



SiPM Characterization for use in HAICU and Assorted Projects for the ALPHA Experiment

Zach Charlesworth

CERN Summer Student Programme

TRIUMF Summer Student Programme

IPP Summer Student Fellowship

McGill University Department of Physics

May 2024 - August 2024

Abstract

This report outlines the work completed during the summer of 2024 at both TRIUMF and CERN by Zach Charlesworth, as part of the IPP Student Fellowship Programme, the TRIUMF Summer Student Programme, and the CERN Summer Student Programme. The first half of the summer was spent working with the HAICU collaboration at TRIUMF in Vancouver, Canada, where room temperature characterization of the Hamamatsu VUV4 SiPM was conducted. This characterization included, but was not limited to, determining the signal-to-noise ratio, system dead time, and dark noise rate. The exact procedures for data acquisition, analysis, and the final results are outlined in this report. The second half of the summer was spent working with the ALPHA collaboration at CERN in Geneva, Switzerland. During this time, the ALPHA2 experiment was in operation, so time was divided between shift work for the ALPHA2 experiment, attending the student lecture programme, and working on projects for the ALPHA-g experiment, which was set to start taking data at the end of the summer. The projects completed in preparation for the ALPHA-g experiment are outlined in this report.

Contents

1	Motivation for SiPM Characterization	2
2	Background Information	3
2.1	Overview of SiPMs	3
2.2	Breakdown Voltage and IV Curves	5
2.3	Afterpulsing	6
3	Setup and Apparatus	7
3.1	General Setup	7
3.2	Triggers	8
3.2.1	Threshold/Self Triggering	8
3.2.2	External Triggering	9
3.3	Data Acquisition and Processing	9
3.3.1	Dead Time	10
4	Characterization	10
4.1	Dark Data	10
4.1.1	Dark Count Rate	11
4.2	Finger Plots	13
4.2.1	LED Data	13
4.2.2	Signal to Noise Ratio	14
5	Assorted Projects for the ALPHA Experiment	16
5.1	ALPHA-g Barrel Veto Cooling Fan	17
5.1.1	Brief Overview of ALPHA-g and the Barrel Veto	17
5.1.2	Cooling Fan	18
5.2	Updating the Temperature Monitor Board	20
5.3	Simulating SVD Trigger Efficiency in ALPHA2	23
6	Appendix	28

TRIUMF

1 Motivation for SiPM Characterization

Phase 1 of the HAICU experiment aims to perform laser spectroscopy on magnetically trapped hydrogen atoms. This is the first step in a long-term goal of creating a hydrogen fountain used for spectroscopy and interferometry. Additionally, the HAICU apparatus is designed to be fully compatible with the ALPHA experiment to make similar measurements with anti-hydrogen.

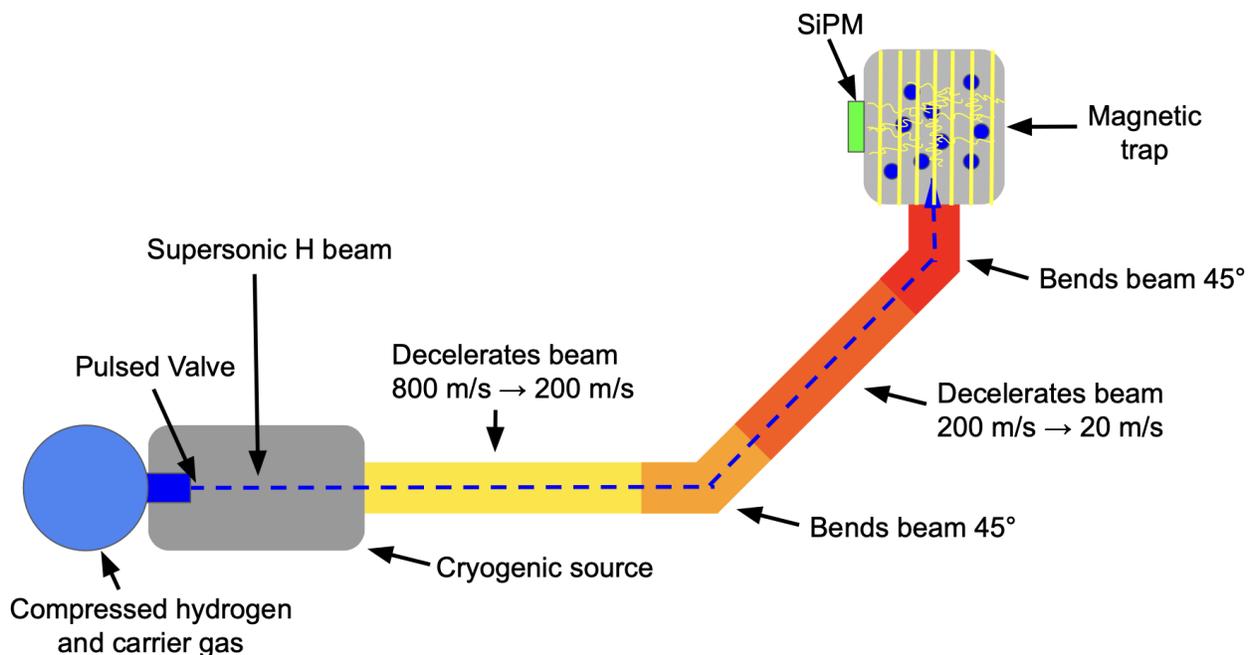


Figure 1: A schematic of HAICU phase 1. Compressed hydrogen and carrier gas exit a pulsed valve to create a supersonic hydrogen beam. This is then decelerated by a series of anti-Helmholtz coils, and steered into the magnetic trap using magnetic benders. Once in the trap, hydrogen is bombarded with lasers and the ultraviolet fluorescence is detected by a SiPM and MCP.

In HAICU, hydrogen atoms will be excited from the $n=1, l=0$ state to the $n=2, l=1$ state by a pulsed 121.6nm laser. When an atom is excited into this state, it will decay immediately, emitting a photon of the same wavelength in a process called fluorescence. To detect the vacuum ultraviolet light emitted during this transition, the Hamamatsu VUV4 SiPM will

be used. This is a new device to the HAICU group at TRIUMF, and in order to test its functionality in conjunction with digitizers, amplifiers, and data acquisition software, room temperature characterization was conducted.

2 Background Information

2.1 Overview of SiPMs

In most simple terms, Silicon Photomultipliers (SiPMs) are extremely sensitive light sensors. To put into perspective the extent of their sensitivity, if you were standing outside on a bright day there would be roughly 10^{14} photons hitting your eye each second. A SiPM, however, can detect a single photon.

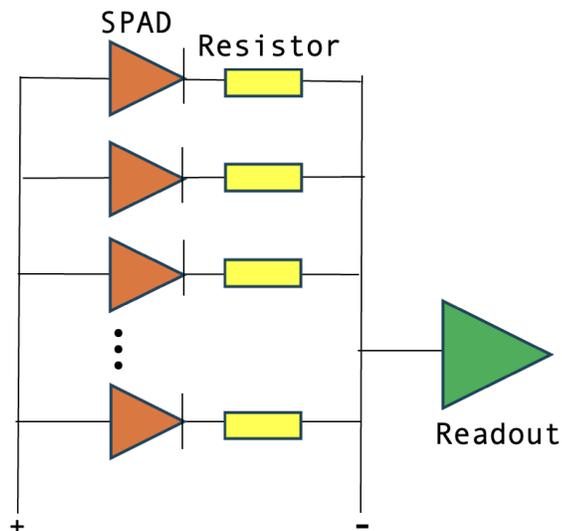


Figure 2: SiPMs are made up of an array of Single Photon Avalanche Diodes (SPADs). Each SPAD is connected to a quenching resistor, then connected in parallel to a single readout

Though the exact layout of a SPAD varies from model to model, in general a SPAD is composed of layers of doped silicon. Silicon can either have n-type doping which is an excess of electrons, or p-type doping which is an absence of electrons (ie, an excess of holes) in its lattice structure [3]. The point of contact between p and n doped silicon forms an

appropriately named "p-n junction", which forms a depletion region along the length of the contact point. Depletion region refers to absence of mobile charge carriers, leaving only the silicon impurities in this zone. Naturally, the p-type and n-type impurities drift towards areas of lower concentrations of them. This creates an electric field across the depletion region. This, along with a reverse bias voltage applied across the SPAD, creates a large electric field at the p-n junction in the SPAD. When a photon is absorbed into this region, an electron-hole pair will be created. The electron will begin accelerating downward due to the electric field, and once it gains sufficient energy, it can create another electron-hole pair. The electrons will continue to drift downwards knocking free more and more electrons, creating an avalanche of electrons. This process results in an output current and voltage drop over the SPAD. The voltage drop results in a diminishing electric field, which in turn stops the avalanche process, allowing the voltage across the SPAD to return to normal so that another photon can be detected [5].

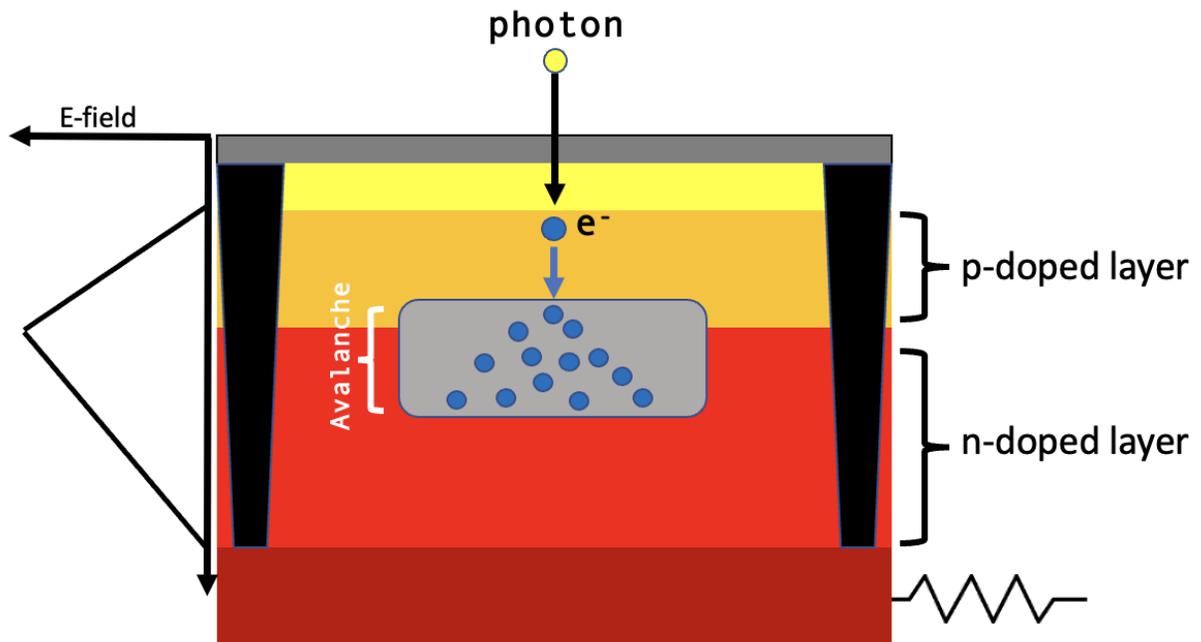


Figure 3: A simple diagram showing a side view of a single SPAD. SPADs are separated by trenches to minimize cross talk between cells.

