Timing Analysis on Fast Pulses in a Picosecond Resolution Photodetector

Aaron Meyer

Department of Physics
University of Chicago
Chicago, IL 60637

May 21, 2012

Abstract

Much work has been done by the Large Area Picosecond Photodetector (LAPPD) Group on research and development of a large area fast photodetector. In this project, I took laser data at Argonne National Laboratory to test the capabilities of the hardware. This involved helping realign optics, ensuring the laser intensity is correct, and interpreting signals in both hardware and software. The data was interpreted using Matlab scripts, in which I calculated the times of the pulses and the associated uncertainties. The calculation was done using a linear least-squares method. Data optimizations are considered and both resolution and goodness of fit are analyzed for the data set.

1 Motivation

With the advent of the LHC era, physicists have a new collection of data to search through for physics beyond the Standard Model. One interesting search that is ongoing is the possiblity of detecting new particles due to models which are otherwise unconfirmed. Such particles could include supersymmetric partners to Standard Model particles, magnetic monopoles, mirror particles, and others [1]. Such particles would be massive, typically predicted to have masses of order $100\text{GeV}/c^2$ or greater, and limits currently exist on their masses [2].

These new particles could be long-lived due to conservation of some undiscovered quantum number. An analogy to compare this to would be the long lifetime of the kaon due to the conservation of strangeness. The new particles, dubbed stable massive particles (SMPs), would behave like heavy muons if generated in a collider [3].

There are many tests which can be used to search for SMPs, but all are subject to troubles which make detection of beyond the Standard Model particles challenging. The first of these techniques includes searches for anomalously large $\frac{dE}{dx}$ in charged particle tracks. If the detector is not designed for a large detection range, SMPs can saturate electronics. Tracks with saturations can give indications that a massive charged particle is present, but reveal little information about

the particle itself. These saturated hits could also be confused with overlapping trajectories. Furthermore, measurements of $\frac{dE}{dx}$ are not normally distributed, making it difficult to determine the average ionization by the charged particle [2].

Another technique that is employed is to look for excesses of particles which have a β small enough to be stopped by the detector. Assuming SMPs are long-lived would mean that their lifetimes could easily be greater than the time window in which the search is done, so no decay products would be seen from a stopped particle. This measurement is limited in its determination of the mass; at low momentum, the uncertainty in momentum becomes too great. This technique can only tell that there is a SMP present and does not give information about the particle.

The technique most relevant to this project is that of time-of-flight analysis. In this technique, the arrival times of particles are compared to the speed of light, and the difference in time is used to determine the velocity. A separate measurement of momentum allows one to calculate the mass of the particle. This technique has the difficulty of being hindered by the Lorentz factor; as the momentum gets larger, an increase in the momentum gives a smaller contribution to the velocity. After some cutoff momentum, a finite resolution detector will not be able to distinguish between the velocity of two particles.

The equation for relativistic momentum is

$$p = \gamma\beta m \tag{1}$$

By plugging in for γ and solving for $\frac{1}{p^2}$ in terms of $\frac{1}{\beta^2}$, we see the linear relationship

$$\frac{1}{p^2} = \frac{1}{m^2\beta^2} - \frac{1}{m^2} \tag{2}$$

evaluation of the percent error associated with mass measurement is given in [1] as

$$\left(\frac{\Delta m}{m}\right)^2 = \left(\frac{\Delta p}{p}\right)^2 + \left(\gamma^2\frac{\Delta\beta}{\beta}\right)^2 \tag{3}$$

Neglecting the uncertainty in momentum and substituting $\frac{\Delta t}{t}$ for $\frac{\Delta\beta}{\beta}$, one can calculate the uncertainty in the mass. Using the quoted time resolution for the ATLAS experiment of 1 nanosecond, a travel distance of 3 meters, and a β of 0.8 yields a 23% uncertainty in the mass, a number comparable to the uncertainty that could be obtained via $\frac{dE}{dx}$ measurements [1].

With the development of faster detectors, masses can be distinguished up to velocities closer to the speed of light. The large area picosecond photodetector (LAPPD) collaboration has developed technology capable of resolving times on the order of ten picoseconds, with capabilities of reaching sub-picosecond resolution within reach. This opens the door to considerably better time-of-flight resolution, achieving the same 23% uncertainty in mass at $\beta = 0.9998$. With this technology, time-of-flight measurements can be reasonably performed on particles of mass $100\text{GeV}/c^2$ or greater for any kinematically allowed momentum capable of being produced at the LHC.

The goal of this project is to convince the reader that hardware capable of picosecond resolution is a reality and to demonstrate that hardware in development is within striking range of this goal. To do so, we subjected the hardware to UV laser pulses at Argonne National Laboratory. Seven runs of 2000 events on two channels worth of data were taken, with the laser pulsing position spaced 3mm apart in successive runs. The data was collected on an 8-inch MCP plate and read using a waveform-sampling ASIC (application-specific integrated circuit) called the PSEC4. Data

was output through USB to a computer and analyzed using Matlab scripts I wrote. From the data, times and the associated uncertainties were collected, as well as the reduced χ^2 goodness-of-fit. A select few data points were checked by hand and different optimization algorithms were considered for finding the most appropriate extraction of timing information from the pulse rising edge.

