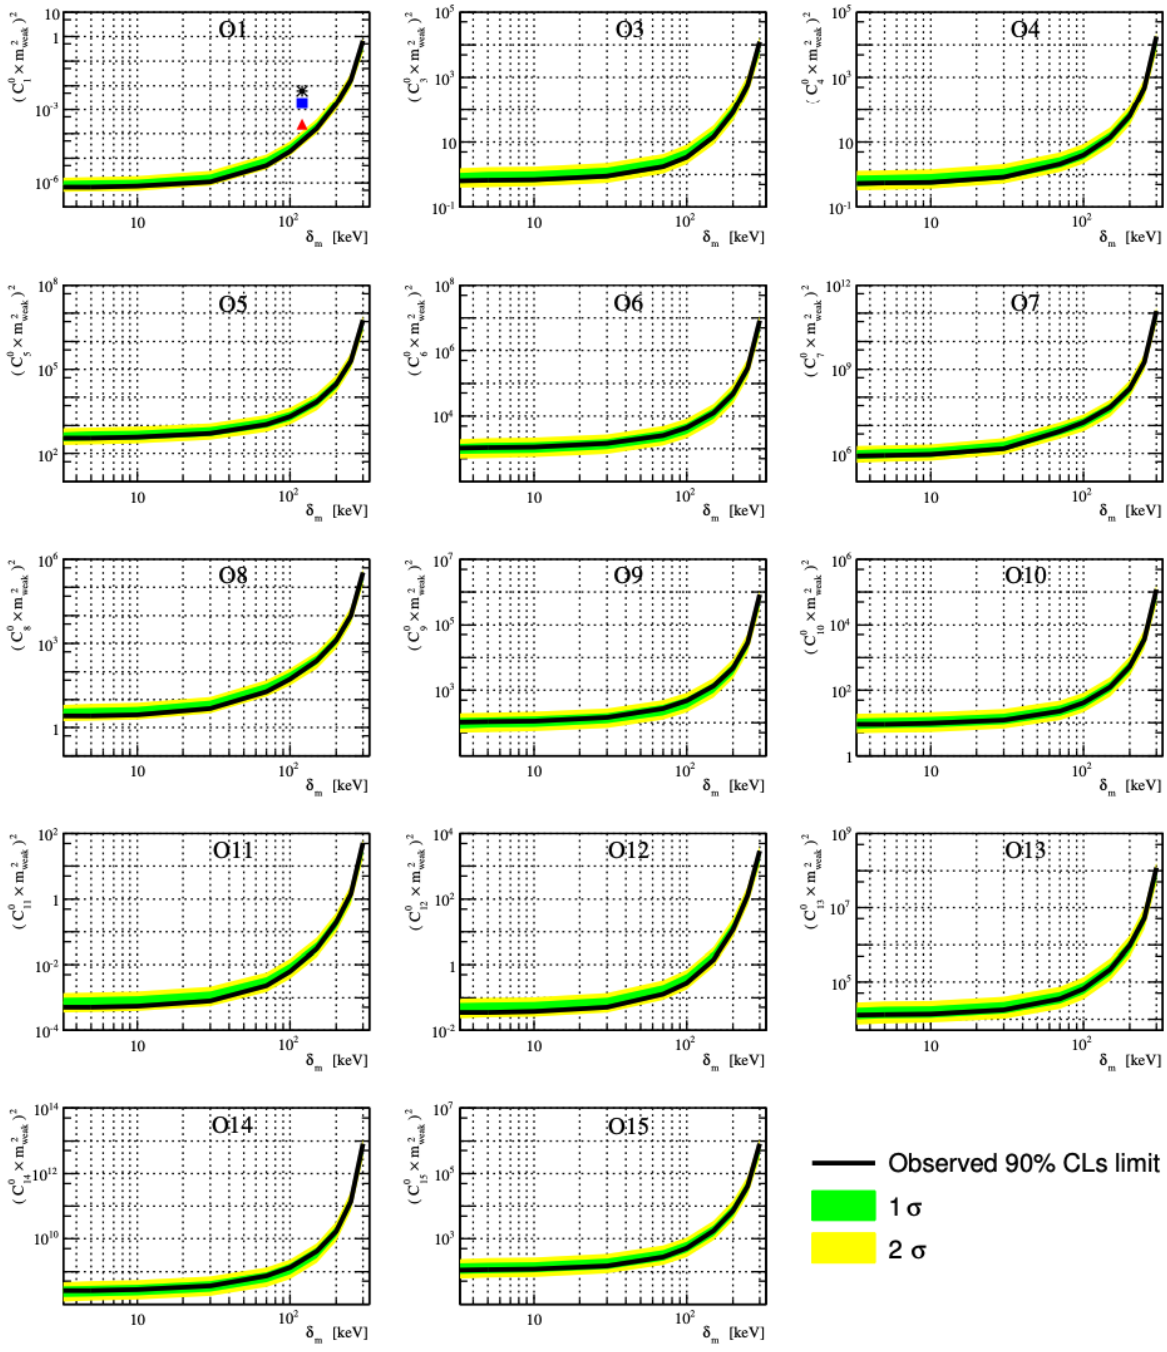

Figure 3.3: Results from XENON100's inelastic EFT analysis [44]. Limits on coupling strength for each operator are shown.

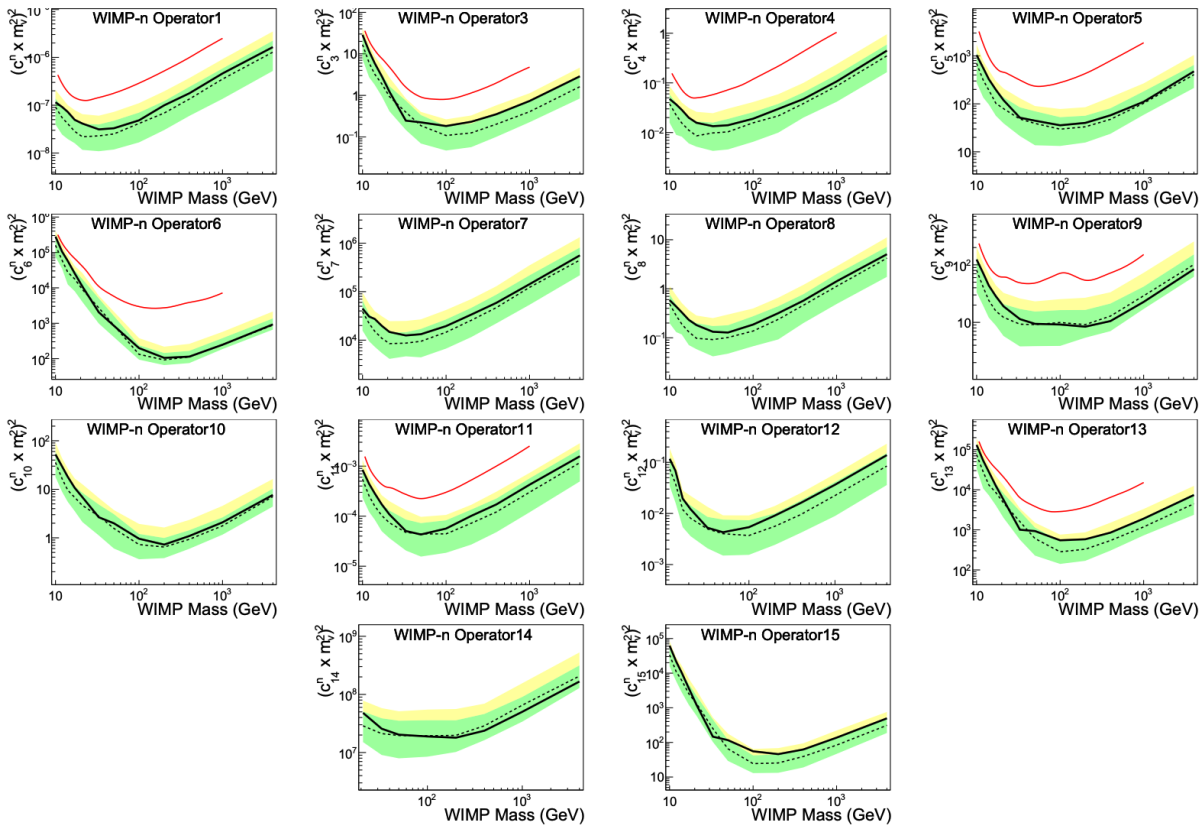

Figure 3.4: Results from LUX's elastic EFT analysis [46]. Limits on coupling strength for each operator (WIMP-n) are shown. Solid black lines represent the limit, dashed black lines indicate the expected limit, and red lines show limits from previous runs. A range of WIMP masses from 10 GeV to 4 TeV are considered for each operator.

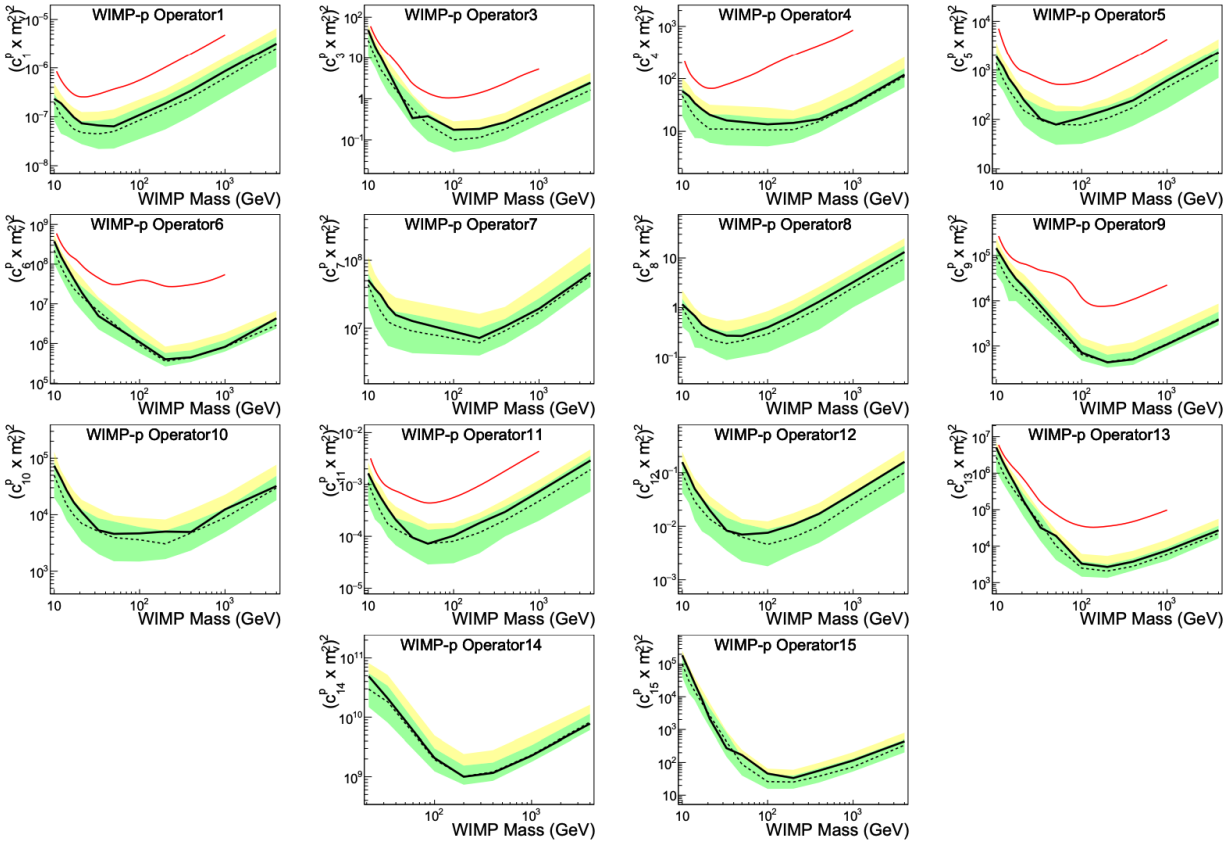

Figure 3.5: Results from LUX's elastic EFT analysis [46]. Limits on coupling strength for each operator (WIMP-p) are shown. Solid black lines represent the limit, dashed black lines indicate the expected limits, and red lines show limits from previous runs. A range of WIMP masses from 10 GeV to 4 TeV are considered for each operator.

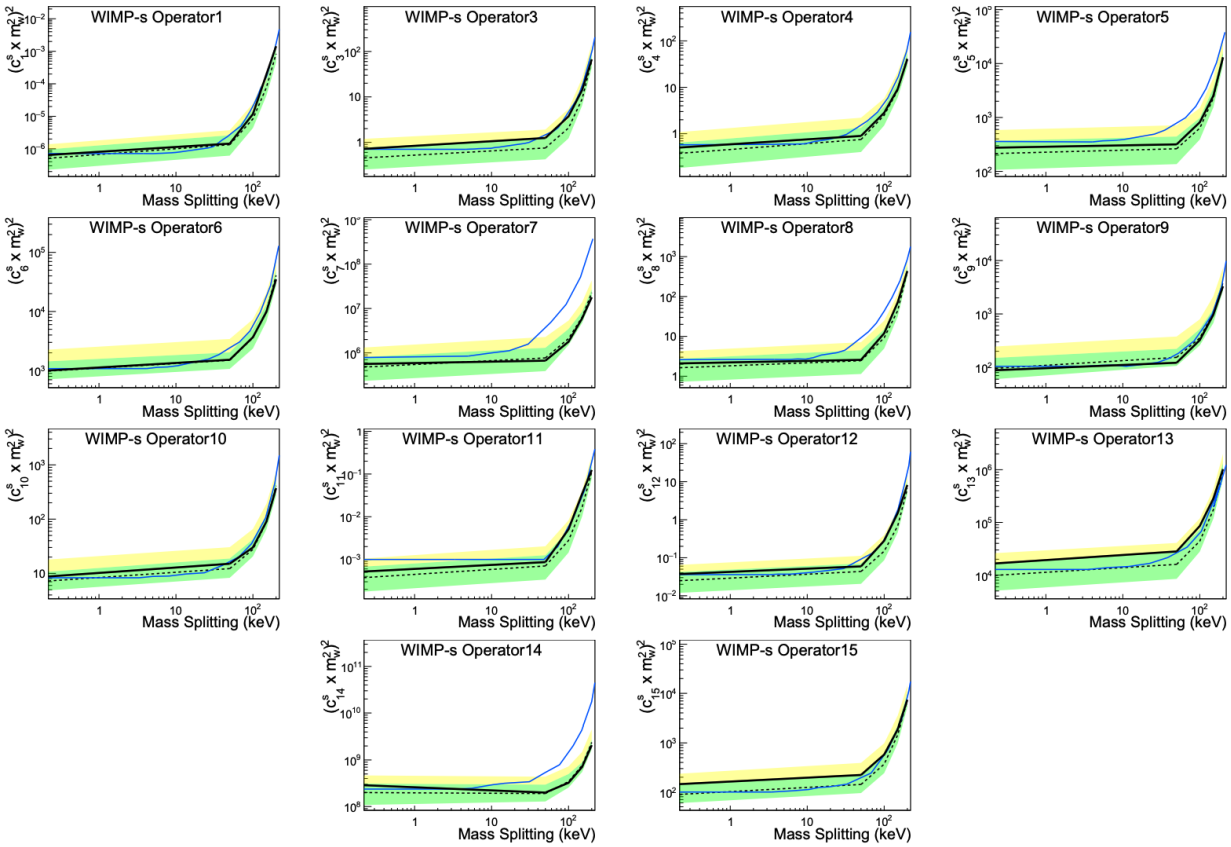

Figure 3.6: Results from LUX's inelastic EFT analysis [46]. Limits on coupling strength for each operator are shown. Solid black lines represent the limit, dashed black lines indicate the expected limt, and blue lines show limits from XENON100, which were also shown separately in Figure 3.3. A range of mass splittings from 50 to 200 keV are considered.

Chapter 4

Inelastic EFT Limits from LZ SR1 Data

We have described the operating principles of time projection chambers (TPCs), specifically the LZ detector, in Chapter 2. In this chapter, we build upon this description and present the steps undertaken in data analysis leading to setting of upper exclusion limits on the interaction strengths for inelastic EFT operators that were derived in Chapter 3. A total of 14 operators and 6 mass-splitting terms from 0 keV to 250 keV for a 1 TeV mass WIMP in the isoscalar basis are considered. We performed hypothesis tests using an unbinned, extended profile likelihood ratio (PLR) method, the details of which are discussed in this chapter. In performing this upper limit setting analysis, we have adopted the conventions used in the previously published limits of the LUX and XENON100 experiments as discussed at the end of Chapter 3. In total, we performed individual limit setting for 84 pairs of operators and

mass-splitting values. The region of interest (ROI) used for the limit-setting analysis in this work is identical to the ROI used in calculating the standard WIMP search upper limits on spin independent and spin dependent cross sections as recently published in [41]. We shall refer to this ROI as the standard WIMP search ROI from now on.

However, an exhaustive EFT based WIMP search will require an extended ROI for performing limit-setting, but is beyond the scope of this work. The reason for this requirement will be explained in Section 4.3.2. The inelastic limits presented in this chapter are intended to be an early look into the LZ science run one (SR1) data. It is meant to exercise the full limit setting analysis procedure using the data collected, but using the standard WIMP search ROI. A corresponding analysis that set limits on elastic interactions has been conducted by the LZ collaboration. In order to perform limit-setting in an extended ROI, there is extensive additional effort required in order to understand new backgrounds, perform detector modelling, and characterize the extended ROI by the LZ collaboration. One such required study evaluates acceptances after imposing event selection criteria in the extended ROI. Details and results of this study are presented while discussing S1- and S2-based data selection cuts (and corresponding acceptances) in Sections 4.2.2 and 4.2.3. Expectations for improvements of the calculated upper limits, moving from the standard WIMP search ROI to the extended ROI, are discussed in Section 4.3.2 and the last section of this chapter.

4.1 Detector Calibrations

Before undertaking the analysis of the data collected for SR1, we need to discuss several detector calibrations that were performed using a variety of radioactive sources to understand and characterize specific detector efficiencies. The full compliment of radioactive sources available for calibration work is detailed in [39] and reproduced below in Table 4.1. In the following subsections, we will summarize the results of key calibrations performed during a period of time pre- and post-SR1.

Nucleide	**Type**	**Energy [keV]**	$\tau_{1/2}$
^{83m}Kr	γ	32.1, 9.4	1.83h
^{131m}Xe	γ	164	11.8d
^{220}Rn	α, β, γ	various	10.6h
^{3}H	β	18.6 endpoint	12.5y
^{14}C	β	156 endpoint	5730y
^{241}AmLi	(α, n)	1,500 endpoint	432y
^{252}Cf	n	Watt spectrum	2.65y
^{241}AmBe	(α, n)	11,000 endpoint	432y
^{57}Co	γ	122	0.74y
^{228}Th	γ	2615	1.91y
^{22}Na	γ	511,1275	2.61y
^{60}Co	γ	1,173; 1,333	5.27y
^{133}Ba	γ	356	10.5y
^{54}Mn	γ	835	312d
^{88}YBe	(γ, n)	152	107d
^{124}SbBe	(γ, n)	22.5	60.2d
^{205}BiBe	(γ, n)	88.5	15.3d
^{206}BiBe	(γ, n)	47	6.24d
DD	n	2,450	N/A
DRef.	n	$272 \rightarrow 400$	N/A

Table 4.1: A summary of the radioactive sources available for the LZ detector calibration work and their key properties, from [39].

4.1.1 Position Correction

One key set of calibrations address spatial variations in signal response across the TPC active volume in (x,y,z) using an injected in-situ ^{83m}Kr source. This isotope produces two conversion electrons with energies listed in Table 4.1; the two decays are temporally separated by a mean time of 154 ns [47]. Since ^{83m}Kr is expected to mix uniformly throughout the TPC volume, it is an excellent source for understanding the spatial variations in signal response. The ^{83m}Kr nuclei are injected into the TPC through the xenon re-circulation system, and decay away with a half-life of 1.83 hours, without leaving behind any residual radio-activity in the longer term.

A special reduced quantity (RQ) variable was developed and defined in order to tag these decays. During dedicated SR1 ^{83m}Kr injections, the tagged events were used to characterize S1 and S2 signal response across the full detector volume. For S1 signals, correction functions were defined in both (x,y) and drift time to normalize S1 response to the geometric center of the TPC. One may recall that drift time is directly proportional to the depth of the event along the z-axis. The corrected S1 signal is referred to as S1c. On average, these corrections were about 9%, primarily due to variations in light collection efficiency across the TPC and the quantum efficiencies of PMTs [41]. On the other hand, for S2 signals, the corrections were normalized to a location at the radial center and at the highest point of the liquid volume in the TPC, which represents the point of shortest drift time.

In analog to S1c, the corrected S2 signal is referred to as S2c. S2 correction size was about

11% on average in (x,y) and 7% in drift time (z) [41]. The difference in correction size for (x,y) vs. z can be attributed to the physical origins of the respective corrections. For (x,y), the primary cause of non-uniformity are non-operational PMTs and radial extraction field non-uniformity. For z, the primary causes are impurities present within the TPC volume, which can capture electrons during drift upwards from the interaction site.

Of course, the final defined correction function will not perfectly remove all non-uniformities in signal response. After the final S1 and S2 correction functions are applied, the final variation in S1c and S2c across the TPC volume (within ^{83m}Kr data sets) is 3%.

4.1.2 ER and NR Band Calibrations

A second key calibration campaign undertaken during SR1 was to understand the response in the ER and NR bands in log10(S2c) vs S1c space. This 2D space is used for ideal separation of ER and NR events due to differences in signal formation in LXe between ER and NR, as outlined in Chapter 2). As discussed, this results in differences of the observed light and charge yields for the two type of recoils. Due to the higher charge yields (S2c area) for ER events, we expect them to populate a band above the NR events. We use the NEST software package v2.3.7 [48] to model and simulate signal formation in LXe for the LZ experiment. The software allows a detector class to be defined, and specific detector and xenon response parameters can then be adjusted (tuned) within this class to match ER and NR calibration data. Two calibration runs were conducted for this purpose. A specially

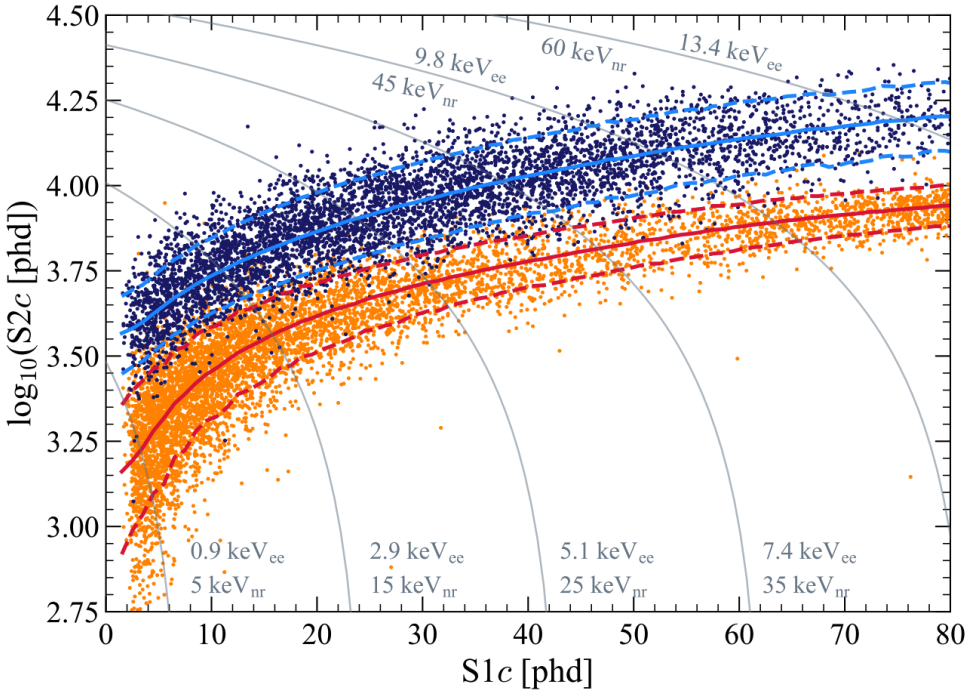

Figure 4.1: Results from LZ's NR and ER band calibration campaign for SR1 are shown. Calibration events in log10(S2c) vs. S1c for the tritium source are the dark blue points (5343 events). The DD neutron source events are shown as orange points (6324 events). Solid blue (red) lines indicate the median of the ER (NR) simulated distributions, and the dotted lines indicate the 10% and 90% quantile region. Thin grey lines show contours of constant nuclear recoil energy (keVnr) and their equivalent electronic recoil energy (keVee) [41].

developed radioactive source, consisting of methane gas (CH_4) in which one or more of the hydrogen nuclei were promoted to tritium in the presence of a neutron flux, was used for calibrations of the ER band mean and widths across log10(S2c) vs S1c space in the standard WIMP search ROI. Tritium has an end-point of 18.6 keV in its well-understood beta decay, which provides an excellent source for low-energy electrons. The details of this type of calibration can be found in [49]. These data are shown in Figure 4.1 as dark blue points.

Another source, consisting of a deuteron-deuteron (DD) neutron generator, was used for NR band calibration. This source accelerates deuteron nuclei under a high electric field, and the resulting DD collisions undergo a fusion process that produces ^{3}He and a nearly monoenergetic neutron. The NR events from this process are shown in Figure 4.1 as orange points. The tuned NEST ER band model parameters were propagated to the associated NR band model as appropriate. Comparisons of the tuned NR model against the DD calibration data showed good agreement for both NR band mean and widths across the standard WIMP search ROI. The details of DD calibrations and their usage in LXe TPC experiments can be found in [50]. Overlays of the respective tuned bands are also shown in Figure 4.1.

The final calculated g1 value is 0.114 ± 0.002 phd/photon and g2 is 47.1 ± 1.1 phd/electron. With these values, along with S1c and S2c, we are able to calculate the reconstructed energy deposition for single scatter events within the TPC. It is important to note, however, that such calibrations will need to be reevaluated in the extended ROI.

4.2 Event Selection

Once detector calibrations are performed, we move on to event selection. Due to the very small weak-scale cross section for the process, a WIMP particle is expected to scatter only once with uniform probability anywhere within the TPC's active volume, and produce an NR event. Thus, events with coincident signals in the LXe Skin or OD veto systems are readily rejected as neutron scatters. Similarly, reconstructed events that contain multiple

scatters within the TPC volume itself are also rejected. As discussed in Chapter 2, in order to mitigate against emissions from the residual radioactivity in detector walls and other detector and cavern origin events, we also defined a fiducial volume within the full TPC active volume (see Figure 4.2).

As discussed in the beginning of the chapter, the extraction of upper limits for the inelastic EFT operators in this work is constrained to the same ROI as the standard WIMP search [41]. This standard WIMP search ROI is defined by minimum and maximum values for S1c and S2c (corresponding to an NR energy range). Explicitly, the ranges were optimized to be 3< S1c <80 phd and 600< S2c <10^5 phd. In general, the primary reason for defining an ROI is to limit the energy range to a region where backgrounds are well-understood and can be modelled accurately. This is a fundamental requirement for using a PLR for hypothesis testing. While the standard WIMP search ROI was selected to provide a good acceptance for the predicted spin-independent WIMP differential rate spectrum, it covers only a fraction of the full differential rate (as a function of recoil energy) for each of the inelastic operators discussed in Chapter 3. This will be explicitly shown later in Section 4.3.2.

In an ideal world, the above are the main cuts required for selection of candidate recoil events as they reflect the underlying physics involved in WIMP-nucleus interactions. However, we do not live in an ideal world. The full SR1 data run was conducted over approximately 4 calendar months. With detector maintenance and calibration (including ^{83m}Kr) periods removed, along with periods of anomalously high trigger rates, and accounting for DAQ dead-time, SR1 consists of 89 live days of data. We also need to account for

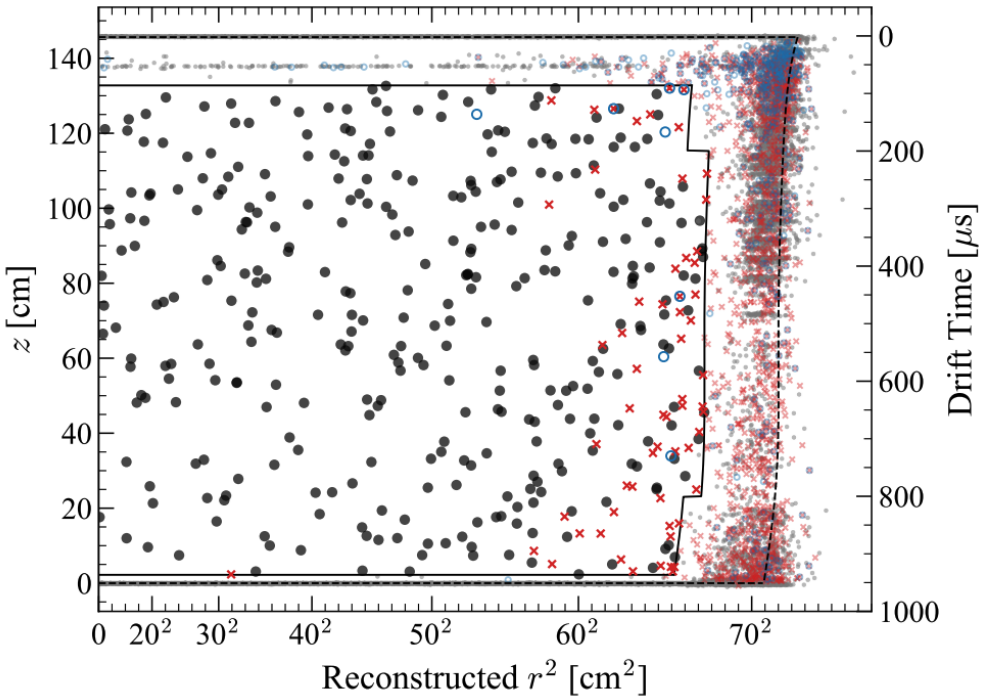

Figure 4.2: Observed data in z versus r^2 space after all SR1 analysis cuts were applied, as discussed in the text. Black (grey) points show the data inside (outside) the FV. Red crosses and blue circles show events vetoed by a prompt LXe skin or OD signal, respectively. The solid line shows the FV definition, and the dashed line shows the extent of the active TPC. Field non-uniformities cause the reconstructed radial position of the active volume boundary to vary as a function of z. Events with drift time of approximately 50 µs are from recoils in the gas which produce S1 and S2 pulses with a fixed time separation [41].

a variety of detector, electronics, event reconstruction and pulse finding inefficiencies by applying data quality cuts in order to obtain the final selection of the standard WIMP search ROI events as shown in Figure 4.2. This led to an $\sim$ 33% reduction in live time resulting in a final SR1 live time of 60 $\pm$ 1 day. These applied data quality cuts are discussed in the following section.

In addition to pure data quality cuts, there are several other cuts applied to SR1 data. These can be classified into cuts based on S1- or S2-based event reconstruction and those targeting accidental S1 and S2 combinations. They do not exclude periods of time within the full SR1 data set, but rather remove specific triggered events from consideration. As a result, we independently calculate the efficiency of the respective S1- and S2-based cuts and quantify their effect on the final NR acceptance curve as a function of energy in the following subsections.

4.2.1 Data Quality Cuts

Most of the cuts that fall into this category are designed to mitigate against pile-up accidentals, which comprise of multiple unrelated single electron pulses that can form an S2-like pulse, and similarly multiple dark count single photon pulses that combine to produce an S1-like pulse. Further, during certain periods of detector operation, which had elevated rates of accidentals, resulted in formation of spurious S1 + S2 coincidence events. These accidental combinations are caused when uncorrelated S1 and S2 pulses accidentally pair up in time to mimic a true single scatter event. Because the pulse finding software that was employed is not perfect, this can lead to a non-zero contribution of events in the NR band in log10(S2c) vs S1c space.

In addition to pile-up accidentals, isolated S1 pulses can also be produced by interactions in charge insensitive TPC regions and from fluorescent light produced by detector materials.

Isolated S2 pulses, on the other hand, can be produced by electron emission from the cathode or gate grids, particle interactions in the gas phase or in the liquid above the gate electrode (resulting in no distinct S1 signal) and from electrons trapped on impurities, which are then released with O(100 ms) time delays [51]. It has been observed in both the LUX [52] [51] and XENON1T experiments [53] that dual-phase Xenon TPCs experience elevated electron and photon rates in "hot spots" or locations on the TPC grids that spuriously (in time) produce high rates of electron emissions. These emissions can also result in elevated rates of accidentals [51].

Cut Name	**Category**	**Description**
Hot Spot Exclusion	electron emission from grids	Periods of time with elevated single electron (SE) and multi-electron emission from grids can be a source of increased rates of accidentals and pile-up. Because these periods are spurious in time, the data set is binned into small time periods and an average SE/multi-electron rate is calculated. If it exceeds the set threshold rate, a fixed window around those elevated rate time bins are removed.
Muon Veto	"after-glow" from TPC crossing muons	When muons cross the TPC, they produce large S2 pulses. This results in elevated rates of accidentals and pile-up for some time afterwards. This cut tags muon events and applies a fixed exclusion period in time after the event.

Electron/Photon Train Veto	"after-glow" from large S2 pulses	In general, all large S2 pulses will be followed by periods of elevated SE/multi-electron rates. The duration of these periods will vary with the size of the S2 pulse. This cut defines an optimized function that defines an exclusion period following every large S2 pulse that scales with the exact S2 area.
High S1 Rate Exclusion	PMT/HV anomalous behaviour	Spurious spikes in the S1 rate are observed. These periods of time are excluded to prevent elevated rates of accidentals. Implementation of the cut is similar to the hot spot exclusion cut. The main difference is that a pre-time bin window is also cut.
Bad Buffer Cut	DAQ inefficiency	The buffers used within the LZ DAQ system can only record for a certain number of us. If the pulses of interest happen to be near the end of the event window (i.e. the actual trigger was on a non-dominant S1 or S2 pulse, etc), then they will be cut off. This cut requires that the DAQ is active for an entire TPC drift time equivalent plus delayed veto windows for the OD/Skin from the start of the prominent pulses in the event.

Excess Area Cut	"after-glow" from "ghost" muons or S2 pulses	If the actual muon or large S2 pulse area containing event occurs during a DAQ dead time (i.e. a "ghost" event), subsequent events can be triggered on during the following period of electron/photon trains. This cut will remove such events by calculating the area of pulses before and between the dominant S1 and S2 pulses and make sure that this cumulative area is less than the dominant S2 pulse area.
Sustained Rate Cut	"after-glow" from "ghost" muons or S2 pulses	If the actual muon or large S2 pulse area containing event occurs during a DAQ dead time (i.e. a "ghost" event), subsequent events can be triggered on during the following period of electron/photon trains. This cut targets such cases in another manner than the above cut for completeness. The event window is subdivided into small time bins and a mean single photoelectron (SPE) rate is calculated. If it is above a fixed threshold, then the event is rejected.
Burst Noise Cut	electronics noise	Electronics burst noise is spuriously observed. Events containing such instances of burst noise are rejected.

Table 4.2: A tabulated list of all the data quality cuts applied during event selection.

In order to tackle these issues, we apply eight data quality cuts to target various subsets of pile-up and other forms of accidentals. A first cut excludes data taken during periods of high grid electron emissions during the full calendar SR1 data run. A second cut, a muon veto, excludes periods of time following a muon-tagged event within the detector. The post-muon period consists of an elevated rate of electron and photon trains, which can cause pile-up and an increased rate of accidentals. In addition, a general electron/photon train veto is also applied to remove a specified period of time following an S2 pulse for the entire SR1 data set. These time periods of removal scale with the calculated S2c pulse area. The S2 pulses considered by this third cut do not have specific origins unlike the previous two cuts. The remaining five data quality cuts target specific short-comings of the electronics/DAQ system (i.e. if a muon passes through the detector during a DAQ dead time period, or if burst (electronics) noise is observed, etc). The first three cuts, explicitly defined here, contribute to nearly all of the $\sim 33\%$ reduction in live time, while the other five cuts cumulatively remove less than 1%. All data quality cuts are explicitly listed in Table 4.2. These data quality cuts are tuned for the standard WIMP search ROI but do not need to be re-evaluated for an extended ROI analysis.

4.2.2 S2c-based Cut Acceptance

Now that we have discussed the pure data quality cuts, we can begin discussing the first sub-group of cuts that focus on inefficiencies in event reconstruction and additional cuts targeting

accidentals (the S2-based cuts). This group of cuts were initially defined and tuned for the standard WIMP search ROI. Unlike the data quality cuts, they need to be re-evaluated for validity in an extended ROI analysis.

We start by defining the cuts. The first S2-based cut removes events that have a non-physical relationship between the S2c pulse's width and the reconstructed event drift time. The cut is implemented based on [54], which established an expected physical relationship between these two variables, based on expected diffusion during the drift. It is defined as a range of allowed widths vs. drift time for various subsets of S2c pulse area. This range is more condensed at higher S2c values.

Another set of cuts within this group targets the origin for a population of isolated S2 accidentals, consisting of events occurring near the liquid surface, where drift times are very short. In these scenarios, the probabilities of either lost S1s or S1s merged into the S2 pulse are non-zero. Another accidental S2 population originates in the extraction region gas. Similar to the previous S2 accidentals population, these events can lead to merged or miss-classified S1 pulses into the associated S2 pulse. Cuts were defined to remove such events. These are the S2 width vs. drift time, narrow S2, S2 rise time, and S2 early peak cuts; all of which are detailed in Table 4.3.

Another cut also targets this population by looking at top-bottom asymmetry (TBA) in the PMT signals of the top and bottom arrays: the S2 TBA cut, which is also defined in Table 4.3. S2 pulses occurring in the gas region (especially above the anode grid events) tend to have higher values for this ratio than valid S2 pulses within the defined fiducial

volume. A final S2 XY quality cut removes reconstructed events at higher radii that have high χ^2/(num. degrees of freedom) values. Overall, as stated, this group of cuts target reconstruction inefficiencies and sources for S2 accidentals populations.

Cut Name	**Category**	**Description**
S2 Width vs. Drift Time	accidentals	A mean S2 width vs. drift time relationship can be calculated using ^{83m}Kr calibration data. +/- 3.5σ bands can also be calculated from the same data set. Events where the calculated S2 width vs. drift time value is above or below these bands are rejected.
Narrow S2	accidentals	A significant portion of accidentals population is from near liquid surface level S2 events. The corresponding S1 pulse is either swallowed or merged into the S2 pulse due to the very short expected drift times. An optimized function is defined in S2 width vs. S2 area phase space to reject this population of events.
S2 Rise Time	accidentals	This cut is targeting the same accidentals population as the cut above, however, it uses area fraction times to do so. An optimized function is defined in S2 rise time vs. S2 area phase space to provide another cut to reject this population of events.

S2 Early Peak	accidentals	Targets accidentals population consisting of merged S1 and S2 pulses due to event occurring in the extraction gas region. An optimized function is defined in full-width-half-maximum (FWHM) vs. S2 area phase space to reject this population of events.
S2 XY Quality	accidentals	The fiducial volume is defined as a function of (r,z). Thus, it is important that the reconstruction (x,y) position of edge events have a good reduced χ^2/d.o.f. value. Events that do not meet a set threshold are removed.
S2 TBA (above-anode gas)	accidentals	Above-anode gas events have higher top-bottom-asymmetry (TBA) than normal S2 pulses originating in the LXe volume. This cut is defined as an optimized function in TBA vs. S2 area phase space to reject such events.

Table 4.3: A tabulated list of all the S2-based cuts applied during event selection.

Each of these cuts was tuned with various skimmed calibration data sets for the standard WIMP search ROI. "Skimmed" in this context is defined as having the previously discussed data quality cuts applied to the complete calibration data. Furthermore, an ER or NR band cut is also implemented depending on the respective calibration source to create the final skimmed data set. The final acceptances, on the other hand, were calculated using a

skimmed file containing a combination of AmLi and tritium calibration data. The reason for this choice is that data sets corresponding to each of these calibrations could not be independently used. AmLi data covers the S2c area of interest (in the standard WIMP search ROI) but has a bad waveform environment due to the presence of many electron and photon trains following large S2 pulses. Tritium data, on the other hand, has a good waveform environment but because it is an ER calibration source, S2 areas are too large to cover the required range of interest for NRs in the standard WIMP search ROI. Thus, single scatter events were stitched together taking the S1 pulse and waveform environment from the tritium data and the S2 pulses from the AmLi calibration data. This data set was used to evaluate overall and individual S2-based cut acceptances as a function of S2c area for the standard WIMP search ROI. The S2c group cut acceptance, in general, was defined as

$$\text{S2c Accept.} = \frac{\text{events passing (livetime cuts + physics bkgd. cuts + S2 group cuts)}}{\text{events passing (livetime cuts + physics bkgd. cuts)}}. \quad (4.1)$$

For the extended ROI analysis, an extended S1c range of 3 phd to 600 phd with the same S2c from 600 phd to 10^5 phd will be used as the ROI. As we shall see in Section 4.3.2, the various inelastic EFT operators' differential rate spectrum extends to several hundred keV NR equivalent energy depositions (especially when considering higher mass-splitting values). Extending the maximum S1c value from 80 phd to 600 phd is an attempt to accept a larger fraction of the total recoil spectrum while balancing against the requirement to both firmly understand the detector and expected background in the extended ROI. One of the many

studies required to extend the ROI so as to accept a larger fraction of the complete inelastic EFT operators' differential rate spectra is to evaluate corresponding S1c and S2c acceptances in this extended ROI.

Due to the limited range of energy depositions in the SR1 AmLi and tritium calibration data sets for $80 < \text{S1c} < 600$ phd and $4.2 < \log10(\text{S2c}) < 5.0$, Rn220 was selected for calculating S1c and S2c acceptances in this part of the extended ROI. Although it is an ER calibration source like tritium, it provides the highest statistics data set among the conducted SR1 calibrations for larger S1c and S2c value recoil events. Unfortunately, a large number of the overall number of triggered events had to be removed by applied data quality cuts (especially the discussed electron/photon train veto) when creating a skimmed Rn220 data set in an analogous procedure to the one used for the stitched AmLi/tritium events that were used to evaluate standard WIMP search ROI S2c acceptances. This is an accepted behavior due to the large number of α, β, and γ decays within the full Rn220 decay chain (including daughters). The expected large S2 pulses lead to many long exclusion periods, which is the cause of the removal of a large number of triggered events. A histogram of events in the Rn220 skimmed file in log10(S2c) vs. S1c space is shown in Figure 4.3.

Although statistics limited, we are nevertheless able to calculate S2c-based group cut acceptances out to the maximum S2c value in the extended ROI. Figure 4.4 shows the final calculated acceptances using Rn220 skimmed events in black. Also overlaid on the figure in red are the calculated S2c-based group cut acceptances using the stitched AmLi and tritium data events (for the standard WIMP search ROI). The fitted standard WIMP search ROI

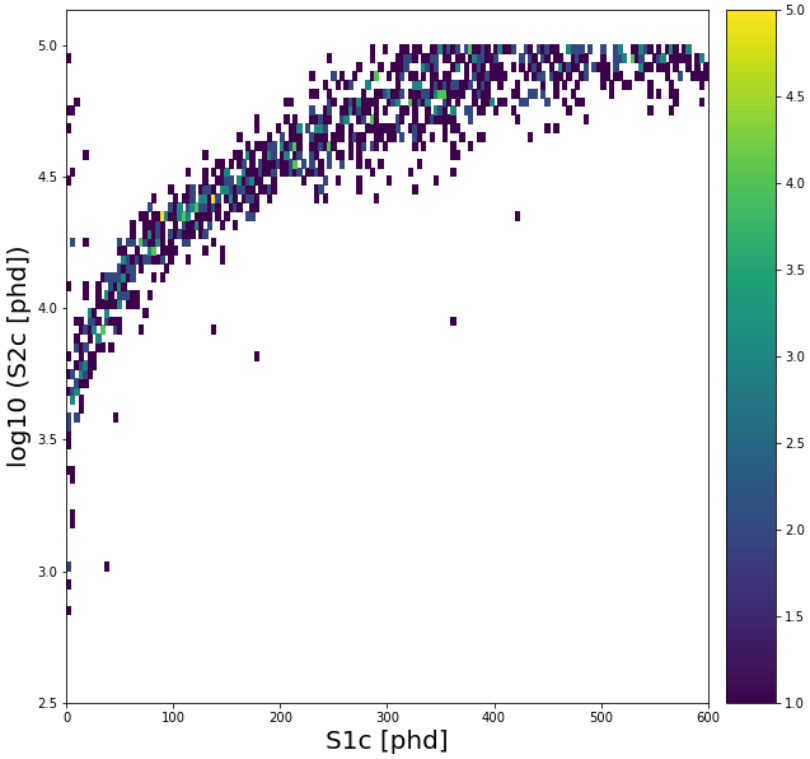

Figure 4.3: A histogram of Rn220 skimmed events in log10(S2c) vs. S1c space. Although there are limited statistics, a clear distribution of events across the entire extended S1c ROI from 3 to 600 phd is seen. Similarly, a clear distribution of events across high log10(S2c) values up to 5.0 are seen.

acceptance curve is plotted in blue (and extended out to log10(S2c) of 5). The main goal of evaluating efficiencies in the extended S2c ROI is to evaluate whether this fitted acceptance curve can continue to be extended to higher S2c values. To this end, a constant value fit was made to the Rn220 acceptances in the log10(S2c) region between 4 and 5. The returned fit is 0.979 ± 0.006. While this is not exactly equivalent to unity (as is the case for the extended fitted curve), it is also important to keep in mind the effect of limited statistics. Figure 4.5 explicitly shows the number of events removed per bin and the corresponding

acceptances using the same x-axis range. With the exception of a single bin where 3 events were removed, in the majority of events within the plateau, a single event at most is removed. The differing effects on the fitted acceptance value explicitly show the statistics limitations of SR1 calibration data available for this analysis. With this extra consideration, we can reasonably conclude, using the Rn220 skimmed data, that the same S2 acceptance curve from the standard WIMP search can be extended out to log10(S2c) of 5. This matches well with expectations since the S2-based cuts were designed primarily to remove low energy accidentals and reconstruction inefficiencies.

4.2.3 S1-based Cut Acceptance

The S1-based cuts, in analog to the S2-based cuts discussed in the previous subsection, target reconstruction inefficiencies and populations of accidental S1s. The cuts are explicitly defined in Table 4.4 and were created and tuned within the scope of the standard WIMP search ROI. Thus, much like with the S2-based cuts, we need to evaluate the validity of extending the calculated cut acceptances in the standard WIMP search ROI to the extended ROI.

We start by defining the S1-based group of cuts. In general, the pulse finding algorithm developed and used by the LZ collaboration prioritized the ability to tag single scatter events, even at the cost of potential background leakage. The S1 prominence cut in the S1-based cuts group targets this background leakage by looking for events with multiple S1s with a

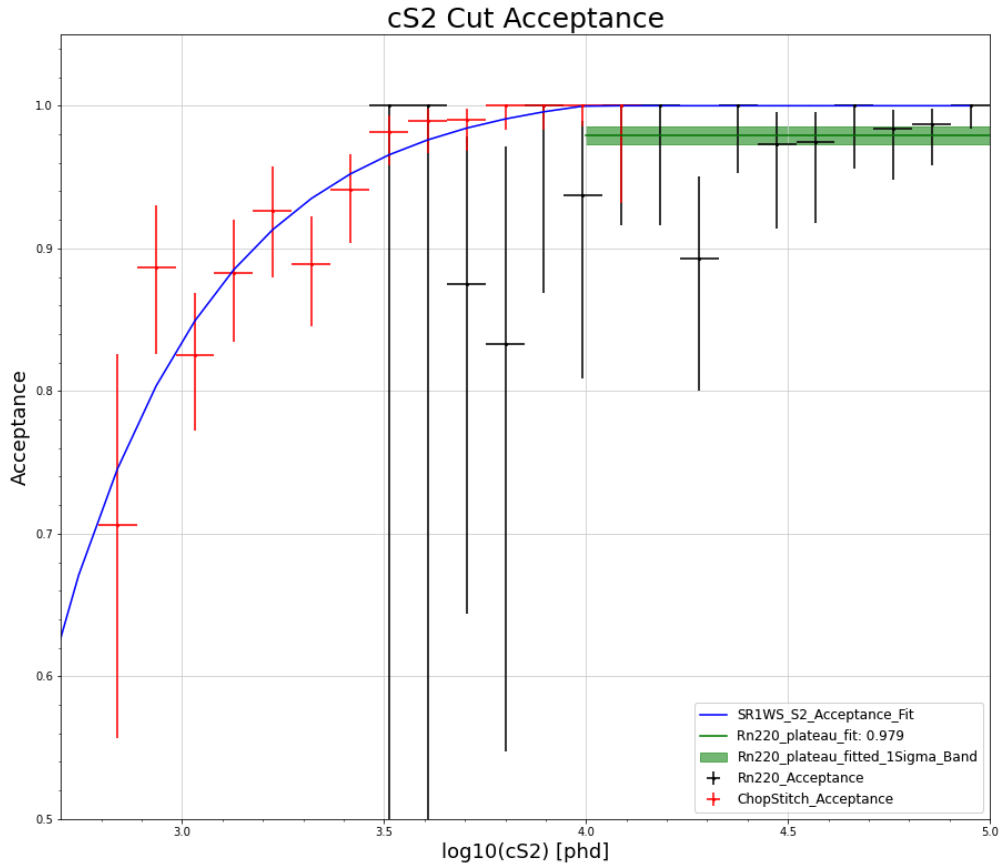

Figure 4.4: An overlay of the S2c-based cuts acceptance for the standard WIMP search ROI (using the AmLi/tritium stitched data set) is shown in red. Rn220 calculated acceptances are shown in black for an extended log10(S2c) region. A fit to this extended region in S2c plus 1σ error bands are shown in green.

specified separation in time within the event window (a single scatter event is defined by a prominent single S1 and S2 pulse within the event window). Another cut, the "stinger" event cut, targets S1 accidentals populations where delayed photons following single electron backgrounds were then paired with a lone S2.

In addition, similar to the S2-based TBA cut, an S1 analog is also implemented. S1

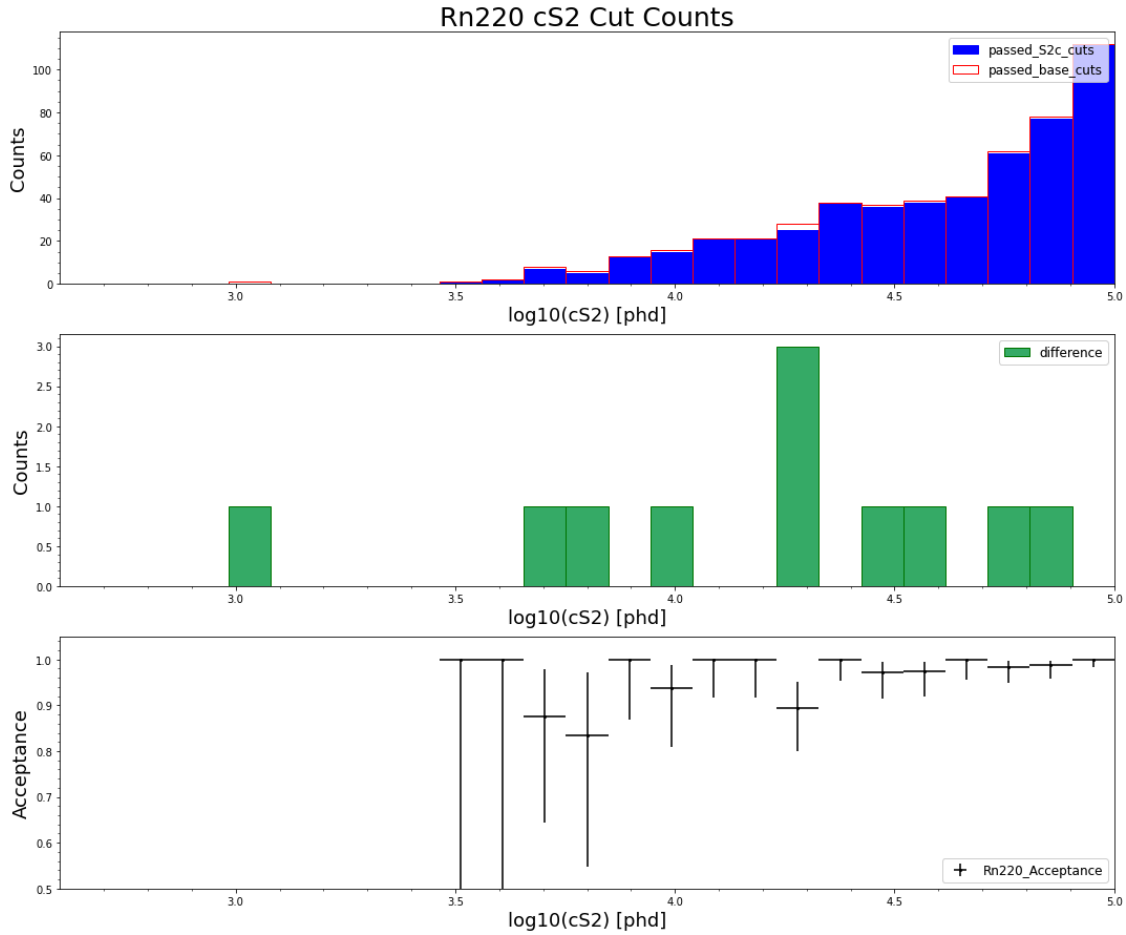

Figure 4.5: In the first row, in red, the original number of events in each log10(S2c) bin is shown. The remaining number of evts/bin is shown in blue after the S2-based cuts are applied. The second row shows the explicit number of events cut per bin. The final row re-plots the Rn220 acceptances from Figure 4.4. The dip in overall S2 acceptance in the plateau is from statistics limitations rather than a true dip in the acceptance.

photons produced are collected differently by the top and bottom arrays of PMTs depending on the location of the event along the z-axis. Events with a calculated ratio vs. drift time falling outside the expected range are removed. Furthermore, a large number of accidental S1s are attributed to events where a large fraction of the total S1 area originates in a single

PMT channel. Such events are, in most cases, traced back to above the anode grid events. The high single channel (HSC) cut is another way to cut the S1 side of this accidentals population. The last pair of cuts (S1 shape and photon timing cuts) target properties of S1 accidental populations. For example, accidental S1s tend to have larger spreads in time for peak S1 size as observed in individual PMT channel waveforms, which contribute to the summed S1 pulse. This can be defined as a spread in relative peak amplitudes for individual PMT channels or by the overall S1 width. All of these cuts are expanded on in Table 4.4.

Cut Name	**Category**	**Description**
S1 Prominence	event reconstruction quality	This cut looks at single scatter events tagged by the pulse and interaction finder software where it falsely classified the event as containing only a single S1 pulse. For all such events, this cut checks if there are S1 pulses before the tagged S1 pulse. If the separation in time is above a certain threshold, such events are cut.

Stinger Event	event reconstruction quality and accidentals	Due to the physical layout and subsequent misalignment of the gate and anode grid wires, there can be cases where some of the photons for single electron pulses are delayed. The pulse finder software incorrectly can label the delayed photons as a standalone S1 pulse. The maximum delay possible is 2us. This cut will remove all events with S1 pulses that are within 2us of a previous SE/S2 pulse and have an area less than that preceding pulse.
S1 TBA vs. Drift Time	accidentals	This cut is analogous to the S2-based equivalent. There is a physical relationship between S1 TBA and drift time. The mean and $\pm$ 3σ bands are fit using ^{83m}Kr calibration data. Any events that fall outside of these bands are rejected.

S1 HSC (high single channel)	accidentals	Several instances of events were observed where a single PMT contained a large fraction of the overall S1 pulse area. The origin for such behavior can be from above anode events and/or events occurring near the edges of the detector. Specific PMTs are also known to be culprits. The cut is defined so as to reject events where a single PMT contribution to the S1 area is above some fraction of the total S1 pulse area. The function to define the threshold fraction depends on drift time and overall S1 pulse area.
S1 Shape	accidentals	Isolated S1 pulses (contributing to accidentals) tend to have distorted pulse shape and slightly longer tails. Such S1 pulses are cut using optimized functions defined in prompt area fraction vs. S1 area phase space.
S1 Photon Timing	accidentals	This cut targets accidental S1s by looking at the spread in times where individual PMTs see peak pulse amplitudes. This will be larger for accidental S1s compared true single scatter S1 pulses. All events with S1 pulses where the spread is longer than a fixed threshold time are cut.

Table 4.4: A tabulated list of all the S1-based cuts applied during event selection.

Overall, the general S1-based cut acceptance is defined as

$$\text{S1c Accept.} = \frac{\text{events passing (livetime cuts + physics bkgd. cuts + S2 and S1 group cuts)}}{\text{events passing (livetime cuts + physics bkgd. cuts + S2 group cuts)}}. \tag{4.2}$$

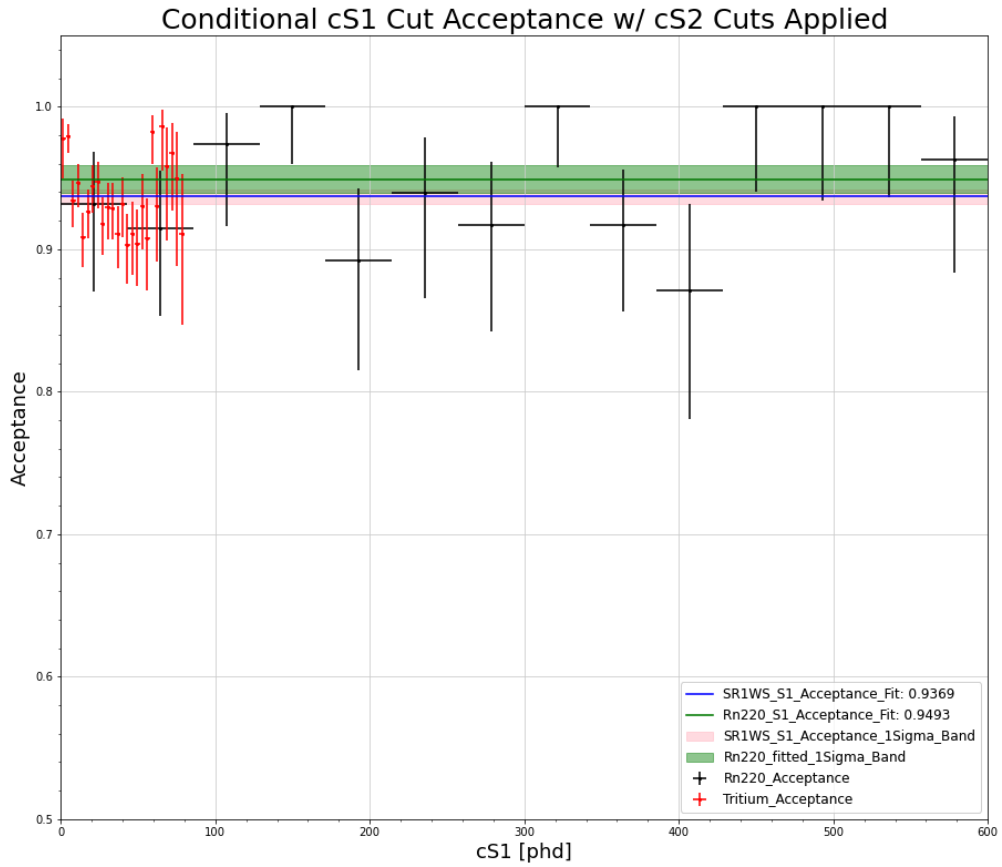

Figure 4.6: An overlay of the S1c-based cuts acceptance for the standard WIMP search ROI (using the tritium stitched data set) is shown in red. Rn220 calculated acceptances are shown in black for an extended S1c region. A constant value fit across the full energy range plus 1σ error bands are shown in green for Rn220.

For the standard WIMP search ROI, a skimmed pure tritium data set was used to evaluate

individual and overall S1c group acceptances. The electron/photon train cut applied to produce the skim file ensures that the skimmed data set has few accidentals. Furthermore, due to the low energy range for the calibration source, we obtain a sufficiently large sample of events with small S1c values. For this same reason, this calibration source cannot be used to study S1c acceptances at higher S1c values. Similar to the S2-based cut acceptance procedure in the previous section, in the extended ROI we use Rn220 for this purpose, because of the relatively large number of events across the extended S1c range (3 to 600 phd). Figure 4.6 shows an overlay of the tritium skim file acceptances (red) and the Rn220 acceptances (black) in this extended ROI. The constant value fit from the tritium skimmed data set is shown in blue (with the pink filled in region showing the $\pm$ 1σ band) and the analogous Rn220 fit is shown in green. While the exact fitted values (0.9369 $\pm$ 0.0048 for tritium and 0.9493 $\pm$ 0.0097 for Rn220) do not match, both fitted values are within 1σ of each other. What looks like large dips in the Rn220 acceptances are once again statistics limitations as seen in Figure 4.7, which clearly shows the number of events removed per bin and the acceptances using the same x-axis range. As in the case of S2-based cuts, we are able to conclude, using this Rn220 skimmed data, that the same standard WIMP search S1 acceptance curve (constant value fit) can also be extended out to 600 phd S1c. Once again, this matches well with expectations since the S1-based cuts were also designed primarily to remove low energy accidental populations and reconstruction inefficiencies.

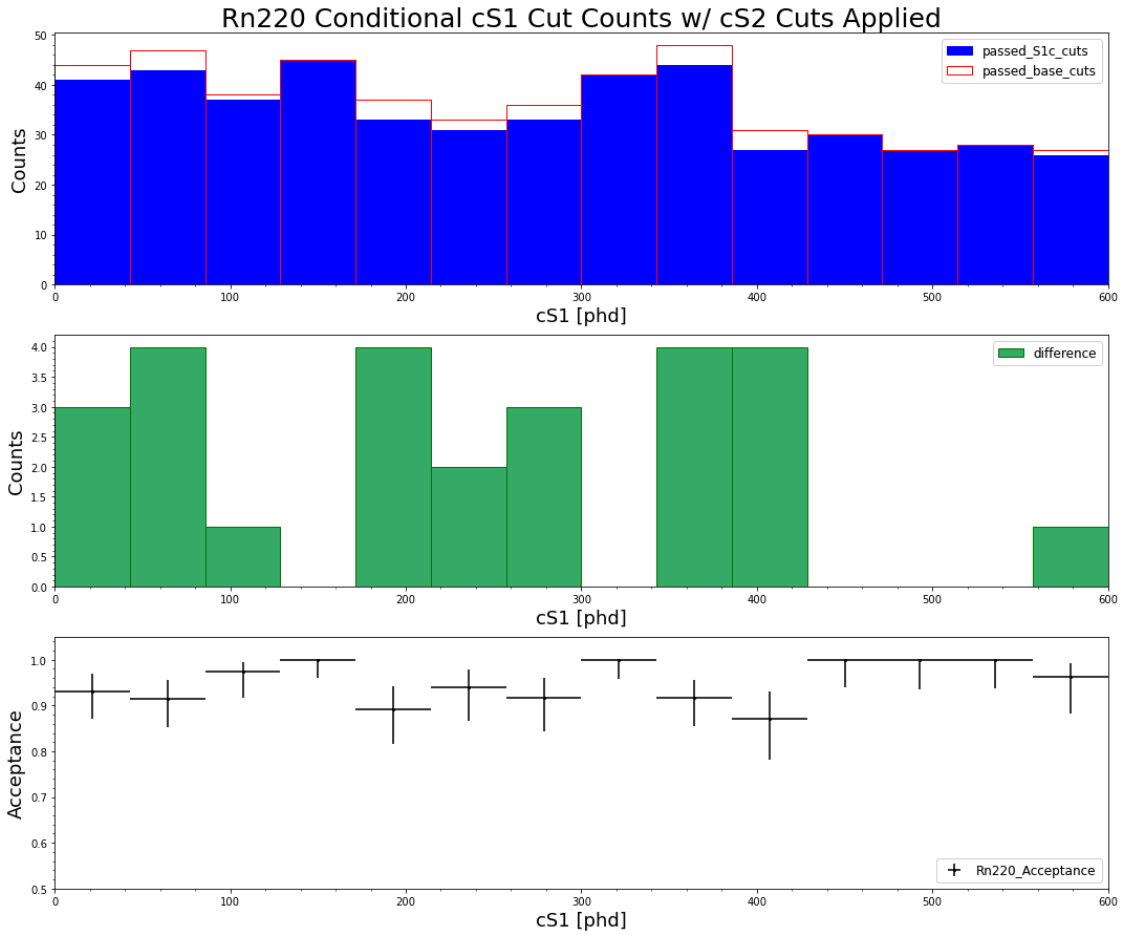

Figure 4.7: In the first row, in red, the original number of events in each S1c bin is shown. the remaining number of evts/bin is shown in blue after the S1-based cuts are applied. The second row shows the explicit number of events cut per bin. The final row re-plots the Rn220 acceptances from Figure 4.6. The first two Rn220 bins match well with the standard WIMP search S1 acceptance fit, validating the choice of Rn220 to evaluate S1-based cuts acceptance for an extended S1c ROI.

