THE UNIVERSITY OF CHICAGO

MEASUREMENTS OF THE DIRECT CP-VIOLATING PARAMETER $\mathrm{Re}(\epsilon'/\epsilon)$ AND THE KAON SECTOR PARAMETERS Δm, τ_S, AND ϕ_{+-}

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

DEPARTMENT OF PHYSICS

BY
JAMES A. GRAHAM

CHICAGO, ILLINOIS
DECEMBER 2001

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

This thesis is an experimental investigation of the violation of the CP symmetry of elementary particle interactions. We have used data from the 1996–1997 run of the KTeV experiment at Fermi National Accelerator Laboratory (Fermilab), to measure the "direct" CP-violation parameter $\mathrm{Re}(\epsilon'/\epsilon)$ in neutral kaon decays. We find $\mathrm{Re}(\epsilon'/\epsilon) = [20.7 \pm 1.5\ (\mathrm{stat}) \pm 2.3\ (\mathrm{syst})] \times 10^{-4}$. In addition, we have made precise determinations of the K_S lifetime τ_S, the K_L - K_S mass difference Δm, and the phase $\phi_{+-} \equiv \arg(\eta_{+-})$, using a data sample of $K \to \pi^+\pi^-$ decays. These measurements yield $\tau_S = [0.8965 \pm 0.0003\ (\mathrm{stat}) \pm 0.0004\ (\mathrm{syst})] \times 10^{-10}\ s$, $\Delta m = [0.5267 \pm 0.0006\ (\mathrm{stat}) \pm 0.0014\ (\mathrm{syst})] \times 10^{10}\ \hbar s^{-1}$, and $\phi_{+-} = [44.12 \pm 0.72\ (\mathrm{stat}) \pm 1.14\ (\mathrm{syst})]^\circ$.

ACKNOWLEDGEMENTS

My first thanks go to my advisor, Edward Blucher. He has been a true advisor and a true friend. It was a pleasure with him, and it didn't feel like working for him. He has taught me how to do analysis and how to enjoy doing it. I also enjoyed running and skiing with him. And he let me stay on his couch when I needed to come back to Chicago.

My second thanks go to Bruce Winstein. He has taught me the correct way to look at many physical problems. He has a keen insight that forced me to think about the clearest and most correct way to present a problem. Long discussions about all topics, physics and non-physics, were always illuminating.

Whoever said you learn more in graduate school from the students than from the professors certainly knew how these things work. Many thanks to the other 'clown prince of KTeV' (his phrase, not mine!), Val Prasad. Half the work (the hard half) presented in this thesis is his [1], and I don't think that I would have survived graduate school without him to deflect most of the heat. I am grateful to those who came before us, allowing us to stand high on shoulders of really smart people. I learned how to write code from Peter Shawhan, and he provided a blueprint for this thesis [2]. For his invention of crazy new variables, and looking at more charged-tracking data than any sane man would, thanks to Colin Bown. To our E799 compatriots: Greg, Eric, Breese and Steve, a big thanks. And further thanks to Elizabeth Turner for coming along and taking the next step for KTeV. I owed much to the post-docs on this experiment; Elliott Cheu for getting me started, Rick Kessler for always having an orthogonal view to any problem, and Aaron Roodman for showing how much information one can extract from data. Sasha Glazov arrived later but contributed as much as anyone to the success of KTeV. At the U. of C., Scott Oser, Miguel Barrio, Craig Wiegert and Jordan Koss get my eternal gratitude for getting me through first-year.

This thesis is based on the collaborative effort of roughly 80 physicists and students. To our collaborators, all the best. At Fermilab, Bob Tschirhart, Bob Hsiung, Peter Shanahan; and extra thanks to Viv O'Dell for leaving me on the 'Current Obstacle to KTeV Success' list for so long. At Arizona, thanks to Sydney Taegar for his work on splitting, and for being a sane individual. At Rutgers, thanks to Sunil Somalwar and Rick Tesarek for getting me started in hardware.

In the non–KTeV world, I've been lucky enough to meet Jordan Koss, Steve Bright, Joseph Biello, Gil Holder and Ken Nollett. Big thanks to my roommates, Gordon Richards and later Matt Blanton, first for wanting to move in with me, and second for always being willing to have a beer when necessary. Extra thanks to the entire Blanton clan (Sean, Ipek, and Matt) for being good friends.

My life in Chicago would have been a terrible time if not for Peggy Hines. I would not have had the stamina to finish if not for her support during two difficult years.

I'd like to thank my friends for supporting me, although they might not understand what I do, or why I do it. In Vancouver, Peter Roeck and Chris Miller, and the boys from the old school; Howie, Bink, Mike, Dave – you know who you are. Also from way back, big thanks to Pat Band; I've come to depend on heavily on your insight and friendship. It's great to be back in the same town with you and Jen. Also in Toronto, thanks go to Hollie Shaw and Matt Syberg-Olsen for being great friends.

I could not have made it through the last years of this thesis without the love of Cindy Newton. Thank you for finding me and for putting up with me while I finished what I started.

My family have been incredibly supportive of my graduate school career. My thanks go to my sister for her support and encouragement. Finally, my most heartfelt thanks to the people most responsible for me getting to this point: my parents. I am incredibly lucky to have parents who encourage and support me to do whatever I think I can accomplish. I cannot thank you enough.

CHAPTER 1
INTRODUCTION

I would rather discover a single fact, even a small one, than debate the great issues at length without discovering anything at all.

Galileo Galilei

The subject of this thesis is symmetries of Nature. Symmetry plays a large role in everyday life, from the complex idea of what we find aesthetically pleasing down to the fundamentals of physical law.

The fundamental symmetries and the accompanying conservation laws of Nature give us insight into the properties and interactions of subatomic particles. Aesthetically, symmetries give beauty to the formal description of physics, and practically, often point to new physics, either by their presence or absence. Some of the greatest insights in physics have arisen from the realization of a hidden symmetry, or the violation of a symmetry when one was believed to exist.

The search for symmetries in the study of elementary particle physics has led to many major discoveries. Of particular interest are the three discrete symmetries of space inversion or parity (P), charge conjugation (C), and time reversal (T). It is known [3–5] that any local field theory that is invariant under proper Lorentz transformations is also invariant under the combined operation of the product CPT. Since the "Standard Model" of particle physics is a local field theory, it is assumed that the physical world is described by a local field theory, and that CPT is an existing symmetry of Nature. The CPT symmetry immediately leads to some very basic predictions, such as the equality of the masses, the lifetimes, and magnitude of the electric charge of a particle and its antiparticle [6].

An elegant way of preserving the CPT symmetry is to have the particle interactions separately invariant under each of the three operators C, P, and T. Since there is strong evidence that the strong and electromagnetic interactions are separately symmetric under C, P, and T, it was generally assumed that so too was the weak interaction. However, in 1956, Lee and Yang [7] proposed that parity violation could explain the so-called "τ-θ puzzle," where two particles seemed to have identical properties (and in fact were the same particle, the K^+), but were thought to be distinct because they decayed into two different final states of opposite parity: a 3π state of parity $P = -1$ and a 2π state with parity $P = +1$. The bias towards conservation of parity in the weak interaction was so strong, however, that even Lee and Yang themselves didn't consider their argument too seriously. They did suggest how to establish parity–non-conservation in the weak interaction experimentally; the observation of an angular asymmetry in the β-decay of spin-oriented nuclei, or observed in an asymmetry in the distribution of angles between the muon and the electron in the decay sequence $\pi \rightarrow \mu + \nu_\mu$ followed by $\mu \rightarrow e\nu_\mu\nu_e$.

The experimental determination of the parity-violating nature of weak interaction was accomplished first by Wu [8] by observing the angular asymmetry of the decay electron in the β-decay of spin-oriented ^{60}Co, and then later in the pion decay experiments of Garwin *et al.* [9] and Friedman and Telegdi [10]. This quickly led to the realization of the $V - A$ nature of the weak interaction.

It was also realized that the aforementioned experiments also demonstrated that charge-conjugation symmetry (C) was violated. Landau [11] had pointed out before the experiments themselves that some of the aesthetic nature of CPT conservation could be restored if the combined operator CP was conserved, such that CP and T were separately conserved. At the time, the experimental evidence supported this notion; the pion decay experiments indicated that the final decay states of the pion had the same value for the multiplicative quantum number CP (although they had different individual values for the quantum numbers C and P). The canonical example of the combined CP symmetry in the Standard Model is the neutrino. The Standard Model has a left-handed neutrino and a right-handed antineutrino (where the "handed-ness" describes its parity). The combined operator CP takes a left-handed neutrino to a right-handed antineutrino, and vice-versa. However, individually, C and P operating on the physical states would create left-handed antineutrinos and right-handed neutrinos, neither of which have been observed.

The belief that CP was a symmetry was shattered by the discovery by Christenson, Cronin, Fitch and Turlay in 1964 [12] of CP violation in decays of the long-lived component of the neutral kaon. The initial experiment showed that a long-lived particle of the same mass as the K_2 decayed to $\pi^+\pi^-$. Initially, some thought that there might be a separate particle of the same mass but of different CP which could decay to $\pi^+\pi^-$ without violating CP [13, 14]. But a follow-up experiment by Fitch *et al.* demonstrated interference in the $\pi^+\pi^-$ decays of the short- and long-lived states downstream of a diffuse beryllium "regenerator", confirming that the K_S and K_L eigenstates decayed to $\pi^+\pi^-$ [15]. Later experiments showed that $\pi^0\pi^0$ decays behave in a similar manner, indicating that the dominant effect is a CP-violating asymmetry in K^0-$\overline{K^0}$ mixing [16, 17].

It has been nearly 40 years since the initial discovery of CP violation in the neutral kaon sector, yet still not much is known about the nature of CP violation. The origin of C and P violation within the Standard Model is not so much of a mystery; it arises through the left-handed nature of the interactions of the W boson. CP violation is accommodated in the Standard Model — the three generations of quark families lead to a natural phase in the weak interactions. The origin of this phase is still unknown. It is possible that some part of CP violation arises from new interactions entirely outside the Standard Model.

However, it is known that CP violation is one of three necessary ingredients for the evolution of the universe. As detailed by Sakharov [18], our current universe can have only arrived in its current state if at one time it was out of thermodynamic equilibrium, had baryon-number violation, and most importantly for this discussion, had manifest CP violation. The level of CP violation seen in the kaon sector is not large enough to drive the predominance of matter in the universe after the Big Bang [19, 20], but provides an interesting "proof of principle" of how this effect may have arisen. At the microscopic level, the kaon sector gives us insight as to the difference between matter and antimatter.

1.1 Kaon Phenomenology

It is now useful to put the concepts of CP violation into more concrete mathematical terms to evaluate the effects that the KTeV experiment is trying to measure. In the course of this discussion, we will cover some interesting topics of both the strong and the weak interactions. There are many fine reviews of K mixing and CP violation [21–25] which provide more detail than presented here.

The strangeness eigenstates of the neutral kaon are the K^0 ($S = +1$) and the $\overline{K^0}$ ($S = -1$). In terms of the Standard Model, these are identified as $K^0 = (d\bar{s})$ and $\overline{K^0} = (\bar{d}s)$; however this identification is not necessary for our discussion. The K^0 and $\overline{K^0}$ mesons are strangeness eigenstates produced through the strong interaction, and are charge conjugates of each other. It was recognized early in the field of kaon physics that the decay of the K^0 to pion final states such as $\pi^+\pi^-$ does not conserve strangeness. More interesting, however, is the fact that the charge conjugate reaction has the same final pion state. Gell-Mann and Pais [26] observed that the weak interaction responsible for the decay of neutral kaons mixes the K^0 and the $\overline{K^0}$, through virtual $\Delta S = 2$ transitions such as $K^0 \leftrightarrow \pi^+\pi^- \leftrightarrow \overline{K^0}$. The mixed states take on a definite physical meaning, as opposed to being simple mathematical constructs. Furthermore, as Gell-Mann and Pais pointed out, the mixed initial states and the final states have definite CP, and hence should have much different lifetimes; the state that decays to $\pi\pi$ is favored by phase-space considerations and should be much shorter-lived than the state that decays to 3π. Hence the names K_S ("K-short") and K_L ("K-long") are quite natural. It is useful to examine this system in more detail, to identify the weak eigenstates, since the weak force is responsible for the decay process which defines the particles with definite masses and lifetimes.

We start with a definition of the CP operator on the K^0 and $\overline{K^0}$ states. Since the strong and electromagnetic interactions are invariant under CP, they do not connect K^0 and $\overline{K^0}$. The weak interaction does connect the two states, but only up to a phase factor $\exp(-i\phi S)$, where S is the strangeness of the kaon:

$$\begin{aligned} CP\left|K^0\right\rangle &= e^{-i\phi S}\left|\overline{K^0}\right\rangle = e^{-i\phi}\left|\overline{K^0}\right\rangle \\ CP\left|\overline{K^0}\right\rangle &= e^{-i\phi S}\left|K^0\right\rangle = e^{+i\phi}\left|K^0\right\rangle. \end{aligned} \tag{1.1}$$

The phase ϕ is not physically observable, so we make the choice $\phi = 0$, such that

$$CP\left|K^0\right\rangle = \left|\overline{K^0}\right\rangle, \tag{1.2}$$

although other choices (notably $\phi = \pi$) are used in the literature.

Since no evidence for CP symmetry breaking was seen until 1964, and was seen to be small at that point, a sensible starting point is to assume that the weak eigenstates are small perturbations away from the CP eigenstates, which are defined as:

$$\begin{aligned} K_1 &= \frac{1}{\sqrt{2}}\left(K^0 + \overline{K^0}\right) \qquad (CP = +1), \\ K_2 &= \frac{1}{\sqrt{2}}\left(K^0 - \overline{K^0}\right) \qquad (CP = -1). \end{aligned} \tag{1.3}$$

Note that if the weak interaction were to conserve CP, these would be the eigenstates of the weak interaction. Due to the CP of the final pion states, the K_1 is expected to decay to two-particle final states $\pi^+\pi^-$ or $\pi^0\pi^0$ (which have definite $CP = +1$); the K_2 is expected to decay to the three-particle states of definite $CP = -1$, such as $\pi^+\pi^-\pi^0$ and $3\pi^0$. It was natural to identify the K_1 as the K_S and the K_2 as the K_L, which was done until the discovery of CP violation in the K sector.

We now consider an effective Hamiltonian and an equation of motion that connects the K^0 state to the $\overline{K^0}$ state. We write a given kaon state as a superposition of the K^0 and $\overline{K^0}$ basis states:

$$K = a_1 K^0 + a_2 \overline{K^0}. \tag{1.4}$$

The Schrödinger-like equation is[1]

$$i\frac{d}{dt}\begin{pmatrix} a_1 \\ a_2 \end{pmatrix} = \mathcal{H}_{\text{eff}}\begin{pmatrix} a_1 \\ a_2 \end{pmatrix}. \tag{1.5}$$

The effective Hamiltonian $\mathcal{H}_{\text{eff}}$ can be decomposed into two sub-matrices

$$\begin{aligned} \mathcal{H}_{\text{eff}} &= \mathbf{M} - \frac{i}{2}\Gamma & (1.6) \\ &= \left[\begin{pmatrix} M_{11} & M_{12} \\ M_{12}^* & M_{22} \end{pmatrix} - \frac{i}{2}\begin{pmatrix} \Gamma_{11} & \Gamma_{12} \\ \Gamma_{12}^* & \Gamma_{22} \end{pmatrix}\right] & (1.7) \end{aligned}$$

where the "mass-matrix" $\mathbf{M}$ and the "decay-matrix" Γ are individually Hermitian, although $\mathcal{H}_{\text{eff}}$ is not. The Hermitian nature of the two matrices requires $\mathbf{M}_{11}$, $\mathbf{M}_{22}$, Γ_{11}, and Γ_{22} to all be real. Γ is the absorptive term that is responsible for the decay of the kaon into physically accessible final states, hence its name. The off-diagonal elements of the mass-matrix would mix the K^0 and $\overline{K^0}$ states in the absence of the decay matrix. From this $\mathcal{H}_{\text{eff}}$, we wish to construct the eigenstates $K_{S,L}$ with definite masses $m_{S,L}$ and lifetimes $\tau_{S,L}$. So far, we have made no assumptions about CPT symmetry. The following derivation can be expanded to allow for possible CPT-violating effects, but for simplicity, we assume CPT. This requires

$$\begin{aligned} \left\langle K^0 \middle| \mathcal{H}_{\text{eff}} \middle| K^0 \right\rangle &= \left\langle \overline{K^0} \middle| \mathcal{H}_{\text{eff}} \middle| \overline{K^0} \right\rangle \\ \left\langle K^0 \middle| \Gamma \middle| K^0 \right\rangle &= \left\langle \overline{K^0} \middle| \Gamma \middle| \overline{K^0} \right\rangle, \end{aligned}$$

where the second constraint arises from its connection to the physical decay amplitudes [22].

[1]One should note that although this appears to be the Schrödinger equation, it is in fact an effective equation of state determined from perturbation theory, which will be discussed in more detail in Section 1.2.

If we assume that the time-dependence of the kaon wavefunction has the form

$$\begin{pmatrix} a_1(t) \\ a_2(t) \end{pmatrix} = \begin{pmatrix} a_1(0) \\ a_2(0) \end{pmatrix} e^{-i\left(m-\frac{i}{2}\Gamma\right)t}, \tag{1.8}$$

and use Equation 1.5, we obtain

$$\mathcal{H}_{\text{eff}}\Big|K_{S,L}\Big\rangle = \left(m_{S,L} - \frac{i}{2}\Gamma_{S,L}\right)\Big|K_{S,L}\Big\rangle, \tag{1.9}$$

where we have now written the eigenstates and eigenvalues as the short and long-lived states. Furthermore, we assume that the short and long-lived states are close to the CP eigenstates K_1 and K_2, such that

$$\begin{aligned} \Big|K_S\Big\rangle &= \frac{1}{\sqrt{1+|\epsilon_K|^2}}\left(K_1 + \epsilon_K K_2\right) \\ &= \frac{1}{\sqrt{2\left(1+|\epsilon_K|^2\right)}}\left[(1+\epsilon_K)\Big|K^0\Big\rangle + (1-\epsilon_K)\Big|\overline{K^0}\Big\rangle\right], \qquad (1.10) \\ \Big|K_L\Big\rangle &= \frac{1}{\sqrt{1+|\epsilon_K|^2}}\left(K_2 + \epsilon_K K_1\right) \\ &= \frac{1}{\sqrt{2\left(1+|\epsilon_K|^2\right)}}\left[(1+\epsilon_K)\Big|K^0\Big\rangle - (1-\epsilon_K)\Big|\overline{K^0}\Big\rangle\right], \qquad (1.11) \end{aligned}$$

with ϵ_K appropriately small. Using this ansatz in Equation 1.9, we obtain for $K_{S,L}$

$$\begin{aligned} \mathcal{H}\left(\left[(1+\epsilon_K)\Big|K^0\Big\rangle \pm (1-\epsilon_K)\Big|\overline{K^0}\Big\rangle\right]\right) &= \left[(1+\epsilon_K)\Big|K^0\Big\rangle \pm (1-\epsilon_K)\Big|\overline{K^0}\Big\rangle\right] \\ &\quad \times \left(m_{S,L} - \frac{i}{2}\Gamma_{S,L}\right). \qquad (1.12) \end{aligned}$$

Using the orthogonality condition $\Big\langle K^0\Big|\overline{K^0}\Big\rangle = 0$ gives:

$$(1+\epsilon_K)\,\mathcal{H}_{22} \pm (1-\epsilon_K)\,\mathcal{H}_{12} = (1+\epsilon_K)\left(m_{S,L} - \frac{i}{2}\Gamma_{S,L}\right). \tag{1.13}$$

Similarly for $\left\langle \overline{K^0} \right|$:

$$(1+\epsilon_K)\,\mathcal{H}_{21} \pm (1-\epsilon_K)\,\mathcal{H}_{22} = \pm\,(1-\epsilon_K)\left(m_{S,L} - \frac{i}{2}\Gamma_{S,L}\right). \tag{1.14}$$

Taking Equation 1.13 minus 1.14 for the m_S solution gives

$$2\epsilon_K\left(m_S - \frac{i}{2}\Gamma_S\right) = (1+\epsilon_K)\,\mathcal{H}_{11} + (1-\epsilon_K)\,\mathcal{H}_{12} - (1+\epsilon_K)\,\mathcal{H}_{21} - (1-\epsilon_K)\,\mathcal{H}_{22}, \tag{1.15}$$

and the combination Equation 1.13 plus 1.14 for the m_L solution gives

$$2\epsilon_K\left(m_L - \frac{i}{2}\Gamma_L\right) = (1+\epsilon_K)\,\mathcal{H}_{11} - (1-\epsilon_K)\,\mathcal{H}_{12} + (1+\epsilon_K)\,\mathcal{H}_{21} - (1-\epsilon_K)\,\mathcal{H}_{22}. \tag{1.16}$$

Defining $\Delta m \equiv m_L - m_S$ and $\Delta\Gamma \equiv \Gamma_S - \Gamma_L$, such that both are positive, and assuming $\epsilon_K \ll 1$, we obtain from the difference of Equations 1.15 and 1.16

$$\begin{aligned} 2\epsilon_K\left[m_L - m_S + \frac{i}{2}\left(\Gamma_S - \Gamma_L\right)\right] &= (1+\epsilon_K)\,\mathcal{H}_{21} - (1-\epsilon_K)\,\mathcal{H}_{12} \\ \epsilon_K &= \frac{\left\langle \overline{K^0}\middle|\mathcal{H}_{\text{eff}}\middle|K^0\right\rangle - \left\langle K^0\middle|\mathcal{H}_{\text{eff}}\middle|\overline{K^0}\right\rangle}{2\Delta m + i\Delta\Gamma}. \end{aligned} \tag{1.17}$$

Using the Hermitian nature of the matrices $\mathbf{M}$ and Γ, we have

$$\mathbf{M}_{12} = \mathbf{M}_{21}^{*} \quad \text{and} \quad \Gamma_{12} = \Gamma_{21}^{*},$$

so we can simplify ϵ_K to

$$\begin{aligned} \epsilon_K &= \frac{\mathbf{M}_{21} + i/2\Gamma_{21} - \mathbf{M}_{12} - i/2\Gamma_{21}}{2\Delta m + i\Delta\Gamma} \\ &= \frac{\mathrm{Im}(\mathbf{M}_{12}) - (i/2)\mathrm{Im}(\Gamma_{12})}{i\Delta m - (1/2)\Delta\Gamma}. \end{aligned} \tag{1.18}$$

Finally, it is possible to show that $\mathrm{Im}(\Gamma_{12})$ is much smaller than $\mathrm{Im}(\mathbf{M}_{12})$ (see [27] and Section 1.2), so ϵ_K is said to characterize CP violation in the mass-matrix. Note that K_S and K_L are not orthogonal;

$$\left\langle K^0 \middle| \overline{K^0} \right\rangle \approx \frac{2\mathrm{Re}(\epsilon_K)}{1 + |\epsilon_K|^2} , \tag{1.19}$$

indicative of the fact that both states can decay to $\pi\pi$.

The previous discussion shows that the physical kaon states need not be CP eigenstates, which provides a nice explanation for the experimental fact that K_L can decay to $\pi\pi$. K_L is mostly K_2, with a small admixture of K_1, and it is the K_1 part of K_L that decays to $\pi\pi$. The size of ϵ_K sets the rate of $K_L \to \pi\pi$. The asymmetry in the composition of the K_L may also be thought of as arising from the fact that the transition rate for $\overline{K^0} \to K^0$ is faster than the rate $K^0 \to \overline{K^0}$, which is the clearest way to see that the mixing is CP-violating. It also means that both the K_L *and* the K_S contain more K^0 than $\overline{K^0}$.

The value of ϵ_K can be measured in a number of ways. The simplest is to measure the ratio of the K_L and K_S decay amplitudes to either $\pi\pi$ final state,

$$\eta_{+-} \equiv \frac{A\left(K_L \to \pi^+\pi^-\right)}{A\left(K_S \to \pi^+\pi^-\right)} \qquad \text{and} \qquad \eta_{00} \equiv \frac{A\left(K_L \to \pi^0\pi^0\right)}{A\left(K_S \to \pi^0\pi^0\right)}, \tag{1.20}$$

both of which should be equal to ϵ_K in the absence of other effects.

The magnitudes of η_{+-} and η_{00} can be independently measured from branching ratios and the K_S and K_L lifetimes, while the measurement of the complex phases ϕ_{+-} and ϕ_{00} requires interference between the K_S and K_L decay amplitudes. Experiments (such as KTeV) which collect both decays to $\pi^+\pi^-$ and $\pi^0\pi^0$ can measure the ratio $|\eta_{00}/\eta_{+-}|$ and the phase difference $\Delta\phi \equiv \phi_{00} - \phi_{+-}$ to high accuracy. From a fit to all available information as of 2000 [28],

$$\begin{aligned} |\eta_{+-}| &= (2.285 \pm 0.019) \times 10^{-3} & \qquad \phi_{+-} &= (43.5 \pm 0.6)^\circ , \\ |\eta_{00}| &= (2.275 \pm 0.019) \times 10^{-3} & \qquad \phi_{00} &= (43.4 \pm 1.0)^\circ . \end{aligned}$$

The errors on $|\eta_{+-}|$ and $|\eta_{00}|$ are correlated.

An independent way to extract ϵ_K is to measure the charge asymmetry in the semileptonic decays $K_L \to \pi^{\pm} e^{\mp} \nu_e$ ("K_{e3}") and $K_L \to \pi^{\pm} \mu^{\mp} \nu_\mu$ ("$K_{\mu 3}$")

$$\delta \equiv \frac{\Gamma\left(K_L \to \pi^- l^+ \nu\right) - \Gamma\left(K_L \to \pi^+ l^- \bar{\nu}\right)}{\Gamma\left(K_L \to \pi^- l^+ \nu\right) + \Gamma\left(K_L \to \pi^+ l^- \bar{\nu}\right)}. \tag{1.21}$$

These decay amplitudes probe the K^0 and $\overline{K^0}$ content of the K_L directly, since a l^+ can come only from a K^0 decay, and likewise a l^- from a $\overline{K^0}$, according to the "$\Delta S = \Delta Q$" rule. Recent results from the CPLEAR group at CERN constrain violation of the $\Delta S = \Delta Q$ rule to a very small level [29], so $\delta = 2\mathrm{Re}(\epsilon_K)$ to an excellent approximation. Averaging results [28] from K_{e3} and $K_{\mu 3}$ modes, $\delta = (3.27 \pm 0.12) \times 10^{-3}$. KTeV has announced a preliminary measurement [30] of $\delta = (3.322 \pm 0.074) \times 10^{-3}$. Using the phase ϕ_{+-} above, the new world average value of δ corresponds to $|\epsilon_K| = (2.278 \pm 0.05) \times 10^{-3}$, in excellent agreement with the measured values of $|\eta_{+-}|$ and $|\eta_{00}|$.

From these measurements, we conclude $|\epsilon_K|$ is about 2.28×10^{-3} and is well known. To explore the reason it is non-zero, we consider the origin of the elements that mix the K_1 and K_2 states, which will lead us to consider if CP violation can occur in the decay process itself.

1.2 Origin of the Off-Diagonal Elements

The equation of motion for the kaon system (Equation 1.5) is a representation of the Wigner-Wiesskopf formalism applied to a perturbative consideration of the weak interaction [21]. We can relate the off-diagonal elements of $\mathbf{M}$ and Γ to specific transition amplitudes:

$$\Gamma_{12} = 2\pi \sum_f \left\langle K^0 \middle| \mathcal{H}_{\mathrm{W}} \middle| f \right\rangle \left\langle f \middle| \mathcal{H}_{\mathrm{W}} \middle| \overline{K^0} \right\rangle \tag{1.22}$$

$$\mathbf{M}_{12} = \left\langle K^0 \middle| \mathcal{H}_{\mathrm{SW}} \middle| \overline{K^0} \right\rangle + \sum_v \int \frac{dE}{m_{K^0} - E_v} \left\langle K^0 \middle| \mathcal{H}_{\mathrm{W}} \middle| v \right\rangle \left\langle v \middle| \mathcal{H}_{\mathrm{W}} \middle| \overline{K^0} \right\rangle, \tag{1.23}$$

where $|f\rangle$ are physical final states and $|v\rangle$ are intermediate states. The $\mathcal{H}_{\mathrm{eff}}$ has been broken into two pieces; the standard weak interaction that changes the strangeness quantum number by one unit $\Delta S = 1$, and a possible "super-weak" interaction $\mathcal{H}_{\rm SW}$ that could change $\Delta S = 2$.

Note that if $\langle f|\mathcal{H}_{\rm W}|K^0\rangle$ and $\langle f|\mathcal{H}_{\rm W}|\overline{K^0}\rangle$ were to have the same phase, then $\langle K^0|\mathcal{H}_{\rm W}|f\rangle\langle f|\mathcal{H}_{\rm W}|\overline{K^0}\rangle$ would be a real number, in which case $\mathrm{Im}(\mathbf{M}_{12}) = \mathrm{Im}(\Gamma_{12}) = 0$ and $\epsilon_K = 0$ and we would not observe CP violation. Thus an imaginary part of $\mathbf{M}_{12}$ and Γ_{12} requires a *phase difference* in the coupling of $\mathcal{H}_{\rm W}$ to K^0 and $\overline{K^0}$. In the following section, we examine the effect of phase differences on the K_S and K_L decay amplitudes.

1.3 K Meson Decay Amplitudes

We wish to consider how the phase difference in the K^0 and $\overline{K^0}$ amplitudes affects the observed $K_{S,L}$ states. Since the final states are pions, we are forced to consider the isospin decomposition of both the initial and final states. The pion final states must be either $I = 0$ or $I = 2$ states.

First we define the decay amplitude for K^0 to a final state with isospin I as

$$\left\langle I \middle| \mathcal{H} \middle| K^0 \right\rangle = A_I e^{i\delta_I}, \tag{1.24}$$

where the additional phase factor $e^{i\delta_I}$ is due to the fact that the final-state pions will interact strongly with each other, so-called final-state interactions. The $\overline{K^0}$ amplitude is

$$\left\langle I \middle| \mathcal{H} \middle| \overline{K^0} \right\rangle = A_I^* e^{i\delta_I}, \tag{1.25}$$

where the interaction part A_I^* is conjugated due to CPT.

We use the definitions of $K_{S,L}$ in Equation 1.10 and 1.11. We consider the general case of η_f, where where $f = \pi^+\pi^-$ or $\pi^0\pi^0$. With a_f defined as the amplitude of $\left|K^0\right\rangle \to f$ and $a_{\bar{f}}$ the amplitude of $\left|\overline{K^0}\right\rangle \to f$, and using the decomposition of $K_{S,L}$

above, we obtain

$$\eta_f = \frac{(a_f - a_{\bar{f}}) + \epsilon_K(a_f + a_{\bar{f}})}{(a_f + a_{\bar{f}}) + \epsilon_K(a_f - a_{\bar{f}})}. \tag{1.26}$$

Let

$$\chi_f = \frac{a_f - a_{\bar{f}}}{a_f + a_{\bar{f}}}, \tag{1.27}$$

so that

$$\eta_f = \frac{\chi_f + \epsilon_K}{1 + \epsilon_K \chi_f} \approx \chi_f + \epsilon_K. \tag{1.28}$$

Now consider two final states $\pi^+\pi^-$ and $\pi^0\pi^0$, which will allow us to calculate η_{+-} and η_{00}. We must consider the isospin decompositions of the two final states, using the Clebsch-Gordan coefficients.

$$\begin{aligned} \Big|\pi^+\pi^-\Big\rangle &= \sqrt{\frac{2}{3}}\Big|I=0\Big\rangle + \sqrt{\frac{1}{3}}\Big|I=2\Big\rangle, \\ \Big|\pi^0\pi^0\Big\rangle &= -\sqrt{\frac{1}{3}}\Big|I=0\Big\rangle + \sqrt{\frac{2}{3}}\Big|I=2\Big\rangle. \end{aligned}$$

This allows us to calculate

$$\begin{aligned} \chi_{\pi^+\pi^-} &= \frac{\Big\langle\pi^+\pi^-\Big|\mathcal{H}\Big|K^0\Big\rangle - \Big\langle\pi^+\pi^-\Big|\mathcal{H}\Big|\overline{K^0}\Big\rangle}{\Big\langle\pi^+\pi^-\Big|\mathcal{H}\Big|K^0\Big\rangle - \Big\langle\pi^+\pi^-\Big|\mathcal{H}\Big|\overline{K^0}\Big\rangle} \\ &= \frac{\sqrt{\frac{2}{3}}\left(\Big\langle 0\Big|\mathcal{H}\Big|K^0\Big\rangle - \Big\langle 0\Big|\mathcal{H}\Big|\overline{K^0}\Big\rangle\right) + \sqrt{\frac{1}{3}}\left(\Big\langle 2\Big|\mathcal{H}\Big|K^0\Big\rangle - \Big\langle 2\Big|\mathcal{H}\Big|\overline{K^0}\Big\rangle\right)}{\sqrt{\frac{2}{3}}\left(\Big\langle 0\Big|\mathcal{H}\Big|K^0\Big\rangle + \Big\langle 0\Big|\mathcal{H}\Big|\overline{K^0}\Big\rangle\right) + \sqrt{\frac{1}{3}}\left(\Big\langle 2\Big|\mathcal{H}\Big|K^0\Big\rangle + \Big\langle 2\Big|\mathcal{H}\Big|\overline{K^0}\Big\rangle\right)}. \end{aligned}$$

We are free to choose the relative phase between our decay amplitudes. We follow

Wu and Yang [31], and choose A_0 real, $A_0 = A_0^*$.

$$\chi_{\pi^+\pi^-} = \frac{\sqrt{2}\left(e^{i\delta_0}A_0 - e^{i\delta_0}A_0^*\right) + \left(e^{i\delta_2}A_2 - e^{i\delta_2}A_2^*\right)}{\sqrt{2}(e^{i\delta_0}A_0 + e^{i\delta_0}A_0^*) + (e^{i\delta_2}A_2 + e^{i\delta_2}A_2^*)} \tag{1.29}$$

$$= \frac{2ie^{i\delta_2}\mathrm{Im}(A_2)}{2\sqrt{2}e^{i\delta_0}A_0 + 2e^{i\delta_2}\mathrm{Re}(A_2)} \tag{1.30}$$

$$= \frac{\epsilon'}{1+\omega/\sqrt{2}}, \tag{1.31}$$

where we use

$$\epsilon' \equiv \frac{i}{\sqrt{2}}\frac{\mathrm{Im}(A_2)}{A_0}\,e^{i(\delta_2-\delta_0)} \qquad \text{and} \qquad \omega \equiv \frac{\mathrm{Re}(A_2)}{A_0}\,e^{i(\delta_2-\delta_0)}\ . \tag{1.32}$$

We assume that ω, which represents the ratio of $\Delta I = 3/2$ transition amplitude to the $\Delta I = 1/2$ transition amplitude, is small. Experimentally $|\omega| \approx 1/22$; the dominance of the $\Delta I = 1/2$ transition is referred to as the "$\Delta I = 1/2$ rule." Hence

$$\eta_{+-} = \epsilon_K + \frac{\epsilon'}{1+\omega/\sqrt{2}}. \tag{1.33}$$

Using the same treatment for the neutral mode,

$$\chi_{\pi^0\pi^0} = \frac{\left\langle \pi^0\pi^0 \middle| \mathcal{H} \middle| K^0 \right\rangle - \left\langle \pi^0\pi^0 \middle| \mathcal{H} \middle| \overline{K^0} \right\rangle}{\left\langle \pi^0\pi^0 \middle| \mathcal{H} \middle| K^0 \right\rangle + \left\langle \pi^0\pi^0 \middle| \mathcal{H} \middle| \overline{K^0} \right\rangle}. \tag{1.34}$$

Using the isospin decomposition of $\left|2\pi^0\right\rangle = -\sqrt{\frac{1}{3}}\left|0\right\rangle + \sqrt{\frac{2}{3}}\left|2\right\rangle$, we obtain

$$\chi_{2\pi^0} = \frac{\sqrt{\frac{2}{3}}\left(\left\langle 2\middle|\mathcal{H}\middle|K^0\right\rangle - \left\langle 2\middle|\mathcal{H}\middle|\overline{K^0}\right\rangle\right) - \sqrt{\frac{1}{3}}\left(\left\langle 0\middle|\mathcal{H}\middle|K^0\right\rangle - \left\langle 0\middle|\mathcal{H}\middle|\overline{K^0}\right\rangle\right)}{\sqrt{\frac{2}{3}}\left(\left\langle 2\middle|\mathcal{H}\middle|K^0\right\rangle + \left\langle 2\middle|\mathcal{H}\middle|\overline{K^0}\right\rangle\right) - \sqrt{\frac{1}{3}}\left(\left\langle 0\middle|\mathcal{H}\middle|K^0\right\rangle + \left\langle 0\middle|\mathcal{H}\middle|\overline{K^0}\right\rangle\right)},$$

$$= -\frac{i\sqrt{2}e^{i(\delta_2-\delta_0)}\mathrm{Im}(A_2)}{A_0 - \sqrt{2}e^{i(\delta_2-\delta_0)}\mathrm{Re}(A_2)},$$

$$= -\frac{2\epsilon'}{1-\sqrt{2}\omega},$$

such that

$$\eta_{00} = \epsilon_K - 2\frac{\epsilon'}{1-\sqrt{2}\omega}. \tag{1.35}$$

The phase shift $\delta_2 - \delta_0$ is determined from other data [32] to be $(-42 \pm 4)^\circ$, so that ϵ' has a phase of $\frac{\pi}{2} + \delta_2 - \delta_0 = (48 \pm 4)^\circ$, which is essentially the same phase as ϵ_K. Conveniently, this maximizes the effect of ϵ' on $K \to \pi\pi$ events.

At this point, we have shown a number of interesting facts. First, η_{+-} and η_{00} can be different, and the difference is parametrized by ϵ'. This difference is dependent on a phase difference between A_0 and A_2 (or $\mathrm{Im}(A_2) \neq 0$ in our phase convention). Second, we can show

$$\begin{aligned} \left\langle \pi^+\pi^- \middle| \mathcal{H}_{\mathrm{W}} \middle| K_2 \right\rangle &= \sqrt{\frac{2}{3}} i e^{i\delta_2} \mathrm{Im}(A_2)\,, \\ &= A_0 \frac{2}{\sqrt{3}} \frac{e^{i\delta_2}}{e^{i(\delta_2-\delta_0)}} \epsilon', \\ &\propto \epsilon', \end{aligned} \tag{1.36}$$

so K_2 ($CP = -1$) can decay to a $\pi\pi$ final state ($CP = +1$). Hence ϵ' describes CP violation in the decay process itself, or "direct" CP violation. This is shown diagrammatically in Figure 1.1.

It is interesting to consider the relationship between $\mathrm{Im}(A_2)$ and $\mathrm{Im}(\Gamma_{12})$ at

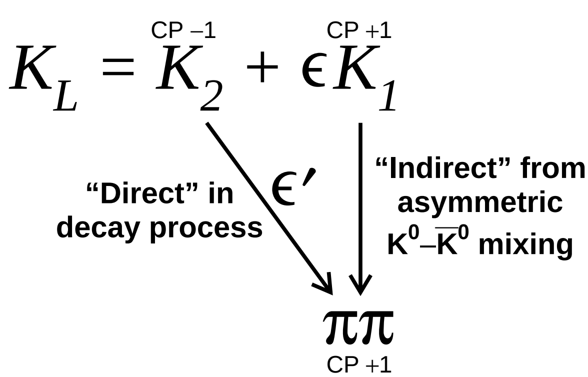

Figure 1.1: The two ways that CP violation can lead to $K_L \to \pi\pi$.

this point. Sachs [22] points out that $\Gamma_{12} = \left\langle K^0 \middle| \Gamma \middle| \overline{K^0} \right\rangle = \sum_f a_f^* a_f$, where a_f is dominated by 2π final states by a factor of ~ 580. We have defined by the phase convention A_0 to be real, so $\mathrm{Im}(\Gamma_{12}) = -2\mathrm{Re}(A_2)\,\mathrm{Im}(A_2)$, and CP violation in the decay matrix is characterized by ϵ'.

Equations 1.33 and 1.35 show us that direct CP violation contributes differently to $K_L \to \pi^+\pi^-$ and $K_L \to \pi^0\pi^0$, so we can measure ϵ' by comparing decay rates. Keeping terms to order $\omega\epsilon'/\epsilon$,

$$\frac{\Gamma\left(K_L \to \pi^+\pi^-\right)/\Gamma\left(K_S \to \pi^+\pi^-\right)}{\Gamma\left(K_L \to \pi^0\pi^0\right)/\Gamma\left(K_S \to \pi^0\pi^0\right)} = \frac{\left|\eta_{+-}\right|^2}{\left|\eta_{00}\right|^2} \approx 1 + 6\left[1 + \sqrt{2}\mathrm{Re}\left(\omega\right)\right]\mathrm{Re}(\epsilon'/\epsilon). \tag{1.37}$$

Using the value of $\mathrm{Re}(\omega) \approx 0.033$, this becomes $1 + (6.28)\,\mathrm{Re}(\epsilon'/\epsilon)$. Traditionally, we ignore ω, and use an *effective* definition of $\mathrm{Re}(\epsilon'/\epsilon)$

$$\mathrm{Re}(\epsilon'/\epsilon) \approx \frac{1}{6}\left[\frac{\Gamma\left(K_L \to \pi^+\pi^-\right)/\Gamma\left(K_S \to \pi^+\pi^-\right)}{\Gamma\left(K_L \to \pi^0\pi^0\right)/\Gamma\left(K_S \to \pi^0\pi^0\right)} - 1\right]. \tag{1.38}$$

We should note that this discussion can be undertaken without assuming the Wu-Yang phase convention. In such a case, we still have $\eta_{+-} \approx \epsilon + \epsilon'$ and $\eta_{00} \approx \epsilon - 2\epsilon'$, however $\epsilon \neq \epsilon_K$ and is not due solely to mixing. However, ϵ' remains a *difference* in the terms $\mathrm{Im}(A_0)$ and $\mathrm{Im}(A_2)$ and is independent of phase convention. The observable $\mathrm{Re}(\epsilon'/\epsilon)$ is independent of phase convention and $\mathrm{Re}(\epsilon'/\epsilon) \neq 0$ is a signal for CP violation in the decay process.

1.4 The Phase of η

It is worthwhile returning to Equation 1.18 and considering the phase of ϵ_K.

$$\begin{aligned}\epsilon_K &= -\frac{\mathrm{Im}(\mathbf{M}_{12}) - (i/2)\mathrm{Im}(\Gamma_{12})}{(1/2)\Delta\Gamma - i\Delta m} \\ &= -\frac{e^{i\phi_{SW}}}{\sqrt{2}\kappa\Delta m}\left[\mathrm{Im}(\mathbf{M}_{12}) - \frac{i}{2}\mathrm{Im}(\Gamma_{12})\right],\end{aligned} \tag{1.39}$$

where

$$\phi_{SW} \equiv \tan^{-1}\left(\frac{2\Delta m}{1/\tau_S - 1/\tau_L}\right) \tag{1.40}$$

and

$$\kappa \equiv \sqrt{\frac{1}{2}\left[1 + \left(\frac{\Delta\Gamma}{2\Delta m}\right)\right]} \simeq 1. \tag{1.41}$$

Thus we see that $\phi_{\epsilon_K} = \phi_{SW}$ as long as we have $\mathrm{Im}(\Gamma_{12})$ small. There are a number of contributions to $\mathrm{Im}(\Gamma_{12})$ as detailed in Reference [33]. If we are willing to assume the $\Delta S = \Delta Q$ rule and assume that $\eta_{3\pi} \simeq \epsilon$ in CP violation in $K \to 3\pi$ decays, then $\phi_{\epsilon_K} \approx \phi_{SW}$.

The next considerations are the phases of ϕ_{+-} and ϕ_{00}. There is the additional contribution to η_{+-} and η_{00} from ϵ', as shown in Equations 1.33 and 1.35. However, we know that $\phi_{\epsilon'}$ is roughly the same as ϕ_{ϵ_K} (as discussed on page 14), and that $\mathrm{Re}(\epsilon'/\epsilon)$ is small. Hence the deviation caused by ϵ' is of order $0.002°$. Note that $\Delta\phi = -3\mathrm{Im}(\epsilon'/\epsilon)$ because ϕ_{SW} cancels. The relation between these numbers is shown in Figure 1.2.

We started with the value of ϵ_K which was derived assuming CPT. Allowing for CPT-violating effects causes ϕ_{+-} and ϕ_{00} to deviate from ϕ_{SW} in different ways. Any deviation from $\phi_{+-} = \phi_{SW}$ at a level greater than expected from ϵ' is a CPT-violating signal. In addition, $\Delta\phi > -3\mathrm{Im}(\epsilon'/\epsilon)$ is also a CPT-violating signal.

1.5 Theoretical Predictions for Re(ϵ'/ϵ)

Almost immediately after the discovery of indirect CP violation, L. Wolfenstein proposed [34] that there might exist a new interaction that was so much weaker than the weak interaction that it was only observable in the K system. The "superweak" interaction is a direct $\Delta S = 2$ interaction, and could account for indirect mixing effects, but since it does not enter in the decay matrix (see Equations 1.22 and 1.23), it predicts $\mathrm{Re}(\epsilon'/\epsilon) = 0$.

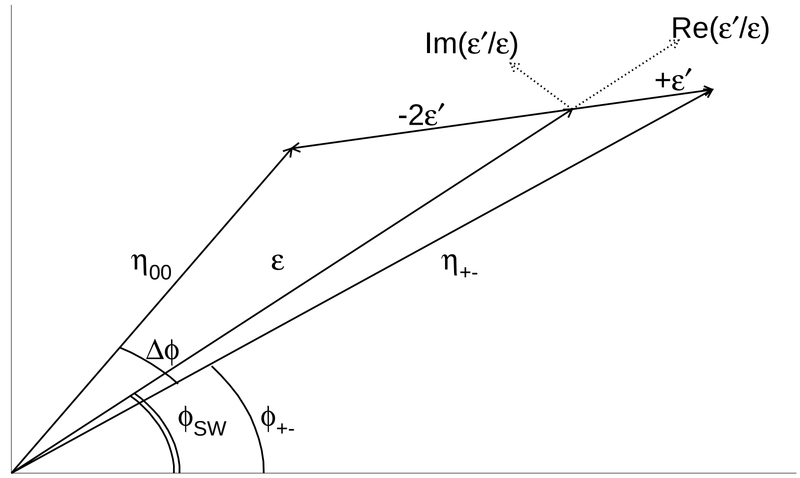

Figure 1.2: The Wu-Yang diagram of the relation between the complex numbers ϵ, ϵ', η_{+-}, and η_{00}, and the phases ϕ_{SW}, ϕ_{+-} and $\Delta\phi$. The phase and magnitude of ϵ' have been greatly exaggerated for the sake of clarity.

By 1974, the particle physics community had come to the picture of "The Standard Model," based on the theoretical work of Glashow, Weinberg, Salam, Gell-Mann, Iliopoulos and others, and much experimental work. The Standard Model includes the strong interaction between quarks, the unified electro-weak interaction mediated by intermediate vector bosons such as the $W^{\pm}$, and the generational model of quarks and leptons.

In the Standard Model, CP violation can arise through the interactions of the quarks with the $W^{\pm}$ gauge bosons via a complex coupling constant. There are many reviews [35, 36] which cover these calculations, and we will only touch on the basics. The weak interaction is not diagonal in the quark flavor basis. The coupling between up-type (u, c, t) and down-type (d, s, b) quarks is given by the 3×3 unitary Cabibbo-Kobayashi-Maskawa (CKM) matrix [37],

$$V = \begin{pmatrix} v_{ud} & v_{us} & v_{ub} \\ v_{cd} & v_{cs} & v_{cb} \\ v_{td} & v_{ts} & v_{tb} \end{pmatrix} \approx \begin{pmatrix} 1-\lambda^2/2 & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1-\lambda^2/2 & A\lambda^2 \\ A\lambda^3(1-\rho-i\eta) & -A\lambda^2 & 1 \end{pmatrix}, \tag{1.42}$$

where the parametrization of the matrix has $A = \mathcal{O}(1)$, $\lambda = \sin\theta_C$, the sine of the Cabibbo angle, and the parameter η sets the level of CP violation in the Standard Model. This parametrization is due to Wolfenstein [38].

While V has in principle 18 real parameters, the unitarity condition and the relative freedom of defining the u and d quark phases reduce the number of free parameters to four. Typically, those parameters are taken as three angles and a complex phase. It is this phase that drives all CP violation within the Standard Model. The diagonal elements of the matrix are close to unity, and the elements become progressively smaller farther from the diagonal. The Wolfenstein parametrization has v_{td} and v_{ub} complex, with all other elements real, to the order in λ shown. One of the six unitarity constraints implies

$$v_{ud}v_{ub}^* + v_{cd}v_{cb}^* + v_{td}v_{vd}^* = 0. \tag{1.43}$$

In the above phase convention, we have $v_{ud} \approx v_{tb} \approx 1$ and $v_{cd}v_{cb}^* = A\sin^3\theta_C$. This unitarity constraint can be represented as a triangle, with its peak at (ρ, η) in the complex plane, as shown in Figure 1.3.

It is interesting to note that at the time of the discovery of CP violation in $K \to \pi\pi$ in 1964, the b and t quarks had not been discovered, and the CKM matrix

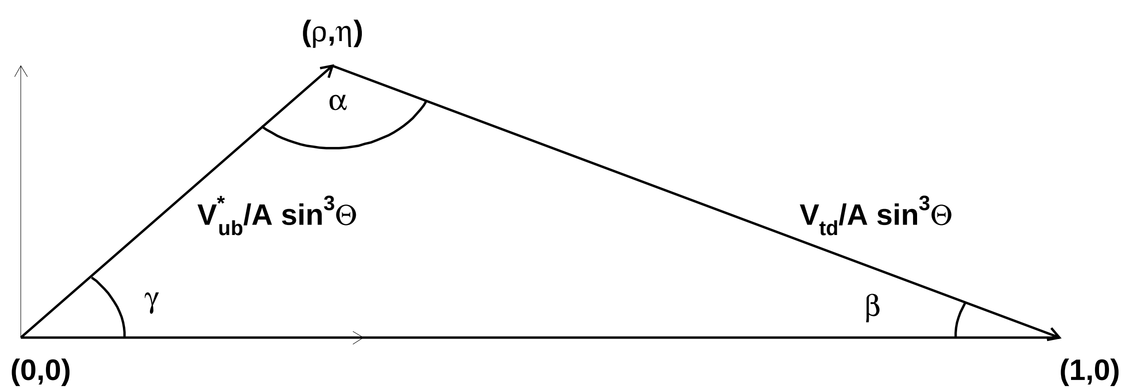

Figure 1.3: The unitarity triangle for the CKM matrix. CP-violating effects in the Standard model are related to the size of the imaginary number η. The sides of the triangle measure elements of the CKM matrix.

was written as the 2×2 Cabibbo mixing matrix. Kobayashi and Maskawa pointed out that the Cabibbo mixing matrix did not have the freedom of a complex phase which could explain the observed CP violation, and that at a minimum, three flavor generations were needed. Hence CP violation was the first sign for the existence of the b and t quarks.

Both indirect and direct CP violation are predicted by a non-zero phase in the CKM matrix. Indirect CP violation arises through "box diagrams" as shown in Figure 1.4. In the box diagrams, there are couplings with v_{td} which give rise to the indirect CP-violating effects in $K^0 \leftrightarrow \overline{K^0}$.

The calculation of ϵ_K from the values of the CKM matrix, or conversely extracting v_{td} from the measured value of ϵ_K, is not trivial. We will not detail these calculations, but refer the reader to References [35, 36]. The problems with these calculations all involve moving from the weak interaction at the quark level (the "short distance" physics) to the physics of mesons, which involves complicated QCD physics (the "long distance" physics), especially at the low energy of the kaon mass. However, one can extract a hyperbolic constraint on the position of the (ρ, η) peak of the unitarity triangle [39]:

$$\begin{aligned} |\epsilon_K| \;&=\; \frac{G_F^2 f_K^2 m_K m_W^2}{6\sqrt{2}\pi^2 \Delta m} B_K \left(A^2 \lambda^6 \eta\right) \\ &\quad \times \left\{ y_c \left[\eta_{ct} f_3(y_c, y_t) - \eta_{cc}\right] + \eta_{tt} y_t f_2(y_t) A^2 \lambda^4 \left(1 - \rho\right) \right\}, \end{aligned} \tag{1.44}$$

where G_F is the Fermi constant, f_K is the kaon decay constant, m_K is the kaon mass, m_W is the W mass, B_K the "bag" factor that parametrizes non-perturbative

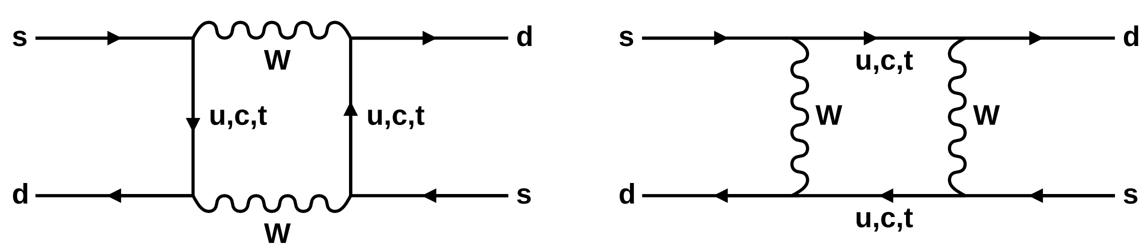

Figure 1.4: The box diagrams that provide the dominant contribution to CP violation in $K^0 - \overline{K^0}$ mixing.

QCD effects, η_{cc}, η_{ct} and η_{tt} are QCD correction constants, $y_t = m_t^2/m_W^2$ and $y_c = m_c^2/m_W^2$, and $f_{2,3}$ are QCD correction functions. Typically the bounds on ρ and η are determined by scanning through the ranges of experimental and theoretical errors.

The complex nature of the CKM matrix also gives rise to direct CP violation within the Standard Model. If we return to our definition of ϵ' in Equation 1.32, but remove the Wu-Yang phase convention ($\mathrm{Im}\,(A_0) = 0$), we have

$$\epsilon' = \frac{ie^{i(\delta_2-\delta_0)}}{\sqrt{2}} \frac{\mathrm{Re}(A_2)}{\mathrm{Re}(A_0)} \left[\frac{\mathrm{Im}(A_2)}{\mathrm{Re}(A_2)} - \frac{\mathrm{Im}(A_0)}{\mathrm{Re}(A_0)} \right]. \tag{1.45}$$

The two terms ($\mathrm{Im}(A_0)$ and $\mathrm{Im}(A_2)$) are due to penguin diagrams in the Standard Model.

The $\mathrm{Im}(A_0)$ term is caused by strong penguin diagrams (see Figure 1.5) and is expected to dominate. However, for certain values of the top quark mass, the electroweak penguin term ($\mathrm{Im}(A_2)$) can become large and cancel the effects of the strong penguin. For certain combinations of masses and Standard Model parameters, $\mathrm{Re}(\epsilon'/\epsilon)$ can be zero, even if CP violation is due to the phase of the CKM matrix.

Much effort has gone into calculating $\mathrm{Re}(\epsilon'/\epsilon)$ in recent years. Again, the diffi-

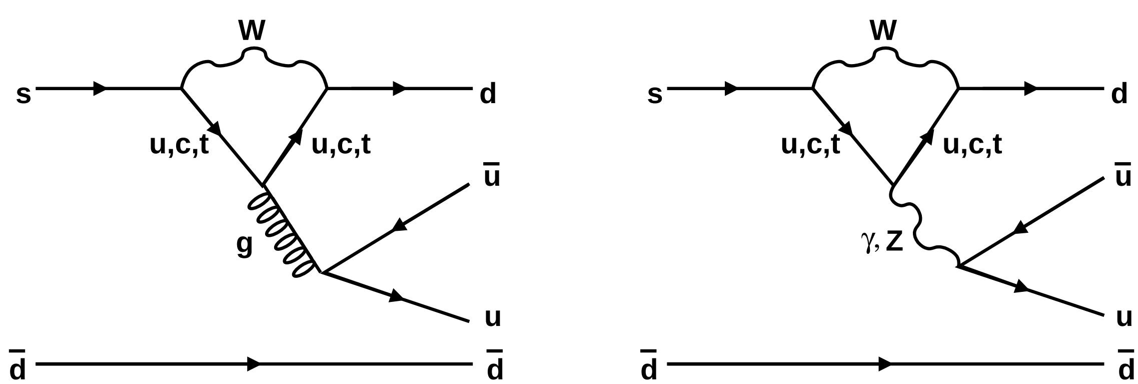

Figure 1.5: Examples of the strong (left) and electroweak (right) penguin diagrams that give rise to ϵ' in the Standard Model.

culty is moving from the quark level theory to an effective theory of mesons at low energies. Table 1.1 lists some recent predictions.

1.6 Other Manifestations of CP Violation

Kaon mixing and decays to $\pi\pi$ are not the only areas in which CP violation is expected. The kaon sector can have other manifestations of CP violation, and the B system is expected to show large CP-violating asymmetries.

Certain rare kaon decays are predicted to have CP-violating effects in addition to $K \to \pi\pi$. For example, the decays $K_L \to \pi^0 l^+ l^-$ occur as penguin diagrams at leading order in the Standard Model. (Replace the $u\bar{u}$ pair with l^+l^- in the electroweak penguin diagram in Figure 1.5 to see this.) These processes are also theoretically much cleaner than ϵ_K or ϵ' since the QCD physics is simplified by the presence of the lepton pair. The decay $K_L \to \pi^0 \nu\bar{\nu}$ and the isospin-conjugate decays $K^\pm \to \pi^\pm \nu\bar{\nu}$ are extremely clean and very attractive places to look for CP violation. Unfortunately, these rates are also extremely rare, at the level of a few $\times 10^{-11}$. One $K^+ \to \pi^+ \nu\bar{\nu}$ event has been observed by the E787 experiment at Brookhaven National Laboratory [49]. $K_L \to \pi^0 \nu\bar{\nu}$ has never been observed, and seems out of experimental reach for the foreseeable future.

Table 1.1: Some recent calculations of $\mathrm{Re}(\epsilon'/\epsilon)$ within the Standard Model.

Method	$\mathrm{Re}(\epsilon'/\epsilon)$ $(\times 10^{-4})$	Reference
Phenomenological	8.5 ± 5.9	Buras [35]
$1/N_C$ expansion	17 ± 9	Pallante, Pich and Scimeni [40]
Chiral Quark Model	11–36	Wu [41]
Phenomenological	$< 22 \pm 9$	Narison [42]
Chiral Quark Model	34 ± 18	Bijnens and Prades [43]
Lattice Gauge	4 ± 7	Ciuchini *et al.* [44]
Chiral Quark Model	< 3.3	Bel'kov *et al.* [45]
$1/N_C$ expansion	7–24.7	Hambye *et al.* [46]
Lattice Gauge	4.6 ± 3.0	Ciuchini [47]
Chiral Quark Model	17^{+14}_{-10}	Bertolini *et al.* [48]

If the CKM matrix is the cause of CP violation, such effects should also appear in the B-meson system. The observed CP-violating effects in the B system will appear to be larger due to the smaller resonance width of the B system. Because of the large number of final states available into which the B can decay, there are no analogous short- and long-lived states in the B system. Instead, CP violation is discussed in terms of asymmetries of decays of B^0 and $\overline{B^0}$ to final states that are CP eigenstates, such as $J/\Psi K_S$. The $J/\Psi K_S$ final state is the cleanest theoretically of the CP eigenstates, and is the first mode where CP asymmetries have been measured in the B system. The discussion of CP violation in the B system is almost always in terms of the unitarity triangle of the CKM matrix. For example, the B asymmetry in decays to $J/\Psi K_S$ measures $\sin(2\beta)$, where β is one of the three angles shown in Figure 1.3.

The B sector is predicted to be such a robust region to look for CP-violating effects that an entire industry has arisen around it. Measurements of CP violation in the B system are possible at the multipurpose $p\bar{p}$ colliders; for example, the CDF collaboration reported [50] a non-zero value for $\sin(2\beta) = 0.79^{+0.41}_{-0.44}$ at the roughly two standard deviation level. However, to get the statistical precision necessary to accurately measure these quantities, dedicated facilities known as "B-factories" were built. Two factories are currently running: the BaBar facility at the Stanford Linear Accelerator Laboratory, and BELLE at the National Laboratory for High Energy Physics (KEK) in Japan. Both groups have reported precise values for $\sin(2\beta)$, BaBar [51] reporting $\sin(2\beta) = 0.59 \pm 0.14$ (stat) ± 0.04 (syst) and Belle [52] reporting $\sin(2\beta) = 0.99 \pm 0.14$ (stat) ± 0.06 (syst), constraining this parameter.

As is the case with the kaon system, measuring only a single asymmetry cannot distinguish between mixing and direct CP violation. To establish direct CP violation in the B sector, the B-factories will need to measure an asymmetry in decays to another final state, such as $B \to \pi\pi$. These additional decay asymmetries are more difficult to measure experimentally than $B \to J/\Psi K_S$, and are fraught with the same types of theoretical issues that make estimations of $\mathrm{Re}(\epsilon'/\epsilon)$ in $K \to \pi\pi$ so difficult. The B-factories are entering an exciting future, but K physics and $\mathrm{Re}(\epsilon'/\epsilon)$

are still a vibrant region to measure CP violation, and offer unique information about CP violation.

1.7 Past Measurements of Re(ϵ'/ϵ)

The search for direct CP violation through the measurement of Re(ϵ'/ϵ), has been ongoing for 35 years, ever since the discovery of CP violation in the K sector. The most recent measurements come from two groups: E731/KTeV at FNAL and NA31/NA48 at the European Center for Nuclear Research (CERN). In 1993, just after E731 and NA31 had published their results, there was some confusion as to the true value of Re(ϵ'/ϵ); the FNAL group favored a lower number than the CERN group. This disagreement lead to a second set of experiments that have just recently finished running. Unfortunately, the published results of the two groups still show a large scatter in the value of Re(ϵ'/ϵ). The history of published results is shown in Figure 1.6. The work that will be detailed in this thesis will further constrain the world average value of Re(ϵ'/ϵ).

1.8 Measurements of Kaon Sector Parameters

The most recent measurements of the kaon sector parameters Δm and τ_S are from the FNAL experiments E731 [57] and E773 [58]; the NA31 [59] experiment at CERN also measured τ_S. The published results for Δm and τ_S are shown in Figures 1.7 and 1.8.

The best measurements of ϕ_{+-} are from the CPLEAR [60] experiment at CERN, using the K^0–$\overline{K^0}$ asymmetry, and the FNAL experiments E731 [57] and E773 [58], using the regeneration technique (see Section 3.2.1), which provides a coherent superposition of K_S and K_L states from which these parameters can be measured. Since these results depend in detail on the values assumed for Δm and τ_S and the fitting technique, we refer the readers to the papers for the reported values.

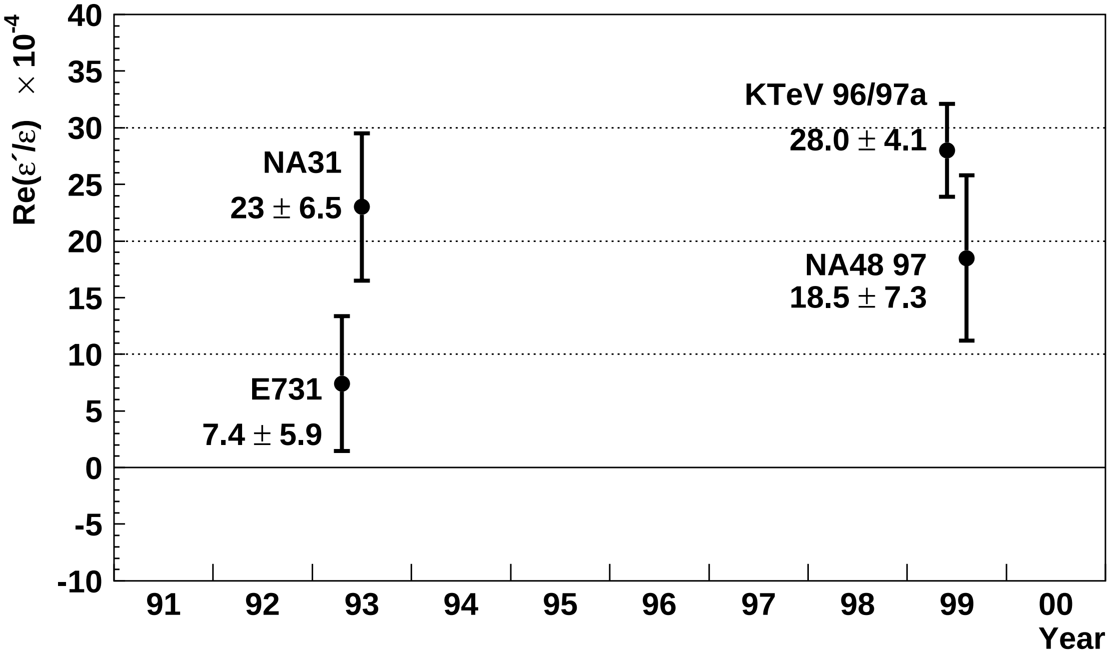

Figure 1.6: The published values for $\mathrm{Re}(\epsilon'/\epsilon)$ up to this thesis. The references are [53–56].

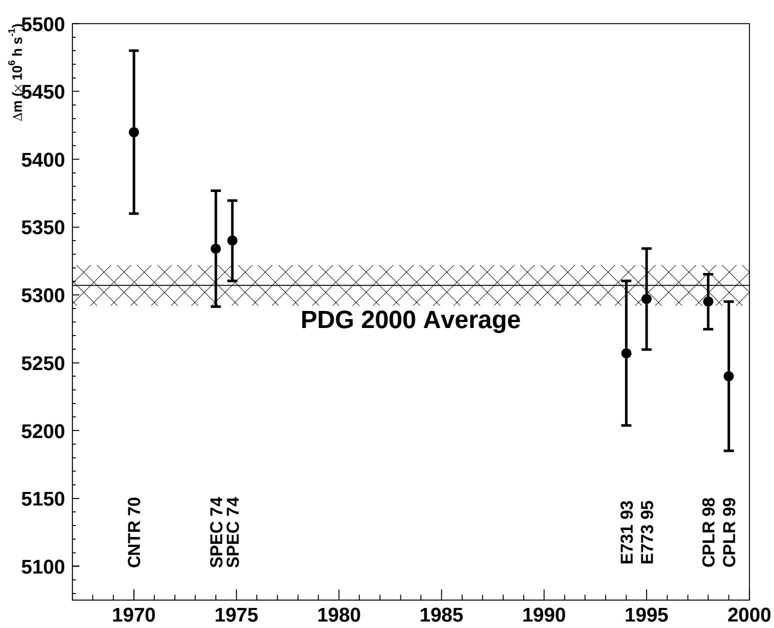

Figure 1.7: The published values for Δm up to this thesis.

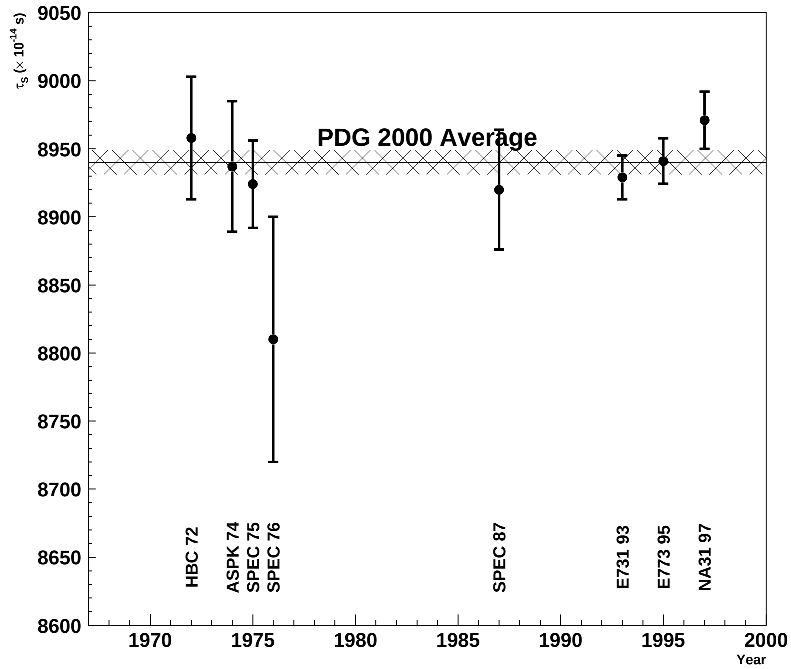

Figure 1.8: The published values for τ_S up to this thesis.

1.9 Overview of this Thesis

In the previous sections, we have discussed many of the historical and phenomenological issues relevant to the measurements we can perform in KTeV. In this thesis, we will present new measurements of $\mathrm{Re}(\epsilon'/\epsilon)$ and the kaon sector parameters Δm, τ_S and ϕ_{+-}. The remaining chapters will describe the measurement technique, with a specific emphasis on $K \to \pi^+\pi^-$. We will make passing reference to $K \to \pi^0\pi^0$, but the reader is referred to the dissertation of Valmiki Prasad [1] for the intricacies of that analysis.

In the next chapter, we provide an overview of the experimental technique, commenting on the strengths and weaknesses of our measurement, and comparing this experiment to others that have measured $\mathrm{Re}(\epsilon'/\epsilon)$. We continue in Chapters 3 and 4 with detailed descriptions of the apparatus and the data we collected. After describing the data samples, we discuss the calibration of the charged-particle spectrometer in Chapter 5, and follow that with a discussion of the analysis technique used to identify $K \to \pi^+\pi^-$ events and reject background events in Chapter 6. Chapter 7 gives an overview of the same issues for the $K \to \pi^0\pi^0$ analysis.

As we will see, our measurement technique requires a detailed simulation of the detector to extract the correct value for $\mathrm{Re}(\epsilon'/\epsilon)$ and other kaon sector parameters. Chapter 8 will detail the Monte Carlo simulation, with a special focus on the parts of the simulation related to $K \to \pi^+\pi^-$. The fitting algorithm used to extract our measurements from the data will be described in Chapter 9, and we will present our results for $\mathrm{Re}(\epsilon'/\epsilon)$, Δm, τ_S and ϕ_{+-}.

The final part of the analysis (Chapter 10) is estimating the systematic uncertainties on each of our measurements. With our measurements and uncertainties in hand, we draw some conclusions in Chapter 11.

CHAPTER 2
THE EXPERIMENTAL TECHNIQUE

The KTeV experiment was designed [61] with the goal of measuring $\mathrm{Re}(\epsilon'/\epsilon)$ at the level of $\sim 10^{-4}$, precise enough to measure a non-zero value even at the low end of theoretical predictions. The experiment must collect enough $K \to \pi\pi$ decays to achieve this statistical precision. More importantly, the design of the experiment must minimize possible biases. KTeV is based on the previous $\mathrm{Re}(\epsilon'/\epsilon)$ experiment E731 [62, 63], but with a new beamline and greatly improved detector. We will comment on the basic design of the experiment below, and list the details in Chapter 3. Chapter 2 concludes with some commentary on other techniques to measure $\mathrm{Re}(\epsilon'/\epsilon)$.

2.1 Overview of the KTeV Apparatus

The neutral kaons necessary for measuring $\mathrm{Re}(\epsilon'/\epsilon)$ are produced by high-energy protons hitting a fixed target. The 800 GeV protons that are extracted from the Tevatron hit a beryllium oxide target, and produce a spectrum of particles, some of which are kaons. A series of sweeping magnets and collimators are used to make two side-by-side beams of neutral particles. These beams travel 90 m from the target before entering the KTeV decay region, allowing short-lived particles such as K_S and hyperons to decay away.

The kaons enter an evacuated decay volume, shown in Figure 2.1. The region extends from 90 m to a large Kevlar-Mylar window 159 m downstream of the target. One of the K_L beams passes through a regenerator, which converts the K_L beam to a coherent superposition of K_L and K_S. The beam with the regenerator is eponymously named. The K_L beam is called the "vacuum beam" because there is nothing in it. This "double-beam" technique allows us to collect K_L and K_S decays simultaneously, which is the most important feature of the KTeV method. Collecting K_L and K_S decays at the same time ensures that our measurement of the ratio of the decay rates is insensitive to the inevitable time variations of beam intensity and trigger/detector deadtime, and has the added benefit that detector inefficiencies and the effects of additional activity in the detector nearly cancel as well. The regenerator alternates between the left and right beams between Tevatron extraction cycles (roughly once per minute) to minimize the effect of any left/right beam or detector asymmetry.

The KTeV detector, located just downstream of the vacuum window at 159 m, has a geometry optimized to reconstruct the decay products. The average laboratory

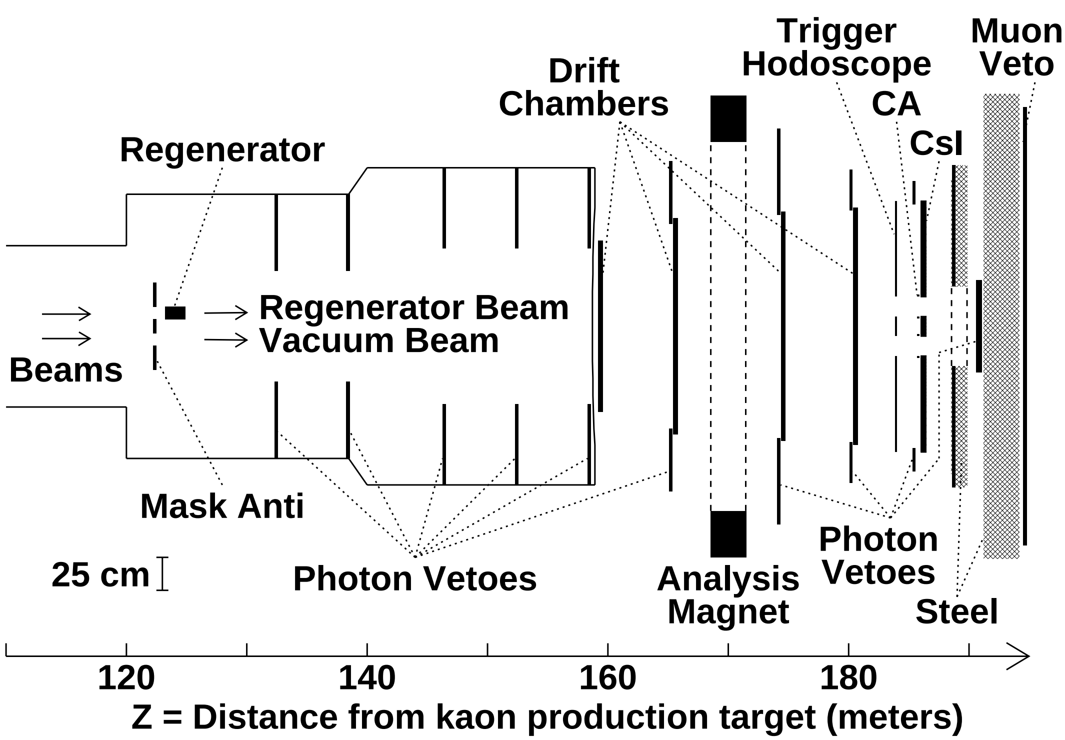

Figure 2.1: Plan view of the KTeV detector as configured to measure Re(ϵ'/ϵ). Figure courtesy of P. Shawhan [2].

energy of a kaon is ~ 70 GeV. The main components of the detector are a charged-particle spectrometer and an electromagnetic (EM) calorimeter.

The KTeV spectrometer has four square drift chambers, two upstream of a dipole magnet and two downstream. The drift chambers have horizontal and vertical sense wires to measure the positions of charged particles passing through them. The full trajectory of a particle is determined by matching segments in the two upstream chambers to segments in the two downstream chambers at the magnet. A particle's momentum is determined from the bend angle in the magnet. The volumes between the drift chambers are filled with helium to reduce scattering of the charged particles.

An electromagnetic calorimeter is located downstream of the spectrometer. It uses pure cesium iodide crystals (CsI) as an interaction and scintillation medium to measure particle energies and positions. The CsI calorimeter was designed to have excellent position and energy resolution for photons from $K \to \pi^0\pi^0$ decays (where each π^0 decays immediately to $\gamma\gamma$). The excellent resolution is necessary because the calorimeter is the only detector in which the $K \to \pi^0\pi^0$ decay products are reconstructed. It is also used to identify charged pions and electrons based on energy deposition.

Due to the finite area of the calorimeter and the drift chambers, the coverage of the detector is not 100%. To catch photons that would miss the CsI, we use a series of photon veto counters arranged in z along the volume of the KTeV detector. The primary use of these veto counters is to catch photons from $K_L \to 3\pi^0$ decays that would miss the CsI. If we were to miss two photons, the $K_L \to 3\pi^0$ could be mis-reconstructed as a $K \to \pi^0\pi^0$ decay, and would be a large source of background. The layout of these detectors can be seen in Figure 2.2, where five are embedded within the vacuum tank, and four are externally mounted on the spectrometer. In addition to the large "photon vetos" on the outer edges of the detector, there are a number of inner photon detectors that define our acceptance. The Mask Anti (MA) sits just upstream of the regenerator and defines the beginning of the decay volume. The Collar Anti (CA), which sits around the beam holes at the CsI, sharply defines the active area around the beam holes at the calorimeter.

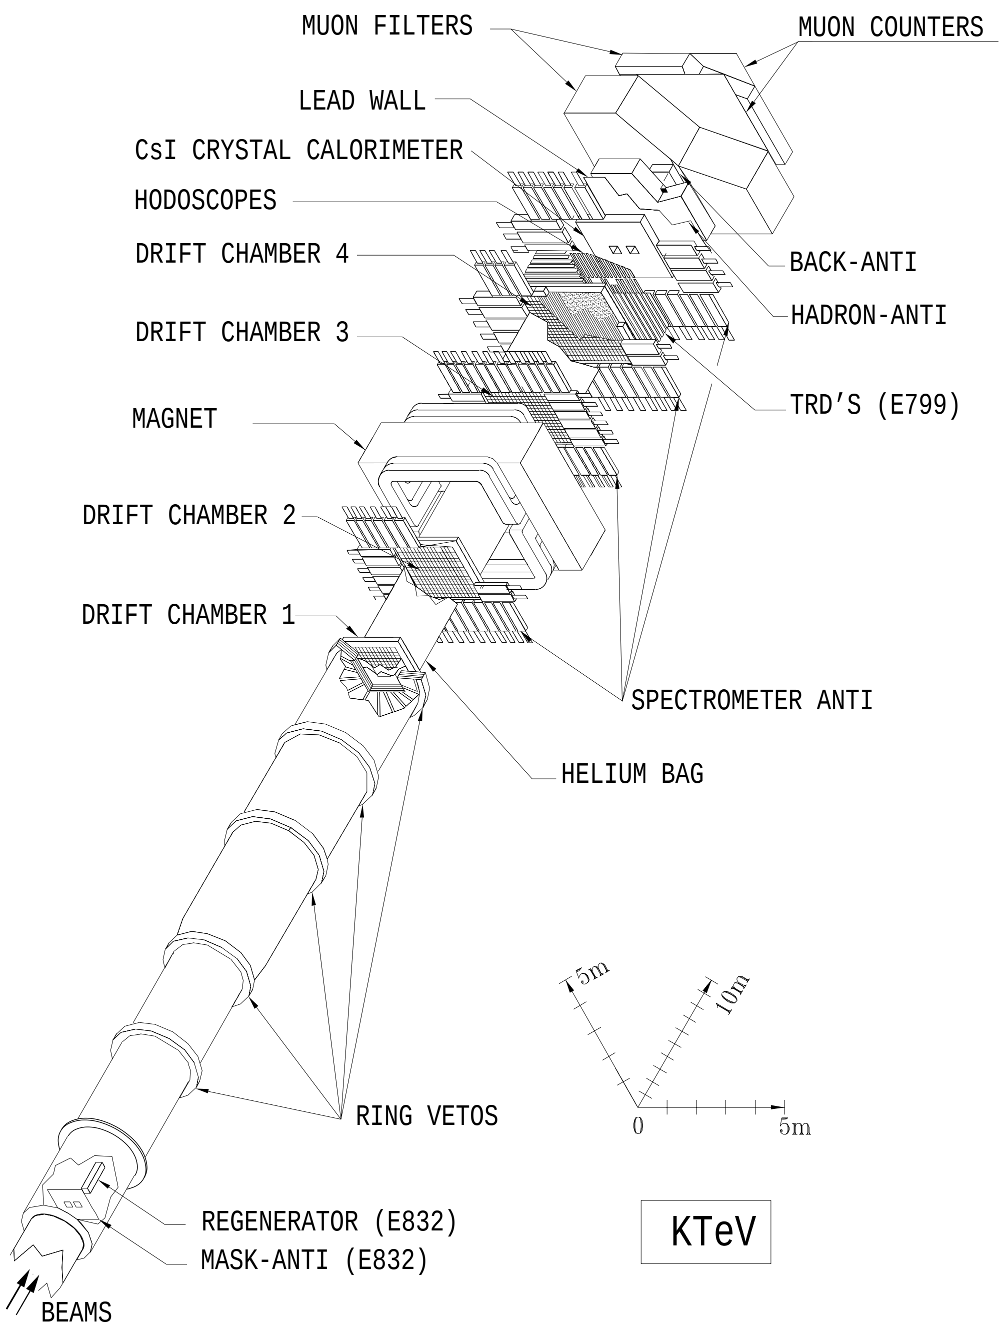

Figure 2.2: Three-dimensional cutaway view of the KTeV apparatus. The components labeled “E799” are not used in this analysis. Figure courtesy of E. Pod.

KTeV uses a series of scintillation counters as trigger elements. Two banks of trigger counters called VV′ form the first level trigger for $K \to \pi^+\pi^-$ decays. Further banks of scintillation counters are used downstream to detect muons.

