

1998

A search for lepton flavor violation in the neutral kaon system

Christopher Henry Hoff

College of William & Mary - Arts & Sciences

Follow this and additional works at: <https://scholarworks.wm.edu/etd>



Part of the [Physics Commons](#)

Recommended Citation

Hoff, Christopher Henry, "A search for lepton flavor violation in the neutral kaon system" (1998).
Dissertations, Theses, and Masters Projects. William & Mary. Paper 1539623930.
<https://dx.doi.org/doi:10.21220/s2-bd5q-yt15>

This Dissertation is brought to you for free and open access by the Theses, Dissertations, & Master Projects at W&M ScholarWorks. It has been accepted for inclusion in Dissertations, Theses, and Masters Projects by an authorized administrator of W&M ScholarWorks. For more information, please contact scholarworks@wm.edu.

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

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

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

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

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

UMI

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

NOTE TO USERS

The original manuscript received by UMI contains pages with indistinct print. Pages were microfilmed as received.

This reproduction is the best copy available

UMI

A Search for Lepton Flavor Violation in the Neutral Kaon System

A Dissertation

Presented to The Faculty of the Department of Physics

The College of William and Mary

In Partial Fulfillment

Of the Requirements for the Degree of

Doctor of Philosophy

By

Christopher H. Hoff

June 1998

UMI Number: 9920301

**UMI Microform 9920301
Copyright 1999, by UMI Company. All rights reserved.**

**This microform edition is protected against unauthorized
copying under Title 17, United States Code.**

UMI
300 North Zeeb Road
Ann Arbor, MI 48103

APPROVAL SHEET

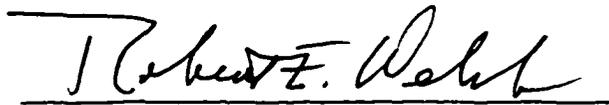
This dissertation is submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy.

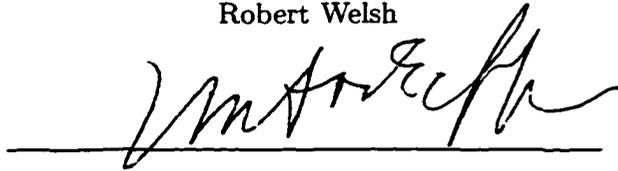


Christopher H. Hoff

Approved, June 1998



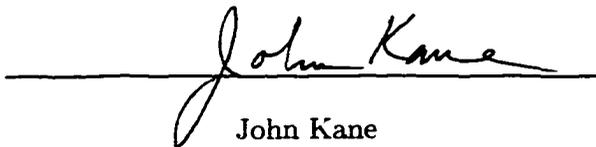
Robert Welsh



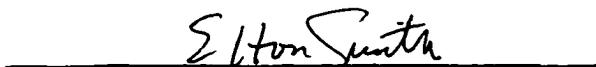
Morton Eckhause



Franz Gross



John Kane



Elton Smith

This effort is dedicated to
my wife, Kay,
my daughter, Christiana,
and my parents, Walter and Connie Hoff.

Contents

List of Tables	vii
List of Figures	viii
Acknowledgments	xi
Abstract	xiv
Chapter 1 Introduction	1
Chapter 2 Theory	5
2.1 Standard Model	5
2.2 Minimal Extensions to the Standard Model	7
2.2.1 Massive Neutrinos	7
2.2.2 Fourth Generation of Leptons	7
2.2.3 Left-Right Symmetric Models	8
2.2.4 Multiple Higgs Doublets	9
2.3 Beyond the Standard Model	9
2.3.1 General Considerations	9
2.3.2 Horizontal Gauge Models	13
2.3.3 Pati-Salam Model	13
2.3.4 Technicolor Models and Extensions	14
2.3.5 Composite Models	16
2.3.6 Supersymmetry	17
2.4 Theory Summary	18
Chapter 3 Experimental Setup	19
3.1 Overview	19
3.2 Beamline	20
3.2.1 Target	20

3.2.2	Beamline Magnets and Collimators	23
3.2.3	Vacuum Decay Region	23
3.3	Spectrometer	24
3.3.1	Straw Chambers	26
3.3.2	Drift Chambers	29
3.3.3	Spectrometer Magnets	29
3.3.4	Beam Plug	31
3.4	Trigger Scintillators	32
3.5	Electron ID	36
3.5.1	Čerenkov Counter	36
3.5.2	Lead Glass Calorimeter	46
3.6	Muon ID	48
3.6.1	Muon Hodoscope	49
3.6.2	Muon Rangefinder	50
3.7	Data Acquisition System	55
3.7.1	Level 0/1 Trigger	55
3.7.2	Data Acquisition Electronics	58
3.7.3	Level 3 Trigger	63
3.7.4	Data Acquisition Computer	64
Chapter 4 Offline Software		67
4.1	Overview	67
4.2	Multiple Scattering Calculation	67
4.3	Pattern Recognition	68
4.4	Track Fitting	71
4.4.1	FT Fitter	71
4.4.2	QT Fitter	73
4.5	Particle Identification	75
4.5.1	General Track Counter Association	76
4.5.2	Electron Identification	76
4.5.3	Muon Identification	78
Chapter 5 Production Analysis		80
5.1	Organization	80
5.1.1	Pass1	81
5.1.2	Pass2	84
5.1.3	Pass3	87

Chapter 6	$K_L^0 \rightarrow \mu e$ Analysis	89
6.1	Backgrounds	91
6.1.1	Pileup	91
6.1.2	Tracking Errors	92
6.1.3	Single and Double Misidentification	92
6.1.4	Pion Decay with Electron Mott scattering	94
6.2	Cuts	95
6.2.1	Event Cuts	97
6.2.2	Vertex Cuts	97
6.2.3	Track Cuts	100
6.2.4	Particle Identification Cuts	100
6.2.5	Cut Results	100
Chapter 7	Analysis Results	103
7.1	Signal Region	103
7.2	Single Event Sensitivity	106
7.2.1	$\mu\mu$ Normalization Sample	106
7.2.2	Experimental Acceptances for $\mu\mu$ and μe	107
7.2.3	Relative Efficiencies for $\mu\mu$ and μe	108
7.3	Final Result	110
Chapter 8	Conclusions	111
8.1	Theoretical Implications and Future Prospects	111
8.2	Individual Contributions	112
Appendix A	Čerenkov Counter Component Tests	114
A.1	Mirror Blank Testing	114
A.2	PMT Rate Tests	116
A.3	Magnetic Shield Tests	117
Bibliography		120
Vita		125

List of Tables

3.1	K_L Decay Transverse Momentum	31
3.2	MHO Panel Dimensions	49
4.1	E871 Spectrometer Radiation Lengths	68
4.2	CER e ID categories	77
4.3	PbG e ID categories	77
4.4	IMF μ ID categories	79
5.1	Pass1 Bits	84
5.2	Pass1 Statistics	85
5.3	Pass2 Bits	86
5.4	Pass2 Statistics	87
6.1	μe Analysis Cuts	96
7.1	Efficiency factors that cancel or largely cancel	109
7.2	Efficiency factors that only affect $\epsilon_{\mu e}$	109
7.3	Efficiency factors that only affect $\epsilon_{\mu\mu}$	109
A.1	CER Mirror Characteristics	115

List of Figures

1.1	E871 Collaboration List	4
2.1	$K_L^0 \rightarrow \mu e$ via Right Handed Bosons	8
2.2	$K^0 \Leftrightarrow \bar{K}^0$ Mixing	10
2.3	$K_L^0 \rightarrow \mu e$ via X boson	10
2.4	$K^+ \rightarrow \mu^+ \nu_\mu$ Decay	11
2.5	$K_L^0 \rightarrow \mu e$ via Y boson	12
2.6	$K_L^0 \rightarrow \mu e$ via Horizontal Gauge Boson	13
2.7	$K_L^0 \rightarrow \mu e$ via Pati-Salam Model	14
2.8	$K_L^0 \rightarrow \mu e$ via ETC boson or technipion	14
2.9	$K_L^0 \rightarrow \mu e$ via pseudoscalar leptoquark or vector leptoquark	14
2.10	Composite model diagram for $K_L^0 \rightarrow \mu e$	16
2.11	Neutralino induced contribution to $K_L^0 \rightarrow \mu e$	17
3.1	E871 Experiment Schematic Drawing	21
3.2	E871 Spectrometer Overview	24
3.3	Wire layout of the straw and drift chambers	25
3.4	SDC Wire Configuration	27
3.5	SDC Amplifier Schematic	28
3.6	SDC Pin Layout	28
3.7	Spectrometer Wire Efficiencies	30
3.8	Spectrometer Resolution	30
3.9	XZ plane cross-sectional view of the beam plug	32
3.10	TSC Slat Mounting and Support	33
3.11	TSC Configuration	34
3.12	Trigger Scintillator Efficiencies	35
3.13	Čerenkov Enclosure	38
3.14	Čerenkov Mirror Configuration	39
3.15	Burle 8854 Response	40

3.16 Čerenkov PMT Base Layout	42
3.17 CER Photomultiplier Assembly	42
3.18 CER ADC response in photoelectrons	44
3.19 CER Efficiency vs. Track Momentum	45
3.20 Lead Glass Array Layout	46
3.21 PbG Performance	47
3.22 IMF Detector Layout	48
3.23 Muon Hodoscope Panel Dimensions and Layout	51
3.24 MHO Efficiencies	52
3.25 MRG Extrusion (End View)	53
3.26 MRG Discriminator Card Schematic	53
3.27 MRG X-measuring plane (Beam view)	54
3.28 Level 1 Trigger	56
3.29 Level 1 Flowchart	57
3.30 Data Acquisition System	59
3.31 Data Pipeline Outline	61
4.1 MRG observed minus expected track length	79
5.1 Pass1 flowchart	82
5.2 FT/QT mass difference plot	88
5.3 FT/QT p_T^2 difference plot	88
6.1 Blind analysis schematic	90
6.2 FT based p_T^2 vs. $M_{\mu e}$ plot from uncut Pass3 output	90
6.3 Effects of particle misidentification in K_{e3} decays on $M_{\mu e}$	93
6.4 Effect of Mott scattering in K_{e3} decay on $M_{\mu e}$	94
6.5 Vertex Z position	98
6.6 Vertex θ_x plot	99
6.7 Vertex θ_y plot	99
6.8 Track front/back momentum cut	101
6.9 Momentum asymmetry cut plot	101
6.10 FT based p_T^2 vs. $M_{\mu e}$ plot, post cuts	102
7.1 Effects of inner bremsstrahlung on μe signal	104
7.2 FT based p_T^2 vs. $M_{\mu e}$ plot, in exclusion region	105
7.3 Plot of p_T^2 background subtraction for $\mu\mu$ normalization.	107
7.4 Plot of $M_{\mu\mu}$ with all relevant μe cuts applied.	108
A.1 Modified E791 CER PMT Base test results	118

A.2 Final CER PMT base design test results	118
--	-----

Acknowledgments

It has been said that a life is a book written by one author but containing an extensive bibliography. While I have written this dissertation myself, its completion would not be possible without the assistance, encouragement and support of many others over the years. I would like to thank many of these people at this time.

Thanks to the members of the E871 Collaboration for making this experiment possible with all of their hard work and dedication. Kudos go to Stan Wojcicki, Bill Molzon, Jack Ritchie and John Kane, the experimental spokesmen, for their overall direction over the past 6 years. Special thanks to Jim McDonough and Mike Hebert for their leadership down in the trenches both in their role as experimental liaisons between E871 and BNL during the construction and running of the experiment and for instructing me, and others, in the finer points of running an experiment and analyzing the data. The camaraderie of fellow on-site collaborators Dave Ambrose, Mark Bachman, Mike Hebert, Jim McDonough, Dave Conner, Roy Lee, Nobu Kanematsu, Andy Milder, Scott Graessle, Vassilis Vassilakopoulos, Brent Ware, Steve Worm, Marize Pommot-Maia, Keith "Casey" Hartman, Karl Ecklund and Rob Martin made life on-site more enjoyable, even when the day's work was long and grueling. Rob Martin deserves special recognition for enduring sharing an apartment with me at BNL over several years and for remaining a good friend while doing so! Mike Hebert also gets a special mention for his patience with my data analysis questions and his numerous efforts

to make me a better physicist. Thanks go also to the other members of the μe analysis team, Mark Bachman, Roy Lee, and Scott Graessle, for their effort and determination to get the best measurement possible. I would like to thank the staff of Brookhaven National Laboratory for their work in this endeavor. In particular, thanks to Ralph Brown, Fred Kobasiuk and the rest of the AGS Mechanical Group, for helping us build and maintain our experiment, Joe Scaduto, for coordinating the work we required from the various BNL shops, the MCR personnel, for getting as much beam to us as we could stand, Peter Takacs and the Metrology Lab staff for the use of their facility during CER mirror testing, and, finally, to Arnie Esper and the people in the Hydrogen Target Group, for designing and building the Čerenkov gas system and letting me scavenge parts from them when needed.

At William and Mary, I would like to thank my advisor, Robert E. Welsh, and his fellow principal investigators, Morton Eckhause and John Kane, for their guidance and advice throughout my graduate career and for their insightful suggestions and comments during the development of this dissertation. I would also like to thank A. Dayle Hancock and Yunan Kuang for their assistance and direction in this endeavor over the years. Fellow graduate students Marco Brown, Mike Seale, Mike West, Dave Meekins, Kelly (Doty) Herbst, Mike Uzzo, Dave Morgan, Alan Coleman, Chris Lasota and Paul King helped make my experience at William and Mary an interesting and enjoyable one. Thanks also to Paula Perry, Sylvia Stout and Dianne Fannin for all of their help with a multitude of administrative tasks.

Closer to home, I would like to thank all my family and friends for their support over the past eight years. They have suffered through a plethora of explanations of my work, rants about why things were taking so long and promises that it would all be over "soon". To all of you who have endured my graduate school tenure, I appreciate your patience! Particular thanks go out to my parents, Walter and Connie Hoff, for encouraging me to

do what I wanted to do, even if they didn't understand it, and for helping me up when I needed it.

Finally, I would like to thank my wife, Kay, for her patience and enduring love, even as I was away for three weeks out of a month for several years. This work would not have been possible without your understanding and your willingness to let me go away sometimes to follow this path to its end. Thank you!

This work was supported in part by the National Science Foundation.

Abstract

The conservation of separate lepton number is an *ad hoc* conservation principle in the Standard Model but several non-Standard Model theories allow violation of separate lepton number conservation. A very sensitive search for lepton flavor violation in the neutral kaon system, via the decay $K_L^0 \rightarrow \mu e$, was undertaken by Brookhaven National Laboratory Experiment 871 (E871) during the 1995 and 1996 AGS HEP running periods. No $K_L^0 \rightarrow \mu e$ decays were observed and a single event sensitivity (SES) of 1.97×10^{-12} was achieved. This single event sensitivity corresponds to an upper limit on the branching fraction of $B(K_L^0 \rightarrow \mu e) < 4.53 \pm 0.21 \pm 0.46 \times 10^{-12}$ at a 90% confidence level.

Chapter 1

Introduction

Particle physics usually follows the “higher energy” path to discovering new physics. This means physicists must build larger and larger accelerators to attain the higher center of mass energies required to directly observe new phenomena. This approach is a good one and has resulted in many new discoveries, *e.g.*, the W , Z and the top quark. However, barring unforeseen breakthroughs in accelerator technology, this approach is ultimately cost prohibitive. Another approach to new physics which has worked well is the “higher precision” path, which allows physicists to use lower center of mass machines to look for small deviations from known physics, *e.g.*, $\sin^2 \theta_W$, CP violation and rare decays. This approach is less expensive, although it requires a large investment in time to complete. One particle system that lends itself well to high precision experiments is the kaon system.

The kaon system is a physics rich system that still deserves study, even 50 years after its discovery. The discovery of the kaon revealed the strange quark, the first quark found that does not exist in “normal” matter. This system is still the only system with observed direct CP violation in the decay $K_L^0 \rightarrow \pi\pi$ [1]. The lower than expected rate of $K_L^0 \rightarrow \mu\mu$ prompted the prediction of the charm quark through the GIM mechanism [2].

This led to the CKM matrix, which describes the weak interaction of quarks in the Standard Model. The advantage of using kaons in rare decay searches is two-fold. First, kaons have a low mass, thus they can be produced in great quantities by lower energy accelerators. Second, the K_L and K^\pm have long lifetimes which allow greater separation between these kaon decays and other prompt processes. This makes a study of the decays of these particles easier to perform.

This dissertation focuses on the search for the lepton flavor violating decay mode of the neutral kaon, $K_L^0 \rightarrow \mu e$, in BNL Experiment 871. E871 was proposed in 1990 and approved in 1991 for over 5,600 hours of running with a high intensity 24 GeV/ c proton beam. The E871 collaboration, listed in Figure 1.1. consisted of 33 members from five institutions participating in the construction and operation of the experimental apparatus as well as in the subsequent data analysis. The construction of E871 began in 1992 and was completed in time for the 1995 running period. The data were acquired during both the 1995 and 1996 AGS running period, or approximately eight months. The data accumulated amounted to 1238 1.4 Gigabyte DAT tapes, for a sum of over 1.7 Terabytes of data. Two other decay modes ($K_L^0 \rightarrow \mu\mu$ and $K_L^0 \rightarrow ee$) were examined as well. The first decay, $K_L^0 \rightarrow \mu\mu$, has already been studied but E871 expected to collect the largest sample of $K_L^0 \rightarrow \mu\mu$ events to constrain the statistical errors on the $K_L^0 \rightarrow \mu\mu$ measurement. $K_L^0 \rightarrow ee$ had not been seen previously but E871 observed four unambiguous $K_L^0 \rightarrow ee$ events, making this decay the rarest particle decay observed to date. The CP violating decay mode, $K_L^0 \rightarrow \pi\pi$, which occurs once in every 500 K_L decays, was collected at a large prescale and used in the normalization of the results from the three decays of interest.

The E871 experimental design limits are a single event sensitivity of approximately 1×10^{-12} , or a factor of 15 improvement over the previous best measurement from its predecessor, BNL E791. The main improvement was expected from a relatively new and

untried concept, the beam plug. E791 was limited for two reasons: the available AGS intensity, and the high neutron and undecayed kaon flux, inevitable in a neutral beam experiment. The intense neutron flux forced their design to use a central corridor between two detectors to avoid overloading the detectors. Unfortunately, the success of this approach was limited because the neutral beam could still interact with the inner edges of the detectors and contribute noise to the desired events. Thus, the beam intensity was kept low, which meant the kaon flux was low as well. The AGS intensity limitation was overcome by the addition of the AGS Booster and the upgrading of the primary beamline magnets, allowing better control of the higher intensity beam. The neutral beam intensity limit in the apparatus was overcome in E871 with the beam plug, which was located in the center of the first spectrometer magnet. This compact beam stop absorbs the neutral beam and allows the downstream detectors to span the beamline without any event degradation due to intense background. The removal of the neutral beam, in turn, allows a higher primary beam intensity without a corresponding increase in detector noise.

This work will discuss the theories that allow $K_L^0 \rightarrow \mu e$ decays, the experimental methods used to search for this decay and the steps taken in the analysis which resulted in a single event sensitivity (SES) of approximately two parts in 10^{12} .

College of William and Mary

M. Eckhause, A.D. Hancock, C.H. Hoff, J.R. Kane,
Y. Kuang, R.D. Martin, R.E. Welsh, E. Wolin

University of California, Irvine

M. Bachman, P. de Cecco, D. Conner, R. Lee,
N. Kanematsu, W.R. Molzon

University of Richmond

P.D. Rubin

Stanford University

C. Arroyo, K. Ecklund, K. Hartman, M. Hebert,
G.M. Irwin, M. Pommot-Maia, S.G. Wojcicki

University of Texas at Austin

D. Ambrose, S. Graessle, G.W. Hoffman, K. Lang,
J. McDonough, A. Milder, P.J. Riley, J.L. Ritchie,
V. Vassilakopoulos, C.B. Ware, S. Worm

Figure 1.1: The E871 Collaboration.

Chapter 2

Theory

2.1 Standard Model

The theory relevant to the decay $K_L^0 \rightarrow \mu e$ begins with the Standard Model. The current electroweak Standard Model (SM) describes the interactions of six quarks, six leptons and the force carrying bosons. Gluons will be ignored in this discussion as they do not play a part in the electroweak portion of the SM. The quarks and leptons can be arranged into three generations in the following manner,

$$\begin{array}{ccc} \begin{pmatrix} u \\ d \end{pmatrix} & \begin{pmatrix} c \\ s \end{pmatrix} & \begin{pmatrix} t \\ b \end{pmatrix} & \begin{array}{l} Q = +\frac{2}{3} \\ Q = -\frac{1}{3} \end{array} \\ \begin{pmatrix} \nu_e \\ e \end{pmatrix} & \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix} & \begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix} & \begin{array}{l} Q = 0 \\ Q = -1 \end{array} \end{array}$$

where Q is the electric charge of the particle. In the quark sector, there is mixing among the mass eigenstates, which can be described by the CKM matrix [3],

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}. \quad (2.1)$$

where (d', s', b') are the weak eigenstates and (d, s, b) are the strong eigenstates of the quarks. Similar mixing in the lepton sector [4.5] is forbidden, however, because the neutrinos are assumed to be massless. The electroweak force-carrying bosons can be arranged by ascending mass in this manner:

$$\begin{pmatrix} \gamma \\ W^\pm \\ Z^0 \end{pmatrix}.$$

The γ is the electromagnetic force boson while the W^\pm and the Z^0 are the bosons associated with the weak force.

While this model is currently compatible with all known experimental observations, it has several shortcomings that make it less appealing as a fundamental theory. The unexplained range of the particle masses, the number of quark/lepton generations, the large number of free parameters and the extreme fine tuning of the Higgs mass required in order that the theory match experimental data all contribute to the impression that there must be a more fundamental theory than the current Standard Model. Many such theories have been proposed, ranging from minimal extensions of the current theory to more exotic supersymmetric (SUSY) models but none of these others has been verified. Nonetheless, they appear to be more fundamental than the SM so they may offer solutions to some of the nagging problems of the SM.

One facet of these other theories is that in most of them, lepton flavor violation is allowed, to varying degrees. An example of lepton flavor violation would be the decay

$\mu \rightarrow e\gamma$, in which the muon would decay into an electron without associated neutrinos, in contrast to the experimentally observed decay $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$. Lepton flavor violation is not explicitly forbidden by the SM provided one or more of the neutrinos is massive. There is some recent evidence for massive neutrinos, based on the measured rate for neutrino oscillation [6]. However, in the SM, the neutrinos have zero mass. Lepton flavor conservation has been added *ad-hoc* to the model, partly because its violation has not been observed and partly because of the assumption that neutrinos are indeed massless. There are several interactions, other than $K_L^0 \rightarrow \mu e$, which can supply constraints to any theory which might allow lepton flavor violation. The decay $\mu \rightarrow e\gamma$ and $K^0 \leftrightarrow \bar{K}^0$ mixing, which gives rise to the $K_L - K_S$ mass difference (Δm_K), are two such interactions which can indirectly limit the viability of a given theory to predict the decay $K_L^0 \rightarrow \mu e$.

2.2 Minimal Extensions to the Standard Model

2.2.1 Massive Neutrinos

If one or more of the neutrinos were massive, as suggested by recent experimental results, a mixing between weak eigenstates, similar to mixing between strong eigenstates via the CKM matrix, would be expected at some level. However, the current limits on the neutrino masses imply $B(K_L^0 \rightarrow \mu e) \approx 10^{-25}$ [7], which is unobservable with present technology. Therefore, any detection of $K_L^0 \rightarrow \mu e$ in this experiment would be a signal of physics beyond the Standard Model.

2.2.2 Fourth Generation of Leptons

Another extension would be to add another generation of leptons. These leptons would both have to be massive and, given current constraints on $\mu \rightarrow e\gamma$, the extension

implies $B(K_L^0 \rightarrow \mu e) < 2 \times 10^{-15}$ [8]. Other estimates, with a “see-saw” mechanism [9, 10] governing the mixing between the standard light neutrinos and a heavy, fourth generation neutrino, have pushed the limit on $B(K_L^0 \rightarrow \mu e)$ down to $\approx 10^{-24}$ [7]. Like the massive neutrino possibility, discussed in Section 2.2.1, testing this extension is also beyond the limits of current experimental measurements. In addition, the possibility of a fourth lepton generation conflicts with data from the measurement of the Z width [11], at least in the case of light leptons.

2.2.3 Left-Right Symmetric Models

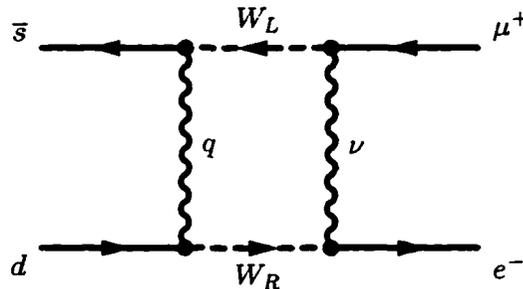


Figure 2.1: $K_L^0 \rightarrow \mu e$ via Right Handed Bosons.

A more promising extension, which would allow $K_L^0 \rightarrow \mu e$ at a measurable level, is the introduction of a right handed symmetry currently omitted from the electroweak model. The standard electroweak model is left handed, that is the W^\pm and Z only couple to left handed interactions. A left-right symmetric model would imply the existence of W_R^\pm and Z_R weak bosons. These new bosons would have to be massive enough (> 300 GeV) to have avoided discovery to this point. If right handed neutrinos existed as well, $K_L^0 \rightarrow \mu e$ could proceed via Figure 2.1 with a branching fraction in the range $10^{-13} - 10^{-15}$ [7, 12]. This value is beyond the capability of E871 but is still of interest for potential future experiments. If additional restrictions, such as horizontal symmetry [13, 14], are placed on

the right handed sector, the restrictions from Δm_K are avoided and the predicted branching ratio can be much higher. One recent calculation [15], which takes Higgs self-energy and vertex correction terms into account in a gauge invariant manner, allows $K_L^0 \rightarrow \mu e$ up to and above the current experimental limit for certain parameter spaces. Therefore, the existing limit of $B(K_L^0 \rightarrow \mu e) < 3.3 \times 10^{-11}$ already puts a severe restriction on these models.

2.2.4 Multiple Higgs Doublets

Adding multiple Higgs doublets [16] is another minimal extension to the Standard Model which can allow lepton flavor violation. Such “extra” Higgs doublets are actually required in left-right symmetric models to break the $SU(2)_R$ symmetry. When additional Higgs doublets are allowed, the new Higgs particles can then mediate $K_L^0 \rightarrow \mu e$. However, they also affect $K^0 \leftrightarrow \bar{K}^0$ mixing [17] and thus Δm_K is a strong constraint on this Standard Model extension.

2.3 Beyond the Standard Model

2.3.1 General Considerations

The expected SES of E871 is on the order of 1×10^{-12} . Considering the limits on the Standard Model extensions set forth above, with the exception of the left-right symmetric model, a positive signal for $K_L^0 \rightarrow \mu e$ would be a signal for physics beyond the Standard Model. Before describing the non-Standard Model theories, it is useful to consider the general types of interactions which may give rise to $K_L^0 \rightarrow \mu e$. Likewise, it is also informative to derive a method for comparing rates for lepton flavor violating theories with known interactions. (The following discussion is taken from [18]). The first

interaction considered is a tree level interaction (Figure 2.3) in which the boson interacts only with quark-quark or lepton-lepton vertices with a coupling of f and f' respectively. This interaction would also mediate $K^0 \Leftrightarrow \bar{K}^0$ mixing, as shown in Figure 2.2, unless another restriction on the theory interfered. $K_L^0 \rightarrow \mu e$, via Figure 2.3, can be compared to $K^+ \rightarrow \mu^+ \nu_\mu$, shown in Figure 2.4, with the assumption that $K_L^0 \rightarrow \mu e$ has the same V-A form as $K^+ \rightarrow \mu^+ \nu_\mu$.

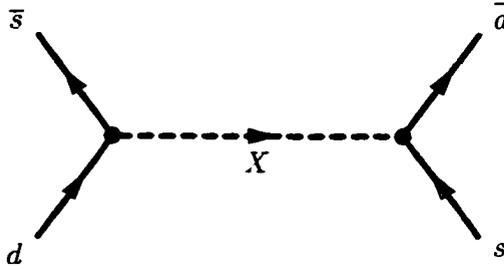


Figure 2.2: $K^0 \Leftrightarrow \bar{K}^0$ Mixing.

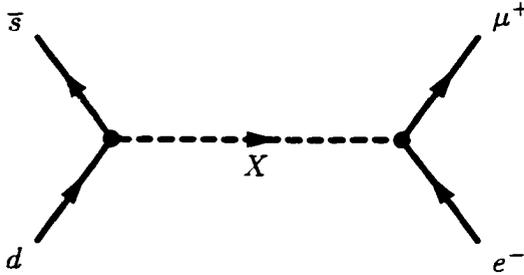


Figure 2.3: $K_L^0 \rightarrow \mu e$ via X boson.

Then, the ratio of the decay rates is:

$$\frac{\Gamma(K_L^0 \rightarrow \mu e)}{\Gamma(K^+ \rightarrow \mu^+ \nu_\mu)} \simeq \left[\frac{f f' / M_X^2}{g^2 \sin^2 \theta_c / M_W^2} \right]^2 \quad (2.2)$$

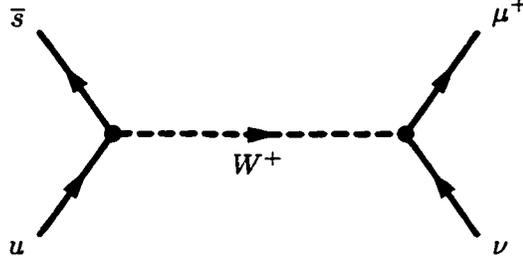


Figure 2.4: $K^+ \rightarrow \mu^+ \nu_\mu$ Decay.

where M_X and M_W are, respectively, the X and W boson mass, g is the electroweak coupling constant and θ_c is the Cabibbo angle. The total branching fraction is

$$\begin{aligned}
 B(K_L^0 \rightarrow \mu e) &= \frac{\Gamma(K_L^0 \rightarrow \mu e)}{\Gamma(K_L^0 \rightarrow \text{all})} \\
 &\simeq \frac{\Gamma(K^+ \rightarrow \mu^+ \nu_\mu)}{\Gamma(K_L^0 \rightarrow \text{all})} \left[\frac{f f' / M_X^2}{g^2 \sin \theta_c / M_W^2} \right]^2
 \end{aligned} \tag{2.3}$$

Putting the known values into Eq. 2.3 gives:

$$B(K_L^0 \rightarrow \mu e) \simeq (1.2 \times 10^{-2} \text{ TeV}^4) \left[\frac{f^2}{M_X^2} \right]^2 \left[\frac{f'}{f} \right]^2 \tag{2.4}$$

Under the assumption that $f = f' = g$, inverting Eq. 2.4 gives:

$$M_X \simeq 220 \text{ TeV} \left[\frac{10^{-12}}{B(K_L^0 \rightarrow \mu e)} \right]^{1/4} \tag{2.5}$$

The effect of Δm_K on $B(K_L^0 \rightarrow \mu e)$ is approximated [19] by

$$\Delta m_K \approx f_K^2 m_K \left[\frac{f^2}{M_X^2} \right], \tag{2.6}$$

where f_K is the kaon decay constant ($\approx 160 \text{ MeV}$) and m_K is the K_L mass. Putting in the known values gives

$$\frac{f^2}{M_X^2} \simeq 2.8 \times 10^{-7} \text{ TeV}^2 \quad (2.7)$$

as a rough estimate. When this value is then put into Eq. 2.4, the branching fraction becomes

$$B(K_L^0 \rightarrow \mu e) \simeq 1 \times 10^{-15} \left[\frac{f'}{f} \right]^2. \quad (2.8)$$

It is possible that $(f'/f)^2$ might increase this fraction several orders of magnitude. However, it is more likely that, if $K_L^0 \rightarrow \mu e$ proceeds via an 'X boson' interaction and there is no mechanism to preferentially suppress $K^0 \leftrightarrow \bar{K}^0$ mixing, $B(K_L^0 \rightarrow \mu e)$ will be below the sensitivity of E871.

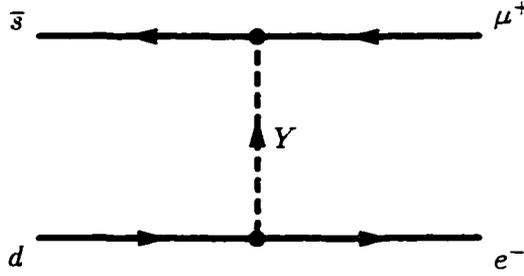


Figure 2.5: $K_L^0 \rightarrow \mu e$ via Y boson.

Another interaction class to consider would be an exchange of a heavy boson, Y , as shown in Figure 2.5, where Y couples to quark-lepton vertices but not to quark-quark or lepton-lepton vertices. These bosons are referred to as 'leptoquarks' and can be of either vector or pseudoscalar bosons. Because a leptoquark only couples to quark-lepton vertices, it is unable to mediate $K^0 \leftrightarrow \bar{K}^0$ mixing and therefore is not constrained by Δm_K . In addition, because $\mu \rightarrow e\gamma$ has no quark contribution, $\mu \rightarrow e\gamma$ is not a constraint on this interaction either. Thus, if interactions of this type exist, it is possible that they can cause

effects in the range to which E871 is sensitive.

2.3.2 Horizontal Gauge Models

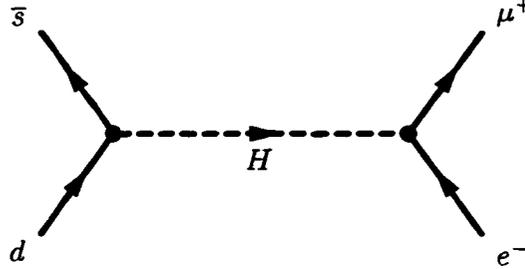


Figure 2.6: $K_L^0 \rightarrow \mu e$ via Horizontal Gauge Boson.