Section 2 of this paper describes both the data acquisition process and the data flow through the hardware chain in the output process. A brief overview of some of the physics involved with the detection process is given. In section 3, a description of the hardware is given. This includes both the laser test bench at Argonne and the LAPPD collaboration hardware. Section 4 discusses the techniques involved with extracting timing information via a process named constant fraction discrimination (CFD) and the associated error analysis. This section also shows the data obtained and interprets the results. Finally, section 5 gives a summary of the results and mentions the capabilities possible from hardware developed by the LAPPD collaboration in the near future.

2 Data Taking

The hardware capable of picosecond resolution makes use of anode striplines to transport information. The striplines carry information about photons incident on the board to analog to digital converting chips, which relay the information to a computer through a series of electronics.

The photodetector operates much like a photomultiplier tube. A photon is absorbed in a photocathode, causing emission of an electron. This electron enters a micro-channel plate, or MCP, stack, which is a resistive substrate with micron-sized pores. There is a high potential difference across

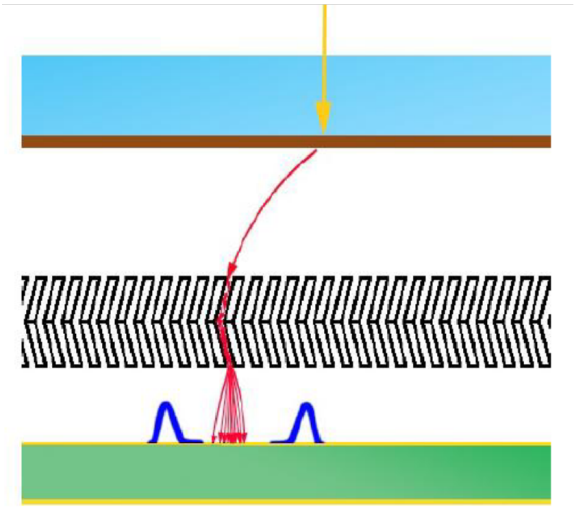

Figure 1: A diagram of the process involved in detecting a photon. The photon (yellow) hits the photocathode, emitting an electron through photoelectric processes. The electron enters the MCP stack, where a high voltage is applied and successive collisions with the walls cause secondary emission. The electrons exit the MCP stack and are collected on the anode striplines, where they are transported as an electrical signal to the readout.

the stack which accelerates the electron enough to induce secondary emission from collisions with the MCP stack. A series of secondary emissions from the electrons create an avalanche of electrons which are collected on anode striplines. To promote these collisions, the MCP pore channels are slightly angled compared to the direction of the applied field, forcing the electrons to hit the MCP instead of traveling straight through.

The use of MCPs as opposed to conventional PMTs increases the spatial, and hence the temporal, resolution significantly. Spatial resolution is governed by the characteristic size of the detector; one can only know that a particle hit the detector, not the location that the detector was hit. For a PMT, the size of the scintillator is about an inch, limiting the time resolution to within 100 picoseconds. An MCP decreases the characteristic size to the order of a few microns, the size of a pore in the MCP stack. This sets the achievable resolution at 1 picosecond, a two order of magnitude increase in resolution.

To test the MCP capabilities, the photodetector is pulsed with UV light from a laser. With the help of Matt Wetstein and Andrey Elagin, we were able to obtain pulser data from the laser testbench at Argonne National Lab. To do so requires realigning the optics if they need to be realigned. This involves tracing the UV light through the optics using a detector card and ensuring that each part of the optical path is seeing the correct light wavelengths make it through the optics unhindered.

Next, we ensure that the laser intensity is high enough to make the journey through the optics, but low enough such that only one photon is hitting the MCP tile per pulse. We do this by hooking up the output of the channels to an oscilloscope and looking at the waveform. When too much light is incident on the MCP, the pulses are large in amplitude and not indicative of a single-photon detector, which is what we would like to test. To little light will reduce the amount of pulses seen in a data set, which is wasteful of data storage.

Following the intensity testing, we connect the waveform digitizing chip and perform a pedestal, which is handled by a script on the computer. The testbench is then ready to collect data. We took several runs of data at different positions on the chip by controlling a translational stage from the computer.

To trigger the electronics, part of the laser light is sent to a photodiode. If the trigger channel is saved, the time of arrival of the pulse on the detector can be inferred and compared to the times that the detector reads a pulse signal. The point of incidence of the laser on the photocathode is set by a mirror on a translational stage controlled by a computer.

The MCP data collection implements a hardware trigger. When a laser pulse is sent through the optics, a portion of the light which would otherwise be unused is extracted and collected in a photodiode. The pulse from this photodiode is added to information about the pulse intensity from another diode and read in through an evaluation board. The evaluation board houses the digitizing electronics and is read out through USB to a computer.

