THE UNIVERSITY OF CHICAGO
THE ENRICO FERMI INSTITUTE

A MEASUREMENT OF THE BRANCHING RATIO OF $\pi^0 \to e^+e^-$ USING $K_L \to 3\pi^0$ DECAYS IN FLIGHT

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
ERIC D. ZIMMERMAN

CHICAGO, ILLINOIS
MARCH 1999

ACKNOWLEDGMENTS

This thesis could not have been produced without the heroic efforts of many people. An incomplete list follows.

Five and a half years ago, I came to the University of Chicago and began my association with Bruce Winstein and the neutral kaon group. I was new to the entire field of high energy physics. Soon, I began working with Yau Wah and learning about the wonders of rare decays and transition radiation. At that time, I also benefited greatly from my discussions and work with Roland Winston, Nick Solomey, and Pat Krolak at Chicago, as well as with Erik Ramberg and John Krider at Fermilab. During this period, I also worked with Dick Armstrong and Gene Beck, neither of whom lived to see the results of the experiment. Chicago and Fermilab miss their presence and expertise.

In the autumn of 1994, I began working on the trigger systems; this became a central focus of my involvement in the experiment. Rick Kessler, my primary collaborator on this project from the beginning, taught me a lot of what I know about assembling a large system of electronics. All of us on the experiment thank him for introducing us to the HIGZ donkey (KTASS).

The previous generation of Chicago kaon students, Debbie Harris, Kevin McFarland, Roy Briere, and Bernhard Schwingenheuer, were most helpful in bringing me up to speed in my early years here, and in giving me an idea of the flavor of experimental kaon physics.

None of the results from E799 or E832 could have come about without the efforts of the current Chicago HEP group: Peter Shawhan, Colin Bown, Val "I'm Stuck!" Prasad, Steve Bright, Greg Graham, Ed Blucher, Rick Kessler, Changqing Qiao, Breese Quinn, Jim Graham, Aaron Roodman, and Elliott Cheu in his years here, as well as Marty Dippel, our system administrator. They have all made Chicago an excellent environment for working and learning.

The gang at the Muon Lab provided intellectual discussion, help with physics and detector problems, and random diversions. These ranged from barbecues to volleyball to comet-watching on freezing nights in the middle of nowhere. Rick Tesarek was an instigator or collaborator on nearly all of them. Rick, Amit Lath, Eva Halkiadakis, Suzanne Averitte, John Belz, Kazu Hanagaki, Peter Mikelsons, Sydney Taegar, Theo Alexopoulos, Ashkan Alavi-Harati, and many others there were important contributors to the experiment and became good friends of mine.

Tsuyoshi Nakaya was the front-end readout expert, and his mastery of all things DYC (Sec. 2.7.4) was accompanied by a remarkable physics acumen. I have also been fortunate to work with the many other people who helped make the experiment a success, including (by no means exhaustively) Viv O'Dell, Hans Kobrak, Herman White, Rick Coleman, John Jennings, Steve Field, Brad Cox, Mike Arenton, Peter Shanahan, Ram Ben-David, Leo Bellantoni, Tony Barker, Albert Erwin, Sunil Somalwar, Katsushi Arisaka, Steve Schnetzer, Bill Slater, Marj Corcoran, Ron Ray, Hogan Nguyen, Bob Tschirhart, Greg Bock, Julie Whitmore, Bob Hsiung.

All of the accelerator staff who kept the Tevatron running through many days and nights (and too many controlled accesses) deserve our thanks.

It has been a great privilege to work closely with Taku Yamanaka, one of the kindest, cleverest, and most dedicated physicists I know. It's difficult to describe the intensity, excitement, and constant stress of commissioning, tuning, and running a kaon experiment. Taku kept it all in perspective while spending long nights at the counting room helping us shape the configuration and maximize the physics reach of E799.

Rick Ford, Doug Jensen, Jim Volk, Taku, and Dave Pushka each played critical roles in getting a magnet designed and built on a ridiculous schedule and near-zero budget. We overcame snakes, flying ferrite, and the skepticism of our colleagues, and somehow pulled it off (see Chapter 7 and Refs. [65, 67]). Sometimes a sledgehammer *is* the right tool for the job.

Bruce Winstein was the first person with whom I came to work at Chicago. His experimental expertise and his focus on the physics enrich all our work. Perhaps more than anything else, I appreciate his ability to ask the right questions of us.

Thanks to Wah for being my advisor for these years, for encouraging me to work on critical parts of the experiment, giving me the independence to pursue them in my own way, and for standing up for me when I ventured into controversy. Wah convinced me that the π^0, which I once considered an eminently dull particle, still held interesting secrets. Now it holds one fewer.

Finally, I thank my family for their support, over all these years, of all my pursuits—even my more whimsical ones (such as a career in physics).

ERIC D. ZIMMERMAN
Chicago, Illinois
December 1998

TABLE OF CONTENTS

LIST OF TABLES

LIST OF ILLUSTRATIONS

ABSTRACT

The branching ratio of the rare decay $\pi^0 \to e^+e^-$ has been measured in E799-II, a rare kaon decay experiment at Fermilab. The π^0's were produced by $K_L \to 3\pi^0$ decays in flight, and the entire K_L decay was reconstructed. We observed 275 candidate $\pi^0 \to e^+e^-$ events, with an expected background of 21.4 ± 6.2 events which included the contribution from Dalitz decays (tree-level $\pi^0 \to e^+e^-\gamma$). We measured

$$\mathrm{BR}\left(\pi^0 \to e^+e^-, \left(\frac{m_{e^+e^-}}{m_{\pi^0}}\right)^2 > 0.95\right) = (6.09 \pm 0.40 \pm 0.24) \times 10^{-8},$$

where the first error is statistical and the second systematic. This result is in agreement with expectations from recent vector meson dominance and chiral perturbation theory models.

CHAPTER 1

THE DECAY $\pi^0 \to e^+e^-$

"...a decay dear to my childhood heart..."

—C. Quigg

Conference summary, Workshop on Heavy Quarks at Fixed Target

Fermilab, October 1998

Neutral pseudoscalar meson decays into lepton pairs have been a rich subject of study over the past four decades. In the Standard Model, there is no light, spinless propagator which can couple to both quarks and leptons; this forbids these decays to occur at tree level. Because one-loop decays are suppressed by two orders of a coupling constant, the dilepton modes are expected to be relatively rare. Depending upon the masses and flavors of the leptons and meson involved, other suppression factors may be present as well, and decays to two leptons can be extremely rare. Rare decays such as these are considered promising avenues to search for new physics, as exotic processes may contribute to these decays at rates significantly above the Standard Model expectations. Within the Standard Model, these decays are a useful tool for constraining form factor models in kinematic regions which cannot be explored in tree-level decays (in particular, momentum transfer greater than the parent meson mass).

The decay $\pi^0 \to e^+e^-$ has received much experimental and theoretical attention since its branching ratio was first calculated by Drell in 1959 [1]. It is expected, within the Standard Model, to occur at a rate less than 10^{-7} that of $\pi^0 \to \gamma\gamma$, the dominant decay

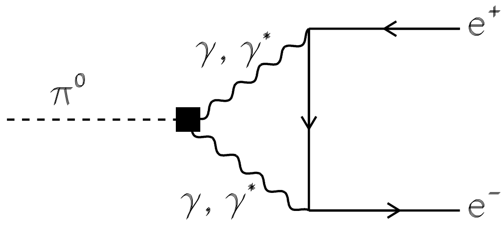

Figure 1. Feynman diagram of the Standard Model electromagnetic contribution to $\pi^0 \to e^+e^-$.

mode of the π^0. Relative to $\pi^0 \to \gamma\gamma$, $\pi^0 \to e^+e^-$ is suppressed by a helicity factor as well as by two orders of the electromagnetic coupling constant α_{EM}. Earlier experiments have produced conflicting measurements of the branching ratio for this mode; this thesis presents a new, precision measurement of BR($\pi^0 \to e^+e^-$).

1.1 Helicity suppression and the unitarity bound

The π^0 is spinless and has negative parity. Because of the pure vector coupling of the electromagnetic interaction, parity is conserved in the decay and helicity is conserved at each vertex. If the electron were massless, the decay $\pi^0 \to e^+e^-$ could not proceed without violating either parity or angular momentum conservation (Fig. 2). The nonzero mass of the electron allows helicity nonconservation to occur; the decay is therefore not forbidden, but suppressed by a factor of $(2m_e/m_{\pi^0})^2$.

1.1.1 QED calculation of $\pi^0 \to e^+e^-$

Assuming that the decay $\pi^0 \to e^+e^-$ proceeds entirely through a two-photon intermediate state, the branching can be calculated in terms of the $\pi^0\gamma^*\gamma^*$ form factor [2]:

$$\frac{\Gamma(\pi^0 \to \gamma^*\gamma^* \to e^+e^-)}{\Gamma(\pi^0 \to \gamma\gamma)} = 2\sqrt{1-\left(\frac{2m_e}{m_{\pi^0}}\right)^2}\left(\frac{\alpha_{\mathrm{EM}}}{\pi}\frac{m_e}{m_{\pi^0}}\right)^2 |R|^2 \tag{1.1}$$

where

$$R \equiv -\frac{2i}{(\pi m_{\pi^0})^2}\int d^4k \frac{q^2k^2-(q\cdot k)^2F(k^2,(q-k)^2;m_{\pi^0}^2)}{[k^2+i\varepsilon][(q-k)^2+i\varepsilon][(k-p)^2-m_e^2+i\varepsilon]} \tag{1.2}$$

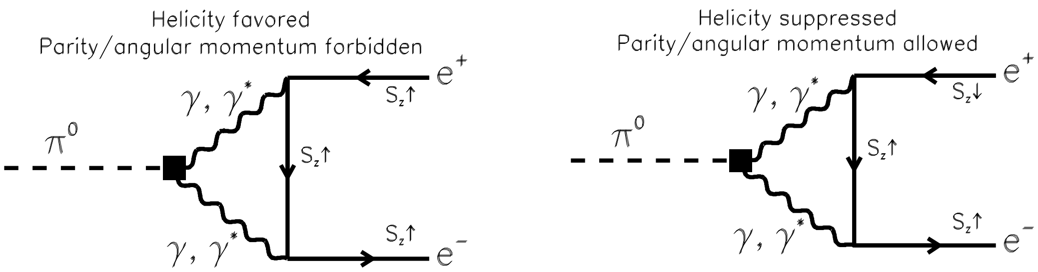

Figure 2. Helicity suppression of the decay $\pi^0 \to e^+e^-$.

and q and p are the four-momenta of the initial π^0 and one outgoing electron, and F is the $\pi^0\gamma^*\gamma^*$ form factor normalized to $F(0,0;m^2_{\pi^0}) = 1$.

In the limit where the form factor $F(k_1^2, k_2^2; m^2_{\pi^0})$ is zero for all nonzero k_1^2 and k_2^2, R reduces to its imaginary part. This corresponds to the contribution from on-shell internal photons, and represents the *unitarity bound* for the decay. Given that $|R|^2 \geq (\mathrm{Im}(R)^2)$, we find that

$$\frac{\Gamma(\pi^0 \to \gamma^*\gamma^* \to e^+e^-)}{\Gamma(\pi^0 \to \gamma\gamma)} \geq 4.75 \times 10^{-8}. \tag{1.3}$$

1.2 Models of the $\pi^0\gamma^*\gamma^*$ form factor and $\pi^0 \to e^+e^-$

In order to predict the rate for $\pi^0 \to e^+e^-$ above the unitarity bound (equivalently, the real part of R), some model of the $\pi^0\gamma^*\gamma^*$ form factor is necessary. Indeed, the assumption of a constant form factor (quite successful in predicting the rate and spectrum of the Dalitz decay $\pi^0 \to e^+e^-\gamma$ and double Dalitz decay $\pi^0 \to e^+e^-e^+e^-$ [3] [4]) is completely unusable in this case, because it causes the rate for $\pi^0 \to e^+e^-$ to diverge logarithmically. In order to produce a sensible result for $\pi^0 \to e^+e^-$, any model must have a form factor "cutoff," a value of k^2 (equivalently, the virtual photon mass) above which F must vanish. Drell interpreted this cutoff as corresponding to a spatial "size" of the π^0.

Various models have been used over the years to predict the shape of the $\pi^0\gamma^*\gamma^*$ form factor and the real part of the amplitude for $\pi^0 \to e^+e^-$. It is expected that many different form factor models produce similar predictions for $\pi^0 \to e^+e^-$, as the decay amplitude is expected to be dominated by its model-independent imaginary part (the contribution from on-shell photons), and the excess depends primarily on the form factor's cutoff and not on the details of its functional form. In particular, Vector Meson Dominance (VMD) models have been shown [2] to be equivalent to Quark Triangle Loop (QTL) calculations such as those in Refs. [5] and [6] with appropriate choices of the quark masses ($m_q \sim 300$ MeV/c^2). Even baryon loop models ([6] and [7]) can be made to produce similar predictions. Therefore, of the form factor models, only a recent VMD result will be considered here in detail. Other calculations prior to 1993 have been thoroughly catalogued in [8].

VMD models have been successful at predicting form factors in other meson decays [9]. In the π^0 system, the assumption of a vector meson dominated form factor has successfully predicted the first-order slope of the form factor $F(k_1^2 < 0, k_2^2 \sim 0; m_{\pi^0}^2)$ as measured in the $e^+e^- \to \pi^0 e^+e^-$ reaction [10]. These predictions are also tested using Dalitz decays, but the precision is low and the results are not conclusive [8] [11] [12]. A 1992 pure VMD calculation by Ametller *et al.* [9] predicts $BR(\pi^0 \to e^+e^-) = (6.33 \pm 0.19) \times 10^{-8}$, with the error dominated by the difference between the ρ and ω masses.

In Chiral Perturbation Theory (χPT) calculations of $\pi^0 \to e^+e^-$, an explicit form factor model is not used; rather, the divergent loop is canceled by an effective pointlike $\pi^0\gamma\gamma$ coupling. This coupling has been taken from the measured $\eta \to \mu^+\mu^-$ branching ratio and used to obtain the prediction $BR(\pi^0 \to e^+e^-) = (7 \pm 1) \times 10^{-8}$ [13]. A more recent $U(3)_L \otimes U(3)_R$ effective theory calculation [14], which also used $\eta \to \mu^+\mu^-$ to fix the local term, predicted $BR(\pi^0 \to e^+e^-) = (8.3 \pm 0.4) \times 10^{-8}$.

1.3 Exotic contributions to $\pi^0 \to e^+e^-$

Although the Standard Model predicts $\pi^0 \to e^+e^-$ to be completely dominated by the two-photon intermediate state, it has been suggested that other, non-Standard Model processes could enhance the decay rate. The observed low $\pi^0 \to e^+e^-$ branching ratio has been used to exclude models involving light axions [15]. Certain leptoquark models [16] predict smaller enhancements to $\pi^0 \to e^+e^-$, but most of these theories are more tightly constrained by experimental limits on $K_L \to \mu^\pm e^\mp$ [17] and muon conversion experiments.

1.4 Other pseudoscalar decays to two leptons

Besides $\pi^0 \to e^+e^-$, the only observed pseudoscalar decays to two leptons are $\eta \to \mu^+\mu^-$ [18], $K_L \to \mu^+\mu^-$ [19], and (recently) $K_L \to e^+e^-$ [20]. None of the current measurements (or limits on currently unobserved modes) shows a rate below the unitarity bound, although an early search for $K_L \to \mu^+\mu^-$ [21] did indicate such a problem.

Recent papers [22] [14] have described calculations of the rate for these and other similar decays in χPT. These models generally use the measured $\eta \to \mu^+\mu^-$ branching ratio as input, and make predictions for the $\pi^0 \to e^+e^-$, $\eta \to e^+e^-$, $K_L \to \mu^+\mu^-$, and $K_L \to e^+e^-$ decay rates.

It is hoped that, eventually, the long-distance contribution to $K_L \to \mu^+\mu^-$ can be understood accurately; any excess from short-distance physics could then be measured. Short-distance physics does not make any significant contribution to decays of the π^0 and η, because these particles can decay electromagnetically. The K_L, however, can only decay through a strangeness-changing interaction: the partial decay width $\Gamma(K_L \to \gamma\gamma)$ is smaller than $\Gamma(\pi^0 \to \gamma\gamma)$ by twelve orders of magnitude. Therefore, the contribution to $K_L \to \mu^+\mu^-$ from a WW intermediate state may be observable as an excess above the rate from the $\gamma\gamma$ intermediate state.

The short distance contribution to $K_L \to \mu^+\mu^-$ is expected to be dominated by a loop containing a top quark, and could provide a measurement of the CKM matrix element V_{td} [23] [24]. At present, the theoretical models of the long-distance contribution are not of sufficient precision to allow such a measurement [22].

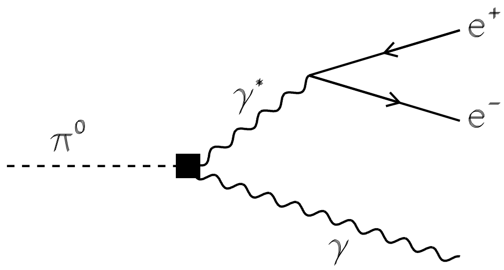

Figure 3. Feynman diagram of the Dalitz decay, $\pi^0 \to e^+e^-\gamma$.

1.5 Radiative corrections to $\pi^0 \to e^+e^-$

Final-state radiative effects introduce a complication for any precise measurement of $\pi^0 \to e^+e^-$. These radiative corrections can cause the total rate for the decay to be different from Eq. 1.1. More problematic, inner bremsstrahlung can cause the $\pi^0 \to e^+e^-$ decay to produce a $e^+e^-\gamma$ final state, with $m_{e^+e^-} < m_{\pi^0}$. Understanding this correction was critical to measurement of $\pi^0 \to e^+e^-$, as we could only accept events inside a small window $m_{e^+e^-} \approx m_{\pi^0}$ to reject background from the copious π^0 Dalitz decay (tree-level $\pi^0 \to e^+e^-\gamma$, see Fig. 3). For the discussion of radiative corrections, we introduce the kinematic variable $x \equiv (m_{e^+e^-}/m_{\pi^0})^2$.

For this measurement, the calculation by Bergström [25] was used to interpret radiative effects. In this paper, which was also used in Refs. [8] and [26] to take account of radiative corrections, a pointlike $\pi^0 ee$ vertex (a γ_5 coupling independent of $m_{e^+e^-}$) was assumed, and first-order radiative effects were calculated (Fig. 4). The pointlike pseudoscalar coupling was argued to be a reasonable approximation for events with high x, which dominate the radiative $\pi^0 \to e^+e^-$ spectrum. This model ignored radiation from the internal electron; this was justified by noting that this electron is far off mass-shell, in a kinematic region where bremsstrahlung is not significant.

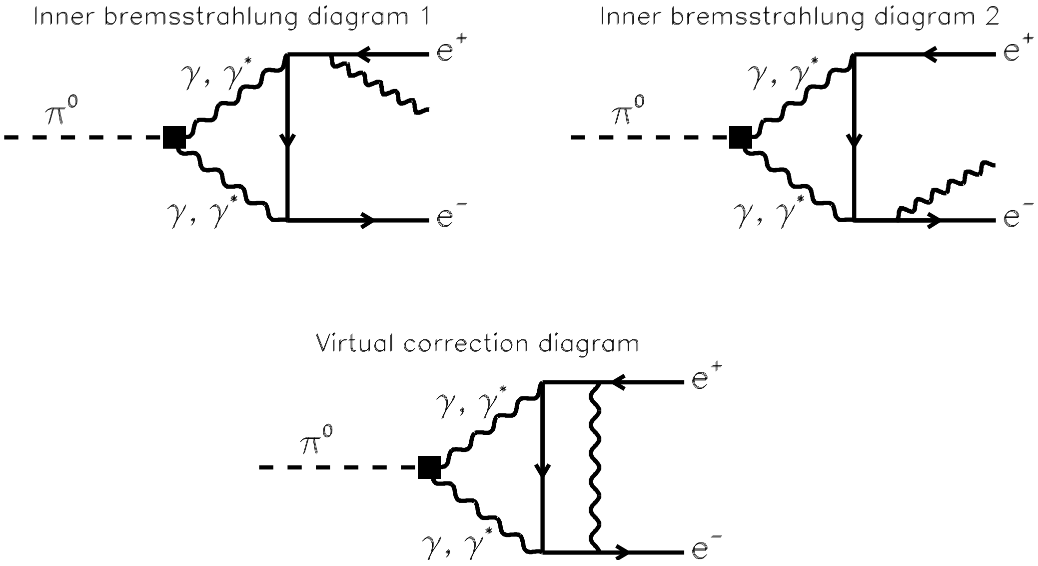

Figure 4. First-order final-state radiation diagrams for $\pi^0 \to e^+e^-$.

1.5.1 Total rate correction

The correction to the total $\pi^0 \to e^+e^-$ rate, from interference between the virtual correction diagram in Fig. 4 and the lowest order $\pi^0 \to e^+e^-$ diagram (Fig. 1) was calculated to be (in Bergström's notation):

$$\frac{\Gamma^{\mathrm{brems}} + \Gamma^{\mathrm{virt}}}{\Gamma^0_{e^+e^-}} = \frac{\alpha_{\mathrm{EM}}}{\pi}\left[\frac{3}{2}\ln\left(\frac{1-v_0}{1+v_0}\right) + \frac{9}{4} + \mathcal{O}\left(\frac{m_e^2}{m_{\pi^0}^2}\right)\right] \tag{1.4}$$

where $\Gamma^{\mathrm{brems}} + \Gamma^{\mathrm{virt}}$ is the total $\pi^0 \to e^+e^-$ rate including inner bramsstrahlung events, $\Gamma^0_{e^+e^-}$ is the rate from the lowest order diagram in the absence of radiative corrections, and $v_0 \equiv \sqrt{1 - 4m_e^2/m_{\pi^0}^2} = 0.99997134$. The modification to the total $\pi^0 \to e^+e^-$ rate is -3.4%.

1.5.2 Inner bremsstrahlung spectrum

Although the correction to the total rate is fairly small, the rate thus obtained is not experimentally measurable because it includes events with a bremsstrahlung photon of significant energy, and therefore $m_{e^+e^-}$ significantly below the π^0 mass. These events

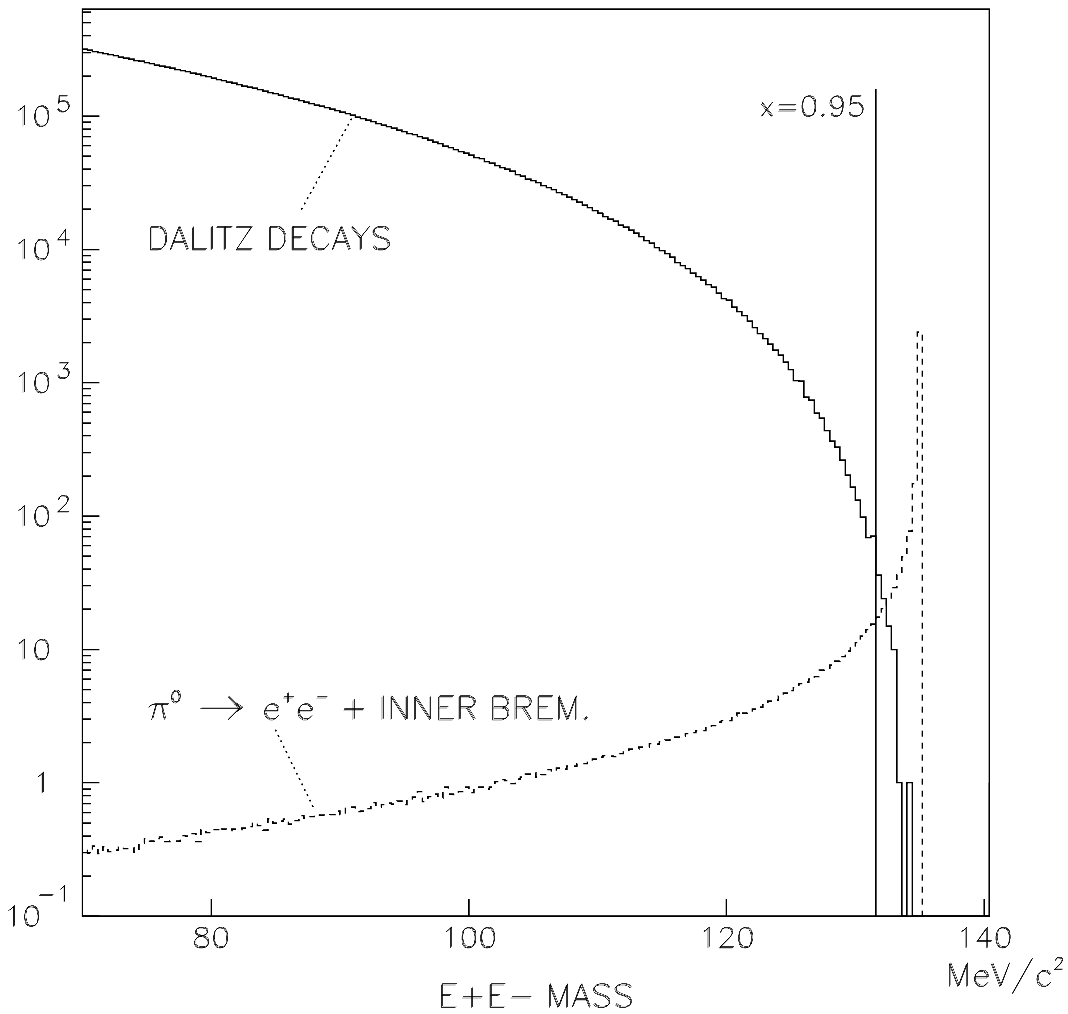

Figure 5. $m_{e^+e^-}$ distribution for Dalitz decays (tree-level $\pi^0 \to e^+e^-\gamma$) and for $\pi^0 \to e^+e^-$ with internal bremsstrahlung correction. Lowest-order $\pi^0 \to e^+e^-$ branching ratio is assumed to be 7×10^{-8}.

are indistinguishable from the much more common Dalitz decays, and thus cannot be observed separately. Indeed, the two processes can actually interfere with each other, although they populate somewhat different regions of phase space and so the magnitude of the interference is expected to be very small at any value of x. The interference term could contribute only a small fraction of one per cent of the decay rate in the region $x > 0.8$ [25].

The inner bremsstrahlung spectrum is given by the differential rate as a function of x:

$$\frac{d\Gamma^{\text{brems}}}{dx} = \Gamma^0_{e^+e^-} \frac{\alpha_{EM}}{\pi} \left[\frac{x^2+1}{1-x} \ln\left(\frac{1-v}{1+v}\right) - \frac{2xv}{1-x} \right] + \mathcal{O}(\alpha^2_{\text{EM}}) \tag{1.5}$$

where $v \equiv \sqrt{1 - 4m_e^2/xm_{\pi^0}^2}$. This expression was integrated up to a cutoff value of x to determine the fraction of decays lost if an experiment is only sensitive to events above the cutoff. It has been noted [27] that this formula may become invalid for cutoffs above $x \sim 0.98$, where higher-order corrections are needed to describe the bremsstrahlung spectrum. However, this corresponds to a cutoff above $m_{e^+e^-}/m_{\pi^0} \sim 0.99$, and our experimental resolution is not sufficient to use a cutoff this high.

For a cutoff value of $x = 0.95$, 89.5% of the bremsstrahlung spectrum is above the cut. This cutoff was used in a previous measurement of $\pi^0 \to e^+e^-$ (Refs. [26] and [8]) which also used Bergström's calculation to interpret radiative effects. Following the same convention, we will quote the branching ratio for events with $x > 0.95$ after subtracting the contribution from tree-level $\pi^0 \to e^+e^-\gamma$ decays; this branching ratio is 13.5% below what the lowest-order rate would be in the absence of radiative corrections. The branching ratio must be scaled upward to compare with theoretical models of $\pi^0 \to e^+e^-$ which neglect final-state radiation.

1.6 Previous measurements of $\pi^0 \to e^+e^-$

The earliest measurements of $\pi^0 \to e^+e^-$ were reported by a Geneva-Saclay group in 1978 [28] and by a LAMPF group in 1983 [29]. The Geneva-Saclay group used π^0's produced by the decay $K^+ \to \pi^+\pi^0$ in flight; the LAMPF experiment used the

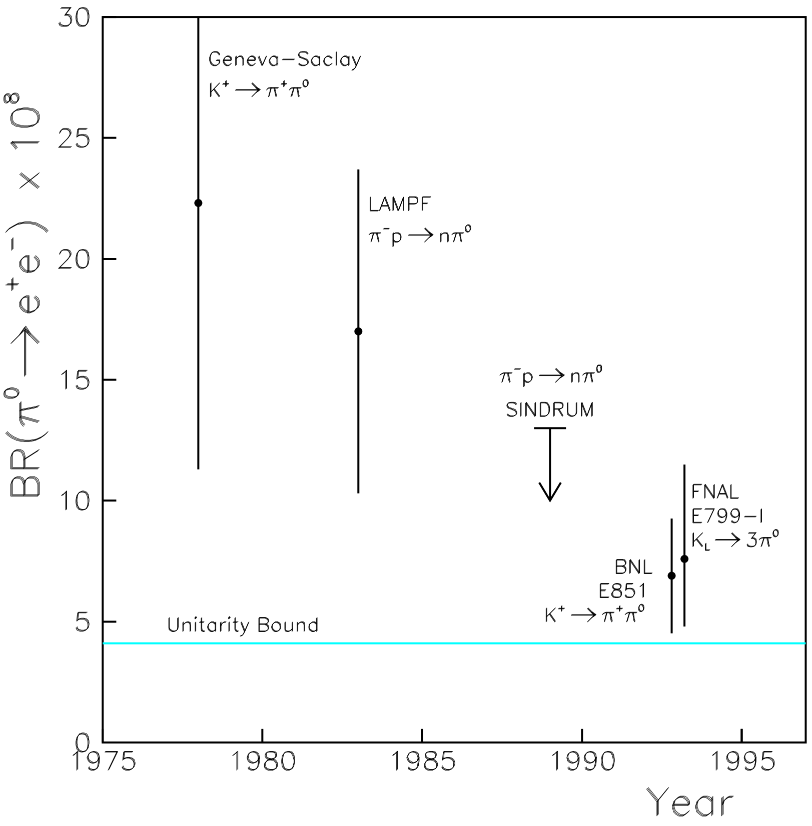

Figure 6. Previous measurements of BR($\pi^0 \to e^+e^-$). A 1987 limit from the OMICRON collaboration is above the vertical scale of the figure.

charge-exchange process $\pi^-p \to \pi^0 n$ from a 300 MeV/c pion beam. Both experiments favored a branching ratio of $\sim 2\times10^{-7}$, which would be very hard to accomodate within the Standard Model. A 1989 search by the SINDRUM collaboration [31], using the charge-exchange reaction with stopped pions, produced a 90% confidence level upper limit of 1.3×10^{-7}, which excluded the central values of both previous measurements. More recently, measurements by BNL E851 [32] and FNAL E799-I [26] have observed the decay at the 5 to 10 $\times10^{-8}$ level, near the Standard Model expectation. The BNL measurement used the $K^+ \to \pi^+\pi^0$ decay in flight technique, while the FNAL experiment used $K_L \to 3\pi^0$ decays in flight as a source of neutral pions. The existing measurements of $\pi^0 \to e^+e^-$ are summarized in Fig. 6 and Table 1.

Collaboration	Technique	Pub. Year	BR($\pi^0 \to e^+e^-$) units of 10^{-8}	Ref.
Geneva-Saclay	$K^+ \to \pi^+\pi^0$ in flight	1978	22.3^{+24}_{-11}	[28]
LAMPF	0.3 GeV $\pi^- p \to \pi^0 n$	1983	$17 \pm 6 \pm 3$	[29]
OMICRON	$\pi^- p \to \pi^0 n$ in flight	1987	< 53 (90% C.L.)	[30]
SINDRUM	Stopped $\pi^- p \to \pi^0 n$	1989	< 13 (90% C.L.)	[31]
BNL E851	$K^+ \to \pi^+\pi^0$ in flight	1993	$6.9 \pm 2.3 \pm 0.6$	[32]
FNAL E799-I	$K_L \to 3\pi^0$ in flight	1993	$7.6^{+3.9}_{-2.8} \pm 0.5$	[26]

Table 1. Previous measurements of BR($\pi^0 \to e^+e^-$). The FNAL E799-I measurement accounts for radiative corrections, and the quoted branching ratio is for $(m_{e^+e^-}/m_{\pi^0})^2 > 0.95$, with the contribution from Dalitz decays subtracted.

1.7 The E799-II measurement technique

The E799-II measurement described in this thesis used the same basic technique as the 1993 E799-I measurement. A high-energy neutral kaon beam was used as a source of tagged π^0's using the $K_L \to 3\pi^0$ decay, where two π^0's decayed to $\gamma\gamma$. The entire K_L decay was reconstructed. This technique has several major advantages over those described in the previous section:

- The $K_L \to 3\pi^0$ decay has a large branching ratio (21.1% [33]). As three π^0's were available in each decay and the $\pi^0 \to \gamma\gamma$ branching ratio is 98.8%, this provided 0.62 usable π^0's per decayed K_L.

- With good electromagnetic calorimetry, the full $K_L \to 3\pi^0$ reconstruction provided redundant kinematic constraints, eliminating all backgrounds from interactions or non-$K_L \to 3\pi^0$ kaon decays.

- The chief backgrounds in searches via the charge-exchange interaction $\pi^- p \to \pi^0 n$ are continuum e^+e^- production and decay product scattering or conversion in the target. These types of background were completely eliminated here because the π^0's were produced and decayed in vacuum.

- No common K_L decays have two electrons in the final state. The continuum process $K_L \to \pi^0\pi^0 e^+e^-$ is expected to have an extremely low branching ratio, and was not a significant background to $K_L \to 3\pi^0$, $\pi^0 \to e^+e^-$. This was in marked contrast to K^+ experiments, where the continuum $K^+ \to \pi^+e^+e^-$ decay is a large background.

CHAPTER 2

THE K_L BEAM AND THE E799-II DETECTOR

E799-II, a fixed-target rare K_L decay experiment, ran in the NM (Neutrino Muon) beamline at Fermilab during two periods in 1997. The two data sets were called the "winter run" and the "summer run."

The basic elements of the experiment were two neutral kaon beams, a long vacuum decay region, and a detector which consisted of a charged spectrometer and an electromagnetic calorimeter. Photon veto systems and additional particle identification detectors were also present.

2.1 The K_L beams

The neutral beams used in the experiment were produced by a primary proton beam which struck a target. Collimators defined the neutral beams, and sweeper magnets removed charged particles. The beams were designed based on a detailed Monte Carlo simulation and on the experience gained from experiments E731, E773, and E799-I. The beam was designed to optimize the kaon flux and the sharpness of the beam edges; it also minimized the neutron contamination in the beams, hadronic "halo" around the beams, and the muon flux through the detector.

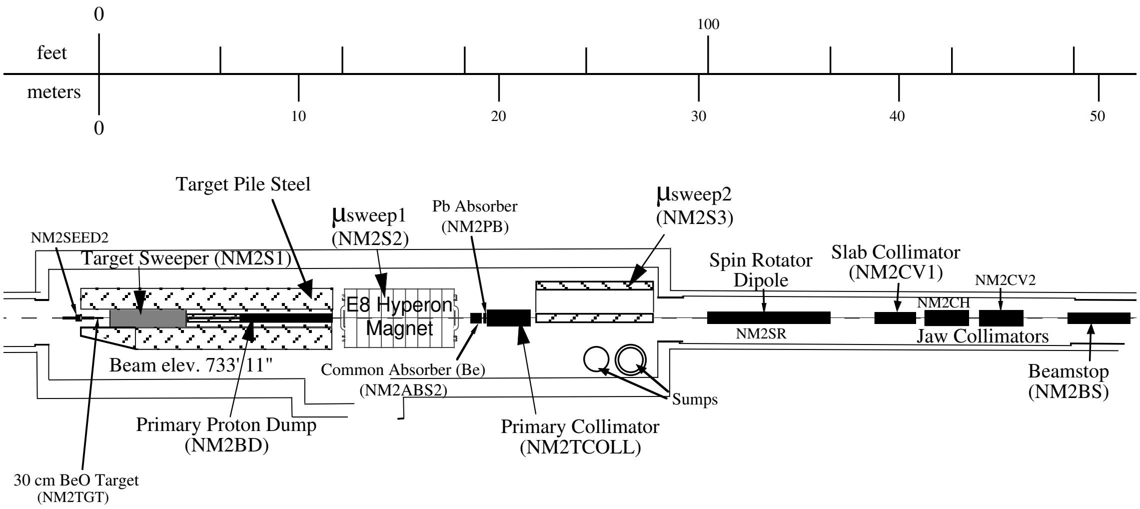

Figure 7. Secondary beam elements in the NM2 enclosure. The common absorber was not used in E799.

2.1.1 The primary proton beam and target

E799-II began with an 800 GeV/c proton beam provided by the Fermilab Tevatron. Protons arrived over a 18-23 second[1] "spill" once every 61 seconds. During E799 running, 2.5×10^{12} to 5×10^{12} protons were delivered per spill. The beam had a 53 MHz microstructure; protons arrived in 1-2 nsec "buckets" once every 19 nsec. A signal (the "Tevatron RF") from the accelerator was synchronized to this RF microstructure, and it was used to synchronize the trigger and all timing measurements in the experiment.

The primary proton beam was nearly horizontal, aimed downward at an angle (the "targeting angle") of 4.8 mrad. The beam was focused on a BeO target which was 3×3 mm transverse to the proton beam, and 30 cm (1.1 interaction lengths) long. The target was in the NM2 beamline enclosure. Beamline components in this enclosure are shown in Fig. 7.

[1]The exact spill duration was dependent on the beam requirements of other experiments.

2.1.2 The experimental coördinate system

The center of the target was the origin of the experiment's coördinate system. We defined the positive Z-axis as the downstream horizontal direction along the beam; the Y-axis was the vertical, with the positive X-axis horizontal and transverse to the beam. The $+Z$ axis was also referred to as "north," although it was closer to northeast in reality. The $+X$ direction was "west."

2.1.3 Sweepers and collimators

Two meters downstream of the target was the "target sweeper," the first of four sweeping magnets. Downstream of the target sweeper, the primary beam was absorbed in a water-cooled copper beam dump. The target, target sweeper, and beam dump were all inside a steel cave called the target pile.

At $Z = 14$ meters, another magnet (the "E-8 hyperon magnet") swept charged particles out of the beam. Two nearly parallel neutral beams were defined roughly by a primary collimator at $Z = 21$ meters; the beam shapes were square, and their centers were separated by 1.6 mrad. The beam centers were at $Y = 0$. A lead absorber at $Z = 19$ m removed photons in the beams. Two different beam sizes were used: "small beams" in the earlier data sets, and "large beams" in the later data sets. Different sets of primary collimator channels were used for the two beam sizes.

Downstream of the primary collimator, the beams entered a vacuum region which extended 138 meters. A third sweeping magnet ("μ-Sweep 2") removed any remaining charged particles from the beam, including upstream decay and interaction products.

At $Z = 40$ meters, a "slab collimator" between the beams was used to absorb kaons with high transverse momentum which could cross between the two beams. The slab collimator was removed when large beams were used. Downstream of the slab collimator, two "jaw collimators" could be used to trim the beams; they did not form limiting apertures.

The beam traveled through an evacuated pipe into the NM3 enclosure, where a tungsten collimator (the "defining collimator") reduced the beams to their final dimensions. The defining collimator was a tungsten assembly with tapered beam holes. It extended from 85.0 to 88.0 meters downstream of the target. The beams emerging from the end of the collimator were nominally square, and each beam subtended a solid angle of 0.25 μstr (small beams) or 0.35 μstr (large beams). The small beams were 4.4 cm across at the exit of the defining collimator; the large beams were 5.2 cm.

Downstream of the defining collimator, a final sweeper magnet removed most charged particles created in upstream decays or interactions in the collimator. This Argonne ZGS dipole extended from $Z = 90$ to $Z = 93$ meters, and provided a transverse momentum kick of 1.1 GeV/c. The vacuum pipe divided into two $4''$ square beam pipes inside this magnet, and the downstream end of the twin-pipe segment formed a limiting aperture.

The beams at this point were composed primarily of neutrons and K_L, with a n/K_L ratio of approximately 3 and a total neutral hadron flux of 25 to 50 MHz. Small numbers of K_S, Λ^0, $\bar{\Lambda}^0$, Ξ^0, and $\bar{\Xi}^0$ remained in the beam as well. Due to their short lifetimes, these particles tended to decay upstream.

2.1.4 Accidental counters

Two sets of counters in NM2 were used to trigger on primary beam activity uncorrelated with detector activity. Events triggered on these counters were used to record accidental detector activity for use in the Monte Carlo simulation.

The primary system used for recording accidental events was the "90° target monitor," which was a series of three counters instrumenting a small hole in the target pile, aimed at 90° with respect to the beam. A second system, the "accidental muon counter," was a two-counter muon telescope designed to detect large-angle muons penetrating the target pile steel. The accidental muon counter signals were found to be more correlated with detector activity, so the 90° target monitor was used instead for Monte Carlo generation.

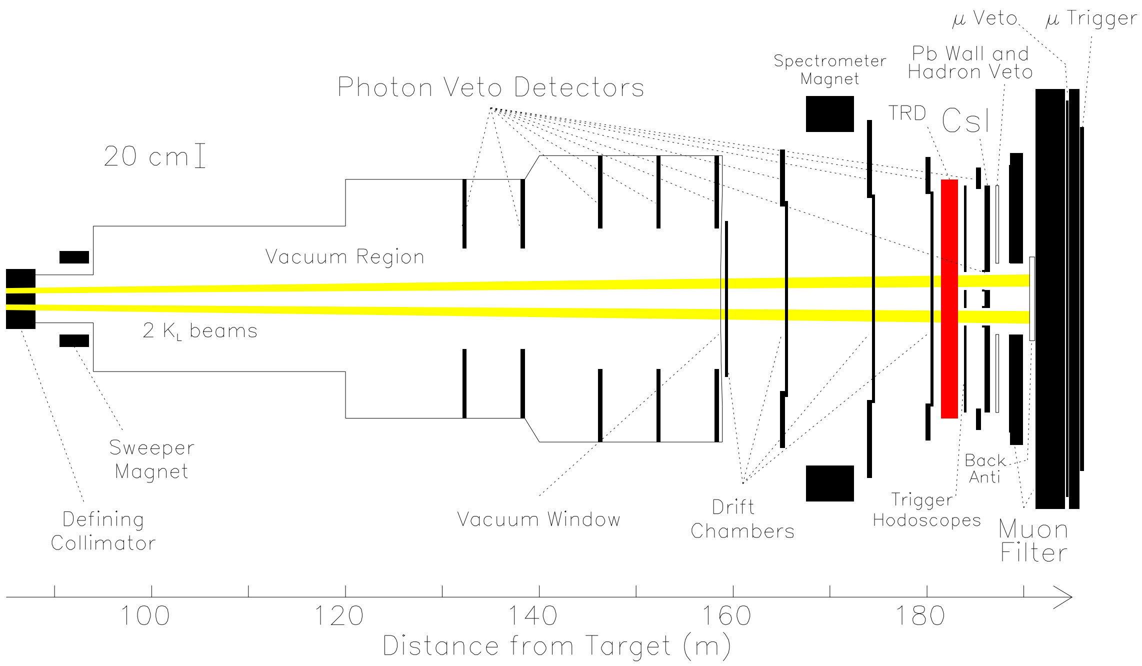

Figure 8. Plan view of the decay volume and the E799 spectrometer. The horizontal scale is compressed.

2.2 The decay volume and vacuum window

The downstream fringe field of the final sweeper dropped below 60 gauss at $Z = 94.1$ meters. At $Z = 93$ meters, the main decay volume began. The vacuum pipe was 18 inches in diameter, and flared out to 96 inches at $Z = 159$ m. The NM3 enclosure ended, and the NM4 enclosure began, at a gate about halfway down the decay volume. Only about 5% of the K_L which entered the vacuum decay region decayed before leaving it; the remainder either decayed in the spectrometer itself or struck a downstream neutral beam dump.

A vacuum window, at $Z = 158.9$ meters, defined the end of the vacuum decay region. This window was a Kevlar weave (for mechanical strength) laminated with a Mylar sheet which formed a gas seal. On the downstream side of the window was a volume, approximately 25 cm deep in Z, which could be filled with helium. This volume ("Helium bag 1a") was bounded on the downstream side by a 1 mil Mylar window.

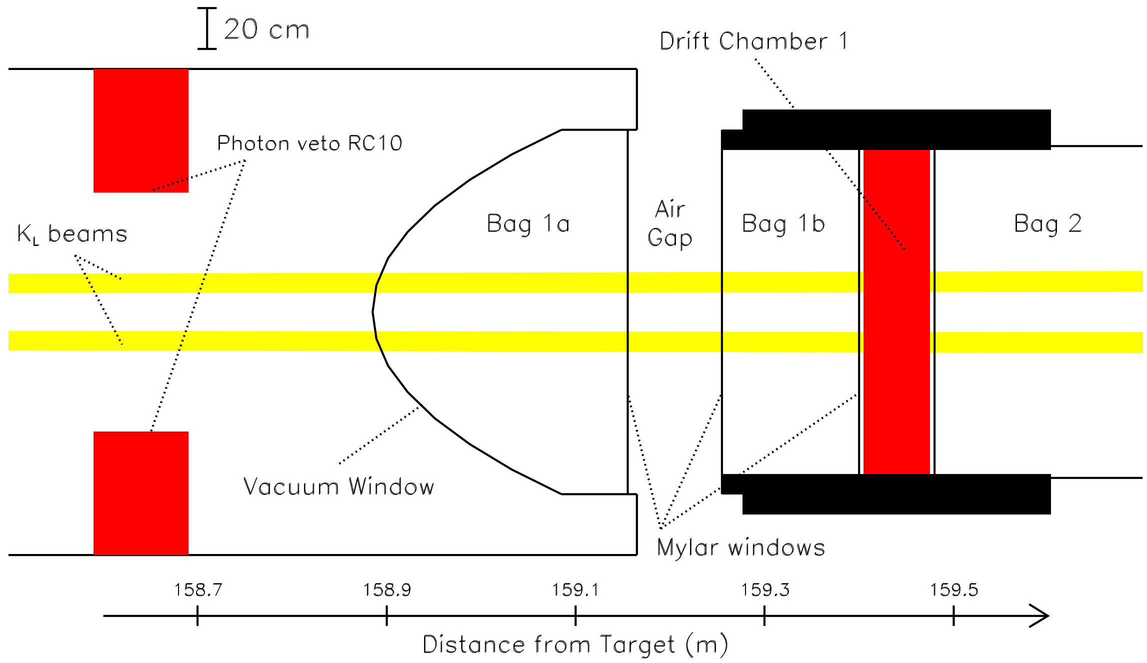

Figure 9. Plan view of the region around the vacuum window. Note the expanded horizontal scale.

Although this bag was to have been filled with helium to reduce multiple scattering and photon conversions, studies (Chapter 7) have shown that this bag was probably filled with air during the run. Figure 9 shows an expanded view of the vacuum window region.

A 10 cm thick air gap followed Bag 1a. This air gap allowed a safety shutter to be lowered in front of the vacuum window during accesses to the detector hall. Downstream of the air gap, another mylar window formed the upstream boundary of Bag 1b, which was filled with helium. A double Mylar window with a narrow CO_2-filled region at $Z = 159.42$ meters formed the downstream end of Bag 1b and the upstream window of the first drift chamber. The CO_2 gap was intended to protect the chamber from helium permeating from the helium bag.

Table 2. Z-position (distance from target) of the midplanes of the drift chambers and the spectrometer magnet.

Chamber	Distance from target (meters)
DC1	159.44
DC2	165.59
Magnet	170.00
DC3	174.61
DC4	180.51

2.3 The charged spectrometer

Charged tracks were detected with a set of four drift chambers (DCs) and a spectrometer magnet. The spectrometer magnet, a large-aperture dipole with a transverse momentum kick of 200 MeV/c, was placed between the second and third DCs. The wire frames for the chambers were used previously in experiments E731, E773, and E799-I, and the details of their wire geometry are described elsewhere [34] [35] [36] [8] [37] [38] [39] [40]. Cathode wires used in this experiment were 4 mil gold-plated aluminum; anode wires were 1 mil tungsten.

The drift chambers recorded hit positions in both X and Y; each view consisted of two overlapping planes offset by a half-cell. The chambers used an 80/20 mix of argon and ethane, and ran at high voltages between 2450 and 2600 V. Helium bags were placed between the drift chambers to reduce multiple scattering, photon conversions, and beam interactions in the spectrometer. The Z-positions of the chambers are shown in Table 2.

The chambers were instrumented with a set of high-gain preamplifiers mounted on the chamber frames, and a rack-mounted postamplifier/discriminator system [41]. The discriminator outputs were fanned out to the trigger systems and a set of LRS 3377 time-to-digital converters (TDCs).

A series of eight transition radiation detector (TRD) modules followed Chamber 4. These detectors discriminated between electrons and pions. Each module consisted of a

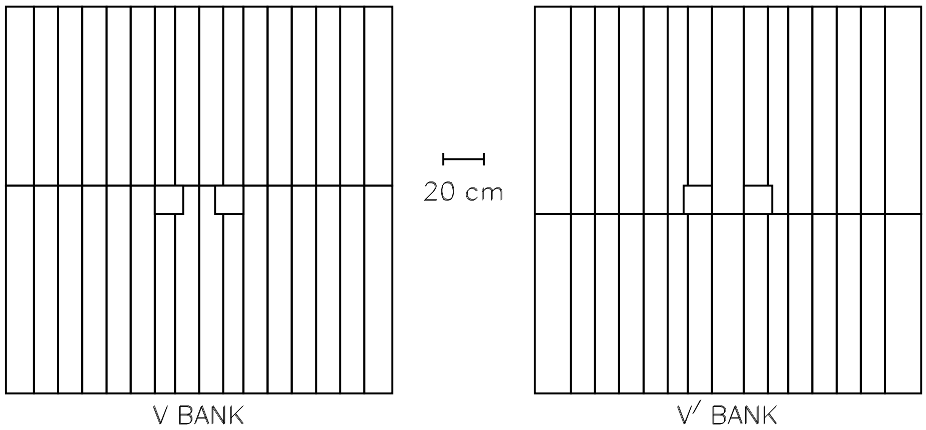

Figure 10. The V and V′ counter banks as viewed from upstream.

polypropylene felt radiator followed by a multiwire proportional chamber. The radiators had holes to allow the neutral beams to pass without interacting. The total radiation thickness of the TRD system was $\sim 0.14X_0$ outside the beam holes. The TRDs were not used for the analysis described in this thesis.

