

A high-precision measurement of the $\pi \rightarrow e\nu$ branching ratio

by

Tristan Sullivan

BSc., University of Victoria, 2009

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

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies

(Physics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

April 2017

© Tristan Sullivan 2017

Abstract

The pion decay branching ratio, $R_\pi = \frac{\Gamma(\pi^+ \rightarrow e^+ \nu_e + \pi^+ \rightarrow e^+ \nu_e \gamma)}{\Gamma(\pi^+ \rightarrow \mu^+ \nu_\mu + \pi^+ \rightarrow \mu^+ \nu_\mu \gamma)}$, is an important observable in the Standard Model of particle physics. The value of the $\pi \rightarrow e \nu$ branching ratio has been calculated within the Standard Model to be $(1.2352 \pm 0.0002) \times 10^{-4}$ [1] [2]. The PIENU experiment at TRIUMF aims to measure this quantity to a precision of $< 0.1\%$. This tests the hypothesis that the leptons have identical weak couplings, known as lepton universality, at the 0.05% level. In addition, it provides stringent constraints on many other extensions to the Standard Model, such as R-parity violating supersymmetry, leptoquarks, and heavy neutrinos lighter than the pion. In certain cases, these constraints can far exceed the reach of direct searches at colliders. Most strikingly, a new pseudoscalar interaction whose energy scale were $\mathcal{O}(1000 \text{ TeV})$ would enhance the branching ratio by $\mathcal{O}(0.1\%)$.

The PIENU data set contains four years of data, taken between 2009 and 2012. The analysis of a subset of the 2010 data was published in 2015 [3]; the precision obtained for the branching ratio was approximately 0.25%. The 2012 data set is roughly five times larger than the 2010 data set, and its analysis is presented here. The statistical error using only 2012 data is 0.09%; incorporating the other data sets will reduce this to 0.07%. The systematic error in the 2012 analysis remains considerably larger than this, 0.27%, and the prospects for reducing it to the level needed to reach the precision goal of the experiment will be presented. The present work has been done as a blind analysis, with an unknown factor masking the branching ratio result; the final value will be obtained when the analysis is completed.

Preface

The PIENU collaboration consists of about twenty people from six different countries. Proposed in 2006, data-taking for the experiment began in the spring of 2009, and continued until December 2012. I joined the experiment as a Masters student in September 2009 and transferred to the PhD program the following year. Throughout the data-taking period, from 2009 to 2012, I took many shifts monitoring the data collection, and during 2011 and 2012 was responsible for the trigger system, including modifying it as necessary and documenting any changes.

I was primarily responsible for measuring the response function of the PIENU calorimeter, which was necessary to obtain the largest correction to the branching ratio, and also expected to be the largest source of systematic uncertainty in the experiment. Special data were taken to evaluate this correction; this was first done at the end of 2009, but the uncertainty was too large, and it was done again in 2011. For this measurement I was responsible for devising a method to reduce the uncertainty and ensuring the quality of the data was sufficient, as well as for its subsequent analysis.

I also contributed to the main branching ratio analysis, which was published in 2015 using part of the data taken in 2010 [3]. I was primarily responsible for the analysis of the 2012 data set, which was approximately five times larger than the 2010 data set. Although the desired level of precision has not yet been achieved, the sources of systematic uncertainty that still need to be reduced, and ways of potentially achieving this, have been identified.

Table of Contents

Abstract	ii
Preface	iii
Table of Contents	iv
List of Tables	viii
List of Figures	x
Glossary	xxii
Acknowledgements	xxx
Dedication	xxxii
1 Introduction	1
1.1 The Standard Model	1
1.1.1 The weak interaction	2
1.2 Motivation for the PIENU experiment	3
1.3 Brief history of pion branching ratio measurements	4
1.4 Experimental technique	5
1.4.1 PIENU experimental technique	9
1.5 Outline of the thesis	11
2 Theory	12
2.1 Pions in the Standard Model	12
2.1.1 Pion decay modes	12
2.1.2 The weak interaction	12
2.1.3 Pion decay rate	14
2.1.4 Corrections	16
2.2 Beyond the Standard Model	19
2.2.1 New pseudoscalar interactions	20

Table of Contents

2.2.2	R-parity violating supersymmetry	21
2.2.3	Charged Higgs	22
2.2.4	Leptoquarks	24
2.2.5	Massive neutrinos	24
2.3	Lepton universality	25
2.4	Physics Reach	28
3	Experimental Setup	29
3.1	Beamline	29
3.2	Detector	30
3.3	Trigger and DAQ	43
3.4	Event types	53
3.4.1	Beam-related background	53
3.5	Monte Carlo	57
4	Data Taking and Processing	58
4.1	Running periods	58
4.1.1	2009	58
4.1.2	2010	59
4.1.3	2011	59
4.1.4	2012	59
4.2	Data processing	59
4.2.1	Plastic scintillators	59
4.2.2	Wire chambers	61
4.2.3	Silicon strips	61
4.2.4	Crystal scintillators	61
4.3	Blinding	61
4.4	Energy calibration	62
4.5	Track reconstruction	67
4.6	Data stability	72
4.7	Timing	74
5	Raw Branching Ratio Extraction	80
5.1	Event selection	80
5.1.1	Good run selection	80
5.1.2	Time spectrum	80
5.1.3	Pion selection	82
5.1.4	Pileup rejection	83
5.1.5	Acceptance cut	91
5.1.6	Minor cuts	92

Table of Contents

5.2	Fitting function	94
5.2.1	Low energy time spectrum	94
5.2.2	High energy time spectrum	94
5.3	Fitting method	99
5.4	Results	101
5.5	Systematic checks	109
5.6	Summary of Chapter 5	113
6	Tail Correction	115
6.1	Introduction	115
6.2	Lower limit	116
6.3	Response function measurement	120
6.3.1	Energy loss processes	120
6.3.2	Detector setup	122
6.4	Data-taking	124
6.4.1	Event selection cuts	125
6.4.2	Muon correction	126
6.5	Other systematics	132
6.5.1	Background	133
6.5.2	Calibration and resolution	133
6.6	Positron beam simulation	134
6.7	Determining the PIENU tail fraction	139
6.8	Tail as a function of R and E_{cut}	145
6.9	Summary of Chapter 6	148
7	Other Corrections	149
7.1	Acceptance correction	149
7.1.1	Pion stopping position	150
7.1.2	Detector geometry	153
7.1.3	Trigger thresholds	153
7.2	Muon decay-in-flight correction	154
7.3	t_0 correction	155
7.4	Stability of the corrected branching ratio	155
7.5	Summary of Chapter 7	158
8	Conclusion	160
8.1	Branching ratio result	160
8.2	Future prospects	162
8.2.1	Systematic uncertainty	162
8.2.2	Limits on new physics	163

Table of Contents

Bibliography 166

Appendices

A Trigger Diagram 172

B Timing Diagrams 173

C Event Selection For Positron Beam Data 179

D Positron Data Systematics 185

List of Tables

1.1	The particle content of the SM; e is the charge of the electron.	1
1.2	Coupling constants for SM interactions.	3
2.1	Pion decay modes and branching ratios.	12
2.2	Corrections to the leading-order value for R_π .	19
2.3	Experimental results on lepton universality tests from studies of pion, kaon, tau, and W decays. Here \mathcal{B} represents the branching fraction of a particular decay mode.	27
3.1	Detector characteristics. The z position given is for the centre of the detector, except as noted for BINA.	33
3.2	The number of events in one 2012 run caused by the physics triggers, and the most important calibration triggers.	46
3.3	Detector readout channels.	49
4.1	Running periods.	58
5.1	The probability of fake hits after the real hit, in each T1 PMT.	89
5.2	Events removed by each cut, with every other cut applied.	93
5.3	Fit parameter list.	100
5.4	Fit parameters for the 2010 and 2012 data sets. Note that the 2012 branching ratios are still blinded by an unknown factor uniformly distributed between $\pm 0.5\%$.	109
5.5	Systematic checks performed on the fit of the 2012 data.	110
6.1	Angles at which positron beam data were taken.	124
6.2	Properties of the muon energy spectra as a function of angle.	130
6.3	ΔT at 0° .	131
6.4	ΔT at 11.8° .	131
6.5	ΔT at 20.9° .	131
6.6	ΔT at 30.8° .	131
6.7	ΔT at 41.6° .	131

List of Tables

6.8	ΔT at 6°	131
6.9	ΔT at 16.5°	131
6.10	ΔT at 24.4°	131
6.11	ΔT at 36.2°	131
6.12	ΔT at 47.7°	131
6.13	The tail fraction as a function of angle, with errors due to statistics and the variation in the muon-corrected values.	132
6.14	Corrections and resulting $\pi^+ \rightarrow e^+\nu_e$ tail fractions as a function of E_{cut}	147
6.15	Tail fractions as a function of the maximum radius in which events are accepted.	147
6.16	Tail fractions as a function of the radius in which events are accepted.	148
7.1	The small corrections that must be made to the branching ratio. The values are multiplied by the tail-corrected branching ratio to give the final result.	149
8.1	Sources of error. The corrected branching ratio is given by the product of the raw branching ratio and the corrections. The errors given for the corrections are the errors on the corrections themselves, not the resulting errors on the branching ratio. The stars indicate that the result is still blinded.	161
D.1	The change in the tail fraction as beam parameters and detector geometry were varied. The values given are the nominal tail fraction minus the new tail fraction (see Section 6.6 for a detailed description of what was changed). Note that the change is given as a fraction of the total spectrum, not the tail. The upper part shows the results for variations that increased the tail, and the lower part shows the results for variations that decreased the tail. The errors are due to Monte Carlo statistics.	185

List of Figures

1.1	The measured value of the branching ratio as a function of time. The dashed line shows the SM prediction. The last point indicates the expected level of precision that will ultimately be achieved by combining the results of the PIENU experiment and the PEN experiment at PSI, as described in the text.	6
1.2	Simplified picture of the PIENU experimental technique. Pions stop in the target and decay into either muons or positrons; muons also stop in the target, and decay into positrons. Positrons, and photons if any are produced, are detected by the crystal scintillator calorimeter.	7
1.3	GEANT4 [4] simulation of the time spectra of $\pi^+ \rightarrow e^+\nu_e$ (red) and $\pi \rightarrow \mu \rightarrow e$ (blue) events. The distributions are normalized to the same height.	8
1.4	GEANT4 simulation of the starting energies of positrons due to $\pi^+ \rightarrow e^+\nu_e$ (red) and $\pi \rightarrow \mu \rightarrow e$ (blue) decays. The distributions are normalized to the same height.	8
1.5	The PIENU detector, with the pion beam incident from the left. The region close to the target is shown in the blowup. Plastic scintillators (polyvinyltoluene) are shown in dark blue, wire chambers in green, silicon strip detectors in orange, and crystal scintillators in light blue (NaI(Tl)) and red (CsI). . . .	10
2.1	Feynman diagram for the decay of the positively charged pion. W^+ represents the positively charged gauge boson mediating the weak interaction, and u and \bar{d} represent an up and an anti-down quark, respectively.	14

List of Figures

2.2	The allowed directions of the spins and linear momenta of the pion decay products, in the pion rest frame. The requirement, from angular momentum conservation, that the spins of the positron and the neutrino must be in opposite directions, leads to the suppression of the positron mode relative to the muon (see text).	16
2.3	Feynman diagrams for the radiative corrections to pion decay, from real (a) and virtual (b) photons. \bar{l}^+ denotes an anti-lepton.	17
2.4	The constraint on the coefficients of R-parity violating interactions from a fit of electroweak observables, including R_π . The blue curve shows the constraint using the PDG value for the branching ratio, and the dashed red curve shows the projected constraint from a 0.1% measurement of the branching ratio, with the same central value as the blue curve. The green curve shows the expected limits with the results of the Qweak experiment, measuring the weak charge of the proton, at Jefferson Lab [5].	23
2.5	The 90% C.L. upper limit on the heavy neutrino mixing parameter, as a function of its mass. The dashed line shows the result from the previous PIENU experiment [6], and the circles and triangles are the limits from a subset of PIENU data, published in 2011 [7]. The circles indicate a restricted angular region was used when constructing the $\pi^+ \rightarrow e^+ \nu_e$ energy spectrum.	26
3.1	A schematic of the M13 beamline.	30
3.2	The end of the M13 beamline, before (left) and after (right) the extension. Part of the detector was in place to measure the particle content of the beam.	31
3.3	The transverse position distributions of pions, muons, and positrons at F3; the particle type was identified using time of flight. The black lines are Gaussian fits to the pion and positron distributions.	32
3.4	One wire chamber plane and its preamplifier board; each chamber consisted of three planes.	34
3.5	Wire chambers 1 and 2 after installation on the beam pipe.	34
3.6	The plastic scintillator readout scheme, for B1, B2, Tg, and T1. The light from the plastic scintillator (purple) was transmitted by four acrylic lightguides (light green) to PMTs (grey cylinders).	35

List of Figures

3.7	The PIENU-I detector subsystem.	35
3.8	The PIENU-II detector subsystem.	36
3.9	The BINA detector on the test bench, with some of its PMTs in place.	36
3.10	BINA and the two CsI rings, with all the BINA PMTs in place.	37
3.11	A Solidworks drawing of PIENU-I and PIENU-II, along with a picture with a human for scale. During data-taking, PIENU- II was rolled forward around PIENU-I.	38
3.12	Wires hit in each plane of the first wire chamber.	39
3.13	Energy deposited in B1 vs. the time of the hit relative to the peak of the cyclotron RF field. The cluster with the most events, labelled π , is caused by pions; the cluster below that, with the same timing but less energy, is caused by pions that decayed in flight prior to reaching B1. The hit is caused by the decay muon. The cluster labelled μ on the far left is due to beam muons, and the low-energy cluster labelled e is due to beam positrons.	40
3.14	The measured energy deposited in B1, B2, and Tg by pi- ons. The additional peak in the target spectrum is caused by events in which the energy of the 4.1 MeV decay muon is also included.	41
3.15	The simulated distance travelled by the 4.1 MeV $\pi \rightarrow \mu\nu$ decay muon in the target.	42
3.16	The simulated pion stopping position along the beam direction.	42
3.17	The simulated energy deposited for $\pi^+ \rightarrow e^+\nu_e$ decays in BINA (black) and BINA and CsI (blue), and for $\pi^+ \rightarrow \mu^+ \rightarrow$ $e^+\nu_e\bar{\nu}_\mu$ decays in BINA and CsI (red). The distributions are normalized to the maximum bin. The low-energy peak is caused by the absorption of a single 511 keV photon from a positron annihilating at rest, in T2 or the front face of BINA.	44
3.18	Time spectra of events recorded by the Early (red) and Prescale (black) triggers. The Prescale trigger events are scaled by a factor of sixteen. The time spectrum is dominated by $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ events. Both the end of the early time window and the end of the trigger window can be seen. . . .	46
3.19	Measured energy deposited in the NaI(Tl) crystal, for events recorded by the BinaHigh trigger. The sharp rise is caused by the threshold of the BinaHigh trigger, and the fall by the endpoint of the Michel distribution. The few events above the Michel edge are mostly due to pileup.	47

List of Figures

3.20	Time spectrum of hits in B1 (top) and T1 (bottom).	50
3.21	The silicon strip readout. Each strip is connected to a capacitor and there is one readout channel per four strips [8].	52
3.22	A muon decaying in the centre of the target and the decay positron going past the edge of T1. The plastic scintillators and BINA are shown; distances are to scale.	55
3.23	Simulated x and y distributions at the first plane of WC3 for $\pi \rightarrow \mu \rightarrow e$ events that deposit energy in BINA and not in T1 (left) and for $\pi \rightarrow \mu \rightarrow e$ events that deposit energy in BINA (right).	55
3.24	The PIENU detector and beamline after the last bending magnet, showing the steel wall used for neutron shielding.	56
4.1	PMT waveform from a pion in B1. The points at the beginning and the end of the waveform are zero-suppressed; the drop around -1380 ns is to the level of the pedestal.	60
4.2	The blinding technique. Events are removed at random in one of two regions of the spectrum of energy deposited in the target counter, corresponding to either $\pi^+ \rightarrow e^+ \nu_e$ or $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events.	62
4.3	The energy deposited in B1 before (upper panel) and after (lower panel) calibration. In the calibrated histogram the MC spectrum is shown in red, and a cut has been made to remove events due to calibration triggers. The difference around 2.5 MeV is due to the requirement of high energy deposit in B1 for physics triggers, which was not included in the MC.	64
4.4	Calibrated energy deposited in T1. The MC spectrum is shown in red.	65
4.5	Energy deposited in BINA + CsI for $\pi^+ \rightarrow e^+ \nu_e$ events. The black histogram, fitted with the black curve, is data, and the red histogram, fitted with the red curve, is MC. The fits are Gaussian; the fitting range is asymmetrical about the peak because the region to the left of the peak is not Gaussian, due to shower leakage.	66
4.6	Energy deposited in BINA + CsI for $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events near the Michel edge. Data is shown in black and MC is shown in red. The lines are fits using Equation 4.1.	66

List of Figures

4.7	The calibrated BINA + CsI spectrum. The spectrum up to 50 MeV is dominated by $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events, and the peak at approximately 65 MeV is due to $\pi^+ \rightarrow e^+ \nu_e$ events. The high-energy tail is mainly due to pileup events, with a small contribution from $\pi^+ \rightarrow e^+ \nu_e \gamma$ events.	67
4.8	The beam spot reconstructed by the WC12 tracker.	68
4.9	Ratios of the x and y momenta to the z momentum (tx and ty), reconstructed by the WC12 tracker.	69
4.10	The beam spot reconstructed by the S12 tracker.	70
4.11	Ratios of the x and y momenta to the z momentum (tx and ty), reconstructed by the S12 tracker. The gaps are due to the track reconstruction algorithm, which uses the centre of the plane hit as the position, leading to some values for tx and ty never occurring.	71
4.12	Decay positron position at the centre (along z) of WC3, reconstructed by the S3WC3 tracker.	72
4.13	Ratios of the x and y momenta to the z momentum (tx and ty) for decay positrons, reconstructed by the S3WC3 tracker.	73
4.14	The run by run variation in the pulse height of the beam positron peak in one BINA PMT. Similar variations were observed for the other PMTs.	74
4.15	The run by run variation in the pion stopping position along the z axis.	75
4.16	COPPER waveform fitted to a template.	76
4.17	Time difference between two T1 PMTs. The red line is a Gaussian fit.	76
4.18	Time of the triggering pulse in the B1 counter.	77
4.19	Reduced χ^2 distribution of the waveform fit in the B1 counter.	77
4.20	Time of the triggering pulse in the T1 counter.	78
4.21	Reduced χ^2 distribution of the waveform fit in the T1 counter.	79
5.1	The time spectra for low energy (upper panel) and high energy (lower panel) events. $t = 0$ is defined by the arrival time of the pion. The repeating peaks are due to beam particles and are separated by the cyclotron RF period (see text).	81
5.2	Number of hits in B1 (left) and B2 (right). The average of four PMTs is taken. The peaks at whole numbers are due to pileup events; the events with extra hits in only some of the tubes are due to noise (see text).	83

List of Figures

5.3	Energy deposited in B1 (left) and B2 (right). The three large peaks in each spectrum are, from left to right, caused by positrons, muons and pions. The smaller peaks are due to events with two particles. The red lines indicate the cut values.	84
5.4	X (left) and Y (right) position at the centre of WC1. The red lines indicate the cut values.	84
5.5	The time spectra after pion selection. The left-hand plot shows the low energy time spectrum and the right-hand plot shows the high energy time spectrum.	85
5.6	High-energy time spectrum without (left) and with (right) the prepileup cut.	86
5.7	The pulse height of the first hit in one of the T1 PMTs divided by the pulse height of the second hit, if one was present. The red line indicates the cut used to select events with fake hits, for plotting the time difference between the initial hit (the real hit) and the fake hit (Figure 5.8).	87
5.8	The time difference between the first and second hits in one of the T1 PMTs, for events with a small second pulse.	87
5.9	The ratio of integrated charge in the T1 PMTs to the fitted pulse height as a function of the fitted pulse height. The red line indicates the cut used to separate real pileup from pileup due to fake hits.	88
5.10	Time spectra after prepileup and T1 pileup cuts. The left-hand plot shows the low energy time spectrum and the right-hand plot shows the high energy time spectrum.	89
5.11	The time difference between the last hit in the T1 VT48 channel and the first hit in the B1 VT48 channel versus the decay time obtained from COPPER.	91
5.12	The distance between the reconstructed positron track and the centre of WC3 (R). The red line indicates the cut value.	92
5.13	Time spectra following all cuts. The left-hand plot shows the low energy time spectrum and the right-hand plot shows the high energy time spectrum. The rise in the high energy spectrum near $t = 0$ at negative times is caused by the integration window of the calorimeter; the closer in time the pileup positron is to the positron from the pion at $t = 0$, the greater the probability that the measured energy in the event will be above E_{cut} .	93

List of Figures

5.14	The time difference between subsequent hits in each T1 PMT; leading times are fitted with an error function. The peak around 30 ns is due to a fake hit at a characteristic time after the real hit.	96
5.15	The shape used in the fit for pileup events that pass the T1 pileup cut due to the double-pulse resolution of the T1 counter.	97
5.16	The shape used in the fit for pileup events where only one positron hit T1.	98
5.17	The shape used in the fit for $\pi \rightarrow \mu\nu\gamma$ events.	99
5.18	The T1res shape when the double-pulse resolution is set to 100 ns.	101
5.19	The fitted T1res amplitude as the double-pulse resolution is increased.	102
5.20	The fit of the high-energy time spectrum with ΔT increased to 100 ns. The red shape is the $\pi^+ \rightarrow e^+\nu_e$ signal, the blue is the $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ background, and the green is the sum of the other backgrounds.	103
5.21	The fitted time spectra from 2012. The left-hand panel shows the low energy time spectrum, fitted with three components: the $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ signal shape, old muon decays, and pion decays-in-flight. The right-hand panel shows the high energy time spectrum, fitted with six components: the $\pi^+ \rightarrow e^+\nu_e$ signal shape, $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ and pion decay-in-flight events promoted to the high energy time spectrum via time independent mechanisms, two mechanisms of old muon pileup, $\pi \rightarrow \mu\nu\gamma$ decays, and old muon decays.	104
5.22	The residuals of the 2010 data set vs. the time of the event. Clockwise from top left, the panels show the residuals for the high energy $t < 0$ spectrum, the high energy $t > 0$ spectrum, the low energy $t > 0$ spectrum, and the low energy $t < 0$ spectrum.	105
5.23	The residuals of the 2011 data set vs. the time of the event. Clockwise from top left, the panels show the residuals for the high energy $t < 0$ spectrum, the high energy $t > 0$ spectrum, the low energy $t > 0$ spectrum, and the low energy $t < 0$ spectrum.	106

List of Figures

5.24	The residuals of the 2012 data set vs. the time of the event, without the T2 pileup cut applied. Clockwise from top left, the panels show the residuals for the high energy $t < 0$ spectrum, the high energy $t > 0$ spectrum, the low energy $t > 0$ spectrum, and the low energy $t < 0$ spectrum.	107
5.25	The residuals of the 2012 data set vs. the time of the event, with the T2 pileup cut applied. Clockwise from top left, the panels show the residuals for the high energy $t < 0$ spectrum, the high energy $t > 0$ spectrum, the low energy $t > 0$ spectrum, and the low energy $t < 0$ spectrum.	108
5.26	The variation of the branching ratio as more pileup events are allowed in the time spectrum, prior to applying the T2 pileup cut. The x axis denotes the time prior to the pion stop in which events with hits in B1, B2, and Tg were rejected. The error bar on the point furthest to the left (with the least pileup) is the error from the time spectrum fit, and the error bars on the other points are the error on the change from the previous point. The points are fitted to a parabola.	111
5.27	The variation of the branching ratio as more pileup events are allowed in the time spectrum, after applying the T2 pileup cut. The x axis denotes the time prior to the pion stop in which events with hits in B1, B2, and Tg were rejected. The error bars represent the statistical variation from the point furthest to the left (with the least pileup).	112
5.28	The fitted branching ratio for groups of 1000 runs, fitted with a flat line.	113
6.1	The total energy deposited in B1, B2, S1, S2, and Tg by $\pi^+ \rightarrow e^+\nu_e$ (black), $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ (red), π DIF- μ DAR (green), and π DAR- μ DIF (blue). The distributions are normalized to the same height. The solid red lines indicate the selected region.	117
6.2	The measured angle between the tracks reconstructed by WC12 and S12 for $\pi^+ \rightarrow e^+\nu_e$ (black) and $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ (red) events after the time and energy loss cuts are applied. The double peak in the $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ distribution is caused by pion decays-in-flight. The solid red line and arrow indicate the selected region.	118
6.3	The measured BINA+CsI spectra as the suppression cuts are applied. The legend indicates the fraction of low-energy events (< 52 MeV) remaining after each cut.	119

List of Figures

6.4	Comparison between measured (filled circles with error bars) and simulated energy spectra. The simulation was performed with (red) and without (blue) hadronic reaction contributions. The histograms are normalized to the same area [9]. . .	123
6.5	Beam particle energy measured by the crystals at 0° with no cuts applied.	125
6.6	Energy measured by the crystals at 0° following event selection cuts (see text).	127
6.7	Energy measured by the crystals at 0° in the RF window corresponding to muons.	127
6.8	Energy measured by the crystals at 0° . The red lines indicate the region defined as the muon peak.	128
6.9	Reconstructed track parameters for the positron beam at 0° . The top-left and top-right panels show the ratio of the x and y momenta to the z momentum, and the bottom-left and bottom-right panels show the reconstructed x and y positions at $z = 0$	135
6.10	Mean of the distribution of the x momentum, normalized to the z momentum, as a function of angle.	135
6.11	The energy spectrum from a 70 MeV positron beam parallel to the crystal axis. Data is shown in black and simulation is shown in red. The histograms are normalized to have the same total number of events. The green line shows the value of E_{cut}	136
6.12	The energy spectrum from a 70 MeV positron beam at 11.8° (top left), 24.4° (top right), 36.2° (bottom left), and 47.7° (bottom right) to the crystal axis. Data is shown in black and simulation is shown in red.	137
6.13	The tail fraction as a function of angle in the positron beam data.	139
6.14	The difference between the tail fraction in the positron beam data and the tail fraction from a simulated positron beam as a function of angle.	140
6.15	The simulated BINA+CsI spectrum from $\pi^+ \rightarrow e^+\nu_e$ decay and $\pi^+ \rightarrow e^+\nu_e\gamma$ decay.	141
6.16	The simulated BINA+CsI spectrum from $\pi^+ \rightarrow e^+\nu_e$ decay and $\pi^+ \rightarrow e^+\nu_e\gamma$ decay, excluding events that underwent Bhabha scattering in the target.	141

List of Figures

6.17	The tail fraction as a function of angle for the positron beam data (left) and Monte Carlo (right), fitted to a fourth-degree polynomial.	142
6.18	Simulated energy spectrum for $\pi^+ \rightarrow e^+\nu_e$ events emitted at small angles.	143
6.19	The tail fraction in the positron beam data minus the simulated tail fraction as a function of angle. The first 8 angles are fitted to a straight line.	144
6.20	The probability distribution obtained by combining the upper and lower limits.	145
7.1	The ratio of $\pi^+ \rightarrow e^+\nu_e$ events to $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ events within different radii of the centre of WC3, as reconstructed by the S3WC3 tracker.	150
7.2	The z coordinate of the reconstructed pion stopping position. Data is shown in black and MC is shown in red.	151
7.3	A comparison of the actual and reconstructed pion stopping positions in MC.	151
7.4	Variation of the acceptance correction with the peak value of the reconstructed pion stopping position. The peak was varied by ± 0.2 mm; the largest variation in the correction was approximately $\pm 0.05\%$, for $R < 90$ mm.	152
7.5	Variation of the acceptance correction with the width of the reconstructed pion stopping position.	152
7.6	Variation of the acceptance correction with the thresholds in the T1 and T2 counters.	153
7.7	The decay time of muons in the target with non-zero kinetic energy at the time of the decay.	154
7.8	Simulated energy spectra measured by BINA+CsI for muon decays-at-rest (black) and decays-in-flight (red).	155
7.9	Variation of the branching ratio as the radius in which events are accepted is varied. The left-hand panel shows the results with the T2 pileup cut, and the right-hand panel shows the results without the T2 pileup cut. The red points show the raw branching ratio, and the black points show the branching ratio after all corrections. The error bars represent the error on the change from the first point.	156

List of Figures

7.10	The branching ratio for statistically independent rings in R, the distance between the reconstructed positron track and the centre of WC3. The left-hand panel shows the results with the T2 pileup cut, and the right-hand panel shows the results without the T2 pileup cut. The red points show the raw branching ratio, and the black points show the branching ratio after all corrections. The value along the x axis is the centre of the ring under consideration; that is, the point at $x = 35$ mm is the branching ratio for events with R between 30 and 40 mm.	157
7.11	The corrected branching ratio for statistically independent rings in R, with (left) and without (right) the T2 pileup cut. This plot is identical to Figure 7.10, but zoomed in on the corrected points.	158
7.12	The change in the branching ratio from the value at $E_{cut} = 52$ MeV.	159
A.1	PIENU Trigger Diagram	172
B.1	A simplified timing diagram for a PIMUE event in which the muon decayed 400 ns after the pion stop.	174
B.2	A simplified timing diagram for an event in which an old muon decays and 200 ns later a pion arrives. Due to the delay between the actual hits in T1 and T2 and the downstream coincidence signal, the event still triggers.	175
B.3	A simplified timing diagram for an event with two muons, in which the positron from the old muon decay completes the trigger.	176
B.4	A simplified timing diagram for an event with two muons, in which the positron from the muon from the decay of the primary pion completes the trigger.	177
B.5	A simplified timing diagram for an event with two muons, in which both decay positrons enter the acceptance, but the event passes the T1 pileup cut because the decays are too close together in time for the separate hits to be resolved. . .	178
C.1	The beam spot in WC1 and WC2 for positron beam data. . .	179
C.2	Energy measured by the crystals following the selection of beam particles.	180

List of Figures

C.3	The time distribution of the first hit in the first plane of each wire chamber.	181
C.4	Energy measured by the crystals at 0° following the removal of events with out-of-time hits.	182
C.5	Time of flight vs. T2 energy.	182
C.6	Energy measured by the crystals at 0° for events with time-of-flight corresponding to positrons (left) and muons (right). .	183
C.7	Energy measured by the crystals at 0° following the removal of events with high energy deposit in T2.	184
C.8	Simulated energy deposit in T2, with (red) and without (black) BINA in place.	184

Glossary

Each entry is followed by (G) if it is a term in general use or (E) if it is specific to the experiment.

Acceptance correction (E)

A multiplicative factor that must be applied to the branching ratio obtained from the time spectrum fit to take into account the difference in geometrical acceptance between the two decay modes.

ADC (G)

Analog-to-digital converter.

B1 (E)

B1 and B2 are plastic scintillators in the PIENU detector.

Bhabha scattering (G)

Electron-positron scattering.

BINA (E)

The NaI(Tl) crystal at the centre of the PIENU calorimeter.

BinaHigh (E)

A trigger used in the PIENU experiment to record events with large energy deposit in BINA and CsI.

Blinding

(G) A procedure whereby the result of an analysis is changed by an unknown random factor, to avoid human bias.

Branching ratio (G)

The ratio of the rate of a specific decay mode to the full decay rate.

Bremsstrahlung (G)

A process by which charged particles emit photons when accelerating.

ChPT (G)

Chiral perturbation theory. A low-energy effective field theory for QCD.

COPPER (E)

A 500 MHz ADC used to digitize the waveforms from some elements of the PIENU detector.

CsI (G)

Cesium iodide, a commonly-used type of crystal scintillator that forms part of the PIENU calorimeter.

Cyclotron (G)

A type of circular particle accelerator with constant magnetic field and variable orbital radius.

ΔT (E)

The time interval below which multiple hits cannot be resolved in T1. A parameter in the fitting function for the high energy time spectrum.

Early (E)

A trigger used in the PIENU experiment to record events shortly after t_0 , the pion arrival time.

E_{cut} (E)

The threshold used to separate low energy and high energy events. Typically set at 52 MeV.

Electroweak (G)

Refers to the combined description of the electromagnetic and weak interactions.

Flavour (G)

Particle species. In the Standard Model, there are three flavours of quarks and three flavours of leptons; for example, the electron and electron neutrino both have electron flavour.

FPGA (G)

Field-programmable gate array. A programmable chip that performs logic operations.

GEANT4 (G)

A software package common in high-energy physics for simulating particles and their interactions with matter.

Generation (G)

In the Standard Model, there are three generations of fermions, each consisting of two quarks, a charged lepton, and a neutrino. The particles in the second and third generations differ from their counterparts in the first generation only by mass. For the quarks and charged leptons, each generation is heavier than the one before; the mass ordering of the neutrinos is unknown.

Kink (E)

The angle between the track reconstructed using WC1 and WC2 and the track reconstructed using S1 and S2. Used to identify π -DIF events.

Lepton universality (G)

The assumption in the Standard Model that the electroweak couplings of the three lepton generations are the same.

Lower limit (E)

The lower limit on the tail correction; obtained from the energy spectrum for pion decay events.

M13 (E)

A secondary beamline at TRIUMF, used to deliver particles to the PIENU experiment.

Michel decay (G)

The decay of a muon to an electron and two neutrinos.

Michel edge (G)

The endpoint of the energy spectrum of electrons produced via muon decay.

MIDAS (E)

A web-based DAQ interface developed at TRIUMF and PSI.

Muon correction (E)

A correction applied to the positron beam data to correct for the presence of muons.

Mu_pie (E)

A term in the fitting function for the high energy time spectrum. Describes events due to old muon decay, with energy added by a mechanism whose timing is independent of t_0 .

Mu_pimu (E)

A term in the fitting function for the low energy time spectrum. Describes events due to old muon decay.

μ -DAR (E)

Muon decay-at-rest.

μ -DIF (E)

Muon decay-in-flight.

μ -DIF correction (E)

A multiplicative factor that must be applied to the branching ratio obtained from the time spectrum fit to take into account muon decay-in-flight events, which have the same time dependence as $\pi \rightarrow e\nu$ events and can have measured energy above E_{cut} .

NaI(Tl) (G)

Thallium-doped sodium iodide, a commonly-used type of crystal scintillator.

NMR (G)

Nuclear magnetic resonance. A technique for measuring magnetic fields based on excitation and relaxation of nuclear spins in a known sample.

oldmu_both (E)

A component in the fitting function for the high energy time spectrum. It describes events where two Michel positrons enter the crystal array but only passes through T1.

Old muon (E)

A muon present in Tg prior to the arrival of the primary pion.

π -DAR (E)

Pion decay-at-rest.

π -DIF (E)

Pion decay-in-flight.

Pion data (E)

Data taken in the usual mode, with a 75 MeV/c pion beam entering the detector.

Positron beam data (E)

Data taken with a positron beam, instead of the usual pion beam. Used to obtain the response function of the crystal array to 70 MeV positrons.

Prescale (E)

One of the triggers used in the PIENU experiment. Also a generic term for storing only one event out of a specified number that would normally trigger, typically to reduce data size. In the PIENU experiment, the prescale factor is sixteen, meaning one event out of sixteen is stored.

proot (E)

Software written for the PIENU experiment to convert raw data into ROOT trees.

PSI (G)

Paul Scherrer Institute. A research facility in Villigen, Switzerland, with a cyclotron similar to the TRIUMF cyclotron.

QCD (G)

Quantum chromodynamics. The theory that describes the strong interaction.

r (E)

A parameter in the fitting functions for the time spectra. Describes the proportion of low energy events promoted to the high energy time spectrum via mechanisms whose timing is independent of t_0 .

R (E)

The distance between the centre of WC3 and the reconstructed positron track.

R_π (E)

The pion branching ratio.

Reduced χ^2

The χ^2 of a fit divided by the number of degrees of freedom of the fit.

RF (G)

Radio-frequency. In this context, refers to the accelerating electric field used in the TRIUMF cyclotron.

ROOT (G)

Object-oriented data analysis software commonly used in high-energy physics.

Scintillator (G)

A material that emits photons of a characteristic wavelength when charged particles pass through it. Common types include crystal, organic, and liquid noble gas scintillators. The PIENU detector includes crystal and plastic scintillators.

Silicon strip (G)

A position-sensitive particle detector, consisting of segmented pieces of silicon with a bias voltage applied.

S1 (E)

S1, S2, and S3 are the three silicon strip detectors used in the PIENU apparatus.

Standard Model (G)

A theoretical description of fundamental particles and their interactions.

t_0 (E)

The pion stop time.

t_0 correction (E)

A multiplicative factor that must be applied to the branching ratio obtained from the time spectrum fit to take into account the dependence of the measured value of the decay time on the decay positron energy.

T1 (E)

T1 and T2 are plastic scintillators in the PIENU detector.

T1res (E)

A component in the fitting function for the high energy time spectrum. It describes events where two Michel positrons pass through T1 into the crystal array sufficiently close together in time that only a single hit in T1 is recorded.

Tail correction (E)

A multiplicative factor that must be applied to the branching ratio obtained from the time spectrum fit to take into account $\pi \rightarrow e\nu$ events whose measured energy was less than E_{cut} .

TDC (G)

Time-to-digital converter (records the time at which a signal went above a given threshold).

Tg (E)

A plastic scintillator in the PIENU detector that functions as the pion stopping target.

TIGC (E)

A VME module used in the BinaHigh trigger from 2010 onwards for summing the pulse height of BINA and CsI, and issuing a trigger if it exceeds a defined threshold.

TOF (E)

Time of flight.

Tpos (E)

The time of the positron hit relative to t_0 . The quantity that is fitted to obtain the branching ratio.

Trigger (G)

A digital logic circuit that takes detector signals as inputs, and sends a signal to the data acquisition system if the event should be stored.

TRIUMF (G)

Canada's national laboratory for particle and nuclear physics. The site of the PIENU experiment.

tx (E)

tx and ty are parameters in the PIENU track reconstruction algorithm. They are the ratios of the x and y momenta to the z momentum.

Upper limit (E)

The upper limit on the tail correction; obtained from positron beam data.

VF48 (E)

A 60 MHz ADC used to record waveforms from some elements of the PIENU detector.

VT48 (E)

A 1.6 GHz TDC used to record signals from some elements of the PIENU detector and trigger.

Wire chamber (G)

A position-sensitive particle detector, consisting of wires in a gas-filled chamber with high voltage applied across it.

WC1 (E)

WC1, WC2, and WC3 are the three wire chambers used in the PIENU detector.

x0 (E)

x0 and y0 are parameters in the PIENU track reconstruction algorithm. They are the x and y positions at $z = 0$.

χ^2 (G)

A method for fitting data to a function based on the squared difference between the data and the function value, divided by the error on the data; χ^2 also refers to the goodness-of-fit parameter in this method.

Acknowledgements

I would like to thank the entire PIENU collaboration for their various contributions to the experimental effort, which are too many to list exhaustively. Thank you to the current and former graduate students, Chloé Malbrunot, Shintaro Ito, Saul Cuen-Rochin, and Dorothea vom Bruch for all the help, discussion, and shared confusion along the way. In particular I would like to thank Chloé Malbrunot for her analysis of the 2010 data, which formed the basis for the analysis of the 2011 data set, performed by Shintaro Ito, and the 2012 data set, the main subject of this thesis. Thank you to the postdocs, Aleksey Sher, Luca Doria, and Dima Vavilov, for their essential contributions to every aspect of the experiment. Thank you to Toshio Numao for his patience and willingness to explain the many things about the experiment I found bewildering over the years. Thank you to Dick Mischke for his thoroughness in checking not only the analysis itself, but also everything written about it. And my biggest thanks go to Doug Bryman, whom I was extremely fortunate to have as my supervisor.

Lastly, I am extremely grateful to Abbra Latta, for all her invaluable help and support.

To my family.

Chapter 1

Introduction

1.1 The Standard Model

The Standard Model (SM) of particle physics is a theory describing the characteristics and interactions of the most fundamental known constituents of matter [10]. The particle content of the SM is divided into fermions, half-integer spin particles, and bosons, integer spin particles. Two forces are contained in the SM, the strong force and the electroweak force. These forces occur through the exchange of spin one particles, called gauge bosons: gluons in the case of the strong force, and photons, W^\pm , and Z for the electroweak force. Gluons and photons are massless, while W^\pm and Z have masses of 80.385 ± 0.0015 GeV and 91.1876 ± 0.0021 GeV, respectively [11]. Below the energy scale of the W^\pm and Z masses the electroweak force appears to separate into two forces, electromagnetism and the weak force.

Fermions in the SM are further divided into two subcategories, quarks and leptons; the former feel the strong force in addition to the electroweak force, while the latter do not. There are three generations of fermions, which have identical properties aside from mass, with particles from higher generations being more massive than the corresponding particles from lower generations. The particle content of the SM is shown in Table 1.1; each column of fermions represents one generation, consisting of two quarks, a charged lepton, and an associated neutral lepton, called a neutrino.

Table 1.1: The particle content of the SM; e is the charge of the electron.

	Fermions			Charge ($ e $)	Bosons	Charge ($ e $)
Quarks	u	c	t	+2/3	γ	0
	d	s	b	-1/3	W^\pm	± 1
Charged leptons	e	μ	τ	-1	Z	0
Neutrinos	ν_e	ν_μ	ν_τ	0	H	0

1.1. The Standard Model

In the SM, these particles are viewed as quantized excitations of fields, in the mathematical sense of a function defined everywhere in spacetime. One more field exists, called the Higgs field; all massive particles in the SM acquire their masses through interactions with the Higgs field. Although this mechanism was proposed as early as 1964 [12], the associated particle was only discovered in 2012, by the ATLAS and CMS experiments at the Large Hadron Collider [13] [14].

1.1.1 The weak interaction

The weak interaction is unique in the SM in that it does not conserve particle type, or flavour. Flavour refers to the generation to which a given fermion belongs. In the strong and electromagnetic interactions, the number of quarks and leptons of each flavour must be the same in the initial state as it is in the final state; in the weak interaction, this is not the case. Without this property, many decays that are allowed in the SM would be forbidden, including that of the pion.

The pion is the lightest particle made of quarks, consisting of a quark anti-quark pair of the first generation. Pions decay either into electrons or muons, which are the second-generation charged lepton, plus the associated neutrinos. Pion decay in the SM proceeds via the weak interaction. Two types of weak interactions exist: charged-current, which occur via W^\pm exchange, and neutral-current, which occur via Z exchange; pion decay is of the former type.

The strength of interactions in the SM is set by the coupling constants α , for the electromagnetic interaction, α_s , for the strong interaction, and g , for the weak interaction (see Appendix A of Reference [10]). The values of these coupling constants are shown in Table 1.2. In fact, these values depend on the energy of the interacting particles; the quoted values apply at low energies, on the order of 100 MeV (the strong coupling, in particular, changes rapidly with energy) [10].

It may seem surprising, given its name, that the coupling for the weak force is so much greater than the coupling for the electromagnetic force. The reason is the mass of the mediating particles, which reduces the coupling at low energies. The “effective” coupling for the weak interaction, the Fermi constant G_F , is related to the “intrinsic” coupling g by $G_F = \frac{g^2}{4\sqrt{2}M_W^2}$.

Table 1.2: Coupling constants for SM interactions.

Interaction	Symbol	Value (at ~ 100 MeV)
Strong	α_s	1.7
Electromagnetic	α	0.0073
Weak	g	0.652

Another unique feature of the weak interaction is the violation of parity symmetry, which is the symmetry of physical laws under a reflection of spatial coordinates. Fermions in the SM are represented by four-component objects called spinors; these can be either left-handed or right-handed, depending on their transformation properties. In technical terms, the two types are different representations of the Poincaré algebra (see Appendix C of Reference [10]). Although the strong and electromagnetic interactions treat left- and right-handed fermions identically, the weak interaction only affects left-handed fermions and right-handed anti-fermions, which leads to parity violation. As will be explained in Chapter 2, this property is responsible for the fact that pions decay into muons overwhelmingly more often than they decay into electrons.

1.2 Motivation for the PIENU experiment

The PIENU experiment aims to make a precise measurement of the decay properties of the pion. The pion branching ratio, R_π , is defined as the rate at which pions decay into electrons and neutrinos divided by the rate at which they decay into muons and neutrinos. Also included are the associated radiative modes, in which a photon is emitted as well:

$$R_\pi = \frac{\Gamma(\pi^+ \rightarrow e^+ \nu_e + \pi^+ \rightarrow e^+ \nu_e \gamma)}{\Gamma(\pi^+ \rightarrow \mu^+ \nu_\mu + \pi^+ \rightarrow \mu^+ \nu_\mu \gamma)}. \quad (1.1)$$

R_π has been calculated within the SM to be $(1.2352 \pm 0.0002) \times 10^{-4}$ [1], an unusual level of precision for a quantity involving quarks. Although strong interaction effects influence the individual decay rates, to first order they cancel in the ratio, and also largely cancel in the higher-order corrections [15]. The small theoretical uncertainty, along with the ability to accurately measure R_π , makes pion decay a sensitive test of the structure of the weak

interaction and the manner in which it connects quarks and leptons, and a powerful tool in the search for new physics.

Many new physics scenarios, such as lepton non-universality [2], R-parity violating supersymmetry [16], leptoquarks [17], and heavy neutrinos [18], have an impact on R_π . In general, new physics at the weak scale is expected to impact R_π on the level of 0.01-1% [1]. By reducing the experimental uncertainty to the level of the theoretical uncertainty, stringent limits can be placed on these theories. The branching ratio calculation, and the potential physics reach, will be discussed in more detail in Chapter 2.

1.3 Brief history of pion branching ratio measurements

The pion was discovered by Cecil Powell, Giuseppe Occhialini, Hugh Muirhead, and Cesare Lattes in 1947 [19]. For the first decade after their discovery, the only known decay mode of charged pions was to muons, $\pi \rightarrow \mu\nu_\mu$; it was unknown at this time what prevented the electron mode.

In 1958, Sudarshan and Marshak [20] proposed that the weak interaction operator has the form of a vector minus an axial vector (V-A), as compared with the pure vector operator of quantum electrodynamics (QED). The resulting theory explained almost all experimental observations of weak processes at the time, with one of the few exceptions being the non-observation of the electronic decay mode of the pion: $\pi \rightarrow e\nu_e$ [21]. This process was predicted by the V-A theory to occur at a rate suppressed by a factor on the order of 10^{-4} relative to the muon mode. At the time, the experimental upper limit on the decay rate was on the level of 2×10^{-5} ; however, the experiments setting these limits proved to be in error, and the decay was finally observed at CERN in 1958 [22]. Within two years, the branching ratio

$$R_\pi = \frac{\pi^+ \rightarrow e^+\nu_e}{\pi^+ \rightarrow \mu^+\nu_\mu} \quad (1.2)$$

was measured in agreement with the predicted value at the 5% level [23], representing another success of the V-A theory. The branching ratio including radiative modes was measured in 1964 to a precision of approximately 2% [24]. The result, $R_\pi = (1.247 \pm 0.028) \times 10^{-4}$, was within 0.5 σ of the theoretical prediction at the time; although it was later revised to

1.4. Experimental technique

$R_\pi = (1.274 \pm 0.024) \times 10^{-4}$ due to a more accurate determination of the pion lifetime, it remained within 2σ [25]. The next measurement was done at TRIUMF in 1983, giving $R_\pi = (1.218 \pm 0.014) \times 10^{-4}$ [26].

The most recent experiments were done at TRIUMF [27] and PSI [28] in the early 1990s. The Particle Data Group average of the last three experiments is $R_\pi = (1.230 \pm 0.004) \times 10^{-4}$, a precision of 0.33% [29]. The latest result, $R_\pi = (1.2344 \pm 0.0023 \text{ (stat)} \pm 0.0019 \text{ (syst)}) \times 10^{-4}$, achieving 0.24% precision, was obtained from a subset of the PIENU data [3].

Despite the precision obtained by the branching ratio measurements, the theoretical uncertainty is smaller by a factor of 15. Nevertheless, the agreement between the theoretical and experimental values of R_π provides the best constraint on the ratio of the electroweak couplings of the electron and the muon, g_e/g_μ (see Section 2.3). The PIENU experiment aims to measure R_π to a precision of $< 0.1\%$, corresponding to a 0.05% measurement of g_e/g_μ . Figure 1.1 shows the uncertainty on the branching ratio achieved by the previous measurements, along with the first PIENU result. Also shown is the projected final result from combining the results of PIENU and PEN [30], the latest PSI experiment. PEN is also attempting to make a $< 0.1\%$ measurement of R_π .

1.4 Experimental technique

Every branching ratio measurement since that by Di Capua in 1964 [24] has used the same fundamental technique: stopping a pion beam in a scintillator target and detecting the decay positrons, which result either from the rare $\pi^+ \rightarrow e^+\nu_e$ process or the far more common $\pi \rightarrow \mu \rightarrow e$ decay chain, in a calorimeter. $\pi^+ \rightarrow e^+\nu_e$ decay is two-body, and produces a positron with a fixed energy of 69.8 MeV. Muon decay, $\mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$, is three-body, and the positron's energy is distributed between 0.511 MeV (its rest energy) and 52.8 MeV (half the muon mass). The energy distribution of positrons from muon decay is referred to as the Michel spectrum, and the endpoint as the Michel edge. The lifetime of the muon, $\tau_\mu = 2.197 \mu\text{s}$, is much longer than the lifetime of the pion, $\tau_\pi = 26.0 \text{ ns}$.

The differing timing and energy distributions in the $\pi^+ \rightarrow e^+\nu_e$ case and the $\pi \rightarrow \mu \rightarrow e$ are used to distinguish the decay modes. The energy of the positrons, and photons if they are produced by either the decaying

1.4. Experimental technique

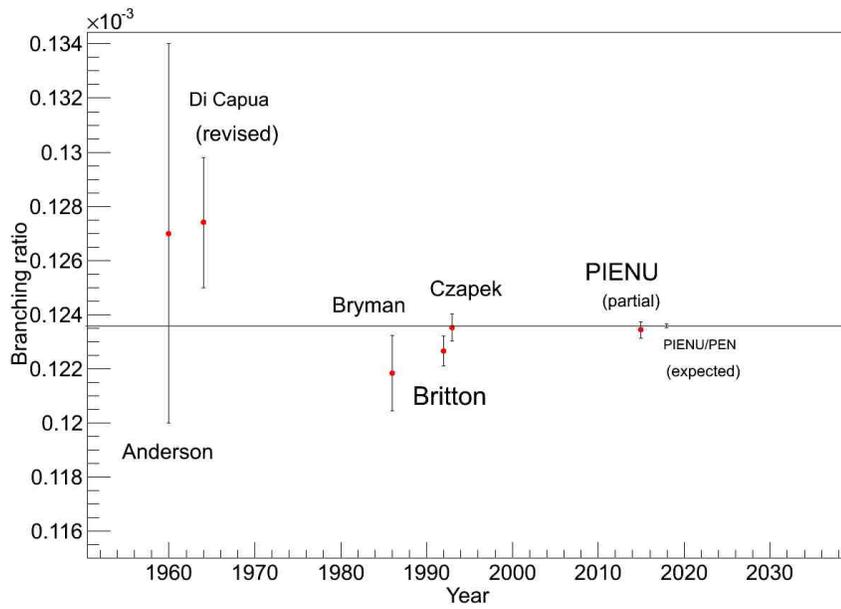


Figure 1.1: The measured value of the branching ratio as a function of time. The dashed line shows the SM prediction. The last point indicates the expected level of precision that will ultimately be achieved by combining the results of the PIENU experiment and the PEN experiment at PSI, as described in the text.

1.4. Experimental technique

pion or muon, is measured using a crystal scintillator calorimeter. The two processes, and the energies and lifetimes of the particles involved, are shown in Figure 1.2.

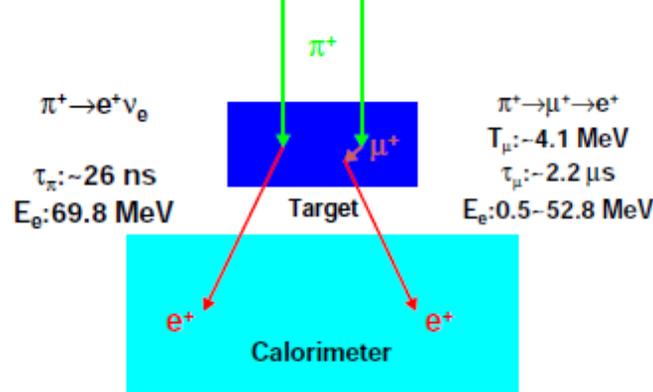


Figure 1.2: Simplified picture of the PIENU experimental technique. Pions stop in the target and decay into either muons or positrons; muons also stop in the target, and decay into positrons. Positrons, and photons if any are produced, are detected by the crystal scintillator calorimeter.

In the PIENU experiment, the branching ratio is obtained by separating the events into high- and low-energy regions, using a threshold just above the Michel edge, and fitting the time spectra in the two regions. These are referred to as the high- and low-energy time spectra. The fitting function used in the PIENU analysis is fully described in Chapter 5.

