

LA-10937-T

Thesis

UC-34

Issued: February 1987

LA--10937-T

DE87 007721

A Search for the Rare Decay $\mu^+ \rightarrow e^+ \gamma\gamma$

David Paul Grosnick

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

MASTER

Los Alamos Los Alamos National Laboratory
Los Alamos, New Mexico 87545

S
DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS	vi
LIST OF TABLES	ix
ABSTRACT	x
Chapter	
I. INTRODUCTION	1
II. THEORY	16
A. The Standard Model	
B. The Standard Model and Lepton-Family Number Nonconservation	
C. Extensions to the Standard Model	
D. Model Extensions Relevant to $\mu \rightarrow e\gamma\gamma$	
III. DETECTOR	50
A. The Muon Beam and Target	
B. The Drift Chamber	
C. Scintillator Hodoscope	
D. NaI(Tl) Calorimeter	
1. Design	
2. Signal Pulse Shaping	
3. NaI(Tl) Electronics	
4. Primary ADC and Time Systems	
5. Pile-up ADC System	
E. I-Counter	
F. Guard Counters	
G. Trigger	
H. Data Acquisition	
I. Liquid Hydrogen Target	
J. Flasher System	
IV. MONTE CARLO PROGRAM, CALIBRATIONS, AND STABILIZATION	89
A. Monte Carlo Program	
B. NaI(Tl) Energy Calibration	
C. Scintillator Energy Calibration	
D. Scintillator Timing Calibration	
E. NaI(Tl) Timing Calibration	
F. Drift Chamber Timing Calibration	
G. NaI(Tl) Position Calibration	

H. NaI(Tl) Energy Stabilization	
I. Timing Stabilization	
V. DATA ANALYSIS	129
A. Pile-up Rejection	
B. First Pass Data Analysis	
C. Second Pass Data Analysis	
D. Maximum Likelihood Analysis	
VI. RESULTS AND CONCLUSIONS	162
APPENDIX	167
REFERENCES	205
ACKNOWLEDGMENTS	215

LIST OF ILLUSTRATIONS

Figure

1. Upper Limits on the Branching Ratio for Several Neutrinoless Muon Decay Modes as a Function of Time	10
2. Schematic View of the Crystal Box Detector	52
3. Diagram of the NaI(Tl) Signal Pulses, Showing the Original Pulse and the Pulses After Signal Processing with a Pole-zero Filter and a Cable Clip	67
4. Schematic Diagram of the NaI(Tl) Array Electronics	70
5. Diagram of the Processed NaI(Tl) Signal Pulse and the Two Current-integrating Time Periods for the Primary and Pile-up ADC Systems	76
6. Logic Diagram for $\mu \rightarrow e\gamma$ and $\mu \rightarrow e\gamma\bar{\nu}$ Triggers	82
7. Photon Energy Spectrum from $\pi^0 \rightarrow \gamma\gamma$ Events	97
8. Photon Energy Spectra from the $\pi^- p \rightarrow n \gamma$ Reaction	99
9. Photon Energy Spectra from $\pi^0 \rightarrow \gamma\gamma$ Events with the Opening Angle Between the Photons Greater Than 175°	101
10. Photon Energy Spectra from a Pu- α -Be Source	104
11. Positron Energy Spectra from Muon Decay	106
12. Double-peaked Time Spectrum from the Upstream Channel of Scintillator Number 6	110
13. Upstream and Downstream Channel Time Difference Spectrum from Scintillator Number 6	113
14. Double-peaked Mean Time Spectrum from Scintillator Number 6	115
15. Measured Time Resolution of the NaI(Tl) Array for Two Photons from the Decay $\pi^0 \rightarrow \gamma\gamma$	120
16. Plot of the Probability of $\mu \rightarrow e\gamma\bar{\nu}\bar{\nu}$ Events as a Function of Energy Threshold for Both Positron and Photon	127

17. Plot of the Ratio of the Energy Deposited in the Pile-up ADC System to the Energy in the Primary ADC System Versus the Energy Measured in the Primary ADC System in Units of ADC Channels for a Typical Crystal	132
18. Plot of the Probability of Pile-up in a NaI(Tl) Crystal Versus the Amount of Piled-up Energy for Photons at 4.83 MHz (Instantaneous) Muon Stopping Rate	135
19. Plot of the Time Differences Between the Positron and the Higher-energy Photon (γ_1), $t_e - t_{\gamma_1}$, and the Positron and the Lower-energy Photon (γ_2), $t_e - t_{\gamma_2}$	144
20. Distribution of 272 Events in the Diagonal Time Projection Variable, τ	148
21. Diagram Showing the Particle Momenta, \vec{p}_a , \vec{p}_b , and \vec{p}_c	150
22. Distributions of the Surviving Nine Data Events in the Four Analysis Variables, E_{Tot} , p_{\parallel} , $\cos \alpha$, and τ	152
23. Monte Carlo-generated Probability Distributions in the Four Analysis Variables, E_{Tot} , p_{\parallel} , $\cos \alpha$, and τ , for the Three-quadrant $\mu^+ \rightarrow e^+ \gamma\gamma$ Events with No Added Piled-up Energy	156
24. Background Probability Distributions in Three of the Four Analysis Variables, E_{Tot} , p_{\parallel} , and $\cos \alpha$	159
25. The Normalized Likelihood Function Versus the Number of $\mu^+ \rightarrow e^+ \gamma\gamma$ Events, n_{evy} , for the Nine Candidate Data Events	161
26. Schematic Diagram of the Flasher System	169
27. The Flashtube Driver Circuit	172
28. Operating Circuit for the Vacuum Photodiode	175
29. Photodiode Monitor High Voltage Plateau Curves	177
30. Operating Circuit for the Silicon Photodiode	179
31. Pulse-shaping Circuit for the Photodiode Monitors	182
32. Raw Data Spectra of a Given Crystal Using the Flasher System	185
33. The Distribution of Energy Peak Positions for a Given Data Run Using the Flasher System	188
34. Plot of the Ratio of the Peak Position from the Vacuum Photodiode to the Peak Position of the Silicon Photodiode as a Function of Time	190

35. Plot of the Average Energy Gain Corrections from the 432 NaI(Tl) ADC Channels Using the Flasher System as a Function of Time	193
36. Photon Energy Spectrum Using the Flasher-corrected Energy Gains 300 Hours After an Energy Calibration	196
37. The Distribution of the Percent Change of the Energy Gains Using the Flasher Corrections and the Gains Determined by an Energy Calibration	198
38. The Distribution of Timing Peak Positions for a Given Data Run Using the Flasher System	201
39. Plot of the Average Timing Peak Positions from the 432 NaI(Tl) TDC Channels Using the Flasher System as a Function of Time	203

LIST OF TABLES

Table

1. The Konopinski and Mahmoud Lepton Number Arrangement	4
2. The Additive Lepton Number Arrangement	5
3. The Multiplicative Lepton Number Arrangement	7
4. Representations of Leptons and Quarks in the Standard Model	28
5. Weaknesses and Possible Extensions to the Standard Model, and the Possible Decay Modes That Can Be Observed	36
6. Values of the Depth Parameter, d, as a Function of Energy and Particle Type	123
7. First Pass Data Reduction	137
8. Second Pass Data Reduction	141

ABSTRACT

An experimental search for the lepton-family number nonconserving decay, $\mu^+ \rightarrow e^+ \gamma\gamma$, has been conducted at the Clinton P. Anderson Meson Physics Facility (LAMPF) using the Crystal Box detector. The detector consists of a modular NaI(Tl) calorimeter, scintillator hodoscope, and a high-resolution, cylindrical drift chamber. It provides a large solid-angle for detecting three-body decays and has good resolutions in the time, position, and energy measurements to eliminate unwanted backgrounds. No evidence for $\mu^+ \rightarrow e^+ \gamma\gamma$ is found, giving an upper limit for the branching ratio of $\Gamma(\mu \rightarrow e\gamma\gamma)/\Gamma(\mu \rightarrow e\nu\bar{\nu}) \leq 7.2 \times 10^{-11}$ (90% C.L.). This result is an improvement of more than two orders of magnitude in the existing limit.

CHAPTER I

INTRODUCTION

The muon is one of a plethora of subnuclear particles found to exist in nature. It is designated a lepton because it is a fermion that does not interact through the strong interaction. Atomic and nuclear physics require only two leptons, the electron and the neutrino. This fact poses the question, What is the purpose of the muon? That same question can be rephrased into the language of modern particle physics, Why is there more than one quark and lepton generation? Experiments studying the muon may provide a clue to answer this basic question.

The muon was discovered in 1936 by Anderson and Neddermeyer¹ observing cosmic rays with cloud chambers and Geiger counters. They determined that these particles could not be protons or positrons from range and energy loss measurements. The mass of this particle was around 200 times the electron mass and positrons appeared in the decay of this positively charged particle.

For over a decade the muon was thought to be the quantum mediating the strong nuclear force as predicted by Yukawa.² An experiment³ performed in 1947 using muons absorbed in dense materials clearly showed however that the muon does not interact via the strong force and hence could not possibly be Yukawa's meson.

The muon interacts with other particles in a manner very similar to that of an electron. It does not appear to have an "anomalous" interaction with matter, but is governed by the electromagnetic and weak forces. The hypothesis that muons have exactly the same electromagnetic and weak couplings as electrons is called muon-electron universality. High precision tests of this hypothesis so far have found no significant differences between muons and electrons.

With the production of muon beams generated by accelerators, systematic and quantitative studies on the properties of the muon proceeded. Fitch and Rainwater⁴ found that the muon mass was about 210 times the electron mass, the spin of the muon was 1/2, and the muon exhibited no anomalous interaction with the nucleus. Experiments⁵ on the $\pi^+ \rightarrow \mu^+ \rightarrow e^+$ decay sequence by two separate groups found that parity is not conserved in muon decay. Muons from pion decay have a strong longitudinal polarization and the angular distribution of the decay electrons has a large asymmetry with respect to the muon spin. At this time, the lifetime, mass, and magnetic moment are all known to high precision.

The muon decays into an electron, a neutrino, and an antineutrino. Only the electron of these three particles is easy to detect, and the broad electron energy spectrum implies at least a three body decay. Michel⁶ characterized the energy spectrum of the decay electron with five parameters. Two of these, ρ and η , determine the shape of the electron momentum spectrum from unpolarized muons; another two, ξ and δ , describe the muon spin - electron momentum correlation; and the final parameter, h , is a measure of the degree of electron polarization. The experimental

values of the first four parameters are obtained from intensity measurements. The present experimental values of these Michel parameters agrees satisfactorily with the V - A theory, but the values still do not give a definitive choice of coupling constants for a more general weak amplitude.

Investigators have searched for alternate decay modes of the muon, such as $\mu \rightarrow e\gamma$, $\mu \rightarrow eee$, and $\mu \rightarrow e\gamma\gamma$. These processes are kinematically possible and obey all conservation principles associated with space-time symmetries. Why are they not observed?

To explain this, Konopinski and Mahmoud⁷ in 1953 proposed the first lepton-number conservation law. Table 1 summarizes their arrangement. The assignments for lepton number are $L = +1$ for the e^- , μ^+ , and ν , while their antiparticles are assigned $L = -1$, and everything else is assigned $L = 0$. The conservation law is stated: $\Sigma L = \text{constant}$. Such a law forbids decays such as $\mu^+ \rightarrow e^+\gamma$ and $\mu^+ \rightarrow e^+e^-e^+$, but allows the decay, $\mu^+ \rightarrow e^+\nu\nu$. This last decay contains identical neutrinos in the final state, and this makes their hypothesis untenable since it implies a value for the Michel parameter ρ that is contrary to experimental evidence.

An improvement, the additive lepton number conservation law with separate lepton-family numbers introduced in 1957 by three separate individuals,^{8,9} Schwinger, Nishijima, and Bludman, uses separate quantum numbers for each of the "families," electron and its neutrino, and the muon and its neutrino. It can be extended to more than two "generations" of particle families, a requirement after the τ lepton discovery.¹⁰ Table 2 lists the values of the assignments for this scheme. The conserved quantities are: $\Sigma L_e = \text{constant}$, $\Sigma L_\mu = \text{constant}$, etc. This now

TABLE 1

THE KONOPINSKI AND MAHMOUD LEPTON NUMBER ARRANGEMENT

L	Particle
+1	e^- , μ^+ , ν
-1	e^+ , μ^- , $\bar{\nu}$
0	All other particles

Conservation Law: $EL = \text{constant}$

Forbids: $\mu^+ \rightarrow e^+ \gamma$

$$H^+ \rightarrow e^+ \gamma\gamma$$

$$u^+ \rightarrow e^+ e^- e^+$$

$$\mu^-Z \rightarrow e^-Z$$

$$K_L^0 \rightarrow \mu^\pm e^\mp$$

Allows: $\mu^- Z \rightarrow e^+ (Z-2)$

$$(\mu^+ e^-) \rightarrow (\mu^- e^+)$$

$$K^+ \rightarrow \mu^+ e^+ \pi^-$$

$$u^+ \rightarrow e^+ \nu \bar{\nu}$$

TABLE 2
THE ADDITIVE LEPTON NUMBER ARRANGEMENT

L_e	Particle	L_μ	Particle	L_τ	Particle
+1	e^- , ν_e	+1	μ^- , ν_μ	+1	τ^- , ν_τ
-1	e^+ , $\bar{\nu}_e$	-1	μ^+ , $\bar{\nu}_\mu$	-1	τ^+ , $\bar{\nu}_\tau$
0	All others	0	All others	0	All others

Conservation Laws: $\Sigma L_e = \text{constant}$

$\Sigma L_\mu = \text{constant}$

$\Sigma L_\tau = \text{constant}$

Forbids: $\mu^+ \rightarrow e^+ \gamma$

Allows: $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$

$\mu^+ \rightarrow e^+ \gamma \gamma$

$\mu^+ \rightarrow e^+ e^- e^+$

$\mu^- Z \rightarrow e^- Z$

$\mu^- Z \rightarrow e^+ (Z-2)$

$\mu^+ \rightarrow e^+ \bar{\nu}_e \nu_\mu$

$(\mu^+ e^-) \rightarrow (\mu^- e^+)$

$\tau^- \rightarrow \mu^- \gamma$

$\tau^- \rightarrow \mu^- e^+ e^-$

$e^- e^- \rightarrow \mu^- \mu^-$

$K_L^0 \rightarrow \mu^\pm e^\mp$

$K^+ \rightarrow \pi^+ \mu^\pm e^\mp$

$\Xi^+ \rightarrow p \mu^\pm e^\mp$

forbids all neutrinoless $\mu \rightarrow e$ processes, but still allows the normal muon decay process.

Feinberg and Weinberg¹¹ proposed yet a third lepton conservation law in 1961. This is a multiplicative scheme where the conserved quantities are $\mathbb{E}L = \text{constant}$ and $\mathbb{M}L_p = \text{constant}$, where L_p is the "muon parity." Table 3 lists the assignments for this method. Again, all neutrinoless processes are forbidden.

The multiplicative and additive lepton number arrangements are equally good a priori, since they both prohibit reactions such as $\mu^+ \rightarrow e^+ \gamma$, $\mu^+ \rightarrow e^+ e^- e^+$, and $\mu^- Z \rightarrow e^- Z$. Each arrangement does have different consequences, however. The additive lepton number conservation law is more restrictive than the multiplicative one. It prohibits muonium - antimuonium conversion and muon decay with inverted neutrino subscripts.

The last two lepton number assignments require separate types of neutrinos for the muon and the electron. When introduced, there was no experimental evidence for distinct types of neutrinos. Pontecorvo and Schwartz¹² both proposed experiments to reveal their existence. In 1962, an important experiment¹³ at the Brookhaven National Laboratory discovered that there were indeed two different types of neutrinos. The consequence of this experiment was to confirm the presence of separate lepton numbers. With this experiment, the search for rare decay modes of the muon essentially ended for about 15 years.

A resurgence of interest in the search for rare muon decays appeared in 1977 when rumors circulated that an experiment¹⁴ at the Swiss Institute for Nuclear Research (SIN) found a signal for the decay, $\mu \rightarrow e \gamma$.

TABLE 3
THE MULTIPLICATIVE LEPTON NUMBER ARRANGEMENT

L	Particle	L_p	Particle
+1	$\mu^-, e^-, \nu_e, \nu_\mu$	+1	$e^+, \nu_e, \bar{\nu}_e$
-1	$\mu^+, e^+, \bar{\nu}_e, \bar{\nu}_\mu$	-1	$\mu^+, \nu_\mu, \bar{\nu}_\mu$
0	All other particles	+1	All other particles

Conservation Laws: $\Sigma L = \text{constant}$

$\Sigma L_p = \text{constant}$

Forbids: $\mu^+ \rightarrow e^+ \gamma$

Allows: $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$

$\mu^+ \rightarrow e^+ \gamma \gamma$

$\mu^+ \rightarrow e^+ \bar{\nu}_e \nu_\mu$

$\mu^+ \rightarrow e^+ e^- e^+$

$(\mu^+ e^-) \rightarrow (\mu^- e^+)$

$\mu^- Z \rightarrow e^- Z$

$e^- e^- \rightarrow \mu^- \mu^-$

$\mu^- Z \rightarrow e^+ (Z-2)$

$K_L^0 \rightarrow \mu^\pm e^\mp$

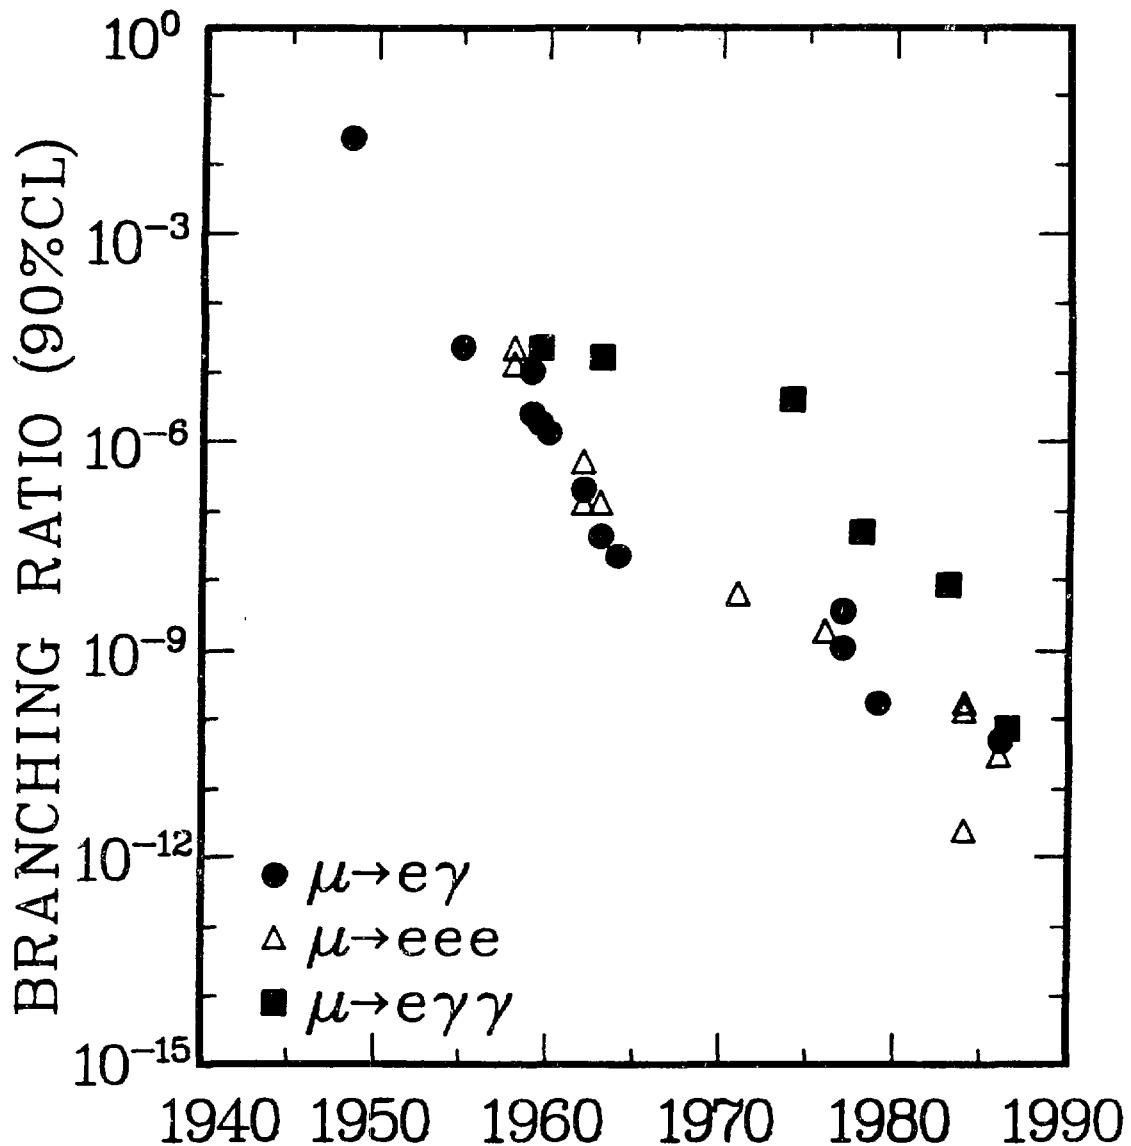
The rumors were later refuted, but not until after a burst of both theoretical and experimental activity. On the theoretical scene, over one hundred papers were published to show that there are ways to produce lepton-family number nonconserving decays. Many theoretical models could account for this phenomenon. Theorists found that suppressing lepton-family nonconserving processes to levels of around 10^{-8} to 10^{-13} was a simple and natural result of the new models.

The experimentalists also did not remain idle. Besides the search at SIN, two other experiments were quickly assembled to search for $\mu \rightarrow e\gamma$. One was located at the Tri-University Meson Facility (TRIUMF) in Vancouver, B.C., and conducted by a collaboration of several Canadian universities. The other experiment at LAMPF was a collaboration from the Los Alamos Scientific Laboratory, Stanford University, and the University of Chicago. The results¹⁵⁻¹⁷ from these experiments showed no evidence for the decay $\mu \rightarrow e\gamma$ at a level of $\sim 10^{-10}$.

Experimentally, the search for neutrinoless decay modes of the muon has spanned almost 40 years. Figure 1 plots the experimental status of the three neutrinoless decay modes as a function of time.¹⁸ Several innovations in muon sources and instrumentation give the groupings of points displayed in the plot. The first few searches¹⁹ used muons from cosmic ray showers and instruments consisting of Geiger-Müller tubes and absorbers. The next series of searches, until about 1964, used accelerators and stopped pion beams as sources of muons; the detectors were of various types, such as scintillation counters, water Cherenkov counters, spark chambers, a freon bubble chamber, and energy measurements with sodium iodide, NaI(Tl), crystals. The recent searches, from the mid

Figure 1. Upper limits on the branching ratio for several neutrinoless muon decay modes as a function of time.

RARE MUON DECAY EXPERIMENTAL HISTORY



1970's until the present, utilized the development of high intensity pion and muon beams from the meson "factories" and better techniques in spectrometer design to search for these rare decays to remarkably low limits. The branching ratio $B_{\mu\gamma\gamma\gamma}$ of the decay $\mu^+ \rightarrow e^+ \gamma\gamma\gamma$, for example, is defined as

$$B_{\mu\gamma\gamma\gamma} = \frac{\Gamma(\mu^+ \rightarrow e^+ \gamma\gamma\gamma)}{\Gamma(\mu^+ \rightarrow e^+ \bar{\nu}_e \nu_\mu)} , \quad (I-1)$$

where $\Gamma(\mu^+ \rightarrow e^+ \gamma\gamma\gamma)$ is the rate of the decay $\mu^+ \rightarrow e^+ \gamma\gamma\gamma$, and $\Gamma(\mu^+ \rightarrow e^+ \bar{\nu}_e \nu_\mu)$ is the rate of the normal muon decay. The present experimental bounds on the branching ratios for rare muon decay modes are:

$$B_{\mu\gamma\gamma\gamma} < 4.9 \times 10^{-11} \quad (90\% \text{ C.L.}) \quad \text{Ref. 20}$$

$$B_{\mu e e e} < 2.4 \times 10^{-12} \quad (90\% \text{ C.L.}) \quad \text{Ref. 21}$$

$$B_{\mu\gamma\gamma\gamma} < 8.4 \times 10^{-9} \quad (90\% \text{ C.L.}) \quad \text{Ref. 22.}$$

York, Kim, and Kernan²³ at the University of Chicago in 1959, first attempted to measure a two photon final state looking for the highly forbidden $\mu - e$ annihilation, $\mu^+ + e^- \rightarrow \gamma + \gamma$. The experiment stopped a beam of positive pions in a copper target, while a two arm spectrometer, each arm consisting of a sodium iodide crystal and scintillator, observed photons originating from the target. Only one event of over 5×10^5 candidate events was of the desired type, and this may have been a background event. This experiment set an upper limit on the branching

ratio for the sum of the processes, $\mu^+e^- \rightarrow \gamma\gamma$ and $\mu^+ \rightarrow e^+\gamma\gamma$, of 1.95×10^{-5} .

The second experiment²⁴ to measure $\mu \rightarrow e\gamma\gamma$ used a double arm spectrometer. The detector consisted of spark chambers, scintillation counters, and sodium iodide crystals set up to detect positrons in one arm and photons in the other. For the $\mu \rightarrow e\gamma\gamma$ search, both photons must be detected in the photon arm. There were no observed $e\gamma\gamma$ events after data reduction, placing an upper limit on the branching ratio of 1.6×10^{-5} . This experiment was conducted at the Lawrence Radiation Laboratory in 1963 by a group from the University of Pennsylvania.

The next attempt to observe $\mu \rightarrow e\gamma\gamma$ was conducted at the Lawrence Berkeley Laboratory in 1974 by a Canadian collaboration.²⁵ The detector involved in this experiment differed from a two arm spectrometer in that it stopped a beam of π^+ within one large sodium iodide crystal. The energy of the decay products of the muon was then measured. Any type of neutrinoless muon decay could be observed in this experiment. No evidence of a peak in the energy spectrum at the muon mass appeared. This results in an upper limit on the branching ratio for $\mu \rightarrow e\gamma\gamma$ of 4×10^{-6} .

In 1978, Bowman, Cheng, Li, and Matis²⁶ calculated an upper limit on the branching ratio for $\mu \rightarrow e\gamma\gamma$ of 5×10^{-8} using data from the experimental search¹⁶ for $\mu \rightarrow e\gamma$ at the TRIUMF facility. Again, the design of the detector used in this experiment was a two arm spectrometer.

The most recent experiment²² attempting to observe the decay $\mu \rightarrow e\gamma\gamma$, performed in 1983 by the TRIUMF group, used the same detector as in their 1977 experiment.¹⁶ They obtained an upper limit on the branching ratio of $B_{\mu\gamma\gamma} < 8.4 \times 10^{-9}$ (90% C.L.).

This thesis reports on the present search for the rare decay, $\mu \rightarrow e\gamma\gamma$, using the "Crystal Box" detector at LAMPF, located in Los Alamos, New Mexico. The experiment involves a collaboration from the Los Alamos National Laboratory, Stanford University, University of Chicago, and Temple University. The proposal²⁷ to search for $\mu \rightarrow e\gamma\gamma$ was submitted in 1979 and appended²⁸ later that same year so that the search for the three decays, $\mu^+ \rightarrow e^+\gamma$, $\mu^+ \rightarrow e^+e^-e^+$, and $\mu^+ \rightarrow e^+\gamma\gamma$, could be performed simultaneously using the same experimental detector. After several delays, the experiment first started taking data in December of 1983 for four weeks and accumulated the remainder of the data from June through October of 1984. The experiment produced a total of 1130 computer tapes; 677 of these contain data, the rest calibration tapes.

The goal of the experiment was a sensitivity in the branching ratio of the $\mu \rightarrow e\gamma\gamma$ decay of a few parts in 10^{11} . Meeting this goal requires an enormous quantity of muons, good measurements of the time, energy, and direction of the decay products so that unwanted backgrounds can be suppressed, and a large solid angle detector to capture the final state particles of a three body decay. The first requirement can be satisfied by running the experiment at a facility, such as LAMPF, that can supply the copious quantities of muons. The second and third requirements were met by designing and building a quality detector, the Crystal Box, that can distinguish the signal from the background events.

The signature for the $\mu^+ \rightarrow e^+\gamma\gamma$ decay is the following: (1) the positron and the two photons must originate at the same location and be in time coincidence; (2) the total energy of the positron and the two photons must equal the muon mass; (3) the vector sum of the momenta of the

positron and the two photons must equal zero. Because it leads to a three-body final state, the decay must be coplanar.

There are two types of backgrounds that can be present in this experiment, prompt and random. A prompt background is the allowed inner bremsstrahlung decay, $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu \gamma\gamma$. However, because this is a five body decay and the experiment measures the sum of the positron and two photon energies close to the muon mass, the probability of observing it is extremely small. The random backgrounds are those processes that mimic the $\gamma\gamma$ signal. Included in these types of backgrounds are: (1) three particles, one positron and two photons, that are accidentally coincident in the detector (triple randoms); (2) an inner bremsstrahlung muon decay, $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu \gamma$, in accidental coincidence with a random photon (double random); (3) one particle misidentified as two separate particles and in accidental time coincidence with a third particle (split particle). Measurements of the time between the three particles, and the energy and direction of each must be very good in order to distinguish the signature of the signal process from the background processes.

The Crystal Box detector is so named because of its large number of thallium-activated sodium iodide, NaI(Tl), crystals arranged in the shape of a box. The detector consists of 396 NaI(Tl) crystals that give the time, position, and energy measurements of the particles, an array of 36 plastic scintillation counters that provide the trigger for the detector as well as the time measurement of the charged particles, and a high resolution cylindrical drift chamber that gives charged particle track information. The design of this detector is a large improvement over the

two arm spectrometers used in the previous experimental searches for $\mu \rightarrow e\gamma\gamma$.

Chapter II briefly describes the standard model, extensions to the standard model accomodating lepton-family number nonconservation, and then specifically the decay $\mu \rightarrow e\gamma\gamma$. Chapter III describes the detector, trigger, and data acquisition system used in this experiment. Chapter IV contains calibration procedures and results, as well as the Monte Carlo computer simulation results needed for the estimates of the resolutions and detector behavior. Chapter V describes the procedures used in the data analysis and the calculations necessary for the results for this experiment. The final chapter presents the results and conclusions from this experiment.

CHAPTER II

THEORY

The past 25 years shows much progress in attempts to understand the interactions and the structure of subnuclear particles. The culmination of this effort, the "Standard Model," describes in terms of gauge theories²⁹ the basic constituents of matter, believed to be quarks and leptons, and three of the four fundamental forces, electromagnetism, weak, and strong interactions. The successes and weaknesses of this model, as well as possible extensions, will be discussed briefly. Finally, some extensions relating specifically to the $\mu \rightarrow e\gamma\gamma$ decay will be presented.

A. The Standard Model

The standard model³⁰ is based on a generalization of the first and most successful quantum field theory, quantum electrodynamics or QED.³¹ The quantum field theories of the electroweak and the strong interactions, both components of the standard model, resemble that of QED. The fourth fundamental interaction, gravity, is not included.

The standard model contains all the phenomenology associated with the fundamental interactions. It describes spinless, spin-1/2, and spin-1 particles, or fields, interacting with one another in a manner determined by a function of these fields, the Lagrangian, which contains the dynamics of the system. It is often possible to find fields that can be

transformed without changing the Lagrangian. Those transformations that leave the Lagrangian unchanged, or invariant, are called symmetries. Symmetries imply the existence of conserved quantities that remain constant during particle interactions. The transformation of the fields that leave the Lagrangian invariant also lead to important predictions in the theory.

The Lagrangian of the standard model is based on two basic ideas besides those necessary for a quantum field theory. One is the concept of local symmetry, which determines the form of the interaction between particles that carry the "charge" of the symmetry. Charge, in this case, not only refers to the electric charge, but has the broader definition of that quantity associated with a given symmetry. The second is spontaneous symmetry breaking, where the vacuum, or state with no particles, has a nonzero "charge" distribution. In the standard model, the nonzero weak interaction charge distribution of the vacuum is the source of most of the particle masses.

The idea of local symmetry has become the guiding principle in the construction of the standard model. The requirement that a theory be invariant under local phase transformations implies the existence of a gauge field. The gauge field has a corresponding integral-spin particle, the vector boson or gauge boson, that mediates the interaction between matter fields. The local symmetry determines the exact form of the interaction. The interaction Lagrangian must be of the form, $e J^\mu(x) A_\mu(x)$, where $J^\mu(x)$ is the current density of the "charged" particles and $A_\mu(x)$ is the field of the vector bosons. The coupling constant e is defined as the strength of the interaction between the

current and the vector boson. The idea that all interactions are mediated by vector bosons, or equivalently, that they originate from local symmetries, has been extended from QED to the weak and strong interactions.

The symmetry associated with quantum electrodynamics is classified as a $U(1)$ symmetry, where the symmetry group $U(1)$ is the set of all unitary 1×1 matrices. The Lagrangian of QED is unchanged by the continuous multiplication of the matter field by an arbitrary phase, that is, by the transformation, $\psi_j \rightarrow e^{iQ\Lambda} \psi_j$, where Λ is an arbitrary real number and Q is the electric charge operator or generator of the $U(1)$ group. This operation is known as a global $U(1)$ transformation. A local $U(1)$ transformation is described by $\psi_j \rightarrow e^{iQ\Lambda_j} \psi_j$, where Λ_j is a different value for every ψ_j . The electromagnetic field precisely compensates for the local $U(1)$ phase transformation in QED and the Lagrangian is left invariant. The local $U(1)$ gauge invariance implies that the photon, the gauge particle that mediates the electromagnetic interaction, is massless because the term in the Lagrangian that would describe a massive photon is not itself invariant under local $U(1)$ transformations.

In QED, the electromagnetic field is a boson field that has the quantum numbers of the photon. Electrically charged particles, such as electrons, are also represented by fields and interact via an exchange of virtual photons. A virtual photon momentarily violates energy and momentum conservation as permitted by the Heisenberg uncertainty principle. The electromagnetic gauge field $A_\mu(x)$ interacts with the electrically charged matter fields through the current $J_\mu(x)$. The coupling strength at each vertex of the interaction is e , the charge of

the electron. Since the virtual photon is emitted and absorbed by the charged particles, the probability for this process is of the magnitude e^2 and corresponds to the dimensionless parameter α , where $\alpha = e^2/4\pi\hbar c \approx 1/137$. Because the coupling strength α is small, the amplitude for processes involving such interactions can be approximated by a perturbation expansion on the field theory. The level of accuracy in the calculation of these processes can then be set by including higher order virtual processes. The agreement of QED calculations with experimental results is excellent.³²

The theory of weak interactions started as a phenomenological treatment³³ by Enrico Fermi to interpret data on nuclear beta decay. After a modification of this theory to include parity violation,³⁴ the so-called V - A structure,³⁵ it contained all the necessary components to describe low-energy weak interactions. The form of the interaction amplitude is a product of two charged weak currents, J^+ and J^- , where each current is a product of fermion fields, and is given by $(G_F/\sqrt{2})(J^+ \times J^-)$, and G_F is the Fermi constant. In low-energy weak processes, the two charge-changing currents appear to interact at a single point. This case is very different from that in QED where two currents interact through the exchange of a virtual photon. Also, unlike the constant α in QED, the weak coupling constant G_F is not a dimensionless constant, but has a value of $(293 \text{ GeV})^{-2}$. Experimentally, the Fermi theory gives accurate predictions of phenomena at low energies, but fails at high energies, that is, energies above $\sim 300 \text{ GeV}$, where it predicts a violation of unitarity.

In order to fix the problem at high energies and to make the form of the weak interaction appear more like the QED theory, an intermediary particle, the W vector boson, was proposed.⁸ The nearly zero range of the weak interaction can be explained by giving this particle a large mass. In matrix elements, the W propagator, $\alpha'/(q^2 + M_W^2)$, appears. In this case, α' is a constant and q is the momentum carried by the W boson. At low q^2 values, the propagator reduces to α'/M_W^2 and is proportional to the Fermi constant G_F . The existence of the W boson thus explains the dimensionality of G_F .

Since both the weak and electromagnetic interactions involve electric charge, these two forces may be different versions of the same basic force. If this is the case, then $\alpha' = \alpha$, and G_F would be proportional to α/M_W^2 . A crude estimate of the W boson mass can then be made at about 25 GeV. This prediction of the existence of the W boson and the theoretical unification of the electromagnetic and weak interactions were made many years before the experimental discovery of the W boson.

Quantum electrodynamics is a gauge theory based on a local $U(1)$ invariance, so it is reasonable to assume that a unified electroweak theory is also a gauge theory. However, a special type of gauge theory is used in the description of the electroweak theory. The symmetry group uses non-Abelian, or non-commuting, algebras and must be broken in order to give the gauge particles their masses. The theory associated with this symmetry group is known as a non-Abelian gauge theory, or Yang-Mills theory.³⁶ In a Yang-Mills theory, the gauge fields interact with each other because they carry the "charge" of the symmetry.

One problem in the electroweak unification was in choosing the correct local symmetry group in order to obtain the correct gauge bosons for the theory. The most natural choice for the local symmetry group was $SU(2)$, the group of 2×2 unitary matrices with $\det U = 1$, because the observed quarks and leptons seemed to form doublets under the weak interaction. The three gauge bosons required to compensate for the three independent phase transformations of a local $SU(2)$ symmetry could be identified as the W^+ , W^- , and photon. Unfortunately, this choice is not correct because it gives the wrong electric charges for quarks and leptons. A way to fix this problem is to expand the local symmetry group. A new $U(1)$ symmetry can be added to form the $SU(2) \times U(1)$ symmetry group, and this group does give the correct quark and lepton charges. There is one type of coupling associated with each part of the $SU(2) \times U(1)$ symmetry, a coupling g for the $SU(2)$ symmetry and a coupling $g'/2$ for the $U(1)$ symmetry. The three gauge bosons associated with the $SU(2)$ symmetry group are the W^+ , W^- , and W_3 bosons and the single boson associated with the $U(1)$ symmetry is the B boson.

Another problem before the electromagnetic and weak interactions can be unified is that of renormalizability. Any theory containing the weak interaction must be renormalizable, that is, the divergences found using higher-order corrections must be eliminated. The QED theory is renormalizable, but the Fermi theory is not. The new unified electroweak theory, including massive intermediate vector bosons, is renormalizable as first proven by 't Hooft.³⁷

The addition of the local $U(1)$ symmetry introduces a new uncharged gauge particle into the theory that produces the weak neutral current interactions. This particle is called the Z^0 . The electromagnetic interaction between neutral currents is mediated by the massless photon. These two physical particles, the Z^0 and the photon, are different from the two neutral gauge particles, W_3 and B , associated with the unbroken $SU(2) \times U(1)$ symmetry, and are linear combinations of the neutral gauge particles: $A = B \cos\theta_W + W_3 \sin\theta_W$ and $Z^0 = B \sin\theta_W - W_3 \cos\theta_W$, where θ_W gives the amount of "mixing" of the W_3 and B states. The mixing of $SU(2)$ and $U(1)$ gauge particles to give the physical particles is a result of the $SU(2) \times U(1)$ symmetry being a broken symmetry.

The weak $SU(2)$ symmetry contains both the charged current and neutral current interactions. The form of the charged current interaction is $(g/\sqrt{2}) J_{\pm}^{\mu} W_{\mu}^{\pm}$, where J_{\pm}^{μ} are the charged currents, W_{μ}^{\pm} are two of the massless gauge fields, and $g/\sqrt{2}$ is the coupling constant. The neutral current interaction has the form $g J_3^{\mu} W_{\mu}^3$, where J_3^{μ} is the current, W_{μ}^3 is the third massless gauge field of the $SU(2)$ symmetry, and g is the coupling constant. The weak $U(1)$ symmetry has only one type of interaction of the form, $(g'/2) J^{\mu} B_{\mu}$, where J^{μ} is the current, B_{μ} is the neutral massless gauge field of the theory, and $g'/2$ is the weak $U(1)$ coupling constant.

Another problem in unifying the weak and electromagnetic interactions is that the gauge particles of the weak interaction are massive, yet the photon of QED is massless. The solution to this problem is contained in the concept of spontaneous symmetry breaking of the $SU(2) \times U(1)$ symmetry. This idea is analogous to the symmetry breaking

found in a superconductor.³⁸ In the weak interaction, the symmetry is assumed to be broken by an asymmetry³⁹ of the vacuum due to an ordered state of electrically neutral bosons, the Higgs bosons, that carry the weak charge. They break the symmetry such that the W^\pm and Z^0 bosons obtain masses while the photon remains massless. This asymmetry of the vacuum is referred to as spontaneous symmetry breaking. It does not destroy the symmetry of the Lagrangian, but it does destroy the symmetry of the states. This mechanism allows the electroweak Lagrangian to remain invariant under local symmetry transformations while the gauge particles become massive.

In the spontaneously broken theory, the electromagnetic coupling e is given by $e = g \sin\theta_W$, where $\sin^2\theta_W = g'^2/(g^2 + g'^2)$, so e and θ_W are an alternate method of expressing the couplings g and g' . The value of e was not determined in QED, and likewise the mixing angle θ_W is not determined in the electroweak theory. The quantity θ_W can be measured⁴⁰ experimentally and has the value, $\sin^2\theta_W = 0.220 \pm 0.009$. The theory does predict the masses of the W and Z bosons as $M_W/M_Z = \cos\theta_W$ and $M_W = (\pi\alpha/\sqrt{2}G_F)^{1/2} (\sin\theta_W)^{-1}$, in good agreement with the experimentally observed values.^{40,41}

Quarks and leptons obtain their masses in the electroweak theory by interacting with the ordered vacuum state. However, the values of their masses are not predicted by the theory, but are proportional to arbitrary parameters related to the strength of the coupling of the quarks and leptons to the Higgs boson. The Higgs boson, in the simplest theory, is a complex SU(2) doublet consisting of four real fields. These fields are needed to transform the massless gauge fields into massive ones. In the

electroweak theory, the W^+ , W^- , and Z^0 bosons absorb three of the four Higgs fields in order to obtain their longitudinal spin components, along with the existing two transverse spin components, and in so doing become massive. The remaining Higgs field is not used in this transformation from massless to massive gauge bosons, and could possibly be observed. The mass of this last Higgs particle is not fixed and so far has not been detected experimentally.

The third interaction described by the standard model is the strong interaction. The early theoretical work used the $SU(3)$ symmetry, the group of all unitary 3×3 matrices U , satisfying $U^\dagger U = I$ and $\det U = 1$, to group the hadrons, those particles that experience the strong force, into multiplets. The finite size⁴² of the hadrons, compared to the leptons, compelled a change to the theory in postulating the existence of quarks. All hadrons at this time could be composed of three types of quarks, designated u , d , and s for up, down, and strange quarks, respectively. The $SU(3)$ symmetry comes from the invariance of the Lagrangian of the strong interaction to rotations among the three quarks. It is an exact global symmetry only if the quarks have identical masses. Since this is not true, the $SU(3)$ symmetry must be a broken global symmetry. It is unknown how the quarks obtain their masses.

Another type of $SU(3)$ symmetry, quite different from the previous one, possessed by the quarks is a local symmetry associated with the strong interaction and is called the $SU(3)$ symmetry of color. The idea of a new degree of freedom⁴³ possessed by each quark, called color, was invented to solve the contradiction of a baryon state violating the Pauli exclusion principle. The theory resulting from this symmetry is called

quantum chromodynamics or QCD, and it is also a Yang-Mills gauge theory. The equivalent of electric charge in QED becomes three "color" charges in QCD. The three color types of each quark form a triplet under the SU(3) local gauge symmetry. A local phase transformation of the quark field can rotate the color and change one colored quark into another. The local gauge transformation of QED only changes the phase of a particle, but in QCD, the color transformations change the entire particle. The local symmetry of the strong interaction theory is not spontaneously broken, unlike the electroweak theory.

The SU(3) local color symmetry implies the existence of eight massless gauge fields, or spin-1 vector bosons, to compensate for the eight independent local symmetry transformations of the colored quark fields. These eight gauge bosons of QCD are called "gluons," and they carry the color charge of the strong interaction. The interaction of quarks and gluons is given by $g_s J_a^\mu A_a^\mu$, where J_a^μ are the quark currents, A_a^μ are the gluon fields, and g_s is the strong interaction coupling constant.

