INFORMATION TO USERS

The most advanced technology has been used to photograph and reproduce this manuscript from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book. These are also available as one exposure on a standard 35mm slide or as a 17" x 23" black and white photographic print for an additional charge.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

U·M·I

University Microfilms International
A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
313/761-4700 800/521-0600

Order Number 8913680

**Results of a neutrino oscillation experiment performed at a
meson factory beam-stop**

Mitchell, Joseph William, Ph.D.

The Ohio State University, 1989

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106

Results of a Neutrino Oscillation Experiment
Performed at a Meson Factory Beam-Stop

A Dissertation

Presented in Partial Fulfillment of the Requirements for
the Degree of Doctor of Philosophy in the Graduate
School of the Ohio State University

By

Joseph William Mitchell, B. S.

*****

The Ohio State University
1989

Dissertation Committee:
T. A. Romanowski
W. F. Palmer
L. S. Durkin
R. D. Kass

Approved by

Adviser
Department of Physics

Acknowledgements

As this document represents the culmination of many years of effort, any list of acknowledgements is certain to be incomplete. If the reader is one whom I have dealt with over the course of my graduate studies and research, and in addition is one who has done me no severe harm during this period, please allow me to hereby express my appreciation. Certain of my friends and colleagues merit special attention. I particularly wish to thank my adviser, Dr. Thomas A. Romanowski, for bringing me onto a fine physics experiment, for providing direction throughout its course, and for his astute criticism of this dissertation. I would also like to thank Dr. T. Y. Ling for his guiding influence and insight, Dr. Stanley Durkin for sharing with me his impressive knowledge of analysis techniques and statistics, and Dr. William Metcalf of Louisiana State for his help in learning the EGS Monte Carlo. Others who have given me invaluable knowledge include Charles Rush (the designer of the E645 data acquisition system), from whom I learned electronics, and John Heimaster (who assembled the core of the E645 data acquisition software), from whom I learned the finer points of computing, as well as (unfortunately) computer repair. I would like to acknowledge the help that I was given by Chris Francis, Joey Donahue and Vern Sandberg in generating or locating a number of the figures used in this document. I would also like to express appreciation to my collaborators, in particular my friends Dr. Elton Smith, Dr. Ron Harper, Dr. Ali Fazely, and Dr. Cathy Choi for making this a bearable and often fun experience. I thank Bill Marterer and Matt Anderson for providing considerable technical assistance as well as companionship and much-needed diversion. And finally, I would like to thank all of the people who have given me support and encouragement, often over long distances, during the preparation of this thesis and the research leading up to it. This includes all of my family, and my friends Joel Miller and Wendy Bachuber, David and Julie Dennison, Michael Dempsey, and in particular my dear friend Diane Conklin, who helped keep me (mostly) sane during the final phases of this work.

Vitae

July 10, 1959	Born – Cleveland, Ohio
June 1977	High School Graduate Boardman High School Boardman, Ohio
June 1979 – September 1981	Undergraduate Research Assistant The Ohio State University Physics Department Columbus, Ohio
March 1981	B.S., Physics The Ohio State University Columbus, Ohio
September 1981 – June 1983	Graduate Teaching Assistant The Ohio State University Physics Department
June 1983 – Present	Graduate Research Assistant The Ohio State University Physics Department

Publications

L. S. Durkin, et. al., *Limits on $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations*, Physical Review Letters, **61**, 1988

Fields of Study

Major Field: Physics

Studies in experimental particle physics – T. A. Romanowski, adviser

Contents

I	Introduction to Neutrino Oscillations	1
1.1	The Experiment E645	1
1.2	History of the Neutrino and Neutrino Oscillations	2
1.2.1	The Neutrino	2
1.2.2	Neutrinos in the Standard Model	7
1.2.3	The Neutrino Oscillation Hypothesis	9
1.3	The Physics of Neutrino Oscillation	13
1.3.1	Dirac Neutrino Theory	14
1.3.2	Majorana Neutrino Theory	17
1.4	Neutrino Oscillation Experiments	20
1.4.1	Experimental Limits on Oscillation: 1988	22
1.5	Physical Implications of Neutrino Oscillation	25
1.5.1	Neutrino Mass — Physical and Astrophysical Implications	25
1.5.2	The Solar Neutrino Question	27
1.5.3	Neutrino Oscillations and Grand Unification	29
II	Design of the Experiment	33
2.1	Searching for Oscillations at a Pi-Meson Factory	33
2.1.1	Neutrino Production at LAMPF	35
2.1.2	Neutrino Physics at a Meson Factory: Advantages and Disadvantages	42
2.2	Detector Requirements	43
2.2.1	Cosmic Ray Rejection	44

2.2.2	Particle Identification & Calorimetry	45
2.2.3	Access to Past and Future Event Information	46
2.3	Components of the Experiment	47
2.3.1	The Central Detector	47
2.3.2	The Active Shield	65
2.3.3	The Physical Site and Shielding	70
2.3.4	Electronics	78
2.3.5	The Data Acquisition System	90
III	Data Reduction	96
3.1	Data Repacking	97
3.2	Calibration	99
3.2.1	Dedicated Calibration Runs	101
3.2.2	Run-to-Run Adjustment of PDT Gains	103
3.3	Structure of the Data Reduction Package	104
3.3.1	First Pass	105
3.3.2	Second Pass	106
3.3.3	Third Pass	107
3.4	Consistency Checks	108
3.5	Track Reconstruction	110
3.6	Stopping Muon Identification	112
3.6.1	First Pass Removal of Stopping Muons	114
3.6.2	Second Pass Stopping Muon Cut	115
3.7	Fiducial Volume Requirements	117
3.8	Particle ID and Energy Cuts	120
3.8.1	Electron Confidence Level	121
3.8.2	Energy cuts	128
3.9	Shield Cuts	134
3.10	Decay-In-Flight Neutrino Candidates	138
3.11	The Final Event Sample	143
3.11.1	Energy Distribution	143
3.11.2	Position Distribution	144
3.11.3	Angular Distribution	146

IV	Neutrino Oscillation Results	152
4.1	Expected Beam-Related Backgrounds and Their Properties	153
4.1.1	The EGS Monte Carlo	153
4.1.2	Backgrounds from Neutrino-Nucleus Scattering . . .	156
4.1.3	$\nu - e$ Elastic Scattering	162
4.1.4	Decay-in-Flight Neutrinos	165
4.1.5	Energy Resolution	167
4.2	Rates and Normalization	171
4.2.1	Detector Composition	171
4.2.2	$\bar{\nu}_e p$ Cross Section	171
4.2.3	Neutrino Flux	174
4.2.4	Trigger and Cut Efficiencies	175
4.2.5	Expected Background Rates and the Observed Excess	176
4.2.6	Systematic Errors	180
4.3	Neutrino Oscillation Limits	182
4.3.1	The Maximum-Likelihood Method	183
4.3.2	Likelihood vs. χ^2	189
4.3.3	Dependence of Answer on Systematic Errors	192
4.3.4	Majorana Oscillation Limits	192
4.4	Conclusion	196
4.4.1	Comparison to Previous Measurements	196
4.4.2	The Future of E645	196
4.4.3	Future Neutrino Oscillation Experiments	198
4.4.4	Summary	201
A	The E645 Triggering System	203
A.1	Introduction	203
A.2	The Fast Logic	204
A.3	The Trigger Processor	208
B	Liquid Scintillator and Crazing	215
B.1	The Physics of Crazing	217
B.2	E645 Crazing Tests	219
B.2.1	Extrusion Stress Testing	220

B.2.2 Annealing	221
B.3 Solution of the Crazing Problem	223

List of Tables

1.1	The known families of quarks and leptons. The tau neutrino and top quark have yet to be observed. Brackets show which vector bosons mediate interactions between which particles.	10
2.1	Expected rates for various muonic backgrounds in the E645 detector	69
2.2	Estimate of neutron attenuation from layers of material between the beam-stop and the detector	76
3.1	Data Flow Through the First Pass Analysis	106
3.2	Data Flow Through the Second Pass Analysis	107
3.3	Data Flow Through the Third Pass Analysis	109
3.4	Summary of acceptance and efficiency of all cuts.	142
4.1	Total neutrino cross sections for E645 signal and background processes integrated over neutrino flux and angle.	162
4.2	Cross section for neutrino-electron scattering for ν_e , ν_{μ} , and $\bar{\nu}_{\mu}$, integrated over angle and flux.	163
4.3	The number of DIF events without an observable muon in the E645 data sample	167
4.4	Materials comprising the E645 detector, their empirical formulas, and their densities.	172
4.5	Number of atoms and electrons in the E645 detector	172
4.6	Compilation of efficiencies for all cuts performed by the data analysis and how these are obtained.	175

4.7	Trigger efficiency, tracking efficiency and fiducial volume predicted by the EGS Monte Carlo for backgrounds and for inverse beta decay.	176
4.8	Observed beam excess and predicted background rates for the 1987 data set	177
4.9	Data & Monte Carlo in 5 MeV Bins	180
4.10	Estimated systematic error in the neutrino cross sections used.	181
4.11	Estimated uncertainty of the efficiencies of the cuts performed in the data analysis	181
A.1	Trigger processor CAMAC commands and their functions. .	212
A.2	Trigger processor mask bit functions.	213
B.1	Physical constants of the liquid scintillator Bicron 517L. . .	216
B.2	Crazing threshold K_m as a function of temperature for PMMA in <i>n</i> -butyl alcohol. Di Benedetto et. al.	219

List of Figures

1.1	Best limits on neutrino oscillations as of 1988	23
2.1	Distances versus neutrino energy for appearance-type oscillation experiments. The dotted lines represent $L/E = \text{constant}$	36
2.2	The Los Alamos Meson Physics Facility	37
2.3	Area A at LAMPF	39
2.4	The A6 Beam-stop. The upper figure shows the A6 target cell. The lower figure shows the isotope production stringers and the beam stop in more detail.	40
2.5	Spectra of $\nu_\mu, \bar{\nu}_\mu, \nu_e$ from the π^+ decay chain.	41
2.6	The E645 central detector and anti-coincidence shield.	48
2.7	Scintillator and PDT planes, and their mountings.	50
2.8	Schematic of the acrylic extrusions used as containers for liquid scintillator	52
2.9	Circuit diagram for the base used with the Hamamatsu R878 photomultiplier tubes.	55
2.10	Light attenuation curve for a typical extrusion	57
2.11	Landau distribution for a typical scintillation counter	58
2.12	Schematic of the E645 proportional drift tube	60
2.13	Distance versus pulse height for the E645 paper PDTs and an aluminum PDT of the same size.	62
2.14	Landau distribution for a typical PDT	64
2.15	Neutron Rate as a Function of Time, Before and After the ^{252}Cf Tagged-Source Trigger	66
2.16	The E645 Active Cosmic Ray Shield	67

2.17	Relative Cosmic Ray Flux as a Function of Overburden . .	73
2.18	Schematic of the E645 Cosmic Ray Shielding	74
2.19	Electron - Proton Separation from the 1984 Beam Test Data	75
2.20	The beam-stop and its position with respect to the detector, showing the neutron beam pipes which were filled with concrete for the 1987 data run.	77
2.21	Block Diagram of the Digitizer Circuit	82
2.22	Block Diagram of the Digitizer Crate Controller	84
2.23	Trigger Processor Schematic	87
2.24	Block Diagram of the Data Acquisition System	92
3.1	Energy spectra of stopping-muon electrons and inverse beta decay electrons, as determined by Monte Carlo.	113
3.2	Time difference between the muon track and the electron track for a sample of stopping muons selected by the first pass.	116
3.3	End point distributions after the first pass, in X, Y, and Z. The cuts are placed at the following distances from the edge X: ± 15 cm., Y: ± 20 cm., Z: $+30$ cm., -20 cm.	119
3.4	dE/dx distributions for stopping-muon electrons in scintillators and PDTs. The X axis is a dE/dx in units of "minimum ionizing particles". The distributions have been normalized so that the peaks fall at 1 MIP.	122
3.5	Confidence level distribution for stopping-muon electrons. The X axis shows scintillator confidence level. The Y axis shows PDT confidence level.	124
3.6	Confidence level distribution for the data set after stopping muon cuts.	125
3.7	Energy spectrum of events with PDT and scintillator confidence levels $< 5\%$	127
3.8	Scintillator confidence level for final sample with (left) and without (right) the removal of events with large associated energy.	129
3.9	Energy deposited in end scintillation counters for the data (after particle ID) and for Michel electrons.	131

3.10	Energy deposited off of the track versus energy deposited on the track for a Michel electron control sample.	132
3.11	Energy deposited off of the track versus energy deposited on the track for the data before the associated energy cut. . .	133
3.12	Energy spectrum of the final sample, before the cut on total energy.	135
3.13	Time difference between trigger time and the time of shield hits which meet the veto condition (pulse sum > 35) for data and random trigger.	139
3.14	The maximum summed pulse height in the shield within 30 μ sec of the trigger, for data and random trigger.	140
3.15	Energy distribution of the final neutrino-induced event candidates: beam on, beam off, and subtracted.	145
3.16	Z (Tunnel Axis) Position distribution of the final neutrino-induced event candidates: beam-on, beam-off, and subtracted	147
3.17	X (South-North) Position distribution of the final neutrino-induced event candidates: beam-on, beam-off, and subtracted	148
3.18	Y (Vertical) Position distribution of the final neutrino-induced event candidates: beam-on, beam-off, and subtracted	149
3.19	Angular distribution of the final event sample with respect to the detector Z axis.	150
3.20	Angular distribution of the final neutrino-induced event sample with respect to the beam-stop: beam-on, beam-off, and subtracted	151
4.1	Differential cross sections for expected E645 neutrino backgrounds, integrated over the beam-stop neutrino spectra. .	159
4.2	a). The Feynman diagrams for a) $\nu_e e$ scattering b) $\nu_\mu e$ scattering c) $\bar{\nu}_\mu e$ scattering.	164
4.3	Energy spectrum for Michel electrons (data points) from stopping muons compared to the spectrum predicted by the EGS Monte Carlo (solid line).	169
4.4	EGS-predicted observed energy spectra for nine different electron energies.	170

4.5	Energy spectra of the predicted E645 backgrounds, and of a maximal oscillation signal	178
4.6	Energy spectra of the combined backgrounds (solid line), the observed beam excess for the 1987 data run (points), and a 5% oscillation signal (dashed line).	179
4.7	90% Confidence level limits on $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations set by E645.	190
4.8	Expected signal from Dirac and Majorana Oscillation assuming 5% oscillation probability, compared with expected E645 background.	194
4.9	90% confidence level exclusion curve for Majorana oscillatons using Particle Data Group renormalization.	195
4.10	90% Confidence level limits on $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations set by E645 and other experiments.	197
A.1	The fast logic amplification and discrimination circuitry for a single channel.	206
A.2	Discriminator threshold versus threshold control voltage for one fast logic channel.	207
B.1	Time to fracture for annealed and unannealed $1\frac{1}{2}$ inch extrusion sections in a bath of Bicron 517L and 517H. Diamond: 517H unanld.; x: 517H anld.; Box: 517L unanld.; Star:517L annld.	222
B.2	Distortion of a $1\frac{1}{2}$ " section of extrusion as a function of the force applied 3" from one end. Diamonds: Unannealed; Squares: Annealed.	224

Chapter I

Introduction to Neutrino Oscillations

1.1 The Experiment E645

This dissertation is intended to detail the design, operation, and analysis of experiment E645, the neutrino oscillation experiment performed at the beam-stop of the Clinton P. Anderson Meson Physics Facility (LAMPF), located at the Los Alamos National Laboratory. This experiment was conducted by a collaboration of institutions including the Ohio State University, Louisiana State University, Argonne National Laboratory, the California Institute of Technology, Los Alamos National Laboratory, and Lawrence Berkeley Laboratory. It was first proposed in 1981. During the next couple of years the elements of the detector were designed and prototyped. Full-scale construction began in 1984, and was completed in time for part of the 1986 beam cycle. The first data revealed a severe background problem due to neutrons from the accelerator. This was remedied by adding extra

shielding around our detector, which eliminated the problem in the 1987 data run. This dissertation uses only the 1987 data. Another data set is being accumulated during the 1988 cycle.

It is important to put this work into historical perspective, and the introductory section of this thesis will attempt to do this. The motivation behind this project results from the culmination of a long study of the neutrino and its interactions, and the desire to probe the standard model of electroweak interactions in regions where its predictions (specifically concerning neutrino masses and lepton family mixing) are not specified. The development of our understanding of the neutrino, the postulation that one type of neutrino can change or “oscillate” into another neutrino type, the corresponding astrophysical and particle physics implications, and the previous searches for this phenomenon will be explained in some detail throughout the rest of this chapter. The rest of this dissertation will be devoted to explaining the workings of this experiment and the results of our search for neutrino oscillation.

1.2 History of the Neutrino and Neutrino Oscillations

1.2.1 The Neutrino

Neutrino history has its genesis in nuclear beta decay[1]. As energy measurements became possible in nuclear physics, the energy spectra of the

products from alpha, beta and gamma decays were determined. Alpha and gamma decays seemed to give off rays of definite energies, while the beta ray spectrum was a distribution that stretched from 0 MeV up to a cut-off. A shocking result of this was that energy, momentum and angular momentum seemed not to be conserved in these reactions. To circumvent this problem, Pauli in 1930 postulated a new particle, without charge and massless, which when emitted in beta decay conserved the necessary kinematic quantities. In 1933, Fermi incorporated this idea into his theory of nuclear beta decay [2]. The detection of this particle was far beyond the technology available at the time, since its interactions with matter were so weak and there was no source copious enough to produce a detectable number of them.

The discovery of the neutrino had to await the advent of the nuclear reactor. In reactors, fission fragments go through several beta decays as they are transformed into longer lived isotopes. Hence they are abundant sources of neutrinos. With such a source, experiments to detect neutrinos directly, through the inverse beta decay process,

$$\bar{\nu} + p \rightarrow e^+ + n,$$

became possible. This was first accomplished by Reines and Cowen [3] in 1959. Their detector consisted of tanks of liquid scintillator separated by tanks of a CdCl₂ solution. The scintillator's purpose was first to detect the positron from the inverse beta decay. The neutron would thermalize in the water and then capture on the cadmium. The cadmium then would

give off gamma rays, which were detected by the liquid scintillator. The “signature” of the event was two pulses in the scintillator separated by several microseconds, which is the thermalization time of the neutron.

Even before this, though, there were attempts to identify neutrino interactions using nuclear reactors. The most famous of these was Davis’s [6] experiment which searched for the reaction

The presence of this reaction would have indicated that neutrinos and antineutrinos were identical particles. His technique was to place a vat of CCl_4 (dry-cleaning fluid) near an operating reactor, then to flush the liquid with helium to carry off any ${}^{37}\text{Ar}$ that may have been formed. No excess of argon was seen above expected background. This was the beginning of the principle known as conservation of lepton number, wherein neutrinos and electrons were given a quantum number of +1, antineutrinos and positrons were given a quantum number of -1. All known reactions conserve this quantum number.

The uniqueness of the neutrino and antineutrino was a concept that is almost as old as the neutrino itself. In 1936 Majorana postulated that the neutrino was a two component spinor, rather than a four component spinor, and that the antineutrino and neutrino were the same state [5]. While Davis laid this to rest, the Majorana neutrino theory still is used to connote a coupling between the particle and anti-particle states of the neutrino.

Such a coupling would, as shown in the next section, allow neutrino — antineutrino oscillations to occur. It would also allow the phenomenon of neutrinoless double beta decay in certain isotopes [7].

Another complication that arose in the understanding of the neutrino was the discovery that more than one “type” of neutrino exists, one type being associated with the muon and one type associated with the electron. This became evident in the early sixties with the introduction of high energy neutrino beams. Observation of a new neutrino was first accomplished by a Columbia – BNL collaboration in 1962 [8]. One of the conclusions reached by these experimenters is that one type of neutrino could create only electrons through its interactions, and the other could create only muons. These became known as the ν_e and the ν_μ . Each has a corresponding anti-particle, as mentioned above.

A major advantage of high energy neutrinos as a probe of the weak interaction is their cross-section increases as energy squared. Coupled with the fact that their reaction products can be identified and have their energies measured makes them almost ideal. Hence it was via neutrino interactions that weak neutral currents were first observed in 1973. This was done in the Gargamelle bubble chamber at CERN [9]. Reactions of the type

$$\nu_\mu + N \rightarrow \nu_\mu + X$$

were searched for, and seen. This observation was a step forward for the Weinberg-Salam model of electroweak interactions, and was the first

evidence of the neutral vector boson — the Z^0 .

In 1975 a new charged lepton, called the τ , was discovered. The presence of a new charged lepton heralds the presence of a new neutrino. Up to this date (1988), however, direct evidence of its production (the creation of τ s from a beam containing ν_τ) has not been seen. This is attributed to experimental difficulties. There is no serious doubt as to the existence of this neutrino.

One question that much theoretical and experimental work has gone into trying to answer has to do with neutrino masses. There is no compelling reason that neutrinos should have zero mass, even though most searches for any exhibition of neutrino mass come out with negative results. The experimental techniques for measuring neutrino masses generally use the process of tritium beta decay, and search for deviations in its endpoint which would be caused by a non-zero electron anti-neutrino mass. One positive result exists; the ITEP group in Moscow claimed in 1980 to have observed a deviation in the tritium beta decay spectrum corresponding to a 33 ± 12 eV/ c^2 mass. They confirmed this result in 1987 [10]. Other experiments tend to disagree, however. A group working at Zürich claims an upper limit of 18 eV/ c^2 for electron anti-neutrino mass, and one at Los Alamos claims a limit of 25 eV/ c^2 [11,12]. Double beta decay experiments place an even more stringent mass limit, however they are subject to nuclear matrix element calculations and are less reliable than direct measurements

[7]. Measurements of muon and tau neutrino masses set much less stringent limits, due to the sensitivity of the apparatus at higher energy scales. The limit on m_{ν_μ} is $250 \text{ keV}/c^2$, set by studying the decay $\pi \rightarrow \mu + \bar{\nu}_\mu$. The best limit on m_{ν_τ} has been set by the kinematics of the reaction $\tau \rightarrow 3\pi^+$. Its value is $70 \text{ MeV}/c^2$. [13]

Having delineated the properties of the neutrino, we digress at this point in order to put the neutrino into its place in the “Standard Model”, which describes the interactions of quarks and leptons.

1.2.2 Neutrinos in the Standard Model

All known interactions of which matter undergoes can be described by four forces. These are the strong, weak, electromagnetic, and gravitational forces. The strong force is responsible for holding the nucleus together, and is responsible for binding the components of the proton and neutron (called quarks) together. It is known as the strong force because it is considerably stronger than the electromagnetic force which would cause nuclei and even protons and neutrons to dissociate due to electrostatic repulsion between like charges. The weak force mediates radioactive decays. All known particles can interact via the weak interaction, but the cross sections are exceedingly small. The neutrino is a particle that interacts via the weak interaction only. As a result, it is exceedingly difficult to detect. The electromagnetic and gravitational interactions are responsible for all of the macroscopic properties of the world so far observed.

Throughout the 1970's and 1980's a "standard model" which explained all particle interactions through the exchange of vector bosons gained wide acceptance in the particle physics community. The "standard model" consists of two separate theories: quantum chromodynamics (QCD) and the Glashow-Weinberg-Salam (GSW) electroweak theory. QCD describes the strong force by introducing a "color" charge which is carried by quarks and their intermediate vector boson, the gluon. Since electrons and neutrinos carry no color charge, they are not affected by the strong force. GSW electroweak theory is the first successful unification of two forces since Maxwell unified the electric and magnetic forces one hundred years ago. This theory states that at a large enough energy scale, the electromagnetic and weak forces are the same. As the energy scale is reduced, however, a "spontaneous symmetry breaking" occurs which creates massive bosons to mediate the weak force (the $W^\pm$ and Z^0) and a massless boson (the photon) which mediates the electromagnetic force.

According to the GSW electroweak theory, only left-handed fermions and right-handed antifermions can couple to the weak force. "Handedness" is a property of the helicity operator:

$$H = \frac{\vec{J} \cdot \hat{p}}{\hbar} \quad (1.1)$$

where $\hat{p}$ is a unit vector in the direction of the momentum. A right handed particle has helicity of +1, whereas a left-handed particle has a helicity of -1. Massive particles always have components of both helicities, since it is

always possible to find a frame of reference in which the particle direction is reversed. Massless particles, as the neutrino may well be, are constrained to have a certain helicity. Hence the right-handed neutrino and the left-handed antineutrino are denied any means of interaction with other matter by the helicity dependence of the weak force. They are called “sterile” particles. They are not produced by weak interactions, and cannot be observed.

There are three known families of quarks and leptons. Each family consists of a quark of charge $+2/3$, a quark of charge $-1/3$, a charged lepton, and a corresponding neutrino. These are shown in table 1.1. No observation has ever been made of a neutrino interacting with a lepton from a different family. This has led to a law of “lepton number conservation”, which is not demanded by the standard model, and is strictly empirical. The potential exists, however, for mixing between lepton families that would defy this empirical law. This would lead to the phenomenon of neutrino oscillation, described in the following section.

1.2.3 The Neutrino Oscillation Hypothesis

In 1955 Gell-Mann and Pais [14] predicted mixing of the neutral kaon and its antiparticle via multiple-pion intermediate states. One implication of this is the creation of a long-lived neutral kaon state, which is prevented from decaying via the 2π mode because it is a $CP = -1$ eigenstate. This state was observed in 1956 by a Columbia-Brookhaven collaboration [15]. The novel concept underlying this prediction and observation is that a neutral particle,

	+2/3 Quark	-1/3 Quark	Charged Lepton	Neutrino
1st Family	u	d	e	ν_e
2nd Family	c	s	μ	ν_μ
3rd Family	t	b	τ	ν_τ

Vector Bosons	Gluons	
	γ^0	
	$W^\pm, Z^0$	

Table 1.1: The known families of quarks and leptons. The tau neutrino and top quark have yet to be observed. Brackets show which vector bosons mediate interactions between which particles.

propagating freely in space, was able to change into its anti-particle. The quantum mechanical explanation of this phenomenon is that the eigenstate present at the creation of the particle is a strangeness eigenstate. This is not the same as a weak eigenstate, and it is the weak eigenstate which is selected out by differential decay rates. The weak eigenstate can be written as an admixture of strangeness eigenstates, and vice-versa. If any state, strangeness or weak, is selected by a physical process and then allowed to propagate, there is a finite probability of it changing into its counterpart. This was called strangeness or hypercharge oscillation.

This discovery prompted Pontecorvo [16] to propose in 1958 that neutrinos could oscillate in the same manner as kaons; with the particle changing into its anti-particle. This can only occur if neutrinos are Majorana particles, in other words if the neutrino field and the antineutrino field are related via the charge conjugation operator [17,18] $|\nu\rangle = C |\bar{\nu}\rangle$. Another requirement is that these particles be massive, as will be shown in the next section. This theory places neutrinos in the unique position of being two component particles, whereas the quarks and leptons are described by four component spinors.

An alternate theory which leads to neutrino oscillation was developed in the mid-nineteen seventies [17,19,20]. It relies on drawing the analogy between the Cabibbo theory for hadronic weak interactions and neutrino theory. Here the neutrinos can be described by Dirac four component

spinors, much like quarks and leptons. Cabibbo theory utilizes a mixing angle (the Cabibbo angle) to specify the charged weak current as it acts in the quark sector:

$$\begin{pmatrix} d' \\ s' \end{pmatrix} = \begin{pmatrix} \cos \theta_c & \sin \theta_c \\ -\sin \theta_c & \cos \theta_c \end{pmatrix} \begin{pmatrix} d \\ s \end{pmatrix} \quad (1.2)$$

The u and c quarks are taken to be unmixed, by convention.

This spawned the idea that mixing could be possible in the lepton sector, as well as the quark sector. Specifically, it meant that there could be oscillation between the muon-neutrino and the electron-neutrino states. In the quark sector, the quantum number that was changed by the mixing of the quark states was strangeness (and later bottom). This was observed in the laboratory. Indeed, the anomalously long decay times of the strange particles, due to the smallness of the mixing, led to the coining of the term “strange” to describe them in the first place. In the case of the neutrinos, however, the quantum number that would be changed via the mixing would be lepton number — and violation of lepton number had never been observed in any physical process.

