RICE UNIVERSITY

Search for the lepton-flavor-number violating decay $K_L \to \pi^0 \mu^\pm e^\mp$ in the full E799II KTeV data set

by

Angela M. Bellavance

A Thesis Submitted
in Partial Fulfillment of the
Requirements for the Degree

DOCTOR OF PHILOSOPHY

Approved, Thesis Committee:

____

Marjorie Corcoran, Chairman
Professor of Physics

____

B. Paul Padley
Assistant Professor of Physics

____

David W. Scott
Noah Harding Professor of Statistics

Houston, Texas

November, 2002

ABSTRACT

Search for the lepton-flavor-number violating decay $K_L \to \pi^0 \mu^\pm e^\mp$ in the full E799II KTeV data set

by

Angela M. Bellavance

Experiments E799 Phase-II (E799II) and E832 of the Kaons at the TeVatron (KTeV) project at the Fermi National Accelerator Laboratory (Fermilab) were designed to study rare and CP-violating kaon decays. The experiments ran first during 1996 and 1997, and then again during 1999. This study focused on the E799II search for the lepton-flavor-violating decay mode $K_L \to \pi^0 \mu^\pm e^\mp$, and on its accompanying background decays, in the full set of 1996-7 and 1999 KTeV data. For the combined data sets, we calculated a new 90% confidence level branching ratio limit BR($K_L \to \pi^0 \mu^\pm e^\mp$) $< 3.37 \times 10^{-10}$.

Acknowledgments

It's done. It's finally freakin' done.

My thanks to my advisor, Dr. Marjorie Corcoran, for supporting me though this whole process and being a good advisor and friend. It's been an adventure, hasn't it? Thank you to my thesis committee for their time and advice. To all my collaborators (listed in Figures 0.1 and 0.2) for the last six and a half years, my thanks and appreciation for all your efforts that made KTeV come together and made this analysis possible. I especially want to thank Leo Bellantoni and Hogan Nguyen for their invaluable help with this analysis. This work was supported in part by the U.S. Department of Energy, The National Science Foundation and The Ministry of Education and Science of Japan.

Dad and Mom, so is being a post-doc considered getting a real job yet? Just kidding. Thanks for your patience and encouragement.

Andrew, thanks for putting up with me while this got finished. You have been a real source of joy, strength and caring.

Angela M. Bellavance
November 8, 2002

The 1997 KTeV Collaboration

E. Cheu, S.A. Taegar
University of Arizona, Tucson, Arizona 85721

K. Arisaka, J. Jennings, W. Slater
University of California at Los Angeles, Los Angeles, California 90095

H.G.E. Kobrak, R.A. Swanson
University of California at San Diego, La Jolla, California 92093

E. Blucher, C. Bown, S. Bright, J. Graham, G. Graham, R. Kessler, E. Monnier†,
V. Prasad, A. Roodman, C. Qiao, B. Quinn, P. Shawhan, N. Solomey, Y.W. Wah,
B. Winstein, R. Winston, E. Zimmerman
The Enrico Fermi Institute, The University of Chicago, Chicago, Illinois 60637

A.R. Barker, J. LaDue, P. Mikelsons, U. Nauenberg, J.-Y. Wu
University of Colorado, Boulder, Colorado 80309

E.C. Swallow
Elmhurst College, Elmhurst, Illinois 60126

L. Bellantoni, R.B. David, G.J. Bock, S. Childress, R. Coleman, M.B. Crisler, R. Ford,
Y.B. Hsiung, D.A. Jensen, P.T. Johnson, H. Nguyen, V. O'Dell, M. Pang, R. Pordes,
E.J. Ramberg, R.E. Ray, P. Shanahan, R. Tschirhart, H.B. White, J. Whitmore
Fermi National Accelerator Laboratory, Batavia, Illinois 60510

K. Hanagaki, M. Hazumi, S. Hidaka, F. Kato, K. Kuribayashi, K. Mori, M. Sadamoto,
K. Senyo, I. Suzuki, M. Yagi, T. Yamanaka
Osaka University, Toyonaka, Osaka 560-0043 Japan

A. Bellavance, M.D. Corcoran
Rice University, Houston, Texas 77005

S.Averitte, J. Belz, E. Halkiadakis, A. Lath, S. Schnetzer, S. Somalwar, R.L. Stone,
R. Tesarek, G.B. Thomson
Rutgers University, New Jersey

M. Arenton, G. Corti, B. Cox, K. Hagan, V. Jejer, A. Ledovskoy, A.P. McManus,
K.S. Nelson
The Department of Physics and Institute of Nuclear and Particle Physics, University of Virginia, Charlottesville, Virginia 22901

A. Alavi-Harati, T. Alexopoulos, A.R. Erwin *University of Wisconsin, Madison, Wisconsin 53706*

†Permanent address: *C.P.P. Marseille/C.N.R.S., France*

Figure 0.1 KTeV collaborators for the 1997 data set.

The 1999 KTeV Collaboration

E. Cheu, J. Hamm, S.A. Taegar, J. Wang
University of Arizona, Tucson, Arizona 85721

K. Arisaka, W. Slater, A. Tripathi, E.T. Worcester
University of California at Los Angeles, Los Angeles, California 90095

H.G.E. Kobrak
University of California at San Diego, La Jolla, California 92093

C.O. Escobar
Universidade Estadual de Campinas, Campinas, Brazil 13083-970

M. Barrio, E. Blucher, C. Bown, S. Bright, A. Glazov, J. Graham, R. Kessler, N. Lai, E. Monnier[†], V. Prasad, B. Quinn, N. Solomey, Y.W. Wah, B. Winstein, R. Winston
The Enrico Fermi Institute, The University of Chicago, Chicago, Illinois 60637

A.R. Barker, H. Huang, J. LaDue, P.A. Toale, M. Wilking
University of Colorado, Boulder, Colorado 80309

E.C. Swallow
Elmhurst College, Elmhurst, Illinois 60126

L. Bellantoni, G.J. Bock, R. Coleman, R. Ford, Y.B. Hsiung, D.A. Jensen, P.L. McBride, H. Nguyen, X.R. Qi, E.J. Ramberg, R.E. Ray, P. Shanahan, R.J. Tesarek, R. Tschirhart, H.B. White, J. Whitmore
Fermi National Accelerator Laboratory, Batavia, Illinois 60510

K. Hanagaki, K. Kotera, K. Senyo, T. Yamanaka
Osaka University, Toyonaka, Osaka 560-0043 Japan

A. Bellavance, M.D. Corcoran
Rice University, Houston, Texas 77005

R.F. Barbosa, P. Gouffon, E. Santos, R.F. Zukanovich
Universidade de Sao Paolo, Sao Paolo, Brazil 05315-970

M. Arenton, B. Cox, A. Golossanov, V. Jejer, A. Ledovskoy, K.S. Nelson, J. Shields
The Department of Physics and Institute of Nuclear and Particle Physics, University of Virginia, Charlottesville, Virginia 22901

A. Alavi-Harati, T. Alexopoulos, A.R. Erwin *University of Wisconsin, Madison, Wisconsin 53706*

[†]Permanent address: *C.P.P. Marseille/C.N.R.S., France*

Figure 0.2 KTeV collaborators for the 1999 data set.

Contents

Chapter 1
Introduction: High Energy Physics and the Standard Model

The field of high energy physics (HEP) deals with the study of sub-atomic particles. HEP is a relatively young science. It began in 1897 with J.J. Thompson's discovery of the electron, took a major step forward with Rutherford's nuclear model in 1911, and moved into the modern era with the construction of the first particle accelerator in 1932. The use of particle accelerators allows experimentalists to briefly turn time back to the high-energy particle soup which followed the Big Bang.

Our current state of knowledge is summarized by the Standard Model for particle physics. The Standard Model indicates that all matter is formed of quarks and leptons, which are listed in Figure 1.1. Each has three generations and six flavors, with two flavors per generation and the mass increasing with each increase in generation. Leptons, which include electrons, can be found free in nature. Quarks exist only in sets of triplets or matter-antimatter pairs, known collectively as hadrons. Quark triplet hadrons are particles called baryons and include the familiar proton and neutron. Mesons are quark-antiquark pair hadrons and include kaons and pions, which are significant in this analysis.

The Standard Model also contains a set of rules by which particles interact. Quarks and leptons feel four forces, but gravity is weak compared to other three

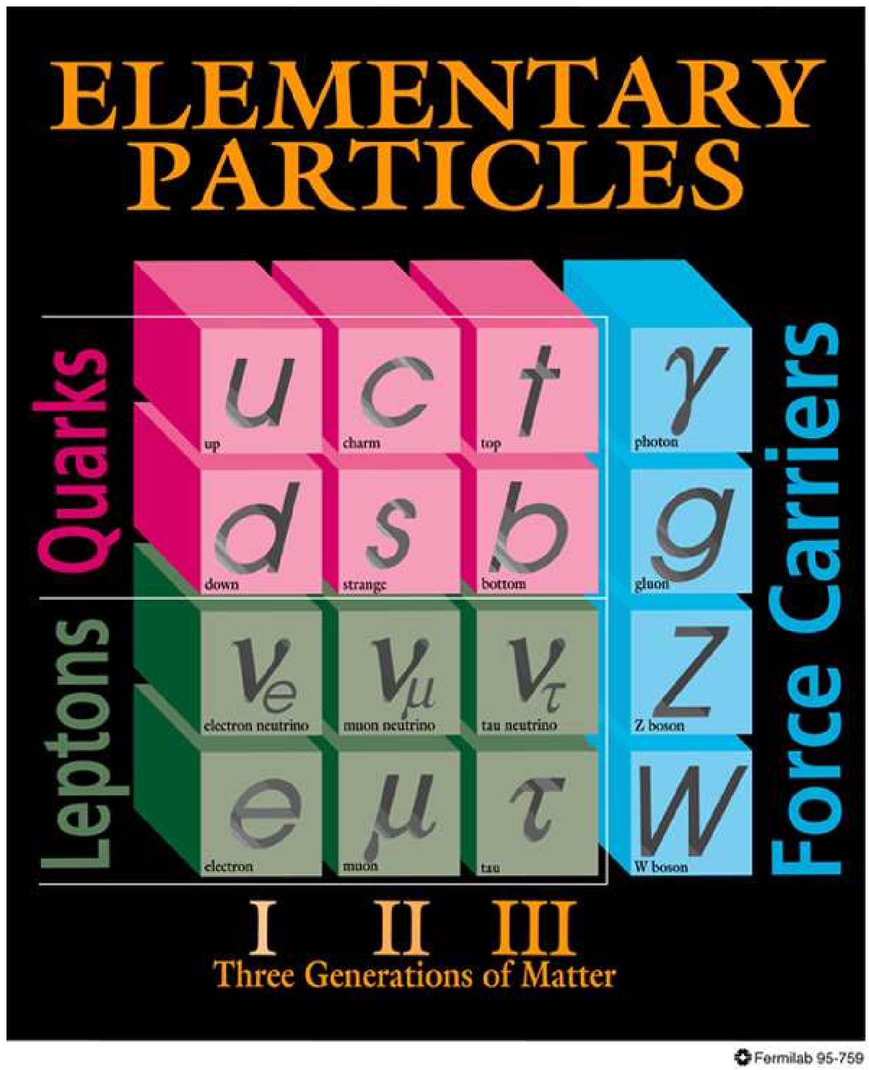

Figure 1.1 A table of particles in the Standard Model[1]. Image courtesy of the Fermilab Photo Archive.

on the scale that we deal with. In high energy experiments, quarks and leptons interact by exchanging force particles called gauge bosons: photons, gluons, W's, and Z's. A photon carries the electromagnetic force. The strong force is carried by gluons. The weak force is carried by W and Z bosons. Beyond the forces, there are other Standard model rules that are based on symmetry or invariance principles: conserved quantities. One oddity of the Standard Model is that some quantities have been shown experimentally to be conserved, but their conservation has no basis in symmetries or invariance principles. Why should these additional quantities be conserved when we see no symmetry for them? Are they actually conserved? If not, at what level are they violated? As one can see, the Standard Model is incomplete, and questions like those above prompt analyses such as this one.

1.1 Symmetries and Conserved Quantities

A basic tenet of classical physics is that our universe is uniform. The laws that govern it are the same at all locations and in all directions. This uniformity leads to conservation laws. Invariance under translations in space leads to the conservation of linear momentum. Invariance under spatial rotations leads to the conservation of angular momentum. Energy conservation is a consequence of invariance under time translations. The one-to-one correspondence between a symmetry and a conserved quantity in classical mechanics is known as Nöether's Theorem.

The relation between symmetries and conserved quantities carries over from the macroscopic classical realm into quantum mechanics. In the Standard Model, the concept of coordinate inversion, or parity, is important. For three dimensions, it is conventional to use a right-handed coordinate system, where $\hat{x} \times \hat{y}$ is in the positive $\hat{z}$ direction. Conceivably one could just as easily use a left-handed coordinate system, where $\hat{x}' \times \hat{y}'$ is in the negative $\hat{z}'$ direction. The two systems are related to each other by a mirror-like reflection in one plane followed by a 180° rotation [2]. No amount of transposition and rotation can make a right-handed system left-handed, or vice-versa, without the mirror reflection. The transformation is achieved via the *parity operator* (P), which reflects the spatial coordinates along their axes, resulting in the transformation

$$P(x, y, z) \;\rightarrow\; (x' = -x, y' = -y, z' = -z). \tag{1.1}$$

If a wavefunction is an eigenstate of an operator, then the operator has the effect of just multiplying the entire wavefunction by a real number, which is called the eigenvalue. The eigenvalues of (P) are +1 and -1. If the parity eigenvalue is +1 (-1), then the wavefunction is said to have *even* (*odd*) parity. The total parity of a system is a multiplicative quantum number, or in other words, it is the product of the parities of the components of the system. Since so many other dimensionally-based conservation laws hold, and since P is a good quantum number in electromagnetic and strong interactions, it was natural to think that parity would also be conserved in weak interactions. We now know it is not.

We can define another symmetry operator which affects some of a particle's properties but not its spatial coordinates. The *charge conjugation operator* (C) reverses the sign of all internal quantum numbers, including electric charge, magnetic moment, baryon number, lepton number, strangeness, charm, beauty, and truth. C does not affect energy, mass, momentum, or spin[3]. In essence, C exchanges a particle for its antiparticle but leaves the other physical properties unchanged. Some examples are:

$$C \mid e^- > = \ a_e \mid e^+ > \tag{1.2}$$

$$C \mid \pi^+ > = \ a_{\pi^-} \mid \pi^- > \tag{1.3}$$

$$C \mid \pi^0 > = \ a_{\pi^0} \mid \pi^0 > . \tag{1.4}$$

The coefficients a_e and a_π are phase factors such that a^2 is 1. As one can see from Equation 1.4, some particles are their own anti-particles. Only particles which are their own anti-particles are eigenstates of C. Like parity, the value of charge conjugation is multiplicative rather than additive, with eigenvalues of +1 or -1. It has been known since 1956, when suggested by T.D. Lee and C.N. Yang[4] and confirmed by others through experiment[5][6], that (P) and (C) are not conserved individually in weak interactions. However, the relationship between these two operators leads one to ask if the combined operation (CP) might be conserved (see Section 1.2).

Quantities which are conserved experimentally but for which there seem to be no underlying symmetry principles include *baryon number* (B) and *lepton number* (L). Both are additive, as opposed to multiplicative, properties. The values for B are +1 for baryons (or +1/3 for each quark), -1 for anti-baryons (or -1/3 for each anti-quark), and 0 for all other particles. The values for L are +1 for leptons, -1 for antileptons, and 0 for all other particles. There is another quantity called *lepton flavor* that is related to lepton number, but makes a distinction between the three generations of leptons. Each flavor has its own number designated by the subscript of the lepton type (L_e for electron flavor, L_μ for muon flavor, or L_τ for tau flavor). Like L, L_l has a value of +1 for each kind of lepton l^- or ν_l, -1 for each kind of antilepton l^+ or $\bar{\nu}_l$, and 0 for all other particles including leptons of other generations. Lepton flavor violation has been observed[7], but only for neutrinos that travel macroscopic distances (on the order of kilometers or more). It has not been observed in any production-and-decay type of experiment. For example, the dominant muon decay mode is

$$\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu \tag{1.5}$$

which has a lepton number of L=1, and lepton flavors of L_μ=1 and L_e=0 in both the initial and final states, and is a well-observed decay. On the other hand, the decay

$$\mu^- \rightarrow e^- \gamma \tag{1.6}$$

is lepton number conserving (again with L=1), but lepton flavor violating (with a L_μ=1, L_e=0 state going into a L_μ=0, L_e=1 state). The decay mode in equation (1.6) has never been observed and has a branching ratio limit of 1.2×10^{-11} [8].

One can compare baryon number, lepton number, and lepton flavor to the charge of the electromagnetic (EM) field. There are underlying symmetries related to EM fields, two of which are global gauge invariance and local gauge invariance. “Gauge” refers to the phase of a particle's wave function $\psi(x)$, so global gauge invariance means that $\psi(x)$ is invariant under the transformation

$$\psi(x) \rightarrow e^{i\alpha}\psi(x) \tag{1.7}$$

where α is a constant, and local gauge invariance means that it is invariant under the transformation

$$\psi(x) \rightarrow e^{i\alpha(x)}\psi(x) \tag{1.8}$$

where $\alpha(x)$ is a function of x. Invariance requires that the Lagrange density ($\mathcal{L}$) of the system is constant, or in other words

$$\delta\mathcal{L} \;=\; 0 \tag{1.9}$$

under a symmetry transformation. Substituting the global gauge transformation 1.7 into equation 1.9 leads to the requirement that charge be conserved [9]. Therefore, when a particle with a given amount of charge is created, other particles with an equal but opposite amount of charge must also be created. In this way, the quantity of charge is analogous to the quantity of B, L and L_l. On the other hand, the condition that $\delta L = 0$ for local gauge invariance requires that there be a massless particle associated with the EM field: the photon. If similar underlying symmetries existed for baryons, for example, we would expect a coupling field associated with baryons and a photon-like baryonic particle. However, no such field has ever been observed and there is no particle analogous to the photon of the EM field for B, L, or L_l [10].

The conservation of baryon number and lepton number implies that matter is only created in conjunction with antimatter. However, our physical universe is composed mostly of matter. The absence of the corresponding antimatter is a major cosmological mystery and is proof that B and L have not always been conserved. CP violation has been known since 1968 to be a necessary condition for generation of the matter/antimatter asymmetry in the early universe [11]. KTeV, a combination of two experiments (E832 and E799II) which share a single detector, was built to answer some of the questions surrounding CP violation in the kaon system. Kaons are mesons which contain strange quarks. They come in four types: K^+ (= the $u\bar{s}$ quark combination), K^- (= $\bar{u}s$) , K^0 (= $d\bar{s}$), and $\bar{K}^0$ (= $\bar{d}s$). KTeV deals only with the neutral kaons K^0 and $\bar{K}^0$. Experiment E832 investigates the amount of direct CP violation in the kaon system through the measurement of ϵ'/ϵ. Experiment E799II studies a wide range of rare kaon decays, including CP violating decays, lepton number violating decays, and lepton flavor violating decays.

1.2 Indirect and Direct CP Violation

All particles have an intrinsic parity (P) eigenvalue, but only particles which are their own antiparticles are also eigenstates of charge conjugation (C). The K^0 and $\bar{K}^0$ particles each have a negative intrinsic parity but individually they are not eigenstates of C. If we take linear combinations of the two particles, then eigenstates of C can be constructed for this system. Parity is not conserved in the weak interaction and since kaons only decay through the weak interaction, the final parity value is undetermined. What about the combined operator CP? Is it conserved in neutral kaon decay? Let's construct CP eigenstates and check the experimental consequences. The phase of C can be chosen so that

$$C \mid K^0\rangle = - \mid \bar{K}^0\rangle \tag{1.10}$$
$$C \mid \bar{K}^0\rangle = - \mid K^0\rangle \tag{1.11}$$

and then

$$CP \mid K^0\rangle = + \mid \bar{K}^0\rangle \tag{1.12}$$
$$CP \mid \bar{K}^0\rangle = + \mid K^0\rangle. \tag{1.13}$$

The CP eigenstates of K^0 and $\bar{K}^0$ are then expressed as K_1 and K_2 through the equations

$$\mid K_1\rangle = \frac{1}{\sqrt{2}}[\mid K^0\rangle + \mid \bar{K}^0\rangle]; \quad CP|K_1\rangle = +|K_1\rangle \tag{1.14}$$
$$\mid K_2\rangle = \frac{1}{\sqrt{2}}[\mid K^0\rangle - \mid \bar{K}^0\rangle]; \quad CP|K_2\rangle = -|K_2\rangle. \tag{1.15}$$

If CP were conserved, then only the following decays into pions would be seen:

$$K_1 \rightarrow 2\pi; \quad CP = +1 \tag{1.16}$$
$$K_2 \rightarrow 3\pi; \quad CP = -1. \tag{1.17}$$

As a reminder, pions are mesons with up and down quark combinations. The π^+ is $u\bar{d}$, π^- is $\bar{u}d$, and π^+ is $\frac{1}{\sqrt{2}}(u\bar{u}\text{-}d\bar{d})$. The K_2 state has a much longer lifetime than the K_1 state because of the small phase space available to this three body decay. By forcing a kaon beam to travel a relatively long distance, the K_1 particles should decay out and leave behind only the K_2 state. When this resulting beam is observed, CP conservation would allow only the 3π decay mode to occur. However, experiments have shown that the 2π decay mode occurs about 0.2% of the time[12], in violation of CP conservation. Since CP is not conserved, the CP eigenstates are not the same as the mass eigenstates of the kaon system.

A more general formalism using an admixture of the two CP eigenstates must be used to find the mass eigenstates. By diagonalizing the Hamiltonian matrix $\mathbf{H} = \mathbf{M}$ - i$\mathbf{\Gamma}/2$, where $\mathbf{M}$ and $\mathbf{\Gamma}$ are 2×2 Hermitian matrices, the mass eigenstates are found to be[13]:

$$\mid K_S\rangle = \frac{1}{\sqrt{(1+\mid \epsilon \mid^2)}}[|K_1\rangle + \epsilon|K_2\rangle] \tag{1.18}$$

$$\mid K_L\rangle = \frac{1}{\sqrt{(1+\mid \epsilon \mid^2)}}[|K_2\rangle + \epsilon|K_1\rangle] \tag{1.19}$$

where ϵ is given by

$$\epsilon = \frac{\langle K^0 \mid \mathbf{H} \mid \bar{K}^0\rangle - \langle \bar{K}^0 \mid \mathbf{H} \mid K^0\rangle}{i(\mathit{\Gamma}_S - \mathit{\Gamma}_L)/2 - (m_S - m_L)} \tag{1.20}$$

with the masses (m_S and m_L) and the decay widths $\mathit{\Gamma}_S$ and $\mathit{\Gamma}_L$ of K_S and K_L being the eigenvalues of $\mathbf{M}$ and $\mathbf{\Gamma}$ respectively. K_S represents a "short-lived kaon" ($\tau_{K_S} = 8.9\times10^{-11}s$) and is the part of the beam which decays out quickly. K_L represents a "long-lived kaon" ($\tau_{K_L} = 5.2\times10^{-8}s$) and makes up the remaining beam when the K_S's are gone. In these eigenstates, the parameter ϵ measures the amount of mixing of CP eigenstates K_1 and K_2 in the wave functions of the mass eigenstates K_S and K_L. Such mixing is referred to as *indirect* CP violation. If ϵ represented the only source of CP violation in this system, then the decay-rate ratios

$$\frac{\Gamma(K_L \to \pi^0\pi^0)}{\Gamma(K_S \to \pi^0\pi^0)} = \mid \eta_{00} \mid^2 \tag{1.21}$$

and

$$\frac{\Gamma(K_L \to \pi^+\pi^-)}{\Gamma(K_S \to \pi^+\pi^-)} = \mid \eta_{+-} \mid^2 \tag{1.22}$$

should be equal. However, the Standard Model permits another source of CP violation in which K_2 decays directly to $2\pi^0$. This source is called *direct* CP violation and is parameterized by ϵ'. The parameter ϵ' is related to η_{00} and η_{+-} by[14]:

$$\eta_{00} \approx \epsilon - 2\epsilon' \tag{1.23}$$

$$\eta_{+-} \approx \epsilon + \epsilon' \tag{1.24}$$

It relates to physically measurable quantities in the following way:

$$\frac{\Gamma(K_L \to \pi^0\pi^0)/\Gamma(K_S \to \pi^0\pi^0)}{\Gamma(K_L \to \pi^+\pi^-)/\Gamma(K_S \to \pi^+\pi^-)} = \mid \frac{\eta_{00}}{\eta_{+-}} \mid^2 \approx 1 - 6Re(\epsilon'/\epsilon) \tag{1.25}$$

If ϵ' is zero then CP violation only occurs indirectly. If ϵ' is not zero then the ratio of charged decays differs from the ratio of neutral decays by some amount. The value of ϵ' can be calculated from the four measured values on the left side of equation (1.25). Currently the Standard Model predicts a positive value for ϵ'/ϵ, but the exact value is difficult to calculate.

Experiments at Fermilab (namely experiments E731, E773, and E799 Phase-I (E799I) at the MCenter facility) have studied neutral kaons in fixed target experiments for many years, but the experimental data were not precise enough to obtain satisfactory measurements of $Re(\epsilon'/\epsilon)$. In the early 1990's, experiment E731 claimed a value of $7.4 \pm 5.2 \pm 2.9 \times 10^{-4}$ for $Re(\epsilon'/\epsilon)$ [15], while competing experiment NA31 (at the CERN laboratory in Switzerland) claimed a value of $23 \pm 3.4 \pm 6.5 \times 10^{-4}$ [16]. Since the results did not agree, it was decided to redo both experiments. The design report for KTeV was proposed and approved in January of 1992, and the KTeV experiment ran from August of 1996 to September of 1997. Following some minor repairs, KTeV ran again from June 1999 to January 2000. The most recent results of the $Re(\epsilon'/\epsilon)$ measurement for E832 were released on June 8, 2001, at a Fermilab Wine & Cheese presentation. Using the full 1996-1997 data set, we measured a value of $20.7 \pm 2.8 \times 10^{-4}$ [17]. The remounted effort at CERN was named NA48 and their most recently measured value is $15.3 \pm 2.6 \times 10^{-4}$ [18].

1.3 Rare Kaon Decays

While experiment E832 studied CP violation, experiment E799II studied rare decay modes of K_L. Decays are considered rare if they have a branching ratio of less than 0.1%, but E799II study modes have even smaller branching ratios. The *branching ratio* of a decay mode is the probability that the particle will undergo that decay, and not any of the other possible decays. The decay modes examined in the E799II experiment have branching ratios of 10^{-7} or less. The focus of this report is a search for the rare decay mode:

$$K_L \rightarrow \pi^0 \mu^{\pm} e^{\mp} \quad (called\ \ Kpi0me) \tag{1.26}$$

which is lepton flavor violating. The neutral pion (π^0), having a very short lifetime (relative to the distance traversed in the detector) will decay almost immediately into a photon pair. Therefore in looking for this event, one would look for a muon (μ), an electron (e), and two photons (γ) to appear in the detector.

There are several theories in which lepton flavor is violated. Evidence from solar and atmospheric neutrino experiments imply that the Standard Model will need to (and can easily) be modified to allow for the possibility of lepton flavor violation through neutrino oscillation. Theoretically, the strange quark in the kaon could decay into an up quark, and the up quark then decay into an anti-down (see Figure 1.2). A W boson would be released in each quark decay. One of the W's can decay into a positive (or negative) muon and a muon (anti-)neutrino, and the other could go to an electron (or positron) plus electron anti-neutrino (or neutrino). The muon neutrino could then oscillate to an electron neutrino and annihilate with the electron

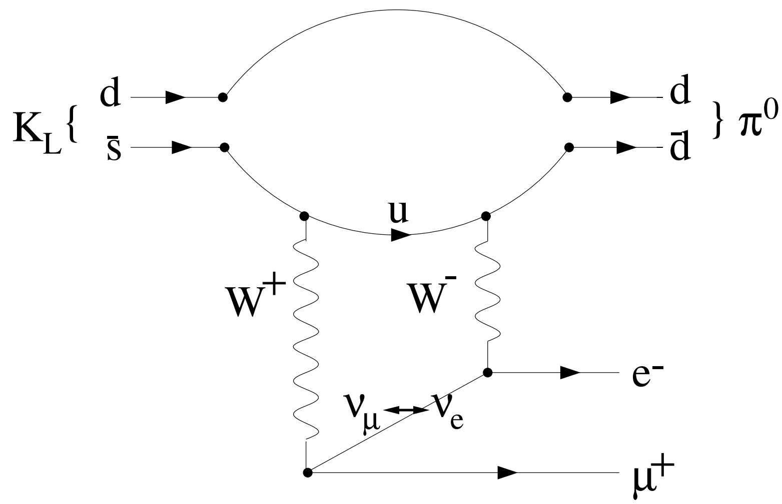

Figure 1.2 One example of the decay $K_L \to \pi^0 \mu e$, with a possible avenue for Lepton Flavor violation.

anti-neutrino in the system. The annihilation energy would be given as momentum to the remaining particles, allowing the kaon's mass to be fully reconstructed. If no oscillation occurs in the decay, then the neutrino and anti-neutrino would escape our detector, making the parent particle's reconstructed mass be less than the kaon mass. Evidence for neutrino oscillation has been seen by the Super-Kamiokande experiment in Japan[19] and the Sudbury Neutrino Observatory (SNO) in Canada[20]. Simple SM extensions with massive neutrinos predict a branching ratio for the $K_L \to \mu e$ decay mode of on the order of 10^{-24}[21]. Since our Kpi0me signal is a three-body decay, its branching ratio through this mechanism should be equivalent or even smaller.

Another extension of the Standard Model that would allow lepton flavor violation is a Supersymmetry model which allows R-parity violation, either through a neutralino decay (see Figure 1.3) or a leptonic squark decay (see Figure 1.4). In the

Supersymmetry model, each Standard Model particle has a supersymmetric partner of the same mass but with a spin that differs by 1/2 unit. No evidence has yet been found for Supersymmetry, but neither has it been ruled out as a possibility. Some supersymmetric models require R-parity be conserved while others allow it to be violated. R-parity (R_P) is a combination of baryon number (B), lepton number (L), and spin quantum number S.

$$R_P \;=\; (-1)^{3B+L+2S} \tag{1.27}$$

R_P is +1 for standard model particles and -1 for Supersymmetric particles. If R-parity is conserved, processes with $\Delta L \neq 0$ and $\Delta B \neq 0$ are suppressed and Supersymmetric particles are always produced in pairs with standard particles. R-parity violation allows the independent creation of a supersymmetric particle, like a single neutralino or a single squark as in our cases. If our search decay occurs through a neutralino, the branching ratio is estimated to be 7.71×10^{-16} in this model and if it occurs through a squark, the branching ratio would be 9.06×10^{-17} using this model[23].

Lepton flavor can be violated through horizontal gauge bosons. In our signal's case, the horizontal gauge boson would be a Z-like neutral particle that can change both quark and lepton flavors (call it V^0). For example, the production of a V^0 would change the kaon's $\bar{s}$ into a $\bar{d}$, and the V^0 itself could decay into a muon and an electron (see Figure 1.5). One horizontal boson model predicts a branching ratio for our signal in the 10^{-7} to 10^{-12} range[24], with the smaller values favored since the

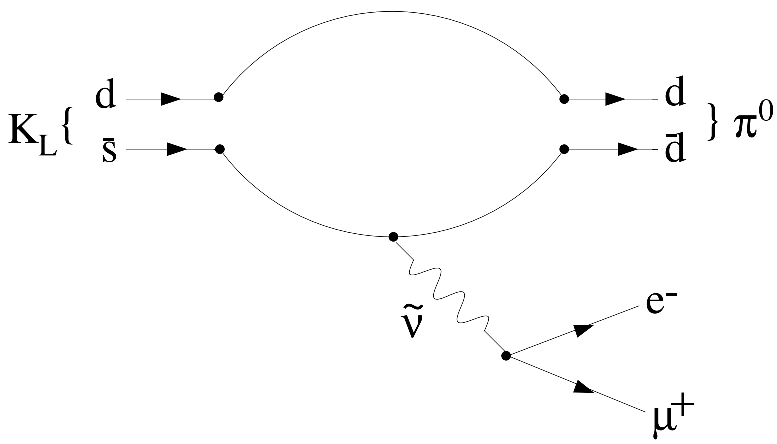

Figure 1.3 Tree-level example of the decay $K_L \to \pi^0 \mu e$, with a supersymmetric neutralino and R-parity violation.

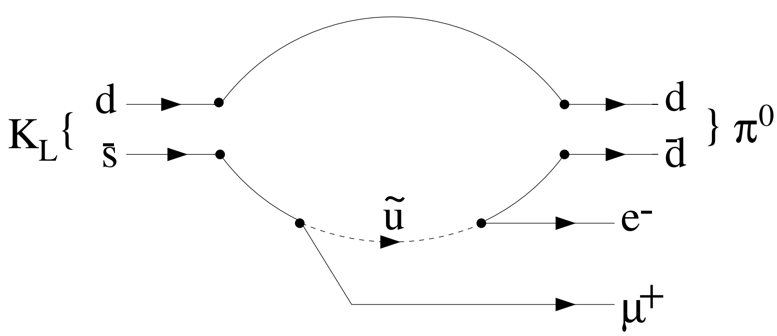

Figure 1.4 Tree-level example of the decay $K_L \to \pi^0 \mu e$, with a supersymmetric squark and R-parity violation.

current limit on $K_L \to \mu e$ is below 10^{-10}[26].

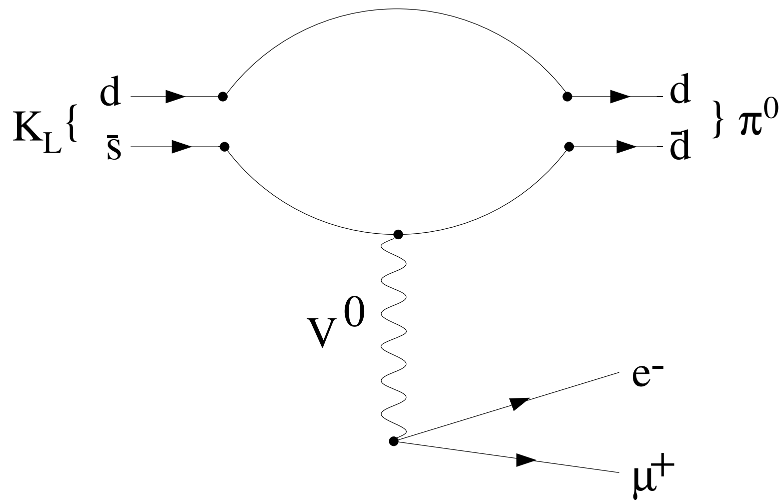

Figure 1.5 A possible avenue of lepton flavor violation for our signal decay involving a new horizontal gauge boson (here labeled V^0).

Still another possibility is the creation of a leptoquark, which carries both lepton number and baryon number[25]. A possible scenario is illustrated in Figure 1.6, with the $\bar{s}$ quark emitting a leptoquark and changing into a muon. The leptoquark then decays into an electron and a $\bar{d}$. The $\bar{d}$ joins the d to form the neutral pion.

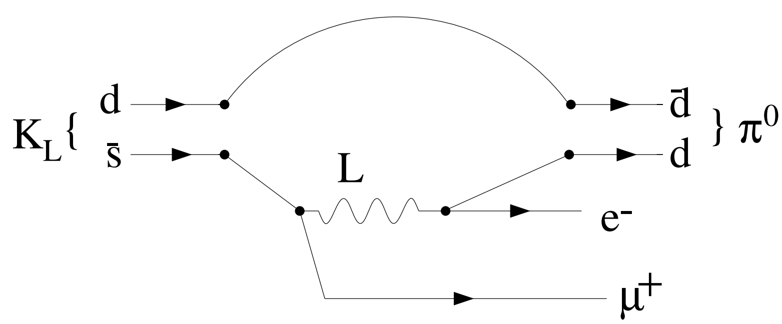

Figure 1.6 A possible avenue of lepton flavor violation for our signal decay involving production of a leptoquark (here labeled L).

If any of these theoretical estimates is correct we have little chance of observing any events as we expect a sensitivity in the range of 10^{-10}. If any signal events are seen in this analysis, they would be clear evidence of *new physics* beyond the Standard Model. The current limit listed in the Particle Data Book was 6.2×10^{-9} [8] from Experiment E799I[27]. I presented a preliminary result of 4.40×10^{-10} [28] at a DPF conference in April of 2000 using part of the 1997 KTeV data set. This study will improve upon the preliminary result, or will be a source of new physics should events be seen.

Decay modes that are not our signal but might leave a similar signature in the detector must be investigated as well. These modes (called *backgrounds*) are listed in Table 1.1 and will be discussed in detail in Section 3.6. Briefly, backgrounds can appear to be signal events through a combination of daughter particle decays, neutrinos going through the detector unseen, extra accidental activity in the detector, and/or particle misidentification.

Decay mode	Common name	Branching ratio
$K_L \to \pi^+\pi^-\pi^0$	K3pi	12.56 ± 0.20 %
$K_L \to \pi^\pm e^\mp \nu_e$	Ke3	38.79 ± 0.28 %
$K_L \to \pi^0\pi^\pm e^\mp \nu_e$	Ke4	$(5.18 \pm 0.29) \times 10^{-5}$

Table 1.1 Significant background decay modes for $K_L \to \pi^0\mu^\pm e^\mp$, and their branching ratios[8].

Chapter 2
Experimental Equipment and Setup

2.1 Beam Acceleration and Structure

Fermi National Accelerator Laboratory (FNAL or Fermilab) is a synchrotron proton accelerator with facilities for both "collider" experiments and "fixed-target" experiments. Collider experiments have a detector which surrounds an interaction region, where beams of particles (protons and antiprotons at FNAL) with opposite momentum are made to cross so they collide nearly head-on. In fixed-target experiments, the beam is directed into target material and the produced particles are Lorentz boosted in the direction of the beam. Therefore fixed-target detectors are laid out linearly with the dimension in the beam direction being much larger than in the transverse direction. The Fermilab beam begins as atoms in a bottle of hydrogen gas. The electrons are stripped off and the protons are injected into the beam line in a stream. The stripped protons are sent through a Cockroft-Walton accelerator, a linear accelerator, and a small booster ring, where a target can be used to produce antiprotons for collider running. During the 1996-7 run, the beam traveled from the booster into the main ring, and finally the TeVatron ring. The main ring and TeVatron ring shared an underground tunnel which has a radius of one kilometer. The 1999 run saw the commissioning of a new "Main Injector" which was built in its own tunnel and replaced the retired main ring. Figure 2.1 shows a schematic of the

Fermilab accelerators as of 1999.

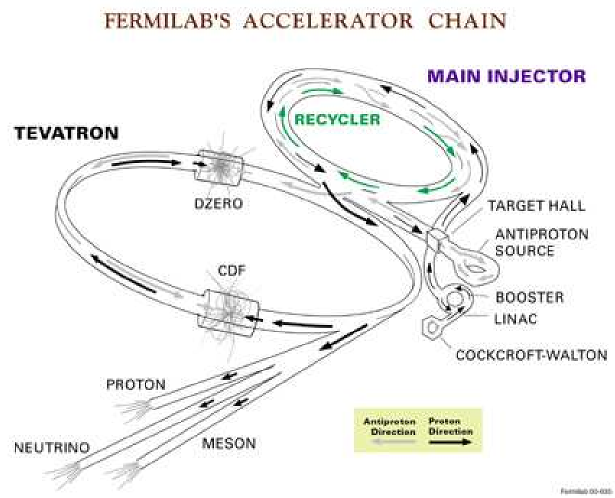

Figure 2.1 A schematic of the Fermilab Accelerator as of 1999 with the new Main Injector[29]. For fixed target running (including KTeV) the Recycler portions in green and anti-proton portions in gray are not used. In the 1996-7 KTeV run the old main ring, housed in the TeVatron tunnel, was the step used after the booster because the Main Injector was not yet built. The KTeV detector hall in on the central ("Neutrino") fixed target line. Image courtesy of the Fermilab Photo Archive.

Acceleration of particles beyond the Cockroft-Walton stage is achieved using radio-frequency (RF) accelerating cavities. The RF cavities divide the stream of particles into bunches or "buckets". Like someone pushing a playground merry-go-round, the RF cavities give each bunch a shove as it goes by. Simultaneously the current (and thereby the field strength) of dipole bending magnets around the ring is increased to keep the accelerating particles in a circular path. Alternating pairs of quadrupole magnets with the dipole magnets keep the protons focused into a beam.

One cycle for fixed-target operation mode consists of (a) particle injection, (b) particle acceleration, (c) a period of beam spillage into the fixed-target areas, and (d) a ramping down of the Tevatron's magnet current (see Figure 2.2). A spill (period (c)) is the diversion of a portion of the beam into the switchyard and out to all fixed target experiments (of which KTeV is one) until the beam is exhausted. This is when signal events occur and when most data is taken. In 1997, a spill occurred for 20 seconds per cycle, with one cycle lasting one minute. In 1999, the spill was 40 seconds long within an 80 second cycle.

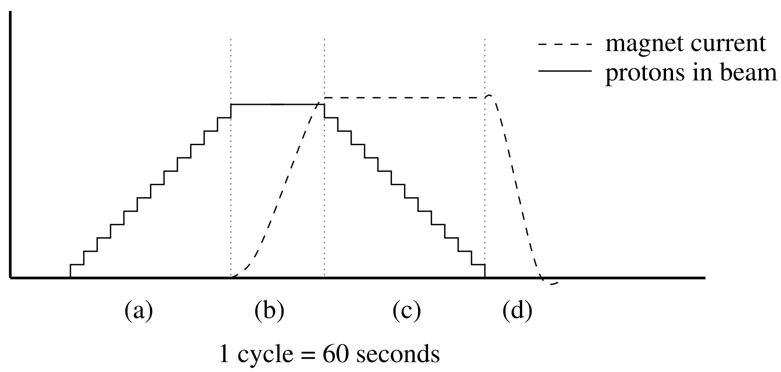

Figure 2.2 A qualitative plot of magnet current and proton beam population in the accelerator over the time period of one cycle. The stages include (a) injection; (b) acceleration; (c) beam delivery to targets; and (d) ramping down of the magnets.

The Main Injector accelerates particles to an energy of $150 GeV$ and the Tevatron takes them to $800 GeV$. Nominal beam intensity for E799II was 5×10^{12} particles on target per spill, and we wrote information to tape from about 15,000 particle decays (called "events") per spill.

2.2 Detector

KTeV was a fixed-target experiment, which means the beam of protons hit the front of a stationary piece of material (the target), producing a spray of many different particles that emerged from the back of the target. In KTeV, the higher energy particles were collimated into a beam which traveled linearly through the elements of our detector: a decay region, a spectrometer, a calorimeter, and finally muon detectors. Figure 2.3 shows a plan view of the KTeV detector as configured for E799II running.

2.2.1 Target and Decay Region

The KTeV production target was made of a $30cm$ length (1.1 interaction lengths) of beryllium-oxide (BeO). The center of the target is used as the origin of KTeV's coordinate system. X is measured in the horizontal (east-west) direction, Y is measured in the vertical (up-down) direction, and the beam traveled in the positive Z direction. The larger the value of Z, the farther "downstream" the location is said to be. Conversely, if one detector element was closer to the target than another, the first is said to be "upstream" of the second. The beam was filtered using a lead filter and sweeping magnets until only neutrons, kaons and neutral hyperons were left. The beam was then collimated into two beams and sent down a hundred-meter long evacuated pipe which allowed nearly all the K_S particles to decay out of the beam. The beams then traveled through a 60 meter evacuated decay region. The evacuated

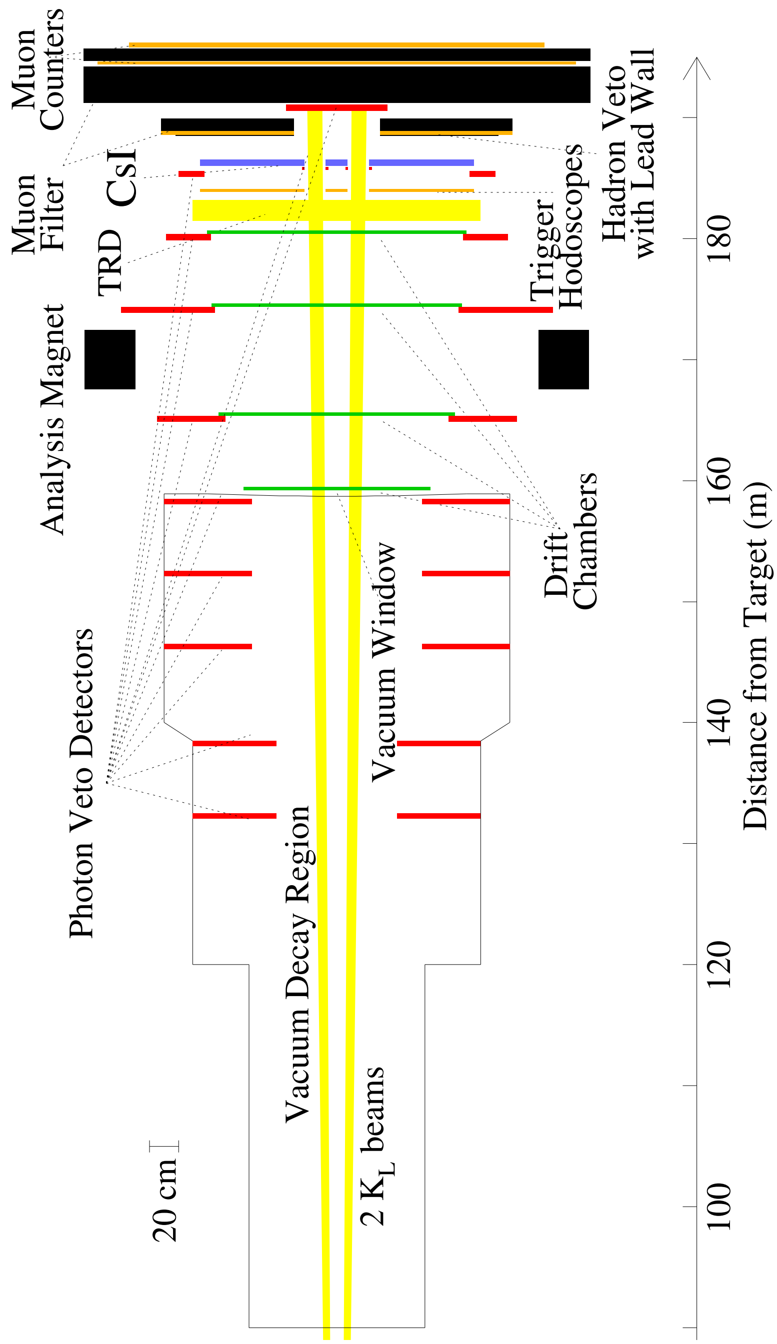

Figure 2.3 A plan view of the KTeV detector for E799II running. The target area is not shown, but the BeO target material is centered at $0m$ on the Z-axis. Note that the scales of X and Z are different, making the detector looked more squished in Z than it really is.

areas were kept at a pressure of 10^{-6} Torr. Finally the beams entered the detector hall through a thin window of Kevlar and aluminized Mylar. The nominal E799II kaon flux during the spill was about $5MHz$, with about 2% of the K_L's decaying in our vacuum region.

2.2.2 Regenerator

One of the two beams could have been sent through a "regenerator" which produced K_S particles from the K_L particles. The regenerator was an active detector consisting of 1.7 meters of scintillator with a beryllium mask upstream. Through strong interactions of the kaons with the regenerator's nuclear matter, the quantum states of the K_L particles were rotated through the K_1 and K_2 phase space to regenerate some component of K_S particles. To illustrate, let us ignore indirect CP violation for the moment and assume a K_L was traveling along as pure K_2 (and hence K_S=K_1). From Equation 1.15 (ignoring the normalization constant) :

$$\psi \;=\; K_2 \;=\; K^0 \;-\bar{K}^0. \tag{2.1}$$

The particle then encountered the material of the regenerator and underwent strong interactions. Kaons feel the strong force as K^0 and $\bar{K}^0$, not K_1 and K_2. The K^0 component would have traveled through the regenerator material with a forward scattering amplitude R. Let's call the $\bar{K}^0$ forward scattering amplitude $\tilde{R}$. R is not equal to $\tilde{R}$ because the K^0 and $\bar{K}^0$ have different constituent quarks. The resulting coherent wavefunction was then:

$$\psi \;=\; RK^0 \;-\; \tilde{R}\bar{K}^0. \tag{2.2}$$

Being a coherent state meant the different parts of the wavefunction acted as a single particle and could have interfered with each other. If the interaction happened to be inelastic, a secondary particle (or particles) would have been produced, and the event would have been vetoed by the regenerator itself. So what was the new wavefunction in terms of decay particles K_1 and K_2? We need to invert Equations 1.14 and 1.15 (again ignoring normalization constants) to get

$$K^0 \;=\; K_1 + K_2 \tag{2.3}$$

and

$$\bar{K}^0 \;=\; K_1 - K_2. \tag{2.4}$$

Substituting into Equation 2.2:

$$\psi \;=\; R(K_1 + K_2) \;-\; \tilde{R}(K_1 - K_2) \;=\; K_1(R - \tilde{R}) \;+\; K_2(R + \tilde{R}) \tag{2.5}$$

Voilà! After traveling through the material, the originally all K_2 particle had a non-zero component of K_1, and hence a non-zero probability of being a K_S. So with the regenerator in a beam line, the detector saw one beam of K_L and one beam of a coherent superposition of K_L and K_S.

For experiment E832, the regenerator was moved from one beam to the opposite beam each spill so that location-specific detector errors would be averaged out. For experiment E799II, the regenerator was completely removed from the beam line, as K_L particles are the focus of the rare decay studies.

2.2.3 Spectrometer

Upon entering the experimental hall through the Kevlar/Mylar window, the beams passed through a spectrometer which consisted of two pairs of drift chambers separated by a magnet. A spectrometer's purpose is to measure the momentum of each charged particle. After exiting the decay region through the vacuum window, the beams traversed two drift chambers separated in Z, each of which measured charged particle X-Y locations. From this information their direction of travel can be found. The particles then passed through the spectrometer magnet in which the path of each charged particle was bent. After the magnet, the beams went through another pair of drift chambers to detect their new direction. Since the amount of momentum translated into the transverse direction in the magnet was known, each particle's momentum could be calculated. Our spectrometer had a momentum-dependent resolution of $\sigma(P)/P = 0.38\% \oplus 0.016\% P(GeV/c)$ [30]. The space between components of the spectrometer were filled in with helium bags to reduce the amount of material the decay products had to traverse.

Multi-wire Drift Chambers

The multi-wire drift chambers for the spectrometer were Rice University's major contribution to KTeV. These drift chambers are an extension of the simple, single wire proportional counter. A single wire proportional counter is a metallic tube with a wire running down its central axis. The tube is filled with a mixture of gases which readily ionize and it is kept at ground potential while the inner wire is raised to a high positive voltage (usually 2000 to 3000 volts). When a charged particle passes through the counter, it ionizes the gas mixture. The freed electrons are then attracted to the wire. As they near the wire (where the electric field gradient is the largest), the free electrons are rapidly accelerated and knock additional electrons out of the gas. The result is a large number of free electrons, called an avalanche, being drawn away through the wire and producing a signal (see Figure 2.4).

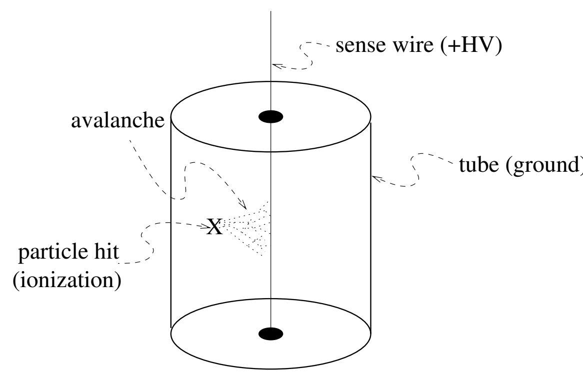

Figure 2.4 Single wire proportional counter (side view).

Multi-wire proportional chambers (MWPCs) work on the same principle, with the single, positive-voltage detector wire being replaced by planes of grounded sense wires interspersed with planes of negative voltage cathode wires. The sense and cathode wires are spaced so as to create a uniformly shaped, mappable electric field within the chamber. A signal on a particular wire gives a volume around that wire, called a *cell*, within which the particle passed. By layering planes of such MWPCs at right angles to each other, the path of a particle can be determined within a range limited by the cell size. For the geometry shown in Figure 2.5, the resolution is rather poor at one-half cell width. The resolution can be improved to one half cell width by using the geometry shown in Figure 2.6. KTeV's drift chambers use this half-cell offset geometry. Still, electrostatic limitations do not permit MWPCs with a sense wire spacing closer than about 1mm.

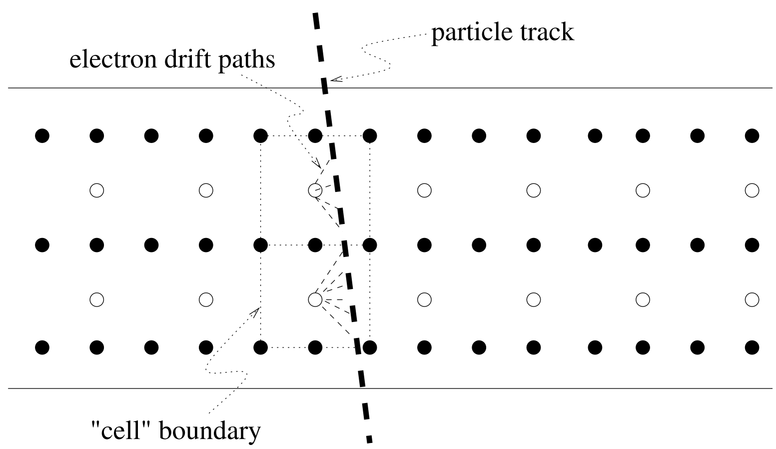

Figure 2.5 End-on view of a typical multiwire proportional counter (MWPC) plane. The black circles are cathode wires and the white circles are sense wires.

To produce an even more precise positioning of the particle track, one must move on to drift chambers. Drift chambers and MWPCs are physically constructed similarly. The major difference is that drift chambers also measure the drift time of the electrons using the properties of the gas filling the chamber. In regular MWPCs, one wants to receive a signal of the particle's passage as quickly as possible. It is therefore advantageous to accelerate the electrons all the way to the sense wire. The terminal

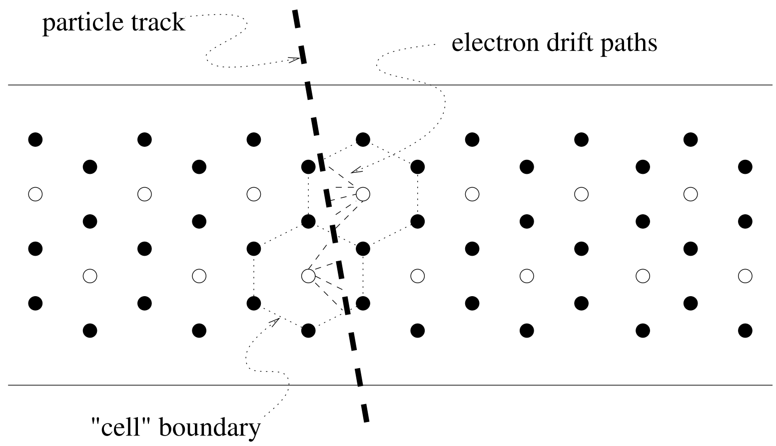

Figure 2.6 An end-on schematic view of wire layout for one plane of drift chambers using a hexagonal offset geometry. The black circles are cathode wires and the white circles are sense wires. KTeV's drift chambers use this geometry.

velocity of the electron should be high enough that it is never reached, and response time is not lost. The terminal velocity of an electron in a gas is determined by the characteristics of the gas. In drift chambers, the mix of gasses is selected so that the gas limits the velocity of the electrons. The strong electrical field ensures that the electron will reach terminal velocity quickly, but the gas ensures it will travel at that constant velocity the majority of the distance to the sense wire. Given an initial time (provided by other parts of the detector) and using the known terminal velocity, the drift time of the electrons can be measured and a radius of origin of the electron shower can be calculated very accurately. Drift times at KTeV were measured to $\pm 0.5 nsec$ by a time-to-digital converter (TDC), giving us a resolution of about $100 \mu m$. If there are two layers of sense wires as in Figure 2.6, one gets two distance measurements per particle. One can see from the figure that if the signals come from a particle passing between the wires roughly perpendicular to the wire plane, then the sum of the distances will add up to the diameter of one cell. Signals from noise do not add up to any particular distance. Requiring the sum-of-distances to be equal to the cell diameter can be used as an additional quality constraint on the track.

The drift chambers used in the KTeV E832 and E799II experiments ranged in size from ($1.251m \times 1.251m$) to ($1.759m \times 1.759m$). There were four chambers (DC1-4) numbered in order of increasing Z. Each contained two layers (or "views") of sense wires in the X-plane (X and X′), and two layers in the Y-plane (Y and Y′) with the layers offset in a hexagonal cell geometry of radius $0.635cm$. Figure 2.6 shows an end-view of the wire layout for one plane, and Figure 2.7 shows the electric field generated by one plane. The cathode wires were gold-plated aluminum with a diameter of $100\mu m$, and the sense wires were gold-plated tungsten with a diameter of $25\mu m$. The wires were individually strung by hand into a "pac-man" system (see Figure 2.8) to ensure the accuracy of the wire's position in the center of the input location. The gas in the chambers was 50% argon, with 50% ethane as the quenching gas. This mixture was bubbled through isopropyl at $-7°C$ on its way into the chamber to prevent polymerization of the gas on the cathode wires. Each wire plane could measure the location of a particle to within $100\mu m$. The analog signals from the chambers were recorded by time-to-digital converters (TDCs) with a resolution of $0.5nsec$ per count. The data-taking time window was set at the maximum drift time

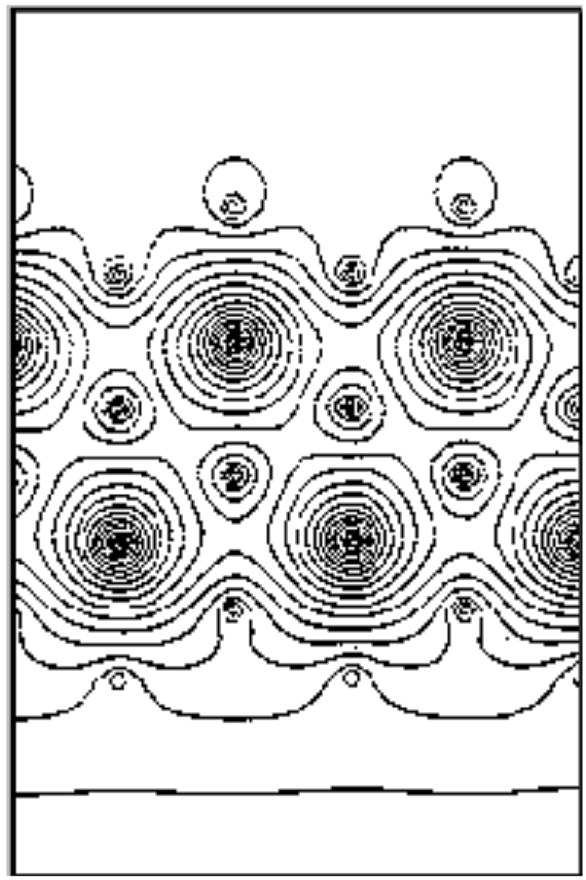

Figure 2.7 The contour diagram (end-on view) of the electric field generated by one plane of the KTeV drift chamber wires.

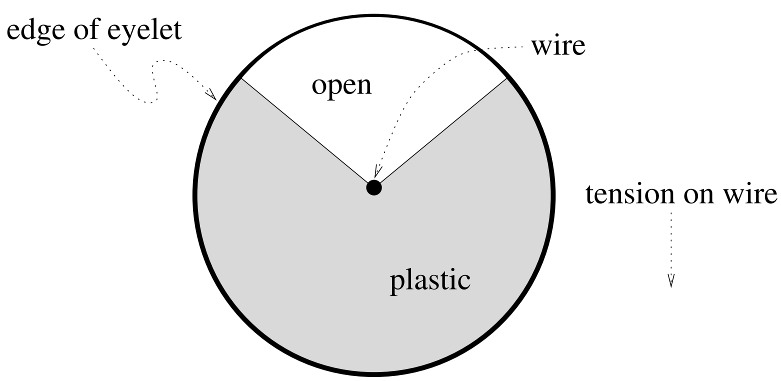

Figure 2.8 The physical support for the drift chamber wires to ensure accurate positioning. The diameter of the eyelet is approximately $0.5cm$.

of 250$nsec$, starting at 400 TDC counts and ending at 650 TDC counts. An exemplary drift-time distribution for the 1997 data run is shown in Figure 2.9. Chamber signals which fell inside the 250$nsec$ window were called "in-time" hits. The sense wire is represented by the sharp drop-off around 650 counts with each X-axis bin equaling the TDC resolution of 0.5$nsec$. The sidebands come from "out-of-time" hits. To be considered part of a good track, chamber hits had to be in-time hit "pairs", which meant the sum-of-distances for two adjoining hits was equal to the cell's diameter, within resolution.

The drift chambers sustained radiation damage during the 1996-7 running period, especially DC1 and DC2. The chamber efficiency slightly declined during running as non-conducting material accumulated on the wires. It was decided to refurbish DC1 and DC2 during the down-time before the 1999 run. All sense wires in DC1 were replaced, as were the central third of the sense wires in DC2 and the central third of the cathode wires in both chambers. The remaining wires were rubbed clean using very fine grit polishing paper glued to a tongue depressor and soaked with ethanol.

Some samples of the removed wires were taken back to Rice and analyzed with a secondary emission electron microscope (SEM) in our Nanotechnology Center, and a detector in our Geology Department usually used for studying rock samples. The SEM allowed us to visually scan the samples and capture the images. The wires exposed to the 1996-7 run clearly showed depositions not present on new wire (see

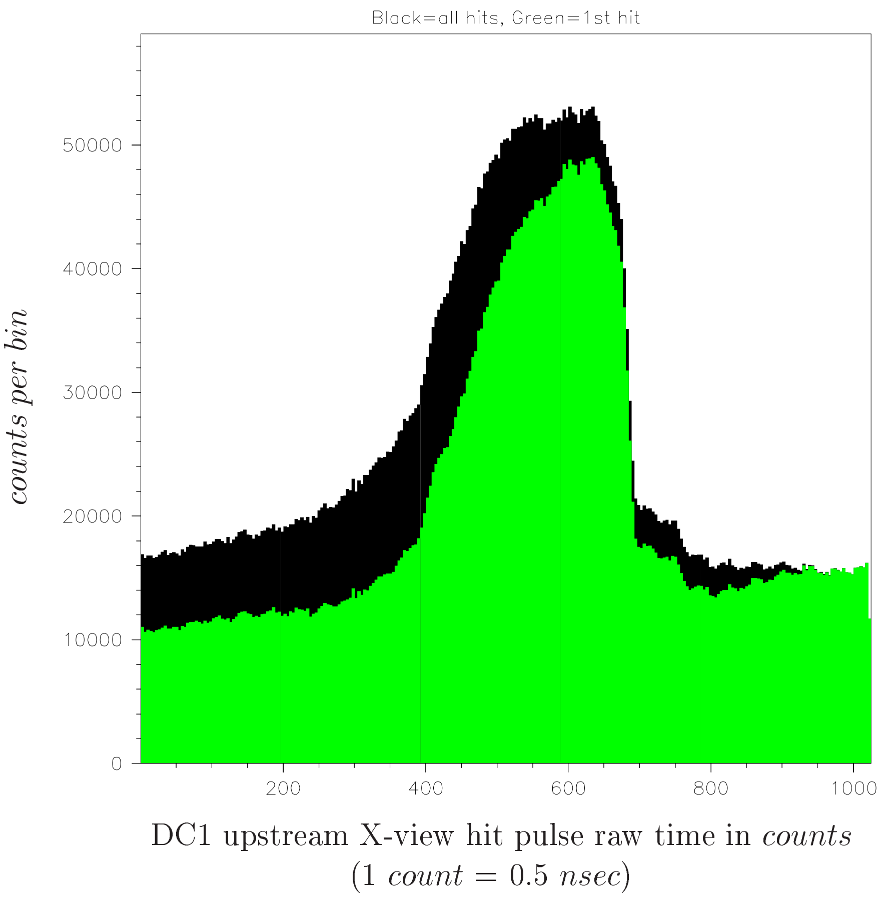

Figure 2.9 Exemplary drift-time distribution (histogram) for one of the drift chambers in KTeV (run #9225). Each count on the X-axis is $0.5nsec$. Signals occurring right on top of the sense wire formed the sharp drop-off near 650 counts. The sidebands were background from out-of-time signals. The black histogram included all hits (up to 16 were possible) that arrived during the trigger window, which was longer than the chamber in-time window. The green histogram included only the first hit to arrive in the trigger window.

Figures 2.10 through 2.15). A wire sample taken from the top of a drift chamber, away from the main beam area, shows only some deposition (Figure 2.12). Geology's device used x-ray fluorescence to determine the samples' composition. In x-ray fluorescence an electron beam about $1\mu m$ wide is used to excite the sample material, and the resulting x-rays are characteristic of the element emitting them. The depositions were composed mostly of silicon, however we could not think of a clear source for this contaminant. We did use room temperature vulcanizing sealant (RTV) on the chambers, which is mostly silicon, but this substance was used outside the gas volume. From talking to other experimenters we learned that this kind of radiation damage has been seen before, but again an obvious source of silicon was not present. One opinion offered was that the contaminant came from silicon-based lubricants commonly used on gas system connectors and mechanisms.