Horizontal Gauge models [20, 21] were developed to explain the observed structure of the quark and lepton generations. These models hypothesize horizontal gauge bosons ($SU(2)_H$ group) which mediate the interactions between the different generations. Each generation is assigned a generation number, $G = 1, 2, 3$, and, if generation number is conserved, the absolute sum $|\Delta G| = 0$. $K_L^0 \rightarrow \mu e$ could be allowed via the exchange of a horizontal gauge boson, H , with a tree level diagram, as shown in Figure 2.6. While $K_L^0 \rightarrow \mu e$ is allowed in this model, $K^0 \leftrightarrow \bar{K}^0$ is suppressed ($|\Delta G| = 2$) as well as $\mu \rightarrow e \gamma$ ($|\Delta G| = 1$), assuming that horizontal symmetry is unbroken. This means that the constraints that these interactions place on $K_L^0 \rightarrow \mu e$ in other models are non-existent in this model. Given the existing limit on $B(K_L^0 \rightarrow \mu e)$, and using the result in Eq. 2.5, the lower limit on the mass of a horizontal gauge boson, M_H , is 90 TeV.

2.3.3 Pati-Salam Model

The Pati-Salam Model [22] is based on the concept that leptons and quarks are part of an $SU(4)$ group with the fourth color being lepton number. In this model, $K_L^0 \rightarrow \mu e$ is mediated by a leptoquark, P , via the diagram in Figure 2.7. Because leptoquarks couple to

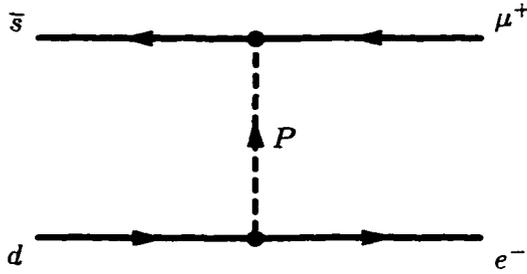


Figure 2.7: $K_L^0 \rightarrow \mu e$ via Pati-Salam Model.

lepton/quark vertices, they allow $K_L^0 \rightarrow \mu e$ but do not affect $K^0 \leftrightarrow \bar{K}^0$, and thus are not constrained by Δm_K . Leptoquarks do not affect $\mu \rightarrow e\gamma$ either and therefore they are not constrained by $B(\mu \rightarrow e\gamma)$. Based on the current limit on $B(K_L^0 \rightarrow \mu e)$, the lower limit on the mass of this leptoquark is 970 TeV [23]. A recent calculation [24] with a modified Pati-Salam model ($SU(4) \otimes SU(2)_L \otimes SU(2)_R$) gives a leptoquark mass > 140 TeV.

2.3.4 Technicolor Models and Extensions

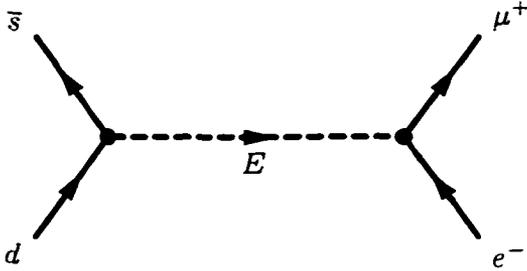


Figure 2.8: $K_L^0 \rightarrow \mu e$ via ETC boson or technipion.

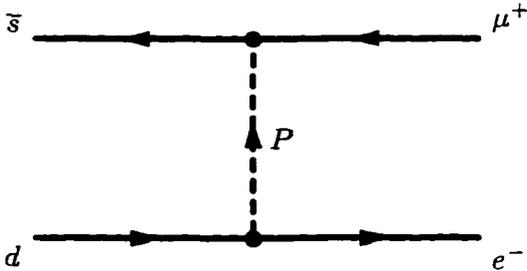


Figure 2.9: $K_L^0 \rightarrow \mu e$ via pseudoscalar leptoquark or vector leptoquark.

One major problem with the Standard Model is that particle masses are generated by adding a scalar term to the Lagrangian. These scalar fields give the Standard Model a large number of arbitrary parameters that must be tuned to extremely high precision. Technicolor (TC) theory does away with these parameters by proposing a dynamic breaking of the $SU(2)_L \otimes U(1)$ electroweak symmetry. In this model, the Higgs is composed of massless fermions bound by a technipion. Unfortunately, while this solves the scalar particle problem, it now removes the mechanism by which quarks and leptons acquire their masses. Extended TC (ETC) models [25, 26] were then proposed to solve the problems with quark and lepton masses in ordinary TC. In this model, the ETC bosons couple normal fermions to technifermions. ETC models predict a wide variety of signatures [27] and allow $K_L^0 \rightarrow \mu e$ via Figure 2.8, where E is an ETC boson or a technipion, or via Figure 2.9 where P is a pseudoscalar or Pati-Salam type vector leptoquark. Several ETC predictions [28–31] place $B(K_L^0 \rightarrow \mu e)$ above 10^{-10} . This presents a problem for TC, not only because this value is above the current experimental limit [32] but because it is difficult to adjust the theory to fit to this new limit. For TC to work properly, the TC scale, Λ_{TC} , must be on the order of 300 GeV. The fermion masses are related to the ETC boson mass by

$$m_f \approx g_{ETC}^2 \Lambda_{TC}^3 / m_{ETC}^2. \quad (2.9)$$

Therefore, it is not viable simply to adjust m_{ETC}^2 higher to account for the lower limit. One suggested solution to this problem has been to allow the coupling constant, g_{ETC} to vary slowly over a large energy scale. Referred to as ‘walking technicolor’ [33–36], this model allows the ETC mass scale to be increased while still generating the correct masses for ordinary fermions. Nevertheless, $B(K_L^0 \rightarrow \mu e)$ still presents a strong constraint to ETC models.

2.3.5 Composite Models

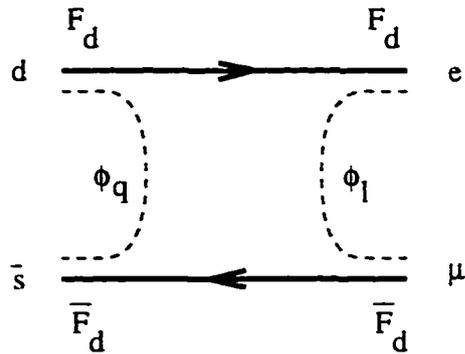


Figure 2.10: Composite model diagram for $K_L^0 \rightarrow \mu e$.

Given the number of particles which must be described by the Standard Model, it is easy to envision that these particles are not fundamental but actually made up of more basic constituents, called preons. One simple composite model [37] contains two types of preons: flavons and chromons. The flavons carry only the flavor characteristics while the chromons carry only the color. These preons are bound by a new strong-like force to form quarks and leptons. In this model, the quarks and leptons in a family contain the same flavon content, but differ in their chromon content. $K_L^0 \rightarrow \mu e$ can occur via flavon and chromon transitions, which are mediated by the strong preonic force.

Another composite model [38, 39] predicts four families of particles ($e^{(0)}$, $\mu^{(0)}$, $\tau^{(0)}$ and $\tau'^{(0)}$) which are bound states of these two types of preons. The two gauge forces of the model, “metacolor” and “hypercolor”, are of vastly different scale with $\Lambda_M \sim 10^{14}$ GeV and $\Lambda_H \sim 10^3$ GeV. Metacolor binds the preons together to make the $e^{(0)}$ and $\mu^{(0)}$ families (the $\psi^{(0)}$ group) as well as “hyperfermions” and “hyperbosons”, which couple to the hypercolor force. The hypercolor force becomes strong at Λ_H and binds the hyperfermions and hyperbosons together to make the higher mass $\tau^{(0)}$ and $\tau'^{(0)}$ families (the $\chi^{(0)}$ group). The physical e , μ , τ and τ' are mixtures of the $\psi^{(0)}$ and $\chi^{(0)}$ groups in this manner:

$$\begin{pmatrix} \psi_{physical} \\ \chi_{physical} \end{pmatrix} = \begin{pmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} \psi^{(0)} \\ \chi^{(0)} \end{pmatrix} \quad (2.10)$$

where $0.04 < \sin \alpha < 0.10$. $K_L^0 \rightarrow \mu e$ can result from mass mixing between the preons, via Figure 2.10. From [39], the comparison with the rate $\Gamma(K^+ \rightarrow \mu^+ \nu_\mu)$ gives

$$\frac{\Gamma(K_L^0 \rightarrow \mu e)}{\Gamma(K^+ \rightarrow \mu^+ \nu_\mu)} = \frac{1}{2} \left[\frac{g_{HC}^2 \sin^4 \alpha}{G_F \Lambda_0^2 \sin \theta_c} \right]^2 \left[\frac{M_K^2}{M_\mu (M_d + M_s)} \right]^2 \quad (2.11)$$

where g_{HC} is the hypercolor coupling constant ($\sim 1 - 10$) and Λ_0 is the mass of the hyperfermion condensates ($\sim 1 - 2$ TeV). Putting in the conservative numbers for these parameters gives $B(K_L^0 \rightarrow \mu e) \approx 3.3 \times 10^{-11}$ which is the current experimentally measured upper bound for this process. Therefore, E871 will present a strong constraint on this model. However, other versions of this model [40] have countered that lepton flavor violating processes can be suppressed by requiring generation number conservation at the preon level, thus moving the limit below the sensitivity of this experiment.

2.3.6 Supersymmetry

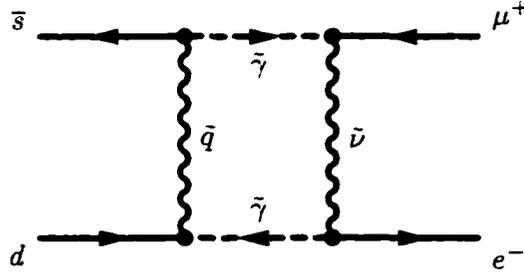


Figure 2.11: Neutralino induced contribution to $K_L^0 \rightarrow \mu e$.

The Supersymmetric (SUSY) models are based on the concept that each known particle has a “super” partner with opposite spin statistics. For example, the electron, a

fermion, has a SUSY partner, the selectron, which is a boson. Likewise, the quark's SUSY partner is called a squark. Lepton flavor violation, shown in the diagram in Figure 2.11, can arise from the mass matrices for these super-partners. Although the total lepton number is conserved, the mass terms are not required to conserve generational lepton numbers. The amount of mixing depends on the size of the squark and slepton mass splitting. Δm_K constrains the magnitude of the squark mass differences while $\mu \rightarrow e\gamma$ constrains the slepton differences. Given the current limit on $\mu \rightarrow e\gamma$ and the value of Δm_K , the upper bound on $B(K_L^0 \rightarrow \mu e)$ is on the order of 10^{-14} [41].

2.4 Theory Summary

The SM does not allow the decay $K_L^0 \rightarrow \mu e$ at any level. Minimal extensions to the SM, with the possible exception of the left-right symmetric model, also do not predict $K_L^0 \rightarrow \mu e$ at any level measurable by E871. If $K_L^0 \rightarrow \mu e$ events were seen at the SES level expected for E871, they would be evidence for physics beyond the SM. Horizontal gauge models, Pati-Salam models, technicolor models and composite models all allow for a $B(K_L^0 \rightarrow \mu e)$ level of 10^{-12} . On the other hand, TC models and composite models would be seriously constrained by the absence of a $K_L^0 \rightarrow \mu e$ signal at that level. While supersymmetry would be unaffected by a negative $K_L^0 \rightarrow \mu e$ result, an observation of $K_L^0 \rightarrow \mu e$ would be a problem for current SUSY models.

Chapter 3

Experimental Setup

3.1 Overview

Experiment E871 was carried out in the B5 beamline at the Alternating Gradient Synchrotron (AGS) of Brookhaven National Laboratory (BNL). A schematic drawing of the apparatus is shown in Figure 3.1. The AGS accelerated up to 60×10^{12} protons per pulse (60 Tp) to an energy of 24 GeV. The beam spill time was 1.4 seconds and the total cycle time was 3.2 seconds, although these values were sometimes increased for power consumption reasons. A good portion of the beam, typically 15 Tp, was extracted down the B5 beamline. Here, the protons struck a rectangular platinum target, resulting in a spray of strongly generated particles. Photons from π^0 's were converted in lead foils and they, along with the rest of the charged beam, were magnetically swept out, leaving the neutral beam to travel down the rest of the beamline. This beam, primarily neutrons and kaons, entered a vacuum decay region and eventually entered the beam plug. A fraction of the kaons (approximately 10%) decayed in the vacuum region and the daughter particles exited the region through a thin Mylar/Kevlar window. These particles then entered the spectrometer

section, where they travelled through the six tracking chambers and two magnets, allowing two measurements of their momenta. The second section of the detector comprised the trigger scintillation counters and the particle identification system, which consisted of a gas Čerenkov counter and a lead glass calorimeter for electron identification and a muon hodoscope and muon rangefinder for muon identification. This provided a twofold redundancy throughout the system with regard both to momentum measurements and to particle identification. This redundancy was crucial in order to reach a single event sensitivity at the 10^{-12} level. The importance of this redundancy will be demonstrated in the analysis section.

3.2 Beamline

3.2.1 Target

The requirements for a target for this experiment were stringent. First, the target must be small to tightly localize the kaon origin. Second, the target must be able to withstand the heat generated by interactions with a very intense proton beam. Third, the amount of material around the target must be low and, preferably, of a low Z material to minimize stray interactions which may make the kaon origin ambiguous. This last requirement had the unfortunate effect of interfering with the pursuit of the second requirement because the connection of the target to a heat sink was made more difficult.

The 1995 target consisted of a 127.0 mm long x 3.15 mm wide x 2.540 mm high platinum piece brazed onto the stem (127.0 mm x 4.83 mm x 9.53 mm) of a water cooled beryllium heatsink. After being brazed onto the beryllium, the platinum was cut into five equal length sections in the z -direction. This was done to minimize thermal stresses due to uneven energy deposition in the target. A thermocouple was placed on the third target

BNL Experiment 871 - The Search for $K_L^0 \rightarrow \mu e$

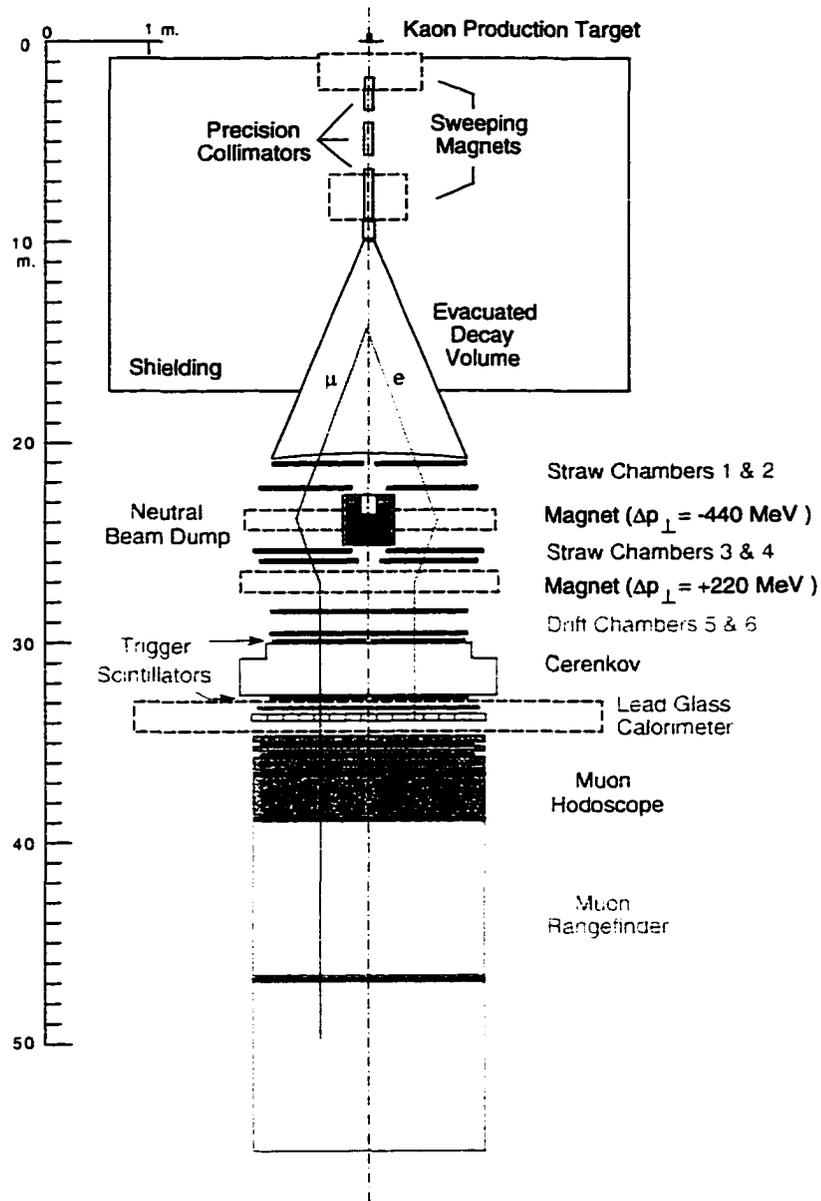


Figure 3.1: E871 Experiment Schematic Drawing.

section to monitor the temperature of the target. The target was cooled with demineralized water pumped through a line cut into the beryllium base. This target was replaced by another identical target after the braze melted during a high intensity beam test and the first two sections separated from the base. In the investigation of the incident, it was determined that the alloy used in the braze (60% Ag, 30% Cu and 10% Sn) formed a Ag-Cu eutectic which had a melting point below either of the two materials alone. From that time on, E871 was limited to an intensity determined from the target temperature. The 1996 target had the same cross sectional dimensions as the 1995 target but was 150 mm long and was cut into fifteen pieces to further help it survive the thermal stresses. The braze used in this target was approximately 92.5% Ag, 7.5% Cu and 0.025% Li. While the tensile strength of the Ag-Cu-Sn alloy is higher than that of the Ag-Cu-Li alloy at room temperature, the relationship is reversed at high temperature, due to the enhanced mobility of the Sn atoms [42]. Thus, the second choice of braze material was better at our target operating temperature. In addition to the new braze, the Be stem connecting the new target to the base was widened to 7.24 mm to increase heat flow to the heat sink. One Pt target piece was broken off during the cutting process but the loss did not greatly affect kaon production in the target.

The target stand was placed in the B5 beamline at an angle of -3.75° to the horizontal with the beam angled to strike normal to the face of the target. This angle was chosen as a compromise between maximizing kaon production and minimizing neutron production down the beamline. The target effective length (after cutting the slots) of 124.50 mm was 1.41 interaction lengths in platinum. With this length, we expected 10^8 kaons and 10^9 neutrons for 20×10^{12} protons, or 20 Tp, incident on the target.

3.2.2 Beamline Magnets and Collimators

The first section downstream of the target contained the first sweeping magnet (B5P4), a set of collimators, and then a second sweeping magnet (B5P5). The sweeping magnets removed charged particles, resulting in a neutral beam within the decay volume. The first sweeping magnet contained a set of eighteen 5 mm lead foils to convert photons generated at the target or from $\pi^0 \rightarrow \gamma\gamma$ decays, into e^+e^- pairs, which were then swept out by the magnets or stopped by the collimators. The precision, lead-lined, brass collimators insured a small divergence of the neutral beam from the beamline. The opening angle was 5 mrad (in X) \times 20 mrad (in Y) allowing a 100 μ steradian beam to exit into the decay region. The second magnet removed K_S decay products and any charged particles generated from the collimators.

3.2.3 Vacuum Decay Region

Because only those decays whose vertices could be traced back to within this vacuum decay region are accepted, this region must be at vacuum low enough to allow the kaons to decay with minimal interference from scattering or production on air molecules. This region must also be large enough to allow a large fraction of the K_L to decay.

This region consisted of an evacuated decay tank (10.9 m in Z; upstream end: 10 cm in X, 16 cm in Y; downstream end: 193 cm in X, 86.4 cm in Y) with a 0.127 mm Mylar/0.381 mm Kevlar combination window on the downstream face. The tank was constructed from 5 cm thick steel and encased in borated concrete, to absorb thermal neutrons generated upstream. The 0.127 mm Mylar was sufficient to make an air tight seal and the Kevlar epoxied to the Mylar provided the strength for the window. The window was held in place by a flange bolted to the tank assembly. The window was replaced after a catastrophic failure at the end of the 1995 running period, in which all of the straw chambers were

seriously damaged or destroyed. The replacement window was 0.127 mm Mylar/ 0.432 mm Ballistic Grade Kevlar. A shutter system was installed after the accident to increase the safety factor when personnel were performing maintenance in the beam cave area. During both running periods, the pressure was maintained at or below 10 $\mu\text{m Hg}$.

3.3 Spectrometer

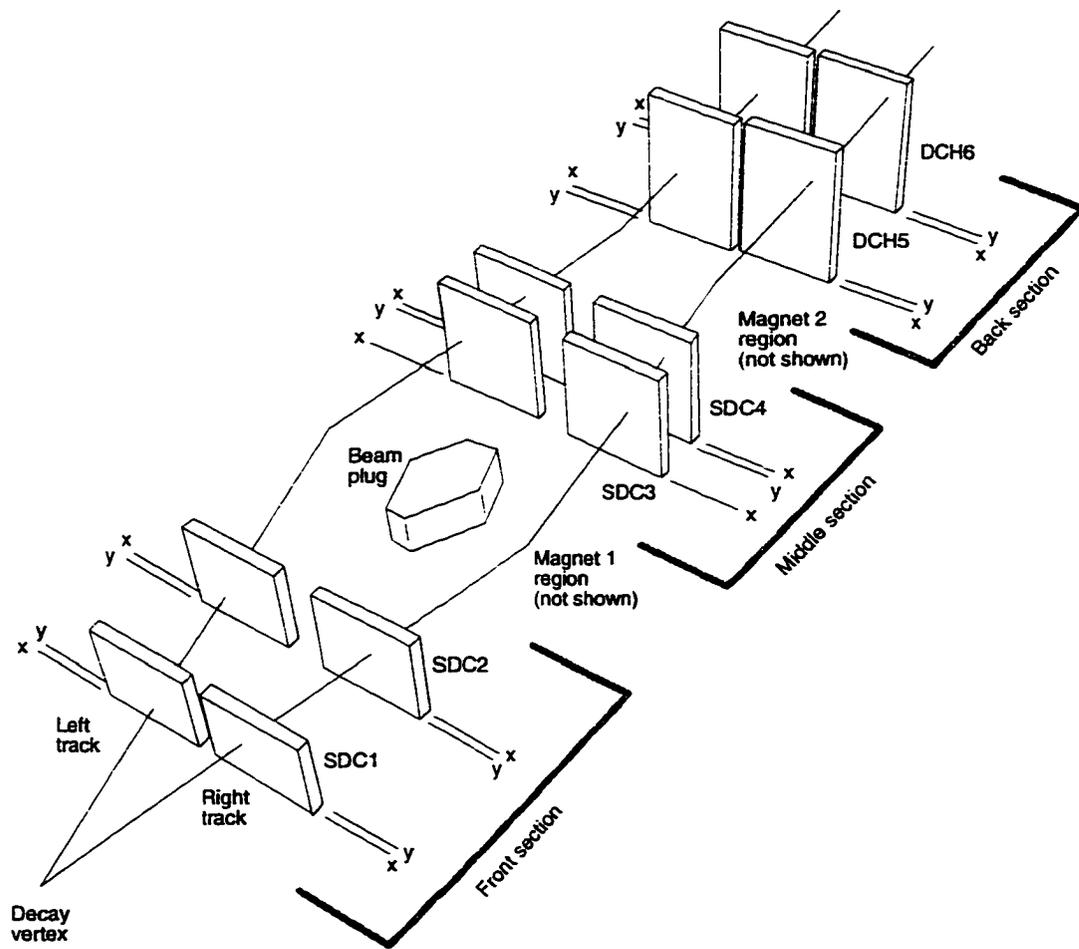


Figure 3.2: E871 Spectrometer Overview.

The front half of the experimental setup was the spectrometer section, shown in

Figure 3.2, with two bending magnets and 22 tracking planes: 14 planes of straw chambers and eight planes of drift chambers. These planes determined particle trajectory with high precision, which, when coupled with a well-defined magnetic field, produced a high precision momentum measurement. The 11 tracking planes on each side of the beamline were distributed among four pairs of straw chamber assemblies and two pairs of drift chamber assemblies. The wire layouts for both types are shown in Figure 3.3.

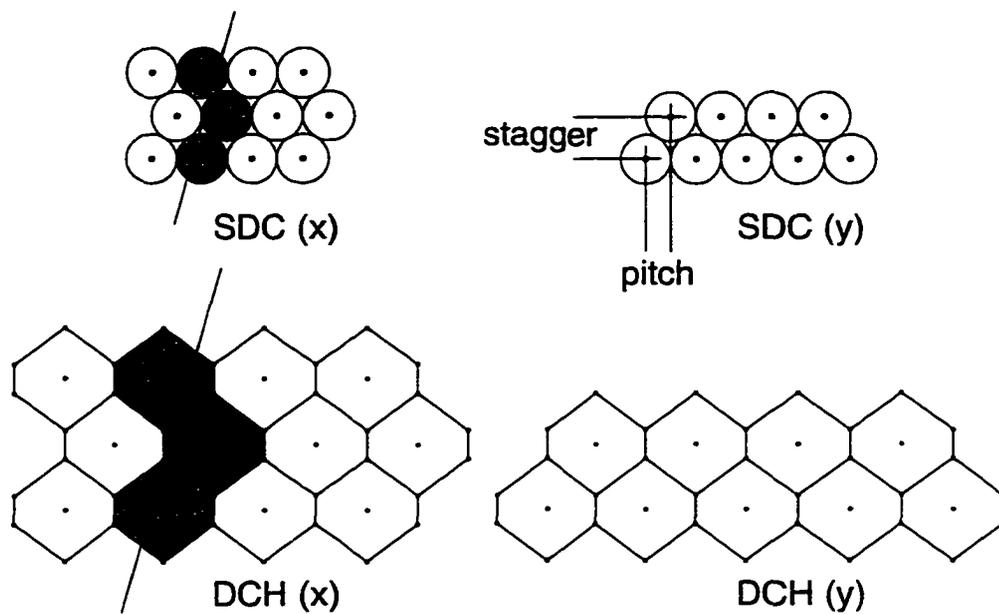


Figure 3.3: Wire layout of the straw (top) and hexagonal-cell (bottom) drift chambers.

Because an accurate momentum measurement was critical for the success of this experiment, the spectrometer was required to be essentially transparent to the traversing particles but segmented enough to lower the occupancy rates arising from stray particles generated in the plug. To that end, the straw material was chosen to minimize particle interactions with the straw wall. In addition, the gaps between the chambers were filled with helium bags to lower the multiple scattering of the traversing particles, further reducing

error in the momentum measurement. The multiple scattering in the helium bags, even when considering the bag material, was much smaller than if air had been allowed to remain in the gaps [43].

3.3.1 Straw Chambers

The interaction of the intense neutral beam with the beam plug was expected to create a large number of low-energy background photons and neutrons resulting in a high occupancy rate in the tracking chambers. It was obvious that normal drift chambers would not work in the E871 configuration and thus the straw drift chambers (SDC) were built. These chambers had a smaller cell size and used a faster drift gas to lower the occupancy in the cells. The SDC were used as tracking chambers 1 through 4, in the high rate region, and were arrayed in the following manner: 1(XY,LR), 2(YX,LR), 3(X,LR), 4(YX,LR).

The straws were 5 mm diameter tubes made from 25 μm thick copper-coated Mylar with a 20 μm diameter gold plated tungsten wire stretched down the center. The straws were copper coated, as opposed to aluminum coated, because 1) copper oxide is conductive and should therefore lessen any cathode field emission effects and 2) the relatively high work function provided some absorption of soft photons in the straw walls [44]. These straws were mounted into holes drilled into the aluminum frames and silver epoxy was used to cement them in place. The tensioned wires were held in the center by a “V” groove on the feedthrough and were fixed in place by brass pins. The pins held the end of the wire against the edge of the plastic feedthroughs as shown in Figure 3.4. These pins were epoxied, for strength, and covered in silicone sealant for the gas seal. Each chamber, with the exception of chamber 3, had three layers of X measuring tubes and two layers of Y measuring tubes. Chamber 3 had the three X layers but did not have the two Y layers. Flowing through the chambers was a 50/50 mix of $\text{CF}_4/\text{C}_2\text{H}_6$, which exhibited a 100 $\mu\text{m}/\text{ns}$ drift velocity. This

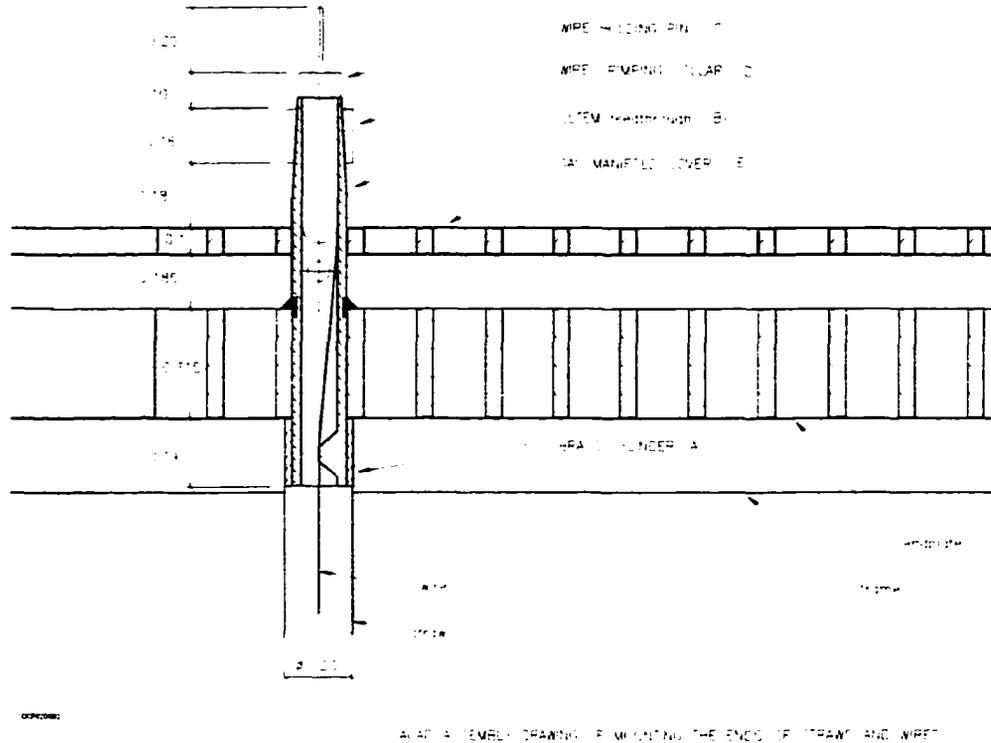


Figure 3.4: SDC Wire Configuration.

large velocity was needed to increase the rate capability in the chambers to accommodate backgrounds from both the neutral beam directly and from particles escaping from the plug.

The chamber electronics started with the amplifier boards, shown schematically in Figure 3.5. These boards connected the wire pins to high voltage, read out the signal via capacitive coupling, discriminated the analog input signal at $\approx 1.5 \mu\text{A}$ and generated a digital output signal 30 ns wide. This signal was sent to 6 bit, 1.25 ns least count TDCs via 32 channel Ansley cables. The high voltage supplied to the wires was generated by a CAEN SY127 power supply, and distributed to the six chamber sections via a HV bus running alongside one side of the pins, as shown in Figure 3.6. This sectioning was done in order to

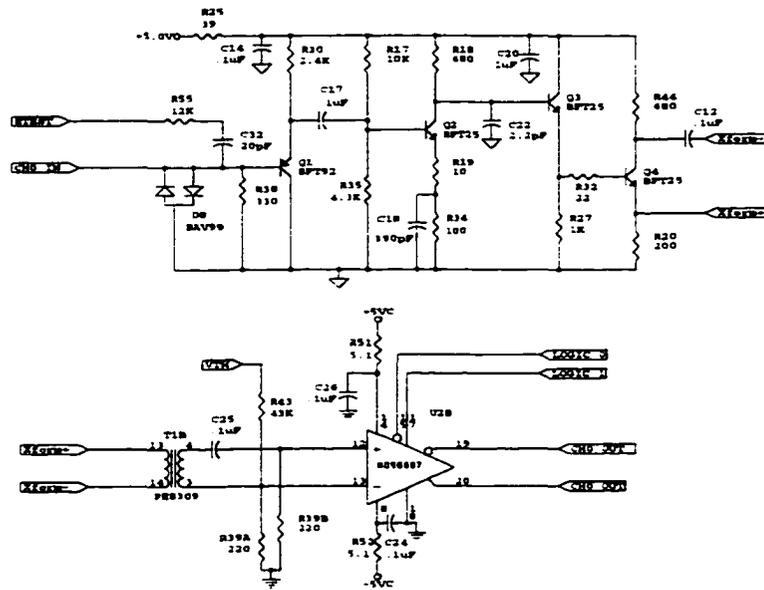


Figure 3.5: SDC Amplifier Circuit Schematic.

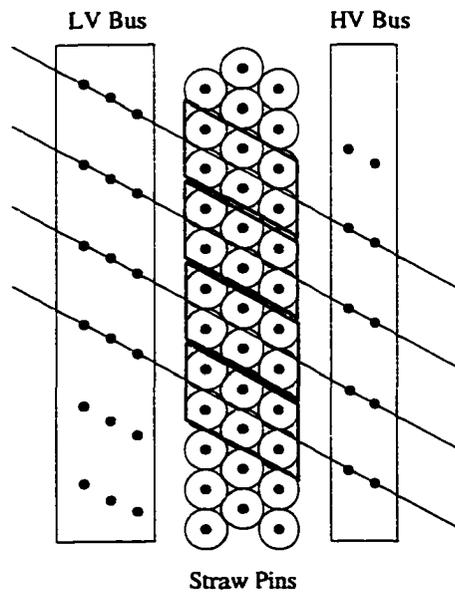


Figure 3.6: Straw Drift Chamber Pin Layout (Not to Scale).

enable easier troubleshooting of a bad wire. In operation, the straw HV was ramped from 1850 V to 1950 V during the spill and back down to 1850 V during out of spill times to reduce the current load from charge accumulation during out of spill times. Low voltage (± 5 V, +6 V) to power the amplifier cards was supplied by rack-mounted power supplies located just to the side of the chambers and was run along a bus on the opposite side of the HV bus, alongside the pins.