Gluons are similar to photons in that they are massless spin-1 particles mediating a force, but differ in that gluons themselves carry the charge of the symmetry, in this case, color. This difference is a result of the color force being a Yang-Mills theory and consequently gluons, unlike photons, can interact among themselves. There are two types of gluon self-interactions, one in which three gluons interact at a single point and whose strength is proportional to g_s , and the other in which four gluons interact at a single point and whose strength is proportional to g_s^2 . The effect of the gluon self-interactions is to

weaken the force of the color charge at short distances, unlike screening effects in QED, which weaken the effective electric charge at long distances. The weakening of the color charge at short distances goes by the name, asymptotic freedom, and it allows the use of perturbation theory to calculate hadronic processes at high momentum transfers, that is, above a few GeV/c , which implies short distances. This phenomenon explains why hadrons at high energies behave as if they were made of almost free quarks, even though no free quark has been experimentally observed.⁴⁴

The gluon self-interaction also explains the apparently permanent confinement of quarks. At long distances it produces a large number of virtual gluons such that the color charge effectively grows without limit, and therefore allows no colored particles to propagate. Only color-neutral states of hadrons, which are complicated composites of quark and gluon fields, are not susceptible to this confinement, thus making the observable hadrons colorless. The baryons are composed of three quarks in equal portions of the three colors, and the mesons are made of quark-antiquark pairs in equal portions of color and anticolor. Unconfined massless gluons, like massless photons, would give rise to a long-range, Coulomb-like force. This is not experimentally observed, therefore the quarks and gluons are believed to be confined.

QCD has been very successful in explaining the short range or high energy behavior of quarks using the strong interaction. Experimental evidence comes from the deep inelastic scattering experiments.⁴⁵ There remains some uncertainty in the QCD theory at low energies where perturbation theory cannot be used. A simple relationship between the

quark and gluon fields and the observed structure of the hadrons does not seem to exist. The quark model has not been rigorously derived from QCD.

Quarks and leptons are believed to be the basic building blocks from which all the observed particles, hence all matter, are made. Leptons are fermions that come in different types, or "flavors," that consist of the electron e , muon μ , and tau lepton τ , along with their associated neutrinos, ν_e , ν_μ , and ν_τ . There also exist antileptons corresponding to each of the lepton flavors. The e , μ , and τ have an electric charge of $-1e$, where e in this context is the electric charge of the electron ($e = 1.6 \times 10^{-19}$ C), and all experience the electromagnetic and weak interactions. The neutrinos do not possess any electric charge and interact only through the weak force. Leptons do not appear to have any inner structure and are point-like particles to distances as small as 10^{-16} cm.⁴⁶ Quarks are fermions that also come in different flavors, up u , down d , strange s , charm c , bottom b , and top t . There is experimental evidence for the existence of the first five quark flavors within hadrons, while the sixth quark is conjectured to exist. The u , c , and t quarks have a fractional charge of $+2/3 e$, and the d , s , and b quarks have an electric charge of $-1/3 e$. The quarks with a common electric charge seem to be distinguished from each other only by their different masses. Quarks experience the strong color, weak, and electromagnetic interactions.

The quarks and leptons can be grouped together into representations, similar to the grouping of elements done by Mendeleev, as shown in Table 4. In the table, the symbols d' , s' , and b' indicate weak interaction mass eigenstates, which are mixtures of the strong interaction

TABLE 4
REPRESENTATIONS OF LEPTONS AND QUARKS IN THE STANDARD MODEL

<u>Family</u>	<u>Leptons</u>	<u>Quarks</u>
I	$\begin{pmatrix} e^- \\ v_e \end{pmatrix}_L$ $(e^-)_R$	$\begin{pmatrix} u_r \\ d'_r \end{pmatrix}_L, \quad \begin{pmatrix} u_b \\ d'_b \end{pmatrix}_L, \quad \begin{pmatrix} u_g \\ d'_g \end{pmatrix}_L,$ $(u_r)_R, \quad (u_b)_R, \quad (u_g)_R,$ $(d_r)_R, \quad (d_b)_R, \quad (d_g)_R$
II	$\begin{pmatrix} \mu^- \\ v_\mu \end{pmatrix}_L$ $(\mu^-)_R$	$\begin{pmatrix} c_r \\ s'_r \end{pmatrix}_L, \quad \begin{pmatrix} c_b \\ s'_b \end{pmatrix}_L, \quad \begin{pmatrix} c_g \\ s'_g \end{pmatrix}_L,$ $(c_r)_R, \quad (c_b)_R, \quad (c_g)_R,$ $(s_r)_R, \quad (s_b)_R, \quad (s_g)_R$
III	$\begin{pmatrix} \tau^- \\ v_\tau \end{pmatrix}_L$ $(\tau^-)_R$	$\begin{pmatrix} t_r \\ b'_r \end{pmatrix}_L, \quad \begin{pmatrix} t_b \\ b'_b \end{pmatrix}_L, \quad \begin{pmatrix} t_g \\ b'_g \end{pmatrix}_L,$ $(t_r)_R, \quad (t_b)_R, \quad (t_g)_R,$ $(b_r)_R, \quad (b_b)_R, \quad (b_g)_R$

mass eigenstates d, s, and b. The amount of mixing of the quark states is given in a 3×3 unitary matrix, called the Kobayashi-Maskawa matrix,⁴⁷ and includes several parameters, three mixing angles and an arbitrary phase. It is customary to assign all the mixing effects to the d, s, and b quarks, leaving the u, c, and t quarks unchanged. Also in the table, the subscripts r, b, and g on the quark flavors denote the three color charges of the quarks, and the subscripts R and L denote the right- and left-chiral projections. A massless particle is said to be left-handed (right-handed) if the direction of its spin vector is opposite (the same as) that of its momentum vector. Chirality is the Lorentz-invariant generalization of this handedness to massive particles and is equivalent to handedness for massless particles. The exact chiral symmetry of the strong and electromagnetic gauge interactions and the nonzero masses of the quarks and charged leptons implies that the right-chiral quarks and charged leptons and their left-chiral partners can be treated as single objects under these interactions. A neutral singlet has no weak charge, so the right-handed component of the neutrino is unaffected by the weak, electromagnetic, and strong interactions. Therefore it can be neglected and does not appear in the table. Each neutral lepton is associated with a specific charged lepton because of the weak SU(2) symmetry. It is easier to think in terms of three sets of quarks, u and d, c and s, and t and b, and three sets of leptons, e and ν_e , μ and ν_μ , and τ and ν_τ , than 33 separate representations. The grouping of the quark and lepton sets by the mass of the particles can be made into quark-lepton families, with each family consisting of successively heavier mass quark and lepton sets. Similar to Mendeleev's situation with the elements, it appears that there

is an orderly grouping of quarks and leptons, but there is no clue as to its basis.

The description of the standard model presented here includes only those features necessary to explain observed phenomena, so this model is called the minimal standard model. An example of this is the low-energy charged-current weak interactions, which are dominated by $V - A$ (vector - axial vector) currents, so only left-chiral W^\pm bosons have been included in the theory. Another example is that neutrinos are assumed to be massless, hence there will be no "oscillations" between different neutrino flavors. The most important assumption given by the minimal standard model relevant to rare muon decays is the exact conservation of individual lepton flavors, that is, electron number, muon number, and tau number, and also the conservation of total lepton number.

B. The Standard Model and Lepton-Family Number Nonconservation

The standard model has been remarkably successful in describing the world of particle physics. One of its major successes is that its description, using symmetries, gauge fields, and invariant Lagrangians, correctly accounts for all elementary particle data. All known experimental phenomena is incorporated somewhere within the standard model. Another success is the mathematical structure of the theory. There are no known mathematical inconsistencies up to very high energies and all components of the standard model use a similar mathematical treatment, namely, the local gauge theories, derived from the principle of local symmetry.

Some of the more notable theoretical successes leading to the development of the standard model have been the prediction⁴⁸ of the Ω^- particle from the SU(3) quark symmetry, the Cabibbo hypothesis⁴⁹ of the mixing of quark states in the weak SU(2) symmetry, the GIM (Glashow - Iliopoulos - Maiani) mechanism,⁵⁰ that generalizes Cabibbo's hypothesis of the hadronic weak interactions, and finally, perhaps the greatest theoretical triumph, the unification³⁰ of the electromagnetic and weak interactions. Some of the experimental successes leading to the development of the standard model have been evidence⁴⁵ for quarks in deep inelastic scattering experiments, the discovery¹⁰ of the τ lepton, the discovery of the c quark⁵¹ and b quark⁵² within new particle states, the discovery⁵³ of the weak neutral current reactions, and, perhaps the greatest experimental triumph, the discovery⁵⁴ of the W^\pm and Z^0 bosons.

Despite all the great predictions, discoveries, and success of the standard model, some difficulties and unanswered questions remain. One is the problem of symmetry breaking. Is the Higgs particle responsible for the mechanism of giving the W^\pm and Z^0 gauge bosons their masses while leaving the photon massless? The procedure of generating the masses of the quarks and leptons is still a mystery. Another question deals with the large number of seemingly arbitrary parameters in the theory. Why are there so many coupling constants and mixing angles, such as θ_W , and can the values only be obtained from experiments? A question confronts the number of observed particles, and the similarities and differences between the quarks and leptons. Are quarks and leptons truly elementary particles, or composites of something more fundamental? Why are there three families of quarks and leptons? Could there be more? The fourth

fundamental interaction is gravity and it has not been incorporated into the standard model. How does gravity fit into the picture? With all of these open questions, it is believed that the standard model may be incomplete.

In the previous section, there was a discussion describing the families of quarks and leptons and these were listed in Table 4. If the decays of the unstable particles are studied, an asymmetry can be observed in the decay patterns. Quarks are observed to change families, but leptons do not. For example, the positive kaon can decay into a positive and a neutral pion, $K^+ \rightarrow \pi^+ \pi^0$. At the quark level, the \bar{s} quark from the K^+ transforms into the \bar{d} quark in the π^+ . The \bar{s} quark is a member of the second family of quarks, while the \bar{d} quark is a member of the first family. The leptons do not seem to follow this pattern. For example, in the normal muon decay, $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$, the μ^- transforms into a ν_μ , and simultaneously the e^- and $\bar{\nu}_e$ appear. The $\bar{\nu}_e$ carries the opposite of whatever family quantum numbers distinguish an e^- from any other charged lepton. Therefore no net "first-familiness" is created and the "second-familiness" of the original μ^- is preserved in the ν_μ . The puzzling question is then, Why can quarks change families while leptons do not? The answer to this question remains a mystery.

Symmetry principles provide important guidelines in constructing theories used in particle physics. This was seen throughout the previous section in building the standard model. There is an important relationship between symmetries and conservation laws, as given in Noether's theorem:⁵⁵ whenever a physical process is invariant under a certain symmetry, there exists a corresponding conservation law.

Likewise, if there appears to be a conservation law at work, a symmetry can usually be found to go along with it.

There exist two types of conservation laws, those related to space-time symmetries and those that are not. The conservation of energy, momentum, and angular momentum are all examples of the first type. The conservation of energy is related by Noether's theorem to the invariance of time displacements; conservation of momentum follows from the invariance of spatial translations. Examples of the second type of conservation laws are conservation of electric charge, lepton number, and baryon number. The first type of conservation law is considered to be fundamental. In fact, Pauli⁵⁶ postulated a new particle, the neutrino, rather than give up conservation of energy in nuclear beta decay. The second type of conservation law is not as dogmatic as the first. The conservation of electric charge is related to the gauge invariance of the exact U(1) symmetry of the electromagnetic field and it has associated with it a massless gauge boson, the photon.

Other examples of the second type are the proposed conservation of lepton number and baryon number. The idea of baryon or lepton number conservation can be associated with a global U(1) symmetry. For example, baryon number conservation is associated with a U(1) phase transformation of all baryon fields by an amount $e^{i\Lambda B}$, where Λ is an arbitrary number and B is the baryon number with values, $B = 1$ for protons and neutrons, $1/3$ for quarks, and 0 for leptons. Analogously, electron number is conserved in muon or beta decay if the field of the ν_e is assigned the same electron number as the field of the electron, and all other fields are assigned an electron number of zero. The same holds true for muon number and tau

number, as given in Table 2. An exact global U(1) symmetry seems to operate on each type of lepton. However, this symmetry is not expected to be exact because that would imply the existence of an additional fundamental force mediated by a massless vector boson. This is the basis of a heuristic argument by de Rújula, Georgi, and Glashow⁵⁷ that lepton-family number is not expected to be an exact global symmetry. There is no widely accepted experimental evidence for the presence of an additional force or massless gauge boson, but this is not proof enough. If we assume such a new force exists, it should add to the gravitational force and therefore be observed in Eötvös-like experiments. Lee and Yang⁵⁸ argued that such a massless gauge boson coupled to lepton family number would violate the gravitational equivalence principle. Lepton family number then is not expected to be exactly conserved, however the experimental evidence shows that any nonconservation of lepton family number must be extremely feeble.

C. Extensions to the Standard Model

The previous section shows that there are compelling reasons for experimentalists to search for the nonconservation of lepton-family number. Different extensions to the standard model attempt to deal with its shortcomings and many of these incorporate lepton-family number nonconservation. Therefore searches for rare decays not conserving this quantity test for possible new physics discoveries beyond that incorporated in the standard model. A number of these extensions are the grand unified theories, or GUTs, composite models, supersymmetry,

technicolor, and left-right symmetric theories. None of these extensions gives truly satisfactory answers to all the unanswered questions of the standard model, but all could provide a basis for the next fundamental theory. More input data, such as that given by rare decays, will help point to the correct picture of particle physics.

Since the standard model explains all observed particle phenomena, no data exists that require extensions. It is difficult for any of these extensions to give predictions of the rates of rare decays because of unknown parameters, such as particle masses and mixing angles. There are many possible rare decay modes or channels through which lepton-family number nonconservation may be seen; examples were given in Table 2. In the various model extensions, different rare decay modes become dominant, due to that particular model's mechanism of nonconserving lepton-family number. Little theoretical guidance is available in making predictions as to which particular rare decay mode is the best to experimentally observe.

A list⁵⁹ of some extensions to the standard model is presented in Table 5. The weaknesses of the standard model are arranged in the first column, the extension to the standard model that deals directly with that problem are in the second column, and the decay modes likely to have the highest rate for a given extension are given in the third column. The third column takes account of the phase space available for the decay, the lifetime of the particle, and the number of vertices in the corresponding Feynman diagram. In this way, an estimate can be placed on the rate of a given decay mode relative to other modes. It becomes quite apparent that only a search for a variety of rare decay modes will restrict the many different models.

TABLE 5

WEAKNESSES AND POSSIBLE EXTENSIONS TO THE STANDARD MODEL,
AND THE POSSIBLE DECAY MODES THAT CAN BE OBSERVED

Theoretical weakness in the standard model	Model extension	Likely rare decay modes
Unknown number of parameters	4 or more generations	$\mu N \rightarrow eN$, $\mu \rightarrow e\gamma$
$V - A$	Left-right symmetry	$\mu \rightarrow e\gamma$, $\mu N \rightarrow eN$
More than one Higgs doublet	Extra Higgs - $\Delta S = 0$	$\mu N \rightarrow eN$, $\mu \rightarrow e\gamma$
	Extra Higgs - $\Delta S \neq 0$	$\mu N \rightarrow eN$, $K_L^0 \rightarrow \mu e$
Gravity	Supersymmetry	$\mu \rightarrow e\gamma$
Too many elementary particles	Composite models	$\mu \rightarrow e\gamma$, $\mu \rightarrow e\gamma\gamma$, $K_L^0 \rightarrow \mu e$
Many undetermined parameters	Horizontal gauge theory	$K_L^0 \rightarrow \mu e$, $\mu \rightarrow eee$
Replication of families	Horizontal gauge theory	$K_L^0 \rightarrow \mu e$, $\mu \rightarrow eee$
Elementary scalar field	Technicolor	$K_L^0 \rightarrow \mu e$, $\mu N \rightarrow eN$
Not a unified theory	GUTs	According to model

A closer look at some of the different models allowing lepton-family number nonconserving decays will now be presented. Exact predictions from the extensions may not be available because of unknown parameters existing in a given model.

The easiest extension⁶⁰ would be to add massive neutrinos into the theory of the standard model. These massive neutrinos would mix and give the lepton-family number nonconserving decays such as $\mu \rightarrow e\gamma$. The branching ratio is given by

$$B_{\mu \rightarrow e\gamma} = \frac{3\alpha}{32\pi} \left| \sum U_{\mu i}^* U_{ei} \frac{m_i^2}{m_W^2} \right|^2, \quad (\text{II-1})$$

where m_i is the mass of the neutrino in the i th generation and U_{ij} is an element of the unitary neutrino mass mixing matrix. If present limits for these values are placed into the formula, the $\mu \rightarrow e\gamma$ branching ratio is conservatively placed at less than 10^{-18} . If a fourth generation of particles is found, and if there was maximal mixing including a fairly heavy neutrino, the branching ratio would be around the present experimental limits.⁶¹

In the above equation, a factor of m_i^4/m_W^4 is included in the branching ratio equation, where m_i is a small mass or small mass difference and m_W is a large mass. This ratio gives a "natural" suppression of the rate for a process, and can possibly account for the very low rates in neutrinoless muon decays. Also, since the branching ratio goes as the fourth power of a particular mass scale, any improvement

in the branching ratio can make only modest improvements in the mass scale.

Other extensions⁶² of the minimal standard model have introduced doubly charged leptons in left-handed triplets or right-handed doublets together with an enlarged set of Higgs bosons. The standard model requires at least one doublet of Higgs particles in order to generate masses for the vector bosons. Enlarging the number of doublets of Higgs bosons can cause lepton-family number nonconservation.⁶³ The two-loop Feynman diagram contribution is larger in this case than the one-loop diagrams, and causes the $\mu \rightarrow e\gamma$ decay rate to be larger than the $\mu \rightarrow eee$ rate. This model contains lepton-family number nonconservation without neutrino oscillations.

The V - A structure³⁵ of weak interactions is not a result of the standard model, but is placed into the theory as an empirical result. Left-right symmetric models,⁶⁴ constructed using the symmetry group, $SU(2)_L \times SU(2)_R \times U(1)$, help explain parity violations in weak interactions and make the left-chiral and right-chiral multiplets balanced. There are now four charged gauge bosons, W_L^+ , W_L^- , W_R^+ , W_R^- , and three neutral bosons, γ , Z_1^0 , Z_2^0 , mediating the electroweak interactions. The extra gauge bosons couple to right-chiral currents and must have a mass of at least 200-400 GeV in order to retain the good features of the minimal standard model. Parity violation is generated dynamically by a spontaneous symmetry breaking of some higher symmetry. Neutrinos can have mass and can cause lepton-family number nonconservation in a similar manner to heavy left-chiral neutrinos.

Horizontal symmetries⁶⁵ have been introduced to understand the number of families. These models bring forth a large number of symmetry groups. In most, the weak symmetry group $SU(2) \times U(1)$ is a subgroup of one of these larger flavor symmetry groups. The larger groups have new gauge bosons with masses larger than the mass of the W^\pm or Z^0 bosons, and therefore can suppress lepton-family number nonconservation to appropriate levels. The mass scale is believed to be on the order of 100 TeV. The decay $\mu \rightarrow e\gamma$ is a radiative correction in these models and would occur at an even lower rate than processes such as $\mu \rightarrow eee$, $K_L^0 \rightarrow \mu e$, or $\mu N \rightarrow eN$.

In the standard model, the weak $SU(2) \times U(1)$ symmetry is broken and masses are given by coupling new scalar bosons, the Higgs bosons, to gauge bosons and fermions. There is some difficulty in coping with this seemingly ad hoc method and with the scalar bosons, so new approaches were developed, namely, technicolor and supersymmetry. These attempts achieve a dynamical symmetry breaking without Higgs scalars, but still generate the masses of the gauge bosons.

Technicolor⁶⁶ uses new fundamental fermions and a new gauge interaction, technicolor, analogous to color QCD, rather than having fundamental scalar Higgs bosons. The bound state "pions" of technicolor are the "Higgs bosons" of the standard model. Unfortunately, technicolor is not able to give masses to the quarks and leptons, therefore it becomes necessary to expand this model requiring another interaction, called extended technicolor.⁶⁷ This interaction has new massive bosons, called ETC bosons, that are constrained in a manner producing the known fermion masses. The new bosons will couple different flavors together, such as the regular fermions with technifermions, and flavor-changing decays are

possible. Since the masses of the new bosons, which are on the order of 10-100 TeV, are not free parameters, an estimate of rare flavor-changing decay rates is on the order of or smaller than present limits. Technicolor theories include as bound states both neutral pseudoscalar bosons and leptoquark bosons that allow flavor-changing decays. Because the mass scale is not very flexible in these theories and experimental bounds may be lower than the definite predictions, the ideas motivating this type of model extension are the important features.

Supersymmetric theories⁶⁸ offer a different approach in dealing with the scalar Higgs bosons in that they treat the scalar bosons on equal terms as the fermions, that is, as fundamental particles. Gravity, the fourth fundamental interaction, may also be included in the supersymmetric theories. In the theory, every known fundamental particle has a supersymmetric partner associated with it, differing in spin by 1/2 unit of \hbar . Some of these partners have light masses and can be included final state products, while others are internal to Feynman diagrams, producing additions to several of the rare decays. These additions may be forbidden or suppressed, depending on the relations between the masses and couplings. The appearance of lepton-family number nonconserving processes comes from the internal intermediate states of the supersymmetric partners. At present, the most stringent limit⁶⁹ on mass differences of the supersymmetric partners of the leptons, the sleptons, comes from the decay $\mu \rightarrow e\gamma$.

The problems with the Higgs particles, the number of families, and also the growing number of fundamental particles, have led to the development of composite models.⁷⁰ Here, the fermions, and possibly the

Higgs and vector bosons, are composite states of a very few fundamental constituents. Composite models must be able to explain how fermions can be bound states of massive fundamental constituents yet remain point-like in size and how these constituents can produce Dirac magnetic moments of the fermions. There are several methods of obtaining rare decays in composite models by rearranging the constituents appropriately. One method includes excited quark and lepton states whose decays are suppressed because of the large masses of the constituents. In a particular composite model, the muon is an excited 2S state of the electron and can decay by both one photon, $\mu \rightarrow e\gamma$, and two photon, $\mu \rightarrow e\gamma\gamma$, transitions.

There have been attempts to unify the strong and electroweak interactions into one basic force. These ideas are known as grand unified theories⁷¹ or GUTs. The simplest are contained in the symmetry group SU(5). The coupling constants of the three interactions merge at an energy of 10^{15} GeV, and it is believed that no new interesting experimental discoveries lie between a few hundred GeV and this ultrahigh energy of 10^{15} GeV. In the theory, neutrinos are massless, the V - A structure is still added empirically, and the conserved quantity of fermions is $B - L$, the baryon number minus the lepton number. A discrepancy with the experimental limit⁷² on the lifetime of the proton exists in the SU(5) model, so a higher symmetry group,⁷³ SO(10), along with the exceptional symmetry group,⁷⁴ E_6 , were developed. After unifying the three basic interactions, the problem of allowing flavor-changing processes still remains because all the GUTs are single family theories and still contain undetermined parameters. Therefore, the GUTs are as yet

incomplete, and so other methods, such as horizontal interactions between families or compositeness, are needed to fill in the shortcomings.

In concluding this section, a comment by Kane and Shrock⁷⁵ seems to be appropriate for which rare decay searches are best.

We would like to stress that experimentalists should search for as wide a range of new physics as possible and should not be deterred from considering a particular decay mode just because some theorist claims that it is predicted not to occur in his set of fashionable models. New physics has sometimes been anticipated by theorists, but at least as often, it has not. Moreover, theorists may not correctly assess the experimental implications of a given model.

D. Model Extensions Relevant to $\mu \rightarrow e\gamma\gamma$

Dreitlein and Primakoff⁷⁶ gave the first expression of the decay rate for $\mu \rightarrow e\gamma\gamma$, based on a local, Lorentz-invariant and gauge-invariant Lagrangian. The matrix element associated with a two photon neutrinoless decay uses an effective Lagrangian, $L(x)$, given by

$$\langle e^\pm \gamma_1 \gamma_2 | \int d^4x L(x) |\mu^\pm \rangle = (2\pi)^4 \delta^4(\vec{p}_{\gamma_1} + \vec{p}_{\gamma_2} + \vec{p}_e - \vec{p}_\mu) \times \langle e^\pm \gamma_1 \gamma_2 | L(0) |\mu^\pm \rangle , \quad (II-2)$$

with

$$L(x) = m_\mu^{-3} \left[g_1 \vec{E}(x) \cdot \vec{B}(x) + (g_2/2)(\vec{E}^2(x) - \vec{B}^2(x)) \right] \left[\bar{\psi}_e(x) \psi_\mu(x) \right] + \text{Hermitian conjugate} , \quad (II-3)$$

where $\vec{E}(x)$ is the photon electric vector, $\vec{B}(x)$ is the photon magnetic vector, $\psi_e(x)$ is the electron spinor quantized field amplitude, $\psi_\mu(x)$ is

the muon spinor quantized field amplitude, and g_1 and g_2 are dimensionless coupling constants. The effective Lagrangian $L(x)$ has the two spinor field amplitudes evaluated at a single space-time point x , therefore $L(x)$ is not the most general effective Lagrangian for the two photon neutrinoless decay.

The decay rate for $\mu^\pm \rightarrow e^\pm \gamma\gamma$ is:

$$\begin{aligned}
 \Gamma(\mu^\pm \rightarrow e^\pm \gamma\gamma) &= 2\pi \int \frac{1}{2} \frac{d\vec{p}_1}{(2\pi)^3} \frac{d\vec{p}_2}{(2\pi)^3} \delta(|\vec{p}_1| + |\vec{p}_2| + |\vec{p}_1 + \vec{p}_2| - m_\mu) \\
 &\quad \times (1/2) \sum | \langle e^\pm \gamma_1 \gamma_2 | L(0) | \mu^\pm \rangle |^2 \\
 &= 2\pi \int \frac{1}{2} \frac{d\vec{p}_1}{(2\pi)^3} \frac{d\vec{p}_2}{(2\pi)^3} \delta(|\vec{p}_1| + |\vec{p}_2| + |\vec{p}_1 + \vec{p}_2| - m_\mu) \\
 &\quad \times \frac{g_1^2 + g_2^2}{m_\mu^6} \left[(1 - \vec{p}_1 \cdot \vec{p}_2) \left(\frac{|\vec{p}_1| |\vec{p}_2|}{4} \right) \right] \\
 &= \frac{(g_1^2 + g_2^2)m_\mu}{128 \times 15\pi} ,
 \end{aligned} \tag{II-4}$$

with a photon-photon angular correlation $\sim (1 - \vec{p}_1 \cdot \vec{p}_2)^2$, a photon momentum distribution $\sim x^3 (1 - 3x/4)$ with $xm_\mu/2 = |\vec{p}_1|$ or $|\vec{p}_2|$, where $0 \leq x \leq 1$, and an electron momentum distribution $\sim y^2(1 - y)^2$ with $ym_\mu/2 = |\vec{p}_e| = |\vec{p}_1 + \vec{p}_2|$, where $0 \leq y \leq 1$.

Bowman, Cheng, Li, and Matis²⁶ present a more general formalism for the $\mu \rightarrow e\gamma\gamma$ decay. The lepton-family number nonconserving amplitude may be described by a local Lagrangian density, L_{eff} , in the case where there is an intermediate particle with a mass much larger than that of the muon. The most general local interaction is given by

$$\begin{aligned}
 m_\mu^3 L_{\text{eff}} = & \bar{\psi}_e (a_S + b_S \gamma_5) \psi_\mu F^{\alpha\beta} F^{\alpha\beta} \\
 & + \bar{\psi}_e (a_p + b_p \gamma_5) \psi_\mu F^{\alpha\beta} \tilde{F}^{\alpha\beta} \\
 & + m_\mu^{-1} \bar{\psi}_e (a_V + b_V \gamma_5) \gamma^\sigma \psi_\mu F^{\alpha\beta} \frac{\partial}{\partial x_\beta} F^{\alpha\beta} \\
 & + m_\mu^{-1} \bar{\psi}_e (a_A + b_A \gamma_5) \gamma^\sigma \psi_\mu F^{\alpha\beta} \frac{\partial}{\partial x_\beta} \tilde{F}^{\alpha\beta} , \tag{II-5}
 \end{aligned}$$

where m_μ is the muon mass, $F^{\alpha\beta}$ is the electromagnetic field tensor with its dual, $\tilde{F}^{\alpha\beta} = \epsilon^{\alpha\beta\gamma\delta} F^{\gamma\delta}/2$, and a_i and b_i are dimensionless and energy independent coupling constants. All fields are evaluated at the same space-time point. This local Lagrangian is an approximation to the actual $\mu \rightarrow e\gamma\gamma$ interaction up to small correction terms, of the order m_μ/M_i , where M_i are the masses of the intermediate particles. Because of the symmetry properties of the indices, the tensor interactions disappear. The case where a_S and a_p are nonzero is given by the equations of Dreitlein and Primakoff. The calculated differential decay distribution for $\mu \rightarrow e\gamma\gamma$ is

$$\frac{d^2\Gamma}{dE_1 dE_2} = \frac{G(a_1, b_1)}{16\pi^3 m_\mu^6} E_e E_1^2 E_2^2 (1 - \cos\theta)^2 , \tag{II-6}$$

where $E_{1,2}$ are the photon energies, $E_e = m_\mu - E_1 - E_2$ is the electron energy, and the angle between the photons is given by $\cos\theta = (E_e^2 - E_1^2 - E_2^2)/2E_1 E_2$. The value of G is

$$G(a_i, b_i) = (4a_S + a_V)^2 + (4b_S - b_V)^2 + (4a_P + 2a_A)^2 + (4b_P - 2b_A)^2. \quad (II-7)$$

These equations give the most general case of the decay $\mu \rightarrow e\gamma\gamma$ that proceeds through an effective local interaction.

Any theoretical model that allows the decay $\mu \rightarrow e\gamma$ will also allow the decay $\mu \rightarrow e\gamma\gamma$. The extra photon in the $\mu \rightarrow e\gamma\gamma$ decay can always be bremsstrahlung of a muon or electron, so the $\mu \rightarrow e\gamma\gamma$ decay rate would be expected to be suppressed by a factor of $\sim \alpha/\pi$ from the $\mu \rightarrow e\gamma$ decay rate. However, there are extensions to the standard model that produce a larger expected rate for $\mu \rightarrow e\gamma\gamma$ than for $\mu \rightarrow e\gamma$. In particular, this occurs for gauge theories in which the $\mu \rightarrow e$ transition is suppressed by the GIM mechanism and the mediating heavy leptons are charged. Another occurrence of an enhanced $\mu \rightarrow e\gamma\gamma$ rate is possible in some composite models.

An example of the first case is the $SU(2) \times U(1)$ theory of Wilczek and Zee,⁶² where two doubly-charged heavy leptons are added to the electron and muon families. The leptons are grouped in triplets,

$$\begin{pmatrix} v_e \\ e \\ h_e \end{pmatrix}_L, \quad \begin{pmatrix} v_\mu \\ \mu \\ h_\mu \end{pmatrix}_L,$$

with $h_e = h_1 \cos\theta + h_2 \sin\theta$ and $h_\mu = -h_1 \sin\theta + h_2 \cos\theta$. The h_1 and h_2 are two doubly-charged heavy leptons with masses m_1 and m_2 . The $\mu \rightarrow e\gamma$ and $\mu \rightarrow e\gamma\gamma$ amplitudes are predicted by this model to be

$$T_{\mu e\gamma\gamma} \sim e G_F m_\mu \cos\theta \sin\theta (m_1^2/M_W^2 - m_2^2/M_W^2) \\ \times [\bar{\psi}_e (1+\gamma_5) \sigma^{\lambda\rho} \psi_\mu] k^\lambda \epsilon^\rho, \quad (II-8)$$

and

$$T_{\mu e\gamma\gamma} \sim e^2 G_F \cos\theta \sin\theta \left[f(m_1^2/k_1 \cdot k_2) - f(m_2^2/k_1 \cdot k_2) \right] \\ \times \left[\bar{\psi}_e (1+\gamma_5) \gamma^\lambda \psi_\mu \right] \left[(k_1 \cdot k_2) \epsilon^{\alpha\rho\sigma\lambda} (k_1 - k_2)^\alpha \right. \\ \left. + k_1^\sigma \epsilon^{\alpha\beta\rho\lambda} k_1^\alpha k_2^\beta - k_2^\rho \epsilon^{\alpha\beta\sigma\lambda} k_1^\alpha k_2^\beta \right] \epsilon_1^\rho \epsilon_2^\sigma, \quad (II-9)$$

where the k_i are photon momenta and ϵ_i are the photon polarizations. The amplitude $T_{\mu e\gamma\gamma}$ corresponds to the a_A and b_A local interactions given previously. The $f(X_i)$ in $T_{\mu e\gamma\gamma}$ are complicated functions that can be evaluated numerically,⁷⁷ and, for a certain range of values of $X_i = m_i^2/k_1 \cdot k_2$, can lead to a less severe GIM suppression for $\mu \rightarrow e\gamma\gamma$, by the factor m_μ^2/M_W^2 than in the amplitude $T_{\mu e\gamma}$. Hence, there is a possibility that

$$\frac{\Gamma(\mu \rightarrow e\gamma\gamma)}{\Gamma(\mu \rightarrow e\gamma)} \sim \frac{\alpha}{\pi} \frac{M_W^4}{(m_1^2 - m_2^2)^2}, \quad (II-10)$$

indicating that the $\mu \rightarrow e\gamma\gamma$ decay rate may be larger than the $\mu \rightarrow e\gamma$ decay rate. This case looks similar to the strangeness-changing neutral current process, where $\Gamma(K_L^0 \rightarrow \gamma\gamma) \gg \Gamma(K_L^0 \rightarrow \mu^+\mu^-)$, even though simple arguments

would lead to the conclusion that these decays would be of equal magnitude.

In a composite model⁷⁸ where leptons and quarks are S states of spin-1/2 and spin-0 particles, the observed leptons e, μ , and τ are the radial-excitation states 1S, 2S, and 3S states, respectively. This situation can be compared to the hydrogen atom, where the 2S state decays to the ground state in two ways.⁷⁹ The predominant decay mode is via a two photon electric dipole emission with a decay time of 1/7 second, and the second decay mode is a one photon emission by a magnetic dipole transition with a slow decay time of about two days. The analogous processes for composite leptons in which the photons are emitted by pure electrodynamic transitions are the decays $\mu \rightarrow e\gamma\gamma$ and $\mu \rightarrow e\gamma$. If the hydrogen 2S metastable state transitions are used as an example, then the $\mu \rightarrow e\gamma\gamma$ decay can occur at a higher rate than the $\mu \rightarrow e\gamma$ decay. The decay rate for $\mu \rightarrow e\gamma\gamma$ is

$$\Gamma(\mu \rightarrow e\gamma\gamma) = \frac{\alpha^2}{12\pi} \omega^3 \left(\frac{1}{M\alpha_c} \right)^4 I, \quad (\text{II-11})$$

where M is the reduced mass of the constituent particles, $\omega = m_\mu - m_e$ is the excitation energy, and I is a parameter with values between 4.4 and 7.1. The quantity $(M\alpha_c)^{-1}$ is a length scale, and the fourth power of that quantity comes from each electric dipole emission contributing a factor $e/M\alpha_c$ in the amplitude. The same arguments can be applied to the decay rate for $\mu \rightarrow e\gamma$,

$$\Gamma(\mu \rightarrow e\gamma) = \frac{128}{91} \alpha \beta^4 \omega^3 \left(\frac{1}{M\alpha_c} \right)^2, \quad (\text{II-12})$$

where β is the relativistic factor, v/c , and v is the velocity of the constituent particle. A determination of the reduced mass of the constituent particles in the composite model using these decay rates gives

$$M = \frac{4m_\mu}{9} \left\{ \frac{1\alpha^4 m_\mu^3}{48\pi \Gamma(\mu \rightarrow e\gamma\gamma) [\Gamma(\mu \rightarrow e\gamma)]^2} \right\}^{1/8} \quad (\text{II-13})$$

In another version⁸⁰ of composite models, the muon is the $2S_{1/2}$ radial excitation state of the electron system, and the $\mu \rightarrow e\gamma\gamma$ decay mode is a double electric dipole transition. This occurs through an intermediate excited lepton that is in a P state. The branching ratio in this case is

$$B_{\mu \rightarrow e\gamma\gamma} = (d^2/M)^2 m_\mu^2 / 64 G_F^2, \quad (\text{II-14})$$

where d is a phenomenological parameter with the meaning of the transition electric dipole moment and M is the mass of the intermediate lepton state.

The reasons for searching for lepton-family nonconserving decays include the possible incompleteness of the standard model, an asymmetry in the behavior of quarks and leptons when changing flavors, and that lepton-family number does not appear to be a fundamental conservation law. Many extensions to the standard model accommodate lepton-family number nonconservation and a search for many rare decays is necessary to choose

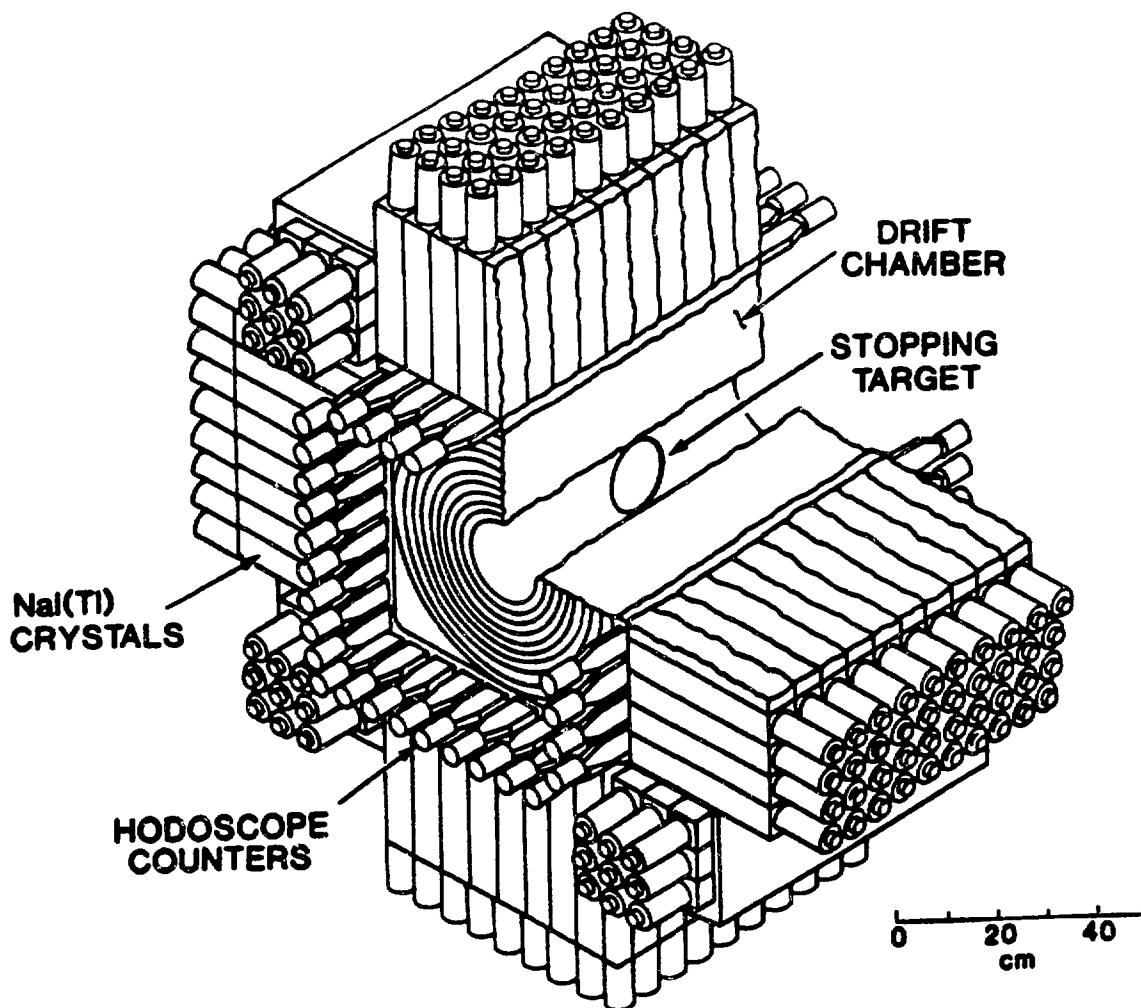
the correct extension. Of the rare decays, a search for $\mu \rightarrow e\gamma\gamma$ can limit some of these extensions.

CHAPTER III

DETECTOR

The detector used in this search for the $\mu \rightarrow e\gamma\gamma$ decay is called the Crystal Box, shown schematically in Figure 2. It consists of a target, drift chamber, scintillation counter hodoscope, and sodium iodide (NaI(Tl)) calorimeter. The target stops low momentum positive muons that subsequently decay. The trajectories of charged particles coming from these decays are measured in a cylindrical drift chamber surrounding the target. A group of 36 plastic scintillation counters cover the area outside the drift chamber. These counters serve as the trigger hodoscope for the detector. A large modular array of NaI(Tl) crystals surrounds the scintillator hodoscope, drift chamber, and target. It is arranged in four quadrants of 90 crystals each and four groups of nine crystals covering the corners. Not shown in Figure 2 are 16 scintillation guard counters, located on the upstream and downstream ends of the detector, and the "I-counter," a special counter used for timing measurements, located behind a small hole in the target. There is no applied magnetic field within the detector. Because the detector searches for three rare decays simultaneously, certain detector elements, sizes, and configurations are chosen to optimize the detector efficiency for one or more of the decay modes. The following sections describe the various components of the detector, the hardware trigger, and data acquisition system.

Figure 2. Schematic view of the Crystal Box detector.



A. The Muon Beam and Target

The muon beam is a very important part of an experiment searching for rare decays. The requirements of such a beam are an intense source of muons and good muon purity.

The first requirement is met by using a beam of surface muons⁸¹ at the LAMPP Stopped Muon Channel.⁸² Surface muons are produced by the decay of pions near the production target surface. Protons from the primary beam are incident on a graphite target and produce copious numbers of pions. Some of these pions stop near the surface of the target and decay into positive muons having a maximum momentum of ~ 30 MeV/c. These muons are transmitted down the beam channel to the experiment. For rare decay experiments, a surface muon beam is superior to the traditional beam from pion decays in flight. Surface beams have very low momentum, can be stopped in thin targets, and have higher fluxes. A surface muon beam does not work for negative muons because the negative pions stopped in the production target are captured in nuclei at a rate much larger than that for $\pi^- \rightarrow \mu^-$ decay.

The second beam requirement is that the contamination of unwanted particles, in this case, pions and positrons, be eliminated or greatly reduced from the muon beam. The Stopped Muon Channel is about 30 m long, so almost all pions within the low momentum acceptance of the channel will decay before they reach the detector. Positrons in the beam come primarily from pair production of π^0 decay photons in the production target. The positron contamination is reduced by using three dipole

magnets to select muons with a given momentum. Another device, the separator, is employed to substantially improve the purity of the beam. The separator has a transverse magnetic field and a strong, transverse electric field. It "separates" muons from positrons because of their difference in velocity. The final purity of the beam, given as the e^+/μ^+ ratio, is measured to be 1/50, an improvement by a factor of 500 over the unseparated beam purity.

The beam spot at the detector is fine-tuned using beam monitors to measure position and intensity. The drift chamber reconstructs tracks from positrons and measures the beam spot at the target. In this manner, the muon beam is studied and, if necessary, changed.

The linear accelerator at LAMPF accelerates protons to a kinetic energy of 800 MeV with an intensity capability of 1 mA. This beam is pulsed, with a pulse length of about 500 - 750 μ s, and runs at a frequency of 120 Hz. The fraction of time that the beam is actually on, defined as the duty factor, varies from about 5% to 10%, averaging 6.8%, during the data-taking periods. The rate of the muon beam stopped in the target, averaged over the data-taking periods, is measured to be 5.4×10^6 muons \cdot s $^{-1}$ (instantaneous rate) or 3.6×10^5 muons \cdot s $^{-1}$ (average rate). This rate is chosen to limit the amount of piled-up pulses in the NaI(Tl) crystals. During the latter portion of the running period, the beam was varied to allow three different muon rates, 4.2×10^6 , 5.2×10^6 , and 8.4×10^6 muons \cdot s $^{-1}$ (instantaneous), into the detector so that the rate-dependent performance could be studied.

The muon target is a thin elliptical piece of polystyrene mounted on a Lucite frame. The target intercepts the incident beam direction at an angle of 45° , thus giving a projected circular profile of 6.25 cm radius facing the beam. A canted target provides a greater surface area in which the muons stop and allows a differentiation of vertex locations from the decay positrons. The target, with a density of $1.057 \text{ g}\cdot\text{cm}^{-3}$, is 0.35 mm thick. The choice of the target material is polystyrene, because it depolarizes the stopped muons. Depolarization gets rid of any spin dependency related to the muon decay. The polarization of the stopped muons is $(19.2 \pm 7.7)\%$ along the beam direction, found by measuring the decay asymmetry within the detector. About 14% of the detectable beam misses the target and has positrons originating from the inner portion of the drift chamber.

