

Improving the Time and Energy Resolution in the PIENU Calorimeter System

by

Kevin Yin

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

BACHELOR OF SCIENCE

in

The Faculty of Science

(Physics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

April 2015

© Kevin Yin 2015

Abstract

The pion decay branching ratio is an important test of the Standard Model, but currently the experimental error in its measurement lags behind the theoretical error by an order of magnitude. Utilizing data collected by the PIENU project at TRIUMF, the effects of pulse shape fitting on the calorimeter waveforms will be examined for their merit in improving the time and energy resolution of the calorimeter system. This in turn will increase the precision of the branching ratio. The results will provide evidence for or place a stringent constraint on non-Standard Model physics related to the assumption of lepton universality.

Table of Contents

Abstract	ii
Table of Contents	iii
List of Tables	v
List of Figures	vi
Acknowledgments	vii
1 Introduction	1
1.1 Motivation	1
1.2 PIENU Experiment	2
1.3 Outline of this Project	2
2 Theoretical Background	6
2.1 The Branching Ratio	6
2.2 Energy of the Decay Products	6
2.3 Expected χ^2	7
3 Waveform Fitting	9
3.1 Overview	9
3.2 Candidate Functions	10
3.3 Averaged Waveforms	10
3.4 Comparing the Fits	12
3.5 Collation	12
3.5.1 Problems with the Averaged Waveforms	12
3.5.2 Process	13
3.5.3 Smoothness Test	14
3.5.4 Spline Fit	14
3.5.5 Iterative Improvements	15
3.6 Time Resolution	16
4 Pileup Detection	18
4.1 Overview	18
4.2 Range Reduction	18
4.3 Pileup Cut	19

Table of Contents

5	Results	21
5.1	Energy Resolution	21
5.2	Time Resolution	21
5.3	Pileup Elimination	21
6	Conclusions	22
6.1	Results	22
6.2	Plausible Improvements	22
	Bibliography	23

List of Tables

5.1	List of α , β , and γ values for fits used in this experiment	21
-----	---	----

List of Figures

1.1	The PIENU apparatus consists of many plastic scintillators and detectors residing in the front, and a large calorimeter in the back.	3
1.2	The PIENU calorimeter is comprised by a large NaI(Tl) crystal surrounded by 97 CsI crystals arranged in two rings.	4
1.3	Energy spectrum of the decay products. The muon decay events are characterized by the large amount low energy events to the left of the red line while the positron decay events are characterized by the small high energy peak to the right	5
1.4	Time spectra of the low energy (left) and high energy (right) decays	5
3.1	A typical waveform produced by the NaI digitizer. The time samples are separated by 33.33 ns giving a total time window of 1333 ns for the entire waveform.	9
3.2	Plot of the 10 averaged waveforms at a fixed time, normalized such that the peak is at 1. The discrepancy between waveforms of different pulse heights is minimal.	11
3.3	Same plot as Figure 3.2, zoomed in at the peak.	11
3.4	χ^2 distribution plotted against pulse height for the Inverse Gaussian function. Note the distinct quadratic dependence of the χ^2 on the pulse height.	12
3.5	Plot of the normalized waveforms referred to earlier with linear interpolation. While the waveforms agree, their interpolations have large discrepancies, specifically at points of high second derivative. This leads to significant differences in the integral.	13
3.6	The collated waveform (in black) along with collated first and second derivatives (in red and green respectively). Notice how the collated derivatives heavily exaggerate any imperfections in the collation	14
3.7	The spline fit (red) along with its control points (blue) and the collated waveform (black)	15
3.8	The second collated waveform (in black) along with its collated first and second derivatives (in red and green respectively). The noise seen in the waveforms is now subdued in comparison to the first collated waveform.	16
3.9	Plot of the time uncertainty (in number of samples) of the fit relative to the pulse height of the waveform.	17
4.1	A typical two pulse pileup event. Note the failure of the fit to match the shape	18
4.2	Distribution of the χ^2 values for various fitting window sizes. Note the stark increase in χ^2 after the range is reduced within ± 5 of the peak.	19
4.3	Distribution of the χ^2 values for various fitting window sizes relative to the amount of samples between the primary pulse peak and the secondary pulse peak	20

Acknowledgments

I would like to thank my supervisor Professor Douglas Bryman for providing me the opportunity, guidance, and resources to pursue this project.

I would also like to express my gratitude towards Luca Doria and the PIENU team for their patience and counsel throughout my work on the project.

Chapter 1

Introduction

1.1 Motivation

The Standard Model is an extensive theory describing three of the four fundamental interactions (electromagnetic, weak, strong, but not gravitational interactions). It has been extremely successful in predicting the behavior of these interactions, most recently with the tentative confirmation of the Higgs boson, first proposed 50 years ago. Since its inception in the mid 1970s, no clear discrepancies between its predictions and experimental data have been found [1], and as a result the formulation has changed very little.

Although it is currently the most accurate description of nature as we know it, the Standard Model is not complete; it depends on 19 empirically determined parameters and a number of core assumptions [2], and as such has spawned many theories attempting to reconcile the massive amount of free parameters. One core assumption is that of lepton universality, which states that all three of the known charged leptons, the electron, the muon, and the tau particle, behave and interact identically under the electroweak force [3]. Should a violation of lepton universality be observed, the Standard Model will need to be amended with additional parameters describing how the electroweak interaction now differs with each of the charged leptons.

The discovery of pion to electron/positron decay in 1958 [4] spawned one of the most exacting tests of lepton universality, wherein the occurrence of the pion to positron decay ($\pi^+ \rightarrow e^+ + \nu_e$) is compared to the occurrence of the pion to muon decay ($\pi^+ \rightarrow \mu^+ + \nu_\mu$). The pion branching ratio, which describes this proportion, is predicted by the Standard Model to depend only on the masses of the positron and the muon, each of which have been measured to a high degree of accuracy. In 1960, the first precise measurement of the branching ratio yielded a value of $1.21(7) \times 10^{-4}$ [5]. Since then, numerous improvements in measurement techniques have gradually increased the precision towards the theoretical limit [6][7]. Currently, the experimental value of the branching ratio is $1.230(4) \times 10^{-4}$ [8], the error of which is still considerably larger than that of the theoretical value $1.2351(2) \times 10^{-4}$ [9]. With a factor of 20 difference in the error, there is significant room for high energy lepton universality violations not explained by the Standard Model.

For the past eight years, the PIENU experiment at TRIUMF has been seeking to improve the experimental error in measuring the branching ratio by about a factor of 5. This would put the error to within the $\pm 0.1\%$ level and within a factor of 4 of the theoretical error [10]. The experiment depends heavily on the readout of the calorimeter system, which measures the energy of the positron produced in both the pion to positron decay and the pion to muon to positron decay. By improving the time and energy resolution in the calorimeter system, and by identifying and reducing background noise, this project will contribute significantly to the precision of the branching ratio.

If the branching ratio agrees with the theoretical value, it will greatly limit the possibilities of non-Standard Model physics and provide solid support for lepton universality. However if the experiment fails to agree with the theory, it will potentially open up a realm of possible new

physics necessary to explain the discrepancy. In either case the result has noteworthy implications on the predictive power of the Standard Model.

1.2 PIENU Experiment

A pion beam of kinetic energy 18.9 MeV from the TRIUMF M13 beamline is fed into the PIENU apparatus (Figure 1.1). The pions stop in a plastic scintillator target, also known as "B3", and subsequently decay into either a muon or a positron (plus assorted neutrinos, which escape the device). In most cases, the significantly more massive muon travels 1 mm in the target before coming to a rest and decaying into a positron and neutrinos as well. Due to the slow down of the muon in the target, the positrons produced by the two types of decays have distinctly different time and energy characteristics. The positrons then hit the calorimeter at the back of the PIENU apparatus, creating a shower of particles. The many plastic scintillators that the particles pass through along the way before hitting the calorimeters are used to calculate the time spectra of the decay.

The two part calorimeter (Figure 1.2) consists of a large 48 cm wide and 48 cm long NaI(Tl) crystal, which captures the bulk of the positrons, and a ring of 97 CsI crystals, which capture some of the particle showers that were not completely contained in the NaI(Tl) crystal. The light produced in the NaI(Tl) crystal is picked up by photomultiplier tubes read out by numerous 30 MHz analog to digital converters (ADCs) while the CsI crystals are read out by a number of 60 MHz ADCs. This is used as the basis for the energy spectra as the light produced is directly proportional to the energy of the particles entering the calorimeter.

The energy spectra of the decay products recorded by the calorimeters (Figure 1.3) is then used to separate the bulk of the high energy positron decays from the low energy muon decays. The raw branching ratio is calculated by simultaneously fitting the time spectra of these two distributions (Figure 1.4) and taking the appropriate ratio. The final branching ratio is then extracted from this value after application of a number of other corrections.

1.3 Outline of this Project

Since the raw branching ratio is calculated from the time and energy spectra of the decay products, improving their resolution in the calorimeters is bound to improve both the cut separating the high and low energy events as well as the process of fitting the time spectra.