Spectrometer Magnet

The magnet gave charged particles a momentum "kick" in the (transverse) X direction. The momentum transfered into the transverse direction was essentially constant for the momentum range of our particles. That is,

$$\frac{dp}{dt} = q\mathbf{v} \times \mathbf{B} = \frac{m\gamma\mathbf{v}^2}{r} \tag{2.6}$$

where r is the radius of curvature due to the constant magnetic field. The distance the particle traveled in the magnetic field was short enough that we can use the small

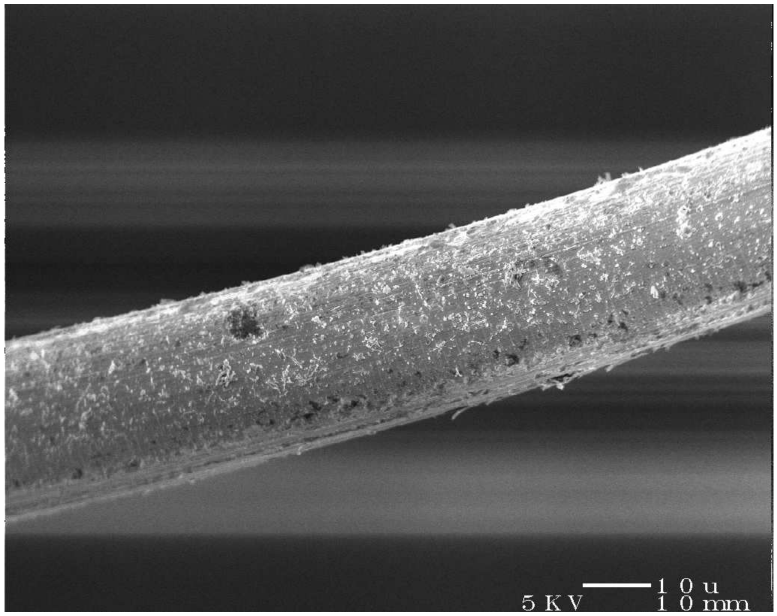

Figure 2.10 An electron microscope photograph of a sense wire removed from the center of a KTeV drift chamber after the 1996-7 run. Notice the blurring of the drawing machine groves as compared to Figure 2.14 from the accumulation of material on the wire surface.

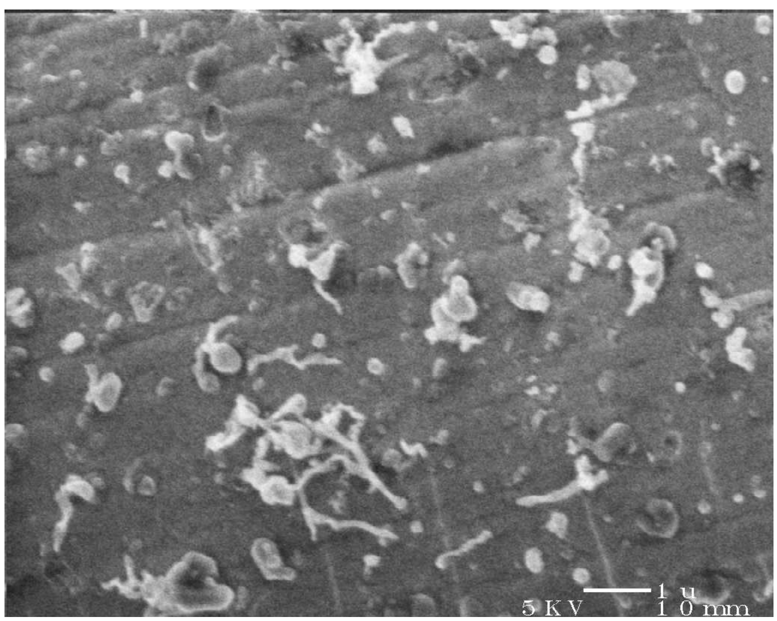

Figure 2.11 An electron microscope close-up of a sense wire removed from the center of a KTeV drift chamber after the 1996-7 run. Some material has coated the wire completely (dark background) and there are additional blob-like deposits.

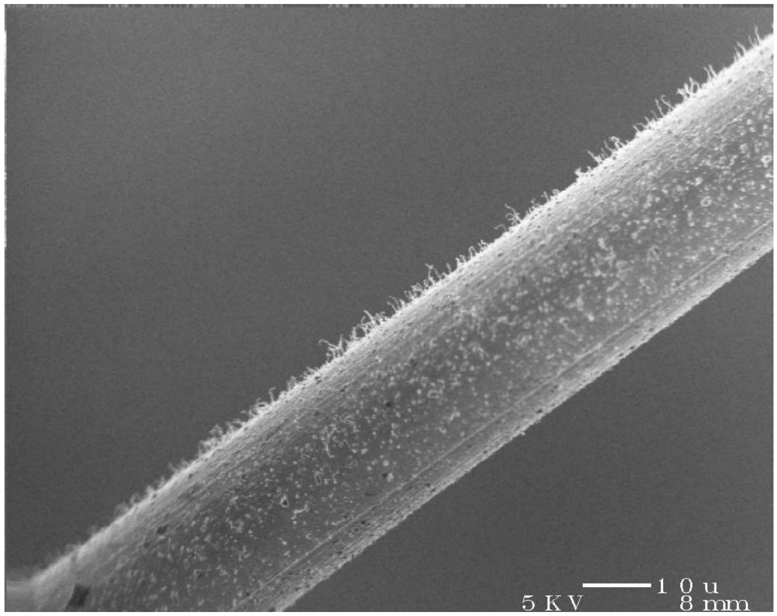

Figure 2.12 An electron microscope photograph of a sense wire removed from the top edge of a KTeV drift chamber after the 1996-7 run. This wire did not see as much radiation as a central wire, but has some deposits growing on it like fuzz.

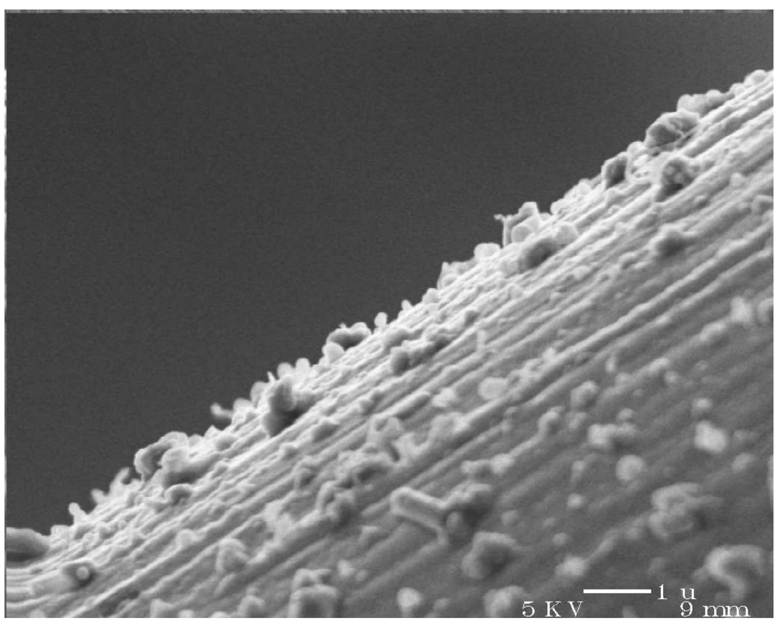

Figure 2.13 An electron microscope close-up of a sense wire removed from the top edge of a KTeV drift chamber after the 1996-7 run. This wire did not see as much radiation as a central wire, but still shows depositions.

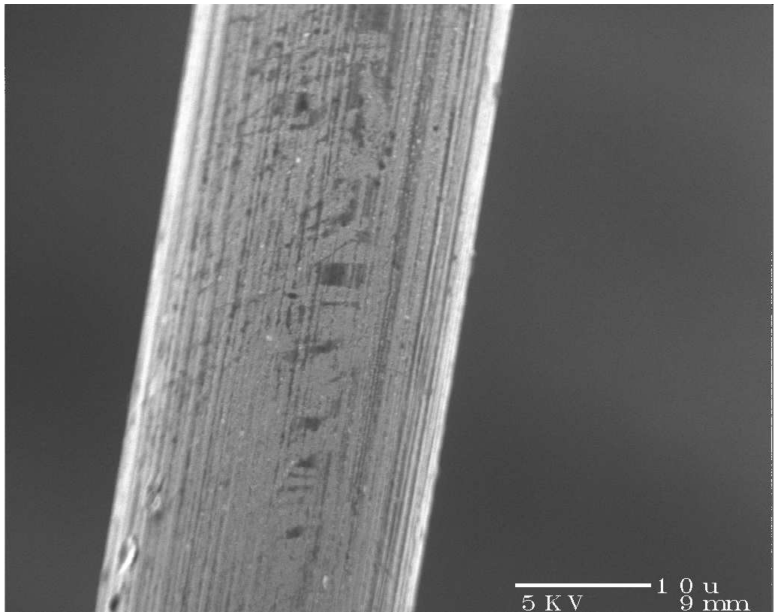

Figure 2.14 An electron microscope photograph of new sense wire taken directly off the factory spool. Notice the clear lines (grooves) from the wire being made by the drawing machine.

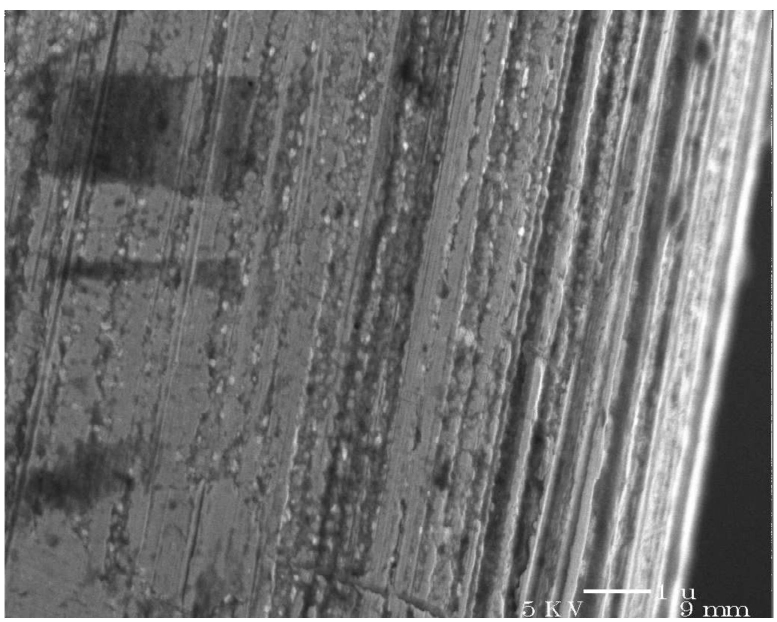

Figure 2.15 An electron microscope close-up of new sense wire taken directly off the factory spool. The grooves from the wire being drawn are crisp, deep, and unblurred by deposits.

angle approximations $l = r\theta$ and $p_T = p\sin\theta = p\theta$, so equation 2.6 becomes:

$$\frac{dp}{dt} = q\mathbf{B} = \frac{m\gamma\mathbf{v}\theta}{l} \tag{2.7}$$

and solving for p_T:

$$p_T = p\theta = q\int B\ dl \tag{2.8}$$

which is a constant. This bending of the charged particle's path allowed the measurement of the momentum and charge of the particle. The magnet consisted of an iron yoke with "race-track" copper coils wrapped around the rectangular poles [31]. The operational aperture of the magnet was ($2m \times 2m$) and the field integral uniformity was $\pm 5\%$ over the aperture. Typical transverse momentum kicks for the magnet in 1996-7 were $400 MeV/c$ for E832 and $200 MeV/c$ for E799II. In 1999 it was decided to lower the kick to $150 MeV/c$ for E799II to try to improve signal acceptance (by reducing losses from particles being bent completely out of the fiducial region of the detector). For this analysis the change may not have been advantageous; more accidental background activity was seen in our 1999 search, and that could have been a result of the lower setting.

Photon Vetos

If not all of the particles from an event are contained within the active part of the spectrometer, then information about the decay will be lost. To prevent useless information from wasting detector readout time or tape space, ring veto counters were placed along the decay region and around the spectrometer apertures. There were 5 ring counters (RC6-RC10) around the evacuated decay region whose outer circular edges varied from $0.75m$ to $1.2m$ in radius. The inner boundaries were rectangular and sized so that the veto counters would not interfere with the active regions of the spectrometer. Three more veto counters, called spectrometer anti's (SA2-SA4), surrounded each of the second, third, and fourth drift chambers. A fourth spectrometer anti, known as the Cesium Iodide Anti (CIA) surrounded the fiducial aperture just in front of the calorimeter. Both the inner and outer boundaries of these counters were rectangular with the inner aperture being the same dimensions as the active region of the corresponding drift chamber or calorimeter.

The veto counters were all constructed of a lead and scintillator sandwich structure with a single photon veto inefficiency better than 10^{-4} including both counter performance and gaps in the physical geometry. While this veto system did not completely cover all possible escape paths of decay particles, it did either hermetically cover kaon decay products with an angular region of less than 80 milliradians, or it forced the decay into a kinematically unfavored state that could be cut by software logic. All numbers cited about system inefficiency and coverage were determined by computer simulation[31].

2.2.4 Transition Radiation Detectors (TRDs)

Immediately following the fourth drift chamber were eight transition radiation detectors (TRD's). TRD's contain MWPC's but have an important additional element.

TRD's contain a layer of material with a high electron density. When a charged particle travels through the material, it emits transition radiation with energy proportional to $\ln(\gamma)$, where γ is the usual $1/(1\text{-}\beta^2)$ and β is velocity over the speed of light. Each KTeV TRD had a layer of polypropylene felt mounted on the upstream face. When relativistic electrons pass through the felt, they produce transition radiation X-rays. Muons and charged pions have lower velocities (due to their higher masses) and only generate ionization energy in the felt. The X-rays then enter the MWPC portion of TRD's, which in our case was filled with mixture of 80% xenon and 20% carbon dioxide. Such mixtures have a very short X-ray absorption length and so TRD's were very efficient at detecting the energy from electrons. As a result an electron has a very different TRD energy spectrum than a pion or muon. For example, a $30GeV$ electron has a γ of about 60,000 while a $30GeV$ pion has a γ of about 200. TRD's are therefore used mostly for particle identification and not tracking.

2.2.5 Trigger Hodoscopes V and V′

For triggering purposes, KTeV used two planes of hodoscopes (V and V′) consisting of 32 scintillating plastic counters each. Each counter was ($13cm$ wide $\times$ $1.10m$ tall $\times$ $5mm$ thick) and they were arranged to cover a cross-section of ($2m$ $\times$ $2m$). The counters were arranged in two rows (one above the other) of 16 counters each, with a seam running horizontally at the midplane of the detector and allowing for gaps around the neutral beams in the center. To ensure complete geometric coverage, one hodoscope was offset $5cm$ vertically and $6.5cm$ horizontally from the other. These hodoscopes served as the primary triggers and timing reference for the entire detector. V and V′ were located after the fourth drift chamber and $1.5m$ upstream of the calorimeter.

2.2.6 CsI Calorimeter

After the spectrometer, the particles pass into a cesium iodide (CsI) electromagnetic calorimeter so that the energies of the photons and electrons could be measured. Calorimeters serve the purpose of “energy collectors”. One kind of electromagnetic calorimeter is made up of blocks of dense, light conducting material which both produces and channels photons to the photo tube collectors. Photons are produced from charged particle scintillation or Cherenkov radiation. The E731 CP violation experiment used a lead glass calorimeter which produced Cherenkov light. KTeV instead chose blocks of pure CsI crystal for its excellent energy resolution, fast scintillation light pulses, and radiation hardness. In many CsI calorimeters, the CsI is doped with thallium to increase the light output. However, thallium produces a slower light pulse with a much longer tail, creating pile-up of pulses and loss of timing information. KTeV’s high event rate meant we were better off having pure CsI to keep the pulses fast and to avoid pile-up rather than doping to increase light yield. To further increase the pulse separation and improve timing, a UG-11 filter was placed on the front of every calorimeter photo tube. UG-11 filters only allow the passage of light in the $300nm$ to $400nm$ range. These wavelengths are present near the peak of a CsI light pulse and not out in the tail, so the pulse was additionally shortened in time. Our calorimeter had an energy resolution ($\frac{\delta E}{E}$) of better than 1% at $10GeV$[32]. In comparison, E731's lead glass calorimeter resolution was 2.8% at $10GeV$.

Even given the resistance of CsI to radiation damage, allowance was made in the center of the calorimeter for the majority of the beam debris (a barrage of neutrons and undecayed kaons) to miss the calorimeter and instead be absorbed by the steel wall behind it (see Figure 2.16). The holes were each ($15cm \times 15cm$) and were separated by $15cm$ [31]. The blocks in the middle of the calorimeter numbered 2,232 and measured ($2.5cm \times 2.5cm \times 50cm$), while the 868 blocks around the outside measured ($5cm \times 5cm \times 50cm$) for a total of 3,100 blocks. The lateral dimensions of the blocks were chosen to balance the spatial resolution of the calorimeter with its construction cost. Data readout for the calorimeter was done in 34 slices of 20 nanoseconds each, although only 4 or 6 slices (including one pre-slice) were saved to tape per event. A *pre-slice* was the information taken just before the event should have reached the CsI according to the trigger hodoscope timing and chamber hits. The pre-slice was kept to study the effect of residual energy from an earlier hit overlapping with energy from an in-time hit.

Most of the energy deposited in our calorimeter came from photons and electrons. The two dominant processes for energy loss by electrons and photons are

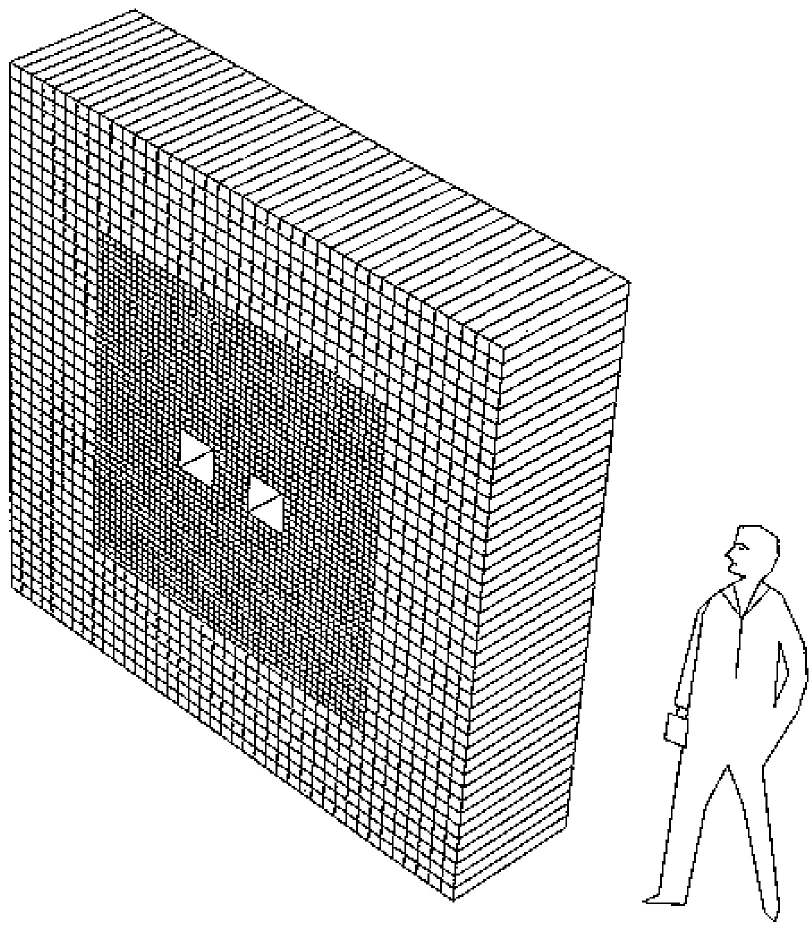

Figure 2.16 Sketch of the CsI electromagnetic calorimeter showing approximate size [31]. The actual calorimeter has one fewer row of crystals around the outside. The two holes in the center allow undecayed beam to pass through the calorimeter without damaging the crystals.

bremsstrahlung radiation and pair production, respectively. Bremsstrahlung radiation involves momentum loss via photon production as the electron travels through dense material. In pair production a photon converts into a pair of electrons in the field of a heavy nucleus. The combined result of these processes is a symmetric cone of energy distribution with a Molier radius of $3.8cm$ for CsI. The *Molier radius* is the lateral radius around the center of the energy distribution which contains 90% of the energy deposited. The distance that an electron travels before its energy drops to $\frac{1}{e}$ of its original value is called a radiation length (X_0), and is characteristic of the material through which the electron travels. The same length multiplied by $\frac{9}{7}$ gives the distance in the material over which the number of photons surviving the journey is $\frac{1}{e}$ of the original number. In short, an electron or photon will lose most of its energy upon traversing about 10 radiation lengths. We chose $50cm$ long pieces of CsI, which translates to 27 radiation lengths. This ensured that all of the electrons and photons in an event deposited all of their energy in our calorimeter. Total energy combined with the energy distribution pattern gave a unique and identifiable electron/photon signature. The lower two clusters of energy in Figure 2.17 show a typical electron/photon signature as detected by KTeV's CsI calorimeter.

In contrast, all other particles in an event, namely pions and muons, often deposited only a minimal amount of energy in the calorimeter. These particles do not radiate easily like electrons, but do ionize the atoms of the material. Pions (which are

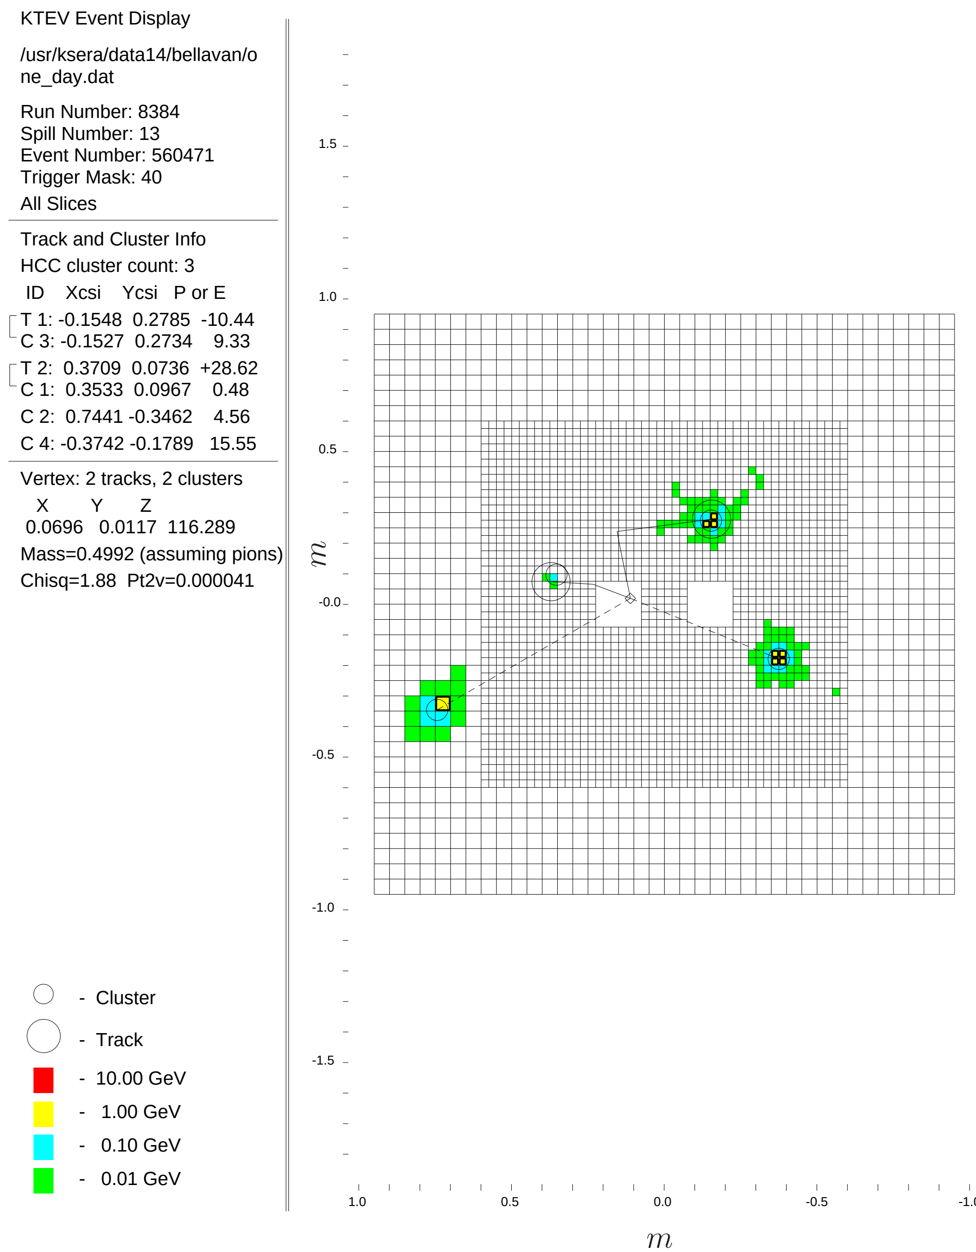

Figure 2.17 Energy deposition in the KTeV CsI calorimeter. The deposition signatures shown here could come from four kinds of particles. On the upper left is a cluster characteristic of a muon or minimum ionizing pion. The upper right cluster is characteristic of a strongly-interacting pion. The two lower signatures are characteristic of an electron or a photon (photon in this case). This is a real $K_L \to \pi^+\pi^-\pi^0$ (K3pi) data event.

hadrons) can also interact strongly with nuclei in the material. Thus the counterpart to the radiation length for leptons is an interaction length (X_I) for hadrons, which is the distance over which 1/e of the particles will have interacted. The probability of collision with a nucleus is so small that interaction lengths are generally much longer than radiation lengths. In this case, the interaction lengths for all remaining particles were greater than the $50cm$ length of the crystals. A hadron could interact strongly to produce an energy shower and/or daughter particles, but it would typically not deposit all of its energy in less than an interaction length and any showers produced tended to be asymmetric. The signature of a pion cluster therefore had a wide range of possibilities. Figure 2.17 shows examples of energy deposition by charged pions in CsI. In the upper left cluster, the pion has left only a trail of minimum ionization in one crystal. In the upper right cluster, a pion's strong interaction with a nucleus is visible as ray-like protrusions of debris coming off of the center cluster of energy. Finally, the pion can sometimes deposit energy in a fairly even pattern, mimicking the signature of an electron or photon (the lower two clusters). Muons interact only electromagnetically and because of their relatively large mass they do not bremsstrahlung at the energies they had in our experiment. Hence their only mechanism of energy deposition in the CsI crystal was minimum ionization, as seen in the upper left cluster of Figure 2.17. Very rarely the energy deposition from a muon could have been wide enough that it looked like a photon or electron cluster.

In order to distinguish between real electrons and a faking pion or muon, the parameter E/p (energy deposited in the CsI divided by the momentum of the charged track) was used. This parameter is described more throughly in Section 3.4.

2.2.7 Collar-Anti, Hadron-Anti and Back-Anti

The Collar-Anti (CA), Hadron-Anti (HA), and Back-Anti (BA) were additional veto detectors placed directly in front of (CA) and behind (HA and BA) the calorimeter. The CA was similar to the RCs and SAs in that it was ring-shaped and covered the inside edge of the beam holes. It was used to veto events with particles that would just scrape the calorimeter crystals around the beam holes. The HA stood downstream of, and covers approximately the same cross-sectional area as, the calorimeter. It's purpose was to detect particles which made it through the calorimeter. These may have been minimum ionizing particles (muons or pions), or pion showers which were not contained in the calorimeter. The BA filled the beam holes behind the calorimeter and completely contained the undecayed neutral beams laterally. It had a cross section of ($30cm \times 60cm$) and was divided into three equal longitudinal sections: two upstream "electro-magnetic" (EM1 and EM2) sections and a downstream "hadronic" section. Each section was a stack of absorber plates and scintillator strips 10 radiation lengths (X_0) deep in 20 layers. Each absorber plate was $0.26cm$ of lead sandwiched between two sheets of 304 stainless steel $0.025cm$ thick. Each scintillator layer was $0.25cm$ thick plastic[33].

The CA was applied as a veto in most trigger settings, including ours. The HA was not in veto for the flux normalization trigger of this analysis, but it was in veto at a high energy threshold (non-minimum-ionizing particles) for the search trigger. The BA was not used in any triggers.

2.2.8 Muon Counters

Downstream of the Back-Anti was a muon filter consisting of approximately $10cm$ of lead ($1.6X_0$, $0.05X_I$) and $5m$ of steel ($36X_0$, $3.7X_I$). This filter served as a beam dump for the neutral beam, and also stopped most charged particles. Only muons should have been able to make it past the lead and steel into the following muon counter detectors. The muon counter detectors were three groups (or "banks") of scintillator labeled Mu2, Mu3X, and Mu3Y. Each bank had $15cm$ wide paddles that were $1.5m$ in length. Mu2 had 56 scintillators segmented along the X-axis so that the paddles stood vertically in two rows of 28 each. Mu3X had 40 scintillators also segmented along the X-axis in two rows of 20. Mu3Y had 40 scintillators segmented along the Y-axis in two stacks of 20.

2.3 Triggers

The most common type of decays were often not wanted and needed to be kept out of the stored data set. The multitude of data sent by the detector components passed through at least four levels of triggers before being saved to tape.

2.3.1 Hardware Triggers

The first three trigger levels (levels 0 through 2) were implemented entirely by hardware logic modules, some of which were located below ground in the experimental hall and the rest of which were located in a ground-floor counting room.

Level 0

Trigger level 0's digital and analog electronics were all physically housed in the downstairs KTeV experimental hall. Multiple copies of the analog signals from each detector were made with at least one copy being sent to the upstairs counting room for recording purposes. The general purpose of level 0 trigger logic was to take the analog signals, discriminate them, and provide at least one level 1 source signal per piece of detector equipment.

The simplest logic was performed on the HA and the MU banks. Analog signals were simply summed by section and discriminated to generate level 1 source signals. For the photon veto detectors (MA, RC, SA, CIA) the analog signals from modules which might share energy were linearly summed, the results discriminated, and a logical OR taken on a detector-by-detector basis (e.g. there were 5 level 1 source signals for the RC's, one each for RC6 through RC10).

Only drift chambers 1 and 2 were used in the level 0 trigger. At this stage each chamber was represented as a collection of "paddles" that were eight and a half cells wide and contained 16 sense wires each, eight from the unprimed view and the corresponding eight from the primed view of the same dimension (see Section 2.2.3 for the definition of the chamber views). If one of these "paddles" was hit, a 90$nsec$ NIM output pulse was produced. A NIM is a Nuclear Instrumentation Module which conforms to industry standards as defined in AEC Report Number TID-20893 [34]. The outputs of a plane were then discriminated to search for a combination of "in-time" hits, with "in-time" being defined as having a TDC value between 400 and 650$nsec$ as in Figure 2.9. This custom discriminator gave outputs for number of hits greater than 0, 1, 2, and 3, for a total of four level 1 source signals per chamber.

The most complicated level 0 logic involved the calorimeter and the trigger hodoscopes. Four different threshold signals for the calorimeter's total energy sum were sent as level 1 source signals. In addition, two threshold signals were generated per calorimeter quadrant's energy total. Signals from individual trigger hodoscope paddles were sent to LeCroy Research Systems' LRS2365 and LRS4508 Programmable Lookup Modules which distinguished between hit combinations that indicate a valid decay in the detector and hit combinations which were unphysical.

Level 1

Some signals sent by the Level 0 trigger stayed active longer than the RF bucket that contained the interesting decay, so the level 1 trigger required coincidence with the beam's RF 19$nsec$ spill signal to create a synchronous bucket-by-bucket system.

Coincidence was made between the trigger hodoscopes and the calorimeter at this level. The level 1 source signals were divided into 24 groups of 8 and each group sent to a look-up table. Each look-up table produced an 8-bit output word. These outputs were compared in two groups of 12 to produce two 8-bit words. The 8-bit words were combined into one 16-bit word that was called the “Level 1 Trigger Mask”. Fifteen of those bits told which of the fifteen different kinds of triggers were satisfied during a particular bucket. The sixteenth bit told if the other fifteen were “beam triggers” or “calibration triggers” and determined how they were handled by the level 2 system. “Calibration triggers” were used to test various detector components or to test trigger hardware and were always accepted. “Beam triggers” contained the physics information during normal operating conditions and went through the entire gamut of acceptance decisions. The trigger mask bits were ORed with each other and with a “busy with previous bucket” from level 2 to generate the level 1 trigger signal (level 2 source signal).

Level 2

The level 2 trigger consisted of trigger processing, trigger control, and readout control. The trigger processing section received a trigger mask from level 1 and received data from the level 2 processors (both commercial programmable boards and custom boards). It then either generated a level 2 trigger mask or aborted the event. The trigger control section vetoed the level 1 signal if there was a BUSY signal coming from the level 2 processors or from level 3. The readout control section created the necessary control signals for reading data from an accepted event into level 3.

Dead time ($\approx 35\%$) first entered at the level 2 trigger because the trigger processing plus detector read-out takes an average of $20\mu sec$, and the acceptance decision rate drops from hundreds of kiloHertz to ten kiloHertz. The main processors at this level were in five groups. Processors in the first group were custom designed VME boards nicknamed "bananas" and "kumquats" that counted the number of hits in the drift chamber signals. Second was the custom-built Hardware Cluster Counter or HCC. The third group did y-track finding in the drift chambers using commercial programmable-logic and memory-lookup modules. The fourth and fifth groups generated information for the stiff-track and TRD triggers. Since the trigger for this analysis (Trigger 7) just used information from the HCC and bananas and kumquats, only they are described in detail here.

Calorimeter energy clusters were classified as either "hardware" (hard) clusters or "software" (soft) clusters. To qualify as a *hard* cluster, the energy pattern must have been identifiable by the Hardware Cluster Counter (HCC) logic at trigger level 2, or in other words it must have been identifiable by KTeV's hardware. A *soft* cluster was an energy pattern that was missed by the HCC, but which, with software calculations and corrections, could be resolved. The HCC dealt with 3,100 bits of information, one signal for each calorimeter crystal being above (on) or below (off) a preset threshold (nominally $1GeV$). The HCC also had strict timing requirements that the bit signals be accepted only during the event's in-time trigger window, so hard clusters were always in-time while soft clusters had no such requirement. A Column Alignment Buffer (CAB) recreated the crystals' stacking pattern and a Cluster Counter Unit (CCU) outputted the number of clusters it found as either 0, 1, 2, 3, 4, 5, 6, 7, or 8 or more. This output was the number of hardware clusters. Various triggers could then require different numbers of hardware clusters. Our search trigger (#7) required 3 or more hardware clusters.

Also done at level 2 was the analysis of the hits in each view of the drift chambers by processors known as "bananas" and "kumquats". The banana processors were so named because of the shape of the time distribution plot they produced. It then became a tradition to name all parts of the same processing system after fruit. The total number of hits was given a value of 0, 1, 2, 3, 4, 5, 6, or 7 or more per view. The logic for a "hit" was such that signals on isolated wires counted as one "hit" and signals on N adjacent wires counted as N-1 "hits." This way a pair of adjacent wires that gave a signal only counted as one "hit."

Due to the multiple processors running in parallel, there was a level 2 controller module to assure the processes integrate smoothly. The level 2 controller received as input the level 2 trigger mask (beam or calibration), signals from each level 2 processor, a BUSY signal from some of the electronics, a RESET signal from the data acquisition system (level 3), and Almost Full signals from level 3 processors. If the trigger mask from level 1 was a "calibration trigger", then the resulting level 2 mask was identical. If the mask was a "beam trigger", then the level 2 output mask was the logical AND of the level 1 beam trigger and the triggers which pass the level 2 processors. In addition to the trigger mask, the readout control section of level 2 generated a sparsification mask that told level 3 what data was to be read out and to which level 3 processors each type of information was to be sent.

2.3.2 Software Triggers

Level 3 was where the trigger logic moved from hardware to software. The final decision of whether the event got written to tape or not was made at this level. Data was read into the VME computer memories to await processing by CPUs. The VMEs communicated with the CPUs through interface cards custom built by Performance Technologies. A total of 24, $200MHz$ processors and 10, $150MHz$ processors ran in parallel on four SGI Challenge servers. The slower processors handled calibration triggers and detector monitoring while the faster processors handled the physics data. Each physics event was sent through software which attempted to trace charged particle tracks in the drift chambers, identify corresponding energy deposition clusters in the calorimeter, and assign a particle type to each coherent trace. Please see Section 2.4.2 for a more complete description of the level 3 trigger.

2.4 Software

As KTeV was going through the planning stages, collaborators wrote and adapted code (mostly from E731, a previous direct CP violation search at Fermilab) for detector simulation, on-line triggers, and off-line event analysis. A Monte Carlo program was needed that would simulate a particle traveling through the various parts of the detector and would simulate how those components would react and/or effect the particle. An analyzing program was also needed to look at both the results of the Monte Carlo program and the results from the run itself. The end product was the programs KTeVMC and KTeVAna.

2.4.1 Monte Carlo simulation program (KTeVMC)

The Monte Carlo (MC) code was written in FORTRAN and made use of a histogram booking package called HBook. It contained detailed information about the structure of the experimental setup and required an input file containing information on what type of decay to simulate. The program stepped each particle through the length of the detector, decided what happened to it at each point, and recorded the detectors' corresponding reactions. Detector material was treated as planes of radiation and interaction lengths in Z and charged particles were allowed to undergo multiple scattering. In addition to scattering, electrons could radiate bremsstrahlung photons according to the Bethe-Heitler cross-section,with the created photon having the same trajectory as the parent electron. Photon conversion was also simulated with a probability of 1 - $e^{-\frac{7}{9}(X/X_0)}$ (where X/X_0 is fraction of radiation length), with a spectrum taken from the Bethe-Heilter formula, and with the opening angle distribution of the electron-positron pair taken from routines in the EGS4 software library.In the calorimeter, particle interaction depended on the particle type. Muons were traced through the calorimeter and a minimum ionizing energy deposit was calculated. Pions were simulated by selecting from a data-generated shower library. Photons and electrons showers were selected from a GEANT-generated shower library created previously. The GEANT showers were smeared to mimic the data resolution.

An event other than the desired event may have deposited information in the detector at the same time. Hence one of the settings in the Monte Carlo program could overlay this extraneous information, called “accidentals”, on the clean event data. The accidental events themselves were collected from data using an accidental trigger (Trigger #14, see Table 2.1). The accidental trigger required at least one hit in the DC Y-views, some minimum energy in the calorimeter, and a hit in either the beam monitor or a muon bank. Detector information from accidentals were added to the MC generated event on a component-by-component basis before threshold and trigger acceptance effects were applied. This was important since it allowed MC background events that would have otherwise been rejected to pass into the triggered data set.

Output from the program included both a data output file which resembled raw data, and a text output (to the screen) of the steps and their results as they were completed. The text output was re-routed into a log file and used as a quick reference for that simulation. The data in the output file was in the same format as data taken during the actual run, so both could be used as input for the analysis program KTeVAna.

2.4.2 Online software triggers (Level 3 trigger)

Early versions of the analysis program were the proto-type for level 3 triggers in the real experiment. All user interface with the trigger was done at this level. Code for many of the basic routines were written by various collaborators or adopted from older experiments and put into a shared library.

The level 3 software was in charge of attempting a preliminary event reconstruction. It searched for charged tracks in the spectrometer and combined this with calorimeter information from the level 2 Hardware Cluster Counter (HCC). Matches were made between a charged track and the nearest calorimeter cluster that it pointed to. As stated before, the HCC timing considerations assured that a hard cluster was in the same bucket as the triggering event, while a soft cluster had no such requirement. Hence all of our signal particles that would leave more than $1GeV$ in the calorimeter, namely photons and electrons, should have each had a matching hard cluster in the corresponding trigger to ensure they were part of a single kaon decay event. Electrons should have also had a matching charged track. For muons, since they deposited little energy in the calorimeter, the only requirement was that the charged track match to any kind of cluster, hard or soft. Charged pions looked like electrons to the level 3 trigger since they deposited hard clusters most of the time and had an associated track. The level 3 software also decided which triggers were enabled for the selected running configuration, and sent control signals to the level 2 hardware about when and where data was to be written to tape.

There were up to sixteen different trigger streams online as supplied by level 2. Table 2.1 contains the trigger numbers and a brief description of each beam trigger used in E799II data taking. The trigger logic for the signal mode $K_L \to \pi^0 \mu e$ (Kpi0me) was called "Trigger 7" or the "EMU trigger". Trigger 2 (or "2-track trigger") events were used for flux normalization.

Since this analysis focused on Trigger 2 and Trigger 7, their requirements are described in more detail here. The more complicated of the two triggers was Trigger 7 and it's acceptance was determined by the logical AND of the following requirements:

- ∘ simple begin readout signal [GATE]
- ∘ 2 hits (allowed for one lost hit) in VV′ [2V]
- ∘ total energy in the calorimeter passed the next to lowest threshold (= $18 GeV$) [ET_THR2]
- ∘ at least 1 hit (in X or Y of) both DC1 and DC2 [DC12]
- ∘ no HA veto (threshold = 7 MIPs) [!HA_DC]
- ∘ 1 or more hits in either X or Y views of MU3 [1MU3]
- ∘ no photon vetos (in RCs or SAs, including SA3) [PHVBAR1]
- ∘ no CA veto [!CA]
- ∘ "loose" 2 hits for all DC Y-views (allowed 1 lost hit in DC1 or DC2) [2HCY_LOOSE]

#	Label	Prescale	Description of event requirements
1	2E-NCLUS	(1)	4+ clusters
2	2TRK	500	2 tracks
3	DALITZ	5	3 or 4 clusters; 1 electron in TRDs; no HA, CA, or photon vetos
4	4TRK	(1)	4 tracks
5	2MU-LD	(1)	2 tracks; 1+ clusters; no photon vetos; 2+ muons
6	2MU-LOOSE	50	2 tracks; 1+ clusters;
7	EMU	1	2 tracks; 3+ clusters; 1+ muons
8	NEUTRAL	40	5+ clusters; no HA, CA or photon vetos
9	2V-MINBI	20000	Simple activity in VV', DC1 or DC2, and CsI
10	HYPERON	1	1 hole track; 1+ clusters; no CA or photon vetos
11	LAMBDA-PPI	50	1 hole track; no photon vetos
12	HYP-MINBI	10000	2+ clusters; no photon vetos
13	ET-MINBI	20000	Energy deposited in CsI calorimeter
14	ACCID	25000	Activity in a chamber Y-view, in calorimeter, and in the beam monitor or any muon bank
15	VETO	(1)	No activity in HA, CA, or photon vetos (no gate)
16	(not used)	(25000)	(not used)

Table 2.1 The various beam triggers used in experiment 799II for data taking in 1997, and a short description of each. A number with a plus sign (+) after it indicates a value greater than or equal to that number. "Tracks" refers to charged tracks in the spectrometer and a "hole track" is a track that projected to one of the beam holes in the calorimeter. "Clusters" refers to hardware clusters in the calorimeter. Unless stated otherwise, each trigger also had an in-time gate requirement. This analysis used triggers 2 and 7, and the 1999 configuration of triggers 2 and 7 were exactly the same as for 1997. Trigger 14 was used to collect "accidental" events for use in MC generation, and it also did not change between runs.

- ∘ 3 or more (hardware) clusters in the calorimeter [HCC_GE3]

In sum, as stated in Table 2.1, the trigger for our signal event required 2 charged tracks, 3 or more hardware clusters, and 1 or more hits in the muon detector banks. Trigger 2 was much looser, requiring only the logical AND of the GATE, 2V, DC12, and 2HCY_LOOSE, plus at least 2 hits in any chamber's X-view and a minimum energy in the calorimeter.

2.4.3 Analysis Program (KTeVAna)

The data output files from KTeVMC and real data were used as input for the KTeVAna analysis program[35]. As one of its outputs, KTeVAna produced a file that included histograms written in the HBook format and a table of parameter values called an *ntuple*. These two information formats are readable by a program called the Physics Analysis Workstation (PAW) which creates graphic displays of the events' information and allows additional cutting of events based on their parameter values. KTeVAna could also produce a file of events which passed the analysis that was written in the same format as the input data file, and also a text output (to the screen) of the steps taken and their results. Again, the text output was re-directed into a file and used as a quick reference to an analysis.

Event display program "kdisp"

A program called "KDISP" could be used to read in a raw data file and display specific events in the format shown in Figure 2.18. The event shown happens to be a Kpi0me MC event. KDISP did enough analysis to identify tracks in the drift chambers, match them to clusters of energy deposition in the calorimeter, and find a preliminary decay vertex. The tracks and clusters were listed numerically (under "Track and Cluster Info" in the figure), both were projected onto the two-dimensional image of the calorimeter, and the tracks were overlaid on a side and top view of the entire detector. Only two of the four tracks in the lower two plots of Figure 2.18 (solid lines) were actually detected in the drift chambers. KDISP assumed these were from charged pions in the rest of its calculations. The straight tracks (dashed lines) were drawn from clusters not having corresponding tracks to the reconstructed charged vertex of the event and were assumed to be photons.

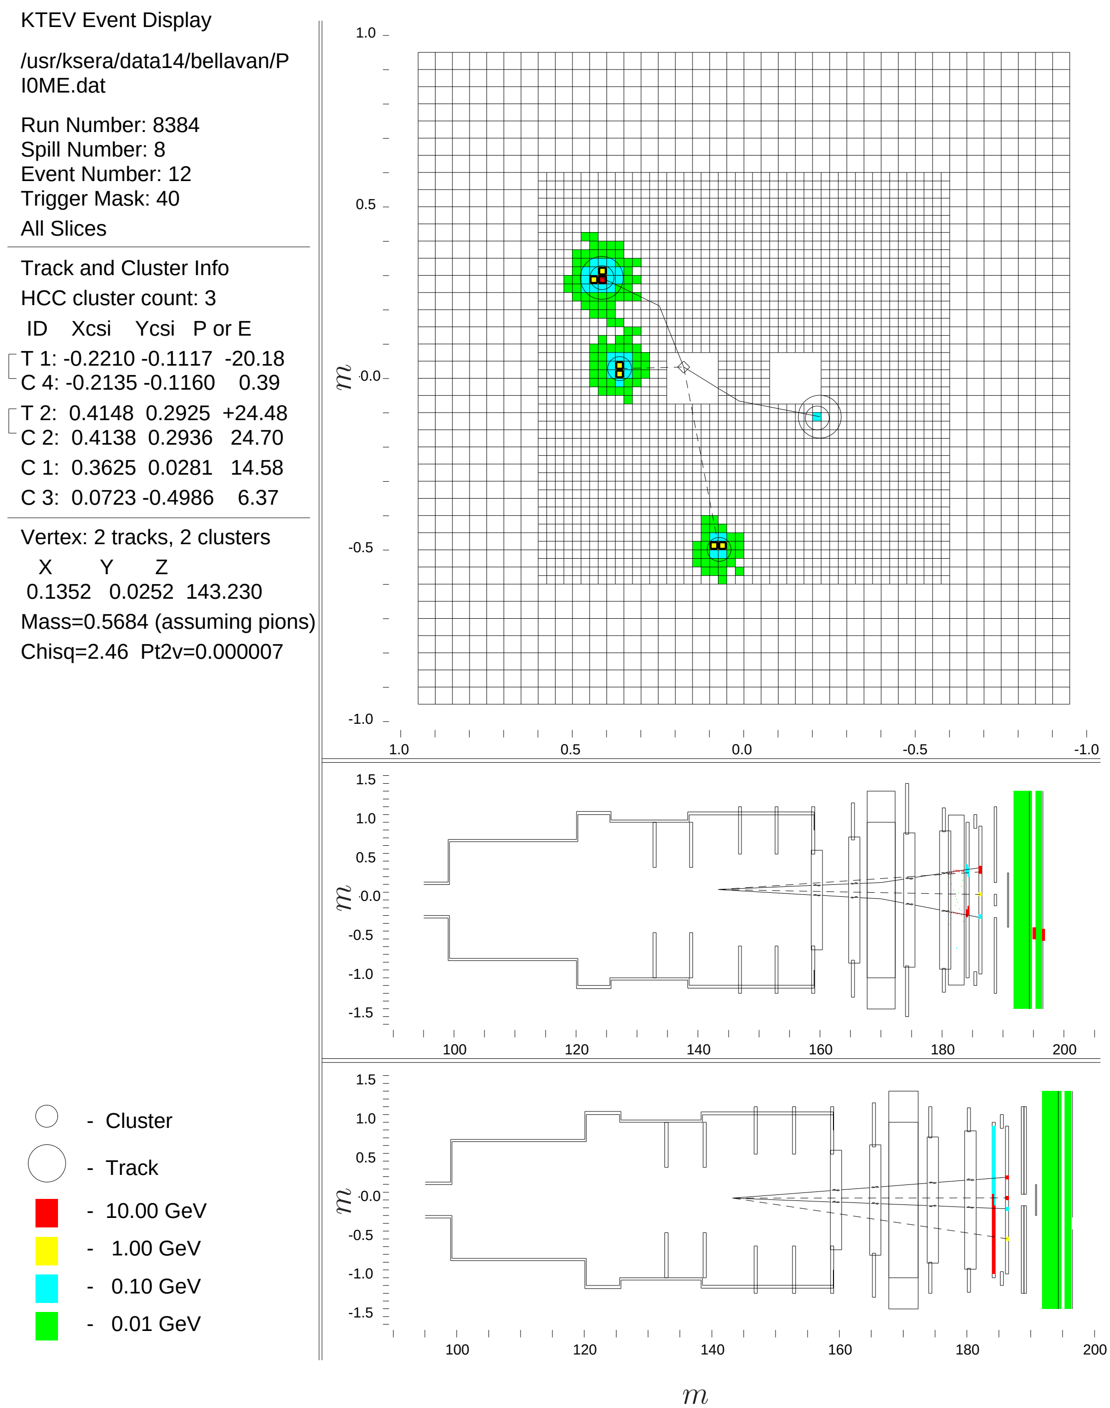

Figure 2.18 Example event (Kpi0me MC) display using "kdisp".

Chapter 3
Analysis Overview

After data was collected using the "Trigger 2" or "Trigger 7" online requirements (see Section 2.4.2), three levels of analysis were done. The parameter cuts for the first two levels were relatively loose. Their purpose was to shrink the number of events down to a manageable size. The first level of cuts was to reduce, or "crunch", the amount of data from more than 100 DLT tapes to 19 DLT tapes (6 for 1997, 13 for 1999), and hence will be called the *crunch cuts*. The second level of cuts will be called *ntuple cuts* because they were made in going from the raw data to filling a PAW ntuple. The third level of cuts were made upon events that reached the ntuple level and consisted of applying a selection file within PAW. Hence this level of cuts will be called *selection cuts*. Details of each level are given below.

It is important to emphasize that we have performed a "blind analysis" (see Section 3.2) on the data using Monte Carlo (MC) simulations of the signal decay, the normalization decay, and the background decays. Possible signal events are removed from the studied data while the analysis evolves to avoid letting human bias affect the cut placement.

Over time since the end of the 1997 and 1999 data-taking runs, there have been periodic updates of the calibration constants used to make the MC detector simulation behave like the real detector, and the analysis software itself has been updated as the performance of the detector is better understood and analysis methods are improved. The plots presented and numbers quoted here were generated using program Version 5.05 of KTeVMC and KTeVAna, unless otherwise noted.

3.1 Analysis Strategy

Since the objective of this study was the search for $K_L \to \pi^0 \mu e$ (Kpi0me), cuts were optimized to find a well-reconstructed Kpi0me decay. Requirements were made on the quality of the data so that all the information needed for a full reconstruction was present. A Kpi0me decay would leave two charged tracks in the spectrometer (the muon and the electron) and a pattern of four or more clusters in the calorimeter (see Figure 2.18). Events with more than two tracks were discarded. The two tracks were each required to match to a cluster in the calorimeter. Three of the clusters were required to be hard clusters (see Section 2.4.2) indicating they were from in-time electrons or photons. Since the electron should have deposited most of its energy in the calorimeter, one of the hardware clusters was required to match to a charged track and this combination was tagged as the "electron". The soft cluster / track combination was tagged as the "muon" since the muon was expected to minimum ionize and deposit less than $1GeV$ of energy ($1GeV$ is the hard cluster threshold). Therefore by definition the muon cluster had less than $1GeV$ of energy. The muon was then required to have a minimum momentum ($10GeV$ for 1997, $7GeV$ for 1999) so that the track was stiff enough to have the opportunity to reach the muon banks without excessive multiple scattering. Additional kinematic variables (described later) were then used to double check that the tagged electron really behaved like an electron, and that the tagged muon really behaved like a muon. The two charged tracks were used to calculate the decay vertex of the event, which was required to lie inside the evacuated decay region. Additional quality cuts (also described later) were made on the charged tracks and the vertex.

Two clusters in the calorimeter not corresponding to charged tracks were assumed to be photons from the immediate decay of the neutral pion. Using the vertex found from the charged tracks, the two photons were reconstructed as having come from one particle decaying at that vertex and that particle's mass must be close to the mass of a neutral pion. If there were more than two clusters not matching to charged tracks, then the combination that resulted in a mass closest to the neutral pion mass was selected.

Eventually the energies and momenta of all four particles were combined to find a parent particle rest mass that hopefully agreed with the mass of the neutral kaon (see Figure 3.1). The K_L mass is $0.49767 GeV$[8] and an ideal kaon decay would have a transverse momentum squared (p_T^2) of zero. To be considered "good" (within the red box in Figure 3.1), the reconstructed event must have had m_K (called `K_mass`) greater than $0.48767 GeV$ and less than $0.50767 GeV$ (a deviation of $10 MeV$ from the nominal kaon mass), and p_T^2 (called `pt2`) less than $0.00025 \frac{GeV^2}{c^2}$. The resolution of our `K_mass` peak was different for the signal mode versus the normalization mode (see Table 3.1), but our signal acceptance box ends up being about $\pm 4.5\sigma$.

Event sample	`K_mass` Gaussian σ	$10MeV$ in `K_mass`
'97 normalization data	$2.11MeV$	4.74σ
'97 normalization MC	$1.96MeV$	5.10σ
'97 signal MC	$3.24MeV$	3.09σ
'99 normalization data	$1.80MeV$	5.56σ
'99 normalization MC	$1.84MeV$	5.43σ
'99 signal MC	$2.98MeV$	3.36σ

Table 3.1 Resolutions of the `K_mass` variable in terms of the sigmas (σ) of Gaussian fits to various event samples. The signal acceptance box width of $\pm 10MeV$ averages to about $\pm 4.5\sigma$.

Figure 3.2 shows a contour plot of signal MC zoomed inside the red acceptance box. The majority of events are clustered around the kaon mass and at a `pt2` of zero, as expected. In retrospect, it might have been better to select a triangular acceptance region instead of a rectangular one.

3.2 Blind Analysis

The availability of KTeVMC code allowed the majority of the analysis development to be done using Monte Carlo (MC) generated signal and background events. Doing so meant that a "blind analysis" could be done with a high level of confidence because detailed comparisons could be made between the MC events and the data events. A *blind analysis* is one in which the majority of the study is done using a

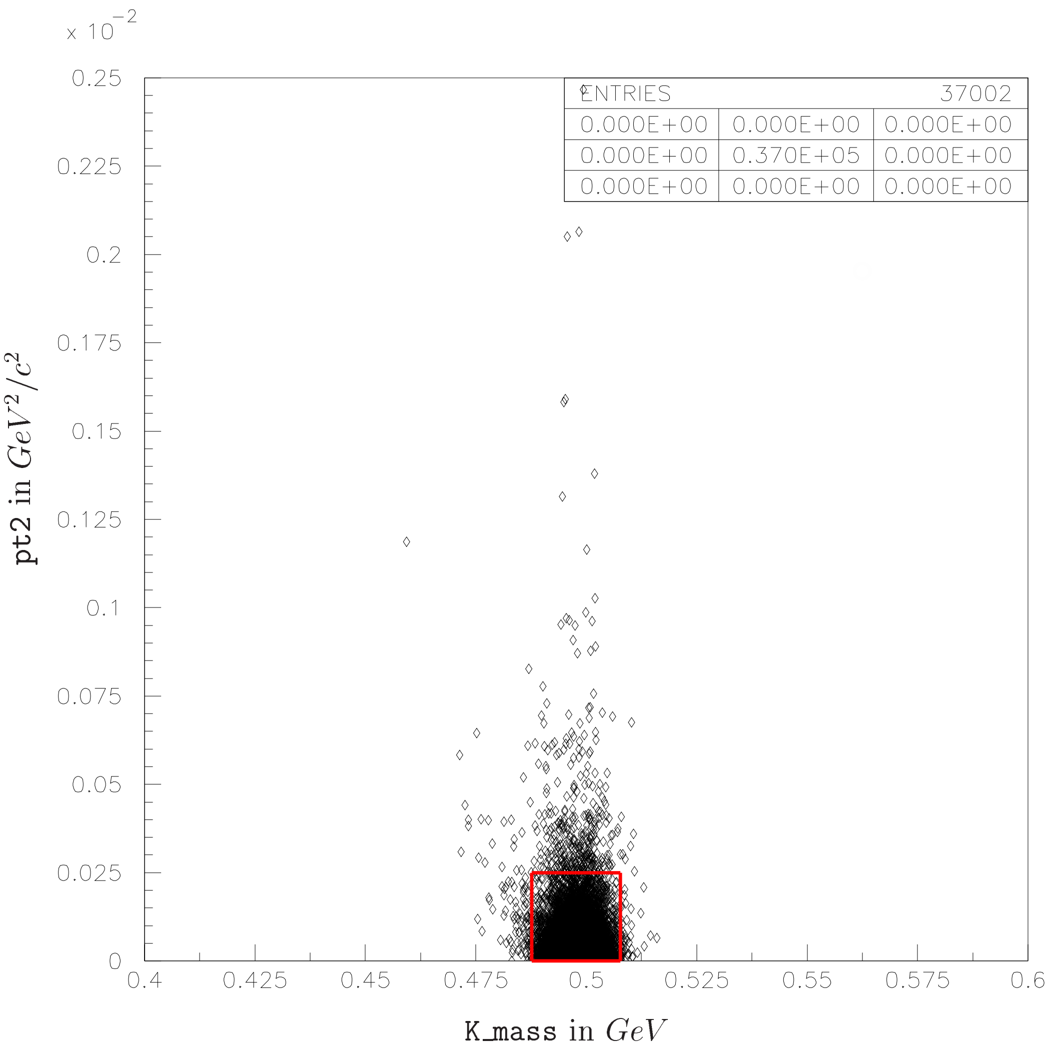

Figure 3.1 Signal Monte Carlo (MC) events for the 1999 data set that passed all cuts except those on the kaon mass (m_K, on the x-axis in GeV) and the transverse momentum squared (p_T^2, on the y-axis in $(\frac{GeV^2}{c^2})$). Events that also passed the m_K and p_T^2 requirements (i.e. are inside the red "box") were considered "good events".

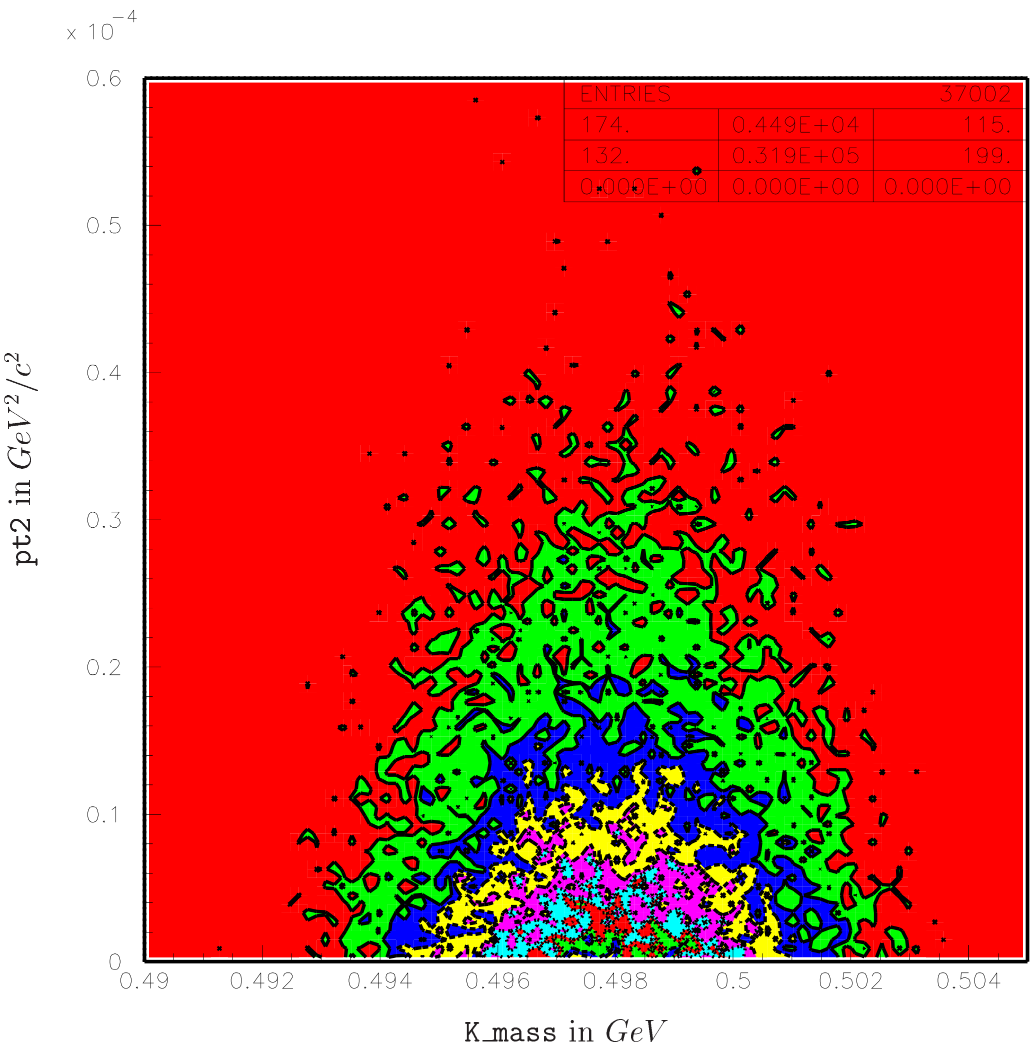

Figure 3.2 Contour distribution of signal MC events for the 1999 data set that passed all cuts, including the signal box. The limits on the axes are actually tighter than the signal box, in order to better show the shape of the distribution. As expected, the majority of events had both a `pt2` value near zero and a `K_mass` value near the mass of the kaon. The contour divisions are each 10%.

data set from which possible signal events have been removed. The advantage of such an analysis is that the experimenter is prevented from artificially biasing the results - an important precaution in a search analysis. If the experimenter believes that no events should be found, it is possible to select cuts that eliminate all the events seen where the signal events are expected to be seen. At the other end of the spectrum, if the experimenter is eager to discover a signal then cuts could be selected that make the data appear to cluster around where the signal events should cluster. However, if one does not see potential signal events until after the cuts are determined then such games cannot be played. A blind analysis *signal region* is defined by putting limits on a few parameters so that events falling within those limits have a reasonable chance of being valid signal events. Since a signal region is often defined by two parameters, it is commonly referred to as the *signal box* or simply the *box*. Typical parameters for defining the signal box include the parent particle's mass, sum of the daughters' transverse momentum, the mass of a daughter (or intermediate decay) particle, and the parent particle's lifetime. The signal box for this analysis was defined by cuts on the parent particle's mass and its transverse momentum squared (see the red outline in Figure 3.1). Any events in the experimental data that fell within the box were masked out of the data set while the analysis evolved. To reflect this, preliminary plots of the data had the signal box blackened in, or *closed*. For MC event sets it was already known what kind of events were present and so keeping the box closed was unnecessary. Once the analysis process and all cut values were finalized using MC to data comparisons, the experimental data set was analyzed one last time with no event masking, the box was opened, and the results interpreted.

A subset of the '97 data was used to perform a blind analysis of the full '97 data set. That analysis was considered completed in 2000, and the box was opened. The results of that analysis are described in Section 4. The full '97 data set was then used as a sample set to perform a blind analysis of the '99 data set. Results of this combined analysis are described in Section 6.

3.3 Normalization Mode

To calculate a branching ratio one must know how many particles were present to begin with. Trying to calculate the kaon flux directly from the number of protons injected results in a great deal of uncertainty. It also does not take into account the inefficiencies of the detector. Instead, while collecting the search data, one simultaneously collects another decay mode whose branching ratio is well known. This second decay channel is called the normalization mode. The branching ratio of the signal data can then be given in proportion to the normalization mode's branching ratio. If the normalization mode chosen is very similar to the signal mode, then some uncertainties and inefficiencies cancel out. We selected the K3pi decay ($K_L \rightarrow \pi^+\pi^-\pi^0$) to use as our normalization mode because it is a three-body decay that includes a neutral pion and has a large branching ratio. The large branching ratio meant that normalization events should have been plentiful even with trigger prescaling.