A 1 cm radius hole is cut in the target, offset from the center by about 2 cm. The I-counter, which will be described later, is located behind this hole. The volume surrounding the target is filled with helium gas to minimize energy loss and background photons.

Located upstream of the detector and surrounding the beam pipe, is a lead collimator to reduce beam-related photons that may be detected in the NaI(Tl). The hole in the collimator has an elliptical shape so that the beam passing through it will image a circular profile on the target. An aluminum sieve, mounted within the lead collimator, transmits $\sim 50\%$ of the muons in the beam. The purpose of the sieve is to reduce the beam illumination at the center of the target.

B. The Drift Chamber

The drift chamber provides information on the trajectories of charged particles within the detector and locates the origin of the decay on the target. The requirements of such a device are: (1) optimal size so that it has a large solid angle, yet reduces the amount of NaI(Tl) crystals needed to surround it, (2) low mass construction to reduce the multiple scattering of particles and positron annihilations in flight, (3) high efficiency of track identification at high rates, and (4) good track direction and resolution in order to determine the vertex of the decay on the target. Much more detailed descriptions of the drift chamber are found in reference 83 for the design and construction and in reference 84 for the performance.

The drift chamber is cylindrical in shape and consists of eight concentric layers of wires surrounding the stopping target. The length of the chamber is 65.0 cm and the radii of the layers are between 10.5 cm and 22.0 cm measured at the chamber ends. The wires of each layer are alternately tilted at angles between 10° and 16° with respect to the beam axis, providing 20° to 32° of stereo angle between the layers. There are a total of 728 drift cells contained in the chamber with each cell measuring 8 mm \times 10 mm at the chamber ends. A drift cell consists of eight field-shaping wires and a sense wire located in the center. The field-shaping wires are 152 μm in thickness and are made of gold-plated copper-beryllium. The sense wires are 25 μm thick and are made of gold-plated tungsten.

The electric field lines within a drift cell are kept as radial as possible so that drift times depend only on the track distance from the sense wire. This is accomplished by setting the voltage on the field-shaping wires at different values. The four corner field-shaping wires of the rectangular cell are set at a nominal value, while the field-shaping wires between between the drift cells of the same layer are set at 90% of this nominal value and field-shaping wires between layers are at 72% of this value. These field-shaping wires between layers and drift cells isolate the different cells from one another. Cylindrical aluminum-coated Mylar sheets are located between the first layer of wires and the target region, and between the last layer of wires and the scintillation counters. These sheets are held at the nominal potential of the corner field-shaping wires so that the first and last layers of wires are electrically the same as the inner layers of wires.

A gas mixture of 49% argon, 49% ethane, and 2% isopropyl alcohol, by volume, is used in the chamber. This gas mixture provides good gain and drift time uniformity within the chamber. The nominal voltage using this gas mixture is -2500 V, giving a drift velocity of $51 \mu\text{m}\cdot\text{ns}^{-1}$ for drift electrons within a cell. The chamber gas is monitored using two small drift chambers, each consisting of three drift cells and a radioactive source. Any change in the chamber gas results in a change in the count rate from the source in the monitor chambers. These chambers can detect 1% changes in the gas mixture.

High voltage to the field-shaping wires is supplied by a programmable power supply.⁸⁵ This power supply is modified to provide protection against voltage surges caused by discharges within the drift

chamber. The voltage and gas mixture for the drift chamber are continuously monitored by a HP-9816 microcomputer.⁸⁶

One of the requirements for the design of the chamber is that it present as small an amount of mass as possible for the particles to traverse. The wires contribute 4.7×10^{-3} radiation lengths, the Mylar sheets 1.3×10^{-3} radiation lengths, and the chamber gas 0.7×10^{-3} radiation lengths to the radial thickness, making a total of 6.7×10^{-3} radiation lengths of material for the particles to cross.

Each sense wire is connected to an amplifier-discriminator channel, designed at LAMPF. There are two parts to the amplifier-discriminator. The first part is an analog part consisting of an amplifier and a pulse differentiator, and the second part is a digital part consisting of a fast comparator and multivibrator. The amplifier portion is based on a high-gain video integrated-circuit amplifier that is designed to reduce crosstalk and to improve the signal-to-noise ratio. The discriminator checks the leading edge of the pulse and has a non-updating one-shot multivibrator and ECL logic outputs. Crosstalk between channels is reduced by using ferrite beads to prevent impedance mismatch between sense wires and signal cable connections and by decreasing the bandwidth of the amplifier-discriminator electronics.

The signals proceed through active delay circuitry⁸⁷ to the input of the LeCroy 4290 time-digitizing system⁸⁸ via twisted-pair cable. This system consists of a time-to-digital converter (TDC) channel for each sense wire controlled by a fast readout pre-processor. The pre-processor rejects null data values given to it by the TDC channels and passes valid data to the main experimental computer. The TDC channels are run in the

common stop mode and a fast inhibit signal is used. Pedestal or null-data values are found by applying a test signal to the sense wires. Gain and pedestal variations in the TDC channels are automatically corrected within the 4290 system.

The track reconstruction computer program allows up to eight charged-particle tracks to be found in three spatial dimensions for each trigger event. The space coordinates are calculated from the times given by the TDC channels measured relative to the common stop time. The drift times are assumed to be proportional to the distance between the sense wire and the track. Deciphering the track locations, including the stereo ambiguity, is a very complex task even for one track, and this complexity increases with the number of tracks. The average track reconstruction time is found to be about 0.2 s per track on a VAX-11/780. A detailed description of the track-finding algorithm is given in reference 84.

The intrinsic spatial resolution per layer of the drift chamber is measured to be $\sigma_{rms} = 130 \mu m$. The overall resolution, which includes the intrinsic resolution plus the multiple scattering of drift electrons, varies from about 150 μm to 250 μm for a given layer of drift cells. This resolution value is determined by the rms position difference of a hit in a given layer and the track fitted to the remaining seven layers. The spatial resolution can be degraded by as much as 30% for higher beam rates due to the increased probability of confusing certain tracks with neighboring tracks. The efficiency of the layers of the drift cells depends on the beam rate, but is typically around 95%. The track reconstruction efficiency for one positron track in the drift chamber is about 97%.

C. Scintillator Hodoscope

The plastic scintillator hodoscope serves several purposes within the detector. These are to identify positrons and veto photons in conjunction with the NaI(Tl) crystals, to be a part of the trigger mechanism for the detector, to measure the arrival time of positrons, and finally, to provide an initial position for the drift chamber track-finding algorithm. Scaler readings of the rates in the middle scintillators of each quadrant are used to count the number of muons stopped within the detector. Because charged particles ionize scintillation material and neutral particles do not, the plastic scintillator can discriminate between positrons and photons. The performance requirements of the scintillators are a high efficiency for detecting positrons, a good time resolution for trigger information, accurate time measurements, and a reasonably good position measurement for use with the track-reconstructing program.

The hodoscope consists of 36 plastic scintillators located between the drift chamber and the NaI(Tl) crystals, and these are arranged in four groups of nine as shown in Figure 2. Each counter is 44.5 cm long, 5.7 cm wide, and 1.27 cm thick and is made of Pilot B scintillator material. The length of the scintillator along the beam direction is chosen to be shorter than the length of the NaI(Tl) crystal array so that any positron traversing the scintillator will have its energy shower contained within the NaI(Tl) array. The thickness of the scintillators is chosen to

produce enough light in order to obtain good timing resolution and the width approximately covers one "row" of NaI(Tl) crystals.

Each scintillator has photomultiplier tubes attached to both ends via a plastic light guide. The signals from each photomultiplier tube proceed into an amplifier, whose outputs are split three ways: 40% of the signal goes toward the trigger and a time measurement, 20% goes toward an energy measurement, and the remaining 40% goes toward the "photon veto," described in more detail with the trigger. The phototube signals for the trigger proceed into a constant fraction discriminator (CFD), designed⁸⁹ especially for the Crystal Box detector. These discriminators have a good time resolution over a large dynamic range of input signals. The amount of energy needed to pass the CFD threshold and to record a hit is 0.5 MeV. The CFD timing signals then proceed into a meantimer, which produces a signal at the mean arrival time of the two pulses put into it and compensates for the length of the counter. Thus the trigger of the detector is started at a common time for all scintillators, regardless of where a particle hits along the length of the scintillator. The meantimers used with the plastic scintillation counters are of a new design⁹⁰ especially built for the Crystal Box detector. The average time resolution for the plastic scintillators is 290 ps (FWHM).

D. NaI(Tl) Calorimeter

The NaI(Tl) crystal array is used to measure the energy deposited in it by positrons and photons, to find the position and arrival time of photons, and to define the trigger of the detector. Good resolution

values in the energy, time, and position measurements are necessary for reducing unwanted background events. The NaI(Tl) energy resolution is $\sim 8\%$ (FWHM) at 50 MeV, the time resolution is ~ 1.2 ns (FWHM), and the position resolution is ~ 4 cm (FWHM).

1. Design

The "Crystal Box," named for the large number of NaI(Tl) crystals constructed in the shape of a box, is shown in Figure 2. There are a total of 360 crystals forming four quadrants of 90 crystals, known as "face" crystals, and 36 crystals forming four sets of nine crystals, known as "corner" crystals. The modular design of many regularly-shaped NaI(Tl) crystals enclosed within a single container is unique to the Crystal Box. A single-container design minimizes the amount of inert material between the crystals. Previous designs have included separate containers for a large set of NaI(Tl) crystals or very special shapes that are very costly. The NaI(Tl) crystals are necessarily placed within a sealed container because of their hygroscopic nature. Having a modular array of NaI(Tl) crystals, such as the Crystal Box, has advantages over a single crystal of NaI(Tl). These include a better determination of the impact point of a particle with the NaI(Tl) crystals, a measurement of time and energy for more than one particle, and a higher rate of particle detection with fewer piled-up pulses. Disadvantages of this design are a loss in the energy resolution, usually on the order of several percent because of the boundary effects between adjacent crystals, and the cost of many more channels of electronics.

Each NaI(Tl) face crystal measures 2.5 in. \times 2.5 in. \times 12.0 in. and these individual crystals are stacked in 9 rows that are transverse to the beam axis, and in 10 columns that are along the beam axis. The cross-sectional dimensions of the crystals are selected to correspond approximately to the size of an electromagnetic shower within a crystal so that shower containment and a position measurement are made without using an extreme number of crystals. The crystal depth of 12 radiation lengths is chosen to contain as much of the shower as possible and to provide good timing characteristics for a fixed cost. Each corner crystal measures 2.5 in. \times 2.5 in. \times 27 in. and are grouped in 3 \times 3 arrays located at the intersections of the four quadrants. These crystals are used to contain the energy from particle showers that may leak into the corners. The 27 in. length that is along the beam axis, covers the distance spanned by 10 columns of face crystals. The crystals are optically polished on all sides to further improve the time measurement. Each crystal is wrapped in a thin, diffuse, millipore paper scatterer and then in a 0.003 in. outer covering of aluminized Mylar. Thus each crystal is optically isolated from its neighbors.

Face crystals are optically coupled to one Amperex XP-2232B photomultiplier tube,⁹¹ while the corner crystals have phototubes on both ends. The particular choice of the XP-2232B photomultiplier tube is made because of its fast timing response and its good anode current linearity. The optical coupling assembly from the NaI(Tl) crystal to the phototube consists of a 0.125 in. elastomer pad against the NaI(Tl) crystal, a fused-quartz glass window through the sealed container, a 0.25 in. elastomer pad, and finally a 2 in. lucite light guide glued to

the phototube. An optical fiber is cemented into a notch cut into the side of the light guide and is used to insert a simulated NaI(Tl) light pulse from a xenon flash lamp.

The phototube base is designed to supply sufficient current to the dynodes of the phototube so that gain variations at high rates are minimized. The construction of the base is such that it provides adequate cooling and contains only passive components in the voltage divider chain. Two separate high voltages are supplied to each phototube. One is at a voltage of -1500 V and the other is at one-half of this value. The latter supplies more current to the last six phototube dynodes. This operation improves the reliability of the phototube and reduces the power dissipated by the base. The high voltage ground and signal ground connections are separated to reduce the adverse effects of "ground loops" on signal pulses.

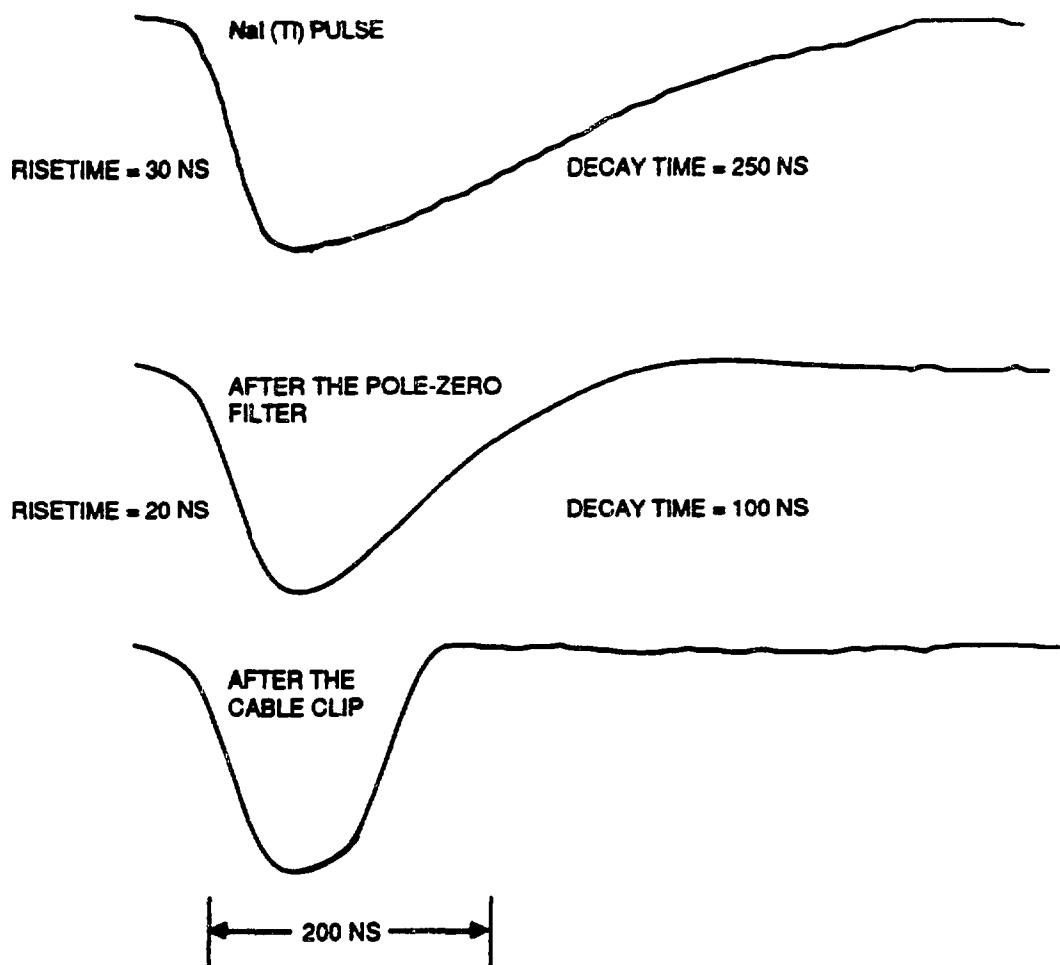
2. Signal Pulse Shaping

Scintillation light is produced in the NaI(Tl) crystals by incident radiation. The energy is given to the host NaI lattice and is efficiently transferred to the impurity thallium ion. The emitted scintillation light is characteristic of the Tl^+ dopant. A more detailed discussion on the NaI(Tl) scintillation process is presented in reference 92. The light reaching the photocathode of the photomultiplier tube consists of photons distributed in time by a combination of the NaI(Tl) scintillation process and the "rattling" of the photons within the crystal before they exit through the window. The resulting signal from the phototube has an

envelope that has a 10% to 90% risetime of 25 ns and a long tail, approximately fit by an exponential function with a 250 ns decay time constant. A sample NaI(Tl) signal pulse is shown in Figure 3. The emitted light of a NaI(Tl) crystal is proportional to the amount of energy deposited in it. The integral of the pulse signal from the phototube is proportional to the emitted light of the crystal. Therefore, the charge in the phototube output pulse is proportional to the energy. A measurement of the charge contained in the pulse should be taken over as long a time period as possible in order to integrate most of the charge. However, a long integration time allows the possibility of other pulses from unrelated events to be included in a particular charge measurement. These pulses are called piled-up pulses. A solution to this problem is to reshape the pulses so that a shorter integration time can be used.

There are two techniques used to reshape the NaI(Tl) signal pulse. A pole-zero cancellation changes the long tail with a 250 ns decay time into a tail with a 100 ns decay time. It also sharpens the leading edge of the pulse and thus improves the time resolution. The second technique "clips" the signal pulse. A portion of the signal pulse is sent down a length of coaxial cable where the pulse is reflected, inverted, and added to the original pulse. The length of the coaxial cable is chosen so that this reflected and inverted pulse adds to the original pulse and cancels the long tail. The resulting pulse after reshaping is about 100 ns wide without the long tail. There is still statistical information in the clipped-out portion of the pulse because of the stochastic nature of the NaI(Tl) signal. This information is not used in the charge integration and consequently degrades the energy resolution. The signal cable is

Figure 3. Diagram of the NaI(Tl) signal pulses, showing the original pulse and the pulses after signal processing with a pole-zero filter and a cable clip.

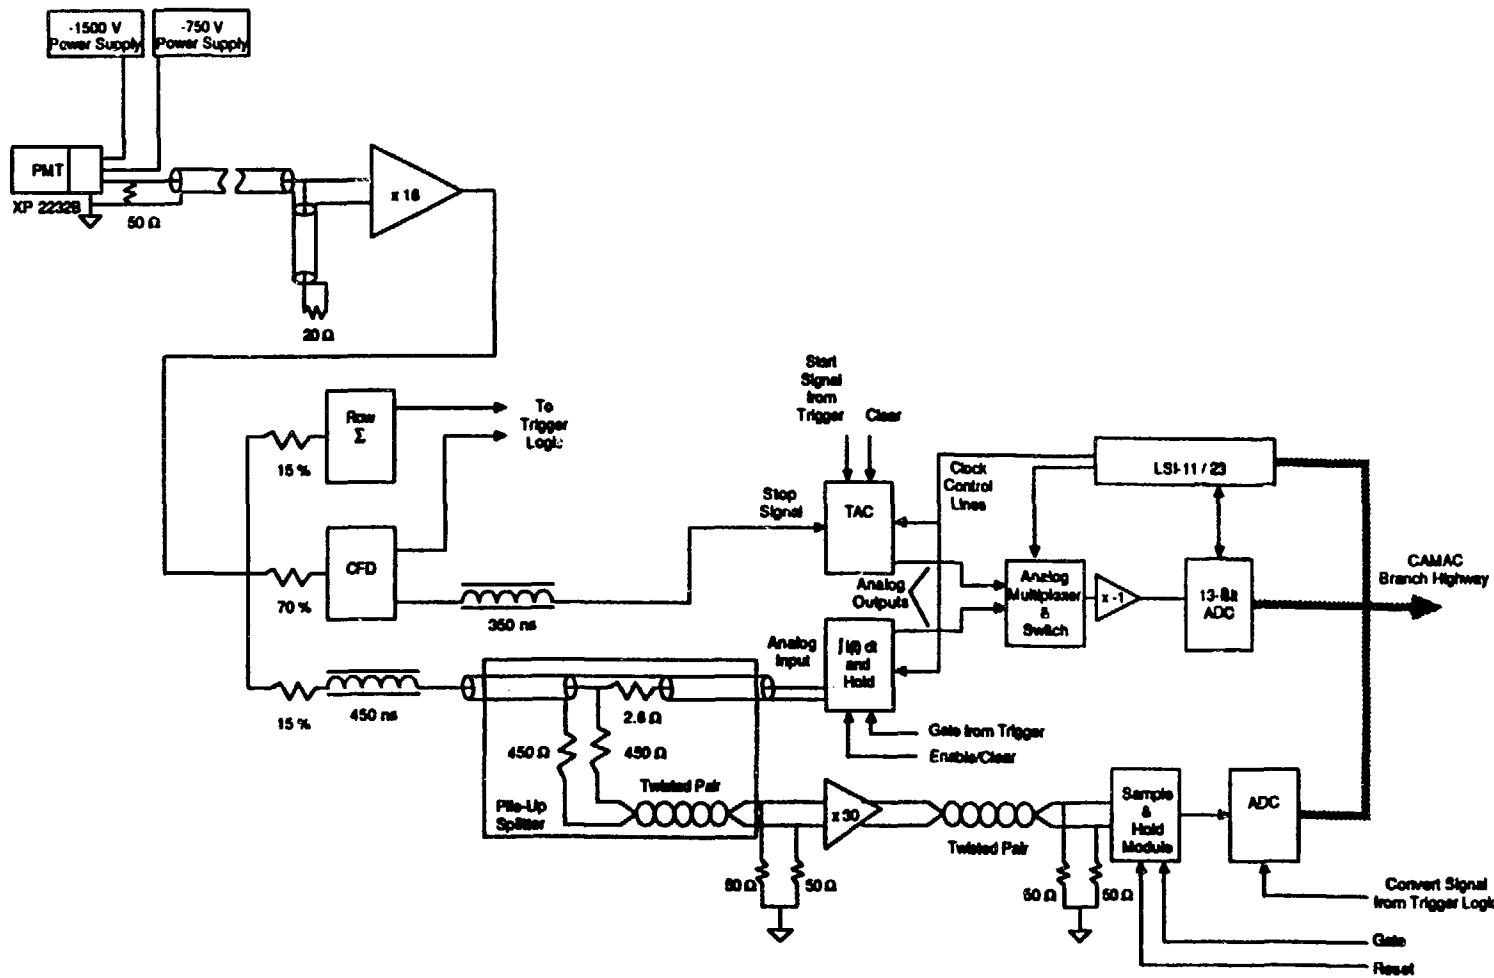


back-terminated at the phototube base to prevent multiple signal reflections along the transmission line. The results of the NaI(Tl) signal pulse shaping are shown in Figure 3.

3. NaI(Tl) Electronics

A schematic diagram of the NaI(Tl) electronics is shown in Figure 4. The photomultiplier tube produces a signal pulse that proceeds via coaxial cable to specially designed amplifiers. These amplifiers not only increase the gain of the signal, but reshape the pulse and passively split the signal into three parts: (1) about 70% of the split signal goes toward the time measurement, (2) about 15% goes toward a measurement of the energy deposited in a given row of NaI(Tl) crystals used in the trigger, and (3) the remaining 15% goes toward the energy measurement. The first part of the split NaI(Tl) signal pulse proceeds into a CFD and then into a time-to-amplitude converter (TAC) module, where the arrival time of the pulse is measured. The third signal part is delayed 450 ns by a distributed delay line, after which the signal is again split into two parts. About 90% of this signal goes into the primary analog-to-digital converter (ADC) system, while the other 10% goes into a second ADC system that measures the charge in piled-up pulses. The primary ADC system consists of an integrate-and-hold (I/H) channel for each NaI(Tl) crystal, four analog multiplexers, and four Tracor Northern TN1213A spectroscopy ADCs.⁹³ The I/H module integrates the current in a pulse and holds the charge until digitized by the ADC. The operation of this system, as well as the timing system, is under the control of four LSI-11/23

Figure 4. Schematic diagram of the NaI(Tl) array electronics.



microcomputers. There are four of these systems corresponding to each of the four quadrants of NaI(Tl) crystals. The second ADC system is very similar to the primary system, containing a differential preamplifier and current-integrating sample-and-hold (SHAM) channel for each NaI(Tl)crystal, with one microprocessor-controlled ADC (BADC) for the entire system. The output of both ADC systems is passed to the main experimental computer.

The NaI(Tl) amplifier consists of two parts. One is a high-frequency amplifier section and the other is a low-frequency stabilization section. The amplifier section is based on a pair of fast transistors and the gain of the amplifier is about 16. The stabilization section uses a comparison of the signals from the original pulse with the output from the high-frequency amplifier. Any difference found between the two signals must be due to drift from the high-frequency amplifier. The output from the stabilization amplifier is then used to correct the drift.

The output ADC signal from the amplifier for a given channel is delayed 450 ns by a Spiradel⁹⁴ distributed delay line. The signal delay is necessary so that the signals can be stored, while the trigger electronics are making their logic decisions. The Spiradel is an alternative method⁸⁷ of delaying signals without using the standard coaxial cable delays. The main advantage of using the Spiradel delays is a substantial reduction in the amount of physical space necessary to hold an equivalent amount of delay using bulky coaxial cables. A Spiradel delay is a disk, 3.75 in. in diameter and 0.438 in. in thickness. In this experiment, the Spiradel delays replaced about 64 km of coaxial cable.

4. Primary ADC and Time Systems

An essential part of the primary ADC system is the integrate-and-hold (I/H) modules⁹⁵ that measure the energy deposited in the NaI(Tl) crystals. The signal from an amplifier channel proceeds into a channel of the I/H module. The current is then integrated and the charge stored on a capacitor. A gate signal, provided by the trigger, is coincident with the pulse from the amplifier and starts the charging of the capacitor. The integration time for the spectroscopy ADC system is 200 ns. The output of the I/H module for zero signal has a small value called the pedestal. The charge on the capacitor is held until it is digitized by the ADC, then the capacitor is discharged. This ADC system is capable of a 13-bit resolution, or 1 part in 8000, in the energy measurement. The performance of this module gives a stability and accuracy of better than 0.1% of full scale and a linearity of 0.2%.

Pulses for the NaI(Tl) time measurement proceed from the amplifiers to constant fraction discriminators (CFD). The CFD output pulses are delayed by 350 ns, again by Spiradel distributed delays, before going into time-to-amplitude converters (TAC).⁹⁵ A Start signal common to all TAC channels and provided by the experimental trigger begins the time measurement. A capacitor within the TAC begins to charge linearly with time. A Stop signal, generated by a pulse from a CFD channel, ends the time measurement and the charging of the capacitor. A null-time value for a channel is produced by automatically discharging the capacitor if no Stop signal is received by the TAC. This minimizes the digitizing time

for the TAC channels. The timing values for a NaI(Tl) crystal are digitized with a resolution of 100 ps per time bin with a total of 1024 bins. The TAC modules are designed to have good stability and charge-integrating linearity.

The same output multiplexing method is used for collecting the data contained in all the I/H and TAC channels. The clock signal output of each module is connected to the clock input of the next module, and so forth, such that the series of modules for a given quadrant are "daisy-chained" together. The analog outputs of these channels are connected in parallel to a common analog output line leading to the spectroscopy ADC. The clock pulses then gate each channel sequentially into a single ADC per quadrant so that it reads the analog signal information from both the I/H and TAC channels for a given event trigger. The analog signal multiplexing is then done with a minimum of connections.

The four LSI-11/23 microcomputers, which consist of a central processing unit and several specialized data interfaces, serve four functions in the NaI(Tl) data acquisition system. They control the multiplexing of the analog signals that are to be digitized by the spectroscopy ADC, acquire the digitized data from the ADC, perform a sparse data scan, and store constants, such as the pedestal values, for each NaI(Tl) channel. The first and second functions have been discussed previously. A sparse data scan involves testing the data values for all channels and compacting the data set by removing those data that contain only pedestal values. After the sparse data scan, the LSI tags each data word with a crystal identification code and signals the main experimental computer that it is ready to transfer the data. The LSI also stores

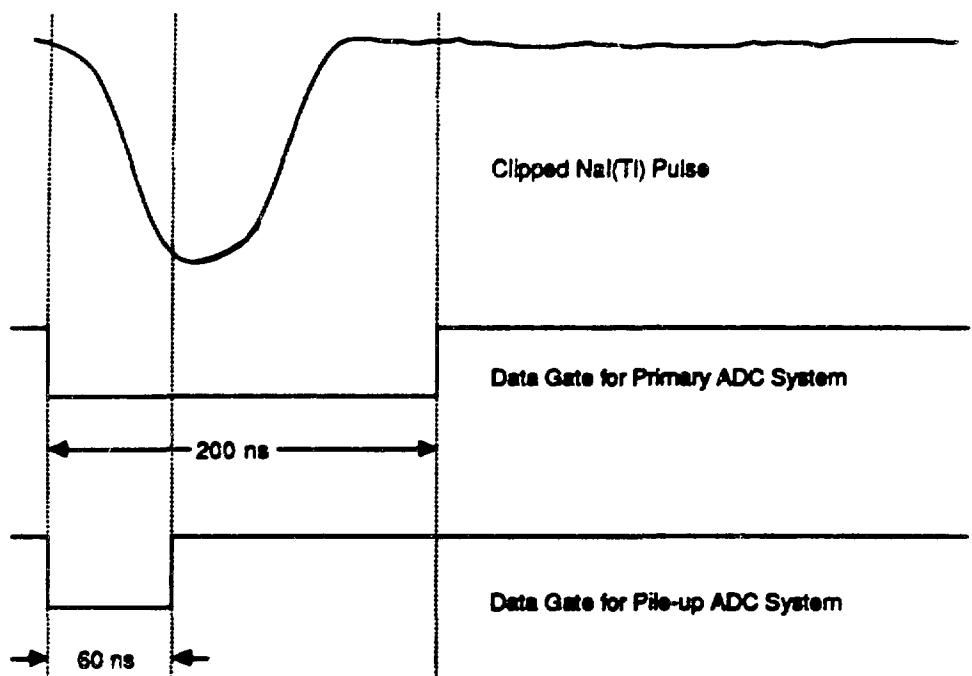
pedestal values for each NaI(Tl) channel and other parameters used to monitor the stabilization of the NaI(Tl) energy and time measurements.

5. Pile-up ADC System

The second ADC system used with the Crystal Box detects the amount of piled-up pulses present in the NaI(Tl) crystals. This system consists of a differential preamplifier, a sample-and-hold (SHAM) circuit for each channel, and an ADC and microprocessor (BADC) that controls the readout of the SHAMs. After amplification, the signal proceeds to the SHAMs, where the pulse is integrated and the charge stored on a holding capacitor. Each channel is read by the BADC, which multiplexes the data acquisition, digitizes the analog information, and compacts the data through a sparse data scan. Pedestal values are automatically updated within the BADC. The BADC and SHAM modules are manufactured and sold commercially.⁹⁶

The integration time for the pile-up ADC system is 60 ns, compared to the 200 ns integration time used with the primary ADC system. Figure 5 shows the integration periods for both ADC systems and a typical NaI(Tl) pulse. The beginning of both periods is coincident, so that the pile-up ADC system integrates the same ratio of charge within the pulse for the first one-third of the spectroscopy ADC integration time. This arrangement of the integration times provides a sensitive detection of piled-up pulses. Details of the two pile-up rejection methods are given in Chapter V.

Figure 5. Diagram of the processed NaI(Tl) signal pulse and the two current-integrating time periods for the primary and pile-up ADC systems.



E. I-Counter

As mentioned in an above section, there is a 1 cm hole in the muon stopping target. Located directly behind this hole is a special counter, called the I-counter. This counter consists of a 1 in. diameter cylindrical aluminum cup, embedded in a 1.5 in. diameter cylindrical scintillator. The scintillator is coupled to a photomultiplier tube by a Lucite light guide. This entire assembly is aligned along the beam axis with the bottom of the aluminum cup located 5.79 in. downstream of the center of the muon target. The location of the I-counter does not interfere with any of the particle trajectories originating from the target and intersecting the NaI(Tl) array.

Muons from the beam pass through the hole in the target and stop in the aluminum cup. The decay positrons, originating from the cup, traverse the scintillator material and cause a pulse in the I-counter.

The purpose of the I-counter is to start a time measurement for the NaI(Tl) crystals and scintillation counters used in the experiment. In this way, the timing of all the counters used in the trigger of the experiment is referenced to a single counter. The time resolution of the I-counter is 350 ps (FWHM). The timing calibration and stabilization for the entire detector depends on the I-counter.

F. Guard Counters

There are 16 plastic scintillation counters located at the upstream and downstream ends of the detector. These guard counters cover the last 1.5 columns of NaI(Tl) crystals whose area is not included in the active region of the scintillator hodoscope. This provides an additional charged particle veto for the trigger. Without these guard counters, positrons from normal muon decay would be misidentified as photons at the outer ends of the detector. There is only one photomultiplier tube attached to these counters. The size of the guard counters are 13.3 cm \times 23.8 cm \times 0.3 cm and they have an average time resolution of 1.7 ns.

G. Trigger

Because of the high rate of muons stopping within the detector and the multitude of accidental, coincident triggers, a sophisticated hardware trigger is required to reduce the amount of data that gets recorded. The overall detector trigger⁹⁷ is complex because the search for the three decays, $\mu \rightarrow eee$, $\mu \rightarrow e\gamma\gamma$, and $\mu \rightarrow e\gamma\gamma$, occurs simultaneously. Only the trigger portion applicable to the $\mu \rightarrow e\gamma\gamma$ decay will be described here.

The plastic scintillators, guard counters, and NaI(Tl) crystals are used in the $\mu \rightarrow e\gamma\gamma$ trigger. Particle identification by the trigger is defined in the trigger logic using the division of the detector into quadrants. A "positron" quadrant is defined as a scintillator meantime signal, requiring signals from both ends of the counter, and a CFD signal

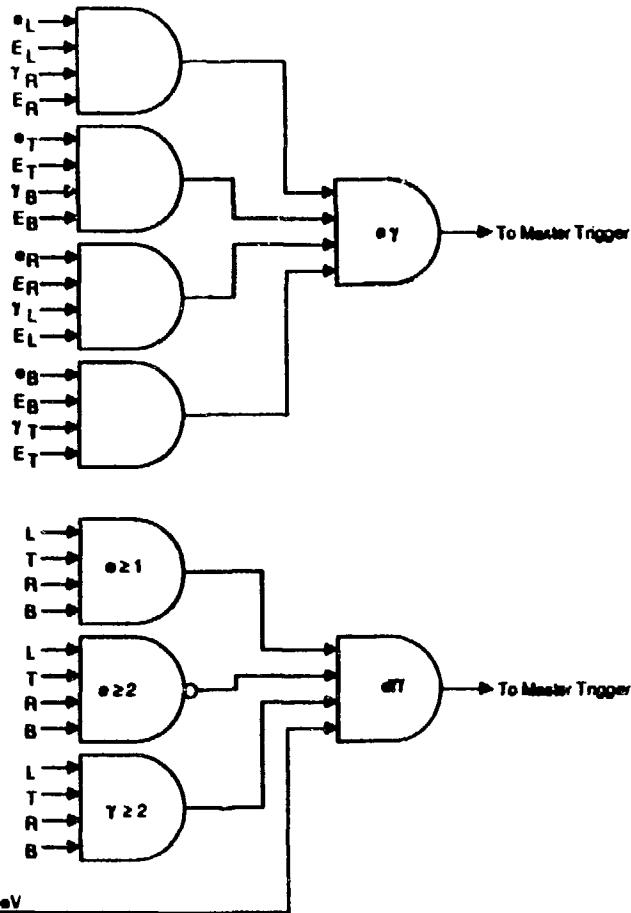
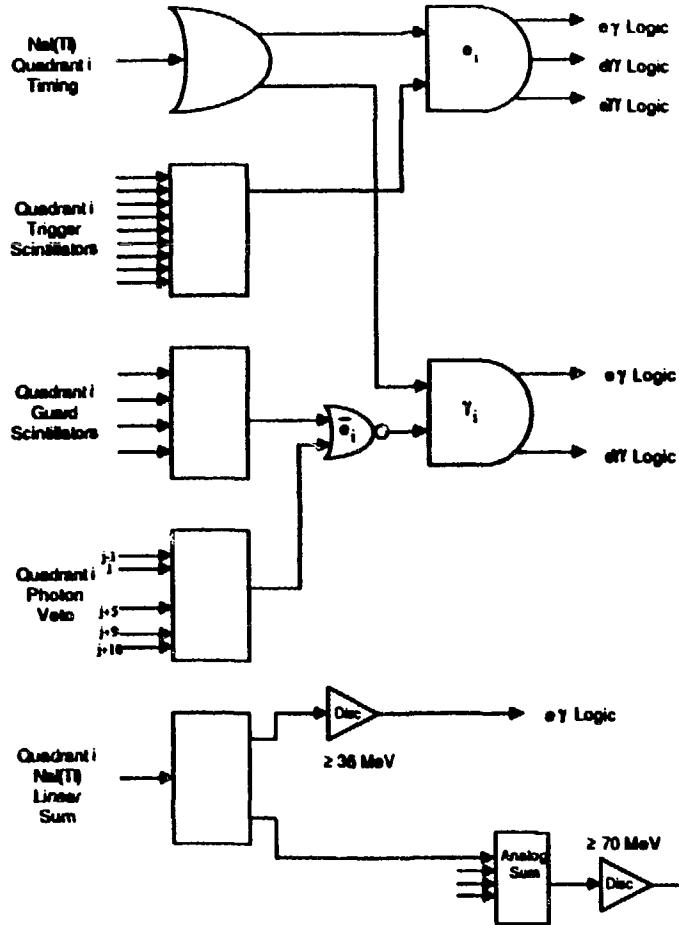
from the NaI(Tl) crystals in that quadrant. The energy threshold in a NaI(Tl) crystal CFD is about 5 MeV with variations as much as 2 MeV. A "photon" quadrant is defined as a NaI(Tl) crystal signal with no signal from the scintillator hodoscope or the "photon veto." The photon veto consists of the plastic scintillation hodoscope counters and the guard counters within a given quadrant, plus one trigger scintillation counter nearest the given quadrant from the adjacent quadrants. This last requirement prevents energy "splash" from a particle in an adjacent quadrant near a corner to be identified as a photon. The photon veto signals originate from leading-edge discriminators because of their better efficiency. Given these quadrant definitions in the trigger, a positron quadrant may contain photons, but no positron can be included in a photon quadrant. Signals from the NaI(Tl) crystals in columns 1 and 10, and in the corner crystals that lie along the central diagonal are not included in the trigger logic. Columns 1 and 10 are located on the far ends of the detector so the probability that a portion of a particle shower is missing can be large. The central diagonal crystals in the corner lie on the border of the definition of a quadrant, so there is an ambiguity as to which quadrant a particle should belong. Therefore these NaI(Tl) crystals, which may lead to problems, are excluded from the trigger logic.

A crude on-line energy measurement is made by simply adding pulses using linear fan-in modules. Signals from each row of ten NaI(Tl) crystals are added together into a row sum, and then the nine row sums are added together to form a face sum. Each set of nine corner crystals are divided into three groups, the three crystals along the diagonal, the three crystals on one side of the diagonal, and the three crystals on the

other side of the diagonal. Signals from the corner crystals in the group closest to a given face are added together along with one-half of the pulses from the diagonal group to form corner sums. The quadrant sums are formed by adding together the face sums and the corner sums. Finally, the four quadrant sums are added together to give a total energy sum. Any on-line restriction of the energy using this method is coarse because the energy limits are set on the basis of pulse height rather than by pulse integration.

An $\text{e}\gamma\gamma$ trigger is defined as: (1) one and only one positron quadrant, (2) at least two photon quadrants other than the positron quadrant, and (3) at least ~ 70 MeV deposited in the total energy sum. The last requirement is met by placing a discriminator on the total energy sum to set a lower bound on the energy, but no upper bound is used because it would tend to eliminate events containing more than three particles. Figure 6 schematically shows the $\mu \rightarrow \text{e}\gamma\gamma$ trigger. If either a positron and a photon or both photons are within the same quadrant, then the second requirement will not be met. In this particular case, there is a possibility of still satisfying the $\text{e}\gamma$ trigger. An $\text{e}\gamma$ trigger is defined as: (1) a positron quadrant, (2) a photon quadrant that is opposite the positron quadrant, (3) at least ~ 30 MeV measured in the positron quadrant sum, and (4) at least ~ 30 MeV measured in the photon quadrant sum, and (5) a modest restriction on the positron and photon geometry using the columns of NaI(Tl) crystals. An $\text{e}\gamma\gamma$ candidate event with two particles in the same quadrant must satisfy these $\text{e}\gamma$ trigger requirements, also shown in Figure 6. An event that satisfies the $\text{e}\gamma\gamma$ trigger is called a

Figure 6. Logic diagram for $\mu \rightarrow e\gamma$ and $\mu \rightarrow e\gamma\gamma$ triggers.



three-quadrant event and one that satisfies the $e\gamma$ trigger a two-quadrant event.

The above trigger requirements identify particles and make crude restrictions on the geometry and energy. Another trigger requirement that is implied within the above set is on the time coincidence of the three particles. The $e\gamma\gamma$ trigger uses both the scintillator hodoscope and the NaI(Tl) crystals, with the trigger time distribution limited by the NaI(Tl) pulses. The time distribution is dispersed by the different times of all the crystals; some of these crystals are a little early and some are a little late in providing a trigger time, so a loose requirement is deliberately set at about 10 - 15 ns. The major source of error for the on-line trigger time comes from using entire quadrants of NaI(Tl) crystals as trigger elements.

H. Data Acquisition

The LAMPF beam structure provides a natural method for processing data. Because of its 120 Hz structure and 6 - 10% duty factor, there is a 7.3 ms break between beam pulses. It is a desirable goal of the experiment to process the data events completely between beam pulses. A study⁹⁷ of the raw trigger rates showed that the highest muon stop rate that can be run with the detector and still maintain the desired data processing goal is 5×10^5 muons \cdot sec $^{-1}$ (average). This condition is imposed mainly to make the $\mu \rightarrow eee$ and $\mu \rightarrow e\gamma$ trigger rates manageable. This limit on the rate is confirmed with tests using a prototype detector.

The number of signals that must be collected at for a given event is about 2000. Energy, time, and pile-up measurements from the 432 channels of NaI(Tl) crystals, time and energy from the 36 plastic scintillation and 16 guard counters, time information from the 728 channels of the drift chamber, and various latch bits and status words are all read during every triggered event. In addition, 120 scaler values are read and recorded every 10 s. A single event can produce 3000 16-bit words of data and this would lead to writing several thousand computer tapes of data.

Data compaction by using a sparse data scan of null data values in the NaI(Tl) crystal and drift chamber channels reduces the data rate by a factor of four. As discussed previously, the LSIs scan the NaI(Tl) data values and the LeCroy 4290 system compacts the drift chamber data. After this compaction, the average data rate is about 750 data words per event indicating that most of the channels are indeed empty.

Even after the sparse data scan, the time for reading and collecting the data takes longer than the 7.3 ms between beam pulses. This adversely affects the dead time of the detector. To minimize the dead time and the consequent loss of beam time, "double buffering" is used to process the data from the NaI(Tl) crystals; other data is collected fast enough not to affect the dead time. Double buffering means that a second data buffer acquires and holds data from a second event while the data from the first event is still being processed. This method gives a more reasonable dead time and accounts for trigger logic that can take data from two, but not necessarily three, consecutive beam pulses.

The Crystal Box data acquisition system is based on the LAMPF standard data acquisition system.⁹⁸ The LAMPF standard system uses a PDP-11 computer, a Microprogrammed Branch Driver (MBD),⁹⁹ and a CAMAC (Computer Automated Measurement and Control) based readout system. The data acquisition and analysis program used at LAMPF is the "Q" program,¹⁰⁰ which provides most of the computer code necessary to run an experiment. It has the capability to record data, analyze and display the data during the data taking, and monitor the detector performance. By using the LAMPF standard data acquisition and analysis program package, the amount of effort, especially in computer programming, is minimized and the expertise and support for such a system is available.

Once the data acquisition system is enabled, a typical event is processed in the following manner. A trigger, such as the `err` trigger, is issued from the hardware logic by a Master Trigger, which in turn signals the Event trigger. This then starts the data collection from the scintillators, drift chamber, and the NaI(Tl) array, and also disables the detector trigger so that no further events can be taken during this time. The MBD runs certain programs that collect data from the scintillator ADC and TDC channels, check the drift chamber 4290 system for finishing the sparse data scan and then acquire its data; and monitor the progress of the NaI(Tl) data acquisition. It is not necessary for the LSI-11/23 microcomputers to have finished processing all the data before the trigger is enabled for the next event, but they do have to finish the data collection from all the I/H and TAC modules. Should a second event occur, these microcomputers are interrupted so that the data are collected from the second event before returning to process the first event. Once the

LSIs are finished with an event, the MBD reads the data into its memory via the CAMAC branch highway. The MBD buffer, which contains all of the data for an event, has its contents transferred into the PDP-11/44 for further processing. In the case of a second event occurring during the processing of the first event, the acquisition system is not enabled until the first event is read by the MBD. This is how data can be taken during the two, but not necessarily three consecutive beam pulses.

The data in the MBD buffer are transferred to the PDP-11/44 through a Direct Memory Access (DMA) transfer that conveys blocks of data from one device to another. Once the data are inside the PDP-11/44, a part of the Q program distributes an event to the Analyzer, a program that checks the data for errors, applies cuts to the data, and if the data pass these cuts, puts it on computer magnetic tape. Besides the processing done by the Analyzer, other computer programs acquire the data from the Analyzer to perform on-line processing of the event and monitoring of the detector performance. This is a useful application because the quality of the data can then be checked. A complete description and account of the performance of the data acquisition system is found in reference 101.

