ORNL/TM-2023/3175

Characterization of the ORNL SiPM PSD ASIC's Basic Properties

James T. Matta
Paul B. Rose Jr.
Lorenzo Fabris

Oct. 2023

OAK RIDGE National Laboratory

ORNL IS MANAGED BY UT-BATTELLE LLC FOR THE US DEPARTMENT OF ENERGY

ORNL/TM-2023/3175

ORNL Physics Division

Characterization of the ORNL SiPM PSD ASIC's Basic Properties

James T. Matta
Paul B. Rose Jr.
Lorenzo Fabris

October 2023

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831
managed by
UT-Battelle LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABBREVIATIONS

ORNL	Oak Ridge National Laboratory
PSD	Pulse Shape Discrimination
SiPM	Silicon Photomultiplier
ASIC	Application Specific Integrated Circuit
FoM	Figure of Merit
TAC	Time to Amplitude Converter
FPGA	Field Programmable Gate Array
ADC	Analog to Digital Converter
PCB	Printed Circuit Board
COTS	Commercial Off The Shelf
LDO	Low-Dropout (Regulator)
RMS	Root Mean Square
FWHM	Full Width at Half Maximum
RF	Radio Frequency

ACKNOWLEDGMENTS

This material is based upon work supported by the Defense Threat Reduction Agency under CONTRACT/VEHICLE 0000-Z340-19.
This research was supported by the Office of Defense Nuclear Nonproliferation R&D in the National Nuclear Security Administration (NNSA), US Department of Energy under contract DE-AC05-00OR22725 with Oak Ridge National Laboratory, managed and operated by UT-Battelle, LLC.

ABSTRACT

ORNL has developed an ASIC intended for reading out arrays of Silicon Photomultipliers coupled to scintillator materials that exhibit neutron/gamma pulse shape discrimination. The ASIC's basic properties: power consumption, pulse shape discrimination, linearity, and timing resolution have all been measured and the results are presented here. As will be shown, while the ASIC has some points that further iteration could improve upon, it has achieved its stated design goals in its first design and fabrication cycle. This is a remarkable feat, as only 26% of ASIC development projects achieve success in their first silicon fabrication[1].

1. INTRODUCTION

The Oak Ridge National Laboratory (ORNL) Silicon Photomultiplier (SiPM) Pulse Shape Discrimination (PSD) Application Specific Integrated Circuit (ASIC) is a hybrid analog/digital ASIC developed to read 64 SiPMs coupled to scintilators that exhibit neutron / gamma ray discrimination by excesses of light in the tail of the pulse. The hybrid design's purely analog front-end allows pulse shape discrimination while minimizing power consumption as full waveform digitization is rendered unnecessary. After the development of testing boards for the ASIC and control firmware for the FPGA controlling the system [2] and after the execution of a study on the ASIC's analog shapers [3], tests of the ASIC's outputs when coupled to ADCs and read out in the usual fashion through the FPGA were devised and have been conducted. These tests show that the ASIC meets or exceeds all of its design goals, despite the presence of a few points that could be improved with iteration. With a first fabrication cycle success rate of only 26 % [1] among ASIC development projects, this is quite remarkable.

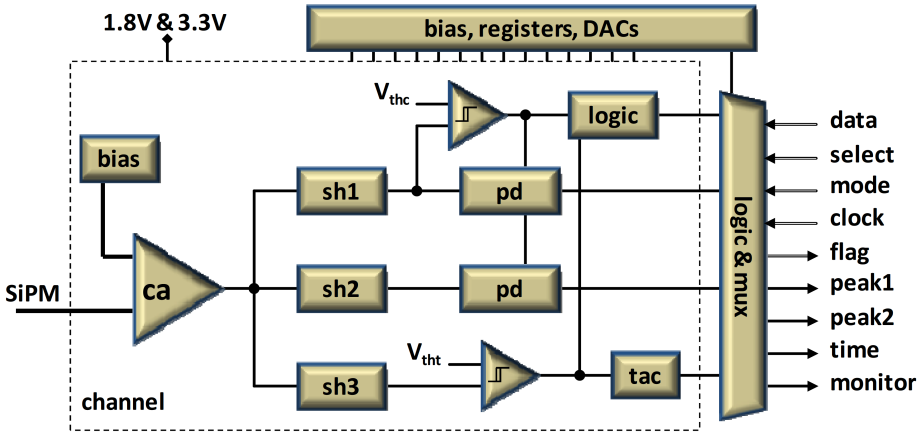

Figure 1. A block diagram of a single ASIC input channel. "ca" is the input stage current amplifier. "bias" is the input bias used to adjust the current amplifier's power consumption. "sh1" is the slow shaper, "sh2" is the fast shaper, and "sh3" is the timing shaper. The two "pd" are peak detect and hold circuits. The two "⎍" are leading edge discriminators. "tac"is the Time to Amplitude Converter (TAC). The "logic" block is the trigger flag and readout logic.

2. ASIC OVERVIEW

Fig. 1 shows a block diagram of the analog front-end of each SiPM readout channel of the ASIC. The SiPM's unbiased side is directly coupled to the ASIC, this allows the ASIC to individually trim the SiPM bias voltages (and thus SiPM gains) by 0.9 V to 2.5 V. As seen in Fig. 1 the current pulse from the SiPM enters a current amplifier (amplifications of, nominally, $\times 1$, $\times 2$, $\times 4$, and $\times 8$ available) and is subsequently passed to three different semi-gaussian shapers.

Each shaper can provide further amplification (with the same selectable gains as the input stage) with variable peaking times. The peak height of each shaper's output is proportional to the total charge that entered the front end amplifier. The slow shaper can have a peaking time of 90 ns, 160 ns, 230 ns, or 300 ns chosen. The fast shaper can have a peaking time of 35 ns, 50 ns, 65 ns, or 80 ns selected. The timing shaper, used to start the channel's TAC, can have a peaking time of 15 ns or 30 ns set. The TACs for each channel have selectable ranges of 500 ns, 1000 ns, 2000 ns, or 4000 ns.

To acquire data the ASIC is placed in acquisition mode. In this mode the ASIC's channels operate independently and send a trigger flag to the central logic when both the slow and timing shaper discriminators have triggered and the peaks of the slow and fast shapers have been detected. The central logic in turn sends a trigger flag to the controlling Field Programmable Gate Array (FPGA) when one or more channels have sent flags to it. To readout the data, the ASIC is transitioned to readout mode. This transition also stops any channel TACs that have been started ramping, providing a value proportional to the time between the channel's trigger and the start of readout mode. For each channel triggered the FPGA then reads the channel number directly from the ASIC as well. The peak heights of the slow shaper and fast shaper peaks as well as the TAC value are analog outputs from the ASIC which are read by the FPGA using separate Analog to Digital Converters (ADCs). Subsequently, the ASIC is returned to acquisition mode to continue the readout of light signals.

3. TEST SETUP OVERVIEW

Several Printed Circuit Boards (PCBs) were built around the ASIC to allow its use and to test its properties. Fig. 2 shows the first of these PCBs, the ASIC support board. This board hosts the ASIC itself, analog buffers for the 4 analog outputs the ASIC produces, high density connectors of the same type as those used for OnSemi SiPM arrays, and pins to connect the support board to a host board. To ease routing from the ASIC to the high density connectors the ASIC input channels are not mapped to corresponding SiPM array channels but instead are connected on the shortest paths possible. The mapping of SiPM channel to ASIC channel is captured in Table 1. It is anticipated that this board (for 3 mm SiPM arrays) or a variant of it (for 6 mm SiPM arrays) will be used in any future development such as large panels, as it is small enough to tile the 3 mm arrays and allows for easy replacements of the ASICs or SiPM arrays in case of problems.

Fig. 3 shows the other PCBs built and purchased for testing the ASIC. To the right of the image is the Commercial Off The Shelf (COTS) FPGA development board (an ALINX XC7A200T FPGA Development Board containing a Xilinx Artix-7 XC7A200T-2FBG484 FPGA). Connected to the top edge of the development board is a USB cable which, in conjunction with a USB to UART adapter chip on the development board, allows $200,000$ byte per second control and readout of the FPGA from a computer. Connected directly to the development board's two forty-pin headers on its left side is the logic level translation board. This board allows the eleven LVDS, one 1.8 V, and two 2.5 V CMOS signals from the ASIC and its ADCs to be translated to the 3.3 V CMOS logic signals needed by the FPGA (due to the

development board's designer's choices for FPGA IO bank power).

In the middle of the picture we see shielded ribbon cables transferring the digital signals to and from the logic level translator board to the ASIC testing board. On the left of the picture we see the test board for the ASIC testing board with the ASIC support board plugged into it. This board serves several purposes analog signal conversion, power supply and filtering, and logic signal buffering.

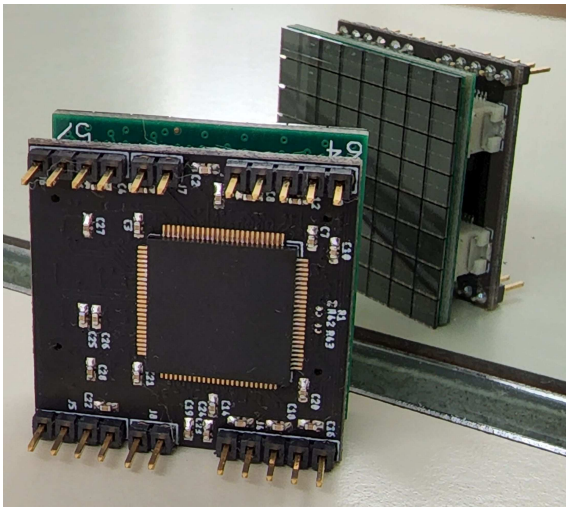

Figure 2. The ASIC support board with SiPM array. In the foreground, the ASIC soldered to its support board and the pins to transfer SiPM bias, ASIC power, digital signals, and analog signals to and from the support board. In the background, the 8×8 array of OnSemi J-Series 3 mm SiPMs, coupled to the support board.