Normalization events were treated to cuts as reasonably similar to the signal selection cuts as possible. The main difference was that some particle ID cuts were removed for the normalization events. The normalization selection cuts used in each analysis are listed in the corresponding analysis chapter. Unlike the signal selection cuts, the normalization selection cuts included the signal box by definition since we wanted to be sure passing events were really K3pi decays. The normalization signal box was identical to the search signal box. The number of normalization data events falling within the signal box was used to find the kaon flux via the following equation:

$$Flux(K_L) = \frac{(Number\,of\;K3pi\,data\,accepted) \times (Prescales\,on\,K3pi\,events)}{(BR(K_L \to \pi^+\pi^-\pi^0)) \times (BR(\pi^0 \to 2\gamma)) \times (K3pi\;acceptance)} \quad (3.1)$$

Normalization data was taken in trigger 2, which was a minimal bias trigger requiring at least two tracks in the spectrometer. It had an online prescale of 500:1. In the 1997 run, our normalization data sample had an off-line prescale of 20:1. In 1999, the off-line prescale was 40:1. The branching ratio for the K3pi decay is 12.56±0.20% and the branching ratio for $(\pi^0 \to 2\gamma)$ is 98.80±0.03%[8]. The acceptance was determined with Monte Carlo normalization events. As described below, this flux only included spills that were considered “good” according to the KTSPILL cut, and so could be lower than KTeV fluxes quoted elsewhere. An example distribution of normalization events is shown in Figure 3.3.

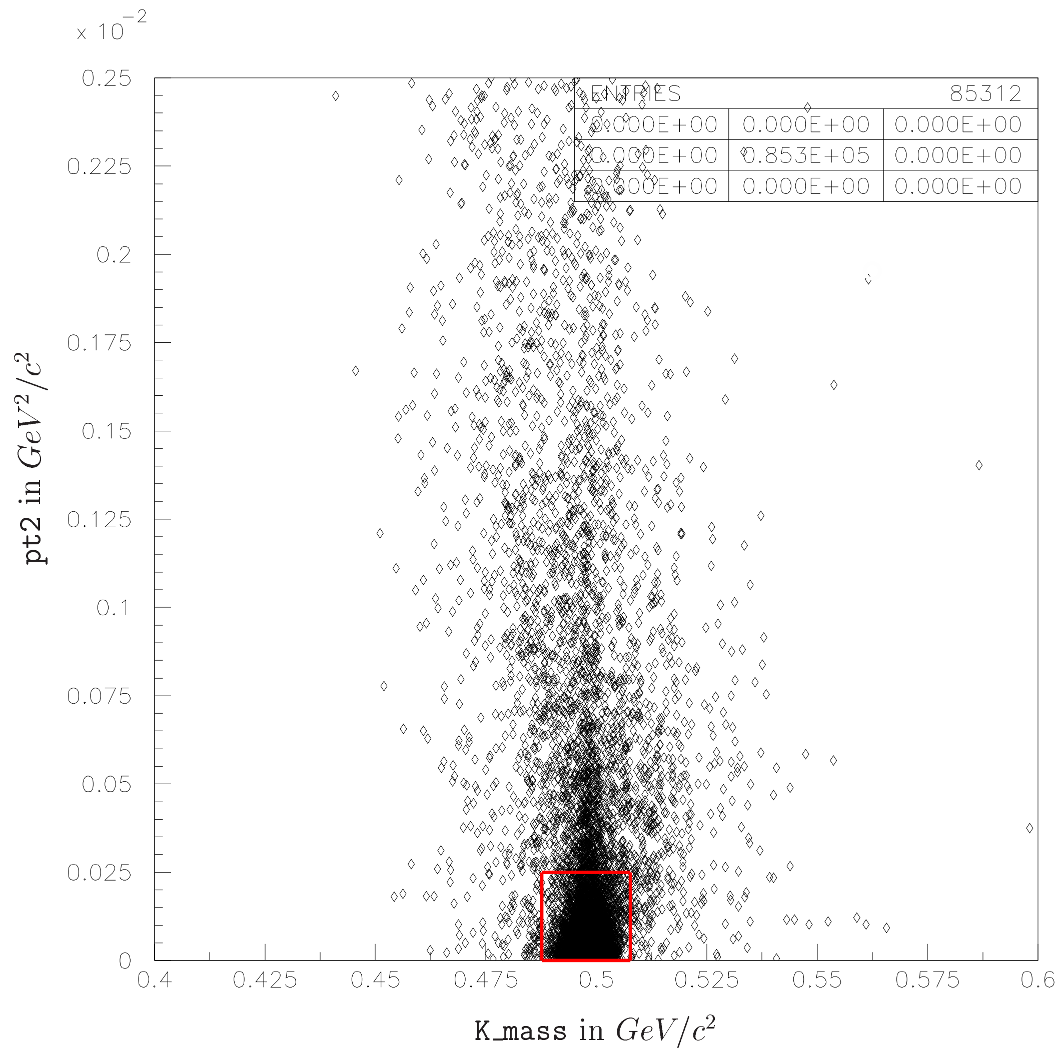

Figure 3.3 1999 trigger 2 data events reconstructed as K3pi's and passing all the normalization selection cuts except the "signal" box. The number of events inside the red signal box (81,949) was used in calculating the 1999 kaon flux.

3.4 Parameter Definitions

Even before the data was collected, a list of possible parameters for the signal was created. Preliminary cutoff values for each parameter were decided upon based on past knowledge and theoretical predictions. As the analysis progressed each parameter cut was examined individually for its signal acceptance and background rejection. Some parameters were added, some were removed, and some had their cutoff values adjusted. Definitions for the parameters used in the final versions of the '97 and '99 analyses are given here, along with justification of the cut locations selected.

In some cases the distributions from the '99 data set did not match those from the '97 data set. If we were satisfied with the effect of the '97 cut location, we chose the '99 cut location based on getting the same acceptance in that parameter across the two years. If the acceptance was deliberately changed, then the '99 cut location was chosen based on our new criteria, and then a "study cut" for the '97 event sets was chosen to match the acceptance of the '99 cut. This study cut was used only for comparison studies and not used in the calculation of a blind limit since it was selected after opening the '97 search data box.

Since most of our background was expected to come from events containing accidental activity, many of the cuts were against such accidentals. They are indicated by the letters "AAP" for "anti-accidental parameters".

3.4.1 KTSPILL definitions

Ideally, once an experiment is running it should be able to continually take good quality data. Real life is never ideal. Occasionally various data-taking systems wrote out garbage or simply failed. A routine called **KTSPILL** was written to look for information about how well each system was working. We cut events which did not have good information taken in the following systems:

- **Triggers** - If the trigger hodoscopes were not working correctly, all the timing information about an event becomes suspect. Timing information was critical in determining if the particles detected were from the same event.
- **Pipeline** - This system was responsible for holding information sent from the calorimeter while the level 3 software analysis package decided to keep the event or throw it away. Problems here indicated lost or misplaced calorimeter information.
- **Global CsI** - Since the calorimeter was an essential part of measuring the energy of particles, events where there was a problem with the calorimeter's hardware systems were flagged.
- **ETOT** - This was a subsystem of the calorimeter readout which did a quick summation of total energy deposited and determined whether to trigger on the event. If this parameter was flagged, there was a problem reading out the amount of energy in the calorimeter.

- **Drift chambers** - If a chamber's high voltage tripped off, or a feedback loop created oscillations on the chamber readout, the chamber's power had to be cycled back on. This process took several spills, and information gathered while a chamber was turning back on was always suspect.

- **Muon banks** - Since one of the major aspects of the trigger for our signal decay was the presence of a muon, reliable data from the muon banks was a necessity.

- **HCC** - HCC stands for Hardware Cluster Counter and was another subsystem of the calorimeter that was also used in the trigger. Energy in the calorimeter needed to be resolvable into clusters to tell were the center of each particle's hit was located. The timing in this system was sharp and it required that a significant amount of energy be deposited to identify a hit as a cluster. Since three of the four particles in the decay were required to deposit hardware clusters (see Section 2.4.2), HCC information must have been complete and correct. If there was trouble with cluster resolution, this bit was flagged.

Obviously, the MC generated events had no such detector-based errors since doing so would have been counter-productive. However, in finding acceptances and total kaon flux using the MC, we did not want to skew the numbers by generating events that fell within the later rejected spills. Therefore, a bit mask corresponding to the cuts applied in KTSPILL was inserted in each MC job so that events were not generated in spills that would have eventually been removed by this subroutine.

3.4.2 KTeVAna subroutine errors

The analysis program was broken down into subroutines that were each responsible for reconstructing some part of the event. If the subroutine had a problem getting the information it needed, or was unsuccessful at performing the task requested, it returned an error. An event was cut if any of the following subroutines returned in error:

- GET_VETO - Information from the veto system was lost or garbled, so it was unknown if all the particles from an event were contained.
- KTCLUS - Similar to the "HCC" tag in KTSPILL; there was a problem resolving information about clusters of energy in the calorimeter; KTCLUS was the routine which did the final calorimeter clustering[36, 37].
- T3MATCH - The program was unable to satisfactorily match the tracks found with the calorimeter clusters that were found[38].
- T3FVTX - The program was unable to resolve a reasonable vertex for the event[38].
- T3FTRK - A more detailed search for particle tracks in the drift chambers discovered information problems[38].
- KTTDC - There was a problem with information conversion in the reading out of the time-to-digital converters (TDCs). This included all TDCs except for the ones associated with the drift chambers.

- T3MASS - There was a problem reconstructing the mass of the charged system from the information known about the charged decay products[38].

- T3FPI0 - The program was unable to reconstruct a π^0 (within 250 MeV of the π^0 mass) from any combination of calorimeter clusters and the charged vertex that was found[38].

3.4.3 Detector Parameters

Parameters listed here and later were investigated offline using a program called the Physics Analysis Workstation (PAW). The parameter name is listed followed by its description and value requirements. As a reminder, parameters with the designation "(AAP)" are anti-accidental parameters.

- **phvbar1** - (AAP) a Level 1 trigger latch logic variable that combined veto information from the ring counters (RCs), spectrometer antis (SAs), and collar anti (CA). It was used as part of the requirements for the signal trigger (Trigger 7) but not for the normalization trigger (Trigger 2). An event was not accepted unless this variable was true. The variable was included in the event selection cuts to get better agreement between signal mode selection and normalization mode selection. All signal mode events automatically passed the cut on this parameter. Although we had separate, tighter cuts on the amount of energy deposited in the RCs and SAs, `phvbar1` included a cut on CA energy that we did not have elsewhere.

- **RCmaxene** - (AAP) the maximum amount of energy deposited in any single Ring Counter (RC) detector. Accepted events were required to have less than or equal to $0.5GeV$ in this variable. For an example distribution, see Figure 3.4.

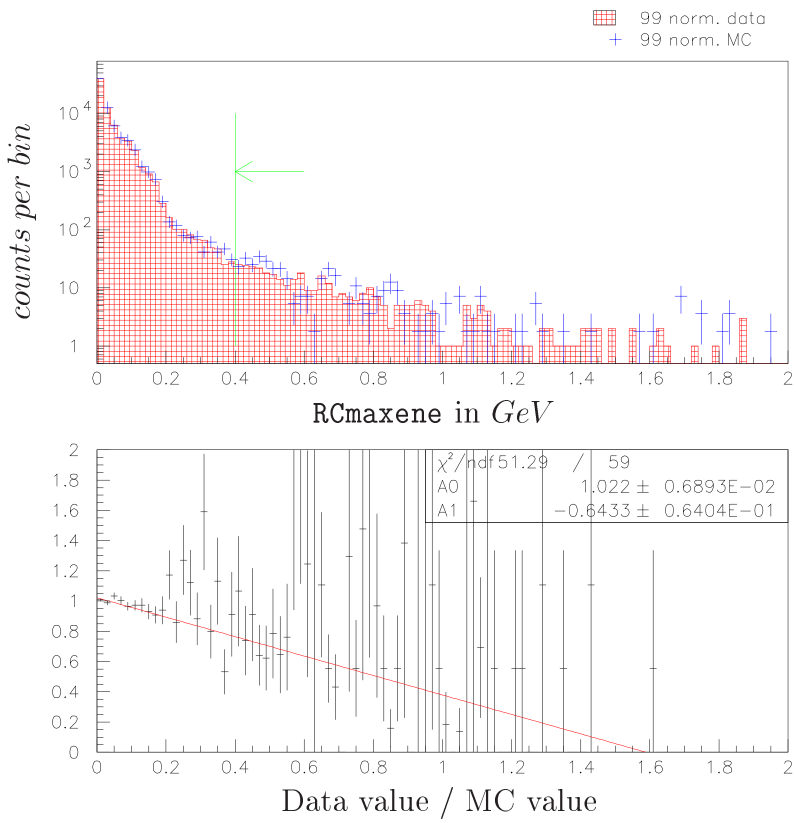

Figure 3.4 Histograms of the maximum amount of energy deposited in the ring counters (`RCmaxene`) for 1999 normalization data and MC events passing all other normalization selection cuts. A cut was placed at $0.5GeV$. The histogram has been normalized to the number of data events. Note the log scale on the Y-axis.

- **SAmaxene** - (AAP) the maximum amount of energy deposited in any single spectrometer anti (SA) detector. Accepted events were required to have less than or equal to $0.4GeV$ in this variable. For an example distribution, see Figure 3.5.

- **EBA1** or **BA1ene** - (AAP) the amount of energy deposited in the first electromagnetic portion of the back anti (BA) detector. This parameter was not used in the '97 blind analysis. The hardware for this detector changed between 1997 and 1999, causing the distributions of this variable to be different. To decide where to place the cut on this variable, we looked at the masked signal search data after all cuts except this one (and the signal box) had been applied (see Figure 3.6).

Note that the cuts applied to the '97 events in this case were the study cuts, not the blind analysis cuts. The `BA1ene` parameter was not cut on in the blind '97 analysis. The most logical place for a `BA1ene` cut seemed to be at 12 for the 97 data set (top plot), and at 5 for the '99 data set (bottom plot). In addition, using a cut at 12 for '97 and at 5 for '99 kept the same amount of signal MC in each case (87.8±0.5% for 1997, 86.5±0.5% for 1999). The second section of the BA was turned off during the 1999 data run (before we realized we wanted to use it in our analysis), so it provided no information for us to cut on.

- **EBA3** - (AAP) the amount of energy deposited in the third (hadronic) portion of the back anti (BA) detector. It was cut on only loosely and only in the '97 blind analysis. Figure 3.7 shows the BA3 energy distribution of the search data with no cuts. The main part of the distribution has a peak at zero and a positive tail. If BA3 detected more than 1024 ADC counts, the variable `EBA3` was assigned the "overflow"

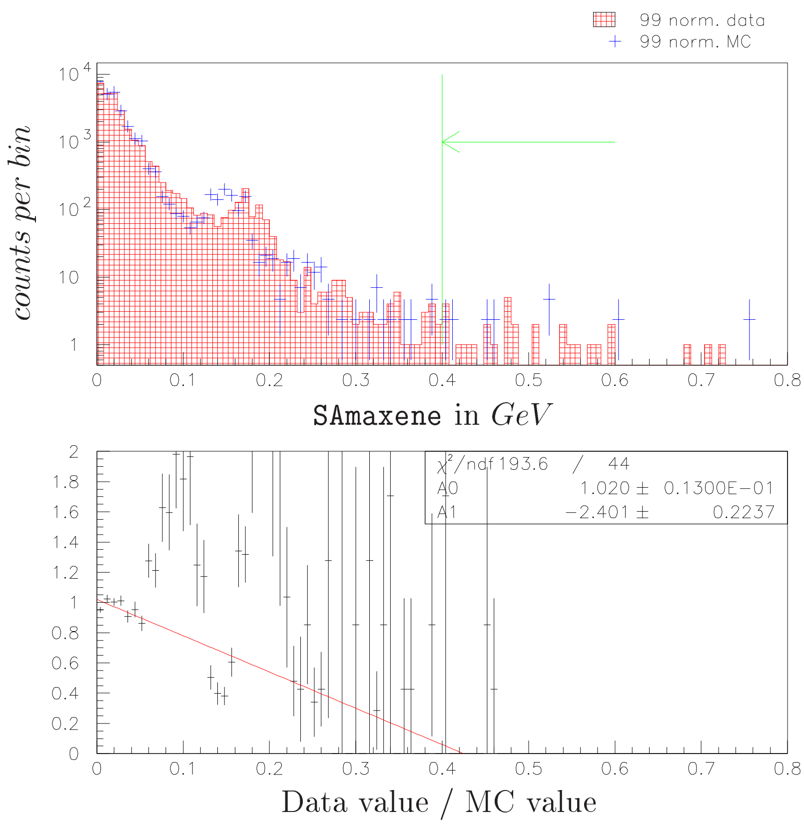

Figure 3.5 Histograms of the maximum amount of energy deposited in the spectrum antis (`SAmaxene`) for 1999 normalization data and MC events passing all other normalization selection cuts. The distributions agree reasonably well and a loose cut was placed at $0.4GeV$. I am unsure what the bump near $1.8GeV$ is from, but it was modeled in the MC even if it was not at exactly the right place. The histogram has been normalized to the number of data events. Note the log scale on the Y-axis.

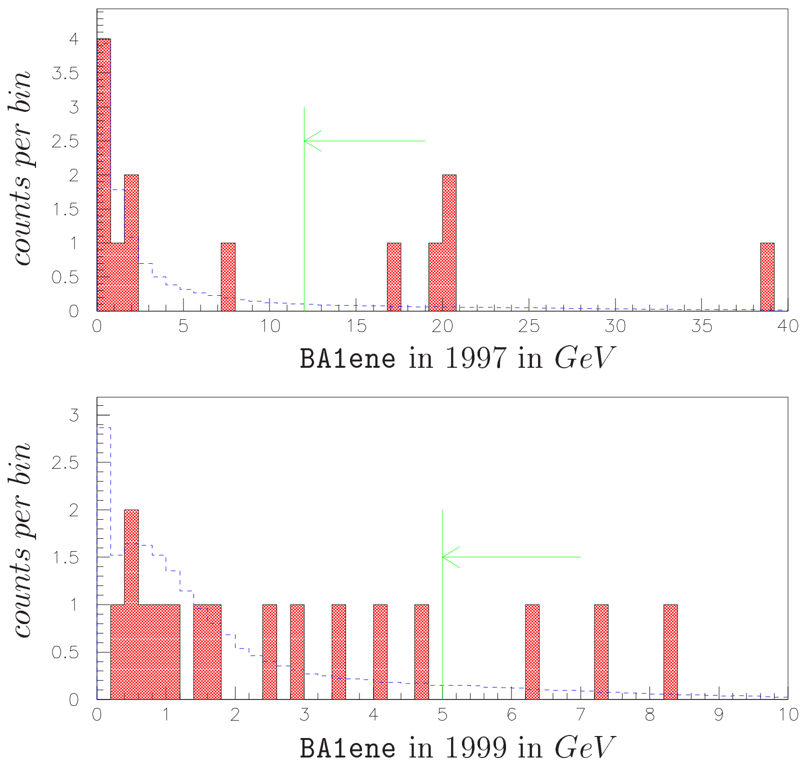

Figure 3.6 Histograms of search data events (solid red histograms) from outside the signal box that passed all cuts except the `BA1ene` cut. The BA1 energy distribution was much more spread out in '97 data (top plot) than in '99 data (bottom plot), so the cut placement had to be different. This parameter was not cut on in the '97 blind analysis, so the '97 event sets shown here were with the study cuts applied. In each case, there seemed to be a clear break in which to place the cut - at 12 for '97 data (again, for use as a study cut) and at 5 for '99 data. This cut location was also reasonably far out on the tail of well-reconstructed events as shown by the normalization data (unfilled blue histograms). The normalization data has been renormalized to fit on the plot.

value of 9999 (the peak at the far right in the plot). If nothing was detected by BA3, `EBA3` was assigned the "underflow" value of -999 (the peak at the far left). Placing our cut at 2000, we removed only events whose energy overflowed BA3.

- **npln_e (npln_m)** - number of TRD planes hit along the "electron" ("muon") track. The transition radiation detectors (TRDs) were used to help distinguish between electrons and charged pions. There were a total of 16 planes of these chambers in our detector. In order to identify a particle, it had to have deposited energy along it's track in at least one of the TRD planes. Hence this number must have been greater than zero for each charged track that we wanted to put a cut on. In '97, we cut on both the electron and muon tracks, so we require both `npln_e` and `npln_m` to be greater than zero. In '99, a cut was only made on `prob_e`, so we only required `npln_e` to be greater than zero.

- **nxtrks (nytrks)** - (AAP) the number of X (or Y) tracks found in the event by routine T3TRAK before any kind of X-Y track matching or vertexing is done[38]. An ideal normalization or signal event would only have a value of 2 for both `nxtrks` and `nytrks`. Cutting at 10 on `nxtrks` and 5 on `nytrks` removed events with excessive accidental activity or drift chamber noise. Plots of nxtrks and nytrks (Figures 3.8 and 3.9) showed that the cut on `nytrks` could be made tighter than that on `nxtrks` without losing too much signal or normalization acceptance. These parameters were only cut on in the '99 analysis.

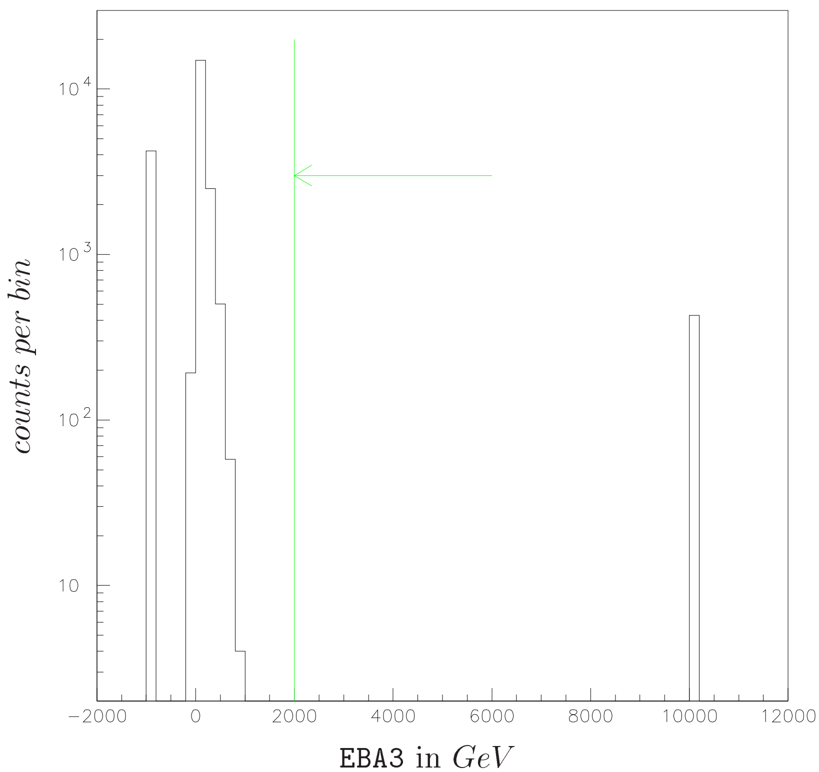

Figure 3.7 Histogram of 1997 search data events with no cuts beyond the ntuple level. The BA3 energy distribution (`EBA3`) had a main distribution between zero and 1024 ACD counts, and two "tag" values. A value of 9999 indicated the ADC went over 1024 counts and was hence called the "overflow" tag. A value of -999 indicated the ADC did not even read out a value of zero, and was called the "underflow" tag. The loose cut, used only in the 1997 analysis, was placed at 2000 to remove only the "overflow" events.

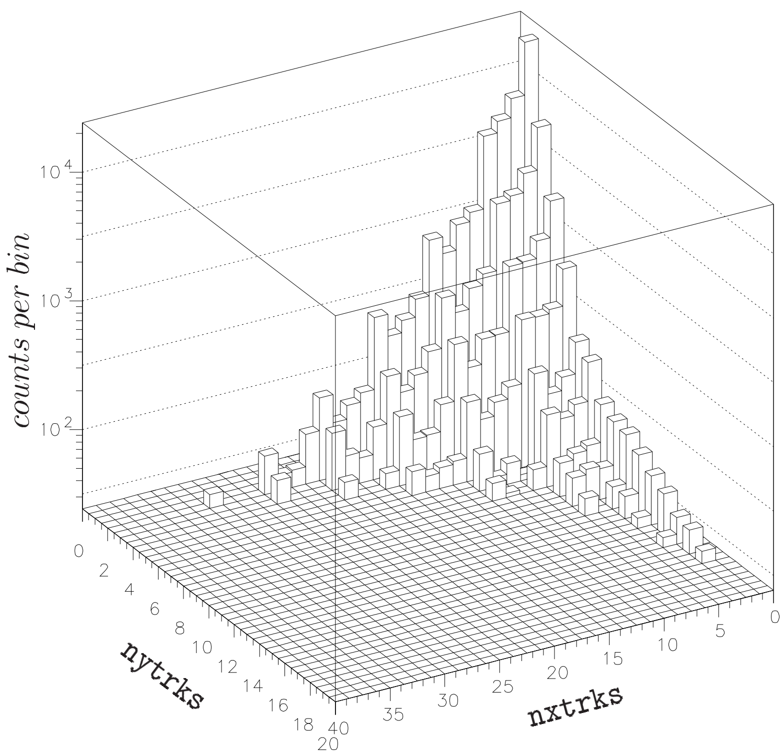

Figure 3.8 Histogram of the number of pre-vertexing Y-tracks (`nytrks`) versus X-tracks (`nxtrks`) for 1999 normalization data with all other cuts. The cuts on these parameters accepted `nxtrks` less than 10 and `nytrks` less than 5.

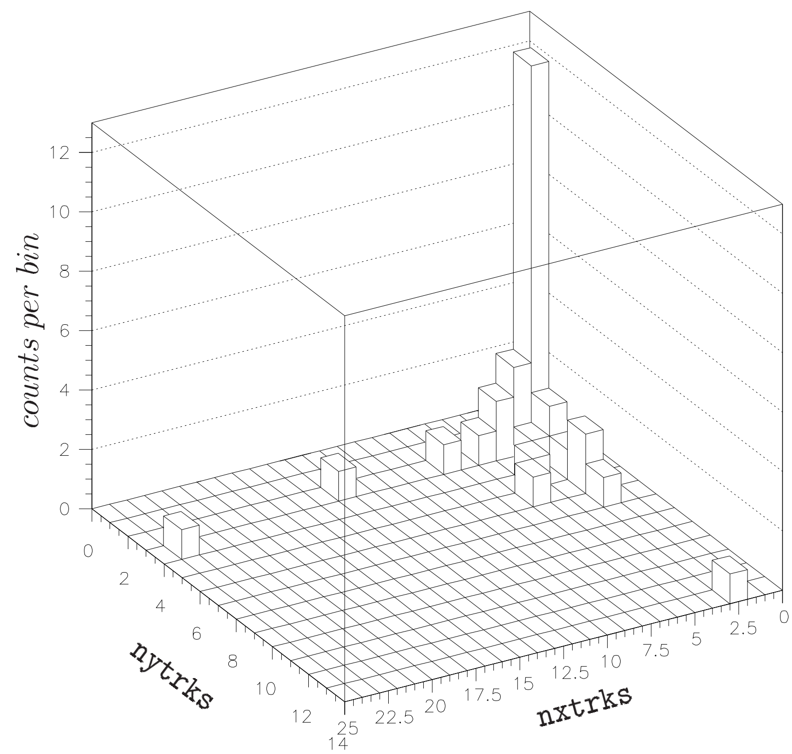

Figure 3.9 Histogram of the number of pre-vertexing Y-tracks (**`nytrks`**) versus X-tracks (**`nxtrks`**) for 1999 search data with all other selection cuts. The cuts on these parameters accepted **`nxtrks`** less than 10 and **`nytrks`** less than 5.

- **npairsup (npairsdn)** - (AAP) the number of in-time hit pairs in the two drift chambers upstream (downstream) of the analysis magnet [DC1 and DC2] ([DC3 and DC4]). Since we wanted two tracks and there were four drift chambers, the cut on npairsup was placed at 8. For an example distribution, see Figure 3.10. The cut on npairsdn was slightly looser, placed at 10. For an example distribution, see Figure 3.11. The discontinuity in the ratio plots of these two parameters shows that these cuts were partially responsible for the artificial loss in flux in 1997 (see Appendix B). These parameters were reworked for the '99 analysis (see `xpairsup` and `xpairsdn`) to reduce this systematic error.

- **xpairsup (xpairsdn)** - (AAP) the number of in-time hit pairs in the two drift chambers upstream (downstream) of the analysis magnet [DC1 and DC2] ([DC3 and DC4]), exclusive of the pairs on a track or within 1.5 cell-widths of a track. Cuts on the parameters `npairsup` and `npairsdn` showed a systematic bias for accepting more MC events than data events. The code made no distinction between pairs that were part of a track, those right next to a track, and those farther away. When a charged particle goes through a drift chamber, it can produce a delta ray. This extra energy in the chamber could have created extra hits adjacent to the track hits that were then reconstructed as good in-time pairs. Our MC did not model delta rays well. The coding of `npairsup` and `npairsdn` included these delta ray pairs when counting, ending up with a systematically higher value for data events than for MC events. The

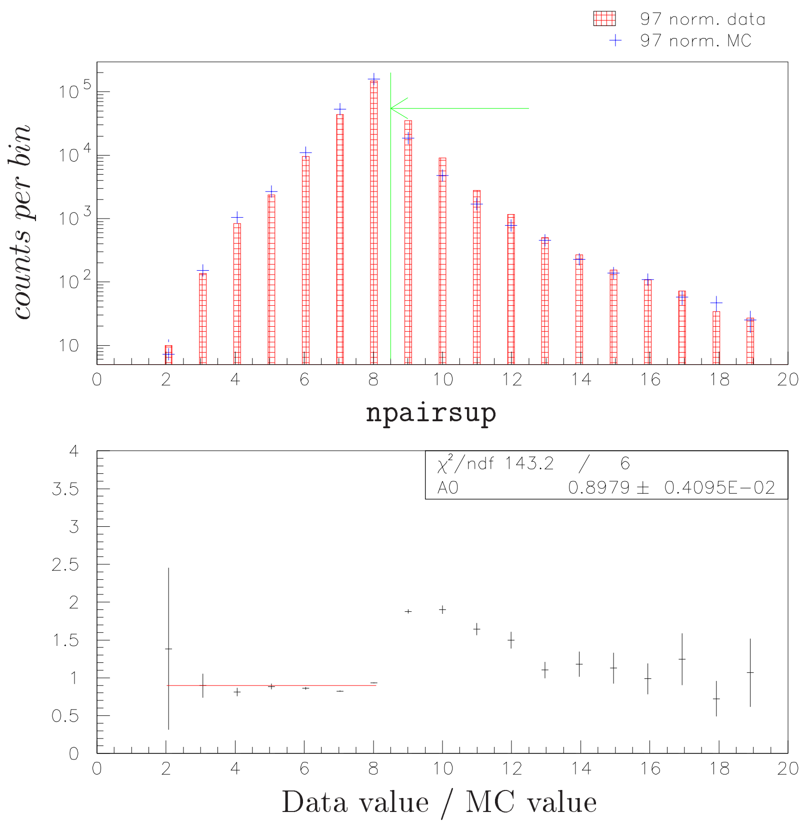

Figure 3.10 Histograms of the number of hit pairs in the two drift chambers upstream of the analysis magnet ("npairsup") for 1997 normalization data and MC events passing all other normalization selection cuts. The selection cuts accepted events with a value of 8 or less (2 tracks × 2 chambers × 2 planes per chamber). The histogram has been normalized to the number of data events. Note the log scale on the Y-axis.

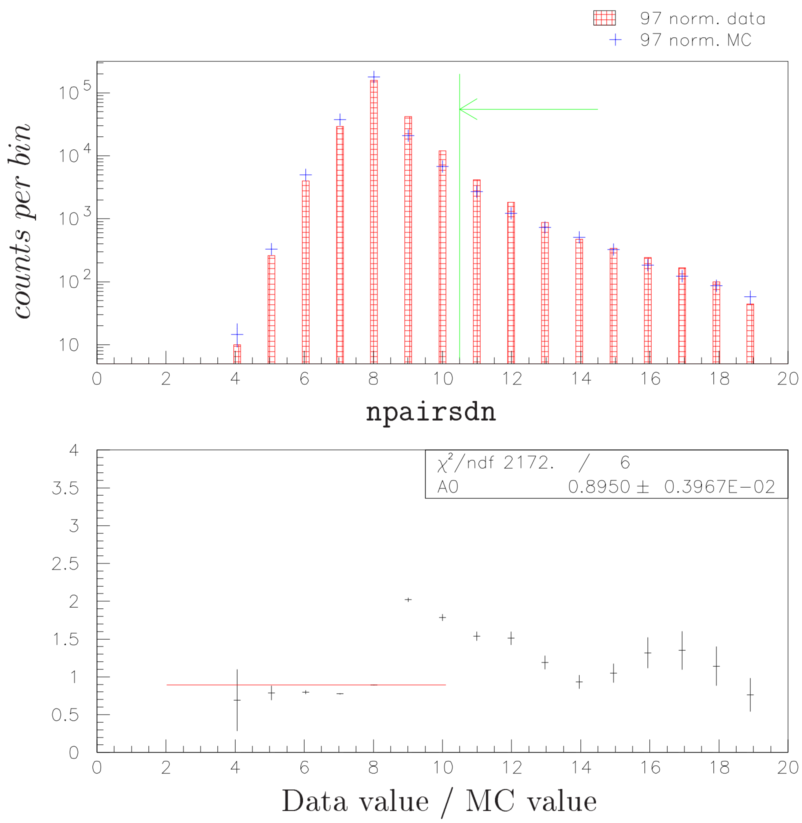

Figure 3.11 Histograms of the number of hit pairs in the two drift chambers downstream of the analysis magnet ("npairsdn") for 1997 normalization data and MC events passing all other normalization selection cuts. The selection cuts accepted events with a value of 10 or less (2 tracks $\times$ 2 chambers $\times$ 2 planes per chamber, plus minor accidental activity). The histogram has been normalized to the number of data events. Note the log scale on the Y-axis.

difference showed up in the normalization mode when adding the cut on `npairsup` or `npairsdn` resulted in a 5% to 10% drop in the flux number. Unfortunately this problem with `npairsup` and `npairsdn` was not realized until after the '97 signal box had been opened, so `xpairsup` and `xpairsdn` was not used in the '97 blind limit; they were only cut on in the '99 analysis. `Xpairsup` and `xpairsdn` were incremented only when an in-time hit was not part of an identified track and was farther than 1.5 cell-widths from a track hit. The cut on `xpairsup` was tight, requiring that there be no extra hits in DC1 and DC2. The cut on `xpairsdn` was looser, accepting events with a value less than or equal to 2. Normalization data versus MC is shown for these parameters in Figures 3.12 and 3.13.

3.4.4 Reconstruction parameters

- **eoffx (eoffy)** - the offset in the X (or Y) direction between the reconstructed electron track upstream of the analysis magnet and the same track downstream of the magnet. Ideally the track segments would have intersected and this value would have been zero. We accepted an absolute value of $0.002m$ or less. Figures 3.14 and 3.15 show example distributions of this variable for a charged pion, which looks very much like an electron.

- **muoffx (muoffy)** - the offset in the X (or Y) direction between the reconstructed muon track upstream of the analysis magnet and the same track downstream of the magnet. Ideally the track segments would have intersected and this

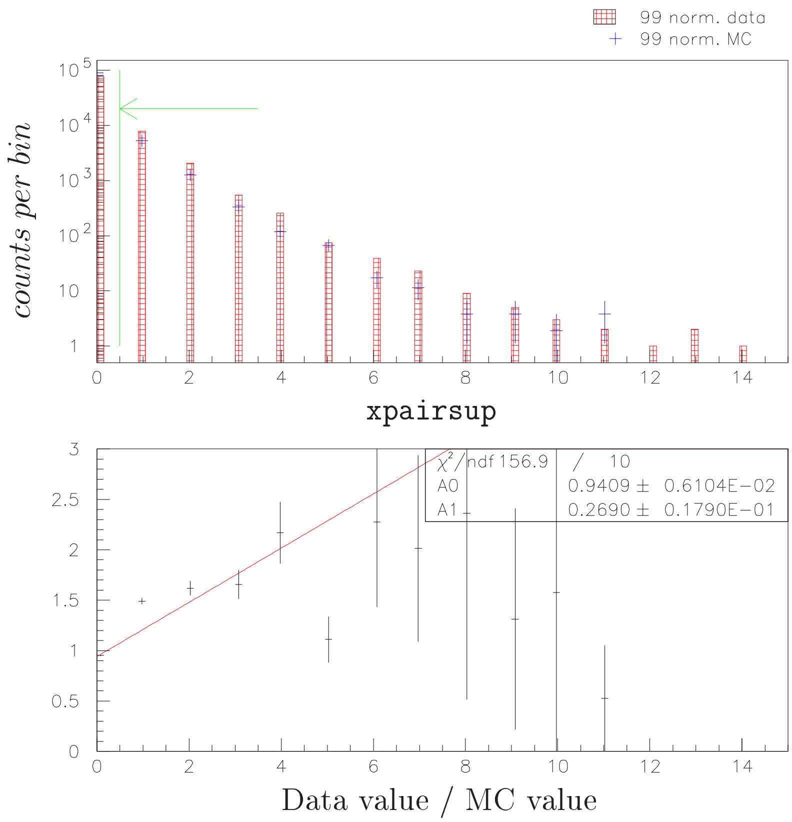

Figure 3.12 Histograms of the number of extra hit pairs in the two drift chambers upstream of the analysis magnet (`xpairsup`) for 1999 normalization data and MC events passing all other normalization selection cuts. This parameter did not count hits that were part of a track or within 1.5 cell-widths of a track. Hits within 1.5 cell-widths of a track hit were likely from delta rays from the charged particle, and were not considered extra chamber activity. The selection cuts only accepted events with a value of 0 (no hits external to the reconstructed tracks). The histogram has been normalized to the number of data events. Note the log scale on the Y-axis.

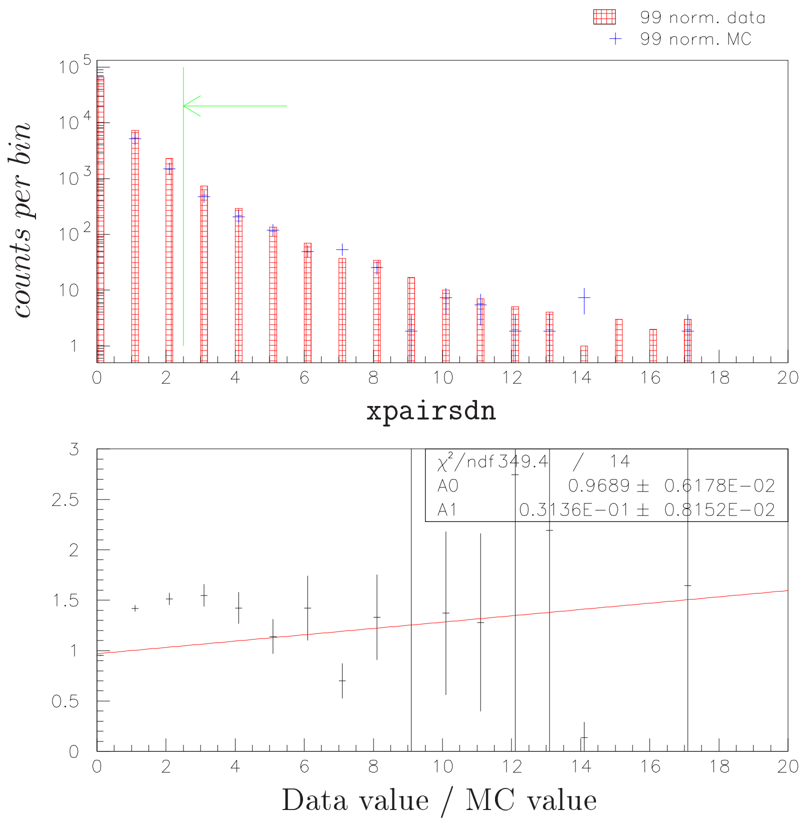

Figure 3.13 Histograms of the number of extra hit pairs in the two drift chambers downstream of the analysis magnet (`xpairsdn`) for 1999 normalization data and MC events passing all other normalization selection cuts. This parameter did not count hits that were part of a track or within 1.5 cell-widths of a track. Hits within 1.5 cell-widths of a track hit were likely from delta rays from the charged particle, and were not considered extra chamber activity. The selection cuts accepted events with a value of 2 or less (allow some minor accidental activity). The histogram has been normalized to the number of data events. Note the log scale on the Y-axis.

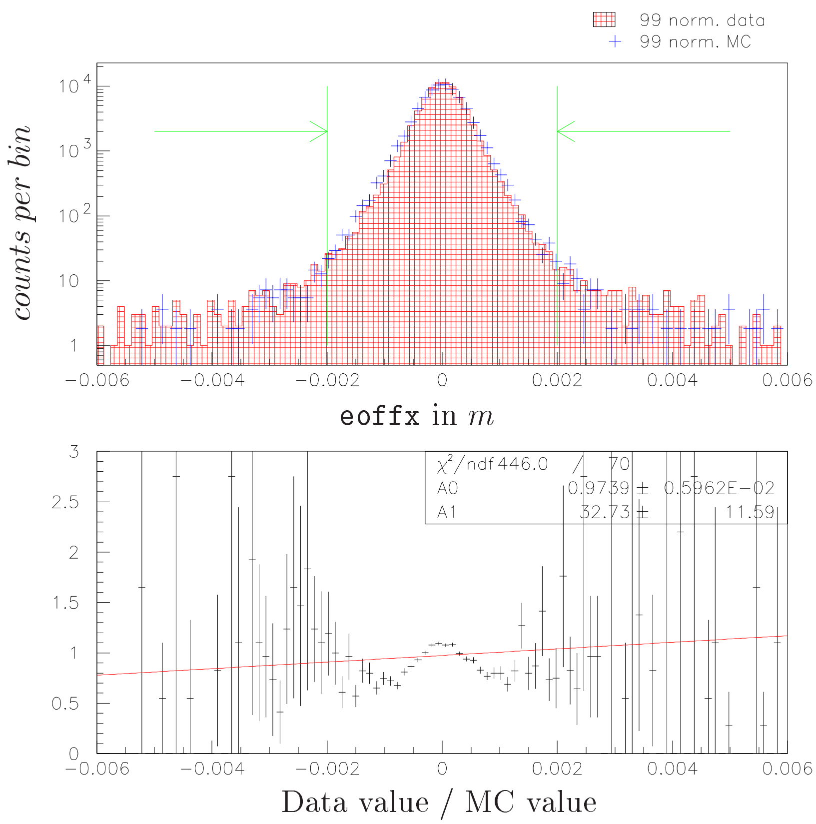

Figure 3.14 Histograms of upstream to downstream track matching in the magnet in X (`eoffx`) for 1999 normalization data and MC events passing all other normalization selection cuts. These histograms are actually of a charged pion, but the distributions for pions, electrons and muons were all similarly shaped. The selection cuts accepted events with an absolute value of $2mm$ or less. The histogram has been normalized to the number of data events. Note the log scale on the Y-axis.

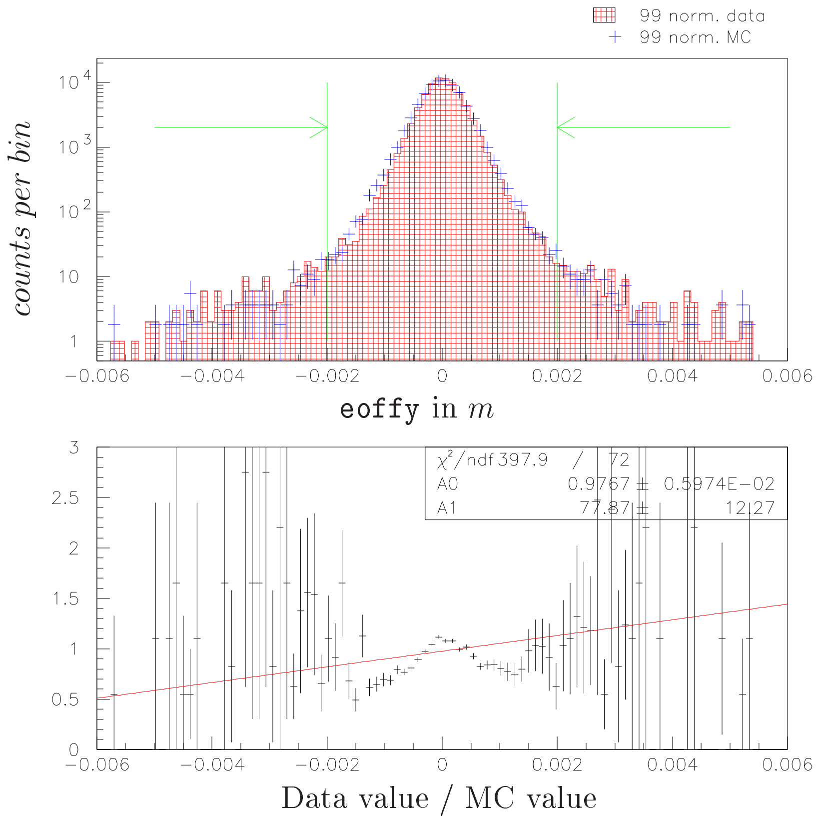

Figure 3.15 Histograms of upstream to downstream track matching in the magnet in Y (`eoffy`) for 1999 normalization data and MC events passing all other normalization selection cuts. These histograms are actually of a charged pion, but the distributions for pions, electrons and muons were all similarly shaped. The selection cuts accepted events with an absolute value of $2mm$ or less. The histogram has been normalized to the number of data events. Note the log scale on the Y-axis.

value would have been zero. Figures 3.14 and 3.15 show example distributions of the related variables `eoffx` and `eoffy` for a charged pion. These variables had much the same distribution, but a tighter cut was placed on the muon track matching to reduce the background from decays in flight. We accepted events with an absolute value of `muoffx` and `muoffy` less than or equal to $1mm$.

- **vtxx, vtxy, vtxz** - the X-, Y-, and Z-position of the vertex reconstructed from the tracks of the two charged particles by the T3FVTX routine[38]. Assuming the parent particle came from the center of the target, the vertex X and Y values were projected down to the calorimeter where they were required to lie inside one of the two beam holes ($0.075m$ to $0.225m$ for $|\texttt{vtxx}|$ and $0m$ to $0.075m$ for $|\texttt{vtxy}|$). See Figures 3.16 and 3.17 for distribution examples. The vertex Z value (`vtxz`) was required to be in the evacuated decay region ($96m$ to $155m$) and an example distribution can be seen in Figure 3.18.

- **vtx_chi** - this parameter was the chi squared divided by the degrees of freedom (χ^2/d.f.) from the fit of a decay vertex to the crossing in three dimensional space of the two charged tracks, as calculated by T3FVTX[38]. The closer the tracks were to meeting, the better the fit for the decay vertex, and the closer this parameter's value was to zero. We cut at a value of 10. For example distributions see Figures 3.19 and 3.20.

- **pt2** - the transverse momentum squared of the event after reconstruction at

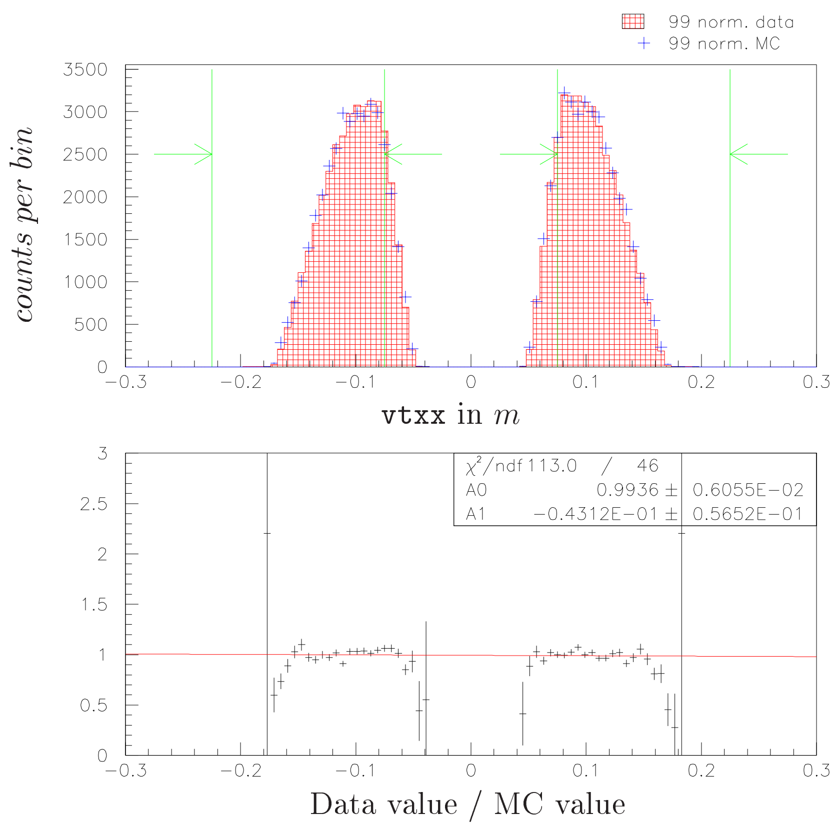

Figure 3.16 Histograms of the X-position of the charged vertex as projected onto the calorimeter face (**vtxx**) for 1999 normalization data and MC events passing all other normalization selection cuts. The selection cuts accepted events within the two calorimeter beamholes ($|$**vtxx**$|$ greater than $0.075m$ and less than $0.225m$). The histogram has been normalized to the number of data events.

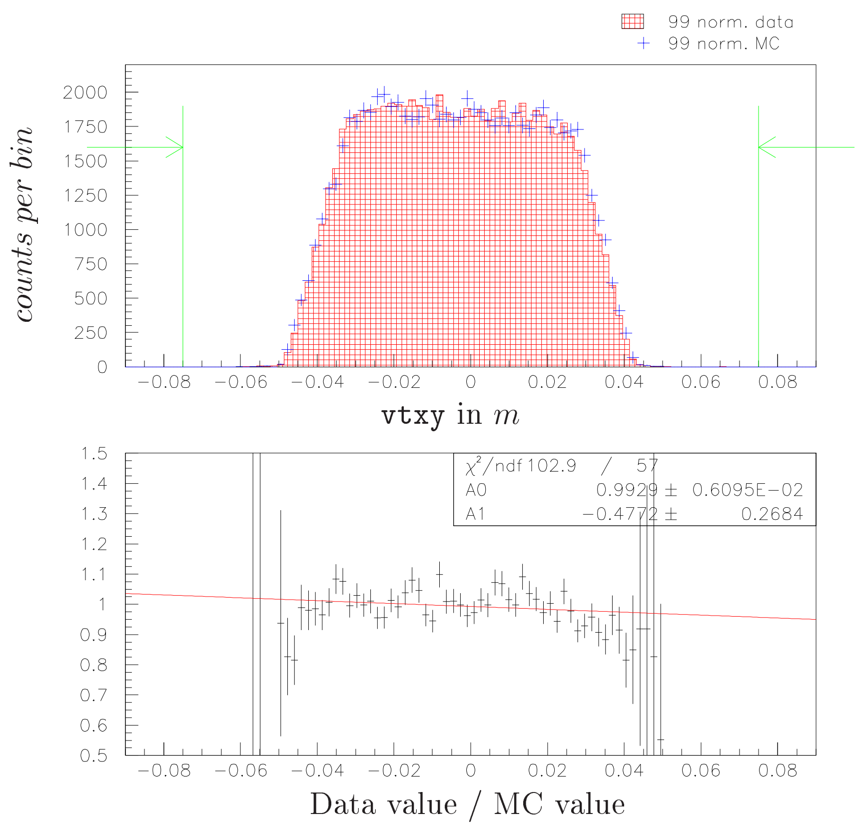

Figure 3.17 Histograms of the Y-position of the charged vertex as projected onto the calorimeter face (`vtxy`) for 1999 normalization data and MC events passing all other normalization selection cuts. The cuts on this parameter (|`vtxy`| less than $0.075m$) lay well beyond the edges of the distribution. The histogram has been normalized to the number of data events.

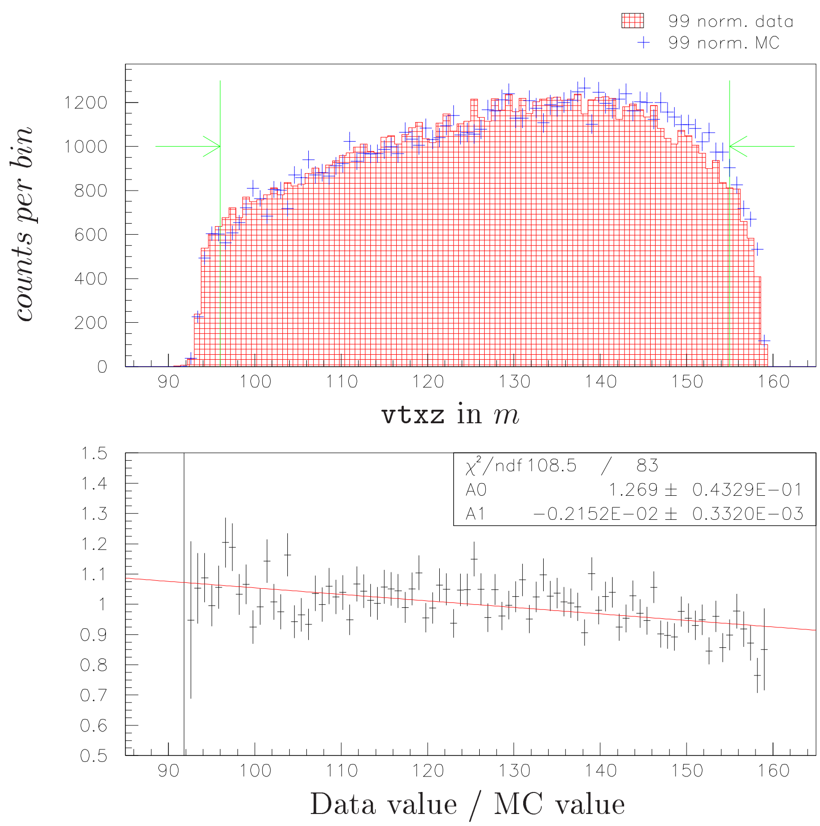

Figure 3.18 Histograms of the Z-position of the charged vertex (`vtxz`) for 1999 normalization data and MC events passing all other normalization selection cuts. The selection cuts accepted events with `vtxz` greater than 96m and less than 155m. The histogram has been normalized to the number of data events.

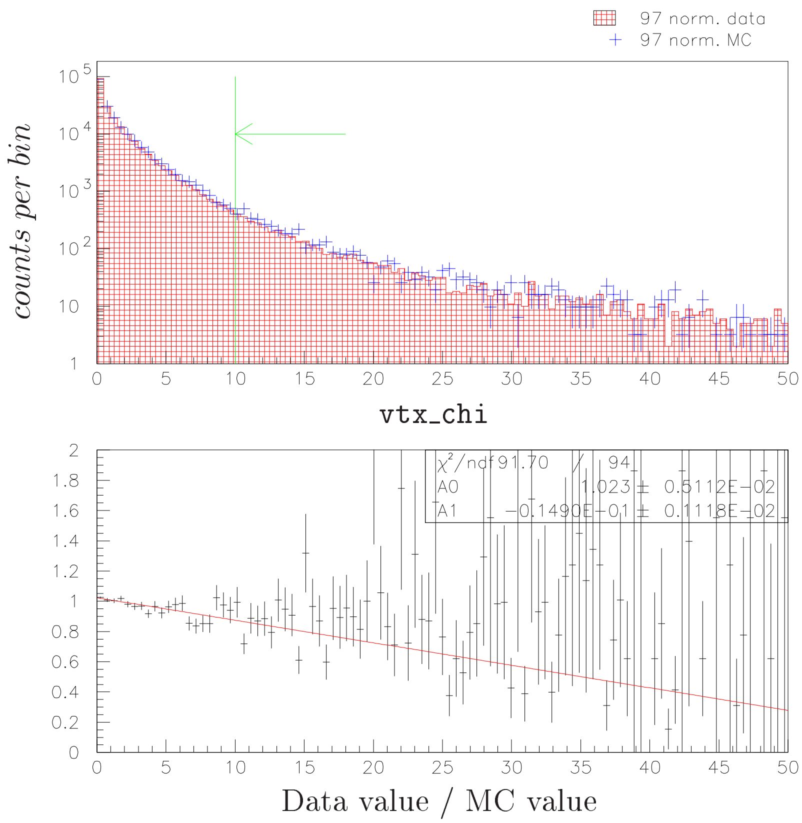

Figure 3.19 Histograms of the chi squared divided by the degrees of freedom (χ^2/d.f.) from the fit of a decay vertex to the crossing in three dimensional space of the two charged tracks (`vtx_chi`) for 1997 normalization data and MC events passing all other 1997 normalization selection cuts. The selection cuts accepted events with a value of 10 or less. The histogram has been normalized to the number of data events. Note the log scale on the Y-axis.

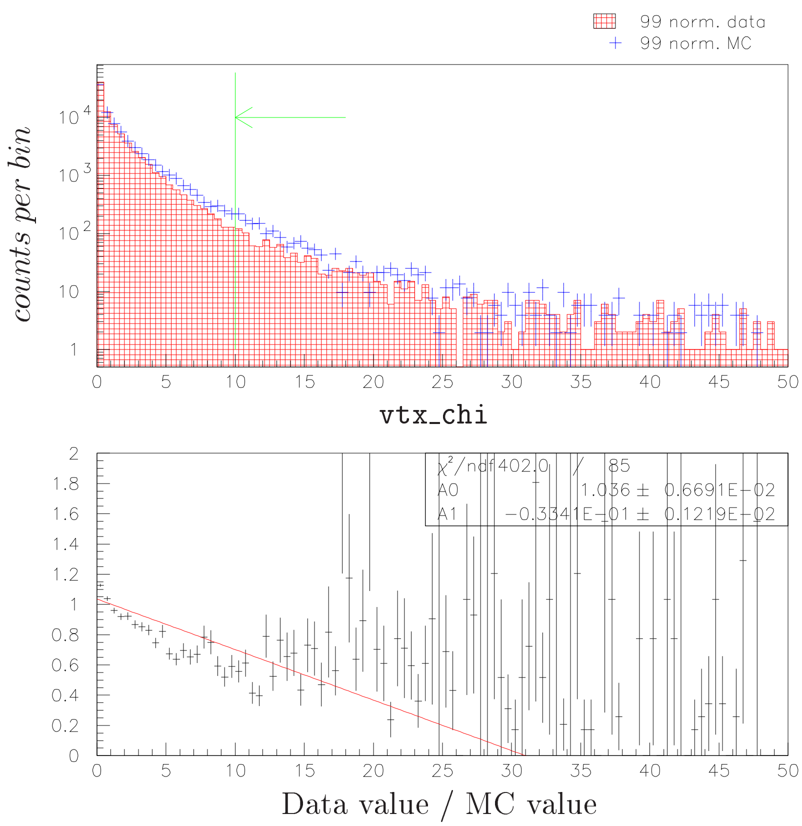

Figure 3.20 Histograms of the chi squared divided by the degrees of freedom (χ^2/d.f.) from the fit of a decay vertex to the crossing in three dimensional space of the two charged tracks (`vtx_chi`) for 1999 normalization data and MC events passing all other normalization selection cuts. The selection cuts accepted events with a value of 10 or less. The agreement may not look very good, but this cut's impact on the overall acceptance of MC versus data is negligible. The histogram has been normalized to the number of data events. Note the log scale on the Y-axis.

the charged vertex, as measured with respect to the flight direction of the kaon. This parameter is calculated within the T3FVTX routine[38]. If all the particles in an event were decay products, and all the decay products were observed, then this parameter would have been zero. This parameter was one dimension of the signal box with a cut value at $0.00025 GeV^2/c^2$. A looser secondary cut was made on `pt2` at $0.0025 GeV^2/c^2$, which marks the top of the study plot and is ten times larger than the signal box cut.

• **K_energy** - the energy (in the lab frame) of the parent particle. Our MC program generated kaons with energies from $20 GeV$ up to $220 GeV$. To be well within the valid energy range for comparisons, we needed to place a cut on this variable for the data. In 1997 analysis, we accepted events with `K_energy` less than $200 GeV$. In the 1999 analysis, we realized we should have also placed a lower bound on this parameter, so we only accepted events with `K_energy` greater than $20 GeV$ and less than $200 GeV$.

3.4.5 Particle ID Parameters

• **pi0_mass** - the mass of the particle reconstructed from two calorimeter clusters not associated with tracks, and using the charged vertex. This parameter is calculated by the T3FPI0 routine[38]. If the event was a Kpi0me decay, then this value should have equaled the mass of a neutral pion ($0.135 GeV/c^2$).

The MC for both 1997 and 1999 was found to have a slightly shifted pi0_mass peak when compared to the data (see Figures 3.21 and 3.23). To compensate for this difference in 1997, cuts were made at $0.132 GeV/c^2$ and $0.138 GeV/c^2$ after a correction constant of $0.35 MeV/c^2$ was added to the MC pi0_mass values (see Figure 3.22).

In 1999, the cuts were changed to be the mean plus and minus 1.25 sigma ($\mu \pm 1.25\sigma$). The mean and sigma were different depending on the event set. For the normalization data, all other cuts were taken and then a Gaussian was fit to the `pi0_mass` distribution. The resulting mean was $135.500 \pm 0.004 MeV/c^2$ and sigma was $1.638 \pm 0.005 MeV/c^2$. Therefore the cuts applied to '99 normalization data were at $0.133453 GeV/c^2$ and $0.137548 GeV/c^2$. The same process was done for '99 normalization MC and '99 signal MC, with the results listed in Table 3.2. Signal MC cuts were also applied to all background MC event sets. The normalization data cuts were applied to the search data. The means and sigmas of the '97 data and MC were also calculated, but this method of cutting was only used for study purposes and not to calculate the '97 blind results.

- **eop_elec** and **eop_mu** - "`eop`" stands for "energy (E) over (or divided by) momentum (p)" and was calculated for the particle tagged as an electron (`elec`) or muon (`mu`) in the event. This calculation was an important check on what kind of particle made the associated track and cluster. The energy in this case was the amount of energy that was deposited in the calorimeter by the particle. The momentum was that calculated from the amount of bending the track underwent through the analysis

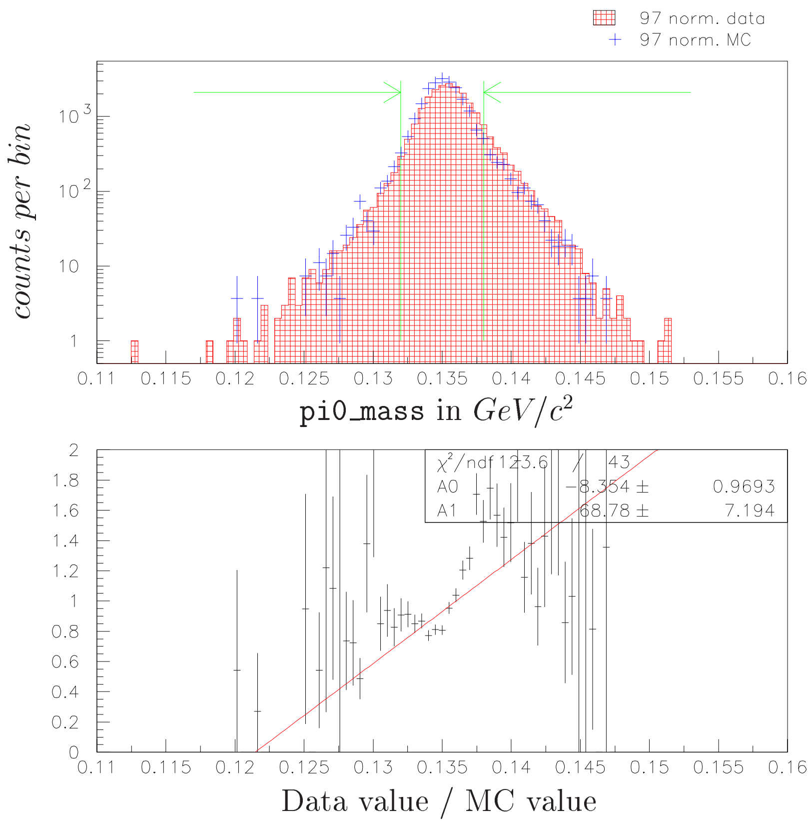

Figure 3.21 Histograms of mass of the neutral pion (M_{π^0} or `pi0_mass`) for 1997 normalization events passing all other normalization selection cuts. There was a shift of $0.35 MeV/c^2$ between the data and the MC which was corrected before cuts were made. Note the log scale on the Y-axis.