4.2.4 Nuclear Recoil Acceptance

We have discussed and explicitly shown the S1- and S2-based cut acceptances in both the standard and extended WIMP search ROIs in the preceding sections. In doing so, we dis-

cussed the results of one of the many studies required for a full extension of the limit-setting analysis presented later in this chapter to the extended ROI. The primary reasons for such an extended ROI has been discussed and the increase in differential rate spectra acceptance will be quantified in Section 4.3.2. However, a complete extended WIMP ROI analysis requires many more detailed studies to be completed by the LZ collaboration. Thus, as discussed earlier, moving forwards in this chapter, we will, by default, assume the standard WIMP search ROI, where S1c is constrained between 3 and 80 phd unless explicitly stated otherwise.

In general, having the S1- and S2-based group cut acceptances is not enough. Our end goal is to evaluate the overall NR acceptance as a function of recoil energy. This is an important component in performing limit setting for the inelastic EFT operators. To get this, there are two other efficiencies that also need to be evaluated. These are the trigger efficiency and the single scatter efficiency. An S2-based DAQ trigger was used for SR1. Its efficiency was determined by comparing the number of external DD generator triggers against the corresponding number of DAQ triggers as a function of reconstructed S2c area. The single scatter efficiency was calculated by manually identifying a large number of true single scatter events and then calculating the pulse finding algorithm's efficiency for tagging all of these manually identified events as single scatter events.

With both of these efficiencies in hand, a final NR efficiency is evaluated using a high statistics (10 million events) simulated flat NR spectrum (uniformly covering the full ROI energy range where the NR band is tuned based on calibrations as discussed). This spectrum was generated using LZ-LAMA, a NEST-wrapper software package, which assigns detector

environment variables for each simulated event. The final NR efficiency can be calculated using the acceptance curves for trigger, single scatter, S1 and S2 acceptances and the generated spectra. The true energy (as per the simulation) values for the spectra allows binned cut efficiencies to be calculated as a function of NR energy. Figure 4.8 plots the cumulative effects of each of the individual efficiency curves with the final fitted black curve showing the overall NR acceptance as a function of recoil energy for the standard WIMP search ROI. The flat-top region of the curve has an average acceptance of approx. 92.5%.

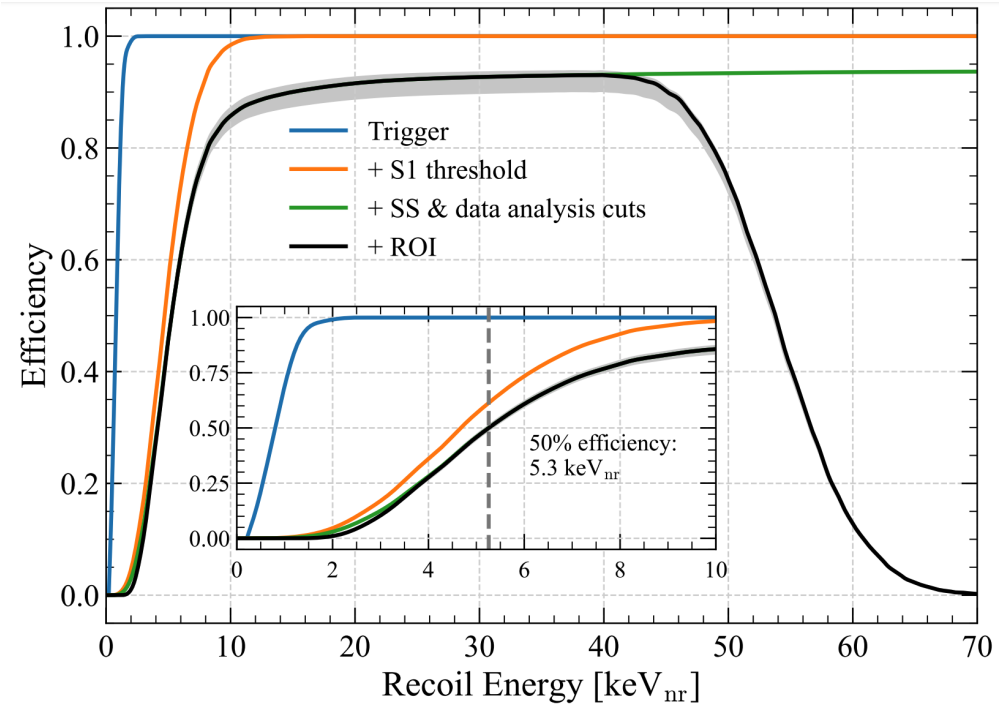

Figure 4.8: Evaluated S2-based trigger and S1 threshold cut acceptances as a function of nuclear recoil energy is shown in blue. In green, the SS and all S1c- + log10(S2c)-based cuts acceptance is plotted. A final S1c and S2c standard WIMP search ROI cut is applied and the final NR acceptance curve is shown in black. The filled grey bands show the 1-σ uncertainty for the final NR acceptance curve [41].

4.3 Background and Signal Models

4.3.1 Backgrounds

As mentioned in Chapter 1, a detailed background model is required in rare event searches. To address this point, the LZ collaboration has performed a detailed assay of all detector materials with results presented in [55]. Similarly, [56] presents results of external background simulations of cosmic-ray muons as well as neutrons and gammas from the underground cavern walls for the experiment. The full backgrounds table consists of several thousand individual contributions from these studies. For SR1, the extremely large backgrounds table is reduced into nine core components. These represent the largest contributions to the expected background for the SR1 standard WIMP search ROI. Several backgrounds that are important for a full 1000 live days science run are negligible for the 60 live days SR1 within this ROI. The vast majority of the important backgrounds are ER events. The largest contributor to the detector ER background are radioactive impurities within the detector xenon volume, Specifically, radon (Rn-220 and Rn-222 daughters), Pb-212 and Pb-214 and Kr-85. They create a cumulative broad, flat energy spectrum across the ROI. For SR1, they were grouped as a β background along with a small γ background contribution from detector materials and the walls of the underground cavern. A separate flat ER background contribution accounts for solar neutrinos within the standard WIMP search ROI.

Furthermore, the xenon used as a target medium for the experiment also contributes an ER background. These backgrounds can be subdivided into two categories. The first

category is formed by backgrounds from naturally occurring xenon isotopes (Xe-124 and Xe-136), which decay via double electron capture and double beta decay respectively. Because the xenon used in the LZ experiment was held at the surface prior to being brought underground, cosmogenic activation of xenon results in two additional primary contaminants with short half-lives, which decayed during SR1: Xe-127 and Ar-37. Xe-127 decays via L- and M-shell electron captures to I-127. If the gamma produced during nuclear de-excitation of this daughter isotope escapes without detection, the detected energy falls within the WIMP search ROI. Details of the cosmogenic production of Ar-37 for the LZ experiment are contained in [57]. Argon is an expected contaminant within the xenon but there is a large uncertainty in the expected number of Ar-37 decays during SR1, which is reflected in the constraint functions used for the PLR.

Expected sources of NR backgrounds are minimal. Radiogenic neutrons are tagged with an 88.5 $\pm$ 0.7% efficiency by the OD as calculated for SR1. Applying this efficiency to the expected neutron rate results in an expected 0 events contribution for 60 live days. This matches the radio-assay performed on detector materials in [56]. Thus, the dominant source of NR backgrounds would be from coherent elastic neutrino-nucleus scattering (CEνNS) from B-8 solar neutrinos. However, because of the S2c $>$ 600 phd requirement, the expected contribution in the standard WIMP search ROI is small. Both the cumulative contribution of these backgrounds and their individual distributions are shown as a function of ER equivalent energy in Figure 4.9. The actual observed data surviving all applied cuts is shown as black dots for each energy bin. The data points are in good agreement with the predictions, as

the eye can see. However, as a cross-check, the observed data was fit to the background only model. The data was found to agree with the background model with a p-value of 0.96.

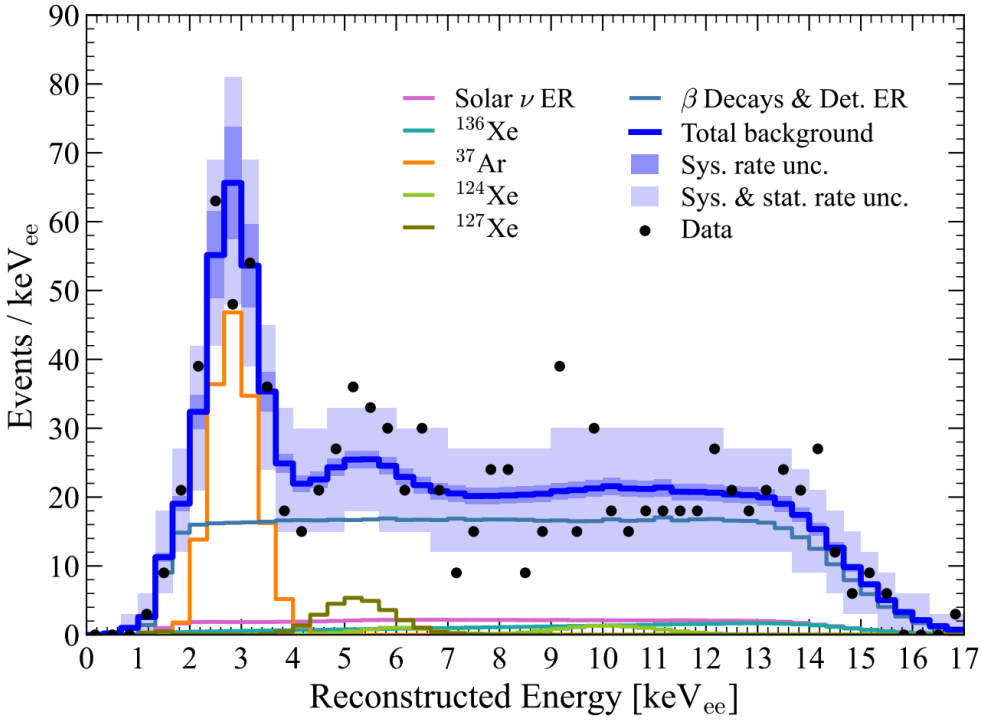

Figure 4.9: The complete background model as a function of reconstructed electronic recoil equivalent energy is shown. The 1-σ systematic rate uncertainty is shown in the filled in dark blue bands and the combined systematic and statistical rate uncertainty is shown by the filled in light blue bands. The individual components of the model are also labeled and plotted individually. The observed data points per energy bin are shown as black dots.

4.3.2 Signal Models

A differential rate spectrum for each of the 14 inelastic EFT operators and its corresponding set of 6 mass splitting terms was generated using the Mathematica package [42] as discussed

in Chapter 3. In order to directly compare limits presented in this dissertation to previous limits from LUX and XENON100 experiments (see Chapter 3), the isoscalar basis was used. In this basis, the nucleon charge densities are averaged so that a WIMP-nucleon interaction is indiscriminate to the type of nucleon involved. For the same reason, a 1 TeV mass WIMP was assumed when generating all 84 operator and mass-splitting differential rate spectra. Because the dependence on target nuclei's spin content varies between the various operators, the natural abundances of the xenon isotopes were folded into the calculation of the differential rate spectra. Additional assumptions made included using a standard halo model (SHM) [58] [59] velocity distribution for WIMPs within the Milky Way with $v_o = 238$ km/s, $v_{earth} = 252.2$ km/s, $v_{esc} = 544$ km/s, and $\rho_o = 0.3$ GeV/cm^3. A procedure that sets only a single operator coefficient to 1, and all others to 0, at a time provides an output for each operator (and mass-splitting choice) spectra, which linearly scales with the number of observed WIMP recoil events.

The spectra for each operator (and the 6 representative mass-splitting terms) are shown in Figures 4.10 and 4.11, while the acceptance as a percentage of the full spectra for each operator and mass splitting pair are listed in Table 4.5. Percentages are reported for the standard WIMP search ROI (S1c between 3 and 80 phd) and for the extended ROI (S1c between 3 and 600 phd). The percentages are calculated by integrating the differential rate spectra from the lower and upper 50% acceptance nuclear recoil energy values and dividing by the full integrated rate from 0 to 1000 keV for each respective ROI. We can see that in general, the extended ROI offers higher acceptances for all operators. For certain operators,

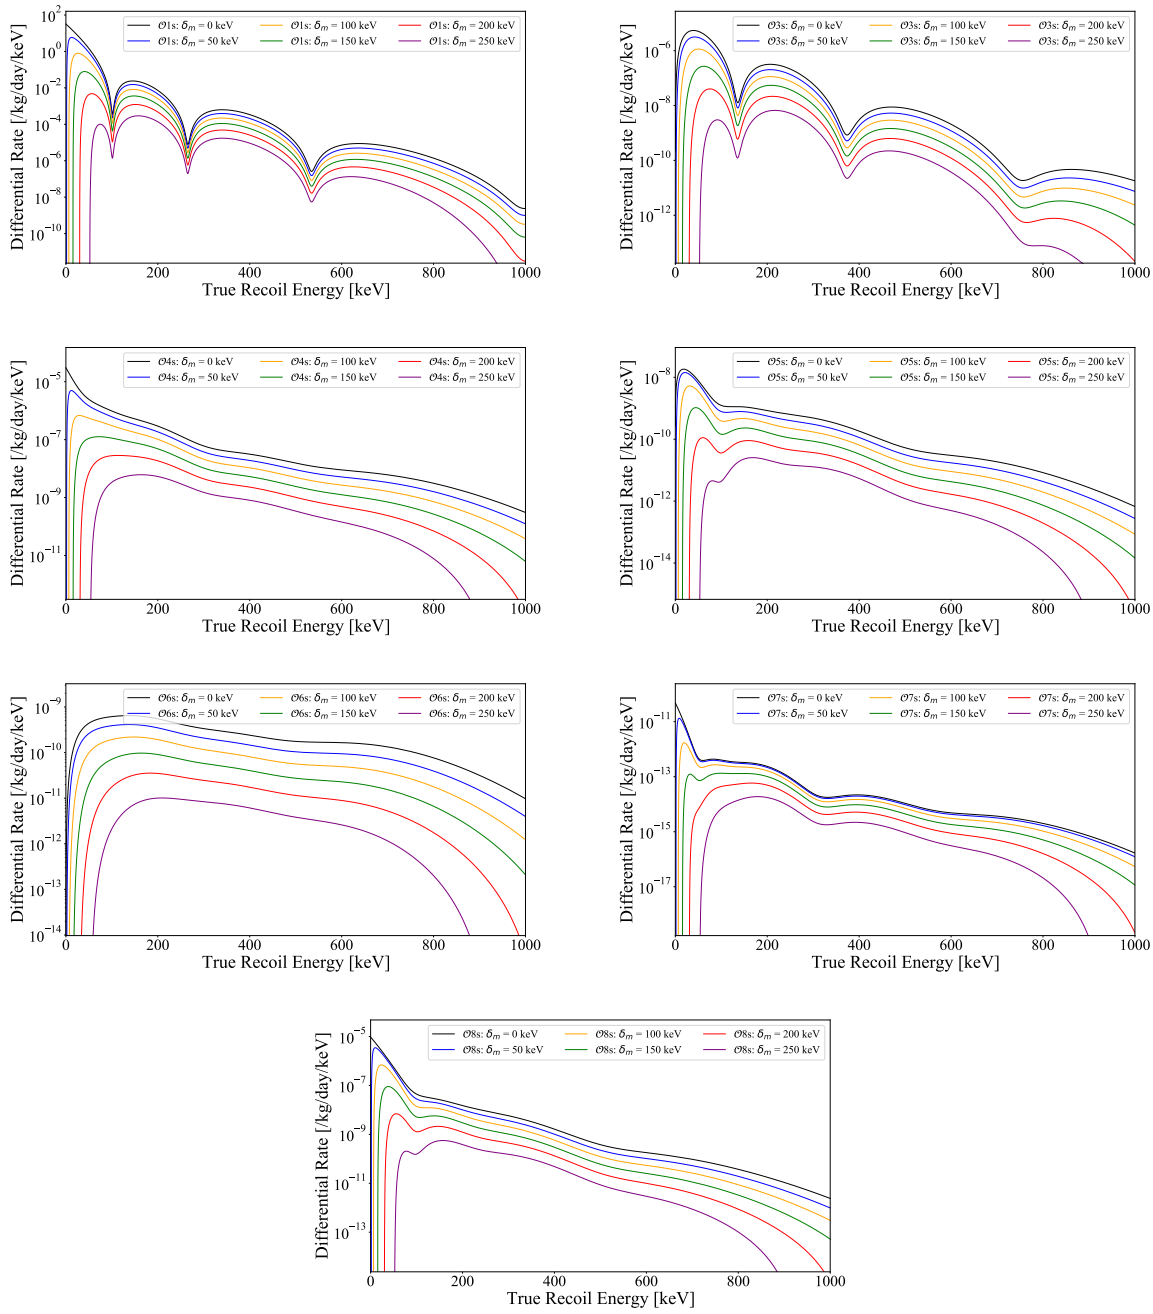

Figure 4.10: The differential rate spectra in units of /kg/day/keV as a function of true recoil energy in the isoscalar basis for a 1 TeV mass WIMP and various mass-splitting values for Operators 1, 2-8.

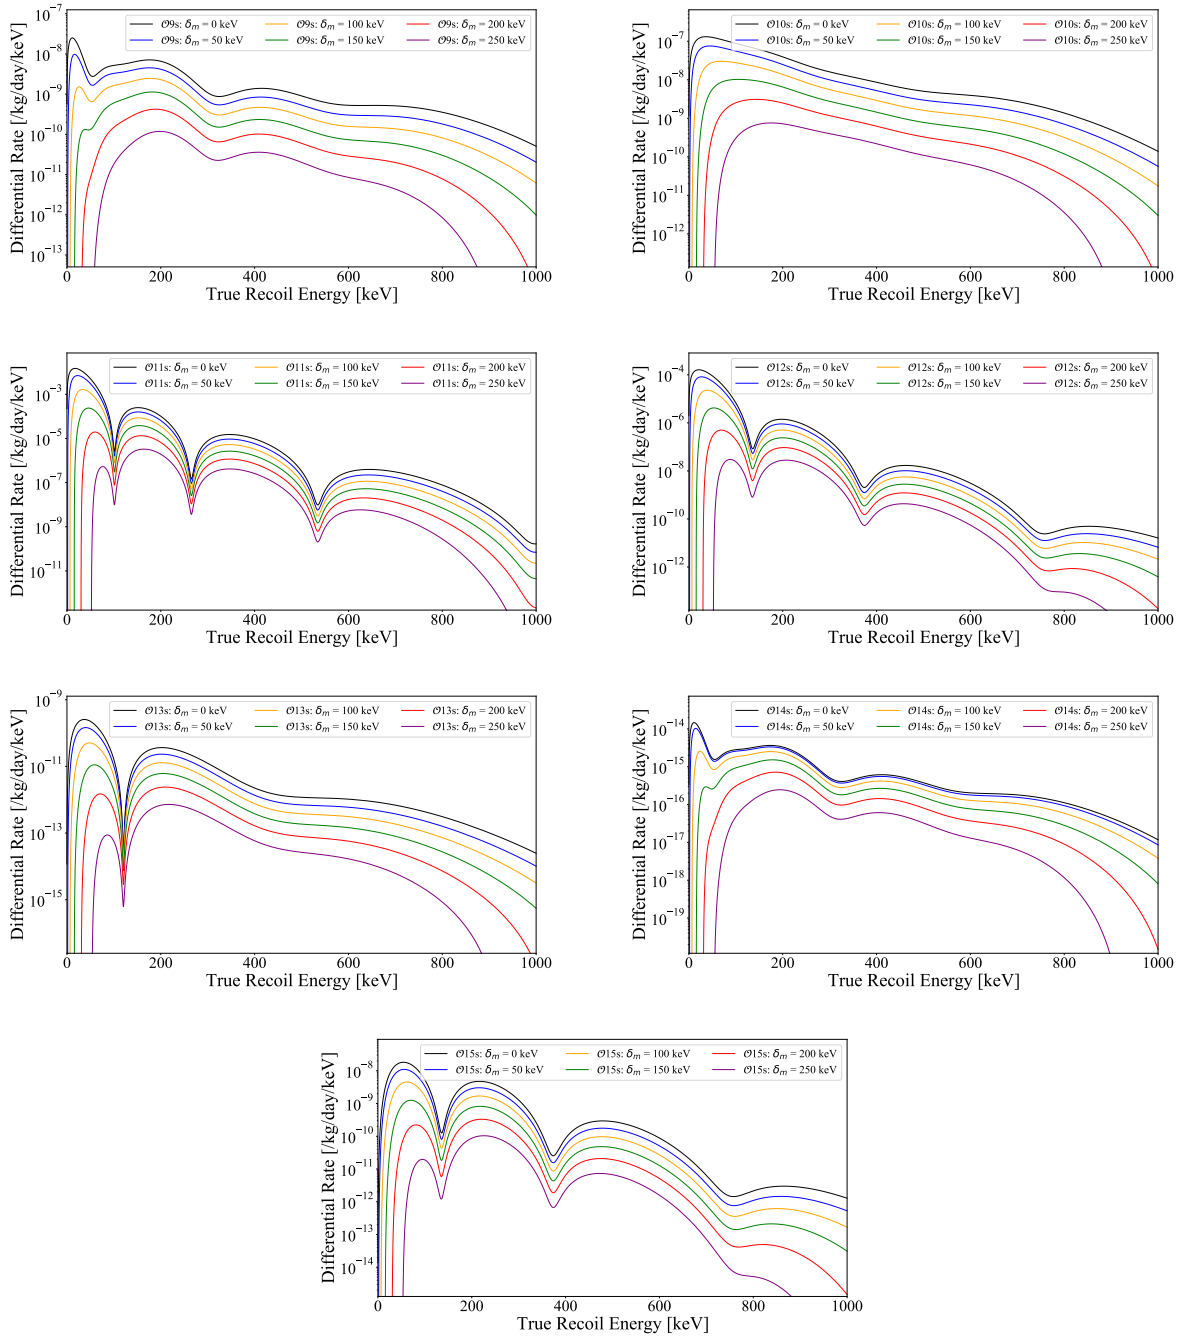

Figure 4.11: The differential rate spectra in units of /kg/day/keV as a function of true recoil energy in the isoscalar basis for a 1 TeV mass WIMP and various mass-splitting values for Operators 9-15.

where the recoil spectra is more spread out across recoil energy, the standard WIMP search ROI offers very low acceptances (for example, operators 6 and 14). Moreover, there is a very sharp drop-off in acceptance at higher mass splitting values for all operators in this ROI. We shall see that this will translate into poor upper exclusion limits at these mass-splitting values for all operators in this analysis work. An analysis using the extended ROI should not observe such behavior as this acceptance drop off is not present. This point is discussed further in the final section of this chapter.

Before continuing from this point, we shall take a brief pause to introduce hypothesis testing using PLRs and discuss how to perform limit setting analysis using this method in the next section.

Operator	**Mass Splitting [keV]**	**Acceptance (standard ROI) [%]**	**Acceptance (extended ROI) [%]**
01	0	65.77	68.00
01	50	85.58	90.01
01	100	80.49	92.42
01	150	60.55	92.14
01	200	19.8	90.60
01	250	2.73e-4	85.95
03	0	51.06	91.52
03	50	47.69	91.68
03	100	35.62	91.32
03	150	18.27	90.27
03	200	3.47	87.99
03	250	4.40e-5	84.47
06	0	4.77	54.18
06	50	4.32	55.82
06	100	2.49	54.95
06	150	0.84	52.25

06	200	0.10	48.37
06	250	1.39e-7	43.64
07	0	51.33	60.79
07	50	68.22	85.37
07	100	43.65	88.60
07	150	10.55	85.13
07	200	0.77	82.17
07	250	7.36e-7	78.16
09	0	25.14	72.54
09	50	19.51	73.85
09	100	7.17	70.88
09	150	1.43	67.93
09	200	0.11	64.59
09	250	1.01e-7	59.85
12	0	70.46	89.77
12	50	68.46	92.13
12	100	55.32	92.32
12	150	33.32	92.04
12	200	8.43	91.07
12	250	1.53e-4	88.42
14	0	28.29	75.96
14	50	25.20	77.37
14	100	11.14	74.84
14	150	2.41	71.99
14	200	0.19	69.03
14	250	1.81e-7	64.64
15	0	26.38	88.04
15	50	24.28	88.06
15	100	16.49	87.05
15	150	6.99	84.88
15	200	1.03	81.69
15	250	1.07e-5	78.21

Table 4.5: The standard WIMP search and extended ROI acceptances for a selection of operators and mass-splitting combinations is tabulated.

4.4 Profile Likelihood Ratio for Limit Setting

In the absence of excesses above the background model, an exclusion upper limit can be set based on the non-observation of a signal. In order to actually calculate such limits for each inelastic EFT operator (for a range of mass splitting terms), we use an un-binned extended profile likelihood ratio (PLR) method. Specifically, in this analysis, we set limits on the coupling strengths of the pure inelastic EFT operators based on the data in the SR1 standard WIMP search ROI.

It is much simpler to build up to a full description of the un-binned profile likelihood ratio (PLR) than to just write out the full equation from the beginning and try to explain it. Specifically, LZ has developed the LZStats package, which was originally written by Ibles Olcina Samblas [60]. The package is the standard within the collaboration to perform all exclusion upper limit setting analyses. We briefly describe the statistical framework of the PLR method here and the key details of using the software package. A more detailed description of the package and the statistical framework can be found in the referenced dissertation.

In general, events for a given experiment are defined by a set of observables, $\boldsymbol{x}_e$, with index e referring to a specific event. For this analysis, $\boldsymbol{x}_e$ is a vector of two observables: (S1c, S2c) for each event. The probability for a particular event to occur can be defined through the use of various probability distribution functions (PDFs),

$$f_i\left(\boldsymbol{x}_e \mid \boldsymbol{\theta}\right), \tag{4.3}$$

where $\boldsymbol{\theta}$ are model parameters and i indices over the many PDFs used to form a complete event probability model consisting of a signal component and several background components. Each of the PDFs have an associated mean, $\mu_i(\boldsymbol{\theta})$. In general, the model parameters, $\boldsymbol{\theta}$, can be subdivided into two categories: a parameter of interest (POI) (what we are interested in measuring) and nuisance parameters (unknown experimental parameters that have no intrinsic value but are nevertheless required to build the complete event probability model). We can write a complete event probability model as a sum over all the signal and background components,

$$f\left(\boldsymbol{x}_e \mid \boldsymbol{\theta}\right) = \sum_{i=1}^{N} \left(\frac{\mu_i(\boldsymbol{\theta})}{\mu(\boldsymbol{\theta})}\right) f_i\left(\boldsymbol{x}_e \mid \boldsymbol{\theta}\right), \tag{4.4}$$

with $\boldsymbol{\mu}(\boldsymbol{\theta})$ the total mean obtained by summing over all N individual component means. Assuming each event is independent of all others, a total probability model for a dataset, D, can be defined as

$$\begin{aligned} f(D \mid \boldsymbol{\theta}) &= \text{Pois}\left(n_0 \mid \mu(\boldsymbol{\theta})\right) \prod_{e=1}^{n_0} f\left(\boldsymbol{x}_e \mid \boldsymbol{\theta}\right) \\ &= \left(\frac{\mu(\boldsymbol{\theta})^{n_0}}{n_0!} e^{-\mu(\boldsymbol{\theta})}\right) \prod_{e=1}^{n_0} f\left(\boldsymbol{x}_e \mid \boldsymbol{\theta}\right), \end{aligned} \tag{4.5}$$

where n_0 is the total number of observed events, and the Poisson term accounts for the fact that we should expect n_0 to fluctuate around the total mean if the experiment is repeated

some number of times. Now, in general, the nuisance parameters for an event probability model can usually be constrained with *a priori* information, whether it comes from other experiments, dedicated measurements, theory, etc. Formally, we refer to this information as global observables, $\boldsymbol{g}_p$, for each nuisance parameter, ν_p. In general, $\boldsymbol{g}_p$ consists of a mean value and associated uncertainty. We can therefore define constraining functions for each of the nuisance parameters, $f_p(\boldsymbol{g}_p \mid \boldsymbol{\nu}_p)$. They are, in essence, PDFs for the $\boldsymbol{g}_p$. With the set of global observables, G, and constraining PDFs included, Eqn. 4.5 can be rewritten as

$$f(D, G \mid \boldsymbol{\theta}) = \left[\left(\frac{\mu(\boldsymbol{\theta})^{n_0}}{n_0!} e^{-\mu(\boldsymbol{\theta})} \right) \prod_{e=1}^{n_0} f\left(\boldsymbol{x}_e \mid \boldsymbol{\theta}\right) \right] \prod_{p=1}^{N_c} f_p\left(\boldsymbol{g}_p \mid \nu_p\right), \tag{4.6}$$

where N_c is the total number of constraining PDFs.

Now, we are interested in defining a profile likelihood ratio as the test statistic used for hypothesis testing (upper limit setting). A likelihood function, $\mathcal{L}$, is defined as a joint probability model of an observed dataset as a function of parameters used to define the statistical model. This is exactly what we have constructed above, because Eqn. 4.6 can be referred to as the likelihood function. However, it is prudent to mention that $\mathcal{L}$ is a function of $\boldsymbol{\theta}$ and it refers to a given experiment's observed dataset. Expressed mathematically,

$$\mathcal{L}(\boldsymbol{\theta}) = f\left(D_{\text{obs}}, G \mid \boldsymbol{\theta}\right). \tag{4.7}$$

Another way to state the role of the likelihood function is that it is used in maximum likelihood estimation (MLE) to find a set of specific values for $\boldsymbol{\theta}$ that maximize the proba-

bility of observing a given dataset. We shall return to this point but we first introduce the hypothesis testing method.

Hypothesis testing can, in general, be used to exclude a specific model or to assign a discovery significance for an observed dataset, when evaluated against a specific model. It is commonly used in direct dark matter searches for the former, and if an excess is observed it can be used for the latter. Most of the time, a frequentist approach is adopted. In this case, two competing hypotheses are compared against each other to determine which is more compatible with a given dataset of observables (S1c, S2c). These two hypothesis are called the null hypothesis (H_0) and the alternate hypothesis (H_1). Generally, H_0 is assumed true, and testing will either reject or fail to reject this hypothesis.

If we want to assign a discovery significance to an observation, H_0 is assumed to be background only, and H_1 is taken to include both a signal and the background model. Thus, the two outcomes are to either reject the background only hypothesis in favor of the signal (plus background) hypothesis or to conclude that there is not enough evidence to reject H_0. For setting upper exclusion limits, the definitions of both hypothesis are switched. The usual method is to test increasing values of the POI (number of WIMP interactions) until a value is found that can exclude the null hypothesis significantly.

In order to actually perform the hypothesis testing for both scenarios, we need to define an appropriate test statistic, t. We then either analytically or numerically calculate test statistic distributions for both hypotheses, $f(t \mid H_0)$ and $f(t \mid H_1)$. Generally, we use the latter for dark matter experiments by employing Monte Carlo techniques to generate a set

of pseudo-experiments. We also need to calculate t for the observed dataset. This is referred to as t_{obs}.

We define a significance α (usually 10%) and an associated critical region w such that if the calculated t_{obs} falls in the critical region, H_0 can be rejected with a confidence level

$$CL = (1 - \alpha)\%. \tag{4.8}$$

We can then calculate a p-value for H_0, defined as

$$\begin{aligned} p &= P\left(t > t_{\text{obs}} \mid H_0\right) \\ &= \int_{t_{\text{obs}}}^{\infty} f\left(t \mid H_0\right) dt. \end{aligned} \tag{4.9}$$

Formally, we are able to reject H_0 with confidence level, CL, if $p < \alpha$. The final step is to define explicitly the test statistic t. We use

$$t_\mu = -2 \log \lambda(\mu) \tag{4.10}$$

where $\lambda(\mu)$ is defined as profile likelihood ratio (PLR):

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\hat{v}})}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{v}})}. \tag{4.11}$$

In this equation, the single caret is referring to the maximum likelihood estimators and the double caret refers to a conditional maximum likelihood estimator for a fixed μ (POI).

Thus, the PLR is a ratio of the conditional maximum likelihood for a given μ to a global maximum likelihood.

In order to set an upper limit, we step through values of the POI, μ, until we find the value for which the null hypothesis can be rejected for a pre-set significance α. Explicitly, the POI is the number of WIMP recoil events within the observed dataset. We note that for inelastic EFT operator upper limits, μ is directly proportional to the coefficients of each respective operator. Thus, we are explicitly setting limits on the strengths of each operator coefficient.

As mentioned earlier, we use LZStats, a software package written in C++ using Root [61] libraries (RooFit and RooStats [62]) to perform upper limit setting. The package requires four main inputs: WIMP mass, data (observables for each event), signal and background; i.e. the total event probability model(s) defined as a RooWorkspace instance, and specific analysis type (upper exclusion limit setting). In general, the package is capable of performing sensitivity, discovery significance and projected discovery significance studies in addition to setting a limit. Further fine-tuning in setting up the PLR is provided via an extensive configuration list that can be modified as required for specific analyses (in this case with detector parameters as calculated for SR1). For the inelastic EFT operator limit setting analysis, workspaces (signal and background models) were generated for a static 1 TeV WIMP mass for each of the 14 operators and 6 mass splitting terms (0 to 250 keV).

4.5 SR1 Initial Inelastic EFT Limits

Returning to where we left off with the differential rate spectra, the observables used for the PLR are S1c and S2c. Thus, we need to translate the differential rate spectra for each operator and mass-splitting pair to a PDF in log10(S2c) vs. S1c space. This is done in a process analogous to generating the NR acceptance curve. We can use LZ-LAMA to sample the differential rate spectra (10 million events) and use the output to generate the PDF in the required 2D parameter space for the standard WIMP search ROI. In doing so, all SR1 S1, S2, trigger and single scatter acceptances are applied. The resulting PDFs for all mass-splittings for two operators (1 and 6) are shown in Figures 4.12 and 4.13. Similar PDFs are generated for the other 72 operator, mass-splitting pairings. The effect of the cut-off at 80 phd S1c is clearly visible in these PDFs. As seen in the differential rate spectra, operator 1 peaks at low recoil energies and thus we should see the majority of the full spectrum contained in the standard ROI in the PDFs for the smaller mass-splitting terms. At higher mass-splitting values, the bulk of the differential rate spectra is pushed to higher recoil energies and thus we will only capture a fraction of the full spectra. This behavior is clearly visible in Figure 4.12. For operator 6, we expect a very different behavior. The spectra are far more diffuse in recoil energy and peak at higher recoil energies as well. This is reflected in the signal PDFs of Figure 4.13, since we capture only the tails of the full spectra in the standard WIMP search ROI. This is especially true for the higher mass splitting values.

Each of these 84 signal PDFs are paired with a combined background model PDF, where

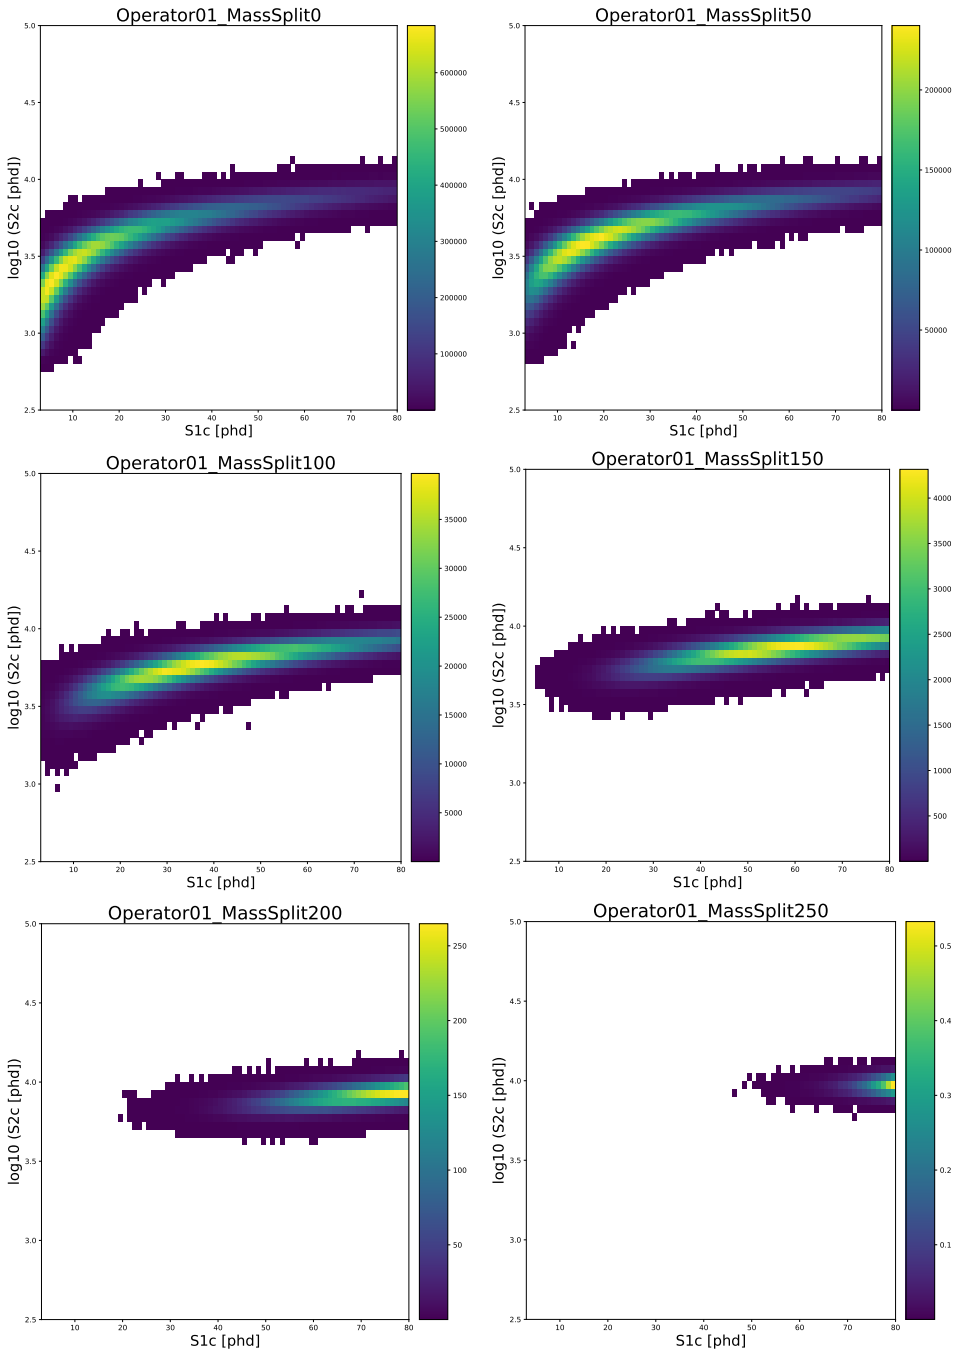

Figure 4.12: The corresponding PDF in log10(S2c) vs S1c space for Operator 1 and various mass-splitting values. All cut efficiencies and the ROI have been folded into the PDFs.

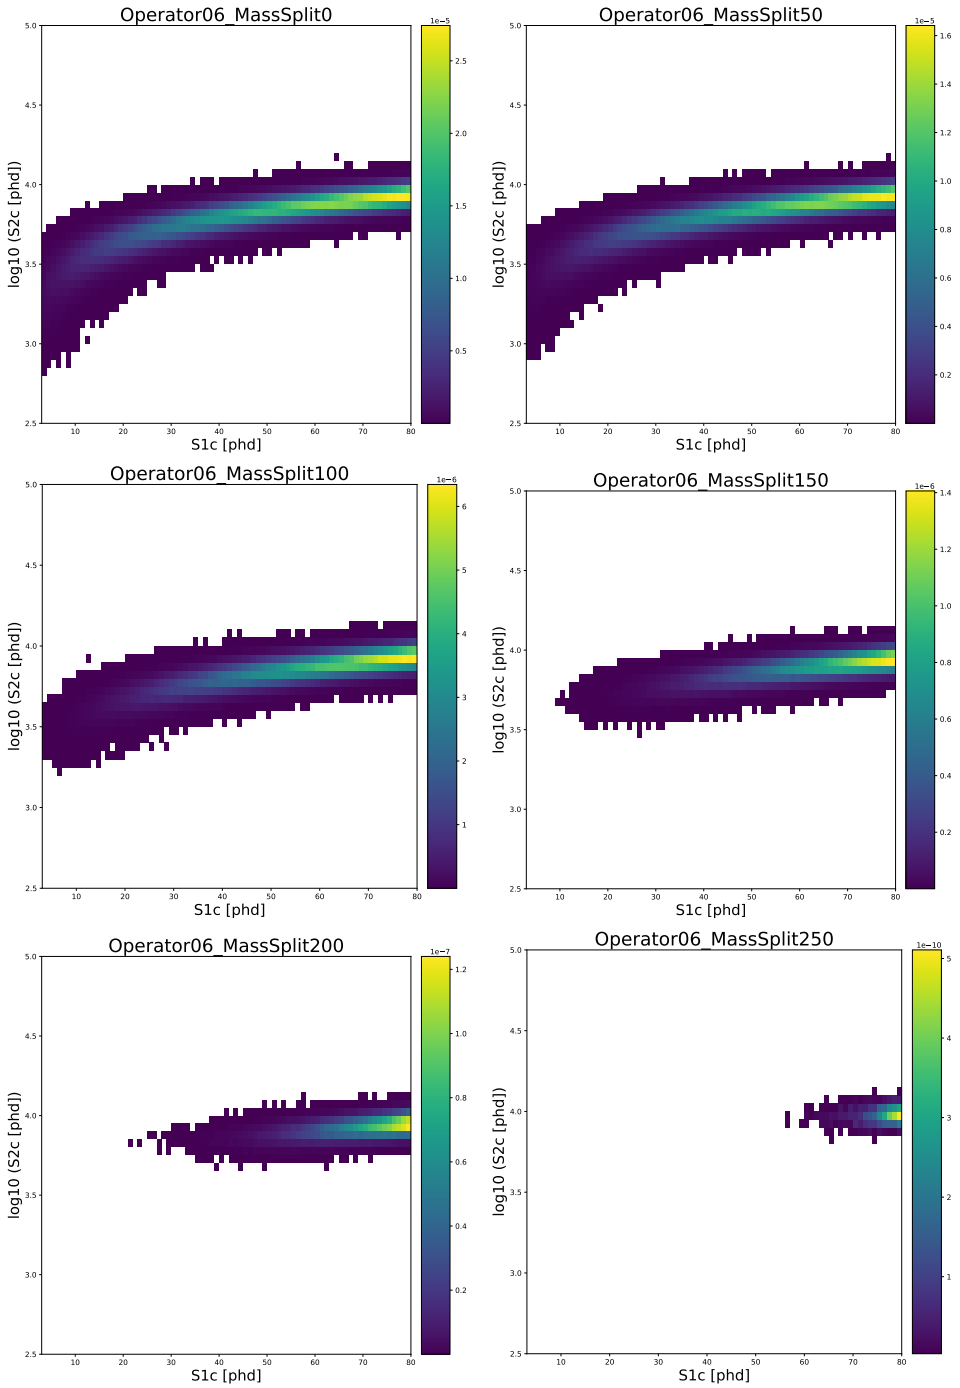

Figure 4.13: The corresponding PDF in log10(S2c) vs S1c space for Operator 6 and various mass-splitting values. All cut efficiencies and the ROI have been folded into the PDFs.

each individual component of the complete background model are simulated (as discussed in Section 4.3.1, and PDFs are generated in an analogous process to the signal PDFs in order to create respective RooWorkpaces for the 84 operator and mass-splitting value pairs. Explicitly, in these workspaces only the signal model changes, while the total background model is fixed. These workspaces are then used as inputs for LZStats in order to perform limit setting as discussed in the previous section. Top-level bash scripts were used in all of the steps to ensure consistent configuration and setting selection across all workspaces and for limit setting using LZStats. Figures 4.14 and 4.15 represent the complete event probability models generated for various mass-splitting values for operator 1 and 6, respectively. In each set of plots, the 1- and 2-σ contours for the combined background model are shown in light and dark grey. The contours for the signal are shown in purple. The NR band median and 10% and 90% quantiles are shown in red. In addition to these, 1- and 2-σ contours for two components of the background model are explicitly shown: B-8 in green and Ar-37 in orange. The observed data points that passed all the cuts are also shown as black dots. One important feature in these figures is that the signal contours, especially at higher mass-splitting values, overlap several observed data points. These data points represent an upward fluctuation of background events within the standard ROI. More specifically, these are most likely observed upward fluctuations of the detector ER background component.

When limit setting is performed using LZStats, a 90% confidence level (CL) upper limit on the number of WIMP recoil events is returned for each operator and mass-splitting pair. These are shown in Figure 4.16 for operators 1 and 6. In this figure, we can clearly see the

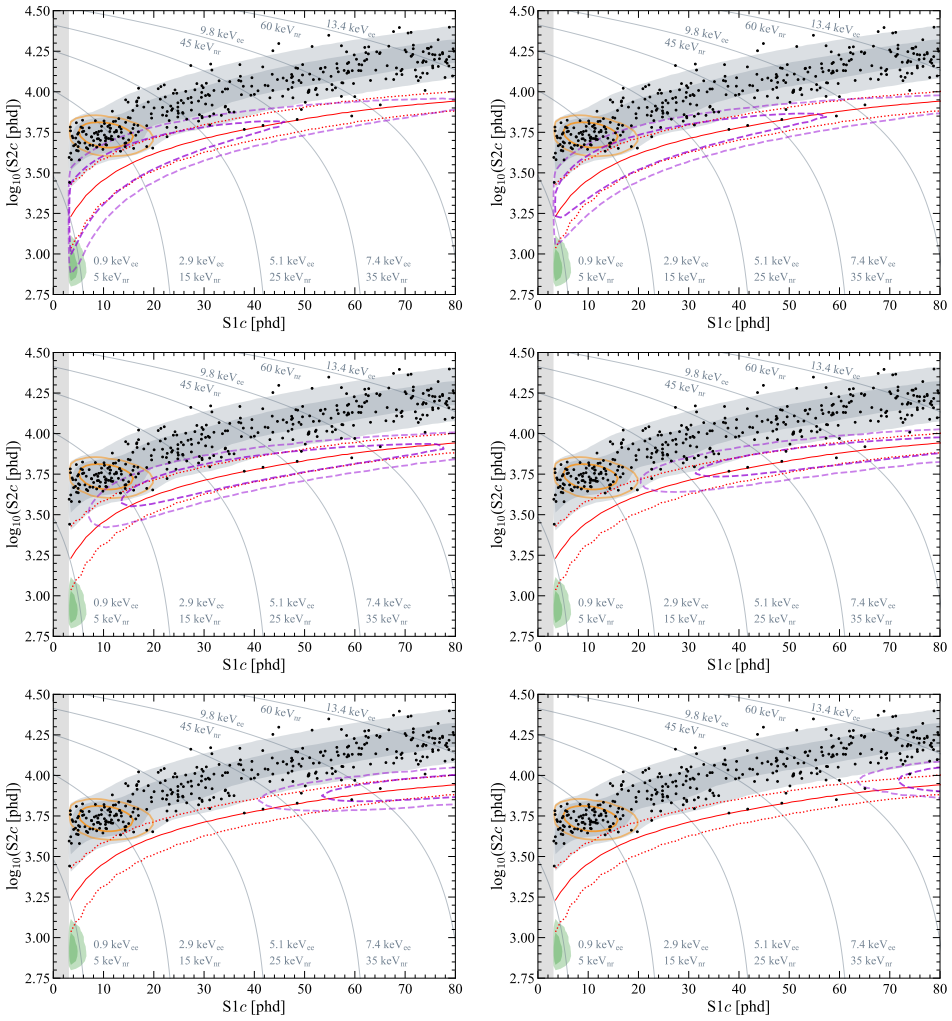

Figure 4.14: Various parts of the input RooWorkspace for LZStats are shown for Operator 1. In purple are the 1- and 2-σ contours for various mass splitting values going from 0 keV in the top left to 250 keV in the bottom right. The complete background model 1- and 2-σ contours are shown in light and dark grey. The tuned NR band median and 90% and 10% quantiles are shown in red. Also shown are the individual 1- and 2-σ contours for Ar-37 (orange) and B-8 (green).

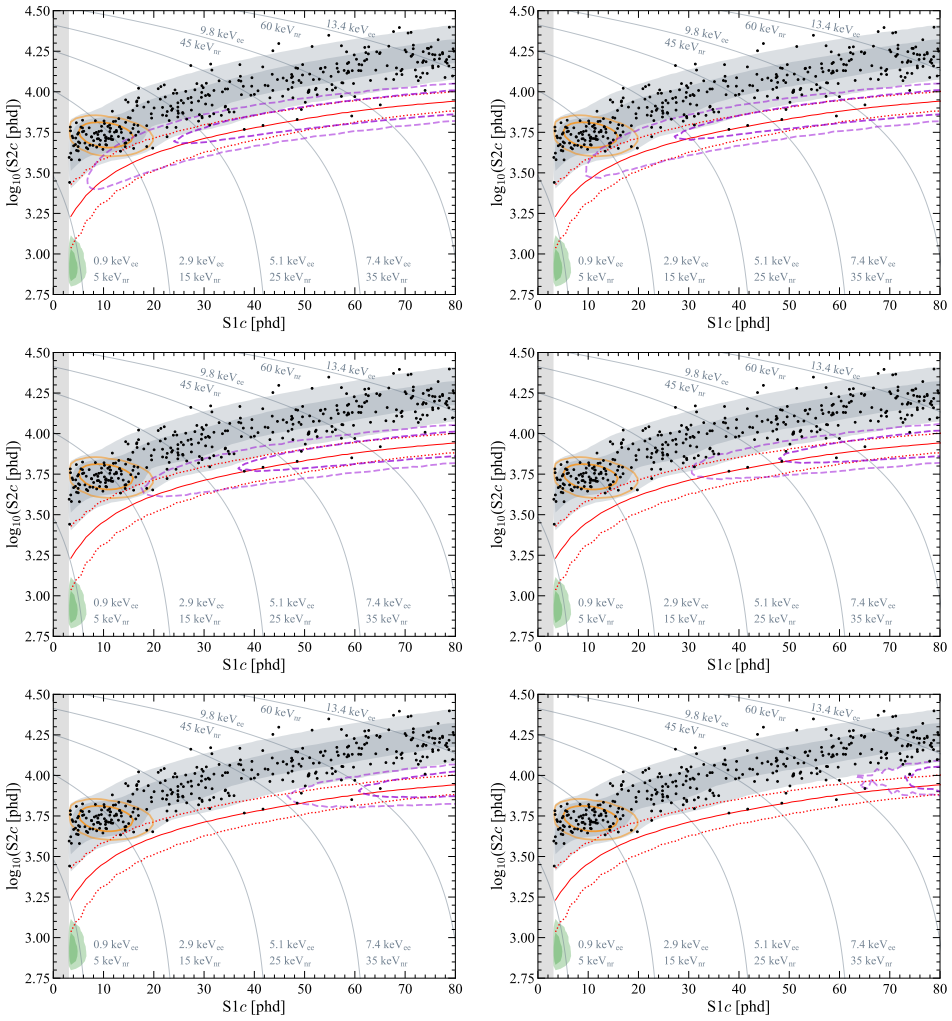

Figure 4.15: Various parts of the input RooWorkspace for LZStats are shown for Operator 6. In purple are the 1- and 2-σ contours for various mass splitting values going from 0 keV in the top left to 250 keV in the bottom right. The complete background model 1- and 2-σ contours are shown in light and dark grey. The tuned NR band median and 90% and 10% quantiles are shown in red. Also shown are the individual 1- and 2-σ contours for Ar-37 (orange) and B-8 (green).

effect of the overlap of the signal region with the upward fluctuation of background events. For operator 1, the 90% CL upper limit is highest for the mass-splitting values where there is the most overlap, while lowest for where there is minimal overlap. For operator 6, the signal region is spread across the ROI from 0 keV mass-splitting and only at the highest mass-splitting values is this not the case. This is reflected in the fact that the 90% CL upper limit is highest at the lower mass-splitting values, and lowest for the higher values. Similar correlations are observed for all operators based on the features of their differential rate spectra and the corresponding signal PDFs. The best fit value on the number of WIMP recoil events follows a similar trend. Non-zero values are returned by the PLR where the 90% CL upper limit is high. However, performing a discovery significance test for these values returns p-values corresponding to at most 2.2σ. This cross-check is consistent with the high p-value of 0.96 when fitting the background only model against the observed data as discussed previously.

In order to extract the corresponding limits on the strength of each operator, and mass-splitting pair coefficient, we need to divide the 90% CL upper limit on the parameter of interest (number of recoil events) by the expected integrated rate. These rates are calculated from the differential rate spectra (generated by assuming a value of one for the coefficient of each respective operator), which also take into account the NR acceptance as a function of recoil energy. The final limits, and the corresponding ones from XENON100 and LUX, are shown in Figures 4.17 to 4.20 for all operators and their corresponding mass-splitting terms. The green and yellow contour bands represent the 1- and 2-σ expectations. As a cross-check,

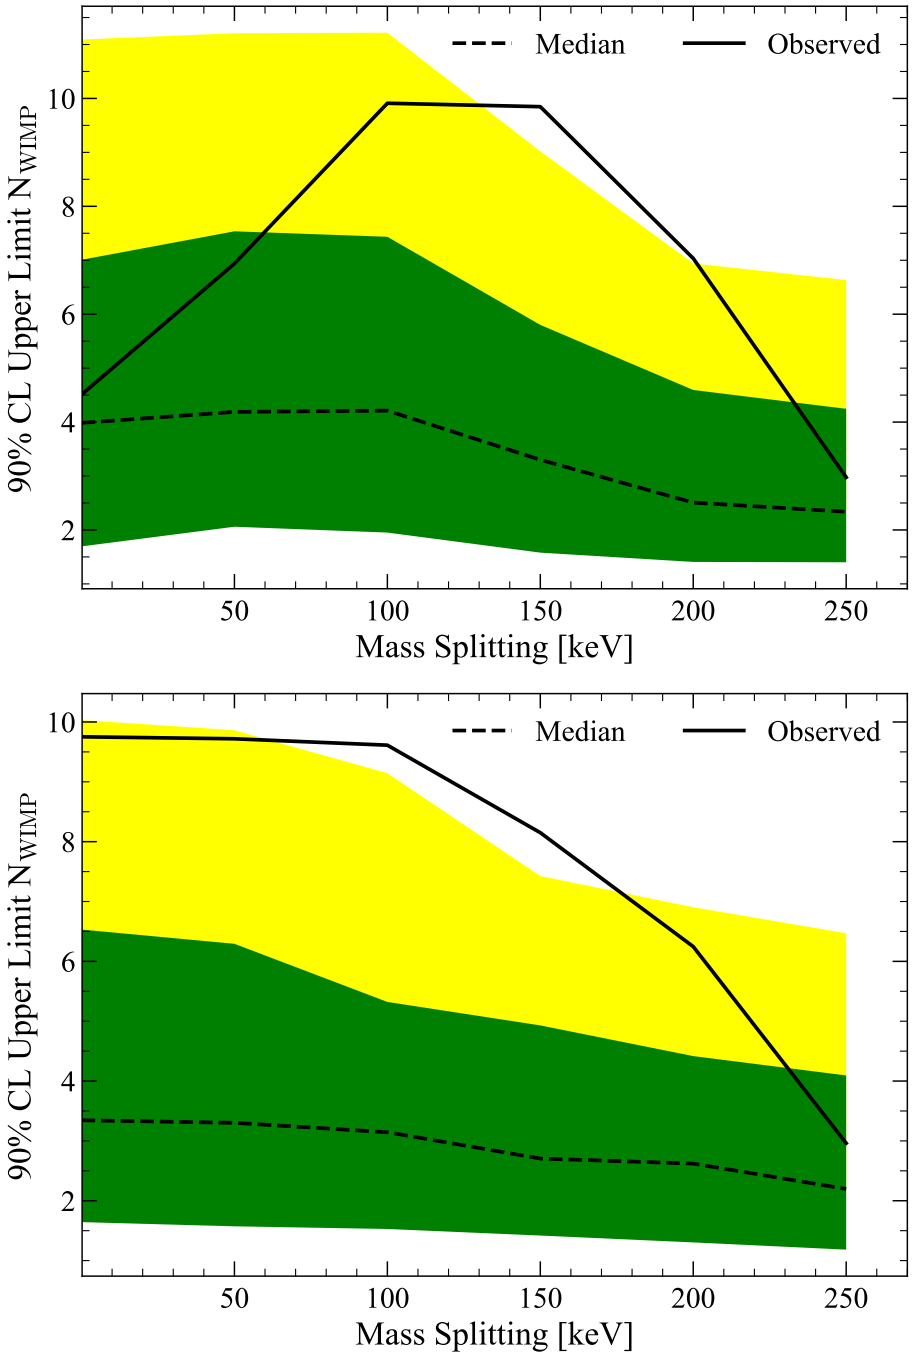

Figure 4.16: The 90% CL upper limit on the POI, μ (the number of WIMP recoil events) as a function of mass-splitting value. Operator 1 is shown in the top figure and Operator 6 is shown in the bottom figure. The dotted line shows the expectation and the green and yellow bands show +/-1 and +2 σ values per mass-splitting term.

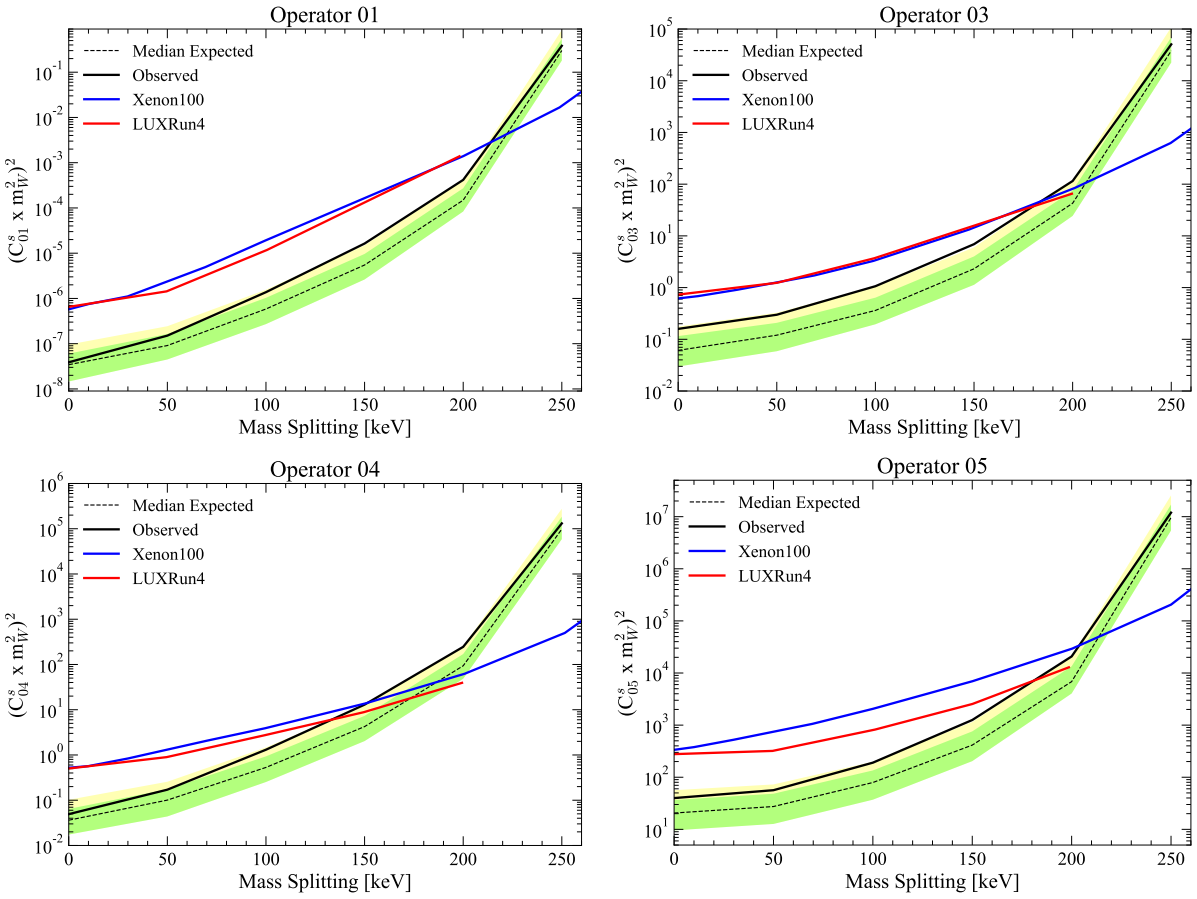

Figure 4.17: The corresponding final exclusion upper limits for the strengths of each operator coefficient is shown as a function of mass-splitting value for Operators 1, 2-5. +/-1 and +2 σ and the median expectation are also shown in green, yellow and the dotted black lines, respectively. XENON100 and LUX experimental limits are shown in blue and red, respectively.

the values for limits produced for the 0 keV mass splitting values for all operators were compared against the analogous elastic operator limits. This point of comparison provides a redundancy between the two analyses. The computed limits match for all the cases, and provide an excellent point of comparison between the two distinct analyses.