The timing distribution of positrons arising from $\pi^+ \rightarrow e^+ \nu_e$ decay is simply an exponential with the pion lifetime. The timing distribution of the positrons from $\pi \rightarrow \mu \rightarrow e$ is determined by the decay time of the pion plus the decay time of the muon, resulting in a more complicated distribution that rises as the pions decay and then falls off with the muon lifetime, $\tau_\mu = 2.197 \mu\text{s}$ [29]. Monte Carlo simulations of the time and energy spectra in the two cases, excluding radiative modes, are shown in Figures 1.3 and 1.4 (for ease of visualization, the decay modes are not shown in the correct proportion).

1.4. Experimental technique

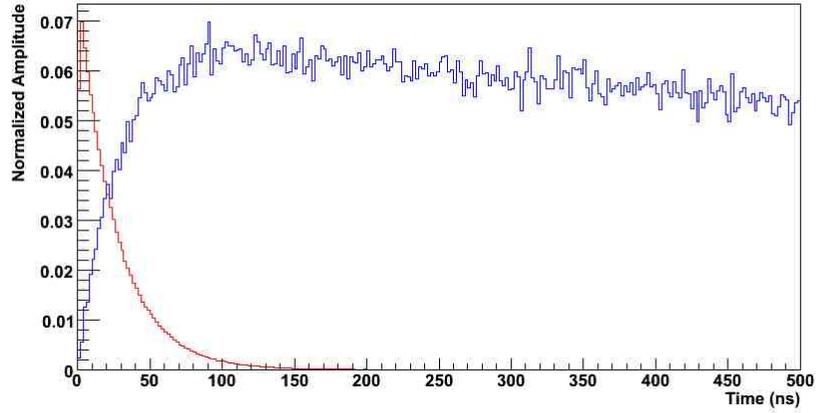


Figure 1.3: GEANT4 [4] simulation of the time spectra of $\pi^+ \rightarrow e^+ \nu_e$ (red) and $\pi \rightarrow \mu \rightarrow e$ (blue) events. The distributions are normalized to the same height.

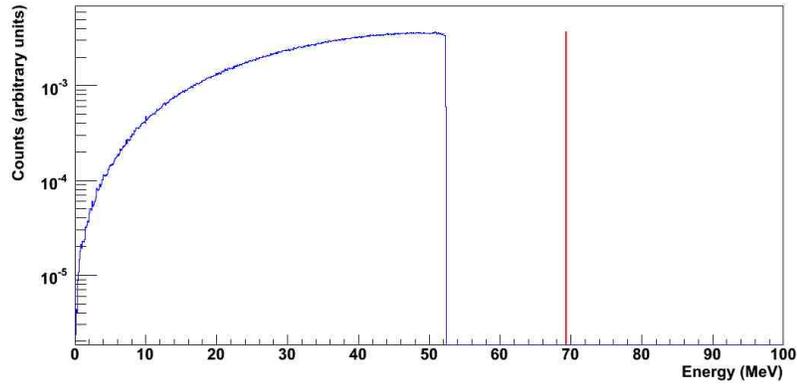


Figure 1.4: GEANT4 simulation of the starting energies of positrons due to $\pi^+ \rightarrow e^+ \nu_e$ (red) and $\pi \rightarrow \mu \rightarrow e$ (blue) decays. The distributions are normalized to the same height.

1.4. Experimental technique

Either pion decay mode, or the decay of the muon, can also result in a photon being emitted. This does not affect the time spectra, but can affect the measured energy. The branching ratio can be affected in two ways. If the pion decays to a muon, a photon emitted during the $\pi \rightarrow \mu \rightarrow e$ decay chain can add sufficient energy for the event to be placed in the high-energy time spectrum. Depending on the timing of the event, it could be misidentified as a $\pi^+ \rightarrow e^+ \nu_e$ decay. On the other hand, if the pion decays to a positron and a photon, the photon can carry away sufficient energy that the event is placed in the low-energy time spectrum, if the photon is not absorbed in the calorimeter. Both of these possibilities are taken into account via Monte Carlo simulation, as described in Chapters 5 and 6.

Measuring the ratio of the decay rates, rather than simply the $\pi^+ \rightarrow e^+ \nu_e$ yield, provides several benefits for reducing the systematic uncertainty. The total number of incoming pions does not need to be known; since positrons are measured regardless of the decay mode of the pion, the efficiencies of the cuts and triggers cancel in the measured ratio of decay rates, to first order; and the geometrical acceptance enters only due to energy-dependent scattering, which results in a small correction to the measured branching ratio. Thus, many sources of systematic error¹ are either reduced or eliminated entirely.

1.4.1 PIENU experimental technique

The PIENU detector (fully described in Chapter 3 and Reference [8]) is shown in Figure 1.5. Wire chambers and silicon strip detectors provided particle tracking; thin plastic scintillators (including the pion stopping target) provided timing information, particle ID, and were used in the trigger; and crystal scintillators provided the energy measurement.

A pion beam with momentum $75 \text{ MeV}/c \pm 1 \text{ MeV}/c$ [31] from the TRIUMF M13 beamline was injected into the detector. The pions were detected by the plastic scintillators B1 and B2, and tracked into the target (Tg) by a pair of wire chambers and a pair of double-sided silicon strip detectors. Decay positrons were detected by two more plastic scintillators, T1 and T2, and tracked by another silicon strip detector and wire chamber, before entering the crystal scintillators.

¹In this thesis, “error” and “uncertainty” are used interchangeably.

1.4. Experimental technique

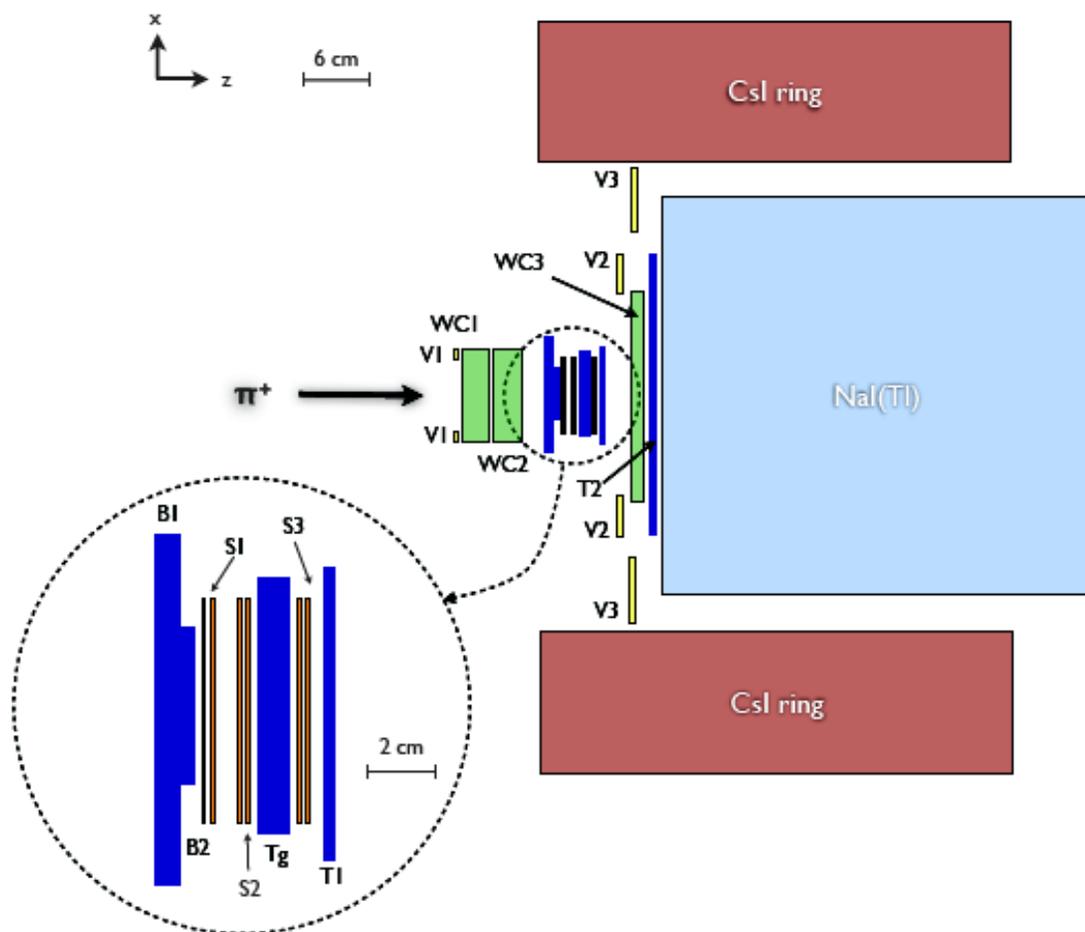


Figure 1.5: The PIENU detector, with the pion beam incident from the left. The region close to the target is shown in the blowup. Plastic scintillators (polyvinyltoluene) are shown in dark blue, wire chambers in green, silicon strip detectors in orange, and crystal scintillators in light blue (NaI(Tl)) and red (CsI).

The central element of the crystal calorimeter array was BINA, a single cylindrical NaI(Tl) crystal borrowed from Brookhaven National Laboratory, 48 cm long by 48 cm in diameter, which was read out by 19 photomultiplier tubes (PMTs). BINA was surrounded by two concentric rings of pure CsI crystals to capture shower leakage. There were a total of 97 individual CsI crystals, which were pentagonal in cross section, 25 cm long, 9 cm wide, and 8 cm high. Each crystal was read out by a single PMT.

The precision of the branching ratio measurement ultimately rests on the extent to which $\pi^+ \rightarrow e^+\nu_e$ events can be separated from background events. Figures 1.3 and 1.4 show that backgrounds with large energy deposit at early time, relative to the stopping time of the pion, have the potential to be misidentified as $\pi^+ \rightarrow e^+\nu_e$ events. The three main sources of background in the PIENU experiment are beam-related background, pileup of muons from previous pion decays, and pileup of neutral particles. These will be discussed in section 3.4.

1.5 Outline of the thesis

In Chapter 2, the theory of pion decay and the sensitivity of the branching ratio to various beyond the Standard Model scenarios are briefly presented. Chapter 3 gives an overview of the experimental apparatus, including the M13 beamline. Chapter 4 describes the data-taking conditions, and describes the data processing. Chapter 5 describes the event selection and the timing fit done to obtain the raw branching ratio. Chapters 6 and 7 describe the corrections that must be applied to the raw branching ratio to obtain the final result, and conclusions are presented in Chapter 8.

Chapter 2

Theory

2.1 Pions in the Standard Model

2.1.1 Pion decay modes

Table 2.1 shows the decay modes of the charged pion listed by the Particle Data Group [11]. The branching ratios quoted are experimental values; the values for the radiative modes apply to a restricted kinematic range for the photon, as indicated in the table.

Table 2.1: Pion decay modes and branching ratios.

Mode	Branching ratio	Notes
$\mu^+\nu_\mu$	0.9998770 ± 0.0000004	
$\mu^+\nu_\mu\gamma$	$(2.00 \pm 0.25) \times 10^{-4}$	$E_\gamma > 1 \text{ MeV}$
$e^+\nu_e$	$(1.230 \pm 0.004) \times 10^{-4}$	
$e^+\nu_e\gamma$	$(7.39 \pm 0.05) \times 10^{-7}$	$E_\gamma > 10 \text{ MeV}, \theta_{e\gamma} > 40^\circ$
$e^+\nu_e\pi^0$	$(1.036 \pm 0.006) \times 10^{-8}$	
$e^+\nu_e e^+e^-$	$(3.2 \pm 0.5) \times 10^{-9}$	
$e^+\nu_e\nu\bar{\nu}$	$< 5 \times 10^{-6}$	

Only the first four decay modes, $\pi^+ \rightarrow \mu^+\nu_\mu(\gamma)$ and $\pi^+ \rightarrow e^+\nu_e(\gamma)$, are relevant to the PIENU experiment. The modes with the π^0 and the extra e^+e^- pair in the final state occur $< 0.01\%$ as often as the $\pi^+ \rightarrow e^+\nu_e$ mode, which is negligible at the level of precision of the PIENU experiment.

2.1.2 The weak interaction

Pion decay to leptons in the SM proceeds via the weak interaction. Since the pion is the lightest hadronic particle, the strong and electromagnetic

2.1. Pions in the Standard Model

decay modes are forbidden, as these interactions conserve flavour. The Lagrangian for the part of the interaction mediated by the W^+ is given by [10].

$$\mathcal{L}_{W^+} = \frac{ig}{2\sqrt{2}} W_\mu^+ (\bar{\nu}_m \gamma^\mu (1 + \gamma^5) e_m + \bar{u}_m \gamma^\mu (1 + \gamma^5) d_m). \quad (2.1)$$

Here there is an implied sum over the index m , which goes from one to three and indicates the generation of the particle involved; that is, $e_1 = e$, $e_2 = \mu$, and $e_3 = \tau$. γ^μ , where μ goes from zero to three, are the Dirac gamma matrices; γ^5 is given by $i\gamma^0\gamma^1\gamma^2\gamma^3$.

This expression is valid in the ‘‘interaction’’ basis, that is, for particles which carry definite flavour; however, in the case of the quarks, this is not the same as the mass basis, which is the physical one. The mass basis and the interaction basis are connected via the Cabbibo-Kobayashi-Maskawa (CKM) matrix, in the following manner:

$$\begin{bmatrix} d' \\ s' \\ b' \end{bmatrix} = \begin{bmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{bmatrix} \begin{bmatrix} d \\ s \\ b \end{bmatrix}. \quad (2.2)$$

Here the primes indicate particles in the interaction basis, while the unprimed vector represents particles in the mass basis. The Lagrangian then becomes

$$\mathcal{L}_{W^+} = \frac{ig}{2\sqrt{2}} W_\mu^+ (\bar{\nu}_m \gamma^\mu (1 + \gamma^5) e_m + V_{mn} \bar{u}'_m \gamma^\mu (1 + \gamma^5) d'_n) \quad (2.3)$$

The term $1 + \gamma^5$ is responsible for the parity-violating nature of the weak interaction. This is most easily seen by writing γ^5 and the spinor representing a fermion in the Weyl, or chiral, basis [32, p. 43-50]:

$$\gamma^5 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}, \psi = \begin{bmatrix} \psi_L \\ \psi_R \end{bmatrix}. \quad (2.4)$$

Here ψ_L and ψ_R are two-component objects, where the components represent the two available spin states of a spin 1/2 particle. In this basis it is clear that $\frac{1+\gamma^5}{2}$ and $\frac{1-\gamma^5}{2}$, when multiplied by a spinor, project out the left- and right-handed components, respectively.

2.1.3 Pion decay rate

A representation of the lowest-order Feynman diagram for the decay of the pion into a lepton and neutrino is shown in Figure 2.1.

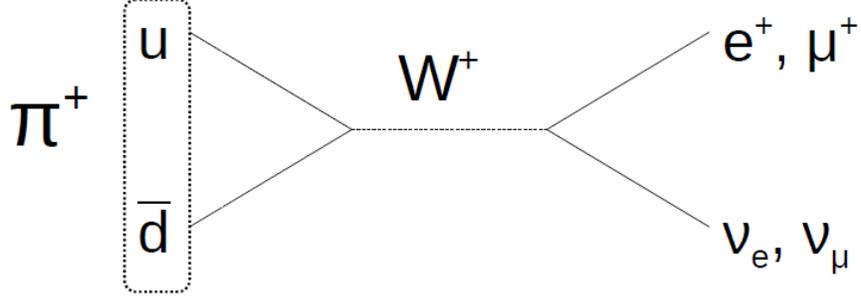


Figure 2.1: Feynman diagram for the decay of the positively charged pion. W^+ represents the positively charged gauge boson mediating the weak interaction, and u and \bar{d} represent an up and an anti-down quark, respectively.

To compute the rate of this process, it is necessary to calculate the matrix element $\mathcal{M} = \langle l^+ \nu_l | \mathcal{L} | \pi^+ \rangle$, where $l = e, \mu$ and \mathcal{L} is the Lagrangian given in Equation 2.3 [10]. Expanding this into a hadronic part and a leptonic part gives

$$\mathcal{M} = \frac{iG_F V_{ud}}{\sqrt{2}} \langle 0 | \bar{d} \gamma^\mu \gamma^5 u | \pi^+ \rangle \bar{l}(p_l) \gamma^\mu (1 + \gamma^5) \nu(p_\nu), \quad (2.5)$$

where p_l and p_ν are the momentum carried by the outgoing anti-lepton and neutrino, respectively. In the hadronic matrix element we have used the fact that only the axial-vector part of the weak interaction, $\gamma^\mu \gamma^5$, contributes to the transition amplitude. The pion itself is pseudoscalar, so the vector part of the weak interaction results in an expression with odd parity; when the integral is done over all space, it therefore vanishes.

The remaining hadronic matrix element is not simple to calculate due to strong interaction effects; however, this term must be a Lorentz four-vector. Since the pion is spinless, the only four-vector that can contribute in general is the momentum transfer q^μ to the virtual W [33]. The most

2.1. Pions in the Standard Model

general expression for the hadronic part of Equation 2.5 can be written as [10]

$$\langle 0 | \bar{d} \gamma^\mu \gamma^5 u | \pi^+ \rangle = i F_\pi q^\mu e^{iqx}, \quad (2.6)$$

where F_π is a constant parameterizing the strong interaction effects, usually called the pion decay constant. Summing over spin states and integrating over outgoing particle energies leads to the following expression for the decay rate [10]:

$$\Gamma_{\pi \rightarrow l}^0 = \frac{G_F^2 V_{ud}^2 m_\pi F_\pi^2 m_l^2}{4\pi} \left(1 - \frac{m_l^2}{m_\pi^2} \right)^2. \quad (2.7)$$

Although this formula contains F_π , the value of which must be taken from experiment, taking the ratio of the electron mode to the muon mode gives the simple expression

$$R_\pi^0 = \frac{\Gamma_{\pi \rightarrow e}}{\Gamma_{\pi \rightarrow \mu}} = \frac{m_e^2}{m_\mu^2} \left(\frac{m_\pi^2 - m_e^2}{m_\pi^2 - m_\mu^2} \right)^2 = 1.283 \times 10^{-4}. \quad (2.8)$$

This expression contains the ratio of the electron mass to the muon mass squared, which is responsible for the small value of the branching ratio. The physical origin of this factor is the $1 + \gamma^5$ term in the electroweak Lagrangian (Equation 2.3), which allows only right-handed anti-leptons and left-handed neutrinos to emerge from this decay. For massless fermions, handedness is equivalent to helicity, the direction of a particle's spin relative to its direction of motion. Left-handed massless fermions have negative helicity, and right-handed massless fermions have positive helicity.

For massive fermions, both spin states are possible, but there is an energy-dependent suppression factor associated with the helicity state that would be forbidden if the fermion were massless. That is, as the energy of a fermion begins to exceed its rest mass, it begins to behave approximately like a massless particle. In the rest frame of the pion, the neutrino and the anti-lepton must emerge back-to-back; since the pion is spin zero, the neutrino and the anti-lepton are forced into the same helicity state. The allowed direction of the spins is shown in Figure 2.2.

For both the positron and the muon decay modes, there is a suppression factor owing to this effect of $1 - \left(\frac{v_l}{c}\right)^2 \approx \left(\frac{2m_l}{m_\pi}\right)^2$, where v_l and m_l are the

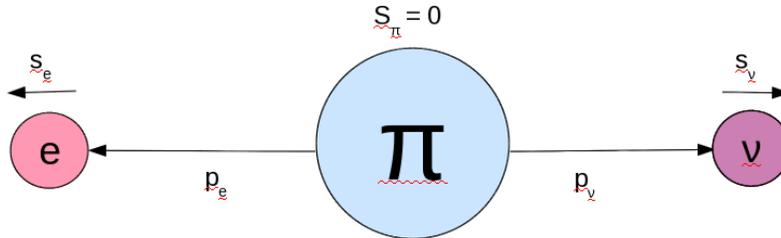


Figure 2.2: The allowed directions of the spins and linear momenta of the pion decay products, in the pion rest frame. The requirement, from angular momentum conservation, that the spins of the positron and the neutrino must be in opposite directions, leads to the suppression of the positron mode relative to the muon (see text).

velocity and mass of the outgoing lepton, m_π is the pion mass, and c is the speed of light [33]. When the ratio of the rates of the two modes is taken, the factor of $\left(\frac{m_e}{m_\mu}\right)^2$ in Equation 2.8 is obtained.

2.1.4 Corrections

Equation 2.8 must be corrected for higher-order effects. Leading-order radiative corrections, still treating the pion as point-like, were first calculated by Berman [34] and Kinoshita [35] in the late 1950s. They considered the influence of diagrams involving both real and virtual photons. Diagrams involving a real photon are referred to as Inner Bremsstrahlung (IB) processes, and diagrams involving a virtual photon are referred to as emission and re-absorption (ER). The diagrams that contribute are shown in Figure 2.3.

Although the calculation of these diagrams requires both infrared and ultraviolet cutoffs to be imposed, their effect on R_π can still be rigorously computed. The term involving the infrared cutoff cancels exactly for IB and ER processes, and the ultraviolet cutoff cancels in the branching ratio, although it affects the individual decay rates. Ultimately, a correction of -3.929% to R_π was obtained.

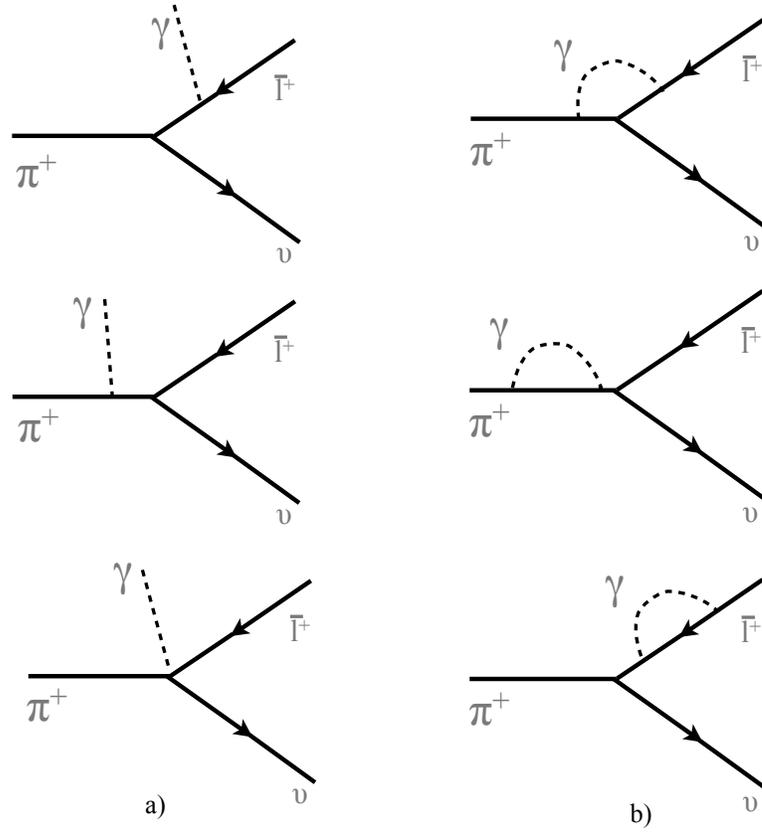


Figure 2.3: Feynman diagrams for the radiative corrections to pion decay, from real (a) and virtual (b) photons. \bar{l}^+ denotes an anti-lepton.

2.1. Pions in the Standard Model

By the 1970s, however, it was clear the pion was not a fundamental particle, but rather a quark anti-quark pair, and the validity of this approach was unclear. In 1976, Marciano and Sirlin expanded the decay rate in powers of lepton mass, m_l , including the contribution of structure-dependent effects. These are diagrams where a photon, either real or virtual, is emitted from the internal structure of the pion. They showed that neither the leading-order m_l term (the leading term in R_π) nor the leading correction term, of order $\alpha \ln(m_\mu/m_e)$, depends on strong interaction effects. They obtained -3.7% for the leading-order correction, and argued that the higher-order corrections were likely to be significantly less than 1% [15]. They repeated the calculation in 1993 with a more rigorous assessment of the theoretical uncertainty, and obtained $R_\pi = (1.2352 \pm 0.0005) \times 10^{-4}$ [36].

The most recent calculation was done by Cirigliano and Rosell [1] incorporating strong interaction effects using Chiral Perturbation Theory (ChPT), a low-energy effective field theory for quantum chromodynamics (QCD), the theory which describes the strong interaction. ChPT allows for an expansion of the decay rates in powers of the pion mass (in this case) and the electromagnetic coupling, through which the uncertainty on the ratio can be tightly constrained. The calculation in Reference [1] is done to $\mathcal{O}(e^2 p^4)$, where e is the electromagnetic coupling constant and p is proportional to the pion mass. The branching ratio can be written as

$$R_\pi = R_\pi^0 [1 + \Delta_{e^2 p^2} + \Delta_{e^2 p^4} + \Delta_{e^2 p^6} + \dots] [1 + \Delta_{LL}]. \quad (2.9)$$

Here $\Delta_{e^2 p^n}$ are the successive terms in the chiral expansion, and Δ_{LL} represents corrections of order $\alpha^n \ln^n(m_\mu/m_e)$. Although the calculation is done to $\mathcal{O}(e^2 p^4)$, this expression contains a term to $\mathcal{O}(e^2 p^6)$. This term arises from the emission of a photon by the decaying pion, which evades the helicity suppression, and so must be taken into account despite being higher-order. Photons emitted at any other part of the pion decay diagram, such as real bremsstrahlung from the decay lepton or a loop starting on the W line, do not affect the helicity suppression. Table 2.2 gives the size of each of these corrections [1].

The result of this calculation is

$$R_\pi = (1.2352 \pm 0.0001) \times 10^{-4}. \quad (2.10)$$

2.2. Beyond the Standard Model

Table 2.2: Corrections to the leading-order value for R_π .

Term	Value (%)
$\Delta_{e^2p^2}$	-3.929
$\Delta_{e^2p^4}$	0.053 ± 0.011
$\Delta_{e^2p^6}$	0.073
Δ_{LL}	0.054

However, it was shown in Reference [2] that $\mathcal{O}(\alpha^2)$ two-loop diagrams contribute an additional 0.0001 to the uncertainty, giving

$$R_\pi = (1.2352 \pm 0.0002) \times 10^{-4}. \quad (2.11)$$

The uncertainty on the SM prediction for R_π is approximately an order of magnitude smaller than the experimental uncertainty. Thus, reducing the experimental uncertainty closer to the level of the theoretical uncertainty will improve the constraints on many extensions to the SM, which predict deviations from the above value. Indeed, the effect of new physics at the weak (\sim TeV) scale on R_π is expected to be in the range of 1% to 0.01% [1]; much of this parameter space will therefore be constrained by the combined PIENU/PEN result.

2.2 Beyond the Standard Model

The SM is a very successful theory, having stood up to decades of experimental scrutiny. Although a number of intriguing experimental anomalies currently exist, such as the discrepancy between the experimental and theoretical values of the anomalous magnetic moment of the muon [37] and the proton radius puzzle [38], no experiment has reached the 5σ threshold necessary to claim a discovery (different techniques for measuring the proton radius differ by 7σ , but no discrepancy of this size exists between an experimental result and the SM prediction).

Despite the many successes of the SM, it is known to be an incomplete description of nature. Cosmological observations give evidence for the existence of dark matter [39] and dark energy [40], neither of which is explained by the SM. The asymmetry between matter and antimatter in the SM, CP

violation, is very small, but the universe appears to be almost entirely composed of matter [41]. That neutrinos have mass is now proven by oscillation experiments, but the mass generation mechanism is not included in the SM. If they obtain mass via the usual Higgs mechanism, there must exist as yet unobserved right-handed neutrino states. Finally, and most obviously, the SM does not describe the gravitational interaction between particles.

In addition to these missing pieces, several aspects of the SM as it stands are considered theoretically unsatisfactory, either because of the necessity of fine-tuning or the large number of parameters whose values have no explanation, and must be determined from experiment. These include the apparent lack of CP violation in the strong interaction (the “strong CP” problem) [42], the existence of three generations of particles that differ only by mass, and the lack of a mechanism to keep quantum corrections to the Higgs mass small, in the presence of fermions (the “naturalness” or “hierarchy” problem) [43].

Given the seeming arbitrariness of several aspects of the SM, and its known shortcomings, it is generally believed that it will eventually be superseded by a more complete theory. However, given the lack of any experimental evidence to indicate the nature of this theory, many extensions to the SM have been proposed; a selection of these with consequences for the predicted value of R_π will now be discussed.

2.2.1 New pseudoscalar interactions

The largest effect on R_π would come from some new interaction that allowed the positron mode without helicity suppression. This could take the form of either a scalar or pseudoscalar operator, although the pion cannot decay directly through a scalar interaction. Since the pion is itself a pseudoscalar particle, an operator of odd parity is required in order for the transition amplitude between the pion and the vacuum to be non-zero.

However, scalar interactions can induce pseudoscalar interactions through electroweak renormalization effects. This possibility is of particular interest, as new scalar interactions appear in several well-motivated extensions to the SM, including extra Higgs multiplets, leptoquarks, and compositeness of quarks or leptons [44].

2.2. Beyond the Standard Model

The transition amplitude for pion decay through a pseudoscalar interaction is given by [44]

$$\langle 0 | \bar{u} \gamma^5 d | \pi \rangle = i\sqrt{2} \frac{f_\pi m_\pi^2}{m_u + m_d} = i\sqrt{2} \tilde{f}_\pi. \quad (2.12)$$

Assuming only left-handed neutrinos are produced by the new pseudoscalar, the Lagrangian for this interaction is given by

$$\mathcal{L}_P = -i \frac{\rho}{2\Lambda^2} [\bar{l}(1 - \gamma^5)\nu_l] [\bar{u}\gamma^5 d], \quad (2.13)$$

where ρ is the coupling constant for the new pseudoscalar and Λ is its mass scale. This leads to a matrix element

$$\mathcal{M}_P = \rho \frac{\tilde{f}_\pi}{\sqrt{2}\Lambda^2} [\bar{l}(1 - \gamma^5)\nu_l]. \quad (2.14)$$

The total matrix element for pion decay then becomes the sum of the SM matrix element and \mathcal{M}_P . After squaring the total matrix element and summing over final states, the branching ratio becomes, assuming lepton universality holds for the new interaction [44],

$$R_\pi = R_\pi^{SM} \left(1 + \sqrt{2} \frac{\tilde{f}_\pi \text{Re}(\rho)}{G_F \Lambda^2 f_\pi V_{ud} m_e} + \mathcal{O}\left(\frac{1}{\Lambda^4}\right) \right). \quad (2.15)$$

Assuming real coupling of approximately the same strength as the weak interaction, this becomes

$$\frac{R_\pi}{R_\pi^{SM}} - 1 \sim \left(\frac{1 \text{TeV}}{\Lambda} \right)^2 \times 10^3. \quad (2.16)$$

This expression applies to a new interaction whose nature is fundamentally pseudoscalar; the effect on the branching ratio from an induced pseudoscalar, from a fundamentally scalar interaction, depends on the details of the new interaction.

2.2.2 R-parity violating supersymmetry

Supersymmetry (SUSY) is widely considered to be one of the most well-motivated extensions to the SM, as it has the potential to solve the Higgs self-coupling problem and explain the nature of dark matter [43]. Versions of SUSY that preserve R-parity are unlikely to have a measurable effect on the branching ratio, at the level of precision of the PIENU experiment. R-parity is a quantum number defined as $P_R = -1^{3B+L+2s}$ [45], where B is

2.2. Beyond the Standard Model

baryon number, L is lepton number, and s is spin. R-parity is conserved in the SM.

In SUSY models that conserve R-parity, there are a variety of processes that could contribute to R_π , but the contribution would not be large except in specific, theoretically disfavoured regions of parameter space. However, if both R-parity and lepton number conservation are violated, tree-level sfermion exchange (the supersymmetric partner to the electron or muon, in this case) would contribute to R_π at a potentially measurable level [16].

The effect on R_π is determined by both the sfermion masses, $m_{\bar{e}}$ and $m_{\bar{\mu}}$, and the coefficients of the R-parity violating interactions in the Lagrangian, λ'_{11k} for the decay to an electron and λ'_{21k} for the decay to a muon. Defining Δ'_{ijk} in terms of the coefficients λ'_{ijk} as

$$\Delta'_{ijk} = \frac{|\lambda'_{ijk}|^2}{4\sqrt{2}G_F m_{\bar{f}}^2}, \quad (2.17)$$

where G is the Fermi coupling constant and $m_{\bar{f}}$ is the sfermion mass, the effect on R_π is given by [16]

$$\frac{\Delta R_\pi}{R_\pi^{SM}} = 2(\Delta'_{11k} - \Delta'_{21k}). \quad (2.18)$$

The allowed regions for Δ'_{11k} and Δ'_{21k} from precision measurements of electroweak parameters are shown in Figure 2.4, at 95% confidence level [16]. Shown are the constraints using the PDG value of the branching ratio $((1.230 \pm 0.004) \times 10^{-4})$ in the fit and, in the dashed red curve, the projected constraint from a 0.1% measurement of R_π , assuming the same central value.

2.2.3 Charged Higgs

The Minimal Supersymmetric Standard Model (MSSM) contains a neutral Higgs doublet, a neutral Higgs singlet, and a charged Higgs doublet [46]. A charged Higgs particle could replace the W in the pion decay diagram (Figure 2.1) [16]. If the coupling of the charged Higgs is $\frac{g}{2\sqrt{2}}\lambda_{ud}$ to pions and $\frac{g}{2\sqrt{2}}\lambda_{l\nu}$ to leptons, the effect on the branching ratio is [2]

$$1 - \frac{R_\pi}{R_\pi^{SM}} = \frac{2m_\pi^2}{m_e(m_u + m_d)} \frac{m_W^2}{m_{H^\pm}^2} \lambda_{ud} \left(\lambda_{e\nu} - \frac{m_e}{m_\mu} \lambda_{\mu\nu} \right). \quad (2.19)$$

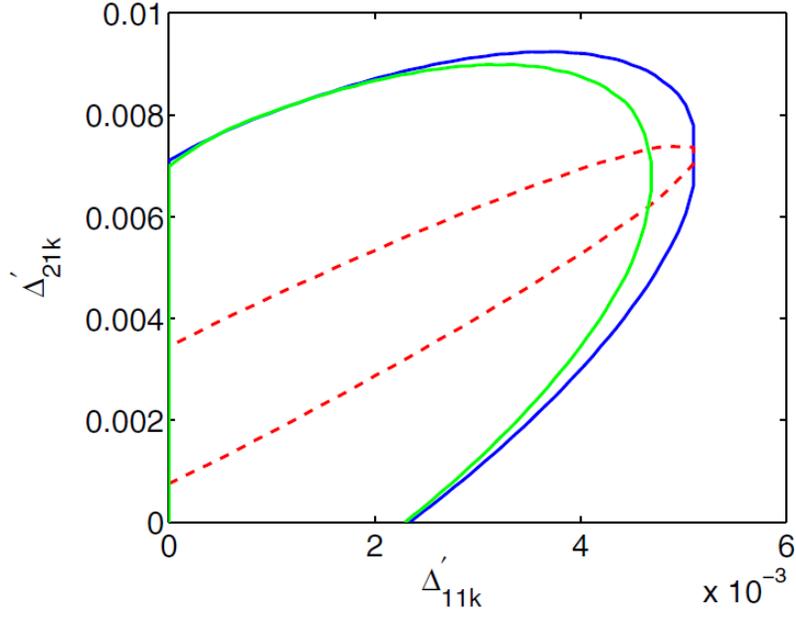


Figure 2.4: The constraint on the coefficients of R-parity violating interactions from a fit of electroweak observables, including R_π . The blue curve shows the constraint using the PDG value for the branching ratio, and the dashed red curve shows the projected constraint from a 0.1% measurement of the branching ratio, with the same central value as the blue curve. The green curve shows the expected limits with the results of the Qweak experiment, measuring the weak charge of the proton, at Jefferson Lab [5].

2.2. Beyond the Standard Model

If, for example, the couplings λ_{ud} , $\lambda_{e\nu}$, and $\lambda_{\mu\nu}$ are all equal to $\frac{\alpha}{\pi}$, where α is the electromagnetic coupling, a 0.1% measurement of R_π corresponds to $m_{H^\pm} \approx 400$ GeV. If, however, $\lambda_{e\nu}/\lambda_{\mu\nu} = m_e/m_\mu$, as for the Standard Model Higgs, the term in the brackets cancels out and R_π is insensitive to the presence of charged Higgs bosons.

2.2.4 Leptoquarks

Leptoquarks are defined by having both baryon and lepton number, and possess a range of other properties in different models. Chiral leptoquarks couple only to left- or right-handed particles, whereas non-chiral leptoquarks couple to particles of either handedness. The bounds from R_π on the latter are very strong: $M_{LQ}^2/(g_L g_R) > (100 \text{ TeV})^2$, where M_{LQ} is the leptoquark mass and g_L and g_R are the couplings. R_π also constrains the M/g ratio for chiral leptoquarks that couple to left-handed particles; the exact constraint depends on the leptoquarks under consideration, but for scalar leptoquarks are in the range of 1-10 TeV [17].

2.2.5 Massive neutrinos

Although the number of stable light neutrinos with SM couplings is restricted to three by the measurements of the Z peak done at LEP in the late 1980s [47], no such restriction exists on heavy or weakly coupled neutrinos. Indeed, these appear in many extensions to the SM; heavy neutrinos are a promising dark matter candidate and could provide a natural way for the SM neutrinos to acquire their small masses [48]. Many experiments currently exist or are under construction to better measure neutrino properties. Some anomalies have been observed, for example by LSND and MiniBOONE, that could be resolved by heavy neutrinos [49].

The PIENU experiment is sensitive to neutrinos in the mass range of 0-130 MeV, and particularly to those above 55 MeV. The presence of any neutrino heavier than a few MeV would reduce the helicity suppression of the $\pi^+ \rightarrow e^+ \nu_e$ mode, and thus enhance the branching ratio. The rate of the decay $\pi^+ \rightarrow e^+ \nu_i$, where ν_i is a heavy neutrino, relative to the rate of the $\pi^+ \rightarrow e^+ \nu_e$ decay, is given by [18]

$$\frac{\Gamma(\pi^+ \rightarrow e^+ \nu_i)}{\Gamma(\pi^+ \rightarrow e^+ \nu_e)} = |U_{ei}|^2 \rho_e. \quad (2.20)$$

2.3. Lepton universality

U_{ei} is the mixing parameter between ν_e and ν_i (analogous to a CKM matrix element) and ρ_e is a function of $\delta_i = m_{\nu_i}^2/m_\pi^2$ and $\delta_e = m_e^2/m_\pi^2$, where m_{ν_i} is the heavy neutrino mass, m_e is the electron mass, and m_π is the pion mass:

$$\rho_e = \frac{[1 + \delta_e^2 + \delta_i^2 - 2(\delta_i + \delta_e + \delta_i\delta_e)]^{1/2}[\delta_i + \delta_e - (\delta_i - \delta_e)^2]}{\delta_e(1 - \delta_e)^2}. \quad (2.21)$$

Using the value for R_π based on the 2010 data, this translates to a limit of $|U_{ei}|^2 < 1.32 \times 10^{-6}$ for $m_{\nu_i} = 50$ MeV, at 90% confidence. The limit increases as m_{ν_i} goes to zero.

Above 55 MeV, the positron energy would be sufficiently reduced that an extra peak would appear in the $\pi^+ \rightarrow e^+\nu_e$ energy spectrum. Figure 2.5 shows the upper limit on $|U_{ei}|^2$ obtained through a search for extra peaks in PIENU data taken in 2009, compared with the limits from the previous PIENU experiment [7] [6]. The same analysis using the full PIENU data set is under way, which is expected to improve the limit by a factor of 3-5.

2.3 Lepton universality

A key feature of the SM is the presence of three generations of fermions, which differ only by mass. The couplings to gauge bosons are assumed to be identical; whether this is exactly true, or only an approximation, is a crucial test of the structure of the SM. Referring to Equation 2.7, the rate of pion decay to either electrons or muons is proportional to g^4 , where g is the weak interaction coupling constant. One factor of g^2 arises from the quark vertex, and the other from the lepton vertex. If the couplings to the W of electrons and muons are not assumed to be equal, R_π is then given by

$$R_\pi = \left(\frac{g_e}{g_\mu}\right)^2 R_\pi^{SM}. \quad (2.22)$$

Constraints on the ratios of the coupling constants come from many different types of measurement, as shown in Table 2.3.

R_π currently provides the most precise test of electron-muon universality, although the branching ratio of tau decays to muons and electrons is close. However, these tests are not exactly equivalent; since the pion is spin zero while the tau is spin 1/2, the mediating W boson in the former case must be in the spin zero state, whereas in the latter case all spin states contribute.

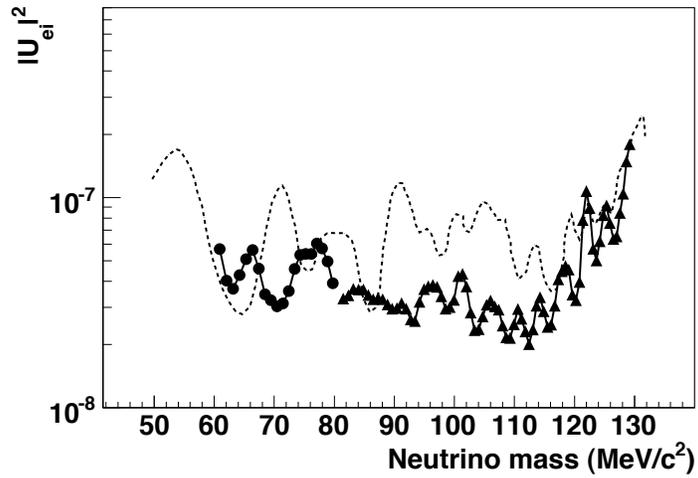


Figure 2.5: The 90% C.L. upper limit on the heavy neutrino mixing parameter, as a function of its mass. The dashed line shows the result from the previous PIENU experiment [6], and the circles and triangles are the limits from a subset of PIENU data, published in 2011 [7]. The circles indicate a restricted angular region was used when constructing the $\pi^+ \rightarrow e^+ \nu_e$ energy spectrum.

2.3. Lepton universality

Table 2.3: Experimental results on lepton universality tests from studies of pion, kaon, tau, and W decays. Here \mathcal{B} represents the branching fraction of a particular decay mode.

Decay mode	g_μ/g_e
$\mathcal{B}_{\pi \rightarrow \mu}/\mathcal{B}_{\pi \rightarrow e}$	1.0004 ± 0.0012 [3]
$\mathcal{B}_{\tau \rightarrow \mu}/\mathcal{B}_{\tau \rightarrow e}$	1.0018 ± 0.0014 [50]
$\mathcal{B}_{K \rightarrow \mu}/\mathcal{B}_{K \rightarrow e}$	0.996 ± 0.005 [51]
$\mathcal{B}_{K \rightarrow \pi \mu}/\mathcal{B}_{K \rightarrow \pi e}$	1.002 ± 0.002 [52]
$\mathcal{B}_{W \rightarrow \mu}/\mathcal{B}_{W \rightarrow e}$	0.997 ± 0.010 [52]
	g_τ/g_μ
$\mathcal{B}_{\tau \rightarrow e \tau_\mu}/\tau_\tau$	1.0011 ± 0.0015 [50]
$\mathcal{B}_{\tau \rightarrow \pi}/\mathcal{B}_{\pi \rightarrow \mu}$	0.9963 ± 0.0027 [50]
$\mathcal{B}_{\tau \rightarrow K}/\mathcal{B}_{K \rightarrow \mu}$	0.9858 ± 0.0071 [50]
$\mathcal{B}_{W \rightarrow \tau}/\mathcal{B}_{W \rightarrow \mu}$	1.039 ± 0.013 [52]
	g_τ/g_e
$\mathcal{B}_{\tau \rightarrow \mu \tau_\mu}/\tau_\tau$	1.0029 ± 0.0015 [50]
$\mathcal{B}_{W \rightarrow \tau}/\mathcal{B}_{W \rightarrow e}$	1.036 ± 0.014 [52]

The most precise tests of electron-tau and muon-tau universality come from $\tau \rightarrow e$ and $\tau \rightarrow \mu$ decay rates at B factories [50].

Other interesting results, that have generated much attention, are the measurements by LHCb of the flavour-changing neutral current processes $B^+ \rightarrow K^+ l^+ l^-$ [53], where $l = e, \mu$, and the charged-current processes $\bar{B}^0 \rightarrow D^{*+} l^- \bar{\nu}_l$ [54], where $l = \mu, \tau$. The first found an excess of 2.8σ in the electron mode, and the second found an excess of 2.1σ in the τ mode. The BaBar collaboration also reported a 2.7σ excess in this mode, and a 2.0σ in the similar $\bar{B}^0 \rightarrow D^+ \tau^- \bar{\nu}_\tau$ [55].

The deviations from universality required to explain these measurements are large compared to the uncertainties quoted in Table 2.3. In order to explain these results in terms of new physics, while remaining consistent with other measurements, generally requires the new physics to couple preferentially to the third generation of particles [56].

2.4. *Physics Reach*

This can be done by assuming the new physics couples only to the third generation in the interaction basis, with effects on the first two generations coming from the mismatch between the mass basis and the interaction basis, or by assuming a mass-dependent coupling for the new physics. In either case, the effect on R_τ would be highly suppressed; the CKM matrix elements linking the first and third generations of quarks, for example, are $V_{ub} = 0.0035$ and $V_{td} = 0.0087$.

2.4 Physics Reach

Analysis of the full PIENU data set will test electron-muon universality at the 0.05% level, improve the limits on the couplings of heavy neutrinos in the mass range from 0-130 MeV, constrain new pseudoscalar interactions up to $\mathcal{O}(1000 \text{ TeV})$, and provide improved constraints on various leptoquark and R-parity violating SUSY models.

Chapter 3

Experimental Setup

3.1 Beamline

The TRIUMF cyclotron accelerates H^- ions to a maximum energy of 520 MeV and a typical intensity of $300 \mu\text{A}$, divided amongst four primary beamlines. The accelerating gradient is provided by a 23.05 MHz 93 kV radiofrequency (RF) field; this corresponds to a bunch spacing of 43.4 ns, with a typical bunch width of 4 ns. Extraction from the cyclotron to the beamlines is accomplished by stripping the electrons off the H^- with a thin foil, thereby reversing the direction of the magnetic steering. The PIENU production target, a 1cm thick piece of beryllium, was along beamline 1A, which sees a typical proton current of about $100 \mu\text{A}$. At the production target, pions, muons, and positrons were produced with a wide range of energies. The M13 secondary beamline selected positively charged particles of momentum 75 MeV/c, with a roughly 1% spread. A schematic of the extended beamline is shown in Figure 3.1 [31]; the extension was added in 2008 for the PIENU experiment. There were three dipole magnets, labelled B1-3, and ten quadrupoles, labelled Q1-10. The last dipole and the last three quadrupoles make up the extension. In addition, there were three sets of adjustable vertical and horizontal slits, an absorber, and a collimator. There were four focus points, labelled F1-4. The last was located at the centre of the pion stopping target.

The first bending magnet (B1) selected particles with a momentum of roughly 78 MeV/c. At this point the beam contained significant numbers of pions, muons, and positrons. While small numbers of muons and positrons were desirable for calibration purposes, the fact that the PIENU detector was in line with the beam meant that positron contamination, in particular, had to be reduced to the level of a few percent. Beam positrons traversed the entire detector and were absorbed in the calorimeter, leaving a similar amount of energy to positrons from $\pi^+ \rightarrow e^+ \nu_e$ decay. Thus, the design of the beamline was intended primarily to eliminate positrons. The end of the beamline before and after the extension is shown in Figure 3.2.

3.2. Detector

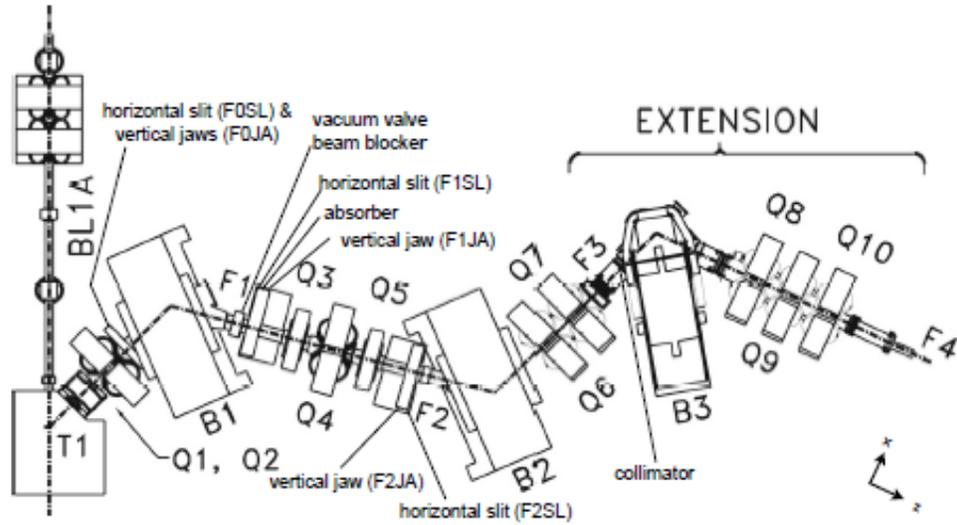


Figure 3.1: A schematic of the M13 beamline.

Positron contamination was reduced through the use of an absorber and collimator. The absorber (a thin piece of Lucite) was placed after the first bending magnet; the three particle species lost different amounts of energy in the absorber, causing them to separate horizontally when they traversed the second bending magnet. The second and third bending magnets selected particles with a momentum of about 75 MeV/c. A collimator was placed near the third focus, and positioned so as to block positrons and some muons. Because of variations in the energy loss in the absorber, many muons went through the collimator as well, along with some positrons. During beam commissioning tests, the particle content of the beam and the horizontal separation of the three particle types were measured at F3; the results are shown in Figure 3.3. The final beam composition at F4 was approximately 85% pion, 14% muon, and 1% positron [31].

3.2 Detector

Following the last set of quadrupoles (Q8-10), the pion beam exited through a stainless steel vacuum window 76.2 microns thick and entered the PIENU

3.2. Detector

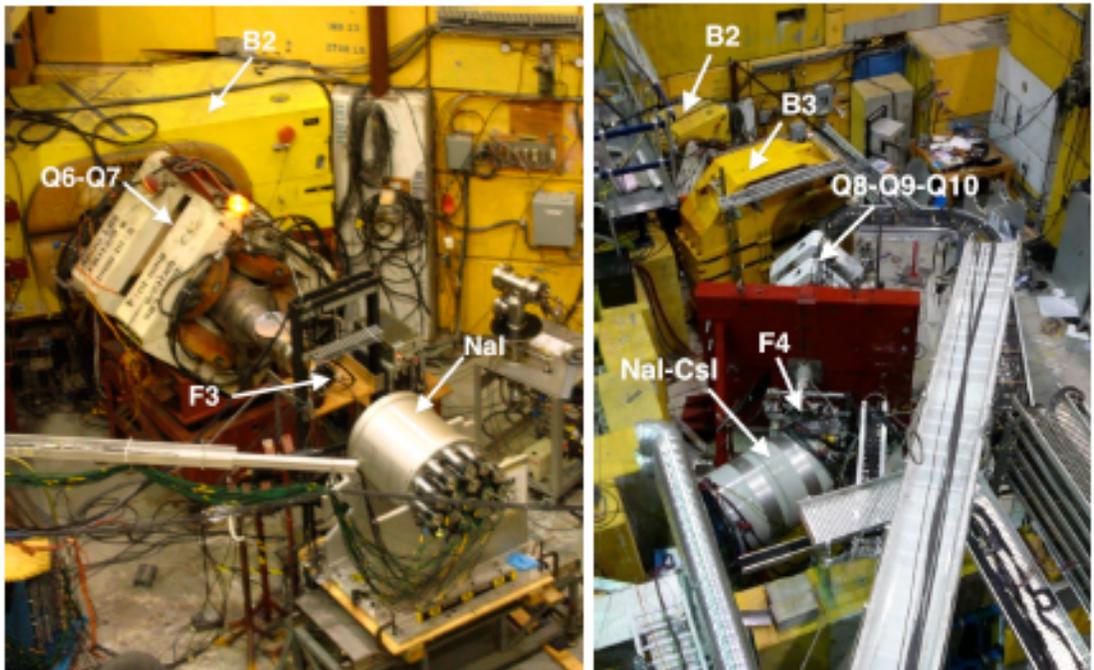


Figure 3.2: The end of the M13 beamline, before (left) and after (right) the extension. Part of the detector was in place to measure the particle content of the beam.