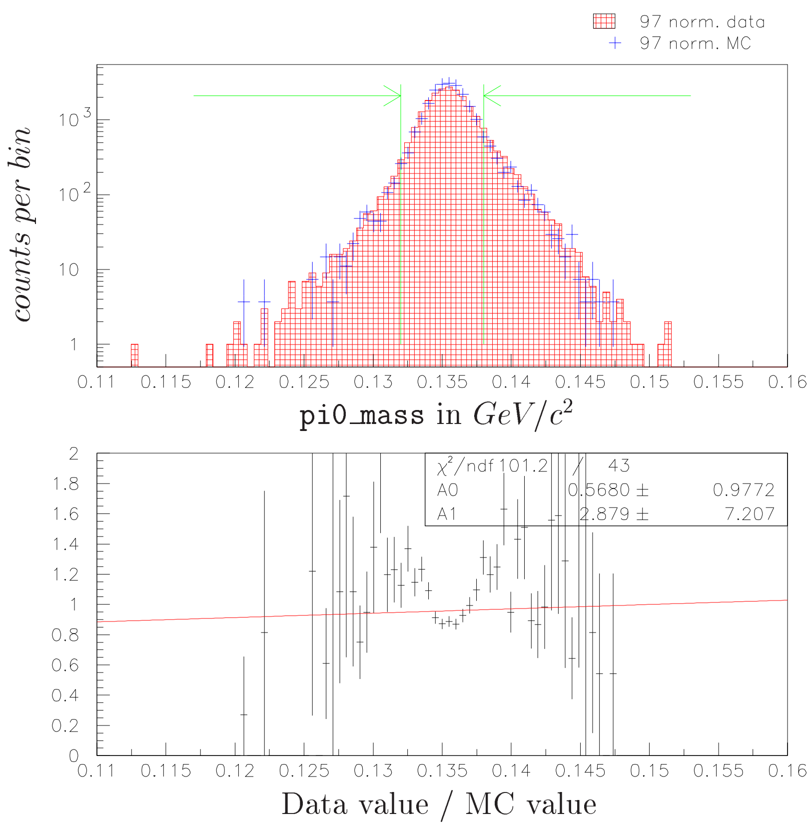

Figure 3.22 Histograms of mass of the neutral pion (M_{π^0} or `pi0_mass`) for 1997 normalization events passing all other normalization selection cuts, and with $0.35 MeV/c^2$ added to the MC values to bring the means into agreement. Note the log scale on the Y-axis.

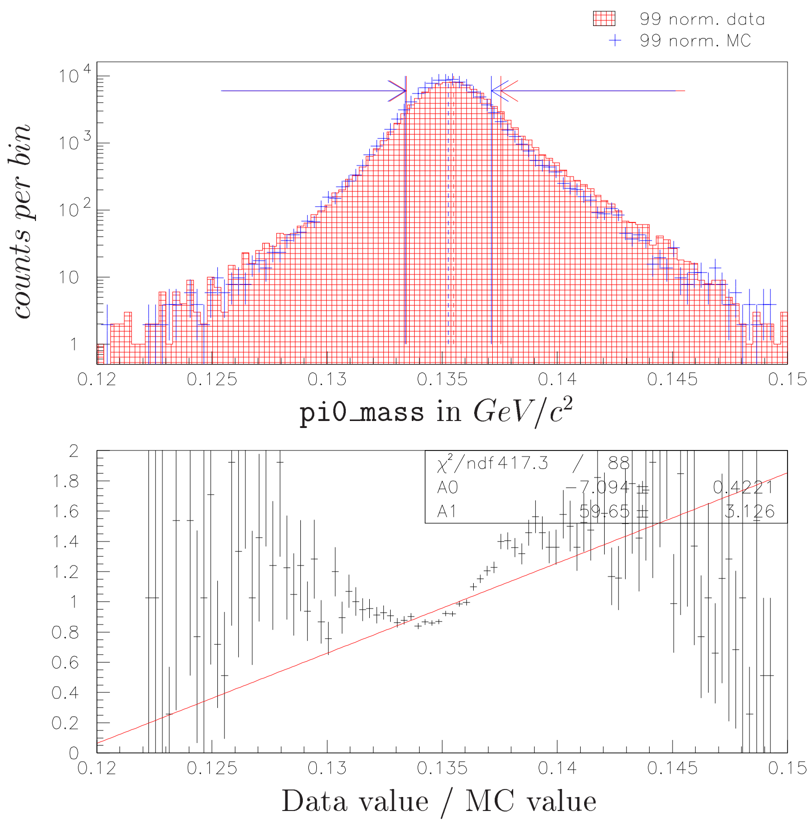

Figure 3.23 Histograms of mass of the neutral pion (M_{π^0} or `pi0_mass`) for 1999 normalization events passing all other normalization selection cuts. There was a shift of $0.23 MeV/c^2$ between the data mean and the MC mean. Cuts were placed at $\pm 1.25\sigma$ based on a Gaussian fit to each curve. The cuts for the '99 normalization data were therefore at $0.133453 GeV/c^2$ and $0.137548 GeV/c^2$ (red lines and arrows, with a mean of $0.135500 GeV/c^2$ marked by the dashed line). For the '99 normalization MC they were at $0.133390 GeV/c^2$ and $0.137130 GeV/c^2$ (blue lines and arrows, with a mean of $0.135260 GeV/c^2$ marked by the dashed line). The cuts were very tight because of the high rate of accidental photons in '99. Note the log scale on the Y-axis.

Event set	Gaussian mean μ (MeV/c^2)	Gaussian sigma σ (MeV/c^2)
'99 normalization data	135.500±0.004	1.638±0.005
'99 normalization MC	135.260±0.005	1.496±0.007
'99 signal MC	135.180±0.012	1.302±0.019
'97 normalization data	135.400±0.003	1.632±0.004
'97 normalization MC	135.050±0.004	1.403±0.006
'97 signal MC	135.010±0.001	1.180±0.015

Table 3.2 The results of Gaussian fits to the `pi0_mass` distribution for '99 and '97 event sets. These numbers were what determined where the cut was placed in 1999 (at $\pm 1.25\sigma$ around the mean). Cuts based on the signal MC μ and σ were applied to the signal and background MC event sets. Cuts based on the normalization MC μ and σ were applied to the normalization MC event sets. Cuts based on the normalization data μ and σ were applied to all data. The '97 values included here were used for the '99-style cut studies.

magnet in the spectrometer. If the particle was a muon, which should have only minimum ionized through the calorimeter, this ratio would have been very close to zero because only a small portion of the particle's momentum would have been lost in the CsI. Electrons would have had a value near one for this parameter, because we expected an electron to lose all of its energy and completely stop within the calorimeter crystal. Pions, with their wide range of energy deposition (see Figure 2.17), had a range of values for $\frac{E}{p}$ between zero and one. Figure 3.24 shows the $\frac{E}{p}$ distributions for all three types of particles. In the top plot, one can see the muons peaked towards zero since they only minimum ionized in the calorimeter. These events were taken from a special muon-only run. The center plot also has a peak at zero from pions that minimum ionized, but it has an additional rounded distribution centered around 0.4 from pions that interacted strongly with the CsI nuclei. These pions were from K3pi data. The electrons in the bottom plot peaked around one, as expected. They were from properly reconstructed Ke3 data events. For search modes, this parameter was called `eop_elec` for "electrons" and `eop_mu` for "muons". For normalization reconstruction, `eop_elec` held the $\frac{E}{p}$ for the pion depositing the larger energy cluster in the calorimeter and `eop_mu` was for the softer pion.

- **prob_e (prob_mu)** - from the energy deposited in the TRDs, a confidence level (CL) was calculated that the track was a result of a charged pion or muon, and not an electron. These parameters were calculated by the TRDANA routine, which used the code written during the TRD offline calibration[39]. A value close to zero in these parameters meant the track probably was caused by an electron. Both `prob_e` and `prob_mu` were cut on at 0.04 in the '97 analysis, with `prob_e` (the electron track) required to be smaller and `prob_mu` (the muon track) required to be greater. The behavior and calibration of the TRDs differed from '97 to '99 however, and so selecting what to do in '99 was not a trivial task. In addition, the MC simulation of the TRDs for '99 was incorrect. Since the cut on `prob_mu` was so loose in '97, it was decided that the best course of action would be to optimize the `prob_e` cut and just drop the `prob_mu` cut (and its associated cut on the number of TRD hits along the muon track `npln_m`). To study the cut placement for `prob_e`, data in trigger 2 was used to calculate the electron acceptance and pion contamination versus the cut

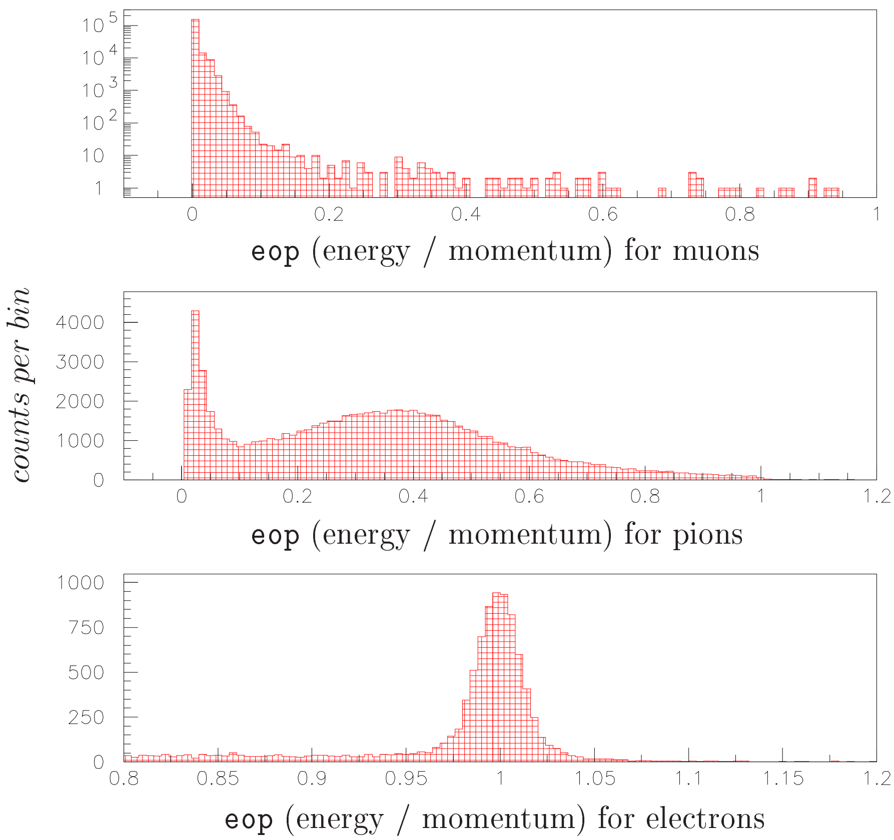

Figure 3.24 Histogram of the energy (in the calorimeter) to momentum (in the spectrometer) ratio ($\frac{E}{p}$) for muons (top plot), pions (middle plot) and electrons (bottom plot). The minimum ionizing muons peak towards zero, the pions peak near zero (minimum ionize) or peak around 0.4 (strong interactions), and the electrons peak around one since they deposited all of their energy in the calorimeter. All distributions are from real data. This information was kept in variables `eop_elec` and `eop_mu`. Note the log scale on the muon plot Y-axis.

value of this TRD parameter (see Figure 3.25). For the electron acceptance, trigger 2 data events were reconstructed as good Ke3 events; the "pion" track had to look like a pion, and the "electron" track had to look like an electron, in the calorimeter. The percentage of events passing was recorded as `prob_e` was stepped from 0. to 0.1. This is the number plotted on the Y-axis in Figure 3.25. For pion contamination, trigger 2 data events were reconstructed as good K3pi events; a good `pi0_mass`, `K_mass`, and `pt2` value was required. The percentage of these events passing was recorded as `prob_e` was stepped, and the results are plotted on the X-axis of Figure 3.25. We wanted good electron acceptance, so we decided to choose the cut such that 98% of the electrons were kept. Placing the cut at "0.192" for '99 data gave an electron acceptance of 98.01±0.04% and a pion contamination of 8.35±0.09%. Placing the cut at "0.049" for the '97 data gave an electron acceptance of 97.99±0.06% and a pion contamination of 4.71±0.05%. The higher pion contamination in '99 versus '97 was not a major concern because this cut was mainly against K3pi background, which naturally reconstructed at a lower `K_mass` than the signal box.

- **fuse3e** - a rating given to the shape of the "electron" cluster using the FUSE3X3 routine. FUSE3X3 takes the center crystal and the eight surrounding crystals of a reconstructed cluster and compares the energy distribution in them crystal-by-crystal to a sample electron/photon energy distribution in a library. The rating given is proportional to the square of the difference between the width (sigma) of the library's

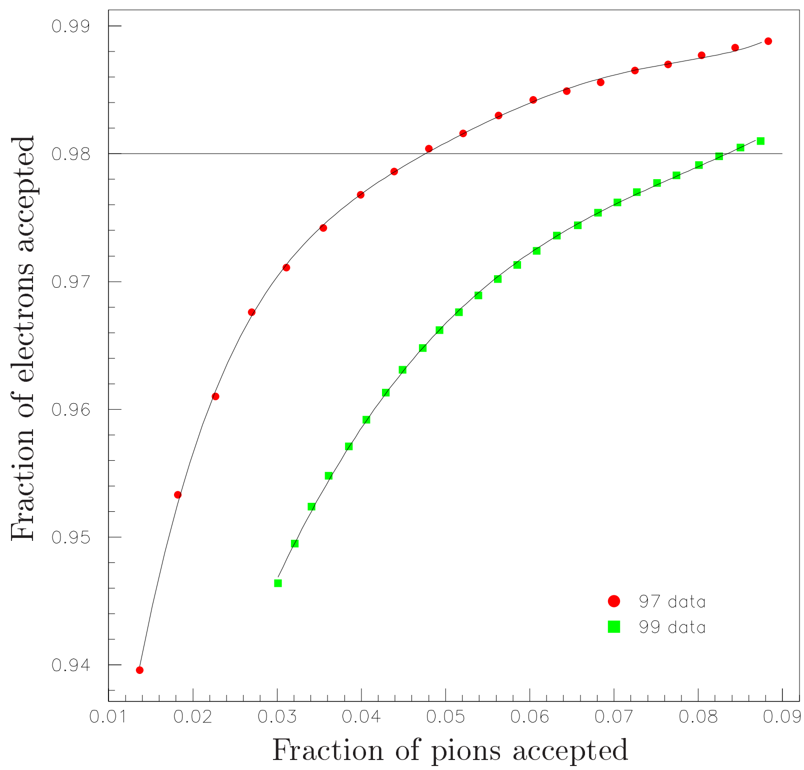

Figure 3.25 TRD electron acceptance versus pion contamination curves for KE3 events in '97 and '99 Trigger 2 data. Cuts were selected to result in a 98% electron acceptance. For '97 the cut value is 0.049 and for '99 the cut value is 0.192.

energy distribution and the width of the reconstructed energy distribution. A true electron or photon should have a rating of 4 or less. Figure 3.26 shows the distribution of this variable for an electron versus that for a pion. We required that `fuse3e` be less than or equal to 4.

In comparing '99 normalization data to MC, we saw a longer tail in the MC distribution (Figure 3.27). If one multiplied the MC values of `fuse3e` by a factor of 0.84, the distributions agreed better (see Figure 3.28). Hence, whenever '99 MC was cut on in this parameter, the value was first corrected by multiplying by 0.84. The '97 MC was not changed for this variable (see Figure 3.29).

- **fuse3p1 (fuse3p2)** - a rating given to the shape of the first (or the second) of the photon clusters from the neutral pion using the FUSE3X3 routine (see `fuse3e` for a description of FUSE3X3). A true electron or photon should have had a rating of 4 or less. See Figure 3.30 for an example distribution of this variable.

In comparing '99 normalization data to MC, we saw a longer tail in the MC distribution. If one multiplied the MC values of `fuse3p1` and `fuse3p2` by a factor of 0.86, the distributions agree better (see Figure 3.31). Hence, whenever '99 MC is cut on in these parameters, the value is first corrected by multiplying by 0.86. The '97 MC was not changed for any of these variables (see Figure 3.32).

- **mu_flag** - a logical tag that indicates if there was at least one hit in any of the muon banks. Such a requirement was part of Trigger 7 (our search trigger), but

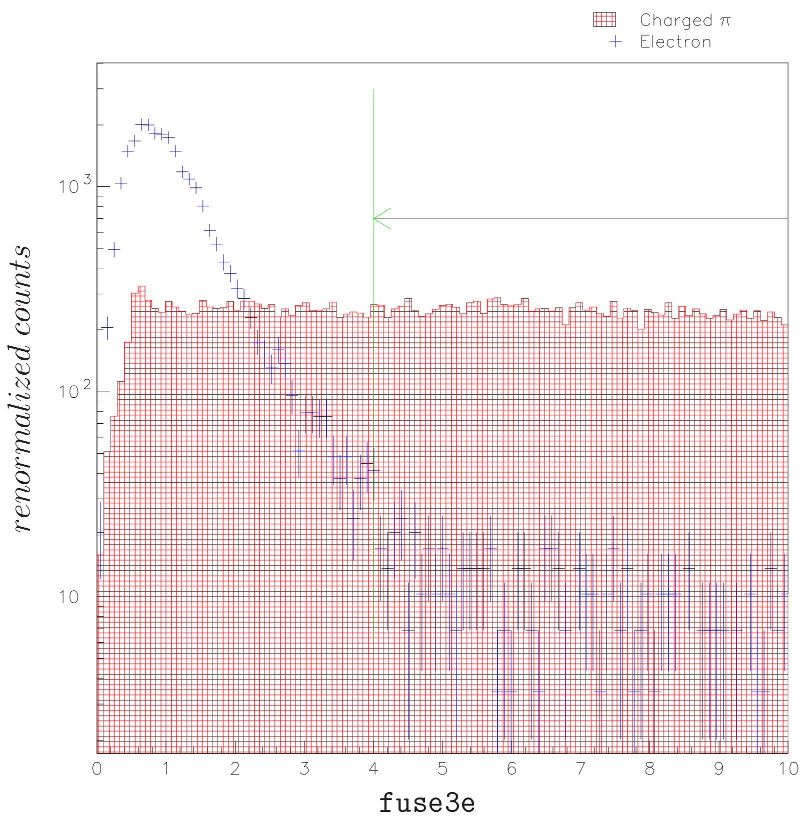

Figure 3.26 Histograms of the rating given to the shape of the "electron" cluster in the calorimeter (`fuse3e`). For the histogram in red, the particle was actually a pion (normalization K3pi 1999 data events were used). For the crosses in blue, the particle was an electron ('99 trigger 2 Ke3 data events were used). The signal selection cut is looking for electrons, and so accepts events with a value equal to or below 4. The histograms have been renormalized to fit on the same plot. Note the log scale on the Y-axis.

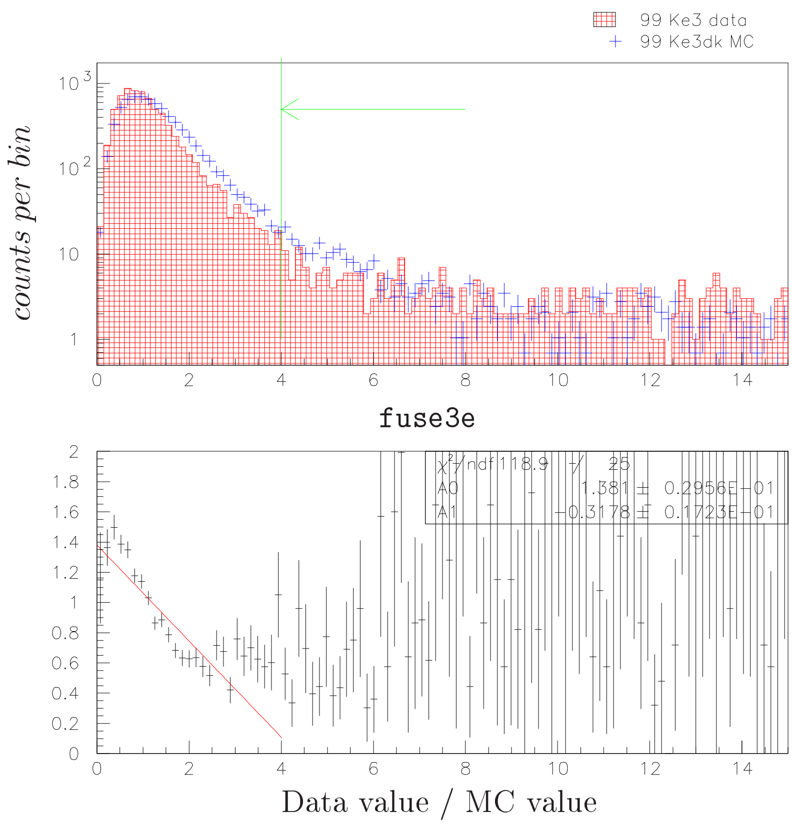

Figure 3.27 Histograms of the rating given to the shape of an "electron" cluster in the calorimeter (`fuse3e`). Ke3 data and MC for 1999 are shown here. Data and MC are not quite in agreement, with the tail of the MC being longer. This disagreement could be corrected somewhat by multiplying the MC by 0.84. The distribution is similar to that for photons (`fuse3p1` and `fuse3p2`). The histogram has been normalized to the number of data events present. Note the log scale on the Y-axis.

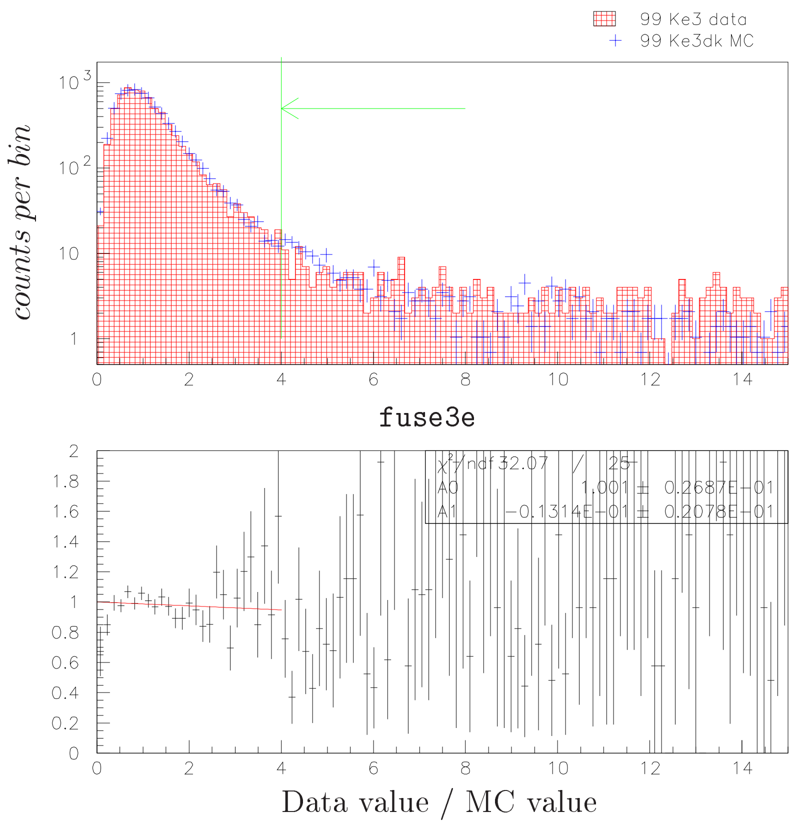

Figure 3.28 Histograms of the rating given to the shape of an "electron" cluster in the calorimeter (`fuse3e`), with MC values corrected by multiplying by 0.84 (i.e. rescaling the MC X-axis). Normalization data and MC for 1999 are shown here. Data and MC are in better agreement than without the correction. The signal selection cut accepted events equal to or below a value of 4 after applying the correction factor to MC when necessary. The histogram has been normalized to the number of data events present. Note the log scale on the Y-axis.

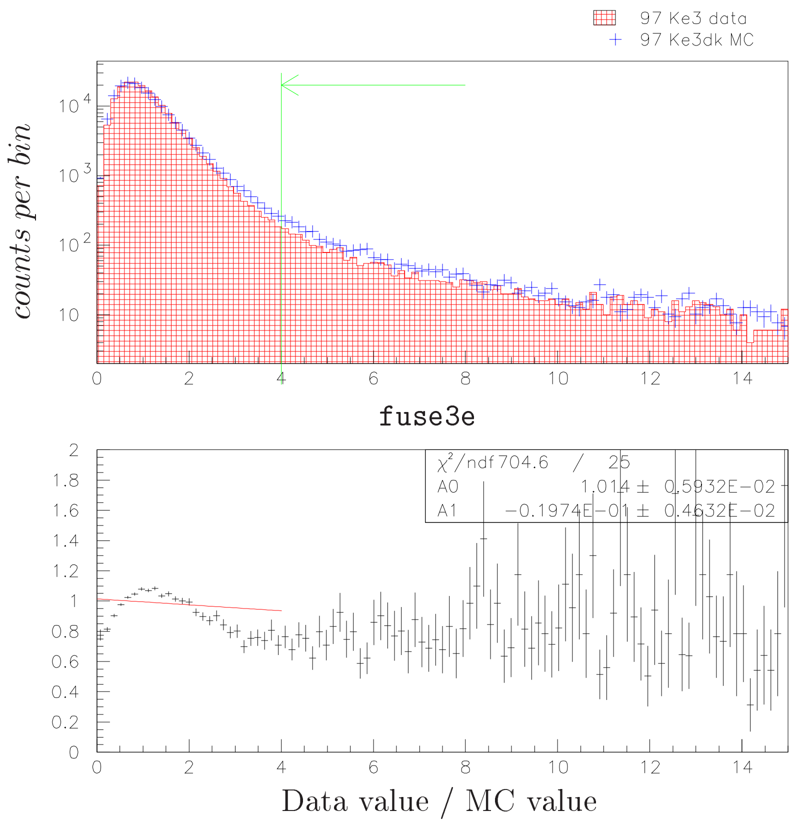

Figure 3.29 Histograms of the rating given to the shape of an "electron" cluster in the calorimeter (`fuse3e`). Ke3 data and MC for 1997 are shown here. Data and MC are in reasonable agreement, and no corrections were made. The histogram has been normalized to the number of data events present. Note the log scale on the Y-axis.

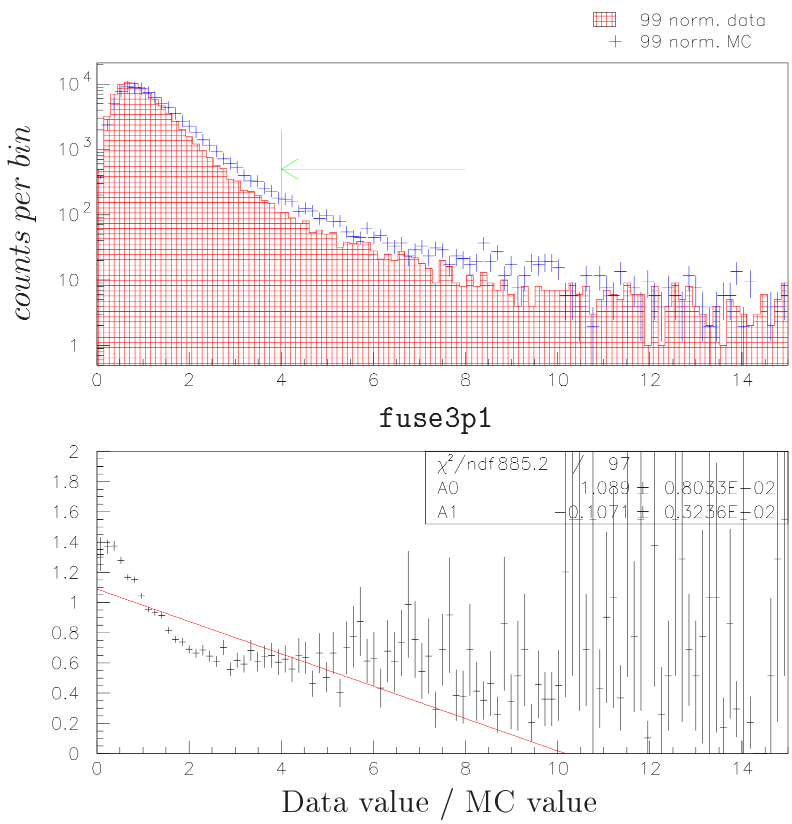

Figure 3.30 Histograms of the rating given to the shape of a "photon" cluster in the calorimeter (`fuse3p1`). The parameter `fuse3p2` distribution looks the same as for `fuse3p1`. Normalization data and MC for 1999 are shown here. Data and MC are not quite in agreement, with the tail of the MC being longer. This disagreement could be corrected somewhat by multiplying the MC by 0.86. The histogram has been normalized to the number of data events present. Note the log scale on the Y-axis.

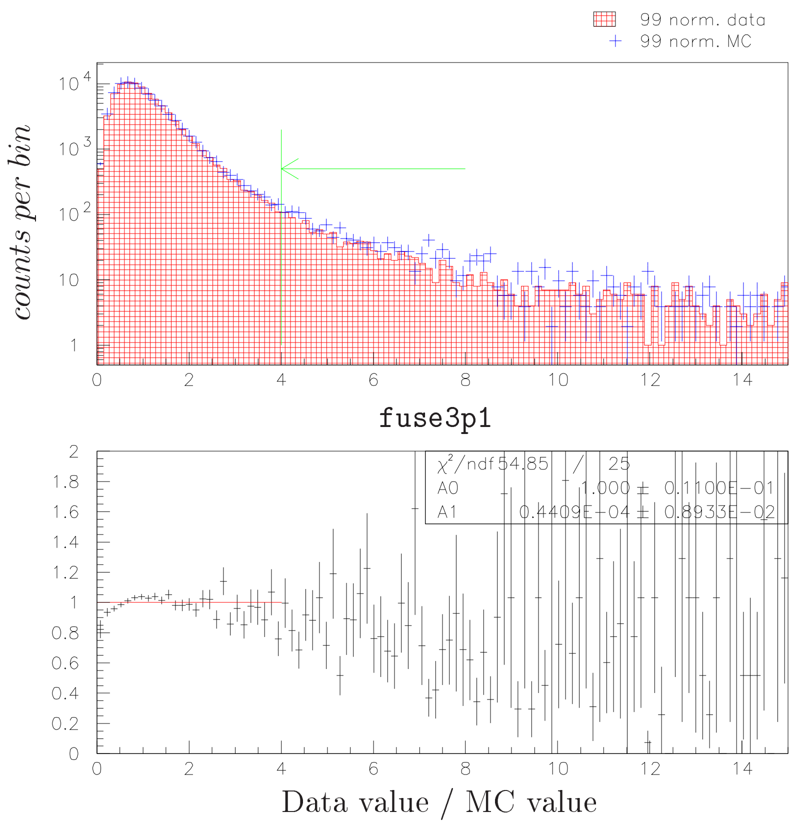

Figure 3.31 Histograms of the rating given to the shape of a "photon" cluster in the calorimeter (`fuse3p1`), with MC values corrected by multiplying by 0.86 (i.e. rescaling the MC X-axis). The parameter `fuse3p2` distribution looks the same as for `fuse3p1`. Normalization data and MC for 1999 are shown here. Data and MC are in better agreement than without the correction. The signal selection cut accepted events with a value of 4 or lower after applying the correction factor to MC when necessary. The histogram has been normalized to the number of data events present. Note the log scale on the Y-axis.

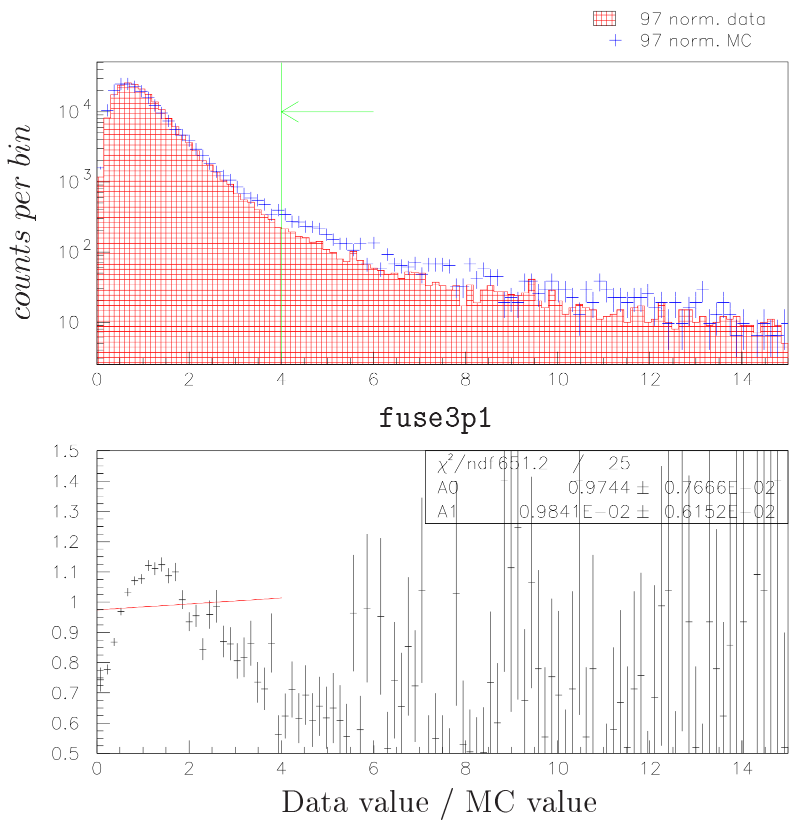

Figure 3.32 Histograms of the rating given to the shape of a "photon" cluster in the calorimeter (`fuse3p1`). The parameter `fuse3p2` distribution looks the same as for `fuse3p1`. Normalization data and MC for 1997 are shown here. Data and MC are in reasonable agreement, and no corrections were made. The histogram has been normalized to the number of data events present. Note the log scale on the Y-axis.

not Trigger 2 (the normalization trigger). We sometimes needed to compare trigger 2 data with trigger 7 data, so this cut was applied to simulate the trigger 2 data having passed the trigger 7 logic.

- **nmumats** - the number of track matches in the muon banks. This parameter was calculated by a subroutine called MUMATCH. MUMATCH took downstream segment of each charged track projected it onto each plane of the muon banks. It then looked for the hit nearest to the projected location in each plane and calculated the distance between them. An expected scattering distance at each plane was also calculated based on the momentum of the track and known multiple scattering properties. If the difference between the projected track and the hit was less than the expected scattering distance, the track was considered "matched" in that plane and the variable `nmumats` was incremented by 1. Since there were three muon banks, the highest possible value for `nmumats` was three.

When the code to do muon track matching was first written, the criteria for a "match" was too strict, causing an unacceptable loss in signal when a `nmumats` cut was applied. This is why there was no cut on the `nmumats` parameter for the '97 analysis. In March of 2001, KTeV collaborator Jason Hamm analyzed data from muon-only runs (with the magnet on) to study the problem. Events were selected which hit locations in the first (vertical) muon bank where the scintillating paddles overlapped. This meant that the track location was known at that plane to within

$2cm$. During normal running, the resolution of the muon banks is half of a paddle width. Then the difference in extrapolated track location and nearest muon bank hit was compared to the corresponding value in Monte Carlo events at the second vertical plane. Since the amount of scattering the muon undergoes is momentum dependent, events were grouped by momentum when plotting the distance between the projected hit location and the recorded hit location. The results are shown in Figures 3.33 and 3.34. As expected, the distributions got narrower as the muons got harder, but the tails persisted all the way out to $20cm$ even at the highest momenta. We re-wrote our own version of MUMATCH based on this study. The maximum expected scattering for a track was calculated to be a scattering coefficient (X_{scat}) divided by the absolute momentum of the track.

$$DX_\mu = \frac{X_{scat}}{\mid \mathtt{pmu} \mid} \tag{3.2}$$

Originally the coefficients were 0.968 for the first bank and 1.67 for the second and third banks. We changed them to 4 for the first bank and 7 for the second and third banks. Since these numbers were based on the $2cm$ resolution, we then increased DX_μ by the width of half a paddle. If DX_μ was less than $0.2m$ at the first bank, then DX_μ is set equal to $0.2m$. At the second bank, the DX_μ minimum was $0.3m$. If DX_μ was greater than the difference between the projected track and the muon bank hit, it was considered matched and `nmumats` was incremented. Therefore, the requirements for a muon bank match were still momentum dependent, but always accepted events falling within $20cm$ at the first muon plane, and within $30cm$ at the second and third planes.

- **pmu** - the momentum of the reconstructed "muon" track. The lower the momentum of a particle, the more it will multiple scatter. We wanted to extrapolate the downstream spectrometer track to hits in the muon banks, but there was a lot of filtering material in between them. By requiring the "muon" to have a minimum momentum, the particle should have been able to travel through all the filtering without an unreasonable amount multiple scattering. A cut was made on the absolute value of `pmu` at $10GeV/c$ in the 1997 analysis and at $7GeV/c$ in the 1999 analysis. The same studies that improved the simulation of the `nmumats` variable also indicated the cut on |`pmu`| could be lowered in '99 to improve acceptance without hurting efficiency.

3.4.6 Signal Parameters

- **chrgmass** - mass of the system reconstructed from only the charged particles detected, by routine T3FVTX[38]. The signal had a very clear kinematic cutoff at $0.364GeV/c^2$ (see Figure 3.36). Such a cut was applied to the normalization mode as well without any loss of efficiency, as shown by Figure 3.37.

- **xbrdist (ybrdist)** - (AAP) the distance in X (or Y) between the projected upstream segment of the electron track and the nearest cluster identified as coming from a photon. The trajectory of the electron track before passing through the anal-

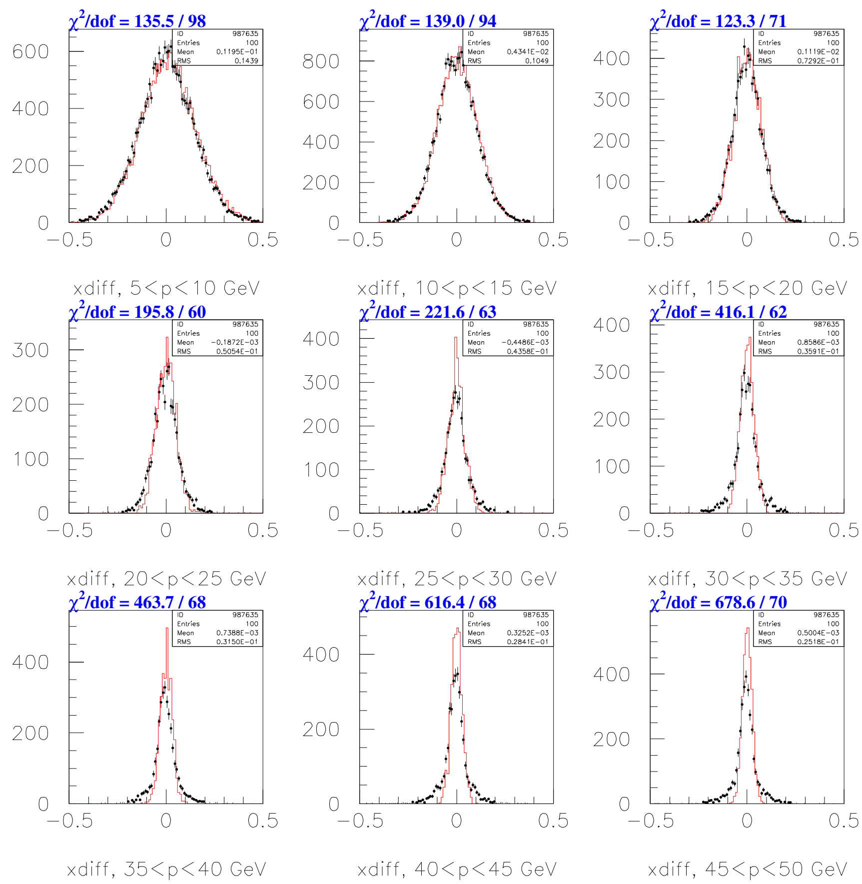

Figure 3.33 Distributions of the difference between the projected muon track and the known hit location for some 1997 muon-only data, separated as a function of momentum. Data are the black points and the MC is the red histogram. As the momentum bin got higher, the original simulations failed to account for the data's sidebands. This plot was generated by Jason Hamm.

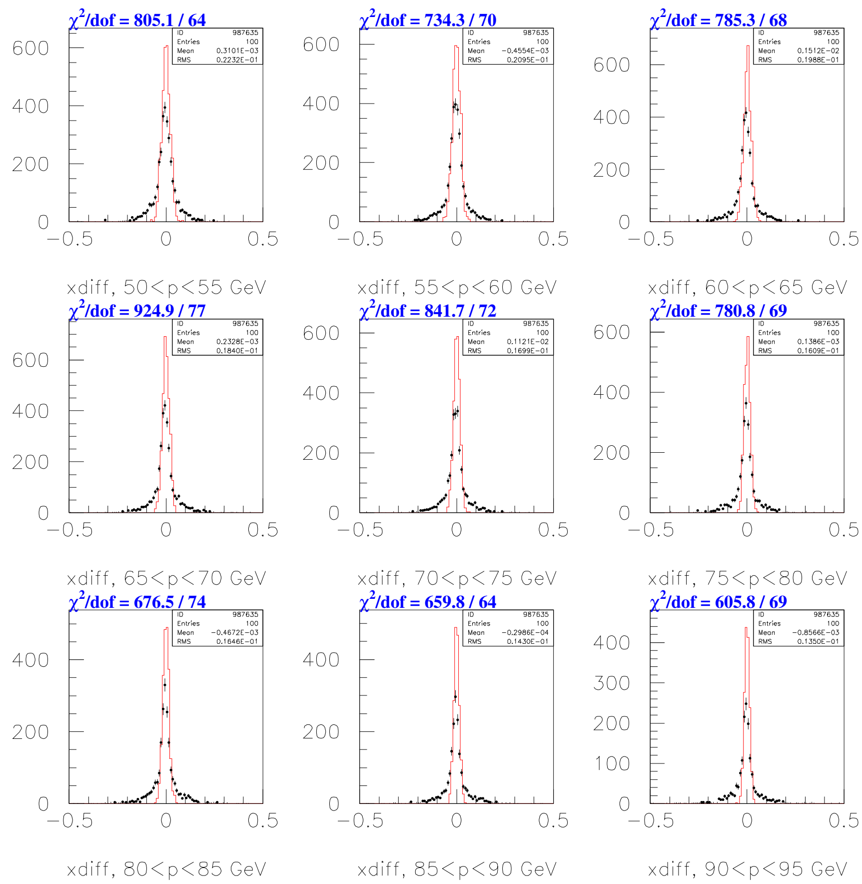

Figure 3.34 Distributions of the difference between the projected muon track and the known hit location for some 1997 muon-only data, separated as a function of momentum. Data are the black points and the MC is the red histogram. Note that even at the highest momenta, there were still tails going out to $20cm$. This plot was generated by Jason Hamm.

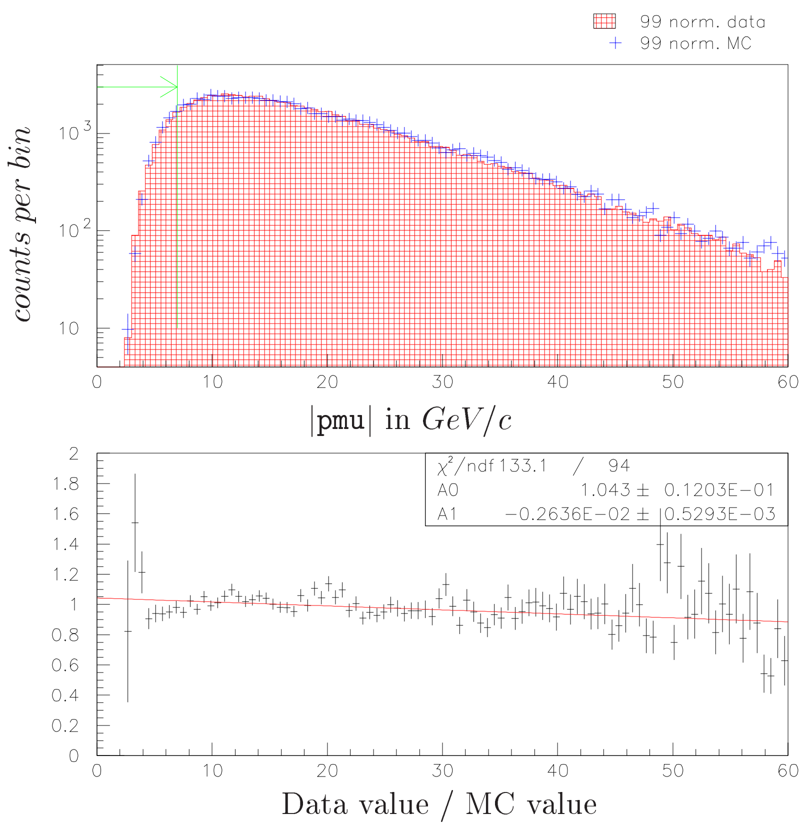

Figure 3.35 Histograms of the absolute momentum of the muon track (|pmu|) for '99 normalization data and MC. No cut was made on this variable in the normalization mode, but for the signal mode, the muon track was required to have a minimum momentum of $\pm 7 GeV/c$ to reduce multiple scattering in the shielding material. In 1997 we used a tighter cut of $\pm 10 GeV/c$. Studies done since the 1997 run indicated we could loosen the cut down to 7 to improve acceptance without losing efficiency. The signal histogram has been normalized to the number of Ke3 events present.

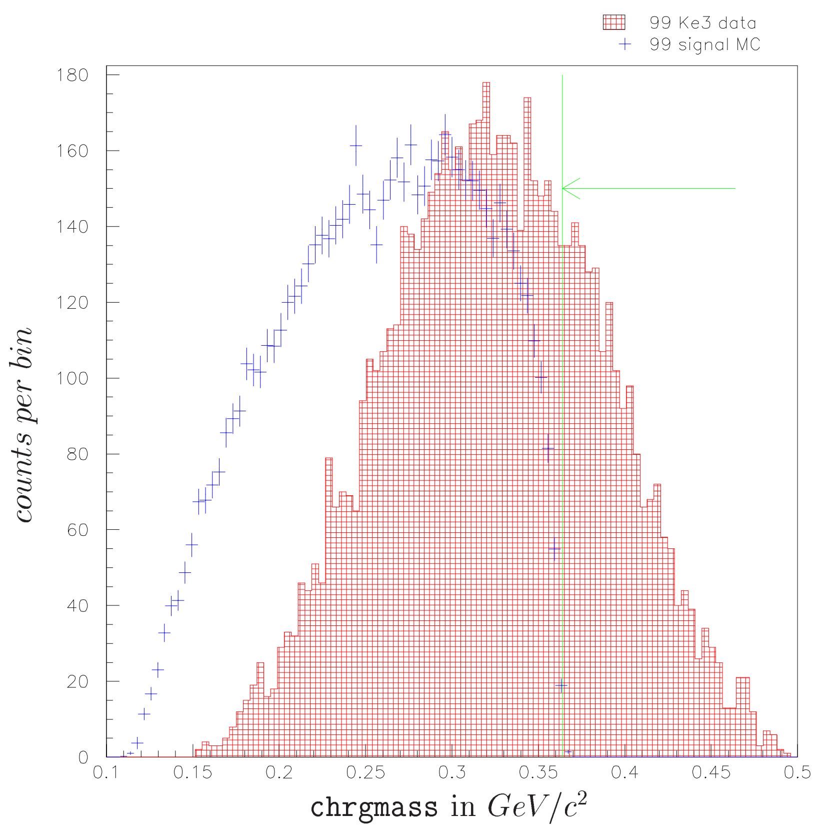

Figure 3.36 Histograms of the reconstructed mass of the charged system (`chrgmass`). Ke3 MC is shown by the red histogram and signal MC are shown by the blue crosses. The signal showed a clear cutoff at $0.364 GeV/c^2$. The signal histogram has been normalized to the number of Ke3 events present.

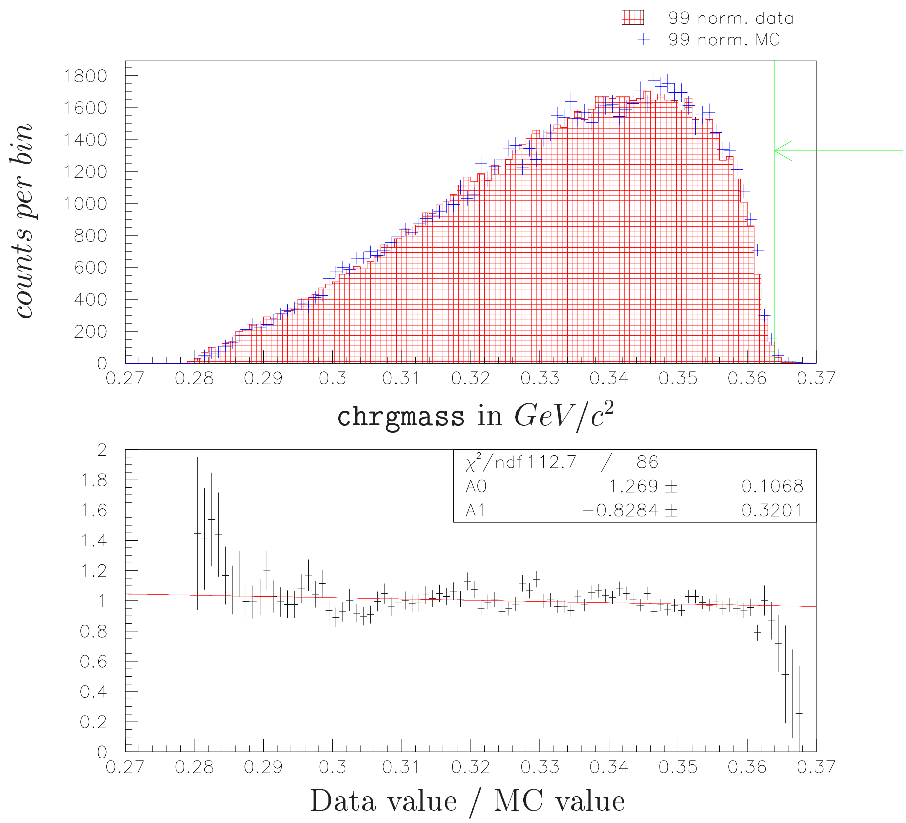

Figure 3.37 Histograms of the reconstructed mass of the charged system (`chrgmass`) for 1999 normalization data and MC. As can be seen, data and MC were in good agreement. The signal and normalization selection cuts accepted events below a value of $0.364 GeV/c^2$. The histogram has been normalized to the number of data events present.

ysis magnet (or "upstream" of the magnet) was extended mathematically to a point on the face of the calorimeter. If the detected photon was a result of bremsstrahlung radiation, it most likely occurred during the passage of the electron through the vacuum window, upstream of the first drift chamber. Such a photon would follow closely the trajectory of its parent electron, and being unaffected by the analysis magnet, would land close to the projected point. Hence, the smaller this number, the more likely the photon was a result of bremsstrahlung radiation off the electron and not from the decay of a neutral pion. A cut was made on these parameters in the 1997 analysis only - an event was removed if |`xbrdist`| and |`ybrdist`| were less than $5mm$. It was later decided that cutting on these parameters was not efficient and so they were not used in the '99 analysis. See Figure 3.38 for an example distribution.

- **pp0_kine** - a kinematic variable involved in reconstructing a $K_L \rightarrow \pi^+\pi^-\pi^0$ (K3pi) event. For this variable, the event was assumed to be a K3pi decay. The two track-cluster combinations were each associated with a charged pion mass and the parent particle was assumed to be a kaon. Using such assumptions, a number proportional to the square of the longitudinal momentum (`p`) for the remaining neutral pion (`p0`) in the kaon rest frame was calculated kinematically (`kine`) using the following equation:

$$\texttt{pp0\_kine} = \frac{(M_K^2 - M_{\pi^0}^2 - M_{\mu e}^2)^2 - 4M_{\pi^0}^2 M_{\mu e}^2 - 4M_K^2 p_{T(\mu e)}^2}{p_{T(\mu e)}^2 + M_{\mu e}^2} \tag{3.3}$$

where M_K is the PDG value for the mass of the kaon, M_{π^0} is the PDG value for the

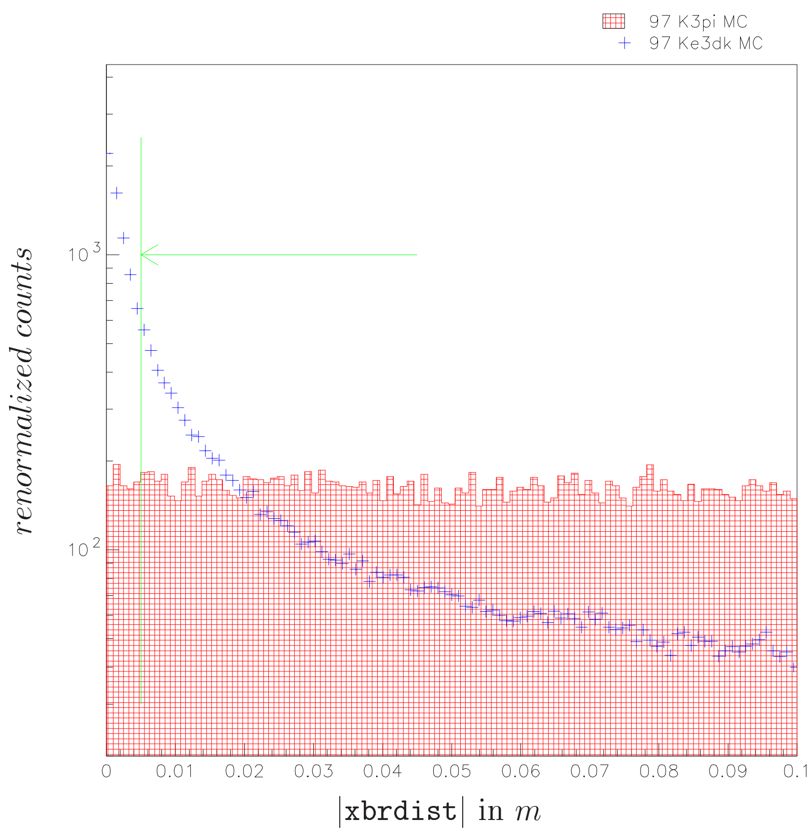

Figure 3.38 Histograms of the magnitude of the distance in X between the upstream electron track projected onto the calorimeter and the nearest hard cluster (|`xbrdist`|) for a '97 K3pi MC pion and a '97 Ke3dk MC electron. Since electrons produce bremsstrahlung radiation, the Ke3dk electron distribution peaks up at zero. In contrast the pion distribution is flat, showing no bremsstrahlung radiation. This distribution in Y (|`ybrdist`|) is similar. The signal and normalization selection cuts rejected events below an absolute value of $5mm$ in '97 only. The histograms have been normalized to fit on the same plot.

mass of the neutral pion, $M_{\mu e}$ is the mass of the charged system only, and $p_{T(\mu e)}$ is the transverse momentum of the charged system only. The value for `pp0_kine` was positive for a true K3pi event. A negative value indicated an incorrect assumption about the charged particle masses or the parent particle mass was made. For the 1997 analysis a cut was made at -0.001 to maximize signal acceptance. For the 1999 analysis, we found we could efficiently reduce background by tightening this cut to -0.06. See Figures 3.39 and 3.40 to see the difference between signal and background in this variable.

- **ke4_kine** - a kinematic variable involved in reconstructing a $K_L \rightarrow \pi^0\pi^\pm e\nu_e$ (Ke4) event. For this variable, the event was assumed to be a Ke4 decay. The photon combination for the neutral pion and the tagged "electron" track-cluster combination were kept, the tagged "muon" track-cluster combination was associated with a charged pion mass, and and the parent particle was assumed to be a kaon. Using such assumptions, the magnitude of the momentum for the undetected neutrino in the kaon rest frame was calculated kinematically using the following equation:

$$\texttt{ke4\_kine} = \frac{M_K^2 - M_{\pi^\pm e^\mp \gamma\gamma}^2}{2M_K} \tag{3.4}$$

where M_K is the PDG value for the kaon, and $M_{\pi^\pm e^\mp \gamma\gamma}$ is the total reconstructed mass. The value for `ke4_kine` was positive for a true Ke4 event. A negative value indicated a real neutrino could not have been produced by this event and therefore the reconstruction assumption was incorrect. This cut eliminated most of the Ke4

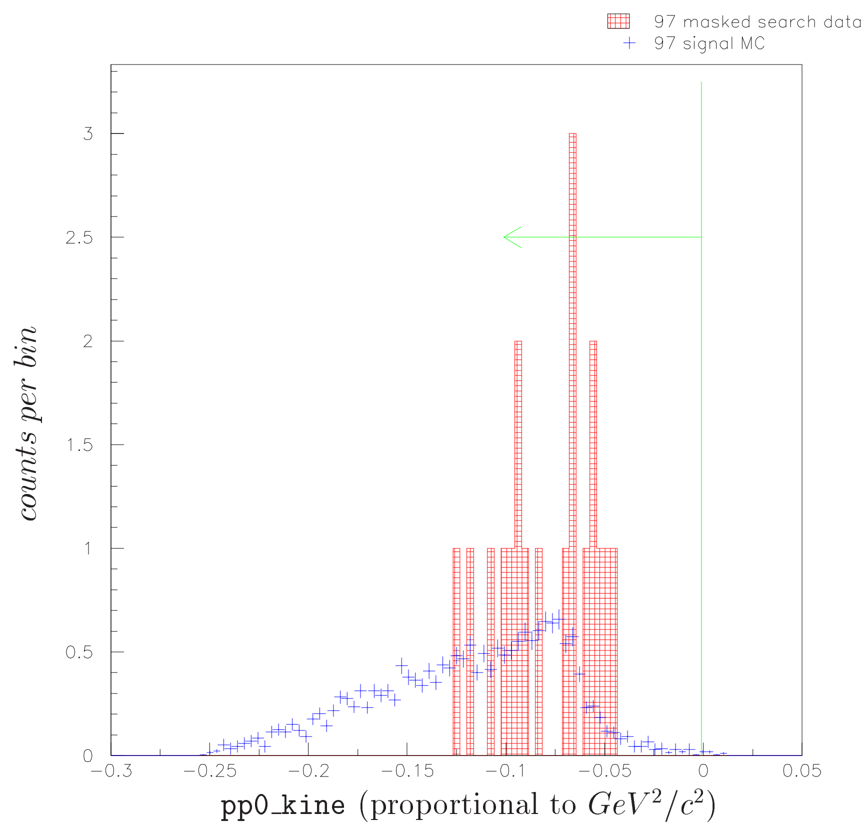

Figure 3.39 Histograms of `pp0_kine`, a value proportional to the longitudinal momentum squared of the neutral pion using the charged system information and assuming a K3pi decay. All other signal selection cuts have been applied. 1997 MC signal events are shown in blue and 1997 data search events (with possible signal events masked out) are shown in red. The '97 signal selection cut accepted events below -0.001, and for this plot a cut at +0.01 had already been made. We chose this cut value based on maximum signal acceptance. The histogram has been normalized to the number of data events present.

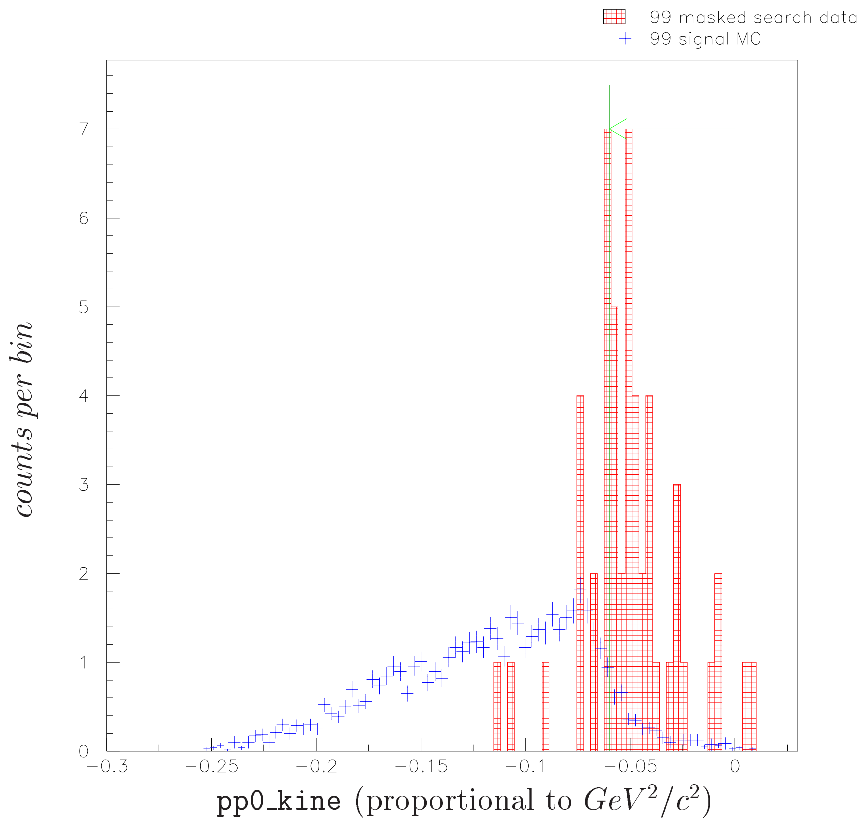

Figure 3.40 Histograms of pp0_kine, a value proportional to the longitudinal momentum squared of the neutral pion using the charged system information and assuming a K3pi decay. All other signal selection cuts except the TRD CL (prob_e) had been applied. 1999 MC signal events are shown in blue and 1999 data search events (with possible signal events masked out) are shown in red. The '99 signal selection cut accepted events below -0.06, and for this plot a cut at +0.01 had already been made. This cut value was chosen based on background rejection. The histogram has been normalized to the number of data events present.

events and was also effective in removing K3pi events. See Figure 3.41 to see the difference between signal and background in this variable.

- **sgp0kine** - a kinematic variable involved in reconstructing the signal event. For this variable, the event was assumed to be a Kpi0me decay. Information about the charged daughter particles (the "electron" and "muon") was used, and the parent particle was assumed to be a kaon. The square of the longitudinal momentum for the neutral pion in the kaon rest frame was then calculated kinematically using the following equation:

$$\texttt{sgp0kine} = \frac{(M_K^2 - M_{\pi^0}^2 - M_{\mu e}^2)^2 - 4M_{\pi^0}^2 M_{\mu e}^2}{4M_K^2} - p_{T(\mu e)}^2 \qquad (3.5)$$

where M_K is the PDG value for the mass of the kaon, M_{π^0} is the PDG value for the mass of the neutral pion, $M_{\mu e}$ is the mass of the charged system only, and $p_{T(\mu e)}$ is the transverse momentum of the charged system only. The value for `sgp0kine` was positive for a true Kpi0me signal event. MC signal events indicated a negative tail for this variable because of resolution (see Figure 3.42), so the cut was placed at -0.002. While this cut may not seem very effective in looking at Figure 3.42, there were events (6 in the 1999 data) that it removed that fell in our `pt2` versus `K_mass` plot that would not have been removed by any of our other cuts.

- **K_mass** - the total reconstructed mass of the event assuming it was the decay of a single particle at the charged vertex. This included the contribution of both charged and neutral decay products. If the event was a Kpi0me decay, then this

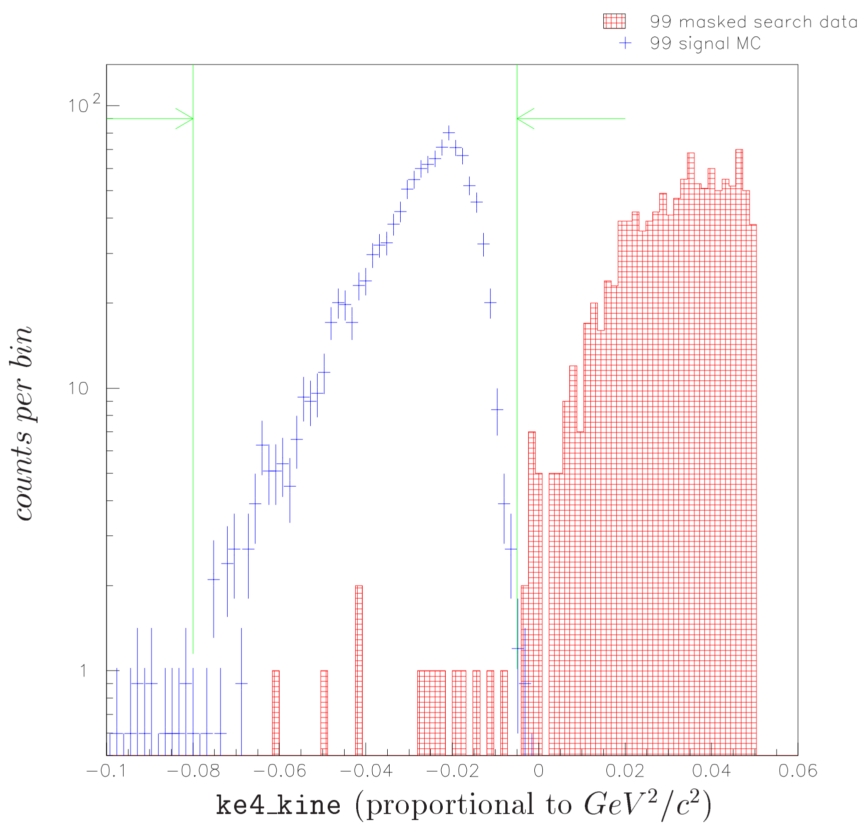

Figure 3.41 Histograms of the magnitude of the momentum squared of the missed neutrino using all daughter particle information and assuming a Ke4 decay (`ke4_kine`). All other signal selection cuts except the TRD CL had been applied. 1999 MC signal events are shown in blue and 1999 search data events (with possible signal events masked out) are shown in red. The signal selection cut accepted events below -0.005 and above -0.08, and for this plot a cut at +0.05 had already been made. The histogram has been normalized to the number of Ke4 events present.