3.3.2 Drift Chambers

The drift chambers were used as tracking chambers 5 and 6, where the background was not as high as at the front section of the spectrometer. They were hexagonal cells (Figure 3.3, lower half) with a cell size of 1.016 cm and were filled with an Ar/C₂H₆ mixture (49/49), with 2% ethanol to keep the wires clean. This mixture of gases resulted in a drift velocity of 50 $\mu\text{m}/\text{ns}$, which was sufficient for the rates found in this region of the spectrometer. The field and guard wires were 109 μm diameter Al wires and the sense wires were 20 μm diameter gold plated tungsten wires. The sense wires were soldered to pins on the outside, which were, in turn, connected to preamplifier cards on the outside of the chamber. These signals were passed on to amplifier cards in racks alongside the drift chambers. Here the signals were sent, via Ansley cables, to 6 bit, 2.5 ns least count TDCs. Because the tracks at this point were roughly parallel, the chambers themselves were angled with respect to the beam axis to resolve possible track ambiguities. The efficiency and resolution of the straw and drift chambers are shown in Figures 3.7 and 3.8.

3.3.3 Spectrometer Magnets

The two spectrometer magnets were a 96D40 and a 100D40 magnet. The two magnets were required in order to generate two momentum measurements but also served a

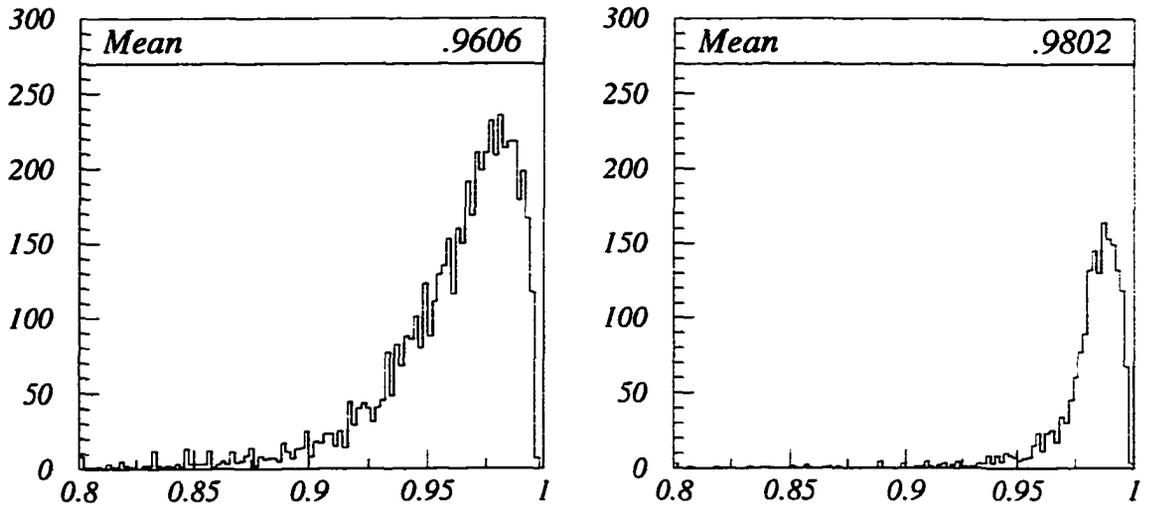


Figure 3.7: Average straw chamber (left) and drift chamber (right) wire efficiencies.

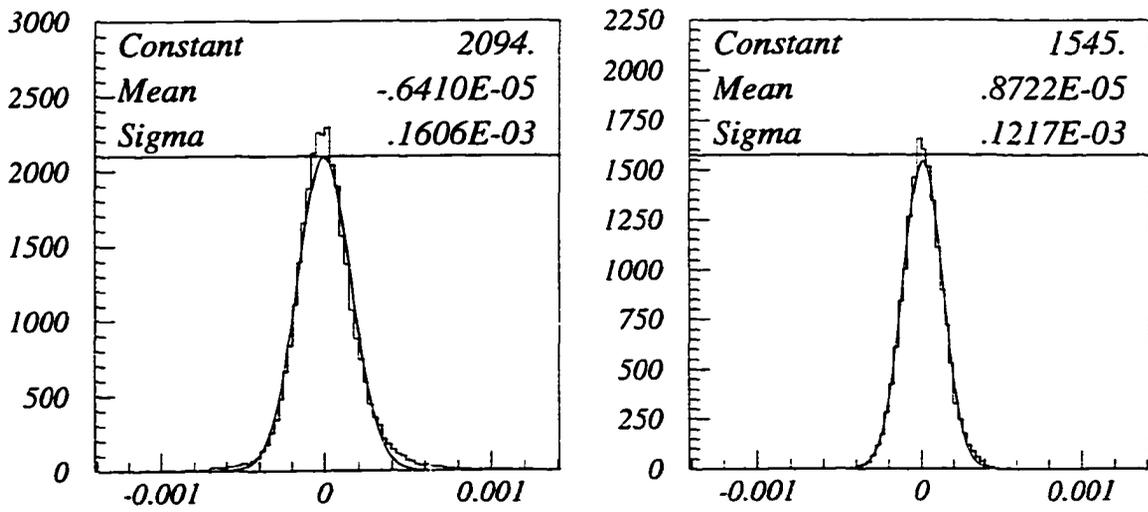


Figure 3.8: Straw chamber (left) and drift chamber (right) resolutions in meters. The resolutions are 160 and 120 μm , respectively.

very useful function for filtering out background events. Listed in Table 3.1 are both the predominant decay modes and the four decay modes of interest to E871. The p_T listed for $K_{\mu 3}$ and $K_{e 3}$ are the maximum allowed p_T for those three-body decays, i.e. $E_\nu \approx 0$, while the p_T for the other decays is the p_T at the Jacobian peak of their decay p_T distribution. The decay modes of primary interest ($K_L \rightarrow \pi\pi$, $\mu\mu$, ee , and μe) all have p_T in the range of 206 - 249 MeV/c. The first magnet was set to give a -440 MeV/c kick to the two tracks while the second magnet gave a +220 MeV/c kick to the tracks, resulting in a net p_T kick of -220 MeV/c. After traversing the field of the second magnet, the typical $K_L^0 \rightarrow \mu e$ decay products would leave the last magnet traveling roughly parallel to each other, while most of the semi-leptonic background decays would not. This track parallelism was required by the trigger system as a necessary condition in order to filter out background events from the data stream. The trigger system will be described in more detail in Section 3.7.2.

Decay Mode	Fraction	Decay p_T (MeV/c)
$K_L \rightarrow \pi^+ \pi^- \pi^0$	$(12.56 \pm 0.20)\%$	133
$K_L^0 \rightarrow \pi^\pm \mu^\mp \nu$	$(27.17 \pm 0.25)\%$	216
$K_L^0 \rightarrow \pi^\pm e^\mp \nu$	$(38.78 \pm 0.27)\%$	229
$K_L^0 \rightarrow \pi\pi$	$(2.067 \pm 0.035) \times 10^{-3}$	206
$K_L^0 \rightarrow \mu\mu$	$(7.2 \pm 0.5) \times 10^{-9}$	225
$K_L^0 \rightarrow ee$	$\approx (9.4^{+5.9}_{-4.6}) \times 10^{-12}$	249
$K_L^r \rightarrow \mu e$	$< 3.3 \times 10^{-11}$	238

Table 3.1: K_L Decay Modes and Decay Transverse Momentum.

3.3.4 Beam Plug

Located in the center of the 96D40 magnet, the neutral beam plug (Figure 3.9) constituted one of the reasons this experiment was able to achieve a factor of 20 in sensitivity over the previously most sensitive experiment. The plug was a structure made of tungsten blocks, copper plates, and polyethylene bricks, sheathed in 2.54 cm lead sheeting, with

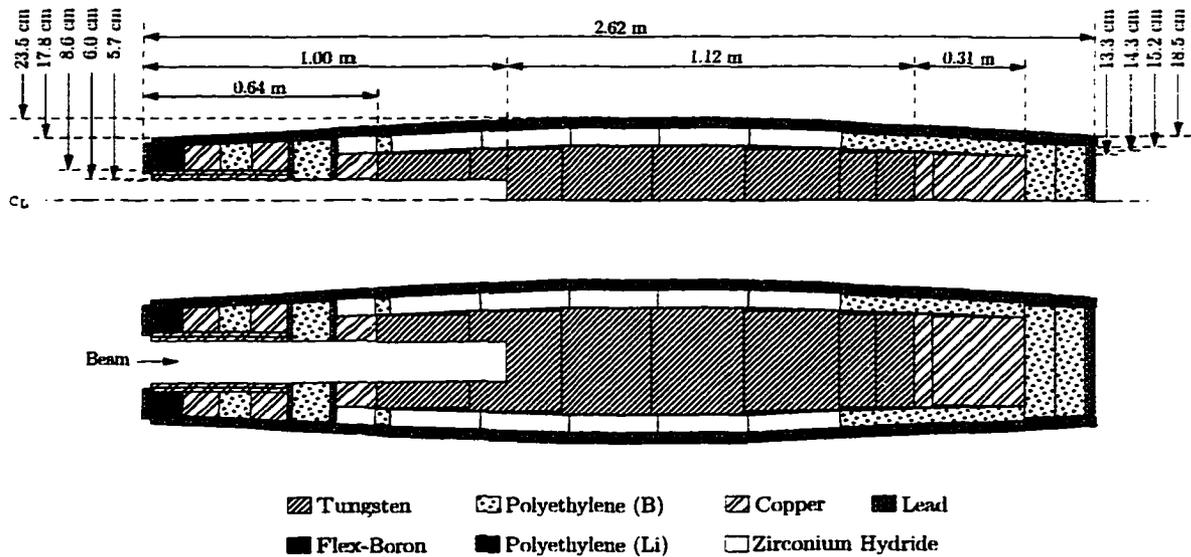


Figure 3.9: An XZ plane cross-sectional view of the beam plug.

zirconium hydride filling the spaces between the bricks. The purpose of the plug was to absorb the intense neutral beam that came down the beampipe. Without the beam plug, the particle identification detectors could not span the beamline as the rates from the neutral beam would then overshadow any true signal from kaon decay. Removal of the neutral beam, in turn, permitted a higher beam intensity without a corresponding increase in detector background noise rates. The beam plug design was based on extensive Monte Carlo simulations of a neutron beam with a momentum distribution peaked at around 9 GeV/c. For more details on the plug performance, see [45].

3.4 Trigger Scintillators

The trigger scintillation counters (TSC) were two sets of plastic scintillator slats which provided the Level 0 (L0) trigger signal that started the triggering process. The first set was located downstream of the sixth drift chamber station and the second set was positioned downstream of the Čerenkov counter, for a separation in the z direction of 2.91 m

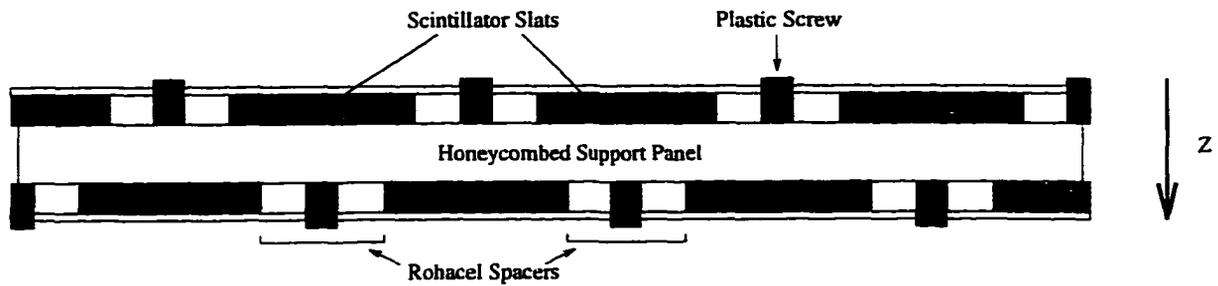


Figure 3.10: Schematic diagram of TSC mounting and support structure.

(see Figure 3.1). The X measuring slats were 32 mm wide \times 5 mm thick \times 1653 mm long and were located in both TSC modules. The Y measuring slats were the same length as the X except they were only 30 mm wide and were located in TSC2 only. While the slats were 30 and 32 mm wide, the center to center separation between adjacent slats was only 27.5 mm. This was accomplished by offsetting adjacent slats in the z direction, as shown in Figure 3.10. This configuration minimized the inactive region between slats that would result if they were positioned edge-to-edge. Each slat was connected to Hamamatsu R1398 photomultiplier tubes (PMTs). The X-view slats had two PMTs connected, one on each end of the scintillator, whereas the Y-view slats had only one tube connected on the outer end of the scintillator. The signals from these tubes were sent over 300 ns of RG-8 coaxial cable to LeCroy 4413 discriminators. The output from the discriminators was then sent to the trigger system and to a set of TDCs. The overview of the TSC configuration is outlined in Figure 3.11 in which the Čerenkov counter has been removed and TSC2 X and Y are separated to show the parts more clearly. The TSC efficiency for each scintillator slat over both running periods is shown in Figure 3.12.

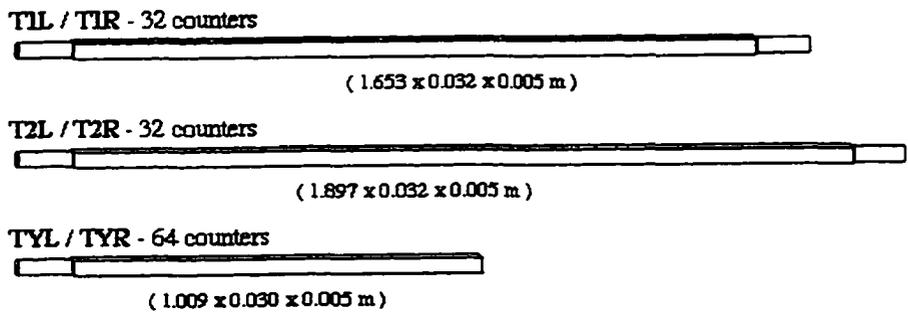
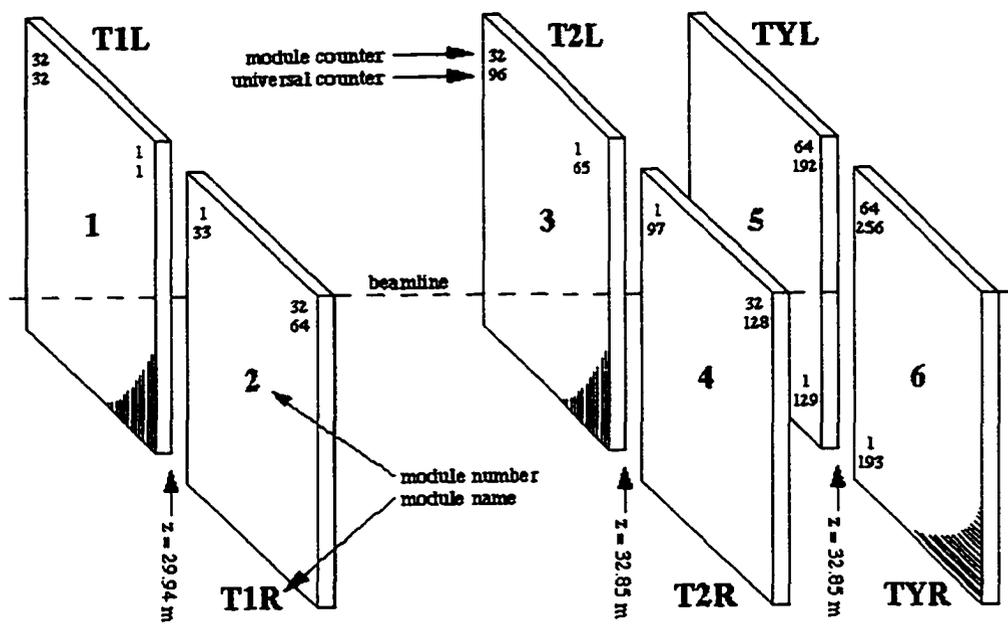


Figure 3.11: Schematic configuration of the TSC.

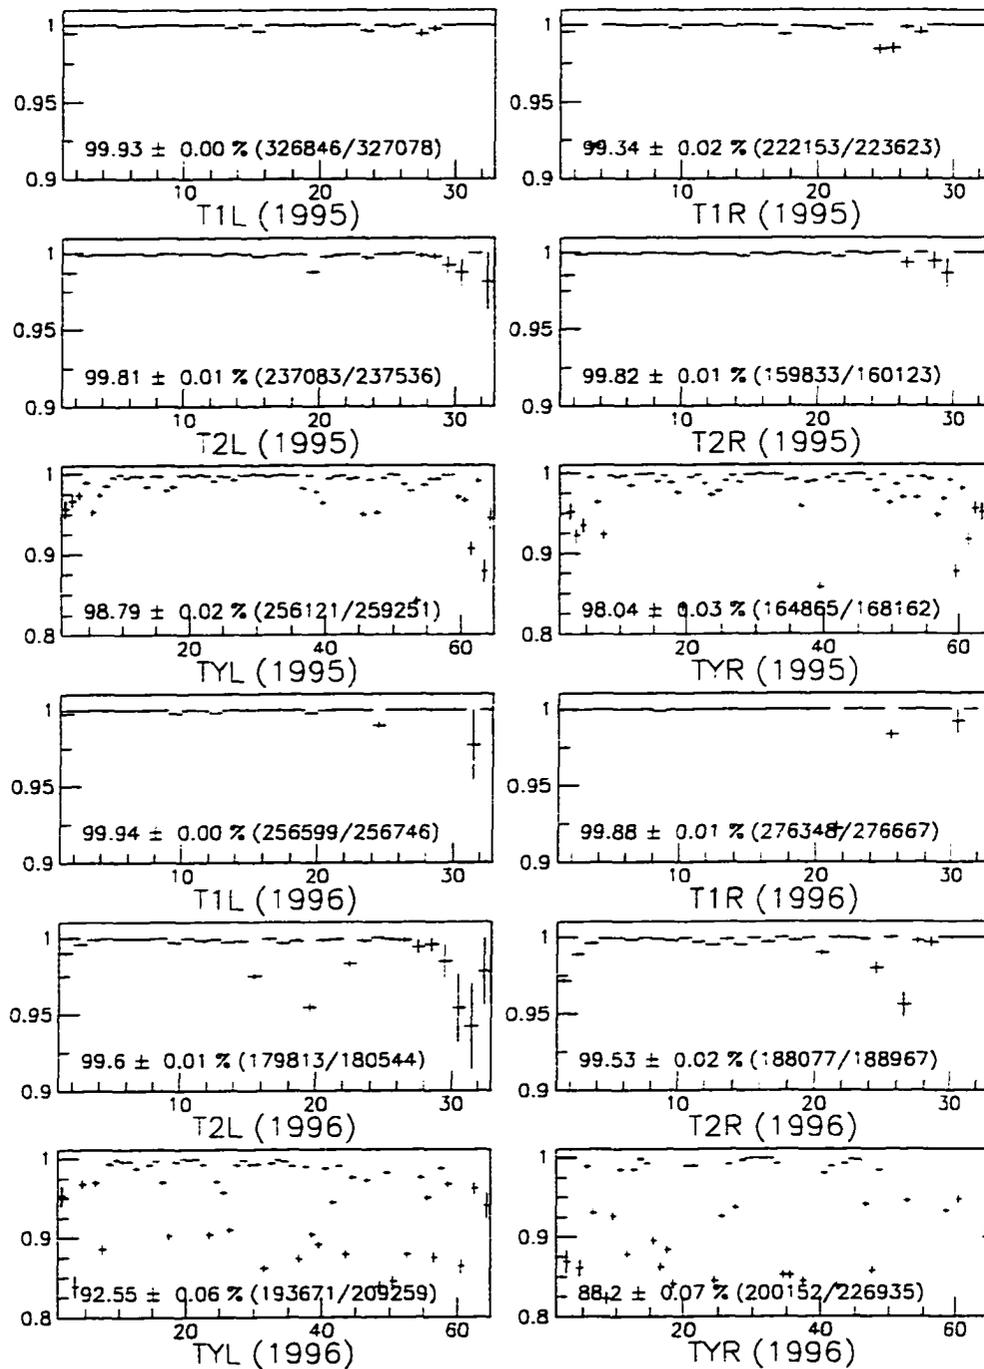


Figure 3.12: TSC efficiencies for each TSC slat.

3.5 Electron ID

The electron identification section of the detector was made up of two components: a hydrogen threshold Čerenkov counter (CER) and a lead glass calorimeter array (PbG). In addition to the obvious redundancy in electron ID, the advantage to using this arrangement is that these detectors perform best in complementary momentum ranges. Electron differentiation in the CER works best at low momentum while, for the PbG, the best range is higher momentum. The CER also provided a signal to the Level 1 trigger system for use with the various electron triggers.

3.5.1 Čerenkov Counter

The main purpose of a threshold Čerenkov counter is to differentiate among different types of particles that have the same momentum. Čerenkov light is generated when a particle is traversing a medium faster than the speed of light in that medium (or $\beta > 1/n$, where n is the index of refraction of the medium). β can be calculated by the equation

$$\beta = \frac{v}{c} = \frac{p}{\sqrt{p^2 + m^2 c^2}}. \quad (3.1)$$

From Equation 3.1, it is apparent that a lower mass particle will emit Čerenkov light at a lower momentum than a higher mass particle. This is key to identifying electrons and excluding the heavier muons and pions.

Another important relation is the differential number of photons generated at a given wavelength. The differential number of photons generated at wavelength λ in a detector of length L at a Čerenkov cone angle of θ_c is

$$\frac{dN}{d\lambda} = 2\pi\alpha L \frac{\sin^2\theta_c}{\lambda^2}. \quad (3.2)$$

Owing to the $1/\lambda^2$ dependence in the number of photons generated, it is very advantageous to have the ability to detect short wavelengths. Therefore, the need for high quality PMTs with good UV response is strong. The problem becomes one of finding the proper medium and apparatus to allow for the differentiation of the particles of interest, in this case the e , μ , and π .

Design Considerations

The requirements for the design of the E871 threshold Čerenkov counter were:

1. Given the momentum range of the accepted kaon daughters, the medium needed an index of refraction close to unity.
2. The medium needed a low Z , to decrease the incidence of “knock-on” electrons.
3. The PMT must have good UV response to maximize photon collection.
4. The PMT magnetic shielding must keep the field below 0.5 Gauss (transverse and longitudinal) for optimal PMT response.
5. The PMT base must be able to handle high rates in the PMT.
6. The active area required twice the segmentation of the E791 CER.
7. High mirror quality was required for good light collection.

To address the first two concerns, H_2 gas at slightly over atmospheric pressure was used as the medium. This complicated the design of the PMT base but that problem was overcome, as described below. The third requirement was met by using a 12 cm diameter Burle 8854 PMT, which has good UV response. The fourth concern led to the development of a multilayer shielding scheme. The fifth issue led to the design of a PMT base with

high rate capability. The sixth and seventh requirements were met by the creation of 32 spherical glass mirrors, each with a radius of curvature and low surface roughness within allowed tolerances.

Čerenkov construction

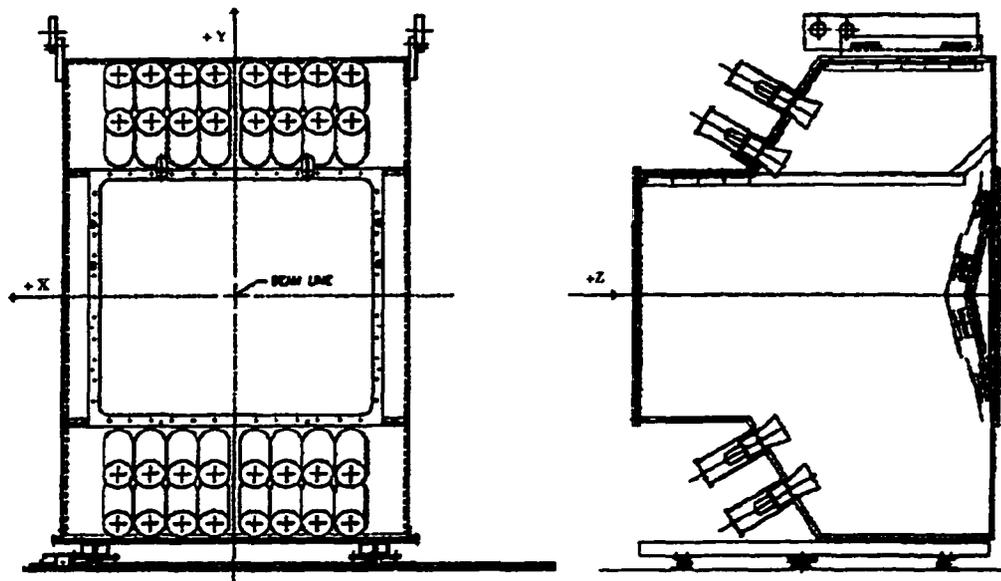


Figure 3.13: Čerenkov Enclosure with Mirrors and Phototubes Shown (Front View and Side View).

The E871 Čerenkov counter, shown in Figure 3.13 was a 18 m^3 aluminum enclosure, constructed by BNL out of 2.5 cm thick aluminum plate. The Čerenkov counter was 2.75 m long in the Z-direction and the K_L decay products traversed an average gas path length of 2.37 m from the front mylar window to the mirror faces. The CER had 32 cells, each cell consisting of a rectangular mirror and a PMT. The mirrors spanned a total area of $\pm 97.7 \text{ cm}$ in X by $\pm 87.5 \text{ cm}$ in Y. The 2.9 m radius, spherical mirror blanks were fabricated by Lancaster Glass Corp.¹ Two types of mirror blanks were produced. The narrow blanks

¹Lancaster Glass Corp., 240 W. Main St., Lancaster OH, 43130

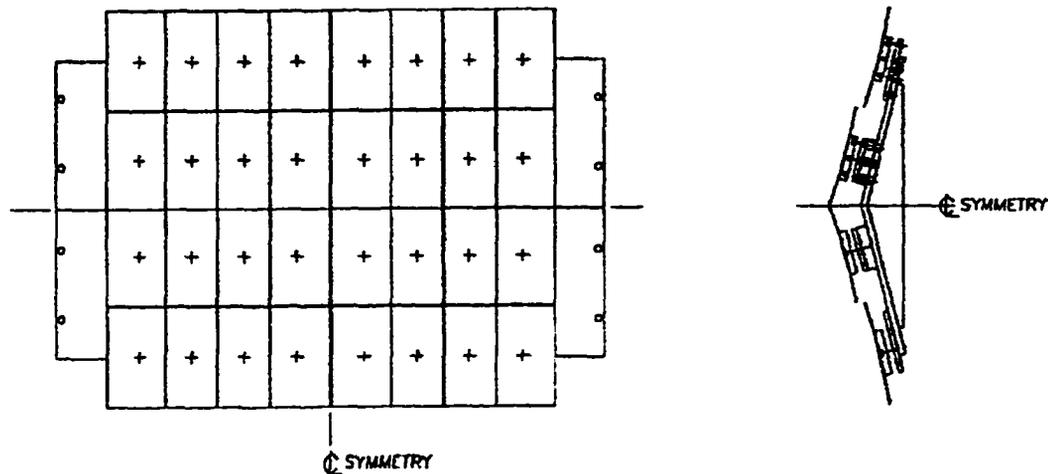


Figure 3.14: Čerenkov Mirror Configuration with Rohacel Mounting Board.

were 0.231 m wide \times 0.457 m high and were located in the inner two columns on each half of the detector (see Figure 3.14). The wide blanks were 0.258 m \times 0.457 m, and were located on the outer two columns of each half of the detector. The mirror blanks were tested at BNL before being delivered to Acton Research Corp.² for coating. All of them were well within the 2.5 - 3.1 m radius tolerance and 30 nm surface roughness that was acceptable for E871 (see appendix A.1). Each mirror blank was coated with a 20 nm thick Al reflective coating and a protective 1.5 nm thick coat of MgF, to prevent oxidation of the aluminum. After coating, the mirrors were brought back to BNL where they were mounted to Rohacel pads with MasterBond EP21 TDCND epoxy. Rohacel, a foam-core laminate material, was used by Composiflex Corp.³ to construct the mirror mounting assemblies in order to minimize the amount of material in the path of the K_L decay products. Embedded in these pads were nylon dowels with screw holes for both the mounting and mirror adjustment screws. The mirror assemblies were each held in place against a Rohacel mounting board by three

²Acton Research Corp., 525 Main Street, Acton MA, 10720

³Composiflex Corp., 2101 Penninsula Dr., Erie PA, 16506

nylon screws. Three other screws allowed adjustment of the mirror angle and distance from the mounting board. At the same time the mirrors were being epoxied to the Rohacel pads, a small grounding wire was attached to the mirror surface at an edge via a minute dab of silver epoxy (Tra-Duct 2902). These grounding wires were used to help reduce dust buildup on the mirrors by minimizing static electricity on the mirror surface.

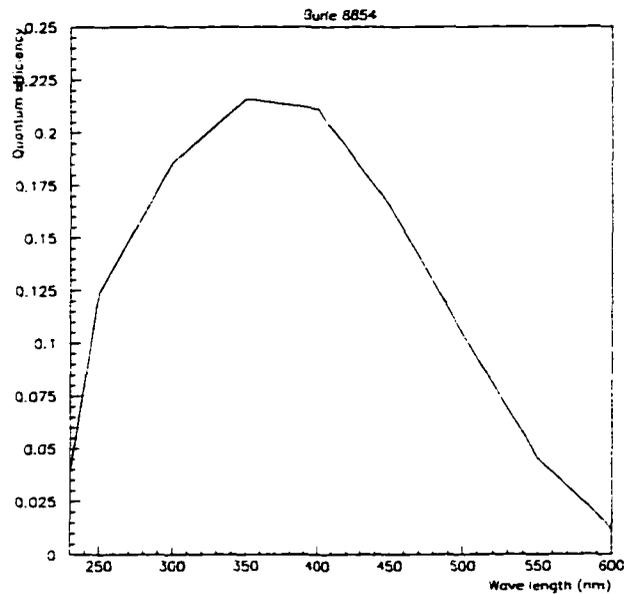


Figure 3.15: Burle 8854 Quantum Efficiency vs. Wavelength.

The PMTs were Burle 8854 Quanticon 12 cm diameter phototubes [46], sensitive to photon wavelengths down to 220 nm. They have a single photoelectron response with a high quantum efficiency of 22.5% at 385 nm (see Figure 3.15). Because the face of the PMT would be in a hydrogen atmosphere, positive HV was required, placing the face of the PMT at ground potential. This HV scheme introduced the possibility of a signal baseline shift when running at high rates, owing to the capacitive coupling on the signal out of the base. However, this was later shown to be a small effect. Another, much larger, effect of high noise rates was a signal shift due to sagging high voltage in the later dynode stages,

which resulted from the higher current present in the PMT at those higher rates. Therefore, the custom-made tube bases, shown schematically in Figure 3.16, were designed with high voltage Zener diodes in the later stages to avoid signal sagging at high rates. The testing of these bases is described in more detail in appendix A.2.

The magnetic shielding, shown in Figure 3.17, was required to compensate for the fringe field from the 100D40 magnet. For a PMT to operate properly, the local axial and transverse magnetic field components must be less than 0.5 Gauss. The fringe field measured in the CER region ranged from 7 to 13 Gauss, which presented a problem. This problem was compounded by the need to have the face of the PMT inside the hydrogen gas and by the design of the gas seal required to keep the hydrogen away from the PMT base, which would be operating at greater than 2000 V. These problems were solved by an unusual set of interlocking shields which allowed the PMT face to sit in the hydrogen but still provided adequate shielding and a gas seal between the hydrogen and HV base. The PMT adapter was constructed from soft iron pipe, which provided magnetic shielding around the photocathode and the first dynode stages, and formed the support for the PMT *in situ*. The PMT was cemented in the PMT adapter with Sylgard 170, which provided both the structural support and the gas seal. This assembly was attached to the CER assembly through portals and the seal for the portal was a radial O-ring located in a groove in the PMT adapter. Over this assembly was placed the next two parts of the shielding. The PMT shield, which slid over the rear half outside the enclosure, and the cone shield, which slid inside the PMT adapter inside the enclosure, were made out of CoNetic material. This metal alloy readily absorbs magnetic fields but saturates fairly quickly. The cone shield was also lined with reflective Mylar, to enhance the collection of photons. On the outside of the enclosure was a thick iron pipe, which encased the rest of the shielding and absorbed most of the magnetic fringe field. Nitrogen flowed through this pipe to flush out any hydrogen that might have leaked

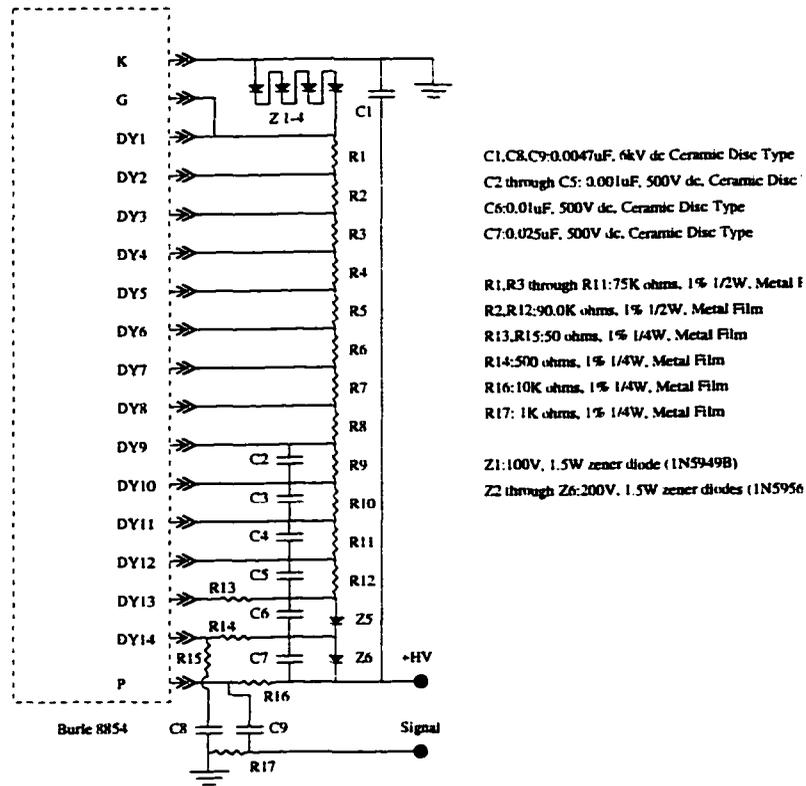


Figure 3.16: Čerenkov PMT Base Schematic.