Discovery of the τ lepton (1975) and the b quark (1977) led to the conclusion that a third family of quarks and leptons exist, including a new neutrino, the ν_τ . This required a generalization of the theory of oscillations [21,17]. Instead of one mixing angle, a mixing angle was postulated for each pair of families. Hence, any of the following oscillations could potentially

occur:

$$\begin{array}{ll}
 \nu_e \rightleftharpoons \nu_\mu & \bar{\nu}_e \rightleftharpoons \bar{\nu}_\mu \\
 \nu_e \rightleftharpoons \nu_\tau & \bar{\nu}_e \rightleftharpoons \bar{\nu}_\tau \\
 \nu_\mu \rightleftharpoons \nu_\tau & \bar{\nu}_\mu \rightleftharpoons \bar{\nu}_\tau
 \end{array}$$

Any observed oscillation of neutrino states would have major theoretical and astrophysical implications. To go into these more explicitly, though, requires some theoretical formalism. This has been neglected up to this point in order to keep the discussion historical in nature. Let us diverge from this approach and explain the theoretical framework in which oscillations are usually discussed.

1.3 The Physics of Neutrino Oscillation

As previously mentioned, there are two possible theoretical constructs which describe the neutrino, the Dirac and Majorana theories. In the standard model, it is aesthetically pleasing to have quark-lepton symmetry. This puts neutrinos on the same theoretical footing as quarks and charged leptons. For this reason Dirac theory is the favorite for describing neutrino oscillations. It will therefore be discussed first, and then Majorana theory will be explained afterward. Where neutrino oscillations fit into the standard model and grand unified models will be the concluding segment of this section.

1.3.1 Dirac Neutrino Theory

Let us begin by assuming that there are two distinct neutrino types, one which couples to electrons and one which couples to muons. These are known as ν_e and ν_μ , as stated in the previous sections. The oscillation of one of these states into another requires two things:

- That there are non-zero, non-degenerate mass eigenstates corresponding each weak eigenstate.
- That these mass eigenstates are not identical to the weak eigenstates.

Let the mass eigenstates be known as ν_1 and ν_2 . The mixing angle θ describes to what degree the eigenstates mix. In matrix form:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} \quad (1.3)$$

This is analogous to the formulation of Cabibbo mixing in quarks, given in equation 1.2. The processes which generate neutrinos act only through the weak force. Hence the initial state is either $|\nu_e\rangle$ or $|\nu_\mu\rangle$. The initial state in the E645 experiment is $|\bar{\nu}_\mu\rangle$. All derivations given work for antiparticles as well. Let us examine the case where we have ν_μ as the initial state.

$$|\nu_\mu(t)\rangle = \cos \theta |\nu_2(t)\rangle - \sin \theta |\nu_1(t)\rangle \quad (1.4)$$

The time evolution of the wave functions is given by

$$| \nu_\alpha(t) \rangle = e^{iHt} | \nu_\alpha(0) \rangle \quad (1.5)$$

The mass eigenstates are the eigenstates of the Hamiltonian. Each one will correspond to a specific energy eigenvalue:

$$e^{iHt} | \nu_\alpha(0) \rangle = e^{iE_\alpha t} | \nu_\alpha(0) \rangle \quad (1.6)$$

where $\alpha = 1, 2$. Because the initial weak eigenstate is specified, the mass eigenstates are also specified. At $t = 0$,

$$\begin{aligned} | \nu_1(0) \rangle &= -\sin \theta | \nu_\mu(0) \rangle \\ | \nu_2(0) \rangle &= \cos \theta | \nu_\mu(0) \rangle \end{aligned} \quad (1.7)$$

Using equations 1.4 and 1.6 we can rewrite the equations for the ν_e and ν_μ eigenstates.

$$| \nu_e(t) \rangle = \cos \theta e^{iE_1 t} | \nu_1(0) \rangle + \sin \theta e^{iE_2 t} | \nu_2(0) \rangle \quad (1.8)$$

$$| \nu_\mu(t) \rangle = -\sin \theta e^{iE_1 t} | \nu_1(0) \rangle + \cos \theta e^{iE_2 t} | \nu_2(0) \rangle \quad (1.9)$$

Now that the states have been specified, we can determine the probability that one weak eigenstate oscillates into another as a function of time. Define this as $P_{\nu_\mu \rightarrow \nu_e}(t)$. We can also define the probability that the state does *not* oscillate into the other state as $P_{\nu_\mu \rightarrow \nu_\mu}(t)$. In the chosen notation,

these are given by

$$P_{\nu_\mu \rightarrow \nu_e}(t) = |\langle \nu_e | e^{-iHt} | \nu_\mu \rangle|^2 \quad (1.10)$$

$$P_{\nu_\mu \rightarrow \nu_\mu}(t) = |\langle \nu_\mu | e^{-iHt} | \nu_\mu \rangle|^2$$

Applying equation 1.8

$$\langle \nu_e | e^{iHt} | \nu_\mu \rangle = \sin \theta \cos \theta (e^{-iE_2 t} - e^{-iE_1 t}) \quad (1.11)$$

$$\begin{aligned} |\langle \nu_e | e^{-iHt} | \nu_\mu \rangle|^2 &= \sin^2 \theta \cos^2 \theta (e^{-i(E_2 - E_1)t} - e^{+i(E_2 - E_1)t}) \\ &= \sin^2 2\theta \sin^2 \left(\frac{(E_2 - E_1)t}{2} \right) \end{aligned} \quad (1.12)$$

Likewise, it is easy to show that

$$P_{\nu_\mu \rightarrow \nu_\mu}(t) = 1 - \sin^2 2\theta \sin^2 \left(\frac{(E_2 - E_1)t}{2} \right) \quad (1.13)$$

Let us now put the energy eigenstates in terms of neutrino mass.

$$E_2 - E_1 = \sqrt{p^2 + m_2^2} - \sqrt{p^2 + m_1^2} \quad (1.14)$$

Assuming $m_\alpha \ll p$, we can expand the equation in a Taylor series.

$$\sqrt{p^2 + m_\alpha^2} \cong p + \frac{2m_\alpha^2}{\sqrt{p^2 + m_\alpha^2}} \quad (1.15)$$

Equations 1.14 and 1.15 then give

$$E_2 - E_1 \cong \frac{2}{p}(m_2^2 - m_1^2) = \frac{2}{p}\Delta m^2 \quad (1.16)$$

Applying this to 1.11 and 1.13 gives

$$\begin{aligned} P_{\nu_\mu \rightarrow \nu_e}(t) &= \sin^2 2\theta \sin^2 \left(\frac{\Delta m^2 t}{p} \right) \\ P_{\nu_\mu \rightarrow \nu_\mu}(t) &= 1 - \sin^2 2\theta \sin^2 \left(\frac{\Delta m^2 t}{p} \right) \end{aligned} \quad (1.17)$$

Assuming the velocity of the neutrino is c , the time variable can be replaced by a distance variable. Putting the “ c ”s back into the equation to make the units physical then gives

$$\begin{aligned} P_{\nu_\mu \rightarrow \nu_e}(t) &= \sin^2 2\theta \sin^2 \left(\frac{1.27 \Delta m^2 L}{E} \right) \\ P_{\nu_\mu \rightarrow \nu_\mu}(t) &= 1 - \sin^2 2\theta \sin^2 \left(\frac{1.27 \Delta m^2 L}{E} \right) \end{aligned} \quad (1.18)$$

where Δm^2 is in units of $(eV/c^2)^2$, L is in units of meters, and E is in units of MeV. (If E is in GeV, the L is in km).

Experiments can measure the probability, P , for either the disappearance of a neutrino type or the appearance of a neutrino type. For any probability, there are a set of values for Δm^2 and $\sin^2 2\theta$ which correspond to it. Experimenters then, can only define a region of $\Delta m^2 \otimes \sin^2 2\theta$ space as being probable if an oscillation signal is observed. One way to obtain unique values of these variables is to do another measurement at a different distance or a different neutrino energy. If oscillations are not observed, limits on Δm^2 and $\sin^2 2\theta$ can be set. These lead to an exclusion curve in $\Delta m^2 \otimes \sin^2 2\theta$ space, as will be shown in the experimental section.

1.3.2 Majorana Neutrino Theory

In the theory presented above, neutrinos are Dirac particles, which can be represented by a 4-vector. This vector represents the helicity states of both neutrino and anti-neutrino. This makes it similar to any other fermion. In standard electroweak theory, only left-handed neutrino fields

and

right-handed anti-neutrino fields interact with matter. Some theories, starting with the original suggestion that neutrinos oscillate, make the supposition that the neutrino fields consist of a two component vector which represents left-handed neutrino and right-handed anti-neutrino states only. This is often called Majorana neutrino theory. The original Majorana theory [5] suggested the possibility that there was no distinct difference between neutrino and anti-neutrino. As stated in the previous section, Davis ruled out this possibility [6]. The possibility still exists, however, that there is some admixture of right-handedness in the weak current. If this is so, and there is also a breaking of lepton conservation, then reactions of the type $\nu \rightleftharpoons \bar{\nu}$ become possible. “Majorana” and “Dirac” become labels which describe parts of the Hamiltonian, with the “Majorana” part describing transitions between particle and anti-particle states.

If there are right-handed weak currents which couple to neutrinos, then it becomes possible to write the equation for the neutrino state as a superposition of helicity eigenstates, rather than the mass eigenstates used in the previous section. Papers by Primakoff and Rosen [22] and Bahcall and Primakoff [23] work out the implications of these assumptions, and write the weak states as

$$\begin{aligned} |\nu_e\rangle &= \alpha |\nu_- \rangle + \beta |\nu_+ \rangle \\ |\bar{\nu}_e\rangle &= \alpha |\bar{\nu}_- \rangle + \beta |\bar{\nu}_+ \rangle \end{aligned} \tag{1.19}$$

$$\begin{aligned}
|\nu_\mu\rangle &= \alpha |\nu_-\rangle - \beta |\nu_+\rangle \\
|\bar{\nu}_\mu\rangle &= \alpha |\bar{\nu}_-\rangle + \beta |\bar{\nu}_+\rangle
\end{aligned}$$

where α and β correspond to the mixing angles in the previous derivation. One unattractive aspect of this formulation is that to include ν_τ requires the existence of a fourth generation neutrino, which has not been discovered and may well not exist. This fourth generation would mix with the ν_τ in the same fashion as the first two generations mix in the equations above. The helicity is defined by

$$\begin{aligned}
h(\bar{\nu}_e) &= -h(\nu_e) = |\alpha|^2 - |\beta|^2 \\
&= \left(\frac{1 - \eta^2}{1 + \eta^2} \right) \frac{|\mathbf{p}|}{(|\mathbf{p}|^2 + m_{\nu_e}^2)^{1/2}}
\end{aligned} \tag{1.20}$$

where η is the admixture of right-handed current. The electron neutrino mass in equation 1.20 is given by

$$m_{\nu_e} = \frac{1}{2}(m_{\nu_1} + m_{\nu_2}) \tag{1.21}$$

where $m_{\nu_{1,2}}$ are the mass eigenstates. From these equations they derive the $\nu_e \rightleftharpoons \bar{\nu}_e$ probability:

$$P(\bar{\nu}_e \rightarrow \nu_e) = (1 - (h(\bar{\nu}_e))^2) \sin^2 \left(\frac{1.27 \Delta m^2 L}{E} \right) \tag{1.22}$$

It is clear that this equation is similar to those derived in the previous section. It is also clear, however, that in the absence of right-handed currents or a massive neutrino that such a phenomenon cannot occur.

1.4 Neutrino Oscillation Experiments

Over the past thirty years, a number of experiments have searched for neutrino oscillation. These searches have occurred at reactors, at high-energy accelerators, and now at a meson factory. The techniques involved in these experiments all involve direct detection of neutrinos. The different energy regimes explored present different types of difficulty and advantage.

The first of these is the variation of neutrino cross-section with energy. Perkins [24] gives the following approximation for inverse beta-decay cross section:

$$\sigma \approx \frac{4G^2 p^2}{\pi} \quad (1.23)$$

This strong cross-section dependence on energy gives an advantage to the high energy regime. There are other factors, however.

One problem with high-energy accelerators is that their flux is low. A reactor, for instance, can generate on the order of 10^{26} neutrinos per second. Compare this with 5×10^{14} neutrinos per second at the LAMPF beam-stop (considerably less at BNL or FNL), and one can see the attractiveness of the reactor as a tool to measure oscillations.

A final consideration is the energy/distance dependence of the oscillation probability (see Equation 1.18). The L/E term implies that one should measure neutrinos at as low an energy as is possible and at distances that are as great as possible in order to reach the most sensitive regions of Δm^2 . Of course the problem with these requirements is that there are

often threshold and energy-dependent efficiency effects which determine the minimum energy that it is practical to measure, and that as one increases L , the distance between the detector and the source, the flux is reduced by $1/L^2$. This is not quite true when neutrino horns are used at high-energy facilities, but is rigorously true at LAMPF and at reactors.

Taking all of these considerations into account, it is surprising that relatively competitive limits can be obtained using all three techniques. The reactors, which can measure neutrino energies on the order of 1 MeV at distances on the order of 10 meters have a definite advantage at setting limits in Δm^2 . Their sensitivity in $\sin^2 2\theta$ is severely limited by the fact that they are “disappearance” experiments rather than “appearance” experiments. A disappearance experiment assumes that the neutrino flux and cross sections are well known, and therefore the rate in the detector is accurately predicted. The experiment then looks for deviations from the expected rate as a signature of oscillation. An appearance experiment looks for the production of particles which can only be produced by a neutrino type other than that which is in the neutrino beam. It is very sensitive to small mixing angles, since it is looking for a forbidden process. This method cannot be employed at reactors (except to look for $\bar{\nu}_e \rightarrow \nu_e$, see Davis [6]), because the source is purely $\bar{\nu}_e$ and 1 MeV $\bar{\nu}_\mu$ s or $\bar{\nu}_\tau$ s are below particle production threshold and therefore unobservable. E645 is an appearance experiment, and we look for the reaction $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$. This will be described

in more detail in the next chapter.

1.4.1 Experimental Limits on Oscillation: 1988

It turns out that the limits reached by different experiments set their best limits in different regions of $\Delta m^2 \otimes \sin^2 2\theta$ space due to the different event rates and different L/E conditions. Therefore, the best limits on oscillation are not set by one experiment. Figure 1.1 shows the best limits on oscillation before E645, with limits set by the SKAT, E734, BEBC, and Gösigen collaborations. A brief description of each of these experiments is given below.

Serpukhov – SKAT [28] The SKAT experiment consists of a heavy-freon (CF_3Br) bubble chamber of 2.2 m^3 fiducial volume. The neutrinos were produced by π s and Ks, which decayed while propagating down a 140 m. decay tunnel. The bubble chamber was located 105 m. from the end of the decay region. The neutrino energies ranged between 2 and 30 GeV. The oscillation mode searched for was $\nu_\mu \rightarrow \nu_e$. SKAT was able to set very good limits in $\sin^2 2\theta$, due to the fact that the neutrinos were high in energy and therefore had a large cross section. The Δm^2 sensitivity was poor compared to other oscillation experiments.

BNL – E734 [25] This experiment was performed Brookhaven National Laboratory's AGS facility. The detector was located about 96 m

Figure 1.1: Best limits on neutrino oscillations as of 1988

from the neutrino source. The portion of the neutrino spectrum used was between 1 and 5 GeV. This was an appearance experiment, which searched for the oscillation $\nu_\mu \rightarrow \nu_e$. To differentiate between electrons and muons, a relatively fine-grained detector was constructed consisting of 112 planes of 8cm thick liquid scintillation counters and x and y proportional drift cells. The fiducial mass of the liquid scintillator was about 100 tons. The relative proximity of their experiment to the source compromised its ability to set competitive limits in Δm^2 , but this and their large fiducial mass allowed it to set excellent limits in $\sin^2 2\theta$ due to the enhanced flux.

CERN – BEBC [26] BEBC is a Ne-H₂ bubble chamber at the CERN PS.

It is located 825 m from the neutrino source, and has a fiducial mass of 14 tons. The neutrino source consisted of π s and Ks decaying in flight inside a 45 m. decay tunnel. The average neutrino energy was 1.5 GeV and ranged from about .5 to 3 GeV. The oscillation mode $\nu_\mu \rightarrow \nu_e$ was searched for. The much larger distance between the source and detector made the limit in Δm^2 much better than that of E734, but it also decreased the flux. The detector also has a much smaller fiducial mass; hence the limits in $\sin^2 2\theta$ were inferior to those set by E734.

Gösgen [27] Reactors are the most abundant source of neutrinos. Only disappearance experiments are possible, though, so a powerful lever

over $\sin^2 2\theta$ is lost. Reactor experiments do tend to set the best Δm^2 limits for large mixing. The best of these to date was performed at the 2.8 GWh Gösgen power reactor. The data was taken at a variety of distances from the reactor core: 40 m, 45 m, and 65 m. The detector consisted of scintillator interspersed with ^3He wire chambers., and required the detection of the positron and the neutron from the inverse beta-decay reaction. The detectable neutrinos varied in energy between about 1 and 5 MeV. Clearly L/E is far superior in this experiment, but the constraint of looking only at disappearance, $\bar{\nu}_e \rightarrow \bar{\nu}_e$, restricts the sensitivity to relatively large mixing angles.

1.5 Physical Implications of Neutrino Oscillation

In the final section of this chapter, the importance of neutrino masses and oscillations in particle physics and astrophysics will be briefly discussed.

1.5.1 Neutrino Mass — Physical and Astrophysical Implications

As mentioned in an earlier section, neutrino oscillation experiments do not provide a direct measurement of the mass of the neutrino. Still, one prerequisite for oscillation is that $\Delta m^2 = m_2^2 - m_1^2$ be non-zero. Thus the observation of oscillation would provide proof, which up until now has *not* come from direct mass measurements, that neutrinos do indeed have mass.

If one assumes that similar mass heirarchies exist in the neutrino sector as exist in the charged lepton and quark sectors (see, for instance, the “see-saw” mechanism in section 1.5.3) one can make the approximation $m_2^2 \approx \Delta m^2$.

The Glashow-Weinberg-Salam model of electroweak interactions assumes a zero neutrino mass. This is not a hard requirement of the theory, however. Indeed, the GSW model includes fermion masses in an arbitrary manner [29]. It invokes a coupling of the fermion fields to the Higgs field, with a specific coupling constant for each fermion which determines its mass. It is possible to include neutrinos among the massive fermions at the cost of extra coupling constants. Mixing between neutrino families, however, is not a part of the standard model, and is not trivially incorporated. It occurs naturally in some of the grand-unified models, as will be seen in section 1.5.3

Massive neutrinos do have cosmological impact. Specifically, neutrinos may determine whether the universe is open (i.e. will continue to expand forever) or closed (i.e. will slow its expansion, stop, and begin to contract), depending on their mass. It is believed that the expansion constant is close to the critical value between expansion and contraction, since most theories have it diverge very rapidly if it is not almost precisely equal to this value. There does not seem to be enough luminous matter in the universe to account for this requirement. So-called “dark matter” has been

postulated for this reason and to explain other observations [30]. One potential candidate for the dark matter constituent is the neutrino. A neutrino mass of about $30 \text{ eV}/c^2$ would be necessary to close the universe. This is consistent with the ITEP mass result, but is significantly more than the limits set by Los Alamos or Zürich. There do exist significant problems with cosmological models which use neutrinos to satisfy the dark matter condition. A completely satisfactory model, however, does not yet exist.

1.5.2 The Solar Neutrino Question

One important example of neutrino astrophysics can be found much closer to home; that is the tremendous flux of neutrinos generated by nuclear reactions in the solar interior. Most of these neutrinos have energies of a few MeV or less. Neutrinos from ^8B , however, have an energy range extending up to 14 MeV. They are above the threshold of such nuclear reactions as

$$\nu_e + {}^{37}\text{Cl} \rightarrow e^- + {}^{37}\text{Ar}.$$

An experiment by Davis [31] was set up in the early 1960s to detect this reaction. It used the same technique as his original reactor experiments (section 1.2.1). Argon at levels above background was detected, confirming the observation of solar neutrinos. The rate, however, turned out to be $\sim \frac{1}{3}$ of the predicted value. One reason quickly suggested was that the neutrinos were oscillating into ν_μ or ν_τ en route from the sun. If true, this would imply that $\sin^2 2\theta \approx 1$, and that $\Delta m^2 \gtrsim 10^{-10} (\text{eV}/c^2)^2$.

Another possibility was devised by Mikheyev and Smirnov [32] based on earlier work by Wolfenstein [33]. The theoretical crux of this work is that ν_e can couple to the electrons in the sun through a charged exchange current, while other species of neutrino cannot. This creates a new term in the Hamiltonian, which can be described as a potential which varies with electron density. If neutrinos are indeed mixtures of mass eigenstates, then a neutrino passing through a region of the sun with a certain critical electron density may find it energetically favorable to oscillate into another neutrino type. This is known as the MSW effect. A number of papers exist which explain the theory underlying this phenomenon [34,35,36]. The result of these calculations is that in order to explain the Davis result via this mechanism, $10^{-4} < \sin^2 2\theta < 1$ and $10^{-8} < \Delta m^2 < 10^{-5}(\text{eV}/c^2)^2$. Observation of an oscillation signal of this magnitude in the laboratory is well beyond the capabilities of current experiments.

New information, however, now throws doubt on the interpretation of the low neutrino flux as a sign of neutrino oscillation. Davis now reports [37] that the neutrino flux shows an anti-correlation with the sunspot cycle. There is no known oscillation-like effect which could cause such a variation. One suggestion is that since sunspots are associated with regions of high magnetic field, perhaps the left-handed neutrinos are being changed into right-handed neutrinos through the interaction of their magnetic moments with the solar magnetic field. These right-handed neutrinos would be

“sterile”; or unable to interact with matter. New limits on neutrino magnetic moment have been set, though, using neutrinos from the supernova 1987A [38,39]. These seem to rule out this possibility. It is likely that the variation, and perhaps the overall flux discrepancy, may be due to poorly understood dynamics of the solar interior. Experiments currently under construction which use gallium as a solar neutrino target will be able to measure neutrinos of much lower energy, and should have a much larger statistical sample. The answer to many of these questions awaits these new experiments.

1.5.3 Neutrino Oscillations and Grand Unification

As a final note on the neutrino, this segment discusses some of the ways in which neutrino masses are incorporated into grand unified theories. In section 1.5.1 it was stated that neutrino mass has to be put in by hand in the standard model, but so do other masses. Mixing between families, on the other hand, is not expected. Many extensions of the standard model have been proposed which incorporate massive neutrinos and mixing. Unfortunately, one of the features of all of these models is the vagueness of the predictions made as to the magnitude of the masses or the size of the mixing. Nonetheless, we present some of the salient features of a couple of them as a demonstration of the state of neutrino theory at this point.

One of the more popular theoretical mechanisms that has been developed over the past few years ties neutrino mass to lepton mass. This is

based on the observation that a heirarchal mass structure exists in both the charged lepton and quark sectors. It seems likely, then, if neutrinos do have mass that $m_{\nu_e} < m_{\nu_\mu} < m_{\nu_\tau}$ (Assume for the time being that the mass eigenstates are the weak eigenstates.) Gell-Mann, Slansky, Ramond, and Yagamida have proposed a model which sets a definite heirarchy of neutrino mass [40]. Called the GRSY or “see-saw” mechanism, it ties the neutrino mass to the mass of the lepton of the same family and to the mass of a hypothetical right-handed W boson:

$$m_{\nu_\alpha} \approx m_{l_\alpha}^2 / M_{W_R} \quad (1.24)$$

Of course the mass scale of the W_R is a matter of conjecture, as there is as yet no evidence for its existence. Different theories invoke different mechanisms to pin down this mass, and hence the neutrino masses. A couple of examples of this (by *no* means an exhaustive list) are given in a paper by Hung and Mohan [41] and another by Kang and Shin [42].

Hung and Mohan work within the framework of an $SU(3)_c \otimes SU(2)_L \otimes SU(2)_R \otimes U(1)$ left-right symmetric model. Among the predictions of this theory is the existence of eight or nine families, not all necessarily having light neutrinos. It can allow for M_{W_R} as low as 8 TeV, and masses as large as

$$m_{\nu_e} : .03 \text{ eV}/c^2 \quad m_{\nu_\mu} : 1.3 \text{ keV}/c^2 \quad m_{\nu_\tau} : 400 \text{ keV}/c^2.$$

Kang and Shin, in contrast, make use of an $SO(10)$ model. This model contains not only neutrino masses, but neutrino mixing angle as well. They

employ a see-saw mechanism to set the neutrino mass heirarchy. To set the scale of the right-handed sector, which in turn sets the scale of the light neutrino masses, they employ the invisible axion scale. The axion is created as a solution of the strong CP problem in the Peccei-Quinn theory [43]. Kang and Shin explain the arbitrariness of the parameters that go into their calculations. Indeed, using their range of acceptable values, one can come up with masses and mixing angles which vary over many orders of magnitude. Using values that are fairly central in their possible range, they obtain the mass range for ν_1 : $.2 \text{ eV}/c^2 < m_{\nu_1} < 100 \text{ eV}/c^2$. Their see-saw heirarchy is different as well. It couples to the quark masses rather than the lepton masses:

$$\begin{aligned} m_{\nu_1}/m_{\nu_2} &\approx \mathcal{O}(m_u/m_c) \\ m_{\nu_2}/m_{\nu_3} &\approx \mathcal{O}(m_c/m_t)\sqrt{m_u m_c/m_t^2} \end{aligned} \quad (1.25)$$

Their mixing angles depend on a heirarchal structure in both the quark and lepton sectors:

$$\begin{aligned} \theta_{e\mu} &\approx | \sqrt{m_e/m_\mu} + e^{i\phi_1} \mathcal{O}(\sqrt{m_u/m_c}) | \\ \theta_{\mu\tau} &\approx | \sqrt{m_\tau/m_\mu} - e^{i\phi_2} \mathcal{O}(\sqrt{m_c/m_t}) | \\ \theta_{e\tau} &\approx \sqrt{m_e/m_\tau} \theta_{\mu\tau} \end{aligned} \quad (1.26)$$

$\phi_{1,2}$ are phases.

The point of all this is to show that while much theoretical effort goes into the prediction of neutrino masses and mixing angles, the uncertainties

in these models renders their predictions somewhat arbitrary. The need for experiment is evident. Setting limits on mixing angles or masses helps to constrain these theories, and can lead to other predictions, such as the mass scale of the right-handed weak currents. An actual measurement of neutrino mass or an observation of neutrino mixing would be fascinating in its own right, in addition to constraining and weeding the plethora of models which currently exist. E645 is an attempt to add to the knowledge of the neutrino and thereby provide a further test of the standard model. This is the motivation for this experiment, and for this dissertation.

Chapter II

Design of the Experiment

2.1 Searching for Oscillations at a Pi-Meson Factory

A number of considerations exist when an experimenter chooses a neutrino source for an oscillation experiment. These are explained in section 1.4 in some detail, but will be briefly reiterated below

Flux Reactors are by far the most intense sources of neutrinos. High energy accelerators, due to their relatively small beam currents, tend to have relatively low neutrino fluxes. LAMPF, with its relatively intense proton beam, has a higher neutrino flux than other accelerators, but not nearly the flux of a reactor.

Distance To maximize flux one would put the detector as close as possible to the neutrino source. Opposing this is the sensitivity of the detector to the oscillation signal. The farther from the source, the greater is

the probability of oscillation. This is in accordance with equation 1.18

$$P_{\nu_\mu \rightarrow \nu_e}(L) = \sin^2 2\theta \sin^2 \left(\frac{1.27 \Delta m^2 L}{E} \right) \quad (2.1)$$

There is a tradeoff between flux and sensitivity in measuring Δm^2 that must be examined during the design of the experiment in order to maximize sensitivity.

Energy A number of different considerations go into the choice of neutrino energy. At very low energies, muon and tau neutrinos are not created and cannot be detected. This means that only disappearance experiments, which measure the difference between predicted and observed neutrino rate, are possible. This type of experiment lacks the statistical power of the appearance experiment, which searches for a neutrino type not present in the incident neutrino flux. At high energies, backgrounds from kaon decays produce neutrinos of the type which could otherwise be produced only by oscillation. Intermediate energy accelerators, such as LAMPF, produce ν_μ and $\bar{\nu}_\mu$ below muon production threshold. This helps to eliminate beam-correlated electrons in the energy range 30-50 MeV, which is where an oscillation signal peaks. Another factor that favors low energy is the L/E dependence of the Δm^2 limit. At low energies, a much shorter baseline can be used to obtain a Δm^2 result than would be required at higher energies.

For a comparison of L/E values (which help determine Δm^2 sensitivity) at high-energy accelerators and at LAMPF, see figure 2.1. E645 is in a position to set somewhat better Δm^2 limits due to the relatively low neutrino energy. In addition, the systematics and background are much different than those encountered at higher energy. The low level of contamination from higher energy neutrinos (in this case, $E > 53$ MeV) and from $\bar{\nu}_e$ adds much to the capabilities of this experiment.