2.2 Acceptance Considerations

Re(ϵ'/ϵ) depends on the decay rates of the four $K \to \pi\pi$ modes, as shown in Equation 1.38. A measurement of the decay rates would require that we know the flux of parent particles and the acceptance of the detector (the fraction of decays that are successfully reconstructed by the detector). Since the measurement of Re(ϵ'/ϵ) involves a ratio of decay rates, and we collect K_S and K_L simultaneously, the double-beam technique causes the flux to cancel; the K_S/K_L flux ratio is determined solely by the material in the regenerator beam, and is the same for the $\pi^+\pi^-$ and $\pi^0\pi^0$ modes. The criteria for triggering, event reconstruction, and event selection are identical between the K_L and regenerator beam, so deadtime and veto-counter occupancy affect both samples in the same way. The detector acceptance varies greatly as a function of kaon decay position and kaon energy. Even though the *local* acceptance at a given (E_K, z) is nearly identical (up to small intensity-dependent effects that will be discussed later) between the two beams, the *global* acceptance between the two beams is manifestly different, due to the large difference in the lifetimes ($\tau_S \ll \tau_L$). This difference is evident in Figure 2.3, which shows the decay vertex position as a function of distance from the target. The regenerator beam has the characteristic exponential decay shape, while the vacuum beam looks roughly flat, due to the long K_L lifetime τ_L. There is no particular cancellation of acceptance effects between the $\pi^+\pi^-$ and $\pi^0\pi^0$ modes, because these modes are measured with different detector systems with different intrinsic acceptances.

There are two separate techniques for accounting for the acceptance differences between the two beams: a Monte Carlo acceptance correction and reweighting. The main KTeV analysis, detailed in this thesis, uses a Monte Carlo (MC) simulation to calculate and correct for the acceptance. The details of this simulation can be found in Chapter 8, but the simulation warrants a few introductory comments.

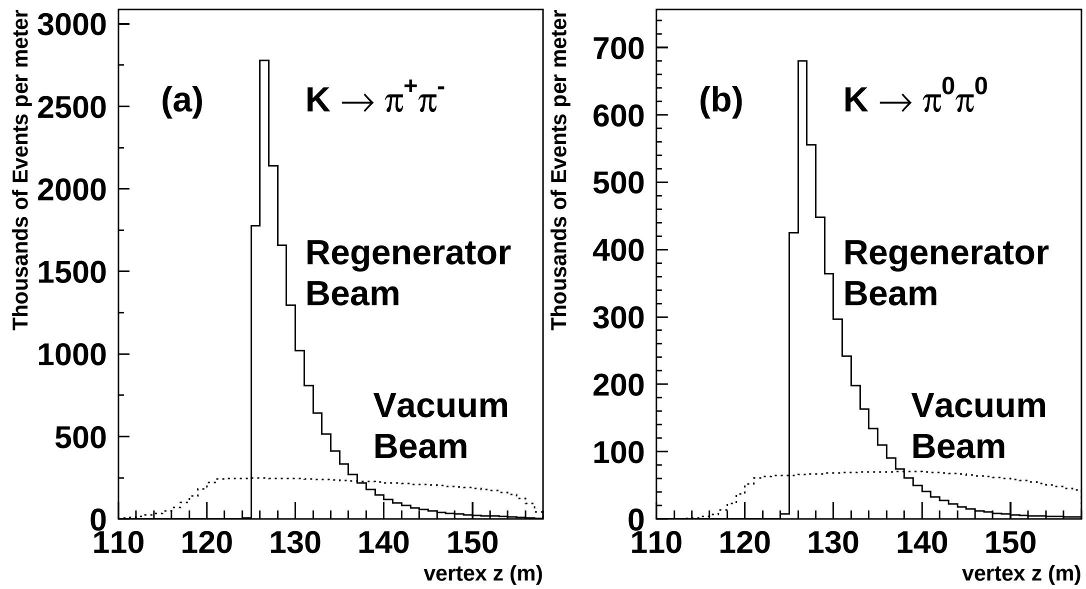

Figure 2.3: Decay vertex distributions for (a) $K \to \pi^+\pi^-$ and (b) $K \to \pi^0\pi^0$ decay modes, showing the difference between the "regenerator" (K_S) and "vacuum" (K_L) beams.

The Monte Carlo technique uses pseudo-random numbers to calculate the integral of a difficult multi-dimensional function: in our case, the acceptance. Standard high-energy physics Monte Carlos are also used to calculate production rates of complicated physics processes, such as $p\bar{p}$ scattering and jet production. The physics processes with which we are concerned ($K \to \pi\pi$) are simple, and do not require a Monte Carlo. The MC is only necessary to calculate the acceptance, and the acceptance for the most part is determined by the K_L and K_S lifetimes and the geometry of the detector, all of which are well-known.

A different technique to correct for acceptance difference between the two beams is a "reweighting technique." Here, the vacuum beam shape is reweighted by a function in (E_K, z) to make it match the regenerator beam (E_K, z) distribution. The reweighting technique assumes that the local acceptance at a given (E_K, z) is the same between the two beams. Because reweighting the vacuum beam reduces the statistical power of the downstream CP-violating $K_L \to \pi\pi$ decays, the statistical error of the measurement increases. We have performed a reweighting analysis, and used it to cross-check our main analysis, as will be detailed in Section 9.5.

2.3 Comments on Other Experiments

As mentioned above, KTeV is essentially an improvement on E731. The main differences are that KTeV collected $\pi^+\pi^-$ and $\pi^0\pi^0$ data simultaneously, whereas for most of E731 data for the two final states were collected in separate periods, and one photon in the $\pi^0\pi^0$ decay was required to convert in a lead sheet placed in the middle of the decay volume. The main systematic errors for E731 were due to calorimeter energy scale, beamline material and accidental activity in the detector, all of which have been addressed in KTeV.

The NA31 experiment at CERN [54] used a different technique to measure $\mathrm{Re}(\epsilon'/\epsilon)$. That experiment collected K_L and K_S decays in alternating periods, although in each period they collected $\pi^+\pi^-$ and $\pi^0\pi^0$ decays simultaneously, such that the integrated kaon flux canceled in the ratio. However, this technique required an absolute normalization between the charged-mode reconstruction efficiency and the neutral-mode efficiency. The K_S beam was produced by using protons incident on a target located within the decay volume. The target could be moved on a track along the decay volume, allowing the shape of the K_S beam to be matched to the shape of the K_L beam, minimizing the need for an acceptance correction. The NA31 detector did not use a magnetic spectrometer; a hadron calorimeter was used to reconstruct $\pi^+\pi^-$ decays. The largest systematic issues for NA31 were backgrounds, accidental activity in the detector and the calorimeter energy scale.

The NA48 experiment [64], which recently finished running at CERN, is the successor to the NA31 program. However, the design of the NA48 experiment uses a technique more similar to KTeV. K_L and K_S decays are collected simultaneously, with the K_S beam produced by a target in the decay volume just above the K_L beam. The K_S beam is angled so that the two beams converge at the center of the calorimeter. $\mathrm{Re}(\epsilon'/\epsilon)$ is extracted using the reweighting technique. NA48 has presented a preliminary result from its full dataset [64] with a total error of 2.6×10^{-4}.

CHAPTER 3
THE KTEV APPARATUS

In this chapter, we will discuss the KTeV beamline and the detector used in the E832 experiment. KTeV is two separate experiments: E832, designed to measure the direct CP-violating parameter $\mathrm{Re}(\epsilon'/\epsilon)$; and E799, designed to measure rare kaon decays, notably $K_L \to \pi^0 e^+ e^-$. The KTeV facility is located on the Neutrino beamline at the Fermi National Accelerator Laboratory (Fermilab) in Batavia, Illinois. The first round of E832 ran between July 1996 and July 1997.

Both experiments use two nearly parallel beams of neutral kaons, produced by primary protons from the Tevatron hitting a beryllium oxide (BeO) target. The two beams are required in E832 to provide a K_L beam and a K_S beam; the K_S beam is created by a regenerator placed in one of the K_L beams. E799 uses both beams as sources of K_L only.

The KTeV beamline and detector (collectively "the KTeV experiment" or simply "KTeV") are large and complicated devices, and much effort went into their design and construction. The following is a description of several of the KTeV subsystems in increasing distance downstream from the primary target.

3.1 The Beamline

Much work has gone into the design of clean, well-defined kaon beams for the KTeV program, based on the knowledge gleaned from the previous series of kaon experiments at Fermilab (E731, E773, and E799-I) [58, 63, 65]. The primary design consideration is to create high-intensity, pure kaon beams that are well-defined to ensure that neither the beams nor the beam halo impinges on the CsI calorimeter. This is accomplished through a well-considered beam line, utilizing sweeping magnets and collimators to define the beam.

3.1.1 Primary Proton Beam and Target

The Tevatron provides up to 5×10^{12} 800 GeV/c protons in a 18–20 second "spill" once every minute. The proton beam has a 53 MHz microstructure such that the protons arrive in 1–2 ns "buckets" once every 19 ns. The KTeV trigger is synchronized to the beam structure by a 19 ns timing signal provided by the Tevatron.

The proton beam is extracted from the Tevatron and directed to the target by a series of magnets called the "primary beam line." The primary beam spot is held to within 250 μm (RMS) in both the x and y directions by the final magnet focusing. Initially, the settings of the magnets within the primary beam line were determined manually; however, for most of the run, an automated program called "Autotune" kept the beam spot on the target.

The target is a narrow BeO rod, 3×3 mm in the transverse dimensions and 30 cm (1.1 interaction lengths) long. The primary beam is tilted downward at a 4.8 mrad angle with respect to the secondary (kaon) beam, as defined by the collimating system. The beam angle is chosen to optimize the kaon-neutron ratio and maximize the kaon flux.

The target area and beamline components are shown in Figure 3.1.

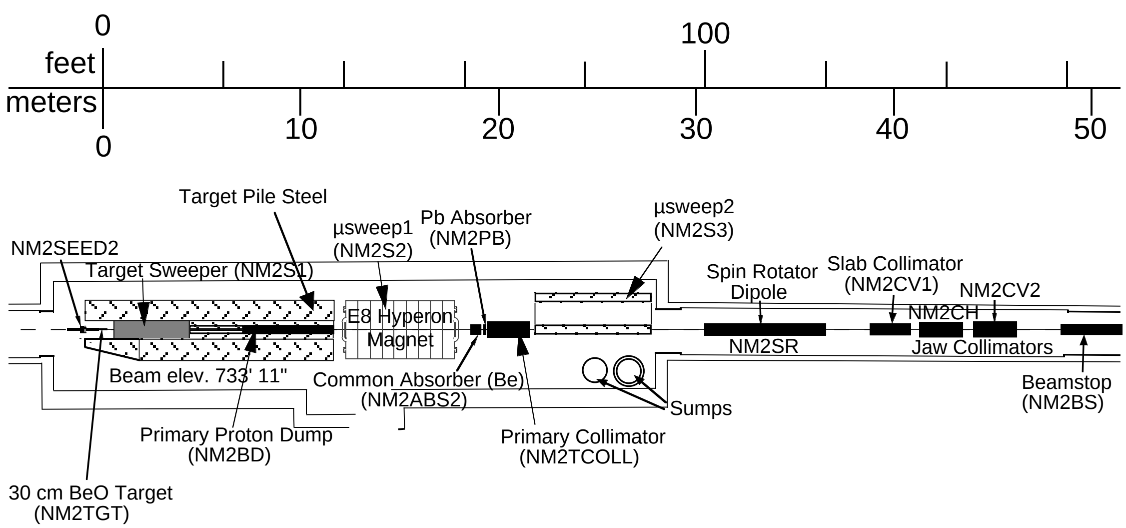

Figure 3.1: Secondary beam elements in the NM2 enclosure. Figure courtesy of R. Ford.

3.1.2 Experimental Coordinate System

KTeV is fixed to a coordinate system where the center of the target is defined as the origin. The positive z direction is defined as the downstream horizontal direction, which roughly points north. The vertical up direction defines the y direction. The x direction is chosen to define a right-handed system, in the horizontal plane, pointing roughly west, or to the left if looking downstream.

Although the target was initially located at the origin of the KTeV coordinate system, later surveys determined that it had sunk by roughly 300 μm in the y direction from its ideal position. We accounted for this movement when aligning the rest of the detector to the target.

3.1.3 Sweepers, Absorbers and Collimators

The kaon beam is created by removing charged particles and other neutral particles from the beam, through a series of sweeping magnets, absorbers and collimators. The first sweeper magnet, the "target sweeper," is 2 m downstream of the target. The main purpose of this magnet is to deflect the remains of the primary proton beam and charged particles produced in the target downwards into a water-cooled copper beam dump. The main goal of the remaining three magnets is to remove muons from the secondary beam. "μ–sweep 1" removes muons from the primary target; "μ–sweep 2" removes charged particles from interactions in the lead absorber and the primary collimator; and the final μ–sweep magnet, located at $z = 90$ m just downstream of the defining collimator, sweeps particles that arise from decays in flight or particles that are produced in the collimators.

The neutral beam is refined by a series of collimators and absorbing elements in the beam. The goal is to produce a beam that is sharply defined in the transverse directions with a high kaon content. The first aim is met by use of redundant collimators, and the second through the use of absorbers placed in the beam, since the neutral beam consists mostly of neutrons and photons. The first elements of the system are the common absorbers, located at $z = 18.5$ m. The first of the two common absorbers is 20 inches of beryllium, which attenuates neutrons more than kaons due to the different interaction cross-sections. The second absorber is 3 inches of lead, which converts photons to e^+e^- pairs, which are absorbed in the lead. An additional beryllium absorber moves between the two beams, shadowing the movement of the regenerator. Due to the movement of this absorber, it is named the movable absorber or the shadow absorber. These terms are used interchangeably. The movable absorber is placed in the beam to reduce the rate of hadronic interactions within the regenerator, which would raise the accidental rate within the detector. The length of the movable absorber is 18 inches long, a length chosen to minimize interactions in the regenerator while not adversely affecting the statistics of the CP-conserving 2π decays in the regenerator beam. At this point, the kaon-to-neutron ratio is roughly 1:1.3 in the vacuum beam and 1:0.8 in the regenerator beam [66].

The primary collimator is a 2 m long brass and steel block with square tapered holes that create the two neutral beams, and is located at $z = 20$ m. The two beam centers are separated by 1.6 mrad relative to the target. The primary collimator reduces the flux on the defining collimator. The next element, at $z = 38$ m, is the slab collimator, which sits between the two beams. The purpose of the slab collimator is to stop scattered particles that could cross from one beam into the other. The final element that defines the beam is the defining collimator, a tungsten block with tapered square holes, located at $z = 85$ m. The final exit of the defining collimator sets the beam shape and the beam divergence of 0.8 mrad. The KTeV collimator system is shown in Figure 3.2.

3.1.4 Accidental Counters

Two sets of counters in NM2 are used to trigger on primary beam activity uncorrelated with activity in the detector. Events triggered by these counters are used in the Monte Carlo simulation to model accidental activity in the detector, which will be discussed in Section 8.6.

The primary accidental counters are the “90° target monitor,” a series of three counters instrumenting a small hole in the target pile oriented 90° from the primary beam. A second set of accidental counters, the “accidental muon counters,” is

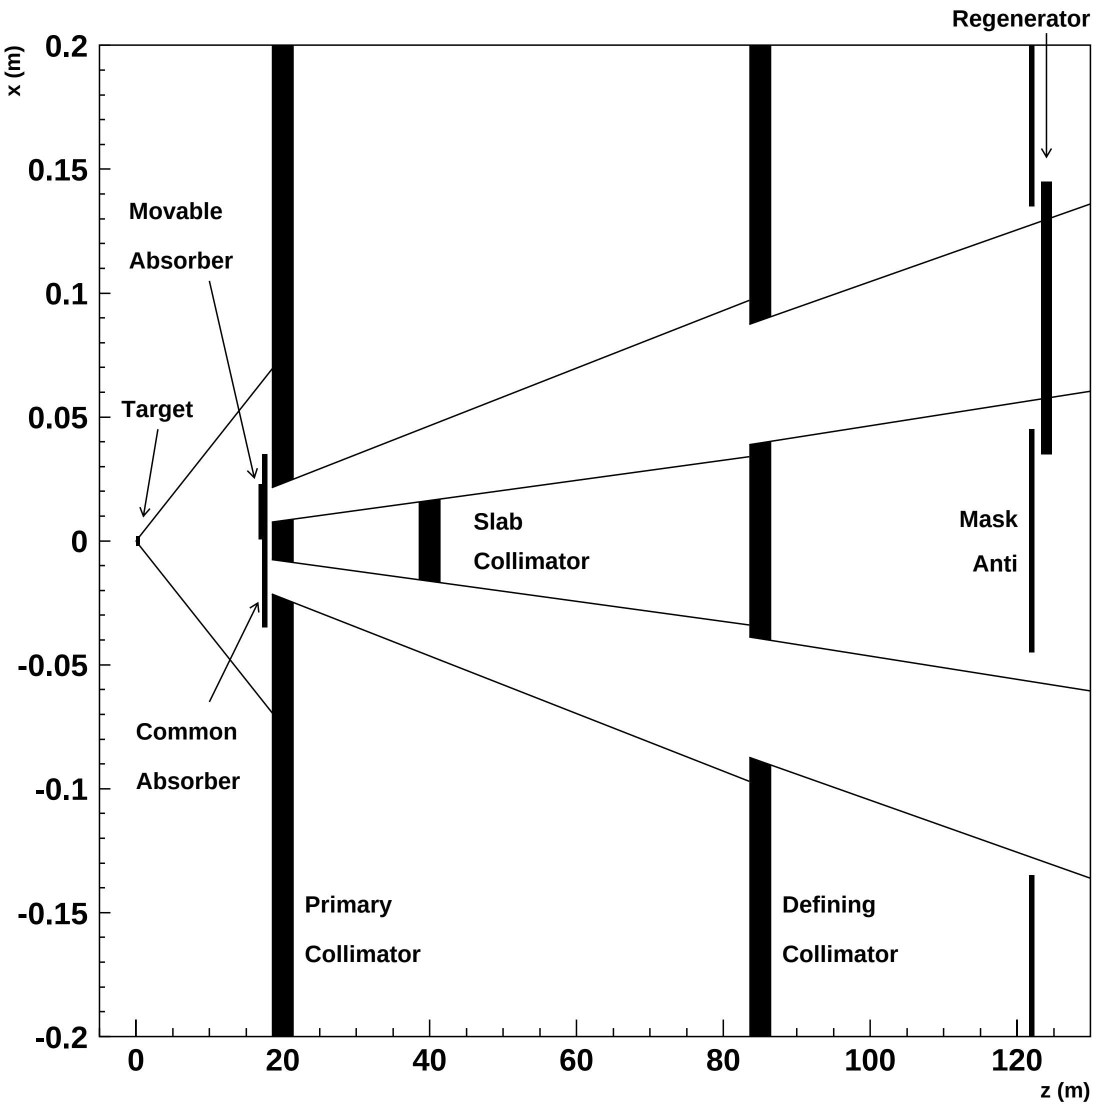

Figure 3.2: The KTeV kaon beam collimation system. After Figure 4.1 of Reference [2].

a system designed to trigger on large-angle muons originating in the target pile. Events recorded from this trigger were found to be slightly correlated with activity in the detector, so only events from the 90° trigger are used.

3.2 The Decay Region

The evacuated region starts at $z = 28$ m, just behind μ-sweep 2. This region is held at a vacuum near 10^{-6} torr, to nearly eliminate interactions between the neutral beams and any surrounding matter, and to eliminate scattering of charged decay products. The evacuated region starts as a 45.73 cm diameter vacuum pipe, which progressively gets larger in diameter. At the end of the decay region, at $z = 159$ m, the decay tank is 243.84 cm in diameter.

Many components of the KTeV detector are located within the decay region. These components are described in the following subsections.

3.2.1 The Regenerator

It is necessary to create a source of K_S from one of the two K_L beams. To this end, we use a device called the "regenerator," which regenerates a small component of K_S from the K_L in the beam. Regeneration occurs when a K_L interacts with hadronic matter. The K_L is a particular composition of K^0 and $\overline{K^0}$. In turn, the K^0 and $\overline{K^0}$, due to their different quark-antiquark contents, have different interaction probabilities with the matter of the regenerator, which contains only valence u and d quarks. The net effect is that the admixture of K^0 and $\overline{K^0}$ in the state is changed due to the interactions within the regenerator. This net change can in turn be re-expressed as an admixture of K_L and K_S states; some amount of K_S has been regenerated.

In KTeV, we use kaons which undergo coherent regeneration, where there is no measurable momentum transfer. The kaon state emerging from the regenerator can be described as $K_L + \rho K_S$, where ρ is a complex number, with a magnitude of ~ 0.03. The fact that the state is mostly K_L is acceptable, because of decays to $\pi\pi$ final states within our fiducial volume, most will be from the K_S component, due to the fact that ϵ_K is small and τ_L is large. Most of the kaons in the beam undergo quasi-elastic or inelastic scattering, with an appreciable momentum transfer. We desire to reduce the effect of these decays in our sample. To this end, the regenerator is active, made of 85 plastic scintillator modules, each viewed by two photomultiplier tubes which detect the recoil products of a non-coherent regeneration event. The signals from the PMTs are integrated and digitized with analog-to-digital converters (ADCs). A number of the signals are also discriminated, and the timing of the signal is used in the trigger system. There remains one type of unavoidable regeneration background: diffractive events leave no observable recoil products. Because we cannot veto these events in the data, we will have to do a background subtraction. However, the length of the regenerator has been chosen to be two interaction lengths to minimize diffractive events [67].

Because the regenerator produces a coherent mixture of K_L and K_S, and both of these states can decay to a $\pi\pi$ final state, interference occurs in the decay process. One can see this interference in the distribution of decays downstream of the regenerator, as shown in Figure 3.3. The decay rate to $\pi\pi$ final states downstream of the regenerator is

$$\begin{aligned} \frac{d\Gamma}{dt} \propto & \Big[|\rho(p)|^2 e^{-t/\tau_S} + |\eta|^2 e^{-t/\tau_L} \\ & + 2\,|\eta\rho(p)| \cos(\Delta m t + \phi_\eta - \phi_\rho) e^{-t(\tau_S/2+\tau_L/2)} \Big] , \end{aligned} \tag{3.1}$$

where p is the kaon laboratory momentum, t is the proper time ($= z/\gamma\beta c$, where z is the distance from the end of the regenerator, γ and β are the usual Lorentz boost factors and c is the speed of light), ρ is the coherent regeneration amplitude and ϕ_ρ is its phase, and η is either η_{+-} or η_{00} as described in Section 1.3 with a phase $\phi_{+-,00}$.

Many decays that we do not want to consider in our final sample occur within the regenerator. To define a sharp end of the regenerator, the last module is composed of a 5.6 mm lead–4 mm scintillator–5.6 mm lead–4 mm scintillator sandwich. For a $\pi^0\pi^0$ decay in the regenerator, there is almost certain to be one photon that converts within the lead, causing the event not to be reconstructed. For a $\pi^+\pi^-$

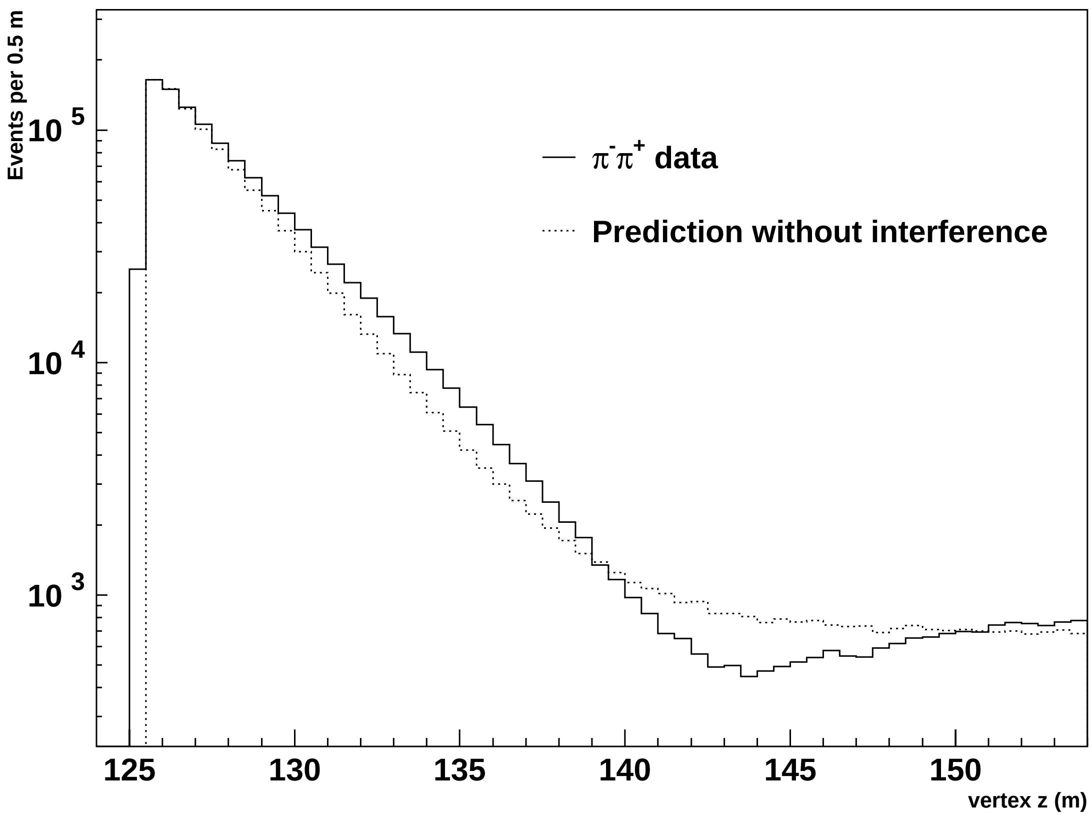

Figure 3.3: Interference in kaon decays downstream of the regenerator for 40 GeV < E_K < 50 GeV.

event decaying in the regenerator, the energy deposited by a pion track will usually cause it to be vetoed by the scintillator; however, the probability of veto drops as the distance of the decay to the end of the regenerator decreases. Hence there is an *effective* edge within the last few millimeters of the last scintillator. This effective edge is different for charged ($z_{\rm eff} - z_{\rm reg} = -1.65$ mm) and neutral decays (-6.2 mm), and is shown in Figure 3.4. Our understanding of this effective edge is a component of our systematic error, and will be discussed in Section 10.9.1.

The last lead piece contributes heavily to regeneration and scattering within the regenerator, and must be included in our models for the Monte Carlo and for the model of background diffractive scattering.

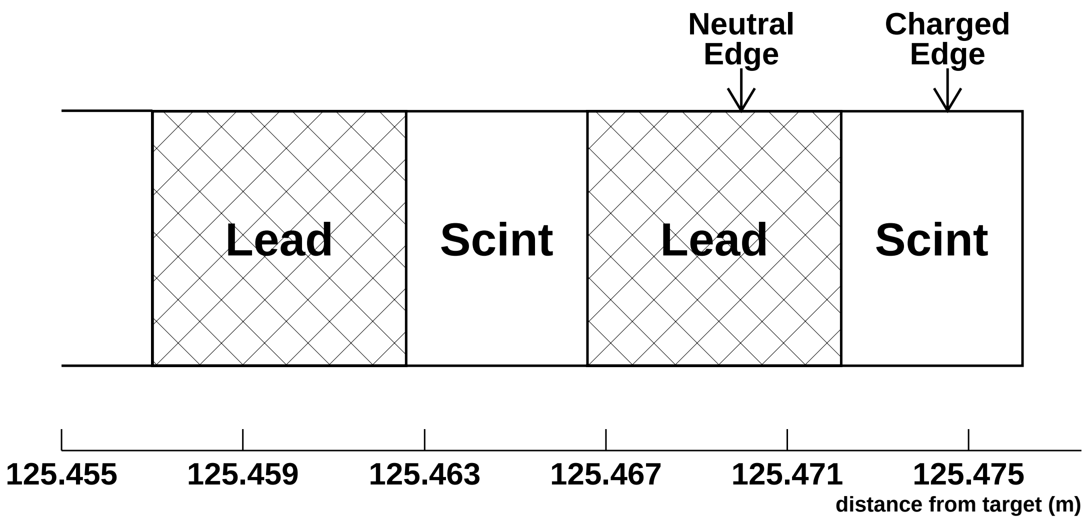

Figure 3.4: The effective regenerator edge. The arrows show the location of the effective edges within the regenerator for $K \to \pi^+\pi^-$ and $K \to \pi^0\pi^0$ decays.

3.2.2 Veto Counters: The Mask Anti

Along the KTeV beamline are a number of counters whose purpose is either to veto events where a decay product would be lost, causing a mis-reconstruction of the event, or to define a sharp edge within the decay volume.

The Mask Anti (MA) sits inside the vacuum decay region, transverse to the beam just upstream of the regenerator, at $z = 122$ m. The purpose of the MA is to veto events that occur upstream, specifically K_L events upstream of the regenerator in the regenerator beam. Hence it defines the upstream acceptance of the detector. The MA has two beam holes slightly larger than the beam profiles.

The MA is composed of a 16 layer lead-scintillator sandwich, 16 radiation lengths long, with fibers embedded within "wedges" of the scintillator. Each fiber is viewed by a PMT, the signal of which is then integrated in an analog-to-digital converter (ADC). The channels are also grouped together and discriminated to form a trigger source, which is used to veto events at the trigger level.

3.2.3 Veto Counters: The RCs

To veto events where a decay product leaves the fiducial volume of the detector, Ring Counters (RCs) are embedded within the decay volume. The counters are composed of 24 layers of Pb-scintillator sandwich (16 radiation lengths), with fibers to read out the scintillator in 12 azimuthal sectors. These vetoes have a circular outer aperture, and a square inner aperture. The inner aperture has a projective size, such that a decay product that misses the inner aperture of a RC will hit the CsI calorimeter downstream. The RCs are bolted into the vacuum tank at points where the tank segments join together.

As with the other veto systems, the fibers are read out by PMTs, which in turn are ADCed; the channels are ganged together to form a trigger level veto.

3.2.4 The Vacuum Window

The decay volume ends at a large window just upstream of the first drift chamber. We cannot use a vacuum region throughout the detector due to difficulty of running the drift chambers near a high vacuum region. Helium, in large “bags”, is used to minimize scattering between the drift chambers, up to the CsI calorimeter.

The “vacuum window” is a Kevlar weave (for strength), with a Mylar backing (for sealing). The window is 7.6 mm thick, and supports a force of 222 kN. It deflects by almost 15 cm from the edge of the vacuum tank to the center of the window. Due to the large force placed on the window and the possibility of catastrophic failure of the window, a steel shutter is placed over the window whenever the KTeV hall is open for access. The shutter lifts vertically out of the way when beam is on the target.

It is important to understand the multiple scattering properties of the window in analyses with charged final states. The window also serves as a production point for hadronic interactions that can be used to cross-check the calibration of the CsI calorimeter.

3.3 The Drift Chambers

The frames of the drift chambers (DCs) were the only component of the detector recycled from E731 [68]. All of the wires in the chambers were restrung. The more-numerous field shaping wires, which had been gold-plated copper-beryllium, were replaced with gold-plated aluminum to reduce multiple scattering from the wires. The sense wires were restrung with tungsten wire.

The DCs use a hexagonal cell geometry, with six field shaping wires providing the field around one sense wire (Figure 3.5). There are two sets of sense wires for each x or y view in the chamber, offset by one half-cell, to resolve the hit ambiguity. In the KTeV terminology, this layout is referred to as the "unprimed-primed" set of wire planes, such that the two planes in the x view are denoted the x–x' planes. Each chamber has four views; the two y views are upstream of the two x views. This layout allows us to resolve on which side of the wire the track passed, as well as giving a clear signal for tracks. Hits from good tracks should have a reconstructed Sum-of-Distance (SOD) between the two wires that is equal to the sense wire separation, 6.35 mm. For tracks with an missing hit, such that only one hit is resolved in the unprimed-primed pair, other tracking information must be used to determine which side of the wire the track passed. This effect will be discussed in more detail under the discussion of the track reconstruction (Section 6.2).

The calibration of the DC system for track reconstruction and use in the Monte Carlo is quite complicated and will be postponed until Chapter 5.

3.4 The Analysis Magnet

KTeV built a large dipole field analysis magnet, capable of providing a large transverse momentum kick (~ 412 MeV/c) to charged tracks. The design of the magnet allowed for an extremely uniform field across the clear face of the magnet ($< 0.25\%$). The uniformity was determined from a precise "zip-track" measurement of the field of the magnet [69], and the exact average field value was initially determined from Hall probe measurements and later refined by calibrating the transverse momentum kick by tuning the field to the known K_L mass in $\pi^+\pi^-$ decays. The magnet also

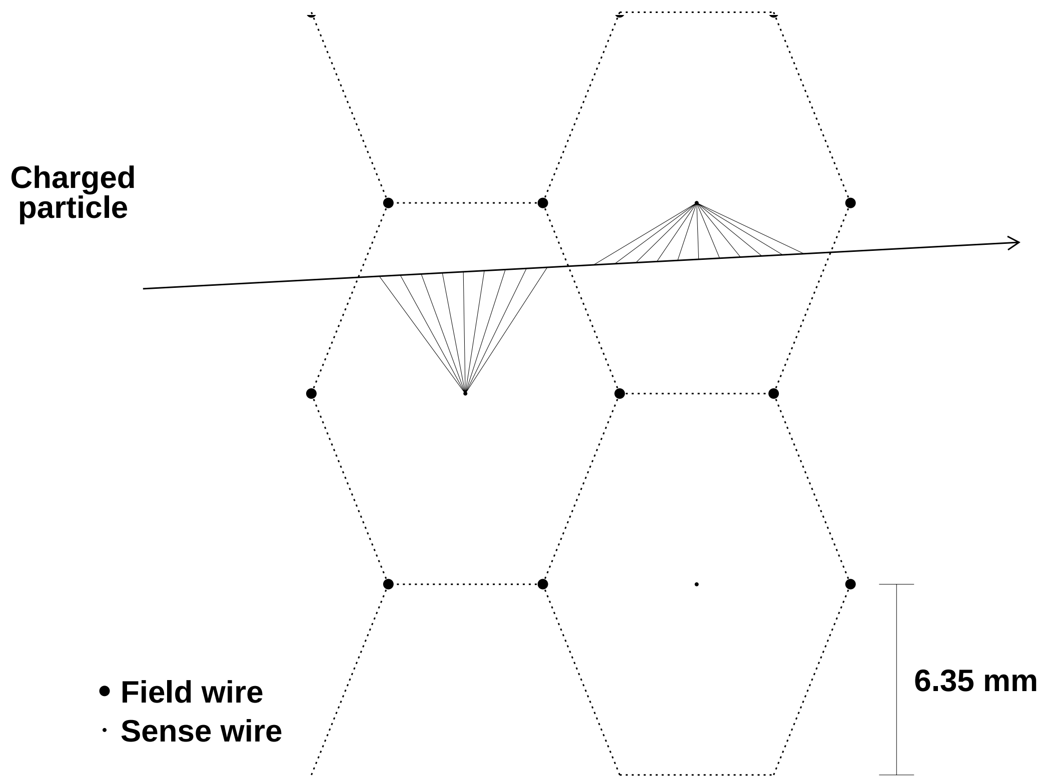

Figure 3.5: The hexagonal drift cell geometry. A typical charged-particle trajectory through the drift chamber is shown, where the thin lines represent the drift paths of ionization electrons toward the sense wires.

has extremely good fringe-field characteristics; the field falls from a maximum field of 3,000 gauss to 60 gauss at the location of DC 2 (4.4 m from the center of the magnet). Even this small fringe field must be considered for a precise determination of the track momentum.

The orientation of the magnet field was reversed every two to three days. Although the orientation of the field is not a systematic issue for $\pi^+\pi^-$ decays, due to the symmetric nature of the final state, it is an issue for asymmetry measurements in $K \to \pi^\pm e^\mp \nu_e$ decays [30, 70].

3.5 Trigger and Veto Counters

The most important set of counters is the VV′ counters, used as the primary trigger for charged particles. The other major counter group is the muon counter system, located at the end of the KTeV detector. In this analysis, the muon counters are used to veto muons from $K_{\mu 3}$ decays, but these counters were used to trigger on muons in certain rare decay searches. Other counters that are used only to veto particles leaving the detector volume, such as the Spectrometer-Antis, are described below.

3.5.1 VV′ Trigger Counters

Two back-to-back banks of scintillator counters are used to provide first-level trigger information for decays with charged particles in the final state. These counter banks are located 183 m from the target, just upstream of the calorimeter. The layout of the two banks of trigger counters is shown in Figure 3.6. Each bank of counters is composed of 30 or 32 counters, oriented vertically. The counters are 11 cm wide and either 88 cm or 110 cm long, and oriented so the counters are effectively split either just above or just below the neutral beam. The vertical orientation with a split at the beam line allows for simple determination of track multiplicity, since the analysis magnet separates tracks primarily in the x view. The second bank of counters has a similar layout, but with the counters offset by one half counter. The offset minimizes the systematic effects of gaps between the counters in a single bank, since the other bank has the center of a counter aligned with the gap between the counters in the other view. Since a large component of the radiation damage in the lead-glass calorimeter in previous experiments was caused by particles produced in interactions of the beam with the trigger counters [61], the KTeV counter banks have specially designed counters to allow for a beam hole, to eliminate material in the beamline.

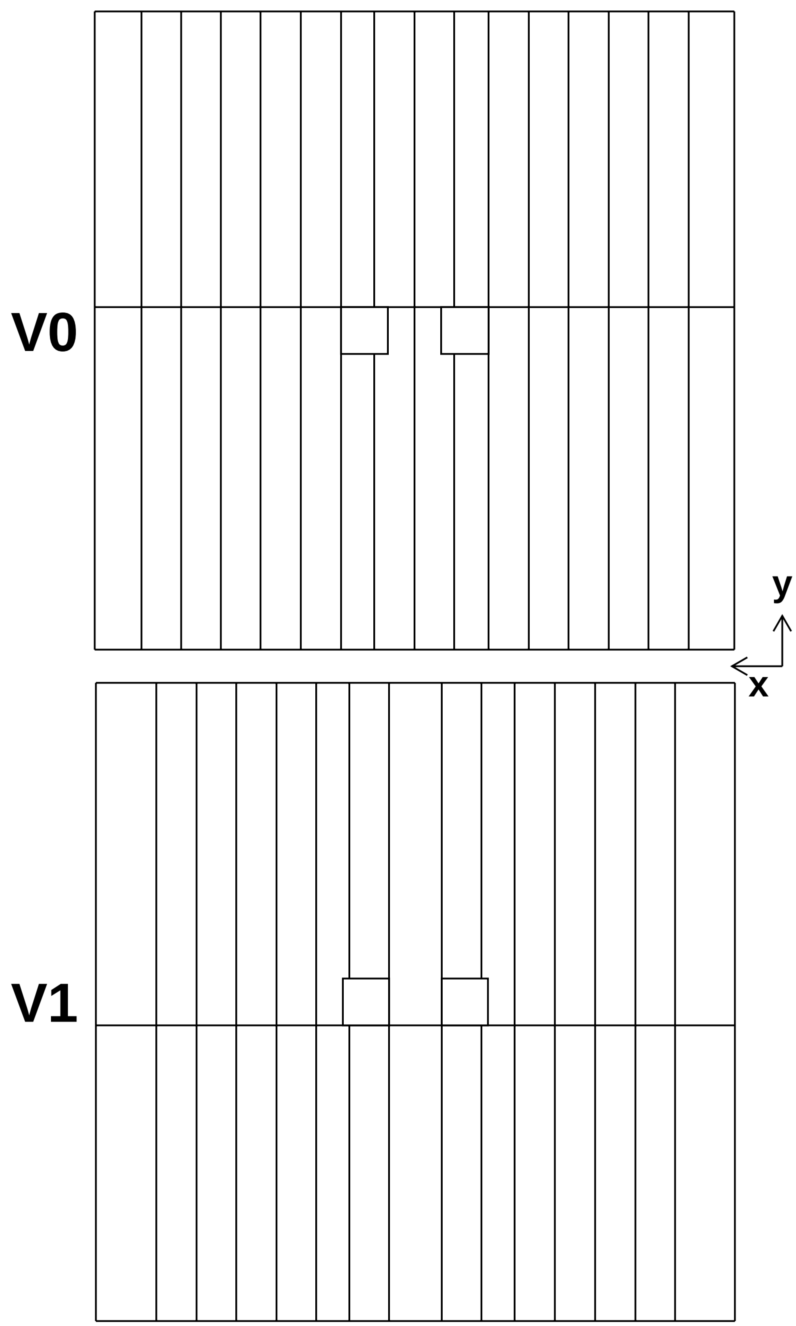

Figure 3.6: The layout of the V0 and V1 trigger hodoscope counters.

3.5.2 SA/CIA Photon-veto counters

The Spectrometer-Antis (SAs) and the Cesium Iodide Anti (CIA) provide additional photon-veto coverage. Each of these detectors is a square annular detector with an inner aperture that projects to the outer edge of the CsI calorimeter. The three SAs are located just upstream of the four drift chambers, while the CIA sits just upstream of the CsI calorimeter itself. The inner aperture of the CIA is small enough to occlude 1/2 of the outer edge of crystals in the calorimeter, to veto events that would be mis-reconstructed in the calorimeter due to energy leakage out of the side of the calorimeter. As with all photon-vetos, the channels are read out by ADCs, as well as used as fast "sources" for the trigger system.

3.5.3 The Collar Anti

The Collar Anti (CA) is intended to detect photons that approach the beam holes at the calorimeter. The CA occludes 1.5 cm of the crystals immediately surrounding the beam holes, thereby vetoing events where it would not be possible to accurately reconstruct the energy of photon. The CA defines the inner aperture of the detector.

The CA consists of two identical square annular detectors that sit directly on the carbon-fiber beam hole pipes in the CsI. Each half detector consists of layers of tungsten and scintillator, read out by long fibers going to PMTs located at the edge of the calorimeter.

3.5.4 The Back Anti

The Back Anti (BA) is a veto element used to identify particles that go down the beam holes. It is another two-part detector that is located directly behind the calorimeter. The BA is also segmented longitudinally to attempt to distinguish electromagnetic showers from hadronic showers. The BA is not used in the analysis with charged final states.

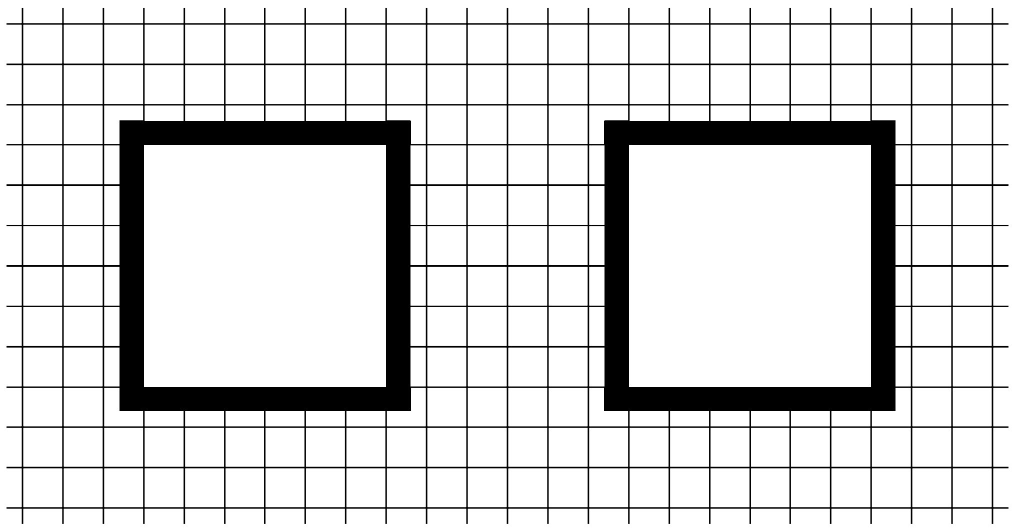

Figure 3.7: Geometry of the CA counters around the beam holes of the calorimeter. The grid indicates the edges of the CsI crystals (2.5 cm square), while the black annular regions indicate the two halves of the CA detector.

3.5.5 Muon Counters

Counters used for muon identification are located behind the calorimeter. The muon identification system starts with a 10 cm Pb wall, used to absorb electro-magnetic showers leaking from the calorimeter and to induce hadronic showers for particles that did not shower in the calorimeter. Those hadronic showers are detected in the Hadron Anti (HA), a bank of 28 overlapping counters. The HA is only used in this analysis in conjunction with the remaining muon system to identify muons for background rejection. Immediately following the HA is the first muon filter steel (MF1), a 1 m steel wall. The purpose of MF1 is to prevent backsplash from the neutral beam dump from firing the HA. The Pb wall, the HA and MF1 each have one rectangular hole through which both neutral beams pass.

The second muon filter (MF2) is 3 m deep, 3.4 m high and 4.3 m wide. MF2 forms the neutral beam dump, and stops most hadronic showers. Following the filter steel is MU2, a set of 56 overlapping counters. MU2 is required to be hermetic because it is used as a veto for events with muons in the final state. The large size of MU2 and the depth of MF2 were chosen to optimize muon rejection efficiency even for muons that undergo large scattering within the steel. Finally, there is an additional set of filter steel and counters, MF3 and MU3. MF3 is 1 m thick, and is designed to stop as much non-muonic activity that leaks through MF2, while retaining a reasonable muon momentum threshold. For the 31 interaction lengths of the Pb wall and the three muon filters, the momentum threshold to reach MU3 is 7 GeV, with a 0.5% probability of the hadronic shower from a 20 GeV pion reaching MU3.

3.6 The Calorimeter

The CsI calorimeter is the centerpiece of the KTeV detector. The lead-glass calorimeter used for E731, E773 and E779-I had been a major systematic limitation, although years of study had characterized it quite well. KTeV needed a new calorimeter with excellent energy and position resolution, fast signals and decent radiation hardness. After experimenting with barium fluoride and lead fluoride crystals, pure cesium iodide (CsI) crystals were chosen.

CsI has two scintillation components, a "fast" component with a decay time of ~ 25 ns at a wavelength of ~ 305 nm, and a "slow" component with a decay time of ~ 1 μs at a wavelength of ~ 480 nm. The fast component has sufficient light at KTeV energies so that the resolution term due to photostatistics is acceptably small. CsI has some properties that make it difficult to handle; it is fairly soft, which makes polishing difficult and makes the crystals susceptible to scratches and bending. It is also hygroscopic, requiring a dry environment for handling and storage. In addition, at the time of the design of the calorimeter, it was unknown if CsI was suitably radiation-hard for the KTeV environment. Crystals near the beam were expected to get a dose of several kilorads over the course of the experiment. Radiation studies were undertaken [2] to show that pure CsI could withstand the dose at KTeV.

The KTeV calorimeter was designed [71] as a square array, 1.9 m on a side, as shown in Figure 3.8. It is composed of 3,100 blocks of two sizes: 2.5 cm $\times$ 2.5 cm in the central region, where position resolution and determining cluster separation are most important; and 5 cm $\times$ 5 cm in the outer region. This arrangement saved a considerable cost, both on the manufacture of the crystals themselves, and the accompanying electronic channels. All of the crystals are 50 cm (27 radiation lengths) long, so that the electromagnetic shower from a high-energy electron or photon is almost fully contained, minimizing the resolution term due to fluctuations in the amount of energy lost out the back of the array. There are two 15 cm square beam holes which are supported by carbon-fiber tubes that allow the neutral beam to pass through the calorimeter.

After their manufacture, all of the crystals were processed at the University of Chicago [2]. The crystals' longitudinal responses were measured, and they were wrapped with mylar, reflective on one side. By tuning where to wrap the crystals with reflective mylar, the longitudinal response was made more uniform. Additional

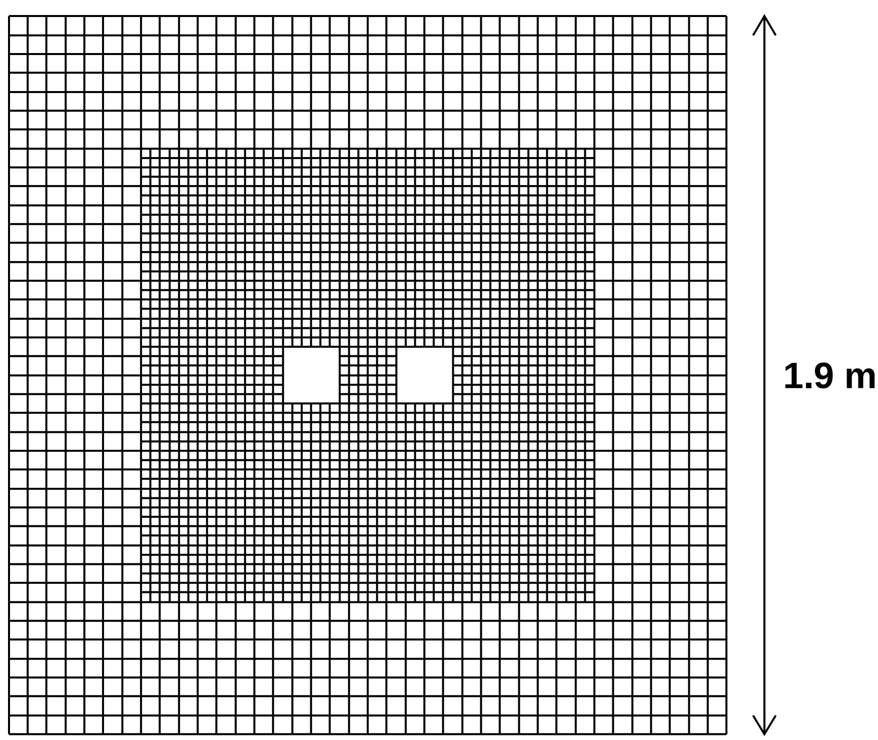

Figure 3.8: The layout of the CsI calorimeter.

parts of the preparation work involved measuring the size of each crystal and gluing a plastic "flange" to the end of the crystal, to which the phototube was attached.

To arrange the crystals into the array, several constraints were imposed. First, crystals in a given row had to have the same height, to make a base for the next row. Second, each pair of small crystal rows had to have the same width, so that they matched up to the large crystals at each end. Third, we required the crystals in the center of the array have good uniformity and radiation hardness. Finally, the crystals with the best uniformity were placed in the center of the array. We stacked a row of crystals per day, and measured their response overnight using cosmic-ray muons[1] before stacking the next row of crystals.

The calorimeter is located within a sealed metal enclosure called "the blockhouse." The blockhouse has its own dehumidifying and air conditioning system to protect the crystals and to cool the electronics. The first stage of the electronics sit on vertical "ribs" directly behind the crystals. The ribs support the electronics, and convey power, digital and analog signals, and an optical fiber to each crystal. The front of the blockhouse has a mylar window which allows particles to enter while keeping the blockhouse dark.

The scintillation light from a crystal is collected by a photomultiplier tube (PMT) coupled to the crystal with a silicone "cookie" and a disk of Schott UG-11 filter glass, which removes most of the slow-component scintillation light. The PMTs were custom-designed for KTeV by Hamamatsu to have improved linearity at the low gains we require. For a number of the crystals, we used a cookie with an annular mask to block some of the light; this allows us to run the PMT at a higher gain where the linearity is better.

The PMT signal is read out by a new device called the "digital photomultiplier tube" (DPMT) [72], where the PMT anode is connected by a short cable to a circuit that digitizes the signal, buffers the value, and transmits on demand. The DPMT is composed of two application-specific integrated circuits (ASICs) and a conventional flash ADC. The "charge integrator and encoder" (QIE) divides the input signal into

[1]We continued to monitor the crystals' response *in situ* throughout the experiment with cosmic-ray muons.

several ranges, with 1/2, 1/4, 1/8, ... of the input signal), accumulates charge on capacitors over a given time interval (a "slice"), selects the range for which the voltage on the capacitor lies within a given window, and outputs that voltage to be digitized by the separate flash ADC chip. The range selected is indicated by the binary "exponent," and the flash ADC indicates the digital "mantissa" of the signal. The second ASIC is the "driver, buffer, clock" (DBC) chip. The DBC provides a FIFO buffer to store the mantissa and exponent from the QIE in each time slice, and transfers the data to an external "pipeline" module that is in turn read out into the KTeV data stream when an event is accepted by the trigger. This system allows for a wide dynamic range with pulse information, and with low noise. The pedestal width (the baseline values recorded in the DPMT when no energy is deposited) is ~ 0.5 MeV).

The PMT dynode for each channel is fed into the "Energy-Total" trigger system (ET). This system provides a quick trigger-level estimate of the energy in each channel, and sums all channels together to estimate the total energy deposited in the calorimeter.

Due to the nature of the QIE, the calibration procedure is quite complex. Cross-calibration of different QIE ranges is done with a laser system which distributes shaped pulses of UV light (from a dye illuminated by the primary laser pulse) along a quartz fiber which connects to the flange on the end of each crystal. The light intensity is scanned through the full dynamic range, and referenced by a PIN photodiode with excellent intrinsic linearity. We also use this system to monitor short-term changes in the PMT gains during normal running.

The charge-to-energy calibration for each channel is determined with electrons from K_{e3} decays, matching the energy measured in the calorimeter to the momentum measured in the spectrometer, through an iterative process. In 1996, we calibrated using 1.9×10^8 electrons, and in 1997 we used 4.24×10^8 electrons, all collected during normal running. Figure 3.9 shows the ratio of the calorimeter energy to the track momentum ("E/p") for the 1997 sample of electrons after all corrections have been applied. Figure 3.10 shows the resolution of the calorimeter itself, after the component due to the momentum resolution of the spectrometer has been subtracted) as a function of energy. The calorimeter has an average energy resolution of 0.72% for photons from $\pi^0\pi^0$ decays (with an average energy of 19 GeV). The x and y position resolutions are roughly 1 mm in the small crystals and 1.8 mm in the large crystals.

3.7 The Trigger System

The Tevatron delivers approximately $3,000$ protons to the KTeV target in each 19 ns bucket. Of those protons, only one in 10^5 produces a kaon that passes through the collimator system and reaches the detector. Of those kaons, a small fraction decay into interesting final states. The KTeV trigger must identify these decays and prompt the detector to be read out to tape.

The KTeV trigger system is built in three levels, each with different deadtime limits and correspondingly different decision strategies. Level 1 uses fast signals from various detector systems, synchronized to the 19 ns beam structure, and simple logic to do basic pattern identification in each bucket. Level 1 is deadtimeless. Level 2 is a set of specialized processors which evaluate the number and pattern of signals on individual drift chamber wires and calorimeter channels. An event that passes Level 2 causes the entire detector to be read out into computer memories. Level 3 is a software trigger running in parallel on 24 CPUs, which reconstructs every $\pi^+\pi^-$ and $\pi^0\pi^0$ event and applies loose kinematic cuts before writing the event to digital tape. Prescaled samples of the more numerous decay modes are also written to tape; these samples are used to calibrate and to cross-check the performance of the detector.

We provide a brief description of the trigger system, with a focus on the trigger elements relevant to $K \to \pi^+\pi^-$. An exhaustive summary of the trigger can be found in Chapter 3 of Reference [73].

3.7.1 Level 1 Trigger

The Level 1 trigger uses fast Boolean “sources” and simple logic to create a trigger decision. A source is a digital signal from one trigger system element which is “on”

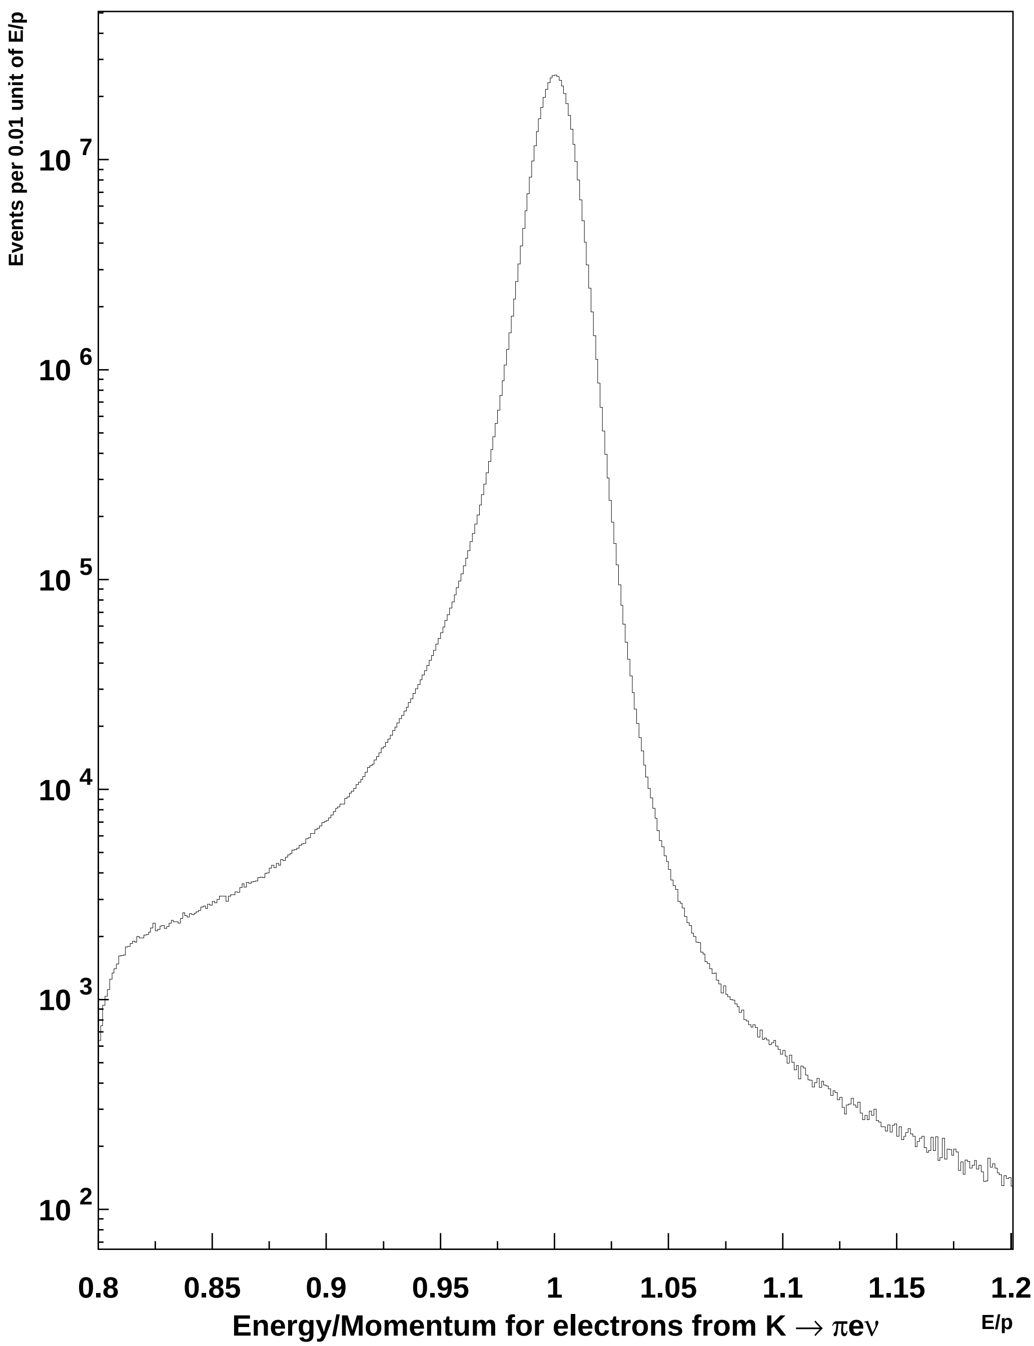

Figure 3.9: The ratio of calorimeter energy to track momentum ("E/p") for electrons from K_{e3} decays, for the 1997 dataset. The plot has 4.24×10^8 events. The resolution on this quantity is 0.72%. The cutoff at the left edge of the plot is due to a Level 3 cut used to select candidate K_{e3} events.

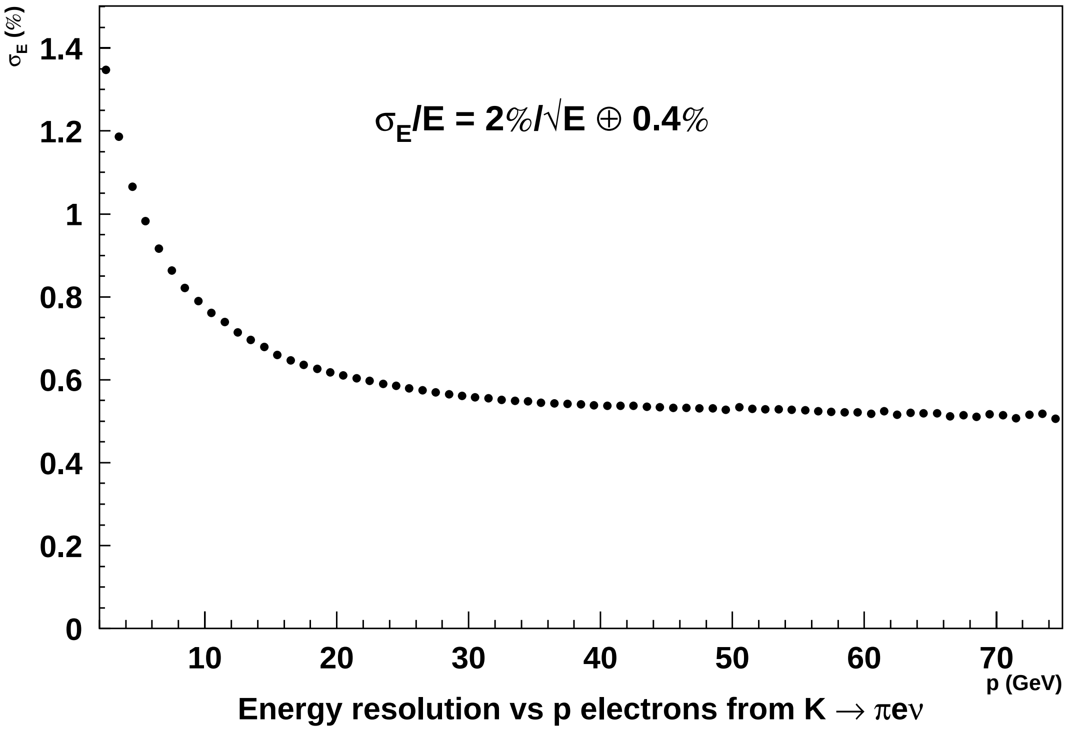

Figure 3.10: The resolution of the calorimeter as a function of energy, measured using electrons. The momentum resolution due to the spectrometer has been subtracted.

when that element is true. For example, one source indicates that at least two counters in the V0 trigger bank were hit; another source indicates that there is at least 24 GeV deposited in the calorimeter, based on an analog sum of PMT dynode outputs. These sources are from fast signals in the detector (usually photomultiplier tube signals), are formed with commercial electronics, and are timed to the beam. Because we use fast signals and simple pattern-matching, the Level 1 trigger is deadtimeless.

To reduce the trigger rate, the Level 1 trigger can be inhibited by fast veto signals from the regenerator, the MA, the photon vetoes, or the muon system.[2] An "accept" decision at Level 1 starts the digitization of the detector, in the form of

[2]Hence, while the Level 1 trigger is deadtimeless, accidental activity in various detectors can veto otherwise good events.

ADCs, TDCs and latch modules, and stores the continuously-digitized calorimeter information in first-in-first-out buffers (FIFOs). The raw Level 1 rate is 60 KHz under normal operating conditions.

3.7.2 Level 2 Trigger

The Level 2 trigger is composed of custom electronics used to do pattern matching at a more complicated level than at Level 1. We are concerned with biases between the vacuum and regenerator beam at the trigger level due to underlying event activity in the detector. Of special concern are hadronic interactions in the regenerator, which could put more activity on one side of the detector. To avoid this problem in the charged-mode analysis, only DC information from the y view is used, meaning that the hits from the same wires are used from both the regenerator and vacuum beams. The Level 2 charged trigger system required there to be enough hits in the y wires to be consistent with a two-track event. This hit-counting system consisted of custom VME boards called "Bananas" and "Kumquats." These boards take fast signals from each wire and count the number of hit-pairs in neighboring wires in each chamber, within a 242 ns timing interval, gated from the Level 1 trigger signal. The Kumquat boards, used for chambers 3 and 4, latch hits within a 220 ns time interval. The Banana boards, used on chambers 1 and 2, have more functionality and can correlate the time information between hits on neighboring wires, and in effect cut on the SOD. However, due to concerns about inefficiency at the trigger level, the Bananas are used in a "wide-open" mode for collecting Re(ϵ'/ϵ) data, and are functionally equivalent to the Kumquats. A full description of the hit-counting system can be found in Reference [74]. The hit-counting system has a decision time of $\sim$ 800 ns.

The "y-Track Finder" (YTF) is used in Level 2 to determine if the hits in the DC system are consistent with two straight tracks in the y view. The YTF is designed for $K \to \pi^+\pi^-$ so it looks for one up-going and one down-going track, with a region in the center of the DC system counting as both up- and down-going. The YTF uses coarse-grained outputs (typically the OR of 16 contiguous wires) from the hit-counting system as its input, and usually returns a decision a few hundred nanoseconds before the hit-counting system. The YTF was implemented with commercial programmable-logic and memory-lookup modules.

Although not used in the trigger decision for $K \to \pi^+\pi^-$, the Hardware Cluster Counter [75] deserves some comment. The HCC was designed to count isolated clusters of energy in the CsI calorimeter, where each cluster is defined to be a set of contiguous channels with energy above a nominal threshold of 1 GeV. Channels touching at a corner are not considered contiguous. The inputs for the HCC are digital bits from the ET boards, which are on if the channel is above the 1 GeV threshold. The algorithm for finding a cluster is quite ingenious. For any contiguous set of crystals that compose a cluster, traversing the perimeter of the crystal clockwise means turning through 360°. Since the cluster is composed of square crystals, 360° can be broken down into four 90° "right-hand turns." So a cluster can be defined as a set of crystals where $N(\mathrm{RHT}) - N(\mathrm{LHT}) = 4$. More importantly, the crystal topologies can be counted in 2×2 groupings in parallel; the number of clusters is $N(C) = 1/4 \sum_{2\times2} [N(\mathrm{RHT}) - N(\mathrm{LHT})]$. There are intricacies involved with the large versus small blocks and the blocks near the beamholes. The HCC has the longest decision time at Level 2, 1.5 μs.

If an event passes Level 2, digitization of the detector is allowed to continue, and the detector is read out into VME memory buffers. Detector read-out takes 18 μs on average, depending on the occupancy in the detector. If the event fails Level 2, the detector front-end modules are cleared and the trigger system is re-enabled. Read-out at Level 2 causes a fractional deadtime of $\sim 35\%$ under normal running conditions. The Level 2 trigger rate is about 10 kHz.

3.7.3 Level 1 and Level 2 Trigger Definitions

Levels 1 and 2 of the trigger use programmable memory lookup units (MLUs) to make trigger decisions based on source inputs. The trigger sources are divided into twelve groups of eight sources each. Each group is fed into one MLU, the output of which is the 16-bit "word" detailing which triggers have sources that are "on" within the group. The final trigger is formed by the AND of the MLU outputs of the twelve groups (done through a three level "tree" of coincidence units). Because the final trigger is the AND of the outputs of the MLUs, care had to be taken to ensure that sources were divided into groups in such a way that sources that needed to be “ORed” together remained in the same group.

A text file is used to specify the trigger configuration. This file contains the Level 1 and 2 trigger definitions, as well as hardware prescales. The file is parsed by an “interpreter” which checks the consistency of the triggers, and generates the maps which are downloaded to the MLUs through the CAMAC hardware protocol. In general, we have attempted to define the loosest possible triggers, in an attempt to attain 100% efficiency for the $\pi\pi$ signal modes, to minimize possible biases on $\mathrm{Re}(\epsilon'/\epsilon)$. For example, we have chosen trigger definitions that allow for at least two hits in one V-bank and one hit in the other, to minimize susceptibility to cracks between the counters, and we allow for 4 or 5 clusters in the HCC, to reduce events lost due to hot channels in the CsI.

In addition to the two main triggers for $\pi\pi$, we define up to 14 additional beam triggers. Most of these triggers are looser versions of the signal triggers, taken with a higher prescale, to allow us to study remaining inefficiencies in the trigger. Other triggers collect events for rare kaon and hyperon decay searches. The final beam triggers are the accidental triggers, that uses hits in the accidental counters to record underlying detector activity for later use in the Monte Carlo simulation.

We also take “calibration” triggers to monitor the performance of the detector. For example, we use a calibration trigger to monitor the response of the CsI calorimeter to the laser-light system, and use a calibration trigger to readout the ADC and DPMT pedestals. Most calibration triggers are taken in the 40 s between Tevatron spills, although some are defined to be accepted continuously.

3.7.4 Level 3 Trigger

If the Level 2 trigger is satisfied, the entire detector is read into VME computer memories. With the reduction we obtain with the Level 1 and 2 triggers, the memories have enough space to store the data from the entire 20 s spill. Hence the 24 CPUs that compose Level 3 have the entire minute to process the data without incurring dead-time. Level 3 applies a minimal set of kinematic and particle identification cuts. Because Level 3 relies on preliminary "on-line" detector calibrations, the cuts must be loose.

For the $\pi^+\pi^-$ trigger, track and vertex candidates are reconstructed first. The track reconstruction is detailed in Chapter 6. If we find a possible vertex candidate, the calorimeter information is unpacked and clusters are found and matched to tracks.[3] From the calorimeter and tracking information, we make loose particle identification cuts, to attempt to reduce the K_{e3} background. The most important feature of Level 3 is that an *inclusive* approach is used: if *any* vertex candidate passes the vertex and particle-identification cuts, it is tagged as a $\pi^+\pi^-$ candidate. The vertex candidate that matches the Level 3 cuts does not have to be the best vertex candidate; the best vertex candidate will be determined in the offline analysis, when the final calibration of the detector is available. The exact details of the Level 3 cuts will be discussed in Chapter 6.

Processing for the $\pi^0\pi^0$ events is more straight-forward. The clustering software must find four clusters corresponding to those found by the HCC. The clusters are paired to make π^0s (see Section 7.3), where the correct pairing is the one that gives the same z-vertex for both π^0s. We also consider events where the second-best pairing of clusters gives a consistent z-decay vertex. Events are written to tape if $m_{\pi^0\pi^0} > 450$ MeV/c^2 for either the best or second-best pairing.

In addition to the signal modes, Level 3 tags and prescales samples of K_{e3}, $K_L \to \pi^+\pi^-\pi^0$, and other decay modes from within the Level 2 triggers. A fraction of all input events are written to tape regardless of the decision of Level 3, although the Level 3 decision is written into the data header; these events allow us to study biases within the Level 3 software.

Any event tagged as belonging to one or more data samples is written to digital linear tape (DLT). Under nominal running conditions approximately 40,000 events are written to tape each accelerator spill. In total, E832 wrote 4.9×10^9 events to 2,645 15 Gb tapes. Dealing with this large sample of data is a complex problem, which we will begin to describe in the next chapter.

---

[3]There is not sufficient CPU time to run clustering for all events, as clustering is much more CPU-intensive than track reconstruction.

CHAPTER 4
THE DATA

In this chapter, we detail how the data were taken, and discuss a number of problems uncovered during the running period. We detail the data reduction from the 2,645 online tapes to a sample of more manageable size, and finally identify the samples used for this analysis.

4.1 Data Collection

KTeV collected data in the 1996–1997 and 1999 Fermilab fixed target runs. A great deal of effort was involved to commission and run the detector before and during the fixed-target run, with around the clock activity starting months before the arrival of beam in Autumn 1996. To set the timeline of progress on KTeV, the civil construction of the hall finished in September 1995, and installation of the detector started immediately. The most difficult installation was the stacking of the CsI calorimeter, discussed in detail in Reference [1]. Stacking of the CsI finished on July 21, 1996. First muon beam was delivered to KTeV on July 31, 1996. We opened the beamstop for hadrons on August 31, and KTeV saw kaons for the first time.

An intensive seven week debugging period followed, during which some problems (to be detailed below) were discovered and remedied. KTeV was determined to be operational on October 24, and the first physics-quality data were taken.

The running time for KTeV was divided between the two physics programs: E832 and E799-II. The division of time is shown in Figure 4.1. The detector configuration is different between the two experiments: E799 removed the beryllium absorbers, the MA and the regenerator, installed Transition Radiation Detectors

(TRDs) downstream of DC 4, and used slightly larger beams for the Summer 1997 and 1999 running. E832 ran in three periods: Winter 1996, Spring 1997, and 1999. The data collecting efficiency improved after 1996, so much so that 3.5 times the number of $\pi\pi$ events were collected in 1997 as in 1996, although the period was only twice as long.

Data collection was a 24-hour job, with teams of three people on shift. Responsibilities of the shift crew included starting new "runs" (which could last up to 10 hours, if the Tevatron ran smoothly), monitoring the detector for various problems, fixing the minor problems and calling in experts to solve major problems.

4.2 Problems Encountered

Even with the best efforts of the shift crews to keep KTeV running smoothly, problems occurred that required a stop in data-taking to fix. These problems ranged

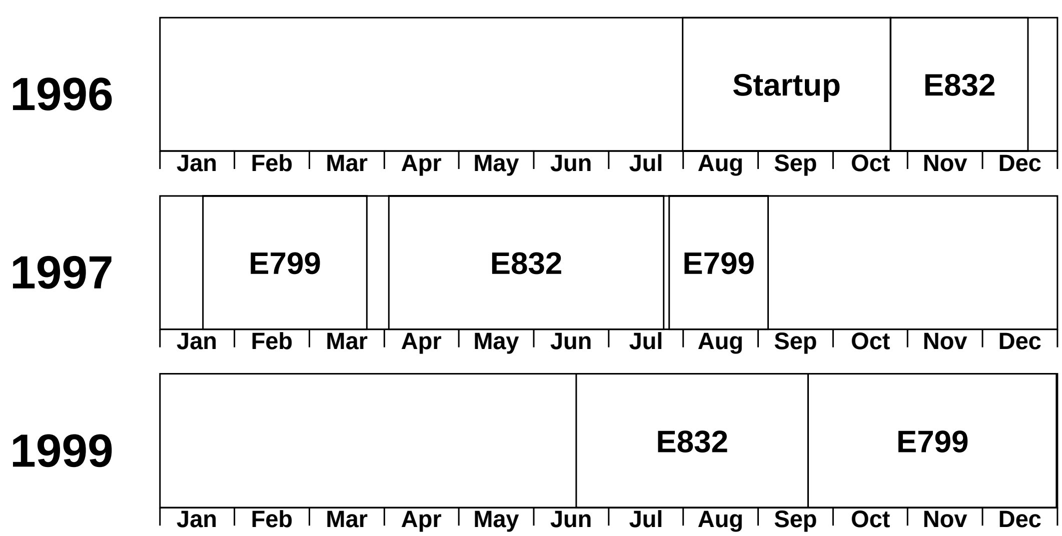

Figure 4.1: Division of running time during the fixed target run. Although E832 data were taken in 1999, they are not used in this analysis.

in severity from minor electronics glitches that only required the detector to be re-initialized, to major equipment failures that required diverting the primary proton beam for long periods to replace parts of the detector. The major category of equipment failure was the custom integrated circuits that composed the CsI readout system, although other systems also had failures. On the whole, KTeV lost 15–20% of otherwise good beam time. Additionally, the Tevatron itself was down at least 10% of each week, with occasionally longer downtimes. We attempt to detail some of the larger problems encountered while taking KTeV data below.

4.2.1 Calorimeter Readout Problems

The first stage of the CsI calorimeter electronics, the QIEs and DBCs, are custom ASIC chips designed to meet the stringent requirements for resolution, linearity, dynamic range, and low-noise operation. They are installed on the DPMTs, mounted directly behind the phototubes. After installation, a number of problems related to lifetime of the chips were discovered. The lifetime of individual channels was supposed to be a number of years, but with 3,100 channels, DMPTs were failing at the rate of a few per day. The main class of problem with the DBC chip was due to traces on the chip failing, causing the incorrect value to be read out of the chip. This mistake was due to a combination of a design flaw coupled to a mistake in the manufacturing process. Once the problem was identified, the entire batch of DBCs was remade, with design modifications to correct for the flaw and to add robustness against the mistake in the manufacturing process. The calorimeter electronics were completely replaced with the new DBC chips in September 1996.

The QIE chips were not immune to problems. In the 1996 running, it was found that one batch of QIE chips (roughly 1/3 of the calorimeter) would read out incorrectly when run at the design frequency of 53 MHz. During the 1996 run, the DPMTs were clocked at "RF/3" (18 MHz), where each DPMT clock cycle would integrate the CsI signal over three beam buckets. The QIEs sometimes failed outright in a manner similar to the DBCs, due to the same flaw in the manufacturing process. These failures necessitated frequent replacements during the data collecting. Roughly half of the QIEs were remade and replaced over the December 1996-January

1997 shutdown, which allowed the DPMT system to be run at 53 MHz for the 1997 period. However, the new batch of QIEs still had the manufacturing flaw, and continued to fail at the rate of one per day throughout the 1997 run.

All the above problems were extensively monitored via software during running, and were flagged within a few minutes. Audible alarms would alert the shift crew, who would begin the process of replacing the DPMT. This process entailed steering away the primary proton beam, putting the beam stop in place, and accessing the CsI blockhouse. The actual replacement time was of the order of 10 minutes; most of the lost time was due to moving the proton beam, and getting it back after the access. We followed the procedure of quick replacement for the entire run, although Monte Carlo studies showed that the bias on $\mathrm{Re}(\epsilon'/\epsilon)$ due to these failure modes would be small, even if the problem persisted for most of the run and were to be ignored in the analysis.

4.2.2 Calorimeter Trigger Problems

The calorimeter trigger also had a number of problems during commissioning of the detector. The ET system uses fast analog signals from each of the 3,100 calorimeter channels (taken from the last dynode stage of each PMT) to generate two types of trigger information. As its name suggests, ET sums the total amount of energy in the calorimeter as a Level 1 trigger source. For Level 2, the ET system discriminates each of the 3,100 channels to produce the input for the HCC. In mid-September 1996, all ET boards had to be removed to replace a set of leaky capacitors. The replacement was done at the same time as the first DBC replacement. In early October 1996, we discovered some HCC input bits were being set even though the channel had no deposited energy. These "hot bits" were discovered to be due to comparators damaged when the air conditioning for the ET system failed. The worst comparators were identified and fixed, and the system was declared operational on October 24, the start of physics-quality data. However, there were low-level "hot bits" throughout the 1996 data period. The $\pi^0\pi^0$ trigger was "opened up" to allow an extra HCC cluster to avoid this problem. The comparators for the E-Total system were completely replaced in the shutdown between 1996 and 1997, and the $\pi^0\pi^0$ trigger required four and only four HCC clusters for 1997.

4.2.3 Drift Chamber Problems

The drift chambers had a number of problems, the most serious of which was not discovered until the end of the 1996 period. The first problem was noise pick-up on the sense wires. The noise was large enough that when amplified by the DC electronics, it would be above the discriminator threshold and would be recorded as "fake" hits. A particular example was the pickup of substantial noise during the asynchronous readout of the digitized TDC information. A number of modifications were made to the amplifier gains and the grounding of the DC system, which reduced the noise to below the discriminator threshold level.

During the 1996 run, it was discovered that a low noise metal-film resistor on the amplifier board could fail through the discharging of the chamber high voltage. During the 1996 run, this failure rate was approximately one channel per week in nominal conditions. The resistors were replaced on the 1972 DC channels during the shutdown between the 1996 and 1997 running periods.

During the commissioning, there were periods in which the DCs would draw excessive amounts of current, sometimes tripping the high voltage supplies. A few of these cases were identified as being due to wire stubs or other objects on a wire within the chamber, and cleaning the wire would fix the problem. With the number of unexplained cases of current-draw before the arrival of high-intensity beam, there was concern about the performance of the chambers under high-intensity beam, and the decision was made to reduce the operating high-voltage set point of the DCs. Even with this modification to the chambers' operating mode, there were times during nominal data-taking when the chambers would draw excessive current. In these cases, the HV to the chambers was cycled, which cured the problem. The HV draw did not seem to adversely affect the performance of the DC system.

In addition to high-current draw, the chambers would, on occasion, "oscillate," causing large numbers of fake hits on many wires within a chamber. The chambers remained in this state until the HV was cycled. The cause of this oscillation was never conclusively determined, but was believed to be due to the high-impedance pickoff of the discriminated signals, which was used as input to the hit-counting trigger system.

The largest problem for the DC system was not discovered until the E799 test period in December 1996. During the debugging of the Level 3 software for E799, a number of events were seen in which the software could not reconstruct two tracks, even though the Drift Chamber system seemed to have the requisite hits. Detailed inspection of these events showed that one of the hit-pairs had a surprisingly high sum-of-distance (SOD), larger than the nominal value of 6.35 mm by more than the 1 mm cut used to classify “good” SODs in the Level 3 tracking code.

The cut of +1 mm used in the Level 3 track reconstruction code was determined from studies of the Drift Chamber system in E773, the immediate predecessor to E832. Although many physical processes can cause *low* SOD pairs, such as delta rays or two tracks within the same cell, high SODs were believed to arise only when hits from real processes were mis-paired with each other (“combinatorics”), and were rejected by the code. The SOD distribution from E773 is shown in Figure 4.2, with the equivalent KTeV distribution overlaid. As can be seen in the figure, KTeV has much larger number of events with high-SODs, and as a result, many more events were rejected by the Level 3 requirement. The essence of the problem is that one of the wires records a hit time which is substantially delayed (by > 20 ns) from the expected drift time, which in turn causes the drift distance to the wire to be overestimated. Once the “high-SOD effect” was discovered in the data, the Level 3 software was modified to allow tracks to contain high-SOD pairs. The exact details of the allowed number of high SOD pairs is detailed in Section 6.2. All E799 data, and all E832 data subsequent to the 1996 period, was collected with Level 3 software suitably modified to allow high SOD pairs, reducing our susceptibility to the high SOD problem to second- or third-order (depending on the position and type of imperfection).