The MCP pulse data is digitized by an ASIC (application specific integrated circuit) capable of high-speed data collection. The ASIC, named the PSEC4, is capable of a 17GHz sampling rate [4]. The PSEC4 has a 256 buffer sampling depth at 12 bit precision and has 1.5GHz bandwidth. It can digitize up to 6 channels simultaneously. A set of digitized pulses from the PSEC4 can be seen in figure 2.

Data collected included 2000 events over two channels for each of seven positions, spaced 3mm apart.

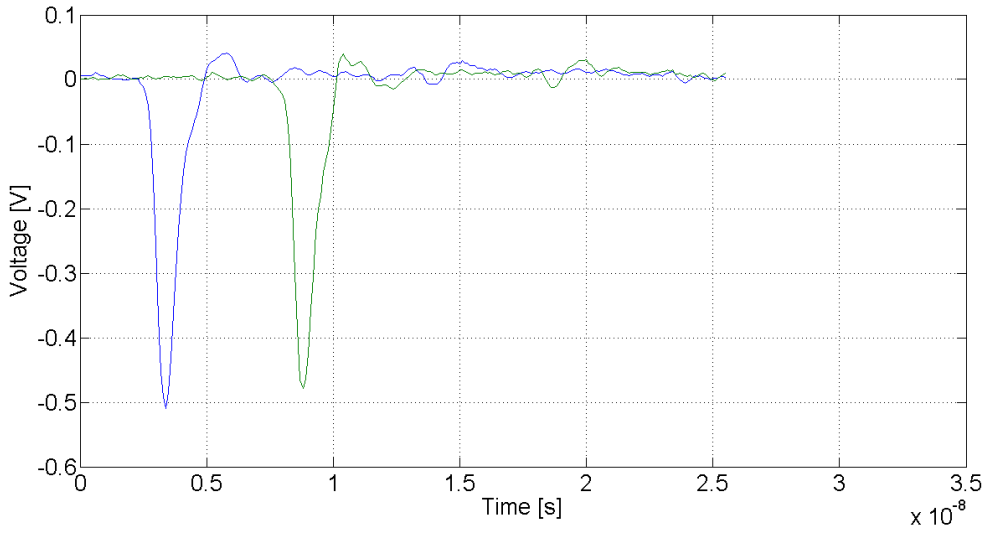

Figure 2: A sample of the readout coming from both sides of a single anode stripline. The x-axis is time in seconds (note the factor of 10^{-8}), and the y-axis is voltage in Volts.

---

3 Hardware Description

The 8"x8" detector panel tiles consist of a photosensitive material laid over two MCPs, aligned to form a chevron MCP stack. A high voltage is applied to the top of the plate to induce secondary emission inside the MCPs. Under the MCP stack is a series of anode striplines which conduct the secondary emission electrons to the readout hardware. Each tile has a set of 30 anode striplines. The tiles can be arranged into a supermodule, increasing the detector area with only a moderate increase in hardware and a minimal loss of resolution.

The anode striplines are directed into the front-end readout electronics, housed on the analog card. The PSEC4 chip reads the signal from the anode striplines and digitizes it, creating a waveform. A single PSEC4 chip can handle 6 anode striplines, with five PSEC4 chips per analog card. For this project, the output digitized waveforms are sent to an evaluation board, which does a conversion from the raw readout to a voltage and buffers the data for output through USB.

Instead of being sent to the computer through the evaluation board, the data can optionally be sent through a series of FPGAs (field programmable gate arrays) which perform analyses on the waveforms. The use of FPGAs has two purposes: data sparsification and reduced analysis time. By running the analysis as the data is sent through the output chain, unnecessary information can be stripped out, leaving only the relevant information behind and thus reducing both the bandwidth and the memory needed to store and transport the information. Furthermore, FPGAs are efficient in that they can process information in parallel, allowing them to perform a series of simple processes considerably faster than a computer.

Two tiers of FPGAs would be used in the front-end data processing. The first is contained on what is called the digital card. This FPGA would perform analyses on only one side of the anode

stripline readouts in what is called first localization. The first localization is the process of collecting information on the integrated charge, the transverse position (to zeroth order, which stripline the pulse hits) and the relative time that the signal pulse reaches the detector. The output data is sent to the central card through 8 LVDS lines. This project is designed to mimic the timing analysis performed by this FPGA.

The second tier FPGA is housed on what is called the central card. Second localization takes place here, where information from the digital card FPGAs on both sides of the anode striplines are combined and processed as a group. The second localization involves converting the relative time information into absolute time information (the time the pulse hits the photodetector plate) as well

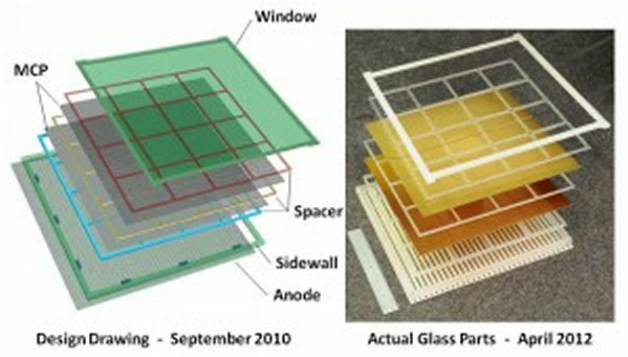

(a) A picture of a single detector tile and an accompanying diagram. The tile has spacers which act to apply the voltage to the MCP stack. The tile is covered by a glass window and the anode striplines lie under the bottom MCP. Photo credit to Rich Northrop.