2.1.1 Neutrino Production at LAMPF

LAMPF is the abbreviation of “Los Alamos Meson Physics Facility”, also known as the Clinton P. Anderson Meson Physics Facility. It is an 800 MeV, high current proton beam. In contrast to most high-energy facilities, which measure beam currents in microamperes, the LAMPF beam can achieve an intensity of over one milliamperere. It is a linear accelerator with an acceleration region over 1/2 mile long (figure 2.2) Its initial stages consist of a Cockroft-Walton accelerator which takes the protons up to .75 MeV, and a drift tube linac which takes them to 100 MeV. The rest of the accelerator is a side coupled linac which brings the beam energy up to 800 MeV. The time macrostructure of the beam consists of 800 μ sec buckets with a repetition rate of 120 Hz. Of these, about 70 buckets per second are delivered to area A. This yields a live fraction of 5.7%. The live time defines the unit of “LAMPF Day (LD)” as 5600 seconds. The beam microstructure has a rate of 200 MHz with a burst width of 250 psec. [44].

Distance vs. Energy for Oscillation Experiments

Figure 2.1: Distances versus neutrino energy for appearance-type oscillation experiments. The dotted lines represent $L/E = \text{constant}$.

Figure 2.2: The Los Alamos Meson Physics Facility

After the beam passes through a switchyard, most of it enters the main experimental hall (Area A) where it passes through a number of targets to produce particles for pion and muon experiments (figure 2.3). The incident current on the beam-stop is reduced by $\sim 18\%$ by the upstream targets, and the energy is reduced to 780 MeV. The protons then strike the A6 beam-stop and are absorbed (figure 2.4,[45]). The A6 target consists of a water target which increases pion production, isotope production stringers, a water-cooled copper beam-stop, and around everything plenty of iron and concrete shielding.

In the beam-stop, the protons interact to produce pions. Negatively charged pions are rapidly absorbed. Positively charged pions produce neutrinos via the decay chain:

$$\begin{aligned}\pi^+ &\rightarrow \mu^+ + \nu_\mu \\ &\hookrightarrow e^+ + \bar{\nu}_\mu + \nu_e\end{aligned}$$

In the absence of π^- , no electron antineutrinos are produced. Estimates of the π^- decay rate place it at $\sim 1/4000$ the π^+ decay rate [46]. This is a factor of 20 below the E645 experimental sensitivity. The spectra of neutrinos emitted by the beam-stop is shown in figure 2.5

Another process that can occur in the A6 target cell is known as the “decay-in-flight” decay. As can be seen in figure 2.4, not all of the space is filled. There is a small gap between the front window and the water degrader. In addition, all stringers are not always in the beam line. Open decay space allows room for pions to decay before coming to rest, which

Figure 2.3: Area A at LAMPF

Figure 2.4: The A6 Beam-stop. The upper figure shows the A6 target cell. The lower figure shows the isotope production stringers and the beam stop in more detail.

Neutrino Energy Spectra

Figure 2.5: Spectra of $\nu_\mu, \bar{\nu}_\mu, \nu_e$ from the π^+ decay chain.

can boost the energy of the emitted neutrino above the muon production threshold. Such events were observed in the detector, as will be discussed in the analysis section.

2.1.2 Neutrino Physics at a Meson Factory: Advantages and Disadvantages

Physics at a pion factory has distinctly different requirements and advantages than physics at high-energy or reactor facilities. The advantage in L/E is shown in figure 2.1. This allows a better Δm^2 limit than can be achieved at higher energy accelerators. Likewise, the fact that it is an appearance experiment helps it to set much better $\sin^2 2\theta$ limits than reactor experiments. While not background-free, it does have the property that contamination by wrong-sign neutrinos is so small as to be ignored. At higher energies, kaon decays create neutrino backgrounds which must be subtracted properly. Being below the muon production threshold is also a great help in minimizing neutrino-induced backgrounds.

A number of challenges are presented by the meson factory environment. First, neutrino-induced event rates tend to be low, due to the small cross-section. This means that a detector for such a facility should have a substantial mass. In case of E645 this is 22.3 metric tons. It also requires that the detector be operated relatively close to the beam-stop. The choice for a distance which optimizes L/E vs. event rate put the detector at 26.8 m from the target. A second obstacle is the relatively long duty cycle of

the accelerator. Because it is live 5.7% of the time, cosmic ray induced backgrounds present a significant problem. A final source of background is due to the copious neutron production at the beam-stop. Fast neutrons can reach a nearby detector within the beam gate, and cause a background which is beam-correlated.

2.2 Detector Requirements

The purpose of the E645 detector is to be able to identify the inverse beta-decay $\bar{\nu}_e + p \rightarrow e^+ + n$. It must be able to distinguish this decay from backgrounds induced by cosmic rays, backgrounds from beam neutrons, and backgrounds from other neutrinos. To attain this, the detector was designed with the following characteristics.

- An active shield to reject cosmic ray muons and other charged cosmic ray components.
- Passive shielding to absorb the hadronic component of the cosmic ray flux.
- Shielding to eliminate the beam-stop neutron flux.
- Tracking and calorimetry to make unambiguous identification of minimum ionizing particles.
- Gadolinium to absorb and tag the neutron from the inverse beta-decay.

- Past and future time history to allow the identification of stopping muons in the past and neutron captures in the future.

A full description of these follows.

2.2.1 Cosmic Ray Rejection

The flux of cosmic rays at Los Alamos (7500' elevation) is roughly 3 times that at sea-level. Cosmic rays can contribute to a number of different backgrounds. They can be divided into two different classes: the “muonic” component and the “hadronic” component. Hadronic in this case can mean a number of different incident particles, such as protons, neutrons, pions, or even heavier nuclei. They interact via the strong force and so are attenuated rather quickly in matter. What remains after these particles have ranged out is the muonic component, mostly muons from pion decays in the upper atmosphere. Their mean energy is about 8 GeV [47]. These may stop in the detector and decay : $\mu^+ \rightarrow e^+ + 2\nu$, The energy spectrum of the electrons is similar to that of the positrons from inverse beta-decay. Luckily, because the muons are charged, they can be rejected by an active veto shield around the detector with high efficiency. The hadronic component, however, has a neutral subset which can penetrate a veto shield. These must be removed through passive shielding. Passive shielding is also critical in removing beam-related neutrons and cosmic-ray photons.

2.2.2 Particle Identification & Calorimetry

Neutrons, from both cosmic rays and the beam-stop can interact via several mechanisms. Higher energy neutrons can interact through elastic np scattering, producing a high energy proton. Because the protons are more highly ionizing than electrons, they will have a shorter track. Hence they will often have the same event topology as an electron track. Separating them out requires accurate determination of dE/dx . To estimate this correctly, one must also be able to determine particle direction. To meet these requirements, the E645 detector was designed to be relatively fine-grained. Its construction utilizes a sandwich structure of 3 cm scintillator tanks interspersed with horizontal (Y) and vertical (X) proportional drift tubes. The scintillator gives an accurate energy measurement, while the PDTs give both energy and position information.

This construction also helps to eliminate other neutron-induced sources of background. High energy neutrons can react on nuclei and generate gamma rays: $n + N \rightarrow n + N' + \gamma$. High energy gamma rays can produce Compton or pair-production electrons which can be mistaken for signal. Being able to eliminate reactions where a considerable amount of energy falls off of the track yields a reduction of this type of background.

Thermal neutrons can capture on Gadolinium, producing gamma rays:

This reaction can occur as a result of thermal neutrons from the beam

or from cosmic rays. The energy of this reaction is far too low for it to be confused with the signal. In addition, these reactions do not usually produce a clean track. Events of this type are common in the data sample, but are easy to eliminate.

Accurate energy measurement is also required if one hopes to distinguish the energy spectrum of the inverse beta-decay positrons from the spectra of other neutrino induced reactions. As will be seen in the analysis chapter, the energy measurement capability was adequate for this task.

2.2.3 Access to Past and Future Event Information

The ability to use timing provides another very important tool in eliminating backgrounds. Both past and future histories can be used for this purpose. Past history allows one to tag events in which energy was deposited in the past which is spatially correlated with the candidate event. An example of this is the stopping muon. In cases where the active shield is inefficient, it falls on the detector to be able to identify the muon in the past which decayed to produce the triggering electron. Another instance in which past information is useful is in identifying “decay-in-flight” events in which muons are produced by high-energy neutrinos. To remove these backgrounds, 57 μsec of past history is digitized for each event.

The window into the future serves the purpose of allowing the neutron from an oscillation event to thermalize, transport, and be captured by the gadolinium. This window goes from the trigger time until 110 μsec

afterwards. It can also be used for identifying decay-in-flight events in which the muon has sufficient energy to trigger.

2.3 Components of the Experiment

In the final section of this chapter, all components of the experiment will be described in some detail. These will be divided up into the following categories: the central detector, the active shield, overburden and neutron shielding, electronics, and the data acquisition system.

2.3.1 The Central Detector

The E645 central detector is a fine-grained, 22.3 ton neutrino detector (see Figure 2.6). It consists of 40 planes of alternating liquid scintillation counters and X and Y (horizontal and vertical) proportional drift tubes. Gadolinium oxide paint has been applied to mylar sheets which have been hung over each scintillation plane. Each plane measures about 3.7 meters $\times$ 3.7 meters in area, and is about 18 cm. thick. The length of the detector is 6.93 m.

The scintillation plane consists of 12 scintillation modules, each of which measures 3.7 m. $\times$ 30 cm. $\times$ 3 cm.. The liquid scintillator provides the main source of hydrogen nuclei for the inverse beta-decay reaction in addition to measuring particle energy. Phototubes at each end convert the emitted light into electrical pulses.

The PDT planes are made up of 5 banks of 9 wires each. The tube

Figure 2.6: The E645 central detector and anti-coincidence shield.

containing each wire measures 7.6×3.8 cm.. These tubes are glued together to make a bank, 3.8 cm. thick and 69 cm. wide. Each tube contains a wire running down the center, and aluminized mylar (held at ground) on the inner surface. The wire is both the signal and high voltage wire. The gas mixture used is P-10; 90% argon and 10% methane.

The PDT planes are mounted on aluminum frames, while the scintillator planes are hung by nylon straps, as shown in figure 2.7. Both hang from rollers on an overhead frame and can be moved back and forth during servicing. The signal and high voltage cables run through trays into an electronics hut, which is mounted on the frame of the detector external to the active shield. The cart upon which the detector is mounted is part of the active shield, and it can move on railroad tracks between the service building and the neutrino tunnel. Each component will now be described in some detail.

Scintillation Counters

There were two primary considerations which lead to the choice of liquid versus solid scintillation counters. First was the abundance of hydrogen in the liquid scintillator. Liquid scintillator has the empirical formula CH_2 , whereas for solid scintillator the formula is closer to CH . For the same number of hydrogen targets, nearly twice the mass of solid scintillator would be required. The second consideration is price; liquid scintillator is cheaper.

Figure 2.7: Scintillator and PDT planes, and their mountings.

A number of different scintillators were tested for light response. The best candidate found was the Bicron 517 series, which uses pseudocumene (1,2,4 tri-methyl benzene) as its active ingredient, dissolved in a light mineral oil base. The final scintillator chosen, 517L, consists of 15% pseudocumene, 1g/l ppo, and .01g/l Bis-MSB. The containers for the scintillators are extrusions of polymethyl methacrylate (Rohm & Hass VM-100). These are illustrated in figure 2.8. Initial tests of the scintillation counters with the scintillator Bicron 517H showed them to be susceptible to stress fractures due to a crazing process. The following measures were taken after studying the crazing problem (Appendix B):

- A different scintillator was chosen. Bicron 517L has half the pseudocumene concentration as does 517H, but the light output is about 80% that of 517H. It was determined that this was adequate. The lower concentration substantially reduced the rate of crazing.
- A central rib was glued into the extrusions to prevent deformation due to hydrostatic pressure. The walls of the extrusions are only 1/8" thick, and so the hydrostatic pressure due to the 1 foot column of liquid scintillator was enough to cause about 8 mm of distortion on each side of a counter which is normally only 3.0 cm thick. Crazing tends to occur at points of high surface stress. By preventing most of the counter deformation, the rib substantially reduced both the rate and probability of crazing.

SCINTILLATION COUNTERS (480)

Figure 2.8: Schematic of the acrylic extrusions used as containers for liquid scintillator

- By shifting the stress to the rib, the region of crazing shifted to the glue joint between the rib and the extrusion. This was remedied by applying a fillet of polyvinyl alcohol (PVA)[48] to the joint.
- Crazing will not occur if stress, solvent concentration, and temperature, fall below a certain threshold. Keeping the detector temperature low helps to preserve the counters.

Despite these precautions, about 20% of the scintillators have experienced some leakage over the lifetime of the experiment, although only a fraction of these experienced any sort of catastrophic failure. Broken scintillators were replaced before the run. At no time during running did the number of leaking counters exceed 10% of the total.

Total internal reflection transports light from its production point, along the scintillator, and into the PMTs. The refractive index of the scintillator and the acrylic plastic of the extrusions is almost the same, so the light is reflected off of the outer surface of the counter. To keep this surface clean and dry, the extrusions were wrapped with an inner layer of white paper. The attenuation length achieved by this technique varies between 2 and 3 meters, with a mean near 2.6 meters. If a scintillator leak occurs, it acts as a light sink. In this case, the attenuation length can drop to less than a meter. To prevent light from entering the extrusions, they were wrapped in Tedlar; a black, opaque plastic wrapping. The ends were fitted with black plastic boots, which were constructed with holes for fittings and for the

PMTs. These were affixed to the Tedlar layer with black electrician's tape. The extrusions were wrapped with a final layer of thick mylar to protect the relatively fragile Tedlar from abrasion.

The photomultiplier tubes were chosen based on two characteristics: low noise and high gain. Timing was not considered a critical issue. Timing resolution is on the order of a few nanoseconds. Low cost was also a leading consideration. Based on these requirements, the Hammamatsu R878 photomultiplier was chosen. It is a ten stage tube with a 2 inch photocathode, and has a quantum efficiency of 25%. Not among the requirements, but surprisingly fortuitous, was the remarkable reliability of these tubes. In the three years of operation there have been no failures aside from a couple of tube breakages during servicing. Bases are constructed from a simple resistor chain with a flat taper (equal voltage drops across the stages), as shown in figure 2.9. They are compact and lightweight due to the minimal current requirements of the tubes. The tubes are run at voltages between +1200 and +1500V. The maximum recommended operating voltage is +1500V. The photocathode is held at ground.

The photomultiplier tubes were extensively tested to determine their gain and noise[49]. The extrusions were tested for physical integrity, defined as the ability to hold 1 ft. H₂O of air pressure without leaking. Prototypes were bench tested for their attenuation length using a cosmic ray telescope as trigger. They performed well in the 1987 run if there was no scintillator

Figure 2.9: Circuit diagram for the base used with the Hamamatsu R878 photomultiplier tubes.

leakage. A sample attenuation length plot and Landau spectrum ¹ are shown in figures 2.10 and 2.11, respectively.

Proportional Drift Tubes

The proportional drift chamber operates by detecting electrons in a gas which are freed by the ionization created by charged particles. An electric field transports them to a signal wire. In many versions of these chambers there are separate "sense" wires, upon which the electrons are collected, and "field" wires, which are responsible for shaping the electric field. In our design these functions are carried out by the same wire. The voltage is chosen so that the device operates in the "proportional" mode, which means that the signal measured is proportional to the number of ion pairs created. At higher voltages, a breakdown occurs and the signal saturates. This is known as the "Geiger" mode. Amplification occurs in the gas by an avalanche effect as electrons are accelerated through the gas toward positive HV. The gas must be chosen so that it is readily ionizable but will not saturate. This is usually accomplished by using a noble gas to provide ionization electrons and a hydrocarbon to absorb electron energy and suppress Geiger breakdown. The E645 PDTs use P-10, a mixture containing 90% argon and 10% methane. It is also desirable that in a proportional

---

¹The Landau spectrum is the energy spectrum obtained from an energy-measuring device in response to normally incident muons. Its general form is that of a Gaussian distribution with a tail going out to high energies. The tail is due to delta rays (high energy electrons from elastic scattering) and is a function of detector thickness as well as the material surrounding the detector.

Figure 2.10: Light attenuation curve for a typical extrusion

Figure 2.11: Landau distribution for a typical scintillation counter

drift tube the “drift time” that it takes the ions to reach the sense wire be proportional to the distance between the sense wire and the ion track.

Proportional drift counters have customarily been made out of aluminum. If such counters were to be used in the LAMPF neutrino beam, they would introduce background from the reaction $\nu_e + {}^{27}\text{Al} \rightarrow {}^{27}\text{Si} + e^-$ which produces an electron similar in energy to the positron of inverse beta-decay. This potential problem was prevented by a rather unconventional approach: making the E645 PDTs out of cardboard [50].

The tubes are Kraft paper laminates. They form rectangular tubes 3.8 cm thick and 7.6 cm. wide, for an aspect ratio of 2:1. On the inside they are covered with aluminized mylar, wrapped helically to make a conducting surface (see figure 2.12). Only 1 mil of aluminum is deposited on this surface, which is held at ground. Nine of these tubes are glued together to make one PDT bank. After drying to remove excess moisture, they are coated with epoxy to seal the Kraft paper from the atmosphere. Aluminum end channels are affixed to the ends, and 50 μm gold-plated tungsten wires are strung through each cell. On one end of the bank, a high voltage board is mounted. On the other is installed, an amplifier board with a ribbon cable which runs to the electronics hut.

The PDT banks were tested in several ways. Early prototypes were tested on a bench which used MWPCs for tracking and scintillators for triggering. In this ideal setup, the spatial resolution (σ of the distance

Figure 2.12: Schematic of the E645 proportional drift tube

distribution) was determined to be 1 mm.. Figure 2.13 shows the distance versus drift time characteristics of the E645 PDTs and of an aluminum PDT of the same size obtained using the MWPC test bench. Another test bench was constructed to test the tracking capabilities of a sample fraction of the PDTs. This used other PDTs for tracking. In this more realistic set-up, spatial resolution was measured to be 2 to 3 mm.

There are 10 PDT banks per plane, 5 horizontal and 5 vertical. Gas is circulated through the 9 PDTs which make up a bank in series, but each bank and plane is fed in parallel. The gas system is constructed to recirculate the P-10 gas through the detector with a relatively slow make-up rate. All PDTs were tested for leaks prior to assembly to minimize the amount of gas required. For a circulation rate of $300 \text{ ft}^3/\text{hr}$, the make-up rate necessary to prevent gas contamination is $30 \text{ ft}^3/\text{hr}$.

The PDTs performed well in the 1987 run, with measured efficiencies of 96% for stopping muon electron tracks. PDTs are generally used in order to obtain drift times, which provide a very precise position measurement. This information was not used, since the cell size is sufficient to provide adequate tracking ability for tracks which go several planes. Pulse height information, went into the particle identification algorithm. While the Landau spectrum is relatively wide (figure 2.14), indicating poor resolution compared to the scintillators, there are twice as many PDT dE/dx samples as there are scintillator dE/dx samples. This adds significantly more information which

Figure 2.13: Distance versus pulse height for the E645 paper PDTs and an aluminum PDT of the same size.

can be used in particle identification.

Gadolinium Sheets and Neutron Detection

The inverse beta-decay reaction creates a free neutron in the detector. If this could be detected in coincidence with the positron from this reaction it would produce a very strong oscillation signature. The coincidence requirement would substantially reduce backgrounds. As mentioned in Chapter 1, a technique of this type was used to confirm the existence of the neutrino. ^{157}Gd is a nucleus with a capture cross-section for thermal neutrons of 26,000 barns, orders of magnitude larger than that of any other isotope. In addition, it gives off a rather large amount of energy in the form of gamma rays upon capture. As mentioned previously

The average number of gammas given off is 4.0, each with an average energy of 1.7 MeV [52]. These interact in the scintillators to give a signal which is above the 200 keV threshold.

The gadolinium is in the form of an oxide, Gd_2O_3 , which is blended into a paint and sprayed onto mylar sheets. The thickness of the gadolinium oxide layer is 10 mg/cm². These sheets are hung over the scintillator plane so that they cover both sides.

To test the efficacy of this scheme, three tagged fission sources are installed in the detector, one at plane 10, one at plane 20, and one at plane

Figure 2.14: Landau distribution for a typical PDT

30. These sources are small ionization chambers which use a ^{252}Cf source as a trigger. The fission fragments, which are highly ionizing, cause the trigger. Each fission gives off an average of 3.7 neutrons [53]. These thermalize and capture on the gadolinium. The time response of the gadolinium to the fission neutrons is shown in figure 2.15. The efficiency of this reaction is rather low, and depends strongly upon the signature required for the event. Accidentals due to natural radioactivity become predominant as the requirements are loosened. Maximum signal to noise is obtained for an efficiency of about 20% and an accidental rate of about 4% [54].

2.3.2 The Active Shield

To remove muons from the data sample requires an active veto, in other words that the detector be uniformly surrounded by an active layer which can tag charged particles. Once again, liquid scintillator has been chosen for the task. It fills a hollow layer in a cylindrical vessel that completely covers the detector, as shown in figures 2.6 and 2.16 [55,57].

The vessel, known as the “shield”, is constructed in three distinct segments. The “cart” is a platform on wheels, upon which the detector is mounted. Wings jut upward at an angle along the side to provide overlap with other layers of scintillator. Photomultipliers are attached on the bottom, and look upward into the scintillation layer through hemispherical glass windows. The scintillator used is a diluted version (Bicron 517P) of

Figure 2.15: Neutron Rate as a Function of Time, Before and After the ^{252}Cf Tagged-Source Trigger

ANL-P-18,843

Figure 2.16: The E645 Active Cosmic Ray Shield

the scintillator used in the detector. The internal walls are painted white to minimize light absorption. Also mounted on the cart is the "blue wall", a circular wall which forms a barrier between the detector and the electronics hut. Around these segments fits the "cylinder", a cylindrical construction which moves on its own set of rails. This unit also contains the forward counterpart of the blue wall. The PMTs on the blue wall and cylinder are mounted on the outside in a triangular grid pattern. There are 46 PMTs on each of the two circular walls, 68 PMTs on the cart, and 200 on the cylinder. When the cylinder is meshed with the cart, it provides very close to 4π active shielding around the detector. A weak area exists around the joint between the blue wall and the cylinder. To compensate, a semi-circular ring of solid scintillator is mounted over this crack.

Active shielding is not sufficient to eliminate all classes of background. There exist a number of processes, such as muon bremsstrahlung and muon capture on nuclei, that produce neutral particles such as photons and neutrons which can penetrate the active shield without triggering it. These neutral particles can produce electrons in the detector, mimicing the oscillation signal. To prevent these backgrounds, another layer of shielding was installed inside the scintillation layer of the cylinder and blue wall. A second hollow region inside the one filled with scintillator, 7" thick and bounded on either side by 1" steel walls, was filled with lead shot. This is equivalent to about 5" of solid lead. On the cart, lead bricks were stacked

and the cracks filled with lead shot for a net thickness of about 3" of lead. The rate of the muonic backgrounds, the attenuation by active and passive shielding, and the final predicted rate of these backgrounds in the detector are given in table 2.1.

Table 2.1: Expected rates for various muonic backgrounds in the E645 detector

Background	# LD ⁻¹	Active Shield	Passive Shield	e or γ Rate LD ⁻¹
$\mu \rightarrow e\nu\nu$	6×10^5	$< 10^{-7}$	—	$< .06$ (e)
$\mu \rightarrow \mu\gamma$	1500	—	3×10^{-6}	$< .005$ (γ)
$\mu \rightarrow e\nu\nu \rightarrow e\gamma$	760	—	3×10^{-6}	.0022(γ)
$\mu^- N \rightarrow {}^0_n X$	550	.8	.14	56 ${}^0_n$'s $\rightarrow \sim .08\gamma$ s

The shield is used in two ways. It provides an active veto, which is defined as a preset number of tubes above a certain threshold. Each segment has an independent threshold. If two or more segments fire, it is assumed that the particle is a through-going muon, and the veto lasts only 1.75 microseconds. Otherwise the particle could potentially be a stopping muon, so the veto lasts 10.75 microseconds. The shield also provides pulse height and timing information for all tubes which is later used in the data reduction routines to help eliminate the muonic backgrounds.

The rate of incident muons through the detector is estimated to be about 3 kHz. The trigger rate with the veto on is .4 Hz. Of these triggers, 70% seem to be muons, either stopping or through-going, giving a muon rate of .28 Hz.. This is a 10^{-4} rate reduction. After application of the shield and

detector information in the analysis, the final background sample contains only 27 events for 269 LD, or about .1 ev/LD. This rate, 1.8×10^{-5} Hz, is a reduction of 6×10^{-9} in overall cosmic ray background.

2.3.3 The Physical Site and Shielding

In order to reduce the rate of backgrounds caused by the neutral cosmic ray component to the level of .1 LD⁻¹, it is necessary to shield the detector with the equivalent of 2000 gm/cm² of overburden. To accomplish this while still allowing access to the detector for servicing, a tunnel has been constructed which extends to within 24 m. of the beam stop. The detector and active shield can be moved between the service building and the tunnel on a set of rails.

The Neutrino Tunnel

Having a tunnel rather than a fixed site for the detector provides the solution to a number of physics and engineering problems. First, the detector could be constructed in an open space, with full support from an overhead crane and fork lifts. This substantially increases the ease and speed of construction. Secondly, it reduces the overall size of the underground chamber required, and thereby reduces the cost. Finally, it provides a way to measure oscillation at two different distances. If a positive oscillation signal was to be observed, the detector could be moved to confirm the measurement and to unambiguously assign values to Δm^2 and $\sin^2 2\theta$.

The tunnel is constructed of arch-plate in a cylinder that extends about 30 m.. There is a concrete slab which runs the length of the tunnel and out to the service building. Upon this are mounted a number of rails upon which the detector and shield can be rolled in or out of the tunnel. The tunnel is off center from the axis of the beam (which runs roughly west to east) by 9° , and 15° below horizontal with respect to the beam stop. The polar angle between the line from the center of the detector to the beam-stop and the axis of the beam is 19° . The tunnel is backfilled with tuff (volcanic ash — the local soil) to the top of the tunnel and covered with steel shielding.

Cosmic Ray Overburden

Measurements of the cosmic ray neutral component as a function of depth reveals that to achieve a background level of .1 electron-like event LD^{-1} the detector has to operate under an overburden of approximately 2500 gm/cm^2 (see figure 2.17). Were this to be in the form of tuff, which has a density of $1.6 gm/cm^2$, a pile of tuff 17 meters high would be required. It was decided to use steel instead. Aberdeen proving grounds provided the roughly 2000 tons of the steel used as cosmic ray overburden in the form of armor plate which had been used in tests of uranium bullets. This is arranged over the detector site in the manner shown in figure 2.18.

In order to provide shielding at the end of the tunnel, a large cylindrical

water tank on wheels is used. This tank is 8 meters long, with a cross section the same shape as the tunnel opening but a few centimeters smaller. Access to the detector is provided by a tunnel through the center of the tank. The tank can be moved on retractable wheels when empty. It sits on its flat bottom when full.

Beam Neutron Shielding

High energy neutrons can cause two classes of background. The first, knock-on protons from np reactions, are only a problem insofar as one can confuse protons with electrons. The ability to separate these particles, based on data taken during the 1984 beam test, ranges from 10^{-3} to 10^{-4} , depending on angle (figure 2.19). Another background is from $n\gamma$ reactions on nuclei, which can produce high-energy gamma rays and therefore high-energy electrons. The cross-section for these reactions is rather low. In either case, a large number of neutrons is needed before the background levels become appreciable.

During the 1986 data run 1370 beam-correlated protons per LAMPF day were observed in the data. Data reduction at the time showed a beam excess rate of 1.5 electron-like events/LD, which is consistent with a proton misidentification rate of 10^{-3} or an $n\gamma/np$ ratio of about 10^{-3} . The proton rate was much larger than that which could be accounted for by neutron leakage through the shielding between the beam-stop and the detector

Cosmic Ray Rate as a Function of Overburden

Figure 2.17: Relative Cosmic Ray Flux as a Function of Overburden

Figure 2.18: Schematic of the E645 Cosmic Ray Shielding

Electron-Proton Separation, 1984 Beam Test

Figure 2.19: Electron - Proton Separation from the 1984 Beam Test Data

(table 2.2) [56]. Investigation of the proton positional and angular dependence implicated two beam pipes which led from the beam stop to near the headwall (and were supposedly filled with shielding) as the probable sources of the beam neutrons.