It seems surprising that such a large effect as shown in Figure 4.2 was not discovered by the online monitoring or periodic checks of the data quality. The intrinsic problem is that SOD plots such as the one shown are typically made only for pairs

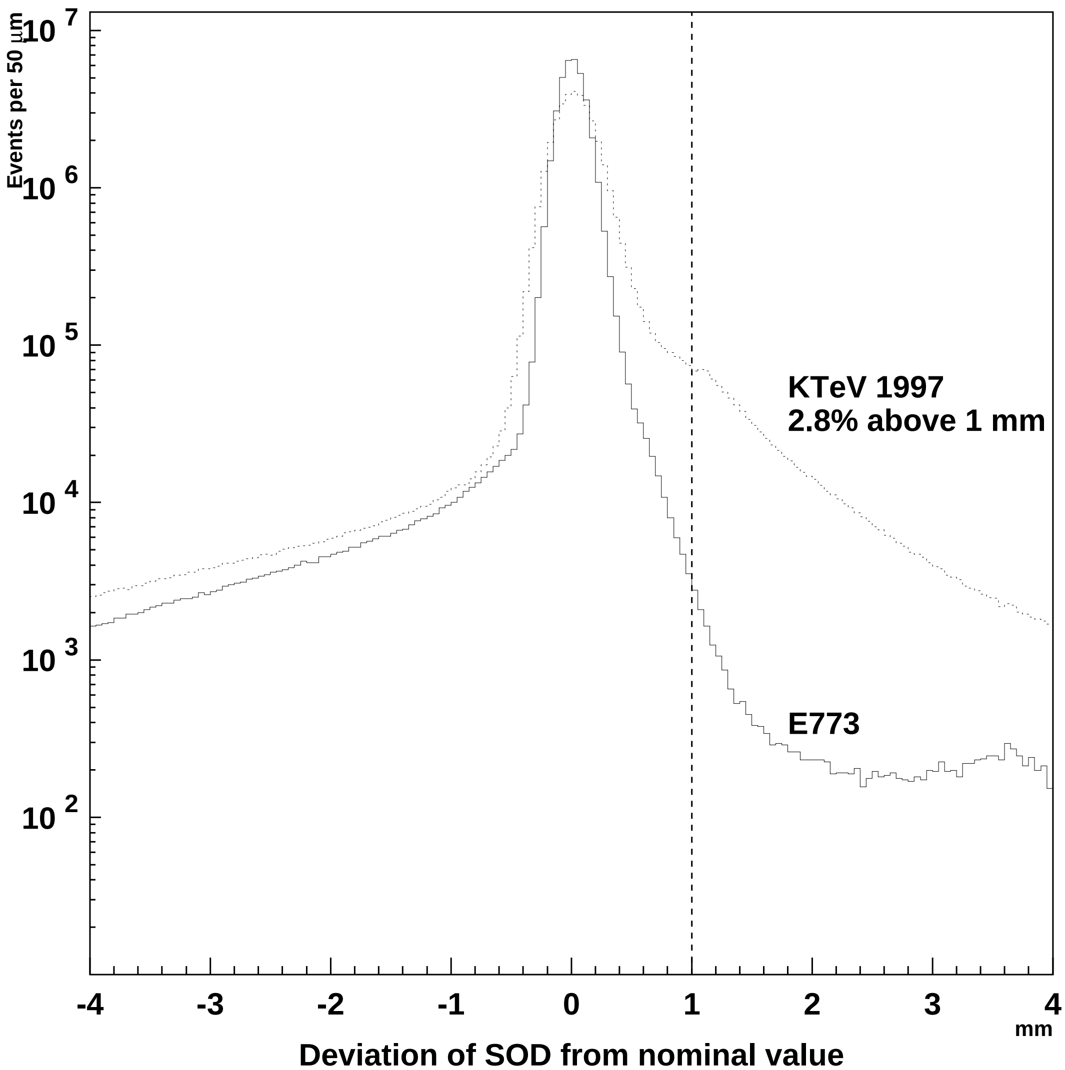

Figure 4.2: Drift chamber sum-of-distances (SOD) distribution for E773 and KTeV. The Level 3 filter software described in the text removed all hit-pairs with a SOD above 1 mm from the nominal value, shown by the vertical line.

of hits on tracks, because the DC system naturally contains pairs from out-of-time tracks that appear as low- or high-SOD pairs, and must be excluded. Unfortunately, this effect also removed any track with a high-SOD from the plot (until the modifications to the tracking code). The problem may have been discovered by scanning events by hand during the commissioning phase; however, the magnitude of the problem was smaller at the start of the run when the most attention was being paid to the detector performance, and became larger when the HV for the chambers was lowered (lowering the gain of the system), without a compensating adjustment of the discriminator thresholds.

All $\pi^+\pi^-$ data written to tape in 1996 was subjected to the requirement that no track contain a high-SOD pair. With the "random accept" Level 3 sample, it was determined that the probability for a high-SOD was less than 1%, but with 32 hits per $\pi^+\pi^-$ event, a total of 22% of otherwise-good events had been discarded. Further study indicated that the differential loss between the vacuum beam and the regenerator beam would lead to a bias on $\mathrm{Re}(\epsilon'/\epsilon)$ of $(10 \pm 5) \times 10^{-4}$, where the error is poorly determined due to the limited number of $\pi^+\pi^-$ events in the random-accept sample.

The high-SOD effect was studied in detail, using the more numerous K_{e3} and $K \to \pi^+\pi^-\pi^0$ events recorded in the random-accept sample. The goal of this study was to understand the high-SOD problem well enough to characterize it in the Monte Carlo to make a correction for the bias, and to reduce the systematic uncertainty below 1×10^{-4}. The study is detailed in Chapter 5 of Reference [2]. The main conclusion is that the high-SOD effect is caused by an inefficiency to the first ionization electron, where subsequent electrons are required to drive the analog pulse above the discriminator threshold. This model explains all the features of the high-SOD problem; the effect is largest near a sense wire, where the drift electrons are spatially separated; the effect depends strongly on the chamber HV, and hence the gas gain of the system; and the effect has a spatial dependence across the face of the drift chamber, both in the beam region, and near isolated "freckles" that have low intrinsic efficiency.

All of the effects listed above were added to the Monte Carlo, and further refinements to the simulation are detailed in Chapter 8. However, with the limited statistics available to monitor the problem, it was not clear how to set the systematic error. Therefore the decision was made *not* to use $\pi^+\pi^-$ data from 1996.

4.2.4 Miscellaneous Other Problems

The previous problems were the largest we faced with the data. We discovered and corrected a large number of minor problems during the data-taking period. A few of these problems are discussed below.

The readout for the digitized detector information is passed to the computer memories by a complicated set of paths known as "streams." KTeV has six data streams in all, four dedicated to the calorimeter and two to the rest of the detector. The data streams read out their information using a token-passing scheme called FERA, a protocol developed by LeCroy for its ADC modules of the same name. This scheme relies on front panel connections that can be quite fragile. In addition, the distance between each read-out station can be quite long, requiring use of "repeater" modules to boost both the digital detector information and the token used to read out the detector. On occasion, the token-passing scheme would become corrupted, causing the readout to hang or the data to be unreadable. In either of these cases, we stopped taking triggers, reinitialized the detector electronics, and re-enabled triggers. These losses are not a concern as a bias for the $\mathrm{Re}(\epsilon'/\epsilon)$ measurement, since they affect both beams identically.

The regenerator moves between the two K_L beams every minute. During the 1996 run, the machinery that moves the regenerator became "sticky" and we changed to moving the regenerator once every three minutes. Even with this change, the regenerator would become stuck on occasion. It is straight-forward to identify and eliminate periods when the regenerator is stuck, since its position, and the position of the Movable Absorber, is recorded into the datastream by using micro-switches. The regenerator mover was redesigned for the 1997 run, fixing this problem.

The KTeV vacuum was held to 1×10^{-5} torr during the 1997 run. During the 1996 run, the vacuum level was roughly 20 times higher, which was larger than expected but still acceptable. At the end of the 1996 run, the vacuum window was replaced and a small tear was found near the edge of the window. The vacuum system worked well after the window replacement, except during one or two failures of the diffusion pumps, at which point we had to stop the run to restore the vacuum.

4.3 Data Reduction

As mentioned at the end of the last chapter, KTeV has a very large data sample to analyze. The data stream was written to tape in the order in which it was processed, which means that the different physics samples are intermingled on every data tape. In the following sections we describe how KTeV dealt with the large datasets.

4.3.1 The Split

It is not practical to read the raw data tapes multiple times, since each user interested in the data would have to read each of the 2,645 tapes, possibly multiple times. It was decided that the data would first be "split" into separate datasets according to the Level 3 tags. For example, all $\pi^+\pi^-$ tags were copied to 277 tapes. During the split process, the calorimeter and HCC information was "squeezed" into a more compact form with no loss of information. The process simultaneously read in the nine raw data tapes (three from each of the three filter computers), and spooled each tape's data to numerous disk files based on the Level 3 trigger tag. An output computer program then combined the disk files into an output tape's worth of data, and spooled the data to the output tape.

The data split took roughly five months to complete, and wrote out 2,897 tapes. The ratio of 2,645 input to 2,897 output tapes shows that most events satisfy only one trigger.

The doubling of the number of tapes and the manpower involved in running the split convinced us to do the split "on-line" for the 1999 part of the KTeV run. Instead of writing the data from Level 3 to raw data tapes, the data were written to large arrays of disks (duplicating the first part of the split). A software "daemon" ran on the main computer, and spooled a full tape of a particular trigger to an output tape, in the format of a split output tape. The on-line split saved the KTeV experiment roughly 3,000 DLT tapes (a cost of $45,000) and six months of down-time when the data would have been split offline.

4.3.2 The Crunch

Even after the split, the dataset is too large to study effectively. We created a number of "data summary tapes" (DSTs) of the $K \to \pi^+\pi^-$ sample in different formats to allow us to access the data in a timely manner.

We reduced the size of the data by 2/3 by removing the raw CsI information. We used the calorimeter calibration to convert the number of counts in each CsI ADC channel to an energy for each CsI cluster. Our first "crunch" used the CsI calibration to remove the raw calorimeter ADC information and save only the CsI cluster information. In addition, we "packed" the information in the veto counters into a more compact format. Finally we placed extremely loose cuts on the track extrapolation at the MA and the CA, and a loose z-vertex cut of $z < 158.5$ m. At this level, the $K \to \pi^+\pi^-$ dataset fit on 19 DLT tapes. Although we had performed a preliminary drift chamber calibration, we kept all the hit-time information from the drift chambers at this stage, allowing us to revisit these 19 tapes after improving the calibration without having to return to the 277 split tapes.

Once the calibration was settled, we created second-level DSTs. These DSTs are composed only of final event information, such as the CsI cluster energies, the charged-particle trajectories, and veto counter and trigger source information. The second-level DSTs fit on disk, allowing us to process the data repeatedly. Running through the data at this stage took roughly 24 hours.

4.4 Data Samples Used for This Analysis

The KTeV group decided before the data-taking began that we would publish a result based on a sub-set of the 1996–1997 data. That analysis was finished in 1999 and published [55]. A more detailed description of the analysis can be found in

Reference [2]. This first result used $K \to \pi^0\pi^0$ data from the 1996 running period and $K \to \pi^+\pi^-$ data from the first part of the 1997 running period.

Subsequent to publishing the result, we decided to focus our attention on the remaining datasets, the 1997 data for $K \to \pi^0\pi^0$ and what we called '1997B' for the remaining part of the 1997 dataset for $K \to \pi^+\pi^-$. Due to the large value of $\mathrm{Re}(\epsilon'/\epsilon)$ that we published [55] ($(28 \pm 4.1) \times 10^{-4}$), we wanted to look at a statistically independent dataset to confirm our published result. The analysis presented in this thesis will focus on the '97B' dataset. The improvements made to the 1997B dataset will then be applied to our published dataset, and we will quote a final value of $\mathrm{Re}(\epsilon'/\epsilon)$ for the entire 1996–1997 KTeV dataset.

CHAPTER 5
SPECTROMETER CALIBRATION

5.1 Introduction

The final goal of the data analysis is the reconstruction, with high efficiency, of $\pi^+\pi^-$ events that decay within our acceptance, with the smallest amount of background contamination. To this end, we need the spectrometer to be calibrated and aligned.

To calibrate the detector, we need to be able to reconstruct tracks within the detector. However, to reconstruct tracks, we need a calibrated detector. To get around this problem, an iterative approach is used, starting with the survey positions of the detector and simplified calibrations. To separate the discussion, this chapter will focus on the calibration, while Chapter 6 will describe the intricacies of track-finding. It should be noted that these two procedures are intertwined, and are separated only for this discussion.

We will present the technique for the internal calibration of each of the drift chambers, and the alignment of the chambers to each other and to the external KTeV detector. We will then discuss the use of the chambers to locate three external systems: the VV$'$ counters, and the Mask Anti and Collar Anti veto systems.

We will not discuss the details of the calibration of the CsI. The interested reader is referred to Chapter 4 of Reference [1] for the complete details.

5.2 Survey Information

KTeV made use of detailed survey information provided by the Geodesy group at Fermilab [76]. Individual detectors were measured to determine the placement of their active elements, a process known as "fiducialization," and each detector was aligned with respect to the KTeV coordinate system before the experiment began. Additionally, the positions of the detectors within the KTeV hall were measured at intervals during the run to cross-check for movement within the hall.

The fiducialization of the drift chambers determined the center of the sets of sense wires as well as the rotation between the x and y sets of sense wires. The level of non-orthogonality between the views is consistent with the levels found when the chambers were constructed [62]. The survey information allows for precise determination of position of the pieces of the detector and is used as input for the calibration procedure.

5.3 TDC Calibration

The KTeV DC system used the LeCroy 3373 multihit Time-to-Digital-Converter (TDC) to measure drift times within the DC system. The TDC has a precision of 0.5 ns and a 15 hit depth. The TDC is operated in "common-stop" mode, where the initial DC hit starts the TDC and trigger signal stops the TDC. This mode means that hits far from the wire have small TDC values, while hits near the wire have large TDC values.

To accurately measure hit positions within the chambers, the timing information from the TDC must be converted to a distance. Due to the properties of the gas mixture, a drift speed of $\sim$ 50 μm/ns is expected; however, we require greater precision.

The TDC calibration is composed of two steps. First the relative timing of a given wire is determined. Due to variations in cable lengths, in the individual electronics channels, and in the propagation time of the global trigger signal sent to the TDCs, a given channel may "turn on" at a different TDC count value than a neighboring channel. The TDC count value corresponding to the "turn-on" of a given channel is called the $t(0)$. Second, the relationship between time and distance must be determined; this relationship is called the $x(t)$ relation.

A sample TDC distribution is shown in Figure 5.1. The common-stop edge is apparent at $\sim$ 680 TDC counts. The distribution between 550 and 680 counts is roughly flat, which as we will see indicates a uniform drift velocity.

5.3.1 *Timing Offsets*

To correctly map a TDC value of a hit in the DC to a drift distance in the cell, it is first necessary to know the TDC value for each wire that corresponds to zero drift distance. The $t(0)$ for this procedure is defined as the TDC count at the 50% point between the zero event point and the "plateau" of the TDC distribution. This corresponds to $\sim$ 680 TDC counts.

There are several mechanisms that could introduce wire-to-wire variations in the

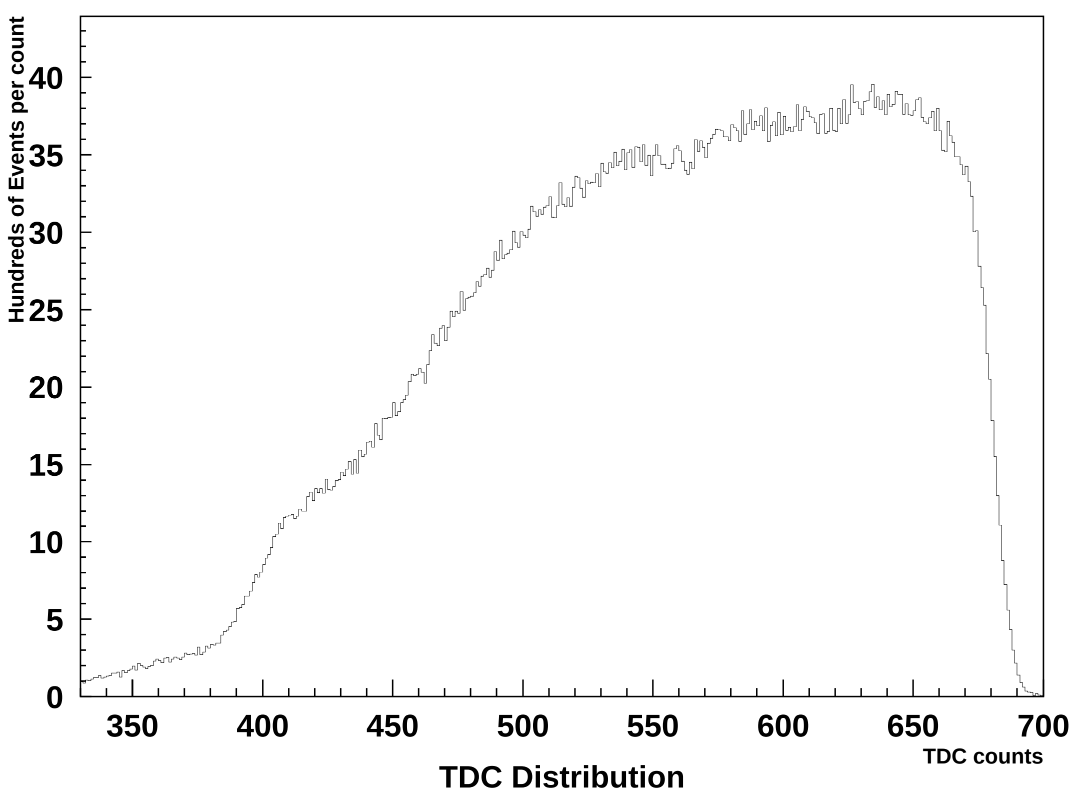

Figure 5.1: Sample TDC distribution.

zero-drift TDC time. For example; varying cable lengths between the pre-amp, post-amp or the TDC; differing pulse heights; channel-to-channel or module-to-module variations in the amplifiers or the TDCs; and finally, the common stop delivered to the TDCs. In order to remove such variations, so that we can combine each wire in a plane to make the $x(t)$s, a simple technique is used to determine each wire's $t(0)$. Each TDC module has 16 inputs, so wires from the DC system are grouped into 16-wire bundles through the pre-amplifier through the post-amplified to the TDC. The procedure to calculate the wire $t(0)$ is to assume that the TDC distribution of the entire 16 wire group is representative of the parent TDC distribution for a wire. The parent distribution is fit in the region around the estimated midpoint of the distribution to determine its global $t(0)$. By sliding the individual wire distributions past the parent distribution and constructing the Kolmogorov-Smirnov test (KS), one can determine each wire's offset with respect to the assumed parent distribution's offset. The distributions were moved in 1 TDC count steps, and the K-S distribution was fit to a Gaussian to find the wire offset to 0.25 TDC counts. For wires that did not have enough statistics for this procedure (typically wires at the edge of a chamber), the offset of the parent distribution was used. The process was iterated, since the assumed parent distribution depends on the determined $t(0)$s. Figure 5.2 shows the final results, which would converge after two to three iterations. All the TDC distributions have been corrected for their individual $t(0)$ offsets and start at 679 TDC counts.

Figure 5.3 shows the $\langle RF - L1 \rangle$ timing versus run number. The largest contributions to shifts in the TDC $t(0)$s are global shifts in the Level 1 common stop delivered to the TDCs by the trigger and global shifts in the beam timing with respect to the RF delivered to KTeV from the Tevatron. In the 1996 run, there were shifts in the timing of 1–2 ns. For the 1997 run, these shifts were reduced to < 0.5 ns. These shifts directly influence the drift times measured in the chambers, and hence directly affect the positions reconstructed in the chambers, through the $x(t)$ map (see the following section). The SOD will be shifted from the nominal value by amounts of order 25–50 μm, since the average drift speed is ~ 50 μm/ns. The $t(0)$s were recalculated after large shifts in the timing. In the intervening runs,

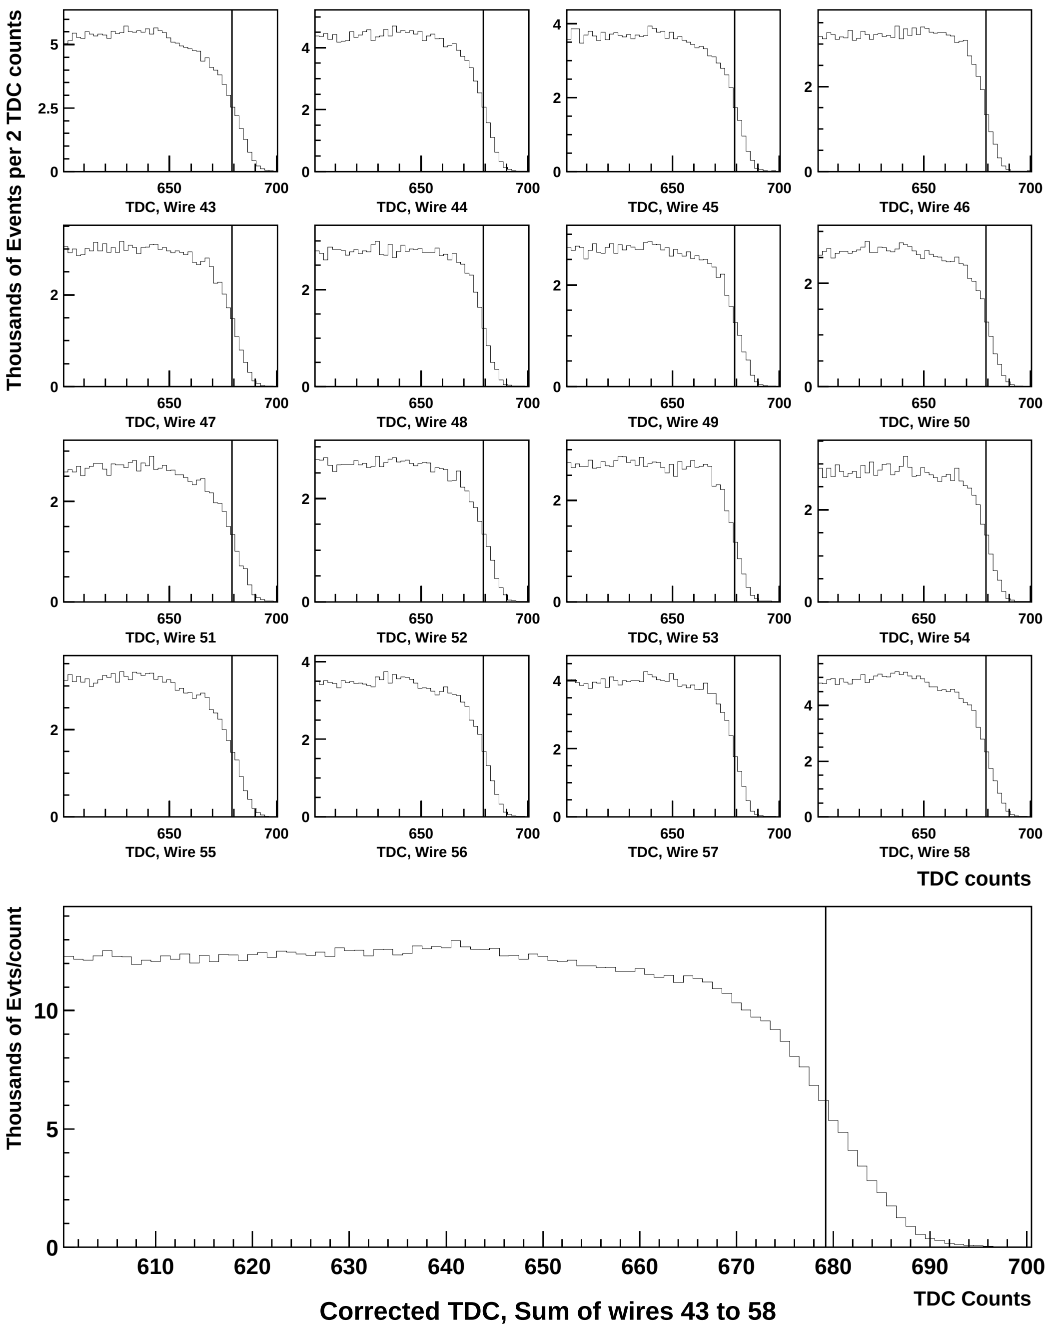

Figure 5.2: Corrected TDC distribution between 600–700 TDC counts for wires 43–58 for plane 1 during Run 9353, and the sum of the 16 wires, after correcting for the individual $t(0)$s. The line shows where the distribution should start. Note that the $t(0)$s are all at 679 TDC counts.

the difference in the $\langle RF - L1 \rangle$ timing of the current run from the run for which the $t(0)$s were generated was used to correct the $t(0)$s.

5.3.2 The $x(t)$ Relation

Once the wire-to-wire timing differences are removed by the previous step, an entire chamber plane can be combined to determine the relationship between drift distances and drift times, commonly referred to as the $x(t)$ relation. This should be a valid procedure, because the parameters that affect this relation, such as cell geometry, gas mixture, and voltage, should be identical for each wire in the chamber. As was detailed in Section 4.2.3, there are position-dependent effects within the chamber; however, these have been ignored for the most part in the determination of the $x(t)$ relation.

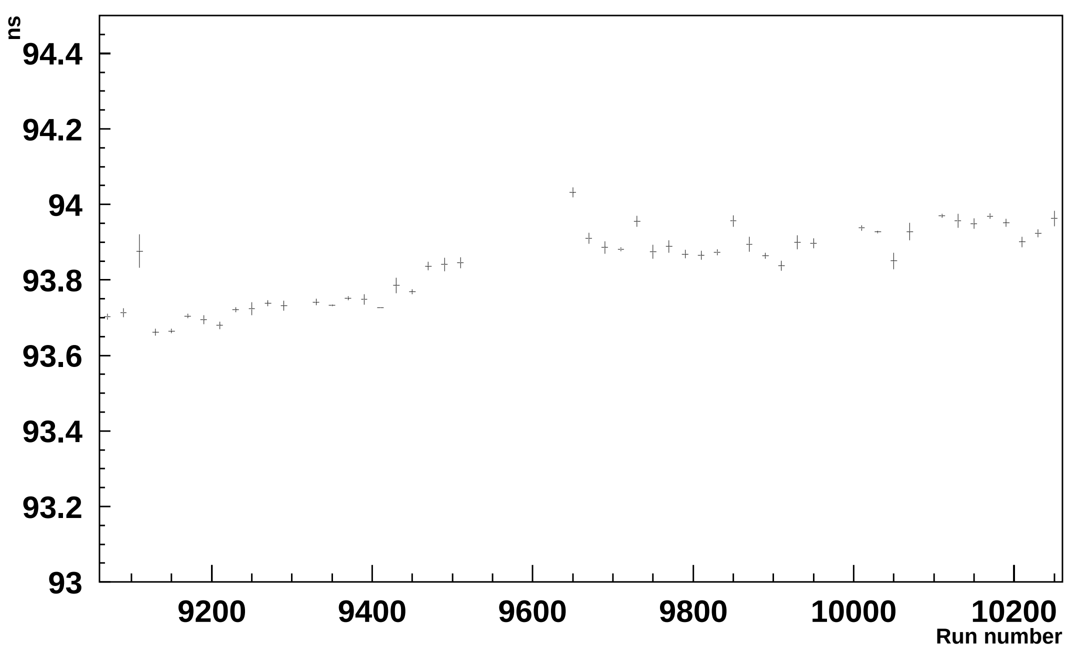

Figure 5.3: The average time between the radio-frequency (RF) strobe from the Tevatron and the Level 1 trigger signal, versus run number.

The procedure relies on the fact that, on average, the illumination of tracks across each drift cell is uniform. One only need assume that the illumination is a slowly-varying function of position in the cell, because all cells within the chamber are summed into one cell, and any local deficit to one side of a cell in a given area of the chamber will be canceled by the abundance on the other side. Given a uniform illumination across the unit drift cell, one can convert the distribution of observed times to a lookup map converting time to distance. The method assumes that the earliest TDC hit comes from right next to the sense wire, the latest TDC hit comes from the edge of the cell, and that distribution is ordered such that larger times correspond to larger distances. The remaining concern is missing hits due to chamber inefficiencies; the method for accounting for missing hits is discussed below.

We used the two-track K_{e3} sample to do the $x(t)$ calibration. Two good tracks that matched to clusters in the CsI calorimeter are required. The tracks are required to point to the trigger counters, and either an in-time latch hit or an in-time TDC hit was required for the counter in question. Approximately 500,000 events are used to determine the $x(t)$ relation in each run.

The TDC distribution for hits on tracks are collected for each of the 16 drift chamber sense planes. The TDC values are corrected for propagation along the wire. Missing hits pose a problem if the inefficiency of the chamber is not uniform across the drift cell. The method to correct for missing hits is to track the number of missing hits in terms of the hit on the complementary plane. The four panels of Figure 5.4 show the steps of this correction; first we are concerned that the raw TDC distribution may have a local deficit of hits in a given region due to non-uniform inefficiencies in the cell (Panel (a)). We track the number of missing hits as a function of the TDC value of the complementary plane (Panel (b)). Panel (b) shows that most of the missing hits are near the complementary plane, or *far* from the plane of interest. Panel (c) shows, for good events, the relation between hit times on the plane of interest and the complementary plane.[1] To correct for the

[1] This plot is known as the "banana" plot because of its shape, and it is from this distribution that the Banana Trigger Boards get their name.

missing hits, we take the slice corresponding to the complementary hit time to get the distribution of hits for the plane of interest. This "slice" is then weighted by the number of missing hits and added to the raw distribution. This is done for all the missing hit bins in the complementary plane; the corrected TDC distribution is shown in Panel (d). In this way, the data are used to correct for inefficiencies in the $x(t)$ map generation.

Once the corrected TDC time distribution is complete, the $x(t)$ relation is given by

$$x(t) = d_{cell} \times \frac{\sum_{t_0}^{t} N(t)}{\sum_{t_0}^{t_m} N(t)}, \tag{5.1}$$

where t_0 and t_m are the minimum and maximum drift times, and d_{cell} is the maximum drift distance of 6.35 mm, and we sum over the number of events $N(t)$ in a 0.5 ns TDC bin. Note that if the maximum drift time t_m comes from hits drifting in from outside the cell, the $x(t)$ map will be biased to smaller drift distances. Figure 5.5b shows a typical $x(t)$ map for the drift chambers. Because the track-finding algorithm depends weakly on the $x(t)$ used, the $x(t)$–finding process was iterated until the $x(t)$ map converged.

5.3.3 Mean Sum-of-Distances, Resolutions, and Final Corrections

A typical sum-of-distance plot is shown in Figure 5.6, before the last correction is applied. Certain distinguishing features are present which deserve comment. First, the fitted peak of the distribution is 6.32 mm, and not the size of the drift chamber cell (6.35 mm). The peak was fit to a Gaussian distribution in the region 6.10–6.5 mm. The mean of the distribution, $\left\langle SOD \right\rangle = 6.344$ mm, is closer to the cell size. Both of these features are typical of all $x(t)$ maps generated. This is a direct result of the fact that the SOD distribution is asymmetric, and that we have a "high-SOD" problem. Currently, all known effects bias the $x(t)$ map to *smaller* mean SODs.

A final correction was applied to the $x(t)$ maps to attempt to move the mean SOD closer to the wire separation distance. It is assumed that the bias toward

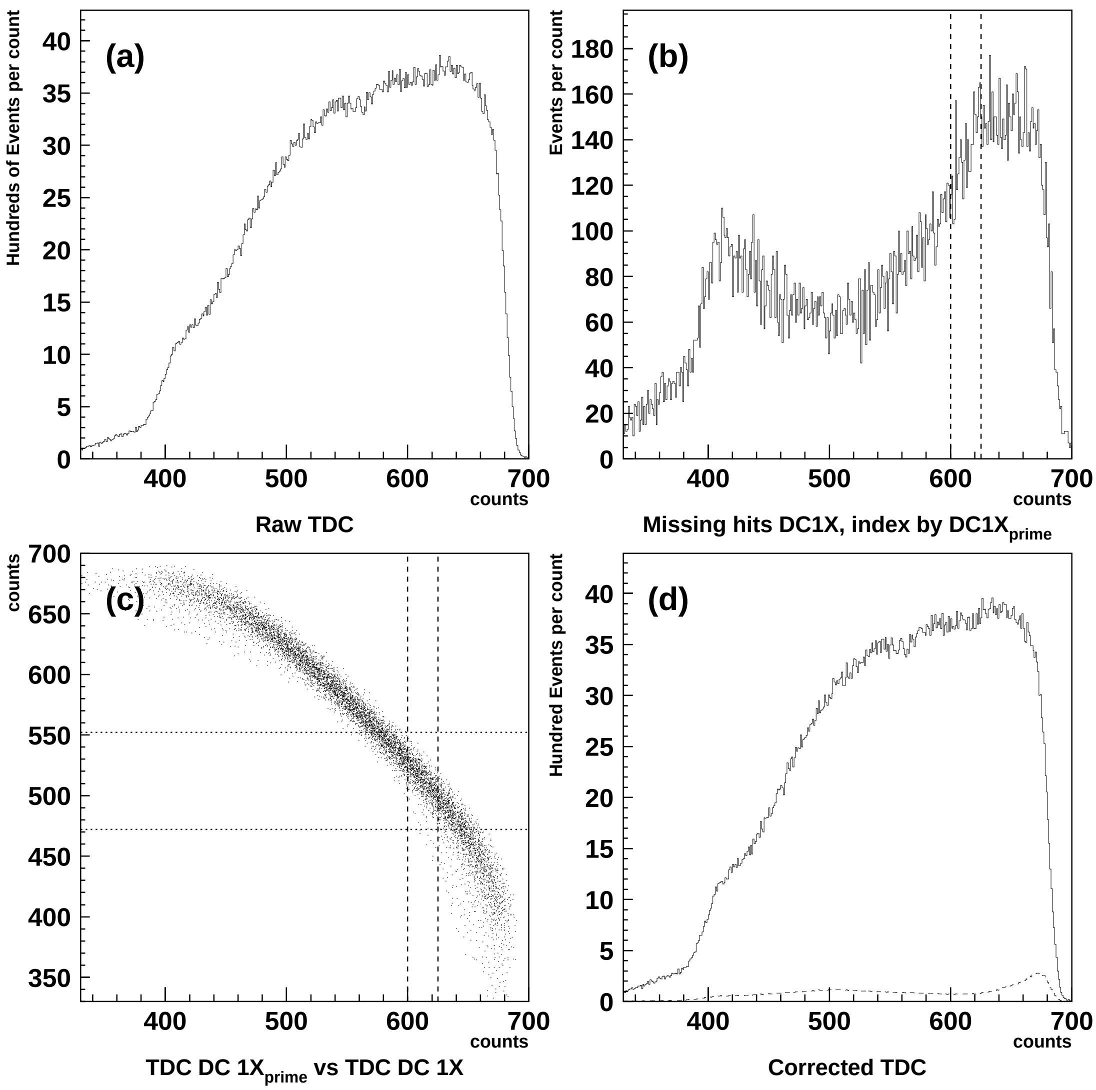

Figure 5.4: The method used to generate the $x(t)$ maps. Panel (a) shows the raw TDC count distribution for a given plane. Panel (b) shows the number of missing hits, as indexed by the TDC count of the complementary plane. Panel (c) shows the relation of good hits between the plane of interest and the complementary plane. Panel (d) shows the adjusted TDC distribution, once corrections using Panels (b) and (c) are made, where the dotted histogram shows the correction made for missing hits.

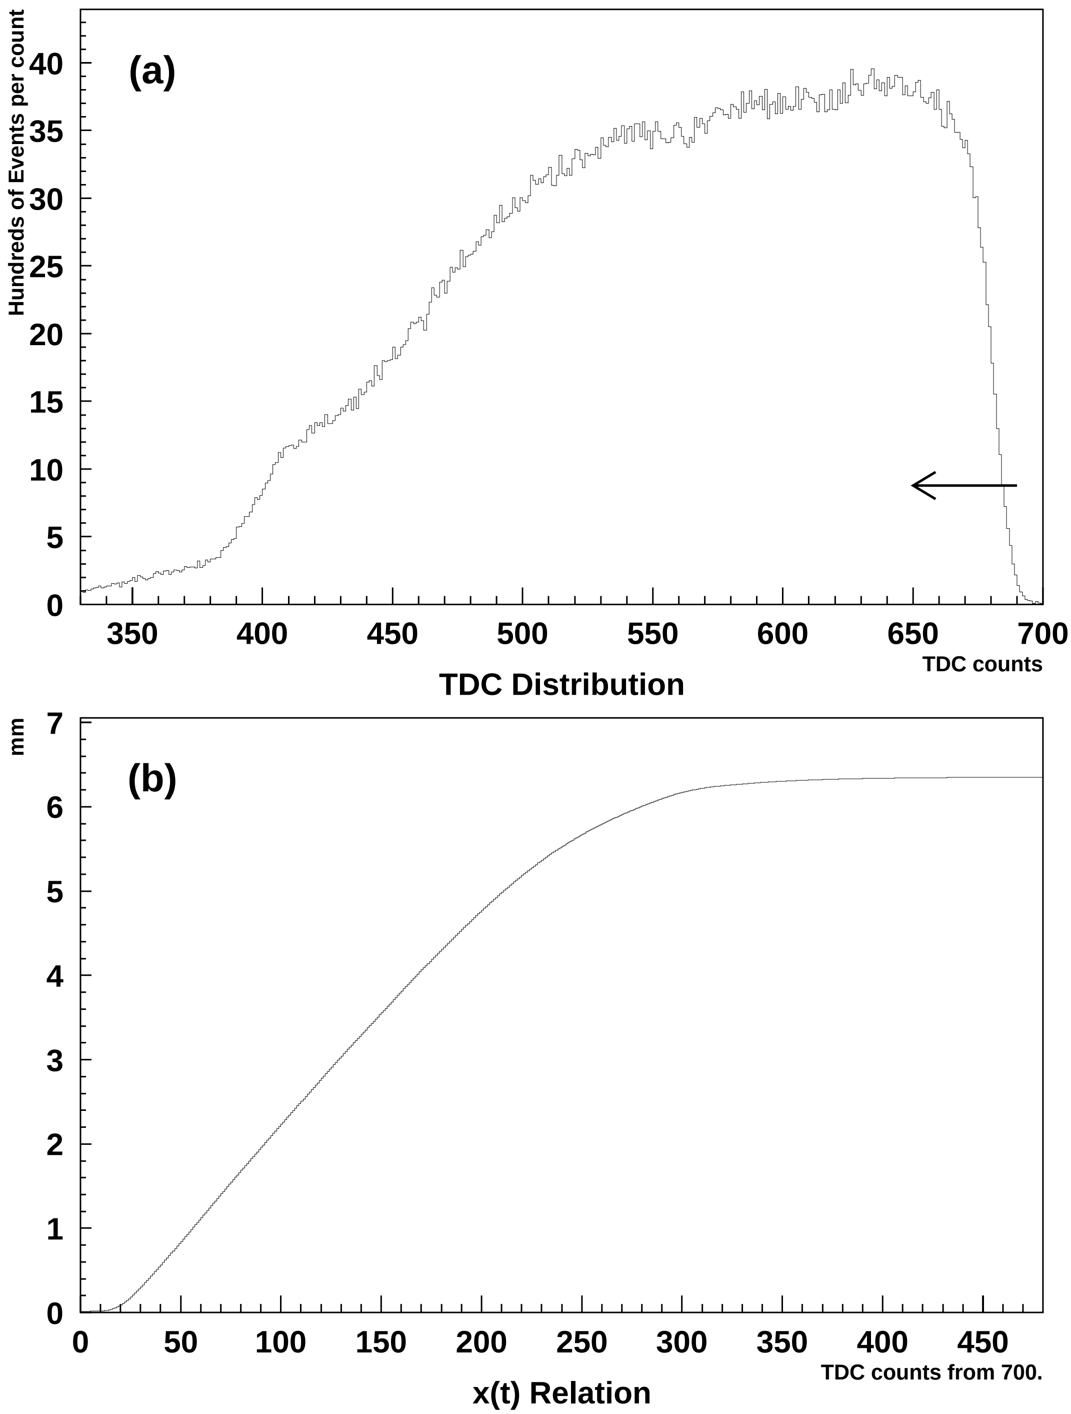

Figure 5.5: TDC distribution for all wires for plane 1 for Run 9244, used to make the $x(t)$ map (Panel (a)). The arrow shows the direction of integration to remove the "common-stop." Panel (b) shows the $x(t)$ map derived from the TDC distribution.

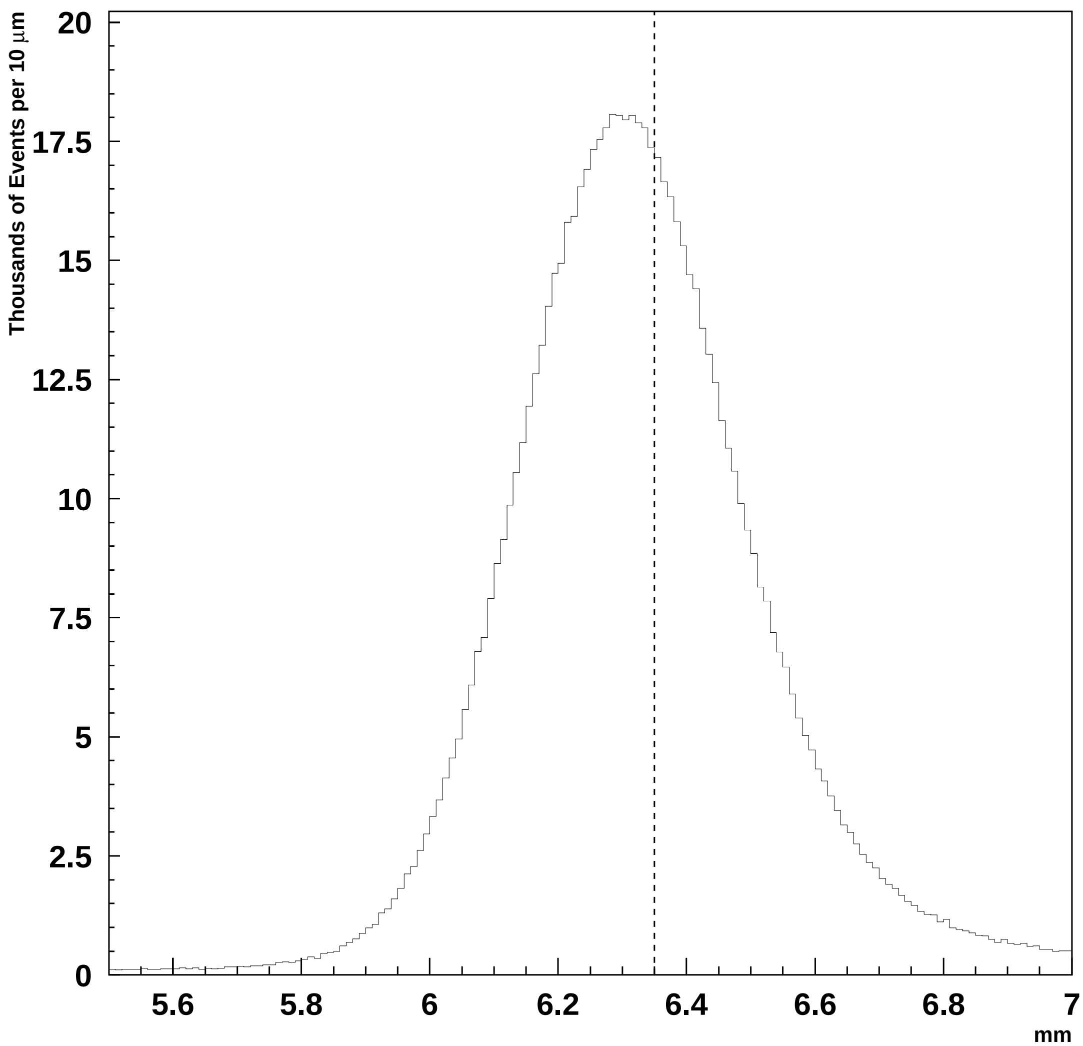

Figure 5.6: Typical Sum-of-Distance for two hits on adjacent wires in a unprimed-primed plane pair. The sum should be the constructed cell size, 6.35 mm, after accounting for angle and resolution effects. The dotted line shows the mean cell size (6.35 mm), indicating that the peak is shifted low.

lower SODs occurs due to effects near the wire, and that the first observed drift electron does not correspond to zero-drift distance; in other words, that the $t(0)$ in fact corresponds to some non-zero distance. The correction is applied by adding to the $x(t)$ maps by starting them a small distance from zero, such that the mean of the SOD distribution is fixed to the cell width. This correction is determined iteratively, and the final correction is of the order of 25 μm. It is important to note that the correction is derived from events with tracks outside the beam region, since this region has other difficulties that will be discussed later.

The $x(t)$ calibrations were performed in time periods corresponding to one to two days. The mean SOD is calibrated to be stable to within ± 10 μm of the cell size. This frequency of calibration improved the stability of DC reconstruction, and marks an improvement over the calibration used in the published result [2, 55].

The single-hit resolution of the chamber system was typically in the range of 100 μm, after correcting for the time delay in propagating the signal along the sense wire, which is not as good as was obtained in E773 [33]. One can clearly see this in the comparison in Figure 4.2.

5.4 Straight-Through Alignment

Once each plane of the drift chamber is calibrated, we can proceed to aligning the four drift chambers within the KTeV coordinate system. We first align the chambers internally using straight tracks, and then align the chambers to the rest of the KTeV detector by pointing tracks to the target and the calorimeter.

The first step is to align drift chambers 2 and 3 in a coordinate system defined by chambers 1 and 4. In this step, chambers 1 and 4 are assumed to define an orthogonal coordinate system. Any misalignment introduced by this assumption will be removed in the next stage of the alignment process.

Muon runs were taken approximately once a day. These runs were performed with the beam-stop in, such that only muons passed into the detector, and with the analysis magnet turned off. Each run had $\sim$ 200,000 events with a one hit coincidence between the trigger hodoscopes and the MU2 counter bank written to tape. In that sample, there were $\sim$ 60,000 single track events. With straight muon tracks, a straight line can be fit to the points in chamber 1 and 4. The difference between the points in chambers 2 and 3 and the fitted line position in one dimension is plotted as a function of the other dimension. For example, Δx is plotted as a function of y, as shown in Figure 5.7. The residuals are stored in a 2-dimensional histogram. Because of residual tails in the scatter plot of residuals versus track coordinate, the two-dimensional plot is not fit directly to a straight line. Instead, the histogram is "sliced" into bins in the abscissa, and each bin is fit to a Gaussian. The Gaussian fit is insensitive to the tails of the distribution. The means of the Gaussians are then fit to a straight line. The intercept of the line gives the offset of the plane, and the slope gives the rotation of the plane. General track quality cuts were made; only events with single tracks were considered, the track was required to have all hits used to make the track, and the track was required to point to a hit trigger bank counter. Due to concerns about residual magnetic fields at the level of 0.5 Gauss, certain cuts were made to select high momentum muons. The track projection was required to match to a hit counter in all three muon banks.

5.5 Corkscrew Alignment

There is one rotation that cannot be determined from single track events. The chamber 1–4 coordinate system is assumed to be orthogonal when determining the chamber 2 and 3 offsets and rotations. It is possible that there is a rotation between chamber 1 and chamber 4. This rotation will introduce a rotation in chambers 2 and 3, because the offsets and rotations for chambers 2 and 3 were determined in a non-orthogonal reference frame. For example, if chamber 4 is rotated by angle ϕ with respect to chamber 1, and chambers 1, 2, and 3 are square, and the drift chambers are equally spaced, chamber 2 will appear to be rotated by $-\phi/3$ and chamber 3 will appear to be rotated by $-2\phi/3$. A rotation about the z–axis, proportional to the z position of the chamber, is introduced. This rotation can be determined from two-track decays coming from a common vertex. Because the two tracks contain a common point (the vertex), they define a plane. If the two tracks are reconstructed

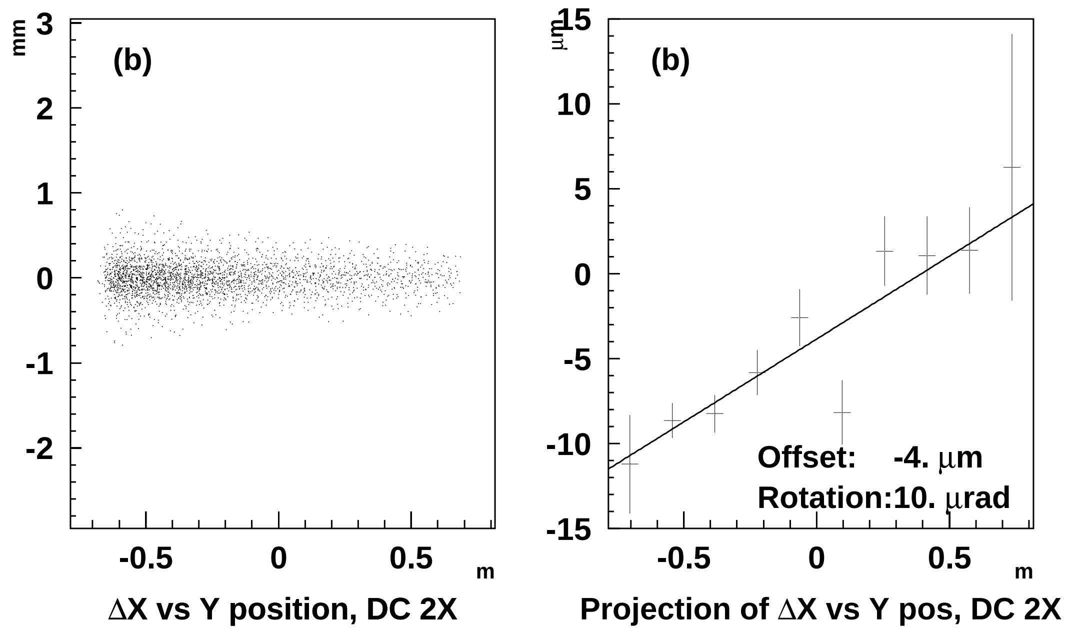

Figure 5.7: Δx residuals versus y track position in Chamber 2 for muons, Run 9059.

in chambers 1 and 2, and determined to be nonplanar, we can determine that there is a residual rotation between chambers 1 and 2. This angle is determined between chambers 1 and 2, and the correction is applied to chambers 2, 3, and 4.

Let $\vec{r}_{1(2)}$ be the line joining the hits of track 1 and track 2 in chamber 1(2) (see Figure 5.8). If the two tracks define a plane, then $\vec{r}_1$ and $\vec{r}_2$ lie in the same plane, and their cross product is zero, $\vec{r}_1 \times \vec{r}_2 = 0$. However, if there is a corkscrew rotation, then

$$\vec{r}_1 \times \vec{r}_2 = |\vec{r}_1|\,|\vec{r}_2| \sin\phi \simeq |\vec{r}_1|\,|\vec{r}_2|\,\phi, \tag{5.2}$$

where ϕ is the corkscrew rotation between chambers 1 and 2. By plotting $|\vec{r}_1|\,|\vec{r}_2|$ versus $\vec{r}_1 \times \vec{r}_2$ one can determine the angle ϕ. Then we can determine the angle per unit z distance. Figure 5.9 shows the distribution of $|\vec{r}_1|\,|\vec{r}_2|$ versus $\vec{r}_1 \times \vec{r}_2$. Instead of fitting this distribution directly to a line to extrapolate ϕ, we again slice the distribution into bins and fit each bin to a Gaussian. Then the means of the fits are in turn fit to a straight line. The slope of the line gives the corkscrew angle between chambers 1 and 2. The corkscrew angle is propagated to chambers 2, 3 and 4 and removed. Once the first large corkscrew rotation ($\sim 450\ \mu$rads) was removed in the initial calibration, corkscrew rotations for subsequent runs were of order $\sim 10\ \mu$rads.

The cuts on events for the corkscrew alignment were relatively simple. The $K \to \pi^{\pm} e^{\mp} \nu_e$ sample is used as the input, and events are required to have two good tracks matched to clusters and making a vertex. No extra tracks are allowed. Also, the two tracks are required to be missing no hits in the first two drift chambers.

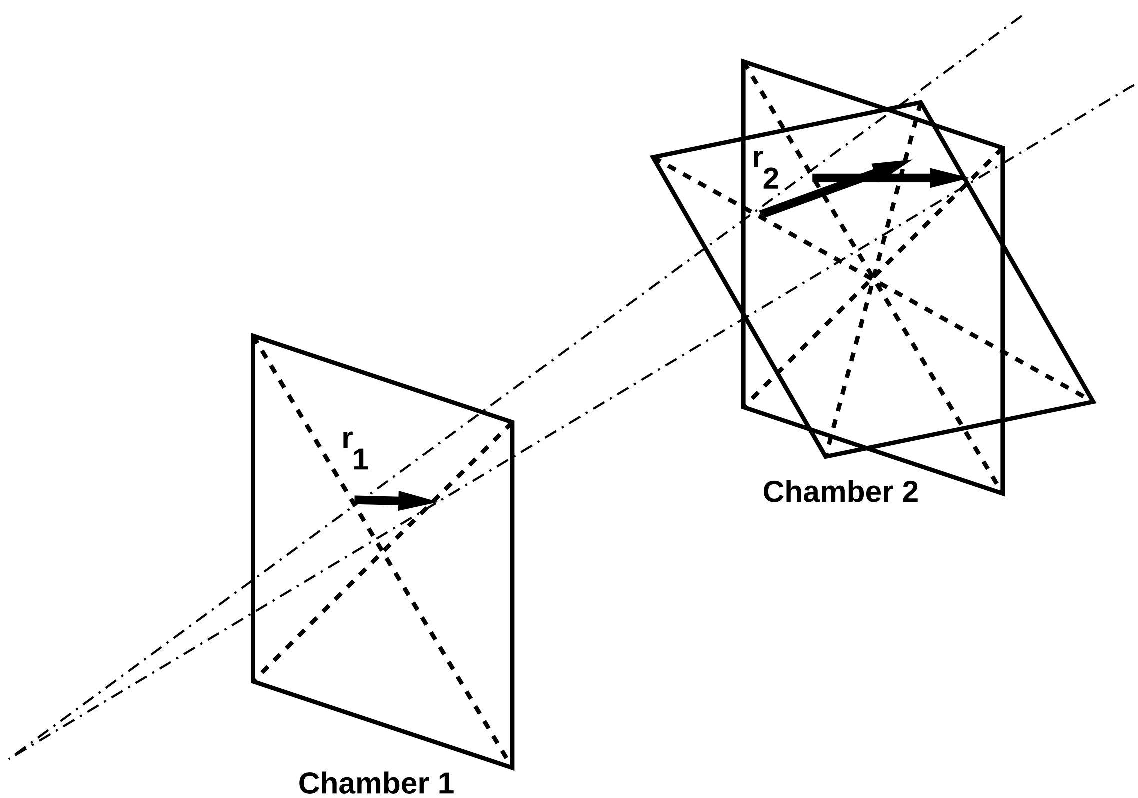

Figure 5.8: Diagram of the Corkscrew rotation as seen in Chambers 1 and 2 if Chamber 4 is rotated with respect to Chamber 1. The reconstructed vectors $\vec{r}_1$ and $\vec{r}_2$ measure this rotation. Figure courtesy of Changqing Qiao.

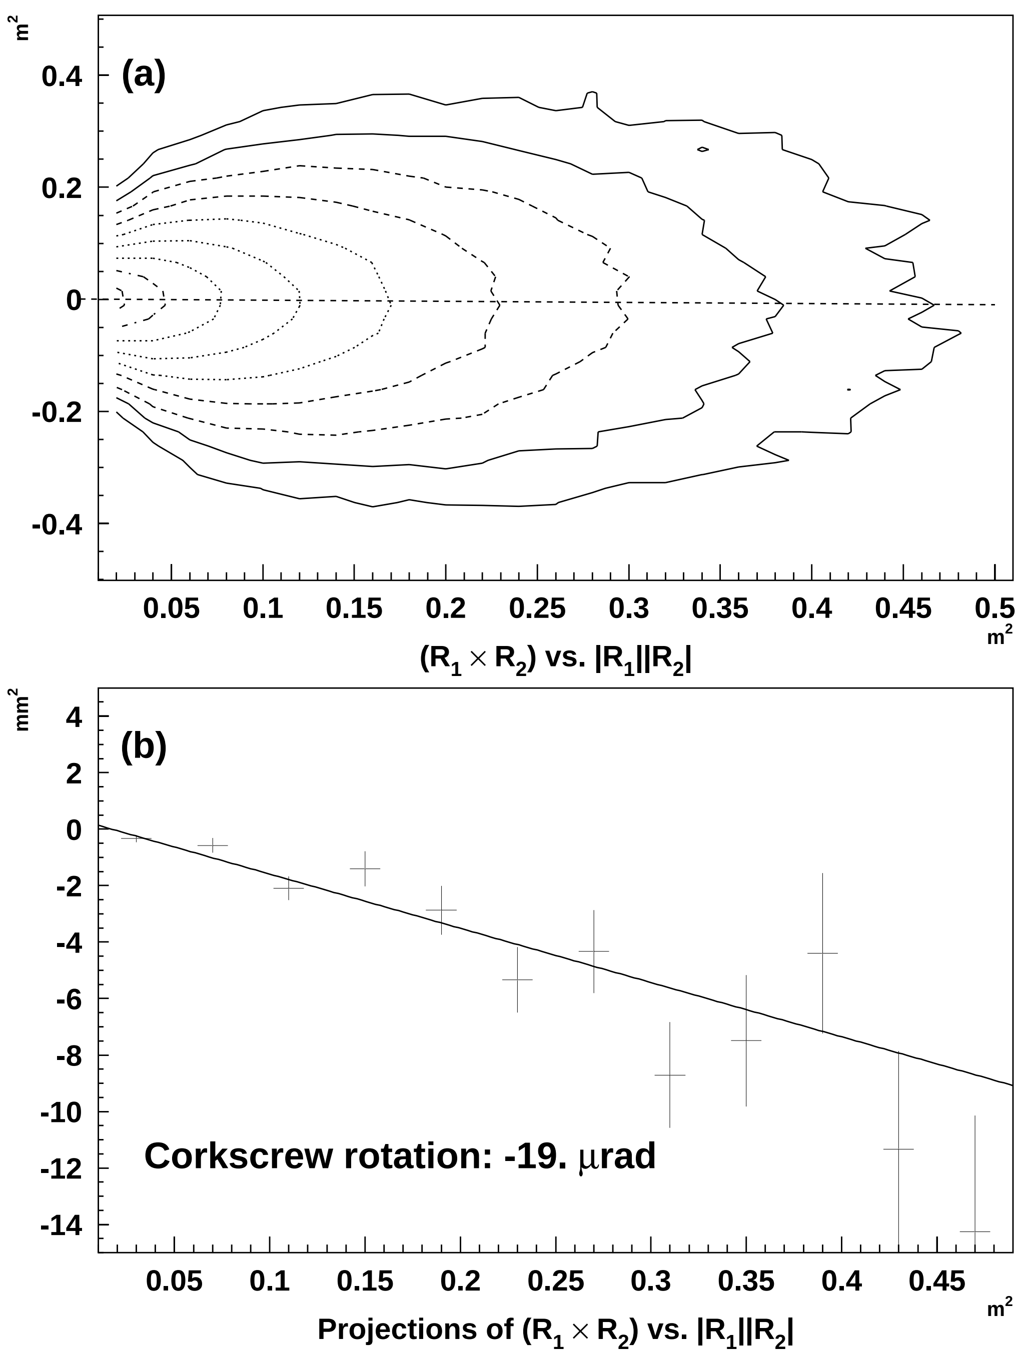

Figure 5.9: The variation of $(\vec{r}_1 \times \vec{r}_2)$ with $|\vec{r}_1|\,|\vec{r}_2|$, used to measure the corkscrew rotation of the drift chamber, for run 9246. Panel (a) shows the logarithmic-scale contour plot. Most events are at small $|\vec{r}_1|\,|\vec{r}_2|$, but there is a clear trend. The bottom panel shows the means and errors on the means of Gaussian fits to projections of the first plot sliced into bands in the abscissa. The slope gives the rotation.

5.6 Alignment to the Target and CsI Calorimeter

The final step of the alignment was to take the orthogonal drift chamber system and align it to the external KTeV coordinate system. This was accomplished by pointing tracks to two external elements: the target and the CsI calorimeter.

The target position can be determined by reconstructing decays and projecting the kaon momentum back to the target z-position. To insure that we select decays where the reconstructed momentum will point to the target, we used $K \to \pi^+\pi^-$ decays from the vacuum beam. One could also use $\Lambda \to p\pi$ decays, either as a cross-check, or as an independent measurement, since $K \to \pi^+\pi^-$ is a signal mode. The procedure was iterated since the target position is used to calculate the transverse momentum of the decay, which is then used to select $K \to \pi^+\pi^-$ decays. Figure 5.10 shows the position of the total momentum in the x–y plane at $z = 0$, and the projections to the x and y axes. The x position was aligned to be $x = 0.0$ m, to within an uncertainty of order 10 μm. The y position was aligned to the survey position of $y = -305$ μm. The position of the CsI calorimeter in the drift chamber system was determined by pointing electrons from $K \to \pi^\pm e^\mp \nu_e$ decays to the calorimeter. The electrons were identified by $E/p > 0.93$. The residuals of the projected tracks and the reconstructed cluster centers were studied as a function of position in the CsI. Figure 5.11 shows the distribution of the x track–cluster difference over the entire CsI calorimeter. After the alignment, the mean of this distribution is less than 1 ± 10 μm. In a fashion similar to that described above for the muon alignment, a common offset and rotation of the drift chambers with respect to the CsI can be determined. Corrections are made to the offsets of each of the chambers such that the target and the CsI will reconstruct to the surveyed positions. Thus the drift chambers are aligned in the external KTeV coordinate system. In aligning to the CsI, some peculiarities of individual crystals' responses were ignored. For example, it is known that currently the center blocks of the CsI have a bias in their x–position cluster reconstruction, which can be clearly seen in Figure 5.11b as a dip in the Δx track–cluster separation versus y. These events were not separated out

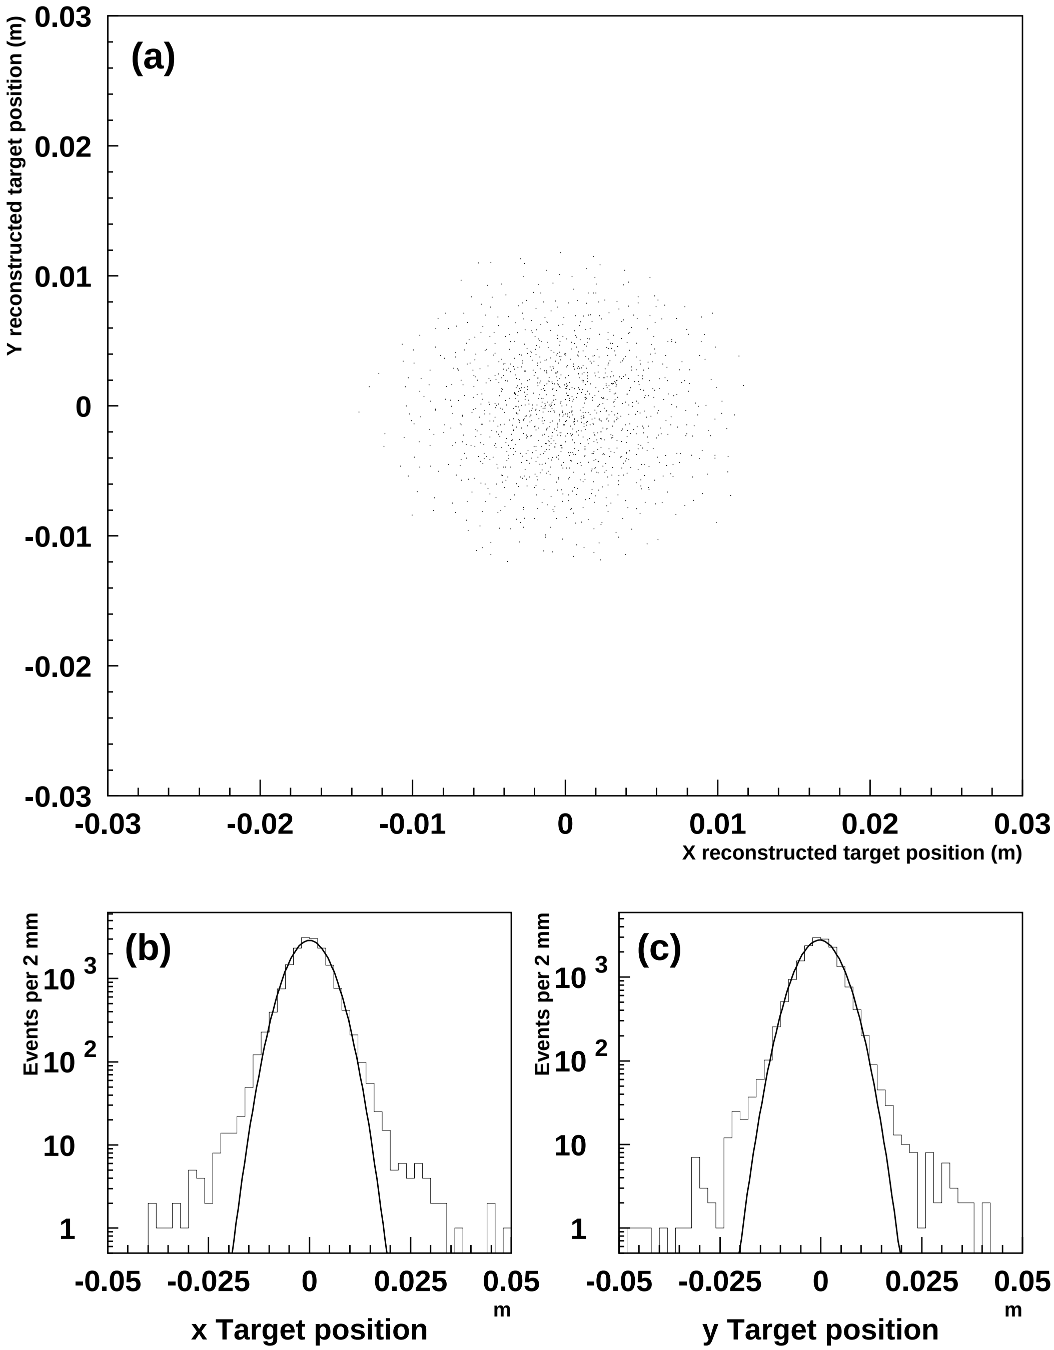

Figure 5.10: The position of the target as imaged in the drift chambers, by pointing the total $\pi^+\pi^-$ momentum vector back to $z = 0$ m. The upper panel is the two-dimensional image, while the lower two panels are the x and y projections of the upper panel.

from Figure 5.11. The rotation in x of the chambers with respect to the CsI was used to correct the rotations of both views of the chambers.

Finally, in this stage of the alignment, we looked for systematic biases that may have been introduced in the muon alignment due to stray magnetic fields in the analysis magnet. We are concerned that the "straight-through" muons may have been bent by a slight magnetic field left in the analysis magnet. We assume that this corresponds to an average bend of $\delta\phi$, introducing offsets of $\delta_{2,3}$ in chambers 2 and 3.

Two tools were useful in determining how well chambers 2 and 3 were positioned. The quantity x_{offmag} is the difference in the positions at the magnet as extrapolated from the upstream and downstream track segments to the middle of the magnet ("the bend plane"). The mean of the x_{offmag} distribution should be 0 to within 1–2 μm. $\langle x_{\text{offmag}} \rangle$ is related to the offsets of chambers 1 and 2 in the following

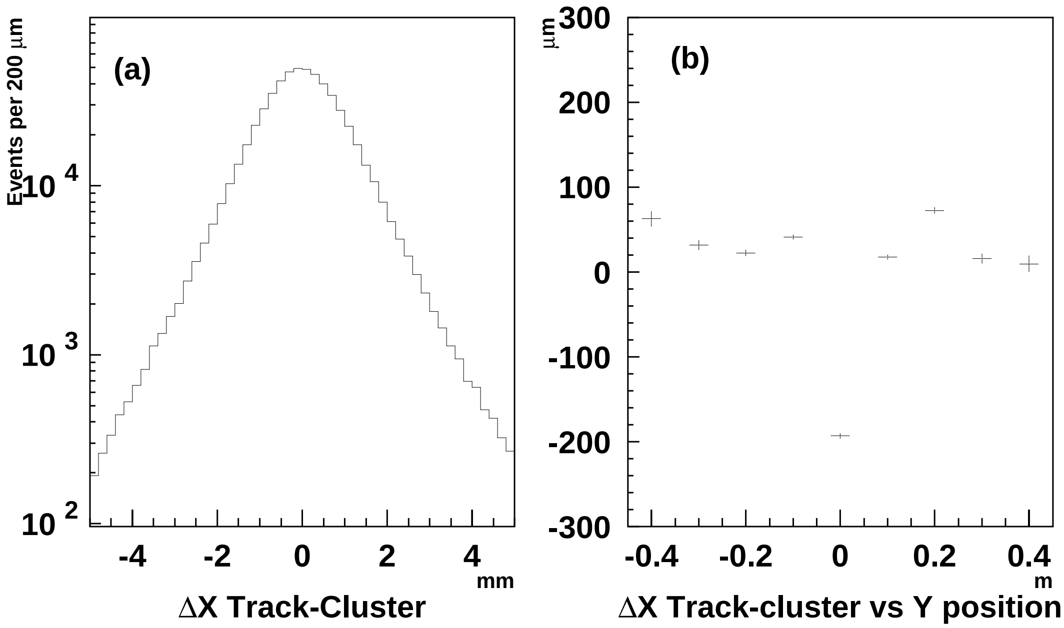

Figure 5.11: Δx track–cluster separation over the entire CsI calorimeter for electrons in $K \to \pi^{\pm} e^{\mp} \nu_e$ decays. The dip at $y = 0$ is a known problem with the cluster reconstruction.

manner:

$$\langle x_{\text{offmag}} \rangle = \delta_3 \frac{z_4 - z_{\text{magnet}}}{z_4 - z_3} - \delta_2 \frac{z_{\text{magnet}} - z_1}{z_2 - z_1}, \tag{5.3}$$

where z_i ($i = 1, 2, 3, 4$, magnet) are the z–positions of the chambers and the magnet, and δ_i, ($i = 2, 3$) are the x-offsets of chambers 2 and 3. The distribution of x_{offmag} after all calibrations is shown in Figure 5.12. For the average alignment job, the mean of x_{offmag} was of order $\sim 50\ \mu$m, before this stage in the process.

The other useful quantity is the slope of $m_{\pi^+\pi^-}$ as a function of the difference of the momentum of the left–bending and right–bending tracks in $K \to \pi^+\pi^-$ decays. Assuming that the muon alignment introduced an angle $\delta\phi$, due to the residual magnetic field, the left-bending track would have a measured angle of $\phi_1 + \delta\phi$, and the right-bending track would have a measured angle of $\phi_2 - \delta\phi$. This relates to the

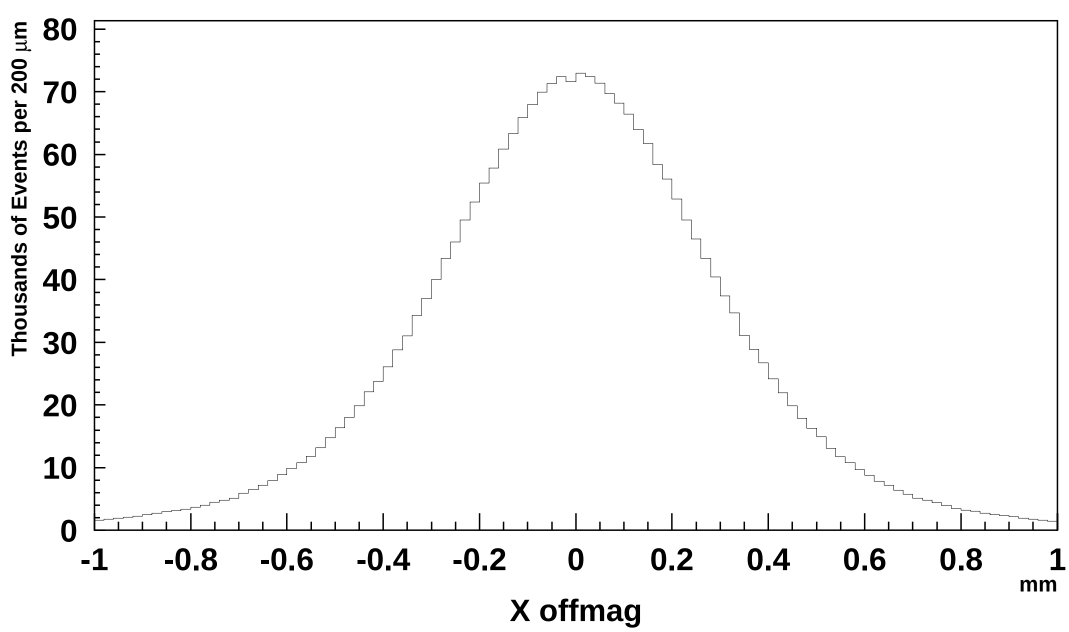

Figure 5.12: x_{offmag} for Run 9278, after the calibration procedure.

mass in the following way.

$$\begin{aligned} m_K^{2\prime} &= 2m_\pi^2 + 2p_1'p_2'(1-\cos\theta) \\ &= 2m_\pi^2 + 2k^2\frac{1}{(\phi_1+\delta\phi)}\frac{1}{(\phi_2-\delta\phi)}(1-\cos\theta) \\ m_K' &\simeq m_K\left[1+\frac{1}{2}\left(1-\frac{2m_\pi^2}{m_K^2}\right)\frac{\Delta p}{k}\delta\phi\right] \end{aligned} \tag{5.4}$$

where k is the magnet kick ($\simeq 410$ MeV/c), $\cos\theta$ is the vertex opening angle, primed quantities are the measured quantities, unprimed quantities are the true quantities, and $\Delta p \equiv p_2 - p_1$. The slope of the measured mass with respect to the difference of the momenta is related to the angle $\delta\phi$ which is in turn related to the average of the offsets of chambers 2 and 3 in the following manner

$$\delta\phi = \frac{\delta_2}{z_2 - z_1} + \frac{\delta_3}{z_4 - z_3}. \tag{5.5}$$

Figure 5.13 shows the kaon mass as a function of the difference of the momenta of the two tracks after the final external alignment was done. Using the above two relations, we can solve for the offsets and correct them. This was done at this stage of the alignment. The corrections to the chamber 2 and 3 offsets using Equations 5.3 and 5.5 are less than 5 μm.

After the entire calibration has been completed, it is worthwhile to note how stable the chamber system was as a function of time. Figure 5.14 shows the variation of the offset of Chamber 2 as a function of run number in the 1997 run. It is clear that after we correct for the above mentioned effects, on average, the chamber moves by less than 50 μm between calibrations. There is a noticeable step after the shutdown period between runs 9500 and 9720, when work was done on the chambers, but within contiguous periods the chamber offsets are quite stable in the x view. The y view movement of Chamber 1 may be related to the movement of the window shutter used to protect the vacuum window during accesses to the KTeV hall. The systematic error is estimated to be 20 μm from the size of corrections due to $x_{\rm offmag}$ and $\delta\phi$.

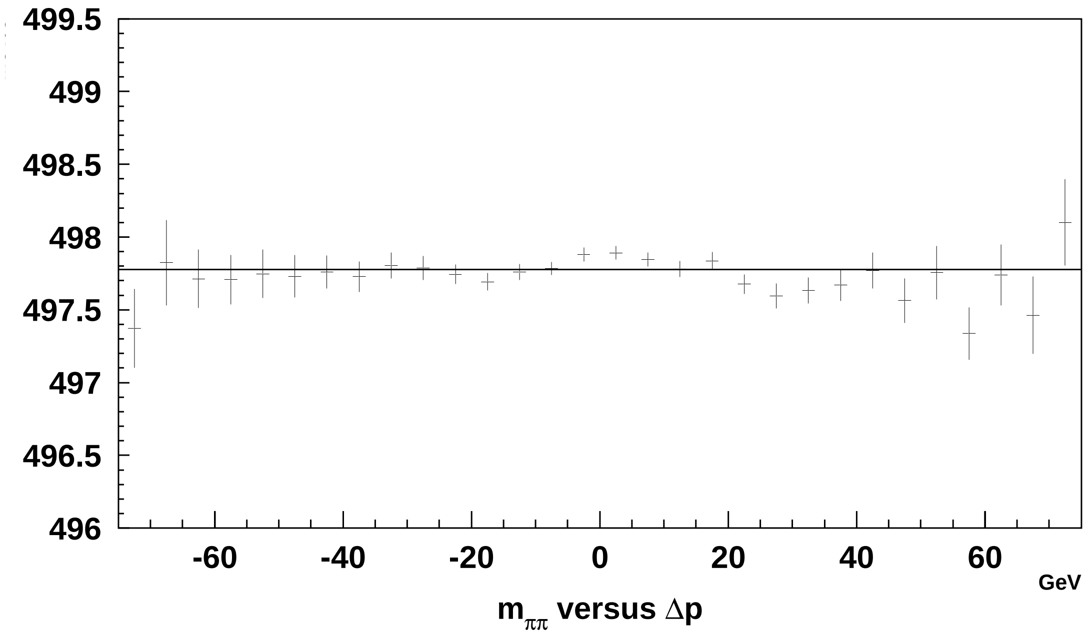

Figure 5.13: The $\pi^+\pi^-$ mass measured as a function of the difference of the left-bending and right-bending pion momenta, after the external alignment procedure.

5.7 Transverse Magnet Kick

The final element that needed calibration was the transverse momentum kick of the analyzing magnet. From detailed field maps, the relation between operating current and average transverse kick was determined. However, for more precision, it was decided to tune the value of the magnet kick from the data.

The magnet kick was set by fully reconstructing $\pi^+\pi^-$ decays, and tuning the kick such that the average $\pi^+\pi^-$ mass equals the kaon mass. A deviation from the kaon mass can be related to a correction in the kick as follows:

$$\begin{aligned} \frac{\Delta p_t}{p_t} &= \frac{\Delta m}{m}\left[1-\left(\frac{m_\pi}{m_K}\right)^2(2+R+1/R)\right]^{-1} \qquad (5.6) \\ &= 1.58\cdot\frac{\Delta m}{m} \end{aligned}$$

where $R = p_1/p_2$ and $\langle R+1/R\rangle \approx 2.7$ for KTeV.

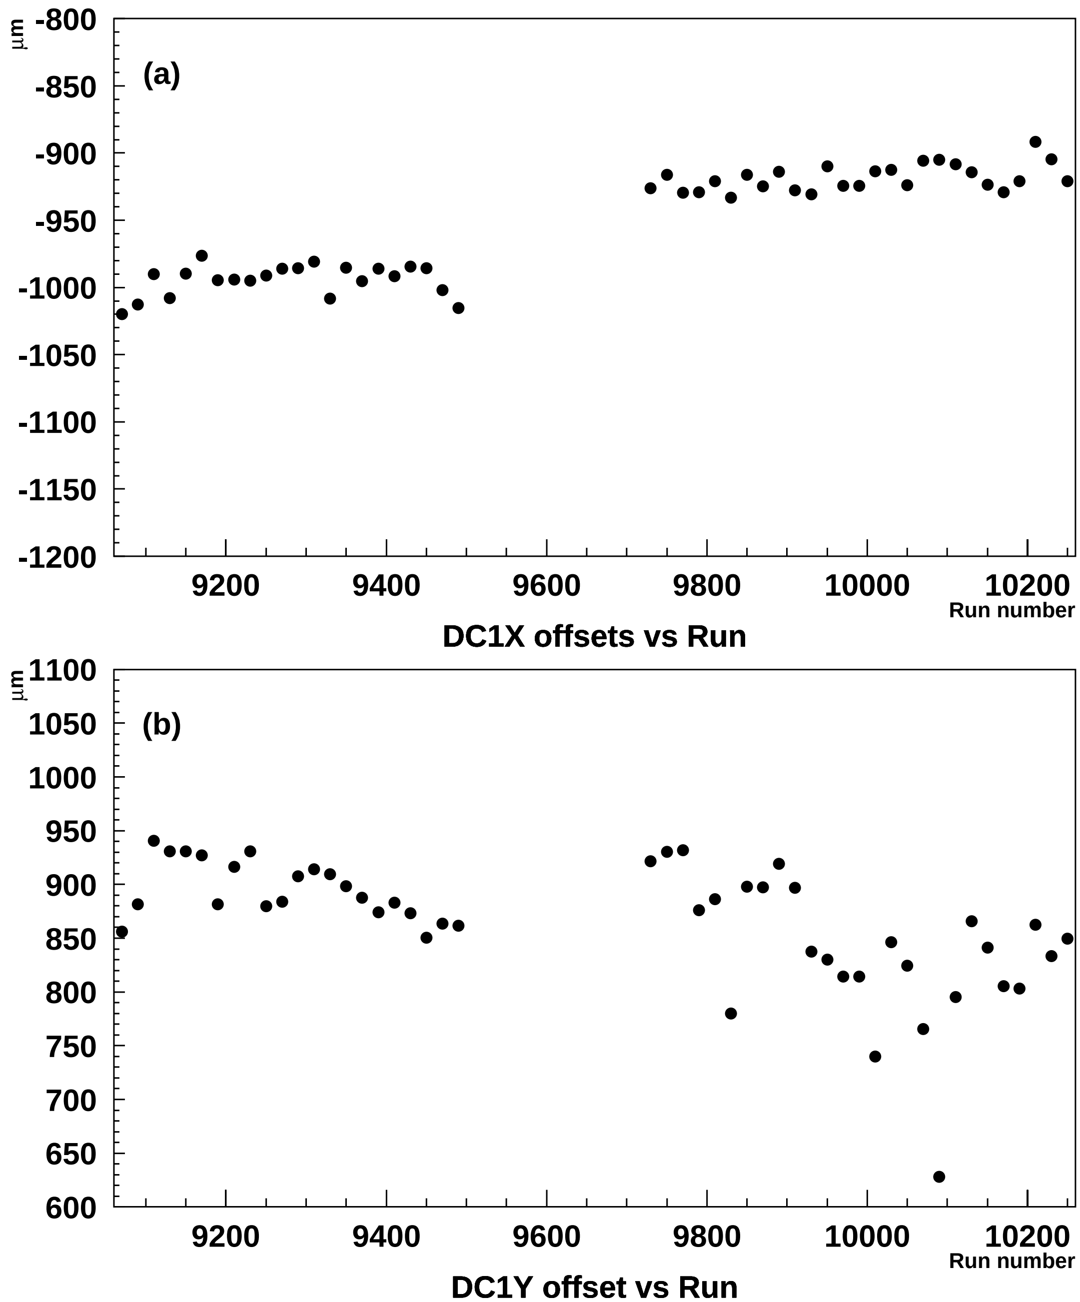

Figure 5.14: x (a) and y (b) offsets of Chamber 1 as a function of Run number.