2.2 Breakdown Voltage and IV Curves

As mentioned above, in order for the avalanche process to happen, a sufficiently large reverse bias voltage must be applied over the SiPM. The minimum voltage at which the avalanche process occurs is called the breakdown voltage. This voltage can be measured by applying different reverse bias voltages to the SiPM, and noting the corresponding current. The point at which the current is no longer constant and begins to increase rapidly, is the breakdown voltage of the device. Note that this process should be done for each individual channel on each individual SiPM since the precise breakdown voltage can vary from channel to channel and device to device.

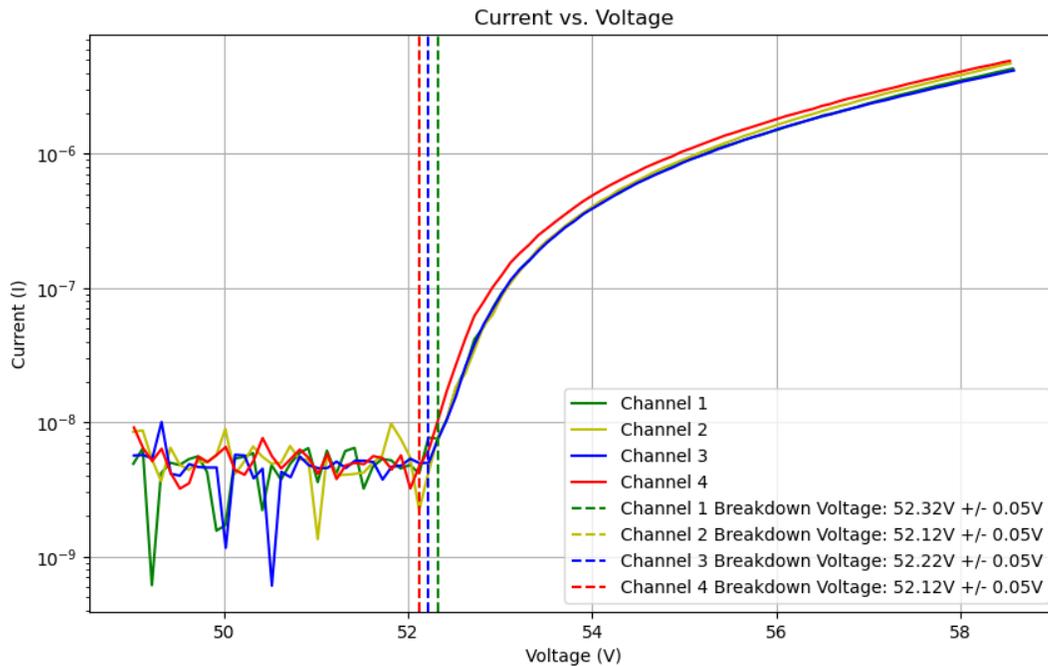


Figure 4: An IV curve of the Hamamatsu VUV4 taken at room temperature. In our measurements, voltages were increased by step sizes of 0.1V which results in breakdown voltage uncertainty of 0.05V. The uncertainty could be reduced by using smaller step sizes, but for our use an uncertainty of 0.05V is just fine.

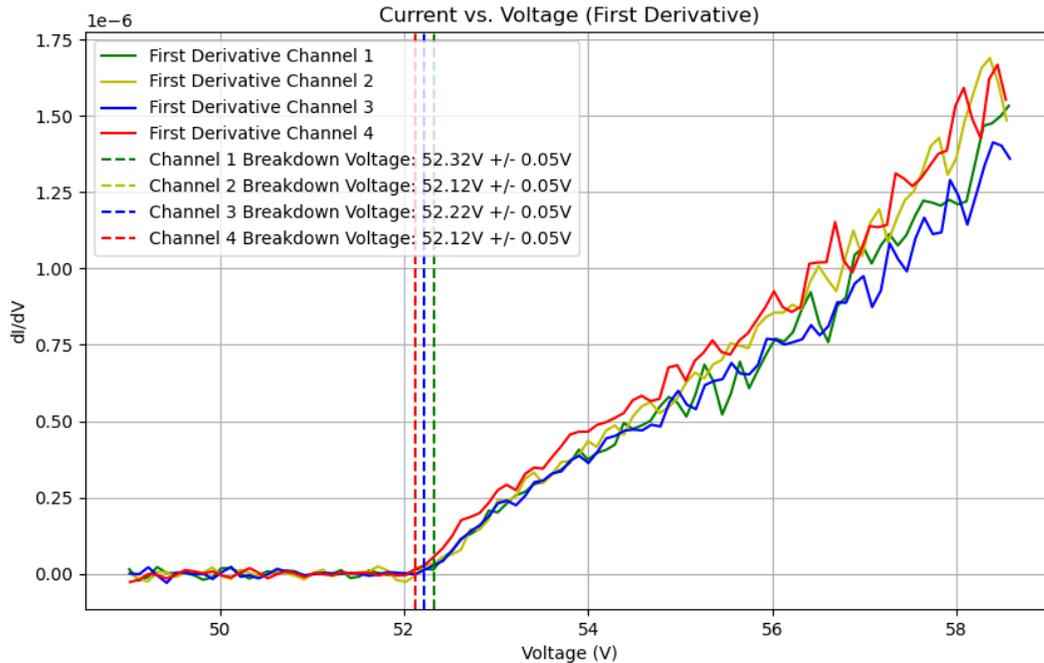


Figure 5: It can be easier to pick out the breakdown voltage by looking at the derivative of the current vs voltage. The point at which the derivative is non-zero is the breakdown voltage.

The measured breakdown voltages for each channel all lie within $52.12\text{V} \pm 0.05\text{V}$ and $52.32\text{V} \pm 0.05\text{V}$. This is well within the value quoted on the Hamamatsu datasheet (which is $53\text{V} \pm 5\text{V}$) [4].

2.3 Afterpulsing

Afterpulsing is when additional SiPM pulses occur directly following the primary avalanche event. This is due to a carrier (an electron or hole) from the primary avalanche getting trapped in a defect state in the silicon lattice, preventing it from participating in the initial pulse. After some time (which depends on the thermal energy of the carriers and the characteristics of the trap), the trapped carrier is released and may initialize a secondary avalanche. This secondary pulse is known as afterpulsing.

3 Setup and Apparatus

Since I am conducting room temperature characterization, the experimental apparatus is relatively simple (ie, no cooling system).

3.1 General Setup

For proper characterization, the SiPM should be tested in a dark environment with minimal to no light. To accomplish this, the SiPM was placed in a three gallon Rubbermaid container sprayed with a black rubber-like coating. Four small wire feed-through holes were drilled on one end of the box to connect the four SiPM channels to the LoLX Board (which is stored outside the box). These wires pass through a 3D printed cable holder to help minimize light entering the box through the wire holes. On the other end of the box, wiring to the LED could be fed from the pulser light-tightly using a PCB plug-in connector. Additional poster board was used to line the box, and a black cloth was placed over the box to further prevent light from entering.



Figure 6: A photo of the SiPM testing setup used for room temperature characterization at TRIUMF.

Device #1, LoLX Board: A custom built voltage control board. It runs using a NanoPi device which controls the voltage to each SiPM channel.

Device #2, CAEN DT5730B Digitizer: Digitizes the SiPMs analog signal to be further processed using a custom MIDAS waveform analysis toolkit.

Device #3, SiPM Box: A light-tight box where the SiPM and LED are held.

LED Specs: Epigap EOLD-255-012 UV LED, 30 Ohm resistor on a breadboard. 5.8V to power the LED, however it was typically run below this voltage to reduce the light produced.

Device #4, BK Precision 4065 Waveform generator: Used to control the voltage supplied to the LED and to provide the external trigger signal.

3.2 Triggers

Every subset of data saved within a run is called an event. The trigger is a function that tells us when to start acquiring data to create an event. The length and location of an event (relative to the trigger) can be controlled using two variables which are set in the ODB (online database):

1. **Buffer Organization:** An integer number relating to the number of samples saved in an event, by the following formula:

let β = buffer organization,

$$\text{num of samples} = 640 \times 2^{10-\beta} - 10$$

Example) A buffer organization value of 5 gives:

$$\text{num of samples} = 640 \times 2^5 - 10 = 40950 \text{ samples}$$

Each sample is 2ns long, so multiply by 2ns to get the length of the event in nanoseconds.

2. **Post Trigger:** The amount of samples saved after the trigger. This, in turn, also determines the number of samples saved before the trigger.

3.2.1 Threshold/Self Triggering

Threshold/self triggering is a trigger system that fires when the waveform passes some pre-defined threshold. In our experimental setup, the threshold value could be set in the ODB

for each individual channel. If any channel passes its set threshold, the event is saved across all channels. This method was mostly used when taking dark data (ie, no LED).

3.2.2 External Triggering

Events across all channels are saved when the external trigger fires. It is activated by a signal sent from the pulser, with 0V and -0.8V as the on and off logic. This trigger was primarily used when taking data with the LED, and was set up to coincide with the LED light pulse.

3.3 Data Acquisition and Processing

Data is collected using MIDAS (Maximum Integration Data Acquisition System) and read using MERCI (Modular Online Multi-threaded C++ based Waveform Toolkit). This allows for real-time waveform display, and for waveforms to be saved in MIDAS files. Waveforms are then fit using an autoregressive filter, timestamps and amplitudes of individual pulses are marked using a custom "HitFinder" module, region of interest can be set, then the processed data is output to ROOT files where further analysis can occur.