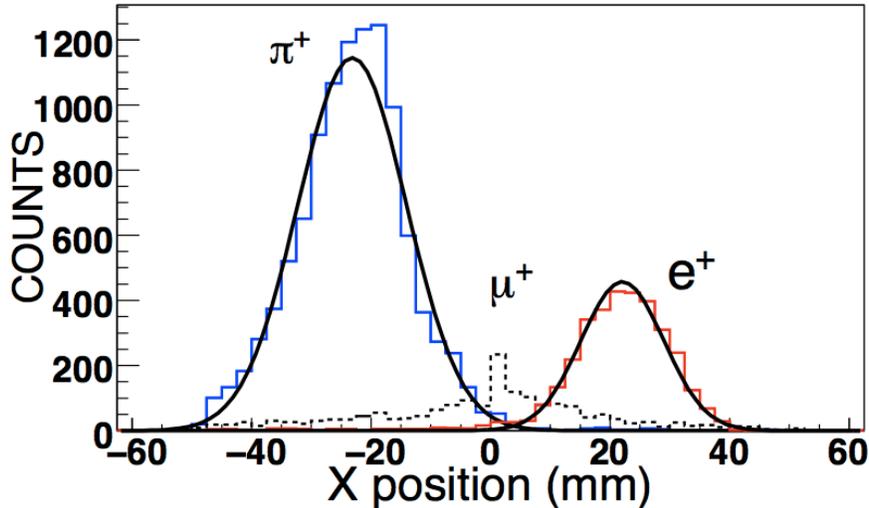


Figure 3.3: The transverse position distributions of pions, muons, and positrons at F3; the particle type was identified using time of flight. The black lines are Gaussian fits to the pion and positron distributions.

detector, which has been briefly discussed in Section 1.4.1. The physical dimensions of the detectors are shown in Table 3.1.

Physically, the detector consisted of three modular subelements: the beam wire chambers, PIENU-I, and PIENU-II. The beam wire chambers, named WC1 and WC2, were mounted on the beam pipe downstream of the vacuum window, and were the first detectors seen by the beam. One wire plane, along with its preamplifier, is shown in Figure 3.4; the chambers in place on the beam pipe are shown in Figure 3.5.

PIENU-I comprised the first four plastic scintillators, B1, B2, Tg, and T1, with their lightguides and PMTs, as well as the three silicon strip detectors, S1-3. PIENU-II comprised the last plastic scintillator, T2, a third wire chamber, WC3, and the NaI(Tl) and CsI crystals, along with their support structures. A schematic of the readout for the plastic scintillators in PIENU-I is shown in Figure 3.6; T2 was read out with optical fibres. PIENU-I and II are shown in Figures 3.7 and 3.8. BINA (the NaI(Tl) crystal) is shown in Figure 3.9, and the full crystal array is shown in Figure 3.10. A Solidworks

3.2. Detector

Table 3.1: Detector characteristics. The z position given is for the centre of the detector, except as noted for BINA.

Name	Z (mm)	Thickness (mm)	Shape	Dimensions (mm)
Plastic Scintillators				
B1	-39.03	6.6	Square	100 x 100
B2	-30.02	3.07	Square	45 x 45
Tg	0.00	8.05	Square	70 x 70
T1	19.92	3.04	Square	80 x 80
T2	72.18	6.6	Circular	171.45 diameter
Name	Z (mm)	Diameter	Wires per plane	Wire spacing (mm)
Wire Chambers				
WC1	-112.55	96.0	120	0.8
WC2	-74.41	96.0	120	0.8
WC3	55.86	230.4	96	2.4
Name	Z (mm)	Dimensions (mm)	Channels per plane	Strip pitch (mm)
Silicon Strips				
S1	-23.54	61 x 61	48	0.32
S2	-11.76	61 x 61	48	0.32
S3	10.50	61 x 61	48	0.32
Name	Z (front face, mm)	Thickness (mm)	Shape	Dimensions (mm)
Crystal scintillators				
BINA		84	480	Cylinder 240 radius
CsI		-	250	Pentagon 90 x 80

3.2. Detector

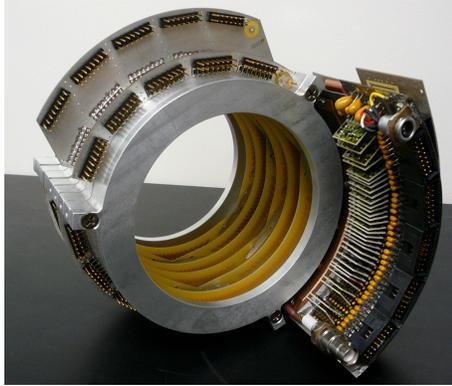


Figure 3.4: One wire chamber plane and its preamplifier board; each chamber consisted of three planes.

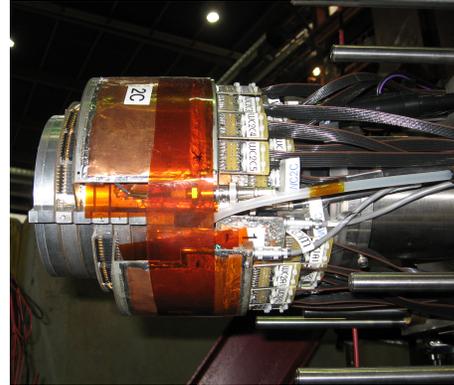


Figure 3.5: Wire chambers 1 and 2 after installation on the beam pipe.

(computer-aided design software) drawing of the full detector, along with a picture, is shown in Figure 3.11.

3.2. Detector

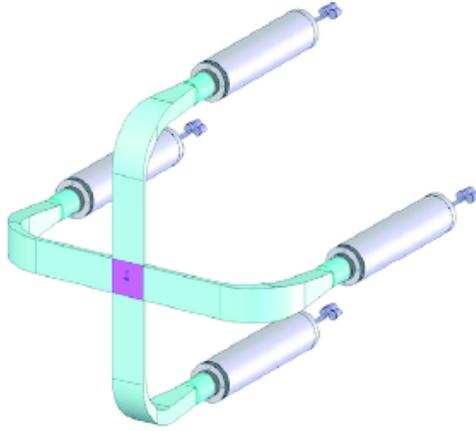


Figure 3.6: The plastic scintillator readout scheme, for B1, B2, Tg, and T1. The light from the plastic scintillator (purple) was transmitted by four acrylic lightguides (light green) to PMTs (grey cylinders).

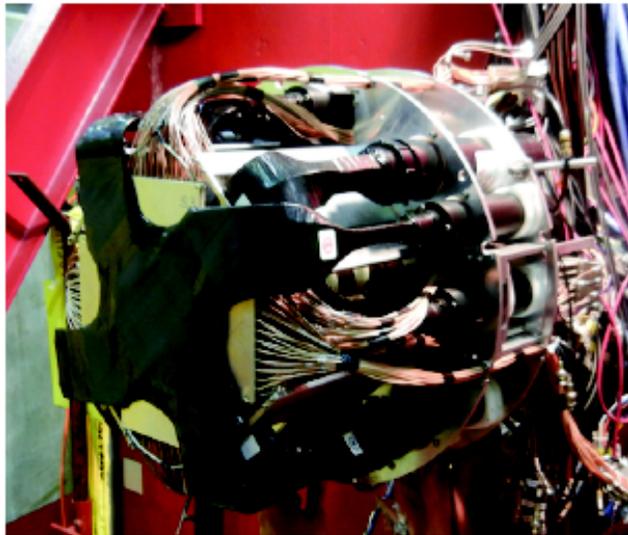


Figure 3.7: The PIENU-I detector subsystem.

3.2. Detector

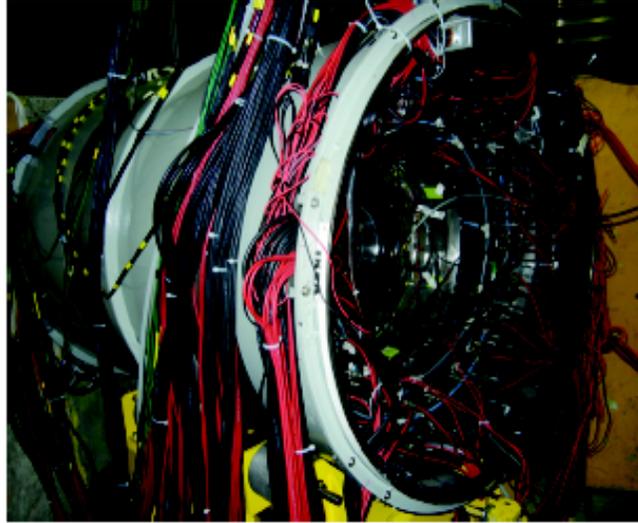


Figure 3.8: The PIENU-II detector subsystem.



Figure 3.9: The BINA detector on the test bench, with some of its PMTs in place.

3.2. Detector



Figure 3.10: BINA and the two CsI rings, with all the BINA PMTs in place.

3.2. Detector

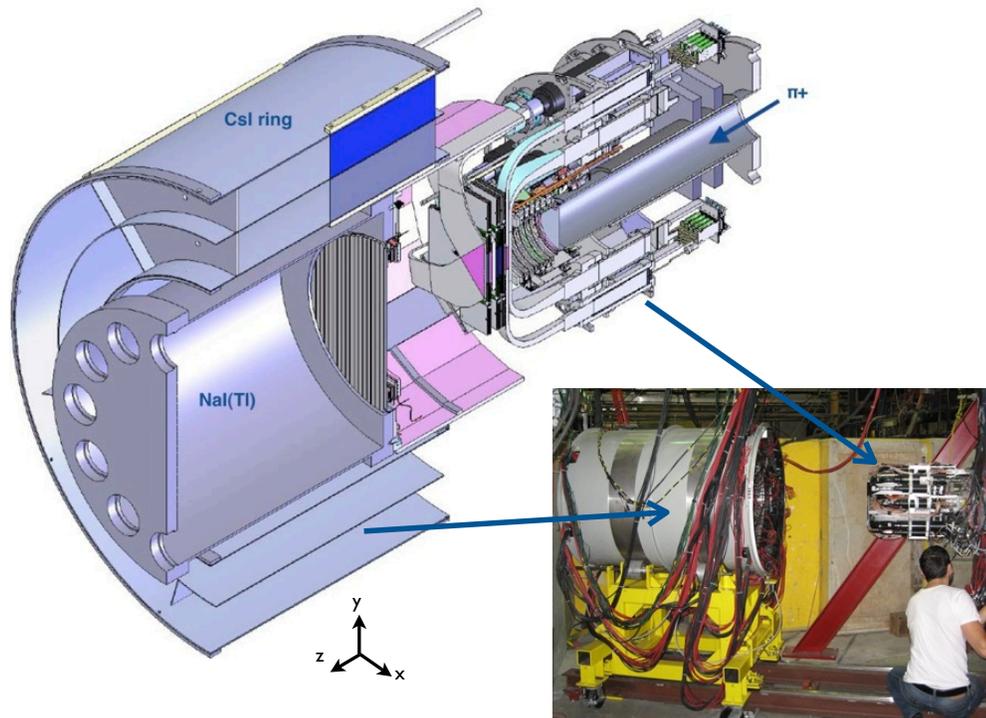


Figure 3.11: A Solidworks drawing of PIENU-I and PIENU-II, along with a picture with a human for scale. During data-taking, PIENU-II was rolled forward around PIENU-I.

3.2. Detector

The beam wire chambers, WC1 and WC2, provided tracking for the incoming pion beam. Each wire chamber contained three planes, inclined at 0, 120, and 240 degrees [8]. Each plane contained 120 wires spaced at 0.8 mm. Since each plane is sensitive to one coordinate, at fixed z , together WC1 and WC2 provided three points with which to reconstruct the particle track. The distributions of wires hit are shown in Figure 3.12; the size of the beam spot is slightly less than 1 cm (FWHM) in each of these dimensions. Events without a track in the beam spot were rejected in the analysis.

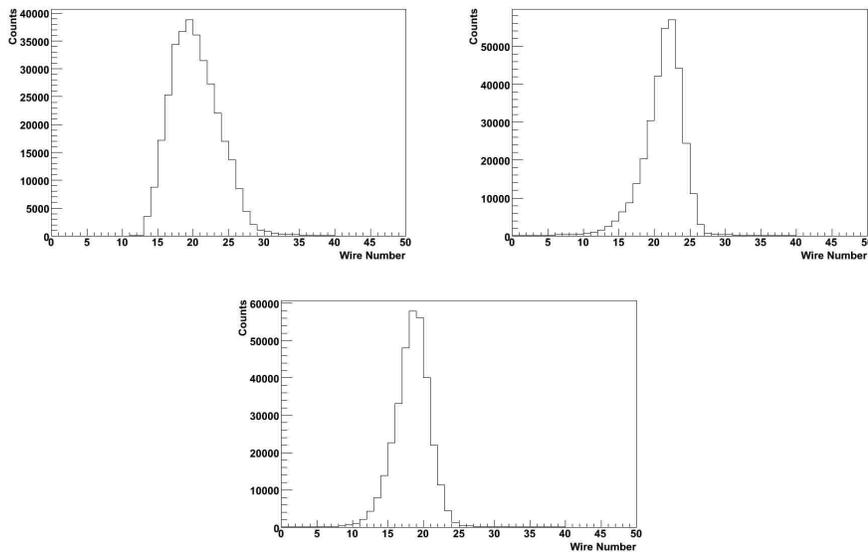


Figure 3.12: Wires hit in each plane of the first wire chamber.

Following the wire chambers, the beam passed through two plastic scintillators (B1 and B2) and two pairs of x-y silicon strip detectors (S1 and S2) [8], before reaching the target (Tg). X and y are defined in the PIENU coordinate system as being the horizontal and vertical dimensions transverse to the beam direction, which points towards positive z . The particle content of the beam can be determined from the correlation of the energy deposited in B1 and the time of the hit in B1 relative to the cyclotron RF; the distribution is shown in Figure 3.13.

3.2. Detector

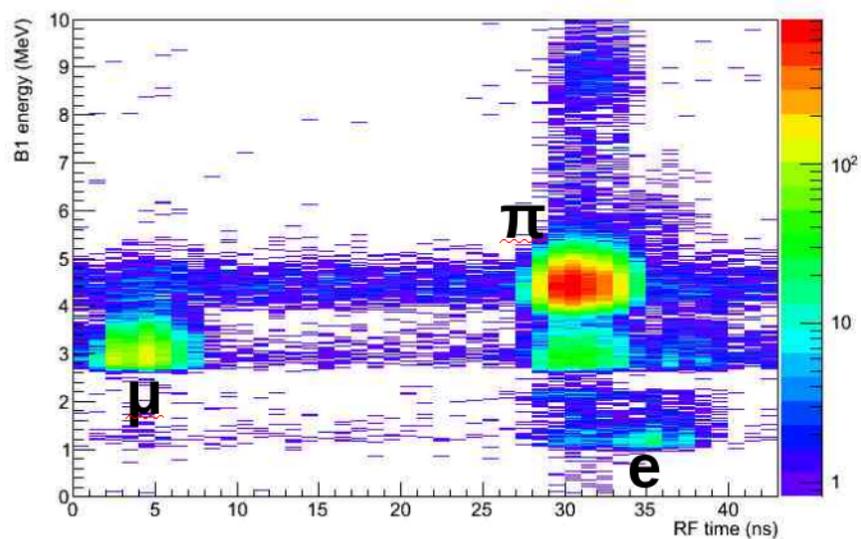


Figure 3.13: Energy deposited in B1 vs. the time of the hit relative to the peak of the cyclotron RF field. The cluster with the most events, labelled π , is caused by pions; the cluster below that, with the same timing but less energy, is caused by pions that decayed in flight prior to reaching B1. The hit is caused by the decay muon. The cluster labelled μ on the far left is due to beam muons, and the low-energy cluster labelled e is due to beam positrons.

3.2. Detector

The measured energy loss distributions for pions in each of the first three plastic scintillators are shown in Figure 3.14. The stopping distribution of the beam was centred in z in the target of thickness 8.05 mm. The distribution was approximately Gaussian, with $\sigma = 0.8$ mm. The range of the decay muon in plastic is less than 2 mm; thus, it was always contained within the target as well. A simulation of the distance travelled by the decay muon is shown in Figure 3.15. The first peak in the right-hand panel of Figure 3.14 corresponds to the full energy of the pion prior to reaching the target, while the second peak is caused by events where the muon emerged within the integration time of the pulse, leading to its energy being included as well.

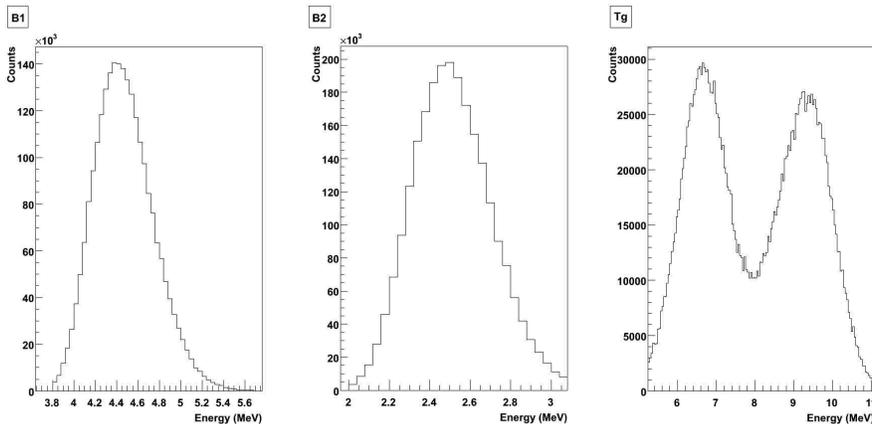


Figure 3.14: The measured energy deposited in B1, B2, and Tg by pions. The additional peak in the target spectrum is caused by events in which the energy of the 4.1 MeV decay muon is also included.

Because of the narrow energy spread of the beam and predictable energy loss in the counters, the pions almost always stopped in the target unless they decayed before reaching it. A Monte Carlo simulation of the pion stopping position along the z axis is shown in Figure 3.16; $z = 0$ is the centre of the target. The fraction of events outside the peak is approximately 0.02%, prior to the application of any cuts.

The PIENU trigger selected events in which the decay positron entered the downstream part of the detector (in the direction of the beam, towards

3.2. Detector

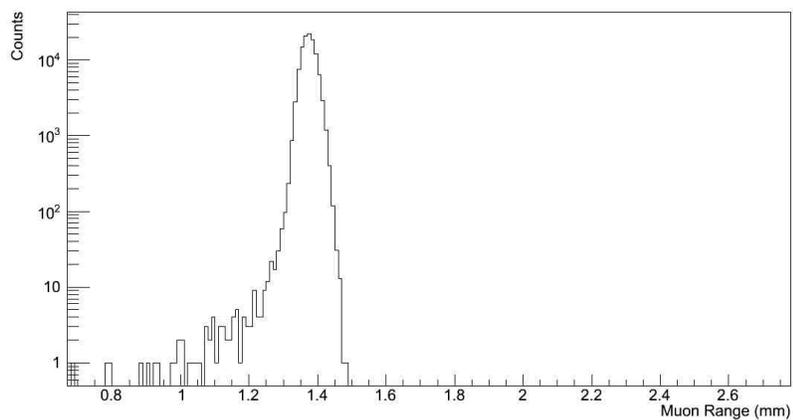


Figure 3.15: The simulated distance travelled by the 4.1 MeV $\pi \rightarrow \mu\nu$ decay muon in the target.

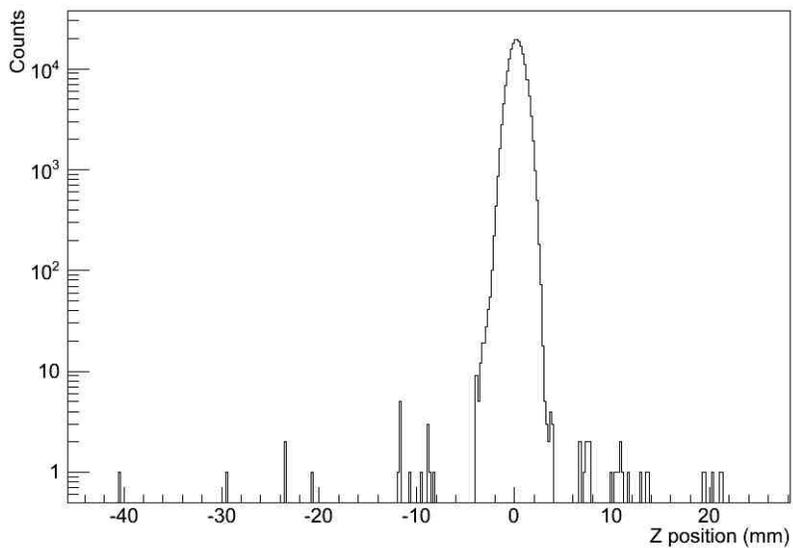


Figure 3.16: The simulated pion stopping position along the beam direction.

3.3. Trigger and DAQ

positive z). Only events with coincident hits in the plastic scintillators downstream of the target, T1 and T2, were recorded. This ensured the decay positron would enter the crystal array. Typically, the full energy of the positron was absorbed by the crystals, although there was some probability of leakage. Since the PIENU analysis hinges on distinguishing $\pi^+ \rightarrow e^+\nu_e$ positrons from $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ positrons based on their energy, detailed characterization of the calorimeter response function was required.

The axis of BINA was aligned with the axis of the beam, corresponding to 19 radiation lengths of NaI(Tl) for a 70 MeV positron entering on-axis. The amount of material along the particle trajectory was much less for positrons that entered at high angles relative to the crystal axis, increasing the probability that part of the electromagnetic shower could escape. The CsI crystals, which provided an additional 13.5 radiation lengths [8], were placed so as to capture much of the leakage from these high-angle events. Figure 3.17 shows a Monte Carlo simulation of the sum of the energy deposited in BINA and CsI by $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ positrons and $\pi^+ \rightarrow e^+\nu_e$ positrons, and the energy deposited solely in BINA by $\pi^+ \rightarrow e^+\nu_e$ positrons. The overlap between the two spectra is much greater without the addition of the CsI energy. The analysis of the crystal response is discussed in detail in Chapter 6.

3.3 Trigger and DAQ

The signals from the plastic scintillator PMTs were used as inputs to the trigger system, a digital logic circuit responsible for identifying events that should be read out and stored. Several triggers were implemented, divided into two categories: physics triggers and calibration triggers. Physics triggers were used to select events where a pion stopped in the target and the decay positron entered the crystal array, while calibration triggers were used to select other types of events to be used for calibration. If an event passed one of the triggers, the signals from all the detectors were digitized and written to disk.

The basis of the logic for physics triggers was a coincidence between B1, B2, and Tg, made by the pion, followed within 540 ns by a coincidence between T1 and T2, made by the decay positron. These were referred to as “upstream coincidence” and “downstream coincidence”, respectively. Events where the positron preceded the pion by up to 300 ns were also

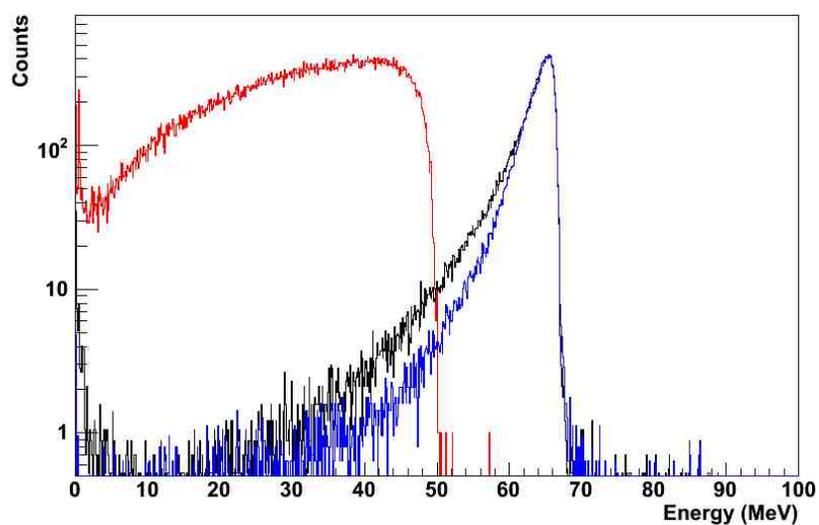


Figure 3.17: The simulated energy deposited for $\pi^+ \rightarrow e^+ \nu_e$ decays in BINA (black) and BINA and CsI (blue), and for $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ decays in BINA and CsI (red). The distributions are normalized to the maximum bin. The low-energy peak is caused by the absorption of a single 511 keV photon from a positron annihilating at rest, in T2 or the front face of BINA.

3.3. Trigger and DAQ

accepted, in order to characterize backgrounds caused by the presence of muons in Tg from previous pion decays. A discriminator was used to select a large pulse in B1, to suppress beam muon and positron events. Three different triggers were used to record pion decay events; one of the triggers had only the requirements of upstream and downstream coincidence with a large pulse in B1, and thus recorded almost entirely $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events. The time window chosen reduced the $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events by approximately a factor of five compared to $\pi^+ \rightarrow e^+ \nu_e$ events, and the trigger was prescaled by a factor of sixteen, meaning that only one event out of sixteen that triggered was actually recorded. This was done to reduce the data size, since it was not necessary to record every $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ event. This trigger was called the Prescale trigger.

Two other, unprescaled, triggers were used to enhance the proportion of $\pi^+ \rightarrow e^+ \nu_e$ events: these were referred to as Early and BinaHigh. The Early trigger fired only if the downstream coincidence was within 5-40 ns of the upstream coincidence, and BinaHigh triggered only for large energy deposit in the calorimeter. Initially this was accomplished by putting an analog sum of the BINA and CsI PMTs into a discriminator, but this trigger was later modified to use a field-programmable gate array (FPGA), which allows logic circuitry to be implemented in software. The energy threshold for this trigger was set a few MeV below the Michel edge, to ensure 100% efficiency for $\pi^+ \rightarrow e^+ \nu_e$ events. The measured time spectra of Prescale events, scaled up by a factor of sixteen, and Early events are shown in Figure 3.18, and the energy spectrum of BinaHigh events is shown in Figure 3.19.

Several calibration triggers were also used. The most important of these were the Positron trigger and the Cosmics trigger. The Positron trigger required a coincidence of B1, B2, and Tg with T1, and small energy deposit in B1, thus recording beam positrons. The Cosmics trigger required energy be deposited in coincidence in the outer and inner layers of the CsI crystals, thereby recording cosmic ray muons. The Positron trigger was used to check the calibration of the scintillators and the NaI crystal, while the Cosmics trigger was used to calibrate the CsI crystals. The full trigger diagram is shown in Appendix A. The proportions of each of these triggers within a typical 2012 run are shown in Table 3.2.

Three different systems were used to digitize the signals from the various detectors. Two of these, VF48 [57] and COPPER (the COmmon Pipelined

3.3. Trigger and DAQ

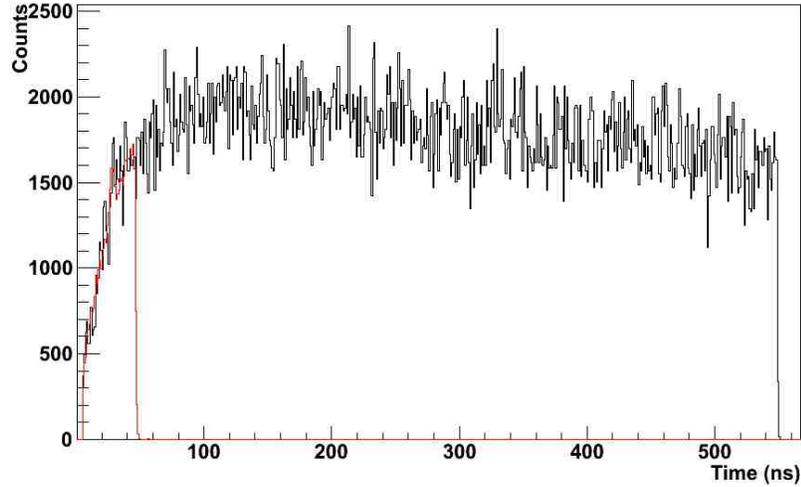


Figure 3.18: Time spectra of events recorded by the Early (red) and Prescale (black) triggers. The Prescale trigger events are scaled by a factor of sixteen. The time spectrum is dominated by $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events. Both the end of the early time window and the end of the trigger window can be seen.

Table 3.2: The number of events in one 2012 run caused by the physics triggers, and the most important calibration triggers.

Trigger	Events	Fraction
Prescale	84,097	29.5 %
Early	52,457	18.4 %
BinaHigh	157,784	55.3 %
Positron	6,683	2.3 %
Cosmics	4,094	1.4 %

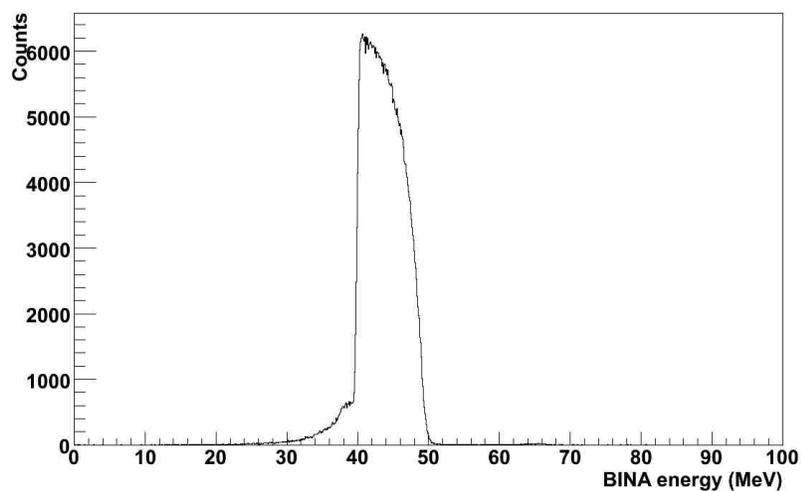


Figure 3.19: Measured energy deposited in the NaI(Tl) crystal, for events recorded by the BinaHigh trigger. The sharp rise is caused by the threshold of the BinaHigh trigger, and the fall by the endpoint of the Michel distribution. The few events above the Michel edge are mostly due to pileup.

3.3. Trigger and DAQ

Platform for Electronics Readout) [58], were analog-to-digital converters with the ability to record waveforms, while the third, VT48 [59], was a time-to-digital converter, which only recorded the time a signal went above a certain threshold. The digitization frequencies were 60 MHz, 500 MHz, and 1.6 GHz, respectively. The VF48 dynamic range was ± 250 mV with 10-bit resolution, and the COPPER dynamic range was ± 500 mV with 8-bit resolution. The signals recorded by each digitizer are shown in Table 3.3. The VT48 also recorded many other signals at various points along the trigger chain; these are shown on the trigger diagram. The VF48 was used to record the signals from the crystals and the silicon detectors, which are much slower than the plastic scintillators. For the scintillators, each channel represents one PMT. COPPER was used mainly to record the plastic scintillators, although analog sums of the BINA and CsI PMTs were digitized by COPPER as well.

The readout window of COPPER was $8 \mu\text{s}$ wide, and defined by the pion time. Approximately $1.3 \mu\text{s}$ after and $6.6 \mu\text{s}$ before the pion time were recorded. This gave three muon lifetimes of pileup coverage before the pion, reducing the background due to old muons (muons left in the target from a previous incident pion) by a factor of about twenty. The VT48 and VF48 readout windows were defined by the positron time, with readout windows of about 1 and $8 \mu\text{s}$, respectively. The times of the first hits in B1 and T1 are shown in Figure 3.20. The times plotted are the averages of the four PMTs reading out each counter. The time spectrum that is fitted to extract the branching ratio is composed of the difference of these two times. Using two detectors triggered by the same signal minimized the impact on the time spectrum of jitter in the timing system.

The large peak in the T1 spectrum, the bottom panel of Figure 3.20, is due to beam particles, primarily muons, which passed through all the scintillators. For such events the T1 and B1 hits were almost coincident. The trigger window of 300 ns before the pion and 540 ns after can be clearly seen in the T1 spectrum, and repeating peaks due to beam pileup can be seen in the B1 spectrum. These are spaced according to the cyclotron RF period of 43 ns, and are caused by events where the first hit in B1 did not trigger. The width of the peak in B1 is set by the width and timing of the gate signal to COPPER. The plots are zoomed around the pion time, and thus do not show the whole readout window.

3.3. Trigger and DAQ

Table 3.3: Detector readout channels.

Detector	Number of Channels
VT48 (1.6 GHz TDC)	
WC1	120
WC2	120
WC3	120
B1	4
B2	4
Tg	4
T1	4
T2	4
V1	1
V2	1
V3	1
VF48 (60 MHz ADC)	
S1	96
S2	96
S3	96
BINA (30 MHz)	19
CsI	97
COPPER (500 MHz ADC)	
B1	4
B2	4
TG	4
T1	4
T2	4
V1	1
V2	1
V3	1
BINA	1
CsI	4

3.3. Trigger and DAQ

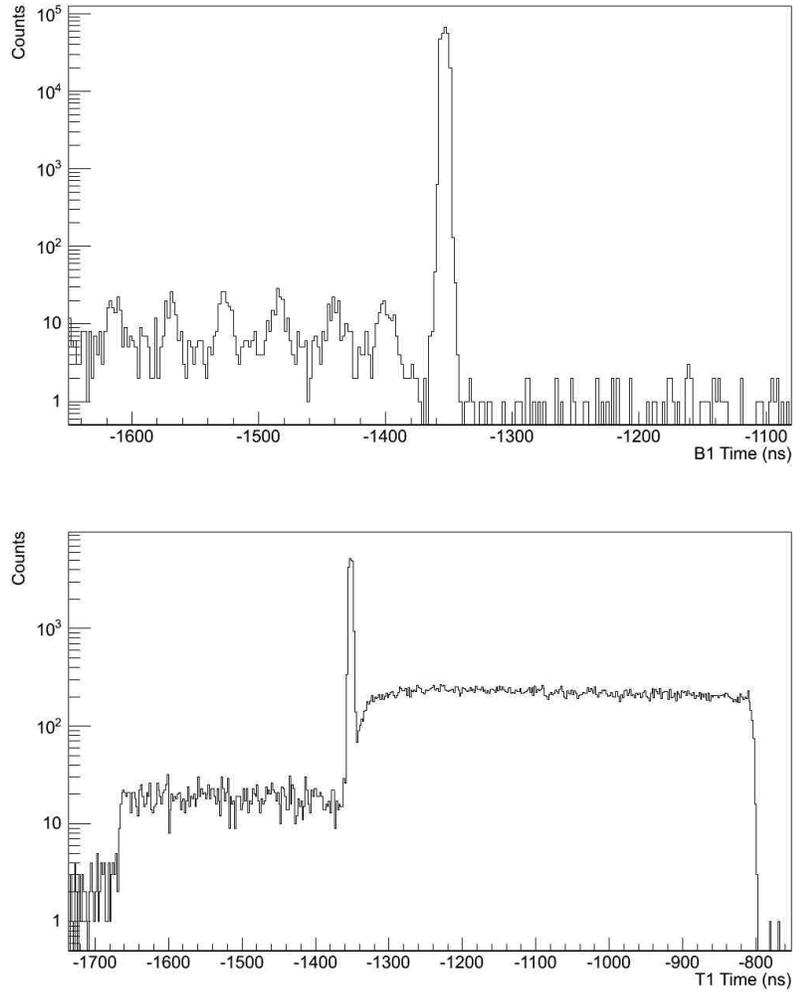


Figure 3.20: Time spectrum of hits in B1 (top) and T1 (bottom).

3.3. Trigger and DAQ

The primary purpose of the VF48s was to digitize the signals from the crystals; each CsI crystal had its own PMT digitized at 60 MHz, whereas BINA had 19 PMTs digitized at 30 MHz, due to the slowness of its response compared to the pure CsI crystals. The CsI and BINA readout windows were 60 and 40 counts long, respectively, corresponding to about 1 and 1.3 μ s. The positron signal was placed near the beginning of the BINA readout window, in order to allow for full integration of the pulse; the positron signal fell roughly in the middle of the CsI readout window.

The silicon strip detectors were also digitized by the VF48. The actual strip pitch was 80 μ m, but this was considerably smaller than the required resolution, so to reduce the number of readout channels four strips were combined into one, leading to a strip pitch of 320 μ m [8]. A further reduction was achieved by connecting each strip to a capacitor and reading out every fourth channel with an amplifier; the charge in neighbouring channels could then be used to determine the position of the hit. A schematic of the readout scheme is shown in Figure 3.21.

In addition to the detectors themselves, several experimental parameters had to be constantly monitored, such as the fields of the bending magnets, the temperature of the detector enclosure and electronics racks, the voltage across the PMTs and wire chambers, and the pressure of the wire chamber gas system. All of this information was aggregated and sent to MIDAS [60], a web-based DAQ interface developed at TRIUMF and PSI. Each subsystem ran its own front end program, which connected MIDAS to the underlying hardware and sent its data to MIDAS, meaning the individual front ends could crash and restart without affecting the entire DAQ. For example, each COPPER board ran its own front end, as did the VME crates housing the VF48 and VT48 modules.

MIDAS then concatenated and compressed the digitized data from the detectors themselves into so-called “raw data” files, and stored all the information from each front end in a text file. The raw data were then processed using a program called “proot”, developed for the experiment, which converted the raw data into ROOT [61] trees.

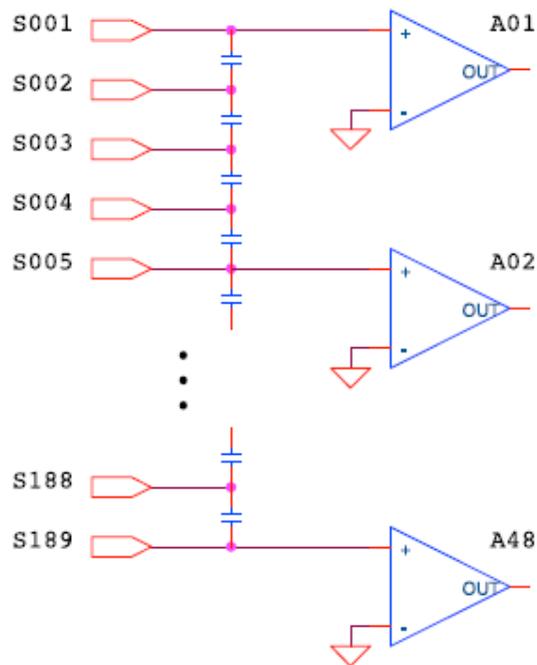


Figure 3.21: The silicon strip readout. Each strip is connected to a capacitor and there is one readout channel per four strips [8].

3.4 Event types

Many types of event could cause triggers other than single $\pi \rightarrow \mu \rightarrow e$ or $\pi^+ \rightarrow e^+ \nu_e$ decays. Some of these backgrounds could be removed with event selection cuts, while others had to be included in the time spectrum fit. In general, background processes resulting in high energy deposit in the crystal calorimeter were the most important to properly characterize, as these had the potential to be misidentified as $\pi^+ \rightarrow e^+ \nu_e$ events. The three main types of background in the PIENU experiment were beam particles, multiple muons decaying in the same event, and neutral pileup.

3.4.1 Beam-related background

The beam delivered to the detector consisted of pions, muons, and positrons, as discussed in Section 3.1. Beam positrons traversed all the plastic scintillators and were absorbed by BINA, while beam muons stopped in either T2 or BINA. Triggers caused by muons or positrons could be removed in offline analysis based on their timing, as the upstream and downstream coincidences came at approximately the same time.

Events where the upstream coincidence was caused by a pion and the trigger was subsequently completed by a muon or, especially, a positron had the potential to bias the branching ratio, as beam positrons deposited a similar amount of energy in the calorimeter as $\pi^+ \rightarrow e^+ \nu_e$ events. These events came at characteristic times relative to the pion, defined by the 43.4 ns RF period of the cyclotron. These events were removed in the offline analysis by rejecting events with extra hits in B1, B2, or T1 in a wide window around the pion. The window used was $-6.6 \mu\text{s}$ to $1.4 \mu\text{s}$, with $t = 0$ defined by the upstream coincidence signal.

Old muon pileup

The rate of particles reaching the PIENU detector was approximately one every fifteen μs , on average. Sometimes, of course, the separation between two particles was considerably less than this. Due to the long lifetime of the muon, $\tau_\mu = 2.197 \mu\text{s}$, there was a significant probability that there would be a muon already in the target when a pion arrived. These were referred to as “old muons”. The cut applied in the offline analysis removing events with extra hits in B1 and B2 mitigated this background, but did not remove

3.4. Event types

it entirely, since approximately 5% of muons took longer than $6.6 \mu\text{s}$ (the pileup inspection window) to decay.

The energy spectrum and the time spectrum for such events were both affected by the presence of the additional muon in the target. The time spectrum was affected because either muon could decay at any given time, resulting in an effectively shorter lifetime for this type of event. The energy spectrum was affected because both muons could decay within the integration time of the calorimeter, and both decay positrons could enter the crystals. This could result in enough energy being deposited to place the event in the high-energy time spectrum.

If both positrons went through T1 into BINA, the event would be removed by the pileup cut in T1, unless the two hits were too close together in time (~ 15 ns) to resolve. A larger contribution came from events where one positron went through T1 into BINA, and the other missed T1 while still entering either BINA or CsI. More solid angle existed for this to occur in CsI, but the narrower integration window of CsI meant that the energy was only added to the event if the two muon decays were less than 80 ns apart. The solid angle for BINA was smaller than for CsI, but the integration time was much longer. A diagram showing a positron emerging from the centre of the target, going past the edge of T1, going through T2, and entering BINA is shown in Figure 3.22.

The proportion of decay positrons that miss T1 and still enter BINA can be determined from Monte Carlo simulation. Figure 3.23 shows the x and y distributions in the first plane of WC3, for $\pi \rightarrow \mu \rightarrow e$ events that deposit energy in BINA, and for $\pi \rightarrow \mu \rightarrow e$ that deposit energy in BINA but not T1. The hole due to T1 in the latter plot is clearly visible. The ratio of events in which the decay positron misses T1 and deposits energy in BINA to the total number of events in which the decay positron deposits energy in BINA was approximately 10%.

Since old muons could not be removed entirely, it was necessary to include their effects in the time spectrum fit. The trigger window was extended into negative times, corresponding to events where the downstream coincidence preceded the upstream coincidence. In the low-energy time spectrum, this region was populated almost entirely by events where an old muon decayed into T1 and T2, and a pion subsequently completed the trigger. The time

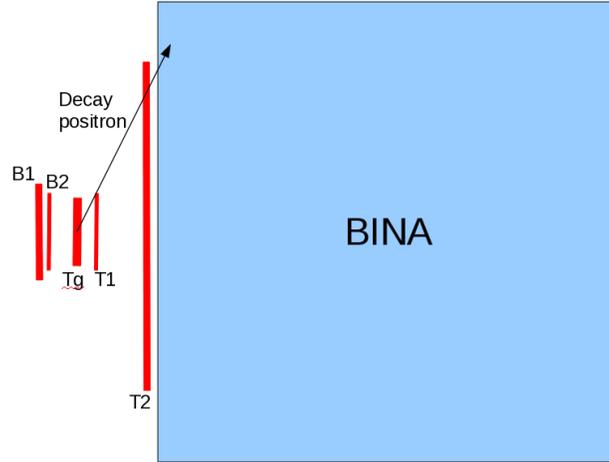


Figure 3.22: A muon decaying in the centre of the target and the decay positron going past the edge of T1. The plastic scintillators and BINA are shown; distances are to scale.

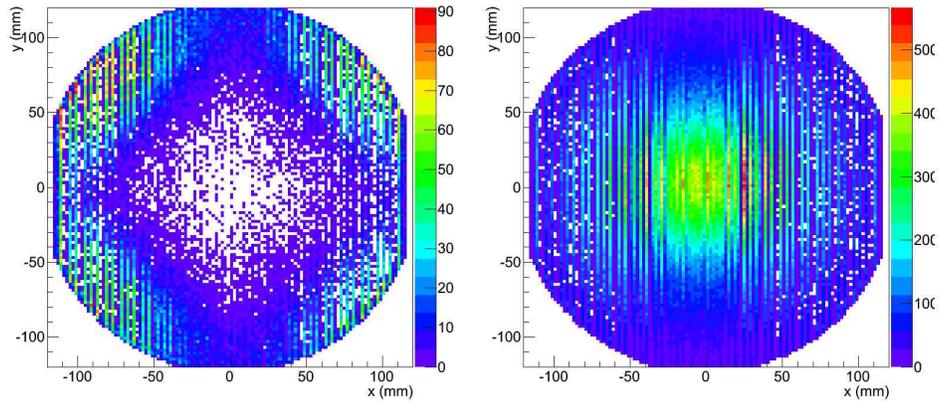


Figure 3.23: Simulated x and y distributions at the first plane of WC3 for $\pi \rightarrow \mu \rightarrow e$ events that deposit energy in BINA and not in T1 (left) and for $\pi \rightarrow \mu \rightarrow e$ events that deposit energy in BINA (right).

spectrum for this component is a simple muon lifetime, which can then be extrapolated to positive times. For events in which both muons decayed, the time spectrum was obtained by a Monte Carlo simulation which included the waveforms and integration windows of BINA and CsI. This, and the other components of the time spectrum fit, are described in detail in Chapter 5. Simplified timing diagrams for a regular $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ event, an event where only the old muon decays, and events where two muons decay can be found in Appendix B.

Neutral pileup

The neutral pileup background was mainly due to neutrons produced at the production target, which then thermalized in the shielding around the beamline before being captured in the calorimeter or its support structure, leading to gamma emission. This background was time-independent relative to the pion decay time, since it took on the order of tens of microseconds for the neutrons to make their way from the primary target to the detector. To mitigate this background, a steel wall was placed between the detector and the beamline, with a narrow opening for the beampipe, as shown in Figure 3.24.

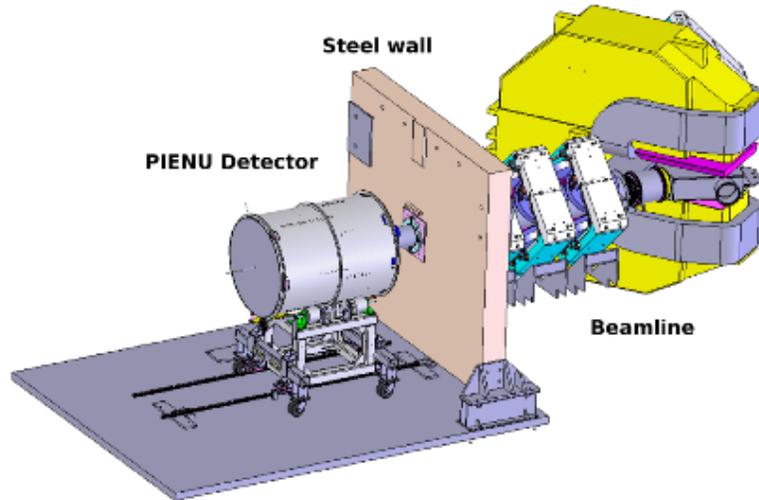


Figure 3.24: The PIENU detector and beamline after the last bending magnet, showing the steel wall used for neutron shielding.

Despite the extra shielding, some neutral pileup remained. This was included in the time spectrum fit by including a parameter representing the proportion of low-energy events promoted to the high-energy time spectrum via time-independent means. Neutral pileup was the primary mechanism for such events, though energy resolution could also contribute in principle.

3.5 Monte Carlo

The full PIENU detector geometry was implemented in GEANT4 [4], including all support structures and dead material. The parameters of the pion beam were taken from data. No waveform digitization was included in the simulation. Gaussian energy resolution was added to each detector; this resulted in good agreement for the plastic scintillators, but the response of the crystal scintillators was taken from data, as this was a crucial aspect of the analysis. The response function of BINA and CsI is described in detail in Chapter 6. Monte Carlo was used to determine the size of several other systematic effects, as explained in Chapter 7.

Chapter 4

Data Taking and Processing

4.1 Running periods

Data-taking began in April 2009 and concluded in December 2012. The TRIUMF cyclotron is shut down at least three months of each year for maintenance, between January and March, so there is a separate data set for each year. The data were divided into runs containing approximately 300,000 events, which took about ten minutes at an incident pion rate of 50-60 kHz. Table 4.1 shows the approximate number of runs taken in each year used thus far in the branching ratio analysis, as well as any special runs. In the following sections, the conditions for each running period are described.

Table 4.1: Running periods.

Year	Runs used	Special runs
2009	0	One week positron beam data
2010	2400	Eight hours per week muon beam data
2011	3600	One month positron beam data
2012	13000	One week beamline tests

4.1.1 2009

A discriminator was used to determine the pulse height of the sum of the NaI(Tl) and CsI PMTs for the BinaHigh trigger, rather than the digital module that was used later. The signal used was an analog sum without gain correction, so the threshold of the trigger was not stable, leading to a potential loss of $\pi^+ \rightarrow e^+\nu_e$ events. Also, the trigger for recording cosmic rays, used to calibrate the CsI crystals, did not exist. The calibration was instead attempted using a xenon lamp with a dedicated trigger; however, this proved inadequate.

4.1.2 2010

The final trigger was in place in 2010; however, until November, the signals from the CsI crystals were out of time with the trigger, and thus not recorded. Consequently, the largest source of systematic error in the experiment, the estimated uncertainty in the low-energy tail of the measured $\pi^+ \rightarrow e^+\nu_e$ energy spectrum, was larger by approximately a factor of two. The November 2010 data were analyzed in [3]; the uncertainty obtained on the branching ratio was 0.24%, with about equal contributions from statistics and systematics.

4.1.3 2011

The cyclotron was shut down until September due to a failure in the vacuum system. In September and part of October, data were taken to measure the response function of the crystals; PIENU data were then taken until the end of the year.

4.1.4 2012

The largest easily useable data set was recorded in 2012. However, at the start of this running period, the energy threshold of the BinaHigh trigger was lowered, to ensure that no $\pi^+ \rightarrow e^+\nu_e$ decays were being missed. This resulted in more $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ events causing BinaHigh triggers. Since these events are not used in the analysis, the number of useable events per run is lower by approximately a factor of 1.5 in 2012 compared with 2010 and 2011.

4.2 Data processing

The raw data from each run were processed into a ROOT [61] tree containing information from each detector. From the point of view of data processing, there were five separate types of detectors that must be considered: the plastic scintillators, the wire chambers, the silicon strip detectors, BINA, and the CsI array. Each of these had a different characteristic response, and different algorithms were employed.

4.2.1 Plastic scintillators

Each plastic scintillator (B1, B2, Tg, T1, and T2) was read out by four PMTs. The signal from each PMT was sent to VT48 and COPPER. De-

4.2. Data processing

spite the higher frequency of the VT48, 1.6 GHz compared to 500 MHz, the COPPER waveform gave not only energy information but also superior timing, since it can be fitted. Thus, for the plastic scintillators, VT48 was used mainly as a diagnostic check of COPPER information. Since the VT48 is a TDC, the only information available to store in the output trees was the time at which the PMT voltage went above a given threshold.

Since COPPER digitized the full PMT waveform, more sophisticated analysis was required. COPPER digitized a time window $8 \mu\text{s}$ wide, with a sample every 2 ns, so to keep the data size manageable it was necessary to zero-suppress the waveform except near hits. A hit was defined by an increase in the waveform from one sample to the next with the waveform at a level of at least three counts; the noise in the ADC was at the level of one count. For each hit, the time and height of the peak was stored, in addition to integrals of the waveform in different time regions, from a narrow region around the peak to the full unsuppressed waveform. The level of the waveform prior to the hit, referred to as the pedestal, was also stored; the zero suppression was removed approximately 30 ns before the hit.

A typical raw PMT waveform (in this case from a pion in B1) is shown in Figure 4.1. The red and black points display the values recorded by two 250 MHz ADCs. The full digitization window is $8 \mu\text{s}$ wide, and the dynamic range of the ADCs is 0 to 255. The noise in the ADCs was at the level of 1 count, and a typical pion signal in B1 was 50-60 counts.

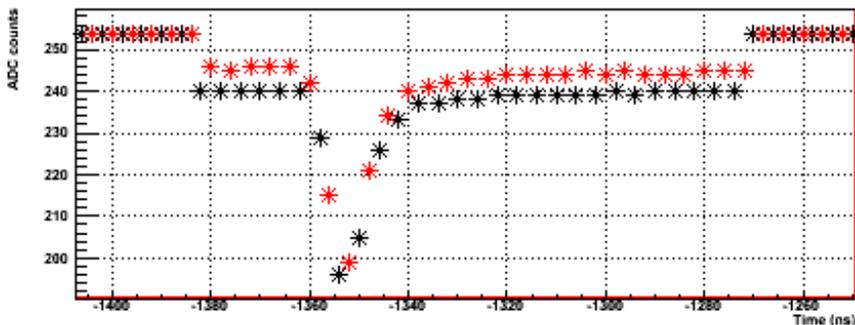


Figure 4.1: PMT waveform from a pion in B1. The points at the beginning and the end of the waveform are zero-suppressed; the drop around -1380 ns is to the level of the pedestal.

4.2.2 Wire chambers

For each wire chamber plane, the wires were bundled in groups of three, and each group was fed into one channel of the VT48; the times of the wires hit were stored in the output trees.

4.2.3 Silicon strips

Each channel of S1, S2, and S3 was digitized by the VF48, a 60 MHz ADC; channels without hits were zero-suppressed. The position of the particle was reconstructed based on a weighted average of the charge deposited in strips with hits.

4.2.4 Crystal scintillators

The scintillation light from BINA was collected by 19 PMTs, each of which was digitized by the VF48. The time and height of the peak sample were stored in the output trees, as well as integrals of both the full waveform and a narrow region around the peak. The former gives the best energy resolution, but is more affected by pileup.

The CsI crystals were read out by the VF48 as well. Similarly to the NaI(Tl) crystal, timing, pulse height, and integrated charge information were stored in the trees.