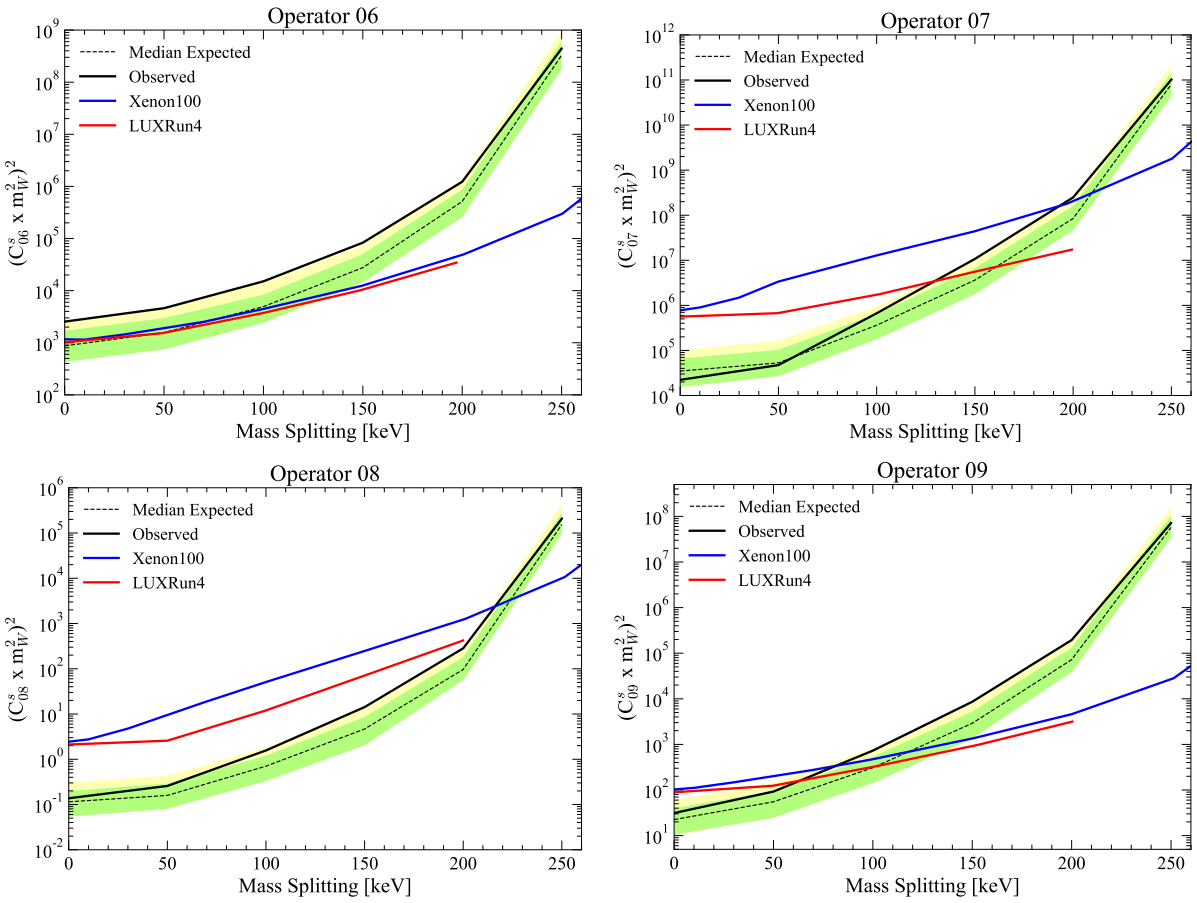

Figure 4.18: The corresponding final exclusion upper limits for the strengths of each operator coefficient is shown as a function of mass-splitting value for Operators 6-9. +/-1 and +2 σ and the median expectation are also shown in green, yellow and the dotted black lines, respectively. XENON100 and LUX experimental limits are shown in blue and red, respectively.

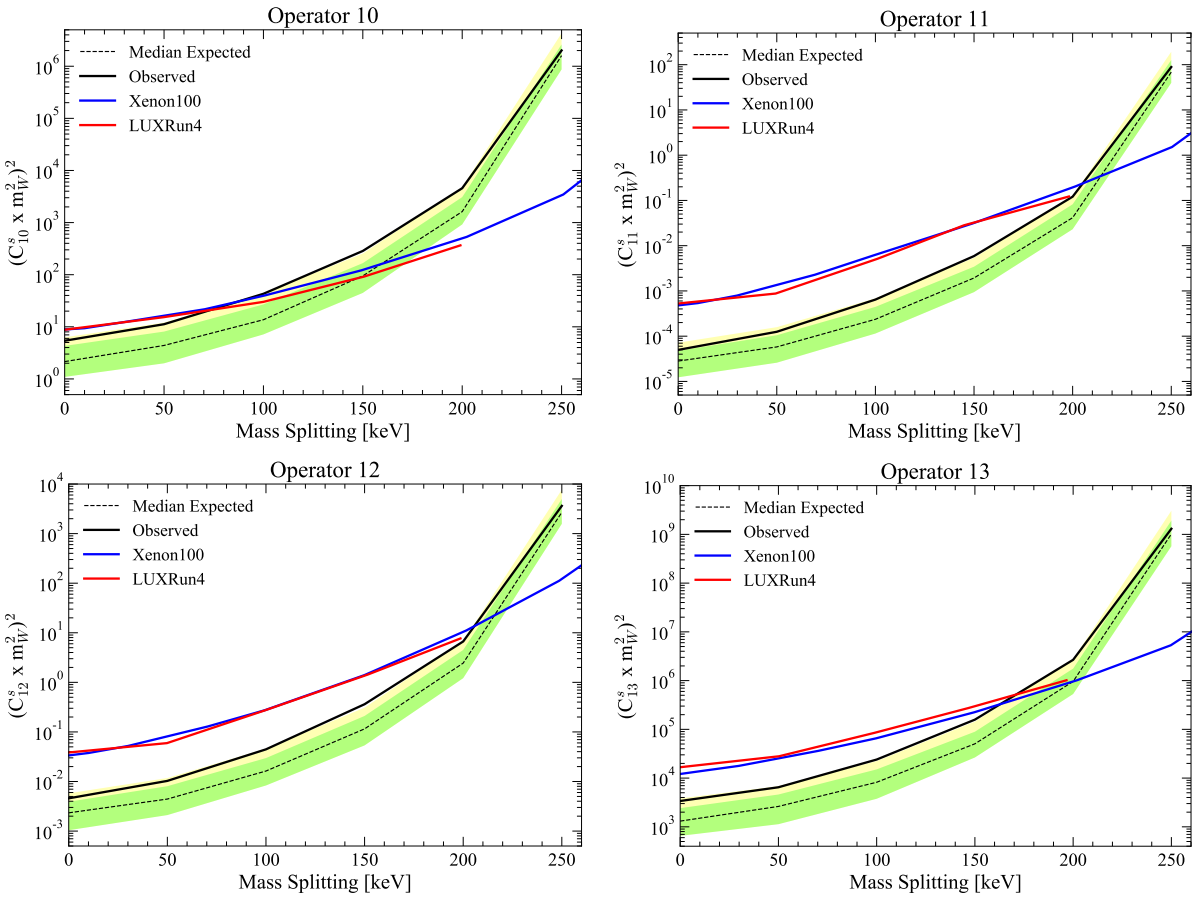

Figure 4.19: The corresponding final exclusion upper limits for the strengths of each operator coefficient is shown as a function of mass-splitting value for Operators 10-13. +/-1 and +2 σ and the median expectation are also shown in green, yellow and the dotted black lines, respectively. XENON100 and LUX experimental limits are shown in blue and red, respectively.

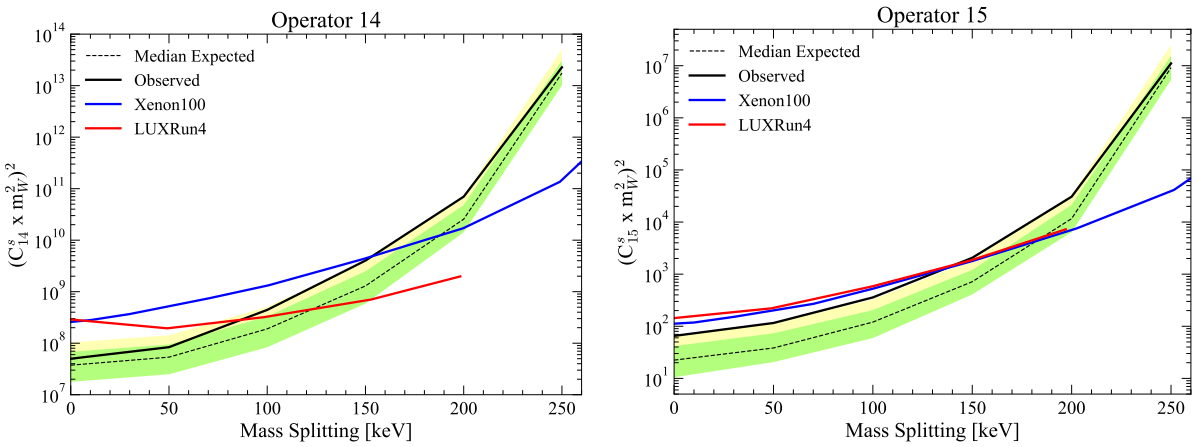

Figure 4.20: The corresponding final exclusion upper limits for the strengths of each operator coefficient is shown as a function of mass-splitting value for Operators 14-15. +/-1 and +2 σ and the median expectation are also shown in green, yellow and the dotted black lines, respectively. XENON100 and LUX experimental limits are shown in blue and red, respectively.

In conclusion, we observe improvements over the limits set by LUX and XENON100 in several regions of the parameter space. These will improve further once the extended ROI is included for the EFT analysis. We should not expect such a steep worsening of the upper limits at higher mass splitting values and in fact, we should expect a much flatter decrease where the ratio between each pair of mass-splitting values will be more in-line with the ratio between 0 and 50 keV mass-splitting values. The higher acceptances for all operators should also lower the upper limits at lower mass-splitting values. Thus, we should expect to see world leading limits across the full mass-splitting range for all inelastic operators in an extended ROI analysis.

Chapter 5

Studies in Discriminated Neutron and Gamma Detection

One of the key aspects of direct dark matter search experiments with dual-phase noble element TPC detectors is the ability to discriminate between electronic (ER) and nuclear recoil (NR) events. In Chapters 2 and 4, we discussed the basis for this discrimination by examining the distribution of events in $\log_{10}$(S2) vs. S1 space. Calibrations of the detector using appropriate radioactive sources enable the calculation of the mean and sigma of $\log_{10}$(S2) across a range of S1 values for both ER and NR events. The mean, 10% and 90% quantile curves in this 2D space define the ER and NR bands, respectively (refer back to Figure 4.1). A fundamental issue in searches for dark matter, however, is the leakage of ER events at lower S2 values into the NR band. In previous generations of dual-phase xenon TPC experiments, this was mitigated by an acceptance cut of only NR events below

the mean of the NR band. With the transition to profile likelihood ratios in calculating limits and sensitivities in the current generation of experiments, this is no longer a hard requirement as we have seen in Chapter 4. Regardless, ER leakage remains a fundamental constraint on sensitivity in this subgroup of direct detection experiments.

In argon (Ar) dual-phase TPC experiments, a crucial tool in this regard is the ability to perform pulse shape discrimination (PSD) to differentiate between nuclear and electronic recoil events [38]. This technique relies on the property that NR and ER induced pulses have measurable differences in their time profile.

For xenon, this is not a viable strategy due to a fundamentally small difference in the decay times of the triplet and singlet excited states compared to Ar, which has an O(us) difference [34]. In xenon based TPCs, especially for large detectors such as LZ, the pulse shape differences between ER and NR are O(ns), which are lost in the light collection process and also in the limited digitization rates in the DAQ.

Even though PSD cannot be utilized effectively in dual-phase xenon TPC dark matter searches, it is an extremely important capability in other applications. One such example is radiation detection in the field of nuclear security and non-proliferation. Although noble element detectors are part of the overall program, much of the emphasis is on smaller, more portable, plastic/organic scintillator based detectors for monitoring ports of entry, borders and for emergency response in the event of a nuclear device being set off. Although primarily focused on detecting and localizing nuclear materials, these detectors can also play a significant role in providing new and improved calibration setups for dark matter experiments.

For example, they can be repurposed as finer resolution neutron tagging systems (with high dynamic range) in nuclear recoil calibration studies. This can allow for improvements in uncertainties and extending the energy ranges of charge and light yield measurements at both high and low energies. [63] and [64] are examples of such calibration experiments. They currently use simple neutron tagging systems, which have a lot of scope for improvements.

In this chapter, we will look at PSD capability in plastic and organic scintillators. We will also introduce another commercially available photodetector technology commonly used as replacements for PMTs in these smaller scale detectors for radiation detection: silicon photomultipliers (SiPMs). We will then discuss a testbed that has been set up to evaluate PSD capability between various commercially available plastic/organic scintillators, within a system that maintains constant response between different scintillator measurements. This testbed will eventually be used to evaluate scintillators that are currently being developed.

Finally, we will discuss initial work in designing a portable, compact, and segmented neutron scatter camera that will optically couple plastic/organic scintillators to SiPMs. For this scatter camera, we have designed a fully custom application specific integrated circuit (ASIC) for front-end read out of SiPMs and to perform real-time PSD between fast neutrons and gammas interacting in the scintillator volume. We will cover the general neutron scatter camera requirements and the design goals and specifications for a first prototype ASIC. We will also introduce a conceptual overview of a full one channel signal chain for the ASIC. Data collected from the testbed have been used to perform a feasibility study for the real-time PSD design. Results from this study will be presented. The next chapter will dwelve

into the finer design choices for the ASIC and provide a detailed look at its various blocks. Also covered will be simulation results used to verify the design. Moving forward, in this chapter and the next, the terms ASIC and chip will be used interchangeably.

5.1 PSD in Organic/Plastic Scintillators

We now introduce PSD capable organic/plastic scintillators. The goal in this section is to briefly cover the origin of PSD capability in this class of scintillator materials. We will then focus on two commercially available scintillators selected for study with the testbed. Both are standard PSD capable organic/plastic scintillators widely used in the radiation detection community: Stilbene and EJ-276. The testbed data used for the feasibility study were taken with these materials, and they are the initial candidates for integration in a complete neutron scatter camera design.

5.1.1 Origin of PSD Capability

A detailed review of the development of organic scintillators and a discussion of the origin of PSD capability within them is given in [65]. Shorter overviews can be found in [66] and [67]. Here, we intend only to present key points.

Organic scintillators generally have prompt and delayed components in their complete scintillation spectrum. The prompt component can be ascribed to transitions from the first singlet excited states, S_1, to the ground states, S_0. The S_1 states are populated by non-

radiative transitions from higher singlet excited states directly excited by interacting particles in the scintillator medium. Similarly, triplet excited states, T_1 are produced by non-radiative transitions down from higher energy excited triplet states, which are directly populated by ionization and recombination resulting from particle interactions in the medium. Because of selection rules, direct transitions from T_1 to S_0 are inhibited. The primary mechanism for triplet state de-excitation is a process known as Triplet-Triplet Annihilation (TTA) followed by the S_1 decay to S_0. The process can be written explicitly as

$$T_1 + T_1 \longrightarrow S_1 + S_0 \tag{5.1}$$

TTA is correlated to the stopping power of an interacting particle since it is a bi-molecular process. A high density of T_1 states is required for this process to occur efficiently. The timing profile of this process is O(100ns) and thus it is identified as the delayed scintillation component of organic scintillators. The correlation between TTA and stopping power of incident particles forms the basis of pulse shape discrimination in these organic scintillators. The resulting differences in intensities of prompt and delayed scintillation light for fast neutrons compared to gammas is directly responsible for differences observed in their pulse shapes.

Historically, PSD capability of organic crystals like Stilbene [68] [69] and liquid scintillators (a modern example is EJ309) were well-established and studied by the 1970s. Plastic scintillators, on the other hand, were thought to have poor or non-existent PSD capability

[70]. It is only within the last decade that PSD capable plastic scintillators, like EJ276 [71], have been commercialized. The reason for poor PSD in plastics is subtle. Liquid scintillators are formed primarily through highly efficient fluorescent dyes dissolved in aromatic solvents [72]. Because liquid scintillators have a proven record of PSD capability, an early idea that crystal structure (as in organic scintillators) is a requirement for PSD is moot. Plastic scintillators also have similar compositions to liquid analogs with the main difference being that solvents are substituted with polymers. Thus, to first order, compositions cannot be the reason for lesser PSD capability either. Recent studies [73] have concluded that the main reason for poor PSD capability is due to excitation traps formed by low concentrations of flourescent dyes used in plastic scintillators. Increasing this concentration provides more suitable conditions for TTA to occur in plastics. Of course, as with most things in life, everything comes with a price. The same studies have also observed that increasing concentrations of the dyes reduces light yield, which directly degrades PSD capability. Thus, the process of creating plastic scintillator compounds involves a careful optimization of dye concentrations and PSD capability.

5.1.2 PSD Results for EJ276 and Stilbene

Having now introduced PSD capability in organic/plastic scintillators, we can take a look at the actual measurements of PSD capability as presented in literature for EJ276 and Stilbene. Table 5.1 summarizes some key properties for these scintillators.

Properties of Stilbene and EJ276			
Parameter	Stilbene	EJ276	Unit
Scintillation Efficiency	14000	8600	ph/1MeV e-
Max. Wavelength of Emission	380-390	425	nm
Density	1.15	1.096	g/cm^3

Table 5.1: A summary of key parameters for Stilbene and EJ276 according to manufacturer datasheets [74] and [75].

In looking at results, we have to be very careful to understand the effects of photodetector selection and readout electronics used. Some papers use PMTs and others use SiPMs; there are also a range of commercial and custom electronics and digitizer setups with varying bandwidths of front-ends, etc. As we will see, for SiPMs, these will affect the single photon shape and, thus, overall pulse shape, which is a convolution of the scintillator's cumulative photon distribution with the single photon response. In fact, this bias is partly why we decided to set up our own testbed. Doing so allows for better accounting of all systematics resulting from setup choices on final PSD capability determination. This also allows for more appropriate direct comparison between different scintillators.

However, it is also initially important to validate results from the testbed by appropriate comparison to results from literature. For that reason, we present some results on PSD capability in EJ276 and Stilbene from [76]. The paper uses a setup where PMTs are used as the photodetector and a CAEN V1761 digitizer, offering 1GHz front-end bandwidth and 4GS/s, was used to record event waveforms. Figure 5.1 demonstrates the resulting pulse shapes for neutrons and gammas using both EJ276 and Stilbene scintillators. It is important to note

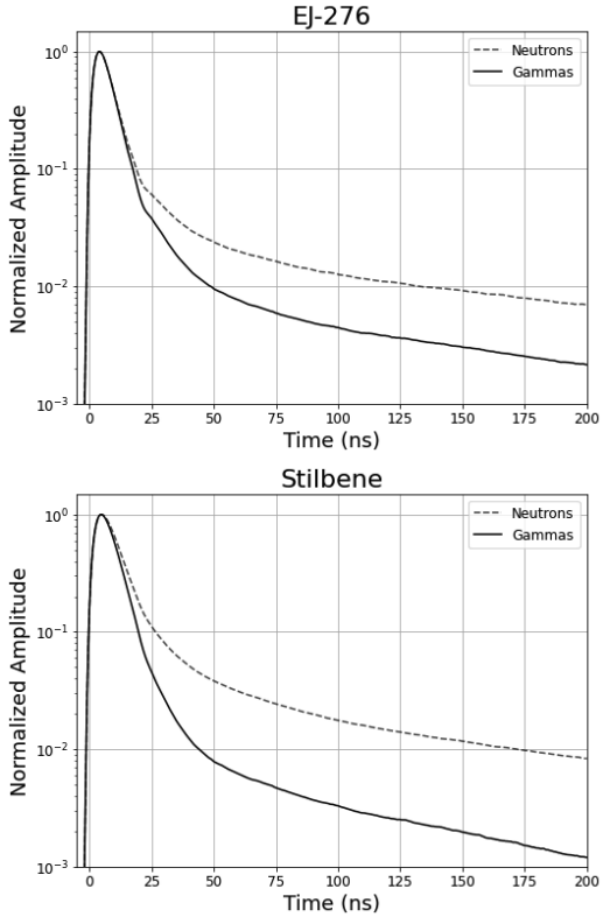

Figure 5.1: From [76], one example of previous measurements of pulse shape differences between fast neutrons and gammas for Stilbene and EJ-276. The paper, however, uses PMTs and not SiPMs for their study. Readout electronics and scintillator geometry is also expected to affect pulse shapes.

that the timescale of these pulses will be different when using a SiPM as the photodetector. Furthermore, it is clear that there should be better PSD capability between neutrons and gammas in Stilbene based on observed differences in the tails of the waveforms.

Figure 5.2 shows the separation between the neutron (top) and gamma (bottom) bands

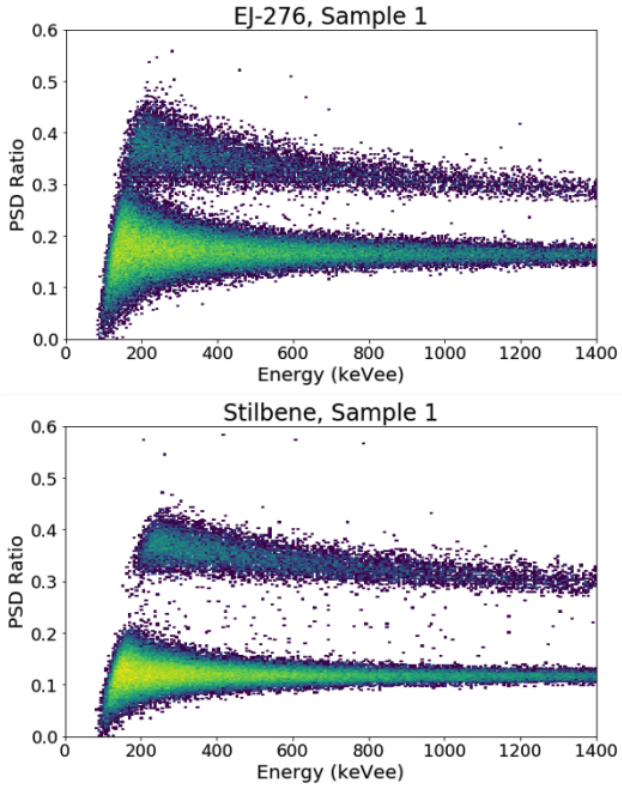

Figure 5.2: From [76], the results of a study into PSD capability of EJ-276 and Stilbene. A distribution of PSD ratio vs energy is plotted for both scintillators. Better separation for Stilbene is observed compared to EJ276. The photo-sensors used in this study were PMTs, not SiPMs.

for both scintillators with energy on the x-axis and PSD ratio on the y-axis. The PSD ratio technique involves performing two integrals on the waveform, one partial to capture the rise of the pulse and the other total to capture the total area under the pulse. The low energy cut-off point and characteristics of band shape is dependent on inherent scintillator properties, electronics readout design, and integration windows chosen for the short and long integrations of the waveforms. Regardless, the expectation that Stilbene offers better PSD

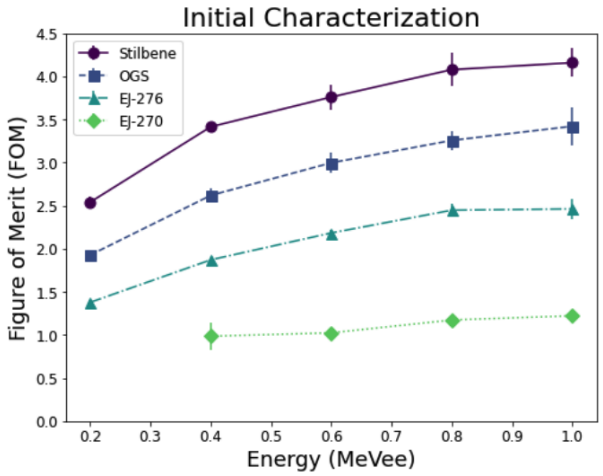

Figure 5.3: From [76], the final figure of merit (FOM) values as a function of energy from a study into PSD capability of EJ276 and Stilbene (and other scintillators). The much better performance of Stilbene across all energies compared to EJ276 is observed. The photo-sensors used were PMTs, not SiPMs.

capability is validated as the band separation is much better when compared against EJ276. In addition, these plots clearly show the effect of higher light yield on reducing widths of the neutron and gamma bands. PSD is quantified in Figure 5.3 by calculating corresponding figure of merit (FOM) values (discussed in more detail in Section 5.6; Eqn. 5.19 explicitly defines FOM) at various energies for both EJ276 and Stilbene. The higher FOM values show that Stilbene has better PSD capability across a wide range of energies.

5.2 SiPMs as Photosensors

We have mentioned that photodetector choice will have an effect on overall pulse shape for the same scintillator. Of course, read-out electronics and digitizer selection will also affect pulse

shape. In this section, we will cover the basics of design, operation, and main performance parameters of SiPMs (including information on the specific SiPM selected for our neutron scatter camera design) and provide some brief points of comparison to PMTs. The overall goal is to expand on the origin of SiPM single photon pulse shape and to understand how to deal with this in the ASIC front-end design.

The SiPM is a solid-state semiconductor alternative to PMTs, which are based on vacuum tube technology. It is formed from a densely packed array of independent single photon avalanche diodes (SPADs) connected in parallel, each with their own quenching resistor. These SPAD sensors are commonly referred to as microcells of the SiPM. Each is essentially a reverse biased p-n junction. The basic operating principle of photon detection in SPADs is that when photons are absorbed within the effective depletion region, they will create electron-hole pairs. The applied reverse bias voltage to the microcells sets up an electric field that accelerates holes towards the anode and electrons to the cathode of the p-n junction. If the reverse bias voltage is high enough (at the edge or slightly past breakdown), the electric field set up will be strong enough that these charge carriers will be accelerated to the point they have enough kinetic energy to generate secondary charge carriers and so on (Geiger avalanche operating mode) [77].

As a result of this basis for photon to electron/current conversion, SiPM gains are comparable to PMTs. However, unlike PMTs, they operate at low voltages, provide insensitivity to magnetic fields, mechanical robustness and excellent uniformity in single photon response [77]. Like PMTs, they provide fast timing and single photon detection capabilities. Silicon

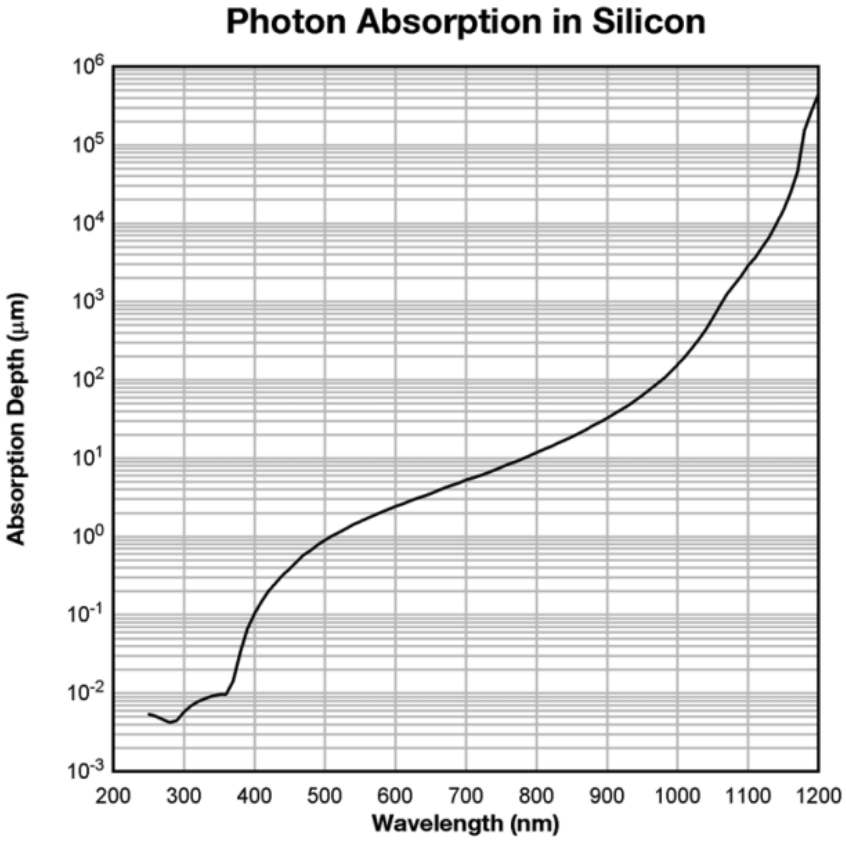

Figure 5.4: From [77], the absorption length for various wavelengths of light in silicon.

can absorb a range of wavelengths within O(10μm) thickness as shown in Figure 5.4. The depths at which it absorbs various wavelengths of light results in a wavelength dependence of photon detection efficiency (PDE) for SiPMs. Thus, an important consideration in optically coupling scintillators and SiPMs is ensuring proper and sufficient overlap of the wavelengths of light produced in a scintillator with the SiPM's photodetection efficiency vs. wavelength curve.

Now, a SiPM wouldn't be a very good photosensor if there is no way to turn off the

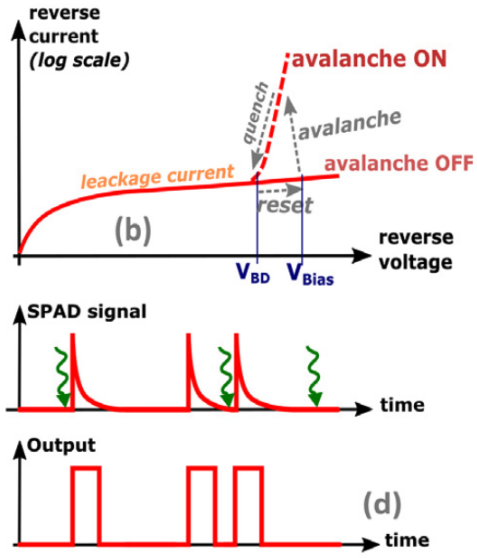

Figure 5.5: The avalanche and reset process for each SPAD that is connected in parallel to form a complete SiPM cell, extracted from [78].

avalanche. This is achieved with the aid of the quenching resistor in series with every microcell. Figure 5.5 shows the reset process visually. In words, as the avalanche is initiated, it causes a substantial increase in current flow across the quenching resistor. This causes a drop in the effective bias voltage to the microcell. When this drop is large enough to bring the effective bias voltage seen by the microcell below breakdown voltage, the avalanche stops. Then, all that is needed is for the effective bias voltage seen at the microcell to return to the nominal applied value. This occurs via charging up the capacitance of the junction again through the quenching resistor. This effective RC charging process is the origin of the characteristic long decay tail in the second row of Figure 5.5. One very important thing to note is that as a result of the SPAD microcell structure of SiPMs, they are only semi-analog

sensors. Multiple coincident photons within the same microcell will result only in a single photon equivalent output. Of course, due to the very small area of a microcell, we would need a very large photon flux in order to expect a non-negligible probability of this occurring.

5.2.1 SiPM Parameters

Having now introduced the SiPM, we move on to the main parameters that define it. These are: **pulse shape, fill factor, breakdown and overvoltage, gain, PDE, dark count rate, optical cross-talk, afterpulsing, dynamic range/linearity and temperature dependence**. The following subsections will address each of these individually.

5.2.1.1 Pulse Shape

The rise time of the current pulse output from a SiPM depends to the first order on the formation time of the avalanche within the microcell. To second order, this will be influenced by transit times of electrons across the sensor area. The decay time of the pulse is the recovery time of the sensor after avalanche is quenched. This recharge time constant can be given by

$$\tau_{RC} = C_d \left(R_q + R_s * N\right) \tag{5.2}$$

with C_d the microcell's effective capacitance, R_q the quench resistor value, N the number of microcells and R_s the sum of all resistance in series with the sensor. From this relationship, it is important to note that microcell size affects C_d and thus the decay time of the output pulse. Similarly, any series resistance will likewise change the decay time of the output pulse

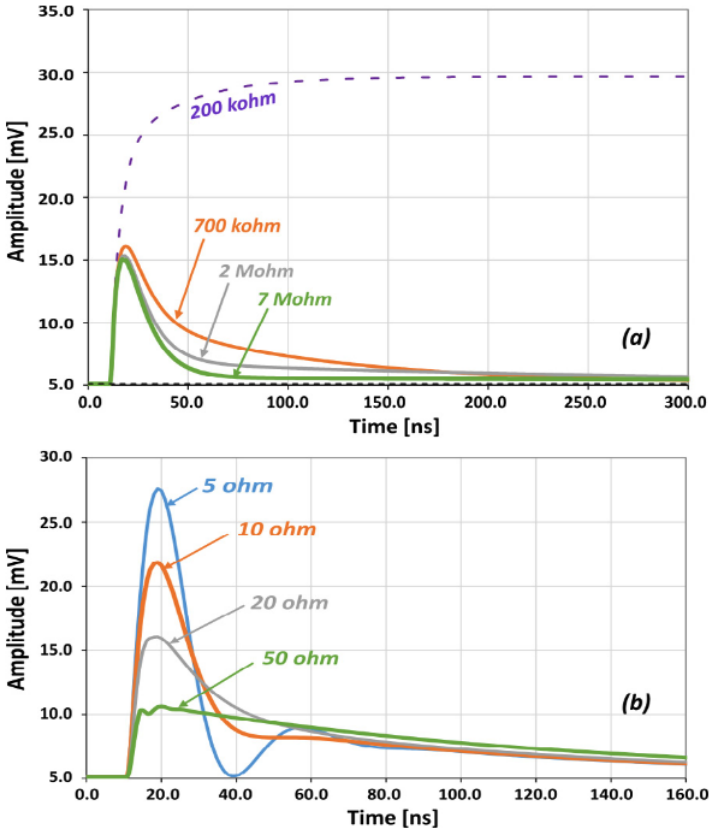

Figure 5.6: Above, the effect of quenching resistor values on overall SiPM output pulse shape. Below, the effect of different front-end input impedances on SiPM pulse shape. Both plots are from [78].

shape. Figure 5.6 shows this effect explicitly. PMTs do not have long decay tails nor do they have a large effective capacitance. As a result, generally speaking, the single photon pulse width in a PMT is O(ns) and it is not quite sensitive to shaping effects from series resistance.

In general, the bandwidth of front-end electronics used for SiPM readout will have shaping (i.e. filtering) effects as well. This is a direct consequence of limited bandwidth amplifiers

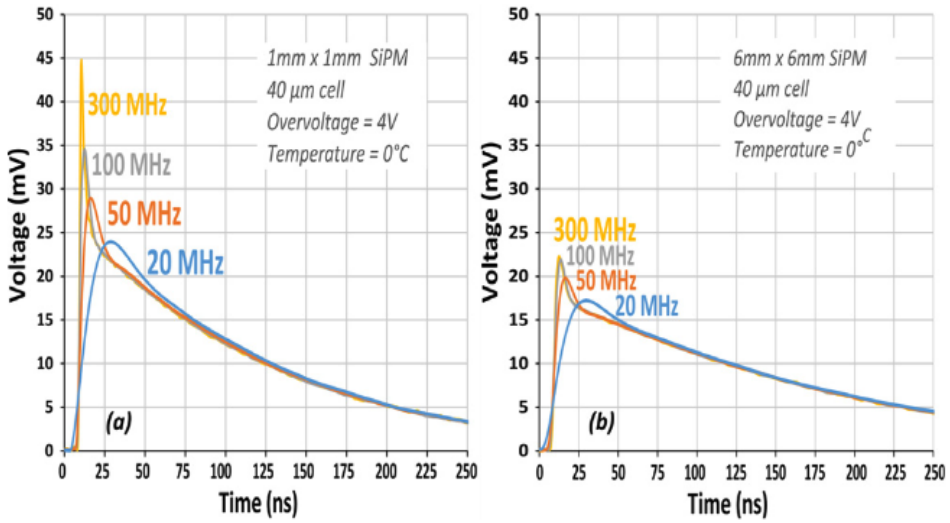

Figure 5.7: From [78], the effect of front-end bandwidth (and SiPM sensor size) on the overall SiPM output pulse shape. Plot (a) describes the effect of front-end bandwidth for a 1mmx1mm SiPM with 40um microcell size and and plot (b) describes the same but for a 6mmx6mm SiPM with the same microcell size.

being unable to resolve the sharp rising edge of the SiPM pulse. This effect is shown in Figure 5.7. Both of these characteristics heavily informs the design of the front-end readout amplifier design on the ASIC.

5.2.1.2 Fill Factor

Fill factor is defined as the percentage of the SiPM surface area that can be considered active (sensitive to light). SiPMs, being composed of microcells, requires both optical and electrical isolation from neighboring microcells. Furthermore, signal tracks and the integrated quenching resistor also contribute to minimum spacing requirements between microcells. SiPMs composed of larger microcell sizes will have higher fill factors because the minimum

spacing requirements do not depend on microcell size. This allows for higher photon detection efficiency but it comes at the cost of higher junction capacitances and thus longer decay (recharge) times.

5.2.1.3 Breakdown Voltage and Overvoltage

Breakdown voltage is defined as the minimum reverse bias voltage required to generate a sufficiently strong electric field for Geiger avalanche mode operation. The overvoltage is the voltage applied in excess of the breakdown voltage. Minimum and maximum overvoltages are listed in SiPM datasheets and are individual to specific SiPMs based on SPAD design by the manufacturer.

5.2.1.4 Gain

SiPM gain is defined as the number of electrons generated per detected/absorbed photon. This parameter is a function of overvoltage and microcell size. The exact definition is given by

$$G = \frac{C_d * \Delta V}{e} \tag{5.3}$$

with C_d the microcell capacitance (increases with larger microcell sizes), ΔV the overvoltage and e the electron charge.

5.2.1.5 Photon Detection Efficiency

PDE is a function of overvoltage, fill factor and wavelength of incident photon. It cannot be treated simply the quantum efficiency of silicon due to the photoelectric effect. The exact definition is given by

$$PDE\left(\lambda, \Delta V\right) = \eta\left(\lambda\right) * \epsilon\left(\Delta V\right) * F \tag{5.4}$$

with λ the wavelength of the incident photon, ΔV the overvoltage, η the quantum efficiency of silicon, ϵ the avalanche initiation probability and F the fill factor. The quantum efficiency, η, is defined as the wavelength dependent probability of a photon being absorbed in the sensitive volume of silicon in the SiPM and generating an electron-hole pair (see Figure 5.4). The avalanche initiation probability, ϵ, is the probability of the generated electron-hole pair going on to initiate an avalanche. This is dependent on the survival probability of an electron to enter the high-field avalanche region.

5.2.1.6 Dark Count Rate

By far, the main sources of noise within a SiPM are dark counts. Dark counts are the result of thermal excitation of electrons into the conduction band, which then initiate an avalanche. It is a function of overvoltage, temperature and active area (fill factor) of the SiPM. Dark count rate increases linearly with active area and non-linearly with temperature. Because of the SPAD substructure of SiPMs, the signal formed in response to these thermally generated

electron-hole pairs and photon-absorption generated pairs are identical. Thus, they create a very high rate of single photon equivalent noise. At room temperature, because of the high rates of dark counts, pileup of dark counts is also a concern and must be accounted for appropriately in trigger threshold selection. Figure 5.9 shows examples of dark count pile-up pulses as a function of time.

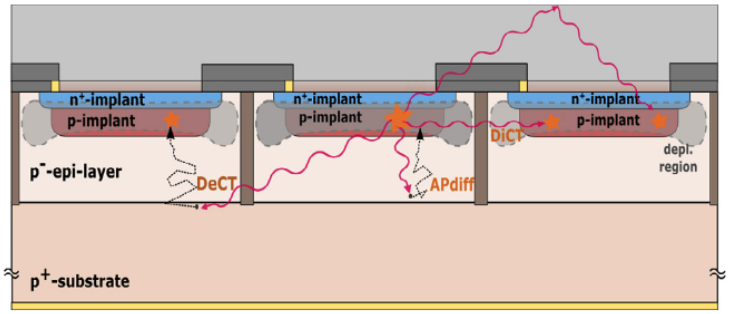

Figure 5.8: From [78], a diagram showing how optical cross-talk and afterpulsing within a SiPM's microcells occur.

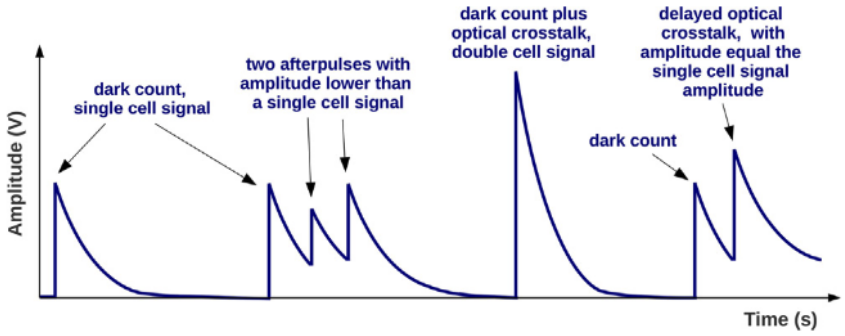

Figure 5.9: From [78], SiPM output pulses from optical cross-talk and afterpulsing.

5.2.1.7 Optical Cross-talk

This phenomenon is an additional source of background noise within a SiPM. Cross-talk is a result of accelerated charge carriers emitting near-infrared (NIR) photons during avalanche. These photons have long photon absorption lengths (refer back to Figure 5.4) in silicon and they can travel to neighboring microcells and initiate a secondary avalanche. Because this occurs near instantaneously, this results in multiple photon equivalent pulses being produced for a single dark count/true photon initiated avalanche. The exact definition for cross-talk is the probability that a microcell undergoing avalanche will cause avalanche in another microcell. The main parameters that affect the probability of its occurrence are overvoltage and fill factor. The probability increases as a function of both. Figure 5.8 visually describes this process. The red lines show possible trajectories of NIR photons that go on to initiate avalanche in other microcells. Figure 5.9 shows examples of optical cross-talk induced pulses.

5.2.1.8 Afterpulsing

In PMTs, afterpulsing occurs when an ion transits back to the photocathode and creates a new photoelectron. In SiPMs, it occurs due to charge carriers getting trapped in defects (in this case within the SPAD p-n junction). These trapped carriers are released after some time, usually on the O(ns) for SensL SiPMs [77] but generally can occur as much as O(us) later. This can potentially initiate avalanche again within the microcell. If the time difference is closer to ns scale, the microcell will not be fully charged back to nominal bias voltage after

the initial avalanche is stopped. In this case, afterpulsing will result in pulses like shown in Figure 5.9. In the case of us time scales, this is another source of single photon background noise. The probability for this phenomenon to occur increases with overvoltage due to the increase in avalanche initiation probability.

5.2.1.9 Dynamic Range and/or Linearity

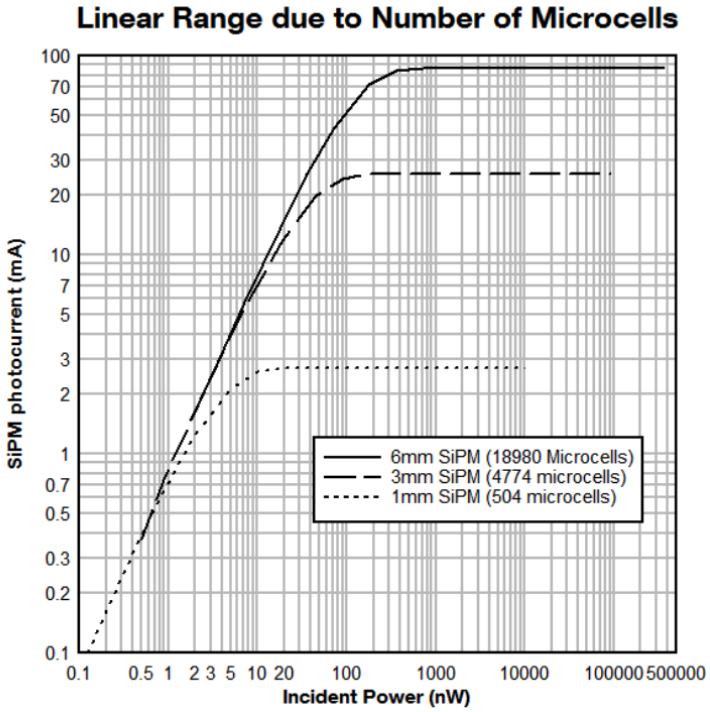

Figure 5.10: From [77], the relationship between microcell size and linear response range of a SiPM.

Linearity can be tied back to an earlier comment we made about the semi-analog nature of SiPMs. At high incident photon fluxes, the response of SiPMs are not expected to be

linear. This can be quantified in the following formula that gives the expected number of microcells fired

$$\mathrm{N_{fired}}\left(\mathrm{N}, \Delta\mathrm{V}, \lambda\right) = \mathrm{N}\left(1 - \exp\left(-\frac{\mathrm{PDE}(\Delta\mathrm{V}, \lambda) \cdot \mathrm{N_{ph}}}{\mathrm{N}}\right)\right) \tag{5.5}$$

In this equation, N is the number of microcells, ΔV the overvoltage, λ the wavelength of incident photons, PDE is defined in Eqn. 5.4, N_{ph} is the number of incident photons. Figure 5.10 plots this behavior for a few different microcell sizes but converts number of incident photons to power and number microcells fired to an equivalent SiPM photocurrent. As expected, at high enough incident power/photon flux, saturation of the SiPM output will occur.

5.2.1.10 Temperature Dependence

There are two primary avenues through which temperature affects SiPM performance. The first is the non-linear increase in dark count rate as a function of temperature (see Figure 5.11). Succinctly put, every 10°C increase in temperature results in a 50% increase in the dark count rate. More specifically, dark counts (and the equivalent dark current) can be ascribed to three sources: surface leakage current, thermal current and tunnelling current [79]. The tunnelling current contribution increases as a function of applied overvoltage but is not dependent on temperature. The effect of the tunnelling current contribution is that the dark count rate is expected to be higher at increasing overvoltage while maintaining the same temperature. However, it is the thermal current contribution that is directly dependent

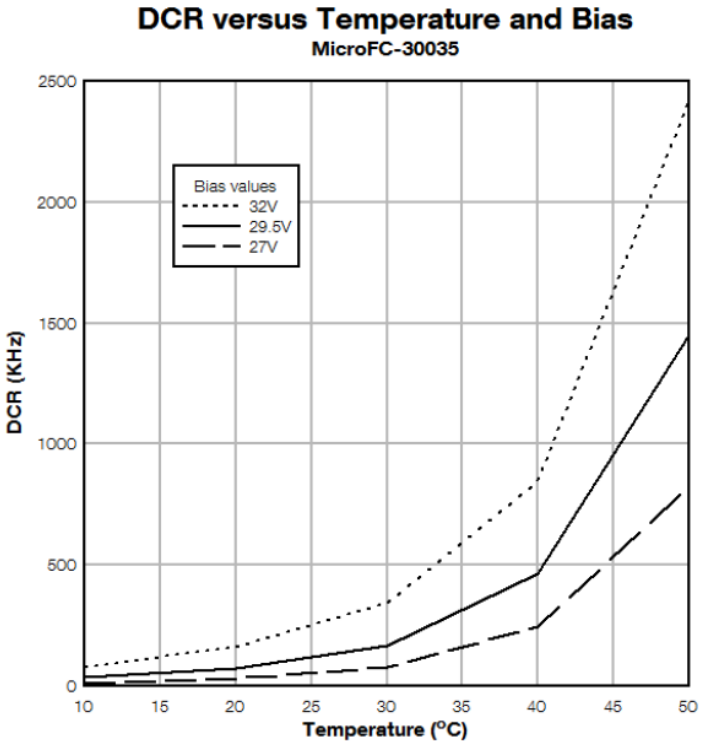

Figure 5.11: From [77], the relationship between temperature and dark count rate in a SiPM.

on temperature and can be modelled through a modified form of the Richardson-Dushman equation:

$$I_{\text{thermal-current}} = AT^2 \exp\left(-E/k_B T\right), \tag{5.6}$$

where A is Richardson's constant, which can be multiplied by a corrective factor that accounts for material type and the conversion from current density to current, T is the temperature in K, E is the energy bandgap, and k_B is the Boltzmann constant.

The second avenue is the linear dependence of breakdown voltage on temperature. While the coefficient is usually small, (21.5mV/°C for SensL SiPMs), if large temperature fluctuations are expected, either thermal stability mechanisms or bias compensation circuits should be included in the overall system design.

5.2.1.11 A Specific SiPM: The SensL J-series

Now that we have introduced SiPMs in general, we can transition to discussing the specific commercially available SiPMs selected for this prototype compact, segmented neutron scatter camera work. We use SensL (now OnSemi) J-series 8x8 SiPM arrays where each pixel in the array is a 6mmx6mm SiPM that consists of 22,292 35um microcells. The overall fill factor for this SiPM is 75%. A summary of values for these and other parameters introduced in the preceding subsections for the SensL J-series SiPMs are given in Table 5.2. In addition, Figure 5.12 shows the specific PDE vs. λ relationship for two different overvoltages for these SiPMs.

An important thing to note is that these SiPM arrays contain two coupled outputs per SiPM pixel. The standard current pulse output, as discussed, is referred to as SOUT. The other output is a fast, capacitively coupled output (FOUT), which is a derivative of the internal fast switching in response to avalanche turn on (photon detection/thermal excitation). Figure 5.13 shows where this output is extracted for each microcell. Because of the ultra-fast characteristics (refer back to Table 5.2) of this pulse compared to the SOUT pulse, this can be, and is, used for timing purposes on the ASIC. Like the standard SOUT output,

SensL J-Series 6mmx6mm (35um microcells) SiPMs			
Parameter	Overvoltage	Value	Unit
No. Microcells	N/A	22292	N/A
Fill Factor	N/A	75	%
PDE (peak λ)	2.5V/6.0V	38/50	%
Dark Count Rate	2.5V/6.0V	50/150	kHz/mm^2
Gain	2.5V/6.0V	2.9/6.3	x10^6
Rise Time (SOUT)	2.5V/6.0V	180/250	ps
Recharge Time Constant	N/A	50	ns
Capacitance (SOUT)	N/A	4140	pF
Capacitance (FOUT)	N/A	160	pF
FOUT FWHM	N/A	3	ns
Cross-talk	2.5V/6.0V	8/25	%
Afterpulsing	2.5V/6.0V	0.75/5.0	%

Table 5.2: This table summarizes performance parameters for the SensL J-series 6mmx6mm (35um microcell) SiPMs according to [80]

the FOUT output is also summed in parallel across all microcells. Thus, besides timing information, the pulse pulse area also provides a quantity proportional to the number of incident photons. It also has a much lower effective capacitance compared to SOUT.

5.3 Studying PSD Capability of Scintillators Coupled to SiPMs

Having covered both PSD capability in organic/plastic scintillators and a detailed description of SiPM functioning, we can now present the testbed that performs characterization of both SiPMs and scintillators. Figure 5.14 provides a system level block diagram of the setup. The

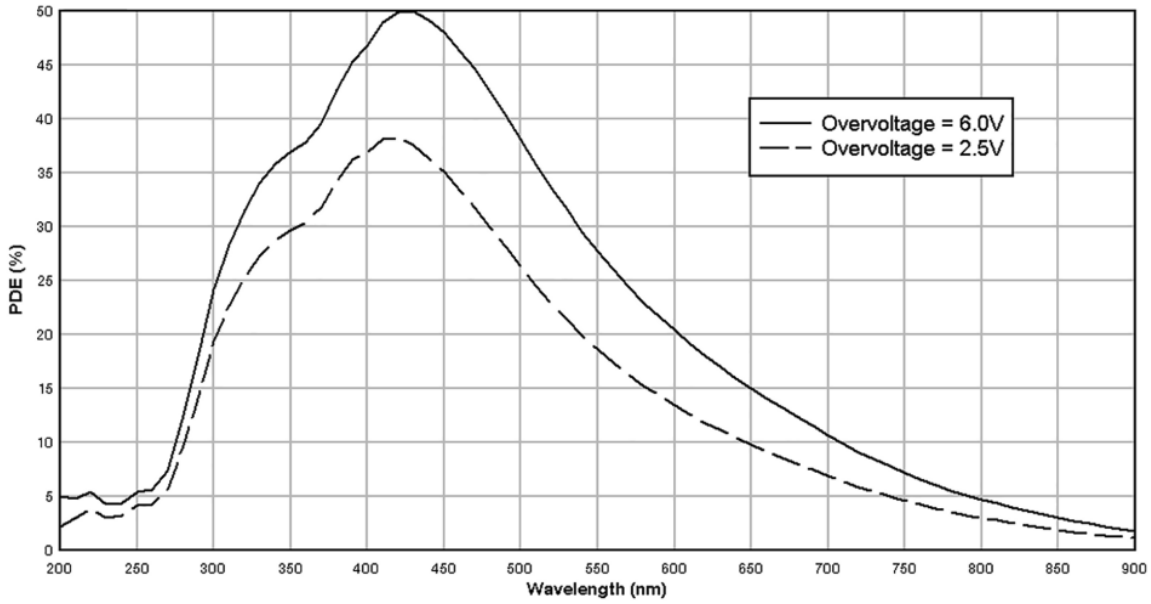

Figure 5.12: From [80], the PDE vs. wavelength relationship for SensL J-series 6mmx6mm (35um microcells) SiPMs at two different overvoltages.

cornerstone of the testbed is a custom dark box design. Pictures of the dark box are shown in Figure 5.15. Metal clasps securely clamp down on detachable top and front side lids. The front side lid, when detached, shows a light maze trap that offsets where power and signal cables are fed into the front side lid and where they enter the dark box. This maze was included into the design to mitigate light leakage into the box.

Continuing with custom parts of the testbed, inside the box, a custom designed fanout printed circuit board (PCB) is placed. Figure 5.17 shows a close up of both sides of the fanout board. It has been designed so that an 8x8 SensL J-series SiPM array can be plugged directly onto it. The board then picks off 18 selected SiPMs from the full array and brings out both the SOUT and FOUT outputs for each SiPM to two respective SMB connectors.

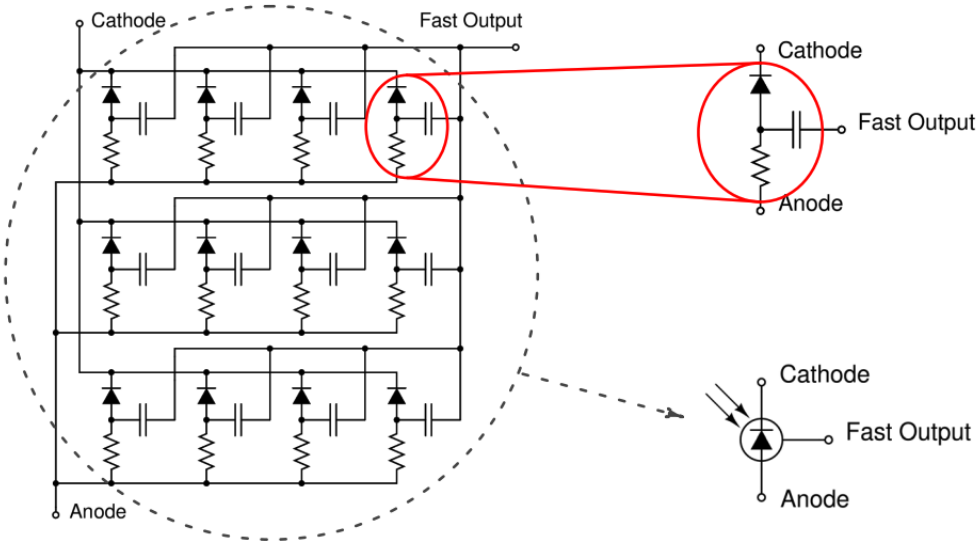

Figure 5.13: A diagram of the substructure of a SensL SiPM, from [80].

A 50Ω resistor is included in parallel to ground with an SMB connector for each SOUT line in order to provide current to voltage conversion and 50x transimpedance gain. Specifically selecting 50Ω resistors was done to provide impedance matching and prevent line reflections as both the SMB connectors and coax cables have 50 Ω characteristic impedance. The downside to this choice, as discussed in Section 5.2.1.1, is a longer decay constant for the SOUT single photon response. Because this is convolved with the full photon distribution generated in response to an interacting particle within the scintillator volume, the overall pulse shape for both fast neutrons and gammas is expected to be affected as well.

The two parallel black connectors in the top picture of Figure 5.17 are mating connectors for the SiPM array's connectors. A top-down view of the SiPM array plugged into the fanout board is shown in the bottom picture of Figure 5.16. We can optically couple a 6mm

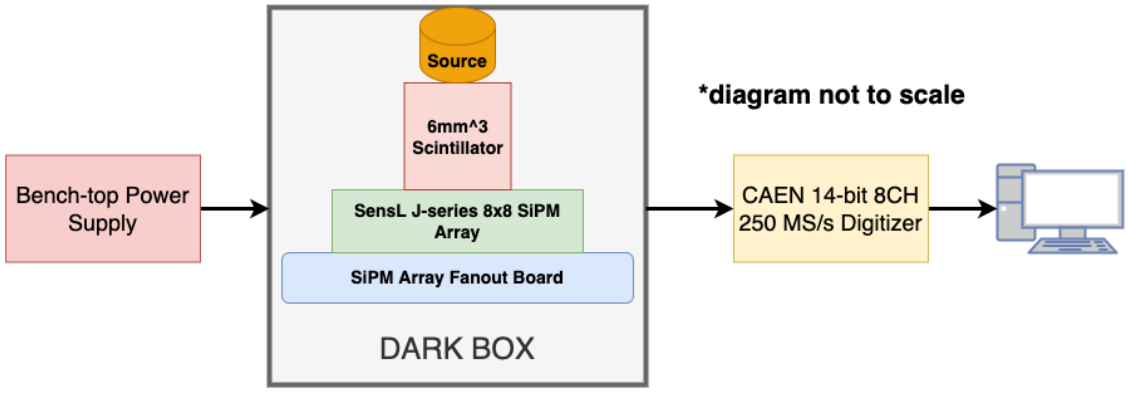

Figure 5.14: A conceptual overview of the SiPM + Scintillator testbed setup. A focus was made to keep the testbed as simple as possible in order to be able to study inherent PSD capabilities of various scintillators coupled to SiPMs without effects of readout electronics.

cube block of either EJ276 or Stilbene (and generally speaking, any scintillator material) to a single 6mmx6mm SiPM on the full array. In the figure, a cube of EJ276 is optically coupled to a SiPM on the full array. Sealed radioactive sources can be placed on top of the cube of scintillator for PSD studies. We can directly connect the FOUT and SOUT outputs from the fanout board to two channels of a CAEN DT5725 digitzer using an SMB $\rightarrow$ MCX coax cable. This digitizer has 14-bit resolution, 8 channels total, and 250 MS/s sampling rate. Two of the eight channels are used to readout out the SOUT and FOUT outputs. A standard bench-top power supply is used to deliver the bias voltage to the fanout board (through the green Phoenix connector), which then routes it to the plugged in SiPM array. The entire setup is relatively simple and kept so on purpose. That's because a primary goal of the testbed is to perform PSD studies using various scintillators coupled to SiPMs without needing to de-convolve the electronics response of a complicated font-end acquisition system

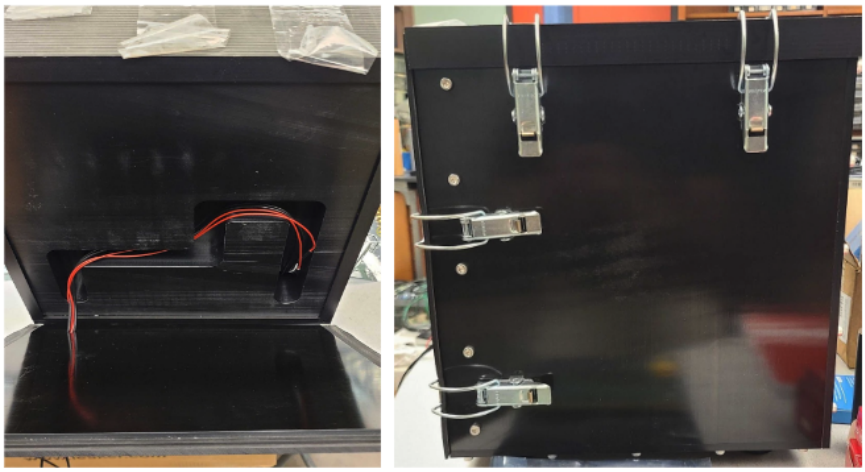

Figure 5.15: Two photographs of the dark box design. A light maze was included as a way to offset the cables coming into the dark box and prevent light leakage. Clasps are included on the sides of the box as an additional precaution against light leakage. They clamp down on the two removable lids of the dark box.

on the pulse shape. This custom testbed can accomplish that goal with a very straight forward and simple to understand signal chain.

Sample FOUT and SOUT waveforms (for gammas from a Cd-109 sealed source) acquired and digitized using the testbed are shown in Figure 5.18. In the plot, each sample corresponds to 4ns. The waveforms show the resulting pulse characteristics when Stilbene is optically coupled to a SensL 6mmx6mm J-series SiPM. As discussed, the SOUT pulse shape is influenced by the 50Ω resistor. On top of that, both SOUT and FOUT pulse shapes are shaped by the 125MHz anti-aliasing filter at the inputs of each channel of the digitzer. This results in a much greater shaping of the FOUT waveform, and it is expected that the SOUT rise time is similarly affected. This plot also shows the minimal noise observed in the

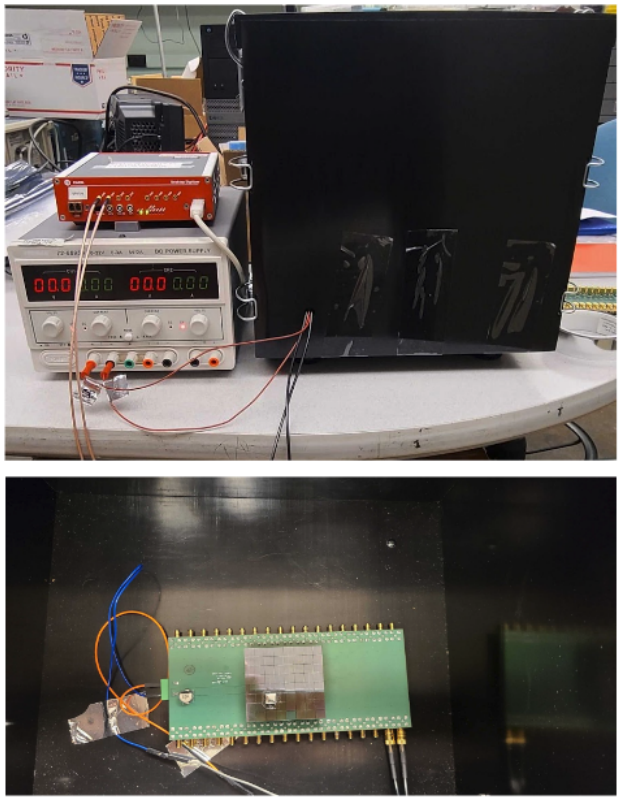

Figure 5.16: Photographs showing all the major components of the SiPM + Scintillator testbed.

testbed. Measured noise of the testbed is approximately 180uV rms. Various baseline noise measurements taken at different bias voltages, with and without various configurations of the fanout board connected to the digitzer inputs, allowed us to determine that the majority noise contributor of the testbed is actually the digitizer itself. This presents an ideal scenario in order to characterize inherent PSD capability of scintillators for small energy depositions.

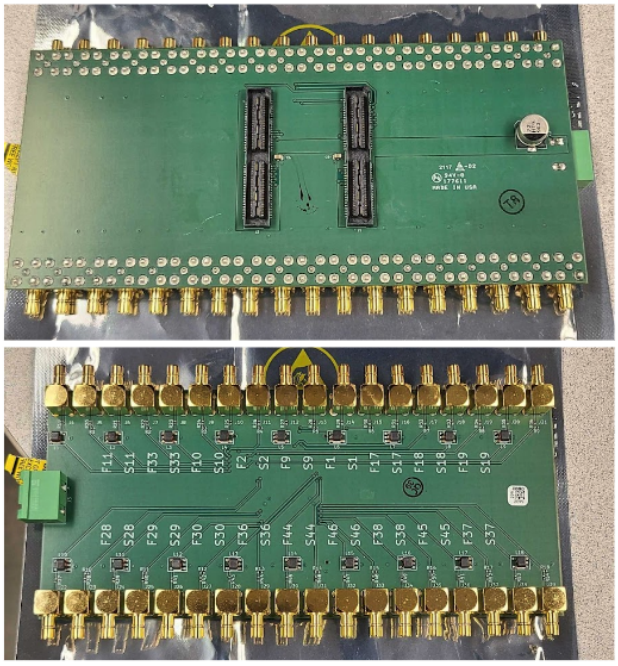

Figure 5.17: A close up of the custom fanout boards designed for use in the SiPM + Scintillator characterization testbed. The board allows for directly plugging in a SensL J-series 8x8 SiPM array and SOUT and FOUT signals of selected channels are fanned-out to respective SMB connectors. The green connector delivers bias to the SiPM array.

5.3.1 Trigger and Energy Calibrations

There are several steps that need to be completed before we are able to directly study PSD capability of our initial selection of scintillators (EJ276 and Stilbene) and perform a validation study of the real-time PSD circuit design. Characterizing baseline noise of the testbed was the first step and was already discussed briefly in the previous section.

