LA--10471-T

DE86 001233

LA-10471-T
Thesis

UC-34D
Issued: July 1985

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

Steven L. Wilson

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

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

Separate muon, electron, and tau numbers are conserved in the minimal standard model of electroweak interactions with massless neutrinos. However, in many extensions to the standard model, separate lepton numbers are not expected to be conserved quantities. A new search for muon number non-conserving processes has been undertaken at the Los Alamos Meson Physics Facility (LAMPF), specifically to look for three neutrinoless decay modes of the muon. The search for the decay of a muon to an electron and a photon is discussed here.

A new detector facility, located in the LAMPF stopped muon channel, was developed for this experiment. This Crystal Box detector consists of a cylindrical drift chamber surrounded by a plastic scintillator hodoscope and a large solid angle, modularized, NaI(Tℓ) calorimeter. The apparatus measures the trajectories, relative timing, and energies of charged particles and photons from the decays of positive muons stopped in a central target. The assembly and calibration of the detector is described, and the procedure for taking data is discussed.

The sample of 1.3 million candidate events, from the first data run of the Crystal Box, was analyzed using a maximum-likelihood method. The upper limit on the branching ratio, relative to normal muon decay, for a muon decaying to an electron and a photon is found to be consistent with previous measurements. With 90 percent confidence, the branching ratio for this neutrinoless decay is observed to be less than 2.8×10^{10}.

GLOSSARY OF SYMBOLS AND ABBREVIATIONS

ACC	Auxiliary Crate Controller
ACP	Auxiliary Control Port
ADC	Analog to Digital Converter
CAMAC	Computer Automated Measurement and Control NIM standard
CERN	European Center for Nuclear Research
CL	Confidence Level
DAP	Dataway Access Port
DEC	Digital Equipment Corporation
DMA	Direct Memory Access
GIM	Glashow-Iliopoulos-Maiani mechanism
GWS	Glashow-Weinberg-Salam theory
HPHC	High Pulse-Height Crystal
LAM	Look-At-Me signal
LAMPF	Los Alamos Meson Physics Facility
LH_2	Liquid Hydrogen
LRS	LeCroy Research Systems
MBD	Microprogrammable Branch Driver
NaI(Tℓ)	Sodium Iodide (Thallium-doped)
NS	Neighbor Set (of crystals)
PHIM	Pulse Height Integration Module
PMT	Photomultiplier Tube
SIN	Swiss Institute for Nuclear Research
TAC	Time to Analog Converter

A SEARCH FOR THE NEUTRINOLESS MUON DECAY $\mu^+ \rightarrow e^+ \gamma$

by

Steven L. Wilson

Chapter I

INTRODUCTION

The theory of quantum mechanics established the concept of quantized conservation laws. Electric charge and spin angular momentum were the first examples of an intrinsic quantum number. Many other intrinsic quantum numbers have been postulated to account for, or classify, the subset of all kinematically allowed interactions which is observed to occur in nature. Such quantum numbers include baryon number, lepton number, quark flavor (strangeness, charm, etc.), color, and, the quantity of interest for this thesis, (lepton) family number.

In recent years, a theoretical consensus has emerged that all fundamental interactions should be described by a non-abelian gauge theory of quantum fields. In this context, aside from quantities whose conservation is derived from a global space-time translation invariance (eg. energy, momentum), only those intrinsic quantum numbers which are associated with the global symmetry of a massless gauge boson (eg. electric charge), are expected to be exactly conserved. It is well known, for instance, that flavor is not conserved. In addition, gauge theories provide a natural explanation for the suppression of extremely rare processes, in virtue of the large masses of the vector bosons which mediate the short range interactions. Consequently, the last decade has seen a revival of interest in sensitive searches, for rare processes, previously abandoned in favor of exact conservation laws.

The copious flux of muons, currently available at meson factories, has attracted experimentalists to look again for muon family number violation. To date, no such neutrinoless decay mode of the muon has been discovered. A 1977 rumor, that the rare decay $\mu^+\rightarrow e^+\gamma$ had been observed at the Swiss Institute of Nuclear Research (SIN), prompted a collaboration of scientists from Stanford University, the University of Chicago, and the Los Alamos Meson Physics Facility (LAMPF) to try to confirm the observation. The purported SIN result was subsequently denied, and the LAMPF experiment went on to set the current upper limit[1]

$$BR(\mu^+\rightarrow e^+\gamma) \equiv \frac{\Gamma(\mu^+\rightarrow e^+\gamma)}{\Gamma(\mu^+\rightarrow e^+\nu_e\bar{\nu}_\mu)} \lesssim 1.7 \times 10^{-10} \text{ (90\% confidence level).} \tag{1}$$

The above experiment was assembled and carried out, with remarkable rapidity, primarily by using existing detectors and equipment. In light of the null result, and the plethora of new models generated in response to the possible observation of $\mu^+\rightarrow e^+\gamma$, the same collaboration decided to design a new apparatus especially suited for a more sensitive search.[2] This thesis describes that Crystal Box experiment, which searches simultaneously for the neutrinoless muon decays $\mu^+\rightarrow e^+\gamma$, $\mu^+\rightarrow e^+e^+e^-$, and $\mu^+\rightarrow e^+\gamma\gamma$. Only the $\mu^+\rightarrow e^+\gamma$ data analysis will be discussed in detail.

The theoretical and experimental background of rare muon-decay physics is discussed in Chapter II. The experiment design is described in Chapter III, and in Chapter IV the experimental procedure is explained. Details of the data analysis are given in Chapter V. A summary of results is presented in Chapter VI.

Chapter II

PHYSICS BACKGROUND

2.1 HISTORY OF THE MUON

The muon was discovered as a component of cosmic radiation, in 1937, by Anderson and Neddermeyer,[3] and Street and Stevenson.[4] It has played a vital role ever since in the understanding of fundamental interactions. Initially considered a candidate for the Yukawa meson, proposed as the mediating particle of the strong interaction, the muon has about the mass predicted to explain the short range of the strong force. However, early cosmic-ray studies showed that the muon does not participate in strong interactions with nuclei,[5] and so could not be the Yukawa meson. In fact, the muon was shown not to be a spin-zero meson at all, but is instead a spin-1/2 fermion.[6] The properties of the muon proved to be very similar to that of the electron, the muon differing only in mass (about 207 times greater than the electron), and in the fact that it is unstable. The electron and muon constituted a new classification of particles called leptons. A third lepton "family", designated tau, has more recently been discovered.[7]

Studies of the decay of cosmic-ray muons identified the charged components to be electrons,[8] showed that the neutral components were not photons,[9,10] and proved that it was not a two-body process.[11] The two neutral components were assumed to be neutrinos, the massless, weakly-interacting particles first postulated to account for the unobserved energy in nuclear beta-decays. Muon decay was therefore not an electromagnetic radiative process ($\mu \rightarrow e\gamma$), but, in strength and structure, was most like nuclear beta-decay. The similarity of the coupling constants for muon decay and nuclear beta-decay, and also muon capture, led to the establishment of the weak interaction as the fourth fundamental force in nature.

In the 1950's, the advent of accelerator sources of muons permitted more detailed muon decay studies. These experiments observed the positrons from μ^+ decay to have an asymmetrical angular distribution,[12],[13] and to be longitudinally polarized.[14-18] This evidence of parity violation in muon decay, along with earlier results on beta-decay (for example[19] in Co^{60}) confirmed the non-conservation of parity as a general feature of all weak processes.

The intermediate vector boson (now known as the $W^{+/-}$) was introduced by Schwinger,[20] in 1957, to remove the violation of unitarity in the current-current model of weak interactions. However, such a nonlocal interaction led to a calculated rate for $\mu \rightarrow e\gamma$ which far exceeded the existing experimental upper limit of 10^{-8} for the branching ratio relative to normal muon decay. This discrepancy motivated the successful search for two types of neutrinos by Danby, et. al., in 1962.[21] The existence of distinct muon and electron neutrinos gave credence to the idea[22] that a new quantum number, called muon number, was carried by the muon and its associated neutrino. The conservation of muon number then explained the non-observation of $\mu \rightarrow e\gamma$ and other neutrinoless muon decays.

The usual assignment of lepton family quantum numbers,[20,22,23] including the tau generation, is given in Table 1. Separate conservation of electron, muon, and tau numbers is assumed. With these assignments, the normal μ^+ decay mode is

$$\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu \ . \tag{2}$$

An earlier scheme which included electrons and muons only (see Table 2), proposed by Konopinski and Mahmoud in 1953,[24] is more compact and less stringent, but cannot be extended to include a third lepton family. A multiplicative scheme suggested by Feinberg and Weinberg in 1961,[25] indicated in Table 3, allows

$$\mu^+ \rightarrow e^+ \bar{\nu}_e \nu_\mu \tag{3}$$

in addition to mode (2). Since these two allowed decays enter this scheme equivalently, they would be expected to occur at equal rates. However, experiments[26,27] have shown the cross section for the related process, $\bar{\nu}_\mu e^- \rightarrow \mu^- \bar{\nu}_e$, to be small compared to that for "normal" inverse muon decay. The most recent result from LAMPF is an upper limit of[27]

$$\frac{\sigma(\bar{\nu}_\mu e^- \rightarrow \mu^- \bar{\nu}_e)}{\sigma(\nu_\mu e^- \rightarrow \mu^- \nu_e)} \lesssim 5 \times 10^{-2} \quad (90\% \text{ CL}). \tag{4}$$

The multiplicative scheme is also difficult to extend to a third lepton family, so the additive scheme of Table 1 appears to be the only viable candidate. By the mid-1960's, the idea of muon number conservation seemed well established, and, for the next decade or so, there was little interest in pursuing more sensitive searches for neutrinoless muon decay.

TABLE 1

Additive Lepton Number Assignments

L_e		L_μ		L_τ	
+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	Everything else	0	Everything else	0	Everything else

Conservations Laws: ΣL_e = Constant
ΣL_μ = Constant
ΣL_τ = Constant

Forbids: $\mu^+ \to e^+\gamma$
$\mu^+ \to e^+e^+e^-$
$\mu^-Z \to e^-Z$
$\mu^-Z \to e^+(Z-2)$
$K^0 \to \mu e$
$\mu^+ \to e^+\bar{\nu}_e\nu_\mu$

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

TABLE 2

Konopinski-Mahmoud Lepton Number Assignments

L	
+1	e^-, μ^+, ν
-1	e^+, μ^-, $\bar{\nu}$
0	Everything else

Conservation Law: $\Sigma L = \text{Constant}$

Forbids:
$\mu^+ \rightarrow e^+\gamma$
$\mu^+ \rightarrow e^+e^+e^-$
$\mu^-Z \rightarrow e^-Z$
$K^0 \rightarrow \mu^\pm e^\mp$

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

TABLE 3

Multiplicative Lepton Number Scheme

$\underline{L}$		$\underline{L_p}$	
+1	μ^-, e^-, ν_e, ν_μ	+1	e^+, ν_e , $\bar{\nu}_e$
-1	μ^+, e^+, $\bar{\nu}_e$, $\bar{\nu}_\mu$	-1	μ^+, ν_μ , $\bar{\nu}_\mu$
0	Everything else	+1	Everything else

Conservation Laws: ΣL = Constant

ΠL_p = Constant

Forbids:		Allows:	
	$\mu^+ \to e^+\gamma$		$\mu^+ \to e^+\nu_e\bar{\nu}_\mu$
	$\mu^+ \to e^+e^+e^-$		$\mu^+ \to e^+\bar{\nu}_e\nu_\mu$
	$\mu^-Z \to e^-Z$		$(\mu^+e^-) \to (\mu^-e^+)$
	$\mu^-Z \to e^+(Z-2)$		
	$K^0 \to \mu e$		

2.2 THEORETICAL ASPECTS OF RARE MUON DECAY

The present interest in rare muon-decays derives largely from the great progress made, over the last twenty years, in unifying the description of the weak and electromagnetic interactions within the framework of Yang-Mills gauge theories.[28] Since they are inherently renormalizable,[29] Yang-Mills theories provide an attractive way of solving the high energy divergence problem of the old V-A Fermi theory of weak interactions. In 1961, Glashow,[30] and, in 1964, Salam and Ward,[31] went further by proposing a combined "electro-weak" model with an SU(2)L X U(1) Yang-Mills gauge structure. The content of this model grouped the two known lepton families into left-handed doublets and right-handed singlets of massless fermions,

$$\begin{pmatrix}\nu_e\\ e\end{pmatrix}_L \quad \begin{pmatrix}\nu_\mu\\ \mu\end{pmatrix}_L \quad (e)_R \quad (\mu)_R \tag{5}$$

where

$$\Big(\;\Big)_{R,L} = (1\pm\gamma_5)\Big(\;\Big) \; . \tag{6}$$

The Yang-Mills construction requires a massless, spin-1 gauge boson for each parameter of the gauge symmetry group. In the present case, the U(1) group requires one such vector boson, B, which is electrically neutral. The three-parameter SU(2) group requires two charged gauge bosons, W^+ and W^-, and a single neutral one, W_3. The familiar photon field is a linear combination of the two neutral boson fields

$$A = \cos\theta_W\, B + \sin\theta_W\, W_3 \tag{7}$$

and the orthogonal combination

$$Z = -\sin\theta_W B + \cos\theta_W W_3 \tag{8}$$

emerges as the mediator of the weak neutral-current interactions which are the most striking new predictions of this model. The Weinberg angle, θ_W, is related to the amount of parity non-conservation present in neutral-current reactions, and along with the electromagnetic coupling constant, α, suffices to determine the two independent coupling constants of the theory.

A major problem with this model is that the weak mediating bosons remain massless, whereas the short range of that force requires them to be massive ($\simeq$ 100 GeV, see below). On the other hand, for the bosons to be massive would break the very symmetry from which their existence was deduced. A solution to this dilemma, provided by Weinberg[32] and Salam,[33] was to add a massive scalar doublet

$$\Phi = \begin{pmatrix} \phi^- \\ \phi^0 \end{pmatrix} \tag{9}$$

which couples to the leptons and gauge bosons, via the Higgs mechanism,[34] to give them appropriate masses. The SU(2)L symmetry is "spontaneously broken" by the configuration of minimum energy for Φ, and not in the equations of motion. At large distances (low energies), the remaining symmetry is the U(1) of electromagnetism.

In 1970, Glashow, Iliopoulos, and Maiani[35] incorporated the hadronic sector into the electroweak model. In the process, the charmed quark was predicted, and a mechanism proposed (GIM suppression) that explains the experimentally observed suppression of strangeness-changing weak neutral-current processes. The currently accepted "minimal" multiplet structure of fundamental fermions required to explain the known particles and processes, including the SU(3)C color symmetry for the QCD description of the strong interaction, is given in Figure 1. The three known lepton families combine with the six quark flavors, grouped in three weak isodoublets (each of which comes in three colors), to form three lepton/quark "generations." The top quark, required to complete the third fermion generation, has very recently been observed at the European Center for Nuclear Research (CERN).[36]

The successes of the Glashow-Weinberg-Salam (GWS) SU(2)L X U(1) electro-weak model have been many. Several weak neutral-current processes have been observed,[37] consistent with a unique value for $\sin^2\theta_W = 0.22$. The existence of the charmed quark has been verified,[38] and, most recently, the $W^{+/-}$ and Z^0 have been "directly" observed at CERN,[39] having masses ($m_W \simeq 81$ GeV, $m_Z \simeq 93$ GeV) in good agreement with GWS predictions using the above value of $\sin^2\theta_W$.

Despite these considerable successes, the GWS model is manifestly incomplete, since the fermion and Higgs masses, and the weak and electromagnetic coupling constants cannot be absolutely calculated. As mentioned above, it is minimal, containing only so many fundamental fermions as are necessary to accommodate the known electromagnetic and weak phenomena, and it provides no explanation of, or prediction for, the number of lepton/quark generations ([e/u,d],[μ/c,s],[τ/t,b],...). The Higgs sector is also minimal, containing only one isodoublet. The V-A structure of weak currents remains phenomenological and is put in

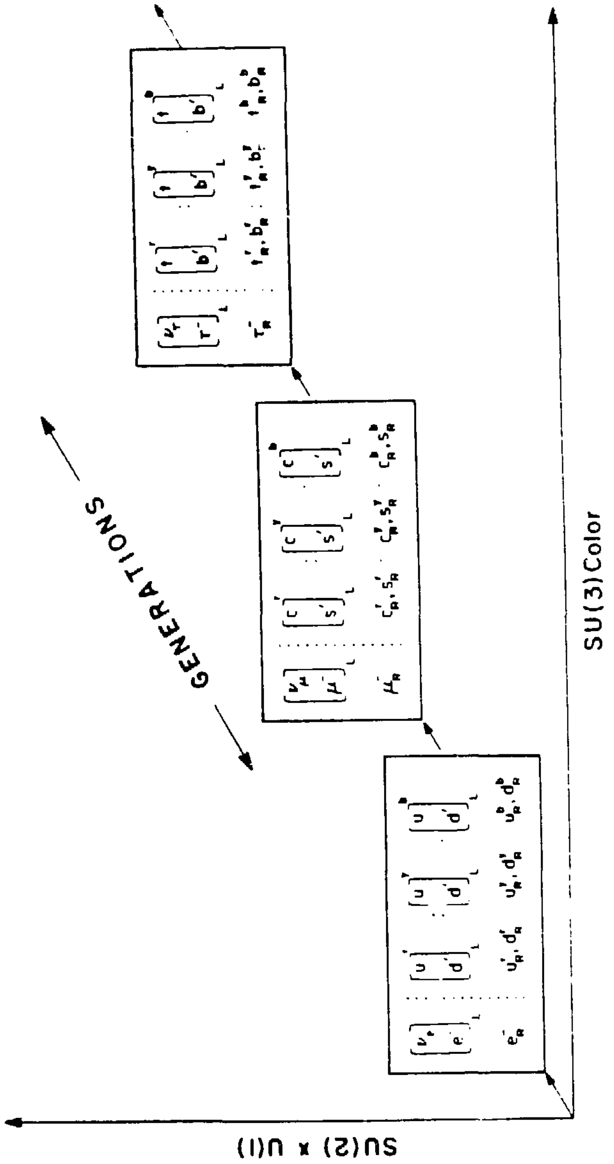

Figure 1: The currently accepted fermion multiplet structure.

The three known fermion generations are shown. The primed quarks indicate the Kobayashi-Maskawa mixing of the flavor eigenstates d, s, and b.

explicitly. Furthermore, neutrinos are assumed to be massless, and lepton number conservation is also assumed, leading necessarily to the separate conservation of lepton family numbers. Thus the observation of muon number non-conservation would imply that some extension of the GWS model is required.

Indeed, the consensus emerged that only those intrinsic quantities, which are associated with the global symmetry of a massless gauge boson, should be absolutely conserved. This idea, and the fact that all spontaneously broken Yang-Mills theories are renormalizable, prompted many authors to examine "non-minimal" models in the early 1970's. Some of these models addressed particular theoretical technicalities, some attempted to fit specific experimental details, and some showed no particular motivation at all.[40] Although it was not difficult to generate models for which approximate symmetries hold to lowest order in the weak coupling constant (G_F), many of the models predicted branching ratios for $\mu \to e\gamma$ which exceeded the existing experimental upper limit of about 10^{-8}. In 1977, Lee and Shrock[41] considered a general SU(2) X U(1) gauge theory and gave a set of criteria which assure "natural" suppression of (muon number) symmetry violation. The criteria are:

1. leptons of a given charge and chirality (handedness) have the same weak isospin T and T_3

2. leptons of a given charge receive their masses from their couplings to a single neutral Higgs field

3. leptons of charge q and leptons of charge q ± 1 do not belong to the same weak isomultiplet for at least one chirality

These authors showed how a generalization of the GIM mechanism leads to an additional suppression of symmetry violating amplitudes by a factor

$$S = \frac{\Delta M_l^2}{M_W^2} \tag{10}$$

where ΔM_l^2 is a "small" difference of squares of masses, and M_W is a heavy mass. The observed branching ratio for the rare process is reduced by the factor S^2, so that it is possible for these processes to be extremely rare and yet remain non-zero. This suppression factor arises directly from the particular group representations of the model, and is independent of the (incalculable) mixing parameters. A good discussion of the representation content, and the form of the calculated branching ratio for $\mu\to e\gamma$, in several models which satisfy the three criteria above, are given in reference 41. It eventually became apparent that the abundant choice of possible models could only be reduced in number with the help of further experimental data.

2.3 PREVIOUS SEARCHES FOR $\mu \to e \gamma$

In Figure 2, due to Hoffman,[42] the upper limits, for the branching ratios of several processes which violate muon number conservation, are plotted versus calendar time. As early as 1948, before the exact nature of normal muon decay was known, cosmic-ray searches for $\mu^+\to e^+\gamma$ set a limit of[10]

$$BR(\mu^+\to e^+\gamma) \lesssim 5 \times 10^{-2} \tag{11}$$

using Geiger-Mueller tubes and absorbers to measure the energy of the decay particles. With the increased muon intensities available from the first particle accelerators, experimentalists steadily reduced the upper limit for BR($\mu^+ \rightarrow e^+ \gamma$) until the early 1960's. Using a variety of detection schemes, including range measurements with scintillators, water Cerenkov counters, spark chambers, a freon bubble chamber, and thallium-doped sodium iodide [NaI(Tℓ)] crystals, these experiments[10,43,44] observed the decays of stopped positive muons. Stopped muons are used because the 2.2 μs lifetime of the muon would require a prohibitively long decay region, or would drastically reduce the useable intensity, for in-flight decay experiments. Positive muons are studied because they are not captured on a nucleus in a stopping target. The experimental signature for $\mu^+ \rightarrow e^+ \gamma$ is the detection of a time coincident electron-photon pair, emitted back-to-back in the lab, each with momentum $m_\mu c/2$ = 52.83 Mev/c. Given the poor energy resolution of the early experiments, the inner bremsstrahlung,

$$\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu \gamma \qquad (12)$$

from normal muon-decay, contributed a background to the above experimental signature, when the neutrinos have little energy and the e^+ and γ are, consequently, almost acollinear. By 1964, the background-limited result of[45] BR($\mu^+ \rightarrow e^+ \gamma$) $\lesssim 2.2 \times 10^{-8}$, consistent with existing detector resolutions, had been achieved. At this point, the apparent verification of muon number conservation, coupled with the need for improved detector technology and/or more intense muon beams to increase experimental sensitivity, served to discourage further searches for nearly a decade. (See Fig. 2)

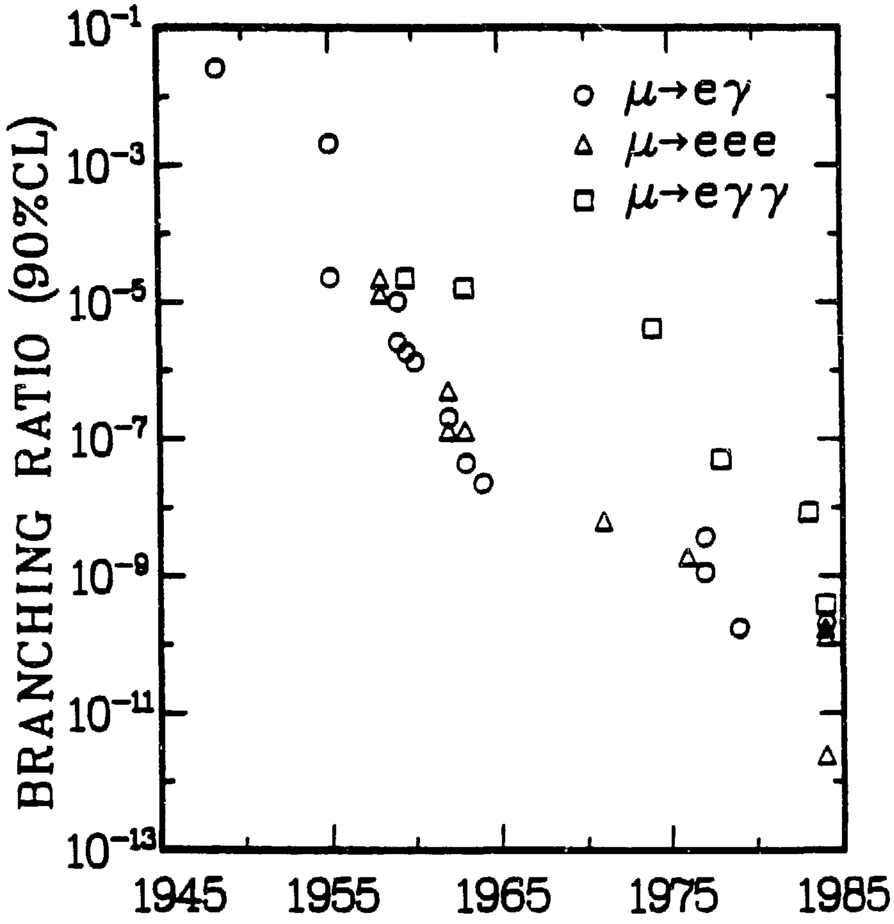

Figure 2: The history of rare muon decay experiments.

The improvement of experimental upper limits on branching ratios for several processes which violate lepton family number.

The improved beams available at meson factories in the early 1970's, and the changing theoretical prejudices, mentioned above in section 2.2, motivated a collaboration at SIN to attempt a new experiment.[46] It was similar in design to the 1963 experiment of Frankel, et.al.,[44] employing two large single crystal NaI(Tℓ) detectors on opposite sides of the stopping target. With a stopping rate of 5 x 10^5 μ/s, an upper limit of

$$BR(\mu^+\rightarrow e^+\gamma) \lesssim 1 \times 10^{-9} \quad (90\%\ CL) \tag{13}$$

was achieved. An essentially identical TRIUMF experiment,[47] using a less intense beam, set a weaker limit of

$$BR(\mu^+\rightarrow e^+\gamma) \lesssim 3.6 \times 10^{-9} \quad (90\%\ CL). \tag{14}$$

A LAMPF experiment[1] set the best limit to date at

$$BR(\mu^+\rightarrow e^+\gamma) \lesssim 1.7 \times 10^{-10} \quad (90\%\ CL). \tag{15}$$

The new techniques employed in this search included a low energy "surface" muon beam (making stopping rates of 2.4 x 10^6 μ/s possible), a magnetic spectrometer to measure the e^+ vector momentum with greatly improved resolution, and a segmented NaI(Tℓ) photon calorimeter, to allow measurement of the photon impact point. The NaI(Tℓ) calorimeter was screened from charged particles by a clearing magnet, to limit signal pile-up to a tolerable level, even at the very high stopping rate. The low energy beam permitted a thin target to be used, thus minimizing the target induced multiple scattering, bremsstrahlung, and annihilation of the decay positron. The prompt (internal bremsstrahlung) background was thereby reduced to the order of 10^{-12}, so that the dominant background was actually due to random coincidences of uncorrelated positrons and photons.

The Crystal Box detector[2] achieves a factor of 3 to 4 better sensitivity in branching ratio over the previous LAMPF experiment. This improvement is primarily due to the twenty-fold increase in solid angle (acceptance) and a two-fold improvement in timing resolution, while running at a ten times less intense stopping rate.

Chapter III

EXPERIMENT DESIGN

The overall layout of the experimental apparatus is shown in Figure 3. The muon beam is stopped in a thin polystyrene target suspended inside the cylindrical drift chamber. A small fraction of the muon beam passes through a hole in the target (see Figure 4) and stops in an aluminum cup surrounded by a hollow cylindrical plastic counter. Muon decays originating in this "I-counter" are used for timing reference and calibration purposes. Surrounding the drift chamber is a plastic scintillator trigger hodoscope which is mounted inside the central cavity of the modularized NaI(Tℓ) calorimeter. The fast trigger logic examines the timing, energy, and geometry of hits in the hodoscope counters and calorimeter modules. When acceptable events are recognized, the data acquisition system is prompted to read out the charged particle track information from the drift chamber, and the timing and energy data from both the plastic hodoscope and the NaI(Tℓ) calorimeter.

For calibration purposes, the drift chamber is removed, a liquid hydrogen (LH_2) target inserted, and the beam switched to negative pions. Gain stabilization and verification of system performance for the NaI(Tℓ) detector is achieved with the aid of light pulses introduced into the readout phototubes by the fiber-optic coupled xenon flashtube system. All of these major systems of the Crystal Box experiment will

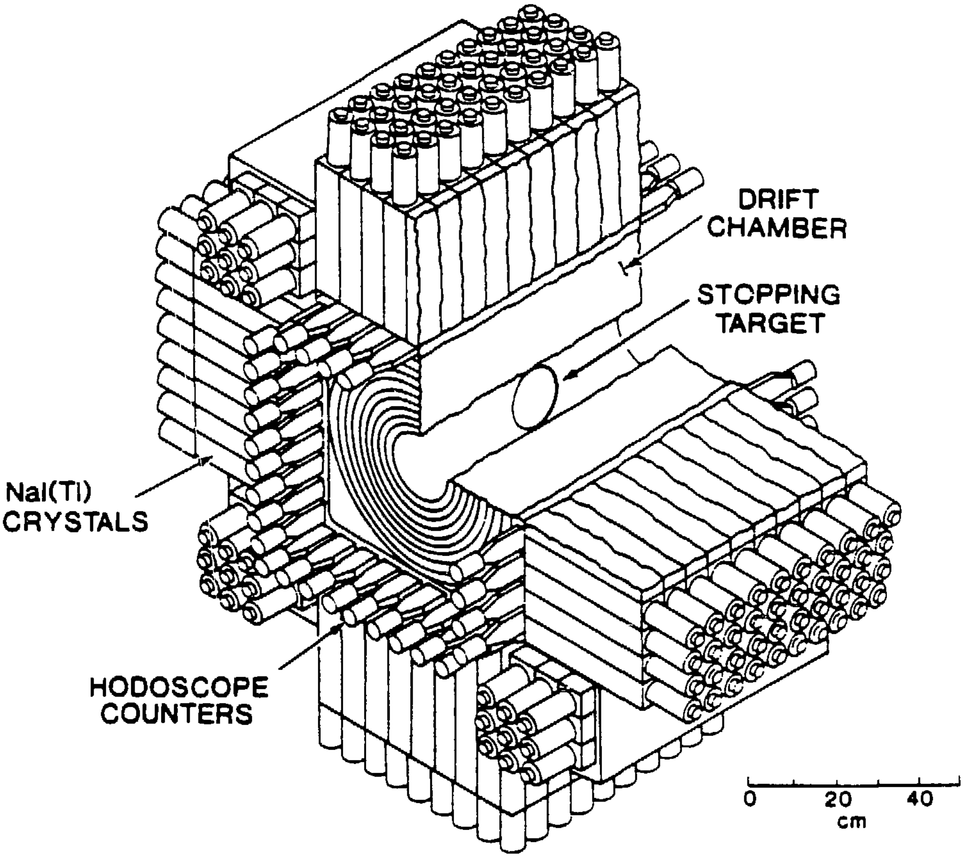

Figure 3: Detailed view of the Crystal Box detector.

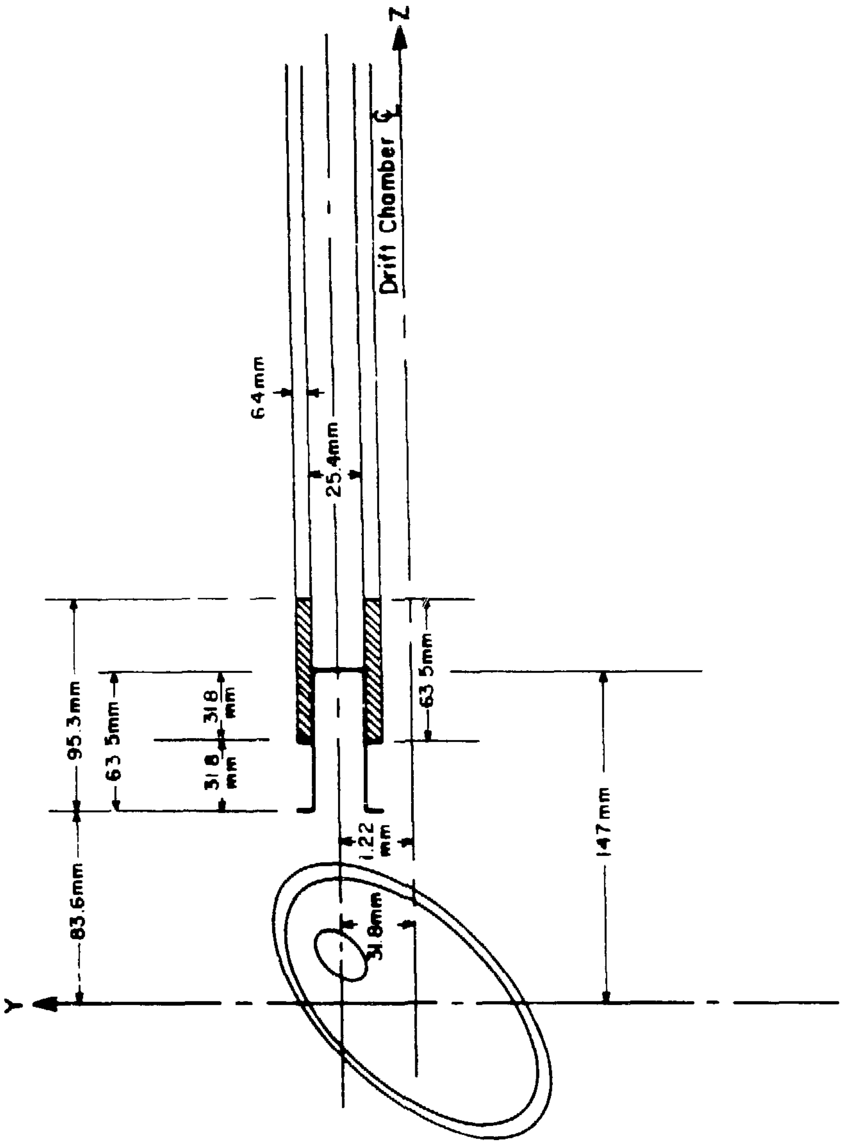

Figure 4: The position of the I-counter relative to the target.

The active plastic scintillator (shaded region) surrounds the downstream half of the Aℓ stopping cup.

be described, in this chapter, in more detail. However, since the author's primary responsibilities involved the operation of the NaI(Tℓ) calorimeter, the description of that system will be emphasized.

3.1 THE BEAM

The experimental source of muons was the LAMPF stopped muon channel (SMC).[48] The general layout of the beam transport is shown in Figure 5. The "surface" muons from the LAMPF A2 target provided the high-intensity beam necessary for a sensitive rare-decay search. The low momentum (26 MeV/c) of the beam permitted the use of a thin stopping target, thus reducing the target-induced photon backgrounds such as external bremsstrahlung. The stopping distribution of μ^+'s on the target was optimized for the $\mu \rightarrow 3e$ reaction, since the requirement of a common origin for three charged tracks provides a strong suppression of random backgrounds for that decay. Therefore the specification for the beam tuning and collimation was to illuminate the entire target area as uniformly as possible, while allowing less than 15% of the intensity to strike the inner foil of the drift chamber. This last requirement was to assure the satisfactory performance of the drift chamber. The beam-spot size and muon-flux intensity distribution achieved are discussed in section 4.1.1 below.