2.4 The trigger hodoscopes

Two banks of scintillation counters at $Z = 183.9$ m were used at trigger level to select events with charged particles. The two planes of the trigger hodoscope system were labeled V and V′. The scintillators were oriented vertically, and extended from a midline to the top and bottom of the fiducial region. Beam holes were cut in the counter banks to reduce rates and radiation damage from interactions. The planes were designed such that cracks between counters in one plane did not overlap with cracks in the other plane. A plane of horizontally oriented counters, labeled H, was constructed but never used for taking physics data. The geometry of the V and V′ counter banks is shown in Fig. 10.

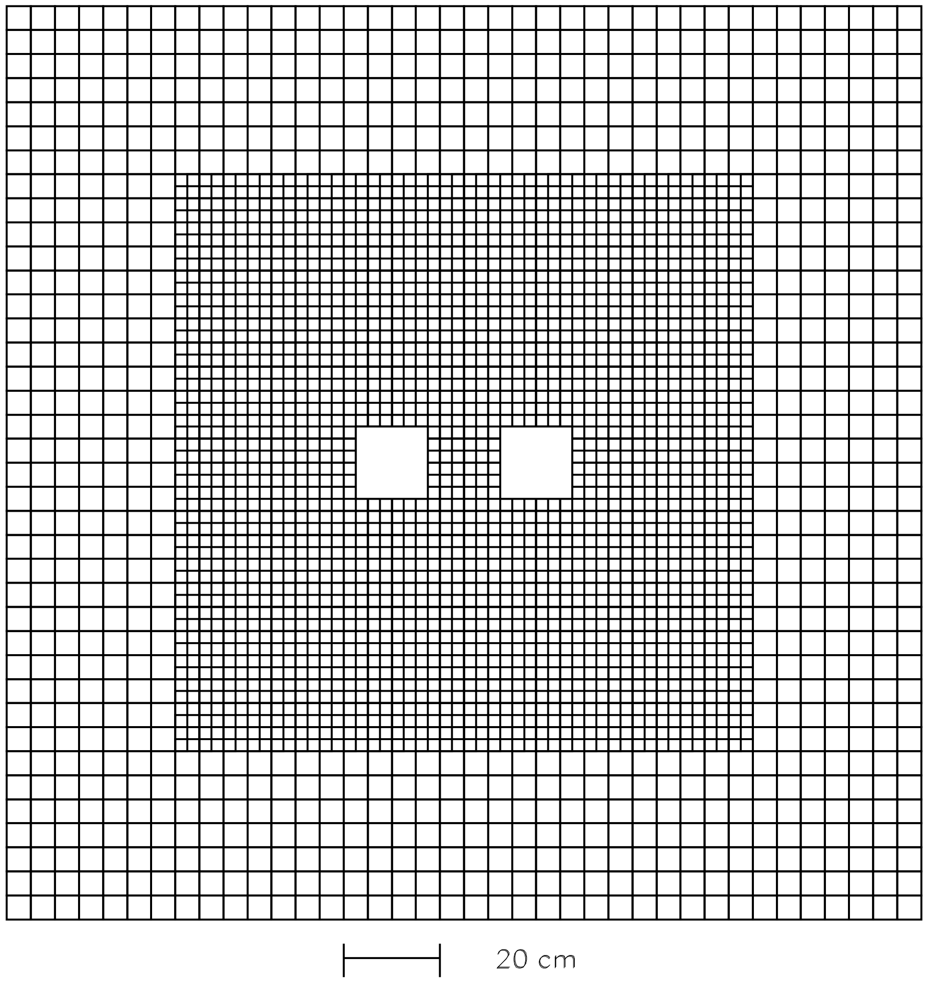

Figure 11. Crystal boundaries in the CsI calorimeter.

2.5 The cesium iodide calorimeter

A 3100-channel pure CsI electromagnetic calorimeter at $Z = 186$ meters recorded photon and electron energies. This calorimeter was 1.8×1.8 meters transverse to the beams, and consisted of 868 “large” ($5.0 \times 5.0 \times 50.0$ cm) crystals and 2232 “small” ($2.5 \times 2.5 \times 50.0$ cm) crystals. The crystals were 27 radiation lengths thick, allowing photon and electron showers to be captured completely. The crystals were 1.4 nuclear interaction lengths; many hadrons passed through the calorimeter without interacting.

The crystal geometry is shown in Fig. 11. Two beam holes, each 15 cm square, allowed the neutral beams to pass through without interacting.

2.5.1 Crystals and photomultipliers

Each crystal was individually wrapped with regions of black and reflective film to compensate for any non-uniformity in the crystal's light output [42]. At the downstream end of the crystal, a "cookie" of transparent RTV rubber provided an optical coupling between the crystal and a photomultiplier tube. Opaque masks in the cookie were used to compensate for the varied light outputs of individual crystals. A filter at the cookie was used to remove the slow component of the CsI scintillation light, which would otherwise have caused significant light pileup from out-of-time events.

Large crystals were instrumented with 6-stage, 1.5″ Hamamatsu R5330 PMTs; 5-stage, 0.75″ R5364 tubes were used for the small crystals.

2.5.2 The CsI readout system

The anode signals from the PMTs were digitized and read out through a digital PMT base (DPMT) system. One DPMT card was used for each CsI crystal. The major components on the card were a charge integrating and encoding (QIE) circuit, an 8-bit flash ADC, and a "driver-buffer-clock" (DBC) circuit [43]. The QIE and DBC devices were custom integrated circuits designed specifically for this calorimeter.

The QIE was a hybrid digital/analog circuit which integrated the charge from the PMT. The current was divided among seven 1 pF capacitors, with each capacitor receiving a different binary fraction of the current ($I/2$, $I/4$, ... , $I/128$), and one 2 pF capacitor which received $I/128$. The voltage levels on the eight capacitors were compared to a common level, and the result of the comparison was used to select the capacitor to digitize. The output of the QIE was the analog level of the appropriate capacitor, and a set of "exponent" bits which identified the sensitivity level selected by the QIE. The QIE contained four identical circuits, which were used in a round- robin sequence.

"Capacitor ID bits" identifying the circuit were also read out from the QIE. The circuit was clocked with the Tevatron RF, and each integration period represented a 19 nsec bucket. The QIE rotated through all four circuits in four RF buckets (96 nsec). Once each RF bucket, the analog output of the QIE was digitized by an 8-bit flash ADC on the DPMT card. These eight bits formed a mantissa which, when scaled by the range, corresponded to the integrated charge from the PMT.

The digital information from the QIE and the FADC were stored continuously in a FIFO buffer in the DBC. Upon receipt of a Level 1 trigger (which was synchronous with the event), 32 slices of this FIFO were transferred to a holding buffer in the DBC. The contents of this buffer were transferred to the "Pipeline" system [44] on receipt of a Level 2 trigger (which could arrive up to several microseconds after the Level 1 signal).

The Pipeline, a custom VME system, was not a conventional readout buffer; its main function was to sparsify the data coming from the 3100 DBCs and reduce the data from an event to a reasonable size. The Pipeline removed channels from the readout list if they did not have energy above a fixed threshold, and could read out data from fewer than the 32 clock cycles passed from the DBC. The energy thresholds and selection of slices to be read out could be set individually for each physics trigger.

2.5.3 Laser calibration system

A dye laser and a fiber-optic light distribution system were used used to inject pulses of light into one end of each CsI crystal for calibration and diagnostic purposes. A set of PIN diodes was used to measure the total light output of each laser pulse. This allowed a direct calibration of the DPMT response to an independently measured light pulse. Laser data were taken in special "scans" during beam-off periods, in which all DPMT ranges were probed by varying the laser light output. Several laser flashes were also taken during each beam spill to allow run-by-run corrections to the calibration. These "quiescent" pulses were not used in the final analysis.

2.6 Photon veto systems

2.6.1 Ring-shaped vetoes

Ten photon veto detectors defined the outer perimeter of the fiducial region transverse to the beam. These detectors could be used to search for and reject events in which extra photons or electrons left the fiducial region.

Five "ring counters" (RCs) [45] were placed inside the vacuum region, and were numbered RC6 through RC10. These detectors had circular perimeters and square inner apertures, and were divided azimuthally into twelve modules. The RCs were each composed of 24 layers of lead and scintillator, and light was transferred through fiber-optic light guides. Each module was read out by a single PMT.

In the spectrometer region, four "Spectrometer Anti" (SA) [46] detectors defined the fiducial region. The first three, SA2, SA3, and SA4, surrounded DC2, DC3, and DC4 respectively. The fourth was between the VV$'$ banks and the CsI, and was called the "Cesium Iodide Anti" (CIA). The SA/CIA detectors were not in the vacuum region, and they had somewhat different geometry. Both perimeters and inner apertures of the SAs/CIA were rectangular, with the inner dimension slightly smaller than the outer aperture of the corresponding drift chamber (or the CsI calorimeter in the case of the CIA). The Z-positions of all RCs and SAs are shown in Table 3.

The energy deposited in each RC, SA, and CIA module was recorded in 10-bit LRS 4300 ADCs. The energy was also discriminated and could be used to veto events at the trigger level.

2.6.2 Collar Anti

A veto detector, the Collar Anti (CA) [40], defined the boundary of the fiducial area at the beam holes of the CsI. The CA was a tungsten-scintillator sandwich detector which was placed immediately upstream of the CsI and shadowed part of the innermost crystals. A diagram of the CA position is shown in Fig. 12.

Table 3. Z-position (distance from target) of the upstream face of all RC and SA detectors.

Photon Veto Detector	Distance from target (meters)
RC6	132.60
RC7	138.60
RC8	146.60
RC9	152.60
RC10	158.60
SA2	165.12
SA3	173.99
SA4	180.02
CIA	185.19

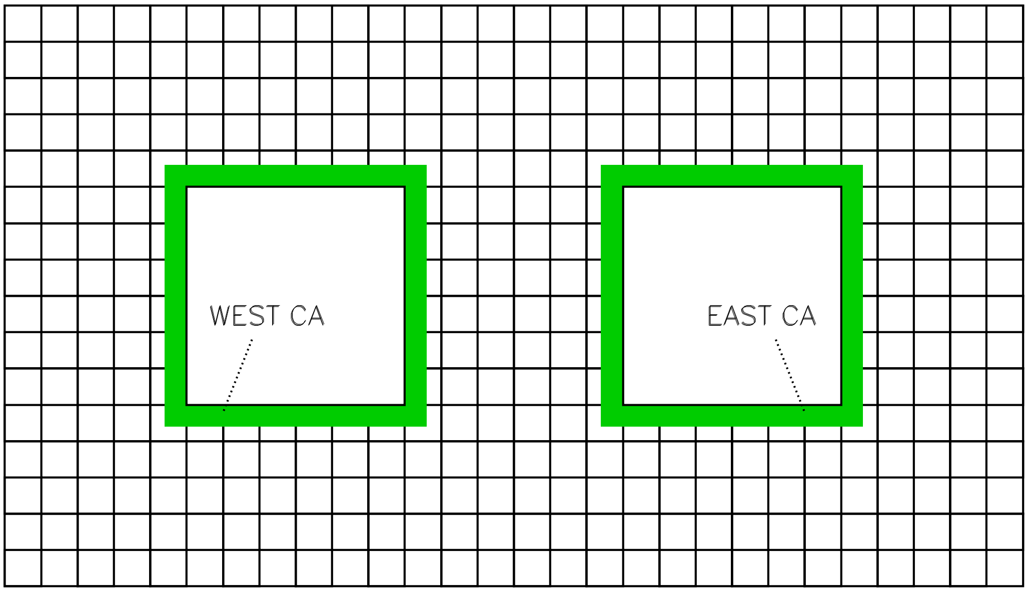

Figure 12. Position of the Collar Anti (CA) at the CsI, as viewed facing downstream. Only the central part of the calorimeter is shown.

Table 4. Z-position (distance from target) of detector components downstream of the CsI calorimeter.

Detector element	Distance from target (meters)
Lead wall	188.53
HA	188.97
Filter steel 1	189.09 to 190.09
BA	191.09
Filter steel 2	191.74 to 194.74
MU2	194.83
Filter steel 3	195.29 to 196.29
MU3	196.36

2.7 The neutral beam dump and hadron/muon vetoes

The components of the detector downstream of the CsI included a hadron veto detector, the neutral beam dump, and the muon identification system. The Z-position of each of these detectors is shown in Table 4.

2.7.1 The lead wall and Hadron Anti

Downstream of the CsI, at $Z = 188.5$ m, a wall of lead bricks 15 cm (0.9 interaction lengths) thick was used to produce hadron showers and range out any electromagnetic showers. The Hadron Anti (HA), a bank of scintillation counters immediately downstream of the lead wall, detected the hadron shower products. One large hole (64 cm in X; 34 cm in Y) allowed both beams to pass through the lead wall and HA. The boundaries of the counters are shown in Fig. 13.

An analog sum of the energy deposited in the HA counters was discriminated and used at trigger level to reject events with hadrons in the final state. In early runs, this sum was performed by a custom AC-coupled summing circuit; due to problems at high rate, this circuit was replaced with a DC-coupled sum for later runs.

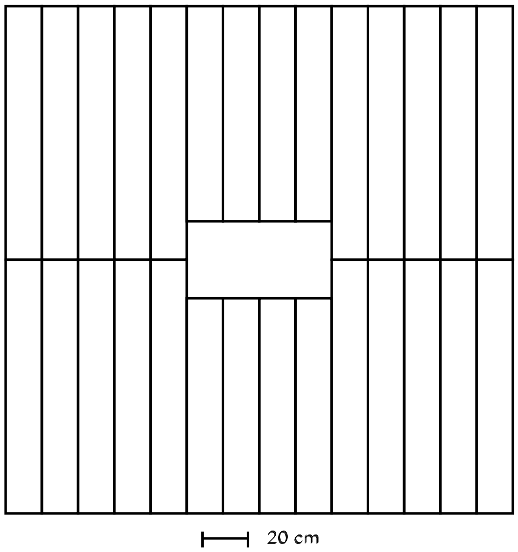

Figure 13. The Hadron Anti scintillator bank.

2.7.2 The Back Anti

Immediately downstream of the HA, the beams passed through a hole in a 1 meter thick steel wall before the beams were dumped in a veto detector called the Back Anti (BA). The steel shielded the HA from backsplash radiation from the neutral beam dump.

The BA was a lead-scintillator sandwich detector with a total thickness of 30 radiation lengths or one interaction length. The detector could be used to veto photons at either the trigger or the analysis level, but it was not used in the analysis described here because the high rate of neutron interactions would have caused substantial signal loss.

2.7.3 Muon identification systems

A large muon veto hodoscope (MU2) sat behind the neutral beam dump and three additional meters of steel. The large transverse dimensions of this detector, 2.99 m $\times$ 3.78 m, were necessary because of its extreme downstream position and the need to catch muons which scattered in the steel. The steel stopped nearly all particles except muons. MU2 could be used to veto events with muons at trigger level, but this requirement was dropped early in the run due to high accidental rates in this detector.

Another set of muon hodoscopes (MU3) was placed behind another meter of steel. These detectors were used at trigger level to select events with two muons for certain rare decay analyses. They were not used for the measurement described here.

2.7.4 Non-CsI readout systems

All detector systems except the CsI calorimeter were read out through a system modeled on the LRS FERA (Fast Encoding and Readout ADC). This architecture allowed fast (10 MHz) data readout through the front panel of a CAMAC module to an intermediate storage buffer in a controller module in the CAMAC crate. This module, called a Damn Yankee Controller (DYC) [47], was designed by the Fermilab Computing Division specifically for high-rate fixed-target experiments. The data and readout control signals were transmitted on two buses (made of flat ribbon cable) which stretched across the front of the crate to the DYC.

When a crate received a trigger, the front-end modules read out their data to the buffer in the DYC using a local token-passing scheme within the crate. The DYCs then transmitted the data to the data acquisition system in sequence, using another token loop. There were two independent streams of data from DYCs on separate token loops, and another four from the Pipeline.

The various modules used in the front end included LRS 4300 ADCs (for the vetoes, VV$'$, and TRD system), LRS 3377 TDCs (used in the drift chambers, VV$'$, and the trigger), custom scaler modules used only in the trigger, and custom Fermilab pipelined latches used in the trigger, certain veto systems, and the VV$'$ and muon counters.

CHAPTER 3

THE TRIGGER

This chapter describes the design philosophy, implementation, and the operation (in E799-II) of the E832/E799 trigger systems.

During the 18-23 second Tevatron spill, the primary proton beam arrived in micropulses ("buckets") 1-2 nsec wide, once every 18.83 nsec. In E799, the kaon decay rate in the detector was ~ 1 MHz, and the neutron interaction rate was several MHz. It was impossible to record events at these high rates, so a trigger system selected events to be digitized and recorded by the data acquisiton system (DAQ). The trigger was designed to select the kaon decay modes of interest, while rejecting the vast majority of uninteresting kaon decays and hadron interactions occurring in the experiment. The trigger was a two-level hardware system, with a software filter ("Level 3") in the DAQ providing additional rejection.

Level 1 ("L1") was a deadtimeless trigger which operated in synchrony with the RF microstructure of the beam. It made decisions based on simple logic operations performed on the fastest signals available from the detector (primarily signals from photomultiplier tubes). When Level 1 accepted an event, the digitization and readout process began, and signals from the chambers, calorimeter, and TRD were made available to the slower Level 2 systems. Level 2 consisted of several hardware processors which made decisions based on these signals. The processors each returned one- to four- bit responses, which generally corresponded to counts of chamber hits, calorimeter energy clusters, or track candidates. The processing times for the various Level 2 systems ranged from 0.8

to 2.5 μsec. During this processing time, the trigger was dead – *i. e.* new events could not be accepted. The Level 2 control system received the processor responses and accepted or rejected events based on them.

When an event was accepted by Level 2, the digitization process was allowed to complete and the event was read out to buffers in the DYCs and Pipeline, a process which took about 12-15 μsec before the trigger could accept another event. When an event was rejected, the readout process was aborted by a fast clear signal which was sent to the various detector readout systems, and the trigger was live again within 0.5 μsec. Because of the dead time resulting from each event processed, the experiment could handle no more than $\sim 10^5$ Level 1 triggers or $\sim 2 \times 10^4$ Level 2 triggers per second without unacceptable ($> 50\%$) dead time.

After the data entered the DYC/Pipeline systems, the readout process through the DAQ was buffered and no additional dead time resulted unless the buffers filled up (the trigger rates and data size were tuned to avoid this happening under most circumstances).

The Level 3 filter was implemented on 24 200 MHz processors and 10 150 MHz processors, which operated in parallel in four SGI Challenge servers. Each server was referred to as a "DAQ plane," and the planes were numbered 1 to 4. Planes 1, 2, and 4 each contained eight 200 MHz processors and were dedicated to filtering physics events; Plane 3, which used the 150 MHz processors, was used for detector monitoring and calibration triggers. The filter code was based a simplified version of the offline analysis code. It performed track and cluster reconstruction, and events could be selected or rejected based on this information. Selected events were written out to up to 10 DLT tape drives running in parallel. Each DLT could hold 10 GB of data; data compression in the drives allowed 14 GB of raw data to fit on a typical tape. The data were buffered in VME memories which could hold an entire pulse's data. Filtering and tape-writing were performed during the spill and the following off-spill period. With all tape drives running in parallel, about 1000 events could be recorded per second; this meant that ~ 60000 events could be written per spill.

Up to sixteen different physics triggers ("beam triggers") could be defined as logic functions of signals from the detector, and each trigger could have independent Level 1, Level 2, and Level 3 requirements. The Level 1 and Level 2 definitions were translated into memory lookup tables, which were downloaded to the hardware. The translation was done through a "trigger file interpreter" [48] which allowed a user to define triggers in a pseudo-programming-language format called a "trigger definition file."

In addition to the beam triggers, up to fourteen "calibration triggers" could be implemented, though these had to be externally generated signals and were not programmable.

3.1 Level 1

This section describes the layout of Level 1 trigger systems, including the processing of signals from the detector and the formation of the Level 1 trigger decision.

3.1.1 Level 1 logic hardware

3.1.1.1 Level 1 sources and lookups

The Level 1 system was designed to allow the use of up to 96 "sources," or input signals from the beam and detector, in the Level 1 decision. These sources were generally NIM pulses which were transported from the detector hall on "hardline" cables. The hardlines were semi-rigid coaxial cables on which signals propagated at $\sim 0.9c$ (versus $\sim 0.6c$ on conventional RG-58 cable).

With the exception of the spill timing signals and the DC OR sources (discussed below), the Level 1 sources were typically no more than one bucket (19 nsec) wide. The signals were synchronized such that the signals from particles on a wave-front moving through the detector at velocity c all arrived at the front of the Level 1 logic simultaneously. Therefore, signals from upstream detector elements went through longer delay cables than signals from downstream. For example, a muon traveling at $\approx c$ would pass the RC6 veto counter (at $Z = 132.6$ m) 213 nsec before it passed the MU3 hodoscope (at $Z = 196.4$ m). The signal from RC6 was delayed by an extra 213 nsec, so the signals from the two detectors arrived at the trigger simultaneously.

Although the L1 sources were synchronized as closely as possible before entering the trigger system, there remained a significant (~ 10 nsec) spread in their arrival times for a given event. Most of this spread was inherent in the sizes of counters (some muon counters were 2 meters long, and signals from the end of the counter near the PMT arrived ~ 6 nsec before signals from the far end) and variations in threshold-crossing time between large and small analog pulses in a given counter. At this stage, copies of all the trigger source signals were fanned out to the trigger TDCs (which are discussed in Sec. 3.3) and to CAMAC scalers which were summed over spills to measure raw counter rates.

To remove the sources' timing spread, all sources were resynchronized to the Tevatron RF clock using LRS 622 coincidence circuits. The source signal was required to be *on* at the leading edge of the RF "strobe" (the RF clock pulse at the coincidence circuit). Therefore, all the source signals were timed such that the leading edge of the source signal always arrived *before* the strobe, and the trailing edge arrived *after* the strobe. Signals such as the muon counters, which had a large spread in arrival time, needed to produce correspondingly wide source pulses. If a source arrived late, it did not get resynchronized properly, and could not be triggered on reliably. The arrival time distribution of a well-timed Level 1 source, showing the RF bucket structure, is shown in Fig. 14.

After this resynchronization step, the timing spread of the Level 1 sources was only 1-2 nsec. At this point, the signals were fanned out, with two copies sent to the Level 1 lookup units, and another copy sent to a set of pipelined latches (also discussed in Sec. 3.3).

The Level 1 sources were arranged into twelve[1] "groups" of eight sources each; each group had an independent memory lookup map. Each group's lookup units had sixteen outputs, one for each beam trigger.[2] The memory lookup was performed once at each

[1] Only eleven groups were used in the final configuration, although the electronics for Group 12 remained operational and were used for test signals.

[2] LRS model 4508 MLU's were used, and these modules had only an 8×8 bit memory. So, for each group, the sources were copied and sent to two modules: one module used a copy to generate Triggers 1-8, and the other generated Triggers 9-16. Logically, however, the system behaved as if each group had eight inputs and sixteen outputs.

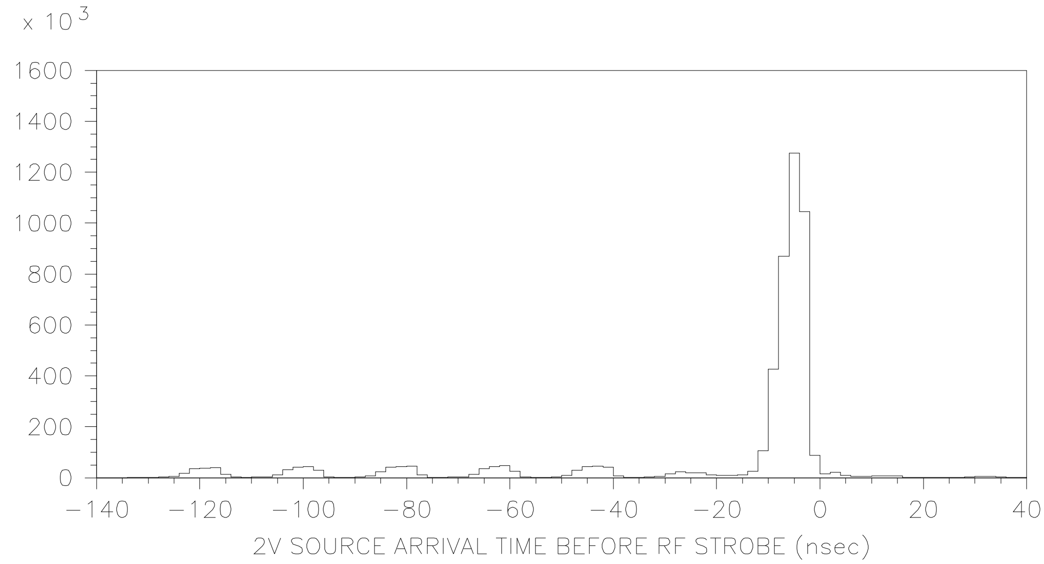

Figure 14. Arrival time before strobe of the Level 1 source corresponding to two hits in the V counter bank. The large peak immediately before zero contains the in-time bucket; the smaller peaks are earlier buckets.

RF clock cycle. A logical diagram of one Level 1 source group is shown in Fig. 15. The output from any given group lookup was the bit-mask of triggers which were satisfied by the input sources from that group. The final set of Level 1 triggers satisfied by a given set of inputs, then, was the trigger-by-trigger AND of the outputs of all twelve groups. This was accomplished using a tree of three levels of sixteen-channel coincidence modules (LRS model 4516). Figure 16 shows a conceptual diagram of this stage of the trigger. A copy of this sixteen-signal trigger "bit-mask" was sent to a pipelined latch and read out into the data.

Because the final trigger was the logical AND of the outputs from each group, it was not possible to define triggers which used 'OR' logic between sources in different groups. This required careful arrangement of the sources such that any two signals which would need to be used in OR logic remained in the same group. The individual Level 1 sources will be described in Sec. 3.1.2.

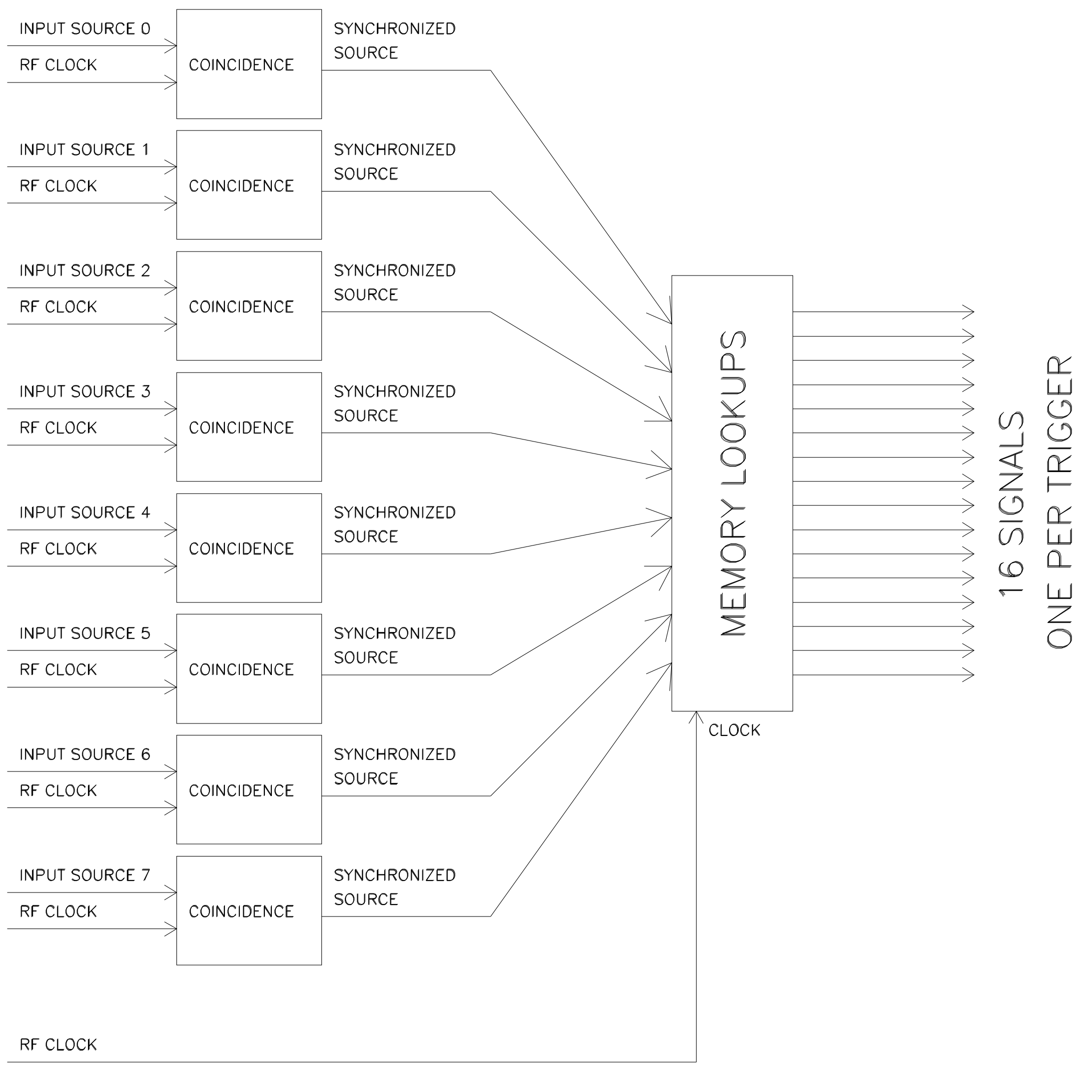

Figure 15. Logical arrangement of one group of Level 1 sources and memory lookup.

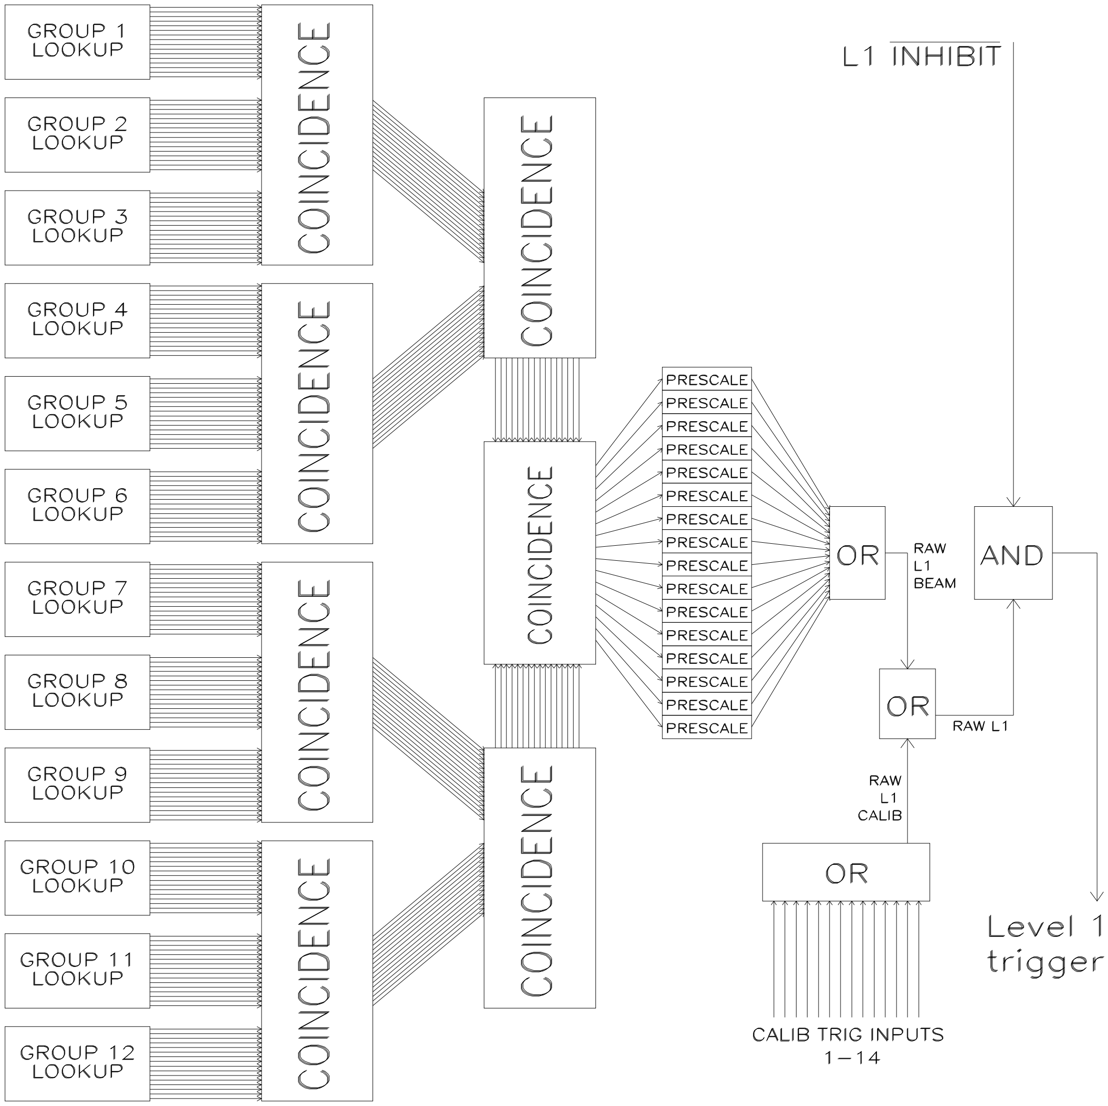

Figure 16. Conceptual diagram of formation of the Level 1 trigger.

3.1.1.2 Prescales

Each of the sixteen Level 1 triggers was sent to a programmable CAMAC prescale module. This allowed certain triggers to be prescaled by integer fractions. The module could implement any prescale factor of the form $1/n$ or $(n-1)/n$. The prescales were declared in the trigger definition file, and the software which downloaded the memory lookup maps was also used to program the prescales. After the prescales, copies of the individual trigger bits were fanned out to a pipelined latch and a set of CAMAC scalers.

After the prescale stage, a grand OR of the sixteen beam trigger signals was formed. This OR was the "Raw L1 Beam Trigger."

3.1.1.3 Calibration triggers and Raw Level 1

In addition to the sixteen beam triggers, up to fourteen calibration triggers could be used. These triggers were generated by different detector systems or function generators, for purposes such as reading pedestals or recording calibration laser flashes. The OR of all the calibration trigger inputs was the "Raw L1 Calibration Trigger." The OR of the Raw L1 beam and calibration triggers was taken; this was the "Raw L1 Trigger".

3.1.1.4 Formation and distribution of the Level 1 trigger

Up to this stage, the trigger system ran continuously, clocked by the Tevatron RF, and independent of the state of the readout or DAQ. After forming the Raw L1 Trigger, a Level 1 trigger was issued if the L1 INHIBIT signal was *not* on. The L1 INHIBIT was set under any one of the following conditions:

- Level 2 trigger decision in progress
- Event readout to DYCs in progress
- Any DYC, Pipeline, or DAQ memory buffers full
- Any trigger error condition (from Level 2 state machine; see Sec. 3.2.1)

- Triggers vetoed from DAQ because no data run in progress

Once a Level 1 trigger was formed, the Level 1 inhibit was set, and the Level 1 trigger bit-mask was latched in a LRS 2371 module for use by Level 2. Each bit of this bit-mask was set if the corresponding beam trigger fired. If a calibration trigger fired, it overrode any beam trigger, and the calibration trigger's bit was set along with the two highest bits (triggers 15 and 16). These two bits, when both set, acted as a flag to indicate that a calibration trigger fired[3].

The Level 1 trigger was distributed to every readout crate in the experiment. It was used as a signal to the Level 2 processors to begin processing, as a gate for all ADCs, and as a common stop signal for all TDCs in the experiment. Level 1 also triggered the DPMTs to clock data into the 32-slice DBC holding buffer (Sec. 2.5.2).

3.1.2 Level 1 sources

This section describes all the detector and beam signals which were used as Level 1 sources.

3.1.2.1 Beam timing signals

Two spill-timing signals from the accelerator were used in the trigger:

- The "SPILL" signal, which was a constant level set for the 23-second portion of the Tevatron cycle when the energy was 800 GeV and beam could be extracted to the experiment.
- The "NEUTRINO PING" signal indicated that fast-spill beam was being delivered to experiment E815, a neutrino-scattering experiment in an adjacent beamline. This signal was used as a veto, because fast-spill beam occasionally leaked into the wrong beamline and could cause very high detector activity rates for a short period.

[3]Note that this required beam triggers 15 and 16 to be defined in such a way that they could never be satisfied simultaneously, or a false calibration trigger flag could be set.

3.1.2.2 Photon veto sources

Discriminated signals from the photon veto detectors were used as Level 1 sources.

There were nine signals from the ring-shaped vetoes (RC6-10, SA2-4, CIA), one signal from each counter bank. The source fired if any counter in the bank registered more than 0.5 GeV of in-time energy.

Two sources were used from the CA: one from the east and one from the west detector. These fired if any PMT registered more than 14 GeV.

The BA used six sources: the east and west sides of the detector were broken up longitudinally into three sections, and an analog sum of each section was discriminated. The digital levels were used as Level 1 sources. The BA activity rates were over 10 MHz in E799, and so the BA was not used in the trigger.

3.1.2.3 Accidental counters

Two Level 1 sources were used for the two accidental counter systems in NM2: the 90° target monitor and the accidental muon counters.

3.1.2.4 Total energy trigger

An AC-coupled analog sum was formed from the dynode signals of all 3100 PMTs in the CsI calorimeter. This analog sum was used to trigger on minimum in-time energy in the calorimeter. A LRS 821 four-channel NIM discriminator was used to set the thresholds.

Four Level 1 sources were used for the total energy trigger; each source corresponded to a different energy threshold. In E799, the nominal thresholds were 10 GeV, 18 GeV, 25 GeV, and 38 GeV, and these were referred to respectively as "E-total (ET) thresholds" 1 through 4.

3.1.2.5 Trigger hodoscope sources

The V and V′ hodoscopes were used extensively in Level 1, and two groups of sources (a total of 14 sources) were devoted to them.

Six hit-multiplicity sources were formed using a series of LRS 4413 16-channel CAMAC discriminators. The analog signals from the photomultiplier tubes were discriminated in these modules, and an analog sum output from the 4413 modules indicated how many counters had fired in each plane. Discriminators on this sum provided separate Level 1 sources for ≥ 1 V, ≥ 2 V, ≥ 3 V, ≥ 1 V′, ≥ 2 V′, and ≥ 3 V′ hits. All VV′ hit multiplicity sources were in the same Level 1 group.

Eight additional sources (the "East-West-Up-Down" logic) were defined for use in two-body kaon decay modes, where the two particles have equal and opposite transverse momentum. These signals fired based on a single hit in the east, west, upper, or lower half of each counter bank. These signals were not widely used in E799; the experiment did not attempt to trigger on many two-body modes.

3.1.2.6 Drift chamber OR sources

A valuable feature for the Level 1 trigger was a set of sources which allowed triggering on upstream drift chamber activity. These sources (the "Drift Chamber ORs") were segmented in regions of 16 wires in the two overlapping planes of a drift chamber. Each 16-wire region was 10.2 cm wide; these groups of wires were called "paddles." The OR of all the wires in a paddle was formed, and a central controller counted the number of hit paddles in a chamber plane.

There were two sources each from the X and Y views of DC1 and DC2, for a total of eight Level 1 sources: one source fired when ≥ 1 paddles were hit, and the other source fired when ≥ 2 paddles were hit. Due to the ~ 150 nsec drift time, the DC OR sources could not have single-bucket timing resolution as most of the other Level 1 sources had. The DC OR sources were set for 90 nsec, timed to cover slightly more than half the drift time. (Because the OR of the two overlapping planes was used, only half the drift time needed coverage because a track could be expected to deposit a hit in one of the two planes in the first half of the drift time.) These eight DC OR sources comprised one Level 1 group.

3.1.2.7 Hadron veto and muon systems

The HA was used at Level 1 to reject events with hadrons in the final state, particularly K_{e3} ($K_L \to \pi^{\pm}e^{\mp}\nu$) and $K_L \to \pi^+\pi^-\pi^0$ decays. The trigger sources from the HA were discriminated thresholds on an analog sum of signals from the 18 counters in the detector.

In the winter run and the very early part of the summer run, this analog sum was performed with a custom AC-coupled summing circuit similar to the boards which performed the CsI calorimeter energy sum. However, it was realized that the high HA activity rates (> 1 MHz) in E799, and occasional very large pulses from high-multiplicity interaction events, caused the baseline of the AC-coupled sum to shift significantly during high-intensity running. This reduced the effectiveness of the veto at high rates, when it was most critical. Therefore, for the summer run, the analog sum was augmented with a DC-coupled sum using LRS 428 analog summing modules.

In the winter run, the trigger sources from the HA were two different thresholds on the AC-coupled sum. One was a moderate threshold equivalent to $\sim$2.5 minimum-ionizing particles (MIPs). This source was called "HA-HI" for historical reasons. The other source was a higher threshold, at about 7 MIPs, called "HA-SUPERHI."

In the summer run, the HA-SUPERHI source was removed and replaced with a threshold similar to HA-HI, on the DC-coupled sum.

The muon counter bank sources were simple counter hit multiplicity thresholds. Each of the three counter banks had sources corresponding to ≥ 1 and ≥ 2 hits. The large MU2 bank, which had overlapping X-view counters and was hermetic, was used as a veto. The two MU3 banks (X and Y views) were used to trigger on one-muon and two-muon signals.

3.1.2.8 Hyperon trigger elements at Level 1

Three special Level 1 systems were used only for triggering on hyperon decays. These decays frequently had a high-momentum proton track through the beam hole in the CsI, and the trigger elements were designed to search for this track. Special DC OR sources were used to trigger on tracks in the beam regions of DC1 and DC2. Downstream of the CsI, a thin scintillation counte called the "hole counter" (HC) was placed in each beam hole, and these signals were used as Level 1 sources. Four counters (the "hole guards" (HG)) were installed immediately upstream of the BA and surrounding the outside of the active area of each HC. Rates in the HG were very high, primarily from backsplash radiation from the BA. The hole guards were never used for physics triggers.

3.2 Level 2

The Level 2 system consisted of the Level 2 processors, a control system driven by a programmable state machine, and the interface to the DAQ. When an event was accepted by Level 1, the Level 2 processors and control system began processing the event. If Level 2 accepted the event, the data were read into the DAQ; if the event was rejected, an ABORT signal was issued and the front end was cleared.

3.2.1 Level 2 control system

The Level 2 control system was centered around the Level 2 state machine. This was a 16×16 bit memory lookup unit (LRS model 2373). The unit was clocked at 10 MHz, and at each clock cycle the unit received a six-bit "input state" which was fed from six "output state" bits from the previous cycle. The input state indicated whether the trigger was IDLE (ready to accept a new Level 1 trigger), waiting for Level 2 processors, waiting for readout to complete, or in STOPPED or ERROR states. When the trigger was running properly, the state machine was in the IDLE state most of the time; when the trigger state was polled, it generally reported that it was IDLE.

Other inputs to the state machine included the Level 1 trigger, signals from the DAQ which indicated beginning and end of a data run, and various "busy," "full," and "done" signals from Level 2 processors and readout crates. Outputs included Level 2 accept, abort, and clear signals, a Level 1 inhibit, and error signals. The operation of the state machine is described in detail in Ref. [49].

The Level 2 logic followed the pattern of the Level 1 logic in using source signals arranged in groups and a memory lookup unit for each group. At Level 2, there were three groups of 16 sources each, and the lookup was performed by an LRS 2373 Memory Lookup Unit. The first two groups were devoted to drift chamber hit counting (detailed in Sec. 3.2.2), and Group 3 comprised the remaining Level 2 processors. The Level 2 groups and trigger mask logic are shown in Fig. 17.

Each Level 2 processor returned one to four data bits and a "done" bit. The "done" bit was set when the individual processor was finished and its data bits were valid.

The 16-bit AND of the three 16-bit output masks from the three memory lookups was taken; this was the mask of triggers which were satisfied by the Level 2 sources. The AND of this and the mask of triggers which satisfied Level 1 was taken, and this formed the Level 2 trigger mask. The lookup maps were programmed such that a trigger could never be vetoed by a processor before its "done" bit was set. The Level 2 mask would become zero if an event was vetoed by a processor which was done.

If the Level 2 trigger mask was nonzero and all necessary processors were not done, the state machine remained in a waiting state. If all processors finished processing and the Level 2 trigger mask was still nonzero, then the state machine issued a Level 2 trigger. If the Level 2 trigger mask dropped to zero, the state machine issued an ABORT and cleared the front end readout crates. The trigger was live (able to acccept another event) ~ 1 μsec after an ABORT. The readout after a Level 2 accept took 12 to 15 μsec before the trigger was live again.

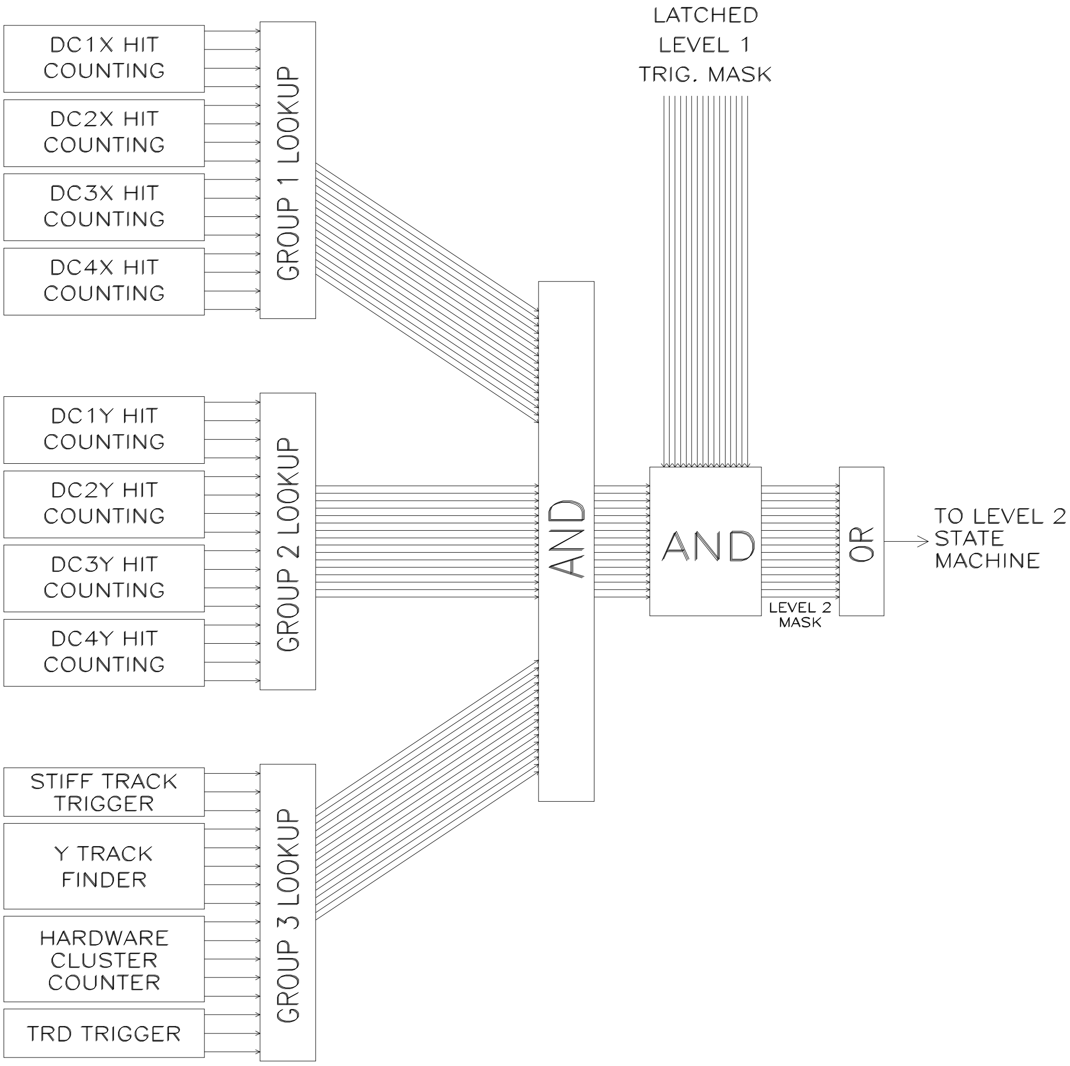

Figure 17. Level 2 groups and Level 2 trigger mask formation.

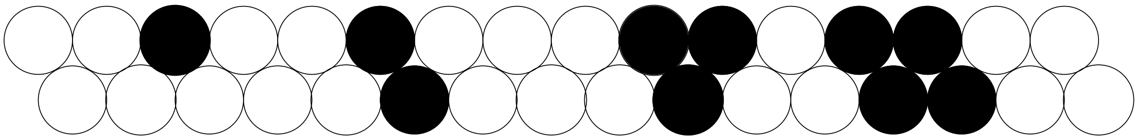

Figure 18. Drift chamber hit topologies used in the hit counting system.