After noise, optimization of trigger threshold selection is an important next step. The choice of trigger threshold for a given data acquisition is an important decision to prevent triggering bias in the collected data. Generally, a static trigger value for data acquisition

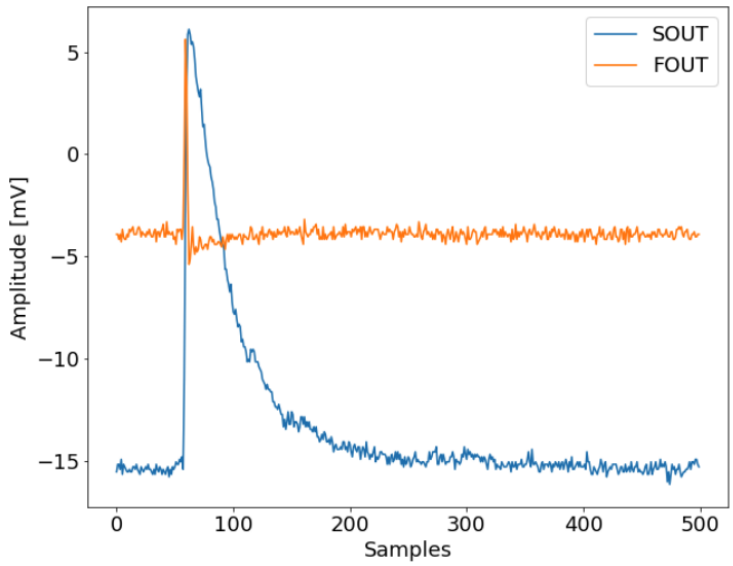

Figure 5.18: Digitized (SOUT in blue and FOUT in orange) waveforms from the testbed are plotted. The sampling rate of the 14-bit resolution CAEN digitizer used is 250MS/s.

is not ideal for many reasons. The 2 channels of the CAEN digitizer, which are used, have small but significant relative baseline offsets between them. Because of the small gain from the 50Ω resistor for SOUT and no amplification of the FOUT signals, even a 1 mV change in triggering threshold will have significant effects. A 1mV relative baseline offset with a static trigger value, has the same effect. Dark counts, discussed already, are the dominant source of noise in SiPMs. Single photon generated pulses (and equivalently, thermally generated dark count pulses) are not discernible above the baseline noise with the current testbed setup. The 50x transimpedance gain is not enough amplification for those purposes. Moreover, as a result of operating the testbed at ambient temperatures (and the resulting high O(100) kHz/mm^2 dark count rate at those temperatures), there are significant rates of double and

triple pile-ups of dark counts. Simple Monte Carlo studies have shown even quadruple pile-ups produce non-negligible rates of O(10-100) mHz depending on the temperature. To prevent dead time issues due to triggering on dark count pileups, we need to set trigger threshold above 4 or 5 photons detected (phd) equivalent amplitude (depending on ambient temperature) to ensure triggered dark count pile up rates don't exceed O(10) mHz. Trying to reduce triggered rates below this will depend on the particular application. Generally, if we try to minimize dark count pile up rate by increasing the triggering threshold, we begin to encroach into the region of true O(10) keVee energy depositions.

In the end, an optimized procedure was decided to dynamically select trigger threshold for each new data acquisition in reference to the measured baseline. After calculating the baseline of the digitizer channel reading in the SOUT input, we add 20 ADC counts (approx. 2.5 mV) to it. This serves as the first guess at the trigger threshold. Generally speaking, this should bring dark count triggering to O(10) mHz, and equivalently, the threshold is above approximately 4 phd mean amplitude. Depending on actual temperature, triggered rates may still be higher than this due to the strong non-linear dependence of dark count rate on temperature. In such cases, we increase triggering threshold by 10 ADC counts (approximately 1.2mV) until the rates fall down to O(10) mHz. Although not an exact mapping due to the uncertainty of the single photon peak amplitude, the procedure roughly equates to placing a triggering threshold above the mean 4-5 phd equivalent amplitude.

The final testbed calibration to be done is the extraction of an expected photons/keV conversion factor. This allows us to convert from a total integrated pulse area to an equivalent

energy for each recorded waveform. We need to calculate this factor for the samples of EJ276 and Stilbene scintillators used in the testbed. It is expected to be larger for Stilbene than for EJ276 due to its higher reported light yield in literature. Although it is simple to frame the task at hand, it is a bit more complicated than expected to actually calculate this factor. Raw digitized waveforms can be converted from ADC counts per sample number for the full event window to mV/ns by using the resolution and sampling rate of the digitizer. The CAEN DT5725 digitizer allows for selecting a 2V input range within an overall -2V to +2 V range. This is done via a 16 bit digital to analog converter (DAC) that specifies a DC offset within the full input range. Setting the DC offset at the midpoint (32768 ADC counts), as was done, results in an analog signal input range for the digitizer to +/- 1V. With the minimal or lack of amplification of SOUT and FOUT respectively, any saturation of digitized waveforms was not observed.

Software integration of the waveforms gives an area in units of [mV*ns] (for a selected integration window). This can be converted to an equivalent integrated charge as

$$Area[C] = \frac{Area[V * s]}{G_{FB}[V/A])} \tag{5.7}$$

with G_{FB} the SOUT transimpedance gain on the fanout board (50x). From there, we convert charge to number of electrons by using the definition of Coulomb

$$C = \frac{-Q(e)}{1.602176634 * 10^{-19}}, \tag{5.8}$$

where Q(e) is the electronic charge. Finally, the number of photons detected [phd] can be calculated using

$$Area[phd] = \frac{Area[electrons]}{G(OV)} \tag{5.9}$$

with G(OV) the gain at a specified overvoltage as obtained from the datasheet [80]. Now, we have a way of obtaining a number of photons detected for any recorded waveform (energy deposition). The final piece of the puzzle is using sealed gamma-ray radioactive sources with known energy peaks. The sources used for this purpose, ^{57}Co, ^{109}Cd, ^{137}Cs, and ^{241}Americium-Berrylium (AmBe), allowed us to perform a linear interpolation. The fitted peaks are described in Table 5.3 with nuclear data obtained from [81].

Calibration Sources and Main Peaks			
Source	Peak [keV]	Type	Note
^{57}Co	14.41, 122.06	gamma	nuclear de-excitation of ^{57}Fe
^{109}Cd	22	x-ray	average X-ray energy from K-shell de-excitations
^{137}Cs	32.2	x-ray	average X-ray energy from K-shell de-excitations
AmBe (^{241}Am)	59.54	gamma	nuclear de-excitation of ^{237}Np

Table 5.3: This table summarizes the main peaks that were used in energy calibrations for EJ276 and Stilbene in the testbed. Nuclear data is obtained from [81].

For each sealed source and scintillator combination, 30 million triggered events were recorded. Afterwards, the conversion to phd was performed as described above, and histograms of areas (in phd equivalent) were generated. Peak fitting was then performed to

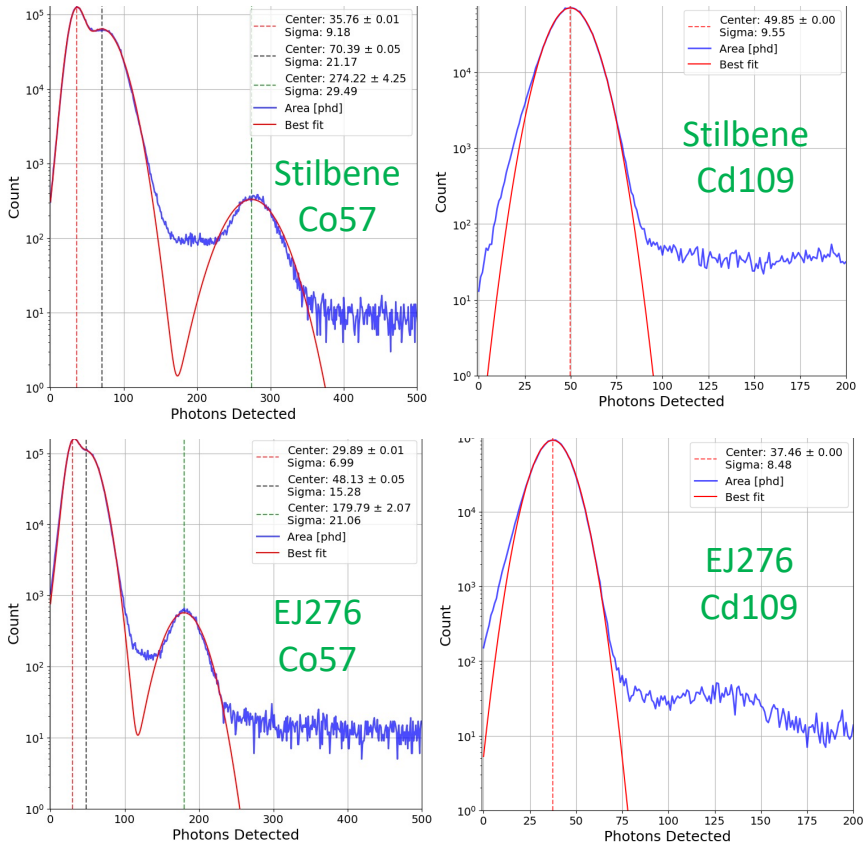

Figure 5.19: The results from fitting known gamma peaks for Co-57 and Cd-109 sealed sources for data acquired from the testbed. Best fits are shown in orange and the centers each fitted peak are shown as dotted vertical lines.

determine the mean for expected gamma-ray emission lines. This procedure was repeated for each source and scintillator type. Figure 5.19 shows the results of this procedure for ^{57}Co and ^{109}Cd. For ^{109}Cd data, the observed Gaussian is the 22 keV average X-ray energy peak for both type of scintillators. For ^{57}Co, the observations are a bit different for each scintillator type. In both cases, we observe two peaks: the 14.41 keV line, and the 122.06 keV line. For Stilbene, there is a clear shoulder before the second peak.

For EJ276, the shoulder is merged into the first peak due to lower light yield. However, there is still a very clear and separated 122.06 keV peak. As a result, our fitted 1σ width for 14.41 keV is wider than the true expected value. In general, and to be conservative with errors across all sources, we have taken the fitted 1σ width (in phd) of each peak as the error on the mean of the peak's fit, which overestimates the error across all sources. A plot of the mean of the fits, and this error, versus energy is shown in 5.20. A straight line fit for both scintillators allowed for the extraction of the slope which provides the phd/keV conversion factor. The calculated factors are 2.20 phd/keV for Stilbene and 1.58 phd/keV for EJ276. These values reflect the lower light yield expected for EJ276 when compared to Stilbene.

Now that trigger selection and energy calibrations have been covered, we will present PSD study results for both EJ276 and Stilbene. This will introduce redundancies with respect to the feasibility study, hence, we pause to describe neutron scatter camera design.

5.4 Neutron Scatter Camera Requirements

As discussed, we are targeting the development of a compact and portable neutron scatter camera. Such an imaging device has several important applications in nuclear security, beam line monitoring, and nuclear recoil calibration systems for dark matter and neutrino physics experiments. The camera design will use (optically isolated) segmented plastic/organic scintillator blocks coupled to individual cells in a full 8x8 SensL J-series SiPM array.

A neutron camera "images" an incident neutron flux using scattering kinematics. The

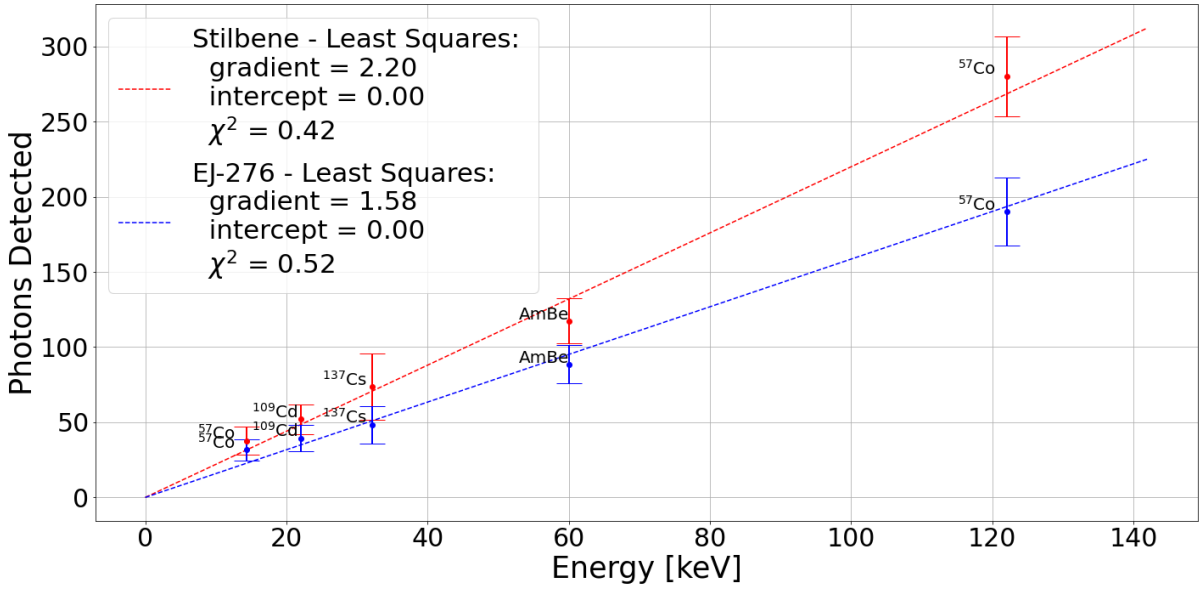

Figure 5.20: All peak fits, with uncertainties, in number of photons detected [phd] are plotted for both Stilbene and EJ-276 scintillators. A straight line was fit for each scintillator, respectively. Stilbene is measured to yield 2.20 phd/keV and EJ-276 yields 1.58 phd/keV.

term imaging is used loosely since the scatter camera is constraining incident neutron direction to a back-projected cone rather than providing fine resolution reconstruction of a given source of neutrons [82],[83]. In general, for such camera designs, two successive interactions of the incident neutron within an active volume are required in order to constrain the angle of incidence. At the same time, single neutron scattering events can be used to measure the total neutron flux and provide count rates. For reconstruction of incident direction using scattering kinematics, as will be shown, we need the location of both neutron scatters within the camera volume, the energy deposition of the initial scatter and the time delay between the two scatters [82]. Figure 5.21 provides a schematic representation of how these quantities

allow for constraint of the initial scatter direction.

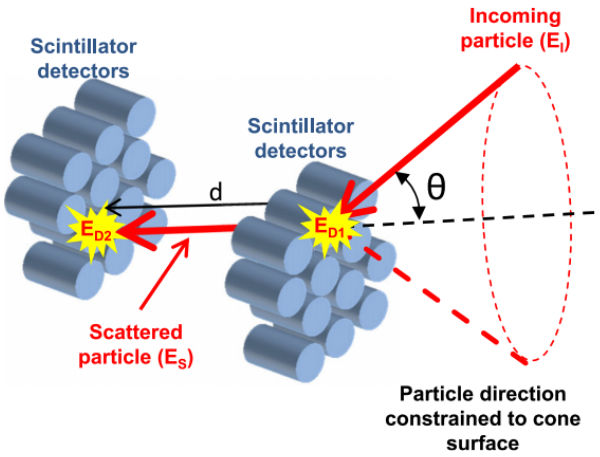

Figure 5.21: A diagram showing the variables required for the reconstruction of the incident abgle of a neutron in a general neutron scatter camera design [83].

To see why these specific quantities are required, we can start by writing the formula for the energy of an elastically scattered neutron in the center of mass (COM) frame as

$$E_{n'} = \frac{(1+\alpha) + (1-\alpha)\cos\theta_{CM}}{2} E_n \tag{5.10}$$

with

$$\alpha = \left(\frac{A-1}{A+1}\right)^2 \tag{5.11}$$

In the above equations (see [82],[84]), A is the mass of the target nucleus, $E_{n'}$ the scattered neutron energy, E_n the incident neutron energy, and θ_{CM} the neutron scatter angle in the COM frame. As discussed, in this work, we will be limiting possible scintillator selection

for a final camera design to plastic/organic materials. Both materials have an abundance of hydrogen nuclei so we can reasonably set A = 1 ($\alpha = 0$). We can therefore rewrite Eqn. 5.10 as

$$E_{n'} = \frac{1 + \cos\theta_{CM}}{2} E_n \tag{5.12}$$

Now, the problem with using Eqn. 5.12 directly is that we won't have access to the true incident neutron energy. However, we can define this to be

$$E_n = E_p + E_{n'} \tag{5.13}$$

with E_n and $E_{n'}$ as defined above and E_p the proton recoil energy when a neutron elastically scatters off a hydrogen nucleus in the scintillator material. The latter can be defined as

$$E_p = E_n \sin^2\theta_L \tag{5.14}$$

with θ_L defined as the angle between the incident and scattered neutron in the lab frame through an appropriate COM to lab frame transformation of a modified Eqn. 5.12. This allows us to change variables between energy of the scattered neutron to energy of the recoiling proton. It should be noted that E_p can also be identified as the energy deposition within the scintillator volume from a neutron scatter. Eqn. 5.14 can be rewritten as

$$\theta_L = \sin^{-1}\left(\sqrt{\frac{E_p}{E_{n'}}}\right) \tag{5.15}$$

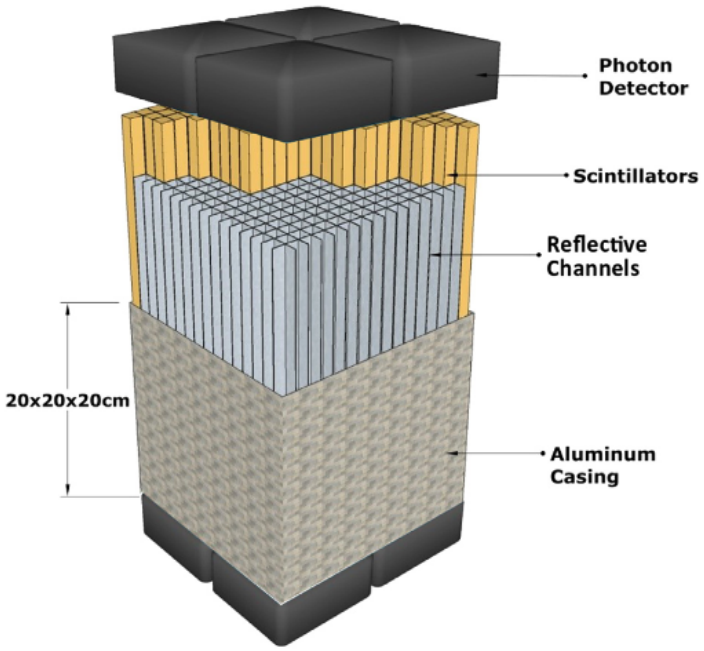

Figure 5.22: From [82], an example of a segmented neutron camera design using plastic/organic scintillator materials coupled to PMTs/SiPM photo-sensors.

So, we observe that knowing the energy deposited by a scattered neutron in the scintillator volume and its remaining energy allows us to estimate the scattering angle. All we need, then, is a formula for the latter. We can now motivate why we need to be able to resolve double scatters of neutrons within the camera volume. Initially, it might be tempting to estimate $E_{n'}$ as the energy deposited in the second neutron scatter. However, for a compact neutron scatter camera design, it is not expected that the neutron will thermalize within the limited scintillator volume. The neutron will most likely exit the camera volume after

the second scattering event, and carry some energy with it. As a result, we need another way to obtain $E_{n'}$ while still taking advantage of double scatters. The key to doing this is to note that the camera is designed to image neutron sources that emit fast neutrons with energies of $O(10MeV)$ or less, and hence they are non-relativistic [84]. Because of this, we can estimate energy of the scattered neutron as

$$E_{n'} = \frac{1}{2} m_n v^2 = \frac{1}{2} m_n \left(\frac{d}{\Delta t} \right)^2 \tag{5.16}$$

with m_n the mass of the neutron, v the velocity of the scattered neutron, d the distance between the two neutron scatters, and Δt the time between the two scatters [82].

With this final formula, We now have a way to determine the incident neutron energy and the scattering angle for neutrons that double scatter within the scintillator volume. To summarize again, the physically measurable quantities that are needed to reconstruct both of these values, as stated initially, are d, Δt, and E_p.

Although a final design for a complete neutron scatter camera will not be presented in this work, Figure 5.22 shows an example of an optically segmented compact neutron scatter camera design that aligns with our goals. We are now well on the way to discuss design requirements for the ASIC. A schematic preview of this is shown in Figure 5.23, which outlines the quantities to be measured by the read out ASIC. Time of arrival (TOA) refers to the time when the first scatter happened, and the time of flight (TOF) is the difference between the first and second TOAs. There are several considerations/specifications that

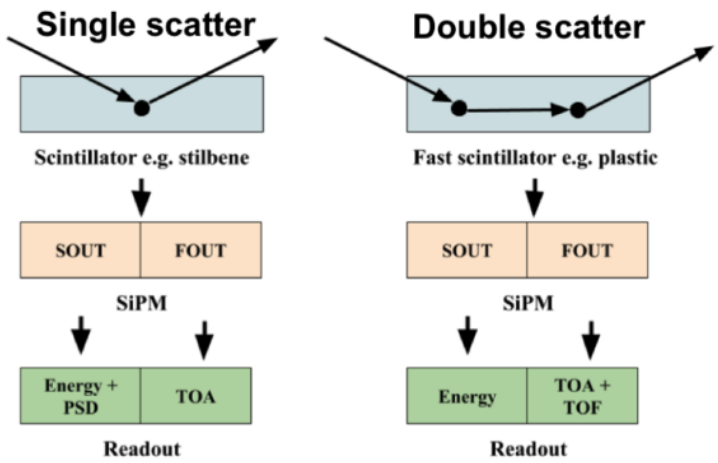

Figure 5.23: A summary of the main signal outputs required for both single and double scatters of fast neutrons within the camera volume for a neutron scatter camera.

need to be discussed both in general for scatter camera design and specifically for our target design.

Identifying fast neutrons from background events (primarily gamma-rays), as hinted at in Figure 5.23, requires the ability to perform pulse shape discrimination (PSD). This ability is crucial to the overall scatter camera design and the efficiency to which this can be done as a function of energy is an important design consideration (for the scatter camera as a whole and the ASIC). One of the main requirements we had for the ASIC design was to include this capability on chip. As mentioned, we have developed a novel real-time PSD circuit to address this need.

PSD relies on the fact that scintillators can have differences in intensities associated with its different decay constants depending on incident particle type (this will be discussed in

much greater detail in the next section). As a result, we expect that these differences will be maintained in the overall pulse shapes as read out by an appropriate optical sensor. The pulse shape will be a convolution of the temporal distribution of the scintillation photons and the single photon response function of the optical sensor. Either photomultiplier tubes (PMTs) or SiPMs are standard optical sensors for such applications. The prototype ASIC is designed for use with the latter and the reasons for such a selection will also be discussed in a later section of this chapter. For the moment, we will summarize by stating that PSD is a powerful tool in allowing us to identify fast neutrons and discriminate against the background gamma-rays. There are several, slightly different, definitions for PSD used in the literature. The definition of PSD as implemented on chip is illustrated in Figure 5.24. It is explicitly defined as

$$PSD = \frac{\int_{t_0}^{t_1} waveform(t)dt}{\int_{t_0}^{t_2} waveform(t)dt} \tag{5.17}$$

with t_0 the start of the waveform, t_1 the length of a partial integration window, and t_2 the length of a total integration window.

5.5 Design Requirements for the ASIC

In the previous section, we discussed the overall scatter camera system into which a final version of PSD_CHIP will be integrated. In this section, we will lay out the requirements for a first prototype version of PSD_CHIP. This chip is implemented in a commercial very

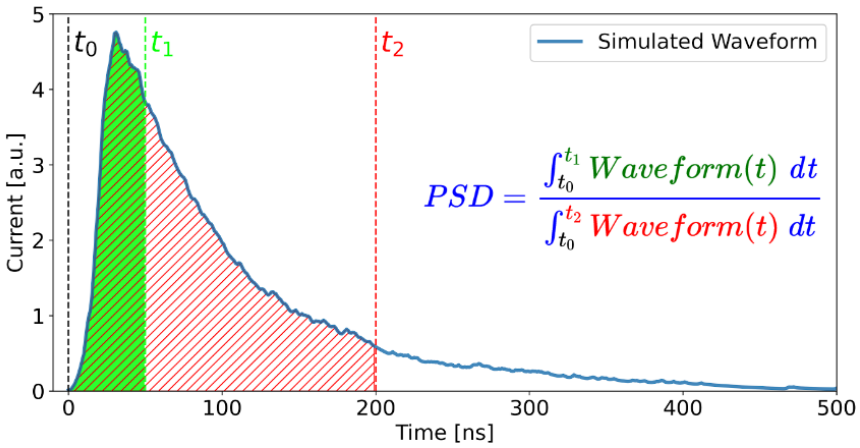

Figure 5.24: A visual definition of PSD ratio as implemented in both versions of the chip.

large scale integration (VLSI) process that supports a minimum feature size of 180nm. This process was selected based on optimizing fabrication costs and ease of learning for a first foray into integrated chip design. To put it into perspective though, even for 180nm (a now more than two decade old process) and minimum reservable area on a multi-project wafer run, the cost was nearly $20k to fabricate the chips.

This prototype is not meant to be a finalized chip design. Rather, it is focused on providing read out of several otherwise inaccessible internal nodes of the chip's full signal chain. Furthermore, the prototype was designed to provide a large degree of tunability for voltage trigger thresholds, baseline reference voltages, current biases, gain, delay lengths, etc. The purpose of focusing on tunability was to provide maximum flexibility in fully exploring neutron scatter camera design parameter space. A more substantive case for needed tunability will be given in tandem to presenting the feasibility study conducted for

key parts of the chip design later in this chapter.

In light of the tunabilty goals for this chip, we decided to only read out only four out of the 64 available cells on the SiPM array. They are enough to measure cross-talk and other potential (noise) couplings between channels either due to design or layout details. A finalized chip design will remove most of the tunability and focus on scaling up channel count to 16, or perhaps 32, channels on a single ASIC. This is expected to also involve porting the design from a 180nm to 65nm process node.

When talking about integrated tunability on chip, one desired design parameter is flexibility in the choice of the scintillator material employed. The primary difference between various PSD capable scintillators, from an electronics perspective, is the intensities associated with the various decay constants. As discussed, this will result in subtle differences in pulse shapes between fast neutrons and gammas. These differences will depend on scintillator properties, either inherent or engineered. In order to maintain chip compatibility with various scintillator materials, we included the ability to adjust partial and total integration windows that are used in calculating the PSD on chip. Several dedicated subsections of the next chapter will discuss the design of these circuits in much more detail.

Zooming in a bit more on the design requirements for the prototype ASIC, a key goal was that the chip needed to be able to read out SensL SiPMs. These chips will be directly coupled to the outputs of a SensL J-series 8x8 SiPM array. As a result, we take in the raw FOUT and SOUT signals per SiPM cell. The chip will need a custom fast front-end pre-amplifier that can deliver a small amount of gain while preserving bandwidth of the FOUT input. For

SOUT, which is a current output, the chip needs to implement a transimpedance amplifier topology that can perform current to voltage conversion while maintaining bandwidth. As shown in Table 5.2, the internal capacitance of the SOUT line is very high, at 4.14nF per 6mmx6mm SiPM cell. This presents a considerable challenge to front-end design. We will discuss this issue and the implemented front-end topologies for both FOUT and SOUT in the next chapter.

The post front-end amplified SOUT waveform will be used to obtain two key parameters of interest needed for neutron scatter cameras: total energy deposition and particle classification as a fast neutron or gamma in the case of single scatters. For the latter, it is unreasonable to expect that we can perform this discrimination between signal and background down to arbitrarily low energy depositions. Although a lofty goal, our target is to provide this capability above 100 keVee energy depositions. To meet this requirement, a novel real-time analog PSD circuit design was implemented on chip. Confirming the feasibility of this design was done with data from the testbed. The details of this are given in subsequent sections of this chapter.

The third key parameter needed is TOA. As discussed in Section 5.4, this is required for neutron double scatter time of flight calculations, which is an input to back projection of a cone around the incident neutron's trajectory. For the prototype, we have implemented a first order approach of directly using the rising edge of the FOUT SiPM signal.

We have now covered the main design requirements for the prototype ASIC above. From these, we can sketch out an initial conceptual overview of a single channel's signal chain.

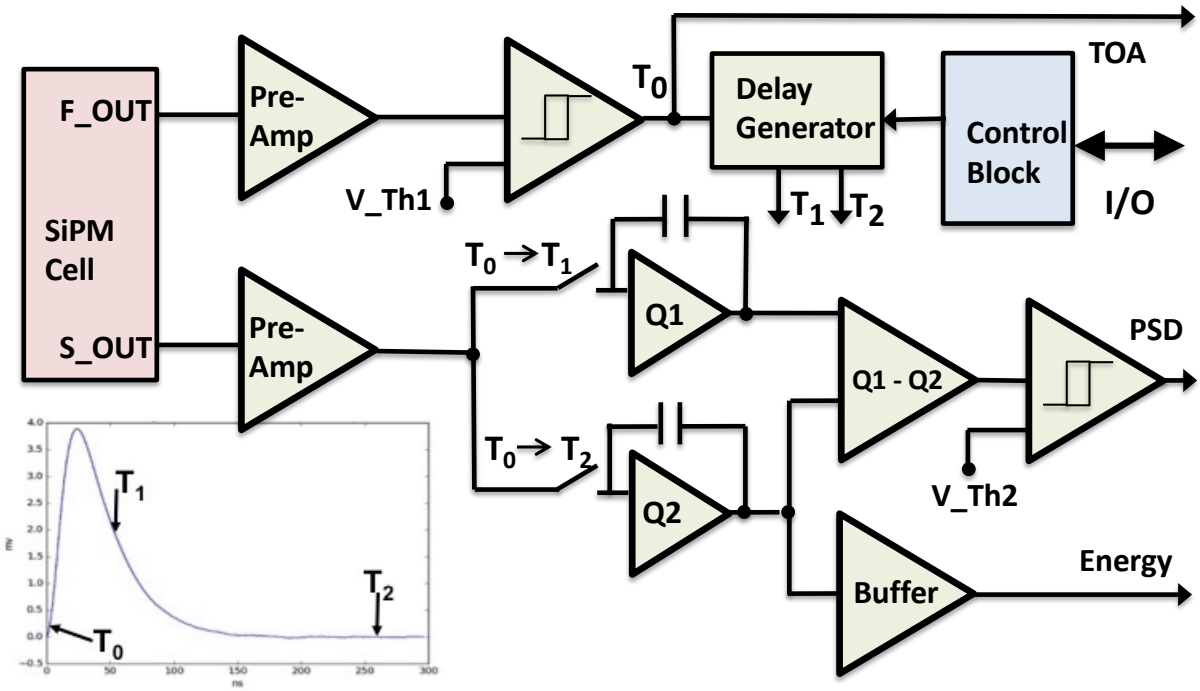

Figure 5.25: A schematic overview of the main blocks of a single channel of the chip.

Figure 5.25 shows the resulting system diagram. Each channel will take in the raw FOUT and SOUT outputs from a SensL SiPM. These are sent into custom pre-amplifier stages, respectively. From there, the FOUT signal is fed into a custom implementation of a discriminator that includes a highly tunable threshold as one of its inputs. The output of the discriminator will be referred to as T0 (or, TOA). In the final design of PSD_CHIP, the TO pulse will be encoded before being sent off chip. For the prototype, a "raw" digital pulse is sent off chip. The T0 pulse is also used internally to generate the integration switches for both partial and total integration of the SOUT input, which are used to calculate an effective

real-time PSD. The actual design of the FOUT signal chain as implemented contains a lot more in-depth configurability, and will be discussed in the next chapter.

The SOUT signal chain is split into parallel partial and total integration stages post front-end. The total integration stage output is buffered and sent off chip to serve as the measured energy deposition of a given interaction. Meanwhile, on chip, a ratio of the two integration outputs will be taken. The details of how this is done forms the backbone of the real-time PSD circuitry and will be discussed further in tandem with the validation study. Afterwards, the output of this stage will be fed into a final classification discriminator, once again, with a highly tunable threshold. This discriminator will trigger on expected differences in this output between fast neutrons and gammas. The digital discriminator output will be a pulse that transitions to HIGH for fast neutrons and remains LOW for gammas.

5.6 Real-time PSD Design Feasibility Study

Now we are ready to return to the testbed and the real-time PSD circuit feasibility study. As mentioned, part of the testbed's goal was to provide an appropriate comparison of PSD capability between EJ276 and Stilbene when coupled to SiPMs. This will be discussed within the scope of the feasibility study.

To start off, let us refer back to Section 5.4 and Figure 5.25. The full details of how an analog-equivalent of the PSD ratio is calculated are provided in the next chapter but a sketch of the principle is given here. As shown in Figure 5.25, there is no explicit ratio

circuit used. Instead, the PSD is implemented via a tunable resistor in the total integration channel. The circuit design for the total integration circuit is shown in Figure 6.14. Without taking into account bandwidth and other performance characteristics, the ideal integration stage transfer function is given as

$$V_{out} = \frac{1}{RC} \int_{t_1}^{t_2} V(t)\, dt \tag{5.18}$$

From Eqn. 5.18, we can see that implementing a tunable resistor (while keeping the capacitor value constant) essentially scales the output of the integration stage. Thus, we can implement a direct analog to the PSD ratio using the ratio between the resistors of the partial and total integration stages. The strategy involves setting the value of the tunable resistor such that the ratio of resistors equals the desired threshold for the PSD ratio. All that is needed then is to quantify the difference from this threshold ratio. For this, we use a subtraction operational amplifier (op-amp) which is fed by the outputs of the two integration stages. Next, the output of the difference op-amp is used to determine PSD, based on whether it is positive or negative. This is converted into a logic pulse to provide the final classification of a waveform as a fast neutron or gamma.

We expect that for input waveforms that have PSD ratios significantly different from the set threshold ratio, the subtraction output will not be near zero. This is the basic principle behind how the real-time analog PSD circuit works on chip but as with most novel designs, it is important to validate this design and understand potential shortcomings and performance

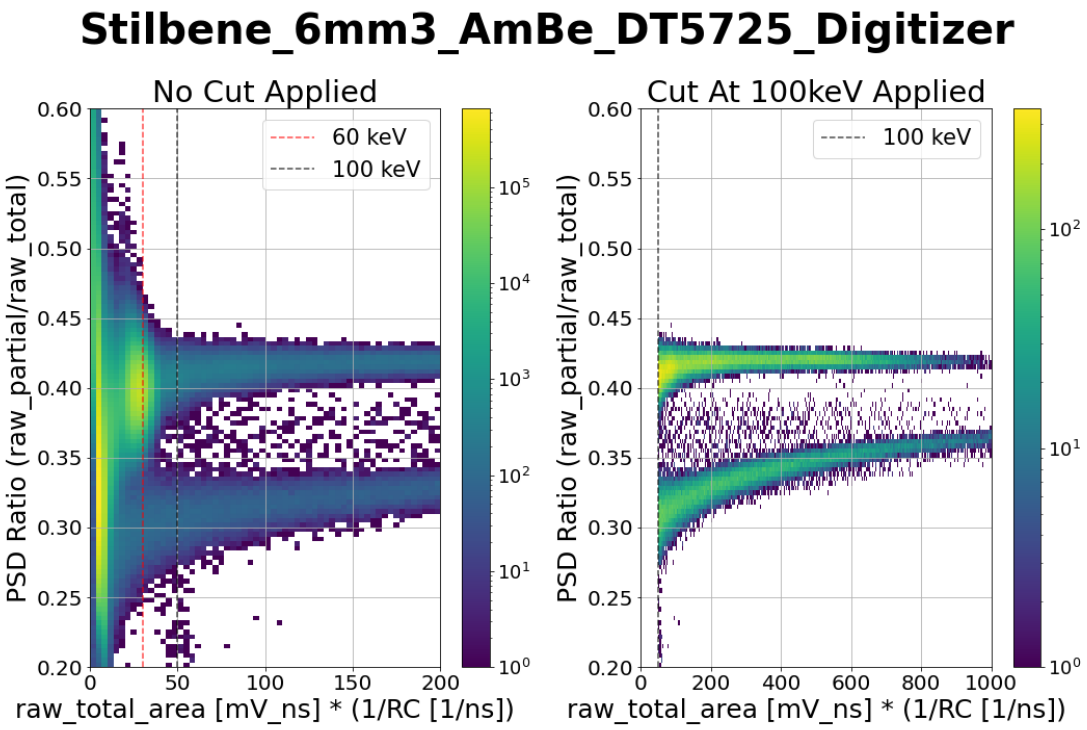

Figure 5.26: PSD ratio versus raw total area scaled by the RC time constant of the integration stage on chip for Stilbene is plotted for 30 million digitized AmBe events. In the right-hand plot, a cut at 100keV was applied and only events passing this cut (shown in the black dotted line) are plotted.

as a function of the deposited energy and scintillation decay constants of the target material. We can do so by processing our digitized calibration data from the testbed through this exact series of steps. It is important to emphasize that in performing this validation study, front-end bandwidth and other electronics response characteristics were not included. This study's main goal is to validate that this real-time analog PSD circuit design is workable.

For this study, conducted for both EJ276 and Stilbene, we used an AmBe calibration source. AmBe is a source that is formed from ^{241}Am and ^{9}Be isotopes. The ^{241}Am emits α

particles which undergo a ^{9}Be(α,n)^{12}C reaction, thus creating a neutron source. However, it also emits a very large percentage of 59.54 keV gammas, which allows us to expose our samples to a simultaneous flux of neutrons and gammas. The AmBe calibration source employed here has a yield of approx. 10^2 neutrons per second, with an average neutron energy between 4-5 MeV, while the maximum energy is nearly 11 MeV [85]. The de-excitation of the daughter ^{12}C produced in the (α,n) reaction produces a 4.44 MeV gamma, which is typically used as a tag of each neutron [86]. This tagging is not a viable method in our case because of the small angular acceptance of our scintillators under test.

Figure 5.26 shows the distributions of calculated PSD ratios versus the raw total area in [mV*ns] for 30 million recorded waveforms for Stilbene. A total integration window of 800ns and a partial integration window of 80ns were used in generating these distributions. A band containing gamma-induced interactions can be seen to be clearly separated and located above a second band that contains neutron events. The bands blend into each other at lower energies due to a combination of dark count pile up effects and electronics noise. On the left, the red dashed line indicates the 60 keV equivalent raw total area and the black dashed line shows the 100 keV equivalent raw total area. An excess of events along the 60 keV line is visible. On the right, the energy axis is expanded to 1000 mV and a cut at 100 keV equivalent is implemented, which reflects our goal for the minimum energy above which we wish to perform PSD with the chip.

Figure 5.27 shows the same but for EJ276. As a result of the lower light yield in EJ276, the neutron and gamma bands are wider and overall, this scintillator maintains a worse

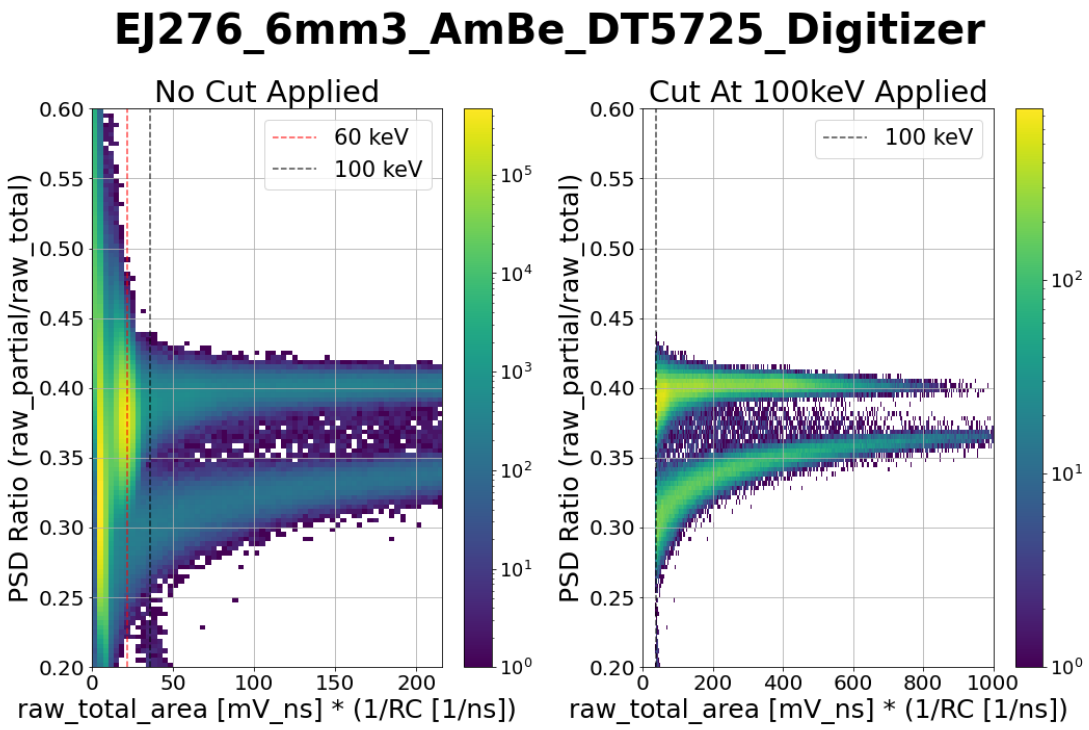

Figure 5.27: PSD ratio versus raw total area scaled by the RC time constant of the integration stage on chip for EJ-276 is plotted for 30 million digitized AmBe events. In the right-hand plot, a cut at 100keV was applied and only events passing this cut (shown in the black dotted line) are plotted.

PSD capability compared to Stilbene, especially at energy depositions near 100keV. The differences in PSD capability will be quantified as we continue progressing through the feasibility study.

We had discussed the fact that various scintillators have differences in the values of decay constants and this can translate into different optimal integration windows for partial and total integration. In order to probe this for EJ276 and Stilbene, we examined the performance of this design for various values of partial and total integration windows. For

each combination of integration windows, we defined a threshold PSD ratio to use. While this is relatively easily to do by hand for any given combination, a procedure will need to be defined so that the selection can be automated as we loop through several combinations. Each pair of integration times is expected to shift the location and widths of the neutron and gamma bands on the PSD ratio axis. Thus, the threshold PSD ratio selection procedure needs to be able to account for this. In the following, we arrived at a cut by projecting all events above 200 keV on to the PSD Ratio axis to create a histogram. Figures 5.28 and 5.29 show the resulting double Gaussian distribution for T_1 of 80ns and T_2 of 800ns for each scintillator. We can fit the distributions with a double Gaussian model and extract the means and σs. 200 keV was used instead of 100 keV to prevent the skew and blending of the bands at lower energies from adversely affecting the ability to fit a double Gaussian in this 1D histogram space. The threshold PSD ratio for this analysis was defined to be the halfway point between the fitted centers of the neutron and gamma Gaussians. The dashed red lines in the 1D histograms show the fit values. These can also be mapped back to the previous 2D histograms as horizontal lines.

Now that a threshold PSD ratio has been defined, the next step is to convert from PSD ratio space to subtraction space in order to model the design of the PSD circuit as implemented in the chip. As a reminder, we set the threshold PSD ratio using the resistor values for the parallel integration stages, but the actual output available to us is the subtraction op-amp output. It is on this output that we need to classify fast neutrons from gammas. In order to convert our 2D distributions to this subtraction space, we scale the x-axis, the

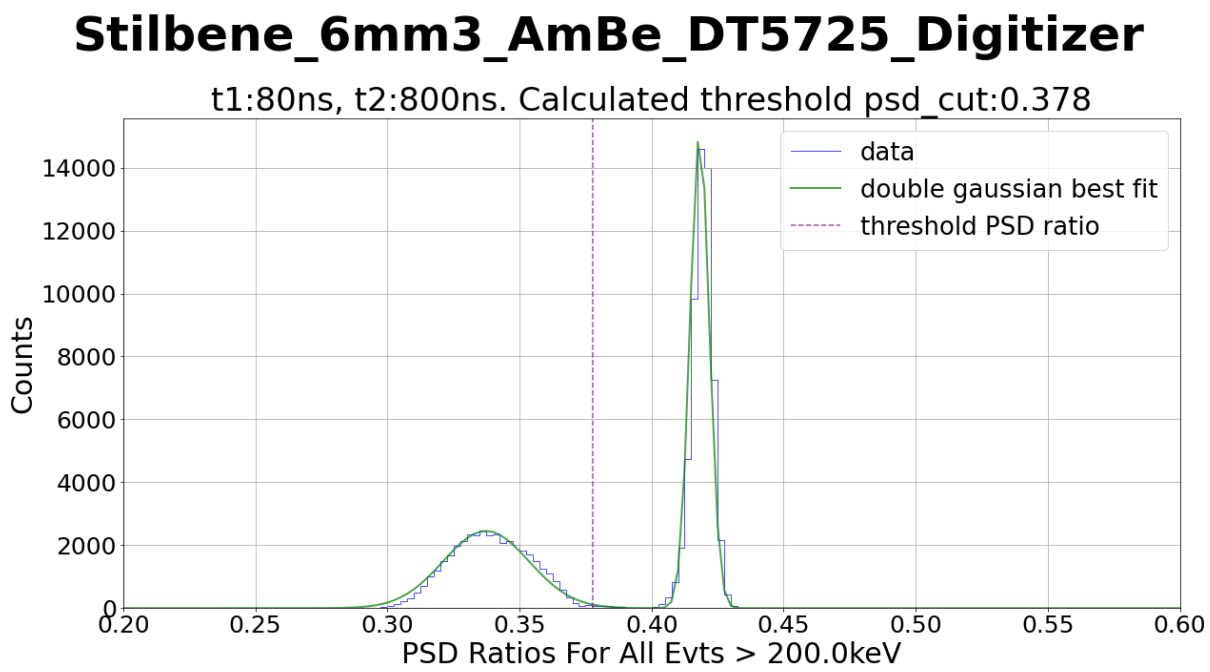

Figure 5.28: After cutting on all events above 200keV equivalent total area, a double Gaussian fit for the two bands was performed for Stilbene. The threshold PSD ratio (dotted purple line) was taken to be the center between the two fitted Gaussian means for the feasibility study.

raw total area by the threshold PSD ratio. On the y-axis, we plot the subtraction output of the partial integration output minus the (scaled) total integration output. For the same partial and total integration windows, these plots for both scintillators are shown in Figures 5.30 and 5.31. The left hand plots in both figures show the separated neutron and gamma bands in this space. The right hand plots show a zoomed in view of the band structure in the region of low scaled areas. The 60 keV excess in the gamma band is clearly visible in both cases, which sets the energy scale. As was the case in the PSD ratio distributions, this new space also has poor PSD capability at lower energies. However, unlike PSD ratio space,

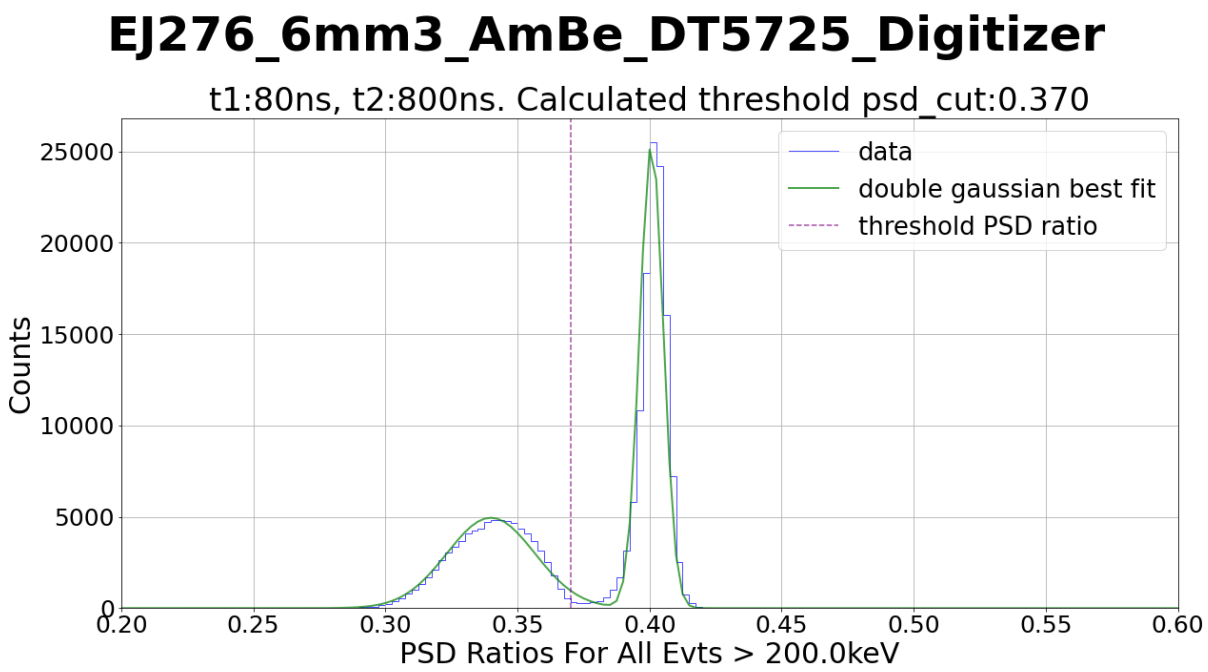

Figure 5.29: After cutting on all events above 200keV equivalent total area, a double Gaussian fit for the two bands was performed for EJ276. The threshold PSD ratio (dotted purple line) was taken to be the center between the two fitted Gaussian means for the feasibility study.

the widths of the bands are much narrower at lower energies and broaden at higher energies. This is because of the nature of subtracting the outputs of the parallel integration stages. For lower energies, the actual integration output values for both channels is expected to be small and thus, the subtraction output will likewise be small. For this reason, the two bands start at a common zero subtraction output and then separate at higher energy depositions.

The broadening of the bands at higher energies can also be explained in terms of subtraction of the two integration outputs. In PSD ratio space, at larger energies, the width of the bands can be mapped back to the choice of integration windows, light yield, and statis-

tical fluctuations in number of photons detected. These widths, combined with the larger subtraction values at higher energies in subtraction space, cause the increasing width of the bands as a function of energy. Overall, the requirement for the chip is to demonstrate PSD capability above 100 keV energy depositions. This is clearly possible for both scintillators with an appropriately chosen final discriminator threshold. Of course, there will be clear differences in the gamma leakage and neutron efficiency between the two that will also depend on chosen integration windows.

Another feature that was not clearly visible, albeit present, in Figures 5.26 and 5.26 is the appearance of extra bands, which are linear and have negative slopes. These can be seen for both scintillators. Further examination of these events revealed that they are due to pile-up of two scattering events within our integration window. These can be easily removed on the basis of the width of the merged waveforms, and do not present a problem.

Since the band structure in subtraction space is not amenable to the collapsing of events above a certain energy cut into a 1D histogram, another standard procedure needs to be selected for proper quantification of PSD capability. These plots clearly show that despite the widening of the bands for higher energy depositions, overall PSD capability (band separation) improves. Thus, it is more important to quantify PSD capability in the region of 100 keV, the minimum energy deposition for which the chip should provide classification between fast neutrons and gammas. In this region, it can be seen that for EJ276 the two bands overlap in the region of interest. It is important to quantify the significance of this overlap. In order to do this, we choose a 100 +/-10 keV window and project all events within that energy range

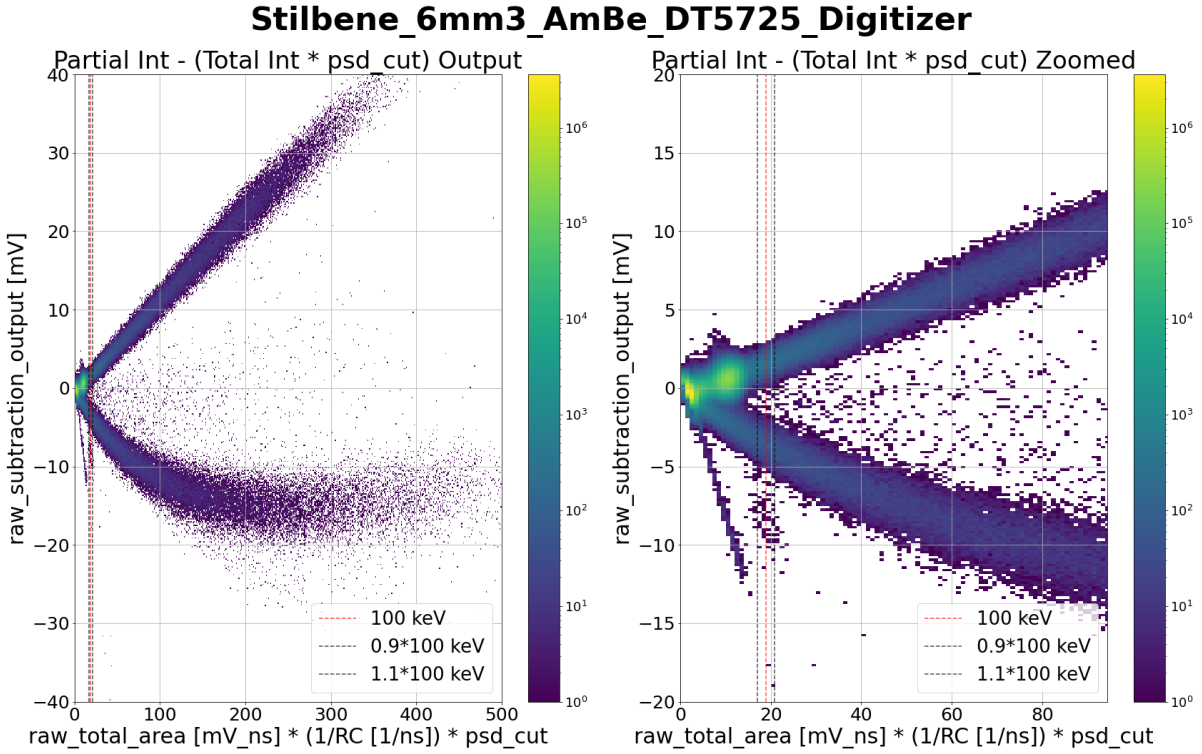

Figure 5.30: Using the extracted threshold PSD ratio, the raw total area scaled by the threshold ratio is plotted vs. the raw subtraction output value for Stilbene. On the right, the plot is zoomed in to show a region around 100 keV. Setting a threshold value for the discriminator following the subtraction stage is equivalent to setting a threshold along the y-axis here.

on to the raw subtraction output axis. The resulting histograms are shown in Figures 5.32 and 5.33. In quantifying the PSD capability of these two scintillators, there are two different metrics that can be used: figure of merit (FOM) and gamma leakage.

Traditionally, the standard FOM metric in the radiation detection community is defined as,

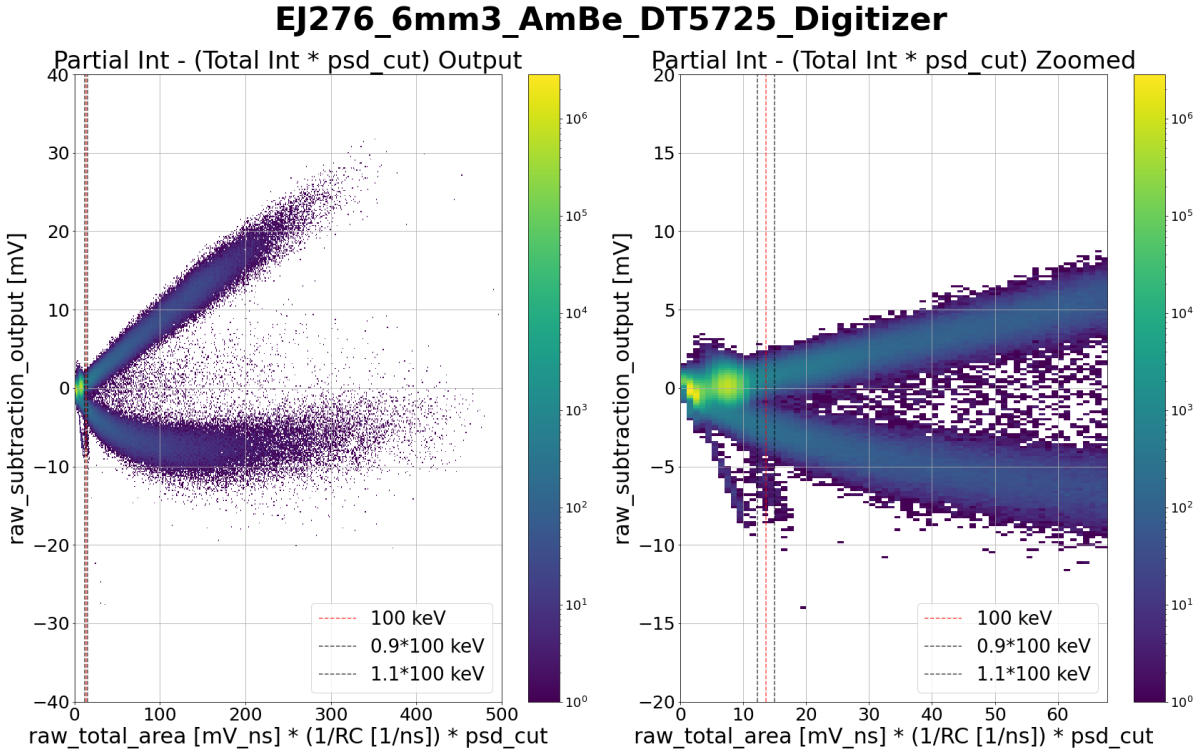

Figure 5.31: Using the extracted threshold PSD ratio, the raw total area scaled by the threshold ratio is plotted vs. the raw subtraction output value for EJ276. On the right, the plot is zoomed in to show a region around 100 keV. Setting a threshold value for the discriminator following the subtraction stage is equivalent to setting a threshold along the y-axis here.

$$FOM = \frac{\mu_\gamma - \mu_n}{2.355 * (\sigma_\gamma + \sigma_n)} \tag{5.19}$$

where μ_γ and μ_n are defined as the means of gamma and neutron Gaussians, respectively, while σ_γ and σ_n are the widths. The metric is a measure of the separation between the two distributions. Gamma leakage, on the other hand, is defined as the percentage of gammas expected to be falsely identified as fast neutrons for a given neutron efficiency. For example,

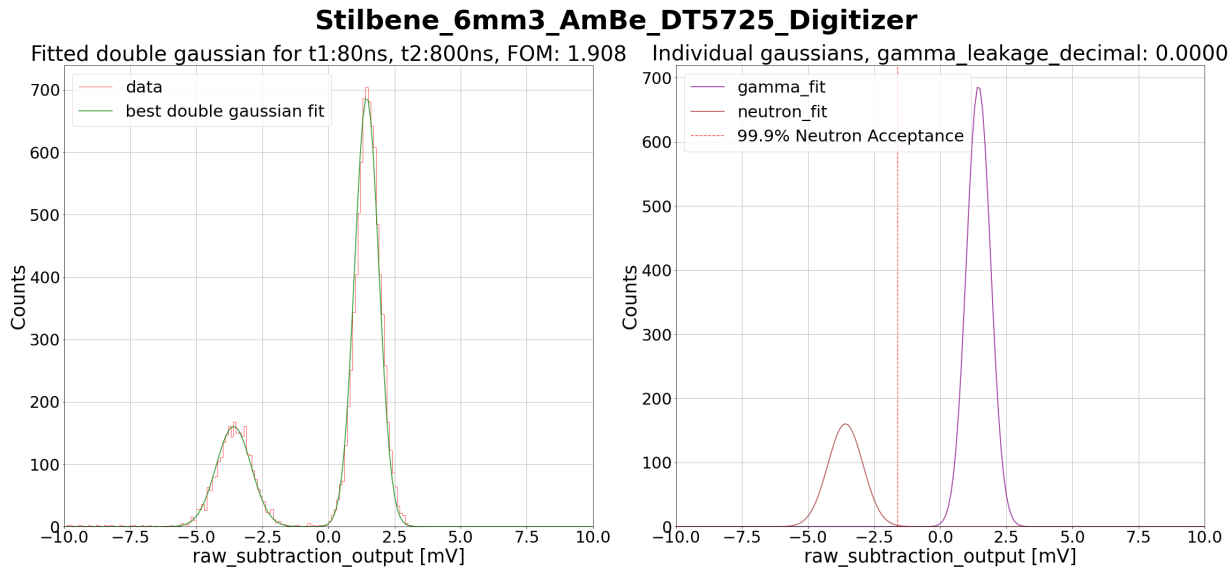

Figure 5.32: On the left, a 1D histogram of all AmBe events within +/-10% of 100keV for Stilbene. A double Gaussian was fit to this data and is overlayed in green. On the right, the separate neutron and gamma Gaussian fits as well as the 99.9% neutron acceptance line are shown.

gamma leakage at 99% neutron efficiency would be identified as

$$Leakage_{\gamma} = \int_{-\infty}^{\epsilon_{99\%}} Gauss(x, \mu_{\gamma}, \sigma_{\gamma})\, dx \tag{5.20}$$

with $\epsilon_{99\%}$ representing the value of x that contains 99% of the area of the neutron Gaussian. In the plots of Figures 5.32 and 5.33, the left hand sides show the actual histogram data in red and the best fit double Gaussian in green. On the right hand sides, the individual gamma and neutron fits are shown in purple and blue respectively. Since the neutron scatter camera will most likely operate in high gamma-ray (background) environments, as an illustrative exercise, we can take a theoretical environment in which we expect a 1MHz

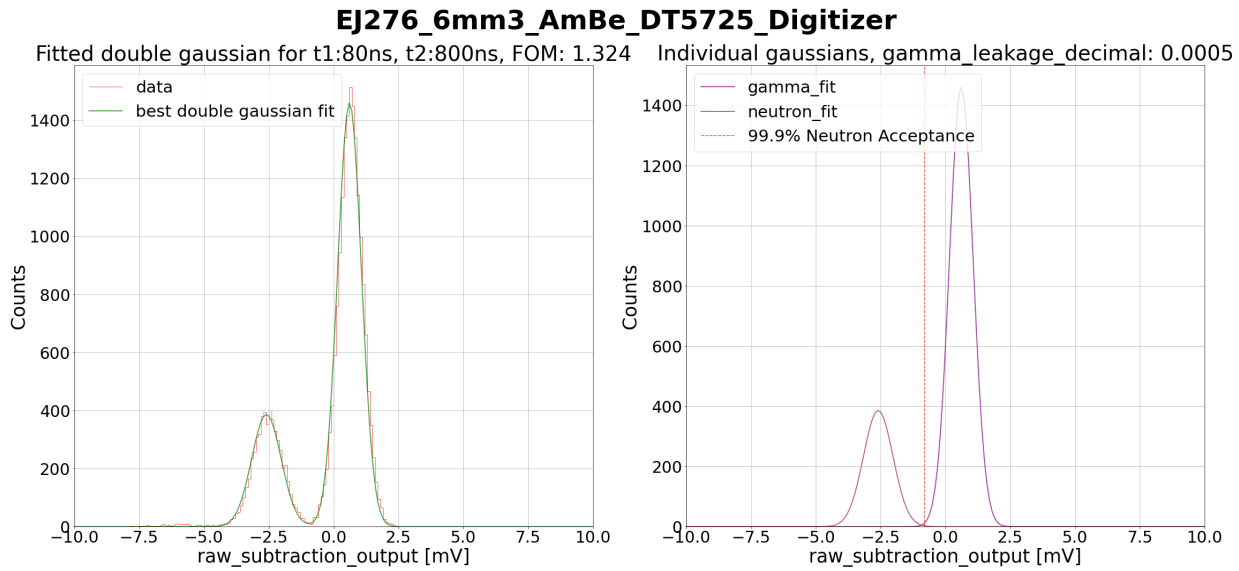

Figure 5.33: On the left, a 1D histogram of all AmBe events within +/-10% of 100keV for EJ276. A double Gaussian was fit to this data and is overlayed in green. On the right, the separate neutron and gamma Gaussian fits as well as the 99.9% neutron acceptance line are shown.

gamma background and 1kHz neutron signal. With 99.9% neutron efficiency, and a 0.1% gamma leakage, we expect to be triggering on 999 neutrons and 1000 gammas per second (not taking into account livetime considerations). In this scenario, in order to ensure only 1% of triggered events are false positives attributed to gamma leakage, we would need to ensure gamma leakage of 0.001% at 99.9% neutron efficiency. Realistically, such leakage rates will be quite hard to achieve and we will need to optimize the threshold in subtraction space. If the desired neutron efficiency is reduced (perhaps to 90% or even as low as 50%), the gamma leakage rate will also reduce drastically. These details will depend on the application.