I. Liquid Hydrogen Target

An important part of any calorimeter system is to obtain a correct calibration of the energy deposited in it. To do this, a thin-walled, liquid hydrogen target replaces the drift chamber in the Crystal Box and the beam channel is tuned to transport negative pions. Two photons are produced in the decay of the π^0 from the single charge exchange reaction,

$\pi^- p \rightarrow \pi^0 n$, with energy values in the range of 55.1 MeV to 82.7 MeV. A single photon is produced in the reaction, $\pi^- p \rightarrow n \gamma$, with a single energy of 129.4 MeV. These photons, of known energy, then calibrate each NaI(Tl) crystal.

The liquid hydrogen target consists of an insulated container for the liquid hydrogen, the associated equipment needed to maintain and cool the hydrogen, and the control electronics for the target. The liquid hydrogen container is made of kapton and is cylindrically-shaped with hemispherical ends. The length of the container is 5 cm and the diameter 3.8 cm. Two plastic scintillation counters are placed upstream of the liquid hydrogen container to define the beam of pions incident on the target.

The container, refrigeration apparatus, and some electronics are mounted together on a stand that replaces the drift chamber and polystyrene target. It takes several hours for the hydrogen gas to cool to a liquid state, change the channel to transport pions, and swap targets inside of the detector.

J. Flasher System

Along with the energy calibration, another important aspect of a calorimeter is stabilization. Once the detector is calibrated, it is very important to monitor and correct any changes in the energy measurement. A method used in the Crystal Box is the flasher system. It consists of a xenon flashtube, two monitors of the light output from the flashtube, and 432 fiber optic cables. Light from the flashtube is conveyed via the

fiber optic cable into the light guide of a NaI(Tl) crystal phototube. The energy of this light is measured exactly as the normal data, and is monitored as a function of time. Changes in the light output from the flashtube are detected by the two monitors, a vacuum photodiode and a silicon photodiode. Any variation detected in the light output or energy calibration is corrected in the off-line analysis.

The other use of the flasher system is to monitor the performance of both the I/H and TAC electronics channels of the NaI(Tl) and flag any channel that appears to malfunction. A complete description of the flasher system and its performance is given in the Appendix.

CHAPTER IV

MONTE CARLO PROGRAM, CALIBRATIONS, AND STABILIZATION

The Crystal Box detector measures the energies, times, and positions of positrons and photons. The detector must be calibrated for any of these measurements. Once calibrated, variations must be monitored as a function of time and corrected in the analysis. The accuracy of these measurements is estimated from Monte Carlo computer simulations. This Monte Carlo program accurately describes the Crystal Box, and its response to simulated events. This chapter presents the Monte Carlo program and calibration data for the experiment.

A. Monte Carlo Program

The program used in the Crystal Box experiments is a sophisticated computer code that simulates events within the detector. It describes the detector completely. All parts of the detector, including the NaI(Tl) array, scintillation counters, drift chamber, guard counters, I-counter, and target, are represented with actual dimensions, material, and arrangement. The response of the detector constituents is accurately reproduced by the computer. The electromagnetic showers produced from particles interacting with material are simulated with the EGS3 shower code.¹⁰² Not included in the EGS3 code, but added to the Monte Carlo program, is the Landau energy loss distribution of the particles through

material. The Monte Carlo also uses the measured beam distribution to reproduce the locations of muons stopped on the target.

The particle trajectories are followed from a selected origin on the target through the various detector elements. Positrons traversing the drift chamber in the Monte Carlo leave tracks through ionization, just as positrons would in a real chamber. The EGS3 shower code keeps an account of the energy loss of a particle through the different materials. Once all the energy is "measured" in the detector, the trigger logic of the experiment is imposed. Those events passing the trigger logic are then recorded in a format similar to the actual data format. The data analysis then treats these simulated events and data events equally.

Different trigger types can be implemented in the Monte Carlo. The $\mu \rightarrow eee$, $\mu \rightarrow e\gamma$, and $\mu \rightarrow e\gamma\gamma$ decays are simulated within the detector to provide distributions of the analysis variables. The prompt background events, $\mu \rightarrow e\gamma\nu\bar{\nu}$, and $\mu \rightarrow eee\nu\bar{\nu}$, can also be generated according to the matrix elements describing these decays. The calibration reactions, $\pi^- p \rightarrow n \pi^0$, $\pi^- p \rightarrow n \gamma$, and $\mu \rightarrow e\nu\bar{\nu}$, are produced in order to determine the resolution functions of the detector.

The computer code generates a $\mu \rightarrow e\gamma\gamma$ event at rest by randomly selecting particle momenta within a given three-dimensional coordinate system, and then by rotating this coordinate system over a 4π solid angle. The magnitude of the positron momentum, $|\vec{p}_e|$, is selected according to the distribution, $y^2(1-y)^2$ where $y = m_\mu/2$ and $0 \leq y \leq 1$ (Equation II-4), and its direction is pointed along axis 1 of the coordinate system. The magnitude of the momentum for one of the photons, $|\vec{p}_{\gamma_1}|$, is randomly picked from the range of available momentum, $0 < |\vec{p}_{\gamma_1}| \leq m_\mu - |\vec{p}_e|$, where m_μ is the mass

of the muon, and the direction of \vec{p}_{γ_1} lies in the plane formed by axis 1 and axis 2. The magnitude of the momentum for the second photon, $|\vec{p}_{\gamma_2}|$, ranges from $0 < |\vec{p}_{\gamma_2}| \leq m_\mu - |\vec{p}_e| - |\vec{p}_{\gamma_1}|$, and is constrained so that the energy of the three particles equals the muon mass. The direction of \vec{p}_{γ_2} is such that momentum is conserved in this decay. The coordinate system is then randomly rotated over a 4π solid angle.

The acceptance of the detector, defined as the fraction of solid angle from the target that the detector subtends, is found using the Monte Carlo program. The efficiency of detecting $\mu \rightarrow e\gamma\gamma$ events, including the data restrictions, is also calculated from the Monte Carlo. The product of the acceptance times the efficiency for detecting $\mu \rightarrow e\gamma\gamma$ events is the number of Monte Carlo events surviving the analysis divided by the total number of Monte Carlo events originating in the detector. The final calculation of the branching ratio uses this product.

B. NaI(Tl) Energy Calibration

The energy calibration¹⁰³ of the NaI(Tl) crystals or equivalently, the energy gains, requires several steps. The NaI(Tl) energy gains are defined as the amount of energy measured for a given ADC value, in units of MeV per ADC bin. The first step is to set the output pulse heights from every NaI(Tl) amplifier to equal values for a given energy. The second step is to measure the response of the ADC system for pulses applied over a wide range of pulse heights. The next step is to perform the actual calibration using photons from the π^0 decay. The final step is to check these calibrations with other energy values.

The output pulse amplitudes for all NaI(Tl) amplifier channels are set to equal values by using a known input signal. Variations of the output pulses among phototubes are then compensated in the amplifiers rather than by changing the high voltages to individual phototubes. A Pu- α -Be source, producing photons with an energy of 4.4 MeV from the reaction, $^{9}\text{Be}(\alpha, n)^{12}\text{C} \rightarrow ^{12}\text{C} + \gamma$, is the pulse standard to set all the amplifier outputs leading to the CFDs such that 1 V output corresponds to 50 MeV. This calibration sets an effective threshold of about 5 MeV in the CFD channels. The corner crystals, different because they have phototubes at both ends, are set so 500 mV corresponds to 50 MeV at the amplifier outputs. This gives a CFD threshold of 10 MeV for the corner crystals. This step can be done with an accuracy of about 10%.

The energy calibration curve turned out to be nonlinear, caused mostly by the nonlinear output of the NaI(Tl) amplifiers. To correct this, the energy response of each NaI(Tl) ADC channel was measured over a wide range of pulse height values using a pulser and precision attenuator. The shape of the nonlinear calibration was fit for each NaI(Tl) channel with the functional form,

$$E_i = A_i (x_i + B_i [1 - \exp(-\alpha x_i)]) , \quad (\text{IV-1})$$

where E_i is the energy equivalent output of the test pulse for a given channel i ; $x_i = \text{PH}_i - \text{Ped}_i$, the measured pulse height minus the measured pedestal value; and A_i , B_i , and α are parameters of the fit. The A_i are the energy gains and the B_i are constants for a given NaI(Tl) crystal. The value of α remains fairly constant for all crystals and is measured as

$\alpha = 0.001$ 11. The values of α and B_i determine the shape and the values of A_i determine the scaling of the calibration curve. The amount of nonlinearity remaining after this correction is less than 0.2 MeV for energies below 100 MeV.

Photons from π^0 decays, produced in the single charge exchange reaction $\pi^- p \rightarrow n \pi^0$, set a calibration point for the energy gains. In the π^0 rest frame, the two photons are produced isotropically, each with an energy of 68.9 MeV. This photon energy is chosen to be the calibration point because of its proximity to the interesting region of 52.8 MeV in muon decays. The energy of the photons, measured in the laboratory frame, ranges from 55.1 MeV to 82.7 MeV, and is called the "box spectrum." The kinematics of the π^0 decay in the lab frame are uniquely determined by the opening angle between the two photons, θ_{12} . The photon energies, E_1^{CM} and E_2^{CM} , in the π^0 rest frame are

$$E_1^{\text{CM}} = \frac{E_1^{\text{Lab}}}{1 + \epsilon(\theta_{12})}$$

$$E_2^{\text{CM}} = \frac{E_2^{\text{Lab}}}{1 - \epsilon(\theta_{12})} \quad (IV-2)$$

determined by measuring θ_{12} , the photon lab energies, E_1^{Lab} and E_2^{Lab} , and

$$\epsilon(\theta_{12}) = \left[1 - \frac{2}{(1 - \cos\theta_{12}) \gamma^2} \right]^{1/2} \quad (IV-3)$$

The value of γ^2 , 1.043, follows from the energetics of the charge exchange reaction at rest.

Each photon deposits its energy into the NaI(Tl) crystal array, but it usually illuminates more than one crystal. The high pulse height crystal is one containing the most energy deposited for a given particle shower. The energy of a particle shower is measured from the sum of the energies deposited in the high pulse height crystal and its 24 nearest neighbors. In a centrally-located face crystal, for example, the sum of the energies from a 5×5 group of NaI(Tl) crystals defines the particle energy. If the high pulse height crystal is located near an edge or a corner of the quadrant, the group of 25 crystals varies from 12 to 20 crystals in the energy sum, depending on the location. Once a crystal is assigned to a particular shower, it is not included with any other particle shower or energy sum.

The calibration procedure requires that at least 50% of the energy in the particle shower be included in the high pulse height crystal. The entire energy sum from the particle shower is then placed into a histogram for that high pulse height crystal. The data peak of the energy spectrum is fit for resolution, gain, and amplitude, and compared with a Monte Carlo-generated energy spectrum for that crystal. In the Monte Carlo program, the crystals are divided into groups according to their position within the array. Using the symmetry of the detector, 31 groups are formed, 25 face crystal groups and 6 corner crystal groups. Crystals within a group have similar photon impact positions and have a similar energy leakage and shower spread. This fitting procedure is reiterated with the new energy constants because the energy distributions in the

surrounding crystals affect the energy in the high pulse height crystal. The energy gains for the crystals usually converge after about 5 iterations.

Figure 7a shows a typical energy spectrum with the energy peak at 68.9 MeV for a centrally-located crystal in the left quadrant. Figure 7b displays the Monte Carlo simulation of the photon energy spectrum for the same crystal with an added Gaussian resolution of 14% (FWHM) measured with respect to the 0.667 MeV photon energy peak of ^{137}Cs . This resolution corresponds to $\sim 5\%$ (FWHM) at 55 MeV for each crystal, when scaled by the $E^{-1/4}$ resolution curve for NaI(Tl) crystals.^{104,105}

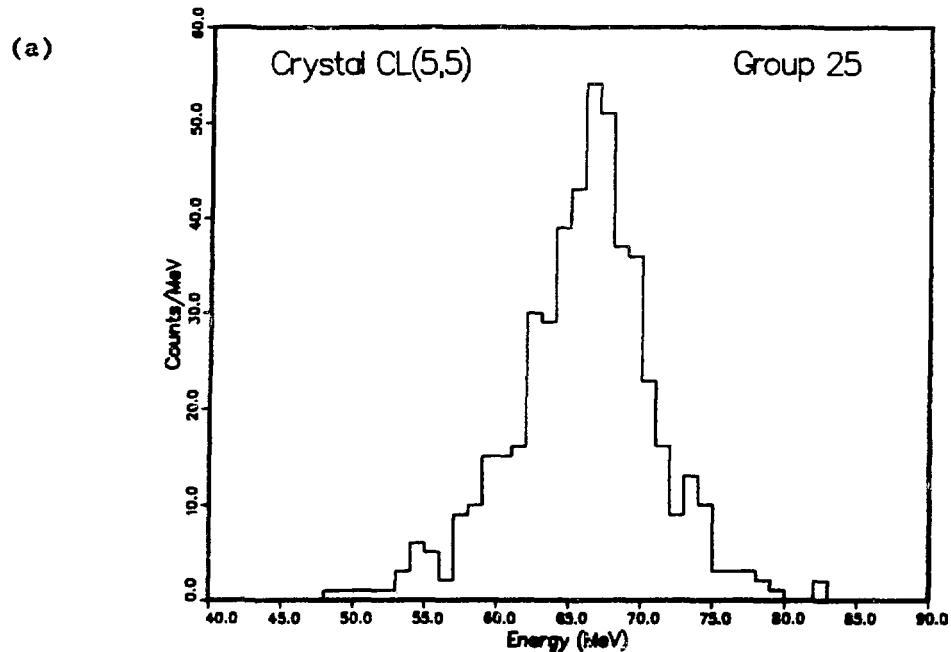
Several checks of the energy calibration are performed using other known energy values. The measurements of the energy used in these checks are taken only from the high pulse height crystals located in the central 30 face crystals for a quadrant. The first check uses the reaction, $\pi^- p \rightarrow n \gamma$, where the single photon energy is 129.4 MeV. The measured energy spectrum is given in Figure 8a and its corresponding Monte Carlo simulation in Figure 8b. The energy resolution is found to be $\sim 7\%$ (FWHM) and the peak positions of the data and Monte Carlo agree to better than 1%.

Another check of the calibration is to examine the two π^0 decay photons at large opening angles. These photons will have energies at the extremes of the box spectrum. The spectra are shown in Figure 9a for the measured energy values and in Figure 9b for the Monte Carlo spectrum. The lower energy peak has an energy resolution of $\sim 8\%$ (FWHM) at 55 MeV, which includes a contribution in the width from the range of measured opening angles.

Figure 7. Photon energy spectra from $\pi^0 \rightarrow \gamma\gamma$ events.

(a) Measured photon energy spectrum for crystal (5,5) in the left quadrant. (b) Predicted Monte Carlo photon energy spectrum for the same crystal.

Detected Central Energy for $\pi^0 \rightarrow \gamma\gamma$



Predicted Central Energy for $\pi^0 \rightarrow \gamma\gamma$

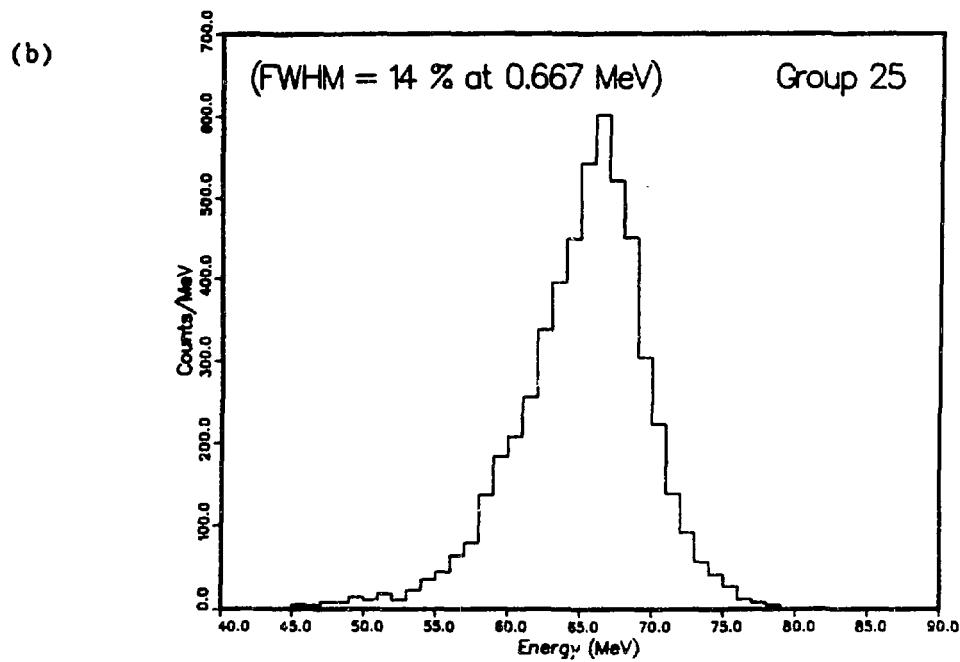
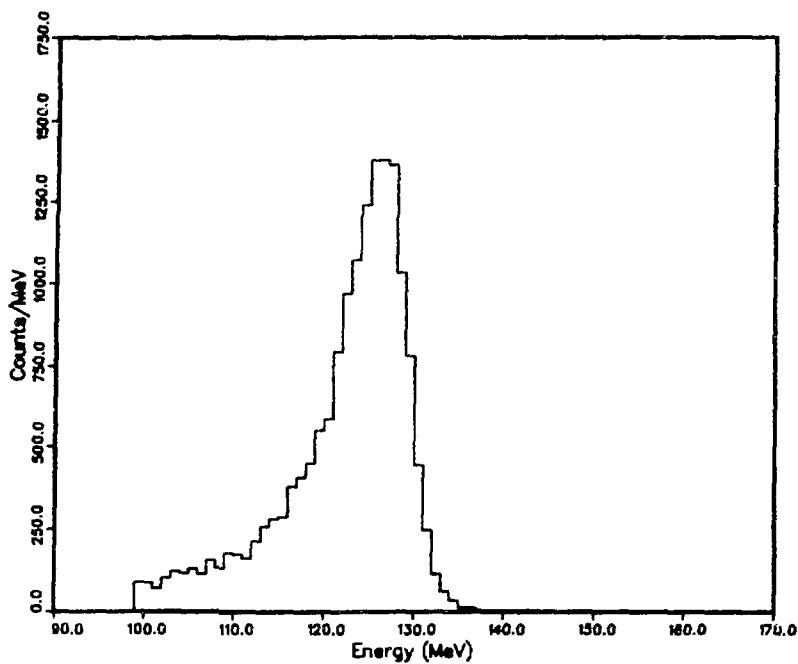


Figure 8. Photon energy spectra from the $\pi^- p \rightarrow n \gamma$ reaction.

(a) Measured and (b) predicted photon energy spectra.

Detected Photon Energy for $\pi^- p \rightarrow n \gamma$

(a)



Predicted Photon Energy for $\pi^- p \rightarrow n \gamma$

(b)

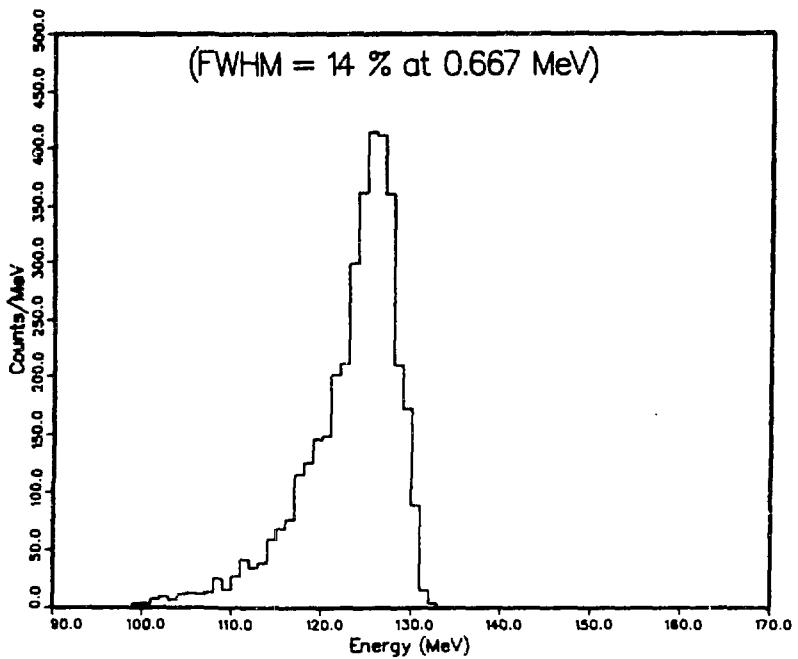
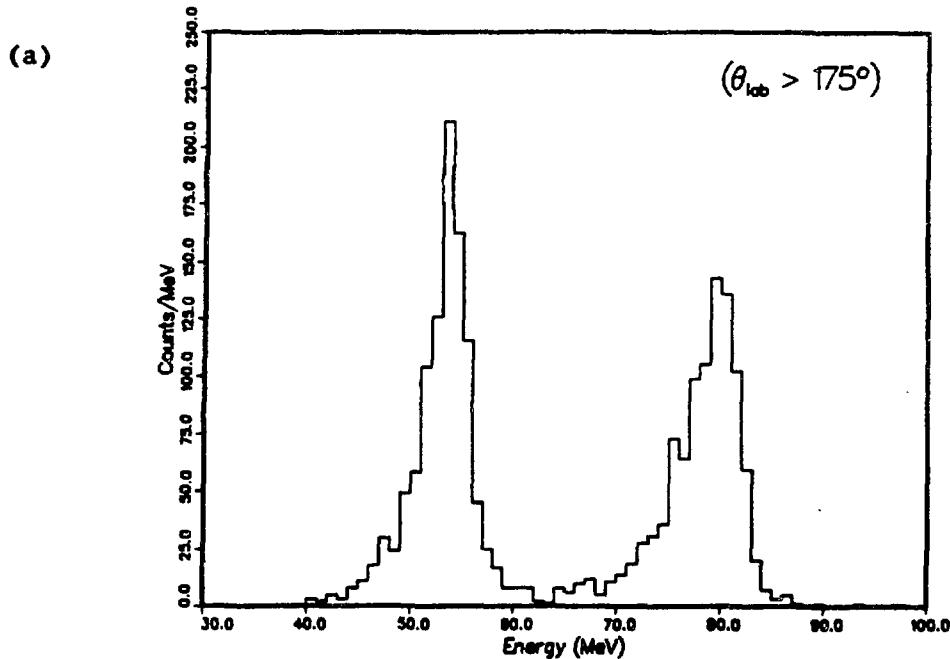
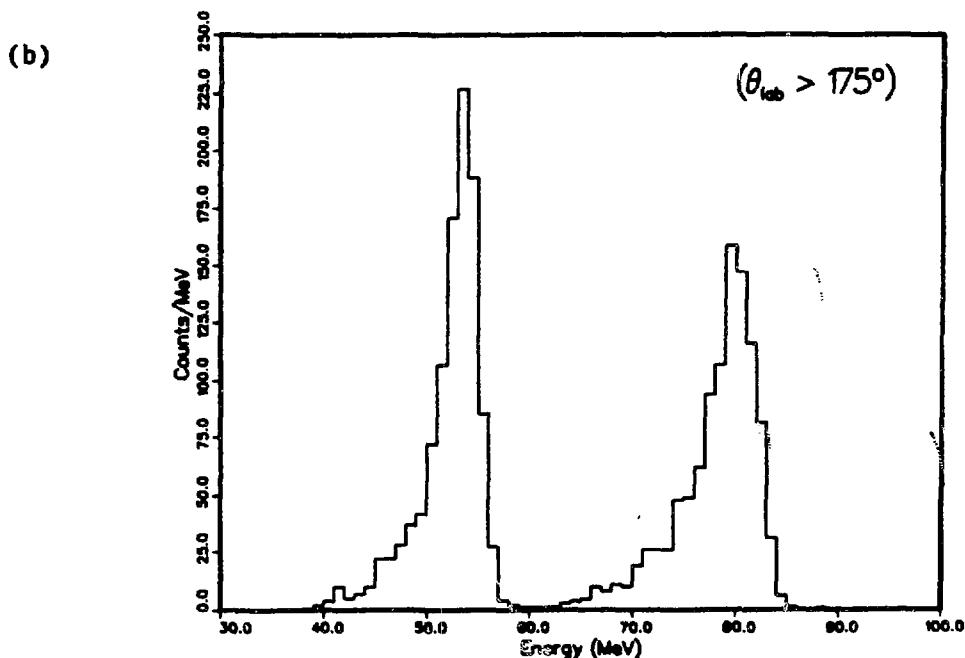


Figure 9. Photon energy spectra from $\pi^0 \rightarrow \gamma\gamma$ events with the opening angle between the photons greater than 175° . (a) Measured and (b) predicted photon energy spectra.

Detected Photon Energies for $\pi^0 \rightarrow \gamma\gamma$



Predicted Photon Energies for $\pi^0 \rightarrow \gamma\gamma$



A low energy calibration check is performed with the Pu- α -Be source. The measured energy spectrum is shown in Figure 10a along with the corresponding Monte Carlo prediction in Figure 10b. The Monte Carlo does not simulate the entire energy spectrum of photons from the source, but only gives the single energy peak at 4.4 MeV. Not included in the Monte Carlo simulation, for example, is the interaction of thermal neutrons captured in ^{127}I from the iodine in the crystals, producing a photon with an energy of \sim 6.5 MeV. The peak positions from the data and the Monte Carlo agree to within 0.2 MeV.

The final check of the calibrations uses the positron energy spectrum $f_1: \text{u}^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$. This has a cutoff point at 52.8 MeV. All previous energy measurements use photons, so the NaI(Tl) crystal response to positrons is determined with this energy spectrum. The positrons are produced in the target with a trigger requirement of one scintillator signal along with some NaI(Tl) energy signal. The results are shown in Figure 11a for the data, with no piled-up pulses included, and in Figure 11b for the predicted Monte Carlo spectrum. The data and the Monte Carlo of the energy calibration agree to within 1%. All of these checks verify that the energy calibration is consistent over a wide energy range.

C. Scintillator Energy Calibration

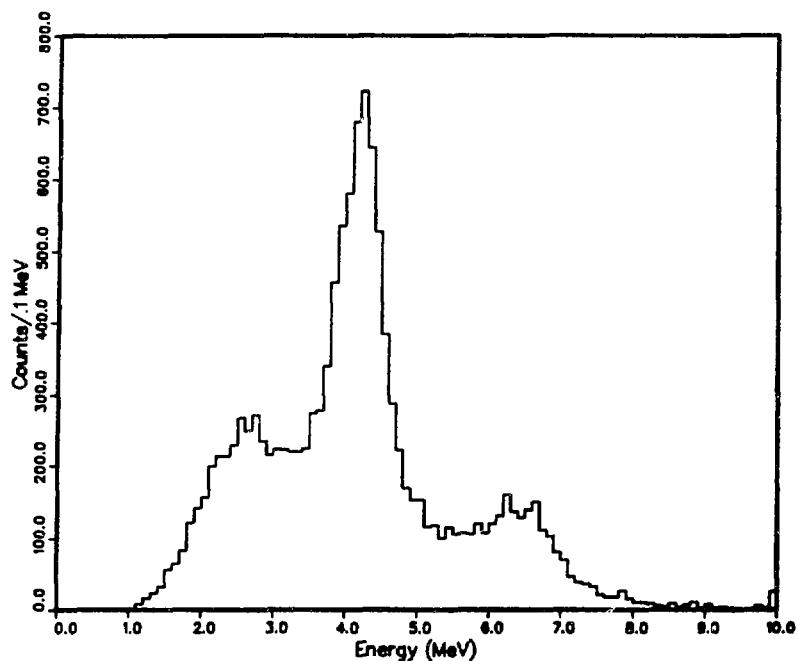
The plastic scintillator hodoscope differentiates between positrons and photons. This differentiation requires a certain amount of energy to be deposited in the scintillator before a particle is called a positron. An ADC channel connected to each phototube measures this amount of energy.

Figure 10. Photon energy spectra from a Pu- α -Be source.

(a) Measured and (b) predicted photon energy spectra.

Pu- α -Be Source Energy Spectrum

(a)



Predicted Detector Response to 4.4 MeV Photons

(b)

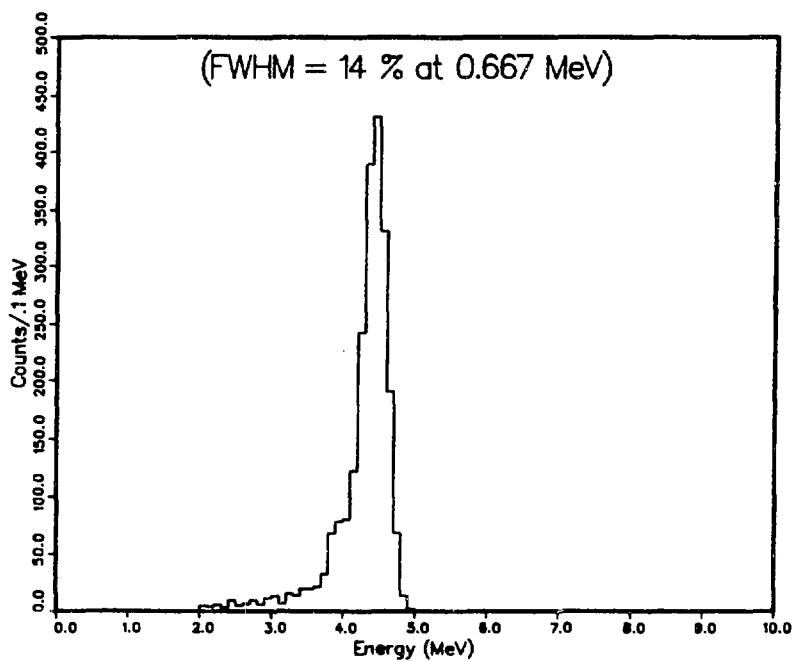
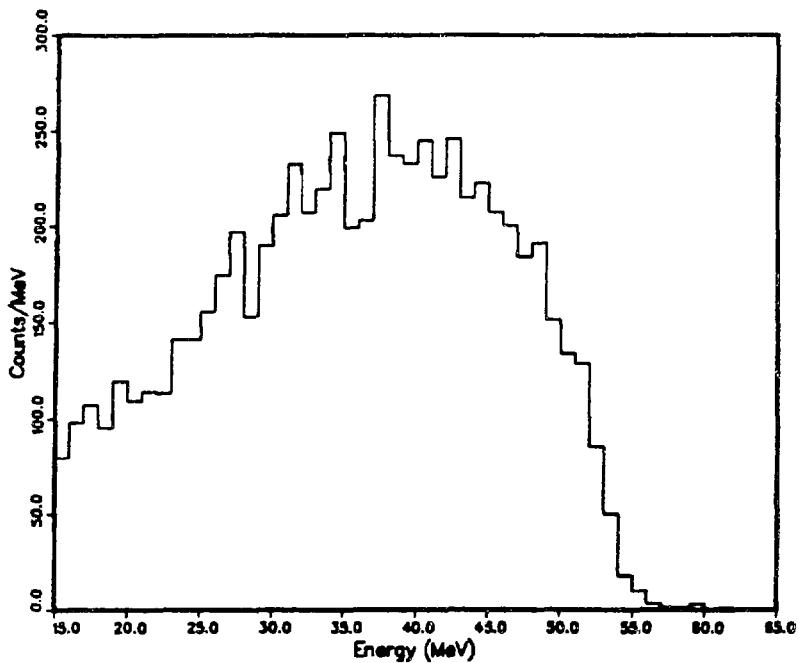


Figure 11. Positron energy spectra from muon decay.

(a) Measured and (b) predicted Michel positron energy spectra.

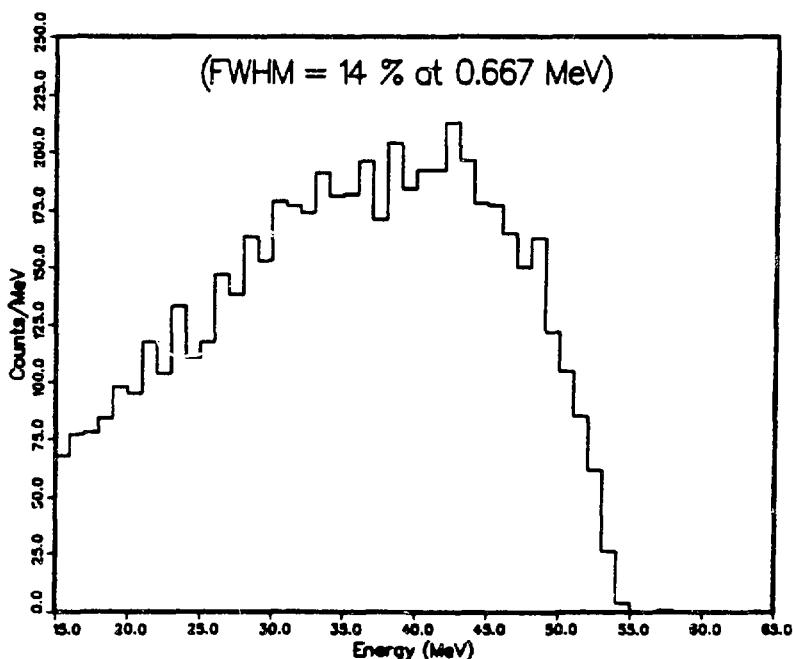
Detected Michel Positron Spectrum

(a)



Predicted Michel Positron Spectrum

(b)



The phototube gains are initially set equal for all scintillators so that the CFD registers a hit for a minimum-ionizing particle above a threshold of 0.5 MeV. A minimum-ionizing particle normally traversing the scintillator deposits about 2.5 MeV. This energy threshold minimizes triggers caused by a shower in the NaI(Tl) crystals that spread into the scintillators. Photons may also spread energy into the scintillators and so can be misidentified as positrons. The photon veto requires no more than 0.25 MeV be deposited in the scintillators.

The scintillator energy gains are determined from an impact position and the energies measured by each phototube. These gains are updated frequently during the data-taking periods, but no significant changes in the values occurred.

The energy gains for the guard counters are set in the same manner as the trigger hodoscope scintillators. The final gains are determined by comparing Monte Carlo and data energy spectra for positrons coming from the target. This method is accurate enough to provide a veto for the trigger logic. No energy sum is used with the guard counter energy values.

D. Scintillator Timing Calibration

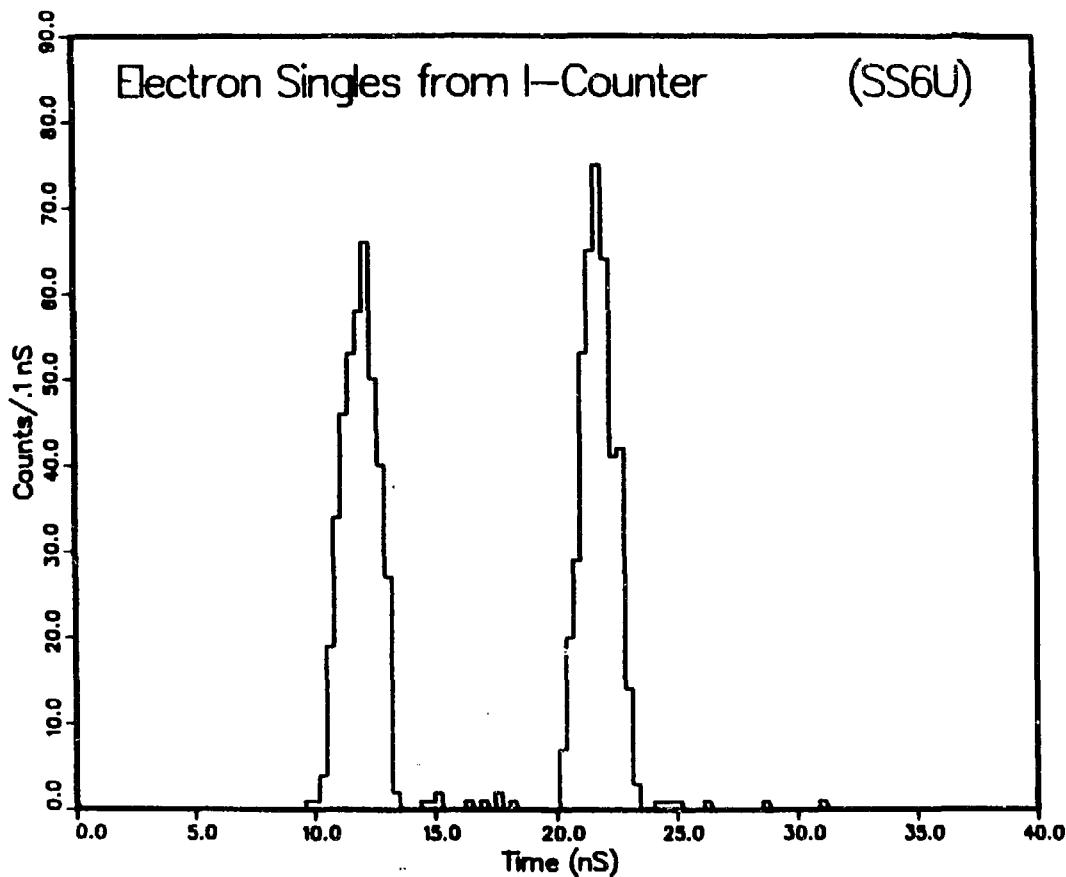
The scintillator timing calibration is performed in several steps. After the initial synchronizing of the meantimer pulses into the trigger logic, two sets of special data are taken and analyzed to determine the timing gains and offsets. The timing or TDC gain is defined as the amount of time measured for a given TDC value, in units of 100 ps per TDC bin.

The timing of all the plastic scintillators into the trigger logic is set relative to the I-counter. The I-counter gives the Start signal and the scintillator meantimer pulse gives the Stop signal from decay positrons that originate from the I-counter. A single TDC channel is used only in this initial step to set the times of all the meantimer pulses into the trigger logic. A set of latches associated with the scintillators indicates which of the scintillators have hits. Data from the TDC channel and latches are both collected and analyzed, so that the timing peaks from all scintillators are set to the same TDC value, within 0.5 ns.

The first of three steps used in the timing calibration is to analyze special timing data and to obtain the timing gains and offsets. The trigger requirements for the special timing data sample are a signal from the I-counter, one and only one scintillator signal, and at least 10 MeV deposited in the NaI(Tl) quadrant containing the scintillator. A 10 ns delay cable is inserted into an alternate signal path to the TDC channels, so that for every other triggered event, the signal pulse used for the time measurement gets switched through the alternate path. The result is a double-peaked time spectrum for each scintillator phototube taken with respect to the I-counter time and separated by a known, fixed time interval. A typical time spectrum is shown in Figure 12 for the upstream timing channel of scintillator 6. Each peak is fit with a Gaussian curve and the difference in the positions of the centroids is used to calculate the TDC gains. The centroid position of the first peak is the first estimate of the time offset for that channel. These timing constants are updated before further analysis is done. A similar analysis

Figure 12. Double-peaked time spectrum from the upstream channel of scintillator number 6.

Trigger Counter Timing Spectrum



procedure is used with the guard counters to find their timing gains and offsets, but there is no further timing analysis after this first step.

The second step in the scintillator timing calibration is to analyze a different set of special timing data in order to find the time differences between upstream and downstream phototubes for a given event. Positrons from muons that have stopped on the target are used to illuminate the scintillators. The trigger requirement is a meantimer pulse from a scintillator and a CFD pulse from a NaI(Tl) crystal that implies ~ 5 MeV deposited in that crystal. A typical time difference spectrum is shown in Figure 13 for scintillator 6. It gives the relative time offsets between the upstream and downstream time channels for a given scintillator. The centers of the distributions are aligned for all scintillators and the timing constants are again updated for the next step.

The third step is to use the special timing data from the first step with the new timing constants. A mean time between the upstream and downstream TDC values is calculated for each scintillator. The resulting mean-time spectrum is also double-peaked and a typical mean-time spectrum for scintillator 6 is shown in Figure 14. The position of the centroid of the first peak, which gives the absolute time offset, is found and adjusted so that all mean times are the same for all scintillators. The peak width is a measure of the time resolution, and the average time resolution for all 36 scintillators is 290 ps (FWHM).

Figure 13. Upstream and downstream channel time difference spectrum from scintillator number 6.