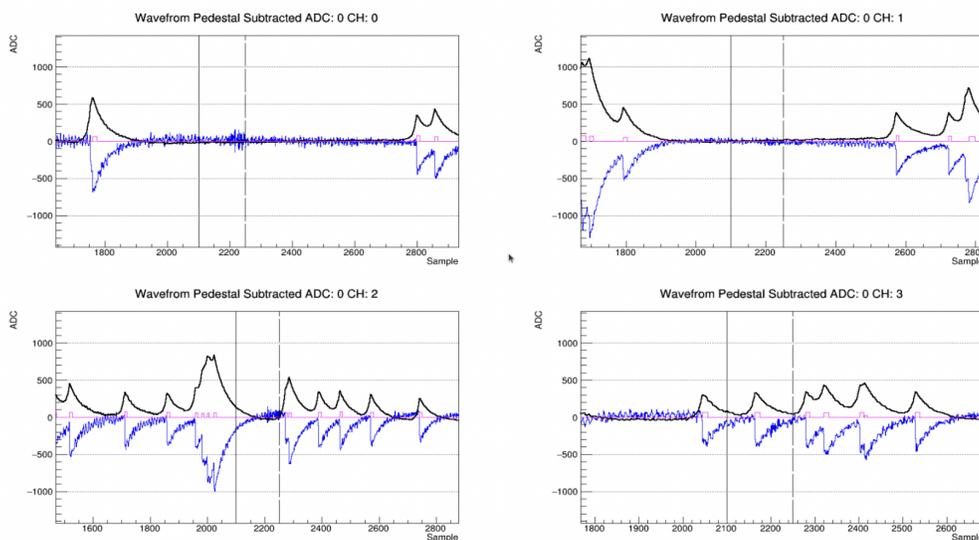


Figure 7: A screenshot of a dark noise waveform for each of the four channels. The inverted blue waveform is the raw unfiltered data. The black waveform is the data with an autoregressive filter. The pink rectangular peaks are pulses detected by the HitFinder module (height and width of the peak is proportional to the pulse size). The black vertical lines mark the beginning and end of the region of interest (irrelevant for this scenario since this is dark data – ie, we are not using an external trigger).

3.3.1 Dead Time

Dead time refers to the period of time immediately following the detection of an event where the detector is unable to record another event. For a SiPM this is the period of time after an avalanche, before the bias voltage returns to above the breakdown voltage. Dead time varies from device to device, and in an apparatus, the device with the longest dead time is what determines the dead time of the whole apparatus. For the particular setup described in this report, the digitizer has the largest dead time due to its limited data transfer bandwidth (ie, the rate at which data can be transmitted) and buffer size (ie, the maximum amount of data that a buffer can hold at any given time before it overflows). To study its dead time, we took dark data at room temperature at various overvoltages, and counted the number of events skipped by the digitizer. This number was then multiplied by the total length of each event to get the total dead time.

VoV	Dead Time [s]	Acquisition Time [s]	Ratio [Dead Time : Acquisition Time]
4	41.4219	601	0.069 → 6.9%
5	41.3940	615	0.067 → 6.7%
6	41.4188	603	0.069 → 6.9%
7	41.4235	602	0.069 → 6.9%

Table 1: Dead time and acquisition time for different overvoltages.

Something to note is that at lower temperatures, such as those intended to be used in the HAICU experiment, the digitizer's dead time should be irrelevant. This is because the SiPM dark count rate should be so small that it would not overwhelm the digitizer. Our measurement of dead time will be particularly useful if the cooling fails while HAICU is taking data, since it will give an estimate of how many events the detector has missed.

4 Characterization

4.1 Dark Data

Dark data refers to the output of the SiPM when no light is hitting its surface. Since a SiPM is a photon counter, you might expect its output to be a flat line (ie, no photons detected)

in the dark. However due to thermal noise, electrons can be excited in the silicon and cause an avalanche without a photon ever hitting the surface of the detector. This effect is known as dark noise.

4.1.1 Dark Count Rate

Dark count rate (DCR) is a measure of dark noise pulses per unit time. To be consistent with literature DCR values, we also normalize this number by the SiPM active area in mm^2 . To begin this calculation, we first histogram the time difference between pairs of pulses over many events. For short timescales (ie, less than 50ns) counts are dominated by non dark pulses (mostly afterpulsing). For this reason, only the tail of this plot (50-2500ns) is fit with a negative exponential.

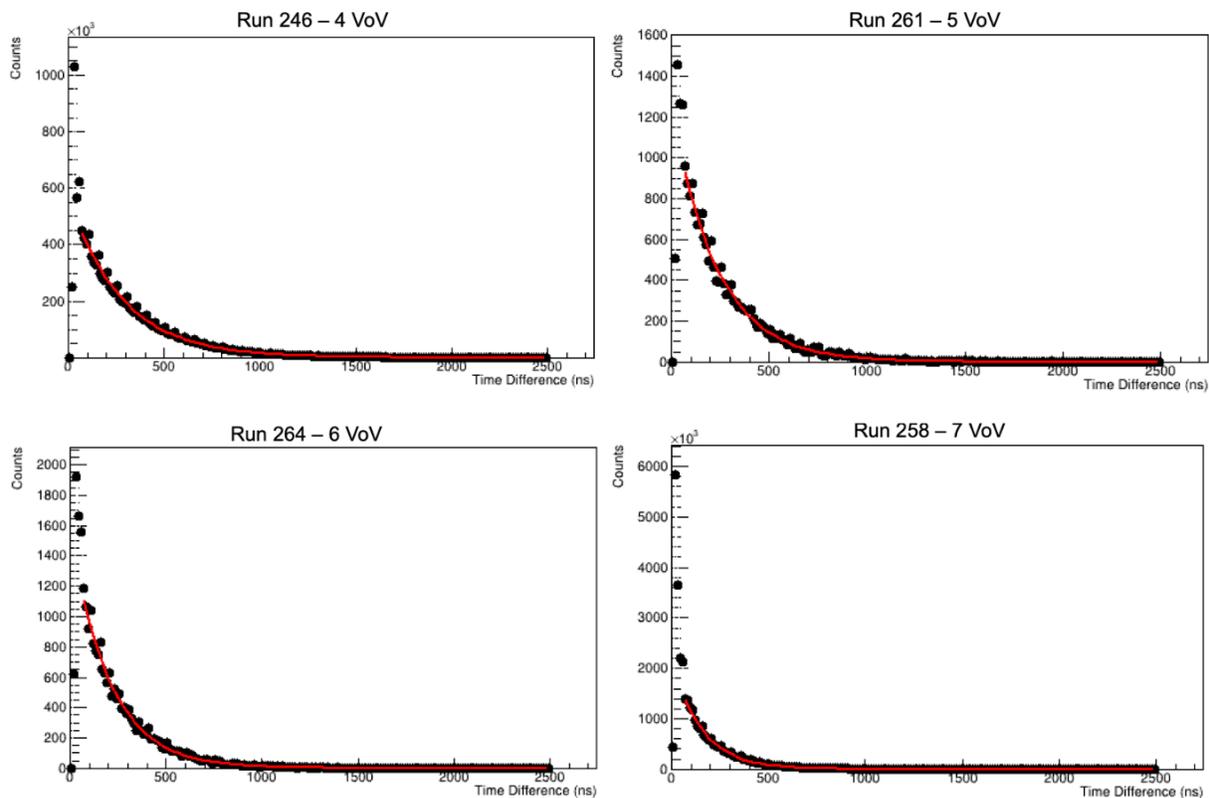


Figure 8: Histograms showing the time difference between pairs of pulses over all events, at different overvoltages.

The reciprocal of the time constant of the negative exponential gives the pulse rate, which is normalized by SiPM area (in our case, the Hamamtsu VUV4 has an active area of 9 mm^2).

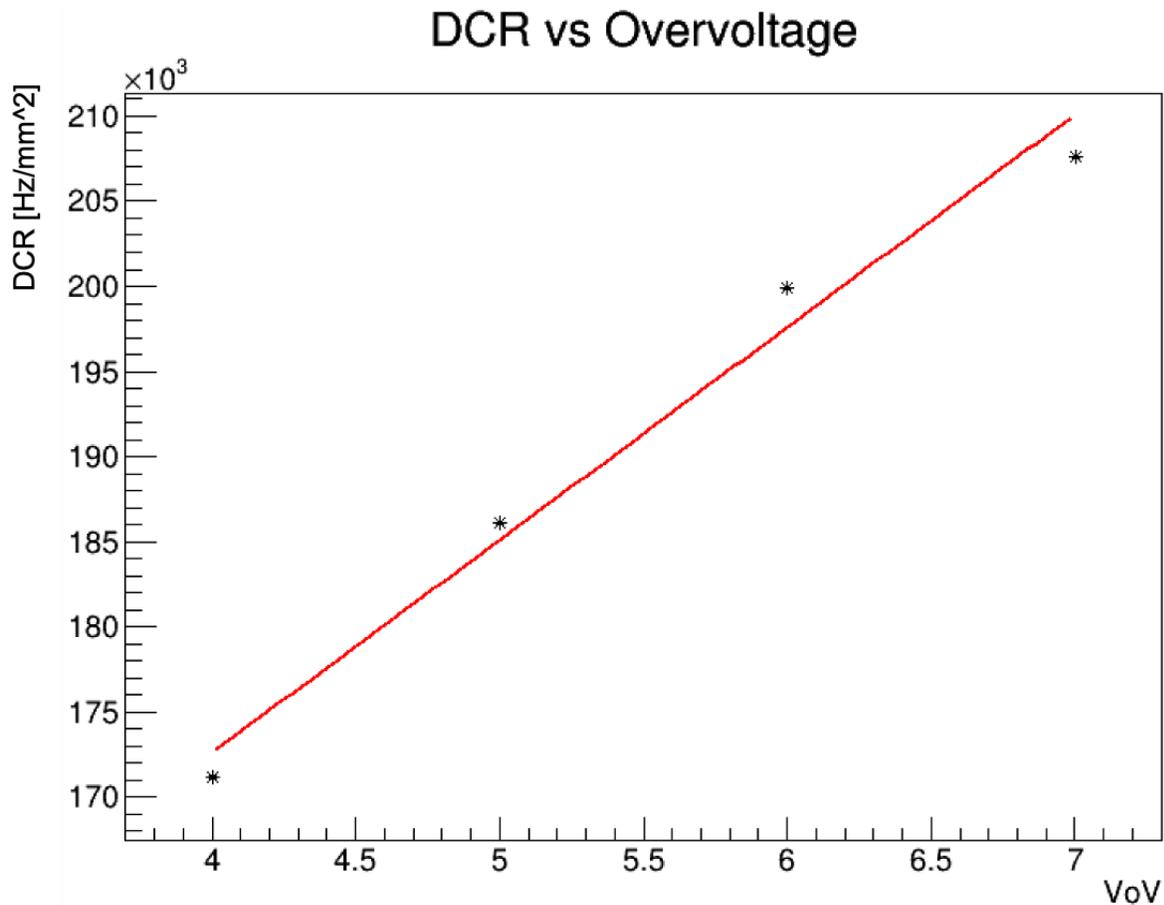


Figure 9: A visual comparison of DCR's at different overvoltages. Note that the relationship between DCR and overvoltage is not linear. The fit is meant to make the relative positions of each point more clear.