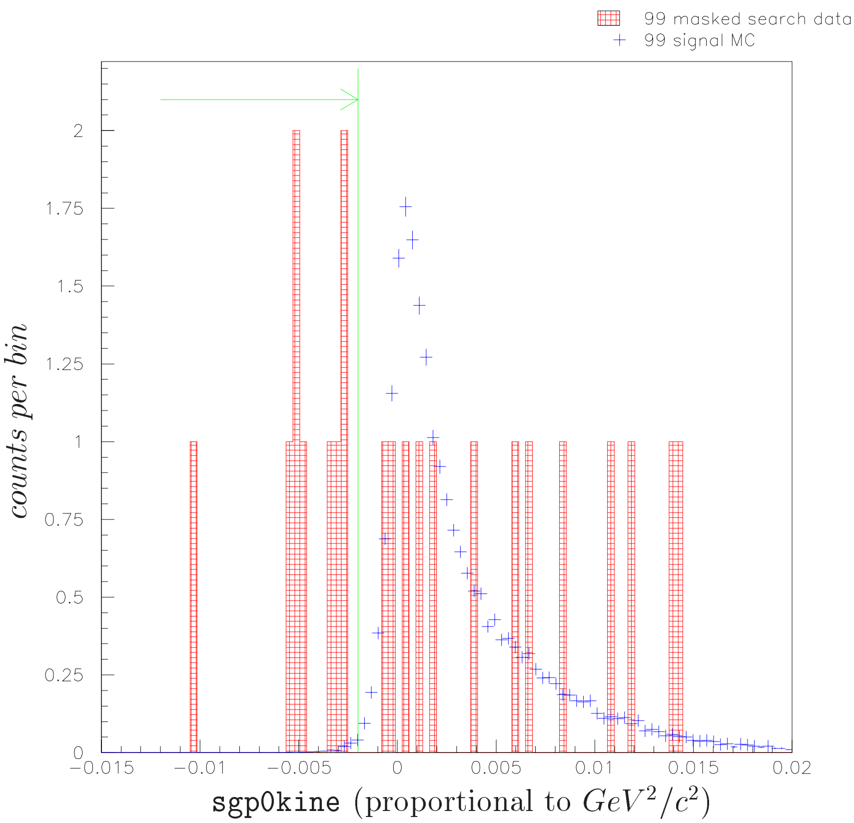

Figure 3.42 Histograms of the longitudinal momentum squared of the neutral pion using the charged system information and assuming a signal Kpi0me decay (`sgp0kine`). 1999 search data background events (with possible signal events masked out) are in red, and MC signal (Kpi0me) events are shown in blue. All signal selection cuts except those on `sgp0kine` and `prob_e` were applied. The signal selection cut accepted events above -0.002.

value would have equaled the mass of a neutral kaon (K_L). Figure 3.43 shows the distribution of `K_mass` for 1999 normalization mode data and MC. This variable was one dimension of the signal box, with the box boundaries being at $0.48767 GeV/c^2$ and $0.50767 GeV/c^2$.

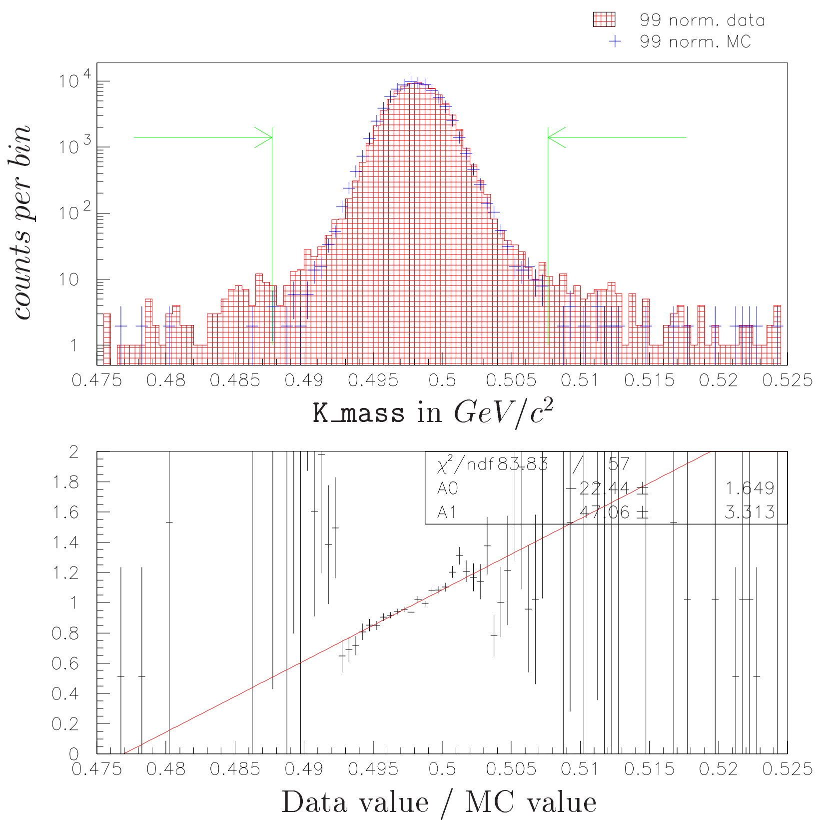

Figure 3.43 Histograms of mass of the parent particle (`K_mass`) for 1999 data and MC normalization events passing all other normalization selection cuts. The distributions agree and are peaked around the kaon mass of $0.498 GeV/c^2$. Cuts were placed at 0.48767 and $0.50767 GeV/c^2$ (the edges of the "signal box"). Note the log scale on the Y-axis.

3.5 Universal Cuts

Many of the loosest cuts were applied to both data sets early on in the analysis process. These "universal cuts" are described here.

3.5.1 KTSPILL and KTeVAna Subroutine cuts

Events were cut if:

KTSPILL was bad
GET_VETO returned in error
KTCLUS returned in error
T3MATCH returned in error
KTTDC returned in error
T3MASS returned in error
T3FPI0 returned in error

Variables that are in all capital letters represent subroutines within the KTeVAna program and were discussed in Section 3.4.2.

3.5.2 Crunch cuts

The intention of the crunch level of cuts was to eliminate badly-recorded events and events that had no chance of reconstructing to a decent Kpi0me. Most of these cuts had to do with analysis routines returning to the main program in an error state, plus a few very basic characteristic or kinematic requirements.

An event was rejected at this level if any of the following were true:

# of clusters, hard or soft, < 4

\# hard clusters < 3 or > 6
\# preliminary tracks < 2
no vertex candidate
`vtx_chi` > 100
Z vertex (`vtxz`) not between target ($0m$) and vacuum window ($165m$)
m_{π^0} (`pi0_mass`) $< 110MeV$ or $> 160MeV$
$\frac{E_e}{p_e}$ (`eop_elec`) < 0.25 or > 1.75
p_t^2 (`pt2`) $> 0.05\frac{GeV^2}{c^2}$

The variables in `typewriter font` are actual parameter names used in the KTeVAna program code. Please see Section 3.4 for their definitions. The other cuts listed above are discussed here.

Clusters were defined as described in Section 2.4.2 with hard(ware) clusters being energy deposits found online with more than $1GeV$ of energy and being in-time with the event trigger. Soft(ware) clusters were found offline with looser energy requirements and no timing requirement. The two photons and one electron from the signal decay should have created hard clusters with the looser cuts allowing extra clusters from accidentals to be present for a total of up to six clusters at this level. With the signal muon creating a soft cluster, a minimum total (hard or soft) cluster number of four was required.

A rough track-finding routine looked for preliminary tracks with loose requirements. If the event had any chance of being reconstructed well as a signal event, at least two charged tracks should have been found in the spectrometer at this level. Later in the analysis program a more exhaustive, CPU intensive track search was done. The preliminary track finding routine was intended to save this processing time by cutting events that obviously had fewer than two tracks.

The segments of the two tracks which were upstream of the analysis magnet were projected back towards the target. If they came reasonably close to meeting at a point, then that point was chosen as the vertex for the decay of the parent particle. Otherwise the event was cut. Also, a χ^2 related to the closeness of the two tracks to the selected vertex was returned as the variable `vtx_chi`. If the fit was very bad (`vtx_chi` > 100), the event was removed. Using the charged track vertex, calorimeter hardware clusters unmatched to tracks were combined in pairs and were reconstructed as having come from the vertex. Assuming these clusters were from the decay of the neutral pion, the pair which best reconstructs to the mass of the neutral pion (m_{π^0}) were tagged. There was then an additional wide quality cut on the pion mass from $110MeV$ and $160MeV$ ($m_{\pi^0} \approx 135MeV$).

As discussed in Section 3.1, one of the charged tracks had to match to a cluster containing less than $1GeV$ of energy. Such a track/cluster combination was assumed (and tagged) to be a muon since it was the particle most likely to minimum ionize in the CsI and we wanted a muon present in our decay. The other charged track/cluster combination was tagged to be an electron. For each of the track/cluster combinations, the ratio was taken of the energy deposited in the calorimeter over the momentum of the associated track ($\frac{E}{p}$). If the $\frac{E}{p}$ of the tagged electron ($\frac{E_e}{p_e}$) was less than 0.25

or greater than 1.75, then it probably was not an electron and the event was cut. Using the vertex found, the decay particles were combined into a parent particle and the parent particle's mass, forward momentum, and transverse momentum were calculated. If the parent particle was a kaon from within the beam then nearly all of its momentum should have been in the forward direction. Hence if the squared transverse momentum (`pt2`) of the parent particle was calculated to be above a loose requirement of 0.05 ($\frac{GeV^2}{c^2}$) at this level, then its reconstruction as a Kpi0me was unreasonable and the event was cut.

3.5.3 Ntuple cuts

The intention of the ntuple level of cuts was to get rid of events that had little chance of reconstructing to a decent Kpi0me. Since the Monte Carlo (MC) events were not put through the "crunch" process, the crunch cut requirements were included again with the ntuple cut requirements and are marked with an dagger (†). Please see the previous section for a description of these cuts. The main difference between this level of cuts and the crunch cuts was that slightly stricter kinematic and characteristic requirements were made.

An event was cut before reaching the ntuple if any of the following was true:

† T3FPI0 returned in error
† number of clusters, hard or soft, < 4
† number hard clusters < 3 or > 5
† number preliminary tracks < 2

† no vertex candidate
† `vtx_chi` > 100
number of final tracks was not equal to 2
both clusters matching to tracks had energy $> 1GeV$
† Z vertex (`vtxz`) not between 0 and vacuum window
† m_{π^0} (`pi0_mass`) $< 110MeV$ or $> 160MeV$
$\frac{E_e}{p_e}$ (`eop_elec`) < 0.8 or > 1.2
† p_t^2 (`pt2`) $> 0.05\frac{GeV^2}{c^2}$

"Final" tracks were required to have in-time hits and reasonable sums-of-distances (equal to the cell diameter, within the range of the resolution) in all four chambers.

If an event passed the ntuple cut level, information about it was saved in a PAW-readable ntuple file consisting of approximately 100 parameters.

3.5.4 Selection cuts

Selection cuts were the final level of cuts made. The number of events which passed these cuts was used in the calculation of signal acceptance and background contamination within the "box". The parameters cut on and the cut values differed some between the different analyses, and so the selection cuts used will be listed in the chapter covering that analysis.

3.6 Background Decays

A crucial part of any analysis is to reduce background as much as possible. To do this we needed to understand how other processes could have looked like a $K_L \rightarrow \pi^0 \mu e$ (Kpi0me) signal decay and how they could have been excluded from the final data sample. There were four backgrounds to the Kpi0me decay, as listed in Table 1.1 at the end of Chapter 1. For the $K_L \rightarrow \pi^+\pi^-\pi^0$ (K3pi) decay mode, the π^0 could have gone to two photons as in the signal mode. A charged pion ($\pi^\pm$) could have decayed into a muon plus a neutrino ($\mu^\pm\nu_\mu$). Alternatively, a charged pion could have interacted very early in its passage through the calorimeter and showered to look like, or *fake*, an electron. If one of the charged pions muon-decayed, and the other faked an electron, the result looked like the particle combination ($\pi^0\mu^\pm e^\mp \nu_\mu$). Since the neutrino would have traveled through the detector without being seen, the final state looked like a Kpi0me decay mode of a K_L. Such an event would have passed triggers and had to be filtered out offline. The K3pi decay mode could have also imitated a signal decay mode in another way. One of the charged pions could have undergone muon decay as before, and the other charged pion could have decayed into a real electron and electron-like neutrino. Missing both neutrinos would have given a Kpi0me-type signature. The electron-mode decay of a pion is suppressed (BR($\pi^\pm \rightarrow e^\pm\nu_e$) / BR($\pi^\pm \rightarrow \mu^\pm\nu_\mu$) $\approx 10^{-5}$), so this type of background was less likely, yet still a concern. In the area around the signal box, the K3pi decay was the primary background before selection cuts were applied.

Several avenues were open for removing K3pi decays from the data set. If one of the charged pions decayed into a muon or electron, the track for that particle would have had a direction change in it from the kinematics of emitting a neutrino. While the spectrometer information was not fine-grained enough to pinpoint the location of a daughter particle's in-flight decay vertex, the track reconstruction was accurate enough that an in-flight decay could have resulted in a mis-match of segments at the magnet within one track, or a mis-match at the vertex of the two charged tracks. If the in-flight decay happened before passing through the analysis magnet, the reconstruction of the original decay vertex could have been bad. Parameters to cut on in that case were the vertex's chi-squared (`vtx_chi`) and the vertex's position (`vtxx`, `vtxy`, and `vtxz`). If the in-flight decay happened after passing through the analysis magnet, the matching of the upstream (pre-magnet) and downstream (post-magnet) segments of the track could have been bad, and/or the X-Y position of the original decay vertex could have been off. Important parameters in this case corresponded to a non-zero value for `muoffx`, `muoffy`, `eoffx`, and/or `eoffy` for a magnet mismatch. In addition to a track mismatch, the production of a neutrino meant that part of the energy of the decay was lost. When the masses and momentums of the detected particles were added up, the total energy would have been less than the mass of the parent kaon. Even if the muon took most of the energy from the decay of the pion, assigning it a muon mass rather than the pion mass would have again caused `K_mass` to be low. Keeping the width of the signal box to $\pm 10 MeV$ should have excluded those events. Since this decay process also involved one type of particle "faking" another, analysis included looking at parameters having to do with particle identification: `eop_elec`, `fuse3e`, `eop_mu`, and `nmumats`. The parameter which helped us reduce the K3pi background the most was `pp0_kine`, which was specifically constructed to kinematically discriminate K3pi decays. Please see the definition of `pp0_kine` on page 126 for a description of how we used it.

The $K_L \rightarrow \pi^{\pm}\mu^{\mp}\nu_{\mu}$ (Kmu3) and $K_L \rightarrow \pi^{\pm}e^{\mp}\nu_e$ (Ke3) decay modes have larger branching ratios than K3pi, but were less likely to satisfy the trigger requirements. In the Kmu3 mode, the charged pion would have had to fake an electron (or have decayed into an electron and neutrino), and the neutral pion would have had to come from other overlapping events (commonly called *accidentals*) or from a real electron undergoing bremsstrahlung radiation. The Ke3's charged pion would have had to fake a muon (or decayed into a muon and neutrino), and again the neutral pion would have had to come from accidentals or from bremsstrahlung radiation.

As in trying to remove K3pi events, particle identification parameters and track matching parameters were important for removing faking events and decays in flight. To identify the two photons as coming from an event outside the kaon decay, we had on-line cuts requiring the photons to be detected within a certain time-window around the detection of the two charged tracks. We also looked at the reconstructed mass of the "neutral pion" (`pi0_mass`). The `pi0_mass` was reconstructed from the two photons using the decay vertex found from the charged tracks. If the photons had not come from the decay of a neutral pion within the kaon decay, `pi0_mass` would have been different from the true π^0 mass. Photons coming from bremsstrahlung radiation were emitted with very little transverse momentum and so traveled a path close to that of its parent electron. If the bremsstrahlung radiation occurred in the vacuum window (as was most likely) then the photon would have continued to follow the electron's original trajectory even though the electron itself was later bent by the analysis magnet. Therefore the electron track segment upstream of the magnet was projected to an X-Y point on the calorimeter and the distance between this point and the nearest photon cluster was calculated (`xbrdist` and `ybrdist`).

Accidentals also resulted from two different kaons decaying within the same time window with some information from one or both of the decays being "lost". If, for example, a K3pi decay happened with a Ke3 decay, failure to reconstruct the K3pi's two charged tracks could have resulted in the neutral pion looking like it came from the Ke3 decay. Hence cuts were placed on the number of in-time hit pairs detected in the chambers (`npairsup` and `npairsdn` or `xpairsup` and `xpairsdn`). A $K_L \rightarrow 3\pi^0$ or $K_L \rightarrow 2\pi^0$ could have lent two of its photons to a Ke3 or Kmu3 decay if the remaining photons escaped. To prevent this background avenue, cuts were placed on the energy detected in the RC and SA veto systems (`RCmaxene` and `SAmaxene`) that were tighter than in the online trigger.

Finally, a $K_L \to \pi^0\pi^\pm e^\mp\nu_e$ (Ke4) decay could have faked a Kpi0me decay if its charged pion faked a muon or decayed into a muon and muon neutrino. This again gave the particle combination ($\pi^0\mu^\pm e^\mp\nu_\mu\nu_e$) where the neutrinos would have been missed. Ke4 gave a final state most like the search target Kpi0me, and so it was one of the greatest concerns for a background. Fortunately it has a very small branching ratio (BR(Ke4) = $(5.18 \pm 0.29) \times 10^{-5}$, see Table 1.1) which kept it from being a dominating background.

To remove this background, we again used the track reconstruction parameters `vtx_chi`, `muoffx`, and `muoffy` to remove decays-in-flight of the charged pion. In addition, for a background event of this type the mass of the parent particle (`K_mass`) should have nearly always reconstructed to a value lower than the actual kaon mass because of the missing neutrino energy. We also constructed a parameter similar to `pp0_kine` for Ke4's called `ke4_kine`. Cuts on `ke4_kine` proved to be very powerful in reducing background. `Ke4_kine` is described on page 128.

Chapter 4
1997 Data Analysis

The first KTeV data to be analyzed in this search was collected in 1997. Three data runs taken on a single day were used as a sample set to tune the analysis parameters. My Master's thesis[40] covers much of what was done during the blind analysis of the 1997 data. However, at the time of defense (December, 1999), there were several problems remaining and the signal box had not yet been opened. The problems were solved, the box opened, and the results presented at the DPF 2000 meeting in August. Between then and this dissertation, more MC statistics was generated and the version of MC and analysis code was updated. I include here a brief overview of that analysis, and an updated branching ratio limit in Chapter 5.

4.1 Selection cuts

After KTSPILL, KTeVAna, Crunch, and Ntuple cuts were applied, the final selection cuts were applied using PAW. The selection cuts used in the 1997 blind analysis were (events pass if true):

- Detector cuts:
  No triggers in veto detectors (`phvbar1` is .TRUE.)
  Ring counter energy (`RCmaxene`) $\leq 0.5 GeV$
  Spectrometer-Anti energy (`SAmaxene`) $\leq 0.4 GeV$
  Back-Anti hadronic section energy (`EBA3`) ≤ 2000 *ADC counts*
  # TRD planes hit along e track (`npln_e`) > 0
  # TRD planes hit along mu track (`npln_m`) > 0

Number of upstream track hit pairs (`npairsup`) ≤ 8
Number of downstream track hit pairs (`npairsdn`) ≤ 10

- Reconstruction cuts:
  hardware clusters (`nhCsIcls) < 4`
  | "electron" track seg. offset in X (`eoffx`) | $\leq 0.0\ 02m$
  | "electron" track seg. offset in Y (`eoffy`) | $\leq 0.0\ 02m$
  | "muon" track seg. offset in X (`muoffx`) | $\leq 0.001m$
  | "muon" track seg. offset in Y (`muoffy`) | $\leq 0.001m$
  Vertex X and Y (`vtxx` and `vtxy`) in beam region
  $96m \leq$ Vertex Z (`vtxz`) $\leq 155m$
  Vertex χ^2 (`vtx_chi`) ≤ 10
  Transverse momentum squared (`pt2`) $\leq 0.0025\frac{GeV^2}{c^2}$
  Energy of kaon (`K_energy`) $\leq 200GeV$
- Particle ID cuts:
  $0.132\frac{GeV}{c^2} \leq$ Two photon mass (`pi0_mass`) $\leq 0.138\ \frac{GeV}{c^2}$
  $0.97 \leq$ Energy over momentum of electron (`eop_elec`) ≤ 1.03
  TRD CL of electron track being from a π or μ (`prob_e`) ≤ 0.04
  TRD CL of muon track being from a π or μ (`prob_mu`) ≥ 0.04
  3x3 fusion χ^2 of electron cluster shape (`fuse3e`) ≤ 4
  3x3 fusion χ^2 of 1st photon cluster shape (`fuse3p1`) ≤ 4
  3x3 fusion χ^2 of 2nd photon cluster shape (`fuse3p2`) ≤ 4
  | Muon momentum (`pmu`) | $\geq 10\frac{GeV}{c}$
- Signal cuts:
  charged system mass (`chrgmass`) $\leq 0.364\frac{GeV}{c^2}$
  γ clusters were not brem. radiation (|`xbrdist`| or |`ybrdist`| $\geq 0.005m$)
  Signal π^0 kinematic (`sgp0kine`) $\geq$ -0.002
  K3pi π^0 kinematic (`pp0_kine`) $<$ -0.001
  -0.08 $<$ Ke4 kinematic (`ke4_kine`) $<$ -0.005

Please see Section 3.4 for the definitions of the parameters used here.

After all cuts were applied, 5,159 signal MC events out of 97,777 fell within the signal box, giving an acceptance of 5.276$\pm$0.071%. These signal MC are shown in Figure 4.1, in a format which I will refer to as the "study plot". It is a plot of `pt2` versus `K_mass` which is 100 times the size of the signal acceptance box, 10 times in the direction of positive `pt2` and 10 times across the range of `K_mass` roughly centered around the box. Normalization events were treated to cuts as reasonably similar to the signal selection cuts as possible. The main difference was that some particle ID cuts were removed for the normalization events. The normalization selection cuts

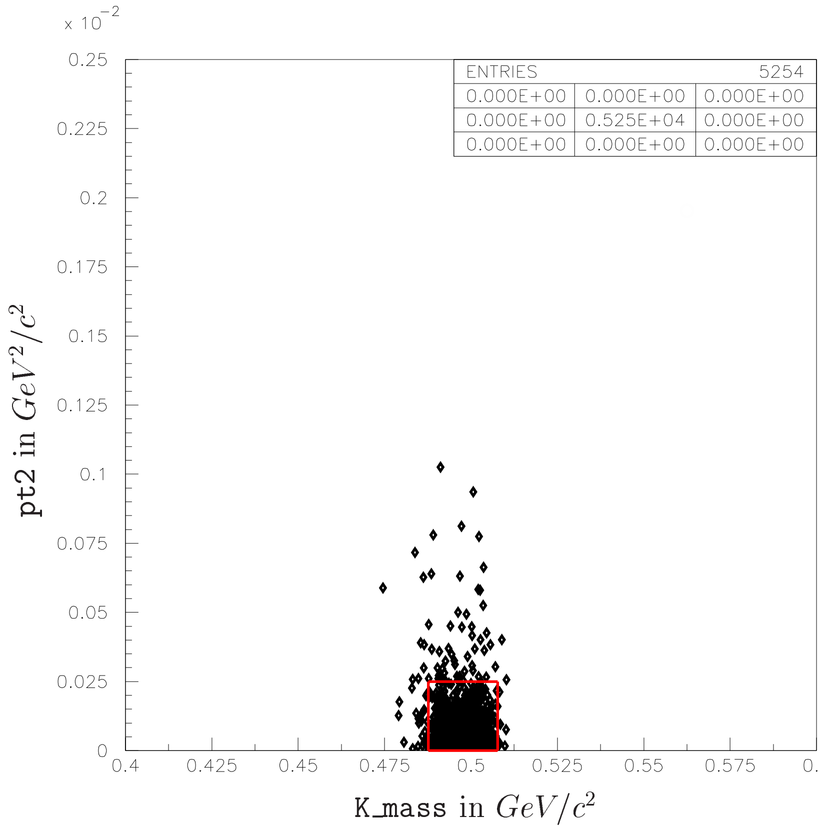

Figure 4.1 The "study plot" of MC signal events passing all cuts (except the signal box). There were 5,159 events inside the red signal box, giving an acceptance of 5.276±0.071%. The area of the plot is 100 times the area of the signal box.

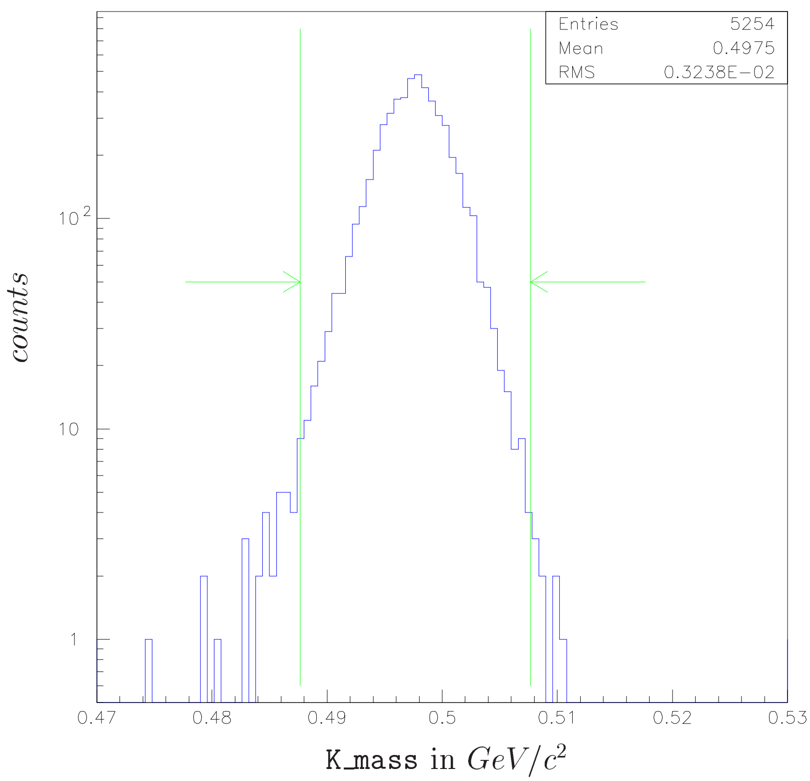

Figure 4.2 A projection onto the X-axis (`K_mass`) of MC signal events passing all cuts (except the signal box) from Figure 4.1. There were 5,211 events inside the signal box `K_mass` limits, shown here by the green vertical lines. Note the log scale on the Y-axis.

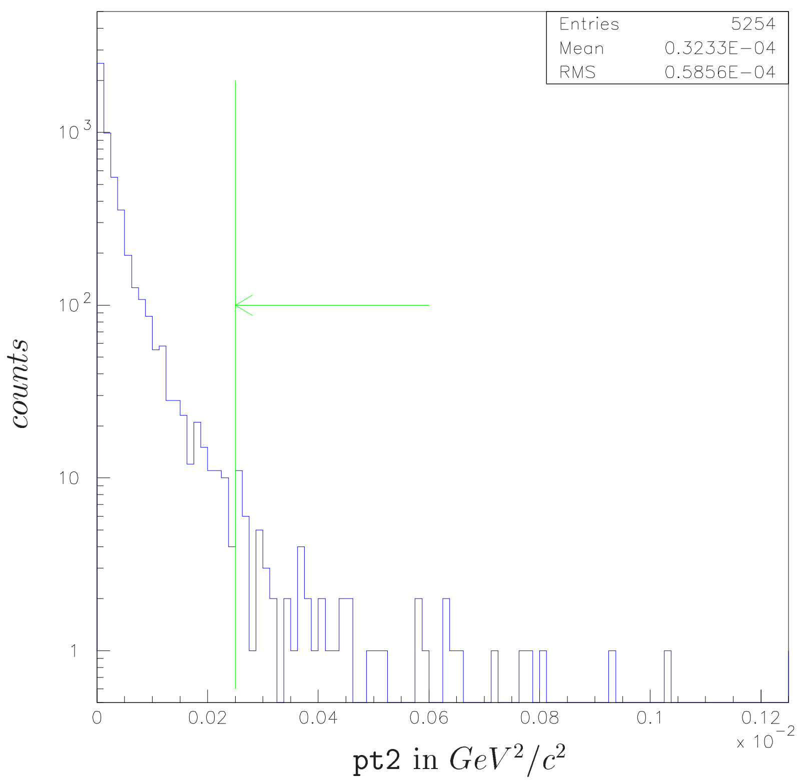

Figure 4.3 A projection onto the Y-axis (`pt2`) of MC signal events passing all cuts (except the signal box) from Figure 4.1. There were 5,192 events inside the signal box `pt2` limit, shown here by the green vertical line. Note the log scale on the Y-axis.

are listed here (events pass if true):

- Detector and Reconstruction cuts:
  same as signal selection cuts
- Particle ID cuts:
  $0.132\frac{GeV}{c^2} \leq$ Two photon mass (`pi0_mass`) $\leq 0.138\ \frac{GeV}{c^2}$
  3x3 fusion χ^2 of 1st photon cluster shape (`fuse3p1`) ≤ 4
  3x3 fusion χ^2 of 2nd photon cluster shape (`fuse3p2`) ≤ 4
- Signal cuts:
  charged system mass (`chrgmass`) $\leq 0.364\frac{GeV}{c^2}$
  neutral clusters aren't brem. radiation ($|$`xbrdist`$|$ or $|$`ybrdist`$| \geq 0.005m$)
  within "signal" box

Unlike the signal selection cuts, the normalization selection cuts included the signal box by definition since we wanted to be sure passing events were really K3pi decays. This normalization signal box is identical to the search signal box. After all cuts were applied, there were 199,756 trigger 2 data events remaining, and 61,820 out of 978,640 MC normalization events remaining. The normalization acceptance was then 6.317±0.025%. The flux was found by

$$Flux(K_L) = \frac{(Number\,of\;K3pi\,data\,accepted) \times (Prescales\,on\,K3pi\,events)}{(BR(K_L \rightarrow \pi^+\pi^-\pi^0)) \times (BR(\pi^0 \rightarrow 2\gamma)) \times (K3pi\ acceptance)} \quad (4.1)$$

Trigger 2 had an online prescale of 500:1, and an off-line prescale of 20:1 was imposed on the data set we used. The branching ratio for the K3pi decay is 12.56±0.20% and the branching ratio for $\pi^0 \rightarrow 2\gamma$ is 98.80±0.03%[8]. The resulting flux was then $2.548{\pm}0.010_{stat}{\pm}0.042_{BRs}{\times}10^{11}$. As a reminder, this flux only included spills that were considered "good" according to the KTSPILL cut, and so can be lower than KTeV fluxes quoted elsewhere.

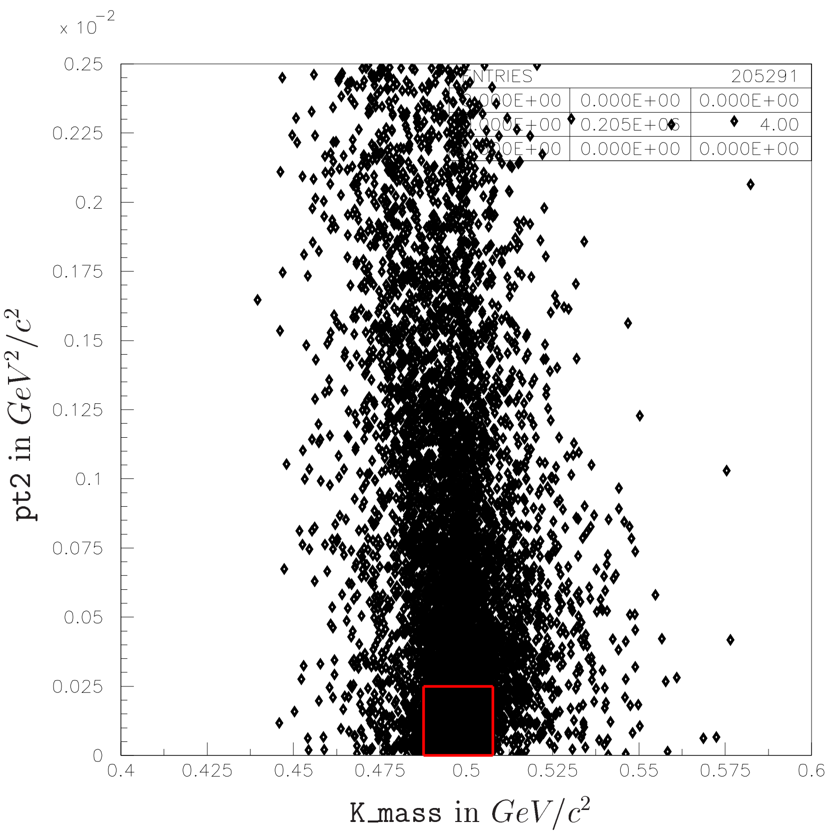

Figure 4.4 1997 trigger 2 data events reconstructed as K3pi's and passing all the normalization selection cuts except the "signal" box. The number of events inside the red signal box (199,756) was used in calculating the 1997 kaon flux.

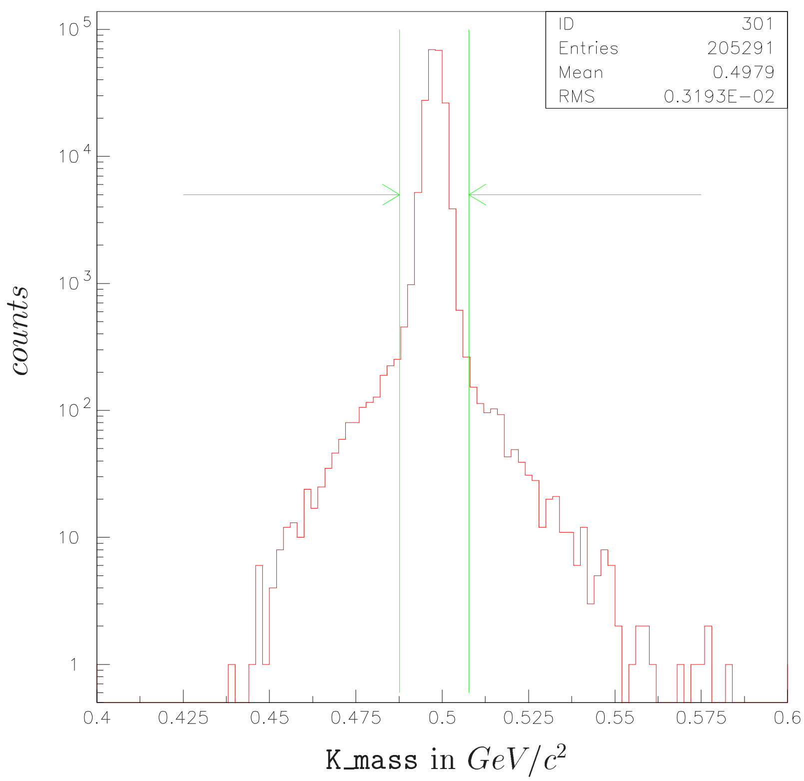

Figure 4.5 A projection onto the X-axis (`K_mass`) of normalization data events passing all cuts (except the signal box) from Figure 4.4. There were 202,978 events inside the signal box `K_mass` limits, shown here by the green vertical lines. Note the log scale on the Y-axis.

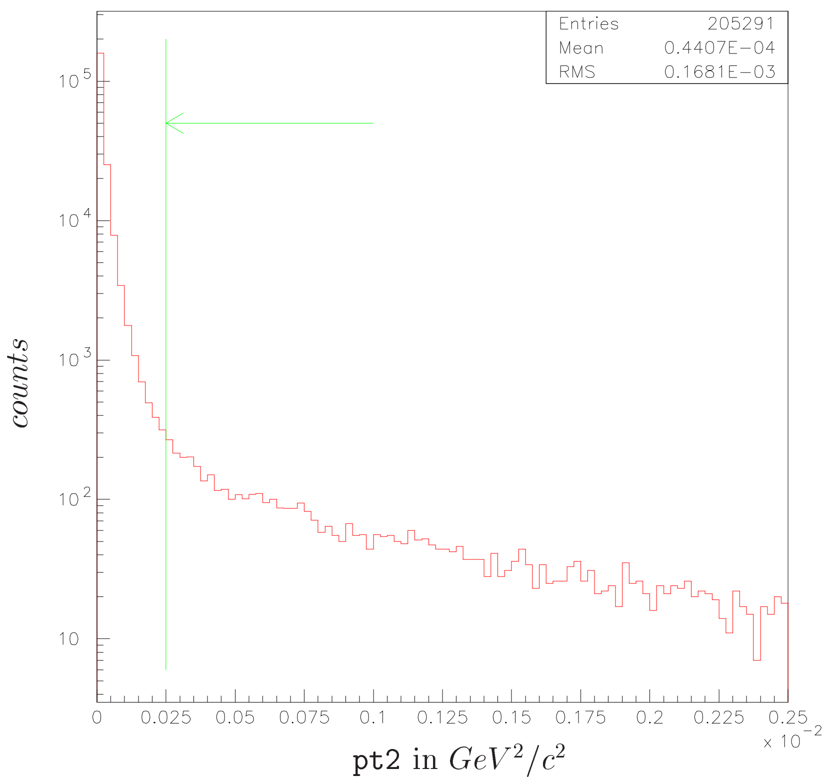

Figure 4.6 A projection onto the Y-axis (`pt2`) of normalization data events passing all cuts (except the signal box) from Figure 4.4. There were 200,212 events inside the signal box `pt2` limit, shown here by the green vertical line. Note the log scale on the Y-axis.

4.2 1997 Backgrounds

Many background modes were investigated during the 1997, but only three were shown to be significant: $K_L \to \pi^+\pi^-\pi^0$ (called K3pi), $K_L \to \pi^0\pi^\pm e^\mp \nu_e$ (called Ke4), and $K_L \to \pi^\pm e^\mp \nu_e$ (called Ke3). Other possible background decay modes included Kmu3 ($K_L \to \pi^\pm \mu^\mp \nu_\mu$), the beta decay of a cascade particle ($\Xi^0 \to \Sigma^+ e^- \bar{\nu}_e, \Sigma^+ \to p\pi^0$), and the occurrence of two kaon decays within a single bucket. All the backgrounds are summarized in Table 4.1. Only the significant backgrounds will be discussed in detail.

Background name	# in study plot (after all cuts)	# in signal box	units of '97 K_L flux used
K3pi ($K_L \to \pi^+\pi^-\pi^0$)	6.8 ± 1.6	negligible	1×10^{-4}
Ke4 ($K_L \to \pi^0\pi^\pm e^\mp \nu_e$)	9.1 ± 4.7	0.10 ± 0.08	31
Ke3 ($K_L \to \pi^\pm e^\mp \nu_e$)	23.7 ± 5.4	0.45 ± 0.12	1
Kmu3 ($K_L \to \pi^\pm \mu^\mp \nu_\mu$)	negligible	negligible	0.1
Cascade beta ($\Xi^0 \to \Sigma^+ e^- \bar{\nu}_e, \Sigma^+ \to p\pi^0$)	$< 2.8 \times 10^{-4}$	negligible	N/A
Double decays			
Ke3 + K2pi0 ($K_L \to 2\pi^0$)	$2.1 \pm 2.1 \times 10^{-3}$	negligible	27
Ke4 + K2pi0	$< 1.5 \pm 1.5 \times 10^{-3}$	negligible	100
Kmu3 + K3pi0 ($K_L \to 3\pi^0$)	$3.5 \pm 3.5 \times 10^{-5}$	negligible	1
Ke3 + K3pi0	0.16 ± 0.16	negligible	1
Ke4 + K3pi0	$6.2 \pm 6.2 \times 10^{-4}$	negligible	10
Ke4 + K3pi	0.13 ± 0.13	negligible	2
Significant background totals	39.6 ± 5.7	0.55 ± 0.14	N/A

Table 4.1 Expected background rates of investigated modes for the 1997 analysis. The far right column gives the approximate size of the event sample used for calculating the given expectation, in units of 1997 K_L flux where 1 = 2.544×10^{11} kaons. For the cascade beta decays, 5×10^8 MC events were generated, far more than was expected to be present in the search data.

4.2.1 K3pi background in 1997

The large branching ratio of K3pi combined with the large acceptance of these events at the trigger level meant that generating enough MC to match the number of events expected in the search data was impossible. Even if we had been able to generate enough MC, our MC did not model hadron showers well, and so the resulting background estimate would not have been very reliable. Instead, we used a modified set of cuts and data events from trigger 2 to estimate the number of K3pi events expected in the trigger 7 study plot. Our normalization data was sent through our signal analysis code to be reconstructed as Kpi0me's. If all the signal selection cuts were then applied, no events passed. Instead, we applied most of the selection cuts, and then individually applied three independent subsets of the remaining selection cuts one at a time. The subsets of cuts used are listed in Table 4.2 and their effects are listed in Table 4.3. The four sets of cuts were called "base", "electron ID", "muon ID", and "kinematic".

The normally ntuple-level cut on the energy of the muon calorimeter cluster was applied here since it was not applied during the reconstruction of the normalization events. The other two unfamiliar items in the muon ID list were cuts that were normally taken online in trigger 7 but not in trigger 2. An extra cut was placed on the reconstructed mass because we were only interested in K3pi's that fell in the study plot, not lower in `K_mass`, and 0.4 was the left edge of the study plot.

Cut subset definitions	
Electron ID cuts	**Kinematic cuts**
0.97 $\leq$ `eop_elec` $\leq$ 1.03	`pp0_kine` $<$ -0.001
`fuse3e` $\leq$ 4	-0.08 $<$ `ke4_kine` $<$ -0.005
`prob_e` $\leq$ 0.04	`sgp0kine` $\leq$ -0.002
Muon ID cuts	**Base cuts**
muon cluster energy $< 1GeV$	All signal selection cuts not in Kinematic, Electron ID, or Muon ID subsets
$\vert$`pmu`$\vert$ $\geq 10GeV$	
`prob_m` $\geq$ 0.04	
a hit in muon bank #3	`K_mass` $\geq 0.4GeV/c^2$ (left edge of study plot)
no hit in hadron anti (HA)	

Table 4.2 The subsets of cuts applied to trigger 2 data to estimate the number of background K3pi events in the study plot. The electron ID, muon ID, and kinematic subsets were assumed to be independent of each other. Each condition had to be true for an event to pass.

Cuts used	# in 1997 trig. 2 data in study plot	reduction factor
Base cuts only	67,575	1
Base + Electron ID cuts	41±6.4	1,648±257
Base + Kinematic cuts	36±6.0	1,877±313
Base + Muon ID cuts	2,112±46	32.00±0.71
Total reduction factor		$9.9 \pm 2.3 \times 10^7$
K3pi's expected in '97 study plot	6.8±1.6	

Table 4.3 The effects of applying subsets of signal selection cuts to trigger 2 normalization data reconstructed as Kpi0me to estimate the expected number of K3pi events passing all signal selection cuts. If all the cuts were applied at once, no events passed, and it was impossible to generate enough MC to get reasonable statistics for this background.

To get the number of K3pi's expected in our study plot, we divided the number passing all but the subset cuts (67,575 events) by the total reduction factor ($1{,}648 \times 1{,}877 \times 32.00 = 9.9 \pm 2.3 \times 10^7$) and multiplied by the prescale factors (500:1 on-line and 20:1 off-line). The result was 6.8±1.6 K3pi background events expected in our study plot. Figure 4.7 shows the study plot with K3pi background data to which all but the kinematic and electron ID cuts had been applied. The events all reconstructed at a mass much lower than the signal box, so the expected background rate for K3pi's in the signal box was negligible. This allows our rough estimation technique to be acceptable.

The expected number of K3pi background events in the study plot was high compared to the final data. Since we used real data, it was entirely possible that there were other decays in the sample besides K3pi's, such as Ke4's and Ke3's. These "extra" passing events would have pushed the background estimate for K3pi's higher than the actual number of trigger 7 K3pi's that passed all the selection cuts. Since the background in the signal box from K3pi's was expected to be tiny, and the error was in the direction of overestimation, we were not too concerned about this discrepancy. That being said, it is encouraging to note that the number of events seen in the data in the K3pi *plus* Ke4 regions is very close to the K3pi estimate (see next Chapter for results).

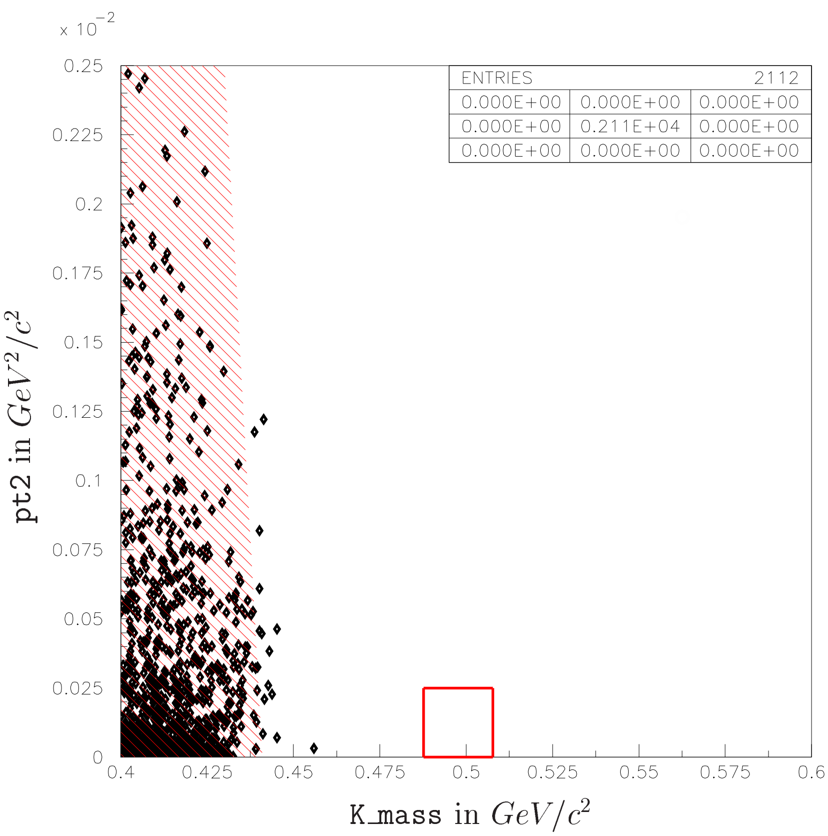

Figure 4.7 Study plot of K3pi data background events passing all cuts except the kinematic and electron ID cuts as defined in Table 4.2. The events reconstructed far below the signal box in `K_mass`, so their expected flux in the box was negligible. The number of events in this plot have not been renormalized to reflect the 1997 K_L flux.

4.2.2 Ke4 background in 1997

The low branching ratio of Ke4's made it possible to generate a MC sample with more than the number of Ke4 decays expected in the 1997 data set. Originally, the form factors used in the Ke4 MC generation code was based on a measurement of $K^+ \to \pi^+\pi^\pm e^\mp \nu_e$. However, the diagrams for the neutral Ke4 decay are quite different, so appropriate form factors were needed. A KTeV undergraduate has started an analysis of KTeV's Ke4 data and was able to provide a preliminary set of form factors, which I used to generate the Ke4 MC samples in this thesis. A full Ke4 analysis will eventually be published. Figure 4.8 shows the Ke4 MC sample after all signal cuts have been applied. There were 284 surviving events out of 4.1×10^8 generated, 31.2 times the number expected in the data set. Scaling the MC by the flux, the expected number of Ke4's in our study plot was 9.1$\pm$0.8. Their distribution in the study plot is shown in Figure 4.8. There were three events which appeared inside the signal box after all cuts were applied. After scaling, this became a small background of 0.10$\pm$0.06 events.

This expectation was higher than the number of events actually seen in the search data study plot by a factor of 2. During analysis of the 1999 data, we discovered that the '97 Ke4 background prediction could be made to agree with the search data if the cut on the `pp0_kine` parameter was moved from -0.001 to -0.06. This indicated that there was a difference in the `pp0_kine` parameter between real Ke4's and MC

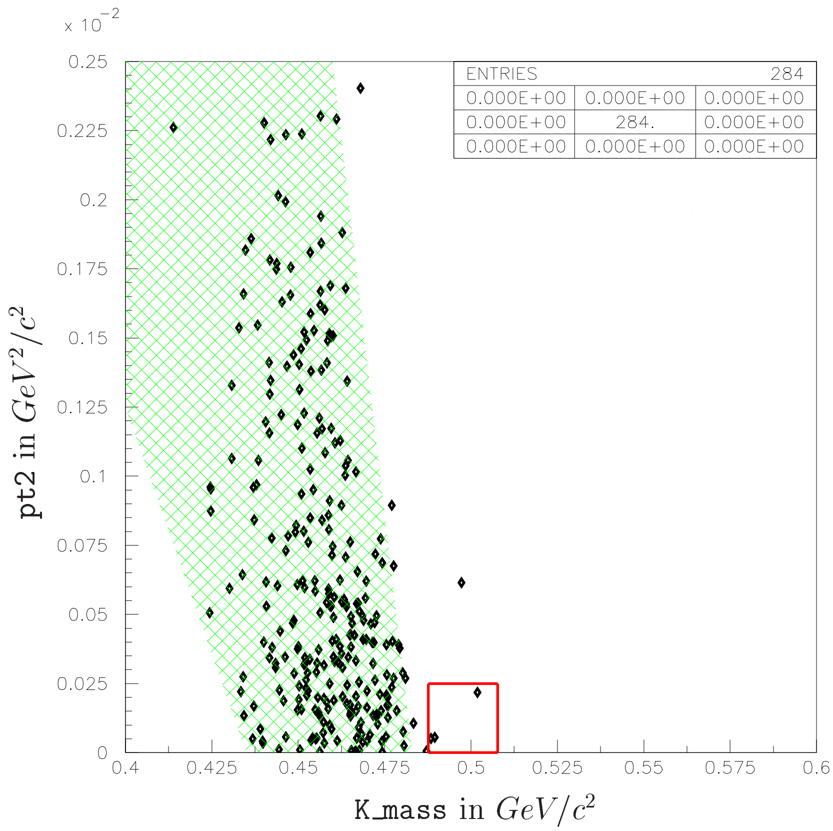

Figure 4.8 Study plot of Ke4 MC background events passing all cuts. This plot represents 31.2 times the 1997 kaon flux. The three events inside the signal box scaled to 0.10±0.06 events expected in the box in the data set.

Ke4's, but my real event sets did not contain enough statistics for my own data versus MC comparison of pure Ke4's. We suspect the discrepancy to be due to the form factors still not being quite correct, so we will add a systematic uncertainty to the background estimations. If the number of events passing all cuts was divided by 2, the background estimates become 4.5±0.5 events in the study plot and 0.05±0.04 events in the box. The systematic uncertainties are therefore (9.1 - 4.5) = 4.6 for the study plot and (0.10 - 0.05) = 0.05 for the box. The final Ke4 background estimations are therefore 9.1±4.7 for the study plot and 0.10±0.08 for the box.

4.2.3 Ke3 background in 1997

The branching ratio for Ke3's is larger than that for K3pi's, but the difficulty of finding two accidental photons that sum to the neutral pion mass suppressed this background. In spite of this suppression, we could not have generated enough statistics of regular Ke3 decays. However, to pass the trigger the pion had to decay into a muon or punch through the detector into the muon banks. So we generated two sets of MC, one in which the charged pion always decayed into a muon (which I will call "Ke3dk") and the other in which the charged pion punched through the filtering steel to make a hit in the muon banks (which I will call "Ke3punch"). Even with the forced decay and punch-through, we needed to remove the `pi0_mass` cut in order to improve our statistics for comparison studies. Also note that since this background required accidental activity to pass the cuts, the estimate was dependent on how well we modeled accidental activity. As mentioned before, accidental activity was incorporated by overlaying data taken with an off-spill trigger so the quality should have been good. On the other hand, we had only a limited sample of such events, with MC generation having to loop several times over the same accidental files. Doing so could have biased our estimate in one direction or the other. Both of these signal-faking backgrounds were significant in the study plot.

Ke3dk's and Ke3punch's

We were able to generate 4.6x10^9 Ke3dk MC and 2.1x10^9 Ke3punch MC for 1997. To determine how often a charged pion from a Ke3 decayed into a muon, we generated a small sample of standard Ke3 MC, which correctly modeled the momentum dependent pion-to-muon decay. During generation we counted the number of times the program entered the pion-decay subroutine. The pion underwent muon decay 5.71$\pm$0.03% of the time. The rate of pion punch-through was not modeled well in our MC, so we had to use another method to determine the Ke3punch scaling factor.

To calculate the expected Ke3punch flux, we looked at the variable which showed the biggest difference in histogram shape between Ke3dk and Ke3punch: `vtx_chi`. The same set of MC events generated to estimate the Ke3dk and Ke3punch background were reconstructed as Ke3 decays and had to satisfy the following cuts to be considered "good":

- Detector and Reconstruction cuts:
  same as signal selection cuts (p. 145) except `vtx_chi` and
  `pt2` cuts removed
- Particle ID cuts:
  $0.97 \leq$ Energy over momentum of electron (`eop_elec`) ≤ 1.03
  TRD CL of electron track being from a π or μ (`prob_e`) ≤ 0.04
  TRD CL of muon track being from a π or μ (`prob_m`) ≥ 0.04
  3x3 fusion χ^2 of electron cluster shape (`fuse3e`) ≤ 4
  at least one muon bank hit (`mu_flag` = .TRUE.)
  | Muon momentum (`pmu`) | $\geq 10\frac{GeV}{c}$
- Signal cuts:
  reconstructed parent mass (`K_mass`) $\leq 0.6\frac{GeV}{c^2}$

Since our normalization trigger was a minimum bias 2-track trigger, our normalization data included Ke3 events as well as K3pi events. We therefore ran Ke3 reconstruction code on the trigger 2 data to get a Ke3 data set to which to compare the MC.

Figure 4.9 shows plots of the parameter `vtx_chi` for the Ke3punch MC, Ke3dk MC, and Ke3 data sets with the cuts listed above. Ke3punch events tended to cluster at low `vtx_chi` values and Ke3dk events had a long tail. The data had some combination of the two. Three different methods were used to determine what fraction of the Ke3 data was Ke3punch.

For the first method, we took the means of the MC distributions and combined them linearly to match the data mean:

$$(\mu_{punch} * f) + (\mu_{dk} * (1 - f)) = \mu_{data} \qquad (4.2)$$

This resulted in 40.6±2.0% Ke3punch and 59.4±2.0% Ke3dk after cuts were made. The combined MC set is shown overlaid on the data in Figure 4.10. To get this same

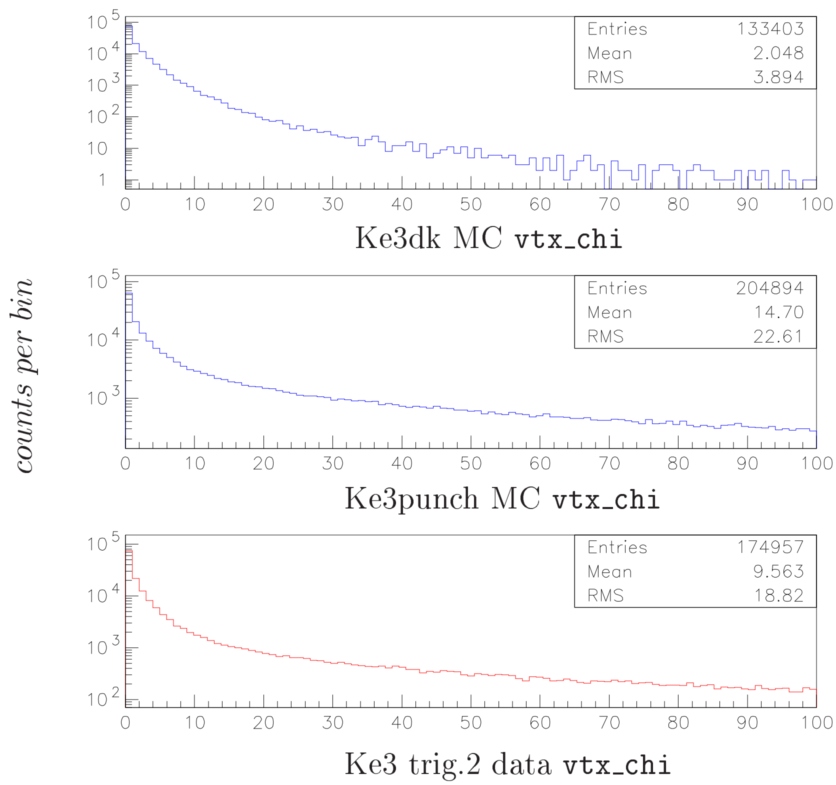

Figure 4.9 Histograms of `vtx_chi` for Ke3punch MC (top plot), Ke3dk MC (middle plot), and Ke3 data (bottom plot). Note the means of each distribution; Ke3punch's tended to have low `vtx_chi` values, Ke3dk's had a long tail, and the data was a mixture of the two.

ratio back at the generation level (so we could predict the background for the signal decay), we had to correct for the acceptance. Ke3punch MC acceptance for the Ke3 cuts was $6.45 \pm 0.02 \times 10^{-5}$. Ke3dk MC acceptance was $4.41 \pm 0.01 \times 10^{-5}$. Since 1.46 Ke3punch's were accepted for every one Ke3dk accepted, the Ke3punch percentage was divided by 1.46. Renormalizing then gave us a generation-level Ke3punch to Ke3dk ratio of 31.8±1.7% to 68.2±1.7%.

The second method looked at the raw number of events in two bins: `vtx_chi` from 0 to 20, and from 30 to 100. Restricting `vtx_chi` from 30 to 100 resulted in nearly all the events being Ke3dk's (36,217 Ke3dk events to 409 Ke3punch events), so I called this the "decay" bin. There were 18,806 Ke3 data events in this bin. The flux correction to the Ke3dk MC was then

$$\frac{\#\ Ke3\ data}{\#\ Ke3dk} = \frac{18,806}{36,217} = 0.519 \pm 0.005 \tag{4.3}$$

In the 0 to 20 `vtx_chi` bin (which I will call the "punch" bin), there were 156,534 Ke3dk's, 132,478 Ke3punch's, and 149,721 Ke3 data events. Of the data events, 156,534×0.519 = 81,282±759 events should have been decays. That left 68,439±852 punch-through data events, giving a flux correction to the Ke3punch MC of

$$\frac{\#\ Ke3\ data}{\#\ Ke3punch} = \frac{68,439}{132,478} = 0.517 \pm 0.007 \tag{4.4}$$

Applying the appropriate flux corrections to the total number of Ke3dk and Ke3punch MC events in the full range of `vtx_chi` (0 to 100) gave 39.3±0.6% Ke3punch and 60.7±0.6% Ke3dk after all cuts (see Figure 4.11). This was in good agreement with

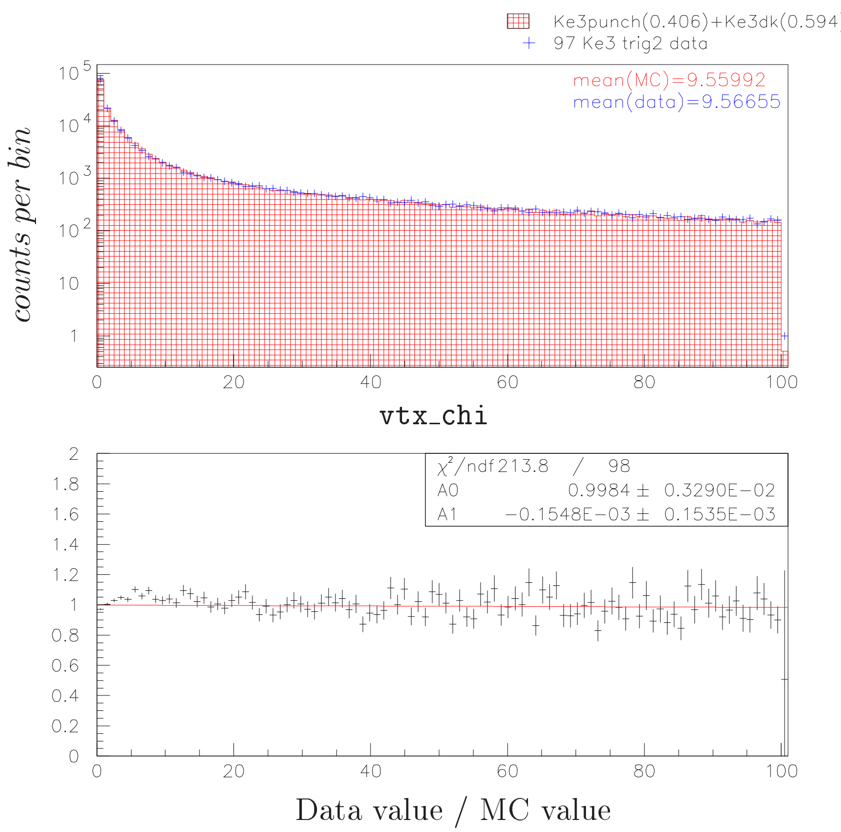

Figure 4.10 Histograms of `vtx_chi` for the Ke3 data and weighted Ke3 MC backgrounds with a punch-to-decay ratio as determined by matching the distribution means. The Ke3punch's are 41% of the MC events,and the Ke3dk's are the other 59%.

the first method.

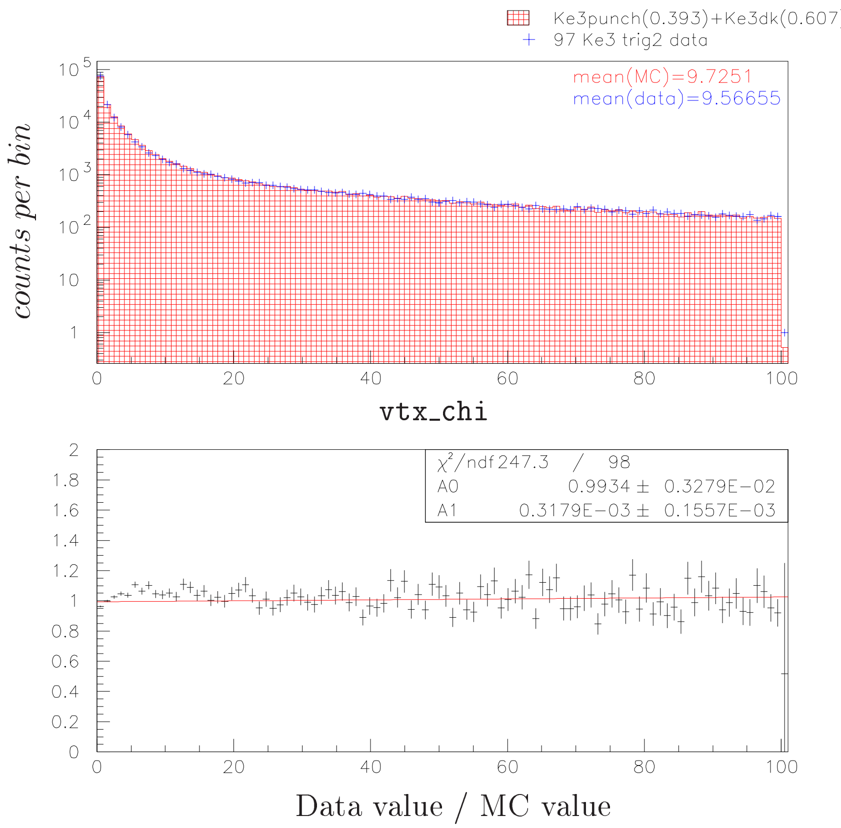

Figure 4.11 Histograms of `vtx_chi` for the Ke3 data and weighted Ke3 MC backgrounds with a punch-to-decay as determined by matching the number of events in the two bins 0 to 20, and 30 to 100. The Ke3punch's are 39% of the MC events,and the Ke3dk's are the other 61%.

There was concern expressed that the normalization mode showed a slope in the `vtx_chi` parameter (see Figure 3.19), but both the mean fitting and raw event number fitting methods resulted in a nearly flat data / MC plot. The third method attempted to take into account the slope seen in the normalization ratio. It used the same two bins in `vtx_chi` as the second method, but this time we looked at the normalization mode in the two bins (0 to 20 or "punch" bin, and 30 to 100 or "decay" bin) and calculated the ratio of MC to data, giving

$$R_p = \frac{\#\ data\ in\ punch\ bin}{\#\ MC\ in\ punch\ bin} = 3.195 \pm 0.018 \tag{4.5}$$

$$R_d = \frac{\#\ data\ in\ decay\ bin}{\#\ MC\ in\ decay\ bin} = 2.745 \pm 0.311 \tag{4.6}$$

From the second method, we knew the raw numbers of events in each bin. We defined the number of Ke3punch events in the punch bin as M_{pp}, the number of Ke3dk events in this bin as M_{pd}, and the number of Ke3 data events in this bin as D_p. Similarly we defined the numbers of raw events in the decay bin as M_{dp} (Ke3punch), M_{dd} (Ke3dk), and D_d (Ke3 data). We then solved the linear system of equations

$$D_p = R_p[\alpha_p M_{pp} + \alpha_d M_{pd}] \tag{4.7}$$

$$D_d = R_d[\alpha_p M_{dp} + \alpha_d M_{dd}] \tag{4.8}$$

for the normalization-`vtx_chi`-slope-corrected flux factors α_p and α_d. The result was α_p=0.164 and α_d=0.187. The αs were weighting factors in the equations, so we had to multiply them by the raw percentages:

$$\%\ Ke3punch\ unnormalized = \alpha_p * \frac{\#\ of\ Ke3punch\ in\ (0-100)}{\#\ of\ Ke3punch\ +\ Ke3dk\ in\ (0-100)} \tag{4.9}$$

$$\%\ Ke3dk\ unnormalized = \alpha_d * \frac{\#\ of\ Ke3dk\ in\ (0-100)}{\#\ of\ Ke3punch\ +\ Ke3dk\ in\ (0-100)} \tag{4.10}$$

and renormalize:

$$\%\ Ke3punch\ after\ cuts = \frac{\%\ Ke3punch\ unnormalized}{\%\ Ke3punch\ unnorm. + \%\ Ke3dk\ unnorm.} \tag{4.11}$$

$$\% \; Ke3dk \; after \; cuts = \frac{\% \; Ke3dk \; unnormalized}{\% \; Ke3punch \; unnorm. + \% \; Ke3dk \; unnorm.} \tag{4.12}$$

which gave 36.2±0.6% Ke3punch and 63.8±0.6% Ke3dk after all Ke3 cuts (see Figure 4.12).