This project accomplishes this feat by fitting the waveforms output by the NaI calorimeter with a well defined analytic function. In this fashion, the time and energy values will be calculated from continuous interpolated data rather than the discrete data given, alleviating the constraints of having limited sample rates and digital resolution. In addition, the resulting fit was used to analyze various properties of pileup, providing methods of marking and eliminating their influence on the energy spectra.

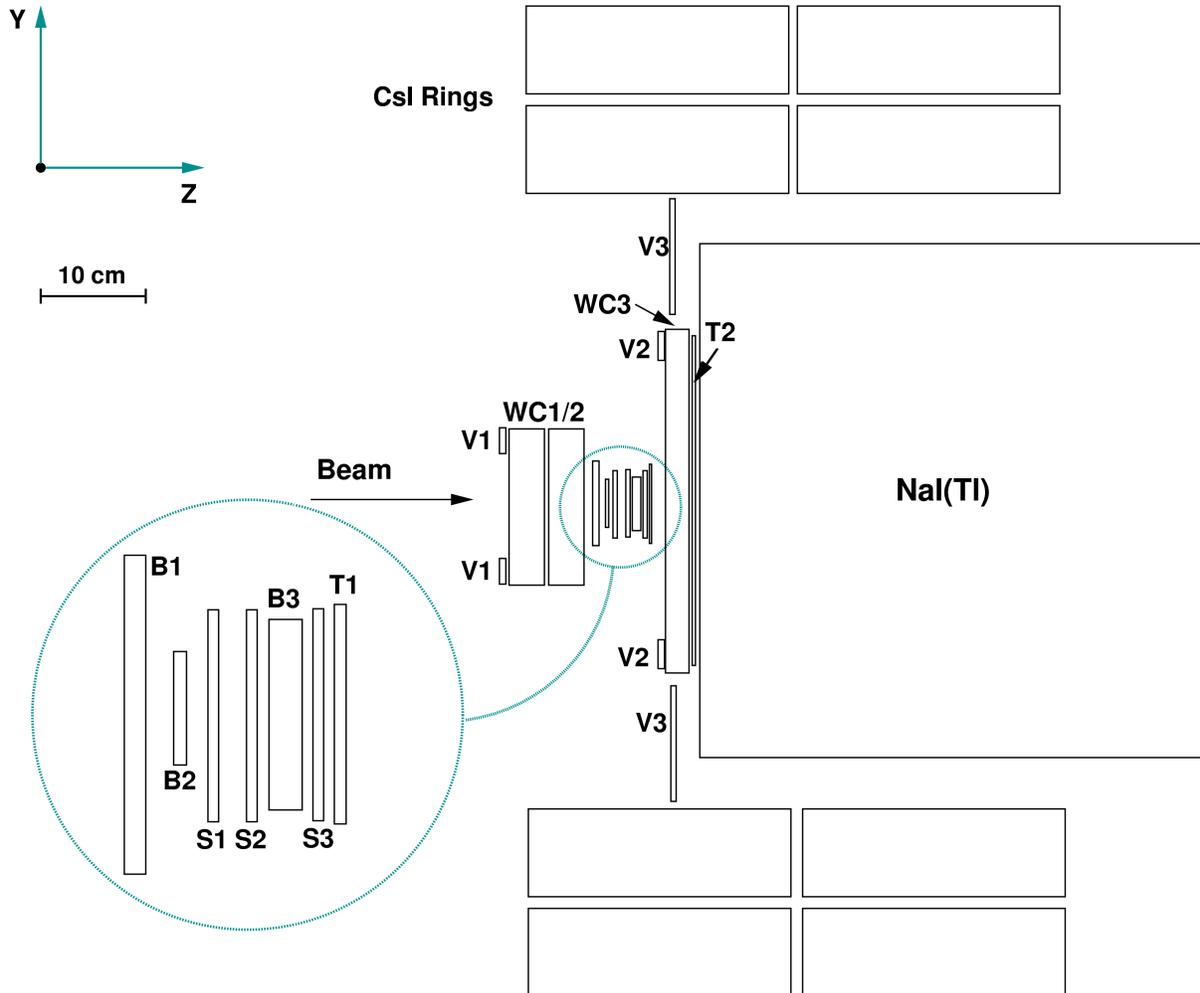


Figure 1.1: The PIENU apparatus consists of many plastic scintillators and detectors residing in the front, and a large calorimeter in the back.



Figure 1.2: The PIENU calorimeter is comprised by a large NaI(Tl) crystal surrounded by 97 CsI crystals arranged in two rings.

1.3. Outline of this Project

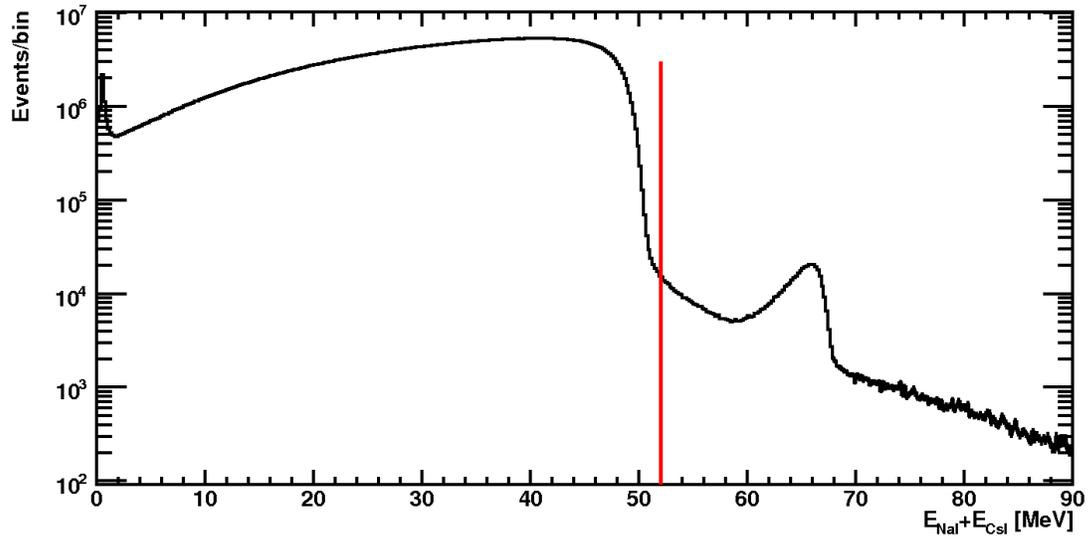


Figure 1.3: Energy spectrum of the decay products. The muon decay events are characterized by the large amount of low energy events to the left of the red line while the positron decay events are characterized by the small high energy peak to the right

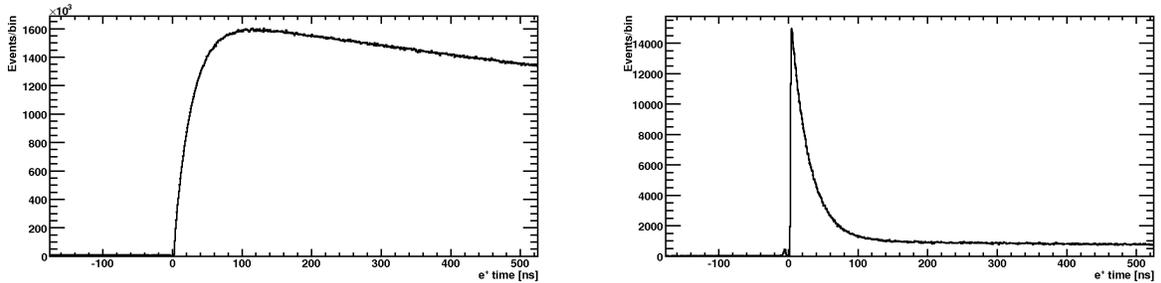


Figure 1.4: Time spectra of the low energy (left) and high energy (right) decays

Chapter 2

Theoretical Background

2.1 The Branching Ratio

To lowest order the rate of decay of a pion to a lepton l is as follows [11]:

$$\Gamma(\pi \rightarrow l + \nu) = \frac{1}{8\pi} G_F^2 f_\pi^2 m_l^2 m_\pi \left(1 - \frac{m_l^2}{m_\pi^2}\right) \quad (2.1)$$

Here G_F is the Fermi coupling constant, and f_π is the theoretically not well known pion decay constant. Assuming lepton universality, f_π will be the same for both muon and positron decays, and so by taking a ratio, the factor of f_π can be eliminated:

$$R_0^\pi = \frac{\Gamma(\pi \rightarrow e + \nu)}{\Gamma(\pi \rightarrow \mu + \nu)} = \frac{m_e^2(m_\pi^2 - m_e^2)}{m_\mu^2(m_\pi^2 - m_\mu^2)} \quad (2.2)$$

Substituting in the current Particle Data Group values for the masses [10], the raw lowest order pion branching ratio evaluates to $1.28336(2) \times 10^{-4}$. With the application of a $-3.761(18)\%$ higher order correction [12] from chiral perturbation theory and radiative corrections, this gives an incredibly precise theoretical prediction of $1.2351(2) \times 10^{-4}$.