Overvoltage [V]	DCR [kHz/mm ²]
4	171.17 ± 0.06
5	186.09 ± 0.06
6	199.88 ± 0.06
7	207.59 ± 0.07

Table 2: Calculated values of DCR at different overvoltages.

These dark count rates are of the expected order of magnitude. A paper from 2019 [2] measured the dark count rate of the Hamamatsu VUV4 at cold temperatures. It can be seen that DCR scales by roughly an order of magnitude for each 20K of temperature increase. At 4V overvoltage and 233K, the DCR was measured to be $\sim 2 \times 10^2$ Hz/mm², so at room temperature (~ 293 K), we'd expect our dark count rate to be about 3 orders of magnitude

higher (ie, $\sim 2 * 10^5$ Hz/mm²). This is quite close to our measured value of $1.71170 * 10^5$ Hz/mm² \pm 60 Hz/mm² at 4V overvoltage.

4.2 Finger Plots

Finger plots are a great way to characterize SiPMs and extract useful properties such as signal to noise ratio, and gain. They are typically made by histogramming either SiPM pulse amplitude or SiPM pulse charge. Due to the discrete nature of low light SiPM outputs, there will be an excess of counts in bins that correspond to one photon pulse size, two photon pulse size, and so on. In dark or low light conditions, it is common to only see one peak corresponding to a one photon pulse size.

4.2.1 LED Data

When taking data with an LED, an external trigger is used. This trigger is synchronized with the pulsing of the LED. This allows us to look for SiPM pulses in a region directly following a light pulse from the LED.

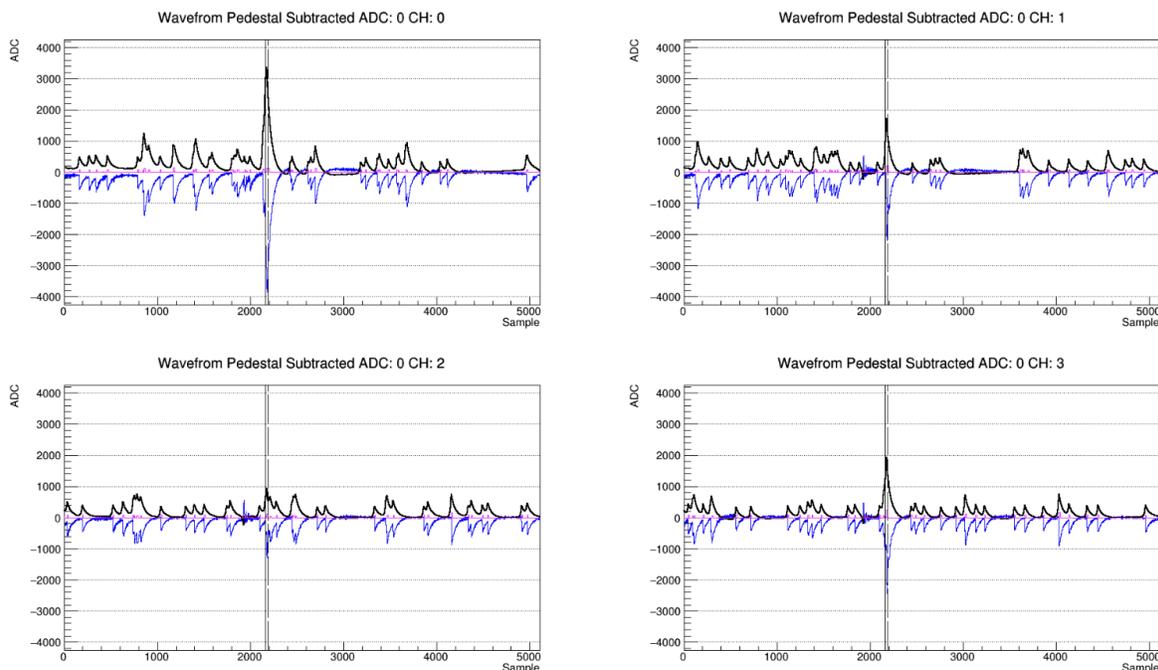


Figure 10: Some sample waveforms from LED triggered data. The vertical lines represent the region in which we are looking for pulses. This region corresponds to the timing of the light from the LED.

Though full LED characterization was cut short due to time constraints, some plots were made which demonstrate the basic qualitative features of our SiPM device.

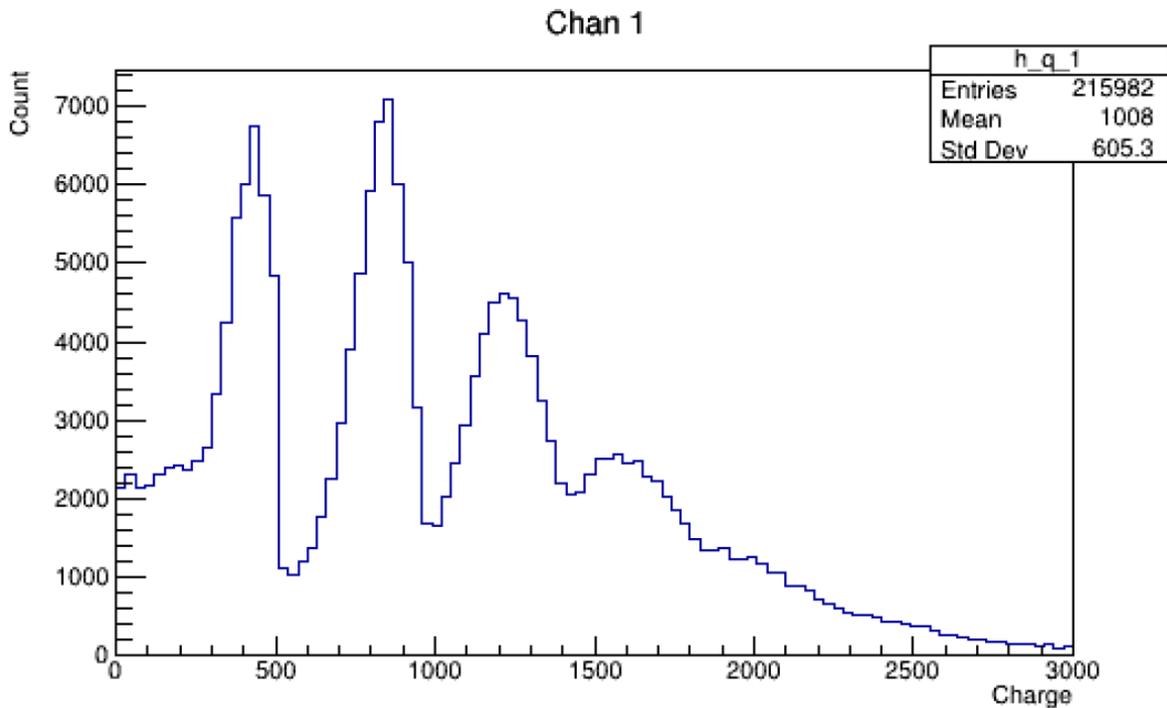


Figure 11: An charge finger plot taken with LED triggered data at 4V overvoltage. Notice the four major peaks, corresponding to the number of photons that our device has detected. This plot only shows data for channel 1 of 4 on our device, though the other channels also have very similar plots.

4.2.2 Signal to Noise Ratio

Calculating the signal to noise ratio of our device is quite simple, however a few careful steps must be taken into consideration when acquiring data. Due to dark noise, the LED is not necessary when taking data. Additionally, all pulses (including electronic noise) must be flagged and measured. In practice, this is done by setting the minimum charge parameter in the HitFinder module to 0. With this done we can histogram the height of all the pulses from our dark data. What we see is two peaks: the first peak corresponding to the amplitude of noise pulses in our waveform, and the second peak corresponding to the amplitude of single photon pulses. To calculate the signal to noise ratio, we fit each peak with a gaussian, determine the mean position of the 1 photon peak as well as the full width half maximum of

the noise peak, then divide the two values. Written out explicitly we get:

$$\text{STNR} = \frac{\text{FWHM}_0}{\mu_1}$$

where STNR is the signal to noise ratio, FWHM_0 is the full width half max of the noise peak, and μ_1 is the mean value of the 1 photon peak.

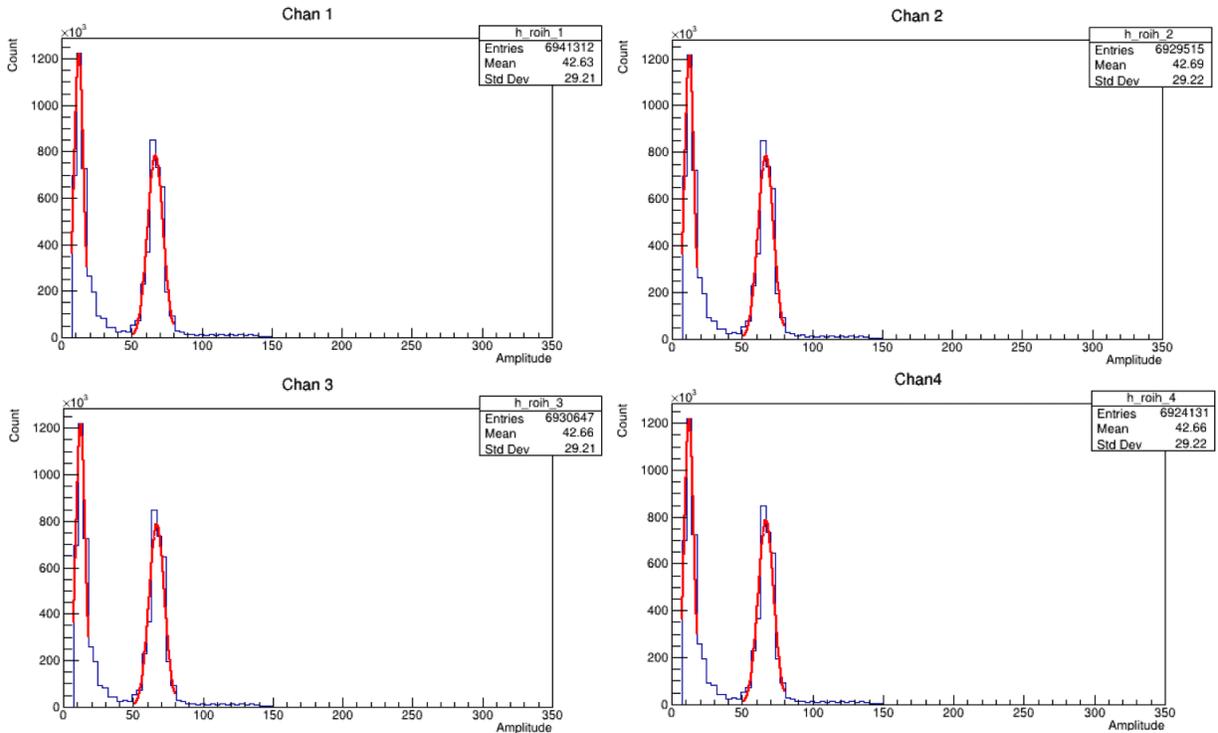


Figure 12: Four finger plots (one for each channel) used to calculate the signal to noise ratio of our device at 4V overvoltage.