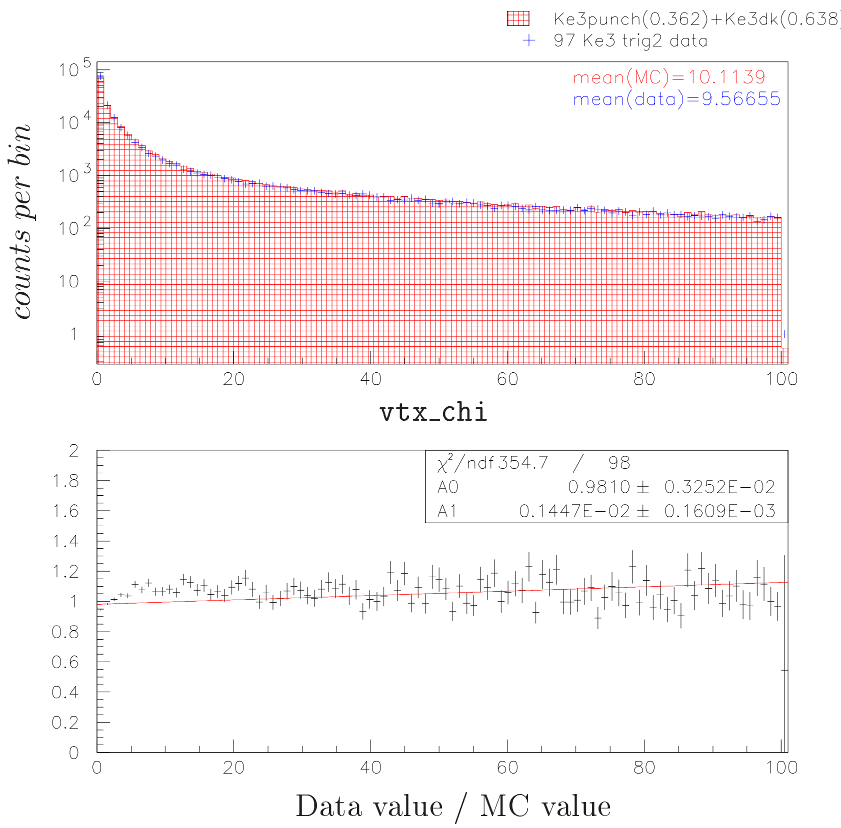

Figure 4.12 Histograms of `vtx_chi` for the Ke3 data and weighted Ke3 MC backgrounds with a punch-to-decay ratio as determined by looking at raw event numbers in two bins and correcting for the slope seen in the normalization data / MC `vtx_chi` ratio plot (Figure 3.19). The Ke3punch's are 36% of the MC events, and the Ke3dk's are the other 64%.

In the end we decided to use the ratio obtained by correcting for the slope seen in

Method of fitting	% of Ke3punch	% of Ke3dk
Mean matching	40.6±2.0	59.4±2.0
Raw number of events	39.3±0.6	60.7±0.6
Correction for K3pi slope	36.2±0.6	63.8±0.6

Table 4.4 Different methods for estimating the ratio of Ke3punch events to Ke3dk events found in Ke3 data after all cuts.

the `vtx_chi` normalization plot. We felt it was the most accurate representation of this background. In the end, the acceptance of Ke3 background events was not very sensitive to this ratio.

Ke3's in study plot

Even with the number of Ke3dk and Ke3punch MC events we were able to generate, we had to remove the cut on `pi0_mass` to get a better picture of their distributions in the study plot (see Figures 4.13 and 4.14). There were 70 Ke3dk events that passed all but the `pi0_mass` cut. If we corrected this number by the expected flux and by the `pi0_mass` cut suppression, this gave 11.4±4.0 Ke3dk events expected to pass all cuts and show up in the study plot. The 76 Ke3punch events in Figure 4.14 likewise scaled to 12.6±3.7 events. This gave a total study plot background expected from Ke3's of 24.0±5.4 events. It turned out that this estimate was high, possibly due to the limited statistics of our accidental overlay events. In the end we chose to use an estimation calculated from the data itself to avoid such problems.

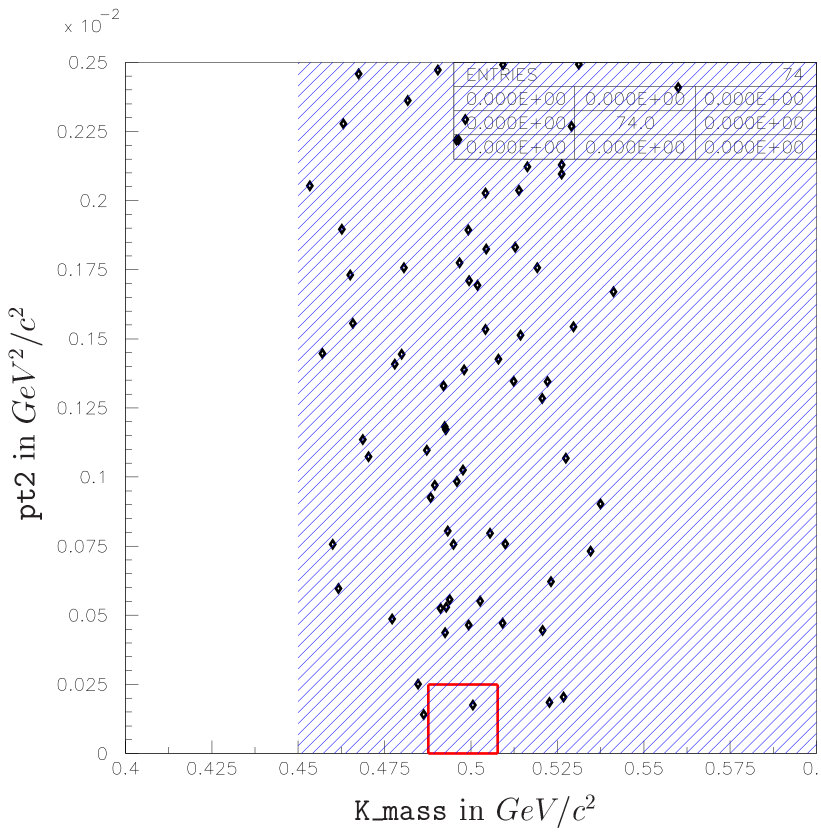

Figure 4.13 Study plot of Ke3dk MC background events passing all but the `pi0_mass` cut. With the `pi0_mass` cut, there were only 8 events in this plot with none in the signal box.

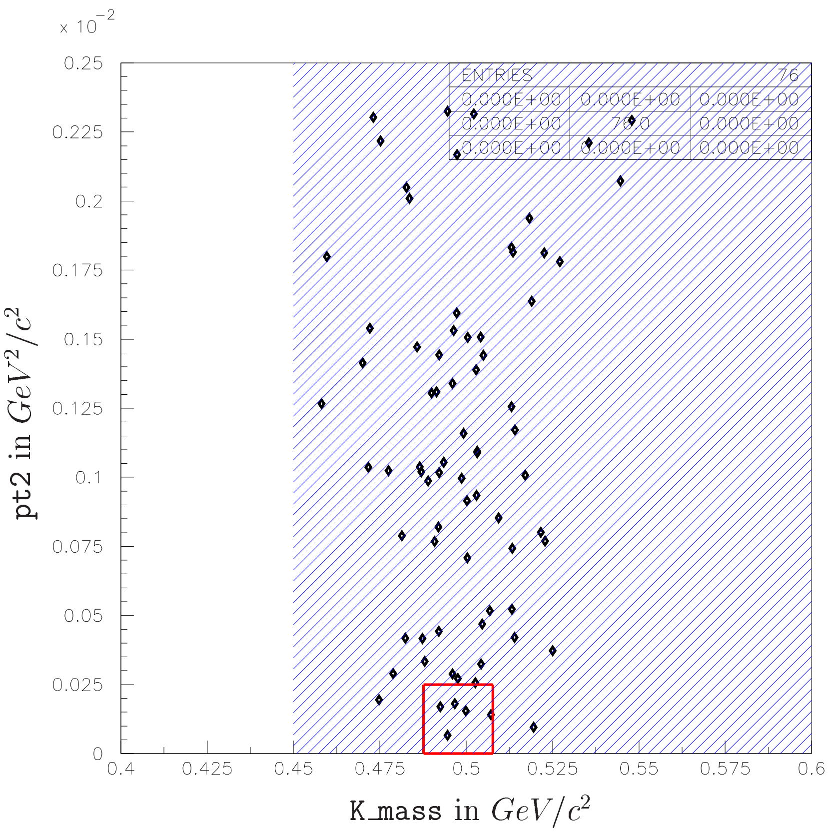

Figure 4.14 Study plot of Ke3punch MC background events passing all but the `pi0_mass` cut. With the `pi0_mass` cut, there were only 12 events in this plot.

Ke3's in Signal Box

There was a good possibility that we could have seen background Ke3 events in our signal box, so it was important to estimate the expected rate as accurately as possible before the box was opened. Several different methods were used to estimate the signal box background as a cross-check of consistency. For one method, removing the `pi0_mass` cut did result in 6 Ke3 MC events in the box, 1 from Ke3dk's and 5 from Ke3punch's. Rescaling of these numbers by the `pi0_mass` reduction factor gave an expected background of 1.0±0.5 events. Another estimation looked at the slice in `K_mass` above the signal box and assumed that the events were uniformly distributed in `pt2`. The slice was ten times larger than the signal box. From Figures 4.13 and 4.14, there were 29 Ke3dk events and 38 Ke3punch events in the slice. Scaling by the expected flux, `pi0_mass` cut suppression, and dividing by 10, gave a signal box background of 0.4±0.2 Ke3dk events and 0.6±0.2 Ke3punch events, for a combined background of 1.1±0.3 events. The slice estimation method was applied to the blind search data as well. Removing the `pi0_mass` cut in the 1997 data resulted in 49 events in the slice above the signal box. Scaling by the appropriate `pi0_mass` cut suppression and dividing by nine this time gave an expected box background of 0.40±0.11 events from Ke3s. When we opened the data's signal box, we could re-do the estimations with the newly available information. That gave a slice estimation of 0.40±0.11 events and a box estimation of 0.87±0.34. These estimations are summarized in Table 4.5.

Method of estimation (pi0_mass cut rem.)	97 events expected in box
Slice in K_mass using MC	1.1±0.3
Events in box using MC	1.0±0.5
Slice in K_mass using blind data	0.45±0.12
Slice in K_mass using opened data	0.51±0.13
Events in box using opened data	0.87±0.34

Table 4.5 Different methods for estimating the expected background in the signal box and their results for the '97 analysis. To call this a blind analysis and get the most accurate estimation, we decided to use the blind data calculation (highlighted in green).

The method with the least uncertainty was the slice method. Since our estimations of the number of events in the study plot has not been very accurate, we decided to use the estimate obtained from the data itself. To call this analysis blind, we must take the estimation made before the box was opened, so we will call our best estimation of the number of Ke3 background events in the box 0.45±0.12.

Since there was not much overlap of the "Ke3" area with the other background areas, we were able to confirm that the background events in the Ke3 area were really Ke3's by comparing data to combined MC. We have included search data / MC comparison plots in Appendix C.

Chapter 5
1997 Data Results

After all cuts have been made there are 19 data events in the region around the signal box for the 1997 data set. Their `K_mass` and `pt2` values are listed in Table 5.1 and Figure 5.1 shows their distribution.

Event label	`K_mass` (GeV/c^2)	`pt2` ($GeV^2/c^2 \times 10^{-4}$)
P01	0.428209	1.73488
P02	0.434315	5.82636
P03	0.439489	1.45711
P04	0.462627	12.2818
P05	0.465229	6.79983
P06	0.473166	5.74874
P07	0.481128	14.0024
P08	0.486848	23.4533
P09	0.491095	10.6743
P10	0.492600	17.0925
P11	0.495497	19.4972
P12	0.498342	23.9350
P13	0.500148	21.4783
P14	0.502267	17.2770
P15	0.511073	5.25049
P16	0.514810	19.5037
P17	0.526813	6.46776
P18	0.535355	8.07041
P19	0.537915	3.99144
B1	0.494102	1.70925
B2	0.499837	1.75208

Table 5.1 A list of the '97 data events passing all the cuts, with their `K_mass` and `pt2` values. Events whose labels start with "P" are in the plot outside of the signal box. Events with "B" labels are inside the signal box. They are ordered left-to-right in Figure 5.1.

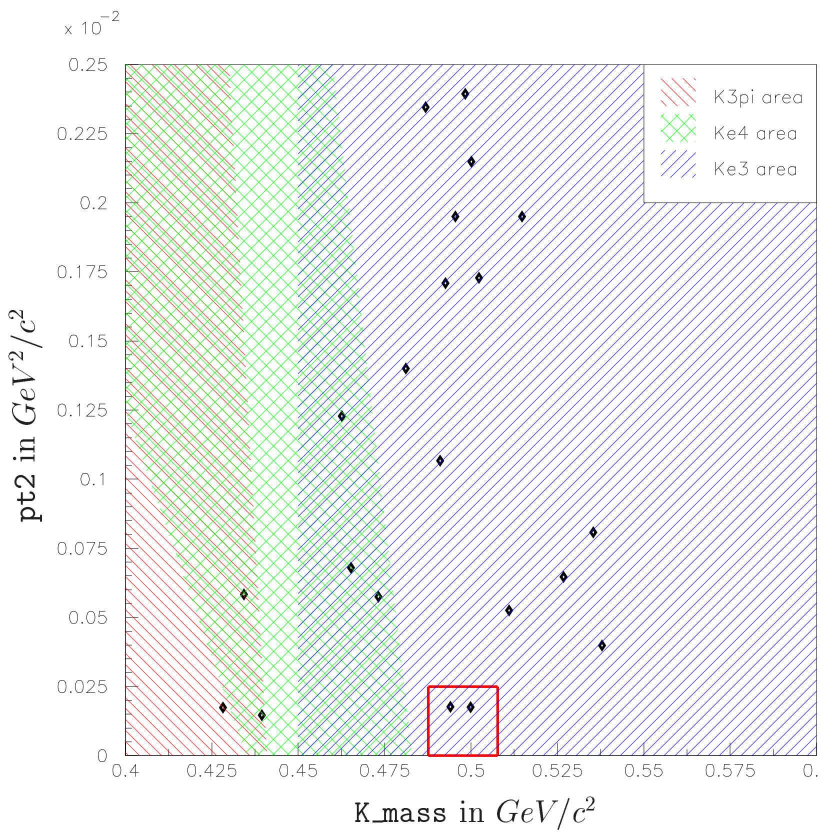

Figure 5.1 Signal search data from 1997. There are 2 events in the K3pi region, 4 events in the Ke4 region, 13 events in the Ke3 region (excluding the box), and 2 events in the signal box.

Using the regions we defined as a guideline, event P01 is a K3pi event. P02 and P03 could be K3pis or Ke4s. Most likely, P03 is a K3pi and P02 with its higher `pt2` is a Ke4 event. Events P04, P05, and P06 fall in the overlap between Ke4s and Ke3s, but they are probably Ke4 events. Events P07 through P19 are all Ke3 events. Such assignments give the breakdown listed in Table 5.2. The estimations all seem to be too high. In the case of the K3pi's, we already discussed that the estimation was derived from data which might have contained other decay processes, pushing the estimation high. It is interesting to note that the number of events seen in the K3pi *plus* Ke4 areas is equal to the K3pi estimation. The reason the Ke4 estimate was high was discovered during the '99 analysis. There was a discrepancy between Ke4 data and Ke4 MC that showed up most in the `pp0_kine` parameter. We changed the cut placement on `pp0_kine` in the '99 analysis and the disagreement went away. Finally, the Ke3 MC estimate was approximate, as discussed before, so it is not that surprising that the numbers do not agree. Since we will use the data itself for the final background estimation, the disagreement was not a major concern.

We saw two data events inside the signal box. This was high compared to the 0.55 events we expected to find (see Table 5.3). The Poisson probability that a value of 0.55 will fluctuate to a value of 2.0 or higher is 10.6%.

The distribution of the data events within the signal box is not consistent with the distribution of the MC signal events. If we project the search data in the box onto

Background in study plot (excluding box)	# of expected events	# of seen events
K3pi	6.8±1.6	2
Ke4	9.1±4.7	4
Ke3	23.7±5.4	13
Total in study plot	39.6±5.7	19

Table 5.2 Comparison between the number of background events expected in the study plot, and what was seen in the search data.

Background in signal box	# of expected events	# of seen events
K3pi	negligible	-
Ke4	0.10±0.08	-
Ke3	0.45±0.12	-
Total background	0.55±0.14	2

Table 5.3 Comparison between the number of background events expected in the signal box, and what was seen in the search data. As mentioned in Chapter 4, the Ke4 estimate was from MC and the Ke3 estimate was from the slice method used on data.

each axis overlaid with the same projection of the signal MC, we can see that the data events do not fall in the peak of the MC distributions (see Figures 5.2 and 5.3). For pt2 especially, the data events are far out on the tail of the signal MC distribution. We can integrate the number of MC events farther from the K_mass mean than each data event. As a percentage of the total number of signal MC in the plot, this gives the probability that a true signal event would fall this far or farther out on the tail of the K_mass distribution. Results for this calculation are given in Table 5.4. A similar calculation can be done for the pt2 projection of Figure 5.3, with results also listed in Table 5.4.

Event label (from left to right in K_mass)	% of MC signal events farther from $\mu_{\texttt{K\_mass}}$	% of MC signal events with greater pt2
B1	23.6±0.6%	1.36±0.16%
B2	42.2±0.7%	1.30±0.16%

Table 5.4 If a given data event were signal, the probability it would fall where it did or farther away from the nominal kaon mass in K_mass or from zero in pt2.

If these two parameters were independent, then these individual probabilities could be simply combined. However, K_mass and pt2 could be correlated. To take into account the possibility of correlations, one could combine the two parameters before comparison. The simplest way was just to multiply their absolute values together after adjusting the K_mass to be centered around zero. Figure 5.4 shows the distributions of the search data and signal MC for $|\ \mu_{\texttt{K\_mass}}$-K_mass$|$*pt2. The percentage of MC

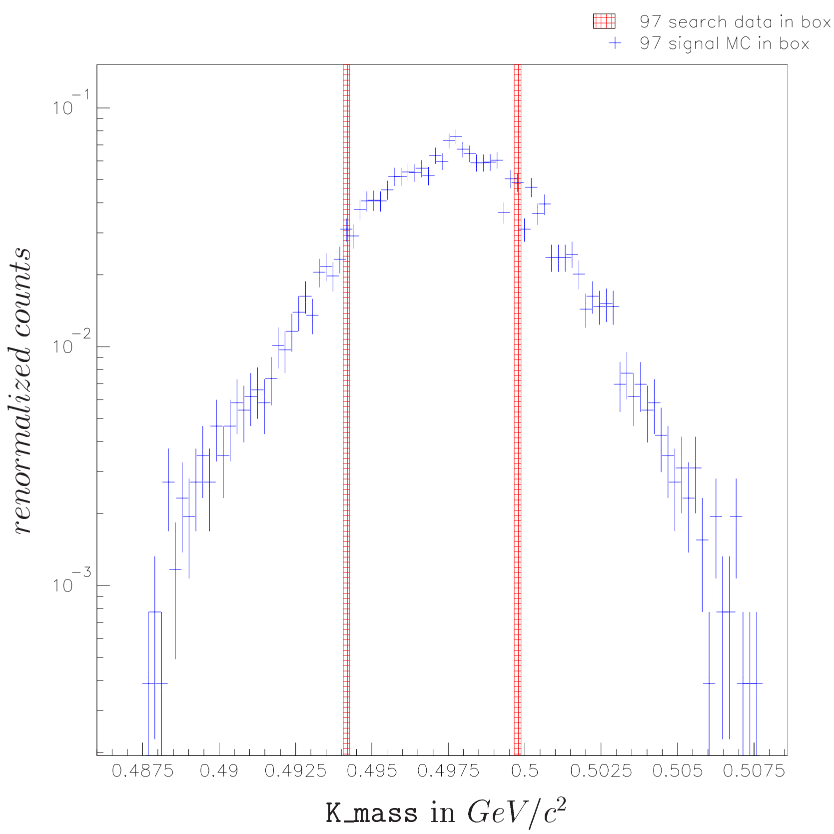

Figure 5.2 1997 data events in the signal box (red) and MC signal events (blue) passing all selection cuts projected onto the `K_mass` axis. The data events do not fall in the center of the signal MC distribution. The MC histogram has been renormalized to fit on the plot.

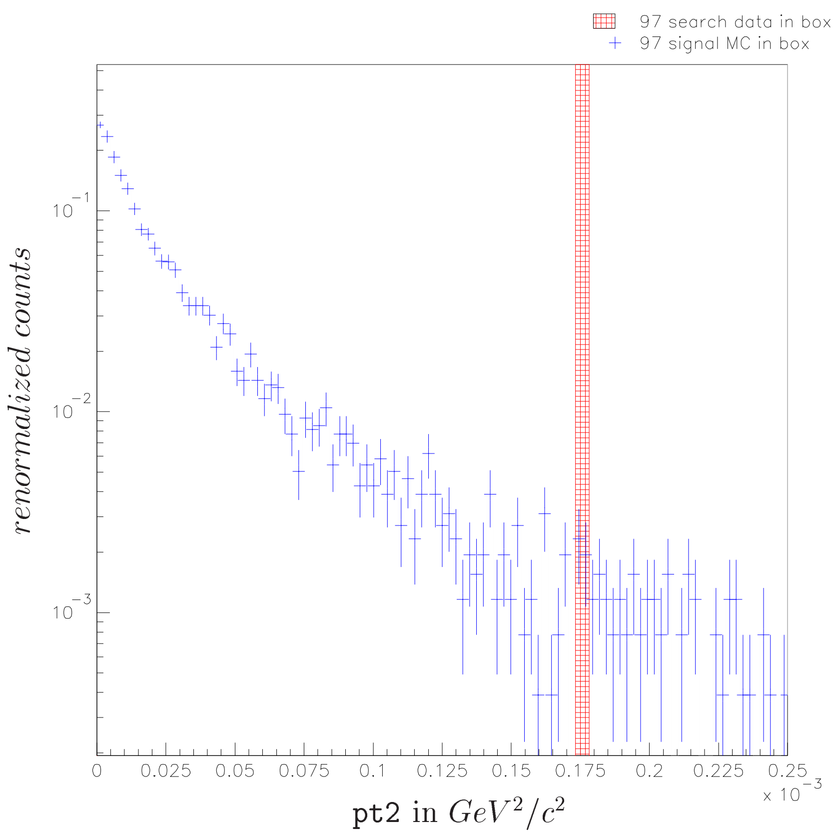

Figure 5.3 1997 data events in the signal box (red) and MC signal events (blue) passing all selection cuts projected onto the `pt2` axis. The data events fall on the tail of the signal MC distribution. The MC histogram has been renormalized to fit on the plot.

events with higher values than the two data points is shown in Table 5.5.

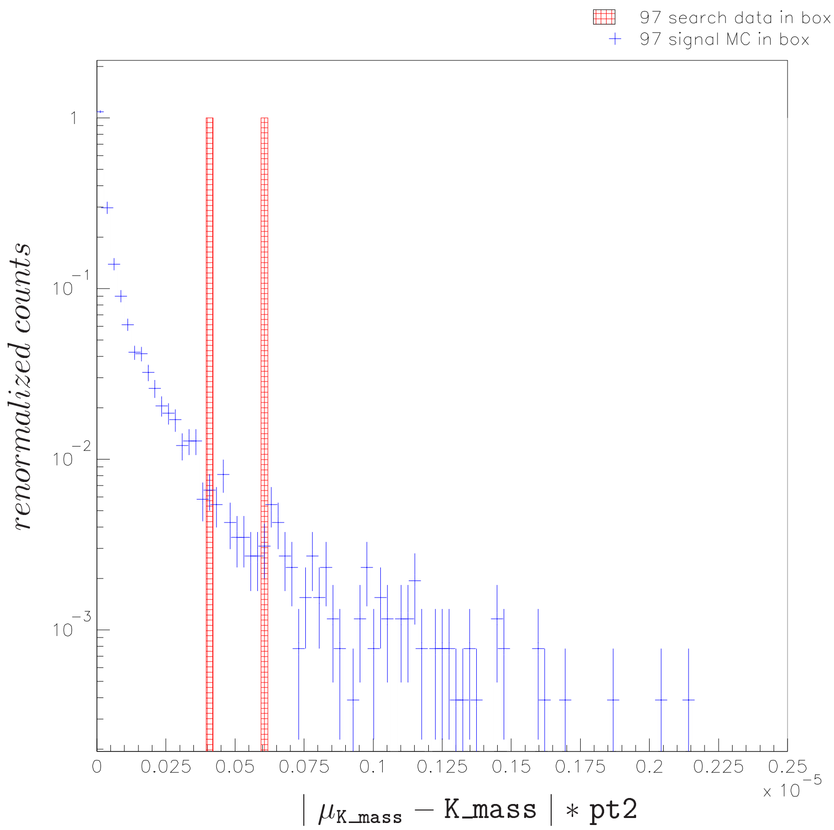

Figure 5.4 1997 data events in the signal box (red) and MC signal events (blue) passing all selection cuts for the combined parameter $|\ \mu_{\mathtt{K\_mass}}$-K_mass$|$*pt2. The data events fall on the tail of the signal MC distribution. The MC histogram has been renormalized to fit on the plot.

Fewer than 1.36% of true signal events would fall at pt2 values as large or larger than the two data events. Fewer than 4.32% would fall at larger values in $|\ \mu_{\mathtt{K\_mass}}$ - K_mass$|$*pt2. We conclude that the distribution of data events is inconsistent with what would be expected from actual signal events. However, we did observe two

Event label (from left to right in K_mass)	% of MC signal events higher in $\mid \mu_{\texttt{K\_mass}} - \texttt{K\_mass} \mid$*pt2
B1	2.50±0.22%
B2	4.32±0.28%

Table 5.5 If a given data event were signal, the probability it would fall where it did or farther away from zero in the combined parameter $\mid \mu_{\texttt{K\_mass}}$ - K_mass$\mid$*pt2.

events in our signal box when we had a background estimation of 0.55±0.14 events. We cannot rule out the possibility that there was an additional source of background that we did not identify. However, we can calculate a conservative 90% confidence level limit on the branching ratio of this decay mode by incorporating both our background estimation and the two seen data events. Such a limit would still be valid if there were unidentified backgrounds.

Using Feldman and Cousins technique [41], assuming a background that is gaussian with mean 0.55 and sigma 0.14, and 2 observed events, the 90% confidence level limit sensitivity is 5.36 events. The branching ratio limit is then:

$$BR(K_L \to \pi^0 \mu^\pm e^\mp) < \frac{5.36}{(Flux\ of\ K_L) * (Signal\ acceptance)} \tag{5.1}$$

Using our K_L flux of 2.55x10^{11} and Kpi0me acceptance of 5.276%, the branching ratio limit for the '97 data alone is:

$$BR(K_L \to \pi^0 \mu^\pm e^\mp(97)) < 3.98 \times 10^{-10} \tag{5.2}$$

We can also calculate a branching ratio limit using Poisson probability density functions (PDFs). We used a piece of code written by one of our collaborators (see

Appendix E) to calculate a 90% CL limit. That branching ratio limit is:

$$BR(K_L \to \pi^0 \mu^\pm e^\mp(97)) < 3.58 \times 10^{-10} \tag{5.3}$$

Chapter 6
1999 Data Analysis

The E799II KTeV data taken in 1999 was meant to be an extension of the 1997 data set, but in reality there were a few important differences. During the downtime, much work was done to reduce the background currents in the drift chambers, and so they were run at a lower threshold in '99. While this improved the acceptances for measurement analyses, it also ended up increasing the background for search analyses such as this one. The same was true of the lowering of the analysis magnet kick from $200MeV/c$ to $150MeV/c$. A few additional cuts had to be added to reduce the '99 backgrounds to acceptable levels.

It was originally hoped to directly combine the 1997 and 1999 data sets and so efforts were taken to make the reconstruction and cuts as similar as possible between the two. It was later decided that the combined 1997+1999 result would not be truly "blind" if we applied the '99 analysis cuts directly to the '97 data set to calculate the final numbers. However, we still used the '97 search data set to study the effects of the cuts we planned to use for '99. When the 1997 data is referred to in this Chapter and the 1999 Results Chapter (Chapter 7), the cuts used were the 1999 analysis equivalent cuts. The changes in conditions and calibrations between the two data sets meant that the numerical values of the 1999 cuts could not always be applied directly to the

1997 data. This has been noted when applicable.

6.1 Selection cuts

After KTSPILL, KTeVAna, Crunch, and Ntuple cuts were applied, the final selection cuts were applied using PAW. The selection cuts used in the 1999 final analysis were:

- Detector cuts:
  No triggers in veto detectors (`phvbar1` is .TRUE.)
  Ring counter energy (`RCmaxene`) $\leq 0.5 GeV$
  Spectrometer-Anti energy (`SAmaxene`) $\leq 0.4 GeV$
  Back-Anti energy (`BA1ene`) ≤ 12 *counts*†
  # TRD planes hit along e track (`npln_e`) > 0
  e track did not pass through any TRD dead region (`etrdhole`=0)
  Number of X track segments (`nxtrks`) < 10
  Number of Y track segments (`nytrks`) < 5
  Number of upstream track hit pairs (`xpairsup`) $= 0$
  Number of downstream track hit pairs (`xpairsdn`) ≤ 2
- Reconstruction cuts:
  exactly 2 charged tracks in the spectrometer
  fewer than 4 hardware clusters
  | "electron" track segment offset in X (`eoffx`) | $\leq 0.002m$
  | "electron" track segment offset in Y (`eoffy`) | $\leq 0.002m$
  | "muon" track segment offset in X (`muoffx`) | $\leq 0.001m$
  | "muon" track segment offset in Y (`muoffy`) | $\leq 0.001m$
  Vertex X and Y (`vtxx` and `vtxy`) in beam region
  $96m \leq$ Vertex Z (`vtxz`) $\leq 155m$
  Vertex χ^2 (`vtx_chi`) ≤ 10
  Transverse momentum squared (`pt2`) $\leq 0.0025\frac{GeV^2}{c^2}$
  $20GeV <$ Energy of kaon (`K_energy`) $\leq 200GeV$
- Particle ID cuts:
  Two photon mass (`pi0_mass`) within 1.25 sigma of mean‡
  $0.97 \leq$ Energy over momentum of electron (`eop_elec`) ≤ 1.03
  TRD CL of the "electron" track being a pion or muon (`prob_e`) ≤ 0.192†
  3x3 fusion χ^2 of electron cluster shape (`fuse3e`) ≤ 4

3x3 fusion χ^2 of 1st photon cluster shape (fuse3p1) ≤ 4
3x3 fusion χ^2 of 2nd photon cluster shape (fuse3p2) ≤ 4
| Muon momentum (pmu) | $\geq 7\frac{GeV}{c}$
muon planes with a matched track hit (nmumats) = 3 (maximum)

- Signal cuts:
  charged system mass (chrgmass) $\leq 0.364\frac{GeV}{c^2}$
  Signal π^0 kinematic (sgp0kine) $\geq$ -0.002
  K3pi π^0 kinematic (pp0_kine) < -0.06
  -0.08 < Ke4 kinematic (ke4_kine) < -0.005

The cuts marked with a dagger (†) have different placements when used to study the '97 data sets versus '99 data sets (the '99 cuts are listed here), and are discussed within the parameter definitions (Section 3.4). For the pi0_mass cut (marked with a double dagger [‡]), the mean and sigma used was taken from the normalization mode data for cutting on data events, and was taken from the normalization mode MC for cutting on MC events. Again, the final requirement was that the events fell within the "signal box" to be considered Kpi0me events. This requirement was *not* considered to be part of the "signal selection cuts" so that we could study background events in the area around the signal box, called the study plot.

The number of MC events which passed all the nominal signal selection cuts was used in the calculation of signal acceptance and expected background within the box. From the plot shown in Figure 6.1, one can see that 36,483 out of 999,800 MC signal events for 1999 fell within the signal box (in red). The TRD cut (prob_e) was not included when looking at 1999 MC events because the 1999 MC behavior of the TRDs for this parameter did not match the data for a real electron. Using Ke3 data, the prob_e cut was selected so as to have a 98.0% electron acceptance (see the definition of prob_e in Section 3.4.5 for a details). Combining the electron acceptance with the number of events in the box in Figure 6.1 led to a signal acceptance of (0.03649×0.980)=3.576±0.018%. If '99-type cuts were applied to signal MC

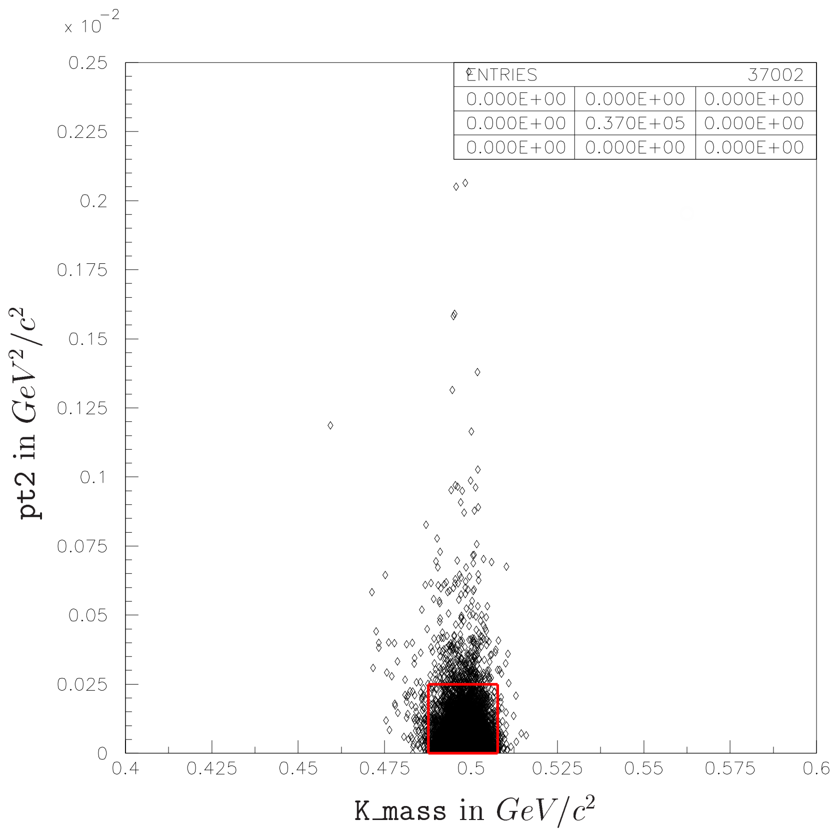

Figure 6.1 The "study plot" of signal MC generated for the 1999 data set. All selection cuts except the TRD CL cut are applied. There are 36,483 events inside the red signal box. The area of the plot is 100 times the area of the signal box.

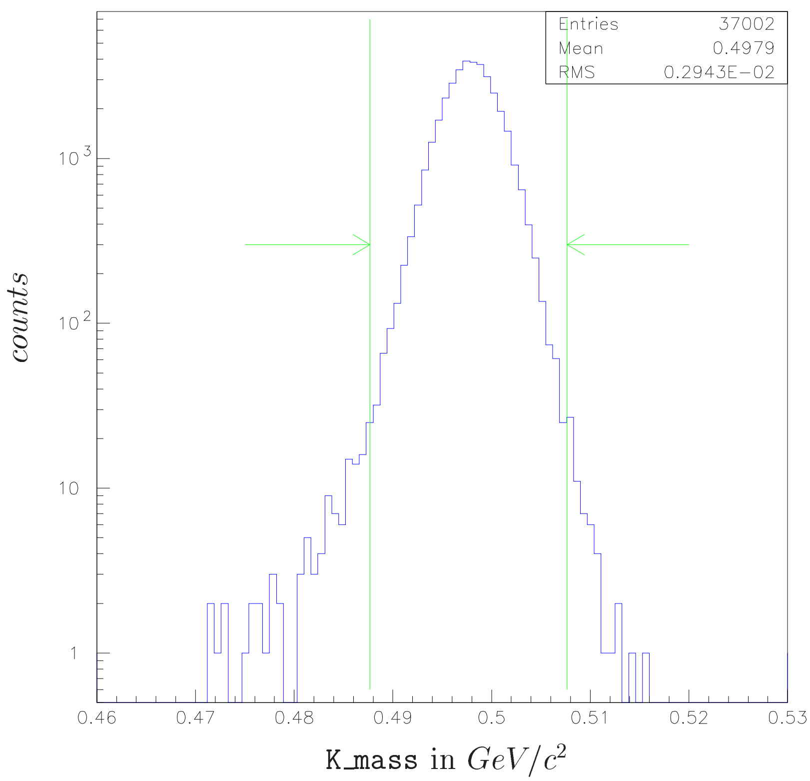

Figure 6.2 A projection onto the X-axis (`K_mass`) of MC signal events passing all cuts (except the signal box) from Figure 6.1. There were 36,828 events inside the signal box `K_mass` limits, shown here by the green vertical lines. Note the log scale on the Y-axis.

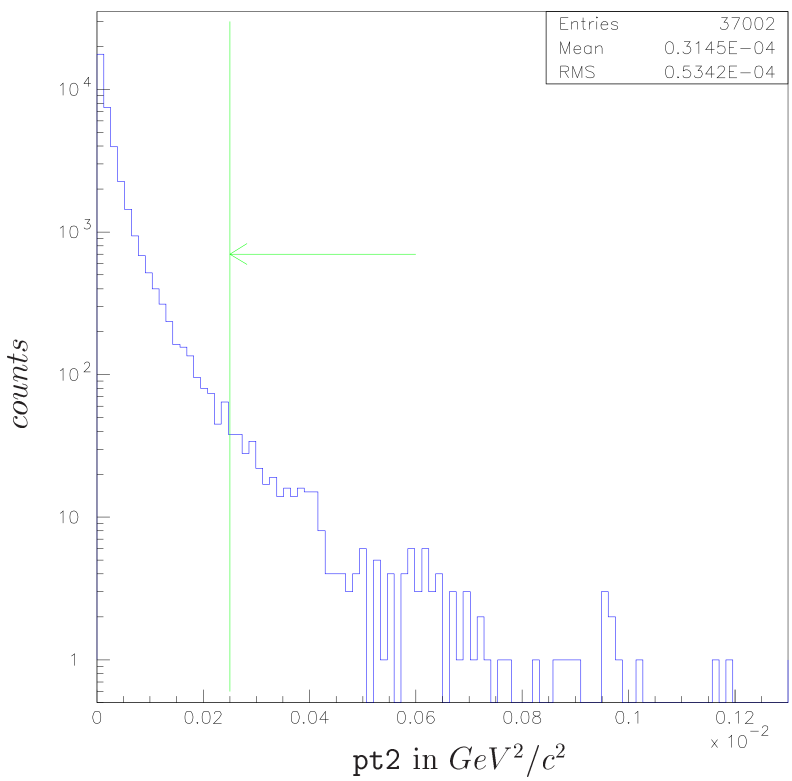

Figure 6.3 A projection onto the Y-axis (`pt2`) of MC signal events passing all cuts (except the signal box) from Figure 6.1. There were 36,624 events inside the signal box `pt2` limit, shown here by the green vertical line. Note the log scale on the Y-axis.

generated for the '97 run, the signal acceptance was 3.751±0.061%. The `prob_e` cut was included this time, but the cut value itself was placed at 0.049 to result in the equivalent 98% electron acceptance. The difference in acceptances was not due to just one or two cuts, but was rather an accumulation of small differences. However, a study of several two-track decay modes show the same discrepancy, and so this effect should cancel in the normalization mode ratio. Please see Appendix B for further details and a list of acceptances cut by cut.

A subset of the signal selection cuts were applied to '99 trigger 2 data to get a sample of well-reconstructed normalization events (K3pis). For normalization events, parameters with a "mu" or "e" distinction refer to one or the other of the charged pions in the K3pi decay with the pion depositing the higher energy calorimeter cluster carrying the "e" label. The "detector" and "reconstruction" cuts given for the '99 signal selection cuts were also used in the '99 normalization selection cuts. The '99 normalization "particle ID" and "signal" cuts are listed here.

- Particle ID cuts:
  Two photon mass (`pi0_mass`) within 1.25 sigma of mean
  3x3 fusion χ^2 of 1st photon cluster shape (`fuse3p1`) ≤ 4
  3x3 fusion χ^2 of 2nd photon cluster shape (`fuse3p2`) ≤ 4
- Signal cuts:
  charged system mass (`chrgmass`) $\leq 0.364 \frac{GeV}{c^2}$
  within "signal" box

The number of normalization data events in the acceptance box was compared to the number of normalization MC events in the acceptance box to get the kaon flux for the data used in this search. Trigger 2 had an on-line prescale factor of 500:1 and our trigger 2 data set had an tape-split prescale factor of 40:1 for 1999. In the MC K3pi events, 42,018 out of 998,070 events fell within the signal box, leading to a K3pi acceptance of 4.210±0.020%. The distribution of '99 normalization data events is shown in Figure 6.4, with projections onto the axes in Figures 6.5 and 6.6. In the data from Trigger 2, 81,949 events fell within the signal box. The flux was found by:

$$Flux(K_L) = \frac{(\#\, of\ K3pi\, data\, accepted) \times (Prescales\, on\, K3pi\, events)}{(BR(K_L \rightarrow \pi^+\pi^-\pi^0)) \times (BR(\pi^0 \rightarrow 2\gamma)) \times (K3pi\ acceptance)} \quad (6.1)$$

The K_L flux for the 1999 data set was then $3.137 \pm 0.015_{stat} \pm 0.052_{BRs} \mathrm{x} 10^{11}$. One of the reasons this flux appears small compared to other KTeV analyses, is that it includes only runs and spills which were considered "good" for this analysis since a mask for bad spills was applied during MC generation. Bad spills removed 6.34±0.01% of the trigger 2 data before any other cuts were applied (details in Appendix D). Also, there were discrepancies in the acceptances for data versus MC which are discussed in Appendix B. A flux corrected for the data/MC mismatch would be $3.437 \pm 0.017_{stat} \pm 0.055_{BRs} \pm 0.148_{syst} \times 10^{11}$. For the '97 data set with '99-type cuts, a MC acceptance of 43,210 out of 978,640 events (4.415±0.021%), a data acceptance

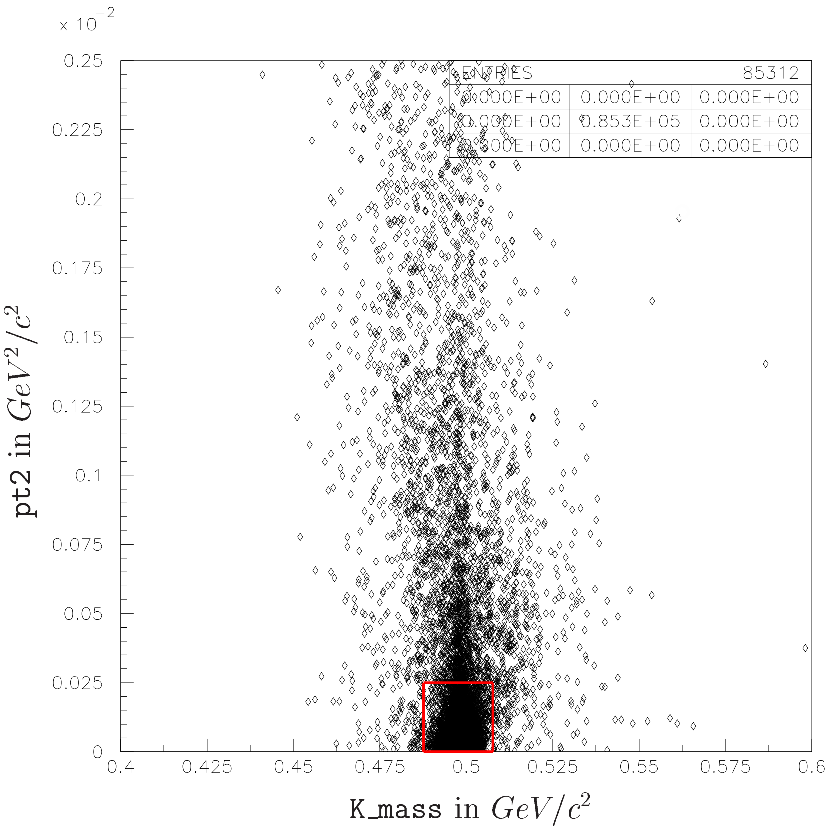

Figure 6.4 1999 trigger 2 data events reconstructed as K3pi's and passing all the normalization selection cuts except the "signal" box. The number of events inside the red signal box (81,949) was used in calculating the 1999 kaon flux.

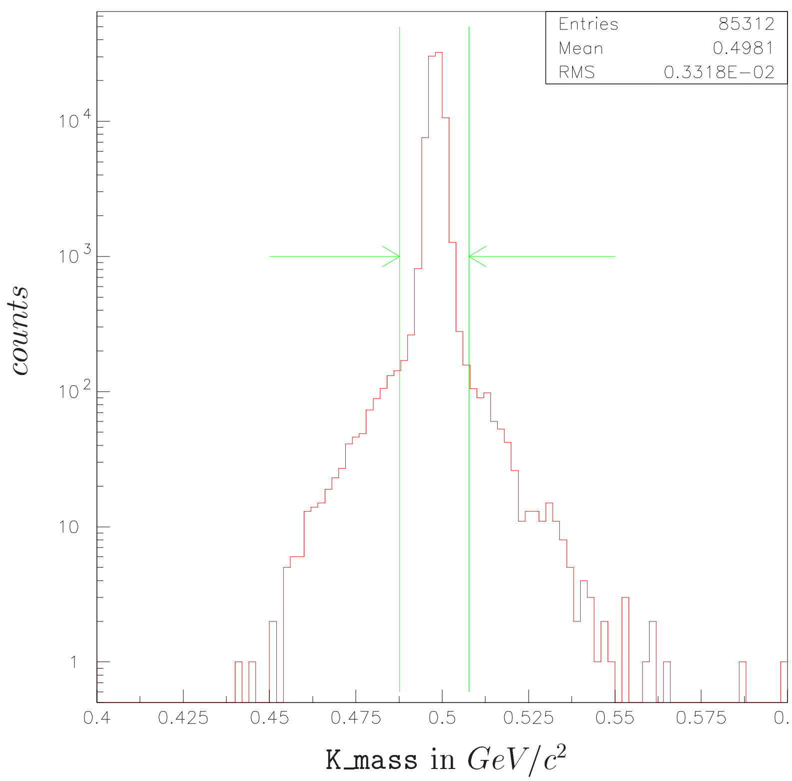

Figure 6.5 A projection onto the X-axis (`K_mass`) of normalization data events passing all cuts (except the signal box) from Figure 6.4. There were 83,923 events inside the signal box `K_mass` limits, shown here by the green vertical lines. Note the log scale on the Y-axis.

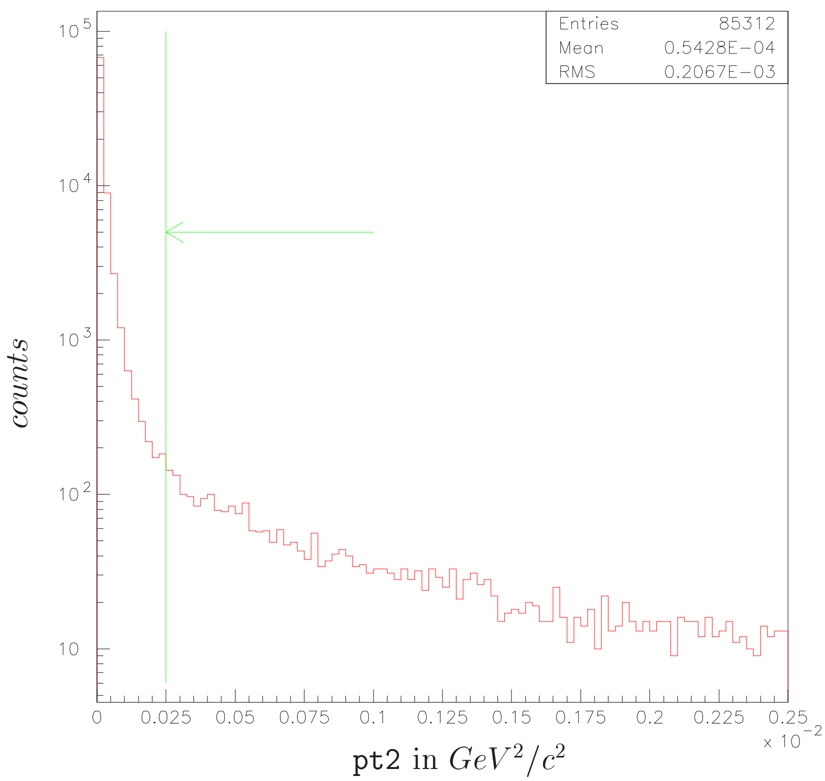

Figure 6.6 A projection onto the Y-axis (`pt2`) of normalization data events passing all cuts (except the signal box) from Figure 6.4. There were 82,212 events inside the signal box `pt2` limit, shown here by the green vertical line. Note the log scale on the Y-axis.

of 145,246 events, and an offline-split prescale factor of 20:1, gave a '97 K_L flux of $2.619 \pm 0.012_{stat} \pm 0.042_{BRs} \mathrm{x} 10^{11}$.

6.2 1999 Backgrounds

Since all background modes except $K_L \to \pi^0\pi^+\pi-$ (K3pi), $K_L \to \pi^0\pi^\pm e^\mp\nu_e$ (Ke4), and $K_L \to \pi^\pm e^\mp\nu_e$ (Ke3) were shown to be negligible in the '97 analysis, we considered only these three backgrounds for the '99 analysis. A summary of the estimates is shown in Tables 6.1 and 6.2.

Background name	# expected in '97 study plot	# expected in '99 study plot
K3pi	1.3±0.7	5.4±3.3
Ke4	2.6±0.3	3.1±0.4
Ke3 (total)	6.9±3.0	9.5±4.2
Ke3dk	5.9±2.9	7.5±4.0
Ke3punch	1.1±0.8	2.0±1.2

Table 6.1 The significant backgrounds considered for Kpi0me and the number of each expected to pass all signal selection cuts to fall in the `pt2` versus `K_mass` plot according to MC. For the Ke3 background, it was time consuming to generate sufficient MC to get any events to pass all cuts, hence the large uncertainties.

6.2.1 K3pi Background in 1999

Once again it was impossible to generate enough correctly modeled MC to simulate the K3pi background satisfactorily, so we again used trigger 2 normalization data reconstructed as Kpi0me events. As described in Chapter 4 for the '97 analysis, the same technique of applying subsets was used for the '99 analysis, with similar kinds

Background name	# expected in '97 box	# expected in '99 box
K3pi	negligible	negligible
Ke4	< 0.08	< 0.09
Ke3	0.24±0.10	0.43±0.13
Total	0.24±0.10	0.43±0.13

Table 6.2 The significant backgrounds considered for Kpi0me and the number of each expected to pass all signal selection cuts to fall within the signal acceptance box. For the Ke3 background, it was time consuming to generate sufficient MC to get any events to pass all cuts, hence the large uncertainties. Since no events appeared in the box for the Ke4 MC samples, the expectation given here is a 90% CL limit and is not included in the total.

of subsets and results. The subsets are listed in Table 6.3 and their effects are listed in Table 6.4. The distribution of K3pi events in the study plot is shown in Figure 6.7. Our estimate of the expected background in the study plot from K3pi decays was 5.4±3.3, with the same caveat as before: this estimate could be pushed high by the presence of decays other than K3pi in the data set used. We did not expect this background to appear in the signal box.

6.2.2 Ke4 Background in 1999

As in the '97 analysis, the low branching ratio of Ke4's made it possible to generate a MC sample (with accidentals) that was larger than the expected number of Ke4 decays in our data set. The plot of 1999 Ke4 MC in Figure 6.8 shows 83 events out of approximately 4.1×10^8 generated events (almost 27 times the expected flux) passed all signal selection cuts except the TRD (`prob_e`) cut. Using this acceptance, the

Cut subset definitions	
Electron ID cuts	Kinematic cuts
$0.97 \leq$ `eop_elec` ≤ 1.03 `fuse3e` ≤ 4 `prob_e` ≤ 0.192	`pp0_kine` $<$ -0.06 -0.08 $<$ `ke4_kine` $<$ -0.005 `sgp0kine` $\geq$ -0.002
Muon ID cuts	Base cuts
muon cluster energy $< 1GeV$ $\lvert$`pmu`$\rvert \geq 7GeV$ a hit in muon bank #3 `nmumats` $= 3$ no hit in hadron anti (HA)	All signal selection cuts not in Kinematic, Electron ID, or Muon ID subsets `K_mass` $\geq 0.4GeV/c^2$ (lower edge of study plot)

Table 6.3 The subsets of cuts applied to trigger 2 data to estimate the number of background K3pi events in the 1999 study plot. The electron ID, muon ID, and kinematic subsets were assumed to be independent of each other.

Cuts used	# 1999 trig. 2 data in study plot	reduction factor
Base cuts only	20,189	1
Base + Electron ID cuts	21±4.6	961.±210.
Base + Kinematic cuts	3±1.7	6,730±3,886
Base + Muon ID cuts	1,752±42	11.52±0.29
Total reduction factor		$7.5 \pm 4.6 \times 10^8$
K3pi's expected in '99 study plot	5.4±3.3	

Table 6.4 The effects of applying subsets of signal selection cuts to trigger 2 normalization data reconstructed as Kpi0me to estimate the expected number of K3pi events passing all signal selection cuts. If all the cuts were applied at once, no events passed, and it was impossible to generate enough MC to get reasonable statistics for this background.

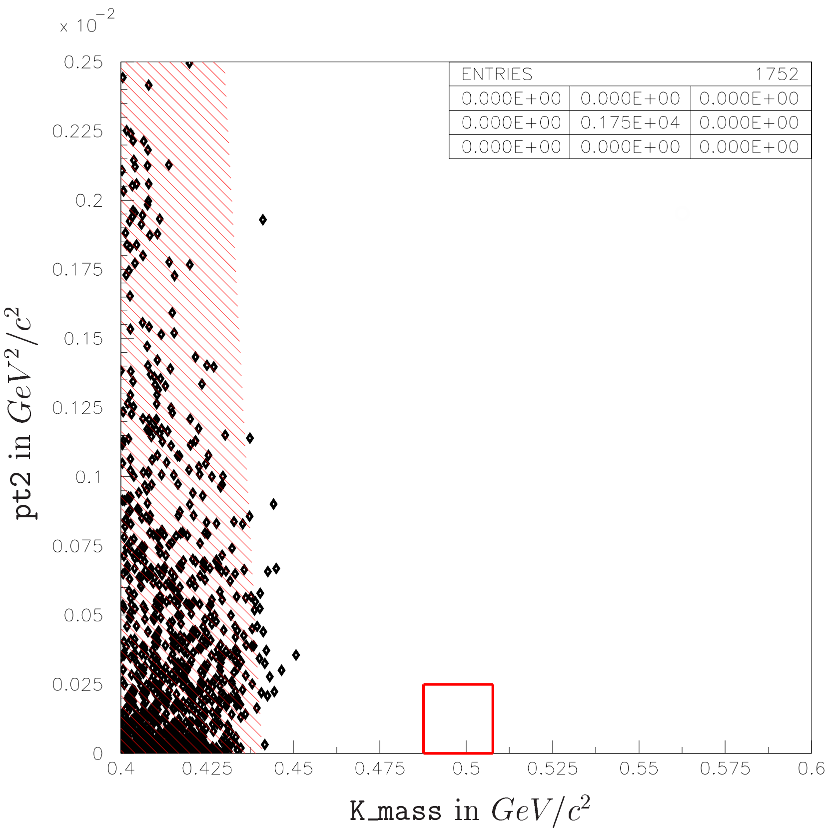

Figure 6.7 1999 K3pi trigger 2 data reconstructed as Kpi0me events. The events shown here did not yet have the "kinematic cuts" or "electron ID cuts" as defined in Table 6.4. As can be seen, these events reconstructed far below the signal box, and so the background from K3pi's in the box was negligible. The red shaded region was defined as the K3pi area of the study plot for looking at search data.

prob_e cut reduction factor (98%), our calculated kaon flux, and the Ke4 branching ratio[8], we expected 3.1±0.4 Ke4 data events to pass all signal selection cuts in 1999. This agreed with what was seen in the '99 search data study plot. No MC events were seen in the signal box. We instead put a 90% CL limit on the number of Ke4 events expected in the box : < 0.09 events.

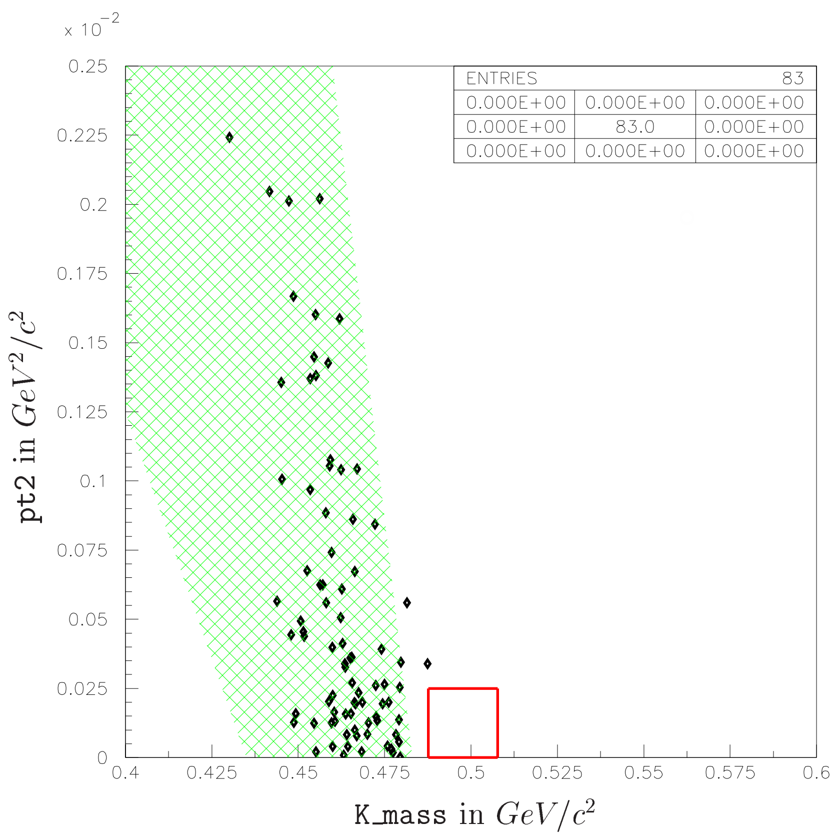

Figure 6.8 Ke4 background MC generated for the 1999 data set. The number of events shown here is for approximately 27 times the expected flux and does not include the TRD CL cut. The green shaded region was defined as the Ke4 area of the study plot for looking at search data.

The plot of 1997 Ke4 MC with '99-like cuts in Figure 6.9 shows 77 out of approximately 4.1x10^8 events (30 times the expected flux) passed all signal selection cuts. After scaling we then expected 2.6±0.3 Ke4 data events to pass all signal selection cuts in 1997. No MC events in the signal box for this sample meant we expected less than 0.08 Ke4 background events in the box for the search data at a 90% CL.

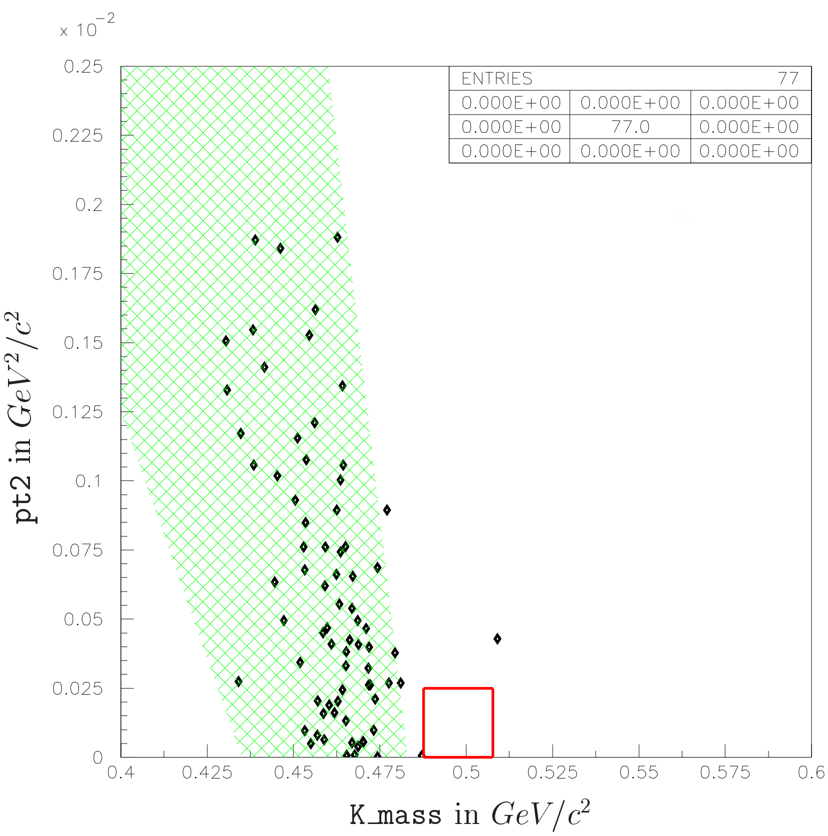

Figure 6.9 Ke4 background MC generated for the 1997 data set with 99-like cuts applied. The number of events shown here is for approximately 30 times the expected flux. The green shaded region is defined as the Ke4 area of the study plot for looking at search data.

6.2.3 Ke3 Background in 1999

The branching ratio of the Ke3 decay is even larger than that for K3pi decays, but the difficulty of finding two accidental photons that add up to the neutral pion mass suppresses this background. Even so, we could not generate MC equal to the flux in our data sample. However, the pion also had to decay into a muon or punch through the detector into the muon banks to satisfy the trigger, so we generated MC in which this decay or punch-through occurred every time. Then by estimating the percentage of Ke3's which should pion-decay or punch-through, we could predict the number of events expected in our study plot. Even with the forced decay and punch-through, we needed to remove the `pi0_mass` cut in order to improve our statistics for comparison studies. Both of these signal-faking backgrounds were significant in the study plot.

Ke3dk's and Ke3punch's

For the '99 data set, we generated about 2x10^9 Ke3 forced decays (which I will call "Ke3dk"s), and about 3.0x10^9 Ke3 forced punch-throughs (which I will call "Ke3punch"s). To determine what fraction of the Ke3's decayed, we generated regular Ke3 MC, which correctly modeled the momentum dependence of the pion decay, and counted the number of times the code entered the decay subroutine during event generation. For both '97 and '99 MC sets, the pion in a Ke3 decayed 5.71$\pm$0.03% of the time. In comparison, a charged pion in a K3pi event decayed 6.34$\pm$0.04% of the time using this method of calculation.

The percentage of time a pion punch-through occurred was not modeled in our MC. To calculate the expected Ke3punch flux, we looked at the variable which showed the biggest difference in histogram shape between Ke3dk and Ke3punch: `vtx_chi`. The same MC events generated to estimate the background were reconstructed as "good" Ke3 decays using the following cuts:

- Detector cuts:
  same as signal selection cuts (p. 186)
- Reconstruction cuts:
  exactly 2 charged tracks in the spectrometer
  fewer than 3 hardware clusters
  $|$ `eoffx` $| \leq 0.002m$
  $|$ `eoffy` $| \leq 0.002m$
  $|$ `pioffx` $| \leq 0.002m$
  $|$ `pioffy` $| \leq 0.002m$
  $96m \leq$ `vtx_z` $\leq 155m$
  `K_energy` $\leq 200GeV$
- Particle ID cuts:
  $0.97 \leq$ `eop_elec` ≤ 1.03
  `fuse3e` ≤ 4
  $|$ charged pion momentum $| \geq 7\frac{GeV}{c}$
  # muon planes with a matched track hit = 3 (maximum)
- Signal cuts:
  total reconstructed mass $\leq 0.6\ \frac{GeV}{c^2}$

The `prob_e` cut was not included since it could not be applied to the MC. Data events from Trigger 2 were reconstructed with the same code and put through the same cuts. Any well-reconstructed Ke3 event should pass most of these cuts. The important cuts were on the momentum of the pion and the number of hits in the muon bank. These ensured that the pion had decayed or punched-through the detector in all the data events that passed. As can be seen in Figure 6.10, the MC Ke3punch's reconstructed much more cleanly than the MC Ke3dk's, and data was somewhere between the two.