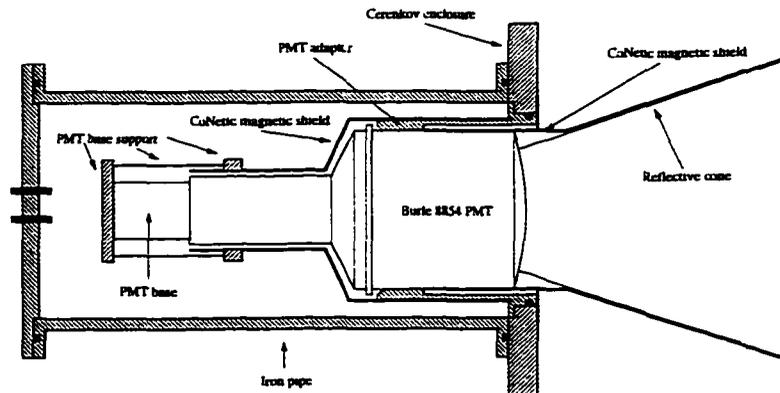


Figure 3.17: CER PMT Assembly.

through and might otherwise collect near the PMT base. It also served to flush out any helium that might leak in and damage the PMTs. The design and testing of this shielding scheme is described in more detail in appendix A.3.

The front window of the CER enclosure was a double window arrangement with a 0.0375 mm black Tedlar sheet, to provide a light seal, and a 0.127 mm translucent Mylar sheet, to provide a gas seal. These two windows were separated by a 2.5 cm gap and a slow flow of nitrogen was sent between these windows to flush out any hydrogen that might have penetrated the Mylar window. The front window was made as thin as possible to minimize the possibility of knock-on electrons, which might then generate Čerenkov light. The downstream window was 1.6 mm thick Al sheet, since the amount of material in the downstream window was not as important as that in the upstream window. There was also a set of Tedlar sheets at beam center that formed a vertical partition between the two sides of the counter. This was done to keep photons from a track on one side from hitting mirrors on the other side and generating a false electron signal.

The HV for the PMTs was generated by a LeCroy 1440 HV power supply. The PMT signals were sent via 300 ns of RG-8 coaxial cable running to a patch panel into the Level 1 trigger room on the lower level of the counting house. Here the signal was split 80/20 with 80% going to the ADCs and 20% going to LeCroy 4413 discriminators and then into both the Level 1 trigger, as the electron PID signal, and the 8 bit TDCs. Figure 3.18 shows the ADC pulse height spectrum for well identified electrons.

With the counter working at an operating pressure of 7.6 cm H₂O overpressure of hydrogen, the momentum thresholds for a CER response for e, μ , and π were 0.03, 6.3 and 8.3 GeV/c, respectively. These thresholds provided very high efficiency for electron detection over the entire accepted momentum range while rejecting the large majority of muons and pions in that same momentum range. The efficiency of the CER for each particle

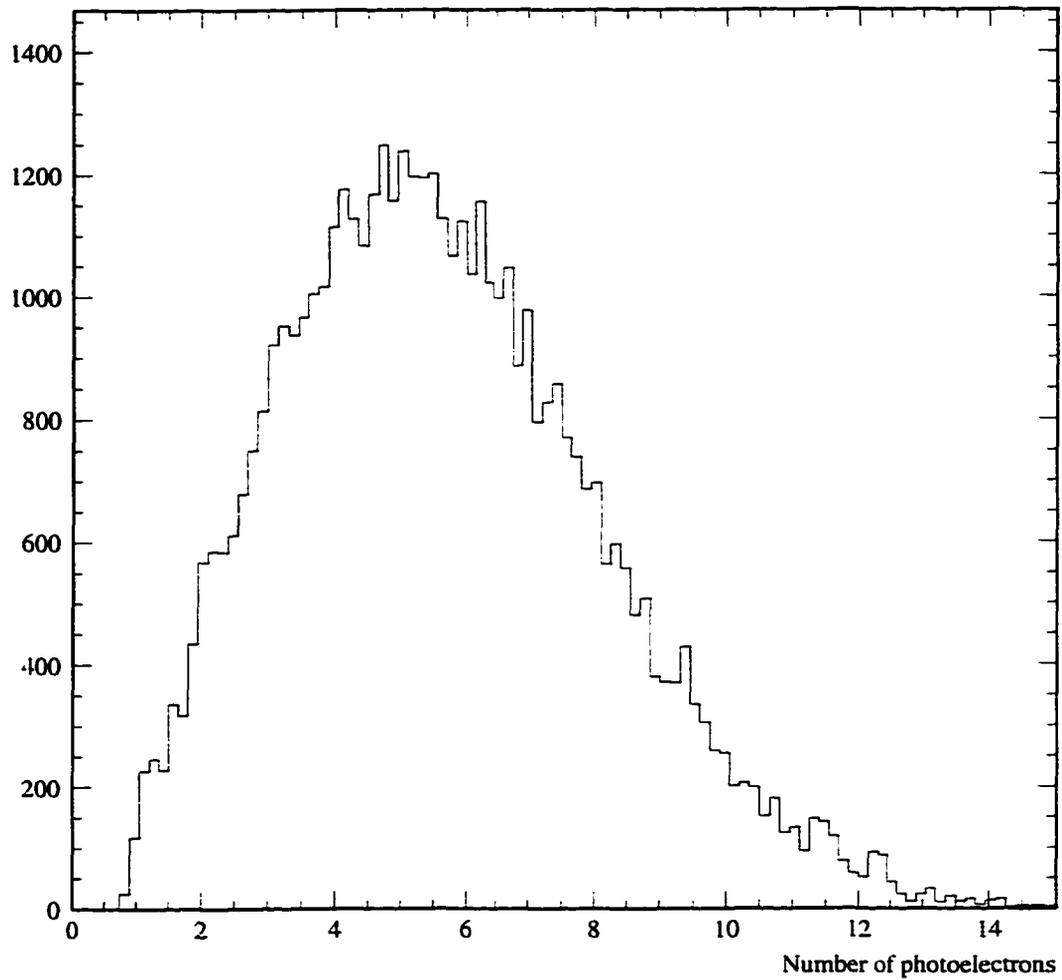


Figure 3.18: CER ADC response plotted in number of photoelectrons.

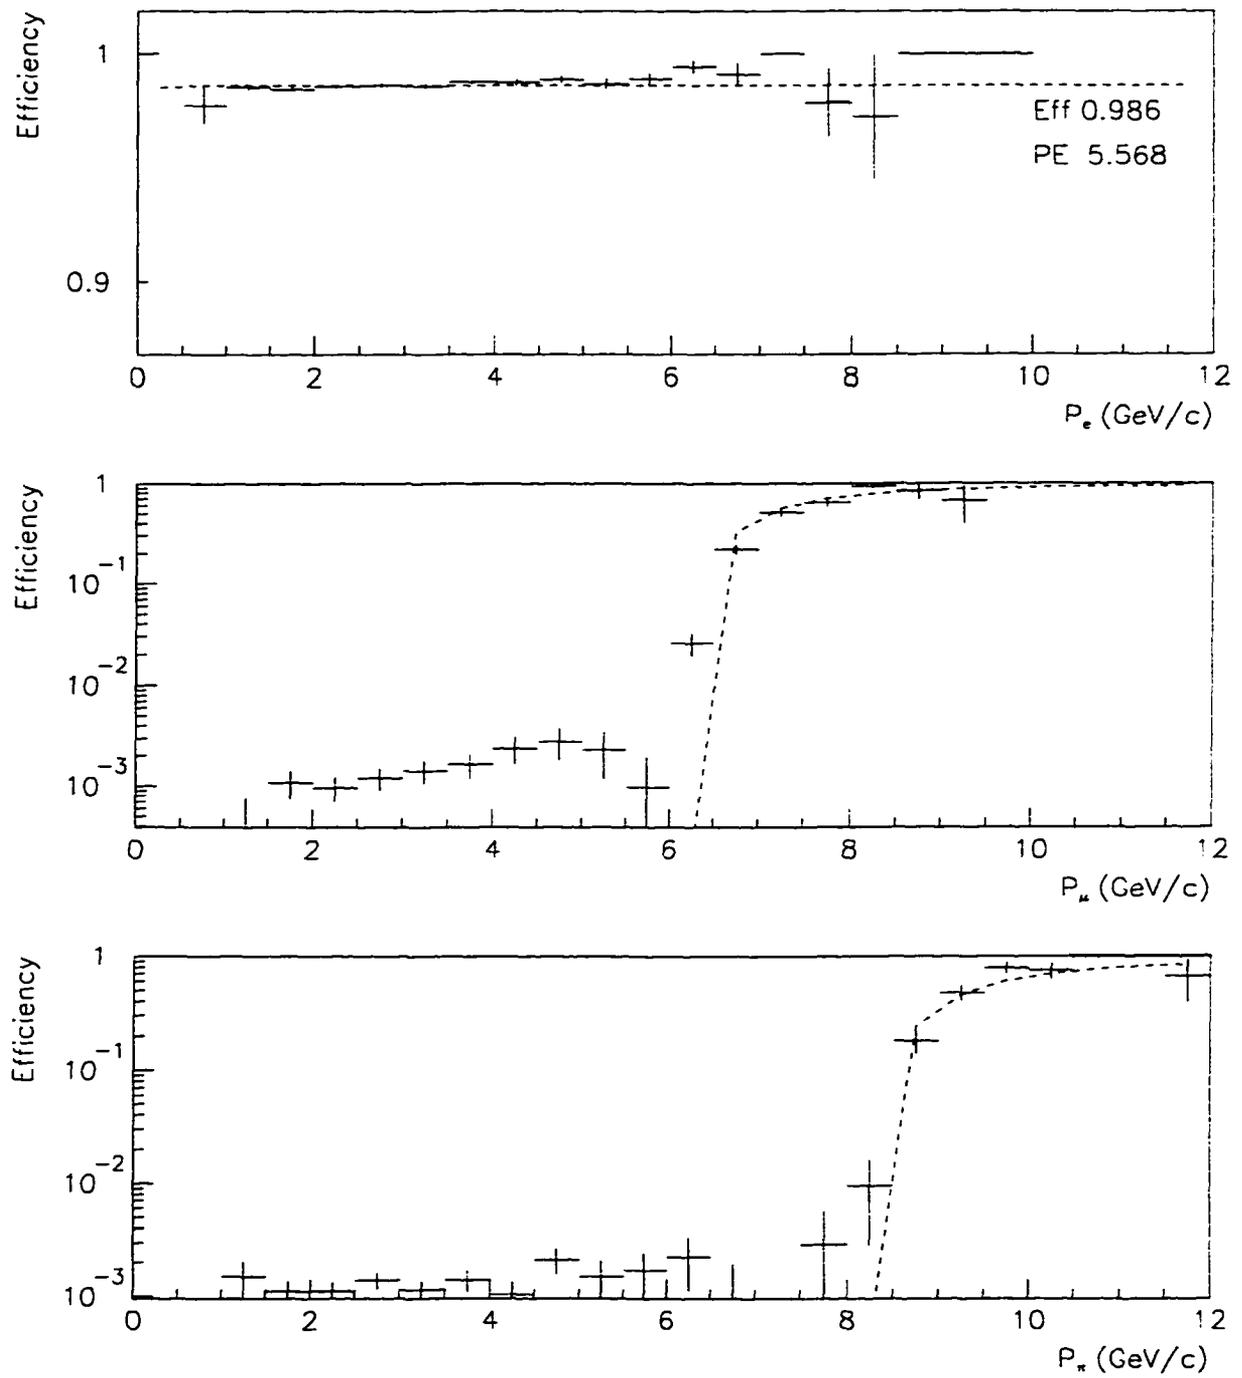


Figure 3.19: CER Efficiency vs. Track Momentum for e , μ , and π .

•

type is plotted over a momentum range of 0-10 GeV/c in Figure 3.19. These plots also clearly show the threshold momentum values for μ 's and π 's.

3.5.2 Lead Glass Calorimeter

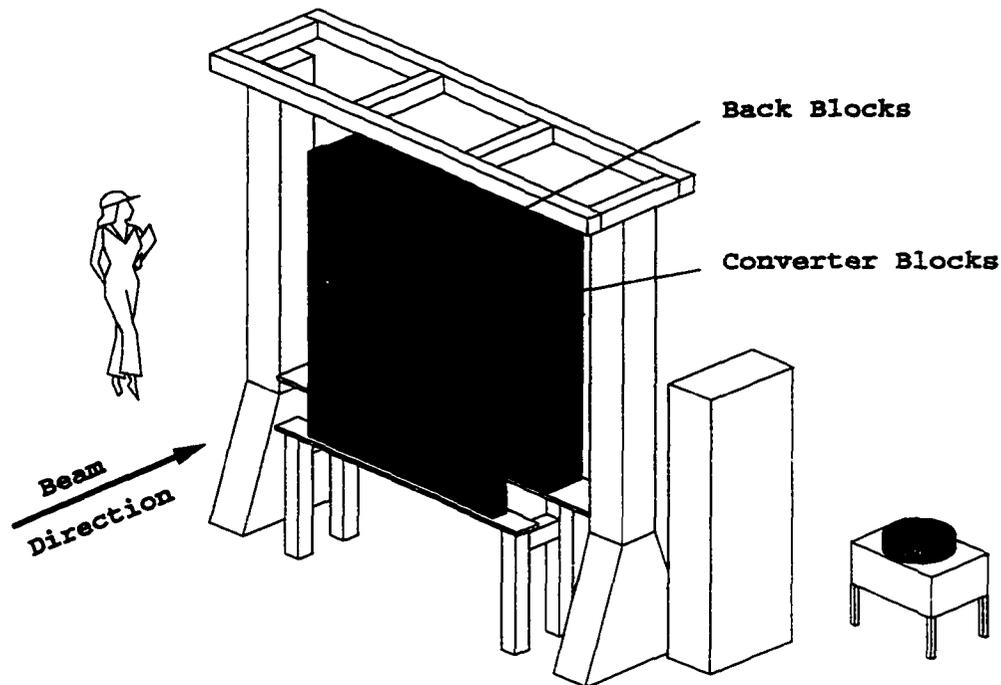


Figure 3.20: PbG Array Layout.

The lead glass array, shown in Figure 3.20, was an arrangement of two sets of lead glass blocks: the converter blocks and the absorber blocks. The converter blocks, with the dimensions 10.9 cm in X by 90.0 cm in Y by 10.0 cm in Z, were placed vertically at the front of the array, two blocks to a column, spanning the Y direction, with 12 cm PMTs at the outer ends of each block. The absorber blocks, with the dimensions 15.3 cm in X by 15.3 cm in Y by 32.2 cm in Z, were stacked in an array behind the converter blocks with the small cross-section normal to the beamline and a PMT attached to the downstream end of each block.

Electrons entering the converters were expected to initiate a large electromagnetic shower. Pions might shower hadronically or electromagnetically, via charge exchange, in the converter or the absorber assembly. Electron identification was then done by calculating the ratio of the total energy deposited to the momentum of the incoming particle (E/p). For electrons, this ratio should be approximately unity, while for pions, the ratio was, on average, substantially lower. Muons in the E871 momentum range are minimum ionizing and thus deposit a very low fraction of their energy in the glass. The E/p ratio range for pions overlapped that for muons and therefore, the PbG gave poor separation between pions and muons. These PbG characteristics for the three particle types are clearly demonstrated in Figure 3.21.

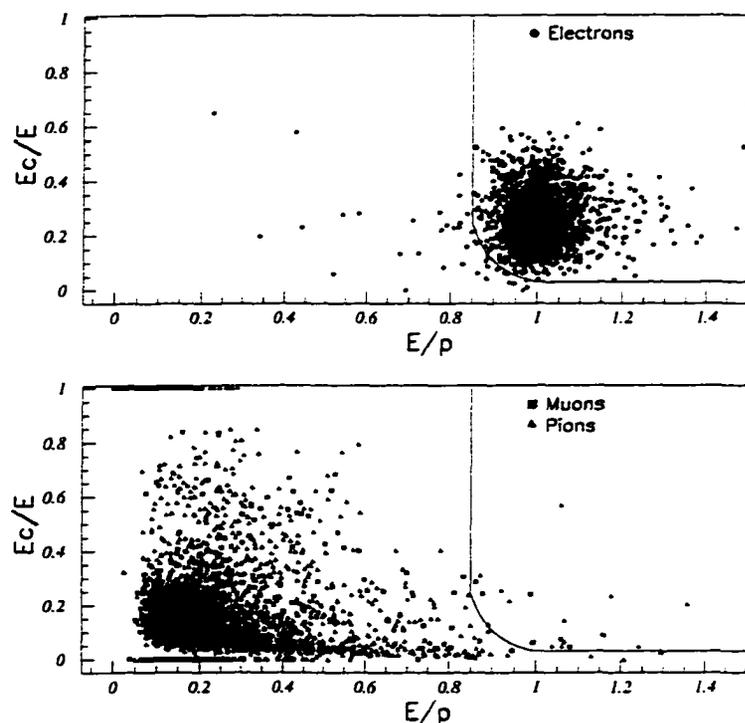


Figure 3.21: PbG performance for well identified e , μ , and π 's in data. e identification contour cut is shown.

3.6 Muon ID

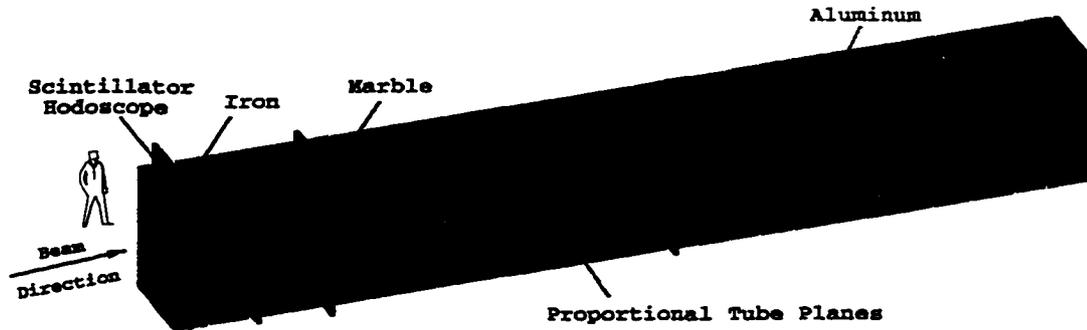


Figure 3.22: IMF Detector layout.

Good muon identification was essential for a $K_L^0 \rightarrow \mu e$ search to avoid single misidentification background ($K_L^0 \rightarrow \pi^\pm e^\mp \nu$ where the π was misidentified as a μ). While background from misidentification can be avoided with excellent invariant mass resolution, it can further be suppressed by proper muon identification. Muon identification was provided by a muon hodoscope (MHO) and a muon rangefinder (MRG), located downstream of the lead glass enclosure. For historical reasons, each system is described separately. However, they were actually integrated into a single instrumented muon filter (IMF) system. The IMF, shown in Figure 3.22, consisted of a horizontal stack of iron, Carrera marble and aluminum plates instrumented to result in a 5% momentum difference between successive active components, either MHO or MRG planes.

The IMF exploited the differences in the energy loss mechanisms among muons and pions or electrons to positively identify muons while rejecting all electrons and most pions. Muons in the momentum range of interest were minimum ionizing and dissipated a known amount of energy when passing through a given material. Therefore, if the range of the track in the IMF was as expected, given the spectrometer-determined momentum and expected variation in track range due to straggling, the track was identified as a muon. Electrons

would be lost to showers in the PbG or, if they penetrated through a crack between PbG blocks, would be stopped by the wall of 30.5 cm thick Fe in front of the first MHO plane. Pions which did not decay would also probably shower in either the PbG or the Fe. However, a small number of undecayed pions, or, more likely, their shower products, would “punch through” the Fe and generate a signal in the first few planes of active material. As with the electrons, pion “punch through” was a problem only in the low momentum regime. Pions which decayed into muons, however, presented the highest possibility for generating a false muon signal in the IMF. Depending on the kinematics of the decay, the decay muon can have from 60% to almost 100% of the original pion momentum in the lab frame. Decay muons that had more than 80% of the pion momentum appeared to be a muon from the original decay, so a K_{e3} decay in which the pion decays was a major source of possible background for $K_L^0 \rightarrow \mu e$.

3.6.1 Muon Hodoscope

MHO Panel Dims.					
Plane Type	X(cm)	Y(cm)	Z(cm)	Z Position (m)	Momentum (GeV/c)
X0	228.6	12.7	1.27	35.265	1.0
Y0	104.14	25.4	1.27	35.355	1.0
X1	266.7	18.73	2.54	34.945	0.75
Y1	228.6	18.73	2.54	36.345	1.5
X2	266.7	18.73	2.54	38.195	2.8
Y2	228.6	18.73	2.54	45.915	5.8

Table 3.2: MHO Panel Dimensions.

The MHO planes, shown in Figure 3.23, consisted of scintillator panels arranged vertically and horizontally to measure the X and Y coordinates, respectively, of a muon passing through the iron and marble stack. Attached to one or both ends of these panels were Phillips XP2262 PMTs to measure the light generated by the scintillator. The signal

from the PMT base was sent along coaxial cable to LeCroy 4413 discriminators in the Level 1 room. The discriminator output was then sent to the 8 bit TDCs. The signals from plane X1 were split before the discriminators and the other part was sent to ADCs to measure the energy deposition. The pulse height measurement was used to help provide better discrimination between muons and pions, because pions might generate particle showers. The second and third MHO planes, numbered X0,Y0, provided a signal to the Level 1 trigger for use with the various muon triggers. All of the MHO plane dimensions are shown in Table 3.2 and the efficiencies of each panel are shown in Figure 3.24

3.6.2 Muon Rangefinder

The MRG planes were two-wire aluminum proportional counters arranged eight cells to an extrusion, as shown in Figure 3.25. These extrusions were bonded edge to edge to cover a cross-sectional area of 2.25 m x 3.05 m normal to the z axis. A finished plane is shown in Figure 3.27.

HV was supplied to the MRG by a Caen SY127 mainframe which provided +2600 V to the HV bus located on the amplifier cards mounted to the MRG plane. The LV for the amplifier and discriminator cards, also mounted on the MRG plane, came from racks of 5 V and 12 V power supplies located on the beam left side of the MRG . Signals from the wires were capacitively coupled to the amplifier cards, amplified in a three stage amplifier and then passed to the discriminator card, shown in Figure 3.26. The discriminator card ORed eight of these channels, i.e. one extrusion, and converted the signal to a differential ECL pulse approximately 160-200 ns wide. This signal was sent via 32 channel Ansley cable to the latch cards in the data acquisition system.

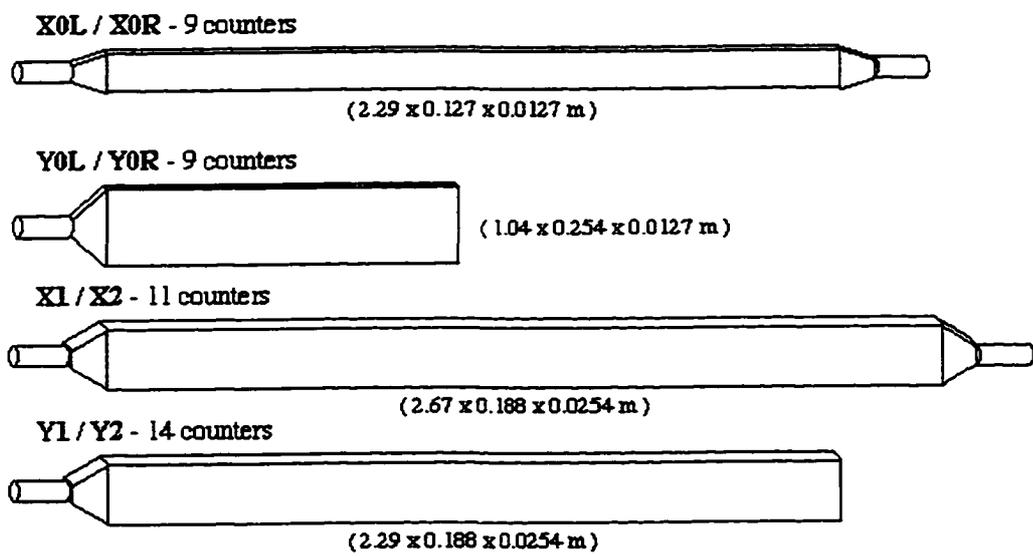
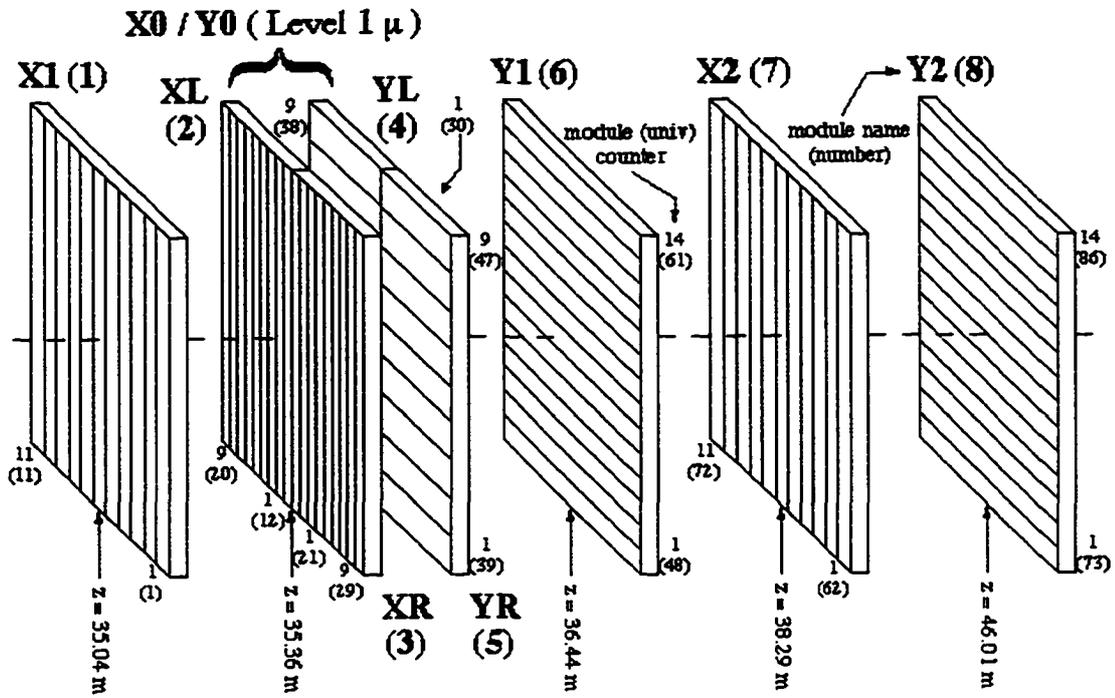


Figure 3.23: MHO Panel Dimensions and Layout. The numbers (1 - 8) represent the sequential plane number.

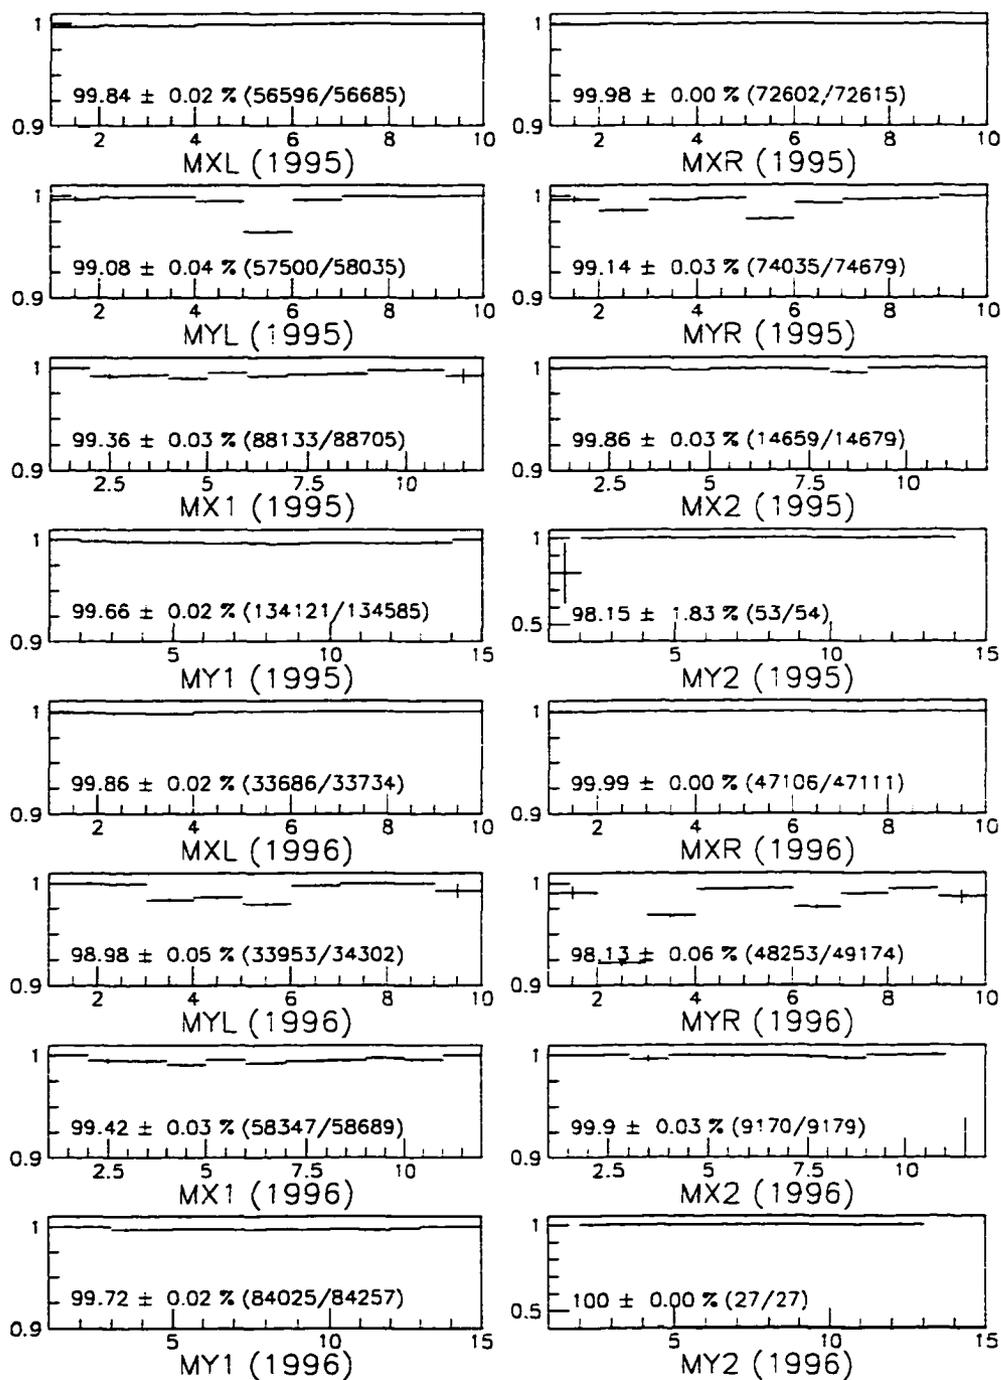


Figure 3.24: MHO Efficiencies by panel.