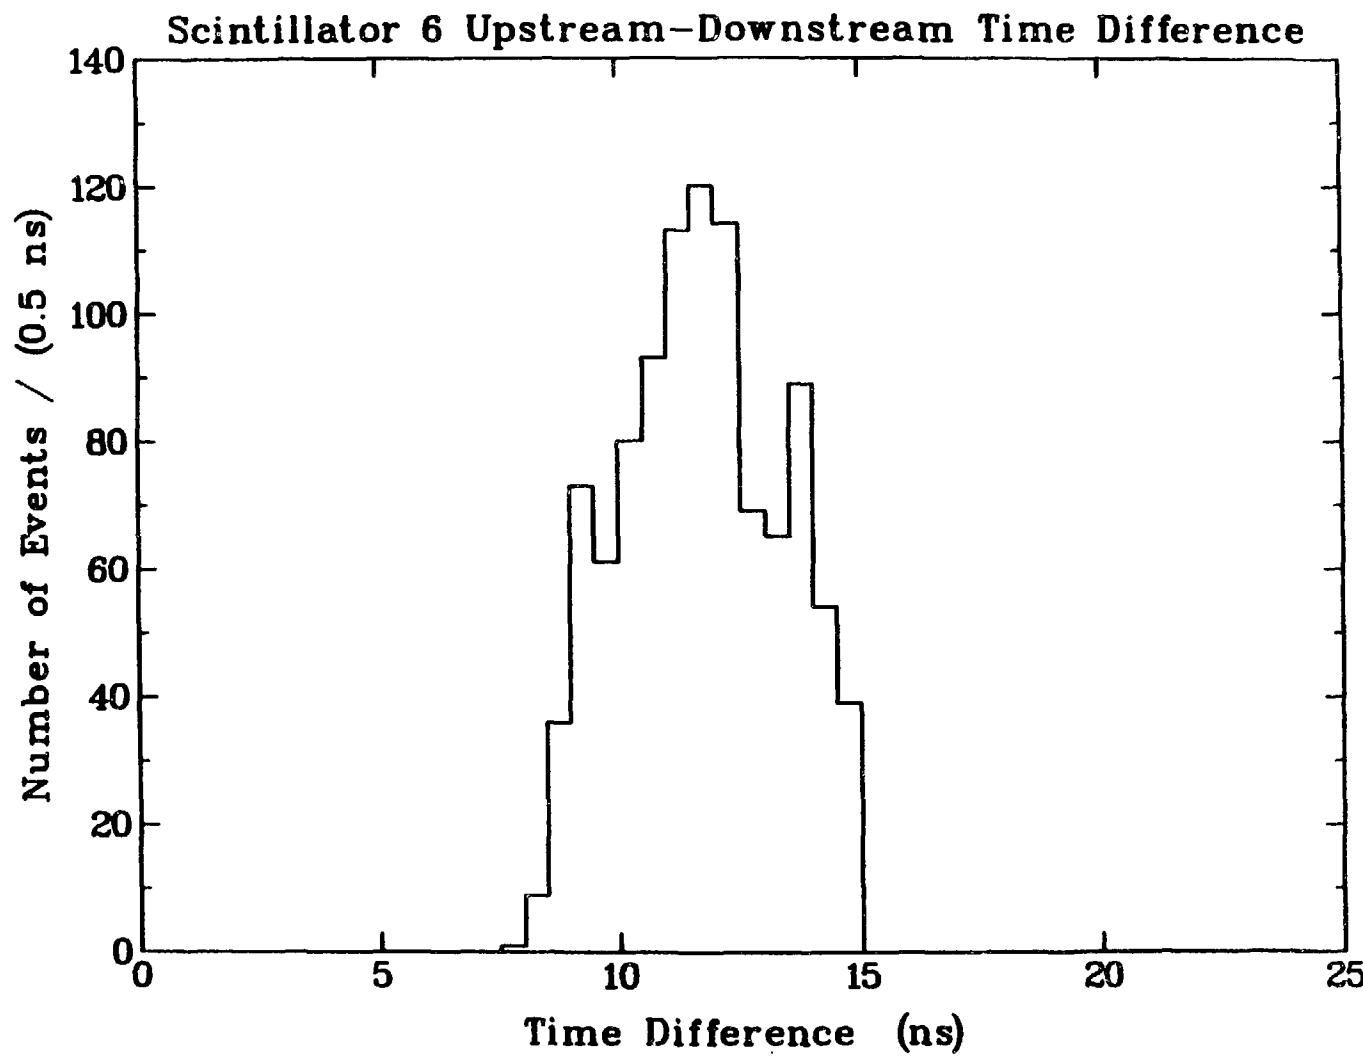
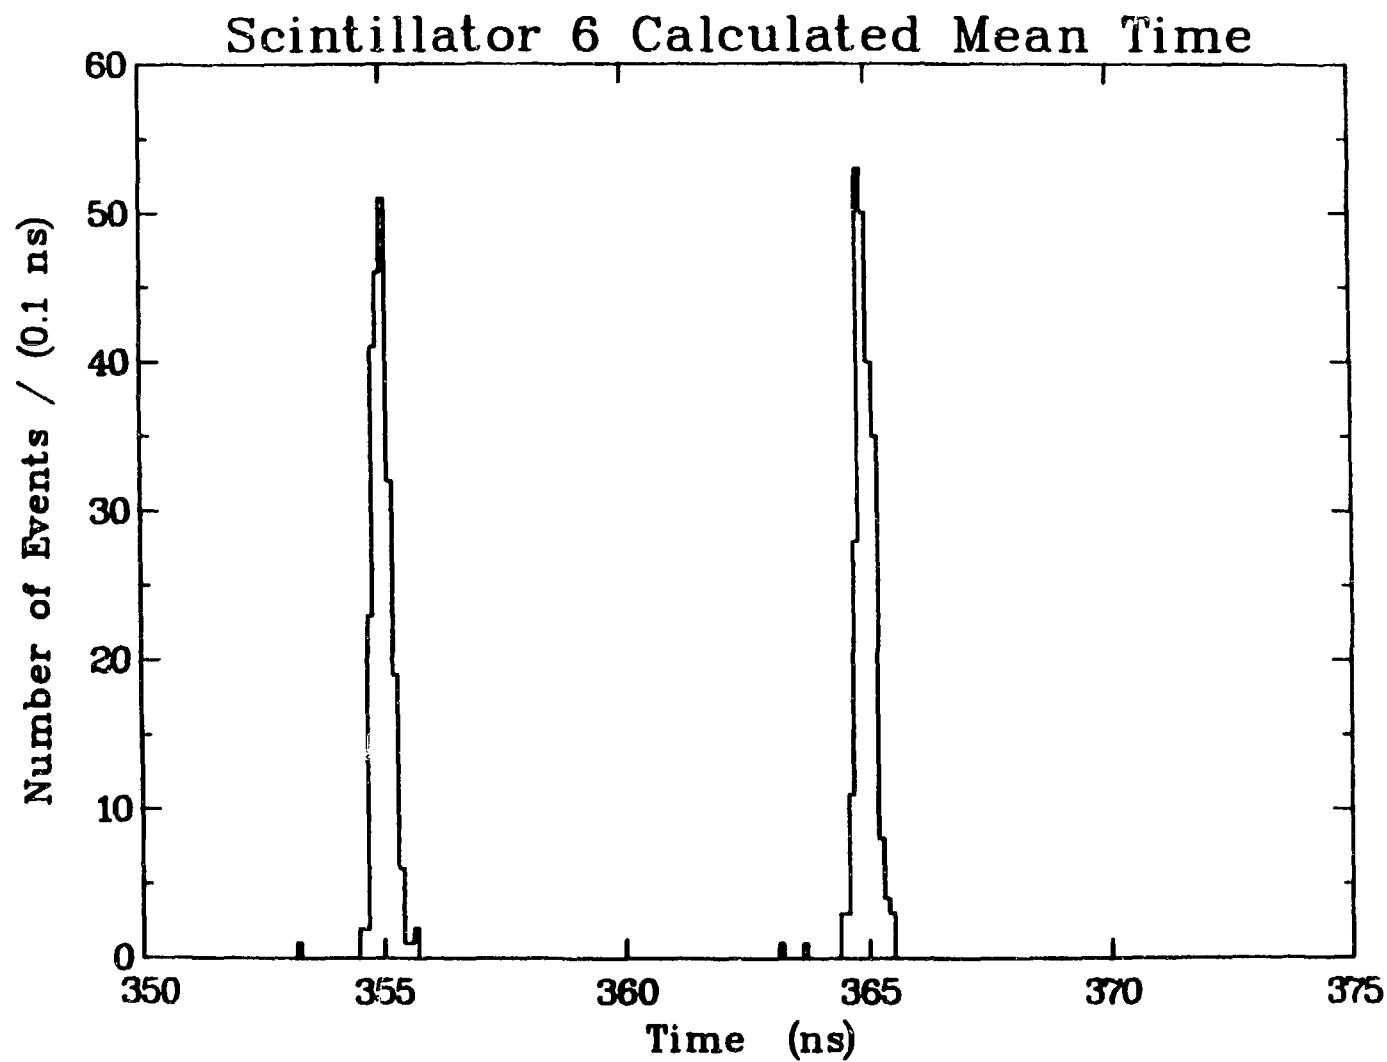


Figure 14. Double-peaked mean time spectrum from scintillator number 6.



E. NaI(Tl) Timing Calibration

The NaI(Tl) crystal timing calibration is done in a similar manner as the scintillator timing calibration, except that fewer steps are involved. All face crystal CFD pulses and the corner crystal meantimer pulses are synchronized into the trigger logic. The double-peaked spectra are used to find the timing gains and offsets, just as in the first step of the scintillator timing calibration. There is an added problem with crosstalk between NaI(Tl) timing channels in the TAC modules, and this is corrected in the calibration analysis.

The timing of all NaI(Tl) crystals into the trigger logic is set relative to the I-counter. Positrons, originating from the I-counter, are used to find the NaI(Tl) timing constants. The I-counter provides the Start signal and the pulses from the face crystal CFDs or the corner crystal meantimers give the Stop pulse. The same set of TDC channels are used in this initial step to set the times of all the NaI(Tl) crystals into the trigger logic. A set of row latches are used to differentiate the rows of NaI(Tl) crystals with hits in them. Only nine crystals, one per row, can have their times set at once because of the row latches. These nine crystals are selected so that there is maximal spacing between the crystals and minimal energy shower spreading into other crystals. Ideally, all crystals would have their times set at the same instant so that the time variations between crystals are minimized. The times of all NaI(Tl) crystals set into the trigger logic are within 0.5 ns.

The times of all NaI(Tl) crystals are set into the I/H modules using the flasher system. This is an important part of the initial set-up because the pulses must be contained within the ADC data gates for proper integration of the pulse and pile-up measurements. The flasher system illuminates all crystals at the same time, so this process is done efficiently. These times are set with an accuracy of ~ 5 ns.

Special timing data are analyzed to obtain the timing gains and offsets. This is done using the same trigger requirements as that for the scintillator timing calibration. Double-peaked spectra for each NaI(Tl) crystal are gathered and each peak is fit with a Gaussian curve. The difference in centroid positions determines the TDC gains. The first peak centroid position gives the time offset for that NaI(Tl) crystal channel.

A correction is made to the timing constants because of crosstalk between adjacent timing channels in the TAC modules. Crosstalk occurs when pulses in one channel induce pulses in another channel. Adjacent channels in the electronics are not necessarily adjacent NaI(Tl) crystals. The effect of crosstalk is a time shift in the time spectrum. The NaI(Tl) timing analysis for a crystal is redone if an adjacent channel in the readout chain has a time pulse in it. The shifts in the time spectra caused by crosstalk varied by as much as 0.5 ns.

The photon time is taken as an energy-weighted average of the NaI(Tl) crystals illuminated by the shower,

$$t = \frac{\sum t_i E_i}{\sum E_i} , \quad (\text{IV-4})$$

where the sum is over all the NaI(Tl) crystals in the shower that have their CFD signal time within 5 ns of the high pulse height crystal time, t_i is the time of crystal i , and E_i is the energy in that crystal. The time resolution for single photons is ~ 1.2 ns (FWHM) with photon energies above 28 MeV, and the resolution is degraded for energies below 28 MeV. Figure 15 shows a time difference plot of the two photons emerging from the π^0 decay. The measured time resolution is 1.8 ns (FWHM).

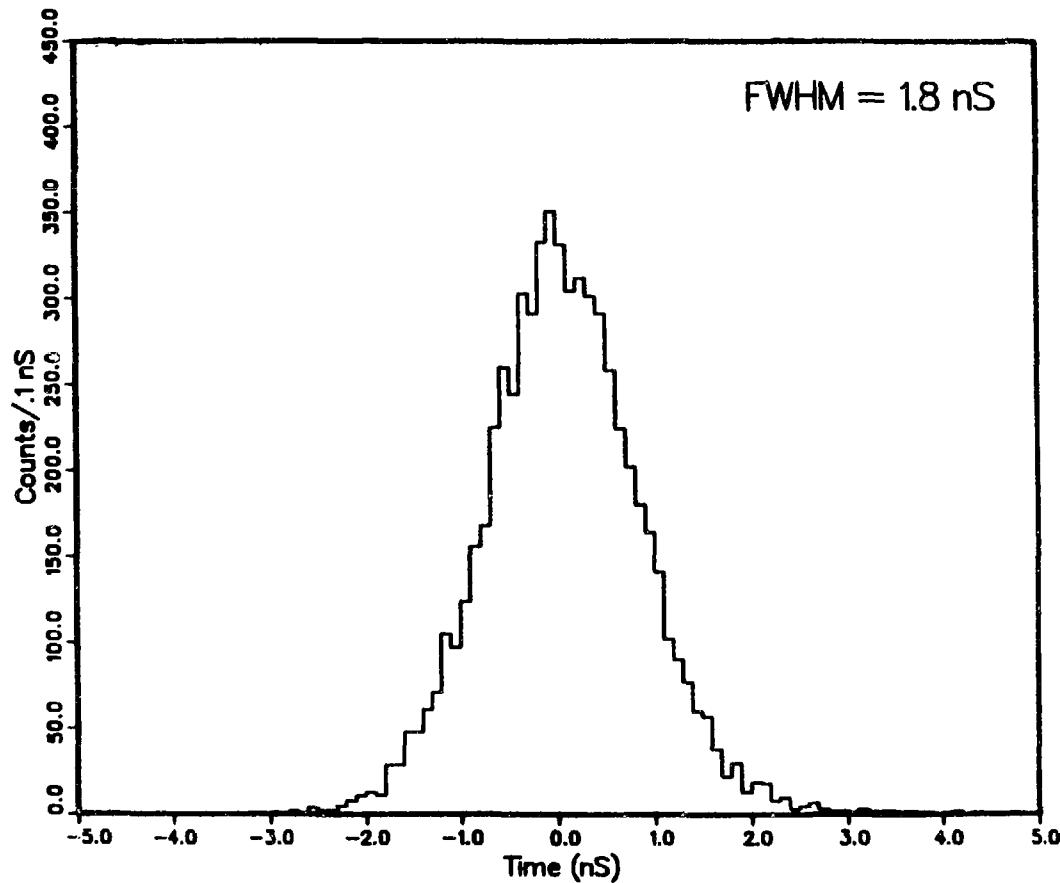
The positron time resolution using the NaI(Tl) crystals is measured to be 1.1 ns (FWHM). Because the NaI(Tl) time resolution is much worse than the time resolution of the scintillators, the positron time is taken from the mean time of the trigger scintillators.

F. Drift Chamber Timing Calibration

The drift chamber timing system is calibrated using positron tracks that originate from the target and contain hits in all eight layers of sense wires. One of the eight wires is removed from the track and the remaining seven sense wires are used to reconstruct the track position. The time of the removed sense wire is compared with the time of the newly reconstructed track and the time difference is plotted. The centroid of the fit is the absolute time offset for the removed sense wire. This procedure is repeated, because the new time offsets affect the reconstructed tracks, until a convergence is reached with the values of the time offsets. This method is used to determine the time offsets for all 728 sense wire channels and these time offsets are used to correct the drift times used in the analysis.

Figure 15. Measured time resolution of the NaI(Tl) array for two photons from the decay $\pi^0 \rightarrow \gamma\gamma$.

Time Difference of Photons for $\pi^0 \rightarrow \gamma\gamma$



G. NaI(Tl) Position Calibration

The position of a particle shower in the NaI(Tl) array is determined by an energy-weighted average of the position of the NaI(Tl) crystals with more than 0.3 MeV deposited in them. The particle position X is given by

$$X = \frac{\sum x_i E_i^\alpha}{\sum E_i^\alpha}, \quad (IV-5)$$

where the sum is over the high pulse height crystal and its eight nearest neighbors, E_i is the energy deposited in crystal i , x_i is the center of crystal i in one of the three spatial dimensions, and α is the weight parameter. The interaction point of the particle with the NaI(Tl) crystal occurs at some distance from the front surface of the NaI(Tl) crystal along the same trajectory. A correction to this position algorithm using the interaction point is made with the quantity d , called the depth parameter.¹⁰⁵

The values of the two parameters, α and d , are found empirically from Monte Carlo studies. Positrons of a fixed energy and origin are generated and uniformly illuminate the detector. The energies, positions, and trajectories are all measured. The energy-weighted position is calculated and because the intersection point on the front face of the NaI(Tl) crystals with the positron trajectory is known, as well as the angle of this trajectory with the NaI(Tl) array, the value of d for a

given energy and particle is determined. The parameter d is energy dependent and is different for positrons and photons. The values of the depth parameter as a function of energy and particle type are given in Table 6. The value of α , the weighting factor, is found by minimizing the position resolution of the reconstructed position with the actual intersection point on the NaI(Tl) crystal face. The value of α is 0.65 for positrons and 0.55 for photons.

The position algorithm is modified when corner crystals are involved. The particle position with high pulse height crystals in rows closest to the corners has the corner crystal position information included. If a high pulse height crystal is located in the corner, the position along the beam direction is obtained from neighboring face crystals, if available. Otherwise, the center of the corner crystal is used for position information.

The position resolution in the NaI(Tl) crystals is measured using positrons and comparing the particle position in the NaI(Tl) array with the extrapolated drift chamber track position. The NaI(Tl) position resolution is 4.1 cm (FWHM) in each of the three spatial directions. The position resolution is found to be relatively insensitive to changes in the value of the weighting parameter, α .

H. NaI(Tl) Energy Stabilization

Once the energy calibration is performed and all the gains are determined, it becomes necessary to monitor and correct variations in the energy gains as a function of time. Drifts in the gains are caused by

TABLE 6

VALUES OF THE DEPTH PARAMETER, d, AS A FUNCTION OF ENERGY AND PARTICLE TYPE

Positrons		Photons	
Energy (MeV)	d (cm)	Energy (MeV)	d (cm)
0.0	0.00	0.0	5.75
9.8	1.44	14.7	6.33
19.4	2.69	31.4	6.98
28.8	3.58	49.1	7.73
44.5	4.92	88.1	8.70
58.7	5.43	126.9	9.43
92.8	6.82		

such factors as temperature fluctuations and electronics instabilities. The flasher system is designed to follow the gain variations, but it is not capable of doing this to the desired accuracy required by this experiment. The decay positron energy spectrum, gathered from the data events, is used to monitor the energy gain fluctuations.

The NaI(Tl) electronics is monitored for changes in the pedestals and gain values approximately every two hours during the data-taking period. The pedestals are checked and updated in the LSI-11/23 microcomputers so that the sparse data scan can be performed properly. The energy gains are updated on-line using the flasher system. The flasher system is used to exercise the system and check that all the channels, both energy and timing, are working properly. Histograms of energy and timing data from each NaI(Tl) crystal using the flasher, as well as pedestal values, are written to computer tape and analyzed off-line.

Later analysis showed that the flasher system is not able to track the energy gain variations averaged over an entire NaI(Tl) quadrant to better than about 5%, but the flasher is used to monitor changes in individual crystal energy gains relative to the quadrant average. The inability of the flasher system to track the energy gain fluctuations to better than several percent is probably due to variations in the outputs of the flashtube monitor photodiodes.

The positron energy spectrum edge at 52.8 MeV is used to monitor and correct the quadrant energy gains, while any change in the individual crystal gains is still corrected by the flasher. The positron energy spectra are collected from the data, using the $\mu \rightarrow e\gamma$ trigger events, and

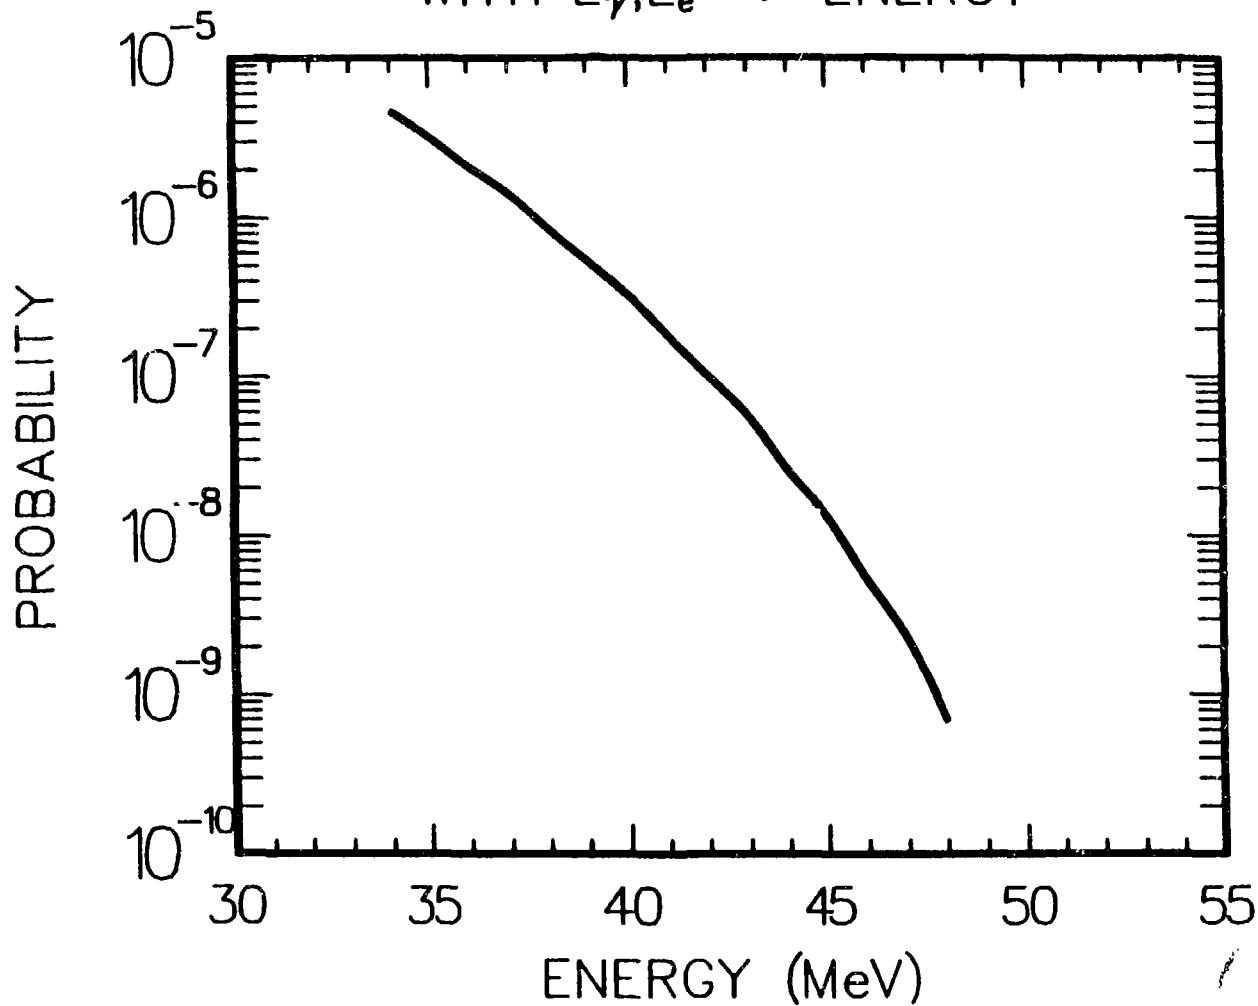
are fit with a fifth order Chebyshev polynomial. Chebyshev polynomials are chosen because they form an orthogonal set and give the smoothest curve through a discrete set of points. The fifth order polynomial is chosen because it gives the best fit to the positron energy edge of a reference spectrum.

The positron energy spectra chosen as references for each NaI(Tl) crystal is taken from a normal data run. That data run is located closest in time to the calibration data because the energy gain variations are assumed to be small. All other positron energy spectra from other runs are referenced to the spectra from this particular run. The energy gain that minimizes the chi-square value of the fit to these spectra is chosen as the new energy gain value. The gain correction is applied on a quadrant-by-quadrant and a run-by-run basis in the data analysis. With this method, the corrected energy gains are stable to about 0.5% over the entire data-taking period.

The final check of the energy calibration and stabilization is to compare the number of $\mu \rightarrow e\gamma\nu\bar{\nu}$ events observed in the $\mu \rightarrow e\gamma$ data sample with the number expected. The expected number of inner bremsstrahlung events is highly dependent on the energy threshold, as shown in Figure 16. In order to observe the expected amount of $\mu \rightarrow e\gamma\nu\bar{\nu}$ events, the energy measurement must be very accurate. The expected number of $\mu \rightarrow e\gamma\nu\bar{\nu}$ events is 3960 ± 90 and this is in agreement with the $3470 \pm 80 \pm 300$ observed events.²⁰ Therefore the absolute energy calibration is known to better than 0.25%.

Figure 16. Plot of the probability of $\mu \rightarrow e\gamma\nu\bar{\nu}$ events as a function of energy threshold for both positron and photon.

PROBABILITY FOR $\mu \rightarrow e\gamma\nu\bar{\nu}$
WITH $E_\gamma, E_e >$ ENERGY



I. Timing Stabilization

The NaI(Tl) timing constants are monitored as a function of time using the flasher system and the special timing data. Double-peaked timing spectra from the flasher are written to computer tape and analyzed off-line to determine the corrections to the time offsets of each NaI(Tl) channel for a given run. The special timing data are taken every three days and analyzed in the same manner as outlined in the previous section. The timing gains and absolute time offsets are updated from the special timing data for use in further data analysis. The timing of the NaI(Tl) crystals is stable to better than 0.25 ns.

The scintillator timing constants are monitored by repeating the timing calibration. The special timing data are taken once per day, and the scintillator timing gains and offsets are updated in the off-line analysis. The scintillator time drifts are about 200 ps and this degrades the time measurement by about 50 ps. The scintillator times are stable to about 80 ps after corrections.

CHAPTER V

DATA ANALYSIS

This chapter presents the data analysis used to find the number of $\mu \rightarrow e\gamma\gamma$ events. The first section describes the two pile-up detection and rejection methods. The next sections detail the procedures used to separate signal from background events. The data are analyzed in several steps. The first pass through the data is an initial reduction using loose requirements on the energy, time, and geometry. The second pass places a series of tighter restrictions on the data, eliminating events that may have a positron annihilating in flight, and applying a cut on energy sums to discard background events. Finally, a maximum likelihood analysis estimates the number of $\mu \rightarrow e\gamma\gamma$ events in the remaining data sample.

The data are arranged in three categories because of the different hardware conditions during the data-taking periods. The January data sample, taken during December 28, 1983 to January 20, 1984, does not have any information available from the pile-up ADC system, so the rejection method uses timing from the NaI(Tl) crystals. The June data sample, taken from June 29, 1984 to July 16, 1984, also does not contain any information from the pile-up ADC system, but includes other detector modifications. The summer data sample, taken from July 17, 1984 to October 2, 1984, contains over two-thirds of the total number of events, and uses the energy measurement from the pile-up ADC system in the rejection criterion.

Because of the different trigger requirements, the two-quadrant $\gamma\gamma$ events are also analyzed separately from the three-quadrant $\gamma\gamma$ events.

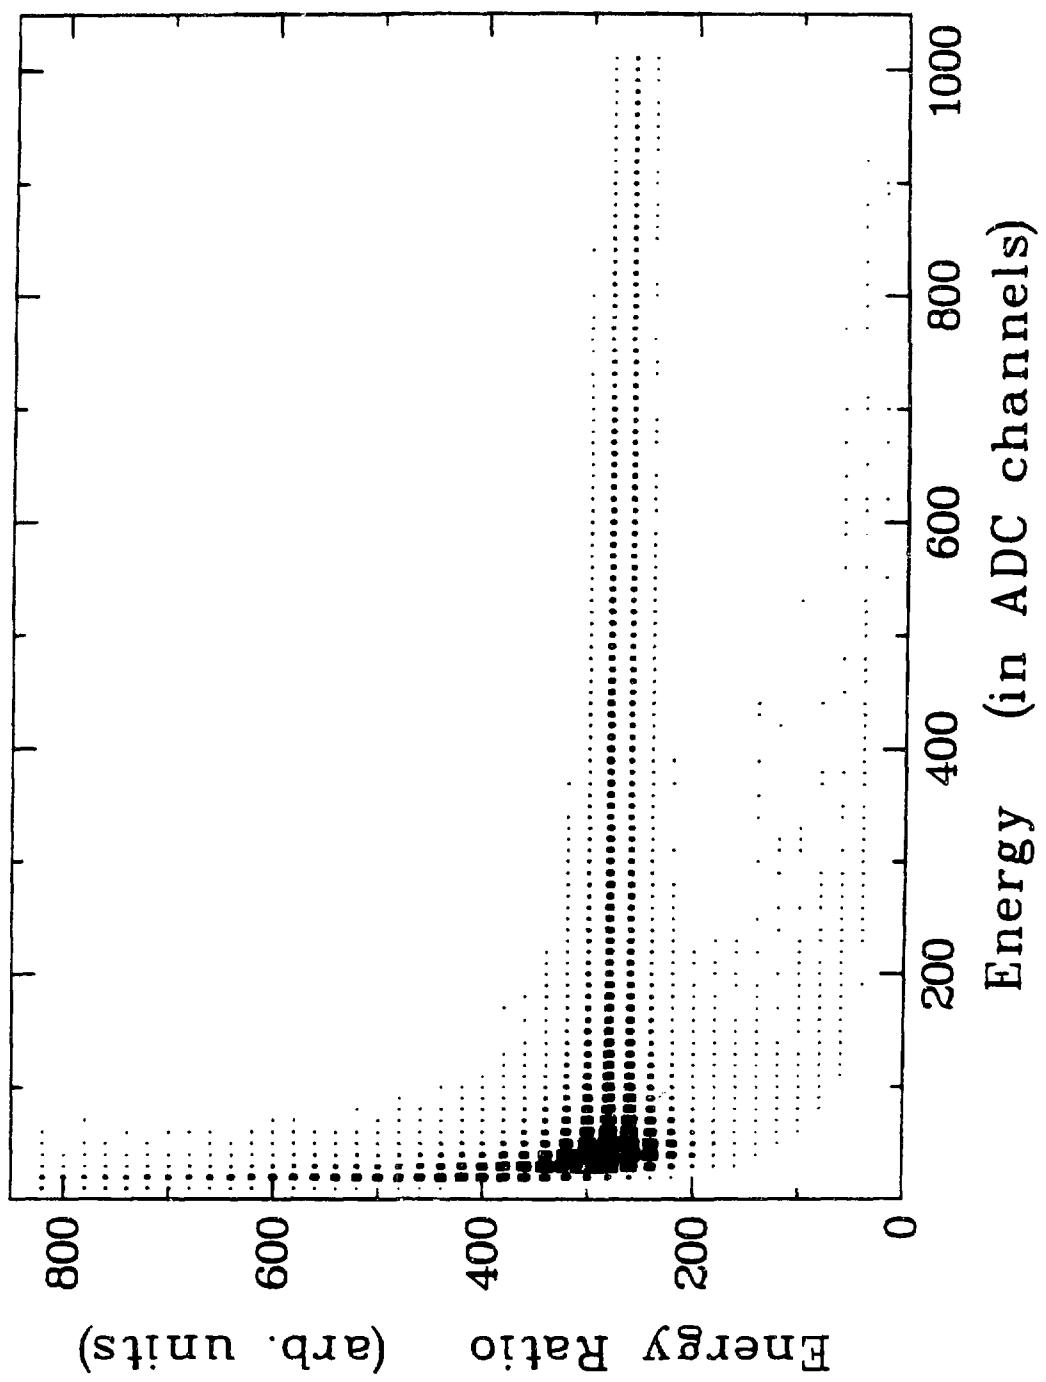
A. Pile-up Rejection

Two distinct methods are used to distinguish piled-up pulses in the NaI(Tl) crystals. One uses the timing of the crystal pulses with at least 5 MeV deposited in them. The other uses the ratio of the energy in the pile-up ADC system to the energy in the primary ADC system for a given NaI(Tl) crystal.

The first method searches for energy deposited in the NaI(Tl) crystals that is out-of-time with the trigger. The time of the NaI(Tl) crystal with the highest pulse height in a particle shower is compared to the times of its nearest 24 neighbors. If the time difference is greater than 5 ns, the NaI(Tl) crystal is considered to be "piled-up." A crystal is also "piled-up" if it has more than 7 MeV deposited in it and has a pedestal value for its time measurement. For any of these cases where pile-up is detected, the "particle" is discarded, and is not considered in subsequent analysis. About 30% of the entire data sample uses this pile-up rejection method.

The second method uses the energy measurements from both ADC systems to detect piled-up energy. The first step in this process is to find the ratio of the energy measured in the pile-up ADC system to the energy in the primary ADC system for each NaI(Tl) crystal. Figure 17 shows this ratio as a function of the energy measured in the primary ADC system, in units of ADC channels, for a typical crystal. No special data

Figure 17. Plot of the ratio of the energy deposited in the pile-up ADC system to the energy in the primary ADC system versus the energy measured in the primary ADC system in units of ADC channels for a typical crystal.



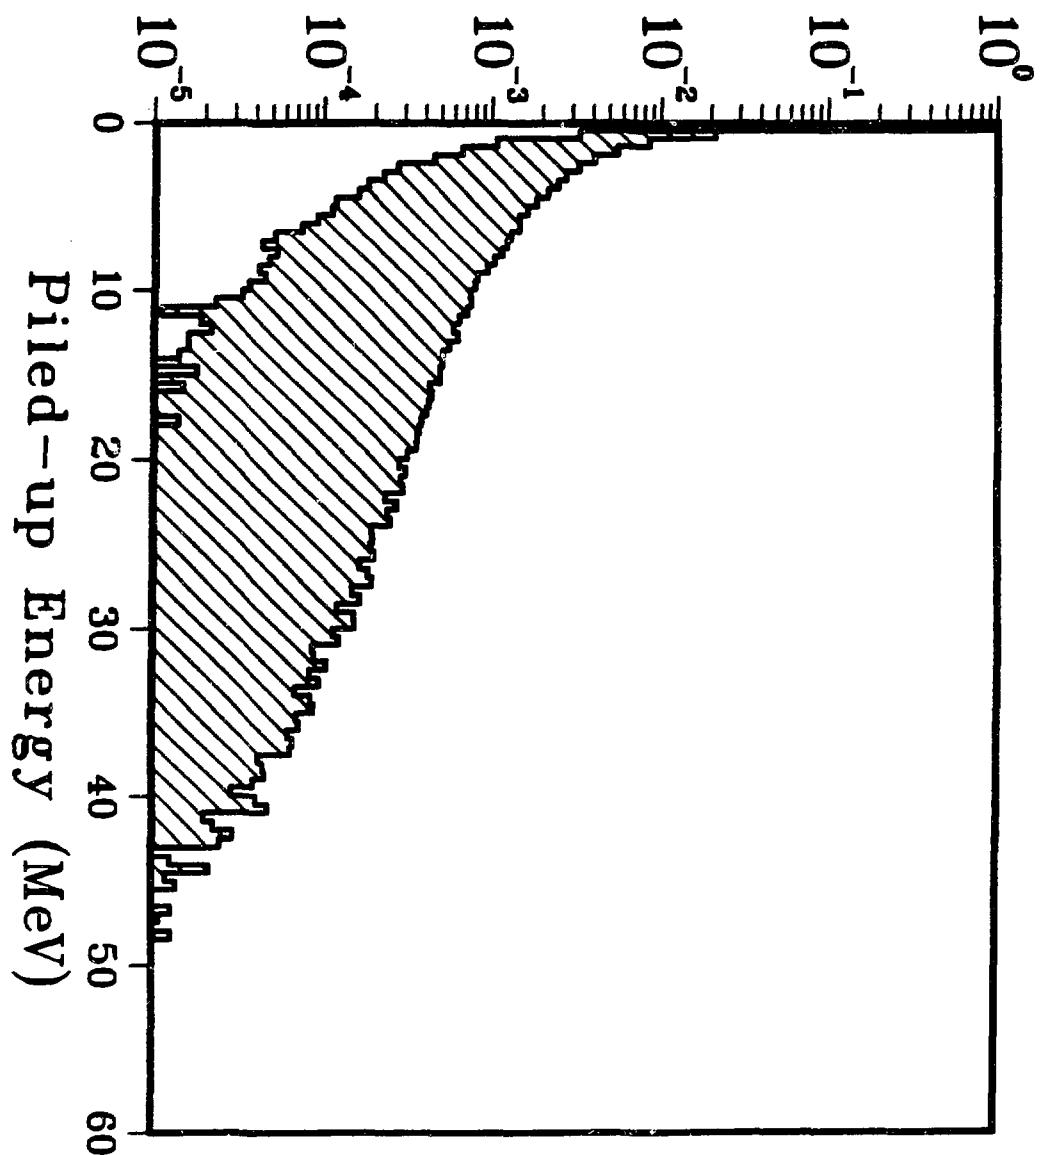
is taken for this ratio determination, and the energy ratios include events with both positrons and photons. Figure 17 shows a narrow band above ADC channel 200 where both ADC systems are presumed to detect no piled-up energy. The second step is to determine the upper and lower bounds of this narrow band for each NaI(Tl) crystal and thereafter to reject as "piled-up" those pulses that fall outside the band. About 95% of the pulses are within the upper and lower bounds.

The amount of pile-up for each high pulse height crystal and its nearest neighbors is found from this energy ratio plot. If the energy ratio is outside the band, then the NaI(Tl) crystal is "piled-up" and one of two possibilities occurs. If the energy in the primary ADC system for a given crystal is greater than 2.5 MeV, then the entire "particle" is eliminated from further analysis. If the energy is below 2.5 MeV, then the energy from that particular NaI(Tl) crystal is thrown out and not included in the energy sum for that particle. This pile-up rejection method is almost 100% efficient for energies above ~ 5 MeV (ADC channel number 100) in the primary ADC system.

The effects of piled-up energy are added to events in the Monte Carlo program. Figure 18 shows a plot of the probability of pile-up in a NaI(Tl) crystal versus the amount of piled-up energy for photons at an instantaneous muon stopping rate of 4.83 MHz. This plot is made separately for both photons and positrons and also for different instantaneous muon stopping rates. It is generated when triggering the Crystal Box detector randomly during the beam gate and measuring the energy in the NaI(Tl) crystals at that time. The darker curve in Figure 18 shows the amount of pile-up remaining after rejecting events

Figure 18. Plot of the probability of pile-up in a NaI(Tl) crystal versus the amount of piled-up energy for photons at 4.83 MHz (instantaneous) muon stopping rate. The second line shows the amount of piled-up energy undetected in the crystals and the shaded region is the amount of pile-up rejected by this method.

Probability/Cluster/.5 Mev



using the two ADC systems. The shaded region is the amount of piled-up energy rejected. For each particle in the Monte Carlo program, the piled-up energy is added according to the probability plot of Figure 18. The amount of piled-up energy is proportional to the muon stopping rate, so the Monte Carlo events must be analyzed separately for each rate.

B. First Pass Data Analysis

The total number of three-quadrant $\gamma\gamma\gamma$ events passed by the hardware trigger is 2.2×10^6 and the number of two-quadrant $\gamma\gamma\gamma$ events is 1×10^5 . The first pass of the data through the off-line analysis reduces the number of candidate events by restricting the values of the analysis variables in a series of cuts. The data cuts become more sophisticated as more information is analyzed from the detector elements. Approximately 95% of the $\gamma\gamma\gamma$ candidate events are eliminated after the first pass analysis. Table 7 summarizes the first pass data reduction.

The first series of data cuts use the time and energy information available from the plastic scintillators. The only NaI(Tl) information used at this point is whether or not a row of NaI(Tl) crystals contains at least one CFD signal. Positron and photon quadrants are identified with this information. For three-quadrant $\gamma\gamma\gamma$ events, there has to be at least one scintillator that registers a hit and at least two nonadjacent NaI(Tl) rows with signals. The plastic scintillators must contain less than 0.25 MeV deposited in them for those quadrants labeled as photon quadrants. The number of surviving events is 1.1×10^6 .

TABLE 7

FIRST PASS DATA REDUCTION
DATA CUTS ARE LISTED WITH THE NUMBER OF SURVIVING EVENTS

Total number of muon stops		1.18×10^{12}
Number of $\mu \rightarrow e\gamma\gamma$ triggers	3 Quadrant	2.2×10^6
	2 Quadrant	1×10^5
Scintillator information		
At least 1 e^+ quadrant		
At least 2 γ quadrants; $E_{\text{scint}} < 0.25$ MeV		1.1×10^6
NaI(Tl) information		
At least 1 e^+ ; at least 2 γ s		
$ t_{\text{scint}} - t_{\text{NaI}} \leq 5$ ns; $ X_{\text{scint}} - X_{\text{NaI}} \leq 15$ cm		
$\sqrt{\sum_i (t_i - \bar{t})^2} \leq 5.5$ ns		2.2×10^5
Drift chamber information		
At least 1 e^+ track reconstructs to a point with a scintillator hit and NaI(Tl) energy		
$ X_{\text{DC}} - X_{\text{NaI}} \leq 15$ cm		
e^+ track $> 3^\circ$ from target plane		2.1×10^5
Calculated variables		
$7 \leq E_e, E_{\gamma_1}, E_{\gamma_2} \leq 65$ MeV		
$40 \leq E_{\text{Tot}} \leq 125$ MeV; $ \vec{\Sigma p} \leq 99$ MeV/c		
$E_{\text{Tot}} + \vec{\Sigma p} \leq 135$ MeV		
$ t_e - t_{\gamma_1} , t_e - t_{\gamma_2} , t_{\gamma_1} - t_{\gamma_2} \leq 5.5$ ns		
$15^\circ \leq \theta_{e\gamma_1}, \theta_{e\gamma_2}, \theta_{\gamma_1\gamma_2} \leq 180^\circ$		
$235^\circ \leq \Sigma \theta_i \leq 365^\circ; 61.5^\circ \leq \alpha \leq 112.5^\circ$		50 920

The next series of cuts use the energy, time, and position of the NaI(Tl) crystals along with the plastic scintillator information. At this point in the analysis, the two-quadrant $\gamma\gamma\gamma$ events are included with the three-quadrant $\gamma\gamma\gamma$ events. Particles are now identified as positrons and photons by groups of NaI(Tl) crystals, rather than by entire quadrants. At least one positron and at least two photons must be detected for an event to pass. The minimum energy allowed for a particle shower is 7 MeV. The difference in time between the NaI(Tl) signal and the corresponding scintillator signal for a positron must be less than 5 ns. The spatial difference between a NaI(Tl) hit and a scintillator hit must be less than 15 cm. The root-mean-square timing of the combinations of three particles is restricted to values less than 5.5 ns. The number of surviving events after this series of cuts is 2.2×10^5 .

The third set of data cuts use the drift chamber track information to identify positrons and photons, in conjunction with the scintillation counters and the NaI(Tl) crystals. At least one positron track must be reconstructed to the point where the scintillator is hit and the NaI(Tl) crystals have energy deposited in them. The spatial difference between the drift chamber reconstructed track position and the NaI(Tl) positron position must be less than 15 cm. A positron track must be greater than 3° from the plane of the target. Approximately 2.1×10^5 events survive these cuts.

The final series of $\gamma\gamma\gamma$ cuts in the first pass use calculated analysis variables in energy, time, and geometry. The energy loss of the positron, as it passes through material, is calculated and added to the energy measured in the NaI(Tl) crystals. Each particle energy is

restricted to values between 7 and 65 MeV. The total energy of the three particles, E_{Tot} , must be between 40 and 125 MeV. The two photon momentum vectors are determined from the energy and the position in the NaI(Tl) array. The positron momentum vector is found from the energy in the crystals and the track-reconstructed position from the drift chamber projected on the NaI(Tl) array. The magnitude of the vector momentum sum, $|\Sigma p|$, must be less than 99 MeV/c. The sum of the total energy and the vector momentum is constrained to be $E_{Tot} + |\Sigma p| \leq 135$ MeV. After time-of-flight corrections from the I-counter to the particle origin on the target, the time differences between any two of the three particles is restricted to be within ± 5.5 ns. The smallest angles between the particle trajectories must be within 15° and 180° . The sum of these three smallest angles between the particles must be between 235° and 365° . An angle that measures the coplanarity of the three particles, which will be described more precisely in the next section, is allowed to have values between 61.5° and 112.5° . The total number of events surviving the first pass is 50 920.

C. Second Pass Data Analysis

The second pass analysis imposes a series of cuts to reject most of the background events without sacrificing much of the detector acceptance. The first of the data cuts eliminates the two-quadrant err triggers where both photons are in the same quadrant, as discussed below. The next series of cuts requires that the events pass software cuts that resemble the hardware trigger using the off-line analysis. The third series of

data cuts discard most of the background events by restricting particle energy sums. Finally, more stringent limits are placed on several analysis variables. The reduction in the number of candidate events during the second pass analysis is shown in Table 8.

Two-quadrant $e\gamma\gamma$ events with both photons in the same quadrant are rejected from further analysis. Most of these events have a single photon that is misidentified as two separate photons. The photon shower "splits" into two distinct parts, separated by at least two NaI(Tl) crystals in order to be identified as two different particles. The $\mu \rightarrow e\gamma\gamma$ rate, as calculated by the most general local interaction²⁶ and given in Equation II-6, is proportional to $(1 - \cos\theta)^2$, where θ is the angle between the two photons. The topology of an event with two photons in close proximity is not a favorable one, so little detector acceptance is expected to be lost. This cut eliminates 16% of the candidate events, leaving 42 624 events, with only a 1% loss in the detector acceptance for $\mu \rightarrow e\gamma\gamma$ events.

The next series of cuts impose the trigger requirements using the off-line analysis. The data are split into two groups, the January data and the summer data, because of different hardware requirements used. The first cut discards events that have muon decays in the I-counter. The next cut restricts events to have their origin within 7.3 cm from the center of the target. The next several data cuts use the energy measurements. The average quadrant CFD thresholds are found from plots of the energy contained in the high pulse height crystal for a given particle shower. For the January data sample, the left, top, and bottom quadrants have CFD thresholds at 7.8 MeV, and the right quadrant threshold is at

TABLE 8

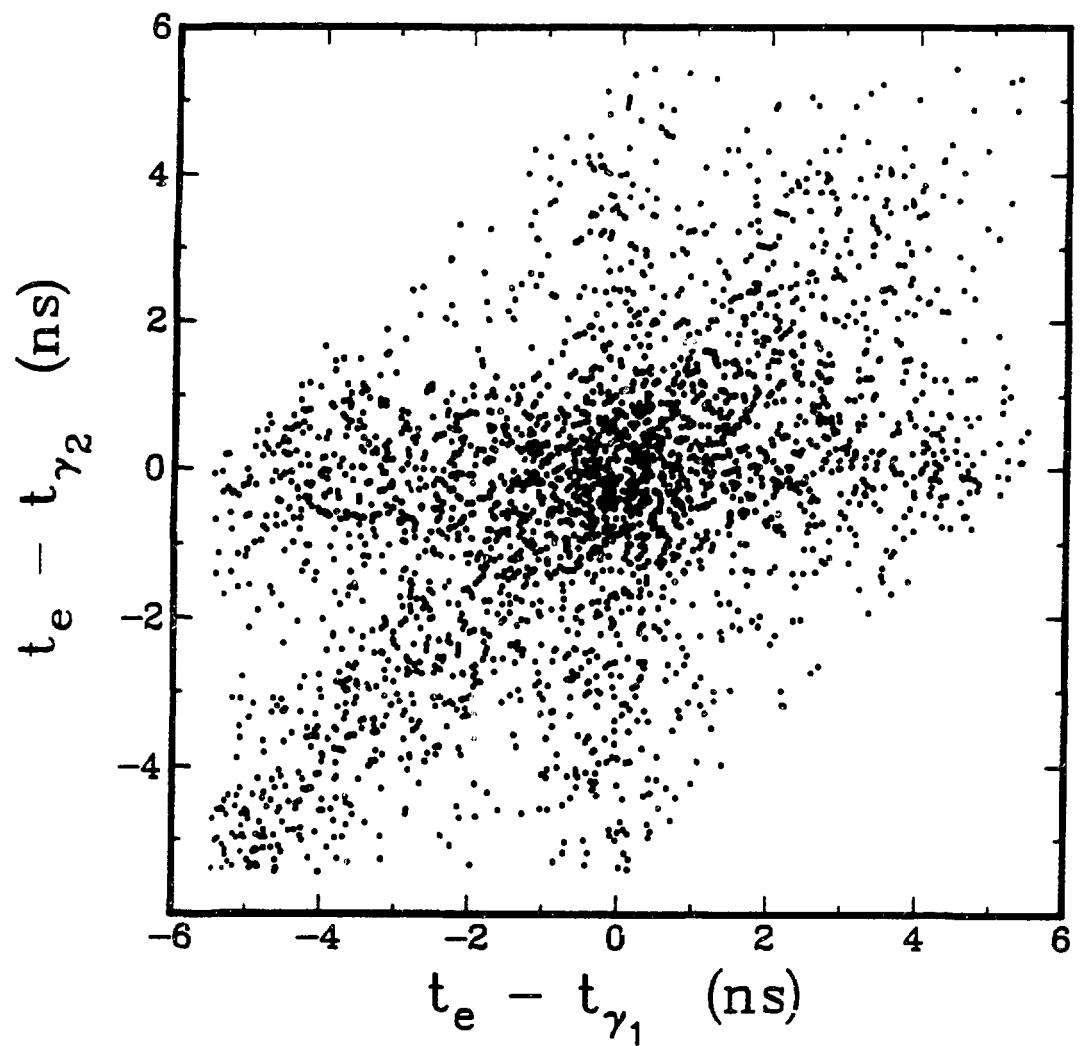
SECOND PASS DATA REDUCTION
THE NUMBER OF EVENTS SURVIVING CUTS ARE SHOWN FOR DATA SAMPLES AND TRIGGERS

Data Sample	Second Pass	Trigger Cuts	Background Cuts	$E_T > 20$ MeV Short Track	Final Cuts
<u>January</u>					
3 Quad	6055	1392	87	37	2
2 Quad	18	9	8	0	0
<u>June</u>					
3 Quad	7710	906	101	49	0
2 Quad	208	94	15	6	0
<u>Summer</u>					
3 Quad	27 891	4269	448	245	5
2 Quad	742	413	52	27	2
Total	42 624	7083	711	364	9

8.2 MeV. In the summer data, the left, top, and bottom quadrants have an energy threshold at 8.0 MeV and the right quadrant has a threshold at 8.2 MeV. The reason for the right quadrant hardware threshold being consistently larger is unknown. All particle energies must be greater than these thresholds. Another threshold is found by adding the NaI(Tl) energy values for the three particles. This threshold is measured to be 70 MeV for the January data and 75 MeV for the summer data. The sum of the energy deposited in the NaI(Tl) array must be greater than this threshold amount. The two-quadrant $e\gamma$ triggers, in satisfying the $\mu \rightarrow e\gamma$ trigger, have an energy threshold of 33 MeV in a quadrant. This implies a minimum photon energy of 33 MeV for the single photon in a quadrant. All of these cuts reduce the data sample to 7083 events.