The rate of spurious hits in the detector, from positrons accepted by the beam transport, was substantially reduced by improving the e^+/μ^+ ratio from its intrinsic value of about 10:1 to a value of 1:50. This reduction was accomplished in two stages. The conventional separation method was to pass the beam through 15 mils of polyethylene degrader

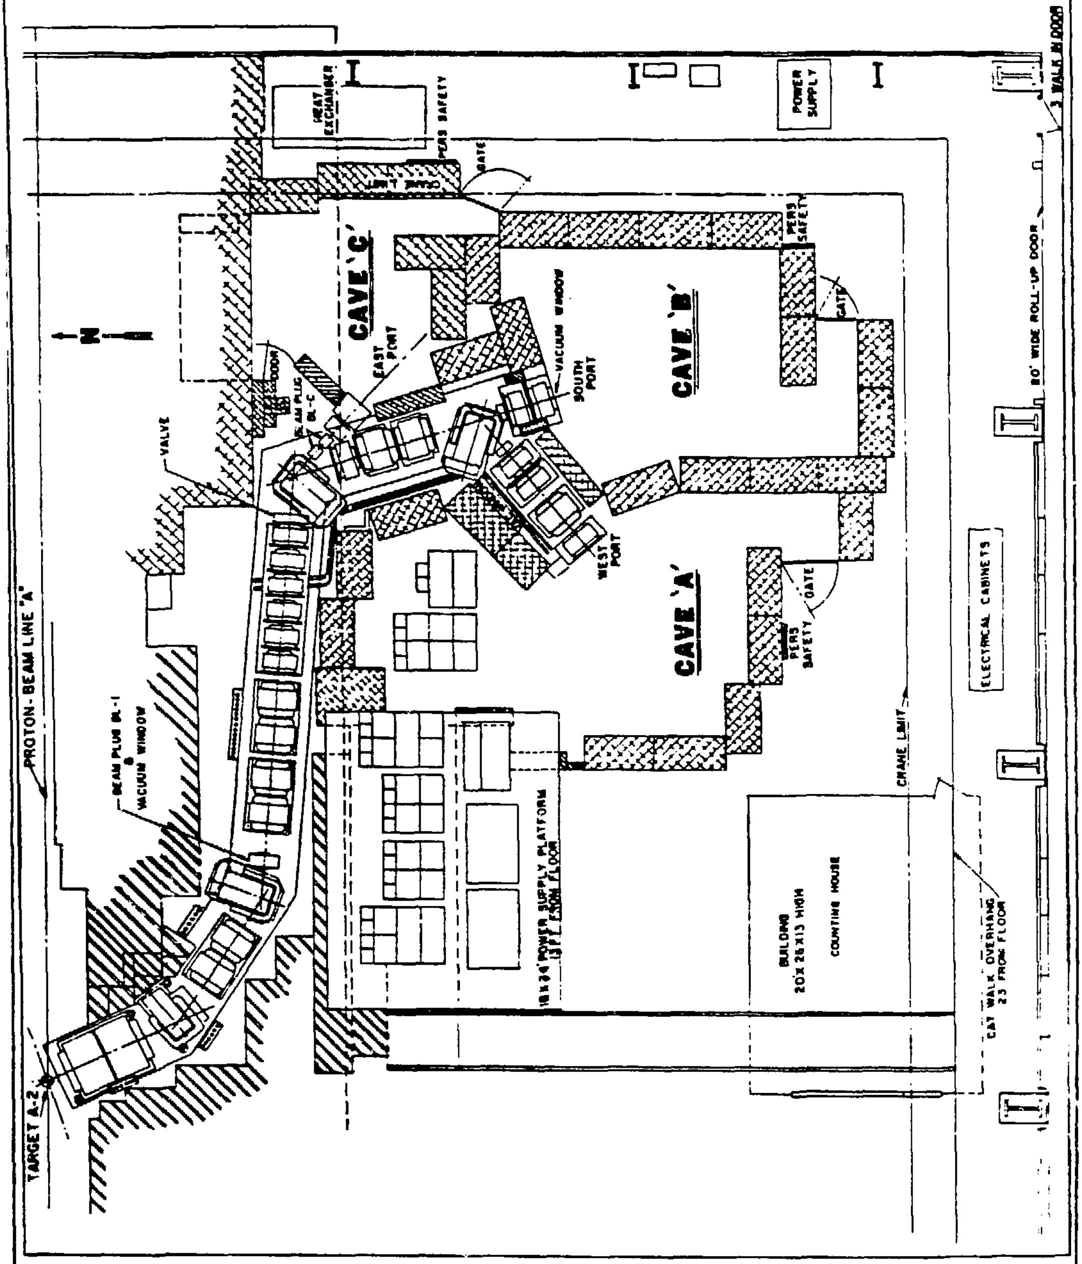

Figure 5: Floor plan of the SMC beamline in LAMPF experimental area A.

located just upstream of the bending magnet BM03 (see Figure 5). This magnet was tuned to accept the muons, which were degraded by about 3.5 MeV/c relative to the positrons. This stage reduces the ratio e^+/μ^+ to 1:1. The second, a velocity-selecting separator, employing crossed electric and magnetic fields, improved the μ^+ purity by another factor of 50. Some of the remaining e^+'s were due to μ^+ decays in flight after the separator.

The pulse structure of the beam consisted of 80 macropulses/s with a duty factor that varied from 6% to 10%. The average duty factor was 7%. Although rates of up to 4.5×10^5 μ^+/s stopping on the target were attainable, considerations of signal pile-up in the NaI(Tℓ) calorimeter limited the useful stopping rate to 3×10^5 μ^+/s.

A high-intensity pion beam was also required for the calibration data to be taken with the LH_2 target. A 140 MeV/c negative pion beam was developed that satisfactorily illuminated the stopping area of the LH_2 target. This beam will also be used for future experiments with the Crystal Box detector, including a search for $\pi^0 \rightarrow 3\gamma$ to be conducted in 1985.

3.2 THE EXPERIMENTAL ENVIRONMENT

The Crystal Box apparatus is located in cave B of the SMC facility (see Figure 5). To minimize the unwanted detector backgrounds from beam related particles, primarily neutrons, and from cosmic-rays, the apparatus is shielded with steel walls on four sides. The upstream side has 40.6 cm of steel shielding, the beam-right wall is 38.1 cm thick, the beam-left wall is 35.6 cm thick, and the top is covered with a 17.8

cm plate. This arrangement was sufficient to reduce the external background to a minimal level. Access to the apparatus was permitted through the unshielded downstream side.

3.3 THE APPARATUS

3.3.1 Target and Calibration I-Counter

The thin polystyrene target for stopping muons, shown in Figure 6, is positioned by an annular lucite holder attached to the drift chamber structure by piano (steel) wire. The target plane is rotated by 45° relative to the beam direction as indicated in Figure 3. The target shape projects a circular profile onto the faces of the Box. Polystyrene was chosen as a target material to depolarize the stopping muons. The residual polarization (along the beam direction) in this material was measured to be (14.6 ± 1.4)%.

A secondary stopping target, as shown in Figure 4, is positioned concentric with and downstream of a hole in the primary target. Muons stopping in the aluminum cup are only slightly depolarized. The cup is surrounded by a hollow cylinder of Pilot B plastic scintillator which is coupled by a cylindrical lucite light guide to an Amperex 2232B photomultiplier tube. The stopping volume of this I-counter was located as near as possible to the center of the Box without intercepting particle trajectories between the primary target and the fiducial area of the Box. The use of muons stopping in this counter for set-up and calibration of the detector is described in Chapters IV and V.

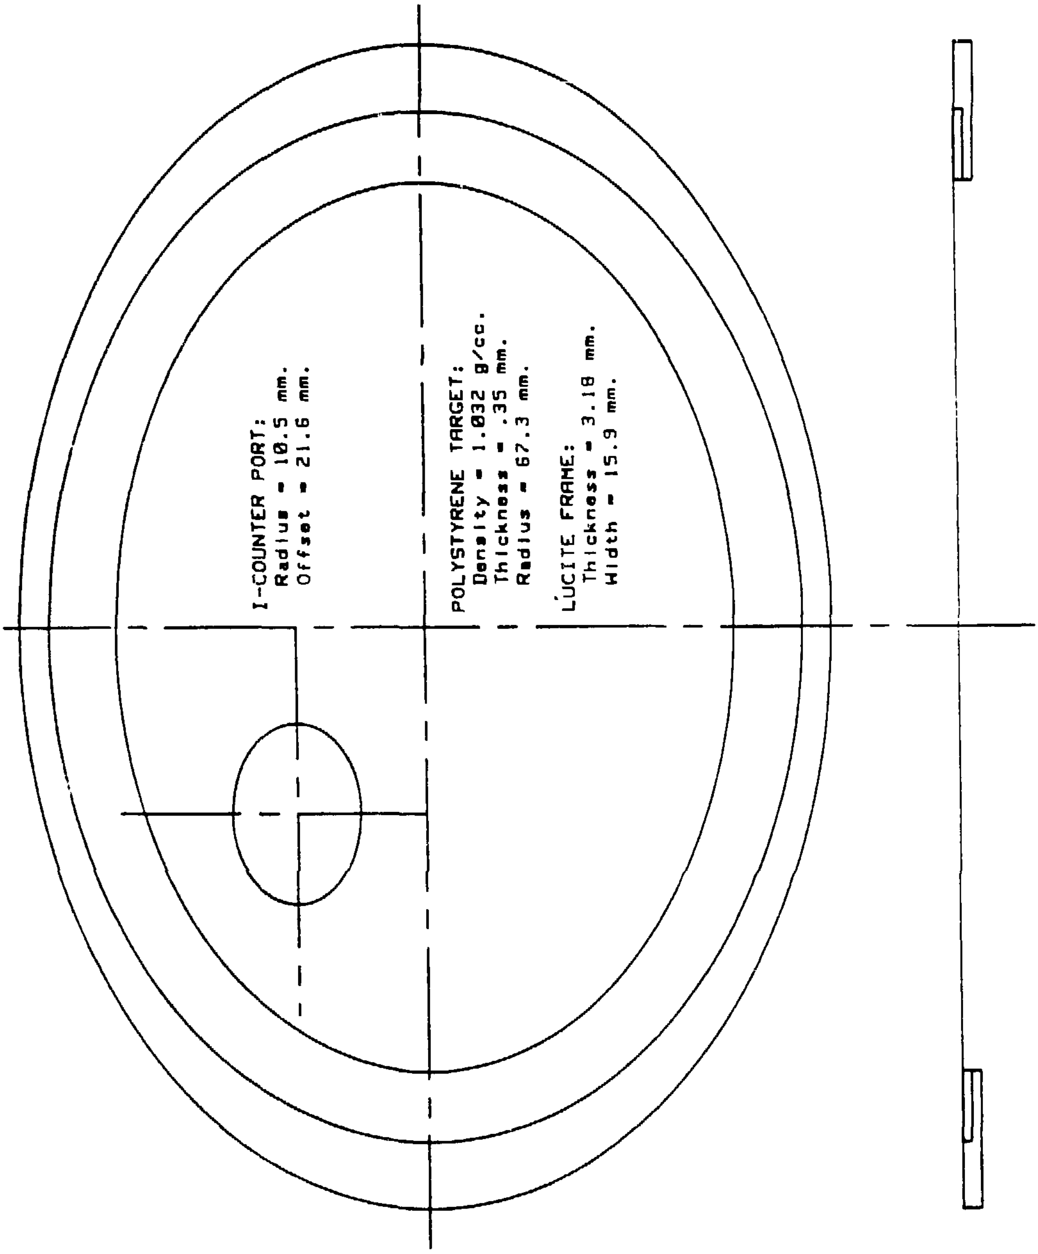

Figure 6: Dimensions and specifications of the primary stopping target.

3.3.2 Drift Chamber

The drift chamber[49] consists of eight concentric cylindrical wire shells employing 10° - 16° stereo angles. The configuration of wire connections in the endplates is illustrated in Figure 7. A photograph of the drift chamber under construction is shown in Figure 8. Each sense wire is surrounded by eight high voltage wires comprising an approximately 8 mm x 10 mm unit cell configuration as shown in Figure 9(a). High voltage is supplied by a Bertan B-Hi-VE controlled by a Hewlett-Packard HP-85 microcomputer. In a unit cell, the corner wires (triangles in Figure 9(a)) are maintained at -2500 volts, while the guard wires between adjacent cells (squares) are held at a 10% lower voltage. Guard wires facing a neighboring plane (circles) are at 72% of the corner wire voltage to provide an island of net positive charge that minimizes the plane-to-plane crosstalk. In addition, the voltages are chosen to keep the field lines as radial as possible (or the equipotentials approximately circular, see Figure 9b) so that the drift time depends only on the distance of the track from the sense wire. The ionization medium is a mixture of 49% argon, 49% ethane, and 2% isopropyl alcohol.

The sense wire signals are processed by the electronics chain as shown in Figure 10. Raw signals pass through approximately 3 m of RG-188 coaxial cable to an amplifier and discriminator system designed and built at LAMPF. The output signals are transferred by twisted-pair ribbon cable to ECL delay-line chips and then through twisted-pair cable again to a LeCroy 4290 time-digitizing system. This system consists of a model 4298 CAMAC crate controller, and 23 model 4291A time-to-digita

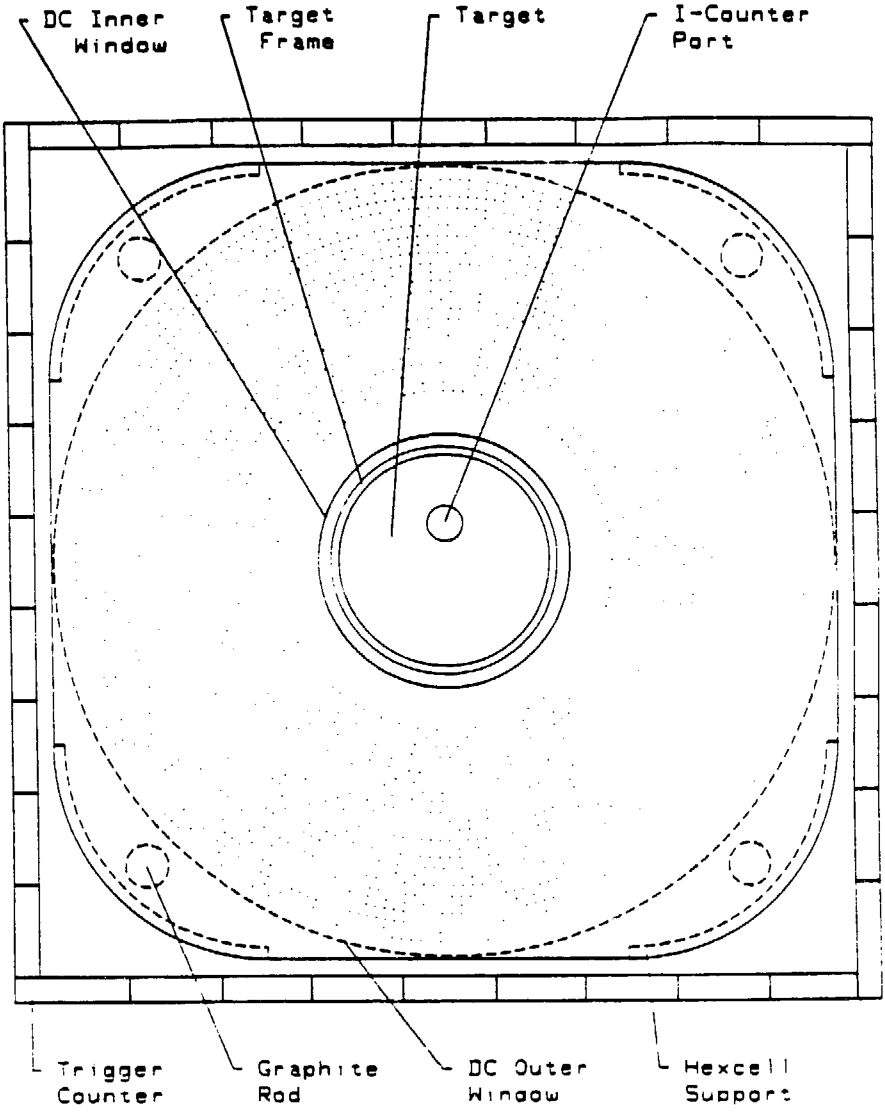

Figure 7: An end view of the Crystal Box detector configuration.

The eight shells of drift chamber wire connections are indicated, showing the rotational offset of unit cells for neighboring shells. This offset reduces the reconstruction ambiguities in the analysis of the charged-particle tracks.

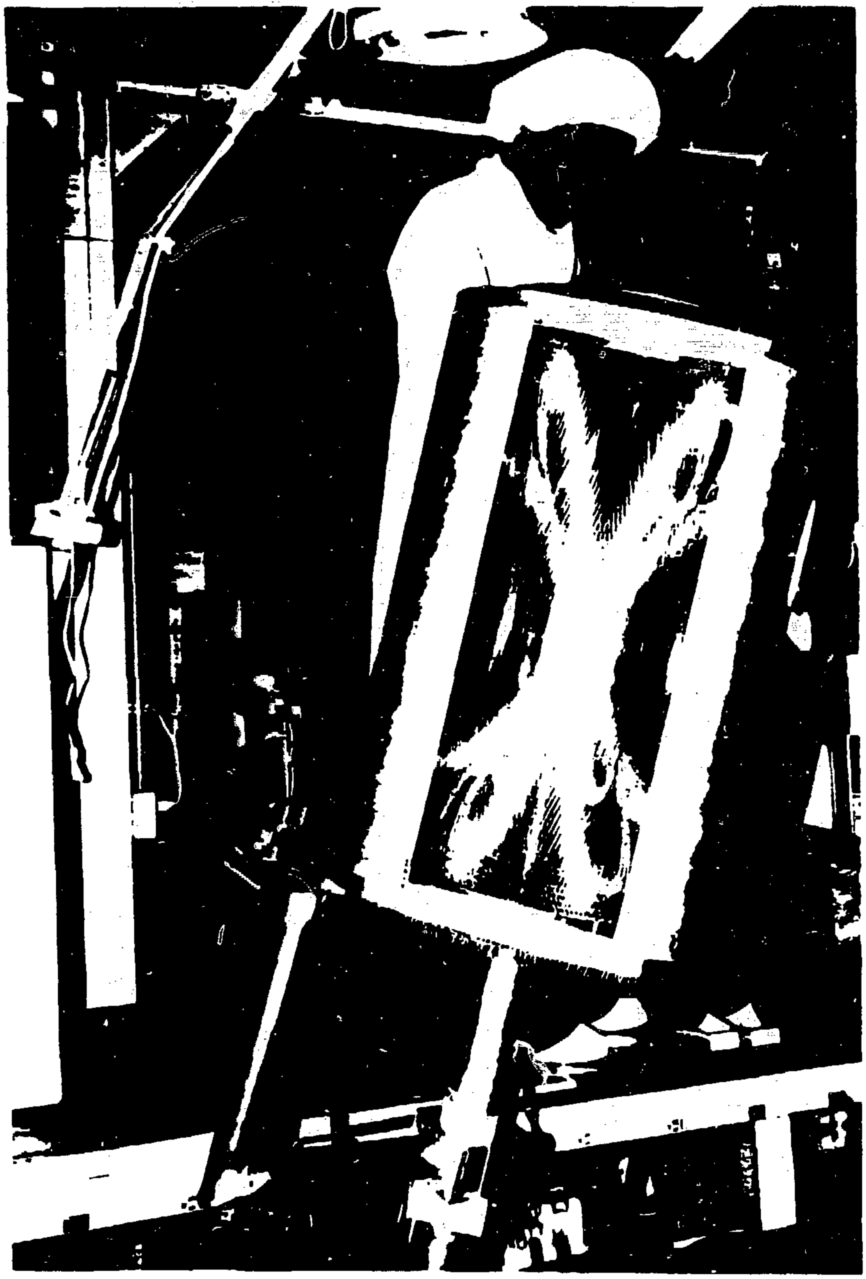

Figure 8: The drift chamber in the clean room during assembly.

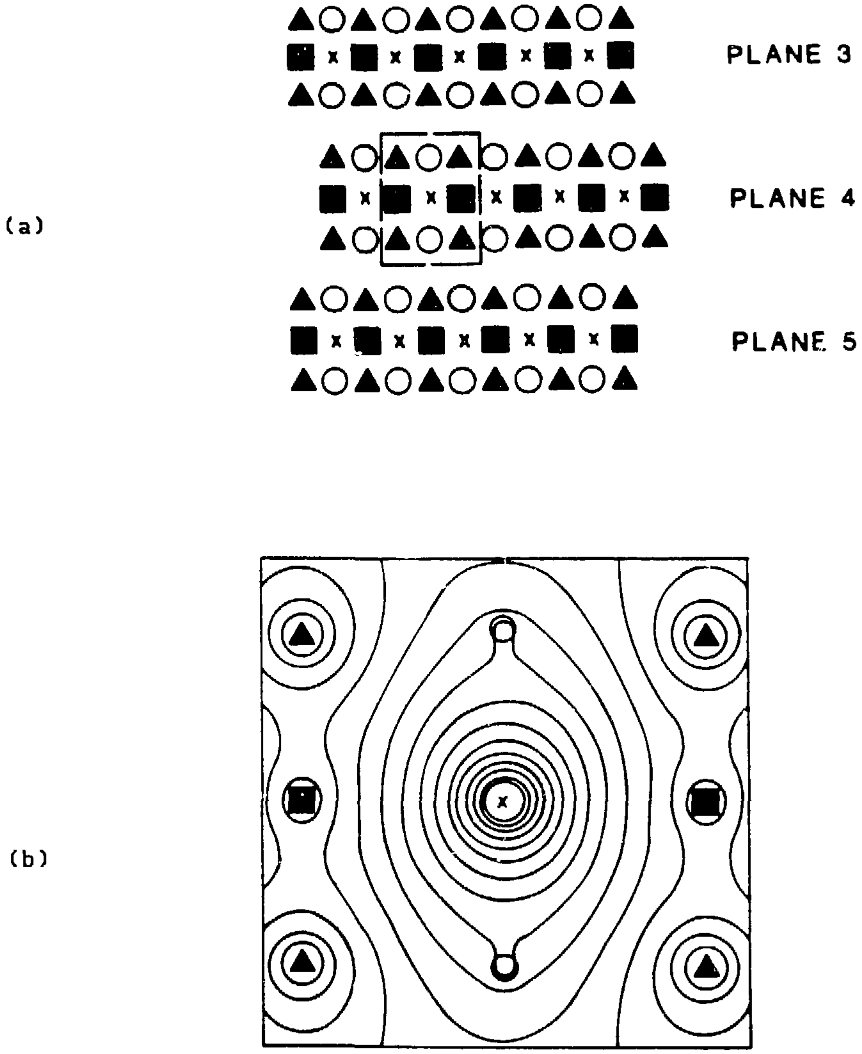

EQUIPOTENTIAL PLOT OF UNIT CELL

Figure 9: Configuration of wires in the drift chamber unit cell.

High voltage wires (circles, triangles, and squares) are 154 µm tungsten, and sense wires (x's) are 25 µm tungsten as shown in (a). The equipotential contours of the drift field are shown in (b).

converters (TDC's), operated in common-stop mode. The data-acquisition system, which controls and receives data from this 4290 system, is described in section 3.5. Each data run was prefaced with a check of the readout timing for each wire. This "autotrim" mode introduced standardized pulses into the readout electronics. The timing of the standard pulses is checked against previous tests to assure the integrity of the readout system. This process is analogous to that of the flasher system (described in section 3.3.5) used to check the NaI(Tℓ) readout system.

3.3.3 Scintillator Hodoscope

Outside the drift chamber and lining the inner walls of the NaI(Tℓ) calorimeter are the plastic scintillator tiles which make up the trigger hodoscope. As shown in Figures 3 and 11, thirty-six counters run longitudinally along the beam, nine adjacent to each side of the Box. The scintillator material, Pilot B, is 1.27 cm thick, 5.7 cm wide, and 44.5 cm long, covering only the central seven crystals in each row. These "trigger counters" provide timing and geometry information about charged particles to the trigger coincidence circuits, as well as fast vetos to the photon detection logic.

The length of the trigger counters was chosen so that trajectories from the target through the ends of these counters would have sufficient length through the NaI(Tℓ) calorimeter to provide good energy containment. Additional charged particle vetoing is provided by the sixteen large (13.3 cm x 23.8 cm x 0.3 cm) "guard counters" also shown in Figure 11. Without these guard counters, the large flux of Michel

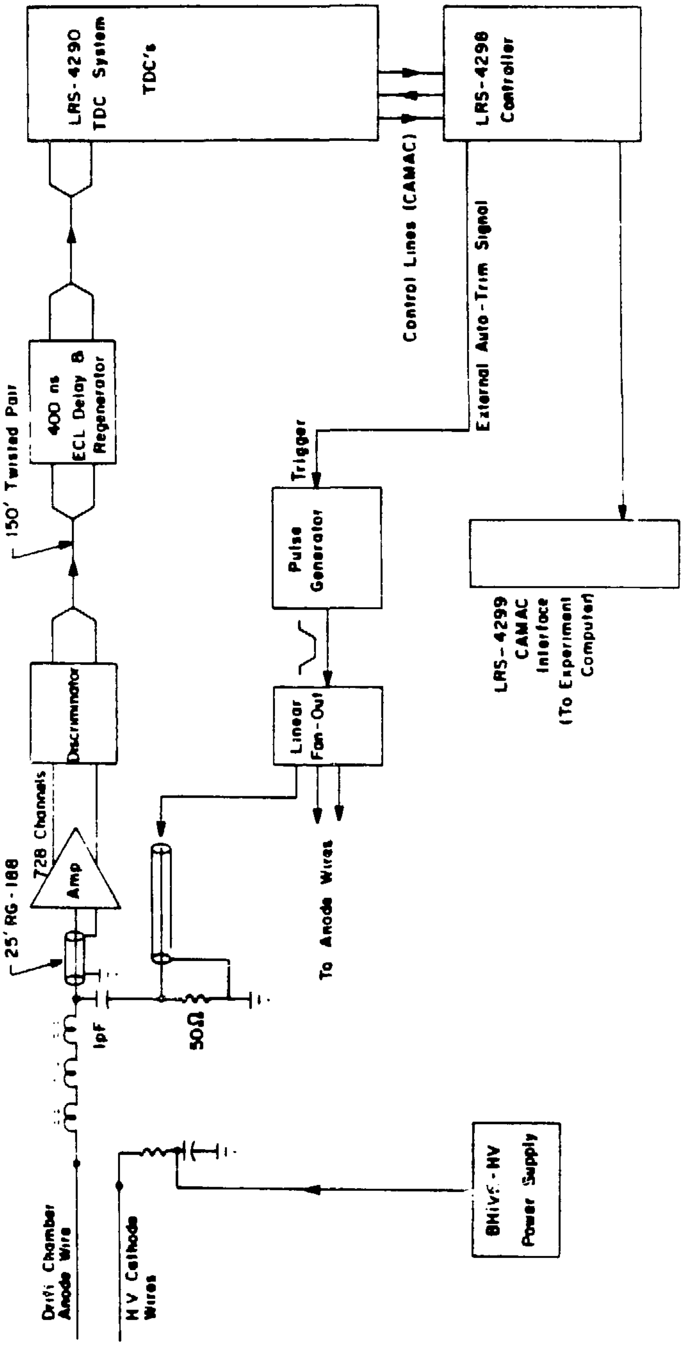

Figure 10: The drift chamber readout electronics chain.

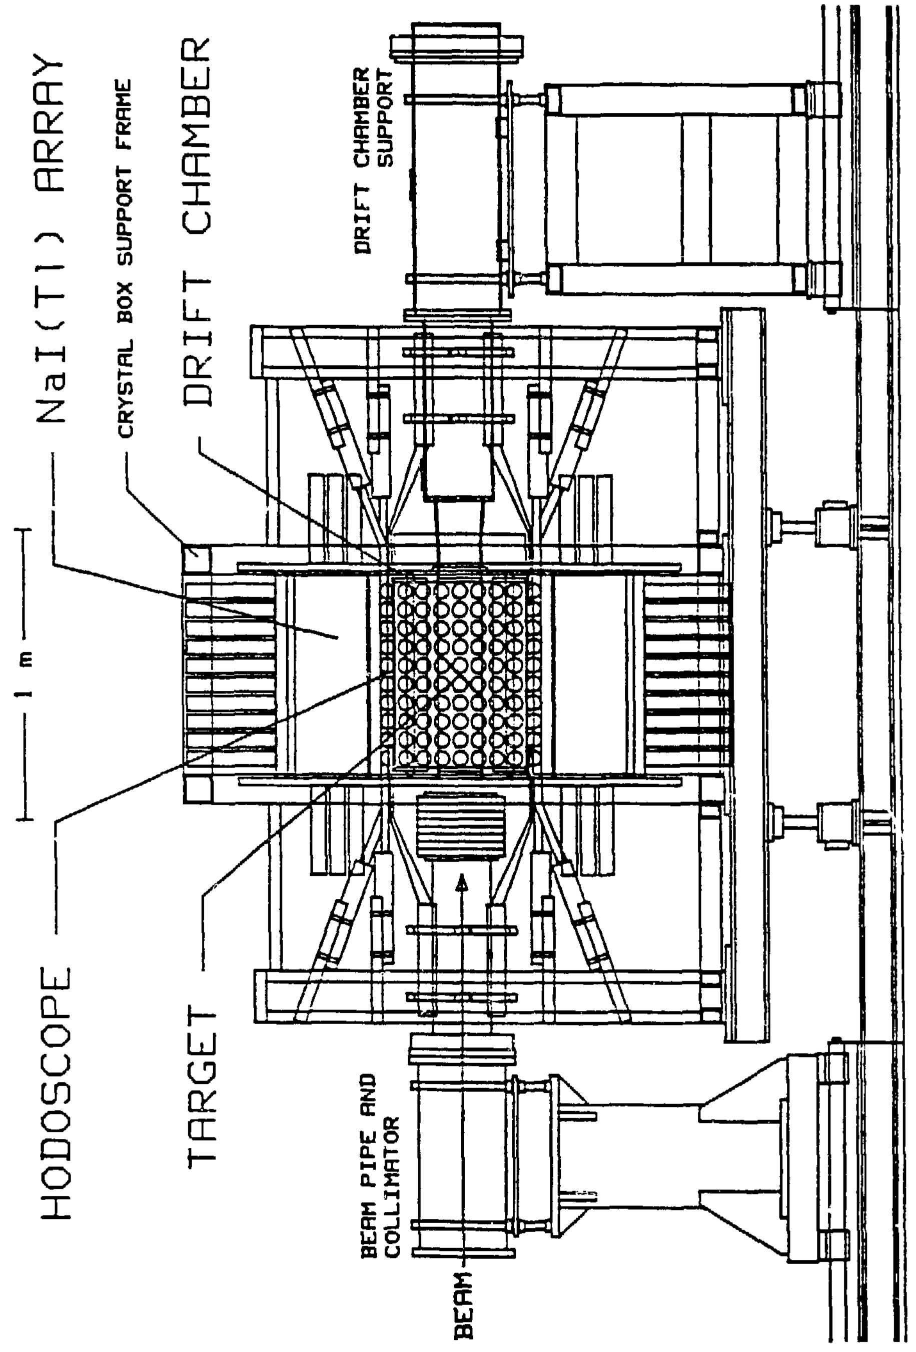

Figure 11: A side view of the complete Crystal Box detector.

The drift chamber mounting is fully retractable to allow substitution of the liquid hydrogen target for calibration runs.

positrons from normal muon decay would generate false photons in the last 1.5 columns of NaI(Tℓ) at the edges of the Box. These counters can also be used in positive logic to illuminate the edge columns of NaI(Tℓ) crystals during calibration and set-up runs.

All the plastic counters are optically coupled by lucite light guides to Amperex 2232B phototubes (PMT's). Such a light guide and PMT is coupled at each end of the trigger counters. The PMT signals are transmitted by RG-58 cable from the experimental area to a counting house, where they are passively split with 25% going to a CAMAC ADC for pulse-height analysis. The remaining 75% is processed by the constant-fraction discriminator (CFD) circuit (shown if Figure 12), which was designed and custom-built at LAMPF.[50] Signals from the CFD's for both ends of a trigger counter go to a specially built "meantimer" circuit which averages the arrival times of the two signals to give an output pulse whose timing is independent of particle impact point on the counter. A schematic of this circuit, also designed and built at LAMPF,[51] is given in Figure 13. The arrival times for particles in the guard counters and the meantimes for trigger counters are digitized by several types of commercial CAMAC TDC's supplied by the LAMPF electronics pool. These TDC's typically have a sensitivity of 100 ps/channel.

3.3.4 NaI(Tℓ) Calorimeter

The Crystal Box calorimeter is the first large solid angle detector to be constructed from regular-shaped polyhedrons of NaI(Tℓ). The earliest large aperture detectors were unmodularized, while the first

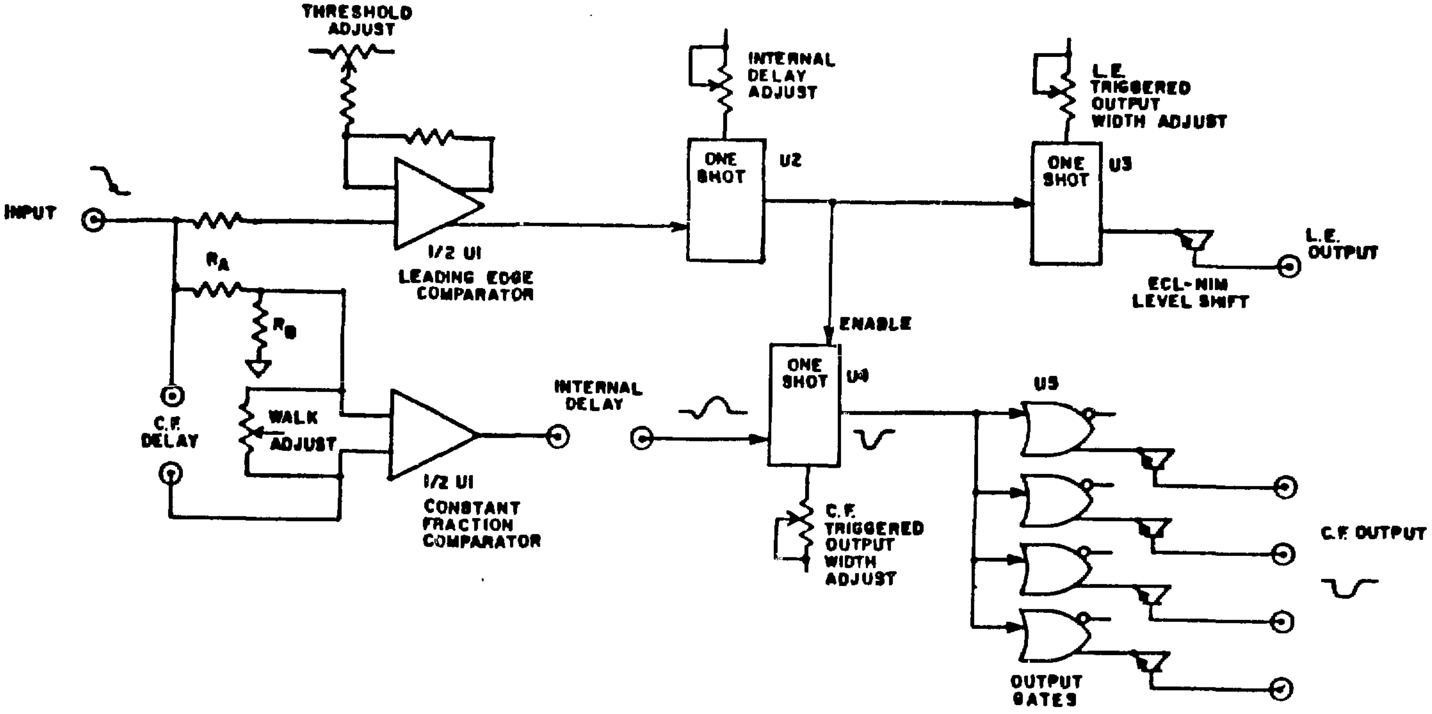

Figure 12: The functional diagram for the LAMPF CFD design.