4.3 Blinding

Since the PIENU experiment aims to make a high-precision measurement which will then be compared to a precise theoretical prediction, it is necessary to mitigate potential biases in the analysis procedure. Therefore, a method for blinding the branching ratio at the stage of raw data processing was developed. The ideal blinding procedure would randomly alter the quantity being measured without affecting the data in any other way; in practice, it is necessary to make use of some quantity that is not used in the analysis but which depends on the quantity being measured. Knowing the exact nature of the dependence is not required.

In the PIENU analysis, a natural choice for a quantity to use for the blinding is the energy deposited in the target counter [62]. The presence or absence of the decay muon causes a significant difference between $\pi^+ \rightarrow e^+\nu_e$

4.4. Energy calibration

and $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events, but the time and energy resolutions of the target are not sufficient to make use of this information at the level of precision required.

In order to alter the branching ratio, an inefficiency function was applied to the target energy. An unknown factor between 0 and 0.5% of events was excluded from the analysis in a region of the target energy spectrum containing mostly either $\pi^+ \rightarrow e^+ \nu_e$ events or $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events; the region in which events were excluded was chosen randomly. The blinded events will be included in the analysis once all the event selection cuts, shapes used in the time spectrum fit, and branching ratio corrections are finalized, the blinded branching ratio is stable as parameters in the analysis are varied, and all systematic errors have been assigned. The blinding procedure does not significantly affect the branching ratio stability tests, as the target energy is almost independent of other quantities used in the analysis, and the blinding factor is small enough that the dependence that may exist is negligible at the level of precision of the experiment. A graphical illustration of the technique is shown in Figure 4.2.

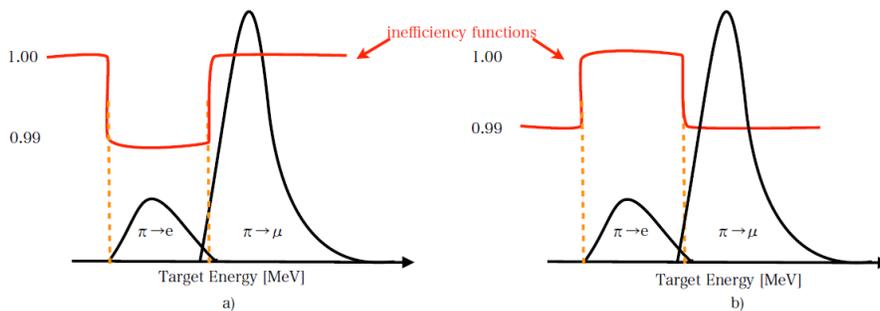


Figure 4.2: The blinding technique. Events are removed at random in one of two regions of the spectrum of energy deposited in the target counter, corresponding to either $\pi^+ \rightarrow e^+ \nu_e$ or $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events.

4.4 Energy calibration

Energy calibration in the PIENU experiment was done for each plastic scintillator as well as the crystals. The energy deposited in the plastic scintillators was obtained from the COPPER waveforms; each PMT was

4.4. Energy calibration

digitized by two 250 MHz ADCs. The offset of each ADC was obtained from the average of the first three non-suppressed samples, just prior to the hit, and the gains were aligned using the peak in the energy deposited by beam muons. The combined waveform from the two ADCs was then integrated in a window around the peak to give the energy in ADC counts. This was converted to MeV using the Monte Carlo, with a Birks' correction included [63]. The consistency of the peaks of the different particle species in each counter provided a test of the accuracy of the calibration. As an example, both the raw and calibrated spectra for the B1 counter are shown in Figure 4.3. In the calibrated histogram, only physics triggers have been included, which removes the positrons visible in the raw histogram. One run of data is shown, along with simulations of beam muons and pions, normalized to the number of events in each peak. The tail on the left is larger in the simulated muon spectrum than the data because the trigger required high energy deposit in B1.

The calibration for the other counters was done similarly, although in T1 and T2 there were only muons and positrons, since pions stopped in the target. The T1 energy from pion decay events is shown in Figure 4.4, along with the simulated spectrum. The energy deposited in the target counter was not used in the branching ratio analysis, to avoid biasing $\pi^+ \rightarrow e^+\nu_e$ or $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ events; the energy spectra of the two decay modes are significantly different, and the response of the target would have to be very well understood to make use of this information without increasing the systematic error to an unacceptable level. Furthermore, the presence of a variable with a strong dependence on the decay mode, that neither is used in the analysis nor has a strong correlation with any variable that is used in the analysis, provides a simple tool for blinding the branching ratio (see Section 4.3).

The calibration of BINA, the NaI(Tl) crystal, was done using the $\pi^+ \rightarrow e^+\nu_e$ peak and the endpoint of the $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ spectrum, with the energy deposited by beam positrons providing an additional check. Because the analysis relies on separating the events into those that deposited less than 52 MeV in the crystals from those that deposited more than 52 MeV in the crystals, the calibration of BINA and CsI was by far the most important of all the detectors. The position of the $\pi^+ \rightarrow e^+\nu_e$ peak was obtained from MC, and the consistency of the Michel edge in data and simulation was used to assign an uncertainty in the energy calibration. The position of

4.4. Energy calibration

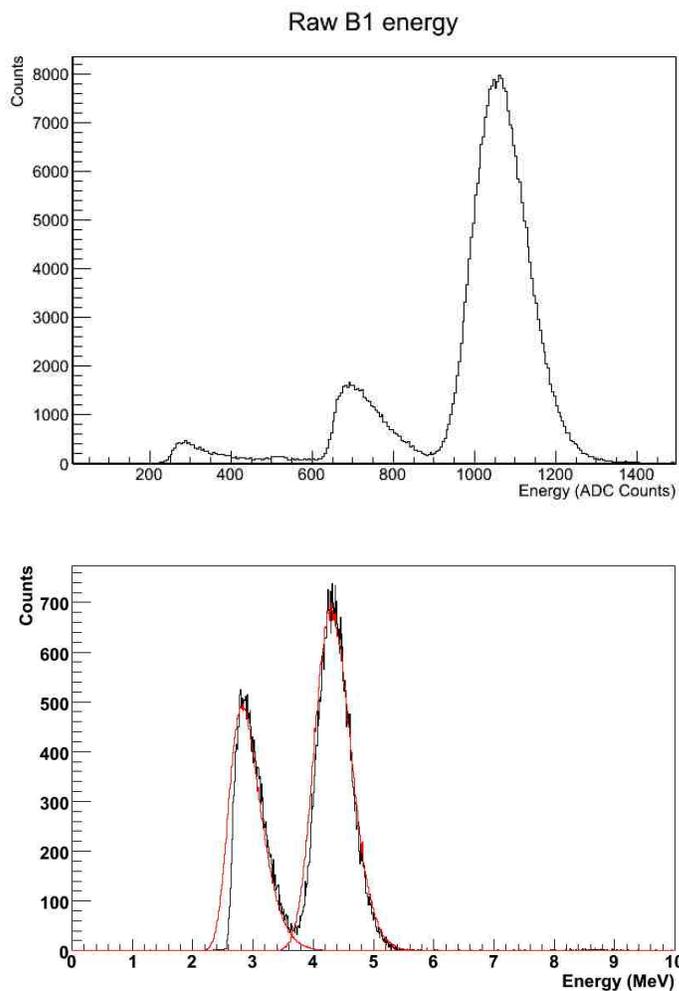


Figure 4.3: The energy deposited in B1 before (upper panel) and after (lower panel) calibration. In the calibrated histogram the MC spectrum is shown in red, and a cut has been made to remove events due to calibration triggers. The difference around 2.5 MeV is due to the requirement of high energy deposit in B1 for physics triggers, which was not included in the MC.

4.4. Energy calibration

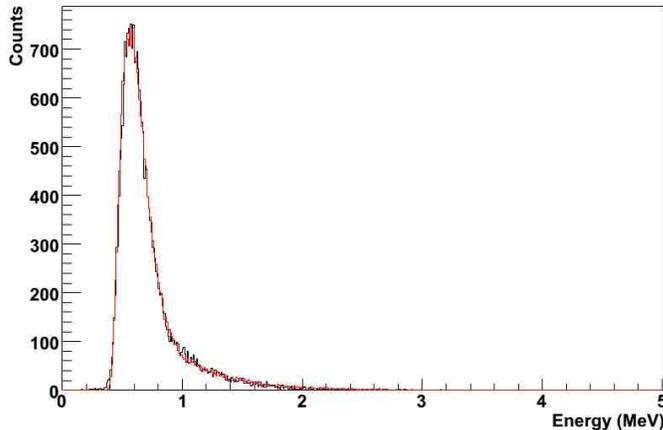


Figure 4.4: Calibrated energy deposited in T1. The MC spectrum is shown in red.

the Michel edge was obtained by fitting the spectra with a convolution of a unit step function with a Gaussian representing the energy resolution. The result of this convolution is an error function:

$$f(E) \propto \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{E - a}{\sigma\sqrt{2}} \right) \right]. \quad (4.1)$$

Here a is the value at which the unit step transitions from 0 to 1, and σ is the standard deviation of the Gaussian resolution function. The mean of the resolution is assumed to be 0. Figure 4.5 shows the measured energy spectrum in BINA + CsI zoomed in on the $\pi^+ \rightarrow e^+\nu_e$ peak for the 2012 data, as well as a simulation of $\pi^+ \rightarrow e^+\nu_e$ positrons in red. Both distributions are fitted with Gaussians. Figure 4.6 shows the measured and simulated energy spectra for $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ positrons, fitted with Eqn 4.1.

When this was done for the 2010 data, the energy calibration was found to be accurate to 0.1 MeV [62]. The results for the edge of the Michel spectrum in data and MC were 48.50 MeV and 48.47 MeV; for the $\pi^+ \rightarrow e^+\nu_e$ peak, the results were 65.49 MeV and 65.38 MeV. Thus, the uncertainty on the energy calibration between 48 and 65 MeV is again taken as 0.1 MeV. The full calibrated energy spectrum is shown in Figure 4.7.

4.4. Energy calibration

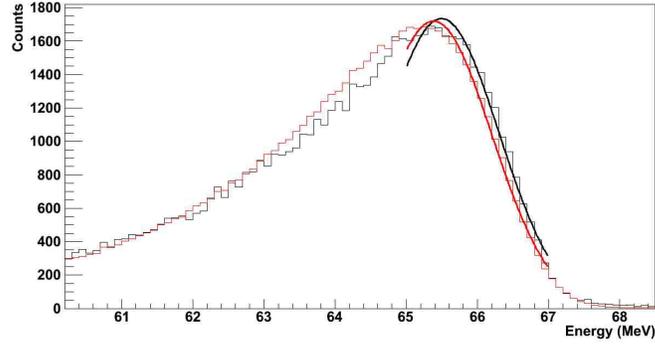


Figure 4.5: Energy deposited in BINA + CsI for $\pi^+ \rightarrow e^+ \nu_e$ events. The black histogram, fitted with the black curve, is data, and the red histogram, fitted with the red curve, is MC. The fits are Gaussian; the fitting range is asymmetrical about the peak because the region to the left of the peak is not Gaussian, due to shower leakage.

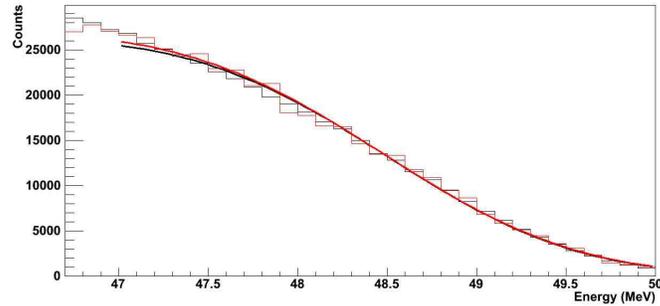


Figure 4.6: Energy deposited in BINA + CsI for $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events near the Michel edge. Data is shown in black and MC is shown in red. The lines are fits using Equation 4.1.

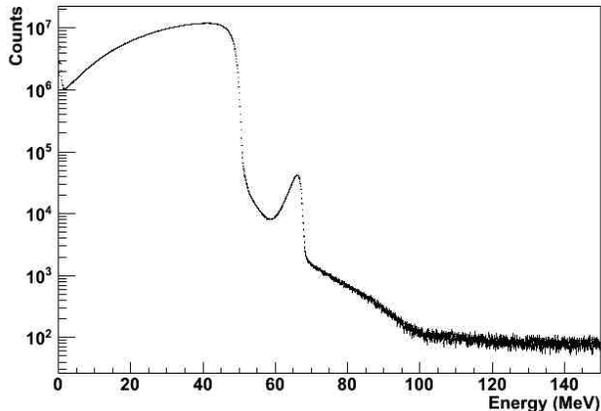


Figure 4.7: The calibrated BINA + CsI spectrum. The spectrum up to 50 MeV is dominated by $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events, and the peak at approximately 65 MeV is due to $\pi^+ \rightarrow e^+ \nu_e$ events. The high-energy tail is mainly due to pileup events, with a small contribution from $\pi^+ \rightarrow e^+ \nu_e \gamma$ events.

4.5 Track reconstruction

The incoming pion track was reconstructed using the beam wire chambers, WC1 and WC2, and the silicon strips near the target, S1 and S2. Each pair was used independently to reconstruct the track, to reduce the impact of multiple scattering in the intervening detectors. Each wire chamber contained 3 planes, each sensitive to a single coordinate; the silicon strips each contained 2 orthogonal planes. The tracking algorithm divided the wire chambers and silicon strips into three groups, called trackers. The first tracker, WC12, contained the beam wire chambers, WC1 and WC2. The second tracker, S12, contained the two silicon strips before the target, S1 and S2. The third tracker, S3WC3, contained the silicon strip after the target and the wire chamber near the crystal face, S3 and WC3. S3WC3 was used to reconstruct the decay positron track.

For a given event in a given tracker, the locations of the hits in each plane were fitted to a straight line. The position of the hit was taken as the centre of the plane. The parameterization used is given in Equations 4.2 and 4.3; z is the coordinate in the direction of the beam, x is horizontal, and y is

4.5. Track reconstruction

vertical. t_x and t_y are the ratios of the x and y momenta to the z momentum, and x_0 and y_0 are the x and y coordinates at $z = 0$, the centre of the target. The tracks are parameterized according to the following equations:

$$x = t_x * z + x_0 \quad (4.2)$$

and

$$y = t_y * z + y_0. \quad (4.3)$$

The reconstructed y vs. x position at the centre of WC1 using the WC12 tracker is shown in Figure 4.8, and the t_x and t_y distributions are shown in Figure 4.9. The same quantities reconstructed by the S12 tracker are shown in Figures 4.10 and 4.11. These distributions are used in the Monte Carlo simulation to provide realistic beam parameters for the incoming pions.

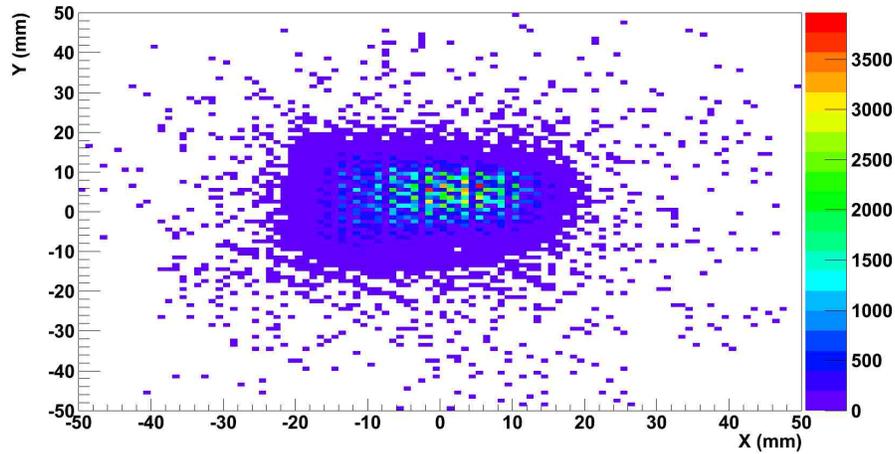


Figure 4.8: The beam spot reconstructed by the WC12 tracker.

4.5. Track reconstruction

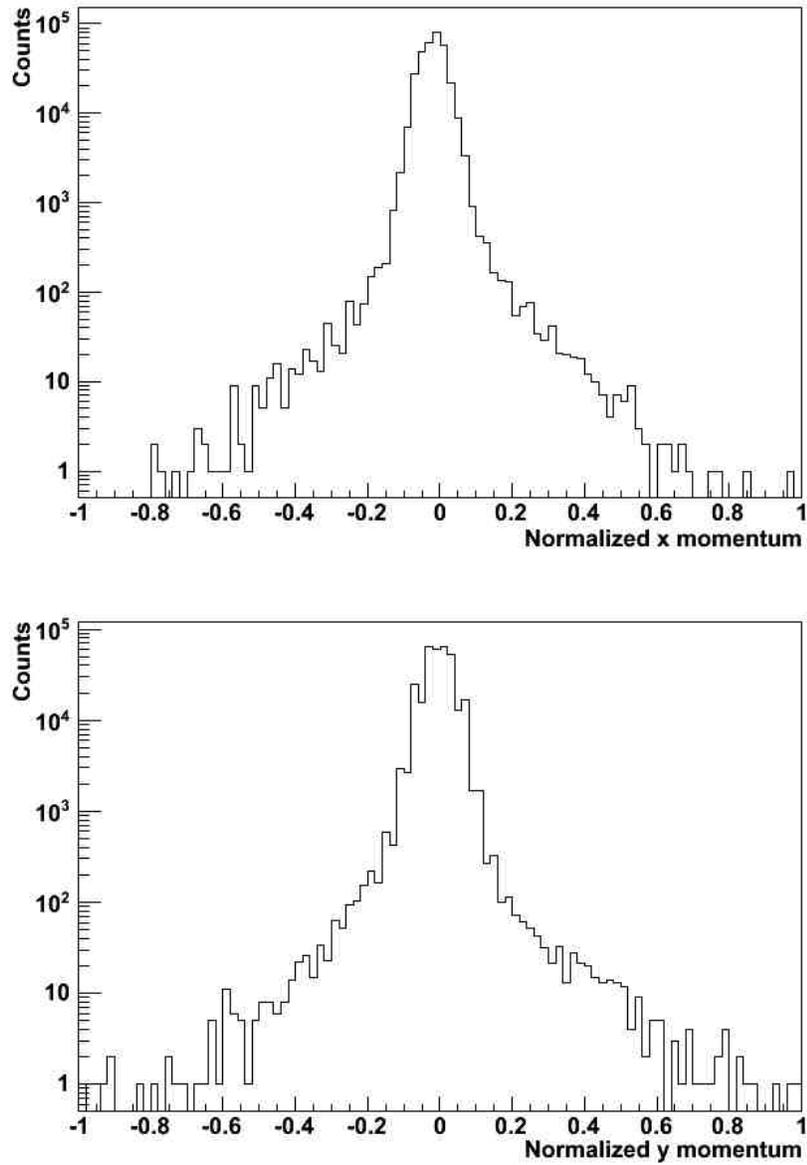


Figure 4.9: Ratios of the x and y momenta to the z momentum (t_x and t_y), reconstructed by the WC12 tracker.

4.5. Track reconstruction

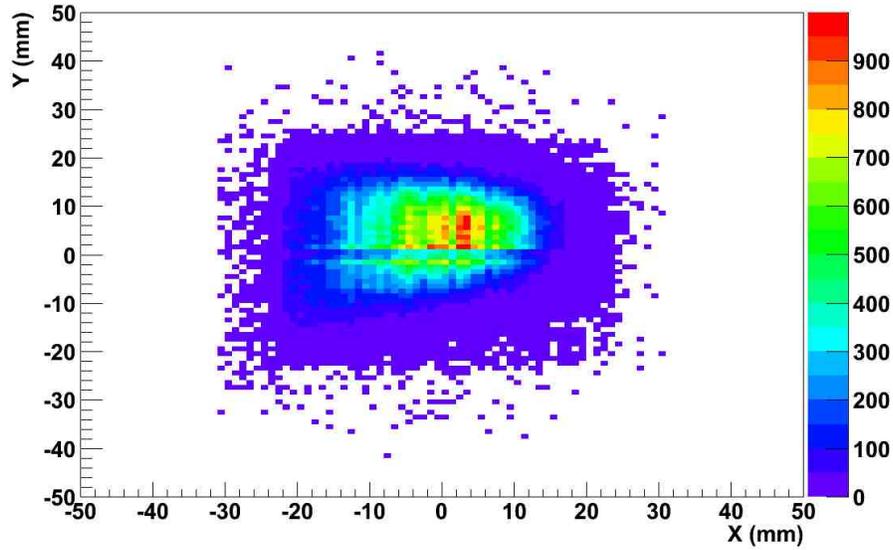


Figure 4.10: The beam spot reconstructed by the S12 tracker.

The decay positron track was reconstructed using S3 and WC3. The y vs x position in the middle of WC3 and the t_x and t_y distributions are shown in Figures 4.12 and 4.13, showing the positrons emerging isotropically rather than the focussed beam seen upstream of the target.

4.5. Track reconstruction

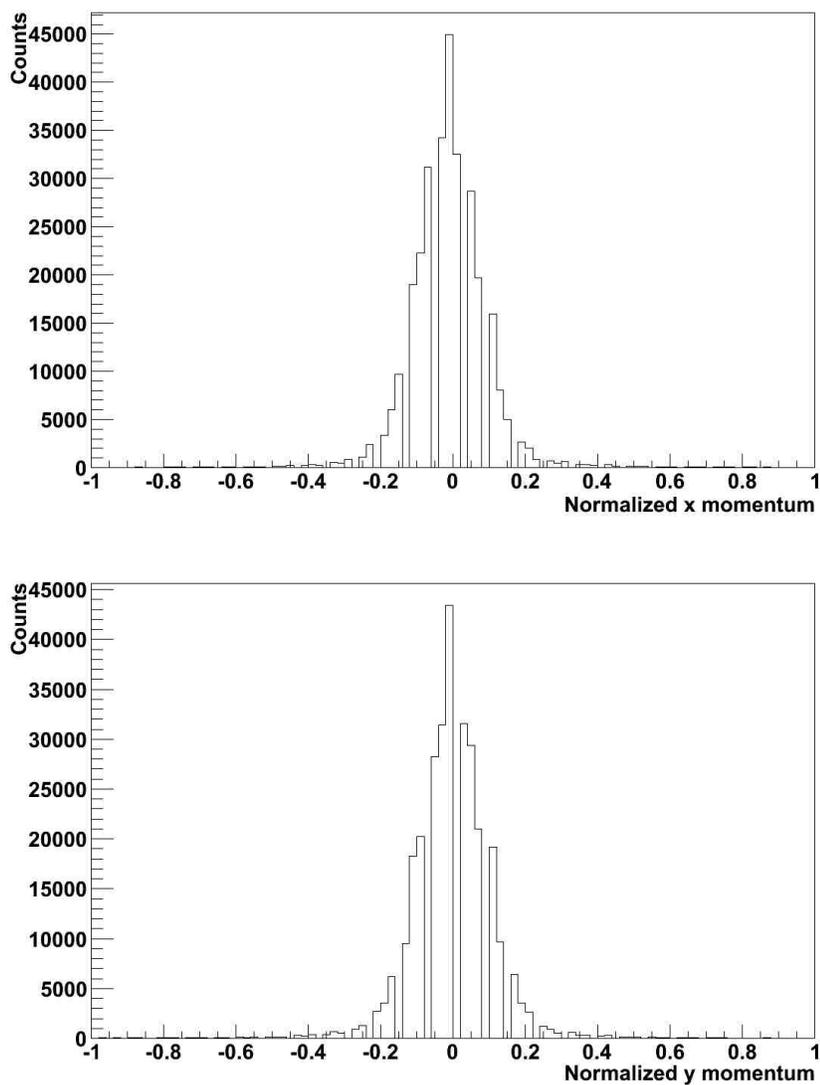


Figure 4.11: Ratios of the x and y momenta to the z momentum (t_x and t_y), reconstructed by the S12 tracker. The gaps are due to the track reconstruction algorithm, which uses the centre of the plane hit as the position, leading to some values for t_x and t_y never occurring.

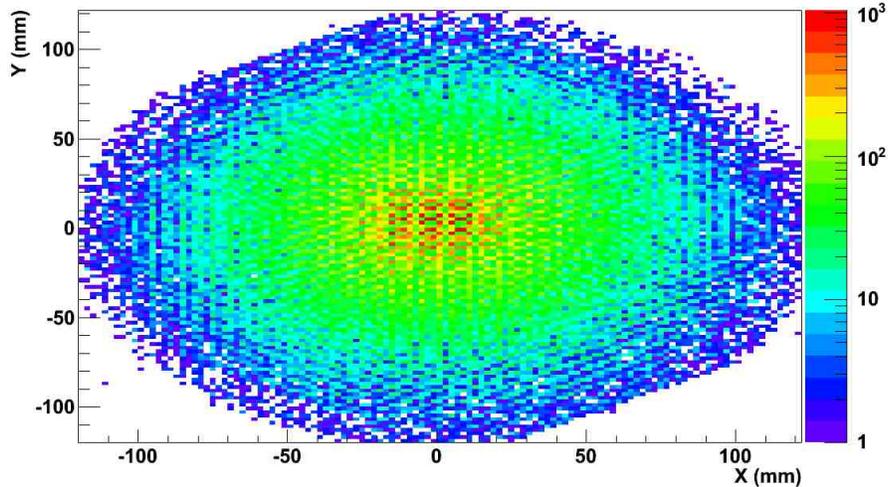


Figure 4.12: Decay positron position at the centre (along z) of WC3, reconstructed by the S3WC3 tracker.

4.6 Data stability

Due to the length of time over which data were taken, compensating for the variation of several experimental parameters was essential. Although the detector enclosure and electronics racks were maintained at roughly constant temperature and humidity, the PMT gains and thresholds, and the offsets in the trigger logic, still varied considerably over the running period. The cyclotron current was often unstable, leading to variations in the beam rate. The dipole magnets were monitored with nuclear magnetic resonance (NMR) probes, which were used to automatically adjust the current to maintain a constant field, so the beam momentum was largely stable. However, this was not done for the quadrupole magnets, meaning that the exact location of the beam was potentially variable.

Although the branching ratio analysis is, in principle, insensitive to most of these changes, the energy measured by the crystals is of crucial importance. Therefore, automatic run-by-run gain correction was performed; in the case of BINA, the beam positron peak in each tube was aligned with a reference run prior to being included in the calibrated sum. This was done

4.6. Data stability

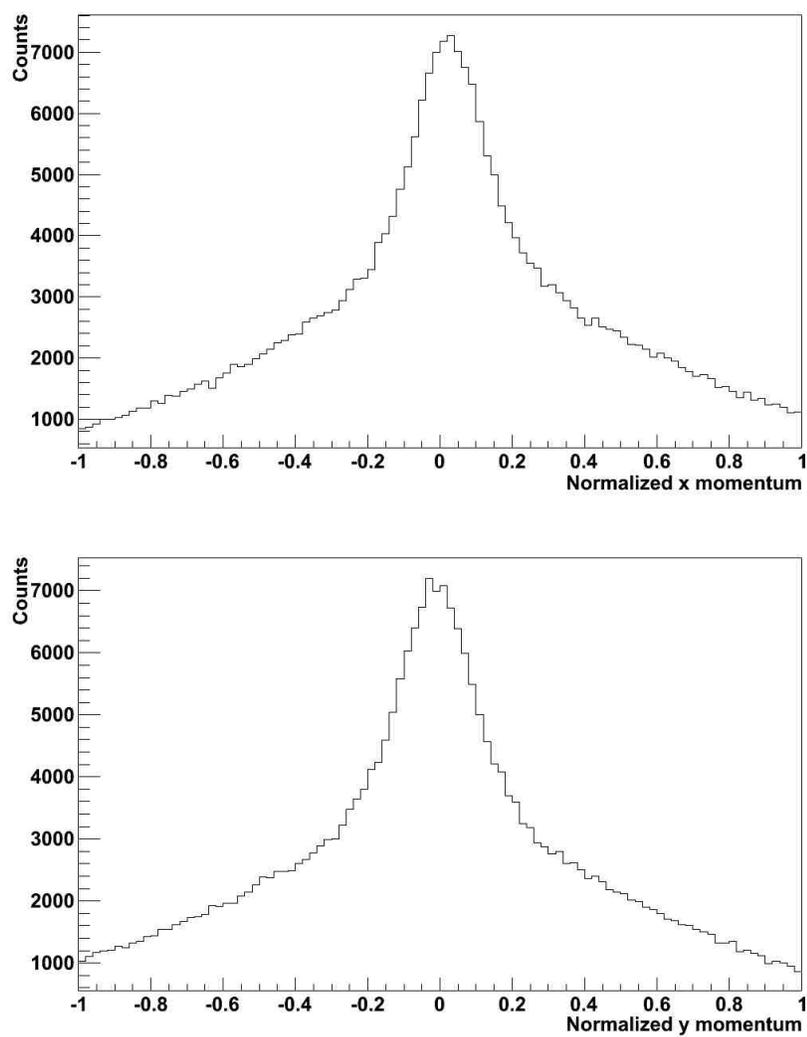


Figure 4.13: Ratios of the x and y momenta to the z momentum (t_x and t_y) for decay positrons, reconstructed by the S3WC3 tracker.

4.7. Timing

at the stage of converting the raw data into a tree. Similar procedures were implemented for the plastic scintillators as well, based on the peaks due to either beam muons or positrons. Figure 4.14 shows the peak of the pulse height distribution of beam positron events in the first BINA tube over the 2012 running period. Significant variation was observed, particularly when the beam resumed after having been off. The gain changes were corrected for on a run-by-run basis, for each PMT individually. Since the runs were short, no significant gain changes within a run were observed, so this level of stability was acceptable for the PMTs.

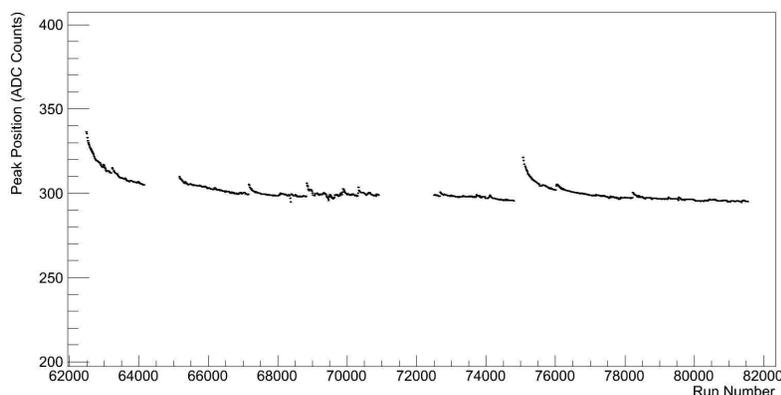


Figure 4.14: The run by run variation in the pulse height of the beam positron peak in one BINA PMT. Similar variations were observed for the other PMTs.

Figure 4.15 shows the peak of the reconstructed pion stopping position distribution over the same period. The stopping position refers to the z coordinate, and thus corresponds directly to the beam momentum. It is obtained from the point of closest approach of the tracks reconstructed by S12 and S3WC3. The gaps are due to regions of bad data.

4.7 Timing

The primary method used to extract timing information was to fit the waveforms produced by the plastic scintillator PMTs, digitized by COPPER.

4.7. Timing

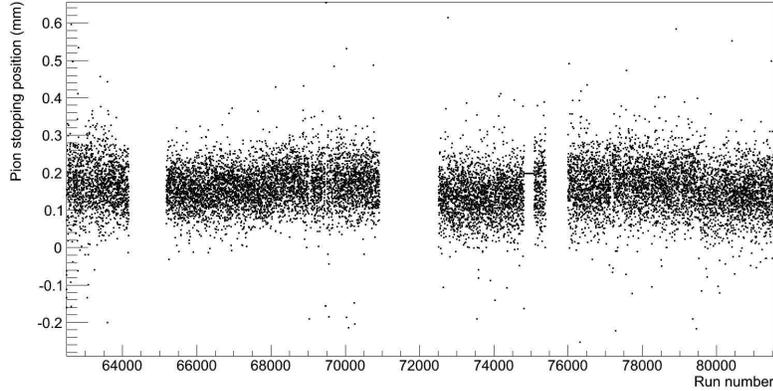


Figure 4.15: The run by run variation in the pion stopping position along the z axis.

Through this technique, sub-ns resolution was obtained. The pion time was taken from the B1 waveform, and the positron time from the T1 waveform. The fitting function for each counter used was a template constructed from the average of many waveforms from that counter; a typical fitted pulse is shown in Figure 4.16. Figure 4.17 shows the time difference between fitted pulses in two of the T1 PMTs, fitted with a Gaussian; the resolution is given by the width of the distribution, which is $\sigma = 0.49$ ns.

The trigger signal going to COPPER comes from the counters upstream of the target, and thus carries the pion time, meaning the B1, B2, and Tg pulses corresponding to the triggering particle are fixed within the readout window. The time of the fitted pulse in B1, averaged over the four PMTs, is shown in Figure 4.18; the width of the distribution is due to the variable arrival time of the trigger signal relative to the gate signal. The reduced χ^2 distribution (χ^2 per degree of freedom) of the waveform fit is shown in Figure 4.19. The number of degrees of freedom varied depending on the number of unsuppressed samples in the waveform, but was typically around 50.

By contrast, the timing of the T1 and T2 pulses, caused by decay positrons, varied within the COPPER window. To account for the variation in the arrival time of the trigger signal, the B1 time was subtracted from the T1

4.7. Timing

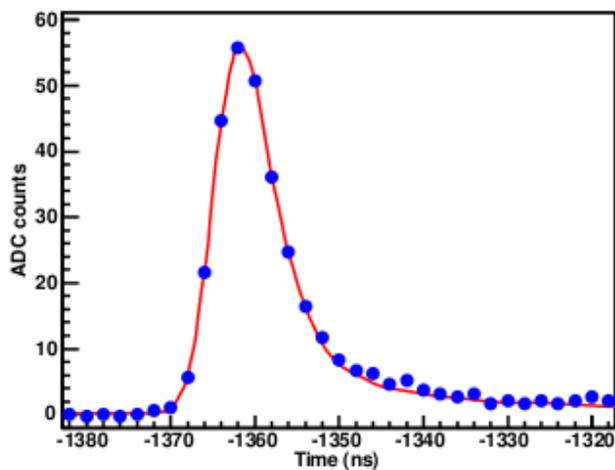


Figure 4.16: COPPER waveform fitted to a template.

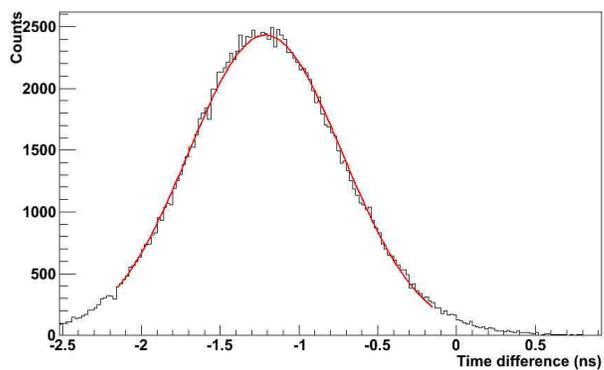


Figure 4.17: Time difference between two T1 PMTs. The red line is a Gaussian fit.

4.7. Timing

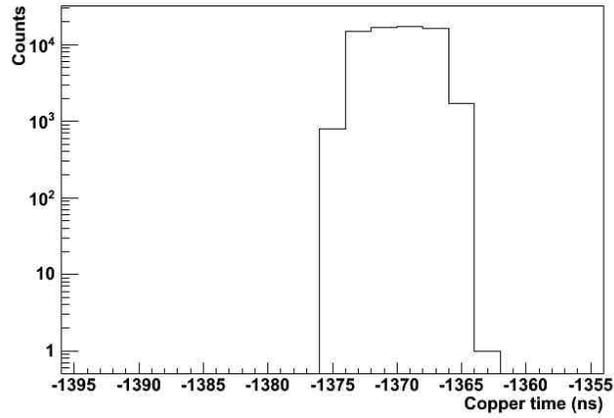


Figure 4.18: Time of the triggering pulse in the B1 counter.

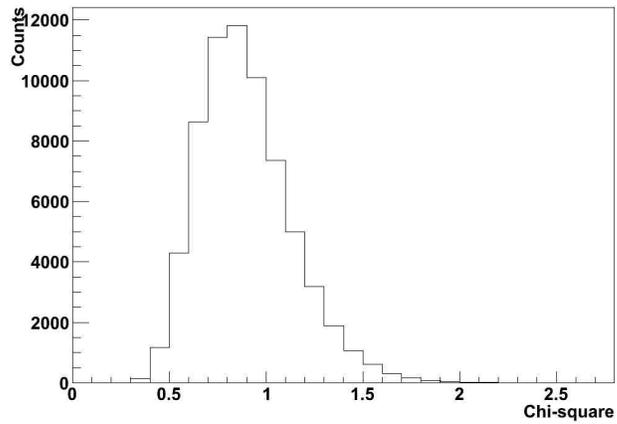


Figure 4.19: Reduced χ^2 distribution of the waveform fit in the B1 counter.

4.7. Timing

time. This quantity is referred to as T_{pos} , and is what is ultimately fitted to obtain the branching ratio. The T1 time and χ^2 distributions are shown in Figures 4.20 and 4.21. The time spectrum shows the size of the trigger window, both before and after $t = 0$, the time of the pion stop. The large peak at $t = 0$ is due to beam muons and positrons traversing the entire detector.

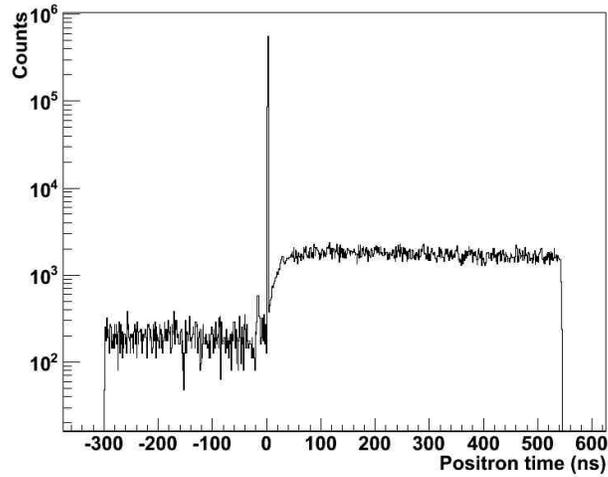


Figure 4.20: Time of the triggering pulse in the T1 counter.

Prior to any event selection, many components exist in the T_{pos} spectrum other than the $\pi^+ \rightarrow e^+ \nu_e$ and $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ signals. Some of these components can be removed or reduced with cuts, whereas some must be modelled and included in the time spectrum fit; this is the subject of the next chapter.

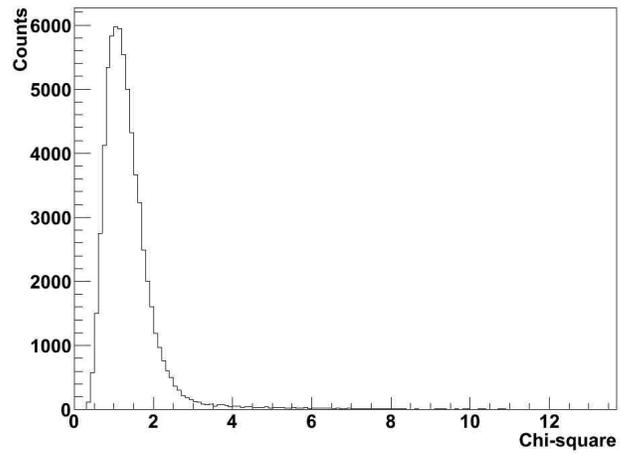


Figure 4.21: Reduced χ^2 distribution of the waveform fit in the T1 counter.

Chapter 5

Raw Branching Ratio Extraction

5.1 Event selection

5.1.1 Good run selection

The first step in the analysis is the removal of runs in which some aspect of the detector, trigger, or DAQ was malfunctioning, or which were not regular data runs. This was done by eliminating the runs with errors based on the logs kept by MIDAS (see Section 3.3), which indicated errors recorded by any of the front end programs monitoring the hardware. Furthermore, runs whose duration fell outside the normal range were excluded; this removed the majority of test runs and runs from periods with no beam, during which cosmic ray data were taken. Finally, the remaining runs taken under special conditions, mainly for evaluating systematic effects, were removed. The entire 2012 data set contains approximately 20,000 runs; 13,210 acceptable data runs were included in the analysis.

5.1.2 Time spectrum

The time of each event is defined as the time of the pulse in the T1 counter minus the time of the pulse in the B1 counter. The events were divided into high and low energy regions based upon the energy deposited in the crystals. The threshold is normally set at 52 MeV, and is referred to as E_{cut} . Both time spectra for the full 2012 data set are shown in Figure 5.1, excluding calibration triggers but with no other cuts.

Several important features in each spectrum can be clearly seen. The large peak at $t = 0$, referred to as the “prompt peak”, is primarily due to beam particles traversing the entire detector; muons, in particular, often fulfilled the trigger requirement of pion-like energy deposit in B1. In the high energy spectrum the same peak is due to beam positrons, which tended to deposit

5.1. Event selection

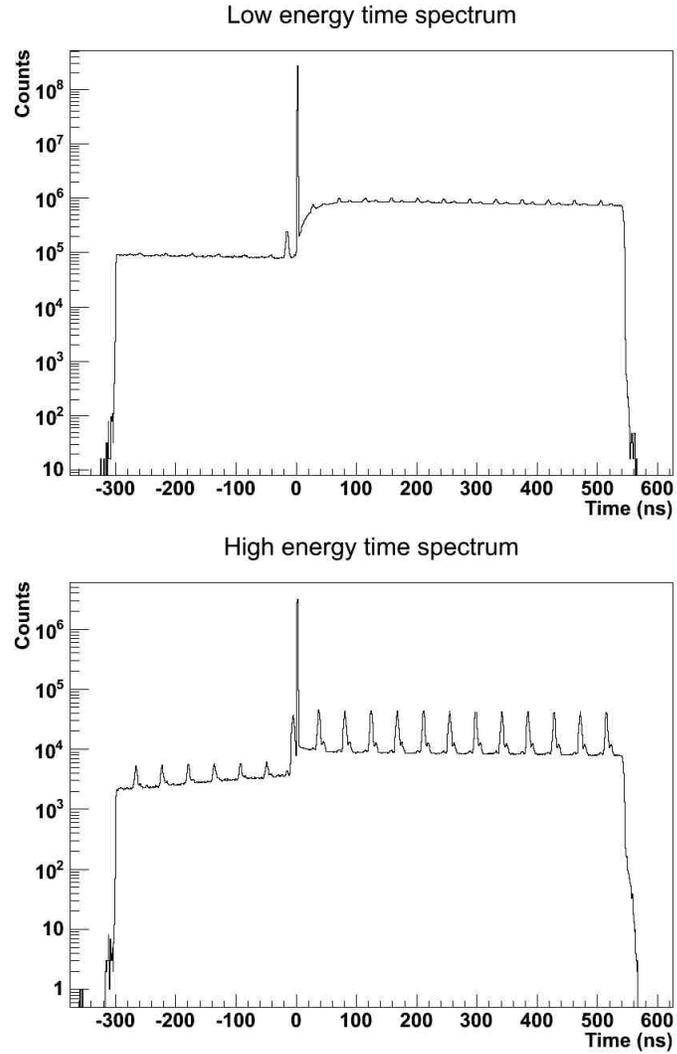


Figure 5.1: The time spectra for low energy (upper panel) and high energy (lower panel) events. $t = 0$ is defined by the arrival time of the pion. The repeating peaks are due to beam particles and are separated by the cyclotron RF period (see text).

5.1. Event selection

about 70 MeV in BINA. The repeating peaks (obvious in the high energy spectrum but also present in the low energy spectrum) were caused by a pion at $t = 0$, followed by a muon (low energy) or a positron (high energy) from a later beam spill. The peaks are separated by approximately 43 ns, which corresponds to the cyclotron frequency of 23 MHz.

The trigger window was extended into negative times as well, corresponding to events where the coincidence in T1 and T2 came before the coincidence in B1, B2 and Tg. Peaks from beam particles can be seen in this region; these occurred either because the beam particle did not deposit enough energy in B1 to trigger, or, more probably in the low energy time spectrum, did not trigger due to the prescale factor of 16. In the high energy spectrum, there is no prescaling, but the beam particles in question were positrons, which did not usually deposit enough energy in B1 to trigger.

5.1.3 Pion selection

Triggers due to beam particles other than pions can be removed from the time spectra with simple cuts. The prompt events can be removed altogether; i.e., when the actual time spectrum fit is performed, the region around $t = 0$ is excluded. Pileup events involving beam particles can be removed by cuts on the number of hits in B1 and B2. This is done by requiring that at least one phototube (out of the four) for each scintillator recorded only one hit, in a window extending approximately $2.2 \mu\text{s}$ before the pion time. The efficiency of the PMTs for real particles hitting the scintillators was very high, $> 99\%$, and due to frequent fake hits (see “T1Pileup cut” in Section 5.1.4) a tighter cut on the number of hits would remove an unacceptably large number of events. The distribution of hits in B1 and B2 is shown in Figure 5.2. The four PMTs are averaged for each counter. One hit in each PMT is the most probable case, but occurs only about half the time. The presence of fake hits can be inferred from the many events that have an extra hit in only one tube.

Two other cuts were done to ensure the selection of only events in which the triggering particle was a pion: the selection of events inside the beam spot in wire chambers 1 and 2, and the selection of events with pion-like energy deposit in B1 and B2. Although the energy deposit requirement for B1 was present in the trigger, an offline version was implemented as well. This was based on the calibrated energy determined from the COPPER waveform, whereas in the trigger it was done by placing a threshold on

5.1. Event selection

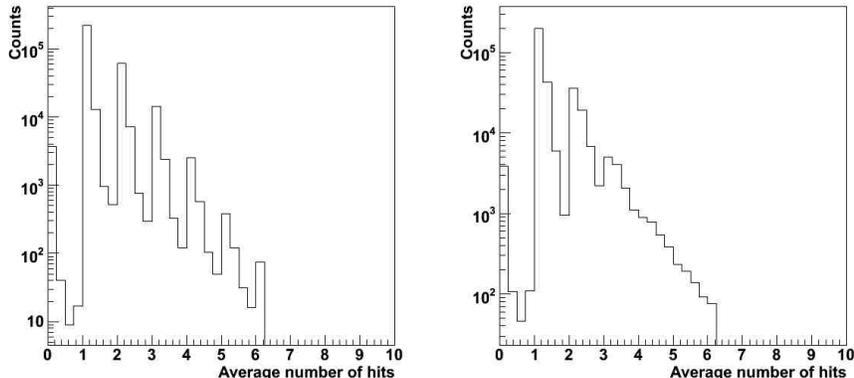


Figure 5.2: Number of hits in B1 (left) and B2 (right). The average of four PMTs is taken. The peaks at whole numbers are due to pileup events; the events with extra hits in only some of the tubes are due to noise (see text).

the raw output of the B1 PMTs, and thus included no gain correction. The distributions on which the cuts were applied are shown in Figures 5.3 and 5.4, with red lines indicating the cut values. The time spectra following all of these cuts are shown in Figure 5.5.

No cuts have been applied to the distributions in Figures 5.1 to 5.4. The three largest peaks in Figure 5.3 are caused by pions, muons, and positrons, respectively. The high energy peak is caused by events in which two pions arrived simultaneously. The requirement in the trigger of large pulse height in B1 is responsible for the difference between the shape of the left-hand side of the muon peak in the two counters (the positrons are present because of a special trigger for selecting them, for calibration purposes).

5.1.4 Pileup rejection

Following the B1B2 pileup, prompt, WC12, and B1B2 energy cuts, events where the trigger involved beam muons or positrons are reduced to a negligible level. Only events with a triggering pion at $t = 0$ survive. However, in a significant fraction of events, a muon was left over from a previous event, due to its long lifetime. This could occur either in the target from a previous pion decay or from a beam muon, which would typically stop in either T2 or the front face of BINA. Events with additional muons had a different time

5.1. Event selection

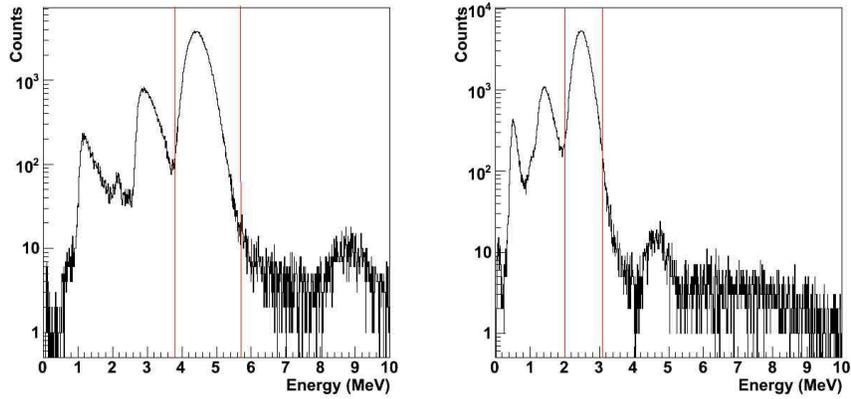


Figure 5.3: Energy deposited in B1 (left) and B2 (right). The three large peaks in each spectrum are, from left to right, caused by positrons, muons and pions. The smaller peaks are due to events with two particles. The red lines indicate the cut values.

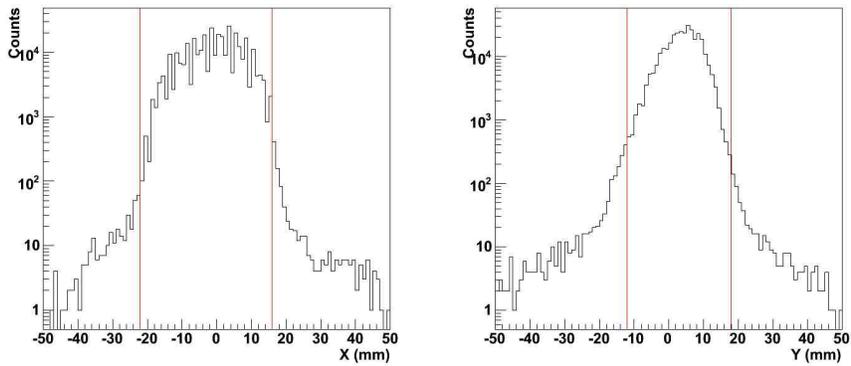


Figure 5.4: X (left) and Y (right) position at the centre of WC1. The red lines indicate the cut values.

5.1. Event selection

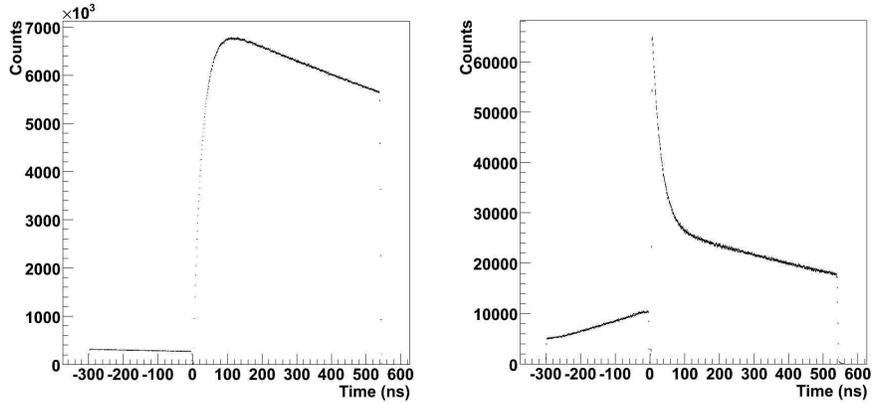


Figure 5.5: The time spectra after pion selection. The left-hand plot shows the low energy time spectrum and the right-hand plot shows the high energy time spectrum.

distribution than those without, since either the muon originating from the primary pion or the pileup muon could potentially decay, causing a trigger; also, if both muons decayed, the event could be boosted from the low energy time spectrum to the high energy spectrum.

Prepileup cut

Most of these “old muons” were rejected by requiring no hits in any of the plastic scintillators in the time range from $6.6 \mu\text{s}$ to $2.2 \mu\text{s}$ before the pion time. $6.6 \mu\text{s}$ corresponds to three muon lifetimes, in which time 95% of muons will decay. The high energy time spectra with and without this cut, referred to as the prepileup cut, are shown in Figure 5.6. The number of events at negative times is much reduced, as expected.