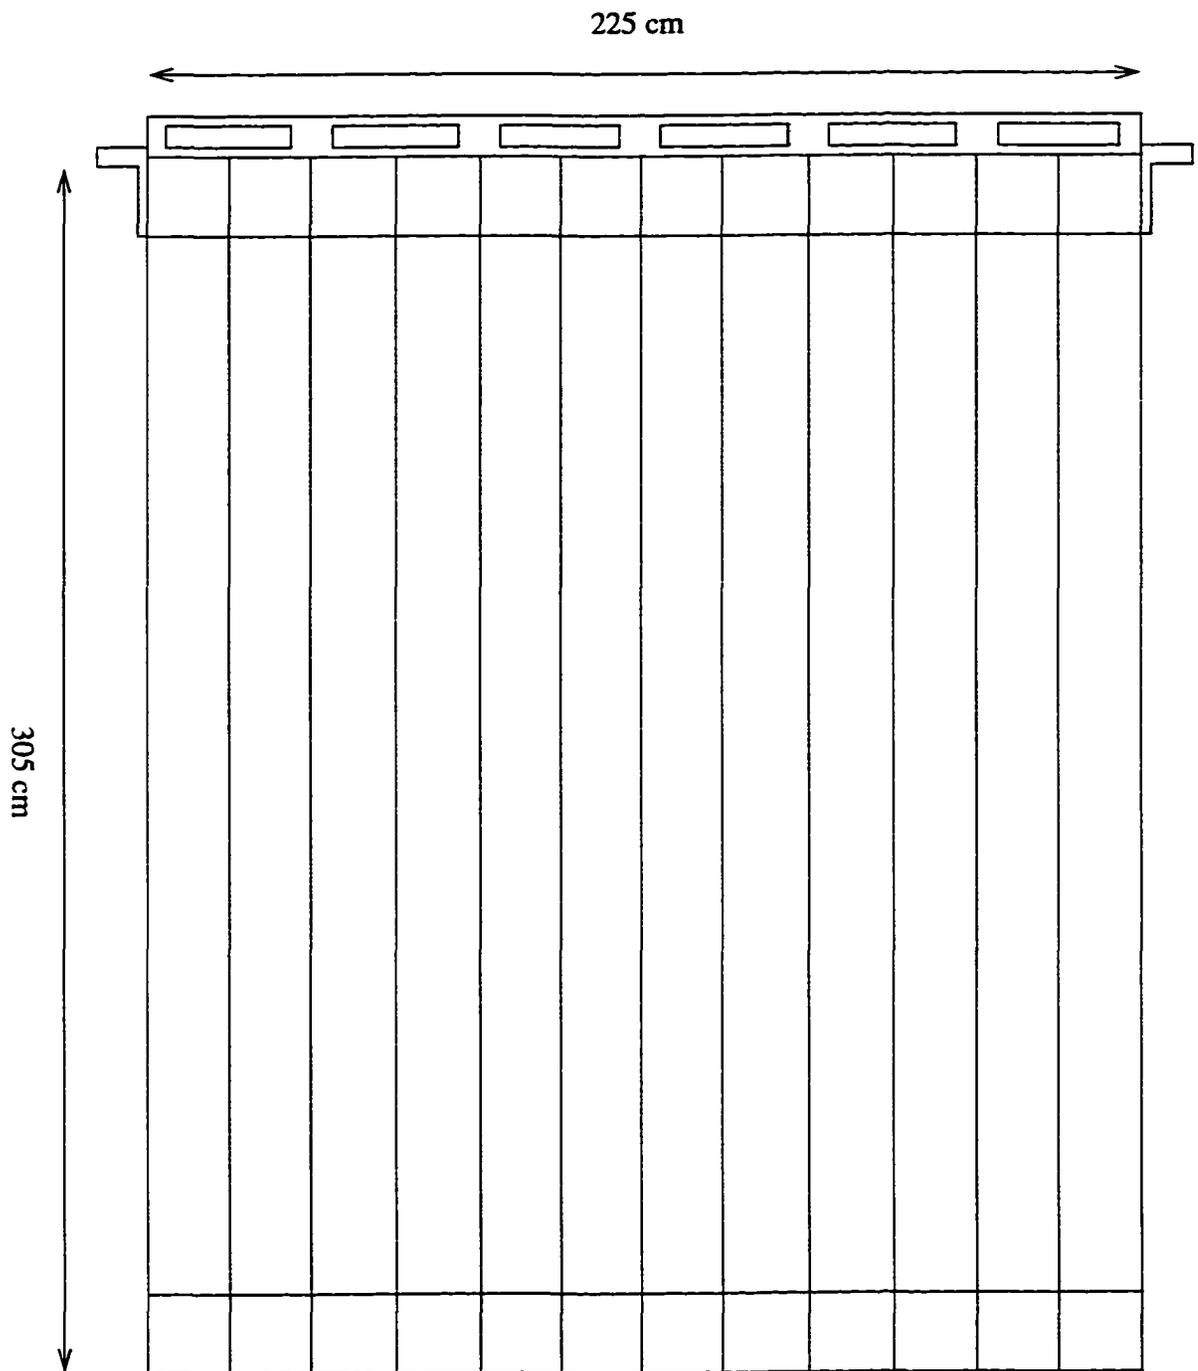


Figure 3.27: MRG X-measuring plane (Beam view).

3.7 Data Acquisition System

In order that E871 reach a single event sensitivity on the order of 10^{-12} , a fast and intelligent trigger and data acquisition system was required. Such a system must be able to deal with high rates and have a reasonably high rejection capability for the primary K_L decays, K_{e3} and $K_{\mu3}$.

In overview, data acquisition started with a Level 0 (L0) trigger. If a L0 trigger passed other requirements, the event fired a Level 1 (L1) trigger. The L1 trigger system provided some signals to the front end electronics and signalled the Readout Supervisor (RS) to start processing the event. The RS then directed the Crate Scanners (CS) either to shift the event in the front end modules internal registers or to transfer the event data to the Dual Port Memories (DPM) for eventual processing by the Level 3 (L3) trigger system. Then, the RS notified the L0/1 system that the data acquisition system was ready to accept another event. When the DPMs were full, the CS notified the RS, which selected a L3 processor to filter the events in the DPM and upload the events which passed the L3 code to the main data acquisition computer for their eventual writing out to tape. This system was able to process an input event rate of about 10^6 Hz while keeping the data-to-tape rate down to approximately 10^2 Hz.

3.7.1 Level 0/1 Trigger

The Level 0/1 Trigger system was the low level hardware trigger, shown in block diagram form in Figure 3.29. Signals from the TSC, CER, and MHO made up the inputs to the L0/1 trigger system. During normal running, the L0/1 trigger system filtered a 1 MHz data rate down to about 10 KHz.

A non-parallel L0 trigger was defined as coincident hits in all six TSC sections (1X/2X/2Y,L/R). This trigger was less important than the parallel version but provided a

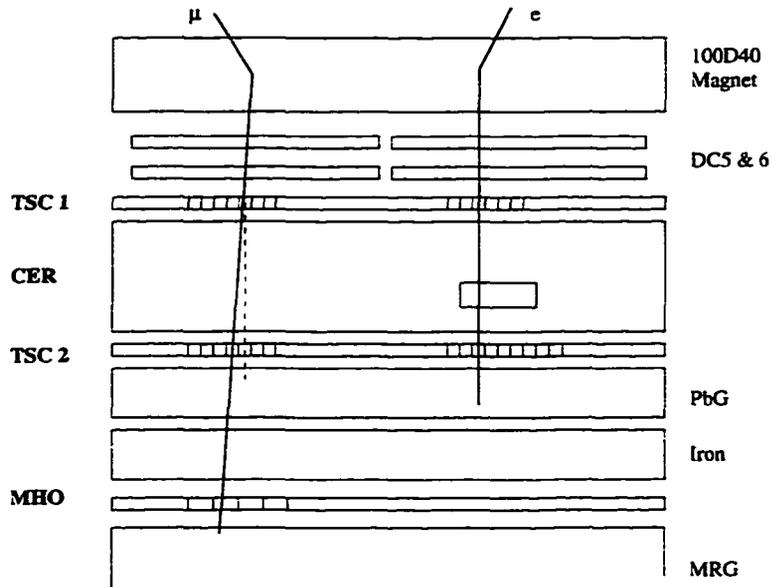


Figure 3.28: Experimental Top View of a Level 1 Trigger Event.

useful benchmark for trigger filtering. The rate of non-parallel L0 triggers during normal running conditions was on the order of a million per spill. The parallel L0 required that the four hits in the X views pass the parallelism condition, usually set at ± 2 TSC slats, corresponding to ± 31 mrad. This was done to take advantage of the parallelism of two track events set up by the spectrometer magnets (see Section 3.3.3). This parallelism reduced the data rate by about a factor of four to approximately 250 kHz.

A L1 trigger, as shown in Figure 3.28, fired on events with either a parallel L0 trigger that had a PID signal of interest (e, μ) or a calibration or $\pi\pi$ normalization event. For the electron trigger, a Čerenkov hit with an ADC charge above the one photoelectron threshold and a space and time match with the TSC hits that formed the parallel L0 trigger was required. For the muon trigger, hits in the MHO X0 and Y0 planes with a space match to the track projections were required. The five physics L1 triggers were $e\mu$, μe , $\mu\mu$, ee and minimum bias (MB). MB triggers were parallel L0 triggers, prescaled by a factor of 1000,

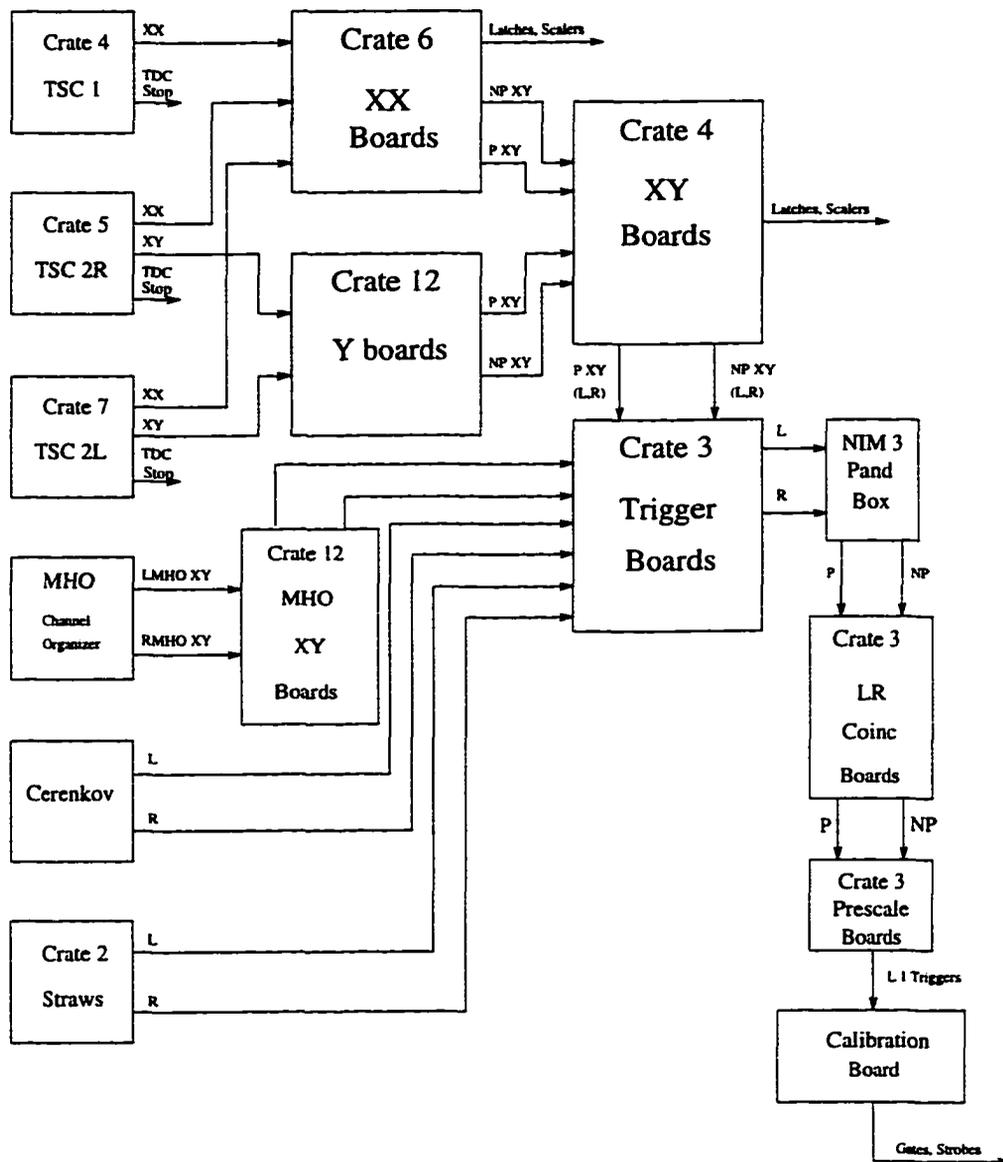


Figure 3.29: Level 1 Flowchart.

and used mainly for normalization. There were several calibration triggers that occurred in each spill, including the 50 Hz trigger, the PbG laser and reference tube triggers, the ADC pedestal trigger and some highly prescaled parallel and non-parallel L0 triggers. The output from L1 during normal running was about 10 kHz.

3.7.2 Data Acquisition Electronics

Much of the design for the data acquisition electronics system, shown in Figure 3.30, was inherited from E791 and most of the electronics were custom made. The main differences between E791 and E871 electronics were the introduction of a lower least count TDC module to read out the straw chambers and the use of faster processors for the L3 trigger. This section is meant as an overview of the hardware used. For more details, see reference [47].

FASTERBUS

FASTERBUS, a custom design for data acquisition was an attempt to overcome the limitations of the data acquisition hardware available for BNL E791. It was a parallel pipelined readout scheme which was capable of transferring up to 250 Megabytes/second into the dual port memory modules to be processed by the L3 processors. The FASTERBUS architecture had four features that contributed to its high performance. First, the data digitization was done in less than 200 ns, the so-called “flash digitization”. Second, the front end modules each had a set of latches to which the data were shifted shortly after digitization (approximately 10 ns) to allow the next event to be processed. Third, a highly parallel pipelined dataway from the modules, through the crate scanner, to the dual port memory (DPM) was used to pass the digitized data along quickly to be processed. Each crate scanner was connected to each DPM via a 17-pair cable, allowing the data to be

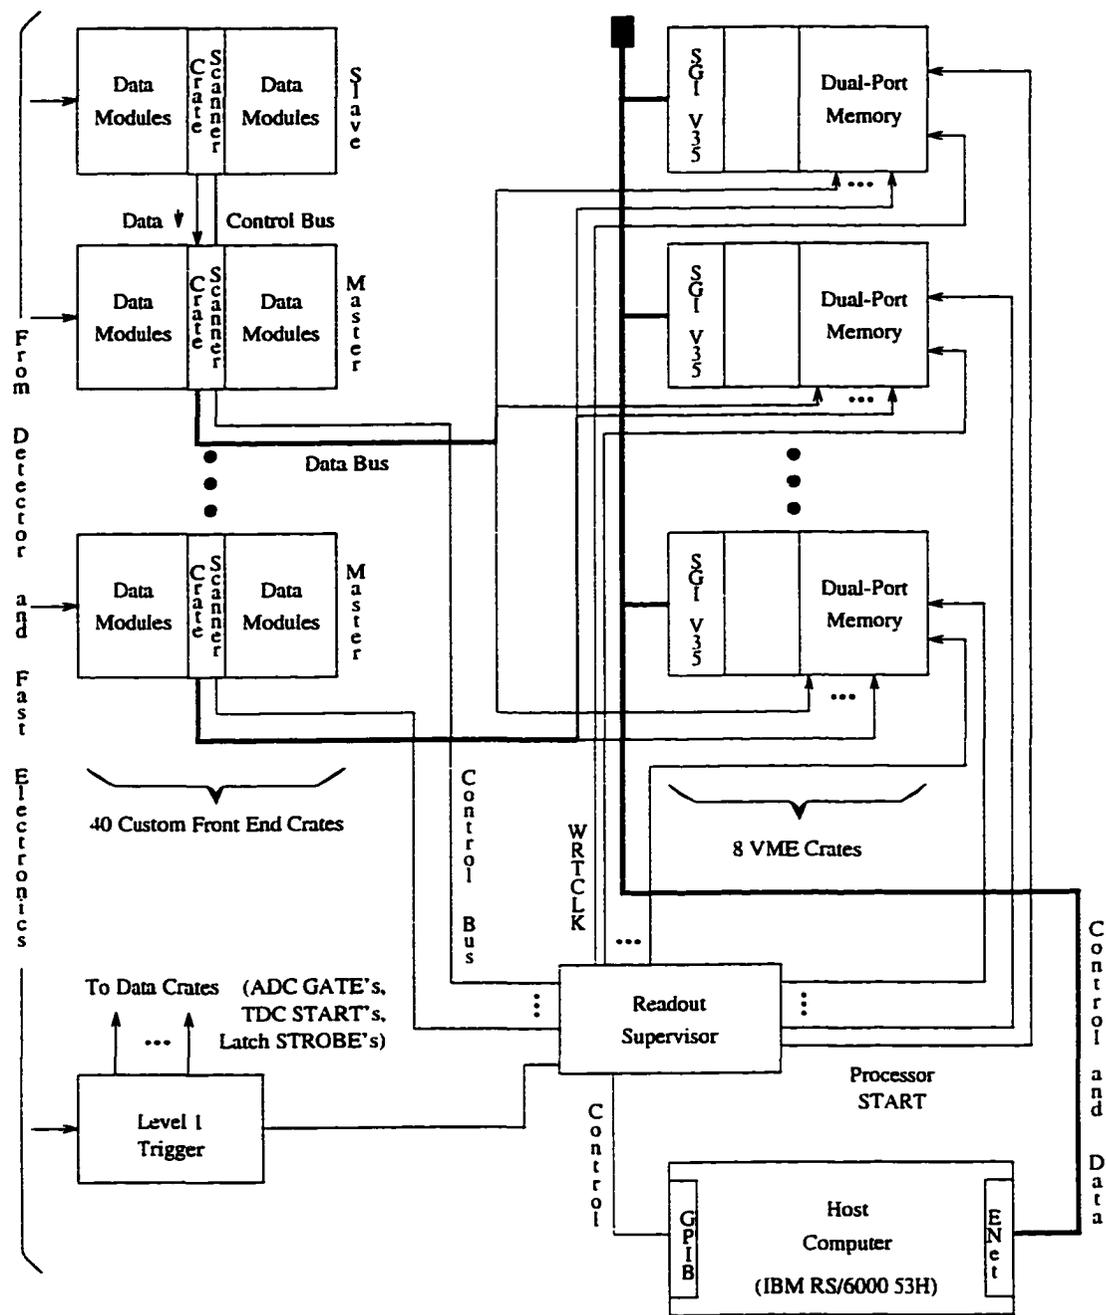


Figure 3.30: Overview of the Data Acquisition System.

quickly transferred directly into whichever processor was available at the time. Fourth, the DPM modules, which were located with the eight SGI V35 processors in eight VME crates, allowed for event buffering to minimize deadtime during a spill.

Readout Supervisor

The Readout Supervisor (RS), as the name implies, controlled the flow of data through the readout system and up to the data acquisition computer. The L1 trigger sent a signal to the RS notifying it that an event was ready for processing. The RS then asserted an RSBUSY (Readout Supervisor Busy) signal to L1 so that L1 would not allow any more triggers until the RSBUSY was clear again. At the same time, the RS sent a SHIFT signal to the front end modules, thus causing them to move the event data from the Stage 1 Flip-Flops (FF) to the Stage 2 FFs. This cleared the front end modules to accept another event. Once the data were SHIFTeD, the RSBUSY was cleared so the experiment could continue to take data. When the RS received a CRINFO (Crate Info) signal, the RS sent a RDCLK (Readclock) signal to the crate scanners to enable the scanners to “clock out” the data in the crate into the DPM. After a small but carefully timed delay from the start of the RDCLK signal, the WRTCLK (Writeclock) signal was sent to the DPM to allow it to “clock in” the data being sent from the crate. This cycle, shown in Figure 3.31 continued until the DPM buffers were full and the CS sent a FULLMEM signal to the RS. The RS kept a record of the available L3 processors, *i.e.* those processors in the system not currently processing an event. Once the DPM buffer was full, the RS sent an Emulator Start (STARTE) signal to the lowest numbered available L3 processor, notifying it to start processing the event.

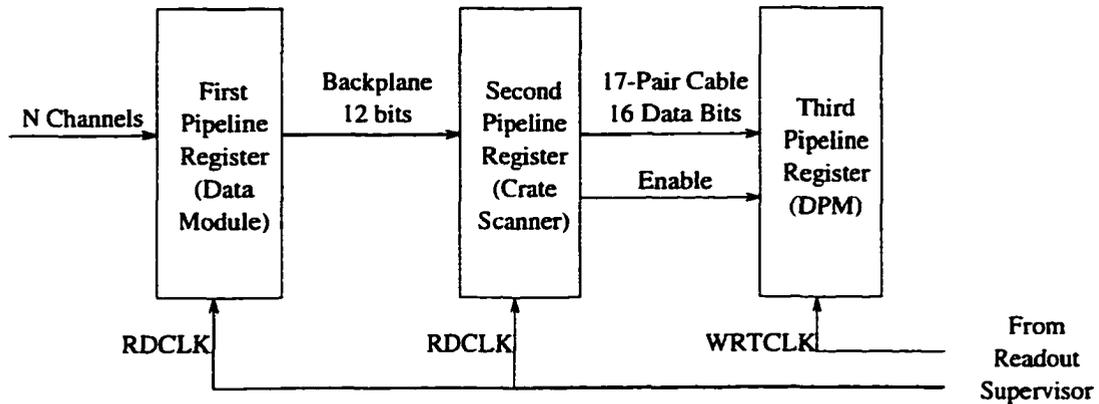


Figure 3.31: Outline of the Data Pipeline.

Crate Scanners

The crate scanners controlled the data flow from the crate to the DPM and were, in turn, controlled by signals sent by the Readout Supervisor (RS). Each crate had a scanner, which was located in the center slot of the crate to minimize the signal travel time over the whole crate. The scanners were capable of operating in a master/slave mode, whereby one crate was “slaved” to another in order to reduce the number of required connections to the DPM. The scanners also generated a clock for use by the six bit TDCs.

The scanner’s function started when the HIT line on the FASTERBUS backplane was asserted by one of the front end modules (TDC, ADC or latch), meaning that data were present in the Stage 2 FFs. The scanner would then assert the CRINFO (Crate Information) line to the RS and keep asserting the line until the crate was completely read out. The RS then sent a stream of pulses, called a RDCLK (Readclock), to the scanner. At each RDCLK, the 12 bits of data, plus four bits for crate and region word, were clocked into a pipeline register. This register was then fanned out to drivers which sent the 16 bit data word to the DPM.

TDC

E871 used three types of TDC module: two types of a six bit module and an eight bit module. The first version of the six bit module had 32 channels with a 1.25 ns least count and a dynamic range of 80 ns. These were used with the straw drift chambers, which used a faster drift gas and a smaller cell size, thus requiring a smaller least count. The second version of the six bit module also had 32 channels but had a 2.5 ns least count and a dynamic range of 160 ns. These were used with the regular drift chambers, which, due to the slower gas mixture and larger cell size, required a larger dynamic range. Both of these versions were all digital and were designed around a six bit gray code clock which was driven by a 100 MHz signal provided by the crate scanner.

The third TDC module used was a 16 channel, analog 8 bit module with 200 ps least count and a dynamic range of 50 ns. These Fast TDCs (FTDC) were used on the PMT detectors (TSC, CER and MHO) which required precision timing. The inputs arrived via the discriminators in the L1 system. At its core, each channel had a Sony CX20052A 20 MHz flash analog-to-digital converter chip. This chip measured the voltage across a capacitor which was charged by a constant current source between the “start” signal, provided by the L1 trigger system, and “stop” signal, provided by the input signal from the detector.

All three versions of TDCs used a sparse data scan so that only the hit channels were read out.

ADC

The ADCs were 12 channel, 8 bit modules which were also designed around a Sony CX20052A ADC chip. They were used to read out the PbG and the 80% split portion of the CER signal. The input circuit was bilinear, i.e. for small signals, each count was 150 fC of charge while, for large signals, each count was 470 fC. With this design, the ADC

exhibited a wide dynamical range while preserving high resolution for small input signals. These modules also used a sparse data scan to produce output only for channels that were actually hit. The 8 bit wide output was generated within 200 ns of the close of the common gate and constituted the Stage 1 Latch.

Latches

The Latches were 96 channel ECL modules that latched in the state of 3 sets of 32 channel inputs on the leading edge of the latch strobe signal. These modules were used to read out the MRG and to latch in the state of each L1 trigger. For the MRG readout, each set of 32 channels read in information from two MRG planes so only nine modules were required for data taking. The input signals represented an OR of the eight cells on an MRG extrusion and either 12 or 16 of these signals comprised the hit information from that plane. Unlike the other modules, there was no sparse data scan on the latches because it was unnecessary. Because each channel required only one bit to store its information, all 96 input channels on a latch were represented by 9 data words read out by the crate scanner.

3.7.3 Level 3 Trigger

The Level 3 Trigger was the software trigger running on the L3 processors. It used spectrometer information from events that passed L0/1 to perform a rough pattern recognition. The processing started when the lowest numbered available L3 processor received a STARTE (START Emulator) signal from the RS. After checking the data integrity in the DPM, the L3 filter was applied to the data. Events flagged as minimum bias events were passed at this point. For the physics trigger events, except *ee* events after mid-run 1995, the L3 filter performed a simple pattern recognition on each event. (For *ee* events after mid-1995, the L3 filter was used and trigger bits were set but the events that failed were

not cut from the data stream.) Starting with clusters in SDC1 and searching for clusters downstream, the algorithm stopped searching after finding one track on each side. These tracks were then extended back into the vacuum decay region and the vertex distance of closest approach, or vertex DOCA, and the vertex p_T were calculated. The invariant mass was also calculated, based on the L1 trigger bits set in the trigger latch module. The following conditions on these parameters were required in order that the event pass the L3 filter:

- Vertex DOCA < 0.5 m
- Vertex p_T < 60 MeV/c
- Invariant mass > 0.465 GeV/ c^2

The $\pi\pi$ mass was calculated for minimum bias events and those passing the mass cut were flagged as $\pi\pi$ candidates. As the buffer of events in the DPM was being processed, the events that passed were extracted and placed in a shared memory segment. Once a full buffer of events had been processed, the shared memory segment was attached to a “result” structure. This structure was then uploaded via a dedicated ethernet connection to the data acquisition computer to be written to tape. The total output rate from the L3 filter was a few hundred events per spill.

3.7.4 Data Acquisition Computer

The data acquisition (DAQ) computer was responsible for uploading the data from the L3 processors, for notifying the RS that a previously busy processor was now available and for writing the uploaded data to tape. In addition to uploading the raw event data, the DAQ computer also uploaded scaler information from the L1 system, controlled and read

out HV settings from over 1100 channels of HV, and monitored several hundred detector operating parameters via scanning ADCs.

The DAQ computer was an IBM RS/6000 53H connected to a dedicated Ethernet line from the L3 processors. The data were uploaded using custom made software designed around the Sun Remote Procedure Call (RPC) library. For more details on the upload software, see Reference [48]. As the data structures were uploaded, they were buffered to disk and the RS was notified, via GPIB calls, that the processor was now available to process more data. The disk buffering was necessary because the tape drives used helical scan tape heads. These tape heads required relatively long repositioning times between tape writes. Therefore, the tape drives would be unable to maintain the required tape writing speed if they wrote out each event individually. However, if the events were written out to tape as a large buffer, the repositioning time was irrelevant as it only factored in once at the beginning of the tape writing. Once the disk event buffer reached 200 MBytes, that data buffer was closed and another buffer was created. When a full tape worth of data, usually seven 200 MByte runs, was on disk, the tape controller started to write the 1.4 GBytes of data to a 4 mm Digital Audio Tape (DAT) for archiving. If that block of data failed to write at any point, the writing of that tape was halted and the remaining data from that tape were appended to the next tapelist to be written out on the next tape. This system allowed the writing of almost 2 Terabytes of data over the span of eight months of data collection.

The DAQ computer also performed scaler readout and detector monitoring. The scalers from the L0/1 system were read out and added to the data stream at the end of every spill. Over 1100 HV channels settings were monitored by the DAQ computer and, in case of a HV trip, the DAQ computer attempted to reset that channel. The detector monitoring was done by reading out a bank of scanning ADCs which were connected to

hundreds of sensors attached to the experimental apparatus. These ADCs read in voltage levels that corresponded to such quantities as temperatures, pressures, flow rates, and LV levels, in the detector systems. Due to the high noise environment in the AGS building which interfered with the proper monitoring of these voltages, a running average of each voltage measurement was calculated. If any of these averages went out of a predefined range, the monitoring software alerted the shift person to take action to correct the problem. Both the HV and ADC measurements were inserted into the data stream at the beginning and at the end of a run.

Chapter 4

Offline Software

4.1 Overview

Much of the offline software was adapted from that used in E791. Those portions which are unchanged will not be discussed here. More information on these topics is found in [49–51]. The areas that have changed, and will be discussed, are the multiple scattering calculation in the offline Monte Carlo, the pattern recognition software, the two track fitting algorithms and the particle identification software for both electron and muon identification.

4.2 Multiple Scattering Calculation

The revised multiple scattering calculation in the Monte Carlo is covered in detail in Reference [43] but the main conclusion will be given here. The Monte Carlo calculates scattering due to material *before* an aperture, for extended media, and *at* an aperture, for thin membranes. For material before an aperture, both the track slopes and the track (X,Y) positions are smeared, while for material at the aperture, only the slopes are smeared. Table 4.1 shows the radiation lengths before and at each aperture.

Aperture Number	Aperture Name	Before Aperture X/X_0	At Aperture X/X_0
1	Target	0.0	0.0
2	Vacuum Window	0.0	0.00117
3	Straw Chamber 1	0.0000322	0.00226
4	Straw Chamber 2	0.000195	0.00226
5	96D40 Magnet	0.000318	0.0
6	Straw Chamber 3	0.000246	0.00146
7	Straw Chamber 4	0.0000869	0.00232
8	100D40 Magnet	0.000230	0.0
9	Drift Chamber 5	0.000226	0.00180
10	Drift Chamber 6	0.000219	0.00169

Table 4.1: E871 Spectrometer Radiation Lengths.

4.3 Pattern Recognition

The pattern recognition algorithm (*patrec*) has the task of sorting through all of the combinations of wire hits in the spectrometer, using them to reconstruct track candidates, determining the kinematic quantities of the candidates and making cuts where appropriate to reduce the number of track candidates. *patrec* then uses the remaining track candidates to form event vertices, sorts these vertices from best to worst and determines an event time, also known as a T_0 .

The first task *patrec* does is look at the number and location of the chamber hits. Events with too many hits overall and events with no hits in one or more planes are rejected. The next step is to determine an approximate event time, based on the time sums from good X-view triplets and Y-view doublets. This is done to reduce trigger jitter, which can be up to a few nanoseconds. Next, the algorithm tries to find “clusters”, or contiguous sets of hit wires in different layers, in each of the chamber planes. An ideal cluster would be a triplet in the X-view or a doublet in the Y-view. However, due to dead or inefficient wires, the

clusters are allowed to have a missing hit, possibly resulting in doublet and singlet clusters in the X and Y view, respectively. Clusters larger than two or three wires, perhaps caused by extra tracks, crosstalk or δ -rays, are also accepted. Clusters larger than 20 wires in X and 10 wires in Y, however, are broken up into two clusters with a two-wire overlap. All of these clusters are then used to develop a set of two-dimensional (2D) track hypotheses. Because of the lower noise rates in the downstream section of the spectrometer, the search algorithm starts with clusters in DC6 and works forward to SDC1, assembling a list of cluster sets along a predetermined road width that might be associated with these two dimensional prototracks. The road width varied for each plane from -180 to +200 mm in X and ± 130 mm in Y for DC5 to ± 55 mm in X and ± 60 mm in Y for SDC1.

Because each cluster can include a large number of wires, it may be difficult to use raw clusters to determine exactly where the track passed through the chamber. Therefore, each cluster in a chamber must be sorted through to find a smaller set of wires which represents the most likely path the track traced through that chamber. This smaller set of wires is referred to as a "segment". When searching for segments, patrec assumes that the track angle determined by the 2D prototracks will not be much different from that determined by the segment. With the track angle fixed, a line with that angle is swept across the cluster and a "score" is calculated at each step. To speed up this process, possible score minima are determined beforehand and the line is swept across regions around the possible minima. This score is different from a simple χ^2 calculation and is calculated to overcome problems with noisy or inefficient wires. In a regular χ^2 calculation for a segment, the distance of closest approach (DOCA) between each wire and the prototrack is calculated given the track angle. Combined with the calculated drift distance, δ_i , of each wire, a "normal" segment χ^2 can be determined from

$$\chi^2 = \sum_{i=1}^{N_w} \frac{(\delta_i - doca)^2}{\sigma_i^2}, \quad (4.1)$$

where N_w is the number of wires in the segment and σ_i is the resolution of the i th wire. Simply minimizing this χ^2 is, in fact, sufficient to find the actual track position for the majority of tracks. However, for some tracks, single wire inefficiencies and noise cause the algorithm to choose the wrong track position because that wrong position had the lowest segment χ^2 . If the algorithm chooses the wrong side of a tracking cell, the track position can be incorrect by as much as 5 mm in the SDC and 10 mm in the DC. Position mismeasurements as large as these can cause the track to fail the quality requirements. Therefore, the concept of a track “score”, which is similar to a negative log likelihood, was developed to take these effects into account. In addition to using the segment χ^2 , the score calculation uses the probability that the wire hit (or miss) is actually

1. a genuine hit,
2. an early noise hit,
3. a late noise hit,
4. or, in the case of a wire that was *not* hit, an inefficient wire that *should* have been hit.

Once the score is calculated, the segment with the lowest score is considered the actual track position. If the difference between the best segment score and another segment score is less than four (which is more than e^{-4} less probable), the other segment is kept as an alternate hypothesis. patrec will keep up to three of these alternates.

The collected segments are then used to form 2D tracks in the X and Y views separately. X view tracks that do not meet the 2D front and back momentum match

requirements (up to 10% difference in momenta) are cut. Y view tracks are required to meet in the center of the magnets to within 50 mm. Once all of the likely 2D tracks are assembled, the X and Y view tracks are paired up to form three-dimensional (3D) tracks. The 3D tracks are required to pass cuts on the segment scores and both the bend angle and Z position of the bend angle in each magnet. This process generates a bank of all of the viable 3D tracks found in the spectrometer.

The next step is to use these 3D tracks to find potential K_L decay vertices. The tracks are projected back into the vacuum decay tank and a “vertex doca” is calculated for that track pair. Any track pair with a vertex doca less than 100 mm and a vertex Z position between 9 and 21 m is considered to form an acceptable vertex. It is possible for tracks to be used in multiple vertices. The two tracks with the best vertex doca are considered the most likely vertex pair and the other vertices are sorted by doca in ascending order. Finally, the track times from the best vertex pair are used to determine the event time, or T_0 .