3.2.2 Kumquats and Bananas

A sophisticated drift chamber hit counting system [51] was used to reject events with insufficient or out-of-time tracks. The hit-counting consisted of two species of FASTBUS modules: the "Kumquat" and "Banana."

The Kumquat was the less complicated of the two modules. Each Kumquat module received cable-delayed inputs from sixteen adjacent drift chamber wires (eight from each of the two overlapping planes), and a 205 nsec gate which was started by the Level 1 trigger. The modules counted the number of adjacent wires with in-time hits (hits which arrived within the gate). An isolated hit wire was counted as one hit. When multiple adjacent wires had in-time hits, the count was the number of adjacent hits minus one. Thus, a hit-pair was considered one hit, and three adjacent hit wires were considered two hits. The criteria for counting hits are shown in Fig. 18.

The hit counts from all the Kumquats in a view of a chamber were added up, and the result was passed to Level 2 as a three-bit number.

Because the in-time gate was much longer than the separation between RF buckets, some out-of-time tracks could form pairs of two in-time hits. The Kumquats could not reject those hits. This duty fell to the Banana modules, which not only latched drift chamber hits, but measured the time of each hit with a 625 MHz TDC. Memory lookup maps for each wire pair were used to classify isolated hits, in-time pairs, and out-of-time pairs.

Similar to the Kumquats' operation, true isolated in-time hits were counted as a

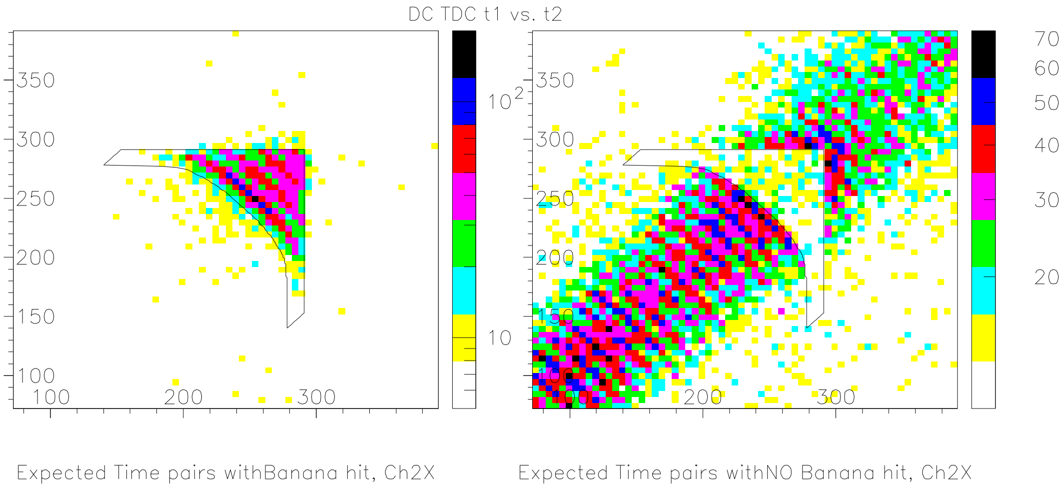

Figure 19. Expected Banana TDC (calculated from Drift chamber TDC) for hit 1 vs. hit 2, for pairs accepted (left) and rejected (right) by the DC2X Bananas. The boundary of the intime-pair allowed region is shown as a black line. The banana-shaped stripes are hit-pairs from different RF buckets.

single hit by the Bananas. Pairs of in-time hits which formed an unphysical sum-of-distances (SOD) were rejected as out-of-time pairs. (See Sec. 6.1.1 for a discussion of the significance of SOD distributions.) When an in-time hit was paired with an out-of-time hit, the pair was only accepted if the difference between the TDC values for the two hits was *greater* than the drift time. This rejected hit-pairs from out-of-time tracks where one hit was within the in-time window, while counting most isolated hits paired with accidental out-of-time hits.

The Kumquats and Bananas completed processing 800 nsec after receiving a Level 1 gate.

The shape of the accepted in-time pair window is shown in Fig. 19. It included the banana-shaped stripe from in-time pairs as well as the triangular low-SOD region. This was done to preserve events with low-SOD pairs. High-SOD pairs were rejected by the Bananas.

The hit-counting configuration changed significantly between the winter and summer runs. In the winter run, the Y-views of all four chambers were instrumented with Kumquats. The X view of Chamber 2 was instrumented with Bananas. The X view of Chamber 1 was instrumented with Bananas for the second half of the winter. The X views of Chambers 3 and 4 were not instrumented with hit counting, except for two Kumquat modules which were used for the Stiff Track Trigger (Sec. 3.2.5). Of the X-view Bananas, only those in Chamber 2 were used for hit counting; the Chamber $1X$ Bananas were only used for the Stiff Track Trigger.

In the summer run, both the X and Y views of Chambers 1 and 2 were fully instrumented with Bananas. The Bananas in the Y view, however, were programmed with large square in-time pair windows, so as to mimic the performance of the Kumquats which had been removed. The Banana maps in the $1X$ and $2X$ views were not changed from the winter run. The X and Y views of Chambers 3 and 4 were fully instrumented with Kumquats. The $1X$, $3X$, $4X$ views were only used for the Stiff Track Trigger.

3.2.3 The Hardware Cluster Counter

The Hardware Cluster Counter (HCC) [50] provided a count of the number of in-time clusters in the CsI calorimeter. The cluster count was based on 3100 ECL discriminated dynode pulses from the calorimeter. The HCC provided a three-bit cluster count plus an overflow bit to Level 2; this allowed triggering on any number of clusters up to 7, or on > 8 clusters.

The HCC received 3100 ECL discriminated signals from the PMT dynode pulses from the calorimeter. The discriminators were tuned to an energy threshold of approximately 1 GeV. Due to the typical transverse energy profile of electromagnetic showers in the CsI, this effectively applied a 2 GeV threshold for electrons and photons. A list of the CsI channels with in-time hits in the HCC was read into the data, and used in the reconstruction.

The HCC was the last Level 2 processor to return a result; its processing time was 2.5 μsec.

3.2.4 The Y-Track Finder

A programmable logic unit was set up to search for hit-patterns in the Y-views of the chambers consistent with tracks in the upper and lower halves of the chambers. The Y-Track Finder (YTF) used the latched trigger wire hit information from the Y-view Kumquats (and, in the summer run, Bananas) as input, and returned four bits of information on the number and position of the hits found. The YTF is described in Ref. [52]. In E799, it was used in a mode where it searched for one track in the upper half and one track in the lower half of the chambers; a track in the beam region counted as both an upper and a lower track.

3.2.5 Stiff Track Trigger

Another programmable logic module was used at Level 2 in the hyperon physics modes to select events with a high momentum track in the beam regions of the chambers. This system used latched information from the Kumquat/Banana system in the beam regions of the X views, and selected events with hit patterns which formed nearly-straight tracks. The Stiff Track Trigger was designed to select tracks with $|p| > 150$ GeV/c. However, as the processor vetoed every event with an extra accidental hit on any beam region wires, its efficiency for signal was below 40%. For the summer run, this processor was reprogrammed simply to require at least one hit in the west beam at each chamber, or at least one hit in the east beam at each chamber.

3.2.6 TRD Trigger

A Level 2 processor used information from the TRD system to select events with electrons in the final state. The algorithm is described in Ref. [53]. The TRD trigger was not used for the analyses described in this thesis because its inefficiency for identifying electrons was difficult to verify. The TRD trigger was more useful in searching for very rare decays which would otherwise have been very difficult to distinguish from K_{e3} decays at trigger level.

3.3 Trigger readout

Information from the trigger control system was read out into the data stream through two readout crates. One crate contained four LRS 3377 TDC modules, which recorded the time when each Level 1 source and each calibration trigger input arrived. The timing of the Level 1 trigger itself with respect to RF was also recorded, as a check on timing within the trigger system.

Another crate contained FNAL pipelined latches, which recorded the Level 1 sources *after* they were synchronized to the RF clock. Latches were also used to record the unprescaled and prescaled Level 1 trigger masks, all the Level 2 sources, and the Level 2 trigger mask.

Also in the latch crate were four-channel 16-bit scaler modules [54]. One module was used to record the spill and event number within each run. This module was only reset once per run. For the event number, which incremented at each Level 2 trigger, two of the four channels in this module were "chained" to form a single 32-bit scaler channel.

Another module was used to record the time of each event within the spill. This "spill clock" was a scaler channel with a 1 kHz clock input derived from the Tevatron RF clock by prescaling it by 53100. The scaler was reset at the beginning of each spill; the spill clock thus read out the time in milliseconds from the beginning of the spill.

Two other scaler modules were reset every event, and counted high-rate detector signals (such as the BA and accidental counters) within a 7.5 μsec gate after the Level 1 trigger. These were implemented in an effort to measure the instantaneous beam intensity in each event, but were untimately not useful; the short gate was required to keep the readout time as short as possible, and no counters had a high enough rate to give statistically useful measurements.

3.4 Filter code

A Level 3 event filter structure was used in the DAQ. Tracks and CsI clusters were reconstructed using slightly simplified versions of the offline analysis code. Events could be selected based on quantities such as number of clusters, charged vertex parameters (number of tracks, invariant mass, vertex position), and particle ID using calorimeter E/p and muon counter hits. Filter routines were written to allow the TRD information to be used at Level 3; no TRD Level 3 cuts were actually implemented online, because electron identification from E/p was sufficient to keep the output data volume acceptable.

Individual physics modes could be prescaled in software at this stage. The system also allowed a trigger-dependent fraction of events to be accepted at random, independent of whether or not they satisfied the normal filter code requirements for that trigger. This was a critical check on the efficiency and operation of the Level 3 software.

In E799-II, about 14% of incoming Level 2 triggers were accepted by Level 3 and written to tape.

3.5 The E799 triggers and destination logic

The E799 triggers are described in this section. The trigger requirements listed below are the ones used in stable running; occasionally, some trigger requirements were dropped when detectors or readouts malfunctioned.

3.5.1 Symbols used for beam triggers

Here we list the symbols used below in the definitions of E799 beam triggers.

Level 1:

- `2V`: Loose VV$'$ requirement: At least 2 hits in one counter bank and at least one in the other.

- `3V`: At least 3 hits in V and 3 hits in V$'$.

- `VEWUD`: One hit in the upper, lower, east, and west halves of either V or V′.
- `ET_THR1-4`: E-total thresholds 1-4 (nominally 11, 18, 25, and 38 GeV).
- `DC12`: At least one paddle hit in all of the DC OR views.
- `2DC12_LOOSE`: At least 2 paddles hit in one view of DC1 OR *and* one view of DC2 OR. Other view of each chamber must have at least 1 paddle hit.
- `2DC12_MED`: At least 2 paddles hit in three of the four DC OR views. Remaining view must have at least 1 hit.
- n`MU3`: At least n counters hit in MU3 X view and n in MU3 Y.
- `!HA`: Veto on the 2.5 MIP threshold on AC-coupled HA sum (winter), DC-coupled sum (summer).
- `!HA_SUPERHI`: Veto on the 7 MIP threshold on AC-coupled HA sum. Winter only.
- `PHVBAR`: No counter above 500 MeV in the RCs, SAs, or CIA.
- `!CA`: No CA counter above 14 GeV.
- `!MU2`: No hits in MU2.

Level 2:

- `2HCY_LOOSE`: Two-track Y view hit counting: at least two hits in each Y view, allowing only one hit in either DC1Y or DC2Y (but not both).
- `1HCY`: At least one hit in each DC Y view.
- `34_HCY`: At least 3 hits each in DC1Y and DC2Y, and at least 4 hits each in DC3Y and DC4Y.
- n`HC2X`: At least n hits in the DC2X Bananas (which used the pair timing requirement shown in Fig. 19.)
- `YTF`: One track candidate in the YTF above and one track below the center (see Sec.3.2.4).

- `HCC_GE`n: At least n clusters in the HCC.
- n`HCC`: Exactly n clusters in the HCC.
- mn`HCC`: m or n clusters in the HCC.
- `TRD_`n`E`: At least n electron candidates in the TRD trigger.
- `STT`: STT observes a stiff track (see Sec. 3.2.5).

3.5.2 Rare decay triggers

E799-II used all sixteen available beam triggers. Of these, the following were primarily for rare decay physics studies:

3.5.2.1 Trigger 1: The 2E-NCLUS trigger

This trigger was used to search for the direct-CP violating mode $K_L \to \pi^0 e^+ e^-$. It was also used to record all decays with at least 4 electromagnetic clusters, at least 2 electrons, and no hadrons in the final state. This included $K_L \to e^+e^-e^+e^-$, and all $K_L \to 3\pi^0$ with either $\pi^0 \to e^+e^-\gamma$, $\pi^0 \to e^+e^-e^+e^-$, or $\pi^0 \to e^+e^-$. The measurement described here uses data from Trigger 1 for both signal and normalization. For most of the data taking, the trigger definition was `2V * ET_THR3 * DC12 * !HA * PHVBAR * !CA * 2HCY_LOOSE * HCC_GE4 * 1HC2X`. At Level 3, events were required to have two electron candidates which formed a vertex, and at least four clusters.

Before run 8577 in the winter, this trigger (along with triggers 3, 8, and 15 – the other pure-electromagnetic triggers) also required no hits in the MU2 bank. This caused a significant loss of signal (5 to 8%) due to the high rates (up to 4 MHz) in this counter bank.

3.5.2.2 Trigger 3: The DALITZ trigger

The DALITZ trigger collected data for a search for $K_L \to \pi^0 \nu \bar{\nu}, \pi^0 \to e^+ e^- \gamma$ [40], as well as $K_L \to e^+ e^- \gamma$ events for normalization. A form factor analysis may be performed on the $K_L \to e^+ e^- \gamma$ mode as well.

This trigger saw some significant changes during the run. The initial trigger definition was `2V * ET_THR2 * DC12 * !HA * PHVBAR * !CA * 1HCY * 1HC2X * 34HCC`. At run 8436, the `TRD_1E` requirement was added to reduce the rate. For the summer run, the E-total requirement was changed to threshold 3 to reduce the rate (this cost about 25% of the $K_L \to \pi^0 \nu \bar{\nu}$ acceptance). Before run 8436, and during subsequent runs when either the TRDs or the TRD trigger were not functioning reliably, the TRD trigger requirement was removed and the Dalitz trigger was prescaled by 3/2 due to its high rate. The Dalitz trigger had the highest rate of any trigger in the experiment over most of the run.

3.5.2.3 Trigger 4: Four-track trigger (4-TRACK)

The 4-TRACK trigger was designed to search for the rare decay $K_L \to \pi^+ \pi^- e^+ e^-$. This data set provided the first observation of this decay mode [55], as well as $K_L \to \pi^+ \pi^- \pi^0, \pi^0 \to e^+ e^- \gamma$ events for normalization. The trigger requirements were `3V * 2DC12_MED * ET_THR1 * !MU2 * PHVBAR * !CA * 34_HCY * HCC_GE2 * YTF_UDO * 3HC2X`. This trigger was dominated at Level 1 by accidental coincidences between out-of-time upstream chamber activity (which fired the DC OR) and either interactions in the TRDs or in-time decays with two tracks plus a photon which converted in the TRDs. The Kumquat hit-counting requirements at Level 2 removed most of these events; the Banana requirement reduced the trigger rate an additional 35%. Note that the muon veto remained in this trigger; although it did cost some signal, it was necessary to keep the rate reasonable while still triggering on pions.

At Level 3, three tracks coming from a common vertex were required.

3.5.2.4 Trigger 5: Two-muon "long-distance" trigger (2MU-LD)

The 2MU-LD (Two-muon "long distance") trigger was set up to look for the direct CP violating mode $K_L \to \pi^0\mu^+\mu^-$, and also accept the muonic Dalitz decay $K_L \to \mu^+\mu^-\gamma$ (hence the "long-distance" trigger name). This trigger required `2V * DC12 * 2MU3 * PHVBAR * 2HCY_LOOSE * HCC_GE1`.

3.5.2.5 Trigger 6: Two-muon "short-distance" trigger (2MU-SD, winter only)

The 2MU-SD trigger was designed exclusively to make a precision measurement of the branching ratio of $K_L \to \mu^+\mu^-$. It required `GATE * 2V * VEWUD * 2MU3 * DC12A * !HA_SUPERHI * 2HCY`. (`VEWUD` was the east-west-up-down requirement in the VV$'$ banks discussed in Sec. 3.1.2.5. Due to the lower-than-expected kaon flux in E799, a competitive measurement was not possible for this mode, and the trigger was dropped for the summer run.

3.5.2.6 Trigger 7: The EMU trigger

E799-II searched for the lepton flavor number violating decays $K_L \to \pi^0 e^\pm \mu^\mp$ and $K_L \to 3\pi^0, \pi^0 \to e^\pm\mu^\mp$ using the EMU trigger. Its Level 1 and Level 2 requirements were `2V * ET_THR2 * DC12 * !HA * 1MU3 * PHVBAR * !CA * 2HCY_LOOSE * HCC_GE3`. At Level 3, a muon hit and an electron candidate were required.

3.5.2.7 Trigger 15: The 2E-2CLUS trigger

An attempt was made to search for the rare decay $K_L \to e^+e^-$ in E799-II, although the kaon flux was not sufficient to make a competitive measurement. The 2E-2CLUS trigger definition was `2V * VEWUD * ET_THR4 * 2DC12_LOOSE * !HA * !CA * PHVBAR * 2HCY * 2HCC * YTF_UDO * TRD_2E`. This trigger was very tight, using the highest E-total threshold, tight DC OR and hit counting requirements, and a tight two-electron requirement in the TRD trigger. This was necessary because the trigger had to allow two-cluster events, which had a very high rate from K_{e3} decays. For runs when the TRD trigger was not functioning properly, this trigger was turned off or heavily prescaled.

3.5.3 Minimum bias and study triggers

The remaining non-hyperon triggers were used for studying systematic errors, obtaining calibration data sets, and collecting large samples of common kaon decay modes:

3.5.3.1 Trigger 2: The two-track trigger (2-TRACK)

Our basic detector calibration and trigger study modes came from the 2-TRACK trigger. This trigger was defined simply as `2V * DC12 * 2HCY_LOOSE`, which was looser than all rare kaon decay triggers. Level 3 accepted any event in this trigger with a two-track vertex candidate. This meant that these events could be used to study all the other trigger elements' performance for the rare decay triggers. The 2-TRACK trigger was prescaled by 500.

This trigger was the source of all common two-track kaon decays. These included K_{e3} events for calorimeter calibration, and $K \to \pi^+\pi^-$ needed for spectrometer alignment and magnet calibration. These two-pion events were also needed for normalization of the $K_L \to \mu^+\mu^-\gamma$ decay. $K_L \to \pi^+\pi^-\pi^0$ and $K_{\mu 3}$ events were collected in this trigger as well. An analysis of the K_{e3} form factors using the 2-TRACK data set is underway.

3.5.3.2 Trigger 6: The 2MU-LOOSE trigger (summer only)

In order to study the muon counter trigger efficiency for $K_L \to \mu^+\mu^-\gamma$ decays, a "loose" dimuon trigger was implemented in the summer run, replacing the 2MU-SD trigger. This was a copy of the 2MU-LD trigger, except that it allowed one hit in either the X or Y view of MU3 to be missing. In order to reduce the rate, E-total threshold 1 was required and the trigger was prescaled by 5.

3.5.3.3 Trigger 8: The NEUTRAL trigger

A neutral trigger, prescaled by 40, was designed to check the efficiency of the 2E-NCLUS trigger using $K_L \to 3\pi^0,\ \pi^0 \to e^+e^-\gamma$ decays. This trigger was defined as `ET_THR3 * !HA * PHVBAR * !CA * HCC_GE5` in the summer run, although the HCC requirement was ≥ 4 or ≥ 6 clusters during parts of the winter run.

Level 3 for the NEUTRAL trigger was the same as for the 2E-NCLUS trigger, accepting any event with at least two electron candidates and at least four hardware clusters. In addition, $K_L \to 3\pi^0 \to 6\gamma$ events reconstructed at Level 3 were written to tape with a software prescale of 10.

3.5.3.4 Trigger 9: The 2V minimum bias trigger (2V-MINBI)

This trigger provided the most unbiased charged track data sample. It required either two hits in the V bank *or* two hits in V$'$. Note that this was looser than the `2V` requirement used in the physics triggers. 2V-MINBI triggers were prescaled by 20000. These events were used mainly for checking the integrity of the trigger systems themselves, rather than studying detector performance (for which the statistics were insufficient).

3.5.3.5 Trigger 13: The E-total minimum bias trigger (ET-MINBI)

The ET-MINBI trigger was a minimum bias sample independent of and complementary to 2V-MINBI. It required only E-total threshold 3, and was prescaled by 20000.

3.5.3.6 Trigger 14: The accidental trigger (ACCID)

Events for use in the Monte Carlo to simulate accidental detector activity (see Sec. 5.4) were collected in the ACCID trigger. This trigger was formed when either the 90° target monitor or the accidental muon counters fired. The accidental trigger was prescaled by 25000 in both the winter and summer runs. It was actually a mistake to keep the prescale the same for the summer run. Because the accidental counters responded to target activity and not neutral beam activity, the ratio of accidental events to recorded kaon decays dropped significantly when the collimator size was increased. This caused there to be fewer accidental events available for summer run Monte Carlo.

3.5.3.7 Trigger 16: The single muon trigger (MUON, winter only)

A trigger requiring only a coincidence between a single V bank hit and a MU2 hit was set up. It was prescaled by 25000, and did not contribute significantly to trigger rates or tape output. No use was ever found for these events. In the summer run, Trigger 16 could not be used due to a hardware failure. This trigger was dropped without hesitation.

3.5.4 Hyperon triggers

Triggers 10, 11, and 12 were dedicated to hyperon physics and acceptance studies for hyperon modes. Trigger 10, the hyperon beta decay trigger (HYPERON-BETA), was designed for a form factor measurement of the decay $\Lambda^0 \to p^+e^-\nu$. This measurement was never performed because the decay was not distinguishable from background. Another hyperon decay, $\Xi^0 \to \Sigma^+e^-\nu$, was observed in the same data set, and the triggers were revised to emphasize that mode. Trigger 11 was a looser trigger, prescaled by 50, which collected normalization modes for the decays in the beta trigger. Finally, Trigger 12 was a "hyperon minimum bias" trigger which required only a coincidence between a V-bank hit and a hole counter hit; Trigger 12 was prescaled by 20000.

Triggers 10 and 11 went through many substantial revisions during the run, the details of which will not be presented here.

3.5.5 Calibration triggers

The following six calibration triggers were used in E799. The rates were negligible for all of them.

1. CsI calibration laser flash. (5 Hz, on- and off-spill)

2. PIN diode pedestal trigger. (< 1 Hz, on and off spill)

3. CsI cosmic ray telescope coincidence. (~ 30 Hz, off-spill only)

4. DPMT pedestal triggers. (3 Hz, on and off spill)

10. FERA pedestal triggers. (off-spill only, every 10 spills)

12. X-ray source trigger for TRD calibration (10 Hz, off-spill)

3.5.6 Trigger rates and dead time

During most of the winter run, the beam intensity was about 4×10^{12} protons per pulse (ppp), and the pulse duration was 19 sec. With the small beams and the standard trigger configuration (after the muon veto was dropped from electromagnetic triggers), approximately 1×10^6 Level 1 triggers and 1.9×10^5 Level 2 triggers were taken per pulse. These numbers are not corrected for dead time. The DAQ typically ran 30% to 40% dead, depending somewhat on total beam rate and instantaneous beam intensity. (The dead time was higher during periods of poor accelerator operation, when the beam was not delivered at a uniform rate throughout the pulse.) Other factors which could affect trigger rate and dead time were the beam targeting angle and the calorimeter PMT high voltage (which could cause the effective E-total and HCC thresholds to change).

After the beam size change and trigger configuration changes in the summer run, the nominal proton beam intensity was lowered to 3.2×10^{12} ppp. With the 40% increase in the beam area, this translated to a 12% increase in the total kaon flux. Because of the trigger changes, we were able to allow this flux increase while keeping the dead time and total Level 1 and Level 2 rates similar to those in the winter. The reconstructed kaon yield per pulse in the summer run was about 10% higher than in the winter run.

The trigger rates in the winter and summer are summarized in Table 5.

Table 5. Level 1 and Level 2 rates of all E799 beam triggers for representative runs in the winter and summer. The figures are corrected for dead time (which was 34% in the winter run shown here and 33% in the summer).

	Winter (run 8597, 3.99×10^{12} ppp)			Summer (run 10703, 3.15×10^{12} ppp)		
Trigger	Prescale	Level 1 per spill	Level 2 per spill	Prescale	Level 1 per spill	Level 2 per spill
2E-NCLUS	1	560 000[a]	52 000	1	651 000[a]	59 000
2-TRACK	500	20 600	16 000	500	24 000	17 800
DALITZ	1	832 000[a]	85 000	1	651 000[a]	60 000
4-TRACK	1	462 000	39 000	1	526 000	53 000
2MU-LD	1	70 000[b]	17 400	1	78 000[b]	19 300
2MU-SD	1	54 000[b]	54 000	—	—	—
2MU-LOOSE	—	—	—	5	10 200[b]	9 600
EMU	1	91 000	11 200	1	106 000	12 400
NEUTRAL	40	52 000	11 000	40	61 000	11 900
2V-MINBI	20 000	1 480	1 480	20 000	1 800	1 800
HYPERON-BETA	2	133 000	9 800	1	215 000	39 104
HYPERON-PPI	50	41 000	2 500	50	23 000	3 900
HYPERON-MINBI	20 000	690	690	20 000	890	890
ET-MINBI	20 000	570	570	20 000	640	640
ACCID	25 000	420	420	25 000	260	260
2E-2CLUS	1	199 000[a]	24 000	1	250 000[a]	16 000
MUON	25 000	360	360	—	—	—
Total		1450 000	284 000		1400 000	269 000

[a,b]These sets of triggers overlapped completely at Level 1.

3.5.7 The online splitting test

For the latter half of the winter E799 run, data were collected using an experimental "online data split." The Level 2 logic had the ability to direct different trigger types to different DAQ planes. Normally, this was only used to direct beam triggers in a round-robin sequence to planes 1, 2, and 4, while calibration triggers and a prescaled set of beam triggers were directed to plane 3.

For runs between 8425 and 8913, different beam triggers were sent to different planes in an effort to reduce the number of tapes necessary for each analysis. The strategy was to separate the triggers into two basic data-stream types: muon triggers and non-muon triggers.

Planes 1 and 2 were used for electromagnetic, four-track, hyperon, and E-total minimum bias triggers. Plane 4 was used for the dimuon and EMU modes. Accidental triggers and the single muon trigger were sent to Plane 3, along with the calibration triggers. The 2-TRACK trigger and 2V minimum bias trigger were distributed among all planes, with 50% to plane 4 and 25% each to planes 1 and 2. This was done so that half of the two-track triggers would be available for normalization each of the two data-stream types.

Trigger overlaps (events where multiple triggers fired) were resolved in the online split using a trigger priority system. Each trigger had a priority assigned, and overlaps were sent to the plane favored for the highest-priority trigger satisfied by the event. The 2V minimum bias trigger had the highest priority, so that its distribution among the planes would not be biased by its overlap with other triggers. The second-highest priority trigger was the 2-TRACK trigger. Next came the dimuon triggers. Following them, the electromagnetic, 4-TRACK, and E-total minimum bias triggers had equal priority. The next priority was the EMU trigger, and the 2E-2CLUS and hyperon triggers had lowest priority. (*i. e.*, in an overlap between a hyperon trigger and a non-hyperon trigger, the event would be sent to the plane favored for the non-hyperon trigger.) The plane logic and trigger priorities are summarized in Table 6.

Although the online split worked, it introduced complications to the data taking and

Table 6. Online split parameters. A larger priority number means that the trigger had *higher* priority.

Trigger	Priority	Favored plane weight		
		1	2	4
2V-MINBI (Trig. 9)	11	25%	25%	50%
2-TRACK (Trig. 2)	9	25%	25%	50%
2MU-LD (Trig. 5)	7	0	0	100%
2MU-SD (Trig. 6)	7	0	0	100%
2E-NCLUS (Trig. 1)	5	50%	50%	0
DALITZ (Trig. 3)	5	50%	50%	0
4-TRACK (Trig. 4)	5	50%	50%	0
NEUTRAL (Trig. 8)	5	50%	50%	0
ET-MINBI (Trig. 13)	5	50%	50%	0
EMU (Trig. 7)	4	0	0	100%
HYP-BETA (Trig. 10)	3	50%	50%	0
HYP-PPI (Trig. 11)	3	50%	50%	0
HYP-MINBI (Trig. 12)	3	50%	50%	0
2E-2CLUS (Trig. 15)	3	50%	50%	0

analysis, and the offline benefits were not substantial. In the end, it was still necessary to split all tapes from a run together. (This was also useful for bookkeeping checks in the offline split; see Sec. 4.2).

The problem of trigger overlaps also made it unattractive to analyze only tapes from the favored planes for a trigger.

In the original design, all trigger overlaps across planes were small: the triggers for the non-muon planes were generally vetoed by muons, which were obviously required for the muon modes. 1/500 of all triggers with two tracks overlapped with the two-track trigger, which had higher priority than any other physics trigger. Therefore, 1/1000 of all two-track physics trigger events were sent to the wrong plane. Of overlaps between other physics triggers, the worst case was the EMU trigger; 3.2% of these events were sent to the wrong plane in the early running. However, after the muon veto was dropped in all electromagnetic triggers due to high accidental losses, the EMU trigger had very large overlap with the electromagnetic triggers. It was sent to the "wrong" planes 61% of the time! While only 0.5-1% of other triggers were misdirected, even this fraction was enough to where it was considered worthwhile to use tapes from all planes for most analyses.

For the summer run, the online splitting was not continued, and all triggers were again sent to all planes in the round-robin scheme.

CHAPTER 4

THE 1997 E799-II DATA SETS

This chapter describes the data sets collected in the two E799-II data runs, and the processing stages between the online data collection and the final reconstruction.

4.1 Winter and summer run ranges

The data used in this measurement came from two runs of E799-II. The "winter run" comprised runs 8245 through 8910; the data were taken between February 1 and March 23, 1997. The "summer run" lasted from July 24 through September 3, 1997 and comprised runs 10463 through 10970.

The chief difference between the winter and summer runs was the beam size (see Sec. 2.1.3). The larger beams used in the summer run provided more kaons per primary proton, and thus allowed us to reduce the primary beam intensity while increasing the kaon yield slightly. This was done because there were more experiments at the lab demanding beam simultaneously in the summer, and the Tevatron could not provide as many protons as requested.

The kaon yield per day (Fig. 20) gives an indication of the reliability of the beam and detector. Most of the long gaps were from problems with the accelerator; specifically, the long shutdowns near the end of the winter run and near the beginning of the summer were due to superconducting magnets in the Tevatron which failed and needed replacement. The gap immediately before the end of the summer run was due to power feeder cable

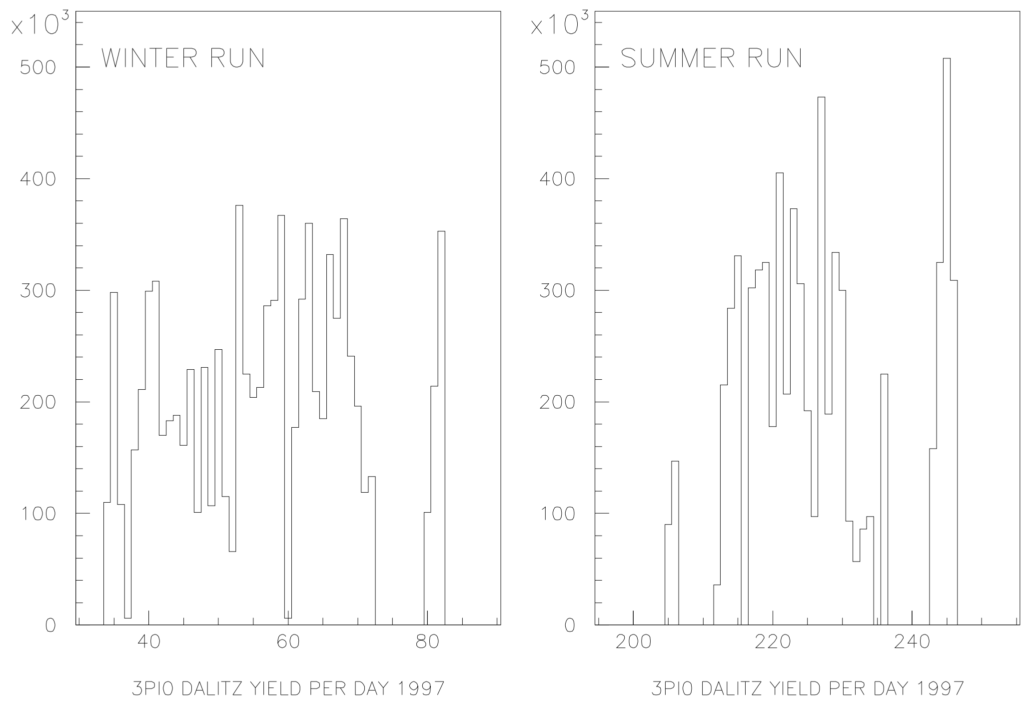

Figure 20. "Online" yield per day of $K_L \to 3\pi^0$, $\pi^0 \to e^+e^-\gamma$ decays in the winter and summer runs (obtained from online monitoring histogram; no $m_{e^+e^-}$ cut).

failures in the Main Ring[1] and some last-minute special runs (including the measurement described in Chap. 7). The very high yield in the last days of the summer run occurred because the Tevatron spill was lengthened from 19 to 23 seconds.

Each run number consisted of data recorded contiguously; typically, runs lasted up to eight hours. A run could be stopped for many reasons; typically, runs ended either when the output tapes were full or when a problem occurred with either the beam or the detector. An effort was made to end runs after eight hours, even in the absence of a known problem, so that monitoring histograms could be reset. About half of usable runs continued until the eight-hour mark or until the tapes were full, while the remainder

[1]The Main Ring was the original high-energy accelerator at Fermilab; it was used in this run to inject 150 GeV protons into the Tevatron. The Main Ring was decommissioned in September 1997, immediately after the data described here were collected.

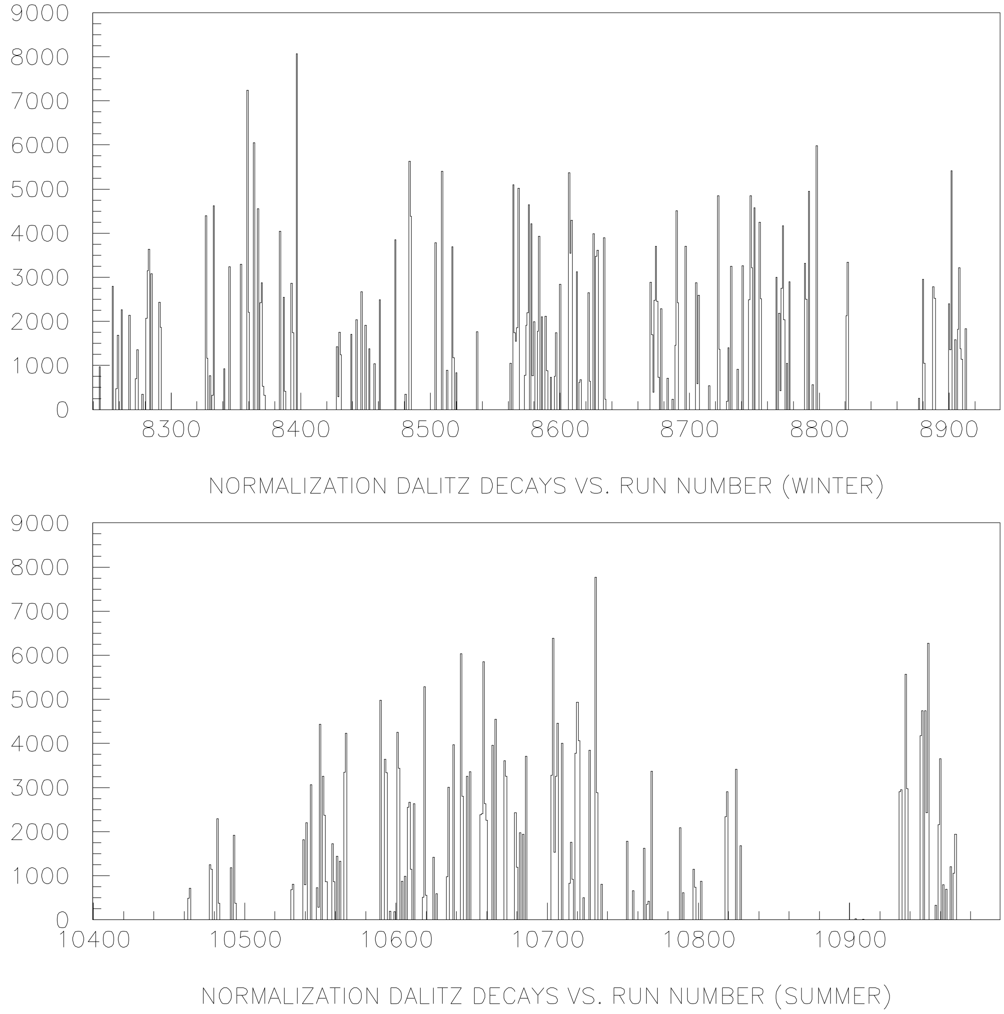

Figure 21. Normalization Dalitz decays per run number for summer and winter data sets.

were stopped because of problems. This is evident from the yield of normalization Dalitz decays per run (Fig. 21). A typical full-length run produced about 4000 normalization decays, although much of the data set is contained in significantly shorter runs. A total of 155 winter runs and 117 summer runs were used in the analysis.

4.2 The E799 data split

After the data were collected, the beam triggers were "split" to 19 different output streams, based on trigger types, Level 3 tags, and, in a few cases, on some simple selection code.

For each run, all the raw data tapes were read in, and the data were split into output streams and "cached" to disk files. For each set of data tapes which were originally written in parallel, all input tapes were read in simultaneously. This provided an important check against incorrect tapes being mounted by the tape operators: all input raw data tapes were expected to begin with the same run number, and all the tapes from that run had to be present.

After the raw data were read in and split, all the disk files for a particular stream were concatenated, and then spooled onto output DLT tapes. Each output tape could contain data from several runs, and a long run could be divided among multiple tapes. For small data streams, where the entire data set only amounted to a few tapes, a single output tape failure would have required a substantial fraction of the split to be re-done. To avoid this, a backup stream was written which duplicated the smallest data sets.

The split streams are described in Table 7. In the winter run, 450 input raw beam trigger tapes produced 494 split output tapes; in the summer, 400 raw beam trigger tapes produced 426 split tapes. The slight increase in tapes was due to writing some events to multiple streams (normalization and calibration modes from the 2-TRACK trigger, as well as the backup stream). All data used in this analysis came from the 2ENEUT stream, which consisted of all 2E-NCLUS and NEUTRAL trigger events.

The calibration triggers (recorded on plane 3) were split separately; each calibration trigger had its own stream.

4.3 The 2E-NCLUS data crunch

The data on the 98 2ENEUT tapes were processed further in a step called the "2E-NCLUS crunch." This step separated the 2E-NCLUS triggered events into nine subsets,

Table 7. Summary of E799 beam trigger data split streams. The 2ENEUT stream provided all data used in the analysis described here.

Stream	Triggers	Tapes out	
		Winter	Summer
2ENEUT	1 + 8	52	46
2TRK	2	82	72
KE3	2[a]	36	35
DALITZ	3 + 1[b] + 2[b]	34	29
4TRK	4	38	44[c]
2MULD	5 + 2[b] + 6[d]	66	64
2MUSD	6 + 2[b]	43	–
EMU	7 + 2[b]	18	15
NEUTRAL	8	10[c]	5
2VMINBI	9	12	13
HYPERON	10 + 11 + 12 + 2[b]	31	60
ETMINBI	13	6	6
ACCID	14	15[e]	3
2E2CLUS	15	4	4
VMU	16	3	–
BACKUP	8 + 13 + 14 + 15 + 16	22	16
PI0EEG	1[b]	4	–
ALIGN	2 ($K_L \to \pi^+\pi^-$ only)	8	7
MUON	ALL (muon runs only)	21	7
Total		505[e]	426

[a] e^+e^- and $\pi^\pm e^\mp$ tags only.
[b] Prescaled or filtered sample for normalization.
[c] Contains significant data from special runs.
[d] Summer data only.
[e] Eleven of these tapes came from the plane 3 (calibration) split.

each of which was used for a different analysis. The crunch was both a split and a filter stage. The code performed a full reconstruction of each events, using offline calibration constants and the algorithms described in Chap. 6. Approximately 70% of the events were thrown out. Many events were written out to more than one stream, as they may have been candidate events in multiple analyses. A tag was written to each event indicating which streams' crunch requirements the event satisfied.

The largest output sample was the "3PI0D" stream, which produced of 14 output tapes. These tapes included all $K_L \to 3\pi^0$, $\pi^0 \to e^+e^-\gamma$ events regardless of $m_{e^+e^-}$. Other streams included $K_L \to \pi^0 e^+e^-$, $K_L \to \pi^0 e^+e^-\gamma$, $K_L \to e^+e^-\gamma\gamma$, and four electron candidates, as well as some smaller samples for systematics studies. A full discussion of the crunch can be found in Ref. [56].

The analyses described here used data from only one stream, the "PI0TEE" stream. In this stream, any event which had at least six clusters and two electron candidates with $m_{e^+e^-} > 70$ MeV$/c^2$ was accepted. An electron candidate was defined as any track with $E/p > 0.9$. This cut could be tighter than the Level 3 cut at 0.75, because the crunch used better calibration and chamber alignment constants than were available at Level 3. The cuts were designed to accept any possible $\pi^0 \to e^+e^-$ candidates, as well as all high-$m_{e^+e^-}$ $\pi^0 \to e^+e^-\gamma$ candidates for the normalization sample.

The PI0TEE crunch sample consisted of four tapes: two from the summer and two from the winter run, containing a total of 6.1 million events.

4.4 Problems in the data

This section describes some of the detector problems encountered during the run, and their disposition offline. Many problems were not deemed serious enough to bias the data, and were ignored in the analysis or treated as systematic errors. Others were considered serious enough to where the data were discarded.

Information on detector problems was stored in a database, where individual spills could be marked with up to 32 "bad spill" tags. Each analysis could, in principle, use a different set of data quality requirements.

4.4.1 Short runs

Runs which produced fewer than 5000 "online" Dalitz decays (identified at Level 3 by loose criteria, and including all values of $m_{e^+e^-}$; equivalent to about 200 normalization events or about 20 minutes of data collection) were not used in the final analysis. These runs amounted to about 1% of the full data set, and nearly all of them were stopped early because of severe problems.

4.4.2 DPMT problems

The CsI calorimeter electronics suffered frequent failures, sometimes several per day. Most of the problems were the result of flaws in the production of the custom integrated circuits in the DPMTs. When the worst varieties of failures occurred, data taking was halted and the offending DPMT was replaced. A few of the more common failure modes are described below.

4.4.2.1 Blown QIE comparators

The comparator circuit on the QIE, which selected the range capacitor for digitization, was subject to failure. When this happened, one range on one of the four rotating capacitor circuits would become inoperative. The next range down was selected instead, and the mantissa registered an overflow. This was not considered a major problem, as it only affected a small fraction of the already small fraction of events where a given crystal was hit. When these problems occurred, data taking continued until a more serious problem required an access to the calorimeter. At that point the DPMT would be replaced. Data taken with blown comparators are used in the analysis.

4.4.2.2 Nonzero pedestal exponent

A "nonzero pedestal exponent" referred to a severe QIE failure where one of the exponent bits became stuck (and hence even pedestal events had a nonzero exponent bit). This caused spurious energy to be seen in the channel. Any channel with this malfunction was swapped immediately, and the spills with this problem were cut from the data sample.

4.4.2.3 Bad/stuck capacitor ID

Another common failure mode was a stuck QIE capacitor. Normally, the signal rotated each bucket among the four QIE capacitors. When the capacitor ID bits became stuck, it was usually because the capacitors had stopped rotating, effectively disabling the channel. This failure was also considered sufficient to halt data-taking and replace the DPMT. Occasionally, a channel's capacitor ID bits became stuck, but the mantissa and exponent data were normal. In this case, the channel was not swapped immediately, and the data from these runs were used in the analysis.

4.4.3 Chamber problems

The most severe problems in the drift chambers were the result of radiation damage, noise, and incorrect gas mixtures and high voltage. These problems varied somewhat in their severity throughout the run, but were always present at some level. These chamber performance issues and their effect on tracking are discussed in Sec. 6.1.1.1 and Sec. 8.5.1.

Certain run-specific problems occurred as well. These included periods of excessive current draw in Chamber 2, voltage trips during high beam intensity, and amplifier oscillations. The last two problems caused all tracks to be unreconstructible. The high current draw had no significant effect on this measurement.

4.4.4 Veto and ADC problems

For much of the winter run, the ADC gate for signals from RC8 was unstable, producing an effective pedestal shift corresponding to ~ 50 MeV. The shifts were observed online, but the cause was unknown and therefore not fixed for a few weeks. While these problems were occurring, RC8 was removed from the trigger requirements due to concerns about our ability to verify its operation. (The pedestal shifts were comparable to the resolution of the counters, so it did not pose a significant problem in the end; events with ADC counts above the nominal trigger theshold were cut offline anyway.)

A more general problem with the FERA ADCs was occasional failure of the pedestal-recording program. This program unsparsified the ADC readouts off-spill, generated pedestal triggers, then resparsified before the next spill began. However, when the computer was heavily loaded, the cycle did not complete before the next spill began, and the sparsification threshold for the ADCs was unknown. As the theshold was subtracted from the raw value to obtain the charge value, the ADC data could not be reconstructed if the threshold was not correct. This affected less than 0.1% of the data, and these spills were removed from the sample.

4.4.5 Kumquat and Banana failures

The Kumquat and Banana readout systems were sometimes unstable, and parts of the system had to be removed from the readout occasionally to be fixed. During these periods, which totaled a few per cent of the data sets, any view which was removed from the readout was also removed from the trigger requirements. This had the effect of slightly ($\sim 1\%$) increasing the dead time.

CHAPTER 5

THE MONTE CARLO SIMULATION

The Monte Carlo was a critical component of the analysis: it was used to determine the acceptance ratio for signal and normalization modes, as well as to estimate the size of the background. Here, the simulation is described from the production of kaons in the target, through the decays of the kaons and π^0's, the tracing of the decay products through the detector, the simulation of accidental detector activity, and the trigger and detector performance.

The Monte Carlo had five basic steps: Kaon production, decay generation, particle tracing, digitization, and trigger selection. The output of the software was an event written in the same format as the real data. This made it easy to apply the same analysis to data and Monte Carlo and make detailed comparisons.

5.1 Kaon production and decay position

The Monte Carlo produced kaons at positions within the target determined by the interaction depth and a transverse smearing parameter; the profiles are shown in Fig. 22. The kaon momentum and angle distribution were selected according to a parameterization by Malensek [57] of measured K^+ and K^- production by 450 GeV protons on a beryllium target.

The production spectrum for neutral kaons was arrived at by breaking down the charged kaon cross-sections into relative production probabilities σ_u, σ_d, and σ_s for $u\bar{u}$,

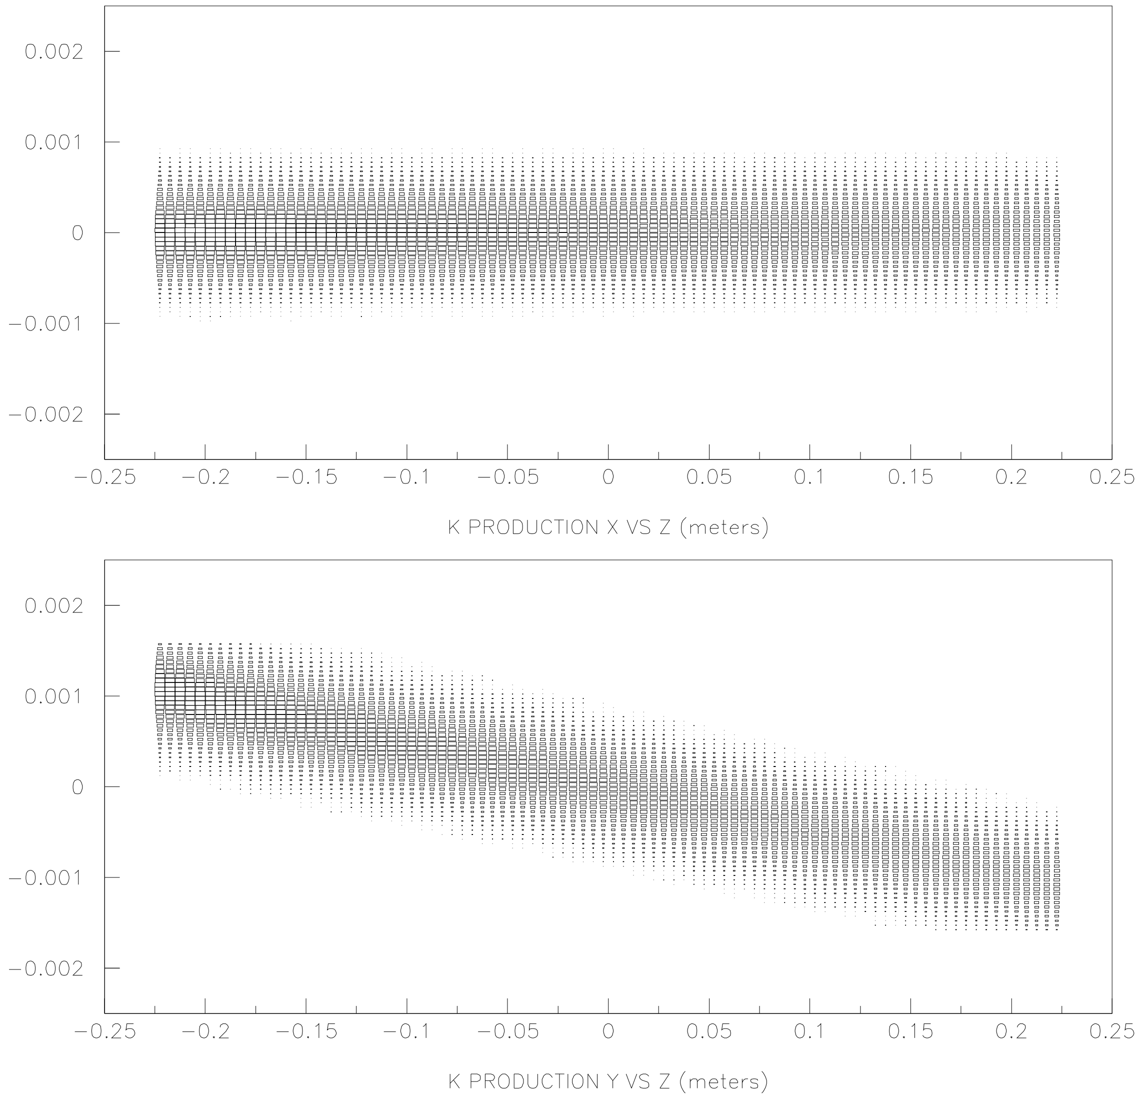

Figure 22. Monte Carlo kaon production position profiles within the target. The slope in the Y view is from the 4.8 mrad targeting angle. The vertical scales are expanded.