T1 pileup cut

The remaining old-muon background was mitigated by a further pileup cut in T1 similar to the pileup cuts for B1 and B2 described above; that is, at least one PMT was required to have only one hit in a time window extending $2.2 \mu\text{s}$ before the pion stop and $1.3 \mu\text{s}$ after. Since the decay positron could emerge at any time up to 540 ns after the pion stop, the pileup rejection window after the trigger depended on the decay time. This

5.1. Event selection

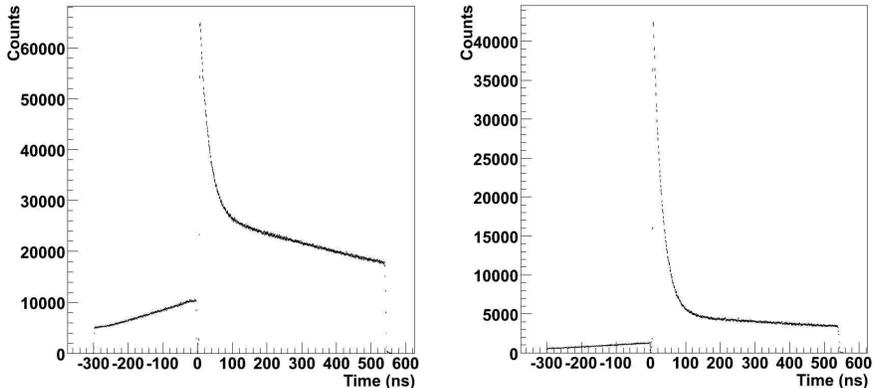


Figure 5.6: High-energy time spectrum without (left) and with (right) the pileup cut.

was potentially problematic due to the presence of fake hits after the real hit in the PMTs, which were observed to occur frequently. Generally, they came soon after the initial hit (within 50-100 ns), but occurred sometimes beyond 1 μ s after the original pulse. The source of the early fake hits was assumed to be in the amplifier or elsewhere in the electronics, but this could not explain fake hits on the time scale of μ s.

PMT afterpulsing is a well-known source of noise in these detectors [64]; one mechanism that has been identified is the ionization of helium, which diffuses into the PMTs over time, by the photoelectrons. The ionization can then be amplified and cause a measurable signal. This cut therefore had to be refined, or it would have the potential to preferentially reject positrons from earlier decays. This could in turn bias the branching ratio, as virtually all pions decay within 100 ns, but most muons do not.

To illustrate the issue, Figure 5.7 shows the ratio of the pulse heights of the first and second hits in one of the T1 PMTs. The ratio is generally close to one for real pileup events, and often much greater than one for events with fake hits. Figure 5.8 shows the time difference between the two pulses, in the region to the right of the line in Figure 5.7, which is populated almost entirely by fake hits.

5.1. Event selection

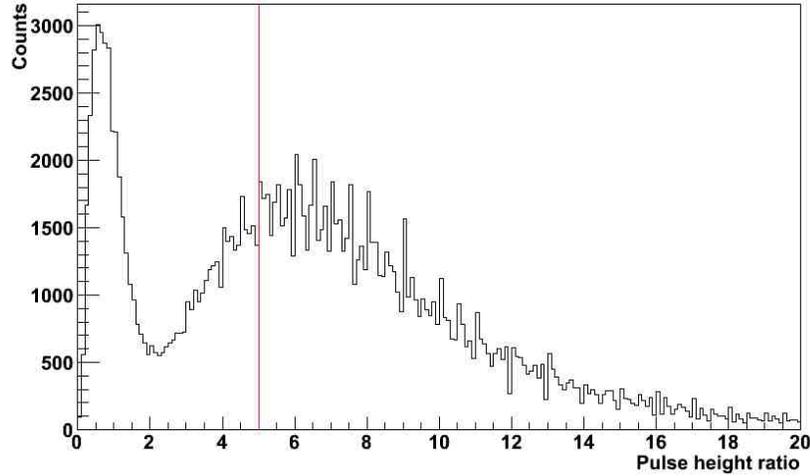


Figure 5.7: The pulse height of the first hit in one of the T1 PMTs divided by the pulse height of the second hit, if one was present. The red line indicates the cut used to select events with fake hits, for plotting the time difference between the initial hit (the real hit) and the fake hit (Figure 5.8).

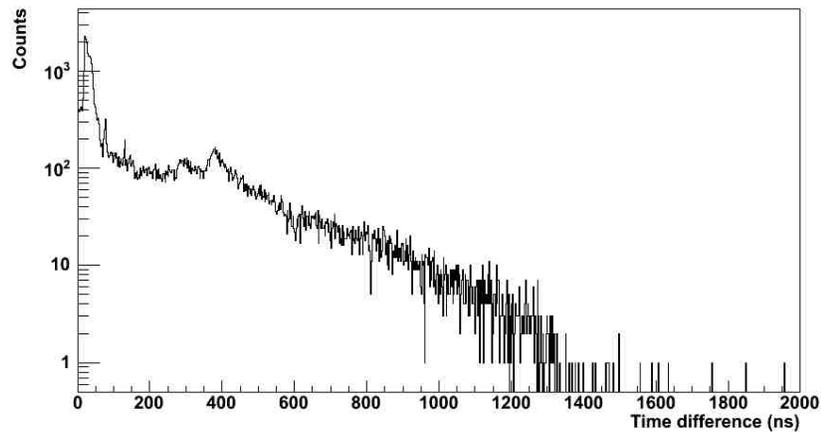


Figure 5.8: The time difference between the first and second hits in one of the T1 PMTs, for events with a small second pulse.

5.1. Event selection

For the latest positrons accepted by the trigger, the post-pileup rejection window is approximately 800 ns, as compared with approximately $1.3 \mu\text{s}$ for the earliest. The presence of fake hits in this time range raises the possibility that the T1 pileup cut will preferentially reject earlier decays, unless an additional requirement is imposed to ensure real pileup. This was done via the ratio of the full integrated charge to the fitted pulse height of the triggering hit. This ratio is shown in Figure 5.9, as a function of the pulse height, for events with multiple hits in every tube.

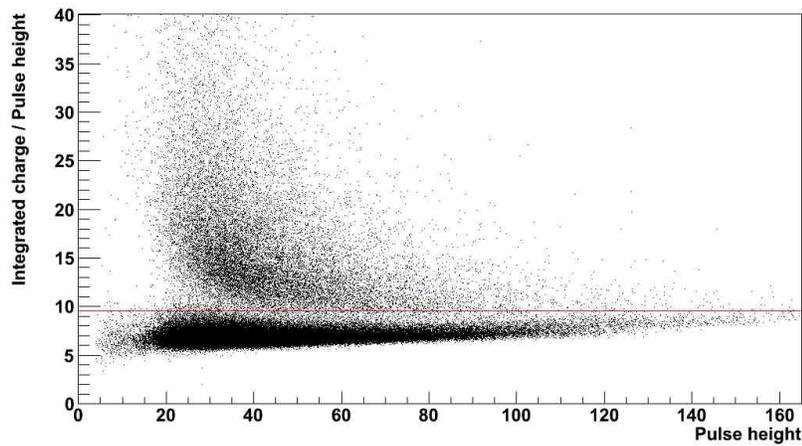


Figure 5.9: The ratio of integrated charge in the T1 PMTs to the fitted pulse height as a function of the fitted pulse height. The red line indicates the cut used to separate real pileup from pileup due to fake hits.

The fake hits and real pileup separate clearly into two bands; by only rejecting events where the ratio of integrated charge to pulse height is higher than the red line, only real pileup will be removed, and events with fake pileup will be preserved. This prevents the probability of being rejected by the pileup cut from depending on the positron decay time. The number of events below the red line in each PMT, divided by the number of events with at least one hit, gives the probability of fake hits for that PMT. The probabilities are shown in Table 5.1, as well as the probability of fake hits in every PMT. This is much higher than the product of the individual probabilities would suggest, indicating that the probabilities are correlated. This may simply be due to the fact that large initial pulses are more likely to

5.1. Event selection

cause fake hits. The time spectra after this cut are shown in Figure 5.10.

Table 5.1: The probability of fake hits after the real hit, in each T1 PMT.

Tube	Probability of fake hits
T1.1	20.7%
T1.2	15.5%
T1.3	15.1%
T1.4	17.5%
All	2.0%

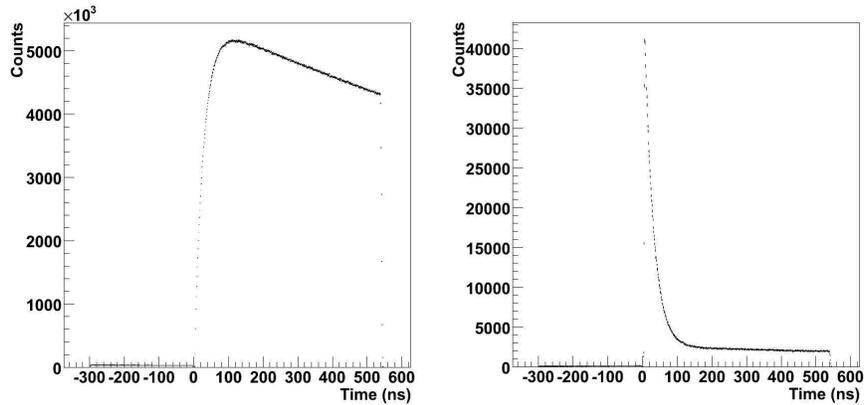


Figure 5.10: Time spectra after prepileup and T1 pileup cuts. The left-hand plot shows the low energy time spectrum and the right-hand plot shows the high energy time spectrum.

T2 pileup cut

A pileup cut can be performed in T2 as well; some solid angle exists for a decay positron emerging from the target to miss T1 but go through T2 and into BINA. Furthermore, beam muons typically stop in T2 or the front face of BINA. The T2 waveforms are substantially noisier than the T1 waveforms, due to the fibre readout of T2, and they are not fitted in the analysis, so distinguishing real pileup from fake pileup is more difficult. The dependence of the probability of fake hits in T2 on the energy deposited in

T2 also presents a more serious problem than in the case of T1; it will be shown in Section 6.4.1 that the energy spectrum of T2 is significantly affected by energy leaking backwards out of BINA. Therefore, a possible distortion of the time spectrum is expected when this cut is applied; however, it is effective in removing pileup. The analysis is done both with and without this cut, and it will ultimately be shown to slightly reduce the systematic uncertainty. Currently, the cut removes any event where the time difference between the first and last hits in an analog sum of the T2 PMTs is more than 100 ns.

Muhit cut

A special VT48 channel was connected to B1 extending approximately 25 μs before the trigger time; this was done to provide extra protection against beam muon pileup. A discriminator was used to require muon-like pulse height in B1. Events with hits in this channel up to 10 μs before the trigger were rejected. Rejecting any event with a hit in the full window was found to provide no benefit relative to the shortened version; its only noticeable effect was to reduce the available statistics.

Postpileup cut

One further minor cut is needed to protect against very late pileup. The BINA integration window extends for approximately 1 μs after the positron time, which, for late-decaying positrons, is slightly past the COPPER window for pileup rejection. Thus, there is a small time window, overlapping with the BINA integration window, in which a positron from old muon decay could enter BINA through T1 and T2 and the event would not be rejected by any of the previous pileup cuts. The VT48, however, covers 4 μs before and after the positron. In general, it was preferred to use the full waveform information provided by COPPER, but in this case it was necessary to resort to VT48 information.

Figure 5.11 shows the time difference between the last hit in T1 and the triggering hit in B1, recorded by the VT48, versus the T1 minus B1 time plotted in the time spectrum, for high energy events. The diagonal structure corresponds to events with no pileup, where the two times are the same, within the resolution. The small cluster of events at the far right, slightly above 1 μs on the y-axis, is due to the mechanism just discussed; the cut removes these events. For earlier decay times, the pileup rejection

5.1. Event selection

covers the full BINA integration window, and thus events with late pileup are rejected. For very late decay times, this is no longer the case, and the extra energy deposited in BINA can result in the event moving to the high energy time spectrum.

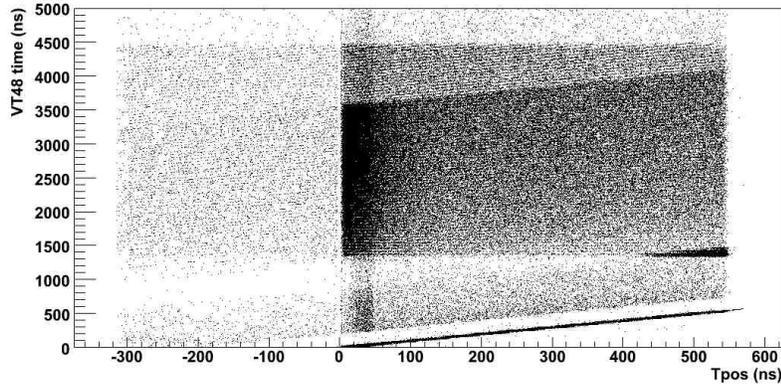


Figure 5.11: The time difference between the last hit in the T1 VT48 channel and the first hit in the B1 VT48 channel versus the decay time obtained from COPPER.

5.1.5 Acceptance cut

The most important cut still to be addressed restricts the radius from the centre of WC3 in which events are accepted, and thus also restricts the angle between the positron track and the crystal axis. The measured energy spectrum is highly dependent on the angle and position at which decay positrons enter BINA (see Chapter 6). The distance between the centre of WC3 and the positron track reconstructed using the S3WC3 tracker, is shown in Figure 5.12. This distance is referred to as R . For events with multiple tracks, the track with the minimum distance from the centre is taken. The usual value of the cut is shown by the red line. The placement of this cut is a tradeoff between the increasing systematic error as the low energy tail of the $\pi^+ \rightarrow e^+ \nu_e$ energy spectrum increases and the decreasing statistical error as more events are included in the analysis.

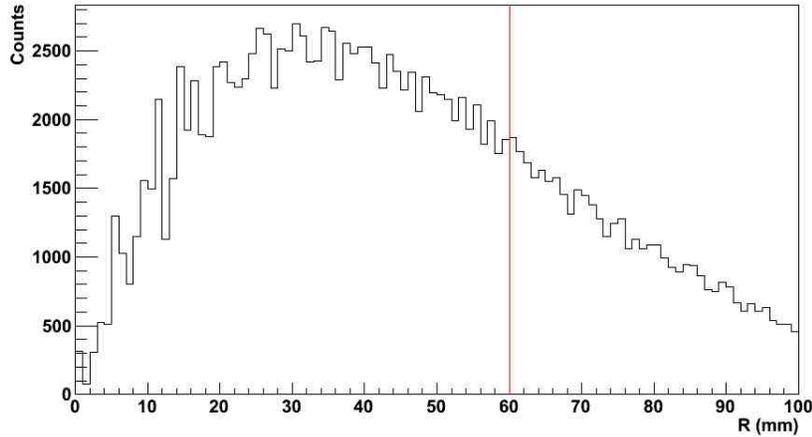


Figure 5.12: The distance between the reconstructed positron track and the centre of WC3 (R). The red line indicates the cut value.

5.1.6 Minor cuts

Four other cuts are applied, which affect a very small proportion of events, but are necessary either to remove rare processes or to ensure data integrity. The first category consists of a cut to remove events triggered by protons emitted when the stopping pion undergoes nuclear reactions in the target, and a cut to remove events in which the pion stopped upstream of the target and the decay muon completed the B1 B2 Tg coincidence. This is referred to as the false trigger cut. The cuts in the second category ensure that the trigger is caused by the first hit in each T1 PMT, and that the triggering hits in T1 and T2 are coincident within 20 ns (the coincidence window in the trigger itself was 100 ns). The proton events are removed based on the energy deposited in T1, T2, S3, and BINA; the events where the pion did not stop in the target are removed based on the energy deposited in the target and the time difference between the hits in B1 and Tg. The time spectra after all cuts are shown in Figure 5.13, and Table 5.2 shows the fraction of events removed by each cut, with all other cuts applied.

5.1. Event selection

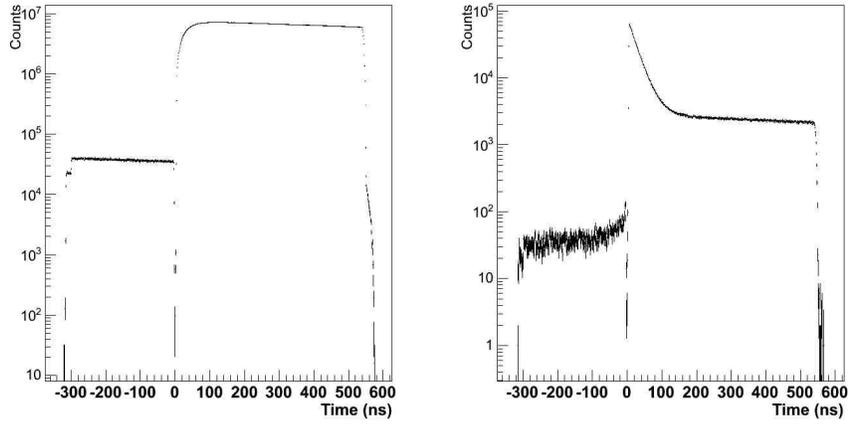


Figure 5.13: Time spectra following all cuts. The left-hand plot shows the low energy time spectrum and the right-hand plot shows the high energy time spectrum. The rise in the high energy spectrum near $t = 0$ at negative times is caused by the integration window of the calorimeter; the closer in time the pileup positron is to the positron from the pion at $t = 0$, the greater the probability that the measured energy in the event will be above E_{cut} .

Table 5.2: Events removed by each cut, with every other cut applied.

Cut name	Low energy	High energy
Prompt	4.2%	42.2%
B1B2 pileup	9.5%	89.2%
WC12	4.7%	4.7%
B1B2 energy	3.1%	5.5%
Prepileup	27.6%	49.2%
T1 pileup	0.03%	7.0%
T2 pileup	1.5%	11.4%
Muhit	0.6%	2%
Acceptance	43.1%	34.7%
Proton	0%	1.9%
False trigger	0.03%	0.03%
T1 trigger	0.1%	0.7%
T1T2 coincidence	0.0%	0.8%

5.2 Fitting function

5.2.1 Low energy time spectrum

Following all of these cuts, only two backgrounds remain in the low energy time spectrum at a non-negligible level: old muon decays and pion decays-in-flight (π -DIF). The time dependence of the former is an exponential with the muon lifetime, starting at the beginning of the trigger window, -300 ns. π -DIF events are only included in the fit if the decay muon stops before T1, which is typically only the case if the pion decays in flight within the target. If the pion decays before the target, the muon will pass through and stop in S3 or T1. If it stops in T1, the event is prompt, and thus outside the fitting range. The only π -DIF events included in the fit are those for which there is a muon in the target or S3 at $t = 0$; in either case, the time dependence of these events is also an exponential with the muon lifetime, but starting at $t = 0$.

The $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ signal shape is the convolution of exponentials with the pion lifetime and the muon lifetime. The fitting function used in the low energy time spectrum is the sum of these three shapes, shown in equation 5.1. H is the Heaviside function, t_0 is the offset in the time spectrum (determined from through-going particles), τ_μ and τ_π are the muon and pion lifetimes, A is the amplitude of the $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ shape, B is the amplitude of the pion decay-in-flight shape, and C is the amplitude of the old muon background. The low energy time distribution is given as follows:

$$f(t) = H(t) \left[A \frac{1}{\tau_\mu - \tau_\pi} \left(e^{-\frac{t-t_0}{\tau_\mu}} - e^{-\frac{t-t_0}{\tau_\pi}} \right) + B \frac{1}{\tau_\mu} e^{-\frac{t-t_0}{\tau_\mu}} \right] + C \frac{1}{\tau_\mu} e^{-\frac{t-t_0}{\tau_\mu}}. \quad (5.1)$$

5.2.2 High energy time spectrum

The high energy time spectrum is significantly more complicated, as several mechanisms can result in extra energy being deposited in the calorimeter (see Section 3.4). These are old muon pileup, which has its own time dependence; neutrons from the cyclotron, which are time independent relative to the decay; energy resolution effects, which are also decay time independent; and radiative pion decay, which has its own time dependence. Furthermore, there are two separate mechanisms by which old muon pileup events can be included in the time spectrum. If both decay positrons hit T1, the event

5.2. Fitting function

is rejected by the T1 pileup cut, unless they are sufficiently close in time, within about 15 ns, to be recorded as a single hit in T1. If the decays are separated in time, one of the positrons must miss T1 for the event to pass the cuts. Simplified timing diagrams for the various types of old muon events can be found in Appendix B. The time dependence of each of these processes will now be evaluated in turn.

Time-independent addition of energy

Time-independent mechanisms by which energy is added to events result in the components of the low energy time spectrum being present in the high energy time spectrum. The term included in the fit is thus Equation 5.1 multiplied by a free parameter, called r .

Old-muon pileup I

Although there is a pileup cut in T1, if two positrons pass through it sufficiently close together in time, the waveforms will overlap, and only a single hit will be recorded. The time spectrum for this component depends on whether the trigger was caused by the positron from the old muon or the positron from the primary pion, since the latter can only occur at positive time but the former can occur at any time. Let ΔT be the minimum time difference for which T1 can resolve hits. The shape for the old-muon trigger case is then given by the product of the amplitude of the old-muon shape and the probability that the $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ positron will emerge within ΔT . This component is called T1res; the expression is

$$f(t) = \begin{cases} 0 & t < -\Delta T \\ \frac{\exp(-\frac{t}{\tau_\mu})}{\tau_\mu} \int_0^{t+\Delta T} \frac{\exp(-\frac{y}{\tau_\mu}) - \exp(-\frac{y}{\tau_\pi})}{\tau_\mu - \tau_\pi} dy & -\Delta T < t < 0 \\ \frac{\exp(-\frac{t}{\tau_\mu})}{\tau_\mu} \int_t^{t+\Delta T} \frac{\exp(-\frac{y}{\tau_\mu}) - \exp(-\frac{y}{\tau_\pi})}{\tau_\mu - \tau_\pi} dy & t > 0. \end{cases} \quad (5.2)$$

If the trigger is caused by the $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ positron, the shape is instead given by the product of the $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ shape and the probability that the old muon will decay within ΔT . This case is given in Equation 5.3:

$$f(t) = \begin{cases} 0 & t < 0 \\ \frac{\exp(-\frac{t}{\tau_\mu}) - \exp(-\frac{t}{\tau_\pi})}{\tau_\mu - \tau_\pi} \int_t^{t+\Delta T} \frac{\exp(-\frac{y}{\tau_\mu})}{\tau_\mu} dy & t > 0. \end{cases} \quad (5.3)$$

5.2. Fitting function

To obtain ΔT , the time difference between subsequent T1 hits was plotted, and the edge fitted with an error function (see Equation 4.1). The fitted spectra for the T1 PMTs are shown in Figure 5.14; the average is $\Delta T = 15.7$ ns. The shape with this value of ΔT is shown in Figure 5.15.

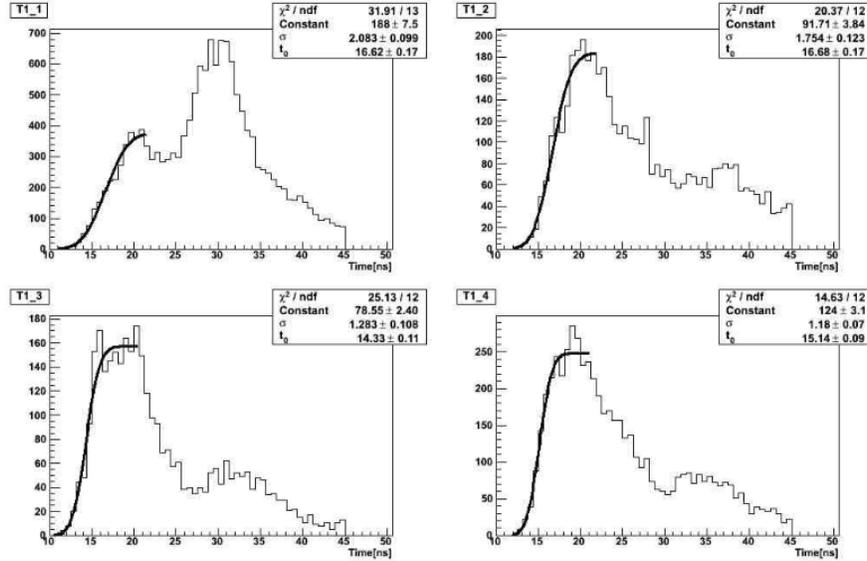


Figure 5.14: The time difference between subsequent hits in each T1 PMT; leading times are fitted with an error function. The peak around 30 ns is due to a fake hit at a characteristic time after the real hit.

Old-muon pileup II

Old-muon pileup events can also appear in the high energy time spectrum if one of the decay positrons misses T1, but still enters the crystal array; some solid angle exists for this to occur in either BINA or CsI (in BINA this is largely possible because of the rotation of T1 by 45° about the beam axis, which was necessary due to spatial constraints). The shape of this component is made more complicated by the BinaHigh trigger requirement in the high energy spectrum. A running sum of the BINA + CsI pulse height was used as an input for this trigger, which would only pass events above a certain threshold. The window in which this was done was 250 ns. However, the integration time used for the calibrated BINA energy was $1 \mu\text{s}$, meaning that if hits were sufficiently separated in time, the calibrated energy could be

5.2. Fitting function

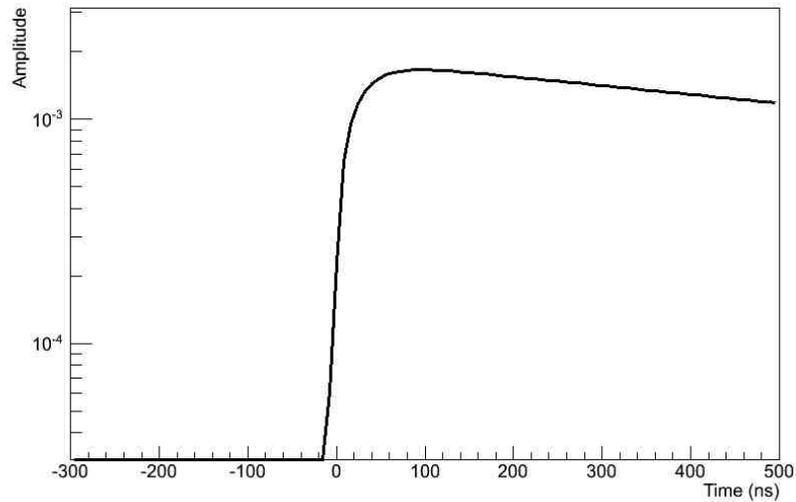


Figure 5.15: The shape used in the fit for pileup events that pass the T1 pileup cut due to the double-pulse resolution of the T1 counter.

above E_{cut} and a BinaHigh trigger would still not be present. Such events are excluded from the time spectrum, so this effect had to be taken into account when determining the shape of this pileup mechanism.

The shape was determined using Monte Carlo. Simulations were done of events caused by a $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ positron and of events caused by a positron from old muon decay. One event was then drawn at random from each simulation, to form one pileup event. If only one positron hit T1, the acceptance cut was passed, and the BinaHigh trigger requirement was met, the time of the event was included in the shape. The presence or absence of the trigger was determined using BINA and CsI waveform templates in conjunction with the energy deposited in the simulation. This term is called `oldmu_both`; its shape is shown in Figure 5.16.

Radiative decay

If the decay positron was produced in association with a photon via $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu \gamma$ the energy spectrum of the positron was altered, but the time dependence was not, and a separate shape is not required. If, however, the

5.2. Fitting function

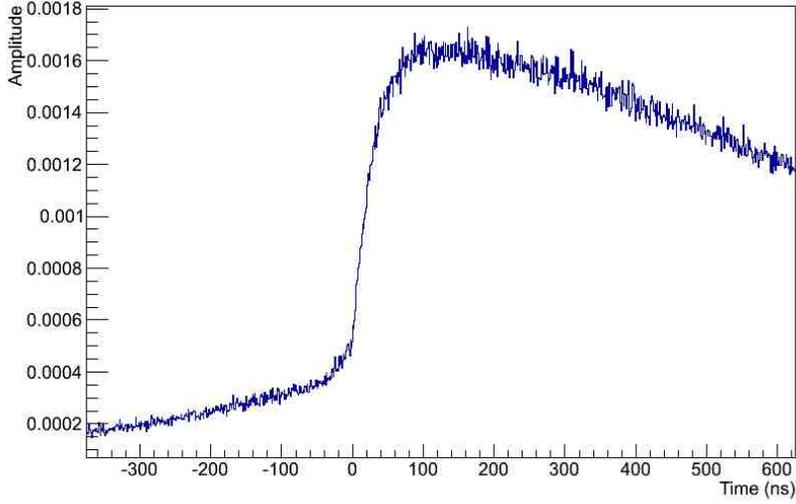


Figure 5.16: The shape used in the fit for pileup events where only one positron hit T1.

pion decayed radiatively to a muon, followed by $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ decay, the measured energy in the event could be above E_{cut} . The probability for this to occur is dependent on the relative timing of the photon and the positron, again because of the BinaHigh trigger requirement. This shape was therefore also taken from simulation; it is shown in Figure 5.17.

The full high energy fitting consists of all of these shapes and the $\pi^+ \rightarrow e^+ \nu_e$ signal shape, the amplitude of which is $(A - A * r) * Br$, where A is the $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ amplitude, r is the proportion of the low energy time spectrum that is present in the high energy time spectrum, and Br is the branching ratio. The fitting function is given in Equation 5.4. \mathcal{F}_1 , \mathcal{F}_2 , and \mathcal{F}_3 are the radiative decay shape and the two old muon pileup shapes, and \mathcal{L} is the low energy fitting function, given in Equation 5.1. The function for the high energy spectrum is given by

$$f(t) = H(t) \left[(A - A * r) * Br * \frac{1}{\tau_\pi} e^{-\frac{t-t_0}{\tau_\pi}} \right] + r * \mathcal{L} + \mathcal{F}_1 + \mathcal{F}_2 + \mathcal{F}_3. \quad (5.4)$$

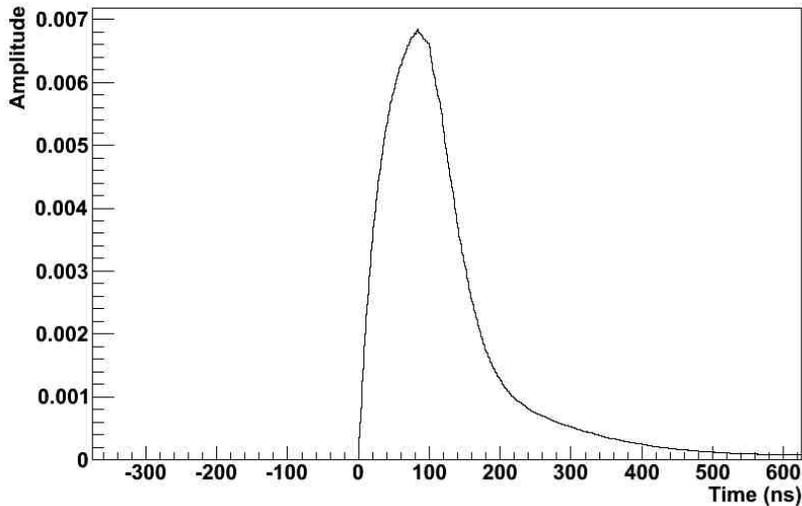


Figure 5.17: The shape used in the fit for $\pi \rightarrow \mu\nu\gamma$ events.

5.3 Fitting method

In all, the functions given in Equations 5.4 and 5.1 contain 12 parameters; their names, the symbols used in the fit, and a brief description of each, are given in Table 5.3. Of these parameters, five are typically fixed: t_0 , τ_μ , τ_π , ΔT (the T1 double-pulse resolution), and \mathcal{F}_2 , the amplitude of the T1res shape.

τ_μ , τ_π , t_0 , and ΔT are all constants, and therefore fixing them in the fit is the natural approach, although the stability of the branching ratio when τ_μ and τ_π are freed in the fit is an important systematic check. The amplitude of T1res is fixed because of the similarity of its shape to \mathcal{F}_3 , the oldmu_both shape, which makes them difficult to fit simultaneously. It was decided to fix the first, rather than the second, because ΔT can be increased arbitrarily by allowing events where successive hits in T1 are separated by less than some time interval. This causes the shape to change significantly, particularly at negative times, making it possible for the fit to determine the correct amplitude. As an example, the shape with $\Delta T = 100$ ns is shown in Figure 5.18. The fit was done for several values of ΔT , and an extrapolation

5.3. Fitting method

Table 5.3: Fit parameter list.

Parameter name	Symbol	Description
A	A	The amplitude of the $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ shape.
r	r	The fraction of events promoted from the low to high energy time spectrum via time independent mechanisms.
t_0	t_0	The pion stop time.
Mu_Pimu	B	The amplitude of the old muon background in the low-energy time spectrum.
Mu_Pie	\mathcal{F}_4	The amplitude of the pure old-muon background in the high energy time spectrum.
π DIF	C	The amplitude of the pion decay-in-flight component.
BR	Br	The branching ratio: the $\pi^+ \rightarrow e^+ \nu_e$ amplitude divided by the $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ amplitude.
τ_μ	τ_μ	The lifetime of the muon.
τ_π	τ_π	The lifetime of the pion.
T1res	\mathcal{F}_2	The amplitude of the old muon pileup background in the high energy time spectrum, for the case where both decay positrons hit T1 too close together in time to resolve as separate hits.
oldmu_both	\mathcal{F}_3	The amplitude of the old muon pileup background in the high energy time spectrum, for the case where one of the decay positrons missed T1.
ΔT	ΔT	The T1 double-pulse resolution.
pimugamma	\mathcal{F}_1	The amplitude of the shape for $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events where the initial pion decay was radiative, and the photon energy moved the measured energy in the event above E_{cut} .

was performed to the true value, $\Delta T = 15.7$ ns. The fitted amplitude as a function of ΔT is shown in Figure 5.19, and the fit of the high energy time spectrum with $\Delta T = 100$ ns is shown in Figure 5.20.

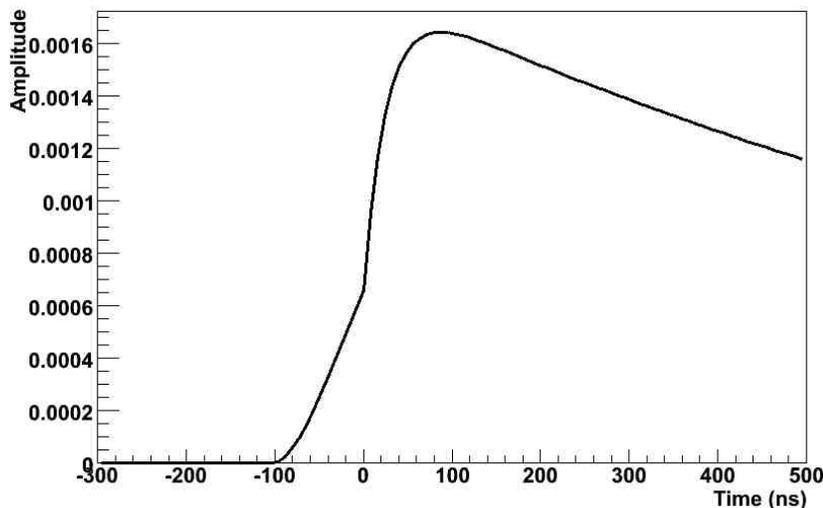


Figure 5.18: The T1res shape when the double-pulse resolution is set to 100 ns.

One other modification had to be made in order to make the fit converge. In principle, the value of the parameter r , describing the fraction of events with low energy timing that are in the high energy time spectrum, should be the same for negative and positive times. However, it was necessary to free the old muon amplitude at negative times in the high energy spectrum to obtain a satisfactory fit; thus, another parameter \mathcal{F}_4 must be added to Equation 5.4. The reasons for this are not currently understood. The additional parameter is referred to as Mu_Pie.

5.4 Results

The fitted 2012 time spectra are shown in Figure 5.21. The fit is performed from -290 to -20 ns, and 10 to 520 ns; the region $-20 \text{ ns} < t < 10 \text{ ns}$ is excluded.

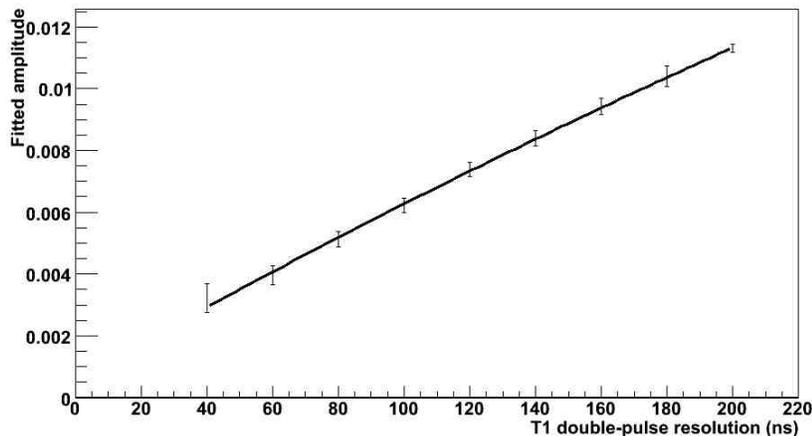


Figure 5.19: The fitted T1res amplitude as the double-pulse resolution is increased.

The values of the fitting parameters and their errors are given in Table 5.4, along with the corresponding values for the 2010 data set. Values are given both with and without the T2 pileup cut (see Section 5.1.4). The total χ^2 / d.o.f. of the fits is 1.39 with the T2 pileup cut and 1.23 without, in each case considerably worse than the values for the 2010 and 2011 data sets. The residuals of the fits of the 2010 and 2011 data sets, and of the 2012 data set with and without the T2 pileup cut, are shown in Figures 5.22, 5.23, 5.24, and 5.25.

Examining the residuals, it is clear that the deterioration in the quality of the fit is driven by the low energy $t > 0$ spectrum, where large, seemingly periodic, distortions can be seen. In principle, a small oscillating distortion of the time spectrum will not affect the branching ratio as long as events are not added or removed, but only shifted. Since the shape of the spectrum is very sensitive to both the pion and muon lifetimes, if these are fixed to the correct values and the integral of the residuals is consistent with zero, the branching ratio should be correct. The result obtained for the integral of the residuals was $551,831 \pm 223,327$, which is approximately 2.5σ away from zero. For reference, the number of events in the low energy $t > 0$ spectrum is 3.269×10^9 . So, although there may be evidence for a slight excess of events in the low energy time spectrum, it is on the level of 1.5×10^{-4} , which will

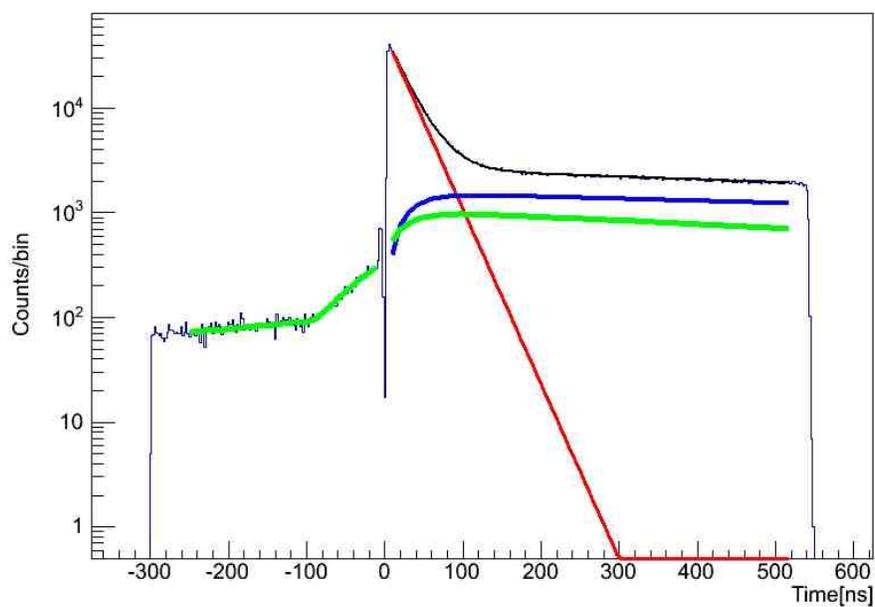


Figure 5.20: The fit of the high-energy time spectrum with ΔT increased to 100 ns. The red shape is the $\pi^+ \rightarrow e^+ \nu_e$ signal, the blue is the $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ background, and the green is the sum of the other backgrounds.

5.4. Results

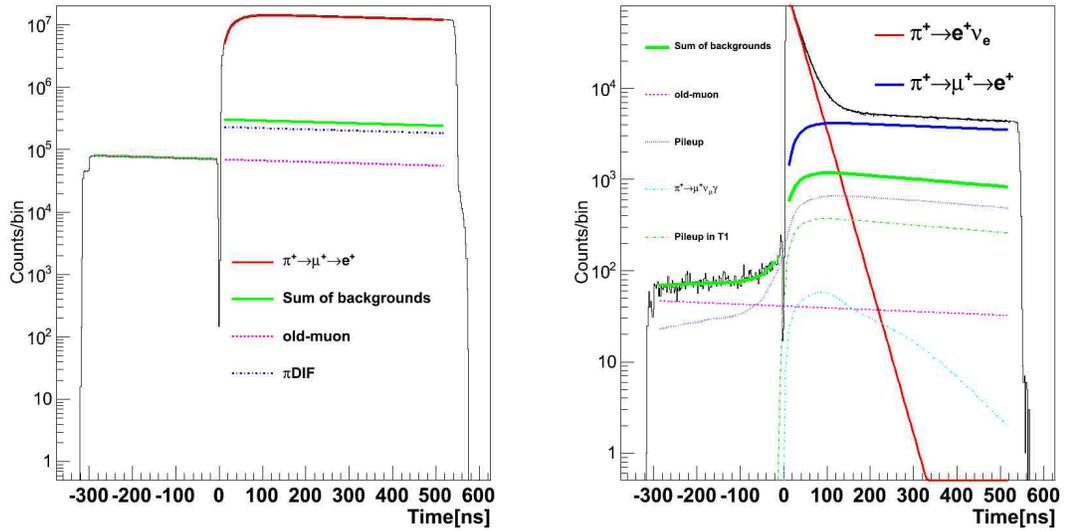


Figure 5.21: The fitted time spectra from 2012. The left-hand panel shows the low energy time spectrum, fitted with three components: the $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ signal shape, old muon decays, and pion decays-in-flight. The right-hand panel shows the high energy time spectrum, fitted with six components: the $\pi^+ \rightarrow e^+ \nu_e$ signal shape, $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ and pion decay-in-flight events promoted to the high energy time spectrum via time independent mechanisms, two mechanisms of old muon pileup, $\pi \rightarrow \mu \nu \gamma$ decays, and old muon decays.

5.4. Results

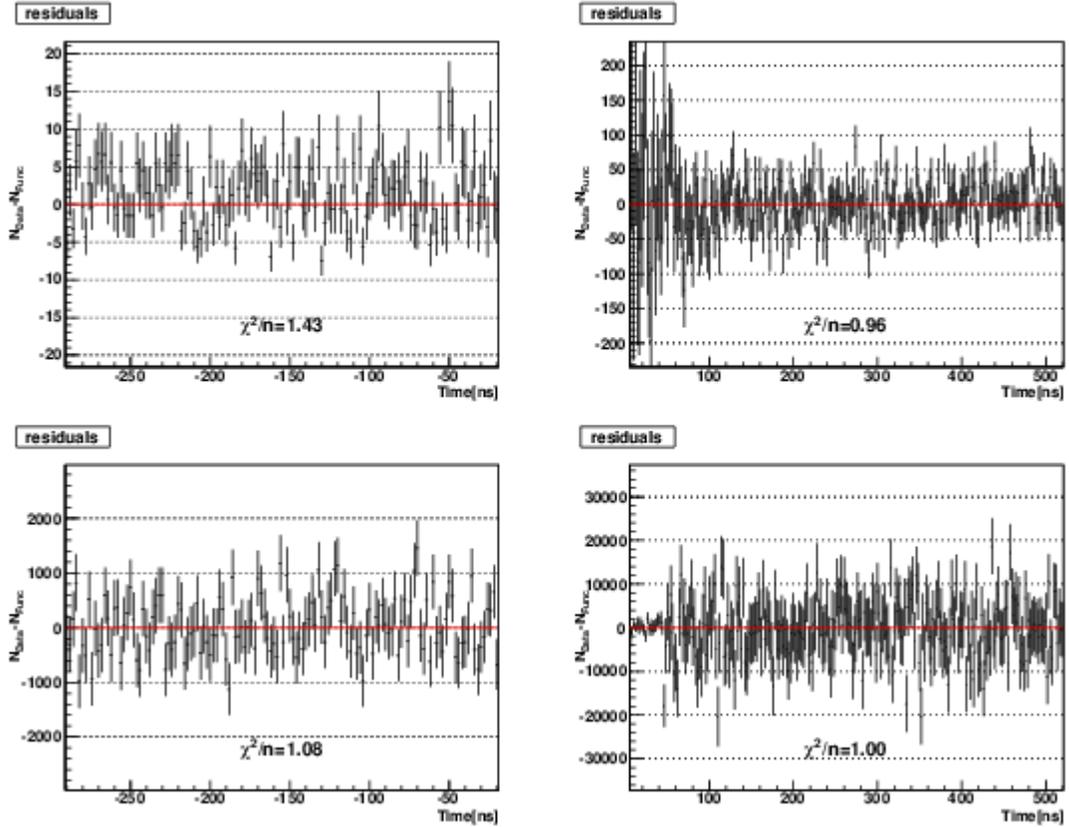


Figure 5.22: The residuals of the 2010 data set vs. the time of the event. Clockwise from top left, the panels show the residuals for the high energy $t < 0$ spectrum, the high energy $t > 0$ spectrum, the low energy $t > 0$ spectrum, and the low energy $t < 0$ spectrum.

5.4. Results

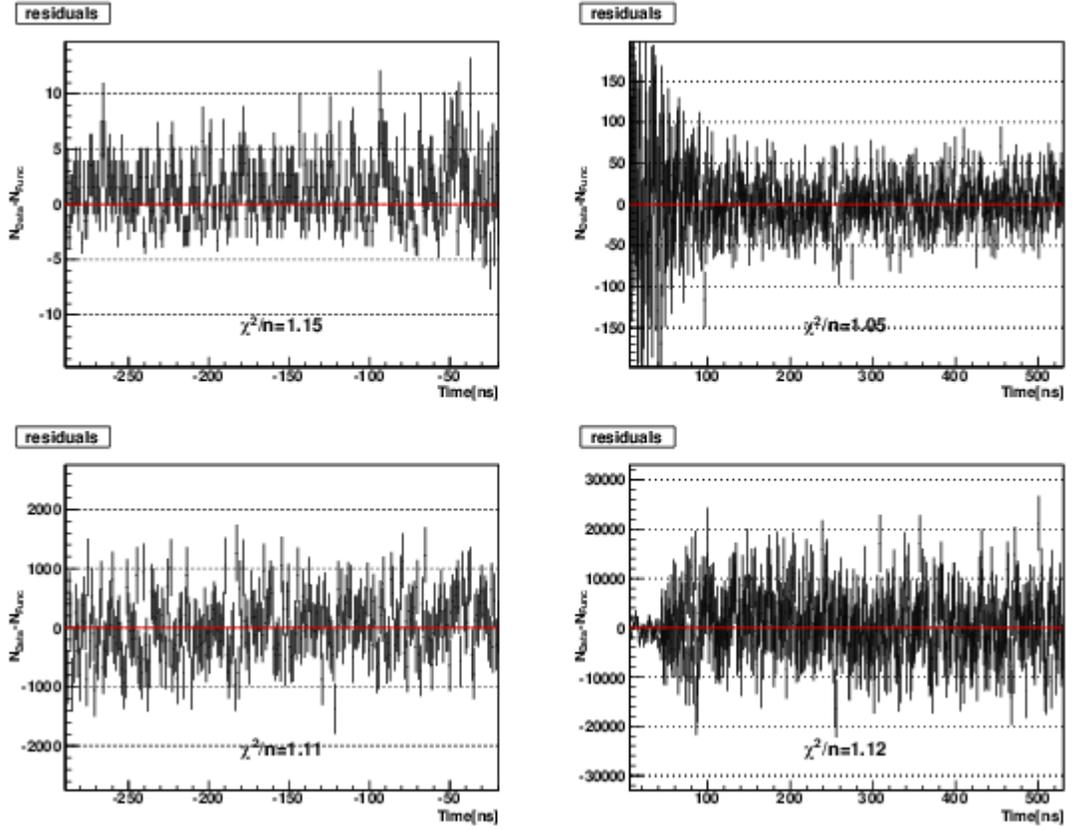


Figure 5.23: The residuals of the 2011 data set vs. the time of the event. Clockwise from top left, the panels show the residuals for the high energy $t < 0$ spectrum, the high energy $t > 0$ spectrum, the low energy $t > 0$ spectrum, and the low energy $t < 0$ spectrum.

5.4. Results

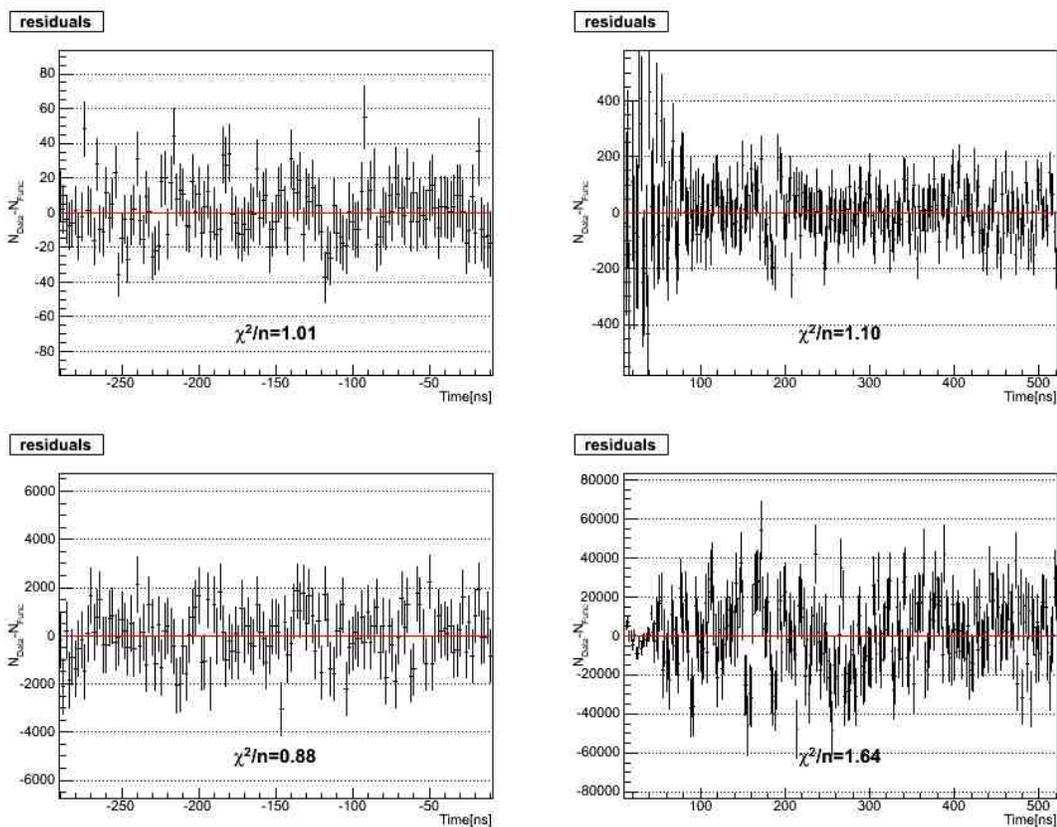


Figure 5.24: The residuals of the 2012 data set vs. the time of the event, without the T2 pileup cut applied. Clockwise from top left, the panels show the residuals for the high energy $t < 0$ spectrum, the high energy $t > 0$ spectrum, the low energy $t > 0$ spectrum, and the low energy $t < 0$ spectrum.

5.4. Results

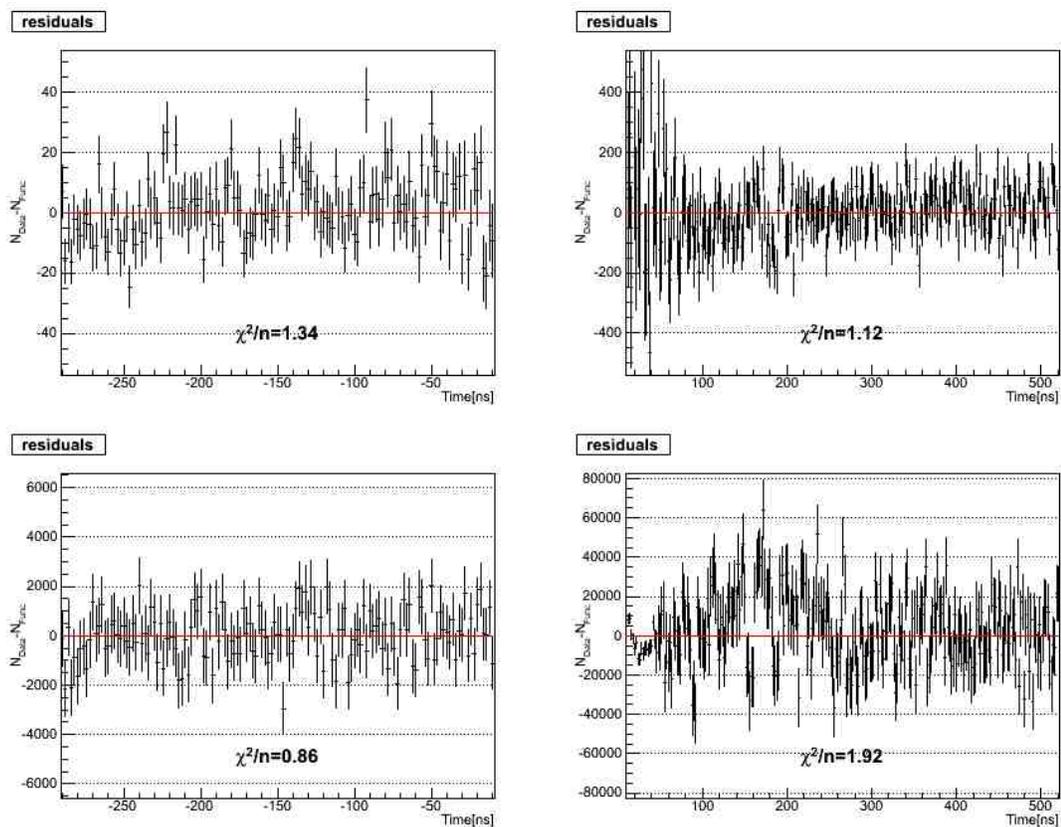


Figure 5.25: The residuals of the 2012 data set vs. the time of the event, with the T2 pileup cut applied. Clockwise from top left, the panels show the residuals for the high energy $t < 0$ spectrum, the high energy $t > 0$ spectrum, the low energy $t > 0$ spectrum, and the low energy $t < 0$ spectrum.