2.2 Energy of the Decay Products

A pion rarely exhibits the decay $\pi^+ \rightarrow e^+ + \nu_e$. Conservation laws and the relativistic energy momentum equation produce the following identities [13]:

$$E_\pi = E_e + E_\nu \quad (2.3)$$

$$\vec{p}_\pi = \vec{p}_e + \vec{p}_\nu \quad (2.4)$$

$$E_\nu^2 - \vec{p}_\nu^2 = m_\nu^2 \quad (2.5)$$

Assuming neutrinos to be massless and letting $\vec{p}_\pi = 0$ (a decay at rest) gives the following equations for the positron energy:

$$(E_\pi - E_e)^2 - (\vec{p}_\pi - \vec{p}_e)^2 = 0 \quad (2.6)$$

$$(m_\pi - E_e)^2 = E_e^2 - m_e^2 \quad (2.7)$$

$$E_e = \frac{m_\pi^2 + m_e^2}{2m_\pi} \approx 69.8 \text{ MeV} \quad (2.8)$$

2.3. Expected χ^2

Alternatively the pion may exhibit the common decay $\pi^+ \rightarrow \mu^+ + \nu_\mu$. The muon then decays further into a positron via the decay $\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu$. The same laws as above give:

$$E_\mu = E_e + E_\nu + E_{\bar{\nu}} \quad (2.9)$$

$$\vec{p}_\mu = \vec{p}_e + \vec{p}_\nu + \vec{p}_{\bar{\nu}} \quad (2.10)$$

$$E_\nu^2 - \vec{p}_\nu^2 = m_\nu^2 \quad (2.11)$$

$$E_{\bar{\nu}}^2 - \vec{p}_{\bar{\nu}}^2 = m_{\bar{\nu}}^2 \quad (2.12)$$

Again assuming neutrinos to be massless and letting $\vec{p}_\mu = 0$, the following equations hold:

$$(E_\mu - E_e - E_\nu)^2 - (p_\mu - p_e - p_\nu)^2 = 0 \quad (2.13)$$

$$m_\mu - E_e - |p_\nu| = |-p_e - p_\nu| \quad (2.14)$$

$$(m_\mu - E_e)^2 \geq E_e^2 - m_e^2 \quad (2.15)$$

$$E_e \leq \frac{m_\mu^2 + m_e^2}{2m_\mu} \approx 52.8 \text{ MeV} \quad (2.16)$$

Conclusively, a positron coming from a decay at rest will have characteristically different energies for these two decays. A cut some amount above 52.8 MeV can thus effectively separate most of the muon and positron decays.

2.3 Expected χ^2

It is useful to see what values of χ^2 would be expected when fitting waveforms. Denote a collection of variables as follows:

- $y(t)$: Amplitude of the waveform at time t
- A : Scaling factor
- N : Number of discrete samples in time
- t_i : Time of sample i
- y_i : Amplitude of the waveform at time t_i including scaling and noise
- $f(t)$: Amplitude of the fitting function at time t
- ϵ : Model of the noise independent of A

The noise ϵ is assumed to have $\mathbb{E}[\epsilon(t)] = 0$ and $\text{Var}(\epsilon(t)) = \sigma_y^2$. In practical application, y_i would represent the data given, and $f(t)$ would represent the function used to fit the data.

Using the typical calculation of $\chi^2 = \sum_{i=0}^{N-1} (y_i - Af(t_i))^2$, the mean of the χ^2 value can be expected to behave as follows:

$$\mathbb{E}[\chi^2] \approx \mathbb{E} \left[\int_0^N (Ay(t) + \epsilon(t) - Af(t))^2 dt \right] \quad (2.17)$$

$$= \alpha^2 A^2 + 2\beta A + \gamma^2 \quad (2.18)$$

2.3. Expected χ^2

where

$$\alpha^2 = \mathbb{E} \left[\int_0^N (y(t) - f(t))^2 dt \right] = \int_0^N (y(t) - f(t))^2 dt \quad (2.19)$$

$$\beta = \mathbb{E} \left[\int_0^N (y(t) - f(t))\epsilon(t) dt \right] = \int_0^N \text{Cov}(y(t) - f(t), \epsilon(t)) dt \quad (2.20)$$

$$\gamma^2 = \mathbb{E} \left[\int_0^N \epsilon(t)^2 dt \right] = \int_0^N \mathbb{E} [\epsilon(t)^2] dt = N\sigma_y^2 \quad (2.21)$$

α represents the L^2 norm of $y(t) - f(t)$ and thus can be used as a measure of goodness of fit in that a good fit will necessarily minimize α . β relates to the covariance of the fitting function and the error, and γ represents a multiple of the variance of the error, which practically represents the digitization noise.

With this information in mind, the upper bound of the discrepancy between the integral suggested by the fit ($Q_f = \int_0^N Af(t)dt$) and the actual value of the integral ($Q = \int_0^N Ay(t)dt$) can be calculated.

$$|\mathbb{E}[Q - Q_f]| = \left| \mathbb{E} \left[\int_0^N (Ay(t) - Af(t)) dt \right] \right| \leq \alpha A \quad (2.22)$$

$$\text{Var}(Q - Q_f) \leq \mathbb{E} \left[\left(\int_0^N (Ay(t) - Af(t)) dt \right)^2 \right] \leq \alpha^2 A^2 \quad (2.23)$$

The inherent variance in Q due to noise is

$$N\text{Var}(\epsilon) = N\sigma_y^2 = \gamma^2 \quad (2.24)$$

Therefore $Q = Q_f \pm \sigma_Q$ where $\sigma_Q = \alpha A + \sqrt{\alpha^2 A^2 + \gamma^2}$. If γ is assumed to be small, then σ_Q can be approximated by $2\alpha A$.

Chapter 3

Waveform Fitting

3.1 Overview

The digitizer for the NaI calorimeter produces a waveform with 40 data points, sampled at 30 MHz (33.33 ns per sample). For convenience "samples" will be used as a unit throughout this document denoting a time interval of 33.33 ns. The shape of the waveform is characterized by a rapid rise as the particle shower enters the calorimeter followed by a slowly decaying tail as the excited electrons recombine (Figure 3.1).

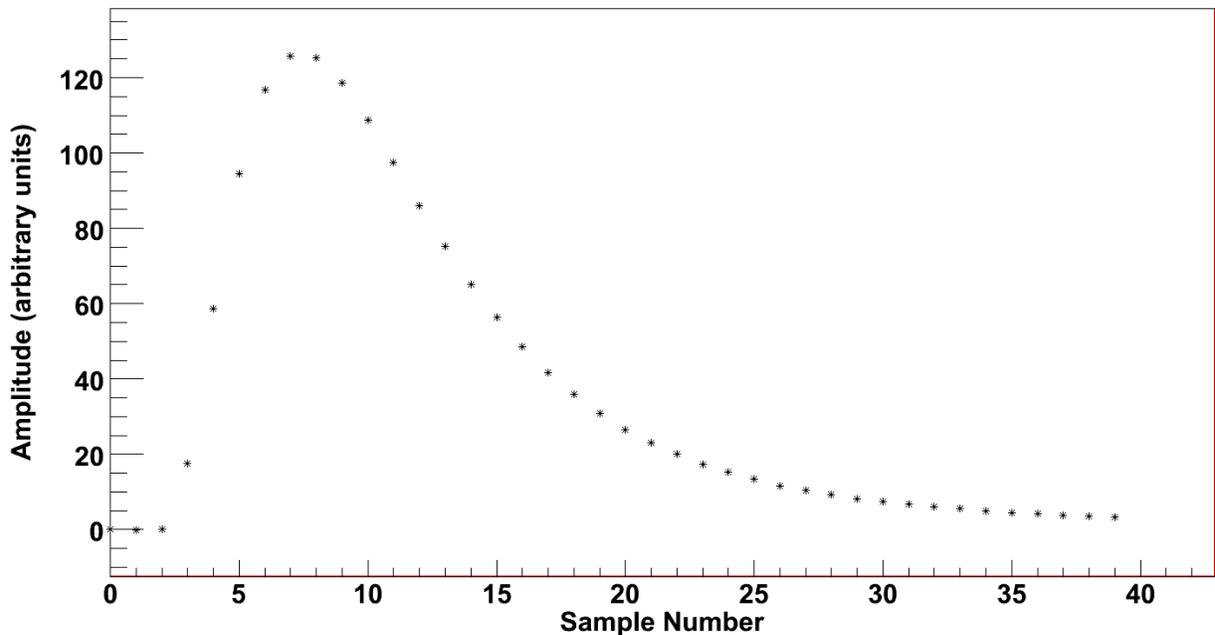


Figure 3.1: A typical waveform produced by the NaI digitizer. The time samples are separated by 33.33 ns giving a total time window of 1333 ns for the entire waveform.