Table 1. Mapping of SiPM Manufacturer Channel Number to ASIC Readout Channel Number

SiPM Ch.	ASIC Ch.	SiPM Ch.	ASIC Ch.	SiPM Ch.	ASIC Ch.	SiPM Ch.	ASIC Ch.
1	39	17	33	33	1	49	6
2	41	18	35	34	3	50	8
3	43	19	37	35	10	51	11
4	45	20	49	36	15	52	14
5	48	21	53	37	19	53	18
6	52	22	57	38	23	54	22
7	56	23	63	39	31	55	26
8	59	24	61	40	29	56	28
9	38	25	32	41	0	57	5
10	40	26	34	42	2	58	7
11	42	27	36	43	4	59	9
12	44	28	47	44	13	60	12
13	46	29	51	45	17	61	16
14	50	30	55	46	21	62	20
15	54	31	60	47	25	63	24
16	58	32	62	48	30	64	27

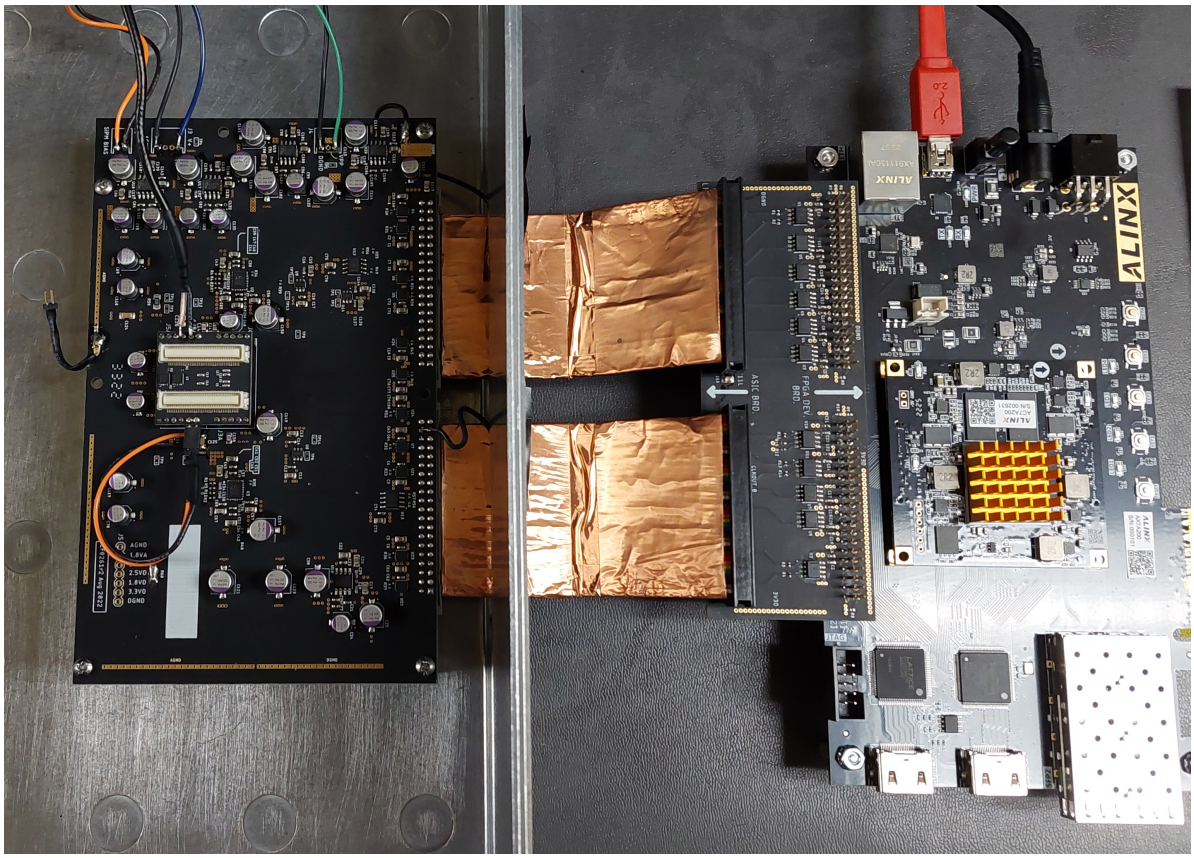

Figure 3. The ASIC testing setup. On the left, the ASIC testing board in a metal project box (sans lid), with the ASIC support board plugged into it (without a coupled array of SiPMs). On the right, the FPGA development board with logic level translator connected to it. In the middle, the shielded ribbon cables connecting the ASIC testing board and logic level translator board passing through a slot cut in the project box.

The three power inputs serve the SiPM bias, analog electronics, and digital electronics. Using a pair of low noise Low Drop-Out (LDO) regulators it converts the 5 V power supplied for analog electronics to the 3.3 V power needed by the ASIC to trim SiPM channels and by the ADCs' analog input stages as well the 1.8 V needed by the ASIC to power the other analog components. Using another pair of LDO regulators the board converts the separate 5 V power supplied for digital electronics to the 1.8 V required by the ASIC's digital components and the 2.5 V power required for the ADCs' digital output. The SiPM bias is provided directly to the ASIC support board where it is filtered and then provided to the SiPM arrays. The aforementioned two ADCs hosted in the testing board, are Analog Devices LTC2323-14 Dual, 14-bit + sign, five mega-sample per second ADCs. Each ADC converts two of the four analog outputs from the ASIC that are routed from the analog buffers on the support board through the pins to the testing board. Finally, the testing board provides digital buffers for the logic signals traveling to and from its components to reduce the noise that could be produced by the ASIC and ADCs having to drive signals strongly enough to cross the ribbon cables.

4. POWER CONSUMPTION

An important consideration for any detector readout electronics is power consumption. The current state of the art for reading arrays of SiPMs, a modern, high channel density 125 MSps waveform digitizer like the CAEN VX2740 draws 1.1 A at 12 V, 2.7 A at 5 V, and 4.9 A at 3.3 V [4] giving a total power consumption of 42.87 W. The VX2740 supports 64 channels being read out simultaneously (much like the ORNL SiPM PSD ASIC) giving a power dissipation per channel of 669.8 mW per channel. Therefore a primary design goal set for the ASIC was to achieve one tenth or less the power of current discrete designs.

4.1 METHODOLOGY

Measurement of the ASIC's power consumption relies on the fact that the current supplied to an LDO regulator is approximately equal to the current it supplies to the components it is powering. Because the power supply providing 5 V analog and digital power to the board allows for measurement of the supplied current with milliamp resolution (tenths of a milliamp resolution for the analog channel), measurement of the current supplied to the ASIC becomes quite simple. First, with the supply off, the ASIC support board is unplugged from the testing board. Then the supply is turned on and the analog and digital supply currents are recorded. Then the supply is turned back off and the ASIC support board is plugged back into the testing board. Finally, the power supply is turned back on and the analog and digital supply currents are recorded again.

The difference in the recorded supply currents is then the currents being delivered by the LDO regulators to the ASIC (and the two low power two-channel analog buffers that are also located on the ASIC support board). From the supply currents, the power dissipated by the ASIC can be quite accurately calculated when combined with knowledge of what voltage those currents are supplied at due to the relationship $P = I \cdot V$ where P is the instantaneous power consumption of a component, I is the current supplied to the component, and V is the voltage drop across the component.

4.2 RESULTS

Without the ASIC support board attached, the analog rail of the power supply was supplying 23.6 mA and the digital rail provided 670 mA. After reattaching the ASIC support board the current supplied by the analog rail was 237.2 mA and the digital rail was supplying 682 mA. The difference of these numbers gives that the ASIC and two analog voltage buffers on the ASIC support board are supplied with 12 mA on the digital rail and 213.6 mA on the analog rail.

As the only digital voltage supplied to the ASIC support board is 1.8 V we can conclude that the digital side of the ASIC consumes $1.8\ \text{V} \cdot 12\ \text{mA} = 21.6\ \text{mW}$ of power. Two different analog voltages are supplied to the ASIC support board. Both the ASIC and the analog buffers draw from the 1.8 V analog rail while only the ASIC draws from the 3.3 V analog rail. To produce the most conservative estimate of the power consumption possible the analog buffers are assumed to draw no current at all. Further, as the total current drawn from the analog rail (the sum of the currents supplied by the 3.3 V and 1.8 V LDO regulators) is less than the sum of the nominal current draws (216 mA, 66 mA and 150 mA for the 3.3 V and 1.8 V rails respectively), the deficit of current is assumed to belong to the 1.8 V rail. With the assumption that the ASIC's analog side draws 66 mA at 3.3 V and 147.6 mA at 1.8 V we get that the most conservative estimate for the ASIC's analog power consumption is $1.8\ \text{V} \cdot 147.6\ \text{mA} + 3.3\ \text{V} \cdot 66\ \text{mA} = 483.48\ \text{mW}$ of power.

Summing the digital and analog power dissipation together, the total power consumption of the ASIC is

505.1 mW. This yields a power dissipation per input channel of 7.89 mW. However, the ASIC cannot be read by itself, it has four analog outputs that need to be read using ADCs. The ADCs used in the current test system draw, nominally 38 mW per channel[5]. As four channels of the ADCs are in use, the total power consumed by the ADCs is assumed to be 152 mW. This gives a total power consumption of the ASIC plus necessary support components to be 656 mW. In stark contrast to the CAEN digitizer's 669.8 mW per channel, the ASIC only consumes 10.27 mW per channel, more than a factor of 65 smaller.

5. PSD

One of the primary goals of the ASIC is to allow the discrimination between neutrons and gamma rays from the amount of excess charge present in the tail of the current pulse from the SiPM. As stated earlier, this is enabled by having a short peaking time semi-Gaussian shaper (fast shaper) and a long peaking time semi-Gaussian shaper (slow shaper) providing, in effect, the early and total integrals of the pulse. From these, the customary PSD parameter is calculated and the ASIC's PSD performance is demonstrated.

5.1 SETUP

To measure the PSD yielded by the ASIC, an adapter board was used to link the 3 mm SiPM array connectors to the style of connector used for 6 mm SiPM arrays. The adapter board is a simple adapter

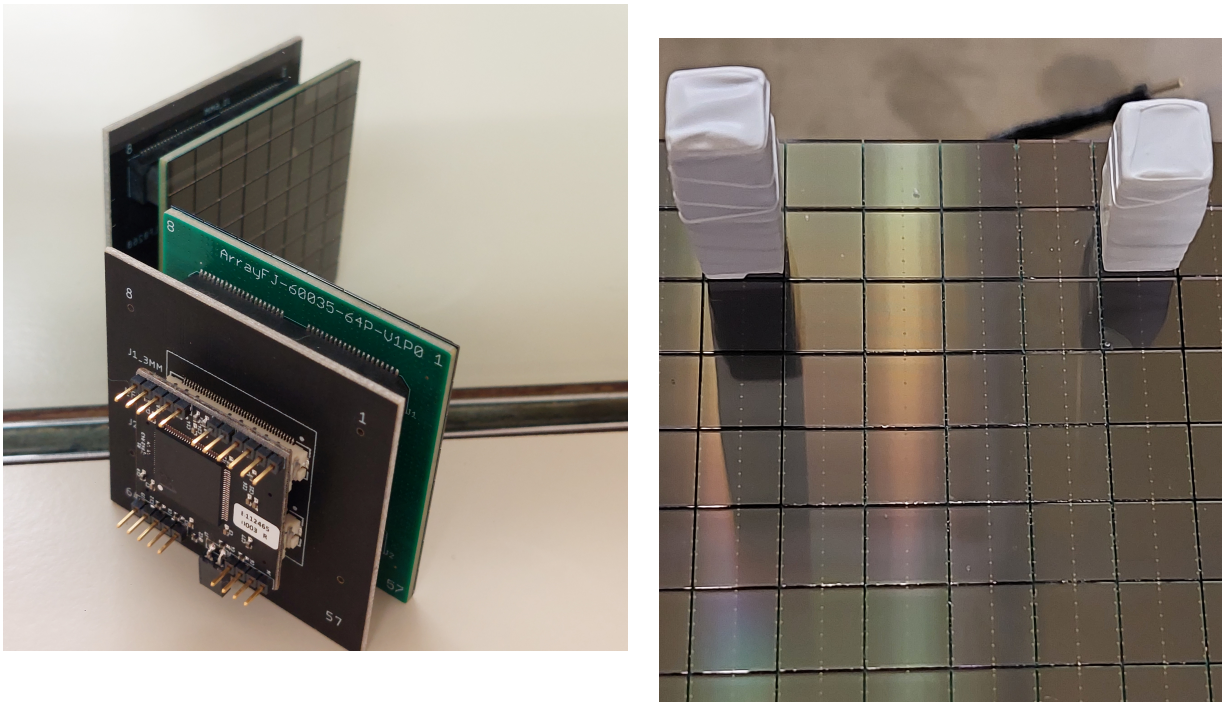

Figure 4. The *right* image shows the two stilbene scintillator crystals atop ASIC channels 8 **(left) and** 40 **(right). The *left* image shows the ASIC support board with large array adapter board and** 6 **mm SiPM array.** In the foreground, the ASIC soldered to its support board and the pins to transfer power and signals to and from the support board. In the background, the 8×8 array of OnSemi J-Series 6 mm SiPMs, coupled to the large array adapter board. In between the two boards the large array adapter board can be seen.