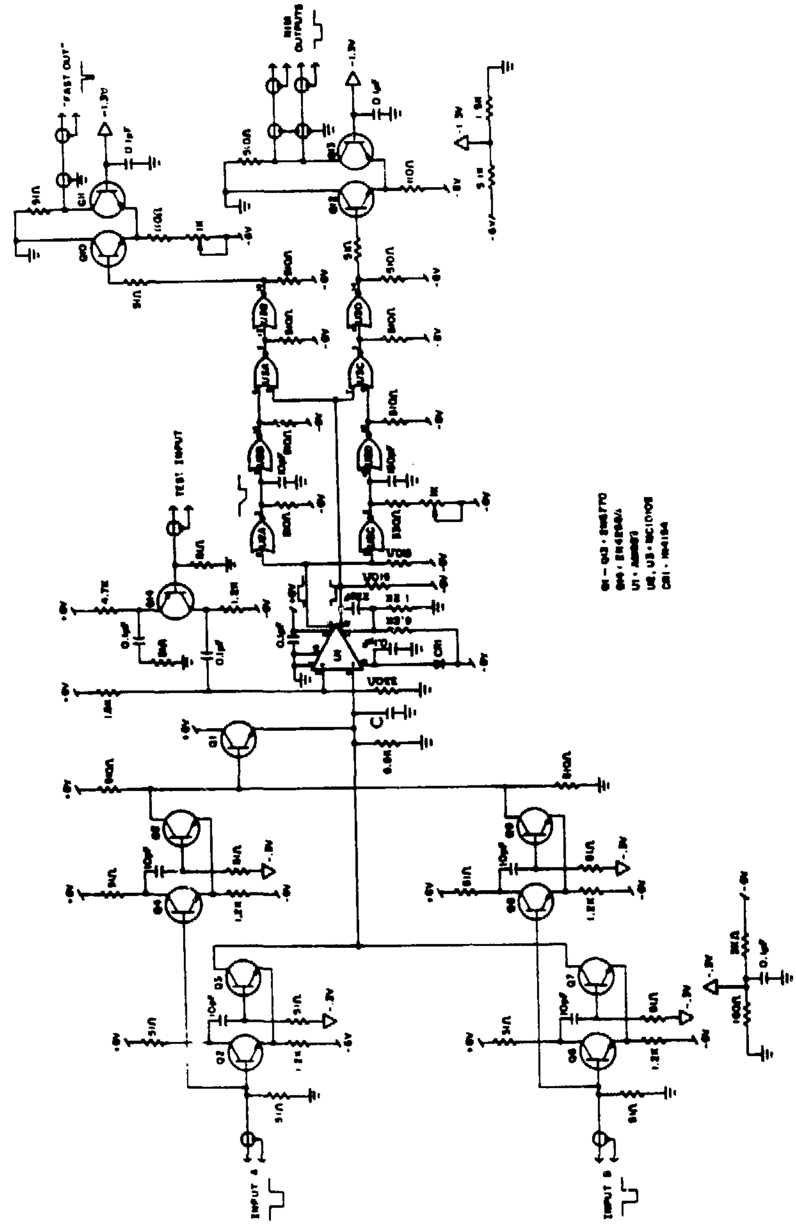

Figure 13: The circuit diagram for the LAMPF meantimer.

modularized versions were assembled from individually-encapsulated modules, either square or hexagonal in cross section. Separate encapsulation led to some degradation in energy resolution due to energy loss fluctuations in the many metallic walls included within the sensitive volume. Later detectors, such as the Crystal Ball,[52] have been constructed of special module shapes, which increases cost of manufacture and precludes the option of reconfiguration for future experiments.

With the expectation of a future reconfiguration for the Crystal Box modules in mind, a rectangular shape was chosen. To limit the potential performance loss, the intervening thicknesses of inert material were minimized and the performance of the individual modules were optimized before insertion into the array. The proposed detector was therefore considered a sufficiently new technology that a prototype 7 x 7 array of 49 modules was assembled.

3.3.4.1 Prototype Design and Construction

The "Array of 49" provided experience in stacking and packaging such regular modules and afforded detector performance tests and electronics development before the full Crystal Box array was available. The depth of each module, 50.8 cm (20"), and the overall transverse acceptance, 43.8 cm x 43.8 cm, were selected so that energy leakage fluctuations do not limit the observed energy resolution if an incident gamma ray enters the detector at any point within the acceptance of the central 3 x 3 group of modules within a reasonable angular range. This permits the exploration of the detector response, including uniformity, over an aperture that includes the boundaries between adjacent modules. The individual module acceptance was chosen to correspond approximately to the transverse size of the electromagnetic shower, and thereby provide isolation between multiple showers and reasonable impact point definition without extreme module density. These modules were rough cut from polycrystalline NaI(Tℓ)[53] extrusions 5″ x 5″ in cross section and up to 50″ in length. The transverse module dimensions were milled to a tolerance of 0.005″, the length to within 0.020″, and all faces to a flatness of 0.005″. Subsequently, the 2.5″ x 2.5″ face adjacent to the optical window was ground to a flatness of 0.001″.

Before assembly, the performance of each module was tested and, if necessary, adjusted in order to meet certain criteria. This test was carried out in a simulated optical environment identical to that of the assembled detector. The crystal module is coupled at one end to a window assembly using a 0.125″ thick transparent elastomer pad. The glass window itself is bonded to a stainless steel cylinder that attaches to the aluminum bulkhead of the detector housing. The phototube couples to the outer face of the window through a 1″ light guide and a second 0.125″ elastomer pad.

Except at the optical window, the module is wrapped with a diffuse millipore paper reflector (2 mil in thickness) which is reinforced by an outer wrapping of aluminized mylar (3 mil in thickness). A polished area of the appropriate diameter is provided on the surface adjacent to the optical window. Elsewhere the surface faces are sanded to maximize the light collection efficiency. The degree of surface sanding is adjusted along the length of the module, usually in no more than two or three discrete steps, to control the uniformity of the light collection efficiency along the module length.

3.3.4.2 Module Performance

The acceptance criteria required the non-uniformity in response to a collimated ^{137}Cs source to be no more than 4% in the region distant 10″ or more from the optical window and no more than 6% in the region distant 2″ or more from the window. No significant variation in the collection efficiency around the major module axis at any longitudinal position was observed for any module. Also measured for each module was the photopeak response to a plane-collimated ^{137}Cs source as it was moved across the 2.5″ x 2.5″ face remote from the optical window. A non-uniformity of less than 2% was required in this area and was easily met by all modules. The resolution in response to a ^{137}Cs source positioned on the major axis approximately 12″ from this face was measured for each module and was typically 12% FWHM.

3.3.4.3 Phototube Selection

The principal requirement of a phototube for use with the Array of 49 modules was its ability to provide the highest possible linear output current for gamma-rays incident in the range 1 MeV to several GeV. Secondary requirements are reasonable gain and moderate price. The output current requirement restricted the search to linear focussed phototubes. Five readily-available candidate phototubes with 2″ photocathodes were tested, and the Amperex 2202B was eventually selected as the optimum phototube for use with the Array of 49.

3.3.4.4 Prototype Test Results

The performance of the Array of 49 was explored in test beams at both the Stanford Linear Accelerator Center (SLAC) and LAMPF. Positrons in the range 750 MeV to 18 GeV were used at SLAC and in the range 20 to 250 MeV at LAMPF. As usual, for NaI(Tℓ) detectors, the energy resolution observed over the range 20 MeV to 18 GeV followed a power law dependence on energy, with an index on the order of -0.25. For a well-collimated beam of 10 GeV positrons incident normally at the boundary between two modules, there was a loss of pulse height of about 1.5% and a corresponding change in resolution (from 1.6% to about 1.9%). The effect is very local, concentrated within ± 2 mm of the boundary, and disappears with non-normal angles of incidence as small as 5°. This boundary effect performance represents a substantial improvement over that for large arrays of separately encapsulated modules.[5] The position resolution of $\sigma_x = 1.5$ cm for positron showers was found to be in agreement with expectations based on a Monte Carlo calculation.

3.3.4.5 Crystal Box

The Crystal Box calorimeter contains 90 optically isolated modules of polycrystalline NaI(Tℓ) in transverse orientation in each of four quadrants, and nine modules in longitudinal orientation in each of four corner arrays. Face crystals are individually optically coupled to Amperex 2232B PMT's at the outside of the Box. Corner crystals have PMT's attached at each end, thus allowing the longitudinal coordinate of incident particles to be computed based on the energy sharing.

This calorimeter incorporates several modifications of the prototype design to improve the timing resolution. First, in part due to considerations of cost, the face modules were shortened from 50.8 cm (20") to 30.5 cm (12"). This adjustment improves the risetime characteristics of the scintillation pulses with an acceptable compromise in the energy leakage from the back of the detector for e^+'s and γ's in the 10-100 MeV range. Second, the modules are optically polished on all faces, further reducing the time dispersion of scintillation light from that of diffusely reflecting sanded surfaces. Third, the 2232B PMT's are intrinsically faster than the 2202's used on the Array of 49. These tubes were instrumented with improved bases designed to accommodate and preserve their superior timing response and to supply the increased current requirements.

The signal chain for the NaI(Tℓ) is shown schematically in Figure 14. The raw signals are amplified and distributed by clip-and-split amplifier circuits specially designed and fabricated at LAMPF. The PMT high voltages can be set uniformly at the maximum rating of the bases for best timing performance while intercalibration, for the trigger energy thresholds, is accomplished through the adjustable gains of the amplifiers. The NaI(Tℓ) signals are split so that 75% goes to the timing arm, 12% to form hardware energy sums for triggering purposes, and 13% to the precision energy-measurement electronics.

The timing signals are discriminated with CFD's similar to those described for the plastic signals but modified to accommodate the slower rise-time characteristics of the NaI(Tℓ) signals. Meantimers average the timing signals from the two ends of the corner crystals

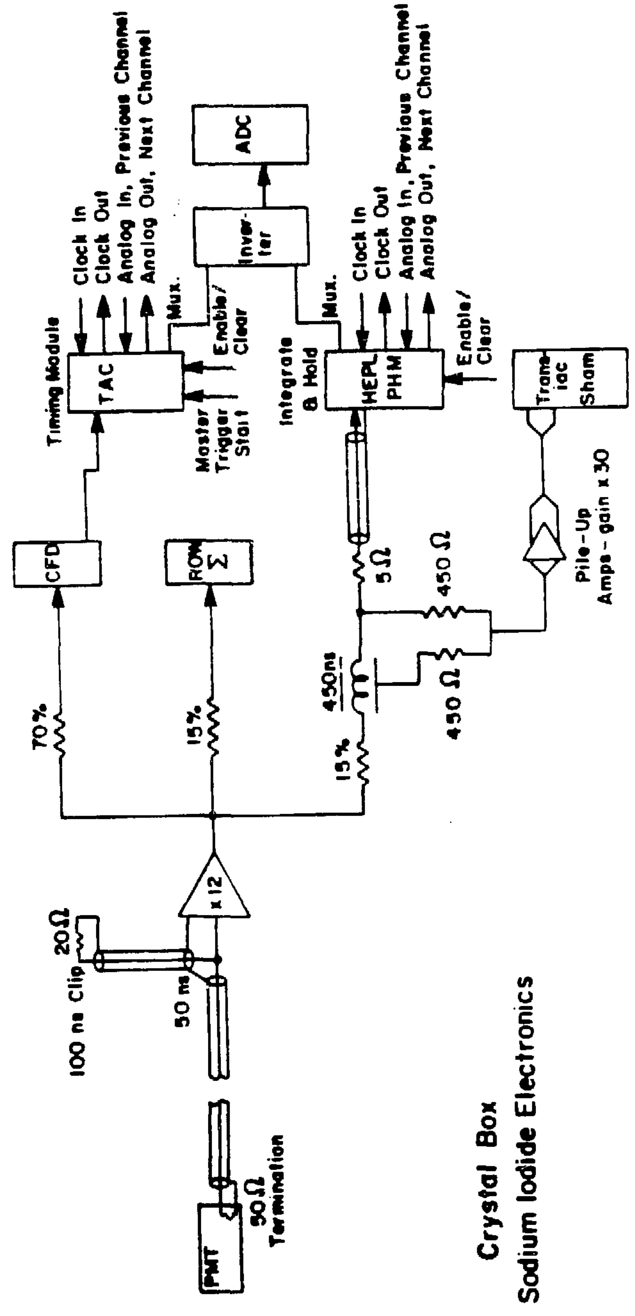

Figure 14: The signal chain for NaI(Tℓ) energy and timing information.

for the trigger, but parallel CFD outputs for the individual PMT's are separately analyzed by the specially designed Time-to-Amplitude Converters (TAC's) provided by Stanford.[55] Since the TAC's operate in a common-start mode, the stop signals from the CFD's must be delayed by the approximately 300 ns to allow time for the trigger logic to generate the start signals. This delay is accomplished by passing the signals through Spiradel transmission lines[56], packaged in 36-channel units, as described by Sandoval, et al..[57]

The pulse height (energy) signals are passively clipped in the amplifier circuits (see Figure 15) from an exponential decay time of 250 ns to a roughly square pulse of about 200 ns width. This pulse shaping allows the use of shorter integration times in the custom-built Pulse Height Integration Modules (PHIM's) (also provided by Stanford[55]), thus reducing susceptibility to signal pile-up. The clipped signals are delayed with Spiradel transmission lines to allow the generation of the required gating signals by the trigger. (It was found that the more expensive Spiradel delay lines were preferable to the lumped constant devices (described in Reference 57) for preserving the clipped trailing edge of the pulse-height signal without unacceptable "ringing".) The same package splits off 10% of the pulse-height signal for analysis by an auxiliary ADC system. This CAMAC based system, produced commercially by Transiac Corporation, will measure out-of-time pile-up directly when fully operational. For the data taking discussed here, this system was not sufficiently developed to provide useful information.

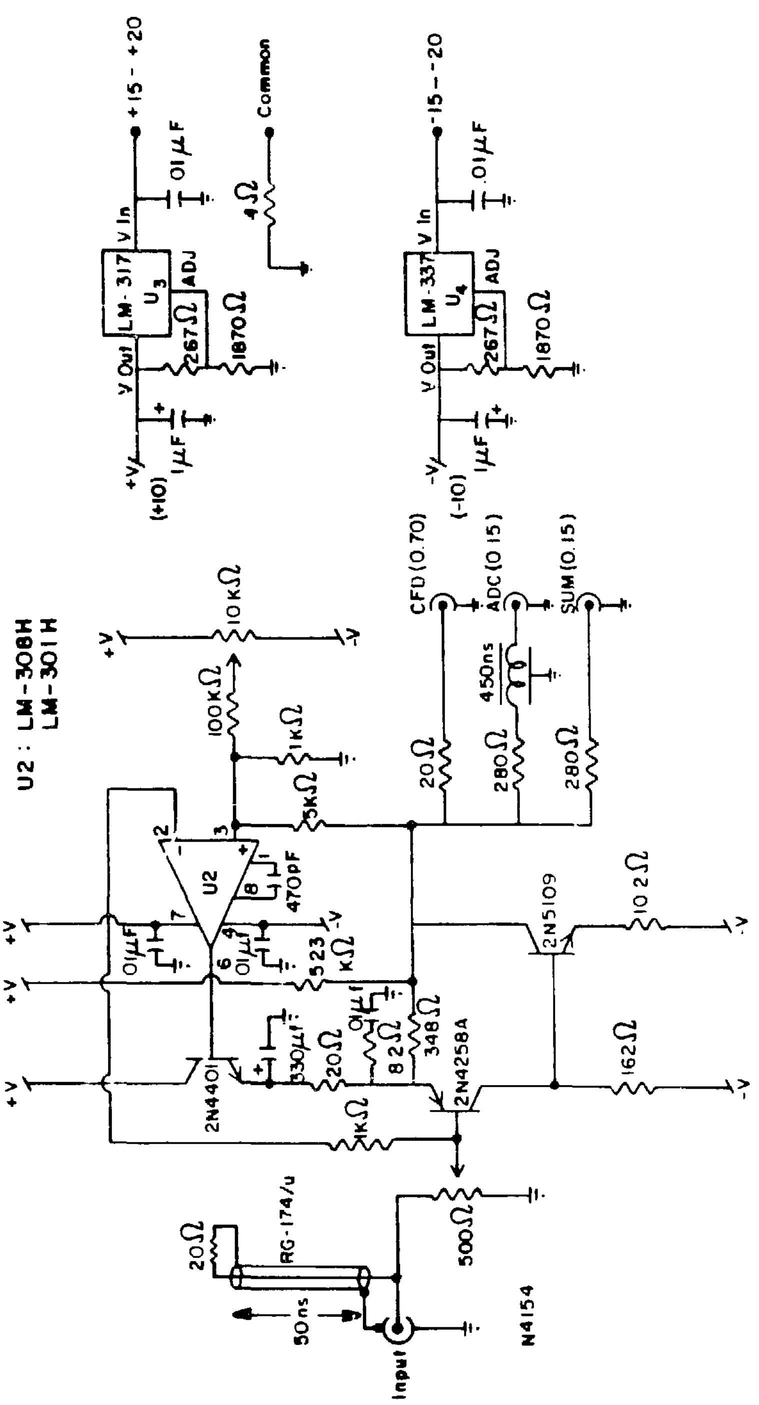

Figure 15: The circuit diagram of the LAMPF clip/split amplifier.

The TAC's and PHIM's store information about the time of arrival and energy deposited in each NaI(Tℓ) crystal in the form of proportional charge on an integrating capacitor. This information is serially digitized by one Tracor-Northern TN1213A Analog-to-Digital Converter (ADC) per quadrant. The multiplexing of ADC inputs and preprocessing of digitized data is performed by a Digital Equipment Corporation (DEC) LSI 11/23 microcomputer for each quadrant.

The constituents of each LSI are detailed in Figure 16. Communication with the ADC is through a custom interface implemented as a circuit board installed directly in the Q-bus backplane of the LSI. Connection to a Bi-Ra model 1150 Auxiliary Control Port (ACP) allows the LSI to initiate CAMAC functions. These functions include trigger control (see Figure 17), collection of data residing in CAMAC modules, and implementation of a handshake protocol with the main data acquisition computer, a DEC PDP 11/44. The completion of this protocol involves a Direct Memory Access (DMA) data transfer from the LSI to the PDP by means of the Bi-Ra model 1151 Dataway Access Port (DAP). Figure 18 gives a schematic representation of the NaI(Tℓ) data collection scheme. The ACP and DAP together comprise the Auxiliary Crate Controller (ACC) which is programmed to respond to specific CAMAC Look-At-Me (LAM) signals (generated by trigger inputs to Jorway model 41 input registers) by interrupting the LSI to start the indicated readout procedure. One set of ACP, DAP, and input register for each LSI resides in Crate 1 of the 6 on the CAMAC branch highway described in section 3.[illegible].2 below.

A flow chart for the LSI data acquisition program is given in Figure 19(a). Each of the four LSI's implements a customized code, named IMP1

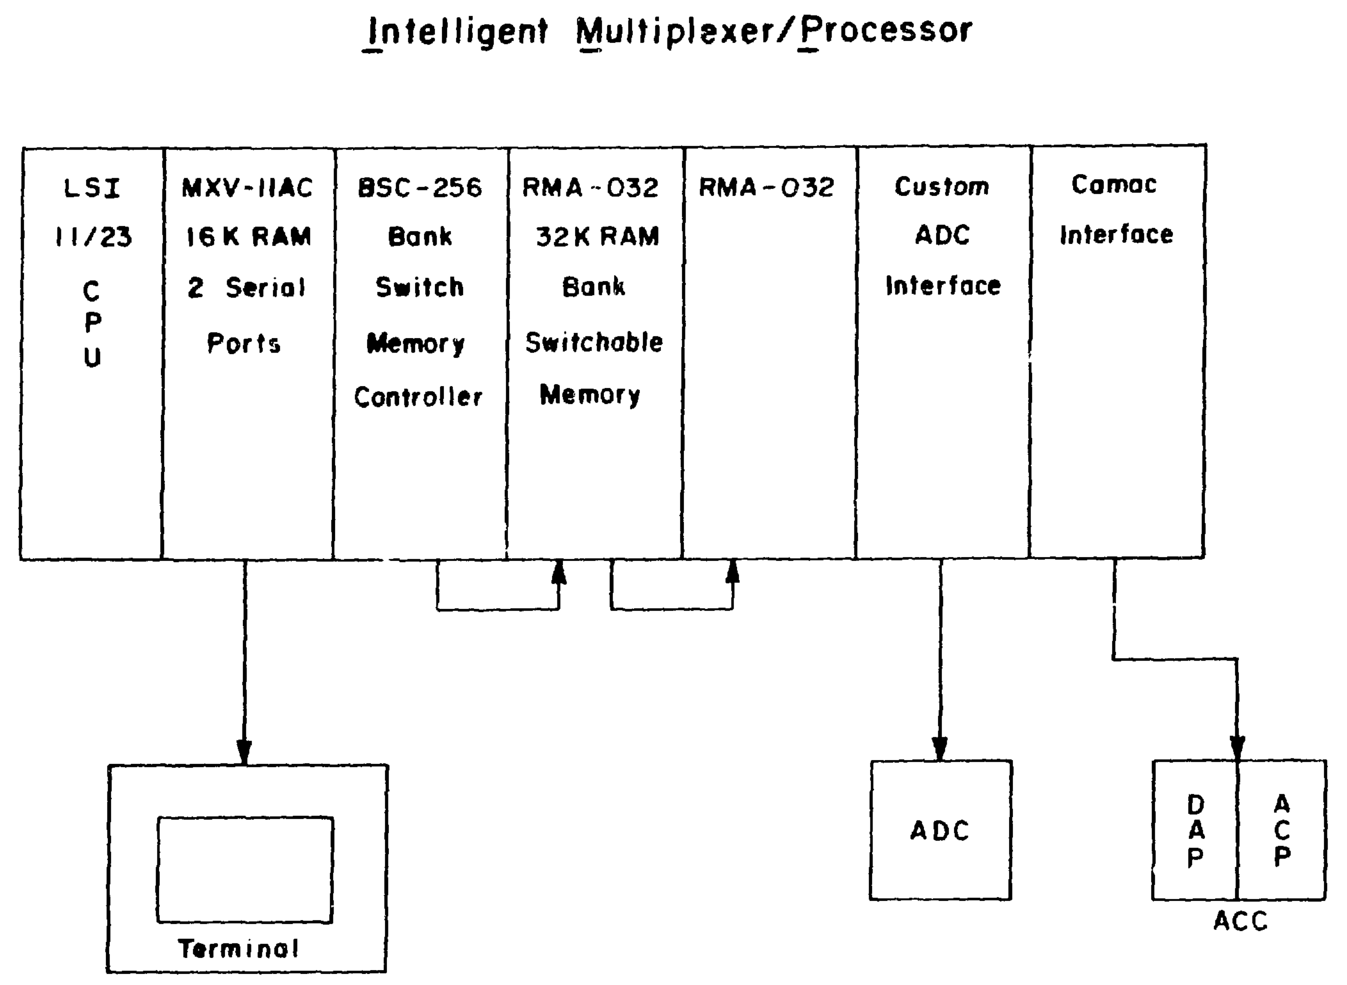

Figure 16: The components of the LSI 11/23 parallel processors. These processors read and compress the NaI(Tℓ) data.

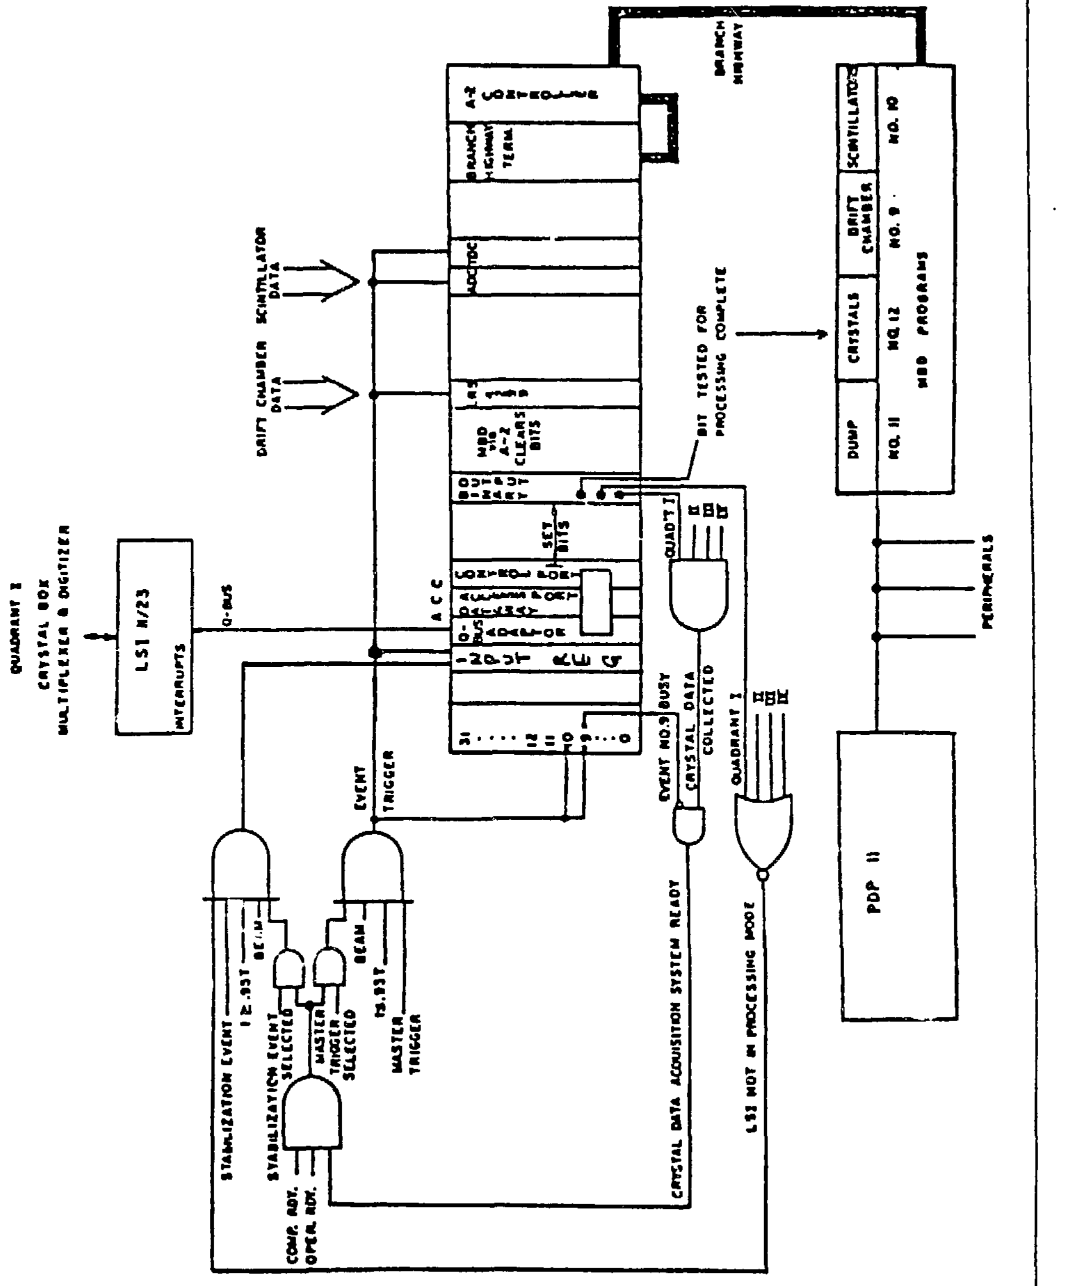

Figure 17: A diagram illustrating the interaction of the LSI's with the trigger control.

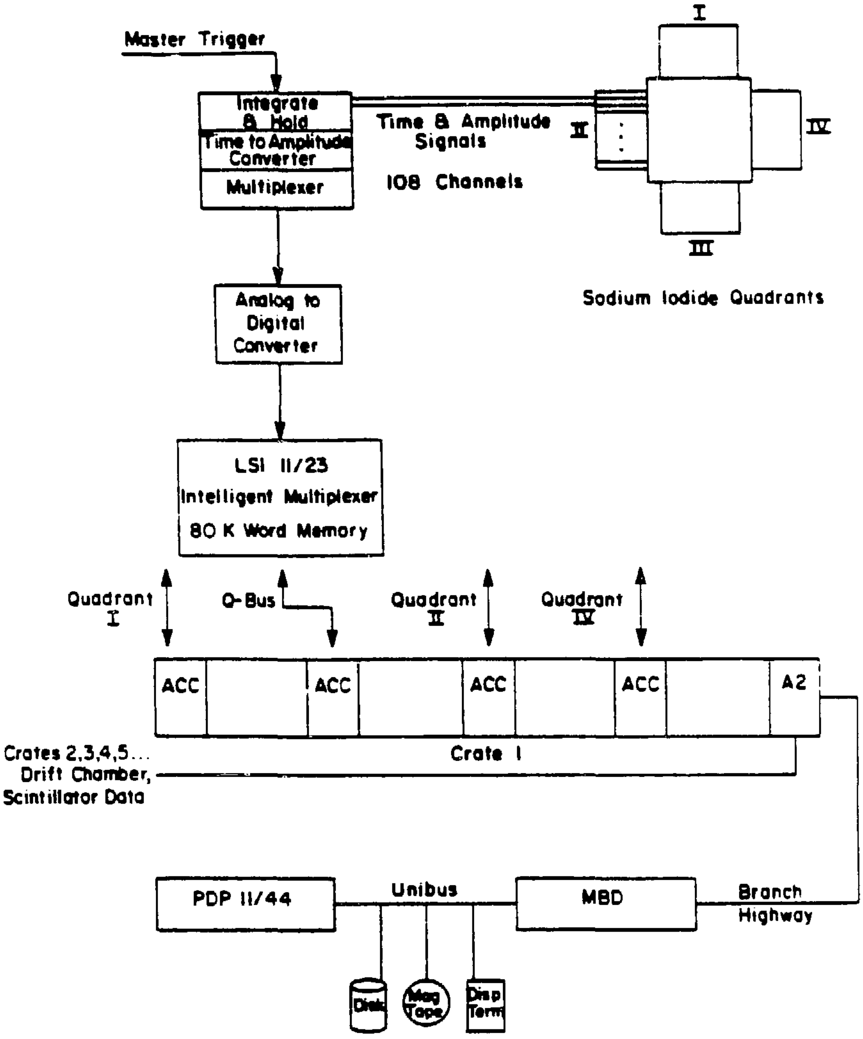

Figure 18: A schematic view of the CAMAC branch highway system.

The communication between the LSI's and the main PDP computer is emphasized.

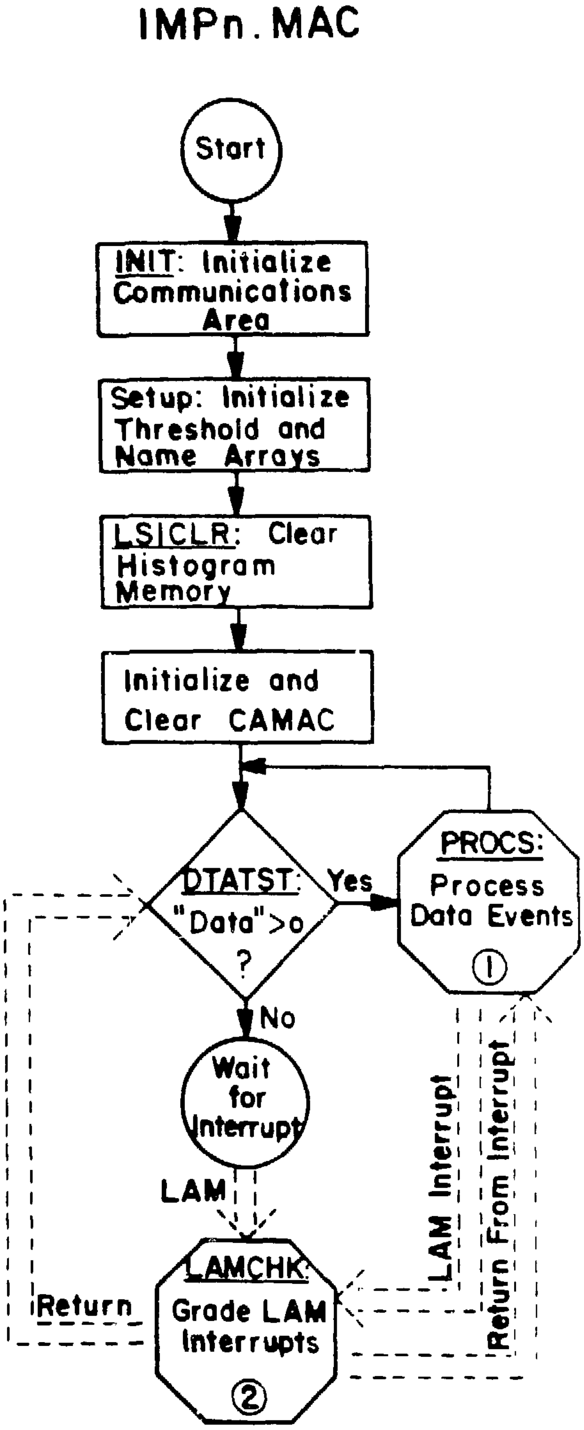

Figure 19: (a) Flow chart of the LSI's data-acquisition code.

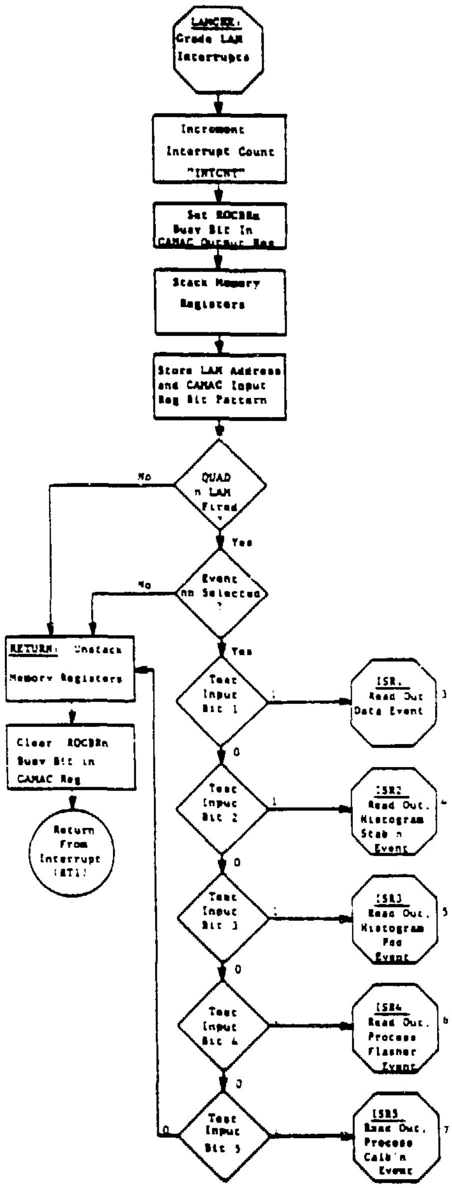

Figure 19: (b) Routine LAMCHK.