The magnet kick was determined for each run period in which the magnet was run in the same polarity. Typically these periods were 10 runs long and lasted for two to three days.

5.8 High-SOD and Inefficiency Maps, and Other Resolution Issues

Although not technically part of the calibration or alignment, it is necessary to determine some quantities from the data for later use in the analysis or in the Monte Carlo simulation.

It is necessary to know the resolution of the DC system for two reasons. The resolutions are used in calculating χ^2s and errors for the vertex position and for the track segment match at the magnet, and are also used to smear the hit positions in the MC. The resolutions are determined in five regions within each chamber, (in each beam, above and below the beams, and outside a beam area) to be used in the Monte Carlo simulation. The average chamber resolution is used in the χ^2 calculations. These resolutions are determined from fitting the bulk of the SOD distribution during the $x(t)$ calibration procedure.

For the MC simulation, the rate of high-SODs and inefficiencies within the chambers as a function of position is needed. The high-SOD rate is the rate a high-SOD is used to make a track over the total number of tracks. The inefficiency rate is the number of times a hit is expected on a wire for a track, but no hit exists at all, divided by the total number of tracks. Tracks with true inefficiencies use isolated singles on the complementary wire within a chamber to construct the track. Two-dimensional "maps" of the position and rate of these two types of inefficiencies were made from the $\pi^+\pi^-$ data samples. Shown in Figure 5.15 is the distribution of the rate of high-SODs in Chamber 1. The rate of high SODs can be seen to be roughly constant at 1–2% across the chamber, and higher within each of the beam regions. In addition, the two beams show slightly different rates of high-SODs; the vacuum beam (on the right in this case) has a higher rate than the regenerator beam, showing that there is a component of the high-SOD rate that is dependent on the rate in the beam. The maps are determined for both regenerator positions separately. The use of these maps will be discussed more in Chapter 8.

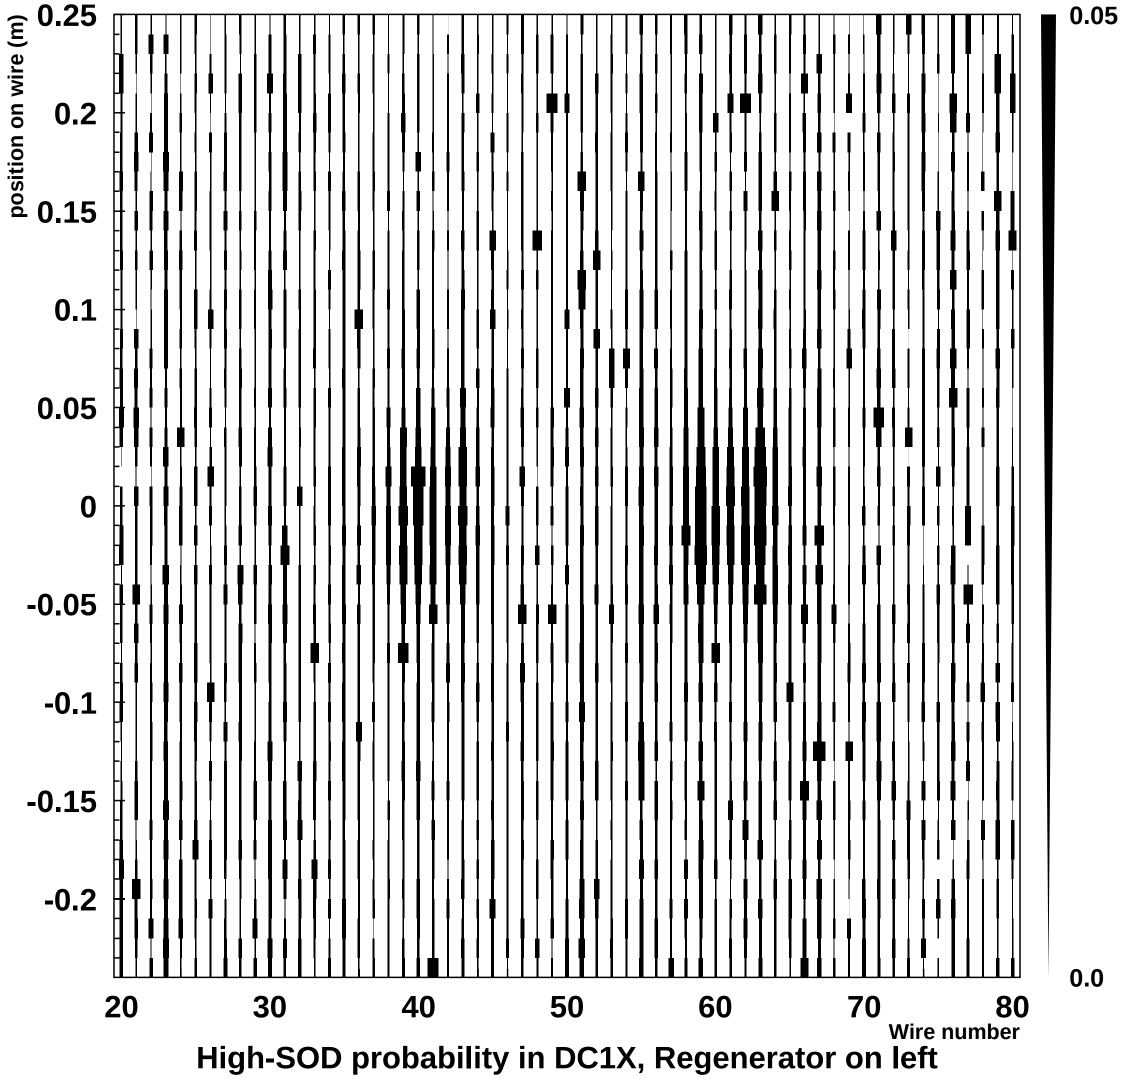

Figure 5.15: Map of high-SOD probability at the face of DC 1.

5.9 VV′ Positions

We need to determine the position of various additional devices with respect to the drift chamber system for use in the Monte Carlo. Of special concern is the location of the VV′ counters, and the location of the gaps between the counters and the location of the beamholes, since placing these counters in the Monte Carlo determine how well the Monte Carlo will model the trigger inefficiency.

During the DC calibration process, we determined the position of the VV′ counters, by determining the position of the gaps between the counters. We looked for two-track events where one track hit a counter in both V and V′, and the other track hit a counter in *either* V or V′, but not both. We have three hits to satisfy the "$2 \otimes 1$" trigger, and a missing hit, where the track passed through a gap. This technique allowed us to determine the position of the gaps between the counters. The long vertical gaps between counters are 500 μm on average, and the horizontal gap between the two sets of counters is 5 mm on average. The pointing resolution of the drift chamber system at the VV′ counters is $\sim$ 150 μm. Knowing the detailed position of the counters and gaps allows the Monte Carlo to reproduce the trigger inefficiency and minimize the systematic uncertainty, as will be discussed in Section 10.4.

To determine the position of the beamholes, we simply reconstruct tracks near the holes. This reconstruction can be see in Figure 5.16 for the V′ right beam-hole. In the figure, we can see the six counters that form the beam-hole, the small gaps between the vertical counters, and the large horizontal gap between the two sets of vertical counters. The beam hole shape is a complicated polygon; the solid black lines show our parametrization of the gaps.

5.10 MA and CA Aperture Positions

The final set of positions that must be reconstructed are the locations of the edges of the Mask Anti and the Collar Anti. These two inner detectors form the defining apertures for the neutral mode reconstruction, and hence we must ensure that the positions in the Monte Carlo simulation match the positions in the data, so that

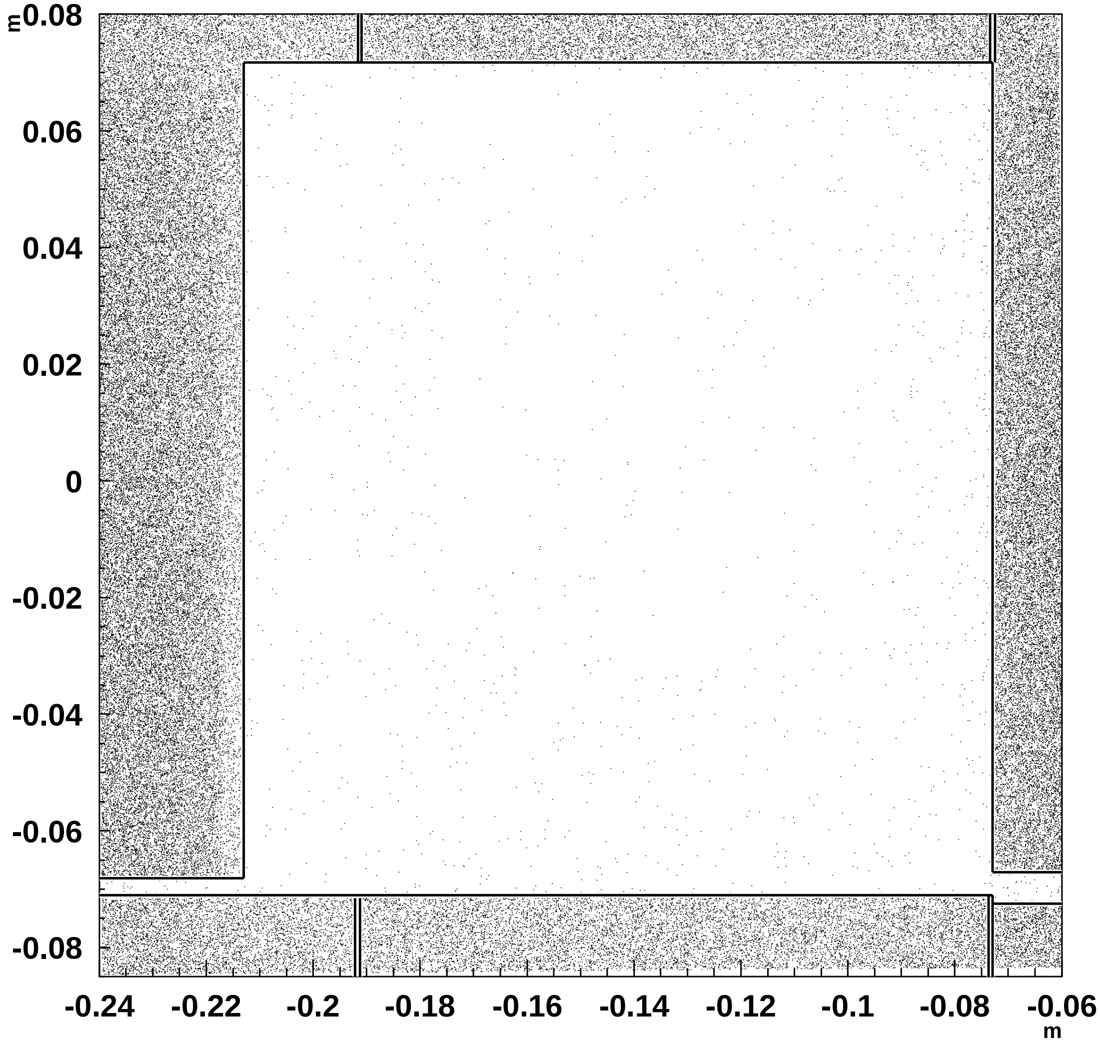

Figure 5.16: Tracks reconstructed around the V′ right beam-hole. The small vertical and larger horizontal gaps are visible. The solid black lines show our determination of the counter positions.

our acceptance is correctly predicted by the Monte Carlo. As will be discussed in Chapter 6, the position of the apertures is not as critical in the charged-mode analysis, since we cut on the *projected* position of the track at the given aperture.

The procedure for determining the effective edge of the apertures is quite simple. For the MA, we reconstruct $e^{\pm}$ from K_{e3} decays where the z-decay vertex is upstream of the MA. We make tight electron-identification, track quality, and track angle cuts to match the electron angle to a typical γ angle. We then reconstruct the distribution of e-tracks near the edge of the aperture in both the data and the MC. An example of this distribution is shown in Figure 5.17a, for the inner x edge of the East beamhole. Next, we "slide" the Monte Carlo distribution in 25 μm steps across the data distribution, and perform the K-S test in each step. The values of the K-S test for each step are shown in Figure 5.17b. The position of the aperture is determined to be the point of the maximal value of the K-S test. The shift is indicated in the figure as (-183 ± 21) μm. We add this value to the edge in the Monte Carlo to match the MC to the data.

This procedure is applied to the eight edges of the MA and the eight edges of the CA. The statistical error on the edges of the MA is ~ 25 μm, while on the CA edges the uncertainty is < 10 μm. These numbers compare to within a few hundred microns on the surveyed sizes of both the CA and the MA [77]. We have performed systematic studies of these edges by varying the cuts, the binning and the data subsets used to perform the edge determination. We arrive at a systematic uncertainty of 100 μm on the MA edges.

5.11 Conclusions

With the detector calibrated, it is now possible to reconstruct charged tracks and identify $\pi^+\pi^-$ decays with minimal bias between the two beams. In the next chapter, the reconstruction of charged tracks and the analysis will be discussed.

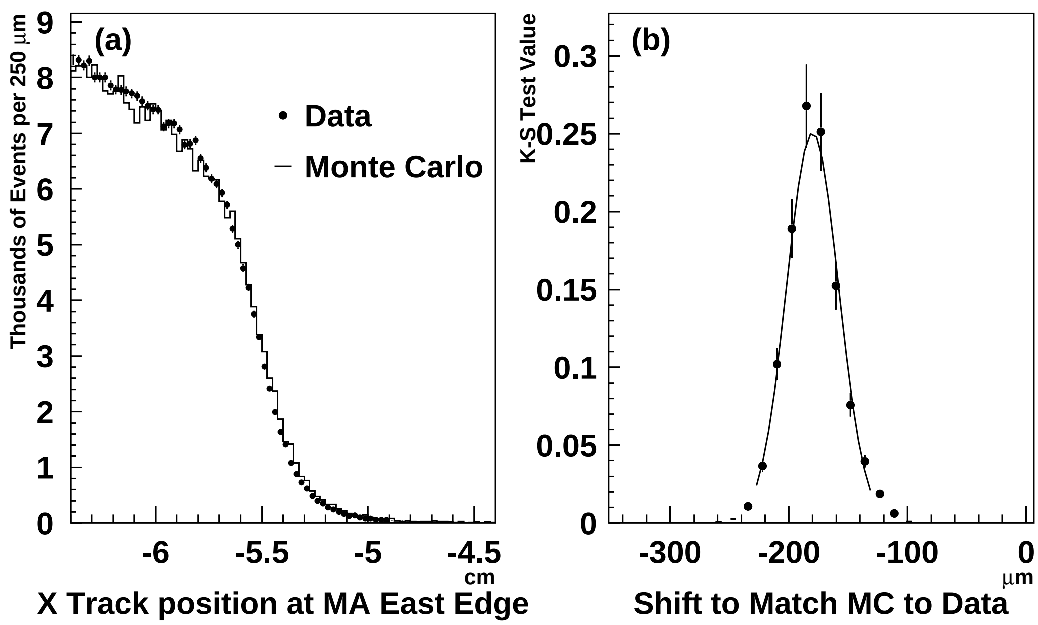

Figure 5.17: Tracks around the x inner edge of the East MA beam-hole.

CHAPTER 6
SELECTION OF THE $\pi^+\pi^-$ SAMPLES

This chapter will describe the selection of the $K \to \pi^+\pi^-$ data samples. We start with the trigger requirements, and then describe how pion tracks are reconstructed in the spectrometer. We will discuss the analysis "cuts" applied to the data samples to define the final samples of events more cleanly. Our discussion ends with a description of the remaining backgrounds in our samples, and the method for subtracting them.

6.1 Trigger Requirements

The main requirement of the trigger is to find kaon decays *consistent* with an unscattered $K \to \pi^+\pi^-$. Every part of the charged trigger attempts to approach 100% efficiency. The purity of the sample is of secondary concern, although an attempt is made to veto semi-leptonic and scattered kaon decays.

Table 6.1 details the trigger requirements for the $K \to \pi^+\pi^-$ sample. The charged trigger is composed of the VV$'$ trigger counters, veto elements, hit-counting in the drift chambers, and simple track identification. The average trigger rate for $\pi^+\pi^-$ decays was 17 KHz at the first level and 7 KHz at the second level of the trigger. Table 6.2 lists the rates of some of the individual sources that are used to make the charged-mode trigger under nominal running conditions. We will detail the elements for the trigger in the following sections.

6.1.1 Trigger Counters

The two trigger hodoscope banks of VV$'$ form the main element of the Level 1 trigger for charged decays. The main charged-mode trigger requires at least two

Table 6.1: Trigger definition for $K \to \pi^+\pi^-$.

Element	Requirement
Level 1 Trigger Elements	
Spill	Beam delivered from Tevatron
VV′ hit counting	2 hit counters in V, 1 hit counter in V′, or vice versa
VV′ East-West-Up-Down	One hit counter in either V-East or V′-East, Idem. for West counters Idem. for Up counters Idem. for Down counters
DC-“OR”	3 DC-“OR” hits in the 4 possible views
Veto counters	No appreciable energy in the regenerator, SAs, CIA or MU2
Level 2 Trigger Elements	
Hit Counting	Require 2 in-time DC hits in at least 3 of 4 of the y views and one hit in the 4th view.
YTF	Require one upward-going and one downward-going y track, with overlap in the center region of the chambers. One track can be of marginal quality.

Table 6.2: Rates of a sample of individual sources in the $K \to \pi^+\pi^-$ trigger for Run 9244.

Element	Rate (kHz)
One or more hits in V0	820
Two or more hits in V0	321
Hit in “east” V0	719
One hit in Chamber 1 y DC-OR	1.5 MHz
One hit in Chamber 2 y DC-OR	1.2 MHz
MU2 veto	384
SA2 veto	199
SA4 veto	189
CIA veto	150
Last module of regenerator	450

counters in the V0 bank and one hit counter in the V1 bank, or the reverse. This "$2 \otimes 1$" definition is looser than the "$2 \otimes 2$" trigger that one could define and still be consistent with two tracks; this is done to improve efficiency. The "$2 \otimes 1$" trigger definition, coupled with the overlapping design of the V0 and V1 counter banks, is mostly insensitive to gaps between the counters, except at discrete points where vertical gaps overlap with the horizontal gap between the upper and lower sets of counters. The location of the gaps is shown in Figure 3.6, and the image of the gaps can be seen in Figure 5.16.

An additional requirement is imposed on the trigger counters; we require that the two hits in the counters be consistent with a decay that has not scattered between the target and the decay vertex, *i.e.*, has no transverse momentum. We require one hit counter (from a track) to be in the east part of the detector, and the other to be in the west part of the detector. Also, one track must be in the upper part of the detector, while the other is in the lower part of the detector. The center set of counters in both the V0 and V1 counter banks count as both "east" and "west" for the purpose of the trigger, and hence the east and west definitions overlap in the center of the detector. The upper and lower trigger definitions overlap due to the layout of the V0 and V1 counter banks: see Figure 3.6. All unscattered $\pi^+\pi^-$ events will pass this trigger requirement; however, some semi-leptonic K_{e3} and $K_{\mu 3}$ decays will fail this requirement, as will some scattered $K \to \pi^+\pi^-$.

6.1.2 DC "OR" Requirements

Since the VV′ counter banks are located downstream of the four drift chambers, it is possible for a kaon to decay to two tracks downstream of DC 1 and still satisfy the trigger hodoscope requirement. Events decaying downstream of DC 1 can only be reconstructed with great difficulty and poor efficiency, because they leave no signals of hits in DC 1; these events are not part of our fiducial sample. To veto events that decay downstream of DC 1, we require that there be signals in DC 1 and DC 2 within 90 ns of the hits in the trigger hodoscope. The signals from groups of 16 DC wires are "OR"ed together by fast electronics that provide the trigger signal. We have instrumented both the x and y views of the two chambers, and we require that 3 of the 4 views have signals.

6.1.3 Veto Sources

The charged trigger is vetoed at Level 1 if there is any significant energy in the regenerator, the photon vetos, or the muon system. The regenerator vetos kaons that scatter inelastically. The veto is the OR of the signals from three modules along the length of the regenerator and the last module that views the lead pieces. The photon vetos reduce the rate from $K \to \pi^+\pi^-\pi^0$ and $K \to \pi^+\pi^-\gamma$ events. The muon system reduces the rate from $K_{\mu 3}$ decays.

6.1.4 Bananas, Kumquats and the YTF

We made a loose requirement on the number of hits in the Drift Chambers at Level 2, requiring that the Banana Trigger boards find at least two hits in three of the four y chamber views, and one hit in the remaining chamber. To lose an event at this stage would require missing both hits in one chamber. We use the y views only so that the same physical wires will be used for both the vacuum and the regenerator beams.

In addition to hits in the y views, we required that the YTF found track candidates in the y view. We required one up-pointing track and one down-pointing track, and tracks in the center region of the chambers counted as both up- and down-going tracks. In addition, one of the two tracks was allowed to be "marginal" and be missing a hit in chamber 1 or 2 (but not in chamber 3 or 4).

6.1.5 Level 3 Requirements

If any possible vertex candidate has an invariant mass $m_{\pi^+\pi^-} > 450$ MeV/c^2, the calorimeter information is unpacked and clusters are found and matched to tracks. For each track within a vertex candidate that matches to a cluster in the calorimeter, the quantity E/p is calculated, and is required to be less than 0.9 for the event to be considered as a $\pi^+\pi^-$ candidate. Note that an *inclusive* approach is used; if *any* vertex candidate has $m_{\pi^+\pi^-} > 450$ MeV/c^2 and two tracks with $E/p < 0.9$, it is tagged as a $\pi^+\pi^-$ candidate. The vertex candidate that matches the Level 3 cuts does not have to be the best vertex candidate; the best vertex candidate will be determined in the offline analysis, when the final calibration of the detector is available.

6.2 Track Finding

As with the trigger definitions, the goal of the track finding algorithms is to reconstruct two-track events consistent with being from a kaon decay. Efficiency is considered more important than purity of the sample; because of the resolution of the detector, we have an excellent intrinsic signal-to-background ratio.

In the following sections, we will discuss how the tracking algorithm reconstructs tracks from signals from the Drift Chamber wires.

6.2.1 Hits, Pairs and SODs

A "hit" in the DC is defined as an analog signal above threshold that fires the TDC for a given wire. To be used for tracking, the TDC signal must occur within an "in–time" window 235 ns from the trigger signal. The $t(0)$ determined in the calibration procedure (Chapter 5) is fixed by the calibration procedure to be 15 ns from the start of the in-time window. The hit represents the signal on the wire from a track passing next to the wire. Moreover, the TDC value of the hit can be related to the distance of the track from the wire through the $x(t)$ calibration as discussed in Chapter 5. The TDCs can store more than one hit per wire, a capability we will use to investigate the chamber efficiency when assigning systematic uncertainties, but only the first hit is used to form pair candidates.

The one-half wire separation between the unprimed and primed wire planes form cells through which the tracks pass. Therefore two hits can be paired together in a cell, to determine on which side of each wire the track passed. Equally importantly, the "pairs" allow for classification of the hits. If the hits come from an in-time event, the sum of the distances of the two hits should be the same as the separation of the two wires that define the cell, 6.35 mm, after accounting for effects such as the track angle and resolution. The resolution can move the SOD by ~ 140 μm. The track angle can have a larger effect; in the downstream chambers, we allow for a deviation from the nominal SOD value of up to 1.5 mm.

The tracking algorithm looks for pairs of hits and classifies them by their SOD value. In a loose sense, SODs can be "low," "good," or "high." Two additional classifications of "pairs" exists, concerning single hits on wires. Isolated hits that have no accompanying hit to make a true pair are still classified as two sets of pairs – one pair for each side of the wire. Singles that have an accompanying hit that is outside the in-time window could be an in-time pair from an earlier event, and strict tracking requirements could reject these events. Our analysis allows for these events to be used to form tracks in the same manner as a true isolated single.

SODs are classified as good if they fall within ± 1 mm of the cell spacing (± 1.5 mm in the downstream x chambers, due to larger track angles caused by the analysis magnet); these SOD pairs are treated as a single unit. Typical hit resolution is 100 μm, so the resolution on the SOD is 140 μm. The position of the track is taken as the average position of the two hits within the pair. Low SODs, less than 5.35 mm, can occur due to physical processes, such as two tracks passing through the same cell, or a delta ray (an energetic electron knocked off an atom). Both hits are listed as "pairs" to be used by the tracking algorithm, but each hit is constrained to be on the side of the wire as determined by the SOD. Because low SODs are caused by physical processes, either hit is equally likely to be the hit associated with the track. High SODs pairs, like low SODs, are constrained to the side of the wire of the SOD pair, and are also split into two pairs. The specific inefficiency that causes the high SOD problem occurs near the sense wire, so that the hit associated with the track is more likely to be the hit farthest from the sense wire.

In the tracking code, pairs are assigned a "quality value" to be used later when assembling tracks. Isolated single hits have a value of 1, low-SODs and high-SODs have a value of 2, and good SOD pairs have a value of 4.

The tracking algorithm finds pairs by "walking" across the unprimed-primed set of wire planes, looking for hit pairs that can be classified as SODs. This procedure is shown in Figure 6.1. The event shown in the figure (with five tracks!) would list the following pairs (from left to right)

- Good SOD pair in cell 2.
- Low SOD pairs in cell 3 (from a combination of other hits) and in cell 4, from a real low-SOD pair.
- High SOD in cell 9.
- Isolated singles in cell 12.

The arrays used in the code to store hit information have limits on the number of pairs that can be stored. To ensure we do not exceed the limit of 256 pairs, and to reduce the combinatorics, pairs are dropped from the list in the following order if the limit is exceeded:

- Singles with an out-of-time partner, possibly a good pair from an earlier bucket.
- Very high-SOD pairs (SOD $>$ 10 mm).
- Isolated singles.
- High SODs.
- Low SODs.

There is concern that this procedure may cause a regenerator-to-vacuum beam bias, because events from one beam may have more underlying event activity and hence more hits than the other. However, systematic studies were undertaken to increase the internal size of the tracking arrays to avoid removing hits, and no bias was seen. Also, we will study the systematic effects of extra hits in Section 10.9.4, and see little effect.

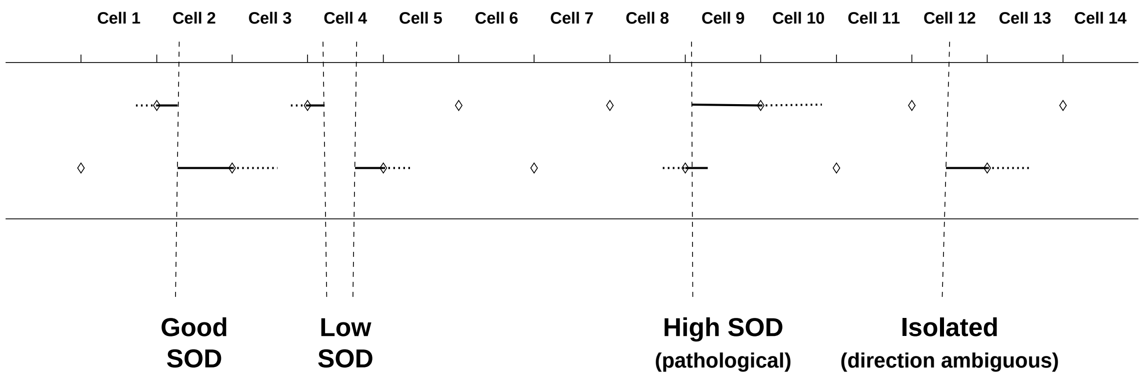

Figure 6.1: Diagram of the types of SODs and hits. The diamonds show the sense wire positions, the dashed vertical lines show the tracks passing through the DC, the solid horizontal lines show the true drift distances, and the dotted horizontal lines show the hit being placed on the other side of the wire. The list of pairs that would be found in this example is described in the text.

6.2.2 The x and y Track Candidates

Once hit pairs are found, the tracking code looks for track candidates. At this stage, no constraints are placed on track candidates; they are not required to come from the beam or point to the CsI calorimeter, and multiple track candidates can use the same hit or hit pair.

Finding y Track Candidates

The tracking code looks for candidates in the y view first, because it is the simpler of the two views. The magnet does not bend the tracks in the y-view, making both the algorithm simpler and the combinatorics easier.

The algorithm searches for y tracks by selecting one pair in DC 1 and one pair in DC 4, and connecting the line between the two. If pairs are found in DC 2 and DC 3 within 5 mm of the line, the “quality values” of the pairs from the four chambers are summed. A perfect track (four good SODs) has a quality sum of 16; the track is kept if the sum is greater than 11, which corresponds to two good SODs, one high or low SOD, and an isolated single hit. This value was chosen as the minimum sum for which there is only one isolated single. It is difficult to determine the track path with more than one isolated single in the track because of sense wire ambiguity. After passing this cut, the track is fit to a line, and kept if the fit χ^2 is acceptable. The fit χ^2 is not normalized to the expected chamber resolution, and the requirement is that the sum of the residuals squared be less than 4×10^{-6} m^2. If we were to normalize to an expected hit resolution of 100 μm, the cut would be $\chi^2 < 400$ for 2 degrees of freedom.

At this stage, the tracking code determines the number of y track candidates that can *coexist* with each other. Track candidates in general are not allowed to share hits, with the exception that tracks can share hits in one chamber if each track uses the hit as part of different good-SOD pairs. This hit ambiguity will be resolved later in the code with final tracks in hand. The code considers which sets of tracks do not share hits, and fills a two-dimensional "co-existence" matrix of allowable track candidate combinations. From this matrix the code determines the number of "exclusive y tracks," which is equal to 0 if there are no y tracks, 1 if there are one or more y tracks but no two coexist, and 2 if there are two or more y tracks that coexist. We require two exclusive y tracks to proceed.

Finding x Track Segments and x Track Candidates

In a manner similar to finding tracks in the y view, we look for track "segments" in the x view. The added complication is the bend in the tracks caused by the dipole analysis magnet; we cannot look for straight tracks between DCs 1 and 4. Track segments are found in DCs 1 and 2, and DCs 3 and 4 by looping over all hit-pair combinations in the chambers, and x track candidates are formed by matching the segments at the midplane of the analysis magnet. The upstream segments in DCs 1 and 2 are only allowed to have a slope with respect to the z axis of 100 μrads, and the sum of "quality value" must be 4 or greater. Similar constraints are placed on the downstream segments; the angle constraint is 150 μrads and the minimum quality sum is 5.

For each combination of upstream and downstream segments, the distance between the two segments' projections to the midplane of the magnet ($z = 170$ m) is calculated. If this separation is less than 6 mm, and the sum of the quality values is at least 11, the two segments are considered an x track candidate. Events without two exclusive x track candidates (defined similar to exclusive y track segments, above) are rejected at this stage.

6.2.3 Vertex Candidates

Once x and y track candidates have been identified, the next stage of the tracking algorithm is to identify tracks coming from a common decay vertex. The z position of the intersection of each pair of y track candidates is calculated, and compared to the same for each pair of x track candidates. A vertex candidate is a set of y track candidates matched to a set of x track candidates consistent with originating at the same z position. At this stage, the requirements for consistency are quite loose, and hence there are many vertex candidates. The two x track candidates in a vertex candidate are required to have opposite bend directions in the analysis magnet, consistent with particles of opposite charge.

It is worth mentioning at this point that since the KTeV DCs only have x and y views, it is impossible to associate an x track with a y candidate without additional information. Even for the case when there are two and only two x track candidates and two and only two y track candidates, there are still two vertex candidates, because each x track candidate can be paired with either of the y track candidates. This ambiguity is resolved by matching tracks to clusters of energy deposited in the CsI calorimeter.

6.3 CsI Energy Reconstruction and Clustering

Finding clusters of energy in the calorimeter consists of measuring the scintillation light in each crystal, translating it to an energy, and summing surrounding crystals into discrete "clusters" of energy. This process is extremely complicated, and the one where the most effort was spent, since the $K \to \pi^0\pi^0$ analysis depends on it completely. We will postpone our discussion of how the energies and positions of clusters are determined until the discussion of the neutral-mode analysis in Section 7.2, and assume we have found clusters in the calorimeter.

6.4 Corrected Tracks and Vertex Finding

After reconstructing clusters within the CsI, the tracking code can match tracks to clusters, apply corrections to the tracks, and find the best vertex from the list of vertex candidates. The procedure loops over vertex candidates found above, determines if the given pairing of x and y track candidates matches to a cluster, and if so, applies final corrections to the tracks.

6.4.1 Track-Cluster Matching

Each pair of x-y track-candidates is considered as a full track candidate. The position of the track candidate is extrapolated to the CsI face, and compared to the x-y position of the clusters. The matching requirement requires the track position to be within 7 cm of the cluster position for a full match. In this analysis, each track is required to match to a cluster in the CsI calorimeter, although the tracking code can be used in a looser fashion where tracks can either match to the beam hole or the edge of the CsI, or not match to clusters at all.

6.4.2 Track Corrections

Corrections to the tracks can be made once x and y tracks are matched to each other. Most of these corrections involve knowing the other coordinate of the track at a given z position; for example, in correcting y tracks, we can apply correction for rotations by knowing the x position of the track. At this stage, we correct for the physical rotation of the drift chambers, individual wire positions, time-delay of pulse propagating along the wire, corrections to the SOD value based on the angle of the track through the drift cell, and making a small correction due to a ~ 60 Gauss (at DC 2) residual magnetic field. The need for this small magnetic field correction was first indicated by a difference between the reconstructed kaon mass in the "inbend" and "outbend" data samples (inbend events have the tracks bent back into the detector by the analysis magnet, as shown in Figure 6.2). The determination of this correction is quite interesting, so we detail it further here.

The large B_y component of the magnetic field moves tracks in the $\pm x$ direction. There is a small part of the field integral between chambers 1 and 2, and between 3 and 4, such that the hits in chambers 2 and 3 are moved from the ideal straight-line trajectory assumed in the track-fitting. Furthermore, this movement is momentum-dependent. We measure this effect by considering the movement of hits in chamber 2 relative to chamber 1 as a function of the azimuthal angle ϕ, where ϕ is the angle of the decay plane with respect to the horizontal plane. We define ϕ such that inbend events that are in the horizontal plane have $\phi = 0°$ and horizontal outbend events have $\phi = 180°$. At chamber 2, the tracks will have been moved in the x direction by the fringe field. This causes a rotation of the plane as a function of the initial angle ϕ; tracks at $\phi = 0°$ have no rotation between chambers 1 and 2 since they are flat in the horizontal plane, while events with $\phi = 90°$ are maximally rotated. This effect is shown in Figure 6.3.

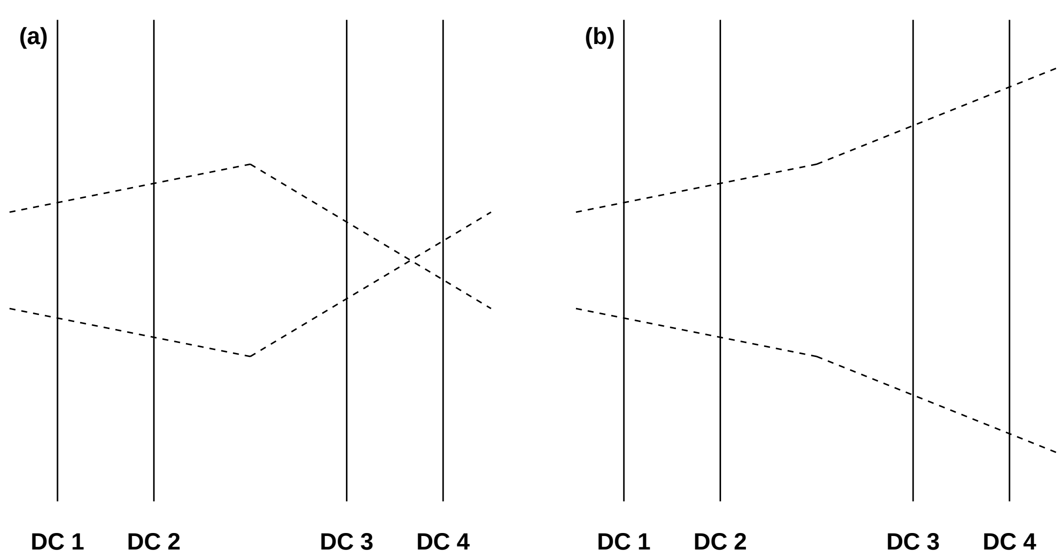

Figure 6.2: Examples of (a) an inbend event and (b) and outbend event.

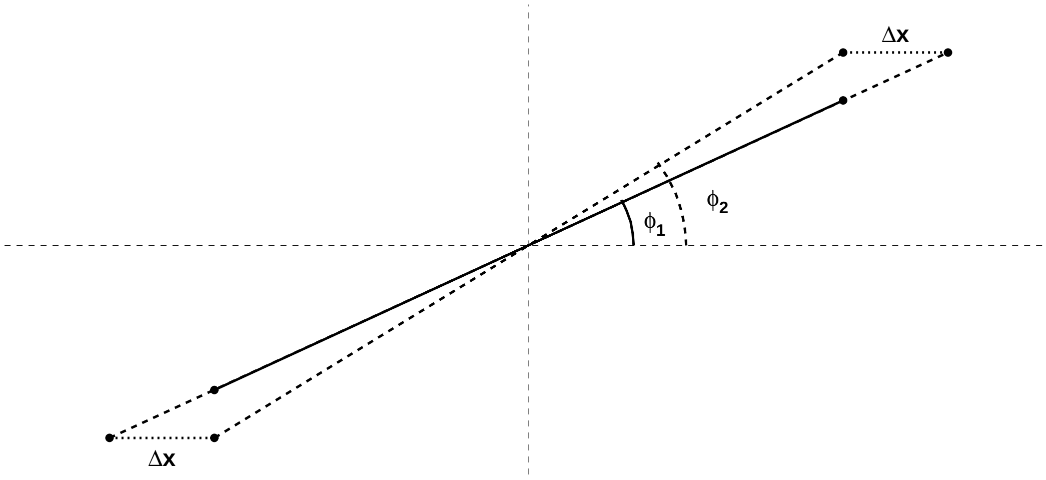

Figure 6.3: Effect of the fringe field on the azimuthal angle ϕ of the decay plane. The fringe field moves the tracks by Δx at DC 2, causing a rotation of the azimuthal plane. The solid line shows the orientation of the decay plane at DC 1, while the dashed line shows the rotated orientation of the plane measured at DC 2.

The rotation of the azimuthal angle can be calculated to be

$$\begin{aligned} \tan(\phi_2 - \phi_1) &= \sin\phi_1 \frac{(\Delta x_2 - \Delta x_1)}{|\vec{r}_2|} \\ \sin\phi_1 &= |\vec{r}_2| \tan(\phi_2 - \phi_1) \frac{C}{1/p_1 + 1/p_2}, \end{aligned} \tag{6.1}$$

where $\phi_{1,2}$ are the plane azimuthal angles measured at DCs 1 and 2, $\Delta x_{1,2}$ are the movements in the x direction of tracks 1 and 2 due to the fringe field, and $|\vec{r}_2|$ is the separation of tracks 1 and 2 at DC 2. In the second equation, we convert the movements Δx to a constant C over the momenta $p_{1,2}$ of the tracks. We plot this dependence in Figure 6.4; the amplitude of the sine function measures the integral of the fringe field between DCs 1 and 2.

With the amount of fringe field determined, we apply corrections to the hits at DCs 2 and 3 to account for the deviation from straight tracks. This correction is shown in Figure 6.5. The solid line shows the true bend of the tracks through the magnet, including a small deviation from straight tracks upstream of DC 2,

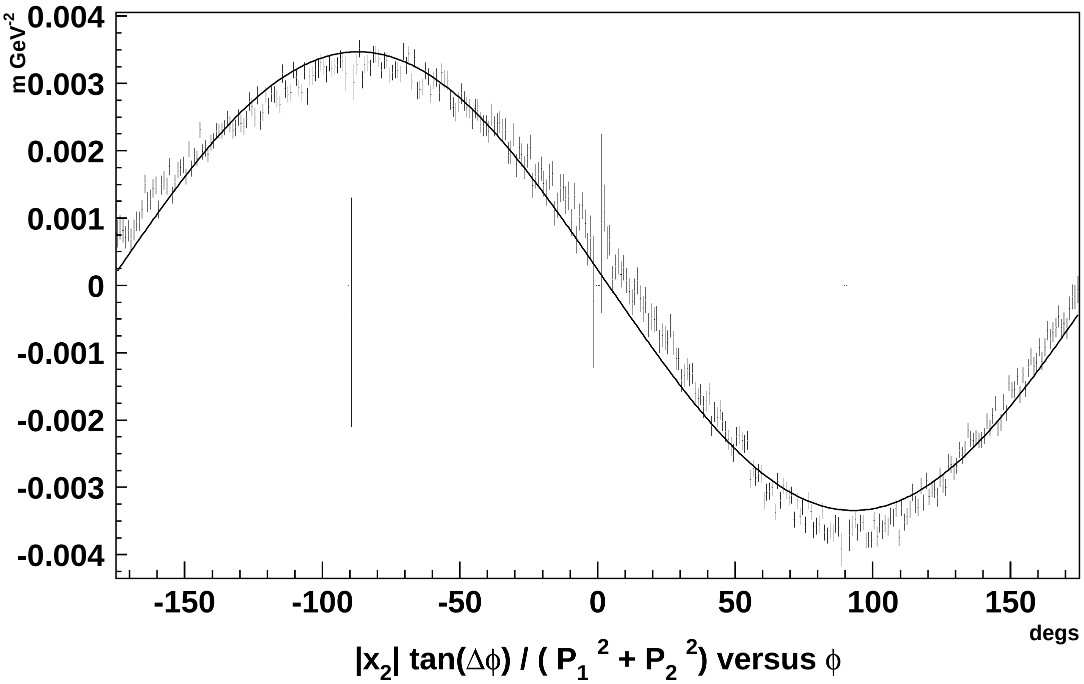

Figure 6.4: The fringe field determination between DCs 1 and 2. We plot the RHS of Equation 6.1 versus the azimuthal angle ϕ measured at DC 1.

and downstream of DC 3. We move the hits in DCs 2 and 3 based on the above correction to line up the hits as if they were on the straight tracks shown as dotted lines. The arrows in the figure show the direction we move the hit positions.

Further corrections are applied to the SODs at this stage to attempt to resolve hits that are used on more than one track, and good-SOD pairs that are more than 600 μm from the cell size. Any hit that is shared by both tracks as part of two good SOD pairs is assigned to the pair with the better SOD value. The other track then uses the other hit from the worse SOD as a single hit. Good SODs that are 600 μm from the nominal SOD value are split into low- and high-SODs pairs, and the hit that matches better to the track (based on extrapolations from the other chambers) is used in the track. At this point we require each track to have at least one good-SOD pair (within the tight 600 μm definition) in both the x and y views.

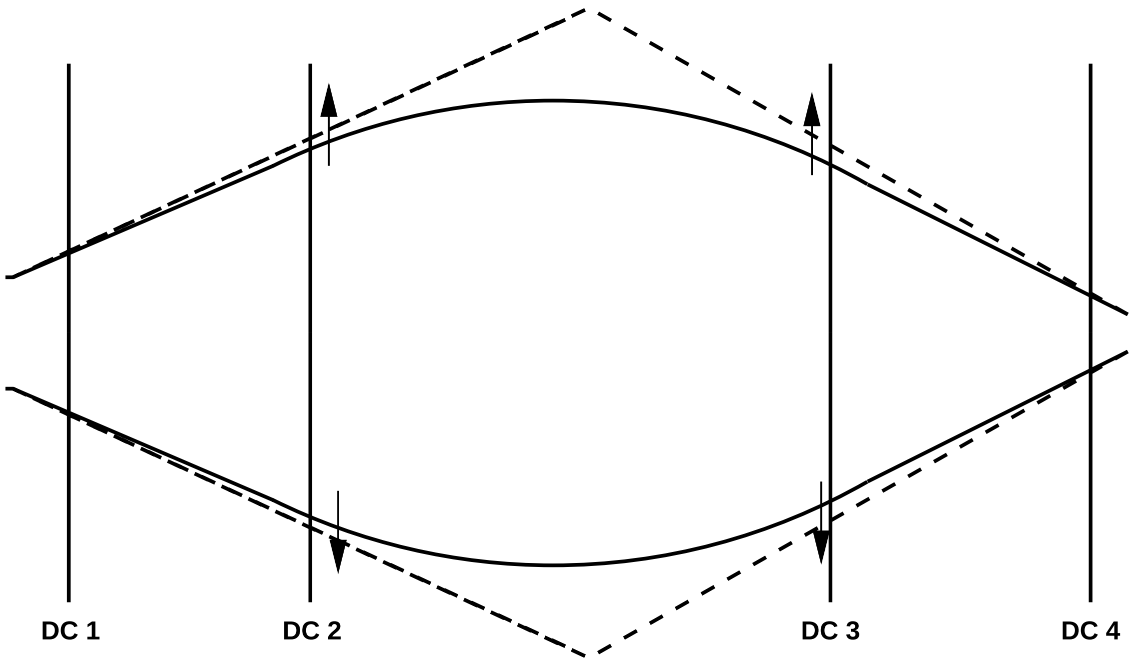

Figure 6.5: The true bend of tracks through the magnet (solid curves) and the assumed straight tracks (dotted lines). We correct the hit positions in DCs 2 and 3 (shown as the arrows) to line up the hit positions to straight lines.

Finally, we refit separate y track segments upstream and downstream of the magnet. Due to the three-dimensional nature of the magnetic field, y tracks will be slightly bent at the magnet.

With all the above corrections, we calculate the "offmag χ^2," which describes how well the upstream track segment matches to the downstream track segment at the bend plane of the analysis magnet. This quantity helps determine the quality of the tracks, and we use it below in our data analysis.

6.4.3 Vertex Finding

The tracking routines can now fit for the best vertex from the list of candidates determined above. Each candidate set of tracks has been matched to clusters and has had final corrections applied to the tracks. Once corrections are applied, the upstream (DCs 1 and 2) track positions in x and y are fit to the constraint that they come from a common (x, y, z) space-point. There are eight data points in the fit (four points in DCs 1 and 2 in each of the x and y views) and seven unknowns (the vertex position and two track slopes in both views), so the fit has one degree of freedom. The χ^2 equation minimized is:

$$\chi^2 = \sum_{i=1}^{2} \sum_{j=1}^{2} \frac{\left[(x_{ij} - x_v) - s_i^x \left(z_{ij}^x - z_v\right)\right]^2}{\sigma_{ij}^{x\ 2}} + \frac{\left[(y_{ij} - y_v) - s_i^y \left(z_{ij}^y - z_v\right)\right]^2}{\sigma_{ij}^{y\ 2}}, \qquad (6.2)$$

where the index i runs over the tracks and the index j runs over the chambers, $(x, y, z)_v$ is the position of the vertex, $z_{ij}^{x,y}$ is the position of the hit at the given x or y chamber plane (which can vary if we use a good-SOD or a single hit), and $s_i^{x,y}$ is the slope of the track. The above equation can *almost* be minimized analytically. With a few mathematical substitutions, the expression simplifies and will converge within a few iterations. The uncertainties for each space-point have a gaussian resolution term and the space points at DC 2 also have a simple model for multiple scattering in DC 1. No effort is made to distinguish single hits, low- or high-SOD or good-SOD pairs, except through the varying resolutions on each space-point.

The best vertex is determined from the χ^2 of the fit, the match between the upstream and downstream segments of each of the tracks at the magnet, and the number of good-SOD pairs used in the vertex candidate. The number of good-SOD pairs is used to disfavor vertex candidates with more isolated hits or bad-SOD pairs, which have increased position errors that may lead to a lowered χ^2.

Figure 6.6 shows a reconstructed two-track event. One can see the nearly straight y tracks passing through the analysis magnet, the x tracks being bent in the analysis magnet, and the tracks matching to clusters of energy in the CsI. The transverse decay vertex position is in the regenerator beam, so this event is a likely $K_S \to \pi^+\pi^-$ event. The figure also shows the invariant mass $m_{\pi^+\pi^-}$ and transverse momentum p_T^2 consistent with the $K_S \to \pi^+\pi^-$ decay. The kinematic properties will be described below.

With the best vertex in hand, we are ready to move on to identifying $K \to \pi^+\pi^-$ events.

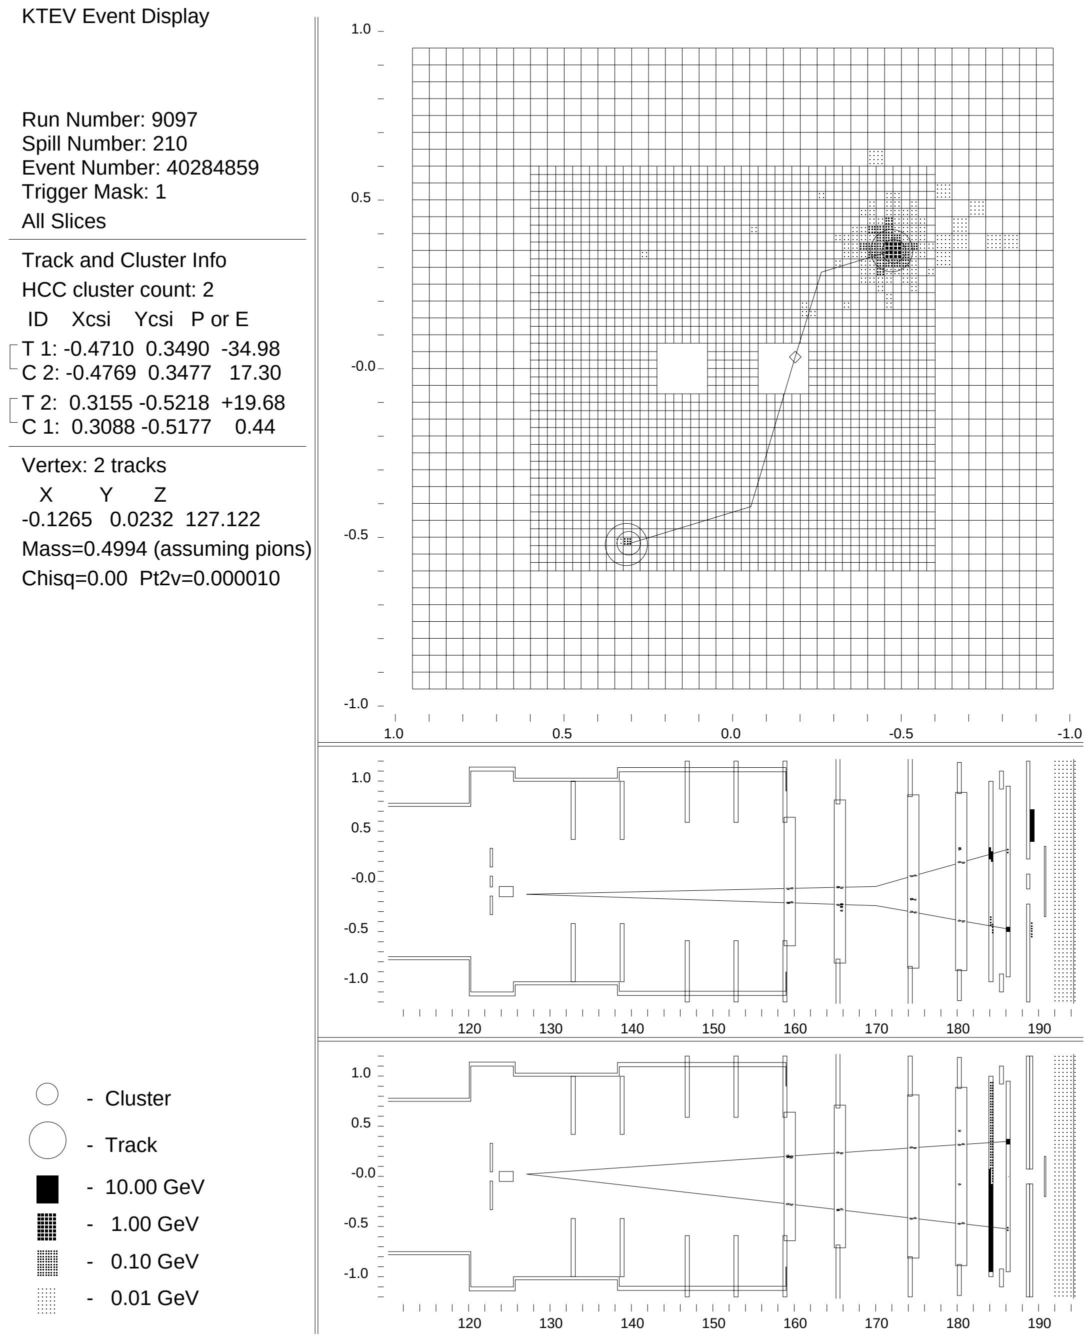

Figure 6.6: Event display for $K \to \pi^+\pi^-$ event. The top panel shows the projection of the tracks matching to clusters of energy in the CsI calorimeter. The middle (bottom) panel shows the x (y) view of the tracks. In the x view, the tracks are bent by the analysis magnet, and are not in the y view. In the middle panel, one can see the regenerator in the right beam.

6.5 Event Reconstruction

From the tracking routines, we have a final vertex, the two track momenta, and the energy of the two clusters corresponding to the tracks. In turn, we can reconstruct three important kinematic quantities of the two-track system: the invariant mass, the lab energy, and the momentum transverse to the beam direction.

We reconstruct the invariant mass assuming both tracks are pions.[1] If the tracks are assumed to be pions, the invariant mass of the two-track system is

$$
\begin{aligned}
m^2_{\pi^+\pi^-} &= 2m^2_\pi + 2\sqrt{m_\pi + p_1^2}\sqrt{m_\pi + p_2^2} - 2\vec{p}_1 \cdot \vec{p}_2 &\quad (6.3)\\
&\sim 2m^2_\pi + p_1 p_2 \theta^2 + m^2_\pi \left(R + 1/R\right), &\quad (6.4)
\end{aligned}
$$

where $R = p_1/p_2$ and θ is the opening angle of the two tracks.

The energy of the system is simply

$$
E_{\pi^+\pi^-} = \sqrt{m^2_\pi + p_1^2} + \sqrt{m^2_\pi + p_2^2}. \quad (6.5)
$$

The final kinematic quantity of interest is the transverse momentum to the kaon direction of flight. For a kaon produced in the target that does not scatter, the sum of the momentum of the daughter particles should point along the line between the target and the decay vertex. If there is appreciable momentum transverse to the kaon direction of flight, either the kaon scattered between the target and the decay vertex, or there is an unmeasured particle in the final state, such as a neutrino from a K_{e3} decay. Even in the regenerator beam, we are only concerned with kaons that undergo *coherent* regeneration, with no momentum transfer. Hence all our events of interest should have no net transverse momentum (p_T).

To better identify kaons that scatter in the regenerator beam, we calculate the p_T^2 assuming that scattering occurred in the regenerator; that is, with respect to

[1] Although the semi-leptonic decays K_{e3} and $K_{\mu3}$ occur at far higher rates, these decays will mis-reconstruct due to the missing neutrino. Furthermore, these backgrounds are suppressed by particle identification in the CsI (K_{e3}) and by the muon trigger hodoscopes ($K_{\mu3}$). The remnant backgrounds due to these decays will be discussed in Section 6.7.

Figure 6.7, we calculate $|\vec{p}|^2 \sin^2\phi_R$ instead of $|\vec{p}|^2 \sin^2\phi_V$, even for events in the vacuum beam. For coherent events, this does not matter, while it allows us to better characterize scattering in the regenerator, which in turn feeds into our understanding of the scattering background in the neutral mode analysis. For signal events, we require the square of the transverse momentum to be less than 250 MeV$^2/c^2$.

Scattering in the regenerator can lead to confusion as to the beam of origin of the kaon. If the kaon scatters in the regenerator with high enough p_T, it can "cross over" into the vacuum beam. To correctly assign these events to the regenerator beam, we project the vertex (x, y) position along the total momentum vector back to the z of the regenerator, and assign the particle to a beam based on that transverse position. With respect to Figure 6.7, we assign the event to a beam based on the position x_R, instead of x_V. Although all of these "cross-over" events are removed by the p_T^2 cut, we need to correctly assign them to the regenerator beam so that we can correctly characterize scattering in the regenerator to be used in predicting

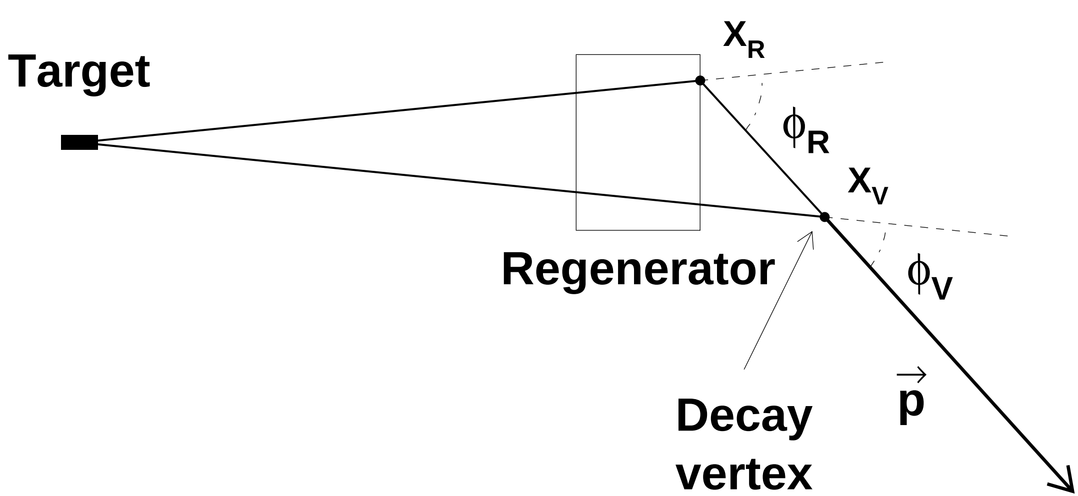

Figure 6.7: Cartoon of the scattering in the regenerator. The kaon scatters in the regenerator at the point x_R, and decays at the point x_V. We use the angle ϕ_R to calculate p_T^2 for all events, even those decaying in the vacuum beam.

the scattering background in the neutral mode, where "cross-over" into the vacuum beam is a large background.

6.6 Analysis Cuts

Analysis cuts are applied to the data to reduce background levels or to allow a cleaner definition of the acceptance. Due to the high performance of the detector, we can make tight analysis cuts on most quantities of interest. For example, due to the performance of the CsI calorimeter, we can make tight particle identification cuts to reduce the K_{e3} background.

In the following sections, we will discuss the cuts used to identify and define the $K \to \pi^+\pi^-$ data samples. The cuts can be broken up into four loose categories: data quality, composed of run selection, trigger verification and track quality cuts; background suppression, composed of veto counter cuts, extra-particle cuts and pion-identification cuts; kinematic cuts, such as energy and invariant mass cuts; and aperture cuts, such as fiducial cuts at the edge of the detector and track-separation cuts. We detail the analysis cuts individually.

6.6.1 Run Selection

After data taking was concluded, certain runs collected were deemed unusable for this analysis. The main class of problems was various malfunctionings of the CsI calorimeter electronics. However, we also eliminated periods when the beam was out of alignment and when the drift chamber voltages sagged.

6.6.2 Trigger Verification

It is possible that some events that would have failed the trigger in the absence of accidental activity (for example, a track passing through a crack in the trigger hodoscope) pass the trigger because the accidental activity aids in satisfying the trigger (*e. g.*, a cosmic ray hits a counter in the hodoscope, satisfying the trigger when combined with the other track). These types of events are called "volunteers."

For this reason, we verify the trigger offline in software, requiring that the two tracks alone satisfy the Level 1 trigger. We do not verify the Level 2 trigger, and we will assign a systematic uncertainty based on this choice in Chapter 10.

6.6.3 Track Quality Cuts

From finding the best vertex, we have calculated two χ^2s. The offset at the magnet χ^2 (offmag χ^2) describes how well the upstream track segment matches to the downstream segment at the bend plane of the magnet; there is one offmag χ^2 for each track. The vertex χ^2 globally describes how consistent the two tracks are with coming from a common vertex.

The χ^2s are calculated using the average tracking resolution. In addition, the vertex χ^2 includes a simple term to describe multiple scattering of the tracks in the vacuum window and DC 1. However, there are non-gaussian tails in the resolutions, mostly caused by accidentals combined with the high-SOD problem described in Chapter 4. The vertex χ^2 distributions for both the regenerator and vacuum beams are shown in Figure 6.8. The large tails are apparent in the figure. Our Monte Carlo predicts a large part of these non-gaussian terms, but we do not want to be too dependent on the specific details of the MC simulation. Hence we cut rather loosely on the χ^2 quantities: vertex $\chi^2 < 500$ and offmag $\chi^2 < 100$.

6.6.4 Veto Counters

We used the veto counters to reduce backgrounds. The main veto cuts were in the muon system, to suppress $K_{\mu 3}$, and in the regenerator, to reduce kaon scatter background. We also cut on the remaining photon vetoes, mostly for consistency with the $\pi^0\pi^0$ analysis.

The muon hodoscopes were in veto in the trigger to remove $K_{\mu 3}$ events. However, the veto was not 100% effective at trigger level. Due to timing jitter in the trigger electronics and the propagation delays along the lengths of the counters, some events that should have be vetoed at trigger level remained in our data sample. We remove events for which there are any hits in the muon hodoscope counters. Cutting on the

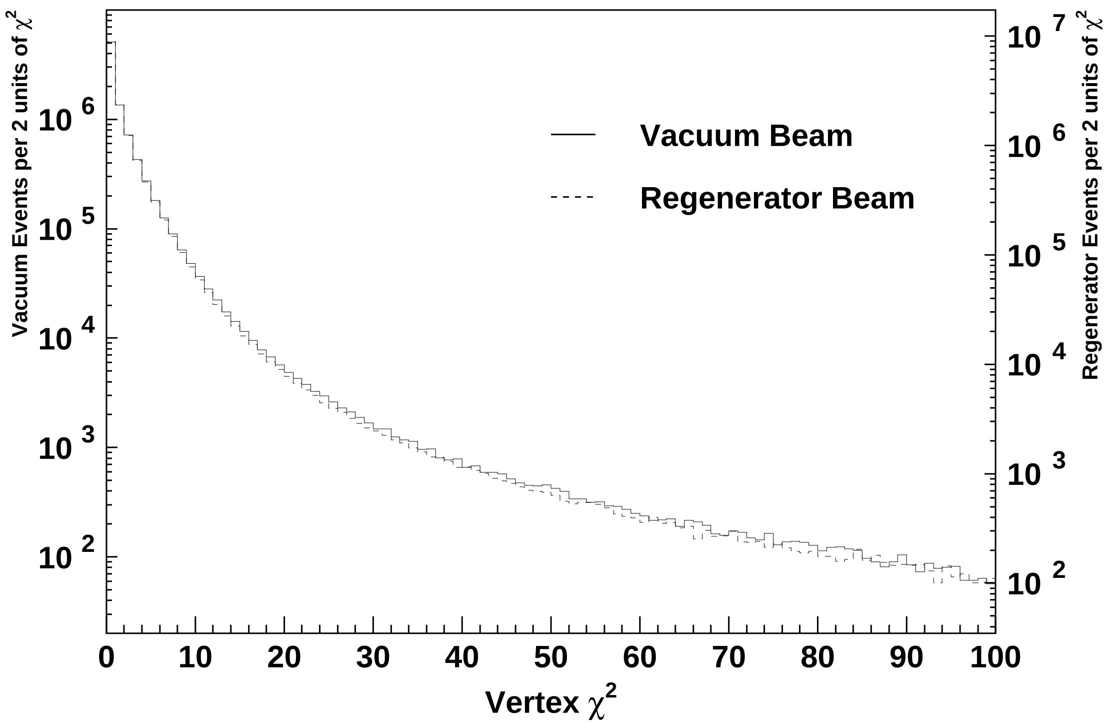

Figure 6.8: Vertex χ^2 for the two beams.

muon hodoscope removed 0.27% (0.11%) of otherwise good events in the vacuum (regenerator) beam. The discrepancy indicates that we are preferentially cutting against $K_{\mu 3}$ events, which are suppressed in the regenerator beam.

To suppress backgrounds from scattering in the regenerator, we cut on the energy deposited in the regenerator. We require that no regenerator module have more than the equivalent energy of two minimum-ionizing particles (MIPs). In addition, we make a tighter requirement on the last module, which is the lead-scintillator sandwich. There is less scintillator material, and the lead causes extra scattering and increases the rate in the final module, which in turn causes large pedestal shifts. We correct for the pedestal shifts by subtracting the signal integrated in a separate ADC gated for the bucket before the in-time bucket. We cut at 0.7 MIPs on the pedestal-corrected energy in the final module. To ensure that the pre-bucket subtraction works correctly, we veto any events in which the regenerator trigger sources fired in the previous or post buckets.

For consistency with the $\pi^0\pi^0$ analysis, we also cut on the "outer" veto counters: the Mask Anti, the Ring Counters and the Spectrometer Antis. These cuts are not strictly necessary, since no additional background to $K \to \pi^+\pi^-$ is suppressed by them. The MA suffered from large pedestal shifts, which we avoided by subtracting the minimum ADC channel from the maximum. We cut on the energy deposition in the MA equal to a 100 MeV photon. In addition, we discard events for which the MA trigger source fires. For the Ring Counters (RCs), we cut on trigger sources only, again because of baseline shifts. The Spectrometer Antis (SAs) did not have large baseline shifts, so we are able to cut on the ADC values, at an energy equivalent to 300 MeV photon.

Cuts on the "outer" photon vetoes cannot bias the $\mathrm{Re}(\epsilon'/\epsilon)$ analysis in the charged-mode analysis, since any real event leaves no energy in these counters. We have to be careful not to cut on veto counters into which energy from the real particles might "splash back." For this reason, we do not cut on the Collar Anti (CA) or the Back Anti (BA), unlike the $\pi^0\pi^0$ analysis, because hadronic interactions from the $\pi^\pm$s can leak into these detectors. This splashback would then have a position dependence, and could bias one beam with respect to the other beam.[2]

6.6.5 Extra-Particle Cuts

We remove events with any extra tracks besides the two from the vertex. In the calorimeter, we cut against extra electromagnetic clusters. We veto any event with an EM cluster above 1.0 GeV. To avoid satellite clusters from hadronic pion interactions, we require the extra cluster to be at least 20 cm from the extrapolated track position at the CsI.

[2] There is some evidence for splashback in the Cesium Iodide Anti (CIA), but it is below the level of the cut.

6.6.6 Pion Identification Cuts

The K_{e3} and $K_{\mu 3}$ decays occur at a much higher rate than the $K \to \pi^+\pi^-$ decay, and would swamp the signal mode unless we took measures to veto these events. Electrons deposit all their energy in the CsI calorimeter, and can be identified as tracks that have an energy to momentum ratio (E/p) near to unity. We identify pions by requiring tracks to have $E/p < 0.85$. Muons from $K_{\mu 3}$ decay leave a signal in the muon hodoscope; however, low momentum muons may range out in the filter steel before depositing energy in the muon hodoscope. We require each track to have a momentum $p > 8$ GeV/c to ensure 100% efficiency for the muon veto hodoscope.

6.6.7 Kinematic Cuts

In addition to kaons and neutrons, the KTeV beams contain neutral lambda particles. These lambdas are a background to $K \to \pi^+\pi^-$ through the decay mode $\Lambda \to p\pi$, where the proton is mis-identified as a pion. The particle identification cuts in the calorimeter do not discriminate against this lambda background because the p interacts similarly to the $\pi^\pm$. Fortunately, it is easy to cut against $\Lambda \to p\pi$ decays kinematically. We reconstruct the two tracks, assuming the higher momentum particle is the p. If the reconstructed $m_{p\pi}$ is within the range 1.112 GeV/c^2 to 1.119 GeV/c^2, we reject the event. The lambda cut is made at all energies, even though the real Λs in the beams are present only at higher energies. In addition, the Λ cut removes some real kaon events.

The final kinematic variables are the invariant mass, $m_{\pi^+\pi^-}$, and the square of the transverse momentum, p_T^2, calculated with respect to the z at the downstream end of the regenerator. Figure 6.9 shows the invariant mass distributions for the two beams. The shapes of the distributions are nearly the same, and have an RMS resolution of $\sim$ 1.6 MeV. The mass distributions are slightly asymmetric on the low side due to the presence of $K \to \pi^+\pi^-\gamma$ events in both beams.[3] The tails in

[3] The decay $K \to \pi^+\pi^-\gamma$ has contributions from decays where the photon is directly emitted from the interaction vertex ("DE") as well as from bremsstrahlung from one of the pions ("IB"). The DE photon is quite stiff, and shifts the $\pi^+\pi^-$ mass outside of the signal region.

the vacuum beam distribution are due to the semi-leptonic backgrounds $K_{\mu 3}$ and K_{e3}, which are suppressed in the regenerator beam due to the dominance of the K_S component. To select good $K \to \pi^+\pi^-$ events, we require that $m_{\pi^+\pi^-}$ be between 488 MeV/c^2 and 508 MeV/c^2.

Most of the remaining events in the data sample are coherent $K \to \pi^+\pi^-$ events. In the vacuum beam, there are some remaining semi-leptonic events. In the regenerator beam, there are still some scattered events not vetoed by the regenerator. It should be noted that without the regenerator veto cuts, the regenerator beam would be dominated by scattered events. As mentioned above, we can gain further discrimination against scattered events by cutting on p_T^2. Figure 6.10 shows the

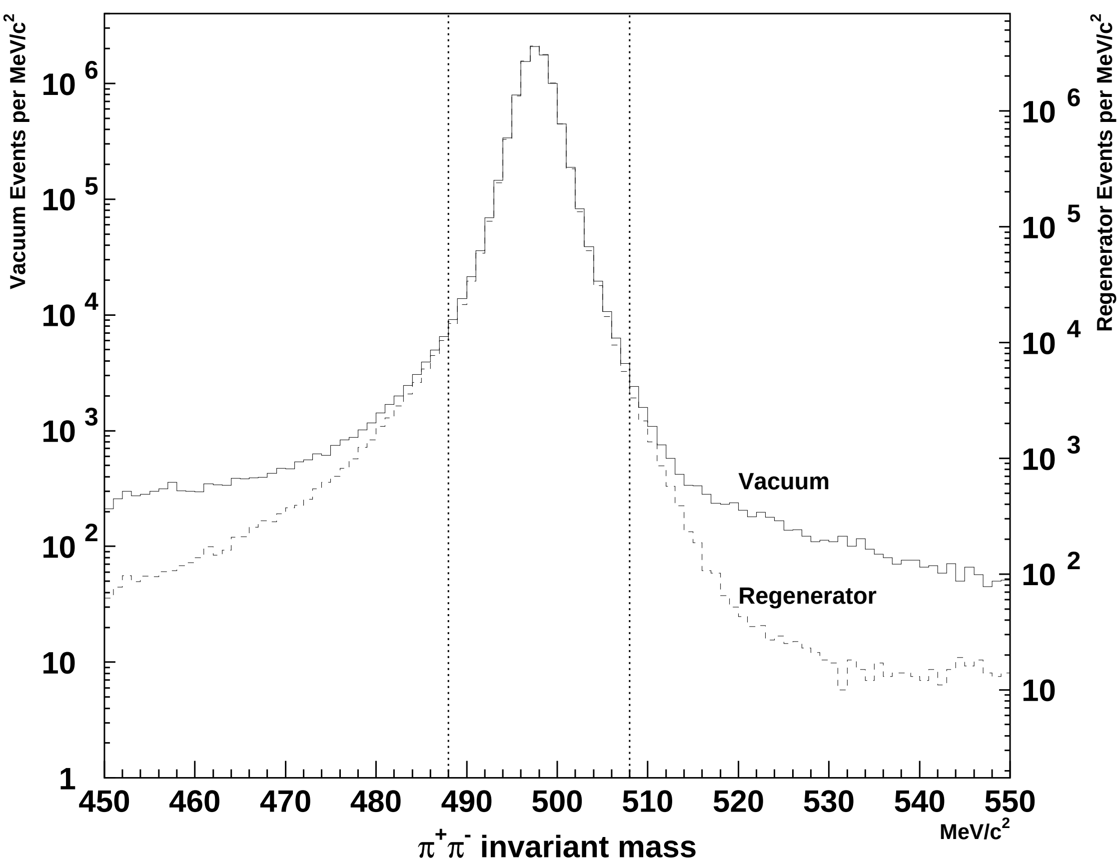

Figure 6.9: $\pi^+\pi^-$ invariant mass distributions after all other analysis cuts are applied. The dotted lines show the values of our cut.

p_T^2 distributions for the two beams overlaid on different scales. The peaks in p_T^2 are broadened due to the finite resolution of the detector, but are similar for the two beams. We cut on $p_T^2 < 250$ MeV$^2/c^2$ to reduce the levels of the various backgrounds. It is clear from Figure 6.10 that this is a tight cut; the resolution tail seems to extend to 500 MeV$^2/c^2$. Hence we will have to assign a systematic uncertainty that covers our understanding of the tails of this distribution.

6.6.8 Aperture Cuts

The KTeV detector has a number of physical apertures that limit the acceptance to $K \to \pi^+\pi^-$ decays. Due to the finite resolution of the tracking system, there is a limited precision to which we can know the location of these apertures. We cannot let the aperture itself define the edge since the pions can pass through a large amount of matter without being destroyed, and we would require a detailed Monte Carlo simulation of the passage of pions through the matter of the outer veto counters. Instead, we cut on the track projection near the aperture edges, since we have a better understanding of the simulation of the DC resolutions in the Monte Carlo.

Near the CA, we require the tracks to project at least 2 mm to the calorimeter side of the edge. Tracks are also required to project inside the outer edge of the CsI calorimeter by 2.9 cm, so the tracks can be matched to clusters. If the vertex position is upstream of the z position of the MA, the tracks are required to pass through the beam hole at least 3 mm from the edge of the MA. The position resolution is poor that far upstream, roughly 2 mm, but this projection cut is preferable to modeling the passage of a pion through the 16 interaction lengths of the MA. Interestingly, there is still a need to model the MA in the Monte Carlo, since it is possible for tracks from upstream decays to scatter in the MA and have the vertex reconstruct *downstream* of the MA z position.

To further define the acceptance to a region of the detector that can be well modeled in the Monte Carlo, we cut away from wires near the edges of the DCs, where the calibration of the chambers is not optimal due to low statistics in the calibration procedure. The outer wire cut varies amongst the chambers, from 12 to

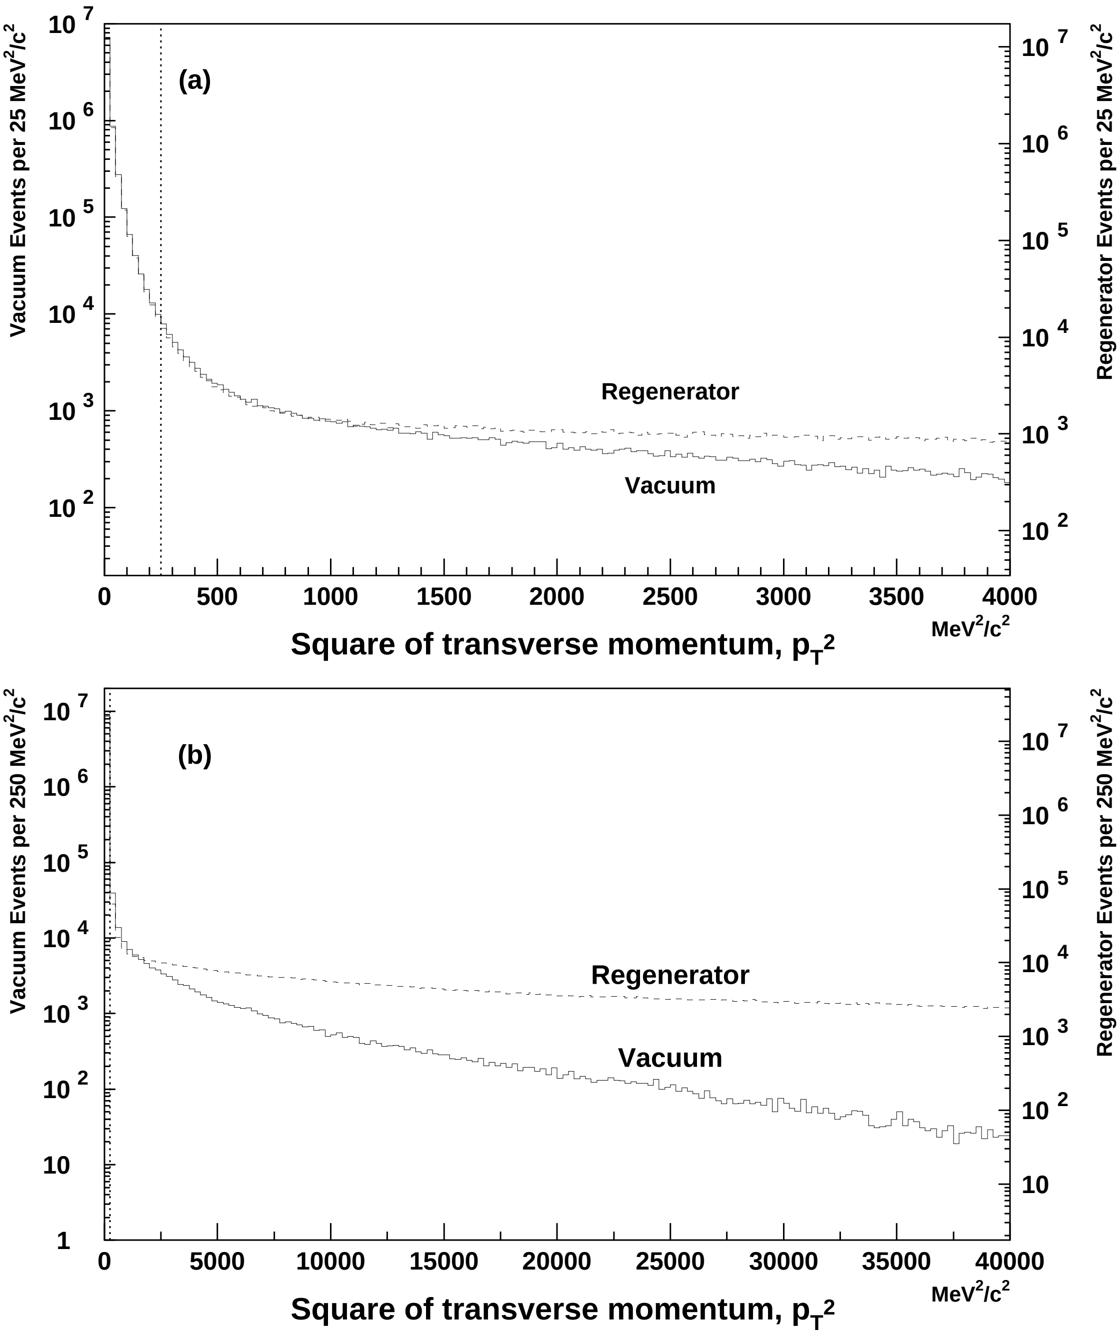

Figure 6.10: p_T^2 distributions for the data after all other analysis cuts have been applied. Panels (a) and (b) show the same data on two different horizontal scales. The dotted line shows our cut at 250 MeV^2/c^2.

40 wires. Drift Chamber 2 extends farther in the x dimension than necessary, and there is no illumination for tracks beyond the 40 wire cut.

When tracks are close to each other in an x or y projection, there is some probability that the x and y track candidates can be mismatched to each other, in which case the E/p of the tracks and the p_T^2 of the event are miscalculated. The miscalculation of E/p increases the K_{e3} background, and the miscalculation of p_T^2 causes the loss of good events. To avoid this problem, we require that the projections of the tracks at the CsI be separated by 3 cm in the x view and 3 cm in the y view.

The final aperture cut is the track separation cut. If tracks approach each other in one of the DCs, specifically within the same cell or adjacent cells, a hit from one track can obscure a hit from the other track. Due to the robust nature of the tracking code, an obscured hit should not cause the event to fail. However, tracks with missing hits are more susceptible to additional imperfections, such as high-SODs and accidental hits. We have made a large effort to model these effects in our Monte Carlo; however, we make a cut to get away from areas of the acceptance that are not modeled as well. Therefore we require a minimum separation between the tracks, in both the x and y views, at every DC. Each track is assigned a "cell" defined with the closest sense wire in center of the cell. We define the cell in this manner so that the boundary between cells is between two sense wires, where the efficiency is the highest and the geometrical symmetry minimizes potential biases. Figure 6.11 shows a diagram of the definition of the cells and the track separation cut. We require that the tracks be separated by 3 cells at each chamber. The track separation cut is the most extreme cut in the analysis, cutting 18.4% (10.2%) of otherwise good events in the vacuum (regenerator) beam. The large difference between the two beams is due to the fact that the vacuum beam has more events decaying downstream than the regenerator beam, and these events tend to have smaller track separations at DC 1. The implementation of this cut has a large effect on the shape of the vertex z distribution at the downstream end of the decay region, as shown in Figure 6.12.

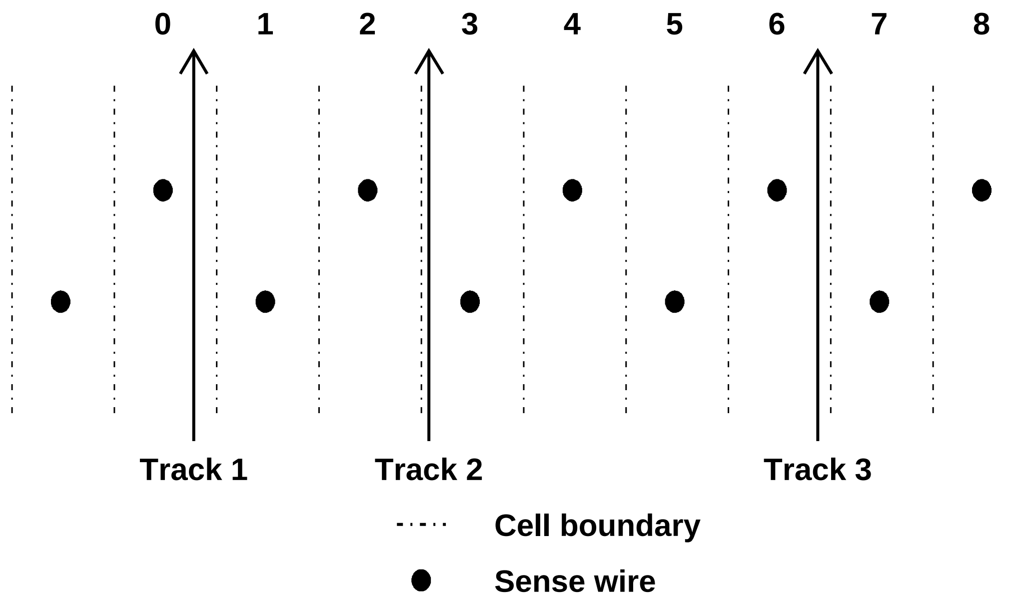

Figure 6.11: Definition of "cells" used in the track separation cut. The cell edges are defined halfway between sense wires. In the diagram, Track 2 is 3 "cells" from Track 1, while Track 3 is 6 cells away.