PCB with the female connectors appropriate to the 8×8 OnSemi J-Series 6 mm SiPM array on one side and the other side has the male connectors for the 3 mm variant of the array. The adapter board can be seen in the left image of Fig. 4 which shows the ASIC support board, the adapter board, and the 8×8 array of 6 mm SiPMs.

Then, without optical grease, two stilbene crystals 6 mm $\times$ 6 mm $\times$ 10 mm and 6 mm $\times$ 6 mm $\times$ 15 mm were placed on SiPM Channels 50 and 10 respectively (corresponding ASIC channels 8 and 40), as can be seen in the right image of Fig. 4. The choice to eschew optical grease was out of practicality for setup modifications and future changes, despite the increased light collection optical grease provides (and improved performance characteristics that improvement yields). With each setup change optical grease needs to be wiped off and reapplied, over time the grease wicks into the area between between the teflon wrapping of the crystal. This degrades light collection and necessitates unwrapping, cleaning and rewrapping crystals, which is a painstaking process with significant chances of breaking the crystals. The reason for using two single crystals is that it was desired to minimize changes required between this test setup and the setup used for timing resolution in Section 7.2. Subsequently the metal lid of the project box was closed and the entire setup covered in several layers of light blocking black felt to ensure light tightness. Afterwards sources were placed on top of the felt, approximately centered over the scintillators within the project box.

Next, measurements were made with a 100 mCi AmBe source to choose slow shaper and fast shaper peaking time parameters that give the best PSD performance. The PSD parameter was calculated using Eq. 1.

$$PSD = \frac{Slow - Fast}{Slow - P} \tag{1}$$

Where PSD is the PSD parameter, $Slow$ is the slow shaper peak value, $Fast$ is the fast shaper peak value, and P is the pedestal subtracted from the slow shaper peak value (since the peak value outputted is the height of the shaper's pulse plus some baseline).

For each combination of peaking times, uncalibrated 2 dimensional PSD spectra (PSD parameter vs. slow shaper peak value) were produced using a pedestal parameter (P) of 1767 (the analog baseline measured using the output monitor) and projections on the PSD parameter axis for various slow shaper peak value ranges. The projections have peaks associated with both neutron and gamma-ray interactions which are fit with Gaussian shapes. For each set of fits the Figure of Merit (FoM, defined as the difference in peak centroid locations divided by the sum of the peaks Full Widths at Half Maximum (FWHM)) is calculated. Then the peaking time parameters for the slow and fast shapers were chosen to maximize the FoM for as many of the projections as possible. The chosen values were 300 ns for the slow shaper peaking time and 65 ns for the fast shaper peaking time.

Afterwards, using the data from the test run with the chosen peaking times, different pedestal values were tested. It was found that a pedestal value of 1200 produced the best FoMs for that dataset and that has been used subsequently. Subsequently, a series of measurements using ^{137}Cs and ^{22}Na sources is run to determine the overall, fast shaper, and slow shaper input gains. The chosen values for the input, slow shaper, and fast shaper gains were chosen to be $\times$8, $\times$1, and $\times$1 respectively.

5.2 ENERGY CALIBRATION

Afterwards, the same sources were used in longer runs to produce spectra for energy calibrating the resulting spectra. Example uncalibrated ^{22}Na spectra can be seen in Fig. 5. The sharp "divots" in the

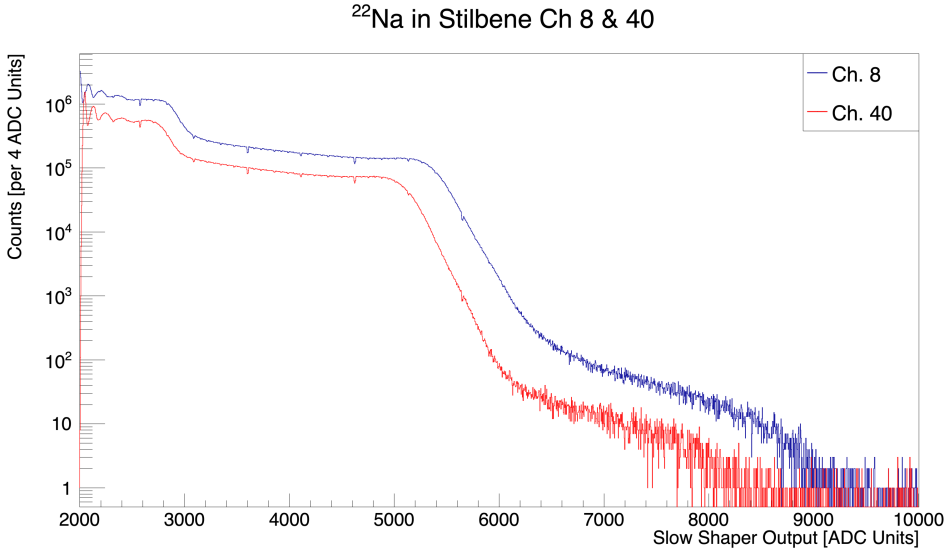

Figure 5. Uncalibrated ^{22}Na spectra for ASIC channels 8 **(blue) and** 40 **(red).** The height differences are due to the differences in crystal volumes and differences in distance between each crystal and the source.

spectra at matching channel numbers are due to the ADC channel that reads the slow shaper analog output having bins whose relative voltage widths are too small. This can be surmised from the positions occurring at identical locations between channels and from the fact that the number of channels separating the divots are powers of two. The cause of the oscillatory behavior on the left that does not quite match between channels is currently unknown, though, as will be seen in Fig. 6, the oscillations match after energy calibration. This hints that it is some sort of common mode effect, either outside the ASIC on the input side, or possibly injected into each channel's input by the ASIC's substrate. Though these oscillations are concerning, they do not significantly affect PSD performance / FoM until below 400 keVee at the gains chosen. As many setups struggle to obtain good PSD below $\sim$ 200 keVee this increased threshold is not terrible. Further, as will be shown later, from overlays of γ-ray PSD plots on AmBe PSD plots these oscillations seem to be approximately constant and an energy dependent PSD cut at low energies could significantly improve PSD performance.

To calibrate the energy scale of the spectra, first the energies of the Compton edges are calculated using the formula in Eq. 2.

$$E_{dep} = E_\gamma - E_{\gamma'} = E_\gamma \frac{\frac{2E_\gamma}{m_e c^2}}{1 + \frac{2E_\gamma}{m_e c^2}} \tag{2}$$

Where E_{dep} is the energy depositted in the scintillator (i.e. location of the Compton edge), E_γ is the incoming γ-ray energy, $E_{\gamma'}$ is the outgoing γ-ray energy, m_e is the mass of the electron, and c is the speed of light in vacuum. Applying this to each source γ ray yields the relationships shown in the first three columns of Table 2.

Table 2. Compton edge locations and their fitted slow shaper peak value locations used for calibration, listed by source and relevant γ ray.

Sources & Energies			Channel 8		Channel 40	
Isotope	γ-ray Energy. [keV]	Edge Energy [keV]	Edge Loc.	Err.	Edge Loc.	Err.
^{22}Na	511.0	340.67	2914.86	0.05	2831.56	0.06
^{137}Cs	661.657	477.34	3386.42	1.11	3285.81	0.97
^{22}Na	1274.537	1061.71	5415.78	0.22	5162.30	0.17

To determine the edge locations each Compton edge in the spectrum was fit with the formula in Eq. 3

$$Counts(x) = \frac{N}{2} * Erfc(\frac{x-\mu}{\sqrt{2}*\sigma}) + A \tag{3}$$

Here x is the slow shaper peak value, $Counts(x)$ is the number of counts as a function of x, N is a height or normalization parameter for the edge, $Erfc$ is the complementary error function, μ is the compton edge location, σ is the width of the edge "smearing" due to the scintillator's intrinsic resolution, and A is the constant component of the background. Each edge was fit using the ROOT Framework's [6] histogram fitting with the fit ranges restricted relatively close to the edges themselves The edge locations and their errors are shown in the last four columns of Table 2.

The edge locations for both channels were fit with linear energy calibrations of the form $E = A + B * x$ where E is the calibrated energy, x is the slow shaper peak value, and A and B are the parameters of the fit. With calibration constants in hand, applying the energy calibration to Fig. 5, yields Fig. 6

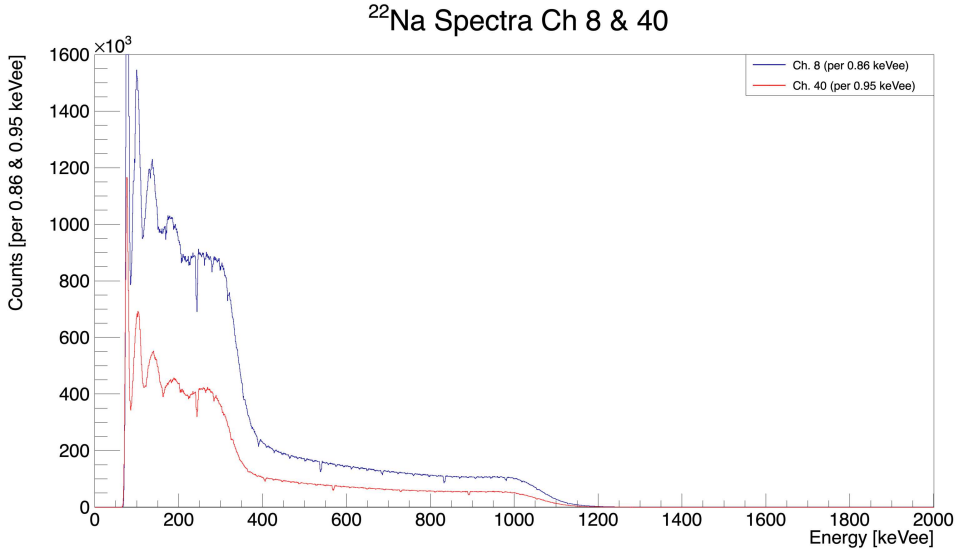

Figure 6. Calibrated ^{22}Na spectra for ASIC channels 8 **(blue) and** 40 **(red).** As for the uncalibrated spectrum the difference in spectral intensities is due to the differences in crystal volumes and differences in distance between the crystal and the source. The previously seen oscillatory behavior on the left now does match in location hinting at some sort of common mode source. The odd bin widths match the energy calibration preventing jaggedness from uneven mapping without applying smoothing.