4.4 Track Fitting

For historical reasons, E871 has two track fitters, referred to as FT and QT . The fitters use the hit information calculated in patrec to determine the most likely trajectory of the particle in the spectrometer and thus determine the track momenta. FT does a true mathematical fit while QT does an iterative fit. It is advantageous to have two fitting algorithms available so each fitter can operate as a check on the other.

4.4.1 FT Fitter

The FT fitter works by fitting the entire track at once and minimizing a calculated vertex χ^2 to find the most likely particle trajectory. This section is an overview of the

principles and algorithms used in FT. More detailed information on the FT fitter is in References [52, 53].

The FT fitting algorithm begins by filtering out tracks that have a large number of alternate hypotheses. This filtering is accomplished by looping over patrec tracks and cutting those tracks whose combined set of track alternates allows more than 32 combinations per track. After filtering, the algorithm starts fitting tracks to this reduced segment set. The track can be described by the 5-vector,

$$\vec{\alpha} \equiv \left(x, y, x', y', \frac{q}{p} \right), \quad (4.2)$$

where (x, y, x', y') are the track positions and X-view slopes at the chambers and $\frac{q}{p}$ is the ratio of the charge to the momentum, which is used because it scales linearly with magnetic bend angle. Next, the elements of a covariance matrix, V^x , are calculated. The elements of this matrix are given by:

$$V_{ij}^x = \left\langle (x_i^{data} - x_i^{theo})(x_j^{data} - x_j^{theo}) \right\rangle, \quad (4.3)$$

where x^{data} is the actual measurement and x^{theo} is the expected measurement at each plane, ignoring multiple scattering or measurement errors. The diagonal terms in this 28×28 matrix are simply the variances of the variables ($V_{ii}^x \equiv \sigma_i^2$) and the off diagonal terms are zero if the errors on each measurement are independent. All of these elements have already been calculated as a function of momentum by using the Monte Carlo. The inverse of V^x is the weight matrix, W^x , which is used to calculate a track χ^2 :

$$\chi^2(\vec{\alpha}) = \sum_i \sum_j (x_i^{data} - x_i^{theo}) W^x (x_j^{data} - x_j^{theo}). \quad (4.4)$$

These sums over i, j are each over the 28 measurements (18 X and 10 Y measurements) in a track. This χ^2 is then minimized by assuming the χ^2 well is parabolic, calculating the first and second derivatives and stepping to the minimum. In this procedure, the track is

constrained to go through the hits in SDC1 and DC6 but the other positions are free to float.

A byproduct of this minimization is a 5×5 curvature matrix which is defined as:

$$W_{ij}^\alpha \equiv \frac{1}{2} \frac{\delta^2 \chi^2(\vec{\alpha})}{\delta \alpha_i \delta \alpha_j}. \quad (4.5)$$

The inverse of this matrix is the covariance matrix V^α for the track parameters $\vec{\alpha}$. The diagonal terms of V^α are the estimated errors of $\vec{\alpha}$.

Once the tracks have been fit, FT tries to fit the event vertex by creating a nine-parameter set, $\vec{\beta}$, from $\vec{\alpha}_1$ and $\vec{\alpha}_2$ such that

$$\vec{\beta} \equiv (x_v, y_v, z_v, x'_1, y'_1, q_1/p_1, x'_2, y'_2, q_2/p_2), \quad (4.6)$$

where the subscripts $v, 1, 2$ are for vertex, track 1 and track 2 quantities from $\vec{\alpha}_{1,2}$. In a manner analogous to Eq. 4.4, χ_{vertex}^2 is calculated from

$$\chi_{vertex}^2(\vec{\beta}) = \sum_{i=1}^{10} \sum_{j=1}^{10} (\alpha_i^{data} - \alpha_i^{theo}) W^\alpha (\alpha_j^{data} - \alpha_j^{theo}). \quad (4.7)$$

The weight matrix, W^α , is a 10×10 matrix generated using the 5×5 track covariance matrices, V^α , with several terms added to account for scattering in the vacuum window, an effect not accounted for in the V^α calculations. Once this matrix is inverted, χ_{vertex}^2 is automatically minimized with no further iteration necessary. The error on the invariant mass is calculated using a propagation of the errors on $\vec{\beta}$:

$$\sigma_M^2 = \sum_{i=1}^{10} \sum_{j=1}^{10} \frac{\delta M}{\delta \beta_i} V_{ij}^\beta \frac{\delta M}{\delta \beta_j}. \quad (4.8)$$

4.4.2 QT Fitter

Usually, a track fitting algorithm attempts to minimize some sort of χ^2 which reflects how well the fit track corresponds to the actual track, *e.g.* the FT fitter. The QT fitter, however, estimates track parameters iteratively as it follows the track through the spectrometer. This iterative approach is based on the observation [54] that

1. the vertex information is largely determined by track information in the first two SDCs and
2. multiple scattering dominates the errors in the momentum determination.

Therefore, attempting a fit of the entire track at once is a more time-consuming process and does not necessarily enhance the measurement. Thus, QT operates on the front and back sections of the spectrometer separately, working from SDC1 to SDC3 in the front and from DC6 to SDC4 in the back. The algorithm adjusts the track parameters with each iteration until the fit positions in each chamber are within $10 \mu\text{m}$ of the average hit position. The tracks fit by QT are still constrained to go through the hits in each chamber except the Y view in chamber 4.

Once the front and back tracks are determined for all of the combinations passed along by patrec, all combinations of two-track sections are paired up one pair at a time and tested. The deviations in

1. front and back momentum match $(\frac{p_f - p_b}{p_{avg}})$,
2. X and Y positions at the third chamber $(\delta x_3, \delta y_3)$,
3. Y position in the fourth chamber (δy_4) , and
4. the X and Y angles at both the third and fourth chamber $(\delta\theta_{x3}, \delta\theta_{y3})$,

are measured for each pair of front/back tracks. The squares of these deviations are used to calculate the QT χ^2 . These front/back track pairs now form a complete set of full tracks, which are then sorted in order of descending reduced χ^2 .

When all of the tracks from each side are fit, QT starts matching left and right tracks together to form a vertex in the vacuum tank. The doca of the two tracks is taken to be the vertex and the vertex χ^2 is calculated from

$$\chi_{vertex}^2 = \frac{\delta_{vertex}^2}{(z_{SDC1} - z_{vertex})^2(\sigma_{\theta L}^2 + \sigma_{\theta R}^2)}. \quad (4.9)$$

where σ_{θ}^2 is the angular resolution due to multiple scattering. If there are multiple vertices, they are sorted by a figure of merit, f , defined as

$$f = \chi_{vertex}^2 + w(\chi_L^2 + \chi_R^2) \quad (4.10)$$

where the weight w is chosen empirically and $\chi_{L,R}^2$ are the reduced QT track χ^2 's. More details are contained in References [54–56].

4.5 Particle Identification

Proper particle identification (PID) is critical in a data analysis of this type. When attempting to measure a decay with a sensitivity of one part in 10^{12} , an otherwise small misidentification probability will have serious consequences. To minimize the possibility of misidentification, the PID system relies on double redundancy for both the e and μ identification systems.

PID starts with the Track-Counter Association (TKC) software, which searches for and catalogs detector hits along identified tracks. These tracks can originate from either *patrec* or one of the fitting algorithms (QT, FT), although in the final data analysis, only the fit tracks are used. After the TKC software is applied to the events, the actual PID routines use the TKC information to determine what type of particle made the track. Information on the performance of the PID detectors and associated software is given in Sections 3.5 and 3.6.

4.5.1 General Track Counter Association

The purpose of the TKC software is to catalog the properties of *all* of the hits found in the TSCs and the PID detectors, while setting a flag to show whether or not a hit is associated with an incoming track. This search is usually done by projecting the track from an upstream location to the detector and looking for hits in the detector that lie along that track with some allowance made for scattering. In most cases, these hits are then categorized by hit location, pulse height and/or timing and that information is stored for use by the PID routines.

4.5.2 Electron Identification

Electron identification is the function of the Čerenkov counter and the lead glass calorimeter software. The CER TKC software starts by projecting the track into the Čerenkov counter, simulating a Čerenkov cone and comparing the hits from the simulated cone with the actual hits. The hit locations, timing and pulse height are all stored in the TKC hit banks. The CER PID routine then determines how well these values match the expected values for an electron. A good space match means the simulated cone hit the same mirror (or mirrors) that the actual track did. A good time match means the CER corrected (after T_0 subtraction) time is ± 4 ns. A ratio of the measured pulse height to the expected pulse height was compared to the same ratio for well determined electrons to calculate a confidence level. The confidence level cut on the pulse height was placed at 0.01. Based on these values, the track was placed in one of two categories, as outlined in Table 4.2. Obviously, a third category is a logical NOT of the “possible” category. The main complication with this method is that if a μ or π has a momentum greater than its Čerenkov threshold momentum, it will mimic an e . This problem is minimized, however, by the PID redundancy provided by the PbG.

Category	Requirements
Possible e	Passes space match cut Passes time match cut
Good e	Passes “possible” cuts Passes pulse height cut

Table 4.2: CER e ID categories.

The PbG PID software, outlined in Reference [57], works by calculating the energy deposited by the particle and the shape of the shower in the PbG array. The latter is indicated by the fraction of energy deposited in the converter blocks. For electrons, the total energy deposited in the PbG should be approximately equal to the track momentum and the fraction of $E_{converter}$ to E_{total} should be greater than 0.06. Actually, these cuts are combined in a slightly more complicated way into a contour cut, as previously shown in Figure 3.21. Another useful value is D^2 , which is the square of the distance between the energy centroid of the shower and the projected hit position. Electromagnetic showers tend to be narrower than hadronic showers and thus they have a smaller D^2 . The PID routine uses these values to put the track into one of several categories, as shown in Table 4.3. Electron ID with the PbG works best at high track momenta, as the E/p fraction for e and (μ, π) are better separated at high momenta than at low momenta.

Category	Requirements
Possible e	Passes E/p cut Passes contour cut
Good e	Passes “possible” cuts Passes D^2 cut
Possible π	Fails E/p cut
Possible μ	Minimum Ionizing energy deposited in blocks

Table 4.3: PbG e ID categories.

4.5.3 Muon Identification

Muon identification is the purpose of the IMF software, which relies on the known energy loss of minimum-ionizing muons to differentiate μ 's from e 's and π 's. The IMF TKC/PID algorithm projects the tracks back into the range stack and looks for hits which might be associated with the muon track. While scanning for possibly associated hits, the algorithm search width must be constrained to avoid associating hits from another track with the current track. The first constraint, a global search width, is defined from the track entry position into the IMF, and limits the search region based on the integrated effects of multiple scattering on the particle up to the current module. The other search limit is a relative limit based on the position of the hit in the last module searched. Combined, these constraints lower the probability that the algorithm will begin to follow other tracks that might emerge from the other side of the detector. As the hits found are examined, the MHO hits are categorized by their space and time match with the projected track. For MHO panels with phototubes on both ends, the time is determined by a weighted average of the PMT times. The MRG hits, however, do not have precise timing information associated with them so only their space match is checked. This search continues through the IMF until the last plane is reached. The track length is determined by the last module hit and the observed range is compared to the expected range, with allowance for straggling in the stack, again based on the depth of the track in the IMF. The IMF PID algorithm then determines in which category, outlined in Table 4.4, the track belongs. If the first three MHO planes (X1 at the front of the range stack or X0,Y0 which comprise the μ trigger planes) contain no track associated hits, the track is considered "not a muon". Any hit in the first three MHO modules qualifies the track as a "possible" muon. If the hits pass the space and time match cuts and the observed minus expected range cut, it is considered a "good" muon. A plot of the observed minus expected range is shown in Figure 4.1. In this plot, the main peak

consists of muons from both K_L and π decays before the spectrometer while the events in the low end tail are from undecayed π tracks.

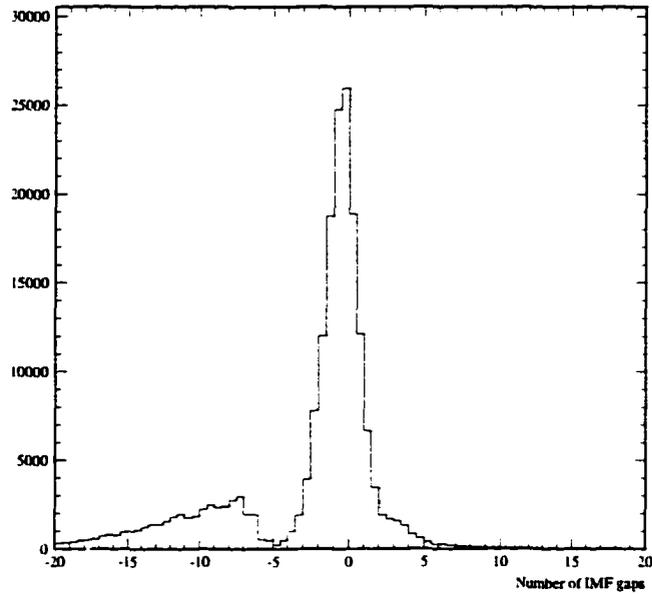


Figure 4.1: Plot of MRG track depth observed minus the track depth expected based on the measured track momentum.

Category	Requirements
Non μ	No hits in X0, Y0 or X1 MHO modules
Possible μ	Hits in X0, Y0 or X1 MHO modules
Good μ	Passes "possible" cuts Hits pass space and time match cuts Passes observed minus expected range cut

Table 4.4: IMF μ ID categories.

Chapter 5

Production Analysis

The offline analysis of $K_L^0 \rightarrow \mu e$ focused on finding good $K_L^0 \rightarrow \mu e$ events in the μe data sets and also on determining the $\pi\pi$ content for use in the normalization of $B(K_L^0 \rightarrow \mu e)$. Due to the large volume of data accumulated (1.7 Terabytes), it was necessary to pass the data through various filters in order to reduce the data set to a manageable size. Once this was done, the individual data stream analyses were done. The analyses of $K_L^0 \rightarrow \mu\mu$ and $K_L^0 \rightarrow ee$ will not be discussed here in any detail other than to say that they followed similar paths.

5.1 Organization

The production analysis was organized into three basic stages called Pass1, Pass2 and Pass3. Pass1 involved running pattern recognition (patrec) software on all of the standard trigger data and making wide cuts to eliminate most of the noise. Pass2 consisted of running the output of Pass1 through the particle identification algorithms to cut events that did not fit any of the following categories: $\pi\pi$, μe , $\mu\mu$, ee , and prescaled minimum bias data. Pass3 was a “fitting” pass, in which all of the Pass2 output data, *i.e.* any vertex which

passed Pass2, were run through the two fitting algorithms and the results written out to tape. While Pass1 and Pass2 were designed to reduce the data set size, the intent of Pass3 was not to reduce the data set size but to provide information from the fitting of the event tracks and vertices. However, the kinematic cuts applied to the fit results were tighter than those used in Pass1 so there was some reduction in the data set size.

5.1.1 Pass1

The first pass on the raw data, Pass1, looked at 1.2 billion events and reduced the set to 192 million physics events written out. Pass1 made wide cuts on patrec quantities in order to decrease noise events in the sample. This effort was carried out on the SLAC Computer Farm, which comprises 35 IBM AIX/RS-6000 Power PC based workstations. There were 20 jobs running, with four jobs pending, at all times, barring problems with the Farm. The input data were on 1238 4mm DAT tapes containing the raw 1995 and 1996 data. The output was written to 384 one Gigabyte (GB) 3290 tape cartridges for ease of use at SLAC. Pass1 ran 24 hours a day, 7 days a week for 9 weeks.

The Pass1 analysis flow is outlined by Graessle [58] and is recreated here. In stage 4, the Level 1 (L1) and Level 3 (L3) trigger bits read in are used to construct a trigger word that is then used in stage 7 to determine patrec success on this event, according to expected collinearity angle (θ_c) or transverse momentum (p_T) and mass for that trigger word. The collinearity angle is defined as the angle between the momentum sum of the two decay products and the vector from the target to the decay vertex. The flow of Pass1 is shown in Figure 5.1.

All the events are processed in the following order. First, events with illegal emulator ID, zero L1 bits or 50Hz bit set are removed. Then, events for which both the physics and calibration bits are set are written out if they pass patrec. Third, every 5000th event is saved

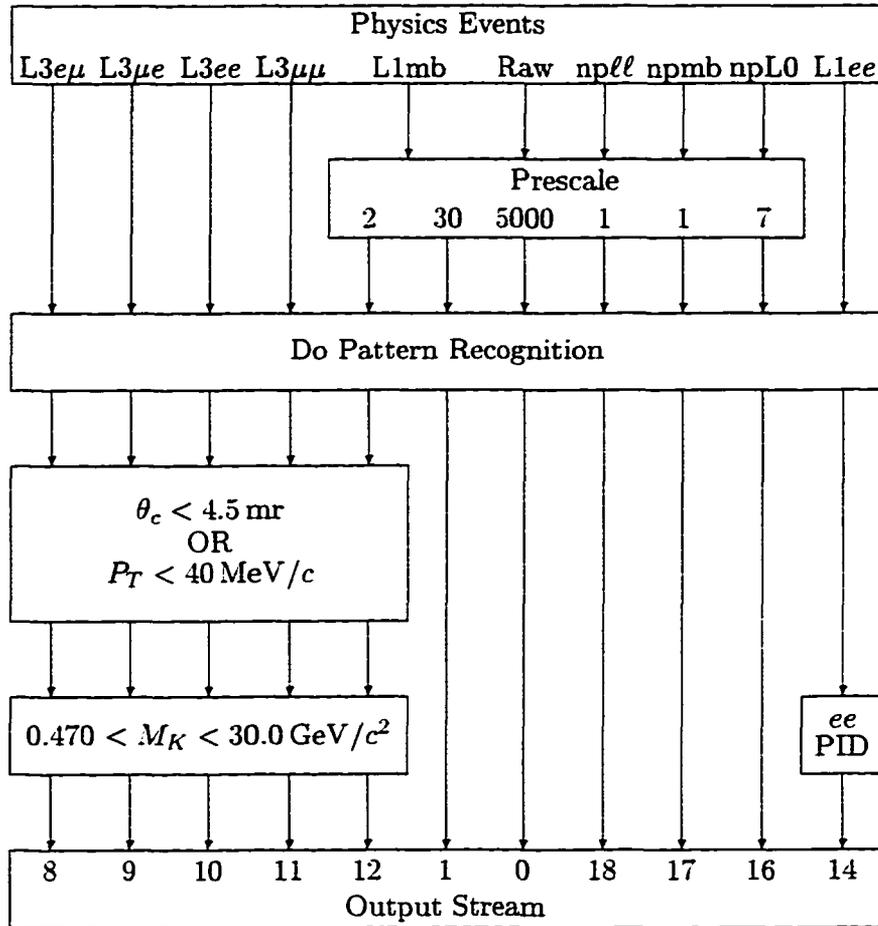


Figure 5.1: Pass1 flowchart for physics events.

if it passes patrec, regardless of trigger type or event kinematics. This is referred to as the “raw event prescale” and is included in the two-body output stream. Fourth, the minimum bias events are dealt with in two ways: 1) Every second minbias trigger is processed by patrec and cut on p_T and $M_{\pi\pi}$. 2) Every thirtieth minbias trigger is saved if it passes patrec. These are referred to as the “two-body minbias prescale” and the “semileptonic minbias prescale”, respectively. Fifth, the non-parallel triggers are considered. All non-parallel minbias and dilepton events are kept, as well as every seventh non-parallel L0 event. Sixth, the L3 dilepton events are required to pass patrec and each event must have a vertex that passes the mode-dependent invariant mass cut,

$$0.470 \text{ GeV}/c^2 < M_K < 30 \text{ GeV}/c^2$$

and either the transverse momentum cut,

$$p_T < 0.040 \text{ GeV}/c$$

or the collinearity cut,

$$\theta_c < 0.0045 \text{ rad.}$$

This option of using either the p_T or θ_c cut was included to remove the momentum dependence from the cut efficiency. Seventh, all L1 e^+e^- events are saved for use in $K_L^0 \rightarrow e^+e^-\gamma$ and $K_L^0 \rightarrow e^+e^-e^+e^-$ studies. For run 16935 and after, the L3 filter was applied to all of the L1 e^+e^- triggers but no cuts were made. Lastly, certain calibration events (ADC pedestals, PbG laser and reference tube) are prescaled and written out. Events that survive these cuts will continue to stage 10, where they are classified according to the reason responsible for their surviving Pass1 and according to which stream bit (Table 5.1) should be set in the event stream control common, `/EVSTRM/`. Table 5.2 shows the event passing statistics for each L3 trigger type.

Bit number	ICLSEK	IEVTEK
0	prescaled raw events	2-body decay
1	prescaled minbias	calibrations
2	calibration	prescaled minbias for semi-leptonics
3		missing crate or phys & cal
4	phys & cal	nonparallel triggers
5	missing crate	prescaled minbias for 2-body
6		$ee\gamma$
7		
8	$e\mu$	
9	μe	
10	ee	
11	$\mu\mu$	
12	minbias passing as $\pi\pi$	
13	L3 $\pi\pi$ passing as $\pi\pi$	
14	$ee\gamma, eeee$	
15		
16	non-parallel L0	
17	non-parallel mb	
18	non-parallel physics	

Table 5.1: Pass1 Bits.

5.1.2 Pass2

The data set surviving Pass1 was large enough, and the fitting routines slow enough, that a further reduction of the data set was needed before proceeding to the fitting stage. For this reason, Pass2 was a PID pass, with all of the Pass1 data filtered through the particle identification (PID) software. Only those events with vertices that passed in one of the relevant categories were saved. It is interesting to note that while Pass2 was intended as a self-contained activity, in practice Pass2 and Pass3 were run concurrently, with Pass2 acting as a prefilter to the Pass3 section of the software.

Pass2 was done in the following manner. First, the Pass1 tapes were read in and all pedestal, laser, reference tube and MRG pulser events were killed. However, the commons

L3 bit	Type	Read	Reconstructed	Written	%out
0	$e\mu$	197,999,685	179,007,065	36,213,949	18.29
1	μe	200,575,046	181,354,653	36,600,412	18.25
2	ee	8,863,452	7,492,132	5,381,688	60.72
3	$\mu\mu$	385,618,825	346,512,761	80,646,168	20.91
4	$\pi\pi$	26,334,104	24,232,075	15,125,307	57.44
5	MB	126,405,973	52,792,809	19,626,627	15.53

Table 5.2: Pass1 Statistics.

/DRVPED/ and /DRVDRF/ (for pedestal and laser/reference data, respectively) are written out for each event. In stage 4, the Pass2 stream bits, listed in Table 5.3, were assigned on the basis of Pass1 stream bits.

In stage 7, the events are run through patrec again, even though they were run through patrec in Pass1. Because the data had already been filtered through a slightly different version of patrec in Pass1, there was concern that the results might be biased. However, the cuts made in Pass1 were shown to be loose enough so as not to affect the outcome of Pass2. After patrec was run, the software looped over each patrec vertex and flagged those vertices which passed the mass cut

$$M > 0.470 \text{ GeV}/c^2$$

and p_T cut

$$p_T < 0.030 \text{ GeV}/c.$$

No mass or p_T cuts were made in the $ee\gamma$ or minbias streams. The event was saved if any vertex passed those cuts. Finally, each flagged vertex was run through PID and if any vertex passed PID in one of the interesting categories, the event was saved. Only in the $e\mu$, $\mu\mu$, ee and $ee\gamma$ ¹ streams were PID cuts applied. The cuts were:

¹ $ee\gamma$ events are L1 ee triggers that did not pass the L3 ee mass and p_T cuts. These events were included in the general data stream when the L3 ee cuts were removed (Section 3.7.3)

Bit number	ICLSEK	IEVTEK
0	prescaled raw events	$e\mu$ and μe
1	prescaled minbias	ee
2	scaler events	$\mu\mu$
3		$\pi\pi$
4	phys & cal	$ee\gamma, eeee$
5		minbias
6		raw events
7		phys & cal
8	$e\mu$	non-parallel
9	μe	scaler
10	$\mu\mu$	μe background
11	ee	
12	$\pi\pi$	
13		
14	$ee\gamma, eeee$	
15		
16	non-parallel L0	
17	non-parallel mb	
18	non-parallel physics	

Table 5.3: Pass2 Bits.

Electron cuts

1. A TDC hit in the CER within ± 10 ns of the event time.
2. PbG measurement of $E/p > 0.5$

Muon cut

1. Either TKXIDR or TKXIDH (MRG and MHO μ ID code, respectively) returns a “possible” μ .

The result of running Pass2 was a reduction of 192 million Pass1 physics events to 54 million physics events to be passed on to Pass3 software. Table 5.4 shows the percentage of events passing Pass2 listed by stream bit.

Stream Name	patrec passing	Mass and p_T	Particle ID
$e\mu, \mu e$	87.4	35.6	17.5
ee	85.1	12.8	12.8
$\mu\mu$	86.6	40.3	25.1
$\pi\pi$	91.0	58.2	58.2
$ee\gamma$	85.7	85.7	42.4
minbias	89.5	89.5	89.5

Table 5.4: Pass2 Statistics.

5.1.3 Pass3

The first two production passes were directed toward data set reduction, i.e. reducing the number of events to a reasonable amount with high efficiency cuts. In Pass3 the goal was not to reduce the size of the data set but to fit all of the Pass2 events, using both FT and QT fitters. Pass3 only fit vertices that passed the Pass2 PID criteria. If the vertex passed either fitter with a mode dependent invariant mass $\geq 0.475 \text{ MeV}/c^2$, the event was accepted and written out, along with the patrec and the QT and FT information. The differences in mass and p_T^2 results between FT and QT are shown in Figures 5.2 and 5.3.

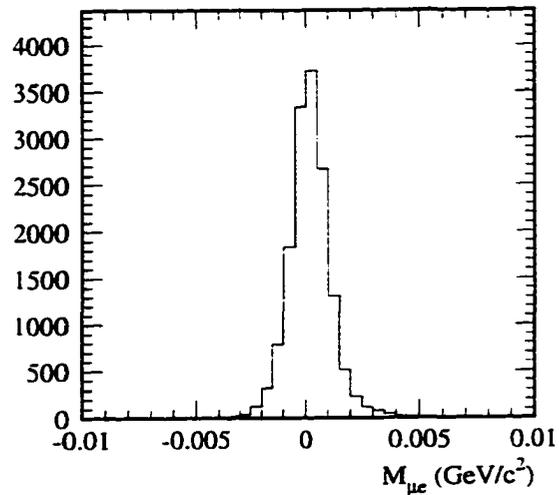


Figure 5.2: Difference in FT/QT masses in the $\mu\mu$ sample.

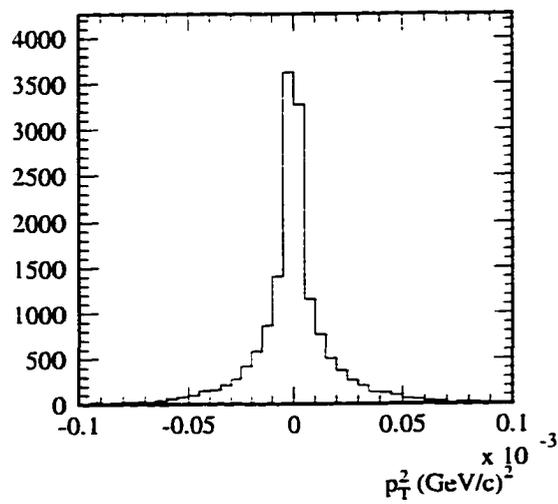


Figure 5.3: Difference in FT/QT p_T^2 in the $\mu\mu$ sample.

Chapter 6

$K_L^0 \rightarrow \mu e$ Analysis

The μe analysis was done using the *blind analysis* protocol, shown schematically in Figure 6.1. A blind analysis attempts to minimize biases by keeping the region in and around the expected signal region “blacked-out” until all of the cuts have been applied to the data. This reduces the temptation to design cuts in a manner that would unreasonably result in a desired, or unfairly exclude an undesired, signal. In a blind analysis, all of the cuts are determined by studying the background region *around* the blacked-out region. Once these cuts are fixed, they are applied to the data *in* the blacked-out region and the results are revealed. Ideally, the results would then stand on their own and whatever signal, or lack thereof, remained would be the final result.

In the $K_L^0 \rightarrow \mu e$ analysis, the raw Pass3 output for the μe data stream is shown in Figure 6.2. The exclusion region from

$$490 \leq M_{\mu e} \leq 505 \text{ MeV}/c^2$$

and

$$p_T^2 \leq 100 \text{ (MeV}/c)^2$$

is blocked out on this plot.

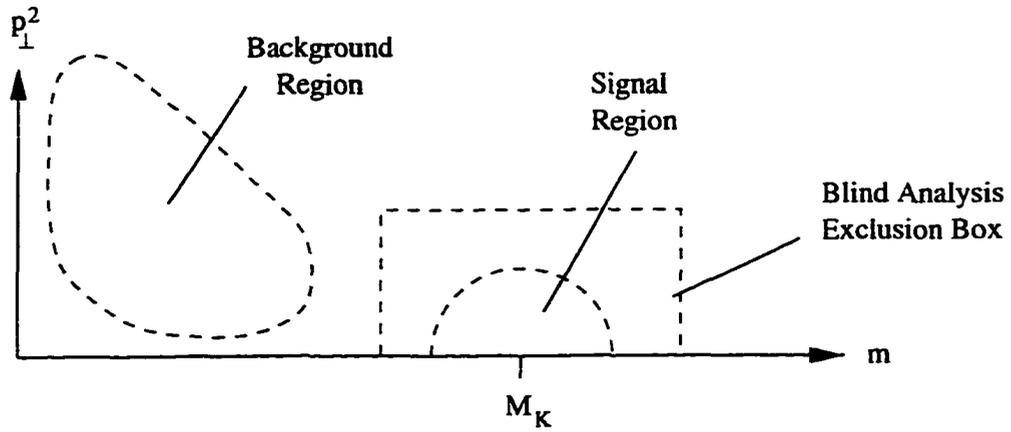


Figure 6.1: Schematic drawing of the blind analysis methodology.

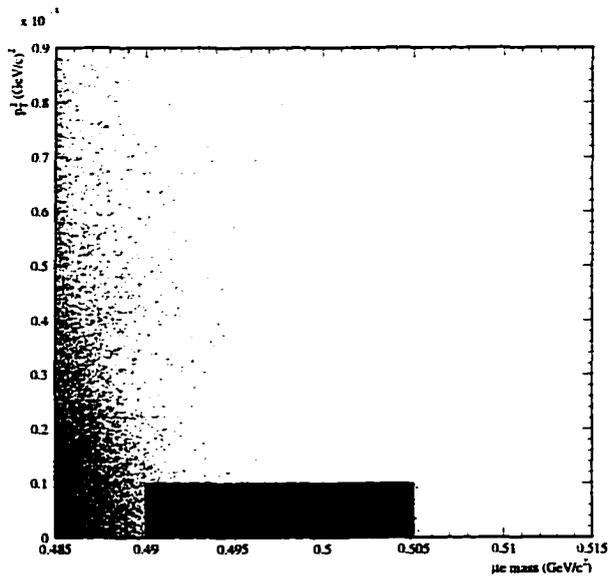


Figure 6.2: p_T^2 vs. $M_{\mu e}$ from uncut Pass3 output with the exclusion region indicated.

6.1 Backgrounds

The $K_L^0 \rightarrow \mu e$ decay is in principle free of physics backgrounds. If the detector systems were 100% efficient and multiple scattering did not occur, no decay would mimic a $K_L^0 \rightarrow \mu e$ event. However, because the detector systems are not perfectly efficient and multiple scattering does happen, kinematic backgrounds become a problem. These backgrounds arise from other K_L decays that get shifted upwards in invariant mass due to spectrometer errors or particle misidentification. The primary goal while designing cuts for the $K_L^0 \rightarrow \mu e$ analysis was not to maximize the signal efficiency but to lower the background contamination in the signal region to an expected value of 0.1 event or less.

The K_L decay that posed the most severe problem for identifying $K_L^0 \rightarrow \mu e$ is the K_{e3} decay. If the π decayed into a μ and the μ retained most of the π momentum, the μ PID software might call that track a μ . In the ideal case, K_{e3} would not be a problem because the kinematic end point (where $E_\nu \approx 0$) is $8.4 \text{ MeV}/c^2$ below the K_L mass and could be excluded with kinematic cuts. However, scattering and detector resolution can contribute to smearing in the mass distribution and result in a mass tail that extends up to and beyond the K_L mass. Particle misidentification can make this problem worse, especially in the case of double misidentification, described below. The main categories of background presented in this section are event pileup, tracking errors, particle type misidentification, and pion decay with Mott scattering.

6.1.1 Pileup

A “pileup” event is one in which a second K_L decays in time (or just slightly out of time) with the primary K_L decay. For example, a K_{e3} and a $K_{\mu 3}$ decay in time with each other could be a dangerous background. Each decay by itself may not be in a kinematic range of interest to be accepted by the data acquisition system. However, one track from

each decay might, by chance, form a good vertex with a good invariant mass and p_T^2 value which would be accepted by the data acquisition system. Even if the other tracks do not interfere with the kinematics of the primary decay, they may leave hits in the PID system that subsequently lead to misidentification of the primary decay products. The “patrec extra track” cut suppresses this background by cutting events that contain an extra, distinct patrec track, *i.e.*, the other track does not share any segments with the primary track. The events with an extra track slightly out of time are not excluded by this cut because patrec does not consider those hits to be part of a real track and thus they are not in the list of tracks.

6.1.2 Tracking Errors

Incorrect identification of drift chamber ambiguities is an example of a tracking error which can contaminate the event sample. By choosing the wrong side of a tracking cell, the fitted track position may be displaced by up to 1 cm from the actual track position. If the hit mismeasurement is in the “wrong” direction, it can result in a higher invariant mass which may cause the event to be mistaken as a $K_L^0 \rightarrow \mu e$ decay. The effect that is most important is a tracking error in SDC1 or SDC2 that increases the opening angle of the K_L decay. Background from this effect is suppressed by vertex and track χ^2 cuts and are not expected to be a significant problem for the $K_L^0 \rightarrow \mu e$ analysis.