Overall, between the two considered metrics in this feasibility study, gamma leakage is

the more important metric for evaluating the design of the real-time PSD circuit on chip. It allows for a conclusive calculation of false positive rates, for a selected neutron efficiency threshold, in different possible operating environments. The traditional FOM metric should not be discarded completely, however. It remains a standard benchmark to compare results against other results published in the literature.

Figures 5.34 and 5.35 show the final calculated values of FOM and gamma leakage for several partial and total integration combinations. The step size in covering range of possible partial and total integration values was mostly informed by computing capabilities of the local HPC cluster at UC Davis. The x-axis for all the plots are various total integration windows probed and the y-axis the partial integration windows probed in this study. The left hand plots show gamma leakage for a 99.9% neutron efficiency threshold. The right hand plots show calculated FOM values.

Although the FOM value is calculated under different circumstances, we can compare these values to Figure 5.3, where the FOM values were calculated in PSD ratio space. In contrast, our calculated FOM values are in subtraction space. It is also important to keep in mind that the data used for the calculation in that paper were collected with PMTs as the photodetector in the setup. The longer single photon response of SiPMs, in comparison, should be expected to reduce PSD capability. However, this seems to not be the case. Furthermore, geometry and size of scintillator is known to have an effect on light yield, an important contributor to measured PSD results.

The minimum energy for which FOM was calculated in the figure was 200 keV. The FOM

values calculated for each pair of short and long integration windows was done in a +/- 10% band around 100 keV in the final heatmaps shown in Figures 5.34 and 5.35 . However, equivalent FOM values for the best integration window combination were also calculated for a +/- 10% band around 200keV to provide a more straight forward comparison. For the most optimal partial and total integration window pairing of 40ns and 800ns, for Stilbene, we have a calculated FOM value of 2.022 at 100 keV. Figure 5.2 shows a value of approx. 2.5 at 200 keV. At 200 keV, we obtain 2.670 with our testbed. For EJ276, we see a reported value of approx. 1.45 at 200 keV. Our calculated FOM for the same pairing of integration values at 100 keV is 1.357. At 200 keV, we obtain a value of 1.751. Although under a different space (PSD vs. subtraction), we have demonstrated improvements in FOM values at the same energies for these two scintillators even when read out by SiPMs.

Furthermore, looking at the calculated gamma leakages at an imposed 99.9% neutron efficiency threshold, we see that for Stilbene, a relatively large number of total and partial integration pairings give us less than 0.005% gamma leakage. For EJ276, the minimum gamma leakage is 0.01% for 40ns, 800ns partial and total integration windows. This would theoretically leave us with a 10% effective gamma leakage under scenario considered previously. We would need to lower the neutron efficiency threshold in order to achieve comparable gamma leakages at a higher neutron efficiency for Stilbene. Regardless of these conclusions, it is important to note that we have demonstrated that theoretically, we are able to maintain PSD capability down to 100 keV energy depositions for both EJ276 and Stilbene even in subtraction space. This provides clear validation of the operating principles for the real-

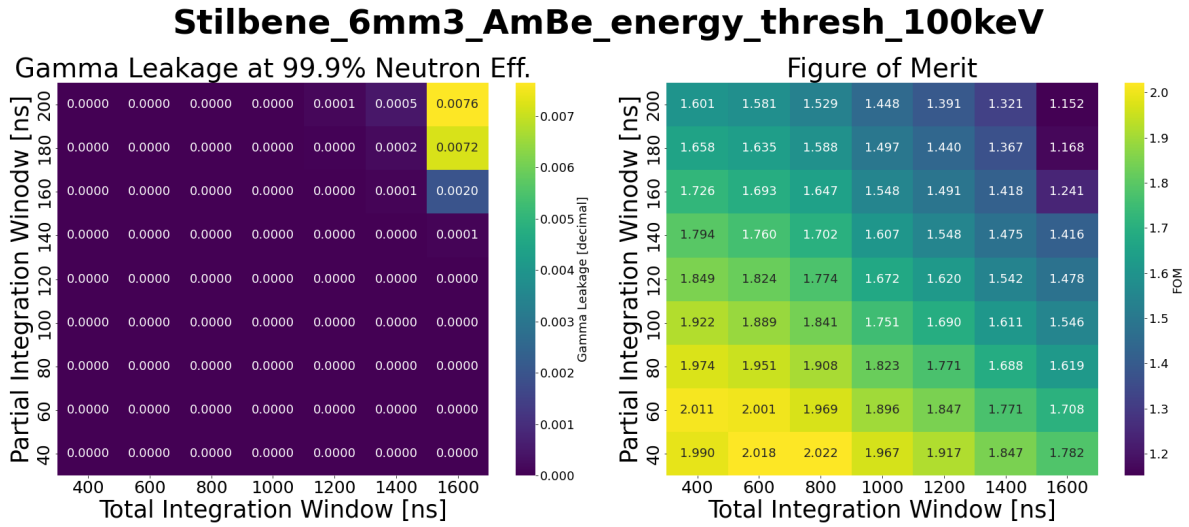

Figure 5.34: Gamma leakage at 99.9% neutron efficiency for a range of partial and total integration windows is shown on the left for Stilbene. Figure of Merit values for the same combinations of partial and total integration windows is shown on the right.

time PSD circuit. Of course, it is also important to restate that this feasibility study was done without taking into account electronics (and bandwidth) differences between the fanout board and the ASIC front end. We will revisit some of the plots after presenting the design of the ASIC and look at simulation outputs using its full signal chain for input waveforms recorded using the testbed.

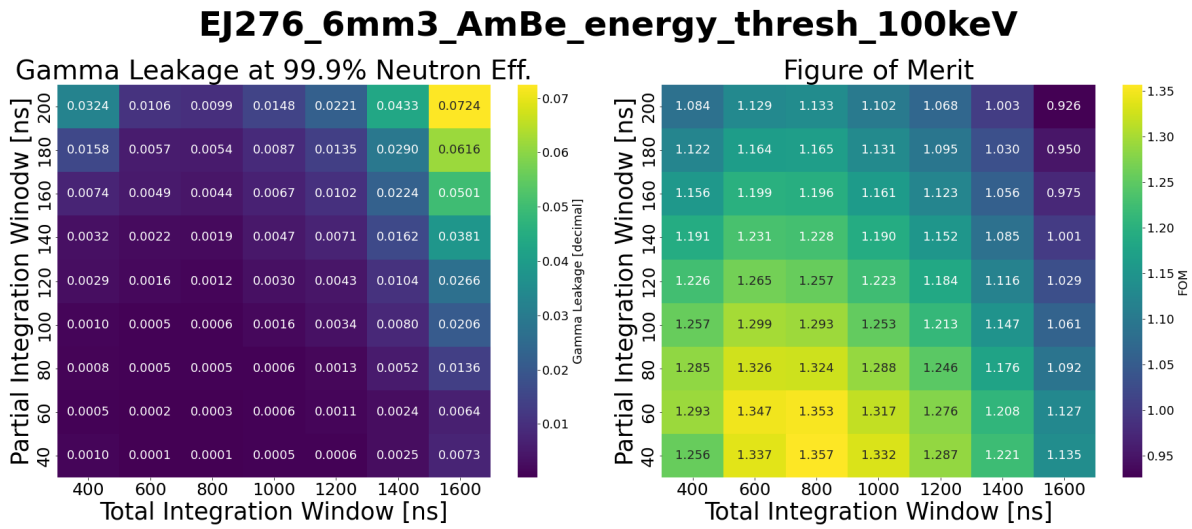

Figure 5.35: Gamma leakage at 99.9% neutron efficiency for a range of partial and total integration windows is shown on the left for EJ276. Figure of Merit values for the same combinations of partial and total integration windows is shown on the right.

Chapter 6

Design, Simulations and Testing of PSD_CHIP

In the previous chapter, we used PSD_CHIP as a general name to refer to the prototype chip designed and fabricated to meet the requirements listed in Section 5.5. However, there are actually two versions of the prototype chip at present. Performance issues observed in V1 from initial testing/verification (originating from layout-related pathologies) led to the design and fabrication of V2. The second version implements a bottom to top improvement in layout (see Subsection 6.2.8). In transitioning to V2, the core design of PSD_CHIP has remained the same. The primary differences, other than layout, between the two versions can be found in the specific implementations of the various hierarchical circuit blocks. V2 provides key improvements in the custom amplifier topologies used for both SOUT and FOUT front-end amplifier design. Furthermore, a redesign of the circuits that generate the

highly tunable partial and total integration enable pulses was also undertaken. Because of the inability to thoroughly test the original design due to the layout issues observed in V1, we implemented both versions in V2 of PSD_CHIP (each tunable delay circuit on two of the four total channels). To prevent redundancy in this chapter when presenting the design of the chip, if specific implementation details for circuit blocks are discussed, we will focus on V2 as appropriate. For blocks with important differences between the two versions, the changes will be explained. Furthermore, simulation outputs (using standard IC design and simulation software from Cadence and Mentor Graphics) will only be presented for the second version of the chip.

Before discussing the chip itself, it is important to mention some general details for ASIC design. At the printed circuit board (PCB) level, due to lower development costs (compared to ASIC design), several rounds of fabrication with design changes (both minor and major) are possible. Furthermore, in-situ board surgery (swapping out resistors and capacitors, cutting traces, adding wire-bridges, switching out IC components, etc) can be performed when feasible and appropriate. This is not the case for ASICs. The fabrication costs, even for 25 year old process nodes like 180 nm, exceed $15k to request a spot on a multi-project wafer (MPW) run (this usually returns 40 bare dies, although additional dies can be requested for much less). Add in the fact that surgery/in-situ fixes are almost never possible for fabricated ASICs, it is not hard to understand why development timescales for complicated chip designs are usually 1-2 years (with multiple engineers), if not longer, and involve at least one prototype fabrication run. In addition, lead times for ASICs are

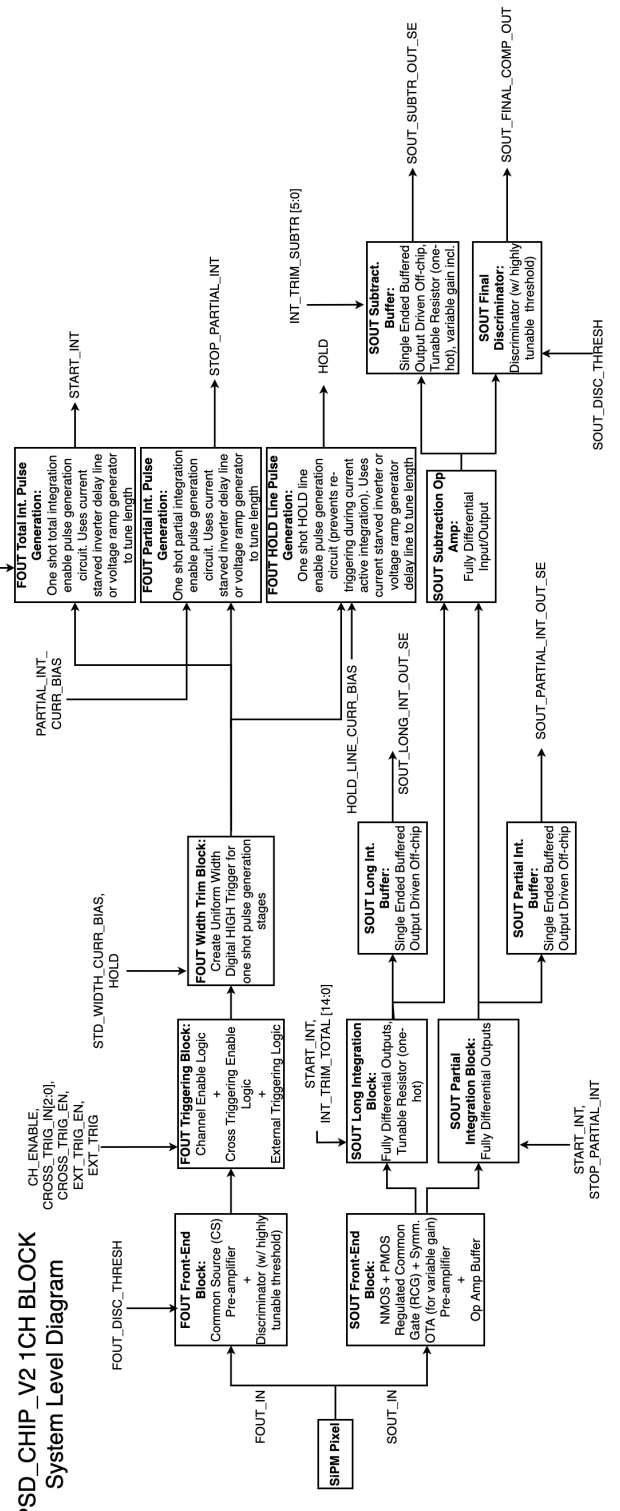

Figure 6.1: An overview of PSD_CHIP_V1's one channel signal chain. Both the FOUT and SOUT signal chains are detailed, including which FOUT signals feed into the SOUT signal chain. Most of the next chapter introduces and describes the various blocks that form the complete one channel signal chain.

roughly 3-4 months after final submission (tape-out) of chip design. There is, of course, the added complication that a custom test PCB must be specified, designed and fabricated in order to allow for initial testing of ASICs. This only adds to costs (in both fabrication and development). Although the design of PSD_CHIP is not overly complex, the base principles of IC design still apply and the same level of detail had to be undertaken to deliver fabricated chips and a custom PCB to perform initial testing/verification.

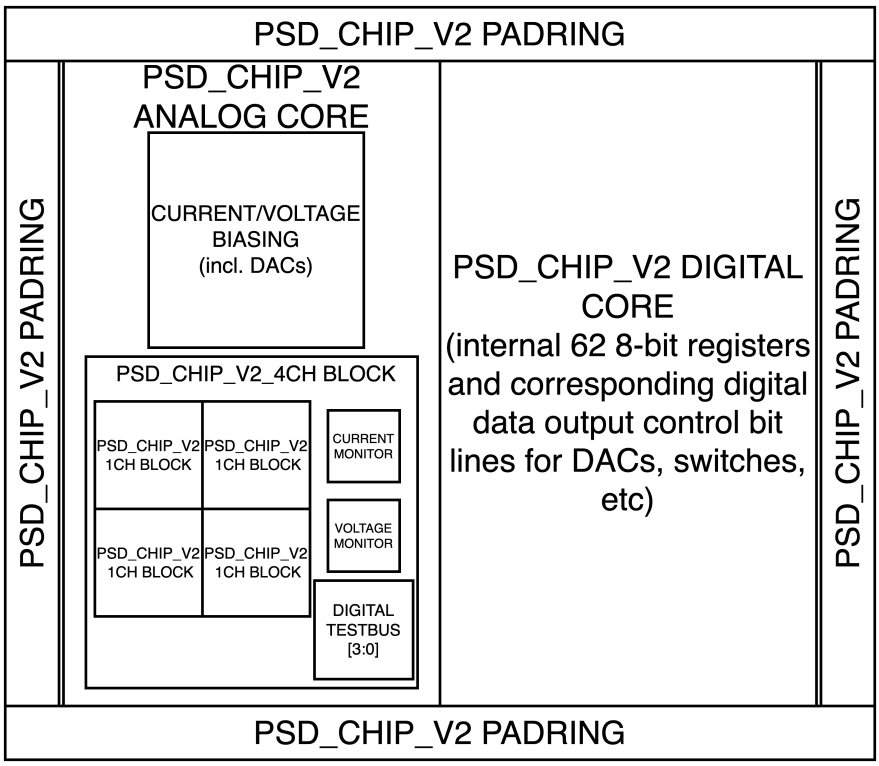

Figure 6.2: A system level overview of the full PSD_CHIP_V2 design. Sizes of blocks in this diagram are not indicative of the physical blocks' sizes on the chip.

There are several ways to approach a description of a full chip design. The approach that we use in this chapter to cover the design of PSD_CHIP is to start by stepping through

the complete one channel signal chain (including both FOUT and SOUT sub-chains) before moving on to the details of scaling to four channels and full chip integration. It is at that point, we detail the implementation of a synthesized digital core to provide channel-level and global tunability of the chip and describe how the chip's global biasing scheme works. Figure 6.2 shows an system level overview of the V2 design. The only difference between the two versions for this figure is that the V1 digital core contains 40 instead of 62 registers. Figure 6.24 explicitly shows the 100 I/O pads of PSD_CHIP (both versions have identical padrings). We shall also include Cadence simulation results at various points in this chapter. Some of these results will be tied back to the results of the feasibility study presented in Chapter 5. It is important to note that datasheets were written for both versions of the chip. These datasheets can be requested from the author of this dissertation. They provide a more detailed look at the design, calibration and operation of the chip(s). This chapter is written as a compliment to the information in the datasheets and provides an overview of why the chip was designed as it was. A concise version of both datasheets has been included as appendices in this work as well (see Appendices A and B).

6.1 The Single Channel Signal Chain

Figure 6.1 offers a detailed overview of PSD_CHIP_V2's one channel signal chain implementation. The goal by the end of this section is to provide a thorough overview of each block in the figure and present Cadence simulation results as appropriate. We start this endeavour by

describing the full one channel chain at a high level and then cover each block individually in dedicated subsections.

Each SensL SiPM (shown in Figure 5.13) has two coupled outputs: FOUT and SOUT. Each channel on the chip takes in these coupled outputs for a single SensL SiPM. The SOUT signal is directly interfaced to the chip and the FOUT signal is connected through a 2nF AC-coupling capacitor. Once on the chip, the signals are directed into customized front-end amplifier stages.

The FOUT signal sub-chain uses a common source (CS) amplifier topology. The main trade-off for this amplifier is between gain and bandwidth. In V1, a more cautious approach led to a higher gain implementation. With the aid of simulations, in V2, the amplifier was specified to have just enough gain (x10) to successfully trigger on 50 keV_{ee} energy depositions. This allowed optimization of the bandwidth for the amplifier. For V2, output pulses with rise times below 1ns was demonstrated in simulations for test square wave inputs. The output of the CS amplifier is fed into a discriminator with a highly tunable (full scale from power to ground) threshold. The threshold sets the minimum FOUT amplitude to generate an output HIGH from the discriminator. Due to a linear correlation between FOUT peak amplitude and total energy deposition, the discriminator threshold is a proxy for selecting a minimum energy deposition to trigger HIGH on.

The SOUT signal sub-chain, on the other hand, has a more involved front-end topology. Due to the SensL SiPM's high junction capacitance and high peak amplitude SOUT current pulse, a specialized topology was selected. Specifically, a regulated common gate (RCG)

amplifier was used (see Figure 6.11). Traditionally, transimpedance amplifiers (TIA) are used in board-level SiPM front-end implementations. The decision to select a RCG amplifier for the chip will be explained later in Subsection 6.1.6. For V1, we specified a RCG with x210 DC transimpedance (current to voltage) gain and single input polarity. Detailed analysis of test bed results post tape-out led to the conclusion that this gain was higher than required. In V2, we implemented a RCG with programmable gain (between. x30 - x180) and dual input polarity capability. This allows V2 to be adaptable to different SiPM biasing setups and maintain dynamic range across various scintillators with a range of light yields. Although the exact dynamic range will depend on scintillator choice, V2 can handle O(MeV_{ee}) energy depositions.

After the respective front-end amplifier stages, the sub-chains for the two inputs are quite different. For the FOUT sub-chain, the output of the discriminator (T0) is fed into a triggering block, which implements channel enable, cross-triggering and external triggering functionalities. The cross-triggering functionality, when enabled, allows a trigger from any of the four channels to trigger all currently enabled channels simultaneously. External triggering allows an external trigger to do the same. When using an external trigger, the discriminator thresholds for all channels are set to the maximum (in reality, minimum since the CS amplifier is inverting) threshold voltage to prevent two concurrent sources of T0 generation. When in normal operating mode, with external triggering and cross-triggering disabled, other than some small amount of propagation delay, the input and output to this block will be the same. Otherwise, the output of the stage will be another channel's TO (from cross-triggering) or

the external trigger.

Regardless, the output of the triggering block proceeds into a width trim block. The reason for this block is subtle. We expect that, for a range of possible energy depositions and incident particle types, the time above (discriminator) threshold (T0 pulse width) can be a range of values. In order to prevent race conditions in the subsequent FOUT blocks that follow the width trim stage, we want to ensure a standard width (in both versions, this width can be specified) pulse as input to these stages. To implement this functionality, a two-stage width trimming circuit was implemented. The first stage trims down the length of any triggers that are longer than the (specified) standard pulse width. The second stage does the exact opposite. It increases the width of any triggers that are shorter than the standard width.

Finally, the now standard width pulse is sent into three parallel programmable delay one-shot generation stages: the total, partial and HOLD blocks. The first two are fairly self-explanatory but the third is designed to prevent re-triggering (of the FOUT sub-chain) during an active integration by ignoring any other triggers that are generated (due to pile-up, etc) in that period of time. The actual T0 pulses are not suppressed and are still sent off-chip (without this capability, we wouldn't be able to externally reconstruct double scatters). It is important to state here that there is nothing within the chip design that prevents this second energy deposition from being integrated in the SOUT sub-chain. There is also nothing on-chip that currently allows for rejection of such pile-up events in the prototype chip, although such functionality can be implemented for a final chip design. Considering that integration

windows are expected to be O(1μs), pile-up is expected to occur in operating environments with > 1 MHz triggering rates per channel (i.e. per SiPM). The design of the programmable delay one-shot generation stages will be covered in detail later in this section. The generated partial and total integration enable pulses from their respective one-shot generator circuits set integration windows in the SOUT sub-chain.

On the SOUT side, after the front-end amplifier, its output is sent into two parallel fully differential integration stages. These stages respectively implement the partial and total integrations required to perform real-time PSD as discussed in Chapter 5. The threshold PSD ratio is set through the use of tunable resistors in the total integration circuit. This ratio is defined by the ratio of the resistors in the partial and total integration stages. In other words, by selecting a value for the tunable resistors, we set the threshold PSD ratio. As mentioned in the feasibility study, we next quantify the difference between the partial and (tuned) total integration outputs using a (fully differential) subtraction stage. The difference is input to a final discriminator, once again, with a highly tunable threshold. Here, the threshold is set to provide minimum gamma leakage at a required neutron acceptance as discussed in the feasibility study of the previous chapter. The output of the discriminator is a single bit that triggers HIGH for a neutron and remains LOW for a gamma/background event.

In a final design of the chip, the three main outputs per channel are: an (analog) energy output (the total integration stage output), a PSD classification bit (the final SOUT discriminator output) and the initial time of arrival trigger (TO; output of the FOUT dis-

criminator). For the prototype, there are several additional outputs per channel including the buffered FOUT and SOUT front-end signals, the partial integration output, and the subtraction output. Furthermore, four digital test buses provide the ability to send several nodes of the FOUT triggering logic off-chip for debugging purposes. These digital test buses are used to provide the T0 output in normal operating mode for chip.

6.1.1 Custom FOUT Front-end Design

The design of the FOUT amplifier was straightforward compared to the equivalent SOUT topology. As already discussed in the previous chapter, the SensL SiPM FOUT output is an ultra-fast voltage pulse with a FWHM of 3ns for a single photon. The input capacitance is also much lower compared to SOUT at only 160 pF (vs. 4140 pF). We had discussed design requirements for the chip and set the minimum energy deposition for which we want to classify fast neutrons from the gamma background at 100 keV_{ee} (with some efficiency). In order to be able to perform studies of PSD capability at that level of energy deposition, we need to be able to trigger on energy depositions that are lower than 100 keV_{ee}. The exact lower triggering threshold is to some level arbitrary and a value of 50 keV_{ee} was selected as a compromise between bandwidth and gain. To be able to trigger on lower energy depositions requires higher gain, and in order to implement higher gain, we would be trading off bandwidth. The details of this are discussed below.

6.1.1.1 The Common Source Amplifier

The common source amplifier is one of the three standard single transistor MOSFET amplifier topologies. A good introduction to MOSFETs and standard amplifier designs can be found in [87]. In this topology, a voltage input is applied to the gate of an NMOS (in general, PMOS can be used as well) transistor and the output voltage is taken at its drain. This is a type of inverting amplifier with the output voltage (with a resistive or active load) negative for a positive input. The DC gain can be approximated as

$$\frac{V_{\text{out}}}{V_{\text{in}}} = -g_m R_L \tag{6.1}$$

with g_m the NMOS transistor transconductance and R_L the total equivalent resistive load seen at the drain. With frequency response included, the transfer function can approximated as

$$\frac{V_{\text{out}}}{V_{\text{in}}} = -\frac{g_m R_L}{1 + sR_L C_L} \tag{6.2}$$

with the new quantity, C_L, the total equivalent capacitive load seen at the output node (drain). A direct trade-off between gain and bandwidth exists in this equation. Bandwidth is primarily determined by the location of the dominant pole (set by $R_L C_L$). For a fixed capacitive load, to increase bandwidth, we need to decrease R_L (reduce gain). Thus, the gain of this amplifier needs to be chosen carefully. We had stated that 50 keV_{ee} equivalent FOUT input waveform amplitude is the minimum signal to trigger on. For this requirement, Cadence

simulations show that a DC gain of $\sim$ x10 is required. Figure 6.3 shows a sample 50 keV$_{ee}$ equivalent FOUT waveform and the output of the CS amplifier (FOUT_FRONTEND_OUT) and the on-chip buffer (FOUT_BUFFER_OUT) that is used to send this waveform off-chip. The shaping of the buffer, simulated with an external 20 pF capacitive load is visible. The DC baseline shift is an expected feature in IC design.

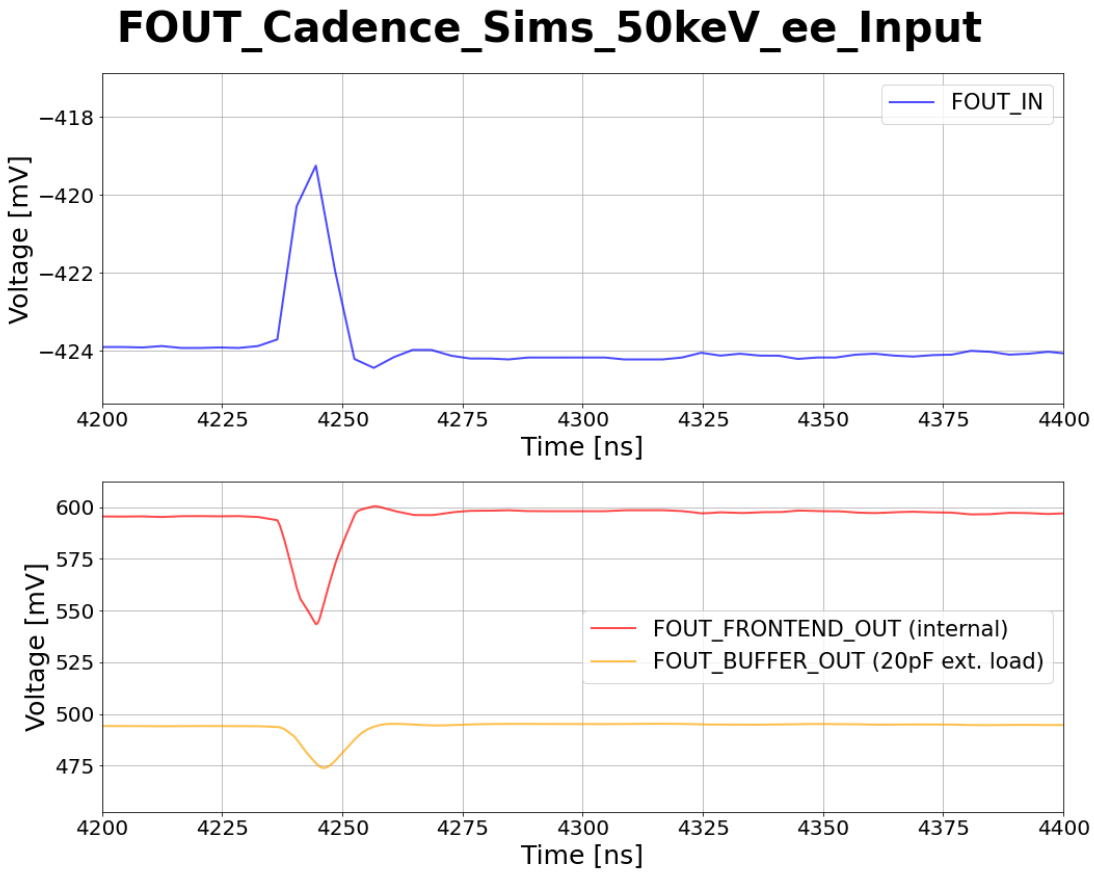

Figure 6.3: A sample (50 keV$_{ee}$, amplitude scaled down from a 500 keV$_{ee}$ pulse) fast neutron FOUT waveform acquired using the test bed was input to the FOUT signal chain in a Cadence simulation (shown in blue). The output from the CS amplifier stage is shown in red. The output post buffer with a 20 pF load is shown in orange.

6.1.2 FOUT Triggering Block

Unlike with the digital core, which shall be discussed towards the end of this section, in order to implement the FOUT triggering block, no synthesized digital block was specified and implemented. Standard digital logic cells provided by the foundry were used to manually construct required functionality of the block. The first feature of the block is the capability to enable/disable each channel individually. It is important to note that this feature only enables/disables the channel within the FOUT sub-chain. Explicitly, SOUT signal chain blocks are not powered down in response to an FOUT channel disable. Although this was the initial plan, the very tight development timelines for both versions of PSD_CHIP meant having to adapt a simplified approach to this functionality. The second feature implemented is a cross-triggering capability to facilitate the study of noise and/or optical coupling between channels either at the scintillator, SiPM, or chip level. Whenever one of the four channels registers a T0 trigger, all enabled channels are triggered simultaneously (within propagation delays). Because this feature is not expected to be used in normal operations of the chip, a separate (from the channel enable/disable functionality) selection of channels to cross-trigger was not implemented. The external triggering feature was implemented in a similar manner. All enabled channels can be triggered through an externally provided differential input low-voltage differential signal (LVDS) trigger (to the chip's relevant pads; see Figure 6.24). This feature is intended to be used to perform signal propagation delay/spread (between channels) studies and, if required, perform general debugging. An appropriate FOUT input test pulse

can also be used for the former if FOUT discriminator time-walk and jitter characterization should also be included in the studies. Although not the primary use case, the external trigger can be used to study scale of variations in the widths of the three one-shot programmable delay lines that generate the partial, total and HOLD enable triggers across channels.

6.1.3 FOUT Width Trim Block

As we move on to the width trim block that follows the triggering block, it is important to first discuss the circuits that we implemented to allow control of digital pulse widths. There are two circuits that are used: the first shortens widths of pulses and the second does the opposite.

We use the circuit shown in Figure 6.4 to perform the former. The input line is fed both directly into one of two inputs of an AND gate and into a delay line (details will be discussed in a subsequent subsection). The output of this delay line is fed into a NOT gate. In the quiescent mode, input A is LOW and B is HIGH. The output, by default, is LOW. When an input digital HIGH pulse arrives, as shown in the timing diagram, input A goes HIGH. However, there is a (programmable) delay before input B is pulled LOW. During this period of time, the output of the circuit is HIGH. Afterwards, the output will return and continue to remain LOW. Thus, for a variable width input pulse, the output will have a fixed output width equal to the length of the set delay in the delay block. If the delay is set to be wider than the width of the initial pulse, the input and output of this circuit will be identical (with

some propagation delay between the two). Similarly, if the input pulse has a width less than the set delay, the output will also be identical.

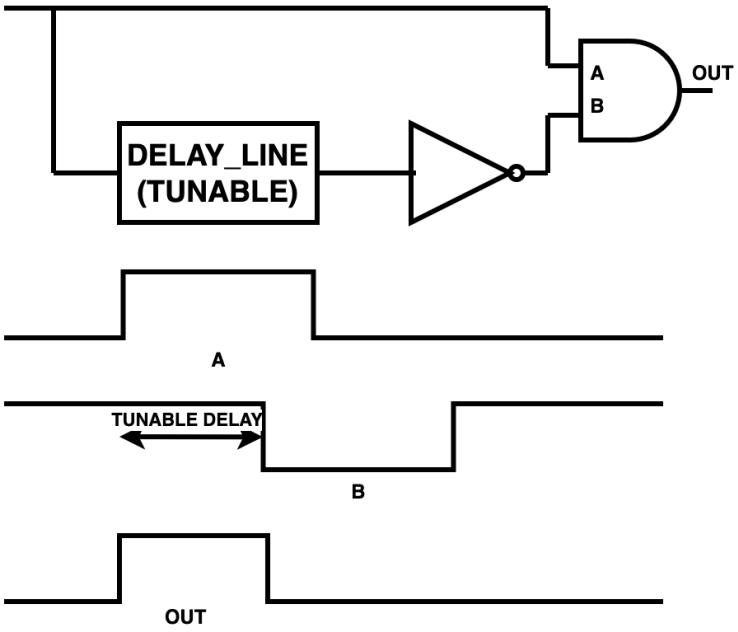

Figure 6.4: The circuit implementation of digital pulse width trimming used on PSD_CHIP. The programmable delay line block is implemented using a chain of current starved inverters (CSI). This circuit is one of two stages used to ensure a standard width for generated T0 triggers before they are fed into the parallel integration enable generation stages.

Figure 6.5 shows an implementation of the opposite circuit. In this circuit, the goal is to take variable width input pulses and increase the width to the set standard width as needed. To do so, we use an SR latch. The truth table for this latch is given in Table 6.1. The input to this circuit is directly connected to the set (S) input and to a programmable delay line. The output of the delay line is connected to a digital buffer. The output from the buffer is fed into the reset (R) input of the SR latch. When the S input is triggered HIGH,

S	R	Q	$\bar{Q}$
0	0	latch	latch
0	1	0	1
1	0	1	0
1	1	illegal	illegal

Table 6.1: The truth table for an SR latch. In the case that both S (set) and R (reset) are LOW/0, the output of both Q and $\bar{Q}$ retain their previous logic states. The opposite is an illegal state and the output will be an undefined state.

Q goes HIGH and $\bar{Q}$ goes LOW. Now, because the SR latch will retain its previous state (latch) after S goes back LOW, the states of Q and $\bar{Q}$ do not change. After a period of time determined by the delay line, the R input is triggered HIGH. Looking at the truth table, we expect then that Q should go LOW and $\bar{Q}$ should go HIGH, thus re-setting the circuit to its quiescent state. In the end, we have a circuit that can increase the width of an input test pulse depending on chosen delay. It is important to note that the reason this circuit follows the previous trimming stage is that if the input pulse is longer than the set delay, then, at a point, both the S and R inputs will be HIGH. This is an illegal state that will lead to an undefined state for the SR latch. To prevent this race condition, the trimming stage was implemented first and the programmed delays (set standard width) is the same between the two circuits.

When the chip is powered on, it is hard to determine in which state these latches power on. As a result, an external SR reset line was implemented (per channel; refer to the chip's padring in Figure 6.24). This line is connected to the R inputs through respective OR gates

along with the output of the NOT gate (for all SR latches used in a channel). The OR gate is not explicitly shown in Figure 6.5. Thus, these external reset lines can be used to force a reset of all SR latches upon powering on the chip or in general debugging scenarios.

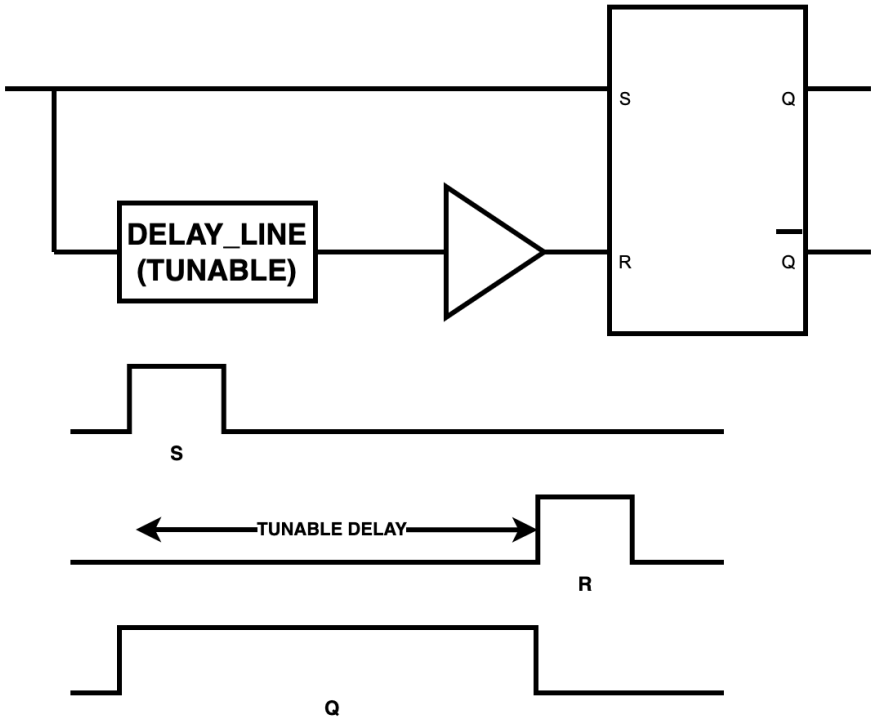

Figure 6.5: The SR latch on-shot generator (with programmable delay) circuit design. The delay line block is implemented through the use of a chain of current starved inverters (CSI) on-chip. This circuit is the core of many parts of the FOUT digital triggering logic. An appropriately selected length of CSIs in series generates the partial, total, and HOLD enable pulses. It is also used as one of two stages to implement a standard width for T0 before it is fed into the parallel integration enable generation stages.

6.1.4 Parallel Integration Enable and Hold Line Generation

The final circuit blocks in the FOUT sub-chain are the programmable one-shot digital pulse generation stages. There are three in parallel as previously discussed. The circuitry for the HOLD and total integration enable stages are identical. The only difference in implementation lies in the programmed delays used. The subtle requirement is that the HOLD line width must always be $>$ the width of the total integration line to prevent race conditions. The partial integration enable generation stage is an altered version of the HOLD/total integration generator circuit since shorter delays are required in comparison. The HOLD line's role was described in the high level overview. It prevents re-triggering during an active integration, which could result in an undefined state for the overall FOUT triggering logic chain. The total and partial integration enable lines are connected to various switches used in the corresponding SOUT integration stages (per channel). Some propagation delay is expected of O(ns) for the signals between the FOUT and SOUT sub-chains.

There are two programmable delay line circuits used in these three one-shot generation blocks: a series of current starved inverters (CSIs; exclusively used in V1 and used on two channels in V2) and the voltage ramp generator (VRG; used on two channels in V2). They will be discussed in the following subsection. For now, it is important to state that they are used to program the delay (and thus the final width) of the pulses from the circuit in Figure 6.5. In other words, the base programmable one-shot generation circuit for these three lines is the same as the width trim block. The main change is the scale of delays

involved, especially for the HOLD and total integration lines. These two lines require delays of > 2 us and as low as a few hundred ns, compared to delays between 10s of ns to a few hundred 100 ns for the partial integration line.

6.1.5 Delay Line Generation

We now discuss the specific programmable delay line implementations. This is one of the key changes between V1 and V2 of the chip. Both versions of the delay line exist in V2 of the chip and so both programmable delay lines are discussed here. The FOUT width trim stages on both versions exclusively use the current starved inverter based delay line. In V2, two channels use the CSI-based delay line and the other two use the VRG-based one for the partial, total integration and HOLD enable lines. In V1, all four channels use the CSI-based delay line. The reason for the switch to the VRG delay line is that the CSI circuit takes up a considerable amount of area and power (during switching) and is prone to degradation in performance due to layout parasitics. On the other hand, the VRG uses a fraction of the area and power and due to the small number of components in the design is not as prone to parasitic-related degredation. We discuss these issues in more detail, including a more detailed comparison, after introducing both programmable delay line designs.

6.1.5.1 Current Starved Inverter Delay Lines

A current starved inverter (CSI) programmable delay line is formed by a number of CSI inverters in series. The natural place to start describing this delay line is with a description

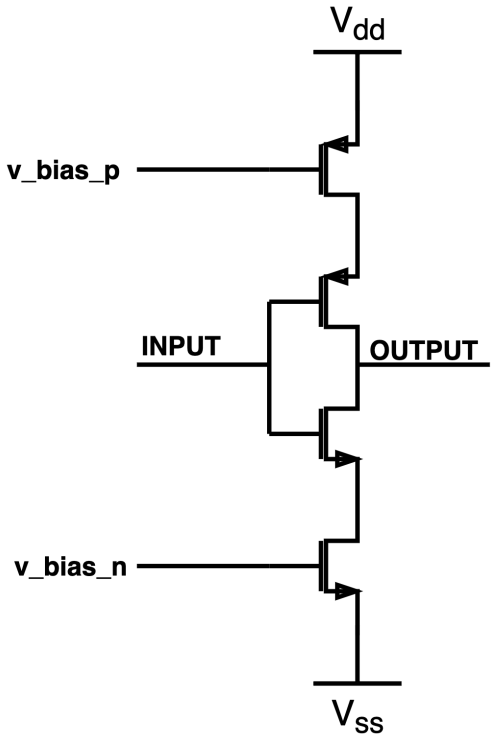

Figure 6.6: A transistor level circuit of a current starved inverter. This inverter is exclusively used as the building block for programmable delay lines in V1 of the chip. It is also used in V2, but not exclusively.

of the base CSI itself. A transistor level CSI circuit is shown in Figure 6.6. This inverter is a modification to the standard CMOS inverter circuit. The core inverter design remains unchanged. However, an additional PMOS and NMOS transistor are included on the top and bottom, respectively in series to VDD (power) and VSSD (ground). The role of these transistors is to "starve" the inverter of current. Without them, there is no limit on current drawn by the central PMOS and NMOS transistors during switching from HIGH $\rightarrow$ LOW and vice-versa. When they are included, the specific values of gate voltages, v_bias_n/p,

correspond to selecting a maximum current that can be drawn during switching. When the inverter is current-limited, it can be said to be slewing (borrowing a term from op amp design). Setting lower maximum current values causes the slewing rate to become slower. This is observed as an overall slower rise/fall time. In reality, we do not set the gate voltages directly. We specify the maximum current bias, which through a current mirror biasing block is translated to appropriate gate voltages for these CSIs. The relationship between supplied current bias and the generated delay for a single pair of CSIs (what we shall call one delay block) is shown in Figure 6.7. We see that the possible range of delays does not follow a linear relationship. This is a limitation to this programmable delay line design. However, such a design works well when a wide range of delays need to be covered with minimal number of bits used.

For an actual programmable delay line implementation, we can vary the current bias for each CSI in an overall chain. However, the range of delays per unit delay block (pair of CSIs) is limited as seen in Figure 6.7. To produce longer delays, additional delay blocks need to be added in series. Implementing a programmable delay line with this circuit requires specifying the total number of delay blocks used, a range of current bias values needed to cover the full range from minimum to maximum delay and a way to select between those different current bias values. For PSD_CHIP, a 5-bit current DAC was used. For the various parts of the FOUT sub-chain, different numbers of delay blocks were used as appropriate. The various numbers of delay blocks used for different parts of the one channel design is listed in Tables A.6 and B.6.

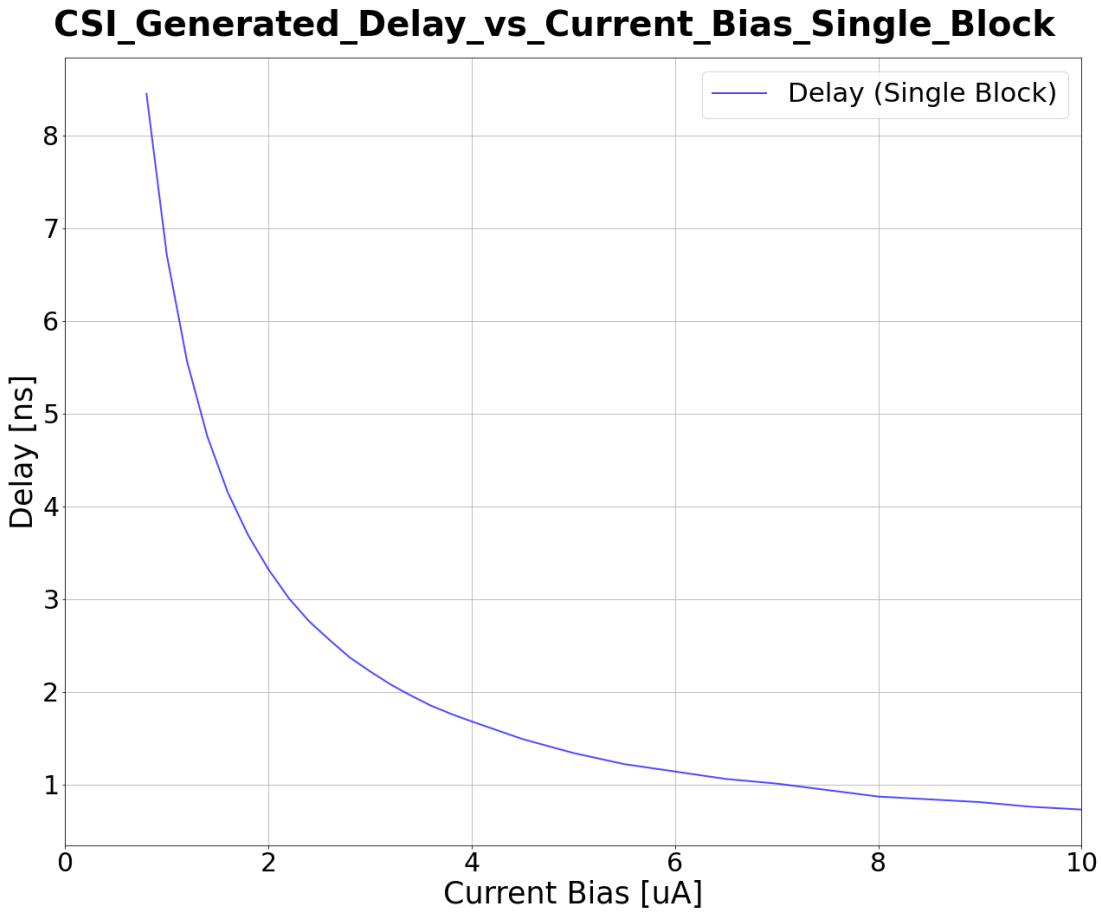

Figure 6.7: The relationship between supplied bias current [uA] and the generated delay for a pair of current starved inverters.

For most CSI-based delay lines in the FOUT sub-chain, a fixed 500nA LSB is used for the associated 5-bit current DAC. This is a 100x division of the 50uA nominal **RBIAS_ANALOG** current bias input to the chip. To allow more flexibility specifically for the total and partial integration enable CSI delay lines, their associated LSBs were made inputs to the chip (**RBIAS_PARTIAL** and **RBIAS_FULL**). The LSB currents can either be

static or tunable through an external DAC setup. In theory, such a setup allows for infinite resolution capability in fine-tuning partial and total integration enable windows. In reality, such resolution is limited by specifics of the external DAC IC and noise. See Figure 6.24 for locations of these three pads on the padring. Tables A.1 and B.1 can be consulted for details on how to interface with these pads to set a specific current bias.

6.1.5.2 Voltage Ramp Generator Delay Lines

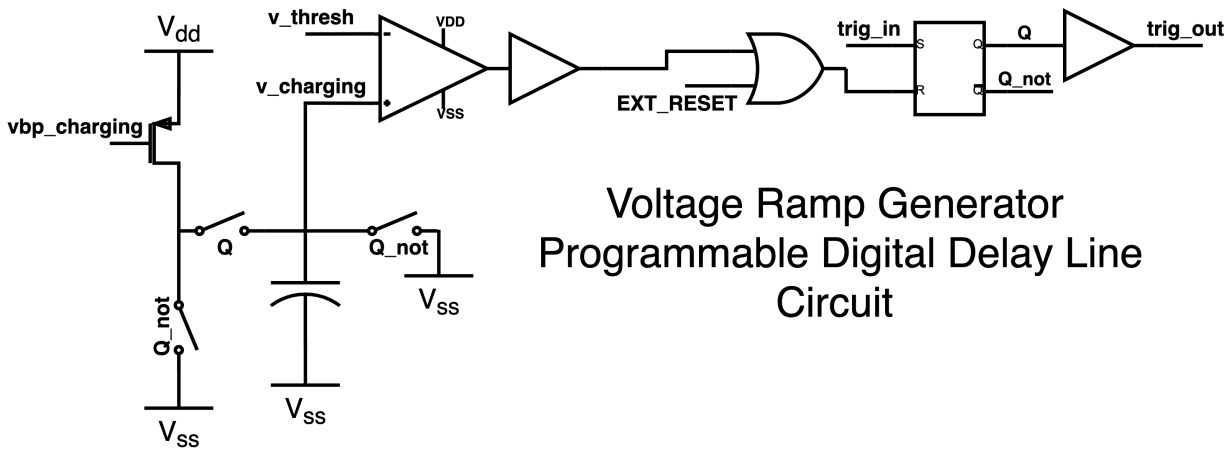

Figure 6.8: A new programmable delay line circuit implemented in PSD_CHIP_V2. This design at the base level is more complicated compared to CSI-based delay lines. However, the overall power and area savings are significant. Details are found in the text.

Figure 6.8 shows a circuit diagram for the voltage ramp generator (VRG) programmable delay line. The input to the circuit is *trig_in*. It is directly connected to the S input of an SR latch. The output, *trig_out*, is a buffered version of the SR latch Q output. Both the Q and $\bar{Q}$ outputs control respective switches shown in the diagram. Like in the case of the

CSI, *vbp_charging* is a gate voltage for the PMOS transistor, which is directly linked to a specified bias current that flows through the transistor. To understand the operation of this circuit, we start by defining the quiescent state. In this state, *trig_in* is LOW and thus, Q (and *trig_out*) is also LOW and $\bar{Q}$ is HIGH. When Q is LOW, the associated switch is open. Similarly, when $\bar{Q}$ is HIGH, the associated switches are closed. Thus, overall, in the quiescent state, the bias current of the PMOS transistor is sunk directly into VSS and the capacitor is not being charged (in fact, both sides of the capacitor are shorted to VSS). This prevents any charge build up on the capacitor during the quiescent state. When *trig_in* goes HIGH, Q also goes HIGH and $\bar{Q}$ goes LOW. As a result, the bias current of the transistor starts to linearly charge the capacitor.

Voltage on a capacitor in general is defined as

$$V = \frac{Q_{cap}}{C}. \tag{6.3}$$

with Q_{cap} the charge on the capacitor and C its capacitance. Taking the derivative of this equation gives a relationship for the change in capacitor voltage as a function of i_{bias}, the bias/charging current:

$$\frac{dV}{dt} = \frac{i_{bias}}{C}, \text{ with } i_{bias} = \frac{dQ_{cap}}{dt}. \tag{6.4}$$

The final piece to the puzzle for this circuit is the reset mechanism to return it to the quiescent state. The voltage of the capacitor is connected to the non-inverting input of a

(PMOS input) discriminator. The other input is a threshold voltage, v_thresh. This voltage, like the other discriminators discussed so far in the chip design, is also highly tunable. The specified threshold directly sets the width of $trig\_out$. This occurs because when the capacitor voltage is >= threshold, the output of the discriminator triggers HIGH. This causes the SR latch's R input to go HIGH, which returns Q LOW and $\bar{Q}$ HIGH (i.e. the quiescent state). A simulation of the circuit is shown in Figure 6.9. Trace names in the figure are directly named after the associated nodes in the circuit diagram.

Overall, $trig\_out$ is a pulse whose width is set by the value of the bias/charging current, which sets $vbp\_charging$ through a corresponding current mirror biasing block, the value of the capacitor and the specified threshold of the discriminator. For the FOUT sub-chain implementations of the programmable VRG digital delay line in V2, appropriate capacitor values are selected and the charging current and threshold is made tunable through respective 5-bit current and 8-bit voltage DACs. This results in 256 **linear** steps to adjust the delay line width via the threshold along with 32 steps to adjust the charging current, i_{bias}. We end up with an effective 8192 (13-bit equivalent) combinations using the two. The 5-bit current DAC for the VRG circuit has an LSB of 100 nA and the 8-bit voltage DAC has an LSB of 3.5 mV. For the HOLD/total integration enable implementations, a capacitor value of 4 pF is used. The resulting full theoretical range of tunability is **4.5 ns to 36 us**. For the partial integration enable implementation, the capacitor value is 1 pF and the associated full theoretical range of tunability is **1.1 ns to 9 us**. Of course, these ranges are not covered uniformly. To get to the maximum delays in total/HOLD and partial integration versions,

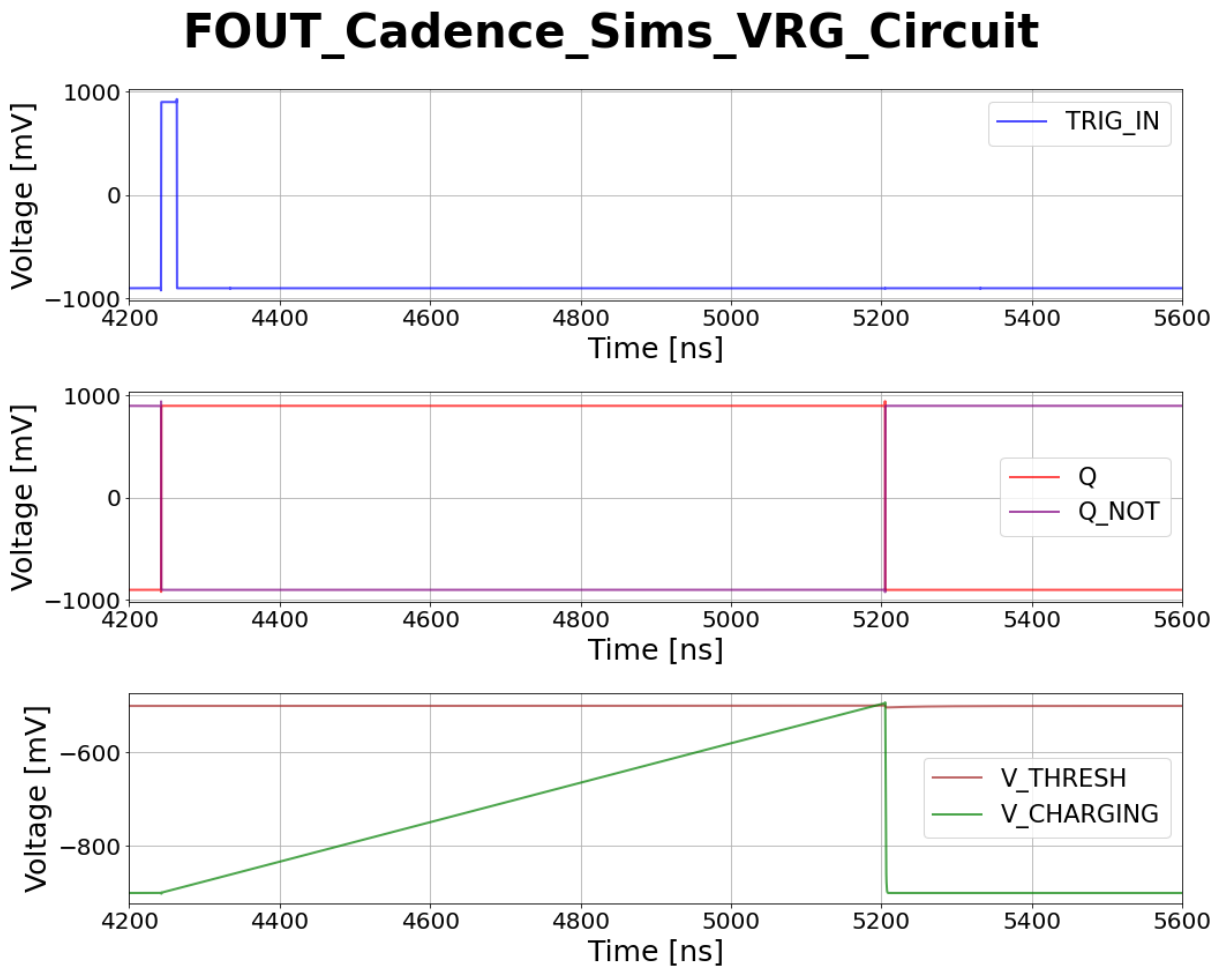

Figure 6.9: A Cadence simulation of the VRG circuit. The different traces represent nodes in the VRG circuit (see Figure 6.8). See text for a thorough explanation.

the steps are quite coarse ($\sim$ 140 and 35 ns, respectively). At the minimum delays, the step size is very fine ($\sim$ 4.5 and 1.1 ns, respectively). More details can be found in Table B.9.

6.1.5.3 Comparison Between the Two Methods

We now conclude our discussion of the programmable digital delay lines with a comparison of pros and cons between the two methods used in PSD_CHIP. The original CSI-based delay line is heavily dependent on the number of CSIs in series to provide large delays. While the circuit is practical to generate delays up to $\sim$ a few 100 ns, the area requirements for longer delays become impractical especially for a final chip with an increased channel count. Physical area is expensive and in IC design, the goal is to optimize area usage even in complicated and high channel count chips. Already, for this reason, keeping the CSI-based design when scaling to 8, 16 or 32 channels would be unfeasible. Another issue is the power draw during switching. We are limiting the current for each of these inverters to a range of O(1 uA) to O(10 uA). This doesn't seem like an issue but considering that this is a per CSI current draw, we have to scale by the total number of CSIs in the delay line (and across all channels and number of implementations per channel). The resulting power draw implications also make this design unfeasible for high channel count and low power applications. The final issue is one of layout parasitics. This is more subtle. In general, the input and output capacitive load of each inverter will add to propagation delay (in addition to whats shown in Figure 6.7). While this delay is a fraction of the delay from starving the inverters of current, having a very long series of CSIs means that there will be a large increase in the observed delay for a specified current bias and number of delay blocks used. While this issue can in theory be mitigated by increasing the current bias, that also increases

power consumption. Layout parasitics will set a lower bound on minimum delay achievable in a real-life implementation of the CSI-based delay line and if we are not careful, this can end up higher than the minimum delay required of an implementation.

The VRG delay line combats all three of these shortcomings of the CSI based design. Power draw is minimal as charging currents are only a few uAs and the current bias for the discriminator is O(100uA). While the SR latch and OR gates used in the design will have high transient current draws of several hundred uAs, average power consumption is minimal. The area to layout the circuit, because there is no repeated part to the design (unlike the series of CSIs), is also a fraction of the CSI delay line. For delays longer than a few hundred ns, the main area consumer is the layout of the capacitor. Furthermore, parasitic effects are also minimal as the effect is constrained mostly to input and output capacitive loads on the single SR latch used in the design. Small propagation delays are expected through the OR gate and SR latch but it is not expected to be an issue. With the added benefit of a dual stage tunability mechanism and linear fine-adjustment steps, the VRG design is an improvement to the CSI in every sense and it is required to be used in a final higher channel count PSD_CHIP design. We have now finished with an overview of the single channel FOUT sub-chain and move on to the SOUT sub-chain.

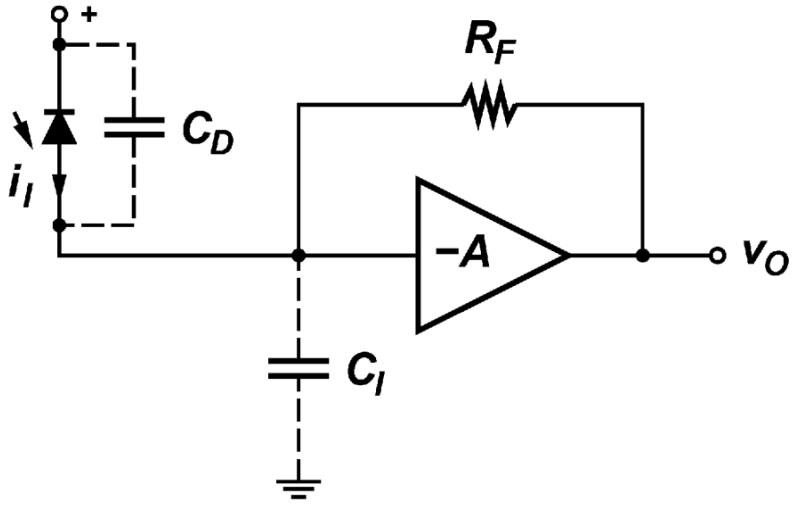

Figure 6.10: A simplified block diagram of a transimpedance amplifier, from [88]. C_D represents the photodiode capacitance, i_I the photodiode current, C_I the input capacitance of the amplifier and R_F a feedback resistor that sets the overall gain of the transimpedance amplifier. v_O is the output node for the amplifier.

6.1.6 Custom SOUT Front-end Design

6.1.6.1 The Issue of Input Capacitance

When selecting a front-end amplifier topology for the SOUT input, the two characteristics to properly account for are the relatively large gain of the SensL SiPM (hence, a low gain requirement for the front-end) and the high effective input capacitance due to the SiPM's junction capacitance (approx. 4.14nF for a J-series SensL 6 mm, 35 um microcell SiPM). Combined, these two factors prevent the easy adoption of the standard transimpedance amplifier (TIA) (see Figure 6.10 for a simplified base circuit design) used for photodiode current readout and amplification for use on a chip. In this circuit, the feedback resistor, R_F, sets

the transimpedance (current to voltage) gain. The input impedance to the circuit scales as $R_F/(1+A)$ with A the open loop gain of the specific op-amp design. A detailed discussion of the limitations of TIAs for SiPM readout can be found in [88] and [89]. To summarize, the main problem is that the input capacitance of the photosensor (SiPM) introduces another pole in the TIA's feedback loop. This results in potentially severe stability issues. Typically, an additional feedback capacitor, C_F, in parallel to R_F is required to ensure stability. The exact value depends on the values of the input and parasitic capacitances as well as the feedback resistor. In practice, for board level designs, the exact value of the required feedback capacitor cannot be determined through first order simulations as there is a level of uncertainty in the true amount of parasitic capacitance of the feedback loop. Trial and error in switching out feedback capacitors for a fabricated board design allows for finding optimal feedback capacitor values. For ASICs, careful design of the op amp used for a TIA implementation is required to prevent these stability issues. Iterative design and (parasitics included) simulations of the front-end and input and output (capacitive and resistive) loads will be needed to converge to a final implementation. The stability of the final design will be strongly influenced by required front-end gain and input capacitance of the SiPM and it is not guaranteed. The final bandwidth of the TIA depends on R_F and C_D.

6.1.6.2 The Regulated Common Gate Amplifier

A commonly used alternative for SiPM readout in front-end IC designs is the regulated common gate (RCG) amplifier (see [89] and [90]). The amplifier topology is shown in Figure

6.11. The input is I_{in} and the SiPM input capacitance is modelled as C_d. I_{BIAS} and I_{BIAS2} provide the required current biases for M_1 and M_6. M_4 is a feedback loop common source amplifier. The overall current gain of this stage can be approximated as unity (it is a current buffer). The current to voltage conversion is set by R_L. Thus, the output voltage is

$$V_{out} = -I_{in}R_L. \tag{6.5}$$

The magnitude of the transfer function is identical to the TIA (the RCG is inverting). In addition, unlike the TIA, there is no direct feedback loop connecting the input and output nodes. Without considering the regulation loop feedback amplifier, M_4, the RCG input impedance can be approximated by the inverse of the transconductance of M_1 (g_{m1}). With the inclusion of M_4, the input impedance is reduced through gain-boosting to $1/(g_{m1}(1+A_o))$ with A_o the gain of M_4. This is a direct analog to the input impedance of the TIA. With the RCG, we have an amplifier topology where we achieve can achieve similar transimpedance gain and input impedance without the direct feedback loop between the input and output nodes. Thus, the RCG amplifier offers an excellent alternative to dealing with the large input capacitance of SiPMs.