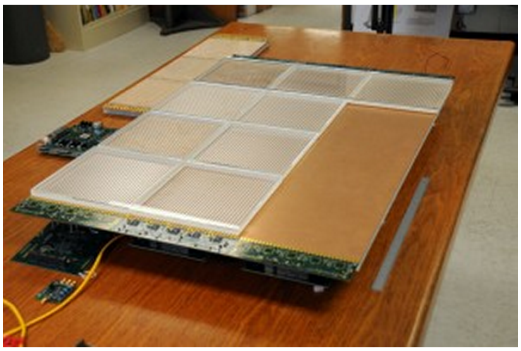

(b) A picture of some detector tiles arranged into a supermodule. In this case, the supermodule will be comprised of a 3 by 4 arrangement of tiles, with 15 PSEC4 chips on each side of the supermodule. A ruler is shown to the right of the supermodule for scale. Photo credit Rich Northrop

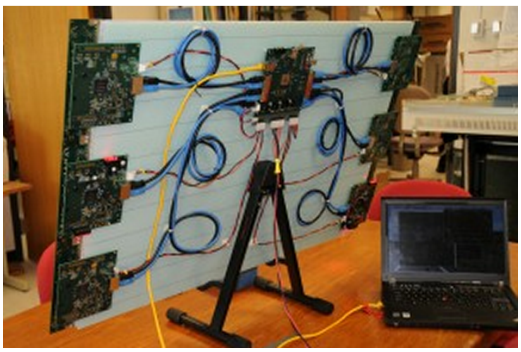

(c) A picture of the supermodule backside. The supermodule has six digital cards, arranged on the sides of the detector, and a single central card in the middle. Photo credit Rich Northrop.

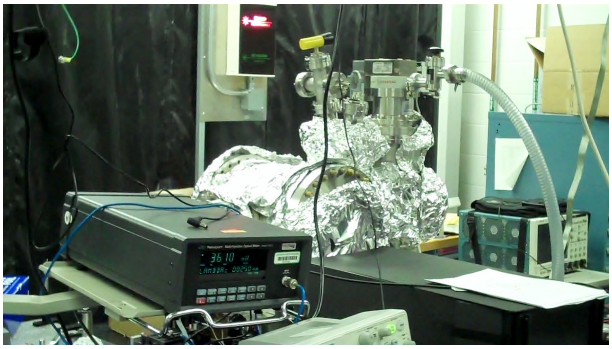

(d) A picture of the vacuum chamber housing the 8 inch MCP tile and part of the laser test bench. The optics for the laser are contained within the black box in the lower right corner of the picture. A power meter can be seen in the foreground on the left. Not pictured is the laser, covering most of the bench beneath the electronics.

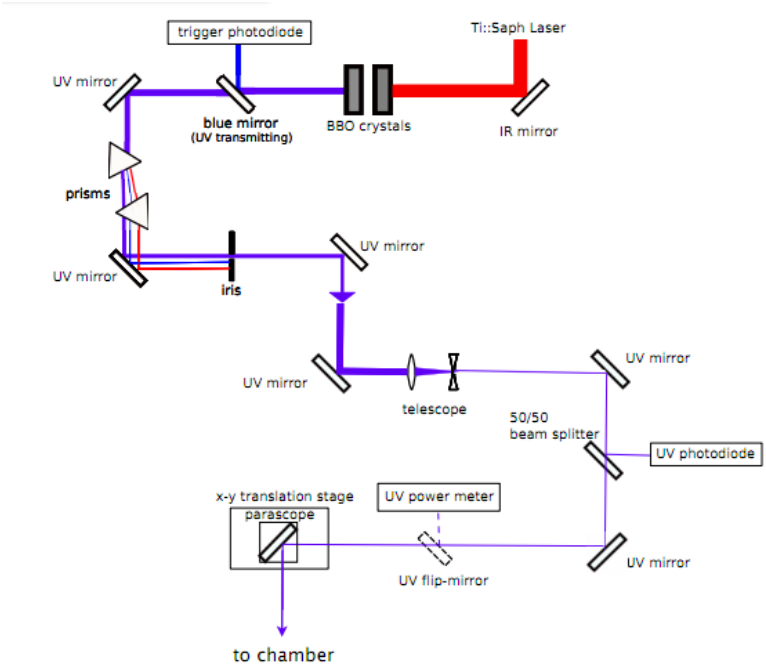

Figure 3: Diagram of the laser testbench setup. The laser test bench has a series of optics to isolate a beam of UV light for pulsing the MCP. The titanium-sapphire laser light is sent through two frequency doubling BBO crystals to generate UV light. The UV beam passes through a blue mirror, and the blue light is used to trigger the readout electronics. The light is further refined with a set of prisms and an iris, then collimated using a telescope. A 50/50 splitter sends part of the beam to a photodiode for a power measurement, and the rest is sent to the MCP in the vacuum chamber. A photodiode is used to trigger the electronics such that they capture the pulses.