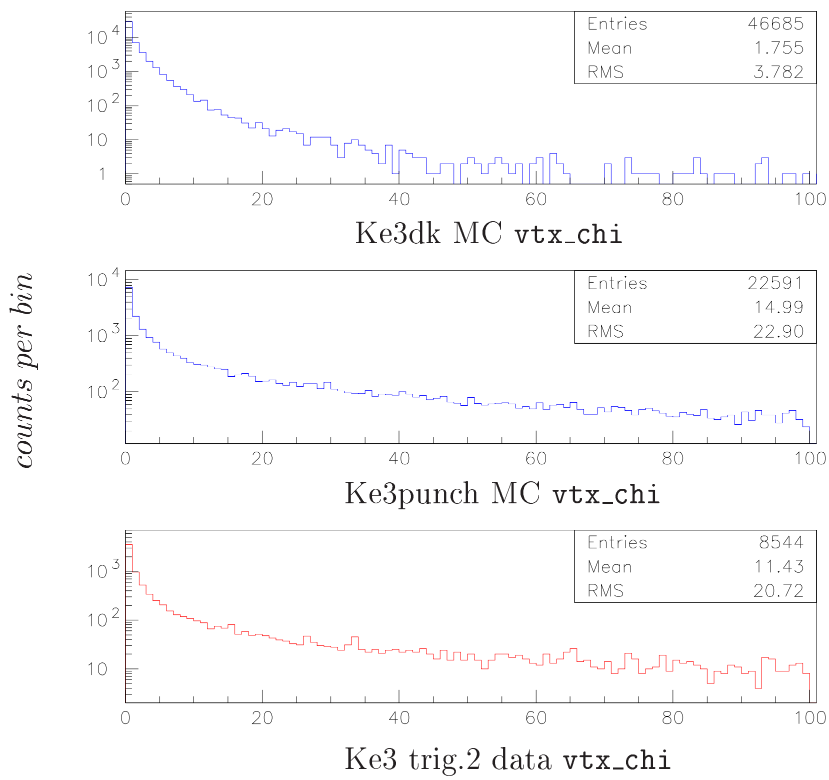

Figure 6.10 Ke3punch MC (top plot), Ke3dk MC (middle plot), and trigger 2 data (bottom plot) from 1999, reconstructed as Ke3 decays with a muon signature. The better the event was able to be reconstructed, the closer to 1.0 this parameter should be. Punch-throughs reconstructed much more cleanly than pion decays, and the data contained some combination of the two.

Three different methods were used to determine what fraction of the Ke3 data was Ke3punch. For the first method, we took the means of the MC distributions and combined them linearly to match the data mean:

$$(\mu_{punch} * f) + (\mu_{dk} * (1 - f)) = \mu_{data} \tag{6.2}$$

This resulted in 26.9±1.9% Ke3punch and 73.1±1.9% Ke3dk for 1999 MC after cuts were made. The combined MC set is shown overlaid on the data in Figure 6.11. To get this same ratio back at the generation level (so we could predict the background for the signal decay), we had to correct for the acceptance. Ke3punch MC acceptance for the Ke3 cuts was $1.572 \pm 0.007 \times 10^{-5}$. Ke3dk MC acceptance was $1.130 \pm 0.008 \times 10^{-5}$. Since 1.39 Ke3punch's were accepted for every one Ke3dk accepted, the Ke3punch percentage was reduced by this factor. Renormalizing then gave us a generation-level Ke3punch to Ke3dk ratio of 20.9±1.5% to 79.1±2.0%. The same process for the '97 data set gave an after-cuts ratio of 23.3±0.4% to 76.7±0.6% and a generation ratio of 18.0±0.5% to 82.0±0.5% Ke3punch to Ke3dk events.

The second method looked at the raw number of events in two bins: `vtx_chi` from 0 to 20, and from 30 to 100. Restricting `vtx_chi` from 30 to 100 resulted in nearly all the events being Ke3dk's (4,076 MC Ke3dk events to 132 MC Ke3punch events). There were 1,140 Ke3 data events in this bin. The flux correction to the Ke3dk MC

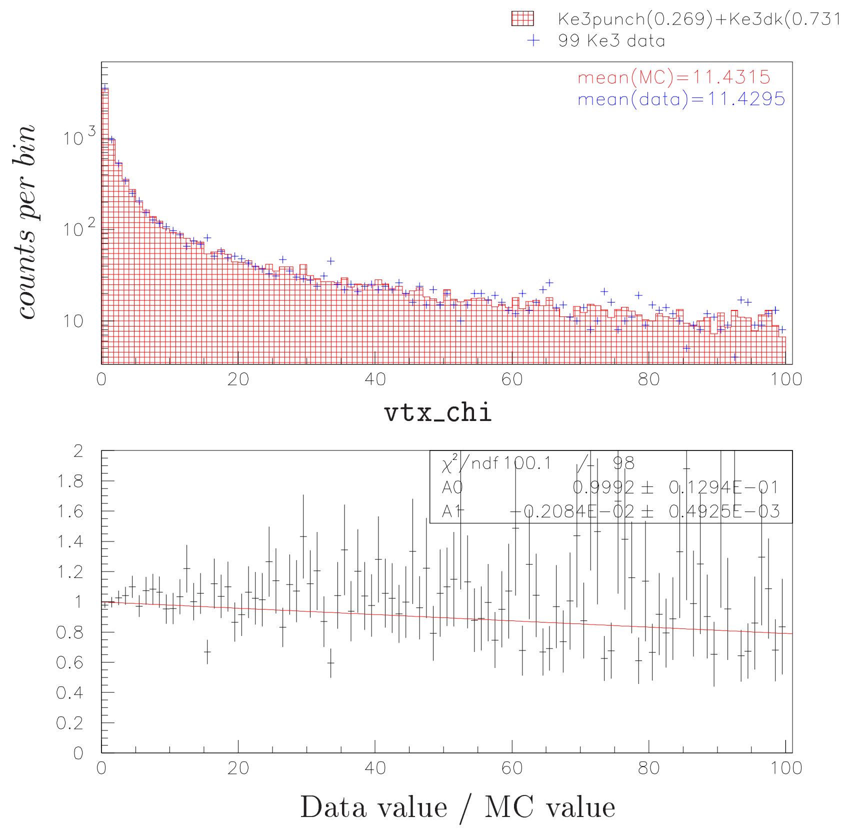

Figure 6.11 The means of the two MC distributions in Figure 6.10 were weighted and combined to equal the mean of the data distribution. The result is 73% Ke3dk and 27% Ke3punch. The MC events were combined in those proportions above and overlaid with the data, showing good agreement. When the ratio was propagated back to the generation level, 209 Ke3punch events occurred for every 791 Ke3dk events.

was then

$$\frac{\#\ Ke3\ data}{\#\ MC\ Ke3dk} = \frac{1,140}{4,076} = 0.280 \pm 0.009 \tag{6.3}$$

In the 0 to 20 `vtx_chi` bin, there were 17,110 MC Ke3dk's, 46,404 MC Ke3punch's, and 7,032 Ke3 data events. Of the data events, 17,110*0.280 = 4,785±164 events should have been decays. That left 2,246±185 punch-through data events, giving a flux correction to the Ke3punch MC of

$$\frac{\#\ Ke3\ data}{\#\ Ke3punch} = \frac{2,246}{46,404} = 0.048 \pm 0.004 \tag{6.4}$$

Applying the appropriate flux corrections to the total number of Ke3dk and Ke3punch MC events in the full range of `vtx_chi` (0 to 100) gave 26.3±2.3% Ke3punch and 73.7±2.3% Ke3dk after all cuts. (see Figure 6.12). For the '97 MC, this method gave 22.8±0.6% Ke3punch and 77.2±0.6% Ke3dk after all cuts. This method was in good agreement with the first method.

The third method used the same two bins in `vtx_chi` as the second method. There was concern expressed that the normalization mode showed a slope in the `vtx_chi` parameter. Both the mean fitting and raw event number fitting methods resulted in a nearly flat data / MC plot. The third method attempted to take into account the slope seen in the normalization ratio plot. We looked at the normalization mode in the two `vtx_chi` bins (0 to 20, and 30 to 100) and calculated the ratio of MC to data, giving

$$R_p = \frac{\#\ data\ in\ punch\ bin}{\#\ MC\ in\ punch\ bin} = 1.937 \pm 0.012 \tag{6.5}$$

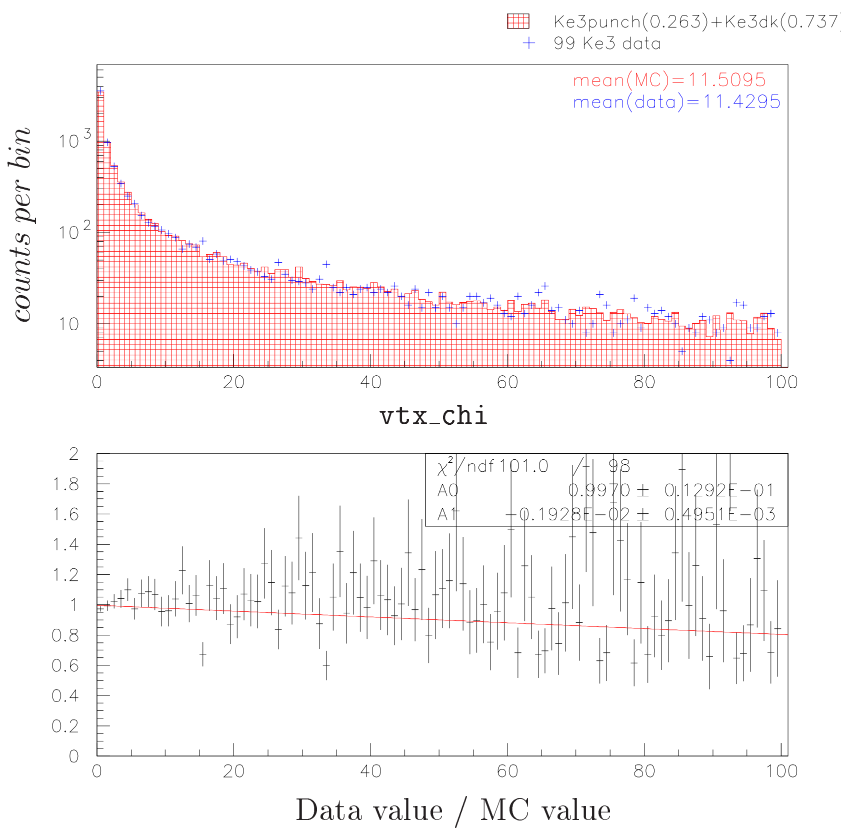

Figure 6.12 Histograms of `vtx_chi` for the Ke3 data and weighted Ke3 MC backgrounds. The Ke3punch's are 26% of the MC events,and the Ke3dk's are the other 74matching the number of events in the two bins 0 to 20, and 30 to 100.

$$R_d = \frac{\#\ data\ in\ decay\ bin}{\#\ MC\ in\ decay\ bin} = 1.16 \pm 0.16 \tag{6.6}$$

From the second method, we knew the raw numbers of events in each bin. We defined the number of Ke3punch events in the 0 to 20 bin as M_{pp}, the number of Ke3dk events in this bin as M_{pd}, and the number of Ke3 data events in this bin as D_p. Similarly we defined the numbers of raw events in the 30 to 100 bin as M_{dp} (Ke3punch), M_{dd} (Ke3dk), and D_d (Ke3 data). We then solved the linear system of equations

$$D_p = R_p[\alpha_p M_{pp} + \alpha_d M_{pd}] \tag{6.7}$$

$$D_d = R_d[\alpha_p M_{dp} + \alpha_d M_{dd}] \tag{6.8}$$

for the weights α_p and α_d. The result was α_p=0.025±0.002 and α_d=0.242±0.035. The αs were weighting factors in the equations, so we had to multiply them by the raw percentages:

$$\%\ Ke3punch\ unnormalized = \alpha_p * \frac{\#\ of\ Ke3punch\ in\ (0-100)}{\#\ of\ Ke3punch\ +\ Ke3dk\ in\ (0-100)} \tag{6.9}$$

$$\%\ Ke3dk\ unnormalized = \alpha_d * \frac{\#\ of\ Ke3dk\ in\ (0-100)}{\#\ of\ Ke3punch\ +\ Ke3dk\ in\ (0-100)} \tag{6.10}$$

and renormalize:

$$\%\ Ke3punch\ after\ cuts = \frac{\%\ Ke3punch\ unnormalized}{\%\ Ke3punch\ unnorm. + \%\ Ke3dk\ unnorm.} \tag{6.11}$$

$$\%\ Ke3dk\ after\ cuts = \frac{\%\ Ke3dk\ unnormalized}{\%\ Ke3punch\ unnorm. + \%\ Ke3dk\ unnorm.} \tag{6.12}$$

which gave 17.7±1.1% Ke3punch and 82.3±1.1% Ke3dk after all Ke3 cuts (see Figure 6.13). For 1997, the after-cuts ratio for this method was 20.0±0.5% Ke3punch to 80.±0.5% Ke3dk events.

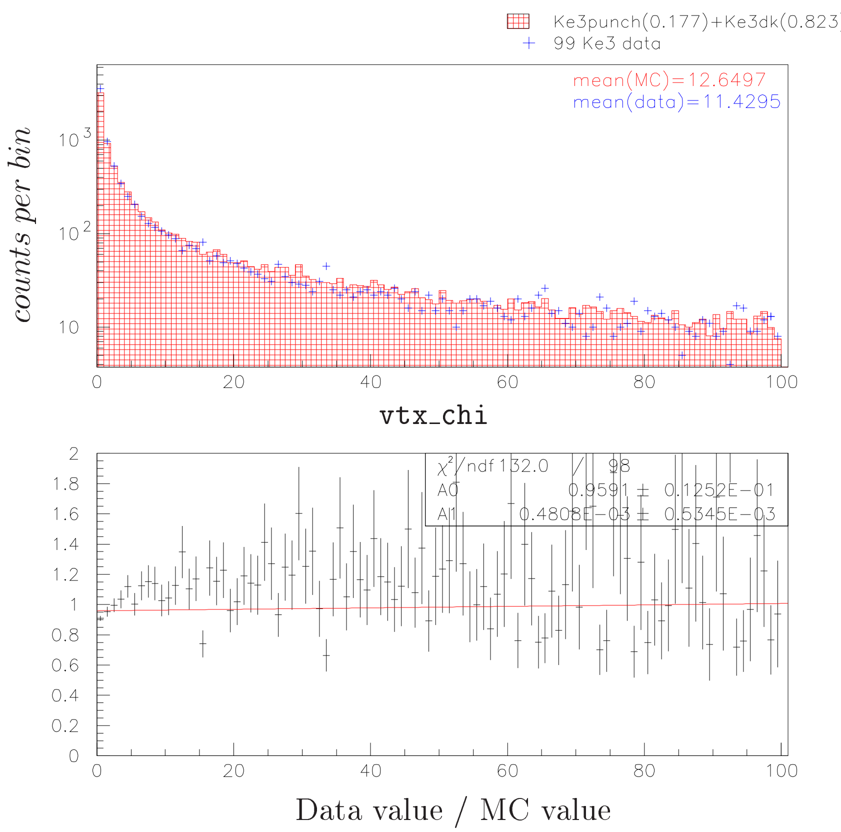

Figure 6.13 Histograms of **vtx_chi** for the Ke3 data and weighted Ke3 MC backgrounds. The Ke3punch's are 18% of the MC events,and the Ke3dk's are the other 82%. This was the ratio of punch-to-decay as determined by looking at raw event numbers in two bins and correcting for the slope seen in the normalization data / MC **vtx_chi** ratio plot.

The third method gave a lower Ke3punch component in the background than the first two methods. It also corrected for the slope in the normalization data/MC ratio plot. We therefore decided to use the numbers generated using the third method to estimate the number of background Ke3's in the study plot. In the end, the MC Ke3 background estimate was not very sensitive to this ratio.

Ke3's in Study Plot

If all cuts were applied to the '99 Ke3dk and Ke3punch MC, then 1 Ke3dk event and 6 Ke3punch events appeared in the study plot. This was not enough statistics for a reasonable background estimate. Since the neutral pion in the Ke3 background events came from accidental activity, we could remove the `pi0_mass` cut and expect the distribution of this background in the study plot to remain the same. This was what was done in Figures 6.14 and 6.15. When the MC was scaled for the expected flux and `pi0_mass` cut suppression, we expected to see 7.5±4.0 Ke3dk events and 2.0±1.2 Ke3punch events in our study plot, for a total of 9.5±4.2 events in the Ke3 area. Figures 6.16 and 6.17 show '97 Ke3dk and Ke3punch MC events without the `pi0_mass` cut. After scaling, we expected 5.9±3.0 Ke3dk events and 1.1±0.8 Ke3punch events in the plot, or combined, 7.0±3.1 events in the Ke3 area. We saw 8 events in the '99 data, and we saw 4 events in the '97 data.

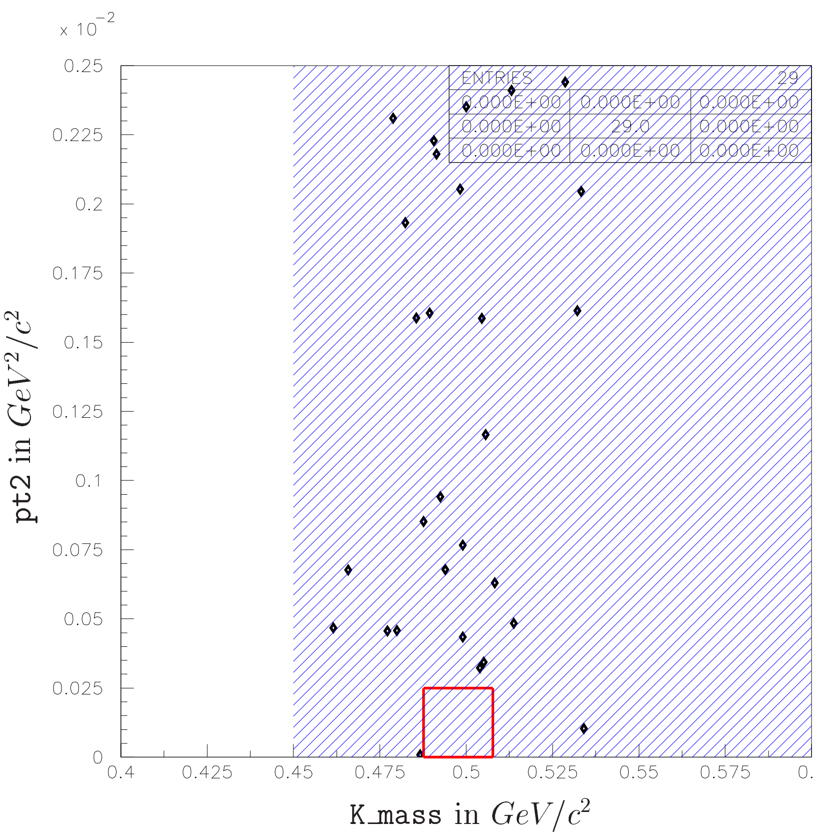

Figure 6.14 Ke3dk background MC generated for the 1999 data set. The number of events shown here is for approximately 32% of the expected flux and does not include the `pi0_mass` cut. If the `pi0_mass` cut was applied, 1 event remained. The blue shaded region was defined as the Ke3 area of the study plot for looking at search data.

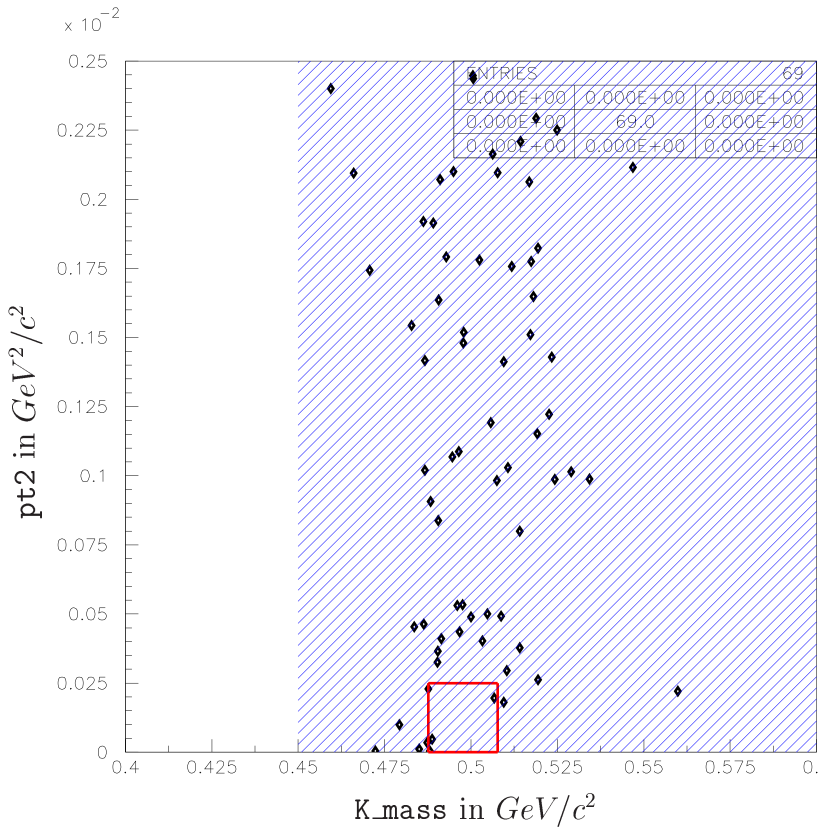

Figure 6.15 Ke3punch background MC generated for the 1999 data set. The number of events shown here is for approximately 310% of the expected flux and does not include the `pi0_mass` cut. With the `pi0_mass` cut, we saw 6 remaining events outside the box. The blue shaded region was defined as the Ke3 area of the study plot for looking at search data.

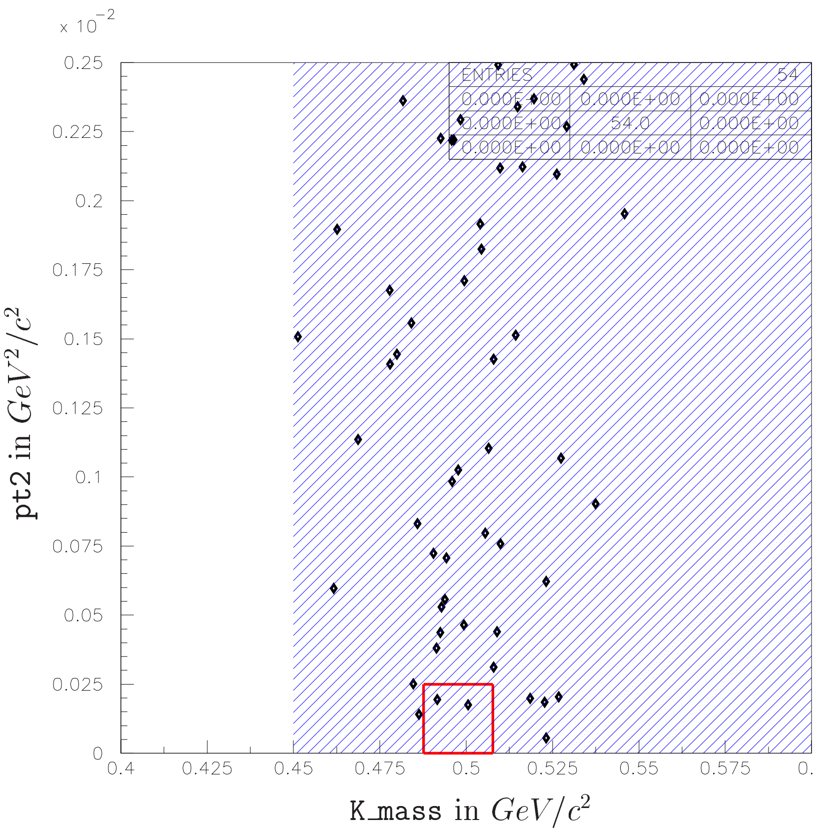

Figure 6.16 Ke3dk background MC generated for the 1997 data set. The number of events shown here is for approximately 67% of the expected flux and does not include the `pi0_mass` cut. Adding the `pi0_mass` cut left 3 events. The blue shaded region was defined as the Ke3 area of the study plot for looking at search data.

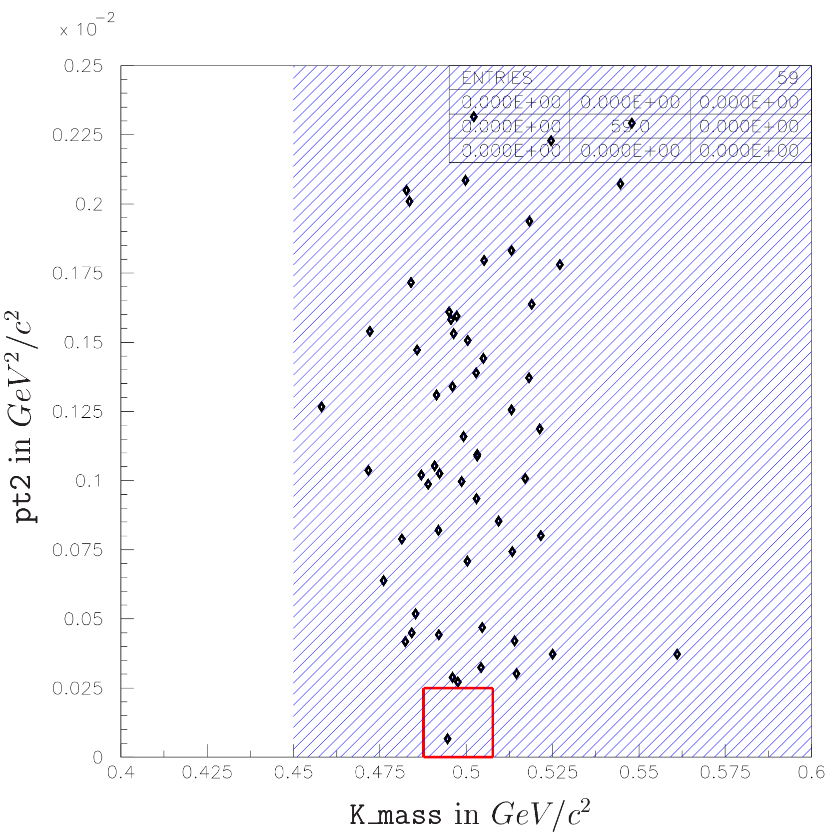

Figure 6.17 Ke3punch background MC generated for the 1997 data set. The number of events shown here is for approximately 182% of the expected flux and does not include the `pi0_mass` cut. Adding the `pi0_mass` cut left 1 event. The blue shaded region was defined as the Ke3 area of the study plot for looking at search data.

Ke3's in Signal Box

There was a good possibility that we could see background Ke3 events in our signal box, so it was important to estimate the expected rate as accurately as possible. Two different methods were used to estimate the signal box background as a cross-check of consistency.

For one method, removing the `pi0_mass` cut resulted in 2 Ke3dk plus 2 Ke3punch MC events in the box for '97 and 0 Ke3dk plus 3 Ke3punch for '99. Rescaling of these numbers by the `pi0_mass` factor gave an expected background of 0.24$\pm$0.18 events for '97. We did not generate enough '99 Ke3dk MC to get any events in the box after all cuts, even with the `pi0_mass` cut removed. Therefore, our expected Ke3 background in the box in '99 for this method is an upper limit. If one Ke3dk event had appeared in the box, the background estimate would have been 0.26$\pm$0.29 Ke3dk events, or 0.34$\pm$30 events combined, so we will use this number as our upper limit.

Another estimation looked at the slice in `K_mass` above the signal box and assumed that the events were uniformly distributed in `pt2` (an assumption supported by Figure C.7). The slice was ten times larger than the signal box. From Figures 6.16 and 6.17, there were 21 Ke3dk events and 29 Ke3punch events in the slice. Scaling by the expected flux, `pi0_mass` cut suppression, and dividing by 10, gave a signal box background of 0.21$\pm$0.12 Ke3dk events and 0.05$\pm$0.04 Ke3punch events, for a combined MC background of 0.27$\pm$0.13 for 1997. From Figures 6.14 and 6.15, the 14

Ke3dk and 29 Ke3punch slice events scaled to an expected signal box background of 0.36±0.21 Ke3dk events and 0.09±0.05 Ke3punch events, for a combined background of 0.45±0.21 for 1999.

The slice estimation method could be applied to the search data as well. Removing the `pi0_mass` cut in the 1997 data resulted in 38 events in the slice, including in the signal box itself. Scaling by the `pi0_mass` cut suppression and dividing by ten gave an expected box background of 0.24±0.09 events. Since the signal box had been opened for the 1997 data set, the box estimation method could be applied as well. 5 events appeared in the box without the `pi0_mass` cut. This scaled to 0.31±0.18 expected background events when all cuts were applied. For the '99 data set, the box remained closed and we could only apply the slice method. 53 data events appeared in the slice without the `pi0_mass` cut. Scaling by the `pi0_mass` cut suppression and dividing this time by 9, gave an expected background of 0.43±0.13 events in the box from Ke3s. These estimations are summarized in Table 6.5. The method with the least uncertainty was the slice method, and since the data knows the real answer, we used 0.43±0.13 as our best estimate of background in the box for '99.

To confirm that the background events in the "Ke3" area were really Ke3's, we compared '97 search data to combined MC in Appendix C.

Method of estimation (pi0_mass cut rem.)	97 Ke3 events expected in box	99 Ke3 events expected in box
Slice in K_mass using MC	0.27±0.13	0.45±0.21
Events in box using MC	0.24±0.18	<0.34±0.30
Slice in K_mass using blind data	0.24±0.10	0.43±0.13
Slice in K_mass using opened data	0.24±0.09	0.53±0.15
Events in box using opened data	0.31±0.18	0.83±0.33

Table 6.5 Different methods for estimating the expected background in the signal box and their results. The "Events in the box using MC" method for '99 is an upper limit because we were not able to generate enough statistics for a Ke3dk event to appear in the box. The number given would be the background estimate had one Ke3dk event passed all cuts and landed in the box. To call this a blind analysis and get the most accurate estimation, we decided to use the blind data calculation (highlighted in green).

Chapter 7
1999 Data Results

After all cuts are made, there are 13 events in the study plot for the 1999 search data set, excluding in the box. When the signal box was opened, we observed 3 events with an expected background of 0.47 events. The Poisson probability of 0.47 events fluctuating to 3 or more events is 1.2%. The data events are listed in Table 7.1 and their distribution can be seen in Figure 7.1.

Event label	`K_mass` (GeV/c^2)	`pt2` ($GeV^2/c^2 \times 10^{-4}$)
P01	0.426976	6.03584
P02	0.437691	2.99804
P03	0.444788	2.35842
P04	0.453020	1.88029
P05	0.459806	13.5133
P06	0.474750	2.56816
P07	0.487699	17.1297
P08	0.488986	15.5335
P09	0.491006	5.28121
P10	0.494097	11.2902
P11	0.506260	21.9497
P12	0.511323	11.0980
P13	0.523066	24.2535
B1	0.489864	1.1239
B2	0.505092	1.4443
B3	0.507146	1.5757

Table 7.1 A list of the '99 data events passing all the cuts, with their `K_mass` and `pt2` values. Events whose labels start with "P" are in the plot outside of the signal box. Events with "B" labels are inside the signal box. The events are labeled left-to-right in Figure 7.1.

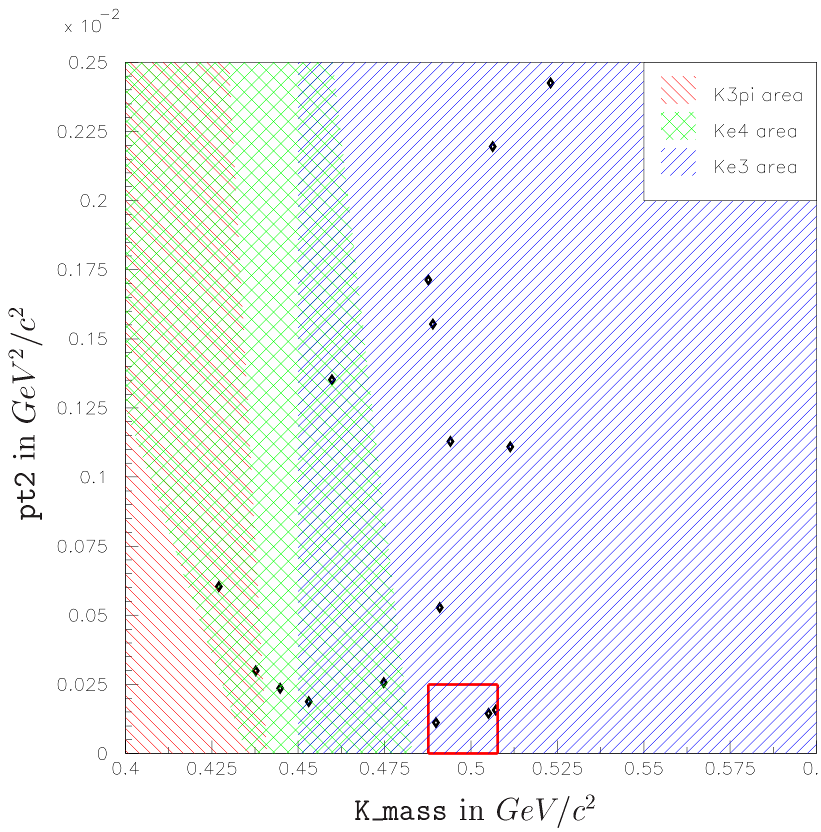

Figure 7.1 Signal search data from 1999. There are 2 events in the K3pi region, 4 events in the Ke4 region, and 7 events in the Ke3 region (excluding the box). The signal box contains 3 events.

Using the regions we defined as a guideline, events P01 and P02 are K3pi events. P03 is a Ke4 event. P04 through P06 are in the overlap of Ke4 and Ke3 area, but are probably Ke4 events. Events P07 and higher are all Ke3 events. A comparison between our expected background values and the search data is listed in Tables 7.2 and 7.3, according to the assignments just given.

Background in study plot (excluding box)	# of expected events	# of seen events
K3pi	5.4±3.3	2
Ke4	3.1±0.4	4
Ke3	9.5±4.2	7
Total in study plot	18.0±5.4	13

Table 7.2 Comparison between the number of background events expected in the study plot, and what is seen in the 1999 search data according to the assignments given in the text.

Background in signal box	# of expected events	# of seen events
K3pi	negligible	-
Ke4	< 0.09	-
Ke3	0.43±0.13	-
Total background	0.43±0.13	3

Table 7.3 Comparison between the number of background events expected in the signal box, and what is seen in the 1999 search data.

If one takes the number of '97 data events passing the '99-like cuts, and then scales up by the higher '99 flux, one would predict 9 events in the '99 study plot. We saw 13 events, more than expected just according to the flux ratio. The reason was that there seemed to be more accidental activity, specifically photons, in the '99 data than in the '97 data. Figure 7.2 shows search data from '97 and '99 with the pi0_mass cut removed. If one takes the ratio of the fluxes, one gets $2.619 \times 10^{11} / 2.915 \times 10^{11} = 0.898$. The ratio of number of events is somewhat lower however, at $222/300 = 0.74$.

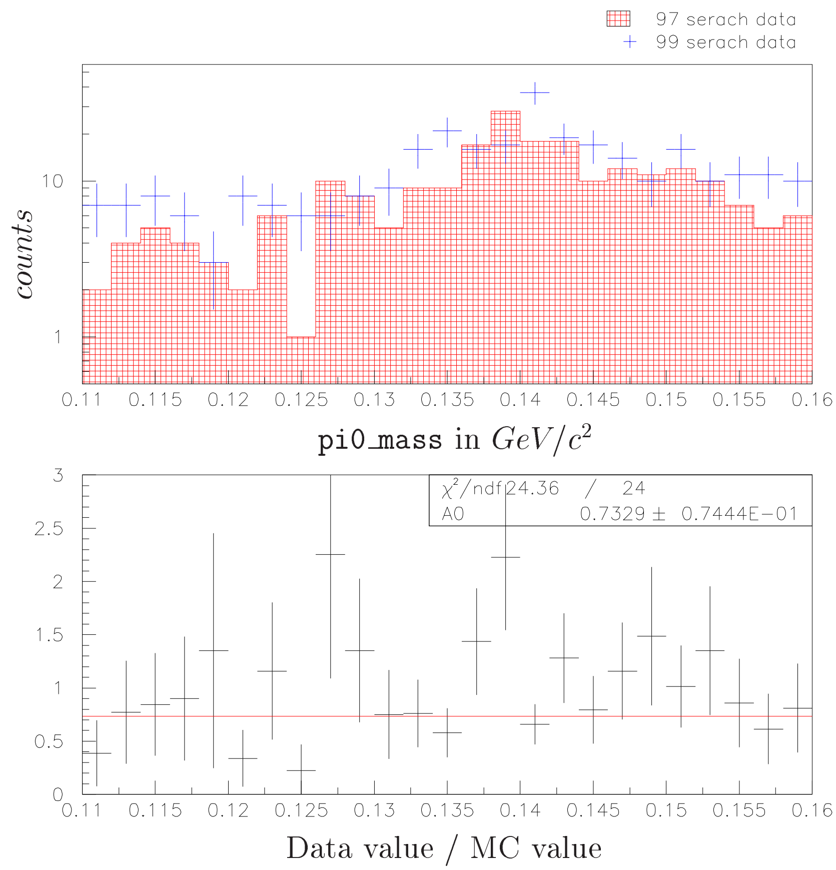

Figure 7.2 Search data with all but the pi0_mass cut. From the flux ratio (2.619/2.915) we would expect the fit constant of the bottom plot to be 0.898. The extra accidental photon activity pushed the line down to around 0.74.

The distribution of the data events within the signal box is not consistent with the distribution of the MC signal events. If we project the search data in the box onto each axis overlaid with the same projection of the signal MC, we can see that the data events fall on the tails of the MC distribution (see Figures 7.3 and 7.4). We can integrate the number of MC events farther from the K_mass mean than each data event. This gives the probability that a true signal event would fall as far or farther from the K_mass mean. Results for this calculation are given in Table 7.4. A similar calculation can be done for the pt2 projection of Figure 7.4, with results also listed in Table 7.4.

Event label (from left to right in K_mass)	% of MC signal events farther from $\mu_{\texttt{K\_mass}}$	% of MC signal events with greater pt2
B1	0.71±0.04%	3.99±0.10%
B2	1.45±0.06%	2.09±0.07%
B3	0.14±0.02%	1.69±0.07%

Table 7.4 If a given data event were signal, the probability it would fall where it did or farther away from the nominal kaon mass in K_mass or from zero in pt2.

If these two parameters were independent, then these individual probabilities could be simply combined. However, K_mass and pt2 could be correlated. To take into account the possibility of correlations, one could combine the two parameters before comparison. The simplest way was just to multiply their absolute values together after adjusting the K_mass to be centered around zero. Figure 7.5 shows the distributions of the search data and signal MC for $|\ \mu_{\texttt{K\_mass}}$-K_mass$|$*pt2. The percentage of MC

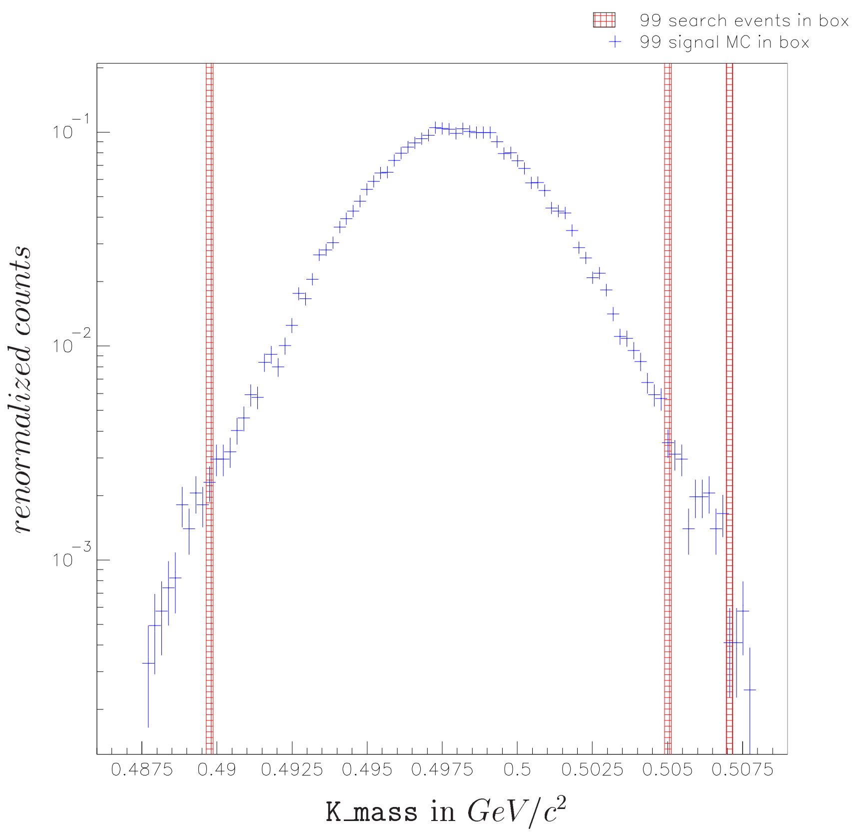

Figure 7.3 1999 data events in the signal box (red) and MC signal events (blue) projected onto the `K_mass` axis. The data events fall on the tails of the signal MC distribution.

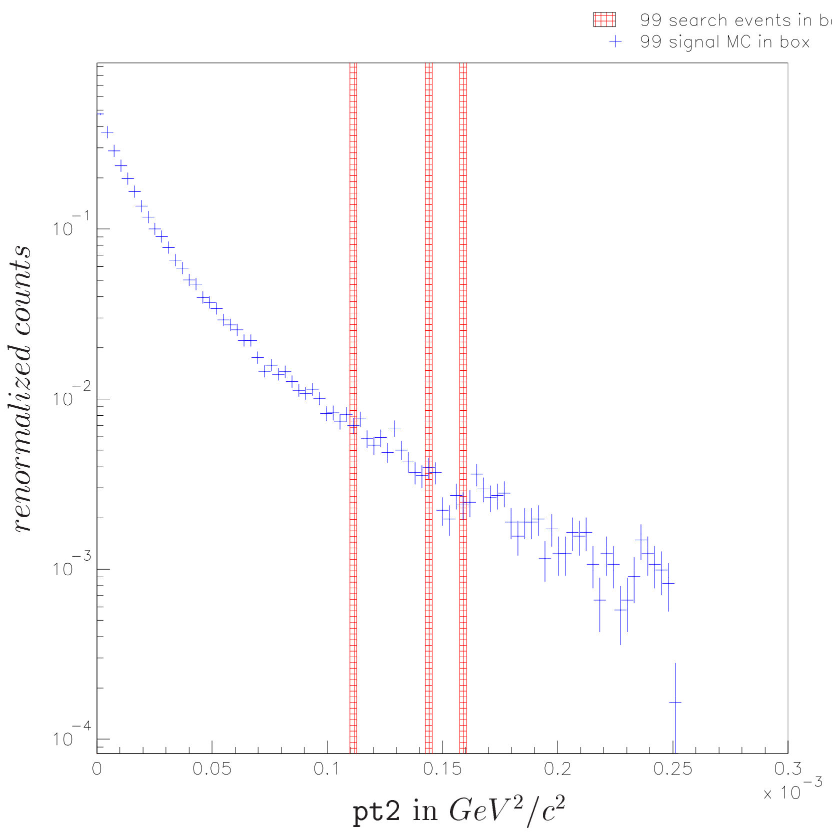

Figure 7.4 1999 data events in the signal box (red) and MC signal events (blue) projected onto the `pt2` axis.

events with higher values than the three data points is shown in Table 7.5.

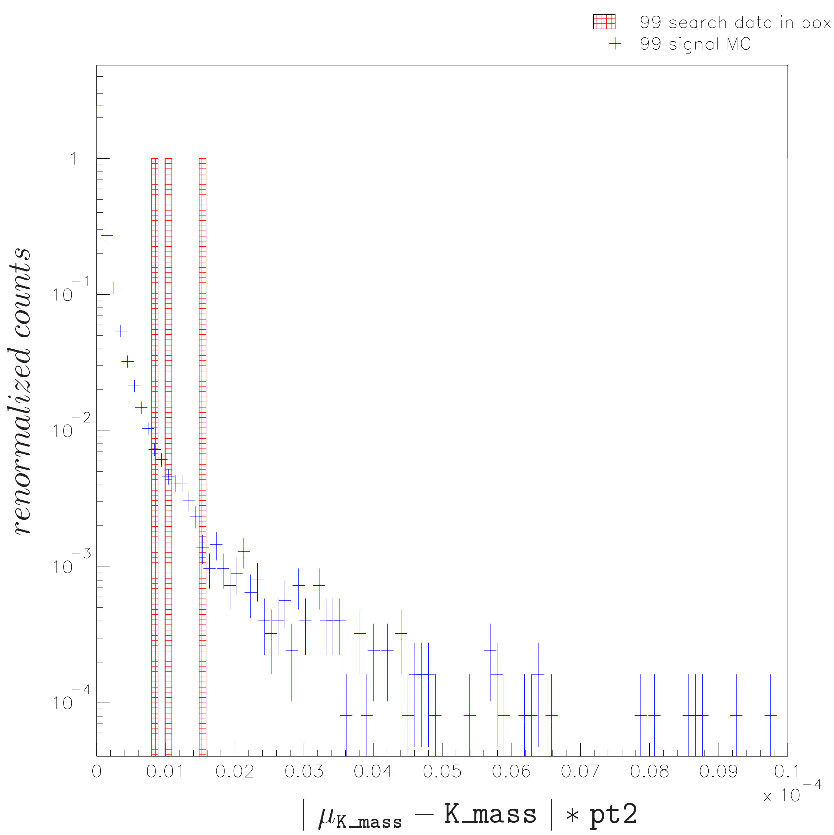

Figure 7.5 1999 data events in the signal box (red) and MC signal events (blue) passing all selection cuts for the combined parameter $\mid \mu_{\mathtt{K\_mass}}$-$\mathtt{K\_mass}|$*$\mathtt{pt2}$. The data events fall on the tail of the signal MC distribution. The MC histogram has been renormalized to fit on the plot.

Fewer than 1.8% in $\mathtt{K\_mass}$, 5.8% in $\mathtt{pt2}$, and 0.54% in $\mid \mu_{\mathtt{K\_mass}}$ - $\mathtt{K\_mass}|$*$\mathtt{pt2}$ of true signal events would fall farther than the three data events we observed. As was the case for the '97 data, the observed events are distributed in a way that is inconsistent with the signal MC. We believe them to be background, but include all

Event label (from left to right in K_mass)	% of MC signal events higher in $\mid \mu_{\texttt{K\_mass}} - \texttt{K\_mass} \mid$*pt2
B1	0.54±0.04%
B2	0.37±0.03%
B3	0.07±0.01%

Table 7.5 If a given data event were signal, the probability it would fall where it did or farther away from the nominal kaon mass in the combined parameter $\mid \mu_{\texttt{K\_mass}}$ - K_mass$\mid$*pt2.

our information in the calculation of a conservative branching ratio limit, as was done in '97.

We observed three events in our signal box and had a background estimation of 0.43±0.13 events. Using this information, we can calculate a conservative 90% confidence level limit on the branching ratio of this decay mode.

Using Feldman and Cousins technique [41], assuming a background that is Gaussian with mean 0.43 and sigma 0.13, and 3 observed events, the 90% confidence level limit sensitivity is 7.00 events. The branching ratio limit is then:

$$BR(K_L \to \pi^0\mu^\pm e^\mp) < \frac{7.00}{(Flux\ of\ K_L) * (Signal\ acceptance)} \tag{7.1}$$

Using our K_L flux of 3.14x10^{11} and Kpi0me acceptance of 3.576%, the branching ratio limit for the '99 data alone is:

$$BR(K_L \to \pi^0\mu^\pm e^\mp(99)) < 6.24 \times 10^{-10} \tag{7.2}$$

We can also calculate a branching ratio limit using Poisson probability density functions (PDFs). We used a piece of code written by one of our collaborators (see

Appendix E) to calculate a 90% CL limit. That branching ratio limit for the 1999 data set is:

$$BR(K_L \to \pi^0 \mu^\pm e^\mp(99)) < 5.57 \times 10^{-10} \tag{7.3}$$

Finally, the PDF technique can be used to combine the results of the two data sets. The code allows for the possibility that the two results could have been from completely different experiments. The combined branching ratio limit for the 1997 plus 1999 data sets is:

$$BR(K_L \to \pi^0 \mu^\pm e^\mp(97 + 99)) < 3.37 \times 10^{-10} \tag{7.4}$$

Chapter 8
Conclusions

We analyzed two KTeV data sets and saw no definitive evidence for the decay $K_L \to \pi^0 \mu^\pm e^\mp$. We set 90% confidence level upper limits for each data set and for the combined set. The best limits were given by the PDF method. These limits were calculated using blind analysis methods and improve on the previous PDG limit by a factor of 20.

Data set	Branching ratio limit (90% confidence level)
1997 only	3.58×10^{-10}
1999 only	5.57×10^{-10}
1997 + 1999	3.37×10^{-10}

Table 8.1 Kpi0me branching ratio limits at the 90% confidence level using the PDF method.

It was hoped the extra data collected in 1999 would greatly improve on the '97 limit, especially with the higher flux. The increased accidental activity really hurt us. We probably could have gotten a much better result if we had had the same running conditions in '99 as in '97. The change could have been due to any number of differences, including poorer beam quality due to Main Injector commissioning, a lower spectrometer magnet setting, lower thresholds on the drift chamber discriminators. Improved photon vetoes, especially an improved BA, would have been helpful in reducing the backgrounds.

To substantially reduce the accidental activity, one would have to reduce the amount of material in the fiducial area. The vacuum window was by far the largest single source of secondary interactions. If the spectrometer, and possibly even the calorimeter, could be placed upstream of the vacuum window (i.e. within the vacuum itself), background from material in the beam would be greatly reduced. The KaMI (Kaons at the Main Injector) experiment that was designed as KTeV's successor had plans to incorporate the drift chambers into the vacuum to reduce this source of background. Unfortunately that project did not have sufficient funding to become reality.

On the analysis end, the greatest improvement would have come from using a differently shaped acceptance area. A triangular acceptance "box" would better fit the distribution of the signal around the kaon mass than did our rectangular box. If we had used a triangle in 1999, then the three events we saw might not have been within the acceptance region and the limit would have been much lower.

Appendix A
Additional Cuts Studied, but Not Used

The abundance of events appearing in the 1999 study plot were a cause of concern for us. Doing a simple ratio of fluxes, one would expect only 8 or 9 events in the 1999 data study plot. We tried several different methods to get rid of the additional background that were unsuccessful. It was finally decided that there was simply more accidental activity in the 1999 and we would just have to deal with seeing more events. However, the unsuccessful attempts are described here for edification.

A.1 Neutral vertexing

In the final analysis it was decided to use the charged vertex to reconstruct the mass of the two photons and then compare the result to the neutral pion mass. We had also investigated the method of assuming the neutral pion mass and calculating a neutral vertex location from the photon information. We applied all of our other cuts and then looked at the difference between the charged vertex Z-location and the neutral vertex Z-location (see Figure A.1). A cut that would remove any additional data background would also cut into the signal acceptance, so no cut is made. Since this cut is somewhat redundant to the pi0_mass cut, it is not surprising that nothing could be gained from it.

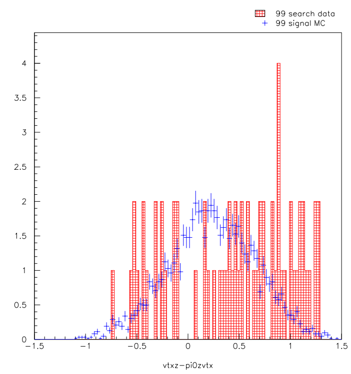

Figure A.1 Histograms of the difference between the charged vertex Z and the neutral vertex Z. Shown here is '99 signal MC (blue crosses) and '99 search data (red histogram) with all signal selection cuts except the TRD CL cut. Any cut placed on this variable would have to cut into the signal acceptance to remove any more data background events, so none was made. The signal MC distribution has been renormalized to fit on the plot.

A.2 Sign of neutral vertex X coordinate

For the Ke3 background, the neutral pion had to come from accidental activity. Since we already had the neutral vertex information (taking the two photon center of energy for X and Y), we investigated making a cut on the sign of the X-coordinate. The kaon beam was collimated into two parts, one with a positive X location and the other with a negative X location. If the neutral pion really came from a neutral kaon decay, then the X-coordinate of the charged vertex and the X-coordinate of the neutral pion should have the same sign. If the photons were accidental, the neutral vertex has a roughly 50% chance of reconstructing to the wrong sign in X. Unfortunately, the signal acceptance loss is too great when such a cut was applied. For example in the 1999 event set, after all the other cuts are made, this cut would remove more signal (roughly 25%) than background (4 of 25 events, or 16%).

A.3 Comparison of photon energies

Studies of other decay modes by other collaborators pointed to investigating the ratio of the energies of the two photons, again to try to distinguish between accidental activity and a neutral pion decay. Theoretically, photons from neutral pion decays should tend to share the energy equally, making this parameter peak near one. If one photon is significantly more energetic than the other, it is less likely they both are from a neutral pion decay. In studying the energy ratio (smaller energy divided by larger energy, called "ggratio"), we did not see a distinction between neutral pion decay photons and accidental photons so this parameter was not cut on (see Figure A.2). We also looked at the asymmetry in the energies (|E1-E2|/E1+E2), which was unhelpful as well (see Figure A.3).

A.4 Elliptical pi0_mass cut

In the analysis of the $K_L \to \pi^0 e^+ e^-$ decay mode done by Leo Bellantoni, the final study plot was of the two-gamma mass versus the parent particle mass. There is a clear correlation between these two parameters for the $K_L \to \pi^0 e^+ e^-$ signal; if the neutral pion mass fluctuates higher, the parent particle mass will be higher as well. In retrospect, he thought that it might have been better to make his acceptance region more like an ellipse than a box. So, we considered applying such a cut to our neutral

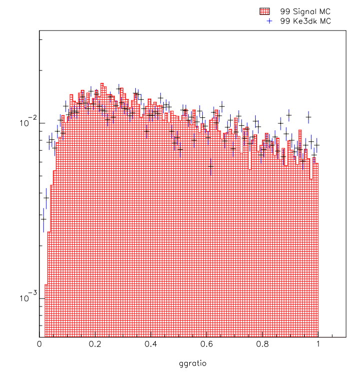

Figure A.2 Histograms of the ratio of energies (E1 / E2, E1 ¡ E2) of the two photons in an event. Shown here is background Ke3dk MC (red histogram) and signal MC (blue crosses) with no cuts applied (to get more statistics). The other background with accidental photons (Ke3punch) was similar to the Ke3dk distribution. There is almost no difference between the two distributions so no cut was made on this parameter. The Ke3dk distribution has been renormalized to the number of signal MC events.

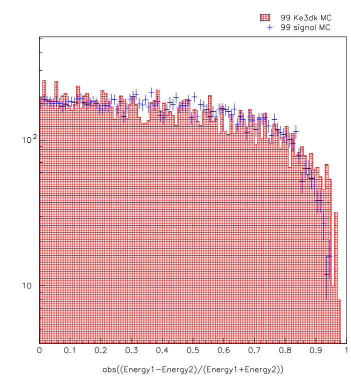

Figure A.3 Histograms of the asymmetry of the energies of the two photons in an event. Shown here is background Ke3dk MC (red histogram) and signal MC (blue crosses) with no cuts applied (to get more statistics). The other background with accidental photons (Ke3punch) was similar to the Ke3dk distribution. There is almost no difference between the two distributions so no cut was made on this parameter. The Ke3dk distribution has been renormalized.

pion mass. Since our final cut on pt2 is linear, it would be best to switch our study plot parameters to pi0_mass versus K_mass (see Figure A.4). We also then removed the pi0_mass cut from the selection cuts and added the pt2 restriction from the old signal box (pt2 $\leq 0.00025\frac{GeV^2}{c^2}$). When we made our new plots, the distributions of the background modes were not quite as nicely sectioned as they had been in our pt2 versus K_mass plot (see Figure A.5). An elliptical cut did not give a better signal acceptance per unit area (ellipse: 5.8% in $114MeV^2$, box: 4.3% in $80MeV^2$ for the 97 signal MC), and we were unsure how to characterize the distributions of the backgrounds and how they approached the final signal “box”. We decided to stay with our rectangular box.

A.5 Decay angle of photons

Yet another attempt was made to cut against accidental photons when a paper by H.B. Greenlee was brought to our attention [42]. Greenlee gives a formula for calculating a parameter he calls y_γ (and I will call “cos_pi0”). It is the cosine of the angle a photon in the neutral pion rest frame makes with respect to the direction of the neutral pion in the kaon rest frame. If the photons really are from a neutral pion decay at the charged vertex, then then this “cos_pi0” should be flat. Looking at Figure A.6, the signal MC (blue crosses) with a neutral pion decay are in fact flat, and the Ke3dk background MC (red histogram) with random accidental photons is skewed towards zero. Ke3punch MC look much like the Ke3dk MC in this variable.

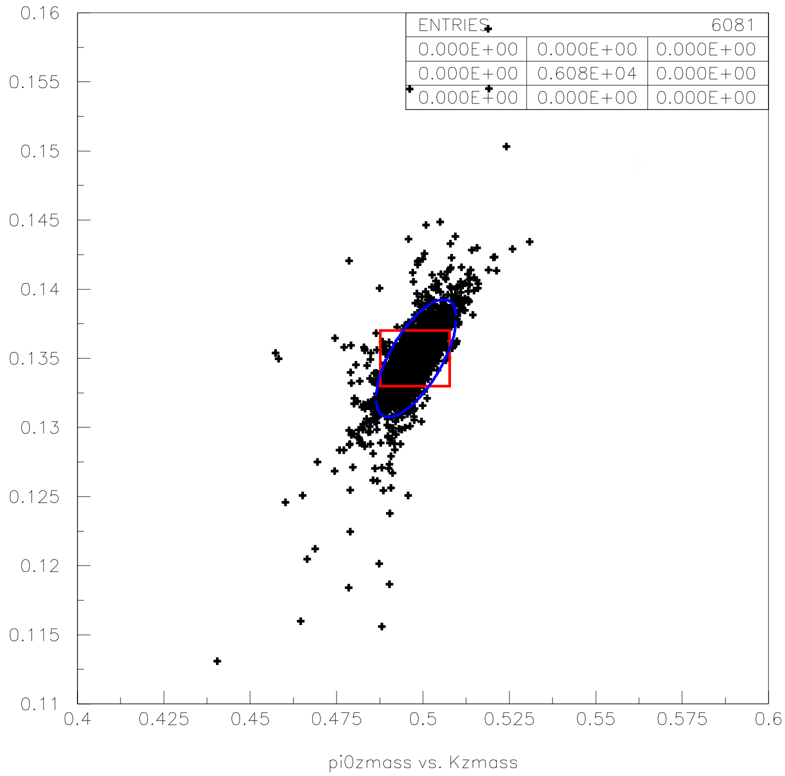

Figure A.4 A plot of pi0_mass versus K_mass for the 1997 signal MC. It was hoped that an elliptical cut (shown in blue) would provide better signal acceptance and background rejection. The difference in results between this new cut and our old cut (box shown in red) was not significant so we did not adopt this method of analysis.

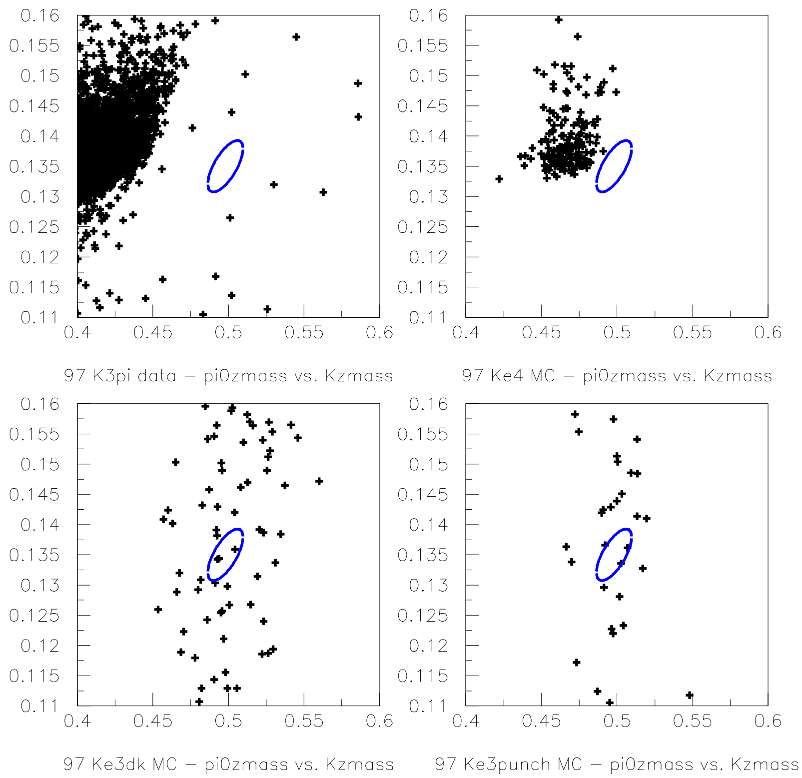

Figure A.5 Plots of pi0_mass versus K_mass for the 1997 backgrounds. The K3pi events are from trigger 2, reconstructed as Kpi0me and subjected only to the "detector" and "reconstruction" cuts. The Ke4's are MC and have passed all cuts. The Ke3's are MC and have passed all but the pt2 cut at $0.00025\frac{GeV^2}{c^2}$.

we did not see a clear enough distinction between them to even try to select a cut location.

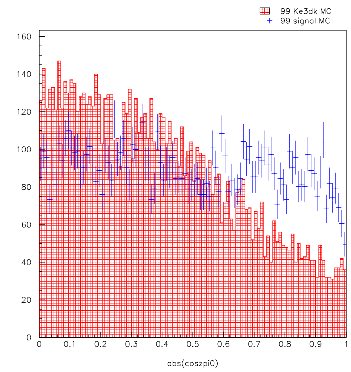

Figure A.6 Histograms of the cosine of the angle between one of the photons and the neutral pion direction in an event. The distribution is symmetric about zero, so the absolute value has been taken. Shown here is background Ke3dk MC (red histogram) and signal MC (blue crosses) with no cuts applied (to get more statistics). There was not enough of a separation between the two distributions to warrant a cut on this parameter. The signal distribution has been renormalized to the Ke3dk area.

A.6 BA source latch

The Back-Anti detector (BA) went through significant changes between the '97 and '99 runs (discussed in Section 3.4). Our goal is to combine the '97 and '99 data sets as seamlessly as possible, so we wanted to implement cuts that would treat '97 and '99 equivalently. To do so in the BA parameters proved difficult. At one point we investigated combining the energy information from the BA with the BA trigger source latch information. The BA source latch had four discrete values: a value of 0 means the latch was not triggered; a value of 1 means a low threshold was satisfied; a value of 3 means a high threshold was satisfied. A value of 2 does not make logical sense (the high threshold was achieved but not the low one), but there are some events occurring with that value. For 1997, plots of the source latch ("BA1src") versus the BA energy looked as expected; the higher energy deposits fell with the higher threshold bits (see Figure A.7). The 1999 data however did not have the expected relationship (see Figure A.8). It was decided we did not understand what was happening in the BA1src parameter in '99 and the idea of such a cut was dropped.

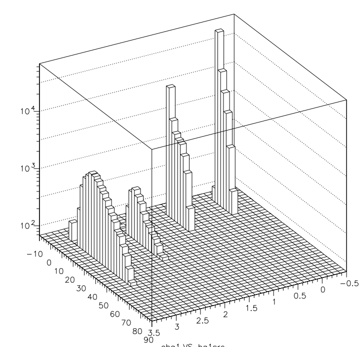

Figure A.7 Distribution of the energy deposited in BA1 ("BA1ene", also called "eba1") versus the value of the BA source latch ("BA1src") for 1997 normalization data. The BA source latch had four discrete values: 0 = no trigger; 1 = low threshold; 3 = high threshold. A value of 2 does not make logical sense (the high threshold was achieved but not the low one), but there are some events occurring with that value. As expected, there is more energy deposited when the BA1src value is higher.

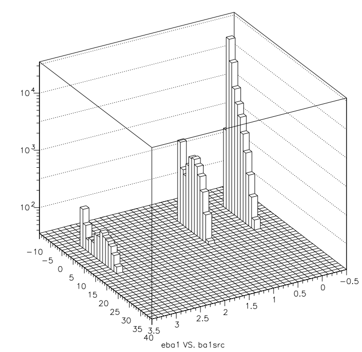

Figure A.8 Distribution of the energy deposited in BA1 ("BA1ene", also called "eba1") versus the value of the BA source latch ("BA1src") for 1999 normalization data. The BA source latch had four discrete values: 0 = no trigger; 1 = low threshold; 3 = high threshold. A value of 2 does not make logical sense, since it would mean the high threshold was passed without passing the low threshold. In this respect, the '99 distribution makes more sense than the '97 one because there are no events with a latch value of 2. Otherwise, this distribution does not follow the pattern expected (higher latch value corresponding to higher energies).

A.7 Number of extra soft clusters

The calorimeter is segmented, with the smallest crystals having a cross-sectional area of $2.5cm$x$2.5cm$. If electrons and/or photons hit the calorimeter near each other, their energies overlap and can result in incorrect reconstruction. Many hits overlapping further complicate the situation. We decided to see if cutting on the total number of soft clusters (we already had a cut on hard clusters) would help remove accidental events that were otherwise accepted because of extra overlapping energy. Unfortunately, a cut on this variable ended up removing more signal than background (see Figure A.9), so it was abandoned.