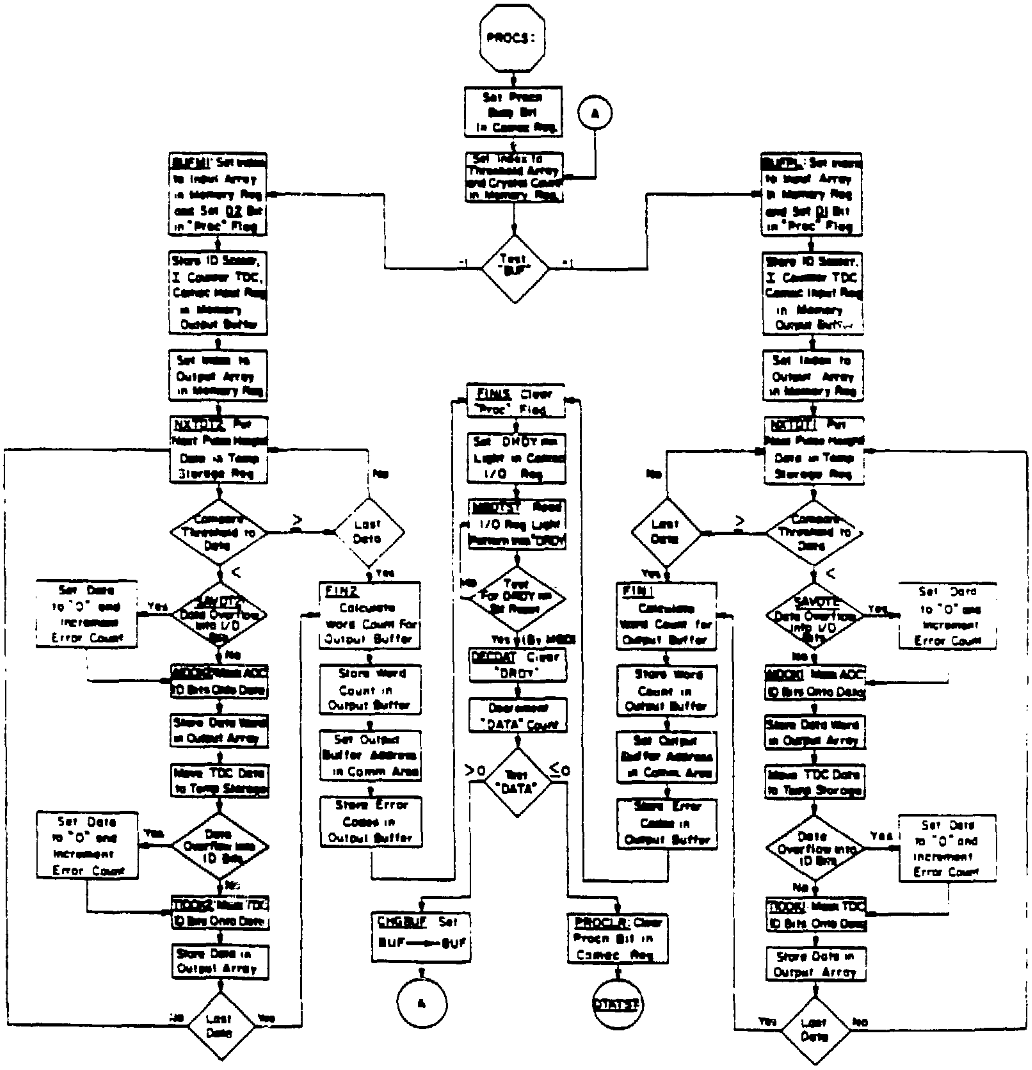

Figure 19: (c) Routine PROCS.

to IMP4, which differ primarily in the CAMAC channels through which they communicate. This main control program is interrupt driven by any of five possible LAM's from the input register. The particular LAM received directs the routine LAMCHK, shown in Figure 19(b), to select the appropriate one of the interrupt service routines, ISR1 through ISR5. Routine ISR1 reads out the acquisition electronics for each data event and resets the system for the next trigger before giving control to the PROCS routine (Figure 19(c)) for background processing of the event. The background processing has the capability to buffer one event of data while performing the sparse-data scan on a previous one. Given the average event rate of the detector, and the time required to process an event, this buffering capability is necessary, to reduce to a tolerable level, the experimental dead time that would occur if events were serially processed. The routines ISR2, ISR3, and ISR4, also charted in Figure 19, provide for the readout and histogramming of stabilization, pedestal, and flasher events respectively. These readout procedures do not communicate directly with the main data acquisition computer. Rather, they collect the individual trigger information into histograms to be analyzed at the end of the data run by the main computer. The pedestal trigger, for instance, was activated immediately before each data run to check the stability of the readout electronics. The final routine, ISR5, provided strictly serial processing of events with different sparse-data-scan thresholds, if desired.

3.3.5 Flasher System

The flasher system, designed and built by collaborators from the University of Chicago, introduces standardized light pulses directly into the light guides of the NaI(Tℓ) PMT's. This capability allows the monitoring of gain drifts in the entire readout electronics chain. The only component of gain drift not directly accessed by the flasher system is a change in the intrinsic scintillation yield of a NaI(Tℓ) crystal. As long as the crystals are maintained in a stable physical environment, this source of gain drift is negligible.

3.3.5.1 Light Pulse Generator

A schematic drawing of the flasher system is given in Figure 20. A Xenon flashtube (EG&G model FX-280) produces an arc of light which is viewed by many hemispherically distributed ports containing fiber optic cables. The light is distributed to the individual crystal light guides by 12 feet of DuPont Crofon 1040, a single fiber plastic cable of 0.040″ diameter, sheathed in a polyethylene jacket giving a total diameter of 0.087″. The intensity loss of this fiber is quoted as 7 - 8% per foot for green light. It accepts incident light in a 64° cone and the output light is forward peaked at less than 20°.

To permit removal of PMT assemblies for replacement or repair, Amphenol fiber optic connectors at the back of the phototube base join the 12′ cable to a 14″ length which is integral to the PMT assembly. This 14″ length injects the light pulse into a 30° slot in the light guide looking away from the phototube. Thus, the integrity of all optical interfaces is required to produce stable reflected intensities at the PMT photocathode.

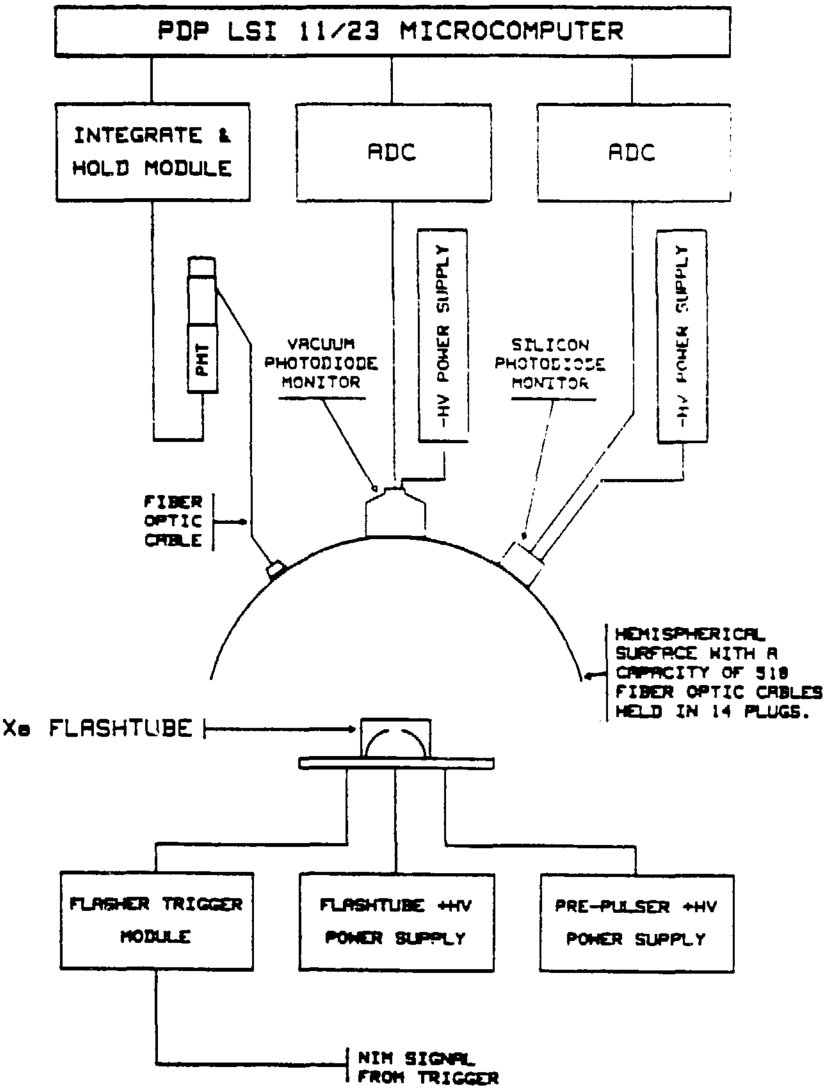

Figure 20: Diagram of the Xe flashtube light pulse injection system.

This system is used to monitor calibration drifts in the NaI(Tℓ) readout electronics.

3.3.5.2 Monitor Electronics

Two different photodiode monitors also view the flashtube directly. One is a vacuum photodiode, an ITT model FW114 biplanar type. The other is an EG&G model SGD-444 solid state silicon photodiode. These monitors provide the capability to compensate for fluctuations in flashtube intensity which is important for the calibration/stabilization procedure discussed in Chapter IV.

3.3.6 The Hydrogen Target

The LH_2 target consisted of a cryostatically cooled vessel, along with its associated control electronics and refrigeration apparatus, which was specially designed and constructed for the Crystal Box experiment, by the MP-7 group at LAMPF. The LH_2 container was cylindrical in shape with hemispherical end-caps. Its overall length was 5 cm and the cylinder radius was 1.9 cm. This target and its attendant apparatus replaced the drift chamber at the center of the Box for the last two days of the January 1984 run. The LH_2 target provided the protons for the two observed calibration reactions. The Panofsky reaction, $\pi^- p \rightarrow n\gamma$ produced "1-γ" events with a γ-line at 129.4 MeV. The π^0's from the single charge exchange reaction, $\pi^- p \rightarrow \pi^0 n$, produced "2-$\gamma$" events. The pion data-taking procedure is described fully in section 4.3 below.

3.4 THE TRIGGER

Given an average of 3×10^5 μ^+/s stopping within the apparatus, it is obvious that not all decays can be recorded. A method is needed to detect the events which are likely to be candidates for neutrinoless muon decay, and trigger the readout of data from the detector. In fact, the data acquisition system was simultaneously sensitive to a separate trigger for each of the three decay modes, $\mu \to e\gamma$, $\mu \to e\gamma\gamma$, and $\mu \to 3e$. Only the $e\gamma$ trigger will be discussed here.

For a two-body decay from rest, the decay particles must be emitted back-to-back with equal momenta. Given the small mass of the electron, the observable signature for $\mu \to e\gamma$ from a stopped muon, is a back-to-back, time-coincident electron-and-photon pair, with (experimentally indistinguishable) energies of 52.8 MeV. The quantities available for construction of a trigger logic are then the geometry, energy, and relative timing of hits in the detector.

3.4.1 Geometry

For the Crystal Box detector, the geometry for $\mu \to e\gamma$ implies that the electron and photon should be observed in opposite "quadrants". After an examination of the expected flux of photons due to background processes, it was decided that no further segmentation of the detector, beyond that of quadrants, was necessary to limit the $\underline{e} \bullet \underline{\gamma}$ triggers to an acceptable rate of a few per second. Therefore, the logic of the $\underline{e} \bullet \underline{\gamma}$ trigger treats each quadrant as a single "particle". A "quadrant" formally consists of:

1. one of the calorimeter face arrays

2. the three crystals on the inward-pointing diagonal and the nearest off-diagonal crystals from each of the two corner arrays adjacent to the face array
3. the nine trigger counters and 4 guard counters in front of the calorimeter face

(Note that the diagonal corner-crystals are each associated with two quadrants.)

3.4.2 Energy

The circuit used to form the quadrant energy sum is shown in Figure 21(a). This sum includes signals only from the NaI(Tℓ) crystals; no energy information from the plastic scintillators is used in the trigger. The energy sum of each individual NaI(Tℓ) row was used only in the 3e trigger. The final sum signal was discriminated at a nominal level of 40 MeV for the quadrant, as shown in Figure 22. (The maximum energy requirement of 65 MeV was not imposed during the data taking discussed here.) The output of the discriminator is designated as the $\underline{SQ}_i$ signal, where i is the number of the corresponding quadrant.

3.4.3 Timing

Figure 21(b) shows the functional diagram of the 90-fold logic fan-in used to form a portion of the quadrant-OR logic signal from the individual NaI(Tℓ) crystal CFD outputs. This fan-in circuit was designed and fabricated for this experiment by collaborators at LAMPF. In the final trigger configuration, the upstream and downstream columns of each NaI(Tℓ) row were removed from this fan-in to reduce the rate of spurious

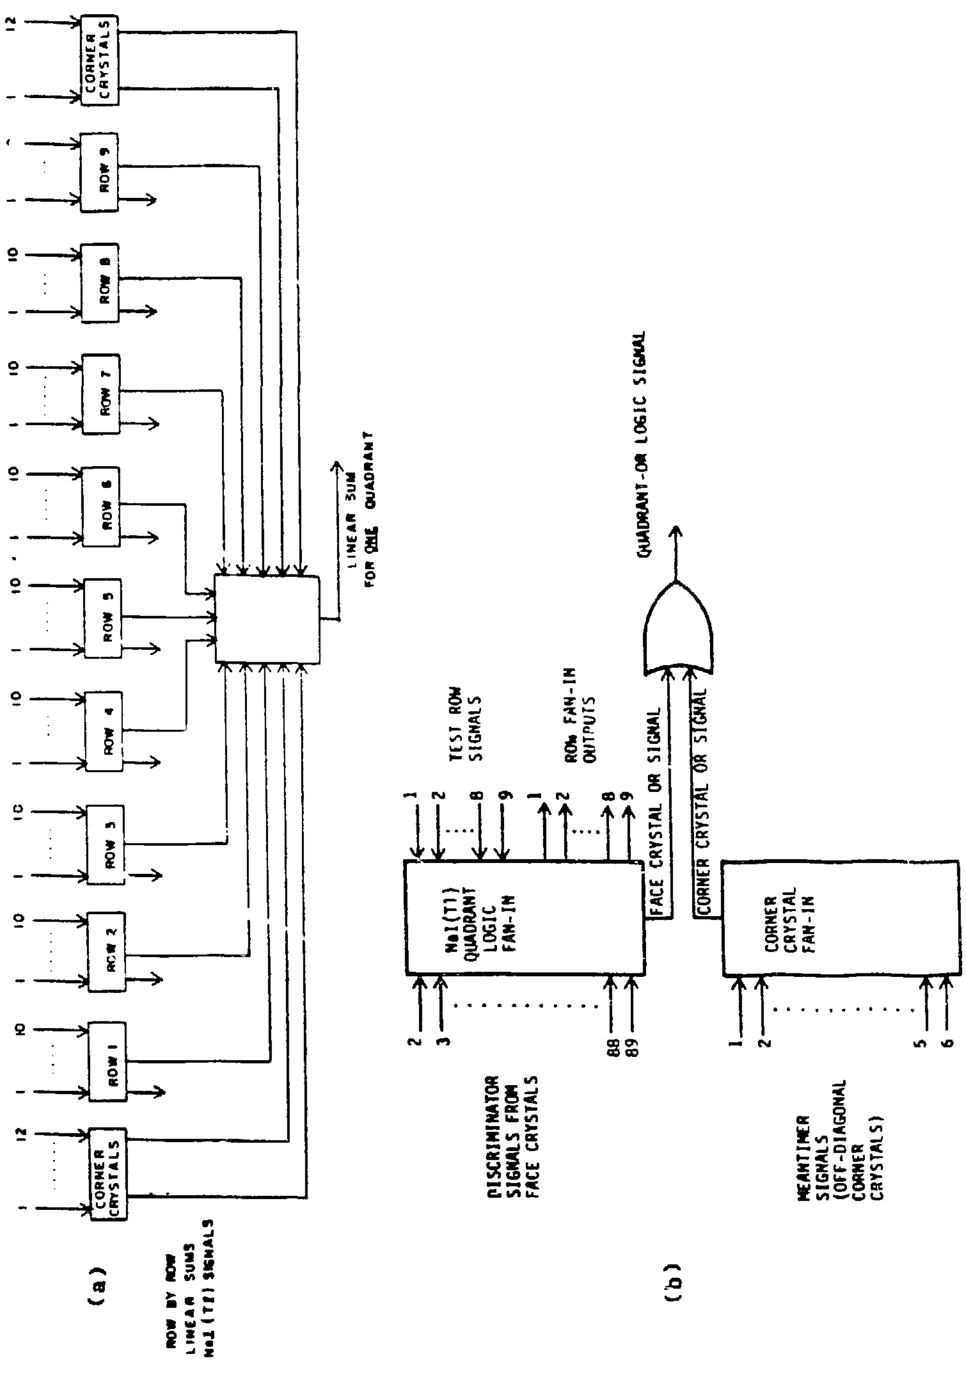

Figure 21: Functional diagrams of the quadrant energy sum circuit (a), and the quadrant-OR logic fan-in (b).

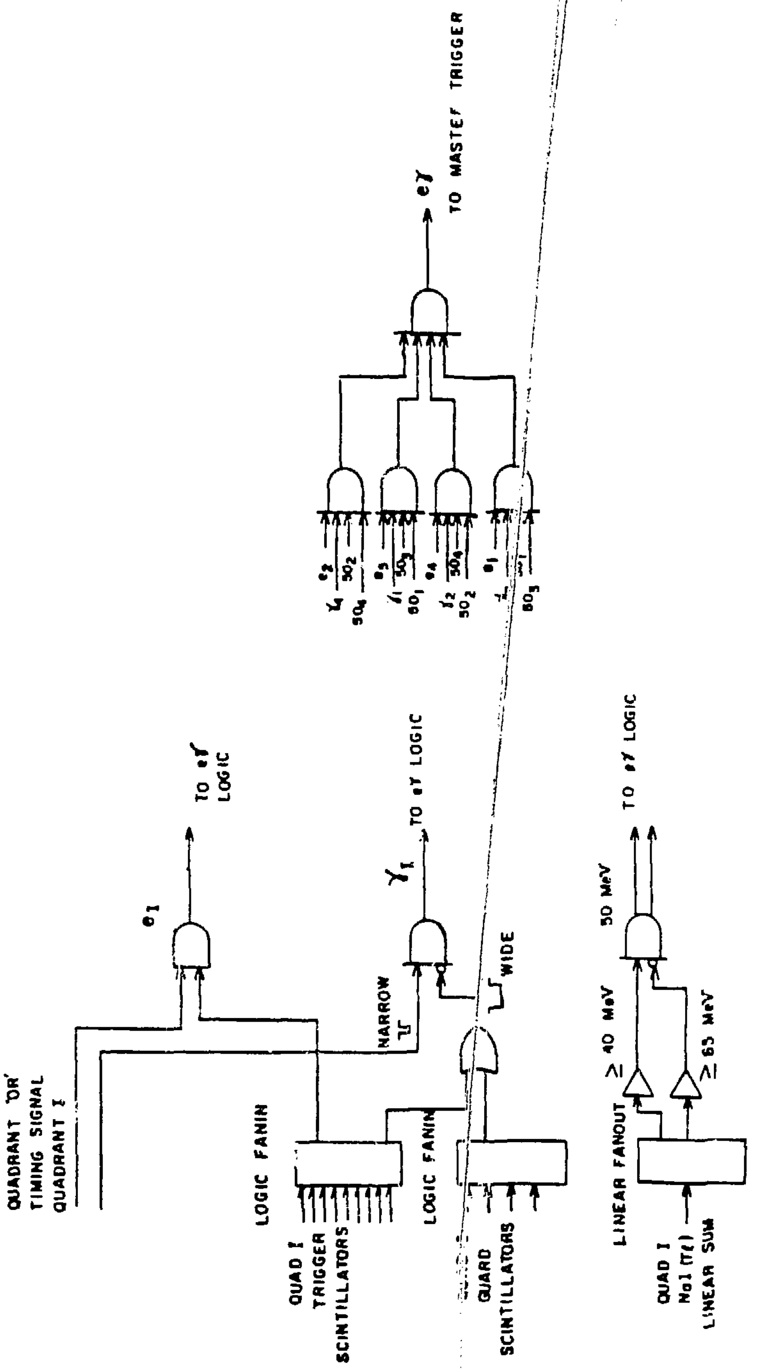

Figure 22: The logic diagram for the eγ trigger.

photons produced by hits at the extreme edges of the Box. The output of the 90-fold fan-in module is later OR-ed with the meantimer signals from the off-diagonal corner crystals in the quadrant to form the quadrant-OR logic signal. The trigger counter CFD signals for each quadrant are logically OR-ed in LeCroy LRS-429 fan-ins, as shown in Figure 22. Guard-counter CFD signals are similarly, but separately, combined.

3.4.4 Trigger Coincidence

For each quadrant i, an electron, $\underline{e}_i$, is defined as the coincidence of the NaI(Tℓ) quadrant-OR signal with that for the trigger counters. A photon, $\underline{\gamma}_i$, is generated by the NaI(Tℓ) quadrant-OR signal, which may be vetoed by both the trigger-counters and the guard counters for that quadrant. An additional photon veto, beyond that shown in Figure 22, is imposed by the nearest trigger counters in the two quadrants adjacent to the photon quadrant. The $\underline{e} \bullet \underline{\gamma}$ trigger is the OR of the four possible coincidences of the form:

$$\underline{e}_i \bullet \underline{50}_i \bullet \underline{\gamma}_j \bullet \underline{50}_j$$

where i and j correspond to opposite quadrants.

3.4.5 Backgrounds

Most of the triggers generated were from random coincidences of Michel positrons and uncorrelated "photons". The Michel spectrum (see section 5.2.2.1) has a large component above 40 MeV. Therefore, any beam-related source of high-energy neutral particles will adversely affect the trigger rate and the experimental background for $\mu \to e\gamma$. Possible sources of "photons" include:

1. $\mu^+ \rightarrow e^+ \underline{\gamma} \nu_e \bar{\nu}_\mu$ (internal bremsstrahlung)

2. $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$

   $e^+ N \rightarrow e^+ \underline{\gamma}$ (external bremsstrahlung)

3. $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$

   $e^+ e^- \rightarrow \underline{\gamma\gamma}$ (positron annihilation)

4. accelerator related neutrons

5. cosmic rays

The internal bremsstrahlung reaction can also produce a time-coincident, or "prompt" background, when the neutrinos are emitted with low energies. Fortunately, the probability for this configuration is small, but the number of such decays observed within the acceptance of the apparatus is an important check of the normalization of the experiment.

3.5 THE DATA ACQUISITION

The initial design of the Crystal Box data acquisition system has been discussed by Butler, *et al*..[58] Modifications to the system have been made since that time and are included in the description below.

3.5.1 Requirements and Specifications

The data for a Crystal Box event includes nearly 2000 signals, 90% of which are analog in origin. Each event generates energy, timing, and pile-up information from the 432 NaI(Tℓ) channels, energy and time information from the 88 scintillator channels, time information with associated identification words for up to 728 drift chamber wires, and additional latch bits and status words. Thus, a single event could produce about 3000 16-bit words of data. Further information was available from 120 CAMAC scaler channels which were recorded only once every 10 seconds. At the proposed trigger rate of 10 Hz, and a total running time of 4 x 10^6 seconds (1000 hours), 3000 words/event leads to writing about 7000 tapes. Experience has shown that 500 tapes is a reasonable limit on the amount of data that can be analyzed with available manpower and computer resources. Therefore, a reduction in the data rate by a factor of 15 is highly desirable.

A factor of 4 data-rate reduction can be obtained by eliminating "empty" channels through a sparse-data scan. In the drift chamber, for example, Monte Carlo simulation indicates that an average of about 70 non-zero signals will be generated for each event. For the the NaI(Tℓ) crystals, Monte Carlo calculations and preliminary experimental tests indicate that typically fewer than 100 crystals are hit, yielding 100 words each of energy and timing information and 200 words of pile-up data. Assuming some form of sparse-data scan, the volume of compressed data from the crystals (conservatively 400 words), plus that from the drift chamber (70 words), and 180 words of scintillator data, gives a total of about 650 words/event. Various status and latch bits bring the grand total to around 750 words/event, a factor of 4 less than the original 3000 words/event. Thus, an additional factor of 4 or 5 reduction in taping rate is required from the higher level on-line analysis.

Another major problem of such a large experiment is to verify that the multitude of individual components are functioning properly, and that the various subsystems are performing their designed tasks. For instance, one must track the individual gains of the NaI(Tℓ) PMT's and their readout electronics often enough to preserve the best possible energy resolution, as well as assuring that the trigger logic elements are all operating properly to give the desired event configurations. A solution to this problem involves a flexibility in on-line histograms and displays to allow the user to monitor the experiment and diagnose problems as they arise.

The required functions of the Crystal Box data acquisition system are summarized as follows:

1. acquire, preprocess, and record the data

2. histogram, display, and/or analyze at least a portion of the data on-line

3. provide for monitoring and exercising the apparatus and for diagnosing faults in it

4. reduce the recorded event rate by a factor of 5

5. perform a sparse-data scan on the drift chamber and crystal data

3.5.2 Choice Of Hardware

To minimize the new programming effort, the LAMPF standard data acquisition system[59] and the associated standard program "Q"[60] were adopted as a basis for the Crystal Box system. To a large extent, this "Q" system automatically satisfies the requirements 1-3 above.

A schematic diagram of the data acquisition hardware, including details of the NaI(Tℓ) crystal readout, is shown in Figure 18. The use of the Q-system required the main computer to be a DEC PDP-11 type CPU. The particular choice was a PDP-11/44 with 1 Megabyte of memory, a 10 Megabyte Winchester-drive fixed disk, and two Kennedy 1600 BPI tape drives. An 11/44 was chosen as an optimum combination of computing speed, memory capacity, and cost.

Interfacing to CAMAC is accomplished through a Microprogrammed Branch Driver (MBD)[61] with several (up to 7) CAMAC crates connected along a serial branch highway. Programs residing in the LeCroy 4298 controller and in the 4 LSI's (both described above) perform sparse-data scans for the drift chamber and NaI(Tℓ) calorimeter, respectively, thereby satisfying requirement 5. They provide additional capabilities for satisfying requirement 3, such as the auto-trim feature of the drift chamber readout and the flasher monitoring system for the NaI(Tℓ) crystals.

3.5.3 Taping Cuts

With regard to requirement 4, a decision was made not to implement on-line cuts of the events during the January 1984 cycle. It was felt that some detailed analysis of uncut data was necessary to define allowable cuts which would not substantially reduce the acceptance of expected prompt events such as $\mu^+ \to e^+ \nu_e \bar{\nu}_\mu \gamma$ and $\mu^+ \to e^+ e^+ e^- \nu_e \bar{\nu}_\mu$. It is vital to the final interpretation of the experiment that these processes be observed at rates that are understood and therefore great care must be taken to avoid systematically excluding them from the data tapes. However, the capability to apply representative cuts was demonstrated and the subsequent effects on system performance were examined for a small sample of events. Without the taping cuts, and running for 6×10^5 live seconds at a conservative muon stopping rate of 3×10^5 μ^+/s (since the pile-up detection system was not fully operational), an average trigger rate of 5 Hz was observed and about 150 data tapes were collected during the cycle. This number of tapes is in agreement with the calculations above.

Chapter IV

EXPERIMENTAL PROCEDURE

The experimental procedure consists of three major phases. The first phase involves the initial set-up of the detector and the associated electronics, including the fast logic trigger. The other two phases, data-taking and calibration/stabilization, are interleaved as necessary for proper monitoring of the experiment.

4.1 SET-UP PHASE

The set-up phase begins even before the beam is available. Electronic gains, thresholds, and timings are roughly set, where possible, to minimize use of beam time for those preliminary procedures requiring muons.

4.1.1 The Beam

After the channel magnets are set to nominal values, a set of special beam monitor counters is used to fine tune the beam position and to measure the intensity profile near the stopping region. A sample of such a scan is given in Figure 23. The drift chamber is then inserted and special runs are taken with electron singles triggers. These tapes are analyzed, and all drift chamber tracks reconstructed to determine the stopping distribution on the target (see Figure 24). Reconstructed tracks that do not intersect the target are assumed to come from the

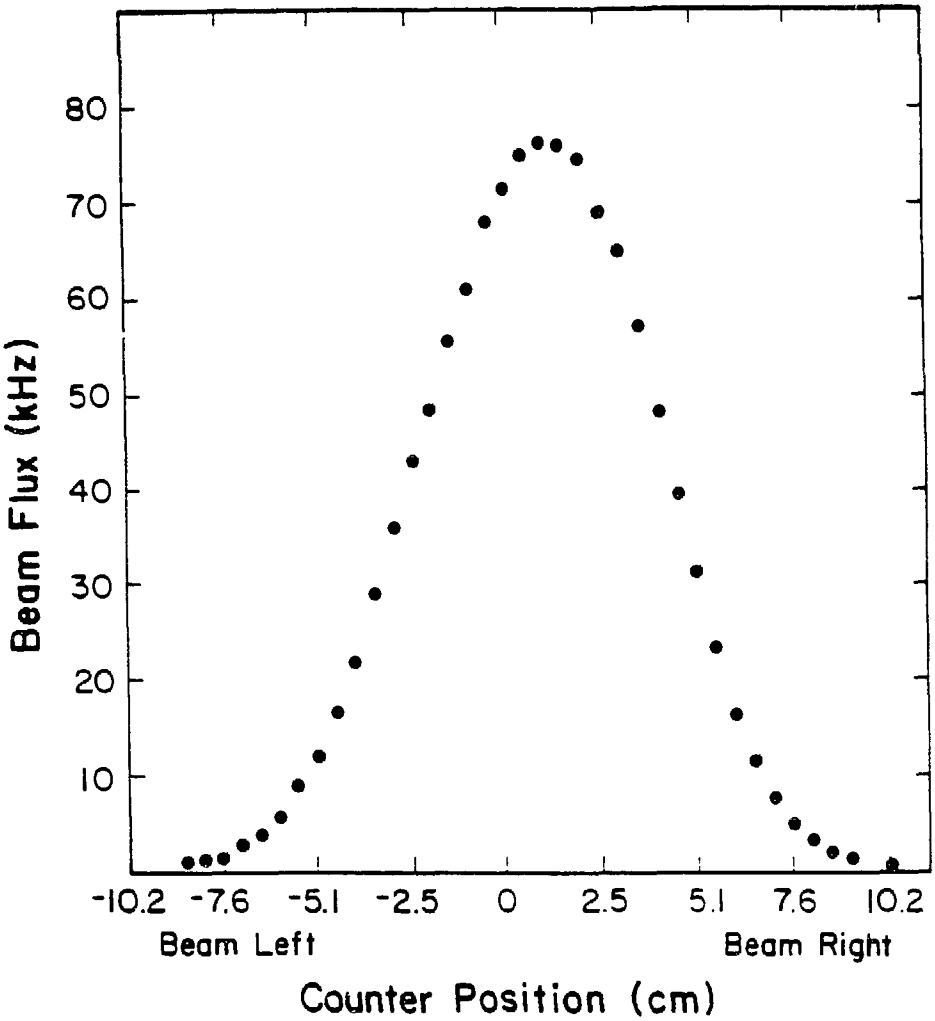

Figure 23: The one-dimensional beam profile.

The profile was measured with a thin plastic counter scanned across the target region before insertion of the drift chamber.

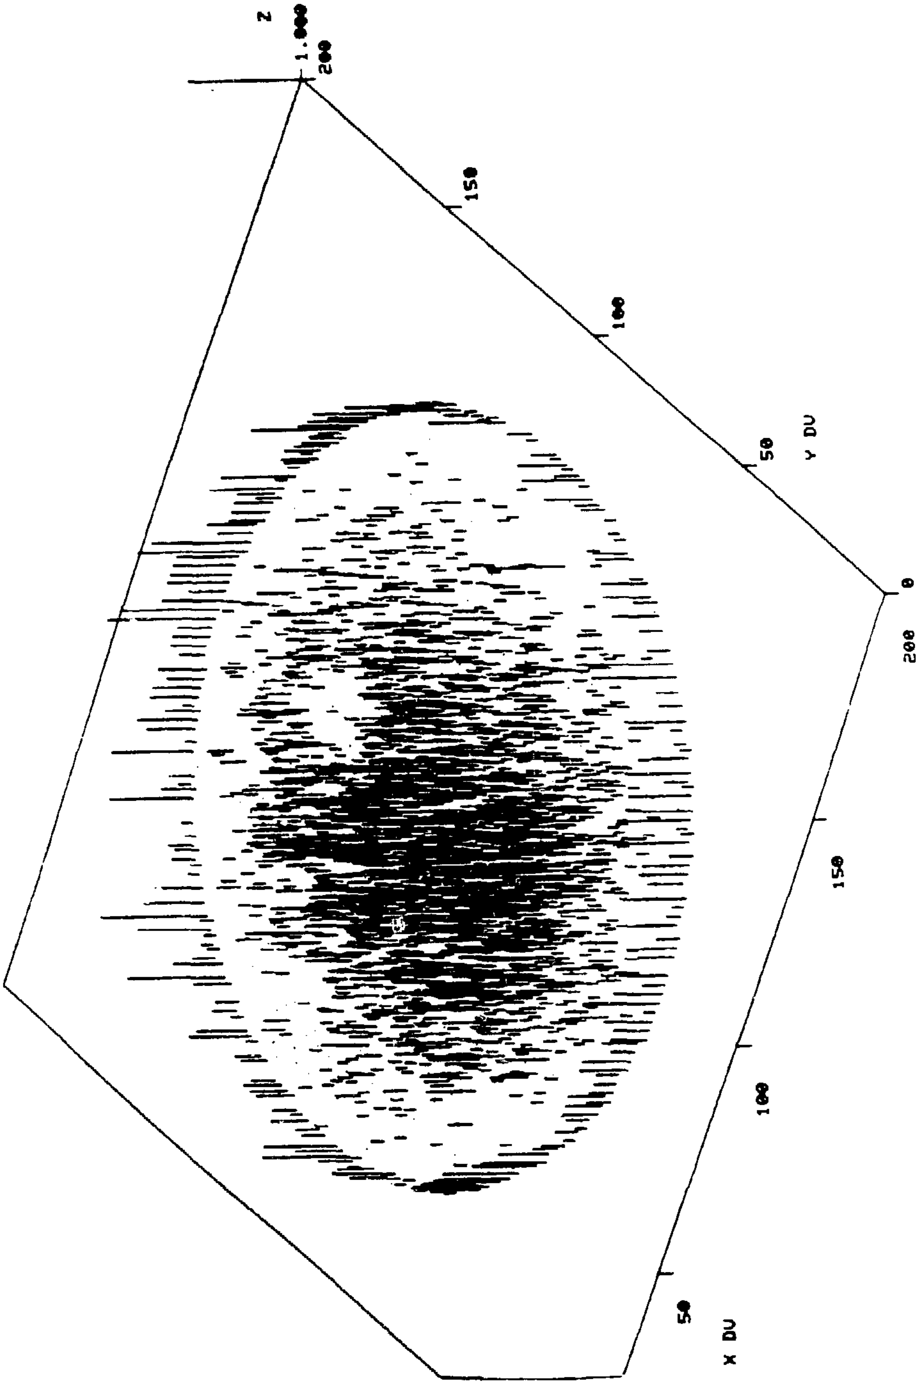

Figure 24: The beam profile on the stopping target.

This profile was determined by reconstructing the distribution of drift chamber track intersections with the target for electron singles triggers.

drift chamber inner foil. This assumption is responsible for the "ring" of events around the target in Figure 24. Scaler readings made during the runs normalize the integrated stopping rate. Problems in the beam configuration indicated by the drift chamber analysis are corrected, when possible.

4.1.2 Energy Calibration

4.1.2.1 Hodoscope

When the beam was available, the hardware gains for the scintillation counters were aligned by adjustment of their PMT high-voltage supplies. On-line ADC histograms, containing the minimum-ionizing energy peaks, were analyzed to determine the necessary voltage adjustment for each counter. The gains were set so that the nominal threshold to fire a trigger-counter CFD was equivalent to 0.5 MeV.

4.1.2.2 Calorimeter

Immediately before the beam is available, the NaI(Tℓ) amplifier gains are roughly calibrated, at the 10% level, using the 4.4 MeV γ-ray line from a Pu-α-Be source. The source is positioned inside the Box (with the drift chamber removed) and centered on the face opposite the quadrant to be calibrated. This arrangement provides the most uniform illumination of the face crystals and the inner corner crystals. The source is moved outside the Box near a corner cluster to properly illuminate the outer corner crystals. The raw output of each clip/split amplifier is then examined on an oscilloscope and the pulse height of the clearly visible photopeak adjusted so that the CFD fraction of the amplifier output had a calibration ≈20 mV/Mev. Thus the CFD threshold of 100 mV corresponds to the desired energy threshold of 5 MeV in the NaI(Tℓ) crystals. Because the corner crystals have two phototubes contributing to the hardware energy sums, the individual gains of the corner crystal amplifiers are set ≈10 mV/MeV. (This means, however, that the corner crystal CFD's have an effective 10 MeV energy threshold.)

After all quadrants are hardware calibrated, one tape of data per quadrant is then taken, triggering on the source photons. In detail, the trigger is:

(exactly one NaI(Tℓ) row hit)•(at least 2 MeV in the calorimeter) .

In addition, the data on each tape was preceded by a set of flasher spectra. The combination of this information permitted a check of the software calibrations of the NaI(Tℓ) readout and established reference values for the flasher gain stabilization system.

4.1.3 Timing

After pulse height gains have been set for the scintillators and NaI(Tℓ) signals, the timing of their signals into the various trigger components must be adjusted. Where possible, signal timings are roughly adjusted using injected test pulses, but detailed timing, such as variations in phototube transit times, can only be adjusted using beam particles.

4.1.3.1 Hodoscope

It was realized early in the planning of this experiment that the individual timing of some 88 plastic counter signals as well as 432 NaI(Tℓ) signals was a task which could reasonably be accomplished only with an automated method. The timing signals for all scintillators (meantimer signals for the double-ended counters) are routed through a single CAMAC TDC. Muons stopping in the secondary target are used as a source of single electron triggers, with the I-counter providing the timing reference. The PDP-11 reads the TDC value and the scintillator latches for each event. The value of the TDC reading is assigned to one of 52 histograms, on the basis of the set latch, for events which have exactly one latch fired. These histograms then constitute a measurement of the relative times of all scintillators. An auxiliary program automatically determines the peak positions in all histograms and calculates a table of delay cable lengths (accurate to the nearest 0.5 ns) to be inserted.