6.6.9 Fiducial Cuts

Figure 6.13 shows the momentum distribution for both beams. The shapes of the two distributions are qualitatively similar, due to the relatively weak dependence on energy of the regeneration amplitude. For the nominal $\mathrm{Re}(\epsilon'/\epsilon)$ data sample, we require the kaon energy to be between 40 GeV and 160 GeV, mostly for consistency with the $K \to \pi^0\pi^0$ analysis.

The z-vertex distributions of the two beams are *not* similar, due to the K_S component in the regenerator beam. Figure 6.14 shows the two vertex distributions. The geometric cuts detailed in Section 6.6.8 largely define the acceptance in the vacuum beam; the sharp turn-on at the upstream end of the decay volume is due to the aperture cuts at the MA edges, while the roll-off at the downstream end is due to the track separation cut at DC 1. However, we make explicit cuts on the

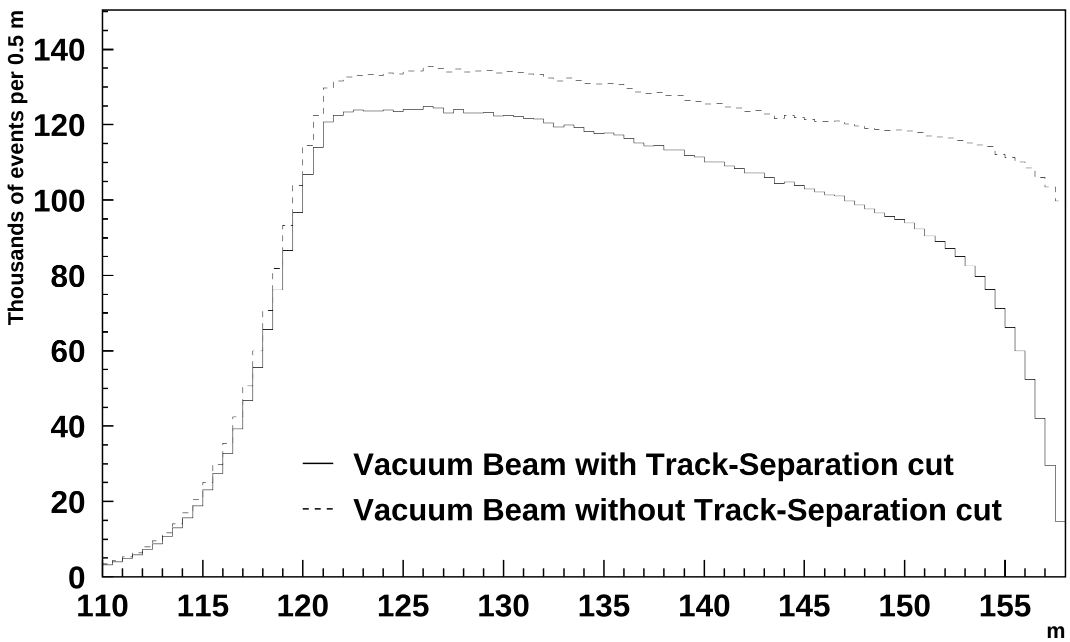

Figure 6.12: Vertex z distribution in the Vacuum beam before and after the track separation cut.

vertex z position in the two beams, requiring 110 m $< z <$ 158 m in the vacuum beam, and 122 m $< z <$ 158 m in the regenerator beam. There are a few events upstream of the MA in the regenerator beam, where the tracks can pass through the MA aperture and miss the front face of the regenerator; these events have no effect on $\mathrm{Re}(\epsilon'/\epsilon)$, but are explicitly removed by the vertex z cut.

Even with all the above cuts applied to identify $K \to \pi^+\pi^-$ events cleanly, we find evidence for a small beam halo, as a number of events reconstruct outside the nominal beams. We remove these events by requiring that the vertex (x, y) position reconstruct within a square of 75 cm^2, at the z of the downstream end of the regenerator, based on the point x_R as shown in Figure 6.7.

Table 6.3 shows the total number of events in our samples for our two time periods in the two beams. In these data samples, there is still a sizable background contribution (as can be seen in Figure 6.10). We must account for these backgrounds to complete our treatment of the data samples.

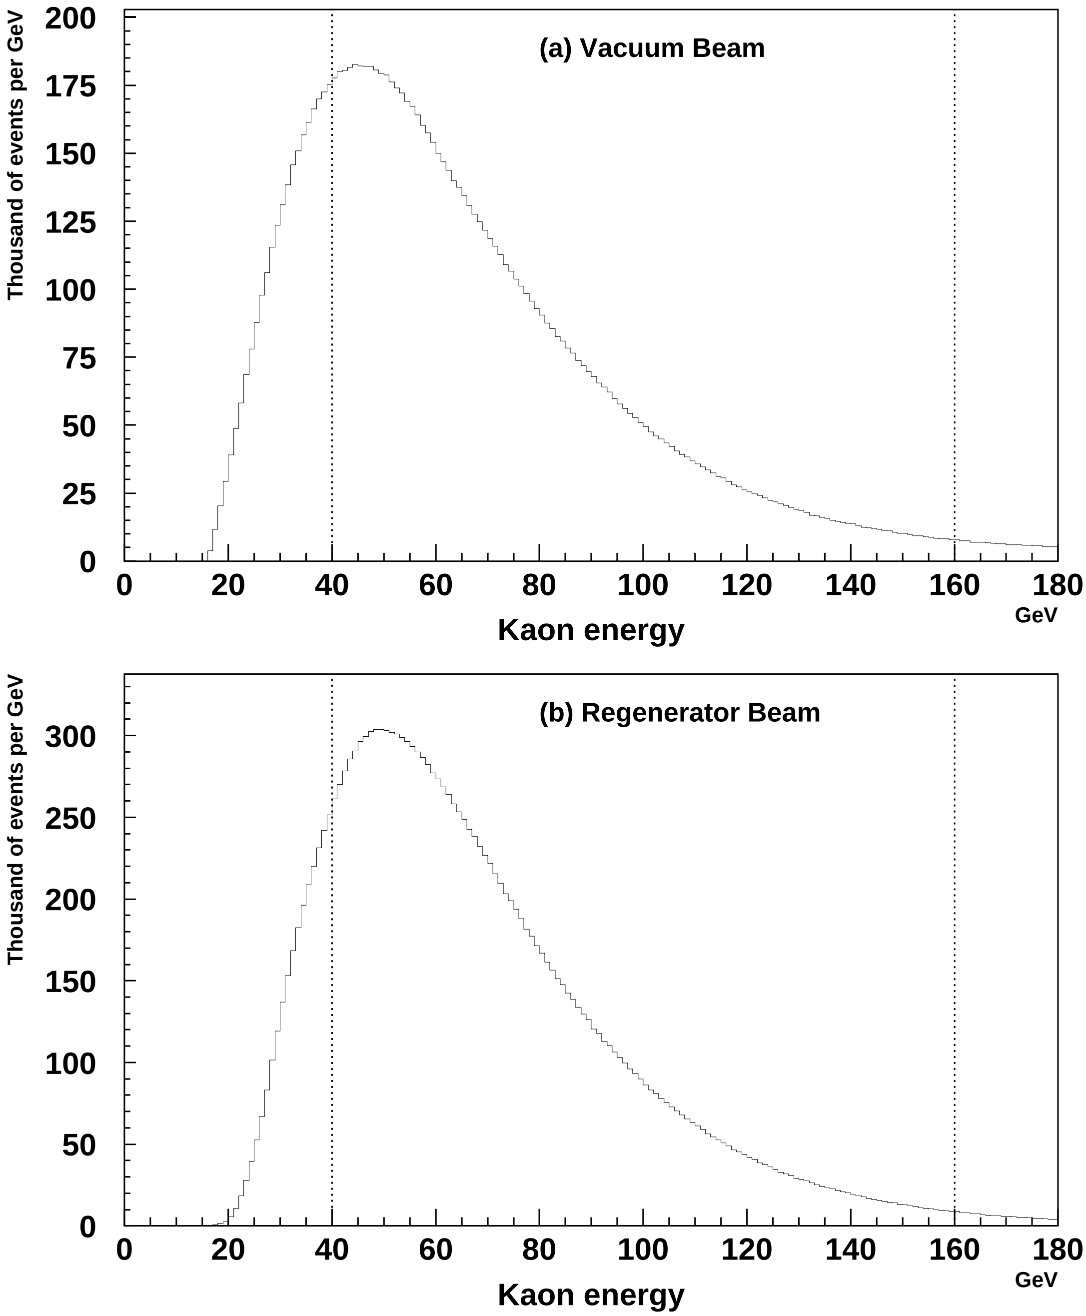

Figure 6.13: Kaon energy distributions for the $\pi^+\pi^-$ data after all other analysis cuts. The dotted lines show the range of energy used in the $\mathrm{Re}(\epsilon'/\epsilon)$ analysis.

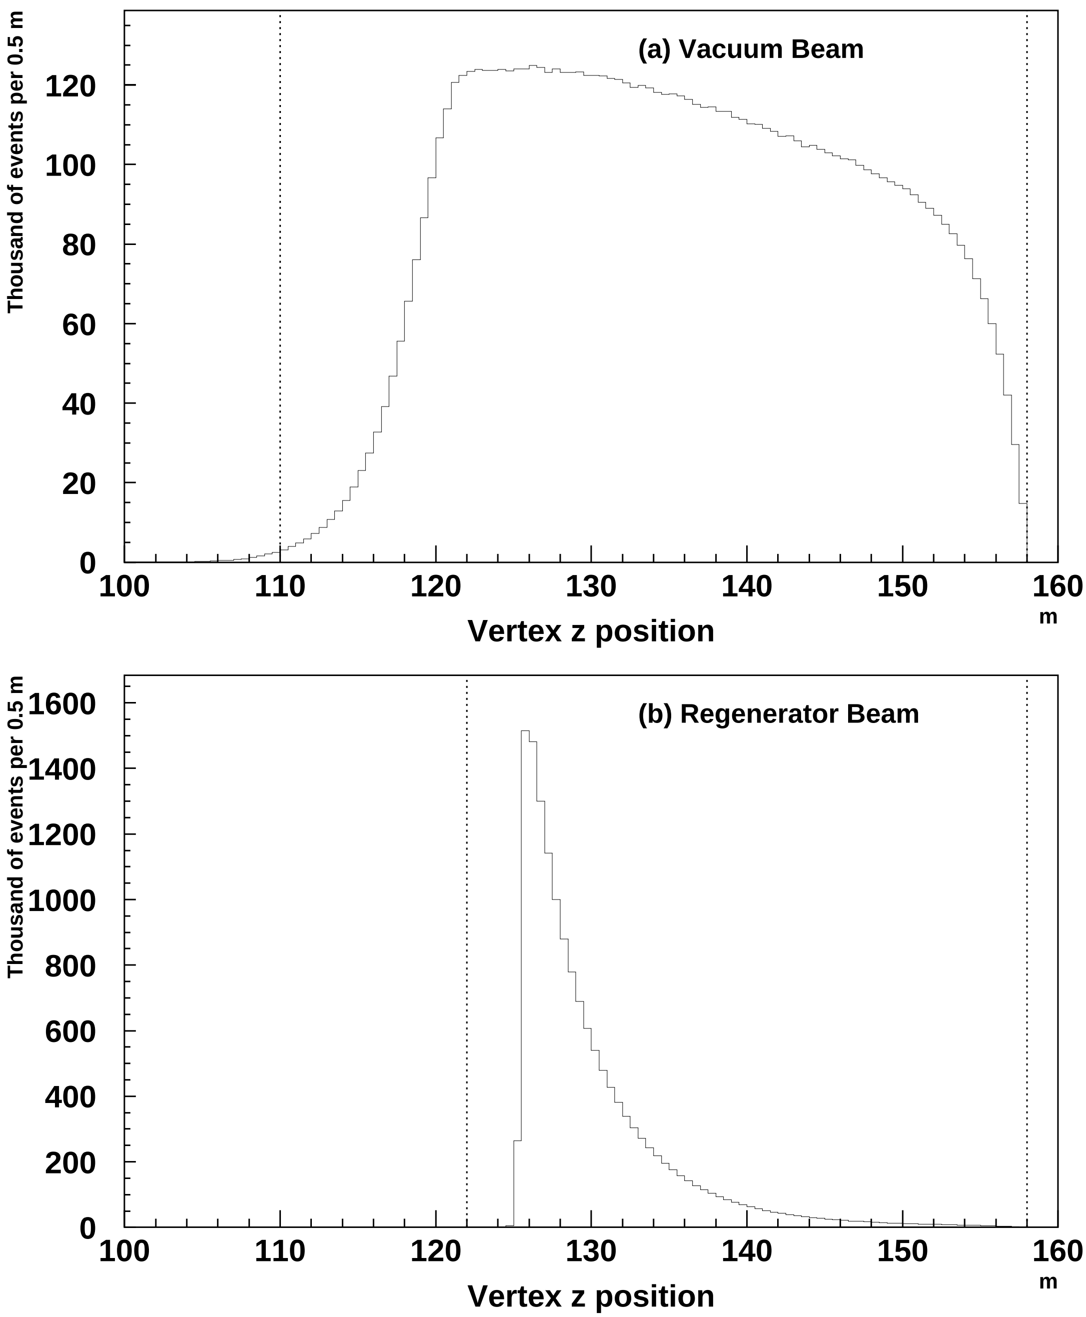

Figure 6.14: The vertex z distributions for the $\pi^+\pi^-$ data after all other cuts have been applied. The dotted lines show the z range used in this analysis.

Table 6.3: Number of $K \to \pi^+\pi^-$ events before background subtraction in the two beams.

Beam	1997A sample	1997B sample	Total
Vacuum beam	2,523,989	8,606,983	11,130,972
Regenerator beam	4,372,532	14,918,614	19,291,146

6.7 Background Subtraction

The well-defined kinematics of $K \to \pi^+\pi^-$ decays make background suppression quite straight-forward. In addition, the excellent pion particle identification (PID) capabilities of the KTeV detector allow for further suppression. However, as can be seen from Figure 6.10, there is a remaining contribution from backgrounds, of order 0.1% in both beams, to our data sample which would affect the value of $\mathrm{Re}(\epsilon'/\epsilon)$ and must be subtracted. An added bonus of understanding the backgrounds in the charged-mode is a better understanding of the backgrounds that contribute to the neutral-mode analysis.

6.7.1 Background Processes

The pion identification cuts remove the majority of the semi-leptonic background events, leaving a background contribution of roughly 0.1% of the signal in the vacuum beam. For a semi-leptonic event to remain in the sample, it must a have a rare interaction (or lack of one) in the detector. Electrons from K_{e3} decays can be identified by an E/p value near 1.0; the nominal electron sample has an RMS on E/p of 0.75%. For an electron to "fake" a pion, it must deposit less than 85% of its energy in the calorimeter. We see evidence for a tail on the electron distribution down to an E/p of 50%, possibly due to rare photo-nuclear interactions. For a muon to fake a pion, it must fail to fire the muon hodoscope veto. The counters of the muon hodoscope run at nearly 100% efficiency; a more likely occurrence is that the muon ranges out in the muon filter steel in front of the muon hodoscopes. A muon above the 8 GeV/c cut should pass through the filter steel, although there are scattering processes than can prevent the muon from reaching the hodoscopes.

We cannot reliably predict the rate of the rare processes that allow semi-leptonic events into our data sample, so we measure the rates of these events from the data. The semi-leptonic events do not contribute greatly to the regenerator beam, since the K_L component is suppressed by roughly a factor of 13.

A common background to both the vacuum and regenerator beams is kaons scattering in the final defining collimator at $z = 86$ m. There is a relatively large flux of kaons on the defining collimator, and although no kaon can pass through the 3 m of steel, there is some chance of a kaon scattering back into the beam hole if it starts out close enough to the edge of the collimator, as shown in the cartoon in Figure 6.15.

These events have been studied by selecting good $K \to \pi^+\pi^-$ events in the vacuum beam with $p_T^2 > 5000$ MeV$^2/c^2$ where the total momentum vector of the kaon points back to the edge of the collimator. At higher energies, the events seem to be K_S events with high p_T^2. At lower energies, there is a K_L component with a lower characteristic momentum transfer. The model of the background is based on

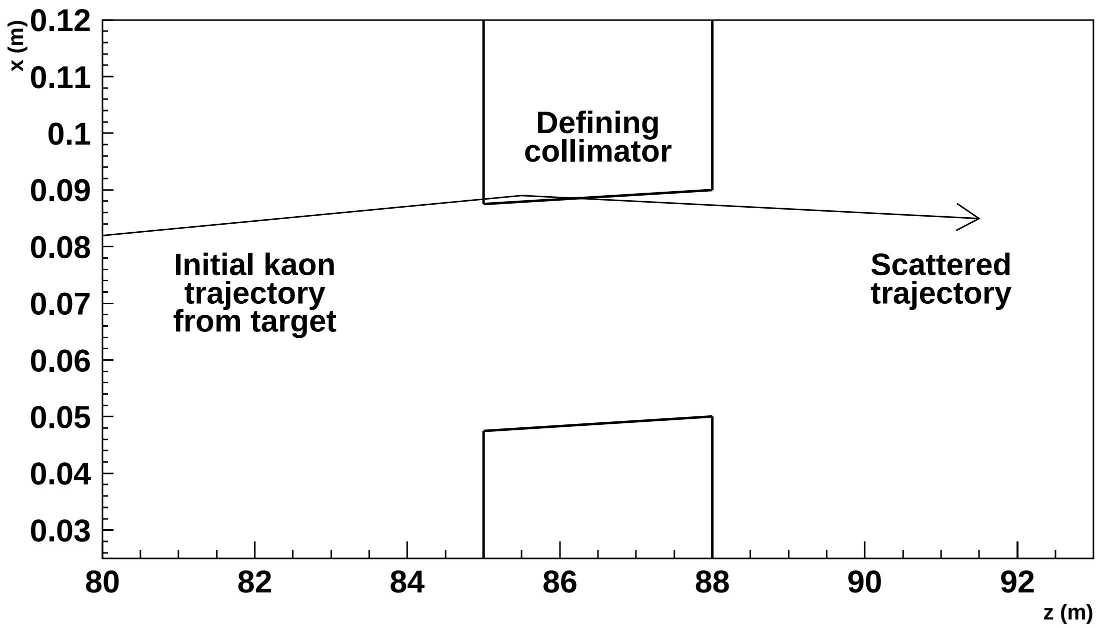

Figure 6.15: Cartoon of kaon scattering in the final defining collimator.

Figure 6.15 and matched to the data. The incident K_L can scatter to a K_L or K_S state, and the p_T^2 imparted to the final state is chosen from an exponential distribution, with different characteristic scales for the K_S and K_L cases. The relative probabilities of the two cases are tuned to match the data.

The dominant background in the regenerator beam is kaons that scatter in the regenerator and decay to $\pi^+\pi^-$. We do not count these events in our sample, since we are interested in coherent events only. We have decided to consider coherent events as our data sample to be able to make use of the relatively tight p_T^2 cut in the charged-mode analysis, helping to remove our semi-leptonic backgrounds. However, due to the method of reconstruction of $\pi^0\pi^0$ decays in the neutral calorimeter (see Chapter 7), we cannot make a similar p_T^2 cut in the neutral-mode analysis, meaning there would be a large number of scattered events in $K_S \to \pi^0\pi^0$ compared to $K_S \to \pi^+\pi^-$. To avoid this problem, we undertake a background subtraction of the scattered $K_S \to \pi\pi$ events in both the charged- and neutral-modes.

A number of physical processes are expected to contribute, such as diffractive scattering from carbon and lead nuclei, multiple diffractive scattering, and some residual inelastic scattering not vetoed by the regenerator. Each of the scattering processes has a different characteristic momentum transfer, and hence a different effect on the kaon state, which can be parametrized as a different effective regeneration amplitude, ρ_{eff}, for each of the terms. These effective regeneration amplitudes, in turn, lead to a complicated distribution of scattered events in kaon energy, proper time or z downstream of the regenerator, and p_T^2. Attempts have been made to quantitatively predict these scattering terms[70] for the K_{e3} asymmetry measurements where the scattering backgrounds are much larger. However, in this case, we cannot hope to predict *a priori* the scattering background to the level necessary, so we fit the data in bins (of p_K, p_T^2, and proper time τ) to a physics-motivated function of the form

$$\begin{aligned}\frac{\mathrm{d}^3N}{\mathrm{d}p_T^2\,\mathrm{d}\tau\mathrm{d}p_K} &= M(p_K)\times T(p_K)\times S(p_K)\\ &\times \sum_{j=1}^{n}\mathbf{A_j}e^{-\mathbf{B_j}\cdot p_T^2}\frac{\mathrm{d}N}{\mathrm{d}\tau}\left(\left|\rho^j_{\text{eff}}\right|,\phi^j_{\text{eff}}\right), \qquad (6.6)\end{aligned}$$

where

$$
\begin{aligned}
M(p_K) &= K^0 + \overline{K^0}\ \text{Energy spectrum} && (6.7)\\
T(p_K) &= \text{kaon transmission} && (6.8)\\
S(p_K) &= \text{absorber scatter correction} && (6.9)\\
\frac{\mathrm{d}N}{\mathrm{d}\tau}\left(\left|\rho^j_{\text{eff}}\right|, \phi^j_{\text{eff}}\right) &= \left|\eta e^{i\Lambda_L \tau} + \rho^j_{\text{eff}} e^{i\Lambda_S \tau}\right|^2, && (6.10)
\end{aligned}
$$

and the index j runs over the different scatter terms. The terms M, T and S were added to the fit since our first result[55], and improve the data-fit match. M accounts for the incident kaon momentum spectrum. T includes two effects: regeneration in the upstream absorber, which enhances higher p_K; and the measured attenuation in the regenerator beam relative to the vacuum beam, which enhances lower p_K. The S term accounts for events that scatter in the absorbers but remain in the regenerator beam sample. The scatter correction is from a Monte Carlo study that compares the p_K distribution after the defining collimator with the p_K distribution at the target. The p_K-dependent factors above are shown in Fig. 6.16 as a function of p_K. The p_K factors are taken to be the same for each scatter term. The combined effect tends to cancel and gives a rather small correction due to kaon transmission. However, there is a need for a final correction to the momentum-dependence of the scattering model to finalize the match between data and the model. This correction is shown in Figure 6.17.

Although the diffractive scattering events form the dominant part of the regenerator beam background, there are still some remaining inelastic scattering events that are not vetoed by the regenerator. The remaining level of inelastic events depends on the performance of the regenerator, and so we use parameters in the model that are measured in the charged-mode for a given period of time. For example, although we do not use the $K \to \pi^+\pi^-$ events from 1996 to measure $\mathrm{Re}(\epsilon'/\epsilon)$, we do use them to measure the regenerator scattering background so that a period-specific parametrization can be used to subtract the regenerator-scattering background in the neutral mode data for the 1996 period.

The improvement of the momentum-dependent terms of the fit function (Equa-

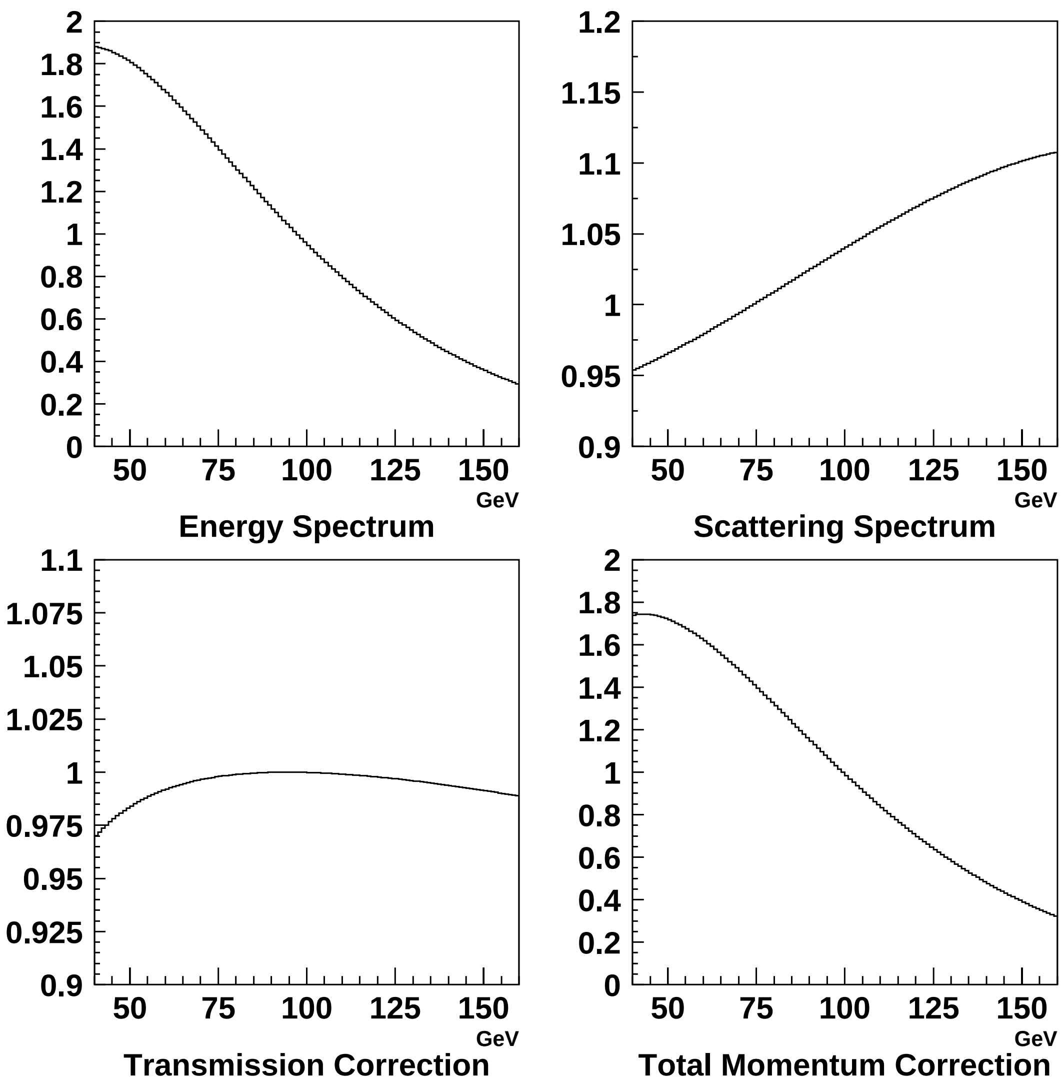

Figure 6.16: The three momentum-dependent terms used in the regenerator scattering model (Equation 6.6), and their product.

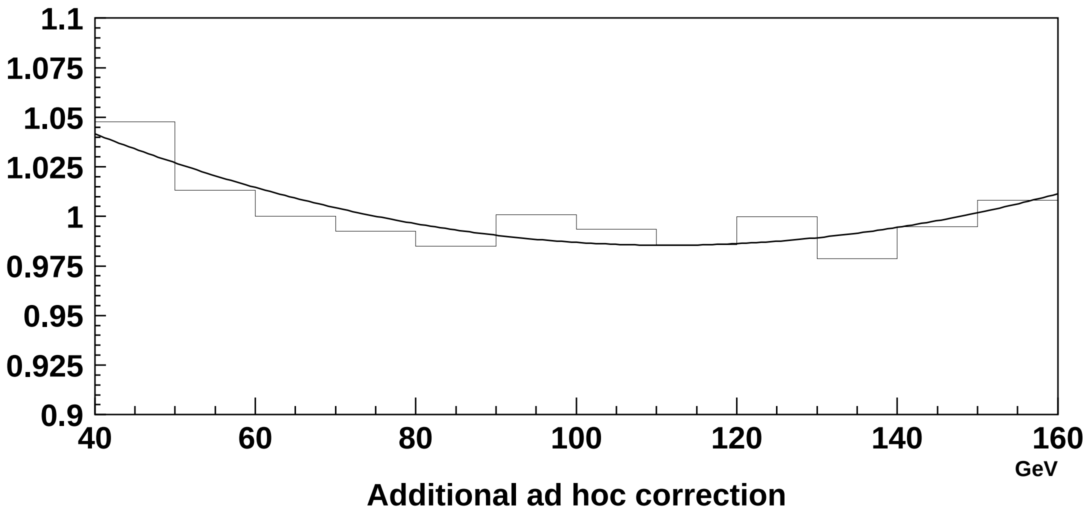

Figure 6.17: Final p_K-dependent correction applied to Equation 6.6 to match the scattered $K \to \pi^+\pi^-$ data.

tion 6.6) constitutes a major improvement in the modeling of the scattering in the regenerator over References [2, 55]. In fact, in determining these improvements, an error was discovered in the application of the previous version of the p_K correction to the model. This error biased the scattering model and introduced a mistake in our estimation of the neutral-mode scattering background, and caused a $+1.7 \times 10^{-4}$ shift in the value of $\mathrm{Re}(\epsilon'/\epsilon)$. This error will be discussed in more detail when we discuss systematic uncertainties in Chapter 10.

The final background we consider is hadronic production in the regenerator, which we call “Regenerator junk.” Two candidate hadronic production mechanisms are $n \to p^{\pm}\pi^{\mp}$ and K^* production and decay. These backgrounds only occur at the end of the regenerator.

6.7.2 Normalization of Background Contributions

Once we have identified the possible background contributions, we need to normalize the contribution of each process within our data sample. As mentioned above, a first-principles calculation of the absolute background rate is difficult due to the rare nature of the processes that the backgrounds must undergo to end up in our sample. Instead of an absolute prediction, we normalize the backgrounds to the data, using regions of the p_T^2 and $m_{\pi^+\pi^-}$ space outside of our cuts (the "sideband" regions) shown in Figure 6.18. The background shape has a few distinguishing features: the collimator scatter background can be seen as a vertical stripe of good $m_{\pi^+\pi^-}$-large p_T^2 events. The semi-leptonic decays form a broad background in both $m_{\pi^+\pi^-}$ and p_T^2. Due to kinematics of the semi-leptonic decays, it is difficult to distinguish K_{e3} from $K_{\mu3}$ decays, although there are more K_{e3} decays at higher $m_{\pi^+\pi^-}$.

One approach to determining the background might be to try to fit the background shape directly to either a 2-dimensional function or to shapes determined from Monte Carlo. We will use the method of fitting directly to MC shapes as a systematic cross-check in Section 10.8. However, to gain additional information, we use the particle identification capabilities of the KTeV detector to determine the background levels, by attempting to measure the background rates by applying cuts to *enhance* a particular background contribution, and then determine the background level in these enhanced samples. For example, to enhance the K_{e3} sample, we require the electron-like particle to have $E/p > 0.75$, and that the pion-like particle deposit at least 1 GeV of energy in the CsI (to ensure it is not a muon). To enhance $K_{\mu3}$ we require the muon-like particle be minimum-ionizing while the other particle has $E/p < 0.5$, to ensure it is not an electron. For the collimator sample, we enhance the two-pion sample by requiring both particles to have $E/p < 0.5$, neither to be minimum-ionizing, and that the total momentum vector of the two-track system point back to the collimator when extrapolated back to $z = 86$ m. Figure 6.19 shows how these cuts work: the left column shows the data with the enhancement cuts applied, while the right column shows the Monte Carlo simulation of the various backgrounds. It is clear from the figure that the enhancement cuts have the desired effect. For example, in the "collimator-scatter" sample (Panel (c)), the semi-leptonic decays are suppressed. The polygons on the figure show the regions used to normalize the background MC samples to the data. The efficiencies of each of the enhancing cuts is applied to each of the normalization factors to extract the normalization of a given background process to the data. Table 6.4 lists

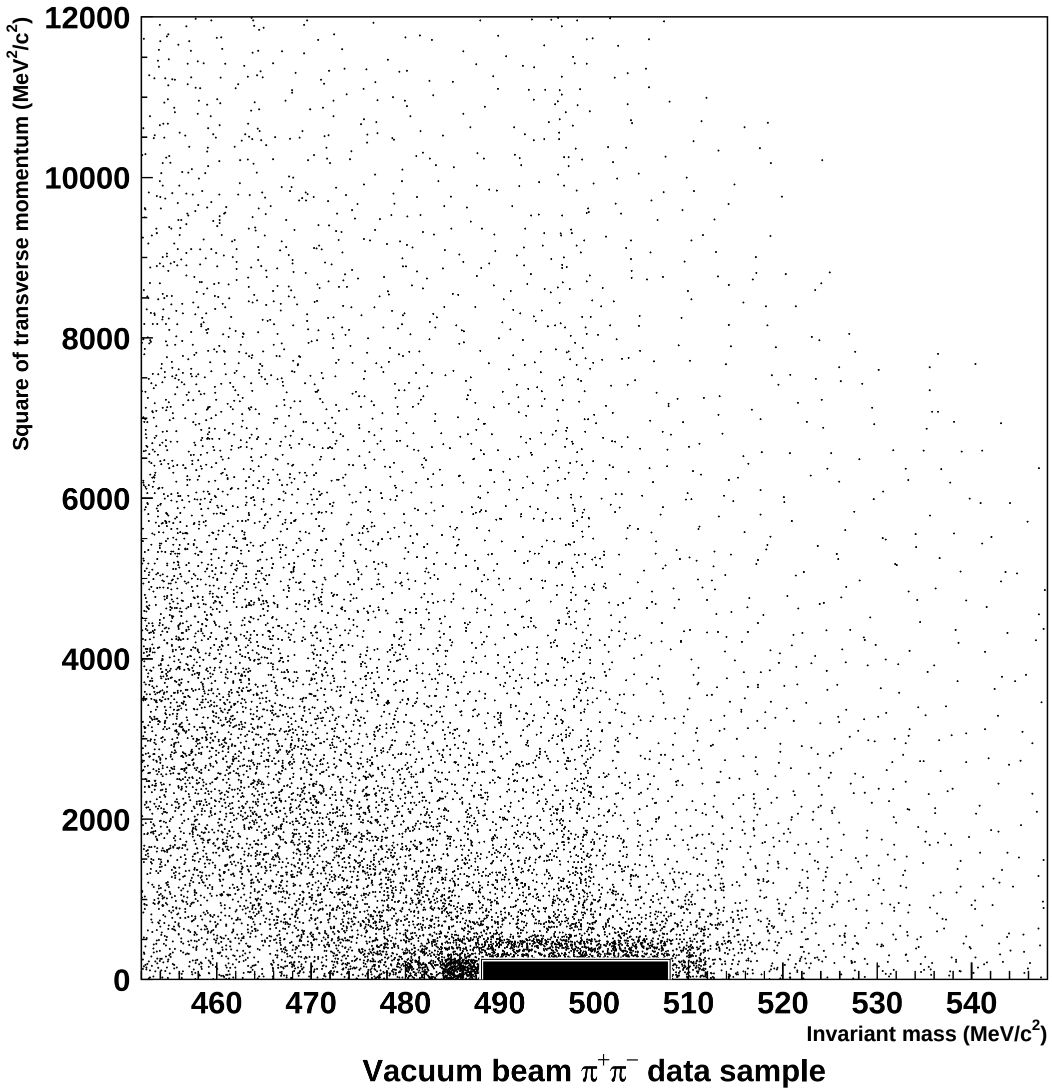

Figure 6.18: Events in the vacuum beam for $m_{\pi^+\pi^-}$ vs p_T^2. Each point represents 60 events. The signal region is delimited by the black box. The semi-leptonic backgrounds can be seen, especially at low $m_{\pi^+\pi^-}$, and the collimator scatter background can be seen as the vertical band in the mass region.

the cuts used to enhance each sample, and the relative efficiencies of these cuts. It should be noted that this method is slightly more robust than simply fitting the background shapes, since it allows us to use the particle-identification capabilities of the detector to identify the various background components, at the cost of some statistical precision. For an extreme example, even if the K_{e3} and $K_{\mu 3}$ backgrounds had the same shape in p_T^2 and $m_{\pi^+\pi^-}$, we could disentangle the relative contributions of each with this method. Such a disentanglement becomes important if the backgrounds had different distributions in z while having similar shapes in p_T^2 and $m_{\pi^+\pi^-}$. We scale the z-distributions of each of the background processes by the determined normalization factor and subtract the distribution from the data.

We do not have a detailed Monte Carlo prediction for the hadronic "junk" background in the regenerator beam, but since these backgrounds are small, we use a direct sideband subtraction from the data. We measure the mass sideband in the regenerator beam with $m_{\pi^+\pi^-} > 510$ MeV/c^2 in the region 124 m $< z <$ 126 m, and subtract the sideband directly from our data.

The full background determination is done in independent 10 GeV energy bins. The $K_{\mu 3}$ normalization scale factor varies by a factor of three in the range 40–160 GeV, indicating that low momentum $K_{\mu 3}$ decays are more likely to have a low-momentum muon that stops in the filter steel. The normalization factors for the K_{e3} background are fairly constant as a function of energy, while the normalization for collimator scattering is flat by construction, since its normalization in the MC is determined from these studies. The normalization factors for K_{e3}, $K_{\mu 3}$ and collimator-scatters determined in the vacuum beam are used to set these levels in the regenerator beam. This procedure is especially important for the collimator scatters,

Table 6.4: Efficiency of background-enhancing cuts.

Background Sample	Cuts to enhance	Combined Efficiency
K_{e3}	Max$(E/p) > 0.75$, $E_\pi > 1.0$ GeV	0.30
$K_{\mu 3}$	$E_\mu =$ MIP, $E/p(\pi) < 0.5$	0.76
Collimator scatters	$(E/p)_{1,2} < 0.5$, $E_{1,2} >$ MIP	
	& p_K point at collimator	0.26

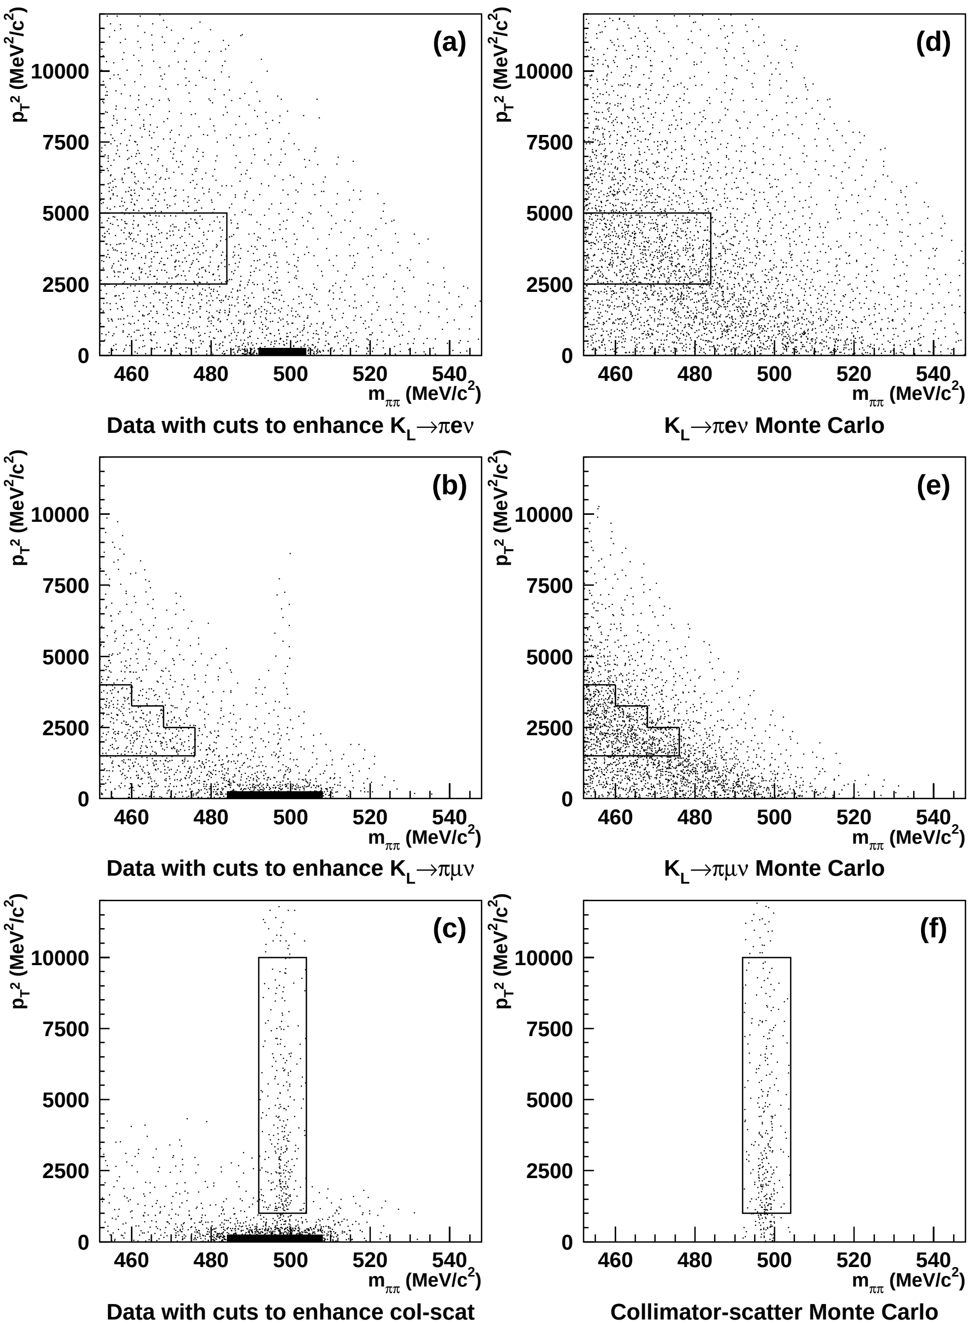

Figure 6.19: Distributions used to normalize background contributions in the vacuum beam $\pi^+\pi^-$ sample. Panels (a), (b), and (c) show the distributions of p_T^2 versus $m_{\pi^+\pi^-}$ for the data, with extra cuts to enhance K_{e3}, $K_{\mu3}$, and collimator-scatter events, respectively. Panels (d), (e), and (f) show the same distributions for the Monte Carlo predictions of these backgrounds. The outlined regions show the sideband areas used to normalize the Monte Carlo prediction to the data.

since there is no way to distinguish between kaons that scatter in the collimator and kaons that scatter in the regenerator. We subtract the collimator-scattered events in the regenerator beam, and the remaining scattered events are used as input to the fitting program to determine the regenerator-scattering parametrization discussed above, which is then used in the Monte Carlo.

Figure 6.20 shows the agreement between the background components and data at large p_T^2 within the mass cut. The fact that the semi-leptonic components match well is encouraging, since these processes were normalized using low mass events that are not in this distribution. In addition, the semi-leptonic events match the mass sidebands in vacuum beam, as shown in Figure 6.21.

The overall fractional background levels are listed in Table 6.5. The statistical error on the background level is quite small, but we assign a relative systematic uncertainty of 10% to account for possible deviations in the p_T^2 shape.

6.8 Summary

Table 6.6 shows the total number of $K \to \pi^+\pi^-$ events in our samples for the two time periods in the two beams. To proceed, we need to know the number of background-subtracted $K \to \pi^0\pi^0$ events in the two beams. Determining these numbers is the subject of the next chapter.

Table 6.5: Background levels for the $\pi^+\pi^-$ samples.

Background process	Vacuum (%)	Regenerator (%)
K_{e3} e mis-identified as π	0.036	0.001
$K_{\mu 3}$ μ mis-identified as π	0.054	0.002
Kaon scattered in collimator	0.010	0.010
Kaon scattered in regenerator	—	0.074
Hadronic “junk” produced at regenerator	—	0.001
Total	0.100	0.088

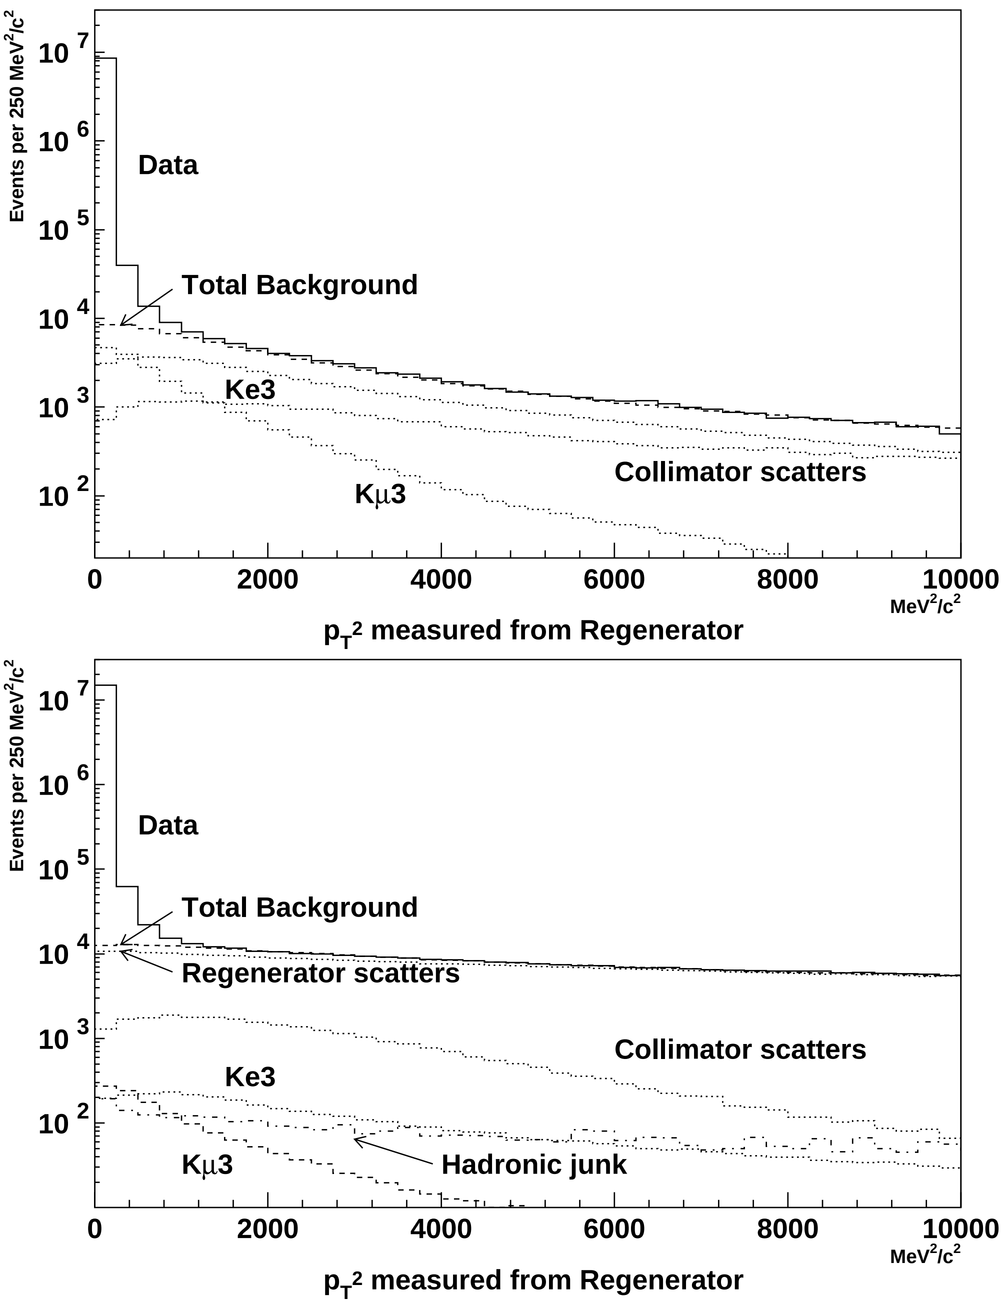

Figure 6.20: p_T^2 distribution after all other cuts for the vacuum beam (a) and regenerator beam (b), with the Monte Carlo prediction of the background levels. The analysis keeps events within the first bin of these plots.

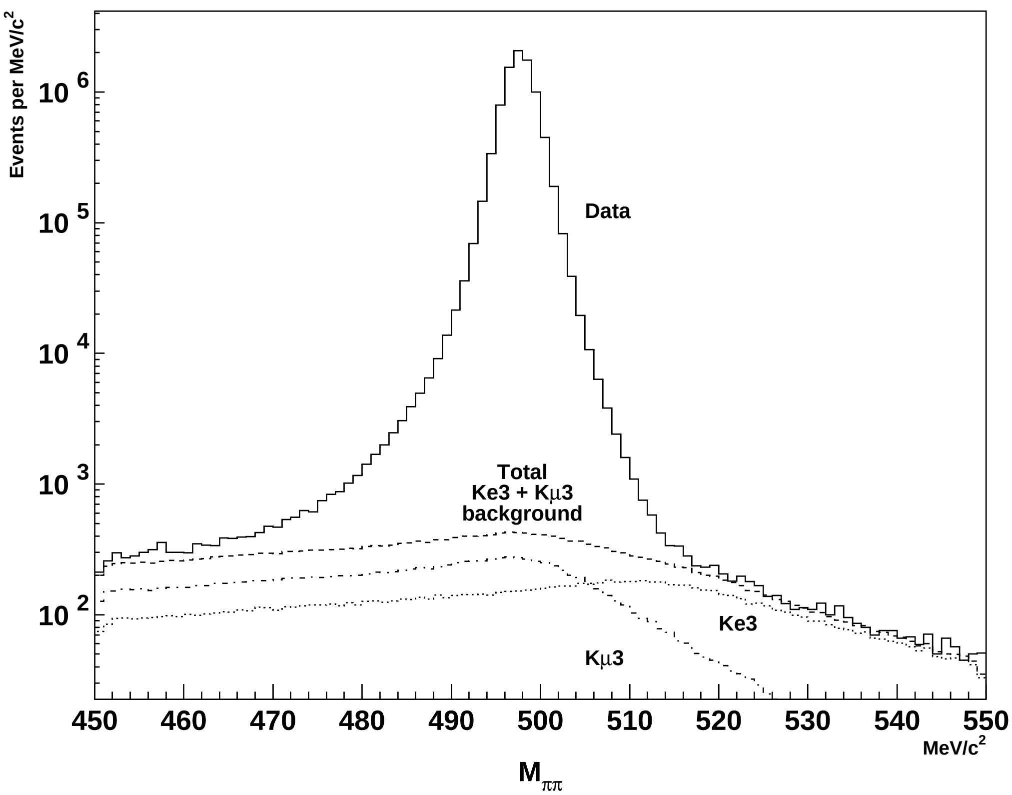

Figure 6.21: Vacuum beam $m_{\pi^+\pi^-}$ distribution for the data, with the semi-leptonic background prediction, after all other cuts have been applied.

Table 6.6: Number of $K \to \pi^+\pi^-$ events after background subtraction in the two beams.

Beam	1997A sample	1997B sample	Total
Vacuum beam	2,532,672	8,593,988	11,126,660
Regenerator beam	4,386,497	14,903,543	19,290,040

CHAPTER 7
SELECTION OF THE $\pi^0\pi^0$ SAMPLES

This chapter will describe how we collect and analyze $K \to \pi^0\pi^0$ events. Only the most basic presentation necessary for an understanding of the $\text{Re}(\epsilon'/\epsilon)$ measurement is presented. This topic is the focus of the dissertation of Valmiki Prasad [1], and any reader wanting to understand our measurement of $\text{Re}(\epsilon'/\epsilon)$ is urged to read that discussion.

The basic event topology is four photons reconstructed in the CsI calorimeter. This topology is required by the trigger, and the offline analysis consists of refining the measured energies and positions of the clusters and reconstructing the $K \to \pi^0\pi^0$ decay. After describing the analysis cuts, we will explain how the absolute energy scale is determined, and then discuss the subtraction of the background events present in this mode. At the completion of this chapter, we will have all the $K \to \pi\pi$ data events necessary to measure $\text{Re}(\epsilon'/\epsilon)$.

7.1 Trigger Requirements

The trigger requirement is to select events with four clusters of energy in the calorimeter consistent with $K \to \pi^0\pi^0$ decays. At Level 1, the trigger requires greater than 26 GeV of total energy deposited in the calorimeter. This sum is produced by the ET system, by summing the signals from the dynodes of all 3,100 channels in the calorimeter. As with the charged-mode trigger, the neutral-mode trigger can be vetoed at Level 1 if there is significant energy in the regenerator, the SA or the muon hodoscope. For part of the run, the Hadron-Anti was also in veto.

The common decay modes $K \to 3\pi^0$, $K \to \pi^+\pi^-\pi^0$, and $K \to \pi^\pm e^\mp \nu_e$ will generally deposit enough energy in the calorimeter to satisfy the neutral-mode trigger at Level 1. The rate is reduced at Level 2 by using the HCC to count the number of clusters of energy in the calorimeter. Once the ET "hot-bit" problem (see Section 4.2.2) was under control, the trigger required four and only four clusters above 1 GeV at Level 2.

At Level 3, we calculate the clusters' energies and positions, and the best and second-best "pairing" of photons to make a $K \to \pi^0\pi^0$. The definition of the "pairing χ^2" that determines the best pairing is given in Section 7.3. We required an invariant mass $m_{\pi^0\pi^0} > 450$ MeV/c^2 and a pairing $\chi^2 < 500$. In addition, we also kept the event if the second-best pairing passed the mass requirement and had a second-best pairing $\chi^2 < 50$.

7.2 Cluster Finding

The neutral analysis begins with finding clusters of energy in the CsI calorimeter and determining their energies and positions. This procedure involves many corrections to the raw digitized phototube signals, which were determined from many *in situ* studies of the calorimeter. The performance of the calorimeter is due to both its intrinsic properties and the large effort invested in understanding its behavior.

In the following sections, we describe how the raw phototube signals are converted to the energy of a cluster. In this chapter, we are concerned with measuring photon energies; however, the discussion will be general enough to apply to hadronic shower energy measurements as well, and differences between the two will be discussed.

7.2.1 Energy in a Crystal

The first step of reconstructing clusters is to unpack the calorimeter data and calculate the energy deposited in each crystal. For each DPMT time slice, the flash ADC value of the scintillation light is converted to an energy value, using a linear conversion for each QIE range. This conversion is determined in a calibration procedure using an adjustable laser pulse.

Studies of electrons from K_{e3} data show that there is a slight non-linearity in the nominal conversion, probably due to differences between the pulse shapes in the laser light and in CsI scintillation pulses. A correction of less than 1%, determined from electrons, is applied to remove this non-linearity. The final energy conversion is derived from K_{e3} events. The energies of the four in-time slices (4×19 ns) are summed together to determine the total energy in the crystal.

A final correction is applied to the raw block energy, based on the current response of the channel to flashes from the laser calibration system, which tracks short-term drifts in the PMT gains.

7.2.2 Cluster Seeds

To construct clusters, we need "seeds" around which to build the cluster. We consider two types of clustering in the calorimeter; "hardware" and "software" clustering, which specifically refer to the type of seed crystals involved.

Hardware clustering is done for events that satisfy triggers that require the HCC, as all neutral-mode triggers do. A block is considered to be a seed if it is the only block with an HCC bit set, or if its energy is greater than any neighbor which also has a set HCC bit.[1] Requiring the HCC bit ensures that the energy is in the in-time bucket. The seed block is also required to have at least 0.1 GeV of energy (far below the ~ 1.0 GeV HCC threshold) to avoid the HCC "hot-bit" problem.

In software clustering, the seed block is simply a block with greater energy than its neighbors. In the neutral mode analysis, we run software clustering after hardware clustering to find additional low-energy clusters. In the charged-mode analysis, we only run software clustering, since pion clusters can often be less than 1.0 GeV, and we do not require the HCC in the trigger.

7.2.3 Clusters

A cluster is defined as a square region around a seed block, either 7×7 small blocks or 3×3 large blocks. The "raw" energy of the cluster is the sum of the block

[1] Blocks touching diagonally at a corner are considered to be neighbors in this case.

energies. The x position of the cluster is determined by summing the block energies in the central column (the one containing the seed block) and in adjacent columns, and comparing the ratio of column energies to look-up tables derived from a sample of real photons. The y position is determined in the same manner for the rows of the cluster.

A series of corrections are applied to a cluster based on its energy and position. The energy is adjusted to correct for the intrinsic nonlinearity of the CsI crystals due to the nonuniformity of light collection along the length of the crystals, measured for each crystal using cosmic-ray muons. If a seed block is located near a beam hole or the edge of the calorimeter, the square cluster region may include non-existent blocks. The expected missing energy (based on the average transverse energy distribution of photon clusters) is added back to the cluster. We also correct for blocks that are present in the cluster but that may not have been read out due to the minimum channel read-out threshold (nominally 4 MeV). Finally, we make an "out-of-cone" correction for energy not included in the square cluster shape, again based on the average transverse energy distribution.

Overlapping clusters have the energy of the shared blocks divided between them based on the expectation from the average transverse energy distribution of a photon, and the energies and positions of the clusters are recalculated iteratively until the process converges to a self-consistent distribution of the shared energy. A small correction is also made for clusters that approach each other but do not overlap in the 7×7 or 3×3 crystal regions, to correct for "leakage" of energy outside of the square region. There is large leakage across the beam holes because shower particles can cross the beam hole and deposit energy on the far side. We apply a special correction derived from the data to handle such cases.

7.2.4 Final Corrections

The final corrections we apply are based on studies of the calorimeter response to electrons as a function of various factors. The largest effect is variation of the scintillation response of the CsI crystals as a function of transverse position. The response is generally lower, by as much as a few percent, near the edges of a given crystal compared to the center of the crystal. Correcting for this effect improves the offline energy resolution of the calorimeter.

Finally, some minor corrections are made to account for slow variations of the crystal response, due to long-time-scale temperature changes and radiation damage of the crystals. After we apply all these corrections, there is a small non-linearity. The remaining *ad hoc* linearity correction is of the order of 0.1%.

7.3 Event Reconstruction

To reconstruct $K \rightarrow \pi^0\pi^0$, we require four in-time hardware clusters in the calorimeter. There are three possible ways that four clusters can be paired into two π^0s. We must determine which is the best pairing. We assume that any two photons a and b come from the same π^0 and constrain the $\gamma\gamma$ invariant mass; from this information, we calculate the opening angle θ_{ab} between the two photons

$$\begin{aligned} m_{\pi^0} &= (E_a + E_b)^2 - |\vec{p_a} + \vec{p_b}|^2 \\ &= 2E_aE_b - \vec{p_a} \cdot \vec{p_b} \\ &= 2E_aE_b\,(1 - \cos\theta_{ab})\,. \end{aligned} \tag{7.1}$$

Since θ_{ab} is small, we make the approximation $(1 - \cos\theta_{ab}) \approx \theta_{ab}^2/2$ and $\theta_{ab} \approx r_{ab}/z_{ab}$, where r_{ab} is the radial separation between the two photons, and z_{ab} is the distance from the π^0 decay position to the CsI calorimeter. Hence

$$z_{ab} \approx \frac{\sqrt{E_aE_b}}{m_{\pi^0}} r_{ab}. \tag{7.2}$$

In general, only one of the ways of pairing the four photons will give a consistent z vertex position for the two π^0s ($z_{ab} \approx z_{cd}$). This is illustrated in Figure 7.1. We use the "pairing χ^2" to determine the best pairing of the photons:

$$\chi^2 = \frac{(z_{ab} - z_{cd})^2}{\sigma_{z_{ab}}^2 + \sigma_{z_{cd}}^2}, \tag{7.3}$$

where σ_z^2 is the position uncertainty due the finite energy and position resolution of the CsI calorimeter. Due to the excellent resolution of the KTeV calorimeter, the probability of incorrectly pairing photons in $K \to \pi^0\pi^0$ decays is only 0.03%.

The z position of the kaon decay vertex is taken as the weighted average of z_{ab} and z_{cd}. The (x, y) position of the decay vertex cannot be determined, since there is no directional information available for the individual γs. The energy centroid of the four photons at the CsI indicates where the kaon momentum vector was pointing when the kaon decayed. We calculate the (x, y) position of the vertex assuming that it lies on the line between the target and the energy centroid. Using this decay position, we calculate the momentum vectors for the photons and the four-photon invariant mass. We will mis-measure the transverse position with this assumption if the kaon scatters before decaying.

Figure 7.2 shows a reconstructed $K \to \pi^0\pi^0$ event. The energy centroid lies within the regenerator beam-hole at the CsI, so this event is most likely a $K_S \to \pi^0\pi^0$ decay.

The initial kaon trajectory is not known, and the energy centroid is the only information available to indicate if the kaon scattered. If the energy centroid lies outside of both of the beams, we know that the kaon scattered. As a measure of the distance of the energy centroid from the center of the beam, we define the "ring number"

$$R = 4 \times \max\left(\Delta x^2, \Delta y^2\right), \tag{7.4}$$

where Δx^2 and Δy^2 are the horizontal and vertical distances to the center of the closest beam, measured in centimeters. The ring number R is the area of the square centered on the beam, on the edge of which the energy centroid lies. The beams are nominally 9.3 cm square at the CsI, so all events with $R < 86.49$ cm^2 are within one of the beams, neglecting resolution effects. The ring number only tells us where the kaon ended up, not from where it came. For example, a kaon can scatter in the regenerator, "cross over" to the vacuum beam, and still have a ring number that

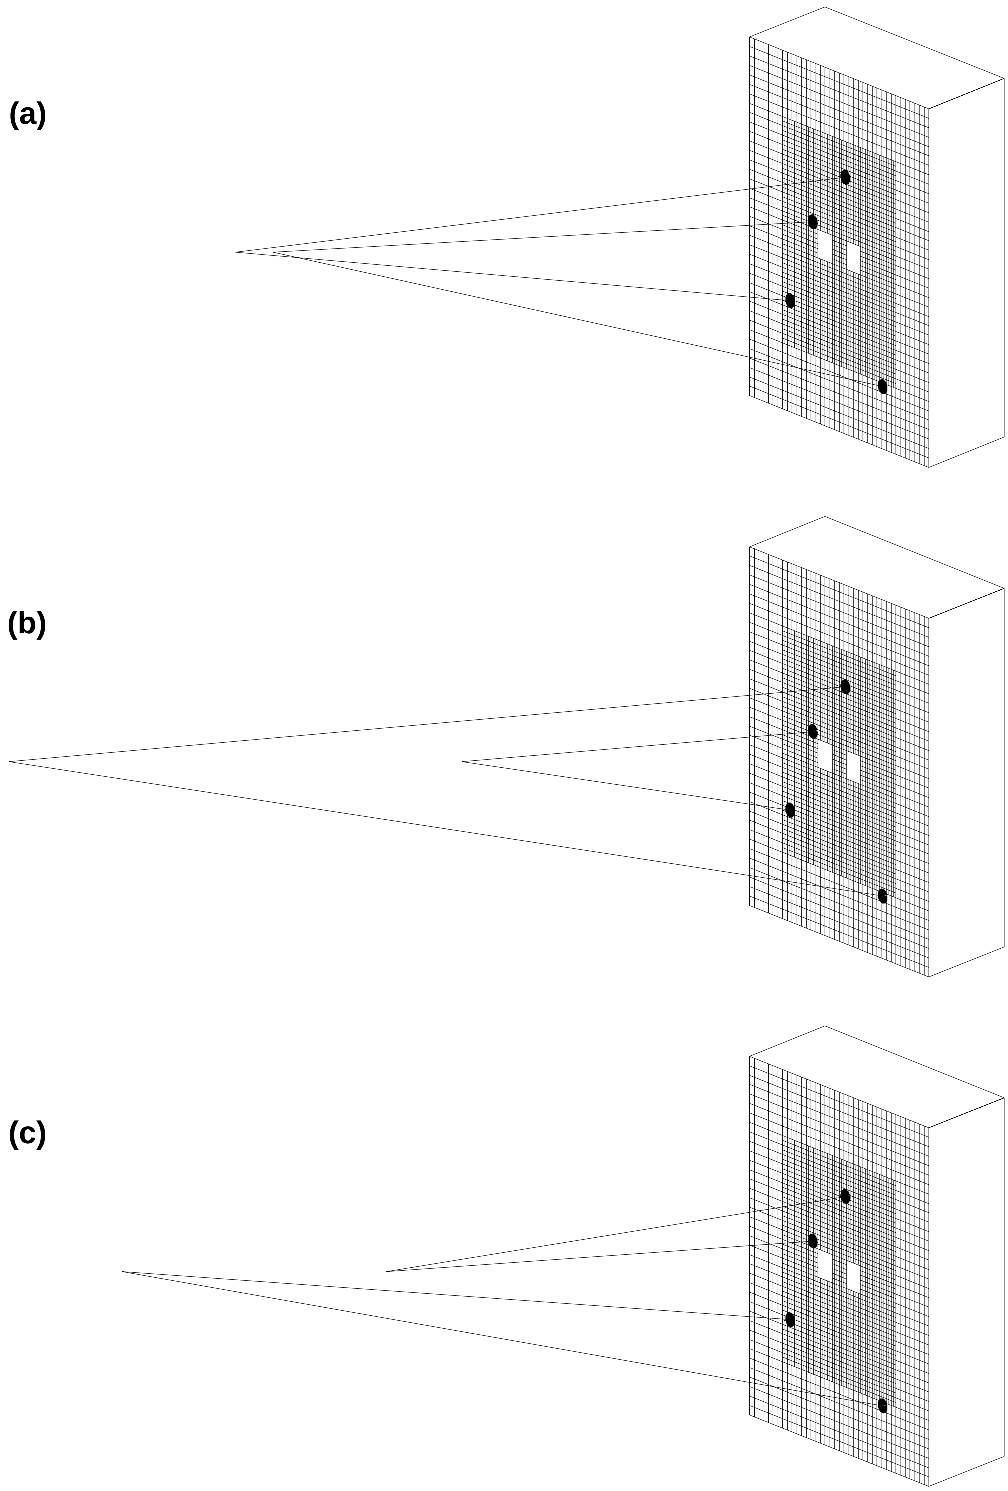

Figure 7.1: The three ways to pair four photons to make two π^0s. The pairing (a) is most consistent with a $K \to \pi^0\pi^0$ decay.

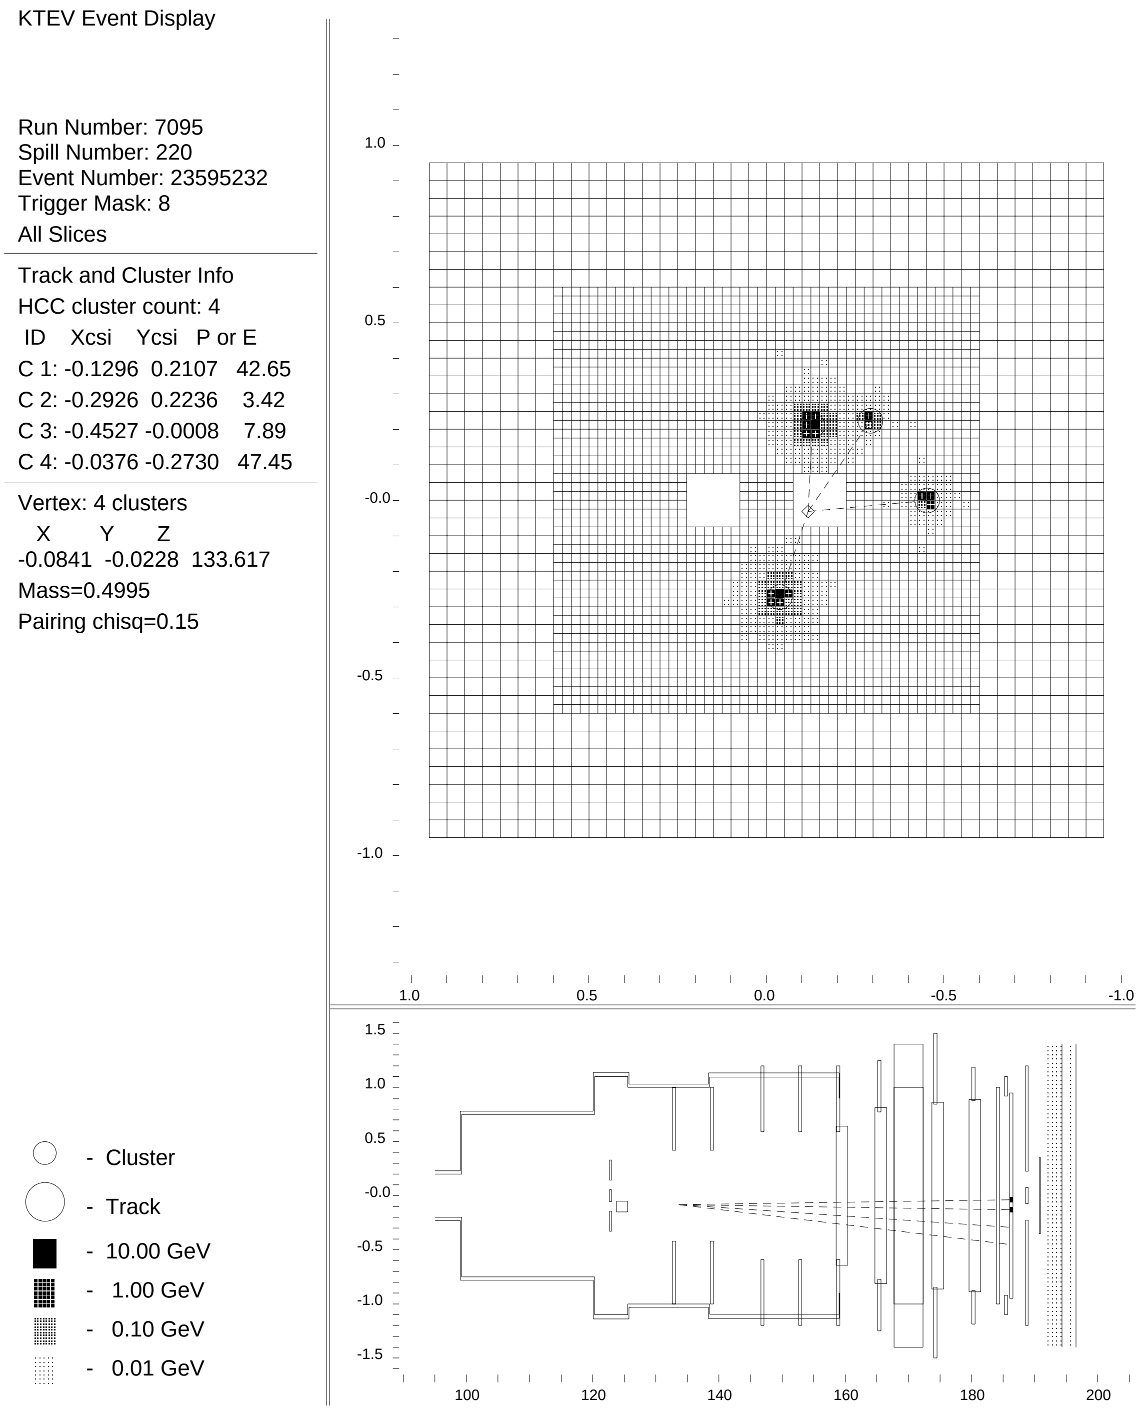

Figure 7.2: Event display for $K \to \pi^0\pi^0$ event. The thick-lined boxes in the centers of the clusters of energy show blocks for which the HCC bit is set. The dashed lines indicate the inferred photon trajectories from the reconstructed decay vertex. The diamond indicates the position of the energy centroid of the four photons.

reconstructs within the vacuum beam, which is a background with which we must be concerned.

7.4 Analysis Cuts

In the following sections, we will discuss the cuts used to identify and define the $K \to \pi^0\pi^0$ data samples. As in the $\pi\pi$ analysis, the cuts can be broken up into four loose categories; data quality, composed of run selection, trigger verification and event quality cuts; background suppression, composed of veto counter cuts and extra-particle cuts; kinematic cuts, such as energy and invariant mass cuts; and aperture cuts, such as fiducial cuts at the edge of the detector and cluster-separation cuts.

7.4.1 Data Quality and Trigger Verification

As with the charged-mode analysis, some runs were determined after-the-fact to be unusable and are removed from the data sample. We also require that the event satisfy the trigger; there must be 28 GeV of total energy in the four clusters and the seed block of each reconstructed cluster must have its HCC bit set.

7.4.2 Event Quality Cuts

The pairing χ^2 determines how consistent the two π^0s are with the hypothesis that they come from the same parent kaon. We require $\chi^2 < 12$. The pairing χ^2 distributions for both beams are shown in Figure 7.3. The response of the HCC trigger is understood and well-modeled for energies above a few GeV. We require all photons to have a minimum cluster energy of 3 GeV.

7.4.3 Veto Counters

The primary purpose of the veto counters is to suppress the $K \to 3\pi^0$ background and inelastic interactions in the regenerator by catching photons that would not hit the calorimeter. Cuts are made on energy deposited in the regenerator, ring

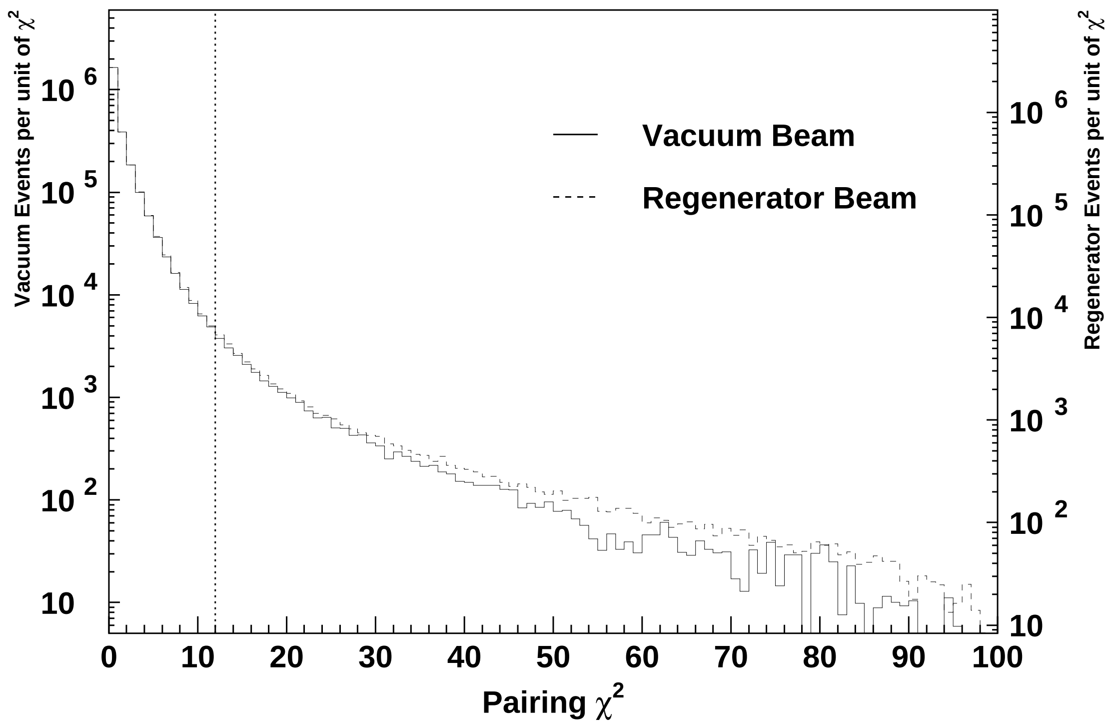

Figure 7.3: Pairing χ^2 distributions after all other analysis cuts have been applied. The dotted line shows the cut on this quantity.

counters, and the Spectrometer Antis and the Mask Anti at a level equivalent to 300 MeV. In addition, we cut an event if there is more that 5 GeV of energy deposited in the front section of the Back Anti, or if any channel of the Collar Anti has more that 1 GeV.

7.4.4 Extra-Particle Cuts

We cannot eliminate $K_L \to 3\pi^0$ decays solely with cuts on the veto counters since it is possible for all six photons to hit the calorimeter and still look like four clusters. The extra photons can be hidden by either having two photons land close enough to each other to be reconstructed as one cluster, or by having low enough energy to be below the 1 GeV threshold for HCC-seeded clusters. To attack the first case, we construct a "shape χ^2" which compares a cluster's transverse shape to that expected from electromagnetic showers. Because the Monte Carlo does not mock up this variable particularly well, we cut loosely at a value of 48. We find low energy clusters by rerunning the clustering algorithm without the HCC seed requirement, and veto events with extra clusters that have an energy of at least 0.6 GeV, as long as the extra cluster is far from other clusters and the beam-holes, has a shape χ^2 consistent with being an EM cluster, and is in-time with the other clusters. We also remove any events with reconstructed charged tracks.

7.4.5 Kinematic Cuts

The four-photon invariant-mass distributions, after all other cuts, are shown in Figure 7.4. The mean mass resolution is ~ 1.5 MeV/c^2 in both beams. The large number of events at high and low reconstructed mass in the vacuum beam are $K_L \to 3\pi^0$ events. We require 490 MeV/$c^2 < m_{\pi^0\pi^0} < 505$ MeV/c^2.

The ring number R can identify most kaons that scatter out of the beams. The data distributions are shown in Figure 7.5. We require $R < 110$ cm^2.

7.4.6 Aperture Cuts

We make aperture cuts on particle positions to cleanly define the acceptance. The two main apertures for the neutral mode analysis are the inner and outer edges of the calorimeter. To define the inner edge near the beam hole, we use the Collar Anti as a physical edge. The CA vetoes photons within 1.5 cm of the calorimeter edge. Photons which hit the CsI just outside the CA, but within the first ring of crystals, are reasonably well constructed, and no additional cut is made based on the reconstructed position.

At the outer edge of the calorimeter, we require that clusters not have their seed blocks within the outer ring of crystals, which maintains a *de facto* 5.0 cm cut from the edge of the CsI. The last veto counter, the CsI-Anti, could have been used to provide this aperture cut. However, since it is located upstream of the calorimeter and the photons have appreciable angles, this cut would not define a very sharp edge.

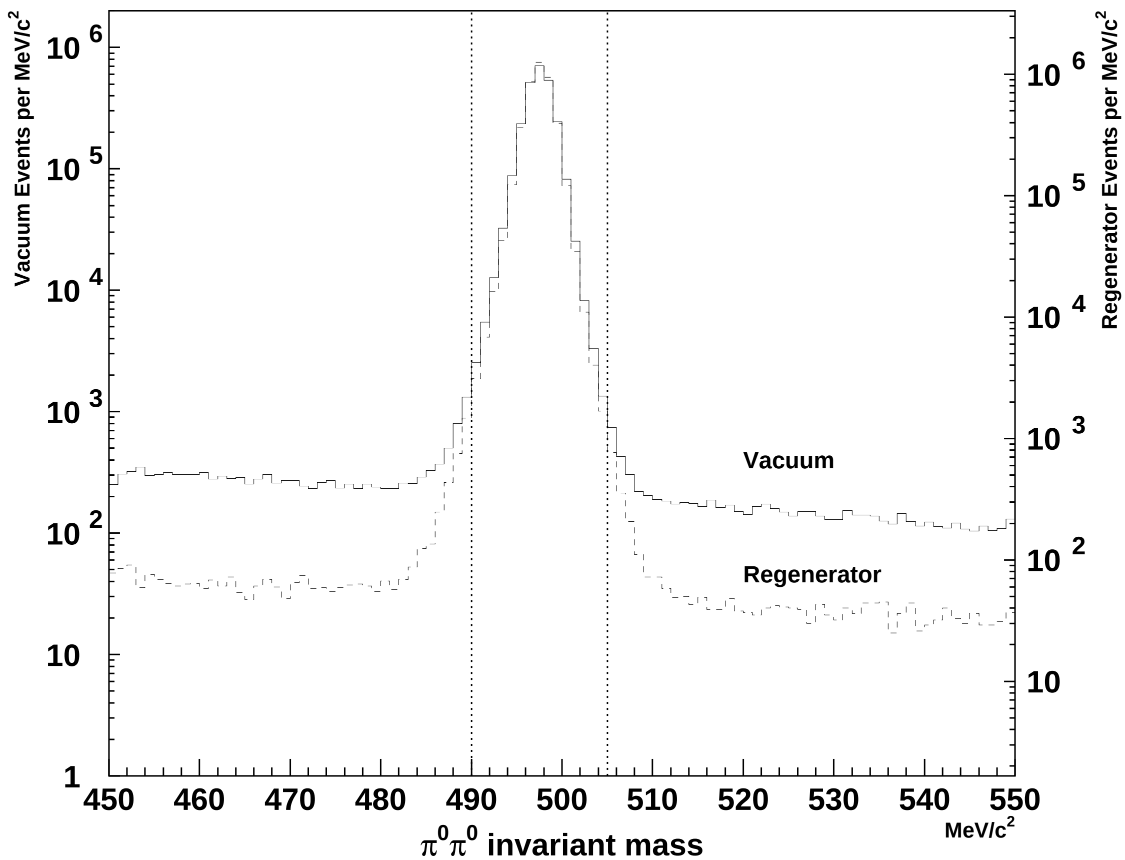

Figure 7.4: The $m_{\pi^0\pi^0}$ invariant mass distribution for the vacuum and regenerator beams, after all other analysis cuts have been applied. The dotted lines show the cuts applied to this quantity.

7.4.7 Fiducial Cuts

As with the charged-mode analysis, we define our final sample to be within the ranges 110 m $< z <$ 158 m and 40 GeV $< E_K <$ 160 GeV. The vertex z distributions for the two beams are shown in Figure 7.6. After all the cuts are applied, we have 2,501,675 events in the vacuum beam and 4,181,311 events in the regenerator beam.