Table 2.2: Estimate of neutron attenuation from layers of material between the beam-stop and the detector

Material	Thickness (gm/cm ²)	Attenuation Length (gm/cm ²)	Attenuation
Cu Beam-stop	523	175	5.04×10^{-2}
Beam-stop Steel	1067	171	1.95×10^{-3}
A6 Vault Wall	218	156	2.47×10^{-1}
Fe B.S Shielding	598	171	3.03×10^{-2}
Tuff Berm	2290	156	4.21×10^{-7}
Tunnel Headwall	182	156	3.11×10^{-1}
Tunnel Steel	1495	171	1.59×10^{-4}
Total			1.54×10^{-17}

After the 1986 run, a number of measures were taken to reduce the neutron rates:

- The neutron pipes were pumped full of concrete.
- Additional steel was buried in front of the headwall.
- Extra cosmic ray shielding was added on top of the detector.

The 1987 beam-excess proton rate was measured to be $.78^{+.44}_{-.44}LD^{-1}$, a thousand-fold reduction in rate. Assuming an electron-proton separation ratio of 1000, one expects a rate for beam-neutron induced electron-like backgrounds of less than $10^{-3}LD^{-1}$.

Figure 2.20: The beam-stop and its position with respect to the detector, showing the neutron beam pipes which were filled with concrete for the 1987 data run.

2.3.4 Electronics

The detector consists of 480 scintillation counters (each with two PMTs), 3700 PDT wires, and 360 shield PMTs. All information from these channels must be recorded from 57 μsec before the trigger occurs to 110 μsec afterwards. Since this was beyond the capability of any commercially available system, it was necessary that the E645 electronics be constructed by the experimenters and their electronics support groups. These systems are described in some detail below.

Scintillation Counter Electronics

As explained in section 2.3.1, each scintillation counter has one Hamamatsu R878 photomultiplier tube on each end, and a base of simple design. Ten Power Designs HV-1547 40 mA high-voltage power supplies feed the distribution panels which power the photomultipliers. The voltage for each channel is adjusted according to its response to through-going muons, using a procedure to be described in the next chapter. Photomultiplier signals are conducted via BNC coaxial cable into the electronics hut and into the shaper/splitter modules.

The PMT shapers are both amplifiers and signal shapers. Signal shaping is necessary due to the relatively narrow PMT signal, which is tens of nanoseconds in width. The digitizer ADCs (described below) sense voltage, so in order to obtain several voltage samples near the peak value, the PMT

pulse is stretched to 2 μ sec in length. It is also desirable to have them similar in shape to the PDT pulses, in order that the same software and digitizer microcode might be used.

Each shaper can handle 24 channels, or one plane. The shaped and amplified outputs go out on ribbon cables to the digitization system. Power for the shapers (± 12 V) is supplied by these cables as well. Another set of outputs, which is not shaped or amplified, goes through coaxial ribbon cable to the trigger system. This is not a direct feed-through; in order to reduce noise from the digitization system which was passing through the shapers into the trigger system, it was necessary to install a decoupling stage, which consists of a gain = 1 amplifier. The next phase of PMT pulse processing will be discussed in the sections on digitizers and triggering electronics.

PDT Electronics

The PDT high voltage is set by two dual-channel, low current HV power supplies. Each of these supplies power for 20 planes. They are connected to a distribution panel which in turn supplies every PDT bank individually. Every bank has a HV board at one end which distributes the HV to each of the nine wires. The PDT wires are held at +1950 V. The current drawn is small, varying between .05 and .2 microamperes per plane under normal running conditions. Variation occurs due to humidity. The high voltage boards absorb water when the air is humid, increasing the currents and also

the noise.

Each of the nine wires runs into an amplifier/shaper board mounted on one end of every nine-bank. There is a built-in test circuit that allows testing of the amplifiers without direct access to the board. Each amplifier board has a connector for a flat ribbon cable. These run from the detector into the electronics hut, where they are fed into the digitizers.

The Digitizer System

A number of requirements are imposed on the data acquisition system.

- It must be able to record the past and future time window around the event.
- It must have good dynamic range. For neutron detection, it is necessary that the scintillation counters be able to record pulses as low as 2 – 3 MeV. A proton at large angle, however, can deposit as much as 40 to 50 MeV of energy in one counter.
- It must have fairly high channel density (the number of channels per crate of electronics). E645 has on the order of 5000 channels. The electronics for these channels can take up no more than seven CAMAC crates, which is the maximum number of crates that CAMAC allows.

The design chosen for the digitizer system solves all of these problems. The key to the design is the use of the CA3300 6 bit flash encoder as the

analogue to digital converter. This device operates on an 81.4 nsec clock. In order to expand the dynamic range, two of these chips are used for each digitizer unit. These are arranged in a bilinear fashion. The channels 0-63 represent an amplifier gain of 27.2 and the channels from 64-127 represent an amplifier gain of 3.4. This arrangement allows sensitivity in the low pulse height region, yet still allows large pulses to be measured without saturating the ADC. See figure 2.21

Two steps are taken to increase channel density and reduce cost. The first step is multiplexing. Trigger requirements bias the detector toward acceptance of forward-going tracks. Since this is so, it is fairly rare for more than two adjacent PDT or scintillation cells to be hit by the same particle, so it is possible to use the same electronics for different channels. Each digitizer sums the signal from nine PDT wires (one bank) or eight scintillation counters (1/3 plane). Each channel has an associated discriminator, set at about 10 mV, so that the identity of the wire that fires is recorded. Occasionally, there is ambiguity in energy reconstruction. As an additional space and cost-saving measure, FASTBUS crates are used, rather than CAMAC crates. This allows five digitizer units to be installed on each card. There were an average of 21 cards per crate in five crates, for a total of 106 cards. Each crate has a controller module, through which data is fed from the digitizer cards into CAMAC.

Figure 2.21: Block Diagram of the Digitizer Circuit

Each digitizer channel, in addition to its flash ADC, has a 2K RAM which records the entire history of the event. The history consists of the pulse height and discriminator mask status during each clock cycle. After a trigger is received, the clock continues to run for 110 μ sec and then stops. The controller in each crate then reads all the data from the digitizer cards, compacts it, and sends it through CAMAC.

Each controller has an AM2914 microprocessor that compacts and re-formats the data in the digitizer memories (see figure 2.22). Compaction of the data is accomplished by imposing an ADC threshold, below which data is ignored. This is set to an ADC value of three channels. The processor assembles the digitizer information into pulses, determines the time, pulse height, and pulse width of each pulse, and tags it with discriminator mask information. The controller is read out through a CAMAC interface module, which also sends a LAM signal to the PDP-11 when data compaction and reformatting is complete. The processor has a 4K working memory. If the data volume becomes too large for this to handle, this triggers an alternate read-out mode known as "flushed event". In this mode, compaction is not completed and the event is sent over in a much less processed form. This can be recovered in software, but it leads to large dead time and event volume. Normal processing of the data by the controller leads to a 3:1 reduction in event size and simplified event format. Hence much effort was directed during the run to minimizing the fraction of flushed events.

Figure 2.22: Block Diagram of the Digitizer Crate Controller

The Triggering System

The triggering system consists of two subsystems which became known as the "fast logic" and the "trigger processor". The fast logic is connected directly to the shapers, described above. Its purpose is fast discrimination of all pulses produced by the PMTs. It performs a logic function of determining whether one tube produces a pulse above threshold (an OR) or whether two tubes on the same scintillation counter detect pulses above threshold (an AND), on a per plane basis. This information is then sent to the trigger processor, which checks the pattern of the incoming AND and OR hits to determine whether a pre-set trigger condition has been satisfied. If so, then the processor checks for veto and gate conditions (such as the beam and the shield), and then issues a trigger for the experiment.

The fast logic operates at a 4.0 MHz clock rate. All PMT signals pass through the shaper unamplified, and enter the fast logic via coaxial ribbon cables. Each cable corresponds to 12 PMTs which represent the south or north side of a plane. Each fast logic card handles two planes of PMTs, or 24 channels. There are therefore 20 fast logic cards, which are mounted in a FASTBUS crate. At the front end of each fast logic channel is a discriminator, set to about 6 mV, which corresponds to an energy deposition of about 1.4 MeV. The status of this discriminator is fed into a latch, where it is held for one clock cycle, then lost. This is necessary due to the delay in getting a trigger signal out in time to stop the clock. The timing

is set so that the contents of the storage register will correspond to the pattern which triggered the detector. A fast logic controller reads out the bit pattern from each card after a trigger occurs. This controller can also write to as well as read from the storage registers, and is therefore also used in diagnostics. Logic on each individual card determines whether the AND or OR conditions for either plane are satisfied during each clock cycle, and a corresponding ECL signal is sent out via coaxial cable to the trigger processor.

The trigger processor is a two stage triggering system. The first stage checks for specific patterns of hits in six-plane segments. The second stage determines whether these six-plane segments fired in a specific pattern. This is possible through the use of 12 bit (4K) memories which are loaded with the trigger patterns (see fig. 2.23). Each of these memories is addressed by 12 lines (corresponding to six planes), which come from the fast logic. Prior to the run, the processor memories are loaded with a "trigger pattern", which means that at all memory locations which correspond to a desired trigger are loaded with "1", and all others are loaded with "0". The standard trigger during the 1987 run was 3 out of 4 possible ANDs. There are 14 of these memories in the trigger processor, each of which handles six planes. The six plane groups overlap with each other by three planes so as to eliminate inhomogeneities at the boundaries between memories.

Outputs from each of these "first level" memories go into a "second

Figure 2.23: Trigger Processor Schematic

level" memory. This is also loaded with a desired trigger pattern at the beginning of the run that controls which segments of the detector are allowed to participate in the trigger. During data-taking mode, it has always been set so that all detector segments can generate a trigger. Different patterns are used during calibration, in order to concentrate cosmic ray data in one region of the detector. If the second level memory has a logical "1" as an output, the primary trigger condition is satisfied, and the processor will generate a trigger contingent upon two other conditions. The first of these is that it be enabled by a NIM level. This allows gating of the processor with the beam gate and associated cosmic-ray gates, and keeps it inoperative while data is being compacted and transferred. Another is that an ECL signal corresponding to a veto from the shield not be present. If all conditions are satisfied, a trigger is issued and the experiment is read out.

The trigger processor resides in a modified CAMAC crate. The first level memory modules are single width CAMAC modules, with the backplane connectors modified to use pins rather than the standard CAMAC card-edge connectors. There are seven of these modules. A specially constructed backplane allows them to communicate with the logic processor, which is a double width CAMAC module. For more details as to the construction, operation and testing of this unit and the fast logic, see Appendix A.

Shield Electronics

The active shield[57] is equipped with electronics similar to those used in the central detector. High voltage is supplied to the photomultipliers by LeCroy LRS 4032 and LRS 1440 power supplies. The BNC signal cables run into a patch panel in the electronics hut. The output from the patch panel is BNC ribbon cable, which then goes into the shield digitizer/veto system.

The shield digitization system bears many similarities to that of the detector. There are important differences, however. Since many PMTs view the same event, multiplexing is not possible. Each channel, therefore, has its own flash ADC. The controller is much less sophisticated as well, not having a microprocessor. Compaction of data is done through a comparator circuit, which leaves out all data below a (digital) threshold and records a time for all data above threshold.

The veto condition is determined via a multiplicity circuit that counts the number of tubes above a discriminator threshold. There are three distinct units in the veto system, one corresponding to the cylinder, one to the blue (rear) wall, and one to the cart. Each of these has an independent, software selectable multiplicity threshold. There is a summing circuit on each card that counts the number of tubes above threshold. This is carried along the backplane to each card, which adds its contribution. The boundaries between the three units are hard-wired on the backplane. If any

segment has a number of tubes firing above threshold greater than the set multiplicity, a veto is issued. There are two veto conditions, long and short veto. The long veto is issued if there is only one segment which satisfies veto conditions. The short veto is issued if there are two or three segments which satisfy the veto condition. The multiplicities of each segment are recorded for every clock cycle, although the identity of each tube that fires is not.

The running conditions for the 1987 run were as follows:

- Cylinder Multiplicity = 5 tubes
- Blue Wall Multiplicity = 4 tubes
- Cart multiplicity = 3 tubes
- Long Veto = $10.75 \mu\text{sec}$
- Short Veto = $1.75 \mu\text{sec}$

2.3.5 The Data Aquisition System

The various components which make up the data aquisition system have been described in the previous section. This section will begin with an overview of the system as a whole, and then explain the functions of the data acquisition computers; the PDP-11 and the VAX 750.

A block diagram of the data aquisition system is shown in figure 2.24. In data-taking mode, there are three sepearate clocks which run the electronics.

The digitizers are run by a 12.3 MHz clock, the trigger logic and shield veto logic by a 4 MHz clock, and the shield digitizers by a 6.7 MHz clock. The trigger logic searches for a pattern of hits that represents a viable trigger pattern. The three additional conditions that must be met to trigger the detector are that no shield veto be present, that the module be enabled by either "beam" or "cosmic-ray gates"² and that the computer be not busy. Upon finding a good trigger, the processor shuts off its clock and latches the trigger pattern into the fast logic storage registers. It sends out a signal to the digitization system to begin compaction and a LAM signal to CAMAC.

The digitizer logic upon receiving this signal generates a 110 μ sec delay in order to record the future history of the event. After the delay, the digitizer and shield clocks are stopped. Compaction and readout then begin. Upon completion of compaction, each digitizer controller sends a LAM to CAMAC, which triggers its read out. After all controllers have transferred their data, the shield digitizers are read. Other devices; the fast logic, scalers, and beam gate condition are also read out during this period. After all devices have transferred their data through the CAMAC system, the clocks are turned back on, and all devices are reenabled.

---

²The beam and cosmic ray gates are synchronized with signals generated by the accelerator control center. The beam gate is on while the beam is active, 6% of the time. The "pre-beam" gate is centered within the period, between beam gates and is 2.5 times as long as the beam gate. The "post-beam" gate is active immediately after the beam gate, and is approximately equal in length to the beam gate. "Cosmic-ray" data refers to the data taken during either the pre-beam or post-beam gates. No differences have been observed between the data taken in the pre-beam and post-beam gates.

The E645 Data Acquisition System

Figure 2.24: Block Diagram of the Data Acquisition System

The PDP-11

The device which drives the CAMAC interface system is a PDP-11, 16 bit computer. The CAMAC interface is a Kinetic Systems 2061. The PDP-11 performs all data aquisition functions, initialization functions, and diagnostics.

Almost all electronic components of the experiment have diagnostic routines that can be run using the PDP-11 and CAMAC. The fast logic has read/write capability in its one bit latches, which is tested by the program FLG. The trigger processor has memory and logic tests which can be run by the program TRG. The CON and DG1 programs test the digitizer controllers and digitizer cards, checking the many memory chips and logic functions in this system. SHIELD tests the shield electronics. All systems utilized the PDP-11 very heavily during the installation and debugging phase of the experiment.

A number of systems require initialization at the beginning of a data run. Trigger patterns must be loaded into the trigger processor first and second level memories. The microcode, which runs the digitizer controller microprocessor, must be loaded from a file on the PDP-11 disk into the processor memory. Multiplicity thresholds and veto lengths must be set in the shield logic. Programs are run on the PDP which perform these functions immediately before the beginning of every data run. These programs, when possible, perform simple diagnostics as well.

A program called BHCD runs the data acquisition, and performs all of the functions described in the previous section. It then packs the data into 5000 word buffers and ships it over to the VAX 750, where it undergoes preliminary analysis and is written to tape.

The VAX 750

The VAX 750 is responsible for accepting the data from the PDP-11. The data is transmitted on a Control Data DR-11W in 5 K word blocks. Three different processes are run on the VAX to control data taking and analysis. The controlling program MCNTL, drives the other processes and coordinates them. The process MCDBU is responsible for taking the event from the buffer, reformatting it, and storing it in a global buffer. Each process takes events off of its own event queue, which is a list of available events created by the previously executed process. Each event is labeled with a pointer to its location in the buffer. MBUTP is the process after MCDBU, so for each event it takes the pointer off of the queue filled by MCDBU and processes it, which in this case means writing the event to tape. MPROC is the next process. It performs analysis of the experiment to verify that all systems seem to be working properly. It analyses only a small fraction of the events because it is rather slow with respect to the data rate. As the global buffer fills, events can be deleted from it, but only after the event is written to tape. This often means that events are deleted from

the MPROC input queue before this process has a chance to analyze them. After MPROC, other analysis routines may have access to the data.

Several means of monitoring the experiment have been installed on the VAX. A continuously updating status display (STATS), determines and displays rates and event type classifications from scaler information taken with each event. This allows many problems to be spotted in a trivial manner. Another useful diagnostic tool is MANAL1, which allows the user to examine any event analyzed by MPROC in detail. It has the ability to display events graphically, to book and display histograms, and to dump event information to the screen or printer. Most electronics failures can be identified by these two tools.

The data reduction necessary to obtain a neutrino candidate sample goes beyond the capabilities of the 750, particularly when it is being used for data acquisition. Instead, the data was analyzed at the LAMPF Data Analysis Center (DAC), which maintains two VAX 8650s. The procedures and algorithms which determine the final sample will make up the bulk of the next chapter.

Chapter III

Data Reduction

The 1987 run consists of three accelerator cycles (cycles 48, 49, and 50) spanning the period from June 15, 1987 to December 1, 1987. During this time, E645 wrote 172 data tapes, with a total of 1.3 million usable events. The expected background rate from neutrino interactions is on the order of 20 events. This means that the data reduction routines have to achieve on the order of 10^{-5} background rejection. The backgrounds fall into three main categories. In order of importance, they are: electrons from stopping muons, protons from neutron elastic scattering on hydrogen, and photons from inelastic neutron scattering on nuclei. All available detector and shield information has been used to eliminate these backgrounds. Backgrounds from neutrino interactions are more subtle in that they are very similar in signature to the oscillation signal. They will be discussed in the next chapter. The thrust of this chapter will be to explain how the data analysis is used to obtain a final neutrino candidate sample.

3.1 Data Repacking

Data are written to tape in a format which maximizes both speed and density of information. This format, called “raw” data format, is not ideal for data analysis. The data are preprocessed prior to analysis to reorder the data into a more convenient form, using the ZBOOK data handling routines [58]. These divide the event information into blocks which delineate each type of data. Associated with each data block is a pointer in the event header corresponding to the address at which the data block begins. The ZBOOK data blocks, in order, are:

1. Header Block: Run name, event number, size, date & time.
2. Fast Logic Latch Block: Latches read out from fast logic at trigger time.
3. CAMAC Scaler Block: Used to obtain rates, beam gate information.
4. Dummy Block: Was to be used for high voltage read in.
5. Detector Block: All scintillator and PDT information.
6. Shield Block: All shield PMT pulse height information.
7. Shield Multiplicity Block: The number of tubes which contribute to shield veto, as a function of time.

8. Dummy Block: Used in Monte Carlo to store generated event parameters.

Of these, the detector and shield blocks merit further notice. In the detector block, pulses have been arranged into time "slices", which contain all pulses occurring within about a microsecond of each other. The slices are ordered such that the slice which occurs at the time of the trigger comes first, with subsequent slices ordered sequentially in time up to the last slice, which contains all hits having no other hits occurring within a microsecond of them. Slices times are defined as the average time of all the scintillator hits occurring within the slice.

Pulse information is ordered differently for scintillators and PDTs. The scintillator variables have information for both north and south PMTs, regardless of whether or not both fired. The information order for each pulse is:

1. Module number: 1-1000 for scintillators, 1001-3000 for Y PDTs, 3001-5000 for X PDTs.
2. Pulse height: This takes up two words for scintillators (south, north) but only one for PDTs. This is given in units of digitizer channels.
3. Time: This is the time of the pulse. There are two words for scintillators, one for PDTs. The units are in microseconds.

4. Sum: This is the integral of the digitized pulse, in digitizer channels. Scintillators use two words.
5. Status: This is a coded version of the digitizer mask. There is only one word used for either data type. It is a numeral which specifies how many masks fired in the digitizer bank shared by the displayed pulse. For scintillators it is a two digit numeral; one digit for south, one for north. For PDTs only one digit is necessary.

The shield block shares a similar structure. The slices, however, are in sequential order, and there is no mask word.

3.2 Calibration

Once a digitized pulse has been obtained from a scintillator or a PDT, the next step is to determine the deposited energy that the pulse represents. This requires that all devices have a measured relation between pulse height and energy deposition. The standard is provided by through going cosmic-ray muons. These particles are sufficiently relativistic so as to have a uniform energy deposition along their track through the detector ($\sim 2 \text{ MeV gm}^{-1} \text{ cm}^2$).

Scintillator and PDTs differ in the way that their signals are normalized. PDTs are the simpler system, having only one pulse per wire corresponding to a muon hit. In such a case, one simply corrects the pulse height for the track angle and histograms the result. The peak of the histogram represents

the most probable energy loss of a minimum ionizing particle traversing a PDT plane. The scintillators are a more complicated in that there are signals from both south and north photomultipliers which correspond to each muon hit. A muon traversing a scintillation counter will deposit a certain amount of energy depending only upon its angle. Light emission in scintillator is directly proportional to energy deposition¹ The problem in estimating light deposition, however, is that the light is exponentially attenuated as it propagates down the counter. The observed pulse height can be expressed

$$PH1 \propto E_d e^{-x/\lambda} \quad (3.1)$$

where $PH1$ is pulse height, E_d is deposited energy, x is the distance from the PMT to the hit, and λ is the attenuation length. On the opposite side of the counter (length L), the pulse height is given by

$$PH2 \propto E_d e^{(x-L)/\lambda} \quad (3.2)$$

A number proportional to deposited energy can be obtained by multiplying both pulse heights together and taking the square root:

$$\begin{aligned} PH1 \times PH2 &\propto E_d^2 e^{-L/\lambda} \\ (PH1 \times PH2)^{1/2} &\propto E_d \end{aligned} \quad (3.3)$$

If only one PMT fires it is still possible to obtain an energy estimate if the

---

¹This is an approximation. There is saturation for large dE/dx . Quantitatively, this is given by Burke's Law: $dL/de \propto \frac{dE/dx}{1+\alpha dE/dx}$ where $\alpha \sim .01 \text{MeV}^{-1} \text{ gm cm}^{-2}$.

attenuation length of the counter is known. These were also determined using cosmic ray muons.

3.2.1 Dedicated Calibration Runs

The active shield, described in the previous chapter, effectively prevents the detector from triggering on through-going muons. Accumulating a muon sample requires a dedicated run in which the shield veto is not enabled. To calculate calibration constants accurately, it is necessary to have at least one hundred events per scintillator or PDT, which allows statistical uncertainties of up to 10%. Since calculating PDT constants for 3700 wires would require seven times the number of events necessary for scintillator calibration, the PDT calibration constants were obtained on a per 9-bank basis rather than a per wire basis.

A special set of triggers is necessary to obtain cosmic rays with a shallow tracks evenly distributed throughout the detector. If a cosmic ray is too steep, there will be loss of precision in dE/dx estimation, since the larger pulse heights will push the digitizers far into the second ADC, which has a lesser precision than the first ADC. Shallow muon tracks can be required by the trigger processor through the use of second-level memory. For instance, one could require that both the first and last six-plane blocks of the detector trigger. By requiring any set trigger, however, there will be some areas of the detector which are heavily biased against. In the example above, the upper and lower counters in the center of the detector would go virtually

unpopulated, because muons going through these regions would have to be almost horizontal in order to pass through the first and last segments of the detector. Hence, the only way to illuminate the detector in a uniform manner is to collect data from a number of runs, each of which uses a different trigger pattern. The trigger patterns used in the calibration runs all require 4 out of 6 hits in two consecutive (but overlapping) processor segments. A series of runs is done which steps through the detector, with each run requiring triggers in a different pair of nine-plane blocks. In this way, all regions are evenly populated with muons, and the two segment requirement keeps the tracks from being too steep.

The first step in processing the calibration run is the construction of a muon track. This specifies the muon angle, and thereby the correction to be made to the pulse height (for the scintillators) or pulse area (for the PDTs). Each PDT or scintillator hit along the track has its pulse height or area histogrammed. There is one histogram for every PMT, and one for every PDT nine-bank. In addition, there is a two-dimensional plot of distance versus pulse height for every scintillation counter, which is used to determine attenuation length. The pulse height or area histograms represent Landau distributions, the peaks of which are used as the calibration constants. The scintillator attenuation lengths, the PMT pulse height peaks, and the PDT pulse sum peaks are all entered into a calibration file which is used by the analysis routines. The mean scintillator pulse height falls

at about 35 digitizer channels. This is near the half-way point of the 64 channel lower ADC.

Calibration runs are done while the accelerator is off, in order not to lose beam data. A full set of calibration data takes about 14 hours to accumulate, and contains approximately 120,000 events. During the 1987 data run, six calibration runs were done, with an average of one month between each run.

3.2.2 Run-to-Run Adjustment of PDT Gains

Monitoring of the PDT gains reveals that they fluctuate as a function of time. A direct correlation is seen between these gain changes and the atmospheric pressure. Because PDT pulse heights play a major role in particle identification, it is important to remove this effect. Using muons for this task is not practical. The veto effectively eliminates almost all through-going muons, and the muons which remain in the form of stopping muons tend to have steep tracks because they do not have constraints on their angle due to the trigger condition. Instead, it was decided to use electrons from the stopping muon events. Calibration constants are determined on a run-to-run basis by taking stopping muon electrons from each run, and determining the most probable pulse areas. This is compared to the PDT information in the most contemporaneous calibration set, and a multiplicative correction factor is obtained, which is written to a file which is used during the analysis to calculate PDT gains.

3.3 Structure of the Data Reduction Package

The evolution of the analysis package was driven by two factors. One was the obvious need to eliminate all backgrounds. The other was the desire to make the analysis as easy and efficient to run on the entire data set as possible. The first requirement is paramount. The order in which data rejection criteria (known henceforth as "cuts") are applied, however, has no bearing on final background rejection. Cut ordering was manipulated in order to enhance ease of execution.

The primary factors in determining ease of execution were CPU limits on batch jobs (90 minutes) and the amount of available scratch disk space on the DAC computers. Another was the "stability" of each cut. After processing of the data, it would become apparent that the criteria for certain cuts needed to be changed. It is very desirable that the entire data set not have to be reprocessed to affect such changes. To achieve all of these goals, the analysis was divided into three distinct stages. The first pass is designed to achieve a five fold reduction in data volume. The entire first pass can be run on the 1987 data sample in one to two weeks. The second pass is designed to create an output file that can fit easily onto available scratch space and can be run by a third pass routine interactively. The second pass routine can be run in two to three days. Often, cuts are repetitious, with variations on the same cut occurring in different passes. This allows gross

reduction of data volume in the initial passes with fine tuning occurring on the remaining sample.

The remainder of this section will be used to describe the structure of each pass: the cuts applied and their effect on the data volume. The subsequent sections will describe each cut in some detail, and will be organised according to type of cut rather than to the order in which cuts are applied.

3.3.1 First Pass

The first pass takes the ZBOOK output and sorts the events into classes, as well as applying cuts. It removes “random” triggers (triggers initiated by a pulse generator) and writes them to an output file. This can be used for estimation of shield noise rates. It does internal consistency checks to make sure that certain test (“fiducial”) pulses exist. A check is done to see if there are any hits in the active shield at the same time as the trigger, which indicates that the trigger was caused by a charged particle entering the detector from the outside. This eliminates muons. Track finding is done in this pass, as is a minimal fiducial volume cut. stopping muons with constructable muon and electron tracks and shield hits are written to a separate data file. These are extremely useful for calibration and efficiency estimation. Events which survive these requirements are written to an output file. The first pass also creates a print file, a histogram file, a CAMAC scaler file, and an event tracing file. The breakdown of events from the 1987 run is shown in table 3.1.

Table 3.1: Data Flow Through the First Pass Analysis

Cut or Condition	Events Remaining
Events Read	1269639
Detector and Shield Fiducial Pulses	1267944
Random Triggers Removed	1150719
In-Time Shield Cut	789844
Track Found	469362
Fiducial Volume Cut	443119
Stopping Muons Removed	204518

3.3.2 Second Pass

The second pass begins with a fiducial volume cut, after which it is clear that most of the remaining events are stopping muon related. A more stringent stopping muon cut is made at this point. Subsequently, most of the remaining events are neutron-induced protons. Particle identification, consisting of an electron confidence level calculation, is used to remove these events from the sample. The electron confidence level (the probability that the event being examined is an electron) is obtained through the comparison of scintillator and PDT dE/dx values for the candidate event with a similar set of values obtained from stopping-muon electrons. A cut is made on the confidence level, but it is not made as stringent as it is in the third pass. A subset of the sample with the lowest probability of being an electron is written out to a “proton” file.

Many of the parameters calculated in the second pass are used to apply future cuts. To prevent the often time-consuming calculations from having

to be repeated, an information file entry is written for all events processed. This contains energy, tracking, particle ID, and other parameters which allow overall second pass performance to be tallied and provide information which the third pass uses to apply cuts. The flow of the second pass through the 1987 data is shown in Table 3.2.