4.1.3.2 Calorimeter

A modification of the above procedure is required for timing the NaI(Tℓ) signals. Since the selection of the histogram to increment is made on the basis of row latch information, only one crystal per row can be timed, in a given run. Timing groups are therefore defined with maximal separation of the crystals from one row to the next. A typical timing group consists of a set of crystals as indicated in Figure 25. Other groups are generated by a circular shift of one column by the crystal in each row. Only the crystals for one timing group are

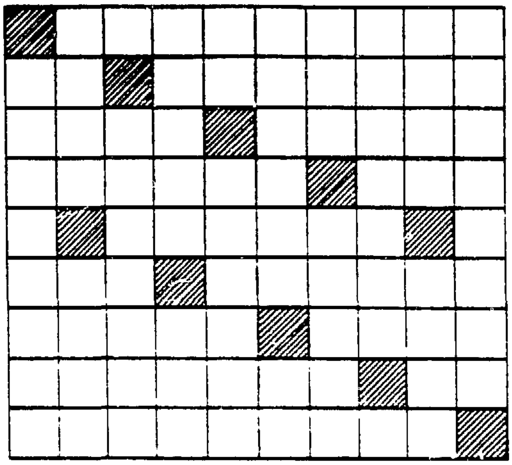

Timing Group 1

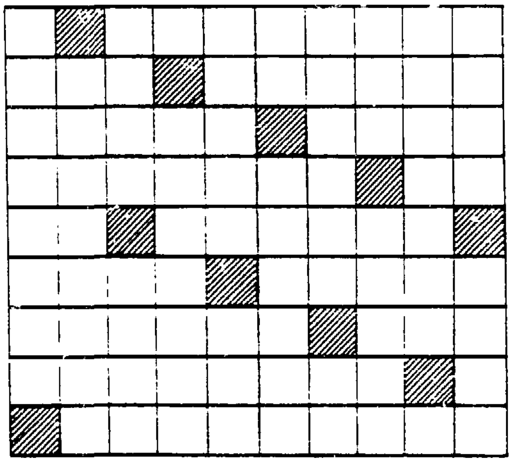

Timing Group 2

Figure 25: Two NaI(Tℓ) timing groups.

The groups are configured to minimize the chance of single electron hits illuminating more than one crystal in a group.

connected to their respective amplifiers for a particular run. The maximal spacing of crystals enhances the efficiency of data-taking by reducing the probability that showers hitting one crystal will fire any other in the timing group. (As for the scintillators, one and only one row latch set is required in order to histogram an event.)

Once the crystal signals are properly timed into the trigger logic, the flasher system can be used to check the timing of pulse height signals into the PHIM's. Since every flasher trigger gives a pulse to every crystal channel, the integration gate versus signal timing can be checked very efficiently. (Of course, these signals need only be timed to a few percent of the integrating time, or within 5 ns.) A systematic offset ≈20 ns must be allowed for when using the flasher events with corner crystals. This offset is attributed to their different geometry compared to that of face crystals.

Once the individual PMT signals have been properly trimmed, the many trigger coincidences must be systematically timed in as well. This procedure, although extremely tedious, is essentially a standard technique and will not be discussed in detail.

4.2 MUON BEAM DATA-TAKING

Because the pile-up detection system was not yet functional, the beam could not be set at the maximum muon stopping rate of 0.45 MHz. At the start of the January data-taking cycle, data runs were taken at gradually increasing rates and were analyzed for pile-up effects to select a muon stopping rate of 0.3 MHz for production runs. The average accelerator duty factor was ≈ 7%.

Before each run, the pedestal position of each NaI(Tℓ) PHIM channel is monitored and written to tape. A flasher run is then taken to monitor the PHIM gains and verify that all NaI(Tℓ) TDC's are functional. The pulse height and timing spectra collected in the LSI's are passed to the PDP-11, analyzed and written to tape for more detailed analysis off-line. Any channels exhibiting unacceptable drifts are noted in a summary printout generated by the PDP, prompting the experimenter to take appropriate action. Based on the updated pedestals and flasher peak positions, new PHIM gains are calculated and sparse data scan thresholds are updated accordingly. Part of the start-up procedure for production data runs automatically performs an auto-trim on the drift chamber readout system as described above. A full tape of nominally 30,000 events is then recorded requiring 1.5 to 2 hours. The plastic scintillator readout is calibrated once a day by taking two tapes of electron-singles triggers, one for muon decays originating in the I-counter, and another for decays from the target. In addition, this data includes a flip-flop in the TDC start signal which alternately switched in and out a fixed 10 ns delay cable, thus producing a double-peaked timing spectrum for all scintillators. The separation of these two peaks thus determined the gain of the TDC channels. This data was analyzed off-line to give the scintillator timing constants and ADC gains. Similar tapes, triggered on single NaI(Tℓ) quadrants, are taken at less frequent intervals to permit off-line calibration of NaI(Tℓ) TDC's and PHIM's with Michel positrons from normal muon decay.

The stability of the beam profile was monitored several times during the cycle by electron-singles tapes of muons stopping in the primary target. The analysis of drift chamber tracks for these tapes verified the distribution of track intersections with the target.

4.3 PION BEAM DATA-TAKING

As described below (in section 5.2.2.1), the LH_2 target, illuminated by a negative-pion beam, provides a means for the absolute calibration of the NaI(Tℓ) calorimeter. For the last two days of the cycle the drift chamber was removed and replaced by the LH_2 target. Each pion data run is preceded by a pedestal and flasher sequence, just as for the muon data. Forty pion calibration tapes containing about 1.5×10^6 events were accumulated at the rate of 1 tape/hour. This sample was estimated to be sufficient to provide a minimum of 1000 calibration events for each crystal.

Chapter V

ANALYSIS OF DATA

The two major divisions of the off-line data analysis were the detector calibration and the reduction of candidate events. For the $\mu \rightarrow e\gamma$ data, a maximum-likelihood analysis of the residual event sample was necessary to arrive at the best possible branching ratio limit. Before describing the calibration procedures and the detailed analysis of the data, a brief description of the relevant computer codes is required. Also, a few concepts must be defined and discussed. The concept of a NaI(Tℓ) "clump" is needed to describe the energy calibration. The "position algorithm" is also essential to the energy calibration procedure.

5.1 COMPUTER CODES AND ANALYSIS CONCEPTS

All of the off-line data analysis was done on a VAX 11/780 computer at the LAMPF Data Analysis Center. The latest version of the codes run on the VMS Version 4.0 operating system.

5.1.1 Monte Carlo Simulation

In an experiment with a detector as complex as the Crystal Box, it is frequently necessary to model the performance of the apparatus through computer simulation. This analysis, for example, enables the experimenter to select the optimum experimental running conditions and to predict the effectiveness of possible data-analysis cuts. A simulation program is essential also for the proper calibration of the Crystal Box detector and for predicting the detector response to previously unobserved phenomena (namely neutrinoless muon-decay).

The "Monte Carlo" modelling program for the Crystal Box detector is implemented on the same VAX 11/780 machine used for all of the off-line data analysis. The simulation of energy deposition was based on the electron-gamma shower code EGS3 developed by Ford and Nelson.[62] An additional program,[63] originally written at LAMPF to aid in the design of the drift chamber, and subsequently used to optimize the track-reconstruction algorithm, provides the capability to simulate electron tracks, including multiple-scattering effects. The merger of these two codes made it possible to simulate the full geometry of the Crystal Box detector and to generate a variety of event types, including $\mu\rightarrow e\gamma$, $\mu\rightarrow e\gamma\gamma$, $\mu\rightarrow 3e$, $\pi^0\rightarrow\gamma\gamma$, $\pi^- p\rightarrow n\gamma$, monoenergetic γ's, and Michel positrons (from either the target or the I-counter). The program output can emulate the information contained in the on-line data tapes, so that the same off-line code may be used to analyze real data and Monte Carlo events.

To match the shape of the observed energy distributions, it is necessary to add to the Monte Carlo simulation an additional component of energy resolution, beyond that predicted by the EGS3 code due to shower fluctuation and energy leakage. This additional component accounts for the degradation in resolution due to all sources other than leakage, including the non-uniformities of light production within any particular crystal, the light collection inefficiencies of the crystal-

PMT system, electronics noise and instability, and residual relative errors in the individual crystal calibrations. The method of introducing this resolution is to multiply the energy deposited in each crystal with a gaussian smearing factor that is centered at unity and has a FWHM which varies as the fourth root of the energy to be smeared. The best agreement with the experimental data is obtained by setting the smearing width to 14% at the 0.667 MeV line of ^{137}Cs. The single crystal resolution for ^{137}Cs photons was observed to range from 12% to 15% during the calorimeter acceptance testing at Harshaw Chemical Company. The leakage contribution to resolution is negligible at this low energy, so the above value for the smearing width seems reasonable.

5.1.2 Data Replay

The data replay code was a modification of the LAMPF standard "Q"-Replay system.[60] In view of the large number of tapes to be analyzed, the non-standard capability to replay data from tapes stored on disk was implemented. The "HPL" histogramming package, available at Los Alamos, was incorporated, as well as specially designed graphics routines, to aid in the visualization of event geometries, energies, and timing. The facility to make and evaluate data reducing cuts was also an essential part of the code.

5.1.3 Track Reconstruction

A detailed description of the drift chamber track-reconstruction (DCR) program is beyond the scope of this thesis. (See reference 63 for a complete explanation). This code provides a reconstructed trajectory, when possible, for electron tracks through the drift chamber which hit seven or more wire planes. For electrons from the target, acceptable tracks are required to make an angle of > 1^{o} with respect to the target surface.

Although the DCR program was initially developed independent of the replay code, the final version of the analysis code includes the track reconstruction capability. However, the average CPU time required by the DCR program is 0.2 seconds per track, so that a substantial reduction of the event sample is required before initiating track reconstruction.

5.1.4 Clump Definition

Although the trigger treats an entire NaI(Tℓ) quadrant as a single "particle", the functional unit corresponding to a particle in the off-line analysis is a "clump". It is possible for several clumps to exist within one quadrant. As many as nine clumps may be defined in a given event. The computer algorithm for finding clumps is as follows:

1. Find the crystal, not in any previous clump, with the largest energy > 5 Mev. This crystal is called the "high pulse-height crystal" (HPHC).

2. Determine the "neighbor set" (NS) of crystals:

   a) If the HPHC is more than one row away from a corner, the NS consists of all crystals within two rows and two columns of the HPHC. (For an HPHC in the central region of a face, this set contains 25 crystals).

b) If the HPHC is a face crystal within two rows of a corner, the NS consists of all crystals in the same quadrant within two rows and two columns of the HPHC, and all of the nearest corner crystals.

c) If the HPHC is a crystal on the inward-pointing diagonal of a corner array, the NS includes all of that corner array, and all of the crystals from both adjoining face rows.

d) If the HPHC is an off-diagonal corner crystal, the NS contains all crystals in that corner, and all crystals in the nearest two rows of the closest quadrant.

3. Scan the NS for all previously unused crystals with energy > 0.1 MeV, add those energies to the clump energy sum, and eliminate those crystals from use in further clumps.

4. Return to step 1 if fewer than nine clumps have been found.

In order to achieve the best possible resolution, electron-clump energies must be corrected for the ionization energy loss in all material along the electron trajectory from the target to the face of the crystals. This correction consists of two components, the measured energy loss in the intercepted scintillator(s), and the estimated energy loss in the inert material traversed. The first component comes directly from the calibrated response of the trigger counter ADC's. The second is derived by tracing the track reconstructed by the drift chamber through the detailed geometry of the Box, collecting the total intersected thickness of each type of inert material, and multiplying each thickness by the average dE/dx in that material. These additional energies are added to the clump energy sum, and this final sum is called the "electron energy", E_e, for that clump. For γ's, the "photon energy", E_p, is identical to the clump energy sum.

5.1.5 NaI(Tℓ) Position Algorithm

An estimate of the particle shower position in the NaI(Tℓ) calorimeter is calculated as the weighted centroid of the energy deposit distribution. The "position algorithm" computes the shower position, $\underline{x}_s$ according to:

$$\underline{x}_s = \frac{\sum \underline{x}_i E_i^{\beta}}{\sum E_i^{\beta}} \qquad (16)$$

where the sum is over all face crystals in the clump with energy > 0.5 MeV, $\underline{x}_i$ is the position vector to the center of crystal i (at its face), E_i is the energy deposited in that crystal, and β is the exponential weighting factor which depends on particle type.

The above computation amounts to a projection of the centroid of the shower envelope to the face of the crystals. For particles with non-normal incidence to the crystals, this projection is not a good representation of the particle impact point, because the particle will have travelled some depth, δ, past the surface, along its original trajectory, before interacting. Thus, a "depth correction" is made to $\underline{x}_s$ to arrive at the final position estimate $\underline{x}_f$. The expected δ for a given hit will depend on the type of particle and on its energy. Given δ, the position $\underline{x}_s$ is projected normally into the calorimeter until the trajectory from the center of the target to the projected point intercepts a path length δ through the NaI(Tℓ) crystals. The intersection point of this trajectory with the face of the crystals is the final clump position $\underline{x}_f$.

The parameters β and δ were optimized by analyzing Monte Carlo electron-singles events and minimizing the resolution of the reconstructed position versus the actual impact point. The best values of β were $\beta_e = 0.65$ for e^+'s, and $\beta_p = 0.55$ for γ's. Table 4 gives the computed energy dependence of δ for both e^+'s and γ's. Two positions are calculated for each clump, one assuming an electron hit, and one assuming a photon hit, since the particle type is determined later in the analysis program. The values of δ used for a particular clump are interpolated from the appropriate column of Table 4 using the clump energy.

Figure 26 illustrates the observed position resolution of the NaI(Tℓ) calorimeter taken from type II events with impact points within the central arrays of the Box. In agreement with the expectations from the Monte Carlo studies, the resolutions are $\sigma_{xy} = 0.72$ cm (transverse to the beam direction), and $\sigma_z = 0.79$ cm (along the beam direction). The one centimeter offset observed in the z-coordinate plot (Figure 26(b)) was traced to an alignment error in the position of the drift chamber. The error was accounted for in the analysis of the $\mu \rightarrow e\gamma$ events.

TABLE 4

Energy dependence of depth parameter, δ, for electrons and photons.

Electrons		Photons	
E_c(MeV)	δ_e(cm)	E_c(MeV)	δ_p(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	5.4[illegible]	126.9	9.43
92.8	6.82	-	-

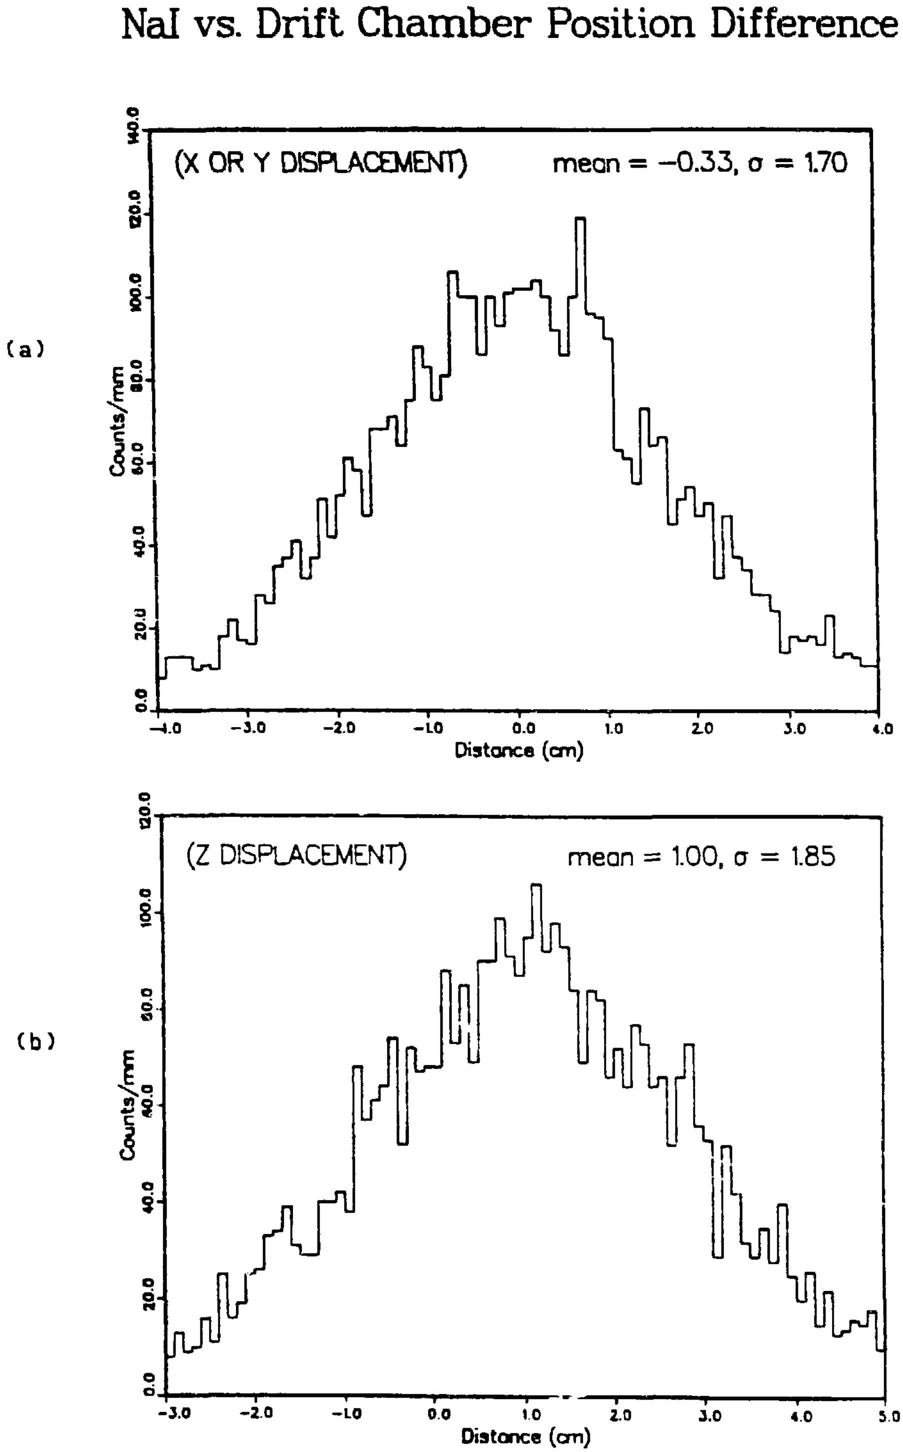

Figure 26: The NaI(Tℓ) transverse (a) and longitudinal (b) position resolution for type II events.

5.2 DETECTOR CALIBRATION

About 100 calibration tapes were accumulated to track the drifts in the detector timing and energy parameters. Two types of calibration tapes were collected, both using an electron singles trigger. Type I tapes observed muons that stopped and decayed in the I-counter. These tapes were used for timing calibration of the plastic scintillation counters and NaI(Tℓ) crystals. Type II tapes, from decays of muons stopped in the target, were also used to define the trigger counter timing constants and those of the drift chamber. A further use of the type II tapes was to check the NaI(Tℓ) energy calibration for electrons. A third type of calibration tape, type III, was triggered by 1-γ and 2-γ events from the hydrogen target. These tapes were used to define the energy calibration constants of the NaI(Tℓ) crystals for γ's and to check the crystal timing resolution for γ's.

5.2.1 Timing

5.2.1.1 Hodoscope

As mentioned above (see section 4.2), the timing constants for the plastic counters were monitored daily using two types of electron calibration tapes. The single-ended veto counters were calibrated using only the type I tapes. For each PMT, a plot of its timing, relative to the I-counter, was collected for events in which only one scintillator discriminator was fired. This selection of events reduces systematic errors due to pile-up and backscatter effects. A typical spectrum is shown in Figure 27. The two peaks are separated by a 10 ns delay, which is alternately switched in and out of the TDC-start signal path during data-taking. A fit to the two peak positions determines the timing calibration constants, as follows:

1. The position of the first peak, t_0, is the TDC channel corresponding to hits which are "in-time" with the I-counter signal.

2. The second peak position, t_1, then determines the TDC gain, g_t in units of 0.1 ns/channel, according to:

$$g_t = \frac{100}{t_1 - t_0} . \tag{17}$$

The timing constants for the trigger counters are derived from the calibration tapes by several steps. First, the type I tapes are replayed to give individual timing plots of all PMT's (both upstream and downstream) with respect to the I-counter TDC. Again, to eliminate pile-up and backscatter effects, events are selected which have only one scintillator fired. Because the trigger counter timing resolution is on the order of one TDC channel, the individual type I spectra, as in Figure 28, are not sufficient to determine the t_0's to the required accuracy. These spectra can only be used to determine the individual gains, g_t's, as above.

The next step is to collect spectra of the time difference between the upstream and downstream PMT's of each counter using the type II tapes. The central origin of electrons for the type II tapes gives a more symmetric time difference spectrum than that of type I tapes. Such a spectrum is shown in Figure 29. The centroid of this spectrum, Δt, determines the relative values of the timing offsets, t_{0u}, t_{0d}:

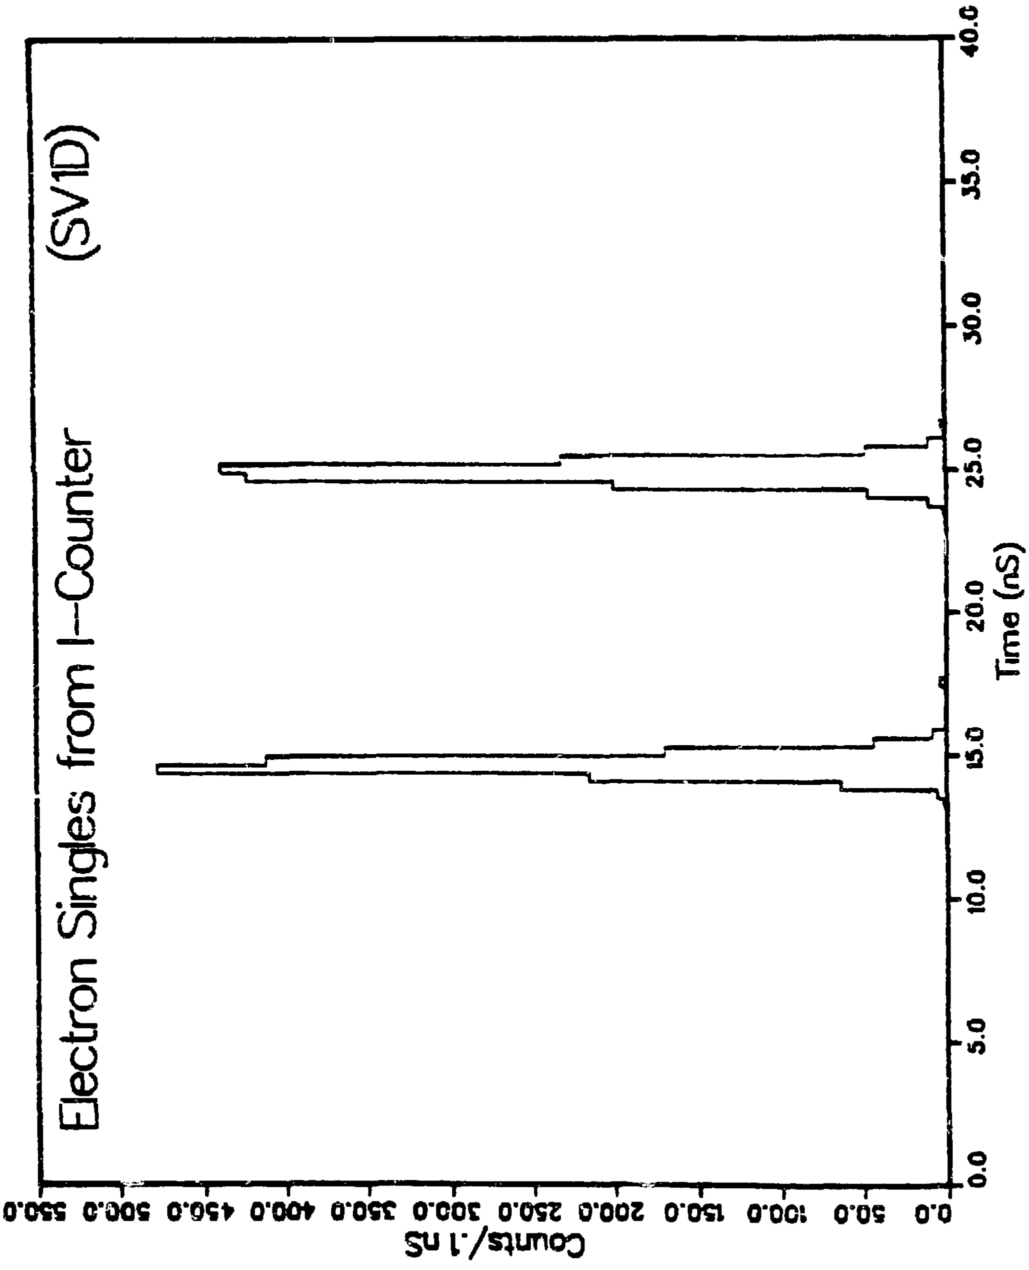

Figure 27: Typical veto counter timing spectrum for type I calibration tape.

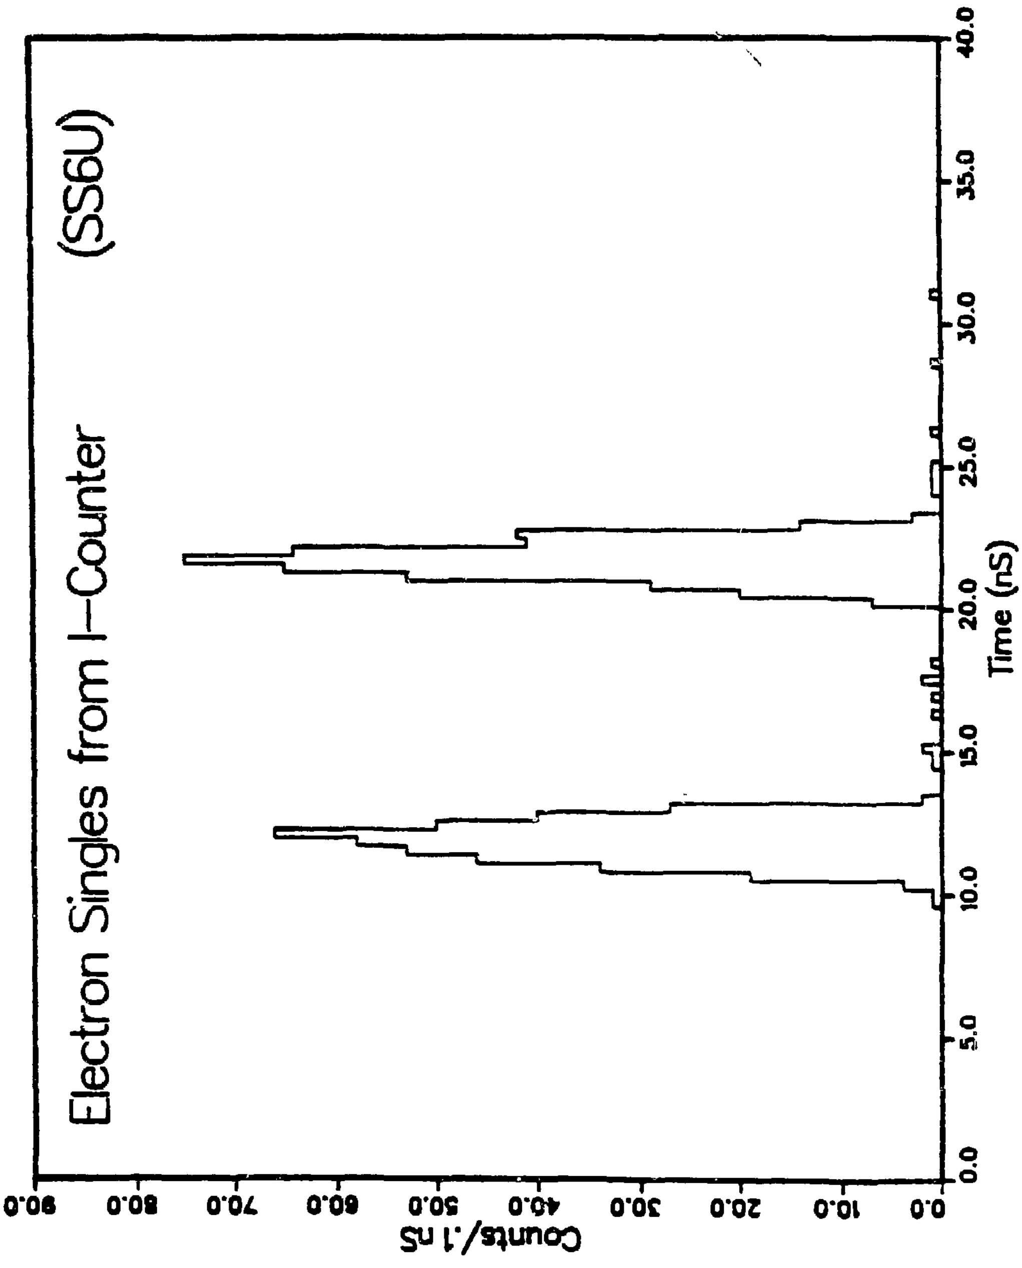

Figure 28: Typical trigger counter timing spectrum for type I calibration tape.

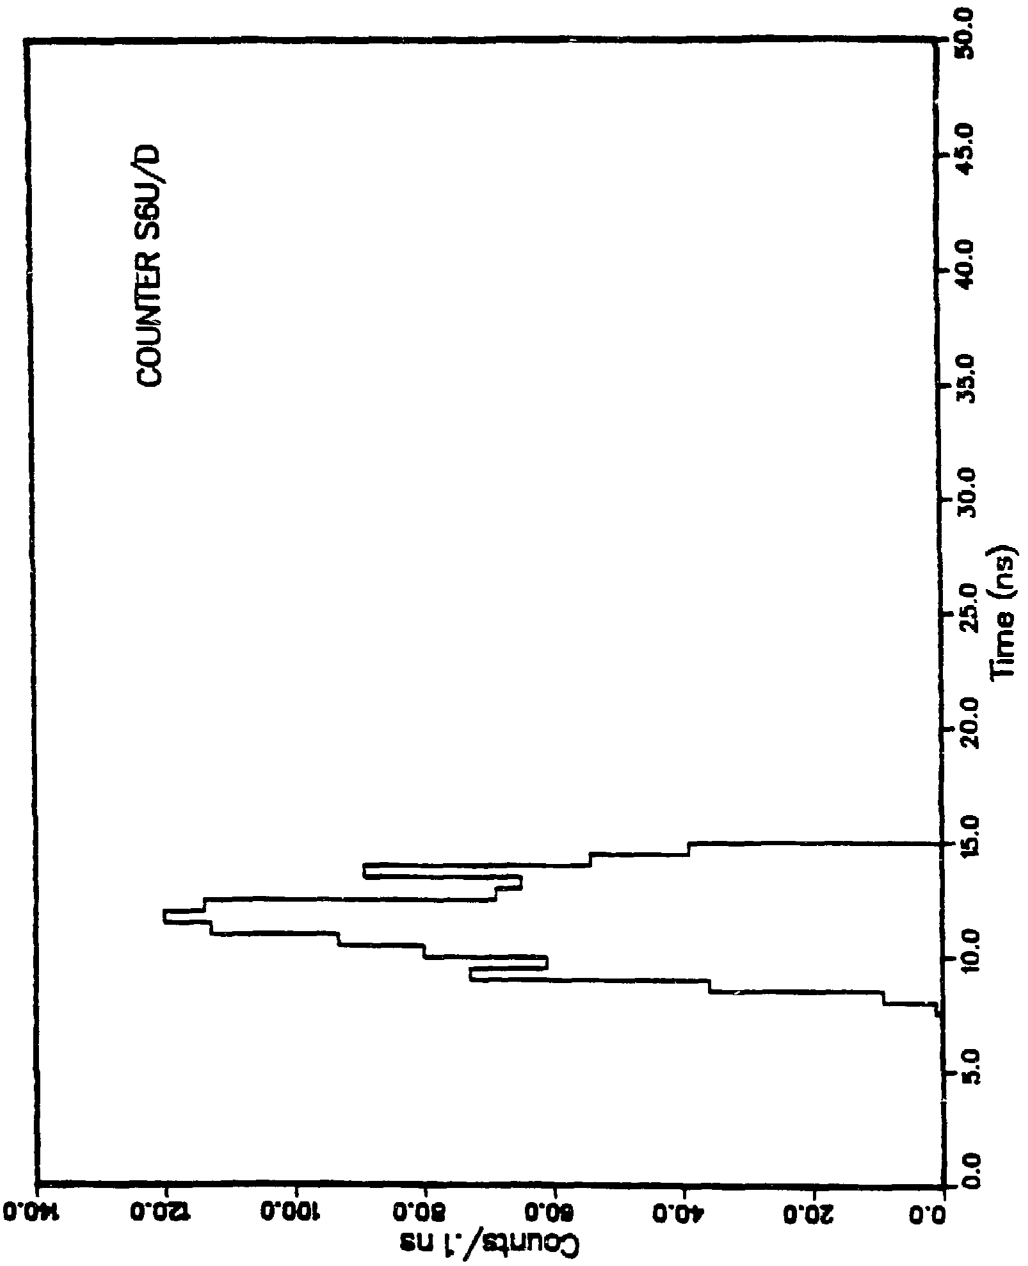

Figure 29: Upstream-minus-downstream time difference for a typical trigger counter using type II tape.

$$t_{0u}-t_{0d} = \Delta t \ . \tag{18}$$

The final step returns to the type I tapes to collect calculated meantime spectra, as shown in Figure 30. The position of the first peak, t_m, determines the absolute offset for both PMT's:

$$t_{0d} = t_m \ , \quad t_{0u} = t_{0d}+\Delta t = t_m+\Delta t \ . \tag{19}$$

The time resolution performance (for type I events), for the 36 trigger counters, was histogrammed as shown in Figure 31. The average timing resolution is seen to be σ_{ts} = 123 ps.