$d\bar{d}$, and $s\bar{s}$ quark pairs. It was assumed that $\sigma_u = \sigma_d$. We considered production of a K^+ to entail *either* the creation of both a $s\bar{s}$ pair and a $u\bar{u}$ pair $(\sigma_s\sigma_u)$, *or* the creation of a $s\bar{s}$ pair and use of one of the two valence u quarks in the proton ($2\sigma_s$ or one "chance" per valence u quark). Thus the K^+ production was proportional to $\sigma_s\sigma_u + 2\sigma_s$. Similarly:

$$\sigma(K^+) \sim \sigma_s\sigma_u + 2\sigma_s \tag{5.1}$$

$$\sigma(K^-) \sim \sigma_s\sigma_u \tag{5.2}$$

$$\sigma(K^0) \sim \sigma_s + \sigma_u\sigma_s = \frac{\sigma(K^+) + \sigma(K^-)}{2} \tag{5.3}$$

$$\sigma(\bar{K}^0) \sim \sigma_s\sigma_u = \sigma(K^-) \tag{5.4}$$

This provided a first approximation of the K^0 and $\bar{K}^0$ production spectra. The spectrum was tuned further, using a third-order polynomial fit in momentum to match the observed $K_{L,S} \to \pi^+\pi^-$ spectrum in the data. A more complete discussion of K^0 and $\bar{K}^0$ production is presented in [38], and the numerical details of the kaon production parameterization used in the E799-II Monte Carlo are given in [40].

After the momentum and angle of the kaon were selected, the kaon was traced to the vacuum decay region. If the kaon hit any collimator along the way, the event was rejected and a new kaon was generated. Simulation of scattering of kaons in the final collimator was not implemented for the Monte Carlo generation described here. (Collimator scattering was a very small effect in $K_L \to 3\pi^0$ decays, but a much larger effect in two-pion decays, because a scattered K_L could become a K_S, which would usually decay rapidly to two pions.)

For $K_L \to 3\pi^0$ decays, the interference between K_L and K_S was taken into account when selecting the decay position. This was not actually necessary: because of the small $K_S \to 3\pi^0$ branching ratio ($< 3.7 \times 10^{-5}$ [33]) and large distance between the target and the decay volume, the Z-distribution was indistinguishable from that obtained using a pure K_L exponential lifetime. The generated kaon momentum and K_L decay position are shown in Fig. 23. Only about 5% of the K_L which entered the vacuum decay region decayed before leaving it; the remainder either decayed in the spectrometer or struck the BA.

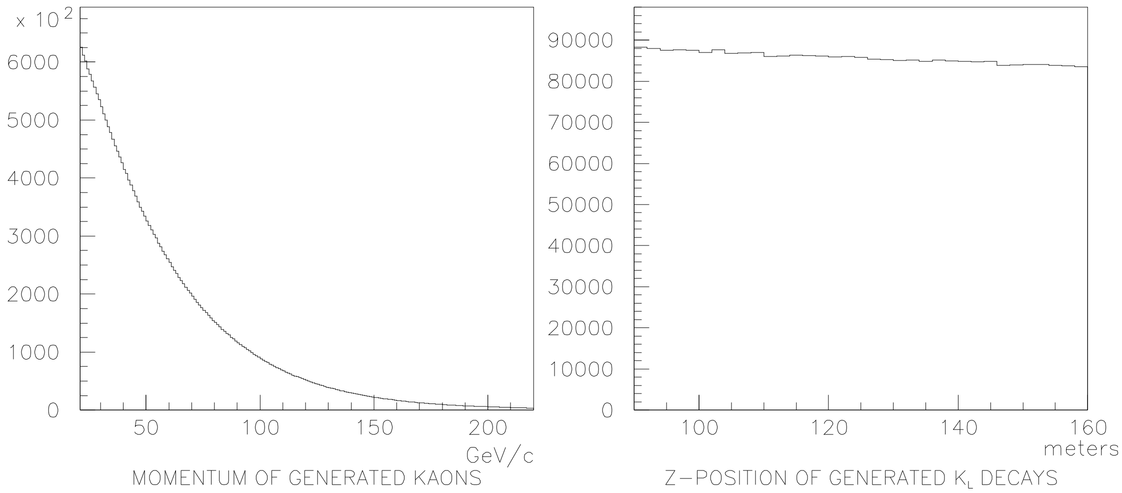

Figure 23. Momentum and K_L decay Z-distribution of Monte Carlo generated kaons.

5.2 Decay Generators

The $K_L \to 3\pi^0$ decay was simulated with a pure phase space distribution. Although E731 has measured a form factor [58], its significance was only 3σ using a sample of 5×10^6 reconstructed decays, an order of magnitude more than the normalization sample here. We did not expect any sensitivity to the form factor.

The π^0 Dalitz decays for normalization and background studies were generated using the form factor measured by the CELLO collaboration [10] and $\mathcal{O}(\alpha_{\mathrm{EM}}^2)$ radiative corrections based on calculations by Mikaelian and Smith [59] [60]. These calculations included inner bremsstrahlung corrections ($\pi^0 \to e^+e^-\gamma\gamma$) and virtual-photon corrections to the $m_{e^+e^-}$ spectrum. The bremsstrahlung correction was implemented with a cutoff of $m_{\gamma\gamma} > 1$ MeV/c^2.

In the simulation of the normalization mode, only the portion of the spectrum with $m_{e^+e^-} > 65$ MeV/c^2 was generated. This $m_{e^+e^-} > 65$ MeV/c^2 tail comprises 3.191% of the total Dalitz decay rate.

For $\pi^0 \to e^+e^-$, the Monte Carlo used the radiative correction model of Bergström

[25] for the inner bremsstrahlung spectrum. The MC simulated the $m_{e^+e^-}$ spectrum up to a cutoff at $(m_{e^+e^-}/m_{\pi^0})^2 = 0.999252$. With this cutoff, 30% of $\pi^0 \to e^+e^-$ events radiated a photon, although most of the photons were very soft and did not affect the event kinematics significantly. We calculated only the distribution of $m_{e^+e^-}$, and not the photon emission angle. In the MC, the angle was assumed to be zero (*i. e.* the photon was collinear with one of the electrons). It should be noted that events with a detectable photon did not pass the signal cuts, and thus the exact photon direction was not relevant.

5.3 Tracing of decay products

After a kaon decay was simulated, the π^0's were assumed to decay immediately upon creation. The decay products were boosted to the lab frame and traced through the detector. All π^0 decay products are stable, so no tertiary decays needed to be considered (although for modes with charged pions in the final state, pion decays to muons were simulated). When particles were traced through active detector elements, the performance of those detectors was simulated. Charged particle behavior in magnetic fields was also simulated.

Certain interactions between particles and detector material were included as well. Charged particles traversing material were allowed to undergo multiple scattering according to a parameterization of the Molière theory [61], which accounted for both a Gaussian scattering angle distribution and a nongaussian single-scattering tail. Electrons, in addition to scattering, were allowed to radiate bremsstrahlung photons according to the Bethe-Heitler cross-section [62]; radiated photons were generated collinear with the electron. Photons were allowed to convert in material with probability $1 - e^{-\frac{7}{9}(X/X_0)}$, where X/X_0 is the fraction of a radiation length the particle traversed. Note that the conversion probability reduces to $\frac{7}{9}X/X_0$ for $X \ll X_0$. The conversion spectrum was taken from the Bethe-Heitler formula, and the pair opening angle distribution from routines in the EGS4 software library [63].

Bremsstrahlung photons and conversion electrons were traced through the detector in the same manner as the decay products.

5.3.1 Final Sweeper simulation

Modeling of the final sweeper magnet was important, because the downstream end of the field defined the upstream end of the fiducial volume of the decay region, and backgrounds could come from decays distorted in the fringe field.

The final sweeper field and beam pipe aperture were simulated using measured field maps and pipe geometry. This simulation was, of course, only applied to decays which occured upstream of the end of the fringe field at $Z = 94.08$ m. The magnet was broken up into 113 $1''$ "slices" in the longitudinal direction, and a transverse momentum kick applied to the particle at each slice. The field was assumed to be purely in the Y direction, and variation of the field in X and Y was not simulated.

The beam pipe was simulated as an aperture at $Z = 93.99$ m, the downstream end of the magnet assembly where the two beam pipes opened into a single vacuum region. Photons and electrons were stopped completely if they hit this aperture; other charged particles were scattered through 3 radiation lengths. Neutral hadrons were assumed to pass through without interacting.

5.3.2 Photon vetoes

Whenever a photon or electron encounted a ring-shaped photon veto detector (RCs, SAs, or CIA), the particle was stopped and all its energy deposited in the detector, with the energy smeared by a Gaussian to simulate the resolution of the detector. For charged pions and muons, the particle was scattered through 16 radiation lengths and a minimum ionizing particle's energy loss was deposited in the detector. If a particle passed outside the outer edge of the veto, the particle was considered lost and not traced further.

The CA and BA were treated in a similar way, although multiple scattering of pions was not done (this was unnecessary due to the far downstream positions of these detectors). Splashback from the CA and BA were not simulated.

5.3.3 Charged spectrometer simulation

5.3.3.1 Passage of particles through the spectrometer

Material in the spectrometer was compressed into planes at certain Z positions, and all classes of interactions were considered to occur only at these planes. For multiple scattering and bremsstrahlung, this was a reasonable approximation; for photon conversions, care was taken not to combine material across an upstream drift chamber or the magnet, as the conversion tracks might not then produce hits in the drift chambers at the correct places. The planes of material and their Z positions are given in Table 8.

The material from the vacuum window to the midplane of Chamber 1 was critical for modeling photon conversions, as any conversion in this region could produce reconstructible tracks in the chambers. The integrated material in this region was measured in a special run, which is described in Chapter 7. The material was broken into three planes in the Monte Carlo: the window itself, an air gap downstream of the window, and the upstream half of the drift chamber (this material was placed at the upstream window of the chamber). The calculated radiation thickness of the vacuum window and the upstream half of the chamber were used for those planes, and the balance of the material measured was assigned to the air gap.

Each drift chamber was modeled as two planes of material, one at the upstream face of the chamber and one at the downstream face. The calculated radiation thickness of the chamber, including all windows and gas, was divided evenly between the two planes. The chamber wires were modeled separately as square in cross-section; each particle was assigned a uniform probability of hitting a wire, independent of the actual transverse position of the particle at the chamber. All scattering, bremsstrahlung, and conversions assumed the wires were at the upstream end of the chamber.

The helium bags upstream of Chamber 1 were included in the "air gap" material. The

Table 8. Planes of material simulated in the charged spectrometer. The helium bag regions are each divided into four equally spaced planes.

Material simulated	Z position (m)	Rad. thickness (X_0)
Vacuum window	158.89	0.00156
"Air gap"	159.25	0.00147
Drift Chamber 1	159.40	0.00050
	159.48	0.00045
Helium Bag 2	159.48 to 165.55	0.00162
Drift Chamber 2	165.55	0.00045
	165.63	0.00045
Helium Bag 3 before magnet	165.63 to 170.01	0.00119
Helium Bag 3 after magnet	170.01 to 173.99	0.00119
Drift Chamber 3	174.65	0.00090
Helium Bag 4	174.65 to 180.02	0.00156
Drift Chamber 4	180.55	0.00090

large helium bags between the drift chambers were each treated as four planes of material; the radiation thickness was calculated assuming the bags were 97% helium and 3% air. These numbers were determined from a study of the matching of upstream to downstream track segments at the magnet plane for very low momentum tracks (< 5 GeV/c), where multiple scattering dominated the tracking resolution.

Finally, the spectrometer magnet was modeled simply by giving each charged particle a transverse momentum kick as it crossed the midplane of the magnet at $Z = 170$ meters. The kick was not dependent on the transverse position of the particle.

5.3.3.2 Drift chamber response

When a charged particle crossed a drift chamber, the transverse position of the particle at the chamber was saved, and the distance between the track and the closest wire in each plane was calculated. A drift time, based on calibrations from the data, was calculated and smeared to simulate 100μm resolution. This time was converted to TDC counts using the offsets from the data calibration. Only the first hit on a wire from a track was simulated.

Low-SOD pairs from δ rays in the chamber cell were simulated, as were multiple tracks in the same cell (the first hit from *either* track on a given wire was recorded). A drift chamber hit inefficiency could be implemented as well, although it was not used for the E799 Monte Carlo. Additional resolution smearing due to effective discrete ionization was simulated in a position-dependent way, which allowed some of the effect of radiation damage in the beam regions of the chamber to be reproduced.

5.3.4 TRDs and trigger hodoscopes

Downstream of Chamber 4, particles were traced through the TRD system. Outside the beam region, where the particle traversed the TRD radiator and the chamber, the TRD system was treated as 0.14 X_0; inside the beams, where there was no radiator, the system was treated as 0.04 X_0. Photon conversions in the TRDs were simulated (due to the large amount of material, these conversions were a major source of hits in the VV$'$ counters). Scattering off the TRD wires was not simulated.

After the TRD system, the particles were propagated to the VV$'$ banks. If a charged particle crossed a counter, a pulse was generated for that counter, and a hit multiplicity count for the trigger was incremented. No counter inefficiencies were simulated (the actual counter inefficiencies were below 0.2%). Particles which hit the VV$'$ hodoscopes were scattered through 0.02 radiation lengths; photons which converted in a counter generated hits in that counter.

5.3.5 CsI calorimeter simulation

When photons, electrons, or pions were propagated to the CsI, if they hit the front face of the calorimeter they were stopped and their showers were simulated. Muons, which only deposited minimum ionizing energy in the calorimeter, were allowed to pass through the calorimeter and a minimum ionizing energy deposit was simulated.

For particles which showered in the CsI, the transverse position of the shower was calculated by extrapolating the particle direction to the depth of the shower mean. For electrons, the depth was $Z_{eff;e} = 0.11 + 0.018 \ln E$, with E in GeV and Z in meters downstream of the front face of the crystal. For photons the depth was $Z_{eff;\gamma} = 0.12 + 0.018 \ln E$, taking into account the ~ 1 cm photons typically traveled in the crystal before converting.

Pion showers were simulated using a shower library taken from the data. Energy in the hadron veto from pion showers was also generated using the pion shower library information.

Photon and electron showers were fully contained in the CsI. The showers in the simulation were selected randomly from libraries of electromagnetic showers previously generated by GEANT code. Each shower in the library contained energies for a square region consisting of 13×13 small (2.5 cm square) crystals. Large crystals were treated as the sum of four small crystals.

Electromagnetic showers were indexed by a position bin number which identified the transverse location of the shower center within the crystal. There were 50 divisions in X and Y across an entire crystal. The bin numbers repeated in each octant of the crystal, so there were only 325 distinct shower position bin numbers. The bins had nonuniform widths, ranging from 0.7 mm in the center of the crystal to 0.2 mm along the crystal boundary. The binning pattern was selected because the distribution of energy among the central crystal and its nearest neighbors was a much more sensitive function of position when the incident particle struck near a crystal boundary. Fig. 24 shows the bin boundaries within a small crystal. Showers were generated for 2, 4, 8, 16, 32, and 64 GeV, and the generated showers were scaled to the appropriate cluster energy in each event.

The showers were segmented longitudinally into 25 slices. To generate the light output for a crystal, the GEANT shower was convolved with any known longitudinally nonuniform CsI behavior.

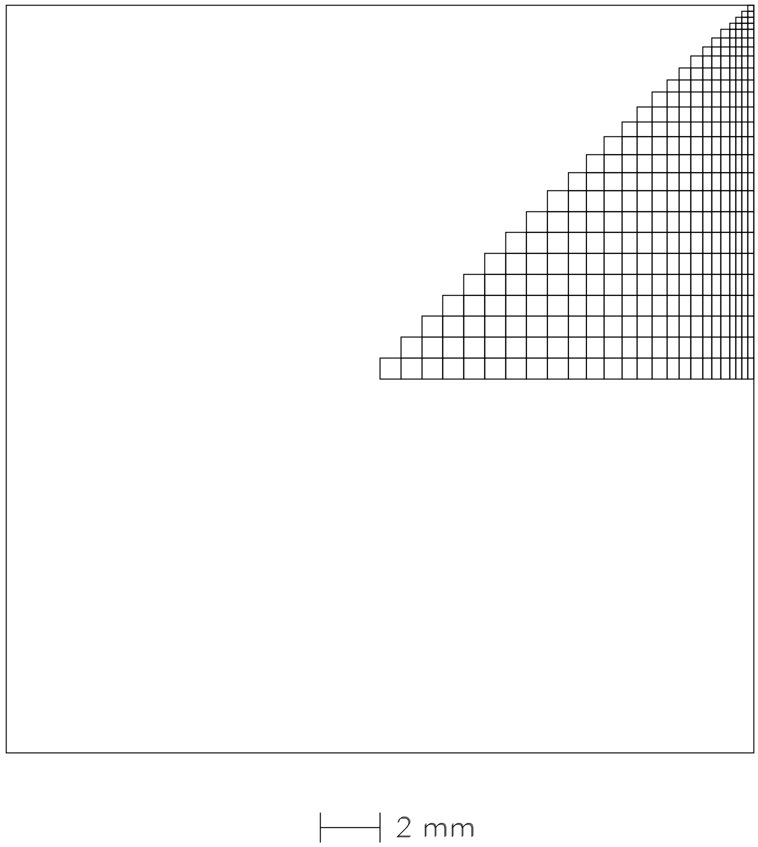

Figure 24. Binning of transverse shower position within a small CsI crystal for showers in the Monte Carlo libraries.

All CsI showers were smeared by a gaussian "resolution fudge" at generation, to account for the fact that the measured resolution in data was somewhat worse than in the GEANT MC. In E799, this fudge was insufficient, so the Monte Carlo cluster energies were smeared again at the reconstruction stage.

After the total light deposited in each crystal had been determined from the selected shower, the time structure of the PMT pulse and the digitization in the DPMT were simulated, with appropriate smearing to reproduce photostatistics effects. The fraction of the total charge which was assigned to each DPMT time slice was taken from a lookup table derived from the shapes of pulses in the data.

5.4 Accidentals and calibration constants

Accidental activity in the detector was simulated by overlaying detector information from accidental-trigger events onto the simulated quantities. Drift chamber TDC hits, veto and trigger hodoscope ADC counts, latch hits, and CsI energies from the accidental event were all added to the Monte Carlo generated event *before* readout thresholds were simulated. This was done in order to reproduce the cases where the "signal" quantities would have been below threshold, but extra accidental energy boosted them above threshold. This was of largest concern in the CsI, where threshold corrections played a major role in the reconstruction. (For this reason, accidental triggers were taken with a near-negligible 2-ADC-count readout threshold.) The FERA ADC thresholds could not be relaxed for recording accidental events, so overlaps between simulated and accidental energies were not necessarily accounted for properly.

In the drift chamber system, any accidental hit less than 50 nsec before a Monte Carlo hit on the *same* wire caused the Monte Carlo hit to be deleted. This was done to simulate the effect of the 50 nsec drift chamber discriminator dead time after a hit.

In the electromagnetic decay modes described here, the signal particles could not produce hits in the hadron veto or the muon counters. Accidental overlays were the only possible source of hits in these systems.

Only accidental events triggered by the 90° target monitor were used in the overlays. The summer and winter run Monte Carlo were generated separately, and accidental events from the appropriate run were used in generating each sample. The run number stamped on each Monte Carlo event was obtained from the accidental event, and this also determined all calibration constants used in MC generation. Events were generated over the entire data sample by weighting generation of events from each run by the total number of accidental events recorded that run. This number was corrected for a change in the accidental trigger prescale, but no "fine-tuning" was done to make the run number distributions in reconstructed data and Monte Carlo events agree perfectly. There were discrepancies in this ratio of up to 20%, mostly due to apparent variations in the efficiency of the accidental counters with beam intensity. As there were no run number-dependent relative acceptance changes between the signal and normalization MC, this variation did not affect the final branching ratio measurement.

5.5 Simulation of the trigger

The trigger was simulated in the Monte Carlo using the same pattern of trigger sources and lookup maps as implemented in hardware. This allowed the Monte Carlo to use the same trigger definition files which were used online in the experiment.

Most of the Level 1 sources were simulated as simple thresholds on counter energies, including both Monte Carlo and accidental energy. The DC OR sources were simulated from the drift chamber hits, assuming no inefficiency. The E-total simulation was somewhat more sophisticated; the threshold was nonsharp and inefficiency at higher energy was simulated as well. The hadron veto simulation summed the Monte Carlo energy (which included pion showers in modes where pions were generated) and accidental energy over all the counters, and compared the result to a threshold. Variation of the threshold with rate was not simulated.

The Level 2 processors were simulated with the exact algorithms used by the hardware, acting on inputs from the Monte Carlo events.

The HCC inputs were modeled from the generated crystal energies. Each crystal's energy threshold and inefficiency were measured in the data, and the nonsharp threshold behavior was parameterized. These constants were used in the Monte Carlo generation of the input bits from the crystal energies, although time-dependent variations over the course of the run were not simulated. HCC input bits from the accidental event were overlaid as well.

In the Kumquats, accidental and Monte Carlo hits in the drift chamber TDCs which fell within the Kumquat gate were added together, and the Kumquats' hit counting algorithm operated on the resulting hit pattern to obtain the Level 2 hit counting sources.

In the Bananas, the drift chamber TDCs were used to generate Banana TDC times. The times were smeared by a Gaussian to take into account the lower resolution of the Banana TDCs. The Banana times were referenced to the Banana lookup maps to generate a list of wires with "valid" Banana hits and the hit counting sources for Level 2. It should be noted that the Banana TDCs were able to observe hits over a larger time range than the drift chamber TDCs. Far-out-of-time hits observed in the Bananas could be used to veto in-time hits on adjacent wires if the two hits formed an out-of-time pair. Because all Banana times in the Monte Carlo were calculated from drift chamber TDC counts, these far-out-of-time hits (which could only come from accidentals) were not simulated. In all modes (except possibly the hyperon triggers), the Banana hit requirements were so loose that these effects were not significant.

The Kumquat and Banana hits were also used to simulate the Y Track Finder and Stiff Track Trigger behavior, as these processors used the hit counting boards for input.

5.6 Monte Carlo samples for this analysis

The background and normalization samples were initially generated to correspond to roughly twice the data. Data were simulated equivalent to 6.3×10^{10} $K_L \to 3\pi^0$ decays for the winter run and 5.5×10^{10} for the summer.

5.6.1 Signal and normalization

The normalization mode, $\pi^0 \to e^+e^-\gamma$, was the largest single sample generated. In order to make detailed data/MC comparisons, over 130 million of these events were generated, of which 10% passed the trigger requirements and were written out for analysis. Only decays with $m_{e^+e^-} > 0.65$ were generated with full statistics, although certain background subsamples with the full $m_{e^+e^-}$ distribution were generated as well.

For the signal mode, 3.00×10^6 $\pi^0 \to e^+e^-$ decays were generated for the winter, and 2.63×10^6 for the summer. The decays were generated with all possible values of $m_{e^+e^-}$.

5.6.2 Backgrounds

All significant backgrounds to $\pi^0 \to e^+e^-$ came from $K_L \to 3\pi^0$ decays. The backgrounds fell into three categories:

- High-mass Dalitz decays: A $\pi^0 \to e^+e^-\gamma$ decay with high $m_{e^+e^-}$ and a soft photon could be misreconstructed as a $\pi^0 \to e^+e^-$ event. These events are from the same decay mode as the normalization, so they were estimated using the normalization MC sample. A separate background sample was not generated.

- "Correctly-paired" 4-track background: These events had four electrons in the final state, but one track of each sign was soft, and was not reconstructed. The remaining tracks were two opposite-sign electrons, and they could reconstruct as $\pi^0 \to e^+e^-$. The three sources of this background were $\pi^0 \to e^+e^-e^+e^-$, $\pi^0 \to e^+e^-\gamma$ where the photon converted at the vacuum window, and $\pi^0 \to \gamma\gamma$ where both photons converted at the window. In all three of these cases, the two electrons observed both came from the *same* π^0, and the four photons observed were correctly reconstructed as two $\pi^0 \to \gamma\gamma$ decays. For the Dalitz decay with the photon converting, only events with $m_{e^+e^-} < 65$ MeV/c^2 were generated, as the higher-$m_{e^+e^-}$ events were already generated as part of the Dalitz background/normalization sample, which allowed photon conversions.

- "Mis-paired" 4-track background: These events also had four electrons in the final state, with only one track of each sign reconstructed. The two observed electrons in this case came from *different* π^0's, and the four observed photons were misreconstructed as two $\pi^0 \to \gamma\gamma$ decays. The sources of this background were $K_L \to 3\pi^0$ where two π^0's Dalitz decay, $\pi^0 \to e^+e^-\gamma$ where a photon from a different π^0 converted, and $K_L \to 3\pi^0 \to 6\gamma$ where two photons from different π^0's converted.

The number of background events generated for each background species is shown in Table 9. Events were generated corresponding to approximately twice the data for both the summer and winter runs.

Table 9. Effective branching ratio and number of MC generated events for contributing $K_L \to 3\pi^0$ background modes. The "branching ratio" is the fraction of $K_L \to 3\pi^0$ events which produced each final state. γ_{ee} indicates a photon which converted at the vacuum window. $m_{e^+e^-}$ cutoffs shown for the Dalitz decays are in MeV/c^2.

Final state			Branching	Generated	Generated
π^0 #1	π^0 #2	π^0 #3	Ratio	(winter)	(summer)
High-x Dalitz decays:					
$\gamma\gamma$	$\gamma\gamma$	$e^+e^-\gamma$ $m_{e^+e^-} > 65$	1.12×10^{-3}	7.047×10^7	6.166×10^7
Correctly-paired 4-track backgrounds:					
$\gamma\gamma$	$\gamma\gamma$	$e^+e^-e^+e^-$	9.2×10^{-5}	5.8×10^6	5.1×10^6
$\gamma\gamma$	$\gamma\gamma$	$e^+e^-\gamma_{ee}$ $m_{e^+e^-} < 65$	9.3×10^{-5}	5.8×10^6	5.1×10^6
$\gamma\gamma$	$\gamma\gamma$	$\gamma_{ee}\gamma_{ee}$	2.2×10^{-5}	1.34×10^6	1.17×10^6
Mispaired 4-track backgrounds:					
$\gamma\gamma$	$e^+e^-\gamma$	$e^+e^-\gamma$	4.3×10^{-4}	2.7×10^7	2.7×10^7
$\gamma\gamma$	$\gamma\gamma_{ee}$	$e^+e^-\gamma$	3.7×10^{-4}	2.3×10^7	2.0×10^7
$\gamma\gamma$	$\gamma\gamma_{ee}$	$\gamma\gamma_{ee}$	8.8×10^{-5}	5.4×10^6	4.7×10^6

CHAPTER 6

TRACK AND CLUSTER RECONSTRUCTION

This chapter describes the procedures used to reconstruct decays. The analysis of each event consisted of two largely-independent algorithms. The charged track-finding used drift chamber hit information to reconstruct the trajectories and momenta of charged particles. The CsI clustering algorithm reconstructed the energies of clusters in the calorimeter, providing electron identification and photon energy measurements.

Each event was reconstructed in three steps: First, track candidates in the X and Y views were found. If at least two track candidates in each view were not found, the event was rejected and processing stopped. After track candidates were found, the calorimeter data were examined and energy clusters were reconstructed. After clustering, corrections were applied to the tracks and kinematic quantities could be calculated.

6.1 Finding track candidates

Track candidates were found in the first stage of the "T3" tracking algorithm [64]. This stage used only information from the drift chambers, and formed track candidates independently in the X and Y views of the chamber system.

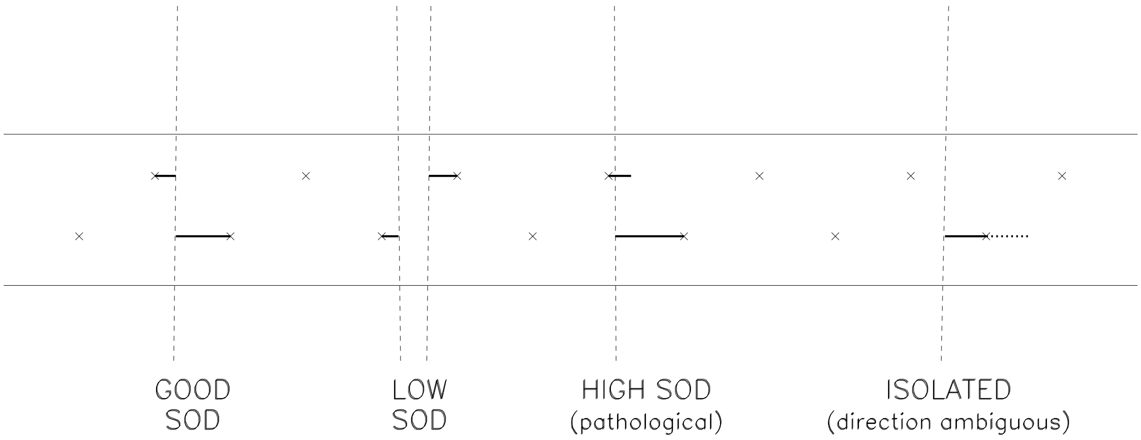

Figure 25. Graphical illustration of the classes of drift chamber hit "pairs." Small crosses indicate the location of sense wires in the two overlapping planes. The dashed lines indicate the paths of charged particles; the bold lines indicate the drift distances reconstructed from the hit times. The drift direction for an isolated hit was ambiguous; a dotted line shows the alternative reconstruction.

6.1.1 Drift distances, hit pairs, and SODs

The first stage of tracking was the translation of raw drift chamber TDC hits into drift distances from wires. This was accomplished by subtracting a wire-dependent time offset from the TDC value, and then using a lookup table to translate the time to a drift distance. Hits outside of a time window 170 nsec wide were considered out of time, and no drift distance was calculated. Although the TDCs recorded multiple hits on a wire, the tracking code used only the first in-time hit on each wire.

Second, a list of "pairs" was formed. A pair consisted of either an isolated hit or a set of hits on adjacent wires in the two complementary planes of a chamber. Pairs were classified according to their "sum of distances" (SOD).

When a single charged particle crossed a drift chamber, it usually left a hit on one wire of each of the two overlapping planes in a view. The drift distances calculated from the two hits should add to the transverse spacing between the wires (6.35 mm), within resolution. Drift distance resolutions were typically ~ 100 μm, so the resolution on the SOD was ~ 140 μm. Pairs where the SOD was within 1 mm of the nominal wire spacing were considered "good-SOD" pairs. Pairs outside this range were considered either low-or high-SOD. The various classes of hit-pair are shown in Fig. 25. It was important to require good-SOD pairs on each track, to reject out-of-time tracks which typically had high or low SOD values.

6.1.1.1 Low and high SODs

Certain mechanisms could produce low or high SOD values for in-time tracks as well. If two tracks were in the same cell, a single low-SOD pair would be produced. Also, a single track which emitted a δ-ray would also produce a low-SOD pair. Inefficiency for observing a hit caused some tracks to leave only an isolated hit in one of the overlapping planes. An isolated hit only carried information on the distance from the track to the wire, and not its direction. Therefore, isolated hits could be used to form a track on either side of the wire.

Finally, pathological behavior in the chambers could cause an in-time track to produce a high-SOD pair. This was a major problem with the data collected in E799-II (and in E832); it resulted from a combination of radiation damage and noise problems in the drift chambers. The chamber gain was too low by of order a factor of two, and the efficiency for observing the first drift electron from a track was significantly below 100% (although the actual inefficiency was unknown). This caused the resolution to be somewhat poorer than in past experiments which used the same chambers (E773 achieved 80 μm resolution in the 1991 run). In addition, there was a large high-side tail in the SOD distribution which was not present in previous experiments. SOD distributions for E832 and E773 are shown in Fig. 26.

In the tracking code, each pair was assigned a "quality value." A good-SOD pair had a quality of 4; each hit in a low-SOD pair had a quality of 2; isolated hits and hits in high-SOD pairs had a quality of 1. For purposes of tracking, at this point each individual hit in a low-SOD pair was considered a separate "pair."

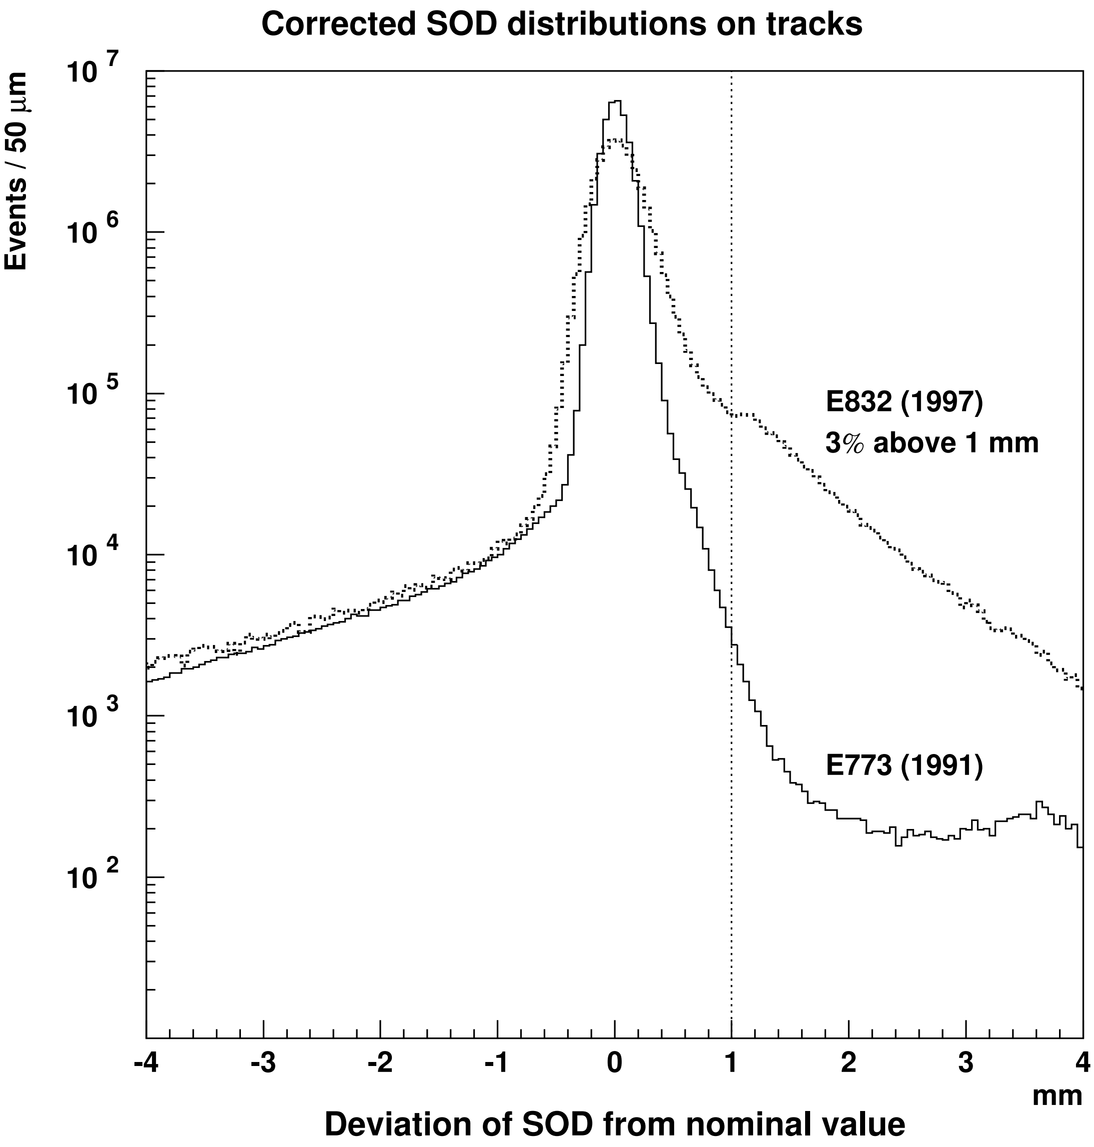

Figure 26. Drift chamber sum-of-distance (SOD) distributions for tracks in experiment E773 (which ran in 1991) and E832 (1997 run) using DC2Y. Deviation between SOD and cell spacing is shown (*i. e.* zero deviation means SOD= 6.35 mm, and negative deviations are low SODs). The nominal tracking cut at Δ(SOD) =1 mm is shown.

6.1.2 X and Y track candidates

Once the pairs had been found, the tracking code searched for track candidates. At this stage, all track candidates were searched for, whether they made any sense or not: they were not required to come from the beam or point to the CsI, and multiple track candidates could use the same hit.

6.1.2.1 Finding Y track candidates

Tracks were searched for first in the Y view. This was expedient, because the lack of a magnet bend made it simpler to search for tracks and generally reduced the number of combinations.

A track candidates was searched for by selecting one pair in DC1 and one pair in DC4 and drawing a line between them. If there were pairs in both DC2 and DC3 which fell less than 0.5 cm from this line, the quality values of the four pairs were added. If the sum was at least 11, then an attempt was made to fit all four pairs to a line. If the fit χ^2 was acceptable, the four pairs were considered a track candidate. This procedure was followed for all combinations of DC1 and DC4 pairs. The pair quality requirement allowed each Y track to have at most two low-SOD hits, or one isolated or high-SOD hit plus one low-SOD hit. Other hits were required to be good-SOD pairs.

Events were rejected at this stage in the offline analysis if they did not have at least two Y track candidates which could coexist. (Online and at the crunch stage, events were allowed to have only one Y track.) Two tracks were allowed to coexist if they did not share any hits or good-SOD pairs, and did not require the use of both hits in a high-SOD pair. Each hit in a low-SOD pair could be used by a different track.

6.1.2.2 X segments and X track candidates

Track candidates in the X view were found by forming "segments" between DC1 and DC2 and between DC3 and DC4, and then searching for intersections at the magnet bend plane between upstream and downstream segments.

A line was drawn between each pair in DC1X and each pair in DC2X. If the line formed less than a 100 μrad angle with the Z axis, and the sum of the two pair quality values was at least 4 (allowing both pairs to be low-SOD hits), the pairs were considered an upstream X segment. A similar procedure was followed for downstream segments between DC3 and DC4, although the maximum angle allowed was 150 μrad and the minimum sum of the pair values was 5.

For each combination of an upstream X segment and a downstream X segment, the distance between the two segments' projections to the magnet bend plane ($Z = 170.0$ m) was calculated. If the distance was less than 0.6 cm, and the sum of the pair values of the two segments was at least 11, the two segments formed an X track candidate. Events without at least two coexisting X tracks were rejected.

6.1.3 Vertex candidates

At this stage, we had an event with at least two track candidates in each view, so we could ask if the tracks were consistent with having come from a vertex in the vacuum decay region. The Z position of the intersection between each pair of Y track candidates was calculated, as well as the intersection of each pair of X track candidates. A vertex candidate consisted of two track candidates in each view, whose intersections in Z were consistent with each other and inside the decay region. All combinations were considered, and the requirements were loose. Some events (especially those with nearly parallel tracks) had many track and vertex candidates at this stage. It should be noted that the two X tracks had to have opposite bend direction (*i. e.* opposite charge). This was imposed at Level 3, which made it impossible to perform background studies later using same-sign vertices.

6.2 CsI energy reconstruction and clustering

This section describes the process of reconstructing electromagnetic clusters using the reconstructed energy in each crystal as well as information from the HCC.

6.2.1 Determining the energy in a crystal

The energy deposited in each crystal was determined from the DPMT readout for each channel in the in-time slice and the three subsequent slices. The energy was rescaled to take into account that the four time-slices used in the analysis only contained 95% of the energy.

The conversion between PMT charge and DPMT counts was calibrated for each channel using data from special laser scan runs described in Sec. 2.5.3. The conversion between charge and energy was calibrated crystal-by-crystal, using electrons from K_{e3} decays collected in the 2-TRACK trigger. The calibration was tuned such that E/p (cluster energy divided by measured track momentum) averaged to 1 for these electrons.

6.2.2 Cluster seeds

The cluster-finding algorithm used in this measurement began by searching all crystals which fired the HCC threshold. A crystal whose energy deposited was greater than any of its neighbors (*i. e.* a local maximum) was considered a "cluster seed." Neighboring crystals were defined as crystals which shared an edge (and not only a corner). Only neighbors which also fired their HCC thresholds were considered.

When a local maximum read out less than 0.1 GeV (symptomatic of a "hot bit" in the HCC, a low-level problem which caused an HCC bit to be set when there was no energy deposited in the crystal), the crystal was not considered a valid cluster seed.

6.2.3 Raw clusters

Once a local maximum was identified, a square block of crystals surrounding it was considered a cluster. In the large-crystal region, a cluster was 3×3 crystals; in the small-crystal region, it was 7×7 crystals. Along the border between the two regions, a cluster was the equivalent of 3×3 large crystals, where four small crystals were equivalent to a large crystal. (Their energies, however, were considered separately.)

The energies in all the crystals in a cluster were summed, forming the raw cluster energy. The ratios of the energy in the seed block to the energies in neighboring blocks were used to calculate the cluster position. The information for clusters in the data was taken from lookup tables generated from the data. The position resolution for electromagnetic clusters was ~ 1 mm.

6.2.4 Final corrections

After the raw clusters were formed, a series of corrections to the cluster energies was performed.

- *Overlap correction:* The overlap correction was used to separate the energies of clusters which shared crystals. This was accomplished by "dividing" the energy in each crystal in the overlap region between the two clusters using the cluster center positions and the expected transverse shapes of the clusters.

- *Neighbor correction:* The "neighbor correction" addressed the problem of the small fraction of energy deposited outside the 3×3 large crystal or 7×7 small crystal region occupied by a cluster. This energy could end up underneath a nearby but not overlapping cluster. The correction subtracted the expected contribution to each cluster from the tails of neighboring clusters.

- *Missing crystal correction:* For clusters which overlapped any edge of the calorimeter (the outer edge or the beam holes), the missing crystal correction added in the energy which would have been deposited if there had been CsI crystals beyond the edge. This energy was extrapolated from the energy in the observed part of the cluster and the expected transverse profile of the shower.

- *Threshold correction:* The outer crystals in a cluster (particularly a low-energy cluster) frequently had so little energy that they were below readout threshold. This energy was missed in the raw cluster energy. The threshold correction added it back in, accounting for pedestal fluctuations which could caused the readout thresholds to be unsharp.

- *Intra-block correction:* A scale correction to each cluster energy was performed based on the location of the cluster within the seed crystal. This was designed to take into account the slightly lower energy response observed in data for showers which originated near the edge of a crystal. As this variation in response was not simulated in the Monte Carlo, the correction was not performed on MC events.

6.3 Corrected tracking and vertex finding

After clusters had been found, the "corrected" tracking code could associate X tracks with Y tracks by projecting the track candidates to the CsI and requiring them to match the same cluster. A track was matched to the cluster closest to the track projection; if no cluster existed within 7 cm of the track, the track candidate was not matched.

For each vertex candidate where both tracks were matched to clusters, the hit positions were corrected for chamber wire alignment. (The chambers were aligned using muons and charged vertices, according to the technique described in Ref. [38].) Corrections for chamber rotations about the Z axis could only be performed after X tracks were matched to Y tracks. "Good-SOD" pairs where the SOD was marginally low (more than 0.06 cm below the nominal value) were "split," and the hit was used which caused the upstream and downstream track segments to intersect better at the magnet plane.

After these corrections, the vertex position in X and Y was redetermined, and a χ^2 was calculated for the hypothesis that the X and Y tracks intersected at the same Z position. A χ^2 for the matching of upstream and downstream segments at the magnet plane was also calculated. This was combined with the Z-position χ^2 and a weighting factor taking into account the number of poor-SOD pairs or single hits used in the two tracks. The result was a "vertex χ^2;" the vertex candidate with the best χ^2 was selected for the analysis.

CHAPTER 7

MEASUREMENT OF MATERIAL IN THE SPECTROMETER

This chapter presents the results of a measurement of the amount of material (measured in radiation lengths) in the upstream part of the spectrometer. This measurement is also described in Ref. [65]. The vacuum window and material associated with the upstream parts of Chamber 1 represented dead material in which photons could convert to electron pairs, and the electron tracks could be reconstructed. This material (Fig. 27) included the Mylar-Kevlar vacuum window, four additional $0.001''$ Mylar windows, half of the drift chamber, the air gap, and Bags 1a and 1b. Although the material was less than $0.01X_0$, these conversions caused significant background to many rare decay modes, particularly those with a two- or four-electron signal. Bremsstrahlung could cause backgrounds to radiative decay modes.

The material in this part of the detector was measured by examining conversions of photons from $K_L \to 3\pi^0$ decays using data from a special run taken near the end of the 1997 E799 run.

7.1 Vacuum window measurement technique

In E799-I, the amount of material in the detector was measured by examining electrons which passed through the detector material, and searching for photons from brems-

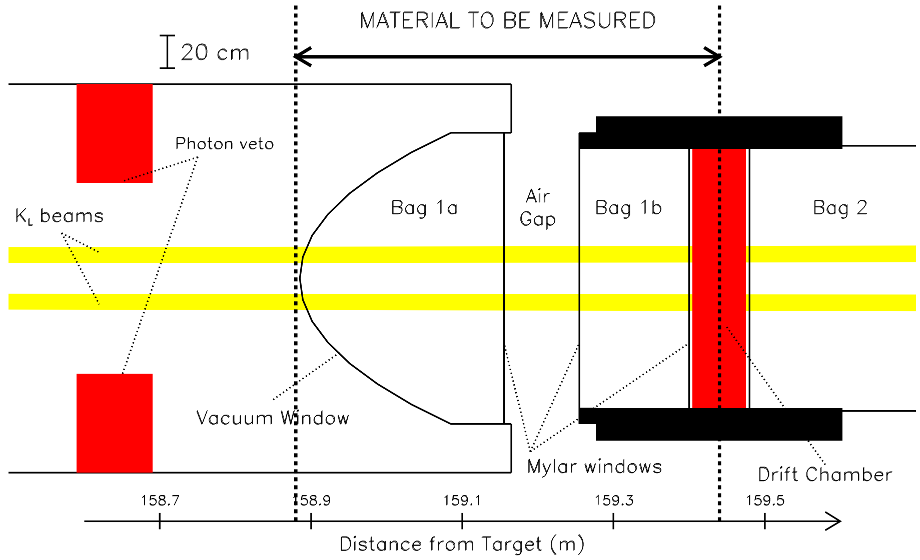

Figure 27. Expanded view of the region of the spectrometer near the vacuum window, with the photon conversion region indicated. Note the compressed vertical scale.

strahlung in the detector [66]. That technique, however, could not distinguish between the vacuum window assembly and material farther downstream. In E799-II, the vacuum window region was instead measured by examining conversions of photons from $K_L \to 3\pi^0 \to 6\gamma$ decays.

In order to make the acceptances similar for decays with a photon conversion and decays with no conversion (all-neutral $K_L \to 3\pi^0$ decays), a special run was taken with the spectrometer magnet off. This took advantage of the narrow opening angle of conversion pairs and the similar response of the CsI calorimeter to electrons and photons. In the energy range above 2 GeV, nearly all conversion pairs [63] had one of two topologies: either the two electrons struck the CsI less than 4 cm apart (and formed a single cluster), or the particles had such asymmetric energies that one electron was far below clustering threshold and could be lost without affecting the event kinematics significantly. The TRDs were removed for this run, to reduce the amount of downstream material where conversion electrons could scatter. Therefore, when using only the calorimeter in the reconstruction, a $3\pi^0$ analysis would find $K_L \to 3\pi^0$ events with six photons as well as those where a photon converted somewhere in the detector and the two converted electrons merged to form a sixth cluster. The large branching ratio and many kinematic constraints on this fully-reconstructed decay made non-$3\pi^0$ backgrounds negligible.

One disadvantage of this configuration was that it was not possible to calibrate the CsI using data from this special run. In normal runs, the calorimeter was calibrated using measured momenta of electrons from K_{e3} decays. In the special run, charged track momenta could not be measured because the spectrometer magnet was off. Calibration constants obtained from nearby runs were used for the special run data, but energy resolution was not optimal when using these constants.

Once the six-cluster final state was reconstructed and gave the kaon mass, the chambers could be examined to see whether or not one of the photons converted. It would be possible to infer the photon positions and directions from the kaon recontruction using only the calorimeter, and one could search for chamber hits on the wires along the reconstructed photon paths. Many hits would indicate a photon conversion. In practice, this did not appear to work well because of accidental chamber hits and multiple-scattering of conversion electrons. Therefore, a loose tracking algorithm was used. In most conversion events, the electrons were too close together to resolve in Chamber 1, so only one of the two tracks was reconstructed (see Fig. 28). Photons which converted downstream of the center of the first chamber did not produce reconstructible tracks, so material downstream of this point was not included in the radiation thickness measurement.