An identification of the background processes can be made from Figure 19, which is a plot of the time difference between the positron and the higher-energy photon, $t_e - t_{\gamma_1}$, versus the time difference between the positron and the lower-energy photon, $t_e - t_{\gamma_2}$. The data displayed in Figure 19 includes an additional constraint that the angle between the two photons be greater than 40° . This constraint is used only for this plot and is not used in any of the data analysis. The reason for this constraint is to show clearly the different types of background processes. Events scattered uniformly over the plot are random coincidences of the three particles. The diagonal band corresponds to two photon coincidences. Contained in this band are random coincidences between a positron and a photon, where the photon is misidentified as two separate, time-correlated photons. The horizontal and vertical bands correspond to positron and photon coincidences. Most of these coincidences are muon

Figure 19. Plot of the time differences between the positron and the higher-energy photon (γ_1), $t_e - t_{\gamma_1}$, and the positron and the lower-energy photon (γ_2), $t_e - t_{\gamma_2}$. An additional cut requiring the angle between the two photons be greater than 40° is imposed to show clearly the different background processes.



inner bremsstrahlung decays, $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu \gamma$, in random coincidence with a second photon. Another possible contribution to the horizontal and vertical bands is the process where the positron is misidentified as both a positron and a photon. The central region is comprised of positron and two photon coincidences, which are combinations of the above background processes plus any signal events from either $\mu^+ \rightarrow e^+ \gamma\gamma$ or $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu \gamma\gamma$.

From the kinematics of the $\mu \rightarrow e\gamma\gamma$ decay, the sum of the energies from any two of the three particles must be at least one-half the muon mass. This is shown using the conservation of energy for this decay,

$$E_1 + E_2 + E_3 = m_\mu, \quad (V-1)$$

where the subscripts indicate any combination of the three particles. The sum of any two of the three,

$$E_1 + E_2 = m_\mu - E_3, \quad (V-2)$$

must have a minimum value at $m_\mu/2$ because the energy E_3 ranges from 0 to $m_\mu/2$. Therefore the sum of any two particle energies must be at least $m_\mu/2$, or 52.8 MeV. This fact can then be used to eliminate background events from all parts of Figure 19. Restricting the sum of the energies so that $E_{\gamma_1} + E_{\gamma_2} \geq 51$ MeV and $E_e + E_{\gamma_2} \geq 51$ MeV eliminates all but 711 events.

The next cut discards events that have a positron that annihilates in flight. The signature is: energy deposited in the NaI(Tl) crystals and drift chamber wires registering hits only in the first few layers whose projection points toward the NaI(Tl) crystal hits. The final cut in this

series requires that the photon energies be greater than 20 MeV. The latter requirement separates background from $\mu \rightarrow e\gamma\gamma$ events, which have a small probability of containing events with low energies.²⁶ Only 364 events survive these cuts.

Figure 20 shows the distribution of events along the diagonal time projection of Figure 19, $\tau = 2t_e - t_{\gamma_1} - t_{\gamma_2}$. This particular variable is chosen because the effects of the background processes are flat projections in τ . Only prompt processes, such as $\mu \rightarrow e\gamma\gamma$, will have a peak centered at $\tau = 0$ with a width of 1.5 ns (FWHM). Events displayed in Figure 20 have $|\tau| < 4.0$ ns, so only 272 of the 364 events are shown. No appreciable coincidence peak can be seen from this plot.

Two variables are defined to measure the coplanarity and conservation of momentum of an event. One variable is a momentum sum within the decay plane, defined as $p_{\parallel} = |\vec{p}_a + \vec{p}_b + \vec{p}_{ab} \times (\vec{p}_c \times \vec{p}_{ab})|$, where \vec{p}_a and \vec{p}_b are the momenta of the particles that are most nearly perpendicular to each other, \vec{p}_{ab} is the unit vector normal to the plane determined by \vec{p}_a and \vec{p}_b , and \vec{p}_c is the momentum vector of the third particle. The other variable is independent of the energy measurement and is the cosine of the planarity angle, $\cos \alpha = \vec{p}_c \cdot \vec{p}_{ab}$. Figure 21 displays the particle momenta and the variables, p_{\parallel} and α .

The final series of cuts in the second pass analysis place restrictions on the two variables, $p_{\parallel} < 14$ MeV/c and $-0.2 \leq \cos \alpha < 0.2$. Further cuts are placed on the total energy, $90 \leq E_{\text{Tot}} < 108$ MeV, and on the time projection, $-2.5 \leq \tau < 2.5$ ns. These cuts leave only nine candidate events, whose values in the four analysis variables are shown in Figure 22.

Figure 20. Distribution of 272 events in the diagonal time projection variable, τ .

Number of Events / (0.1 ns)

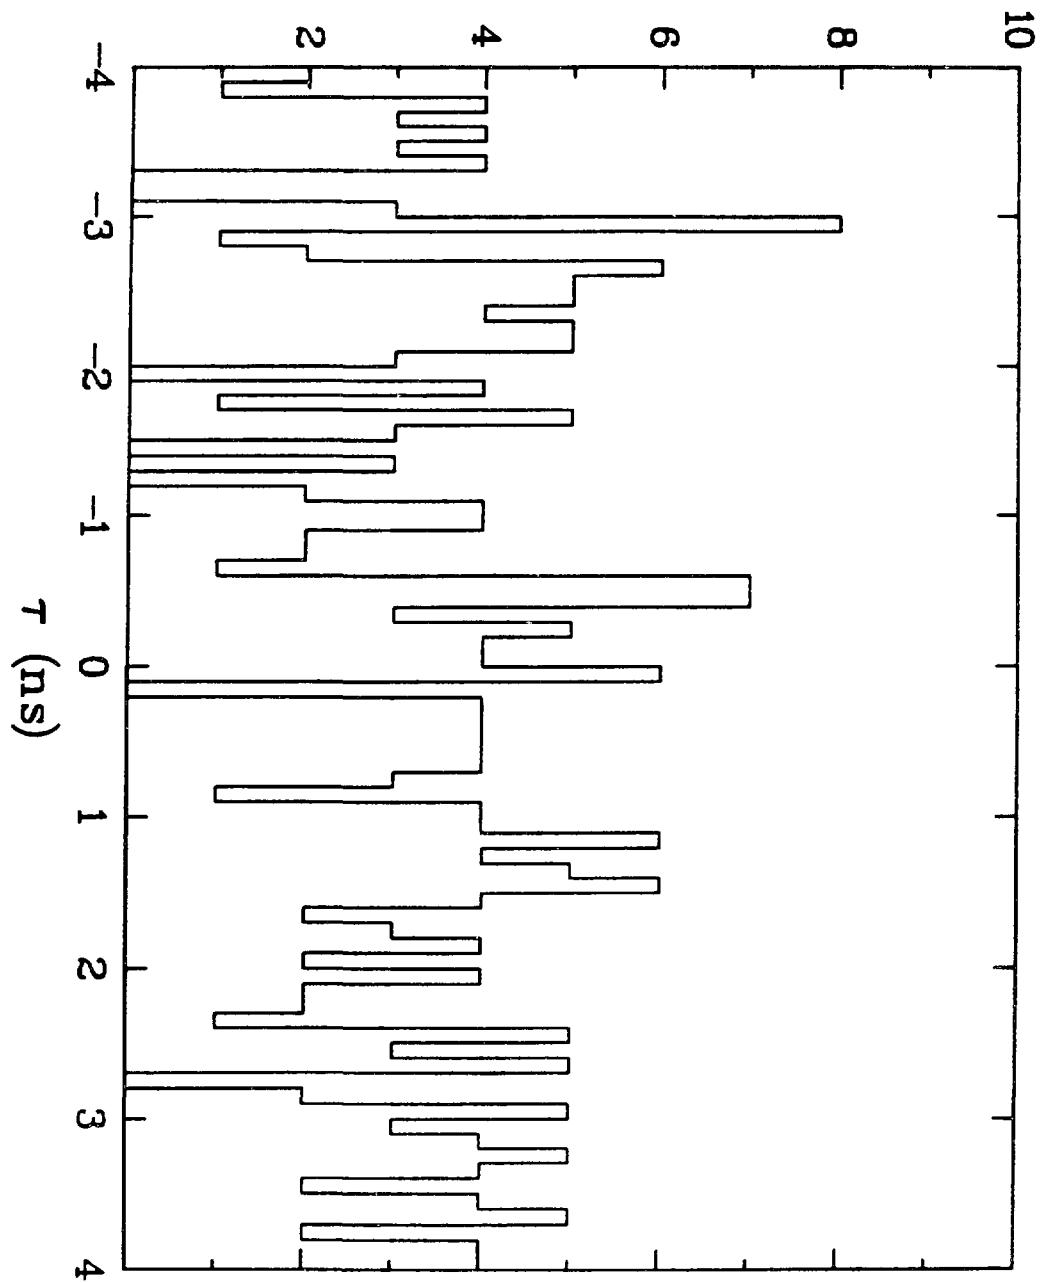
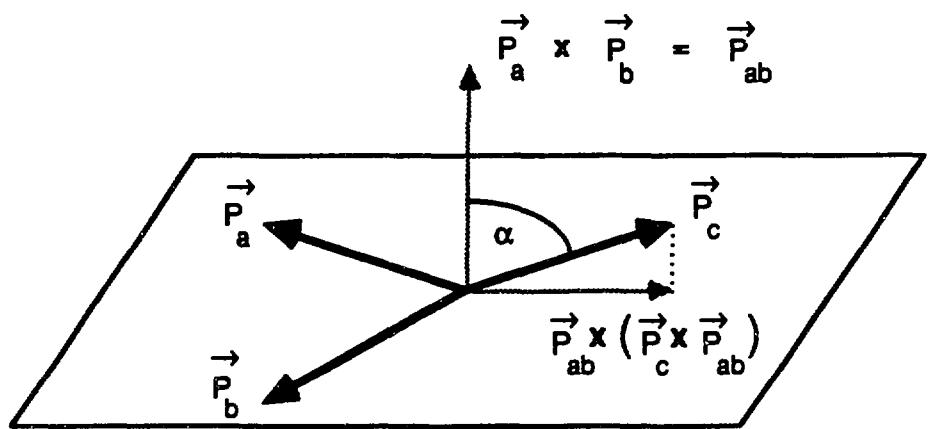
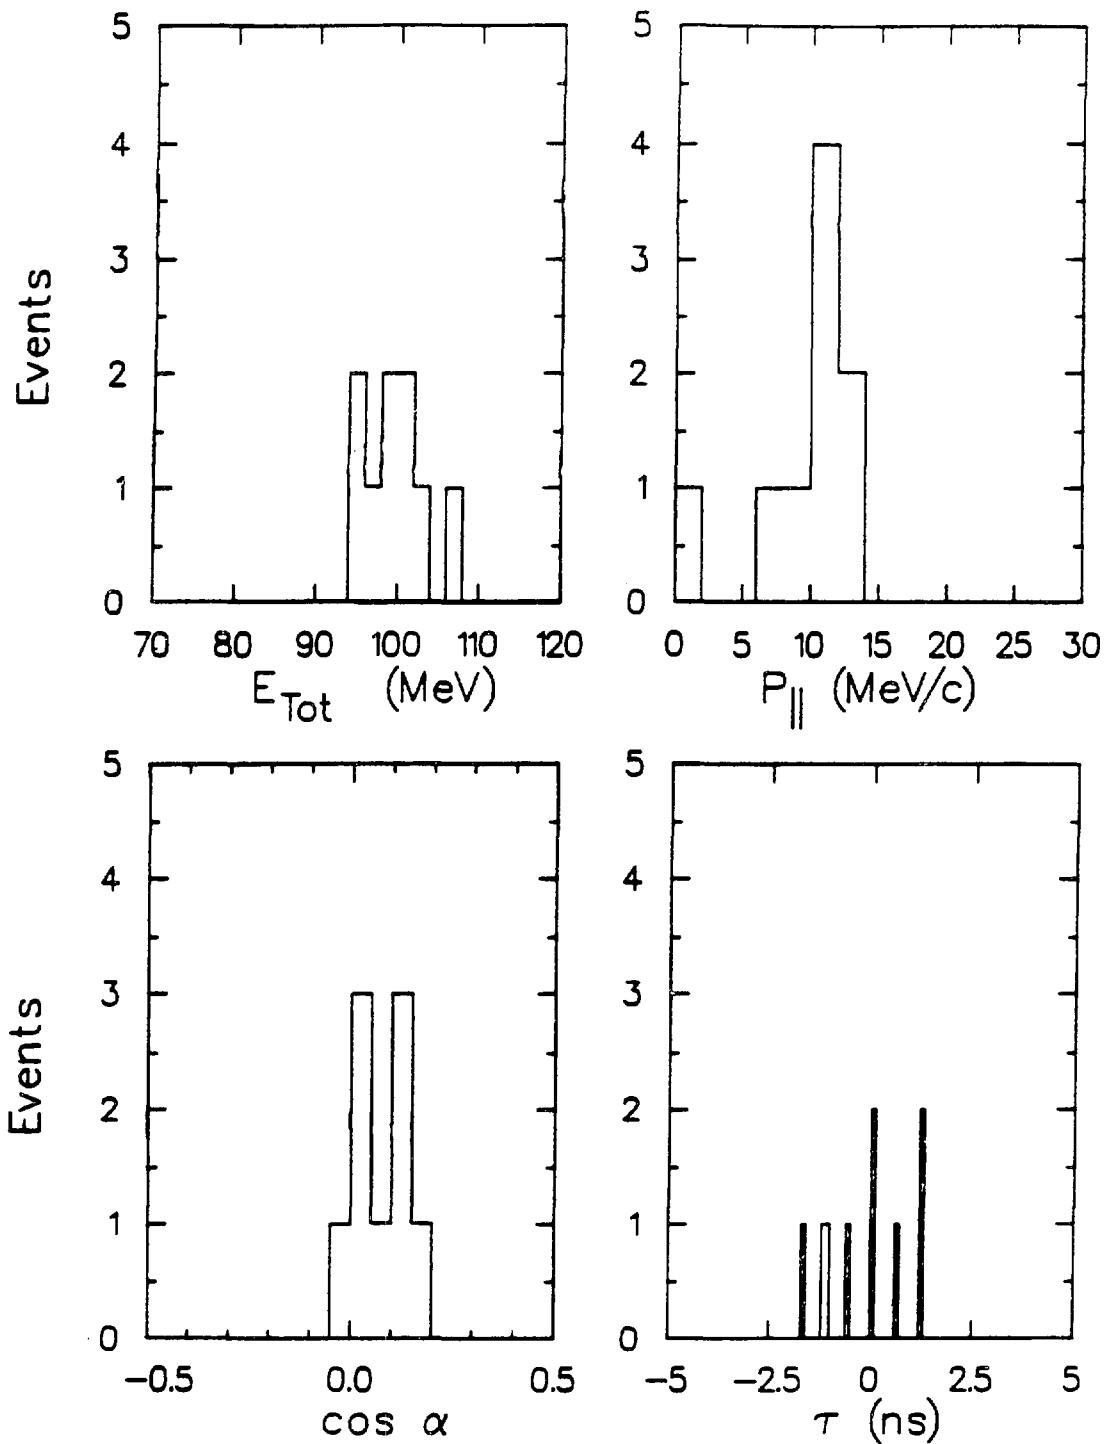


Figure 21. Diagram showing the particle momenta, \vec{p}_a , \vec{p}_b , and \vec{p}_c . The vector sum of the momenta in the decay plane is given by p_{\parallel} , and the planarity angle by α .



**Figure 22. Distributions of the surviving nine data events in
the four analysis variables, E_{Tot} , p_{\parallel} , $\cos \alpha$, and τ .**



D. Maximum Likelihood Analysis

The analysis of the last nine candidate events uses the maximum likelihood method to estimate the number of $\mu \rightarrow e\gamma\gamma$ events. The method uses probability distributions in four variables that are generated from the Monte Carlo program and the data.

The probability of obtaining an event i with measured parameters \vec{x}_i , is given by

$$P(n_{e\gamma\gamma}, \vec{x}) = \frac{n_{e\gamma\gamma}}{N} Q(\vec{x}_i) + \left(1 - \frac{n_{e\gamma\gamma}}{N}\right) R(\vec{x}_i), \quad (V-3)$$

where \vec{x}_i is a vector with four measured variables as the components, $n_{e\gamma\gamma}$ is the estimate of the number of $\mu \rightarrow e\gamma\gamma$ events, N is the total number of events of the data sample within a defined range of \vec{x}_i values, $Q(\vec{x}_i)$ is the normalized probability density for $\mu \rightarrow e\gamma\gamma$ events, and $R(\vec{x}_i)$ is the normalized probability density for background events. The likelihood function is defined as

$$L(n_{e\gamma\gamma}) = \prod_{i=1}^N P(n_{e\gamma\gamma}, \vec{x}) = \prod_{i=1}^N \left[\frac{n_{e\gamma\gamma}}{N} Q(\vec{x}_i) + \frac{n_B}{N} R(\vec{x}_i) \right], \quad (V-4)$$

where $n_B = N - n_{e\gamma\gamma}$ is the estimate of the number of events due to background processes. The vector \vec{x} has four components, which are the variables, E_{Tot} , p_T , $\cos \alpha$, and τ . The best estimates for $n_{e\gamma\gamma}$ and n_B are

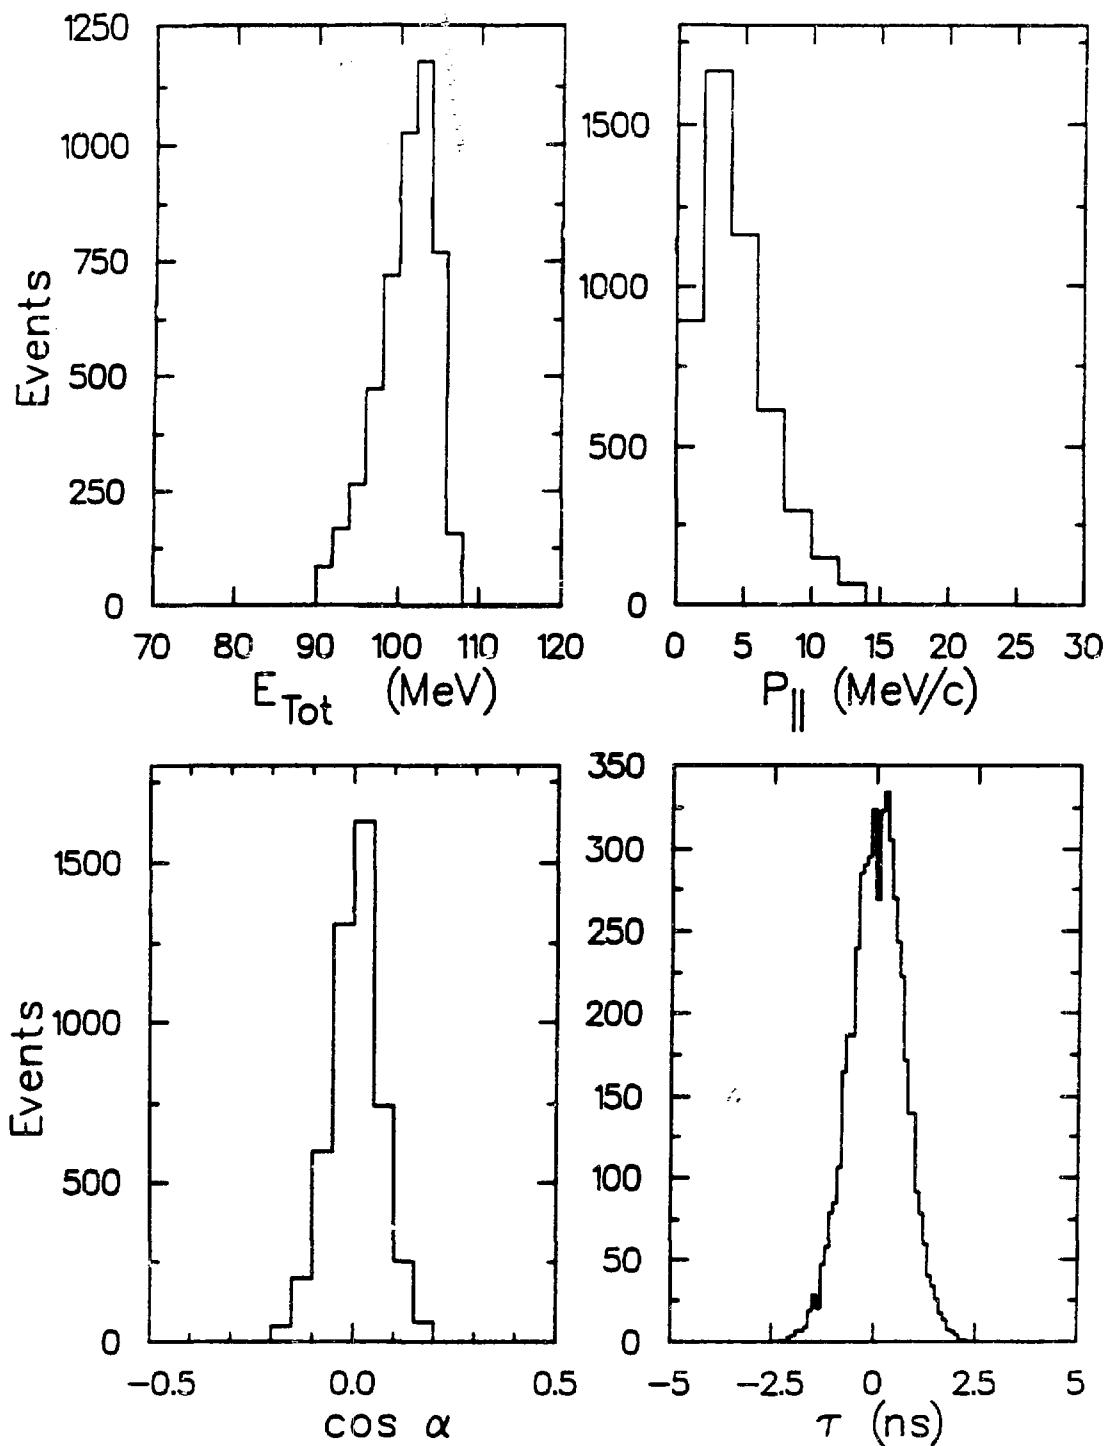
those values that maximize the likelihood function, for positive values of $n_{e\gamma\gamma}$ and n_B .

The remaining data sample is divided into several parts for the analysis. The January and summer data are separated because of the different hardware configurations and pile-up rejection methods. Two-quadrant and three-quadrant events are also kept separate because of the different hardware requirements. Separating the data sample means that different Q and R probability distributions are used for the different parts.

The Q probability distributions for $\mu \rightarrow e\gamma\gamma$ are generated by the Monte Carlo simulations as discussed previously. The probability distributions for these events are somewhat different for the three-quadrant events and the two-quadrant events. The variables, E_{Tot} , p_{μ} , and $\cos \alpha$, are weakly correlated to each other, so a correlated probability distribution is generated using these variables. The Monte Carlo-generated projections of Q in the four analysis variables are shown in Figure 23 for the sample of three-quadrant events with no piled-up energy included.

The data is used to find the amounts of piled-up energy to be added to the particle energy in Monte Carlo events. There are essentially two muon stopping rates in the January data, so the halfway point is chosen to divide the rates. Data events taken at an instantaneous rate lower than 4.7 MHz are counted as a 4.2 MHz rate and data events higher are counted with a 5.22 MHz rate. The number of triggers at each of these two rates is then the relative number of events in the Monte Carlo with piled-up energy at the 4.2 MHz and at the 5.22 MHz rates. The actual difference in

Figure 23. Monte Carlo-generated probability distributions in the four analysis variables, E_{Tot} , $p_{||}$, $\cos \alpha$, and τ , for the three-quadrant $\pi\pi\pi$ events with no added piled-up energy.



the amount of added piled-up energy between the two rates is fairly small. The Monte Carlo events used in conjunction with the summer data contain no added piled-up energy because no evidence of residual pile-up is found between low and high muon stopping rates.

The R probability distributions for background events are taken from the data. Those events with an out-of-time τ value, $2.5 < |\tau| < 5.0$ ns, from the sample containing 364 events comprise the background distributions. The background distributions are not separated into different processes and no distinction is made between data samples, trigger types, or rate. The reason for this is to accumulate a reasonable amount of statistics in each of the probability distributions. The background probability distribution for the time variable has a constant value over the region of interest. The reason for a constant value of τ is that the time projection for background processes is a flat distribution in τ . The background projections in three of the four variables are shown in Figure 24.

The normalized likelihood function is shown in Figure 25 for the nine data events. The function peaks at a value of $n_{\text{e}\gamma\gamma} = 0$. A 90% confidence level (C.L.) for the number of $\mu \rightarrow \text{e}\gamma\gamma$ events is found from the shape of the likelihood function. The value of $n_{\text{e}\gamma\gamma}$ that contains 90% of the area under the curve is n_{90} , and has the value, $n_{90} = 2.9$ events for the entire data sample. This result is insensitive to reasonable variations made to the probability distributions. The quantity n_{90} is used in the calculation of the branching ratio.

**Figure 24. Background probability distributions in three of
the four analysis variables, E_{Tot} , p_{\parallel} , and $\cos \alpha$.**

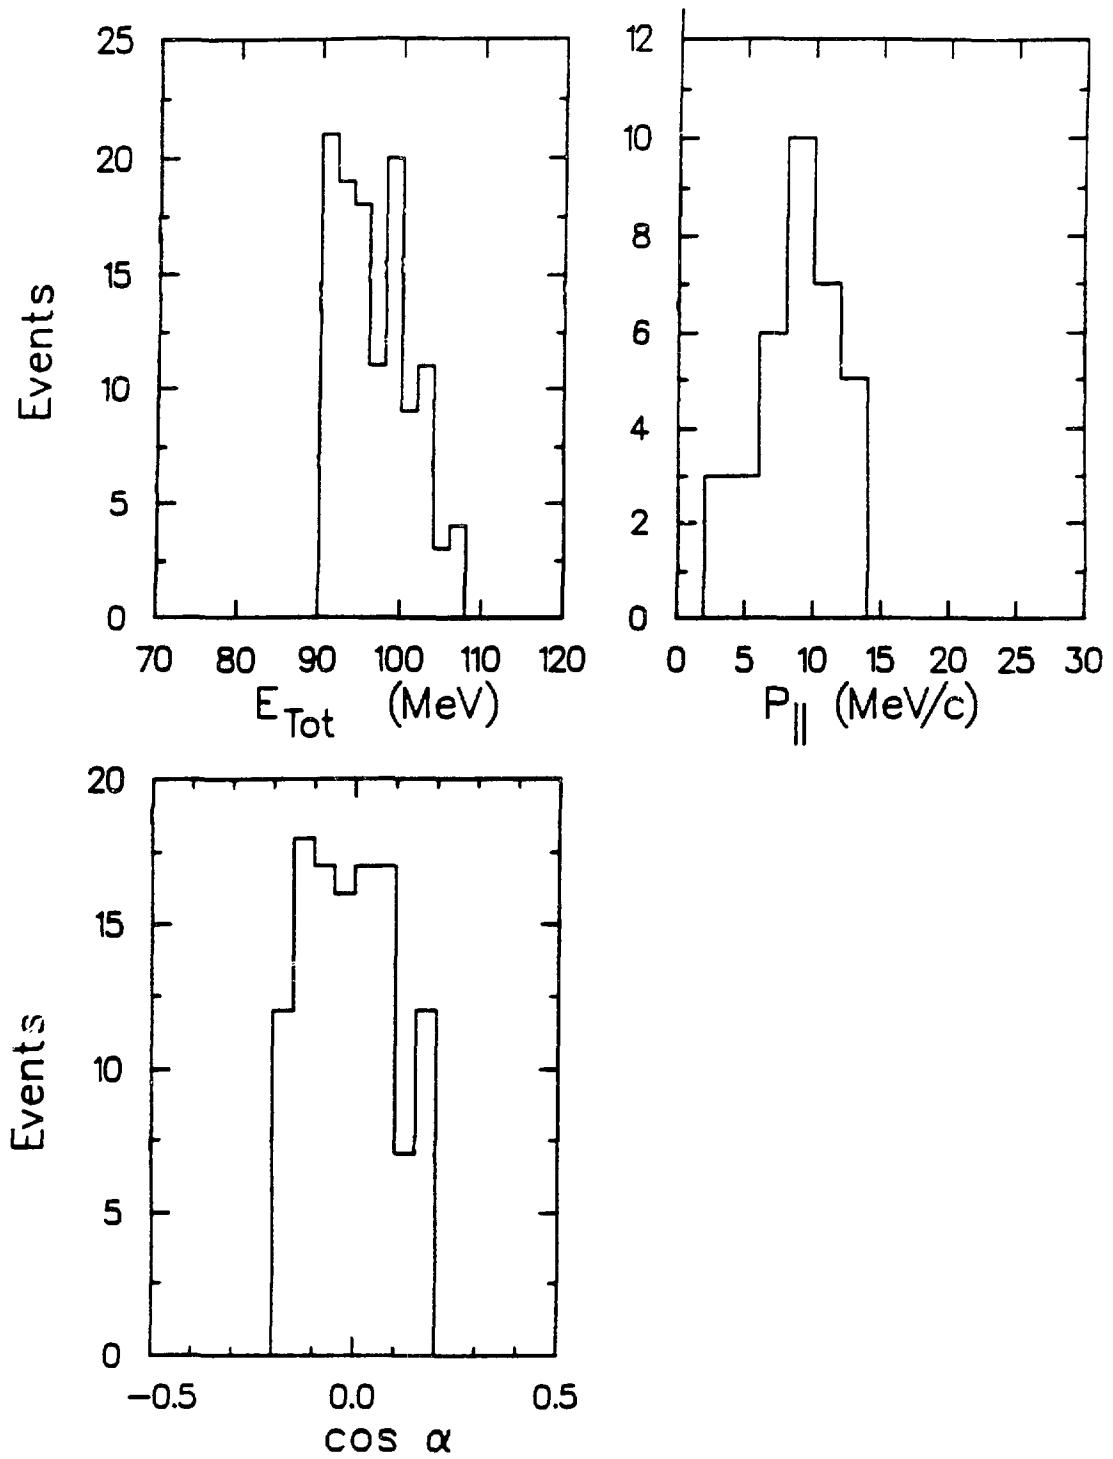
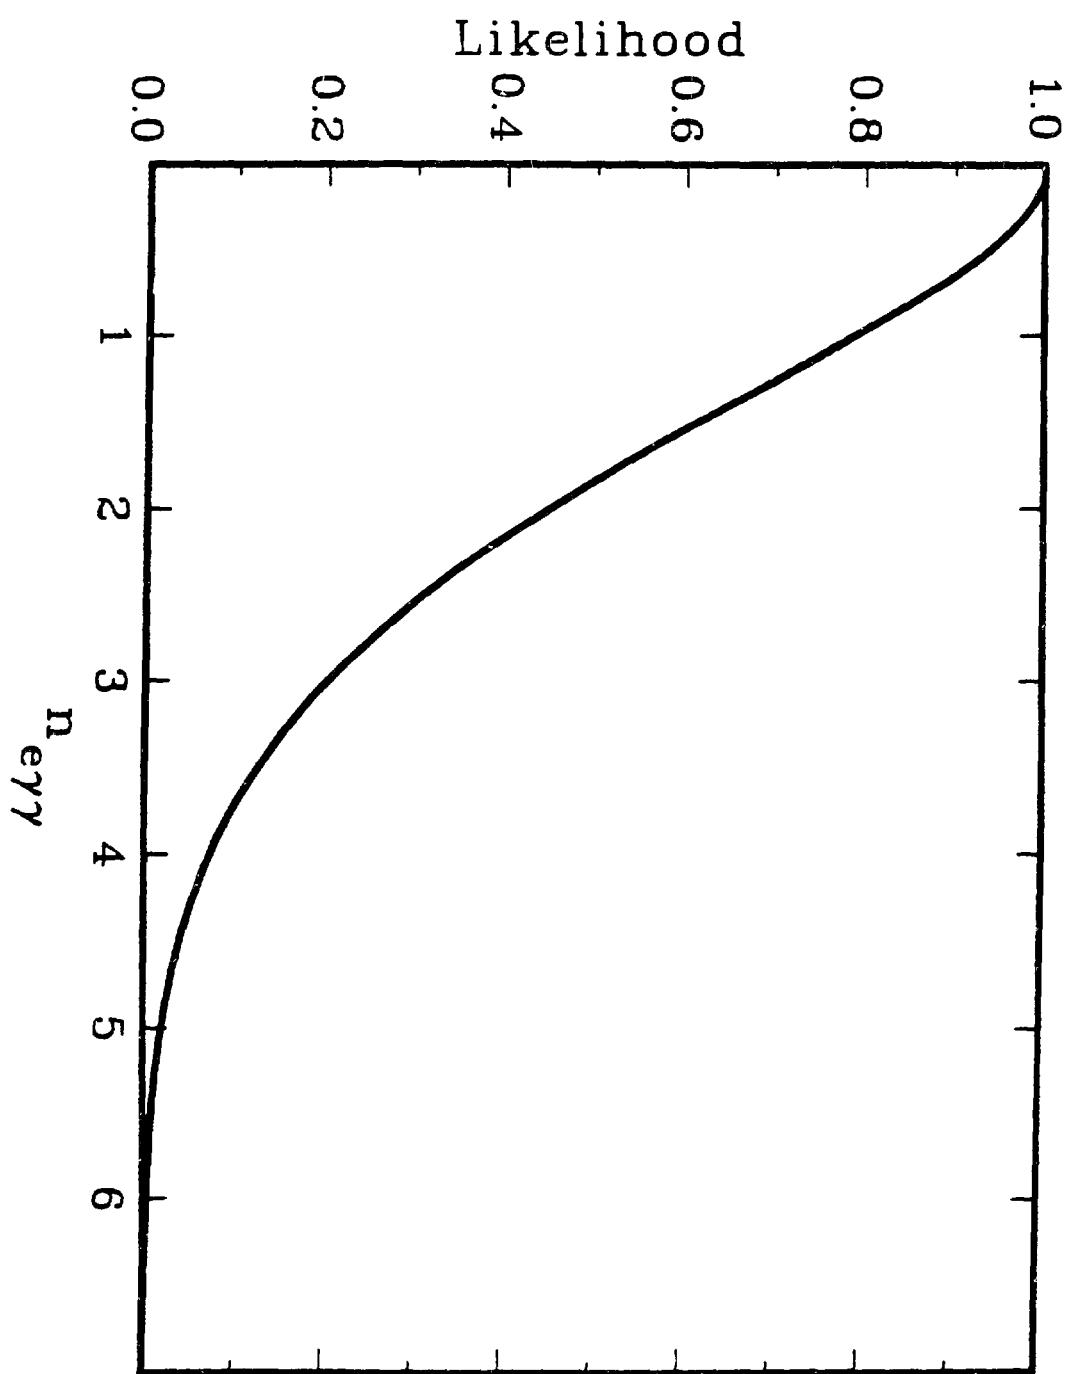


Figure 25. The normalized likelihood function versus the number of $\mu^+ \rightarrow e^+ \gamma\gamma$ events, $n_{e\gamma\gamma}$, for the nine candidate data events.



CHAPTER VI

RESULTS AND CONCLUSIONS

The 90% confidence level upper limit for the $\mu^+ \rightarrow e^+ \gamma\gamma$ branching ratio is calculated from the relation,

$$B_{e\gamma\gamma} \leq \frac{n_{90}}{N \frac{\Omega}{4\pi} \epsilon_C \epsilon_T \epsilon_D} , \quad (VI-1)$$

where n_{90} is the number of events at the 90% confidence level, N is the total number of muons stopped during the live time of the detector, $\Omega/4\pi$ is the solid angle of the detector for $\mu \rightarrow e\gamma\gamma$ events subtended from the target, ϵ_C is the efficiency of the data cuts, ϵ_T is the efficiency of the trigger requirements, and ϵ_D is the efficiency associated with the detector.

The total number of muons is calculated for a given run by taking the average of the rates in the four plastic scintillators located in the center of each quadrant and then correcting this value for the acceptance of these counters and the live time of the detector. The live time is defined as the time when the detector is actually available for collecting data and is 86.0%, averaged over the entire data-taking period. The number of muons stopped during the January period was 2.25×10^{11} , the June period contained 1.35×10^{11} stopped muons, and the summer period had 8.15×10^{11} stopped muons, all corrected for the detector live time. The

total number of muons stopped during the live time of the detector was $N = 1.18 \times 10^{12}$.

The product, $\Omega/4\pi \cdot \epsilon_C \cdot \epsilon_T$, is found from the Monte Carlo program using the method described previously in Chapter IV. This product is calculated for the conditions in each of the January, June, and summer data samples, and the average value is 6.4%. The product for three-quadrant $\epsilon_{\gamma\gamma}$ triggers is 5.1% and 1.3% for two-quadrant $\epsilon_{\gamma\gamma}$ triggers, averaged over the entire data sample.

Several detector efficiencies are included in the value of ϵ_D . Most of these are dependent on the instantaneous muon stopping rate. The quantity ϵ_D is also different for two-quadrant and three-quadrant $\epsilon_{\gamma\gamma}$ trigger types, and is calculated for each data run. The total detector efficiency is given by

$$\epsilon_D^{2Q} = \epsilon_\gamma \cdot (\epsilon_{\gamma p})^2 \cdot \epsilon_{ep} \cdot \epsilon_{DC} \cdot \epsilon_m \quad (VI-2)$$

for two-quadrant $\epsilon_{\gamma\gamma}$ triggers and by

$$\epsilon_D^{3Q} = (\epsilon_\gamma)^2 \cdot (\epsilon_{\gamma p})^2 \cdot \epsilon_{ep} \cdot \epsilon_{DC} \cdot \epsilon_m \quad (VI-3)$$

for three-quadrant $\epsilon_{\gamma\gamma}$ triggers. The first term, ϵ_γ , is the efficiency of identifying a photon quadrant with no extraneous scintillator signal that would veto a valid $\epsilon_{\gamma\gamma}$ trigger. This quantity depends on the instantaneous beam rate. The value of ϵ_γ is obtained from the number of photon quadrants without scintillator signals when randomly triggering the Crystal Box during the beam gate. An approximate value for ϵ_γ is 80.0% at the weighted-average instantaneous muon stopping rate of 5.5 MHz, taken

over the entire data-taking period. The square of ϵ_γ appears in Equation VI-3 because two different photon quadrants are required for three-quadrant $\gamma\gamma\gamma$ triggers. The next two quantities, $\epsilon_{\gamma p}$ and ϵ_{ep} , are the efficiencies for discarding valid events with photon and positron piled-up energy over 2.5 MeV, respectively. The efficiency is different for the January and June data samples and also for the summer data sample because of the different pile-up rejection methods. The efficiencies are determined by finding the percentage of photons or positrons with piled-up energy greater than 2.5 MeV. The value for $\epsilon_{\gamma p}$ at 5.5 MHz is 95% and 90% for ϵ_{ep} at the same rate. The quantity ϵ_{DC} is the rate-dependent efficiency of the drift chamber sense wires detecting a particle trajectory and the track-reconstruction program locating it. The value for ϵ_{DC} is found from the percentage of single positron tracks reconstructed in the drift chamber for several beam rates. It has a value of 95.9% at the 5.5 MHz rate. The final quantity, ϵ_m , takes into account several miscellaneous efficiencies, such as the probability the I-counter vetoes a valid event originating from the target and the algorithm for finding positrons annihilating in flight vetoing a good event. The efficiency ϵ_m has a rate-independent value of 98%. The average detector efficiency for two-quadrant triggers, ϵ_D^{2Q} , weighted by the number of muons per data sample, is 61.8%, and the efficiency for three-quadrant triggers, ϵ_D^{3Q} , is 49.6%.

The denominator of Equation VI-1 is 4.02×10^{10} . Combining this value with $n_{90} = 2.9$, the upper limit for the $\mu^+ \rightarrow e^+ \gamma\gamma$ branching ratio is

$$\mathcal{B}_{\gamma\gamma\gamma} \leq 7.2 \times 10^{-11} \text{ (90% C.L.)} . \quad (\text{VI-4})$$

This result is an improvement of more than two orders of magnitude over the previous limit for this muon decay mode.

The Crystal Box experiment searching for $\mu \rightarrow e\gamma\gamma$ is the first to use a large solid-angle detector that can identify three separate particles. Previous searches^{22,24-26} used two arm spectrometers in which both photons were detected in a single arm, and they observed only a small fraction of the possible decay topologies. This particular topology is eliminated in the data analysis for the present search for $\mu \rightarrow e\gamma\gamma$. In the Crystal Box analysis, the positron energy is greater than ~ 8 MeV and each photon energy is greater than 20 MeV. In the previous searches, the particle energy thresholds were about 30 MeV. With the lower energy thresholds more area of the Dalitz plot for $\mu \rightarrow e\gamma\gamma$ can be explored than in previous experiments.

The new limit for the $\mu \rightarrow e\gamma\gamma$ branching ratio can place bounds on some theoretical models involving the composite structure of quarks and leptons. A bound on the reduced mass of the constituent particles in the composite model,⁷⁸ where muons are radially-excited states of the electron, can be calculated using Equation II-13 and the upper limits on the rates for $\mu \rightarrow e\gamma$ ²⁰ and $\mu \rightarrow e\gamma\gamma$. The lower bound on the mass of this constituent particle is 6.9×10^7 GeV. Another composite model⁸⁰ has an intermediate excited lepton state mediating the $\mu \rightarrow e\gamma\gamma$ decay. A lower limit on the mass of this intermediate particle is set at 1.2 TeV, by scaling the previous mass limit with the new $B_{e\gamma\gamma}$ value.

In retrospect, one may ask how this experiment could have been improved so that an even more sensitive search could be performed. This experiment was on the verge of being background-limited, so it becomes

necessary to understand what background processes still existed at the final stages of the analysis. A plot similar to Figure 19 of the time differences for the 364 events reveals that most of these events are contained in the horizontal and vertical bands. This implies coincidences between inner bremsstrahlung decays and random photons. An improvement in the detector resolutions of the energy, time, and position measurements would help differentiate these background events from $\mu \rightarrow e\gamma\gamma$ events. The background from random coincidences would be reduced by using a direct-current, instead of pulsed, beam.