5.5. Systematic checks

Table 5.4: Fit parameters for the 2010 and 2012 data sets. Note that the 2012 branching ratios are still blinded by an unknown factor uniformly distributed between $\pm 0.5\%$.

Parameter	2010 value	2012 value (T2 cut)	2012 value (no T2 pileup cut)
A (10^9)	3.2903 ± 0.0010	16.227 ± 0.002	16.449 ± 0.002
r (10^{-4})	2.458 ± 0.0047	2.94 ± 0.02	3.40 ± 0.02
Mu_Pimu (10^7)	1.533 ± 0.004	7.65 ± 0.01	8.03 ± 0.01
Mu_Pie (10^3)	7.18 ± 0.65	44.7 ± 1.7	38.8 ± 5.5
π DIF(10^7)	3.821 ± 0.067	25.3 ± 0.2	27.6 ± 0.2
BR (10^{-4})	1.1971 ± 0.0022	1.2005 ± 0.0011	1.2027 ± 0.0011
Oldmu_both (10^4)	2.93 ± 0.39	8.92 ± 0.4	18.1 ± 0.5

not be negligible when the error on the branching ratio reaches the 0.1% level, but will not contribute greatly to the systematic error, since it would decrease the branching ratio by $< 0.02\%$.

5.5 Systematic checks

Several different tests were performed to ensure the stability of the fit: the fitting range was changed, the bin size was changed from the usual value of 2 ns, τ_μ and τ_π were made free parameters, a flat component was added, and the amplitudes of the fixed shapes were changed. A summary of these systematic checks is given in Table 5.5. By far the largest variation arises when the lifetimes are freed, if the T2 pileup cut is applied.

Another important diagnostic, particularly with regard to the fitting function, is the stability of the raw branching ratio when more pileup events are allowed into the time spectra. The prepileup cut normally rejects events with hits in any of the plastic scintillators in a window 6.6 μ s to 2.2 μ s before the pion stop (see Section 5.1.4). The fit was done for several different prepileup rejection windows in B1, B2, and Tg, from the full cut to no prepileup cut. The branching ratio as a function of the left-hand edge of the prepileup rejection window, before and after applying the T2 pileup cut, is shown in Figures 5.26 and 5.27, respectively.

Finally, the stability of the branching ratio as a function of time was tested. The fit was performed for groups of 1000 runs, and the result fitted

5.5. Systematic checks

Table 5.5: Systematic checks performed on the fit of the 2012 data.

Test	$\Delta\text{BR} (10^{-8})$	$\Delta\text{BRerror} (10^{-8})$	T2pileup
Fitting range			
$t < 420$	0.1	0.0	No
$t < 420$	2.0	0.1	Yes
$12.5 < t$	-3.7	0.5	No
$12.5 < t$	-8.4	0.5	Yes
$t < 7.5$	-2.6	-0.5	No
$t < 7.5$	2.3	-0.4	Yes
$t < -30$	0.0	0.0	No
$t < -30$	-0.8	0.0	Yes
$-250 < t$	0.0	0.0	No
$-250 < t$	0.2	0.0	Yes
Bin size			
Bin = 1 ns	-2.4	0.3	No
Bin = 1 ns	-4.5	0.2	Yes
Bin = 4 ns	3.0	0.0	No
Bin = 4 ns	1.4	0.0	Yes
Lifetime			
τ_μ and τ_π free	-6.1	6.8	No
τ_μ and τ_π free	43.9	6.8	Yes
Additional background			
Flat component	1.4	0.0	No
Flat component	2.6	0.0	Yes
Fixed parameter			
$\pm 20\%$ pimugamma	± 2.6	0.0	No
$\pm 20\%$ pimugamma	± 2.6	0.0	Yes

5.5. Systematic checks

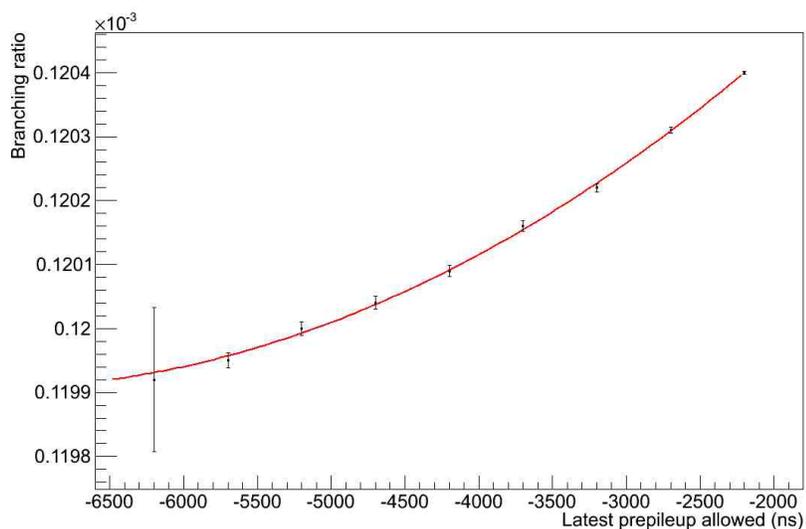


Figure 5.26: The variation of the branching ratio as more pileup events are allowed in the time spectrum, prior to applying the T2 pileup cut. The x axis denotes the time prior to the pion stop in which events with hits in B1, B2, and Tg were rejected. The error bar on the point furthest to the left (with the least pileup) is the error from the time spectrum fit, and the error bars on the other points are the error on the change from the previous point. The points are fitted to a parabola.

5.5. Systematic checks

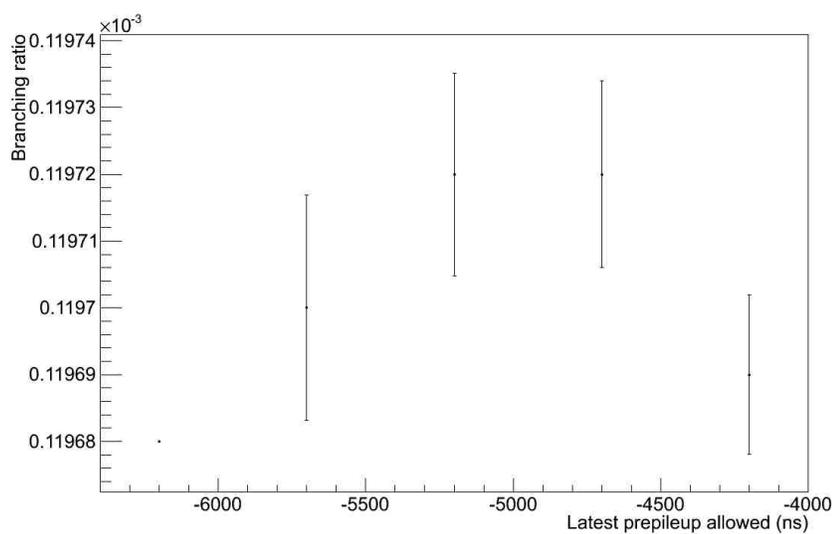


Figure 5.27: The variation of the branching ratio as more pileup events are allowed in the time spectrum, after applying the T2 pileup cut. The x axis denotes the time prior to the pion stop in which events with hits in B1, B2, and Tg were rejected. The error bars represent the statistical variation from the point furthest to the left (with the least pileup).

5.6. Summary of Chapter 5

to a straight line, as shown in Figure 5.28. The p-value of the fit is 40.7%, meaning there is no evidence for time dependence.

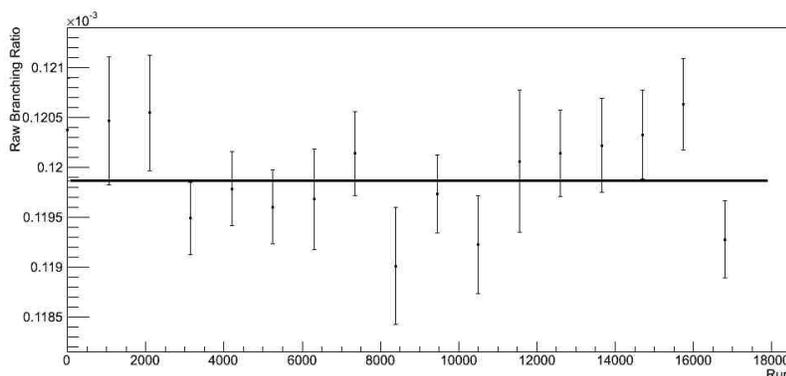


Figure 5.28: The fitted branching ratio for groups of 1000 runs, fitted with a flat line.

5.6 Summary of Chapter 5

In the absence of the T2 pileup cut, the branching ratio varies considerably as pileup is added to the spectrum, indicating the presence of an incorrect shape or missing component in the time spectrum fit. However, the impact on the branching ratio becomes almost negligible as the cut approaches its nominal value. The fit shown in Figure 5.26 was extrapolated to its minimum; the corresponding branching ratio was within 2×10^{-8} of the lowest point on the graph. Variations at this level are negligible for the present analysis; see Chapter 8.

The other large variation in the branching ratio occurs when the T2 pileup cut is applied and τ_μ and τ_π are made free parameters in the fit. It is possible that the value of the branching ratio when the lifetimes are freed is closer to the correct value; the change in the lifetimes may compensate for whatever distortion exists in the time spectrum. However, the lifetimes are highly correlated with the branching ratio, and, particularly in the presence of distortions, spurious minima may exist in the χ^2 landscape. In the absence of a reason to prefer one value of the branching ratio over the other, the

5.6. Summary of Chapter 5

average is taken, and half the variation assigned as the 1σ error; in this case $\sigma = 21.5 \times 10^{-8}$. This is to be compared with a statistical error of 11×10^{-8} .

Several corrections must be applied to the raw branching ratio in order to obtain the final result. These are the subjects of the next two chapters; the final systematic error will be given in Chapter 8.

Chapter 6

Tail Correction

6.1 Introduction

By far the largest correction to the branching ratio arises from the tail of the measured $\pi^+ \rightarrow e^+\nu_e$ energy spectrum below the cutoff value E_{cut} between the low and high energy time spectra. The standard condition for the analysis is $E_{cut} = 52$ MeV. Let the measured $\pi^+ \rightarrow e^+\nu_e$ energy spectrum be denoted by $N(E)$. The tail fraction T is defined as the proportion of this spectrum below the cutoff; that is,

$$T = \frac{\int_0^{E_{cut}} N(E)dE}{\int_0^{\infty} N(E)dE}. \quad (6.1)$$

The size of this tail will be shown to be approximately three percent. The amplitude of the $\pi^+ \rightarrow e^+\nu_e$ component in the low energy time spectrum is equal to the amplitude of the $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ component, multiplied by T and the branching ratio. This is too small to fit with the PIENU data set; neglecting other corrections, the raw branching ratio obtained from the fit is thus related to the true branching ratio by

$$(1 - T)BR_{true} = BR_{raw}. \quad (6.2)$$

Since this component cannot be fitted, the response of the crystal calorimeter to a 70 MeV positron beam had to be determined independently, and the raw branching ratio multiplied directly by $\frac{1}{1-T}$. If ΔT is the uncertainty on T , this results in an uncertainty on the branching ratio of $\frac{\Delta T}{1-T} * BR$; in order to meet the precision goal of the experiment, $\frac{\Delta BR}{BR} < 0.1\%$, ΔT must itself be substantially less than 0.1%.

Two fundamentally different approaches were used to determine T . One was to inject a positron beam with almost the same energy as $\pi^+ \rightarrow e^+\nu_e$ decay positrons, about 70 MeV, into BINA at several angles and directly measure the proportion of the spectrum below E_{cut} . The other was to, as far

6.2. Lower limit

as possible, suppress $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events in the data, and attempt to deduce T from the energy spectrum itself, using the known shape of the Michel spectrum.

The first of these techniques must be regarded as giving an upper limit on T , because of the potential for the positrons to scatter in the beamline, giving a low momentum tail to the beam itself. The second uses the approximation that the $\pi^+ \rightarrow e^+ \nu_e$ tail is zero at very low energies, which leads to a slight over-subtraction of $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events from the measured energy spectrum. This results in an underestimation of T , so this technique is regarded as giving a lower limit. The upper and lower limits are then combined to give the best estimate of T . For a detailed description of the method for obtaining the lower limit, see [65].

6.2 Lower limit

The first step in the lower limit procedure was to obtain an energy spectrum with as many $\pi^+ \rightarrow e^+ \nu_e$ events and as few $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events as possible remaining, and other backgrounds suppressed as far as possible without distorting the energy spectrum for $\pi^+ \rightarrow e^+ \nu_e$ positrons. Five additional cuts were imposed to accomplish this: a timing cut, a cut on the energy deposited in all detectors up to and including the target, a cut on the angle between the tracks reconstructed by WC12 and S12 (to remove pion decays-in-flight), a cut on the energy deposited in S3, and a cut on the shape of the Tg waveform.

The timing cut removed any event outside the time window 7-42 ns; this window contains 80% of $\pi^+ \rightarrow e^+ \nu_e$ events, and less than 2% of $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events. The energy loss cut selected events depositing between 15.5 and 16.5 MeV in total in B1, B2, S1, S2, and Tg. Simulations of the energy deposited by the four relevant processes are shown in Figure 6.1: $\pi^+ \rightarrow e^+ \nu_e$ decay-at-rest (π -DAR), $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ decay-at-rest (π -DAR μ -DAR), π decay-at-rest followed by μ decay-in-flight (π -DAR μ -DIF), and π decay-in-flight followed by μ decay-at-rest (π -DIF μ -DAR).

The energy loss cut was very effective in removing $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events where both particles decayed at rest. However, contamination remained from events where either particle decayed in flight. To suppress π -DIF events, a cut was done on the angle between the tracks reconstructed

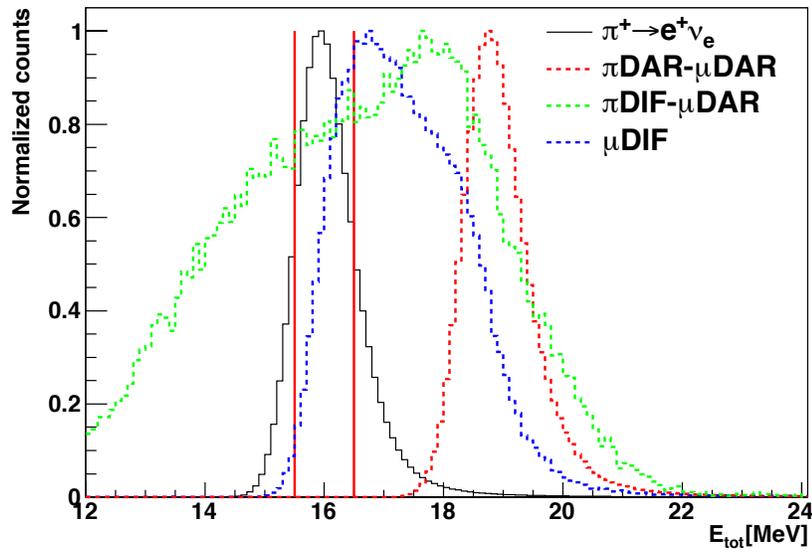


Figure 6.1: The total energy deposited in B1, B2, S1, S2, and Tg by $\pi^+ \rightarrow e^+ \nu_e$ (black), $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ (red), π DIF- μ DAR (green), and π DAR- μ DIF (blue). The distributions are normalized to the same height. The solid red lines indicate the selected region.

6.2. Lower limit

by WC12 and S12. The distributions in the two cases are shown in Figure 6.2.

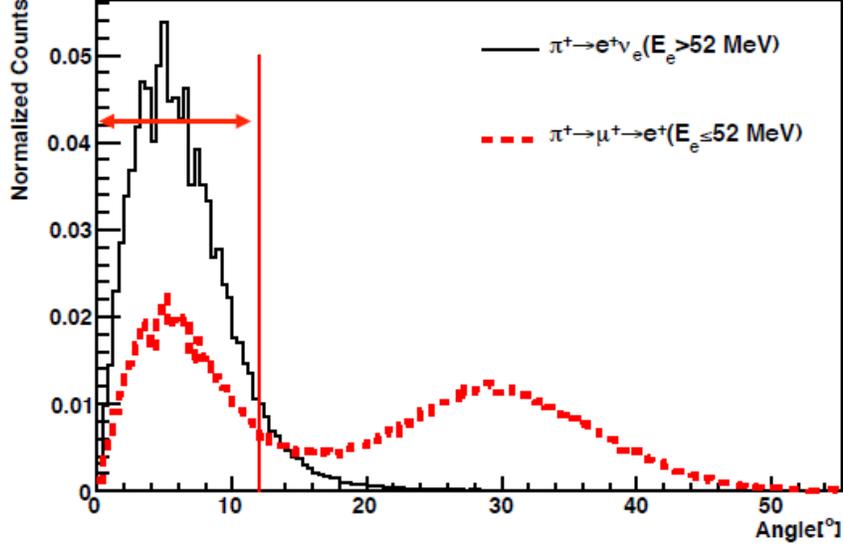


Figure 6.2: The measured angle between the tracks reconstructed by WC12 and S12 for $\pi^+ \rightarrow e^+ \nu_e$ (black) and $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ (red) events after the time and energy loss cuts are applied. The double peak in the $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ distribution is caused by pion decays-in-flight. The solid red line and arrow indicate the selected region.

π -DIF events were further suppressed using the energy deposited in S3. About 30% of muons arising from π -DIF stopped in S3 rather than the target, resulting in larger energy deposit than was typical for positrons. Finally, the target waveform was fitted with a template containing two peaks and a template containing three peaks, corresponding to $\pi^+ \rightarrow e^+ \nu_e$ and $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events, respectively. Events were removed if the χ^2 of the three-pulse fit was better than the two-pulse fit. This was done to ensure as few as possible $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events remained, but after all the other suppression cuts the effect was almost negligible. The effect of each cut in turn upon the energy spectrum is shown in Figure 6.3; the

6.2. Lower limit

spectrum after all cuts is referred to as the suppressed spectrum.

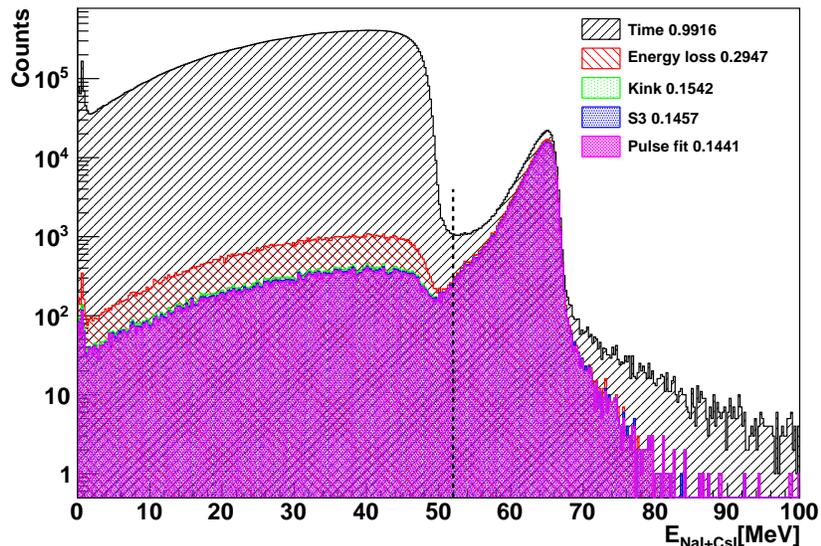


Figure 6.3: The measured BINA+CsI spectra as the suppression cuts are applied. The legend indicates the fraction of low-energy events (< 52 MeV) remaining after each cut.

There are expected to be three components below 52 MeV in Figure 6.3: the $\pi^+ \rightarrow e^+\nu_e$ tail, μ -DIF events, and π -DIF μ -DAR events. The last have essentially the same energy spectrum as normal $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ events, whereas μ -DIF events are somewhat higher energy. The energy spectrum for μ -DIF events was obtained from Monte Carlo. The energy spectrum for π -DIF μ -DAR events was taken from data, assuming the same distribution as $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ events. These spectra were then combined to give the total background spectrum. The number of background events was estimated by extrapolating from the number of events in the suppressed spectrum at low energy, where the $\pi^+ \rightarrow e^+\nu_e$ tail is very small. This estimate of the number of background events was then subtracted from the total number of events below 52 MeV in the suppressed spectrum, and the remainder was taken as the lower limit on the $\pi^+ \rightarrow e^+\nu_e$ tail.

A further correction had to be applied because of the cut on the total energy deposited in the detectors up to the target, which tended to remove $\pi^+ \rightarrow e^+\nu_e$ events that underwent Bhabha (electron-positron) scattering. This correction was taken from Monte Carlo. The final result for the lower limit was $T_{LL} = 2.95\% \pm 0.07\%(\text{stat}) \pm 0.08\%(\text{syst})$ for the 2010 data and $T_{LL} = 3.22\% \pm 0.09\%(\text{stat}) \pm 0.05\%(\text{syst})$ for the 2011 data. However, the 2011 lower limit analysis was done with a modified procedure where part of the $\pi^+ \rightarrow e^+\nu_e$ energy spectrum was taken from Monte Carlo and used in a fit of the measured energy spectrum. In order to ensure the crystal response is determined entirely from data, the 2010 value was used in combination with the upper limit for this analysis. Analysis of the 2012 suppressed spectrum is underway.

6.3 Response function measurement

6.3.1 Energy loss processes

To obtain the upper limit on the $\pi^+ \rightarrow e^+\nu_e$ tail, a positron beam was injected into the crystal array. The PIENU beamline was designed to be able to deliver either a pion, muon, or positron beam, through the adjustment of the absorber and collimator configuration and the fields of the bending magnets. Special data were taken in 2009 and again in 2011 with the beamline tuned for positrons at 70 MeV. The 2011 data were of higher quality, and are all that will be discussed here. There were two primary contributions to the low energy tail in the measured $\pi^+ \rightarrow e^+\nu_e$ energy spectrum: electromagnetic shower leakage and energy loss upstream of the calorimeter. Another small contribution arose from photonuclear interactions within BINA [9].

Electromagnetic showers

When a 70 MeV positron enters a material, in this case a NaI(Tl) crystal, it immediately begins to lose energy via bremsstrahlung and ionization. The bremsstrahlung photons in turn undergo pair production, resulting in electron-positron pairs that themselves undergo bremsstrahlung, ionization, and, in the case of positrons, annihilation. This process, referred to as an electromagnetic shower, continues until the initial positron runs out of energy and annihilates. Eventually, the energy thus absorbed by the crystal is converted into scintillation light, which can be read out by PMTs.

For any detector of finite size, some of the shower will occasionally escape, resulting in less measured light. If enough energy escaped from BINA, the measured energy could fall below E_{cut} , putting the event in the tail. Because the shower mainly proceeded via bremsstrahlung and pair production, it was highly forward-peaked, so the amount of material along the path of the initial positron essentially determined the probability that less than the full positron energy would be absorbed. Thus, it was essential to measure the response of the crystals as a function of the angle and entrance position of the positron on the front face of BINA. Because the orientation of the beamline was fixed, the crystal array was rotated relative to the beam.

Upstream interactions

During normal PIENU data-taking, the pions stopped near the centre of the target, approximately 8 cm upstream from the front face of BINA. If the pion underwent $\pi^+ \rightarrow e^+\nu_e$ decay, a positron with kinetic energy of 69.3 MeV was emitted isotropically from the centre of the target. Positrons further than 6 cm away from the centre of WC3, 5.5 cm downstream of the centre of the target, were rejected in the analysis; this corresponded to an angle of slightly more than 45° relative to the axis of BINA. Prior to reaching BINA, the positrons traversed half of Tg, and all of S3, T1, and T2. For a positron travelling on-axis, this corresponded to about 1.4 cm of plastic scintillator and 0.6 mm of silicon, or about 3 MeV of energy loss on average. However, sometimes significantly more energy was lost, raising the probability that the measured energy would fall below E_{cut} .

The effect of the material in the path of the $\pi^+ \rightarrow e^+\nu_e$ positrons could not be replicated using the positron beam. Even if the correct amount of material were placed in the path of the beam at each angle, which due to mechanical constraints was not straightforward, one significant difference would still remain. $\pi^+ \rightarrow e^+\nu_e$ positrons emitted at high angles relative to the axis of BINA would not normally trigger; however, a non-negligible proportion of such positrons underwent Bhabha scattering while exiting the target. The scattered, generally low energy, electron could then trigger and enter the calorimeter, producing an event with low measured energy and $\pi^+ \rightarrow e^+\nu_e$ time dependence. If the same thing happened to a $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ event, it was included in the branching ratio, so $\pi^+ \rightarrow e^+\nu_e$ events that underwent this process had to be included as well.

The effect of high angle positrons scattering electrons into the detector could not be replicated using the positron beam. However, the interactions of positrons in matter are well-understood, and can be precisely calculated using GEANT4. The electromagnetic shower inside BINA, on the other hand, was sensitive to its real properties, such as defects or inhomogeneities. Thus, the decision was made to remove as much material as possible in front of BINA in order to obtain as accurate a measurement as possible of the response function, and rely on MC for the contribution to the tail of energy loss upstream of BINA. This meant the value of E_{cut} for the response function measurement was effectively different from its value for the $\pi^+ \rightarrow e^+ \nu_e$ case, as the peak of the energy spectrum was shifted.

Photonuclear effects

When the PIENU detector was initially being commissioned and characterized, the response of BINA to a positron beam parallel to the crystal axis was measured. In addition to the expected peak around 70 MeV, the initial positron energy, with a low energy tail, additional small peaks were seen at 58 and 50.5 MeV. After ruling out explanations such as scattering in the beamline, the experimental configuration was simulated using GEANT4, and peaks were seen at similar energies, but only if photonuclear interactions were included in the simulation. By examining the events in these peaks, it was determined that they were caused by either one or two neutrons being emitted from iodine and escaping BINA. Although GEANT4 did not precisely reproduce the shape of these peaks, it confirmed their existence, and the mechanism of their creation. The energy spectrum obtained during commissioning, and simulations both including and not including hadronic interactions, are shown in Figure 6.4 [9]. The fact that the second peak was near the nominal value of E_{cut} , and that the MC did not accurately reproduce the size of the peak, was an additional reason to rely on data to determine the true crystal response.

6.3.2 Detector setup

The PIENU-I subassembly, comprising B1, B2, S1, S2, Tg, S3, and T1, was removed from the detector, leaving only the wire chambers, T2, and the calorimeter. This reduced the momentum and position divergence of the beam, allowing for a more accurate measurement of the crystal response. Finally, the PIENU-II subassembly (WC3, T2, and the crystals) was removed from the rails to which it was normally affixed so that it could be

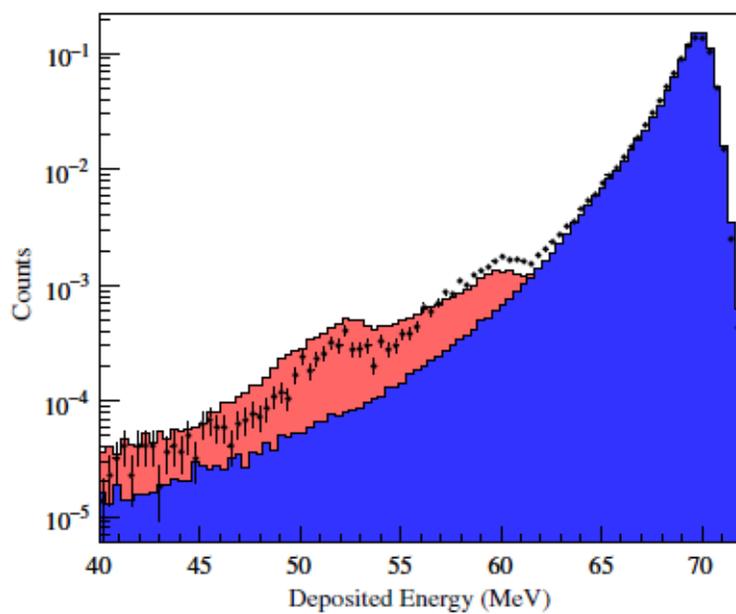


Figure 6.4: Comparison between measured (filled circles with error bars) and simulated energy spectra. The simulation was performed with (red) and without (blue) hadronic reaction contributions. The histograms are normalized to the same area [9].

rotated with respect to the beam.

The amount of material in the path of the beam depended not only on the angle θ between the beam and the crystal axis, but also on the centre of rotation. In the $\pi^+ \rightarrow e^+\nu_e$ case, the centre of rotation of the positron tracks was the centre of the target. In order to mimic this configuration using a straight beam and rotated detector, the centre of rotation had to be the same distance from the front face of BINA as the centre of the target. To this end, a shaft was attached to the cart on which PIENU-II sat, at the same distance in z from the front face of BINA as the centre of the target.

6.4 Data-taking

The angles at which data were taken, and the number of good positron events recorded at each angle, are shown in Table 6.1. To accurately determine the angle, markers were placed along the beamline and along a bar parallel to the crystal axis. A theodolite was then used to measure the position of these markers each time the detector was rotated. The accuracy of the theodolite was found to be 0.25 mm over the distance used; the targets were 30 cm apart, resulting in an error in the angle measurement of less than 0.1° . This was the primary improvement made during 2011 data-taking relative to 2009.

Table 6.1: Angles at which positron beam data were taken.

Angle ($^\circ$)	# of Events (10^6)
0.0	5.973
6.0	11.65
11.8	6.509
16.5	7.253
20.9	6.532
24.4	6.278
30.8	5.779
36.2	5.861
41.6	6.429
47.7	9.061

6.4.1 Event selection cuts

Since T2 was the only plastic scintillator present for these data, it was the only detector required in the trigger. To help reduce background, approximately half the 43 ns cyclotron RF period was vetoed in the trigger; this region contained almost all the pions in the beam, and very few positrons. Due to the looseness of this requirement, many types of event besides beam positrons could cause triggers. The energy spectrum in BINA + CsI prior to any cuts is shown in Figure 6.5. The data shown were taken with the beam aligned with the crystal axis, that is, at 0° .

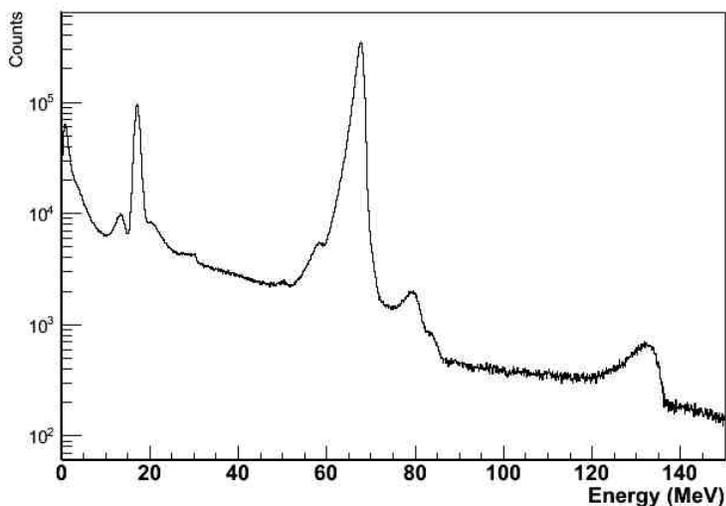


Figure 6.5: Beam particle energy measured by the crystals at 0° with no cuts applied.

The spectrum contains many backgrounds that can be removed with simple cuts. The low energy peak contains events without hits in the beam wire chambers in the region corresponding to the beam spot, suggesting they are not due to beam particles; they are possibly due to muons decaying backwards out of BINA. The peaks at 14 MeV and 18 MeV are due to beam pions and muons, respectively. The slight ridge around 30 MeV appears to be caused by pions decaying in flight; this was determined by Monte Carlo. The slight peak at 50 MeV and the clear shoulder near 60 MeV are the

photonuclear peaks. The main peak just below 70 MeV is due to beam positrons. The peak immediately to the right of the main peak is due to a positron and pion arriving simultaneously, and the shoulder to the right of that peak is due to a positron and muon arriving simultaneously. This was determined by the correspondence of the peaks to the sum of the individual particle peaks, and by the energy deposited in T2. The pion peak is larger than the muon peak because pions did not generally trigger, as they arrived in the region vetoed by the RF. Muons, however, overlapped with positrons in RF time. The peak around 130 MeV is due to events with two positrons.

The beam wire chambers, WC1 and WC2, were used to remove most events not due to beam particles. The incoming positron track was reconstructed and events outside the beam spot were removed. The background was further reduced by eliminating events with out-of-time hits in the wire chambers, and selecting the RF window for positrons.

Following these cuts, the spectrum contained events due to beam positrons and beam muons, which were reduced but not entirely eliminated by the RF cut. It was possible to remove them entirely with a cut on the energy deposited in T2; however, many positrons were removed as well. Any cut removing positrons in a way that depends on the energy deposited in the crystals could bias the response function measurement. In the absence of leakage from the crystals, the initial energy of the positron is equal to the energy deposited in T2 plus the energy deposited in BINA. In addition, GEANT4 simulation revealed that a significant fraction of events with high energy deposit in T2 were due to shower leakage backwards out of BINA. More detail is provided in Appendix C. The energy spectrum following the wire chamber position cut, the wire chamber timing cut, and the RF cut is shown in Figure 6.6.

6.4.2 Muon correction

Although the dependence of the response function on the T2 energy cut could perhaps be taken into account by Monte Carlo, it would be safer not to use any information so closely related to the BINA response. Thus, an alternative means to remove the muon contamination from the energy spectrum was developed. The energy spectrum in the RF window corresponding to muons is shown in Figure 6.7.

6.4. Data-taking

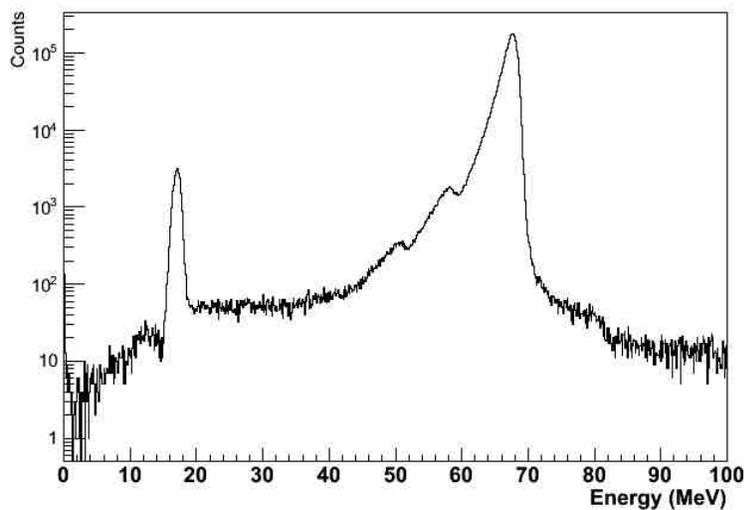


Figure 6.6: Energy measured by the crystals at 0° following event selection cuts (see text).

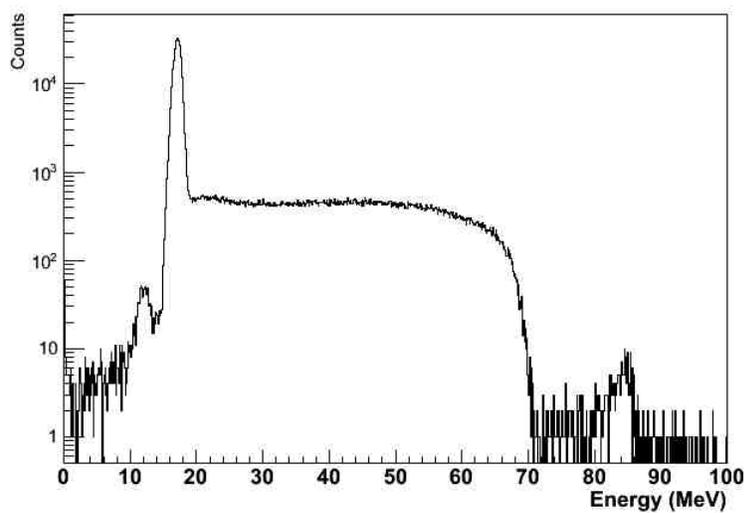


Figure 6.7: Energy measured by the crystals at 0° in the RF window corresponding to muons.

This spectrum is almost devoid of positrons, by the absence of any visible peak around 70 MeV. The large, almost flat distribution between 20 and 70 MeV arises from the long integration time of the BINA pulse, which is about $1 \mu\text{s}$. There was a significant probability that the muon would decay in this time, resulting in all or part of the decay positron energy being added to the muon energy. The cut-off of the energy spectrum corresponds to the energy of a beam muon plus the highest energy Michel positron.

The fraction of positrons below the muon peak is very small, on the order of 10^{-5} (see Appendix C, Figures C.4 and C.7). The BINA + CsI spectrum obtained by selecting muons by RF time can be used to determine the ratio of muon events in the peak to the total number of muon events. Since there are essentially no positron events in the energy region containing the muon peak, the total number of muons in the spectrum can be accurately determined from the size of the muon peak. This allows the muon component to be subtracted from the energy spectrum, giving the tail due to positrons only. Figure 6.8 shows the muon energy spectrum zoomed in on the peak, with red lines at 15 and 19 MeV defining the region considered as the peak. Due to energy loss in T2, these values change as a function of the angle between the beam and the axis of BINA.

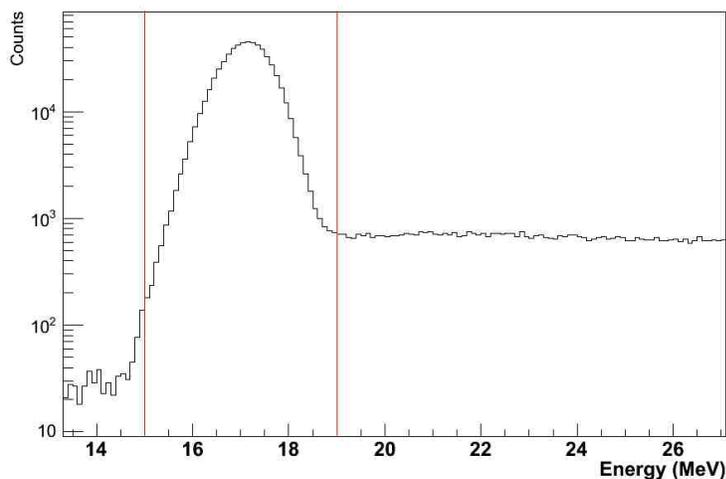


Figure 6.8: Energy measured by the crystals at 0° . The red lines indicate the region defined as the muon peak.

6.4. Data-taking

To correct Equation 6.1, both the number of muons in the tail and the total number of muons in the spectrum must be known; these quantities are linearly related to the number of events in the muon peak. From the BINA + CsI spectrum containing only muons, the scaling factor can be obtained. Let $N_{RF}^\mu(E)$ denote the energy spectrum in the RF window corresponding to muons, N_{peak}^μ be the number of events in the peak (defined by the red lines in Figure 6.8), N_{tail}^μ be the number of events in the tail (below E_{cut}), and N_{total}^μ be the total number of muon events in the spectrum. Then,

$$N_{peak}^\mu = \int_{15MeV}^{19MeV} N_{RF}^\mu(E) dE, \quad (6.3)$$

$$N_{tail}^\mu = \int_{0MeV}^{E_{cut}} N_{RF}^\mu(E) dE, \quad (6.4)$$

and

$$N_{total}^\mu = \int_0^\infty N_{RF}^\mu(E) dE. \quad (6.5)$$

Now considering the energy spectrum $N(E)$ with the RF cut corresponding to positrons, Figure 6.6, let N_{peak} be the number of events in the muon peak, N_{tail} be the number of events in the tail, and N_{total} be the total number of events in the spectrum; i.e.,

$$N_{peak} = \int_{15MeV}^{19MeV} N(E) dE, \quad (6.6)$$

$$N_{tail} = \int_{0MeV}^{E_{cut}} N(E) dE, \quad (6.7)$$

and

$$N_{total} = \int_0^\infty N(E) dE. \quad (6.8)$$

Let C_1 denote the number of muons present in $N(E)$ and C_2 denote the number of muons present below E_{cut} in $N(E)$; these are given by

$$C_1 = N_{peak} \times \frac{N_{total}^\mu}{N_{peak}^\mu} \quad (6.9)$$

and

$$C_2 = N_{peak} \times \frac{N_{tail}^\mu}{N_{peak}^\mu}. \quad (6.10)$$

6.4. Data-taking

The formula for obtaining the tail fraction following the muon correction, T_{corr} , is then

$$T_{corr} = \frac{N_{tail} - C_2}{N_{total} - C_1}. \quad (6.11)$$

This formula is valid so long as the shape of the muon energy spectrum in $N(E)$ is identical to the shape of the muon energy spectrum in $N_{RF}^\mu(E)$. The validity of this assumption can be determined by varying the RF window used for $N(E)$, from selecting as few muons as possible (Figure 6.6) to including the entire muon region in $N(E)$. The stability of T_{corr} as the RF window is varied can then be used to estimate the uncertainty of this correction. Table 6.2 shows the ratios of N_{peak}^μ to N_{tail}^μ and N_{peak}^μ to N_{total}^μ for each angle when the RF window corresponding to muons is selected.

Table 6.2: Properties of the muon energy spectra as a function of angle.

Angle ($^\circ$)	μ RF window (ns)	μ peak (MeV)	Peak / tail	Peak / total
0.0	13-16	15-19	0.734	0.688
6.0	12-15	15-19	0.731	0.686
11.8	12-15	15-19	0.731	0.686
16.5	11-14	15-19	0.730	0.683
20.9	11-14	15-19	0.730	0.688
24.4	11-14	15-19	0.728	0.684
30.8	10-13	14-18	0.721	0.673
36.2	13-16	14-18	0.710	0.678
41.6	13-16	13-17	0.711	0.686
47.7	13-16	12-16	0.670	0.660

These numbers, along with the total number of muons in the peak, can then be used in Equation 6.11 to calculate the tail fraction. Tables 6.3 to 6.12 show the results for the tail fraction after the muon correction, for each angle, as the RF window is broadened to include more muons.

Considering the size of the correction, the tail fraction is quite stable as the RF window is broadened. The largest relative variation was at 36.2° , where the spread was slightly under 2% of the tail at that angle. At each angle, the tail fraction was taken to be the average of the highest and lowest values in the tables shown above, and the systematic error due to the correction

6.4. Data-taking

Table 6.3: ΔT at 0° .

RF window	Tail fraction (%)
5-12	0.590
5-13	0.581
5-14	0.577
5-15	0.584
5-16	0.590

Table 6.8: ΔT at 6° .

RF window	Tail fraction (%)
3-11	0.602
3-12	0.603
3-13	0.601
3-14	0.601
3-15	0.602

Table 6.4: ΔT at 11.8° .

RF window	Tail fraction (%)
3-11	0.630
3-12	0.631
3-13	0.632
3-14	0.628
3-15	0.630

Table 6.9: ΔT at 16.5° .

RF window	Tail fraction (%)
1-10	0.674
1-11	0.680
1-12	0.675
1-13	0.673
1-14	0.674

Table 6.5: ΔT at 20.9° .

RF window	Tail fraction (%)
3-10	0.776
3-11	0.773
3-12	0.770
3-13	0.776
3-14	0.776

Table 6.10: ΔT at 24.4° .

RF window	Tail fraction (%)
3-10	0.874
3-11	0.884
3-12	0.882
3-13	0.873
3-14	0.874

Table 6.6: ΔT at 30.8° .

RF window	Tail fraction (%)
2-9	1.189
2-10	1.199
2-11	1.189
2-12	1.186
2-13	1.189

Table 6.11: ΔT at 36.2° .

RF window	Tail fraction (%)
6-12	1.386
6-13	1.378
6-14	1.342
6-15	1.365
6-16	1.386

Table 6.7: ΔT at 41.6° .

RF window	Tail fraction (%)
6-12	1.987
6-13	1.984
6-14	1.950
6-15	1.967
6-16	1.987

Table 6.12: ΔT at 47.7° .

RF window	Tail fraction (%)
6-12	3.316
6-13	3.313
6-14	3.258
6-15	3.278
6-16	3.316

6.5. Other systematics

was assigned to be the variation from the average. The tail fractions with statistical error and this source of systematic error are shown in Table 6.13. Note that the errors are given as percentages of the total number of events; i.e., they are the absolute errors on the tail fraction, not the relative errors.

Table 6.13: The tail fraction as a function of angle, with errors due to statistics and the variation in the muon-corrected values.

Angle (°)	Tail fraction (%)	Statistical error (%)	Muon correction error (%)
0.0	0.583	0.008	0.006
6.0	0.603	0.004	0.001
11.8	0.630	0.012	0.002
16.5	0.677	0.012	0.004
20.9	0.773	0.009	0.003
24.4	0.879	0.006	0.005
30.8	1.193	0.005	0.006
36.2	1.364	0.005	0.022
41.6	1.969	0.006	0.018
47.7	3.287	0.006	0.029

6.5 Other systematics

Many systematic effects besides the muon correction could potentially contribute to the error on the tail fraction. These come in two varieties: actual distortions in the energy spectra measured with the positron beam, and uncertainties in experimental parameters that must be known in order to translate the tail fractions in the positron beam data into the predicted tail fraction for the $\pi^+ \rightarrow e^+\nu_e$ energy spectrum. The latter can be evaluated only through simulation, and will be addressed in Section 6.6.

Since none of the cuts make any reference to what is measured by the crystals, they should leave the energy spectrum for good positron events undistorted. Thus, the only errors on the measured tail fraction *per se* arise from the presence of backgrounds and uncertainty on the measured energy itself (i.e. the BINA and CsI resolution and calibration).

6.5.1 Background

The level of pileup background can be obtained from Figure 6.6. The events to the right of the main peak were caused by a normal positron event happening at the same time as a Michel positron entered BINA. This is clear from the fact that the endpoint of the spectrum is at approximately 120 MeV, which is very close to the sum of the main peak and the highest-energy Michel positron. This could affect the tail fraction only in the case where an event that would have been in the <52 MeV region ended up above 52 MeV due to the presence of the extra particle. Events where a Michel positron made the trigger were almost completely removed by the cut requiring hits consistent in time in all three wire chambers. The fraction of events to the right of the main peak is approximately 0.2% of the main peak itself; therefore, since the tail itself will be seen to be approximately 3%, the impact on the tail will be less than 0.01%, which is negligible for the purposes of this analysis.

Other backgrounds could only be significant if they added events to the tail, since they would have to be present at a much higher level to be relevant in the rest of the spectrum. If muons are suppressed either by a T2 energy cut or a tighter RF cut, the very low energy part of the zero-degree spectrum shows that such backgrounds, if present at all, must be present only at a negligible level. The fraction of events under 10 MeV at 0 degrees, for example, is less than 10^{-5} . Thus, a flat background would contribute at a level of much less than 10^{-4} , and indeed the only type of background that could be present would have a very similar shape to the positron energy spectrum, which is implausible.

The possibility of scattering in the beamline leading to a low momentum tail in the beam momentum spread is taken into account by regarding this analysis as providing an upper limit on the tail due to the response function of the calorimeter.

6.5.2 Calibration and resolution

Using the normal data, the BINA calibration has been established to be accurate within 0.1 MeV (see Section 4.4). This uncertainty is taken into account by varying the value of E_{cut} by ± 0.1 MeV. The energy resolution could only affect the tail fraction insofar as it could move events below E_{cut} above it, and vice versa, so the actual error arises from the difference

between events moved into the tail region and those moved out of the tail region. Fitting the 0° peak with a Gaussian gives $\sigma = 0.6$ MeV, which includes the beam momentum spread. Even if this were entirely due to the energy resolution, its effect on the tail fraction could be neglected, since the energy range in which this can happen is only about 1 MeV. This was verified by adding Gaussian resolution to the simulated BINA energy; no significant effect on the tail fraction was observed.

6.6 Positron beam simulation

Since there are contributions to the $\pi^+ \rightarrow e^+\nu_e$ tail that were not present in the positron beam data, namely interactions upstream of the crystals and radiative decay, the ideal outcome of the response function analysis would be the development of a simulation that could reproduce the data at each angle, which could then be used to simulate $\pi^+ \rightarrow e^+\nu_e$ decay. Also, several sources of systematic uncertainty, such as the characteristics of the beam and the detector geometry, could only be assessed via simulation. Thus, the normal PIENU Monte Carlo was modified to match the conditions of the positron data-taking. This required the removal of all detectors from the simulation except for the three wire chambers, T2, and the crystals, and the replacement of the pion beam with a positron beam.

The beam parameters for the 0° data, reconstructed by the WC12 tracker, are shown in Figure 6.9; the quantities plotted are x , y , tx , and ty (see Section 4.5). x and y are the positions at $z = 0$, and tx and ty are the ratios of the x and y momenta to the z momentum. As is done for the normal pion beam, the correlations between these quantities were determined, and the Cholesky decomposition of the correlation matrix was multiplied by a vector of these values sampled independently, thereby reproducing the correlations [66].

The beam parameters changed considerably over the running period; the mean of the tx distribution for the data set taken at each angle is shown in Figure 6.10. Although the ideal procedure would be to use a separate beam at each angle in the simulation, the tail fraction is insensitive to the beam parameters except at high angles. Thus, one beam was used for the first seven angles, and a different beam was used for the 36.2° , 41.6° , and 47.7° angles.

6.6. Positron beam simulation

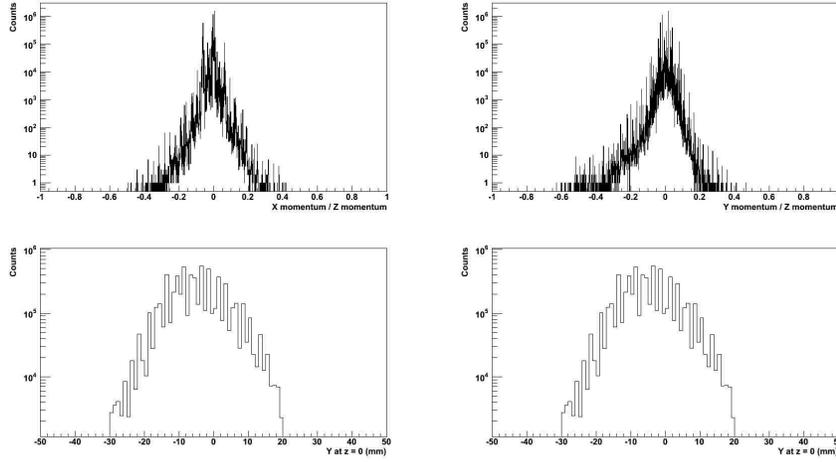


Figure 6.9: Reconstructed track parameters for the positron beam at 0° . The top-left and top-right panels show the ratio of the x and y momenta to the z momentum, and the bottom-left and bottom-right panels show the reconstructed x and y positions at $z = 0$.

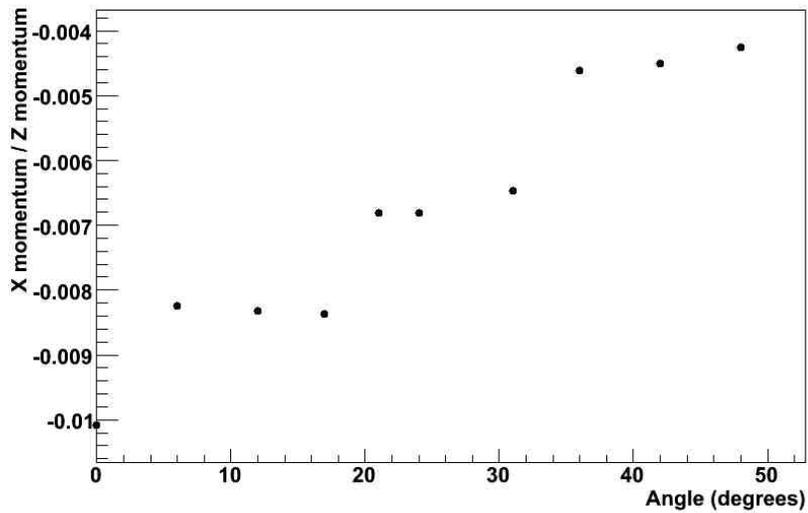


Figure 6.10: Mean of the distribution of the x momentum, normalized to the z momentum, as a function of angle.