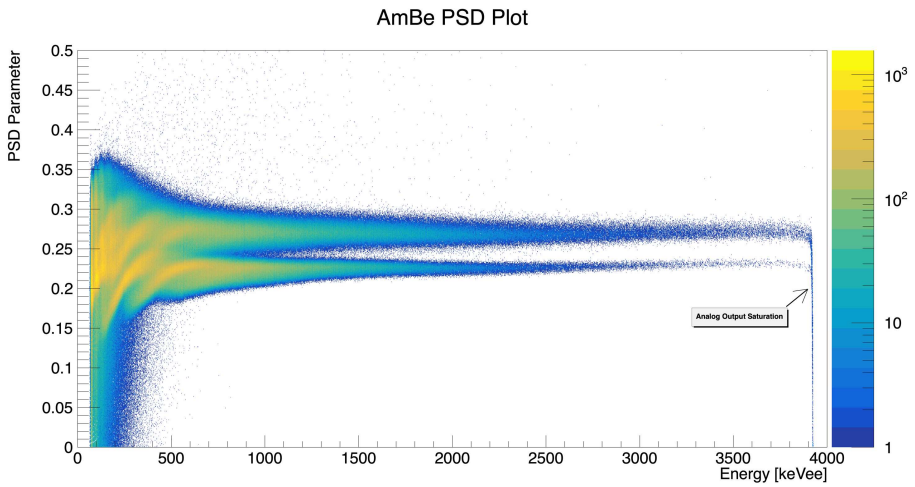

Figure 7. AmBe PSD Plot for ASIC Channel 40. The oscillatory behavior seen in the slow shaper spectra now appears as odd structures in the low energy region of the PSD spectrum. At the high energy part of the spectrum the effects of saturation (because the ASIC cannot produce values above 1.8 V) can be seen.

5.3 PSD PERFORMANCE

Finally, with all the parameters chosen and energy calibrations complete a long measurement of the 100 mCi AmBe source was taken to characterize the ASIC's PSD performance. The PSD parameter vs measured energy plot can be seen in Fig. 7. The oscillatory behavior at low energy seen earlier in Figures 5 and 6 can be seen in the low energy structure. Simultaneously, saturation of the analog outputs of the ASIC can be seen at high energy, as expected. When the slow shaper's peak height reaches 1.8 V it cannot rise

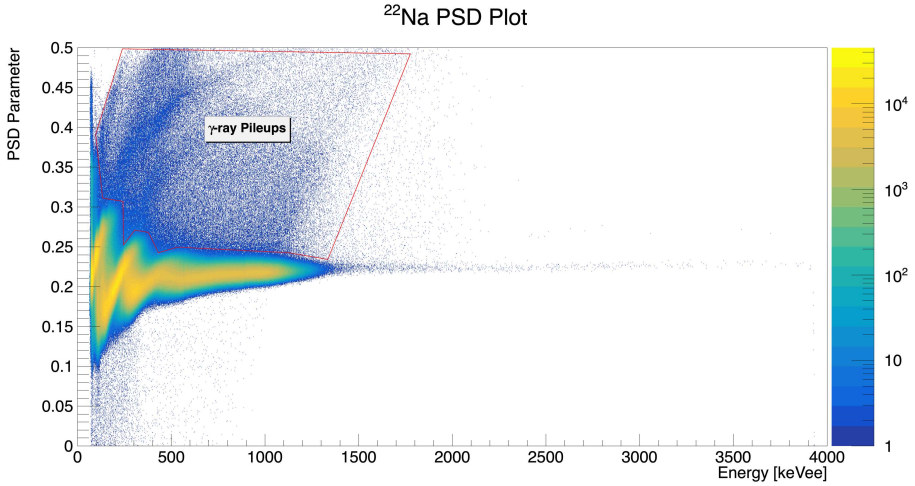

Figure 8. ^{22}Na PSD Plot for ASIC Channel 40. As can be seen the low energy structure essentially perfectly overlays a portion of the structure seen in Fig. 7. The region seemingly covering the neutron interaction region is due to the high intensity of the ^{22}Na source used which produced significant pileup counts. Using a longer measurement with a lower intensity source, an energy dependent PSD parameter cut can be developed and used to achieve better separation than FoM or eye might suggest.

Table 3. PSD FoM values for ASIC Channel 40 for various energy ranges.

Energy Region	FoM	γ ray Centroid γ ray FWHM	Neutron Centroid Neutron FWHM
250 to 400 keVee	0.672 ± 0.002	C: 0.22394 ± 0.00003 W: 0.01361 ± 0.00003	C: 0.26636 ± 0.00003 C: 0.13187 ± 0.00005
400 to 750 keVee	1.268 ± 0.003	C: 0.22602 ± 0.00002 W: 0.00624 ± 0.00002	C: 0.27084 ± 0.00002 C: 0.00877 ± 0.00003
750 to 1500 keVee	1.273 ± 0.001	C: 0.22585 ± 0.00001 W: 0.00545 ± 0.00001	C: 0.27056 ± 0.00001 C: 0.00947 ± 0.00001
1500 to 3500 keVee	1.619 ± 0.002	C: 0.22586 ± 0.00001 W: 0.00412 ± 0.00001	C: 0.26955 ± 0.00002 C: 0.00734 ± 0.00001
2500 to 3900 keVee	1.939 ± 0.009	C: 0.22825 ± 0.00004 W: 0.00331 ± 0.00003	C: 0.26950 ± 0.00004 C: 0.00573 ± 0.00003

any higher even if a larger pulse comes in. However, the fast shaper's peak height is generally smaller then the slow shaper's and can continue rising with increasing pulse size, resulting in the downward "streak".

Of particular interest, given the structure of the PSD plot at low energy, is that structural deviations. Therefore to obtain improved neutron and γ ray separation it is possible to measure a γ ray spectrum in the range that covers the region with structure. Then one would use it to map the PSD vs Energy shape at low energies and produce an energy dependent PSD parameter cut. This would accurately extend the neutron and γ ray separation far lower and with far greater accuracy than eye inspection of Fig. 7 or the FoM values given in Table 3 would suggest. This can be seen in Fig. 8 which mimics Fig. 7 in setup but was produced using a γ ray-only ^{22}Na source.

Afterwards projections onto the Y axis of Fig. 7 are taken for 5 energy ranges, the energy ranges and FoM

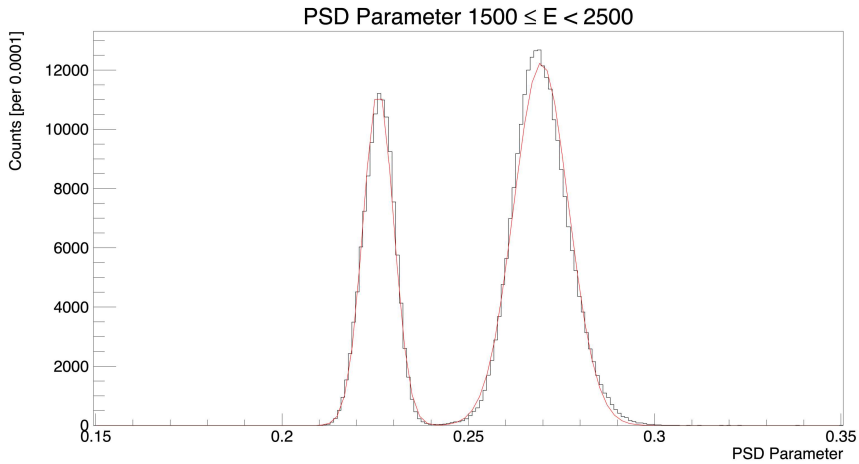

Figure 9. PSD FoM plot for 1500 to 3500 keVee of ASIC Channel 40. The high PSD parameter peak is due to neutron interactions in the scintillator while the lower peak is due to γ rays. The peaks are not quite Gaussian in shape, but Gaussian is a close enough approximation to measure FoM.

results are listed in Table 3. An example of one of these projections can be seen in Fig. 9.

6. LINEARITY

In the ASIC design brief, no statement was made as to linearity of the ASIC. In general this was considered unimportant as semi-Gaussian shapers with carefully selected poles produce little deviation from linearity unless near threshold or saturation. However, in discussions, a few voiced their concern was raised that the structures seen in the PSD spectra may be due to non-linearity effects. To address this a sequence measurements using an arbitrary waveform function generator to test the linearity of two of the ASICs channels (the same two used in subsections 5.1 and 7.1)

6.1 SETUP

To measure the ASIC's linearity current pulses resembling those of an SiPM needed to be injected into the ASIC's channels. As arbitrary waveform function generators output voltage waveforms they need to be capacitively coupled to the ASIC's inputs to produce the requisite current pulse from the voltage pulse. A small board (seen in Fig. 10) was developed to capacitively couple two outputs of an arbitrary waveform generator to two selectable ASIC input channels. The board connects directly to the high-density SiPM array connectors available on the ASIC support board. As the chosen arbitrary waveform pulser (a Keysight 81150A Pulse Function Arbitrary Noise Generator [7]) can output pulses ranging from −5 V to 5 V the ceramic coupling capacitor was chosen to be 680 pF so as to be able to inject enough charge to drive the ASIC through its whole dynamic range, even at the lowest gain.

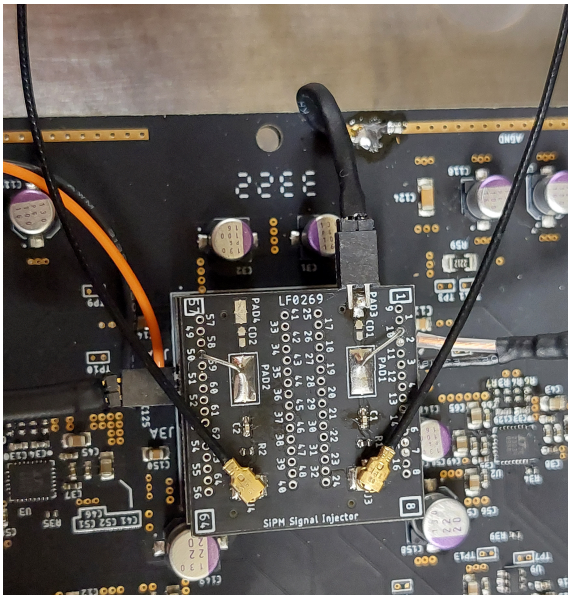

Figure 10. The ASIC pulse injection board. The top (visible) side of the injection board shows the two UFL RF connections with the fine coaxial cables, of UFL to BNC (not in the image) are attached to the metallic box. The capacitors coupling the pulse inputs to the ASIC input channels are visible immediately next to them. Finally, the wires connecting the "output" side of the capacitors to the chosen ASIC input channels are also visible. The numbers on the board next to each channel's connection hole correspond to the SiPM array connector channel. Table 1 shows how to map those to ASIC channels.

Capacitive coupling to convert a voltage pulse into a current pulse introduces a derivative. The current pulse injected into the ASIC is $I(t) = C \times \frac{dV(t)}{dt}$: where $I(t)$ is the current into the ASIC as a function of time, C is the capacitance of the injection capacitor, and $V(t)$ is the input voltage as a function of time. The J-Series SiPM datasheet [8] provides a plot of an idealized pulse from a J-Series SiPM discharged through a resistive load (so voltage is proportional to current). This pulse was digitized and the resultant points fit with:

$$I(t) = a \cdot e^{-\frac{c_1^2 \ln\left(\frac{\left(\left(\frac{d}{r}\right)^{c_2} - \left(\frac{d}{r}\right)^{c_3}\right)}{d} t\right)^2}{\ln\left(\frac{d}{r}\right)^2}} \tag{4}$$

Where $I(t)$ is the current pulse as a function of time, a is the amplitude, r is 10% to 90% peaking time of the pulse, d is the 90% to 10% decay time of the pulse, t is time, and the constants c_1, c_2, and c_3 are defined by Eq. 5.

$$\begin{aligned} c_1 &= \sqrt{\ln(10)} + \sqrt{\ln\left(\frac{10}{9}\right)} \\ c_2 &= \left(1 + \sqrt{\frac{\ln\left(\frac{10}{9}\right)}{\ln(10)}}\right)^{-1} \\ c_3 &= \left(1 + \sqrt{\frac{\ln(10)}{\ln\left(\frac{10}{9}\right)}}\right)^{-1} \end{aligned} \tag{5}$$

The the fitted rise and decay times were 2.13 ns and 176 ns respectively. The fitted amplitude was discarded as, for the pulser, the pulse shape's scale is renormalized to the desired fraction of the pulsers 10 V rail to rail range.