---

as finding the position along the stripline based on the arrival time of the pulses in the readouts on either side. The data would be collected and sent to the computer from the central card either through USB or through ethernet.

For testing, the laser testbench at Argonne National Laboratory was used. A titanium-sapphire laser beam is frequency doubled and sent through a series of optics to isolate a beam of UV light. A portion of the unused beam is read with a photodiode and used as a trigger for the electronics. The beam is then focused with a telescope and measured for intensity as it enters the vacuum chamber which houses the detector tile.

4 Analysis

Several factors play into what determines the best resolution possible for a given detector system. As stated in [5], the time resolution depends on the signal to noise ratio, analog bandwidth, sampling

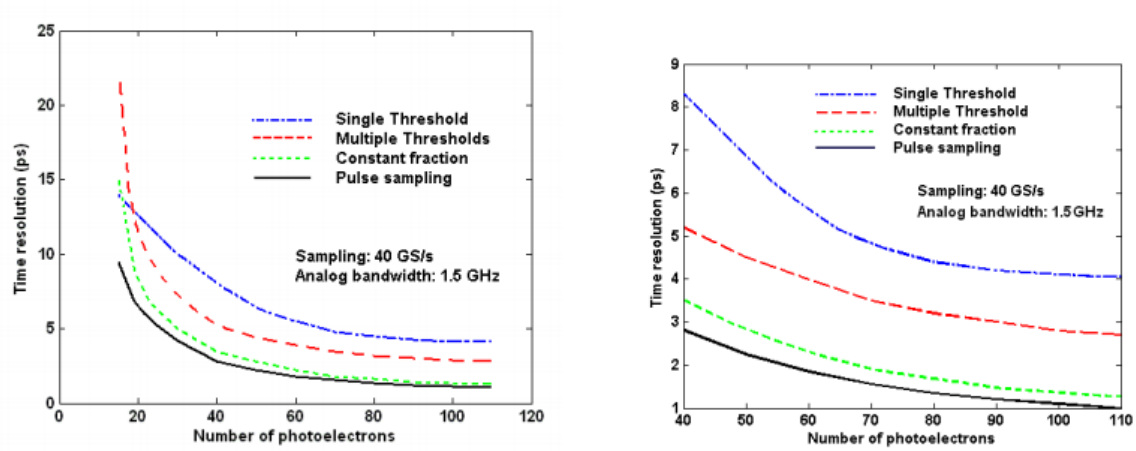

Figure 4: Simulation results of the number of photoelectrons versus time resolution for various fitting algorithms. The algorithms used are single threshold, multiple threshold, constant fraction discrimination, and a spline fit over the waveform. Single threshold takes only the time that the threshold is crossed. Multiple threshold does a fit over the times when the voltage crosses several threshold values and extrapolates a time from the fit. Only the constant fraction discrimination is used in this project. [7]

---

frequency, and signal statistics. The signal to noise ratio, analog bandwidth, and sampling frequency are controlled by the properties and quality of the detector. The signal statistics can be pushed slightly depending on how many data are available to extract information from. We can achieve high resolution by finding a good fit for the shape of our waveform and determining the relevant properties from our fit.

4.1 Constant Fraction Discrimination

More than one algorithm exists to calculate the time of the pulse. These algorithms make use of the long rise time of the rising edge of the pulse. A linear fit over the rising edge is sufficient to extract the timing information necessary for analysis. However, the exact approach to use in creating this fit is ambiguous. Simulations show that not all algorithms are created equal, as in figure 4.

One algorithm that is shown to be robust is the constant fraction discrimination (CFD) method. This algorithm compares the waveform to the maximum amplitude of the pulse and finds the time when the pulse rising edge reaches a constant fraction of the maximum height. Unlike threshold-crossing timing, constant fraction discrimination has the advantage of being mostly independent of pulse amplitude and threshold voltage. Pulses with shapes that are scaled only in amplitude will result in the same times when CFD is implemented.

Applying CFD to a real pulse waveform requires interpolation between data points. This is accomplished by a least-squares linear fit over points in the rising edge of the pulse. The linear

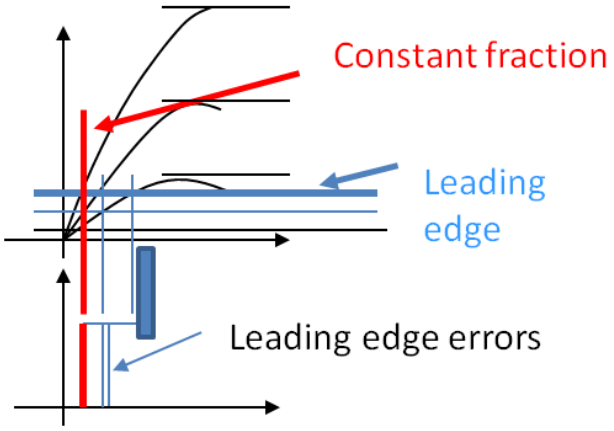

Figure 5: A demonstration of the advantage of the constant fraction discrimination method versus leading edge threshold timing. The time and error in leading edge discrimination varies with pulse amplitude, but not for CFD. [8]