Unfortunately, just like the TIA, we cannot escape the dependence of the amplifier bandwidth on the input capacitance. In general,

$$\frac{V_{out}}{I_{in}} = \frac{R_L}{\left[1 + \frac{s(C_d + AC_{gs1})}{g_{m1}A}\right][(1 + sR_LC_L)]} \tag{6.6}$$

can be used to provide a simplified approximation for the RCG transfer function. We can clearly see the effect of C_d in this equation. For the same amplifier design, bandwidth is directly constrained by the input capacitance. Thus, the design of an RCG amplifier must be done with the expected input capacitance accounted for. Figure 6.12 shows gain-frequency Cadence simulation outputs for the RCG implemented in PSD_CHIP_V2 using various input capacitance values.

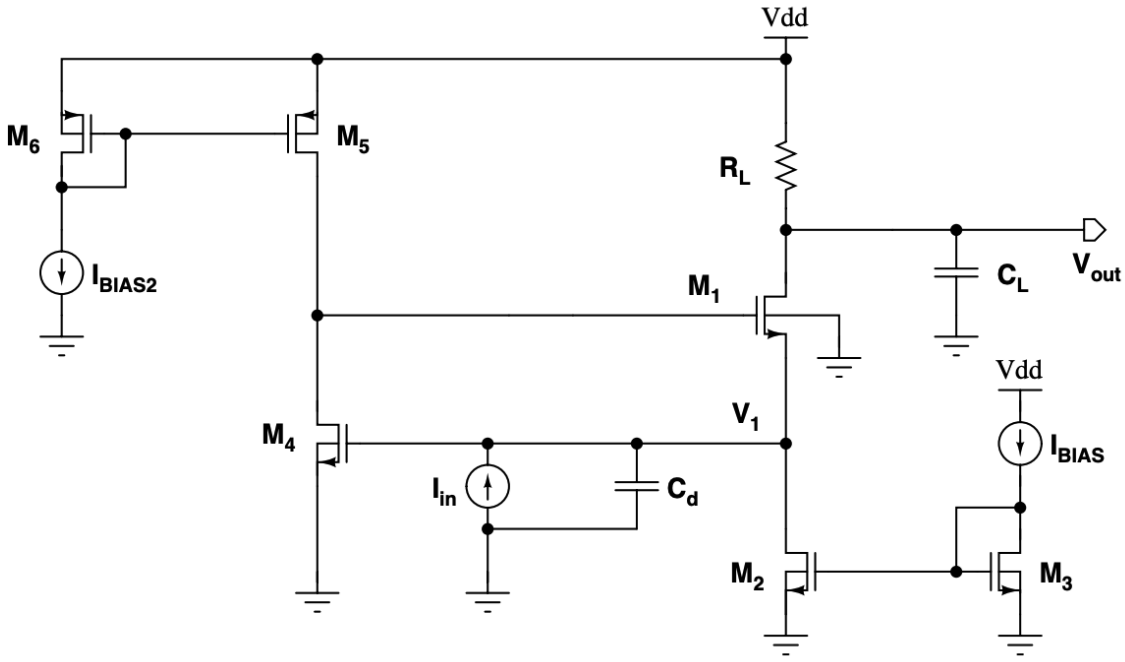

Figure 6.11: A transistor level breakdown of the regulated common gate (RCG) amplifier, from [87]. This circuit is used for the SOUT front-end. Included in this diagram are the current mirror biasing and main capacitive elements at the input and output nodes that affect bandwidth and performance for this circuit.

A subtle feature of the RCG is its asymmetric source/sink capabilities. In general, max current sourced from (sunk into) an NMOS (PMOS) input is limited by (W/L) dimensions of

the MOSFET and max current sunk into (sourced from) NMOS (PMOS) input is constrained by its selected bias current value. In situations where the peak current amplitudes are expected to be small, it is possible to pick either an NMOS or PMOS input stage for the RCG. However, in the case of SiPMs, due to the range of energy depositions expected for a neutron scatter camera, peak current amplitudes can be quite large. While this issue was not fully appreciated in V1, for V2 we designed a dual input polarity capable RCG amplifier. In order to provide dual-polarity capability, a merged NMOS+PMOS input stage design is implemented.

A final note about RCGs is another subtle feature. The input impedance is frequency dependent. The gain-boosting (impedance lowering) effects of the feedback common source amplifier are diminished at higher frequencies and in general can approach $1/g_m$ in the high frequency limit. Thus, there is a subtle fundamental limit on rise-time achievable for the sharp rise and slow decay SiPM pulses that we amplify using the RCG.

In V2, we have implemented programmable gain in the SOUT front-end as mentioned in the high level overview. Symmetrical operational transconductance (OTA) amplifiers with tunable feedback resistors was used to implement this feature. The full range of programmable gain in V2 depends on whether the NMOS or PMOS input stage is active (see the datasheet for further details). A range from unity up to 6x for an active NMOS input and up to 5x for an active PMOS input is achievable. The difference in maximum gain comes from the choice of an inverting or non-inverting symmetric OTA topology for the second stage. Because of the inverting nature of the RCG, these two setups are required to ensure

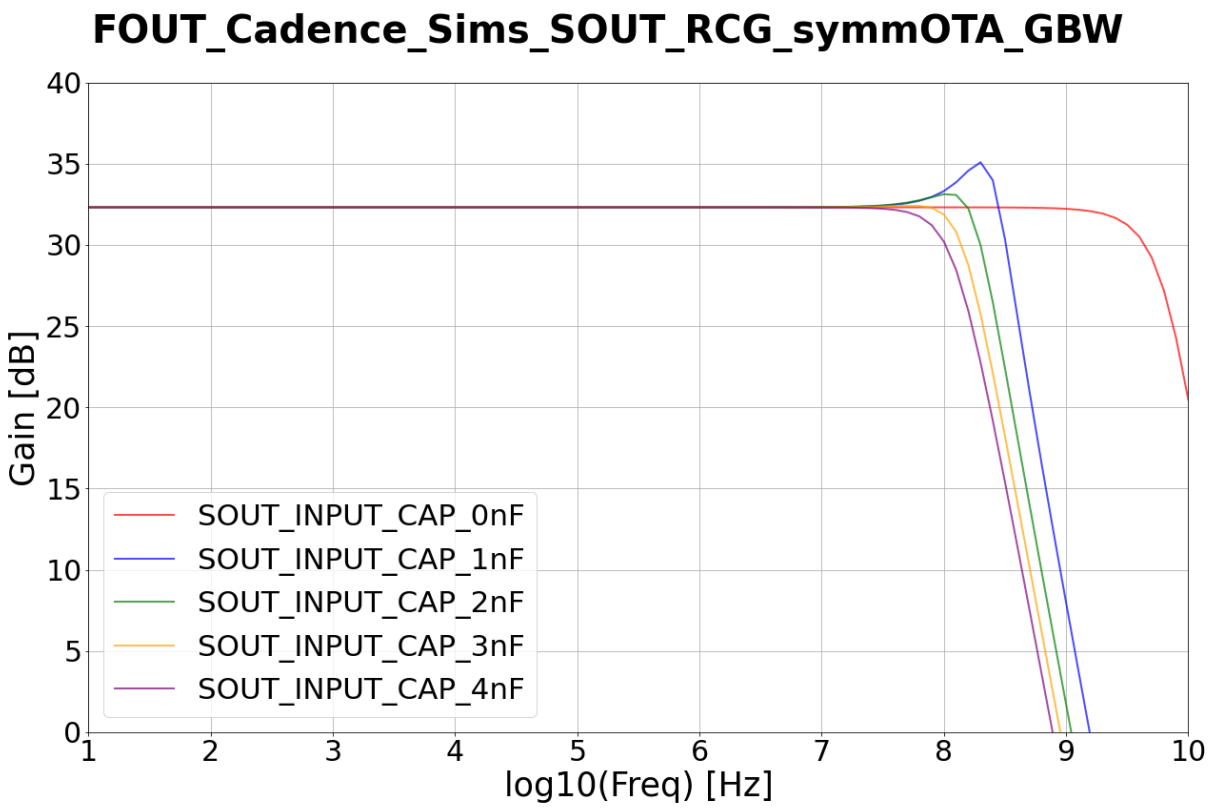

Figure 6.12: Gain-bandwidth (GBW) plots for various input capacitances is shown. The symmetrical OTA is not included in these simulations. A clear dependence between input capacitance and RCG bandwidth is observed.

a positive output voltage polarity in both scenarios.

In general, the approximate transfer functions defined above only provide a first order estimate for the behavior of deep sub-micron RCG implementations. While these equations can help with general trends between different design parameters, Cadence simulations with detailed transistor models from the foundry must be used to inform final design parameters. A sample simulation of the RCG showing linearity across a broad range of input square

wave test pulses is shown in Figure 6.13. The first row shows the input test current pulses. The second row shows the output of the RCG + symmetric OTA front-end stage for these inputs. The final row shows the pulse outputs after the on-chip buffer before it is sent off-chip through a dedicated pad (see Figure 6.24).

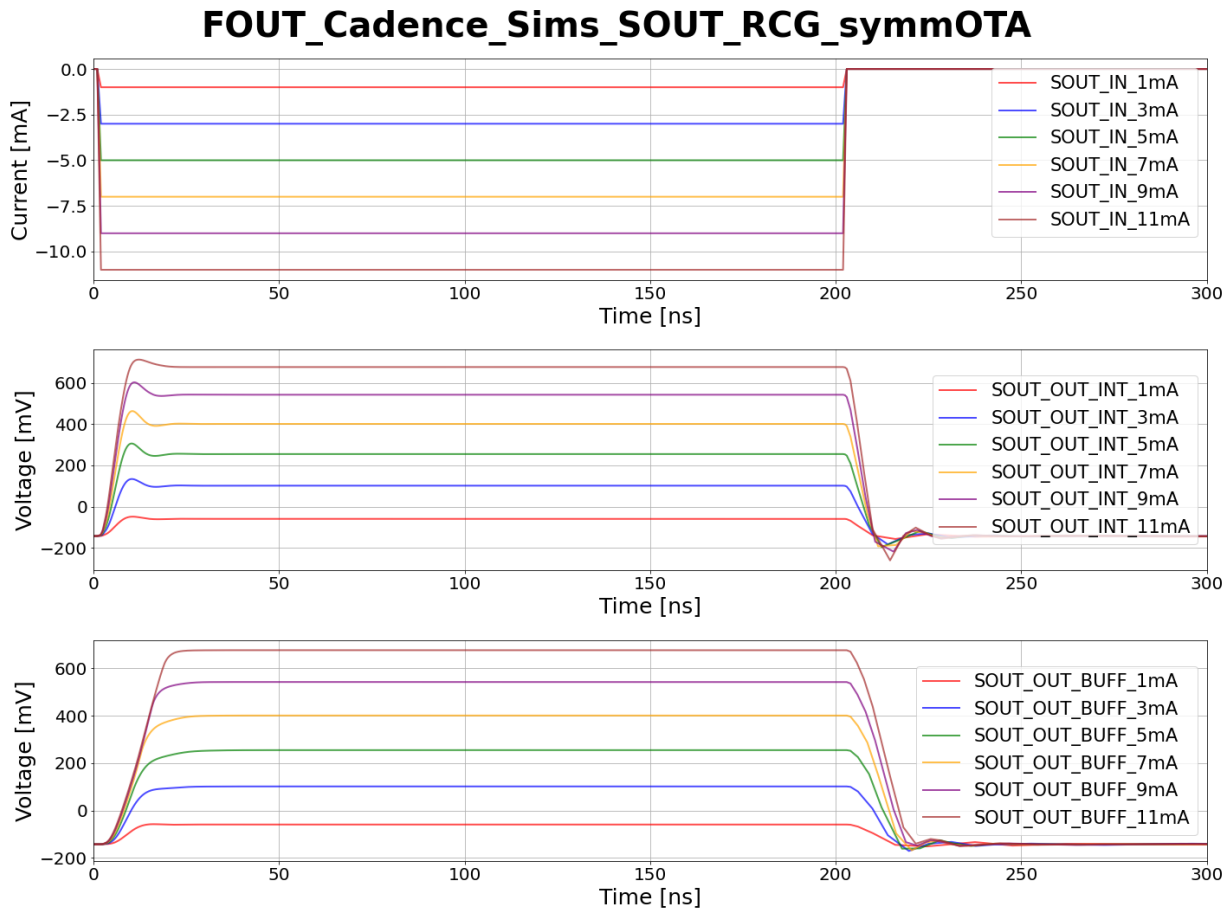

Figure 6.13: A square wave input waveform with varying current amplitudes is used as a test input to the SOUT front-end (first row). The output waveform from the RCG + symmetrical OTA front-end (with the gain of the OTA set to unity) is shown in the second row. In the third row, the waveform post-buffer is shown.

6.1.7 Parallel Partial and Total Integration Stages

After the RCG + symmetrical OTA stage front-end, the now amplified voltage SOUT waveform is split into parallel partial and total integration stages. These stages provide two of the three outputs needed for a neutron scatter camera: total energy deposited and the core of the real-time PSD implementation on-chip (refer back to 5.4 for further details).

6.1.7.1 A Fully Differential Total Integration Stage

Figure 6.14 shows the circuit diagram for the total integration stage. On top of the feedback capacitors and (tunable) resistors, feedback switches are included to reset the stage for a subsequent waveform integration. Both the partial and total integration stages were implemented as fully differential to mitigate against switching noise and layout parasitics. This is accomplished through common-mode noise rejection (i.e. switching noise). This helps to prevent spurious switching in the final SOUT discriminator, which classifies an incident particle as fast neutron or gamma/background based on single threshold operation.

In Chapter 5, we had wrote out the explicit transfer function for the ideal integration stage in Eqn. 5.18. This equation does not reflect reality in the slightest. For a real-life implementation of the integration stage, there are several sources of non-ideal behavior. The first is the effect of a non-zero baseline offset. If the baselines between the differential op amp inputs are non-zero, this difference will be included in the integration. In the world of chip design, 0 V baselines are uncommon. The specialized supporting circuitry required to allow

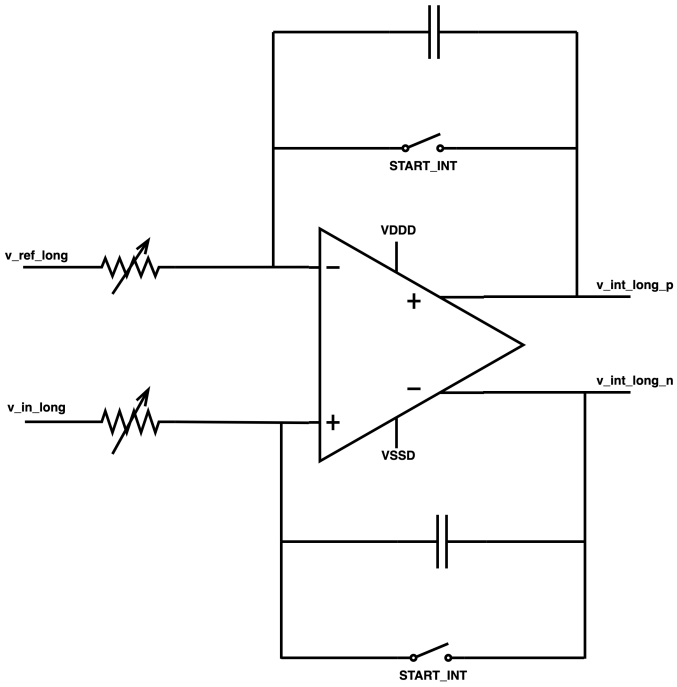

Figure 6.14: A fully differential total integration stage implementation. The tunable resistor, as described in Chapter 5, is implemented in this stage and forms a core part of the real-time PSD design.

for 0 V baselines to an op amp with dual-polarity power supplies is not usually included for internal chip stages without good reason. In most cases, we do expect DC baselines to be roughly at the midpoint between the power and ground rails used (but not exactly). The input common mode range (ICMR) and output common mode range (OCMR) is an important specification of op amp design and explicitly states the expected baseline values.

In general, we need to match or account for the baseline reference voltage. For PSD_CHIP, we use an 11-bit equivalent voltage DAC + buffer (to provide sufficient current load capability) set up to tune the reference voltage (v_ref_long) to match the observed baseline of the

output from the two-stage SOUT front-end system. Further details are found in Tables A.5 and B.5. Over full scale (power to ground; 1.8 V), we obtain sub-mV resolution in matching the DC baseline.

The second non-ideal behavior is that integration stages do not retain integrated charge in perpetuity. In reality, we only require that charge be held until discharging/resetting the capacitors. However, even this requires optimization of the explicit resistor and capacitor values through simulations to understand discharge/leakage through parasitic paths from the exact implementation of the integration stage (including modeling input and output loads). For V2, the end result was using 75 kΩ as the base resistor value (16 values ranging from 75 kΩ to 250 kΩ are actually selectable) and 2 pF for the feedback capacitors.

The tunable resistor is meant to set a threshold PSD ratio as discussed in detail in the previous chapter and in the high level overview of this chapter. This ratio is the ratio of the respective partial and total integration resistors. The partial integration stage resistor, as we will see, features a fixed 75 kΩ resistor. Thus, a range of threshold ratios are covered between 0.3 and 1.0. The lower bound was selected to match feasibility study results. The upper bound was kept at 1.0 to ensure the ability to perform a non-scaled partial and total integration of an incident waveform. The timing diagram for the stage is found in Figure 6.16. It will be discussed jointly with the partial integration stage timing later. Further details on the exact ratios implemented are found in Sections A.8 and B.9.

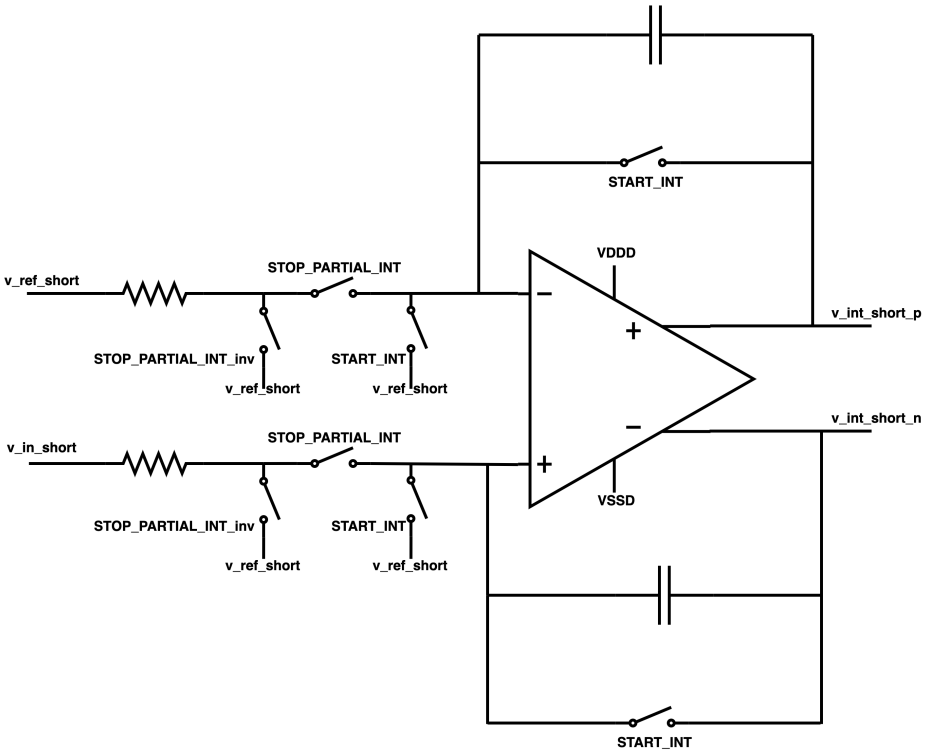

Figure 6.15: A fully differential partial integration stage implementation. This stage contains only fixed resistors and several series and feedback switches. The timing diagram for these switches is in Figure 6.16.

6.1.7.2 A Fully Differential Partial Integration Stage

The partial integration stage, at its core, is a modified version of the total integration stage. Figure 6.15 shows the circuit diagram. At the channel level, the baseline reference, v_ref_short is the same as v_ref_long. Although the resistor values are fixed at 75 kΩ, the feedback capacitors are still 2 pF. The primary change is the increased number of switches used to actually implement a partial integration and synchronous reset with the total integration stage. The timing diagrams for the various switches are discussed in tandem

with the total integration stage logic in the following subsection.

6.1.7.3 Integration Enable Pulse Logic

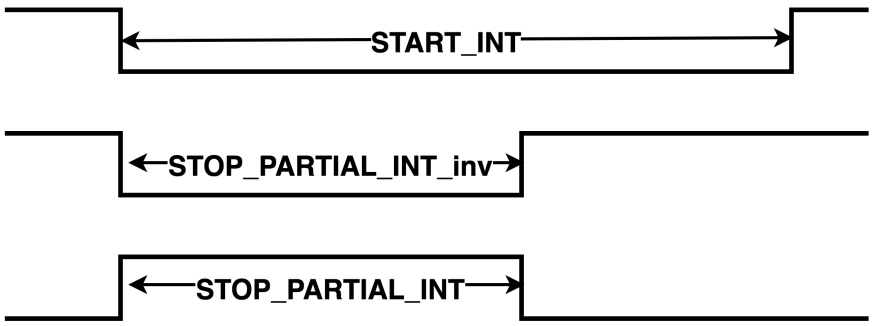

Figure 6.16: A timing diagram for the various switches used in the parallel partial and total integration stages. The actual partial and total integration enable pulses are generated in the FOUT signal chain using either a current starved inverter (CSI) or voltage ramp generator (VRG) based design. They are then transformed in a manually implemented digital block into the versions seen in this diagram. The actual STOP_PARTIAL_INT digital testbus output is a "delayed" START_INT pulse, where the delay in going low, compared to START_INT, represents the width of STOP_PARTIAL_INT shown in this timing diagram. Further details are found in the appendices.

The timing diagram for both partial and total integration stages is shown in Figure 6.16. The one-shot pulse generator circuits discussed as part of the FOUT sub-chain provide a straight-forward way to generate programmable width enable pulses for the partial and total integration stages. However, these pulses are not used independently for the parallel integration stages. Overall, we need to ensure a synchronous start, stop and reset cycle for every parallel integration performed. With an internal clock, this is more straight-forward to implement. However, the chip features an asynchronous design, so we had to be a little more

creative. A manually implemented digital block transforms the FOUT sub-chain generated pulses into the versions seen in the timing diagram. Further details are in the appendices.

For the total integration stage, the feedback switches are controlled by the **START_INT** line. The quiescent state for these switches is that they are closed (**START_INT** is HIGH). Thus, the capacitors are being continually discharged when not integrating an incident waveform above FOUT threshold. This is a direct analog to the design of the VRG delay circuit. The benefit of this is that no charge (voltage) build-up occurs. Explicitly, the (FOUT) total integration enable line is inverted to form the **START_INT** line shown in the timing diagram. When an FOUT trigger starts integration, this causes the feedback switches to open (**START_INT** goes LOW) for the duration of the selected enable pulse width. The same logic applies for the partial integration stage feedback switches.

The other partial integration switches in series to the op amp inputs are used to properly implement a stop to the partial integration and a synchronous integration reset with the total integration stage. In the quiescent state, **STOP_PARTIAL_INT** is LOW (the switch is open). **STOP_PARTIAL_INT_inv** is the inverse and it is HIGH (switch is closed). This ensures that the input to the stage is shorted to v_ref_short by default for both inputs. To prevent a floating input to the op amp during this period (and to provide a discharge path for the capacitor), the third switch, also controlled by **START_INT** is separately connected to v_ref_short. When an integration starts, the middle switch is closed (**STOP_PARTIAL_INT** goes HIGH), allowing the input signal to be integrated on the feedback capacitors for the total width of the partial integration enable line from

FOUT. While partial integration is complete but total integration is ongoing, the middle switch opens (**STOP_PARTIAL_INT** goes back LOW) and the left switch closes (**STOP_PARTIAL_INT_inv** goes HIGH). Thus, until total integration is finished, the partial integration input waveform is sent straight to v_ref_short while the charge (voltage) on the capacitor is held. When total integration has finished, the capacitor is allowed to discharge (**START_INT** returns HIGH). At this point, the circuit is reset back to its quiescent state and then waits for another integration trigger.

6.1.8 A Fully Differential Subtraction Stage

The outputs of the parallel integration stages are sent to a fully differential subtraction stage as discussed in the high level overview. Explicitly, the fully differential signal path (for each SOUT sub-chain) is provided from the parallel integration stages to right before the final SOUT discriminator. The subtraction stage is a standard implementation of a fully differential subtraction op amp circuit with the exception that the feedback resistors are tunable. This allows a gain from unity to x15 to be selected for the stage to provide flexibility with different ranges of expected energy depositions (depending on operating environment). More details can be found by referring back to the feasibility study in Chapter 5. Explicitly, the possible gain settings are: x1, x2, x3, x5, x10, x15. As discussed in both the feasibility study and the high level overview of this chapter, for a specified total integration resistor value (PSD threshold ratio), the output of this stage quantifies the difference from that value

for an incident waveform.

6.1.9 Final Classification Discriminator

The final block in the SOUT sub-chain is the SOUT discriminator used to classify incident particles as fast neutrons or gamma/background. Prior to this stage, the subtraction stage output is converted from differential to single-ended. This single-ended difference is input to the discriminator. This discriminator, like all others in the chip design features a highly tunable threshold. As seen in the feasibility study, this threshold is the equivalent of setting a threshold in the y-axis of Figures 5.31 and 5.30 between the neutron and gamma bands in subtraction space. The output of this stage is a single PSD bit that goes HIGH for neutrons and remains low for gammas. This signal is sent off-chip as an LVDS differential output per channel. See Figure 6.24 for the explicit corresponding pads. It is possible to select a standard width for this PSD bit (details can be found in the appendices). Furthermore, another manually constructed digital block ensures that the output of this stage can only trigger HIGH during the latter half of an active parallel integration to prevent early on spurious false triggers in the integration cycle.

6.1.10 Full Channel Simulation

We have the ability to select "raw" waveforms corresponding to different energy depositions for both EJ276 and Stilbene (acquired from the test bed) and use them as inputs

into full channel simulations of the implemented PSD_CHIP design. The plots presented in the feasibility study (refer back to Chapter 5) are shown in the appropriate (mV) units for direct comparison to the chip's outputs. In this subsection, we present results of such simulations with all layout parasitics included. Figure 6.17 shows a sample output for representative 500keV$_{ee}$ gamma and fast neutron depositions in EJ276. The first row shows input current pulse waveforms for each particle. The second row shows the output of the two-stage SOUT front-end. The third row shows actually generated **STOP_PARTIA_INT** and **START_INT** pulses (refer back to the timing diagram in Figure 6.16). The fourth row shows the outputs of the partial and (tuned) total integration stages for the incident waveforms. The fifth row shows the output of the subtraction stage and the set final SOUT discriminator threshold. Finally, the sixth row shows the standard width HIGH pulse for the fast neutron waveform and not for the gamma.

Particle	Energy	EJ276	Stilbene
Neutron	100	-17.24mV	-18.06mV
Neutron	150	-1.43mV	6.89mV
Neutron	500	36.02mV	67.93mV
Gamma	100	-23.79mV	-26.33mV
Gamma	150	-24.64mV	-29.55mV
Gamma	500	-30.10mV	-57.79mV

Table 6.2: A table of the subtraction stage outputs for a range of energy depositions of fast neutrons and gammas for Stilbene and EJ276.

We can repeat these simulations for a range of corresponding deposition energies for both EJ276 and Stilbene. The resulting subtraction stage output values are tabulated in

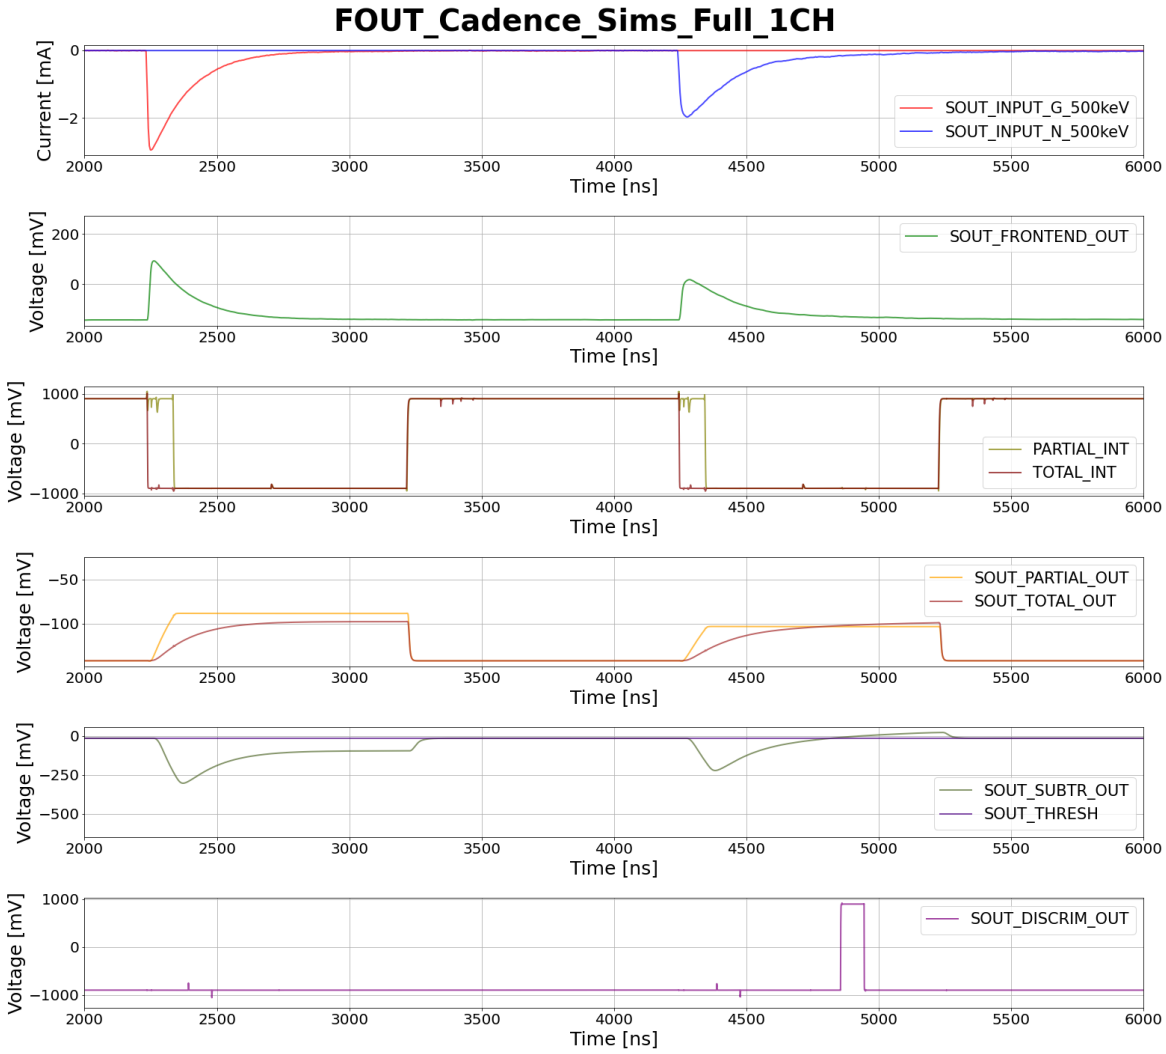

Figure 6.17: A full one channel simulation output (run with all layout parasitics included using Cadence). The outputs of the main SOUT sub-chain stages are shown along with the relevant partial and total integration enable lines. Further details on all stages are found in the text.

Table 6.2. Figure 6.18 represents the table visually. The absolute difference between the fast neutron and gamma subtraction outputs for each considered energy deposition is taken and plotted as a function of energy for both scintillators.

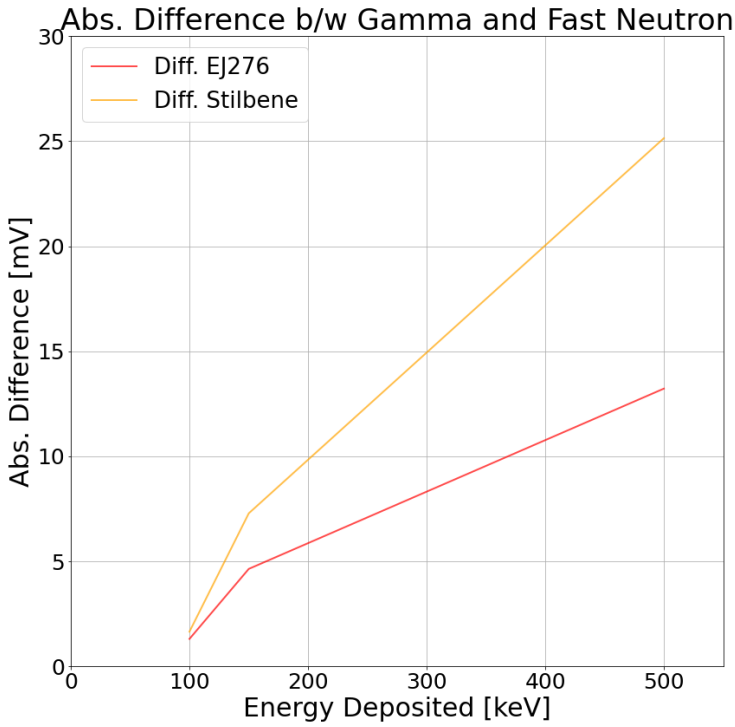

Figure 6.18: A plot of absolute differences between the fast neutron and gamma subtraction output values at various energy depositions is shown for both scintillators.

There are some caveats that prevent these results from being directly compared to the feasibility study results. The SOUT input waveforms (even with the transimpedance gain of the fanout boards divided out to convert to units of current) contain shaping effects of the 50 Ω resistor on the fanout board. Both SOUT and FOUT waveforms are processed through the 125 MHz front-end unity-gain buffer (plus a subsequent anti-aliasing filter) on the CAEN DT5725 digitizer. The input signals for the Cadence simulations are, as a result, not an exact reflection of the true inputs directly from the SensL J-series SiPM coupled outputs. In addition, the feasibility study made no attempt to model bandwidth (signal shaping) in

PSD_CHIP's one channel signal chain. This affects the final separation and widths of bands in both PSD and subtraction space. Nevertheless, we are able to verify that the implemented PSD_CHIP 1CH signal chain follows the expected trends.

6.2 Full Chip Integration

In discussing the single channel design of PSD_CHIP, we have made several references to the integrated tunability and programmability of several parts of the design (i.e. tunable resistors, programmable gain, reference voltages, thresholds, partial/total integration pulse widths, etc.). In this section, we cover how this was integrated into the full chip design. We also discuss considerations in scaling from one to four channels. This resulted in including several debugging tools such as digital test buses, voltage and current monitors, and isolated analog and digital power/ground rails. Scaling to a full chip design also requires a well thought out proper biasing framework for the chip. All of these are discussed generally in dedicated sections below. Concrete values and further details are provided in the appendices.

6.2.1 The Digital Core

Much of the integrated programmability on-chip is controlled through a synthesized digital core (written in SystemVerilog). Communication with this digital block is controlled through a custom clocked UART implementation. UART (Universal Asynchronous Receiver/Transmitter) is a standard serial digital communication protocol between chips, which

features customizable data formats and variable speed, non-clocked transfer of data. The digital core is composed of 62 8-bit registers in V2 (40 in V1). Each bit of each register directly acts as control bits for various switches, tunable resistors and current/voltage DACs integrated on-chip. Thus, setting the values of register bits will allow for control over various features such channel enable/disable, specific resistor value selection in the various tunable resistors, specific output currents and voltages, etc. The full list of register (and bit) names, groupings and conversion formulas to set appropriate values are given in Tables A.3 and B.3.

6.2.1.1 Communicating with the Digital Core

Bit Range	Contents	Comments
[0]	WRB	1:write, 0:read
[8:1]	Data	8-bit data word
[16:9]	Address	8-bit address
[17]	Parity	0:odd num. of 1s in [16:0] and 1: even num. Ignored for write

Table 6.3: The breakdown of the 18-bit data word structure for the custom UART implementation on PSD_CHIP.

In order to communicate with the digital core, an external microcontroller (such as Raspberry Pi, BeagleBone, etc) or FPGA is required. We have used a Raspberry Pi for the custom test boards designed and fabricated to perform initial verification and characterization of the chip. Because a standard UART protocol implementation does not feature a dedicated CLK line, custom interface software was written in order to integrate the over-sampled clock (and custom reset line) into the UART protocol. The software was written hierarchically

using inherited classes so that individual functions (i.e. adjust a specific threshold voltage, enable/disable channels) are already implemented. Figure 6.19 showcases how data is sent/received. The chip has four lines corresponding to the custom implementation: PISO, POSI, RST_N, CLK. Figure 6.24 can be referred to for locations of the relevant pads in the padring. PISO (primary in, secondary out) and POSI (primary out, secondary in) are the standard UART receive/transmit lines. The chip is always considered secondary and the external microcontroller or FPGA is considered primary. In the quiescent state, PISO, POSI and RST lines are HIGH and the CLK line is held LOW. In order to initiate a read/write request, the CLK line must be producing an active CLK signal. The CLK frequency divided by 16 sets the data transfer rate. The oversampling CLK is used to prevent spurious logic transitions from being interpreted as actual bits of the data word. While the CLK is running, an 18-bit data word can be sent (either a read or write request). The structure of this word is described by Table 6.3. The data transmission starts with a START BIT that brings the POSI line LOW. Afterwards, the 18 bits are sent in serial LSB:MSB format. At the end, a STOP BIT brings the POSI line HIGH again. If the command was a read request, after a delay of $\sim$ two oversampled CLK cycles, a digital word is sent back from the chip in the same format on the PISO line.

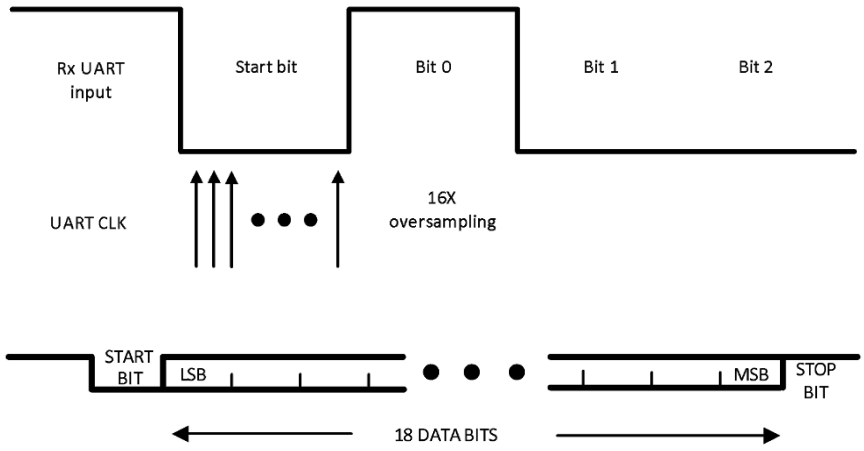

Figure 6.19: A diagram providing details on the custom UART implementation that allows for communication to/from the chip. Above, the details of the oversampled CLK feature is described. Below, a breakdown of how the 18-bit word is sent/received from the chip.

6.2.2 Scaling to Four Channels

Although seemingly simple, scaling from one to four channels comes with its own host of considerations. As the different channels should have identical performances and characteristics (ideally), care needs to be taken in the layout.

6.2.3 Digital Test Buses

Across the four channels, each with its own copy of the FOUT sub-chain, there are tens of potential nodes in the respective signal chains that can be made available for debugging purposes. Unlike on a PCB, we cannot simply probe a trace or IC pin to track/debug signal progression. If we want such capability in a chip, it needs to be integrated into the full chip

design. We used four 32x1 digital MUXes to provide this functionality for PSD_CHIP. We selected eight nodes in the FOUT signal chain to make accessible externally for each channel. Instead of providing the capability to select between the eight nodes for each channel using 8x1 digital MUXes, the 32x1 version also allows for viewing four sequential nodes within a single channel. Thus, we have the overall ability to either access any of the 8 nodes for all four channels or view a subset of the nodes for a single channel. The latter functionality is immensely useful for debugging. In addition, both capabilities are useful for timing and propagation delay studies. In normal operation of the chip, the digital test buses are set to output the T0 (output of the FOUT discriminator) pulse for each channel. The details of mapping register bits to output specific signal nodes can be found in Tables A.7 and B.7.

6.2.4 Voltage/Current Monitors

We discussed integrating "spy nodes" for the FOUT sub-chain into the full chip design above. We also want such a capability for certain nodes in the SOUT signal chain. Unlike the FOUT chain, we already send off the main analog block outputs to dedicated pads (see Figure 6.24). Thus, the requirements for SOUT "spy nodes" are much simpler. We implemented an 8-bit analog MUX that operates on the one-hot principle (only one bit can be HIGH at a time) to serve as an internal voltage monitor for both versions of the chip. The actual voltages that can be externally observed are different between the two versions. For both versions, we selected a subset of the internally generated DC baseline and threshold voltages to ensure

that the internal voltage DACs are working as expected. Details on the exact voltages read out for both versions can be found in Sections A.7.1 and B.7.1.

6.2.5 Two-Stage Global/Coarse Tunability

Due to process variations and layout-related effects, we have implemented a two-stage global and coarse adjust capability for the majority of tunable voltage thresholds and references on the chip. This allows for flexibility in expected minor baseline shifts between each channel's signal chain for respective signal blocks. A combination of 3-bit string and 8-bit R2R voltage DACs allow for full-scale selection of reference and threshold voltages between -0.9 V and +0.9 V. A global 225 mV range that is common to all four channels can be selected with the string DAC and then a fine adjust can be performed within this range per channel with the R2R DAC. This allows for full-scale fine-tuning of sensitive thresholds between channels while limiting redundancy in the DAC setup. Further details on the exact thresholds/references and the formulas for converting between setting register bits in the digital core and the mapping to an exact voltage are found in Sections A.5.1 and B.5.1.

6.2.6 Digital I/O Using LVDS

Other than the digital UART lines that allow for communication with the digital core, a number of other digital I/O pads exist. These pads (external SR latch reset lines, SOUT final discriminator PSD bits, external trigger, and the digital test buses) use the low-voltage

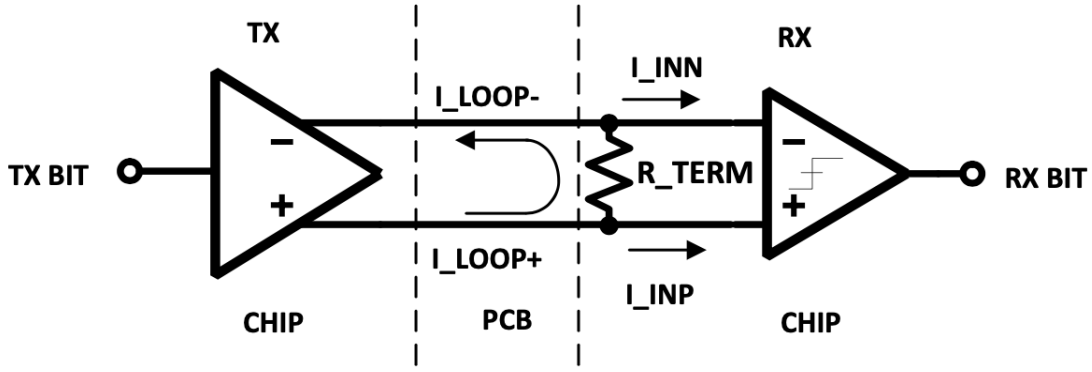

Figure 6.20: A diagram of the LVDS digital I/O circuitry, including the interface between the mated TX and RX components. Diagram courtesy of the Berkeley Lab IC Design Group.

differential signaling (LVDS) standard. A circuit diagram for this standard is provided in Figure 6.20. The transmitted bit is converted to a differential current at the transmitter chip and sent to the receiver chip. At the receiver chip, the differential current is converted to a differential voltage by a termination resistor (which may be internal or external to the receiving chip). In the receiving chip, an amplifier senses the differential voltage and outputs single-ended CMOS logic-levels. on-chip, either the TX or RX circuit is implemented depending on whether the signal is an input or output. A corresponding TX/RX IC or FPGA interface must be used to send or read these signals to/from the chip. Standard LVDS logic levels are CMOS (3.3 V). Thus, dedicated level-shifters shift up/down digital I/O signals as needed.

Specification	Padring Name	Value	Note
Analog Power	VDDA	+0.9V nominal	Only used for front-ends
Analog Ground	VSSA	−0.9V nominal	Only used for front-ends
Digital Power	VDDD	+0.9V nominal	Used for everything else
Digital Ground	VSSD	−0.9V nominal	Used for everything else
LVDS Power	VDD_IO	+2.4V nominal	Only used for digital I/O
LVDS Ground*	VSS_IO/VSSD	−0.9V nominal	VSS_IO is shorted to VSSD

Table 6.4: The various power and ground rails on PSD_CHIP.

6.2.7 Power Rails and Voltage/Current Biasing

In PSD_CHIP, we have used three separate power rails and two (with VSSD and VSS_IO pads in Figure 6.24 shorted internally) ground rails. These are listed in Table 6.4. The analog VDDA/VSSA rails are used exclusively to isolate the SOUT and FOUT front-ends across all four channels from common substrate noise (especially digital noise pickup) for the raw input signals. The rest of the chip is powered on the VDDD/VSSD digital rails. Explicitly, the global p-substrate of the chip is tied to VSSD. Finally, a separate LVDS power rail (VDD_IO) is used for CMOS logic levels in LVDS digital I/O. For V2, first order simulations have shown that we should expect 86 mW analog and 154 mW digital DC power dissipation. This results in a total of 240 mW total power dissipation with 75 mW from biasing circuitry and 41 mW per channel.

Unlike in PCB designs using discrete ICs, we must directly provide appropriate current biases for the various op amps, front-end topologies, buffers, digital delay line circuits, etc on-chip through dedicated biasing modules. Referring to Figures 6.21 and 6.22, the top halves

of the chip layout shown is where these modules are implemented. The overall biasing set-up is quite complicated and involved and will not be covered here. In general, much like for simple ICs, we ideally only want to provide input voltages for the (multiple) power and ground rails to a chip's pads. However, for the prototype, we have included 5 external **RBIAS** pads: **RBIAS_ANALOG**, **RBIAS_PARTIAL**, **RBIAS_FULL**, **RBIAS_RCG**, and **RBIAS_CS**.

Two of these, **RBIAS_PARTIAL** and **RBIAS_FULL**, were included to provide external LSB selection capability for the CSI delay lines as already discussed. The other three directly control current biases on-chip. **RBIAS_RCG** and **RBIAS_CS** respectively set the bias for the SOUT RCG front-ends and FOUT CS front-ends for all four channels. **RBIAS_ANALOG** sets appropriate current biases for everything else. The **RBIAS_ANALOG** current bias is scaled up or down as required for all of the various components used on-chip through a series of current mirrors in the biasing modules. Nominal current biases for all these pads are found in Tables A.4 and B.4.

6.2.8 A General Overview of Layout Considerations

Design and simulation of the full ASIC isn't the end for the IC design procedure. Layout plays a significant role in signal characteristics through parasitics. These parasitics correspond to trace resistances and capacitances from layout choices as well as MOSFET capacitances that can affect signal bandwidth and power distribution. Resistances of traces can cause issues such as IR drop and ground bounce. Parasitic capacitances can cause unwanted ringing

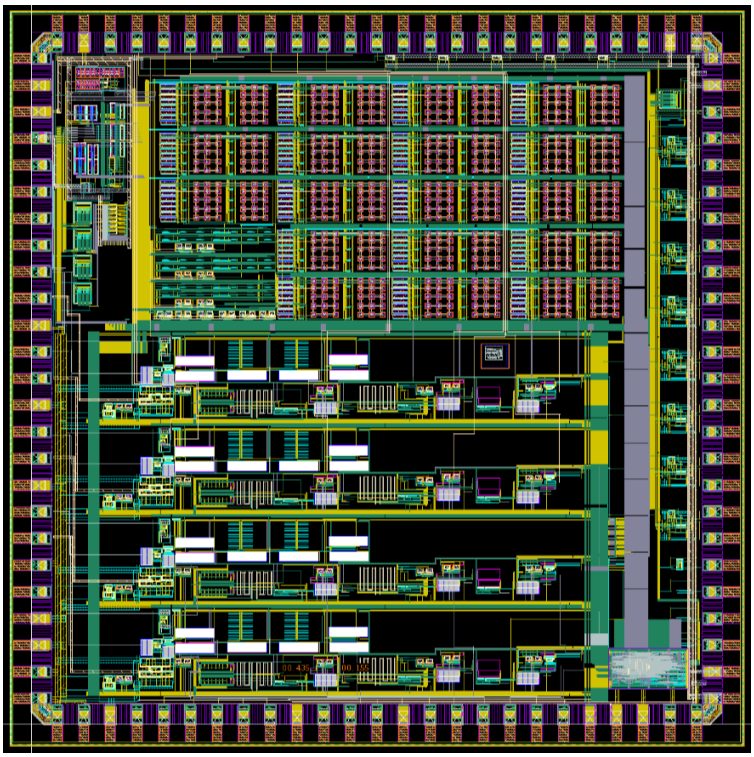

Figure 6.21: A picture of the layout of PSD_CHIP_V1. Each color in this figure represents a different metal or transistor layer in the overall chip design. Black represents an absence of any layer. A thorough walk through of the mapping of this layout to the system level overview is given in the text and not repeated here.

and oscillation from shifting (and creating new) poles and zeros. Electromigration from improperly set widths for power busses can cause the slow and eventual degradation of these traces. For layout of resistors and capacitors, process variations can cause gradients in final wafer that can cause, among other effects, the sheet resistance across the chip's surface area to vary. For sensitive parts of the chip design, the layout of these components have to be optimized to the extent possible against these effects. For op amp topologies with differential pair inputs on-chip, matching of input transistors through the use of special layout geometries

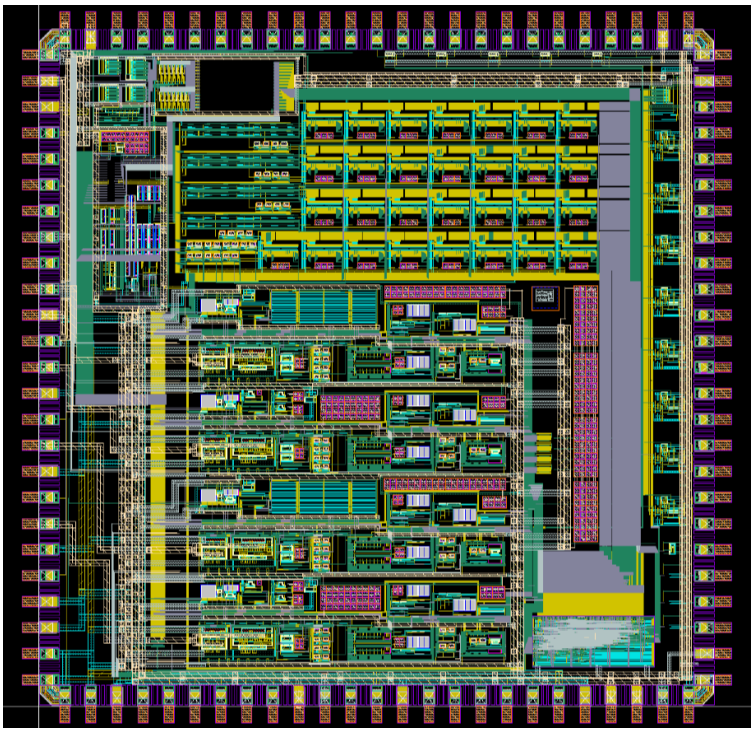

Figure 6.22: A picture of the layout of PSD_CHIP_V2. Each color in this figure represents a different metal or transistor layer in the overall chip design. Black represents an absence of any layer. A thorough walk through of the mapping of this layout to the system level overview is given in the text and not repeated here.

is required in order to ensure general characteristics such as offset, gain, bandwidth, etc are maintained. Specialized parasitic extraction and simulation software for IC design helps keep all these tasks manageable to a degree. We shall not go into detail about layout in this work and will only offer one specific example of optimization in the SOUT and FOUT front-end layout. They are placed in deep n-wells with surrounding guard rings to allow for decoupling of the NMOS transistors' p-channels from the global p-substrate and noise isolation. This is what allows for a separate and isolated analog power supply for the front-ends.

Figures 6.21 and 6.22 showcase the layouts of PSD_CHIP V1 and V2 respectively. Each color in the layout represents a different layer in the N-dimensional IC layout design space. On top of the base p-substrate of a wafer, NMOS, PMOS transistors are fabricated (with W/L ratios for each having to be specified). Then comes the poly-silicon and multiple metal layers. Due to limited area for chip layout, routing involves several vertical metal layers on top of the 2D surface routing for each layer. Overlaps between them create unwanted parasitic capacitive coupling that also need to be taken into account. Lengths and widths for each trace need to be carefully specified since they directly contribute to overall parasitic resistances. In general, IC design must include an additional round of parasitics included simulations and layout optimization, without which there is a large risk that the chip will not work the way it was intended to.

Due to time constraints, for V1, a lot of the parasitics included checks were not able to be conducted. As a result, it is not wholly unexpected that we had several issues with chip performance stemming from layout parasitics. For V2, power busses, inter-channel and intra-channel routing needed to be redone and optimized. This is obvious in comparing the layouts of V1 (Figure 6.21) and V2 (Figure 6.22). A much more careful study of parasitics was conducted all around to prevent the appearance of the same issues. We are able to take microscope pictures of the returned bare dies from the foundry. Such a picture for V1 is shown in Figure 6.23. Only the top metal layer and pads are visible in the photograph.

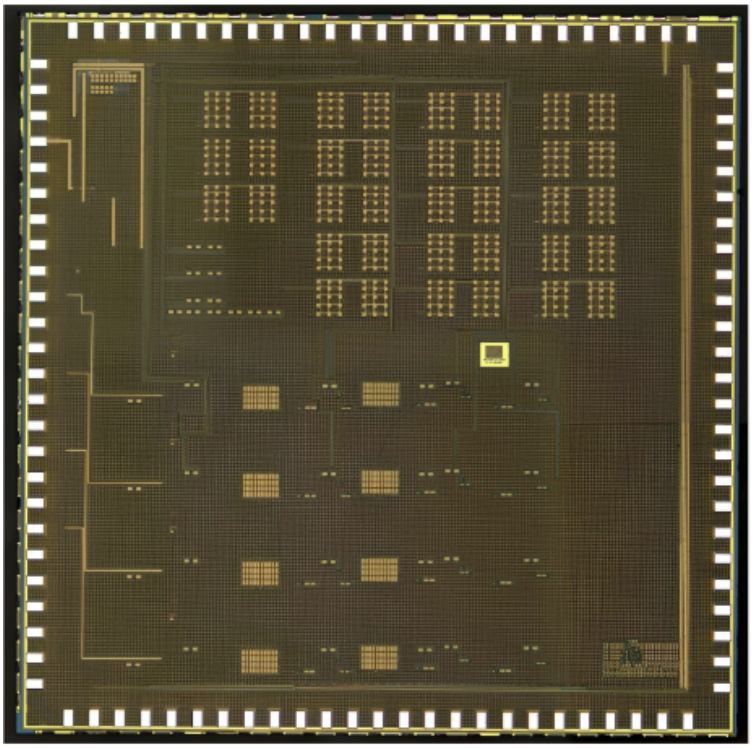

Figure 6.23: A microscope picture taken of the PSD_CHIP_V1 bare die as received from the foundry is shown. Only the top metal layer traces are visible in the bare die photograph.

6.3 Initial Testing of PSD_CHIP

6.3.1 Custom Test Board Design

The bare dies of PSD_CHIP_V1 were wire-bonded onto quad flat-pack (QFP-100) ceramic packages and then soldered onto a landing padring on a custom designed initial verification test board. A photograph of this board is provided in Figure 6.25. The test board was designed to be flexible for both versions of PSD_CHIP so that it can be reused in the future initial testing and verification of PSD_CHIP_V2.

The board features a three-pin Phoenix connector through which +/- 5 V and a GND connection are delivered. On the board, the +/-5 V rails deliver power to a mating connector. This connector allows for a Raspberry Pi to be directly interfaced with the test board. In addition, these power rails are inputs for 5 LDO regulators, which generate the various power/ground rails for PSD_CHIP and the other discrete components on the test board. Each LDO is also configured with a mechanical pot, which can be used to fine-adjust the output voltages of these ICs, if needed. The regulators feature a enable pin, which is directly tied to one of the Raspberry Pi's general purpose in/out (GPIO) pins. This allows for control over when power is delivered by the regulators to PSD_CHIP and other components on the board independently of when the board is powered on. Thus, initially, when the board is powered on, only the Raspberry Pi will receive power. The LDOs also provide noise immunity from the main power inputs, which is an important consideration since the various on-chip voltage DACs use VDDD/VSSD as the positive and negative references. For a similar reason, VDDA/VSSA and VDDD/VSSD are independently supplied by different regulators to prevent digital noise injection into the analog front-ends of the chip's four channels.

The four SOUT and FOUT signal inputs to the chip are brought in through SMB connectors. Each of these connectors are 50 Ω terminated to prevent reflection issues. In series with the four SOUT inputs are 1kΩ resistors. For initial verification, the test input signals for both SOUT and FOUT are generated through an arbitrary waveform generator (AWG). Thus, the series resistors provide a simple voltage to current conversion for the SOUT signals.

One important thing to note is that these boards can also be interfaced directly with the previously discussed SiPM test bed. Instead of feeding the SOUT and FOUT inputs from the fanout board directly to the CAEN digitizer, they can be connected to the relevant SMB connectors on the test board. In this scenario, all that needs to be done is to have the series resistors switched out to 50 Ω resistors instead. For the FOUT inputs, a 2nF AC coupling capacitor is connected in series with the respective SMB connectors before the connection to the relevant input pads of the chip. This was done to decouple the common source pre-amplifier's DC operating point from the baseline of the input AWG signal. If bypassing of the FOUT inputs is desired, an SMB connector is included for providing a global external trigger for the four FOUT channels.

The board also features dedicated SMB connectors for each of the six main outputs per channel of PSD_CHIP along with four SMB connectors, which read out the digital testbuses. The six outputs are the front-end outputs for SOUT and FOUT, the partial and total integration stage outputs, the difference output and the final PSD classification discriminator output. For each of the chip's various analog outputs, a buffer is included on the test board to provide the capability to drive the capacitance of cables and 50 Ω loads externally (i.e. when viewing traces on an oscilloscope).

The Raspberry Pi requires +3.3 V/0 V digital logic levels and the chip requires +2.4 V/-0.9 V logic levels. As a result, level shifters are included in series as appropriate to all Raspberry Pi/external to chip I/O lines. Furthermore, an on-board two-channel ADC is used to convert the output of the chip's current and voltage monitors and allow for the

signal to then be read by the Raspberry Pi.

Finally, the current biases for the chip (other than RBIAS_TOTAL and RBIAS_PARTIAL) are set with the aid of mechanical potentiometers. An 8-bit external digital potentiometer (with a maximum resistance of 50 kΩ) is used to set the LSB current bit for the partial and total integration one-shots. For both, the maximum values of the potentiometers were selected so that it covers the range of biasing differences between the two versions of the chip as detailed in the appendices.

6.3.2 Initial Results

As previously discussed, there are several layout-related issues with PSD_CHIP_V1. This prevents any quantitative characterization of the V1 chip's performance. However, the test board has allowed for qualitatively verifying the full 1CH signal chain. Successful communication with the digital core registers through the custom UART line via custom written software in Python was demonstrated. A successful test of the digital core reset line was also performed, where the values of all the registers were verified to match the preset defaults. Initially, square wave pulses were used for both FOUT and SOUT inputs to perform an initial verification of the full 1CH signal chain. A full exercise in adjustment of all discriminator thresholds, baselines, current biases for the various current starved inverter based delays, and externally setting the LSB for the partial and total integration windows were undertaken with CH0 of the chip. Afterwards, an average representation of equal area fast

neutron and gamma waveforms were generated using the test bed AmBe dataset and used as input waveforms through the AWG. The results of propagating these waveforms through the 1CH signal chain (with all thresholds, baselines and delays set appropriately) is shown in Figure 6.26. This figure can directly be compared to Figure 6.17, which shows the full 1CH simulation for PSD_CHIP_V2.

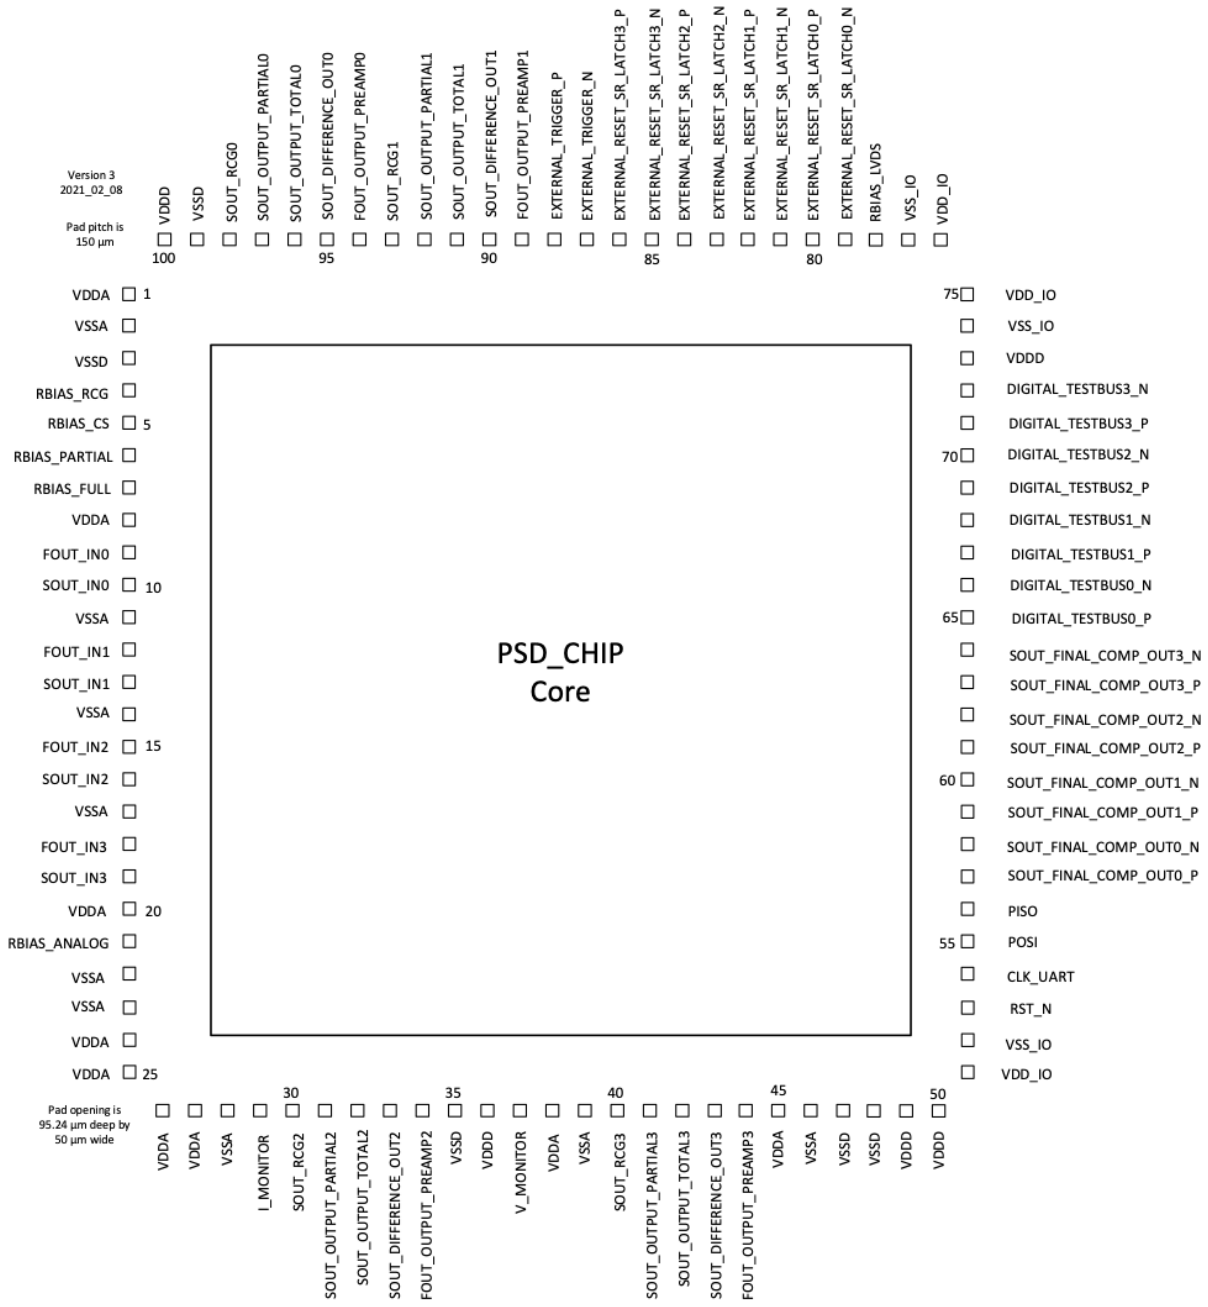

Figure 6.24: A more detailed look at the 100 pads that form the complete padring of both versions of PSD_CHIP. Pad pitch is 150um and individual pad sizes are 95um x 50um.