In summary, the rare neutrinoless muon decay, $\mu^+ \rightarrow e^+\gamma\gamma$, is not observed in this search using the Crystal Box detector. A new upper limit for the branching ratio is set at $B_{e\gamma\gamma} \leq 7.2 \times 10^{-11}$ (90% C.L.), a more than two orders of magnitude improvement over the previous limit. No evidence for the nonconservation of lepton-family number is as yet observed in muon decays and there is still no evidence for new physics beyond the standard model.

APPENDIX

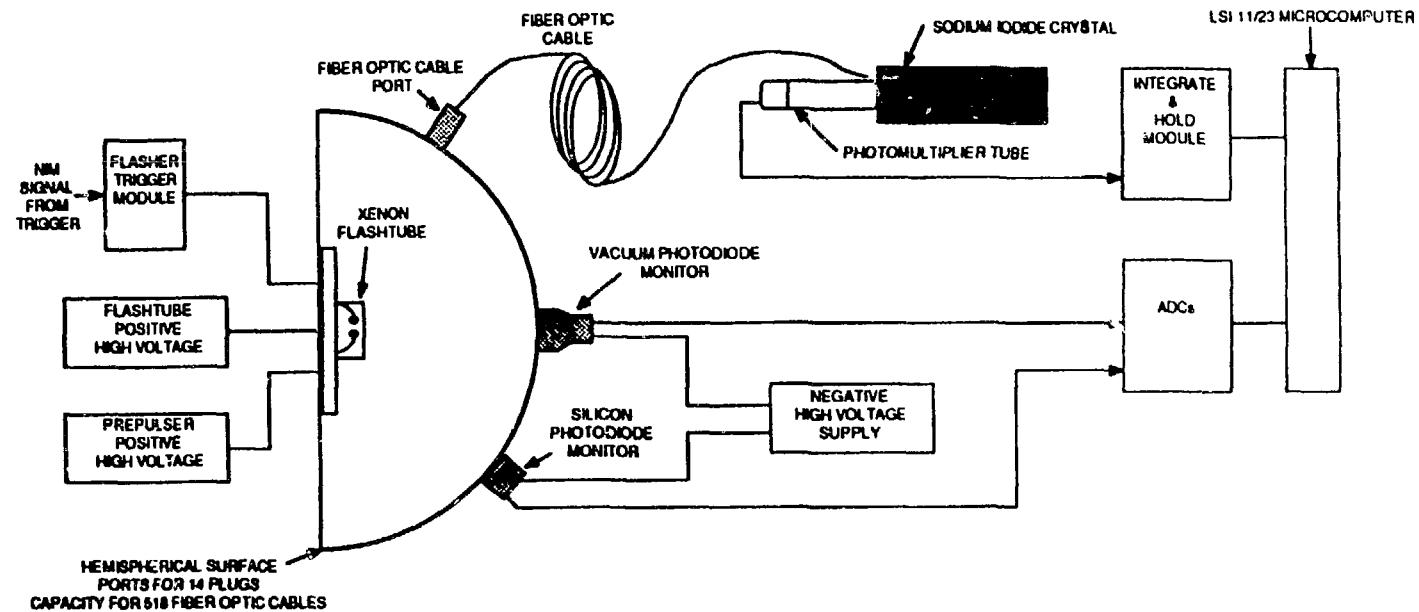
THE FLASHER SYSTEM

A reliable system for monitoring changes in the energy and time measurements is a necessity for any detector using calorimetry. Such a system can be based on two stabilization-monitoring techniques. One measures a known physical process, such as particle masses or energies, within the detector. The other uses sources of light or charge produced externally from the detector. The first technique is more desirable because actual events are measured in the detector, but it may not always be practical. The second technique may use radioactive sources, light-emitting diodes, or externally-produced light sources, such as spark gaps or lasers, coupled to the detector. This technique has the advantages of being flexible and producing large amounts of light, but the disadvantage of being mechanically difficult to implement.

The requirements for a monitoring system are reliability, flexibility, speed, and stability. The system must be able to function well for long periods of time, be easy to use, and measure both energy and time variations. It should also operate reasonably fast as not to erode data-taking time. Finally the monitoring system itself has to be stable in order to spot variations in the detector as a function of time.

The energy and time stabilization monitoring system for the Crystal Box detector, shown in Figure 26, consists of a xenon light source, two photodiodes to monitor variations in the light source output, and a

Figure 26. Schematic diagram of the flasher system.



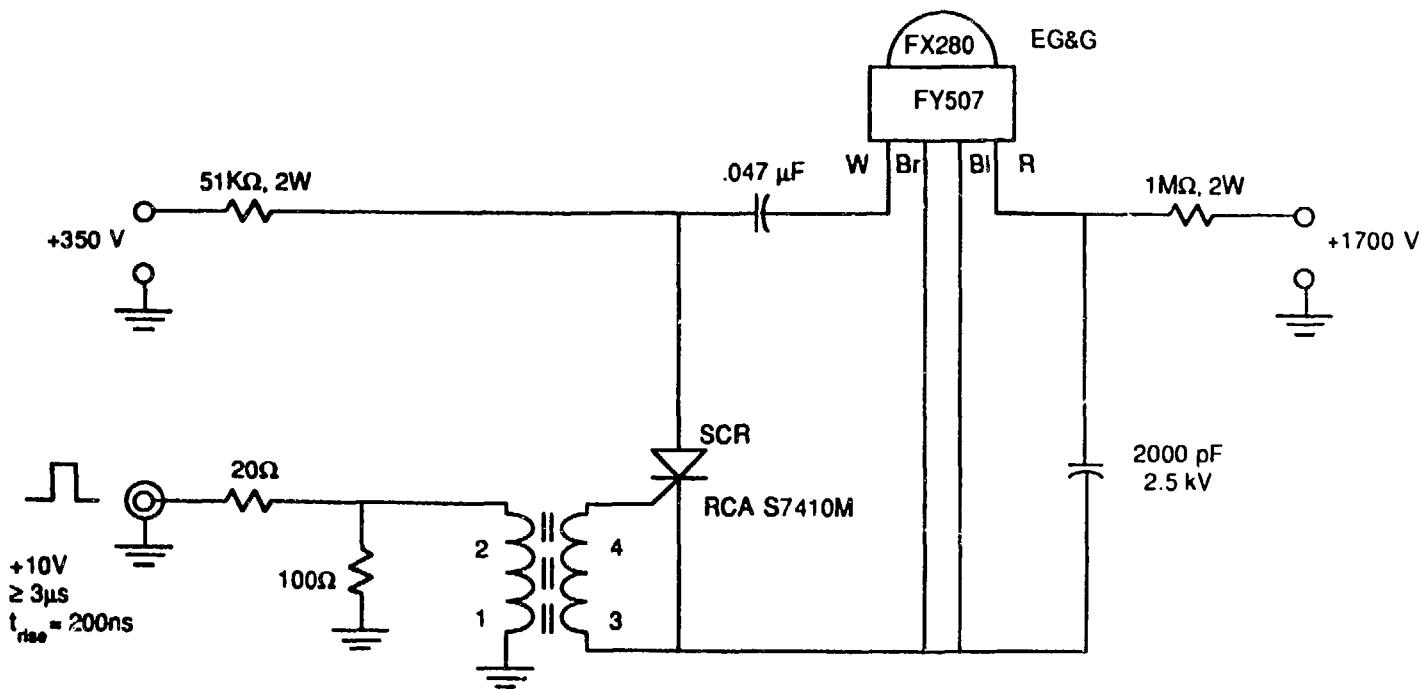
network of fiber optic cables to transport the light to each NaI(Tl) crystal. This light flasher system is also used to check the performance of each NaI(Tl) electronics channel.

A. Apparatus and Method

The light source for the flasher system is a xenon flashtube that is able to generate short durations of light at high radiance. The flashtube is a bulb-type EG&G model FX-279¹⁰⁶ with the capability of producing an output energy of several joules. Only several millijoules are used for this application. At this energy output, the lifetime of the bulb exceeds 10^9 flashes. The distance between electrodes producing the discharge arc is 1.5 mm. The spectrum of the flashtube light has wavelengths ranging from 300 nm to 1100 nm and is peaked around 430 nm. Time jitter is rated at less than 200 ns and the light output variations are less than 10%.

The flashtube driver circuit is shown in Figure 27. A NIM-level signal from the experimental electronics is converted into a special pulse that triggers the flashtube. The special pulse is a positive square wave, 10 V high, 3 μ s wide, and with a risetime of 200 ns. This input pulse triggers a prepulse of +350 V through a high voltage transformer and this initiates a high voltage arc across the electrodes in the flashtube driven by a potential difference of +1700 V. The light output is shaped with a long tail in the driver circuit to imitate a NaI(Tl) pulse. The voltages of the flashtube are set at these values to minimize the time jitter of the flash pulse during the NaI(Tl) signal charge integration period.

Figure 27. The flashtube driver circuit.



Two devices, a vacuum photodiode and a silicon photodiode, monitor the flashtube light output. The vacuum photodiode is an ITT biplanar-type phototube model FW114.¹⁰⁷ It has a Cs-Sb photocathode with an S4 spectral response. The measured peak radiant sensitivity is 35 mA/W at the spectral peak of 400 nm. The photocathode has an active area of 2.4 in.². The anode, a Ni mesh with an 80% transmission, is 0.25 in. distant from the cathode. The risetime of the output pulse from the vacuum photodiode is less than 500 ps.

Figure 28 shows the operating circuit for the vacuum photodiode. The circuit contains isolated signal and high voltage ground connections to reduce the effects of noise and ground loops. Figure 29a shows a plateau curve of the high voltage applied to the vacuum photodiode. The normal operating voltage is -2.5 kV and is set at this voltage to ensure a fast output pulse risetime. The fast risetime is necessary to provide the Start pulse for the timing measurements made with the flasher system.

The silicon photodiode is a solid-state PIN device that serves as a second monitor of the flashtube light output and a check of the vacuum photodiode performance. The silicon photodiode is an EG&G SGD-444 PIN photodiode¹⁰⁶ with an active area of 1.0 cm². It has a spectral range of 350 nm to 1150 nm, with the peak sensitivity located in the "red" portion of the spectrum. The risetime is rated at 10 ns. Figure 30 shows the operating circuit for the silicon photodiode. The ground connections are again isolated from one another. The normal operating voltage is -180 V, as determined by the voltage curve in Figure 29b.

The output signals from both of the photodiodes do not go through amplifiers, but proceed directly into the NaI(Tl) integrate-and-hold (I/H) modules. No amplification of these pulses is necessary and the

Figure 28. Operating circuit for the vacuum photodiode.

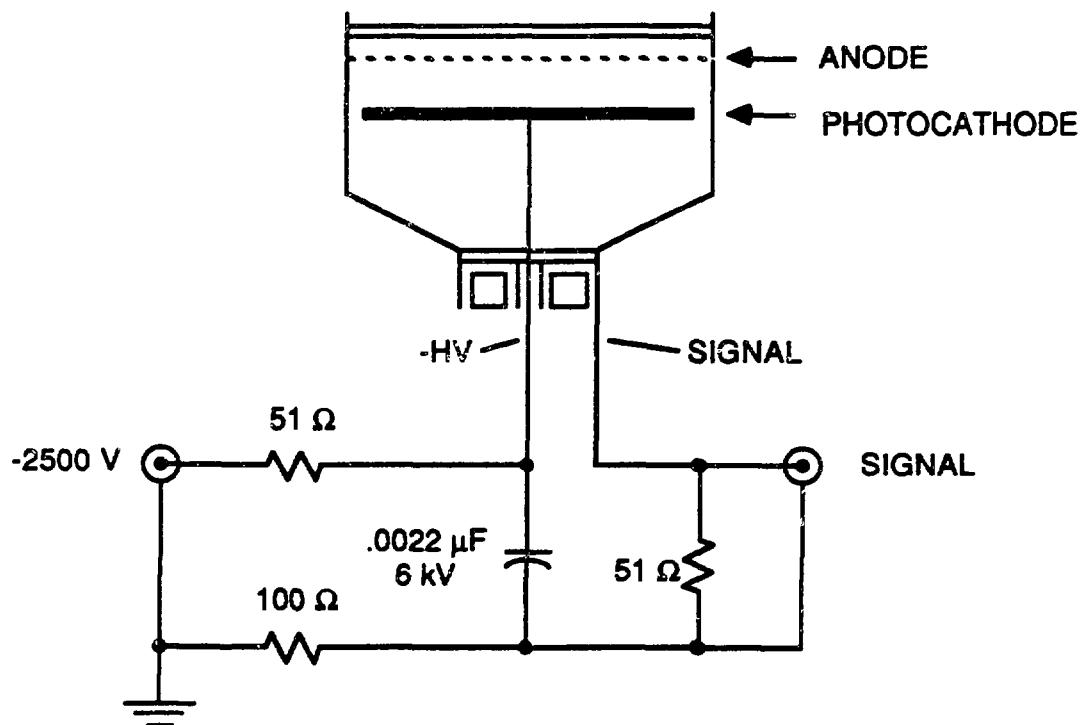


Figure 29. Photodiode monitor high voltage plateau curves.
Plots of the ADC response as a function of voltage for (a) the vacuum
photodiode and (b) the silicon photodiode.

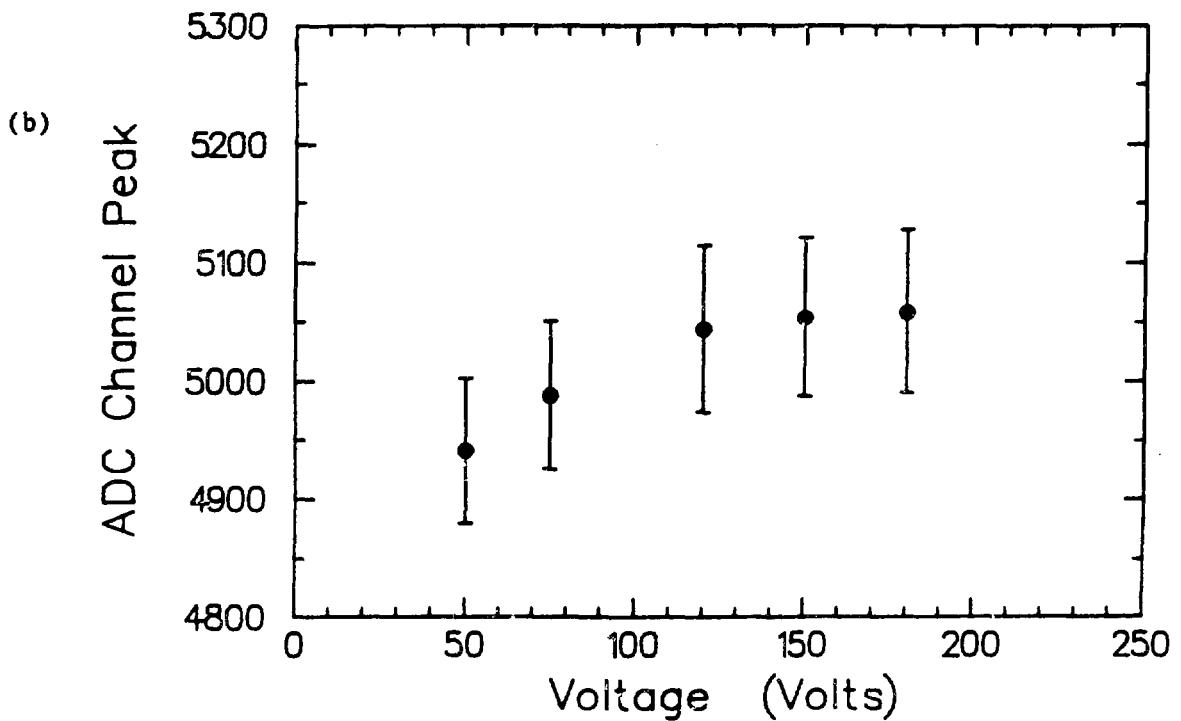
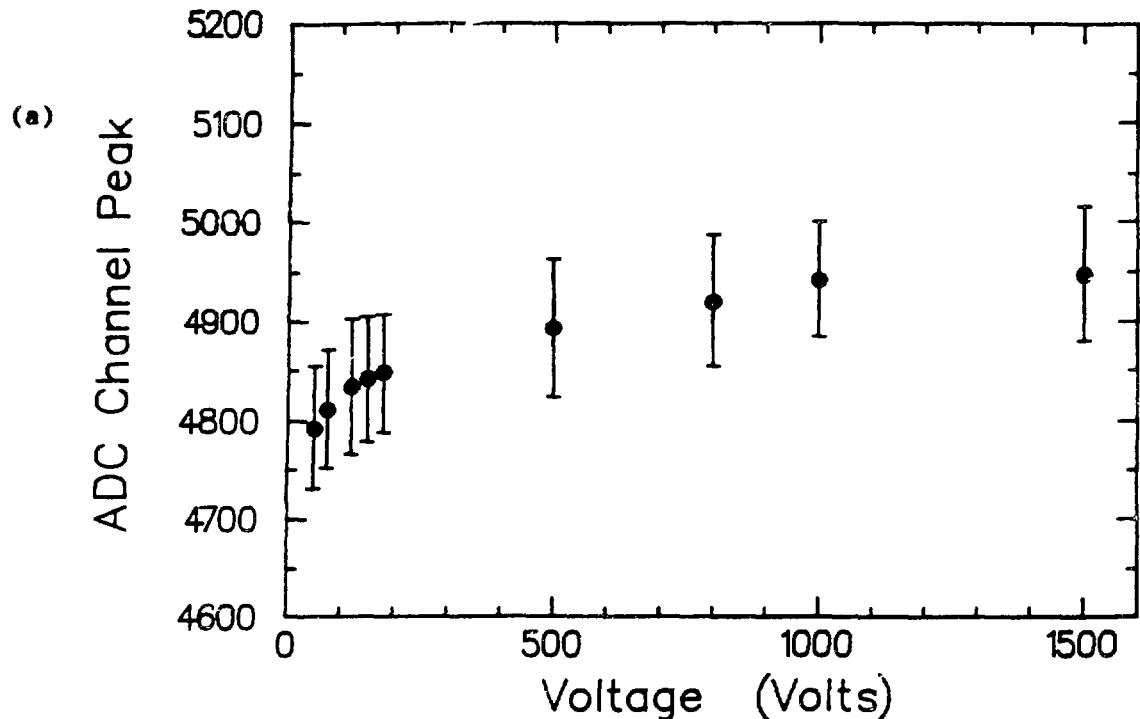
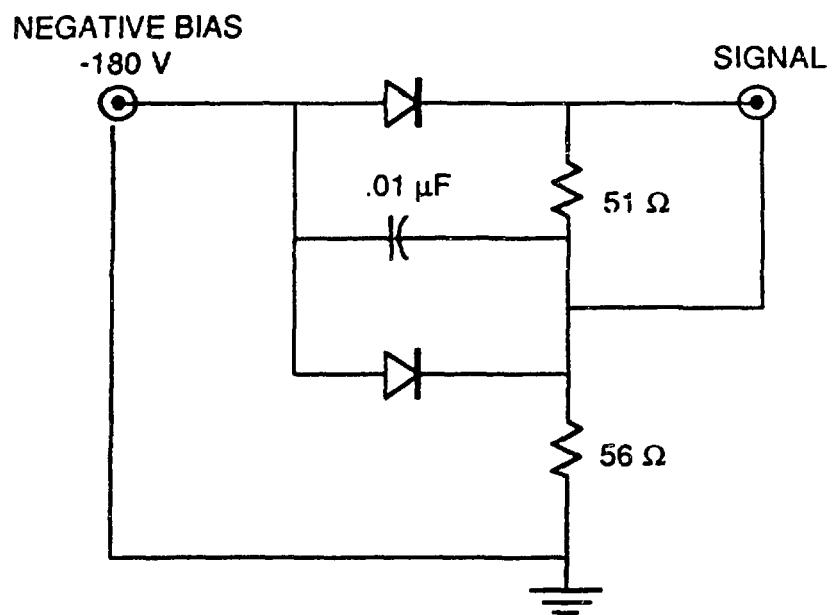


Figure 30. Operating circuit for the silicon photodiode.



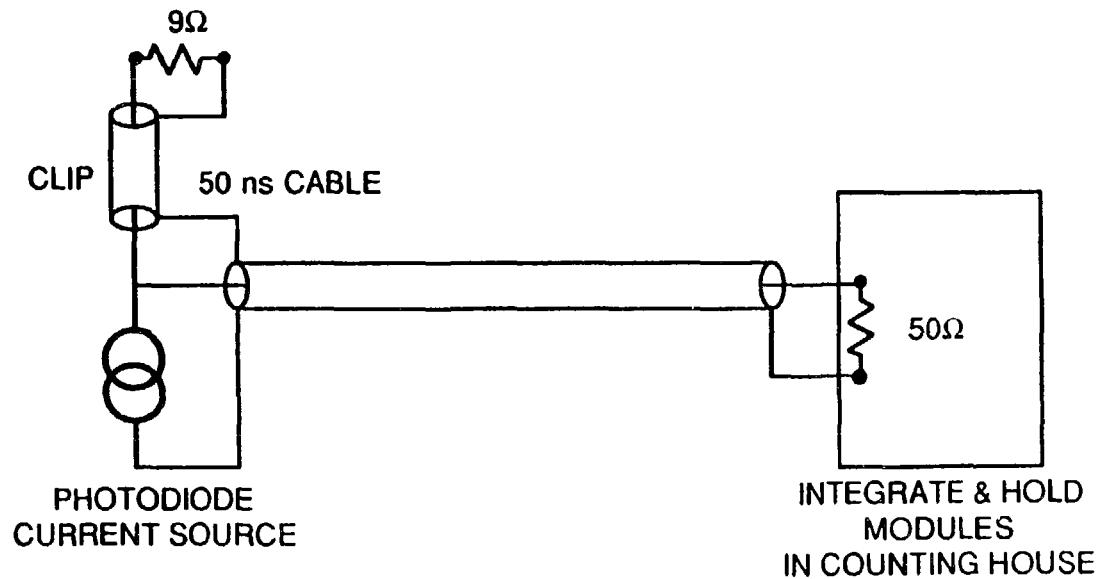
pulse-shaping is made through cable clips located at each photodiode output, as shown in Figure 31. The signals are treated exactly like the NaI(Tl) crystal pulses after the I/H modules.

The light from the flashtube is conveyed to each NaI(Tl) phototube through a single, plastic optical fiber. The fiber optic cable is Dupont CROFON 1040¹⁰⁸ with a fiber diameter of 0.040 in. and sheathed in a polyethylene jacket. The angle of acceptance for this optical fiber is 64.2° with the light output direction peaked forward at 20°. Light loss in the fiber is cited as 7-8% per foot at the green spectral region and is mainly due to absorption.

One end of the optical fibers are situated on a spherical surface such that each cable directly views the flashtube light source. Two lengths of fiber optic cable carry the light from the source to the phototube. One length is 12 ft., from the spherical surface to the phototube base, and the other length is 14 in., from the phototube base to a slot cut at 30° in the phototube light guide. The segmentation of the fiber optic cable is designed for the easy removal of the phototube, but the disadvantage is a significant loss of light through this connection. All of the connectors are manufactured commercially.¹⁰⁹ The connection at the light guide has the optical fiber glued into the slot pointing toward the NaI(Tl) crystal. The physical arrangement of the phototubes necessitate this design.

The flasher system is operated approximately every two hours to monitor both energy and time measurements between data-taking runs. It follows an update of the pedestal values in the NaI(Tl) crystals. The flashtube trigger rate is 60 Hz and "flashes" for a minute to accumulate sufficient statistics within histograms for the stabilization

Figure 31. Pulse-shaping circuit for the photodiode monitors.



measurements. The data from the flasher system are collected by the four LSI-11/23 microcomputers and written to computer tape for off-line analysis. All NaI(Tl) electronics channels are checked on-line for proper operation by the flasher system.

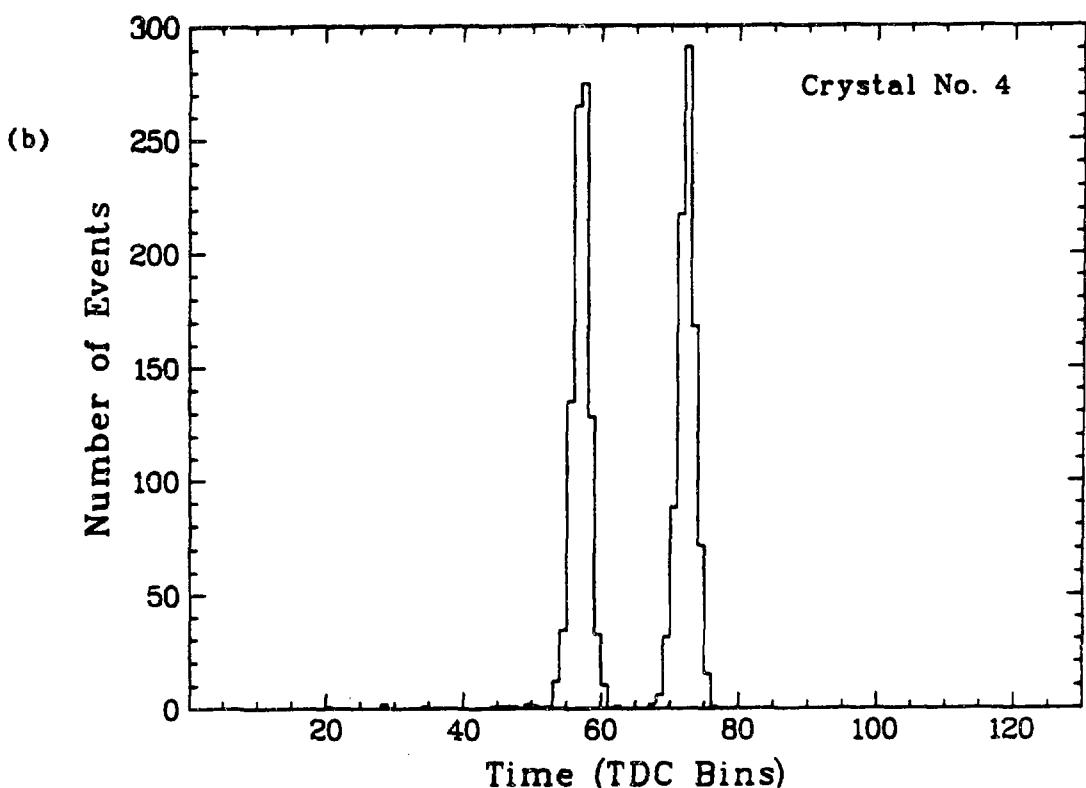
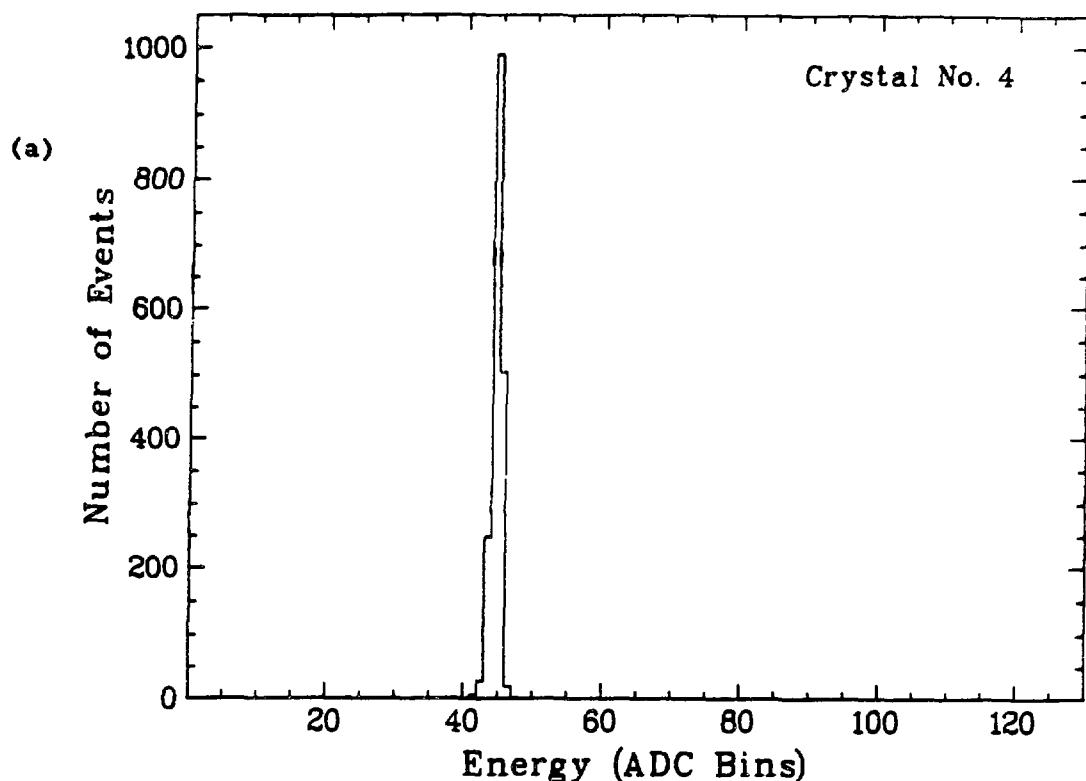
The LSIs store histograms containing flasher data for each NaI(Tl) ADC and TDC channel. Each energy ADC histogram for the flasher consists of 128 bins with 32 channels per bin. Each TDC histogram also has 128 bins but has 4 TDC channels per bin. The number of bins per histogram is limited by the storage capacity within the LSI. The LSI histogram can be moved within an absolute scale by obtaining the location of the first bin with respect to this scale. These threshold values are also stored in the LSIs. The number of channels per bin is chosen to accommodate data shifts within the histogram. Figure 32 shows the raw ADC and TDC flasher data for a given NaI(Tl) crystal. A peak position is fit from the data. The peak positions from the monitors are found in exactly the same manner. Reference values of the peak positions are chosen from the flasher data taken during the calibration of the detector. All subsequent flasher data values are compared to these reference values.

The relative energy gain correction g_i for a crystal i is computed from

$$g_i = \frac{1}{2} \left[\frac{\frac{PP_i}{PPVP}}{\left(\frac{PP_i}{PPVP} \right)_{ref}} + \frac{\frac{PP_i}{PPSP}}{\left(\frac{PP_i}{PPSP} \right)_{ref}} \right], \quad (A-1)$$

where PP_i is the corrected peak position, $PPVP$ is the corrected peak

Figure 32. Raw data spectra of a given crystal using the flasher system. (a) Raw energy spectrum of crystal number 4. The histogram contains 128 bins with 32 ADC channels per bin. (b) Raw double-peaked timing spectrum of the same crystal. The histogram has 128 bins with 4 TDC channels per bin.



position of the vacuum photodiode monitor, PPSP is the corrected peak position of the silicon photodiode monitor, and "ref" indicates the reference peak position ratio. In this manner, the variations in the light output from the flashtube are removed from the energy gain corrections. The peak position ratios are averaged over both monitors in case of drifts in either monitor output. These gain corrections modify the energy values for a given set of data.

The timing peak positions are found from double-peaked spectra, generated from the 10 ns alternate path to the trigger. The position of the first peak gives the time offset for a given NaI(Tl) channel. This peak position corrects the time value for each channel and data run. Timing gains can be computed with the flasher system, but are not used for a timing calibration.

B. Results

The distribution of energy peak positions from the flasher system is shown in Figure 33. The amount of light measured by each NaI(Tl) crystal is different; no attempt is made to make the input light uniform. The average amount of light corresponds roughly to 20 MeV. Light loss occurs primarily at the junction of the two fiber optic cables and where the optical fiber views the light source.

The photodiode monitors are an essential part of the energy gain corrections. Any variations in these devices affect the gain value, so it is important that the monitors be stable. Figure 34 shows the ratio of the peak positions of the vacuum photodiode to the silicon photodiode as a

Figure 33. The distribution of energy peak positions for a given data run using the flasher system.

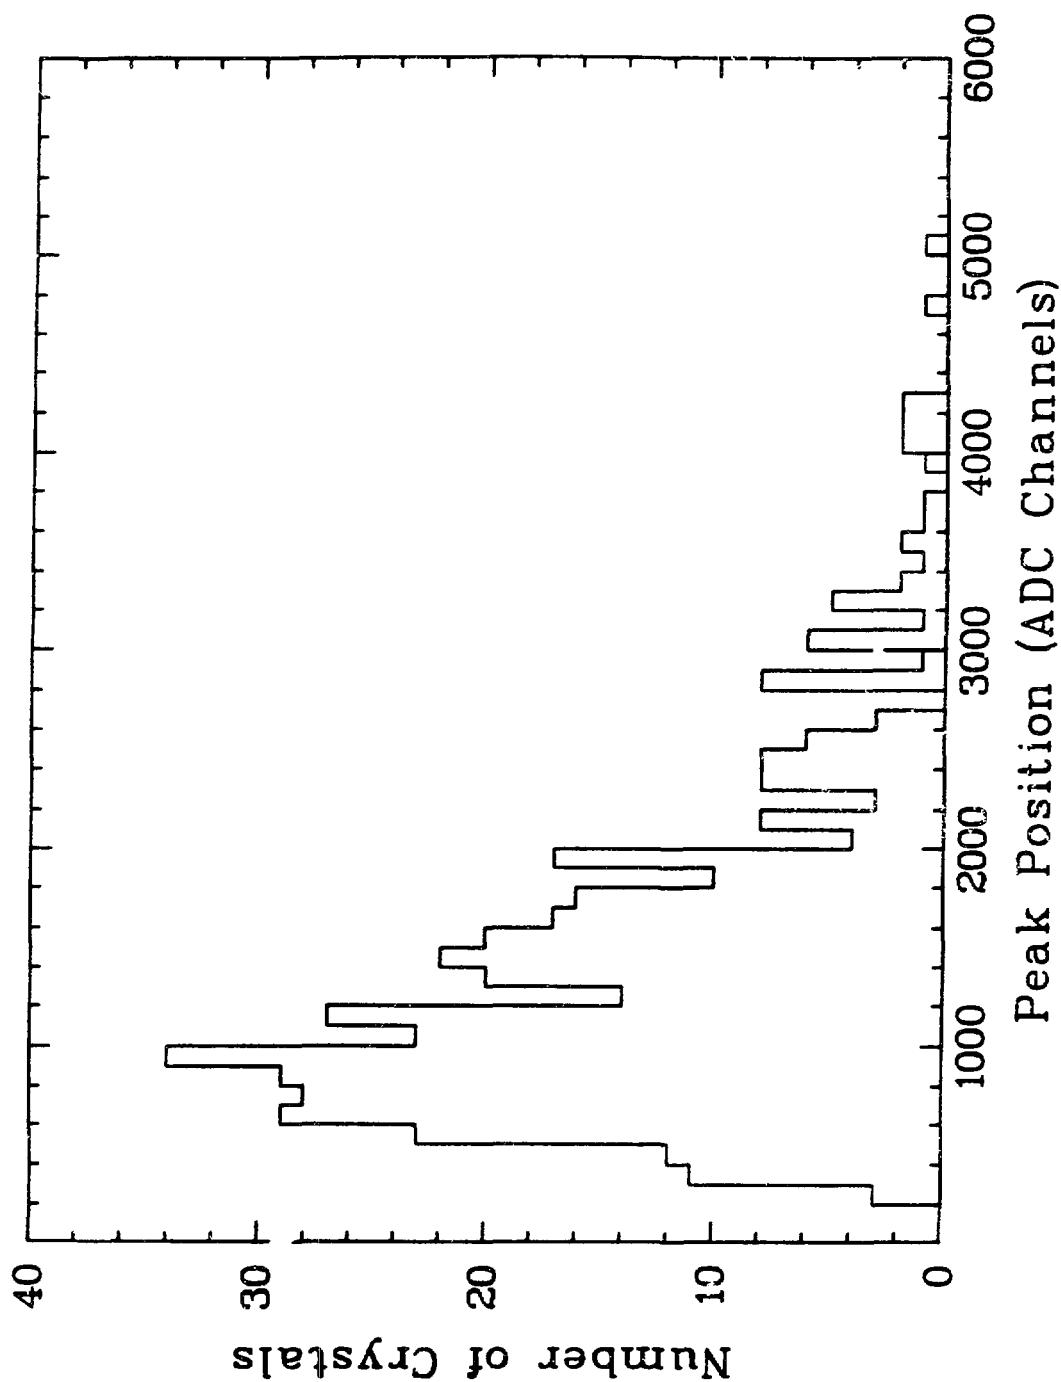
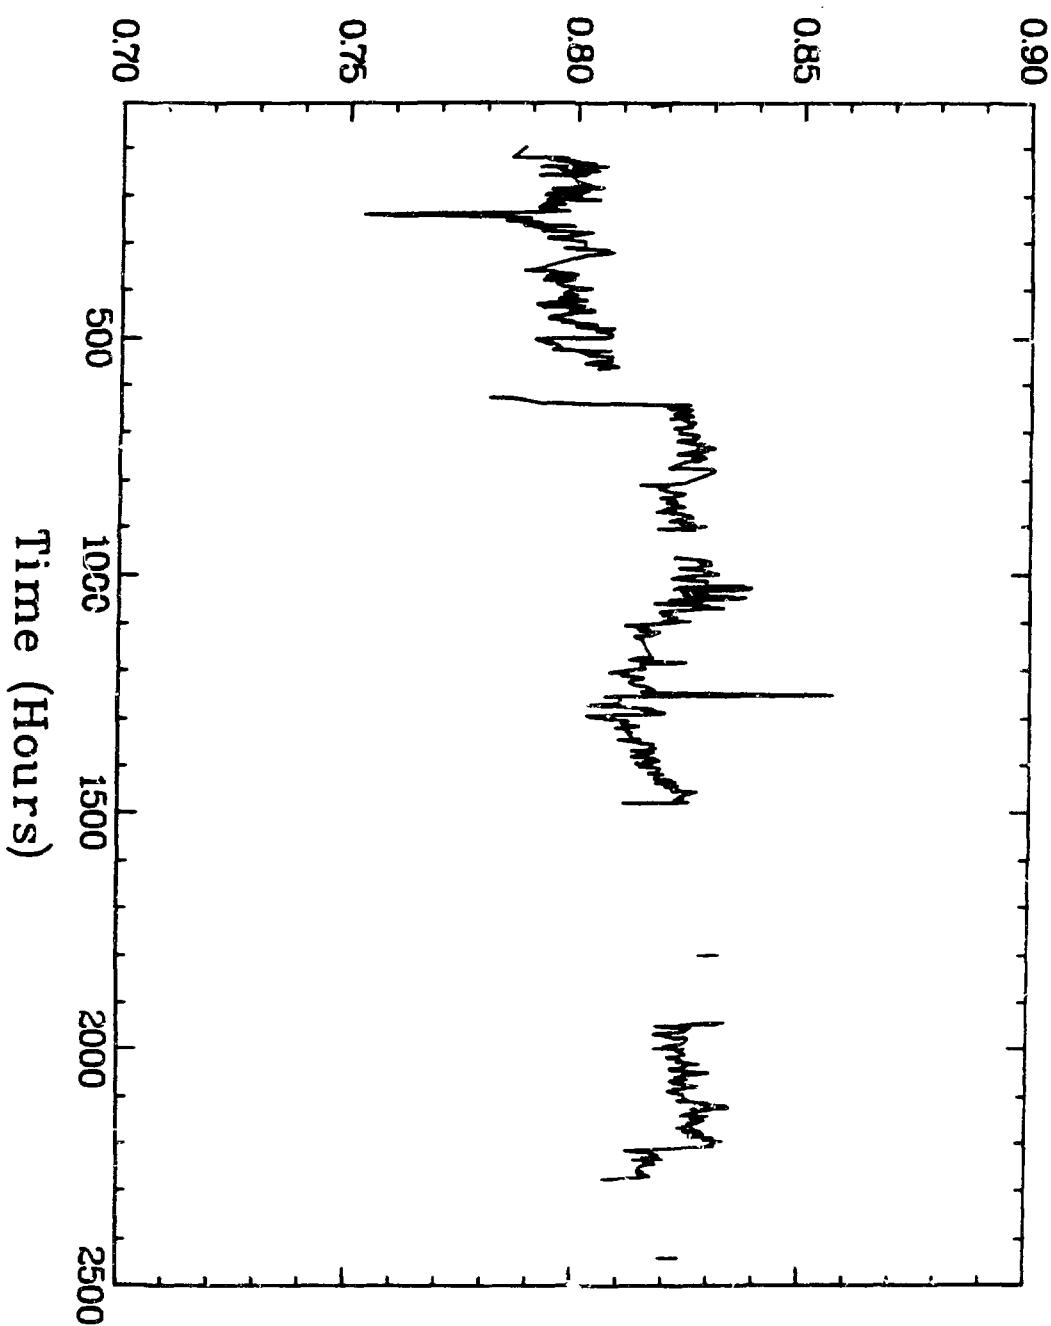


Figure 34. Plot of the ratio of the peak position from the vacuum photodiode to the peak position of the silicon photodiode as a function of time.

Ratio of Photodiode Peak Positions



function of time. If both devices are stable, the ratio should be a constant value. The discontinuity at hour 650 is due to a hardware change in the flasher system. Typical variations in this ratio are 1.6%.

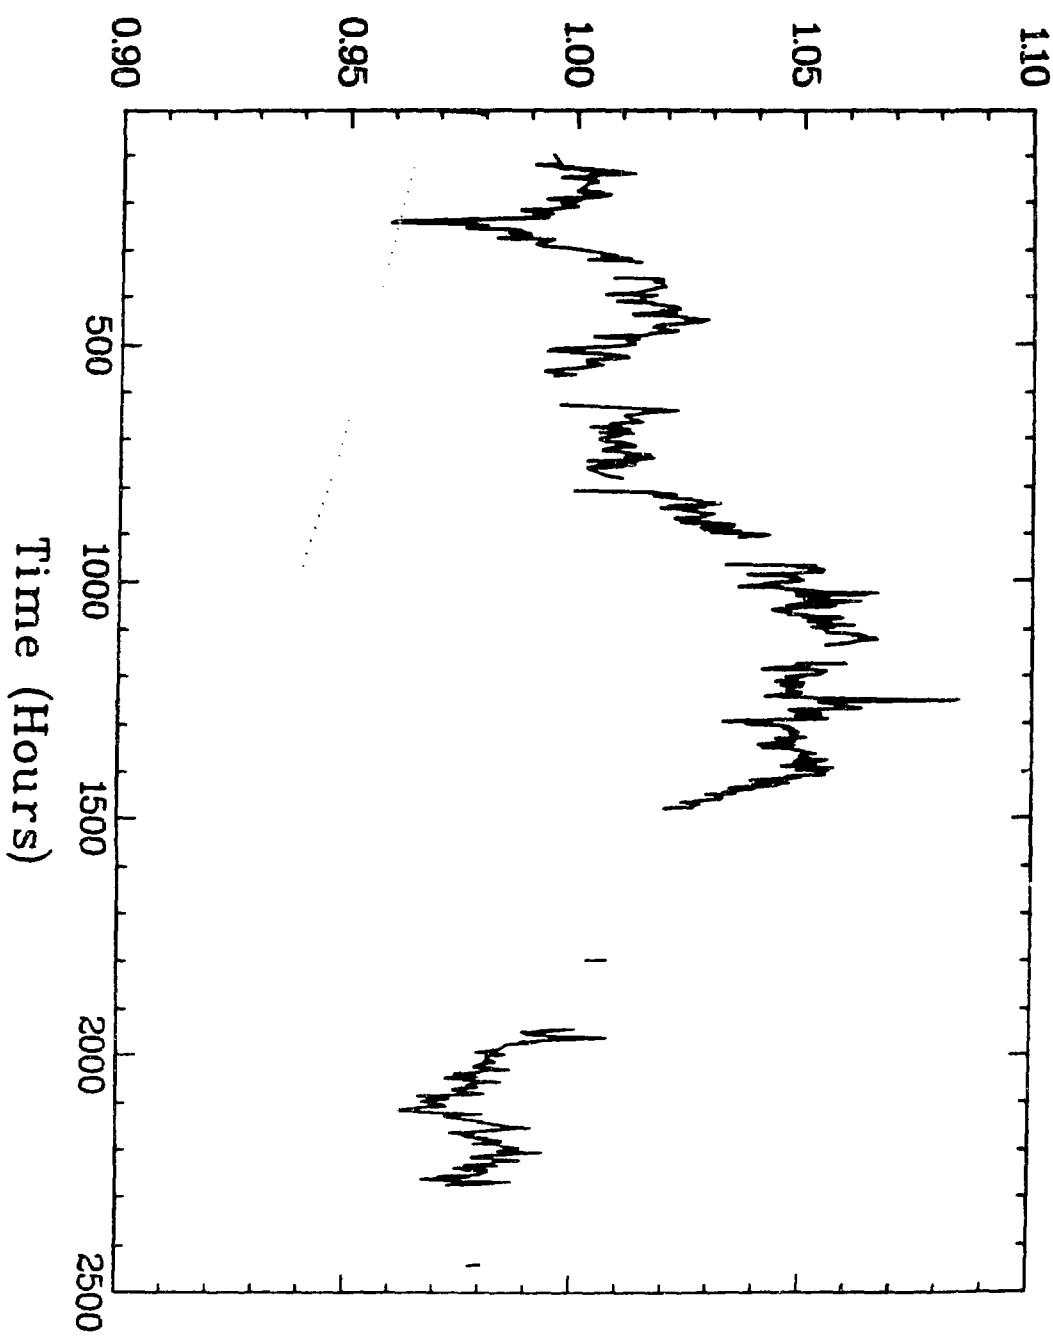
The energy gain corrections are calculated using Equation A-1 for each crystal and run. Figure 35 shows the average correction of the 432 NaI(Tl) ADC channels as a function of time. It is quite apparent that the flasher system detects changes in the average energy gains, some as high as 3% between data runs.