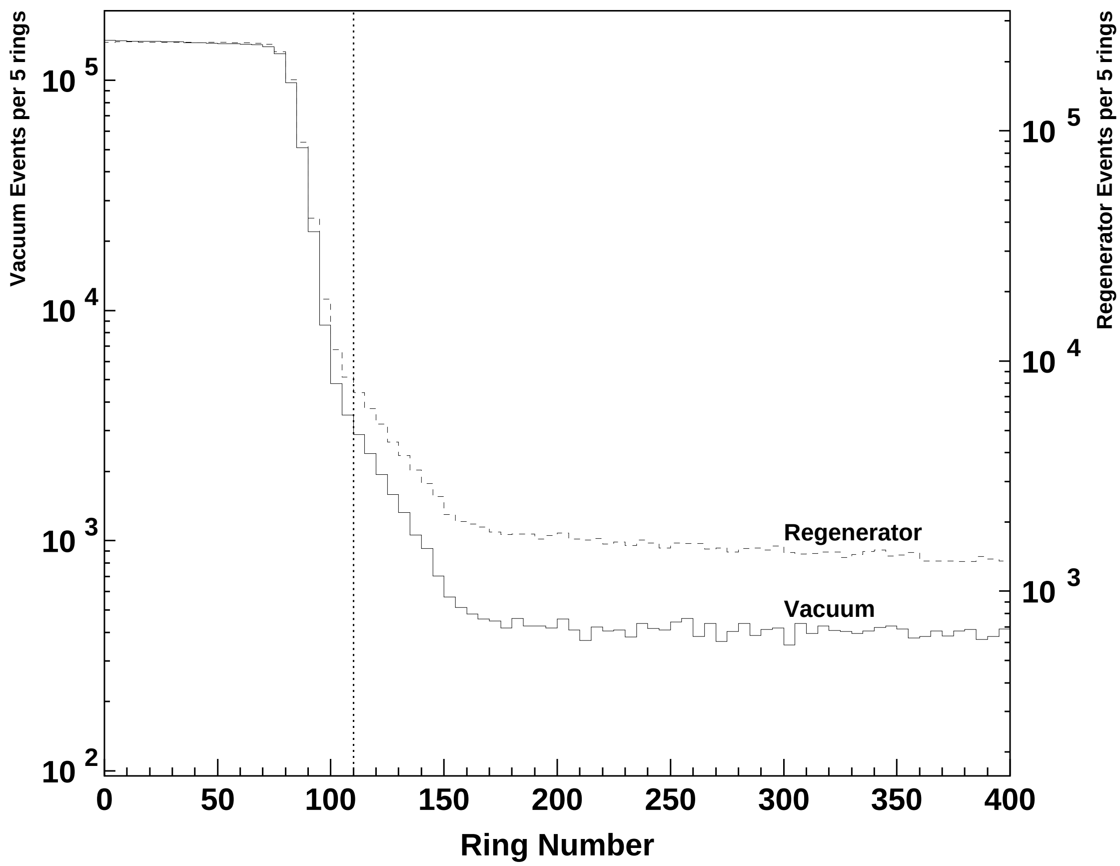

Figure 7.5: The neutral-mode ring distribution for the vacuum and regenerator beams, after all other analysis cuts have been applied. The dotted line shows the cut applied.

7.5 The Absolute Energy Scale

Reconstruction of both E_K and z in the neutral mode depends critically on setting the absolute energy scale of the calorimeter. The calorimeter is initially calibrated with electrons, and we make a final correction to move from the electron calibration to a calibration used for photons. We expect a small 0.1% difference based on the slightly different electromagnetic shower profiles in the CsI blocks, coupled to non-uniformity of the light collection in the crystal, for which we account in the calibration procedure.

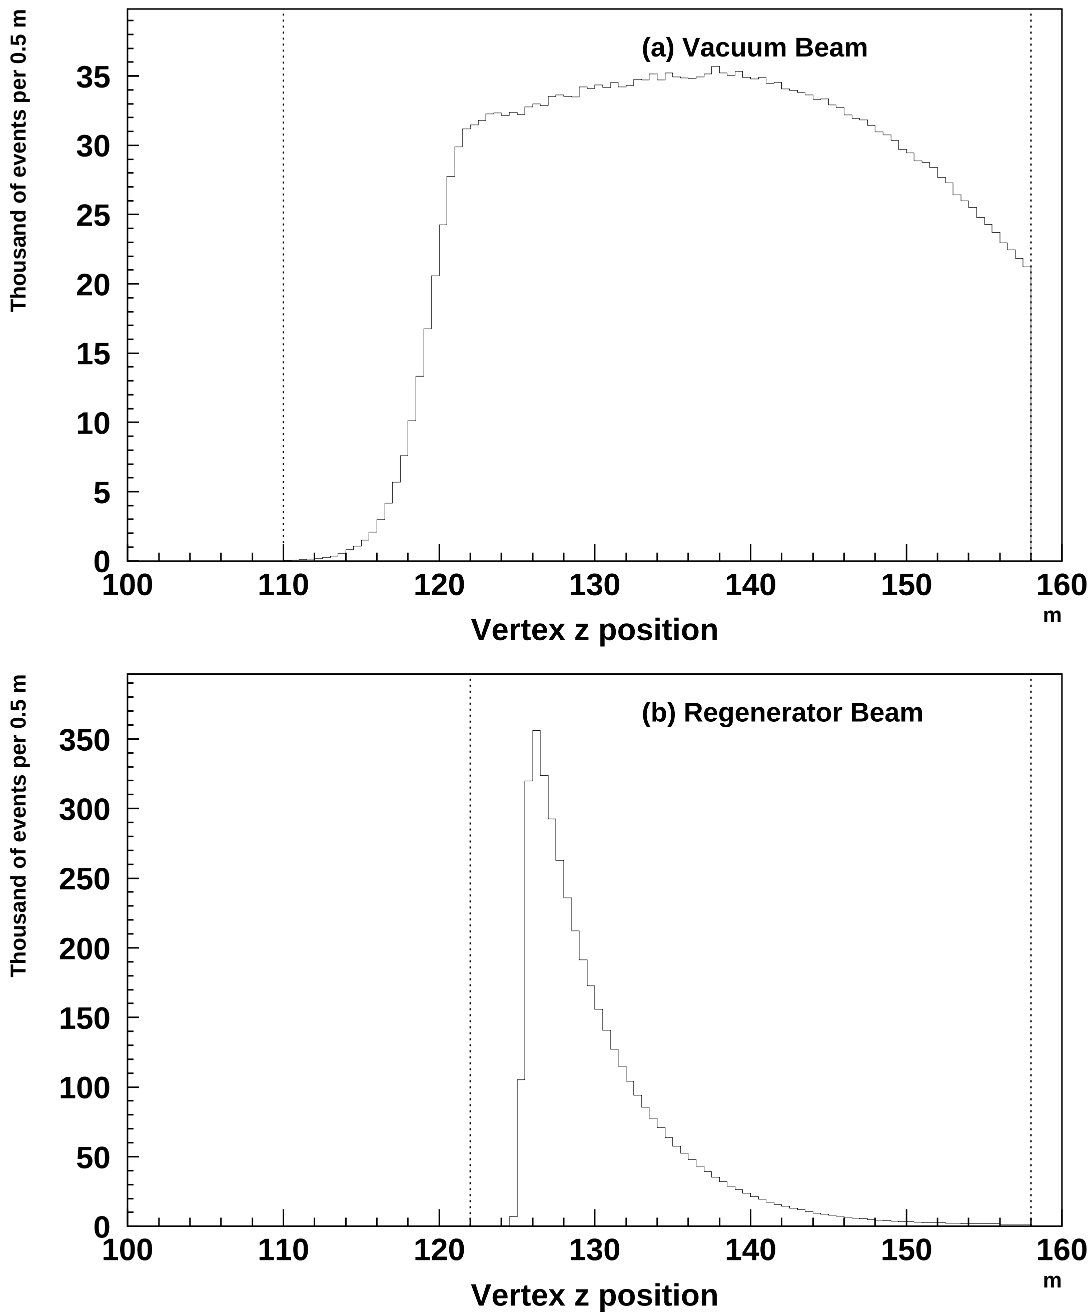

Figure 7.6: Vertex z distributions for the neutral-mode data after all analysis cuts (including the cut on the vertex z) have been applied. The dotted lines show the cut made.

We check the overall energy scale by examining decays in the regenerator beam near the downstream edge of the regenerator. As shown in Equation 7.2, the reconstructed z position of the vertex is directly related to the measured photon energies. We have a prediction of the position of the edge from the Monte Carlo, based on the surveyed position of the regenerator and the simulated response to photons of the CsI calorimeter. Any remaining energy scale problem with our data calibration will be apparent when we compare the shape at the edge of the regenerator between data and Monte Carlo. Figure 7.7a shows that there is initially a shift of ~ 7.5 cm between the data and the MC, which corresponds to an average energy scale of 0.125%. We measure the shift and apply the correction in the same 10-GeV kaon energy bins that will eventually be used to fit for $\mathrm{Re}(\epsilon'/\epsilon)$. After the correction is applied to the data, the match at the regenerator (shown in Figure 7.7b) is quite good. The neutral energy scale will be one of the larger systematics we must evaluate in Chapter 10.

7.6 Background Subtraction

There are four major categories of backgrounds in the neutral analysis: $K_L \to 3\pi^0$, kaons that scatter in the defining collimators and decay to $\pi^0\pi^0$, kaons that scatter in the regenerator and decay to $\pi^0\pi^0$ (which is a background in *both* the regenerator and the vacuum beams), and hadronic interactions in the regenerator itself. The backgrounds are estimated using Monte Carlo simulations, and normalized to the data in sideband regions of $m_{\pi^0\pi^0}$ and ring number. There are other smaller backgrounds, such as $K \to \pi^+\pi^-\pi^0$, $K_L \to \pi^0\gamma\gamma$, production of neutral particles in the material of the drift chambers, and $\Sigma^0 \to \Lambda\pi^0$ with $\Lambda \to n\pi^0$. These are small enough that we do not consider them here.

There is a final class of events that we do not consider a background. Occasionally the photons from a real $K \to \pi^0\pi^0$ event will be paired incorrectly, yielding an event with a good ring number but an incorrect reconstructed mass. We consider this part of the signal, and as stated above, the level of the "mis-pairs" is absolutely predicted by the Monte Carlo to be 0.03%.

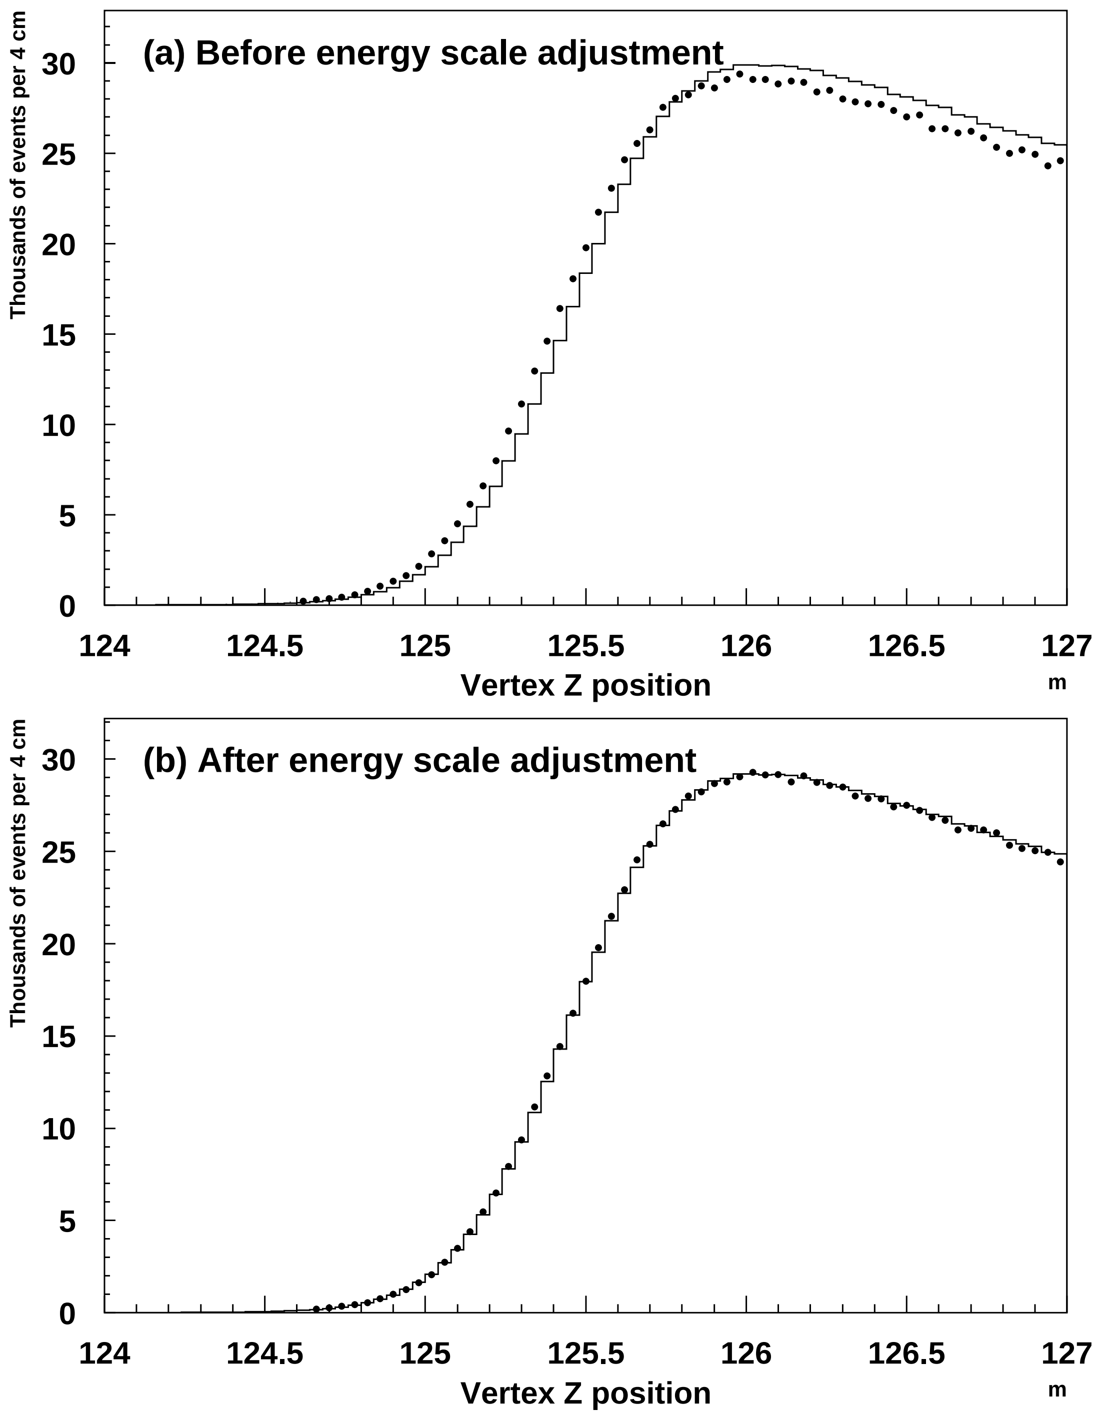

Figure 7.7: $K \rightarrow \pi^0\pi^0$ decays at the regenerator edge, used to match the energy scale. The dots are data; the histogram is Monte Carlo.

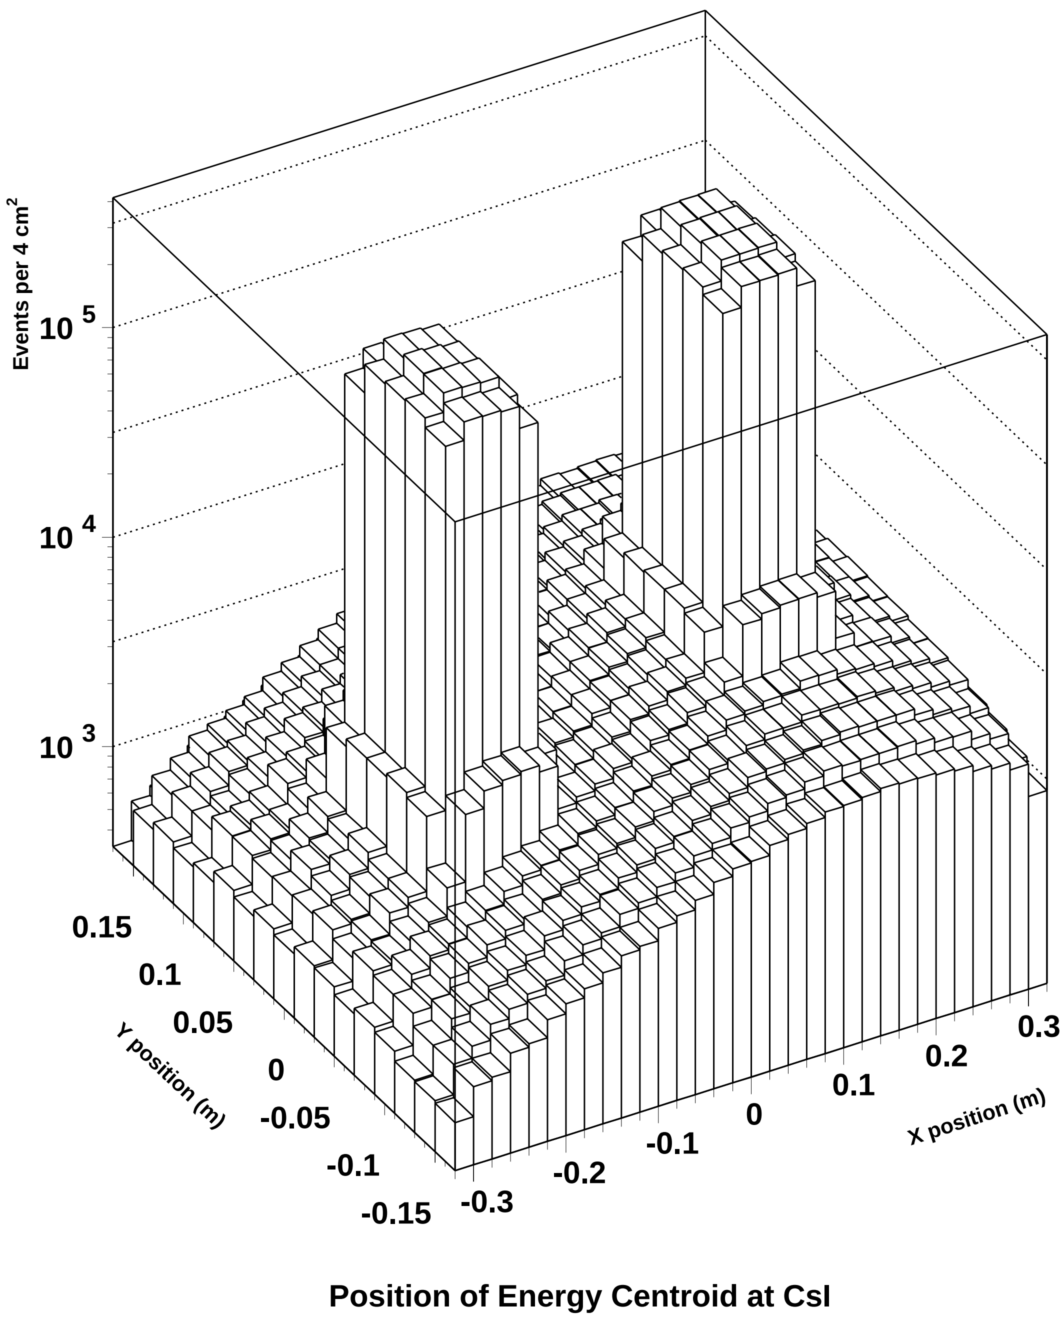

Figure 7.8: Position of the energy centroid at the CsI, showing regenerator scattering extending under the vacuum beam. This plot is made using $K \to \pi\pi$ events for which the regenerator is on the $+x$ side of the detector, but without the final offline cut on energy deposited in the regenerator. Hence the scattering background is larger in the plot than in the final data sample.

7.6.1 Normalization of the Backgrounds

The various background contributions are normalized and subtracted in a manner that decouples them as much as possible. The $3\pi^0$ background is normalized using the sidebands of the mass distribution (separately in the vacuum and regenerator beams, and after accounting for the $\pi^0\pi^0$ mis-pairing contribution to the sidebands). This is done using events with $z > 140$ m, where the $3\pi^0$ background is the largest, and there is no contribution from hadronic interactions in the regenerator. The match between the data and $3\pi^0$ background in the mass sidebands is shown in Figure 7.9.

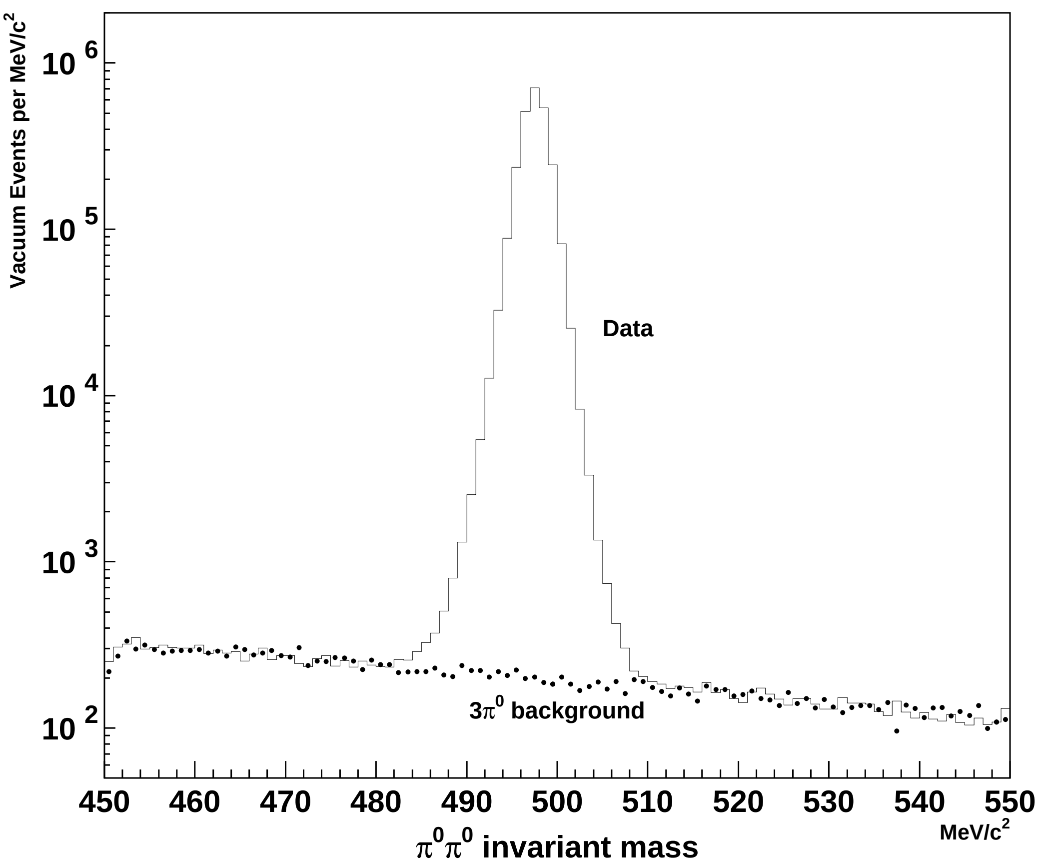

Figure 7.9: Vacuum-beam mass distribution for data and $K_L \to 3\pi^0$ background Monte Carlo after all analysis cuts have been applied.

We next subtract the hadronic "junk" produced in the regenerator. We normalize this background to the tails of the mass distribution for regenerator-beam events in the range 124 m $< z <$ 126 m, where essentially all of these background events reconstruct.

The next background subtracted is the collimator-scatter background. We cannot cleanly normalize this background to $\pi^0\pi^0$ data, since it reconstructs in a range of $m_{\pi^0\pi^0}$ and ring number similar to the regenerator scatter background. Instead, the absolute normalization is taken from the Monte Carlo, which in turn was tuned to match $K \to \pi^+\pi^-$ data.

Finally, after subtracting the collimator-scatter background, we assume the remaining ring-number tail background is due to regenerator-scattered events. We normalize the MC simulation of the scattering to the data using the ring-number range from 300 to 800. The normalization is done separately in the vacuum and regenerator beams, and the normalization factors agree fairly well.

The sum of all background contributions does a reasonable job of reproducing the data ring-number distributions, as shown in Figure 7.10. The fractional background contributions for each process are listed in Table 7.1.

7.7 Summary

After background subtraction, we have 4,130,372 $K \to \pi^0\pi^0$ events in the regenerator beam and 2,489,522 $K \to \pi^0\pi^0$ events in the vacuum beam. The events in the vacuum beam limit our statistical error for the measurement of $\mathrm{Re}(\epsilon'/\epsilon)$. To extract

Table 7.1: Background levels for the $\pi^0\pi^0$ samples.

Background process	Vacuum (%)	Regenerator (%)
$K_L \to 3\pi^0$, lost or fused photons	0.107	0.003
Kaon scattered in collimator	0.123	0.094
Kaon scattered in regenerator	0.252	1.130
Hadronic "junk" produced at regenerator	0.002	0.008
Total	0.484	1.235

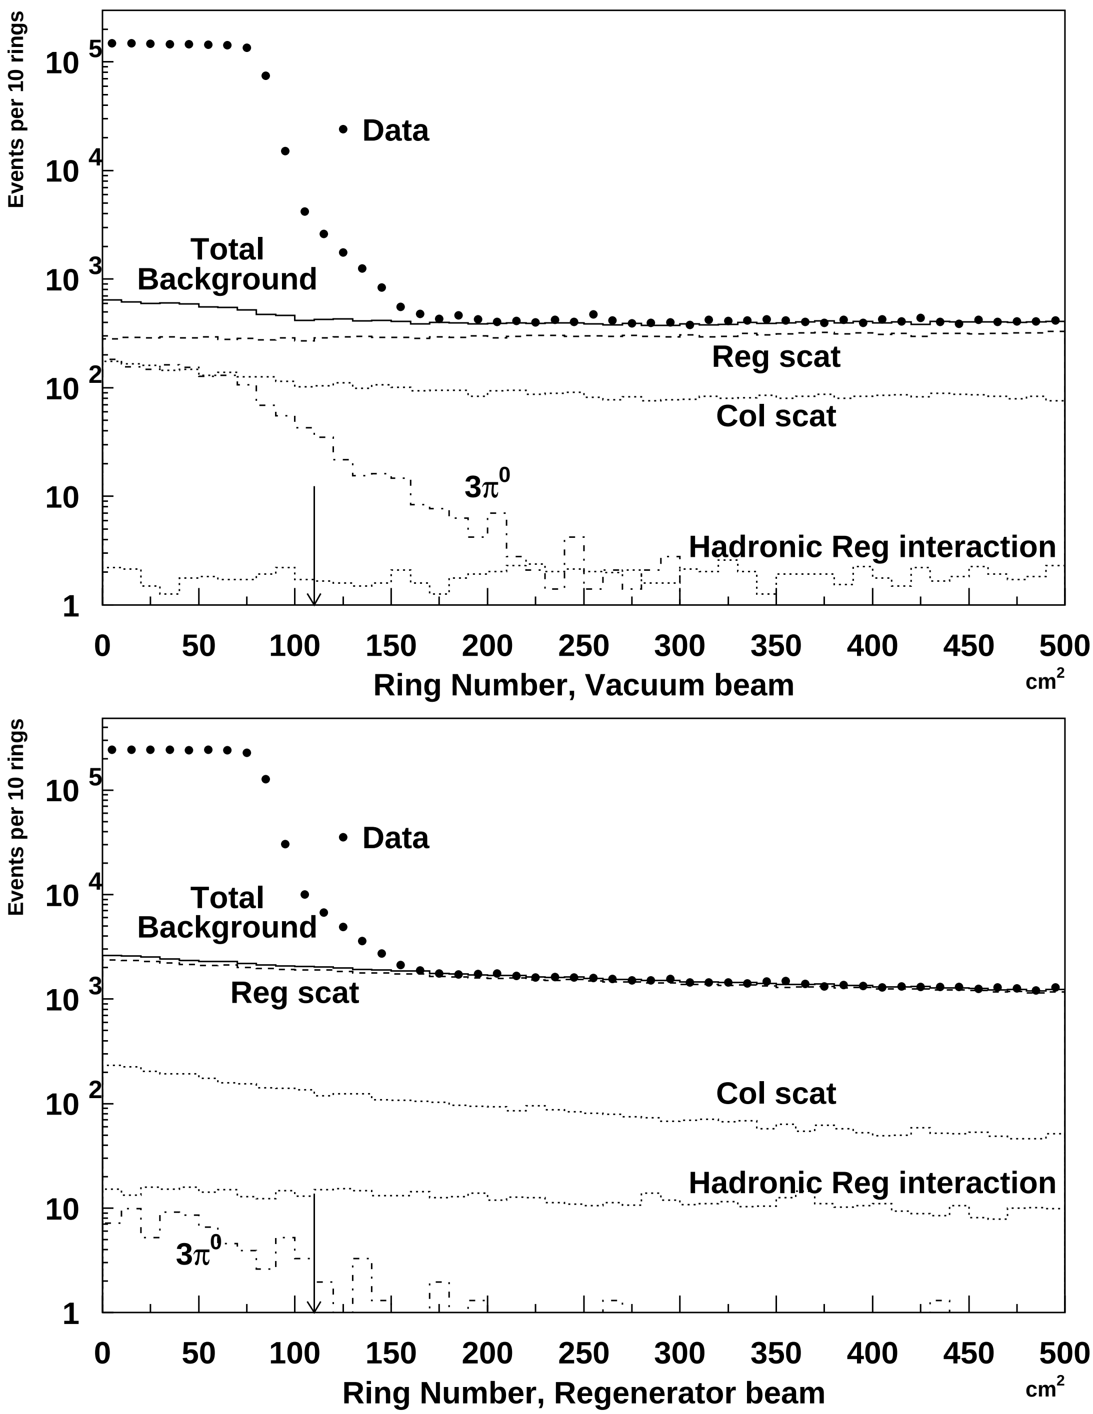

Figure 7.10: Ring number distributions for data and background contributions after all other cuts have been applied. The arrow shows the value of the cut on this quantity.

Re(ϵ'/ϵ), we must next understand the acceptance of the KTeV detector, for which we require a detailed Monte Carlo simulation. The Monte Carlo simulation is the subject of the next chapter.

CHAPTER 8
THE MONTE CARLO SIMULATION

We have written a detailed simulation of the KTeV detector, called "KTEVMC". KTEVMC simulates in detail kaon production, transport and decay, as well as the measurement of the decay products within the KTeV detector. The decays are simulated with the same energy and vertex z distributions as observed in the data, and the output of the program is events in the same format as KTeV data, with additional information detailing the true kinematics of the decay. The Monte Carlo events are analyzed as if they were data, with the same reconstruction code (with some minor exceptions detailed in Section 8.8) and analysis cuts.

As detailed in the previous chapter, we have used KTEVMC to predict the backgrounds to the signal modes. However, the main need for the Monte Carlo is to predict the acceptance of the detector; that is, given the number of reconstructed kaon decays seen in the detector, what was the number of true kaon decays? Due to geometry, the detector is limited in its coverage, and cannot reconstruct all kaon decays. Furthermore, the resolution of the detector may cause good events that decay within the detector to be lost. We require a detailed Monte Carlo to determine the level of these losses.

Considerable effort has been invested in KTEVMC to ensure that it correctly models our detector and predicts our acceptance. The initial geometry of the detector was determined from the precision survey of the KTeV hall, and further refinements made to the geometry based on the data calibrations described in Chapter 5. In fact, most of the acceptance can be predicted from simple geometry, ignoring the corrections due to resolution and efficiency. This "geometry-only" Monte Carlo is described in Section 8.3. However, the nominal KTEVMC models all the resolution and efficiency effects in considerable detail. Whenever possible, we use the high statistics modes K_{e3} and $K_L \to 3\pi^0$ to measure the response of the detector.

The Monte Carlo is also useful for furthering our understanding of the detector. We can modify the detector response in a controlled way and determine the effect on $\mathrm{Re}(\epsilon'/\epsilon)$. With knowledge of the true kinematics of the decay, stored from the Monte Carlo event generation, we can study biases in our reconstruction techniques.

The following sections will detail how KTEVMC works, from producing a kaon state, to tracing the decay products through the detector, to finally simulating the response of the detector to the decay products. At the end of this chapter, we will present measurements of the KTeV detector acceptance.

8.1 Kaon Production, Evolution and Decay

For kaon decays, the simulation code tracks the quantum state of the kaon as it propagates down the beamline, such that the decay distribution includes any interference effects.

8.1.1 Kaon production

The code first chooses the beam of origin for the kaon (left or right and vacuum or regenerator). If we are overlaying accidental data events on the Monte Carlo event (see Section 8.6), the regenerator position is fixed in the data, so we pick a regenerator or vacuum beam event and then pick the left or right beam based on the position of the regenerator in the accidental event. The kaon energy and production angle are chosen from a combined spectrum motivated by a parametrization due to Malensek [78] of K^+ and K^- production by 450 GeV protons incident on a beryllium target. Using the valence quarks in the proton as a guide, we assume the production of neutral kaons is related to charged kaon production through the following:

$$K^0 \sim \left(K^+ + K^-\right)/2 \qquad\qquad \overline{K^0} \sim K^- \tag{8.1}$$

We arrived at this model by breaking up the charged kaon production into relative probabilities σ_u, σ_d and σ_s for $u\bar{u}$, $d\bar{d}$ and $s\bar{s}$ quark pair production. We assume $\sigma_u = \sigma_d$. K^+ production is considered to be either the creation of both an $s\bar{s}$ pair and a $u\bar{u}$ pair ($\sigma_s\sigma_u$) or the creation of an $s\bar{s}$ pair and the use of one of the two valence u quarks in the proton ($2\sigma_s$). Hence K^+ production is proportional to $\sigma_s\sigma_u + 2\sigma_s$. Similarly, we find:

$$K^+ \sim \sigma_s\sigma_u + 2\sigma_s, \tag{8.2}$$

$$K^- \sim \sigma_s\sigma_u, \tag{8.3}$$

$$K^0 \sim \sigma_s + \sigma_s\sigma_u = \left(K^+ + K^-\right)/2, \tag{8.4}$$

$$\overline{K^0} \sim \sigma_s\sigma_u = K^-. \tag{8.5}$$

The energy distribution is modified by a polynomial correction such the observed energy spectrum in the Monte Carlo matches the energy spectrum seen in the data. This correction differs from unity by ±8% over the energy range of interest (40–160 GeV).

The finite neutral beam size results in a range of kaon production angles, from 4.55 mrad at the bottom of the beam to 5.05 mrad at the top. The production model predicts a variation of the flux and energy distribution over this range, but it was found that we needed to add a small linear correction of ±0.8% at the bottom (top) of the beam to the production angle distribution to get it to match the data.

8.1.2 Kaon Transport

The kaon is assigned an initial position within the production target. The transverse position is based the size of the primary proton beam and the longitudinal position is determined from an exponential profile based on the interaction probability of protons on beryllium. The kaon is propagated along the beamline, downstream toward the detector. The kaon passes through the absorbers, including the movable absorber if the kaon is in the regenerator beam. There is some probability that the kaon scatters in one of the absorbers, where it is given a transverse momentum chosen from an exponential distribution based on lead or beryllium scattering probabilities, and continues down the beamline. It should be noted that kaons that scatter through a small angle within the absorbers and pass through the defining collimator become part of our *coherent* data sample.

The kaon is propagated through the neutral beam collimators. The primary beam collimators are treated as perfectly absorbing in the Monte Carlo. A kaon that hits the defining collimator is given some probability to scatter back into the beam, according to the model described in Section 6.7.1.

8.1.3 Evolution of the Kaon Quantum State

As a kaon propagates down the beamline, its quantum state changes from the initial K^0 or $\overline{K^0}$ into a state that is almost completely K_L, except at high kaon energy.[1] KTEVMC uses K_S and K_L as its basis states, and calculates the exact transformation matrix. In vacuum, the matrix is diagonal, with terms that simply represent the decays of the two states. When the kaon traverses matter, such as the absorbers, additional diagonal terms are generated, as well as off-diagonal terms that mix K_S and K_L. These off-diagonal terms characterize regeneration in our beamline. Regeneration is considered for all material in the beam. Regeneration is not particularly important in the absorbers, but does modify the energy spectrum in the vacuum beam slightly. Clearly the most important regeneration effect to consider is in the regenerator itself. Our Monte Carlo acceptance correction is done in 2 GeV $\times$ 2 m bins (see Section 9.2.1), so our acceptance prediction from Monte Carlo needs only be accurate over these scales, which would not necessitate a detailed description of regeneration. However, for the best match between data and Monte Carlo, we attempt to model regeneration as well as possible. The bulk of the regenerator is plastic scintillator; the regeneration amplitude used in the MC is based on preliminary fits to the KTeV data, and has a power-law dependence on kaon energy. The lead pieces at the end of the regenerator, used to convert γs from $K \to \pi^0\pi^0$ decays

[1] The remnant K_S component is the only indicator of whether a K^0 or $\overline{K^0}$ was the initial kaon state.

within the regenerator, contribute significantly to the total regeneration. The behavior of the lead regeneration amplitude is based on parameters measured in an earlier experiment [79]. The resulting decay distribution shows the characteristic interference between the K_L and K_S amplitudes in the regenerator beam.

Regeneration in the MC is due to coherent forward scattering only; our MC treats all other scattering as a loss. However, we attempt to model our regenerator scatter background in the MC. The hardware veto cuts and trigger requirements remove most of the inelastic scatter background in the regenerator beam, leaving diffractive scattering events from carbon and lead as our largest background. We model the remaining scatter background as described in Section 6.7.1. In the Monte Carlo, the quantum state of the scattered kaon is changed depending on the effective regeneration amplitude of the scattering term of the model.

The kaon-nucleon cross-sections that determine the regeneration also determine the attenuation of kaons in the regenerator, so our Monte Carlo does predict some level of attenuation in the regenerator beam.

8.1.4 Kaon Decay

Each kaon is forced to decay within a specified z range downstream of the target. This range is slightly larger than the nominal fiducial region, to allow for resolution smearing. The z position of the decay is chosen based on the appropriate z distribution for the kaon state, including interference between K_L and K_S if necessary. This technique makes efficient use of CPU time.

Each type of kaon decay has a different decay generator, with varying levels of complexity. $K \to \pi^+\pi^-$ decays include the appropriate amount of $K \to \pi^+\pi^-\gamma$ due to inner-bremsstrahlung. The $\pi^\pm$ is allowed to decay via $\pi \to \mu\nu$ within the detector. Three-body decays such as K_{e3} have the appropriate form-factor within the decay matrix element. The neutral decay $K \to \pi^0\pi^0$ forces the π^0 to decay immediately, either to $\gamma\gamma$ or possibly to the Dalitz decay $e^+e^-\gamma$ with full radiative corrections, or to other rare decay modes. After the decay products are determined, they are boosted to the lab frame.

8.2 Tracing of Decay Products

All decay products, except neutrinos, are traced through the detector. All particles are traced to the end of the Back-Anti, except for muons, which are traced through the calorimeter, Pb wall, the filter steel and the muon-veto counters. All charged particles have transverse momentum imparted to them by the analysis magnet. The value of the transverse kick depends on the position of the charged particle at the center of the magnet, as determined from the zip-track measurement of the magnetic field. Photons which hit a photon veto detector and deposit enough energy can veto an event at the "tracing" stage, reducing CPU time usage. Particles that leave the detector are no longer traced; the user can veto events that have these "missing" particles to further reduce the use of CPU time. The position of all traced particles is stored at various points during tracing to be used later when simulating the response of the detector.

Electro-magnetic interactions in the material of the detector upstream of the CsI are treated in a simple manner. Electrons can undergo bremsstrahlung in the material of the detector, and photons can convert to e^+e^-. Charged particles can multiple-scatter; some effort was invested to improve this part of the simulation, and these improvements are discussed below.

8.2.1 Multiple Scattering

The current Monte Carlo simulation has been improved to use the GEANT package [80] to model the multiple scattering of charged particles interacting with the matter of the detector. These improvements involve generating a distribution of scattering angles for the different materials in the detector in GEANT, and using that distribution to pick the scattering angle. We simulate seven different materials in the KTeV detector, and use the scattering distributions within KTEVMC. For the generated GEANT scattering particle, we used 10 GeV muons; within KTEVMC, we scaled the scattering angle according to

$$\theta = \theta_0 \sqrt{p/10 \text{ GeV}}, \tag{8.6}$$

where θ_0 is the average multiple-scattering angle for a muon through a given material at 10 GeV.

We use the following seven groups of materials in the KTEVMC simulation.

- The Vacuum Window is comprised of 584 μm of Kevlar and 127 μm of Mylar.
- Helium Bag 1A, which was shown in Reference [73] to be mostly air. We simulated air and Mylar.
- DC gas and windows are simulated as a sandwich of mylar, N_2 gas, and 50–50 Ar-ethane gas.
- The helium volumes between the drift chambers.
- 4 mil gold-coated aluminum sense wires.
- 1 mil tungsten field-shaping wires.
- 5 mm of plastic scintillator.

Before the use of the GEANT simulation, the MC used a parametrization of multiple scattering which over-predicted the average scattering angle when compared to the current GEANT-based simulation. Figure 8.1a shows the comparison for scattering in He, and Figure 8.1b shows the comparison for scattering in the vacuum window. The match for He is good, while scattering is over-predicted in mylar in the previous simulation.

The improvements we have made with the treatment of the multiple scattering at the vacuum window can be seen when we compare the distance of closest approach between the two tracks in the vertex. Multiple scattering will tend to increase this distance, and decrease with the momentum of the track. We define the variable "scaled-z" ($\Delta z'$):

$$\Delta z' = \frac{z_x - z_y}{(z_{vtx} - z_P)\sqrt{\sum_i 1/\Delta M_i}}, \tag{8.7}$$

where $z_{x,y}$ are the vertex z positions measured in the x and y views, z_P is the z-position half-way between DC 1 and 2, ΔM_i is the difference of slopes of the

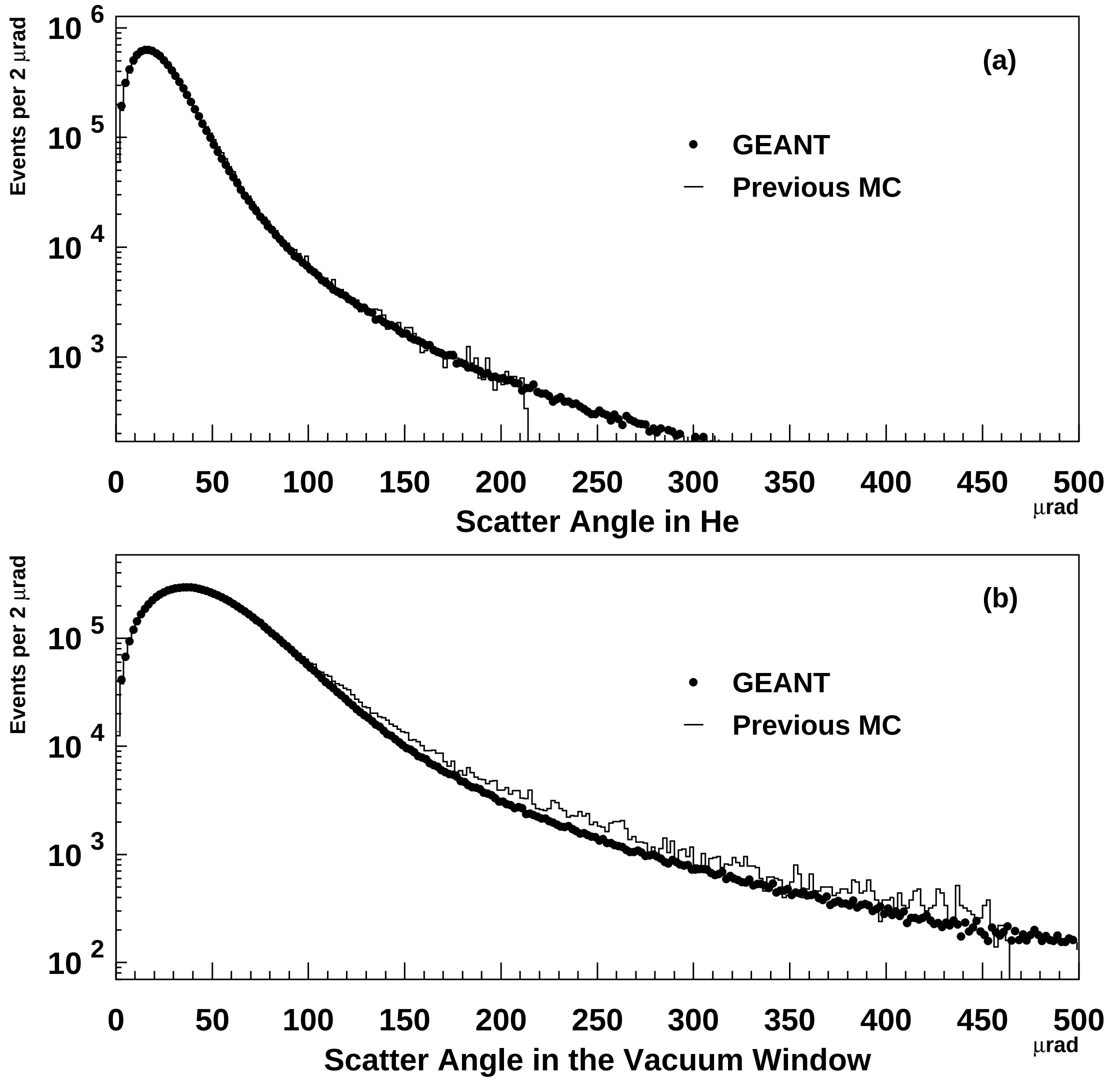

Figure 8.1: Comparisons of multiple scattering treatments. Panel (a) shows the scattering angle for 10 GeV muons in He, where both treatments give similar responses. The cut-off at 215 μrad in the previous KTEVMC treatment is artificially introduced by the generating function. Panel (b) shows the space angle for scattering through the material in the vacuum window. The cutoff at 460 μrad is due to a cutoff in the generator in the previous MC. Here the previous treatment introduced too many large scatter events.

two tracks as measured in $i = x$ or y, and p_i are the track momenta. Tracks that multiple-scatter in the vacuum window will tend to have larger $(z_x - z_y)$, and this is dependent on the momentum of the tracks. The RMS of $(\Delta z')$ has a term proportional to scattering at the vacuum window:

$$\sigma^2\left(\Delta z'\right) = A'\bar{\sigma}^2 + B'\left(p_1^{-2} + p_2^{-2}\right), \tag{8.8}$$

where A' and B' are constants, and $\bar{\sigma}^2$ is the average hit resolution of the chambers. If we plot $\sigma^2(\Delta z')$ versus $(p_1^{-2} + p_2^{-2})^{-1}$, we should see a linear dependence, where the slope is proportional to the material in the vacuum window. This is shown in Figure 8.2. Improving the scattering treatment in the MC allows us to reduce the systematic uncertainty due to data-MC mismatches in the DC simulation.

8.3 "Geometry-Only" Monte Carlo

Although we are concerned with the details of the response of the detector, such as the resolutions and efficiencies, we can demonstrate that the bulk of our acceptance correction can be determined from a Monte Carlo simulation that *does not* contain any resolution information.

We have added the functionality to KTEVMC to remove all resolution, scattering, inefficiency, and conversion details from the MC. The remaining MC simply traces the direct kaon decay products (either $\pi^{\pm}$ or $\gamma\gamma$), and rejects events solely based on the position of the decay particles within the detector. We apply our standard set of cuts as detailed in Chapters 6 and 7, using the positions and energies of these perfectly reconstructed decays.

The final goal of the Monte Carlo is to predict the acceptance of the detector. The distribution of reconstructed events in the vacuum beam gives us a good insight into the *shape* of the acceptance, since we know the shape of the underlying true number of events as a function of kaon energy from the K_L lifetime. In loose terms, if our detector had perfect acceptance, the reconstructed z distribution would have

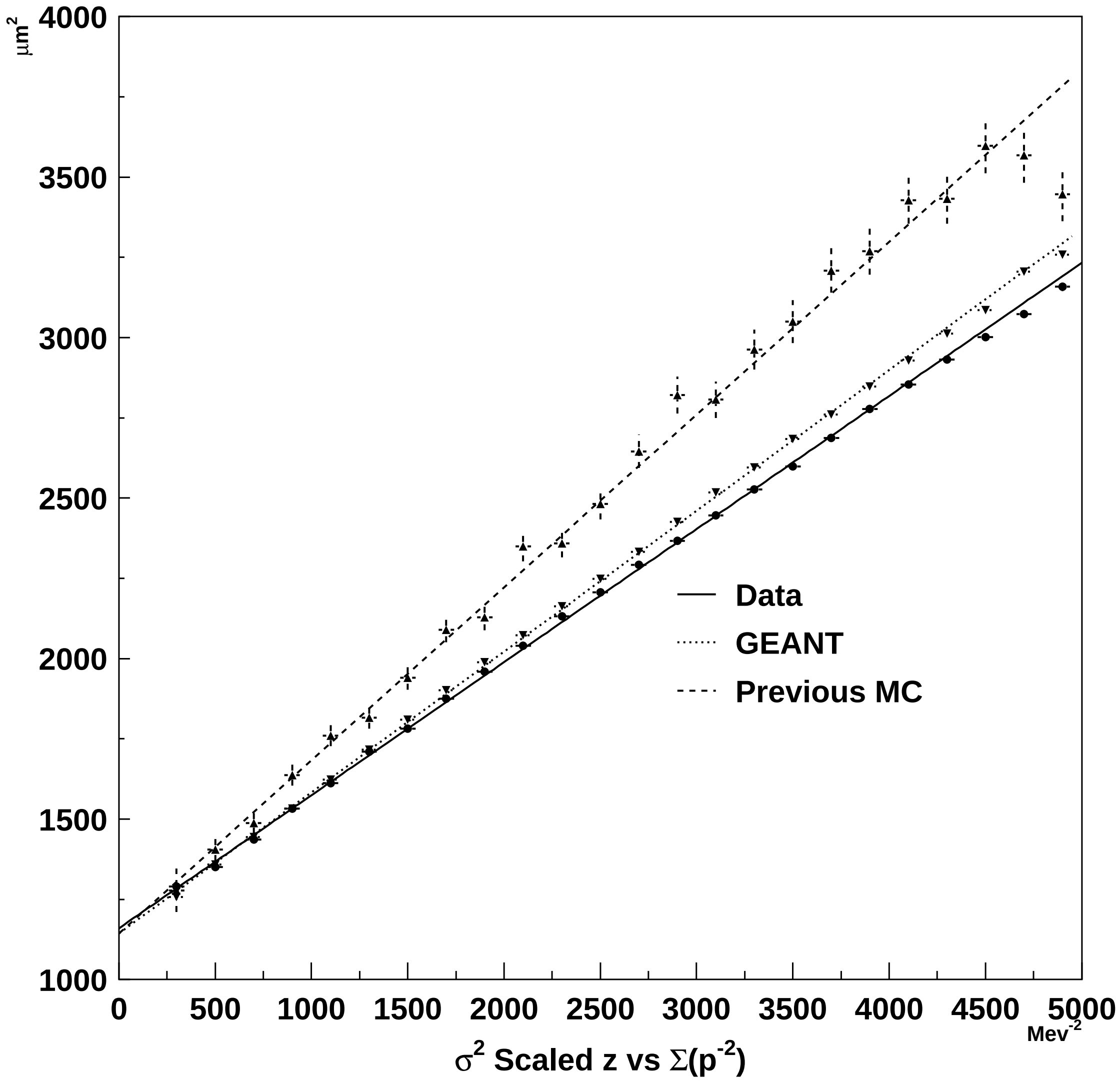

Figure 8.2: Comparison of σ^2(“scaled Δz”), defined in the text, versus $\left(p_1^{-2} + p_2^{-2}\right)$. The solid line is 1997 data, the dashed line is previous MC, and the dotted line in the GEANT-based MC. The value of the intercept is proportional to the average DC resolution, while the slope was related to the amount of material upstream of Chamber 1. The fact that the slopes between the data and previous MC do not match shows there was too much scattering in the old MC when compared to the data. The match between data and the GEANT-based MC shows the improvement.

an exponential shape with the τ_L lifetime. Clearly this is not the case, and the shape tells us the relative acceptance.

If the Monte Carlo is predicting the correct acceptance, then the shape of the reconstructed z distribution in the MC should match to the data. Any deviation in the match indicates that the MC is mis-predicting the acceptance of the detector. Furthermore, any mistake in predicting the acceptance as a function of z will bias the vacuum beam decays relative to the regenerator beam decays because of the vastly different z distributions in the two beams, leading to a bias on $\mathrm{Re}(\epsilon'/\epsilon)$.

One can show to a good approximation that an acceptance slope s between data and MC affects the measured value of $\mathrm{Re}(\epsilon'/\epsilon)$ as $s\Delta z/6$, where Δz is the difference of the mean z values for the vacuum and regenerator beams, and the factor 6 arises from converting a bias on $|\eta_{+-}|^2$ to a bias on $\mathrm{Re}(\epsilon'/\epsilon)$. $\Delta z = 5.6$ m in the charged mode.

To evaluate the accuracy of this geometry-only acceptance correction, we compare the global z-vertex distribution in our two modes to the Monte Carlo prediction, shown in Figure 8.3. From the figure, we can see that in the charged mode, we have a data-MC disagreement in the slope at the level of 10×10^{-4} m^{-1}, which would translate into an acceptance bias on $\mathrm{Re}(\epsilon'/\epsilon)$ of 10×10^{-4}. The acceptance bias is smaller in the neutral mode. We will return to the the discussion of the accuracy of this acceptance correction when we discuss the global acceptance systematic uncertainty in Section 10.9.5.

8.4 Simulation of the Drift Chambers

The drift chamber simulation works to turn a hit position as stored by the tracing routines into a digitized TDC time. In the process, we must correctly deal with the chamber resolutions, inefficiencies to true hits, the effects of overlapping hits, and the delayed-hit inefficiency, which causes the high-SOD problem.

The chamber resolutions were measured in the calibration stage of the analysis. A small component, 15 μm in quadrature, is subtracted from the measured tracking

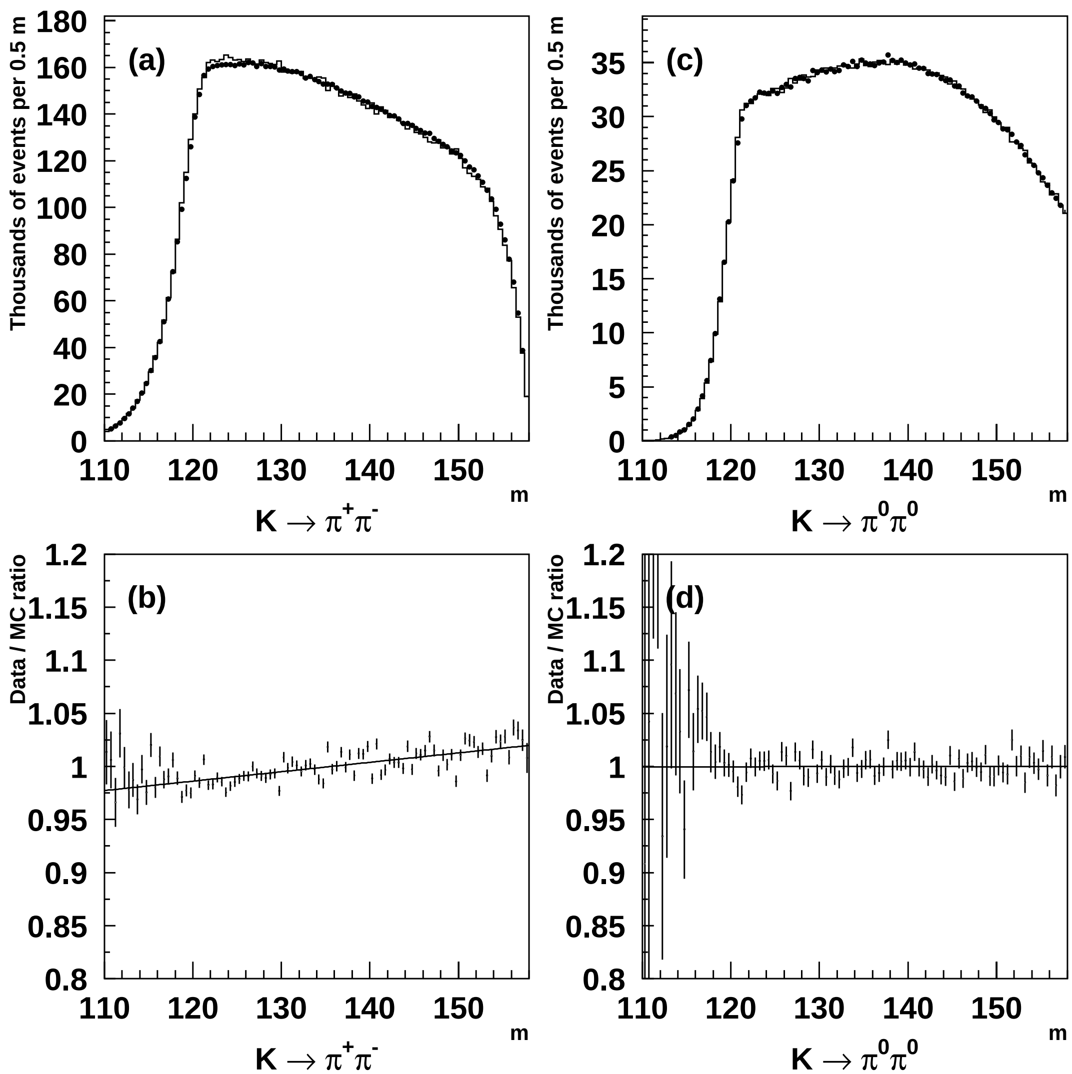

Figure 8.3: Comparison of the z-vertex distributions for data and the "geometry-only" MC, for $\pi^+\pi^-$ decays ((a) and (b)) and $\pi^0\pi^0$ decays ((c) and (d)). Panels (a) and (c) show the overlay, while panels (b) and (d) show the data/MC ratio.

resolution to account for jitter in trigger that is simulated at the trigger stage. The resolutions were used to smear the track position within the cell.

The average single hit efficiency for a given wire was high, in excess of 99%. However, it was determined that on some wires, the inefficiency was localized to specific spots on the wire, called "freckles." Wires with freckles were later analyzed with electron microscopy, and the freckles were determined to be silicon compounds. To model the spatial dependence of the inefficiency, we measured the inefficiency across the face of each chamber to create a 2D "map," shown in Figure 8.4. Furthermore, we measured the profile of the hit inefficiency as a function of the position of the track within the cell. Figure 8.5 shows the relative inefficiency as a function of distance from the sense wire for a typical wire. We have found that the DC system is more inefficient when the track passes far from the wire.

Our DC simulation would not be complete without a simulation of the high-SOD effect. The previous version of our MC parametrized the high-SOD effect in terms

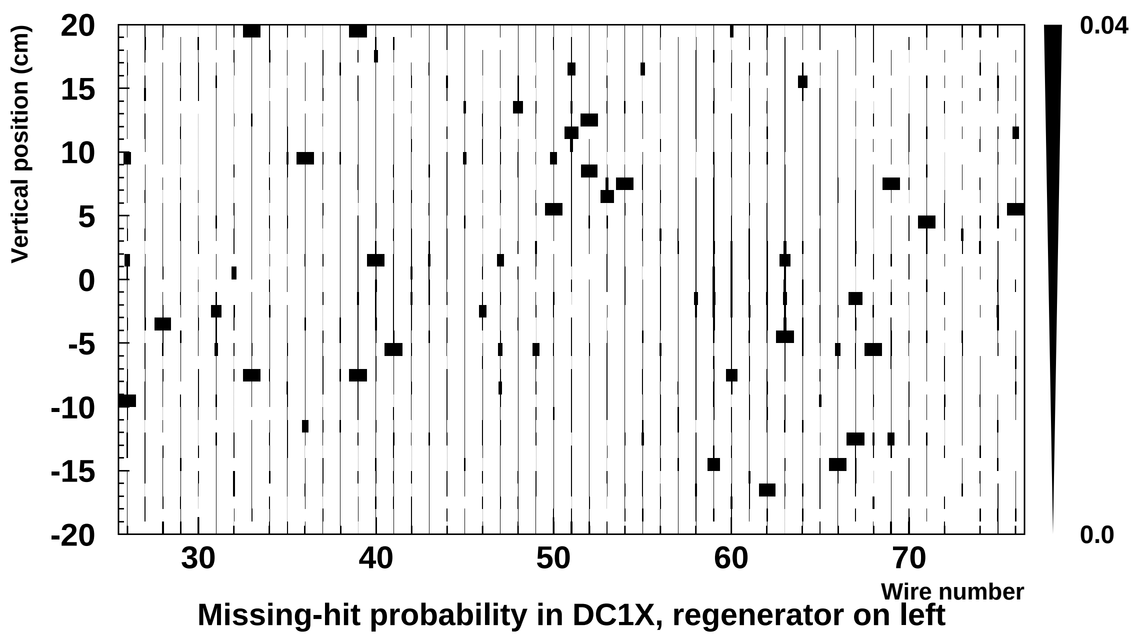

Figure 8.4: Map of missing-hit inefficiency for chamber 1 x view. The inefficiency is low across the face of the chamber except in isolated regions called "freckles."

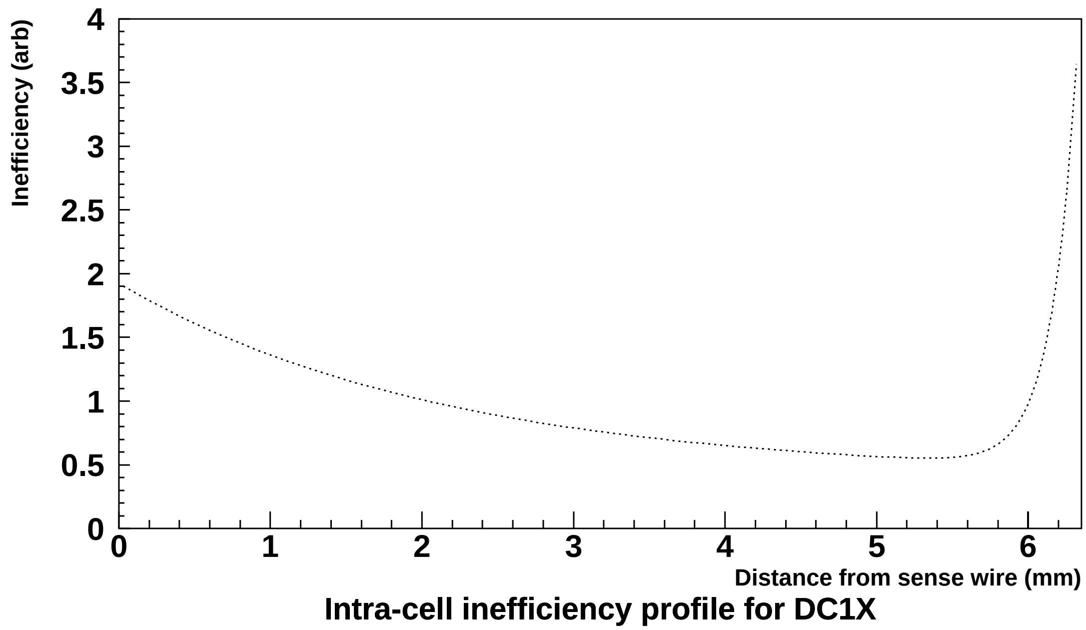

Figure 8.5: Profile of missing-hit inefficiency as a function of distance from the sense wire.

of the rate, the distance from the wire, and the amount by which the SOD was high, all measured from the data. Although this model could accurately reproduce the data in some respects, it was not tied directly to the inefficiency that we believe causes high-SODs. We have written a new simulation that uses pulse information and a variable threshold to incorporate high-SODs into our Monte Carlo simulation. This work was undertaken to try to incorporate the the high-SOD effect at a basic level so that any correlated effects between rate, distance from the wire and SOD value would be modeled correctly.

8.4.1 The High-SOD Model

High-SODs are caused by an inefficiency to the first drift electron from a track. To model this effect in our MC, we must have a measure of the threshold to pulses, a measure of the pulse shape, and a measure of the position of the primary ionization sites along the track.

In our model, primary ionization sites are distributed along the track, with the separation of each site from its neighbor picked from an exponential distribution, since ionization is a Poisson process [81]. The mean separation of ion sites for argon is $\lambda = 340\ \mu$m [81]. In addition, the number of primary ionizations per unit length is roughly linear with average atomic number, and ethane and argon have approximately the same atomic number, so an inter-ion spacing of $\lambda = 340\ \mu$m is used for the KTeV gas mixture. The inter-ion spacing also depends on pressure and temperature, but this is ignored in the simulation.

The Ideal Pulse Shape

To accurately sum many drift electrons together into a reasonable composite pulse, one needs to know the pulse shape for a single drift electron. Unfortunately, it is difficult to observe such a single pulse. Attempts have been made to look at pulses from ^{55}Fe sources [82], as well as a `SPICE` [83] simulation of the pulse shape [84]. Based on the above, an estimate for a reasonable pulse shape is parametrized as the convolution of a gaussian with $1/(1 + (t/t_0)^{3/2})$:

$$p(t) = \frac{1}{1 + \left(\frac{t}{t_0}\right)^{\frac{3}{2}}} * \exp\left(-\frac{t^2}{2\sigma^2}\right) \quad . \tag{8.9}$$

The pulse shape is shown in Figure 8.6.

There is some freedom in the choice of the pulse shape, since a change in the pulse shape can be accommodated by a compensating change in the threshold function. However, a minimum amount of 'tail' in the pulse (described by the $(1+(t/t_0)^{-3/2})^{-1}$ term) is required to allow the pulse to go above threshold. This parametrization was chosen to match as closely as possible the ^{55}Fe and `SPICE` information.

Primary Ionization

Twenty-six ions are distributed along the track segment passing a sense wire. This number was determined by increasing the number of ions until the slope of the high-SOD tail did not change (for a given threshold function). After the primary

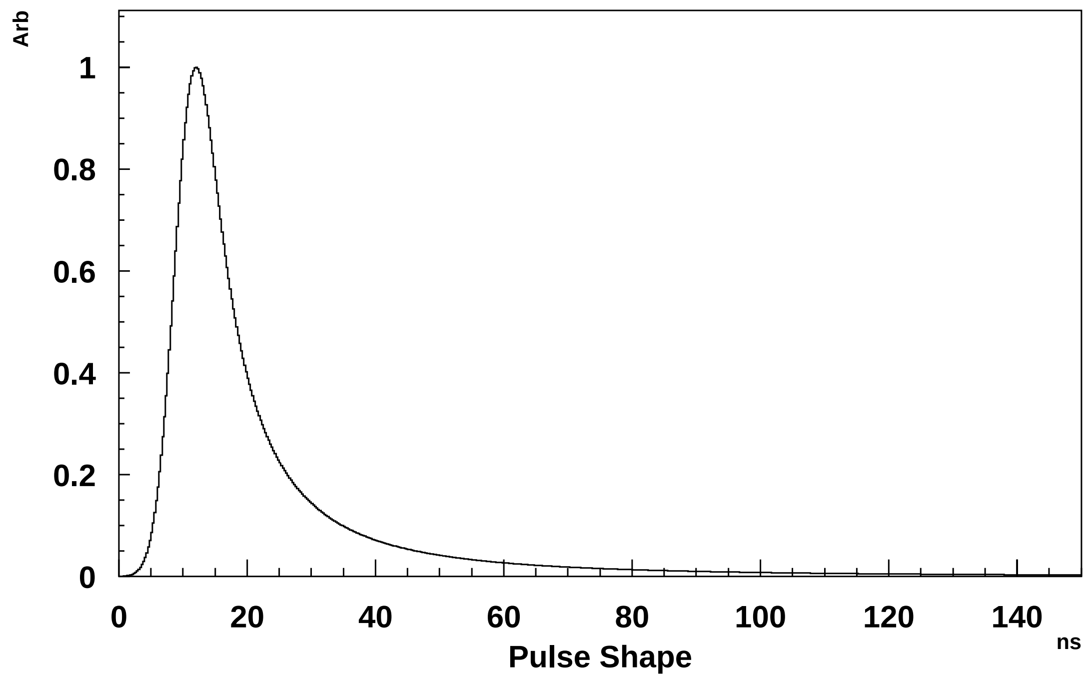

Figure 8.6: The ideal pulse shape used in the drift chamber simulation.

ionization sites are distributed along the track, the drift distance is determined from geometry, and the drift electrons are sorted by drift distance. The drift distances are then converted into drift times by using the $t(x)$ maps which are made by inverting the standard $x(t)$ maps.

With the drift times in hand, it is possible to create the composite pulse reaching the discriminator. The single drift electron pulse is offset by the drift time, and added to the composite pulse. This is done for all drift electrons (although short-cuts are taken in the code to improve the speed). For the most part, only a few drift electrons are necessary to go above threshold. However, for some pathological cases, it may take many drift electrons to go above threshold; this reproduces the high-SOD problem. Pulse shapes for the two different cases are shown in Figure 8.7. In the first panel, most of the drift electrons are separated by small drift times, and the pulse appears to be a scaled-up version of the single electron pulse. In the second panel, there are late-arriving drift electrons, causing the secondary peaks. If the threshold is high enough, the drift time may be significantly delayed from the time of the first drift electron.

Threshold Function

Within the MC, the threshold curve is used to set pulse threshold as a function of the observed high-SOD rate, which in turn comes from the 2D high-SOD maps measured from the data. It was initially determined from a toy MC that simulated one cell, with perpendicular tracks, and no additional track defects. The threshold was varied in 0.25 drift electron steps, and the high-SOD rate was determined. This function was then inverted to determine the threshold as a function of high-SOD rate.

The initial "toy MC-based" threshold function and the final function are shown in Figure 8.8. The main feature is that the threshold is greater than one drift electron in all cases, which is necessary to create the high-SOD problem. The curves are qualitatively similar; the main difference is a constant offset of 0.75 drift electrons in the threshold between the toy MC and the full simulation. This may not be too surprising, as accidental effects (included in the full simulation) can cause high-SODs not seen in the toy MC.

The curves in Figure 8.8 are parametrized by the following function:

$$T(h) = A\sqrt{h} + B \cdot h + C \cdot h^2 + D, \tag{8.10}$$

where T is the threshold in drift electrons, and h is the high-SOD rate.

Additional Corrections

There are extra corrections used to make the SODs in the MC match the data. The first correction is to determine the average time offset caused by the threshold. In the calibration, the average delay of the first drift electron is removed, and this delay is a function of the average high-SOD rate in the chamber. Hence the drift delay that causes the high-SOD effect can be seen as a difference between the local

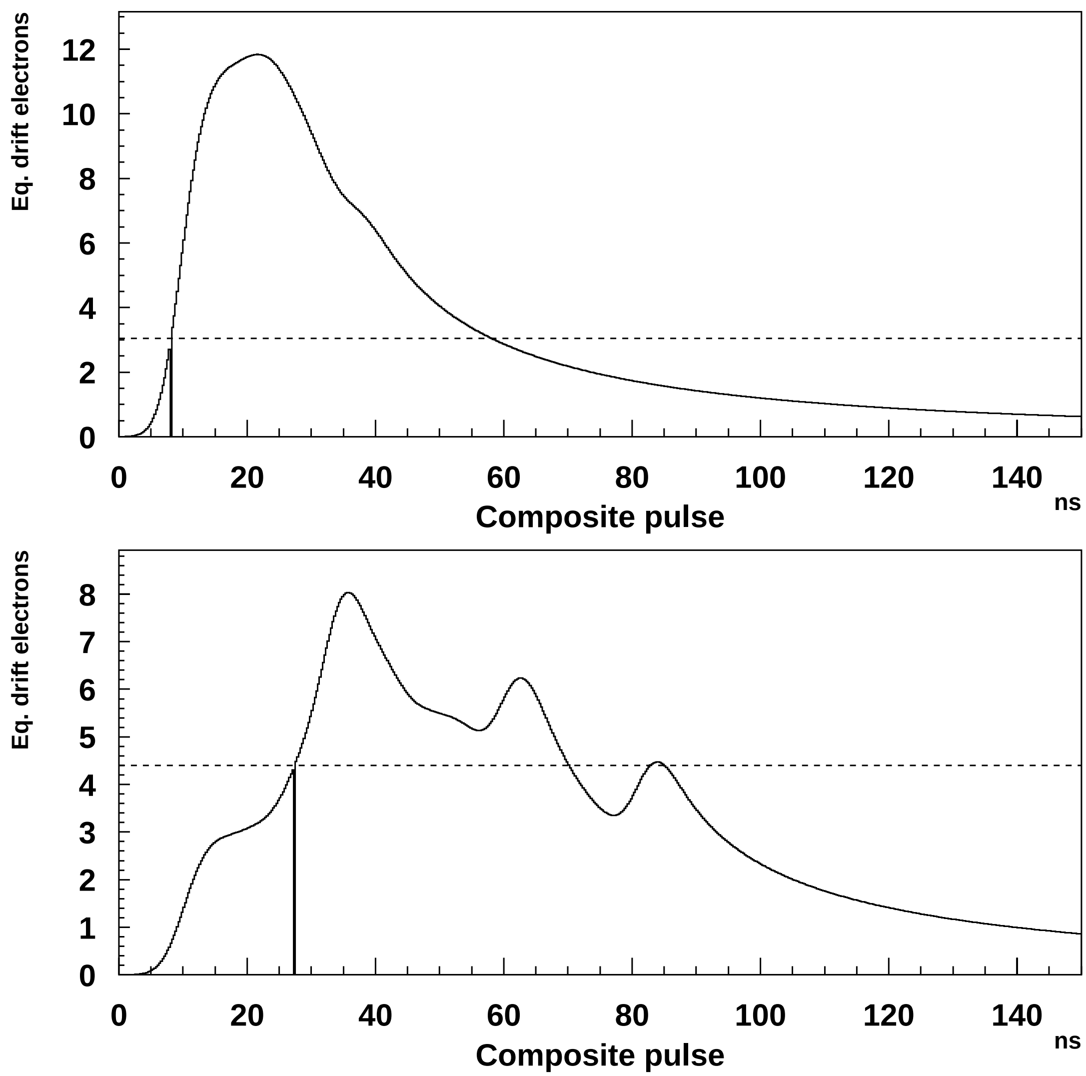

Figure 8.7: Composite pulses for two cases. The top panel shows the sum for a track far from the wire where the drift electrons are close together in drift times. The second panel shows the case when the individual drift electrons are separated in time, either due to discrete ionization or geometric effects, or more likely both. The vertical line shows the point at which the pulse went above the local threshold. In the first panel, this is near the first drift electron. In the second panel, the pulse crosses the threshold roughly 12 ns or 600 μm after the first drift electron, leading to a high-SOD.

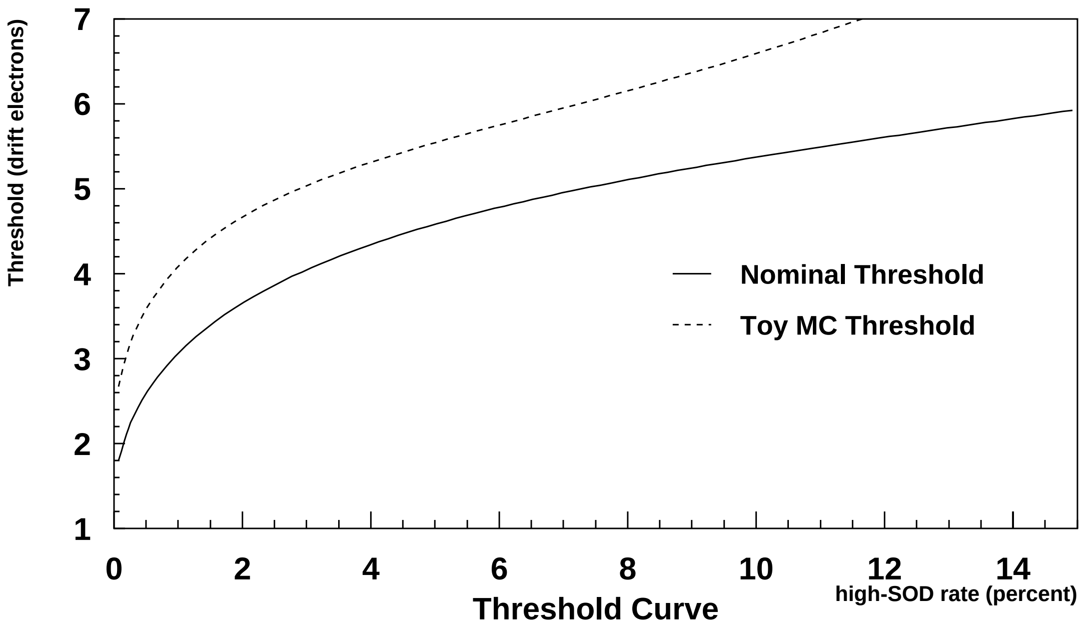

Figure 8.8: The threshold curve as a function of the observed high-SOD rate. The solid curve is used in the simulation; the dashed curve was determined from a toy MC. The main difference below 0.1 is the constant offset.

threshold and the average threshold across the chamber. This correction maps out this average threshold. This effect was parametrized as a linear effect in the average plane high-SOD rate. The high-SOD rate is $\sim 1\%$ on average across DC 1, and is ten times smaller in DC 4. This threshold correction corresponds to a shift in the mean SOD of ~ 70 μm between these two cases (*i. e.*, if we did not account for this effect, the SODs in Chamber 1 would be shifted by 70 μm while the those in DC 4 would reconstruct correctly).

Even with the above correction, the value of the mean SOD varies across the cell much more in the MC than seen in the data. An additional polynomial correction is made to the average threshold as a function of position within the cell, to flatten this walk. This correction to the mean of the SOD is 200 μm at the sense wire and falls to no correction half-way between sense wires. The need for this correction is not well understood.

Resolution Issues

Ideally, a full pulse-based Monte Carlo model should mock up the resolutions. In general, the resolution is due to the discrete ionization, diffusion, and the interplay between the threshold and the gas gain. The current model includes the ionization and some part of the gain (through the threshold). However, without an additional correction, the model underestimates the resolutions.

To correctly model the resolutions, we use the resolutions as measured in the five regions of the drift chambers, and subtract off a correction term due to the pulse model. The correction corresponds to subtracting 80 μm in quadrature from an average resolution of $\sim$ 110 μm.

Results and Data-MC Comparisons

In general, the new pulse model reasonably matches the data. As shown in Figure 8.9, the Monte Carlo correctly models the both the level and the shape of the high-SOD tail. Although there was some tuning of the threshold curve to set the level correctly, the shape of the high-SOD tail in the MC comes from the model.

As a cross-check, one can look at the distribution of high-SODs within the cell. As stated above, we see more high-SODs near the wire in the data, and this model should simulate this feature. As shown in Figure 8.10, the data and the MC match resonably, although the Monte Carlo shape does not exactly predict the data shape. Previously, the intra-cell behavior was another free parameter measured from the data. Here it is predicted after fixing the high-SOD level.

To get the best possible match to the data, it was necessary to recalibrate the MC $x(t)$s. After this was done, the data-MC comparisons in variables that depend on the resolution look very good.

The final cross-check for the high-SOD model is the comparison of p_T^2 between the data and the Monte Carlo, as shown in Figure 8.11. The presence of the high-SOD simulation clearly improves the data-MC agreement, although there is still some residual difference in the tail of the distribution above our cut. This variation will set our systematic uncertainty on p_T^2 (see Section 10.7).

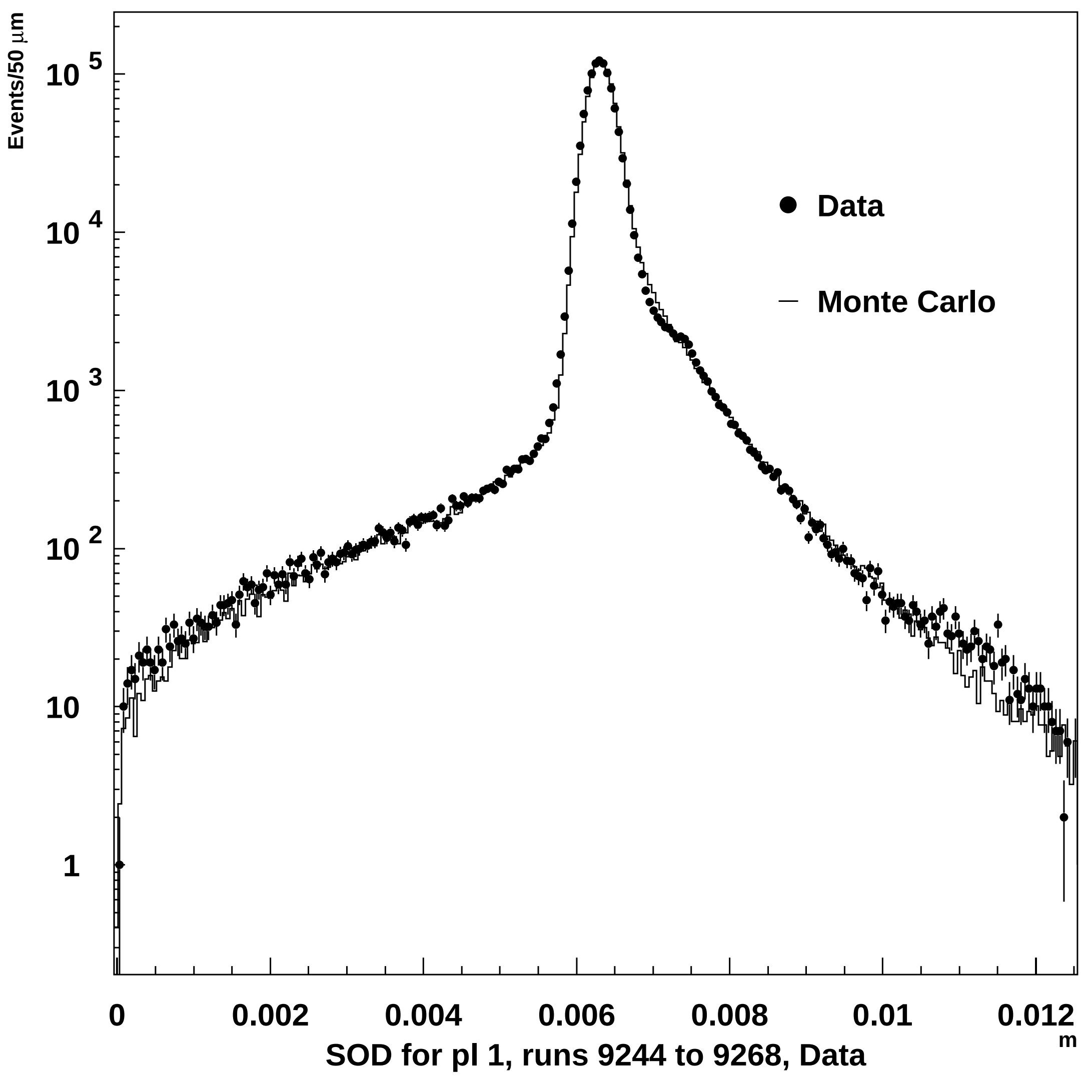

Figure 8.9: Overlay of Data (dots) on MC (histogram) for the SOD distribution for the chamber $1x$ plane-pair for a given set of runs. The MC predicts both the shape and level of the high-SOD problem.