There was a potentially large background to conversions from the π^0 Dalitz decay. This mode has a large branching ratio (1.2%), and strongly favors small e^+e^- invariant mass (and therefore small opening angle). When the electrons were too close together to reconstruct a vertex, this decay was not easily distinguishable from a window conversion. With three π^0 decays in each event, the probability of one Dalitz decay was 3.5%, which was significantly higher than the probability of a photon conversion in the thin window. Although requiring the K_L decay to occur far upstream gave a longer flight distance for Dalitz electrons to separate, the background remained significant and difficult to model; Monte Carlo (MC) studies of this background showed it to be critically dependent on the simulation of multiple scattering in the detector and energy deposition in the CsI.

To reduce the Dalitz background, a special separator magnet [67] was used in the

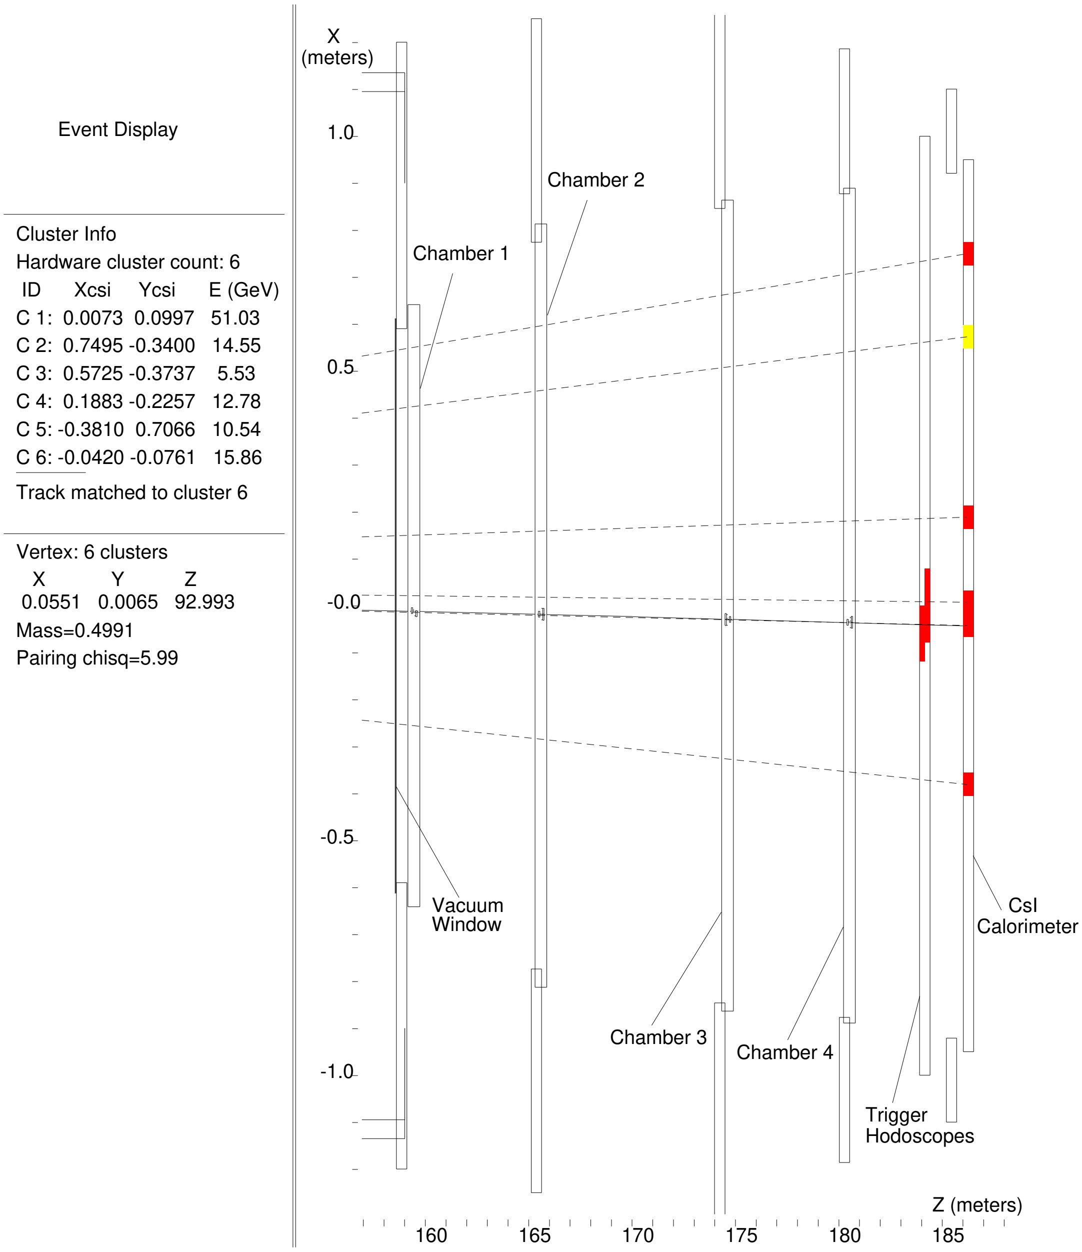

Figure 28. Display of a candidate photon conversion at the vacuum window. The display shows a plan view of the charged spectrometer including the four drift chambers, trigger hodoscopes, and calorimeter. In-time hits are shown in all detector elements. The downstream direction is to the right. Dashed lines indicate photon directions reconstructed assuming $3\pi^0 \to 6\gamma$; the solid line indicates a reconstructed charged track. Chamber wires with in-time hits are shown as small vertical lines in the two overlapping views of each chamber. The invariant mass of the three reconstructed π^0's is equal (within resolution) to the K_L mass.

upstream part of the decay region. This low-field dipole provided a transverse momentum kick of 10 MeV/c, which separated most Dalitz pairs enough to where they formed two separate tracks in the chambers and two clusters in the calorimeter. The magnet components were outside the vacuum pipe, so the field covered the entire cross-section of the vacuum and did not form an additional aperture. Only kaon decays which reconstructed upstream of the separator magnet were used in the analysis. All material being measured was well downstream of this magnet.

7.2 Magnet design and construction

The electron separator magnet was a C-magnet providing a dipole field of 400 gauss in the Y-direction. It was constructed from Sr-ferrite permanent magnet bricks of the type used for the FNAL recycler ring [68], with steel pole faces and flux return. The magnet was positioned around the downstream end of the 18" vacuum pipe; the center of the magnet was at $Z = 97.60$ meters. This position was selected because it was the farthest downstream position where the decay tank was stainless steel— the steel tanks farther downstream would have attenuated the field, requiring that the magnet be placed *inside* the vacuum. This was judged impractical.

The field was in the positive Y-direction, selected to match the polarity of the final sweeper magnet (NM3S), which ended four meters upstream of the separator magnet.

7.3 Special runs

7.3.1 Beam and detector configuration

The radiation length analysis used four data runs: 10886, 10887, 10891, and 10893. These were taken on August 27 and 28, 1997 under the following beam conditions:

- Primary beam intensity: 4×10^{11} protons per pulse

- Slow spill duration: 23 seconds

- Beryllium absorber: none
- Collimators: Summer E799 beams (11×11 cm at CsI)
- Regenerator and mask: Out
- TRD: TRDs 1-8 out; TRD 9 in
- Magnet settings:
  - Spectrometer magnet: Off
  - Final sweeper magnet: 2000 A
  - Separator magnet at 97.60 meters; P_T kick 10 MeV

7.3.2 Triggers

There were four significant triggers used in the radiation length run. No Level 3 filtering was done, although standard E799 Level 3 code was run and produced histograms.

- Trigger 3 (`DC2-LOOSE` or "charged trigger") was the primary trigger for collecting vacuum window conversions. At Level 1, it required E-total threshold 3 and a hit in either the X or Y views of the DC2-OR. All the ring counters, SAs, and the CA and HA were in veto. At Level 2, at least 6 HCC clusters were required, as well as at least one Y-view Kumquat/Banana hit in chambers 2, 3, and 4. This trigger was not prescaled.
- Trigger 4 (`DC4-LOOSE`) was a trigger similar to Trigger 3, with all DC activity requirements removed except for the Chamber 4 Kumquat requirement. Trigger 4 was taken to allow radiation lengths downstream of DC1 and DC2 to be measured; this analysis has not yet been done. Trigger 4 was prescaled by a factor of 2.
- Trigger 8 (`NEUTRAL`) was the same as trigger 3, but with all chamber requirements removed. It was used to collect all-neutral $3\pi^0$ decays for normalization. Trigger 8 was prescaled by a factor of 20.

- Trigger 14 (ACCID) was the standard accidental trigger from E799. All Monte Carlo studies described here include accidental overlays from this trigger, collected in the radiation length runs.

7.4 Analysis

7.4.1 Calorimeter analysis

The $K_L \to 3\pi^0 \to 6\gamma$ decay was reconstructed using its six photon signature in the CsI calorimeter. The calorimeter analysis was the same for both trigger types.

Events were required to have exactly six hardware clusters, with a minimum cluster energy of 2 GeV. The six clusters were arranged into pairs, and a "vertex" distance z_{12} from the calorimeter was calculated for each pair using the π^0 mass as a constraint (the reconstruction technique is described further in Sec. 8.2.1):

$$z_{12} = \frac{\sqrt{E_1 E_2 r_{12}}}{m_{\pi^0}}$$

E_1 and E_2 are the cluster energies and r_{12} is the distance between the clusters in the calorimeter. Of the 15 possible pairing combinations, we selected the one with the best χ^2 for the hypothesis that the three pairs all came from π^0 decays at the same Z-position. This χ^2 was required to be less than 25 (with two degrees of freedom). The requirement was loose because cluster position and energy resolution was degraded if the photon converted far upstream, and this component of the resolution was not taken into account when calculating the χ^2. In addition, the reconstructed $3\pi^0$ invariant mass was required to be within 20 MeV/c^2 of the K^0 mass (Fig. 29), and the center of energy at the calorimeter was required to be within the beams. This last cut removed kaons which scattered in the absorber or collimators.

Only events with a reconstructed decay position between 91 and 97 meters from the target were accepted. This short range, in the far upstream end of the vacuum decay region, accepted only decays upstream of the e^+e^- separator magnet.

MC studies showed that the efficiency of the $K_L \to 3\pi^0$ reconstruction for events with a photon conversion at the vacuum window was 0.94 of the efficiency for events

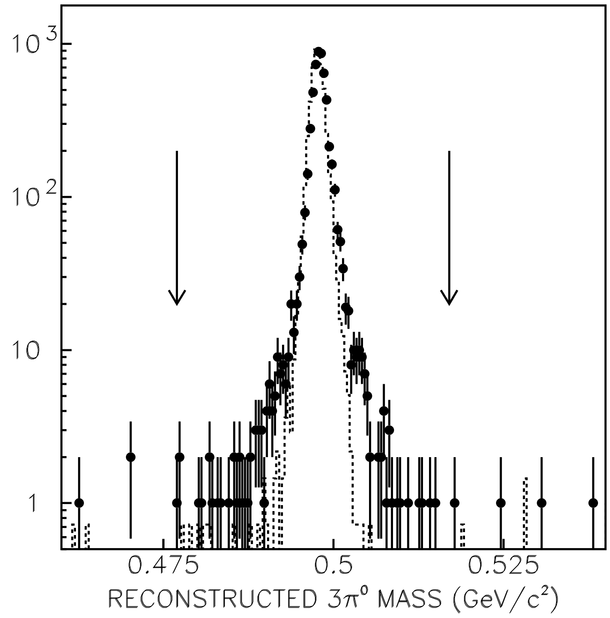

Figure 29. Reconstructed $3\pi^0$ mass for photon conversion candidates which passed all other cuts. Points with errors are data; dashed histogram is Monte Carlo. The vertical arrows indicate the cut boundaries.

without one. The acceptace loss for window conversions was caused by events with a seventh cluster above 2 GeV or a bad pairing χ^2.

After the calorimeter analysis, 17559 events remained in the data sample taken from the neutral trigger. Correcting for the prescale, acceptance, and the $\pi^0 \to \gamma\gamma$ branching ratio of 0.988, this yielded a total sample of 1.44×10^7 $K_L \to 3\pi^0$ decays (or 6.8×10^7 total K_L decays) in the region of interest.

In the charged trigger, 17228 events remained at this stage of the analysis. Most of these did not have a vacuum window conversion, but satisfied the trigger due to accidental drift chamber activity.

7.4.2 Charged analysis

Conversions were identified using a special version of T3 tracking, which allowed many poor-SOD pairs to be considered for tracks. This was necessary because the very close tracks in conversion events could cause many low-SOD pairs. Standard tracking allowed each X or Y track candidate to include one isolated hit or two low-SOD pairs. The special tracking allowed each X or Y track candidate to have up to four low-SOD hits, or three low-SOD hits with an isolated hit, or up to three isolated hits with a good-SOD pair. Allowing these extra poor-SOD hits increased the efficiency for finding at least one track in a MC conversion event to greater than 99.9%. There were no additional restrictions on the number of good-SOD pairs in corrected tracking.

Two tracking topologies were accepted for conversion candidates:

7.4.2.1 One-track events

"One-track" events accounted for 94.9% of the MC conversion events which passed the calorimeter analysis cuts. These events had a single reconstructed track, which pointed to one of the six clusters in the CsI. In these events, either the two tracks were too close to be resolved in either the chambers or the calorimeter, or the low-energy track scattered out downstream and produced a cluster below threshold. 5318 events in the charged trigger data sample satisfied this cut.

7.4.2.2 Two-track events

Only 1.8% of MC conversion events were accepted as "two-track" events. These events had two reconstructed tracks, both of which pointed to the *same* cluster in the CsI. Because events were required to have only six clusters in the calorimeter analysis, single-conversion events where the two electrons formed two different clusters — approximately 3% of all conversions — were rejected earlier, at the calorimeter analysis stage. 304 charged trigger events satisfied the two-track criteria. Adding the one-track events gave a total of 5622 photon conversion candidates.

Table 10. Acceptance categories for six-cluster MC conversion events which passed the calorimeter analysis.

One track, matched to cluster (accepted)	94.9%
Two tracks, matched to same cluster (accepted)	1.8%
Two tracks, matched to different clusters (double conversion topology; rejected)	1.5%
Two tracks, one not matched to any cluster (accidental or low momentum track scattered; rejected)	1.2%
Three or more tracks (accidental track or multiple conversion; rejected)	0.6%

7.4.2.3 Conversions which fail the charged analysis

Of MC conversion events which passed the calorimeter analysis, 96.7% also passed either the one-track or two-track charged analysis. On examination of the 3.3% of MC conversion events which passed the neutral cuts but failed to satisfy either of these two topologies, 80% had two reconstructed tracks; the rest had three from either multiple conversions or accidental tracks.

Of the two-track events, about half had two photons convert (this was consistent with the expected radiation thickness of the vacuum window). The remaining events had two tracks, only one of which was matched to a CsI cluster because the lower-energy track from the conversion was below the threshold for forming a cluster in the CsI. This event loss depended sensitively on how multiple scattering was modeled in the MC, and uncertainty in this inefficiency dominated the systematic error in the vacuuum window measurement (as discussed below).

The various topologies of MC conversions are summarized in Table 10.

7.4.3 Dalitz background

The largest irreducible background to conversions came from the π^0 Dalitz decay ($\pi^0 \to e^+e^-\gamma$). Both the final sweeper and the electron separator magnet fields affected the spectrum of Dalitz background; both magnets were simulated using measured field maps. Studies of fully reconstructed Dalitz decays from the radiation length run were used to confirm the separator field simulation, and studies of upstream $K_S \to \pi^+\pi^-$ and $\Lambda^0 \to p\pi^-$ decays in E799 data verified the final sweeper simulation.

MC studies showed the expected background from Dalitz decays in the signal region to be 183 ± 11 events (3.4% of conversion candidates) in the one-track case and 37 ± 5 events (12% of candidates) in the two-track case, giving a total of 220 ± 12 events (3.9% of the total conversion candidates). The errors shown for the Dalitz background are from MC statistics. Although this background could be reduced somewhat by cutting events with extra in-time activity in the first chamber, the efficiency of this cut for signal events was difficult to model. Therefore, no further attempt was made to reduce Dalitz background. Fig. 30 shows the Z-distribution of data points along with MC predictions from photon conversions and Dalitz background. As can be seen in the figure, the separator magnet improved the signal-to-background ratio from ~ 1 to about 25.

7.4.4 Out-of-time track background

Although requiring tracks to match in-time hits in the trigger hodoscopes (VV$'$ hit verification) reduced the out-of-time track background greatly, a very small background remained due to accidental in-time hodoscope hits and hits from photons which converted downstream of the first drift chamber. Some of these events could be identified by pointing the track back to the reconstructed Z-position of the decay (see Fig. 31). An accidental track generally did not point back to the transverse position of the decay, while most genuine conversion tracks pointed back to within a few centimeters (though this was somewhat dependent on multiple scattering modeling).

74 events in the data had a track projection more than 40 cm from the reconstructed

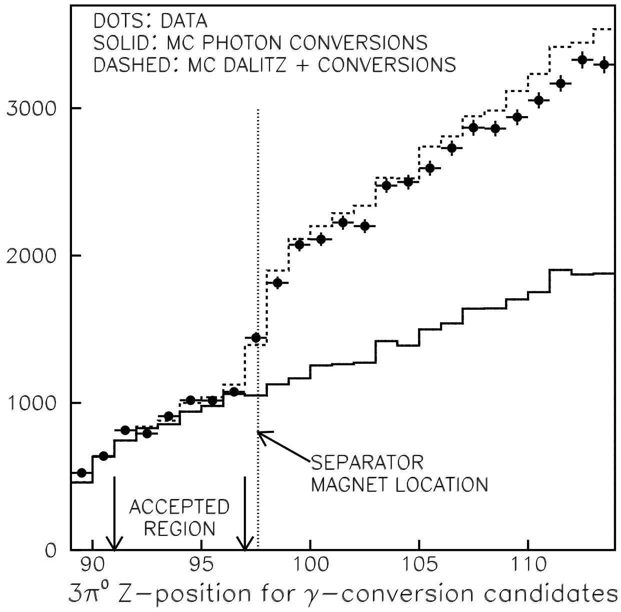

Figure 30. Z-position (meters from target) of candidate $K_L \to 3\pi^0 \to 6\gamma$ events with a photon conversion. The dots with errors are data; the solid histogram is the MC prediction from conversions alone, normalized to the background-subtracted data points in the accepted region from $Z = 91$ to $Z = 97$ meters. The dashed histogram is the same MC prediction with Dalitz decay background added. The Dalitz background was normalized using the measured rate of $K_L \to 3\pi^0$ decays and the $\pi^0 \to e^+e^-\gamma$ branching ratio. The electron separator magnet at $Z = 97.6$ meters greatly suppressed the Dalitz background upstream of it. The Dalitz background downstream of the magnet does not match the MC simulation perfectly, because it was critically sensitive to the modeling of multiple scattering in the detector.

decay position; the Monte Carlo predicted only about 4 signal events in this region, so these events were mostly background. Of these 74 events, 8 survived VV$'$ verification. Assuming these were all signal gave a background efficiency of zero; assuming they were all background gives a background efficiency of 0.11. For purposes of determining systematic error and residual background, the background efficiency was taken as 0.06 ± 0.06.

Similarly, we can examine events where the track pointed closer than 1 cm to the decay, to enhance the signal content and determine the signal efficiency of the VV$'$ cut. Of the 2771 events in this region, 2761 satisfied the cut. Assuming that all the lost events

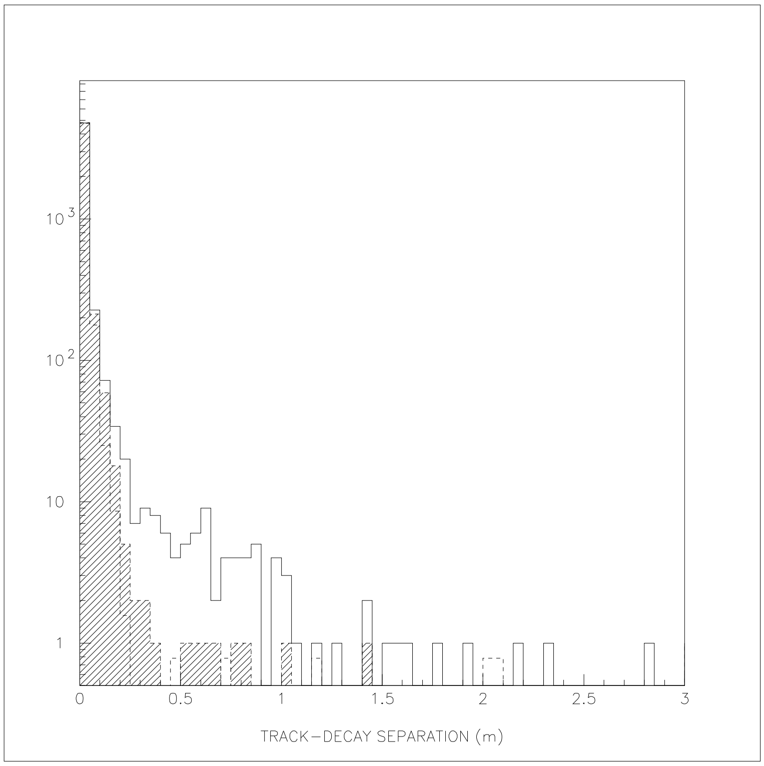

Figure 31. Distance between track projection and decay center of energy. Solid: Data without VV′ hit verification. Dashed: Monte Carlo vacuum window conversions. Striped: Data after VV′ hit verification.

were signal, the cut had a signal efficiency of 0.996; assuming they were all background gave a signal efficiency of unity. The signal efficiency was taken as 0.998 ± 0.002.

To estimate the total expected out-of-time track background remaining after VV′ hit verification, we noted that the cut rejected 165 events. Assuming these events were all background, and that the cut accepts 0.06 ± 0.06 of the background, there were 11 ± 11 remaining background events (systematic error only) to be subtracted.

7.4.5 Extraction of the photon conversion probability

The background-subtracted conversion signal was 5391 ± 78 events, where the error is statistical only, and includes statistical error on the subtracted backgrounds. Dividing by the MC acceptance gave a total of $(2.29 \pm 0.033) \times 10^5$ $K_L \rightarrow 3\pi^0 \rightarrow 6\gamma$ decays in the acceptance region where at least one photon converted at the vacuum window.

We divided by the measured number of $K_L \rightarrow 3\pi^0$ decays and corrected for the $\pi^0 \rightarrow \gamma\gamma$ branching ratio, yielding a probability of 0.01645 ± 0.00024 that at least one of six photons converted, or an individual photon conversion probability of $(2.761 \pm 0.040) \times 10^{-3}$ (the error is statistical only).

7.4.6 Other systematic errors

The radiation length analysis measured the ratio of events with tracks to events without tracks. Tracking efficiency, therefore, entered the calculation at first order. Studies in normal data showed the single-track-finding efficiency to be better than 0.99. Although in principle we had two tracks in the event, only one was usually reconstructible. As there was no clean way to measure inefficiency at this level in the radiation length data set, a systematic error of $\pm 1\%$ was assigned.

A larger systematic error came from modeling of multiple scattering in the detector. This was a dominant source of the acceptance loss for conversions relative to all-neutral decays. The model used in the MC was a parameterization of the Molière theory [61]; all scattering was assumed to occur at the vacuum window, the four chamber planes, and four equally-spaced slices of material which represented the helium bags (as described in Chap. 5). Much of the Dalitz background and some of the signal, according to the MC, had at least one track below 0.5 GeV. Typical scattering angles for these tracks in the vacuum window and the chambers were greater than 1 mrad, producing position shifts greater than one cell in the downstream drift chambers. Under these conditions, tracks in the spectrometer may not have remained straight enough to be reconstructed, and the approximation that scattering occured only at the chamber planes could become invalid. This is seen in the disagreement between the data and MC in the acceptance for Dalitz decay background downstream of the separator magnet (Fig. 30), as well as in the poor prediction of the ratio of two-track to one-track conversion candidates (5% in data after background subtraction, versus 2% in MC).

Table 11. Summary of sources of error in radiation length measurement

Conversion statistics	1.5%
Tracking efficiency	1.0%
Trigger hodoscope efficiency	0.2%
Out-of-time track background	0.2%
Pairing χ^2 cut	2.7%
Multiple scattering modeling	3.7%
Total	4.9%

The effect of multiple scattering on the radiation length measurement was studied by varying all material in the MC scattering simulation upward by a factor of two, and downward to zero. Removing scattering effects caused the reconstructed conversion probability to drop by 0.9% of itself. Doubling the scattering had more effect; it caused the measured conversion probability to increase by 3.7%. We conservatively assigned a $\pm 3.7\%$ systematic error to multiple scattering effects.

The pairing χ^2 cut in the calorimeter analysis was an additional source of systematic error. This error was primarily due to miscalibration of the CsI calorimeter, which caused cluster resolutions to be poorer in the data than in the MC. The error was estimated by changing the value of the cut fron 6.25 to 200 (the nominal cut is at 25). The range was chosen to be more than twice the discrepancy between the median χ^2 in the data and MC photon conversion samples. This caused a variation of $\pm 2.7\%$ in the measured conversion probability, and this was taken as the systematic error.

7.4.7 Calculating the radiation length

The probability for a high-energy photon to convert in a thin ($\ll X_0$) window is 7/9 of the thickness of the window in radiation lengths. The error on the measurement of conversion probability was 4.9% of the total (see Table 11), so the conversion probability

Table 12. Calculated sources of material up to midplane of DC1.

Material	X_0	Contribution ($10^{-4}X_0$)
Vacuum Window [69]		15.60
25 cm Helium (Bag 1a)	5.68×10^5 cm	0.44
1 mil Mylar	28.7 cm	0.89
10 cm air gap	3.04×10^4 cm	3.29
1 mil Mylar	28.7 cm	0.89
15 cm Helium (Bag 1b)	5.68×10^5 cm	0.26
1 mil Mylar	28.7 cm	0.89
0.9 cm "N_2" gap ($Ar+CO_2$)	1.78×10^4 cm	0.51
1 mil Mylar	28.7 cm	0.89
4.1 cm DC gas	1.78×10^4 cm	2.30
6 cathode planes + window wires		0.84
2 anode planes		0.11
Total calculated		26.91
Total measured		35.5 ± 1.7

was $(2.76 \pm 0.13) \times 10^{-3}$. Scaling by 9/7, the amount of material upstream of the midplane of the first drift chamber was measured to be $(3.55 \pm 0.17) \times 10^{-3} X_0$.

7.5 Comparison to calculated value

What was measured in this exercise was the probability of a photon conversion in material upstream of the mid-plane of DC1. This region included the vacuum window, Helium bags 1a and 1b, an air gap for the vacuum safety shutter, and the "N_2 gap" (which was really $Ar+CO_2$, due to chamber overpressure and a bulging window) at the upstream face of DC1. Adding up the calculated radiation length of all these components yielded $2.69 \times 10^{-3} X_0$ (See Table 12).

The measured value of 35.5×10^{-4} was significantly higher than that number. One possible explanation was that the vacuum window itself, which had been expected to account for 60% of the material in the region, was thicker than calculated. The basket-weave of the window material, as well as the Mylar-Kevlar composition, complicated the calculation. However, the calculation was expected to be accurate to 10% ($1.6 \times 10^{-4} X_0$). The discrepancy between the measured and calculated thickness was five times this uncertainty. It would be difficult to explain this large an error in the vacuum window calculation.

A smaller uncertainty in the calculation was the treatment of the "N_2 gap" region. This was originally a gap between the helium bag and the chamber gas, and had CO_2 flowing through it. However, the chambers were operated at a large overpressure, and this caused the chamber window to bulge until it touched the helium bag window. So, the "gap" region was mostly chamber gas, and indeed, the chamber window may even have bulged very slightly into the helium region. The calculation assumed the 0.9 cm region was filled with chamber gas. Even assuming it was all CO_2, however, only affected the result by $0.3 \times 10^{-4} X_0$.

Finally, we addressed the question of whether the helium bags were actually filled with helium. If Bag 1a were assumed to be all air, the calculated material would total 34.7×10^{-4}, which is in good agreement with the measured value. This suggestion was later confirmed by a study in E832 of the Z-position distribution of vertices from hadronic interactions (Fig. 32), which indicated that the rate of interactions in "Helium bag" 1a was similar to that in the air gap, and much higher than in Bag 1b, which was filled with helium.

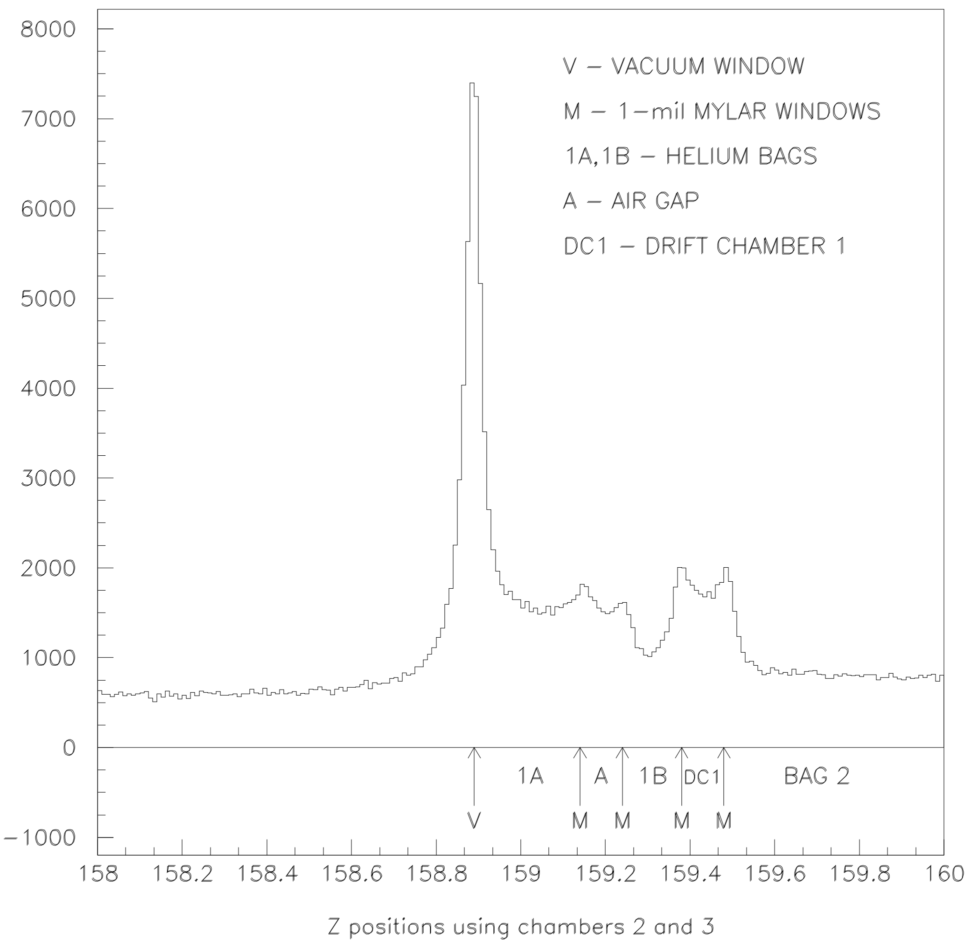

Figure 32. Z-position of vertices in the region of the vacuum window, reconstructed using drift chambers 2, 3, and 4 only. The positions of detector elements are shown. The helium in Bag 1b reduced the number of interactions in that region, but the expected reduction is not seen in Bag 1a. Kaon decays produced a flat background. (Data provided by P. Shawhan)

CHAPTER 8

ANALYSIS OF THE SIGNAL AND NORMALIZATION SAMPLES

"Error analysis: Errors are random and systemic (sic)."

—*Anonymous Physics 121*[1] *laboratory report, Autumn 1993*

This chapter describes the analysis of the $\pi^0 \to e^+e^-$ and $\pi^0 \to e^+e^-\gamma$ samples in the data, and the calculation of the branching ratio of $\pi^0 \to e^+e^-$.

The branching ratio of $\pi^0 \to e^+e^-$ was measured by normalizing it to the high-m_{ee} tail of the $\pi^0 \to e^+e^-\gamma$ decay; the actual measured quantity was the ratio

$$\frac{\mathrm{BR}(\pi^0 \to e^+e^-, x > 0.95)}{\mathrm{BR}(\pi^0 \to e^+e^-\gamma, x > 0.232)}.$$

The cutoffs in x correspond to:

$$x > 0.95 : m_{e^+e^-} > 0.1316 \text{ MeV}/c^2$$

$$x > 0.232 : m_{e^+e^-} > 0.0650 \text{ MeV}/c^2$$

8.1 Acceptance at trigger level

In the winter run MC with the muon veto in place but no Collar Anti (CA) veto, the acceptance at Level 2 was 10.37% for Dalitz decays, and 13.20% for $\pi^0 \to e^+e^-$. In the

[1] An introductory physics course at The University of Chicago, primarily for pre-medical students.

summer run, with the CA veto but no muon veto, the acceptance was 9.30% for Dalitz decays and 12.25% for $\pi^0 \to e^+e^-$. Note that the $\pi^0 \to e^+e^-$ acceptance is quoted here for decays generated with *all values* of $(m_{e^+e^-}/m_{\pi^0})^2$, not the $(m_{e^+e^-}/m_{\pi^0})^2 > 0.95$ region for which we will eventually quote a branching ratio.

The CA rejected 17.6% of Dalitz decays (these decays were generally unreconstructible) and 16.9% of $\pi^0 \to e^+e^-$. The muon veto rejected 7.4% of all data in the winter MC. As the generated decays could not produce a muon hit, all these losses came from the accidental overlays and were the same in the signal and normalization modes.

8.2 Signal and normalization reconstruction

The analysis for both signal and normalization events required that a good two-track vertex be reconstructed, and that both tracks were matched to clusters in the CsI. Any events with more than two tracks were rejected. All calorimeter clusters without tracks pointing to them were assumed to be photons.

8.2.1 $\pi^0 \to \gamma\gamma$ reconstruction

For each pair of photons, we could assume they came from a $\pi^0 \to \gamma\gamma$ decay, and calculate the decay distance from the calorimeter, using the π^0 mass as a constraint and assuming a small opening angle between the photons:

$$Z_{12} = Z_{\mathrm{CsI}} - \frac{\sqrt{E_1 E_2} r_{12}}{m_{\pi^0}}$$

where E_1 and E_2 are the photon energies and r_{12} the distance between the clusters at the calorimeter. Given two pairs of photons, we could calculate a Z position for each pair, and then calculate the χ^2 for the hypothesis that the two Z-positions agreed:

$$\text{Pairing } \chi^2 = \frac{(Z_{12} - Z_{34})^2}{\delta^2_{Z_{12}} + \delta^2_{Z_{34}}}$$

For four photons, there were three possible $2\pi^0$ pairings, as illustrated in Fig. 33. For five photons, there were fifteen possible $2\pi^0$ pairings.

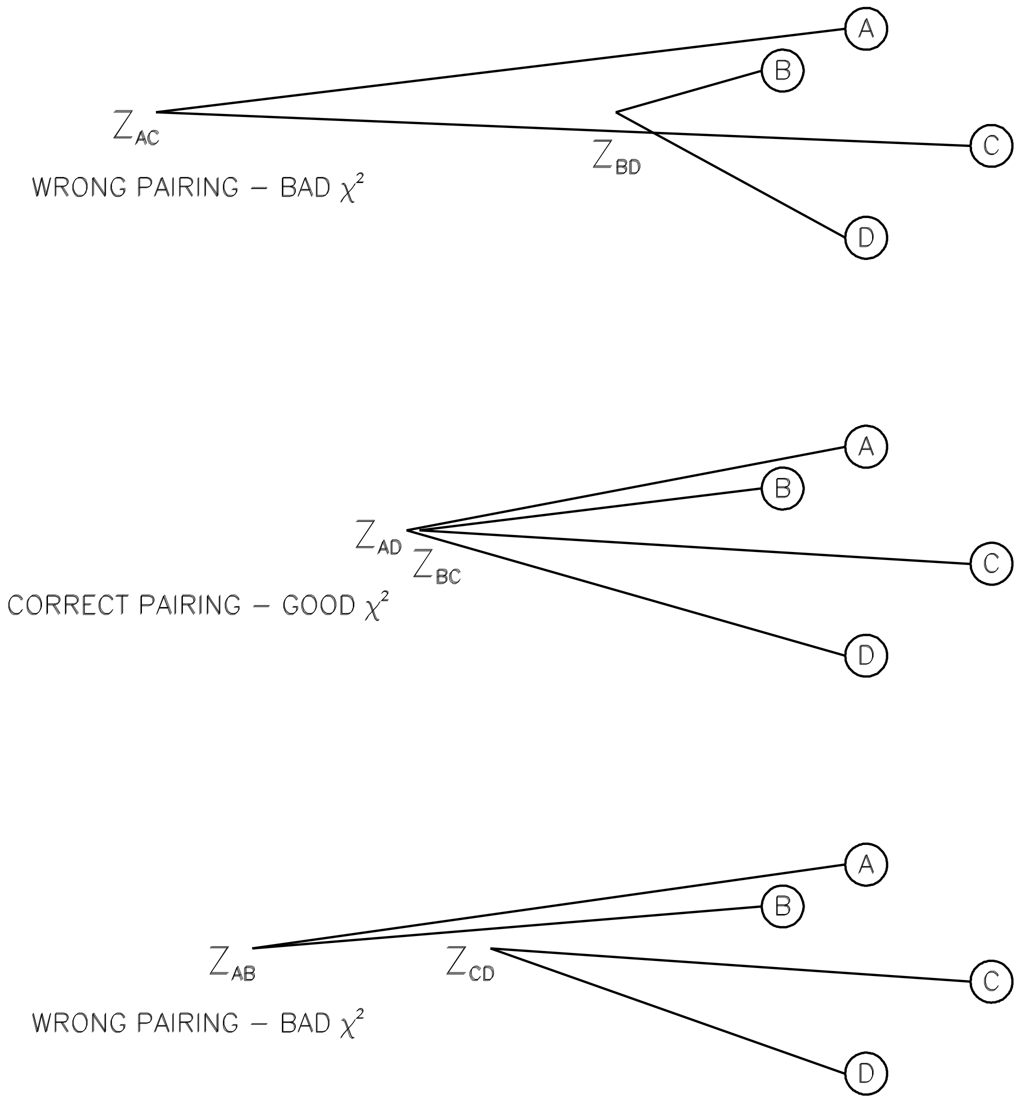

Figure 33. The three possible $2\pi^0$ pairing combinations for four photon clusters. The second case is the correct pairing.

Six-cluster, two-electron events were reconstructed as $K_L \to 3\pi^0$, $\pi^0 \to e^+e^-$ candidates. The four photon clusters were paired as two π^0's, and the pairing with the best χ^2 was selected. The momentum four-vector of each photon was calculated assuming it originated at the Z-position of the $2\pi^0$ pairing and the X- and Y-position of the two-electron charged vertex.

Seven-cluster, two-electron events were reconstructed as $K_L \to 3\pi^0$,$\pi^0 \to e^+e^-\gamma$ candidates. The five photon clusters were paired as two π^0's and one extra photon; the pairing with the best χ^2 was selected and the extra photon was assigned to the $\pi^0 \to e^+e^-\gamma$ decay. As in the signal analysis, the photon four-momenta were calculated assuming the photons originated from the $2\pi^0 Z$-position and the charged vertex X- and Y-position.

8.2.2 Basic reconstruction cuts

The following cuts were applied in the early stages of the analysis. Except where noted, these cuts were common to both the signal and normalization analyses.

8.2.2.1 Two-electron requirement

Both tracks were required to be electron candidates: the final analysis required $0.92 > E/p > 1.08$ (E/p distribution for data and MC are shown in Fig. 34). This cut was nearly 100% efficient for reconstructible signal and normalization decays.

8.2.2.2 Vertex position cuts

In E799-II, unlike in E832, the upstream end of the fiducial decay region was not defined by an aperture, but instead by the end of the field region of the final sweeper magnet at about $Z = 94$ meters. This was a problem, because decays which occured in the downstream fringe field of this magnet were distorted by the magnetic field. This could fake a vertex somewhat farther downstream, and also change the invariant mass of the vertex. These decays could become very serious backgrounds, as a low-mass electron

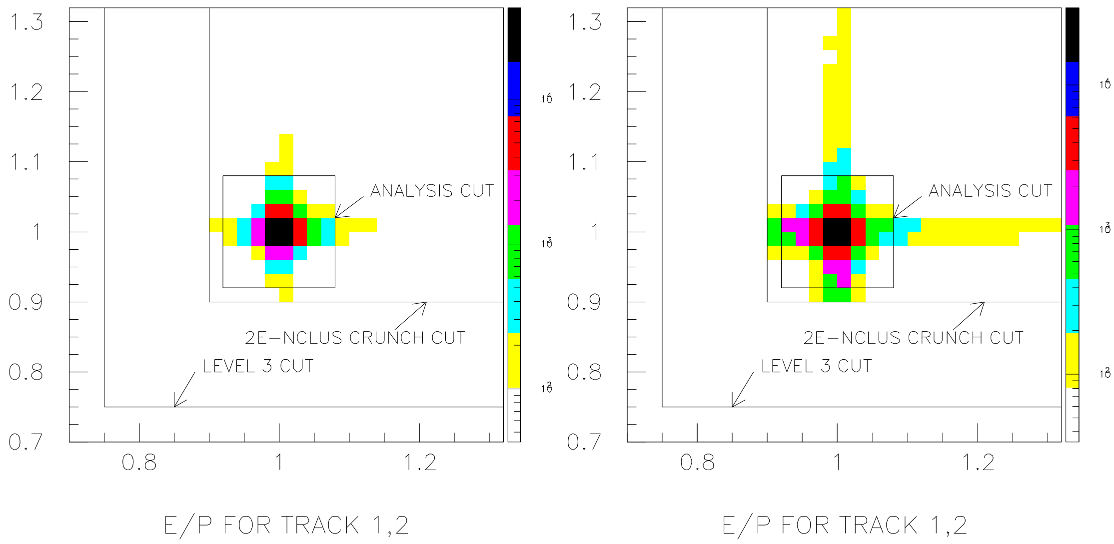

Figure 34. E/p vs. E/p for the two reconstructed tracks in Dalitz MC (left) and 2E-NCLUS crunch output (right). Note that events with $E/p < 0.9$ in the data plot were already removed by the crunch code. The events with $E/p \gg 1$ may have had accidental clusters matched to the tracks (which may or may not have been electrons).

pair (say, from a Dalitz decay) could accidentally be kicked up to form a π^0 mass. To reject these decays, all events with a charged vertex upstream of $Z = 97$ meters were cut from the sample.

On the downstream end, the boundary of the fiducial region was the vacuum window at $Z = 158.89$ meters. The vertex position was required to be upstream of $Z = 158$ meters, rejecting vertices which came from conversions or interactions in the window. It should be noted that after making this cut, there was still background from vacuum window conversions in both the Dalitz and $\pi^0 \to e^+e^-$ samples, but the remaining events had a four-track topology and the vertex reconstructed well within the vacuum region. These events are discussed in Sec. 8.3.

The vertex Z distribution at this stage is shown in Fig. 35. Backgrounds from the final sweeper at $Z \approx 94$ m and from the vacuum window at $Z = 159$ m are clearly visible. About 70% of the events which survived the Z-cut were well-reconstructed Dalitz decays.

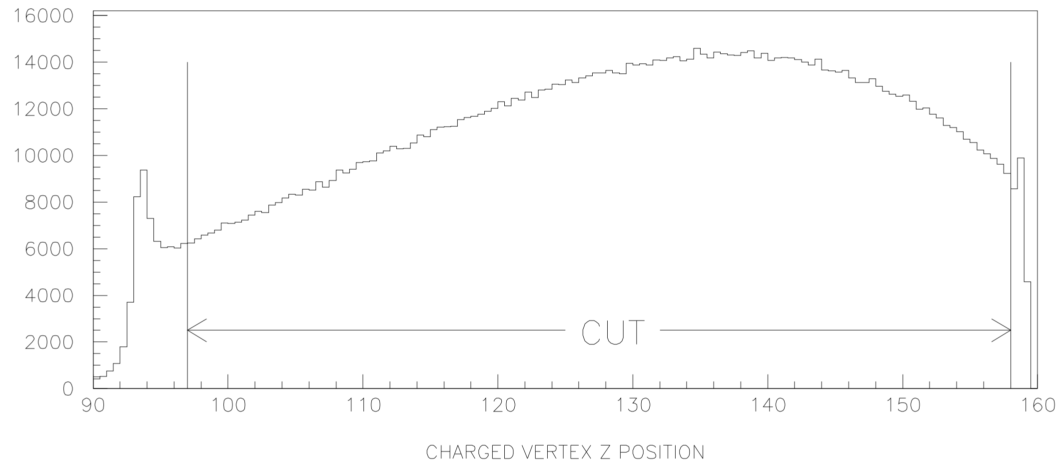

Figure 35. Charged vertex Z position (meters) for the normalization sample after reconstruction and E/p cuts. The Z cut is shown.

8.2.2.3 Trigger verification

Certain cuts were made on the sample to ensure that the decay products were responsible for firing the Level 1 and Level 2 trigger. Some events may have passed the trigger due to extra activity in the detector from accidentals, even if the signal event on its own would not have. These "volunteer" events could not be simulated reliably in the MC, and we did not use them in the final sample.

- The tracks were required to project at least 1 cm from the nearest VV$'$ beam hole. This ensured that both tracks actually passed through the counters and that the trigger was generated by the signal tracks in the event, as opposed to accidental downstream tracks or photon conversions in the TRDs.

- The total kaon energy was required to be above 40 GeV. This cut was made to remove events which required extra accidental energy to pass the energy sum trigger at Level 1. This cut was well above the nominal threshold of 25 GeV, removing any possible effects from the modeling of the threshold. 5% of signal MC was lost by this cut.

- The minimum energy for any of the reconstructed clusters in the CsI was required to be above 1.5 GeV. This cut was imposed to reduce dependence on the simulation of the HCC thresholds.

- Events where the two X tracks shared a Y track were rejected in the final reconstruction. The drift chamber hit counting at Level 2 required multiple Y-view hits, and thus events where the two Y-tracks overlap could only have passed this requirement if extra activity had been present in the chambers.

- The maximum energy in any counter in the RCs, SAs, or CIA was required to be below 0.4 GeV, and the maximum counter in the CA below 14 GeV. These thresholds were slightly tighter than the online thresholds used in the trigger, and removed any events which may have sneaked through because of malfunctioning trigger sources or (more important) runs where certain vetoes were not in the trigger.

8.2.2.4 Mass and transverse momentum cuts

The total invariant mass of the electrons and photons in the event was compared to the K_L mass, and was required to be within 0.05 GeV/c^2. This cut was very loose, to reduce any bias from misunderstanding of the mispairing and radiative tails in the mass distribution.

We also placed a cut on the transverse momentum ($|p_T|$). This quantity was defined as the modulus of the total momentum of the reconstructed decay products transverse to a line between the decay vertex and the target. $|p_T|$ was required to be below 30 MeV/c. This helped reduce background from events with missing particles and accidental energy.

8.2.2.5 Dalitz π^0 mass and e^+e^- mass

For the seven-cluster normalization mode only, a loose cut was made on the $e^+e^-\gamma$ invariant mass: $|m_{e^+e^-\gamma} - m_{\pi^0}| < 30$ MeV/c^2.

Also, events with $m_{e^+e^-} < 0.070$ GeV/c^2 were cut from the sample. This was done to remove effects from e^+e^- mass resolution (as MC events were only generated with $m_{e^+e^-} > 0.065$ GeV/c^2).

8.2.3 Acceptance after reconstruction and basic cuts

In the winter MC, 1.615% of generated Dalitz decays were reconstructed as 2-electron, 5-photon events. 4.11% of the generated $\pi^0 \to e^+e^-$ events were reconstructed as 2-electron, 4-photon events. In the summer, the numbers were 1.731% and 4.36% respectively. In the Dalitz case, the acceptance rose by 7.2%; in the $\pi^0 \to e^+e^-$ case, it rose by 6.1%. Assuming the muon veto to be the only change in the MC, the acceptances for both modes should have increased by 7.9%.

The basic cuts described in this section were 70.6% efficient for normalization MC, and 75.9% efficient for signal MC. At this stage, the acceptance for Dalitz decays was 1.139% (winter), 1.223% (summer). For $\pi^0 \to e^+e^-$, the acceptance was 3.106% (winter), 3.325% (summer).

8.3 Backgrounds and final sample selection

Here we consider the backgrounds to $\pi^0 \to e^+e^-$, and the final cuts which were implemented to reduce the background level in the $\pi^0 \to e^+e^-$ sample. The discussion of backgrounds will concentrate on the three classes of backgrounds: the high-x Dalitz decays, the correctly-paired four-track backgrounds, and the mispaired four-track backgrounds. The three individual modes within each class of four-track background are discussed in Sec. 5.6.2, and will not be treated separately here. They had very similar shapes and the relative contributions among them were already implemented in the MC.