Channel	Signal to Noise Ratio
1	8.45 ± 0.01
2	8.42 ± 0.01
3	8.45 ± 0.01
4	8.44 ± 0.01

Table 3: Signal to noise ratios for different channels at 4V overvoltage.

The signal to noise ratio plots and exact values for other overvoltages can be found in the appendix.

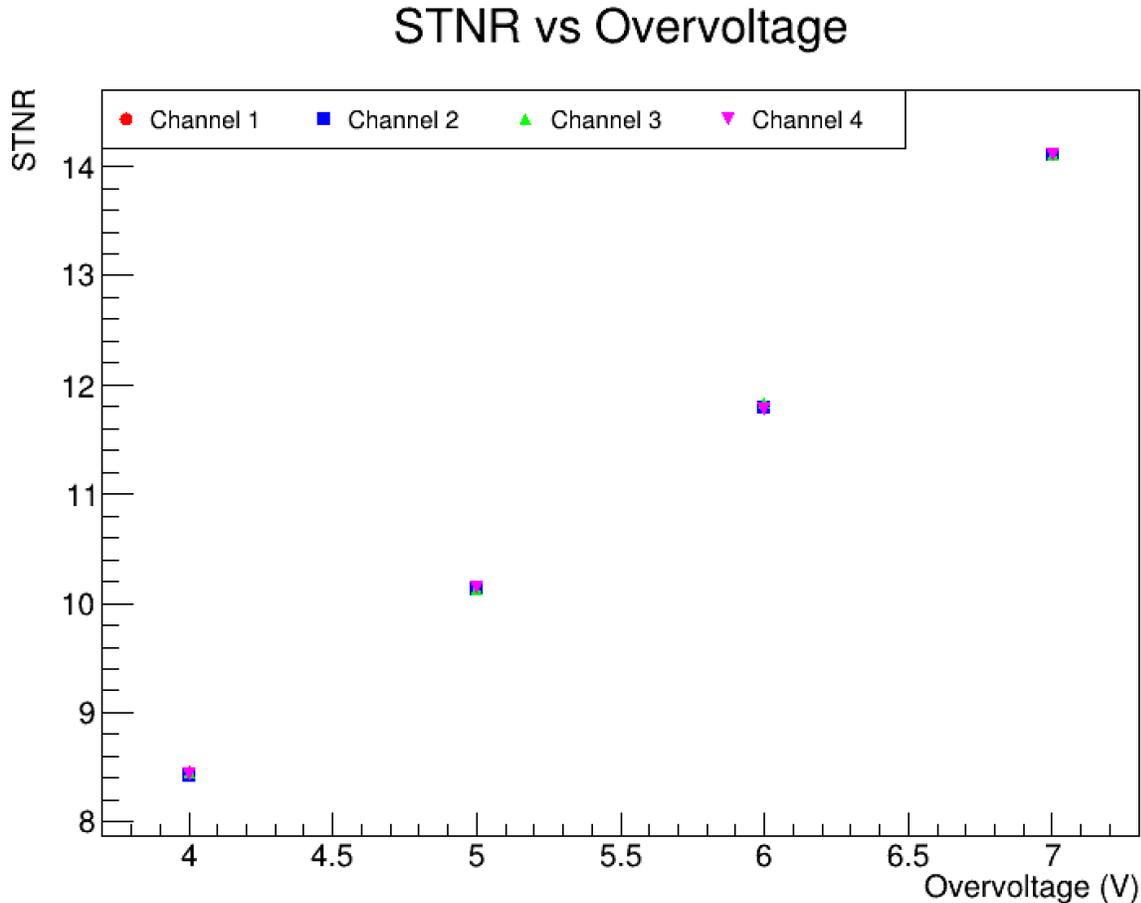


Figure 13: An overview plot showing the signal to noise ratios for each overvoltage and channel.

CERN

5 Assorted Projects for the ALPHA Experiment

My time at CERN was split between attending the student lecture program, helping take data for the ALPHA2 experiment, and working on small projects for use in the ALPHA-g experiment. The results of the ALPHA2 experiment are ongoing and confidential so they cannot be discussed in this report. Below focuses on a handful of the projects I worked on during my stay at CERN.

5.1 ALPHA-g Barrel Veto Cooling Fan

5.1.1 Brief Overview of ALPHA-g and the Barrel Veto

ALPHA-g is an experiment located in the anti-matter factory at CERN. Its goal is to verify the weak equivalence principle with antimatter, through studying the effect of gravity on anti-hydrogen. This is done by catching anti-hydrogen in a vertical magnetic trap, then ramping down the magnetic field allowing the atoms to free fall. Due to the kinetic energy each atom possesses, some atoms will overcome earth's gravitational field, travel upwards, and annihilate against the wall higher in the trap. However, the majority of atoms fall downwards.

When anti-hydrogen annihilates on the detector walls, charged pions are produced. By detecting the pions and their trajectories, the original annihilation location can be reconstructed. The pions are tracked and detected using a Time Projection Chamber and backgrounds are removed using the Barrel Veto.

1. Time Projection Chamber (TPC)

This detector records the path of the annihilation products. In ALPHA-g, the TPC is filled with a gaseous mixture of argon and carbon dioxide and an electric field, perpendicular to the detector axis, is present. When a charged particle traverses the TPC, electron-ion pairs are produced from interactions with the gas. The electrons drift towards the anode wires and the ions drift towards the cathode. The electrons undergo a charge multiplication as they near the anodes, which is read out from our detector. Since the drift time of each electron is proportional to the distance it has travelled in the detector, it is possible to determine the spatial position of the interaction and to reconstruct the vertex of the original annihilation.

2. Barrel Veto (BV)

The main goal of the Barrel Veto is to measure the time of flight of particles traversing the TPC. This is particularly useful for distinguishing between cosmic ray and charged pion tracks in the TPC. Though both pions produced from an anti-hydrogen annihilation and cosmic rays will leave similar tracks in the TPC, the timing of the hits on

the BV will differ. The pions produced from an annihilation near the center of the TPC will be detected nearly coincidentally, whereas a cosmic ray passing through the detector will produce two charge avalanches separated by roughly the time of flight of the cosmic ray in our detector. By measuring the timing of avalanches in the TPC, it is possible to distinguish between background events (ie, cosmic rays) and meaningful annihilation events.

The Barrel Veto is made from 64 plastic scintillator bars surrounding the TPC, each bar with a SiPM at either end [6]. When ionizing radiation interacts with a bar, light is produced isotropically within the bar. This light is detected by the SiPMs, and thus the timing of particles passing through the detector can be inferred.

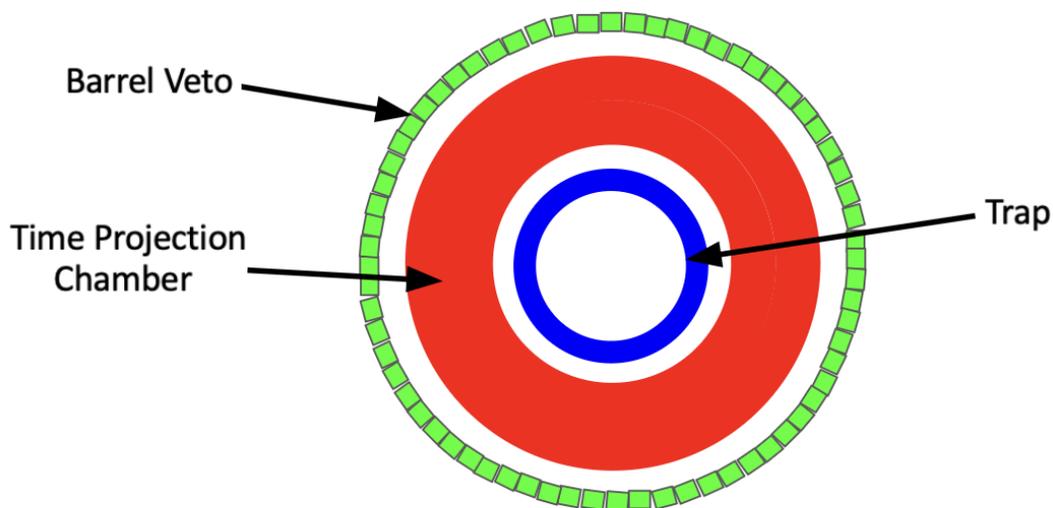


Figure 14: A simplified figure showing the relative component positions in the ALPHA-g detector.

5.1.2 Cooling Fan

To read out the SiPM waveforms from the Barrel Veto, a set of Amplifier-Shaper-Discriminator (ASD) cards are used. These cards can become quite warm during operation, and thus need regular airflow to be cooled. This is done by connecting a computer fan to two pipes which enter the Barrel Veto and lead to an ASD card. In order to attach a single fan to both pipes,

an airflow splitter was designed in CAD. The final design is roughly 12cm long, and connects an 8cm diameter computer fan to two pipes each with diameter 4cm.