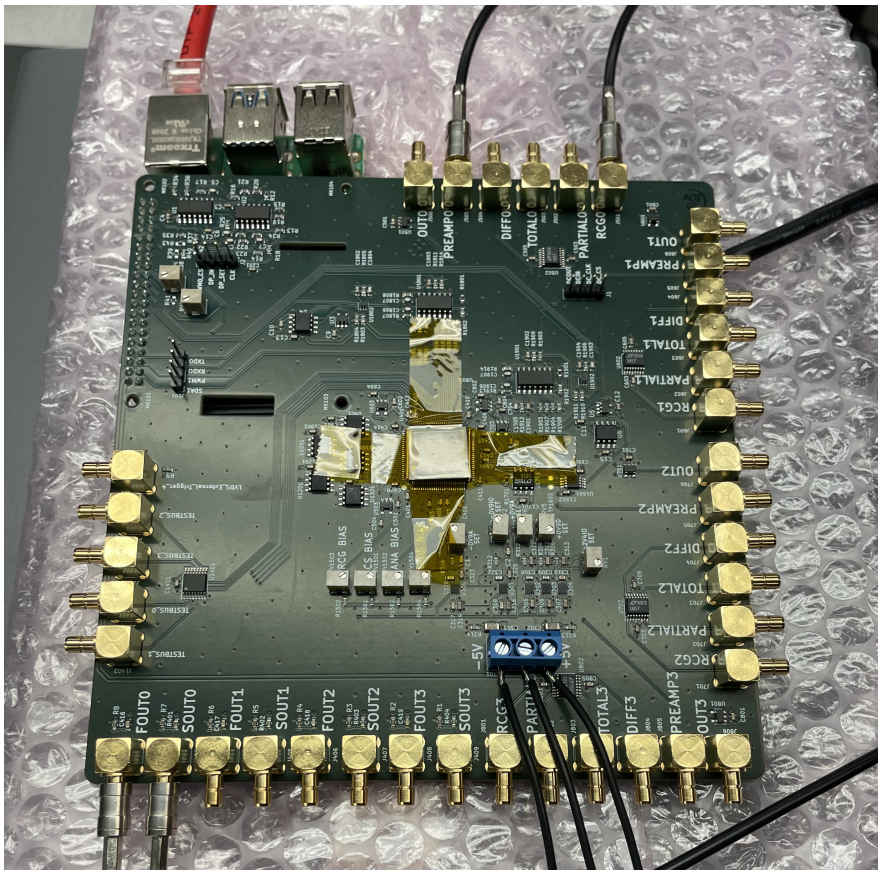

Figure 6.25: A photograph of the custom test board designed for initial verification of the PSD_CHIP bare dies. A breakdown of features of the test board is given in the text.

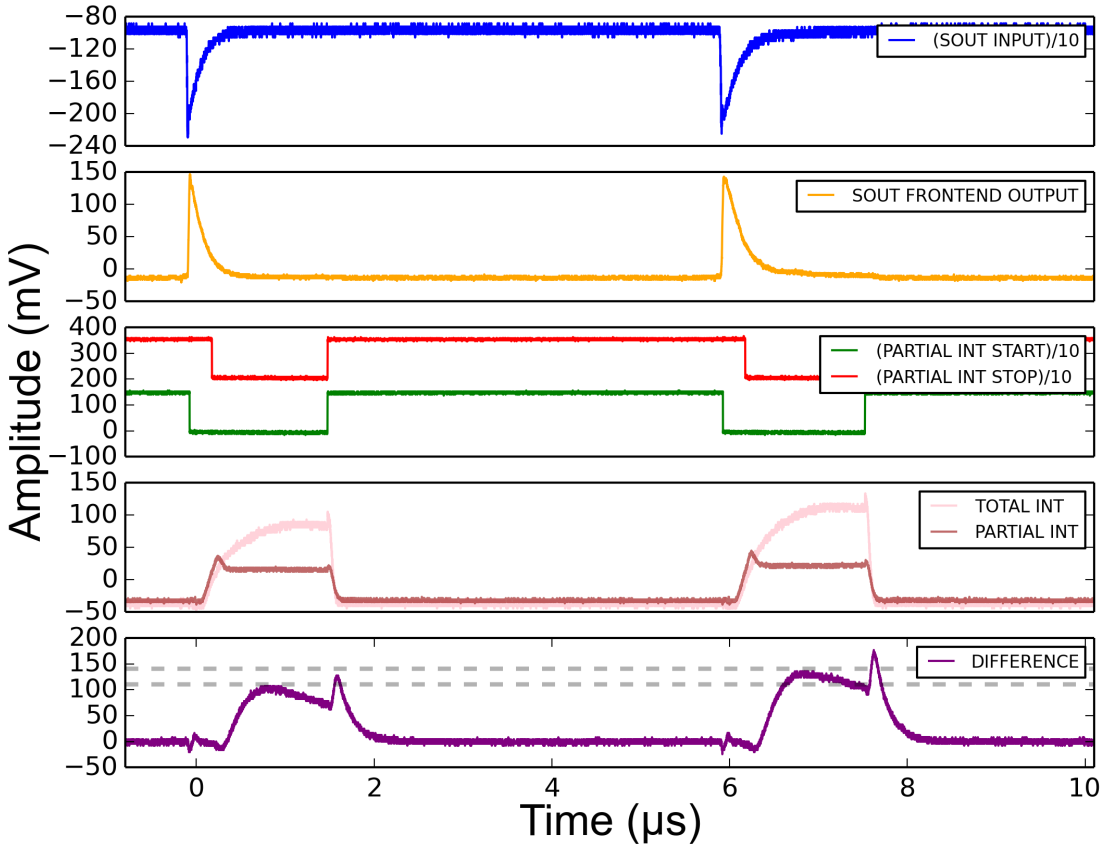

Figure 6.26: Scope waveform captures of various parts of the 1CH signal chain. An average waveform representation of equal area fast neutron and gamma SOUT waveforms from the test bed datasets were used as input signals. In the final row, the dashed line represents the peak of the subtraction op amp stage output and shows a clear difference in the values for the fast neutron and gamma waveform.

Chapter 7

Summary/Conclusions

This dissertation presents the results of a first attempt at setting upper exclusion limits on the inelastic EFT operators based on data collected by the LZ experiment. These 14 operators form a generalized interaction basis for non-relativistic WIMP-nucleon interactions. Upper limits for 6 mass-splitting terms per operator (a total of 84 limits) are shown in Chapter 4. To build up to the results, the latter half of Chapter 1 and Chapter 2 detail the design of LZ, currently the world's largest direct dark matter detection experiment. Work performed on simulating, testing and commissioning the analog signal chain for the experiment is emphasized. The beginning of Chapter 4 details the SR1 calibrations and subsequent data selection procedure. Results of a study looking at evaluating the latter for an extended ROI analysis are presented as well. Further, the validity of extending the standard WIMP search ROI cut efficiencies for the S1- and S2-based cuts is established.

World leading limits for the coupling strengths of all but operator 6 were established for

a zero mass-splitting value. Depending on the operator, the cross over point where upper limits were no longer world leading varied between mass-splitting values of 50 to 200 keV. The analysis was performed in the standard WIMP search ROI. The collaboration is currently working to extend the ROI for S1c from 80 phd to 600 phd, which allows for acceptance of NR events up to approximately 300 keV. The net effect will be much higher acceptances on the full differential rate recoil spectra for the higher mass-splitting terms across all operators. As a result, the limits presented in this dissertation are expected to show improvements. This work has provided for the first exercise of the full analysis procedure, and should serve as a reference for all future efforts.

In the second half of this dissertation, Chapter 5 presents a very thorough overview of PSD capable plastic and organic scintillators and SiPMs. It is followed by a discussion of a testbed involving SiPMs that were optically coupled to scintillators, which was designed and set up at UC Davis. Using data taken with the testbed and a sealed AmBe source, measurements of PSD capability in EJ276 and Stilbene are presented. Results are presented using the standard FOM metric as well as gamma leakage, both in PSD and subtraction space. Also covered in this chapter are the requirements for a compact, portable, and segmented neutron camera design. Further, an overview of a fully custom ASIC is presented. This ASIC has been developed to serve as a front-end for SensL (OnSemi) J-series SiPMs and provides real-time PSD capability. An analog to the chip's logic was implemented in software and a feasibility study using AmBe data was undertaken. The results of this study are discussed as well.

Finally, Chapter 6 presents a detailed discussion of the design of a prototype ASIC, which was thoroughly examined via simulations. An in depth overview of the one channel signal chain is provided, along with details of the full chip integration. Explicit details for both fabricated versions of this prototype chip are given in Appendices A and B. Simulation results using Cadence software are presented throughout the chapter and initial results from testing of V1 of the chip are also discussed at the end. The group at UC Davis will characterize and test V2 of the chip design and eventually integrate the chip into a prototype demonstration of the full neutron scatter camera electronics chain.

Appendix A

PSD_CHIP_V1 Data Sheets

This appendix is primarily included to fill in details on how to operate PSD_CHIP V1. It lists out specific details of the design that were not included in Chapter 6. The appendix is written to serve as a user guide to communicate required information for initial testing/characterization of the bare dies. The information is also applicable for the eventual integration of PSD_CHIP as part of a full prototype neutron scatter camera system. Both versions of PSD_CHIP were fabricated in a commercial 180 nm CMOS process node.

A.1 Pad Functions

The function of each I/O pad is described in more detail in Table A.1, including information on pad interfacing requirements. The digital I/O LVDS circuits use the VDD_IO (+2.4 V) and VSS_IO (**VSSD**) (-0.9V) power rails. Explicitly, for V1, even though VSS_IO and

VSSD are not shorted internally, these rails should be, externally, on any interface PCB. The internal digital core, FOUT sub-chain digital logic blocks and most SOUT analog cells use the +0.9 V VDDD and -0.9 V VSSD power lines. Only the SOUT and FOUT front-ends use an isolated VDDA (+0.9 V) and VSSA (-0.9 V) power lines.

Pad Name	**Direction**	**Type**	**Description**	**Num. Pads**	**Add. Info.**
SOUT_IN [3:0]	IN	ANALOG	Nominally, a current pulse input from SiPM. However, this can be any test current pulse as well.	4	Direct connection from SiPM output to pad.
FOUT_IN [3:0]	IN	ANALOG	Nominally, a voltage pulse input from the SiPM. This can be a test voltage pulse as well. For SiPMs other than SensL's, this should be an external T0 trigger.	4	2nF cap in series from SiPM output to pad is needed for decoupling of DC operating point of internal CSA from external ground referenced FOUT signal.
VDDA	IN/OUT	POWER	+ 0.9 V (nominal) Analog Power.	8	Voltage should not exceed +10% of nominal.
VDDD	IN/OUT	POWER	+ 0.9 V (nominal) Digital Power.	5	Voltage should not exceed +10% of nominal.
VSSA	IN/OUT	POWER	- 0.9 V (nominal) Analog Ground.	10	Voltage should not exceed +10% of nominal.
VSSD	IN/OUT	POWER	- 0.9 V (nominal) Digital Ground.	5	Voltage should not exceed +10% of nominal.

VDD_IO	IN/OUT	POWER	+ 2.4 V (nominal) Digital I/O Power.	3	Voltage should not exceed +10% of nominal.
VSS_IO /**VSSD**	IN/OUT	POWER	- 0.9 V (nominal) Digital I/O Ground.	3	Voltage should not exceed +10% of nominal. VSS_IO should be shorted to VSSD at the board level (it is not on chip for V1, but should have been).
POSI	IN	DIGITAL	Primary-Out-Secondary-In. Digital output of FPGA, etc to chip (change on negative edge).	1	Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.
PISO	OUT	DIGITAL	Primary-In-Secondary-Out. Digital output to FPGA, etc from chip (change on negative edge).	1	Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.
CLK_UART	IN	DIGITAL	Digital clock input for UART communication. This is a 16x oversampling CLK. As a result, data transfer rates are only 1/16 of the CLK frequency.	1	Only needs to be active during data transmission. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.

RST_N	IN	DIGITAL	Reset all internal 8-bit registers to preset programmed default values. These values are found in the table listing all digital core registers.	1	Digital core is pre-programmed with default values for all 8 bits of all registers. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.
EXTERNAL_ TRIGGER_ $\{P/N\}$	IN	DIGITAL	Differential, current-mode external trigger that can be used to start integration of incident waveform independent of FOUT triggering. NOTE: this is a global trigger for all enabled channels.	2	Digital I/O. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.
V_MONITOR	OUT	ANALOG	Monitor port for various internal analog DC voltage references. Select specific voltage by setting relevant digital core register's bits.	1	Full output range is nominally +/-0.9V (VDDD/VSSD).
DIGITAL_ TESTBUS_ $\{P/N\}$ [3:0]	OUT	DIGITAL	Monitor ports for four of any eight internal nodes of the FOUT triggering logic chain (per channel). Each test bus is a 32x1 digital MUX.	8	Digital I/O. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.

I_MONITOR	OUT	ANALOG	Monitor port for verifying internal current DACs. Must be terminated with appropriate resistor to (nominally) -0.9V (VSSD). The value of interest is the current across the termination resistor.	1	Max. expected output currents expected ∼16uA.
RBIAS_RCG	IN/OUT	ANALOG	Set the internal current bias for the SOUT RCG front end amplifiers (all channels). See Table A.4 for actual resistor value needed. The internal current bias will be the current across the resistor.	1	Resistor should deliver 10 mA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.
RBIAS_CS	IN/OUT	ANALOG	Set the internal current bias for the FOUT CS front end amplifiers (all channels). See Table A.4 for actual resistor value needed. The internal current bias will be the current across the resistor.	1	Resistor should deliver 100 uA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.

RBIAS_ PARTIAL	IN/OUT	ANALOG	Set the LSB for an internal current DAC (all channels), which sets the current bias for the current starved inverter delay line that controls the partial integration enable line. See Table A.4 for actual resistor value needed. The internal LSB current will be the current across the resistor divided down internally by x100.	1	Resistor should deliver b/w 20 - 100 uA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.
RBIAS_FULL	IN/OUT	ANALOG	Set the LSB for an internal current DAC (al channels), which sets the current bias for the current starved inverter delay line that controls the total integration enable line. See Table A.4 for actual resistor value needed. The internal LSB current will be the current across the resistor divided down internally by x100.	1	Resistor should deliver b/w 20 - 100 uA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.

RBIAS_ ANALOG	IN/OUT	ANALOG	Sets the internal current bias that will be scaled internally as required to properly bias all op amps, discriminators, etc (all channels). See Table A.4 for actual resistor value needed.	1	Resistor should deliver 50 uA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.
RBIAS_ LVDS	IN/OUT	ANALOG	Sets the internal current bias for the LVDS TX and RX modules on chip (all channels). See Table A.4 for actual resistor value needed.	1	Resistor should deliver 50 uA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.
EXTERNAL_ RESET_SR_ LATCH_{P/N} [3:0]	IN	DIGITAL	Optional reset of all SR latches in the FOUT triggering digital logic chain for a given channel.	8	Digital I/O. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.
SOUT_RCG [3:0]	OUT	ANALOG	Output of the SOUT waveform post RCG front-end amplification for each channel. A board level buffer is required to preserve signal bandwidth before driving small resistive loads/cables/etc.	4	Full range of signal is (nominally) +/-0.9V (VDDD / VSSD), with a baseline set by internal transistor DC operating points.

SOUT_ OUTPUT_ TOTAL [3:0]	OUT	ANALOG	Output of the SOUT total integration stage for each channel. A board level buffer is required to preserve signal bandwidth before driving small resistive loads/cables/etc.	4	Full range of signal is (nominally) +/-0.9V (VDDD / VSSD), with a baseline set by internal transistor DC operating points.
SOUT_ OUTPUT_ PARTIAL [3:0]	OUT	ANALOG	Output of the SOUT partial integration stage for each channel. A board level buffer is required to preserve signal bandwidth before driving small resistive loads/cables/etc.	4	Full range of signal is (nominally) +/-0.9V (VDDD / VSSD), with a baseline set by internal transistor DC operating points.
SOUT_ DIFFERENCE_ OUT [3:0]	OUT	ANALOG	Output of the SOUT subtraction op amp stage of partial and total integration outputs for each channel. A board level buffer is required to preserve signal bandwidth before driving small resistive loads/cables/etc.	4	Full range of signal is (nominally) +/-0.9V (VDDD / VSSD), with a baseline set by internal transistor DC operating points.

FOUT_ OUTPUT_ PREAMP [3:0]	OUT	ANALOG	Output of the FOUT waveform post front-end CS amplification for each channel. A board level buffer is required to preserve signal bandwidth before driving small resistive loads/cables/etc. This buffer show be placed as close to pad as possible to prevent capacitive loading.	4	Full range of signal is (nominally) +/-0.9V (VDDD / VSSD), with a baseline set by internal transistor DC operating points.
SOUT_FINAL_ COMP_OUT_ $\{P/N\}$ [3:0]	OUT	DIGITAL	Final SOUT discriminator output (PSD) bit classifying a pulse as either fast neutron/gamma for each channel.	8	Digital I/O. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.

Table A.1: List of all PSD_CHIP V1 I/O pads and description of their functionality and interface requirements.

A.2 Digital I/O Pads

There are several LVDS digital I/O pads included as part of the PSD_CHIP V1 design. Each LVDS signal is differential and as a result, requires two pads. These are listed below in Table

A.2.

Signal/Pad Name	**Description**	**Function**	**Add. Info.**
EXTERNAL_ TRIGGER_ $\{P/N\}$	Input differential external trigger that can be used to trigger the one-shot pulse generation stages in lieu of an FOUT input. **EXT_TRIG_EN** must be set HIGH in the digital core.	A nominally 20 - 100 ns wide input digital pulse that is fed directly into the FOUT width trim block, bypassing the FOUT front-end stage. See Section 6.1.2 for more details.	This is a global trigger for all enabled channels.
DIGITAL_ TESTBUS_ $\{P/N\}$	Four output differential digital lines, each for one of 32 (8 per FOUT sub-chain) possible signals.	Each testbus' output is selected by setting the 5 bits for each of **digital_ testbus_sel** [3:0] in the digital core.	The mapping of values for the 5 bits to test bus outputs is given in Section A.6.
EXTERNAL_ RESET_SR_ LATCH_ $\{P/N\}$	Four input differential external reset digital lines, one for each channel. Will trigger the reset lines for all the SR Latches in that channel.	A nominally 20-100ns length digital pulse input will be fed directly to all the Reset lines of the SR latches for each channel.	
SOUT_ FINAL_COMP_ OUT_ $\{P/N\}$	Four output differential digital lines, one for each channel, will be triggered HIGH if the subtraction op amp output goes above the user-set threshold.	The line will trigger HIGH at the moment during the integration window that the threshold is first crossed and remain HIGH until the end of the integration window, unless **sel_std_ sout_comp_width** is enabled in the digital core. See Section 6.1.9 and Table A.6 for more details.	

Table A.2: LVDS digital I/O pads are listed. A more detailed description and explanation of each signal's functionality is provided.

A.3 Digital Core Registers

The digital core of PSD_CHIP V1 contains 39 internally connected 8-bit registers and two spare 8-bit registers, which are not internally connected. This block is configured to be written to/read from using a simple serial communication protocol implemented using UART (described in more detail in Section 6.2.1.1). It also connects all digital bit data lines for each register to the appropriate current and voltage DACs, switches, and tunable resistors internal to the chip. The bit lines act as control bits for all of these blocks. Several bits also control additional functionality, such as enabling external or cross triggering, enabling/disabling individual channels, etc. Table A.3 lists all of the on-chip registers, along with descriptions and default values for each.

The defaults are based on Verilog HDL notation. The number before the ' indicates the number of bits or hex digits in the number, the letter b or h after ' indicates whether the number is in binary or hexadecimal, and a single value after the letter indicates all bits or hex digits are identical. A single 0 or 1 means the bit is repeated. For example, 8'b0 means 0000_0000 (or eight individual bits set to zero) and 8'h03 means the 8-bit number 0000_0011.

Signal Name	Description	Register [bits]	# bits	Default; Default Decimal	Add. Info.
disable_ channel	0: disable, 1: enable	0 [3:0]	4	4'h0; 0	All channels disabled at power up by default
enable_ replica	0: disable buffer reference DACs and uses internal replica bias circuits per channel; 1: enable	0 [7:4]	4	4'hF; 15	Replica bias circuits are enabled by default for all channels
v_ buff_ ref_ global	Sets global threshold range for the buffer reference voltage for all 4 channels (implemented using string DAC).	1 [2:0]	3	8'h07; 7	See Section A.5.1 for more details, including mapping of digital 3-bit values to global ranges.
v_ buff_ ref_ fine0	Sets fine-tuned threshold for the buffer reference voltage for channel 0 (implemented using 8- bit R2R Voltage DAC).	2 [7:0]	8	8'h77; 119	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ buff_ ref_ fine1	Sets fine-tuned threshold for the buffer reference voltage for channel 1 (implemented using 8- bit R2R Voltage DAC).	3 [7:0]	8	8'h77; 119	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_ buff_ ref_ fine2	Sets fine-tuned threshold for the buffer reference voltage for channel 2 (implemented using 8- bit R2R Voltage DAC).	4 [7:0]	8	8'h77; 119	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ buff_ ref_ fine3	Sets fine-tuned threshold for the buffer reference voltage for channel 3 (implemented using 8- bit R2R Voltage DAC).	5 [7:0]	8	8'h77; 119	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ fine_ refp	Sets positive voltage reference for associated v_ref_comp0 thru v_ref_comp3 8-bit voltage DACs.	6 [7:0]	8	8'h8E; 142	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ fine_ refn	Sets negative voltage reference for associated v_ref_comp0 thru v_ref_comp3 8-bit voltage DACs.	7 [7:0]	8	8'h70; 112	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_ ref_ comp0	Baseline reference voltage for op amp that converts differential outputs of subtraction stage to single ended before SOUT final discriminator for channel 0 (implemented using 8- bit R2R Voltage DAC).	8 [7:0]	8	8'h81; 129	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ ref_ comp1	Baseline reference voltage for op amp that converts differential outputs of subtraction stage to single ended before SOUT final discriminator for channel 1 (implemented using 8- bit R2R Voltage DAC).	9 [7:0]	8	8'h81; 129	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ ref_ comp2	Baseline reference voltage for op amp that converts differential outputs of subtraction stage to single ended before SOUT final discriminator for channel 2 (implemented using 8- bit R2R Voltage DAC).	10 [7:0]	8	8'h81; 129	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_ ref_ comp3	Baseline reference voltage for op amp that converts differential outputs of subtraction stage to single ended before SOUT final discriminator for channel 3 (implemented using 8- bit R2R Voltage DAC).	11 [7:0]	8	8'h81; 129	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ comp_ thresh_ sout_ global	Sets global threshold range for the SOUT discriminator threshold across all 4 channels (implemented using string DAC).	12 [2:0]	3	8'h07; 7	See Section A.5.1 for more details, including mapping of digital 3-bit values to global voltage range.
v_ comp_ thresh_ sout_ fine0	Sets fine-tuned threshold for the SOUT discriminator threshold for channel 0 (implemented using 8-bit R2R Voltage DAC).	13 [7:0]	8	8'hF2; 242	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ comp_ thresh_ sout_ fine1	Sets fine-tuned threshold for the SOUT discriminator threshold for channel 1 (implemented using 8-bit R2R Voltage DAC).	14 [7:0]	8	8'hF2; 242	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_comp_thresh_sout_fine2	Sets fine-tuned threshold for the SOUT discriminator threshold for channel 2 (implemented using 8-bit R2R Voltage DAC).	15 [7:0]	8	8'hF2; 242	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_comp_thresh_sout_fine3	Sets fine-tuned threshold for the SOUT discriminator threshold for channel 3 (implemented using 8-bit R2R Voltage DAC).	16 [7:0]	8	8'hF2; 242	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_comp_thresh_fout_global	Sets global threshold range for the FOUT discriminator threshold across all 4 channels (implemented using string DAC).	17 [2:0]	3	8'h02; 2	See Section A.5.1 for more details, including mapping of digital 3-bit values to global voltage range.
v_comp_thresh_fout_fine0	Sets fine-tuned threshold for the FOUT discriminator threshold for channel 0 (implemented using 8-bit R2R Voltage DAC).	18 [7:0]	8	8'h1C; 28	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_comp_thresh_fout_fine1	Sets fine-tuned threshold for the FOUT discriminator threshold for channel 1 (implemented using 8-bit R2R Voltage DAC).	19 [7:0]	8	8'h1C; 28	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_ comp_ thresh_ fout_ fine2	Sets fine-tuned threshold for the FOUT discriminator threshold for channel 2 (implemented using 8-bit R2R Voltage DAC).	20 [7:0]	8	8'h1C; 28	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ comp_ thresh_ fout_ fine3	Sets fine-tuned threshold for the FOUT discriminator threshold for channel 3 (implemented using 8-bit R2R Voltage DAC).	21 [7:0]	8	8'h1C; 28	See Section A.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
i_ bias_ total_ int	Control bits for CSI delay line for total integration width for ch0.	22 [4:0]	5	8'h10; 16	See Section A.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
i_ bias_ partial_ int	Control bits for CSI delay line for partial integration width for ch0.	23 [4:0]	5	8'h15; 21	See SectionsA.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.

i_ bias_ hold	Control bits for CSI delay line for HOLD enable width for ch0.	24 [4:0]	5	8'h19; 25	See Section A.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
i_ bias_ hold_ delay	current DAC for current starved inverter delay line for delay of hold line trigger. Sets the feedback delay before HOLD line is fed back into FOUT triggering logic to prevent race conditions (all channels).	25 [4:0]	5	8'h1B; 27	See Section A.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
i_ bias_ disc_ sout_ delay	current DAC for current starved inverter delay line for delay of SOUT final discriminator output logic. Sets a delay period during which final discriminator will not trigger HIGH (prevent spurious switching during initial integration period (all channels).	26 [4:0]	5	8'h19; 25	See Section A.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.

i_ bias_ width_ trim	current DAC for current starved inverter delay line that sets the standard width of digital pulses going to parallel total, partial, HOLD one-shot generation circuits (all channels).	27 [4:0]	5	8'h19; 25	See Section A.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
i_ bias_ sout_ comp_ width	current DAC for current starved inverter delay line that sets the standard width of SOUT discriminator output (all channels).	28 [4:0]	5	8'h19; 25	See Section A.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
sel_ std_ sout_ comp_ width	1: select std. width option for SOUT discriminator; 0: no std. width (all channels).	28 [5]	1	1'b0; 0	
tunable_ res_ total_ int	Adjust resistor value in total integration channel (one-hot) 8 MSBs, so corresponds to **int_ trim_ total [14:7]** (all channels).	29 [7:0]	8	8'h00; 0	See Section A.8 for more details.
tunable_ res_ msbs	Despite the name, 7 LSBs for tunable resistor of total integration stage. Corresponds to **int_ trim_ total [6:0]** (all channels).	30 [6:0]	7	8'h10; 16	See Section A.8 for more details.

tunable_ res_ subtr_ gain	Adjusts feedback resistor value (and thus gain) for subtraction stage output. Corresponds to **int_ trim_ subtr [5:0]**. (all channels).	31 [5:0]	6	8'h01; 1	See Section A.8 for more details.
digital_ testbus0_ sel	Select which FOUT triggering logic node is routed to digital_testbus0.	32 [4:0]	5	8'h00; 0	See Section A.6 for more details.
digital_ testbus0_ pd	1 to power down digital_testbus0.	32 [5]	1	1'b1; 1	
digital_ testbus1_ sel	Select which FOUT triggering logic node is routed to digital_testbus1.	33 [4:0]	5	8'h08; 8	See Section A.6 for more details.
digital_ testbus1_ pd	1 to power down digital_testbus1.	33 [5]	1	1'b1; 1	
digital_ testbus2_ sel	Select which FOUT triggering logic node is routed to digital_testbus2.	34 [4:0]	5	8'h10; 16	See Section A.6 for more details.
digital_ testbus2_ pd	1 to power down digital_testbus2.	34 [5]	1	1'b1; 1	
digital_ testbus3_ sel	Select which FOUT triggering logic node is routed to digital_testbus3.	35 [4:0]	5	8'h18; 24	See Section A.6 for more details.
digital_ testbus3_ pd	1 to power down digital_testbus3.	35 [5]	1	1'b1; 1	
sout_ comp_ out_ pd	1 to power down channel (SOUT PSD bit) LVDS output.	36 [3:0]	4	4'b0; 0	
lvds_ rx_ pd	1 to power down all other LVDS outputs across all channels.	36 [4]	1	1'b1; 1	See Section 6.2.6 for more details.

lvds_ high_ cm	1 to use higher LVDS common mode (might be needed depending on external LVDS IC used).	36 [5]	1	1'b1; 1	See Section 6.2.6 for more details.
lvds_ loop-back_ enable	1 to loop RX back to TX (one channel).	36 [6]	1	1'b0; 0	See Section 6.2.6 for more details.
external_ trigger_ enable	1 to enable external trigger.	37 [0]	1	1'b0; 0	See Section 6.1.2 for more details.
cross_ trigger_ enable	1 to enable cross triggering functionality.	37 [1]	1	1'b0; 0	See Section 6.1.2 for more details.
current_ monitor	1 to enable current monitor; used to verify performance of internal 5-bit current DACs.	38 [0]	1	1'b0; 0	See Section A.7.2 for more details.
voltage_ monitor	1 to enable voltage monitor (one-hot); used to monitor eight internal DC voltages.	39 [7:0]	8	8'h00; 0	See Section A.7.1 for more details.
spare0	spare config. register	40 [7:0]	8	8'h00; 0	No internal connections from digital core to rest of the chip.
spare1	spare config. register	41 [7:0]	8	8'h00; 0	No internal connections from digital core to rest of the chip.

Table A.3: All registers of the digital core for PSD_CHIP V1 are listed. Information on what the bits from each register control on-chip are provided along with their pre-programmed default values.

A.4 RBIAS Pads

The required resistor values are listed in Table A.4 for all RBIAS pads. A description of these pads is provided in Section 6.2.7.

RBIAS Pad Name	**Nominal Current Bias**	**Equivalent Nominal Resistor Value**
RBIAS_ANALOG	50 uA	$\sim$ 24.6 kΩ
RBIAS_LVDS	50 uA	$\sim$ 24.6 kΩ
RBIAS_RCG	10 mA	$\sim$ 114.7 Ω
RBIAS_CS	100 uA	$\sim$ 11.9 kΩ

Table A.4: List of all V1 RBIAS I/O pads with nominal current biases and the equivalent external termination resistor values needed for each.

A.5 On Chip Voltage and Current DACs

A.5.1 Voltage DACs

There are several highly tunable DC voltage references and thresholds that need to be supplied to relevant parts of the SOUT and FOUT sub-chains. These voltages are set using a two stage coarse/fine tuning process. A coarse global voltage range is defined for all four channels for every tunable voltage. This is either set using a 3-bit string DAC or a standard 8-bit R2R DAC. For the 3-bit string DAC, setting the control bits is relatively simple. Setting

the bits to 000 will set the positive reference voltage to +900 mV (VDDD) and negative reference voltage to +675 mV (VDDD - (VDDD - VSSD)/8). 001 sets the positive reference range to +675 mV (VDDD - (VDDD - VSSD)/8) and negative reference range to +450 mV (VDDD - 2(VDDD - VSSD)/8) and so on. Thus, 111 sets the positive reference voltage to -675 mV (VDDD -7(VDDD - VSSD)/8) and the negative reference voltage to -900 mV (VSSD). The coarse global voltage range can span the full 1.8 V (nominal) power supply swing. This global range sets the positive and negative reference voltages for a subsequent 8-bit R2R DAC that fine-tunes the tunnable voltage for each respective channel. The equation to set the values for the control bits for the 8-bit R2R DACs is:

$$V_{out} = \frac{V_{ref;p} - V_{ref;n}}{256} * v_{fine} < 7:0 > [\text{decimal}] + V_{ref;n}, \tag{A.1}$$

where $v_{fine} < 7:0 >$ [decimal] is the decimal equivalent value for the 8 control bits that make up $v_{fine} < 7:0 >$. $V_{ref;p}$ and $V_{ref;n}$ are the selected positive and negative global reference range.

The fine-tuned voltage is buffered before being routed to the relevant internal circuitry. Table A.5 lists the tunable voltages for PSD_CHIP V1. The relevant digital core register is named in parenthesis (some of the names are identical).

Name (Digital Core Equiv.)	**Description**	**Max Val.**	**Min. Val**	**Step Size**	**Num. Bits**

V_BUFF_ REFERENCE (v_buff_ref_ global)	Set the global range for the integration stages DC baseline (to match baseline of buffer output).	900mV/ VDDD	-900mV/ VSSD	225mV/ (VDDD-VSSD)/8	3 (000: 675 mV to 900 mV; 111: -900 mV to -675 mV)
V_BUFF_ REFERENCE_ FINE [3:0] (v_buff_ref_fine [3:0])	Allows fine-tuning and selection of the DC baseline reference for the integration stages per channel.	225mV/ (VDDD - VSSD) /8	0V	0.88mV/ (Max. Val)/ 256	8 (see equation in Section A.5.1)
V_REF_COMP [3:0] (v_ref_comp [3:0])	Allows for channel level fine-tuning and selection of DC baseline for differential to single ended conversion of subtraction op amp stage output.	100mV (nominally)	0V (nominally)	0.39 mV/ (Max. - Min. Val.)/ 256	8 (see equation in Section A.5.1)
V_REF_ FINE_{P/N} (v_fine_ref{p/n})	Sets the global reference range for DC baseline of op amp that converts differential outputs of the subtraction op amp to single ended	900mV/ VDDD	-900mV/ VSSD	7.03 mV (over full-scale for each)	8 (see equation in Section A.5.1)
V_COMP_ THRESH_ SOUT_ GLOBAL (v_comp_ thresh_sout_ global)	Set the global range for the SOUT final discriminator threshold for all four channels	900mV /VDDD	-900mV /VSSD	225mV/ (VDDD -VSSD)/8	3 (000: 675mV - 900mV; 111: -900mV to -675mV)
V_COMP_ THRESH_SOUT_ FINE [3:0] (v_comp_ thresh_sout_ fine [3:0])	Allows for channel level fine-tuning and selection of SOUT discriminator threshold.	225mV/ (VDDD - VSSD) /8	0V	0.88mV/ (Max. Val.) /256	8 (see equation in Section A.5.1)

V_COMP_ THRESH_FOUT_ GLOBAL (v_comp_ thresh_fout_ global)	Set the global range for the FOUT discriminator threshold for all four channels	900mV / VDDD	-900mV /VSSD	225mV/ (VDDD -VSSD) /8	3 (000: 675mV-900mV; 111: -900mV to -675mV)
V_COMP_ THRESH_FOUT_ FINE [3:0] (v_comp_ thresh_fout_ fine [3:0])	Allows for channel level fine-tuning and selection of FOUT discriminator threshold.	225mV/ (VDDD - VSSD) /8	0V	0.88mV/ (Max. Val.) /256	8 (see equation in Section A.5.1)

Table A.5: List of all programmable voltages and their functionality for PSD_CHIP V1. All voltages listed are generated through internal voltage DACs with the control bit values set through the digital core.

A.5.2 Current DACs

Several 5-bit current DACs are included as part of the PSD_CHIP design. They require as input, a LSB current. This LSB current, depending on the specific current DAC's role is either externally tunable or a fixed current input. The current DACs take the currents generated at either of RBIAS_TOTAL, RBIAS_PARTIAL or RBIAS_ANALOG pads as the input LSB. The LSB for the DACs that set delay (width) for the START_INT and STOP_PARTIAL_INT enable lines are set externally using the current supplied (internally divided by x100) to the RBIAS_PARTIAL and RBIAS_TOTAL pads. All other current starved inverter delay line

LSBs are fixed at 500nA by a 100x internal division of the nominal 50uA RBIAS_ANALOG current input.

Depending on the specified 5-bit value, the output of a current DAC will be the input LSB current multiplied by the decimal-equivalent 5-bit value, with the caveat that control lines for the DAC are active LOW. Explicitly, the output current of the DAC is:

$$\text{Output Current } [A] = LSB[A] * (31 - \text{5-bit value [decimal]}). \tag{A.2}$$

Table A.6 lists all current DACs internal to PSD_CHIP V1. The number of CSI pairs used is different for each of the delay lines needed on chip and is specified in the table. More detail is found in Section 6.1.5.1.

Name	**Description**	**Max Val.**	**Min. Val.**	**Step Size**	**External LSB?**
CURR_BIAS_ TOTAL_INT	Supplied bias current for delay line that determines total integration length. 400 CS-inverter pairs in delay chain.	32uA	0A	variable b/w 200nA to 1uA	Yes; scaled down by 100 from external value set by RBIAS_ TOTAL pad.
CURR_BIAS_ PARTIAL_INT	Supplied bias current for delay line that determines partial integration length. 30 CS-inverter pairs in delay chain.	32uA	0A	variable b/w 200nA to 1uA	Yes; scaled down internally by 100 from external value set by RBIAS_ PARTIAL pad.

CURR_BIAS_ HOLD	Supplied bias current for delay line that determines length of the HOLD line. The length should always be >= total integration length. 400 CS-inverter pairs in delay chain.	15.5uA	0A	500nA	No
CURR_BIAS_ HOLD_DELAY	Supplied bias current for delay line that determines delay in feeding HOLD line back to RESET for SR latch to prevent re-triggering during ongoing active integration. The length should always be >= standard width of FOUT trigger pulses. 10 CS-inverter pairs in delay chain.	15.5uA	0A	500nA	No
CURR_BIAS_ SOUT_DELAY	Supplied bias current for delay line that determines delay in allowing SOUT final discriminator to trigger HIGH after start of integration. The length should always be >partial integration length by approx. 1uA. 10 CS-inverter pairs in delay chain.	16uA	500nA	500nA	No

CURR_BIAS_ FOUT_WIDTH_ TRIM	Supplied bias current for delay line that determines the standard width of FOUT trigger pulses to ensure clean input to one shot generation stages. 10 CS-inverter pairs in delay chain.	16uA	500nA	500nA	No
CURR_BIAS_ SOUT_COMP_ STD_WIDTH	Supplied bias current for delay line that determines the standard width of SOUT final discriminator output pulses to allow for uniform triggered HIGHs. 30 CS-inverter pairs in delay chain.	16uA	500nA	500nA	No

Table A.6: List of the various current DACs (and associated biases) for the respective CSI delay chains on PSD_CHIP V1.

A.6 Digital Test Buses

The digital test buses probe various points in the FOUT sub-chain. A detailed overview can be found in Section 6.2.3. Eight locations can be probed for each of the four channels. Thus, each digital test bus essentially functions as a 32x1 digital MUX. The buses can read out up to four probe locations for a given channel. They can also be used to read out the same location across all four channels. In normal operating mode, it is expected that each bus will read out the FOUT discriminator output (this will serve as the T0 pulse for PSD_CHIP). The

eight probe nodes are listed in Table A.7 below. The control bits for digital_testbus0_sel[4:0] to digital_testbus3_sel[4:0] need to be set using the appropriate decimal equivalent values in this table in order to output the correct node.

Name	**Description**	**Dec. Equiv. of 5-bit Val. [CH0,CH1, CH2,CH3]**	**Additional Comments**
HOLD [3:0]	This pulse is generated along with total and partial integration one-shot enable pulses. It will disable the FOUT triggering logic during an ongoing active integration to prevent re-triggering. The length of the HOLD pulse should be >= START_INT pulse.	0,8,16,24	see Section A.5.2 for actual characteristics of pulse.
HOLD_delayed [3:0]	The HOLD pulse, except with implemented delay in order to ensure no race condition. The length of this delay should be >the standard width selected for the width trim stages.	1,9,17,25	see Section A.5.2 for actual characteristics of pulse.
START_INT [3:0]	This pulse is generated to enable total and partial integration of an incident pulse.	2,10,18,26	see Section A.5.2 for actual characteristics of pulse.
STOP_PARTIAL_INT [3:0]	This pulse is generated to enable and set partial integration length of incident pulse.	3,11,19,27	see Section A.5.2 for actual characteristics of pulse.
CROSS_TRIG_OUT [3:0]	This is the digital pulse that is generated by the FOUT discriminator once incident waveform is above threshold.	4,12,20,28	This signal should be nominally output during actual data-taking runs to serve as T0.

TRIGGERING_OUT [3:0]	This is the pulse post FOUT triggering block.	5,13,21,29	see Section A.5.2 for actual characteristics of pulse.
width_trim_stg1_out [3:0]	The output of the first of two width trim stages that ensures a standard width digital pulse input to the one-shot pulse generation stages.	6,14,22,30	see Section A.5.2 for actual characteristics of pulse.
width_trim_out [3:0]	The final output of the two width trim stages that ensures a standard width digital pulse input to the one-shot pulse generation stages.	7,15,23,31	see Section A.5.2 for actual characteristics of pulse.

Table A.7: List of all probe nodes that are selectable outputs through the four digital test buses.

A.7 Voltage and Current Monitors

A.7.1 Voltage Monitor

The internal voltage monitor functions equivalently to an 8x1 analog MUX. However, eight bits are needed to address the monitor since it works on the one-hot principle (i.e. only one bit can be HIGH at a time). All four baseline reference voltages for the associated SOUT integration stages and the SOUT and FOUT discriminator thresholds for channels 0 and 1 are selectable. These analog voltages are used to verify the performance of the on-chip

coarse/fine voltage DACs set-up. The relevant digital bits in the digital core to set are **voltage_monitor** [7:0].

Name	**Description**	**Dec. Equiv. One-Hot Val.**
v_buff_bias0	Reference voltage for SOUT integration stages (channel 0)	0
v_buff_bias1	Reference voltage for SOUT integration stages (channel 1)	1
v_buff_bias2	Reference voltage for SOUT integration stages (channel 2)	2
v_buff_bias3	Reference voltage for SOUT integration stages (channel 3)	3
v_SOUT_ comp_thresh0	SOUT discriminator threshold for channel 0.	4
v_SOUT_ comp_thresh1	SOUT discriminator threshold for channel 1.	5
v_FOUT_ comp_thresh0	FOUT discriminator threshold for channel 0.	6
v_FOUT_ comp_thresh1	FOUT discriminator threshold for channel 1.	7

Table A.8: List of all voltage monitor outputs and the associated bit that needs to be HIGH to select a given output.

A.7.2 Current Monitor

There is only one possible current monitor output. This monitor's main purpose is to verify the performance of the on-board 5-bit current DACs that output bias currents for the various CSI-based delay lines. The associated bits in the digital core are **current_monitor** [7:0]. Although there are 8 bits, only **current_monitor** [0] is relevant. Setting the bit HIGH

activates the current monitor output and LOW disables it. The current monitor output must be terminated to -0.9 V (VSSD) through an appropriate resistor. The relevant signal is the current through this resistor.

A.8 Tunable Resistors

A.8.1 Total Integration Tunable Resistor

The total integration stage features a programmable resistor. The selectable values are listed below in Table A.9. See Section 6.1.7.1 for further details.

Digital Bit (one-hot)	**Resistor Value**	**Equiv. Partial/Total Ratio**
int_trim_total [0]	5 kΩ	1
int_trim_total [1]	5.5 kΩ [+0.5 kΩ]	0.9
int_trim_total [2]	6 kΩ [+0.5 kΩ]	0.83
int_trim_total [3]	6.5 kΩ [+0.5 kΩ]	0.76
int_trim_total [4]	7 kΩ [+0.5 kΩ]	0.71
int_trim_total [5]	7.5 kΩ [+0.5 kΩ]	0.67
int_trim_total [6]	8 kΩ [+0.5 kΩ]	0.62
int_trim_total [7]	10 kΩ [+2 kΩ]	0.5
int_trim_total [8]	11 kΩ [+1 kΩ]	0.45
int_trim_total [9]	12 kΩ [+1 kΩ]	0.41
int_trim_total [10]	15 kΩ [+3 kΩ]	0.33
int_trim_total [11]	17 kΩ [+2 kΩ]	0.29
int_trim_total [12]	20 kΩ [+3 kΩ]	0.25
int_trim_total [13]	25 kΩ [+5 kΩ]	0.20
int_trim_total [14]	30 kΩ [+5 kΩ]	0.16

Table A.9: List of selectable resistor values in the long integration stage and the equivalent effective threshold PSD ratio.

A.8.2 Subtraction Stage Tunable Resistor

The subtraction stage features a programmable feedback resistor. The selectable values are listed below in Table A.10. See Section 6.1.8 for further details.

Digital Bit (one-hot)	**Resistor Value**	**Equivalent Gain**
int_trim_subtr [0]	25kΩ	x1
int_trim_subtr [1]	37.5kΩ [+12.5kΩ]	x1.5
int_trim_subtr [2]	50kΩ [+12.5kΩ]	x2
int_trim_subtr [3]	75kΩ [+25kΩ]	x3
int_trim_subtr [4]	100kΩ [+25kΩ]	x4
int_trim_subtr [5]	125kΩ [+25kΩ]	x5

Table A.10: List of resistor values for selecting the gain of the buffer that converts the differential subtraction stage output to a single-ended signal.

Appendix B

PSD_CHIP_V2 Data Sheets

This appendix is primarily included to fill in details on how to operate PSD_CHIP V2. It lists out specific details of the design that were not included in Chapter 6. The appendix is written to serve as a user guide to communicate required information for initial testing/characterization of the bare dies. The information is also applicable for the eventual integration of PSD_CHIP as part of a full prototype neutron scatter camera system. Both versions of PSD_CHIP were fabricated in a commercial 180 nm CMOS process node. Most of the data tables in this appendix are similar, if not identical, to the tables in Appendix A for PSD_CHIP V1. This reflects the fact that the core design of PSD_CHIP has not changed between versions. However, because enough of the specifics have changed, separate appendices are provided for each versions to avoid confusion.

B.1 Pad Functions

The function of each I/O pad is described in more detail in Table B.1, including information on pad interfacing requirements. The digital I/O LVDS circuits use the VDD_IO (+2.4 V) and VSS_IO (**VSSD**) (-0.9V) power rails. Explicitly, for V2, VSS_IO and VSSD are shorted on chip, thus these rails should also be on any interface PCB. The internal digital core, FOUT sub-chain digital logic blocks and most SOUT analog cells use the +0.9 V VDDD and -0.9 V VSSD power lines. Only the SOUT and FOUT front-ends use an isolated VDDA (+0.9 V) and VSSA (-0.9 V) power lines.

Pad Name	**Direction**	**Type**	**Description**	**Num. Pads**	**Add. Info.**
SOUT_IN [3:0]	IN	ANALOG	Nominally, a current pulse input from SiPM. However, this can be any test current pulse as well.	4	Direct connection from SiPM output to pad.
FOUT_IN [3:0]	IN	ANALOG	Nominally, a voltage pulse input from the SiPM. This can be a test voltage pulse as well. For SiPMs other than SensL's, this should be an external T0 trigger.	4	2nF cap in series from SiPM output to pad is needed for decoupling of DC operating point of internal CSA from external ground referenced FOUT signal.
VDDA	IN/OUT	POWER	+ 0.9 V (nominal) Analog Power.	8	Voltage should not exceed +10% of nominal.

VDDD	IN/OUT	POWER	+ 0.9 V (nominal) Digital Power.	5	Voltage should not exceed +10% of nominal.
VSSA	IN/OUT	POWER	- 0.9 V (nominal) Analog Ground.	10	Voltage should not exceed +10% of nominal.
VSSD	IN/OUT	POWER	- 0.9 V (nominal) Digital Ground.	5	Voltage should not exceed +10% of nominal.
VDD_IO	IN/OUT	POWER	+ 2.4 V (nominal) Digital I/O Power.	3	Voltage should not exceed +10% of nominal.
VSS_IO /**VSSD**	IN/OUT	POWER	- 0.9 V (nominal) Digital I/O Ground.	3	Voltage should not exceed +10% of nominal. VSS_IO should be shorted to VSSD at the board level.
POSI	IN	DIGITAL	Primary-Out-Secondary-In. Digital output of FPGA, etc to chip (change on negative edge).	1	Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.
PISO	OUT	DIGITAL	Primary-In-Secondary-Out. Digital output to FPGA, etc from chip (change on negative edge).	1	Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.

CLK_UART	IN	DIGITAL	Digital clock input for UART communication. This is a 16x oversampling CLK. As a result, data transfer rates are only 1/16 of the CLK frequency.	1	Only needs to be active during data transmission. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.
RST_N	IN	DIGITAL	Reset all internal 8-bit registers to preset programmed default values.	1	Digital core is pre-programmed with default values for all 8 bits of all registers. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.
EXTERNAL_TRIGGER_$\{P/N\}$	IN	DIGITAL	Differential, current-mode external trigger that can be used to start integration of incident waveform independent of FOUT triggering. NOTE: this is a global trigger for all enabled channels.	2	Digital I/O. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.
V_MONITOR	OUT	ANALOG	Monitor port for various internal analog DC voltage references.	1	Full output range is nominally +/-0.9V (VDDD / VSSD).

DIGITAL_ TESTBUS_ $\{P/N\}$ [3:0]	OUT	DIGITAL	Monitor ports for four of any eight internal nodes of the FOUT triggering logic chain (per channel). Each testbus is a 32x1 digital MUX.	8	Digital I/O. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.
I_MONITOR	OUT	ANALOG	Monitor port for verifying internal current DACs. Must be terminated with appropriate resistor to (nominally) -0.9V (VSSD). The value of interest is the current across the resistor.	1	Max. expected output currents expected 20uA.
RBIAS_RCG	IN/OUT	ANALOG	Set the internal current bias for the SOUT RCG front end amplifiers (all channels). See Table B.4 for actual resistor value needed. The internal current bias will be the current across the resistor.	1	Resistor should deliver 5 mA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.

RBIAS_CS	IN/OUT	ANALOG	Set the internal current bias for the FOUT CS front end amplifiers (all channels). See Table B.4 for actual resistor value needed. The internal current bias will be the current across the resistor.	1	Resistor should deliver 500 uA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.
RBIAS_ PARTIAL	IN/OUT	ANALOG	Set the LSB for an internal current DAC (channels 0 and 2), which sets the current bias for the current starved inverter delay line that controls the partial integration enable line. See Table B.4 for actual resistor value needed. The internal LSB current will be the current across the resistor divided down internally by x100.	1	Resistor should deliver b/w 20 - 100 uA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.

RBIAS_FULL	IN/OUT	ANALOG	Set the LSB for an internal current DAC (channels 0 and 2), which sets the current bias for the current starved inverter delay line that controls the total integration enable line. See Table B.4 for actual resistor value needed. The internal LSB current will be the current across the resistor divided down internally by x100.	1	Resistor should deliver b/w 20 - 100 uA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.
RBIAS_ ANALOG	IN/OUT	ANALOG	Sets the internal current bias that will be scaled internally as required to properly bias all op amps, discriminators, etc (all channels). See Table B.4 for actual resistor value needed.	1	Resistor should deliver 50 uA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.
RBIAS_ LVDS	IN/OUT	ANALOG	Sets the internal current bias for the LVDS TX and RX modules on chip (all channels). See Table B.4 for actual resistor value needed.	1	Resistor should deliver 50 uA nominally. Terminate this pad to (nominally) - 0.9V (VSSD) through the resistor.

EXTERNAL_ RESET_SR_ LATCH_{P/N} [3:0]	IN	DIGITAL	Optional reset of all SR latches in the FOUT triggering digital logic chain for a given channel.	8	Digital I/O. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.
SOUT_RCG [3:0]	OUT	ANALOG	Output of the SOUT waveform post RCG front-end amplification for each channel. A board level buffer is required to preserve signal bandwidth before driving small resistive loads/cables/etc.	4	Full range of signal is (nominally) +/-0.9V (VDDD / VSSD), with a baseline set by internal transistor DC operating points.
SOUT_ OUTPUT_ TOTAL [3:0]	OUT	ANALOG	Output of the SOUT total integration stage for each channel. A board level buffer is required to preserve signal bandwidth before driving small resistive loads/cables/etc.	4	Full range of signal is (nominally) +/-0.9V (VDDD / VSSD), with a baseline set by internal transistor DC operating points.

SOUT_ OUTPUT_ PARTIAL [3:0]	OUT	ANALOG	Output of the SOUT partial integration stage for each channel. A board level buffer is required to preserve signal bandwidth before driving small resistive loads/cables/etc.	4	Full range of signal is (nominally) +/-0.9V (VDDD / VSSD), with a baseline set by internal transistor DC operating points.
SOUT_ DIFFERENCE_ OUT [3:0]	OUT	ANALOG	Output of the SOUT subtraction op amp stage of partial and total integration outputs for each channel. A board level buffer is required to preserve signal bandwidth before driving small resistive loads/cables/etc.	4	Full range of signal is (nominally) +/-0.9V (VDDD / VSSD), with a baseline set by internal transistor DC operating points.

FOUT_ OUTPUT_ PREAMP [3:0]	OUT	ANALOG	Output of the FOUT waveform post front-end CS amplification for each channel. A board level buffer is required to preserve signal bandwidth before driving small resistive loads/cables/etc. This buffer show be placed as close to pad as possible to prevent capacitive loading.	4	Full range of signal is (nominally) +/-0.9V (VDDD / VSSD), with a baseline set by internal transistor DC operating points.
SOUT_FINAL_ COMP_OUT_ $\{P/N\}$ [3:0]	OUT	DIGITAL	Final SOUT discriminator output (PSD) bit classifying a pulse as either fast neutron/gamma for each channel.	8	Digital I/O. Must be +2.4V/-0.9V digital logic as seen at the pad; level shifters may be required.

Table B.1: List of all PSD_CHIP V2 I/O pads and description of their functionality and interface requirements.

B.2 Digital I/O Pads

There are several LVDS digital I/O pads included as part of the PSD_CHIP V2 design. Each LVDS signal is differential and as a result, each requires two pads. These are listed below in

Table B.2.

Signal/Pad Name	**Description**	**Function**	**Add. Info.**
EXTERNAL_ TRIGGER_ $\{P/N\}$	Input differential external trigger that can be used to trigger the one-shot pulse generation stages in lieu of an FOUT input. **EXT_TRIG_EN** must be set HIGH in the digital core.	A nominally 20 - 100 ns wide input digital pulse that is fed directly into the FOUT width trim block, bypassing the FOUT front-end stage. See Section 6.1.2 for more details.	This is a global trigger for all enabled channels.
DIGITAL_ TESTBUS_ $\{P/N\}$	Four output differential digital lines, each for one of 32 (8 per FOUT sub-chain) possible signals.	Each testbus' output is selected by setting the 5 bits for each of **digital_ testbus_sel** [3:0] in the digital core.	The mapping of values for the 5 bits to test bus outputs is given in Section B.6.
EXTERNAL_ RESET_SR_ LATCH_ $\{P/N\}$	Four input differential external reset digital lines, one for each channel. Will trigger the reset lines for all the SR Latches in that channel.	A nominally 20-100ns length digital pulse input will be fed directly to all the Reset lines of the SR latches for each channel.	
SOUT_ FINAL_COMP_ OUT_ $\{P/N\}$	Four output differential digital lines, one for each channel, will be triggered HIGH if the subtraction op amp output goes above the user-set threshold.	The line will trigger HIGH at the moment during the integration window that the threshold is first crossed and remain HIGH until the end of the integration window, unless **std_width** is enabled in the digital core. See Section 6.1.9 and Table B.6 for more details.	

Table B.2: LVDS digital I/O pads are listed for V2. A more detailed description and explanation of each signal's functionality is provided.

B.3 Digital Core Registers

The digital core of PSD_CHIP V2 contains 62 internally connected 8-bit registers and four spare 8-bit registers, three of which are not internally connected. This block is configured to be written to/read from using a simple serial communication protocol implemented using UART (described in more detail in Section 6.2.1.1). It also connects all digital bit data lines for each register to the appropriate current and voltage DACs, switches, and tunable resistors internal to the chip. The bit lines act as control bits for all of these blocks. Several bits also control additional functionality, such as enabling external or cross triggering, enabling/disabling individual channels, etc. Table B.3 lists all of the on-chip registers, along with descriptions and default values for each.

The defaults are based on Verilog HDL notation. The number before the ' indicates the number of bits or hex digits in the number, the letter b or h after ' indicates whether the number is in binary or hexadecimal, and a single value after the letter indicates all bits or hex digits are identical. A single 0 or 1 means the bit is repeated. For example, 8'b0 means 0000_0000 (or eight individual bits set to zero) and 8'h03 means the 8-bit number 0000_0011.