5.2.1.2 Calorimeter

The timing calibration of the NaI(Tℓ) crystals is closely analogous to that of the plastic scintillators, with the procedure for face crystals following that for guard counters, and the method for corner crystals closely resembling that used for the trigger counters. The modification to the above procedures is the selection of appropriate type I events with slightly different cuts. Face crystal calibration events pass the following cuts:

1. only one (trigger or guard) scintillator CFD fired
2. the trigger counter hit is within one row of the hit crystal
3. the z-position of the counter hit (as estimated by the counter PMT time difference) is within 12.7 cm of the hit crystal

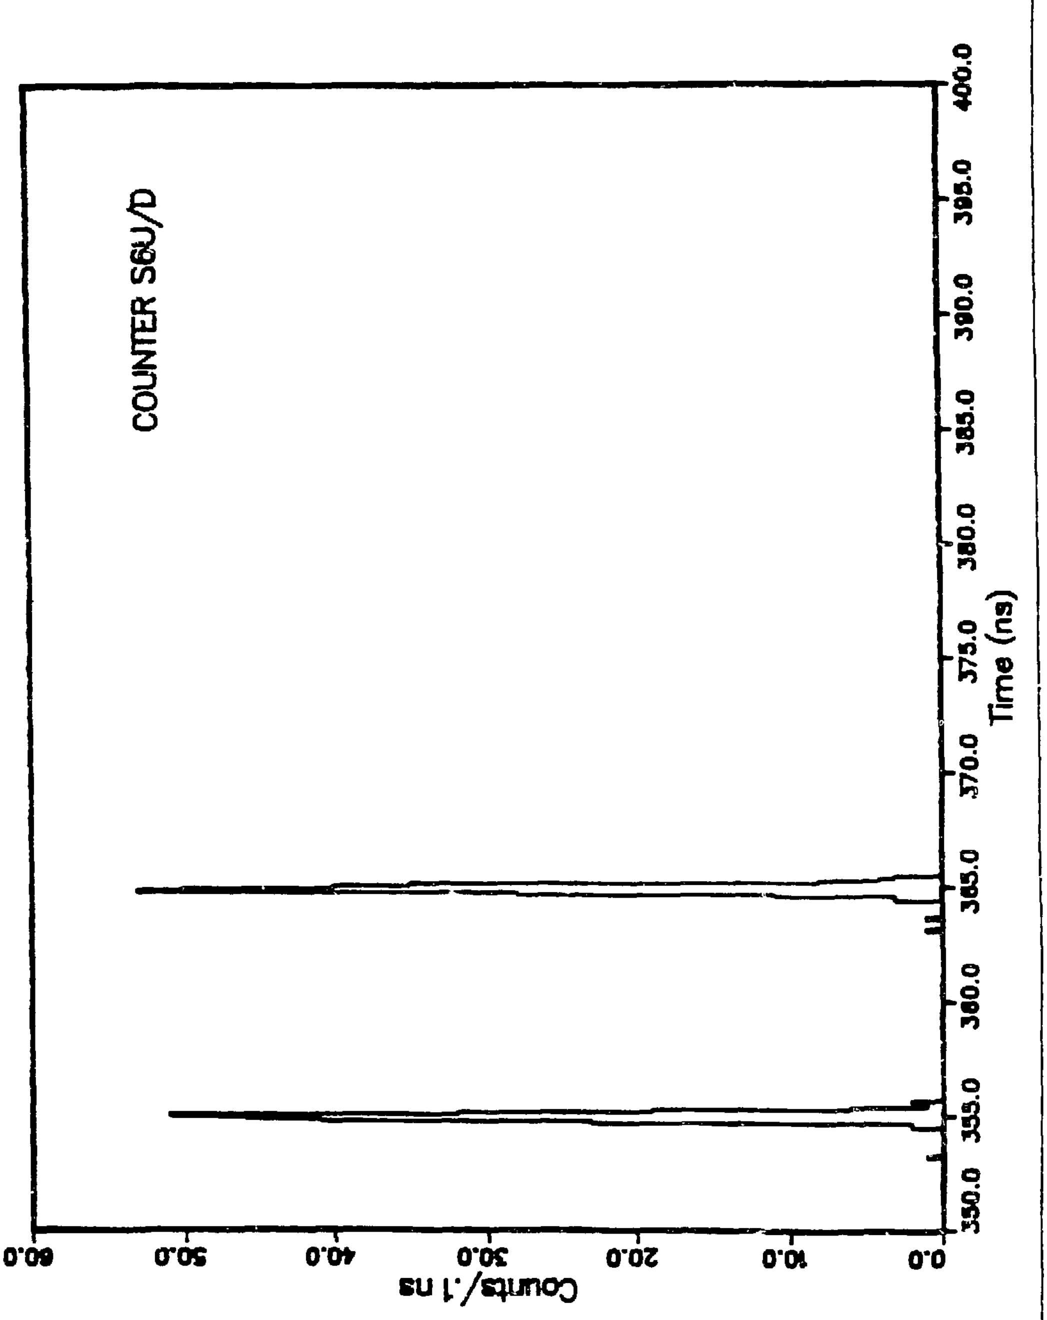

Figure 30: The distribution of calculated meantimes for a trigger counter using type I tape.

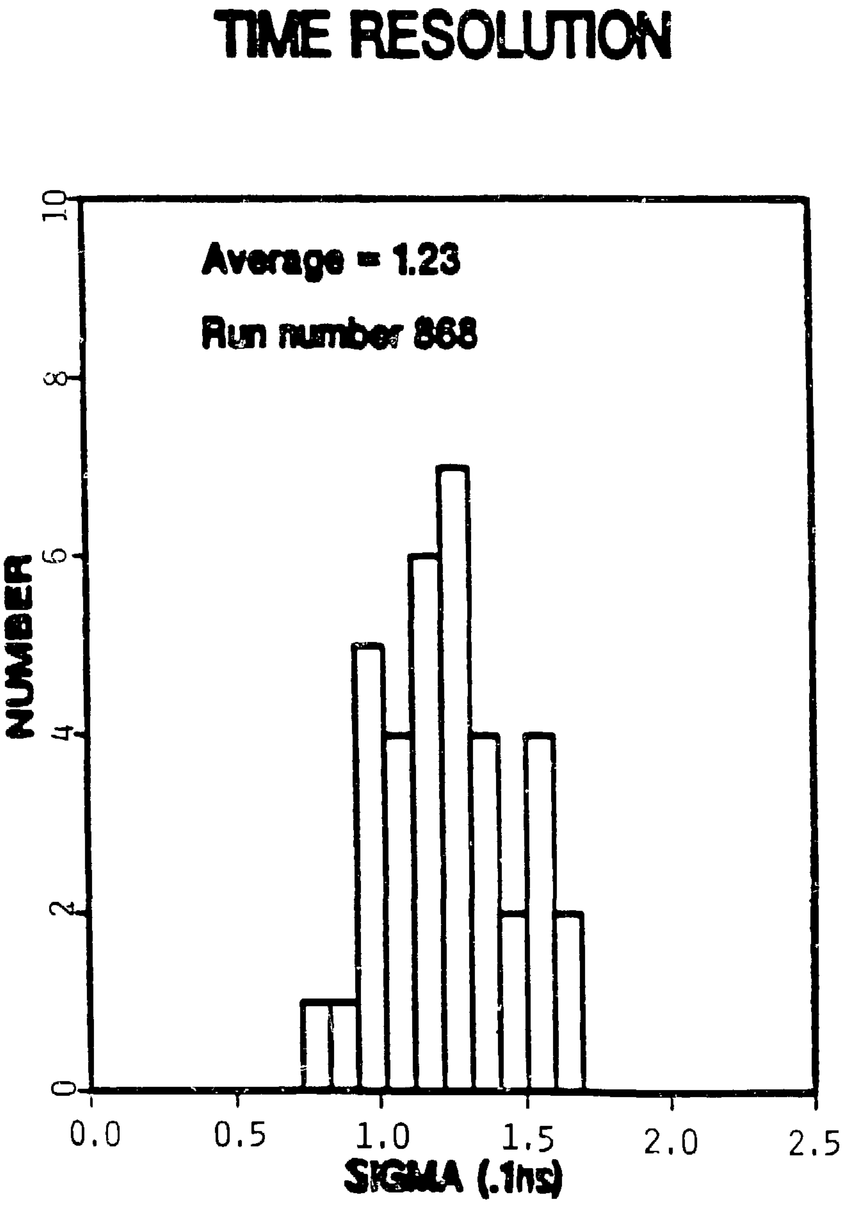

Figure 31: The distribution of timing resolution performance for the 36 trigger counters.

Corner crystal events are selected similarly, except that requirement 3 is not imposed.

The (electron) timing resolution of individual face crystals versus the I-counter is measured to be $\sigma_{tel} \approx 470$ ps. However, the superior timing resolution of the trigger counters is always used to define the time of an electron clump. The NaI(Tℓ) TDC's are only used to define the photon-clump time, and to detect energy pile-up (see section 5.3.1.1). The photon-clump time is an average over the clump crystals:

$$t_{avg} = \frac{\sum t_i E_i}{\sum E_i} \tag{20}$$

where the sum is over all crystals in the clump that fired a CFD within 5 ns of the HPHC time (the calculated meantime is used for the corner crystal contribution), and t_i, E_i, are the time and energy, respectively, for crystal i of the clump. As a check of the photon timing resolution, the measured time difference of coincident γ's from π^0 decays is histogrammed as shown in Figure 32. The FWHM of the distribution is 1.8 ns, which corresponds to $\sigma_{tp} = 540$ ps per photon. Evidently, the residual error in the relative calibrations of the individual crystals, and any difference in response to photons vis-a-vis electrons, contributes an additional resolution component, $\sigma_{err} \approx 270$ ps, to the single crystal resolution for electrons. Nevertheless, this timing performance is certainly noteworthy for such a large NaI(Tℓ) array.

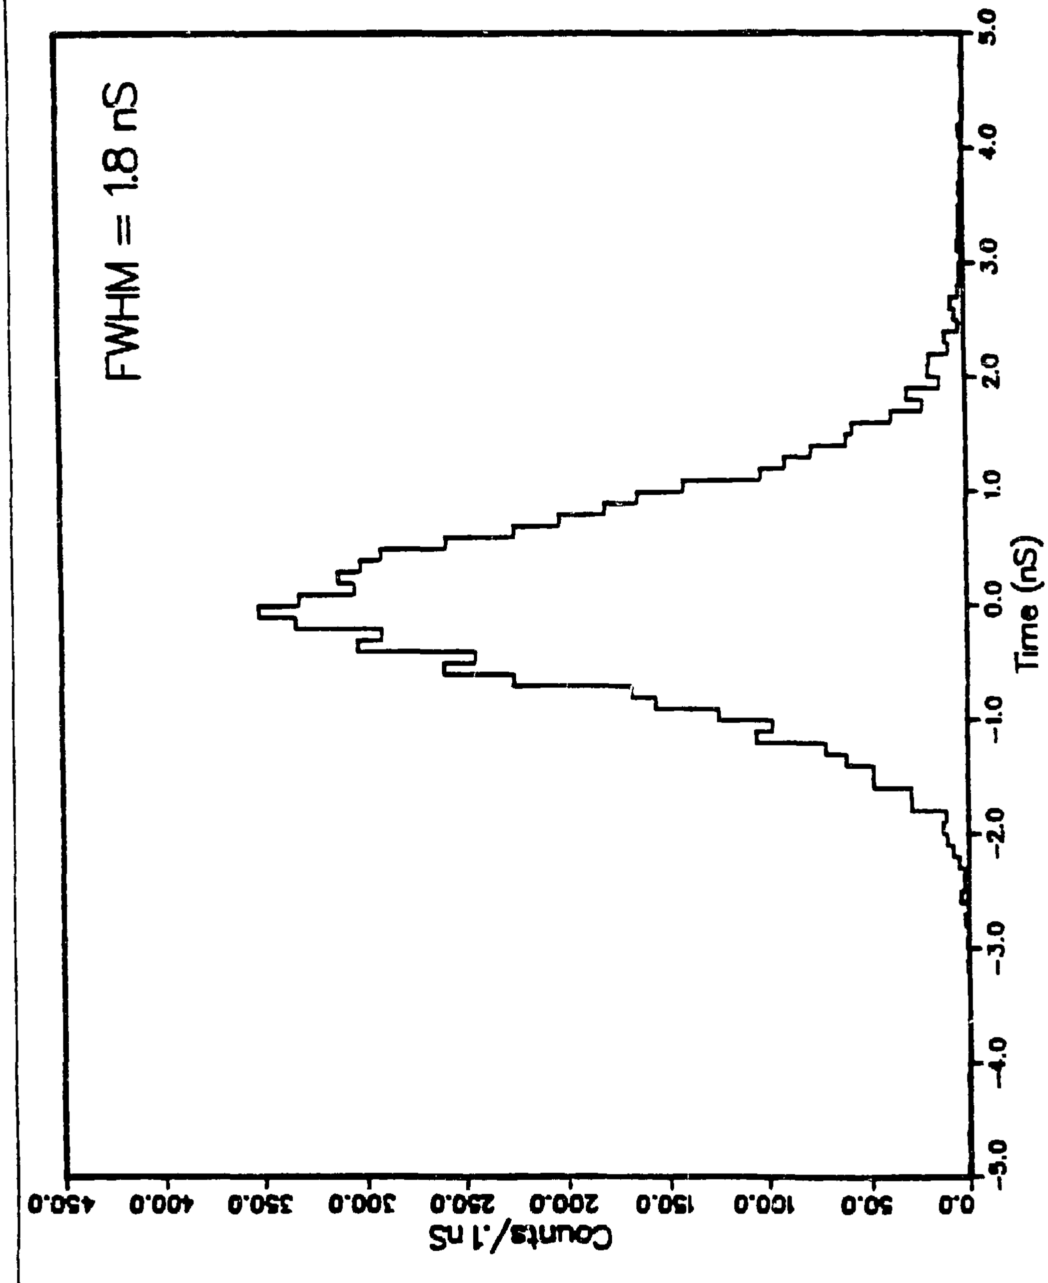

Figure 32: Timing resolution of NaI(Tℓ) for $\pi^0 \to \gamma\gamma$.

5.2.1.3 Drift Chamber

The drift chamber timing readout is calibrated with electron tracks (from type II tapes) which have hits in all eight wire planes. For each such event, the eight possible subsets of seven wire-hits are used to reconstruct a track. The timing difference between the excluded wire-hit and that expected from the reconstructed track is histogrammed. The centroid of each histogram determines an absolute offset to correct the timing of the corresponding wire. Because these adjustments affect, to a slight degree, the reconstructed tracks for each event, this procedure must be iterated for several passes through the data until a satisfactory convergence of the timing offsets is achieved.

The drift chamber performance was checked with cosmic-ray events which passed through the center of the detector, thereby simulating a two-particle event with a common origin and an opening angle of 180°. Such events were analyzed for reconstructed opening angle and distance of closest approach between the two tracks. The opening angle was found to have a $\sigma_{ang} = 0.5^0$. The resolution of the minimal distance was $\sigma_{min} = 0.4$ mm. Comparison of the chamber performance with the Monte Carlo simulation showed it to be consistent with an intrinsic spatial resolution per wire of $\sigma_{res} = 130$ μm, and a reconstruction efficiency of 97% for the low rate type II events. For the high beam rate used to collect data tapes, these performance parameters degrade to no worse than $\sigma_{res} = 150$ μm and 95% efficiency for eγ triggers.[63]

5.2.2 Energy

5.2.2.1 Calorimeter

As described above, in section 4.1.2.2, the NaI(Tℓ) energy-readout hardware is roughly calibrated by equalizing the response of each channel to the 4.4 MeV photopeak from a Pu-α-Be source. The absolute calibration of each channel is adjusted off-line in several steps. First, the shape of the non-linear calibration function is derived for each readout channel by analysis of a special set of test data. The absolute scale of the calibration function is then set by the analysis of γ spectra (from π^0 decays), collected in the type III tapes during the hydrogen target runs. Finally, the energy scale is checked by the 4.4 MeV source line and the monoenergetic γ from $\pi^- p \rightarrow n\gamma$. A check of the charged particle response is provided by the Michel spectrum from type II tapes.

Initial efforts to derive consistent linear calibration functions, using both the source data and pion data, revealed the NaI(Tℓ) readout system to be significantly non-linear. The NaI(Tℓ) clip/split amplifier was found to be the source of this non-linearity. A testing circuit was assembled, which utilized an Ortec pulse generator and a Tektronix precision variable-attenuator, to inject precisely graduated pulses into the NaI(Tℓ) readout channels. A typical response curve is shown in Figure 33(a), and Figure 33(b) shows the discrepancy between the actual response and the expected response extrapolated from a linear fit to the large pulse-height region of the curve. The non-linear behavior of the readout chain is well represented by the functional form:

$$E_{xtl} = B(p-p_0) + C(1-\exp(\alpha(p-p_0))) \qquad (21)$$

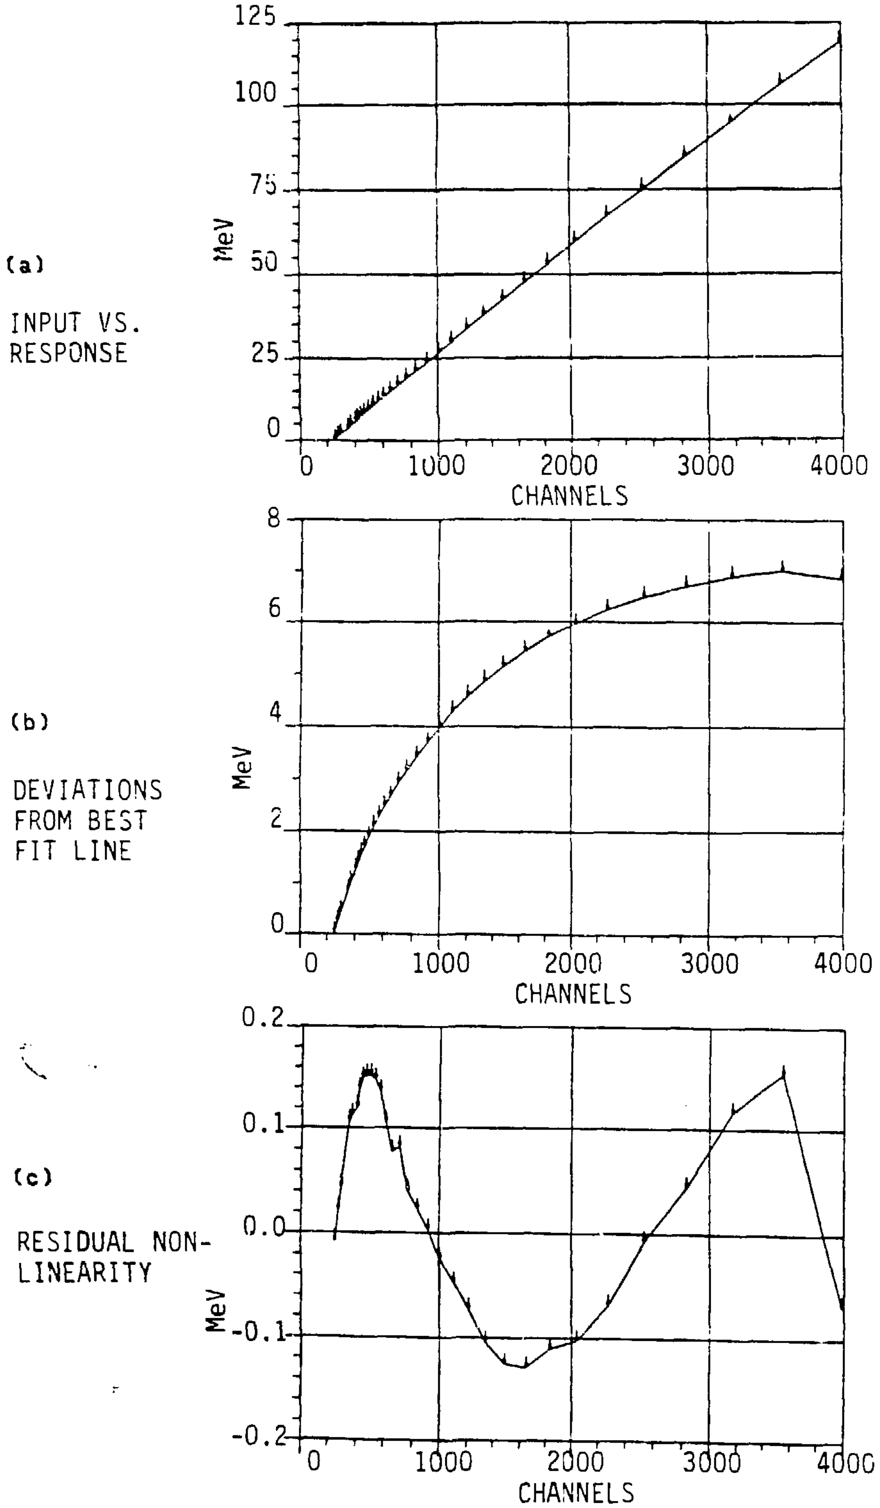

Figure 33: Results of the non-linearity correction analysis.

where E_{xtl} is the energy equivalent of the test pulse, p is the measured pulse-height, p_0 is the readout pedestal, and α, B, C are the parameters for a least-squares fit. It was found that the range of fit values for α is narrow, and that the quality of fit is insensitive to α within that range, so it is sufficient to set $\alpha = 0.00111$, the average value over a representative set of readout channels. This simplification allows the reduction of the fitting problem to a simple linear least-squares fit for B and C by an appropriate logarithmic transformation. The ratio B/C defines the shape of the calibration curve, but the absolute scale must be set by reference to a known physical process detected in the calorimeter. In order to provide a reasonable set of starting parameters for the refinement of the absolute scale, the B and C parameters were normalized to match the previously mentioned linear calibration at 70 MeV. Figure 33(c) shows that the residual non-linearity, after this exponential correction, is typically $\lesssim 0.2$ MeV for all energies less than 100 MeV.

Once the shape and approximate scaling of the calibration curves are determined, the final scale may be set by referring one point on the curve to an absolute standard. The "central" energy of the π^0-decay γ spectrum (68.9 MeV) was chosen as the reference point because it was the calibration energy closest to the 52.8 MeV region of interest for $\mu\to e\gamma$.

When the LH_2 target is illuminated with the π^- beam, π^0's are produced through the charge exchange reaction:

$$\pi^- \, p \to \pi^0 \, n. \qquad (22)$$

The recoil momentum, q_0, of the produced π^0's is 28.0 MeV/c. The π^0's subsequently decay (98% of the time) to two photons. Because the π^0 does not decay from rest, the produced photon energies are distributed with the so-called "box spectrum", which ranges from 55.1 to 82.7 MeV. The minimum opening angle between the the γ's is 156.5°. In the Crystal Box detector, these π^0 decays are thus observed as coincident photons, in opposite quadrants, each within the above energy range.

In the calibration procedure that follows, the relevant measured quantity is the clump energy sum, and specifically is not the individual energy of the crystals. The final calibration spectra are collected in a set of histograms containing one histogram associated with each crystal. A given histogram is incremented when the corresponding crystal is the HPHC of a clump. The fraction of the clump energy contained in the HPHC is required to be > 70%, so that the peak position of a calibration spectrum is essentially determined by the gain of the HPHC. The appropriate adjustment of gain for each crystal is then derived from the histogram associated with it. Because adjustments to the gains of its neighboring crystals subsequently affect the clump sum for a crystal, the above procedure must be iterated several times until a satisfactory convergence of the gain constants is achieved. The method given immediately below is that required for each of the passes through the data.

Each π^0 decay contributes a count to the box spectrum collected for each of two NaI(Tℓ) crystals. Unfortunately, only those counts near the endpoints of the box spectrum contain information about the absolute energy calibration. However, the observation of θ_{pp}, the opening angle between the photons, uniquely determines the two energies E_p^+, E_p^-, as follows:

$$E_p^+(\theta_{pp}) = E_0(1+\epsilon(\theta_{pp})) \tag{23}$$

$$E_p^-(\theta_{pp}) = E_0(1-\epsilon(\theta_{pp})) \tag{24}$$

where

$$E_0 = \frac{\sqrt{(m_0)^2 + (q_0)^2}}{2} \tag{25}$$

is the "central energy" of the box spectrum, m_0 is the π^0 mass, and

$$\epsilon(\theta_{pp}) = 1 - \frac{2}{(1-\cos\theta_{pp})\gamma^2} . \tag{26}$$

The γ^2 for the decaying π^0 is 1.0433. One may then construct

$$E_0^+ = E_p^+(1-\epsilon(\theta_{pp})) \tag{27}$$

and

$$E_0^- = E_p^-(1+\epsilon(\theta_{pp})) \tag{28}$$

for each event. Then all events contribute to a peak at E_0 for each of two crystals. Thus, by using the position resolution of the NaI(Tℓ) calorimeter, one makes more efficient use of the calibration information contained in the π^0 decay data.

A further consideration is the fact that the detected energy, for a given incident photon energy, will depend on the impact point of the photon due to shower leakage variations. This effect means that the expected central energy will be a function of the photon impact point. Therefore, the calorimeter crystals were assigned to 31 groups, corresponding to the 31 geometrically distinct crystal positions relative to a LH_2 target placed exactly at the center of the Box. To assure optimum agreement of the Monte Carlo events with the data, the observed central energy spectrum, for each crystal in a symmetry group, is matched to the predicted spectrum for that group. The matching was done by scaling the B and C calibration parameters to align the observed peak with the predicted one, while keeping the ratio B/C fixed. Figure 34 shows the Monte Carlo prediction of the central energy spectrum for group 25 and the agreement, after calibration, of the observed spectra for two crystals in that symmetry group.

There are several checks available to assure that the final energy calibration is consistent over the dynamic range of interest. In all of the comparisons below, events were selected for which the HPHC of the clump is within the central 5 X 6 array of one of the quadrants. Figure 35 shows the predicted and measured detector response to the monoenergetic γ at 129.4 MeV from $\pi^- p \rightarrow n\gamma$. Although the measured resolution of 7% (FWHM) is slightly worse than predicted, the absolute peak position is clearly in agreement to better than 1%.

Figure 36 compares the Monte Carlo prediction to the observed spectra for π^0 decay events with an opening angle > 175°. In these events, the two photons are almost parallel and anti-parallel, respectively, to the

(a) Predicted Central Energy for $\pi^0 \rightarrow \gamma\gamma$

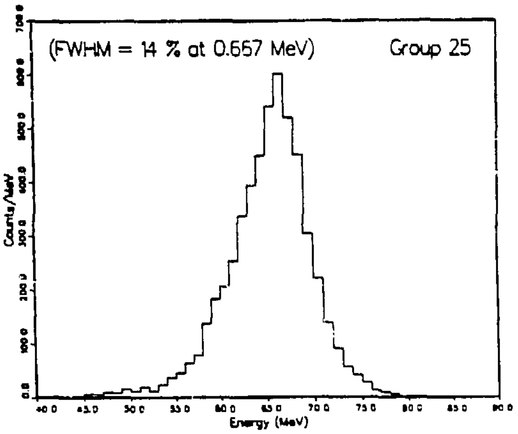

(b) Detected Central Energy

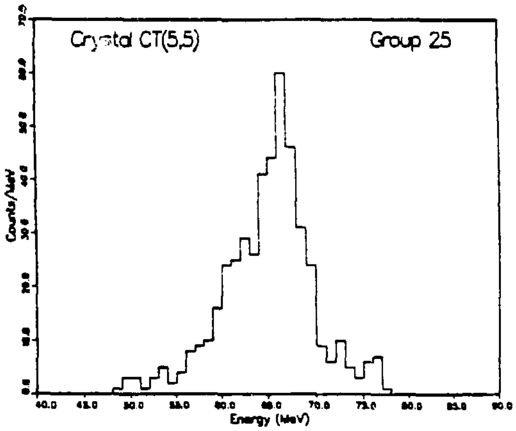

(c) Detected Central Energy

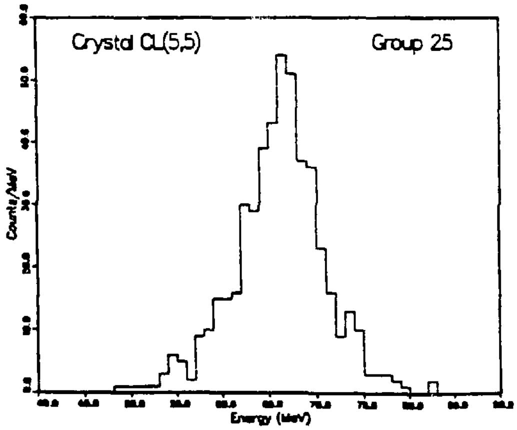

Figure 34: The predicted central energy (a) compared with the observed spectra for two group 25 crystals (b),(c).

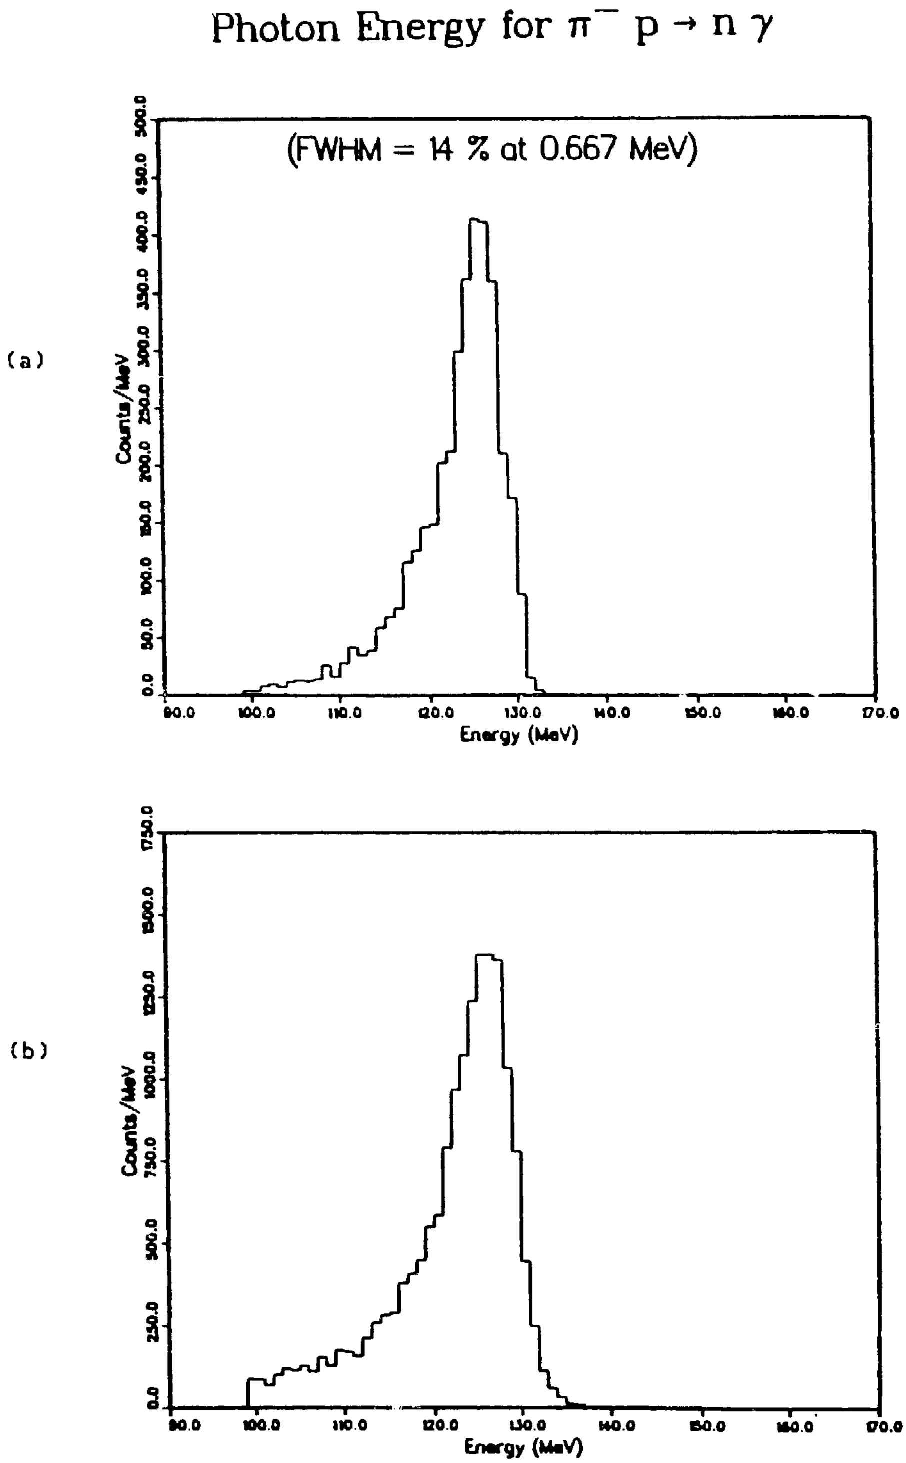

Figure 35: Predicted (a) and measured (b) γ spectra for $\pi^- p \to n\gamma$.

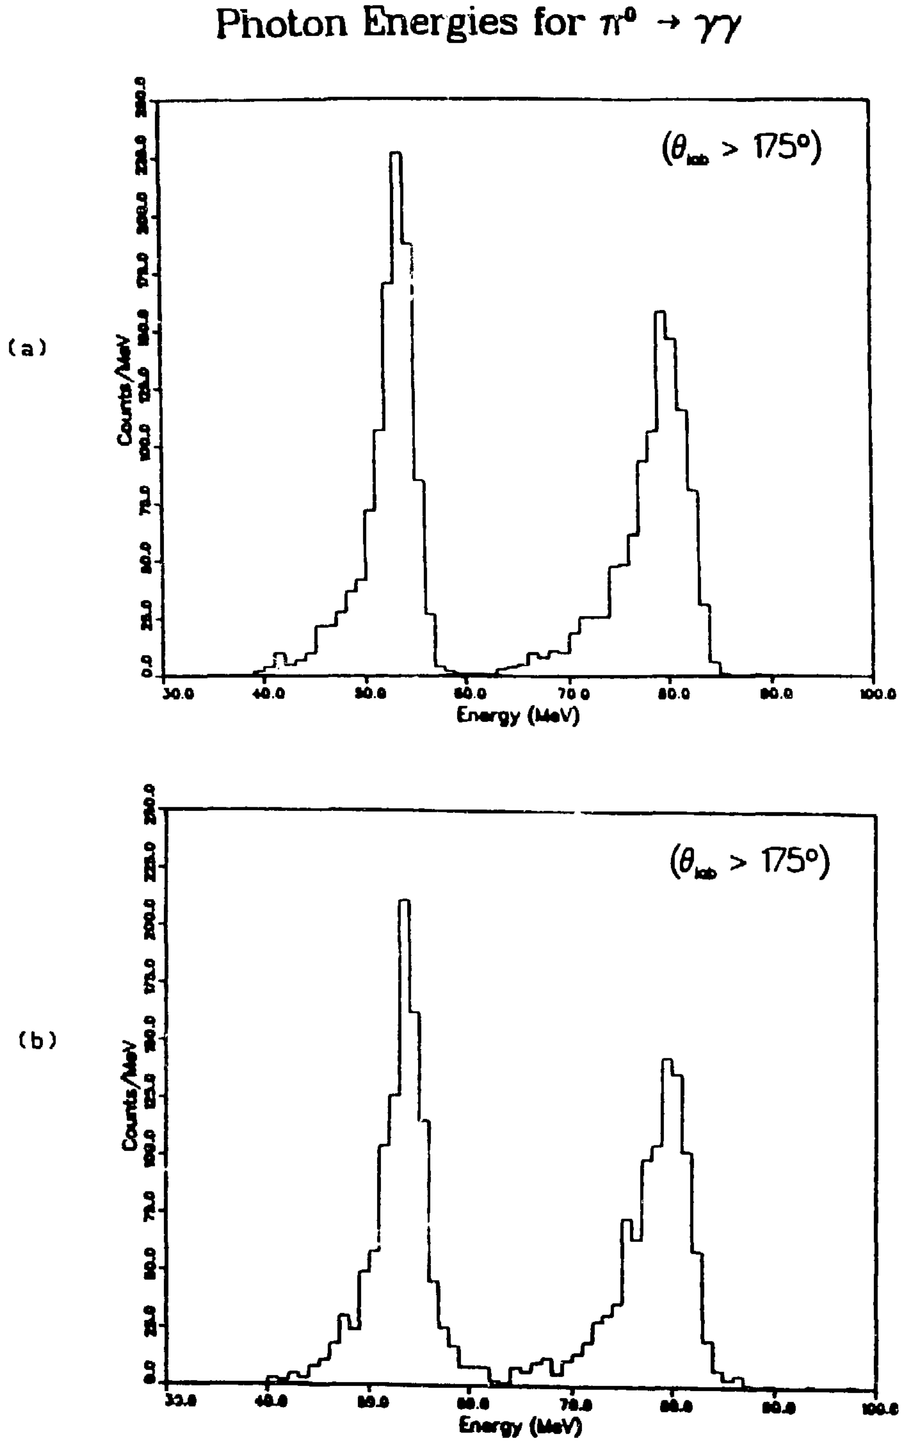

Figure 36: Predicted (a) and measured (b) γ spectra for π^0 decays with large opening angles.

π^0 momentum, and thus differ maximally in energy. This selection of events leads to an observed energy spectrum with two peaks near the endpoints of the box spectrum. Of particular interest for the $\mu \to e\gamma$ experiment is the fact that the lower peak, near 55 MeV, is quite well simulated in position and shape. The upper edge of this peak is broadened due to the finite angular acceptance of the sample. Therefore, the peak width of 8% (FWHM) may be taken as an upper limit to the detector resolution at 55 MeV.

The low energy photon response is checked by reference to the 4.4 MeV line from the Pu-α-Be source. The Monte Carlo program did not fully simulate the spectrum of radiation produced by the source. Instead, monoenergetic photons were thrown at 4.4 MeV. The comparison of the predicted 4.4 MeV photon response to the observed source spectrum is given in Figure 37. The peak positions agree to within 0.2 MeV, and the photon calibration is therefore seen to be satisfactorily consistent over a dynamic range of 30.