6.1.3 Single and Double Misidentification

The invariant mass of a K_{e3} decay (with $E_\nu \approx 0$) can be approximated by [59]

$$M_{\pi e}^2 = m_\pi^2 \left(1 + \frac{p_e}{p_\pi}\right) + m_e^2 \left(1 + \frac{p_\pi}{p_e}\right) + p_e p_\pi \alpha^2, \quad (6.1)$$

where α is the angle between the π and the e momenta. If the particle type is misidentified, the calculated mass will be different from the actual mass. If there is also a mistake in

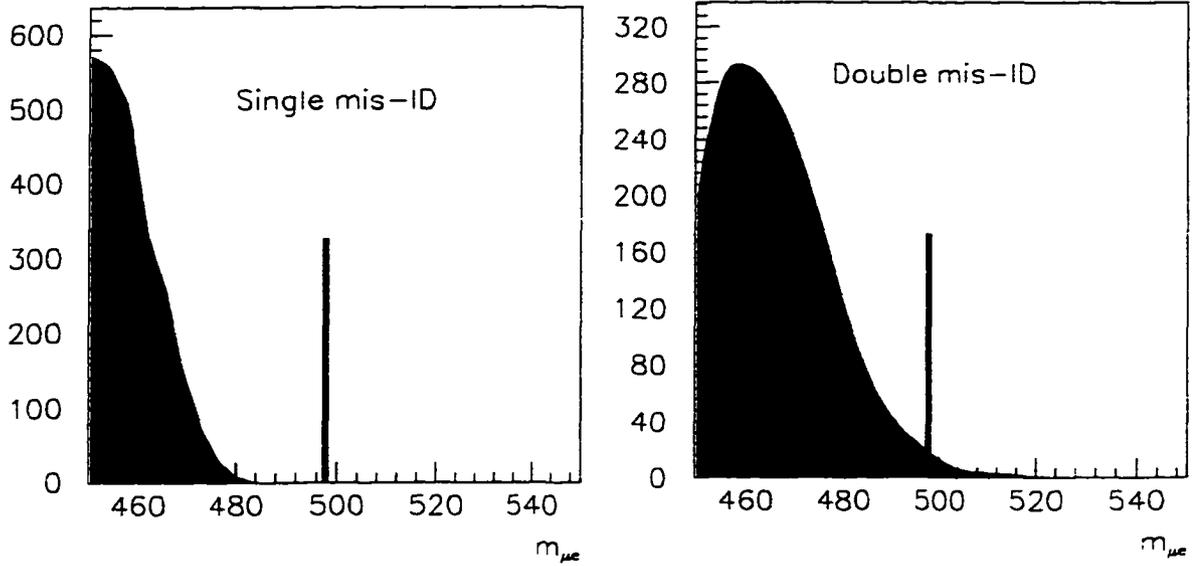


Figure 6.3: The effects of single (left) and double (right) misidentification in K_{e3} decays on $M_{\mu e}$, measured in MeV/c^2 .

measuring the kinematics, it is possible that the calculated mass equals or exceeds the K_L mass.

The main source of misidentification (mis-ID) results from misidentifying the π as a μ in a K_{e3} event. This will most likely arise via π decay but can also occur from low momentum π 's with a coincident noise track in the IMF system. Fortunately, for the standard case of π decay in a K_{e3} event, the endpoint of the invariant $M_{\mu e}$ distribution is 8.4 MeV lower than the K_L mass (Figure 6.3 (left)). With a mass resolution on the order of $1.4 \text{ MeV}/c^2$, this limit is 6σ below M_K and is not considered a serious problem for the $K_L^0 \rightarrow \mu e$ analysis.

A less likely, but potentially more dangerous, source of background is a double mis-ID where, in a K_{e3} decay, the π is identified as a e and the e is identified as a μ . The loss in mass due to the $\pi \rightarrow e$ mis-ID is countered somewhat by the $e \rightarrow \mu$ mis-ID but the real mass shift comes from the switched track momenta. With the right kinematics (within the range allowed by E871), it is possible to have a μe invariant mass that ranges from below

to well above the K_L mass (Figure 6.3(right)). Fortunately, due to the double redundancy in the PID detectors, the probability of a double mis-ID is extremely small.

6.1.4 Pion Decay with Electron Mott scattering

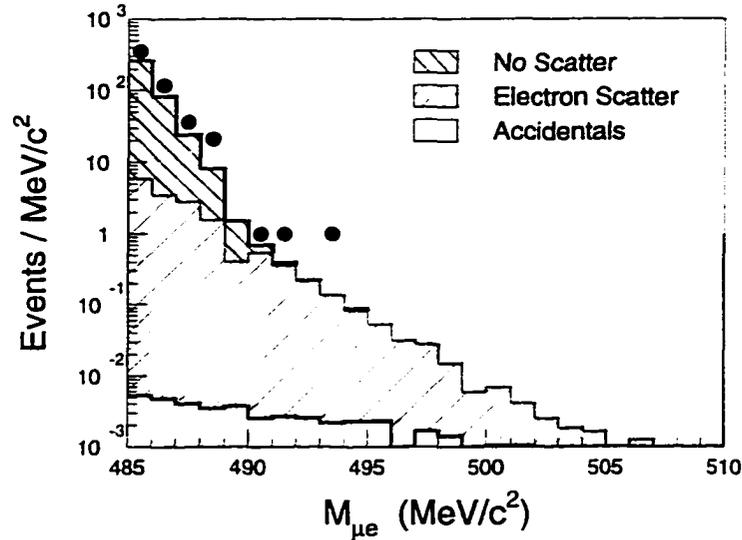


Figure 6.4: Monte Carlo $M_{\mu e}$ invariant mass plot of expected background from K_{e3} decays with pion decay. The solid lines represent the cumulative effect of three types of background and the dots represent low p_T^2 (< 20 $(\text{MeV}/c)^2$) background events from the μe data set. (Note: The three events shown with $M_{\mu e} > 490$ MeV/c^2 actually lie within the exclusion region. They are shown here merely to illustrate the integrity of the Monte Carlo generated background. The presence of these three events was ascertained only after the exclusion region was opened, as described in Chapter 7.) [60]

The most serious background for $K_L^0 \rightarrow \mu e$ is from a K_{e3} decay, in which the π decayed in the vacuum chamber and the e underwent a Mott scatter in the vacuum window or the first straw chamber window. With the π decaying in the vacuum chamber, there is no way to discern that the track was ever anything but a μ . There would be no characteristic “kink” in the spectrometer track, which usually happens when a π decays inside the spectrometer system. The muon PID software would call the track a μ because it met

the range requirements determined by the spectrometer momentum. Likewise, for electron Mott scattering, the electron would otherwise behave normally in the spectrometer and the electron PID software would recognize it as such. If either of these processes happened individually, this event would be removed by a p_T cut. However, if these two processes are simultaneous and both particle trajectories are pushed away from the beamline almost equally in opposite directions, this event will appear to be a well identified μe pair with a large invariant mass (Figure 6.4), due to the increase in the opening angle.

6.2 Cuts

While there was no expected physics background to $K_L^0 \rightarrow \mu e$ decays, kinematic backgrounds were expected to move events closer to the signal region. Therefore, the μe analysis used the FT fitter information instead of QT information because the FT mass resolution was better than that of QT. Most of the cuts were determined simply by studying their effect on $\pi\pi$ data or μe Monte Carlo information. However, the following track and vertex cuts were optimized simultaneously to achieve the highest efficiency reasonably attainable:

- χ_{vertex}^2
- $\frac{\chi_{track}^2}{\nu}$
- $\frac{\delta p}{p}$
- $\left| \frac{p_\mu - p_e}{p_\mu + p_e} \right|$
- p_e
- $|p_T^\parallel|$

- Signal region
- FT \leftrightarrow QT match

The final cuts, listed in Table 6.1, were applied to the μe data stream. The efficiency shown for a particular cut is measured by applying that cut to the test data stream listed [60, 61].

Category	Cut Applied	Efficiency	Data Stream
Event	Tracks passed both fitters	0.9345 ± 0.0003	$\mu\mu$ data
	No extra patrec tracks	0.7554 ± 0.0030	$\mu\mu$ data
	Fiducal volume cut	0.9881 ± 0.0001	$\mu\mu$ data
	$ \delta M_{FT \leftrightarrow QT} \leq 3 \text{ MeV}/c^2$	0.9856 ± 0.0010	$\mu\mu$ data
	$ \delta p_{T, FT \leftrightarrow QT}^2 \leq 70 \text{ (MeV}/c)^2$	0.9758 ± 0.0012	$\mu\mu$ data
	$ p_T^{\parallel} \leq -0.0025 + 0.003 p_e \text{ GeV}/c$	0.9495 ± 0.0014	μe MC
Vertex	$21 \geq z \geq 9.75 \text{ m}$	Acceptance	
	FT $\chi_v^2 \leq 8.0$	0.8398 ± 0.0027	$\mu\mu$ data
	$-2.00 \text{ mR} \leq \theta_x \leq 2.75 \text{ mR}$ (1995)	0.9698 ± 0.0017	$\mu\mu$ data
	$-2.50 \text{ mR} \leq \theta_x \leq 2.50 \text{ mR}$ (1996)	0.9768 ± 0.0021	$\mu\mu$ data
	$-10.0 \text{ mR} \leq \theta_y \leq 10.0 \text{ mR}$	0.9979 ± 0.0004	$\mu\mu$ data
Track	FT $\chi_v^2 \leq 10.0$ (μ)	0.9349 ± 0.0014	μ data
	FT $\chi_v^2 \leq 40.0$ (e)	0.9955 ± 0.0003	e data
	$ \frac{p_\mu - p_e}{p_\mu + p_e} \leq 0.500$	0.9621 ± 0.0013	μe MC
	$ \frac{p_{front} - p_{back}}{p} \leq 0.05$	0.9845 ± 0.0007	$\mu\mu$ data
	$ \delta t_{SPEC} \leq 2.0 \text{ ns}$	0.9918 ± 0.0007	$\mu\mu$ data
	$ \frac{dx}{dz} _{DC6} \leq 35 \text{ mR}$	0.9997 ± 0.0001	$\mu\mu$ data
	$1 \leq p_\mu \leq 8 \text{ GeV}/c$	0.9999 ± 0.0001	μ data
	$1 \leq p_e \leq 8 \text{ GeV}/c$	0.9993 ± 0.0001	e data
PID	Good e ID (CER)	0.9740 ± 0.0010	minbias data
	Good e ID (PbG)	0.9862 ± 0.0007	minbias data
	Good μ ID (IMF)	0.9580 ± 0.0012	minbias data
	PbG e veto on μ side	0.9958 ± 0.0004	minbias data

Table 6.1: Table of μe analysis cuts and corresponding efficiencies.

6.2.1 Event Cuts

This category of cuts is a catch-all for cuts that do not fit into the other categories. First, the “pass both fitters” cut was simply a cut on the quirks of each fitting algorithm. If the event was not of good enough quality to pass both fitters, it was rejected. The second cut, the “no extra patrec tracks” cut, was used to suppress the background from pileup events. The cut worked by searching for separate and distinct patrec tracks, *i.e.*, tracks that did not share segments with any of the primary tracks, in both views. If any extra patrec tracks were found, the event was cut. The third cut was a fiducial volume cut, which required that the tracks pass through all well-understood regions of the spectrometer. This cut had no serious effect on the data but was applied because it was also applied to the Monte Carlo simulations, for which it did have an effect. Another event cut used was a cut on the FT \leftrightarrow QT mass and transverse momentum match. This cut forced the two fitting routines to agree with each other to within 2.5σ of their resolutions.

6.2.2 Vertex Cuts

The first vertex cut made is the cut on the Z position of the vertex. This cut is made because the fringe field of the downstream sweeping magnet widens the opening angle for decays that occur in the upstream portion of the tank. This fringe field boosts otherwise low mass events above the invariant mass cutoff and causes them to be accepted by the production analysis. This effect, shown in Figure 6.5, appears as a disproportionate number of events in the upstream section of the vacuum tank. No efficiency is calculated for this cut as it affects the acceptance of all K_L decay modes equally.

The cut on the beam divergence in the X and Y view, defined as

$$\theta_x = \frac{x_{vertex} - x_{target}}{z_{vertex} - z_{target}} \quad (6.2)$$

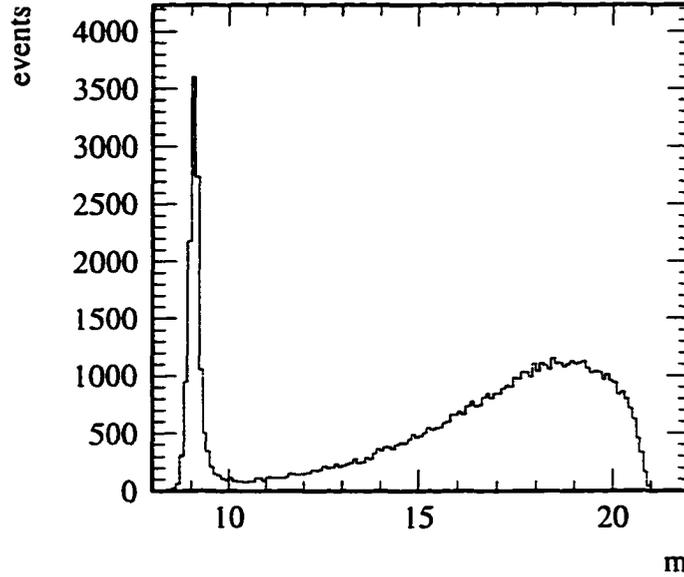


Figure 6.5: A plot of the number of events vs. the Z position of the vertex.

and

$$\theta_y = \frac{y_{vertex} - y_{target}}{z_{vertex} - z_{target}}, \quad (6.3)$$

requires that the event originate from somewhere in the neutral beam, which is limited by the collimators upstream. This requirement removes off-axis K_L decays and events with poorly reconstructed vertices. The cuts applied in the X view differ from 1995 to 1996 because the beam profile changed slightly between running periods, as shown in Figure 6.6. The Y view beam divergence did not change appreciably between running periods (Figure 6.7) and thus the cut on θ_y remained the same for both years.

The cuts on χ_{vertex}^2 and $|\delta t_{SPEC}|$ ensured that the chosen track pair indeed originated from the same decay vertex and thus helped to lower the rate from accidentals. The cut on p_T^{\parallel} was developed to suppress background from K_{e3} decays where the electron underwent Mott scattering in the vacuum window.

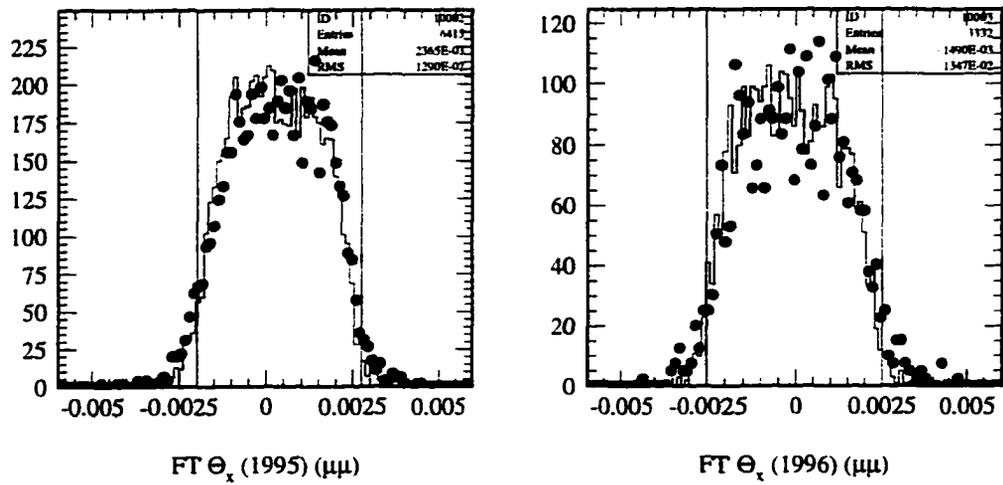


Figure 6.6: A plot of the beam divergence in X for 1995 (left) and 1996 (right). $\mu\mu$ data are plotted as a solid line while the μe background events are denoted by dots.

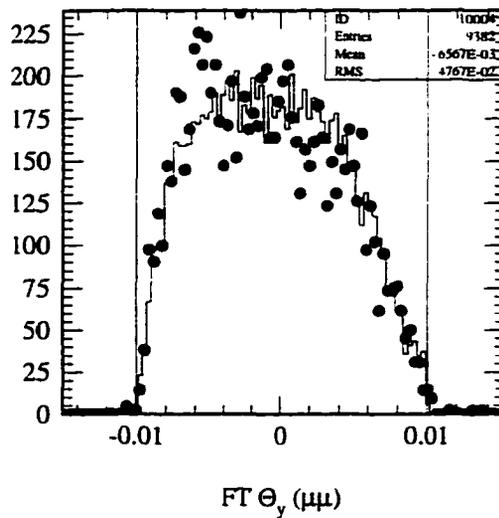


Figure 6.7: A plot of the beam divergence in Y . $\mu\mu$ data are plotted as a solid line while the μe background events are denoted by dots.

6.2.3 Track Cuts

The $\frac{\chi_{\text{track}}^2}{\nu}$ cut was set to be rather loose on the side with the identified e , to allow for more multiple scattering expected for the electron, and tight on the μ side, to guard against π decay in the spectrometer. The cut on $|\frac{\delta p}{p}|$, shown in Figure 6.8, guards against π decay as well but also suppresses tracking errors by comparing the front and back measured momenta.

The momentum asymmetry cut in Figure 6.9 was applied to remove high mass K_{e3} events and suppress possible mis-ID background. This cut is useful in rejecting semileptonic background due to the “leading pion” argument, presented in Reference [62]. Briefly stated, the invariant mass of a K_{e3} decay misidentified as a $K_L^0 \rightarrow \mu e$ gets closest to $(M_K - 8.4 \text{ MeV}/c^2)$ when the πe momentum asymmetry is high in the direction of the π . Therefore, a cut on events with high asymmetry is a cut on K_{e3} decays which are misidentified as $K_L^0 \rightarrow \mu e$ decays.

6.2.4 Particle Identification Cuts

The PID cuts applied required that the CER and PbG both positively identify the e and the IMF positively identify the μ . The good TSC L0 cut insured that the two tracks would have fired the L0 trigger. Finally, for safety, a very efficient veto for e 's on the μ side was applied to further suppress the double misid background.

6.2.5 Cut Results

When all of these cuts are applied to the events which pass Pass3 cuts, the plot shown in Figure 6.10 results. The background at low $M_{\mu e}$ is a combination of the tail expected from normal K_{e3} decays and from K_{e3} decays in which the π decays into a μ upstream of SDC1.

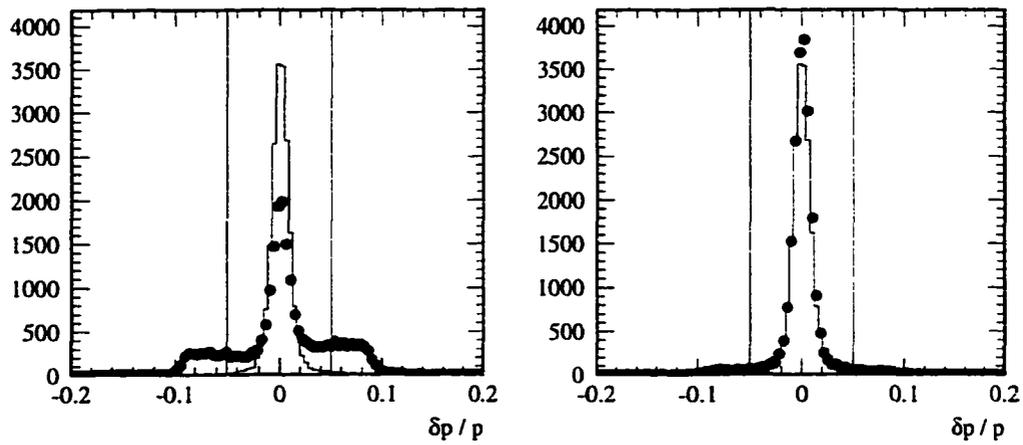


Figure 6.8: A plot of the front/back momentum comparison between μ 's (left) and e 's (right) and the $\mu\mu$ sample. $\mu\mu$ data are plotted as a solid line while the μe background events are denoted by dots. The wide shoulders on the μ side result from π decays in the spectrometer.

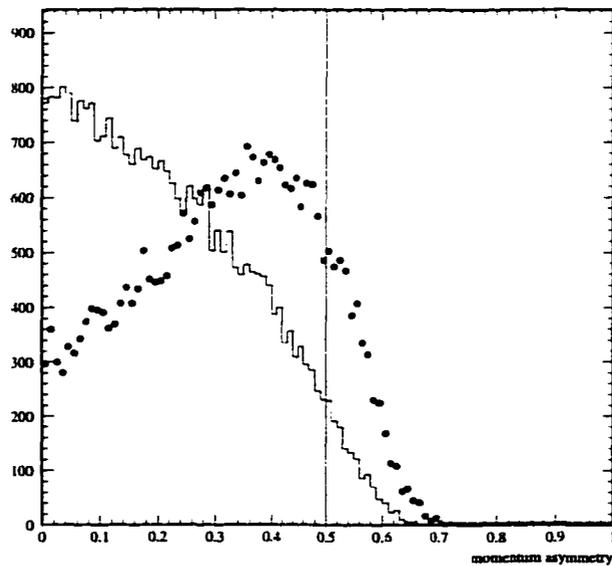


Figure 6.9: A plot of the momentum asymmetry between the two tracks. μe MC event asymmetry is plotted as a solid line while the μe background events are denoted by dots.

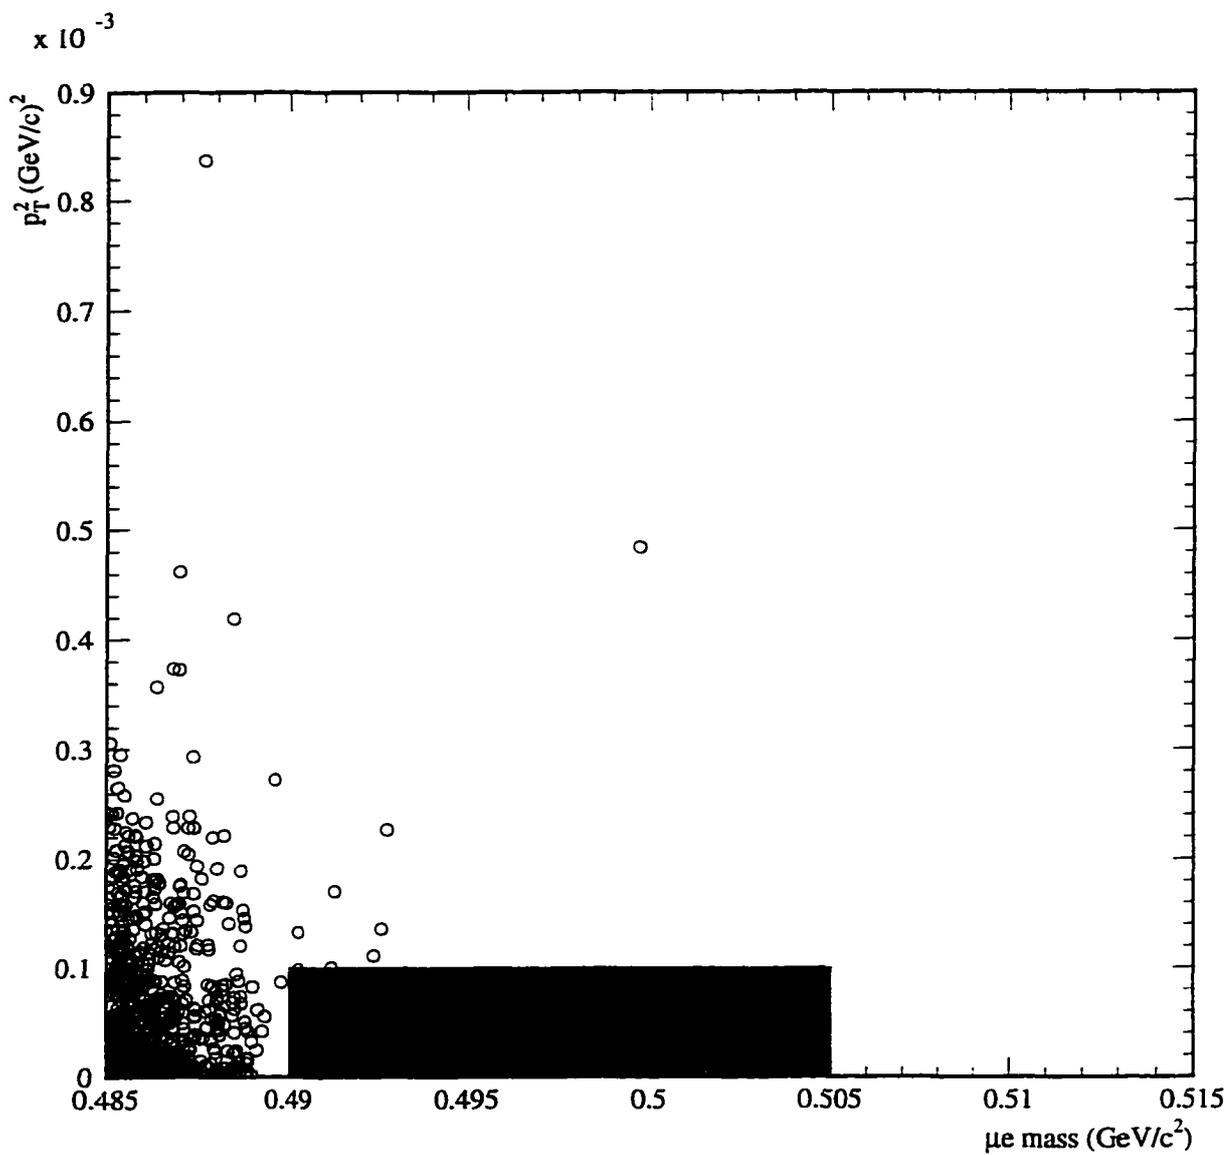


Figure 6.10: FT based analysis background plot of p_T^2 vs. $M_{\mu e}$ with exclusion region shown.

Chapter 7

Analysis Results

7.1 Signal Region

The last step in the blind analysis, after cuts have been established, is to define a signal region within which any signal would be considered a positive result. This must be done with the exclusion region veto still in place to insure the integrity of the blind analysis. Deciding upon a signal region after removing the exclusion region veto might result in a signal region chosen simply to include or exclude certain events.

The $K_L^0 \rightarrow \mu e$ signal region was determined by examining the backgrounds outside the exclusion region and attempting to minimize them while optimizing the signal region. The simplest signal region to use was a box in $(M_{\mu e}, p_T^2)$ space, but such a shape might have allowed more background into our signal region. The most problematic background for this experiment was the electron Mott scattering in the vacuum window coupled with simultaneous π decay in the vacuum tank. Because such events would tend to have higher p_T^2 , the lower edge of the signal box was shaped like a half-ellipse to reduce the effect of the Mott scattering events. Therefore, the signal box was defined in two parts.

For $M_{\mu e} < M_K$,

$$\frac{(M_K - M_{\mu e})^2}{(2.4 \text{ MeV}/c^2)^2} + \frac{(p_T^2)^2}{(20 \text{ (MeV}/c)^2)^2} \leq 1, \quad (7.1)$$

while for $M_{\mu e} \geq M_K$,

$$|M_K - M_{\mu e}| \leq 4.0 \text{ MeV}/c^2$$

and

$$p_T^2 \leq 20 \text{ (MeV}/c)^2.$$

Based on Monte Carlo distributions, shown in Figure 7.1. the efficiency of these signal cuts is 88.85%. These same distributions predict the additional signal loss for $K_L^0 \rightarrow \mu e$ due to inner bremsstrahlung [63] to be 4.67%.

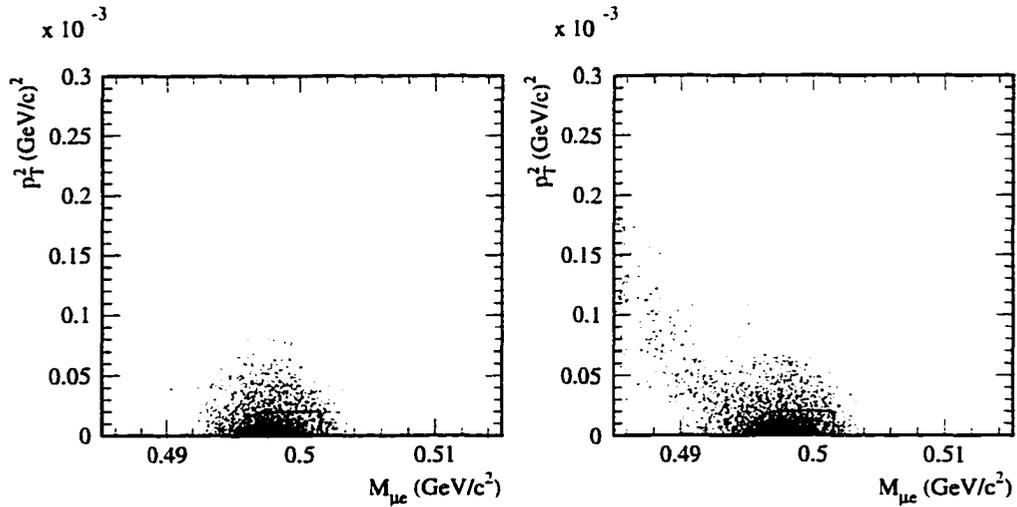


Figure 7.1: Plot of p_T^2 vs. $M_{\mu e}$ for Monte Carlo μe events without (left) and with (right) inner bremsstrahlung. The signal region as defined in Eq. 7.1 is shown for reference.

Once the signal region was determined, the exclusion region was opened. This opening took place on May 3, 1998. No events were observed in or significantly close to the signal region (Figure 7.2).

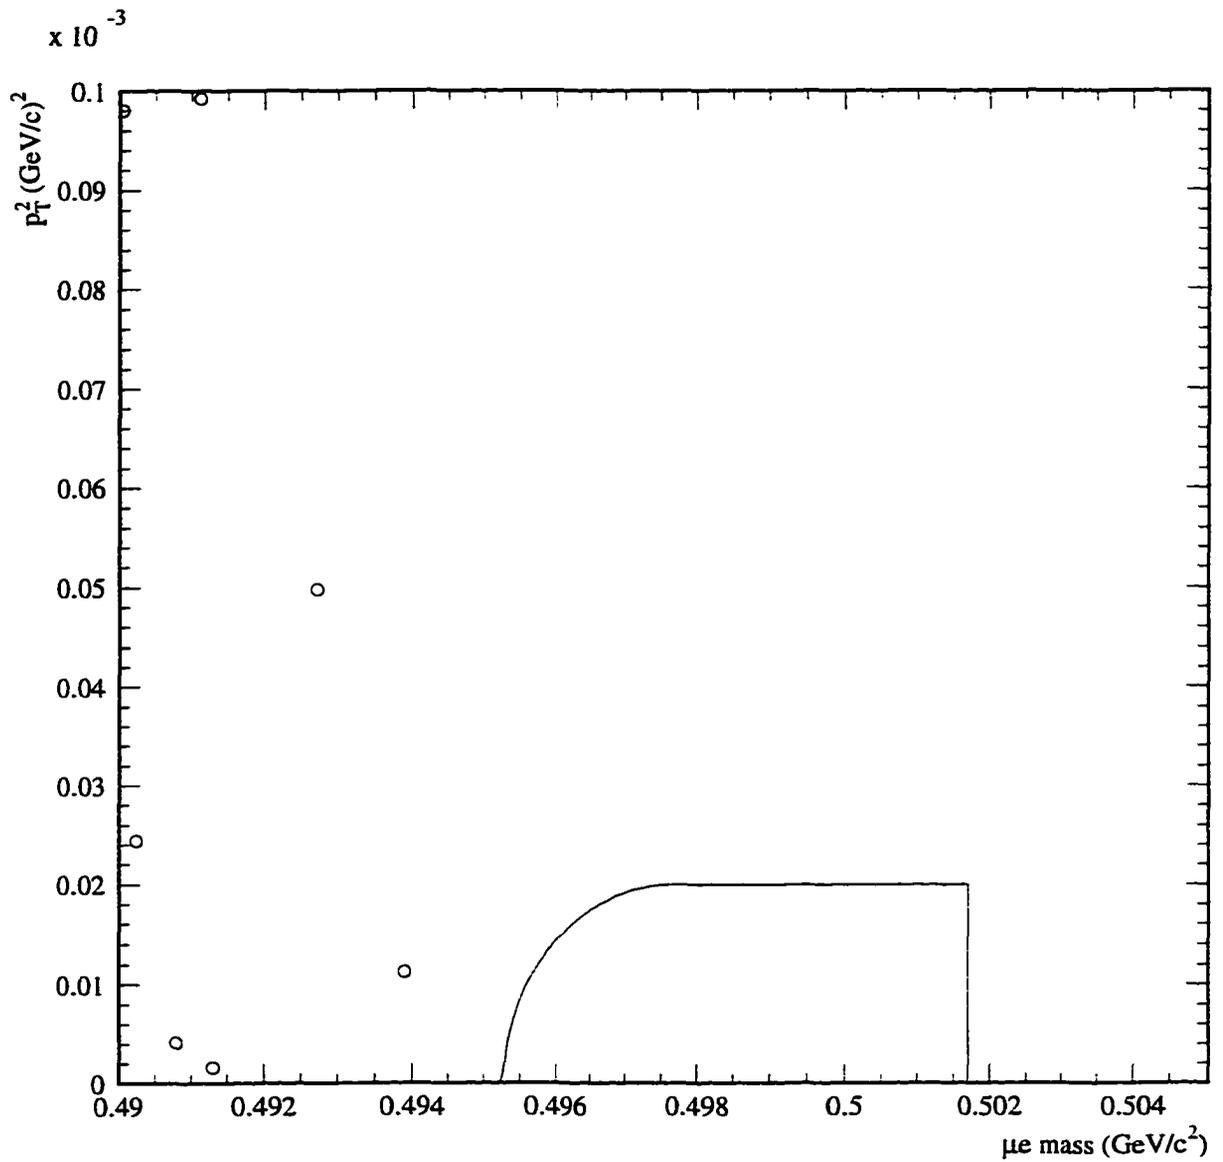


Figure 7.2: FT Based analysis plot of p_T^2 vs. $M_{\mu e}$ inside the exclusion region. The signal box is drawn for reference.