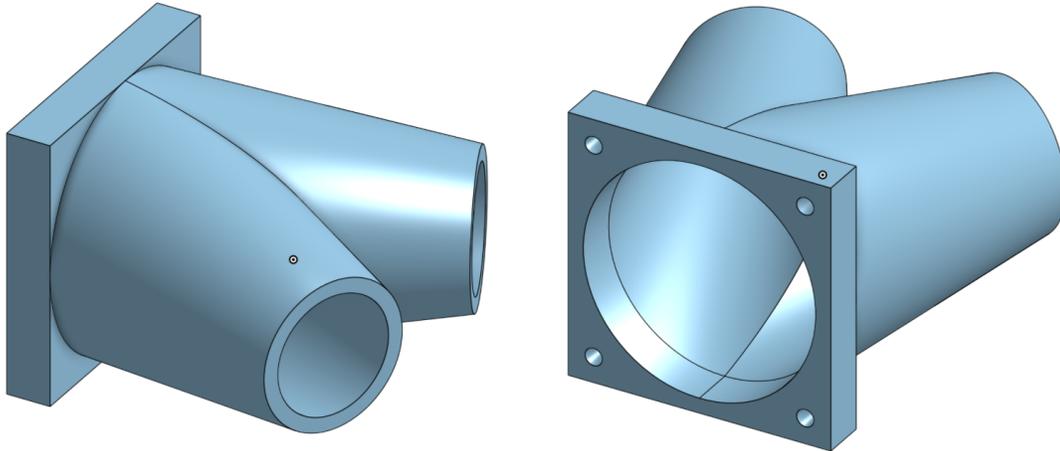


Figure 15: Two angles of the airflow splitter CAD design. Holes for threaded inserts were added so the fan can be screwed to the apparatus.

Once the CAD design was printed, the fan could be screwed to the airflow splitter and hooked up to the 12V power supply. The wires from the fan needed to be soldered to the wire from the power supply, then wrapped in electrical tape for insulation. Lastly, the fan could be powered up and connected to the tubes on the Barrel Veto – just in time for ALPHA-g to begin taking data!



Figure 16: Images of the cooling fan attached to the Barrel Veto. The pipes lead to 2 separate ASD cards.

5.2 Updating the Temperature Monitor Board

The temperature monitor board is used to give real-time temperature readings along the length of the ALPHA-g TPC. A custom made circuit board is used in conjunction with a Raspberry Pi device to read the temperatures and send them to a live display. Though this board has two ADC chips allowing for 32 separate channel readouts, only 16 of them were being used. Both software and hardware updates were needed for all 32 channels to be used properly.

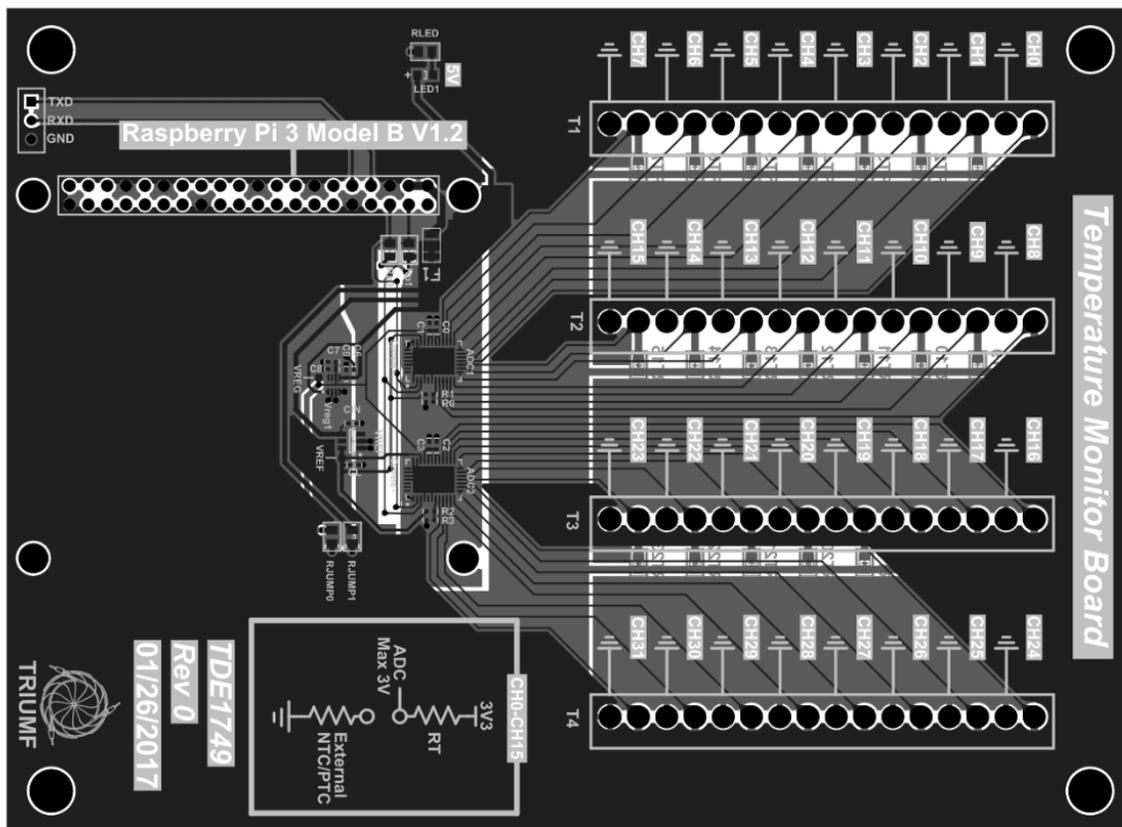


Figure 17: A schematic showing the layout of the temperature monitor board.

The code that reads the resistance of each channel, converts the voltage to a temperature, then sends the data to a live display was originally written by Lars Martin (TRIUMF). In brief, the script functions as follows:

1. Defining Addresses

In order to read out data, the I2C address of each ADC must be defined. This can be done by:

```
#define ADC_ADD1 ADC1_I2C_Address
```

```
#define ADC_ADD2 ADC2_I2C_Address
```

Additionally, each input channel address must be defined and mapped to an input on the ADC. For example, the following two lines map ADC input 8 to channel 0:

```
#define IN8 ADC_IN8_Address
```

```
#define CH0 IN8
```

2. Main Loop

At the beginning of the `Main()` loop, the I2C connection is reinitialized:

```
system("sudo raspi-config nonint do_i2c 0");
```

This will reinitialize the I2C connection every time the front end program is restarted.

The main loop uses a "while loop" to continuously read data from the ADCs unless a shutdown request is received:

```
while(!mfe->fShutdownRequested) { ... }
```

To read out data, each channel must be looped over. Depending on the ADC of the current channel, the I2C interface with the current ADC must be initialized. This is done by the following "if" statement:

```
if (i<16) {fd = wiringPiI2CSetup(ADC_ADD1);
```

```
else{fd = wiringPiI2CSetup(ADC_ADD2);}
```

The next block of code writes to the ADC to initiate a conversion, then reads the result.

The buff array stores the data bytes read from the ADC.

```
write(fd, buff, 2);
```

```
read(fd, buff, 3);
```

Next, the data read from the ADC is bit shifted and combined into a single integer value:

```
x=buff[0];  
y=buff[1];  
z=buff[2];  
data=((x<<10)|(y<<2)|(z>>6));
```

The `data` value read from the ADC is converted to an actual voltage:

```
NTC_voltage = data * (vref / (2 * gain)) / 65536;
```

Where `vref` is the maximum voltage the ADC can measure, `gain` is the factor by which the input signal was amplified before being digitized by the ADC, and `65536` is the maximum possible value for a 16-bit ADC (2^{16}). In our case, `vref` is 2.5V, and `gain` is 256. The voltage is converted to a resistance using the voltage divider formula, where `rt` is the resistance of the resistor attached to the connector (49.9 k Ω in our case):

```
NTC_resistance=(rt*NTC_voltage)/(vref-NTC_voltage);
```

Now, the resistance can be converted to a temperature value. The equation for this was derived from the PS103J2 NTC reference table:

```
NTC_temp[readnum][i]=-21.67*log(NTC_resistance)+224.62;
```

Note that the array `NTC_temp` holds multiple temperature values at once in the `readnum` slot. This is used for the moving average calculation, where the average is calculated over a customizable number of samples. Lastly, the temperature and moving average temperature data is sent to MIDAS using:

```
WVD(mfe, eq, "Cooling T", N_ACTIVE_CHANNELS, NTC_temp[readnum]);  
WVD(mfe, eq, "Cooling avgT", N_ACTIVE_CHANNELS, NTC_avge);
```

Where `mfe` is an instance of the `TMFE` class and is used to connect to the MIDAS front end, and `eq` is an instance of the `TMFEEquipment` class which specifies the specific equipment used.

In terms of hardware, two new connectors needed to be soldered to the 2nd ADC, as well as 16 surface mounting resistors. Once both the software and hardware were up and running,

a live display of temperature reading across the Barrel Veto were displayed on a custom MIDAS page.

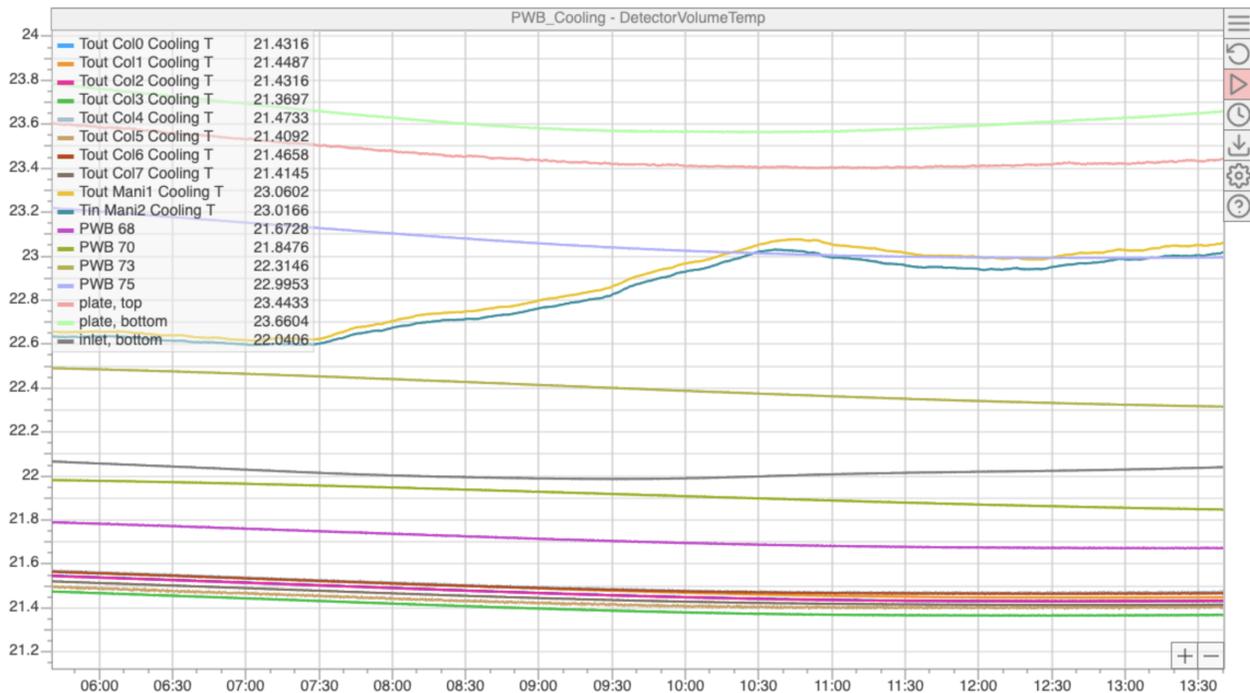


Figure 18: A screenshot of the live display of temperatures read out from the temperature monitor board.