The rationale behind why fitting is necessary is due to the large uncertainties posed by the number of samples. The digitizer itself only provides 1024 possible values for the amplitude. Since data sets with no waveform at all have an average value of about 420 (this is referred to as the "Pedestal"), this leaves effectively 600 possible values to describe the waveform.

An estimate based purely on the discrete data gives one of 600 possible values for the pulse height, and one of only 40 values for the peak time. This is a particular problem in waveforms like the one shown in Figure 3.1, where the two highest points of the waveform are of approximately equal value, hence suggesting a true peak time somewhere in between.

3.2 Candidate Functions

The first task undertaken in this experiment was to find a set of candidate functions that fit the waveform. Each function was normalized so that the peak is at a value of 1 instead of the typical normalization of having the integral at 1. The function was then transformed ($Af(t - T) + P$). This is done so that fitting will immediately produce values for the pulse height (A) and peak timing (T), which are easily verifiable by visual inspection compared to the less tangible integral.

$$\text{CrystalBall}(t; \sigma, n, \alpha) = \begin{cases} e^{-\frac{t^2}{2\sigma^2}} & \text{if } \frac{t}{\sigma} \leq \alpha \\ \left(1 - \frac{\alpha^2}{n} + \frac{\alpha t}{n\sigma}\right)^{-n} e^{-\frac{\alpha^2}{2}} & \text{if } \frac{t}{\sigma} > \alpha \end{cases}$$

$$\text{InvGamma}(t; p, b) = \left(\frac{p}{b}t\right)^{-p} e^{-\frac{b}{t}+p}$$

$$\text{InvGauss}(t; a, b) = \left(\frac{2a}{-\frac{3}{2} + \sqrt{\frac{9}{4} + 4ab}}t\right)^{-3/2} e^{-at - \frac{b}{t} + \sqrt{\frac{9}{4} + 4ab}}$$

$$\text{DoubleExponential}(t; B, \tau_1, \tau_2) = e^{-t/\tau_1} + Be^{-t/\tau_2}$$

To quickly reject poor functions, the functions were tested against the first 1000 waveforms provided in the data file, and then ranked based on their average χ^2 value from fitting the waveforms. The four most promising functions found are shown above.

3.3 Averaged Waveforms

Each fit gives a easily calculable value for the pulse height and the peak timing. To check how consistent the waveforms were, these values were used to segregate the waveforms into 10 different bins for pulse heights, and 50 different bins for peak timings. The waveforms belonging to each bin were then averaged, and then compared at fixed times to examine the dependence of the shape on the pulse height. The results are highlighted in Figure 3.2.

3.3. Averaged Waveforms

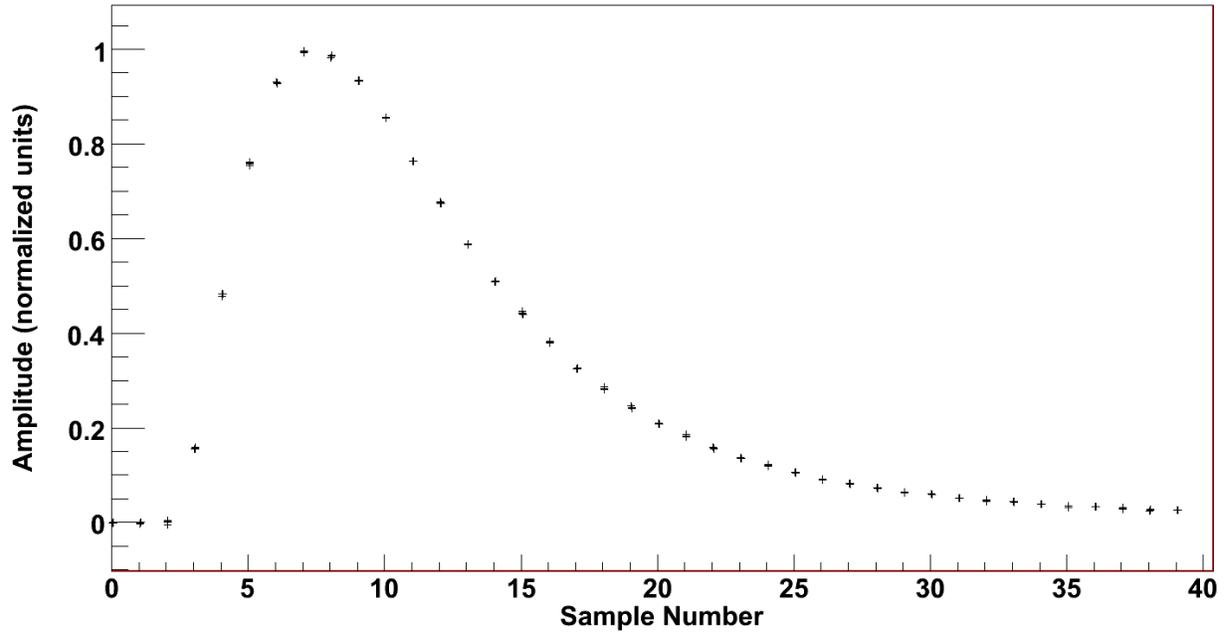


Figure 3.2: Plot of the 10 averaged waveforms at a fixed time, normalized such that the peak is at 1. The discrepancy between waveforms of different pulse heights is minimal.

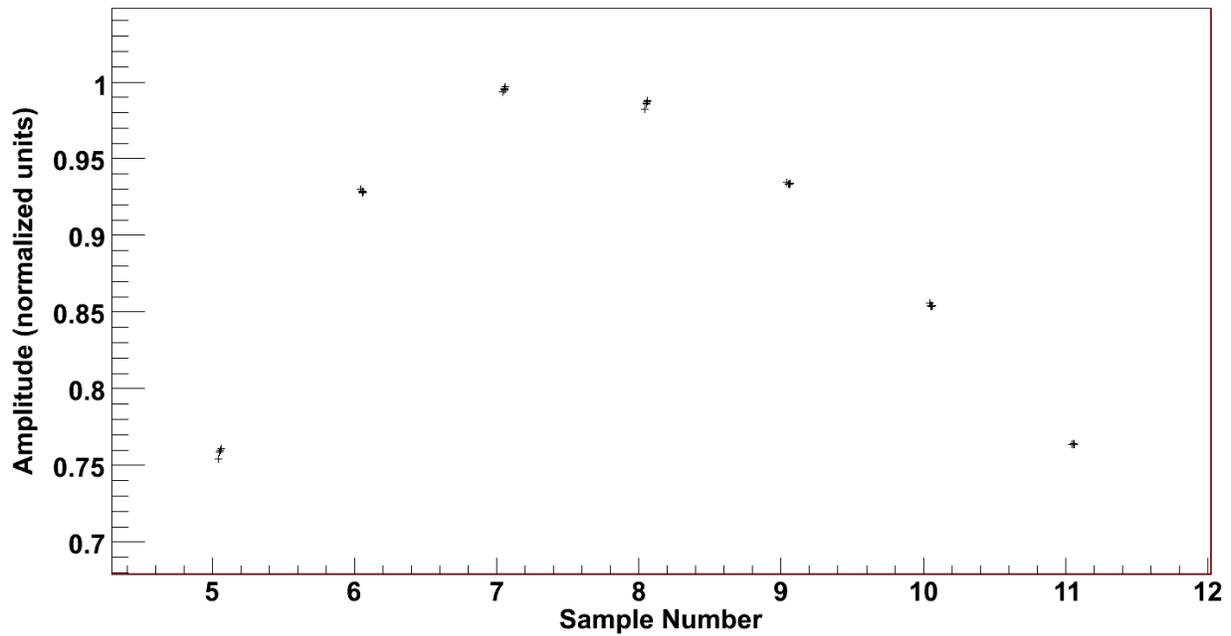


Figure 3.3: Same plot as Figure 3.2, zoomed in at the peak.

3.4 Comparing the Fits

As Figure 3.2 suggests, there is no dependence on the waveform shape on the pulse height. The reasoning in Section 2.3 now applies and the quality of each fit can be compared by calculating α in Equation 2.18 by fitting the distribution shown in Figure 3.4. However there is a difficulty in directly using the mean of the χ^2 , as it is sensitive to the massive values the fit produces for pileup events. Specifically, the presence of pileup events effects the shape of the fit, and thus poses additional quadratic factors in the mean of the χ^2 , artificially inflating α and thus making it unsuitable to use except for very rough upper bounds.

For a fixed pulse height, the χ^2 values follow a χ^2 distribution, save for large values caused by pileup. While these outliers effect the mean, they do little to change the mode of the distribution, which is a known difference from the mean. The value of α was therefore calculated using the mode of the distribution instead, eliminating the effect of pileup. Out of all the candidate functions trialled, the Inverse Gaussian function performed the best, with $\alpha \approx 0.029$. While this alone would give improved energy resolution, there were clear areas of improvement. Examining the residues of the fit have shown hints that the shape of the waveform and the Inverse Gaussian function had significant differences.