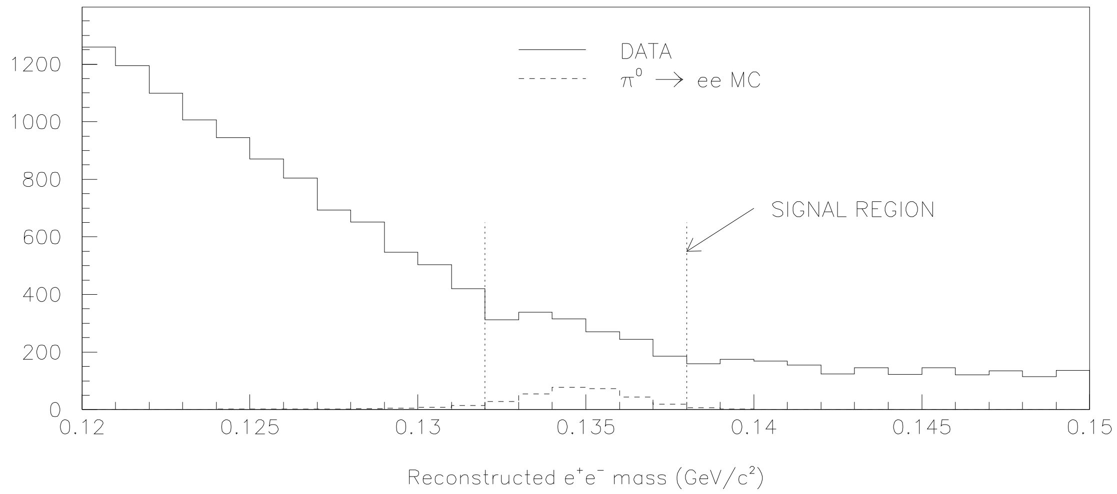

Figure 36. $m_{e^+e^-}$ distribution for $\pi^0 \to e^+e^-$ sample after basic cuts. $\pi^0 \to e^+e^-$ MC distribution is overlaid, normalized to an assumed branching ratio of 7×10^{-8}.

8.3.1 Cuts against mispaired backgrounds

After the loose cuts listed above, the $\pi^0 \to e^+e^-$ sample was dominated by background (Fig. 36). The background shape was, however, well modeled by the MC (Fig. 37), although the absolute levels of the three background components were only predicted to $\sim 15\%$.

At this point, the mispaired four-track background was dominant; the MC predicted 1134 of these events in the signal region. As the four reconstructed photons in these events did not originate from two $\pi^0 \to \gamma\gamma$ decays, we could reduce this background greatly by cutting on the $2\pi^0$ pairing χ^2 (Fig. 38). A cut at $\chi^2 < 4.5$ removed 88% of this mispaired background, with a signal loss of 10%. This cut had little effect on other backgrounds where the $2\pi^0$ pairing was, in general, correct. The data and background MC after this cut are shown in Fig. 39 and Fig. 40. An excess above the background from $\pi^0 \to e^+e^-$ is clearly visible in the signal region.

Because the mispaired backgrounds were nearly flat in both $m_{e^+e^-}$ and reconstructed kaon mass, they could be reduced further by tightening the $|m_{\pi^0\pi^0 ee} - m_{K_L}|$ cut from

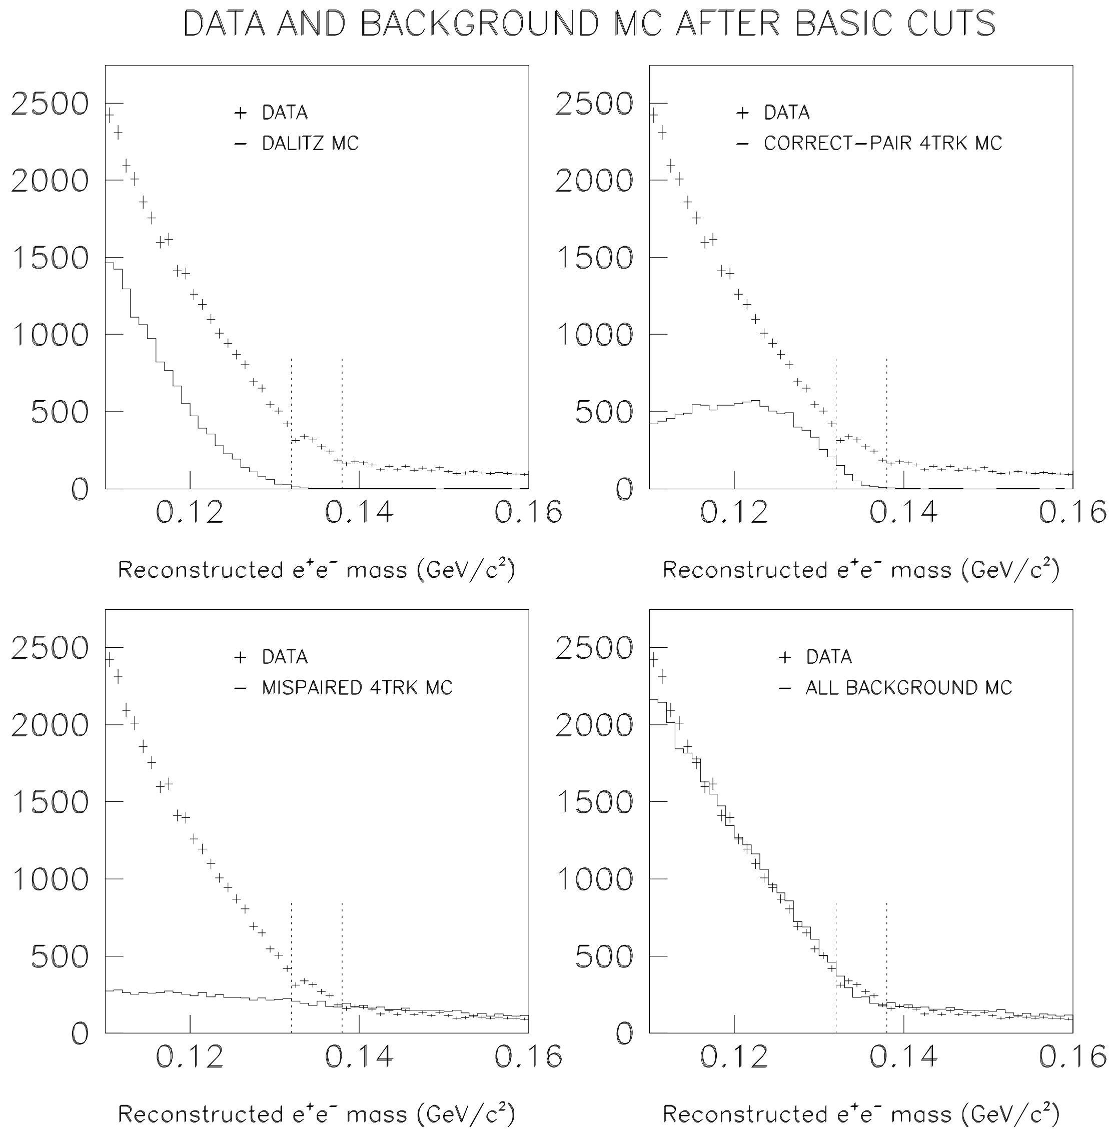

Figure 37. $m_{e^+e^-}$ profiles of the $\pi^0 \to e^+e^-$ data and the three major background categories after basic cuts described in Sec. 8.2.2. The backgrounds are normalized to the flux as measured from the normalization Dalitz sample.

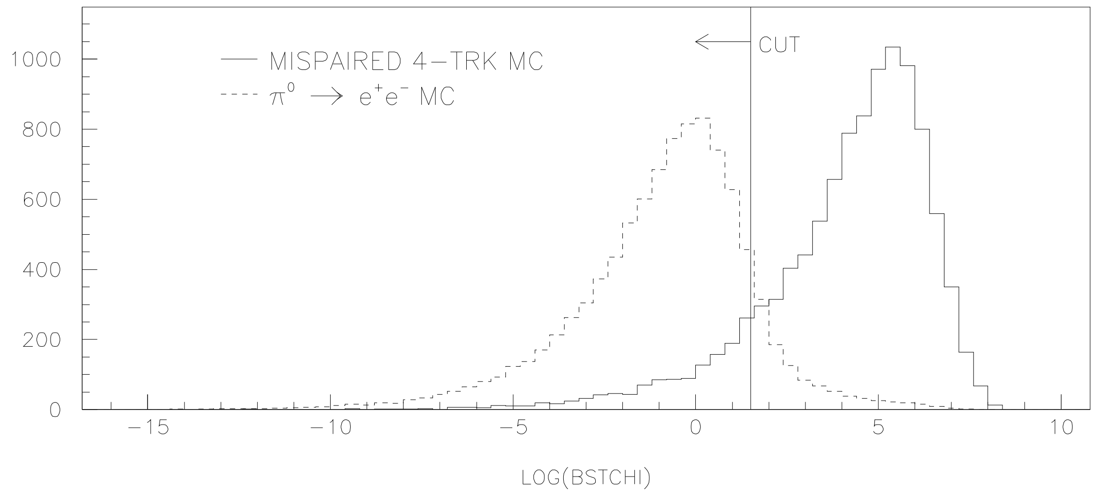

Figure 38. Natural logarithm of the lowest $2\pi^0$ pairing χ^2 for $\pi^0 \to e^+e^-$ MC and mispaired four-track background MC. The vertical scales are arbitrary.

0.050 to 0.010 GeV. Only 0.7% of the $\pi^0 \to e^+e^-$ signal MC was lost by this cut. The cut reduced the mispaired background by a further 80%, to an estimated 29 ± 4 events (the error is from MC statistics only). As a cross-check on the level of background from this source, we looked at the sideband region $0.145 < m_{e^+e^-} < 0.165$ GeV/c^2, which was expected to be completely dominated by mispaired four-track events. Here, the MC predicted 69 ± 6 events, and in the data we observed 56 events. The data and MC predictions after the tight kaon mass cut are shown in Fig. 41.

The other background sources were not significantly affected by the kaon mass cut. The dominant backgrounds at this stage were the correctly-paired four-track modes, and these could not be reduced much further by cuts on kinematics or photon pairing. The MC predicted 300 events from these modes in the signal region.

8.3.2 The DC2 activity cut

The strategy for removing the remaining four-track backgrounds was to search in the drift chamber system for evidence of the two unreconstructed tracks. The strategy was

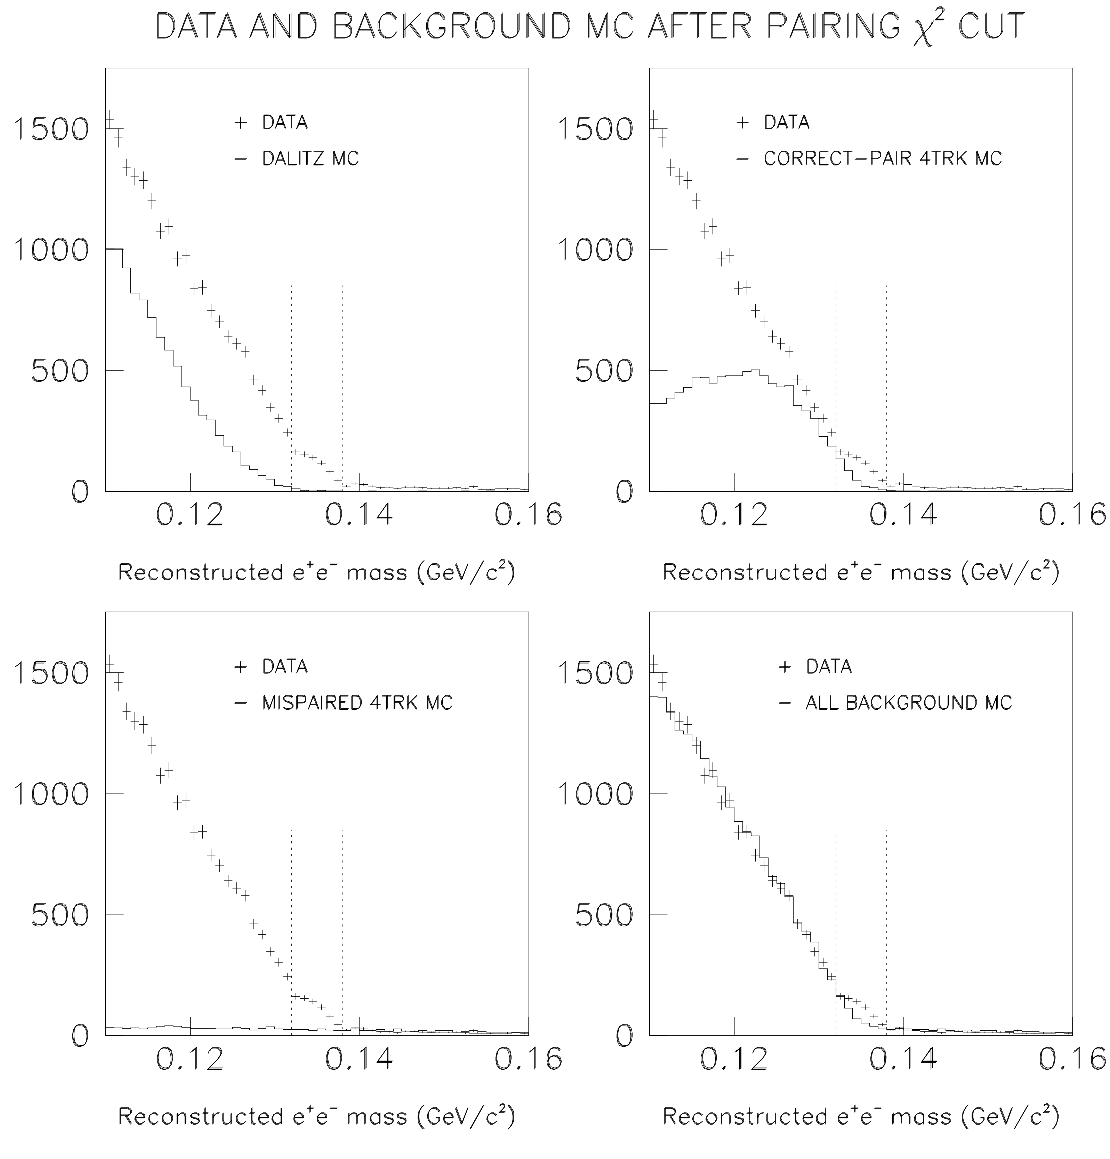

Figure 39. $m_{e^+e^-}$ profiles of the $\pi^0 \to e^+e^-$ data and the three major background categories after basic cuts and a cut on the best $2\pi^0$ pairing χ^2.

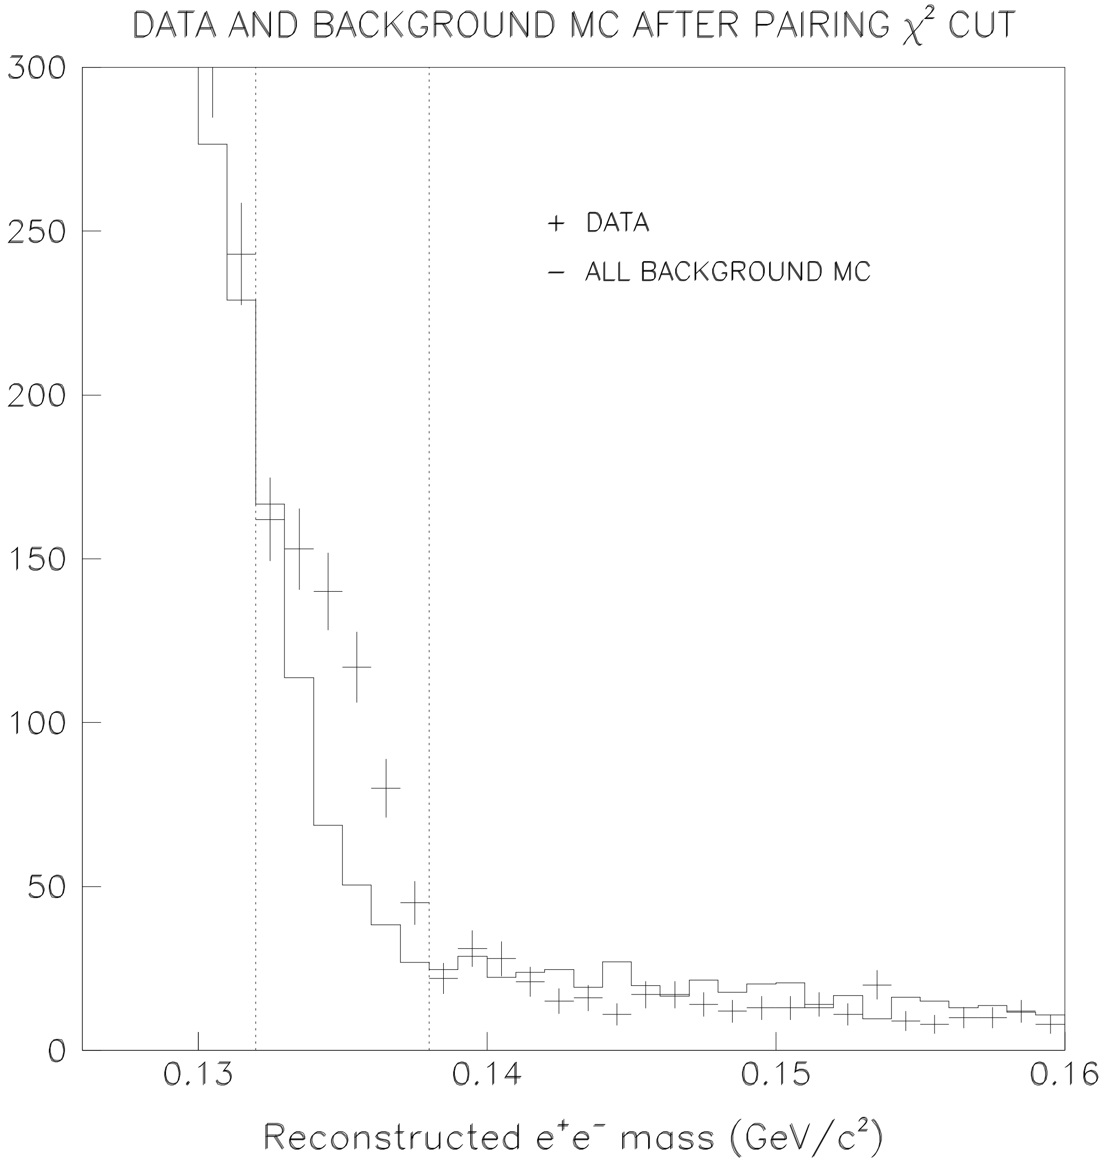

Figure 40. An expanded view of the signal region and high sideband from Fig. 39. Dotted lines indicate the boundaries of the signal region.

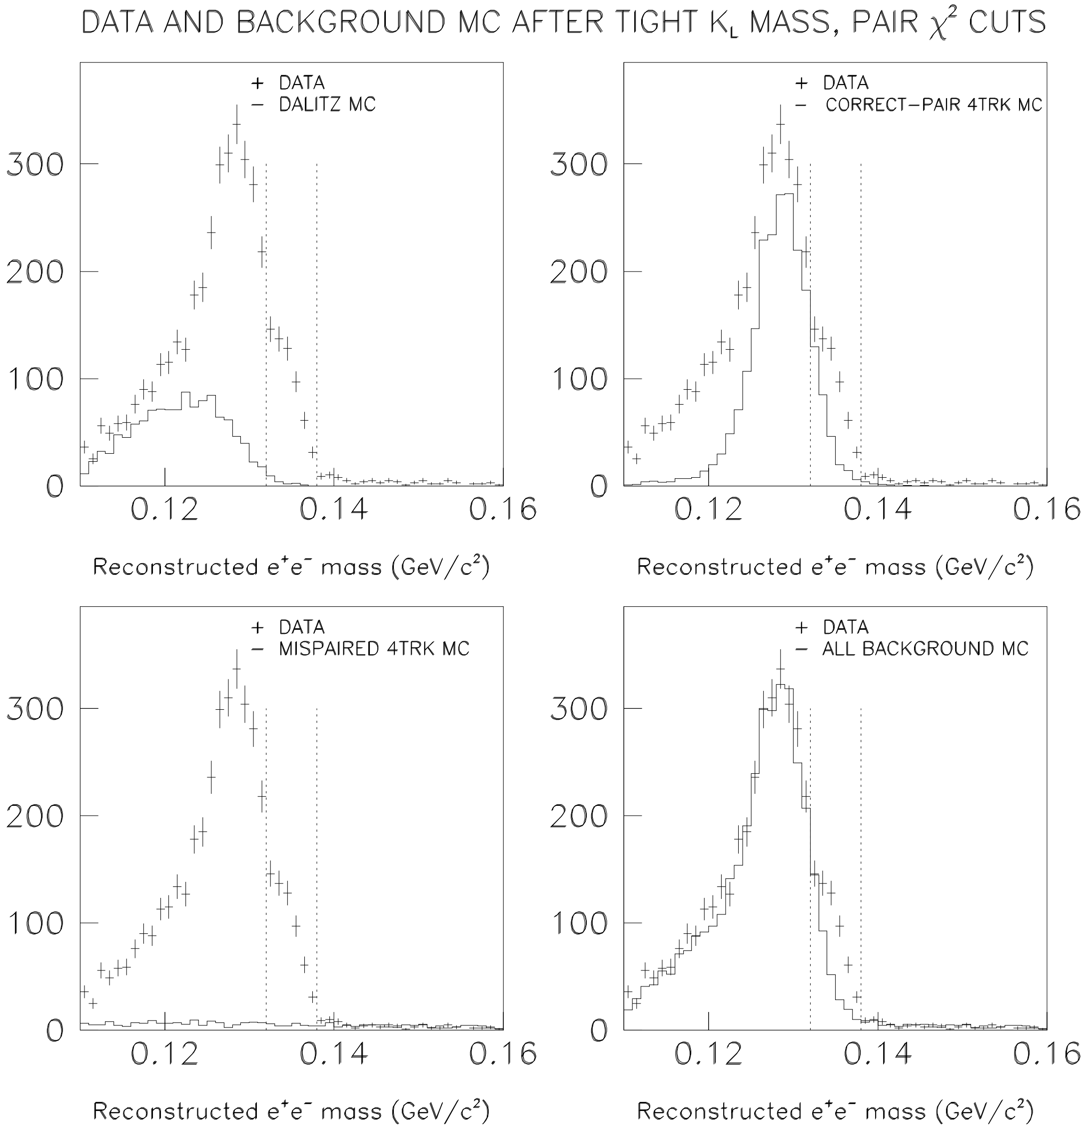

Figure 41. $m_{e^+e^-}$ profiles of the $\pi^0 \to e^+e^-$ data and the three major background categories after basic cuts, the cut on the best $2\pi^0$ pairing χ^2, and a tighter cut (± 0.01 GeV/c^2) on the K_L mass.

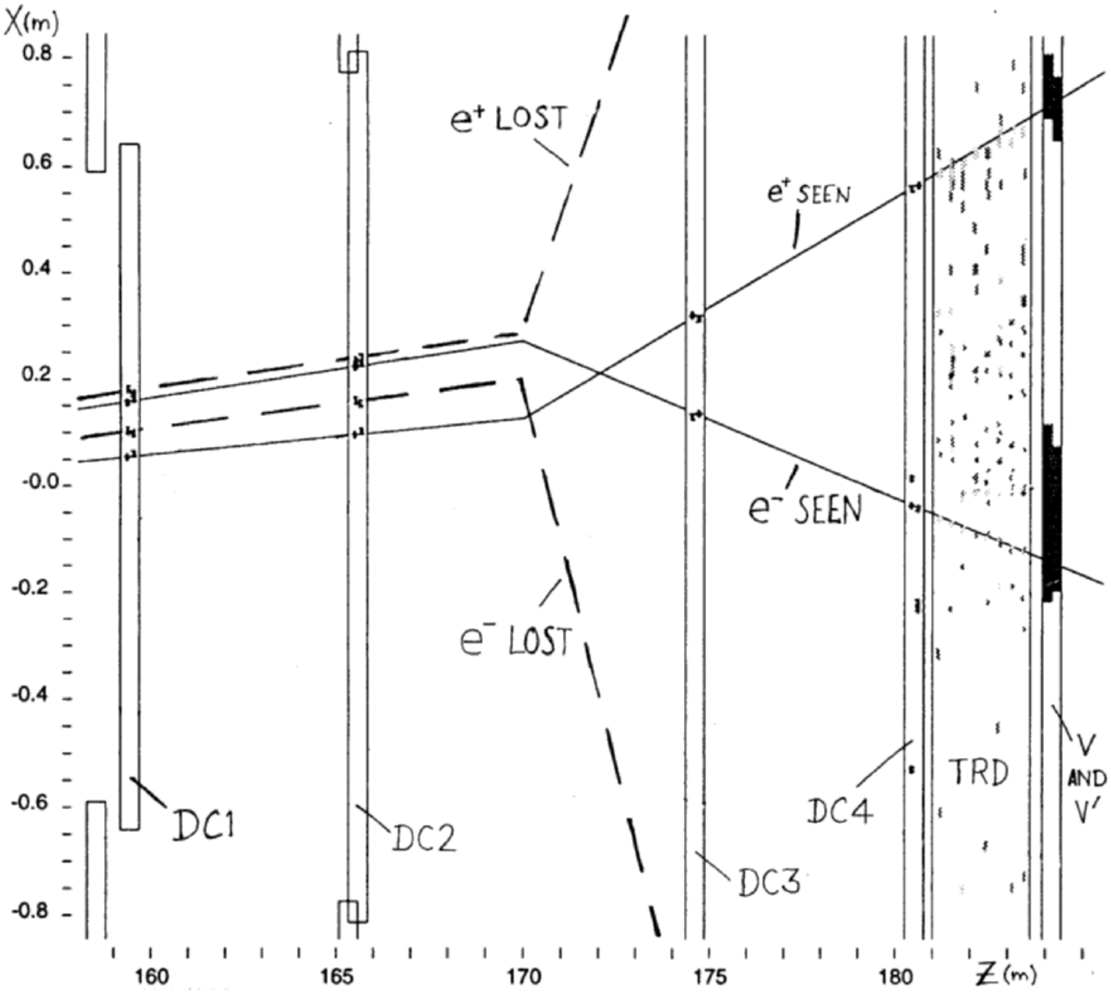

Figure 42. Drift chamber X-view display of a typical four-track background event from the data. The two solid tracks are the reconstructed e^+ and e^- tracks; the dashed lines are possible paths for the two extra tracks which were not reconstructed.

similar to that used in Ref. [8], but the cuts here did not need to be as severe. The superior resolution of the CsI calorimeter allowed us to remove most of the mispaired backgrounds with the cuts on pairing χ^2 and the total invariant mass, so the remaining four-track background level was smaller, and dominated by the correctly-paired component.

Fig. 42 shows a four-track background event. It is clear that there was more activity in the upstream chambers than in the downstream chambers. In particular, there are two in-time hit pairs visible in DC2 which were not part of the two reconstructed tracks.

(Possible paths of the extra tracks have been drawn in.) To reject events of this topology, we vetoed any event with in-time pairs in DC2 not associated with reconstructed tracks. This cost 5% of signal MC (though there is indication that the loss was closer to 8% in the data; this is discussed further in Sec. 8.5.6).

The DC2 activity cut removed $(98.4 \pm 1.6)\%$ of the remaining mispaired four-track background, leaving a residual background of 0.5 ± 0.5 events (this is based on one observed MC event). The correctly-paired four-track background was reduced by $(99.0 \pm 0.4\%)$, leaving 2.8 ± 1.1 expected events in the signal region. The total background remaining from all four-track sources was 3.3 ± 1.2 events. As a cross-check on the rejection of four-track backgrounds by the DC2 actvity cut, we looked again at the high-mass sideband events between $0.145 < m_{e^+e^-} < 0.165$ GeV/c^2, after removing the tight kaon mass cut and pairing χ^2 cuts. In this region, 130 events survived the DC2 activity cut in the data; the MC predicted 114 ± 8 events. This 14% discrepancy is not statistically significant, and would correspond to fewer than 0.5 events in the final background estimate.

8.3.3 Final background levels

The background MC and data after all cuts are shown in Fig. 43. The evidence for the $\pi^0 \to e^+e^-$ signal is clear. The dominant background after all cuts came from the high-$m_{e^+e^-}$ Dalitz decays, which were expected to contribute 18.1 ± 3.0 events. This uncertainty is from MC statistics alone. The disturbing aspect of this calculation was that, unlike the four-track backgrounds where there were high-mass events to study, the Dalitz background was a prediction of the extreme tail of the misreconstructed mass distribution, which was completely hidden beneath the $\pi^0 \to e^+e^-$ signal. It was not obvious that the MC could predict this tail properly. In order to know how badly we misreconstructed the events which contributed to the $\pi^0 \to e^+e^-$ background, we looked at the difference between the generated and reconstructed $m_{e^+e^-}$ in the Dalitz MC events which fell inside the $\pi^0 \to e^+e^-$ signal region. This is shown in Fig. 44. As seen in the

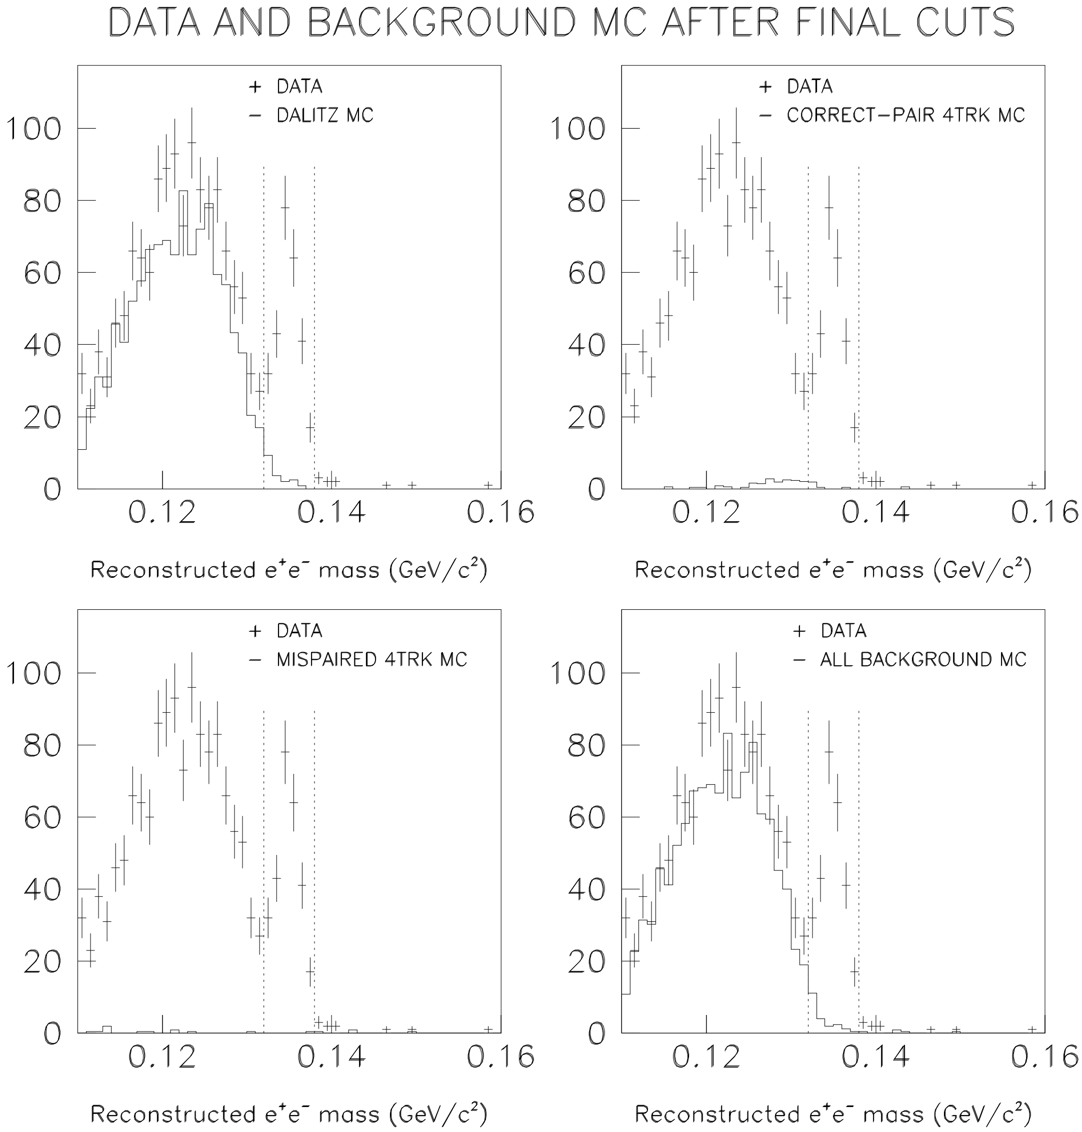

Figure 43. $m_{e^+e^-}$ profiles of the $\pi^0 \to e^+e^-$ data and the three major background categories after basic cuts and a cut on the best $2\pi^0$ pairing χ^2, the tight K_L mass cut, and a cut on in-time activity in DC2.

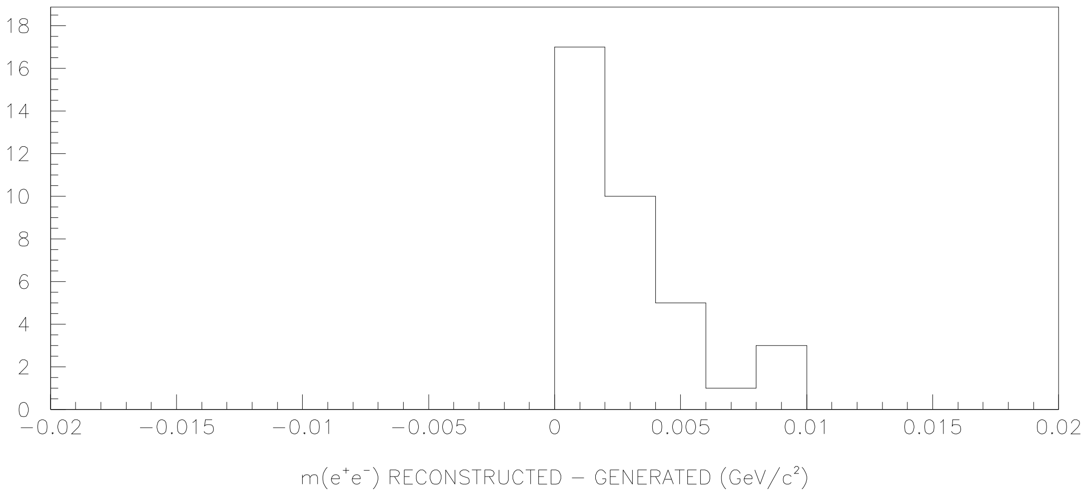

Figure 44. Difference between MC generated and reconstructed e^+e^- invariant mass for MC Dalitz decays which reconstruct in the $\pi^0 \to e^+e^-$ signal region.

figure, the background events had $m_{e^+e^-}$ reconstructed between 0 and 10 MeV/c^2 higher than generated.

As a cross-check on the background estimate, then, we looked at the comparably misreconstructed high-$m_{e^+e^-\gamma}$ events in the tail of the normalization Dalitz decay distribution. In the Dalitz decay case, the high tail of the $m_{e^+e^-\gamma}$ distribution came from two sources: tails in the charged vertex resolution (which also caused the $\pi^0 \to e^+e^-$ background) and selection of the wrong photons to form the two $\pi^0 \to \gamma\gamma$ decays in the event. To isolate the vertex contribution, we looked only at events where the pairing was unambiguous: the selected pairing had to have a good χ^2 (< 4.5) and the second-best pairing a very poor χ^2 (> 50). The distribution of $m_{e^+e^-\gamma}$ for these events is shown in Fig. 45. In the range of interest (up to 10 MeV/c^2 above the center of the mass peak), the data and MC agreed to within 20%. Therefore, we could be confident that the Dalitz background to $\pi^0 \to e^+e^-$ was modeled this well. Assigning a 20% systematic error to the Dalitz background estimate brought this background estimate to 18.1 ± 4.7 events. Adding the four-track backgrounds, we obtained a total background estimate of 21.4 ± 4.9 events, including only the systematic errors mentioned so far.

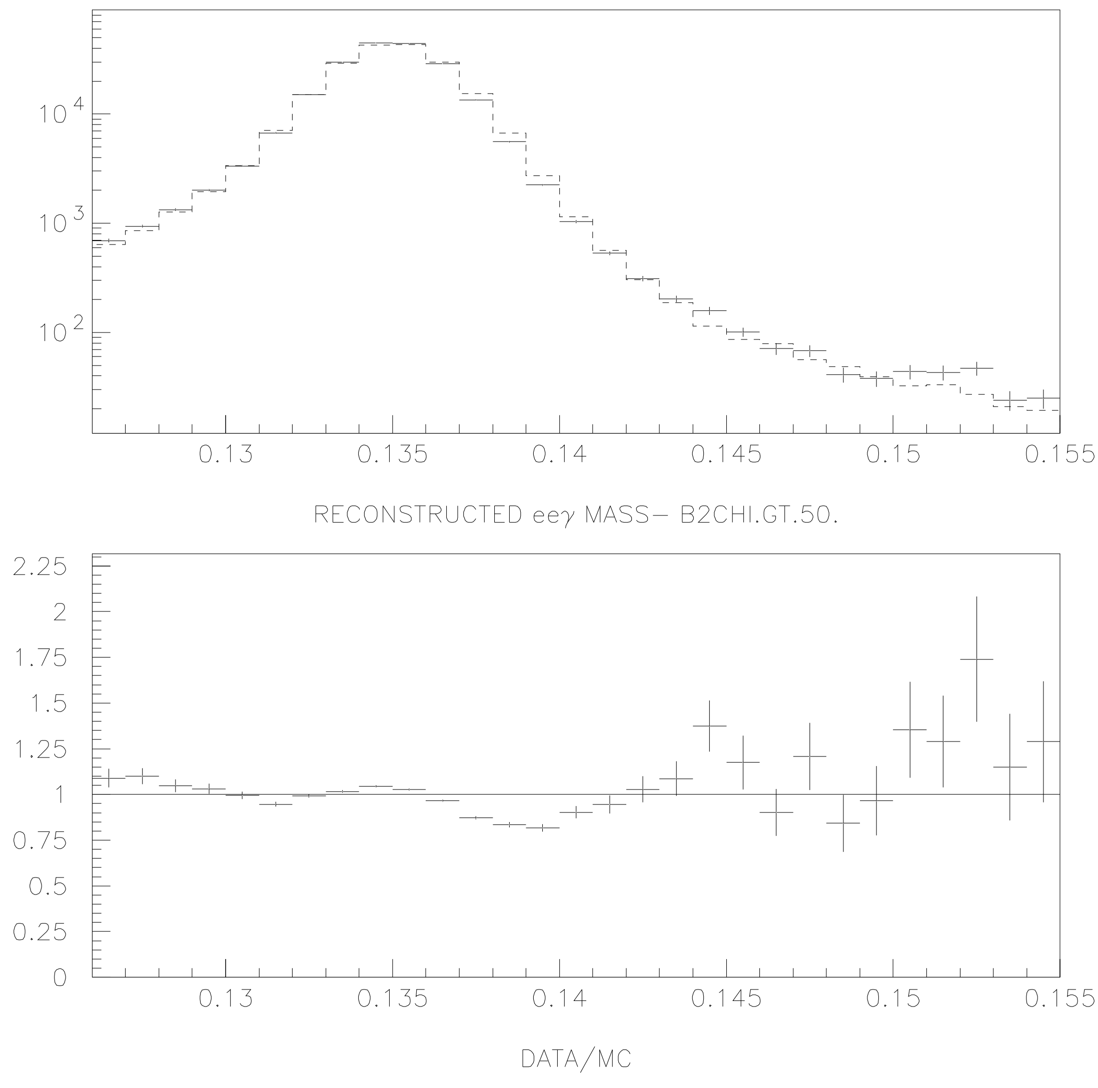

Figure 45. Data (crosses) and MC (dashed) distributions of the $e^+e^-\gamma$ invariant mass in normalization Dalitz decays where the best $2\pi^0$ pairing χ^2 was selected to be below 4.5, and the second-best pairing χ^2 greater than 50. Lower plot is data/MC ratio.

Table 13. Events observed in the $\pi^0 \to e^+e^-$ spectrum in the range $0.110 < m_{e^+e^-} < 0.125$ GeV, compared to background MC. Error on background is from MC statistics only.

	Winter	Summer	TOTAL
Background MC	444 ± 15	344 ± 12	788 ± 19
Data	540	388	928
Excess	$(22 \pm 6)\%$	$(13 \pm 7)\%$	$(18 \pm 5)\%$

An additional systematic error on the background was estimated from the disagreement between the actual and predicted number of events in the low-mass sideband (which was dominated by Dalitz decays). The region between $0.110 < m_{e^+e^-} < 0.125$ GeV$/c^2$ was used because the contribution from $\pi^0 \to e^+e^-$ was expected to be negligible in this region. In the data, 18% more events were seen than expected from the MC (see Table 13). Scaling this excess to the expected contamination in the signal region, we obtained an excess of 3.8 events, which was taken as a systematic error on the background estimation. The total background estimate, with all systematic errors included, was 21.4 ± 6.2 events.

8.4 Acceptance and data after final cuts

In the winter MC, the acceptance after all cuts was 1.007% for Dalitz decays and 2.21% for $\pi^0 \to e^+e^-$ decays. In the summer, the acceptance was 1.061% for Dalitz decays and 2.31% for $\pi^0 \to e^+e^-$. The acceptance increased in the summer MC by 5.4% for $\pi^0 \to e^+e^-\gamma$ and 4.5% for $\pi^0 \to e^+e^-$. The acceptance statistics are summarized in Table 14.

After all cuts, the total number of events in the normalization sample was 650264 and the total in the signal sample was 275. The breakdown between the runs is shown in table 15.

The shape of the mass peak (Figs. 46 and 47), and the observed correlation between the $m_{e^+e^-}$ and m_K distributions (Fig. 48) agreed well with the expectation from the

Table 14. Number of generated signal and normalization MC events and acceptance statistics at each stage of processing.

	Winter $\pi^0 \to e^+e^-\gamma$ MC	Winter $\pi^0 \to e^+e^-$ MC	Summer $\pi^0 \to e^+e^-\gamma$ MC	Summer $\pi^0 \to e^+e^-$ MC
Generated	7.05×10^7	3.00×10^6	6.17×10^7	2.63×10^6
Pass L2 Trig.	7.31×10^6	396112	5.73×10^6	321616
Reconstructed	1.138×10^6	123210	1.067×10^6	114324
Pass Basic Cuts	802510	93186	754236	87281
Pass Final Cuts	709669	66374	653825	60644

signal and background Monte Carlo. Note that the background Monte Carlo in Fig. 46 is normalized to the absolute flux, while in Fig. 47 it is normalized to the low-mass sideband.

Table 15. Final data sample sizes for signal and normalization modes.

	Normalization Data		Signal Data	
	Winter	Summer	Winter	Summer
Events observed	373777	276487	159	116
Total	650264		275	

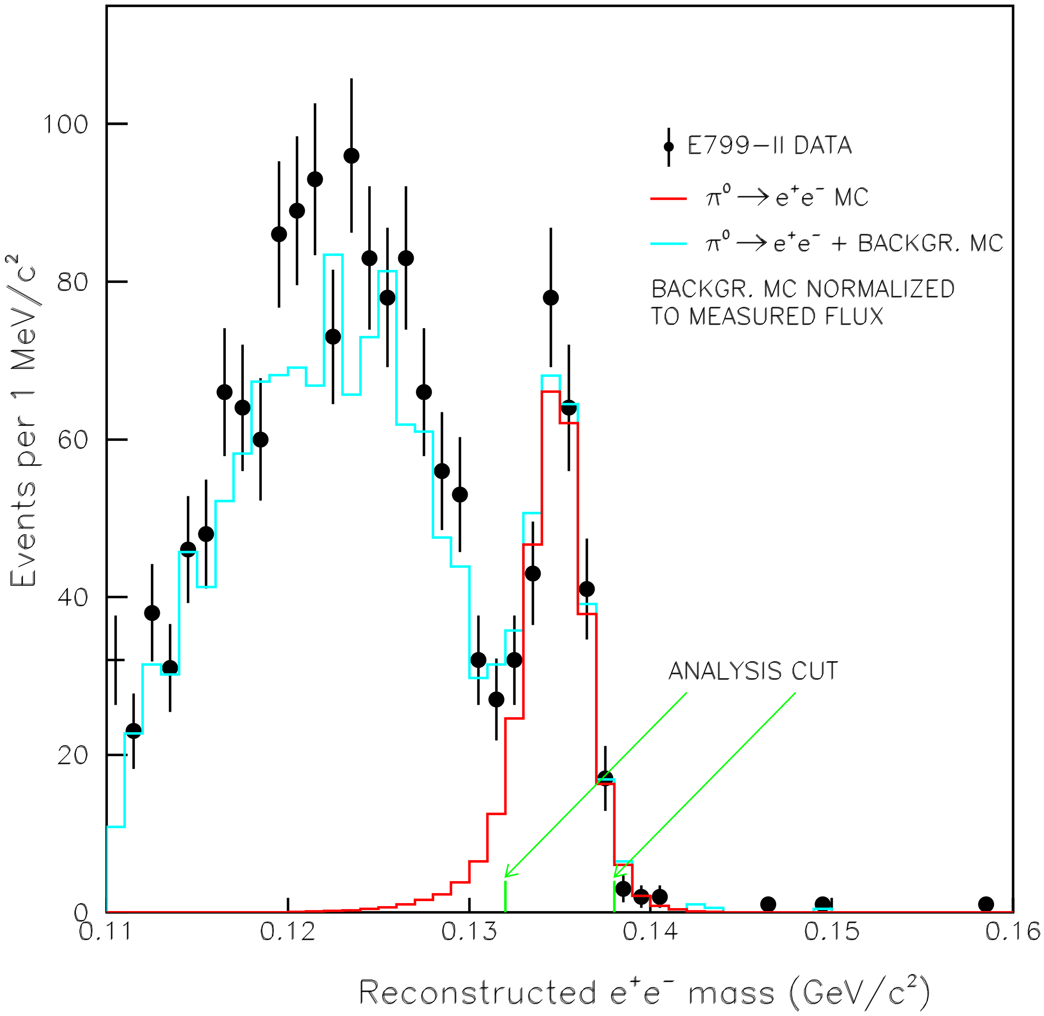

Figure 46. The $m_{e^+e^-}$ spectrum of $\pi^0 \rightarrow e^+e^-$ candidates, after all cuts. The points with errors are data. The lower histogram is $\pi^0 \rightarrow e^+e^-$ MC; the upper histogram is the $\pi^0 \rightarrow e^+e^-$ MC added to background MC. The Dalitz MC is normalized to the measured flux; the $\pi^0 \rightarrow e^+e^-$ MC is normalized to the background-subtracted signal. There are 275 events in the signal region in the data.

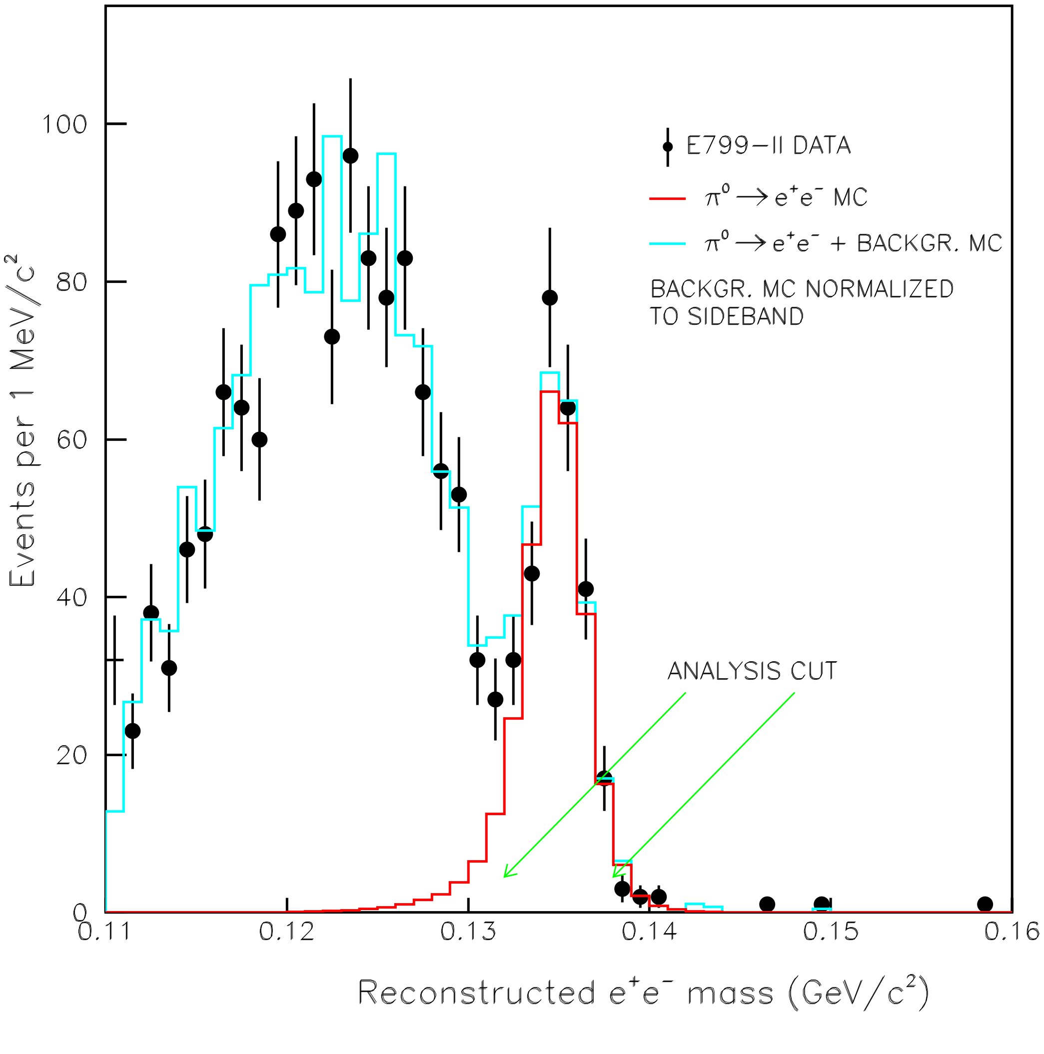

Figure 47. The $m_{e^+e^-}$ spectrum of $\pi^0 \rightarrow e^+e^-$ candidates, after all cuts. The points with errors are data. (The data points are the same as in Fig. 46.) The lower histogram is $\pi^0 \rightarrow e^+e^-$ MC; the upper histogram is the $\pi^0 \rightarrow e^+e^-$ MC added to background MC. The Dalitz MC is normalized to the $0.110 < m_{e^+e^-} < 0.125$ region. The $\pi^0 \rightarrow e^+e^-$ MC is normalized to the background-subtracted signal.