Table 3.2: Data Flow Through the Second Pass Analysis

Cut of Condition	Events Remaining
Events Read	204518
Fiducial Volume Cut	111788
Stopping Muon Cut	21979
Degree of Freedom Cut	20033
Particle ID Cut	3984
End Counter Cut	3018

3.3.3 Third Pass

The final pass of the analysis produces the neutrino candidate sample. It begins by applying a further fiducial volume cut. The next cut is on electron confidence level. This is a much more severe cut than was done in the previous pass. After this cut, the confidence level is not flat. The events responsible for this tend to have deposited much of their energy off of the track, and so are probably photon or neutron induced. They can be eliminated by putting a cut on off-track ("associated") energy. After this, most events are electron like, and the likelihood levels are flat. There are a large number of events remaining. Most have shield hits occurring some time

before the track. From this, one concludes that they are stopping muons and that somehow the muon is not seen. The only recourse in this case is to apply a cut which uses only the information in the active shield. This cut reduces the data volume by a factor of 10. The remaining events are neutrino-like: contained electron-like events with no evidence of a cosmic ray muon. At this point an attempt is made to remove “decay-in-flight backgrounds”, which are high energy ν_μ interactions, by identifying the muon generated by the neutrino interaction. The final cut constrains the measured energy to be less than 60 MeV. What remains are candidates for neutrino-induced events. The step-by-step breakdown of the third pass is shown in Table 3.3. Each cut is broken down into beam-on and beam-off categories.

At this point the discussion diverges from the chronological order in which the analysis is done and looks in more detail at the cuts described in this and the previous sections. The last section of this chapter will concern itself with the final sample of neutrino-like events.

3.4 Consistency Checks

A number of steps are taken to assure that the data are not corrupted and that there were no hardware problems present when data were taken. This is done throughout the earlier parts of the analysis package.

The first set of consistency checks is applied during the ZBOOK refor-

Table 3.3: Data Flow Through the Third Pass Analysis

Cut	Total Events	Beam On	Beam Off
Total Events Processed	3018	690	2328
Fiducial Volume Cut	2873	653	2220
Particle ID:			
1 Scint. on Track	2752	630	2122
Confidence Level	2066	476	1590
2 or More Hits Discd.	1987	454	1533
Associated Energy	1739	392	1347
Shield Cuts	147	59	88
Past Hit Near Track	99	45	54
Decay-In-Flight Removed	77	30	47
Total Energy (Final Sample)	47	20	27
Decay-In-Flight Cand.	22	15	7

matting routine. There is sometimes a problem in reconstructing flushed events (section 2.3.1). Pointers which define the location of the data segments are checked to be sure that they are within certain limits and sequential, otherwise the event is discarded. This leads to a loss of 1% of the data.

Another check is a hardware cross-check known as the "fiducial pulse". Fiducial pulses are generated at trigger time by the detector electronics. Six pulses are sent out of a fan-out. Five of these go to digitizer cards in each of the five crates. The sixth goes into the shield digitizers. The first pass ascertains whether these pulses were received by the digitizers. The absence of these pulses indicates that pulse reconstruction is not working properly. There were short periods throughout the run when not all fiducial pulses were present. During these it was necessary to bypass this cut. When in use, the fiducial requirement eliminated only .1% of the data.

3.5 Track Reconstruction

Electrons in the energy range from 20 to 50 MeV tend to have short, well defined tracks. This energy range includes both electrons from stopping muon decay and from inverse beta decay. A precisely known track direction makes accurate dE/dx measurements possible, and is therefore critical for good particle identification. This section describes the algorithm which this analysis uses to construct particle tracks.

One standard method for constructing a track is a least-squared fit to the points along the track, otherwise known as a minimum χ^2 fit. This technique utilizes all hits and finds the line which best fits all points. If there are hits off of the track or multiple tracks, however, the track found will have a very poor fit. Hits off of the track are common in the E645 data sample. These can be caused by two processes. First, cross-talk can occur in the digitizers. This takes the form of multiple discriminator masks firing on the same pulse, thus making it appear that the particle passed through more than one chamber. Another source is low-energy bremsstrahlung emitted by the electron, converting in nearby chambers. For this reason, it was decided not to use a standard χ^2 track construction algorithm, which would include extraneous hits at the cost of a poor fit, but to instead construct a track which passes through the most devices possible and ignores extraneous hits.

The algorithm is as follows: A track is first constructed using the vertical (Y) PDTs, then using the scintillators, then using the horizontal (X) PDTs. First, the maximum extent of the track is found in the Z (beam axis) direction. The program starts at the endpoints and draws a line between them. The number of hits that occur within 8 cm. of this line must be greater than 75% of possible hits for short (≤ 4 planes) tracks or 25% of possible hits for longer tracks. If this is not true for the track endpoints, the algorithm loops over the other hits, using their coordinates as the track endpoints and determining whether the line between each set of endpoints

meets the tracking criteria. This requirement is applied to both X and Y PDTs. For scintillators, 50% of possible hits must be on track regardless of track length. If no track can be constructed for which these conditions are met, the event is cut. The tracking algorithm then loops over all possible tracks and selects the one with the most hits within 8 cm.. Once the best fit is found, slopes and intercepts are calculated in both horizontal and vertical views.

The efficiency of this algorithm for constructing electron tracks was determined by a hand scan of electrons from stopping muons to be $98 \pm 2\%$. The track reconstruction efficiency predicted by Monte Carlo for these events is 96.3%.

3.6 Stopping Muon Identification

As mentioned several times previously, stopping muons constitute a major E645 background. The reasons are twofold. First, their energy spectrum closely resembles the one which would be expected from inverse beta decay, being only a few MeV lower in energy (see figure 3.1). Second, there are many of them. The stopping muon rate in the detector is on the order of 50 Hz, or 3×10^5 per LAMPF day. The active shield prevents most of these from triggering, but the veto is only 10 μ sec long. The stopping muon trigger rate is about .2 Hz, or 1000 LD^{-1} . Most of these events can be removed from the sample in a straightforward manner, as will be described

Monte Carlo Electron Energies

Figure 3.1: Energy spectra of stopping muon electrons and inverse beta decay electrons, as determined by Monte Carlo.

in the remainder of this section.

3.6.1 First Pass Removal of Stopping Muons

The first pass removes a substantial fraction of the stopping muons, taking the easily identifiable, or “pristine”, stopping muons out of the data and writing them to a separate file. Because electrons from muon decay bear a close resemblance to the signal we are trying to detect, they are ideal for such things as a particle identification reference and to check efficiencies of various cuts.

A pristine stopping muon is defined by the following properties:

- A track can be constructed in the triggering (in-time) time slice.
- A track can be constructed in a past time slice within 30 μ sec of the trigger.
- One of the endpoints of the past-slice track is within 50 cm. of one of the endpoints of the in-time track. The other endpoint of the past-slice track is outside the detector.
- The supposed muon track points to within 50 cm. of the centroid of a cluster of hits in the active shield when extended backward.

The tracking, end-point finding, and shield hit algorithms are described in other sections. Those events satisfying these requirements are written out to a file separate from event candidates. Of the events that are written,

virtually 100% are stopping muons, since there is no other background which could fake such a signature. As can be deduced from Table 3.1, about 240,000 such events were extracted during the 1987 run. Figure 3.2 shows the time between the muon track and the electron track for a sample of these events. The measured muon lifetime is $2.05 \pm 0.06 \mu\text{sec}$. The predicted lifetime in carbon is $2.11 \mu\text{sec}$

3.6.2 Second Pass Stopping Muon Cut

The first pass stopping muon cut is designed more to select stopping muons than to eliminate them. Many muons exist in the sample which do not meet the tracking criteria for this cut, but still must be removed. An additional cut in the second pass takes this process one step further by applying a less stringent muon track-finding algorithm, which is able to identify the tracks made by steep muons which are not found in the first pass.

This algorithm uses the endpoint coordinates of the "in-time" (triggering) track and the positions of all hits occurring off of the track. It then goes into the past and searches for detector hits and shield hits which occur within $1 \mu\text{sec}$ of each other. Once these are found, it tries to draw connecting lines between track endpoints or hits which occur off of the track, all past detector hits, and the hits in the active shield. For single hits, three-dimensional position information is not available so tracks are constructed in both X and Y projections. Each track is drawn from the endpoint or off-track hit to all possible past hits which satisfy the condition

Figure 3.2: Time difference between the muon track and the electron track for a sample of stopping muons selected by the first pass.

of having time-correlated shield hits. Tracks are drawn to all hits which occur contemporaneously. Once each track is constructed in the detector, it is projected onto the shield. If it falls within 50 cm. of shield hits, and the summed pulse height of all shield hits at that time exceeds 10 digitizer channels (less than 10% of the energy deposited by a through-going muon), then the event is labeled a stopping muon.

As seen in Table 3.2, this cut is very effective at eliminating stopping muons. Its efficiency was checked by writing out a file of endpoint and associated (off-track) hit information from a stopping muon file. This was then read in and processed with random trigger events. The random events provide a measure of shield and detector noise. 93.7% of the events passed this cut.

Unfortunately, the stopping muons are not so easy to eliminate. So far, both cuts have required detector-shield coincidences. Events exist, however, in which either shield or detector information is lacking or is insufficient to allow the event to be rejected. Cuts must be made on the shield alone, or on the detector alone. These will be described in sections 3.9 and 3.10.

3.7 Fiducial Volume Requirements

Muons and other charged particles entering the detector from the outside leave a track whose endpoints will lie near the edge of the detector. To verify that neutrino candidates are actually generated within the detector, rather

than enter from outside, it is necessary to impose a fiducial volume cut. The justification for this can be seen in figure 3.7, which shows the spatial distribution of track endpoints before a fiducial volume cut is applied. The figures show that many of the tracks do seem to be generated outside of the detector.

The fiducial volume cut requires that both of the track endpoints occur a certain distance in from the edge of the detector. The track endpoints are computed by the tracking algorithm. The positions of the edges of the detector, measured in a detector survey, are contained in the same constant file which contains the calibration constants. A logical variable is set for each endpoint, whose value indicates whether or not that endpoint falls within a set distance from the edge.

A fiducial volume cut is applied in each pass, each imposing a more severe cut than the previous one. The first pass cut is set to remove muons with endpoints outside the specified edge of the detector. Few are removed in this fashion. The most effective fiducial volume cut is applied at the beginning of the second pass. Here, the distances inward defining the fiducial volume are 15 cm. in the X direction (south-north), 20 cm. in the Y direction (up-down), and 20 cm. in the Z direction (east-west; the beam axis). This cut removes a considerable fraction of the events. The third pass imposes a cut along the +Z axis 30 cm. from the edge. This is necessary because of the weak spot in the shield at the junction of the blue wall

Track Endpoints After Muon Cut

Figure 3.3: End point distributions after the first pass, in X, Y, and Z. The cuts are placed at the following distances from the edge X: ± 15 cm., Y: ± 20 cm., Z: $+30$ cm., -20 cm.

and the cylinder.

To calculate the fraction of randomly distributed events remaining after the fiducial volume cut a Monte Carlo is used. When electrons with the energy spectrum expected from inverse beta decay are generated uniformly throughout the detector, the fraction predicted to fall within the fiducial volume is 69.3%.

3.8 Particle ID and Energy Cuts

Once the stopping muon generated electrons have been removed from the sample and the fiducial volume cut has eliminated hits near the edge of the detector, another background becomes apparent. Short but highly ionizing tracks become a major component of the sample. These are protons, given their energy through np elastic scattering. The high energy neutrons are a component of the cosmic ray flux (see section 2.3.3). To remove these events, a particle identification algorithm which utilizes a confidence level calculation is applied.

The events remaining after this cut tend to be more distributed in space than do electrons from stopping muon decay. In addition, the confidence level in scintillators is skewed towards low confidence levels. The interpretation given to this information is that these electron-like events are generated by photons which either produce e^+e^- pairs or Compton scatter. The photons themselves can be created by neutron inelastic scattering on

nuclei, either in the detector or in the shield. This class of background can be removed by cutting on the amount of energy deposited off of the track.

The final cut in the analysis requires total energy to be less than 60 MeV. At this point in the analysis, all backgrounds with past hits in the detector or the active shield have been removed from the sample. This helps to remove misidentified protons and some of the remaining photon-induced events.

3.8.1 Electron Confidence Level

In order to ascertain the confidence level that a particle is an electron, it is necessary to compare it in some manner to a standard set of electrons. The quantity compared is measured dE/dx per module, for both scintillators and PDTs. The standard set of electrons is provided by the stopping muon file created in the first pass.

The following method is used: The electron dE/dx , corrected for angle, is histogrammed for a subset (of about 2000 events) of the stopping muon data set (see figure 3.4, which shows scintillator and PDT dE/dx for a larger data sample). The PDTs must be corrected for short term gain drifts, so the second pass uses a PDT calibration file which contains a gain adjustment for each run. The subset of events used is a randomly selected sampling of 1/100 stopping muon events processed. The area under the histogram is normalized to 1 so that it represents a dE/dx probability distribution. This distribution is read in at the beginning of the second pass analysis.

Figure 3.4: dE/dx distributions for stopping muon electrons in scintillators and PDTs. The X axis is a dE/dx in units of "minimum ionizing particles". The distributions have been normalized so that the peaks fall at 1 MIP.

A particle undergoing identification should have dE/dx measurements in both scintillator and PDTs. If it does, each hit is compared to the standard histogram, and a differential probability is assigned to that hit. These probabilities are multiplied together to yield an unnormalised probability for both scintillators and PDTs. To normalize these probability distributions, a Monte Carlo is used. One thousand events are generated for each event analyzed. The Monte Carlo uses the probability distribution obtained from the stopping muon histogram to create an "event" having the same number of PDTs and scintillators hit as the actual event. The differential probabilities for each PDT or scintillator hit in the Monte Carlo events are multiplied together just as those for the events are. This provides a probability measurement which can be directly compared to the one generated for the event. The fraction of Monte Carlo events with probability less than that of the measured event defines the "confidence level" that the track is that of an electron. Figure 3.5, displays a confidence level distribution for both PDTs and scintillators using stopping muon electrons. As can be seen, the distribution is flat in both scintillator and PDT confidence levels.

When the particle ID is applied to the data sample, the result is quite different, as can be seen in figure 3.6. Here one can see an accumulation of events in the region of very low PDT and scintillator confidence levels. These are protons. A sample of data having both PDT and scintillator confidence levels $< 5\%$, is written to a "proton" output file. Of course, a

Figure 3.5: Confidence level distribution for stopping muon electrons. The X axis shows scintillator confidence level. The Y axis shows PDT confidence level.

Figure 3.6: Confidence level distribution for the data set after stopping muon cuts.

low confidence level can be caused by either hits that are highly ionizing or hits that are below minimum ionizing. The energy distribution of these events indicate that the former is the case, that these events deposit large amounts of energy (figure 3.7).

The cut on the data to remove the proton like events is applied in two stages, once again for ease in computing. The second pass requires the product of the scintillator and PDT confidence levels to be greater than .01. The third pass imposes the requirement that both scintillator and PDT confidence levels be greater than 12%.

A number of other cuts are associated with the particle ID. The first is the requirement that there be something energy deposited in both PDTs and scintillators. Certain counters have been defined as "bad", with low or undefined calibration constants. These are not used in the particle ID. It is possible, therefore, to have events where no usable PDTs or no usable scintillators. These events are rejected. Other hits which are discarded are above 4 MIP (1 MIP is the energy deposited by a normally incident muon passing through the center of a counter.) in energy (figure 3.4). Hits of very large pulse height are rejected because the pulse height distributions have low statistics in the high energy region. If more than one such hit occurs in a given track, however, the event is cut in the third pass. The efficiencies of all cuts are measured with a stopping muon sample. The results are listed in Table 3.4.

Proton Sample Energy Spectrum

Figure 3.7: Energy spectrum of events with PDT and scintillator confidence levels $< 5\%$.

3.8.2 Energy cuts

Upon removal of proton contamination, another class of background emerges which causes a scewing of the scintillator confidence level, although not the PDT confidence level (figure 3.8). Closer examination of these events reveals that many of their hits are separated in space from the track, and that energy deposition along the track is not uniform. It is believed that these events are the result of $n\gamma$ reactions taking place on the nuclei in the detector. High energy cosmic ray neutrons can interact on carbon or other nuclei to produce photons which in turn can create electrons through pair production or Compton scattering. There are several characteristics which allow such events to be identified. First, the photon does not generally deposit all of its energy in one place. Much energy is deposited in extra hits scattered throughout the detector. Also, if an interaction occurs on a nucleus, it can emit considerable energy, thus generating a hit with very large pulse height. Lastly, the energy of these events is determined by the energy of the incoming neutron, which can be very large. The neutrino candidates, on the other hand, have an upper limit on their event energy of 52.8 MeV. These characteristics allow cuts to be imposed which eliminate this type of background.

Figure 3.8: Scintillator confidence level for final sample with (left) and without (right) the removal of events with large associated energy.

End-Point Energy

An $n\gamma$ event is generated when a high-energy neutron interacts on a nucleus. This can create an electron-like track. It often happens that the counter at the end of the track is the one in which the interaction occurred. If this is the case, extra energy due to fragmentation or proton emission from the nucleus can be observed. The end counters, however are often excluded from the confidence level calculation, which requires that all hits used not be at the end-point of the track. Rather than ignore the information from these counters, the analysis makes an energy cut on them, which is applied immediately after the confidence level calculation. Figure 3.9 shows the energy distribution in the end counters for Stopping muons and for the data after the confidence level calculation. As can be seen, there is a tail extending to large energies in the data sample due to the process described above. A cut is placed at 1.5 times minimum ionizing which is 92.4% efficient when applied to stopping-muon electrons.

Associated Energy

Figure 3.10 shows a plot of associated (off-track) energy versus energy on the track for stopping-muon electrons. The the sum of the two is kinematically constrained to be less than 52.8 MeV. Comparison of this plot with the same plot of the data after the confidence level cut (figure 3.11) reveals that a number of events in the data sample have very large

Figure 3.9: Energy deposited in end scintillation counters for the data (after particle ID) and for Michel electrons.

Figure 3.10: Energy deposited off of the track versus energy deposited on the track for a Michel electron control sample.

Figure 3.11: Energy deposited off of the track versus energy deposited on the track for the data before the associated energy cut.

associated energies. A cut is applied at 12 MeV in off-track energy to eliminate the anomalous events. The effect of this cut can be seen by examining figure 3.8, which shows the scintillator confidence level of the final event sample without an associated energy cut versus the confidence level of the final sample with the cut. The efficiency of this cut is 97.9%, determined by the Michel electron control sample.

Total Energy

After removal of the neutron induced events, the preponderance of remaining events are electrons from stopping muons. These are removed by cuts in the shield and detector which will be described in the next sections. The remaining energy spectrum is shown in figure 3.12. The cut-off for oscillation (inverse beta decay) events is 52.8 MeV. A cut is placed at 60 MeV to remove the high energy events. Many of these appear to be misidentified protons, due to their relatively short track length. The excess of beam-on events above 60 MeV is believed to be due to decay-in-flight ν_μ interactions (section 3.10). The efficiency of this cut, determined by applying it to stopping-muon electrons, is 99.7%.

3.9 Shield Cuts

The first cut that utilizes the shield alone is applied in the first pass. This is a cut which removes all events in which two or more spatially correlated

Figure 3.12: Energy spectrum of the final sample, before the cut on total energy.

shield photomultiplier tubes fire within one microsecond of the trigger time. The purpose of this cut is to remove through-going muons from the sample. Its efficiency is 95%.

The events remaining after the confidence level and associated energy cuts are predominantly low-energy electrons with short tracks. These occur at a rate of about 4.8 LD^{-1} , and show no correlation with the beam gate. Clearly, a cosmic-ray induced background remains. Evidence that this background is related to cosmic ray muons is found in figure 3.13, which shows the timing distribution of shield hits occurring before the trigger. A large number of shield hits are seen at about $10 \mu\text{sec}$ before the trigger, which is when the long veto turns on. Only muons can cause a delayed trigger in the detector. Either the muons are entering the detector unobserved, or they are stopping outside the detector and the decay electron is producing a hard photon which converts in the detector and triggers. In either case, the only information which can be used to eliminate these events comes from the active shield.

Several different methods are used to quantify the amount of energy deposited in the active shield. One is to take the maximum pulse height in the shield. While this is a relatively poor measure of overall energy deposition, it is useful in special cases in which a muon passes through a shield PMT. Because the hemispherical surfaces of these PMTs project into the liquid scintillator, the active shielding is thinnest in the region of the

PMTs. This leads to events where very few, or only one shield PMT fires, with the one PMT which the muon passed through exhibiting a very large pulse height. Another method utilizes multiplicity in a shield "cluster", which is defined as a set of shield tubes which fire at the same time and are spatially correlated. All shield PMTs which have nonzero pulse heights and are within 1.75 meters of other firing PMTs are grouped into clusters. Clusters which are within 3.5 meters of each other are combined. Another method entails summing the pulse height within the shield clusters. Cuts are applied to all of these variables. They were studied in order to determine the cuts which yield the best separation between random noise and actual muon hits in the shield.

It was found that the most sensitive measure of energy was the sum of the pulse heights in each shield cluster. When this cut is optimized, it removes virtually all events which would have been removed by the other algorithms. Hence the final version of the shield cut utilizes only this algorithm. The time at which the hits occur is also taken into account. Because the muon lifetime is only $2.1 \mu\text{sec}$, most of the events which are muon induced will have the time of the muon near the time at which the long veto turns off, $10 \mu\text{sec}$ before the trigger. Optimization leads to a cut-off of the shield software veto at $30 \mu\text{sec}$. All of these decisions were made by comparing the data sample to a control sample of random-trigger data (figure 3.14). The cut was set at a pulse sum of 35. Typical muons have

a pulse sum of 200 to 400. Cut efficiency is 75.9 %, determined by passing the random data through the shield cuts.

3.10 Decay-In-Flight Neutrino Candidates

At this stage in the analysis, the sample consists of electron-like events with have passed the shield cut. Examination of these events reveals that many still seem to have muons spatially correlated with and occurring before the electron. These muon-induced events fall into two classes. In one class it is clear that the muon originated outside of the detector. These are cosmic ray muons which did not fire enough PMTs in the shield to cause the event to be rejected. In the other class of muon-induced events, the muon is contained within the fiducial volume. These events are generated by decay-in-flight (DIF) neutrinos: ν_μ or $\bar{\nu}_\mu$ created by the decay of high energy π^+ or π^- .

These two event classes can be separated by the fact that a cosmic-ray muon must enter from the outside of the detector. The neutrino induced muons will be spread evenly throughout the detector, with a $1/R^2$ dependence where R is the distance from the beam-stop. The algorithm begins by screening all past hits and selecting those which meet the following criteria:

- The past hits occur within 20 μ sec of the trigger.
- At least one past hit occurs within 35 cm. of one of the track endpoints.

Figure 3.13: Time difference between trigger time and the time of shield hits which meet the veto condition (pulse sum > 35) for data and random trigger.

Figure 3.14: The maximum summed pulse height in the shield within 30 μsec of the trigger, for data and random trigger.

- If these conditions are satisfied, then at the time specified by the past hit, there must be at least one scintillator with a pulse height in channel 4 (about .8 MeV energy deposition) or higher, or 2 PDTs with any energy deposition.

After these events are found, the past hits are checked to ascertain whether all hits at the specified time occur within the fiducial volume. If they do not, it is assumed that the event is cosmic ray induced, and it is discarded. Otherwise the event is classed as a decay-in-flight event. Decay-in-flight events are potential backgrounds. If the muon is created and decays without passing through sufficient live material, the electron produced could be mistaken for oscillation signal. This background will be discussed at greater length in the next chapter. If the muon causes the trigger, or if the muon decays quickly so that its energy is included with that of the electron, it will create an event with large energy. It is believed that this is the source of the beam excess events above 60 MeV (figure 3.12). Most DIF candidates with muons in the past occur during the beam gate, with 15 events occurring in the beam-on gate (75.48 LD live time) and 7 events occurring in the pre-beam and post-beam gates (265.8 LD live time). This yields a beam excess of 13.0 ± 3.9 events, and a rate of $.17 \pm .05$ ev/LD.

Table 3.4: Summary of acceptance and efficiency of all cuts.

Cut or Condition	Events	Acceptance(%)	Efficiency(%)
Events Read	1269639		
Detector and Shield	1267944	99.9	99.9
Fiducial Pulses			
Random Triggers Removed	1150719	90.8	100.0
In-Time Shield Cut	789844	68.6	98.0
Track Found	469362	59.4	96.3
Fiducial Volume Cut I	443119	94.4	100.0
Stopping Muons Removed	204518	46.2	100.0
Fiducial Volume Cut II	111788	54.7	69.3
Stopping Muon Cut	21979	19.7	93.7
Degree of Freedom Cut	20033	91.2	99.2
Confidence Level I	3984	19.9	94.7
End Counter Cut	3018	75.8	92.4
Fiducial Volume Cut III	2873	95.2	
1 Scint. on Track	2752	89.5	99.5
Confidence Level II	2066	75.1	83.9
2 or More Hits Discd.	1987	96.2	99.0
Associated Energy	1739	87.5	97.9
Shield Cuts	147	8.5	75.9
Past Hit Near Track	99	67.4	98.0
Decay-In-Flight Removed	77	77.8	
Total Energy	47	61.0	99.0

3.11 The Final Event Sample

After all cuts have been applied a small (47 event) sample of events remains. 20 events occur in the beam gate (74.5 LD live time) and 27 events in the beam-off gates (265.8 LD live time). This gives a beam excess of 12.3 ± 4.9 events, and a rate of $.16 \pm .07$ events/LD. Results from the 1986 run, which had a high beam-neutron rate, compared with the very low beam-neutron rate of the 1987 run, lead to the conclusion that neutrons from the beam-stop cannot account for the final event sample, as shown in section 2.3.3. Decay-in-flight events cannot either, as will be discussed in some detail in the next chapter. These events are believed to be neutrino reactions in the detector, mostly ν_e on ^{12}C and electron elastic scattering. The rest of this chapter briefly presents the properties of the final data set.

3.11.1 Energy Distribution

Figure 3.15 shows the energy distribution of the final data set broken into beam-on, beam-off, and cosmic-ray background subtracted samples. The negative going point in the 30 MeV bin is a statistical fluctuation, caused by a large number of beam-off events in that bin with no corresponding events occurring in the beam gate. The salient feature of the background subtracted plot is that the energy of the beam correlated neutrino candidates is lower than that which would be expected from either inverse beta decay or stopping muon decays. As will be shown in the next chapter, this

is very similar to the energy spectrum from the reaction $\nu_e \text{ }^{12}\text{C} \rightarrow e^+ \text{ }^{12}\text{N}$.

3.11.2 Position Distribution

Neutrino events should be distributed with at $1/R^2$ dependence, where R is the distance to the beam-stop. Figure 3.16 shows the distribution of the final candidate sample along the Z (east-west) axis. This is shown for those events occurring in the beam gate, in the beam-off gate, and the beam gate with background subtraction. An excess is seen near the front ($-Z$) of the detector for the beam-correlated data and near the rear ($+Z$) of the detector for the cosmic-ray data. Two effects explain the cosmic-ray data. First, the passive shielding is somewhat thinner near the rear of the detector, so more backgrounds can be expected in this region. Second, the active shield has a weak spot at the junction of the cylinder and the blue (rear) wall which can allow charged particles to penetrate without vetoing. In either case, this effect can be removed through subtraction of the cosmic ray background.

The beam excess seems to be concentrated in the front half of the detector, as would be expected for a $1/R^2$ distribution. The lack of events near the front of the detector is caused by a low efficiency in plane 6, which had a scintillator leak. The the rear half of the detector seems to have fewer events than would be expected from a $1/R^2$ distribution. This deficit is believed to be the result of a statistical fluctuation — an excess of cosmic-ray induced events falling in the rear half of the detector during the beam-off

Figure 3.15: Energy distribution of the final neutrino-induced event candidates: beam on, beam off, and subtracted.

gate but not during the beam-on gate. There are no other systematic effects which could account for this distribution.

The X and Y (south-north and vertical) distributions of the final event sample are shown in figures 3.17 and 3.18. Neither of these is as sensitive to the $1/R^2$ dependence as the Z distribution. No correlation with the direction of the beam-stop (top and north) is apparent.

3.11.3 Angular Distribution

The angular distribution of neutrino events, both from backgrounds and from signal, is expected to be essentially isotropic. The exception to this is neutrino-electron scattering, which is very strongly peaked in the forward direction ($\cos\theta \geq .96$). Trigger efficiency for a particle of given energy and range, however, is a strong function of angle. Acceptance is much larger for particles traveling along the detector axis. This is evident in figure 3.19, which shows the cosine of track angle with respect to the detector Z axis for the final event sample. Figure 3.20 shows the angular distribution of the final sample with respect to the line between the forward vertex of the event and the beam-stop. This is shown for beam-on, beam-off, and cosmic-ray subtracted data. Forward peaking occurs because the beam axis differs by only 17° from the detector axis. There appears to be no evidence for any strongly forward peaked component of the beam-on sample.