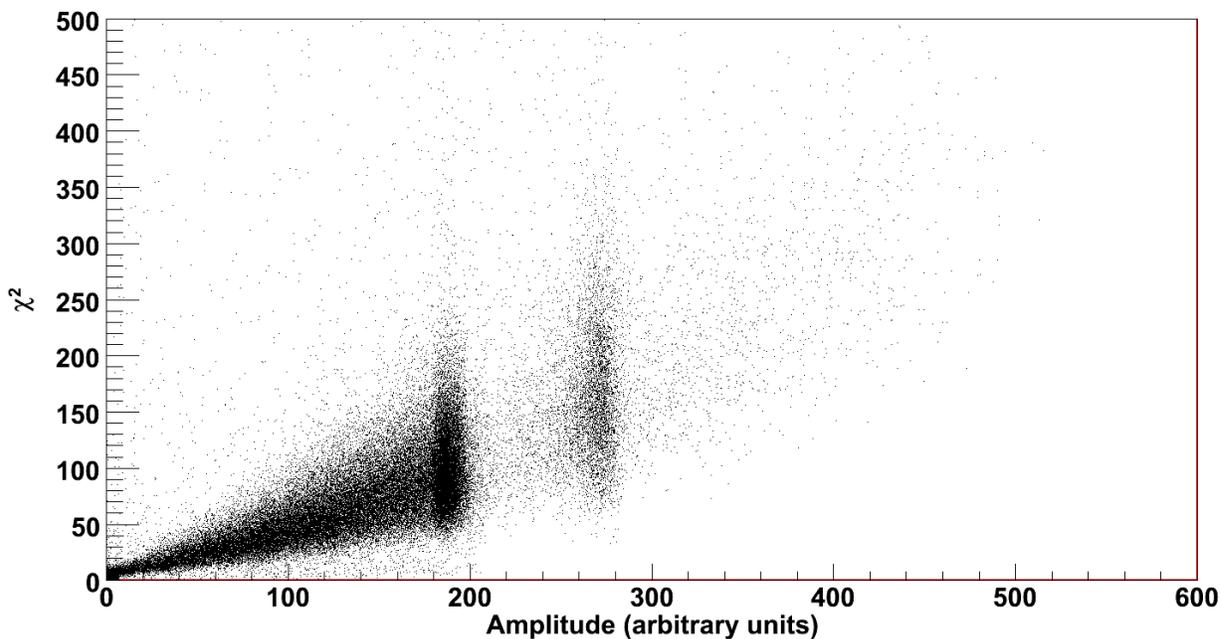


Figure 3.4: χ^2 distribution plotted against pulse height for the Inverse Gaussian function. Note the distinct quadratic dependence of the χ^2 on the pulse height.

3.5 Collation

3.5.1 Problems with the Averaged Waveforms

Although the NaI waveforms were shown to be consistent across pulse heights as shown earlier, the discretization of the 40 given samples can give a grossly inaccurate estimation of the actual

integral, especially at areas of high curvature in the waveform (Figure 3.5). Consequently choosing just one of the averaged waveforms, or using the average of all waveforms as a fit is insufficient for obtaining a good energy resolution.

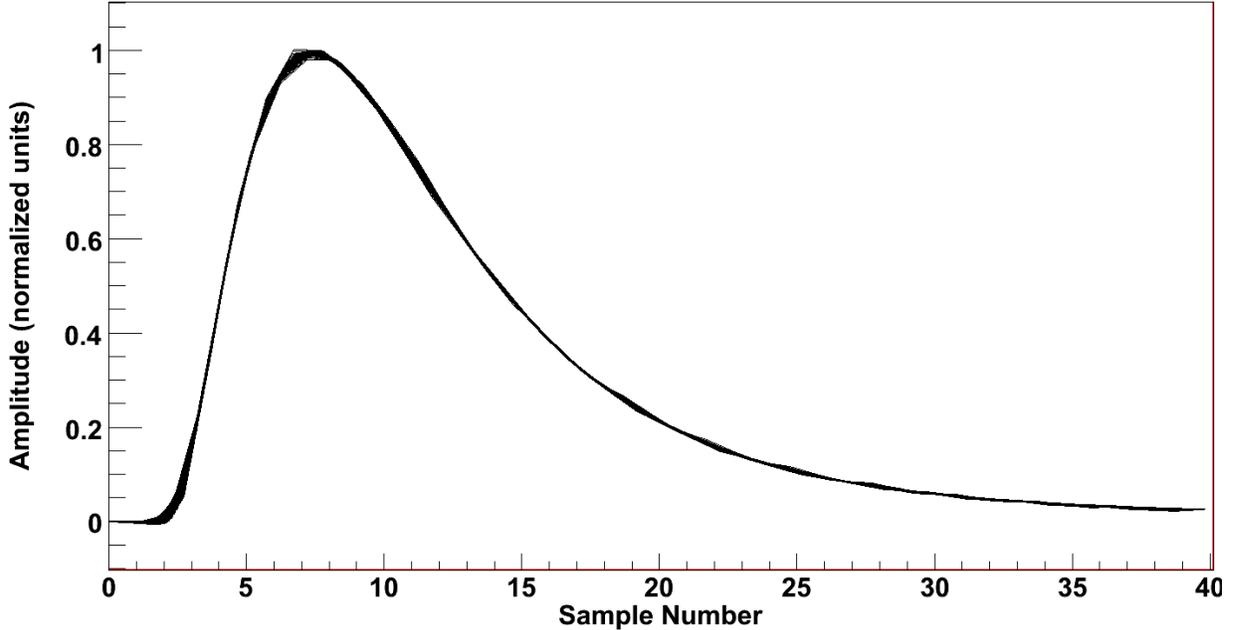


Figure 3.5: Plot of the normalized waveforms referred to earlier with linear interpolation. While the waveforms agree, their interpolations have large discrepancies, specifically at points of high second derivative. This leads to significant differences in the integral.

Noticing that the other waveforms provide the necessary missing data between samples of a specific waveform, the solution employed in this project was to collate all of the averaged waveforms, resulting in a singular waveform with effectively many times the original amount of samples.

3.5.2 Process

To prevent favoring any discretization, each averaged waveform was linearly interpolated, and then used to fit each other averaged waveform. From the fits 3 parameters regarding the relative amplitude, time, and pedestal differences were extracted. Each averaged waveform was associated with a count, i.e. the number of waveforms averaged together to create the averaged waveform. The averaged waveform with the highest count was chosen as the anchor. For each specific waveform, the differences between the waveform and the fitting waveform were combined with the differences between the fitting waveform and the anchor, essentially yielding the differences between the waveform and the anchor via the fitting waveform. By the Central Limit Theorem [14] the standard deviation of each averaged point is proportional to $n^{-1/2}$ where n is the sample size. With this in mind, the differences were then averaged across all the fitting waveforms, weighted appropriately by the counts.

In this fashion, each waveform is assigned a set of 3 differences from the anchor, with the result having appropriate contributions from every other waveform. The waveforms were then

transformed using the 3 fitted differences to the same location as the anchor, resulting in the collated waveform.

3.5.3 Smoothness Test

Rather than undergoing the computationally expensive process of deriving a fit from the collation and testing it against all of the data, a method of testing the collation quality was developed. A well collated waveform necessarily ensures that the derivative is smooth as well. Using this knowledge, the discrete first and second derivatives of each average waveform was calculated, and then collated using the same parameters as the parent waveform. Some internal parameters in the collation process were subsequently tuned as to minimize the artifacts produced in the collated first and second derivatives Figure 3.6.

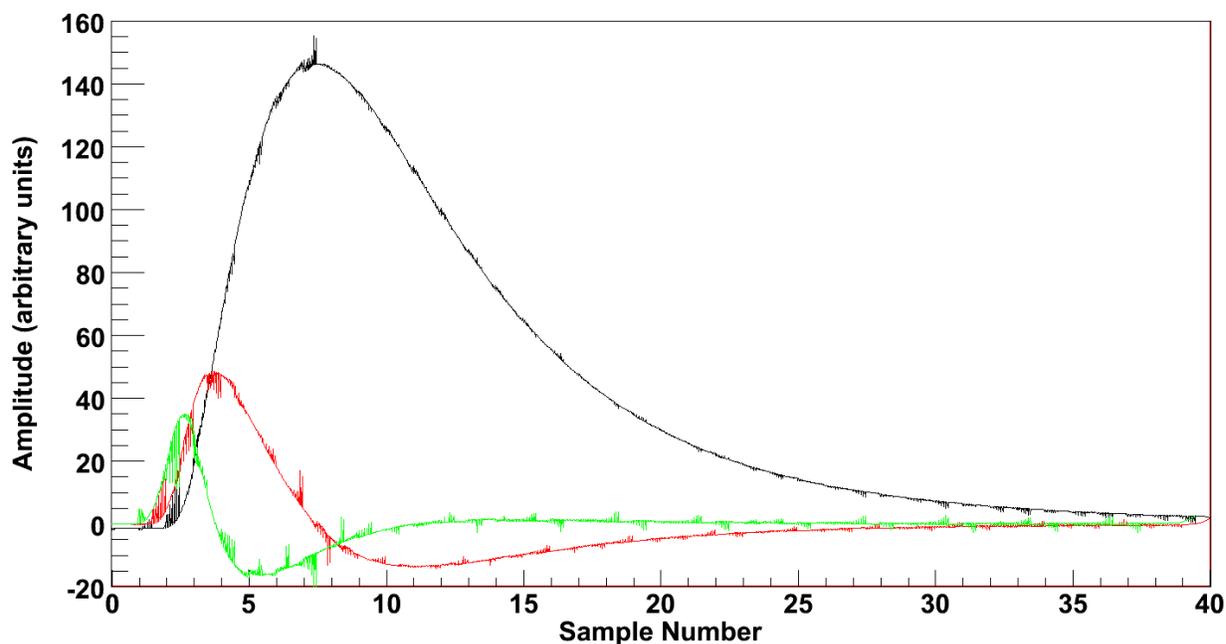


Figure 3.6: The collated waveform (in black) along with collated first and second derivatives (in red and green respectively). Notice how the collated derivatives heavily exaggerate any imperfections in the collation