---

least-squares equations are given as follows:

$$m = \frac{N\sum_i t_i - \sum_i t_i \sum_j V_j}{N\sum_i t_i^2 - \left(\sum_i t_i\right)^2} \tag{4}$$

$$b = \frac{\sum_i V_i \sum_j t_j^2 - \sum_i t_i \sum_j t_j V_j}{N\sum_i t_i^2 - \left(\sum_i t_i\right)^2} \tag{5}$$

and the associated uncertainties are

$$(\Delta m)^2 = \frac{\sigma_V^2}{\sigma_t^2 N} \tag{6}$$

$$(\Delta b)^2 = \frac{\sigma_V^2}{N} \tag{7}$$

Where m is the slope and b is the intercept [6].

From here, we can invert the linear equation to find the required time for our analysis:

$$t_{CFD} = \frac{V - b}{m} \tag{8}$$

The uncertainty in time is calculated using standard error propagation, which reduces to:

$$(\Delta t_{CFD})^2 = \frac{1}{m^2}\left((\Delta V)^2 + (m\Delta t)^2 + (\Delta b)^2 + (t_{CFD}\Delta m)^2\right) \tag{9}$$

where the normal error in voltage $(\Delta V)^2$ is replaced by $(\Delta V)^2 + (m\Delta t)^2$ to account for the uncertainty in time of the data. The values used were $\sigma_t = \Delta t = 10\text{ps}$ and $\sigma_V = \Delta V = 0.0072V$.

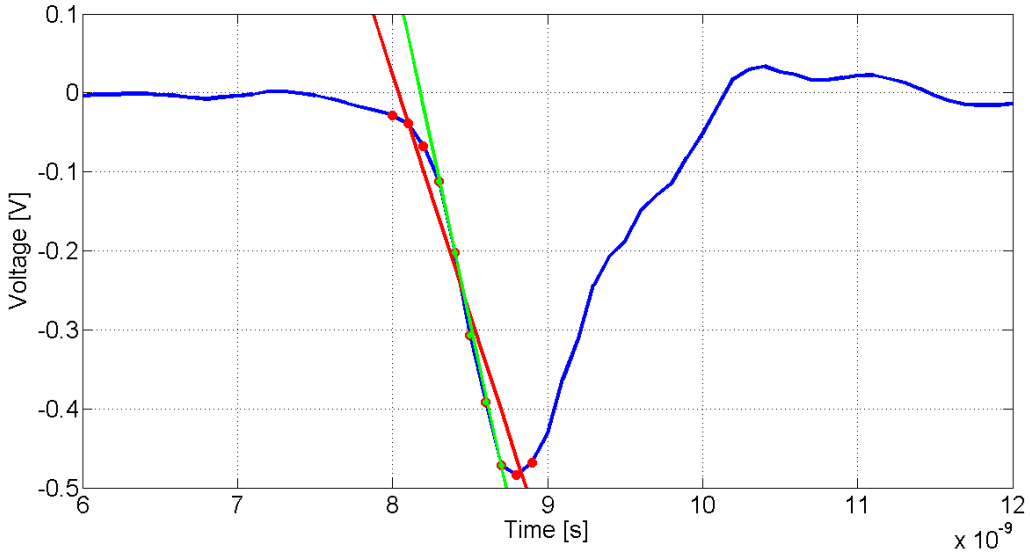

Figure 6: A sample linear fit over a pulse. The shape of the pulse is outlined in blue. Two different optimizations were used, as is described in Section 4.2. A fit over the weighting function (N=5) described in Equation 11 is shown in green, while the optimization over uncertainty only (N=10) is shown in red. For both of these optimizations, the data points on the pulse that were used in each of the linear fits are also shown. Although the optimization in uncertainty claims to have a better resolution than the weighting function (21ps versus 31ps), edge effects of the pulse in the linear fit have a considerable impact on the goodness of fit. The reduced χ^2 for the weighting and uncertainty optimizations are 1.12 and 41.5 respectively. Note that all of the data points used in the weighting optimization are also used in the uncertainty optimization.

To get the goodness of fit, the reduced χ^2 value was calculated as well. This is simply χ^2 divided by the number of degrees of freedom, or

$$\chi^2_{red} = \frac{\chi^2}{\nu} = \frac{1}{N-2}\sum_i \frac{V(t_i) - V_i}{\sigma_V} \tag{10}$$

where N is the number of data points used in the least-squares fit. A sample linear fit can be seen in figure 6.

4.2 Assumptions and Approach

In analyzing the pulse waveforms, some simplifying assumptions were made:

- No relative timing information was available. This includes the trigger and the wrap-around point of the data. All data was assumed to be contiguous and in chronological order. Without this data, the absolute arrival time of the pulse cannot be measured, only the relative time. For the extent of this project, the absolute time is unnecessary.