As stated earlier the capacitive coupling introduces a derivative, so Eq. 4 is integrated and the boundary condition that the amplitude at $t = 0$ must be zero applied to get:

$$V(t) = \frac{d e^{\frac{\ln\left(\frac{d}{r}\right)^2}{4c_1^2}} \ln\left(\frac{d}{r}\right) \sqrt{\pi}}{2c_1\left(\left(\frac{d}{r}\right)^{c_2} - \left(\frac{d}{r}\right)^{c_3}\right)^2} \cdot Erfc\left(\frac{\ln\left(\frac{d}{r}\right)}{2c_1} - \frac{\ln\left(\frac{\left(\frac{d}{r}\right)^{c_2} - \left(\frac{d}{r}\right)^{c_3}}{d} t\right)}{\ln\left(\frac{d}{r}\right)}\right) \tag{6}$$

Where $V(t)$ is the voltage as a function of time, $Erfc$ is the complementary error function, and r, d, c_1, c_2, and c_3 all retain their meanings from Eq. 4

With the preperatory work complete a script was written in the Python programming language to control the pulser, setting it to generate pulses of the shape defined in Eq. 6 at 5000 pulses per second on each of the two output channels. Additionally, another python script was written to orchestrate the ASIC DAQ as well as the pulser control script. The second script chooses a pulse amplitude and ASIC gain and shaping time parameters, executes the pulse control script to set the new pulse waveform, and transmits new configuration parameters to the ASIC DAQ. Then the orchestration script monitors the ASIC DAQ output file size, if it does not grow at the expected rate then it assumes the pulse amplitude is low enough the ASIC cannot trigger and ceases the run sequence to move on to the next. Otherwise the orchestration script ceases the run when the file size indicates enough triggers have been recorded and begins the next run.

The orchestration script tests 28 different ASIC configurations 7 different *total* gain values multiplied by 4 different shaping time configurations. The shaping time configurations set the slow shaping and fast shaping times to be one of the following pairs: 300 ns and 80 ns, 230 ns and 65 ns, 160 ns and 50 ns, or 90 ns and 35 ns. The total input gain values of ×1, ×2, ×4, ×8, ×16, ×32, ×64 are achieved by setting the input gain and both the fast and slow shaper gains to one of the following pairs: ×1 and ×1, ×2 and ×1, ×4 and ×1, ×8 and ×1, ×8 and ×2, ×8 and ×4, or ×8 and ×8.

For each ASIC configuration the orchestration script attempts to run 210 different pulse amplitudes from largest to smallest. The largest pulse amplitude of the sequence is 10 V when the total gain is ×1. For sequences with higher gains the largest pulse amplitude of the sequence is $10/(g - 0.1)$ V where g is the total input gain. All other pulse amplitudes in a sequence are scaled down from the maximum amplitude (A_{max}) in steps of $A_{max}/210$. This allows each configuration to have approximately 210 points evenly distributed between too small to trigger and saturation.

6.2 RESULTS

Once all the runs are complete the raw data for each run is converted to a ROOT TTree and for each tree the single peak for the slow and fast shapers from each run is fit and the fit parameters and other relevant information are stored in a separate TTree. To render the results comparable instead of working with just the injected pulse amplitude the scaled injected pulse amplitude (the product of the injected pulse amplitude and the total ASIC gain) is used. Plots of the fitted slow and fast shaper peak positions for channel 8 vs the scaled pulse amplitude can be seen in Fig. 11

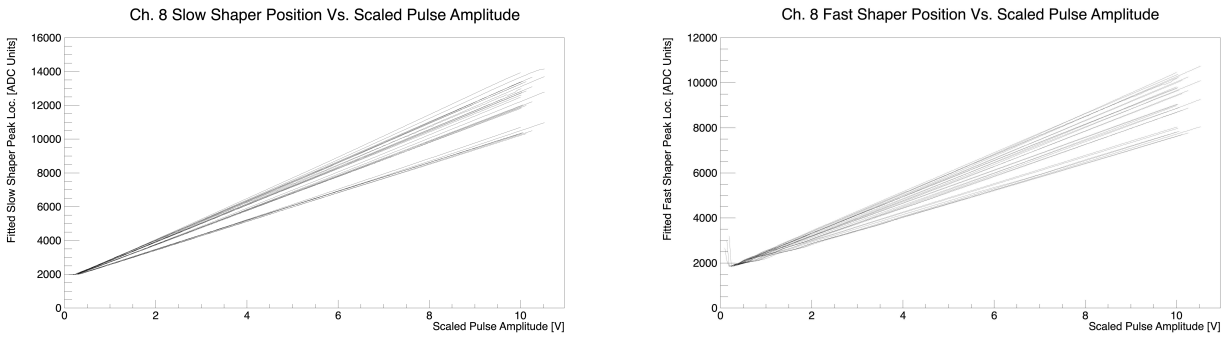

Figure 11. ***Left*** **A plot of fitted peak position vs scaled injected pulse amplitude for the ASIC Channel** 8 **slow shaper for all** 28 **configuration combinations.** ***Right*****A plot of fitted peak position vs scaled injected pulse amplitude for the ASIC Channel** 8 **fast shaper for all** 28 **configuration combinations.** For both plots some "jitter" in the vertical position of the lines can be seen especially at high pulse position. This is can be attributed to the automated nature of the peak fitting process and thus some fits may be imperfect. For the fast shaper, the lowest scaled pulse amplitude points for three of the four ×64 gain runs show an abrupt shift upwards. This is due to conversion noise from the pulser causing a few counts to be broadly scattered, which in turn badly disturbed the automated fit range, preventing the fit from finding the peak. Finally, the fast shaper appears at low scaled pulse amplitude to have some small regions where the slope of the line increases before returning to its prior value. While this may contribute to the structure seen in PSD plots like Fig. 7 it seems unlikely to be the the the full cause, and perhaps not even a significant contributor as this pattern is not seen for the slow shaper; despite the structure being clearly visible in slow shaper histograms like Fig. 6

Table 4. Nonlinearity ratio percentages for ASIC Channel 8.

Peak Time	Slow Shaper				Fast Shaper			
[ns] / Gain	90	160	230	300	35	50	65	80
1	1.39	2.34	2.62	2.71	1.21	2.11	2.48	2.70
2	1.33	2.07	2.26	2.28	1.42	2.12	2.35	2.43
4	0.45	0.97	1.25	1.36	0.87	1.63	1.86	1.92
8	0.22	0.44	0.68	0.85	0.55	1.04	0.58	0.48
16	0.38	0.092	0.29	0.42	0.28	0.063	0.17	0.26
32	0.86	0.23	0.026	0.17	0.63	0.46	0.44	0.40
64	1.16	0.40	0.23	0.22	0.84	0.66	0.62	0.68

Further analysis of the data was carried out in bulk using mathematica, importing CSV formatted X-Y pairs that had been generated in the analysis process. For each configuration each of the sets of peak position vs scaled pulse height data for ASIC Channel 8 was fit with a linear and a quadratic polynomial of the forms:

$$LinPos(x) = A + Bx \tag{7}$$

$$NonLinPos(x) = A + Bx + Cx^2 \tag{8}$$

Here $LinPos$ is the linear fit to peak position, $NonLinPos$ is the quadratic fit to peak position, x is the scaled pulse amplitude and A, B, and C are constants of the fit. For each fit the non-linear contribution, determined by ratio $|\frac{NonLinPos(x)-LinPos(x)}{LinPos(x)}|$ expressed as a percentage, with x chosen so as to maximize the ratio, is stored. This method gives a worst case measure of the integral non-linearity by comparing how different the, somewhat better, non-linear fit is from the linear fit at point of maximum difference. The results can be found in Table 4.

As can be seen in Table 4, even in the worst case, the nonlinearity ratio never exceeds 2.71 %. This data demonstrates that the ASIC's response is highly linear, as expected of semi-gaussian shapers.

7. TIMING RESOLUTION

Another goal of the ASIC is to achieve a timing resolution less than 1 ns Root Mean Square (RMS, i.e. σ, the standard deviation of a gaussian distribution). It would also be preferable for the ASIC's timing resolution to be comparable to the timing resolution of a recent state of the art Associated-Particle Neutron Radiography measurement which was found to have an FWHM ($2 \cdot \sqrt{2 \cdot ln(2)} \times \sigma$) of 1.7 ns[9]. The method to measure timing resolution is somewhat convoluted and so has been divided into two components: Relative Timing Calibration and Resolution Measurement.

7.1 RELATIVE TIMING CALIBRATION

The first step in measuring the time resolution is to produce a relative timing calibration of the two ASIC channels that the measurement will be performed with. As, in principle, any pair of channels should exhibit similar results, ASIC channels 8 and 40 were chosen. For the relative timing calibration, the ASIC support board was coupled to an 8×8 array of OnSemi J-Series 6 mm SiPMs via the 3 mm array to 6 mm array interposer board as in the left image of Fig. 4. Aftwards, bias was applied to the SiPM array (the capacitance of an unbiased SiPM is so high that the ASICs internal pulser cannot trigger a channel).

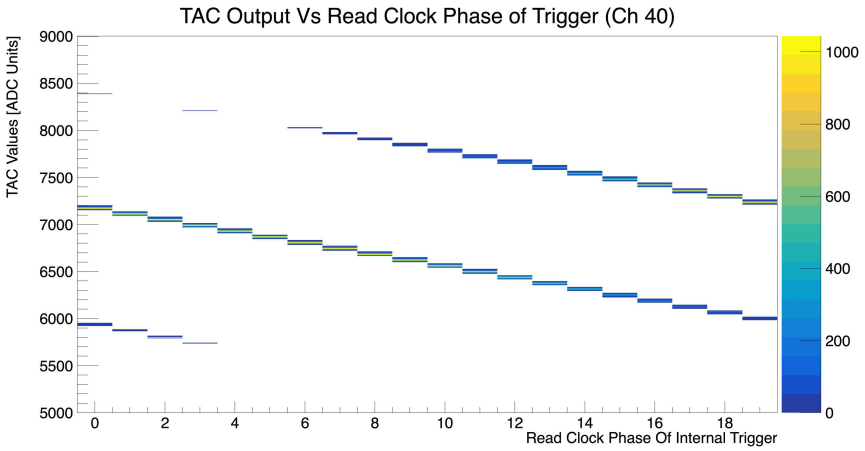

Figure 12. Uncalibrated TAC output of ASIC channel 40 **vs phase of the FPGA's ASIC readout clock.** As can be seen, each phase has two or more TAC peaks. This is due to some non-determinism in when exactly a channel's peak is marked found and the flag sent to the FPGA. This in turn causes the same phase to give multiple peaks, spaced 100 ns apart.