A method of obtaining the collated waveform via fitting the collated derivative instead was briefly explored, but was deemed unsuitable for use due to difficulties in predicting the constants produced by integration as well as problems with the increased complexity of the function after integration.

3.5.4 Spline Fit

With the creation of the collated waveform, the reason why the candidate waveforms used earlier were inadequate became immediately evident. While they had delivered some good results in fitting individual waveforms, none of them matched the shape of the collated waveform. Trial

and error methods of producing a function that will fit the collated waveform was likely to result in failure.

A new fit was produced by using splines to approximate the collated waveform. While interpolation would have offered a more accurate representation of the collated waveform, it would have also incorporated some of the imperfections highlighted by the smoothness test. Uniform B-Splines were chosen for creating the new fit due to their low descriptonal complexity. The decreased degrees of freedom due to low numbers of control points mitigate possible overfitting of the collated waveform.

Ways of automatically producing the spline such as knot reduction was explored, but did not converge quickly enough for low amounts of control points. Fitting the waveform directly with the spline was also not an option due to the copious amounts of parameters and local minima. In the end, the spline was produced by manually positioning the control points, and then optimizing the result slightly with the fitter (Figure 3.7).

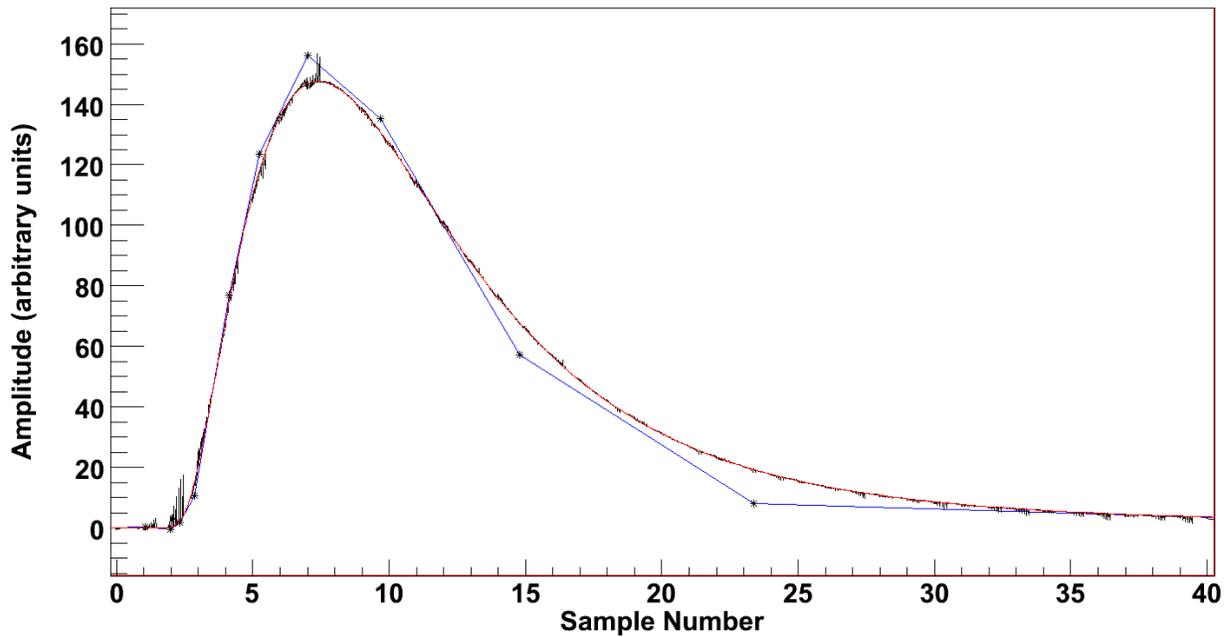


Figure 3.7: The spline fit (red) along with its control points (blue) and the collated waveform (black)

3.5.5 Iterative Improvements

The spline fit was implemented on the main data set, immediately showing an improvement over the Inverse Gaussian fit ($\alpha \approx 0.011$). The averaging process was redone, this time with 200 bins for time and one bin for pulse height. Since digital noise become less significant as the pulse height for the waveform increases, using just one bin for the pulse height means that all waveforms of a certain time range are summed, which places heavier weight on the waveforms with larger pulse heights.

The collation process was also edited to use the spline fit as the fitting waveform. The resulting

second collated waveform saw massively decreased amounts of distortion. With this, a second spline fit was also constructed, and then tested on the main data set.

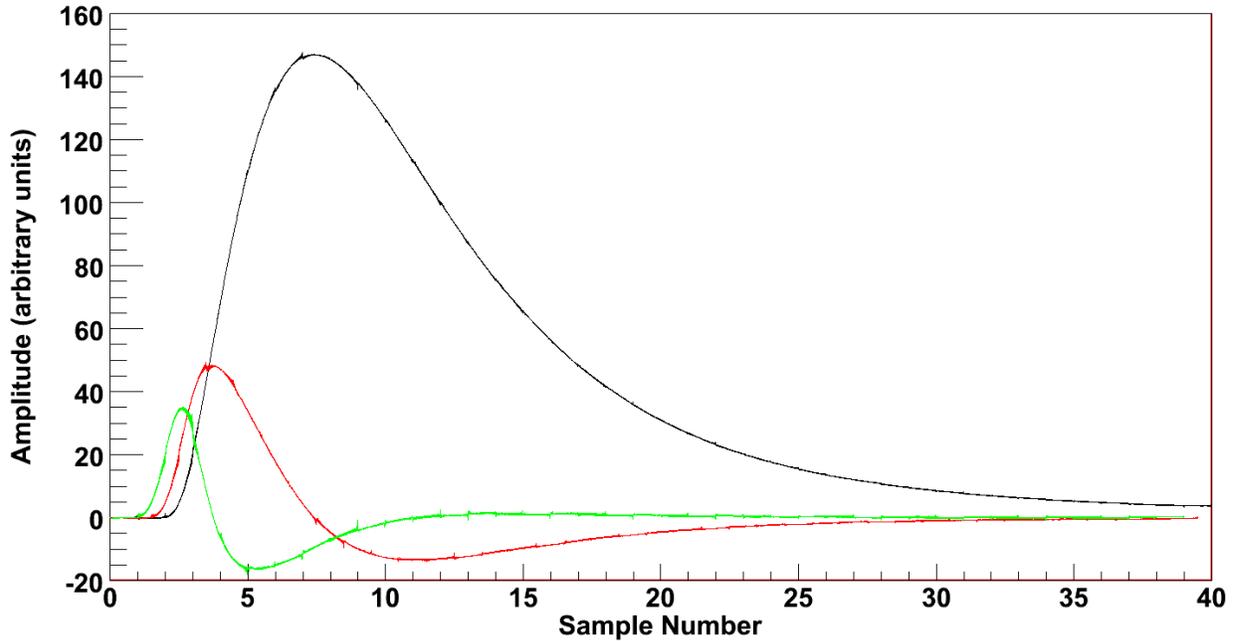


Figure 3.8: The second collated waveform (in black) along with its collated first and second derivatives (in red and green respectively). The noise seen in the waveforms is now subdued in comparison to the first collated waveform.

3.6 Time Resolution

While the uncertainty introduced by the fit in the energy (i.e. the integral) of the waveform is well defined by Equation 2.24, there is no equivalent expression for the uncertainty in time. To quantitatively measure the uncertainty in time, waveforms representing the data were simulated using the spline fit. The difference between the spline fit and the true waveform was compensated for using γ in Equation 2.18 to provide an estimate for the digitization error, and α to add an additional error to each point of the waveform. Since α is a generous upper bound on the difference between the spline fit and the true waveform, this additional factor will eliminate any biases from using the simulated data. These waveforms were then fitted with the spline fit and the difference between the waveform peak and the fitted peak was recorded.