- Pedestal values were not calculated. The pedestal was assumed to be consistent over the entire channel and to have a value equal to the median of all of the buffer values on the channel for that event.

- Uncertainty in voltage was unknown for these runs. However, recent work has shown that the uncertainty is around 0.1% of the peak maximum. The uncertainty for this analysis was found by taking the RMS voltage over an unused output channel for all of the events. This uncertainty was found to be 7.2mV, or approximately 2% of the peak maximum. The increase in the uncertainty can be explained by electromagnetic interference from correlation between channels. This uncertainty is still used in the analysis as a "worst case" scenario.

- The time between consecutive buffers is assumed to be 100 picoseconds ±10%, consistent with the hardware sampling rate for the data runs. Again, this is a "worst case" assumption; the hardware is capable of reaching times of 60ps ±10%.

- The original data was found to have severe problems with noise, due to electronics interfering with each other and possibly a ground loop issue. This data was thrown away in favor of data taken by Eric Oberla, a month prior. Analysis was done on this older data, which was obtained using the same methods mentioned in section 2.

The times and uncertainty in times were calculated for all events on all channels using Matlab scripts implementing the constant fraction discrimination method. To ensure the best fit possible in the linear least squares, combinations of consecutive data points for each waveform were tried and tested for both their uncertainty and reduced χ^2. The number of data points used in the linear fit was allowed to range from 4 to 10. Cuts were made to ensure that the slope of the calculated linear fit was negative and that the median of the data points used in the fit were above threshold, thereby ensuring that the linear fit was run over the rising edge of the pulse.

The resulting uncertainties and χ^2_{red} were tested using an arbitary weighting function for the "optimal" fit. The weighting function was invented by hand-picking the "best" choices from a select set of events. For most of the analyses, the weighting function

$$W = -sign(m)\frac{N-3}{\left(\chi^2_{red} + \frac{1}{\chi^2_{red}}\right)^2} \tag{11}$$

was used, where m is the slope of the fit and N is the number of points used in the fit. This weighting function favors those with $\chi^2_{red} \approx 1$ and larger N. For completeness, this optimization algorithm is compared to an algorithm which favors uncertainty in time only.

A few cuts needed to be made on the data to ensure reasonable results. The first cut was a standard threshold cut: if peaks from both channels after pedestal subtraction did not have an amplitude of at least 50mV, the event was thrown out. Another cut was required to remove events which did not trigger, in which case the voltage was at a constant -1.75V. A cut requiring the pedestal to have an amplitude below 1V was sufficient to remove these outliers.

4.3 Results

A histogram of the calculated optimum uncertainties in time can be seen in figure 7. This result uses the weighting function described in equation 11. This shows that the calculated error in uncertainty peaks somewhere around 55-60 picoseconds.

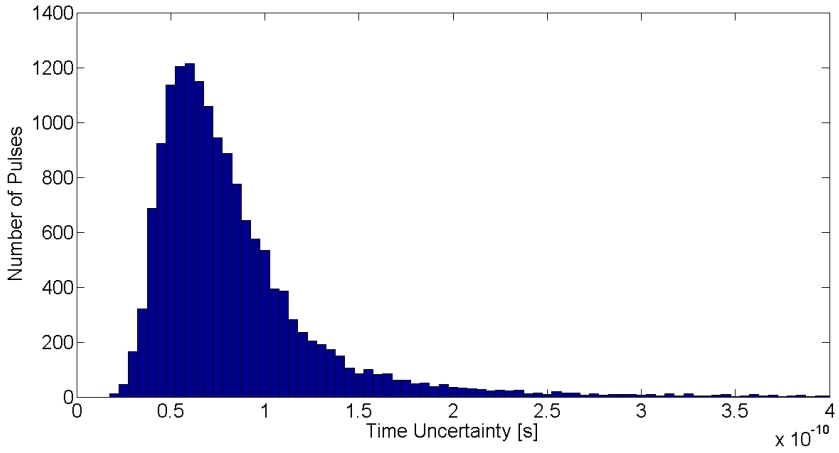

Figure 7: Histogram of the observed uncertainty in time. This includes all runs over all events and both channels, totaling 15972 pulses after pedestal and threshold cuts. No distinctions between runs were made because position in the lateral direction along an anode stripline would not be available for the front-end electronics. Bins are in increments of 5 picoseconds with the lowest filled bin centered around 20ps.

One might ask how these uncertainties compare to what would be obtained if, rather than using this arbitrary weighting function, the choice with minimum uncertainty is used. As is expected, the calculated uncertainties in time decrease. With the weighting function, the average uncertainty (with a cut requiring the uncertainty to be less than 300 picoseconds) was calculated to be 74 picoseconds and the best uncertainty achieved was 21 picoseconds. With an uncertainty optimization, the average uncertainty reaches 60 picoseconds with a best uncertainty of 12 picoseconds.