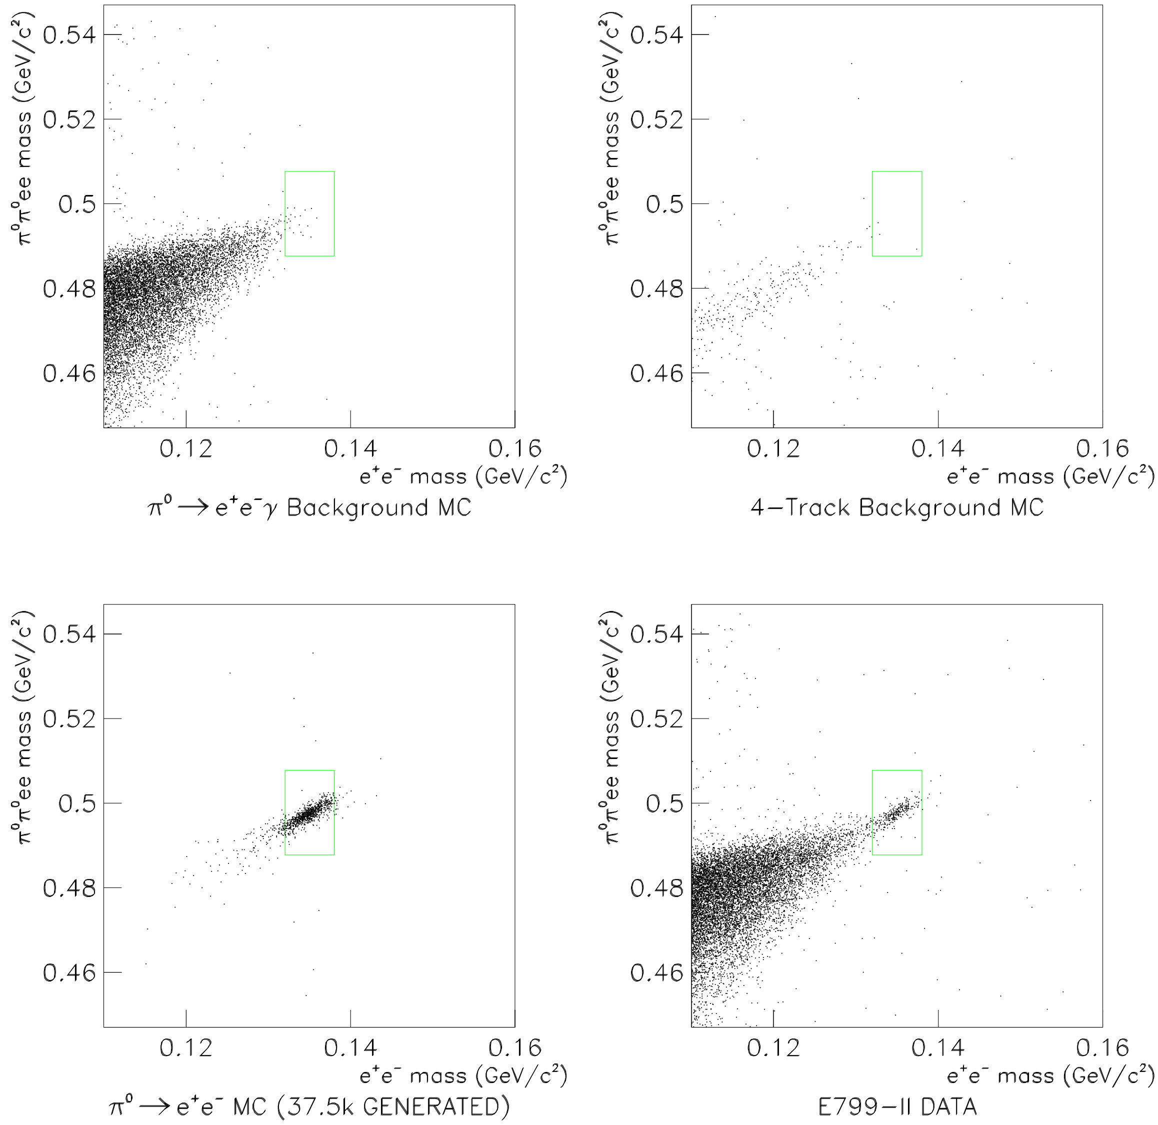

Figure 48. Reconstructed K_L mass vs. e^+e^- mass for Dalitz and four-track backgrounds (background statistics correspond to approximately 1.1 times the data), $\pi^0 \to e^+e^-$ signal (37500 generated events– approximately 5.5 times the data), and E799-II data. The box in the middle of the figures shows the accepted region.

8.5 Systematic error studies

Before calculating the $\pi^0 \rightarrow e^+e^-$ branching ratio, we considered additional systematic errors beyond the background subtraction. These errors included tracking inefficiencies in the beam regions of the chambers, biases from data/MC disagreements in $2\pi^0$ pairing and reconstructed mass resolution, and uncertainties in the branching ratio and form factor of the π^0 Dalitz decay.

8.5.1 Chamber problems in beam regions

The high-side tail in the drift chamber SOD distribution (see Sec. 6.1.1.1) caused a significant ($\sim$1%) inefficiency for finding tracks. This inefficiency was concentrated in the regions of the chambers where the neutral beams passed, and was a problem primarily in the upstream chambers (1 and 2).

To study the effect of this inefficiency, the distribution of an integer variable called INBEAMS was studied in Dalitz decays. INBEAMS was the number of upstream beam regions crossed by tracks in an event, where the beam region was defined as the exact sizes of the neutral beams at the chambers. The beam regions in the summer were therefore larger than in the winter.

The INBEAMS variable ranged from 0 to 4 for a two-track decay (there were two tracks, and each track could potentially cross two beam regions). When, for example, both tracks were outside the beams at both chambers, then INBEAMS=0. When one track was completely outside the beams, but the other track was in a beam at one chamber, then INBEAMS=1. When both tracks in the event were in a beam region at Chamber 1 but were outside the beam at Chamber 2, then INBEAMS=2. When one track was in the beam region at both chambers, and the other track was outside the beams at both chambers, then INBEAMS=2 as well. The cases where INBEAMS=3 generally had both tracks in the beam at Chamber 1, but one track outside the beam at Chamber 2. When both tracks were in beam regions at both chambers, then INBEAMS=4 (this was a rare occurence in high-$m_{e^+e^-}$ events).

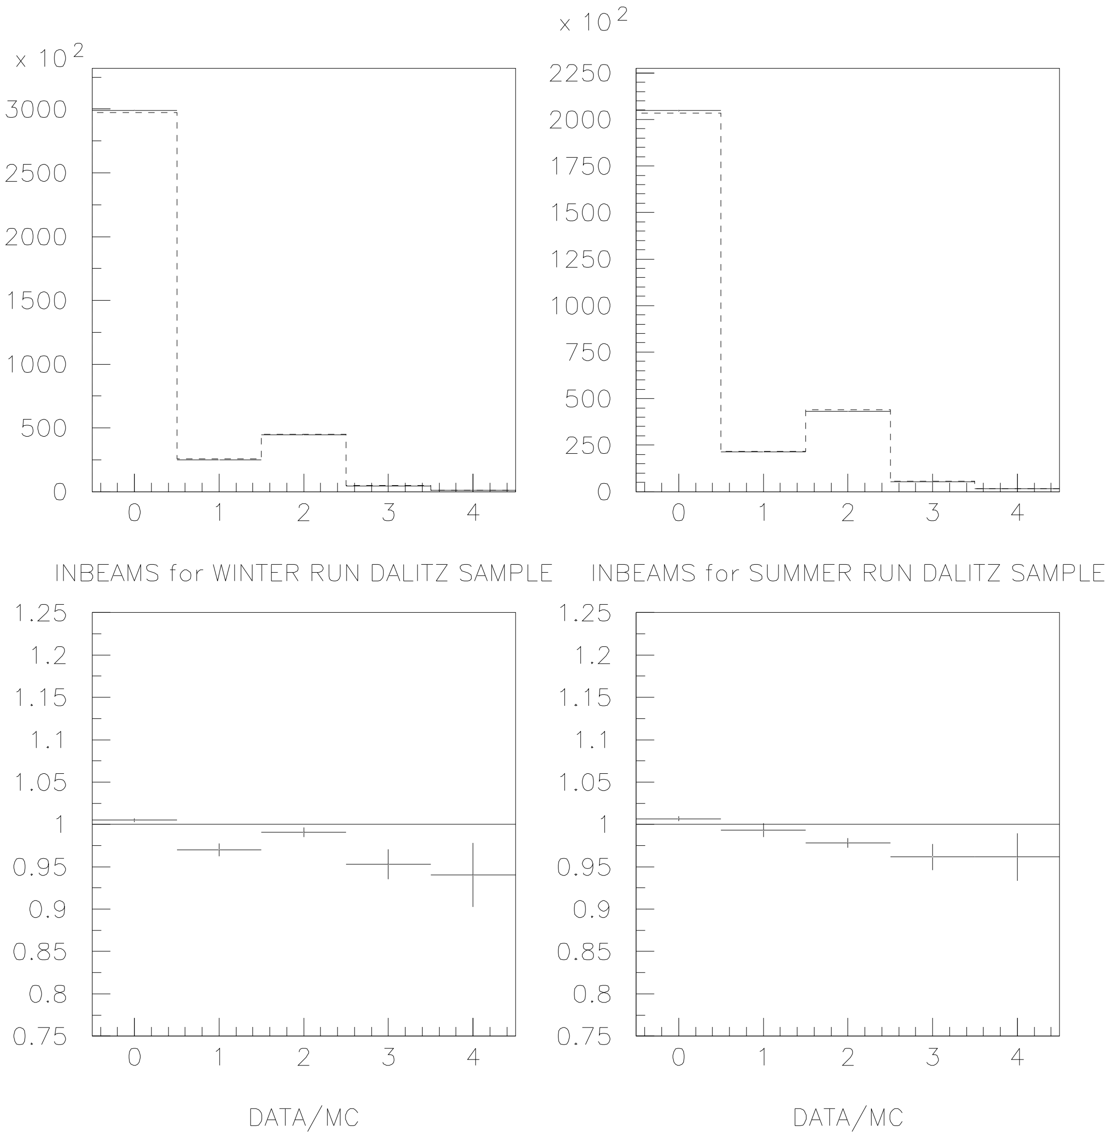

Figure 49. INBEAMS variable for Dalitz decays. Winter run shown on left; summer run on right. Data and MC are normalized to the same total number of events.

Table 16. The bias on a $\pi^0 \to e^+e^-$ measurement resulting from beam-region inefficiencies in Chambers 1 and 2.

Track-finding inefficiency in beams (%)	Normalization MC loss (%)	Signal MC loss (%)	$\pi^0 \to e^+e^-$ bias (%)
0.5	0.24	0.22	0.02
1.0	0.48	0.43	0.05
1.5	0.72	0.65	0.07
2.0	0.96	0.86	0.10
3.0	1.44	1.28	0.16
5.0	2.47	2.12	0.35

The distribution of INBEAMS is shown in Fig. 49. There was a significant deficit of events where tracks passed through multiple beam regions. (The apparent excess of events with INBEAMS=0 was due to the normalization of the plots to the same total number of events.) A fit to the distribution indicated that the loss was $1.12 \pm 0.26\%$ per beam region in the winter, and $1.28 \pm 0.25\%$ in the summer.

In order to study the bias in the $\pi^0 \to e^+e^-$ branching ratio measurement, we looked at the illumination of signal and normalization mode tracks in the chambers (Fig. 50), and artificially applied inefficiencies in those regions in the MC. It should be noted that studying the problem in this way made the dubious assumption that all upstream beam regions had the same inefficiency; in reality, the problems were somewhat worse in Chamber 1. Therefore, inefficiencies larger than those observed were studied. The results are summarized in Table 16. The $\sim$1% inefficiencies observed in the data corresponded to less than a 0.1% bias on the $\pi^0 \to e^+e^-$ measurement. Even a 5% inefficiency, much larger than *any* observed effect, would only introduce a 0.35% bias. These effects were negligible in this measurement.

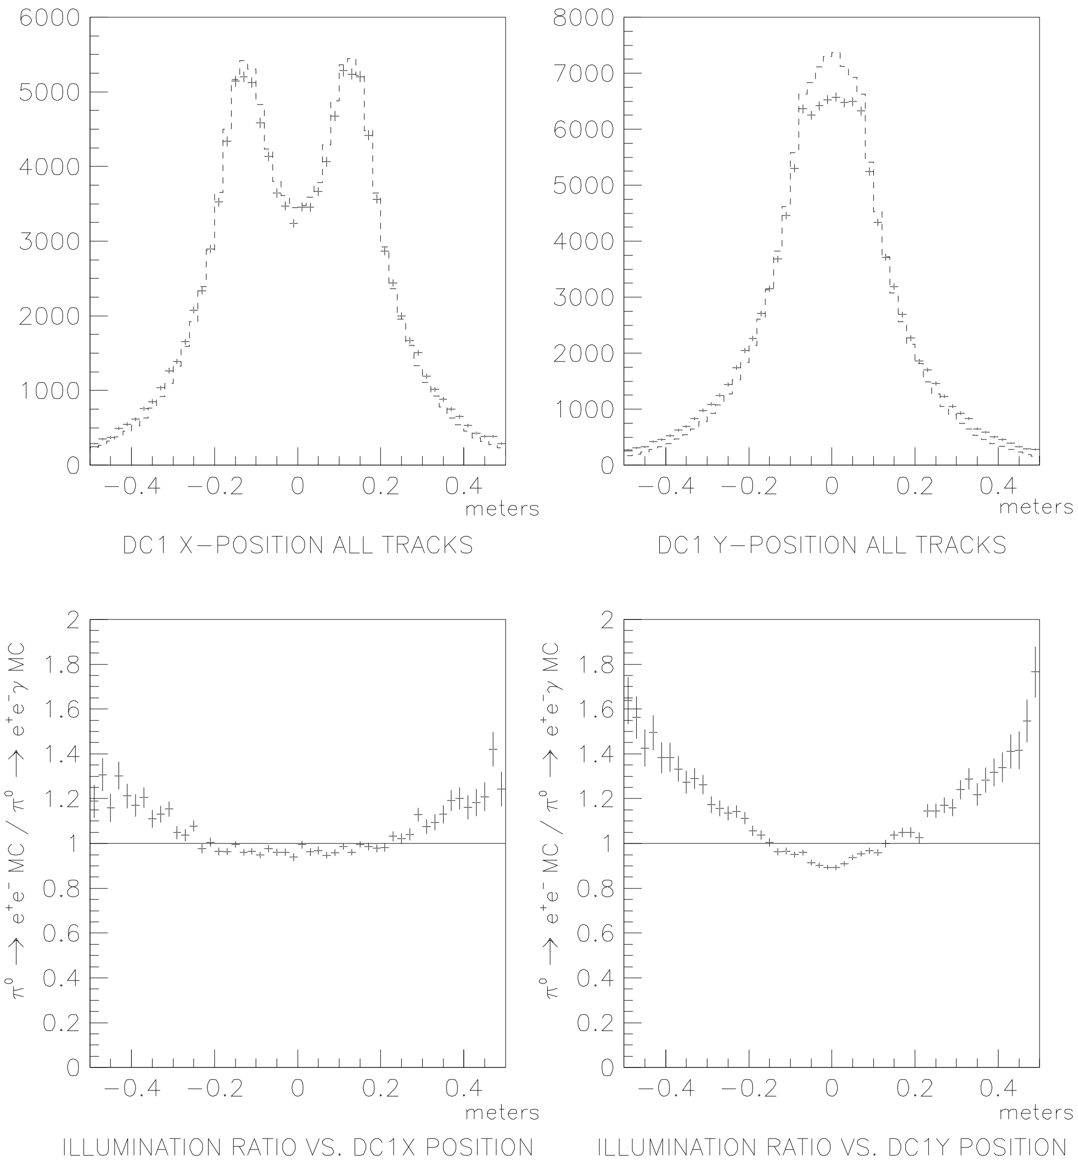

Figure 50. Upper plots: DC1 X- and Y-view illumination for $\pi^0 \to e^+e^-$ MC (crosses) and $\pi^0 \to e^+e^-\gamma$ normalization MC (dashed histogram) after all cuts. Lower plots: Ratio of illumination ($\pi^0 \to e^+e^-$ MC / $\pi^0 \to e^+e^-\gamma$ MC). Plots are normalized to the same number of reconstructed events.

8.5.2 Pairing χ^2 cut efficiency

The $2\pi^0$ pairing χ^2 distribution (Fig. 51) agreed well between MC and data except for the extreme high tail, where there was an excess in the data. To reject mis-paired four-track backgrounds to $\pi^0 \to e^+e^-$, we placed a cut on this variable at $\chi^2 = 4.5$ ($\log(\chi^2) = 1.5$). In order to study the effect of this cut, we varied the cut and observed its efficiency for data and MC Dalitz decays over the range where the data-MC disagreement was significant. Varying the cut from $\chi^2 = 4.5$ down to 2.2 and up to 20 caused the measured flux to change by 0.8% in the winter data and 1.2% in the summer.

The efficiency of the cut in MC for events which pass all other cuts was 0.927 for Dalitz decays, and 0.901 for $\pi^0 \to e^+e^-$. If the data/MC disagreement were merely a resolution problem, it would have canceled to a high degree in the ratio. However, the disagreement was only in the high tail, and as such may have been due to a discrete misreconstruction, such as mispairing, which could have been much more likely in the 5-photon Dalitz mode than in the 4-photon $\pi^0 \to e^+e^-$ mode. Therefore, we conservatively considered the entire "walk" of the flux measurement to be a systematic error. Taking a weighted average of the summer and winter run discrepancies, the systematic error in the final result was estimated to be 1.0%.

8.5.3 Mass resolution

A significant systematic error came from modeling of the resolution of the π^0 mass resolution in the Dalitz mode. The cut on $m_{e^+e^-}$ in the signal mode had to be very tight to reject background. The cut at $|m_{e^+e^-} - m_{\pi^0}| = 3$ MeV/c^2 was at 1.8σ, while the equivalent cut on the normalization mode ($|m_{e^+e^-\gamma} - m_{\pi^0}| = 30$ MeV/c^2) was much looser (17σ). This cut was loose in order to include most Dalitz decays where the wrong $2\pi^0$ pairing was selected (this resulted in a large nongaussian tail in the normalization mode mass spectrum). The effect of mispairings can easily be seen in the distribution of $m_{e^+e^-\gamma}$ (Fig. 52) and in a scatter plot of $m_{e^+e^-\gamma}$ vs. $m_{\pi^0\pi^0e^+e^-\gamma}$ (Fig. 53). The shape

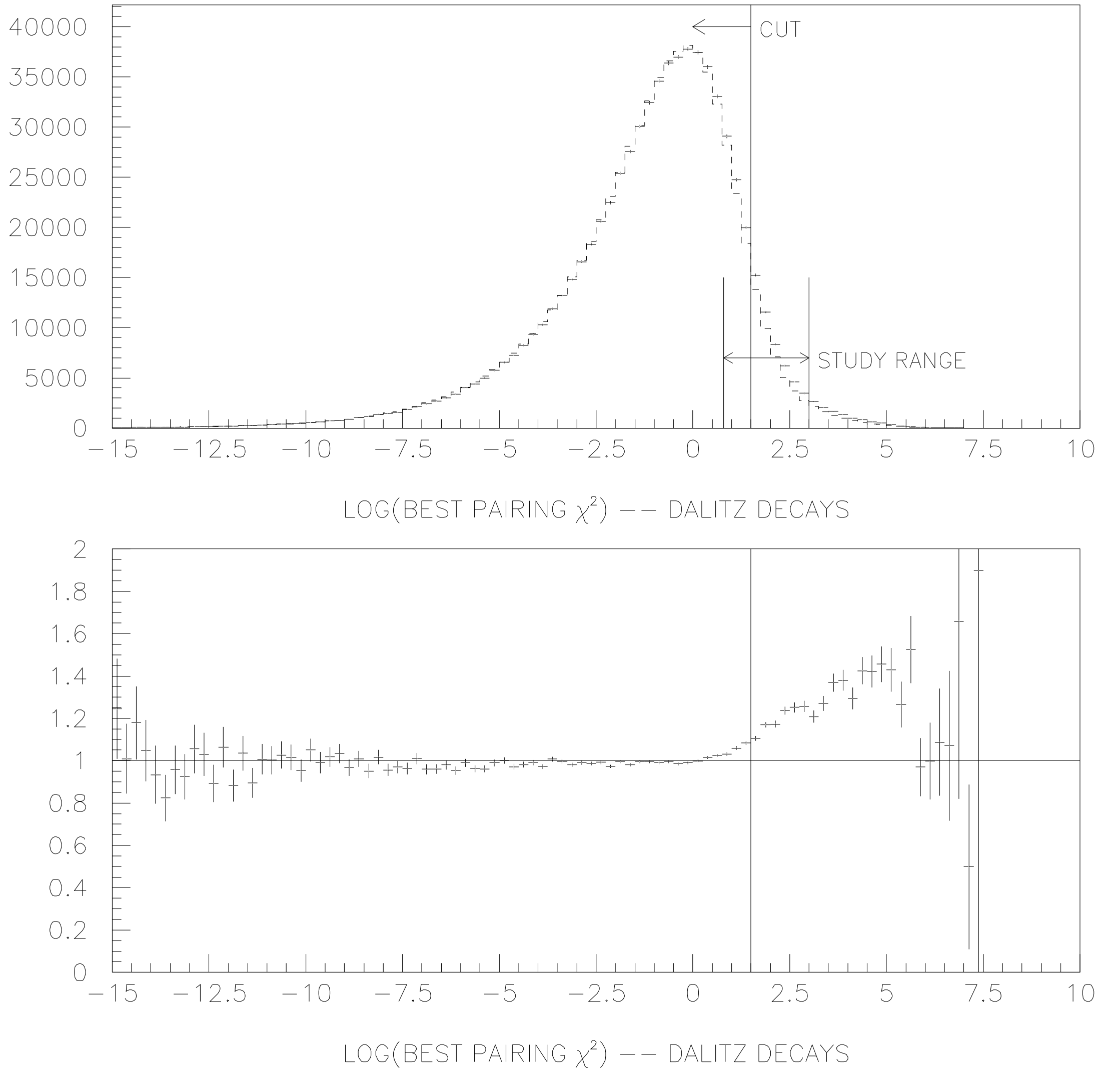

Figure 51. Natural log of best $2\pi^0$ pairing χ^2 for normalization Dalitz decays which pass all other cuts. Winter and summer data are combined. Data are crosses; MC is dashed histogram. Cut and cut study regions are shown. Lower plot is data/MC ratio.

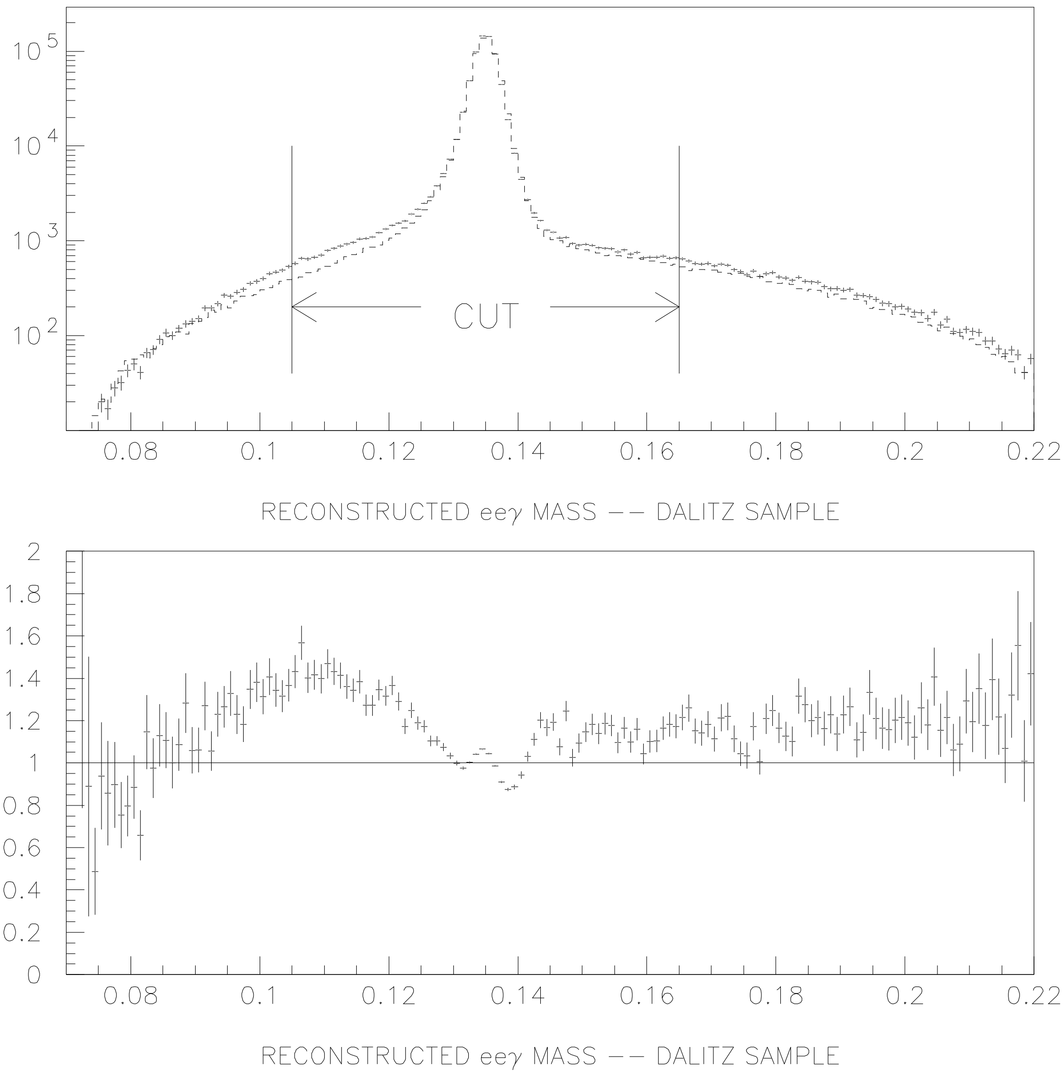

Figure 52. Reconstructed π^0 Dalitz mass ($m_{e^+e^-\gamma}$) in GeV/c^2 for $\pi^0 \to e^+e^-\gamma$ Monte Carlo (dashed histogram) and data events (crosses) which pass other normalization cuts. Lower plot is data/MC ratio.

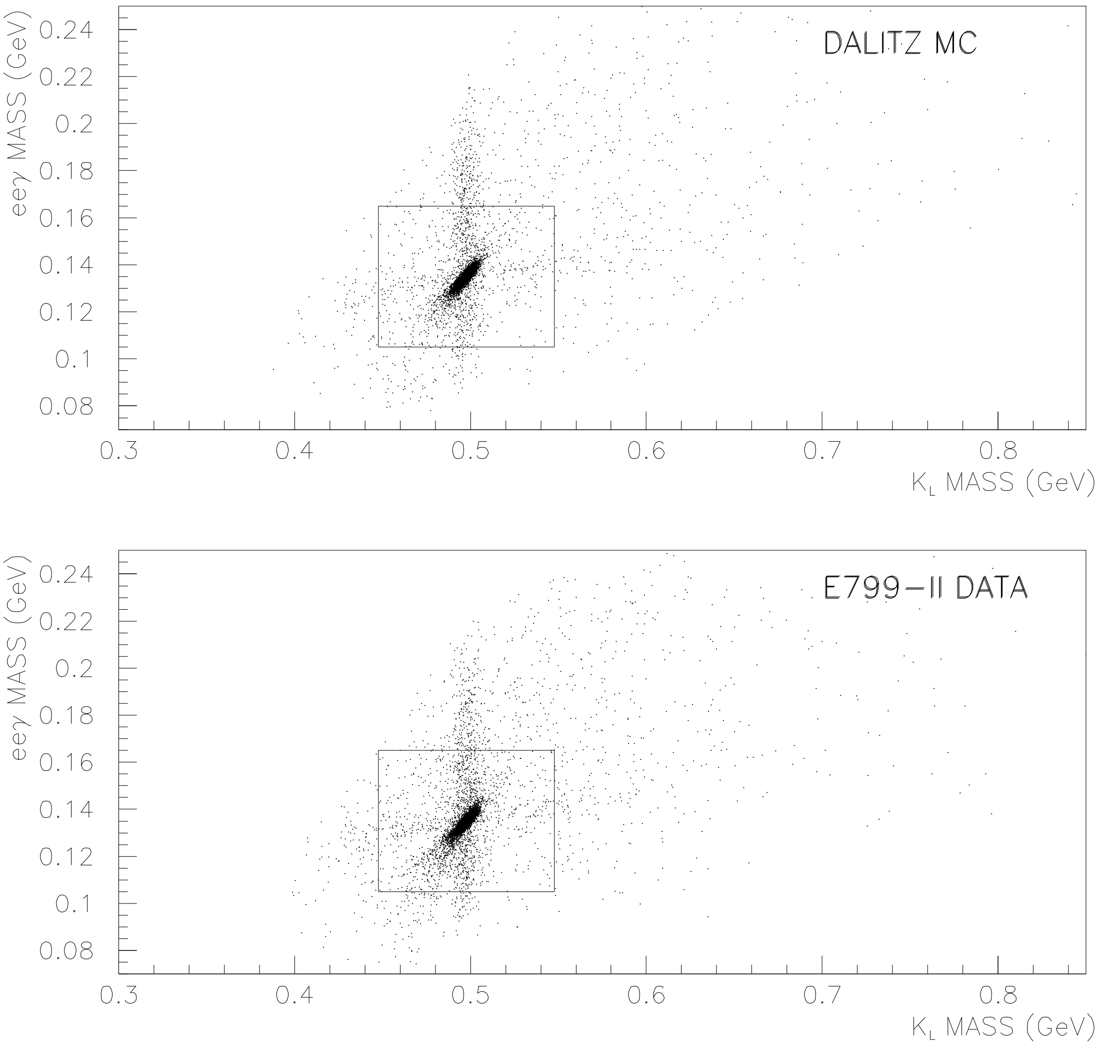

Figure 53. Reconstructed π^0 Dalitz mass ($m_{e^+e^-\gamma}$) vs. K_L mass ($m_{\pi^0\pi^0e^+e^-\gamma}$) for $\pi^0 \to e^+e^-\gamma$ Monte Carlo and data events which pass other normalization cuts. Scatter plot contains about 1/45 of the data. The tilted "galaxy-shaped" distribution contains the correctly paired events (which dominate the statistics); the vertical and near-horizontal bands are events where the wrong $2\pi^0$ pairing was selected. The normalization mode mass cuts are shown as a box.

of the high tail (almost entirely due to mispairing) was well modeled, but its absolute level was off by $\sim$15%.

Other data/MC disagreements were in the low (radiative) tail and the central peak. The radiative tail was larger in the data than in the MC, and the resolution of the central mass peak was somewhat narrower in data than in MC. A Gaussian fit of the region $0.130 < m_{e^+e^-\gamma} < 0.140$ showed the resolution to be $\sigma = 1.86$ MeV/c^2 in MC and $\sigma = 1.79$ MeV/c^2 in data. To study the effect of these discrepancies on the measurement, the π^0 mass cut was varied from no cut at all to ± 2 MeV/c^2 (1.1σ). Over the full range studied, the measured flux varied by $\pm 1.2\%$ in both the winter and summer data.

The component of this discrepancy from charged vertex resolution effects would be expected to cancel somewhat in the branching ratio measurement, if the mass cut were similar between the two modes. However, it was much tighter in the signal mode than in the normalization. In addition, there was no mispairing tail in the signal mode because the e^+e^- mass came entirely from the charged tracks (Fig. 54). Therefore, because the π^0 mass determination was so different between the two modes, we could not expect the data/MC discrepancies to cancel in the ratio, and the $\pm 1.2\%$ variation in the Dalitz flux was treated as a systematic error.

After the charged mass cuts, the final cuts on K_L mass for both signal and normalization modes were very loose (5.6σ for signal, 21σ for normalization). The cut was nearly 100% efficient for signal MC. For the normalization mode, the cut efficiency was 98.69% for MC and 98.50% for data (Fig. 55). As the cut was very efficient and the data/MC discrepancy was so small, no additional systematic error was calculated for this cut.

8.5.4 Transverse momentum resolution

The cut on transverse momentum rejected 0.6% of the signal MC and 1.9% of normalization MC. The transverse momentum distribution (Fig. 56) agreed well between data and MC in the normalization mode over most of the range. The mean of the p_T distribution was actually 5% smaller in the data than in the MC. As the cut is loose, this disagreement produced no significant bias.

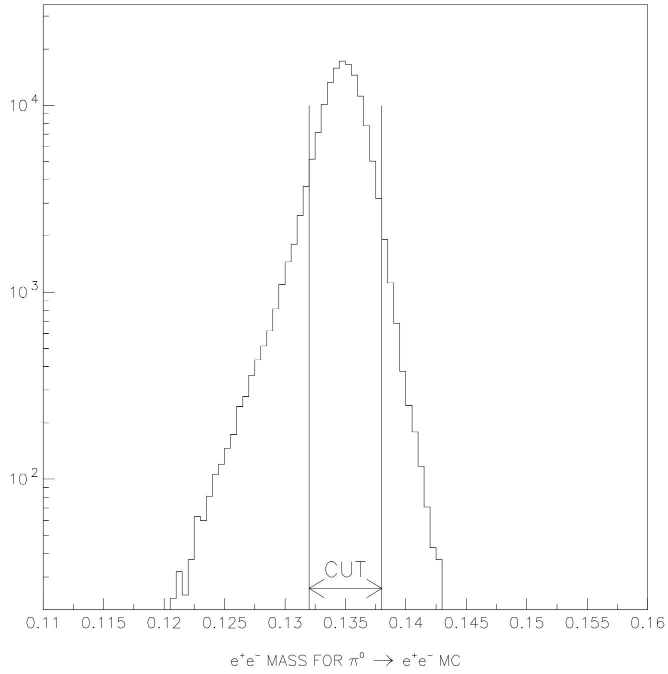

Figure 54. Reconstructed $m_{e^+e^-}$ for $\pi^0 \to e^+e^-$ Monte Carlo; the final mass cut is shown.

However, In the data, there was also an excess of events at high $|p_T|$ ($\ln|p_T| > -3.5, |p_T| > 0.03$ GeV/c). These were mostly from four-track background events (the same four-track backgrounds which were a background to $\pi^0 \to e^+e^-$), and they were removed by the p_T cut. Adding the four-track background to the MC (Fig. 57) brought the data and MC into better agreement in the high-p_T tail. The remaining integrated discrepancy was 0.5%. Most of this excess consisted of high-p_T, good-mass events which came from decays of kaons which scattered in the defining collimator. These events were equally present in both the signal and background modes, so cutting them out caused no bias in the ratio. We did not expect the p_T cut to cause a significant systematic error.

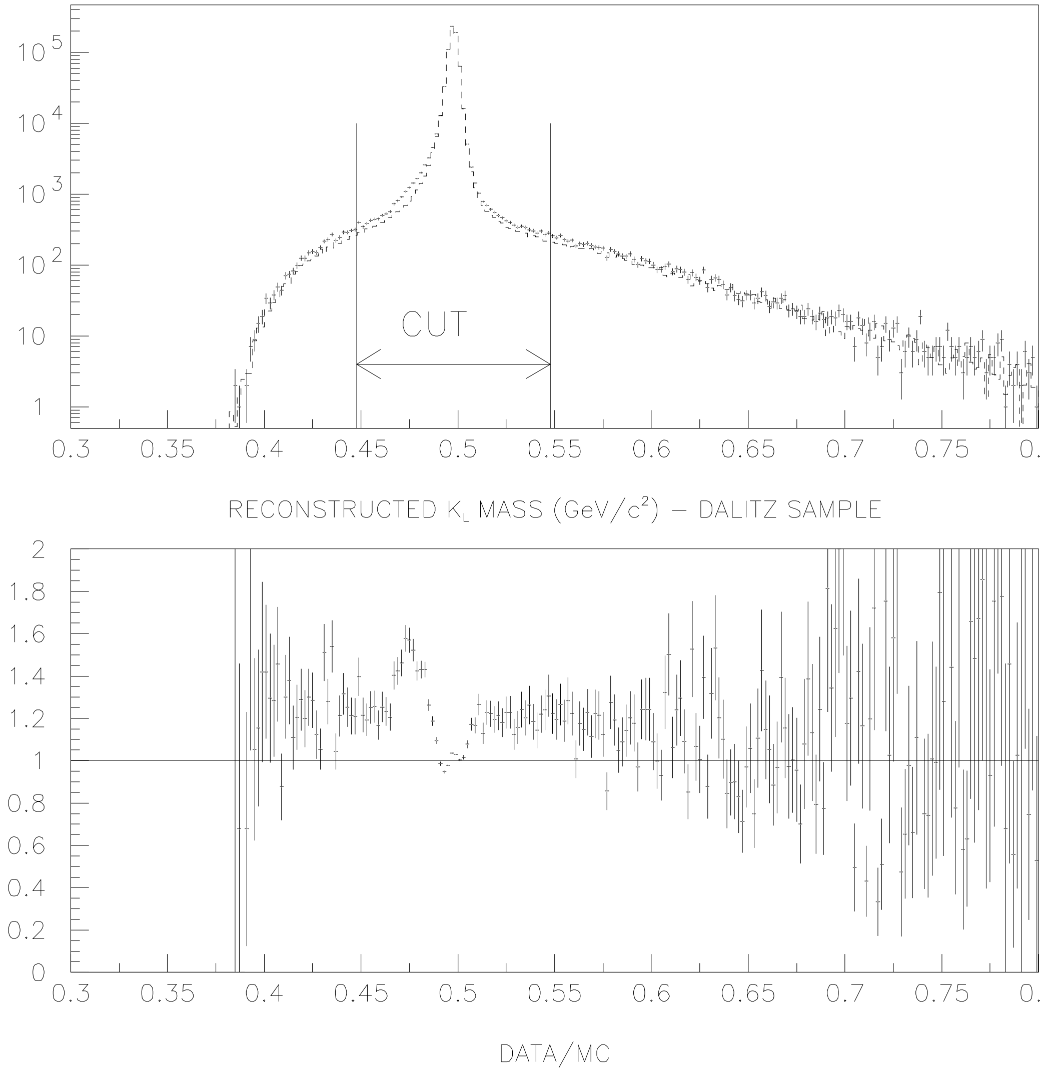

Figure 55. Reconstructed K_L mass for $K_L \to 3\pi^0$, $\pi^0 \to e^+e^-\gamma$ decays which pass all normalization cuts. Data are shown as crosses with errors; MC is dashed histogram. Lower plot is data/MC ratio.

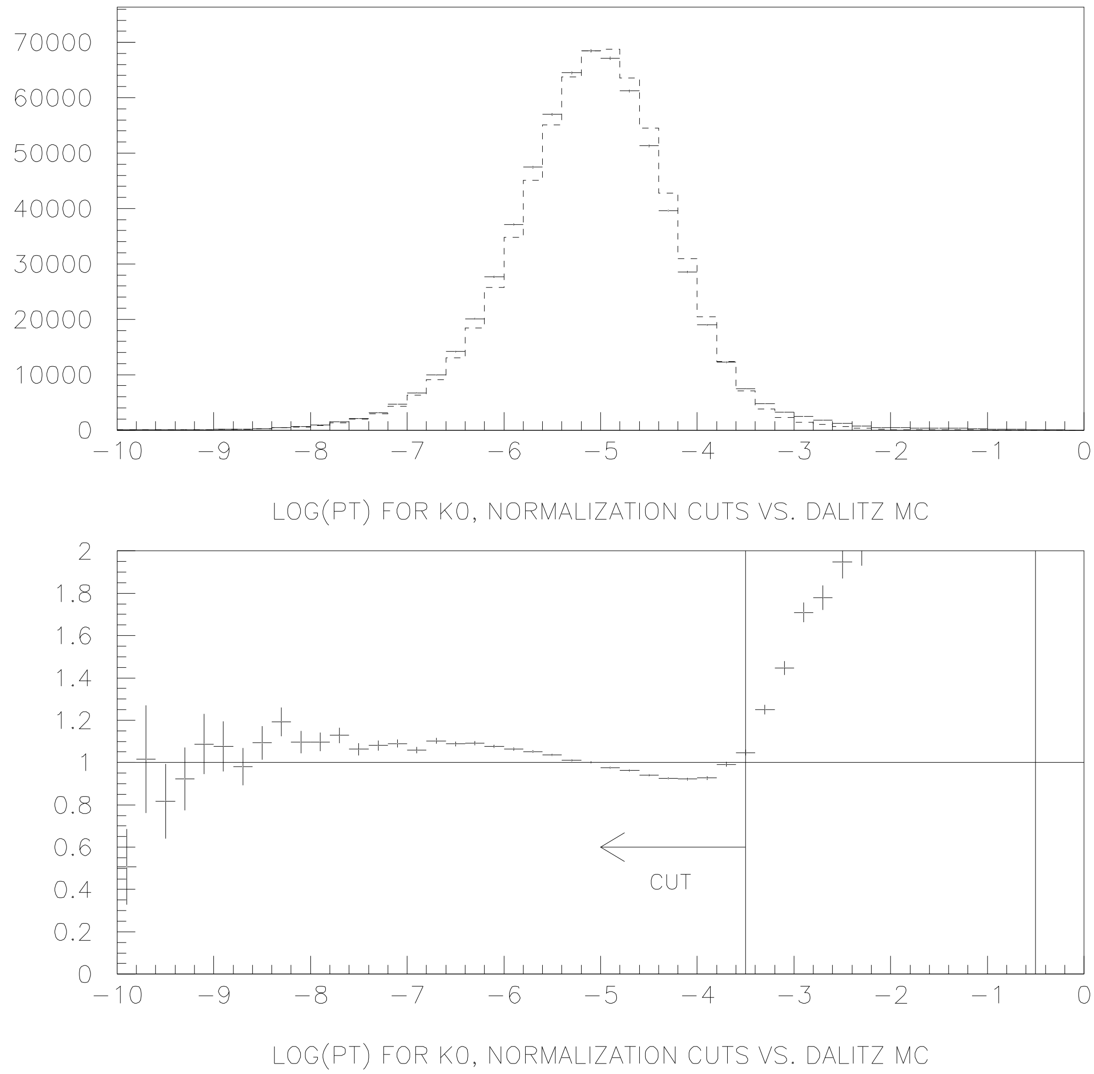

Figure 56. Natural logarithm of transverse momentum (GeV/c) of normalization Dalitz decays. The upper plot shows the data and Dalitz MC spectra (crosses are data; dashed histogram is MC); the lower plot shows the Data/MC ratio. The location of the analysis cut is shown.

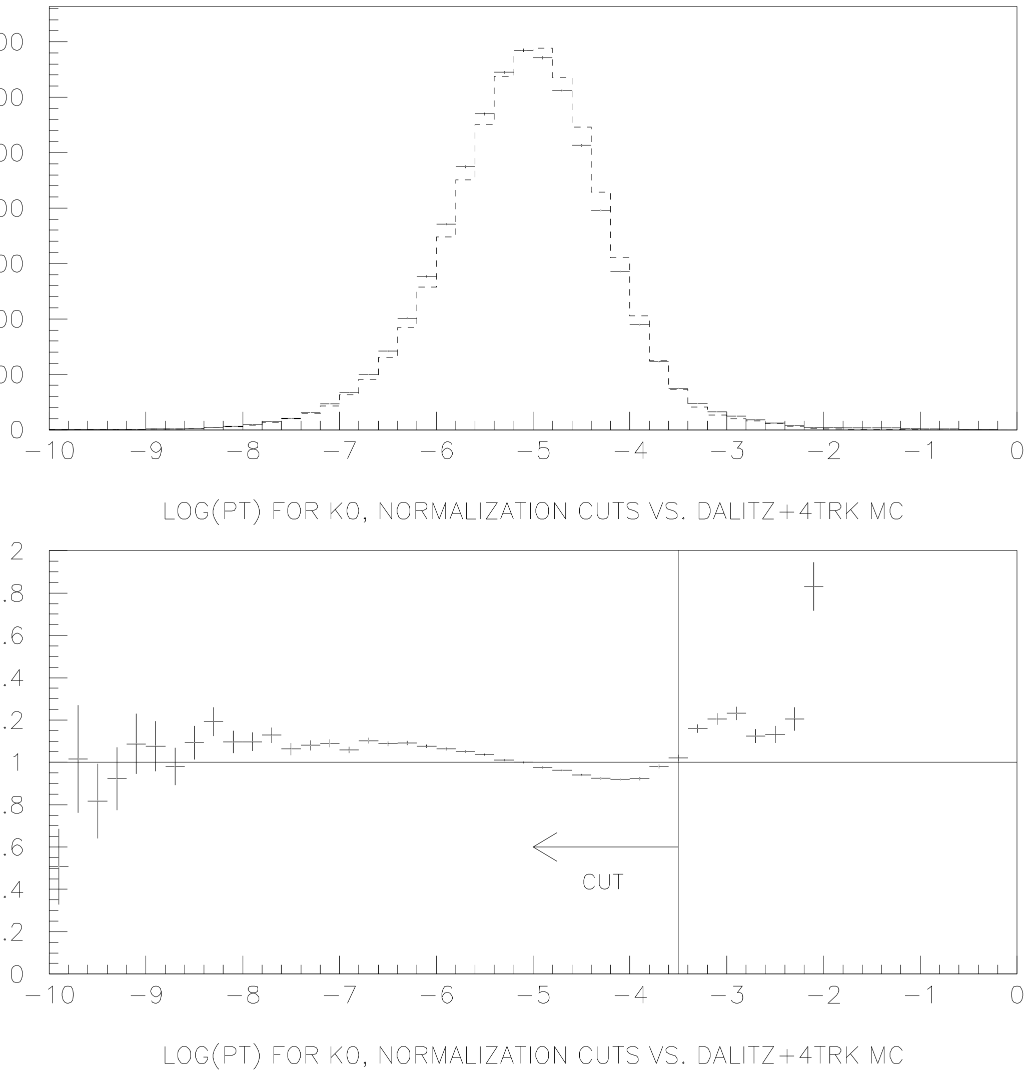

Figure 57. Natural logarithm of transverse momentum (GeV/c) of normalization Dalitz decays, shown with MC simulation of Dalitz decays *and* all four-track background sources summed. Data are crosses; dashed histogram is Monte Carlo. The agreement at the high end of the spectrum is significantly better than in Fig. 56, where the four-track backgrounds are not accounted for in the MC. (After the p_T cut, only ~500 background events are expected to remain in the Dalitz sample, a contamination well below one part in 10^3.)

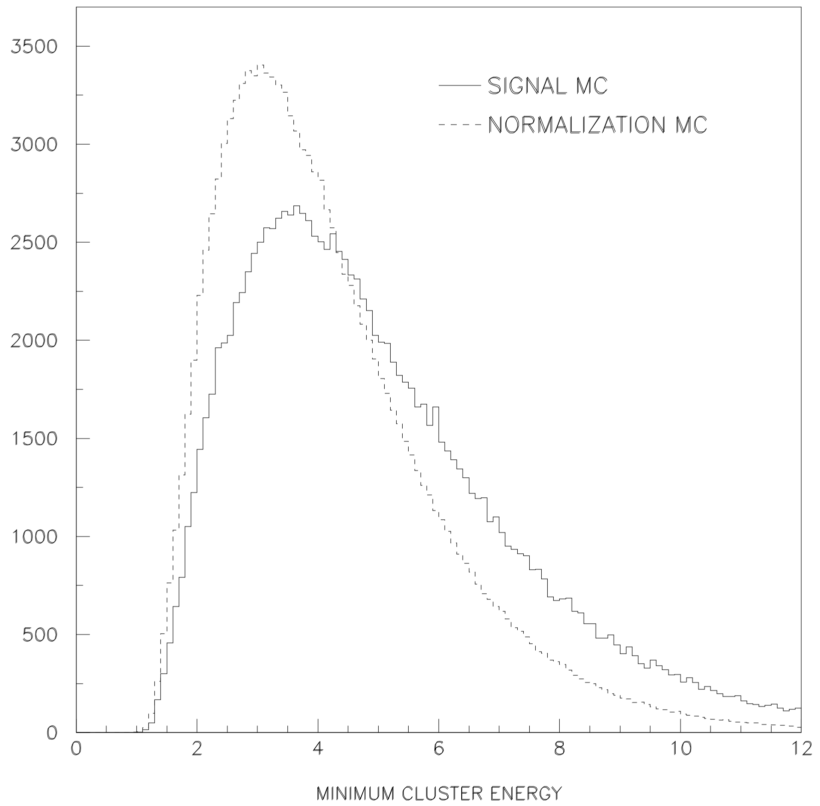

Figure 58. Minimum cluster energy (GeV) for MC $\pi^0 \to e^+e^-$ (solid histogram) and MC $\pi^0 \to e^+e^-\gamma$ normalization events (dashed histogram).