When a TAC is triggered by a leading edge discriminator there can be a significant difference in the timing of the pulse depending on the amplitude relative to the trigger threshold of the ASIC. Pulses whose amplitude is closer to trigger level start the TAC later than those that have a higher amplitude. To get different pulse heights from the ASIC's internal pulser which always generates pulses of approximately identical heights, all the configuration parameters are held constant from those chosen in the 5.1 with the exception of the timing shaper's current mirror gain. The nominal timing shaper gain, irrelevant prior to this measurement, was chosen to be ×4 henceforth. This, in turn, allowed multiple relative time calibration runs with the internal pulser to give different pulse amplitudes in the timing shaper by running the

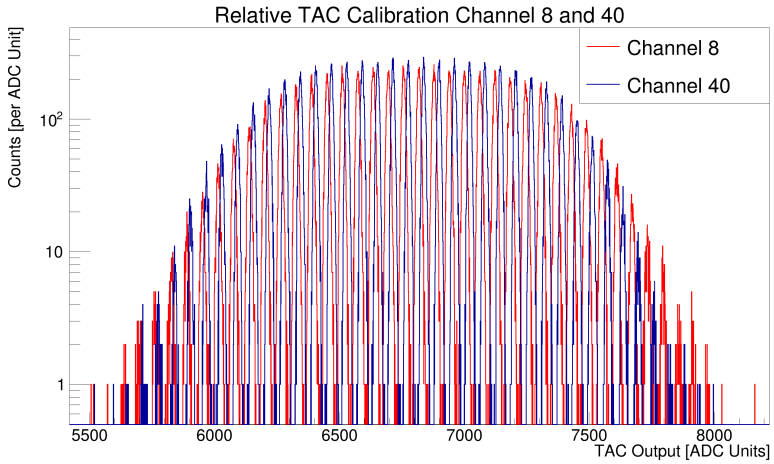

Figure 13. Uncalibrated TAC output of ASIC channels 8 **and** 40**.** The observed, shifting, offset between the two channels is caused by manufacturing variations in the ASIC which slightly vary the actual current output of the TAC ramping mechanism.

calibration runs with the timing shaper gain at ×2 and ×8 in addition to the nominal ×4 to be used for the actual resolution measurement.

For each timing shaper gain, the relative time calibration function on the FPGA was run 2048 times, yielding 4096 pulses per channel per 5 ns phase of the the 100 ns period ASIC readout clock. This in turn gives a sequence of gaussians in the TAC output value whose centroids are separated by 5 ns. A plot of the TAC out values vs readout clock phase for a single channel of the ×4 timing calibration run can be seen in Fig. 12. The overlay of the 1D histograms of uncalibrated TAC values for ASIC channels 8 and 40 can be seen in Fig. 13 and exemplifies the necessity of the relative timing calibration step to correct for manufacturing variations between each channel's TAC.

For the relative timing calibration an arbitrary (but consistent between channels and runs) TAC peak must be chosen to be $t = 0$ ns. For this work, the lower TAC peak associated with read clock phase 13 was chosen because it was generally the lowest TAC peak that contained more counts than it's upper TAC peak equivalent.

From this point every peak more than 20 counts in height was fit and its centroid TAC value was associated with a relative time set by the number of 5 ns intervals it was separated from the zero-point by. The TAC centroids and relative times used in this work can be seen in Appendix A. in Tables 8 and 9 for ASIC channel 8 and 40 respectively. Additionally, for each of the runs the slow shaper peak was fit for ASIC channels 8 and 40, yielding the data in Table 5.

As the slow shaper gain did not change, the centroids in Table 5 are relatively unvarying. Since the purpose of performing these calibrations at multiple different gains is to get some handle on the timing walk due to varying pulse amplitudes. Under normal circumstances the slow shaper pulse amplitude should be proportional to the timing shaper amplitude which is in turn the independent to the function that determines the amount of shift due to amplitude walk. Unfortunately, the slow shaper parameters cannot be varied for these runs because the selected gain for energy calibration was ×1 so while it can go to double its normal value, it cannot go to half its value. Therefore, in order to calculate the "correct" slow shaper values for each of the runs the following is done. The analog baseline of the ASIC is measured using the monitor output of the ASIC and the error weighted average of the centroid was found to be 1767.0 ± 0.02 across all the runs (the variance between runs was extremely small so a single value was used instead of a value per run). Then for each run the analog baseline was subtracted from the slow shaper centroid, the result

Table 5. Fitted Slow Shaper Peaks For Timing Calibration

Gain Setting	Ch 8 Cent.	Ch 8 Cent. Error	Ch 40 Cent.	Ch 40 Cent. Error
×2	2398.75	0.050902	2387.64	0.038932
×4	2398.22	0.051698	2387.02	0.039092
×8	2399.25	0.050446	2388.50	0.038778

Table 6. Modified Slow Shaper Peaks For Timing Calibration

Gain Setting	Ch 8 Cent.	Ch 8 Cent. Error	Ch 40 Cent.	Ch 40 Cent. Error
×2	2082.88	0.027345	2077.32	0.021884
×4	2398.22	0.051698	2387.02	0.039092
×8	3031.50	0.102855	3010.00	0.080093

Table 7. Timing Calibration Fit Parameters

	Ch. 8 Parameters		Ch. 40 Parameters	
Parameter	Value	Error	Value	Error
a	-149.097	0.000140838	-149.044	0.000151543
b	0.076173	0.000498192	0.0748547	0.000553594
c	4.06056×10^{-7}	3.57107×10^{-8}	4.66437×10^{-7}	4.00571×10^{-8}
d	1.74101	0.0120611	1.72286	0.0131099
e	0.00113912	2.61093×10^{-6}	0.00113554	2.73176×10^{-6}
f	-1.93346×10^{-7}	5.06064×10^{-10}	-1.93874×10^{-7}	5.32192×10^{-10}

multiplied by the gain of the run relative to the gain of the ×4 run (0.5, 1.0, and 2.0 for ×2, ×4, and ×8 respectively). Then the result of that multiplication had the analog baseline added back in. This yielded the data in Table 6.

With that complete the slow shaper peak heights are mixed with the relative timing calibration information and fitted to the function seen in Eq. 9.

$$RelativeTime(S,T) = a\left(d + eS + fS^2\right) + bT + cT^2 \tag{9}$$

Where S is the slow shaper amplitude for the event, T is the TAC output for the event, and a, b, c, d, e, and f are fit parameters. In essence, the slow shaper height, gently, shifts the offset of the timing calibration. This matches how amplitude causes timing to walk, though the polynomial fit to the slow shaper is not necessarily ideal; but, it is simple enough that the limited height data we have can be used. The resulting fits yielded the parameters and errors in Table 7.

Due to the relatively small variations in pulse height, the walk correction this yields is only valid for pulses

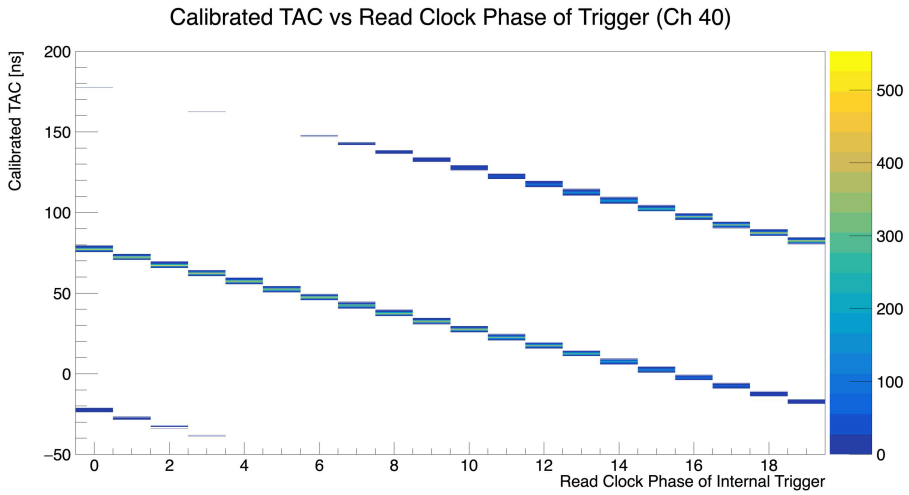

Figure 14. Calibrated relative time of ASIC channel 40 **vs phase of the FPGA's ASIC readout clock.** As can be seen, each phase has two or more TAC peaks. This is due to some non-determinism in when exactly a channel's peak is marked found and the flag sent to the FPGA. This in turn causes the same phase to give multiple peaks, spaced 100 ns apart.

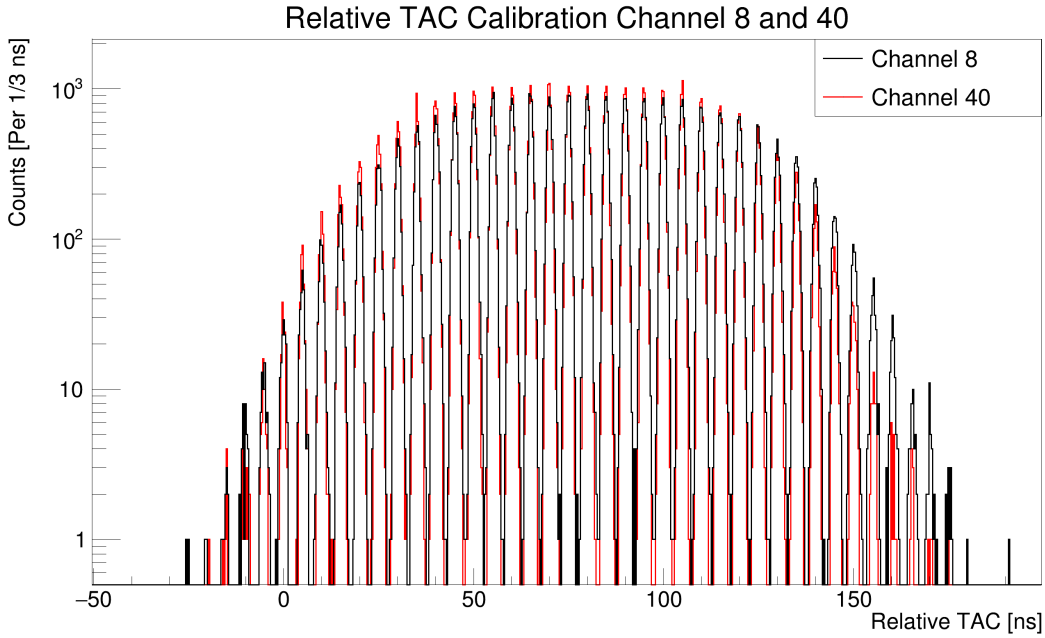

Figure 15. Calibrated relative time of ASIC channels 8 **and** 40 **from a timing calibration run.** As can be seen, post relative timing calibration, the peak centroids overlay perfectly. The difference in peak heights at the sides can be attributed to differences in flag release times and flag signal propagation delays within the ASIC on a channel by channel basis.

with relatively small amplitudes. Fortunately, as will be described in the upcoming timing resolution measurement section, the signals used to measure timing resolution will yield small pulses. A calibration technique utilizing an external programmable pulser will be necessary to produce a timing calibration that is valid for the entire gamut of pulse heights. Happily, such a method would allow the timing shaper gain to be set to ×8 at all times which will yield the least timing walk at low signal amplitudes.

After applying the timing calibrations to the timing calibration data we find that Fig. 12 now yields Fig. 14. As a demonstration of the calibrations success Fig. 13 was regenerated with the relative time instead of TAC and that results in Fig. 15.