6.6. Positron beam simulation

The energy spectrum at 0° is shown in Figure 6.11, from both data and Monte Carlo. For this plot, events depositing more than 1.6 MeV in T2 have been removed, so as to show clearly the shape of the energy spectrum due to leakage, and the photonuclear peaks.

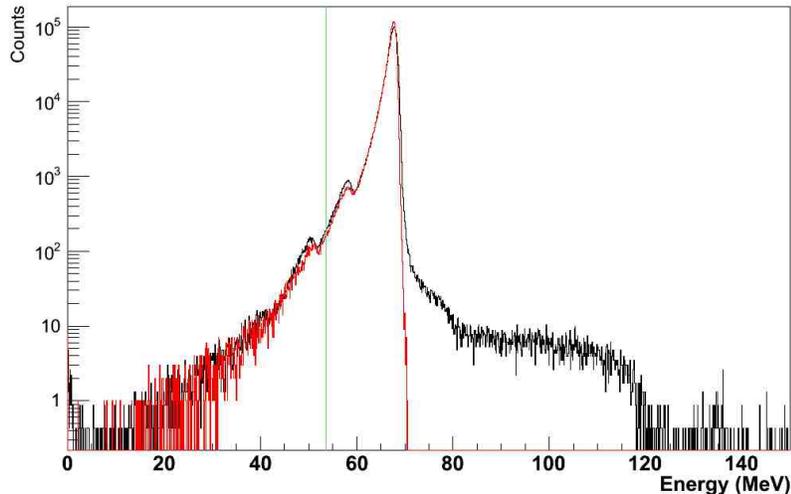


Figure 6.11: The energy spectrum from a 70 MeV positron beam parallel to the crystal axis. Data is shown in black and simulation is shown in red. The histograms are normalized to have the same total number of events. The green line shows the value of E_{cut} .

Several differences between data and simulation are apparent. There is no pileup in the simulation, resulting in the difference to the right of the main peak. There is also no energy resolution applied to the simulated BINA energy, although the simulated CsI energies have a Gaussian resolution applied; the main peak in the simulation is thus sharper than in the data. Finally, the photonuclear peaks (see Section 6.3.1) are larger in the data than in the simulation. Only the last disagreement has any significant impact on the tail, as pileup and crystal resolution are negligible. The position of the second photonuclear peak just below the standard value of E_{cut} means that its amplitude must be correct for the tail to be properly reproduced by the simulation. It in fact contains the majority of the tail at 0° , although its importance diminishes at large angles. The energy spec-

6.6. Positron beam simulation

tra at 11.8° , 24.4° , 36.2° , and 47.7° are shown in Figure 6.12. The same T2 energy cut is applied, scaled by the path length difference through T2 as it is rotated. Qualitatively, the differences between data and simulation at 0° persist as the angle increases.

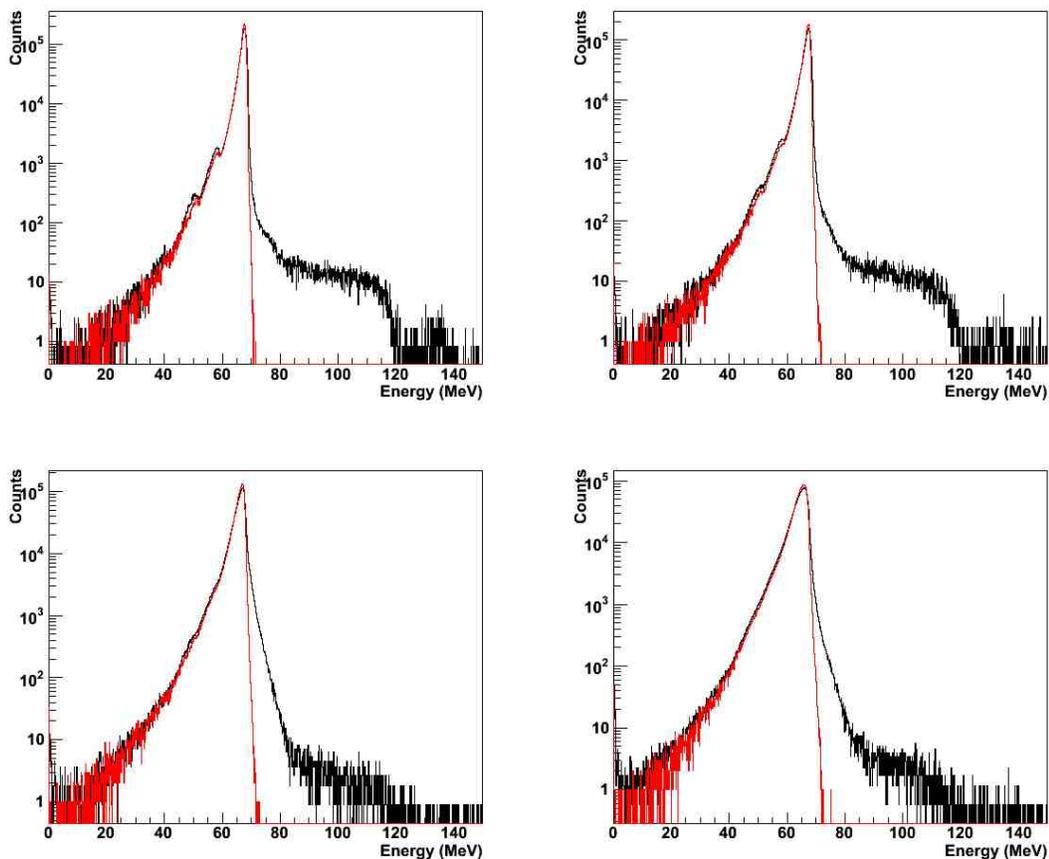


Figure 6.12: The energy spectrum from a 70 MeV positron beam at 11.8° (top left), 24.4° (top right), 36.2° (bottom left), and 47.7° (bottom right) to the crystal axis. Data is shown in black and simulation is shown in red.

The effect of uncertainty on the angle between the crystal axis and the beam, the centre of rotation of the crystal array, the momentum of the beam, and the x and y divergences of the beam were assessed via simulation. For

the most part, these errors were negligible, although at the highest angles they contribute significantly. Appendix D shows the complete results.

The error on the crystal-beam angle was taken as $\pm 0.1^\circ$, determined from a test of the accuracy of the measurement system (see Section 6.4). The position of the centre of rotation was fixed by the shaft attached to the cart on which the crystals sat; this was positioned according to alignment markings made by the TRIUMF beamlines group, so an error of ± 1 mm was assigned. The uncertainty in the divergence of the beam was taken from the spread in Figure 6.10. For the first seven angles, the mean of the tx distribution was varied by ± 0.02 , and for the last three, it was varied by ± 0.004 .

In principle, the beam momentum was fixed by the fields of the bending magnets, which were monitored by NMR probes to less than 0.1% and maintained at a constant value. Thus, the momentum should be known to high accuracy, and was nominally set to 70 MeV/c for this data. However, the ratio of the positions of the positron and muon peaks was inconsistent with this value, and instead suggested a momentum of approximately 74 MeV/c. This degree of deviation from the field determined by the NMR probes was not plausible, but in light of this fact an error was assigned of 1%, or 0.7 MeV/c. That the usual settings for pion data-taking correspond to 75 MeV/c was verified to within a few hundred keV by the pion stopping position (see Section 7.1.1). The uncertainty in the beam momentum is the largest uncertainty in the positron beam data.

The tail fraction as a function of angle is shown in Figure 6.13, and the difference between the MC and data tail fractions is shown in Figure 6.14. The uncertainties from the various sources of error have been added in quadrature for visualization purposes, but note that the errors at each angle due to beam momentum, beam divergence, and the position of the centre of rotation are not independent; they will shift each point in the same direction.

Some difference is expected due to the photonuclear peaks (see Figure 6.6); this difference should increase as a function of angle, since the tail itself increases, making it more likely that an event in the first photonuclear peak will be in the tail. However, the size of the difference at the last two angles cannot be explained by this effect. The difference in counts at 0° in the region containing both photonuclear peaks is approximately 0.25% of the

6.7. Determining the PIENU tail fraction

total spectrum, and most of the first peak remains out of the tail even at 47.7° .

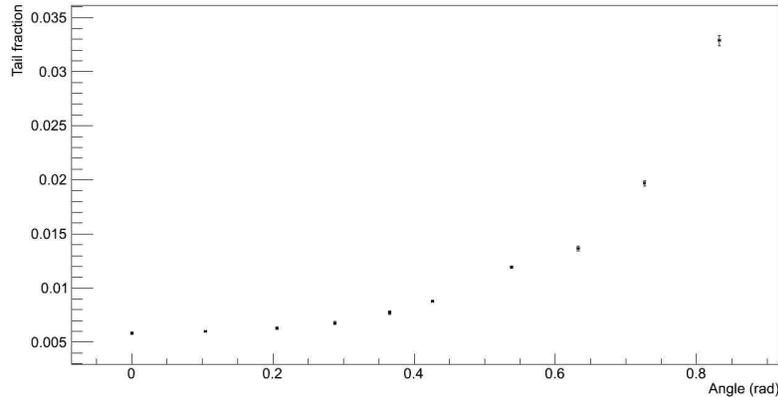


Figure 6.13: The tail fraction as a function of angle in the positron beam data.

Currently, the cause of the discrepancy at the two highest angles is unknown. Two possibilities must be considered: that it is due to a genuine feature of the calorimeter response that is not included in the simulation, or that it is due to some unconsidered or underestimated systematic error in the data. In the latter case, the tail fraction for $\pi^+ \rightarrow e^+\nu_e$ events can be taken from simulation, and corrected based on the difference measured for the first eight angles, which is consistent with being entirely due to the photonuclear effect. In the former case, the simulated $\pi^+ \rightarrow e^+\nu_e$ tail must be corrected for the difference measured at all angles. In order to ensure that the uncertainty was properly covered, corrections for both cases were determined, the average used as the correction, and half the difference taken as the 1σ error.

6.7 Determining the PIENU tail fraction

The simulated $\pi^+ \rightarrow e^+\nu_e$ energy spectrum is shown in Figure 6.15; the normal acceptance cut was applied, requiring at least one reconstructed track to be within 60 mm of the centre of WC3, but no other cuts were

6.7. Determining the PIENU tail fraction

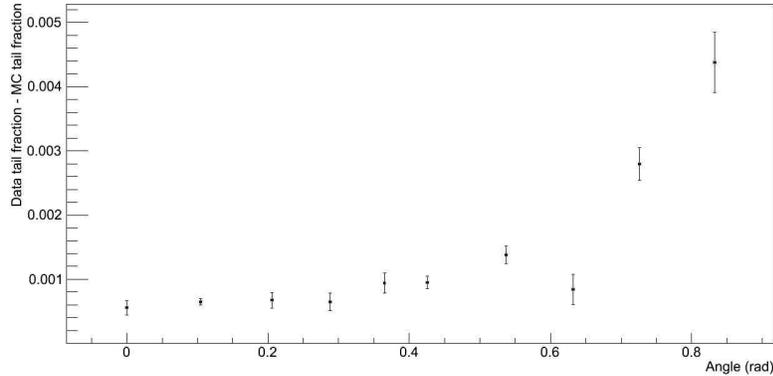


Figure 6.14: The difference between the tail fraction in the positron beam data and the tail fraction from a simulated positron beam as a function of angle.

applied besides the trigger condition (energy in T1 and T2) and that the pion decayed at rest within the target. The high energy tail is caused by $\pi^+ \rightarrow e^+\nu_e\gamma$ events. The proportion of the spectrum less than 52 MeV is 2.99%, almost as high as the tail at the largest angle at which positron beam data were taken. The reason for this, and the long, almost flat, part of the spectrum extending down to zero, is interactions in the scintillators upstream of BINA, particularly the target. Thus, the tail fraction must be extracted from a combination of the $\pi^+ \rightarrow e^+\nu_e$ simulation and the positron beam data; it cannot be obtained from the latter alone. The spectrum without events undergoing Bhabha scattering in the target is shown in Figure 6.16; the tail fraction is 1.82%.

To obtain a correction for the simulated $\pi^+ \rightarrow e^+\nu_e$ tail fraction, the tail fraction in the positron beam data must be averaged as a function of angle. Because the tail fraction is a smooth function of the angle, it can be fitted with a polynomial. This is shown in Figure 6.17 for both data and Monte Carlo. For this plot, the correlated errors (from point to point) have been omitted. The average of the fitted function can then be taken, weighted for

6.7. Determining the PIENU tail fraction

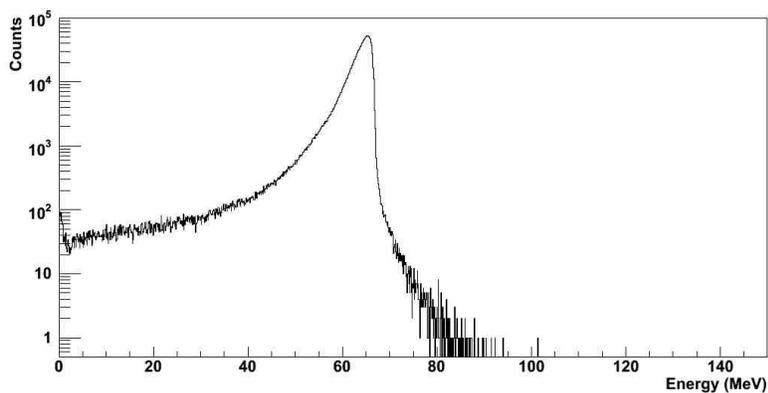


Figure 6.15: The simulated BINA+CsI spectrum from $\pi^+ \rightarrow e^+\nu_e$ decay and $\pi^+ \rightarrow e^+\nu_e\gamma$ decay.

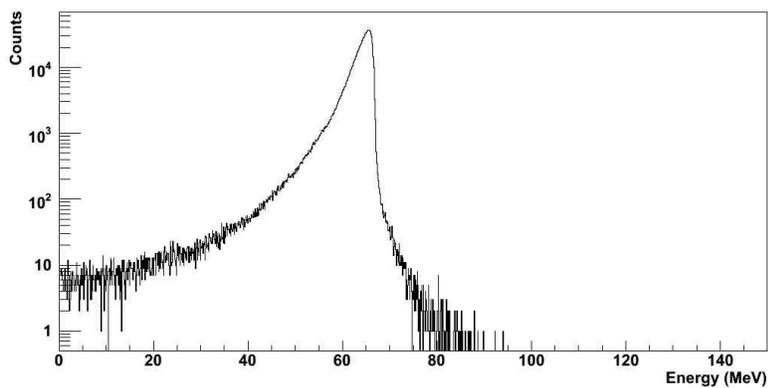


Figure 6.16: The simulated BINA+CsI spectrum from $\pi^+ \rightarrow e^+\nu_e$ decay and $\pi^+ \rightarrow e^+\nu_e\gamma$ decay, excluding events that underwent Bhabha scattering in the target.

6.7. Determining the PIENU tail fraction

an isotropic distribution, according to the formula

$$T_{average} = \frac{\int_0^{\theta_{max}} w(\theta) f(\theta) d\theta}{\int_0^{\theta_{max}} w(\theta) d\theta}. \quad (6.12)$$

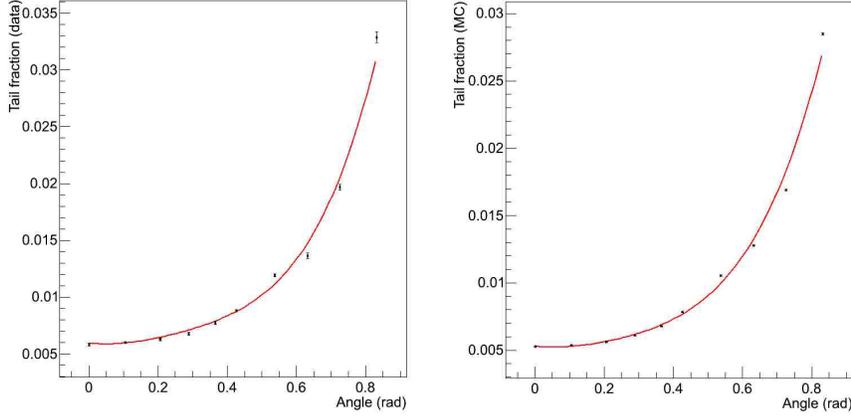


Figure 6.17: The tail fraction as a function of angle for the positron beam data (left) and Monte Carlo (right), fitted to a fourth-degree polynomial.

Here $w(\theta)$ is simply $\sin\theta$. The results are $T_{average}^{data} = 1.42\% \pm 0.01\%$ and $T_{average}^{MC} = 1.26\% \pm 0.01\%$; the uncertainties are those returned by the polynomial fit on the weighted average, due to the uncertainty in the parameters. In order to take the correlated errors into account, each point was shifted by the 1-sigma error and the fit was redone, giving $1.45\% \pm 0.01\%$ and $1.39\% \pm 0.01\%$. The correction is obtained by subtracting the MC tail from the data tail; this gives a correction of $0.16\% \pm 0.03\%$.

In this case $E_{cut} = 53.7$ MeV, which takes into account the difference in peak position between the $\pi^+ \rightarrow e^+\nu_e$ case and the positron beam case. This difference is primarily due to the extra material present in the $\pi^+ \rightarrow e^+\nu_e$ case, although the starting energy is slightly different as well. The peak for simulated $\pi^+ \rightarrow e^+\nu_e$ events emitted at small angles is shown in Figure 6.18; the difference between this peak and the 0° peak is 1.7 MeV. Thus, an E_{cut}

6.7. Determining the PIENU tail fraction

value of 52 MeV in the $\pi^+ \rightarrow e^+\nu_e$ case corresponds to 53.7 MeV in the positron beam case.

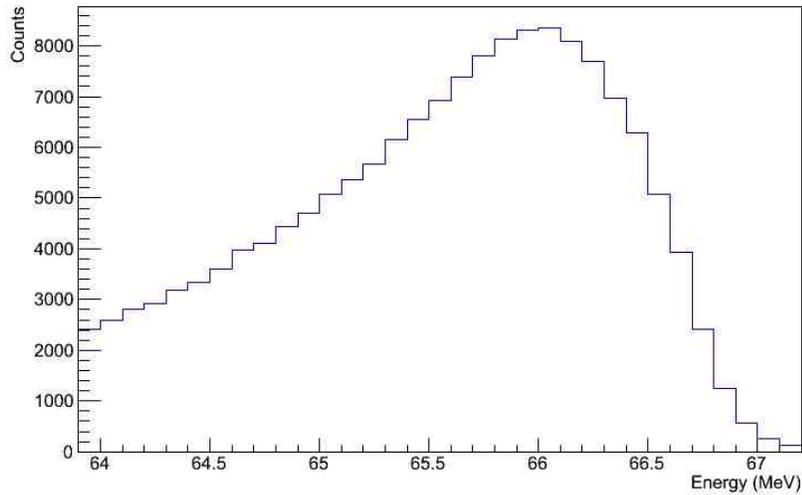


Figure 6.18: Simulated energy spectrum for $\pi^+ \rightarrow e^+\nu_e$ events emitted at small angles.

The difference between data and Monte Carlo, with the first eight points fitted to a straight line, is shown in Figure 6.19. Averaging this, again using $w(\theta) = \sin \theta$, gives $0.11\% \pm 0.01\%$. Again, the correlated errors were used to shift each point by the same amount; the resulting averages were $0.14\% \pm 0.01\%$ and $0.08\% \pm 0.01\%$.

The final correction is obtained by averaging the highest and lowest corrections, and taking half the difference as the error. The largest correction comes from the method with the polynomial fit, and is 0.19%. The smallest comes from the method with the straight line fit, and is 0.08%. The result is therefore a final correction to the $\pi^+ \rightarrow e^+\nu_e$ tail of $0.14\% \pm 0.06\%$. The simulated $\pi^+ \rightarrow e^+\nu_e$ tail fraction was 2.99%; thus, the corrected value is $T = 3.13\% \pm 0.06\%$.

Since the correction takes into account the difference between the response function obtained from data and the simulated response function, the only

6.7. Determining the PIENU tail fraction

uncertainties that must be taken into account in the $\pi^+ \rightarrow e^+\nu_e$ simulation itself are those arising from the detector geometry and the pion stopping position. The uncertainties on the detector positions and thicknesses are very small [8], and the simulated pion stopping distribution agrees very well with the data (see Section 7.1.1). The variation in the simulated tail fraction from these effects was found to be $< 0.01\%$, which is negligible compared to the error on the correction.

The only further uncertainty entering into the upper limit on the tail fraction comes from the uncertainty on the energy calibration itself (see Section 4.4) of 0.1 MeV between E_{cut} and the peak of the $\pi^+ \rightarrow e^+\nu_e$ energy spectrum. This was obtained simply from the tail < 51.9 MeV and the tail < 52.1 MeV, which varied by 0.04% from the tail < 52.0 MeV. This increased the error on the upper limit from 0.06% to 0.07%.

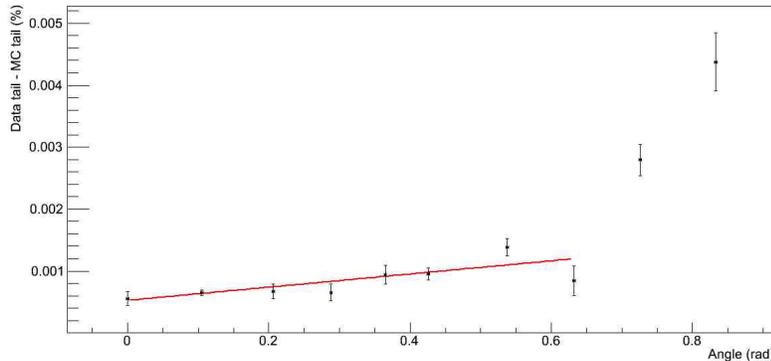


Figure 6.19: The tail fraction in the positron beam data minus the simulated tail fraction as a function of angle. The first 8 angles are fitted to a straight line.

This result for the upper limit on the tail correction is combined with the lower limit to give the best estimate of the $\pi^+ \rightarrow e^+\nu_e$ tail. The limits are combined in the following way: a probability distribution is constructed by assuming the tail is equally likely to be any value above the lower limit and any value below the upper limit, and that the errors on each quantity are Gaussian. The peak of the resulting probability distribution is then taken

as the central value for the tail. This procedure is graphically illustrated in Figure 6.20, for the lower limit obtained from the 2010 data.

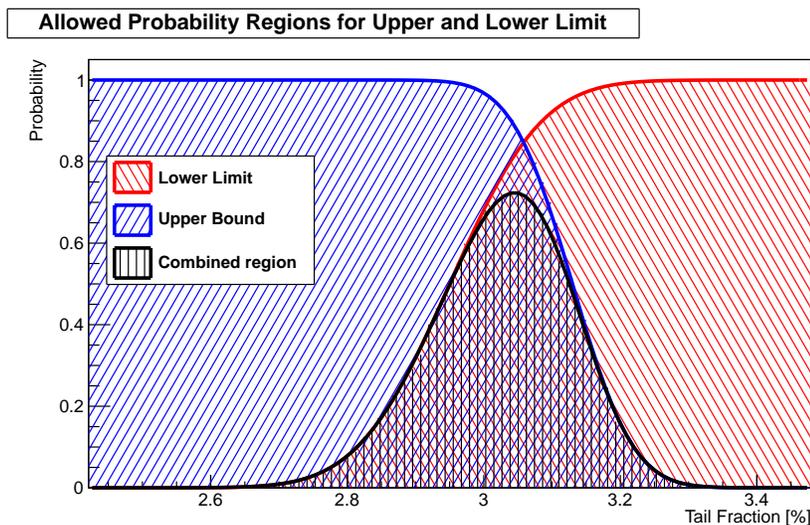


Figure 6.20: The probability distribution obtained by combining the upper and lower limits.

6.8 Tail as a function of R and E_{cut}

The stability of the branching ratio as the acceptance cut and E_{cut} are varied is a crucial systematic check of the analysis. The tail correction depends strongly on both of these parameters. Obtaining corrections as a function of E_{cut} is straightforward; the only modification is the value up to which the spectra are integrated. Obtaining corrections for acceptance cuts less than 60 mm is also straightforward, as the positron beam tail fractions must simply be averaged up to a smaller angle. However, since the largest angle at which positron beam data were taken corresponds to $R = 62$ mm, obtaining corrections above this value is not straightforward.

The obvious approach is to extrapolate the function obtained from fitting the tail fractions and average that; however, when this approach was tested via simulation, it was found to significantly underestimate the tail for angles above those that were included in the fit. When the function in the right-

hand panel of Figure 6.17 was extrapolated to 52.0° , 54.0° , and 56.0° , the tail fractions obtained were 3.63%, 4.14%, and 4.72%. When these angles were simulated directly, the tail fractions were 6.40%, 10.26%, and 16.2%. For reference, $R = 80$ mm corresponds to 55.1° ; to verify the analysis, extending the acceptance cut to 90 mm is desirable.

To a reasonable approximation, the size of the correction to the tail for a given acceptance cut should be the same as the size of the correction for a value of E_{cut} for which the uncorrected size of the tail is the same. The validity of this approximation can be tested using acceptance cuts < 60 mm (this is the reason for including results for the upper limit for $E_{cut} < 50$ MeV).

Consider the uncorrected tail fractions for $R < 30$ mm, 40 mm, and 50 mm; these are approximately equal to the tail fractions at 48 MeV, at 49.5 MeV, and halfway between 50.5 MeV and 51 MeV. The corrections for these values of R are $0.09\% \pm 0.02\%$, $0.10\% \pm 0.03\%$, and $0.12\% \pm 0.04\%$; for these values of E_{cut} , they are $0.05\% \pm 0.05\%$, $0.07\% \pm 0.05\%$, and $0.10\% \pm 0.06\%$ (averaging 50.5 MeV and 51 MeV). These are consistent within errors. The uncorrected tail fractions for $R < 70$ mm, 80 mm, and 90 mm are 3.46%, 3.99%, and 4.63%. These are approximately equal to the uncorrected tail fractions for 53 MeV (3.40%), 54 MeV (3.92%), and 55 MeV (4.58%); the same corrections are therefore applied.

It is also desirable to check the stability of the branching ratio in statistically independent regions of R ; to this end, the tail correction in 10 mm slices is given in Table 6.16. For large R the tail grows quickly, and the correction cannot be obtained by using a value of E_{cut} for which the tail is comparable. Instead, the tails for large R slices were estimated using the behaviour of the tail as larger R values are included. For example, the tail in the 60-70 mm slice is equal to the average of the 0-60 mm tail and the 60-70 mm tail, weighted by the number of events in each slice. The relative error of the correction is taken to be the same as the 50-60 mm slice.

6.8. Tail as a function of R and E_{cut}

Table 6.14: Corrections and resulting $\pi^+ \rightarrow e^+ \nu_e$ tail fractions as a function of E_{cut} .

E_{cut} (MeV)	Correction (%)	Upper limit (%)	Combined tail (%)
48.0	0.05 ± 0.05	2.02 ± 0.05	
48.5	0.05 ± 0.05	2.11 ± 0.05	
49.0	0.06 ± 0.05	2.22 ± 0.05	
49.5	0.07 ± 0.05	2.34 ± 0.05	
50.0	0.08 ± 0.05	2.46 ± 0.06	2.34 ± 0.10
50.5	0.09 ± 0.05	2.60 ± 0.06	2.49 ± 0.10
51.0	0.11 ± 0.06	2.76 ± 0.07	2.65 ± 0.10
51.5	0.12 ± 0.06	2.93 ± 0.07	2.83 ± 0.10
52.0	0.14 ± 0.06	3.13 ± 0.07	3.03 ± 0.10
52.5	0.16 ± 0.07	3.34 ± 0.08	3.25 ± 0.10
53.0	0.18 ± 0.07	3.58 ± 0.08	3.50 ± 0.10
53.5	0.21 ± 0.08	3.86 ± 0.09	3.78 ± 0.11
54.0	0.24 ± 0.08	4.16 ± 0.10	4.09 ± 0.12
54.5	0.28 ± 0.09	4.51 ± 0.11	4.45 ± 0.13
55.0	0.33 ± 0.09	4.91 ± 0.12	4.85 ± 0.13

Table 6.15: Tail fractions as a function of the maximum radius in which events are accepted.

Max R (mm)	Correction (%)	Upper limit (%)	Combined tail (%)
20	0.07 ± 0.02	1.79 ± 0.03	
30	0.09 ± 0.02	1.99 ± 0.03	1.91 ± 0.08
40	0.10 ± 0.03	2.31 ± 0.04	2.21 ± 0.09
50	0.12 ± 0.04	2.69 ± 0.05	2.58 ± 0.10
60	0.14 ± 0.06	3.13 ± 0.07	3.03 ± 0.10
70	0.18 ± 0.07	3.64 ± 0.08	3.55 ± 0.10
80	0.24 ± 0.08	4.23 ± 0.10	4.12 ± 0.15
90	0.33 ± 0.09	4.96 ± 0.11	4.84 ± 0.17

6.9. Summary of Chapter 6

Table 6.16: Tail fractions as a function of the radius in which events are accepted.

R range (mm)	Correction (%)	Upper limit (%)	Combined tail (%)
20-30	0.10 ± 0.03	2.18 ± 0.04	2.03 ± 0.13
30-40	0.12 ± 0.04	2.84 ± 0.05	2.63 ± 0.18
40-50	0.16 ± 0.08	3.70 ± 0.09	3.60 ± 0.12
50-60	0.23 ± 0.14	4.88 ± 0.15	4.77 ± 0.18
60-70	0.42 ± 0.26	6.48 ± 0.27	6.52 ± 0.26
70-80	0.64 ± 0.39	8.76 ± 0.41	8.99 ± 0.36
80-90	1.2 ± 0.73	12.3 ± 0.75	12.5 ± 0.60

6.9 Summary of Chapter 6

The combined tail fraction is used to correct the branching ratio. The combined tail fraction is defined as the peak of the probability distribution obtained by combining the upper and lower limits, as shown in Figure 6.20. The formula for obtaining the corrected branching ratio is

$$BR_{tailcorrected} = \frac{1}{1-T} BR_{raw}. \quad (6.13)$$

For the nominal values of $R = 60$ mm and $E_{cut} = 52$ MeV, the combined tail fraction is $T = 3.03\% \pm 0.10\%$, resulting in a multiplicative correction of 1.0312 ± 0.0011 . The raw branching ratio with the T2 pileup cut applied, given in Table 5.4, is $BR_{raw} = (1.2005 \pm 0.0011) \times 10^{-4}$. Applying the tail correction gives

$$BR_{tailcorrected} = (1.2380 \pm 0.0011(\text{stat}) \pm 0.0013(\text{tail})) \times 10^{-4}. \quad (6.14)$$

This result is still blinded by an unknown factor between $\pm 0.5\%$, so it cannot be directly compared to the theoretical prediction. The corrections that must be applied to the tail-corrected branching ratio are described in the next chapter.

Chapter 7

Other Corrections

Three other, much smaller, corrections must be made to the raw branching ratio. These are referred to as the acceptance correction, the muon decay-in-flight correction, and the t_0 correction. The first corrects for the difference in acceptance between $\pi^+ \rightarrow e^+\nu_e$ and $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ events, which arises primarily from the energy dependence of multiple scattering. The second corrects for the presence of muon decay-in-flight events in the high energy time spectrum, which are indistinguishable from $\pi^+ \rightarrow e^+\nu_e$ events based on their timing. The last corrects for the slight difference in the measured time of a very low energy positron and a higher energy positron. Their values are shown in Table 7.1.

Table 7.1: The small corrections that must be made to the branching ratio. The values are multiplied by the tail-corrected branching ratio to give the final result.

Correction	Value	Error
Acceptance	0.9991	0.0003
μ -DIF	0.9983	< 0.0001
t_0	1.0006	0.0003

7.1 Acceptance correction

Two effects could potentially change the ratio of $\pi^+ \rightarrow e^+\nu_e$ events to $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ events within the geometrical acceptance: the extra spread in the starting position distribution of the decay positron caused by the distance travelled by the 4.1 MeV muon, and energy dependent interactions upstream of BINA. Both of these can be assessed with Monte Carlo, since they depend on well-understood electromagnetic physics.

Each of the 2010, 2011, and 2012 data sets required its own correction, since the beam momentum (and thus the pion stopping position) and de-

7.1. Acceptance correction

tector geometry were slightly different. For each case, 1 billion of each decay were simulated, and the ratio of events within different acceptance regions was calculated. The ratio of $\pi^+ \rightarrow e^+\nu_e$ to $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ events within different radii of the centre of WC3 is shown in Figure 7.1. The errors are due to Monte Carlo statistics. The systematic error on the correction was obtained by varying several parameters in the simulation: the position and width of the pion stopping distribution, the positions and thicknesses of various detectors, and the trigger thresholds in T1 and T2.

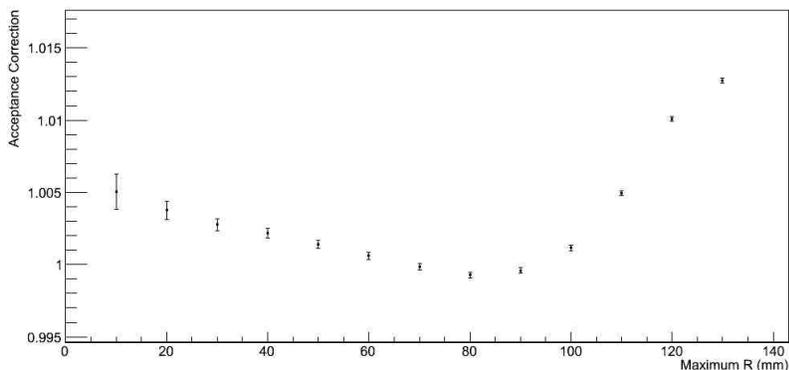


Figure 7.1: The ratio of $\pi^+ \rightarrow e^+\nu_e$ events to $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ events within different radii of the centre of WC3, as reconstructed by the S3WC3 tracker.

7.1.1 Pion stopping position

The pion stopping position was calculated from the point of closest approach in the tracks reconstructed by the S12 and S3WC3 trackers. This was done in both data and Monte Carlo; a comparison of the two distributions for the 2012 data set is shown in Figure 7.2. Figure 7.3 shows the difference between the reconstructed stopping position and the actual stopping position in the simulation, Figure 7.4 shows the variation in the correction with the mean of the reconstructed stopping position, and Figure 7.5 shows the variation in the acceptance correction with the width of the reconstructed stopping position.

7.1. Acceptance correction

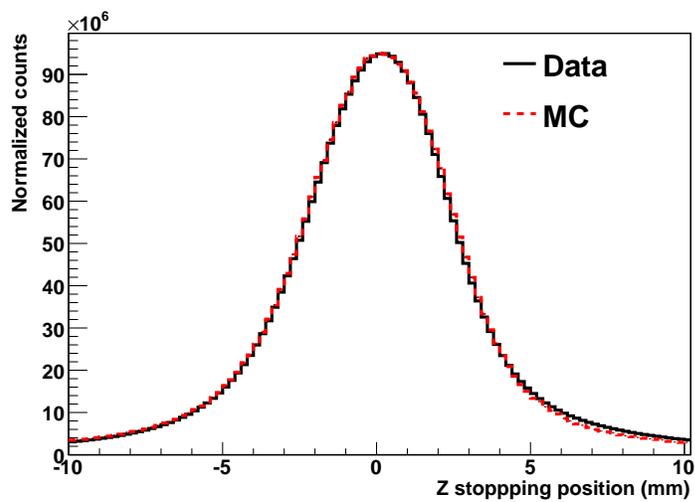


Figure 7.2: The z coordinate of the reconstructed pion stopping position. Data is shown in black and MC is shown in red.

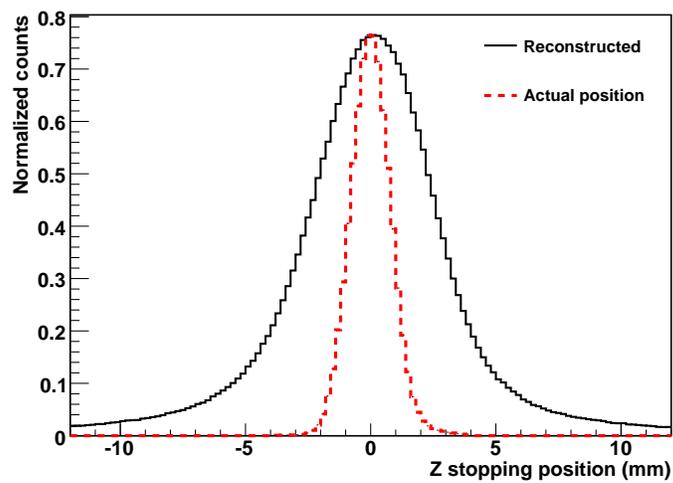


Figure 7.3: A comparison of the actual and reconstructed pion stopping positions in MC.

7.1. Acceptance correction

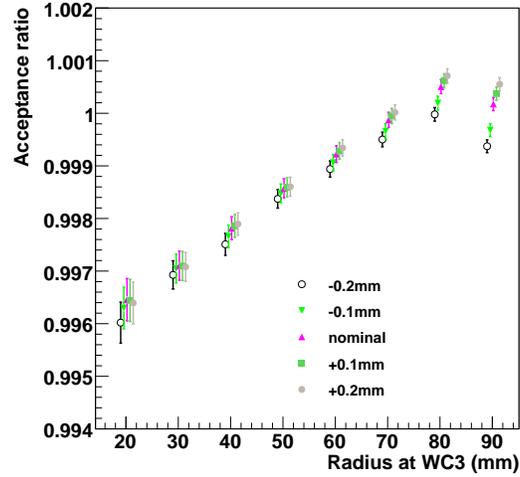


Figure 7.4: Variation of the acceptance correction with the peak value of the reconstructed pion stopping position. The peak was varied by ± 0.2 mm; the largest variation in the correction was approximately $\pm 0.05\%$, for $R < 90$ mm.

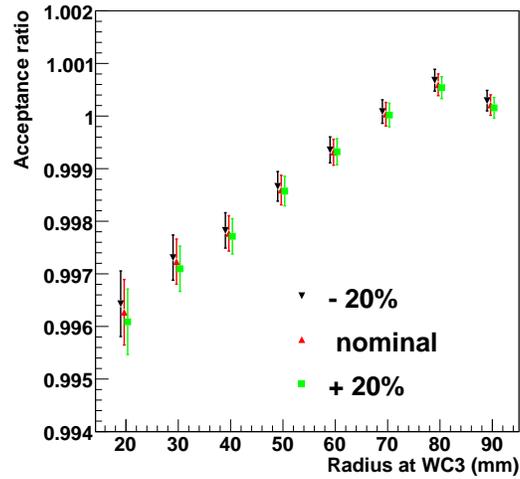


Figure 7.5: Variation of the acceptance correction with the width of the reconstructed pion stopping position.

7.1.2 Detector geometry

The relevant detector parameters for the pion stopping position are the thicknesses of Tg, S3, and T1 and the positions of Tg, S1, S2, S3, and WC3; these are known to the level of microns and tens of microns, respectively. The effect on the correction of varying each of these parameters within their uncertainties was found to be much less than 0.0001.

7.1.3 Trigger thresholds

The energy thresholds in T1 and T2 were each approximately 0.1 MeV. Since the acceptance correction could be sensitive to the exact value of the thresholds, they were varied by ± 25 keV. The effect on the acceptance correction is shown in Figure 7.6. At $R = 60$, it was found to be 0.0001. These uncertainties are all negligible in the present analysis.

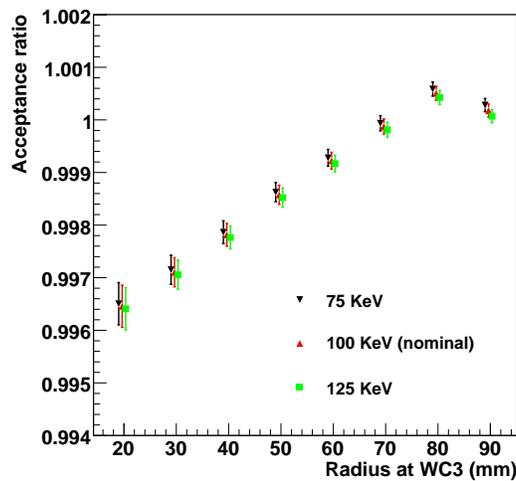


Figure 7.6: Variation of the acceptance correction with the thresholds in the T1 and T2 counters.

All the uncertainties on this correction, both statistical and systematic, are on the level of a few parts in 10^{-8} , which is negligible for the purposes of this analysis (see Chapter 8). The variations are at the same level as the 2010 analysis, for which the error was 0.0003 at the nominal value of R .

7.2 Muon decay-in-flight correction

Despite the very low energy of muons arising from $\pi \rightarrow \mu\nu$ decay, the influence of decays-in-flight upon the branching ratio is not negligible. Figure 7.7 shows a simulation of the decay time of muons that were not at rest when they decayed; the stopping time is approximately 19 ps. The probability of this occurring can be approximated by the proportion of the muon decay curve between 0 and 19 ps; i.e.

$$1 - e^{\frac{-0.019}{2197}} = 8.3 \times 10^{-6}. \quad (7.1)$$

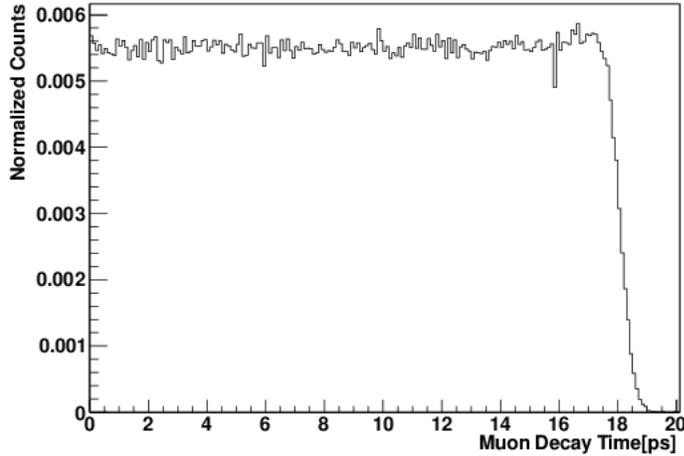


Figure 7.7: The decay time of muons in the target with non-zero kinetic energy at the time of the decay.

Since the time of these events is distributed according to a pion lifetime, any event with measured energy above E_{cut} will be misidentified as a $\pi^+ \rightarrow e^+\nu_e$ event by the time spectrum fit. Figure 7.8 shows simulations of the measured energy of both muon decay-in-flight and muon decay-at-rest events. The proportion of the decay-in-flight spectrum above 52 MeV was found to be 2.37%, giving a total correction of $0.0237 \times 8.3 \times 10^{-6} = 1.97 \times 10^{-7}$. Given the level of agreement in the measured energy spectra between Monte Carlo and data for both $\pi^+ \rightarrow \mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ events and the positron beam, the relative error on the proportion of the spectrum above

52 MeV is on the order of a few percent, which results in an uncertainty on the correction of less than 10^{-8} , which is negligible for this analysis.

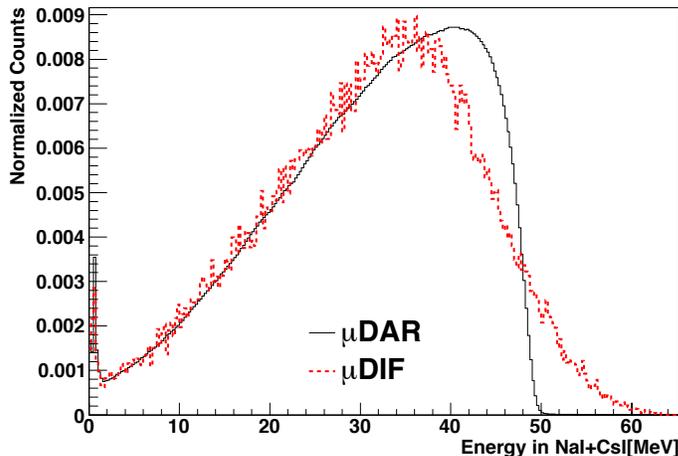


Figure 7.8: Simulated energy spectra measured by BINA+CsI for muon decays-at-rest (black) and decays-in-flight (red).

7.3 t_0 correction

If the shape of the T1 waveform were to depend on the positron energy, there could also be a dependence of the fitted time on the positron energy. This would result in an effectively different value of t_0 for $\pi^+ \rightarrow e^+ \nu_e$ and $\pi^+ \rightarrow \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ events. To determine if this effect exists, the time spectra for different energy regions were plotted and t_0 was obtained by fitting the edge with a step function with Gaussian resolution. The correction obtained using 2010 data was 1.0004 ± 0.0005 [3]. It was done again using 2011 data, obtaining 1.0006 ± 0.0003 [65]. The error is due to statistics; it could be reduced still further using 2012 data. For this analysis, the 2011 value was used.

7.4 Stability of the corrected branching ratio

Two crucial checks on the analysis are the stability of the branching ratio when the radius in which events are accepted is varied, and when the value

7.4. Stability of the corrected branching ratio

of E_{cut} is varied. The fitting functions and the corrections both change with these parameters, so any variation is an indication of a flaw in the analysis. Figure 7.9 shows the raw and corrected branching ratios for different acceptance cuts; in each case the results are plotted both with and without the T2 pileup cut applied. The error bars represent the error on the change from the first point; for the statistical error, this is done by taking the square root of the squared sum of the difference of the statistical error at each point. For the tail correction, the difference in the error from the first point is taken as the error on the change: if the tail is above its central value at the first point, it will also be above its central value at each subsequent point.

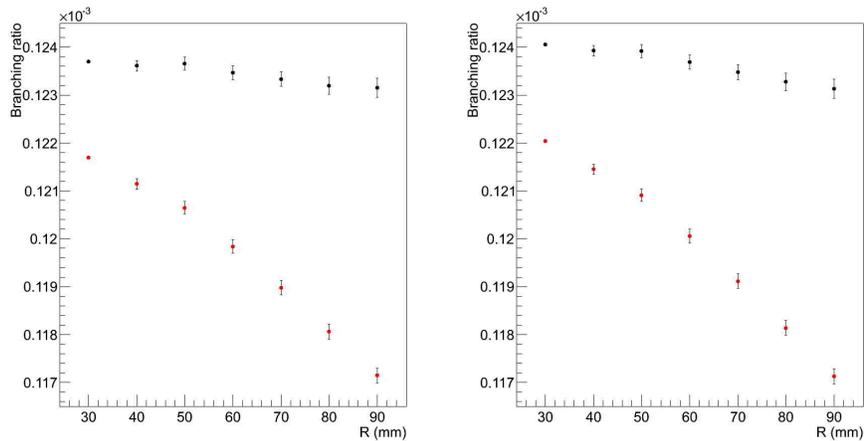


Figure 7.9: Variation of the branching ratio as the radius in which events are accepted is varied. The left-hand panel shows the results with the T2 pileup cut, and the right-hand panel shows the results without the T2 pileup cut. The red points show the raw branching ratio, and the black points show the branching ratio after all corrections. The error bars represent the error on the change from the first point.

In both cases a trend is apparent from $R > 60$ mm onwards, although it is reduced when the T2 pileup cut is applied. The fit was also done for events in 10 mm rings from the centre of WC3; these have the advantage of being statistically independent, although the errors are large. The error on the tail is again taken as the difference in the error from the first point; it is then added in quadrature with the statistical error. The results are

7.4. Stability of the corrected branching ratio

shown in Figure 7.10. Figure 7.11 is the same plot, zoomed in on the region containing the corrected points. The branching ratio appears to take on one value for $R < 50$ mm, and another, lower, value after, thus causing the trend in Figure 7.9.

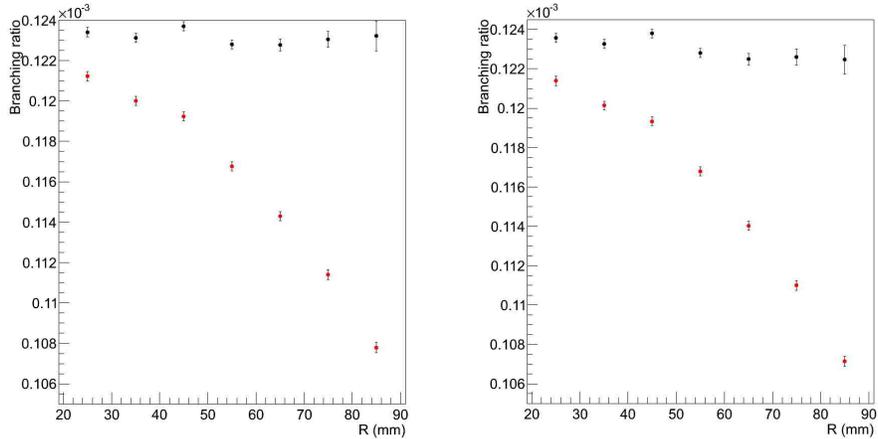


Figure 7.10: The branching ratio for statistically independent rings in R , the distance between the reconstructed positron track and the centre of WC3. The left-hand panel shows the results with the T2 pileup cut, and the right-hand panel shows the results without the T2 pileup cut. The red points show the raw branching ratio, and the black points show the branching ratio after all corrections. The value along the x axis is the centre of the ring under consideration; that is, the point at $x = 35$ mm is the branching ratio for events with R between 30 and 40 mm.

Since the downward trend in the branching ratio only becomes statistically significant after $R > 60$ mm, which is beyond the usual acceptance cut, it could perhaps be argued that the behaviour of the branching ratio in this region does not affect the systematic error on the result for $R < 60$ mm. This would be the case if, for example, the tail correction at large R were underestimated. However, if instead there were a systematic effect in the track reconstruction by which $\pi^+ \rightarrow e^+\nu_e$ events were shifted to lower radii, the branching ratio would be overestimated for small R , and underestimated for large R . Without knowing the source of the R dependence, the true branching ratio could be anywhere within the observed range, so the

7.5. Summary of Chapter 7

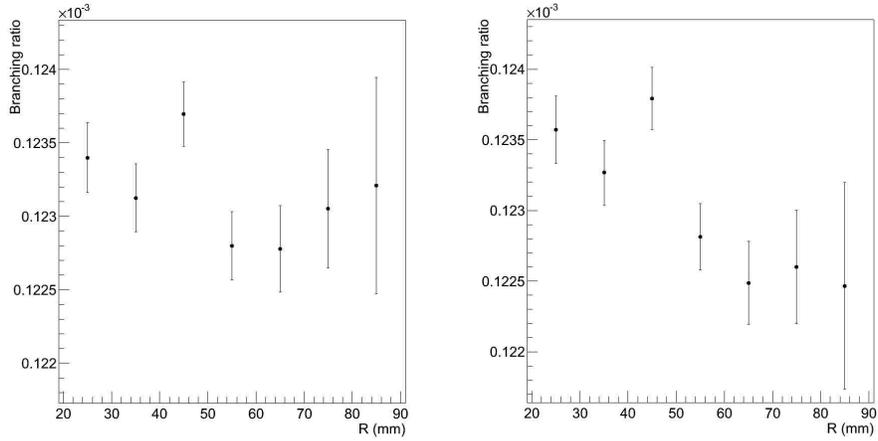


Figure 7.11: The corrected branching ratio for statistically independent rings in R, with (left) and without (right) the T2 pileup cut. This plot is identical to Figure 7.10, but zoomed in on the corrected points.

variation is included in the systematic error.

The change in the branching ratio as E_{cut} is varied is shown in Figure 7.12. Aside from the point at 50 MeV, no significant variation is observed. The signal to background ratio in the high energy spectrum decreases rapidly starting near 50 MeV, and the quality of the fit at this point is poor. In light of the stability above 50 MeV, no systematic error is assigned for the selection of E_{cut} .

7.5 Summary of Chapter 7

The systematic uncertainties on the three small corrections are all less than 0.05% (see Table 7.1). The dependence of the corrected branching ratio on the radius at WC3 in which events are accepted, however, is much larger, and contributes significantly to the total uncertainty on the branching ratio (see Figure 7.9). The size of each source of error, and the total error, are given in the next chapter.

7.5. Summary of Chapter 7

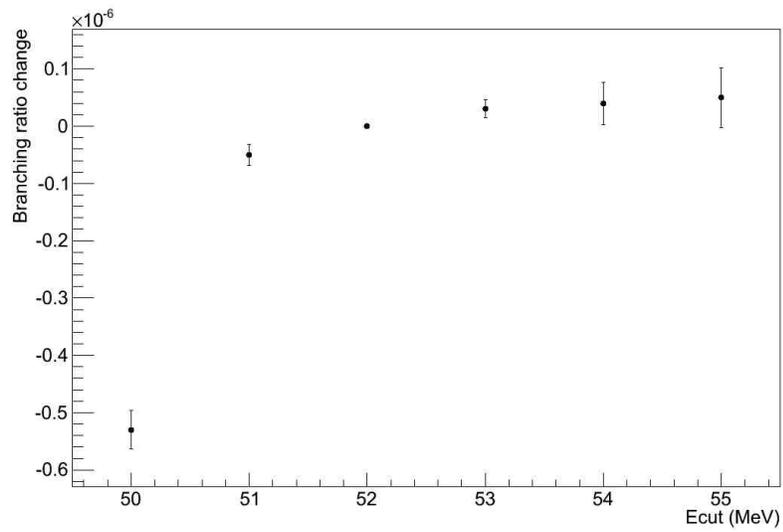


Figure 7.12: The change in the branching ratio from the value at $E_{cut} = 52$ MeV.