Signal Name	**Description**	**Register [bits]**	**# bits**	**Default; Default Decimal**	**Add. Info.**
enable_channel	0: disable, 1: enable	0 [3:0]	4	4'hF; 15	All channels enabled at power up by default
frontend_ polarity_ sel	0: enable NMOS RCG input; 1: enable PMOS RCG input	0 [4]	1	1'b1; 1	PMOS RCG input enabled by default
frontend_ output_ sel	0: enable RCG output; 1: enable Symm. OTA output (global output)	0 [5]	1	1'b1; 1	Symm. OTA output enabled by default
int_ polarity_ sel	0: expected positive polarity for subtr. stage; 1: vice-versa. Discr. thresh. can be adjusted for either, but subsequent triggering logic needs specific polarity to output clean final SOUT classification bit	0 [6]	1	1'b1; 0	
v_ buff_ ref_ global	Sets global threshold range for the buffer reference voltage for all 4 channels (implemented using string DAC).	1 [2:0]	3	8'h04; 4	See Section B.5.1 for more details, including mapping of digital 3-bit values to global ranges.
v_ buff_ ref_ fine0	Sets fine-tuned threshold for the buffer reference voltage for channel 0 (implemented using 8- bit R2R Voltage DAC).	2 [7:0]	8	8'h61; 97	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_ buff_ ref_ fine1	Sets fine-tuned threshold for the buffer reference voltage for channel 1 (implemented using 8- bit R2R Voltage DAC).	3 [7:0]	8	8'h61; 97	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ buff_ ref_ fine2	Sets fine-tuned threshold for the buffer reference voltage for channel 2 (implemented using 8- bit R2R Voltage DAC).	4 [7:0]	8	8'h61; 97	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ buff_ ref_ fine3	Sets fine-tuned threshold for the buffer reference voltage for channel 3 (implemented using 8- bit R2R Voltage DAC).	5 [7:0]	8	8'h61; 97	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ fine_ refp	Sets positive voltage reference for associated v_ref_comp0 thru v_ref_comp3 8-bit voltage DACs.	6 [7:0]	8	8'h8E; 142	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_ fine_ refn	Sets negative voltage reference for associated v_ref_comp0 thru v_ref_comp3 8-bit voltage DACs.	7 [7:0]	8	8'h70; 112	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ ref_ comp0	Baseline reference voltage for op amp that converts differential outputs of subtraction stage to single ended before SOUT final discriminator for channel 0 (implemented using 8- bit R2R Voltage DAC).	8 [7:0]	8	8'h81; 129	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ ref_ comp1	Baseline reference voltage for op amp that converts differential outputs of subtraction stage to single ended before SOUT final discriminator for channel 1 (implemented using 8- bit R2R Voltage DAC).	9 [7:0]	8	8'h81; 129	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_ ref_ comp2	Baseline reference voltage for op amp that converts differential outputs of subtraction stage to single ended before SOUT final discriminator for channel 2 (implemented using 8- bit R2R Voltage DAC).	10 [7:0]	8	8'h81; 129	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ ref_ comp3	Baseline reference voltage for op amp that converts differential outputs of subtraction stage to single ended before SOUT final discriminator for channel 3 (implemented using 8- bit R2R Voltage DAC).	11 [7:0]	8	8'h81; 129	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ comp_ thresh_ sout_ global	Sets global threshold range for the SOUT discriminator threshold across all 4 channels (implemented using string DAC).	12 [2:0]	3	8'h03; 3	See Section B.5.1 for more details, including mapping of digital 3-bit values to global voltage range.
v_ comp_ thresh_ sout_ fine0	Sets fine-tuned threshold for the SOUT discriminator threshold for channel 0 (implemented using 8-bit R2R Voltage DAC).	13 [7:0]	8	8'h0A; 10	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_comp_thresh_sout_fine1	Sets fine-tuned threshold for the SOUT discriminator threshold for channel 1 (implemented using 8-bit R2R Voltage DAC).	14 [7:0]	8	8'h0A; 10	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_comp_thresh_sout_fine2	Sets fine-tuned threshold for the SOUT discriminator threshold for channel 2 (implemented using 8-bit R2R Voltage DAC).	15 [7:0]	8	8'h0A; 10	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_comp_thresh_sout_fine3	Sets fine-tuned threshold for the SOUT discriminator threshold for channel 3 (implemented using 8-bit R2R Voltage DAC).	16 [7:0]	8	8'h0A; 10	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_comp_thresh_fout_global	Sets global threshold range for the FOUT discriminator threshold across all 4 channels (implemented using string DAC).	17 [2:0]	3	8'h01; 1	See Section B.5.1 for more details, including mapping of digital 3-bit values to global voltage range.
v_comp_thresh_fout_fine0	Sets fine-tuned threshold for the FOUT discriminator threshold for channel 0 (implemented using 8-bit R2R Voltage DAC).	18 [7:0]	8	8'h93; 147	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_ comp_ thresh_ fout_ fine1	Sets fine-tuned threshold for the FOUT discriminator threshold for channel 1 (implemented using 8-bit R2R Voltage DAC).	19 [7:0]	8	8'h93; 147	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ comp_ thresh_ fout_ fine2	Sets fine-tuned threshold for the FOUT discriminator threshold for channel 2 (implemented using 8-bit R2R Voltage DAC).	20 [7:0]	8	8'h93; 147	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ comp_ thresh_ fout_ fine3	Sets fine-tuned threshold for the FOUT discriminator threshold for channel 3 (implemented using 8-bit R2R Voltage DAC).	21 [7:0]	8	8'h93; 147	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ baseline_ RCG_ global	Global adjustment of SOUT RCG output baseline voltage range implemented using 3-bit string DAC.	22 [3:0]	3	8'h04; 4	See Section B.5.1 for more details, including mapping of digital 3-bit values to global voltage range.
v_ baseline_ RCG_ fine0	Fine adjustment of SOUT RCG output baseline reference voltage implemented using 8-bit R2R voltage DAC for channel 0.	23 [7:0]	8	8'h6E; 110	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_ baseline_ RCG_ fine1	Fine adjustment of SOUT RCG output baseline reference voltage implemented using 8-bit R2R voltage DAC for channel 1.	24 [7:0]	8	8'h6E; 110	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ baseline_ RCG_ fine2	Fine adjustment of SOUT RCG output baseline reference voltage implemented using 8-bit R2R voltage DAC for channel 2.	25 [7:0]	8	8'h6E; 110	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ baseline_ RCG_ fine3	Fine adjustment of SOUT RCG output baseline reference voltage implemented using 8-bit R2R voltage DAC for channel 3.	26 [7:0]	8	8'h6E; 110	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ thresh_ mdpt_ VRG	We use PMOS input discriminators for all VRG set-ups. Thus, should only use -VSSD to 0V as full threshold range. This register sets that midpt value (while also allowing it to be adjustable if needed; implemented using 8- bit R2R Voltage DAC).	27 [7:0]	8	8'h80; 128	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_ thresh_ total_ VRG_ ch1	Set voltage threshold of VRG discriminator, which sets width of total int pulse for ch1 (implemented using 8-bit R2R Voltage DAC).	28 [7:0]	8	8'h73; 115	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ thresh_ partial_ VRG_ ch1	Set voltage threshold of VRG discriminator, which sets width of partial int pulse for ch1 (implemented using 8-bit R2R Voltage DAC).	29 [7:0]	8	8'h23; 35	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ thresh_ hold_ VRG_ ch1	Set voltage threshold of VRG discriminator, which sets width of HOLD pulse for ch1 (implemented using 8-bit R2R Voltage DAC).	30 [7:0]	8	8'h82; 130	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ thresh_ total_ VRG_ ch3	Set voltage threshold of VRG discriminator, which sets width of total int pulse for ch3 (implemented using 8-bit R2R Voltage DAC).	31 [7:0]	8	8'h73; 115	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.

v_ thresh_ partial_ VRG_ ch3	Set voltage threshold of VRG discriminator, which sets width of partial int pulse for ch3 (implemented using 8-bit R2R Voltage DAC).	32 [7:0]	8	8'h23; 35	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
v_ thresh_ hold_ VRG_ ch3	Set voltage threshold of VRG discriminator, which sets width of HOLD pulse for ch3 (implemented using 8-bit R2R Voltage DAC).	33 [7:0]	8	8'h82; 130	See Section B.5.1 for more details, including mapping of digital 8-bit values to fine-tuned voltage output.
i_ bias_ fast_ buffer	Control the bias current supplied for the FOUT fast buffers used for all channels. When all bits are 0, output current bias is 1mA. When all bits are 1, output current bias is 3.5mA.	34 [5:0]	5	8'h03; 3	The total bias current is set by 1mA + (num_ bits_ high * 0.5mA). The actual bits that are set to 1 to set a bias current doesn't matter.
i_ bias_ total_ int_ CSI_ ch0	Control bits for CSI delay line for total integration width for ch0.	35 [4:0]	5	8'h10; 16	See Section B.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.

i_ bias_ partial_ int_ CSI_ ch0	Control bits for CSI delay line for partial integration width for ch0.	36 [4:0]	5	8'h15; 21	See Section B.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
i_ bias_ hold_ CSI_ ch0	Control bits for CSI delay line for HOLD enable width for ch0.	37 [4:0]	5	8'h19; 25	See Section B.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
i_ bias_ total_ int_ CSI_ ch2	Control bits for CSI delay line for total integration width for ch2.	38 [4:0]	5	8'h10; 16	See Section B.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
i_ bias_ partial_ int_ CSI_ ch2	Control bits for CSI delay line for partial integration width for ch2.	39 [4:0]	5	8'h15; 21	See Section B.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.

i_ bias_ hold_ CSI_ ch2	Control bits for CSI delay line for HOLD enable width for ch2.	40 [4:0]	5	8'h19; 25	See Section B.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
i_ bias_ total_ int_VRG_ ch1	Total int VRG charging current value for ch1.	41 [4:0]	5	8'h0F; 15	See Section 6.1.5.2 for more details on the voltage ramp generator (VRG) circuit. See Section B.5.2 for mapping of digital 5-bit values to fine-tuned charging current output.
i_ bias_ partial_ int_VRG_ ch1	Partial int VRG charging current value for ch1.	42 [4:0]	5	8'h0F; 15	See Section 6.1.5.2 for more details on the voltage ramp generator (VRG) circuit. See Section B.5.2 for mapping of digital 5-bit values to fine-tuned charging current output.

i_ bias_ hold_ int_VRG_ ch1	HOLD line VRG charging current value for ch1.	43 [4:0]	5	8'h0F; 15	See Section 6.1.5.2 for more details on the voltage ramp generator (VRG) circuit. See Section B.5.2 for mapping of digital 5-bit values to fine-tuned charging current output.
i_ bias_ total_ int_VRG_ ch3	Total int VRG charging current value for ch3.	44 [4:0]	5	8'h0F; 15	See Section 6.1.5.2 for more details on the voltage ramp generator (VRG) circuit. See Section B.5.2 for mapping of digital 5-bit values to fine-tuned charging current output.
i_ bias_ partial_ int_VRG_ ch3	Partial int VRG charging current value for ch3.	45 [4:0]	5	8'h0F; 15	See Section 6.1.5.2 for more details on the voltage ramp generator (VRG) circuit. See Section B.5.2 for mapping of digital 5-bit values to fine-tuned charging current output.

i_ bias_ hold_ int_VRG_ ch3	HOLD line VRG charging current value for ch3.	46 [4:0]	5	8'h0F; 15	See Section 6.1.5.2 for more details on the voltage ramp generator (VRG) circuit. See Section B.5.2 for mapping of digital 5-bit values to fine-tuned charging current output.
i_ bias_ hold_ delay	current DAC for current starved inverter delay line for delay of hold line trigger. Sets the feedback delay before HOLD line is fed back into FOUT triggering logic to prevent race conditions (all channels).	47 [4:0]	5	8'h1B; 27	See Section B.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
i_ bias_ disc_ sout_ delay	current DAC for current starved inverter delay line for delay of SOUT final discriminator output logic. Sets a delay period during which final discriminator will not trigger HIGH (prevent spurious switching during initial integration period (all channels).	48 [4:0]	5	8'h19; 25	See Sections B.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.

i_ bias_ fout_ width_ trim	current DAC for current starved inverter delay line that sets the standard width of digital pulses going to parallel total, partial, HOLD one-shot generation circuits (all channels).	49 [4:0]	5	8'h19; 25	See Section B.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
i_ bias_ sout_ comp_ width	current DAC for current starved inverter delay line that sets the standard width of SOUT discriminator output (all channels).	50 [4:0]	5	8'h19; 25	See Section B.5.2 for more details, including mapping of digital 5-bit values to the fine-tuned current bias output.
sel_ std_ sout_ comp_ width	1: select std. width option for SOUT discriminator; 0: no std. width (all channels).	50 [5]	1	1'b0; 0	
tunable_ res_ total_ int	Adjust resistor value in total integration channel (one-hot) 8 MSBs, so corresponds to **int_ trim_ total [14:7]** (all channels).	51 [7:0]	8	8'h00; 0	See Section B.9 for more details.
tunable_ res_ msbs	Despite the name, 7 LSBs for tunable resistor of total integration stage. Corresponds to **int_ trim_ total [6:0]** (all channels).	52 [6:0]	7	8'h10; 16	See Section B.9 for more details.

tunable_ res_ subtr_ gain	Adjusts feedback resistor value (and thus gain) for subtraction stage output. Corresponds to **int_ trim_ subtr [5:0]**. (all channels).	53 [5:0]	6	8'h01; 1	See Section B.9 for more details.
tunable_ res_ RCG_ gain	Select gain of the RCG front end; for NMOS, gain goes from 2x-6x and for PMOS, gain goes from 1x-5x due to differences of transfer function for inverting vs non-inverting symm. OTA setup. Corresponds to **int_ trim_ RCG [2:0]**. (all channels).	54 [2:0]	3	8'h01; 1	This is **int_ trim_ RCG [2:0]**. **Spare0 [1:0]** is used as **int_ trim_ RCG [4:3]** due to mistranslation of how many bits were required for this register when writing digital core code.
digital_ testbus0_ sel	Select which FOUT triggering logic node is routed to digital_testbus0.	55 [4:0]	5	8'h00; 0	See Section B.6 for more details.
digital_ testbus0_ pd	1 to power down digital_testbus0.	55 [5]	1	1'b1; 1	
digital_ testbus1_ sel	Select which FOUT triggering logic node is routed to digital_testbus1.	56 [4:0]	5	8'h08; 8	See Section B.6 for more details.
digital_ testbus1_ pd	1 to power down digital_testbus1.	56 [5]	1	1'b1; 1	
digital_ testbus2_ sel	Select which FOUT triggering logic node is routed to digital_testbus2.	57 [4:0]	5	8'h10; 16	See Section B.6 for more details.
digital_ testbus2_ pd	1 to power down digital_testbus2.	57 [5]	1	1'b1; 1	

digital_ testbus3_ sel	Select which FOUT triggering logic node is routed to digital_testbus3.	58 [4:0]	5	8'h18; 24	See Section B.6 for more details.
digital_ testbus3_ pd	1 to power down digital_testbus3.	58 [5]	1	1'b1; 1	
sout_ comp_ out_ pd	1 to power down channel (SOUT PSD bit) LVDS output.	59 [3:0]	4	4'b0; 0	
lvds_ rx_ pd	1 to power down all other LVDS outputs across all channels.	59 [4]	1	1'b1; 1	See Section 6.2.6 for more details.
lvds_ high_ cm	1 to use higher LVDS common mode (might be needed depending on external LVDS IC used).	59 [5]	1	1'b1; 1	See Section 6.2.6 for more details.
lvds_ loopback_ enable	1 to loop RX back to TX (one channel).	59 [6]	1	1'b0; 0	See Section 6.2.6 for more details.
external_ trigger_ enable	1 to enable external trigger.	60 [0]	1	1'b0; 0	See Section 6.1.2 for more details.
cross_ trigger_ enable	1 to enable cross triggering functionality.	60 [1]	1	1'b0; 0	See Section 6.1.2 for more details.
current_ monitor	1 to enable current monitor; used to verify performance of internal 5-bit current DACs.	61 [0]	1	1'b0; 0	See Section B.7.2 for more details.
voltage_ monitor	1 to enable voltage monitor (one-hot); used to monitor eight internal DC voltages.	62 [7:0]	8	8'h00; 0	See Section B.7.1 for more details.

spare0	spare config. register	63 [7:0]	8	8'h00; 0	spare0 [1:0] is used as **int_trim_RCG [4:3]** due to mistranslation of required number of bits for register 54 when writing digital core code. See Section B.9 for more details.
spare1	spare config. register	64 [7:0]	8	8'h00; 0	No internal connections from digital core to rest of the chip.
spare2	spare config. register	65 [7:0]	8	8'h00; 0	No internal connections from digital core to rest of the chip.
spare3	spare config. register	66 [7:0]	8	8'h00; 0	No internal connections from digital core to rest of the chip.

Table B.3: All registers of the digital core for PSD_CHIP V2 are listed. Information on what the bits from each register control on-chip are provided along with their pre-programmed default values.

B.4 RBIAS Pads

The required resistor values are listed in Table B.4 for all RBIAS pads. A description of these pads is provided in Section 6.2.7.

RBIAS Pad Name	**Nominal Current Bias**	**Equivalent Nominal Resistor Value**
RBIAS_ANALOG	50 uA	$\sim$ 24.6 kΩ
RBIAS_LVDS	50 uA	$\sim$ 24.6 kΩ
RBIAS_RCG	5 mA	$\sim$ 171.0 Ω
RBIAS_CS	500 uA	$\sim$ 1.7 kΩ

Table B.4: List of all V2 RBIAS I/O pads with nominal current biases and the equivalent resistor values needed for each.

B.5 On Chip Voltage and Current DACs

B.5.1 Voltage DACs

There are several highly tunable DC voltage references and thresholds that need to be supplied to relevant parts of the SOUT and FOUT sub-chains. These voltages are set using a two stage coarse/fine tuning process. A coarse global voltage range is defined for all four channels for every tunable voltage. This is either set using a 3-bit string DAC or a standard 8-bit R2R DAC. For the 3-bit string DAC, setting the control bits is relatively simple. Setting the bits to 000 will set the positive reference voltage to +900 mV (VDDD) and negative reference voltage to +675 mV (VDDD - (VDDD - VSSD)/8). 001 sets the positive reference

range to +675 mV (VDDD - (VDDD - VSSD)/8) and negative reference range to +450 mV (VDDD - 2(VDDD - VSSD)/8) and so on. Thus, 111 sets the positive reference voltage to -675 mV (VDDD -7(VDDD - VSSD)/8) and the negative reference voltage to -900 mV (VSSD). The coarse global voltage range can span the full 1.8 V (nominal) power supply swing. This global range sets the positive and negative reference voltages for a subsequent 8-bit R2R DAC that fine-tunes the tunnable voltage for each respective channel. The equation to set the values for the control bits for the 8-bit R2R DACs is:

$$V_{out} = \frac{V_{ref;p} - V_{ref;n}}{256} * v_{fine} < 7:0 > [\text{decimal}] + V_{ref;n}, \tag{B.1}$$

where $v_{fine} < 7:0 >$ [decimal] is the decimal equivalent value for the 8 control bits that make up $v_{fine} < 7:0 >$. $V_{ref;p}$ and $V_{ref;n}$ are the selected positive and negative global reference range.

The fine-tuned voltage is buffered before being routed to the relevant internal circuitry. Table B.5 lists the tunable voltages for PSD_CHIP V2. The relevant digital core register is named in parenthesis (some of the names are identical).

Name (Digital Core Equiv.)	**Description**	**Max Val.**	**Min. Val**	**Step Size**	**Num. Bits**
V_BUFF_ REFERENCE (v_buff_ref_ global)	Set the global range for the integration stages DC baseline (to match baseline of buffer output).	900mV/ VDDD	-900mV/ VSSD	225mV/ (VDDD-VSSD)/8	3 (000: 675 mV to 900 mV; 111: -900 mV to -675 mV)

V_BUFF_ REFERENCE_ FINE [3:0] (v_buff_ref_fine [3:0])	Allows fine-tuning and selection of the DC baseline reference for the integration stages per channel.	225mV/ (VDDD - VSSD) /8	0V	0.88mV/ (Max. Val)/ 256	8 (see equation in Section B.5.1)
V_REF_COMP [3:0] (v_ref_comp [3:0])	Allows for channel level fine-tuning and selection of DC baseline for differential to single ended conversion of subtraction op amp stage output.	100mV (nominally)	0V (nominally)	0.39 mV/ (Max. - Min. Val.)/ 256	8 (see equation in Section B.5.1)
V_REF_ FINE_{P/N} (v_fine_ref{p/n})	Sets the global reference range for DC baseline of op amp that converts differential outputs of the subtraction op amp to single ended	900mV/ VDDD	-900mV/ VSSD	7.03 mV (over full-scale for each)	8 (see equation in Section B.5.1)
V_COMP_ THRESH_ SOUT_ GLOBAL (v_comp_ thresh_sout_ global)	Set the global range for the SOUT final discriminator threshold for all four channels	900mV /VDDD	-900mV /VSSD	225mV/ (VDDD -VSSD)/8	3 (000: 675mV - 900mV; 111: -900mV to -675mV)
V_COMP_ THRESH_SOUT_ FINE [3:0] (v_comp_ thresh_sout_ fine [3:0])	Allows for channel level fine-tuning and selection of SOUT discriminator threshold.	225mV/ (VDDD - VSSD) /8	0V	0.88mV/ (Max. Val.) /256	8 (see equation in Section B.5.1)
V_COMP_ THRESH_FOUT_ GLOBAL (v_comp_ thresh_fout_ global)	Set the global range for the FOUT discriminator threshold for all four channels	900mV / VDDD	-900mV /VSSD	225mV/ (VDDD -VSSD) /8	3 (000: 675mV- 900mV; 111: -900mV to -675mV)

V_COMP_ THRESH_FOUT_ FINE [3:0] (v_comp_ thresh_fout_ fine [3:0])	Allows for channel level fine-tuning and selection of FOUT discriminator threshold.	225mV/(VDD - VSSD) /8	0V	0.88mV/ (Max. Val.) /256	8 (see equation in Section B.5.1)
V_BASELINE_ RCG_GLOBAL (v_baseline_ RCG_global)	Global adjustment of SOUT RCG output baseline reference voltage	900mV / VDDD	-900mV / VSSD	225mV/ (VDDD -VSSD) /8	3 (000: 675mV-900mV; 111: -900mV to -675mV)
V_BASELINE_ RCG_FINE [3:0] (v_baseline_ RCG_fine [3:0])	Fine adjustment of SOUT RCG output baseline reference voltage	225mV/ (VDDD - VSSD) /8	0V	0.88mV/ (Max. Val) /256	8 (see equation in Section B.5.1)
V_THRESH_ MIDPT_VRG (v_thresh_ midpt_VRG)	Since we are using PMOS input discriminator for all VRG set ups, we should only work on using -VSSD to 0V as threshold range. This sets that midpoint value (and makes this value adjustable if needed)	900mV / VDDD	-900mV / VSSD	7.03mV/ (VDDD -VSSD) /256	8 (see equation in Section B.5.1)
V_THRESH_ TOTAL_VRG_ CH [3,1] (v_thresh_ total_VRG_ ch[3,1])	set voltage threshold of discriminator in VRG setup which sets width of total int pulse	0mV (nominally)	-900mV (VSSD)	3.52mV (nominally)	8 (see equation in Section B.5.1)

V_THRESH_PARTIAL_VRG_CH[3,1] (v_thresh_partial_VRG_ ch[3,1])	set voltage threshold of discriminator in VRG setup which sets width of partial int pulse	0mV (nominally)	-900mV (VSSD)	3.52mV (nominally)	8 (see equation in Section B.5.1)
V_THRESH_HOLD_VRG_CH[3,1] (v_thresh_hold_VRG_ ch[3,1])	set voltage threshold of discriminator in VRG setup which sets width of hold line pulse	0mV (nominally)	-900mV (VSSD)	3.52mV (nominally)	8 (see equation in Section B.5.1)

Table B.5: List of all programmable voltages and their functionality for PSD_CHIP V2. All voltages listed are generated through internal voltage DACs with the control bit values set through the digital core.

B.5.2 Current DACs

Several 5-bit current DACs are included as part of the PSD_CHIP design. They require as input, a LSB current. This LSB current, depending on the specific current DAC's role is either externally tunable or a fixed current input. The current DACs take the currents generated at either of RBIAS_TOTAL, RBIAS_PARTIAL or RBIAS_ANALOG pads as the input LSB. The LSB for the DACs that set delay (width) for the START_INT and STOP_PARTIAL_INT enable lines (for channels 0 and 2) are set externally using the current supplied (internally divided by x100) to the RBIAS_PARTIAL and RBIAS_TOTAL pads. All other CSI delay line (see Section 6.1.5.1 for further details) LSBs are fixed at 500 nA by a 100x internal

division of the nominal 50 uA RBIAS_ANALOG current input. In addition, several current DACs set the charging current for VRG-based programmable delay lines (see Section 6.1.5.2 for further details). These have a fixed LSB of 100 nA.

Depending on the specified 5-bit value, the output of a current DAC will be the input LSB current multiplied by the decimal-equivalent 5-bit value, with the caveat that control lines for the DAC are active LOW. Explicitly, the output current of the DAC is:

$$\text{Output Current } [A] = LSB[A] * (31 - \text{5-bit value [decimal]}). \tag{B.2}$$

Table B.6 lists all current DACs internal to PSD_CHIP V2. The number of CSI pairs used is different for each of the CSI delay lines on chip and is specified in the table. More detail is found in Section 6.1.5.1.

Name	**Description**	**Max Val.**	**Min. Val.**	**Step Size**	**External LSB?**
CURR_BIAS_ TOTAL_INT_ CSI_CH0	Supplied bias current for delay line that determines total integration length. 400 CS-inverter pairs in delay chain.	32uA	0A	variable b/w 200nA to 1uA	Yes; scaled down by 100 from external value set by RBIAS_ TOTAL pad.

CURR_BIAS_ TOTAL_INT_ CSI_CH2	Supplied bias current for delay line that determines total integration length. 400 CS-inverter pairs in delay chain.	32uA	0A	variable b/w 200nA to 1uA	Yes; scaled down by 100 from external value set by RBIAS_ TOTAL pad.
CURR_BIAS_ PARTIAL_INT_ CSI_CH0	Supplied bias current for delay line that determines partial integration length. 30 CS-inverter pairs in delay chain.	32uA	0A	variable b/w 200nA to 1uA	Yes; scaled down internally by 100 from external value set by RBIAS_ PARTIAL pad.
CURR_BIAS_ PARTIAL_INT_ CSI_CH2	Supplied bias current for delay line that determines partial integration length. 30 CS-inverter pairs in delay chain.	32uA	0A	variable b/w 200nA to 1uA	Yes; scaled down internally by 100 from external value set by RBIAS_ PARTIAL pad.
CURR_BIAS_ HOLD_CSI_CH0	Supplied bias current for delay line that determines length of the HOLD line. The length should always be >= total integration length. 400 CS-inverter pairs in delay chain.	15.5uA	0A	500nA	No

CURR_BIAS_ HOLD_CSI_CH2	Supplied bias current for delay line that determines length of the HOLD line. The length should always be >= total integration length. 400 CS-inverter pairs in delay chain.	15.5uA	0A	500nA	No
CURR_BIAS_ TO- TAL_VRG_CH1	Voltage ramp generator total int charging current bias value for ch1. Nominal is 1.5uA.	3.1uA	0A	100nA	No
CURR_BIAS_ TO- TAL_VRG_CH3	Voltage ramp generator total int charging current bias value for ch3. Nominal is 1.5uA.	3.1uA	0A	100nA	No
CURR_BIAS_ PAR- TIAL_VRG_CH1	Voltage ramp generator partial int charging current bias value for ch1. Nominal is 1.5uA.	3.1uA	0A	100nA	No
CURR_BIAS_ PAR- TIAL_VRG_CH3	Voltage ramp generator partial int charging current bias value for ch3. Nominal is 1.5uA.	3.1uA	0A	100nA	No
CURR_BIAS_ HOLD_VRG_CH1	Voltage ramp generator hold line charging current bias value for ch1. Nominal is 1.5uA.	3.1uA	0A	100nA	No
CURR_BIAS_ HOLD_VRG_CH3	Voltage ramp generator hold line charging current bias value for ch3. Nominal is 1.5uA.	3.1uA	0A	100nA	No

CURR_BIAS_ HOLD_DELAY	Supplied bias current for delay line that determines delay in feeding HOLD line back to RESET for SR latch to prevent re-triggering during ongoing active integration. The length should always be >= standard width of FOUT trigger pulses. 10 CS-inverter pairs in delay chain.	15.5uA	0A	500nA	No
CURR_BIAS_ SOUT_DELAY	Supplied bias current for delay line that determines delay in allowing SOUT final discriminator to trigger HIGH after start of integration. The length should always be >partial integration length by approx. 1uA. 10 CS-inverter pairs in delay chain.	16uA	500nA	500nA	No
CURR_BIAS_ FOUT_WIDTH_ TRIM	Supplied bias current for delay line that determines the standard width of FOUT trigger pulses to ensure clean input to one shot generation stages. 10 CS-inverter pairs in delay chain.	16uA	500nA	500nA	No

CURR_BIAS_ SOUT_COMP_ STD_WIDTH	Supplied bias current for delay line that determines the standard width of SOUT final discriminator output pulses to allow for uniform triggered HIGHs. 30 CS-inverter pairs in delay chain.	16uA	500nA	500nA	No

Table B.6: List of the various current DACs (and associated biases) for the respective CSI delay chains on PSD_CHIP V2.

B.6 Digital Test Buses

The digital test buses probe various points in the FOUT sub-chain. A detailed overview can be found in Section 6.2.3. Eight locations can be probed for each of the four channels. Thus, each digital test bus essentially functions as a 32x1 digital MUX. The buses can read out up to four probe locations for a given channel. They can also be used to read out the same location across all four channels. In normal operating mode, it is expected that each bus will read out the FOUT discriminator output (this will serve as the T0 pulse for PSD_CHIP). The eight probe nodes are listed in Table B.7 below. The control bits for digital_testbus0_sel[4:0] to digital_testbus3_sel[4:0] need to be set using the appropriate decimal equivalent values in this table in order to output the correct node.

Name	**Description**	**Dec. Equiv. of 5-bit Val. [CH0,CH1, CH2,CH3]**	**Additional Comments**
HOLD [3:0]	This pulse is generated along with total and partial integration one-shot enable pulses. It will disable the FOUT triggering logic during an ongoing active integration to prevent re-triggering. The length of the HOLD pulse should be >= START_INT pulse.	0,8,16,24	see Section B.5.2 for actual characteristics of pulse.
HOLD_delayed [3:0]	The HOLD pulse, except with implemented delay in order to ensure no race condition. The length of this delay should be >the standard width selected for the width trim stages.	1,9,17,25	see Section B.5.2 for actual characteristics of pulse.
START_INT [3:0]	This pulse is generated to enable total and partial integration of an incident pulse.	2,10,18,26	see Section B.5.2 for actual characteristics of pulse.
STOP_PARTIAL_INT [3:0]	This pulse is generated to enable and set partial integration length of incident pulse.	3,11,19,27	see Section B.5.2 for actual characteristics of pulse.
CROSS_TRIG_OUT [3:0]	This is the digital pulse that is generated by the FOUT discriminator once incident waveform is above threshold.	4,12,20,28	This signal should be nominally output during actual data-taking runs to serve as T0.
TRIGGERING_OUT [3:0]	This is the pulse post FOUT triggering block.	5,13,21,29	see Section B.5.2 for actual characteristics of pulse.

width_trim_stg1_out [3:0]	The output of the first of two width trim stages that ensures a standard width digital pulse input to the one-shot pulse generation stages.	6,14,22,30	see Section B.5.2 for actual characteristics of pulse.
width_trim_out [3:0]	The final output of the two width trim stages that ensures a standard width digital pulse input to the one-shot pulse generation stages.	7,15,23,31	see Section B.5.2 for actual characteristics of pulse.

Table B.7: List of all probe nodes that are selectable outputs through the four digital test buses.

B.7 Voltage and Current Monitors

B.7.1 Voltage Monitor

The internal voltage monitor functions equivalently to an 8x1 analog MUX. However, eight bits are needed to address the monitor since it works on the one-hot principle (i.e. only one bit can be HIGH at a time). All four baseline reference voltages for the associated SOUT integration stages and the SOUT and FOUT discriminator thresholds for channels 0 and 1 are selectable. These analog voltages are used to verify the performance of the on-chip coarse/fine voltage DACs set-up. The relevant digital bits in the digital core to set are **voltage_monitor** [7:0].

Name	Description	Dec. Equiv. One-Hot Val.
v_buff_bias0	Reference voltage for SOUT integration stages (channel 0)	0
v_buff_bias1	Reference voltage for SOUT integration stages (channel 1)	1
v_buff_bias2	Reference voltage for SOUT integration stages (channel 2)	2
v_buff_bias3	Reference voltage for SOUT integration stages (channel 3)	3
v_SOUT_ comp_thresh0	SOUT discriminator threshold for channel 0.	4
v_SOUT_ comp_thresh1	SOUT discriminator threshold for channel 1.	5
v_FOUT_ comp_thresh0	FOUT discriminator threshold for channel 0.	6
v_FOUT_ comp_thresh1	FOUT discriminator threshold for channel 1.	7

Table B.8: List of all voltage monitor outputs and the associated bit that needs to be HIGH to select a given output.

B.7.2 Current Monitor

There is only one possible current monitor output. This monitor's main purpose is to verify the performance of the on-board 5-bit current DACs that output bias currents for the various CSI-based delay lines. The associated bits in the digital core are **current_monitor** [7:0]. Although there are 8 bits, only **current_monitor** [0] is relevant. Setting the bit HIGH activates the current monitor output and LOW disables it. The current monitor output

must be terminated to -0.9 V (VSSD) through an appropriate resistor. The relevant signal is the current through this resistor.

B.8 VRG Delay Line Details

The VRG-based programmable delay lines are quite flexible. Table B.9 lists the nominal maximum delay possible for various current biases (charging currents) and associated capacitor value. Also listed is the nominal step size at each combination. See Section 6.1.5.2 for further details.

Current Bias Selected	**Nominal Max. Delay Possible**	**Capacitor Value**	**Nominal Step Size**
0.1uA	36.00 us	4pF	140.6 ns
1.5uA	2.40 us	4pF	9.4 ns
3.1uA	1.16 us	4pF	4.5 ns
0.1uA	9.00 us	1pF	35.2 ns
1.5uA	0.60 us	1pF	2.3 ns
3.1uA	0.29 us	1pF	1.1 ns

Table B.9: A list of the nominal maximum delay possible for each current bias (charging current) and capacitor value pair. Nominal step sizes are also given for the pairs.

B.9 Tunable Resistors

B.9.1 RCG Tunable Resistor

The SOUT front-end symmetric OTA for V2 features a programmable feedback resistor providing adjustable gain. The selectable values are listed below in Table B.10. See Section 6.1.6 for further details.

Digital Bit (one-hot)	**Resistor Value**	**Gain (PMOS/NMOS)**
int_trim_RCG [0]	5 kΩ	1/2
int_trim_RCG [1]	10 kΩ [+5 kΩ]	2/3
int_trim_RCG [2]	15 kΩ [+5 kΩ]	3/4
int_trim_RCG [3]	20 kΩ [+5 kΩ]	4/5
int_trim_RCG [4]	25 kΩ [+5 kΩ]	5/6

Table B.10: List of selectable resistor values in the SOUT front-end stage and the equivalent adjustable gain value for active PMOS/NMOS input.

B.9.2 Total Integration Tunable Resistor

The total integration stage features a programmable resistor. The selectable values are listed below in Table B.11. See Section 6.1.7.1 for further details.

B.9.3 Subtraction Stage Tunable Resistor

The subtraction stage features a programmable feedback resistor. The selectable values are listed below in Table B.12. See Section 6.1.8 for further details.

Digital Bit (one-hot)	**Resistor Value**	**Equiv. Partial/Total Ratio**
int_trim_total [0]	75 kΩ	1
int_trim_total [1]	90 kΩ [+15 kΩ]	0.83
int_trim_total [2]	100 kΩ [+10 kΩ]	0.75
int_trim_total [3]	110 kΩ [+10 kΩ]	0.68
int_trim_total [4]	120 kΩ [+10 kΩ]	0.63
int_trim_total [5]	125 kΩ [+5 kΩ]	0.6
int_trim_total [6]	130 kΩ [+5 kΩ]	0.58
int_trim_total [7]	135 kΩ [+5 kΩ]	0.56
int_trim_total [8]	140 kΩ [+5 kΩ]	0.54
int_trim_total [9]	150 kΩ [+10 kΩ]	0.5
int_trim_total [10]	160 kΩ [+10 kΩ]	0.47
int_trim_total [11]	180 kΩ [+20 kΩ]	0.42
int_trim_total [12]	200 kΩ [+20 kΩ]	0.38
int_trim_total [13]	225 kΩ [+25 kΩ]	0.33
int_trim_total [14]	250 kΩ [+25 kΩ]	0.3

Table B.11: List of selectable resistor values in the long integration stage for V2 and the equivalent effective threshold PSD ratio.

Digital Bit (one-hot)	**Resistor Value**	**Equivalent Gain**
int_trim_subtr [0]	25kΩ	x1
int_trim_subtr [1]	50kΩ [+25kΩ]	x2
int_trim_subtr [2]	75kΩ [+25kΩ]	x3
int_trim_subtr [3]	125kΩ [+50kΩ]	x5
int_trim_subtr [4]	250kΩ [+125kΩ]	x10
int_trim_subtr [5]	375kΩ [+125kΩ]	x15

Table B.12: List of resistor values for selecting the gain of the buffer that converts the differential subtraction stage output to a single-ended signal in V2.

Bibliography

[1] N. Aghanim et al. "Planck 2018 results". In: *Astronomy & Astrophysics* 641 (Sept. 2020), A6.

[2] Pavel Motloch. "Testing consistency of Ω b h^2 in the Planck data". In: *Physical Review D* 101.12 (2020), p. 123509.

[3] Jin Koda et al. "Approximately a thousand ultra-diffuse galaxies in the Coma cluster". In: *The Astrophysical Journal Letters* 807.1 (2015), p. L2.

[4] Matthias Bartelmann. "Gravitational lensing". In: *Classical and Quantum Gravity* 27.23 (2010), p. 233001.

[5] Julio F Navarro. "The structure of cold dark matter halos". In: *Symposium-international astronomical union*. Vol. 171. Cambridge University Press. 1996, pp. 255–258.

[6] Tjeerd S van Albada et al. "Distribution of dark matter in the spiral galaxy NGC 3198". In: *The Astrophysical Journal* 295 (1985), pp. 305–313.

[7] Douglas Clowe et al. "A direct empirical proof of the existence of dark matter". In: *The Astrophysical Journal* 648.2 (2006), p. L109.

[8] Richard G Strom. "How was atomic HI (λ= 21 cm line) in space discovered?" In: *International Journal of Modern Physics: Conference Series*. Vol. 23. World Scientific. 2013, pp. 472–477.

[9] Wayne Hu and Scott Dodelson. "Cosmic microwave background anisotropies". In: *Annual Review of Astronomy and Astrophysics* 40.1 (2002), pp. 171–216.

[10] Bruce Bassett and Renée Hlozek. "Baryon acoustic oscillations". In: *Dark energy: observational and theoretical approaches* (2010), p. 246.

[11] Marc Schumann. "Direct Detection of WIMP Dark Matter: Concepts and Status". In: *J. Phys. G* 46.10 (2019), p. 103003.

[12] Gianfranco Bertone and Dan Hooper. "History of dark matter". In: *Reviews of Modern Physics* 90.4 (2018), p. 045002.

[13] G. Arcadi et al. "The Waning of the WIMP? A Review of Models, Searches, and Constraints". In: *Eur. Phys. J. C* 78.3 (2018), pp. 1–57.

[14] Howard Baer et al. "Dark matter production in the early Universe: beyond the thermal WIMP paradigm". In: *Physics Reports* 555 (2015), pp. 1–60.

[15] Teresa Marrodan Undagoitia and Ludwig Rauch. "Dark matter direct-detection experiments". In: *Journal of Physics G: Nuclear and Particle Physics* 43.1 (2015), p. 013001.

[16] Jatan Buch, Shing Chau John Leung, and JiJi Fan. "Using Gaia DR2 to constrain local dark matter density and thin dark disk". In: *Journal of Cosmology and Astroparticle Physics* 2019.04 (2019), p. 026.

[17] MG Aartsen et al. "Improved limits on dark matter annihilation in the Sun with the 79-string IceCube detector and implications for supersymmetry". In: *Journal of Cosmology and Astroparticle Physics* 2016.04 (2016), p. 022.

[18] Nahee Park. "Performance of the VERITAS experiment". In: *arXiv preprint arXiv:1508.07070* (2015).

[19] Michele Doro. "A review of the past and present MAGIC dark matter search program and a glimpse at the future". In: *arXiv preprint arXiv:1701.05702* (2017).

[20] Aharonian Abramowski et al. "Search for dark matter annihilation signatures in HESS observations of Dwarf Spheroidal Galaxies". In: *Physical Review D* 90.11 (2014), p. 112012.

[21] Dan Hooper and Lisa Goodenough. "Dark matter annihilation in the Galactic Center as seen by the Fermi Gamma Ray Space Telescope". In: *Physics Letters B* 697.5 (2011), pp. 412–428.

[22] J Buckley et al. "Cosmic Frontier indirect dark matter detection working group summary". In: *arXiv preprint arXiv:1310.7040* (2013).

[23] Morad Aaboud et al. "Search for dark matter and other new phenomena in events with an energetic jet and large missing transverse momentum using the ATLAS detector". In: *Journal of High Energy Physics* 2018.1 (2018), pp. 1–53.

[24] Emidio Gabrielli et al. “Dark photons and resonant monophoton signatures in Higgs boson decays at the LHC”. In: *Physical Review D* 90.5 (2014), p. 055032.

[25] David Curtin et al. “Illuminating dark photons with high-energy colliders”. In: *Journal of High Energy Physics* 2015.2 (2015), pp. 1–45.

[26] Joss Bland-Hawthorn and Ortwin Gerhard. “The galaxy in context: structural, kinematic and integrated properties”. In: *arXiv preprint arXiv:1602.07702* (2016).

[27] G. Barello, Spencer Chang, and Christopher A. Newby. “A model independent approach to inelastic dark matter scattering”. In: *Phys. Rev. D* 90 (9 Nov. 2014), p. 094027.

[28] Gerard Jungman, Marc Kamionkowski, and Kim Griest. “Supersymmetric dark matter”. In: *Physics Reports* 267.5-6 (1996), pp. 195–373.

[29] Richard H Helm. “Inelastic and elastic scattering of 187-Mev electrons from selected even-even nuclei”. In: *Physical Review* 104.5 (1956), p. 1466.

[30] Helen Quinn. “The CP puzzle in the strong interactions”. In: *arXiv preprint hep-ph/0110050* (2001).

[31] Masahiro Kawasaki et al. “Oscillons of Axion-Like Particle: Mass distribution and power spectrum”. In: *Journal of Cosmology and Astroparticle Physics* 2021.01 (2021), p. 061.

[32] Stephen J Asztalos et al. “SQUID-based microwave cavity search for dark-matter axions”. In: *Physical review letters* 104.4 (2010), p. 041301.

[33] Alexey Boyarsky et al. “Sterile neutrino dark matter”. In: *Progress in Particle and Nuclear Physics* 104 (2019), pp. 1–45.

[34] E Aprile and T Doke. “Liquid xenon detectors for particle physics and astrophysics”. In: *Reviews of Modern Physics* 82.3 (2010), p. 2053.

[35] Carl Eric Dahl. “The physics of background discrimination in liquid xenon, and first results from Xenon10 in the hunt for WIMP dark matter”. PhD thesis. Princeton University, 2009.

[36] BJ Mount et al. “LUX-ZEPLIN (LZ) technical design report”. In: *arXiv preprint arXiv:1703.09144* (2017).

[37] A Manzur et al. "Scintillation efficiency and ionization yield of liquid xenon for monoenergetic nuclear recoils down to 4 keV". In: *Physical Review C* 81.2 (2010), p. 025808.

[38] WH Lippincott et al. "Scintillation time dependence and pulse shape discrimination in liquid argon". In: *Physical Review C* 78.3 (2008), p. 035801.

[39] DS Akerib et al. "The lux-zeplin (lz) experiment". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 953 (2020), p. 163047.

[40] A Liam Fitzpatrick et al. "The effective field theory of dark matter direct detection". In: *Journal of Cosmology and Astroparticle Physics* 2013.02 (2013), p. 004.

[41] J Aalbers et al. "First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment". In: *arXiv preprint arXiv:2207.03764* (2022).

[42] Nikhil Anand, A Liam Fitzpatrick, and WC Haxton. "Model-independent WIMP scattering responses and event rates: A Mathematica package for experimental analysis". In: *arXiv preprint arXiv:1308.6288* (2013).

[43] Billy Boxer. *LUX-ZEPLIN Dark Matter Searches: Development of Optical and Source Based Calibration Systems*. The University of Liverpool (United Kingdom), 2021.

[44] Elena Aprile et al. "Effective field theory search for high-energy nuclear recoils using the XENON100 dark matter detector". In: *Physical review D* 96.4 (2017), p. 042004.

[45] DS Akerib et al. "Effective field theory analysis of the first LUX dark matter search". In: *Physical Review D* 103.12 (2021), p. 122005.

[46] DS Akerib et al. "Constraints on effective field theory couplings using 311.2 days of LUX data". In: *Physical Review D* 104.6 (2021), p. 062005.

[47] LW Kastens et al. "Calibration of a liquid xenon detector with Kr 83 m". In: *Physical Review C* 80.4 (2009), p. 045809.

[48] M. Szydagis et al. *Noble Element Simulation Technique*. Version v2.3.7. May 2022.

[49] DS Akerib et al. "Tritium calibration of the LUX dark matter experiment". In: *Physical Review D* 93.7 (2016), p. 072009.

[50] DS Akerib et al. "Low-energy (0.7-74 keV) nuclear recoil calibration of the LUX dark matter experiment using DD neutron scattering kinematics". In: *arXiv preprint arXiv:1608.05381* (2016).

[51] DS Akerib et al. “Investigation of background electron emission in the LUX detector”. In: *Physical Review D* 102.9 (2020), p. 092004.

[52] DS Akerib et al. “Results from a search for dark matter in the complete LUX exposure”. In: *Physical review letters* 118.2 (2017), p. 021303.

[53] XENON Collaboration et al. “Dark matter search results from a one ton-year exposure of xenon1t”. In: *Physical review letters* 121.11 (2018), p. 111302.

[54] Peter Sorensen. “Anisotropic diffusion of electrons in liquid xenon with application to improving the sensitivity of direct dark matter searches”. In: *arXiv preprint arXiv:1102.2865* (2011).

[55] DS Akerib et al. “The LUX-ZEPLIN (LZ) radioactivity and cleanliness control programs”. In: *The European Physical Journal C* 80.11 (2020), pp. 1–52.

[56] David Woodward, LUX Collaboration, LZ Collaboration, et al. “Simulations of external backgrounds at SURF for the LUX and LZ experiments”. In: *APS April Meeting Abstracts*. Vol. 2018. 2018, J10–005.

[57] J Aalbers et al. “Cosmogenic production of Ar 37 in the context of the LUX-ZEPLIN experiment”. In: *Physical Review D* 105.8 (2022), p. 082004.

[58] Andrzej K Drukier, Katherine Freese, and David N Spergel. “Detecting cold dark-matter candidates”. In: *Physical Review D* 33.12 (1986), p. 3495.

[59] JD Lewin and PF Smith. “Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil”. In: *Astroparticle Physics* 6.1 (1996), pp. 87–112.

[60] Ibles Olcina Samblas and Henrique Araújo. “A statistical framework for the characterisation of WIMP dark matter with the LUX-ZEPLIN experiment”. PhD thesis. Imperial College London, 2019.

[61] Ilka Antcheva et al. “ROOT—A C++ framework for petabyte data storage, statistical analysis and visualization”. In: *Computer Physics Communications* 180.12 (2009), pp. 2499–2512.

[62] Lorenzo Moneta et al. “The roostats project”. In: *arXiv preprint arXiv:1009.1003* (2010).

[63] Teal Pershing et al. "Calibrating the scintillation and ionization responses of xenon recoils for high-energy dark matter searches". In: *arXiv preprint arXiv:2207.08326* (2022).

[64] Brian Lenardo et al. "Measurement of the ionization yield from nuclear recoils in liquid xenon between 0.3–6 keV with single-ionization-electron sensitivity". In: *arXiv preprint arXiv:1908.00518* (2019).

[65] FD Brooks. "Development of organic scintillators". In: *Nuclear Instruments and Methods* 162.1-3 (1979), pp. 477–505.

[66] M Sénoville et al. "Neutron-γ discrimination with organic scintillators: Intrinsic pulse shape and light yield contributions". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 971 (2020), p. 164080.

[67] Joanna Iwanowska-Hanke et al. "Comparative study of large samples (2" $\times$ 2") plastic scintillators and EJ309 liquid with pulse shape discrimination (PSD) capabilities". In: *Journal of Instrumentation* 9.06 (2014), P06014.

[68] Natalia Zaitseva et al. "Scintillation properties of solution-grown trans-stilbene single crystals". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 789 (2015), pp. 8–15.

[69] Marc Lavi Ruch, Marek Flaska, and Sara A Pozzi. "Pulse shape discrimination performance of stilbene coupled to low-noise silicon photomultipliers". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 793 (2015), pp. 1–5.

[70] Natalia Zaitseva et al. "Plastic scintillators with efficient neutron/gamma pulse shape discrimination". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 668 (2012), pp. 88–93.

[71] M Grodzicka-Kobylka et al. "Fast neutron and gamma ray pulse shape discrimination in EJ-276 and EJ-276G plastic scintillators". In: *Journal of Instrumentation* 15.03 (2020), P03030.

[72] John Betteley Birks. *The theory and practice of scintillation counting: International series of monographs in electronics and instrumentation*. Vol. 27. Elsevier, 2013.

[73] Natalia Zaitseva et al. “Pulse shape discrimination in impure and mixed single-crystal organic scintillators”. In: *IEEE Transactions on Nuclear Science* 58.6 (2011), pp. 3411–3420.

[74] *PSD Plastic Scintillator EJ276*. Rev. July 2021. Eljen Technology. July 2021.

[75] *Scintinel Stilbene for Fast Neutron Detection*. Rev. July 2020. Inrad Optics. July 2020.

[76] MI Pinilla-Orjuela et al. “Performance Evaluation of Pulse Shape Discrimination Capable Organic Scintillators for Space Applications”. In: *arXiv preprint arXiv:2111.08845* (2021).

[77] *Introduction to the Silicon Photomultiplier (SiPM)*. AND9770/D. Rev. 7. OnSemi. July 2018.

[78] Fabio Acerbi and Stefan Gundacker. “Understanding and simulating SiPMs”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 926 (2019), pp. 16–35.

[79] Fang Liu et al. “Characterization of a Mass-Produced SiPM at Liquid Nitrogen Temperature for CsI Neutrino Coherent Detectors”. In: *Sensors* 22.3 (2022), p. 1099.

[80] *Silicon Photomultipliers (SiPM), High PDE and Timing Resolution Sensors in a TSV Package; J-Series SiPM Sensors*. MICROJSERIES/D. Rev. 7. OnSemi. Aug. 2021.

[81] C. Dulieu et al. “Reference decay data, tools and applications”. In: *2009 1st International Conference on Advancements in Nuclear Instrumentation, Measurement Methods and their Applications*. 2009, pp. 1–4.

[82] Kyle Weinfurther et al. “Model-based design evaluation of a compact, high-efficiency neutron scatter camera”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 883 (2018), pp. 115–135.

[83] John EM Goldsmith, Mark D Gerling, and James S Brennan. “A compact neutron scatter camera for field deployment”. In: *Review of Scientific Instruments* 87.8 (2016), p. 083307.

[84] Glenn F Knoll. *Radiation detection and measurement*. John Wiley & Sons, 2010.

[85] United States Nuclear Regulatory Commission. “Neutron Sources”. In: *Basic Health Physics* (2010).

[86] Isao Murata et al. “Neutron and gamma-ray source-term characterization of AmBe sources in Osaka University”. In: *Progress in nuclear science and technology* 4 (2014), pp. 345–348.

[87] Angelo Rivetti. *CMOS: front-end electronics for radiation sensors*. CRC press, 2018.

[88] Eduard Säckinger. “The transimpedance limit”. In: *IEEE transactions on circuits and systems I: regular papers* 57.8 (2010), pp. 1848–1856.

[89] Luis B Oliveira, Carlos M Leitão, and M Medeiros Silva. “Noise performance of a regulated cascode transimpedance amplifier for radiation detectors”. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 59.9 (2012), pp. 1841–1848.

[90] Manuel de Medeiros Silva and Luis B Oliveira. “Regulated common-gate transimpedance amplifier designed to operate with a silicon photo-multiplier at the input”. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 61.3 (2013), pp. 725–735.

[91] Steve Ritz et al. *Building for Discovery: Strategic Plan for US Particle Physics in the Global Context*. 2014.

[92] Jonathan L. Ouellet et al. “First Results from ABRACADABRA-10 cm: A Search for Sub-μeV Axion Dark Matter”. In: *Phys. Rev. Lett.* 122.12 (2019), p. 121802. arXiv: 1810.12257 [hep-ex].

[93] M. Battaglieri et al. “US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report”. In: 2017. arXiv: 1707.04591.

[94] J.I. Read. “The Local Dark Matter Density”. In: *J. Phys.* G41 (2014), p. 063101.

[95] P.F. de Salas. “Dark matter local density determination based on recent observations”. In: *J. Phys. Conf. Ser.* 1468 (2020), p. 012020.

[96] B. Holdom. “Two U(1)’s and $\in$ charge shifts”. In: *Phys. Lett. B* 166.2 (Jan. 1986), p. 196.

[97] T.G. Rizzo. “Kinetic mixing and portal matter phenomenology”. In: *Phys. Rev. D* 99.11 (June 2019), p. 115024.

[98] A. Drlica-Wagner et al. “Probing the Fundamental Nature of Dark Matter with the Large Synoptic Survey Telescope”. In: *arXiv e-prints* (Feb. 2019). arXiv: 1902.01055.

[99] T. M. P. Tait. “Dark matter candidates: status and perspectives”. In: *34th International Cosmic Ray Conference (ICRC2015)*. Vol. 34. International Cosmic Ray Conference. July 2015, 5, p. 5.

[100] G. Alonso-Álvarez et al. "On the Wondrous Stability of ALP Dark Matter". In: *arXiv e-prints*, arXiv:1911.07885 (Nov. 2019), arXiv:1911.07885. arXiv: 1911.07885 [hep-ph].

[101] M. Tanabashi et al. "Review of Particle Physics". In: *Phys. Rev. D* 98 (3 Aug. 2018), p. 030001.

[102] P. W. Graham et al. "Parametrically enhanced hidden photon search". In: *Physical Review D* 90.7 (Oct. 2014).

[103] A. Wagner et al. "Search for hidden sector photons with the ADMX detector". In: *Phys. Rev. Lett.* 105.17 (2010), p. 171801.

[104] N. Du et al. "Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment". In: *Phys. Rev. Lett.* 120 (15 Apr. 2018), p. 151301.

[105] T. Braine et al. "Extended Search for the Invisible Axion with the Axion Dark Matter Experiment". In: *Phys. Rev. Lett.* 124 (10 Mar. 2020), p. 101303.

[106] Brian D. Fields et al. "Big-Bang Nucleosynthesis after Planck". In: *Journal of Cosmology and Astroparticle Physics* 2020.03 (Mar. 2020), pp. 010–010.

[107] Ryan J Cooke, Max Pettini, and Charles C Steidel. "One percent determination of the primordial deuterium abundance". In: *The Astrophysical Journal* 855.2 (2018), p. 102.

[108] J. L. Feng et al. *Planning the Future of U.S. Particle Physics (Snowmass 2013): Chapter 4: Cosmic Frontier*. 2014. arXiv: 1401.6085 [hep-ex].

[109] Liliya LR Williams and Paul L Schechter. "The quest for the golden lens". In: *Astronomy & Geophysics* 38.5 (1997), pp. 10–14.

[110] J Anthony Tyson, Greg P Kochanski, and Ian P Dell'Antonio. "Detailed mass map of CL 0024+ 1654 from strong lensing". In: *The Astrophysical Journal* 498.2 (1998), p. L107.

[111] I. Stern. *ADMX Status*. 2017. arXiv: 1612.08296 [physics.ins-det].

[112] Masahiro Kawasaki et al. "Oscillons of axion-like particle: mass distribution and power spectrum". In: *Journal of Cosmology and Astroparticle Physics* 2021.01 (Jan. 2021), pp. 061–061.

[113] John Preskill, Mark B Wise, and Frank Wilczek. "Cosmology of the invisible axion". In: *Physics Letters B* 120.1-3 (1983), pp. 127–132.

[114] Laurence F Abbott and P Sikivie. "A cosmological bound on the invisible axion". In: *Physics Letters B* 120.1-3 (1983), pp. 133–136.

[115] Michael Dine and Willy Fischler. "The not-so-harmless axion". In: *Physics Letters B* 120.1-3 (1983), pp. 137–141.

[116] Joel R. Primack. *Whatever Happened to Hot Dark Matter?* 2001. arXiv: `astro-ph/0112336 [astro-ph]`.

[117] David N Spergel and Paul J Steinhardt. "Observational evidence for self-interacting cold dark matter". In: *Physical review letters* 84.17 (2000), p. 3760.

[118] J. Choi et al. "CAPP-8TB: Axion dark matter search experiment around 6.7 ev". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1013 (Oct. 2021), p. 165667.

[119] LM Capparelli et al. "Axion-like particle searches with sub-THz photons". In: *Physics of the Dark Universe* 12 (2016), pp. 37–44.

[120] S. J. Asztalos et al. "SQUID-Based Microwave Cavity Search for Dark-Matter Axions". In: *Physical Review Letters* 104.4 (Jan. 2010).

[121] P. Sikivie. "Experimental Tests of the "Invisible" Axion". In: *Phys. Rev. Lett.* 51 (16 Oct. 1983), pp. 1415–1417.

[122] Gary Steigman and Michael S. Turner. "Cosmological Constraints on the Properties of Weakly Interacting Massive Particles". In: *Nucl. Phys. B* 253 (1985), pp. 375–386.

[123] Jonah Herzog-Arbeitman et al. "Empirical Determination of Dark Matter Velocities Using Metal-Poor Stars". In: *Physical Review Letters* 120.4 (Jan. 2018).

[124] Francesco Capozzi et al. "Global constraints on absolute neutrino masses and their ordering". In: *Physical Review D* 95.9 (2017), p. 096014.

[125] Marco Drewes. "The phenomenology of right handed neutrinos". In: *International Journal of Modern Physics E* 22.08 (2013), p. 1330019.

[126] A. A. Aguilar-Arevalo et al. "Significant Excess of Electron like Events in the MiniBooNE Short-Baseline Neutrino Experiment". In: *Physical Review Letters* 121.22 (Nov. 2018).

[127] Peter B Denton. "Sterile Neutrino Searches with MicroBooNE: Electron Neutrino Disappearance". In: *arXiv preprint arXiv:2111.05793* (2021).

[128] Sebastian Böser et al. “Status of light sterile neutrino searches”. In: *Progress in Particle and Nuclear Physics* 111 (Mar. 2020), p. 103736.

[129] John N Bahcall and Raymond Davis Jr. “Solar neutrinos: a scientific puzzle”. In: *Science* 191.4224 (1976), pp. 264–267.

[130] Y. Fukuda et al. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Physical Review Letters* 81.8 (Aug. 1998), pp. 1562–1567.

[131] Q. R. Ahmad et al. “Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by ^8B Solar Neutrinos at the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 87 (7 July 2001), p. 071301.

[132] Laura Baudis and. “DARWIN dark matter WIMP search with noble liquids”. In: 375.1 (July 2012), p. 012028.

[133] J Hart et al. “The telescope system of the macho program”. In: *Publications of the Astronomical Society of the Pacific* 108.720 (1996), p. 220.

[134] C. Alcock et al. “MACHO Project Limits on Black Hole Dark Matter in the 1–30 $10^{-6}M_{\odot}$ Range”. In: *The Astrophysical Journal* 550.2 (Apr. 2001), pp. L169–L172.

[135] C. Alcock et al. “EROS and MACHO Combined Limits on Planetary-Mass Dark Matter in the Galactic Halo”. In: *The Astrophysical Journal* 499.1 (May 1998), pp. L9–L12.

[136] KIM GRIEST. “The Search for the Dark Matter: WIMPs and MACHOs”. In: *Annals of the New York Academy of Sciences* 688.1 (June 1993), pp. 390–407.

[137] Andrew Gould. “Extending the MACHO search to about 10^6 solar masses”. In: *The Astrophysical Journal* 392 (1992), pp. 442–451.

[138] Claudio Corianò and Paul H Frampton. “Does CMB Distortion Disfavour Intermediate Mass Dark Matter?” In: *arXiv preprint arXiv:2012.13821* (2020).

[139] Katherine Freese, Mariangela Lisanti, and Christopher Savage. “Annual modulation of dark matter: a review”. In: *arXiv preprint arXiv:1209.3339* (2012).

[140] Virginia Trimble. “History of Dark Matter in Galaxies”. In: *Planets, Stars and Stellar Systems: Volume 5: Galactic Structure*. Ed. by Gerard Gilmore. New York, NY: Springer, 2013, pp. 1093–1116.

[141] Georges Lemaitre. “Republication of: A homogeneous universe of constant mass and increasing radius accounting for the radial velocity of extra-galactic nebulae”. In: *General Relativity and Gravitation* 45.8 (2013), pp. 1635–1646.

[142] Edwin Hubble. “A relation between distance and radial velocity among extra-galactic nebulae”. In: *Proceedings of the national academy of sciences* 15.3 (1929), pp. 168–173.

[143] Fritz Zwicky. “On the Masses of Nebulae and of Clusters of Nebulae”. In: *The Astrophysical Journal* 86 (1937), p. 217.

[144] Harold Irving Ewen and Edward Mills Purcell. “Observation of a Line in the Galactic Radio Spectrum: Radiation from Galactic Hydrogen at 1,420 Mc./sec.” In: *Nature* 168.4270 (1951), pp. 356–356.

[145] V.C. Rubin, Jr. Ford W.K., and N. Thonnard. “Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa -¿ Sc.” In: *APJL* 225 (Nov. 1978), pp. L107–L111.

[146] Paris Pismis. “ON THE WAVY NATURE OF ROTATION CURVES IN GALAXIES”. In: *Sbc* 9600 (1965), p. 700.

[147] Morton S Roberts. “A high-resolution 21-CM hydrogen-line survey of the Andromeda nebula”. In: *The Astrophysical Journal* 144 (1966), p. 639.

[148] Vera C Rubin and W Kent Ford Jr. “Rotation of the Andromeda nebula from a spectroscopic survey of emission regions”. In: *The Astrophysical Journal* 159 (1970), p. 379.

[149] Kenneth C Freeman. “On the disks of spiral and S0 galaxies”. In: *The Astrophysical Journal* 160 (1970), p. 811.

[150] Andrew R Liddle. “Acceleration of the universe”. In: *New Astronomy Reviews* 45.4-5 (2001), pp. 235–253.

[151] Charles L Bennett et al. “Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: final maps and results”. In: *The Astrophysical Journal Supplement Series* 208.2 (2013), p. 20.

[152] Charles H Lineweaver. “Inflation and the cosmic microwave background”. In: *The New Cosmology*. World Scientific, 2005, pp. 31–65.

[153] David Valls-Gabaud. "The conceptual origins of gravitational lensing". In: *AIP Conference Proceedings*. Vol. 861. 1. American Institute of Physics. 2006, pp. 1163–1163.

[154] Albert Einstein. "Lens-like action of a star by the deviation of light in the gravitational field". In: *Science* 84.2188 (1936), pp. 506–507.

[155] Tilman Sauer. "Nova Geminorum 1912 and the origin of the idea of gravitational lensing". In: *Archive for history of exact sciences* 62.1 (2008), pp. 1–22.

[156] Dennis Walsh, Robert F Carswell, and Ray J Weymann. "0957+ 561 A, B: twin quasistellar objects or gravitational lens?" In: *Nature* 279.5712 (1979), pp. 381–384.

[157] Sam Palermo. *ECEN 474/704 Lab 7 Handout*.

[158] Peter F Smith and John David Lewin. "Dark matter detection". In: *Physics Reports* 187.5 (1990), pp. 203–280.

[159] Peter F Bloser et al. "Scintillator gamma-ray detectors with silicon photomultiplier readouts for high-energy astronomy". In: *UV, X-Ray, and Gamma-Ray Space Instrumentation for Astronomy XVIII*. Vol. 8859. International Society for Optics and Photonics. 2013, 88590A.

[160] Rhys M Preston, John E Eberhardt, and James R Tickner. "Neutron-gamma pulse shape discrimination using organic scintillators with silicon photomultiplier readout". In: *IEEE Transactions on nuclear science* 61.4 (2014), pp. 2410–2418.

[161] Viacheslav A Li et al. "A prototype for SANDD: A highly-segmented pulse-shape-sensitive plastic scintillator detector incorporating silicon photomultiplier arrays". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 942 (2019), p. 162334.

[162] Lorenzo Fabris. "Novel readout design criteria for SiPM-based radiation detectors". In: (2016).

[163] Manuel D Rolo et al. "A 64-channel ASIC for TOFPET applications". In: *2012 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*. IEEE. 2012, pp. 1460–1464.

[164] Motaz M Elbadry et al. "Design technique for regulated cascode transimpedance amplifier using Gm/ID methodology". In: *Microelectronics Journal* 95 (2020), p. 104676.

[165] W Cheng et al. "A mixed-signal large dynamic range front-end ASIC for high capacitance detectors". In: *Journal of Instrumentation* 14.08 (2019), P08013.

[166] BYT Kamath, Robert G Meyer, and Paul R Gray. "Relationship between frequency response and settling time of operational amplifiers". In: *IEEE Journal of Solid-State Circuits* 9.6 (1974), pp. 347–352.

[167] Melinda Sweany et al. *Design and Evaluation of a Pixelated PSD-capable Scintillator Detector with SiPM Readout.* Tech. rep. Sandia National Lab.(SNL-CA), Livermore, CA (United States); Sandia National ..., 2019.

[168] Marvin L Roush, MA Wilson, and William F Hornyak. "Pulse shape discrimination". In: *Nuclear Instruments and Methods* 31.1 (1964), pp. 112–124.

[169] Azaree T Lintereur et al. *Neutron and gamma ray pulse shape discrimination with polyvinyltoluene.* Tech. rep. Pacific Northwest National Lab.(PNNL), Richland, WA (United States), 2012.

[170] Cláudio Frederico Pascoal da Silva. "Dark matter searches with LUX". In: *arXiv preprint arXiv:1710.03572* (2017).

[171] E Bodnia et al. "The electric field dependence of single electron emission in the PIXeY two-phase xenon detector". In: *Journal of Instrumentation* 16.12 (2021), P12015.

[172] Shahram Hatefi Hesari, Mohammad Aminul Haque, and Nicole McFarlane. "A Comprehensive Survey of Readout Strategies for SiPMs Used in Nuclear Imaging Systems". In: *Photonics*. Vol. 8. 7. Multidisciplinary Digital Publishing Institute. 2021, p. 266.

[173] J Liu et al. "Neutron calibration sources in the Daya Bay experiment". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 797 (2015), pp. 260–264.

[174] Rafael F Lang et al. "A 220Rn source for the calibration of low-background experiments". In: *Journal of Instrumentation* 11.04 (2016), P04004.

[175] NP Zaitseva et al. "Recent developments in plastic scintillators with pulse shape discrimination". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 889 (2018), pp. 97–104.

[176] Richard S Woolf et al. "Comparing the response of PSD-capable plastic scintillator to standard liquid scintillator". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 784 (2015), pp. 80–87.

[177] MP Taggart, C Payne, and PJ Sellin. "Neutron-gamma discrimination via PSD plastic scintillator and SiPMs". In: *Journal of Physics: Conference Series*. Vol. 763. 1. IOP Publishing. 2016, p. 012007.

[178] M Grodzicka-Kobylka et al. "Comparison of detectors with pulse shape discrimination capability for simultaneous detection of gamma-rays, slow and fast neutrons". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1019 (2021), p. 165858.

[179] Glenn R Jocher et al. "miniTimeCube as a neutron scatter camera". In: *AIP Advances* 9.3 (2019), p. 035301.

[180] Joshua Braverman et al. "Single-volume neutron scatter camera for high-efficiency neutron imaging and spectroscopy". In: *arXiv preprint arXiv:1802.05261* (2018).

[181] Nick Mascarenhas et al. "Development of a neutron scatter camera for fission neutrons". In: *2006 IEEE Nuclear Science Symposium Conference Record*. Vol. 1. IEEE. 2006, pp. 185–188.

[182] Micah Folsom. "A Compact Neutron Scatter Camera Using Optical Coded-Aperture Imaging". In: (2020).

[183] Nicholas Mascarenhas et al. "Results with the neutron scatter camera". In: *IEEE Transactions on Nuclear Science* 56.3 (2009), pp. 1269–1273.