5.3 Simulating SVD Trigger Efficiency in ALPHA2

The SVD is a particle detector that tracks anti-hydrogen annihilation products within the ALPHA2 apparatus. It is comprised of 72 silicon modules that surround ALPHA2 in three layers. Each silicon module allows for precise readout of the annihilation position. Each pixel functions similar to a SiPM, where a reverse bias voltage is applied to a P-N junction allowing for charge multiplication upon interaction with ionizing radiation [1]. As a charged particle traverses the 3 layers of the silicon detector, its path can be reconstructed. Doing this for all the annihilation products permits the reconstruction of the anti-proton annihilation position, called the "vertex". The SVD trigger is designed to minimize the number of events due to cosmic rays which only reads out events that have at least 2 hits in the inner layer of the detector, and 1 hit in each of the other two detector layers.

The goal of the simulation I ran was to determine the efficiency of the Silicon Vertex Detector

(SVD) for anti-proton annihilation, in a specific region of interest. The region of interest was selected to match the location of anti-hydrogen annihilations in the "slow dump" experiment. This experiment consisted of recatching ionized anti-hydrogen atoms (ie, anti-protons), then slowly releasing them from the trap and allowing them to be detected by the SVD. To find the region of interest, a histogram showing the z-position of annihilations in our detector was fit with a double gaussian (see Fig. 19). Then, the region of interest was taken to be from minus 3 sigma of the mean of the first gaussian, to plus 3 sigma from the mean of the second gaussian:

$$\text{Region of Interest} = (\mu_1 - 3\sigma_1, \mu_2 + 3\sigma_2)$$

where μ_1, μ_2 is the mean of each gaussian, and σ_1, σ_2 are the standard deviation of each respectively.

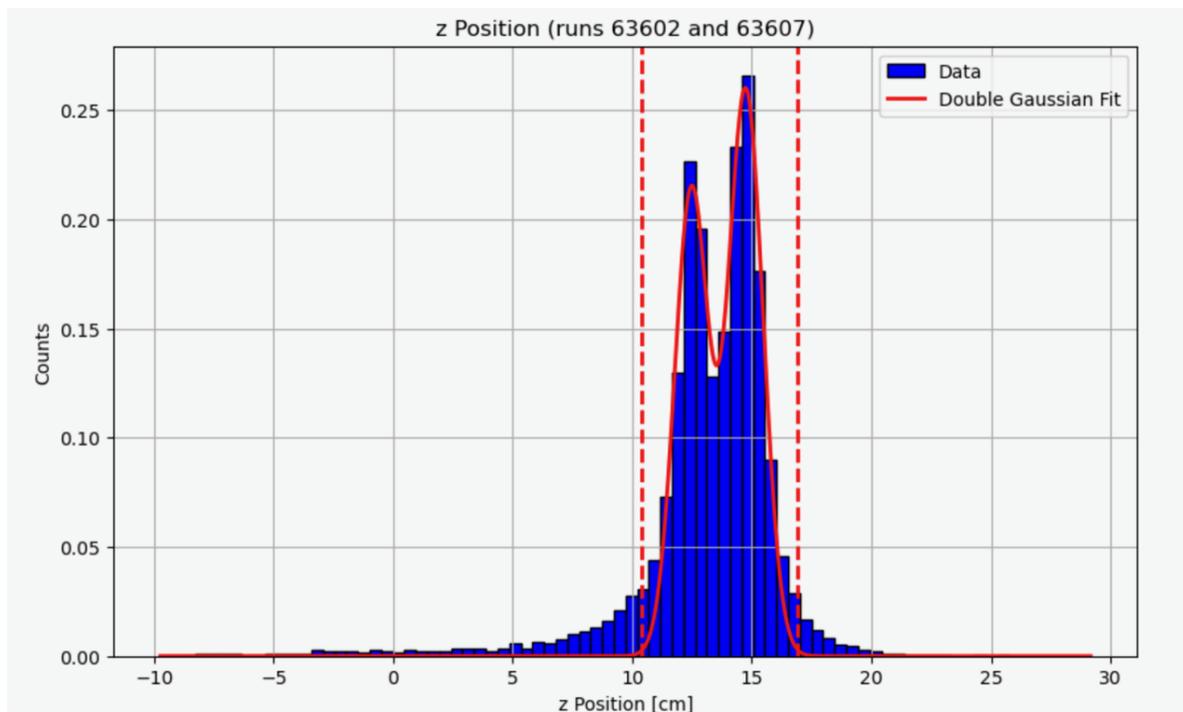


Figure 19: The histogram used to find the region of interest used in the simulation. The area between the dashed red lines corresponds to the region of interest.

With the region of interest set, a Monte Carlo simulation could be run. A uniform distribution of 1 million annihilation events were simulated within the region of interest, and the trigger efficiency was histogrammed against the true z-position in the detector. Trigger efficiency at each z-position was calculated as follows:

$$\text{Trigger Efficiency} = \frac{\text{Number of Triggers}}{\text{Number of annihilations}}$$

Error bars for each point were calculated using Poisson statistics, where:

$$\text{Error} = \frac{\sqrt{\text{Number of Triggers}}}{\text{Number of annihilations}}$$

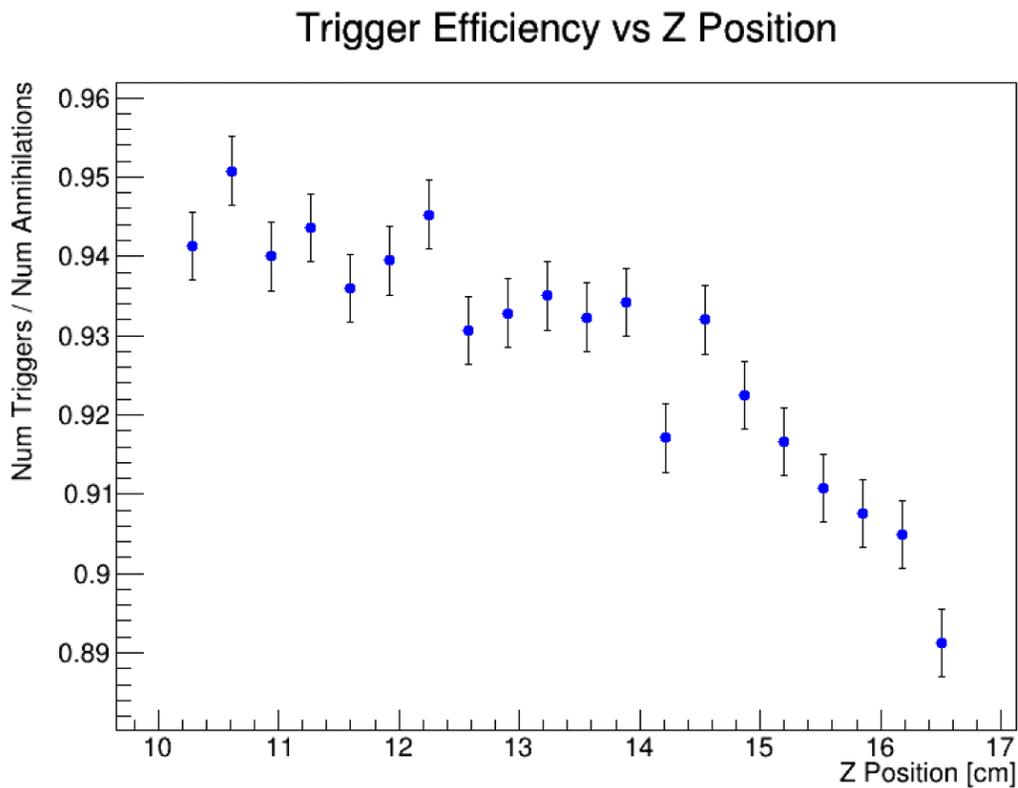


Figure 20: A plot showing the SVD trigger efficiency vs z-position in the detector.

Lastly, the mean efficiency within the region of interest could be calculated. The error on this value was taken to be highest and lowest values within this region of interest (ROI).

$$\text{SVD Trigger Efficiency (within ROI)} = 0.928^{+0.023}_{-0.037}$$

Acknowledgments

Special thanks to my supervisors (Andrea Capra, Ina Carli, and Makoto Fujiwara) at both TRIUMF and CERN for their support and guidance over the past four months. Thanks to all the researchers, engineers, graduate and undergraduate students in both the HAICU and ALPHA collaborations for providing support when needed, and creating an enriching work environment. Thank you to the IPP Student Fellowship Programme, the TRIUMF Summer Student Programme, and the CERN Summer Student Programme for this opportunity and all the memories I've made. Lastly, thanks to my friends and family for all the long distance phone calls, text messages, and general good vibes during my stay at CERN.