Chapter 8

Conclusion

8.1 Branching ratio result

The major sources of error on the branching ratio obtained from the 2012 data are as follows:

- **Statistics:** 11×10^{-8} (see Section 5.4)
- **Low-energy tail:** 13×10^{-8} (see Section 6.8)
- **R dependence:** 48×10^{-8} (see Section 7.4)
- **R dependence:** 29×10^{-8} (T2 pileup cut applied) (see Section 7.4)
- **τ_μ, τ_π dependence:** 22×10^{-8} (T2 pileup cut applied) (see Section 5.5)

Since this analysis represents a work in progress, a conservative approach to the final systematic error is taken; namely, the full variations due to the R dependence and the lifetime dependence are taken as the 1σ range of the branching ratio. For example, with the T2 pileup cut applied, the variation as a function of R is 58×10^{-8} ; thus, a systematic uncertainty of $\pm 29 \times 10^{-8}$ is assigned. The impacts of the small ($< 0.05\%$) sources of systematic uncertainty on the raw branching ratio are calculated in the same way (see Table 5.5). They do not affect the final error.

Adding the systematic uncertainties in quadrature gives

$$R_\pi = (1.2^{***} \pm 0.0011(\text{stat}) \pm 0.0040(\text{syst})) \times 10^{-4}. \quad (8.1)$$

Adding the statistical and systematic uncertainties in quadrature gives 41×10^{-8} , or 0.33% of the theoretical value. For reference, the blinded values are $R_\pi^{\text{raw}} = 1.2005 \times 10^{-4}$ and $R_\pi^{\text{corr}} = 1.2355 \times 10^{-4}$.

8.1. Branching ratio result

Table 8.1: Sources of error. The corrected branching ratio is given by the product of the raw branching ratio and the corrections. The errors given for the corrections are the errors on the corrections themselves, not the resulting errors on the branching ratio. The stars indicate that the result is still blinded.

	Values	Uncertainties	
		<i>Stat</i>	<i>Syst</i>
$R_{\pi}^{raw}(10^{-4})$ § 5.4	1.2***	0.0011	
π, μ lifetimes §5.5			0.0022
Fitting range §5.5			0.0004
Bin size §5.5			0.0003
Fixed parameters §5.5			0.0001
Additional components §5.5			0.0001
Corrections			
Low energy tail §6.8	1.0312		0.0011
t_0 §7.3	1.0006		0.0003
Acceptance §7.1	0.9991		0.0003
μ -DIF §7.2	0.9983		< 0.0001
R dependence §7.4			0.0029
$R_{\pi}^{corr}(10^{-4})$	1.2***	0.0011	0.0040

8.2 Future prospects

8.2.1 Systematic uncertainty

Clearly, the largest issue with the analysis is the R dependence. The fact that the dependence is reduced when the T2 pileup cut is applied indicates that some pileup mechanism not included in the fit could be responsible. If this is the case, identifying and characterizing the component, and including it in the fit, would remove the dependence. Alternatively, if the T2 pileup cut could be improved in such a way as to not distort the time spectrum, it could perhaps be tightened, which might also remove the dependence.

The time spectrum distortion caused by the cut is likely due to the removal of events with fake hits; this effect could be reduced or removed if the cut were made more sophisticated, by comparing the pulse height of the triggering hit to the pileup hit, for example. It is worth mentioning that the R dependence is not present if the T2 pileup cut is applied and the lifetimes are freed in the fit; however, at large R the fitted lifetimes are several standard deviations away from their accepted values.

Another possibility to reduce the influence of pileup is to evaluate the BINA energy based on the pulse height of the waveform, rather than the integrated charge in a wide gate. Although the energy resolution may suffer slightly, this will not appreciably affect the result at least as long as the resolution remains Gaussian. This was done for the 2011 data, but is currently being implemented for the 2012 data.

Once the systematics due to the R dependence and the lifetime dependence are removed or reduced to an acceptable level, the largest remaining systematic error will be due to the low energy tail; this was the largest source of systematic error in the 2010 and 2011 analyses. The error coming from the tail correction could be reduced if the influence of beamline scattering in the response function measurement could be either characterized or established as negligible; then what is now considered the upper limit on the tail correction could be taken as the true tail correction. This would reduce the error from 13×10^{-8} to 8×10^{-8} , even if no other improvements in the tail analysis were made. An effort to simulate the positron transport down the beamline, and evaluate the low momentum tail of the beam entering the detector, is ongoing, but has not yet produced conclusive results.

Ultimately, the analysis of at least the 2010, 2011, and 2012 data will be combined to obtain a single result for the branching ratio. The statistical error using the current cuts will be 8×10^{-8} . If the systematic error can be reduced to 9×10^{-8} or less, the goal of 0.1% uncertainty on the branching ratio will have been achieved. If the systematic error can be reduced below this level, the data from 2009 and the beginning of 2010 could be incorporated as well, reducing the statistical error further. The different data-taking conditions for these running periods would make this a challenge, but if the result from the easily-useable data turns out to be statistics-limited, it could be attempted.

8.2.2 Limits on new physics

If there were a difference in the couplings of the W to the electron and muon, g_e and g_μ , the branching ratio would be related to the Standard Model prediction by $R_\pi = (g_e/g_\mu)^2 R_\pi^{SM}$ (see Section 2.3). Thus, achieving a precision of 0.1% on the branching ratio would correspond to a 0.05% level test of lepton universality.

This would make pion decay once again the most sensitive test of lepton universality, and improve the already stringent constraints on models attempting to explain the hints of possible lepton nonuniversality seen by the LHCb [53] [54] and BaBar [55] experiments. Essentially, the models must include the property that the mechanism that couples differently to the different generations be greatly enhanced for the third generation [56].

The limits placed on specific processes are calculated using formulae found in Section 2.2. The unblinded result of the 2010 analysis was

$$R_\pi^{2010} = (1.2344 \pm 0.0030) \times 10^{-4}. \quad (8.2)$$

To translate this into an upper limit on the branching ratio, the Feldman-Cousins technique [67] will be used. This provides a frequentist confidence interval based on an ordering of likelihood ratios, avoiding such problems as non-physical confidence regions. Notably, it also provides consistent treatment when the data itself is used to make the decision to assign an upper limit or a two-sided interval. Consulting Table X in Reference [67], the upper limit on the branching ratio at 95% confidence is 1.67 standard deviations above the SM prediction, or

$$R_\pi^{UL} = 1.2402 \times 10^{-4}. \quad (8.3)$$

New pseudoscalar interactions

Substituting the SM prediction and the value from Equation 8.3 into Equation 2.16 gives

$$\frac{1.2402}{1.2352} - 1 = \left(\frac{1 \text{ TeV}}{\Lambda} \right)^2 \times 10^3, \quad (8.4)$$

giving

$$\Lambda = 497 \text{ TeV}. \quad (8.5)$$

Thus, the mass scale of a new fundamental pseudoscalar, with the same coupling strength to quarks and leptons as the weak interaction, must be $> 500 \text{ TeV}$ at 95% C.L. A 0.1% measurement, with the same central value, would give $\Lambda > 880 \text{ TeV}$.

R-parity violating SUSY

The relationship between R_π and the R-parity violating parameters Δ'_{11k} and Δ'_{21k} (see Section 2.2.2) is

$$\frac{\Delta R_\pi}{R_\pi^{SM}} = 2(\Delta'_{11k} - \Delta'_{21k}). \quad (8.6)$$

R_π itself does not provide any constraint on the size of Δ'_{11k} and Δ'_{21k} in the case where they are equal in value. In the case where $\Delta'_{11k} = 0$, $\Delta'_{21k} < 0.0020$ at 95% C.L.; in the case where $\Delta'_{21k} = 0$, $\Delta'_{11k} < 0.0028$ at 95% C.L.

Charged Higgs bosons

As discussed in Section 2.2.3, if the coupling of the charged Higgs boson to leptons is proportional to the lepton mass, as with the SM Higgs boson, R_π is unaffected by the presence of a charged Higgs boson. However, if the coupling is independent of the lepton mass, this is no longer the case. For a coupling of α/π , the limit at 95% C.L. is

$$M_{H^\pm} > 144 \text{ GeV}. \quad (8.7)$$

Massive neutrinos

The limits obtained for massive neutrino mixing, for neutrino mass in the range 55-130 MeV, using a search for extra peaks in the $\pi^+ \rightarrow e^+ \nu_e$ energy

8.2. Future prospects

spectrum with PIENU data taken in 2009, were shown in Figure 2.5. A similar analysis using the 2010, 2011, and 2012 data is underway, which is expected to improve the limits by up to a factor of 5. The limit below 50 MeV was obtained from R_π itself; this provides a considerably worse limit than the peak search, by more than an order of magnitude. For a 0.1% measurement of R_π , it is expected to improve by a factor of 2-3.

Bibliography

- [1] Vincenzo Cirigliano and Ignasi Rosell. $\pi/K \rightarrow e\nu$ branching ratios to $O(e^2 p^4)$ in Chiral Perturbation Theory. *JHEP*, 10:005, 2007. arXiv: 0707.4464, doi:10.1088/1126-6708/2007/10/005.
- [2] D. A. 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*, 61:331–354, 2011. doi:10.1146/annurev-nucl-102010-130431.
- [3] A. Aguilar-Arevalo et al. (PIENU Collaboration). Improved measurement of the $\pi \rightarrow e\nu$ branching ratio. *Phys. Rev. Lett.*, 115:071801, 2015. doi:10.1103/PhysRevLett.115.071801.
- [4] J. Sulkimo et al. GEANT4 a simulation toolkit. *Nucl. Inst. & Meth. in Phys. Res. A*, 506:250–303, 2003.
- [5] R. Carlini, J.M. Finn, S.Kowalski, and S. Page, spokespersons. JLab Experiment E-02-020.
- [6] D.I. Britton et al. Improved search for massive neutrinos in $\pi^+ \rightarrow e^+\nu$ decay. *Phys. Rev. D*, 46:R885–R887, Aug 1992. doi:10.1103/PhysRevD.46.R885.
- [7] M. Aoki et al. Search for massive neutrinos in the decay $\pi \rightarrow e\nu$. *Phys. Rev. D*, 84:052002, 2011. doi:10.1103/PhysRevD.84.052002.
- [8] A.A.Aguilar-Arevalo et al. Detector for measuring the $\pi^+ \rightarrow e^+\nu_e$ branching fraction. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 791:38–46, August 2015.
- [9] A.A.Aguilar-Arevalo et al. Study of a large NaI(Tl) crystal. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 621:188–191, September 2010.
- [10] Cliff Burgess and Guy Moore. *The Standard Model: A Primer*. Cambridge University Press, 2007.

- [11] C. Patrignani et al. Particle data group. *Chin. Phys. C*, 40, 2016.
- [12] Peter W. Higgs. Broken symmetries and the masses of gauge bosons. *Phys. Rev. Lett.*, 13:508–509, 1964.
- [13] G. Aad et al. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, 716:1–29, 2012. doi:{10.1016/j.physletb.2012.08.020}.
- [14] S. Chatrchyan et al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B*, 716:30–61, 2012. doi:10.1016/j.physletb.2012.08.021.
- [15] W. J. Marciano and A. Sirlin. Theorem on π_{l2} decays and electron-muon universality. *Phys. Rev. Lett.*, 36(24):1425–1428, Jun 1976. doi:10.1103/PhysRevLett.36.1425.
- [16] Michael J. Ramsey-Musolf, Shufang Su, and Sean Tulin. Pion Leptonic Decays and Supersymmetry. *Phys. Rev.*, D76:095017, 2007. arXiv:0705.0028, doi:10.1103/PhysRevD.76.095017.
- [17] Miriam Leurer. Comprehensive study of leptoquark bounds. *Phys. Rev. D*, 49:333–342, 1994.
- [18] Robert E. Shrock. General theory of weak processes involving neutrinos. ii. pure leptonic decays. *Phys. Rev. D*, 24:1275–1309, Sep 1981. doi:10.1103/PhysRevD.24.1275.
- [19] C. M. G. Lattes, H. Muirhead, G. P. S. Occhialini, C. F. Powell. Processes involving charged mesons. *Nature*, 159:694–697, 1947.
- [20] E. C. G. Sudarshan and R. E. Marshak. Chirality invariance and the universal fermi interaction. *Phys. Rev.*, 109(5):1860–1862, Mar 1958. doi:10.1103/PhysRev.109.1860.2.
- [21] R.P. Feynman and M. Gell-Mann. Theory of the fermi interaction. *Physical Review*, 109(1):193–198, 1958.
- [22] T. Fazzini et al. Electron Decay of the Pion. *Phys. Rev. Lett.*, 1:247–249, 1958.
- [23] H.L. Anderson et al. Branching Ratio of the Electronic Mode of Positive Pion Decay. *Phys. Rev.*, 119:2050–2067, 1960.

Bibliography

- [24] E. Di Capua, R. Garland, L. Pondrom, and A. Strelzoff. Study of the decay $\pi \rightarrow e + \nu$. *Phys. Rev.*, 133(5B):B1333–B1340, Mar 1964. doi:10.1103/PhysRev.133.B1333.
- [25] Douglas Bryman and Charles Picciotto. Revised value for the $\pi \rightarrow e\nu$ branching ratio. *Phys. Rev. D*, 11:1337, 1975. doi:10.1103/PhysRevD.11.1337.
- [26] D. A. Bryman, R. Dubois, T. Numao, B. Olaniyi, A. Olin, M. S. Dixit, D. Berghofer, J. -M. Poutissou, J. A. Macdonald, and B. C. Robertson. New measurement of the $\pi \rightarrow e\nu$ branching ratio. *Phys. Rev. Lett.*, 50:7–10, 1983. doi:10.1103/PhysRevLett.50.7.
- [27] D. I. Britton, S. Ahmad, D. A. Bryman, R. A. Burnham, E. T. H. Clifford, P. Kitching, Y. Kuno, J. A. Macdonald, T. Numao, A. Olin, J-M. Poutissou, and M. S. Dixit. Measurement of the $\pi^+ \rightarrow e^+\nu$ branching ratio. *Phys. Rev. D*, 49(1):28–39, Jan 1994. doi:10.1103/PhysRevD.49.28.
- [28] G. Czapek, A. Federspiel, A. Flükiger, D. Frei, B. Hahn, C. Hug, E. Hugentobler, W. Krebs, U. Moser, D. Muster, E. Ramseyer, H. Scheidiger, P. Schlatter, G. Stucki, R. Abela, D. Renker, and E. Steiner. Branching ratio for the rare pion decay into positron and neutrino. *Phys. Rev. Lett.*, 70(1):17–20, Jan 1993. doi:10.1103/PhysRevLett.70.17.
- [29] K. A. Olive et al. Particle data group. *Chin. Phys. C*, 38, 2014 and 2015 update. doi:10.1088/1674-1137/38/9/090001.
- [30] Andreas van der Schaaf, Dinko Počanić, spokespersons. PSI Experiment R-05-01.
- [31] A.A.Aguilar-Arevalo et al. High purity Pion Beam at TRIUMF. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 609:102–105, October 2009.
- [32] Micheal E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Westview Press, 1995.
- [33] James D. Bjorken and Sidney D. Drell. *Relativistic Quantum Mechanics*. Mcgraw-Hill Book Company, 1964.
- [34] S. M. Berman. Radiative corrections to pion beta decay. *Phys. Rev. Lett.*, 1(12):468–469, Dec 1958. doi:10.1103/PhysRevLett.1.468.

Bibliography

- [35] Toichiro Kinoshita. Radiative corrections to $\pi - e$ decay. *Phys. Rev. Lett.*, 2(11):477–480, Jun 1959. doi:10.1103/PhysRevLett.2.477.
- [36] William J. Marciano and A. Sirlin. Radiative corrections to π_{l2} decays. *Phys. Rev. Lett.*, 71(22):3629–3632, Nov 1993. doi:10.1103/PhysRevLett.71.3629.
- [37] G.W. Bennett et al. Measurement of the Negative Muon Anomalous Magnetic Moment to 0.7 ppm. *Phys. Rev. Lett.*, 92:161802, 2004. doi:10.1103/PhysRevLett.92.161802.
- [38] Randolph Pohl, Ronald Gilman, Gerald A. Miller, and Krzysztof Pachucki. Muonic Hydrogen and the Proton Radius Puzzle. *Annual Review of Nuclear and Particle Science*, 63:175–204, 2013. doi:10.1146/annurev-nucl-102212-170627.
- [39] Troy A. Porter, Robert P. Johnson, and Peter W. Graham. Dark matter searches with astroparticle data. *Annual Review of Astronomy and Astrophysics*, 49:1–580, 2011.
- [40] Joshua A. Frieman, Micheal S. Turner, and Dragan Huterer. Dark energy and the accelerating universe. *Annual Review of Astronomy and Astrophysics*, 46:1–572, 2008.
- [41] Edward W. Kolb and Michael S. Turner. Grand unified theories and the origin of the baryon asymmetry. *Annual Review of Nuclear and Particle Science*, 33:645–696, 1983.
- [42] Jihn E. Kim and Gianpaolo Carosi. Axions and the strong cp problem. *Rev. Mod. Phys.*, 82:557–601, Mar 2010. doi:10.1103/RevModPhys.82.557.
- [43] Jonathan L. Feng. Naturalness and the status of supersymmetry. *Annual Review of Nuclear and Particle Science*, 63:351–382, 2013.
- [44] Bruce A. Campbell and David W. Maybury. Constraints on scalar couplings from $\pi^+ \rightarrow l^+ + \nu_l$. *Nucl. Phys.*, B709:419–439, 2005. doi:10.1016/j.nuclphysb.2004.12.015.
- [45] R. Barbier et al. R-Parity-violating supersymmetry. *Physics Reports*, 420:1–195, 2005.
- [46] Ian Hinchliffe. Supersymmetric models of particle physics and their phenomenology. *Annual Review of Nuclear and Particle Science*, 36:505–543, 1986.

Bibliography

- [47] D. DeCamp et al. Determination of the number of light neutrino species. *Phys. Lett. B*, 231:519–529, 1989.
- [48] Alexander Kusenko. Sterile neutrinos: The dark side of the light fermions. *Physics Reports*, 481:1–28, 2009.
- [49] R. Adhikari et al. A White Paper on keV Sterile Neutrino Dark Matter. 2016. [arXiv:1602.04816](https://arxiv.org/abs/1602.04816).
- [50] Lusiani, Alberto. Lepton universality and lepton flavour violation tests at the b-factories. *EPJ Web of Conferences*, 118:01018, 2016. doi:10.1051/epjconf/201611801018.
- [51] C. Lazzeroni et al. Test of lepton flavour universality in $K^+ \rightarrow l^+\nu$ decays. *Physics Letters B*, 698(2):105 – 114, 2011. doi:10.1016/j.physletb.2011.02.064.
- [52] A. Pich. Tau Physics: Theory Overview. *Nucl. Phys. Proc. Suppl.*, 181-182:300–305, 2008. [arXiv:{0806.2793}](https://arxiv.org/abs/0806.2793).
- [53] The LHCb Collaboration. Test of lepton universality using $B^+ \rightarrow K^+l^+l^-$ decays. *Phys. Rev. Lett.*, 113(15):151601, Oct 2014. doi:10.1103/PhysRevLett.113.151601.
- [54] The LHCb Collaboration. Measurement of the ratio of branching fractions $\mathcal{B}(\overline{B}^0 \rightarrow D^{*+}\tau^-\overline{\nu}_\tau)/\mathcal{B}(\overline{B}^0 \rightarrow D^{*+}\mu^-\overline{\nu}_\mu)$. *Phys. Rev. Lett.*, 115(11):111803, Sep 2015. doi:10.1103/PhysRevLett.115.111803.
- [55] J.P. Lees et al (BaBar Collaboration). Evidence for an excess of $\overline{B} \rightarrow D^{(*)}\tau^-\overline{\nu}_\tau$ decays. *Phys. Rev. Lett.*, 109:101802, Sep 2012. doi:10.1103/PhysRevLett.109.101802.
- [56] Ferruccio Feruglio, Paride Paradisi, and Andrea Pattori. Revisiting lepton flavor universality in b decays. *Phys. Rev. Lett.*, 118:011801, 2017. doi:10.1103/PhysRevLett.118.011801.
- [57] J.-P. Martin and P.-A. Amaudruz. A 48 channel pulse shape digitizer with dsp. *IEEE Transactions on Nuclear Science*, 53(3):715 – 719, 2006. doi:10.1109/TNS.2006.875049.
- [58] Y. Igarashi, H. Fujii, T. Higuchi, M. Ikeno, E. Inoue, T. Murakami, Y. Nagasaka, M. Nakao, K. Nakayoshi, M. Saitoh, S. Shimazaki, S.Y. Suzuki, M. Tanaka, K. Tauchi, T. Uchida, and Y. Yasu. A common

- data acquisition system for high-intensity beam experiments. *IEEE Transactions on Nuclear Science*, 52(6):2866–2871, 2005.
- [59] C. Ohlmann, TRIUMF, Vancouver, Canada. VT48 Rev.A, 48-Channel Time-to-Digital Module. User’s Manual, 2007.
- [60] S. Ritt and P. A. Amaudruz. The midas data acquisition system. Last retrieved September 26/16. Available from: `midas.triumf.ca`.
- [61] R. Brun and F. Rademakers. ROOT - An Object Oriented Data Analysis Framework. *Nucl. Inst. & Meth. in Phys. Res. A*, 389:81–86, 1997.
- [62] Chloé Malbrunot. *Study of $\pi^+ \rightarrow e^+\nu_e$ decay*. PhD thesis, The University of British Columbia, 2012.
- [63] J. B. Birks. *The Theory and Practice of Scintillation Counting*. Pergamon, 1964.
- [64] K.J. Ma, W.G. Kang, J.K. Ahn, S. Choi, Y. Choi, M.J. Hwang, J.S. Jang, E.J. Jeon, K.K. Joo, H.S. Kim, J.Y. Kim, S.B. Kim, S.H. Kim, W. Kim, Y.D. Kim, J. Lee, I.T. Lim, Y.D. Oh, M.Y. Pac, C.W. Park, I.G. Park, K.S. Park, S.S. Stepanyan, and I. Yu. Time and amplitude of afterpulse measured with a large size photomultiplier tube. *Nucl. Inst. & Meth. in Phys. Res. A*, 629:93 – 100, 2011. doi:10.1016/j.nima.2010.11.095.
- [65] Shintaro Ito. *Measurement of the $\pi^+ \rightarrow e^+\nu_e$ Branching Ratio*. PhD thesis, Osaka University, 2016.
- [66] Dariusz Dereniowski and Marek Kubale. *Cholesky Factorization of Matrices in Parallel and Ranking of Graphs*, pages 985–992. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. doi:10.1007/978-3-540-24669-5_127.
- [67] Gary J. Feldman and Robert D. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, 1998. doi:10.1103/PhysRevD.57.3873.

Appendix A

Trigger Diagram

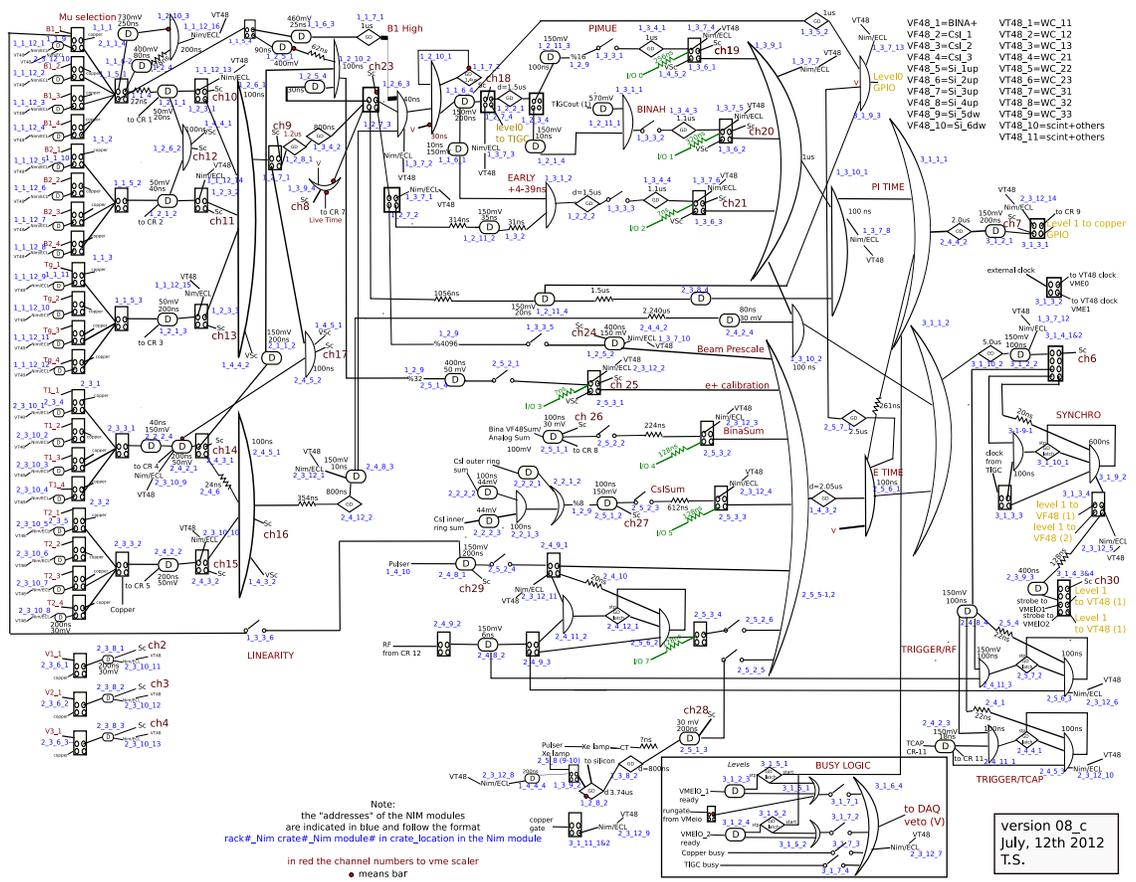


Figure A.1: PIENU Trigger Diagram

Appendix B

Timing Diagrams

This appendix contains simplified timing diagrams for various types of event, depicting the pion stop time, the upstream coincidence signal, the muon decay time, the downstream coincidence signal, and the trigger signal. The events shown are a regular PIMUE event, an event where an old muon decays and the muon from the primary pion decay does not, an event where both muons decay and the positron from the old muon decay completes the trigger, an event where both muons decay and the positron from the decay of the muon from the decay of the primary pion completes the trigger, and an event where both muons decay sufficiently close together in time that only one hit is recorded by T1. The only timing offset shown is the 300 ns delay between the upstream coincidence signal and the pion stop time.

Appendix B. Timing Diagrams

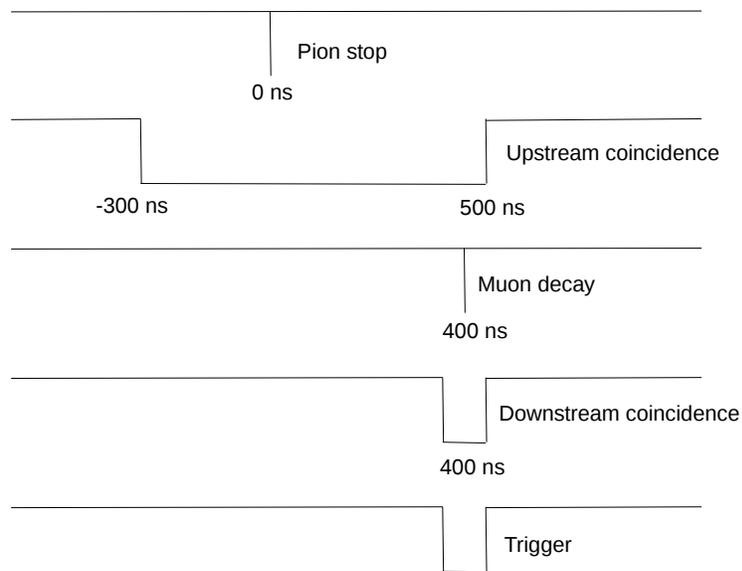


Figure B.1: A simplified timing diagram for a PIMUE event in which the muon decayed 400 ns after the pion stop.

Appendix B. Timing Diagrams

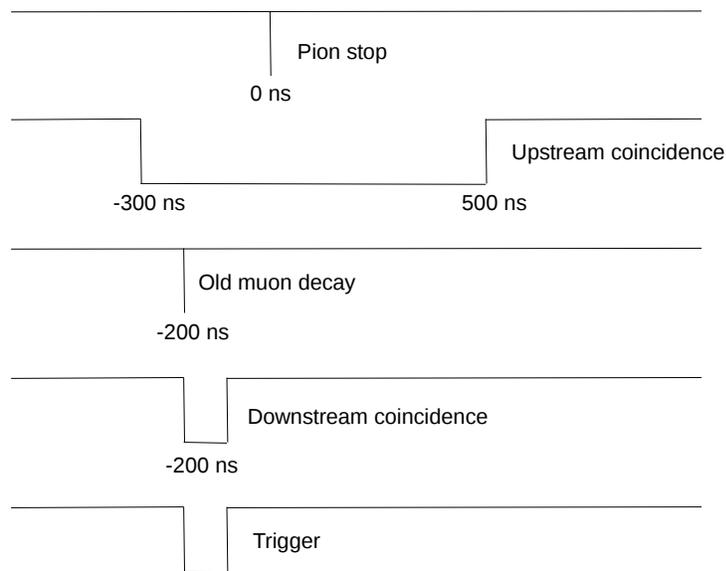


Figure B.2: A simplified timing diagram for an event in which an old muon decays and 200 ns later a pion arrives. Due to the delay between the actual hits in T1 and T2 and the downstream coincidence signal, the event still triggers.

Appendix B. Timing Diagrams

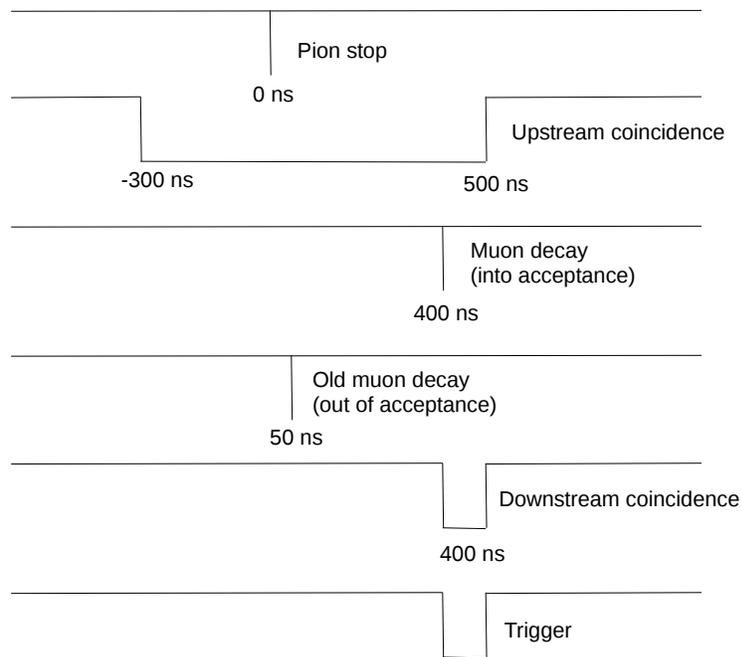


Figure B.3: A simplified timing diagram for an event with two muons, in which the positron from the old muon decay completes the trigger.

Appendix B. Timing Diagrams

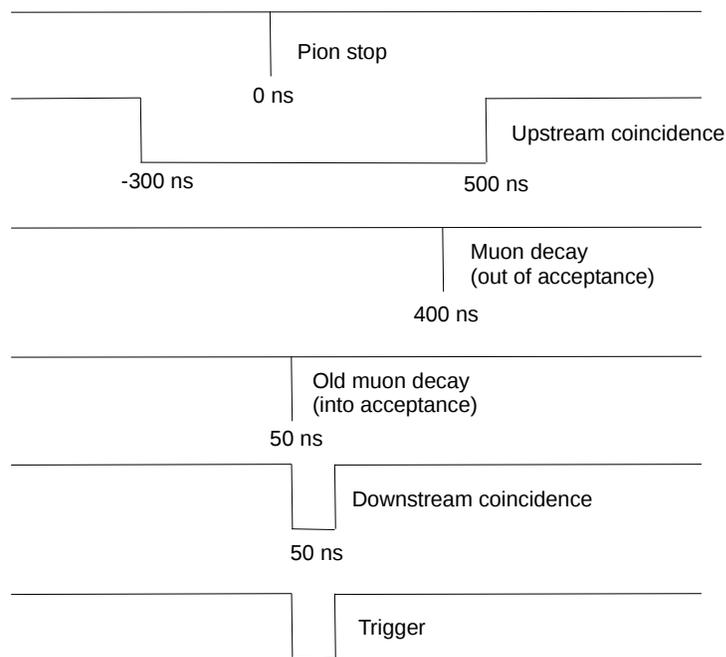


Figure B.4: A simplified timing diagram for an event with two muons, in which the positron from the muon from the decay of the primary pion completes the trigger.

Appendix B. Timing Diagrams

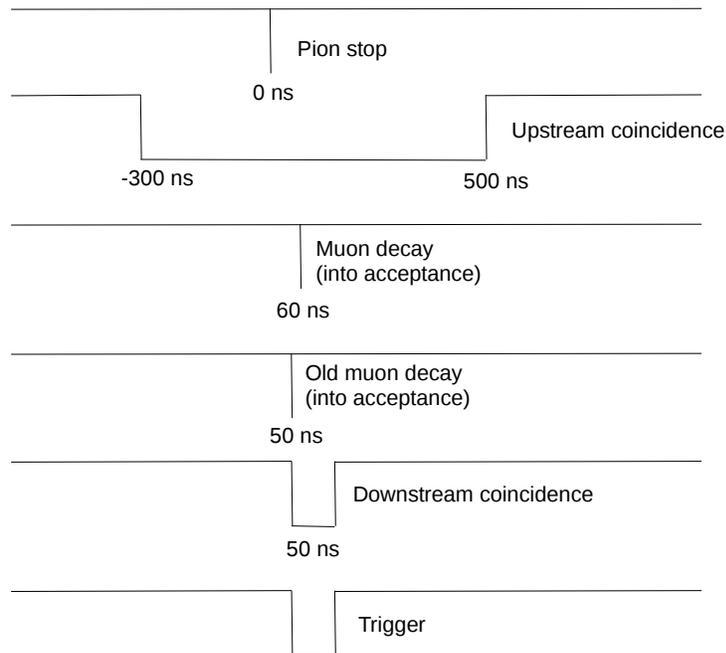


Figure B.5: A simplified timing diagram for an event with two muons, in which both decay positrons enter the acceptance, but the event passes the T1 pileup cut because the decays are too close together in time for the separate hits to be resolved.

Appendix C

Event Selection For Positron Beam Data

This Appendix contains the details of the cuts used to produce the energy spectrum used for the upper limit analysis (see Chapter 6). The purpose of the analysis was to determine the response of the crystal calorimeter array to a 70 MeV positron beam, as a function of angle. The distributions shown were taken with the beam axis aligned with the crystal axis. The reconstructed x and y distributions in the WC12 tracker are shown in Figure C.1; events outside the beam spot were removed. The energy spectrum following this cut is shown in Figure C.2.

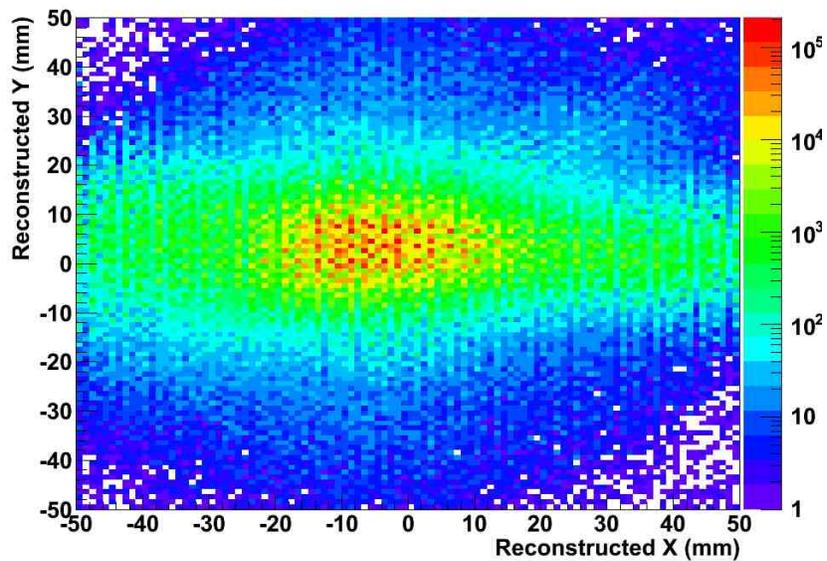


Figure C.1: The beam spot in WC1 and WC2 for positron beam data.

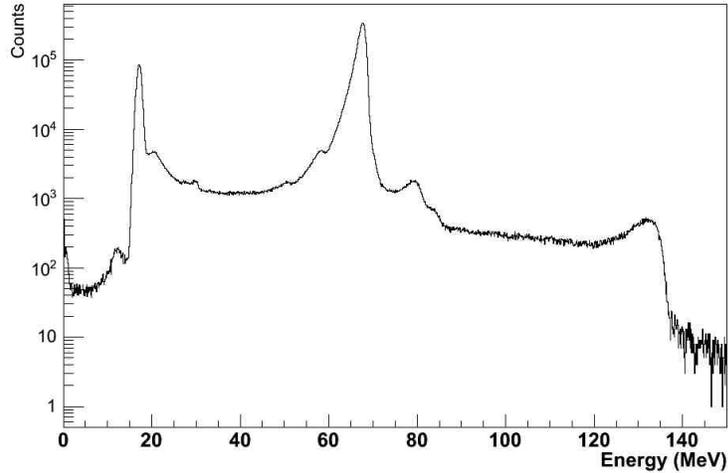


Figure C.2: Energy measured by the crystals following the selection of beam particles.

The timing distributions in the first plane of WC1, WC2, and WC3 is shown in Figure C.3. Event with times outside the peak in any of the distributions were removed. The energy spectrum following this cut is shown in Figure C.4.

The two main features remaining in the energy spectrum are the peaks due to beam positrons and beam muons. A time-of-flight cut was used to mitigate the muon background. A plot of the time of the hit in T2 relative to the RF versus the energy deposited in T2 is shown in Figure C.5. The y-axis is the time difference between the peak of the 23 MHz cyclotron RF and the hit in T2; because the particles travelled the entire length of the beamline before reaching the detector, the three particle species separated in time. The x-axis is the integral of the T2 waveform in a region around the peak, which is proportional to the energy deposited. The two dark blobs are positrons and muons; the RF region vetoed in the trigger can be seen. The energy spectra in the region from 3-12 ns and 12-16 ns, corresponding to positrons and muons, respectively, are shown in Figure C.6.

Appendix C. Event Selection For Positron Beam Data

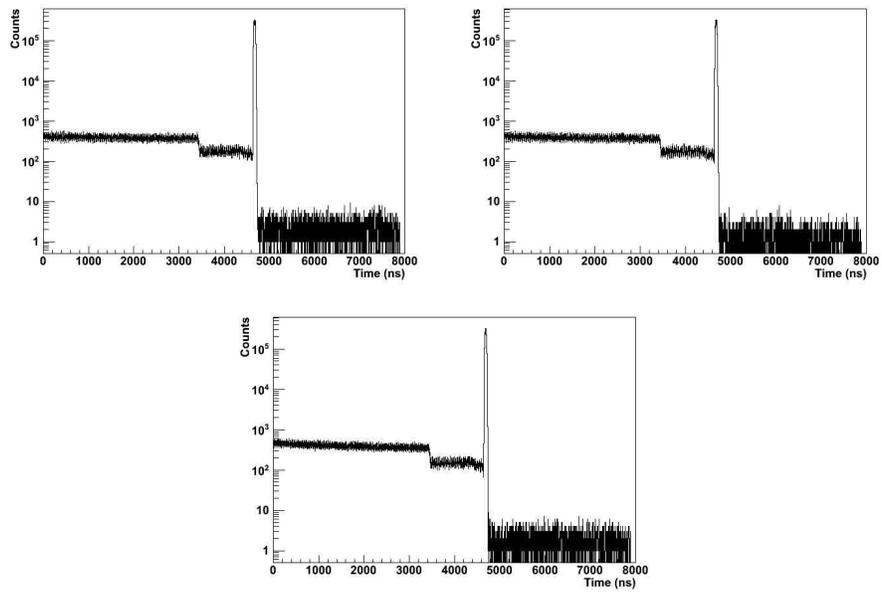


Figure C.3: The time distribution of the first hit in the first plane of each wire chamber.

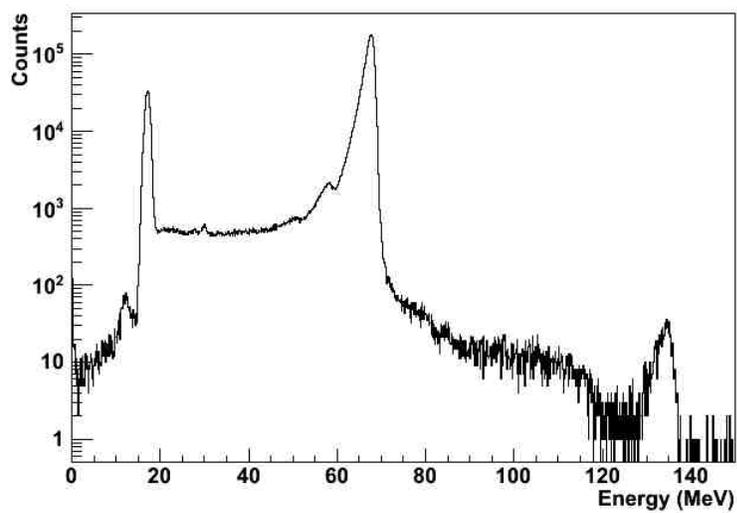


Figure C.4: Energy measured by the crystals at 0° following the removal of events with out-of-time hits.

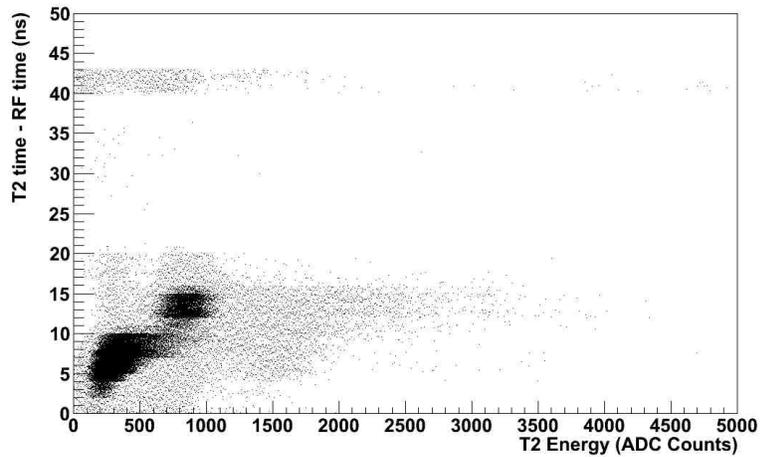


Figure C.5: Time of flight vs. T2 energy.

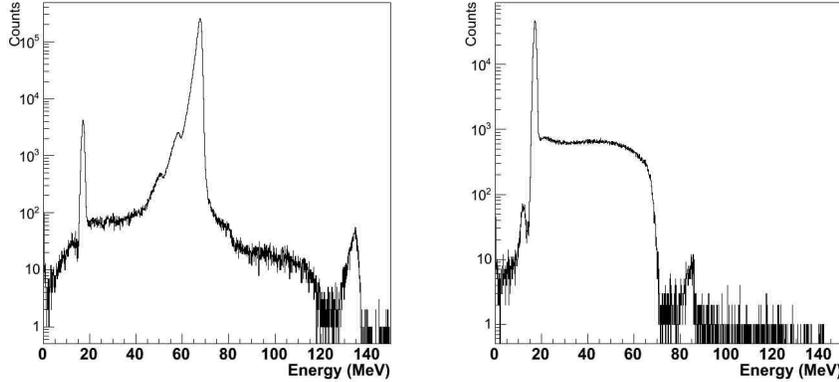


Figure C.6: Energy measured by the crystals at 0° for events with time-of-flight corresponding to positrons (left) and muons (right).

Although selecting the RF window corresponding to positrons substantially reduced the muon background, it did not eliminate it entirely. This could be accomplished with a cut in the T2 energy, as shown in Figure C.7. In this plot, any event with a total T2 energy greater than 400 ADC counts was rejected (see Figure C.5).

This cut is effective in removing muons; no trace of the peak around 18 MeV remains. However, inspection of Figure C.5 shows that positrons which deposited an unusually large amount of energy in T2 were removed as well. This could bias the response function measurement, since it alters the energy distribution of the positrons entering BINA. Furthermore, shower leakage backwards out of BINA affected the T2 energy spectrum, meaning a cut in the T2 energy would preferentially remove events with shower leakage, worsening the potential for bias. Figure C.8 shows simulated T2 energy spectra with and without BINA in place. Ultimately, no cut was done on the energy deposited in T2.

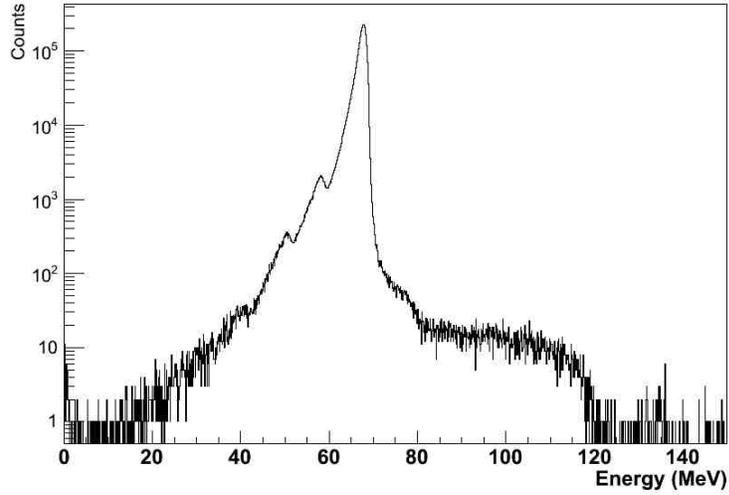


Figure C.7: Energy measured by the crystals at 0° following the removal of events with high energy deposit in T2.

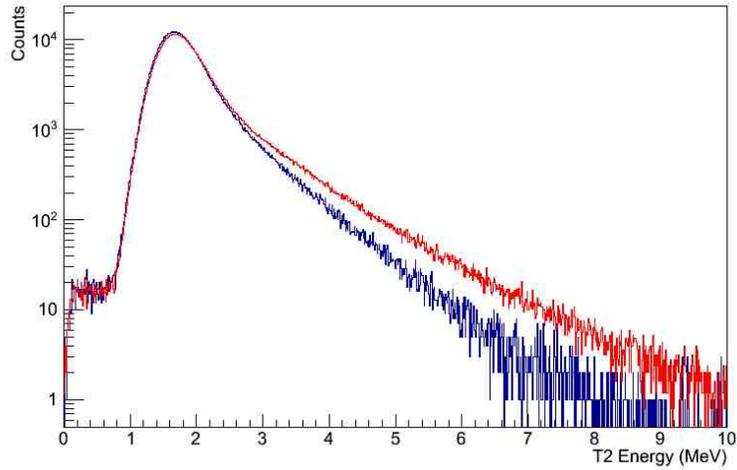


Figure C.8: Simulated energy deposit in T2, with (red) and without (black) BINA in place.

Appendix D

Positron Data Systematics

Table D.1: The change in the tail fraction as beam parameters and detector geometry were varied. The values given are the nominal tail fraction minus the new tail fraction (see Section 6.6 for a detailed description of what was changed). Note that the change is given as a fraction of the total spectrum, not the tail. The upper part shows the results for variations that increased the tail, and the lower part shows the results for variations that decreased the tail. The errors are due to Monte Carlo statistics.

	Crystal-beam angle	Centre of rotation	Beam momentum	Beam divergence
0°	$(-5.2 \pm 3.1) \times 10^{-5}$	$(-3.8 \pm 4.1) \times 10^{-5}$	$(-11 \pm 4) \times 10^{-5}$	$(-3.1 \pm 4.1) \times 10^{-5}$
6°	$(-0.7 \pm 3.2) \times 10^{-5}$	$(-2.2 \pm 4.1) \times 10^{-5}$	$(-13 \pm 4) \times 10^{-5}$	$(5.8 \pm 4.1) \times 10^{-5}$
11.8°	$(-1.7 \pm 3.3) \times 10^{-5}$	$(-7.4 \pm 4.2) \times 10^{-5}$	$(-13 \pm 4.0) \times 10^{-5}$	$(-3.6 \pm 4.2) \times 10^{-5}$
16.5°	$(3.4 \pm 3.4) \times 10^{-5}$	$(7.9 \pm 4.3) \times 10^{-5}$	$(-2.0 \pm 4.4) \times 10^{-5}$	$(1.5 \pm 4.4) \times 10^{-5}$
20.9°	$(-12 \pm 4) \times 10^{-5}$	$(-12 \pm 5) \times 10^{-5}$	$(-12 \pm 4) \times 10^{-5}$	$(-12 \pm 5) \times 10^{-5}$
24.4°	$(-4.5 \pm 3.9) \times 10^{-5}$	$(-12 \pm 5) \times 10^{-5}$	$(-18 \pm 5.0) \times 10^{-5}$	$(-11 \pm 5) \times 10^{-5}$
30.8°	$(-11 \pm 5) \times 10^{-5}$	$(-21 \pm 6) \times 10^{-5}$	$(-13 \pm 6) \times 10^{-5}$	$(-19 \pm 6) \times 10^{-5}$
36.2°	$(-3.9 \pm 6.4) \times 10^{-5}$	$(-3.3 \pm 6.5) \times 10^{-5}$	$(-11 \pm 6) \times 10^{-5}$	$(-3.9 \pm 6.4) \times 10^{-5}$
41.6°	$(-8.1 \pm 7.7) \times 10^{-5}$	$(-12 \pm 8) \times 10^{-5}$	$(-27 \pm 8) \times 10^{-5}$	$(-4.8 \pm 7.7) \times 10^{-5}$
47.7°	$(-36 \pm 10) \times 10^{-5}$	$(-37 \pm 10) \times 10^{-5}$	$(-49 \pm 10) \times 10^{-5}$	$(-37 \pm 10) \times 10^{-5}$
0°	$(3.1 \pm 3.1) \times 10^{-5}$	$(2.9 \pm 4.0) \times 10^{-5}$	$(7.2 \pm 3.8) \times 10^{-5}$	$(-0.6 \pm 4.1) \times 10^{-5}$
6°	$(1.3 \pm 3.1) \times 10^{-5}$	$(-1.5 \pm 4.1) \times 10^{-5}$	$(13 \pm 4) \times 10^{-5}$	$(-3.8 \pm 4.1) \times 10^{-5}$
11.8°	$(0.1 \pm 3.3) \times 10^{-5}$	$(-2.4 \pm 4.2) \times 10^{-5}$	$(6.3 \pm 4.0) \times 10^{-5}$	$(2.1 \pm 4.2) \times 10^{-5}$
16.5°	$(5.8 \pm 3.4) \times 10^{-5}$	$(10 \pm 4) \times 10^{-5}$	$(13 \pm 4) \times 10^{-5}$	$(-1.0 \pm 4.4) \times 10^{-5}$
20.9°	$(-4.0 \pm 3.6) \times 10^{-5}$	$(-8.1 \pm 4.6) \times 10^{-5}$	$(4.9 \pm 4.4) \times 10^{-5}$	$(3.9 \pm 4.6) \times 10^{-5}$
24.4°	$(-5.7 \pm 3.9) \times 10^{-5}$	$(1.5 \pm 4.9) \times 10^{-5}$	$(10 \pm 5) \times 10^{-5}$	$(1.5 \pm 4.9) \times 10^{-5}$
30.8°	$(-1.5 \pm 4.5) \times 10^{-5}$	$(-1.8 \pm 5.7) \times 10^{-5}$	$(10 \pm 6) \times 10^{-5}$	$(10 \pm 6) \times 10^{-5}$
36.2°	$(6.1 \pm 6.4) \times 10^{-5}$	$(6.0 \pm 6.5) \times 10^{-5}$	$(22 \pm 6) \times 10^{-5}$	$(7.2 \pm 6.4) \times 10^{-5}$
41.6°	$(17 \pm 6) \times 10^{-5}$	$(20 \pm 8) \times 10^{-5}$	$(12 \pm 8) \times 10^{-5}$	$(4.4 \pm 7.7) \times 10^{-5}$
47.7°	$(14 \pm 10) \times 10^{-5}$	$(44 \pm 10) \times 10^{-5}$	$(56 \pm 10) \times 10^{-5}$	$(31 \pm 10) \times 10^{-5}$