One of the key questions in monitoring and correcting the energy gains is whether or not the flasher system correctly tracks the changes in the NaI(Tl) gains. A direct comparison of the average energy gain of all the crystals is made with the flasher system versus the average energy gain measured by positron energy spectra. These calibrations, at a known energy value, show that the flasher gain changes by as much as 8% over the data. A quadrant average of the energy gains, determined from the positron energy spectra, indicates that the flasher quadrant averages change by several percent with respect to each other. The positron energy spectra from a single data run can determine the average energy gain of all the crystals to about 0.25%, a quadrant average to about 0.7%, and single crystal gains to about 1.5% for a face crystal located near the center of a quadrant. Determination of the individual crystal gains is limited by the statistics accumulated for each crystal.

The evidence indicates that the changes in the gains corrected by the flasher system are caused by drifts. The statistical spread of the quadrant averages at times nearest to the energy calibration are smaller than averages at times further away from the calibration. The quadrant

Figure 35. Plot of the average energy gain corrections from the 432 NaI(Tl) ADC channels using the flasher system as a function of time.

Ave. Gain Corrections



averages have systematic differences, such as the right quadrant average energy gain having a consistently lower value than the other quadrant averages. The systematic variations are within the quadrant gains and not within the fitting procedure, as evidenced by the agreement in the average energy gains between the even-numbered crystals and the odd-numbered crystals in a given quadrant. The individual crystal energy gains appear to be stable within the 1.5% measurement error per run. From these results, the flasher system can be used to correct individual crystal gains, but the average quadrant gains must be corrected using a different method in order to achieve the desired 0.5% experimental accuracy in the energy values.

The stability of the energy gains can be checked by comparing the flasher-determined gains versus the "true" gains as measured in an energy calibration. The energy calibration is the monochromatic 129.4 MeV photon spectrum from the reaction, $\pi^- p \rightarrow n \gamma$, and is measured in the centrally-located face crystals. The histogram in Figure 36 shows the energy spectrum with the energy gains corrected by the flasher system, 300 hours after the calibration. The data points are from the original energy calibration. The flasher system improves the individual crystal gains by a few percent, several hundred hours after an energy calibration.

Another indication of the flasher system performance in correcting the energy gains is shown in Figure 37. It contains the distribution of the percent change of the energy gains using the flasher corrections and the gains as determined by an energy calibration at that same time. This occurs at 300 hours after the original energy calibration. The peak is centered around zero with an uncertainty of 2.2%.

Figure 36. Photon energy spectrum using the flasher-corrected energy gains 300 hours after an energy calibration. The data points use the original energy calibration gains.

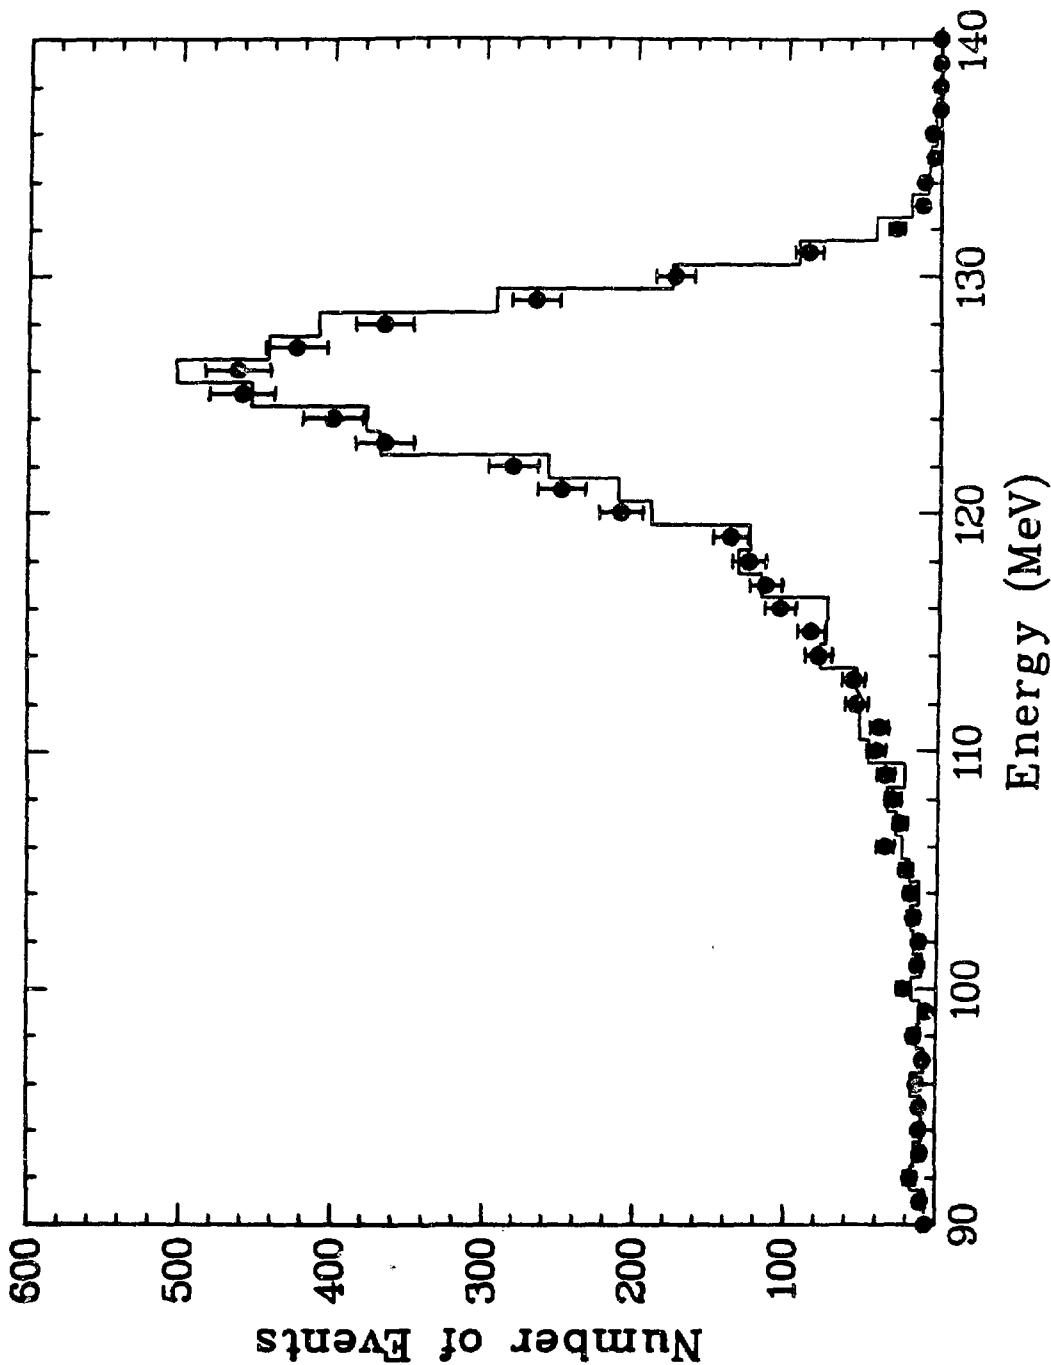
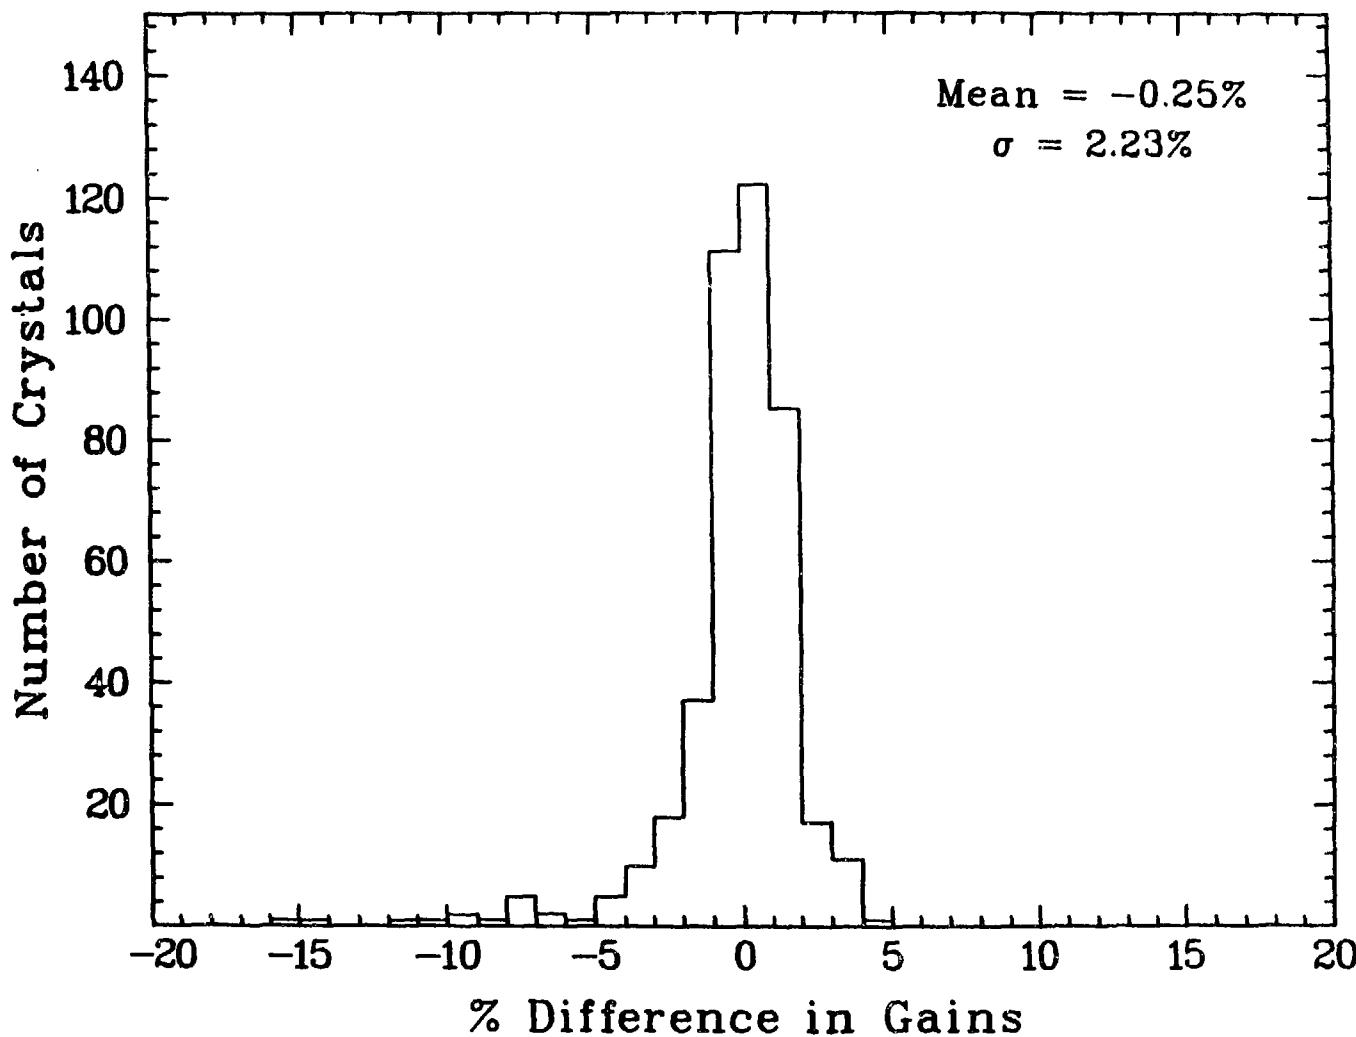


Figure 37. The distribution of the percent change of the energy gains using the flasher corrections and the gains determined by an energy calibration. This comparison occurs 300 hours after an initial energy calibration.



The second function of the flasher system is to monitor the time offsets on a run-by-run basis. The distribution of the first peak positions in the double-peaked spectra for all crystals is displayed in Figure 38 for a given data run.

Figure 39 shows the average timing peak position from all the crystals as a function of time during the data-taking period. Variations in this average value are less than 1%. The discontinuity in the distribution after 1300 hours is due to a hardware change in the flasher system. The time offsets, as determined by the peak positions, are corrected for each data run using the flasher system. The quadrant average of the time peak positions without flasher corrections vary by about 500 ps between runs during the course of a month. Some of these variations are as much as 1 ns. These time differences diminish to ± 250 ps using the flasher corrections. These corrections make a large improvement in the time measurement between calibrations, which occur approximately every three days using the method outlined in Chapter IV.

C. Conclusions

The flasher system is a useful tool in monitoring the performance of all NaI(Tl) electronics channels. Bad channels can easily be found and replaced. The correction of the NaI(Tl) timing offsets for each run using the flasher system improves the time measurement, keeping variations to within 250 ps. The flasher system is able to monitor variations in the individual crystal energy gains, but could not detect changes in the quadrant averages to better than several percent.

Figure 38. The distribution of timing peak positions for a given data run using the flasher system.

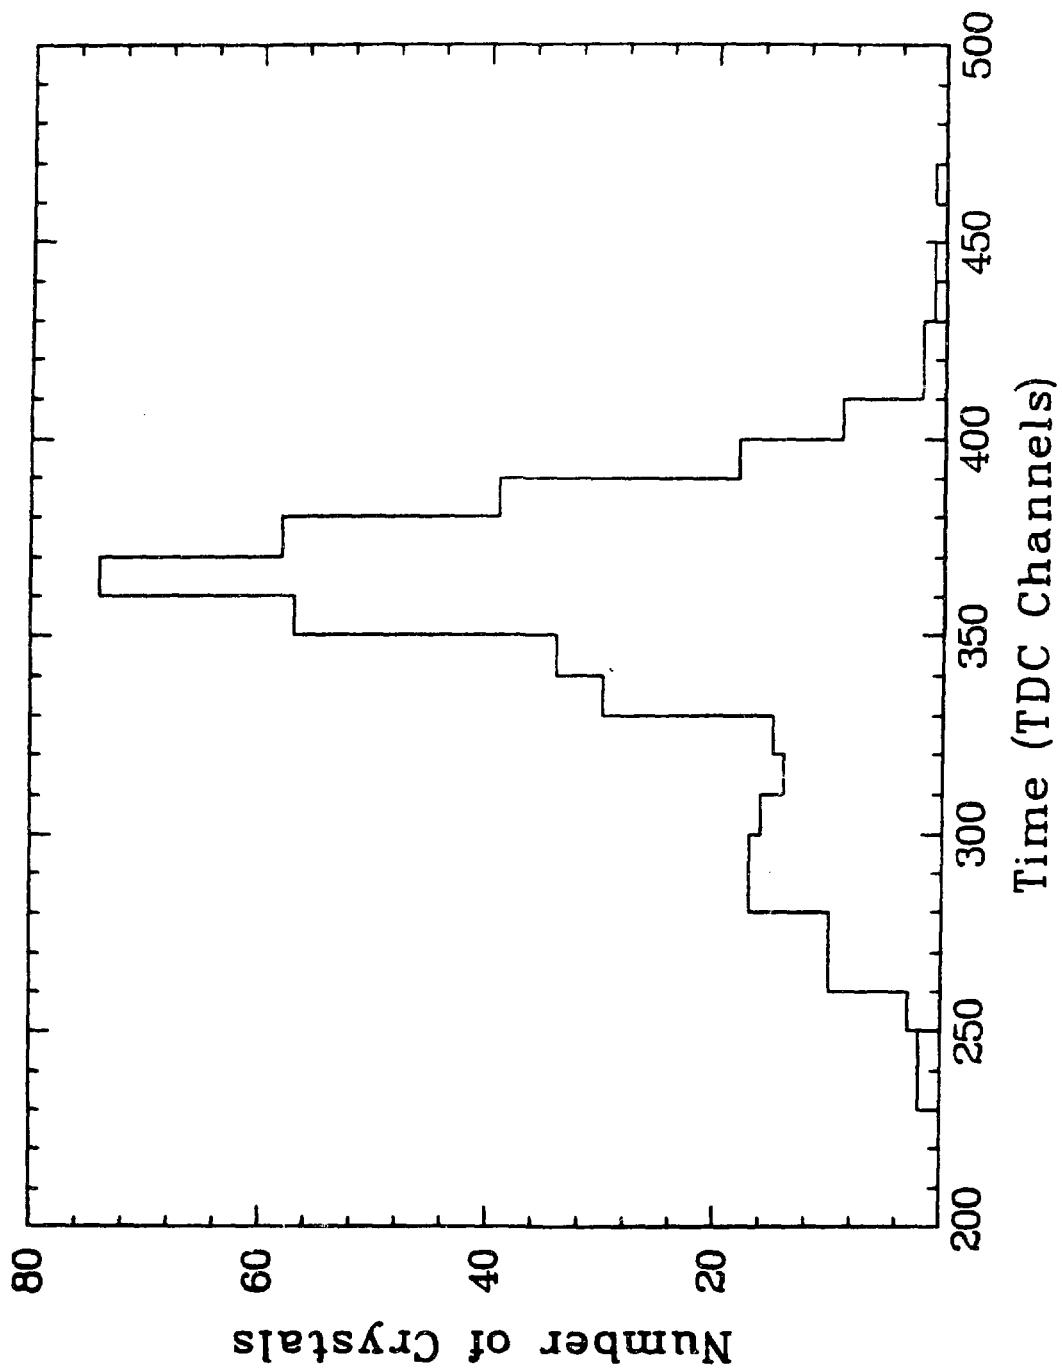
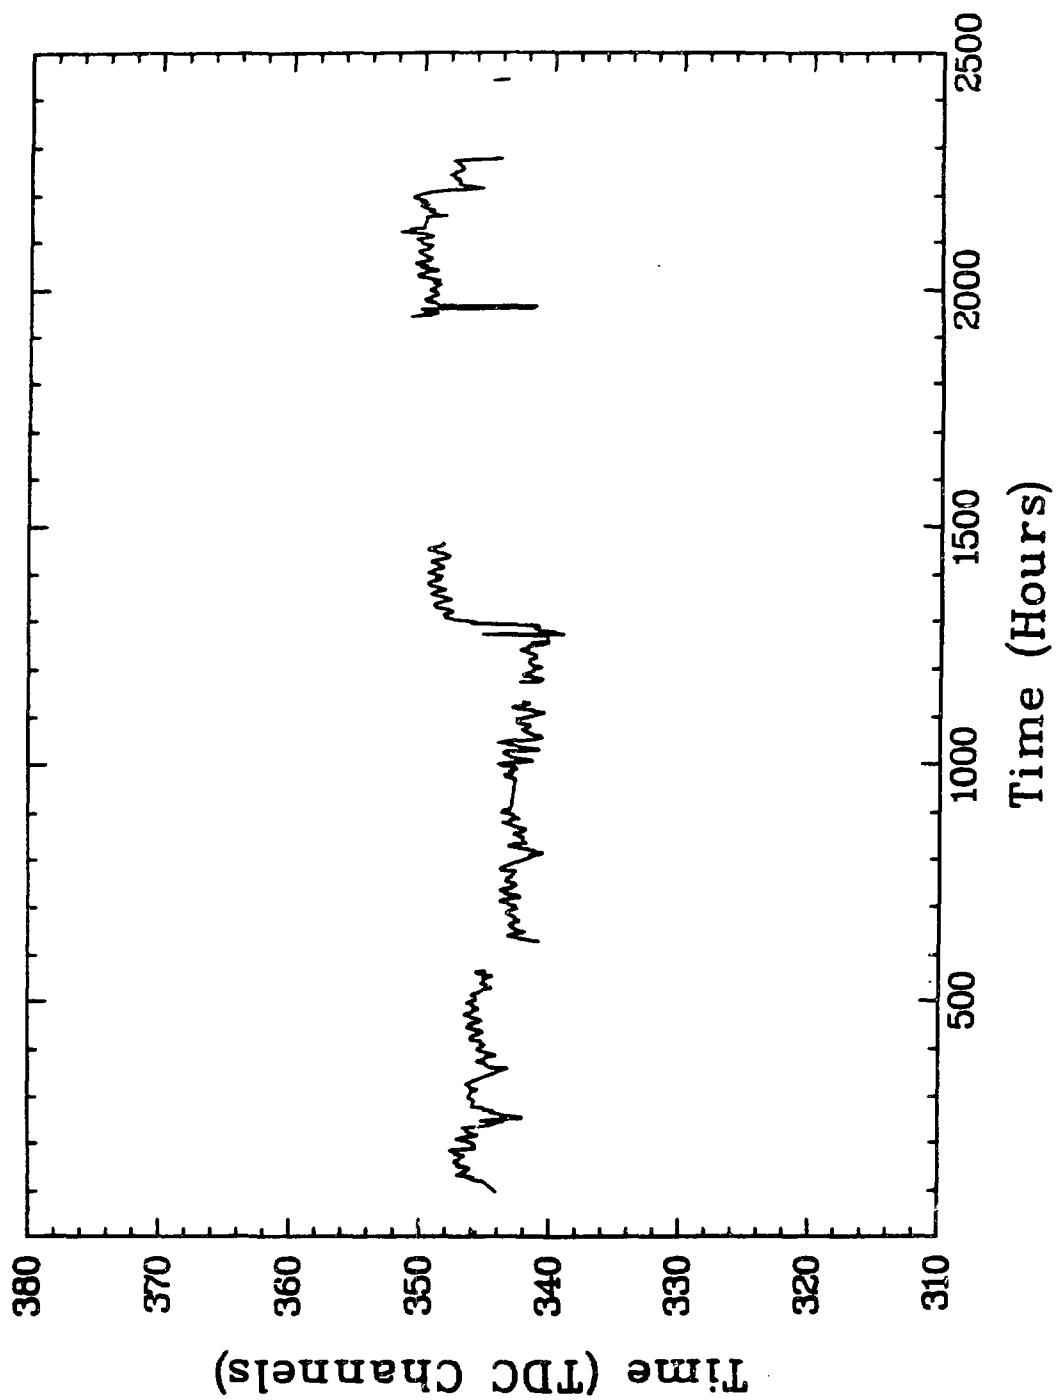


Figure 39. Plot of the average timing peak positions from the 432 NaI(Tl) TDC channels using the flasher system as a function of time.



Several changes to the flasher system could improve its performance. The shape of the light pulse is similar to that of the NaI(Tl) pulse, but it is not an exact duplicate. For low levels of light entering a crystal, the constant fraction discriminator channel does not trigger properly, because of the difference in the light pulse shape. The amount of light loss could be reduced by using a single length of fiber optic cable instead of two, and by mounting the optical fiber viewing the light source differently. The energy gain measurements indicate that the monitors may not have been as stable as desired, but it is not obvious how they can be improved.

REFERENCES

1. C.D. Anderson and S.H. Neddermeyer, Phys. Rev. 50, 263 (1936); S.H. Neddermeyer and C.D. Anderson, Phys. Rev. 51, 884 (1937), Phys. Rev. 54, 88 (1938), and Rev. Mod. Phys. 11, 191 (1939); see also J.C. Street and E.C. Stevenson, Phys. Rev. 51, 1005 (1937).
2. H. Yukawa, Proc. Phys. Math. Soc. Japan 17, 48 (1935).
3. M. Conversi, E. Pancini, and O. Piccioni, Phys. Rev. 71, 209 (1947).
4. V.L. Fitch and J. Rainwater, Phys. Rev. 92, 789 (1953).
5. R.L. Garwin, L.M. Lederman, and M. Weinrich, Phys. Rev. 105, 1415 (1957); J.I. Friedman and V.L. Telegdi, Phys. Rev. 105, 1681 (1957).
6. L. Michel, Proc. Phys. Soc. A63, 514 (1950); C. Bouchiat and L. Michel, Phys. Rev. 106, 170 (1957); C. Jarlskog, Nucl. Phys. 75, 659 (1966).
7. E.J. Konopinski and H.M. Mahmoud, Phys. Rev. 92, 1045 (1953).
8. J. Schwinger, Ann. Phys. (N.Y.) 2, 407 (1957).
9. K. Nishijima, Phys. Rev. 108, 907 (1957); S.A. Bludman, Nuovo Cimento 9, 433 (1958).
10. M.L. Perl et al., Phys. Rev. Lett. 35, 1489 (1975).
11. G. Feinberg and S. Weinberg, Phys. Rev. Lett. 6, 381 (1961).
12. B. Pontecorvo, Zh. Eksp. Teor. Fiz. 37, 1751 (1959) [Sov. Phys. JETP 10, 1236 (1960)]; M. Schwartz, Phys. Rev. Lett. 4, 306 (1960).

13. G. Danby et al., Phys. Rev. Lett. 9, 36 (1967).
14. W. Sullivan, New York Times, p. A24 (February 9, 1977); and Tribune de Genève, p. 7 (February 10, 1977).
15. H.P. Povel et al., Phys. Lett. 72B, 183 (1977); A. van der Schaaf et al., Nucl. Phys. A 340, 249 (1980).
16. P. Depommier et al., Phys. Rev. Lett. 39, 1113 (1977).
17. J.D. Bowman et al., Phys. Rev. Lett. 42, 556 (1979); W.W. Kinnison et al., Phys. Rev. D 25, 2846 (1982).
18. C.M. Hoffman, in Fundamental Interactions in Low-Energy Systems, edited by P. Dalpiaz, G. Fiorentini, and G. Torelli (Plenum, New York, 1985), p. 138.
19. E.P. Hincks and B. Pontecorvo, Phys. Rev. 73, 257 (1948); R.D. Sard and E.J. Althaus, Phys. Rev. 74, 1364 (1948).
20. R.D. Bolton et al., Phys. Rev. Lett. 56, 2461 (1986).
21. W. Bertl et al., Nucl. Phys. B 260, 1 (1985).
22. G. Azuelos et al., Phys. Rev. Lett. 51, 164 (1983).
23. C.M. York, C.O. Kim, and W. Kernan, Phys. Rev. Lett. 3, 288 (1959); ibid. 4, 320(E) (1960).
24. S. Frankel et al., Nuovo Cimento 27, 894 (1963).
25. J.M. Poutissou et al., Nucl. Phys. B 80, 221 (1974).
26. J.D. Bowman, T.P. Cheng, L.-F. Li, and H.S. Matis, Phys. Rev. Lett. 41, 442 (1978).
27. LAMPP Research Proposal 445, H.S. Matis and J.D. Bowman, spokespersons, Los Alamos National Laboratory, 1978 (unpublished).

28. Addendum to Research Proposals 400/445, M. Duong-van and C.M. Hoffman, and H.S. Matis and J.D. Bowman, spokespersons, Los Alamos National Laboratory, 1979 (unpublished).
29. For a description of gauge theories, see E.S. Abers and B.W. Lee, *Phys. Rep.* 9, 1 (1973); M.A.B. Beg and A. Sirlin, *Ann. Rev. Nucl. Sci.* 24, 379 (1974); *Phys. Rep.* 88, 1 (1982); E. Leader and E. Predazzi, An Introduction to Gauge Theories and the 'New Physics' (Cambridge University Press, Cambridge, England, 1982); T.P. Cheng and L.-F. Li, Gauge Theory of Elementary Particle Physics (Oxford University Press, New York, 1984).
30. S.L. Glashow, *Nucl. Phys.* 22, 579 (1961); A. Salam, in Elementary Particle Theory: Relativistic Groups and Analyticity (Nobel Symposium No. 8), edited by N. Svartholm (Almqvist and Wiksell, Stockholm, 1968), p. 367; S. Weinberg, *Phys. Rev. Lett.* 19, 1264 (1967); *ibid.* 27, 1688 (1971).
31. For an introduction to quantum electrodynamics, see R.P. Feynman, Quantum Electrodynamics (Benjamin/Cummings, Reading, Mass., 1961). For a collection of papers on the subject, see J. Schwinger, Selected Papers on Quantum Electrodynamics (Dover Publications, New York, 1958).
32. P.B. Schwinberg, R.S. Van Dyck, Jr., and H.G. Dehmelt, *Phys. Rev. Lett.* 47, 1679 (1981); J. Bailey et al., *Nucl. Phys. B* 150, 1 (1979).
33. E. Fermi, *Z. Phys.* 88, 161 (1934) [F.L. Wilson, translator, *Am. J. Phys.* 36, 1150 (1968)]; *Nuovo Cimento* 11, 1 (1934).

34. T.D. Lee and C.N. Yang, Phys. Rev. 104, 254 (1956); C.S. Wu et al., Phys. Rev. 105, 1413 (1957).
35. S.S. Gershtein and I.B. Zel'dovich, Zh. Eksp. Teor. Fiz. 29, 698 (1955) [Sov. Phys. JETP 2, 576 (1956)]; R.P. Feynman and M. Gell-Mann, Phys. Rev. 109, 193 (1958); E.C.G. Sudarshan and R.E. Marshak, Phys. Rev. 109, 1860 (1958); J.J. Sakurai, Nuovo Cimento 7, 649 (1958).
36. C.N. Yang and R.L. Mills, Phys. Rev. 96, 191 (1954).
37. G. 't Hooft, Nucl. Phys. B 33, 173 (1971); ibid. 35, 167 (1971).
38. Y. Nambu, Phys. Rev. Lett. 4, 380 (1960); Y. Nambu and G. Jona-Lasinio, Phys. Rev. 122, 345 (1961); ibid. 124, 246 (1961); J. Goldstone, Nuovo Cimento 19, 154 (1961); J. Goldstone, A. Salam, and S. Weinberg, Phys. Rev. 127, 965 (1962).
39. P.W. Higgs, Phys. Lett. 12, 132 (1964); Phys. Rev. Lett. 13, 508 (1964); Phys. Rev. 145, 1156 (1966).
40. C. Rubbia, Rev. Mod. Phys. 57, 699 (1985).
41. G. Arnison et al., Phys. Lett. 129B, 273 (1983); P. Bagnaia et al., Phys. Lett. 129B, 130 (1983).
42. R. Hofstadter and R.W. McAllister, Phys. Rev. 98, 217 (1955).
43. O.W. Greenberg, Phys. Rev. Lett. 13, 598 (1964); M.Y. Han and Y. Nambu, Phys. Rev. 139B, 1006 (1965).
44. For a review of free quark searches, see L.W. Jones, Rev. Mod. Phys. 49, 717 (1977); L. Lyons, in Progress in Particle and Nuclear Physics, edited by D. Wilkinson (Pergamon Press, Oxford, 1981), Vol. 7, p. 157.

45. W.K.H. Panofsky, in Proceedings of the 14th International Conference on High-Energy Physics, Vienna, 1968, edited by J. Prentki and J. Steinberger (CERN, Scientific Information Service, Geneva, 1968), p. 23; M. Breidenbach et al., Phys. Rev. Lett. 23, 935 (1969). For a review of deep inelastic lepton scattering, see J.I. Friedman and R.W. Kendall, Ann. Rev. Nucl. Sci. 22, 203 (1972).
46. D.P. Barber et al., Phys. Rev. Lett. 43, 1915 (1979).
47. M. Kobayashi and T. Maskawa, Prog. Theor. Phys. 49, 652 (1973).
48. M. Gell-Mann, in Proceedings of the 1962 International Conference on High-Energy Physics at CERN, edited by J. Prentki (CERN, Scientific Information Service, Geneva, 1962), p. 805; R. Behrends et al., Rev. Mod. Phys. 34, 1 (1962); S.L. Glashow and J.J. Sakurai, Nuovo Cimento 25, 337 (1962).
49. N. Cabibbo, Phys. Rev. Lett. 10, 531 (1963).
50. S.L. Glashow, J. Iliopoulos, and L. Maiani, Phys. Rev. D 2, 1285 (1970).
51. J.J. Aubert et al., Phys. Rev. Lett. 33, 1404 (1974); J.E. Augustin et al., Phys. Rev. Lett. 33, 1406 (1974).
52. S.W. Herb et al., Phys. Rev. Lett. 39, 252 (1977).
53. F.J. Hasert et al., Phys. Lett. 46B, 138 (1973).
54. G. Arnison et al., Phys. Lett. 122B, 103 (1983); M. Banner et al., Phys. Lett. 122B, 476 (1983).
55. E. Noether, Nachr. kgl. Ges. Wiss. Göttingen, 235 (1918); see also E.L. Hill, Rev. Mod. Phys. 23, 253 (1951).

56. W. Pauli (private correspondence with O. Hahn and L. Meitner, 1931); W. Pauli, in Proc. VII Solvay Congress, Brussels (Gauthier-Villars, Paris, 1933), p. 324.
57. A. De Rújula, H. Georgi, and S.L. Glashow, Phys. Rev. D 12, 147 (1975).
58. T.D. Lee and C.N. Yang, Phys. Rev. 98, 1501 (1955).
59. LAMPF Research Proposal 969, M. Cooper, spokesperson, Los Alamos National Laboratory, 1985 (unpublished).
60. T.P. Cheng and L.-F. Li, Phys. Rev. D 16, 1425, 1565 (1977); Phys. Rev. Lett. 38, 381 (1977); ibid. 45, 1908 (1980); B.W. Lee, S. Pakvasa, R.E. Shrock, and H. Sugawara, Phys. Rev. Lett. 38, 937, 1230(E) (1977); B.W. Lee and R.E. Shrock, Phys. Rev. D 16, 1444 (1977); S.B. Treiman, F. Wilczek, and A. Zee, Phys. Rev. D 16, 152 (1977); G. Altarelli, L. Baulieu, N. Cabibbo, L. Maiani, and R. Petronzio, Nucl. Phys. B 125, 285 (1977); H. Fritzsch, Phys. Lett. 67B, 451 (1977); J.D. Bjorken, K. Lane, and S. Weinberg, Phys. Rev. D 16, 1474 (1977); V.J. Marciano and A.I. Sanda, Phys. Lett. 67B, 303 (1977); S.M. Bilenky, S.T. Petcov, and B. Pontecorvo, Phys. Lett. 67B, 309 (1977); V. Barger and D.V. Nanopoulos, Nuovo Cimento 44A, 303 (1978).
61. A. Barroso, G.C. Branco, and M.C. Bento, Phys. Lett. 134B, 123 (1984).
62. F. Wilczek and A. Zee, Phys. Rev. Lett. 38, 531 (1977).
63. J.D. Bjorken and S. Weinberg, Phys. Rev. Lett. 38, 622 (1977); G.C. Branco, Phys. Lett. 68B, 455 (1977); W. Konetschny and W. Kummer, Phys. Lett. 70B, 433 (1977); A. Zee, Phys. Lett. 93B, 389 (1980).

64. J.C. Pati and A. Salam, Phys. Rev. D 10, 275 (1974);
R.N. Mohapatra and J.C. Pati, Phys. Rev. D 11, 566, 2558 (1975);
R.N. Mohapatra and R.E. Marshak, Phys. Rev. Lett. 44, 1316,
1644(E) (1980); R.N. Mohapatra and G. Senjanović, Phys. Rev. D 23,
165 (1981); Riazuddin, R.E. Marshak, and R.N. Mohapatra, Phys.
Rev. D 24, 1310 (1981); P.M. Fishbane, K. Gaemers, S. Meshkov, and
R.E. Norton, Phys. Rev. D 32, 1186 (1985).
65. R.N. Cahn and H. Harari, Nucl. Phys. B 176, 135 (1980); G.L. Kane
and R. Thun, Phys. Lett. 94B, 513 (1980); O. Shanker, Phys. Rev. D
23, 1555 (1981); Nucl. Phys. B 185, 382 (1981); D.R.T. Jones,
G.L. Kane, and J.P. Leveille, Nucl. Phys. B 198, 45 (1982).
66. For a review of technicolor, see E. Farhi and L. Susskind, Phys.
Rep. 74, 277 (1981). See also J. Ellis, D.V. Nanopoulos, and
P. Sikivie, Phys. Lett. 101B, 387 (1981); J. Ellis, M.K. Gaillard,
D.V. Nanopoulos, and P. Sikivie, Nucl. Phys. B 182, 529 (1981).
67. S. Dimopoulos and L. Susskind, Nucl. Phys. B 155, 237 (1979);
S. Dimopoulos and J. Ellis, Nucl. Phys. B 182, 505 (1981);
68. For a review of supersymmetry, see P. Fayet and S. Ferrara,
Phys. Rep. 32, 249 (1977); S. Ferrara, Phys. Rep. 105, 5 (1984);
P. Fayet, Phys. Rep. 105, 21 (1984); C.H. Llewellyn Smith, Phys.
Rep. 105, 53 (1984); D.V. Nanopoulos, Phys. Rep. 105, 71 (1984);
A. Savoy-Navarro, Phys. Rep. 105, 91 (1984); J. Ellis, Phys. Rep.
105, 121 (1984).
69. J. Ellis and D.V. Nanopoulos, Phys. Lett. 110B, 44 (1982).

70. For a review of composite models, see M. Peskin, in Proceedings of the 1981 International Symposium on Lepton and Photon Interactions at High Energies, edited by W. Pfeil (Physikalisches Institut, Universität Bonn, Bonn, 1981), p. 880; L. Lyons, in Progress in Particle and Nuclear Physics, edited by D. Wilkinson (Pergamon Press, Oxford, 1983), Vol. 10, p. 227.
71. H. Georgi and S.L. Glashow, Phys. Rev. Lett. 32, 438 (1974); J. Ellis and M.K. Gaillard, Phys. Lett. 88B, 315 (1979); R. Barbieri, J. Ellis, and M.K. Gaillard, Phys. Lett. 90B, 249 (1980). For a review, see P. Langacker, Phys. Rep. 72, 185 (1981).
72. R.M. Bionta et al., Phys. Rev. Lett. 51, 27, 522(E) (1983).
73. H. Fritzsch and P. Minkowski, Ann. Phys. (N.Y.) 93, 193 (1975); H. Georgi and D.V. Nanopoulos, Nucl. Phys. B 155, 52 (1979).
74. R. Barbieri and D.V. Nanopoulos, Phys. Lett. 91B, 369 (1980); F. Gürsey, P. Ramond, P. Sikivie, Phys. Lett. 60B, 177 (1975); O.K. Kalashnikov and S.E. Konstein, Nucl. Phys. B 166, 507 (1980); H. Ruegg and T. Schucker, Nucl. Phys. B 161, 388 (1979).
75. G.L. Kane and R.E. Shrock, in AIP Conference Proceedings No. 102, Intense Medium Energy Sources of Strangeness, UC - Santa Cruz, 1983, edited by T. Goldman, H.E. Haber, and H.F.-W. Sadrozinski (American Institute of Physics, New York, 1983), p. 123.
76. J. Dreitlein and H. Primakoff, Phys. Rev. 126, 375 (1962).
77. M.K. Gaillard and B.W. Lee, Phys. Rev. D 10, 897 (1974).
78. Y. Tomozawa, Phys. Rev. D 25, 1448 (1982).

79. J. Sucher, Rep. Prog. Phys. 41, 1781 (1978).
80. V. Višnjić-Triantafillou, Phys. Lett. 95B, 47 (1980).
81. A.E. Pifer, T. Bowen, and K.R. Kendall, Nucl. Instrum. Methods 135, 39 (1976); T. Bowen, Phys. Today, 38 (7), 22 (1985).
82. H.-W. Reist et al., Nucl. Instrum. Methods 153, 61 (1978); P.A. Thompson et al., Nucl. Instrum. Methods 161, 391 (1979).
83. R.D. Bolton et al., Nucl. Instrum. Methods 188, 275 (1981); R.D. Bolton et al., Nucl. Instrum. Methods 217, 173 (1983).
84. R.D. Bolton et al., Nucl. Instrum. Methods A241, 52 (1985).
85. Bertan Associates, Inc., 3 Aerial Way, Syosset, NY 11791.
86. Hewlett-Packard Co., 3404 East Harmony Road, Fort Collins, CO 80525.
87. J.P. Sandoval et al., Nucl. Instrum. Methods 216, 171 (1983).
88. LeCroy Research Systems Corporation, 700 South Main Street, Spring Valley, NY 10977.
89. G.H. Sanders et al., Nucl. Instrum. Methods 180, 603 (1981).
90. V.D. Sandberg et al., Nucl. Instrum. Methods A234, 512 (1985).
91. Amperex Electronic Corp., Hicksville Div., 230 Duffy Ave., Hicksville, NY 11802.
92. R.L. Heath, R. Hofstadter, and E.B. Hughes, Nucl. Instrum. Methods 162, 431 (1979), and references therein.
93. Tracor Northern Inc., 2551 W. Beltline Hwy., Middleton, WI 53562.
94. Allen Avionics Inc., Catalog 24D, 224 E. Second St., Mineola, NY 11501 (1982).
95. J. Rolfe et al., IEEE Trans. Nucl. Sci. NS-30 (1), 202 (1983).

96. DSP Technology, Inc., 48500 Kato Road, Fremont, CA 94538-7338.
97. G.H. Sanders et al., Los Alamos National Laboratory document LA-UR-81-1323, 1981 (unpublished).
98. D.G. Perry, IEEE Trans. Nucl. Sci. NS-26 (4), 4422 (1979).
99. L.R. Biswell and R.E. Rajala, Los Alamos National Laboratory document LA-5144, 1973 (unpublished); R.F. Thomas, Jr., Los Alamos National Laboratory document LA-5404-MS, 1973 (unpublished).
100. M. Kellogg et al., Los Alamos National Laboratory document LA-7001-M, 1978 (unpublished); D.G. Perry, IEEE Trans. Nucl. Sci. NS-26 (4), 4494 (1979).
101. H.S. Butler et al., IEEE Trans. Nucl. Sci. NS-28 (5), 3863 (1981).
102. R.L. Ford and W.R. Nelson, Stanford Linear Accelerator Center Report No. SLAC-PUB-210, 1978 (unpublished).
103. S.L. Wilson, Ph.D. thesis, Stanford University, Los Alamos National Laboratory Report LA-10471-T, 1985.
104. R.L. Ford et al., IEEE Trans. Nucl. Sci. NS-24 (1), 264 (1977); Y. Chan et al., IEEE Trans. Nucl. Sci. NS-25 (1), 333 (1978).
105. R.L. Carrington et al., Nucl. Instrum. Methods 163, 203 (1979).
106. EG&G Electro-Optics, 35 Congress St., Salem, MA 01970.
107. ITT, Electro-Optical Products Division, 3700 E. Pontiac St., Fort Wayne, IN 46803.
108. E.I. duPont de Nemours Co., Inc., Plastic and Resins Dept., Wilmington, DE 19898.
109. AMP Incorporated, Harrisburg, PA 17105.

ACKNOWLEDGMENTS

I thank my advisor, Courtenay Wright, for his support over the past years. With his help, I had the opportunity to participate in an exciting and interesting research project.

I especially thank Cy Hoffman, Gary Hogan, and Vern Sandberg for the many conversations regarding particle physics and in particular, the Crystal Box experiments. I greatly appreciate their patience and helpful guidance. Peter Heusi, Aksel Hallin, and Leo Piilonen deserve special thanks for their hard work on the data analysis and their useful suggestions.

Many thanks go to the rest of the Crystal Box collaboration: Rick Bolton, Dave Bowman, Martin Cooper, Jim Frank, Fasseha Mariam, Howard Matis, Dick Mischke, Darragh Nagle, Gary Sanders, Urs Sennhauser, Dick Werbeck, and Bob Williams, all from Los Alamos; Steve Wilson, Robert Hofstadter, Barrie Hughes, and Mike Ritter, from Stanford University; and Virgil Highland and Jim McDonough, from Temple University. All of their devoted efforts and contributions to the success of this experiment are deeply appreciated.

John Sandoval, Scott Chesney, Greg Hart, Mike Dugan, Camilo Espinoza, George Krausse, Jim Sena, Linda Bayliss, and Teresa Gordon gave the needed technical assistance. Bob Damjanovich and Val Hart, of Los Alamos, and Bob Parks and Joe Rolfe, of Stanford, gave engineering assistance. Jan Novak and his staff provided support for the liquid hydrogen target. John Conway, Camilo Espinoza, Joan Harshman, Billy

Martinez, and Dave Sanchez painstakingly assembled the drift chamber. I am indebted to all these people for contributing their expertise.

I am grateful to Gary Hogan, Vern Sandberg, Steve Wilson, and Leo Piilonen for their help in providing figures for this thesis, and to Cy Hoffman for his comments on the content.