7.2 Single Event Sensitivity

Because no events appear in the signal region, the next step is to determine the sensitivity of the experiment to $K_L^0 \rightarrow \mu e$ given the analysis cuts applied. The single event sensitivity (SES) is the level at which one event is expected to be observed. In standard practice, the decay $K_L^0 \rightarrow \pi\pi$ is used to normalize rare neutral kaon decays and, indeed, this will be done for the final report of E871 results. However, when using $K_L^0 \rightarrow \pi\pi$, there are a number of effects, such as pion absorption and K_S contamination, that must be considered. These effects can be ignored when using $K_L^0 \rightarrow \mu\mu$ as a normalization sample. Given the large (~ 6500 pre-cut) number of $K_L^0 \rightarrow \mu\mu$ events seen by E871, the statistical error obtained from both normalizations is comparable. Therefore, this dissertation will present the branching ratio limit calculated relative to the rate for $K_L^0 \rightarrow \mu\mu$. The calculated SES for μe is given by

$$\text{SES}(\mu e) = \left(\frac{B(K_L^0 \rightarrow \mu\mu)}{N(\mu\mu)} \right) \left(\frac{\epsilon_{\mu\mu}}{\epsilon_{\mu e}} \right) \left(\frac{A_{\mu\mu}}{A_{\mu e}} \right) \left(\frac{1}{\kappa_{IB}} \right), \quad (7.2)$$

where $\epsilon_{(\mu\mu,\mu e)}$ and $A_{(\mu\mu,\mu e)}$ are the relative efficiencies and experimental acceptances, respectively, for those decays. κ_{IB} is the correction factor for inner bremsstrahlung and was determined from previous Monte Carlo calculations to be 0.9533. The determination of these numbers is presented in the following sections.

7.2.1 $\mu\mu$ Normalization Sample

The $\mu\mu$ normalization sample was created by applying the relevant μe kinematic cuts, *i.e.*, those cuts that apply to both $\mu\mu$ and μe data with little or no difference, and $\mu\mu$ PID cuts to the $\mu\mu$ stripped data sample. Because the $K_{\mu 3}$ background is roughly flat in p_T^2 space (Figure 7.3), the background subtraction can be carried out by fitting a straight line to the higher p_T^2 region, from 100 to 400 (MeV/c)². Then, extrapolating this line into the

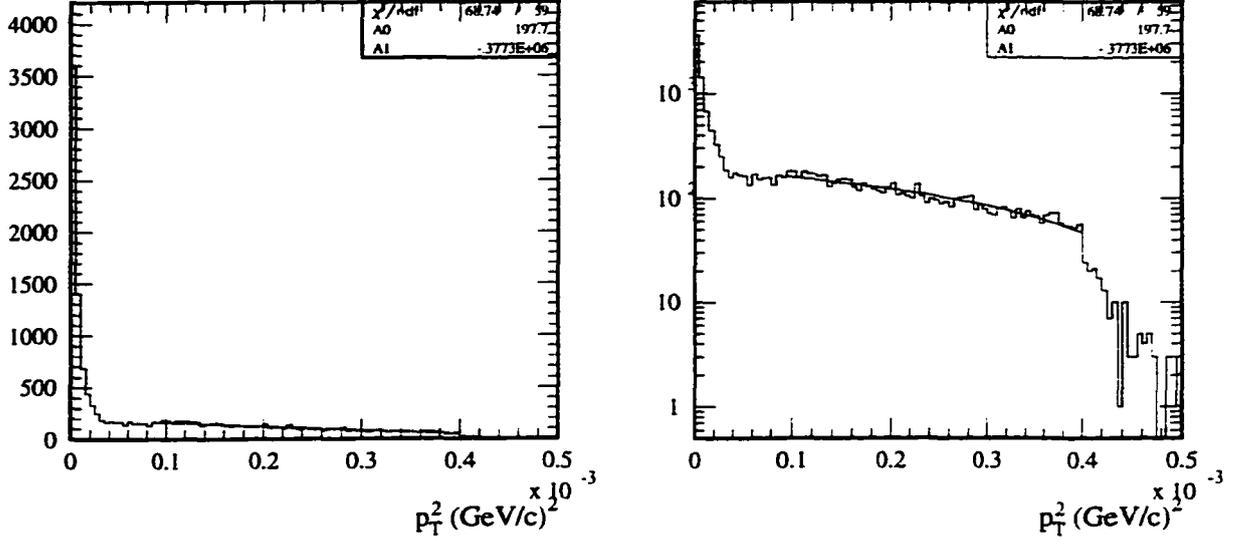


Figure 7.3: Plot of p_T^2 for the $\mu\mu$ data sample showing the fit for background subtraction on a linear (left) and logarithmic (right) scale.

lower p_T^2 region and subtracting the data from the fit one obtains the number of $\mu\mu$ events remaining in the sample. Using this method, 5470 ± 70 $\mu\mu$ events survive the relevant μe cuts. Figure 7.4 shows the $\mu\mu$ distribution in mass space.

7.2.2 Experimental Acceptances for $\mu\mu$ and μe

The experimental acceptances, $A_{\mu\mu}$ and $A_{\mu e}$, are determined using the E871 Monte Carlo software. These acceptances are based on the ratio of the number of events passing the parallelism condition (Stage 6) to the number of events generated (Stage 1). This ratio takes into account the geometric acceptance of the apparatus as well as the decay mode dependence of the trigger parallelism. This dependence arises due to the differences in the p_T of each decay mode compared to the spectrometer magnet settings of a net -220 MeV/c p_T kick. For E871, the acceptance ratio is [60]:

$$\frac{A_{\mu\mu}}{A_{\mu e}} = \frac{0.0285}{0.0234} = 1.22 \quad (7.3)$$

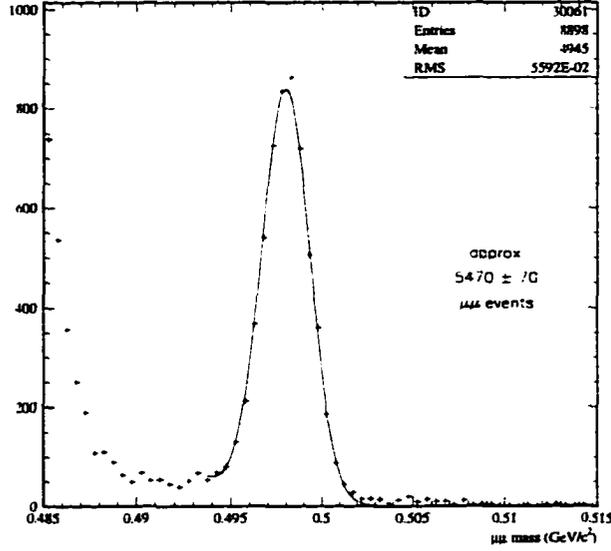


Figure 7.4: Plot of $M_{\mu\mu}$ with all relevant μe cuts applied.

7.2.3 Relative Efficiencies for $\mu\mu$ and μe

By using a $\mu\mu$ sample to normalize the μe result, many efficiency factors cancel or have minimal effect in the ratio of $\frac{\epsilon_{\mu\mu}}{\epsilon_{\mu e}}$. Therefore, the efficiency ratio can be written as

$$\frac{\epsilon_{\mu\mu}}{\epsilon_{\mu e}} = \frac{\epsilon_{\mu\mu}^C}{\epsilon_{\mu e}^C} \times \frac{\epsilon'_{\mu\mu}}{\epsilon'_{\mu e}} \quad (7.4)$$

where $\frac{\epsilon_{\mu\mu}^C}{\epsilon_{\mu e}^C}$ consists of terms that are expected to cancel and $\frac{\epsilon'_{\mu\mu}}{\epsilon'_{\mu e}}$ represents factors that differ between the two modes. Table 7.1 lists the cuts whose efficiencies either cancel exactly or very nearly cancel, implying $\frac{\epsilon_{\mu\mu}^C}{\epsilon_{\mu e}^C} \approx 1$. Tables 7.2 and 7.3 list the remaining factors that do not cancel and thus affect the branching ratio calculation. One finds that the efficiency ratio for the $\mu\mu$ to μe is

$$\frac{\epsilon_{\mu\mu}}{\epsilon_{\mu e}} \simeq 1.173. \quad (7.5)$$

Cut	Effect on ratio $\frac{\epsilon_{\mu\mu}}{\epsilon_{\mu e}}$
Neutral beam definition cuts	Exact cancellation
Vertex quality cuts	Minor systematic effect
Track quality cuts	Minor systematic effect
Extra track cut	Exact cancellation
TSC Good L0	Minor systematic effect
Spec track time difference cut	Exact cancellation
Fiducial volume cut	Exact cancellation
Parallelism cut	Exact cancellation
Level 1 trigger efficiency	Minor systematic effect
Level 3 trigger efficiency	Minor systematic effect

Table 7.1: Efficiency factors that are expected to cancel or have only a minor effect ($\sim 1-2\%$) on the ratio of $\mu\mu$ to μe .

Cut	Efficiency
FT $\frac{\chi^2}{\nu}$ (single μ)	0.9349 ± 0.0014
FT $\frac{\chi^2}{\nu}$ (single e)	0.9955 ± 0.0003
CER e efficiency (single e)	0.9740 ± 0.0010
PbG e efficiency (single e)	0.9862 ± 0.0007
IMF μ efficiency (single μ)	0.9580 ± 0.0012
PbG e veto on μ side	0.9958 ± 0.0004
Momentum asymmetry	0.9621 ± 0.0013
p_T^{\parallel} cut	0.9495 ± 0.0014
e momentum	0.9993 ± 0.0001
Signal region cuts	0.8885 ± 0.0021
Total efficiency	0.6917 ± 0.0027

Table 7.2: Efficiency factors that only affect $\epsilon_{\mu e}$.

Cut	Efficiency
FT $\frac{\chi^2}{\nu}$ (double μ)	0.8740 ± 0.0020
IMF μ efficiency (double μ)	0.9282 ± 0.0025
μ momentum	0.9999 ± 0.0001
Total efficiency	0.8112 ± 0.0029

Table 7.3: Efficiency factors that only affect $\epsilon_{\mu\mu}$.

7.3 Final Result

Combining the ratios previously calculated gives a final SES of

$$\left(\frac{7.2 \times 10^{-9}}{5470}\right) \left(\frac{0.8112}{0.6917}\right) \left(\frac{0.0285}{0.0234}\right) \left(\frac{1}{0.9533}\right) = 1.97 \times 10^{-12} \quad (7.6)$$

To calculate an upper limit on $B(K_L^0 \rightarrow \mu e)$ at a 90% confidence level, the SES must be multiplied by a factor of 2.3. This results in an upper limit of

$$B(K_L^0 \rightarrow \mu e) < 4.53 \pm 0.21 \pm 0.46 \times 10^{-12} \quad (7.7)$$

where the first error is the calculated statistical error and the second error is the estimated systematic error of 10% for this preliminary result. This value represents a factor of seven improvement in the branching ratio limit over the E791 result.

Chapter 8

Conclusions

8.1 Theoretical Implications and Future Prospects

E871 has not observed a $K_L^0 \rightarrow \mu e$ signal at a SES of 1.97×10^{-12} . Therefore, the preliminary branching ratio limit is $B(K_L^0 \rightarrow \mu e) < 4.53 \pm 0.21 \pm 0.46 \times 10^{-12}$ based on normalizing to $K_L^0 \rightarrow \mu^+ \mu^-$ decays. This result is a factor of seven lower than the previous limit set by E791 and, even though E871 did not achieve the expected factor of 20 improvement, it represents a significantly stricter constraint on any theory that predicts the decay $K_L^0 \rightarrow \mu e$. We can draw several conclusions from this result.

First, the SM, along with most of its minimal extensions, is unaffected by this result. The only SM extension affected is the left-right symmetric model, which now has some of its available parameter space excluded but is still a viable theory. Second, for the non-SM theories, the non-observation of $K_L^0 \rightarrow \mu e$ implies no non-SM processes were observed. The SES reached by E871 places severe restrictions on TC and composite models, requiring these models to invoke other mechanisms, such as slowly varying couplings and generation number conservation, to remain viable. Horizontal and Pati-Salam models can still explain

this result, provided the respective boson masses are greater than 129 TeV and 1400 TeV. The null result from E871 has no effect on SUSY models, as they predict a branching ratio two orders of magnitude lower than this limit.

At this time, there are no near term plans to search for $K_L^0 \rightarrow \mu e$. Indeed, for the E871 apparatus, any further measurements would be background limited by K_{e3} decays at about the level of 10^{-13} . However, the main limitation to this measurement, as well as to any other rare decay experiment, is beam intensity. Even if K_{e3} decays were not an issue, the time required to accumulate enough K_L decays (at present intensities) to have a measurable impact on $B(K_L^0 \rightarrow \mu e)$ is substantial. If that were not enough, the time required to reach the sensitivities required by some models would be absolutely prohibitive. In summary, any future measurements of $K_L^0 \rightarrow \mu e$ must be able to accommodate a substantially higher intensity and must further suppress K_{e3} decays in some manner if they are to have any significant effect on $B(K_L^0 \rightarrow \mu e)$.

8.2 Individual Contributions

This experiment would not have been possible without the combined effort of the members of the E871 collaboration. Just the scale and complexity of the equipment used in high energy physics require that many highly skilled people be involved in the construction, operation and data analysis efforts of these experiments. Because of this, it is possible to lose track of individual contributions to the experimental effort. Therefore, the author will discuss here the contributions that he made to E871.

The E871 construction started in 1992 with the decommissioning of the E791 apparatus and the refurbishing of the MRG planes and electronics. The author tested and repaired the MRG electronics and assisted in the cleaning, repairing and rigging of the MRG planes. When the CER construction got underway, he was the lead graduate student on the

project. He spent much of his time performing quality assurance checks on incoming parts, measuring mirror blank and PMT characteristics and coordinating CER related activities, such as rigging and installing the gas handling system, with BNL personnel. Once the 1995 running period started, he, with a few other William and Mary colleagues, maintained the MRG and CER systems and assisted others with maintenance of their systems when needed. During the straw chamber rebuilding effort, following the vacuum window implosion of June 1995, the author became the on-site manager of the BNL-local straw chamber rebuilding effort as the original experts were drawn into the main rebuilding effort at UCI and UTA. From that point on he was a *de facto* straw expert as well.

The author's contributions to the data analysis effort were the development of the MRG muon identification software, the " $\pi\pi$ scaling" software, which allows the Monte Carlo to sample events properly over both running periods, and participation in the μe analysis effort. In addition, he was an active participant in all three production analysis efforts.

Appendix A

Čerenkov Counter Component Tests

A.1 Mirror Blank Testing

The defining characteristics of the CER mirrors depend on the radius of curvature and the surface roughness. The first category was a concern because of the large distance from the mirror to the PMT. A small deviation in the angle of reflection from the mirror could result in the reflected photon's missing the PMT. Monte Carlo testing determined the allowable range in radius of curvature to be 86 to 106% of the ideal radius (2.9 m). The second category, surface roughness, was a constraint because a high surface roughness results in dispersion of the incoming photons. The original allowable limits on surface roughness came from an SLD experiment [64] which faced similar concerns. They determined that an rms surface roughness of 3 nm was their limit. Further Monte Carlo tests showed that we could accept 10 times that surface roughness and still meet our design goals.

Before being coated, the mirror blanks were tested in the Metrology Optics Lab at

BNL. Initially, the radius of curvature was tested by shining a laser, with an expanding lens attached, at the mirror blank and finding the distance to the minimum reflected spot size. This test was abandoned for a reason given below. The surface roughness test was done on a Micromap Interferometer, which generated a surface plot approximately $2.5 \times 2.5 \text{ cm}^2$ with a feature resolution 10^{-2} nm . This test covered only a small area and therefore five locations were tested on each mirror to get a representative sampling. The interferometer also determined the radius of curvature in that region. Because the interferometer results for the radius of curvature were consistent with the laser test results for the first few mirrors, it was decided to use only the interferometer to test the mirror quality. The results for the mirror blanks used in the CER are listed in Table A.1.

Mirror	Radius (m)	Roughness (nm)	Mirror	Radius(m)	Roughness (nm)
1	2.90	3.64	17	2.85	3.91
2	2.81	4.18	18	2.90	3.68
3	2.60	5.37	19	2.71	4.67
4	2.81	11.8	20	2.72	8.86
5	2.85	3.09	21	2.82	2.78
6	2.82	2.58	22	2.84	3.65
7	2.80	3.76	23	2.74	3.57
8	2.84	3.85	24	2.70	4.77
9	2.77	3.14	25	2.78	3.00
10	2.85	3.51	26	2.80	2.86
11	2.91	4.55	27	2.81	2.76
12	2.80	4.71	28	2.78	4.00
13	2.91	3.95	29	2.82	3.74
14	2.67	4.40	30	2.90	4.43
15	2.67	5.95	31	2.78	4.92
16	2.78	3.98	32	2.85	4.22

Table A.1: CER Mirror Characteristics listed by location.

A.2 PMT Rate Tests

When PMTs are operated in a high rate environment, the voltage differences between the last few stages of the tube can be seriously affected by the high current passing through the dynodes. This phenomenon of “voltage sagging” can lower the gain of the PMT considerably. Because the experiment was designed to run at 15 - 20 Tp per spill (3 to 4 times the E791 intensity), the operating rate for the CER was expected to be high. The PMT base voltage divider network was designed to accommodate high rates.

The initial base design testing was done at William and Mary [65] on three initial base designs. The bases were

1. a slightly modified E791 base (with polarity switched)
2. a base borrowed from the University of Colorado, and
3. a base modified to reduce the gain in the first half of the dynode chain.

The tests were done by pulsing the tube with a light pulser that generated a train of two light pulses 400 ns apart (the transit time in the PMT is approximately 70 ns). The first pulse produced about 6 photoelectrons while the second produced fewer so that the single photoelectron peak was clearly visible. By gating on this second pulse, the systematic effects of the pulser instability are minimized. The frequency of the double pulser was then varied to study the sag as a function of pulse rate.

The first base started to sag at 300 kHz (see Figure A.1). Therefore, performing minor modifications to the existing E791 bases was not a satisfactory solution to the problem. The Colorado base was designed to handle high rates and therefore performed well in this test. However, the average pulse height for this base design was rather low (about a 30 mV single photoelectron at 2900 V). Thus, we could not simply copy this design for our CER

PMT bases. After working with a few other modifications, such as tuning the individual dynode gains, the idea of using Zeners in the final two stages to inhibit voltage sagging was implemented. This design accommodated high rates without much loss of pulse height (see Figure A.2). It was then used to manufacture 32 bases for the CER .

A.3 Magnetic Shield Tests

Because of the distance between the photocathode and the first dynode on the 8854 (approximately 120 mm), it was essential that both the transverse and longitudinal magnetic fields in this region be small (< 0.5 Gauss). Magnetic shield testing was done at William and Mary [66] with both computer simulation and bench testing with a prototype shield arrangement and a Helmholtz coil.

The first tests showed that the prototype shielding scheme was very effective against a field transverse to the axis of the PMT, reducing a 32 Gauss applied field down below 0.2 Gauss. The longitudinal field, however, proved to be more of a problem. Initially, it was thought that a coil around the PMT might be needed to reduce the longitudinal field. To test the effects of the coils, simulation software was written using *Poisson*, a magnetic field simulation library developed at Los Alamos National Laboratory. When simulating the original shielding scheme, it was determined that the field lines had a peak in the photocathode-to-first-dynode region, thus confirming the experimental results. Several simulations with active coils were done and tested with the experimental setup. The results showed that the best location for a coil was around the PMT itself, covering the photocathode/first dynode region. However, the use of an active shielding method was problematic. Not only would it require constant monitoring and adjustment to be effective but it would also present an engineering and safety problem of a coil operating in a hydrogen atmosphere. Replacing the Al PMT adapters with Fe adapters was proposed and further simulation and testing

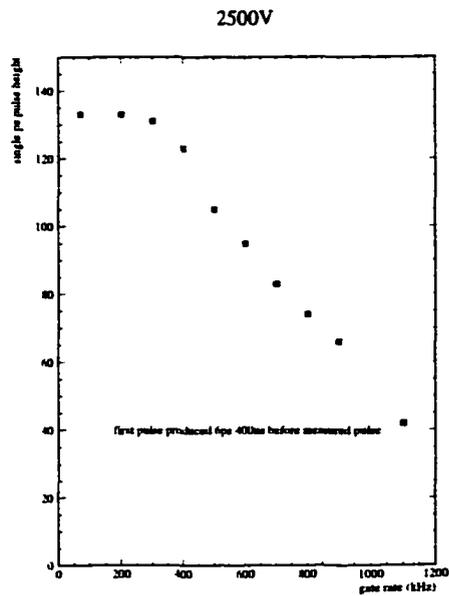


Figure A.1: The results from a rate test on the modified E791 CER base.

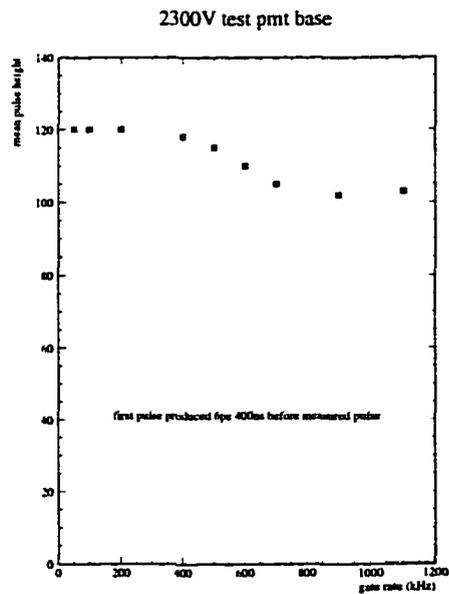


Figure A.2: The results from a rate test on the final CER base design.

was done. In the simulation, the Fe adapter had a large effect, with the magnetic field lines all passing through the Fe, instead of inside the PMT. In bench tests, however, the results, although promising, were not that dramatic, reducing a 6.75 Gauss field (with the Al adapter) to 1.75 Gauss for a 16 Gauss applied field. Several tests with additional Netic and CoNetic magnetic shielding metal were done with some success but the results were still not satisfactory. Finally, another Fe PMT adapter was constructed with a smaller inner diameter. This design reduced the residual field to 1.25 Gauss, which was comparable to the improvement attained by adding Netic and CoNetic material. However, when the external field was reduced to 8 Gauss, the field at the PMT was 0.3 Gauss. The fringe field in the CER region from the 100D40 was measured to be from 7-13 Gauss, but for the most part was below 10 Gauss. Therefore, it was decided that the original shielding with the new Fe PMT adapter was adequate for E871.

Bibliography

- [1] J. Christenson, J. Cronin, V. Fitch, and R. Turley, *Physical Review Letters* **13**, 138 (1964).
- [2] S. Glashow, J. Iliopoulos, and L. Miani, *Physical Review D* **2**, 1285 (1970).
- [3] M. Kobayashi and T. Maskawa, *Prog. Theor. Phys.* **49**, 652 (1973).
- [4] T. Cheng and L. Li, *Physical Review D* **16**, 1425 (1977).
- [5] B. Lee and R. Shrock, *Physical Review D* **16**, 1444 (1977).
- [6] Y. Fukada *et al.*, the Super-Kamiokande Collaboration, hep-ex/9805006 (unpublished).
- [7] P. Langacker, S. Sankar, and K. Schilcher, *Physical Review D* **38**, 2841 (1988).
- [8] A. Acker and S. Pakvasa, *Mod. Phys. Lett. A* **7**, 1219 (1992).
- [9] M. Gell-Mann, P. Ramond, and R. Slansky, in *Supergravity* (North Holland Publ. Co., Amsterdam, 1979), p. 315.
- [10] T. Yanagida, *Prog. Theor. Phys.* **64**, 1103 (1980).
- [11] R. M. Barnett *et al.*, *Physical Review D* **54**, 1 (1996), and 1997 off-year partial update for the 1998 edition available on the PDG WWW pages (URL: <http://pdg.lbl.gov/>).

- [12] A. Barroso, G. Branco, and M. Bento, *Physics Letters B* **134**, 123 (1984).
- [13] W. Hou and A. Soni, *Physical Review Letters* **54**, 2083 (1985).
- [14] W. Hou and A. Soni, *Physical Review D* **35**, 2776 (1987).
- [15] Z. Gagyi-Palffy, A. Pilaftsis, and K. Schilcher, *Nuclear Physics B* **513**, 517 (1998).
- [16] B. McWilliams and L. Li, *Nuclear Physics B* **179**, 62 (1981).
- [17] G. Mainland and K. Tanaka, *Physical Review D* **20**, 1241 (1979).
- [18] J. L. Ritchie and S. G. Wojcicki, *Rev. Mod. Phys.* **65**, 1149 (1993).
- [19] G. L. Kane and R. Thun, *Physics Letters B* **94**, 513 (1980).
- [20] R. Cahn and H. Harari, *Nuclear Physics B* **176**, 135 (1980).
- [21] O. Shanker, *Nuclear Physics B* **185**, 382 (1981).
- [22] J. Pati and A. Salam, *Physical Review D* **10**, 275 (1974).
- [23] N. Deshpande and R. Johnson, *Physical Review D* **27**, 1193 (1983).
- [24] R. Foot, *Physics Letters B* **420**, 333 (1998).
- [25] S. Dimopoulos and L. Suskind, *Nuclear Physics B* **155**, 237 (1979).
- [26] E. Eichten and K. Lane, *Physics Letters B* **90**, 125 (1980).
- [27] E. Eichten, I. Hinchliffe, K. Lane, and C. Quigg, *Physical Review D* **34**, 1547 (1986).
- [28] S. Dimopoulos and J. Ellis, *Nuclear Physics B* **182**, 505 (1981).
- [29] S. Dimopoulos, S. Raby, and G. Kane, *Nuclear Physics B* **182**, 77 (1981).
- [30] B. Holdom, *Physics Letters B* **143**, 227 (1984).

- [31] S. King, *Physics Letters B* **195**, 66 (1987).
- [32] T. Akagi *et al.*, *Physical Review D* **51**, 2061 (1995).
- [33] T. Appelquist, D. Karabali, and L. Wijewardhana, *Physical Review Letters* **57**, 957 (1986).
- [34] T. Appelquist and L. Wijewardhana, *Physical Review D* **36**, 568 (1987).
- [35] S. King and D. Ross, *Physics Letters B* **228**, 363 (1989).
- [36] S. King, *Nuclear Physics B* **320**, 487 (1989).
- [37] J. Pati, *Physics Letters B* **59**, 265 (1975).
- [38] J. Pati, *Physical Review D* **30**, 1144 (1984).
- [39] J. Pati and H. Stremnitzer, *Physics Letters B* **172**, 441 (1986).
- [40] O. Greenberg, R. Mohapatra, and S. Nussinov, *Physics Letters B* **148**, 465 (1984).
- [41] B. Mukhopadhyaya and A. Raychaudhri, *Physical Review D* **42**, 3215 (1990).
- [42] P. Gill, private communication.
- [43] K. Hartman *et al.*, E871 internal note KL-488 (unpublished).
- [44] S. Graessle *et al.*, in *Proc. Of The 7th Vienna Wire Chamber Conference* (Vienna, Austria, 1995).
- [45] S. Worm, Ph.D. thesis. The University of Texas at Austin, 1995.
- [46] *8854 Photomultiplier*, Burle Electron Tubes, datasheet for the Burle 8854 Photomultiplier Tube.

- [47] R. Cousins *et al.*, Nuclear Instruments and Methods in Physics Research **A277**, 517 (1989).
- [48] R. Martin, A. Hancock, D. Connor, and P. Rubin, in *Proc. Conf. on Computing in High Energy Physics* (San Francisco, CA, 1994).
- [49] J. W. Belz, Ph.D. thesis, Temple University, 1993.
- [50] M. T. Witkowski, Ph.D. thesis, The College of William and Mary, 1993.
- [51] D. Wagner, E791/E871 internal note KL-173 (unpublished).
- [52] P. Mèlèse, E791/E871 internal note KL-145 (unpublished).
- [53] R. Cousins and T. Kaarsberg, E791/E871 internal note KL-248 (unpublished).
- [54] W. Molzon and S. Imlay, E791/E871 internal note KL-237 (unpublished).
- [55] R. Lee, E871 Internal Note Kl-472 (unpublished).
- [56] R. D. Martin, Ph.D. thesis, The College of William and Mary, 1998.
- [57] A. Milder, E871 Internal Note (unpublished).
- [58] S. Graessle, E871 internal note KL-499 (unpublished).
- [59] D. Connor and W. Molzon, E871 Internal Note Kl-390 (unpublished).
- [60] R. Lee, private communication.
- [61] S. Graessle, private communication.
- [62] R. Cousins and D. Murphy, E791/E871 Internal Note Kl-005 (unpublished).
- [63] L. Bergström, Z. Phys. C. – Particles and Fields **20**, 135 (1983).

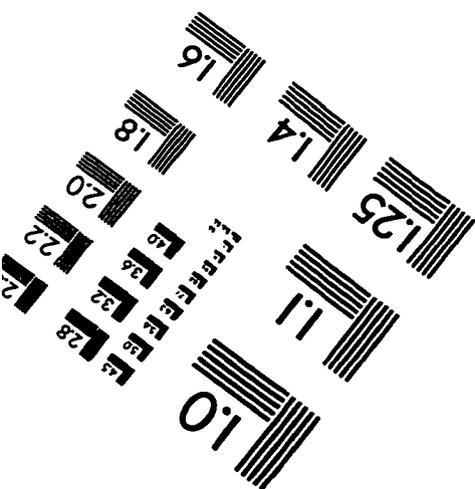
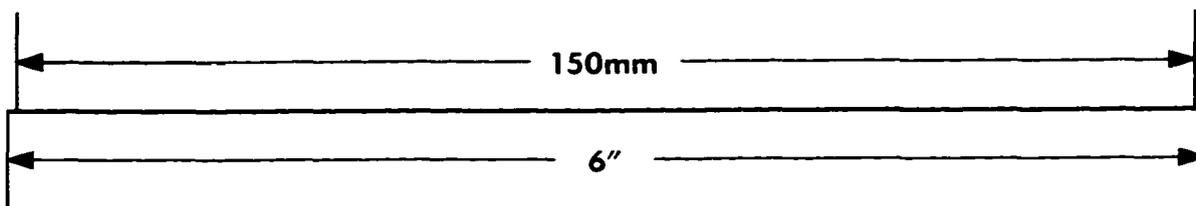
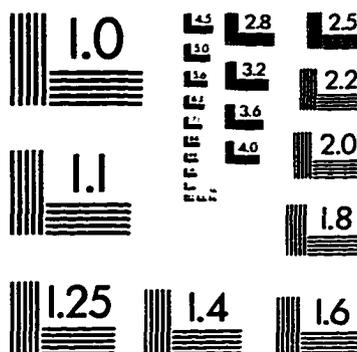
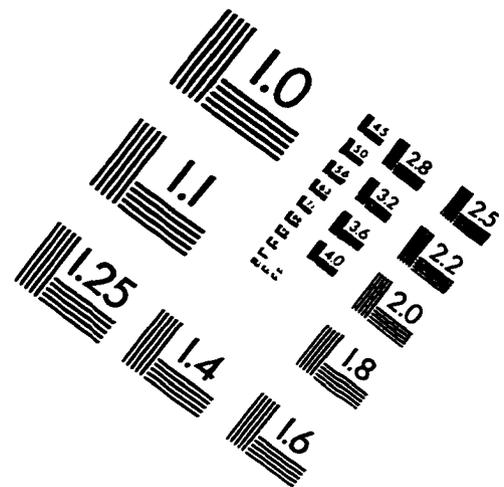
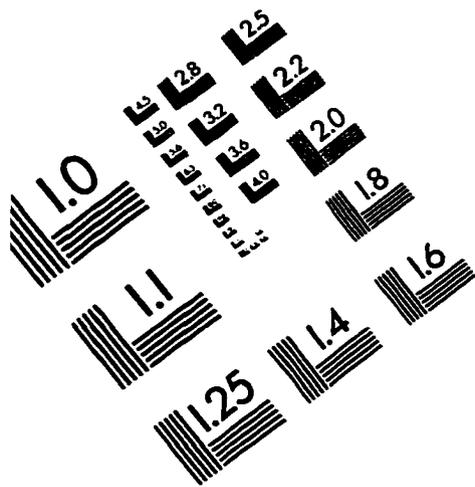
- [64] K. Abe *et al.*, Nuclear Instruments and Methods in Physics Research **300**, 501 (1991).
- [65] Y. Kuang, private communication.
- [66] W. Hall, William and Mary Senior Honors Thesis, 1993 (unpublished).

Vita

Christopher Henry Hoff

Born on May 14, 1967 in Greensboro, North Carolina. Graduated from First Colonial High School in Virginia Beach, Virginia in May 1985. Entered Virginia Polytechnic Institute in Blacksburg, Virginia in September 1985. After taking seven months for active duty with the US Navy Reserves from June 1986 - January 1987, graduated with a B.S. in Physics in May 1990. Entered The College of William and Mary graduate program in June 1990 and started in the Ph.D. program in January of 1992. Currently working as a programmer/analyst at 3G-International in Williamsburg, Virginia.

IMAGE EVALUATION TEST TARGET (QA-3)



APPLIED IMAGE . Inc
 1653 East Main Street
 Rochester, NY 14609 USA
 Phone: 716/482-0300
 Fax: 716/288-5989

© 1993, Applied Image, Inc., All Rights Reserved