7.2 TIMING RESOLUTION MEASUREMENT

To measure the timing resolution of the ASIC a source that can emit two photons *simultaneously* is necessary. Fortunately, the dual 511 keV γ rays emitted by positron annihilation fulfull this requirement perfectly. The two single stilbene crystals were placed on ASIC Channels 8 and 40 on the 6 mm SiPM array. Then the 3 mm SiPM array was placed SiPM side down on the 6 mm array so that the two high density connectors could serve as a holder for the cylindrical 50 μCi ^{22}Na source. This can be seen in Fig. 16

With the source in place a long measurement (approximately 24 hours) was taken resulting in 635773940 single trigger events. Another script was then written to filter these events for events that occurred in coincidence between ASIC channels 8 and 40. This script also applied the relative timing calibrations to

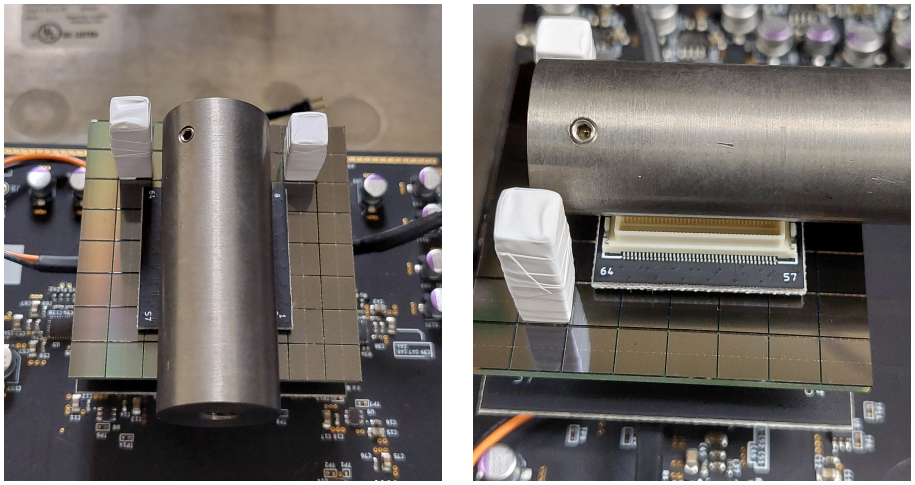

Figure 16. ***Left*** **the cyldrical ^{22}Na source resting on the upside down 3 mm SiPM array seen from above.** ***Right*** **a side view of the same setup.** The source material is located close to one end of the cylinder, primarily in the portion between the two scintillator crystals.

the TAC outputs and put the coincident pairs of trigger data into a new tree which contained 986979 coincidence events. Taking this filtered coincidence tree the relative time values of the two coincident events were plotted against eachother to yield Fig. 17.

Many of the counts in the diffuse region outside the "coincident streak" of Fig. 17 are random coincidences, γ rays from two unrelated nuclear decays that happened to occur close enough together. Due to the nature of a ^{22}Na source, within the streak there are two categories of coincidences, synchronized and unsynchronized. Synchronized coincidences occur when both 511 keV photons from from the positron

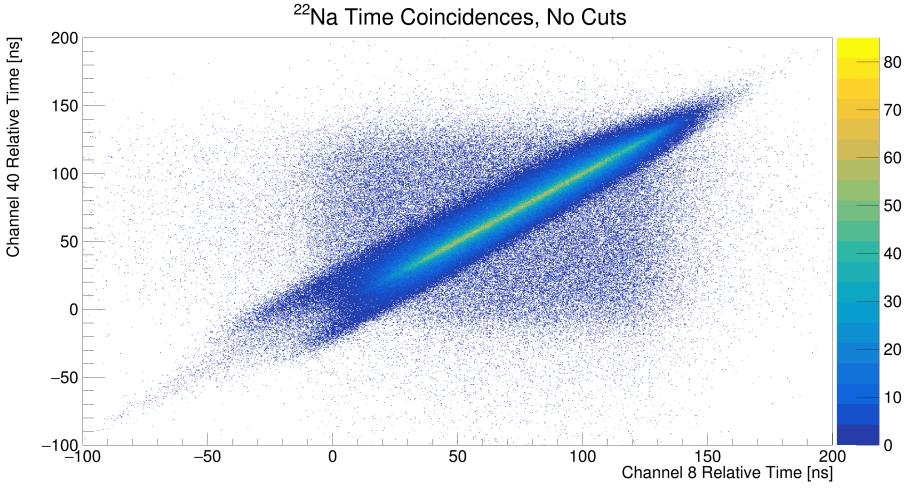

Figure 17. Relative times of coincident triggers ASIC channel 40 **vs** 8**.** The low density triangular regions above and below the "coincident streak" are random coincidences. Much of the breadth of the streak is due to true, but unsychronized coincidences between one of the two 511 keV photons from positron annihilation and the 1274.5 keV photon emitted by the ^{22}Ne daughter nucleus.

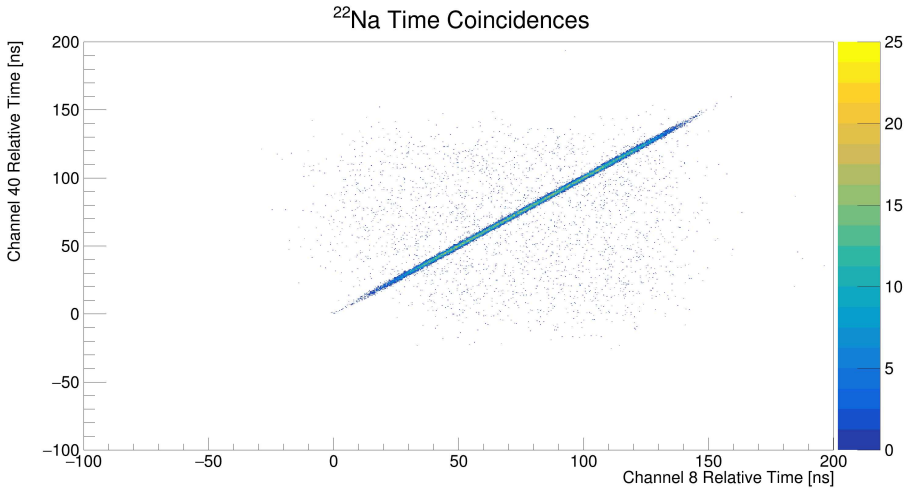

Figure 18. Relative times of coincident triggers ASIC channel 40 **vs** 8 **with energy cuts.** Here are plotted only the timing coincidences whose slow shaper values lie within a narrow region of the 511 photon's Compton Edge. As can be seen, relative to Fig. 17 the overwhelming majority of the random coindinces and most of the unsynchronized coincidences have been removed.

annihilation (generated at exactly the same instant by the annihilation process) interact, one in each crystal. Unsynchronized coincidences occur when one 511 keV photon from the annihilation interacts in one crystal, and the 1274.5 keV photon emitted by the ^{22}Ne daughter nucleus interacts in the other crystal. Both the amount of time for the positron to slow to rest and annihilate, as well the amount of time for the the excited daughter nucleus to emit it's γ ray are distributions (though the latter is quite narrow). Because of this, these unsynchronized coincidences form a timing peak significantly wider than the ASIC's intrinsic resolution.

To remove the overwhelming majority of both the random coincidences as well as the unsynchronized coincidences energy cuts are placed on the slow shaper peak values. A narrow range around the Compton Edges of the two channels is selected, approximately 273 to 366 keV for both channels. This has several desirable effects, it significantly attenuates random and unsynchronized coincidences. Additionally, it selects for the brightest of the synchronized coincidences and restricts their amplitude range which in turn decreases the importance of the walk correction. A verion of Fig. 17 with the cuts applied can be seen in Fig. 18

At this point, to measure the timing resolution a histogram of the difference between the two relative time values is generated. The histogram from this effort can be seen in Fig. 19. Fitting the coincidence peak with a Gaussian yields a centroid of -37 ± 6 ps and an RMS width of 980 ± 4 ps. However, this timing peak is not merely the ASIC's intrinsic timing resolution T_{ASIC}. It is broadened by the jitter in the FPGA's clock tree J_{FPGA} and the jitter added to signal by the logic translators and buffers J_{Tr}. As the peak's width is simply the sum in quadrature of its three contributors can write out that relationship and solve it for T_{ASIC} as seen in Eq 10:

$$\begin{aligned} (980 \pm 4)^2 &= T_{ASIC}^2 + J_{FPGA}^2 + J_{Tr}^2 \\ T_{ASIC} &= \sqrt{(980 \pm 4)^2 - \left(J_{FPGA}^2 + J_{Tr}^2\right)} \end{aligned} \tag{10}$$

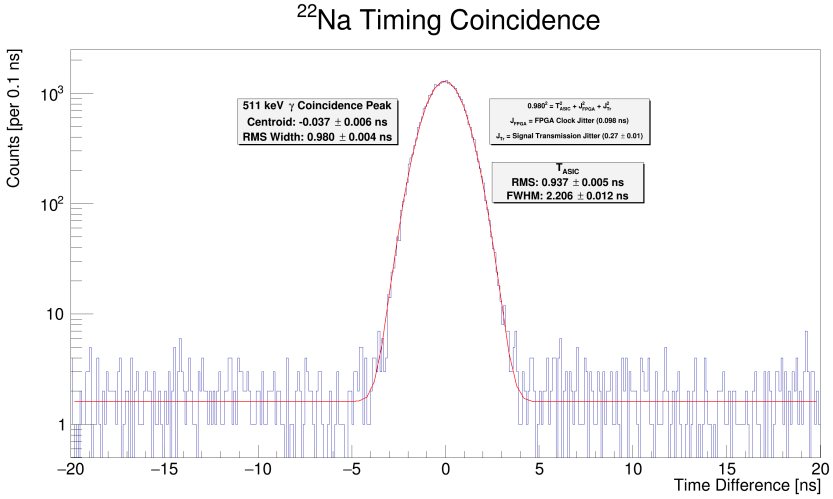

Figure 19. Difference spectrum of the ASIC TAC outputs which gives timing resolution. As the insets in the figure explain, while the fitted width is proportional to the ASIC's intrinsic timing resolution there are some corrections to remove the effects of jitter from various FPGA effects. Such effects can be removed with a more careful design involving externel retiming flip-flops, but introducing that in the test board would have complicated the design of the test board and possibly introduced additional errors to debug.

Ideally $\sqrt{J_{FPGA}^2 + J_{Tr}^2}$ could be measured directly be triggering an oscilliscope on the signal immediately prior to entering the ASIC while another channel reads the 200 MHz clock source (itself having approximately 1.4 ps of timing jitter, small enough to neglect). However, the contacts for the FPGA development board's clock source are impossible to reach. Happily, the FPGA's clocking resource tools estimate the clock jitter introduced by having the clock signal traverse the FPGA clocking fabric. The estimate for the input clock's jitter is 98 ps, no error bar provided. Additionally, J_{Tr} can be measured by triggering an oscilliscope on the signal as it enters the ASIC and reading the signal that is leaving the FPGA. This measurement was carried out with a Rhode & Schwarz RTO6 oscilliscope and yielded $J_{Tr} = 270 \pm 10$ ps. Combining these numbers with Eq. 10 yields an intrinsic ASIC timing resolution of 937 ± 5 ps RMS (or 2.206 ± 0.012 ns FWHM). This value is well within the original design goal of < 1 ns *RMS*

It should be noted that the signal sizes for the measurement were quite small. Performing this same measurement with no amplitude walk correction in the relative timing gives an RMS timing resolution of approximately 1.47 ns. This means that the amplitude walk correction is playing a signifcant role. Improved methods for determining the walk correction could reduce the timing resolution further. Simultaneously larger pulses would also improve the timing resolution, giving two paths to improving the timing resolution without a redesign of the ASIC.

8. CONCLUSION

In short, the ASIC meets or exceeds its stated design goals and performs well, even in areas not stringently specified.