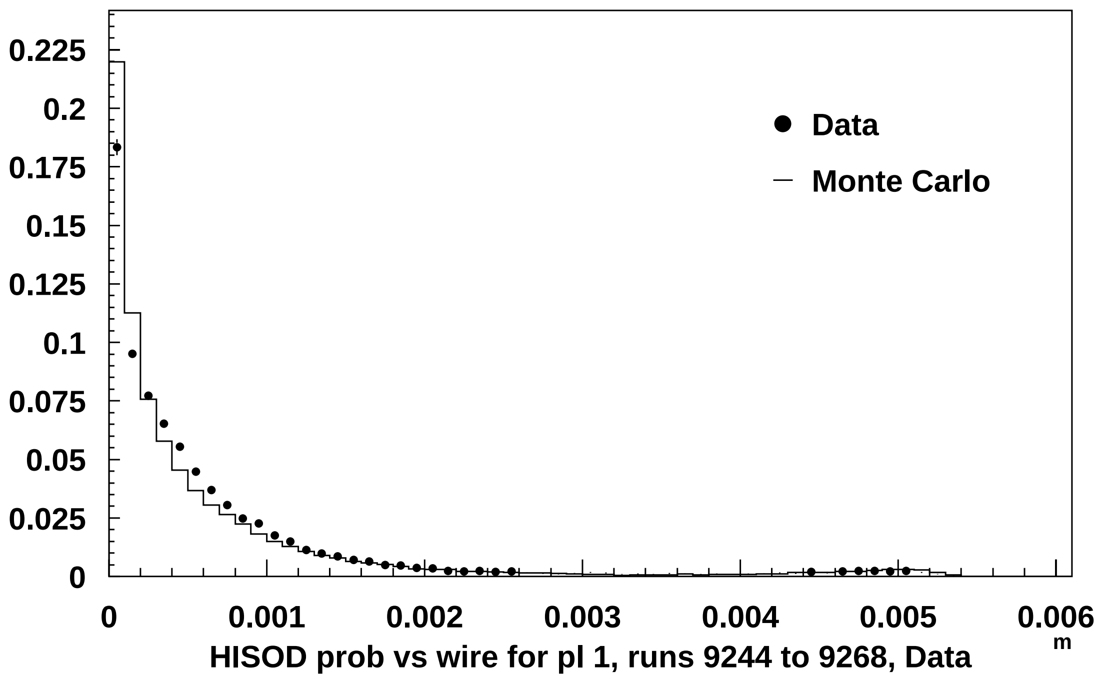

Figure 8.10: Overlay of Data (dots) on MC (histogram) for the High-SOD rate as a function of distance from the wire.

8.4.2 Early Hit Inefficiency

The final consideration for the DC modeling is the effect of earlier hits obscuring later hits. Even though we only use the first hit on a wire in the tracking code, we are concerned with early accidental activity in the drift chambers causing the loss of hits associated with real tracks.

The early hit inefficiency arises from two causes. First, the digital discriminator was designed to have a deadtime of about 50 ns, and the wire is 100% inefficient for this period of time. The second cause of inefficiency is long analog pulses that stay above the discriminator threshold for more than 50 ns. The shape of the analog pulse was measured from the data. Initially, it appeared that the distribution of pulse widths could be described by an exponential distribution. The combination of the discriminator deadtime and the exponential distribution is shown in Figure 8.12. With more detailed study, it became apparent that there is an additional component

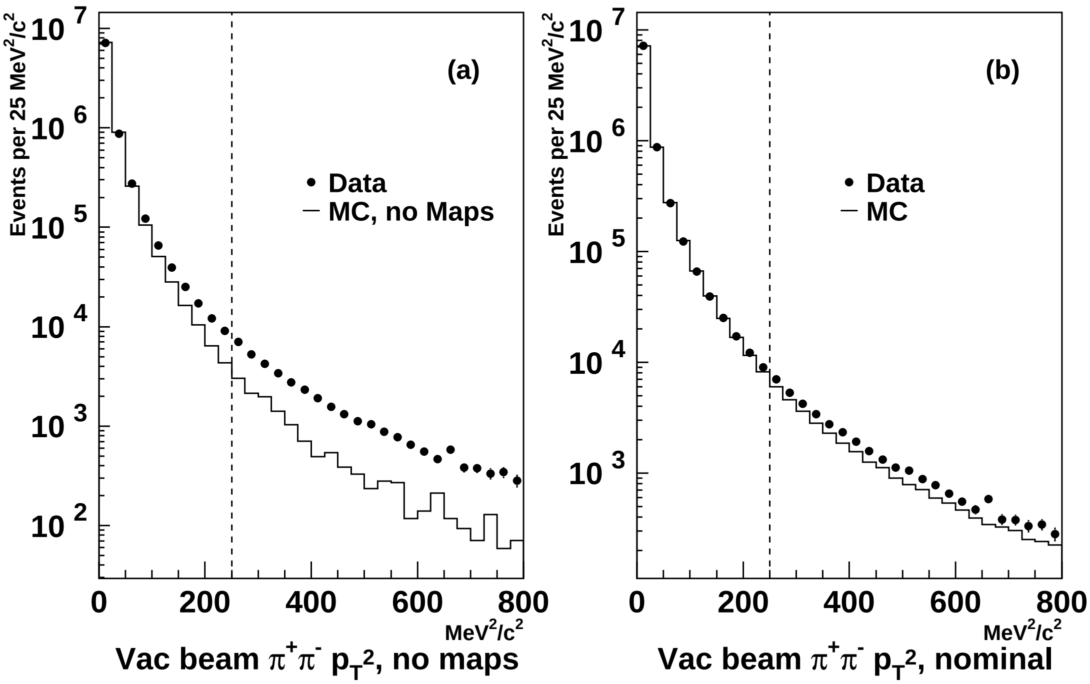

Figure 8.11: The effect of the high-SOD simulation on the p_T^2 distribution. Panel (a) shows the vacuum beam data (dots) compared to a Monte Carlo run without the high-SOD simulation (histogram). Panel (b) shows the same data compared to a MC with the high-SOD simulation. The dotted line shows the value of the cut at 250 MeV$^2/c^2$.

that could be modeled by the geometry of the cell. If there are two early accidentals on neighboring wires, one can determine the position of the early track within the cell in a procedure analogous to finding the SOD. From the position of the early track, one can determine the geometric path-length of the track within the cell, and from that, determine the length of the pulse. This effect is shown in Figure 8.13. For large pulses, the pulse will stay above threshold for the full amount of time the early track traverses the cell.

It was found that this "geometric" early accidental effect occurs for $\sim 40\%$ of early accidentals in the chambers. For the remaining cases, we used the exponential model.

It should be noted that the timing window for in-time hits in the drift chamber

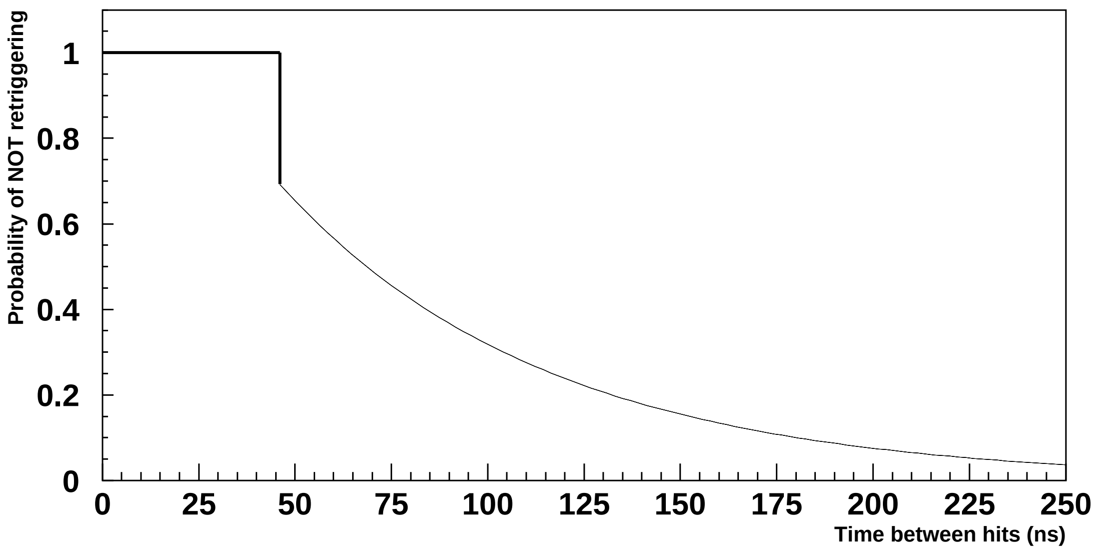

Figure 8.12: The exponential model for early accidental inefficiency. A wire in the chamber is inefficient after an early hit for 42 ns due to the discriminator pulse-width, and has a varying inefficiency at later times.

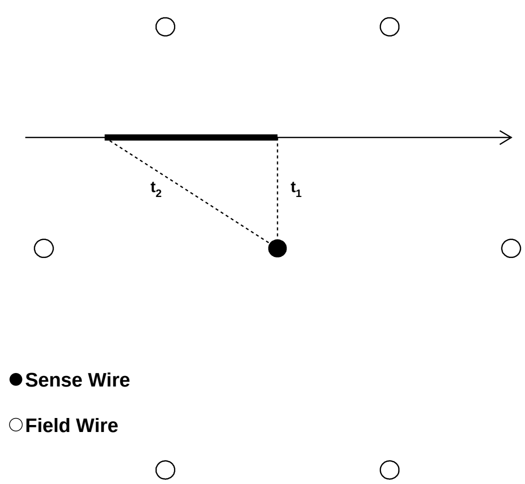

Figure 8.13: The "geometry" model for early accidental hits. By measuring the early hits on neighboring wires, we can determine the position of the track within the cell. In this model we assume the pulse is above threshold for the entire time it traverses the cell, indicated by the darker line, or time $t_2 - t_1$.

system is centered on the timing range of the TDC. The in-time window is 220–700 TDC counts. We can only measure early hits if they occur within the window of 700–1024 TDC counts (since we are in common-stop mode). Hits earlier that this window are not recorded by our drift chamber electronics, but still may affect later in-time pulses. We will set a systematic uncertainty due to this effect in Section 10.9.4.

8.5 Simulation of the Calorimeter

The full simulation of the calorimeter is discussed in Reference [1]. Here we detail some of the pertinent facts about photon and hadronic pion interactions in the calorimeter.

For the neutral-mode analysis, we are concerned with the full non-Gaussian behavior and the tails of the response of the calorimeter, such that our simulation matches the data as accurately as possible. To that end, we use the GEANT package [80] to simulate all the relevant interaction processes in the CsI crystals, including pair production and bremsstrahlung. We trace the shower particles down to energies of a few MeV. This simulation is quite time-consuming, and we cannot generate electromagnetic showers for each event. Instead we create a library of pre-generated showers that are accessed by the Monte Carlo as necessary.

Hadronic interactions for $\pi^{\pm}$s are also simulated using a library of showers generated with the GEANT package. Pions of various energies are propagated through the VV$'$ counters and the CsI and allowed to interact. The showers are stored in 41×41 crystal array. If a pion showers in VV$'$, it is possible that the hadronic clusters in the calorimeter will be more than 7 cm from the extrapolated track position; this is a source of track loss.

8.6 Accidental Overlays

The high flux of kaons and neutrons in the KTeV apparatus means that events that have a real kaon decay can also have significant underlying activity within the detector. The most significant source is the activity from hadronic interactions in the regenerator. The beam impinging on the regenerator has an interaction rate of ~ 1.5 MHz. Although most of these interactions cause the trigger system to be vetoed, there may be remaining activity in the detector a few buckets after the vetoed event, which may be reconstructed as early or in-time activity for a later real event. In addition, stray particles from upstream kaon and hyperon decays or from beam interactions with the material of the detector can contribute to underlying activity without firing a veto counter. We are concerned that the presence of additional activity can bias the reconstruction of the event, for example by depositing extra energy under real CsI clusters, or causing there to be too many hits in the drift chamber system. Furthermore, there may be a positional dependence on the accidental activity that could cause a beam-to-beam difference, biasing the acceptance and our result for $\mathrm{Re}(\epsilon'/\epsilon)$.

To model these effects, we have recorded "accidental" events to tape, taken concurrently with our nominal dataset, as described in Section 3.1.4. The trigger for these data was a coincidence of counters near to and oriented 90° from the target. This trigger allows us to record data accidental events that have the correct distribution of activity in the detector, since the trigger is correlated with beam activity, which would not be possible with a random trigger.

We use the data accidental event in the generation of the Monte Carlo event. First, the position of the regenerator is fixed in the MC event to match the position of the regenerator in the accidental event. Note that this does not determine if the event is a vacuum beam or regenerator beam event. Second, we overlay the recorded accidental activity in each detector system. For example, we add energy to each Monte Carlo cluster if the accidental event has energy in the channels used by the cluster. The readout threshold of the detector is lower for accidental events, so we can correctly account for below-threshold effects in the Monte Carlo. Energy is overlaid in each of the veto counters, if present, and the trigger sources are also overlaid. Either extra energy or a trigger bit could cause the trigger to veto the Monte Carlo event. As discussed, the drift chamber simulation (see Section 8.4) has a model for the width of each analog pulse, which may obscure subsequent hits.

This model is applied to accidental hits, which may obscure real Monte Carlo hits, and cause event loss.

The effect of accidentals on the acceptance is rather small, at the level of 1×10^{-4}. We will evaluate the systematic uncertainty due to accidentals in Chapter 10.

8.7 Simulation of the Trigger System

The KTeV trigger system, as described in Section 3.7, is based on digital logic, and is, for the most part, easy to simulate. The efficiencies of the counters are nearly 100%, so there is no model for loss due to counter inefficiency. The simulation accounts for time slewing in the long vertical counters of VV′ and the muon counter system, and has a straight-forward model for the coincidence logic of these trigger elements. The drift chamber trigger elements (YTF, bananas/kumquats, and DC-ORs) are based on the recorded digital hits and the width pulses from the model described above.

The neutral mode trigger elements have more intricate models of the response of the trigger to the analog input signals. The low-energy threshold curve for each channel of the ET system was measured in data and used as input for the trigger simulation in the Monte Carlo. There is an additional inefficiency to high-energy pulses that stay above the ET threshold, and this effect is also modeled in the Monte Carlo [1].

8.8 Generation and Analysis of the Monte Carlo Samples

Due to the high-statistics nature of the KTeV experiment, we were required to generate large amounts of Monte Carlo data to calculate our acceptance. We generated $> 900 \times 10^6$ $K \to \pi^+\pi^-$ decays, on a “farm” of 15 dual-processor 500 MHz Intel Pentium III Linux workstations, corresponding to 4.4 times the $K \to \pi^+\pi^-$ data sample. For $K \to \pi^0\pi^0$, we generate $> 700 \times 10^6$ events (12 times the data sample) on the same Linux farm.

These MC samples were analyzed in the same manner as the data. All of the Monte Carlo was calibrated and analyzed int the same manner as the data; for the charged-samples, the biggest effect is using MC-specific $x(t)$ calibrations. In the neutral reconstruction, the transverse look-up maps were generated from Monte Carlo, since it is known that the width of the showers in the MC is narrower than in the data. Finally, specific corrections made to the data that are not simulated in the MC are not corrected in the Monte Carlo. The only correction not applied in the track reconstruction is the fringe-magnetic field correction. In the neutral mode reconstruction, we do not correct for DPMT non-linearities or time dependence of the channel gains.

In addition to the nominal analysis cuts, we also ran the Level 3 software trigger, using the correct "online" database constants for a given run. The Level 3 trigger decision was stored in the event header, as it was in data. We have found a $\sim 0.8\%$ loss in both beams, consistent with the loss in data. However, there is no apparent bias between the two beams ($0.06 \pm 0.05 \times 10^{-4}$), and our nominal acceptance MC sample does not include the cut on the Level 3 trigger bit.

The MC produces a table indicating how many events were generated in each beam as a function of (E_K, z). The analysis of the MC samples produces a similar table of how many events are reconstructed and pass the selection cuts. The ratio of these two tables is the acceptance. It is possible for events generated in one (E_K, z) bin to migrate into another bin; we still calculate the acceptance without an explicit correction for migration, since we assume that the data migrate in a similar fashion. The average acceptance, for the region of interest ($40\text{ GeV} < E_K < 160\text{ GeV}, 110\text{ m} < z < 158\text{ m}$) is shown in Table 8.1. The fitting program does not use these average numbers, for reasons that will be detailed in Section 9.2.1, but makes full use of the (E_K, z) tables to calculate the acceptance in small (E_K, z) bins. The acceptance as a function of (E_K, z) for the vacuum beam for $K \to \pi^+\pi^-$ and $K \to \pi^0\pi^0$ is shown in Figure 8.14. It should be noted that the acceptance in a given (E_K, z) bin is almost identical between the vacuum and regenerator beams. The absolute acceptance is not critical for this measurement; only the relative acceptances in each beam as a function of position and energy matter. For example, the cut on "bad runs," as detailed in Section 6.6.1, can affect our calculation since we will keep events in those runs in the "generated" total, but throw them out in the "accepted" total.

Table 8.1: Average acceptance for each of the $\pi\pi$ samples.

Sample	Average Acceptance (%)
Vacuum-beam $\pi^+\pi^-$	20.77
Regenerator-beam $\pi^+\pi^-$	27.62
Vacuum-beam $\pi^0\pi^0$	10.29
Regenerator-beam $\pi^0\pi^0$	12.98

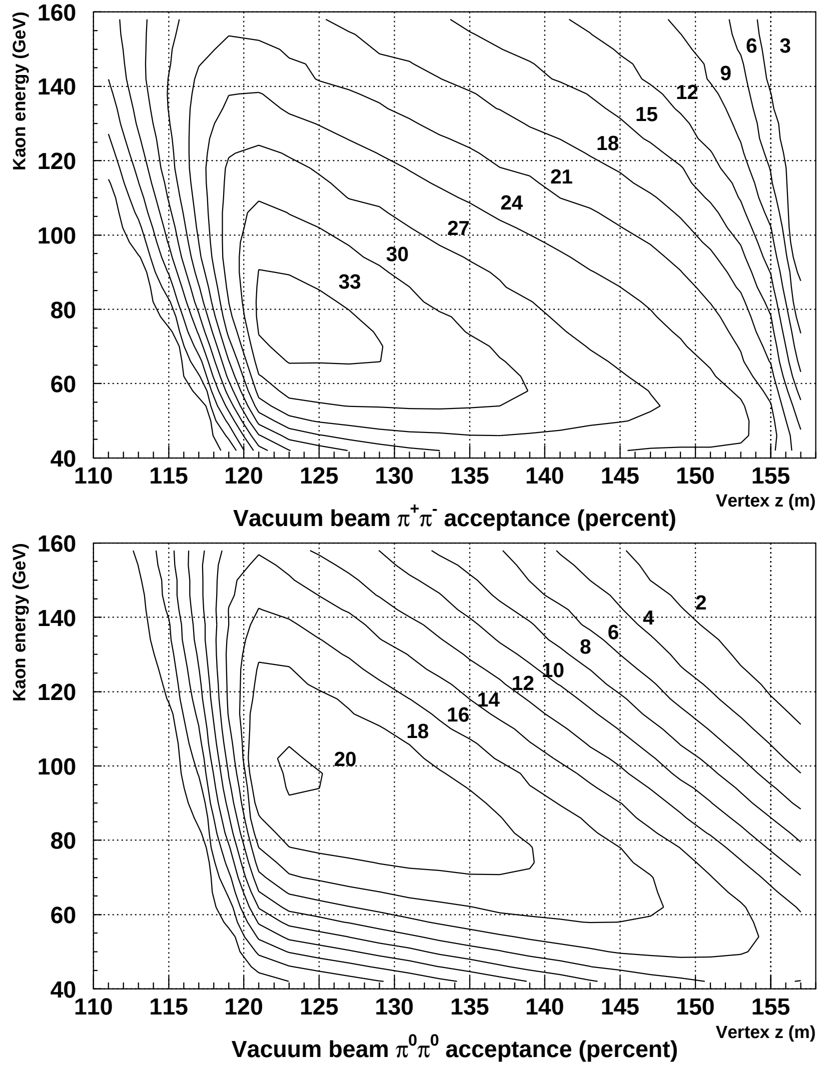

Figure 8.14: Acceptance for the $K \to \pi^+\pi^-$ (top) and $K \to \pi^0\pi^0$ (bottom) decays in the vacuum beam as a function of (E_K, z). From the 1997b sample.

CHAPTER 9
EXTRACTION OF THE KAON SECTOR PARAMETERS AND RE(ϵ'/ϵ)

With the work presented in previous chapters, we have the number of observed events in each of the four $\pi\pi$ samples, and the acceptance for each. Naïvely, one might think we could determine Re(ϵ'/ϵ) from a grand ratio:

$$\mathrm{Re}(\epsilon'/\epsilon) \stackrel{?}{\approx} \frac{1}{6}\left[\frac{\frac{N(\mathrm{Vac}\ \pi^+\pi^-)}{A(\mathrm{Vac}\ \pi^+\pi^-)} \Big/ \frac{N(\mathrm{Reg}\ \pi^+\pi^-)}{A(\mathrm{Reg}\ \pi^+\pi^-)}}{\frac{N(\mathrm{Vac}\ \pi^0\pi^0)}{A(\mathrm{Vac}\ \pi^0\pi^0)} \Big/ \frac{N(\mathrm{Reg}\ \pi^0\pi^0)}{A(\mathrm{Reg}\ \pi^0\pi^0)}} - 1\right]. \tag{9.1}$$

There are two important reasons that this approach will not work. First, the regenerator beam is not a pure K_S beam; there is significant K_S-K_L interference and transmission of the K_L component, such that $N(\mathrm{Reg} \to \pi\pi)$ is not equal to $N(K_S \to \pi\pi)$. If the acceptance between $\pi^+\pi^-$ and $\pi^0\pi^0$ decays were the same, this effect would merely dilute the measured value of Re(ϵ'/ϵ) by a multiplicative factor. Given the fact that the acceptances are not the same, there can be an additive bias.

The second reason that Equation 9.1 is not precise enough to determine Re(ϵ'/ϵ) is that it would require the global acceptance for the four modes to be known extremely well. That requirement, in turn, would mean that the Monte Carlo simulation would have to match the data exactly. Although our simulation matches the data quite well, as shown in Chapter 8, we have ignored certain disagreements between the data and the Monte Carlo. The largest example is the kaon energy distribution; we have a simple quark-counting model relating production of $K^\pm$ at lower energies to K^0-$\overline{K^0}$ production at 800 GeV (Section 8.1.1), but this does not match the Monte Carlo to the data exactly. Furthermore, the details of the energy distribution mismatch are different between the $\pi^+\pi^-$ and $\pi^0\pi^0$ modes, indicating different energy-dependent effects between the two modes. Were we to use a global acceptance, we would incorrectly weight the acceptance from different kaon energies when calculating the average acceptance. In the regenerator beam, we have an additional problem due to the uncertainties in the kaon sector parameters Δm and τ_S, and the regeneration parameters, which can couple to variations in the acceptance in z and affect the average acceptance. We have tuned the Monte Carlo to be as close to the data as possible before generating the large MC acceptance datasets, so we don't expect these effects to be large. However, if we were to use the ratio calculation only, we would have to modify the Monte Carlo acceptance sets to even test for these effects. Instead, we choose to bin the acceptance in (E_K, z), so that we can ignore these effects for reasons that will be detailed in Section 9.2.1.

We have developed a program called KFIT[1] as a tool to extract $\mathrm{Re}(\epsilon'/\epsilon)$ and other physics parameters from our dataset. The basic operating principle of KFIT is to calculate kaon decay distributions using the physical parameters such as Δm, τ_S, and $\mathrm{Re}(\epsilon'/\epsilon)$ and a regeneration model, then multiply these decay distributions by the acceptance determined by KTEVMC to produce a *prediction function* for the number of events we expect to see in a given (E_K, z) bin. The prediction function is compared to the *observed* number of events. The quantities of interest (such as $\mathrm{Re}(\epsilon'/\epsilon)$ and Δm) are varied to minimize the χ^2 of the comparison between the prediction and the observed data. A flow-chart of the operation of the fitter is shown in Figure 9.1.

We will describe the different types of fits we can perform with KFIT, and then describe the details of the fitting software. We will then present our results for fits to the kaon sector parameters Δm, τ_S and ϕ_{+-}, which we will then use as inputs to our fit for $\mathrm{Re}(\epsilon'/\epsilon)$. This is followed by a summary of some of the cross-checks to our result we have performed.

[1] KFIT is derived from an earlier fitting program used in E773 [33], which in turn was descended from the fitting program used in E731.

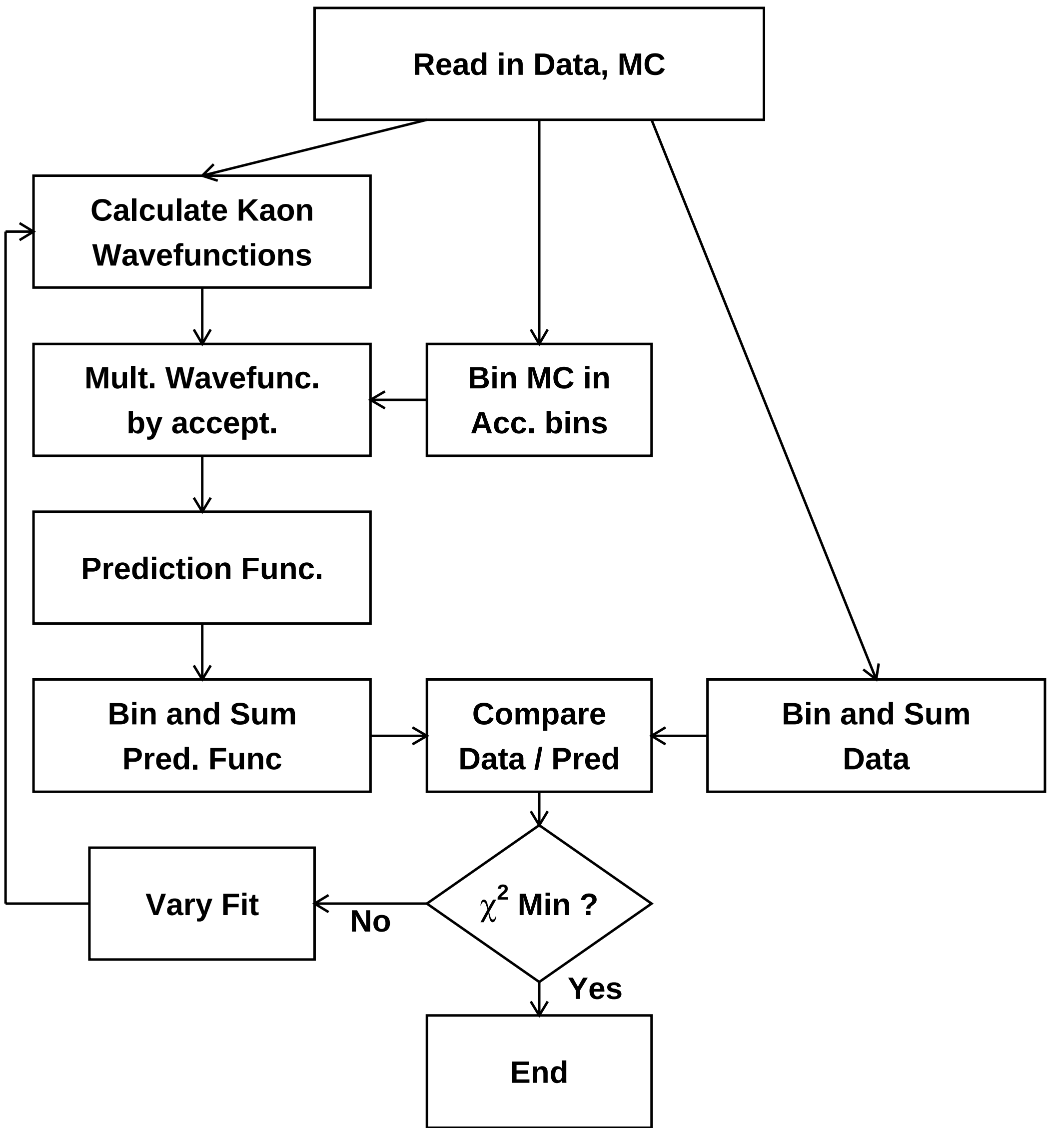

Figure 9.1: Flow chart of KFIT operation.

9.1 Types of Fits

The fits performed on the E832 data break up into two large classes: unbinned and binned fits, where we are referring specifically to binning in the decay z position. For the measurement where we are interested in the *rate* of a given process, such as $\mathrm{Re}(\epsilon'/\epsilon)$, we perform a fit in a single z bin, since we are only concerned with the total number of events. By doing the fit in one large z-bin fit, we minimize our exposure to resolution and migration effects, especially in the regenerator beam, in which the decay distribution varies rapidly with z. More specifically, we minimize our exposure to a mismatch in the migration in the data as compared to the Monte Carlo acceptance (and hence KFIT's prediction function). The specifics of the unbinned fit for $\mathrm{Re}(\epsilon'/\epsilon)$ will be discussed below.

For measurements where we are concerned with *shape* of the decay distribution, we have to use a z-binned fit. τ_S, Δm and ϕ_{+-} can be measured through the shape of regenerator beam decay vertex distribution; τ_S through the steeply falling exponential decay, and Δm and ϕ_{+-} through the interference between the K_S and K_L terms (see Equation 3.1). The phase ϕ_{+-} should be related to Δm and $\tau_{S,L}$ through the superweak phase ϕ_{SW}, as detailed in Section 1.4.

9.2 The Fitting Software

9.2.1 Inputs and Binning

The main KFIT inputs are 2-dimensional (E_K, z) tables of background-subtracted data, reconstructed coherent Monte Carlo events, and the underlying generated MC events. Each of these tables exists for $\pi^+\pi^-$ and $\pi^0\pi^0$ decays in each of the vacuum and regenerator beams, for each regenerator position (left or right). Our nominal fits combine the left and right beam tables with a simple sum, which is valid because the KTeV beam intensities are equal to a fraction of a percent. We kept the left and right distinction when analyzing the data for historical reasons; E731 combined the left and right beam tables using a geometric mean to reduce their sensitivity to an 8% difference in their beam intensities [62]. However, having the left and right beam information will allow for a data cross-check.

As discussed in Chapters 6 and 7, we use data and Monte Carlo tables in the ranges 40 GeV $< E_K <$ 160 GeV and 110 m $< z <$ 158 m. The z range is defined by the geometry of the detector. The energy cuts are defined by the energy spectrum of the beam. There is a significant amount of charged mode data below 40 GeV, and these data also exhibit the largest interference effects in the regenerator beam, due to the low boost factor, and including these data would improve the statistical determination of Δm and τ_S. However, that region of kaon energy has systematic issues with the treatment of screening and transmission effects in the regenerator. To avoid these problems and for consistency, we apply the same E_K and z cuts in both the neutral and charged-mode analyses for all the fits we perform.

We bin the data into "fit" bins. Fit bins are uniformly 10 GeV wide in E_K. For unbinned z fits, such as $\mathrm{Re}(\epsilon'/\epsilon)$, we sum all the data in one E_K bin into one z bin. For z-binned fits, such as those for the kaon sector parameters, we divide the regenerator beam data in 2 m z-bins, to allow us to extract the shape of the vertex decay distribution. In the z-binned fits, the vacuum beam determines the kaon flux, and has no dependence on the kaon sector parameters. Hence we do not bin the data in z in the vacuum beam.

The use of small acceptance bins allows us to sidestep the issue of data/MC mismatches in the overall E_K distribution and mismatches in the z distribution in the regenerator beam due to using incorrect values for Δm, τ_S, and the regeneration parameters when generating the Monte Carlo. The important factor is that the data/MC ratio does not vary greatly *within* a small acceptance bin; the fact that the data and Monte Carlo E_K and z distributions do match globally quite well means this will be true. Effectively, KTEVMC averages the acceptance over the bin with the proper weight within the bin. We use 2 GeV acceptance bins in E_K in both beams. In the regenerator beam, where the kaon sector parameters couple to the z distribution, we use 2 m acceptance bins. The vacuum beam has no dependence on the kaon sector parameters; hence we use one large z-bin to avoid migration issues. KFIT rebins the reconstructed and generated MC tables into the appropriate bins, and then divides the two to determine the acceptance and its error in each bin.

The fit-bin and acceptance-bin schemes are shown in Figure 9.2.

9.2.2 Calculation of Decay Distributions

On each iteration of the fit, KFIT propagates K^0 and $\overline{K^0}$ states from the target to the beginning of the decay region for each beam ($z = 110$ m in the vacuum beam and the downstream end of the regenerator in the regenerator beam). The code used to calculate this propagation is very similar to that in KTEVMC. The propagation

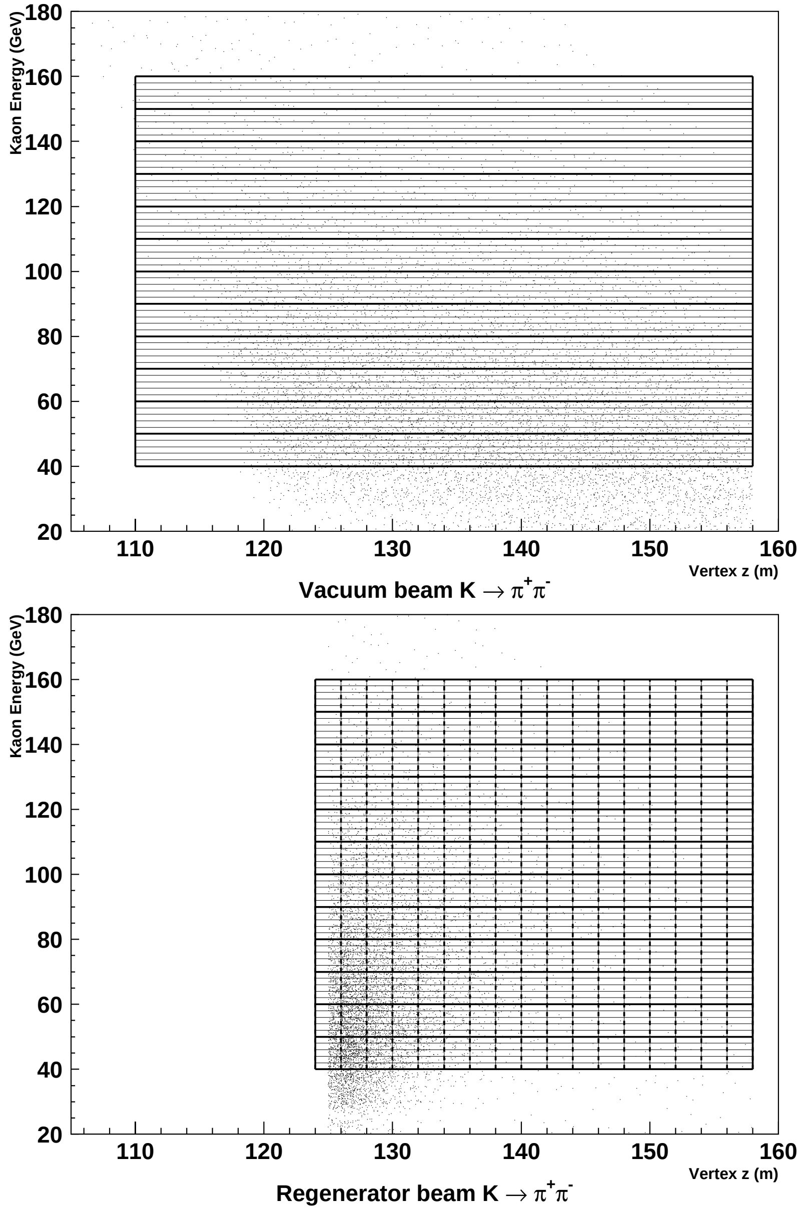

Figure 9.2: Binning used for the fit. The thin lines show the acceptance bins, and the thick lines show the fit bins. The top panel shows the vacuum beam, for which the binning is the same regardless of the type of fit we perform. The bottom panel shows the regenerator beam; the acceptance binning (thin lines) is same across fit types, while the fit binning in z depends on the type of fit. The dots represent the $K \rightarrow \pi^+\pi^-$ data; each dot in the top panel represents 1200 events, while each dot in the lower panel represents 2400 events.

calculation is done for kaons in an energy range of 40.5 to 159.5 GeV in 1 GeV steps, and uses several parameters that may be fixed or floated in the fit, such as Δm, τ_S, and the regeneration parameters, creating the need to recalculate these functions during each iteration of the fit. At high energies, a significant "primary" K_S component from the target can reach the decay region and affect the decay distributions. This K_S component is visible at high energies in the vacuum beam in both $K_L \to \pi^+\pi^-$ and $K_L \to \pi^0\pi^0$ decays. As a cross-check, we can float the amount of primary K_S to match our data; the systematic uncertainty from this effect will be discussed in Section 10.10.

In a manner similar to KTEVMC, the energy dependence for kaon production is taken from the Malensek spectrum, suitably modified by the polynomial correction (Section 8.1.1). The energy spectrum is further modified to account for kaons scattering in the absorbers which remain in our beam and are counted as coherent events. The correction is implemented as a polynomial determined from a Monte Carlo study, and is different for the two beams, due to the presence of the movable absorber in the regenerator beam. Even with all these corrections, we do not claim to understand the energy spectrum of kaon production fully. Furthermore, there may be issues with our simulation of inefficiencies in the detector that may affect the neutral and charged mode reconstructions differently, resulting in different reconstructed energy distributions in the two modes. For example, the hadronic pion showers in the VV$'$ cause a slight energy-dependence in track-cluster matching efficiency that we do not absolutely predict with the Monte Carlo, and this is affects the charged-mode but not the neutral-mode. Therefore, we float the kaon flux normalization (a multiplicative correction to the nominal energy distribution) independently in each 10-GeV kaon energy bin, with separate parameters for charged and neutral decays. However, the parameters are the same within a decay mode for the regenerator and vacuum beams, and hence cancel in the double ratio.

A further input to KFIT is the net transmission of the kaon beam through the regenerator, *i.e.*, the ratio of kaon flux between the regenerator and vacuum beams. We do not know the kaon-nucleon cross sections well enough to predict the transmission accurately, so we measure the transmission from the data sample.

Using $K_L \to \pi^+\pi^-\pi^0$ decays, we measure the average transmission in the regenerator beam (at 70 GeV) to be $(7.70 \pm 0.023)\%$, with an energy dependence of $(-0.0040 \pm 0.0009)\%$ GeV^{-1}. Considerable effort has been invested in understanding the shape of the transmission below 40 GeV, where predictions indicate that it should start to deviate from a linear slope [77], as shown in Figure 9.3. We have considered models that allow for both elastic and inelastic scattering within the regenerator. The predictions of these models for the transmission are quite similar. Each of these models also predicts effects in due to screening in regeneration, and we will discuss them in more detail in that context. Ideally, we should be able to use these models to predict the transmission of the regenerator; however, with the uncertainty in the kaon-nucleon cross-sections, we cannot predict the shape of the transmission below 40 GeV. It is for this reason that we do not use charged-mode data below 40 GeV in our fits.

The final input to KFIT is the model for coherent regeneration. Regeneration arises through the difference between the forward kaon-nucleon scattering amplitudes for K^0 and $\overline{K^0}$. The quantity of interest is

$$f_- \equiv \frac{f(0) - \bar{f}(0)}{k}, \tag{9.2}$$

where $f(0)$ and $\bar{f}(0)$ are the forward scattering amplitudes for K^0 and $\overline{K^0}$, respectively, and k is the kaon wave number. f_- is related to the regeneration parameter ρ via

$$\rho = i\pi N L f_- g(L), \tag{9.3}$$

where N is the number density of nuclei, L is the length of the regenerator, $\hbar k$ is the kaon momentum and $g(L)$ is the geometric correction for the time evolution of the regenerated K_S relative to the K_L as they traverse the regenerator. To first order, for the energy range of interest, Regge theory [85] predicts that the magnitude of f_- should vary with kaon energy as a power law. We express this dependence with

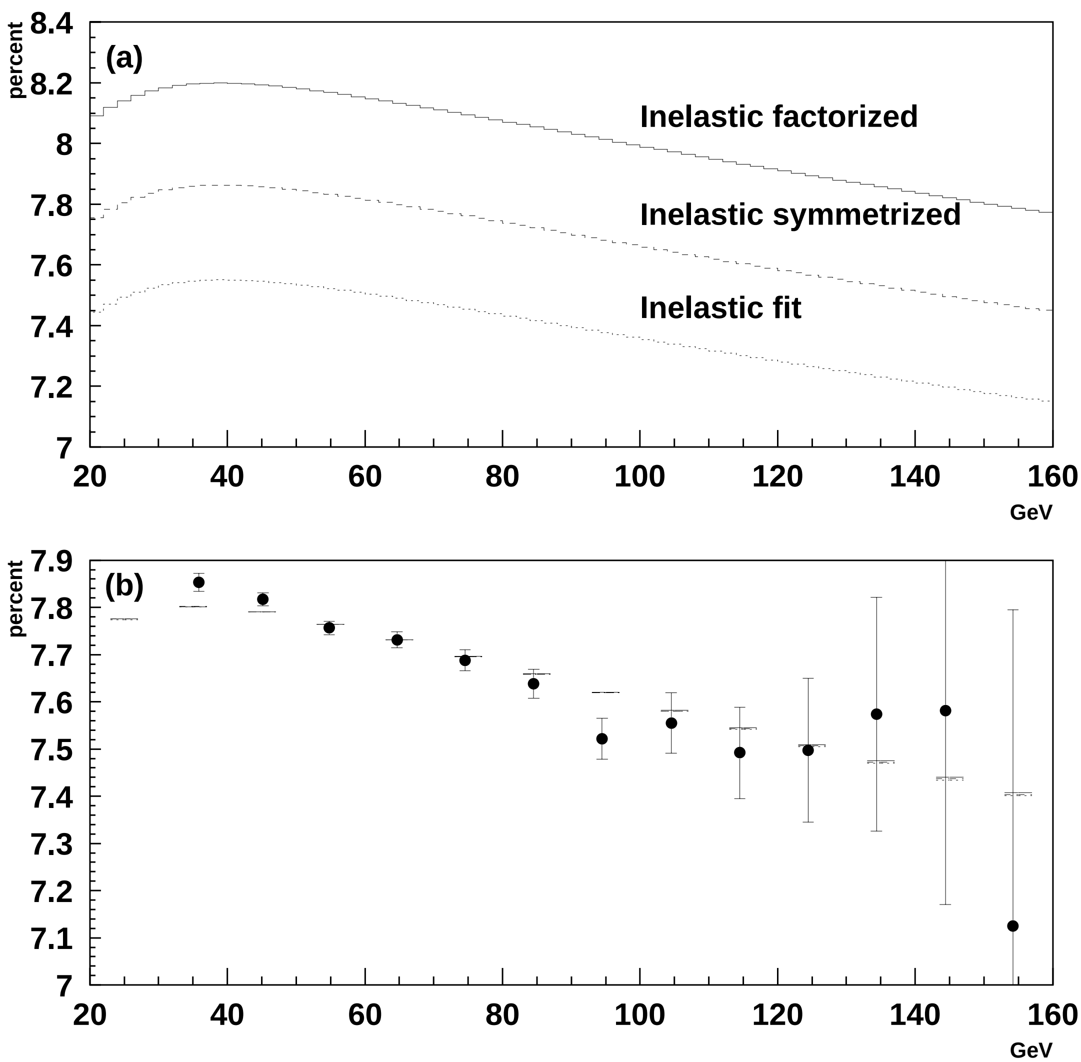

Figure 9.3: Transmission in the regenerator. Panel (a) shows predictions for the transmission under various scattering assumptions. Panel (b) shows the transmission as measured in $K_L \to \pi^+\pi^-\pi^0$ decays, with the three predictions normalized to the data. The normalization puts the three models nearly on top of each other. The disagreement between the data and the predictions below 45 GeV indicates the predictions are not reliable enough allow us to use data below this point.

respect to our nominal energy

$$|f_-(E_K)| = |f_-(70\text{ GeV})| \left(\frac{E_K}{70\text{ GeV}}\right)^{\alpha}. \tag{9.4}$$

Figure 9.4 shows a power-law fit to the values of $|f_-(E_K)|$ as measured in KTeV $\pi^+\pi^-$ data, showing the good agreement with the power-law. The complex phase of f_- can be determined from its energy dependence through an integral dispersion relation, with the requirement that the forward scattering amplitudes be analytic functions. For a pure power-law energy dependence, this yields a constant phase,

$$\arg(f_-) = -\frac{\pi}{2}(2+\alpha). \tag{9.5}$$

In addition to the pure power-law behavior, we make a small correction for nuclear screening effects in the regenerator, including both elastic and inelastic terms. The correction affects both the power-law and the phase ϕ_ρ to preserve the analyticity relationship. This screening correction maps our pure-power law to a phenomenological variable related to the momentum dependence of Regge ω exchange, allowing for better comparison to direct regeneration measurements such as Reference [86].

The nuclear screening calculations are based on computer codes detailed in Chapter 9 and Appendix H of Reference [33]. The inputs to the screening calculations are KN cross-section data and the nuclear densities of the material in the regenerator. Regge-inspired functions of the form detailed in Reference [87] are used in the fits. The bulk of the screening effect is due to elastic screening. Inelastic scattering, where the incoming K scatters into an intermediate K' state which then reforms into a K after a subsequent scatter, is evident in fits to total cross-section data [87]. Our fits allow for both C-even terms, as determined from the total cross-section data [88], and C-odd terms [89]. The level of the C-odd terms is varied amongst the models.

We have chosen to use the screening model based on the inelastic calculation with the level of C-odd terms set by factorization. This choice was motivated by preliminary fits to the deviation of the regeneration parameters from a pure power-law in $K \to \pi^+\pi^-$ data. However, the low-energy data in our $K \to \pi^+\pi^-$ measurements

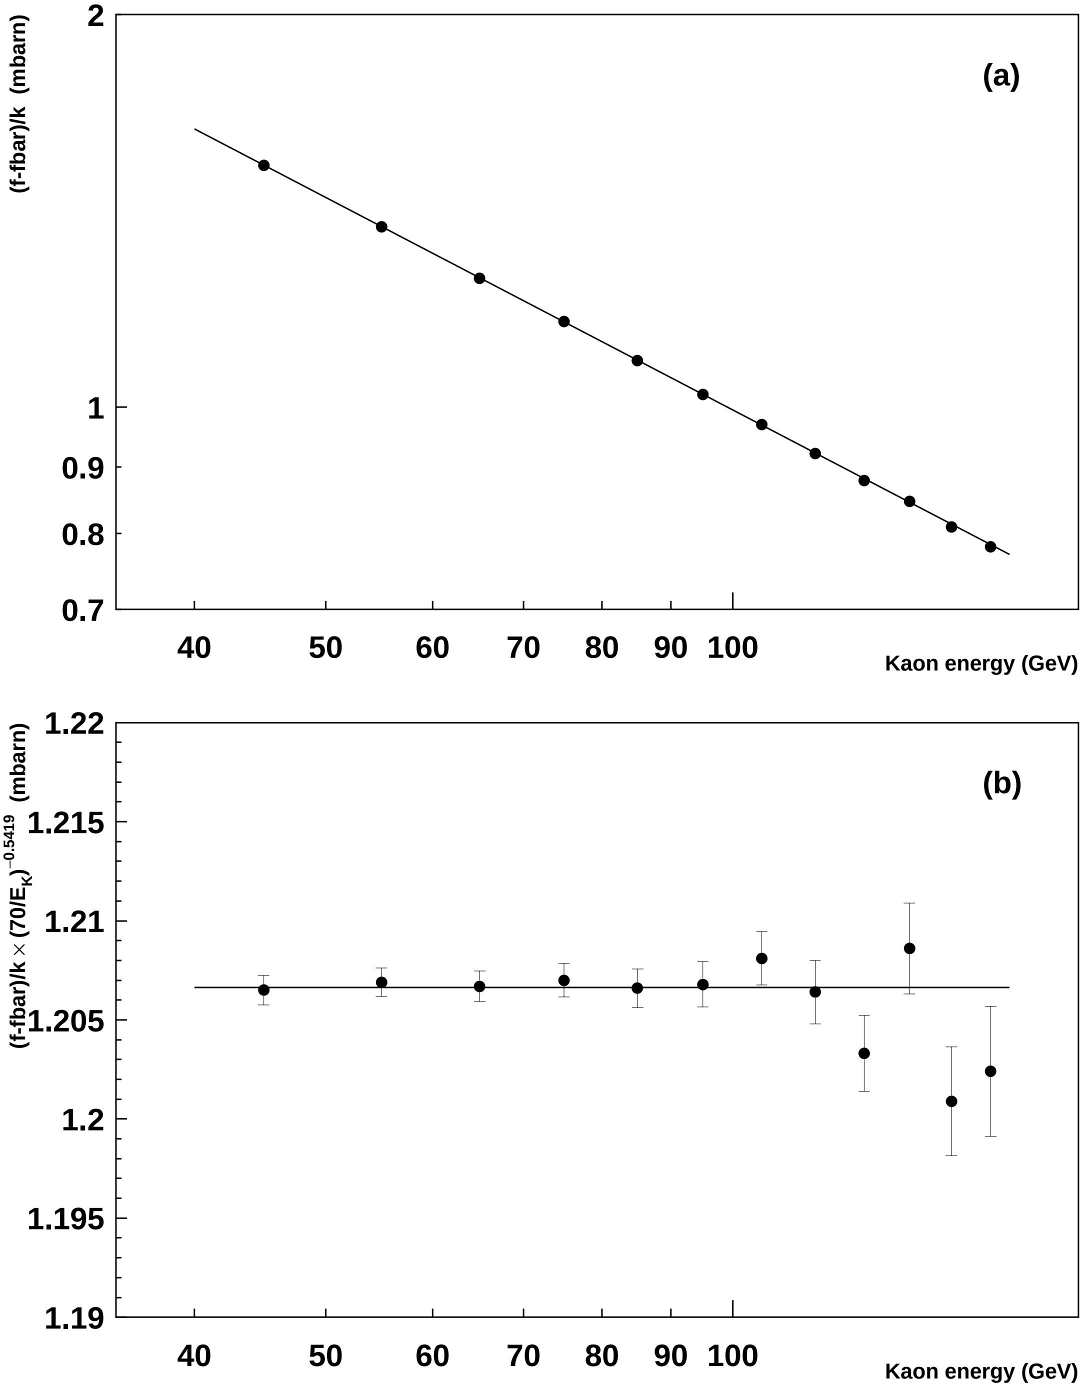

Figure 9.4: Power law fit to $\left|\left(f(0) - \bar{f}(0)\right)/k\right|$ measured from KTeV $\pi^+\pi^-$ data.

depend on our measurement of the transmission of the regenerator, as shown in Figure 9.3, where we currently do not have a measurement below 40 GeV. However, the calculation that predicts the nuclear screening correction also predicts the transmission of the regenerator. In the future, we should be able to simultaneously fit the transmission and regeneration data to determine the best parametrization of the inelastic terms in the regenerator. Note that the inelastic terms do not contribute greatly to transmission (Figure 9.3) but they do contribute to regeneration. Because of the impact of the screening correction on the phase ϕ_ρ, there is a systematic uncertainty due to the screening correction, especially in our measurement of ϕ_{+-}. This will be discussed in Section 10.2.2.

The main component of the regenerator is plastic scintillator. We allow the values of $|f_-(70\text{ GeV})|$ and α for plastic scintillator to float in our fits. For the lead piece of the regenerator and the beryllium and lead in the absorbers, we use fixed values measured from previous experiments [79]. The propagation of the kaon wavefunction through these objects is a full quantum-mechanical matrix transformation, which calculates regeneration exactly. No assumptions are made about the size of the object relative to $c\tau_S$.

The end result of the calculation is a kaon wavefunction (with K_S and K_L amplitudes $a_S^{V,R}$ and $a_L^{V,R}$) at the start of the decay region, either in the vacuum or regenerator beam. The $\pi\pi$ decay rate as a function of z, in a manner similar to Equation 3.1, is

$$\begin{aligned}\frac{d\Gamma}{dt} \;\propto\;& \left|a_S^{V,R}\right|^2 e^{-t/\tau_S} + \left|\eta a_L^{V,R}\right|^2 e^{-t/\tau_L} \\ & +2\left|a_S\right|\left|\eta a_L\right| \cos\left[\Delta m t + \arg(a_S^{V,R}/a_L^{V,R}) - \phi_\eta\right] e^{-t/(\tau_S/2+\tau_L/2)} \; . \end{aligned} \quad (9.6)$$

The vacuum beam is effectively all K_L, and there are no interference effects; the terms are of the same order in the regenerator beam. Note that $a_{S,L}^{V,R}$ are independent of the decay mode ($\pi^+\pi^-, \pi^0\pi^0$), while η is different between the two modes. The functional definition for the $\mathrm{Re}(\epsilon'/\epsilon)$ fits in KFIT is

$$\eta_{+-} = \epsilon\left[1 + \mathrm{Re}(\epsilon'/\epsilon)\right] \quad (9.7)$$

$$\eta_{00} = \epsilon\left[1 - 2\mathrm{Re}(\epsilon'/\epsilon)\right], \quad (9.8)$$

as detailed in Equation 1.38 as the effective definition of $\mathrm{Re}(\epsilon'/\epsilon)$, assuming ϵ_K and ϵ' have the same phase. KFIT attempts to vary the value of $\mathrm{Re}(\epsilon'/\epsilon)$ to simultaneously match the prediction function to the reconstructed charged- and neutral-mode data.

9.2.3 Calculation and Minimization of the Fit χ^2

KFIT integrates Equation 9.6 analytically in z, and sums over the 1-GeV E_K steps in the function, to calculate the number of kaon decays in each acceptance bin. The bin containing the regenerator edge is treated as a special case; the integral begins a certain distance *inside* the regenerator, since decays in the last module of the detector can escape being vetoed. This distance is fixed at -1.65 mm in the charged mode and -6.2 mm in the neutral mode, as will be discussed in Section 10.9.1.

The number of events in each acceptance bin is then multiplied by the acceptance in each bin to determine the expected number of *observed* events within the acceptance bin. The acceptance bins are summed together to form a fit bin. KFIT then compares the number of predicted events within a fit bin to the number of data events within the fit bin. The fit parameters are modified to minimize the global χ^2 of the fit

$$\chi^2 = \sum_{i=\text{Fit bins}} \frac{(N_i - P_i)^2}{\sigma_{N_i}^2 + \sigma_{P_i}^2}, \quad (9.9)$$

where N_i is the number of data events observed and P_i is the prediction. KFIT calculates the kaon wavefunction exactly, so the error σ_{P_i} arises solely from the statistical error on the Monte Carlo samples used to determine the acceptance. The statistical error on acceptance is calculated as the binomial error of the reconstructed MC events over the generated MC events within an acceptance bin, which ignores migration effects. Typically one would use $\sigma_{N_i} = \sqrt{N_i}$ as the predictor for the statistical error on N_i. However N_i is subject to statistical fluctuations; the estimator $\sigma_{N_i} = \sqrt{P_i}$ is more accurate because we typically generate five to ten equivalent datasets in our Monte Carlo, reducing the fluctuations of $\sqrt{P_i}$.

The χ^2 is minimized using the MINUIT package [90], which is a standard HEP package for minimizing a function of several variables. Standard fits run in a few minutes.

9.3 The Fits for Kaon Sector Parameters

We need to know the value of the kaon sector parameters Δm and τ_S to predict the shape of the decay distribution in the regenerator beam correctly. The KTeV data can be used to measure these parameters to a high precision, and we would like to use our best-fit values of these parameters as input to our fit for Re(ϵ'/ϵ). We first present our measurements of Δm and τ_S, for which we assume CPT, and then present our measurement of ϕ_{+-}, which is a test of CPT-violation.

9.3.1 The Fits for Δm and τ_S

The first kaon sector fit is a simultaneous fit to Δm and τ_S. For the regeneration model, we assume the power-law model, and relate the regeneration phase to the power law through the analyticity relation (Equation 9.5), with additional corrections due to nuclear screening. Furthermore, we assume CPT symmetry, such that ϕ_η is set to the superweak phase dynamically determined from the fit values of Δm and τ_S. The acceptance uses the "acceptance binning" described above, and the "fit bins" are (10 GeV $\times$ 2 m) in the regenerator beam, and (10 GeV $\times$ 48 m) in the vacuum beam.

Besides fitting for Δm and τ_S, we float the mean regeneration amplitude (ρ), the power law exponent (α), and the flux normalization in 10 GeV bins. In the charged mode, we fix the regenerator edge at $z = -1.65$ mm relative to the downstream edge of the regenerator and float a linear stretch in z between chambers 1 and 2, to allow for some possible remaining systematic issues with the location of the chambers.

The nominal fit values are:

$$\begin{aligned}\Delta m &= (0.5267 \pm 0.0006) \times 10^{10}\ \hbar s^{-1}, \\ \tau_S &= (0.8965 \pm 0.0003) \times 10^{-10}\ s, \\ \chi^2/\nu &= 210.3/198, \end{aligned} \tag{9.10}$$

where the errors are statistical. The superweak phase for these values of Δm and τ_S is

$$\phi_{SW} = 43.41^{\circ}. \tag{9.11}$$

For the nominal fit, the charged z-stretch is $(-0.7 \pm 1.2) \times 10^{-3}$ %, indicating we understand the regenerator edge and see no evidence for a stretch between the chambers.

We have performed a series of cross-checks to the Δm and τ_S results, such as breaking the data into time-bins, energy-bins, and data sub-sets. None of the data subdivisions show a statistically significant difference from the nominal fit. The top panels of Figures 9.5 and 9.6 show the variation of Δm and τ_S with kaon momentum, respectively.[2] Neither plot indicates any problem with our technique as a function of energy. The bottom panels of Figures 9.5 and 9.6 show the results of Δm and τ_S after breaking the data into various independent sub-samples; all are consistent.

9.3.2 The Fit for ϕ_{+-}

As detailed in Section 1.4, CPT symmetry implies $\phi_{+-} \approx \phi_{SW}$ up to small effects at the level of $\mathrm{Im}(\epsilon'/\epsilon)$. We can test this relation by fitting for ϕ_{+-} and comparing to the value of ϕ_{SW} calculated from our values of Δm and τ_S. In the ϕ_{+-} fit, we float all the parameters that are floated in the Δm–τ_S fits, as well as ϕ_{+-}. To be explicit, we float Δm, τ_S, and ϕ_{+-}, in addition to the regeneration parameters and

[2] Figures 9.5 and 9.6 show the cross-checks for only the 1997B dataset. The full result is completely consistent with this subset. The same caveat holds true for the cross-check plots for ϕ_{+-} in Section 9.3.2.

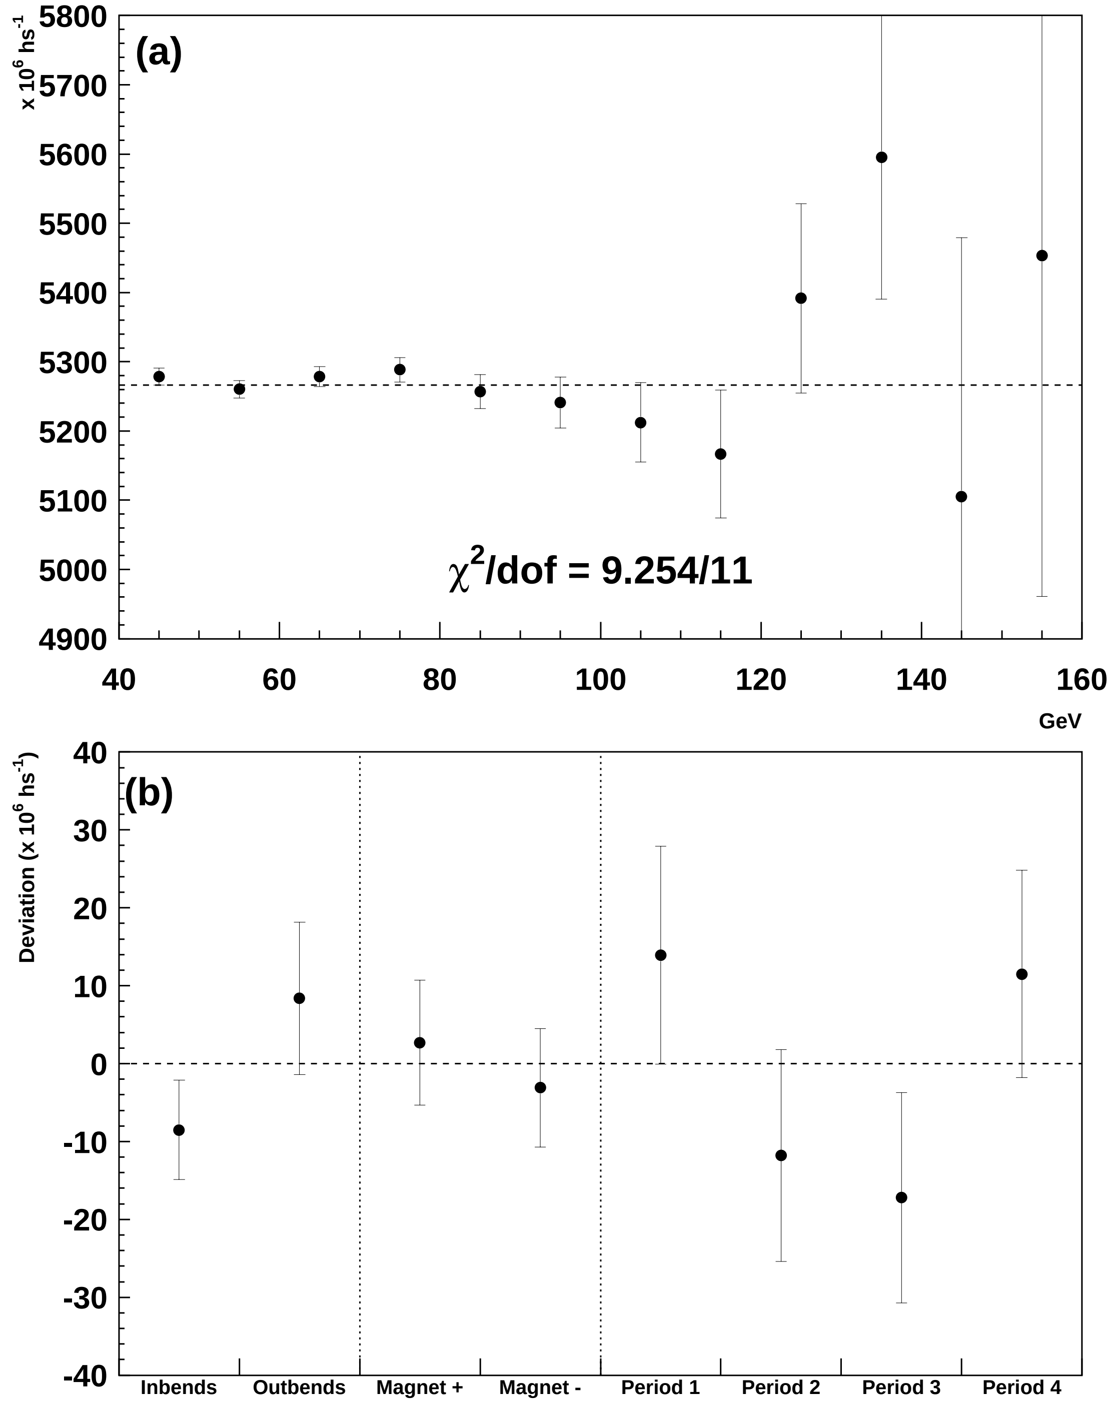

Figure 9.5: Variation of Δm with kaon momentum (a) and data subset (b) for the charged mode data. For Panel (b), the data sets between dotted lines are statistically independent.

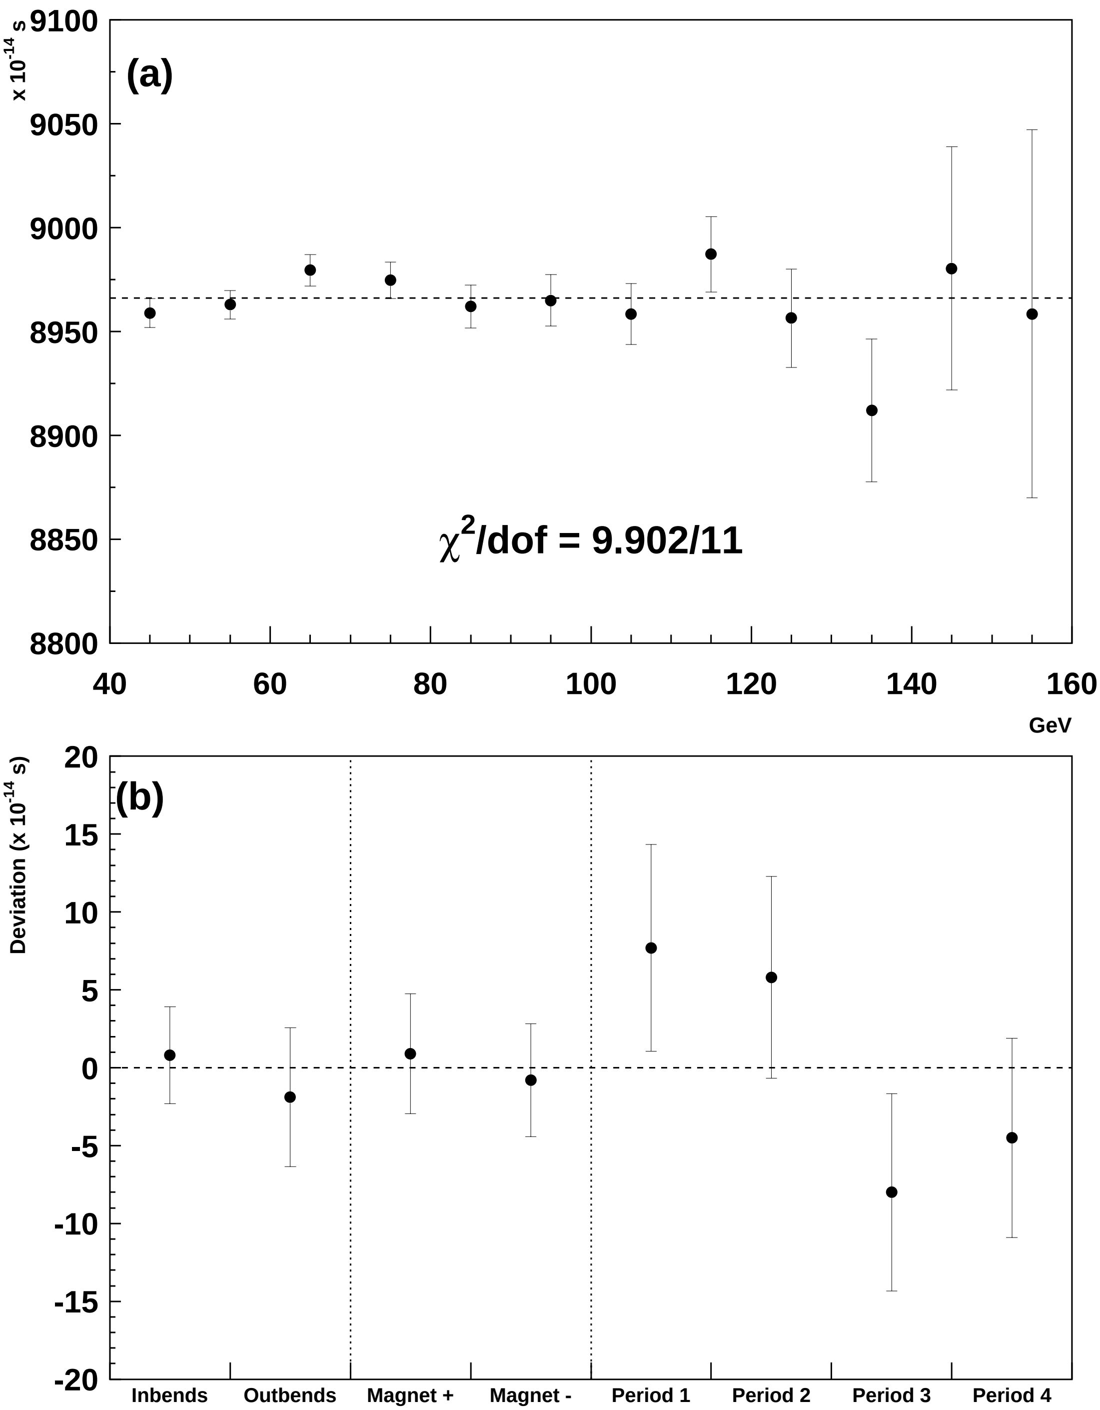

Figure 9.6: Variation of τ_S with kaon momentum (a) and data subset (b) for the charged mode data. For Panel (b), the data sets between dotted lines are statistically independent.

flux factors. Our statistical measurement of ϕ_{+-} would improve if we could fix τ_S, say to the PDG world-average, but we must be careful to not use a value and error for τ_S that was derived by assuming CPT invariance through the ϕ_{SW} assumption, as the previous FNAL result has done [58]. Measurements of τ_S that do not make a CPT assumption are statistically less precise than allowing τ_S to float in our fits.

In Equation 9.12, we show the variation of ϕ_{+-} with Δm and τ_S.

$$\begin{aligned}\phi_{+-} &= 44.12^\circ \pm 0.13^\circ(\text{stat}) \\ &\quad +0.53^\circ\left(\frac{\Delta m - 0.5288}{0.0023}\right) - 0.71^\circ\left(\frac{\tau_S - 0.8958}{0.0008}\right),\end{aligned} \tag{9.12}$$

where Δm and τ_S are in units of 10^{10} $\hbar \text{s}^{-1}$ and 10^{-10} s respectively. When we float all three parameters (Δm, τ_S and ϕ_{+-}) simultaneously, we obtain:

$$\begin{aligned}\phi_{+-} &= (44.12 \pm 0.72)^\circ, \\ \Delta m &= (0.5288 \pm 0.0023)\ \times 10^{10}\ \hbar \text{s}^{-1}, \\ \tau_S &= (0.8958 \pm 0.0008)\ \times 10^{-10}\ \text{s}, \\ \chi^2/\nu &= 223.6/197,\end{aligned}$$

where only statistical errors are shown. In fits where τ_S is fixed but Δm floats, the statistical error is $\pm 0.30^\circ$. One can use Equation 9.12 to determine ϕ_{+-} for a given value of τ_S.

In Figure 9.7, we show the variation of ϕ_{+-} versus kaon momentum and various ways of splitting the dataset. No significant variations are seen. In Figure 9.8, we show the one-sigma contours ($\Delta\chi^2 = 1$) for Δm and τ_S with ϕ_{+-} and with each other in the grand fit. For comparison, we show the correlation of Δm with τ_S for the nominal $\phi_{+-} = \phi_{SW}$ (superweak) fit. It is clear that relaxing the superweak fit causes Δm and τ_S to become much more correlated with each other, and hence allowing them to float simultaneously is the correct treatment.

A final fit to consider is fitting the difference between ϕ_{+-} and ϕ_{SW} as dynamically calculated in the fit. Here, Δm, τ_S, and $\delta\phi = \phi_{+-} - \phi_{SW}$ float. We obtain a

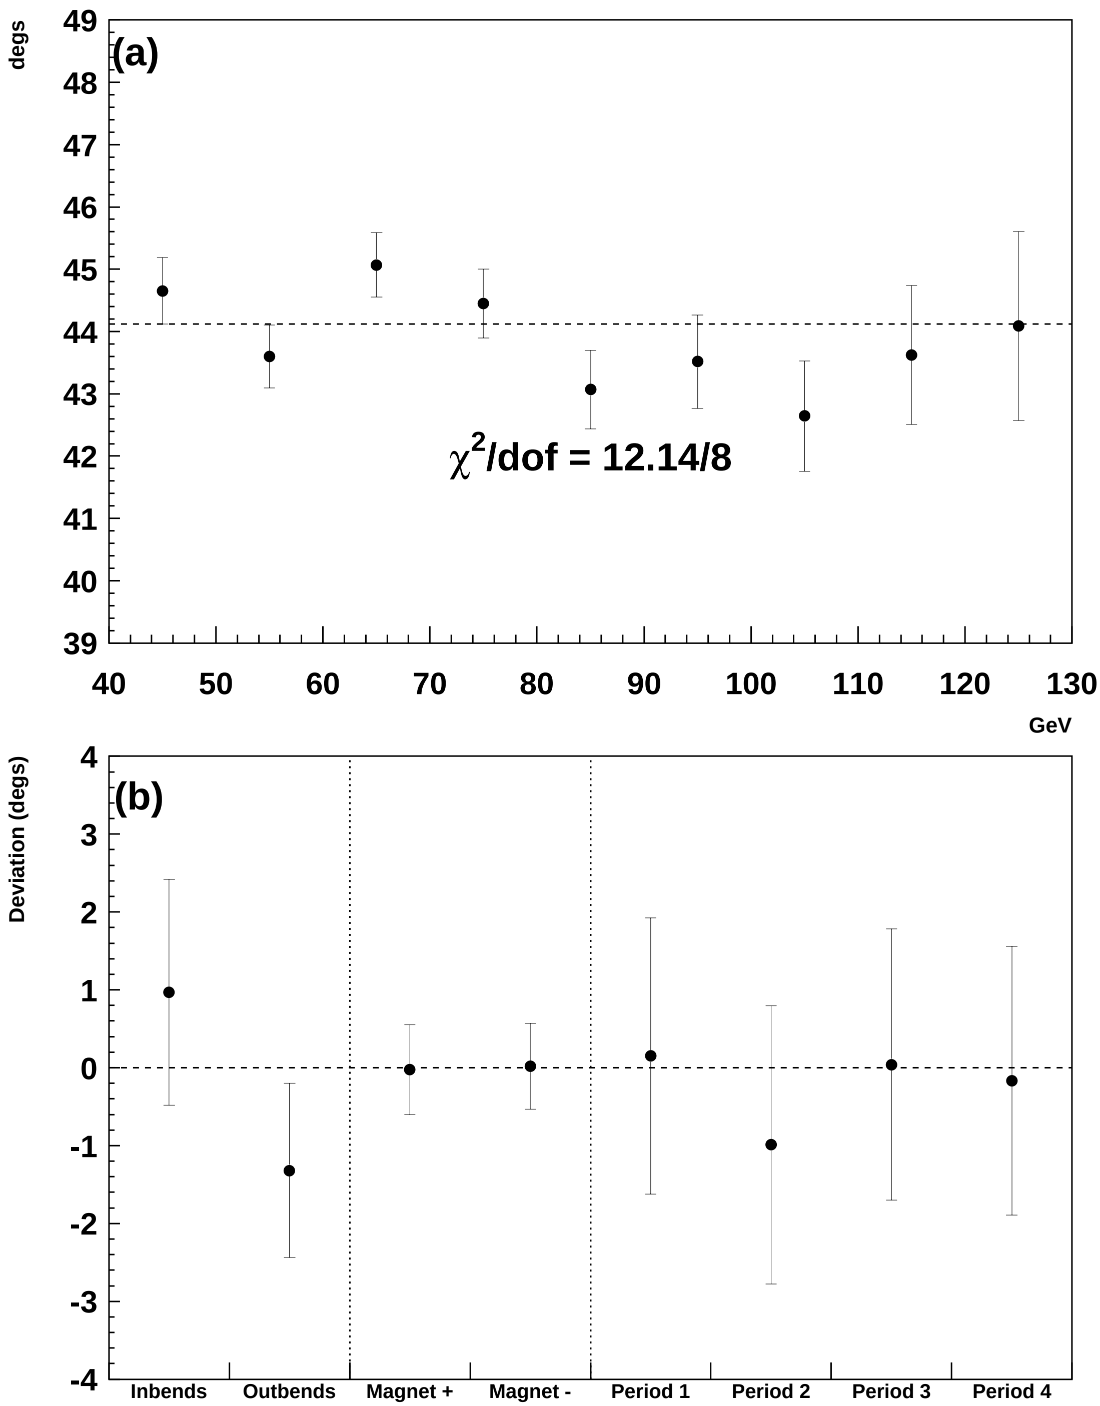

Figure 9.7: Variation of ϕ_{+-} with kaon momentum (a) and data subset (b) for the charged mode data. For Panel (b), the data sets between dotted lines are statistically independent.

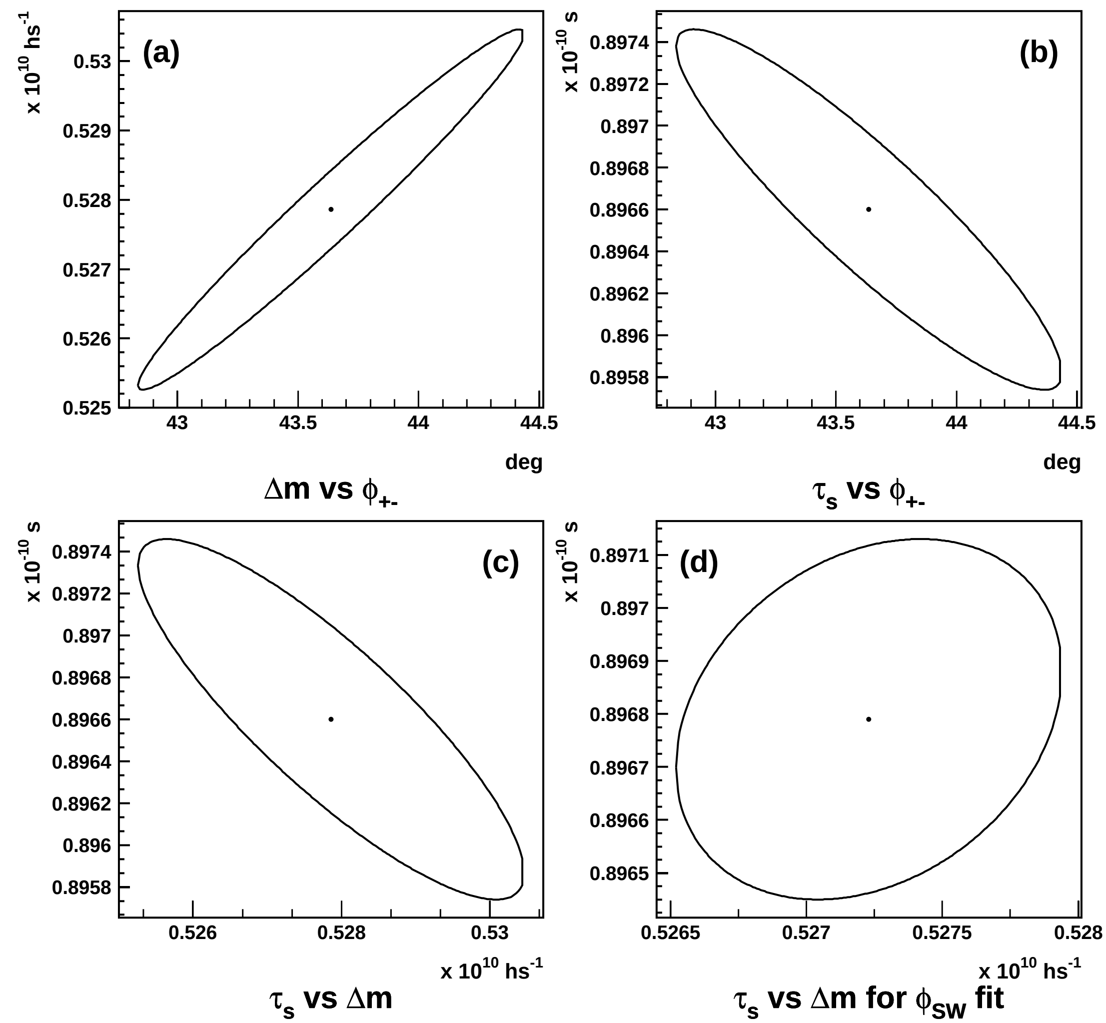

Figure 9.8: One sigma contours for Δm with ϕ_{+-} (a), τ_S with ϕ_{+-} (b), and τ_S with Δm (c) in the grand fit. For comparison, the correlation between τ_S and Δm is shown (d) for the nominal fit assuming CPT symmetry.

slightly better error estimate on the phase *difference* than on the full phase ϕ_{+-}:

$$\delta\phi = (\phi_{+-} - \phi_{SW}) = (0.61 \pm 0.62)^\circ \text{ (stat)}. \tag{9.13}$$

With the full results for Δm and τ_S, we are ready to fit $\text{Re}(\epsilon'/\epsilon)$ using our best fit numbers for Δm and τ_S.

9.4 The Fit for Re(ϵ'/ϵ)

The standard fit for $\text{Re}(\epsilon'/\epsilon)$ has 12 10-GeV energy bins for each of the four $\pi\pi$ samples, for a total of 48 bins. The kaon flux normalizations for $\pi^+\pi^-$ and $\pi^0\pi^0$ are floated in each energy bin. The regeneration parameters $|f_-(70 \text{ GeV})|$ and α are floated. Finally, $\text{Re}(\epsilon'/\epsilon)$ is floated, for a total of 27 floating parameters. The fit has 21 degrees of freedom.

We fix the values $\Delta m = 0.5262 \times 10^{10}\ \hbar\text{s}^{-1}$ and $\tau_S = 0.8964 \times 10^{-10}$ s, which is the average of our charged- and neutral-mode results for these quantities (see Section 11.1.1). The value of ϵ is fixed to the mean of the Particle Data Group (PDG) [28] values for $|\eta_{+-}|$ and $|\eta_{00}|$, 2.28×10^{-3}. We assume CPT symmetry, such that the phases ϕ_{+-} and ϕ_{00} are equal to the superweak phase ϕ_{SW}, as defined in Equation 1.40.