The only well-defined energy spectrum available, in the data, to check the electron calibration is that of the Michel positrons. Type II events were selected because the stopped muons are largely depolarized in the target material, thus minimizing any asymmetry in the angular distribution of the positron energy. Figure 38 shows that the simulation and the data are in agreement to 1% for the electron calibration. Note that some residual pile-up is evident in the actual data spectrum, while no pile-up was added to the events in the predicted spectrum. This residual pile-up effect was simulated in the spectra, described below, used to analyze the $\mu \to e\gamma$ data.

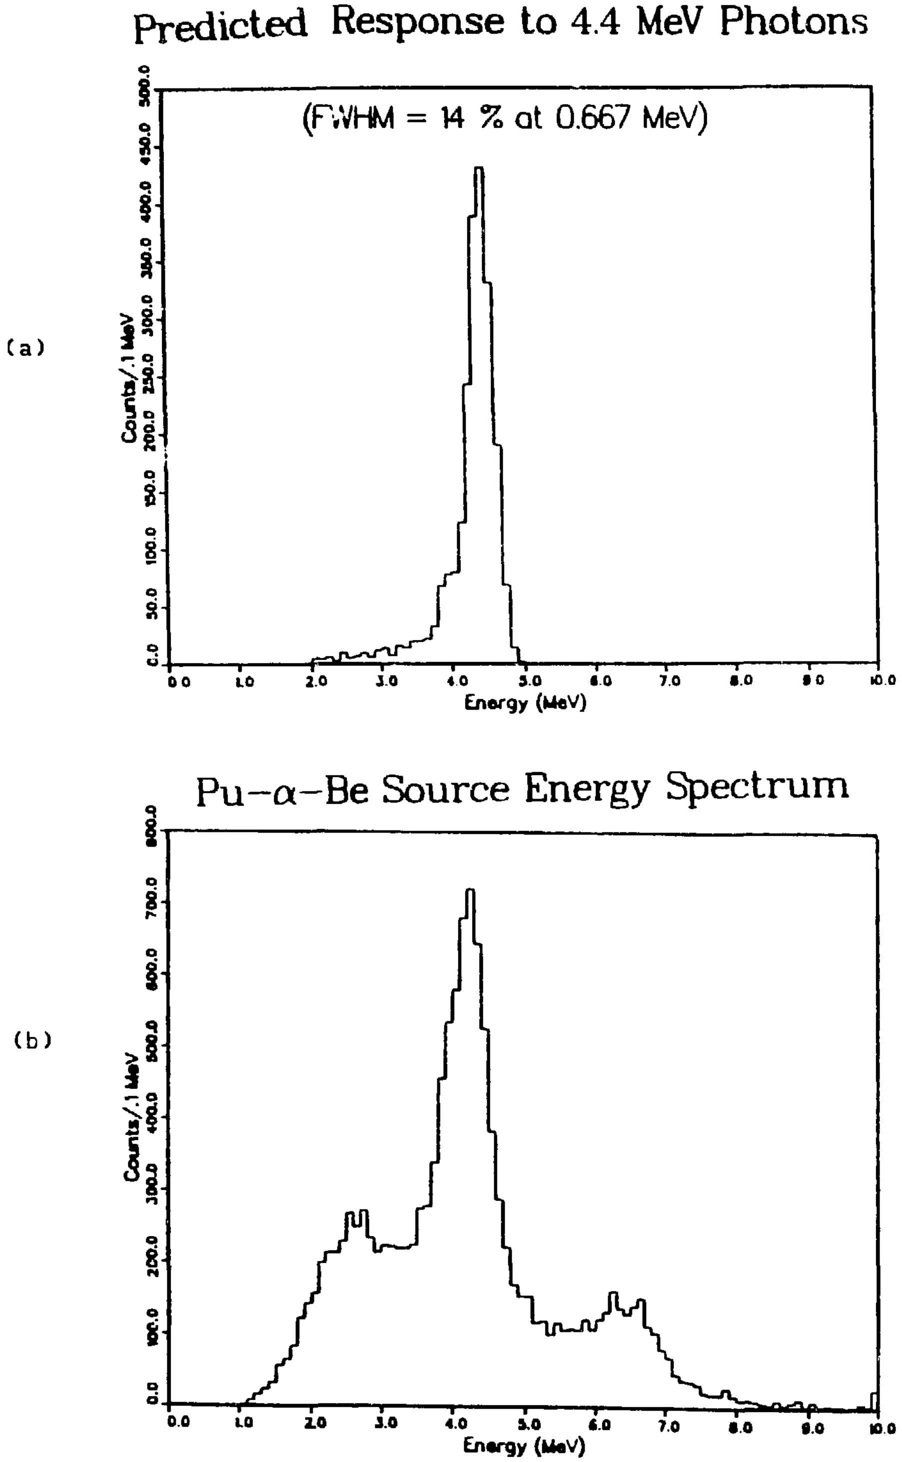

Figure 37: Predicted response of calorimeter to 4.4 MeV γ (a) and measured Pu-α-Be source spectrum (b).

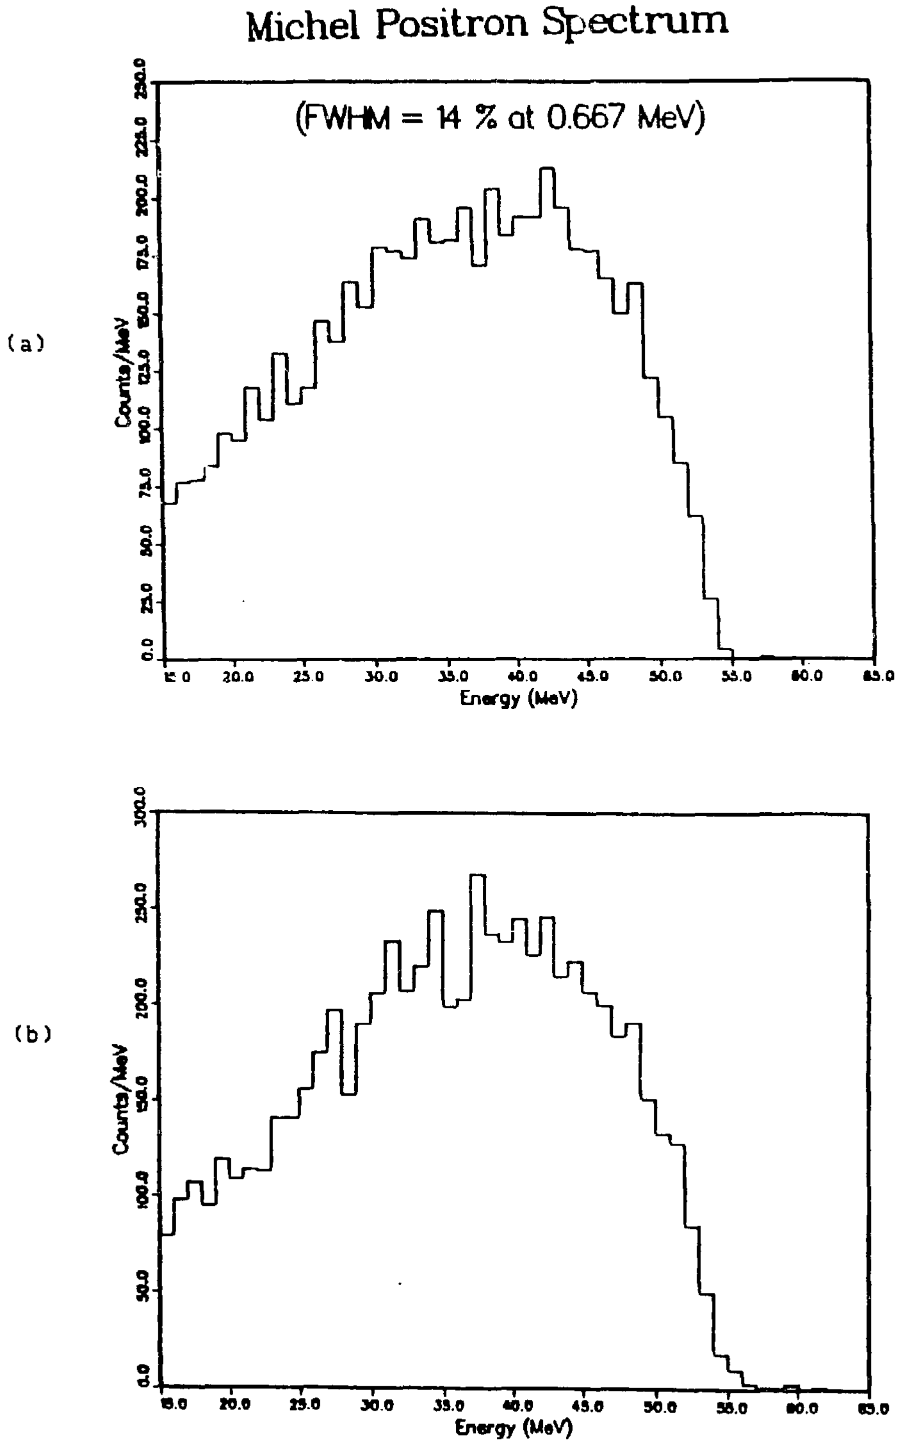

Figure 38: Comparison of predicted (a) and detected (b) Michel positron energy spectra from type II tapes.

5.2.2.2 Hodoscope

The energy calibration of the plastic trigger counters is accomplished using the type I electron tapes. Such a calibration tape is first analyzed to determine the pedestals of all scintillator readout channels. The electron trajectory through the scintillator, for each single-hit event, is estimated from the time difference between the two PMT's connected to that scintillator, assuming the electron originated in the center of the I-counter. This estimate determines the impact position (on the counter), z, and the path length (through the counter), l, of the electron track. The pedestal-subtracted responses of the upstream and downstream PMT's are then taken to be of the form:

$$P_u = G_u'E(l)\exp(-\lambda(\frac{L}{2}+z)) \tag{29}$$

$$P_d = G_d'E(l)\exp(-\lambda(\frac{L}{2}-z)) \tag{30}$$

where $z = 0$ at the center of the counter, L is the length of the counter, $E(l)$ is the ionization energy loss of the electron, and λ is the light attenuation coefficient. The factor $\exp(-\lambda L/2)$ can be absorbed into G_u', G_d':

$$P_u = G_uE(l)\exp(-\lambda z),\ P_d = G_dE(l)\exp(+\lambda z) \tag{31}$$

Over the range of path lengths possible in the Crystal Box geometry, the energy deposited may be represented as

$$E(l) = E_0 l \tag{32}$$

where E_0 is the dE/dx loss per unit path length, determined by Monte Carlo calculations to be $\simeq 1.74$ MeV/cm. It is easily seen that a histogram of the the combination

$$\frac{(P_u P_d)^{1/2}}{l} = E_0 (G_u G_d)^{1/2} \tag{33}$$

gives a position-independent determination of the quantity $(G_u G_d)^{1/2}$. Similarly, the quantity

$$\frac{P_u}{P_d} = \frac{G_u}{G_d} \exp(-2\lambda z) \tag{34}$$

may be plotted versus z, so that a fit to the curve determines λ and G_u/G_d. From knowledge of $(A)^{1/2} \equiv (G_u G_d)^{1/2}$ and $B \equiv G_u/G_d$, the individual gains are determined as:

$$G_u = (AB)^{1/2}, \; G_d = (A/B)^{1/2} \tag{35}$$

The values of λ for a given scintillator, as determined by successive calibration tapes, give a quantitative monitor of any deterioration of the light transmission properties of that scintillator. The quantity

$$\left(\frac{P_u P_d}{G_u G_d}\right)^{\frac{1}{2}} = E(l) \tag{36}$$

is then used as the measure of the energy deposited in a trigger counter which is independent of the electron impact position.

To define the calibration of the guard counters, the expected energy spectrum due to Michel positrons from the target is generated using the

Monte Carlo program. The peak of the observed spectrum for type II tapes is then matched to that of the predicted spectrum, thereby establishing the calibration of each guard counter readout channel. This method is sufficiently accurate for the guard counter calibration because these counters are only used to veto prospective photon quadrants and they never contribute to the energy sum of an electron hit.

5.3 MUON-DECAY DATA

Approximately 150 data tapes were collected, containing 1.3×10^6 candidate $\mu \rightarrow e\gamma$ events, from a total of 2.2×10^{11} muons stopped in the target during the experimental live time. The event sample was reduced in two phases. The 2.84×10^4 events remaining after the first phase were used to characterize the experimental backgrounds. A maximum-likelihood analysis was performed on the 192 candidates surviving the second phase to calculate the most probable number of $\mu \rightarrow e\gamma$ events in the residual sample.

5.3.1 Reduction of Candidate Events

As explained above in section 3.4.5, the experimental signature for a $\mu \rightarrow e\gamma$ event is a time coincident electron-photon pair, each with 52.8 MeV, and an opening angle of 180°. Therefore, the reduction of candidate events consists of cuts on the observed timing, energy, and relative angle of prospective e-γ pairs, as well as cuts to eliminate undesirable instrumental effects. A preliminary description of the off-line pile-up rejection is first presented, followed by the data-reduction procedure.

5.3.1.1 Pile-up Rejection

The on-line pile-up detection system was not useful for the January 1984 run cycle, so the pile-up was eliminated, as much as possible, by off-line software cuts. The NaI(Tℓ) timing information was used to look for any indication of out-of-time energy deposition. A clump was eliminated as a candidate particle if any of the following conditions were true:

1. the HPHC did not fire its CFD
2. a crystal in the neighbor set fired its CFD more than 5 ns out of time with that of the HPHC
3. an NS crystal with > 7 MeV did not fire its CFD within 5 ns of the HPHC

The effect of these cuts on the experimental acceptance are discussed in section 5.6 below.

5.3.1.2 Phase I Reduction

The phase I data reduction is summarized in Table 5. The cuts are conceptually described as follows:

1. Cut 1 required less than 0.25 MeV for the energy sum of all the photon-quadrant plastic counters, including the guard counters and the nearest counter from both neighboring quadrants. It also required the time difference between the $\underline{e}$ and $\underline{\gamma}$ signals, used in the $\underline{e} \bullet \underline{\gamma}$ trigger coincidence, to be positive and less than 4 ns.

TABLE 5

Summary of Phase I Data Reduction

2.2×10^{11} μ-stops 1.3×10^{6} eγ triggers

Data Cuts	Number of Surviving Events
1. scintillator veto $0 < T(\gamma)-T(e) < 4$ ns	4.8×10^{5}
2. $E_c(e) \geq 35$ MeV $E_c(\gamma) \geq 35$ MeV $\lvert\Delta t_{ep}\rvert \leq 4.5$ ns no pile-up	1.2×10^{5}
3. electron track reconstructed to target	9.4×10^{4}
4. $E_e \geq 35$ MeV $E_p \geq 35$ MeV $E_e + E_p + \lvert p_e+p_p\rvert \leq 135$ MeV	5.95×10^{4}
5. $\theta_{ep} \geq 140^0$	4.08×10^{4}
6. no reconstructed drift chamber track for γ	2.84×10^{4}

Because the timing of the $\underline{\gamma}$ signal was explicitly set later than that of the $\underline{e}$ signal when the trigger was set up, the second requirement merely reduces the window for accidental coincidences, by a factor of two, without eliminating any real coincidences.

2. Cut 2 required the NaI(Tℓ) energy sums, for both the electron and photon clumps, to be at least 35 MeV with no detected pile-up. In addition, the time difference between the electron and photon clumps is required to be no more than 4.5 ns. (Note that this requirement is made on the clumps, while the timing requirement in cut 1 is made on the trigger).

3. With the data sample reduced by a factor of 10, the track-reconstruction program may be used to eliminate electron clumps. A track pointing within 15 cm of the electron clump was successfully reconstructed back to the target for about 90% of the remaining events.

4. Cut 4 required a minimum of 35 Mev for the electron and photon clump energies, E_e, and E_p, as well as a maximum of 135 MeV for their sum plus the "missing momentum" of the two particles. (Monte Carlo predictions show that the latter requirement on the residual vector momentum of the event is quite conservative and will not cut any real $\mu \to e\gamma$ events).

5. Cut 5 required the opening angle, θ_{ep}, to be at least 140°.

6. Finally, cut 6 searched for evidence that the photon was in fact a misidentified electron. The acceptable χ^2 of fitted tracks was relaxed, and the angle of the track with the target surface was allowed to go to 0°, to find any reconstructable tracks pointing within 15 cm of the photon clump position.

A sample of 28,350 events remained after these cuts were imposed.

Figure 39 shows the distribution of the electron-minus-photon time difference for this sample. A prompt peak, clearly visible on top of a flat spectrum due to random coincidences, is consistent with the photon timing resolution of the NaI(Tℓ) calorimeter. The peak is from inner-bremsstrahlung events that satisfy the above cuts, and any $\mu\to e\gamma$ events that may be in the sample. The "shoulders" of the flat background are rounded for two reasons. First, the timing requirement in cut 1 removed events, near the upper and lower edges of the spectrum, that would have passed the cut 2 timing requirement. Second, a "time-of-flight correction" was made, after cut 2, to the timing of all clumps. Because the timing of all plastic counters and NaI(Tℓ) crystals was calibrated relative to the I-counter, an adjustment to the clump times must be made, to account for the difference in distance from the clump position to the I-counter, as opposed to that from the clump to the reconstructed electron origin on the target. The electron origin was not available to the analysis until the track reconstruction was turned on in cut 3. After the time-of-flight correction was made, some events had a value for $|\Delta t_{e\gamma}|$ exceeding 5 ns. The extreme out-of-time events from this sample were used to characterize the random background for the maximum likelihood analysis.

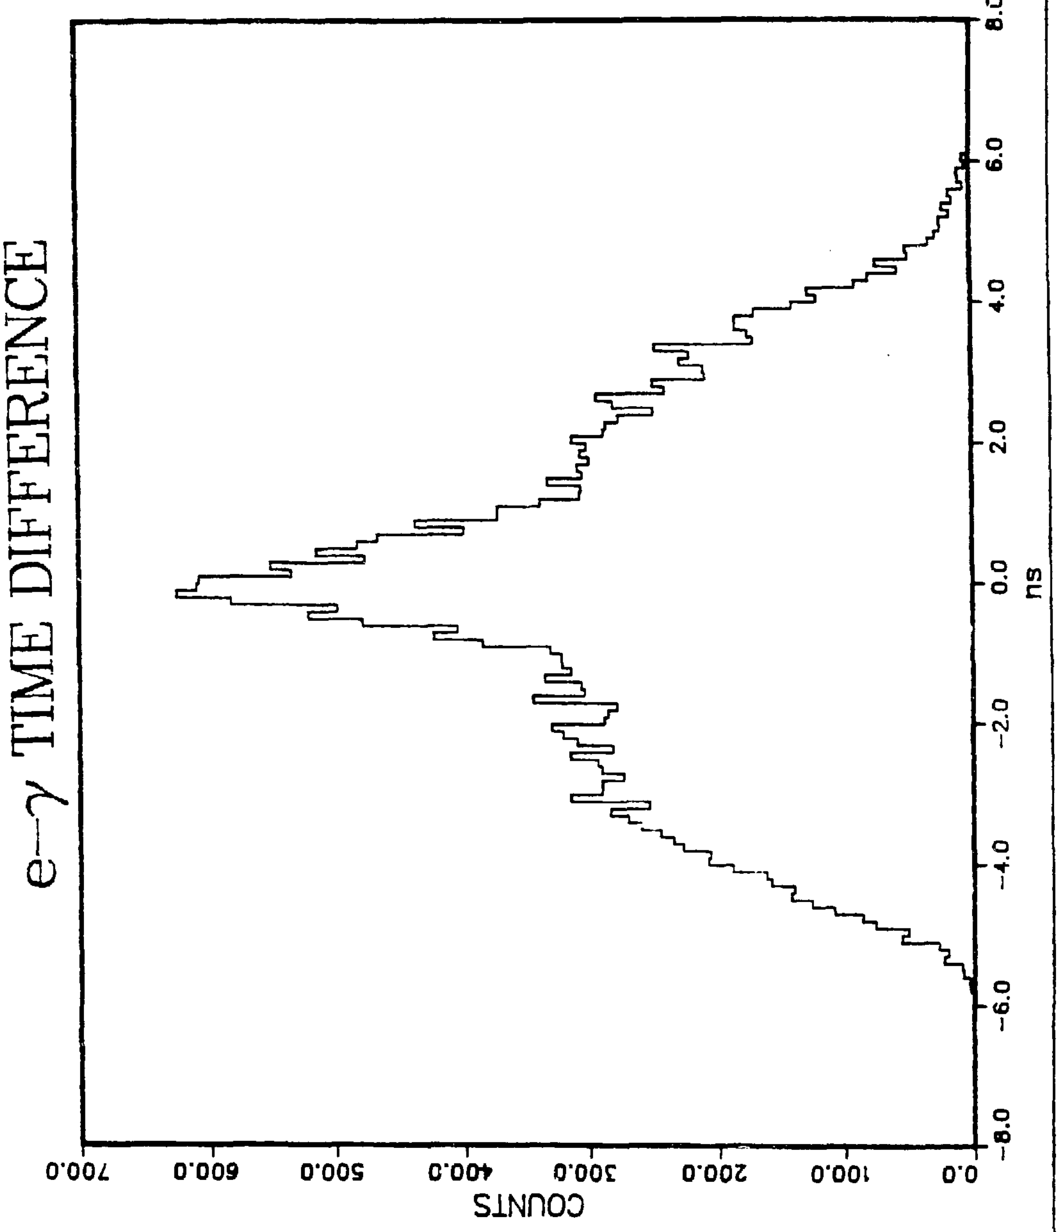

Figure 39: Time difference of the electron-photon pairs remaining after phase I cuts.

5.3.1.3 Phase II Reduction

Once a suitable set of background events had been selected, the data sample could be further reduced in preparation for the final statistical analysis. The minimum opening angle accepted was increased to $\theta_{ep} \gtrsim 166^0$ and the relative timing requirement was tightened to $|\Delta t_{ep}| \lesssim 2$ ns. The phase II energy cuts for the photon were:

$$47.5 \text{ MeV} \lesssim E_p \lesssim 58 \text{ MeV} .$$

The electron energy, E_e, was also restricted to a maximum of 58 MeV. The minimum threshold of 44 Mev was imposed on the electron NaI(Tℓ) clump energy. The lower cut was applied in this manner, for technical reasons, to provide a consistent NaI(Tℓ) quadrant threshold for use in the calculation of the experimental acceptance. A sample of 192 events remained after the phase II cuts.

5.3.2 Maximum Likelihood Analysis

Given the energy, timing, and position resolutions of the Crystal Box detector, the data sample collected from 2 x 10^{11} μ-stops is expected to contain random coincidence events, and (possibly) inner-bremsstrahlung events, that are indistinguishable from μ→eγ within the limitations of these resolutions. Such a data sample is said to be "background-limited". In principle, the Monte Carlo model could be used to predict the number of backgrounds events expected to survive the data reducing cuts. However, the statistical uncertainty in the predicted background level will restrict the result of such an analysis to an upper limit for the number of observed events. To set the most stringent upper limit hat is rigorously allowed, the statistical analysis technique of "maximum likelihood" is employed.

The residual data sample is presumed to consist mostly of random coincidence events, with small admixtures of prompt background events ($\mu \rightarrow e\gamma\nu\bar{\nu}$) and $\mu \rightarrow e\gamma$ events. Each event is characterized by a set of $n = 4$ observables, o_i, and each event type is defined by its probability density distributions for these observables. For this data, the observables are:

$$o_i \in \left\{\theta_{ep}, \Delta t_{ep}, E_e, E_p\right\} \tag{37}$$

and the event types are associated with the following distributions: $P_i(o_i)$ for $\mu \rightarrow e\gamma$, $Q_i(o_i)$ for $\mu \rightarrow e\gamma\nu\bar{\nu}$, and $R_i(o_i)$ for random coincidences. The products

$$P(\{o_i\}) = \prod_1^4 P_i(o_i) \tag{38}$$

$$Q(\{o_i\}) = \prod_1^4 Q_i(o_i) \tag{39}$$

$$R(\{o_i\}) = \prod_1^4 R_i(o_i) \tag{40}$$

give the probability for an event of the corresponding type to be observed with the particular set $\{o_i\}$. For a sample of N events, the likelihood function

$$L(\alpha,\beta) = N^{-} \prod_j \left[\alpha P(\{o_i\}_j) + \beta Q(\{o_i\}_j) + (N-\alpha-\beta) R(\{o_i\}_j) \right] \tag{41}$$

is the probability that the sample consists of α $\mu\to e\gamma$ events, β $\mu\to e\gamma\nu\bar{\nu}$ events, and $(N-\alpha-\beta)$ random events. The values α_{max} and β_{max} that maximize the likelihood function determine the most probable content of the respective event types in the sample. For simplicity of calculation, it is clearly sufficient to maximize the reduced likelihood function

$$L_0(\alpha,\beta) \equiv \frac{L(\alpha,\beta)}{L(0,0)} \tag{42}$$

because $L(0,0)$ is a constant scale factor.

5.3.2.1 Generation of Probability Density Distributions

The ingredients required to carry out this prescription are the P_i's, Q_i's, and R_i's corresponding to the three event types. The P_i's, for $\mu\to e\gamma$, and the Q_i's, for $\mu\to e\gamma\nu\bar{\nu}$, were generated from two large samples of Monte Carlo events. These events were analyzed with the same cuts as the data. The resulting probability distributions for $\mu\to e\gamma$ are shown in Figure 40, and those for $\mu\to e\gamma\nu\bar{\nu}$ are shown in Figure 41. The probability distributions of Δt_{ep}, for both event types, were assumed to be a gaussian of a width consistent with that of the prompt peak observed in the phase I data sample (Figure 39). The R_i's (Figure 42), for the random background, were taken from those phase I events with $|\Delta t_{ep}| >$ 1.2 ns. The timing distribution for random events was assumed to be independent of Δt_{ep}.

(a) Predicted Electron Energy
for $\mu \to e\gamma$

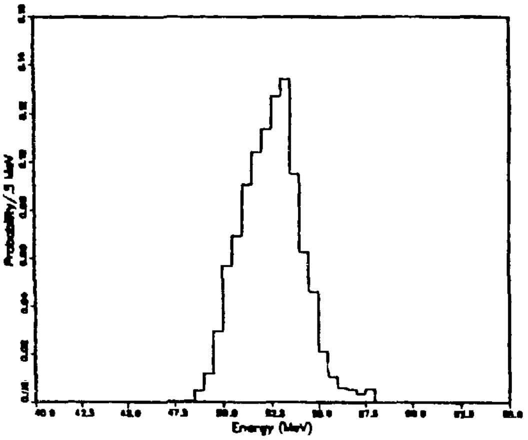

(b) Predicted Photon Energy

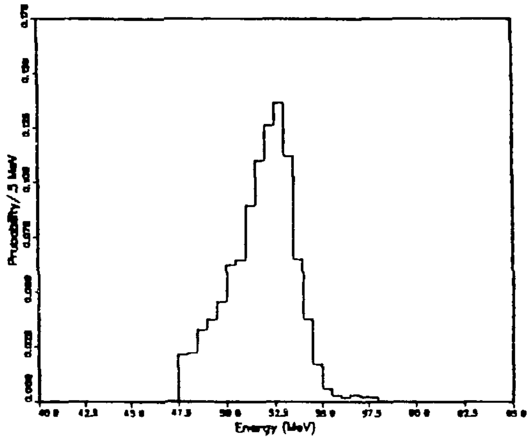

(c) Predicted Opening Angle

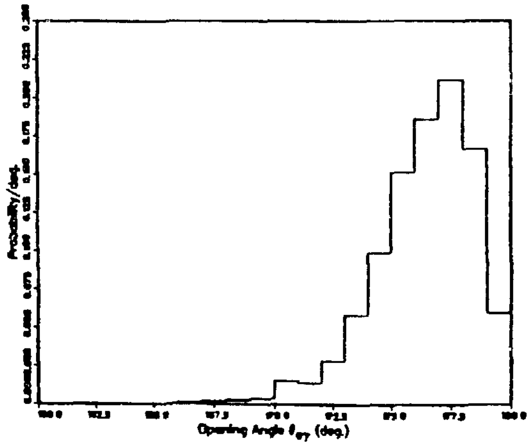

Figure 40: Probability distributions used to characterize eγ events in the maximum-likelihood analysis.

(a) Predicted Electron Energy for $\mu \to e\gamma\nu\bar{\nu}$

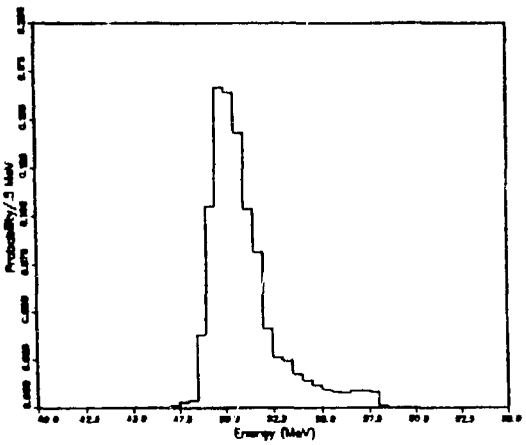

(b) Predicted Photon Energy

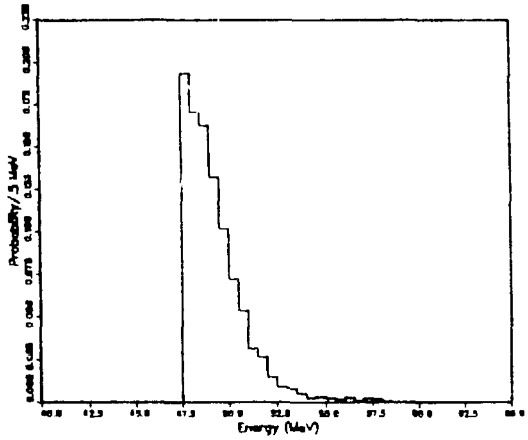

(c) Predicted Opening Angle

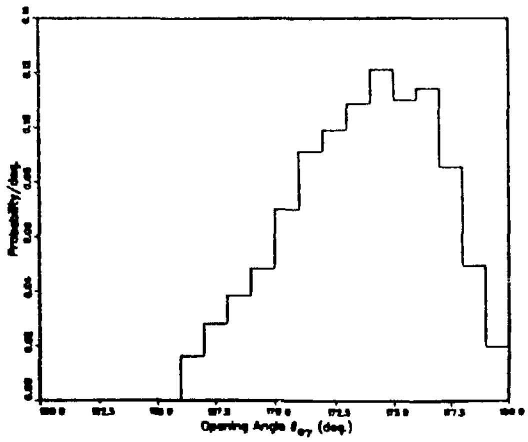

Figure 41: Probability distributions used to characterize $e\gamma\nu\bar{\nu}$ events in the maximum-likelihood analysis.

(a) Measured Electron Energy for Randoms

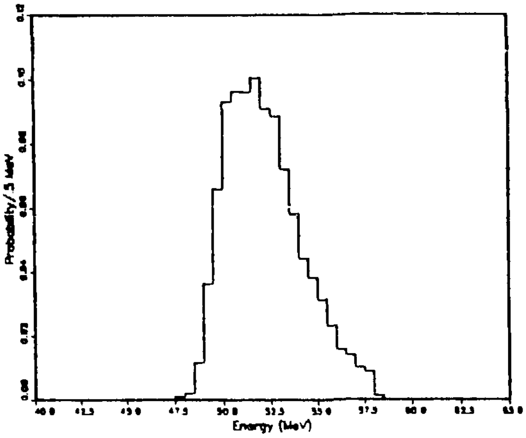

(b) Measured Photon Energy

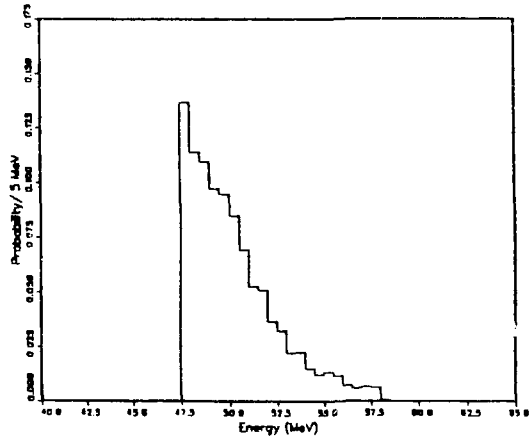

(c) Measured Opening Angle

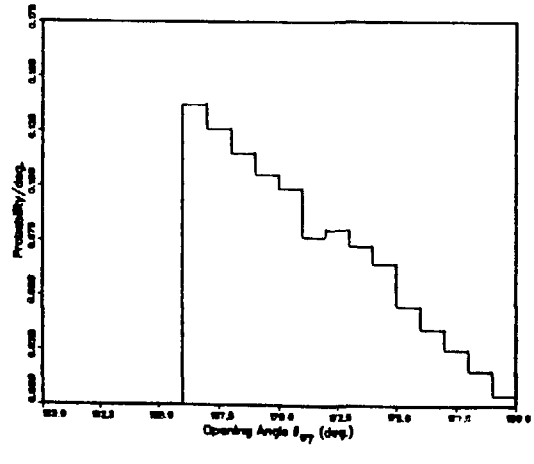

Figure 42: Probability distributions used to characterize random events in the maximum-likelihood analysis.

5.3.2.2 Calculation of α_{max} and β_{max}

The calculated distributions of P, Q, and R for the Monte Carlo $\mu\rightarrow e\gamma$ events are shown in Figure 43. Figure 44 shows the same three plots for Monte Carlo $\mu\rightarrow e\gamma\nu\bar{\nu}$ events. The abscissa of each plot is the negative logarithm of the probability product for the particular event, so that the events increase in probability as the abscissa decreases. These plots provide a check of the computer algorithm used to calculate the likelihood function. Roughly speaking, the $\mu\rightarrow e\gamma$ events are 10 times more likely to be $\mu\rightarrow e\gamma$ events than to be either $\mu\rightarrow e\gamma\nu\bar{\nu}$'s or randoms, and an analogous statement holds for the $\mu\rightarrow e\gamma\nu\bar{\nu}$ events. The distributions of P, Q, and R for the final data sample of 192 events are shown in Figure 45. As expected, very few events have a significant probability to be $\mu\rightarrow e\gamma$'s as compared to the distribution in Figure 43(a). A greater number have probabilities Q near that expected for $\mu\rightarrow e\gamma\nu\bar{\nu}$'s, and the R distribution is strongly peaked to the left, indicating the sample consists mostly of randoms.

The individual probability products were appropriately combined, and the reduced likelihood function explored as a function of α and β, with the result given in Figure 46. The maximum value of L_0 is obtained for α_{max} = 2.5, and β_{max} = 16. The 90% confidence level (CL) for α is the value α_{cl} satisfying

$$\int_0^N \int_0^{\alpha_{cl}} L_0 \, d\alpha \, d\beta = 0.9 \times \int_0^N \int_0^N L_0 \, d\alpha \, d\beta \ . \qquad (43)$$

The first integral was approximated by a sum of L_0 over β from 0 to 40, in steps of one, to form the projected L_0 distribution shown in Figure

(a) Predicted $\mathcal{P}$ for $\mu \to e\gamma$ Events

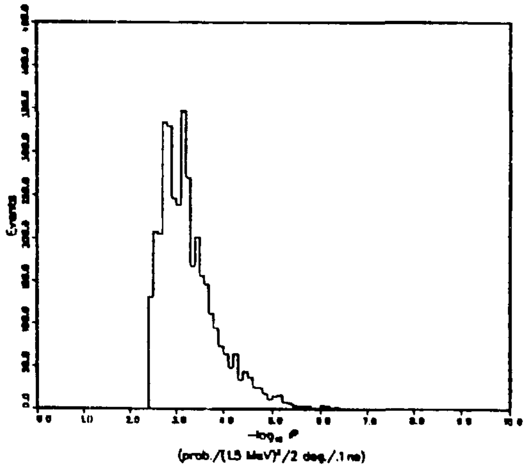

(b) Predicted $\mathcal{Q}$

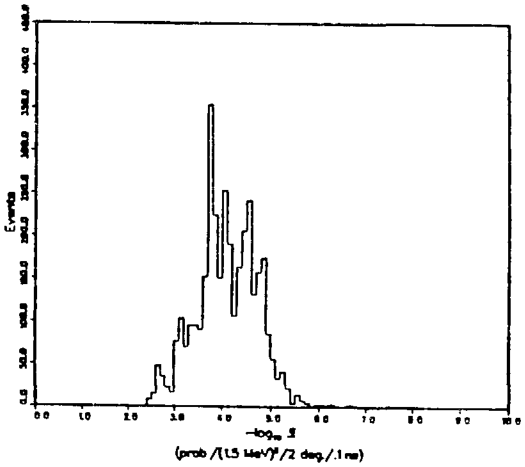

(c) Predicted $\mathcal{R}$

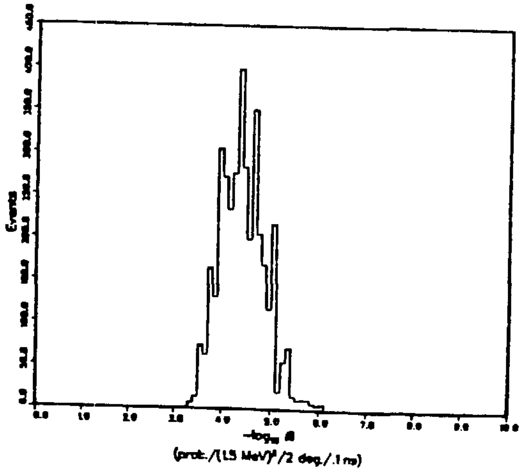

Figure 43: Probability-product distributions for Monte Carlo $e\gamma$ events.