The ASIC's power consumption, with necessary support components added in, is 10.27 mW per channel in

the worst case. The true value is likely less if a more fine-grained measurement of current consumption could be performed. Regardless, the ASIC's power consumption per channel is more than a factor of 65 smaller than current discrete designs; significantly exceeding the design goal of one tenth.

At the gain settings chosen, above 400 keVee the ASIC's PSD performance is excellent. Below 400 keVee the situation is more nuanced due to low energy structure induced in the PSD spectrum, the source of which is currently unknown. The properties of the structure between channels suggests that it may be some common mode effect, being injected into all the channels from the SiPM array, or injected into the channel inputs through the ASIC substrate. Further efforts are underway to determine the source of this structure to attempt to mitigate it. Regardless of the structure's source, there is evidence to suggest that the structure is largely constant in shape. Due to this, a careful measurement of a γ-ray source whose spectrum extends past the structured region could be used to develop an energy dependent PSD cut. This energy dependent PSD cut could in turn do a significantly better job of separating neutron and γ-ray interactions than the FoM figures and PSD plot may initially indicate.

The ASIC's linearity has been tested and it was found to be extremely linear in its response, never exceeding 2.7 % integral nonlinearity even at worst, and frequently less. If points close to the trigger threshold and close to saturation, problematic regions for any system, are not included, the linearity improves further.

The ASIC's intrinsic timing resolution was found to be 937 ± 5 ps RMS even for pulses that are comparatively quite close to the timing shaper TAC trigger threshold. This meets the stated timing resolution design goal of < 1 ns RMS. Further, even for pulses close to the threshold, the FWHM timing resolution of 2.206 ± 0.012 ns approaches the 1.7 ns FWHM timing resolution achieved by a 250 MSps waveform digitizers in state of the art neutron imaging measurements with extremely bright (high above threshold) proton recoils from 14.1 MeV neutron interactions[9]. If timing resolution were measured with significantly larger pulses (as would be the case with proton recoils from 14.1 MeV neutron interactions) then it is expected that the timing resolution would get significantly better, likely equaling or exceeding the timing resolution found in Ref. [9].

In conclusion:

- The ASIC's power consumption is outstanding, beating similar discrete solutions by a factor of 65.
- The ASIC's PSD performance is good, the figures of merit exceed 1.2 in all but the lowest energy region.
- The ASIC's linearity is excellent, even at its worst it never exceedes 2.7 % integral non-linearity.
- The ASIC's timing resolution exceeds the design goals and further measurements with the right setup and improved calibration methods can shrink that further still.

While the ORNL SiPM PSD ASIC has some points that further iteration could improve upon, it has achieved its stated design goals in its first design and fabrication cycle. As only 26% of ASIC development projects achieve success in first silicon[1], this is a remarkable achievement.

9. REFERENCES

References

[1] H. D. Foster, *2022 Wilson Research Group IC/ASIC Functional Verification Trends*, Tech. Rep. (Siemens Digital Industries Software, 2022).

[2] J. T. Matta, L. Fabris, and P. B. Rose, "Development of a Testing Setup and FPGA Firmware for the ORNL SiPM PSD ASIC," (2023), *Unpublished*.

[3] J. T. Matta, L. Fabris, and P. B. Rose, "Analog Shaper Characterization of the ORNL SiPM PSD ASIC," (2023), *Unpublished*.

[4] *2740/2745 Digitizers User Manual*, CAEN S.p.A. (2023), rev. 8.

[5] *Dual, 14-Bit + Sign, 5Msps Differential Input ADC with Wide Input Common Mode Range*, Analog Devices (2014), lT 1017 REV B.

[6] R. Brun and F. Rademakers, Nucl. Inst. & Meth. in Phys. Res. A **389**, 81 (1997).

[7] *Keysight 81150A and 81160A Pulse Function Arbitrary Noise Generator*, Keysight Technologies (2014), edition 2.0.

[8] *Silicon Photomultipliers (SiPM), High PDE and Timing Resolution Sensors in a TSV Package, J-Series SiPM Sensors*, OnSemi (2021), mICROJ-SERIES/D, REV 7.

[9] M. R. Heath, J. Daughhetee, P. Hausladen, J. Newby, and J. Matta, *Gantryless Associated-Particle Neutron Radiography*, Tech. Rep. (Oak Ridge National Laboratory, 2022) oRNL/TM-2022/2710.

APPENDIX A. Relative Timing Calibration Data

This appendix has the tables of raw relative timing calibration data for ASIC channels 8 and 40 extracted from the three timing calibration runs performed with the timing shaper gain set to ×2, ×4, and ×8.

Table 8. Relative time calibration data for ASIC channel 8. For each shaper gain setting the "Time" column shows the assigned relative time, "TAC Cent." shows the fitted centroid of the TAC, and "Cent. Error" shows the error in the fitted centroid of the TAC.

Timing Shaper Gain = ×2			Timing Shaper Gain = ×4			Timing Shaper Gain = ×8		
Time	TAC Cent.	Cent. Error	Time	TAC Cent.	Cent. Error	Time	TAC Cent.	Cent. Error
-30.	5828.8	0.920602	-30.	5989.1	1.06579			
-25.	5889.9	0.481938	-25.	6051.08	0.528547			
-20.	5953.02	0.393427	-20.	6114.48	0.380589	-20.	6218.42	0.444782
-15.	6016.17	0.278059	-15.	6176.94	0.304324	-15.	6281.65	0.347376
-10.	6078.1	0.203766	-10.	6239.41	0.230742	-10.	6343.32	0.313843
-5.	6140.88	0.190526	-5.	6301.36	0.180507	-5.	6404.48	0.300217
0.	6203.7	0.163576	0.	6363.08	0.151656	0.	6466.71	0.190266
5.	6266.41	0.150422	5.	6425.23	0.143781	5.	6528.04	0.1731
10.	6328.34	0.135617	10.	6486.64	0.136731	10.	6589.79	0.188699
15.	6389.86	0.127753	15.	6548.16	0.127345	15.	6650.92	0.144711
20.	6452.02	0.126347	20.	6609.09	0.123383	20.	6711.47	0.140873
25.	6513.24	0.115533	25.	6670.15	0.117488	25.	6772.78	0.139537
30.	6574.79	0.111054	30.	6730.86	0.121904	30.	6833.99	0.141898
35.	6635.98	0.10964	35.	6791.84	0.113238	35.	6894.56	0.135401
40.	6696.59	0.114732	40.	6852.99	0.113875	40.	6955.68	0.136254
45.	6757.42	0.110621	45.	6913.96	0.114246	45.	7016.67	0.122232
50.	6818.57	0.089511	50.	6974.92	0.129966	50.	7077.69	0.133332
55.	6879.56	0.106912	55.	7035.53	0.112519	55.	7138.26	0.13889
60.	6940.6	0.111152	60.	7096.52	0.111257	60.	7200.05	0.136009
65.	7001.24	0.113774	65.	7157.14	0.112027	65.	7261.12	0.127872
70.	7062.6	0.115663	70.	7219.05	0.118882	70.	7322.28	0.124483
75.	7123.29	0.114878	75.	7280.32	0.122395	75.	7383.04	0.130548
80.	7184.5	0.122052	80.	7341.13	0.123194	80.	7443.94	0.120519
85.	7246.15	0.13397	85.	7401.94	0.126935	85.	7504.89	0.129812
90.	7307.1	0.131621	90.	7462.8	0.135492	90.	7565.71	0.137737
95.	7368.11	0.13972	95.	7523.43	0.144547	95.	7626.16	0.120698
100.	7428.66	0.164771	100.	7584.76	0.120601	100.	7687.41	0.169756
105.	7489.78	0.189477	105.	7645.49	0.157468	105.	7748.75	0.210966
110.	7550.18	0.218855	110.	7705.96	0.119496	110.	7810.38	0.248717
115.	7611.17	0.277514	115.	7767.6	0.264438	115.	7870.26	0.260889
120.	7672.42	0.377981	120.	7829.22	0.406425	120.	7929.89	0.420909
			125.	7889.95	0.541587	125.	7989.88	0.411062
			130.	7949.93	0.863787	130.	8048.7	0.602766

Table 9. Relative time calibration data for ASIC channel 40**.** For each shaper gain setting the "Time" column shows the assigned relative time, "TAC Cent." shows the fitted centroid of the TAC, and "Cent. Error" shows the error in the fitted centroid of the TAC.

Timing Shaper Gain = ×2			Timing Shaper Gain = ×4			Timing Shaper Gain = ×8		
Time	TAC Cent.	Cent. Error	Time	TAC Cent.	Cent. Error	Time	TAC Cent.	Cent. Error
-30.	5841.91	0.595198	-30.	5999.41	0.659196	-30.	6104.08	0.676904
-25.	5904.86	0.379606	-25.	6062.35	0.382469	-25.	6166.34	0.341355
-20.	5967.96	0.278437	-20.	6124.31	0.25604	-20.	6228.75	0.338945
-15.	6030.83	0.196533	-15.	6187.9	0.200084	-15.	6291.43	0.288268
-10.	6093.43	0.166019	-10.	6251.06	0.170018	-10.	6353.42	0.189925
-5.	6156.45	0.144174	-5.	6313.2	0.140391	-5.	6415.87	0.163192
0.	6219.93	0.134066	0.	6375.5	0.122419	0.	6477.82	0.143892
5.	6282.49	0.114297	5.	6437.51	0.112923	5.	6539.49	0.132454
10.	6344.9	0.105698	10.	6499.42	0.107559	10.	6601.43	0.121974
15.	6406.8	0.106183	15.	6561.37	0.105584	15.	6662.83	0.111618
20.	6469.06	0.101983	20.	6623.13	0.100657	20.	6724.14	0.113806
25.	6530.72	0.0987	25.	6683.81	0.100799	25.	6785.64	0.113014
30.	6592.59	0.097457	30.	6745.2	0.100239	30.	6846.85	0.110517
35.	6653.96	0.094331	35.	6807.07	0.099034	35.	6908.33	0.112554
40.	6715.14	0.109511	40.	6868.17	0.097298	40.	6969.54	0.109108
45.	6776.49	0.093871	45.	6929.87	0.102142	45.	7030.94	0.110656
50.	6837.97	0.08445	50.	6990.91	0.100536	50.	7092.01	0.106182
55.	6899.41	0.097676	55.	7051.9	0.098182	55.	7152.97	0.114088
60.	6961.04	0.099765	60.	7113.26	0.090115	60.	7215.19	0.110807
65.	7021.57	0.093192	65.	7174.02	0.082311	65.	7276.8	0.105463
70.	7083.16	0.100631	70.	7236.35	0.099568	70.	7338.54	0.109071
75.	7144.22	0.101529	75.	7297.98	0.099339	75.	7399.54	0.107984
80.	7206.28	0.106337	80.	7359.24	0.105315	80.	7460.83	0.104338
85.	7268.05	0.113493	85.	7420.26	0.093962	85.	7521.81	0.111466
90.	7329.5	0.121733	90.	7481.7	0.126906	90.	7583.31	0.135808
95.	7390.65	0.141648	95.	7542.7	0.135359	95.	7644.11	0.153415
100.	7451.75	0.161017	100.	7603.74	0.160962	100.	7705.19	0.17581
105.	7513.23	0.191951	105.	7665.04	0.215275	105.	7767.52	0.227988
110.	7573.91	0.265547	110.	7726.39	0.299747	110.	7829.54	0.289435
115.	7635.27	0.369964	115.	7788.21	0.354824	115.	7889.72	0.408587
120.	7696.64	0.685561	120.	7847.81	0.607573	120.	7949.45	0.718562