Figure 3.16: Z (Tunnel Axis) Position distribution of the final neutrino-induced event candidates: beam-on, beam-off, and subtracted

Figure 3.17: X (South-North) Position distribution of the final neutrino-induced event candidates: beam-on, beam-off, and subtracted

Figure 3.18: Y (Vertical) Position distribution of the final neutrino-induced event candidates: beam-on, beam-off, and subtracted

Track Cosine With Respect to Detector Axis Final Sample

Figure 3.19: Angular distribution of the final event sample with respect to the detector Z axis.

Track Cosine with Respect to the Beam Stop Final Sample

Figure 3.20: Angular distribution of the final neutrino-induced event sample with respect to the beam-stop: beam-on, beam-off, and subtracted

Chapter IV

Neutrino Oscillation Results

The last chapter described how a final sample of neutrino-like events was obtained from the 1987 data run. This sample showed a definite excess of events occurring in correlation with the beam. The conclusion reached is that these events are neutrino induced. The main question that must be answered at this point is whether these events represent a signal from neutrino oscillation or whether they are the result of other neutrino processes. The first segment of this chapter deals with the various backgrounds which are correlated with the beam and their respective contributions to the E645 data sample. The latter section of this chapter presents the conclusions reached by this experiment, based on an analysis which takes into account the relevant backgrounds.

4.1 Expected Beam-Related Backgrounds and Their Properties

Of all the particles produced within the beam-stop, only neutrinos and neutrons can mimic a neutrino oscillation signal. The beam-stop and its immediate shielding effectively prevent any charged particles from escaping. Neutrons are strongly attenuated by the shielding between the beam-stop and the detector. During the 1986 run they contributed substantially to the background levels, as described in section 2.3.3. Extra shielding eliminated all beam-related proton reactions in the detector. As a result, beam neutrons are no longer expected to contribute to E645 backgrounds.

Neutrinos of the types ν_e , ν_μ , and $\bar{\nu}_\mu$ can interact in the detector and mimic the oscillation signal ($\bar{\nu}_e$). These are produced in copious numbers by the decay of pions and muons in the beam-stop. This section describes the predicted magnitude of neutrino backgrounds and how well the data sample fits this prediction.

4.1.1 The EGS Monte Carlo

In order to determine the acceptance, the tracking efficiency, and the fiducial volume of the detector for both signal and background, it is necessary to model the detector using Monte Carlo techniques. A sophisticated simulation is necessary to model electrons, since they lose a considerable portion of their energy through radiation. This radiation can convert back into

electrons via the photoelectric effect, Compton scattering, and pair production. The standard electron Monte Carlo used for high energy electrons is EGS4[59], developed at SLAC. Its range of applicability spans from 1 MeV to the TeV region.

EGS transports electrons or photons and any secondary particles by generating them randomly in a specified region, and gives them a random direction. The detector modeling is done prior to the execution of the EGS program. The detector is broken down into 232 regions defined by 225 planes. Some simplifications are involved. These regions can be grouped into plane units consisting of four subregions. Each plane unit contains a rectangular slab of liquid scintillator, a rectangular slab of P-10 gas, a rectangular slab of air, and a dead material combination consisting of the acrylic scintillator walls, the cellulose PDT walls, and the aluminum layer on the inner surface of the PDTs. Determination of the scintillator tank number and the PDT number in each plane is done externally to the EGS routine. Including each scintillator and PDT, or breaking the dead material composite into its components would have created an immense number of regions. It was determined that the number of regions used provided precision comparable to the more complex geometry. One exception to the uniform slab model was imposed for the scintillator. Dead space, consisting of acrylic and air, was included between the individual scintillator modules in a given plane.

To calculate spectra using EGS, a flat energy spectrum between 0 and 60 MeV was used. The events thus generated are weighted by a probability distribution which represents the desired energy spectrum. Thus any energy spectrum which has an endpoint energy of less than 60 MeV can be modeled using the same Monte Carlo electron file.

Muons had to be generated for decay-in-flight neutrino studies and for cosmic ray studies. This is done by a separate program. This did not need to be nearly as elaborate as EGS. Muons lose energy almost solely through ionization. One necessary correction is the generation of delta-rays: high energy electrons from $\mu - e$ scattering. If one of these particles was produced, the EGS routine was called to handle it. The same procedure was followed if the muon stopped and decayed.

The Monte Carlo has several functions.

- It predicts the spectrum for all ν_e induced backgrounds. This, in conjunction with the overall normalization, allows backgrounds to be incorporated into the likelihood calculation of the final oscillation limit.
- It can give an estimate of the level of decay-in-flight (DIF) ν_μ backgrounds based on the number of DIF electrons detected in the event sample.
- It is used to calculate the spectrum of inverse beta-decay electrons which in turn is used to set oscillation limits.

- It is used to determine the normalization of observed energy to deposited energy.

The remainder of this section gives the result of these calculations.

4.1.2 Backgrounds from Neutrino-Nucleus Scattering

The flux of electron neutrinos through the detector equals the flux of $\bar{\nu}_\mu$. One might wonder, then, why the interactions of electron neutrinos on nuclei do not create a signal which would completely swamp any but a near maximal oscillation signal. A general answer, with details differing from nucleus to nucleus, is that the ν_e must interact only on neutrons within the nuclei. This interaction transforms these neutrons into protons and creates an electron which can mimic the positron from inverse beta-decay. The final state nucleus must have an available state for the new proton. In nuclei where $N = Z$, especially those with “magic numbers” of neutrons or protons (2, 8, 20, 50, 162) the energy levels of available states tend to be much higher than those in which there are one or more “quasi-free neutrons”, or $N > Z$. The reactions with higher excitation energies tend to have smaller cross-sections due to the smaller available phase space. For nuclei with quasi-free neutrons the reaction cross-sections can be written

$$\sigma = \frac{M_F^2 + M_{GT}^2}{\pi} G_F^2 p^2 \quad (4.1)$$

where M_F is the Fermi transition matrix element, M_{GT} is the Gamow-Teller matrix element, G_F is the Fermi constant, and p is the neutrino momentum. The “sum-rule” states that

$$M_F^2 \approx N - Z \quad (4.2)$$

$$M_{GT}^2 \approx 3(N - Z)$$

so the larger the neutron excess, the larger the neutrino cross section. Two of the nuclei which are of concern to E645, ^{12}C and ^{16}O , have $N = Z$, while two others, ^{13}C and ^{27}Al , have $N - Z = 1$. The $N = Z$ nuclei, while abundant in the detector, have very small cross sections. The $N - Z = 1$ nuclei make up a very small fraction of the mass of the detector, but have relatively large cross section and so contribute 6.0% of the total background.

The calculations of the cross sections used for various nuclei will be shown below. One addition is the background due to neutrino-electron elastic scattering, which is not a nuclear reaction but can be predicted accurately and contributes a substantial fraction of the expected background rates.

Carbon 12 and Oxygen 16

Electron neutrinos interact with ^{12}C and ^{16}O via the reactions

$$\nu_e + {}^{12}\text{C} \rightarrow e^- + {}^{12}\text{N}$$

$$\nu_e + {}^{16}\text{O} \rightarrow e^- + {}^{16}\text{F}$$

The ^{12}C and ^{16}O cross sections have been calculated by Donnelly [65]. He uses an analogy between the electromagnetic and semi-leptonic weak interactions (justified by the conserved-vector-current hypothesis, later shown to be a consequence of the Weinberg-Salam electroweak model) which allows the use of electron scattering and gamma decay data to calculate the neutrino cross sections. These estimates show that for ^{12}C the transition that contributes the most to the cross section is between the carbon ground state ($J^\pi T = 0^+0$) and the nitrogen ground state ($J^\pi T = 1^+1$), about 80% of the total cross section. Oxygen, however, is a nucleus of magic number, so there are fewer states of small Q value available. Hence the cross section is smaller. A number of oxygen states contribute to the cross section, leading to its somewhat more complicated structure (figure 4.1).

These cross sections were parameterized[60] to yield the following spectrum from ^{12}C :

$$\frac{d\sigma_{\nu e}^{12\text{C}}}{dE_e d\theta_\nu} = \kappa \frac{d\sigma}{d\theta_\nu} \epsilon (12\epsilon^2 - 12\epsilon^3) \quad (4.3)$$

$$\frac{d\sigma}{d\theta_\nu} = 1.13 - .5 \cos \theta_\nu \quad (4.4)$$

$$\kappa = 4.7 \times 10^{-44} \text{cm}^2 \text{MeV}^{-1} \text{strd}^{-1}$$

$$\epsilon = E_e / E_{\text{max}}$$

$$E_{\text{max}} = 35.0 \text{MeV}$$

For ^{16}O , the spectrum is not nearly as smooth in variation with E or θ as the ^{12}C spectrum, due to the number of nuclear states which contribute

Figure 4.1: Differential cross sections for expected E645 neutrino backgrounds, integrated over the beam-stop neutrino spectra.

substantially to the cross section:

$$\frac{d\sigma_{\nu_e \text{ } ^{16}\text{O}}}{dE_e d\theta_\nu} = \kappa' \frac{d\sigma}{d\theta_\nu} \frac{d\sigma}{dE_e} \quad (4.5)$$

$$\begin{aligned} \frac{d\sigma}{dE_e} = & 4.91 - .019E' - .051E'^2 + .00051E'^3 \\ & + .00020E'^4 - .0000015E'^5 - .00000015E'^6 \end{aligned} \quad (4.6)$$

$$\begin{aligned} \frac{d\sigma}{d\theta_\nu} = & .468 - .468\cos\theta_\nu + .096\cos^2\theta_\nu - \\ & .048\cos^3\theta_\nu \end{aligned} \quad (4.7)$$

$$\kappa' = 10^{-44} \text{cm}^2/\text{MeV}/\text{strd}$$

$$E' = E_\nu - 22\text{MeV}$$

$$5\text{MeV} \leq E_\nu \leq 38\text{MeV}$$

Carbon 13 and Aluminum 27

The interaction of neutrinos with ^{13}C and ^{27}Al takes the same form as seen in the previous section:

$$\nu_e + ^{13}\text{C} \rightarrow e^- + ^{13}\text{N}$$

$$\nu_e + ^{27}\text{Al} \rightarrow e^- + ^{27}\text{Si}$$

The approach to calculating the ^{13}C and ^{27}Al cross sections involves a comparison to the ^{12}C calculation. Specifically, this means that the assumption was made that the bulk of the cross-section was due to the transition to one final state. This yields a spectral shape that is the same as that for ^{12}C , with only the end-point energies (E_{max}) differing between

these reactions. For both of these reaction, the final state is assumed to be the ground state of the product nucleus. The magnitude of the cross section can be calculated from the beta-decay matrix elements:

$$\sigma = \frac{10^{-40} \text{MeV}^{-2} \text{cm}^2 \text{sec}}{ft_{1/2}} S(E_\nu - Q)^2 \quad (4.8)$$

$$S = \frac{2J_f + 1}{2J_i + 1} \quad (4.9)$$

The Q values are 4.8 MeV for ^{27}Al and 2.2 MeV for ^{13}C . For both nuclei, $J_f^\pi = J_i^\pi = 1/2$. The result of this calculation is

$$^{13}\text{C} : \sigma = 2.0 \times 10^{-40}$$

$$^{27}\text{Al} : \sigma = 2.2 \times 10^{-40}$$

Unfortunately, the calculation of the total cross section for these nuclei is not so straightforward. Both nuclei have excited states which contribute to the size of the overall cross section. Judging from the relative strengths of transitions to these states versus the ground states observed in electron and proton scattering, a rough estimate can be made of their contribution. This leads to a cross section of about 4×10^{-40} for both of these nuclei. The uncertainties are relatively large when compared to those of ^{12}C . The parameterized differential cross sections, integrated over flux, are shown in figure 4.1. Total cross sections are shown in table 4.1.

Table 4.1: Total neutrino cross sections for E645 signal and background processes integrated over neutrino flux and angle.

Reaction	Cross Section ($\times 10^{-42} \text{cm}^2$)
$\bar{\nu}_e \text{p}$	102.0
$\nu_e {}^{12}\text{C}$	14.6
$\nu_e {}^{13}\text{C}$	40.0
$\nu_e {}^{27}\text{Al}$	40.0
$\nu_e {}^{16}\text{O}$	5.2
νe	4.0

4.1.3 $\nu - e$ Elastic Scattering

Another oscillation-mimicing background is caused by neutrino-electron elastic scattering. The magnitude of this background can be calculated from the Weinberg-Salam theory of weak interactions. All neutrinos produced in the beam-stop ($\nu_e, \nu_\mu, \bar{\nu}_\mu$) contribute to the cross section, which has been explicitly calculated for pion decays at rest by Kayser et. al.[62]. The flux-averaged differential cross section has been parameterized [60] in the following form:

$$\begin{aligned}
 \frac{dP}{dE_e}(10^{-45} \text{cm}^2/\text{MeV}) &= .778 - .597\epsilon + .721\epsilon^2 - 2.409\epsilon^3 + \\
 &\quad 1.604\epsilon^4 \quad (\epsilon \leq \epsilon_0) \quad (4.10) \\
 &= .661 - .423\epsilon + .560\epsilon^2 - 2.409\epsilon^3 + \\
 &\quad 1.604\epsilon^4 \quad (\epsilon \geq \epsilon_0)
 \end{aligned}$$

where $\epsilon = 2E_e/M_\mu$ and $\epsilon_0 = .564M_\mu$, the muon neutrino energy. This equation represents the contributions from all neutrino flavors. The integrated

cross section for each neutrino type, and the total, is given in table 4.2. The energy distribution is shown in figure 4.1. All of these cross sections were

Table 4.2: Cross section for neutrino-electron scattering for ν_e , ν_{μ} , and $\bar{\nu}_{\mu}$, integrated over angle and flux.

Neutrino Type	ν e Cross Section ($\times 10^{-45} \text{cm}^2$)
ν_e	300.5
ν_{μ}	46.7
$\bar{\nu}_{\mu}$	49.2
Total	396.4

calculated using the world average for the Weinberg angle; $\sin^2\theta_W = .23$ [63]. The electron neutrino has a larger cross section than the muon neutrino and antineutrino because it can interact via a charged exchange current as well as a neutral current. The muon neutrino and antineutrino are constrained to interact only via neutral current, as shown in figure 4.2. This interaction has been measured by a LAMPF experiment[64]. The statistical and systematic error on the result is on the order of 50%, whereas electroweak theory calculates the cross section to within 5%, with the uncertainty in the Weinberg angle being the limiting factor. Hence we use the calculated rather than the measured result.

One notable property of neutrino-electron elastic scattering is its strong peaking in the forward direction. For neutrino energies greater than 10 MeV, the largest scattering angle is 16° . Since multiple scattering of the electron in three scintillators can be up to 10° , Monte Carlo simulation assumes a flat distribution in $\cos\theta$ for $\cos\theta \geq .96$, which drops to 0 for

Figure 4.2: a). The Feynman diagrams for a) $\nu_e e$ scattering b) $\nu_\mu e$ scattering c) $\bar{\nu}_\mu e$ scattering.

$$\cos\theta < .96.$$

4.1.4 Decay-in-Flight Neutrinos

“Decay-in-Flight” neutrinos are neutrinos created when pions decay before coming to rest. If the pion has an energy greater than 50 MeV, it is possible for the muon neutrino or antineutrino to have an energy greater than the muon production threshold. Muons produced by neutrino interactions in the detector can create electron-like background if the muon is created and decays without passing into live material. Prediction of the decay-in-flight background level presents problems. First, the energy spectrum and flux of the neutrinos is a strong function of the pion production, absorption, and reaction in the beam stop. These cross sections are not known to a high degree of precision. Secondly, the beam-stop geometry must be modeled with all empty spaces (potential π decay regions) taken into account. Due to the removal and insertion of the isotope production stringers (see figure 2.4), the beam-stop geometry was not constant. Lastly, the cross section for muon production increases very steeply in the energy region near the production threshold. These factors make it very difficult to calculate the muon production rate and energy spectrum accurately.

Another approach to estimating the decay-in-flight (DIF) rate uses a Monte Carlo of the DIF events normalized to the size of the observed DIF sample to estimate the number of DIF events expected. This still is sensitive to the shape of the neutrino energy spectrum, but removes the uncertainty

in the overall normalization. Hence, rates can be calculated for various reasonable DIF neutrino spectra. These spectra are parameterized as a function of muon kinetic energy in the following form[60]:

$$\frac{dN}{dT_\mu} = \sigma_0 T_\mu^\alpha e^{-bT_\mu^2} \quad (4.11)$$

where

$$\sigma_0 = \frac{2b^{\frac{\alpha+1}{2}}}{\Gamma(\frac{\alpha+1}{2})}$$

$$b = \frac{\alpha}{2T_{max}^2}$$

$$\alpha = 3$$

Three cases were studied; $T_{max} = 20, 60$, and 100 MeV. The shape of the energy distribution was used to weight a muon Monte Carlo. The Monte Carlo events were then processed to see if they triggered the detector within the fiducial volume, whether the muon decay time was within detector time resolution, and whether the muon was identifiable via a large energy deposition. The results are shown in table 4.3.

As can be seen from this table, the expected number of DIF background events varies between about 1 and 4, depending on the energy spectrum assumed. Our measured result finds ≤ 1 event in the region above 30 MeV, where one expects to find the majority of Michel electrons. Due to the uncertainties in the calculation, the DIFs are treated as irreducible background, and no attempt is made to subtract them from the signal.

Table 4.3: The number of DIF events without an observable muon in the E645 data sample

Cut or Condition	$T_{max} = 20$	$T_{max} = 60$	$T_{max} = 100$
Events Processed	4908.7	5005.3	4837.1
Trigger, Fid. Vol., Tracking	1322.8	1688.1	1274.8
Electron Triggers	1083.1	508.0	245.0
Muon Triggers	209.6	1072.7	926.9
"Invisible Muon"	211.0	29.6	8.81
Tagged DIF Events	630.8	243.0	101.4
Ratio: Tagged/Invisible	2.99	8.21	11.5
Expected Background for 13.0 $\pm$ 3.9 Observed DIF events	4.35 $\pm$ 1.30	1.58 $\pm$ .48	1.13 $\pm$ .34

4.1.5 Energy Resolution

The EGS Monte Carlo is also used to determine the energy resolution by comparing the spectrum of observed Michel electrons with the spectrum predicted by the Monte Carlo. This is shown in figure 4.3. The Monte Carlo energy was normalized with respect to the observed energy so that the peaks were in alignment. After this adjustment, the energy distributions are very similar.

Once agreement exists between data and Monte Carlo, the Monte Carlo can be used to determine the mapping between actual particle energies and energies observed in the detector, as shown in figure 4.4. This shows the observed energy predicted by EGS for monoenergetic electrons generated in the detector. Two salient features are evident. One is that the energy resolution, measured as full width at half maximum (FWHM), is on the

order of 20-30%. The FWHM increases as a function of energy due to the increased probability of bremsstrahlung at higher energies. The second feature is that the peak of the measured energy tends to be lower than the actual energy by about 10%.

Oscillation probability is a function of the particle energy. Therefore, to calculate oscillation limits requires a means of incorporating energy offsets and smearing. This is accomplished by constructing a matrix which maps real neutrino energy into measured neutrino energy. To generate this matrix, a flat electron energy spectrum is thrown by EGS. The proper kinematics are imposed on the positrons, and the events are weighted by the $\bar{\nu}_\mu$ energy spectrum. A simulated trigger condition is then imposed, as well as tracking and fiducial volume requirements. Events which satisfy these criteria have their energies calculated, and the weighted value for the event is entered into a two dimensional matrix. One dimension represents the actual neutrino energy, and the other the energy measured by the detector.

To use this matrix, a vector is calculated corresponding to the probability of oscillation for a given Δm^2 and $\sin^2 2\theta$. This is then multiplied by the matrix, and by the maximum rate for full oscillation. The resulting vector is the predicted signal in the detector, broken into energy bins, for the value of Δm^2 and $\sin^2 2\theta$ chosen. This vector can then be used to calculate a likelihood or χ^2 , as described in section 4.3.1.

Total Energy, Monte Carlo Michel Electrons and Stopping Muon Electrons

Figure 4.3: Energy spectrum for Michel electrons (data points) from stopping muons compared to the spectrum predicted by the EGS Monte Carlo (solid line).

E-645 Detector Response to Electrons (MC Uniform Angular Distribution)

Figure 4.4: EGS-predicted observed energy spectra for nine different electron energies.

4.2 Rates and Normalization

The rate of any neutrino background is directly proportional to the neutrino flux, the cross section for the given background, the trigger efficiency, and the number of targets upon which the neutrino can interact within the fiducial volume. How these quantities are obtained will be explained in the following sections.

4.2.1 Detector Composition

The number and type of nuclei in the detector is one factor determining the overall neutrino background rate and the potential level of oscillation signal for a given oscillation probability. An estimation [61] of the number of each type of nucleus in the detector can be found in tables 4.4 and 4.5. Hydrogen and carbon are the main components of the detector. The number of electrons is also included in these calculations, since they contribute to the background via neutrino-electron elastic scattering.

4.2.2 $\bar{\nu}_e p$ Cross Section

The derivations of the neutrino cross sections for our energy spectra are shown in some detail in the previous section. These are listed in table 4.1 and shown as a function of energy in figure 4.1. The one cross section calculation excluded up to this point has been that of $\bar{\nu}_e p$, inverse beta decay. As this determines the rate that would be seen due to an oscillation

Table 4.4: Materials comprising the E645 detector, their empirical formulas, and their densities.

Material	Formula	Atomic Weight	Density (g/cm ²)
Liquid Scintillator (30% 1,3,5 Trimethyl Benzene 70% Light Mineral Oil)	CH ₂	14.03	.858
Lucite (Methyl Methacrylate)	C ₅ H ₈ O ₂	100.13	1.16
Mylar (Polyethylene Terephthalate)	C ₁₀ H ₈ O ₄	192.17	1.39
Kraft Paper (Cellulose)	C ₆ H ₁₀ O ₅	162.14	.90
Aluminum	Al	26.92	2.70
Gadolinium Oxide	Gd ₂ O ₃	362.50	7.41

signal, it is important that it be estimated as precisely as possible.

The cross section for inverse beta decay has been calculated by O'Connell [67]. Our estimate adapts the notation of Fayans [68,66] to yield the following differential cross section:

Table 4.5: Number of atoms and electrons in the E645 detector

Material	e ×10 ³⁰	H ×10 ³⁰	C ×10 ²⁹	O ×10 ²⁹	Al ×10 ²⁷	Gd ×10 ²⁶
Scintillator	4.45	1.12	5.55			
Lucite	1.46	0.22	1.36	0.54		
Mylar	0.17	0.01	0.17	0.07		
Paper	0.77	0.09	0.53	0.44		
Aluminum	0.08				6.53	
Gadolinium Oxide	0.02			0.00		2.05

$$\begin{aligned}
\frac{d\sigma}{d\Omega}(\bar{\nu}_e p \rightarrow e^+ n) &= \frac{1}{2\pi} \frac{G^2}{\pi} E_e^2 \times \frac{1}{1 + 2(E_\nu/M_p) \sin^2 \frac{\theta}{2}} \times \quad (4.12) \\
&\quad \left[(g_V^2 + g_A^2) \cos^2 \frac{\theta}{2} + 2g_A^2 \sin^2 \frac{\theta}{2} + 2g_A g_M \alpha_r \frac{q}{M_p} \sin \frac{\theta}{2} \right] \\
g_V^0 &= 1 \\
g_A^0 &= -1.261 \\
g_M^0 &= 3.706 \\
\alpha_r &= .93 \\
\frac{G^2}{\pi} &= \frac{\sigma_0}{g_V^{0^2} + g_A^{0^2}} = 1.658 \times 10^{-44} \text{ cm}^2 \\
\sigma_0 &= \frac{2\pi^2 (\hbar c)^3}{m_e^5 c f \tau_n} = 9.565 \times 10^{-44} \text{ cm}^2 / \text{MeV}^2 \\
f &= 1.6857 \\
\tau_n &= 900 \text{ sec}
\end{aligned}$$

The total cross section is obtained by integrating this quantity with respect to Ω and averaging it over the $\bar{\nu}_\mu$ flux:

$$\sigma_t(\bar{\nu}_e p \rightarrow e^+ n) = 2\pi \int_{M_e}^{M_\mu/2} \int_{-1}^1 \frac{d\sigma}{d\Omega} \Phi(E_\nu) dE_\nu d(\cos\theta) \quad (4.13)$$

$$\Phi(E_\nu) = \frac{3E_\nu^2}{8M_\mu^4} \left(\frac{1}{2} M_\mu - \frac{2}{3} E_\nu \right) \quad (4.14)$$

Parameterizing these equations and integrating numerically gives

$$\sigma_t(\bar{\nu}_e p \rightarrow e^+ n) = 1.02 \times 10^{-40} \text{ cm}^2$$

The differential cross sections as a function of energy for this and all potential backgrounds are shown in figure 4.1.

4.2.3 Neutrino Flux

The neutrino flux is determined by two factors, the distance between the beam-stop and the detector, and the neutrino production rate in the beam stop. The proton flux incident on the beam-stop is known with a high degree of precision. The largest unknown is the pion production rate. This has been the object of intensive study.

The experiment of Cochran et. al.[69], measured the pion production cross sections at $E_p = 730$ MeV, for pion energies from 30 to 553 MeV and angles from 15° to 150° , and for a number of different targets. This, and the experiment of Crawford et. al.[70], with $E_p = 585$ MeV, constitute the only relatively complete set of pion production cross sections available. Based on these results and a Monte Carlo analysis of the LAMPF beam-stop, the pion per incident proton ration is estimated to be .089[64].

To obtain a better estimate of the pion to proton ratio, an experiment, E866[71], was proposed and carried out at LAMPF. This experiment used a copper-scintillator sandwich construction to estimate the spatial distribution of pion decays in the beam-stop. Copper was used because this is the material of which the beam-stop is constructed. The data from this experiment are currently being analyzed[72,73].

Distances from the beam-stop are determined by survey. Benchmarks are set along the floor of the tunnel, and the distances between these and the beam-stop have been precisely estimated. The detector was surveyed with

respect to these benchmarks. The distance from the center of the neutrino production region to the center of the cylinder has been determined to be $26.96^{+0.04}$ m, or 26.80 m if one weights for the $1/R^2$ dependence of the neutrino flux.

4.2.4 Trigger and Cut Efficiencies

The previous chapter dealt with the data analysis, and described the various cuts in some detail. It gave the efficiency of each cut, and how this efficiency is estimated. These efficiencies are compiled in tabular form in table 4.6.

Table 4.6: Compilation of efficiencies for all cuts performed by the data analysis and how these are obtained.

Cut/Requirement	Efficiency (%)	Checked by:
Random Trigger Removal	100.0	—
In-Time Shield Cut	98.0	Random Triggers
Track Reconstruction	96.6	Monte Carlo
Stopping Muon Cut I	~ 100.0	—
Fiducial Volume	69.3	Monte Carlo
Stopping Muon Cut II	93.7	Random Trigger (Shield) + Stopping Muon (Track)
Particle ID	65.3	Stopping Muons
Off-Track Energy Cut	96.9	Stopping Muons
Shield/Detector Veto	75.9	Random Trigger (Shield) Protons (Detector)
Decay-in-Flight Cut	99.0	Protons
Total Energy Cut	99.0	Stopping Muons
Total Analysis Efficiency	28.9	

The overall analysis efficiency was found to be 28.9% for inverse beta decay positrons. Particle energy has a role in determining the trigger

efficiency, and to a lesser extent the tracking and fiducial volume efficiencies. These were estimated for inverse beta decay and all background processes using the EGS4 Monte Carlo, as described in section 4.1.1. The results are shown in table 4.7.

Table 4.7: Trigger efficiency, tracking efficiency and fiducial volume predicted by the EGS Monte Carlo for backgrounds and for inverse beta decay.

Reaction	Trigger Efficiency (%)	Tracking Efficiency (%)	Fiducial Volume	Total
$\bar{\nu}_e p$	33.3	96.6	69.3	22.3
$\nu_e {}^{12}\text{C}$	12.1	94.2	75.7	8.5
ν_e	21.2	98.8	73.7	14.9
$\nu_e {}^{27}\text{Al}$	19.5	96.2	73.0	13.8
$\nu_e {}^{13}\text{C}$	26.3	96.0	70.2	17.7
$\nu_e {}^{16}\text{O}$	8.6	94.3	75.1	6.1

4.2.5 Expected Background Rates and the Observed Excess

Using the above information, rates were calculated for all backgrounds. ${}^{12}\text{C}$ is the largest expected contributor to the background, followed by ν_e scattering. The overall expected rates are shown in table 4.8, as well as the observed beam excess and the size of a signal expected for a full neutrino mixing.