(a) Predicted $\mathcal{P}$ for $\mu \rightarrow e\gamma\nu\bar{\nu}$ Events

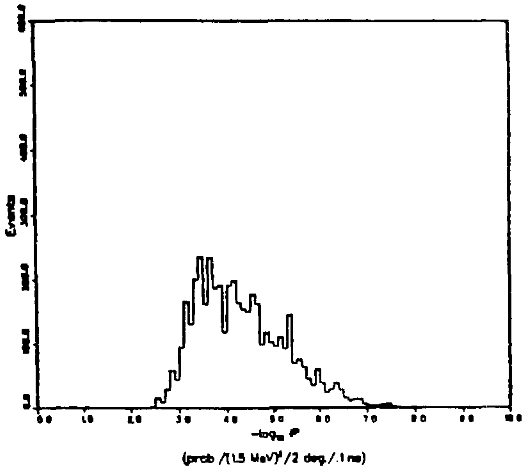

(b) Predicted $\mathcal{Q}$

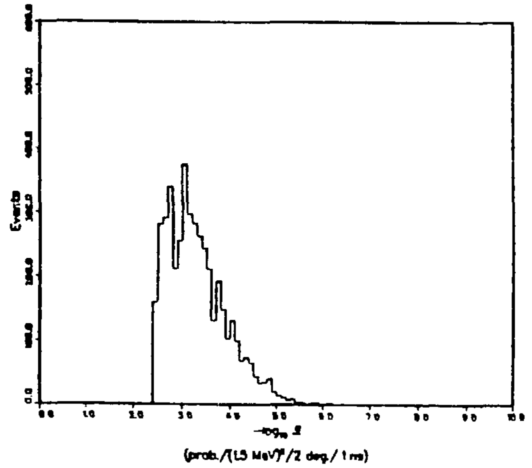

(c) Predicted $\mathcal{R}$

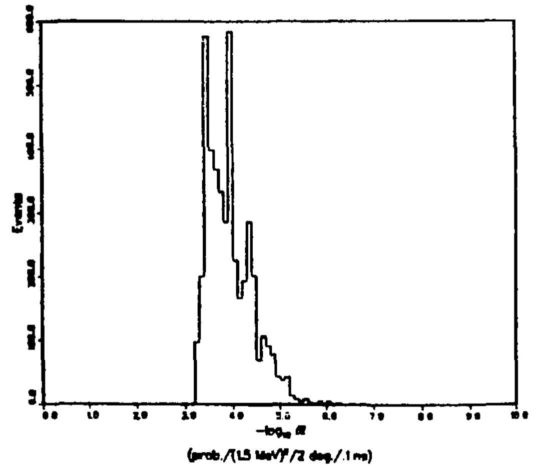

Figure 44: Probability-product distributions for Monte Carlo $e\gamma\nu\bar{\nu}$ events.

(a) Measured $\mathcal{P}$ for Final Data Sample

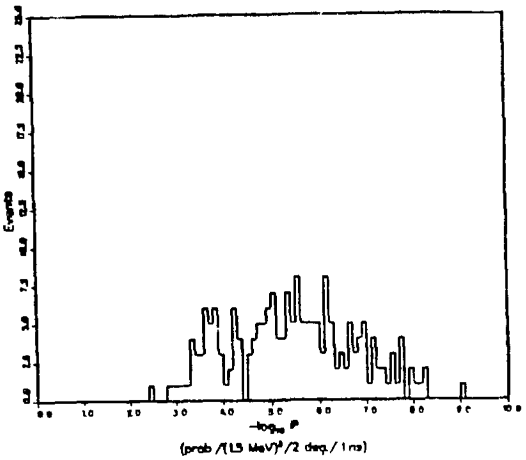

(b) Measured $\mathcal{Q}$

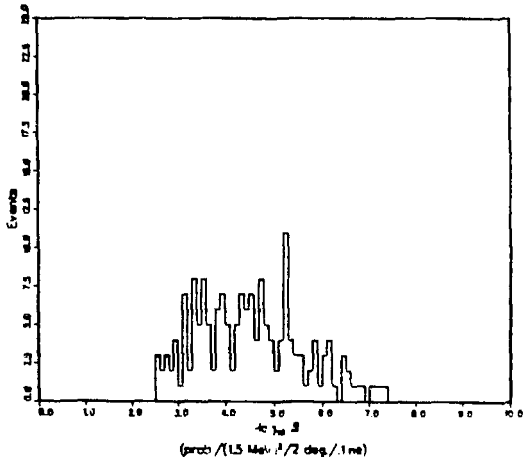

(c) Measured $\mathcal{R}$

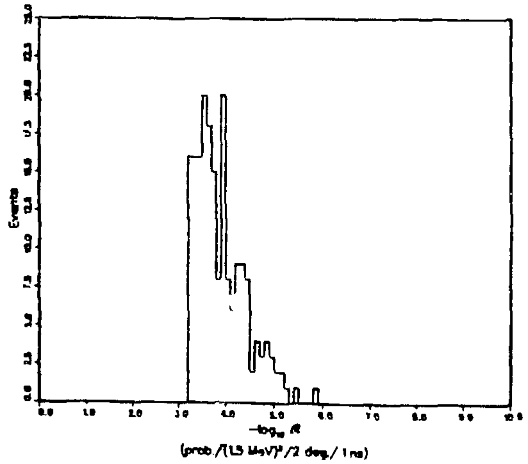

Figure 45: Probability-product distributions for events surviving the phase II data cuts.

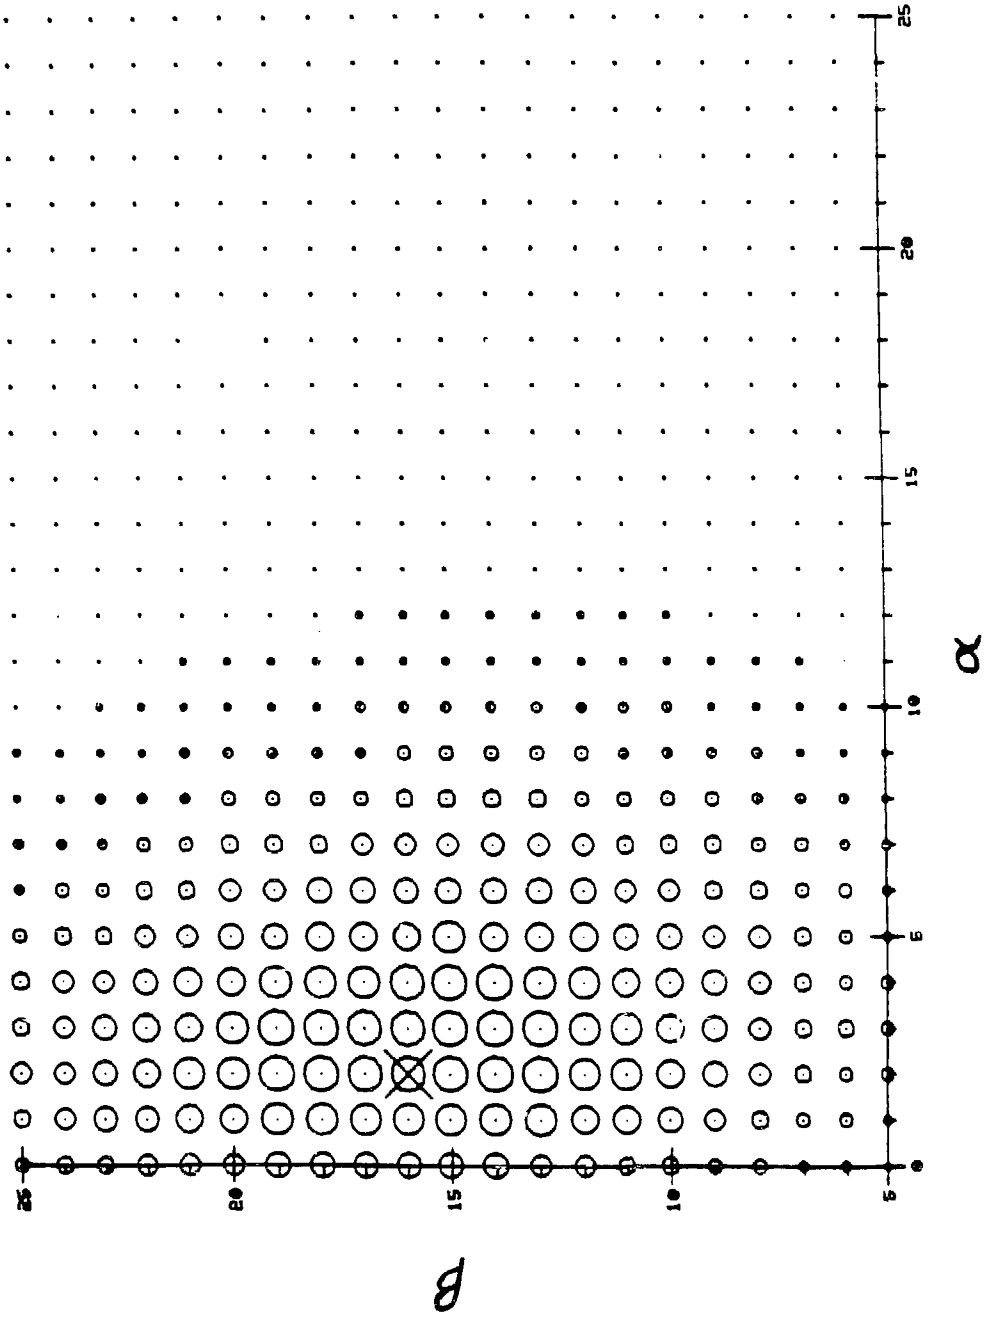

Figure 46: Plot of the maximum-likelihood function, L_0, as a function of α and β.

47. The L_0 function is effectively zero for all $\beta > 40$, and it varies smoothly over the entire range of beta, so this approximation is reasonable. The projected distribution was then analyzed to find the minimum upper limit, for a sum over α, that would include 90% of the sum over α from 0 to 25. Again the projected L_0 function is smoothly varying and goes to zero rapidly for $\alpha > 25$, so the approximation is a good one. The 90% CL for α, obtained in this manner, is $\alpha_{cl} = 9.5$. From the projection of L_0, summed over α from 0 to 25, shown in Figure 48, the 90% CL for β is seen to be $\beta_{cl} = 26$.

5.4 EXPERIMENT NORMALIZATION AND ACCEPTANCE

The experimental acceptance was calculated from an analysis of the Monte Carlo $\mu\to e\gamma$ events, by applying the same cuts that were imposed on the data. The fraction of events that passed the trigger requirements, the phase I cuts, and the phase II cuts, was f = 18.5%. An additional fraction of the events would actually be excluded from the data sample due to pile-up effects and due to the fact that photon quadrants could be vetoed, in the trigger, by an accidental coincidence with a Michel positron. To estimate these effects, a tape of data was analyzed for which the trigger was generated randomly within the macropulses of the beam and the stopping rate was 3×10^5 μ^+/s. The result of this analysis was a reduction of the acceptance by another 17% to the final value of

$$\epsilon = 0.83f = 0.83(18.5\%) = 15.4\% \,. \qquad (44)$$

A check of the normalization is provided by the number of $\mu\to e\gamma\nu\bar{\nu}$ decays observed in the data sample. The sample of Monte Carlo $\mu\to e\gamma\nu\bar{\nu}$

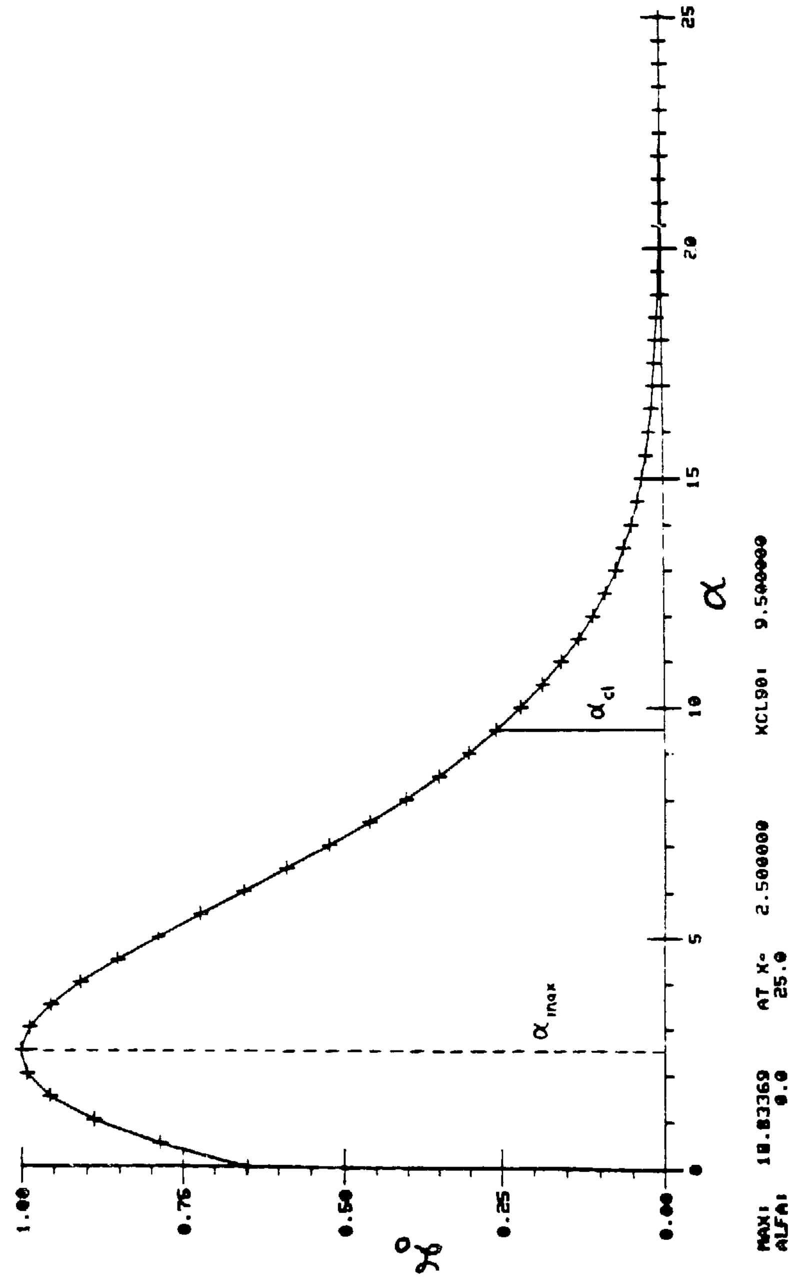

Figure 47: Plot of L_0, summed over $\beta = 0$ to 40, as a function of α.

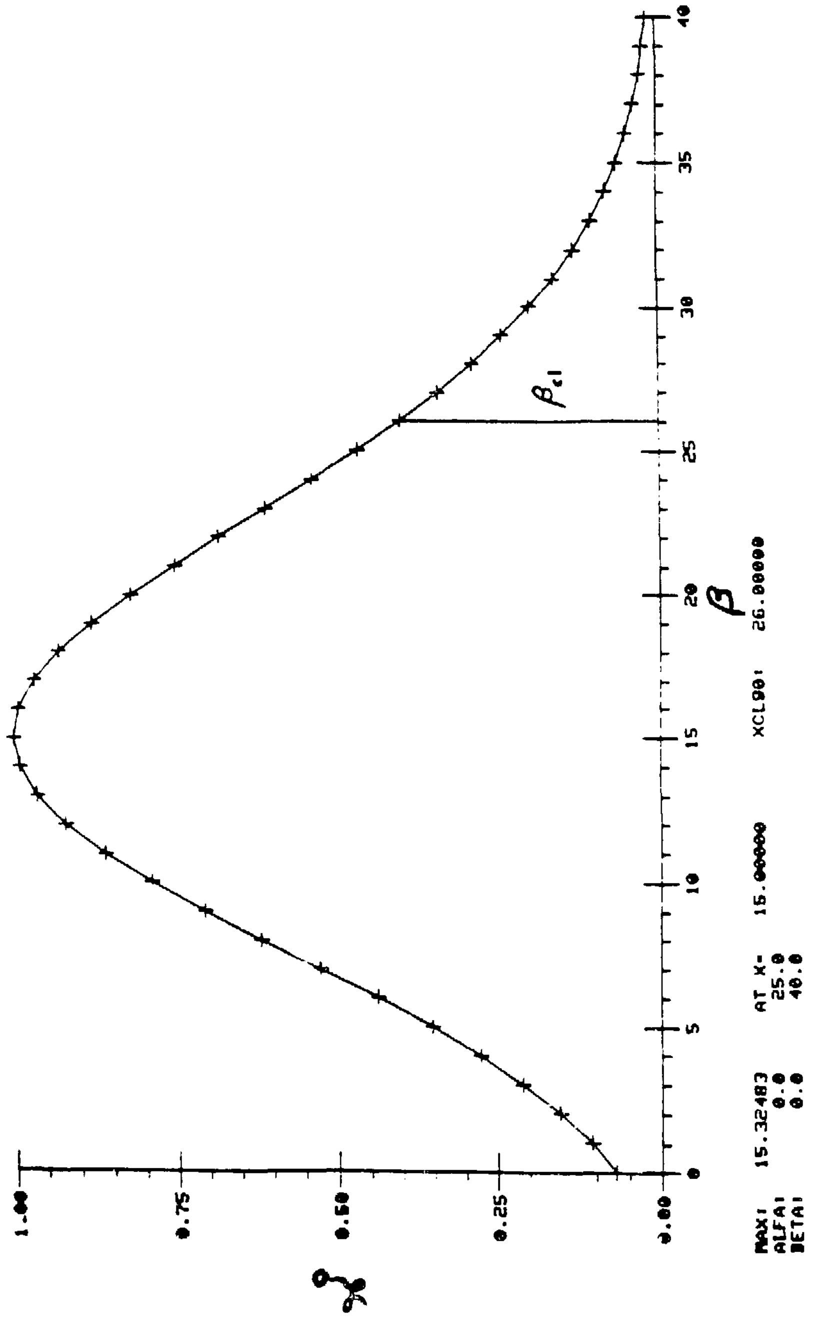

Figure 48: Plot of L_0, summed over $\alpha = 0$ to 25, as a function of β.

events was generated according to the standard QED matrix-element. When these events were analyzed with the same cuts as the data, and the number of events surviving the cuts was normalized to 2.2 x 10^{11} muon decays, the expected number in the final data sample was 10 ± 1. The most probable number observed, according to the maximum likelihood analysis above, is 16 ± 8, in reasonable agreement with the expected number.

Chapter VI

RESULTS AND CONCLUSIONS

6.1 CALCULATION OF BRANCHING RATIO UPPER LIMIT

The (90% CL) upper limit for the branching ratio BR($\mu^+ \to e^+\gamma$) is calculated as follows:

$$BR(\mu^+ \to e^+\gamma) \leq \frac{\alpha_{cl}}{(\#\ \mu\ \text{stops}) \cdot \epsilon} = \frac{9.5}{(2.2\times10^{11})(0.154)} . \qquad (45)$$

The final result for this data sample is then

$$BR(\mu^+ \to e^+\gamma) \lesssim 2.8 \times 10^{-10} \quad (90\%\ \text{CL}) . \qquad (46)$$

This result, from the first data-taking run of the Crystal Box detector, does not improve upon the existing upper limit for $\mu^+ \to e^+\gamma$ quoted in Chapter 2. However, this limit is a confirmation of the previous result by a new experiment with a different apparatus.

6.2 RECENT IMPROVEMENTS

The data-taking from the recently completed Crystal Box run, conducted during the summer of 1984, observed five times more muons with improved detector performance. The major changes to the apparatus are listed below.

1. The on-line pile-up detection system was operational. It has already been shown to reduce the number of high-energy photons contributing to the background spectrum by a factor of 3.

2. A new system for monitoring and controlling the temperature was installed in the electronics racks to reduce pedestal instabilities in the NaI(Tℓ) PHIM's and thereby improve the energy resolution.

3. Stabilizing circuitry was added to the inputs of the quadrant energy-sum modules in the trigger, to remove bothersome (60 Hz) AC pick-up. This addition was to assure consistent thresholds and subsequently allow less restrictive off-line energy cuts, implying a larger experimental acceptance for this data.

4. The temperature of the detector environment was better controlled to reduce timing fluctuations due to drifts in the PMT's.

5. Possible refinements to the NaI(Tℓ) position algorithm that may improve the resolution of θ_{ep} are being examined.

Given these improvements, the analysis of the new data sample is expected to reduce the limit achieved in this thesis by as much as a factor of five.

6.3 FUTURE EXPECTATIONS

The Crystal Box detector was originally conceived as a means to search for the decay $\mu^+\to e^+\gamma\gamma$. An essentially simultaneous proposal to look for $\mu^+\to e^+e^+e^-$ was also approved by the LAMPF Program Committee, with the suggestion that they be run concurrently. It was only later realized that an improved $\mu^+\to e^+\gamma$ experiment could also be simultaneously done. For these reasons, the roughly cylindrical symmetry of the detector is best suited to look for 3-particle decays, and therefore is not the optimum configuration for $\mu^+\to e^+\gamma$. In addition, the logistics of searching for three rare reactions at once requires compromises in the running conditions which are not most appropriate for any of the three. The design of a dedicated $\mu^+\to e^+\gamma$ experiment is now being pursued by some of the Crystal Box collaborators. The goal of this design effort is an experiment sensitive to a branching ratio for $\mu^+\to e^+\gamma$ at the level of a few times 10^{-14}.

6.4 SUMMARY

The motivations to make a sensitive search for neutrinoless muon decays have been discussed. The design and fabrication of the Crystal Box detector to perform a simultaneous search for three such modes has been described. The data-taking procedures, and the subsequent detector calibration and analysis of the $\mu^+\to e^+\gamma$ events have been presented. The result from this first Crystal Box data run of

$$BR(\mu^+\to e^+\gamma) \leq 2.8 \times 10^{-10} \quad (90\%\ CL) \tag{47}$$

confirms the existing limit of

$$BR(\mu^+\to e^+\gamma) \leq 1.7 \times 10^{-10} \quad (90\%\ CL)\ . \tag{48}$$

An improved result, from recent experimental runs with five-fold greater statistics, is expected in the near future.

REFERENCES

1. J. D. Bowman <u>et al</u>., Phys. Rev. Lett. <u>42</u>, 556 (1979), and W. W. Kinnison <u>et al</u>., Phys. Rev. D <u>25</u>, 2846 (1982).

2. LAMPF Experiments 400/445, C. M. Hoffman, J. D. Bowman, and H. S. Matis, spokesmen. First results for the $\mu^+ \rightarrow e^+e^+e^-$ search are given in R. D. Bolton <u>et al</u>., Phys. Rev. Lett. <u>53</u>, 1415 (1984).

3. C. D. Anderson and S. H. Neddermeyer, Phys. Rev. <u>51</u>, 884 (1937), and <u>54</u>, 88 (1938).

4. C. Street and E. Stevenson, Phys. Rev. <u>51</u>, 1005 (1937).

5. M. Conversi, E. Pancini, and O. Piccioni, Phys. Rev. <u>71</u>, 209 (1947).

6. M. Schein and P. S. Gill, Rev. Mod. Phys. <u>11</u>, 267 (1939).
   R. F. Christy and S. Kusaka, Phys. Rev. <u>59</u>, 414 (1941).
   J. Matthews, Phys. Rev. <u>102</u>, 270 (1956).
   A. N. Mitra, Nucl. Phys. <u>2</u>, 283 (1956/57).
   S. Hirokawa, H. Komori, and S. Ogawa, Nuovo Cimento <u>4</u>, 736 (1956).
   S. Hirokawa and H. Komori, Nuovo Cimento <u>7</u>, 114 (1958).

7. M. L. Perl <u>et al</u>., Phys. Rev. Lett. <u>35</u>, 1489 (1975).

8. E. P. Hincks and B. Pontecorvo, Phys. Rev. <u>77</u>, 102 (1950).

9. L. W. Nordheim, Phys. Rev. <u>59</u>, 554 (1941).

10. E. P. Hincks and B. Pontecorvo, Phys. Rev. <u>73</u>, 257 (1948).
    R. D. Sard and E. J. Althaus, Phys. Rev. <u>74</u>, 1364 (1948).

11. J. Steinberger, Phys. Rev. <u>75</u>, 1136 (1949).
    R. B. Leighton, C. D. Anderson, and A. J. Seriff, Phys. Rev. <u>75</u>, 1432 (1949).

12. R. L. Garwin, L. M. Lederman, and W. Weinrich, Phys. Rev. <u>105</u>, 1415 (1957).

13. J. Friedman and V. Telegdi, Phys. Rev. <u>105</u>, 1681 (1957).

14. P. C. Macq, K. M. Crowe, and R. P. Haddock. Phys. Rev. <u>112</u>, 2061 (1958).
    G. Culligan, S. G. F. Frank, and J. R. Holt, Proc. Phys. Soc. <u>73</u>, 169 (1959).

15. L. Dick, L. Feuvrais, and M. Spighel, Phys. Lett. <u>7</u>, 150 (1963).

16. A. Buhler *et al.*, Phys. Lett. 7, 368 (1963).

17. S. Bloom, L. Dick, L. Feuvrais, G. R. Henry, P. C. Macq, and M. Spighel, Phys. Lett. 8, 87 (1964).
J. Duclos, J. Heintze, A. DeRujula, and V. Soergel, Phys. Lett. 9, 62 (1964).

18. D. M. Schwartz, Phys. Rev. 162, 1306 (1967).

19. C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, Phys. Rev. 105, 1413 (1957).

20. J. Schwinger, Ann. Phys. 2, 407 (1957).

21. G. Danby, J.-M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger, Phys. Rev. Lett. 9, 36 (1962).

22. K. Nishijima, Phys. Rev. 108, 907 (1957).

23. S. Bludman, Nuovo Cimento 9, 433 (1958).

24. E. J. Konopinski and H. M. Mahmoud, Phys. Rev. 92, 1045 (1953).

25. G. Feinberg and S. Weinberg, Phys. Rev. Lett. 6, 381 (1961).

26. S. E. Willis *et al.*, Phys. Rev. Lett. 44, 522 (1980).

27. F. Bergsma *et al.*, Phys. Lett. 122B, 465 (1983).

28. C. N. Yang and R. L. Mills, Phys. Rev. 96, 191 (1954).

29. G. 't Hooft, Nucl. Phys. B33, 173 (1971), and B35, 167 (1971).
G. 't Hooft and M. Veltman, Nucl. Phys. B50, 318 (1972).
B. W. Lee, Phys. Rev. D 5, 23 (1972).
B. W. Lee, J. Zinn-Justin, Phys. Rev. D 5, 3121, 3137, 3155 (1972), and 7, 1049 (1973).

30. S. L. Glashow, Nucl. Phys. 22, 579 (1961).

31. A. Salam and J. C. Ward, Phys. Lett. 13, 168 (1964).

32. S. Weinberg, Phys. Rev. Lett. 19, 1264 (1967).

33. A. Salam, in *Elementary Particle Theory*, ed. N. Svartholm, Stockholm: Almquist and Wiksell, 367 (1968).

34. P. W. Higgs, Phys. Lett. 12, 132 (1964), and Phys. Rev. Lett. 13, 508 (1964), and Phys. Rev. 145, 1156 (1966).
T. W. B. Kibble, Phys. Rev. 155, 1554 (1967).

35. S. L. Glashow, J. Iliopoulos, and L. Maiani, Phys. Rev. D 2, 1285 (1970).

36. The New York Times, E7 (July 1, 1984).
J. Maddox, Nature 310, 97 (July 12, 1984).

37. For a review of the status of weak neutral-current experiments, see S. M. Bilensky and J. Hosiak, Phys. Rep. 90(2), 73 (1982).

38. J. J. Aubert et al., Phys. Rev. Lett. 33, 1404 (1974).
J. E. Augustin et al., Phys. Rev. Lett. 33, 1406 (1974).

39. G. Arnison et al., Phys. Lett. 122B, 103 (1983), and 126B, 398 (1983).
P. Bagnaia et al., Phys. Lett. 129, 130 (1983).

40. M. A. B. Beg A. Sirlin, Ann. Rev. Nucl. Sci. 24, 379 (1974).

41. B. W. Lee and R. E. Shrock, Phys. Rev. D 16, 1444 (1977).
J. D. Bjorken and S. Weinberg, Phys. Rev. Lett. 38, 622 (1977).

42. C. M. Hoffman, in Proceedings of the 4th Course of the International School of Physics of Exotic Atoms on Fundamental Interactions in Low Energy Systems, Erice, Italy, (1984). (Los Alamos National Laboratory document LA-UR-84-1327).

43. S. Lokanathan and J. Steinberger, Phys. Rev. 98, 240(a) (1955).
H. F. Davis, A. Roberts, and T. F. Zipf, Phys. Rev. Lett. 2, 211 (1959).
D. Berley, J. Lee, and M. Bardon, Phys. Rev. Lett. 2, 357 (1959).
T. O'Keefe, M. Rigby, and J. Wormaid, Proc. Phys. Soc. 73, 951 (1959).
V. Krestnikov, IX Annual International Conference on High Energy Physics, Kiev (1959), unpublished; J. Askin et al., Nuovo Cimento 14, 1266 (1959).
S. Frankel, V. Hagopian, J. Halpern, and A. L. Whetstone, Phys. Rev. 118, 589 (1960).
R. R. Crittenden, W. D. Walker, and J. Ballam, Phys. Rev. 121, 1823 (1961).
S. M. Korenchenko et al., Yad. Fiz. 13, 341 (1971).

44. S. Frankel et al., Nuovo Cimento 27, 894 (1963).

45. S. Parker, H. L. Anderson and C. Rey, Phys. Rev. 133, B768 (1964).

46. H. P. Povel et al., Phys. Lett. 72B, 183 (1971), and A. Schaaf et al., Nucl. Phys. A 340, 249 (1980).

47. P. Depommier, Phys. Rev. Lett. 39, 1113 (1977).

48. H.-W. Reist et al., Nucl. Instr. Meth. 153, 61 (1978).

49. R. D. Bolton et al., Nucl. Instr. Meth. 188, 275 (1981).

50. G. H. Sanders, G. W. Hart, G. E. Hogan, J. S. Frank, C. M. Hoffman, H. S. Matis, and V. D. Sandberg, Nucl. Instr. Meth. 180, 603 (1981).

51. V. D. Sandberg, L. Bayliss, M. Dugan, J. S. Frank, T. Gordon, G. Hart, C. M. Hoffman, G. E. Hogan, H. S. Matis, and G. H. Sanders, "A Fast Analog Meantimer," to be published in Nuclear Instruments and Methods.

52. G. I. Kirkbride _et al_., IEEE Trans. Nucl. Sci. _NS-26_, 1535 (1979).

53. Polycrystalline NaI(Tℓ) is a proprietary development of Harshaw Chemical Co.

54. R. I. Carrington _et al_., Nucl. Instr. Meth. _163_, 203 (1979).

55. J. Rolfe _et al_., IEEE Trans. Nucl. Sci. _NS-30_, 202 (1983).

56. Allen Avionics, Inc., Catalog 24D, 224 E. 2nd St., Mineola, New York (1982).

57. J. P. Sandoval, L. Bayliss, T. Gordon, G. Hart, C. M. Hoffman, G. E. Hogan, V. D. Sandberg, G. H. Sanders, and S. L. Wilson, Nucl. Instr. Meth. _216_, 171 (1983).

58. H. S. Butler, M. D. Cooper, R. A. Williams, E. B. Hughes, J. R. Rolfe, S. L. Wilson, and H. D. Zeman, Topical Conference on Computerized Data Acquisition in Particle and Nuclear Physics, Oak Ridge National Laboratory, Oak Ridge, Tennessee (May 1981).

59. H. L. Anderson _et al_., Los Alamos National Laboratory report LA-4504-MS (August 1970).
R. F. Thomas, Jr. _et al_., Los Alamos National Laboratory report LA-6543-MS (October 1976).
D. G. Perry _et al_., Los Alamos National Laboratory report LA-7034-MS (November 1977).
D. G. Perry, IEEE Trans. Nucl. Sci. _NS-26_, 4422 (1979).

60. M. Kellog _et al_., Los Alamos National Laboratory report LA-7001-M (June 1978).
D. G. Perry, IEEE Trans. Nucl. Sci. _NS-26_, 4494 (1979).
J. F. Harrison _et al_., IEEE Trans. Nucl. Sci. _NS-28_, (1981).

61. L. R. Biswell and R. E. Rajala, Los Alamos National Laboratory report LA-5144 (May 1973).
R. F. Thomas, Jr., Los Alamos National Laboratory report LA-5404-MS (September 1973).

62. R. L. Ford and R. W. Nelson, SLAC Report 210 (June 1978).

63. R. D. Bolton, M. D. Cooper, J. S. Frank, V. E. Hart, H. S. Matis, R. E. Mischke, V. D. Sandberg, J. P. Sandoval, and U. Sennhauser, "Performance of the Cylindrical Drift Chamber for the Crystal Box at LAMPF," to be published in Nuclear Instruments and Methods.

ACKNOWLEDGMENTS

I wish to express my sincere thanks to all of the collaborators on the Crystal Box experiment for their kind help in making this thesis possible. I must especially thank C. M. Hoffman for his advice and attention whenever it was required. I am indebted to A. Hallin for his help in creating many of the illustrations in this work, and for always sharing a smile. R. Bolton was also gracious with his time in preparing some of the figures. Thanks go to R. Williams for his patience in completing the on-line data-acquisition system software, and to V. Sandberg for his valuable assistance in debugging the readout electronics. The efforts of S. Chesney, M. Dugan, G. Hart, and J. Sandoval to keep the electronics in working order were greatly appreciated. Many enlightening discussions were shared with M. Cooper, and G. Sanders was always an inspiration to get things done. Thanks to J. Rolfe and L. "Dusty" Doster for overseeing the fabrication of hundreds of channels of electronics, and to R. Parks and V. Hart for attending to the construction of "the Box." I appreciate the help of D. Bowman and H. Matis in understanding the behavior of NaI(Tℓ) crystals. To my fellow graduate students, D. Grosnick and J. McDonnough, go many thanks for wading through all the calibration procedures with me. My appreciation is extended to R. Hofstader for his patience with me as the thesis neared completion, and to D. Nagle for his support where it was sometimes needed most. I am thankful that R. Mischke was willing to be the software "Czar."

Of course, this experiment would never have been "fired-up" without the trigger, for which we are all indebted to G. Hogan (and his Magic Light Show). Special thanks to P. Heusi for pushing through the data analysis. Finally, I would like to express deepest appreciation to E. B. Hughes for his efforts in the tedious task of proof-reading and polishing this text, thereby teaching me a bit about how to write.

U S GOVERNMENT PRINTING OFFICE 1985— 57[illegible] 034 / 20115