While the error in time quickly tapers off as the pulse height increases, it is extremely large for waveforms of low pulse heights. This is to be expected as if the pulse height was zero, there would be no indication whatsoever to where the waveform is.

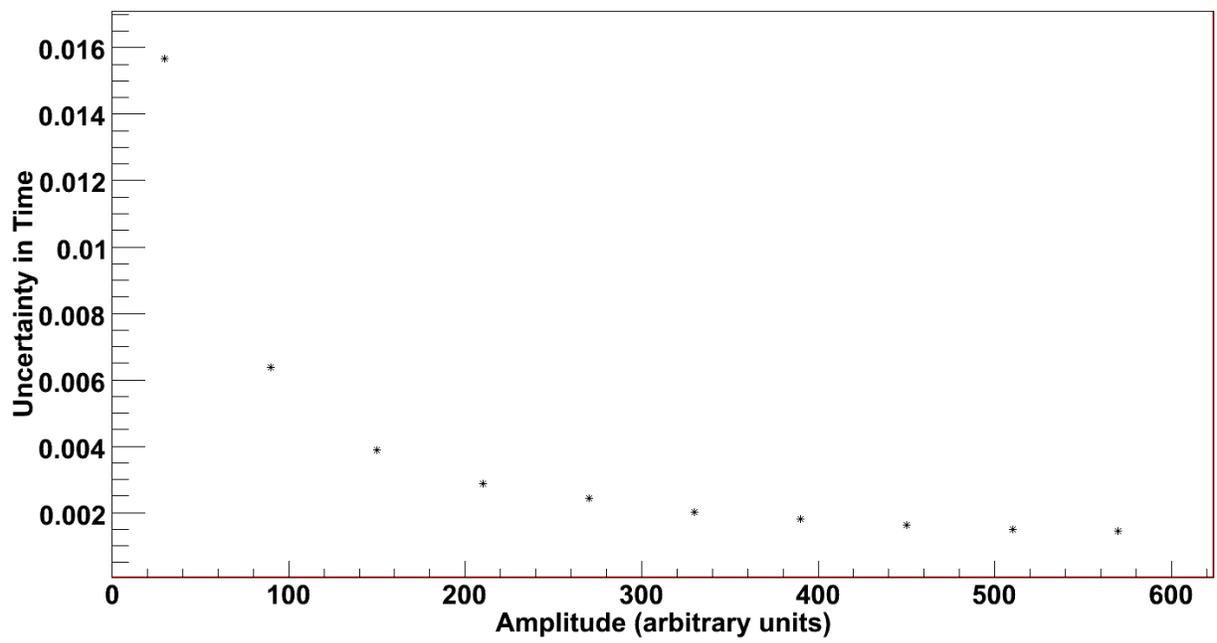


Figure 3.9: Plot of the time uncertainty (in number of samples) of the fit relative to the pulse height of the waveform.

Chapter 4

Pileup Detection

4.1 Overview

A secondary function of fitting was to eliminate pileup events. Due to their rigorous optimization in matching the shape of a single pulse waveform, the characteristic multi pulse pileup waveforms will on average show a large χ^2 value compared to typical waveforms as in 3.1. This chapter aims to illustrate some of the limitations the fit of the waveform has in detecting pileup events as well as some improvements that can be made to mitigate these limitations.

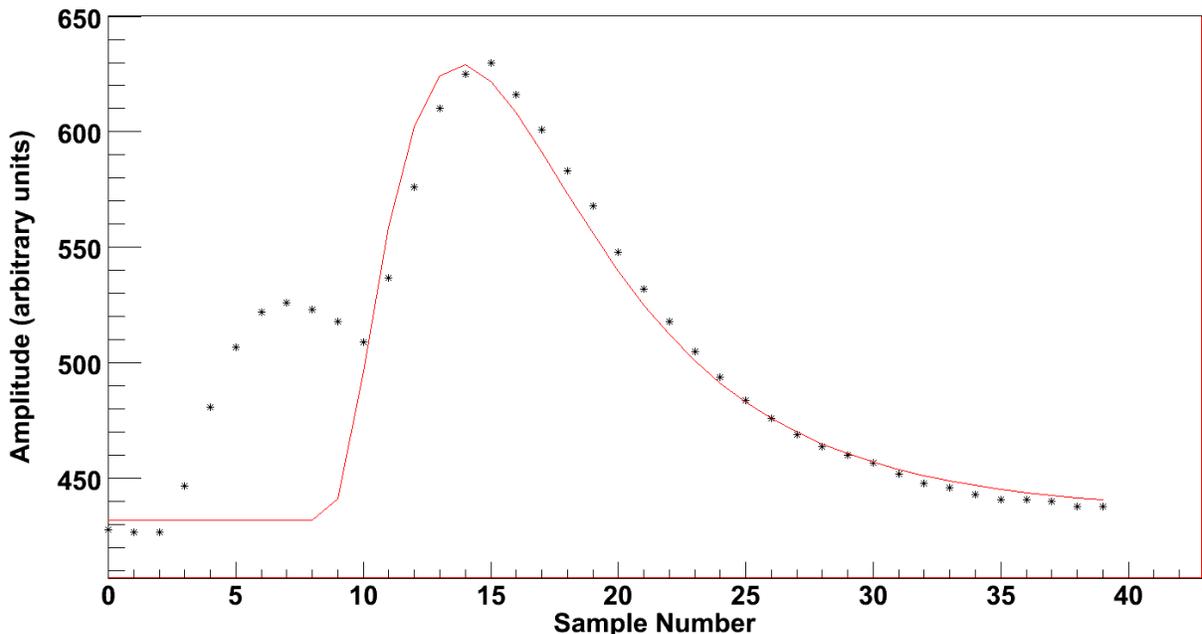


Figure 4.1: A typical two pulse pileup event. Note the failure of the fit to match the shape

4.2 Range Reduction

One way of eliminating pileup events is simply not to integrate them. This can be done by restricting the integrating window to within a certain range of the peak of the waveform. The extent of this method was tested by simulating a number of waveforms and then fitting the waveform based on data from within a certain range. The χ^2 value for the entire data set was then calculated and the results are shown in Figure 4.2.

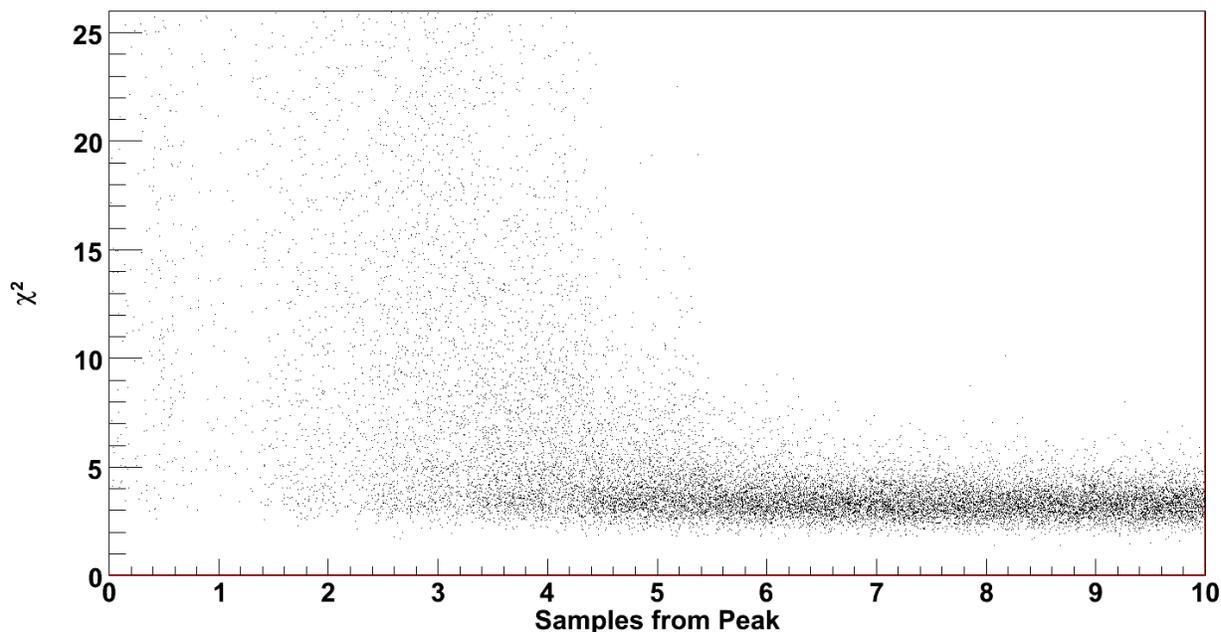


Figure 4.2: Distribution of the χ^2 values for various fitting window sizes. Note the stark increase in χ^2 after the range is reduced within ± 5 of the peak.

With this result, not only has limiting the range enable the use of some previously unusable waveforms, the fact that it uses a quarter of the samples given presented sizable computational speedups for the fitting algorithm.