References

- [1] ALPHA Experiment. Silicon vertex detector. <https://alpha.web.cern.ch/component/silicon-vertex-detector-svd>, 2018. Accessed: 2024-08-20. 23
- [2] G. Gallina, P. Giampa, F. Retière, J. Kroeger, G. Zhang, M. Ward, P. Margetak, G. Li, T. Tsang, L. Doria, and S. Al Kharusi. Characterization of the hamamatsu vuv4 mppcs for nex0. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 940:371–379, October 2019. 12
- [3] Stefan Gundacker and Arjan Heering. The silicon photomultiplier: fundamentals and applications of a modern solid-state photon detector. *IOP Science*, 65(17TR01), 2020. 3
- [4] Hamamatsu. Hamamatsu vuv-mppc datasheet. [https://hamamatsu-su/files/uploads/pdf/3_mppc/s13370_vuv4-mppc_b_\(1\).pdf](https://hamamatsu-su/files/uploads/pdf/3_mppc/s13370_vuv4-mppc_b_(1).pdf), 2017. Product Flyer. 6
- [5] SensL. An introduction to the silicon photomultiplier. https://websites.umich.edu/~ners580/ners-bioe_481/lectures/pdfs/2017-02-SensL_IntroSiPM.pdf, 2017. Technical Note. 4
- [6] Gareth Smith. *Characterization and analysis development for the ALPHA-g barrel scintillator*. PhD thesis, University of British Columbia, 2021. 18

6 Appendix

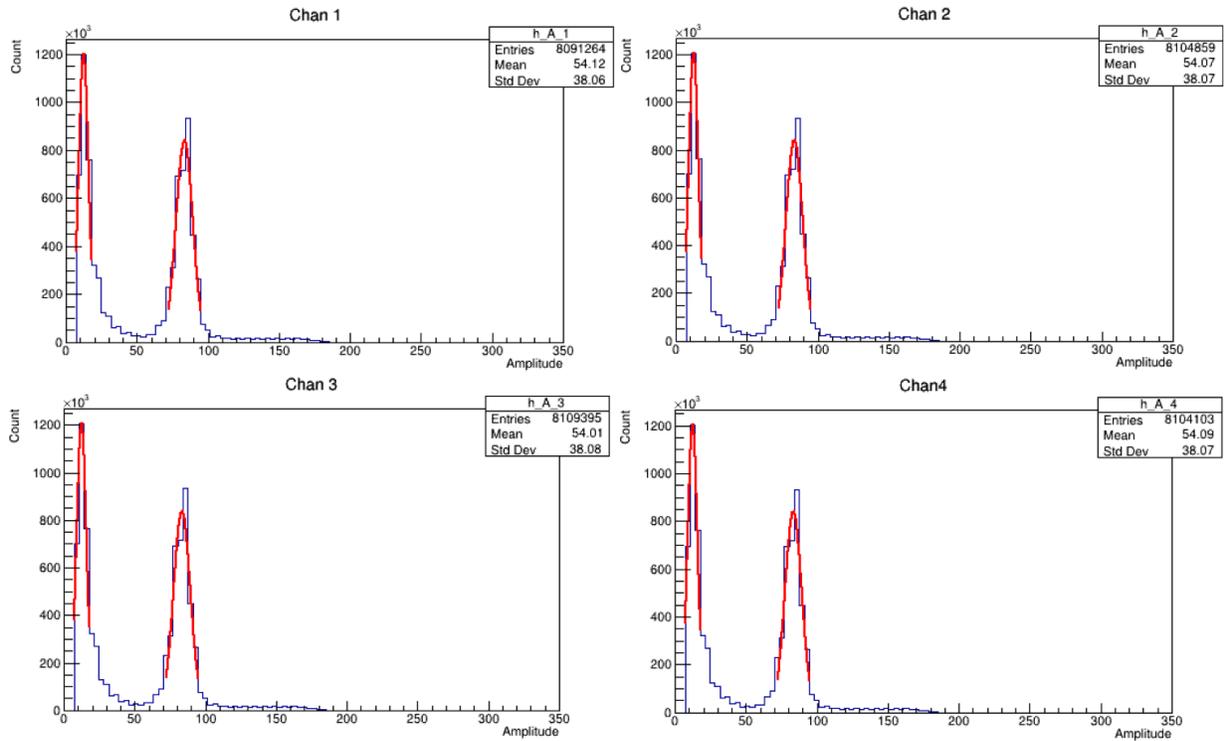


Figure 21: Four finger plots (one for each channel) used to calculate the signal to noise ratio of our device at 5V overvoltage.

Channel	Signal to Noise Ratio
1	10.12 ± 0.01
2	10.13 ± 0.01
3	10.12 ± 0.01
4	10.14 ± 0.01

Table 4: Signal to noise ratios for different channels at 5V overvoltage.

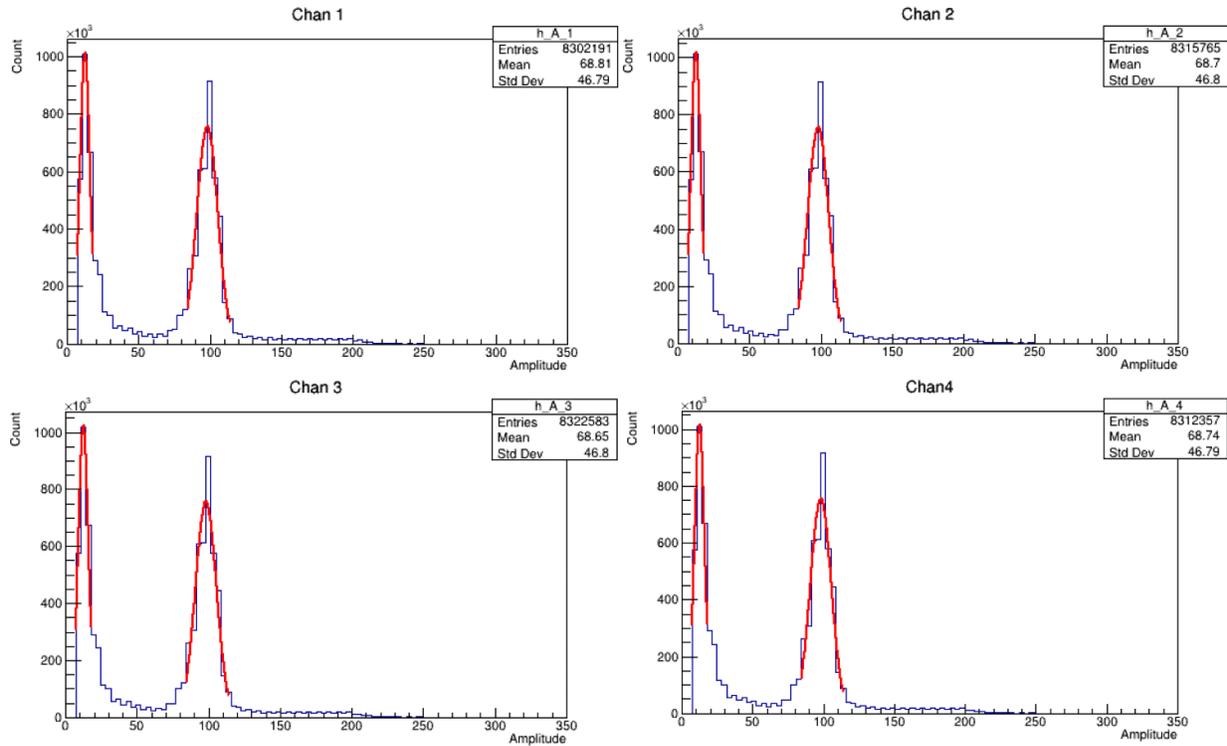


Figure 22: Four finger plots (one for each channel) used to calculate the signal to noise ratio of our device at 6V overvoltage.

Channel	Signal to Noise Ratio
1	11.79 ± 0.02
2	11.79 ± 0.02
3	11.83 ± 0.02
4	11.78 ± 0.02

Table 5: Signal to noise ratios for different channels at 6V overvoltage.

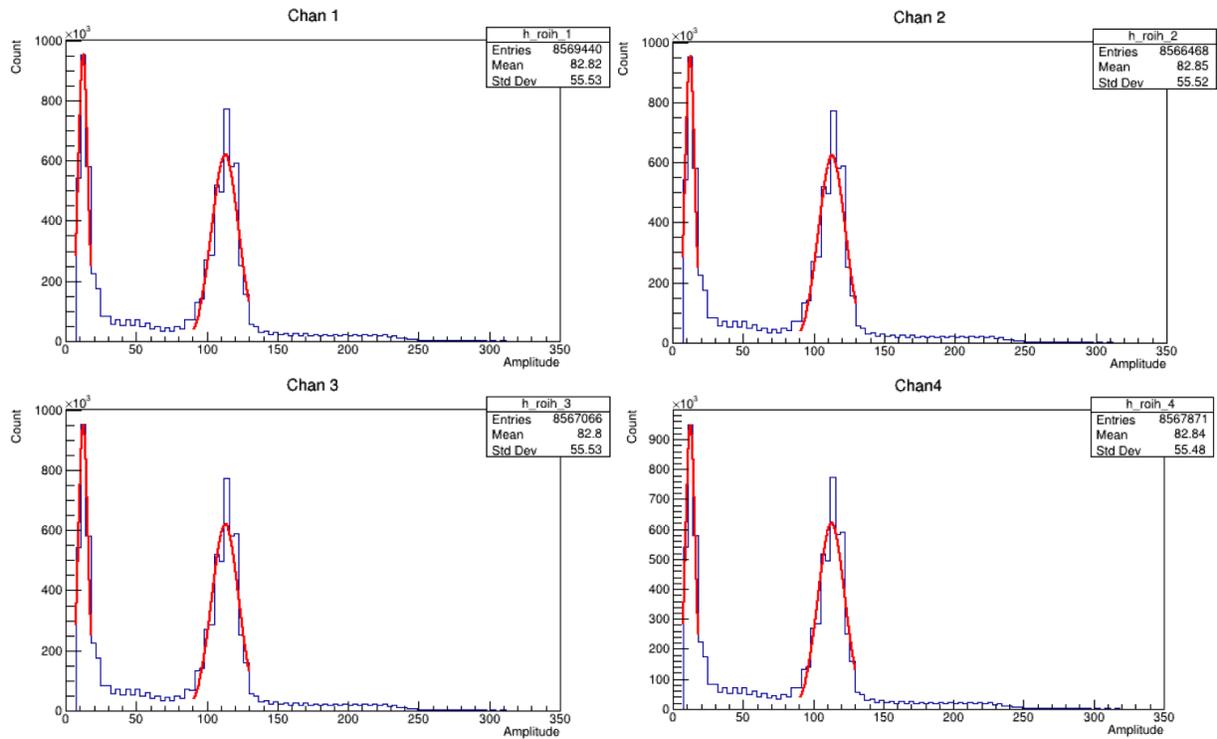


Figure 23: Four finger plots (one for each channel) used to calculate the signal to noise ratio of our device at 7V overvoltage.

Channel	Signal to Noise Ratio
1	14.11 ± 0.02
2	14.11 ± 0.02
3	14.11 ± 0.02
4	14.10 ± 0.02

Table 6: Signal to noise ratios for different channels at 7V overvoltage.