The feature that allows the separation of the backgrounds from a possible oscillation signal is the energy dependence of these backgrounds. As mentioned in the sections on background, most of them tend to have low

Table 4.8: Observed beam excess and predicted background rates for the 1987 data set

Observed Beam Excess	12.3 ± 3.8
Maximum Oscillation	472.31
Signal, $\bar{\nu}_e$ p	
ν_e ^{12}C	14.30
ν_e	6.24
ν_e ^{27}Al	.49
ν_e ^{13}C	.87
ν_e ^{16}O	.62

end-point energies, whereas the inverse beta-decay positrons will tend to have larger energies. This is displayed in figures 4.5, 4.6, and in table 4.9. Figure 4.5 shows the spectra of all expected backgrounds individually compared to a maximal oscillation signal. ^{12}C , which is the largest contributor to the backgrounds, can be seen not to contribute substantially in the region of oscillation sensitivity. In figure 4.6, the combined background rates are shown versus the observed beam excess, and versus an oscillation signal which would be seen if 5% of the $\bar{\nu}_\mu$ oscillate into $\bar{\nu}_e$. The observed beam excess seems to be fit fairly well by the predicted background levels, although the overall rate is to be low. This can be accounted for by uncertainties in the cross section and by other systematic errors, as will be explained in the next section. Table 4.9 divides the beam-on, the beam-off, the predicted background, and a 5% oscillation signal into 12 5 MeV bins. It is this information that is used to obtain neutrino oscillation limits.

Expected Event Rates for 5100 C on Target

Figure 4.5: Energy spectra of the predicted E645 backgrounds, and of a maximal oscillation signal

Beam Excess, Predicted Neutrino Background,
and 5% Oscillation Signal

Figure 4.6: Energy spectra of the combined backgrounds (solid line), the observed beam excess for the 1987 data run (points), and a 5% oscillation signal (dashed line).

Table 4.9: Data & Monte Carlo in 5 MeV Bins

MeV	0-5	5-10	10-15	15-20	20-25	25-30	30-35	35-40	40-45	45-50	50-55	55-60
Beam-on	0.0	0.0	1.0	5.0	5.0	4.0	0.0	1.0	2.0	0.0	1.0	1.0
Beam-off	0.0	1.0	1.0	3.0	2.0	3.0	10.0	2.0	1.0	0.0	4.0	0.0
Excess	0.0	-.28	.72	4.15	4.44	3.15	-2.82	.44	1.72	0.0	-.13	1.0
^{12}C	0.0	.03	.40	2.73	5.52	4.42	1.17	.04	0.0	0.0	0.0	0.0
νe	0.0	.01	.20	1.04	1.81	1.48	.97	.51	.19	.04	0.0	0.0
^{13}C	0.0	0.0	0.0	.04	.12	.19	.25	.18	.08	.02	0.0	0.0
^{16}O	0.0	0.0	0.0	.07	.23	.20	.11	.02	0.0	0.0	0.0	0.0
^{27}Al	0.0	0.0	0.0	.02	.07	.11	.15	.10	.04	.01	0.0	0.0
Total	0.0	.04	.61	3.89	7.75	6.39	2.65	.83	.33	.06	0.0	0.0
Excess - Bkgds	0.0	-.32	.11	.26	-3.31	-3.23	-5.47	-.40	1.40	-.06	-.13	1.0
1% Osc.	0.0	0.0	.01	.09	.32	.76	1.09	1.12	.91	.35	.06	0.0

4.2.6 Systematic Errors

The largest contributions to systematic error on this experiment come from the estimations of the neutrino background cross sections. The error for these quantities are listed in table 4.10. The ^{12}C cross section is known the most precisely, since 80% or so of the strength of the reaction goes to the ground state of ^{12}N , which is well known from beta decay. The other backgrounds have large contributions to the cross sections due to transitions to excited states which are not nearly as well measured. The estimates in the table come from Donnelly[74].

Large contributions to the systematic error also come from uncertainties in the efficiencies of the cuts performed in the data analysis, as shown in table 4.11. The largest of these uncertainties is in the efficiency of the particle ID. This is due to energy dependence, particularly of the “end energy”

Table 4.10: Estimated systematic error in the neutrino cross sections used.

Reaction	Systematic Error (%)
$\bar{\nu}_e p$	3
$\nu_e {}^{12}\text{C}$	20
$\nu_e e$	2
$\nu_e {}^{13}\text{C}$	30
$\nu_e {}^{27}\text{Al}$	30-50
$\nu_e {}^{16}\text{O}$	30

cut. Other sources of error are shown in the second part of the table. The largest uncertainty in any quantity is that of beam stop flux. This dominates all other backgrounds. When the errors are added in quadrature, the total estimated systematic uncertainty is 15.9%.

Table 4.11: Estimated uncertainty of the efficiencies of the cuts performed in the data analysis

Cut	Systematic Uncertainty (%)	Estimated by:
Particle ID	5	Energy dependence of cut Gain variations with time
Fiducial Volume	2	Geometry dependence
Shield	1	Statistics on Random sample
Associated Energy	1	Energy dependence of cut
Total Energy	.5	Energy dependence of cut
Number of Atoms	.5	Measurement Error
Trigger Efficiency	5	Variation of Threshold
Neutrino Flux	14	Estimated uncertainty in π^+/p
Total	15.9	

4.3 Neutrino Oscillation Limits

The following conclusions can be reached concerning the final data sample.

- It is possible to see neutrino reactions from background processes. The energy spectrum of the observed beam excess is very similar to that of the predicted backgrounds.
- Most of the background events are predicted to originate from the reaction $\nu_e {}^{12}\text{C} \rightarrow e^- {}^{12}\text{N}$. The second most prevalent background should be νe elastic scattering. Study of the angular distribution of the beam excess events reveals no evidence for the νe process.
- The observed level of beam excess is approximately 55% of that predicted for background processes. One standard deviation, determined by taking the uncertainty in the cross sections for background processes and the uncertainty in the overall normalization in quadrature is equal to 26%. Hence, the observed signal is small by 1.8 standard deviations.
- There is no observed beam excess above 35 MeV, which is the region where an oscillation signal would be observed. Hence, there is no evidence for neutrino oscillation to within the limits of experimental sensitivity.

Given that no oscillation signal is observed, it is possible to set limits on Δm^2 and $\sin^2 2\theta$. The standard procedure in previous oscillation ex-

periments is to set 90% confidence level limits in $\Delta m^2 \otimes \sin^2 2\theta$ space. The most precise method of doing this utilizes a maximum likelihood technique to compute the 90% confidence level.

Maximum likelihood techniques are generally used in particle physics, and provide the most powerful statistical means of treating data [75,76]. A good likelihood fit to the data should meet a number of criteria:

- The answer obtained should not depend on the algorithm used.
- It should be able to deal with the negativity of our overall signal.
- It should take all systematic errors into account in a proper manner.

These points will be covered in the following sections.

4.3.1 The Maximum-Likelihood Method

The likelihood function is defined as the joint probability density of obtaining a given set of measurements for a given hypothesis assuming a normalized distribution function[76]:

$$\mathcal{L}(\alpha) = \prod_{i=1, N} f(\alpha; x_i) \quad (4.15)$$

Here, $\mathcal{L}$ is the likelihood function, α is a continuous set of hypotheses, $x_i, i = 1, N$ represents a set of measurements, and $f(\alpha; x_i)$ is the normalized probability distribution representing the probability of observing x_i given

that hypothesis α is true. The maximum likelihood solution for α is found by maximizing the log of $\mathcal{L}$ with respect to α :

$$\frac{d(\ln \mathcal{L}(\alpha))}{d\alpha} \big|_{\alpha=\alpha^*} = 0 \quad (4.16)$$

α^* is the value of α which maximizes the likelihood function. In our case the probability distributions for the data are described by Poisson statistics:

$$f(\alpha; x) = \frac{\alpha^x}{x!} e^{-\alpha} \quad (4.17)$$

and the probability distributions for the systematic errors are described by Gaussian statistics:

$$f(\alpha; x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\alpha)^2}{2\sigma^2}} \quad (4.18)$$

The likelihood function can now be constructed:

$$\mathcal{L} = \prod_{i=1}^{\mathcal{N}} \frac{(N_i^{on})^{n_i^{on}} e^{-N_i^{on}}}{n_i^{on}!} \frac{(N_i^{off})^{n_i^{off}} e^{-N_i^{off}}}{n_i^{off}!} e^{\frac{x^2}{2\sigma_i^2}} e^{\frac{t^2}{2\sigma_t^2}} e^{\frac{u^2}{2\sigma_u^2}} \quad (4.19)$$

where

$N_i^{on} \equiv$ predicted number of beam-on events in energy bin i .

$n_i^{on} \equiv$ measured number of beam-on events in energy bin i .

$N_i^{off} \equiv$ predicted number of beam-off events in energy bin i .

$n_i^{off} \equiv$ measured number of beam-off events in energy bin i .

$\mathcal{N} \equiv$ number of bins in the energy histogram.

$\alpha \equiv$ beam-on/beam-off normalization = .282

$C_i \equiv$ calculated number of ^{12}C events.

$s \equiv$ the fraction by which this cross section is incorrect.

$\sigma_s \equiv$ the systematic error on this cross section.

$t \equiv$ the fraction by which overall normalization is incorrect.

$\sigma_t \equiv$ the systematic error on overall normalization.

$D_i \equiv$ calculated number of ^{13}C , ^{27}Al , and ^{16}O events.

$u \equiv$ the fraction by which these cross sections are incorrect.

$\sigma_u \equiv$ the systematic error on these cross sections.

$B_i^{e,l} \equiv$ calculated number of νe events.

$T_i(\delta m^2, \theta) \equiv$ predicted oscillation signal.

These variables satisfy the condition

$$N_i^{on} - \alpha N_i^{off} = (1 + t) (B_i^{e,l} + (1 + s)C_i + (1 + u)D_i + T_i(\delta m^2, \theta)) \quad (4.20)$$

Certain quantities are unknown in our estimates. The “theoretical” beam off rate and the value of the actual systematic error (as opposed the *uncertainty* in the systematic error) are not known. To remove these from the likelihood function, these quantities are integrated over their range of possible values. The errors for ^{13}C , ^{27}Al , and ^{16}O are combined into one integration. Otherwise, it would be necessary to integrate over each of them independently, making a numerical integration prohibitively expensive in computer time. Neutrino-electron scattering is not integrated over. The systematic error in this quantity compared to the others is small enough

that it can be assumed to be a constant. The resulting function to be evaluated is:

$$\mathcal{L} = \int_0^\infty \int_0^\infty \int_0^\infty \prod_{i=1}^{\mathcal{N}} \int_0^\infty \frac{(N_i^{on})^{n_i^{on}} e^{-N_i^{on}}}{n_i^{on}!} \frac{(N_i^{off})^{n_i^{off}} e^{-N_i^{off}}}{n_i^{off}!} dN_i^{off} \quad (4.21)$$

$$e^{\frac{s^2}{2\sigma_s^2}} ds e^{\frac{t^2}{2\sigma_t^2}} dt e^{\frac{u^2}{2\sigma_u^2}} du$$

To find the point of maximum likelihood, a two-dimensional grid search is done over Δm^2 and $\sin^2 2\theta$ space. It turns out that this value would fall in the unphysical region (Δm^2 or $\sin^2 2\theta < 0$), but because the fit is constrained to be within the physical region it converges to 0 for both variables. Thus our value of maximum likelihood is given by $\mathcal{L}(\Delta m^2 = 0, \sin^2 2\theta = 0)$. This corresponds to a negative oscillation result.

Once the value of maximum likelihood is determined, the 90% confidence level can be defined as a fraction of the maximum likelihood, or as a difference in the log likelihoods ($\Delta\mathcal{L} \equiv \log\mathcal{L}_{max} - \log\mathcal{L}_{90\%}$). The procedure that specifies this fraction consists of a Monte Carlo which throws beam-on events, beam-off events, and systematic errors according to a procedure explained in the following section. The fraction is entered as a parameter in our computation of the 90% confidence level curve. To calculate the position of this curve, Monte Carlo events are thrown over Δm^2 and $\sin^2 2\theta$ space in a random but logarithmically weighted manner. A two dimensional grid search is performed to find the nearest point which has a value of likelihood corresponding to the the 90% confidence level. This point is then recorded. Enough events are thrown to define the curve.

Likelihood Renormalization

When data lie totally in the physical region the $\Delta\chi^2$ corresponding to 90% confidence level can be looked up in tables. If the data lie partially outside the physical region choosing the curve for a given confidence level is not quite so straightforward. For instance if a Gaussian peaks at zero and one is measuring a positive definite quantity, the 90% confidence level given by a table corresponds to the 80% confidence level of the Gaussian because the answer is constrained to be at one end. There are basically two standard treatments of this renormalization problem.

One method is favored by the Particle Data Group[77]. Here one sets the limit on number of oscillations events corresponding to 90% by renormalizing the confidence level to the physical region. It is easily shown that this method is very conservative in that the final confidence interval is larger than 90%. For instance, in the example given above, the 90% confidence level would include 95% of the events under the Gaussian curve, as opposed to 90%. Using this renormalization technique the 90% confidence levels at large mixing and at large Δm^2 are given by:

$$\begin{aligned}\Delta m^2 &= .135, \sin^2 2\theta \sim 1 \\ \sin^2 2\theta &= .024, \Delta m^2 \sim 100eV^2\end{aligned}$$

An alternate method of renormalizing the probability is to throw Gedanken experiments about one's experimental data. To do this one generates Monte Carlo experiments consistent with the measured beam-off and beam-on

data. The generator assumes a Poisson distribution, with the measured value in each bin as the median. A problem occurs for bins with no entries, as is common in the beam-on data (see table 4.8). It is not clear whether a 0 in a histogram bin is there because of a statistical fluctuation or whether the probability of having an event fall into that bin is particularly low. Indeed, *in the absence of a priori information concerning the mean of the distribution in a bin containing zero events, it is impossible to correctly generate events in that bin!* Luckily, *a priori* information does exist in the form of beam-off data. It is proper to assume that the data in the beam-off gate represent the cosmic ray background much more precisely than the beam-on data, since 3.5 times as much beam-off as beam-on data was accumulated. If the data in the beam gate in a given bin fell below what would be expected from the cosmic ray (beam-off) gate, the observed value in that bin during the cosmic ray gate was used as the median around which the Poisson distribution was thrown. The beam-off data were averaged over 15 MeV bins to eliminate zeros and reduce statistical fluctuations.

Systematic errors are also thrown about their central values. The generating function for these is a Gaussian distribution. Once all free parameters have been generated by the Monte Carlo, the resulting Gedanken experiment is subjected to the maximum likelihood fit. The best value of $\sin^2 2\theta$ and Δm^2 are evaluated for each experiment. One then chooses a confidence interval which contains 90% of the Gedanken experiments. This

procedure is equivalent to putting all Gedanken experiments in the unphysical region at zero Δm^2 . The Particle Data Group's method is equivalent to ignoring all events in the unphysical region.

This technique is correct, but it is difficult to separate fit quality from experimental sensitivity (e.g. The present data exclude the physical region at 62% confidence level). Using this technique the 90% confidence levels for large Δm^2 and $\sin^2 2\theta = 1$ are:

$$\begin{aligned}\Delta m^2 &= .11, \sin^2 2\theta \sim 1 \\ \sin^2 2\theta &= .014, \Delta m^2 \sim 100eV^2\end{aligned}$$

The curve representing the 90% confidence level is shown in figure 4.7.

4.3.2 Likelihood vs. χ^2

The minimum χ^2 technique is not well suited to this type of analysis because it relies on Gaussian statistics, which are not followed due to our limited sample size. Nevertheless, a parallel analysis was run which utilized this method. The reasoning behind this was to provide a cross-check on the likelihood fitter and to have a program which would generate results very quickly (because it avoids the multiple integrations of the maximum likelihood). It also gives a goodness-of-fit estimate that is readily understandable.

The implementation of this procedure was identical to that of the maximum likelihood, with the exception that the calculated quantity that was

Figure 4.7: 90% Confidence level limits on $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations set by E645.

minimized in $\Delta m^2 \otimes \sin^2 2\theta$ space was

$$\chi^2 = \sum_{i=1}^{\mathcal{N}} \sum_{j=1}^{\mathcal{N}} \left(X_i^{meas} - T_i(\delta m^2, \theta) \right) A_{ij}^{-1} \left(X_j^{meas} - T_j(\delta m^2, \theta) \right) \quad (4.22)$$

where

$$X_i^{meas} = n_i^{on} - \alpha n_i^{off} - D_i - C_i - B_i^{el}$$

A_{ij} is the covariance matrix, which takes into account the statistical and systematic errors. Statistical errors occur only on the diagonal, while the off diagonal elements are calculated using the systematic errors. The matrix is inverted and used in the above calculation. This provides a number equivalent to (but two times larger than) the likelihood. The best fit ($\chi^2 = 18.7$ for 12 degrees of freedom) corresponds to negative values of Δm^2 and $\sin^2 2\theta$, just as in the likelihood fit. The 90% confidence level curves are obtained for $\Delta\chi^2 = 1.76$:

$$\begin{aligned} \Delta m^2 &= .075, \text{ for } \sin^2 2\theta \sim 1 \\ \sin^2 2\theta &= .0075, \text{ for } \Delta m^2 \sim 100 eV^2 \end{aligned}$$

This gives a considerably better answer than the likelihood. To test whether this is legitimate, the likelihood routine was altered so that it used Gaussian rather than Poisson statistics. The result of this test was nearly identical to that of the χ^2 calculation. One can conclude that the difference between the results is due to the statistics used, and that since Poisson statistics are proper in this case, the likelihood gives the more accurate answer.

4.3.3 Dependence of Answer on Systematic Errors

The sensitivity of the binned likelihood fit to systematic errors has been examined carefully. The limits are not very sensitive to the size of these errors. This can be demonstrated by setting all systematic errors to zero. The large mass and mixing limits then are calculated to be

$$\Delta m^2 = .087 \text{ eV}^2$$

$$\sin^2 2\theta = .0104$$

The error having the largest bearing on the final answer is the error in overall normalization. It was found that the limits vary linearly with this error. The limit is much less sensitive to variation of the error in the cross section.

Another integration loop was added to determine the sensitivity of the result to the energy dependence of the cross sections. The result was indistinguishable from the result without that particular systematic added.

4.3.4 Majorana Oscillation Limits

E645 can also set limits on Majorana type oscillations, $\nu_e \rightarrow \bar{\nu}_e$. The positrons from inverse beta decay assuming Majorana oscillations have a lower energy than those which would be produced from the inverse beta decay assuming $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ (Dirac) oscillations, due to the higher energy of the incident ν_μ . The Monte Carlo calculated energy spectra of the positrons from both of these reactions are shown in figure 4.8. The expected results

for 5% oscillation probability in the case of both Majorana and Dirac oscillations are compared to the expected background level. Aside from the overall shift downward in energy for the Majorana spectrum, there is also a difference in normalization between the two cases. Since the trigger is energy dependent, positrons produced by Majorana oscillations will have a lower trigger efficiency than those produced by Dirac oscillations. The number of events expected for Dirac oscillations assuming full mixing is 472. For Majorana events this number is 402, corresponding to a trigger efficiency of 27.4%.

A maximum likelihood fit was done for the Majorana oscillation hypothesis, analogous to the one done assuming Dirac oscillations. The fit in this case was much worse however. The normalization procedure used in the Dirac case excludes the physical region at 95% confidence level when applied to the Majorana case. The only recourse under these conditions is to apply the Particle Data Group normalization[77], which ignores all events in the unphysical region. The resulting limits in Δm^2 and $\sin^2 2\theta$, respectively, are .14 eV² and .030. The 90% confidence level exclusion curve is shown in figure 4.9.

Figure 4.8: Expected signal from Dirac and Majorana Oscillation assuming 5% oscillation probability, compared with expected E645 background.

Figure 4.9: 90% confidence level exclusion curve for Majorana oscillations using Particle Data Group renormalization.

4.4 Conclusion

4.4.1 Comparison to Previous Measurements

The final limits on $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations set by the E645 oscillation experiment are shown in figure 4.7. To put these limits into perspective, they are compared with the results from BNL734, BEBC, Gsngen, and SKAT in figure 4.10. E645 does not perform as well as reactor experiments at large $\sin^2 2\theta$, or as well as high energy accelerator experiments at large Δm^2 , but it seems to outperform all other experiments in the $\Delta m^2 = .1 \text{ eV}^2$, $\sin^2 2\theta = .01$ region.

4.4.2 The Future of E645

E645 has accumulated data through the 1988 cycles, cycles 51 and 52. In addition, there are plans to continue the experiment through 1989. Cursory examination of the 1988 data reveals that the observed neutrino rates are consistent with those obtained during cycles 49 and 50. The limits will not improve linearly with running time. The $\sin^2 2\theta$ limit improves as the square root of running time, whereas the Δm^2 limit improves only as a fourth root of running time (the Δm^2 term is inside a squared *sin* function). One reaches the point of diminishing returns rather quickly. It is clear that new strategies are needed to push oscillation limits considerably lower than where they currently sit.

Figure 4.10: 90% Confidence level limits on $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations set by E645 and other experiments.

4.4.3 Future Neutrino Oscillation Experiments

The experimental conditions which determine the sensitivity of an experiment to neutrino oscillation were described in chapter 1. To briefly reiterate, the conditions favorable to neutrino oscillation experiments are:

- High neutrino flux.
- Low cosmic ray background rate.
- Large detector volume.
- An optimized energy, which balances the L/E dependence of the Δm^2 limit (better at low energies) with the energy dependence of the neutrino cross section, which rises with energy.
- An optimized source-to-detector distance, which balances the L/E dependence of the Δm^2 limit (better at long distances) with the $1/R^2$ dependence of neutrino flux.

No new high flux accelerators are currently scheduled for construction. Hence, we will explore only the possibilities for physics with currently existing accelerators and those under construction. There are two problems with large detectors. First, they are very expensive. Second, the rates from cosmic ray backgrounds increase with detector volume. The cosmic ray problem can be offset with a short accelerator duty cycle.

The potential for substantial improvement in oscillation limits using high energy accelerators is meager. Although the cross section increases with energy, the distance between the detector and source must become very large in order to measure Δm^2 in a competitive manner. Suggestions have been made to use baselines of up to 1000 km.[21]. One current proposal for a two component detector to be constructed at Brookhaven National Laboratory (BNL 846)[78], claims that it can set limits at the level of 10^{-3} in $\sin^2 2\theta$ and .05 eV². in Δm^2 . These estimates assume a factor of 20 increase in the neutrino flux at the Brookhaven neutrino channel, and two detectors: a 110 ton detector at 150 m. and one a 220 ton detector at 1000 m.. If these limits were to be attained, they would be marginal improvements on limits that already exist. We must ask if there is any means by which order of magnitude improvements in current limits can be obtained using currently available acceleration and detection technologies. One such proposal exists. It would utilize the Large Čerenkov Detector which is proposed for construction at LAMPF[79].

The proposed experiment would be constructed at the proton storage ring (PSR) at LAMPF, if approved. Its stated purpose is to do a 1% measurement of $\sin^2 \theta_W$ by doing a precise measurement of νe scattering. It consists of a large tank of water, 9500 metric tons in all, with photomultipliers around the periphery to observe the Čerenkov light given off by relativistic particles. The large fiducial mass of this detector could create

severe cosmic ray background problems. The PSR, however, would run at 12 Hz, with 270 nsec. beam spills, so the live fraction is 3×10^{-6} . This effectively eliminates the cosmic ray background. Its planned current is 100 μ A, 1/10 of the present LAMPF intensity. The beam-stop is located directly under the experiment itself. One might expect that the L/E of such an arrangement would limit the sensitivity in Δm^2 . There are two factors which overcome this problem. First, the detector is very large. Its distance from the beam-stops varies between 8 and 28 meters. Second, the expected neutrino rate is so large that it is able to improve the Δm^2 limits through statistics alone. The estimated oscillation sensitivity for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ is 4×10^{-4} in $\sin^2 2\theta$ and $<.01 \text{ eV}^2$ in Δm^2 .

The Δm^2 limits could be improved substantially by constructing a secondary beam-stop at a larger distance from the detector, at the cost of weaker $\sin^2 2\theta$ limits. The disadvantage of this arrangement is that the proposed oscillation experiment can run parasitically with the νe scattering experiment. A secondary beam-stop experiment would have to be run at the expense of νe scattering. If, on the other hand, analysis of the LCD data reveals an oscillation signal, then the construction of a secondary beam-stop would unambiguously specify Δm^2 and $\sin^2 2\theta$, as well as provide a check on some systematics.

4.4.4 Summary

Presently, the best neutrino oscillation limits are set by a number of different experiments. Each of these has a different region of maximum sensitivity which is determined by the design parameters of both the detector and the neutrino source. Experiments at high energy accelerators exploit a large neutrino cross section to set the best limits in $\sin^2 2\theta$. Reactor experiments utilize a high neutrino flux and a large source-to-detector to neutrino energy ratio to set the best limits in Δm^2 . Intermediate energy experiments do not have the advantage of large cross section found at higher energies, or the flux of the low energy reactor experiments. Yet they have much higher flux and a larger L/E ratio than experiments performed at high energy accelerators, and they are able to do the more statistically powerful appearance experiments which are not possible at reactor facilities. Hence they are best able to set limits in the intermediate region, between very low Δm^2 and very low $\sin^2 2\theta$.

E645 has utilized the advantages of the intermediate energy meson factory environment to set limits of $\Delta m^2 = .11 \text{ eV}^2$ and $\sin^2 2\theta = .014$. In the region of $\Delta m^2 = 1 \text{ eV}^2$ and $\sin^2 2\theta = .01$, it sets limits in a region not excluded by previous experiment. Further running of E645 will improve these limits somewhat. The elements of the experiment which allow detection of neutrinos while suppressing backgrounds are a 4π anticoincidence cosmic-ray shield, fine grained construction to allow good particle identification,

and the recording of event history.

The overall mass of the E645 detector is 20 tons, mostly consisting of CH_2 . Substantial improvements can be made in oscillation sensitivity by constructing detectors of much larger mass. To counter the cosmic ray problem, a short accelerator duty cycle must be used. The proposed LCD experiment at the LAMPF proton storage ring meets both of these requirements. Order-of-magnitude improvements in current oscillation limits await the construction of this or a similar facility.

Appendix A

The E645 Triggering System

A.1 Introduction

The E645 detector is designed to select and identify electrons with energies between 10 and 50 MeV, which usually have tracks extending no more than five or six planes. It is also necessary to calibrate the detector with stopping muons which have a shallow angle of incidence, so the trigger must also be able to select longer tracks. At the time that the trigger was designed, the exact trigger configuration that would be necessary to maximize efficiency while minimizing triggers from noise or natural radioactivity was not known. Therefore, the trigger system had to be built to be as versatile as possible.

There are two systems making up E645 trigger: the “fast logic” and the “trigger processor”. The fast logic processes the raw information from the PMTs and deduces whether each plane has any PMTs with pulse heights above a set threshold. It also checks for end-to-end coincidences in scintil-

lation counters. A single hit is classed as an “OR”, and a coincidence as an “AND”. The AND and OR output from each plane is fed into the trigger processor. Patterns representing all possible trigger configurations are loaded into the trigger processor memory prior to each data run. The AND and OR lines from the fast logic address these memories. When the pattern of hits generated in the fast logic corresponds to one of the preloaded patterns, a trigger is issued by the processor if no veto from the anticoincidence shield is present. Descriptions of both the fast logic and the trigger processor systems are given in the subsequent sections.

A.2 The Fast Logic

The fast logic is a 4 MHz updating discrimination and pattern recognition system designed to process all PMT signals in the detector and determine whether single hits or end-to-end coincidences occur in the scintillation counters on a per plane basis. It utilizes ECL (emitter-coupled logic) circuitry to achieve high speed. FASTBUS cards are used to allow a larger number of channels per crate. Instrumentation for all 40 scintillator planes fits onto 20 FASTBUS cards, with one more card acting as a crate controller and interface to CAMAC. CAMAC commands are sent via an IO612 module, which is a dual 24 line read/write register.

PMT signals enter the electronics hut through 50 Ω BNC cables. These feed into the shaper modules, where the signals are split. Shaped and

amplified PMT pulses go into the digitizer system, and unamplified signals are fed into the fast logic system. Shielded 50Ω ribbon cables carry the PMT signals between the shaper and fast logic system. Each of these carries 12 signal lines, corresponding to either the north or the south side of a scintillator plane.