8.5.5 Cluster threshold studies

The main topological difference between the signal and normalization mode was the presence of an extra photon in the normalization mode. This photon was typically soft in the high-$m_{e^+e^-}$ events which were used for the normalization, so the minimum cluster energy in an event tended to be somewhat lower in the normalization than in the signal mode (Fig. 58). Therefore, the branching ratio measurement was sensitive to the cluster energy threshold. The minimum reconstructible cluster energy was set by the thresholds on the HCC, and these thresholds were not perfectly modeled in the MC. The cut on the minimum cluster energy at 1.5 GeV was required to remove events in the range where

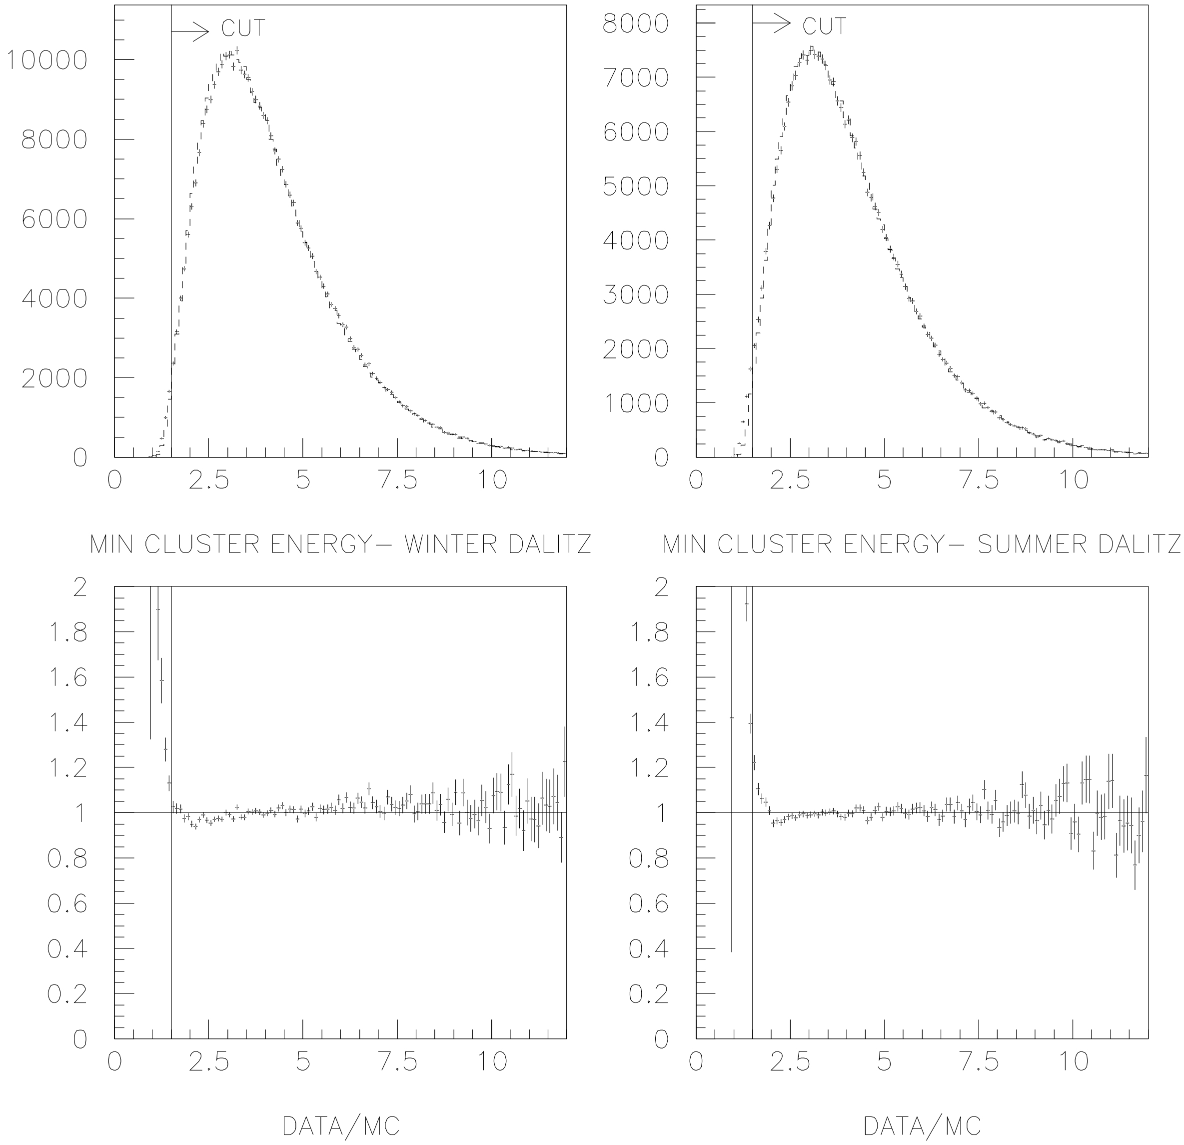

Figure 59. Minimum cluster energy (GeV) for reconstructed Dalitz decays. Left plots are winter data; right plots are summer data.

the efficiency was very poorly modeled (see Fig. 59). This cut cost 0.7% of Dalitz mode MC and 0.3% of signal mode MC.

After this cut, the remaining bias was studied by imposing altered reconstruction efficiency in 0.5 GeV bins in the MC up to 3.0 GeV, where the HCC efficiency was near unity. The efficiency change was applied to both signal and normalization, and a “bias

Table 17. $\pi^0 \to e^+e^-$ biases from bins in the minimum cluster energy distribution.

	1.5 – 2.0 GeV	2.0 – 2.5 GeV	2.5 – 3.0 GeV
Bias per % efficiency change	+0.023%	+0.036%	+0.043%
Largest observed discrepancy	+10%	−5%	−4%
Expected bias	+0.23%	−0.18%	−0.17%

per % efficiency change" was calculated for each bin. This was the fractional change in the number of $\pi^0 \to e^+e^-$ MC events minus the fractional change in the number of $\pi^0 \to e^+e^-\gamma$ MC events when the efficiency for that bin was increased by 1% of itself. The results are shown in Table 17. The bias from any bin was significantly less than 0.5%, and the biases from different bins tended cancel somewhat. This was reasonable, as the thresholds were not uniformly high or low in the data, but rather had a larger spread than in the MC. The efficiency was thus higher than expected below about 2 GeV and lower than expected between 2 and 3 GeV. Taking into account this cancellation, the expected net bias was therefore much smaller than 0.5%, and therefore negligible for this measurement.

8.5.6 Chamber activity cut

The chamber activity cut was consistently 2.5% less efficient in normalization data than in Dalitz MC (Table 18). A discrepancy was somewhat expected, as certain effects which could cause extra chamber hits were not simulated (such as δ-rays which crossed chamber cell boundaries). As expected, the MC predicted nearly identical efficiency in signal and normalization. The same cut was made in signal and normalization, and was effectively a uniform prescale on both modes.

Although there did not appear to be any mechanism to cause a significant variation of the cut efficiency with $m_{e^+e^-}$, studying this was an important cross-check. Three-quarters of the normalization sample had $m_{e^+e^-} < 90$ MeV/c^2, while the signal was clustered at $m_{e^+e^-} \approx 135$ MeV/c^2. The efficiency of the cut for data and Dalitz MC

Table 18. Fraction of events passing DC2 activity cut in MC and data. Errors are statistical only.

	Efficiency of DC2 activity cut	
Sample	Winter run	Summer run
$\pi^0 \to e^+e^-$ MC	0.9552 ± 0.0010	0.9377 ± 0.0010
Dalitz MC only	0.9530 ± 0.0003	0.9362 ± 0.0003
Dalitz + four-track MC	0.9495 ± 0.0003	0.9326 ± 0.0003
Normalization data	0.9285 ± 0.0004	0.9112 ± 0.0004

appears in Fig. 60. Although it appeared that the cut was less efficient in the high-$m_{e^+e^-}$ region of the Dalitz data sample, this was actually the effect of real four-track backgrounds in the Dalitz sample before the cut. If the four-track modes were included in the MC, the efficiency profile matched the data well (Fig. 61). The chamber activity cut was not expected to contribute any significant systematic error in the branching ratio.

8.5.7 The Dalitz branching ratio and π^0 form factor

Normalizing the $\pi^0 \to e^+e^-$ rate to $\pi^0 \to e^+e^-\gamma$ introduced systematic errors associated with the branching ratio and form factor of the Dalitz decay. The $\pi^0 \to e^+e^-\gamma$ branching ratio is $(1.198 \pm 0.032) \times 10^{-2}$ [33]; this has a relative error of 2.7%. This was the largest systematic uncertainty in the $\pi^0 \to e^+e^-$ measurement.

Another source of error was possible misunderstanding of the π^0 form factor in the Dalitz decay. Because the normalization mode was the high-x tail of the Dalitz distribution, an incorrect form factor used in the Monte Carlo would have caused us to extrapolate incorrectly from the high tail to the entire $\pi^0 \to e^+e^-\gamma$ branching ratio. We assumed a form factor with the functional form:

$$F(xm_{\pi^0}^2, 0; m_{\pi^0}^2) = 1 + ax$$

For the slope, we used the measurement $a = 0.033 \pm 0.003$ by the CELLO collaboration

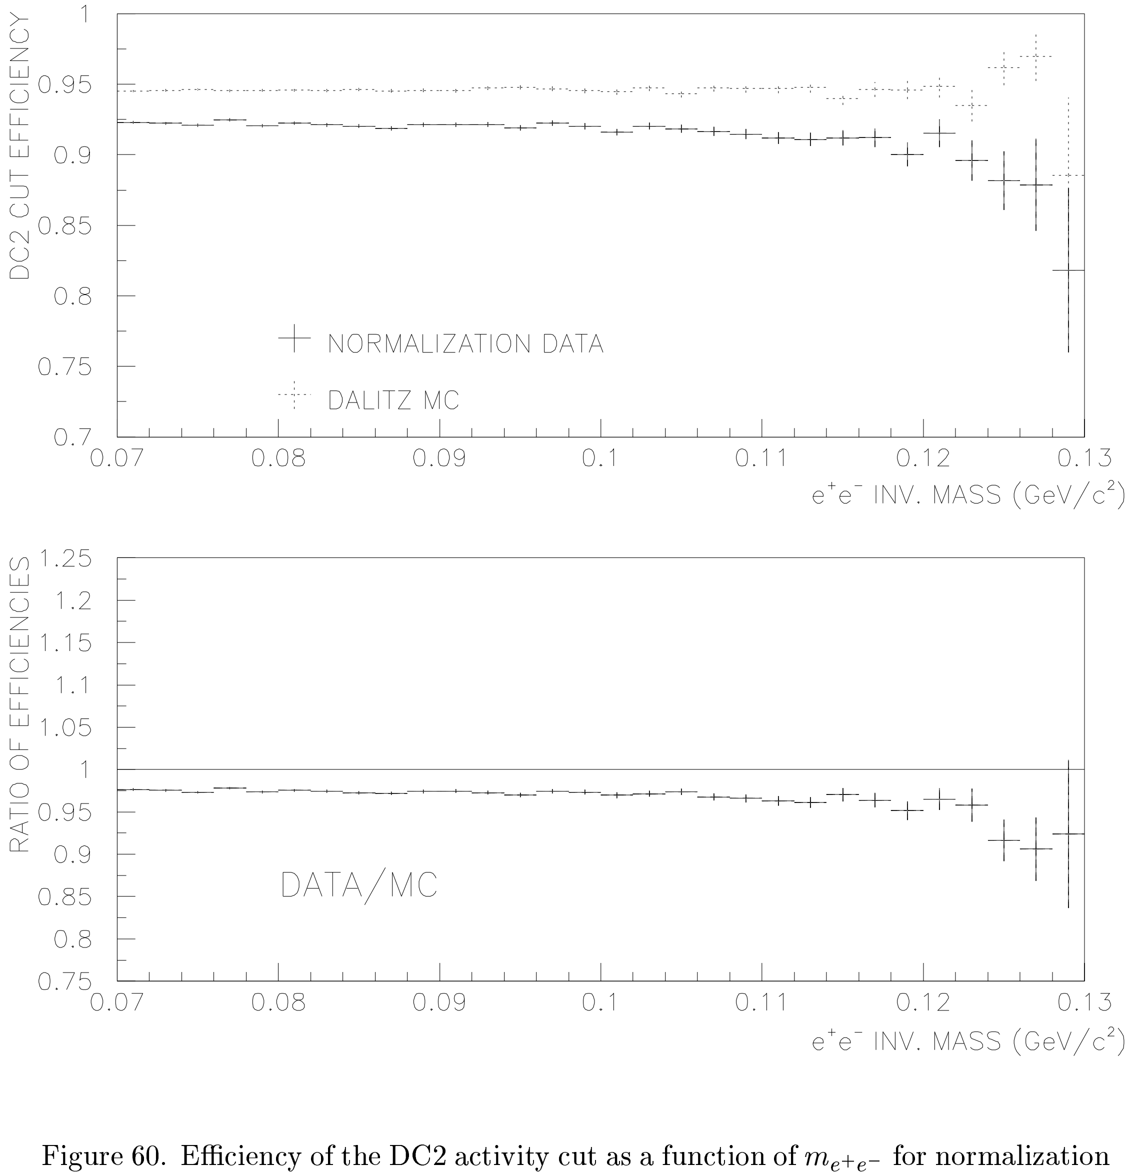

Figure 60. Efficiency of the DC2 activity cut as a function of $m_{e^+e^-}$ for normalization data and $\pi^0 \to e^+e^-\gamma$ MC.

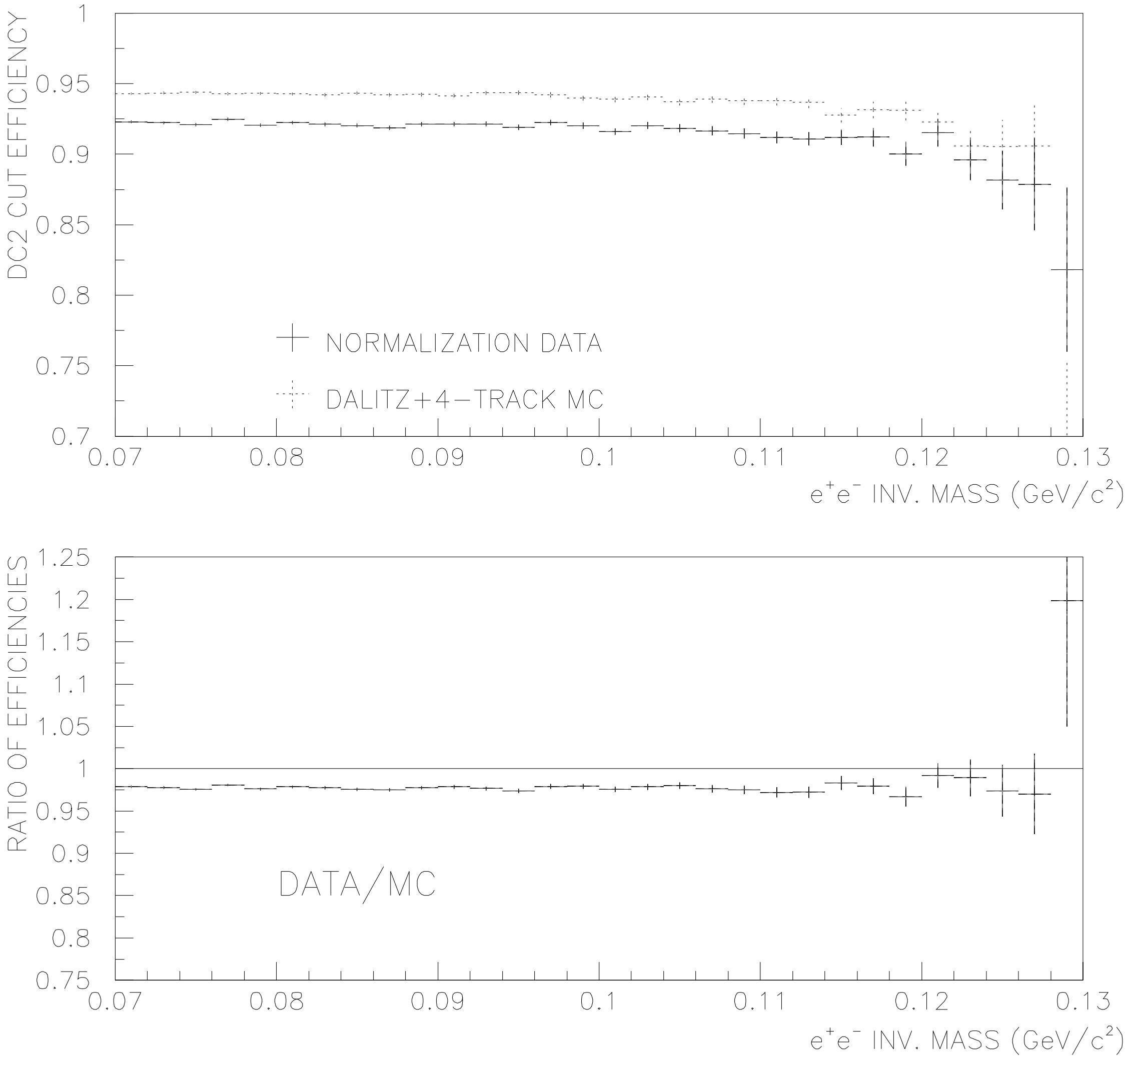

Figure 61. Efficiency of the DC2 activity cut as a function of $m_{e^+e^-}$ for normlization data and $\pi^0 \to e^+e^-\gamma$ MC with four-track background MC added.

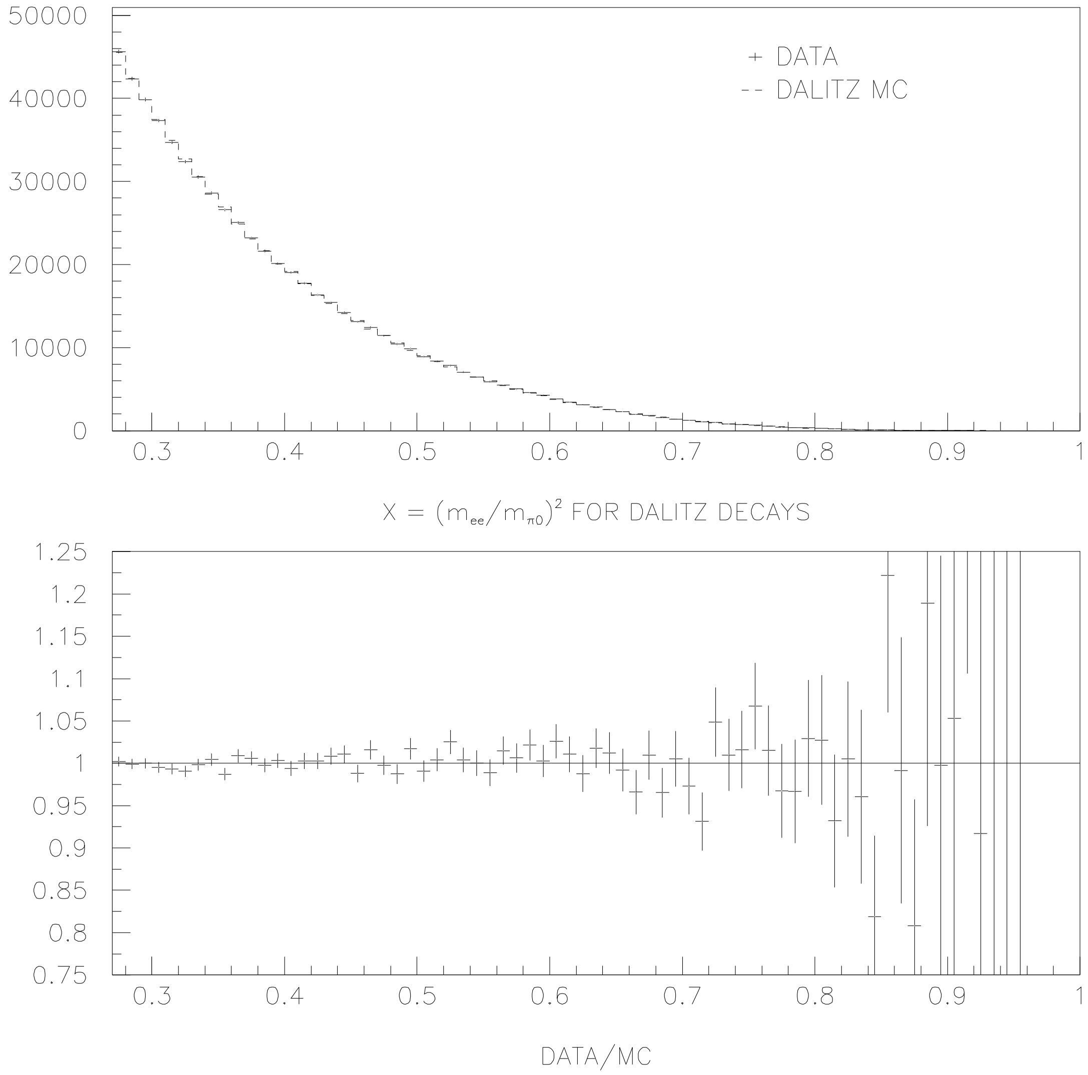

Figure 62. x distribution for data and MC Dalitz decays.

[10]. This measurement was performed in the reaction $e^+e^- \to \pi^0 e^+e^-$, in a region of spacelike momentum transfer, and the slope was extrapolated to the kinematic region of the Dalitz decay assuming vector meson dominance. Direct measurements from the Dalitz decay itself are consistent, but of much lower precision; errors are > 0.03 [11] [12]. We placed a limit on any disagreement in the form factor by looking for a data/MC disagreement in the x distribution of the observed Dalitz decays. A discrepancy δa in the form factor translated (for small a) into the following multiplicative factor in the x distribution:

$$\frac{d\Gamma}{dx} = (1 + 2(\delta a)x)\left[\frac{d\Gamma}{dx}\right]_{\text{expected}}$$

This would manifest itself as a slope in x of $2\delta a$ in the data/MC ratio. A linear fit in x of this ratio (Fig. 62) over the range $0.27 < x < 0.90$[2] gave a slope of 0.0144 ± 0.0146, or $\delta a = 0.0072 \pm 0.0073$. The fit had a χ^2 of 48.0, with 60 degrees of freedom. Assuming no slope ($\delta a = 0$) gave $\chi^2/n = 49.0/61$, which was not significantly different.

Although there was no evidence for a discrepancy in the form factor, our sensitivity was only 0.0073, and we took that as our uncertainty in the form factor. Taking a form factor change δa of this size and weighting our events by $(1 + \delta a x)^2$, we found the integrated Dalitz flux (and therefore the measured $\pi^0 \to e^+e^-$ branching ratio) changed by 0.5%; this was taken as the systematic error associated with the π^0 form factor.

8.6 Branching ratio calculation

In the signal sample, there were 275 candidate events in the data, and MC background estimates totaled 21.4 ± 6.2 events. Considering only statistical errors (excluding MC statistics, which were considered a systematic error), the background-subtracted signal was 253.6 ± 16.6 events. The Dalitz normalization sample was 650264 events, with negligible ($< 0.08\%$) background from four-track and radiative $\pi^0 \to e^+e^-$ events.

In order to extract a branching ratio for $\pi^0 \to e^+e^-$, we used the Monte Carlo samples generated for the winter and summer run separately, and weighted the samples by the

[2]The fit was cut off at $x = 0.9$ because the statistics in this region were very low, and there was a contribution (~ 10 events) from radiative $\pi^0 \to e^+e^-$ in the Dalitz sample at high x.

observed normalization sample sizes. The method was to weight the MC samples such that one normalization data event was equivalent to one effective (weighted) MC event. We could then simply divide the observed number of events in the data by the *weighted* number of generated events to obtain an acceptance. For either the summer or winter run alone, this would have been an unnecessary complication. However, the samples for the summer and winter runs were generated before the ratio of sample sizes was known, so we needed to go through this reweighting step to produce acceptance numbers for the data sets when combined.

For the winter run, 7.047×10^7 $\pi^0 \to e^+e^-\gamma$ MC events were generated (see Table 9), 709669 of which were accepted after final cuts. 373777 events were observed in data; the ratio between data and MC was 0.5267. Multiplying the number of generated events by this factor, we get the weighted number of generated events, or 3.712×10^7. (This corresponds to the number of $\pi^0 \to e^+e^-\gamma$ decays with $m_{e^+e^-} > 0.065$ GeV/c^2 in the winter run after correcting for acceptance.) In the summer run, 6.166×10^7 MC events were generated, of which 653825 passed the final cuts. 276487 were observed in the data, yielding an MC weight of 0.4229 and a total number of 2.608×10^7 $\pi^0 \to e^+e^-\gamma, m_{e^+e^-} > 0.065$ GeV/c^2 decays in the summer run.[3]

After weighting the MC events by the data/MC ratio for the winter and summer runs, the MC samples for the two runs could simply be added: the effective number of generated Dalitz decays for the two runs combined was thus 6.319×10^7. The 650264 observed data events were equivalent to the same number of weighted MC events; the combined acceptance was the ratio of those two numbers:

$$A(\text{normalization}) = \frac{N(\text{Accepted weighted MC})}{N(\text{Generated weighted MC})} = \frac{650264}{6.319 \times 10^7} = 0.01029$$

The $\pi^0 \to e^+e^-$ MC samples were generated with the same ratio of winter to summer

[3]From these numbers and the known $K_L \to 3\pi^0$ and $\pi^0 \to e^+e^-\gamma$ branching ratios, we could calculate the total "flux," or number of K_L decays in the decay volume in the generated momentum range when the detector was live. This number was 1.57×10^{11} for the winter run and 1.10×10^{11} for the summer. However, the error on those numbers may be potentially large, as many systematic effects may be present which cancel in the branching ratio we measure. The "flux" was only consistent to $\sim 15\%$ when measured using decays to different types of particles.

events as the normalization MC, so the same MC sample weights could be used. After weighting, the effective numbers of generated MC $\pi^0 \to e^+e^-$ events from the winter and summer runs were 1.580×10^6 and 1.110×10^6 respectively, or a total of 2.690×10^6. Of these, 34959 winter and 25646 summer weighted MC events passed the final cuts, or a total of 60605. The combined acceptance for $\pi^0 \to e^+e^-$ was therefore

$$A(\pi^0 \to e^+e^-,\ \text{all } x) = \frac{N(\text{Accepted weighted MC})}{N(\text{Generated weighted MC})} = \frac{60605}{(2.690 \times 10^6)} = 0.02253$$

The acceptance was calculated from a total of 1.3×10^5 accepted (unweighted) MC events, so the error from signal MC statistics was a negligible 0.3%. Using the result from Sec. 1.5.2 that

$$\frac{\text{BR}(\pi^0 \to e^+e^-, x > 0.95)}{\text{BR}(\pi^0 \to e^+e^-,\ \text{all } x)} = 0.895$$

we found

$$A(\pi^0 \to e^+e^-, x > 0.95) = \frac{1}{0.895} A(\pi^0 \to e^+e^-,\ \text{all } x) = 0.02517.$$

Knowing the combined acceptances, we could then calculate the relative branching ratio:

$$\begin{aligned}
\frac{\text{BR}(\pi^0 \to e^+e^-, x > 0.95)}{\text{BR}(\pi^0 \to e^+e^-\gamma, m_{e^+e^-} > 0.065\ \text{GeV}/c^2)} &= \frac{N_{\text{signal}}/A(\text{signal})}{N_{\text{normalization}}/A(\text{normalization})} \\
&= \frac{(253.6 \pm 16.6)/0.02517}{650264/0.01029} \\
&= (1.594 \pm 0.104(\text{stat})) \times 10^{-4}.
\end{aligned}$$

To calculate the systematic error on this ratio, we added in quadrature the acceptance-related and background-related errors discussed in Sec. 8.5 and Sec. 8.3.3, obtaining a total systematic error of 2.9%. The sources of systematic errors are summarized in Table 19. This gave the result:

$$\frac{\text{BR}(\pi^0 \to e^+e^-, x > 0.95)}{\text{BR}(\pi^0 \to e^+e^-\gamma, m_{e^+e^-} > 0.065\ \text{GeV}/c^2)} = (1.594 \pm 0.104(\text{stat}) \pm 0.038(\text{syst})) \times 10^{-4}.$$

The high-$m_{e^+e^-}$ events used for the normalization mode comprised 3.191% of the total Dalitz decay rate (see Sec. 5.2). This number was subject to the fractional error

Table 19. Systematic errors in the $\pi^0 \to e^+e^-$ branching ratio.

Source of error	Fractional error on branching ratio		
Chamber inefficiencies in beams	$< 0.1\%$ (ignored)		
Pairing χ^2 cut efficiency	1.0%		
e^+e^- mass resolution	1.2%		
p_T resolution	$< 0.2\%$ (ignored)		
Cluster energy threshold	$\ll 0.5\%$ (ignored)		
Total acceptance-related errors		1.6%	
Background uncertainty (6.2 events)		2.4%	
Total from acceptance and background			2.9%
$\pi^0 \to e^+e^-\gamma$ branching ratio	2.7%		
π^0 form factor	0.5%		
Total error from π^0 parameters			2.7%
TOTAL SYSTEMATIC ERROR			4.0%

of 0.5% from the π^0 form factor. From this, we calculated the ratio between the signal and the entire Dalitz decay rates (incorporating this additional systematic error):

$$\frac{\mathrm{BR}(\pi^0 \to e^+e^-, x > 0.95)}{\mathrm{BR}(\pi^0 \to e^+e^-\gamma)} = (5.085 \pm 0.333(\mathrm{stat}) \pm 0.150(\mathrm{syst})) \times 10^{-6}$$

Finally, to obtain the absolute branching ratio of $\pi^0 \to e^+e^-$, we needed to multiply this number by the $\pi^0 \to e^+e^-\gamma$ branching ratio. The $\pi^0 \to e^+e^-\gamma$ branching ratio is only known to 2.7% of itself, and this had to be included as an additional systematic error, bringing the total systematic error to 4.0%. The final result for the $\pi^0 \to e^+e^-$ branching ratio is:

$$\mathrm{BR}(\pi^0 \to e^+e^-, x > 0.95) = (6.09 \pm 0.40(\mathrm{stat}) \pm 0.24(\mathrm{syst})) \times 10^{-8}.$$

CHAPTER 9

CONCLUSIONS

We have measured the branching ratio of the decay $\pi^0 \to e^+e^-$ to a precision of 7.7%. From the previous chapter,

$$\mathrm{BR}\left(\pi^0 \to e^+e^-, \left(\frac{m_{e^+e^-}}{m_{\pi^0}}\right)^2 > 0.95\right) = (6.09 \pm 0.40(\mathrm{stat}) \pm 0.24(\mathrm{syst})) \times 10^{-8}.$$

This represents a major improvement over all previous results, and the first significant observation of excess rate above the unitarity bound. Fig. 63 shows this result along with previous measurements. This result confirms the lower values of the branching ratio reported by BNL E851 and FNAL E799-I in 1993 (Refs. [32] and [26]), and does not agree with the high average value of the results reported before 1985 (Refs. [28] and [29]).

9.1 The excess above unitarity

To calculate the excess rate above the unitarity bound, we divided out the radiative correction described in Sec. 1.5.2, extrapolating from our measurement of the $x > 0.95$ branching ratio to the lowest-order rate for the $\pi^0 \to \gamma^*\gamma^* \to e^+e^-$ diagram. For this rate, we found

$$\frac{\Gamma_{e^+e^-}^{\text{lowest order}}}{\Gamma_{\text{all}}} = (7.04 \pm 0.46(\mathrm{stat}) \pm 0.28(\mathrm{syst})) \times 10^{-8}$$

which is over four standard deviations above the unitarity bound

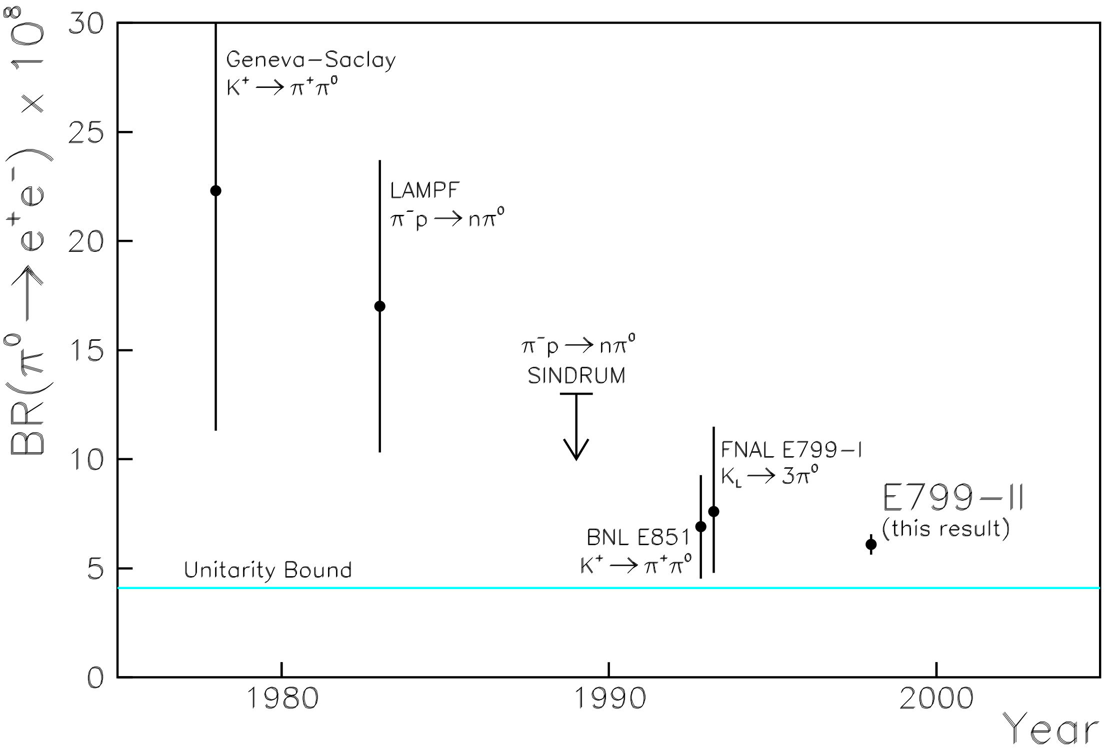

Figure 63. Measurements of the branching ratio of $\pi^0 \rightarrow e^+e^-$, including this result.

$$\frac{\Gamma_{e^+e^-}^{\text{lowest order}}}{\Gamma_{\text{all}}} \geq 4.69 \times 10^{-8}$$

obtained from Eq. 1.3 and the PDG value BR($\pi^0 \rightarrow \gamma\gamma$) = 0.988 [33]. Expressing our branching ratio as a fraction of the unitarity bound, with the statistical and systematic errors combined in quadrature, we can write:

$$\frac{\Gamma(\pi^0 \rightarrow e^+e^-)}{\Gamma^{\text{unitary}}(\pi^0 \rightarrow e^+e^-)} = 1.501 \pm 0.115$$

or, in the notation of Eq. 1.1, where $\text{Im}(R) = -17.5$ is an exact and model-independent result:

$$|\text{Re}(R)| = 12.40 \pm 1.42$$

and

$$\left|\frac{\text{Re}(R)}{\text{Im}(R)}\right| = 0.708 \pm 0.081.$$

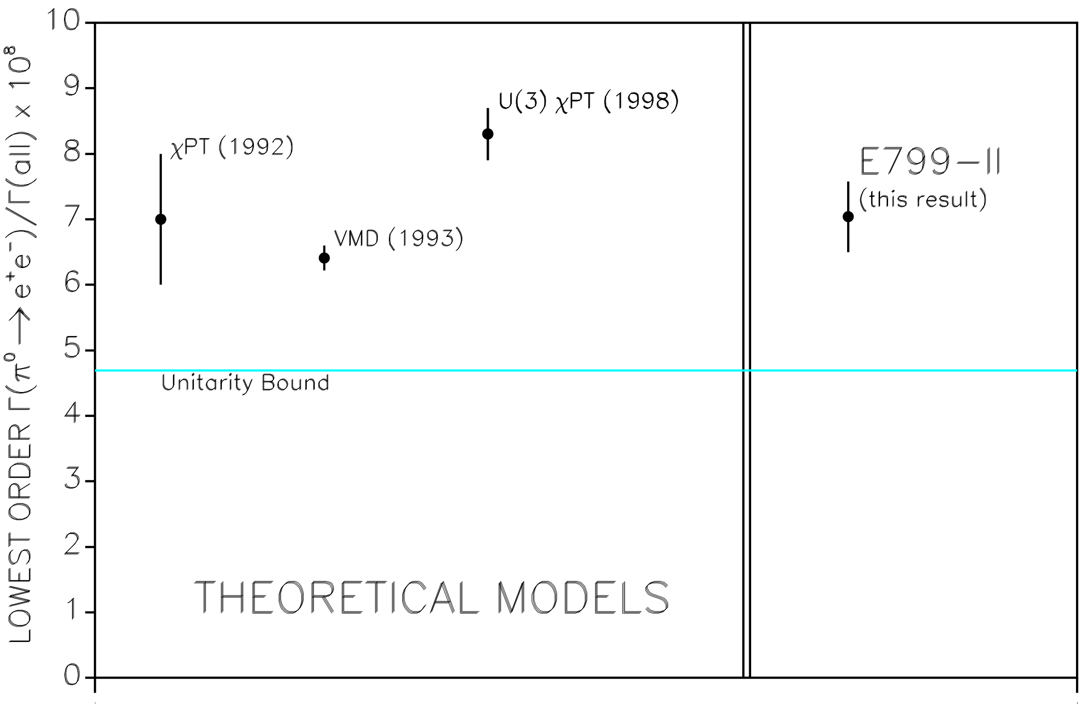

Figure 64. Lowest order $\pi^0 \to e^+e^-$ branching ratio for various models, compared with this result. Theory points are, from left to right, from Refs. [13], [9], and [14].

9.2 Comparison to models

Given the precision of this result, we could make some statements about its consistency with recent theoretical models. Our result for the lowest order branching ratio is shown in Fig. 64, along with recent models. Comparing our value for the lowest order rate for $\pi^0 \to e^+e^-$ with that predicted by the models, we found that the χPT model in Ref. [13] was clearly in the best agreement with our data. This was, of course, somewhat unsurprising because of the large errors in this model. Our measurement agrees with the pure VMD model in Ref. [9] at the 27% confidence level, and with the $U(3)_L \otimes U(3)_R$

χPT model in Ref. [14] at the 6% confidence level, treating the theoretical uncertainties as 1σ statistical errors[1].

Clearly, our result is somewhat less favorable to the last model, although it is only "excluded" at the 2σ level.

9.3 Remarks and future prospects

This measurement represents a major step in our understanding of the decay $\pi^0 \to e^+e^-$. The decay rate has now been firmly established to be above the unitarity bound. A significant measurement of the excess above unitarity has been measured for the first time, and was found to be fully consistent with Standard Model expectations. Indeed, this result happens to fall in between predictions from different models (and is perhaps the only possible result which could be consistent with the predictions in both Refs. [9] and [14]).

A Fermilab fixed-target run is scheduled for 1999, and E799-II could then produce two to four times the statistics reported here. If this high yield materializes, the measurement of $\pi^0 \to e^+e^-$ using the current technique will become limited by the uncertainty in the Dalitz decay branching ratio. It is ironic that we can contemplate the prospect of a rare decay measurement being limited by our knowledge of a normalization mode which is more common by over five orders of magnitude.

This uncertainty could be avoided by either improving the Dalitz decay measurement itself, or by developing a technique for normalizing $\pi^0 \to e^+e^-$ directly to $\pi^0 \to \gamma\gamma$. Either prospect would entail normalization of a charged decay to a neutral decay: this would require very detailed understanding of the absolute track-finding efficiency, as tracking effects would not cancel in the ratio. The problems with drift chamber performance in the 1997 data set made this a daunting prospect, although modifications

[1] The χPT models have errors derived from the measured $\eta \to \mu^+\mu^-$ branching ratio, so the gaussian error assumption is reasonable. The error in the VMD model comes from the difference between the ρ^0 and ω masses, and is thus not statistical in nature.

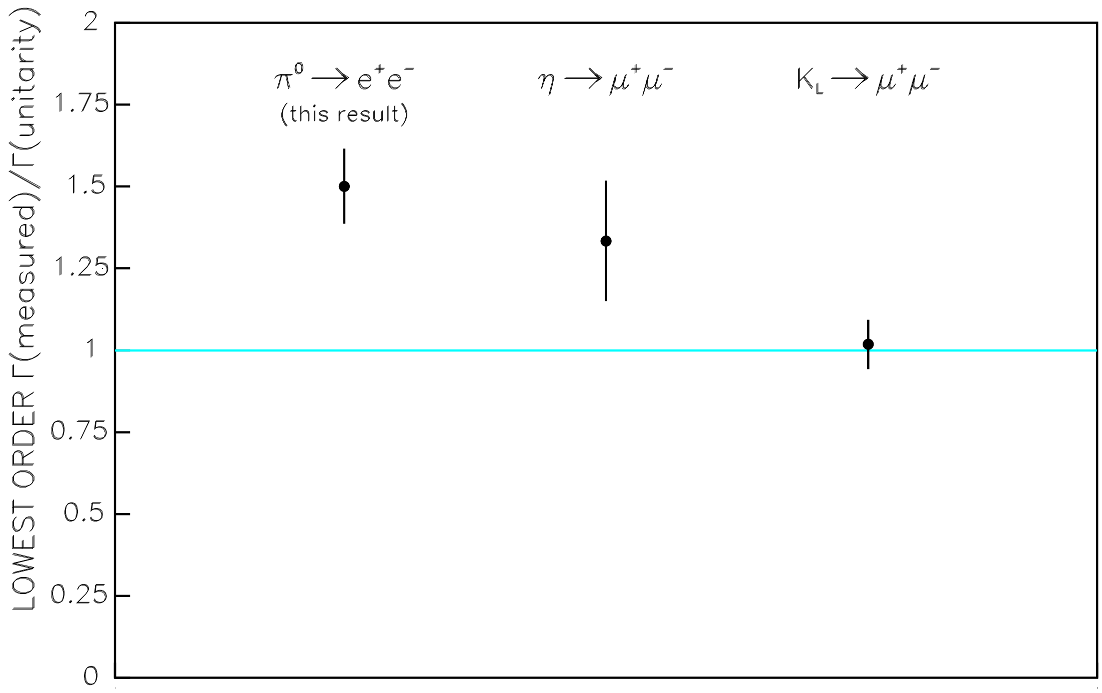

Figure 65. Measured excess above unitarity bound for pseudoscalar meson decays to lepton pairs. $\eta \to \mu^+\mu^-$ and $K_L \to \mu^+\mu^-$ results use 1998 PDG [33] values for branching ratios.

currently being made to the chamber system should alleviate the problems for the 1999 run.

Beyond the 1999 run, a consistent picture of the decays of pseudoscalar mesons into lepton pairs will demand better measurements of the $K_L \to \mu^+\mu^-$ and $\eta \to \mu^+\mu^-$ branching ratios. BNL E871 will soon produce a significantly improved measurement of $K_L \to \mu^+\mu^-$ [70].

As can be seen in Fig. 65, $\eta \to \mu^+\mu^-$ has now replaced $\pi^0 \to e^+e^-$ as the least precise result among the three experimentally accessible pseudoscalar-to-dilepton decay modes. It is hoped that this result on $\pi^0 \to e^+e^-$ may be able to constrain predictions for it, and that experiments will be able to improve on the existing measurements of $\eta \to \mu^+\mu^-$.

$\pi^0 \to e^+e^-$ is now the only one of these three decays with a well-measured excess above the unitarity bound. The good agreement between the measurement and the models is encouraging; further theoretical and experimental work on $\pi^0 \to e^+e^-$ and, in particular, on $\eta \to \mu^+\mu^-$ could significantly improve our understanding of these decays.

REFERENCES

[1] S. Drell, *Nuovo Cim.* **XI**, 693 (1959)

[2] L. Bergström *et al., Phys. Lett.* **126B**, 117 (1983)

[3] N. Kroll and W. Wada, *Phys. Rev.* **98**, 1355 (1955)

[4] T. Miyazaki and E. Takasugi, *Phys. Rev.* **D8**, 2051 (1973)

[5] A. Pich and J. Bernabéu, Z Phys. **C22**, 197 (1984)

[6] Ll. Ametller *et al., Nucl. Phys.* **B228**, 301 (1983)

[7] M. Pratap and J. Smith, *Phys. Rev.* **D5**, 2020 (1972)

[8] K. S. McFarland, PhD. Thesis, The University of Chicago (1994)

[9] Ll. Ametller, A. Bramon, and E. Massó *Phys. Rev.* **D48**, 3388 (1993)

[10] H. Behrend *et al., Z. Phys.* **C49**, 401 (1991)

[11] F. Farzanpay *et al., Phys. Lett.* **B278**, 413 (1992)

[12] R. Meijer Drees *et al., Phys. Rev.* **D45**, 1439 (1992)

[13] M. J. Savage, M. Luke, and M. B. Wise *Phys. Lett.* **B291**, 481 (1992)

[14] D. Gomez Dumm and A. Pich, *Phys. Rev. Lett.* **80**, 4633 (1998)

[15] E. Massó, *Phys. Lett.* **B181**, 388 (1986)

[16] J. Pati and A. Salam, *Phys. Rev.* **D11**, 1137 (1975)

[17] J. Rosner, private communication (1998)

[18] R. Abegg *et al., Phys. Rev.* **D50**, 92 (1994)

[19] A. P. Heinson *et al., Phys. Rev.* **D51**, 985 (1995)

[20] D. Ambrose *et al., Phys. Rev. Lett.* **81**, 4309 (1998)

[21] A. R. Clark *et al., Phys. Rev. Lett.* **26**, 1667 (1971)

[22] G. Valencia, *Nucl. Phys.* **B517**, 339 (1998)

[23] R. E. Schrock and M. B. Voloshin, *Phys. Lett.* **87B**, 375 (1979)

[24] T. Inami and C. S. Lim, *Prog. Theo. Phys.* **65**, 297 (1981)

[25] L. Bergström, *Z. Phys.* **C20**, 135 (1983)

[26] K. S. McFarland *et al., Phys. Rev. Lett.* **71**, 31 (1993)

[27] G. Triantaphyllou, *Mod. Phys. Lett.* **A8**, 1691 (1993)

[28] J. Fischer *et al., Phys. Lett.* **73B**, 364 (1978)

[29] J. S. Frank *et al., Phys. Rev.* **D28**, 423 (1983)

[30] A. G. Zephat *et al., J. Phys.* **G13**, 1375 (1987)

[31] C. Niebuhr *et al., Phys. Rev.* **D40**, 2976 (1989)

[32] A. Deshpande *et al., Phys. Rev. Lett.* **71**, 27 (1993)

[33] Particle Data Group, *Eur. Phys. J.* **C3**, 1 (1998)

[34] M. Woods, PhD. Thesis, The University of Chicago (1988)

[35] J. R. Patterson, PhD. Thesis, The University of Chicago (1990)

[36] L. Gibbons, PhD. Thesis, The University of Chicago (1993)

[37] D. Harris, PhD. Thesis, The University of Chicago (1994)

[38] R. Briere, PhD Thesis, The University of Chicago (1995)

[39] B. Schwingenheuer, PhD. Thesis, The University of Chicago (1995)

[40] K. Hanagaki, PhD Thesis, Osaka University (1998)

[41] E. Bartz *et al.,* "KTeV Drift Chamber Electronics Design Report" KTeV Internal Note 190 (1994)

[42] A. Roodman, "The KTeV pure CsI calorimeter," *Proceedings of the VII International Conference on Calorimetry, Tucson, Arizona,* World Scientific (1998)

[43] P. Shanahan, Proceedings of ICHEP, Frascati, Italy (1996)

[44] R. Ben-David, private communication (1998)

[45] Y. B. Hsiung, "Vacuum Ring Veto Geometry and Design" *KTeV Internal Note* 121 (1993)

[46] A. Barker and U. Nauenberg, "Spectrometer Anti Geometry and Design" *KTeV Internal Note* 132 (1993)

[47] "DYC Write-Up V1.0," FNAL Computing Division, document HN134 (1996)

[48] E. D. Zimmerman, "KTeV Trigger File Syntax Rules" *KTeV Internal Note* 309 (1995)

[49] R. S. Kessler, "KTeV Trigger Control Logic" *KTeV Internal Note* 361 (1996)

[50] C. Bown *et al.*, *Nucl. Inst. Meth.* **A 369**, 248 (1996)

[51] A. R. Barker, "Trigger TDC/Hit-Counting Design Report" *KTeV Internal Note* 204 (1994)

[52] P. S. Shawhan, "A Y Track Finder for KTeV" *KTeV Internal note* 257 (1994)

[53] M. Arenton *et al.*, "KTeV TRD e^+e^- Trigger Design Report III" *KTeV Internal Note* 286

[54] J. Kubic, "KTeV FERA Scaler CAMAC Module" *KTeV Internal note* 371 (1996)

[55] J. Adams *et al., Phys. Rev. Lett.* **80**, 4123 (1998)

[56] P. Mikelsons, "The 2E-NCLUS Crunch" *KTeV Internal Note* 500 (1988)

[57] A. Malensek, FNAL Preprint FN-341 and FN-341a (errata) (1981)

[58] S. Somalwar *et al., Phys. Rev. Lett.* **68**, 2580 (1992)

[59] K. O. Mikaelian and J. S. Smith, *Phys. Rev* **D5**, 1763 (1972)

[60] K. O. Mikaelian and J. S. Smith, *Phys. Rev* **D5**, 2891 (1972)

[61] H. Bethe, *Phys. Rev.* **89**, 1256 (1953)

[62] J. D. Jackson, *Classical Electrodynamics* New York: J. Wiley (1975)

[63] W. R. Nelson *et al.*, "The EGS4 code system," SLAC Report 265 (1985)

[64] P. Shawhan, "T3 Tracking Algorithm," *KTeV Internal Note* 328 (1995)

[65] E. D. Zimmerman, "Use of the $K_L \to 3\pi^0$ decay as a tagged photon source to measure material thickness in a neutral kaon beam," EFI preprint 98-43 (to appear *Nucl. Inst. Meth*) (1998)

[66] T. Nakaya, PhD. Thesis, Osaka University (1995)

[67] R. Ford *et al.*, "A Low Cost, 400 Gauss 18$''$ Gap C-Magnet Using Permanent Magnet Technology," FNAL Pub-98/271-E (to appear in *Nucl. Inst. Meth.*) (1998)

[68] G. Jackson, "Fermilab recycler ring technical design report, Rev. 1.2," FNAL-TM-1991 (1996)

[69] D. Jensen, private communication (1997)

[70] D. Ambrose, talk given at *Workshop on Heavy Quarks at Fixed Target (HQ98),* Batavia, Ill. (1998)