4.3 Pileup Cut

By cutting the χ^2 at a certain value, many plausible pileup events can also be eliminated. Pileup was simulated by adding secondary pulses with a pulse height 100 times smaller than the primary pulse height and a varying peak time. The pileup waveforms were then fitted and used to qualitatively examine the efficiency of the fit in eliminating pileup.

A properly triggered event will have a peak somewhere between sample number 6 and 8, so the points of interest in Figure 4.3 lie between -8 and 34 samples. In this region, the cut at $1.5E[\chi^2]$ eliminates all events with secondary pulse peaks outside of ± 1.5 samples from the peak. Unfortunately, pileup with pulse heights over 200 times smaller than the primary pulse height start to become virtually undetectable by the fitter.

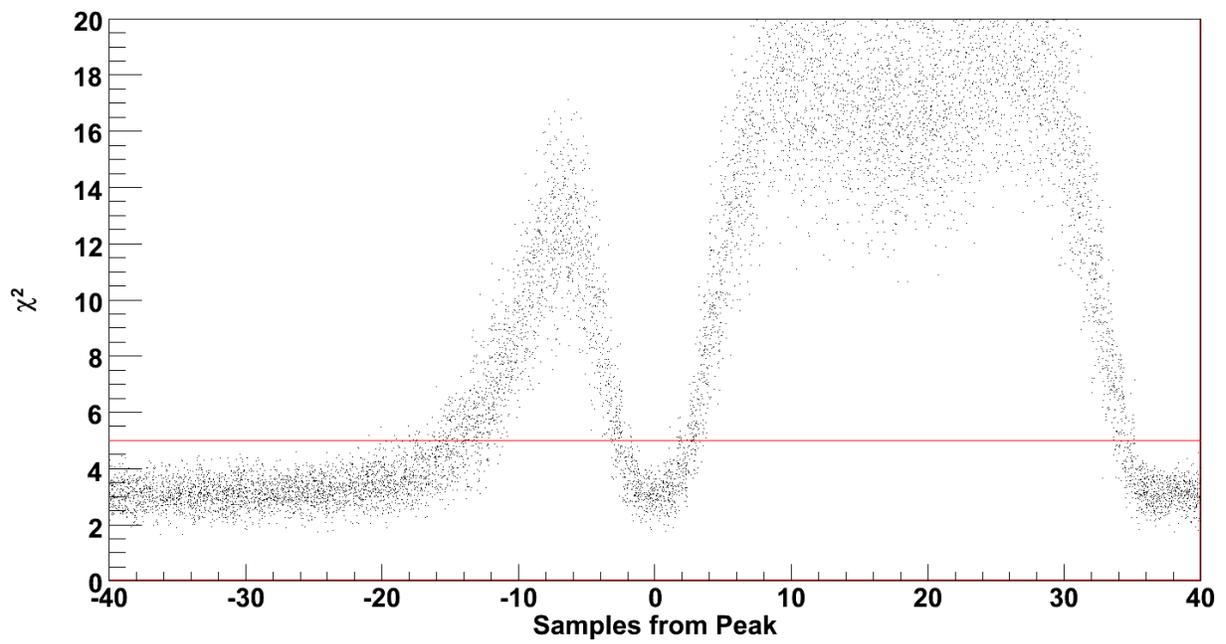


Figure 4.3: Distribution of the χ^2 values for various fitting window sizes relative to the amount of samples between the primary pulse peak and the secondary pulse peak

Chapter 5

Results

5.1 Energy Resolution

Table 5.1 shows the α , β , and γ values for a selection of fits used in this experiment. The Linear fit is based on the linear interpolation of an averaged waveform and is included for completeness to show the error expected from trapezoidal integration.

Fit	α	β	γ	σ_Q/Q
Linear	0.041	0.25	2.3	0.7%
InvGamma	0.036	0.20	2.0	0.6%
InvGauss	0.029	0.19	2.0	0.5%
Spline (v1)	0.011	0.22	2.1	0.2%
Spline (v2)	0.0097	0.19	2.3	0.16%

Table 5.1: List of α , β , and γ values for fits used in this experiment

The spline fit developed for this project yields an uncertainty of at most 0.16% for the energy. Compared to 0.7% for using the linear fit, this marks an improvement by a factor of 4 for the energy resolution.

5.2 Time Resolution

Disregarding digitization noise, the difference in time from taking the maximum point as the peak would be uniformly distributed from -0.5 to 0.5 samples. This corresponds to a standard deviation of 0.144 samples as the lower bound for the uncertainty in time. As Figure 3.9 suggests, for waveforms of pulse heights above 20, the uncertainty in time is at most 0.016 samples, which is a factor of 9 better than just using the discrete time. For pulse heights above 150, the uncertainty in time is reduced to at most 0.004 samples, a factor of 36 better than using discrete time.

5.3 Pileup Elimination

According to Section 4.2, the integration range can be reduced to as much as 5 samples from the peak time without any adverse effects on the quality of the fit. In addition, Section 4.3 suggests that by cutting out waveforms with $\chi^2 > 1.5\mathbb{E}[\chi^2]$, the spline fit can effectively eliminate all pileup events with secondary pulse peaks at least 1.5 samples away from the primary pulse peak and above 1% of the primary pulse height. Subsequent experimentation has shown that below 0.5% this method becomes ineffective for eliminating pileup as the secondary pulse heights start to become comparable to noise.

Chapter 6

Conclusions

6.1 Results

The fit produced from this work brings about at least a factor of 4 improvement in the energy resolution, and at least factor of 9 improvement in the time resolution. While this does not directly improve the uncertainty in the branching ratio, it does help improve the effectiveness of the myriad of different cuts used to calculate the branching ratio. Additionally, the collation process shown in this project is trivially generalizable to all discrete waveforms with a pulse height and peak time independent shape.

6.2 Plausible Improvements

Although the collated waveform itself is smooth, there is significant amounts of noise evident in its derivatives (Figure 3.8). This suggests that there is room for improvement in constructing the collated waveform. The averaging process currently averages all waveforms within a certain time bin. This introduces an error in the peak time proportional to the bin size and causes the averaged waveform to be slightly wider than the true waveform. While this does not effect the results in any way, as the spline fit induced is still tested on the data itself, some method of interpolating waveforms within a bin such that they all have the same peak time may improve the quality of the collated waveform.

The process of procuring the spline from the collated waveform is also manual. Based on visual inspection however, it is debatable whether switching to an automatic method will bring significant improvements to the spline fit. While further iterative improvements in the spline fit may show better results, the improvement in the second iteration of the spline fit as seen in Table 5.1 is not as stark as the difference between the first iteration of the spline fit and the Inverse Gaussian fit.

Bibliography

- [1] G. Altarelli and M. W. Grünewald. Precision electroweak tests of the standard model. *Physics Reports*, 403404(0):189 – 201, 2004. {CERN} - the second 25 years.
- [2] M. K. Gaillard, P. D. Grannis, and F. J. Sciulli. The standard model of particle physics, 1998.
- [3] B. Martin and G. Shaw. *Particle Physics*. Manchester Physics Series. Wiley, 2013.
- [4] T. Fazzini, G. Fidecaro, A. W. Merrison, H. Paul, and A. V. Tollestrup. Electron decay of the pion. *Phys. Rev. Lett.*, 1:247–249, Oct 1958.
- [5] H. L. Anderson, T. Fujii, R. H. Miller, and L. Tau. Branching ratio of the electronic mode of positive pion decay. *Phys. Rev.*, 119:2050–2067, Sep 1960.
- [6] E. Di Capua, R. Garland, L. Pondrom, and A. Strelzoff. Study of the decay $\pi \rightarrow e + \nu$. *Phys. Rev.*, 133:B1333–B1340, Mar 1964.
- [7] D. A. Bryman, R. Dubois, J. A. MacDonald, T. Numao, B. Olaniyi, A. Olin, J.-M. Poutissou, and M. S. Dixit. Measurement of the $\pi \rightarrow e\nu$ branching ratio. *Phys. Rev. D*, 33:1211–1221, March 1986.
- [8] D. E. Groom et al. Review of particle physics. *The European Physical Journal*, C15:1+, 2000.
- [9] L. Doria. The pienu experiment: a precision measurement of the pion branching ratio, 2010.
- [10] C. Malbrunot et al. The pienu experiment at triumf : a sensitive probe for new physics. *Journal of Physics: Conference Series*, 312(10):102010, 2011.
- [11] K. A. Olive and Particle Data Group. Review of particle physics. *Chinese Physics C*, 38(9):090001, 2014.
- [12] D. Bryman, W. J. Marciano, R. Tschirhart, and T. Yamanaka. Rare kaon and pion decays: Incisive probes for new physics beyond the standard model. *Annual Review of Nuclear and Particle Science*, 61(1):331–354, 2011.
- [13] J. B. Hartle. *Gravity: An Introduction to Einstein’s General Relativity*. Addison-Wesley, 2003.
- [14] Rick Durrett. *Probability: Theory and examples*, 2011.