When optimizing over uncertainty, the data favor more points in the linear fit. With the weighting function, the average number of points used in linear fits was 6.56; the average with uncertainty optimization is 8.92. This increase in data points allows the linear fit to cover the entire rising edge of the pulse. However, the increase in coverage also adds the fringe effects of the beginning and peak to the linear fit, as is seen in Figure 6. This increases the reduced χ^2, as is seen in figure 8. When the number of points used in the linear fit is held fixed, optimization over uncertainty and optimization over the weighting function (at fixed number of points, the weighting function becomes optimization over χ^2_{red}) become synonymous.

5 Conclusions

Data was collected by firing a beam of UV light at an MCP detector tile. Seven positions relative to an origin on the stripline and both sides of a single stripline were used in the data collection. The data was read out through USB from an evaluation card and saved to a computer. A matlab analysis was run over the events to find the time and uncertainty in time for each pulse.

When the fits are optimized over uncertainty, linear fits generated from more data points are

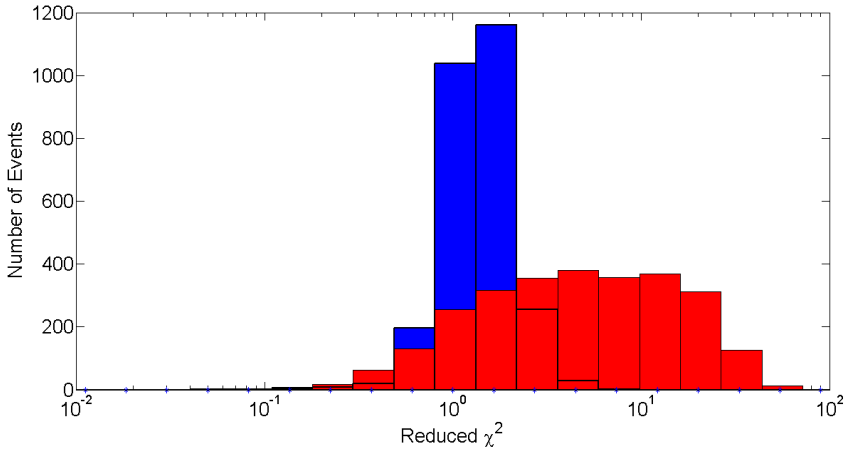

Figure 8: Histogram of the reduced χ^2 for both the weighting function optimization (blue) and the uncertainty optimization (red) over a single data run. Optimization over uncertainty leads to an observable shift in the average χ^2_{red}, which is due to the fringe effects of the linear fit.

---

preferred. These fits use the entire rising edge to calculate the time of the pulse, making the linear fit less susceptible to fluctuations due to noise and decreasing the uncertainty in time. Allowing more points in the fit is susceptible to fringe effects from the start and peak of the pulse, making the linear fit a less-preferred method for describing the rising edge.

The mean uncertainty in time over a single run with the arbitrary weighting function given in equation 11 was calculated to be 74 picoseconds after cutting uncertainties greater than 300 picoseconds. The best uncertainty 21 picoseconds. This is a conservative estimate though, as these uncertainties were calculated using "worst-case" scenario values, and on hardware that is not running at full speed nor optimized for the highest time resolution possible.

With a more realistic uncertainty, a significant improvement in time can be expected. Using the quoted value of 0.1% of the average peak height as the voltage uncertainty, the mean uncertainty drops to 17 picoseconds with a best uncertainty of 10 picoseconds. This is still under the restriction that the hardware is not running at top speed. With optimized hardware and a faster speed, this estimate can be improved even more.

Acknowledgements

I would like to thank Henry Frisch for guiding me through this project in its entirety and making it possible in the first place, Eric Oberla for all his advice and aid in getting the hardware working, Matt Wetstein and Andrey Elagin for operating the laser and showing me the ropes in data collection, Craig Harabedian and Mircea Bogdan for their help with learning to program hardware, Mary Heintz for making sure the computers are available and in working condition for my use, and Max Hutchinson for his help in coding and moral support through many a long night.

References

[1] M. Fairbairn, A. Kraan, D. Milstead, T. Sjostrand, P. Z. Skands *et al.*, "Stable massive particles at colliders," *Phys.Rept.*, vol. 438, pp. 1–63, 2007.

[2] S. Chatrchyan *et al.*, "Search for heavy long-lived charged particles in pp collisions at sqrt(s)=7 TeV," 2012, submitted to Physics Letters B.

[3] S. Tarem and S. Bressler, "Search for a heavy stable charged particle with the atlas detector," 2007. [Online]. Available: vipac.desy.de/Common/miniworkshopDM/tarem-StableStau.ppt

[4] E. Oberla, "Electronics for a position & time sensing large area photo- detector system," 2011.

[5] H. Frisch, "Overview and a little history of lappd," 2011.

[6] P. H. Richter, "Estimating Errors in Least-Squares Fitting," 1995, tDA Progress Report 42-122.

[7] J.-f. Genat, G. Varner, F. Tang, and H. Frisch, "Signal processing for pico-second resolution timing measurements," *psec.uchicago.edu*, vol. 607, no. 2, pp. 387–393, 2009. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S016890020901167Xhttp://psec.uchicago.edu/Papers/NIM_v8b_final.pdf

[8] M. Bogdan *et al.*, "A 20 gs/s sampling asic in 130nm cmos technology," 2010.