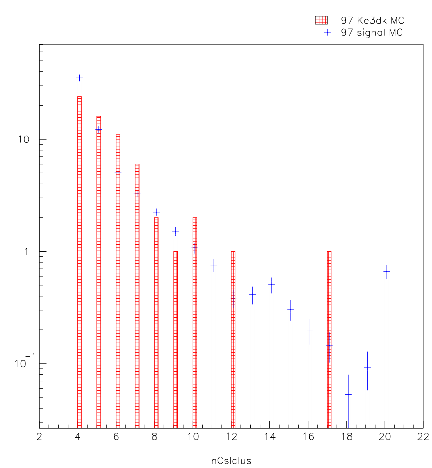

Figure A.9 Distribution of the number of soft clusters in 1997 Ke3dk background MC and signal MC, with the pi0_mass cut removed for improved statistics. A cut placed on this parameter to remove background events (at 10, for example) would remove too much signal to make it advantageous.

A.8 Soft cluster energy and timing

Since simply cutting on the number of extra soft clusters was unsuccessful, we attempted to combine the soft cluster energy information with the cluster timing. We were not as concerned with soft clusters that arrived out of time with the decay event, but wanted to cut on events with in-time soft clusters. When we examined the timing parameter of our hard cluster “electron” and hard cluster photons, the plots looked promising (see Figure A.10). However, as the cluster energy decreased, the disagreement between data and MC became significant (see Figures A.11 and A.12). This discrepancy might have been correctable, so we tried applying some cuts to the '97 search data. We set up the additional cuts to remove an event if the most energetic extra cluster energy was above $0.3GeV$, its timing chi-square parameter was less than 5, and it was not within $20cm$ of a charged track cluster (less than $20cm$ and it could be a satellite of the charged cluster). From Figure A.11 one can see the a cut at 5 is getting down to where the data / MC disagreement is worst. Adding the new cuts did not remove any '97 data events. Even if BA energy cut was removed, only one out of 22 data events was removed by the new cut combination. The same new cuts did remove a couple '99 data events but the corresponding loss in acceptance was discouraging. In the end it was decided that the timing parameter did not show enough potential to justify trying to calibrate it and use it.

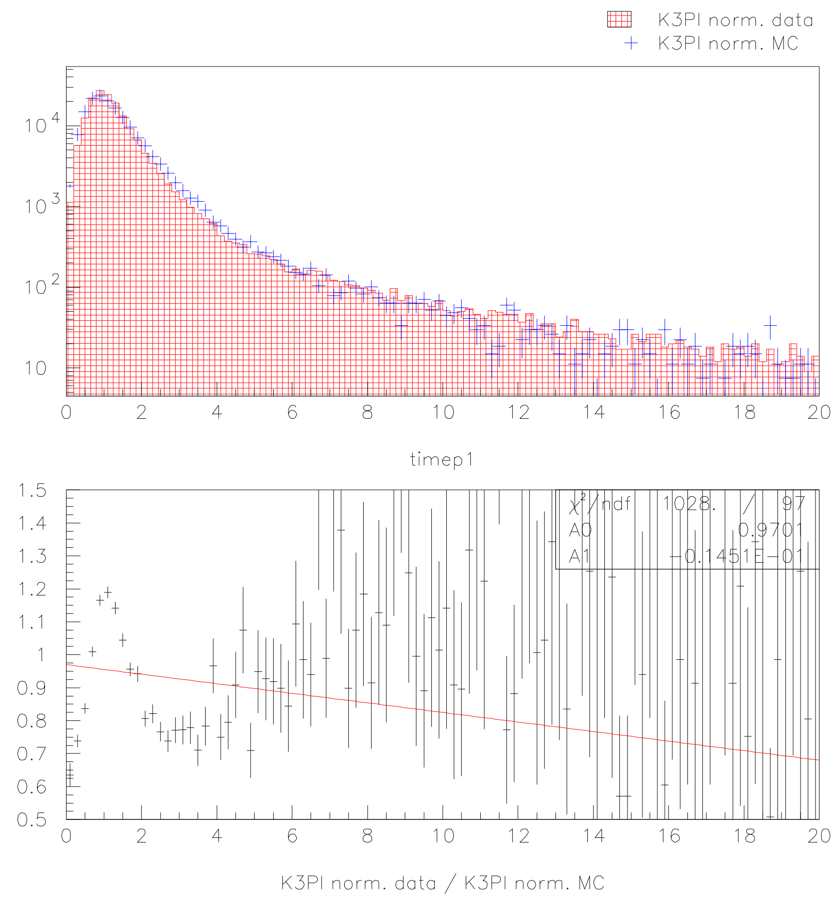

Figure A.10 Distribution of the timing chi-square parameter for the hardware cluster of one of the photons in a normalization event (1997 K3pi). The agreement between MC and data is reasonable, however it was not this cluster that we were interested in cutting on.

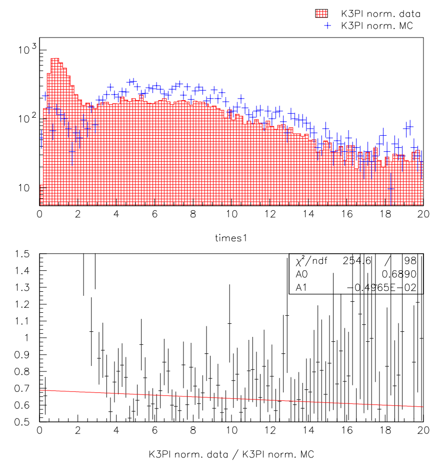

Figure A.11 Distribution of the timing chi-square parameter for a software cluster in 1997 normalization data and MC events. This cluster was selected by looking for the highest energy software cluster that was not associated with a track.

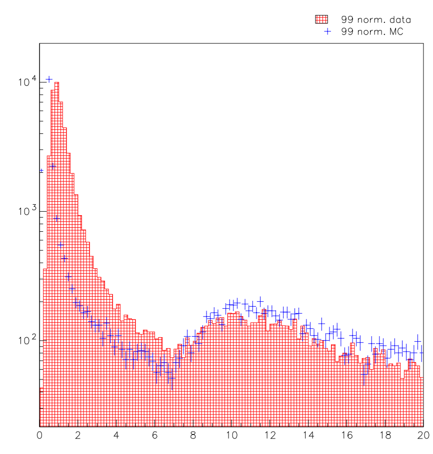

Figure A.12 Distribution of the timing chi-square parameter for a software cluster in 1999 normalization data and MC events. This cluster was selected by looking for the highest energy software cluster that was not associated with a track.

Appendix B
Acceptance and flux studies

Differences in acceptances between 1997 and 1999 brought concerns that just a few cuts could be responsible and perhaps should be removed. Comparisons were made between '97 and '99 for the signal MC and the normalization data as each cut was added cumulatively. Tables B.1 and B.2 show the results of this process. At the time this was done, we were still using the parameters `npairsup` and `npairsdn` for both the '97 and '99 analyses. In fact, the results of this study were what lead us to investigate switching to `xpairsup` and `xpairsdn` in the '99 analysis. Other than the `npairs` cuts, there did not seem to be any outstanding trouble makers, but rather a slightly greater loss in acceptance in each cut for '99.

We also investigated the acceptance as each cut was removed, leaving all the other cuts in place. This time the comparison was between data and MC for the normalization sets to see where flux might have been artificially lost. Tables B.3 and B.4 show the acceptance for each removed cut. For this study the baseline (100% acceptance) was defined as having passed the analysis code and the `vtxx` and `vtxy` cuts. The ratio of acceptances was taken and compared to the ratio for having all cuts applied to indicate where the largest differences were between '97 and '99. Ranking the cuts by this "Δ ratio", cuts were then applied cumulatively starting with the

Cut added	Signal MC			
	% '97 accept	Δ %	% '99 accept	Δ %
none	100.00%		100.00%	
analysis+vtxx,y	10.36±0.10%	89.64±0.10%	10.58±0.10%	89.42±0.10%
phvbar1	10.36±0.10%	0.00±0.14%	10.58±0.10%	0.00±0.14%
RCmaxene	10.31±0.10%	0.05±0.14%	10.52±0.10%	0.06±0.14%
SAmaxene	10.30±0.10%	0.01±0.14%	10.51±0.10%	0.01±0.14%
EBA1	8.92±0.09%	1.38±0.13%	8.96±0.09%	1.55±0.13%
npln_e	8.88±0.09%	0.04±0.13%	8.93±0.09%	0.03±0.13%
npairsup	7.96±0.09%	0.92±0.13%	7.76±0.08%	1.18±0.12%
npairsdn	7.79±0.08%	0.17±0.12%	7.57±0.08%	0.19±0.12%
nxtrks	7.41±0.08%	0.38±0.12%	7.14±0.08%	0.44±0.12%
nytrks	7.05±0.08%	0.35±0.12%	6.75±0.08%	0.39±0.11%
nhCsIcls	7.05±0.08%	0.00±0.12%	6.75±0.08%	0.00±0.11%
eoffx	7.01±0.08%	0.04±0.12%	6.69±0.08%	0.06±0.11%
eoffy	6.97±0.08%	0.04±0.12%	6.65±0.08%	0.04±0.11%
muoffx	6.81±0.08%	0.16±0.11%	6.48±0.08%	0.17±0.11%
muoffy	6.70±0.08%	0.12±0.11%	6.32±0.08%	0.15±0.11%
nmumats	6.69±0.08%	0.00±0.11%	6.32±0.08%	0.01±0.11%
vtxz	6.29±0.08%	0.40±0.11%	5.90±0.07%	0.41±0.11%
vtx_chi	6.13±0.08%	0.16±0.11%	5.75±0.07%	0.15±0.10%
pt2	6.13±0.08%	0.00±0.11%	5.73±0.07%	0.03±0.10%
K_energy	6.11±0.08%	0.01±0.11%	5.72±0.07%	0.01±0.10%
pi0_mass	5.19±0.07%	0.93±0.10%	4.67±0.07%	1.05±0.10%
eop_elec	5.04±0.07%	0.15±0.10%	4.49±0.07%	0.18±0.09%
fuse3e	4.95±0.07%	0.09±0.10%	4.43±0.07%	0.06±0.09%
fuse3p1	4.85±0.07%	0.10±0.10%	4.35±0.06%	0.07±0.09%
fuse3p2	4.76±0.07%	0.09±0.10%	4.28±0.06%	0.07±0.09%
pmu	4.41±0.07%	0.35±0.09%	3.88±0.06%	0.39±0.09%
chrgmass	4.41±0.07%	0.00±0.09%	3.88±0.06%	0.01±0.09%
pp0_kine	4.09±0.07%	0.31±0.09%	3.60±0.06%	0.28±0.08%
ke4_kine	4.07±0.07%	0.02±0.09%	3.58±0.06%	0.02±0.08%
sgp0kine	4.06±0.07%	0.01±0.09%	3.56±0.06%	0.02±0.08%
signal box	4.02±0.06%	0.05±0.09%	3.49±0.06%	0.07±0.08%

Table B.1 Acceptance of MC signal events for 1997 and 1999 as each cut was added cumulatively. There was not an obvious culprit for the difference in the final acceptance.

Cut added	Normalization data			
	% ’97 accept	Δ %	% ’99 accept	Δ %
none (split tapes)	100.000%		100.000%	
analysis +vtxx,y	4.222±0.004%	95.778±0.004%	4.203±0.006%	95.797±0.006%
phvbar1	3.461±0.004%	0.761±0.006%	3.842±0.006%	0.361±0.008%
RCmaxene	3.423±0.004%	0.038±0.005%	3.798±0.006%	0.044±0.008%
SAmaxene	3.419±0.004%	0.004±0.005%	3.771±0.006%	0.027±0.008%
EBA1	2.748±0.003%	0.671±0.005%	2.881±0.005%	0.890±0.008%
npln_e	2.631±0.003%	0.117±0.005%	2.806±0.005%	0.075±0.007%
npairsup	2.106±0.003%	0.525±0.005%	2.135±0.004%	0.671±0.007%
npairsdn	2.023±0.003%	0.083±0.004%	2.038±0.004%	0.097±0.006%
nxtrks	1.858±0.003%	0.166±0.004%	1.847±0.004%	0.191±0.006%
nytrks	1.666±0.003%	0.192±0.004%	1.687±0.004%	0.160±0.006%
nhCsIcls	1.125±0.002%	0.541±0.003%	1.263±0.003%	0.423±0.005%
eoffx	1.115±0.002%	0.010±0.003%	1.251±0.003%	0.013±0.005%
eoffy	1.108±0.002%	0.007±0.003%	1.242±0.003%	0.009±0.005%
muoffx	1.043±0.002%	0.064±0.003%	1.158±0.003%	0.084±0.005%
muoffy	1.007±0.002%	0.036±0.003%	1.114±0.003%	0.044±0.004%
vtxz	0.934±0.002%	0.074±0.003%	1.034±0.003%	0.080±0.004%
vtx_chi	0.880±0.002%	0.054±0.003%	0.964±0.003%	0.070±0.004%
pt2	0.867±0.002%	0.012±0.003%	0.943±0.003%	0.021±0.004%
K_energy	0.866±0.002%	0.001±0.003%	0.942±0.003%	0.001±0.004%
pi0_mass	0.620±0.002%	0.246±0.003%	0.660±0.002%	0.282±0.004%
fuse3p1	0.610±0.002%	0.010±0.002%	0.647±0.002%	0.013±0.003%
fuse3p2	0.600±0.002%	0.010±0.002%	0.635±0.002%	0.012±0.003%
chrgmass	0.598±0.002%	0.001±0.002%	0.634±0.002%	0.002±0.003%
signal box	0.584±0.002%	0.015±0.002%	0.609±0.002%	0.024±0.003%

Table B.2 Acceptance of data normalization events for 1997 and 1999 as each cut was added cumulatively.

smallest Δ ratios. The biggest culprits in the 1997 data set seemed to be the `pi0_mass` and `npairsup` cuts (see Table B.5). For 1999, it was the `EBA1` and `npairsup` cuts (see Table B.6). This discrepancy between data and MC was also in evidence visually in the discontinuities of the ratio plot of 1999 normalization data vs. MC for `npairsup` (Figure 3.12). We were cutting out more data than MC, making the acceptance too large and the flux too small. Although the MC acceptance was too large, the same effect should have been present in the signal MC as well. We did a brief study of several other two-track decay modes, namely Ke3's, Kmu3's, and $K_L \to 3\pi^0$ Dalitz, in addition to the normalization mode and they all showed a similar data vs. MC mismatch. We concluded therefore that our signal mode would have behaved the same way and the effect should have canceled when the ratio of signal acceptance to normalization acceptance was taken in calculating the branching ratio limit.

If the cuts for '97 were loosened to `npairsup` at 10 and `pi0_mass` from $130 MeV$ to $140 MeV$, the acceptance was 5.93% and the flux became 2.68×10^{11}, comparable to other KTeV analyses flux numbers. For '99, if `npairsup` was again loosened to 10, and the `EBA1` cut removed entirely, the acceptance was 5.33% and the flux became 3.212×10^{11}. If we then factored back in the bad spill losses (6.34%), the comparative flux was 3.429×10^{11}, again in much better agreement with other KTeV analyses flux numbers.

Cut removed	1997 Normalization % MC accept (±0.1%)	% Data accept (±0.04%)	% MC / % Data (±0.001)	Δ ratio from no cuts removed
none	25.5%	13.83%	1.843	0.
phvbar1	25.5%	13.89%	1.838	0.005
RCmaxene	25.6%	13.88%	1.843	-0.0005
SAmaxene	25.5%	13.84%	1.843	0.0001
EBA1	29.4%	16.21%	1.812	0.03
npln_e	26.0%	14.40%	1.805	0.04
npairsup	27.6%	16.56%	1.668	0.17
npairsdn	25.8%	14.11%	1.829	0.01
nxtrks	26.7%	14.69%	1.818	0.02
nytrks	26.7%	14.88%	1.794	0.05
nhCsIcls	35.0%	19.31%	1.815	0.03
eoffx	25.6%	13.88%	1.841	0.002
eoffy	25.6%	13.88%	1.841	0.002
muoffx	26.3%	14.30%	1.841	0.002
muoffy	26.1%	14.22%	1.839	0.004
vtxz	27.3%	14.81%	1.841	0.002
vtx_chi	26.1%	14.12%	1.848	-0.005
K_energy	25.5%	13.84%	1.842	0.0004
pi0_mass	31.4%	18.14%	1.732	0.11
fuse3p1	26.1%	14.06%	1.857	-0.01
fuse3p2	26.1%	14.06%	1.856	-0.01
chrgmass	25.5%	13.83%	1.843	-0.000005
signal box	26.0%	14.17%	1.831	0.01

Table B.3 Acceptance of the normalization cuts as individual cuts were removed, using normalization MC and trigger 2 data for 1997. While the acceptances were expected to be somewhat different (since there were more than just K3pi events in the data set, but only K3pi's in the MC set), the ratio of acceptances helped to point to where flux was being lost because of disagreement between data and MC. The difference of a cut's ratio from the "none" ratio (no cuts removed) is shown in the last column. The larger this "δ ratio", the more of an effect the cut had in artificially changing the flux. For the purpose of this table, "100%" acceptance was defined as having passed the analysis code and `vtxx` and `vtxy` cuts.

Cut removed	1999 Normalization % MC accept (±0.1%)	% Data accept (±0.05%)	% MC / % Data (±0.001)	Δ ratio from no cuts removed
none	23.5%	14.50%	1.620	0.
phvbar1	23.6%	14.67%	1.608	0.01
RCmaxene	23.6%	14.57%	1.619	0.0009
SAmaxene	23.5%	14.51%	1.620	-0.0004
EBA1	27.5%	18.45%	1.488	0.13
npln_e	23.9%	14.78%	1.615	0.005
npairsup	26.1%	18.08%	1.445	0.18
npairsdn	23.8%	14.86%	1.604	0.02
nxtrks	24.8%	15.59%	1.593	0.03
nytrks	24.7%	15.54%	1.592	0.03
nhCsIcls	31.5%	19.21%	1.622	-0.002
eoffx	23.6%	14.56%	1.618	0.002
eoffy	23.6%	14.56%	1.617	0.003
muoffx	24.3%	15.03%	1.619	0.001
muoffy	24.2%	14.96%	1.618	0.002
vtxz	25.1%	15.49%	1.622	-0.002
vtx_chi	24.1%	14.73%	1.633	-0.01
K_energy	23.5%	14.09%	1.668	-0.04
pi0_mass	30.2%	19.32%	1.561	0.06
fuse3p1	24.0%	14.77%	1.628	-0.008
fuse3p2	24.0%	14.78%	1.626	-0.006
chrgmass	23.5%	14.51%	1.621	-0.0005
signal box	24.0%	15.08%	1.594	0.03

Table B.4 The same calculations as Table B.3, but for 1999 data. Note that this study was done when we were still using the `npairs` parameters, before switching to the `xpairs` parameters.

Cuts added	97 norm. MC accept	97 flux ($\pm$ stat. only)	Δ flux
analysis code vtxx,y phvbar1 RCmaxene SAmaxene eoffx,y muoffx,y vtxz vtx_chi K_energy chrgmass	13.40±0.03%	$3.16 \pm 0.1 \times 10^{11}$	
signal box	12.72±0.03%	$3.02 \pm 0.1 \times 10^{11}$	4.3% drop
fuse3p2	12.30±0.03%	$3.03 \pm 0.01 \times 10^{11}$	0.4% rise
fuse3p1	11.89±0.03%	$3.05 \pm 0.01 \times 10^{11}$	0.5% rise
npairsdn	11.57±0.03%	$3.02 \pm 0.01 \times 10^{11}$	1% drop
nxtrks	10.72±0.03%	$2.97 \pm 0.01 \times 10^{11}$	1.7% drop
nCsIcls	7.64±0.03%	$2.89 \pm 0.01 \times 10^{11}$	2.4% drop
BA1ene	6.61±0.03%	$2.85 \pm 0.01 \times 10^{11}$	1.5% drop
npln_e	6.48±0.02%	$2.79 \pm 0.01 \times 10^{11}$	2.1% drop
nytrks	6.11±0.02%	$2.71 \pm 0.01 \times 10^{11}$	2.8% drop
pi0_mass	4.95±0.02%	$2.55 \pm 0.01 \times 10^{11}$	5.9% drop
npairsup	4.56±0.02%	$2.31 \pm 0.01 \times 10^{11}$	9.5% drop

Table B.5 Change in the calculated flux as cuts were applied cumulatively to normalization MC and trigger 2 data for 1997. The flux was expected to be too high for the first set of cuts, since there were more than just K3pi events in the trigger 2 data and the first group of cuts would not have removed all the background. The fact that the flux mainly continued to drop as cuts were added reflected the disagreement between MC and data acceptance. The largest drops came from the `pi0_mass` and `npairsup` cuts, so these cuts were loosened to get a better approximation of the real flux.

Cuts added	99 norm. MC accept	99 flux (±stat. only)	Δ flux
analysis code vtxx,y RCmaxene SAmaxene npln_e nhCsIcls eoffx,y muoffx,y vtxz fuse3p1,2 chrgmass	9.24±0.03%	$3.73 \pm 0.01 \times 10^{11}$	
phv1bar	9.17±0.03%	$3.70 \pm 0.01 \times 10^{11}$	0.7% drop
vtx_chi	8.77±0.03%	$3.67 \pm 0.01 \times 10^{11}$	1% drop
npairsdn	8.49±0.03%	$3.62 \pm 0.01 \times 10^{11}$	1.4% drop
signal box	8.09±0.03%	$3.52 \pm 0.01 \times 10^{11}$	2.8% drop
nxtrks	7.38±0.03%	$3.44 \pm 0.01 \times 10^{11}$	2.3% drop
nytrks	6.93±0.03%	$3.36 \pm 0.01 \times 10^{11}$	2.1% drop
K_energy	6.92±0.03%	$3.36 \pm 0.01 \times 10^{11}$	0% drop
pi0_mass	5.38±0.02%	$3.24 \pm 0.01 \times 10^{11}$	3.7% drop
BA1ene	4.59±0.02%	$2.98 \pm 0.01 \times 10^{11}$	8% drop
npairsup	4.13±0.02%	$2.66 \pm 0.01 \times 10^{11}$	10.8% drop

Table B.6 Change in the calculated flux as cuts were applied cumulatively to normalization MC and trigger 2 data for 1999. The flux was expected to be too high for the first set of cuts, since there were more than just K3pi events in the trigger 2 data and the first group of cuts would not have removed all the background. The fact that the flux continued to drop as cuts were added reflected the disagreement between MC and data acceptance. The largest drops came from the `BA1ene` and `npairsup` cuts, so these cuts were loosened to get a better approximation of the real flux. Note that this study was done when we were still using the `npairs` parameters, before switching to the `xpairs` parameters.

Appendix C
Ke3 MC Versus Data

Since there was not much overlap of the "Ke3" area with the other background areas, we were able to confirm that the background events in the Ke3 area were really Ke3's by comparing data to combined MC. We have included search data / MC comparison plots in Figures C.1 through C.7. The events in these plots were reconstructed as Kpi0me's, and the appropriate ratio of punch-through to decay MC was applied. The standard signal selection cuts were applied with the following modifications: the `prob_e`, `pt2`, `pi0_mass`, and `pp0_kine` cuts were removed to improve statistics; a requirement that `K_mass` be greater than $0.48767 GeV/c^2$ was added to restrict the events to the Ke3-only area of the study plot. The ratio plots are not shown because statistics were so low they were not informative. The `vtx_chi` plot demonstrates that we had a reasonable ratio of Ke3dk's and Ke3punch's combined. Figure C.2 (`eop_elec`), shows that there was in fact a real electron in the data, as was expected from Ke3's. The `chrgmass` distribution shows that the data followed the shape of a Ke3 charged system with an electron plus a pion. The `pi0_mass` distribution of Figure C.4 was flat for both data and MC, as expected since the photons were from accidental activity and not really from a neutral pion decay. The parent particle mass (`K_mass`) and `pp0_kine` reconstructions also agreed between data and MC

(Figures C.5 and C.6). Finally, the distribution of `pt2` in the region which included the study plot was indeed flat, so our assumption for the slice method of background approximation was supported. All these distribution agreed well enough that we felt the background in the study plot at higher `K_mass` was accounted for by Ke3's with accidental photons.

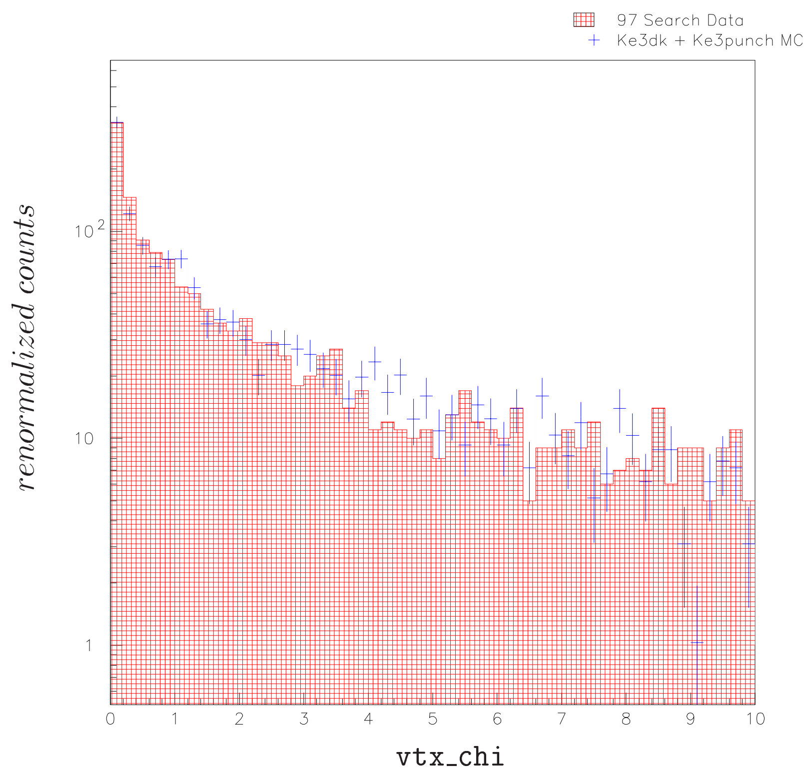

Figure C.1 A comparison of `vtx_chi` between 1997 search data in the Ke3 region and Ke3 MC. These curves agreed well enough to support our assertion that these background events were due to Ke3's plus accidentals.

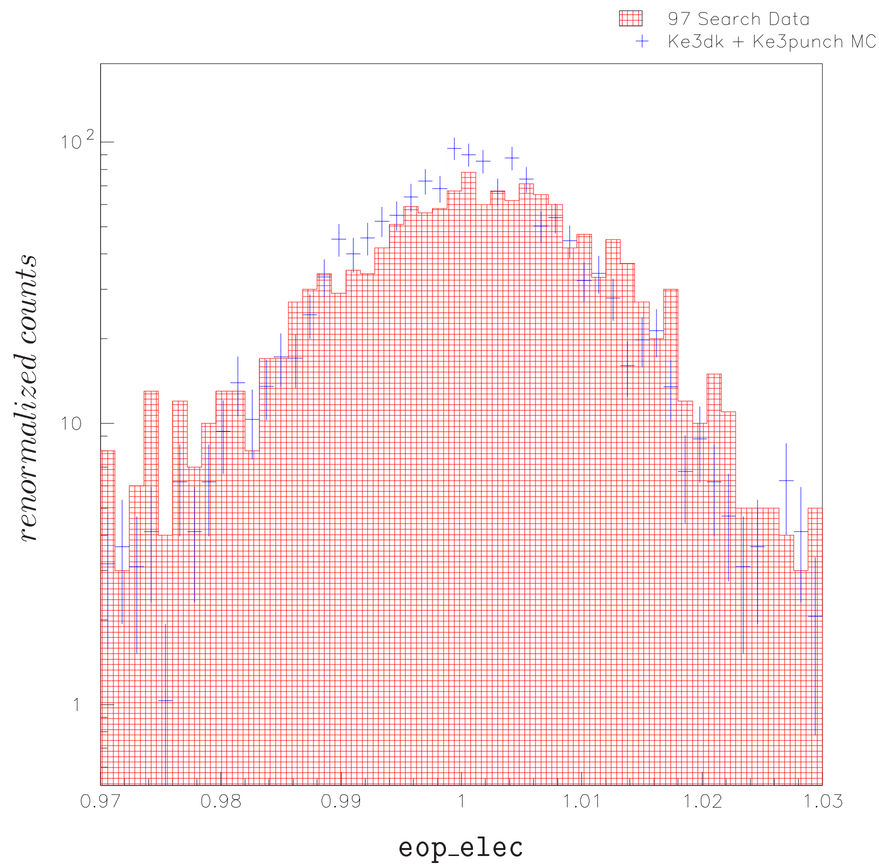

Figure C.2 A comparison of `eop_elec` between 1997 search data in the Ke3 region and Ke3 MC. These curves agree with the identification of this track as being from an electron, supporting that the events were Ke3s.

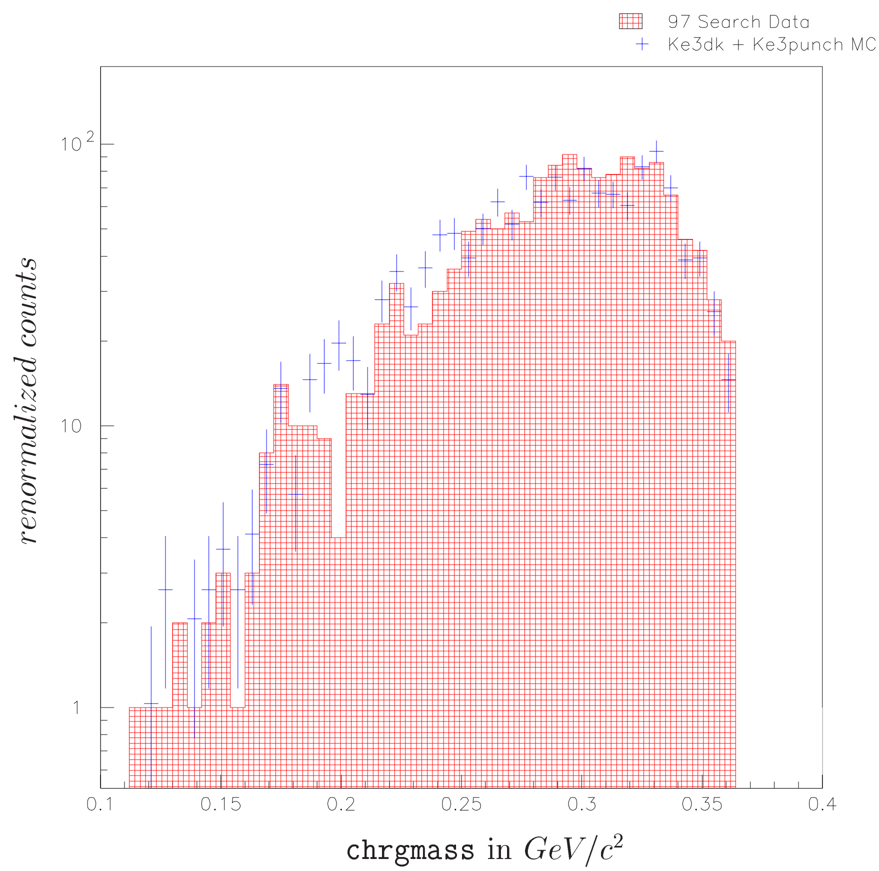

Figure C.3 A comparison of `chrgmass` between 1997 search data in the Ke3 region and Ke3 MC. These curves agreed well enough to support our belief that this subset of search data was Ke3s.

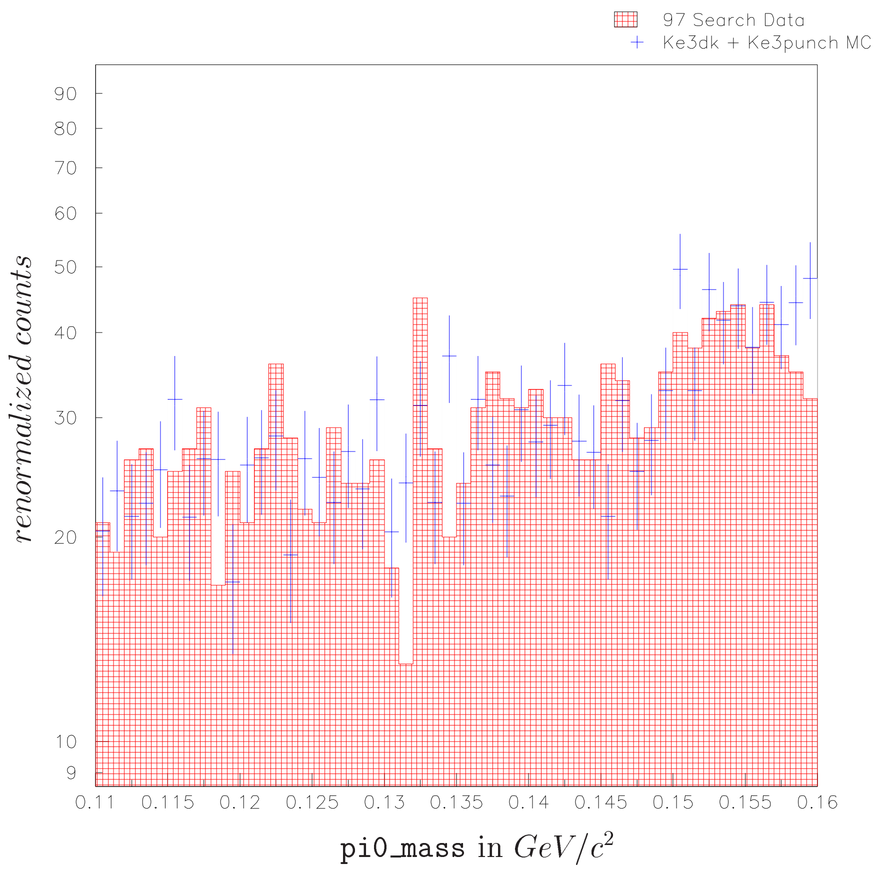

Figure C.4 A comparison of `pi0_mass` (two photon mass) between 1997 search data in the Ke3 region and Ke3 MC (with the correction factor of 2.79×10^{-4} applied). These curves indicate the photons in the events were purely random and not from a neutral pion decay, supporting our belief that this subset of search data was Ke3s.

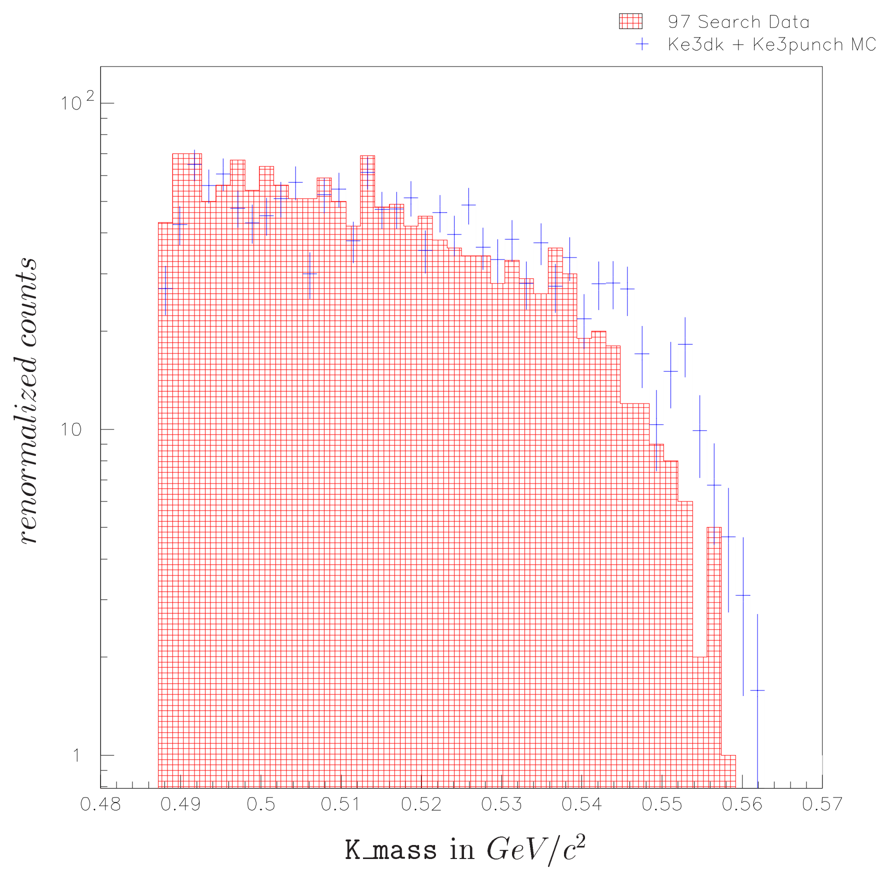

Figure C.5 A comparison of `K_mass` between 1997 search data in the Ke3 region and Ke3 MC. These curves agreed well enough to support our belief that this subset of search data was Ke3s.

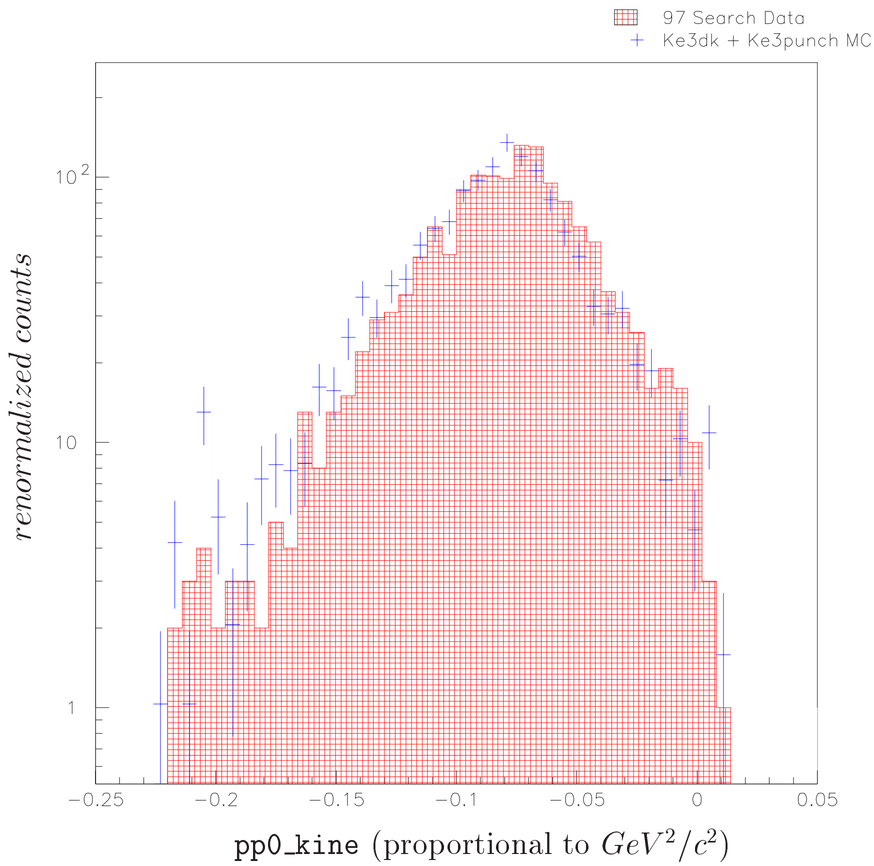

Figure C.6 A comparison of `pp0_kine` between 1997 search data in the Ke3 region and Ke3 MC. There was a slight difference in the data and MC position of the peak, but these curves still agreed well enough to support our belief that this subset of search data was Ke3s.

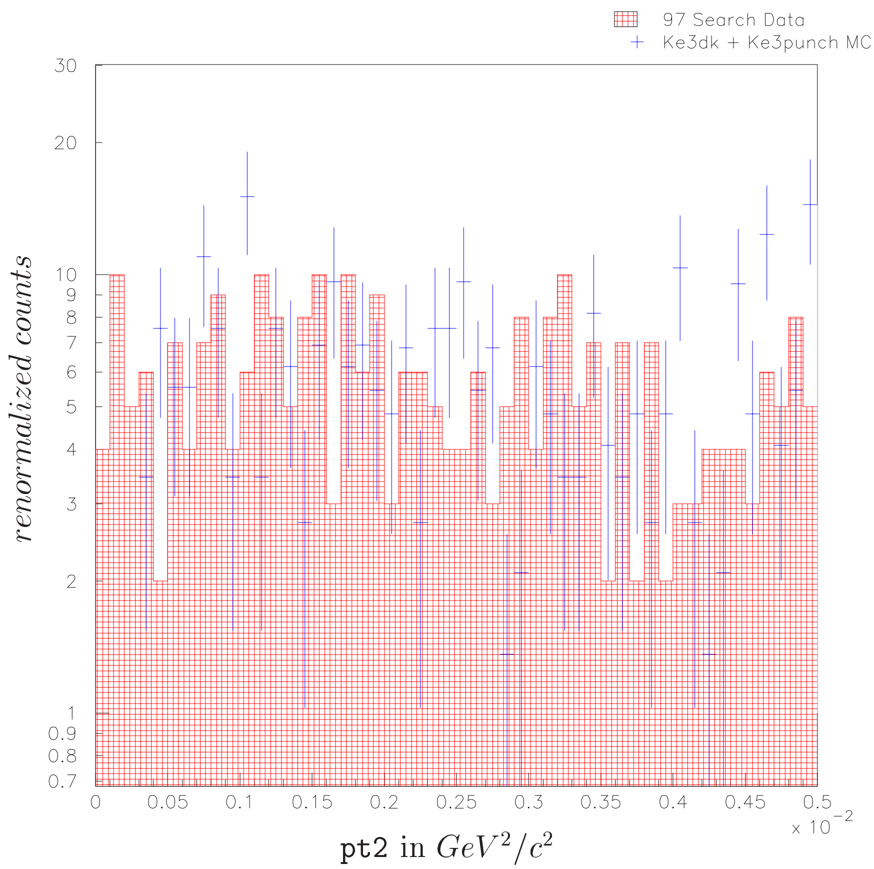

Figure C.7 A comparison of `pt2` between 1997 search data in the Ke3 region and Ke3 MC. These curves agreed well enough to support our belief that this subset of search data was Ke3s.

Appendix D
Bad spills in 1999

Part of the flux lost in the 1999 data set was because of the bad spill cuts applied. Events were cut if any of the following tags were turned on:

Triggers (tag 1)
DPMT ped exp > 0 (tag 2)
Bad DPMT cap (tag 3)
Blown QIE comp. (tag 4)
Misc dead PMT (tag 5)
Broken dynode (tag 8)
Pipeline (tag 9)
Global CsI (tag 10)
ETOT (tag 11)
FERA ADC (tag 12)
Drift chambers (tag 13)
Misc VETO (tag 14)
Muon banks (tag 16)
DAQ/L3 (tag 21)
not 799 run (tag 22)
more than 1 TRD plane dead (tag 26)
Severe TRD prob (tag 28)
Beam problems (tag 29)

These tags added up to the bit mask of 439402399. The only exception was for the 1997 summer data, where tag 3 was not cut on, making the mask 439402395. The tables shown here give some additional detail about where events were lost in 1999. The worst loss was from tag 21 (DAQ/L3). Total badspill loss was 6.34% for trigger 2 and 6.097% for trigger 7.

Tape	tag 2	tag 9	tag 10	tag 13	tag 17	tag 21	tag 22	tag 26	tag 28	tag 29
UPF001	0	0	0	0	0	37118	36435	0	12	0
UPF002	0	0	0	0	2919	0	0	8	0	0
UPF003	0	0	0	0	2316	0	39866	10	6	0
UPF004	0	11109	0	0	0	454936	6312	0	9	0
UPF005	0	1741	0	0	5926	82387	172	0	89	0
UPF006	0	0	0	49805	0	0	0	0	0	0
UPF007	0	11128	0	0	0	0	0	0	0	0
UPF008	0	0	0	0	0	0	491299	0	6	0
UPF009	0	63393	0	0	0	0	16646	0	0	0
UPF010	0	0	0	0	0	103832	0	0	0	0
UPF011	0	0	0	0	0	491860	0	0	62	13589
UPF012	0	0	0	0	7689	232322	0	17	0	0
UPF013	0	0	0	0	3967	131189	0	0	0	0
UPF014	0	0	0	0	0	0	0	0	0	0
UPF015	0	0	0	0	0	124628	0	0	0	0
UPF016	0	0	0	0	0	0	0	171160	0	0
UPF017	0	31426	0	0	0	0	11075	4667	11076	0
UPF018	0	0	0	0	0	41297	54465	58019	302	0
UPF019	0	0	0	0	7501	0	0	0	0	0
UPF020	0	0	0	0	0	0	0	4	4	0

Table D.1 Bad spill bits tagged for trigger #7 in the 1999 on-split data tapes, divided up by tag (tag = bit + 1) number. Tags 1,3,4,5,8,11,12,14, and 16 were also searched for, but no events were found with these tags set.

Tape	tag 2	tag 9	tag 10	tag 13	tag 17	tag 21	tag 22	tag 26	tag 28	tag 29
UPF021	0	0	0	0	0	0	0	0	0	0
UPF022	0	0	0	0	0	56025	0	0	0	0
UPF023	0	16794	0	0	0	0	0	11	3	0
UPF024	0	0	0	0	0	56449	0	0	0	0
UPF025	0	0	0	0	0	32109	0	0	0	0
UPF026	0	0	0	0	0	0	0	0	0	0
UPF027	0	2184	0	0	7589	0	176924	18	20	0
UPF028	0	0	0	0	0	0	0	11	1	0
UPF029	0	0	0	579	28	0	0	55	665	0
UPF030	0	8898	0	1173	19318	0	13798	60	1983	378
UPF031	0	0	0	0	0	0	20166	0	0	0
UPF032	0	0	0	0	0	0	0	0	0	0
UPF033	0	4873	0	0	0	0	0	0	0	0
UPF034	0	0	0	0	0	0	0	0	0	0
UPF035	0	0	0	0	10597	0	0	0	0	0
UPF036	0	0	0	0	0	0	45	26	4845	0
UPF037	0	0	0	0	1545	32403	45195	12281	894	0
UPF038	0	0	0	0	0	38033	0	36187	0	0
UPF039	0	0	0	0	168337	0	0	16876	1022	0
UPF040	0	0	0	0	27917	0	0	8	934	0

Table D.2 (cont.) Bad spill bits tagged for trigger #7 in the 1999 on-split data tapes, divided up by tag (tag = bit + 1) number. Tags 1,3,4,5,8,11,12,14, and 16 were also searched for, but no events were found with these tags set.

Tape	tag 2	tag 9	tag 10	tag 13	tag 17	tag 21	tag 22	tag 26	tag 28	tag 29
UPF041	0	0	0	0	0	0	40672	0	0	0
UPF042	0	1626	0	0	53274	0	0	0	0	0
UPF043	0	0	0	1248	0	0	0	0	0	0
UPF044	0	0	0	0	0	0	0	0	2091	0
UPF045	0	0	0	0	0	0	0	0	2259	0
UPF046	0	0	0	0	0	0	0	0	440	42390
UPF047	0	0	0	0	0	0	0	0	0	0
UPF048	0	0	0	0	0	0	0	0	0	0
UPF049	0	6682	0	0	6860	54733	199	0	17	0
UPF050	0	0	0	0	0	0	0	300727	2	561
UPF051	0	0	0	0	0	0	0	0	5146	0
UPF052	0	0	0	0	0	0	0	0	3	0
UPF053	0	0	0	0	0	0	0	0	0	0
UPF054	0	0	0	0	0	0	27	22	5	0
UPF055	0	2079	0	0	0	0	0	0	1248	0
UPF056	0	0	0	0	0	0	0	0	4692	119
UPF057	0	0	0	0	0	0	0	10	993	0
UPF058	0	0	0	0	4849	0	0	0	2383	0
UPF059	0	19611	0	0	0	0	18796	0	20112	0
UPF060	0	0	0	0	0	91011	0	0	0	0

Table D.3 (cont.) Bad spill bits tagged for trigger #7 in the 1999 on-split data tapes, divided up by tag (tag = bit + 1) number. Tags 1,3,4,5,8,11,12,14, and 16 were also searched for, but no events were found with these tags set.

Tape	tag 2	tag 9	tag 10	tag 13	tag 17	tag 21	tag 22	tag 26	tag 28	tag 29
UPF061	0	0	0	0	0	0	0	0	82	0
UPF062	0	0	0	0	0	0	0	0	24	0
UPF063	0	11696	0	0	28	0	0	0	1258	0
UPF064	0	0	0	0	0	0	0	4890	0	87
UPF065	0	0	0	0	1659	0	36271	11	0	0
UPF066	534868	0	0	0	4331	0	0	0	43	0
UPF067	486812	0	0	0	0	0	0	0	0	0
UPF068	0	0	0	0	0	27282	0	35	90	0
UPF069	0	19733	0	0	0	0	0	0	29	0
UPF070	0	0	0	0	0	0	0	112	0	0
UPF071	0	4305	0	0	258	23367	0	25	14	0
UPF072	0	0	0	0	0	0	0	0	0	0
UPF073	0	0	2047	0	0	0	0	0	812	0
UPF074	1971	0	0	46838	0	0	0	0	1173	0
UPF075	0	0	0	0	0	30996	0	0	0	0
UPF076	0	12533	0	0	7083	0	0	0	0	0
UPF077	0	0	0	0	46655	43636	0	15797	9	0
UPF078	0	0	0	0	0	0	0	0	63	0
UPF079	0	6760	0	0	0	0	0	0	1804	0
UPF080	0	0	0	0	0	0	0	0	0	0
Total	1023651	236571	2047	99643	390413	2185613	1008363	621037	65732	57849

Table D.4 (cont.) Bad spill bits tagged for trigger #7 in the 1999 on-split data tapes, divided up by tag (tag = bit + 1) number. Tags 1,3,4,5,8,11,12,14, and 16 were also searched for, but no events were found with these tags set.

Tape	% Trigger 2 events cut	% Trigger 7 events cut
UPF001	4.05±0.03%	6.08±0.02%
UPF002	0.44±0.01%	0.263±0.005%
UPF003	5.89±0.05%	4.37±0.02%
UPF004	46.3±0.1%	48.00±0.05%
UPF005	8.62±0.07%	8.11±0.03%
UPF006	4.66±0.06%	4.59±0.02%
UPF007	1.01±0.03%	1.00±0.01%
UPF008	44.8±0.1%	44.70±0.05%
UPF009	6.93±0.07%	6.97±0.02%
UPF010	9.38±0.08%	9.04±0.03%
UPF011	47.8±0.1%	46.53±0.05%
UPF012	22.5±0.1%	21.77±0.04%
UPF013	13.4±0.1%	11.98±0.03%
UPF014	0.%	0.%
UPF015	11.56±0.09%	10.95±0.03%
UPF016	15.4±0.1%	15.08±0.03%
UPF017	5.65±0.06%	5.31±0.02%
UPF018	14.6±0.1%	14.17±0.03%
UPF019	0.701±0.024%	0.686±0.008%
UPF020	0.003±0.002%	0.0007±0.0003%
UPF021	0.%	0.%
UPF022	5.43±0.06%	5.11±0.02%
UPF023	1.51±0.03%	1.52±0.01%
UPF024	5.38±0.06%	5.06±0.02%
UPF025	2.94±0.05%	2.81±0.02%
UPF026	0.%	0.%
UPF027	9.32±0.08%	15.62±0.03%
UPF028	0.020±0.004%	0.0010±0.0003%
UPF029	0.679±0.021%	0.111±0.003%
UPF030	2.38±0.04%	3.83±0.02%

Table D.5 Percentages of events cut by bad spill tags in the 1999 data.

Tape	% Trigger 2 events cut	% Trigger 7 events cut
UPF031	1.81±0.04%	1.73±0.01%
UPF032	0.%	0.%
UPF033	0.403±0.017%	0.420±0.006%
UPF034	0.%	0.%
UPF035	0.924±0.026%	0.911±0.009%
UPF036	0.615±0.021%	0.403±0.006%
UPF037	7.65±0.07%	7.92±0.03%
UPF038	6.54±0.07%	6.25±0.02%
UPF039	15.8±0.1%	15.47±0.03%
UPF040	2.35±0.04%	2.36±0.01%
UPF041	3.27±0.05%	3.32±0.02%
UPF042	4.42±0.05%	4.33±0.02%
UPF043	0.108±0.009%	0.104±0.003%
UPF044	0.173±0.011%	0.172±0.004%
UPF045	0.190±0.012%	0.190±0.004%
UPF046	3.78±0.05%	3.60±0.02%
UPF047	0.%	0.%
UPF048	0.%	0.%
UPF049	5.83±0.06%	5.34±0.02%
UPF050	24.4±0.1%	24.52±0.04%
UPF051	0.425±0.018%	0.442±0.006%
UPF052	0.%	0.0003±0.0001%
UPF053	0.%	0.%
UPF054	0.034±0.005%	0.0046±0.0006%
UPF055	0.611±0.030%	0.57±0.01%
UPF056	0.417±0.017%	0.395±0.006%
UPF057	0.101±0.009%	0.045±0.001%
UPF058	0.667±0.023%	0.612±0.007%
UPF059	5.86±0.07%	4.90±0.02%
UPF060	8.25±0.08%	7.63±0.02%

Table D.6 (cont.) Percentages of events cut by bad spill tags in the 1999 data.

Tape	% Trigger 2 events cut	% Trigger 7 events cut
UPF061	0.%	0.0067±0.0007%
UPF062	0.%	0.0019±0.0004%
UPF063	1.02±0.03%	1.006±0.009%
UPF064	0.391±0.016%	0.444±0.006%
UPF065	3.12±0.04%	2.92±0.01%
UPF066	38.6±0.1%	40.89±0.04%
UPF067	37.1±0.1%	38.02±0.04%
UPF068	2.38±0.04%	2.19±0.01%
UPF069	1.55±0.03%	1.56±0.01%
UPF070	0.007±0.002%	0.0088±0.0008%
UPF071	2.77±0.04%	2.20±0.01%
UPF072	0.%	0.%
UPF073	0.291±0.016%	0.220±0.004%
UPF074	3.96±0.05%	3.85±0.02%
UPF075	2.63±0.04%	2.43±0.01%
UPF076	1.07±0.03%	1.52±0.01%
UPF077	5.54±0.06%	8.20±0.02%
UPF078	0.0007±0.0007%	0.0049±0.0006%
UPF079	0.656±0.022%	0.683±0.007%
UPF080	0.%	0.%
All tapes	6.34±0.01%	6.097±0.003%

Table D.7 (cont.) Percentages of events cut by bad spill tags in the 1999 data.

Appendix E
Code for PDF 90% confidence level limit

The FORTRAN code used to calculate the probability density function (PDF) branching ratio limits at a 90% confidence level (CL) is included here. The code was used as written for the combined 1997 and 1999 limit. To calculate the separate limits for 1997 and for 1999, the `do` loop involving the variable `j` was changed to (`do j=1,1`) for 1997 and (`do j=2,2`) for 1999. The code was written by KTeV collaborators Hogan Nguyen and Leo Bellantoni based on a statistical methods book[43].

```
c-------------------------------------------------------------
      subroutine lfv

c-----Combine 97 and 99 results and set the limit on the
c     Branching Ratio given the following

      integer n(2)             ! number of events observed
      real  b(2), a(2), fkl(2) ! background, acceptance,
                               !   and flux estimates
      real eb(2),ea(2),efkl(2) ! errors for the above

c-----relevant numbers for 97

      n(1)   = 2               ! number of events observed
      b(1)   = 0.50            ! background estimate
      a(1)   = 5.28e-2         ! signal acceptance
      fkl(1) = 2.55e11         ! flux for this mode
      eb(1)  = 0.13            ! gaussian errors
      ea(1)  = 0.07e-2         !       "
      efkl(1) = 0.050e11       !       "

c-----relevant numbers for 99
```

```
      n(2)   = 3               ! number of events observed
      b(2)   = 0.48            ! background estimate
      a(2)   = 3.576e-2        ! signal acceptance
      fkl(2) = 3.137e11        ! flux for this mode
      eb(2)  = 0.14            ! gaussian errors
      ea(2)  = 0.018e-2        !        "
      efkl(2) = 0.054e11       !        "

c-----book histograms to keep the PDF and its integral

      nbin = 1000
      brmin = 0
      brmax = 1.5e-9
      bw2 = (brmax-brmin)/nbin/2

      call hdelet(101)
      call hbook1(101,'Bayesian estimate PDF',nbin,brmin,brmax,0.)
      call hdelet(102)
      call hbook1(102,'Integral of PDF',nbin,brmin,brmax,0.)

c-----To account for the uncertainties in the background,
c     acceptance, and other nuisance parameters, we shall
c     use a monte carlo technique to integrate over them.
c     NSAMPLE is the number of times to sample these PDFs.
c     NSAMPLE has to be reasonably large.
c     The uncertainties are assumed to be gaussian.

      nsample = 300

c-----loop over each possible value of BR and calculate
c     its likelihood, and then fill the histograms

      ytot = 0
      do i=1,nbin
         call hix(101,i,x)
         x = x + bw2
c--------integrate over the all the uncertainties by randomly
c        sampling over the PDF (assumed to be gaussian).
         do isample = 1,nsample
c----------loop over the two experiments 97 and 99
```

```
          yjoint = 1.    ! the joint PDF
          do j=1,2
c------------xtot is the expected number of events given
c             br=x, background = b(j), and gaussian
c             uncertainties
              a_rndm   =   a(j) +   ea(j)*gauss()
              fkl_rndm = fkl(j) + efkl(j)*gauss()
              b_rndm   =   b(j) +   eb(j)*gauss()
              xtot = x * a_rndm * fkl_rndm + b_rndm
c------------y is the likelihood to observe n, given xtot
              y = p(xtot,n(j))
              yjoint = yjoint * y
          enddo
c----------fill the hists with the joint PDF and its integral
          ytot = ytot + yjoint
          call hf1(101,x,yjoint)
        enddo
        call hf1(102,x,ytot)
      enddo

c-----now calculate the 90% UL

      do i=1,nbin
         call hix(102,i,x)
         x = x + bw2
         if(hi(102,i)/ytot.gt..9)then
            print *,'90% UL is ',x
            return
         endif
      enddo

      end

c---------------------------------------------------

      real function p(r,n)

c-----the poisson probability

      real r     ! mean
```

```
      integer n  ! observation

      ifact = 1
      do j=1,n
         ifact = j*ifact
      enddo
      p = exp(-r)*r**n/ifact

      end

c---------------------------------------------------

      real function gauss()

      Implicit none

c-----gaussian number generator (got from Leo)

      Real v1,v2,r,fac,rndm

   9  v1  = 2*rndm()-1.
      v2  = 2*rndm()-1.
      r   = (v1**2 + v2**2
      if (r.gt.1.0) goto 9
      fac = sqrt(-2.0*log(r)/r)
      gauss = v2*fac

      Return
      end
```

References

1. Fermilab Photo Archive #95-759D.

2. Tipler, Paul A., *Modern Physics*. Worth Publishers, New York: 1978. pp. 456-7.

3. Griffiths, David, *Introduction to Elementary Particles*. Harper & Row, New York: 1987. p. 128.

4. T. D. Lee and C. N. Yang, "Question Of Parity Conservation In Weak Interactions," Phys. Rev. **104**, 254 (1956).

5. C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes and R. P. Hudson, "Experimental Test Of Parity Conservation In Beta Decay," Phys. Rev. **105**, 1413 (1957).

6. R. L. Garwin, L. M. Lederman and M. Weinrich, "Observation Of The Failure Of Conservation Of Parity And Charge Conjugation In Meson Decays: The Magnetic Moment Of The Free Muon," Phys. Rev. **105**, 1415 (1957).

7. Y. Fukuda *et al.* [Super-Kamiokande Collaboration], "Evidence for oscillation of atmospheric neutrinos," Phys. Rev. Lett. **81**, 1562 (1998) [arXiv:hep-ex/9807003].

8. K. Hagiwara *et al.* [Particle Data Group Collaboration], "Review Of Particle Physics," Phys. Rev. D **66**, 010001 (2002).

9. Halzen, Francis, and Alan D. Martin. *Quarks and Leptons: An Introductory Course in Modern Particle Physics.* John Wiley & Sons, New York: 1984. p. 315.

10. Perkins, Donald H. *Introduction to High Energy Physics.* Addison-Wesley Publishing Company, Reading, Massachusetts: 1982. p. 17.

11. Sakarov, A.D., "Violation of CP invariance, C asymmetry, and the baryon asymmetry of the universe," JETP Lett. **6**, 24 (1967).

12. J. H. Christenson, J. W. Cronin, V. L. Fitch and R. Turlay, "Evidence For The 2 Pi Decay Of The K(2)0 Meson," Phys. Rev. Lett. **13**, 138 (1964).

13. Gottfried, Kurt and Victor F. Weisskopf, *Concepts of Particle Physics, Volume I.* Oxford University Press, New York: 1986, p.153.

14. L. K. Gibbons *et al.* [E731 Collaboration], "CP and CPT symmetry tests from the two-pion decays of the neutral kaon with the Fermilab-E731 detector. (Revised version)," Phys. Rev. D **55**, 6625 (1997).

15. L. K. Gibbons *et al.*, "Measurement Of The CP Violation Parameter Re (Epsilon-Prime / Epsilon)," Phys. Rev. Lett. **70**, 1203 (1993).

16. G. D. Barr *et al.* [NA31 Collaboration], "A New measurement of direct CP violation in the neutral kaon system," Phys. Lett. B **317**, 233 (1993).

17. A. Alavi-Harati [KTeV Collaboration], "Measurements of direct CP violation, CPT symmetry, and other parameters in the neutral kaon system," arXiv:hep-ex/0208007. Submitted to Phys. Rev. D, August 6, 2002.

18. A. Lai *et al.* [NA48 Collaboration], "A precise measurement of the direct CP violation parameter Re(epsilon'/epsilon)," Eur. Phys. J. C **22**, 231 (2001) [arXiv:hep-ex/0110019].

19. S. Fukuda *et al.* [Super-Kamiokande Collaboration], "Constraints on neutrino oscillations using 1258 days of Super-Kamiokande solar neutrino data," Phys. Rev. Lett. **86**, 5656 (2001) [arXiv:hep-ex/0103033].

20. Q. R. Ahmad *et al.* [SNO Collaboration], "Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory," Phys. Rev. Lett. **89**, 011301 (2002) [arXiv:nucl-ex/0204008].

21. P. Langacker, S. Uma Sankar and K. Schilcher, "K-L $\to$ Mu E In SU$(2)_L$ X U(1) And SU$(2)_L$ X SU$(2)_R$ X U(1) Models With Large Neutrino Masses," Phys. Rev. D **38**, 2841 (1988).

22. H. K. Dreiner, "An introduction to explicit R-parity violation," arXiv:hep-ph/9707435.

23. A. Belyaev, M. Chizhov, A. Dorokhov, J. R. Ellis, M. E. Gomez and S. Lola, "Charged-lepton flavour violation in kaon decays in supersymmetric theories," Eur. Phys. J. C **22**, 715 (2002) [arXiv:hep-ph/0008276].

24. W. S. Hou and A. Soni, "Experimental Consequences Of A Horizontal Gauge Model For CP Violation," Phys. Rev. Lett. **54**, 2083 (1985).

25. L. J. Hall and L. Randall, "CP Violation From Scalar Leptoquarks," Nucl. Phys. B **274**, 157 (1986).

26. K. Arisaka *et al.*, "Improved Upper Limit on the Branching Ratio B($K_L^0 \to \mu^\pm e^\mp$)," Phys. Rev. Lett. **70**, 1049 (1993).

27. K. Arisaka *et al.*, "Search For The Lepton-Family Number Violating Decays K(L) → Pi0 Mu+- E-+," Phys. Lett. B **432**, 230 (1998).

28. A. Bellavance, "Search For K(L) → Pi0 Mu+- E-+," Int. J. Mod. Phys. A **16S1B**, 651 (2001).

29. Fermilab Photo Archive #00-635D.

30. A. Affolder *et al.* [The KTeV E832/E799 Collaboration], "Observation of the decay Xi0 → Sigma+ e- anti-nu/e," Phys. Rev. Lett. **82**, 3751 (1999).

31. K. Arisaka *et al.*, *KTeV Design Report*. FERMILAB-FN-580, Fermi National Accelerator Laboratory, Batavia, IL: Jan. 22, 1992.

32. KTeV internal note #0518, *The KTeV Pure CsI Calorimeter*.

33. KTeV internal note #0187.

34. DOE/ER-0457T, U.S. NIM committee, May 1990; Standard NIM Instrumentation System, NTIS, U.S. Department of Commerce, Springfield, Virginia 22161.

35. KTeV internal note #0374, *Guide to Accessing KTEVANA Data*.

36. KTeV internal note #0317, *CsI Clustering*.

37. KTeV internal note #0577, *CsI Clustering and Corrections*.

38. KTeV internal note #0328, *T3 Tracking Algorithm User's Manual*.

39. KTeV internal note #0555, *About the TRD offline calibration*.

40. A. M. Bellavance, "Ktev E799II Search For The Lepton-Flavor-Number Violating Decay Kaon(Long) Going To Neutral Pion Charged Muon Charged Electron," UMI-13-99250. Masters Thesis.

41. G. J. Feldman and R. D. Cousins, "A Unified approach to the classical statistical analysis of small signals," Phys. Rev. D **57**, 3873 (1998) [arXiv:physics/9711021].

42. H. B. Greenlee, "Background To K0(L) → Pi0 E E From K0(L) → Gamma Gamma E E," Phys. Rev. D **42**, 3724 (1990).

43. W. T. Eadie, *Statistical Methods in Experimental Physics*. American Elsevier Pub. Co., New York: 1971.