The front end of the fast logic circuit, which accepts the PMT signals, consists of an AC coupled amplification and discrimination circuit, as shown in figure A.1. Amplification is performed by three 10116 amplifier circuits arranged in series. The ECL discriminator is a 10130 with a threshold set by an applied voltage. There is one adjustable voltage circuit for every 12 channels, settable by a potentiometer. The discrimination level does not vary linearly with the voltage setting, as is apparent in figure A.2. The threshold control voltage was set so as to give a discrimination level of 6 mV.

The discriminators are reset every 250 nsec.. The information from the previous clock cycle is held in a latch consisting of an ECL 10135 flip-flop. This allows the discriminator configuration that existed at the time of the trigger to be read out through CAMAC. Simple logic is performed on the discriminator outputs to determine whether one of two conditions occurred. The first condition, known as an "OR", implies that one discriminator has fired anywhere in a plane. The second condition, "AND", is satisfied if any two discriminators fired that corresponded to PMTs on the opposite

Figure A.1: The fast logic amplification and discrimination circuitry for a single channel.

Figure A.2: Discriminator threshold versus threshold control voltage for one fast logic channel.

sides of the same scintillation counter. Each fast logic card processes the information for two planes, so each card has two AND and two OR outputs. These are ECL signals, fed out through LEMO connectors and cables into the trigger processor system.

Read and write functions are performed by the fast logic controller card. This communicates with CAMAC via a Jorway IO612 module. The IO612 has two 24 line read/write registers: J1 and J2. J1 is used for data transfer. This allows data to be written from CAMAC to the fast logic latches for diagnostic purposes, and allows the latches to be read either for diagnostic purposes or to obtain the triggering information. J2 is used as the control register. Patterns corresponding to read, write, and clear commands are sent via J2 to the controller card, where they are decoded and the appropriate function is performed. A counter on the controller card addresses each card sequentially when read or write functions are being executed. The address of each fast logic card is set by a DIP switch. The data is carried through the crate on a standard FASTBUS backplane, with buffered connections to the J1 register via the fast logic controller card.

A.3 The Trigger Processor

The trigger processor is a pattern recognition device which uses preprogrammed memories to search for coincidences between discriminator firings over the forty planes of the detector. It utilizes a two tier memory

system. The first level of memory is programmed for coincidences which occur over a six plane block. There are 14 of these blocks, each of which overlaps with its neighboring blocks by three planes. Output from the first level memories feed into the second level memory, which is programmed to look for correlations which occur between the six plane blocks. The trigger processor also accepts an ECL veto signal from the shield logic, and additionally requires a NIM level enable signal in order to produce a trigger. The trigger processor is housed in a modified CAMAC crate with a specially constructed backplane. There are seven first level memory modules, each containing two memories in a single width CAMAC module. The logic processor, which contains second level memory plus enabling and veto circuitry, is a double width CAMAC module. There is also a terminator card to set the proper impedances and voltage levels on the backplane. Backplane connections use FASTBUS-type pins and connectors rather than the standard CAMAC card-edge connectors. The crate also contains a power supply which provides -5 V and -2 V to the cards via bus lines on the backplane.

Fast logic AND and OR signals (described in the previous section) enter the processor via LEMO cables. As in the fast logic, ECL circuitry is used due to its fast response time. Each memory card accepts 12 AND and OR signals, corresponding to six planes. These signals are gated and latched by a D type flip-flop, then go to address the first level memories in the

manner shown in figure 2.23. This arrangement is necessary to eliminate any bias in the trigger at the edge of each memory segment. Each input line addresses two memories. The "primary" memory is addressed in such a way that each of the twelve input lines corresponds to a memory address line. The "overlap" memory on a card is addressed by the last six input lines only. These address the lower order bits of the memory. The higher order bits are addressed by the first six lines of the *next* memory module, which are carried on the backplane. Hence every input addresses two memories, except for those coming from the first three planes. Tracks which span the boundaries of two primary memories are then contiguous in the overlap memories.

The first level memories are 10470 ECL 4096 $\times$ 1 RAMs. Each is loaded individually with the selected trigger pattern. This is accomplished through the use of a counter which sequentially loads either a "1" or "0" into each memory location. Memories may be selected and read in a similar manner for diagnostic purposes. Setting read, write, or trigger modes for these memories is done in the logic processor under command from CAMAC.

The central unit of the trigger processor system is known as the logic processor. This unit communicates with a CAMAC interface module via a ribbon cable. 16 lines are used for data transfer, and seven are used for control functions. The logic processor also houses the second level memory, as well as enabling and veto logic.

Outputs from the first level memories are carried through the backplane into the logic processor unit. There they are latched until the next clock cycle, at which time they address the second level memory, a 10414 ECL 256 $\times$ 1 RAM (see figure 2.23). The second level memory allows correlation between detector segments to be required. It is loaded and read in the same manner as the first level memories. The output from this memory specifies whether or not the trigger conditions have been satisfied in the detector. At this point, other requirements may be imposed before a trigger is issued. A NIM enable is required to allow disabling of the processor while event read out is in progress. The absence of an ECL shield veto signal is also required during vetoed running. Both of these signals enter the logic processor through LEMO ports on its front panel. If all conditions are met, a trigger is issued and the 4 MHz clock is turned off. This freezes the discriminator pattern responsible for the trigger into the fast logic latches.

Control of all trigger processor functions is accomplished using CAMAC commands issued through the interface module. These include read and write commands for both memories, as well as enabling write, read and trigger modes. The standard CAMAC commands used are shown in table A.1. One of the CAMAC commands, F17A0, sets a 16 bit mask word. Each of the 16 bits corresponds to a processor command, as shown in table A.2. The first four bits control the trigger requirement: whether a positive or negative condition is required for the shield veto, and whether

the NIM enable is implemented. Other functions performed by the mask word put the memory address registers into either parallel (running) mode or increment (read or write memory) mode, certain diagnostic functions, enabling or inhibiting the trigger mode, and resetting counters and flip-flops.

Table A.1: Trigger processor CAMAC commands and their functions.

CAMAC Command	Function
F0A0	Read and Increment Memories
F1A0	Register Test (Diagnostic)
F8A0	Set Write Mode Flip-Flop
F9A0	Reset Write Mode Flip-Flop
F16A0	Write and Increment Memories
F17A0	Write Mask Word

Trigger patterns are generated by a FORTRAN program which simulates all possible fast logic outputs and determines whether a given bit pattern satisfies the desired trigger configuration. Depending on whether or not the trigger condition is satisfied, a 1 or a 0 is written to an output file in a position in the file corresponding to the memory address. Similar files are created for both first and second level memories. The files are loaded into the trigger processor prior to the start of each data run by another FORTRAN routine which executes the proper CAMAC commands. This routine also reads the processor memories and compares the results to the input file to verify that the hardware is operating correctly.

The standard running configuration during data acquisition consists of a

Table A.2: Trigger processor mask bit functions.

Bit	Name	Function
1	DETM	Enables second stage memory output to act as trigger.
2	SHIELDM	Requires shield veto for trigger.
3	SHIELD-M	Requires absence of shield veto for trigger.
4	NIM	Requires NIM enable.
5		Unused
6	ENA MASK LATCH	Required for loading first four mask bits.
7	DISABLE CNT1	Puts 1st level memory into parallel (trigger) mode.
8	ENABLE CNT1	Puts 1st level memory into increment (load/read) mode.
9	DISABLE CNT2	Puts 2nd level memory into parallel (trigger) mode.
10	ENABLE CNT2	Puts 2nd level memory into increment (load/read) mode.
11	RESET LOOP	Inhibits loop mode (diagnostic).
12	SET LOOP	Sets loop mode (diagnostic).
13	RESET INHB	Enables trigger mode.
14	SET INHB	Inhibits trigger mode.
15	RELOAD-	Resets memory addresses.
16	START	Resets event flip-flop & enters trigger mode.

“3 AND” trigger with the detector fully enabled. The first level memories are loaded with a pattern that allows three out of four consecutive ANDs anywhere in the detector to constitute a trigger. The second memory is set such that all detector segments are enabled. In addition, the shield mask is set to inhibit the trigger if a shield veto is received, and the NIM enable mask is set in order to allow all the data acquisition electronics to be in an active mode before the trigger is enabled.

For calibration runs a different configuration is set. The first level memory pattern is set to require 4 out of 6 possible ANDs. In addition, two adjacent first level memory blocks are required to fire before a trigger is issued. This condition is imposed in order to favor muons which pass through many planes. Each calibration run enables a different segment of the detector, thus allowing muon data for a given plane to be concentrated into two or three runs. The shield masks are disabled, though the NIM enable is still required.

Appendix B

Liquid Scintillator and Crazing

The fiducial mass of the E645 detector consists mostly of liquid scintillator, specifically Bicron 517L. There are two factors which make the use of liquid scintillator more attractive than solid scintillator. First, it is less expensive. Second, it has the empirical formula CH_2 , whereas solid scintillator has the empirical formula CH . The hydrogen in the scintillator provides the target proton for the inverse beta decay reaction, which is the signature of neutrino oscillation. Furthermore, electron neutrinos from the LAMPF beam-stop interact on carbon and produce background. To maximize sensitivity to the oscillation signal and minimize background, one should choose the scintillator with the largest hydrogen to carbon ratio. The properties of Bicron 517L are shown in table B.1

The tanks to hold liquid scintillator must be designed with a number of factors in mind. First, they must be made of a low A (atomic weight) material. Heavy nuclei tend to have extra neutrons, which substantially

Table B.1: Physical constants of the liquid scintillator Bicorn 517L.

Density	.858
Refractive Index	1.47
Boiling Point	350°C
Light Output	40% Anthracene
Decay Constant	4 nsec
Wavelength of Maximum Emission	420 nm
H/C Ratio	2.0

increases their ν_e cross sections. In addition, it must be inert with respect to the scintillator — a chemical reaction between the container and scintillator can “poison” the scintillator and destroy its light producing capability. Aluminum and copper cannot be used for these reasons. Stainless steel can be, but it has a large number of free neutrons. Glass satisfies both of the above conditions, but it is expensive, dense, and difficult to work with. Plastics, made primarily of hydrogen, carbon and oxygen, seem to satisfy all the above requirements. They provide more hydrogen targets, and the carbon and oxygen nuclei have $N = Z$, and therefore small cross sections for neutrino background. Plastics do not poison scintillator, are light (to minimize “dead” material in the detector), relatively cheap, and easy to mold and glue.

Scintillator containers for this experiment are made from the plastic polymethyl methacrylate (PMMA), molded by an extrusion process. The dimensions of the extrusions, are 3.7 m. $\times$ 30 cm. $\times$ 3 cm., with a wall

thickness of .32 cm.. Acrylic end plates are attached to the end of the extrusions by applying a solvent to melt the end of the extrusion and then affixing the end plate. A gluing jig maintains the contact between the two. A strong seal is formed in this way. Acrylic cement is applied as a fillet around the joint between the end plate and extrusion to seal any gaps and to provide further mechanical support.

Early tests of the scintillation counters revealed a severe problem. The counters tended to crack a few days to a few weeks after being filled with scintillator. This was due to a crazing process. This appendix will first explain the nature of the crazing process, and then the steps that were taken in order to alleviate the problem for this experiment.

B.1 The Physics of Crazing

Crazing is a process which occurs in plastics under stress or in contact with solvents. It is exhibited as small parallel cracks at the surface of the plastic. These can grow and penetrate into the plastic as a function of time, substantially decreasing the load-bearing capability of the plastic. If the plastic is under stress, cracking will occur.

Theoretical models of crazing [80,82] suggest that in the case of solvent-induced crazing it is caused by capillary pressure and hydrodynamic transport of the solvent into the crack. This is confirmed by experiment [81,82]. The crazing model is as follows: Crazes occur at small flaws in the surface

of the plastic. The fluid adheres to the surface, and is drawn by capillary pressure into the flaw. Under the correct physical conditions (temperature, viscosity, adhesion, surface stress, as will be discussed below) the capillary force at the innermost surface of the flaw is enough to break apart the fibrils of polymer making up the plastic. At this point, a craze begins to propagate. Fluid flows into the craze, drawn in by capillary pressure, and forms a liquid wedge which drives itself through the plastic. The plastic is not completely split. Some fibrils remain, spanning the crack. In this way crazed plastic can support up to half the load that uncrazed plastic can. Under stress and in continuous contact with solvent, however, voids will form in the crazes and the plastic will inevitably rupture. Crazes do not form under all conditions, though, which makes it possible to control their occurrence.

The variables that control craze growth are solvent viscosity, the elasticity and yield strength of the plastic, the stress that the plastic is under, the external pressure, and the temperature. Under the condition that the fluid can enter the craze only at the surface of the plastic (and not around the sides), the equation for craze growth is given by [81]

$$x = C \left(\frac{P}{\sigma_y E \eta} \right)^{\frac{1}{2}} (K_0 - K_m) t^{\frac{1}{2}} \quad (\text{B.1})$$

where x is the craze length, C is a constant, P is external pressure, σ_y is the yield strength of the plastic and E is its elasticity, K_0 is the "stress-intensity factor" (given in units of $\text{N mm}^{-3/2}$), K_m is a threshold value of the stress

intensity factor, and t is the time. K_m is notable in that it implies that there is a stress threshold below which crazing *will not occur*. This threshold is a function of temperature. Di Benedetto et al. [81] measured the crazing threshold for PMMA in n -butyl alcohol at three different temperatures, shown in table B.2.

Table B.2: Crazing threshold K_m as a function of temperature for PMMA in n -butyl alcohol. Di Benedetto et. al.

$T(^{\circ}\text{C})$	$K_m (\text{N mm}^{-3/2})$
8	3.82
15	2.66
35	1.04

Crazing can be prevented by keeping the stress in a plastic in contact with a solvent below a certain threshold. Applying this information to the design of the E645 scintillation counters led to the solution of the crazing problem.

B.2 E645 Crazing Tests

The plastic used in the construction of the extruded scintillation counters is polymethyl methacrylate (PMMA), manufactured by Rohm & Hass, type VM-100. The scintillator used in the experiment was Bicron 517L, which is 15% pseudocumene (1,2,4 tri-methyl benzene). Bicron 517H, which is 30% pseudocumene was used in the early phases of the experiment. As mentioned above, this scintillator can cause craze-induced fractures of the

PMMA containers. A number of tests were run in order to study the relation between stress and crazing, and to thereby devise methods of preventing the stress fracturing of the counters.

B.2.1 Extrusion Stress Testing

It was observed that counters which were laid on their side and filled with scintillator were immune from craze fracturing. Indeed, fracturing was first observed when 20 extrusions were hung in the vertical position (with the 30 cm. side vertical) for an extended period of time. Counters hung in this way show a distortion of 8 mm. over their normal width of 3.0 cm. This distortion induces a great deal of stress in the counter, and as mentioned in the previous section, both the likelihood and speed of crazing are determined by stress.

The stress-crazing relation for the extrusions was determined by subjecting $1\frac{1}{2}$ " slices of an extrusion to various distortions and placing them in a bath of liquid scintillator. The distortions were caused by sections of a 1" hexagonal Teflon rod, which was cut into sections of varying length and then shaped so that the force was distributed evenly over the contact surface. These rods were placed 3" from one end of the test piece. The pieces were then immersed in a scintillator bath. The time to fracture was measured. Pieces were observed for several weeks.

Tests were performed in both Bicron 517L and 517H. The results are shown in figure B.1 as the diamonds and boxes. Several features are evident.

- The time to fracture seems to vary exponentially with distortion.
- There is evidence for a “crazing threshold”, a distortion below which crazing will not occur. This is much more evident in the case of 517H, which shows a discontinuity in the slope between 37 and 40 mm. of distortion.
- Pieces immersed in 517H (30% pseudocumene) fractured much more quickly than those immersed in 517L (15% pseudocumene).

This information led to the choice of Bicron 517L as the scintillator used for the experiment.

B.2.2 Annealing

One common procedure used to make plastics more craze resistant is annealing [80]. In annealing, the plastic piece is held at an elevated temperature (below its melting point) for an extended period of time. Several test pieces, identical to the ones mentioned above, were annealed in an oven at 76° C for 24 hours. These were then subjected to the same distortion tests as the unannealed pieces. The results of these tests are shown on the lifetime plot as the x's and stars. There seems to be no difference in the lifetimes between annealed and unannealed pieces at equal distortions. To determine whether annealed and unannealed test pieces respond in the

Figure B.1: Time to fracture for annealed and unannealed $1\frac{1}{2}$ inch extrusion sections in a bath of Bicron 517L and 517H. Diamond: 517H unanld.; $\times$: 517H anld.; Box: 517L unanld.; Star: 517L anld.

same way to applied stress, a distortion versus stress test was conducted by hanging weights from an extrusion test piece (3" from the end as in the lifetime tests) and measuring the distortion. The results for annealed and unannealed pieces are shown in figure B.2. Annealing seems not to affect the elastic properties of the plastic in a noticable way. One can conclude that annealing is not an effective way to combat the crazing problem.

B.3 Solution of the Crazing Problem

A number of solutions were suggested as possibilities for eliminating the crazing problem. A number of these were pursued. One possibility entailed coating the inside of the counter with a substance which is crazing resistant. Two coatings were tried: an ultraviolet-cured plastic which is used in coating plastic lenses to make them scratch resistant, and poly-vinyl alcohol (PVA). The ultraviolet-cured coating significantly degraded the reflectivity of the counter surface, and so was not considered. The PVA coating tended to pull away from the counter surface when it was distorted, and was also regarded as unworkable.

Since crazing is a function of the surface stress inside of a filled counter, it follows that by physically constraining the counter one may be able to relieve the stress to the point where crazing does not occur. Constraining the counters from the outside would require a massive frame, and would lead to a large fraction of dead material inside the detector. Instead, an

Figure B.2: Distortion of a $1\frac{1}{2}$ " section of extrusion as a function of the force applied 3" from one end. Diamonds: Unannealed; Squares: Annealed.

internal support consisting of an acrylic plastic rib which runs the length of the counter along the center line was tried. This rib is glued in place with the acrylic solvent mentioned in the first section, which provides a very solid bond between the rib and the extrusion. Distortion of a ribbed counter when it is filled is reduced to the order of a millimeter. This configuration transfers stress to the joint between the rib and the extrusion, and it was found that crazing occurred along this joint. This is remedied by applying a fillet of PVA along the rib-extrusion joint. The distortion at this joint is not enough to pull the PVA layer from the surface, so the PVA protects the joint from crazing.

The median time to failure of the counters used during the 1984 beam test, which were filled with 517H, was on the order of two weeks. The counters used in the detector have been filled since May 1986. In the period between May 1986 and October 1988, about 20% of the counters have leaked, and in only about half of these was the leak due to a failure of the type described above. Clearly, substantial progress has been made in controlling the crazing problem. One final weapon used against crazing has been temperature control. As shown in table B.2, the stress threshold for crazing is a function of temperature. Low temperatures raise the crazing threshold, so it is desirable to keep the counter temperature as low as possible. Air conditioning units cool both the neutrino tunnel and the inside of the active shield. This has also helped to preserve the counters.

Bibliography

- [1] F. Reines, Ann. Rev. Nuc. Sci. **10**,1 (1960)
- [2] E. Fermi, Z. Physik **88**, 161 (1934)
- [3] F. Reines et. al., Phys. Rev **117**, 159 (1960)
- [4] G. Feinberg and L. Lederman, Ann. Rev. Nuc. Sci. **13** ,431 (1963)
- [5] E. Majorana, Nuovo Cimento **14**, 171 (1937)
- [6] R. Davis, Phys. Rev. **97**, 766 (1955)
- [7] S. P. Rosen, Comments Nucl. Part. Phys. **18**,31 (1988)
- [8] G. Danby, et. al. Phys. Rev. Lett. **9**, 36 (1962)
- [9] F. J. Hasert, et. al., Phys. Lett. **46B**, 138 (1973)
- [10] S. Boris et. al., Phys. Rev. Lett. **58**, 2019 (1987)
- [11] M. Fritschi et. al., Phys. Lett. **173B**, 485 (1985)
- [12] J. F. Wilkerson et. al., Phys. Rev. Lett. **58**, 2023 (1987)
- [13] Particle Data Group, *Review of Particle Properties*, Physics Letters B **170B** pp. 97-98, 104 (1986)
- [14] M. Gell-Mann and A. Pais, Phys. Rev. **97**,1387 (1955); A. Pais and O. Piccioni, Phys. Rev. **100**,1487 (1955)

- [15] K. Lande et.al., Phys. Rev. **103**,1901 (1956); Phys. Rev. **105**, 1925 (1957)
- [16] B. Pontecorvo, JETP **34**, 247 (1958)
- [17] S. Bilenky and B. Pontecorvo, Phys. Reports **41**, 225 (1978)
- [18] R. E. Marshak, Riazuddin, C. P. Ryan *Theory of Weak Interactions in Particle Physics* Monographs and Texts in Physics and Astronomy, v. 24, Wiley Interscience (1969), p. 64
- [19] S. Bilenky and B. Pontecorvo, Phys. Lett. **61B**, 248 (1976)
- [20] S. Eliezer and A. Swift, Phys. Rev. D **10**, 3088 (1974)
- [21] A.K. Mann and H. Primakoff, Phys. Rev. D **15**, 655 (1977)
- [22] H. Primakoff and S. P. Rosen, Ann. Rev. Nucl. Part. Sci. **31**, 145 (1981)
- [23] J. N. Bahcall and H. Primakoff, Phys. Rev. D **18**, 3463 (1978)
- [24] Donald H. Perkins, *Introduction to High Energy Physics*, 2nd ed., Addison-Weseley Publishing Co., p. 227 (1982)
- [25] L. A. Ahrens et. al., Phys. Rev. D **31**, 2732 (1985)
- [26] C. Angelini et. al., Phys. Lett. B **179**, 307 (1986)
- [27] V. Zacek et. al., Phys. Lett. **164B**, 193 (1985)
- [28] V. V. Ammosov et. al., presented at the XXIV International Conference on High Energy Physics, Munich 1988.
- [29] A good description of the GSW model can be found in Ch. 11 of I. J. R. Aitchison and A. J. G. Hey *Gauge Theories in Particle Physics, A Practical Introduction*, p.239. Adam Hilger Ltd., Bristol (1982)

- [30] B. Sadoulet, *Proceedings of the Fifteenth SLAC Summer Institute on Particle Physics*, p. 277. Eileen Brennan, ed. SLAC Report No. 328 (1987)
- [31] J. N. Bahcall, et. al., *Astrophys. J.* **292**, 279 (1985)
- [32] S. P. Mikheyev and A. Yu. Smirnov *Proceedings of the Tenth International Workshop on Weak Interactions*, Savlonia, Finland, June 1985
- [33] L. Wolfenstein, *Phys. Rev. D* **16**, 2369 (1978)
- [34] H. A. Bethe, *Phys. Rev. Lett.* **56**, 1305 (1986)
- [35] S. P. Rosen and J. M. Gelb, *Phys. Rev. D* **34**, 969 (1986)
- [36] W. C. Haxton, *Phys. Rev. Lett.* **57**, 1271 (1986)
- [37] A. K. Mann, Presented at the 13th International Conference on Neutrino Physics and Astrophysics, Boston MA (June 5-11, 1988)
- [38] J. M. Lattimer and J. Cooperstein, *Phys. Rev. Lett.* **61**, 23 (1988)
- [39] R. Barbieri and R. Mohapatra, *Phys. Rev. Lett.* **61**, 27 (1988)
- [40] M. Gell-Mann, P. Ramond, and R. Slansky, *Supergravity*, eds. P. van Neuwenheusen and D. Freedman, North-Holland, Amsterdam (1979)
T. Yanagida, *Prog. Theor. Phys. B* **135**, 66 (1978)
- [41] P. Q. Hung and S. Mohan, *Phys. Rev. Lett.* **61**, 31 (1988)
- [42] K. Kang and M. Shin, *Phys. Lett. B*, **185**, 163 (1987)
- [43] R. Peccei and H. Quinn, *Phys. Rev. Lett.* **38**, 1440 (1977)
- [44] *LAMPF Users Handbook*, (MP-DO-3-UHB (Rev.)), p. 1.-3 Los Alamos Natl. Lab. (1984)
- [45] D. R. Davidson, et. al., submitted to *Effects of Radiation on Materials 13th Symposium, Seattle, Washington, 1986*

- [46] E. Smith, Private Communication.
- [47] Particle Data Group, *Review of Particle Properties*, Physics Letters B **170B** p. 43 (1986)
- [48] V. Sandberg, E645 Internal Memo (Sept., 1984)
- [49] R. Miranda, *Instrumentation Studies for a High-Energy Physics Experiment*, Masters Thesis, Ohio State University, p. 47 (1983)
- [50] J. Fitch, et. al., Nucl. Instr. & Meth. **215**, 71 (1984)
- [51] J. Mitchell, *Using the Cosmic Ray PDT Tester*, E645 Internal Memo (June, 1985)
- [52] L. V. Groshev, et. al., *Nuclear Data Tables A5*, 1-245 (1968)
- [53] B. Fujikawa, *Neutron Detection - E645*, E645 Internal Memo (1986)
- [54] B. Fujikawa, *Search for Neutrion Oscillation Events (NOES)*, E645 Internal Memo (1988)
- [55] S. J. Freedman, et. al., Nucl.Instr. & Meth. **215**, 71 (1983)
- [56] E. Smith *E645 Backgrounds — 1986*, E645 Internal Memo (Mar. 1987)
- [57] J. J. Napolitano, et. al., *Construction and Performance of a Large Area Liquid Scintillator Cosmic Ray Anticoincidence Detector*, ANL Report PHY-5277-ME-88, Submitted to Nucl. Inst. and Methods. (1988)
- [58] R. Brun, et. al., *ZBOOK User Guide and Reference Manual, vers. 2*, CERN Data Handling Division, DD/US/73 (1982)
- [59] W. R. Nelson, et. al., *The EGS4 Code System*, Stanford Linear Accelerator Center, Document # SLAC-265 (1985)

- [60] E. S. Smith, *Neutrino Backgrounds to the Searches for the Appearance of $\bar{\nu}_e$ at LAMPF*, E645 Internal Memo, (Nov. 1987)
- [61] E. Smith *Parameters that Determine Neutrino Oscillation Limits*, E645 Internal Memo (Mar. 1988)
- [62] B. Kayser, et. al., Phys. Rev. **D20**, 87 (1979)
- [63] U. Amaldi, et. al., Phys. Rev. **D36**, 1385 (1987)
- [64] R. C. Allen, et. al., Phys. Rev. Lett. **55**, 2401 (1985)
- [65] T. W. Donnelly, et. al. Phys. Lett. **43B**, 93 (1973)
- [66] E. Smith, *Cross Section for $\bar{\nu}_e p$* , E645 Internal Memo, (Nov. 1987)
- [67] J. S. O'Connell, Proc. of the Los Alamos Neutrino Wkp., June 8-12, 1981, p.44
- [68] S. A. Fayans, Sov. J. Nucl. Phys. **42**, 590 (1985)
- [69] D. R. F. Cochran, et. al., Phys. Rev. **D6**, 3085 (1972)
- [70] J. F. Crawford, et. al., Phys Rev. **C22**, 1184 (1980)
- [71] R. Allen and J. B. Donahue, Spokespersons, LAMPF Proposal 866, *Calibration of the Neutrino Flux at the LAMPF Beam-Stop*
- [72] R. L. Burman and E. S. Smith, *Parameterization of Pion Production and Reaction Cross Sections at LAMPF Energies*, E866 Internal Memo (Apr. 1988)
- [73] R. L. Burman, M. E. Potter, and E. S. Smith, *A Monte Carlo Simulation of Neutrino Production at a Medium Energy Proton Beam-Stop*, To be Published.
- [74] T. W. Donnelly, Private Communication
- [75] W. T. Eadie, et. al. *Statistical Methods in Experimental Physics*, North-Holland Publ. Co., New York, (1971)

- [76] J. Orear, *Notes on Statistics for Physicists, Revised*, Cornell Univ., CLNS 82/511 (1982)
- [77] Particle Data Group, *Review of Particle Properties*, Phys. Lett. B **170B**, 55 (1986)
- [78] M. Murtagh and W. Lee, Spokespersons, *Neutrino Oscillation Experiment at BNL*, Brookhaven Proposal P-846 (1988)
- [79] D. H. White, Spokesperson, *A Proposal for a Precision Test of the Standard Model by Neutrino-Electron Scattering (The Large Čerenkov Detector Project)*, Los Alamos Proposal LA-11300-P (1988)
- [80] *Encyclopedia of Plastic and Polymer Science*, **4**, 294 (1964)
- [81] Di Benedetto, et. al., J. Mat. Sci. **16**, 2310 (1981)
- [82] E. J. Kramer and R. A. Bubek, J. Polymer Sci. **16**, 1195 (1978)