We are ready to look at the value of $\text{Re}(\epsilon'/\epsilon)$. For the 1997B dataset, independent of the previous published result [55], the fit determines

$$\text{Re}(\epsilon'/\epsilon) = (19.8 \pm 1.85) \times 10^{-4}, \tag{9.14}$$

where the error is statistical only. The regeneration parameters found by the fit are $|f_-(70 \text{ GeV})| = 1.2067 \pm 0.0003$ mb and $\alpha = -0.5419 \pm 0.0009$. The fact that α disagrees with the value reported in our first measurement [55] ($\alpha = -0.5903 \pm 0.0015$) is due to the new treatment of kaon-nucleon scattering. If we do not use a screening correction, our “pure power-law” value is $\alpha = -0.5875 \pm 0.0008$. The kaon flux normalization parameters vary by about 4% for charged-mode and 0.7% for neutral-mode. The variation of these parameters is shown in Figure 9.9. The normalization factors for the charged-mode parameters indicates that the prediction function over-predicts $K \rightarrow \pi^+\pi^-$ events at high energy (in both the regenerator and vacuum beams) and we apply the normalization factors to match the data. The χ^2 is 30 for 21 degrees of freedom.

The statistical error on $\mathrm{Re}(\epsilon'/\epsilon)$ is due to both the statistics of the data and of the Monte Carlo acceptance correction. We have determined from a separate fit setting the acceptance uncertainty to zero ($\sigma_P \equiv 0$) that the error from the data alone is $\sigma_N = 1.73 \times 10^{-4}$ and that $\sigma_P = 0.80 \times 10^{-4}$ of the error is due to the finite amount of Monte Carlo generated to make the acceptance correction.

Refitting the previously published 1997A sample using all the improvements discussed in this thesis gives

$$\mathrm{Re}(\epsilon'/\epsilon) = (23.2 \pm 3.1) \times 10^{-4}, \tag{9.15}$$

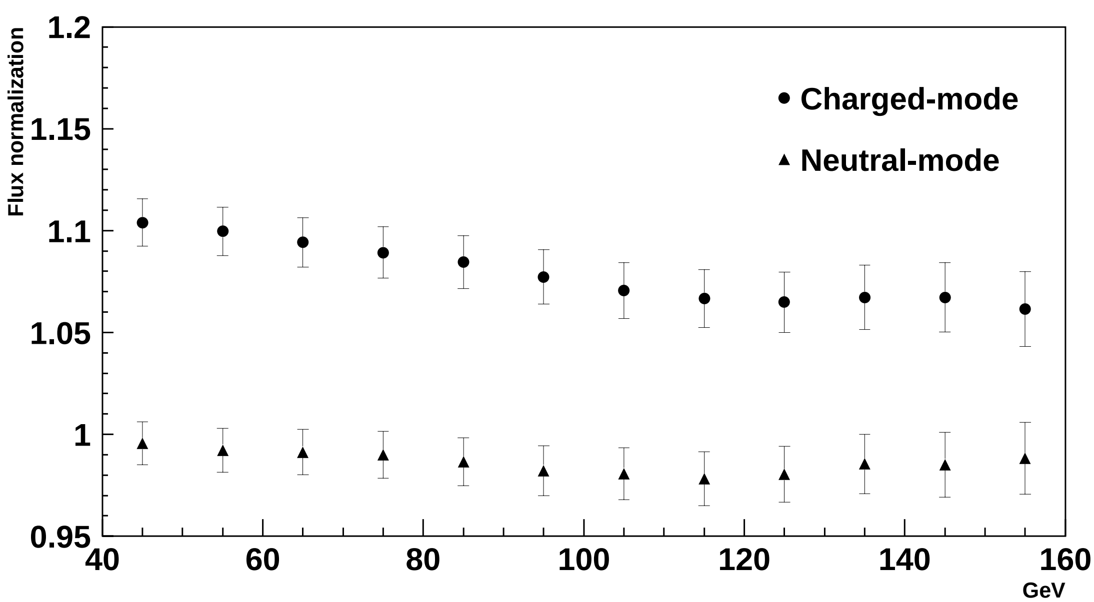

Figure 9.9: Fit flux factors. The overall normalization is not important. The error bars are those returned from the fit, and are correlated. The neutral-mode shows little variation as a function of momentum. The charged-mode shows some variation due to the new pion shower MC.

where 3.1 is the statistical error.

Interestingly, the statistical error from the data is roughly 14% larger than one would expect from straight $\sqrt{N}$ statistics. This is because the K_L component in the regenerator beam increases the uncertainty on the K_S decay amplitudes, and also because the charged and neutral data samples have somewhat different energy distributions.

9.5 Fitting Cross Checks

We have performed a number of cross-checks to the main result for $\mathrm{Re}(\epsilon'/\epsilon)$, which we will now detail. First we relate some of the studies done to check the fitter itself. Then we discuss fits done to check the assumptions used in the fitter. We will touch upon the alternate "reweighting" technique, and finish with a sample of tests that used subsets of our full dataset.

For the first result, a number of cross-checks of the internal consistency of KFIT itself were performed [2]. The basic operation of KFIT has not changed in the intervening time, and those tests remain valid. The two most important tests were to use Monte Carlo generated with $\mathrm{Re}(\epsilon'/\epsilon) = 10 \times 10^{-4}$ as data, and fit it in exactly the same way as the nominal fits. The result was consistent to within the 2×10^{-4} statistical uncertainty of the fit. The test was limited by the amount of "large $\mathrm{Re}(\epsilon'/\epsilon)$" Monte Carlo generated. To avoid this problem, we used a simplified acceptance model to calculate an exact parent distribution, and applied Poisson statistical fluctuations to generate 100 independent "fake" datasets. The change in the mean value $\mathrm{Re}(\epsilon'/\epsilon)$ from the generated value was $\Delta\mathrm{Re}(\epsilon'/\epsilon) = (0.11 \pm 0.28) \times 10^{-4}$.

As mentioned above, we do not have perfect data-MC agreement in our energy distributions, a problem that is side-stepped by using independent kaon flux factors in each 10-GeV bin. Many inefficiencies will cause an energy-dependent bias. We fit $\mathrm{Re}(\epsilon'/\epsilon)$ in 10 GeV kaon energy bins, as shown in Figure 9.10, and see consistent results. We have relaxed the power-law constraint on the regeneration amplitude ρ, fitting the value of the amplitude in each 10 GeV bin. Within each bin, the regeneration amplitude is still assumed to follow the power law with $\alpha = -0.5419$, and the regeneration phase is determined from the analyticity relationship for that power-law. This fit found $\mathrm{Re}(\epsilon'/\epsilon) = 19.66 \times 10^{-4}$, with a χ^2 of 22 for 11 degrees of freedom. The statistical error increased slightly to 1.85×10^{-4}. Furthermore, we have fit the data fixing the same normalization parameters in both charged and

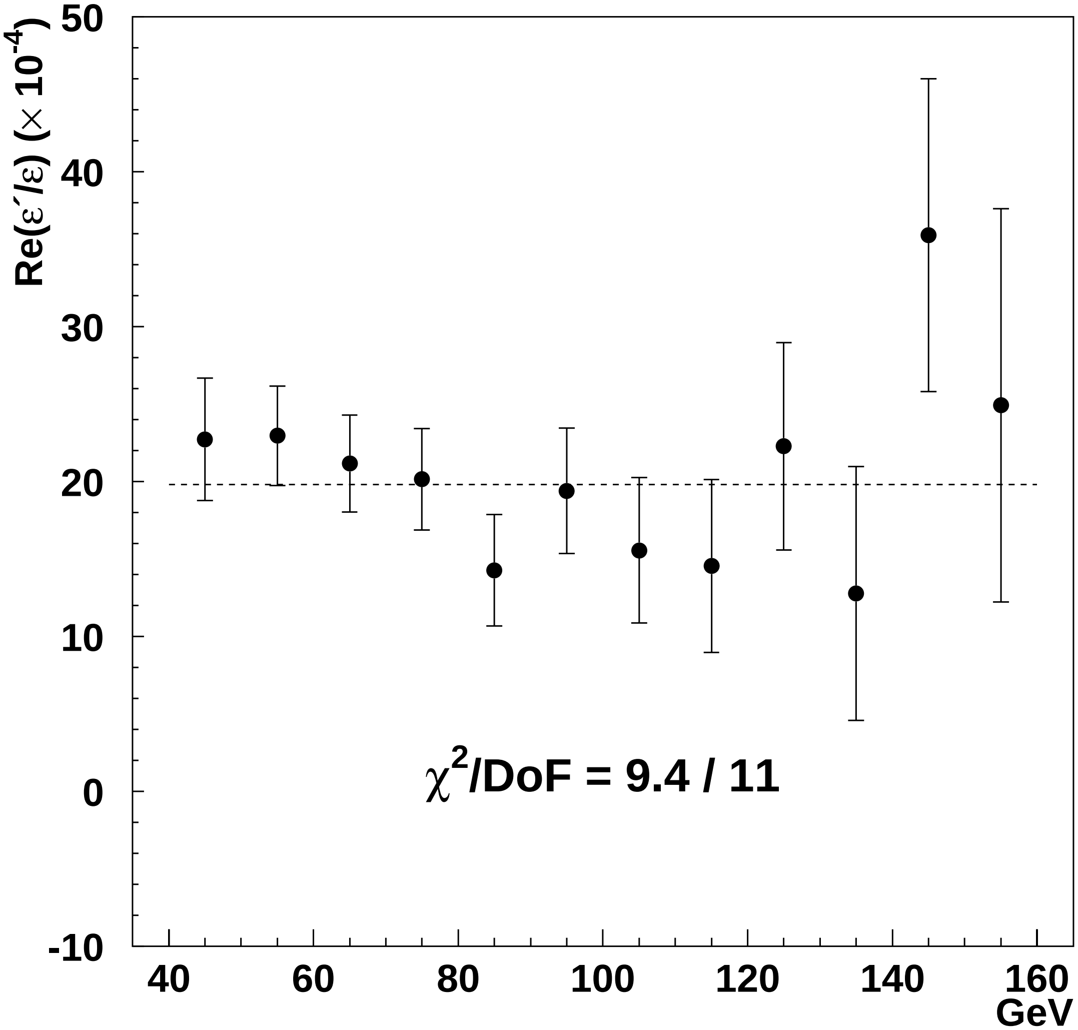

Figure 9.10: Fitted values of $\mathrm{Re}(\epsilon'/\epsilon)$ in 10 GeV kaon energy bins.

neutral modes. This fit has a very large $\chi^2 = 228$, indicating our original treatement is favored by the data. The shift in the value is $\Delta\mathrm{Re}(\epsilon'/\epsilon) = (1.8 \pm 1.3) \times 10^{-4}$.

Fitting the data using the "geometry-only" Monte Carlo acceptance correction (see Section 8.3) leads to a change of $\Delta\mathrm{Re}(\epsilon'/\epsilon) = +12.7 \times 10^{-4}$. However, we have seen that the global z-acceptance slope is $10 \times 10^{-4}\ \mathrm{m}^{-1}$, and hence predicts most of the bias that we see in the fit.

KTeV has also developed a reweighting technique for the KTeV dataset [91]. The KTeV data are ideal for the reweighting technique, since the use of the regenerator to generate K_S ensures that the two beams are left-right symmetric with respect to the detector, and that the acceptances for the two beams at a given (E_K, z) will match. This is in contrast to the NA48 technique, where the K_S and K_L beams originate from different targets, in turn requiring a small acceptance correction. The reweighting function is the ratio of the predictions for the kaon wavefunction in the regenerator beam over the same for the vacuum beam, determined in a manner extremely similar to the methods used in both KTEVMC and KFIT. We are reweighting to the *predicted* shape of the two distributions; whether the two beams match after reweighting is an indication of the similarity of the local acceptances between the two beams. This technique reweights events in the vacuum beam to have the same distribution in (E_K, z) as the regenerator beam, effectively making downstream K_L events "count" for less, meaning the downside of the reweighting is the loss of the statistical power of the CP-violating $K_L \to \pi\pi$ decays. The effect of reweighting is shown in Figure 9.11.

The reweighting technique has a larger statistical uncertainty, as well as different systematic issues from the main analysis, since the focus is not to match the Monte Carlo to the data, but to understand beam-to-beam effects. The result of the reweighting fit to the 1997B dataset is

$$\mathrm{Re}(\epsilon'/\epsilon) = (21.7 \pm 2.9) \times 10^{-4}, \tag{9.16}$$

where only the statistical error is shown. The difference between the two analyses is $\Delta\mathrm{Re}(\epsilon'/\epsilon) = (1.5 \pm 3.0) \times 10^{-4}$, where the statistical uncertainty on the difference

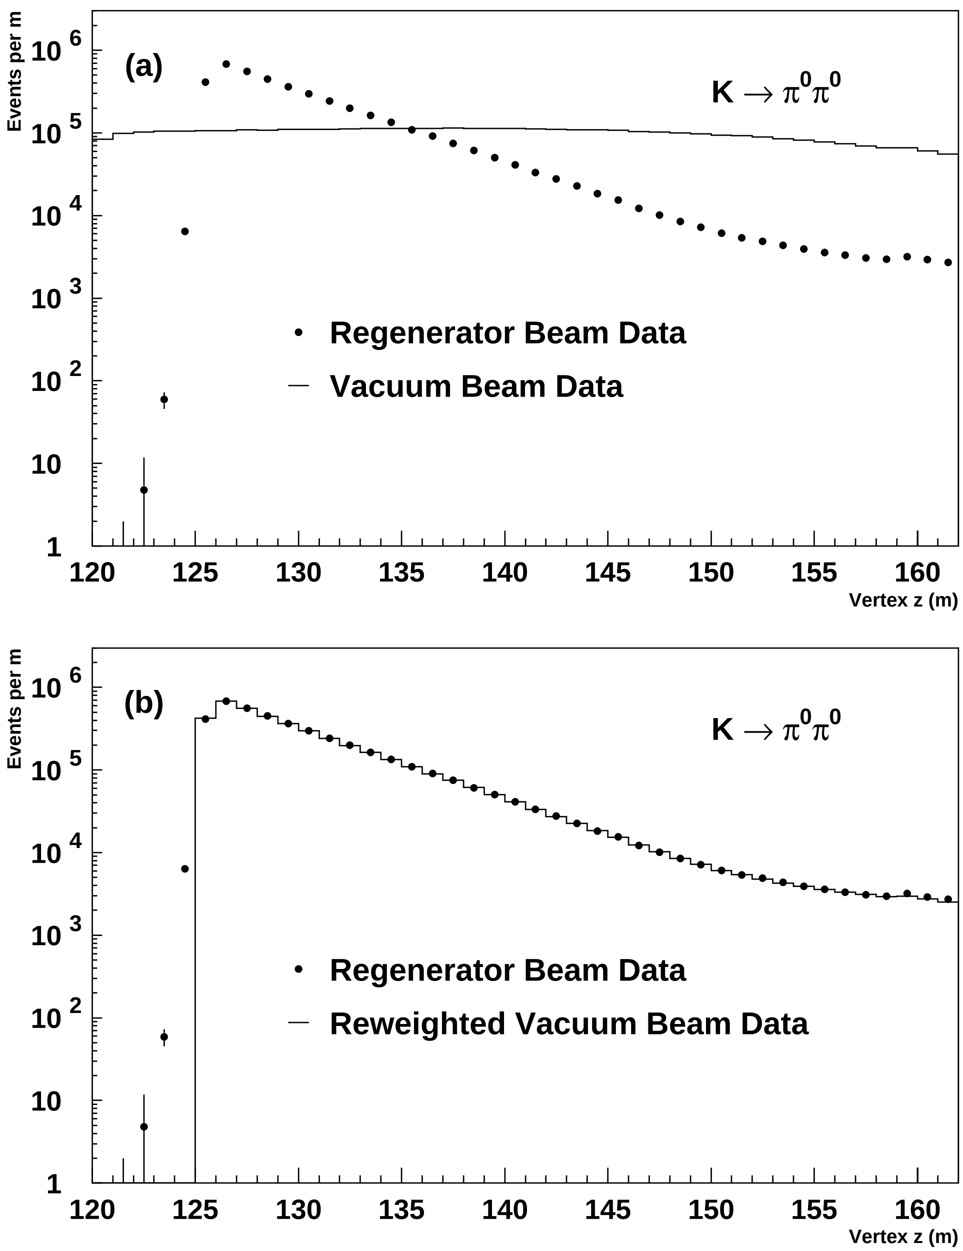

Figure 9.11: The effect of reweighting the vacuum beam to the regenerator beam. Panel (a) shows the overlay of the regenerator beam data on the vacuum beam for $K \to \pi^0\pi^0$. Panel (b) shows the same after the vacuum beam data has been reweighted by the prediction function. The match between the two shows the local acceptance is similar in both beams.

has been estimated using a large number of "fake datasets." In addition, there are uncorrelated components of the systematic uncertainty, all of which indicate that the reweighting analysis confirms the main analysis result.

A further set of cross-checks is to break the data into subsets of the full dataset. In Figure 9.12, we show the comparison of these sub-sets to our nominal result. The data have been broken into five time periods, into left and right beam, and into magnet polarity. We expect no change due to the variations, and see no statistically significant differences.

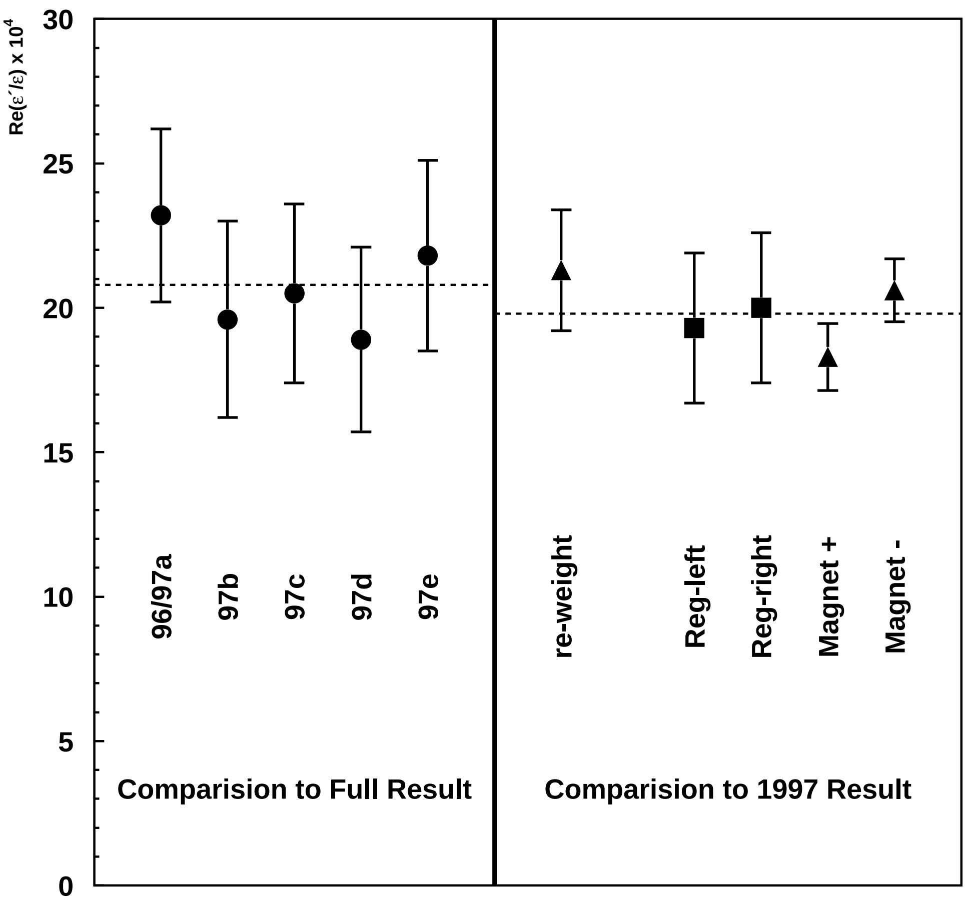

Figure 9.12: Cross checks of sub-sets of the $\mathrm{Re}(\epsilon'/\epsilon)$ data.

9.6 Summary

We have extracted Re(ϵ'/ϵ) from the data, and have cross-checked the result with a number of alternate techniques. We now change our focus to evaluating the systematic uncertainty.

CHAPTER 10
SYSTEMATIC UNCERTAINTIES

The hard work involved in measuring $\mathrm{Re}(\epsilon'/\epsilon)$ is the effort to minimize and understand the systematic uncertainties of the measurement. The focus of the effort is any effect that may bias K_S relative to K_L. The effort begins with the conceptual design of the experiment; the use of two beams originating from the same target to simultaneously collect K_S and K_L decays, a cleanly defined beamline, a high-precision detector and high-efficiency triggers. The reconstruction software has high efficiency and is robust, and the analysis software is tuned to affect the K_S and K_L samples in the same manner. However, significant differences result from the different decay distributions for K_L and K_S. A great deal of effort is invested in the Monte Carlo simulation to make sure it accurately reproduces the detector response and acceptance.

Even with all this effort, we must consider a variety of possible remaining biases, and work to constrain them, generally based on studies of the data. The size of the uncertainty of some of these effects is determined by the amount of suitable data we have to study them. Other cases are not statistically limited, and after a number of years, we have decided not to invest further effort in understanding these remaining effects. In general, we would like to reduce the size of systematic uncertainty to be of order the size of the statistical uncertainty on $\mathrm{Re}(\epsilon'/\epsilon)$, 1.50×10^{-4}.

In this chapter we will describe the studies involved in determining our systematic uncertainty, and quote the uncertainty for our measurements of the kaon sector parameters and $\mathrm{Re}(\epsilon'/\epsilon)$. For these effects, we quote a symmetric uncertainty around our nominal value, even if the study in question indicates a bias. To use a continuous limit on quoting systematic uncertainties between the two statistical extremes listed above, we quote a systematic uncertainty based on the following: let $b \pm \Delta b$ be the measured bias in some parameter. We assign a symmetric systematic uncertainty a such that $\sim 68\%$ of the area of a Gaussian with $\mu = b$ and $\sigma = \Delta b$ is included in the range $[-a, a]$:

$$\frac{1}{\sqrt{2\pi}\sigma} \int_{-a}^{a} \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right] dx = 0.683. \tag{10.1}$$

If the measured bias is zero with some error $(0 \pm \Delta b)$, the quoted uncertainty will be the error: $a = \Delta b$. For a case in which the measured bias is very large $(b \gg \Delta b)$, this procedure yields an error equal to the bias plus $\sim 0.48\sigma$: $a = b + 0.48 \cdot \Delta b$.

This chapter describes the studies used to determine the systematic uncertainties on the kaon sector parameters Δm, τ_S, and ϕ_{+-}, and on $\mathrm{Re}(\epsilon'/\epsilon)$. We present the components of the charged-mode uncertainty, and remaining uncertainties associated with KFIT, kaon sector parameters and the behavior of the regenerator. We refer the reader to Reference [1] for the corresponding work associated with the neutral-mode uncertainty, although we quote the results of those studies here. Although the systematic studies are divided into kaon sector parameters and $\mathrm{Re}(\epsilon'/\epsilon)$, it should be noted that we performed all the systematic studies detailed below for all the fits presented, but only detail the dominant contributions to each result.

Since we use our values for Δm and τ_S in determining $\mathrm{Re}(\epsilon'/\epsilon)$, we present the estimation of the systematic uncertainties on these parameters first.

10.1 Systematic Uncertainties on Δm and τ_S

We perform several systematic tests on our measurement of Δm and τ_S. For the systematic tests, we fix ϕ_{+-} to be the superweak phase ϕ_{SW} as dynamically calculated for that particular fit.

10.1.1 Analysis Cuts, Backgrounds, and Resolutions

We determine the systematic uncertainty due to analysis cuts by measuring Δm and τ_S as we vary each analysis cut. Background normalizations are determined separately for each cut value, in the same manner as was done for $\mathrm{Re}(\epsilon'/\epsilon)$. The results of the variations of Δm and τ_S are presented in Tables 10.1 and 10.2, which list the variations of the cuts that we considered. Most of the cut variation studies show no systematic effect outside the statistical error on the difference.

We also consider the effect of varying the background level subtracted from the data. All the backgrounds are varied within $\pm 10\%$ of their respective levels. Also, we test an alternative technique of floating the various backgrounds without making the enhancement cuts discussed in Section 6.7.2. The results of varying the background contribution are shown in Table 10.2.

Lastly, we consider the effects of resolution and the high-SOD rate. We use correlated sets of Monte Carlo to evaluate the difference between our nominal set and sets with resolutions altered. The effects of resolutions and the high-SOD rate are shown in Table 10.3.

Because we see no significant changes with analysis cuts and the resolutions, we consider only a small number of the fit variations as systematics. We use the following analysis variations as systematic uncertainties for Δm, in units of $10^6\ \hbar\text{s}^{-1}$ (in which $\Delta m = 5267$):

- Track separation $x > 5$: $\Delta(\Delta m) = -3.0 \pm 3.6$

Table 10.1: The cut variations considered for systematic study of Δm and τ_S in $K \to \pi^+\pi^-$ decay modes.

Variable	Change of Variable	$\Delta(\Delta m)$ ($\times 10^6\ \hbar\text{s}^{-1}$)	$\Delta(\tau_S)$ ($\times 10^{-14}$ s)
Offmag χ^2	$\chi^2 < 30$	0.7 ± 1.1	0.3 ± 0.5
Vertex χ^2	$\chi^2 < 30$	0.0 ± 0.4	0.1 ± 0.2
Trk sep CsI	$x_{\text{sep}} > 0.06$ m	-2.1 ± 1.5	0.2 ± 0.8
Track sep	$x \geq 0$	-3.7 ± 3.3	-3.1 ± 1.3
Track sep	$x \geq 5$	3.0 ± 3.6	3.1 ± 1.4
Min p	$p_{\text{min}} > 9.0$ GeV/c	-0.4 ± 1.0	-1.0 ± 0.4
Vtx z	126 m$< z <$158 m	0.1 ± 0.0	-0.1 ± 0.0
p_T^2	$p_T^2 > 150$ MeV$^2/c^2$	-1.0 ± 0.6	0.4 ± 0.3
Λ Mass	$\lvert M_\Lambda - 1.1105\ \text{GeV}\rvert < 0.0065$ GeV	-0.2 ± 4.1	0.20 ± 1.4
E/p	$E/p < 0.875$	1.5 ± 0.6	-0.2 ± 0.3
N extra wires	NWire< 100	0.0 ± 2.3	-2.6 ± 1.1

Table 10.2: The background variations considered for systematic study of Δm and τ_S in $K \to \pi^+\pi^-$ decay modes.

Change of Background	$\Delta(\Delta m)$ ($\times 10^6\ \hbar\mathrm{s}^{-1}$)	$\Delta(\tau_S)$ ($\times 10^{-14}$ s)
Float backgrounds without enhancement cuts	−0.2	0.0
+10% K_{e3} background (both beams)	−0.1	0.0
−10% K_{e3} background (both beams)	0.0	0.0
+10% $K_{\mu3}$ background (both beams)	−0.1	0.1
−10% $K_{\mu3}$ background (both beams)	0.0	0.0
+10% collimator background (both beams)	0.0	0.0
−10% collimator background (both beams)	0.0	0.0
+10% Regenerator scatter background (reg. beam)	0.9	0.2
−10% Regenerator scatter background (reg. beam)	−0.0	−0.1
+10% all backgrounds (both beams)	0.8	0.2
−10% all backgrounds (both beams)	−0.9	−0.2

Table 10.3: The variation of Δm and τ_S due to changes in resolutions, using correlated MC sets for the charged mode data.

Change of DC property	$\Delta(\Delta m)$ ($\times 10^6\ \hbar\mathrm{s}^{-1}$)	$\Delta(\tau_S)$ ($\times 10^{-14}$ s)
High SOD 80%	−0.5	1.3
High SOD 90%	−1.3	0.6
High SOD 110%	−0.5	−0.3
High SOD 120%	0.1	−1.1
Resolution 90%	−0.5	1.3
Resolution 95%	−0.5	0.7
Resolution 105%	1.9	0.5
Resolution 110%	1.7	1.6

- $p_T^2 > 150\ \mathrm{MeV}^2/c^2$ cut variation: $\Delta(\Delta m) = -1.0 \pm 0.6$
- $E/p < 0.875$ cut variation: $\Delta(\Delta m) = -1.5 \pm 0.6$
- Track separation at CsI, $x_{\mathrm{sep}} > 6$ cm: $\Delta(\Delta m) = -2.1 \pm 1.5$
- Background estimates, (±10% all backgrounds): $\Delta(\Delta m) = \pm 0.8$
- Resolutions and High-SODs: we consider the quadrature sum of the errors for High-SOD ±10% and resolutions ±5%: $\Delta(\Delta m) = 2.3$.

For τ_S, we obtain similar results from the cut variations, where all the following are in units of 10^{-14} s:

- $p_T^2 > 200$ MeV$^2/c^2$ cut variation: $\Delta(\tau_S) = -0.4 \pm 0.2$
- Background estimates, ($\pm 10\%$ all backgrounds): $\Delta(\tau_S) = \pm 0.2$
- Resolutions and High-SODS: we consider the quadrature sum of the errors for High-SOD $\pm 10\%$ and resolutions $\pm 5\%$: $\Delta(\tau_S) = 0.9$.

We assign systematic uncertainties of

$$\begin{aligned} \sigma_{\text{cuts, bkgd, res}}(\Delta m) &= 4.5 \times 10^6\ \hbar \text{s}^{-1}, \\ \sigma_{\text{cuts, bkgd, res}}(\tau_S) &= 1.2 \times 10^{-14}\ \text{s}, \end{aligned} \tag{10.2}$$

due to analysis cut variation, backgrounds and resolutions in the charged-mode data.

10.1.2 Screening and Transmission

The systematic error from screening is determined by varying the screening calculation amongst reasonable models. Table 10.4 shows the variation of our fit values with different models of the regenerator screening correction and the transmission models.

In the case of the various screening models, we ignore the case for no screening, since it is neither well-motivated nor supported by the fit χ^2. Because the both

Table 10.4: Change of Δm and τ_S with screening model and transmission. The various screening models are discussed in Section 9.2.2.

Screening	$\Delta(\Delta m)$ ($\times 10^6\ \hbar\text{s}^{-1}$)	$\Delta(\tau_S)$ ($\times 10^{-14}$ s)	χ^2 (for 198 d.o.f.)
None	2.2	-2.6	226
Inelastic fit	-0.3	-0.9	211
Inelastic symmetric	2.2	-2.0	214
Linear transmission ($\pm\sigma$ on energy dependence)	$\mp 10.$	± 2.0	208 & 212

screening models have similar χ^2s to the nominal fit, and similar ranges of error, we take

$$\begin{aligned} \sigma_{\text{screening}}(\Delta m) &= 3 \times 10^6\ \hbar\text{s}^{-1}, \\ \sigma_{\text{screening}}(\tau_S) &= 2 \times 10^{-14}\ \text{s}. \end{aligned} \tag{10.3}$$

The biggest change in Δm from the above-considered effects is from the energy-dependence of transmission. Since we have elected to use our measurement of the transmission, we must consider the uncertainty in the determination of the energy dependence.

$$\begin{aligned} \sigma_{\text{transmission}}(\Delta m) &= 10 \times 10^6\ \hbar\text{s}^{-1}, \\ \sigma_{\text{transmission}}(\tau_S) &= 2 \times 10^{-14}\ \text{s}. \end{aligned} \tag{10.4}$$

Offsetting the ϕ_ρ analyticity assumption by 0.25° causes systematic uncertainties on Δm and τ_S:

$$\begin{aligned} \sigma_{\text{analyticity}}(\Delta m) &= 8 \times 10^6\ \hbar\text{s}^{-1}, \\ \sigma_{\text{analyticity}}(\tau_S) &= 3 \times 10^{-14}\ \text{s}. \end{aligned} \tag{10.5}$$

It has been argued [92] that the analyticity assumption is good to 0.35° for the E773 experiment, which included lower kaon energies where more of a deviation is expected.

The total systematic uncertainties for Δm and τ_S are:

$$\begin{aligned} \sigma_{\text{syst}}(\Delta m) &= 14 \times 10^6\ \hbar\text{s}^{-1}, \\ \sigma_{\text{syst}}(\tau_S) &= 4 \times 10^{-14}\ \text{s}. \end{aligned} \tag{10.6}$$

10.2 Systematic Uncertainties on ϕ_{+-}

The systematic issues for ϕ_{+-} are similar to those for Δm and τ_S. The analysis cut variations show a small effect, while we have to deal with larger effects due to the screening models in the regenerator.

10.2.1 Analysis Cuts, Backgrounds, and Resolutions

We performed systematic studies similar to those listed in Tables 10.1 and 10.2. Only a few of the studies show any statistically significant variation:

- Track Separation cut $x \geq 2$: $\Delta(\phi_{+-}) = (-0.63 \pm 0.27)^\circ$
- $p_T^2 \geq 500$ MeV$^2/c^2$: $\Delta(\phi_{+-}) = (-0.18 \pm 0.05)^\circ$
- 492 MeV$/c^2 < m_{\pi^+\pi^-} < 504$ MeV$/c^2$: $\Delta(\phi_{+-}) = (0.23 + 0.10)^\circ$.

Background variations change ϕ_{+-} by less than 0.008°, so we ignored this systematic uncertainty.

Using the same correlated MC samples described in Section 10.1.1, we have determined the systematic uncertainty due to resolution and high-SOD mismatch. We quote an uncertainty of $\sigma_{\text{res}} = 0.08^\circ$.

We quote a systematic error due to cuts, backgrounds and resolutions of

$$\sigma_{\text{cuts, bkgd, res}}\,(\phi_{+-}) = 0.83^\circ. \tag{10.7}$$

10.2.2 Screening, Transmission, and Fitting Procedure

We vary the screening and transmission assumptions for the ϕ_{+-} fits in the same manner as the charged-mode Δm and τ_S fits. The largest change comes from the screening model. Changing to the symmetrized inelastic model from the factorized inelastic model changes ϕ_{+-} by -0.75°, with an accompanying increase of the χ^2 by of 7 units. Changing the linear transmission model by $\pm 1\sigma$ causes $\Delta\phi_{+-} = 0.07^\circ$, and offsetting the analyticity assumption by 0.25° causes a commensurate 0.25° change in ϕ_{+-}.

We have found that the varying the screening models has a large effect on ϕ_{+-} in fits where τ_S is allowed to float in conjunction with ϕ_{+-}. In fits where τ_S is fixed, ϕ_{+-} varies by only 0.28° when we scan our three screening models. However, when τ_S is allowed to float, the relative error due to the screening nearly triples. The change in ϕ_{+-} is accompanied by a similar change in τ_S from the nominal value, such that most of the change in ϕ_{+-} can be calculated from Equation 9.12. However, it is probable that this systematic uncertainty is conservative, since we have allowed the maximum C-odd terms in the inelastic scattering correction as determined from factorization [89]. This systematic uncertainty can be reduced with further study of the transmission and regeneration in the KTeV regenerator. Also, fixing τ_S to a value measured without a CPT constraint, such as Reference [59], reduces the systematic uncertainty.

For the screening and transmission, we quote a systematic uncertainty based on inelastic-symmetric model and the change in the transmission slope:

$$\begin{aligned} \sigma_{\text{screening}}(\phi_{+-}) &= 0.75^\circ, \\ \sigma_{\text{trans}}(\phi_{+-}) &= 0.07^\circ, \\ \sigma_{\text{analyticity}}(\phi_{+-}) &= 0.25^\circ. \end{aligned} \tag{10.8}$$

We obtain our final systematic on ϕ_{+-} as the quadrature sum of the above:

$$\sigma_{\text{syst}}\left(\phi_{+-}\right) = 1.14^\circ. \tag{10.9}$$

10.3 Introduction to the Systematic Uncertainties on Re(ϵ'/ϵ)

In the following sections, we will detail the studies we use to set the systematic uncertainty on Re(ϵ'/ϵ). We consider uncertainties in the $K \to \pi^+\pi^-$ analysis due to trigger inefficiencies, the spectrometer alignment and calibration procedure, the analysis cuts, the background-estimation procedure, and the detector acceptance. After discussing the acceptance correction, we will revisit the acceptance uncertainty for our published result [55]. Re(ϵ'/ϵ) depends on external parameters, so we set systematic uncertainties based on these parameters. We then quote uncertainties from the neutral-mode analysis, although we do not present the studies of these effects. In the remaining sections, all the quoted systematic uncertainties are on Re(ϵ'/ϵ).

10.4 Trigger Inefficiencies

As mentioned in Section 3.7.3, we collected 14 triggers in addition to our two main triggers. These triggers have looser requirements and are prescaled. From these trigger samples, we can study small inefficiencies and biases in the main triggers, both in the Level 1 and 2 hardware trigger, and in the Level 3 software trigger. Since the trigger comes in three levels, it is appropriate to discuss the trigger inefficiency and bias at each level.

The Level 1 trigger inefficiency and bias is evaluated using K_{e3} events collected in a special trigger that only required energy deposited in the calorimeter, and specifically *none* of the first level charged trigger elements. The two-track topology is chosen such that the event should have satisfied the Level 1 trigger hodoscopes, specifically the EWUD logic. The overall inefficiency of the trigger is measured to be 0.48%, most of which was due to the EWUD logic failing at specific points where the vertical gaps between the trigger counters overlapped with the large horizontal gap between the upper and lower counter sets (see Figure 3.6). Because this large effect is due to the geometry of the counters, it is modeled by the Monte Carlo, and the data/MC ratio of the EWUD gap inefficiency is 1.08 ± 0.05. From the size of the data-MC discrepancy, we assign a systematic of 0.09×10^{-4} for this particular effect. When we consider the remaining inefficiencies, such as inefficiency from the "$2 \otimes 1$" trigger and the DC-OR sources, we obtain a systematic uncertainty of 0.17×10^{-4}.

There are a class of Level 1 failures at the 0.04% level that are not understood. This inefficiency is consistent with being the same in the vacuum and regenerator beams; however, we assign the statistical error on this check (0.28×10^{-4}) to be conservative.

We studied the Level 2 trigger inefficiency by using another pre-scaled trigger that required the nominal Level 1 trigger elements, but no Level 2 information. The overall inefficiency at Level 2 was 0.08%, due mostly to a glitch in the signal path to the hit-counting system. This glitch is expected to affect both beams equally and not cause a bias. The Banana/Kumquat modules had to miss two hit-pairs within a chamber y-view in order for an event to fail the trigger; this was extremely rare. However, events with a particular topology (both tracks either above or below the center of Chamber 4) could fail the YTF requirement if only one hit was missed by the hit-counting system. We determined that this topology is more common for vacuum beam events than for regenerator beam events. We measure the bias directly from the data to be 0.07×10^{-4}.

Finally, we studied the "B1-Random" sample of events to set the systematic uncertainty on the Level 3 trigger. The B1-Random sample is collected with the same Level 1 and 2 trigger as our nominal sample, but is not subjected to Level 3 triggering. The systematic is 0.54×10^{-4}.

10.5 Energy Scale

The absolute energy scale in the charged-mode analysis is set by fixing the analysis magnet's momentum kick, which in turn is set by the known kaon mass. Although this energy scale does not affect reconstruction at the level that it does in the neutral mode, there is a slight effect in the fit because we assume a power-law for the regeneration amplitude, which slightly ties different energy bins together.

The kaon mass is known to 0.031 MeV [28] or fractionally to 6.2×10^{-5}, and our data-MC and beam-to-beam mass distributions agree to within that accuracy. Based on that uncertainty in the absolute energy scale, we assign a systematic of 0.16×10^{-4}.

10.6 Spectrometer Alignment and Calibration

We calibrated and aligned the charged-mode spectrometer to obtain the most precise determination of track positions and momenta. There is a limit of accuracy to which we determined the calibration, and a mis-calibration at this level may add a bias to our determination of $\mathrm{Re}(\epsilon'/\epsilon)$.

To estimate the bias in $\mathrm{Re}(\epsilon'/\epsilon)$ from the calibration procedure, we first estimate the level of the miscalibration and misalignment from the data. Second, the drift chamber information in the Monte Carlo is corrupted by the amount measured from the data. Then the Monte Carlo was reanalyzed in the same manner as the data to ensure that the correct offsets and rotations could be extracted from the Monte Carlo. Finally, the $K \to \pi^+\pi^-$ signal Monte Carlo was analyzed with the full charged mode analysis, with the corruption applied to the drift chamber parameters. The reanalyzed Monte Carlo was then compared to Monte Carlo analysed with the correct drift chamber constants to determine a systematic bias on $\mathrm{Re}(\epsilon'/\epsilon)$.

10.6.1 Miscalibration and Misalignment in the Data

A miscalibration in either the $t(0)$s or $x(t)$s manifests itself in a misreconstruction of the drift time of hits. One can see this effect in either reconstructed SODs or data-MC illumination overlays. The two sigma limit on the deviation from the mean of the SOD is ± 50 μm. With an average drift velocity of 50 μm/ns, this 100 μm uncertainty corresponds to a ± 2 TDC count uncertainty. From the $t(0)$ calibration, it is possible to place a 1 TDC count error on wire timing offsets.

Muon runs give some insight into the size of the uncertainty of the offsets of the chambers. The standard offsets and rotations of chambers 2 and 3 are determined from the residuals of the track$-$hit as plotted versus the orthogonal coordinate. The residuals plotted versus the same coordinate (such as Δx versus x or Δy versus y) should also give the chamber offset. Any deviation from zero after the calibration procedure is a measure of the uncertainty in the offset determination. The variation from one side of the chamber to the other in these types of plots can limit the uncertainty in the chamber offsets to the order 20 μm.

For rotations, two general classes of rotations have been investigated. We have revisited the corkscrew rotation in more discrete time periods than the drift chamber calibration to look for any remaining rotations. And we have looked for a residual rotation between the first two and last two chambers, which can be determined by looking at the separation of the projected upstream and downstream track segments at the magnet bend plane versus the orthogonal coordinate. Even with a perfectly aligned drift chamber system, this latter rotation can be a real effect, because the non-zero B_z component of the magnetic field can cause an apparent rotation between upstream and downstream track segments. Hence this latter rotation is an overestimate of the possible systematic uncertainty in the chamber rotations. We conservatively place a limit of 50 μrads on the uncertainty of rotations in the drift chamber system.

The drift chamber position offsets were fixed such that the target reconstructed at the surveyed positions of $x = 0.0$ m and $y = -300$ μm. The target alignment is most critical for determining the p_T^2 of the decay. For a 70 GeV kaon with no intrinsic p_T^2, one would need to move the target by 2.8 cm to see a change in p_T^2 of 250 MeV$^2/c^2$. The target position has an uncertainty of 100 μm in x and 150 μm in y.

There are systematic effects in the CsI cluster reconstruction that bias the position in certain regions of the calorimeter affecting the track-cluster separation by more than $\sim$ 150 μm; that is the maximum bias that was used. For example, Figure 5.11 shows the track-cluster separation in x as a function of y. There is a systematic dip near $y = 0$ m, of order 200 μm.

The transverse magnet kick is determined by fitting the $K \to \pi^+\pi^-$ mass peak, and fixing that to the measured kaon mass. The uncertainty in the reconstructed kaon mass is 0.033 MeV. The resulting uncertainty on the transverse kick is then 0.050 MeV.

The z positions of the chambers are additional possible sources of systematic error, and are very difficult to measure in the data. By reconstructing the sharp edge of the z distribution in the regenerator beam and comparing to Monte Carlo, we can determine that we understand the shape at the edge of the regenerator to within 1 mm, which in turn corresponds to a $\sim$ 200 μm uncertainty on the z-separation between chambers 1 and 2. The uncertainties in drift chamber alignment numbers that can be estimated from the data are listed in Table 10.5.

Table 10.5: Estimates from data of the uncertainties in given DC calibration and alignment parameters.

Parameter	Estimate of uncertainty	Method
Hit Times	< 2 counts	Mean SODs
Offsets	10 μm	Muon runs
Offsets	20 μm	x_{offmag} and $m_{\pi^+\pi^-}$ vs Δp
Rotations	17 μrads	Corkscrew method
Rotations	50 μrads	Offmag vs orthogonal coord at magnet
Target x	100 μm	Total momentum
Target y	150 μm	Total momentum
CsI position	150 μm	Track-cluster position
Magnet kick	0.050 MeV	Reconstructed mass
z positions	< 200 μm	Regenerator edge

10.6.2 Monte Carlo Reproduction of Misalignment

The next step in determining the bias from the Monte Carlo is to ensure that the Monte Carlo reproduces offsets and rotations if they are introduced. The concept here is to introduce standard offsets and rotations to the chamber positions from which the Monte Carlo was generated, and run the same analysis as above, to ensure that the correct offsets and rotations can be extracted.

In general, the features as introduced are completely measurable by analyzing the Monte Carlo. Table 10.6 lists a subset of the changes made to the Monte Carlo, and the results extracted by the techniques described above. For all cases except the chamber offsets, the correct bias could be extracted from the Monte Carlo. The one failure will be discussed below.

For illustration, two plots derived from the Monte Carlo are shown. Figure 10.1 shows the corkscrew angle as measured between Chambers 1 and 2 in the Monte Carlo, when a corkscrew rotation of 20.47 μrads was introduced between the two Chambers. The extracted result is (21.28 ± 3.3) μrads. Figure 10.2 shows the average x_{offmag} versus y at the magnet, and vice versa, for the case where a 50 μrad rotation was introduced between the upstream and downstream chambers. Here the extracted results are (42 ± 4) μrads and (47 ± 4.2) μrads for the two cases.

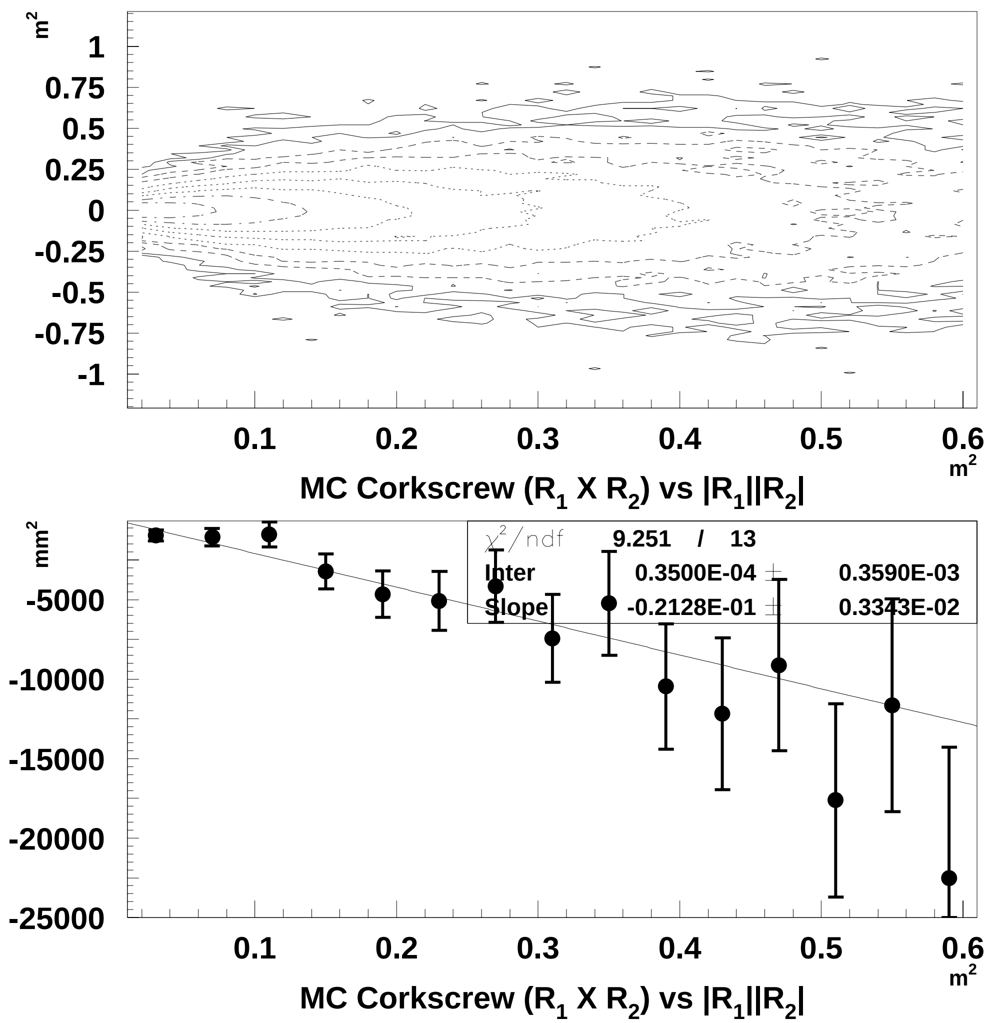

Figure 10.1: The corkscrew rotation measured in the Monte Carlo, which has had a corkscrew rotation of 20.47 μrads introduced between Chamber 1 and Chamber 2. The measured value is 21.28 ± 3.3 μrads.

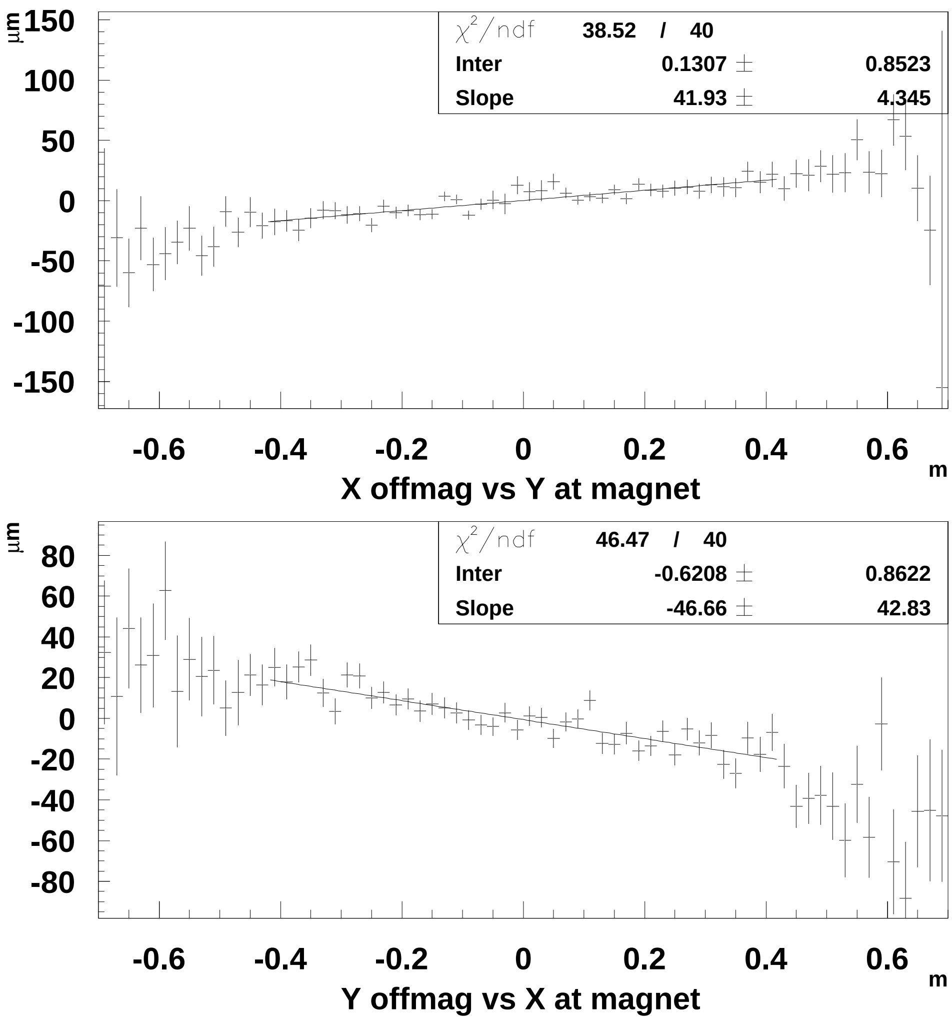

Figure 10.2: y_{offmag} versus x at the magnet (top panel), and x_{offmag} versus y at the magnet, for Monte Carlo with a rotation of 50 μrad introduced between the upstream and downstream chambers. The measured results are 42 ± 4 μrads and 47 ± 4 μrads.

Table 10.6: Amount of misalignment/miscalibration introduced into the Monte Carlo, and the measurement of the misalignment extracted by the techniques described in the text.

Parameter	Introduced Offset	Measured Offset
Rotation in x or y	50 μrads	49.3 ± 5.7 μrads
Rotation in both x, y	50 μrads	47.7 ± 3.3 μrads
Rotation between Ch 2 & 3	50 μrads	46.5 ± 6.5 μrads
Corkscrew	3.33 μrads/m	3.46 ± 0.5 μrads/m
SOD	-2 counts $\sim$ 100 μm	$\Delta(SOD) = 93$ μm
SOD	$+2$ counts ~ -100 μm	$\Delta(SOD) = -90$ μm
Transverse Kick	0.050 MeV	0.051 ± 0.012 MeV

10.6.3 Systematic Uncertainties Measured from Monte Carlo

With an understanding that the Monte Carlo correctly introduces the DC calibration and alignment biases as seen in the data, one can then turn to estimating the bias on Re(ϵ'/ϵ) by introducing these biases on top of the constants used to generate the Monte Carlo, and seeing how the number of events in each beam changes as compared to the default. This analysis was done with the constants changed by the two sigma amounts shown above, and again with the constants changed by 10 times those values. In all cases the systematic biases on Re(ϵ'/ϵ) are small. The bias on Re(ϵ'/ϵ) was calculated by the direct change in the vacuum beam–regenerator beam ratio.

$$\Delta\mathrm{Re}(\epsilon'/\epsilon) = \frac{1}{6}\frac{\Delta\left|\eta_{+-}\right|^2}{\left|\eta_{+-}\right|^2} \qquad (10.10)$$

The fact that the events are the same and just migrate near cuts gives this method its statistical power. Table 10.7 lists the estimated systematic bias for the misalignments and miscalibrations introduced to the Monte Carlo. The biases are calculated for values of the parameters measured from the data, as presented above.

The total bias is estimated to be less than 0.18×10^{-4}. This small systematic error is understandable; to cause a bais, the calibration would need to either favor

Table 10.7: Bias as measured from Monte Carlo $K \rightarrow \pi^+\pi^-$ events, for sample with offsets introduced of order the size measured from the data.

Action	Bias on $\mathrm{Re}(\epsilon'/\epsilon)$ ($\times 10^{-4}$)
Offset CH1X by 20 μm	0.006
Offset CH1Y by 20 μm	-0.047
Rotate CH1X by 50 μrads	0.075
Rotate CH2X by 50 μrads	0.055
Rotate CH1Y by 50 μrads	0.072
Rotate CH2Y by 50 μrads	0.058
Corkscrew rotate by 3.33×10^{-6} rads/m	-0.030
Move Target by 300 μm in x and y	0.036
Move CsI by 500 μm in x and y	-0.004
Offset CH2X and CH3X by 20 μm	-0.03
Offset CH2Y and CH3Y by 20 μm	-0.004
Rotate CH3 & 4 by 50 μrads wrt to CH1 & 2	-0.022
Shift SODs high by ~ 100 μm	0.140
Shift SODs low by ~ 100 μm	-0.090
Offset CH2X by 20 μm, CH3X by -40 μm	-0.051
Offset ptkick by 0.050 MeV	-0.032
Rotate CH1X & CH1Y by 50 μrads	0.089
Rotate CH3X & CH3Y by 50 μrads	0.118
Shift SODs by ~ 100 μm for beam region	-0.052
Move CH1 z position by 200 μm	0.040

one beam over the other, which is difficult to do with global positions and rotations, or to push events systematically outside of cuts.

10.7 Analysis Cuts

We varied the analysis cuts discussed in Chapter 6 to look for potential biases in the reconstruction or in the data-Monte Carlo agreement. The only two that give any cause for concern are the p_T^2 cut and the track separation cut, which we detail below.

Doubling of the p_T^2 cut causes a variation of $(-0.23 \pm 0.05) \times 10^{-4}$. The two background subtracted p_T^2 plots are shown in Figure 10.3. As can be seen from the figure, our Monte Carlo does not model the tails of these distributions down three orders of magnitude from the signal region, and furthermore there is a slight difference between the two beams. There is no further statistically significant variation if we increase the cut beyond 500 MeV$^2/c^2$. We assign a systematic of 0.25×10^{-4}. This systematic is slightly reduced from the value used in the first result [55] because improved the modeling of scattering in the absorbers that contributes to the tails of these distributions. We have considered a possible variation of both the p_T^2 cut (to 500 MeV$^2/c^2$) and the mass cut (to 470 MeV/c^2 $< m_{\pi^+\pi^-} <$ 526 MeV/c^2). This "widening of the box" allows for signal events that are mis-reconstructed, which tend to move events in both mass and p_T^2. We have seen that data events are misreconstructed more often that MC events. The shift due to opening the signal region leads to a change of $(-0.40 \pm 0.08) \times 10^{-4}$; however, there is an accompanying change in the final acceptance systematic (Section 10.9.5) which covers this change.

A large concern in the Monte Carlo is the modeling of tracks that approach each other. We have included an aperture cut at the track separation value of 3 cells to exclude this region. We are concerned whether the MC predicts the distribution of close tracks in each beam. We address this systematic by varying the track separation cut. The effect on Re(ϵ'/ϵ) of varying this cut is shown in Figure 10.4. The MC does not mock up close tracks, and hence we expect a variation in Re(ϵ'/ϵ) below the cut of 3. However, the final acceptance systematic uncertainty when we make no track separation cut is much larger because the final vertex-z distributions do not agree when we do not make the track separation cut. The uncertainty that we would have to quote on the acceptance would cover the difference in Re(ϵ'/ϵ) that we see as we lower the track separation cut from 3 to 0.

In addition to the acceptance, the MC may mispredict the distribution of close tracks if the invariant mass distribution differs between the data and MC, which can occur if the relative amount of $K \to \pi^+\pi^-\gamma$ is incorrect in the MC. In Figure 10.5, we show the agreement between the data and MC in the mass distribution. It is difficult to disentangle the relative contributions of resolutions and the $K \to \pi^+\pi^-\gamma$, but based on the ratios shown in Figure 10.5, we seem to have limited the possible error in the amount of $K \to \pi^+\pi^-\gamma$ to $\sim 20\%$.

Above the track separation cut of 3, we believe we should be insensitive to

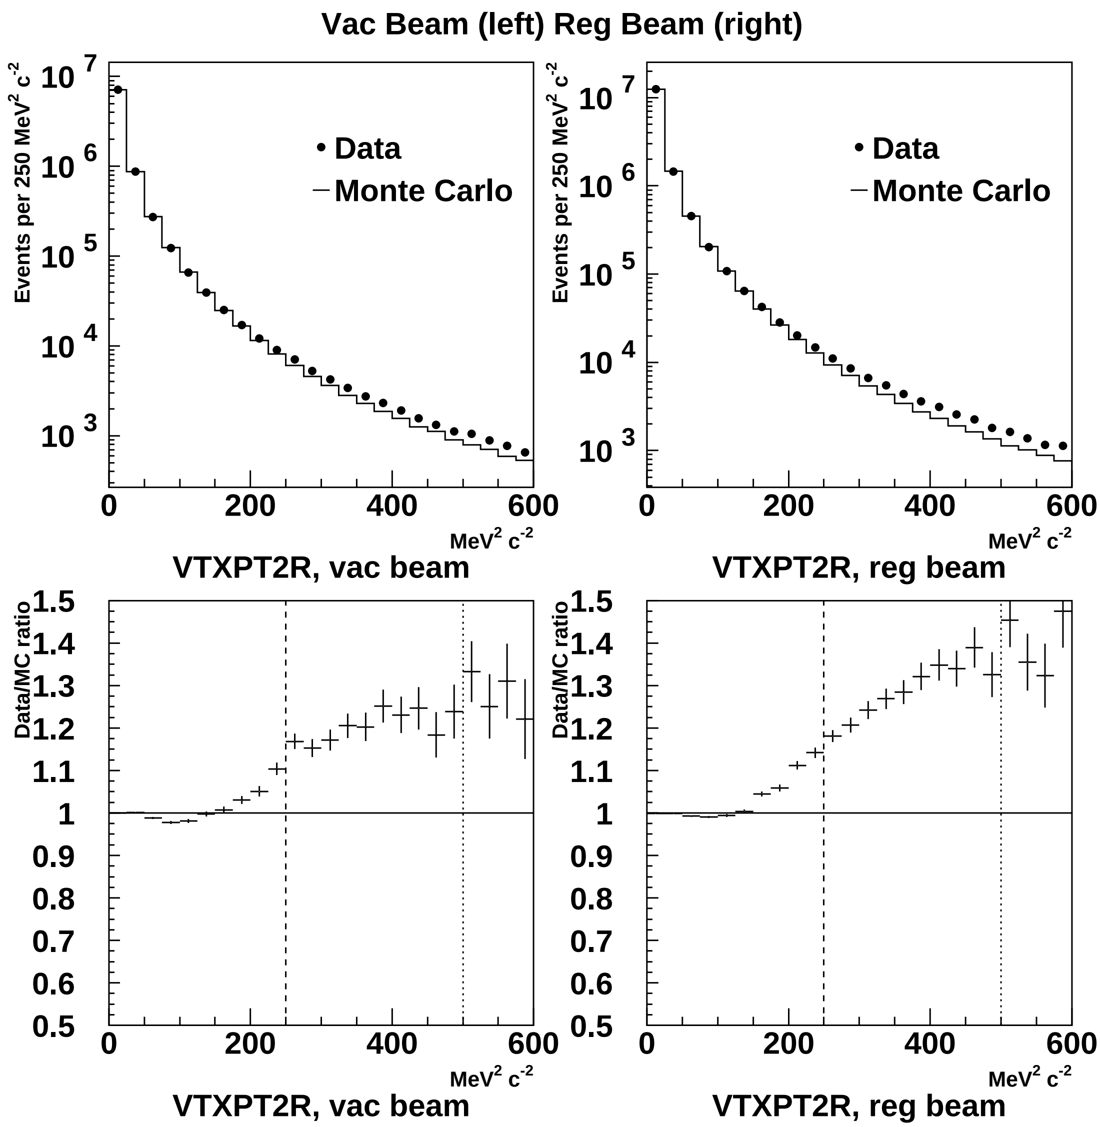

Figure 10.3: Data/MC comparisons of the background-subtracted data (dots) to the coherent MC (histogram), for the Vac beam (left) and the Reg Beam (right). The nominal cut is at 250 MeV$^2/c^2$. Doubling the cut causes a shift of $\mathrm{Re}(\epsilon'/\epsilon)$ of $\sim 0.25 \times 10^{-4}$.

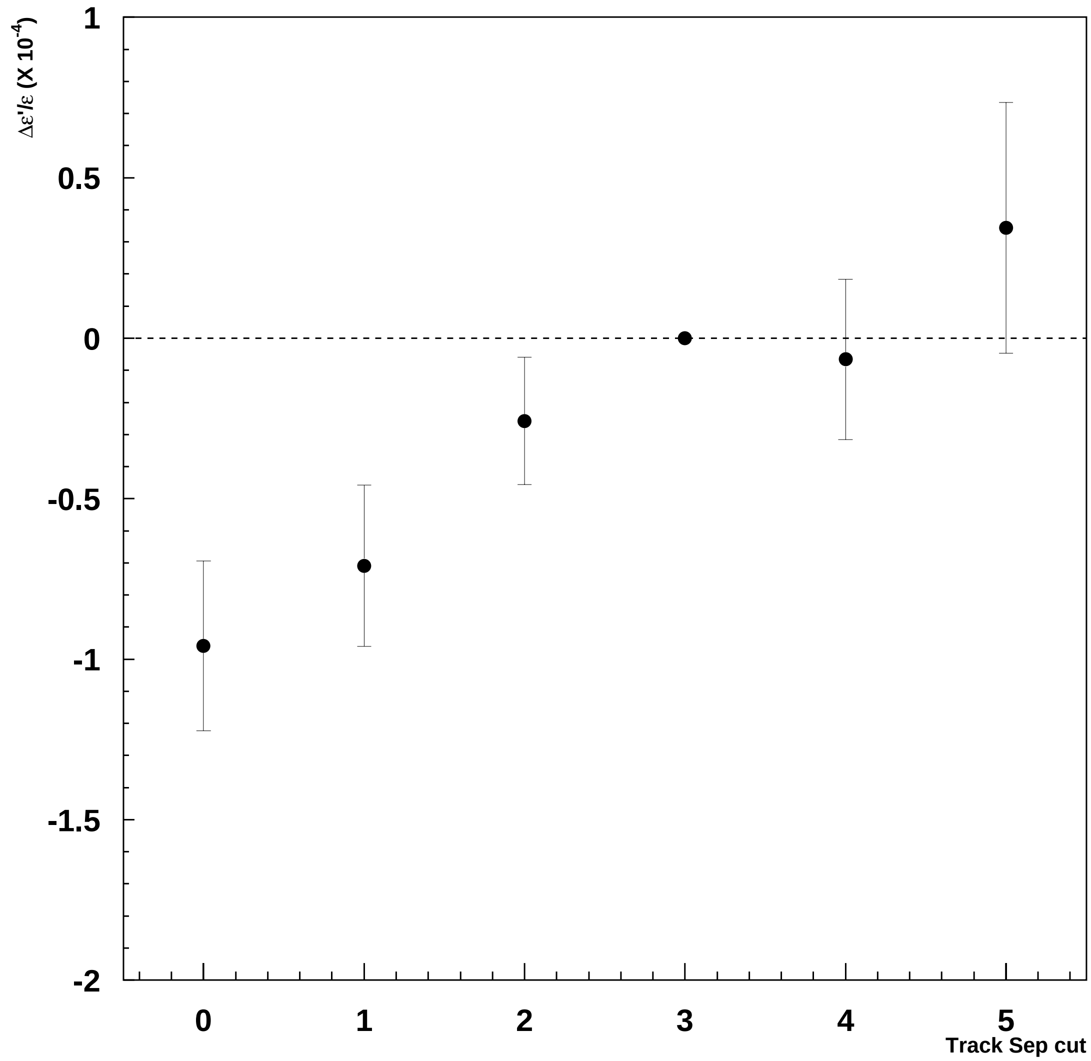

Figure 10.4: Variation of $\mathrm{Re}(\epsilon'/\epsilon)$ with wire-centered cell track separation cut. The error on a given point is the statistical error on the change between that point and the default cut at 3. The MC does not mock up close tracks below a cut of 3.

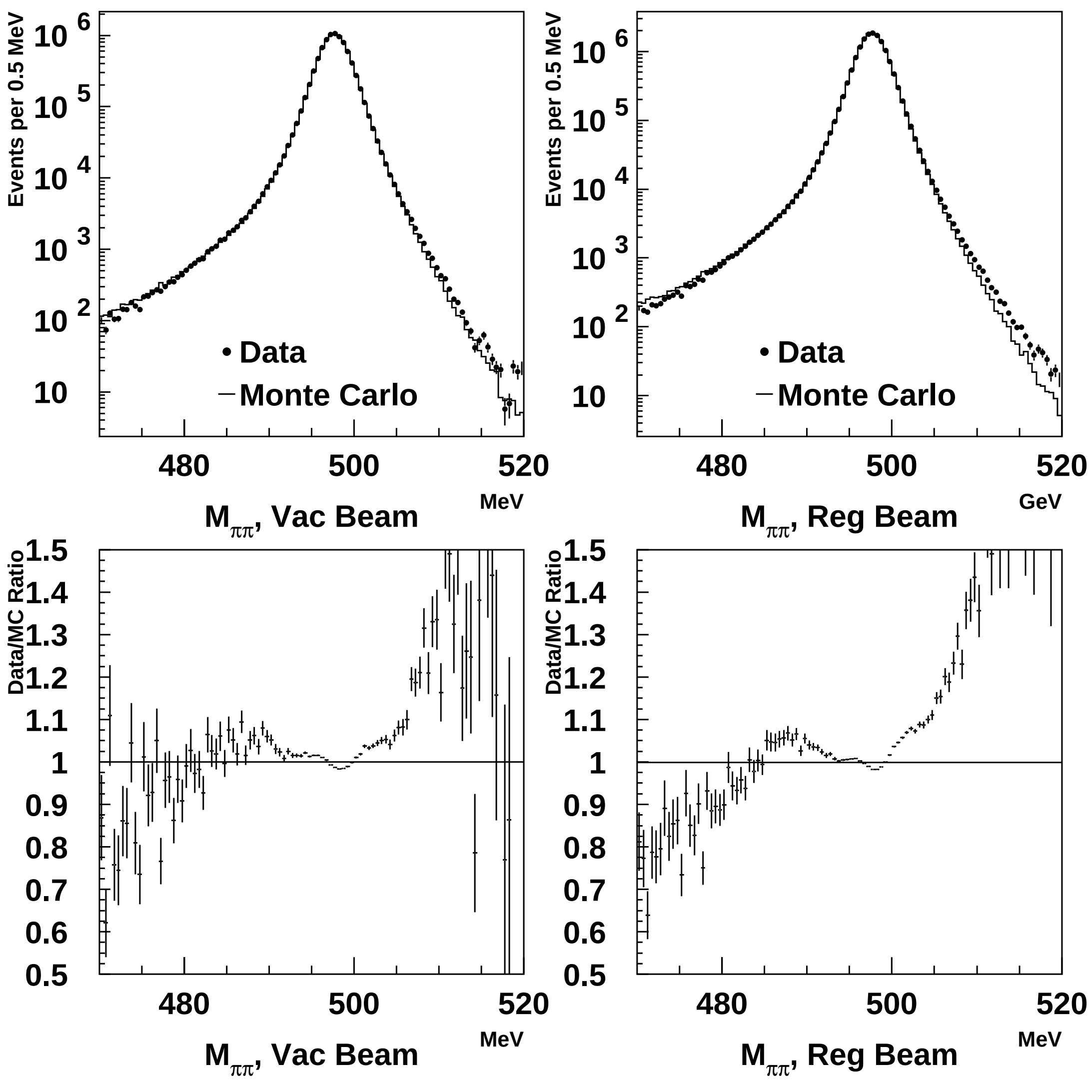

Figure 10.5: Overlay of data to MC for the background-subtracted mass distributions in the vacuum (L) and regenerator (R) beams.

inefficiencies caused by close tracks, and we see no significant sign of a bias above the cut. To be conservative, we assign a systematic of 0.26×10^{-4} based on the $(-0.06 \pm 0.25) \times 10^{-4}$ variation between the nominal cut at 3 and the cut at 4.

10.8 Backgrounds

In general, the backgrounds are well-understood (see Section 6.7). The background levels are 0.098% in the vacuum beam and 0.081% in the regenerator beam. The background rate is slightly higher in the vacuum beam than in Reference [55] due to a less-effective $K_{\mu 3}$ veto. We added an additional cut on MU3 and removed run 10059, reducing the background level from 0.14% to 0.098%.

We vary the background levels ±10% and see no variation above 0.10×10^{-4}. We cross-check our result by floating the background contributions directly, without using the enhancement cuts described in Section 6.7.2, which changes our result by 0.07×10^{-4}. We have tested the background subtraction procedure by opening up the mass cut to 484 MeV/c^2 $< m_{\pi^+\pi^-} <$ 512 MeV/c^2, which changes the backgrounds to 0.128% (0.0822%) in the vacuum (regenerator) beam, and changes $\mathrm{Re}(\epsilon'/\epsilon)$ by $(-0.04 \pm 0.04) \times 10^{-4}$. Additionally, we have varied the minimum p cut (varying the $K_{\mu 3}$ background) and the E/p cut (varying the K_{e3} background); neither of these cuts has an appreciable effect. We conservatively assign a systematic error of 0.20×10^{-4}.

10.9 Detector Acceptance

The Monte Carlo prediction of the acceptance depends on a number of input parameters having the correct values. These include positions and sizes of limiting apertures, resolutions and inefficiencies. We consider these issues in the following sections, and give a final check of the acceptance using the distribution of the vertex z positions in the vacuum beam.

10.9.1 Limiting Apertures

For the charged-mode analysis, there are two important aperture effects not modeled in the Monte Carlo. The track separation cut defines an inner aperture, and couples to possible deviations from the ideal positions of the wires in the drift chambers. The edge of the regenerator defines the upstream aperture in the regenerator beam, and the Monte Carlo does not simulate decays within the regenerator itself, so we must understand the effective edge of the regenerator as an aperture.

Wire Positions

The track-separation cut forms a limiting inner aperture, and is dependent on the position of each wire within the drift chambers. Although the wires positions were fixed to within 20 μm on average [68], some wires are offset on the drift chamber frames. These offsets are measured in the data by determining the deviation of the SODs for neighboring sets of wires; if a wire is out of position, the SOD to one side will be high, while on the other it will be low. In the MC, the wires are placed on an ideal grid. To determine the effect of this mismatch on $\mathrm{Re}(\epsilon'/\epsilon)$, the track illumination is convoluted with the wire-cell size to determine the number of events that migrate across the track separation cut in the data but not in the Monte Carlo, which could lead to a bias on $\mathrm{Re}(\epsilon'/\epsilon)$. The measured bias from the cell size component is $(-0.16 \pm 0.12) \times 10^{-4}$, leading to a systematic of 0.22×10^{-4}.

Effective Regenerator Edge

As mentioned in Section 3.2.1, the two pions in $K \to \pi^+\pi^-$ can pass through some amount of the scintillator at the downstream edge of the regenerator and enter our data sample. We measure the effective charged mode regenerator edge to be (-1.65 ± 0.45) mm (upstream) of the physical regenerator edge, using single muon tracks that pass through the regenerator in nominal hadron running. This value depends on the level of the regenerator veto in the data, which in turn has an online trigger component and the offline 0.7 MIPs cut. The trigger source threshold is $(0.6{\pm}0.1)$ MIPs, lower than the offline cut. The average amount of material traversed leading to the 0.6 MIP energy deposition is calculated from a Landau distribution, integrating the probability distribution over the length of the last scintillator piece to find the effective edge.

$\mathrm{Re}(\epsilon'/\epsilon)$ varies as $\Delta x \cdot (0.45 \times 10^{-4})\ \mathrm{mm}^{-1}$, so the 0.45 mm uncertainty on the effective edge leads to a systematic error of 0.20×10^{-4} on $\mathrm{Re}(\epsilon'/\epsilon)$.

10.9.2 Detector Resolutions

The spectrometer resolution affects the widths of kinematic distributions such as $m_{\pi^+\pi^-}$ and p_T^2. As can be seen from Figures 10.3 and 10.5, the MC models these distributions well. The resolutions also can cause net migration of events across the E_K and z cuts, changing the number of events in the final sample. We measure the resolution of the drift chambers from data and use that as input to the MC, and cross-check the inputs using higher-level distributions. One useful distribution is the RMS of the match between upstream and downstream segments at the magnet. The momentum dependence of this quantity probes multiple scattering in the chambers and the intercept probes the average resolution of the chambers. From Figure 10.6, we see that we understand the resolution of the drift chambers very well on average. We conservatively assign a 5% uncertainty in the resolutions based on comparisons of SOD plots.

Using correlated MC samples, we determine that the slope of the change of $\mathrm{Re}(\epsilon'/\epsilon)$ as a function of the resolution scale factor is $(3.0 \pm 0.4) \times 10^{-4}$. Since we claim to understand the resolution to 5% of itself, the systematic error is 0.15×10^{-4}.

10.9.3 Drift Chamber Modeling

The drift chamber maps for the rates of high-SODs and inefficiencies are determined from $K \to \pi^+\pi^-$ data, and are used as input to the MC. For the first result [2], we weighted the maps by a factor of 1.25, because this was found to improve the data/MC agreement in the tails of the p_T^2 distribution. Due to simulation and calibration improvements, we now find that kinematic distributions match well with the nominal map factor of 1.0. To set a systematic uncertainty, we compare the rate

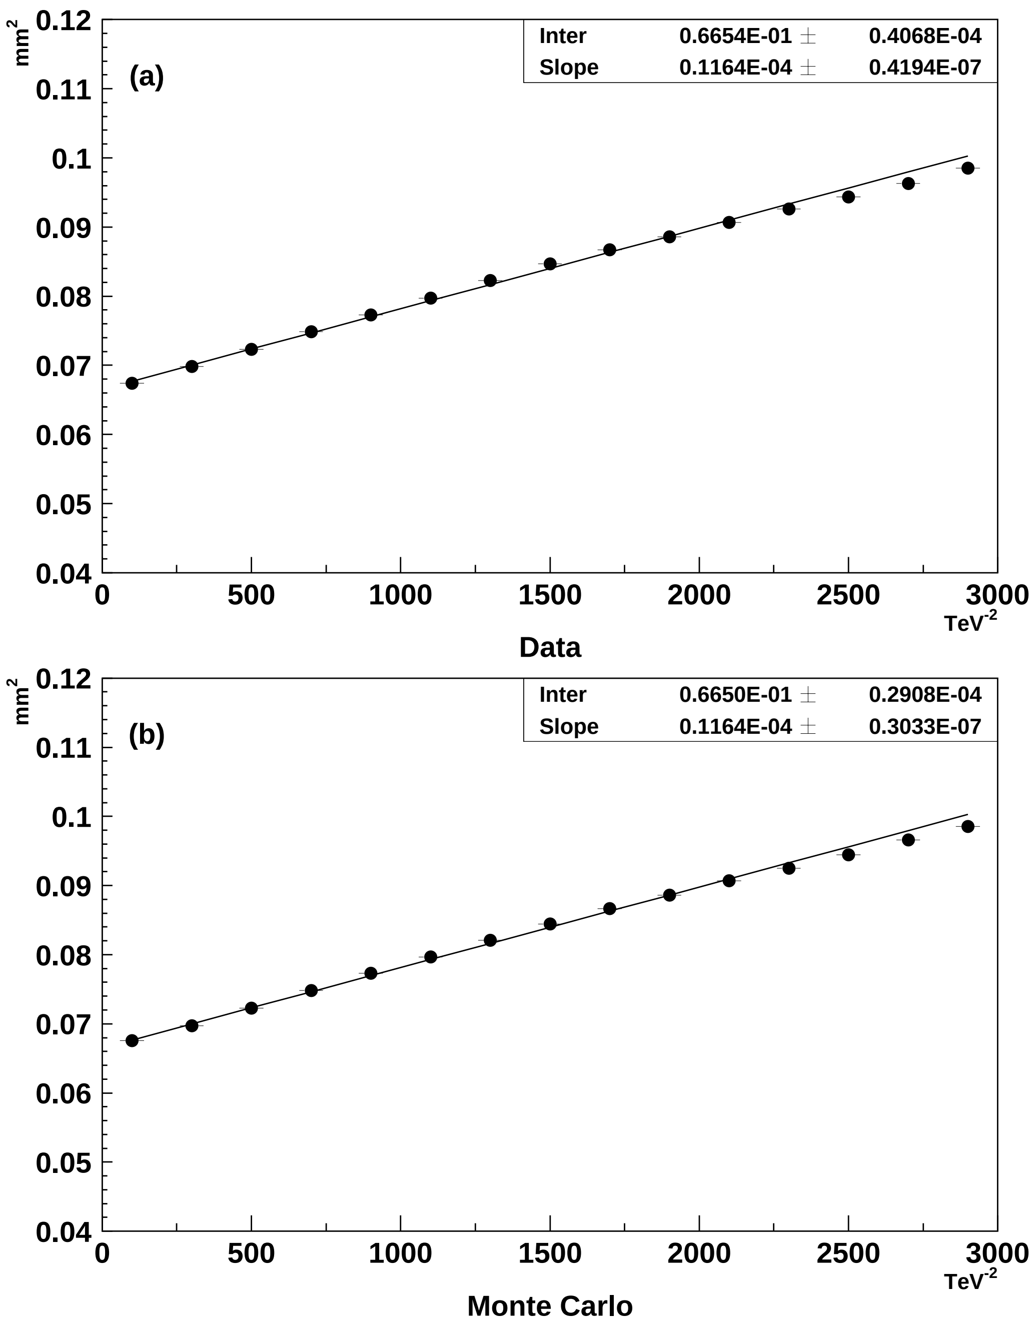

Figure 10.6: σ^2 of x_{offmag} versus p^{-2}. The intercept is a measure of the average resolution of the drift chambers, while the slope measures the multiple scattering between chambers 1 and 4. Panel (a) shows the data, while panel (b) shows the Monte Carlo. The agreement in both the intercept and slope is a cross-check of the resolutions and multiple scattering within the Monte Carlo.

of high-SODs in the data in various regions of the chambers and time periods to the same rates in the MC; the ratio of these rates is 1.0 ± 0.1, so we claim to understand the drift chamber map factor to within 10%. The map scale factor is varied within the MC from 0.8–1.2, and the variation is shown in Figure 10.7. The rate of change of $\mathrm{Re}(\epsilon'/\epsilon)$ with the scale factor is $(-1.96 \pm 0.38) \times 10^{-4}$. With the 10% possible discrepancy between data and MC, the systematic uncertainty is 0.21×10^{-4} from the drift chamber map simulation.

10.9.4 Accidental Activity

The model for the inefficiency due to early hits in the drift chamber system is described in Section 8.4.2. The systematic uncertainty due to this model is assessed by varying the parameters within the model. We have varied cuts on quantities directly related to accidentals and see little change in $\mathrm{Re}(\epsilon'/\epsilon)$. Cutting on extra hits not associated with tracks in the drift chambers has no effect on $\mathrm{Re}(\epsilon'/\epsilon)$. Requiring < 100 extra hit wires, which reduces the sample size by $\sim 7\%$, changes $\mathrm{Re}(\epsilon'/\epsilon)$ by $(0.28 \pm 0.27) \times 10^{-4}$.

We have studied various early-accidental models with correlated samples of MC. Each set of MC is roughly the size of one charged-mode dataset. The relative effects of these various models are summarized in Table 10.8. The models presented are

- Reducing the early-hit window from 300 counts to 150 counts.
- Using the 42 ns discriminator digital deadtime only.
- Using the 42 ns discriminator digital deadtime plus the “geometry” inefficiency model.
- Doubling the relative content of the “geometry” inefficiency model.
- Doubling the length of the exponential tail in the inefficiency model.
- No accidentals.

The largest change comes from the case where we do not use accidental information in the Monte Carlo at all, corresponding to a $(0.93 \pm 0.42) \times 10^{-4}$ change

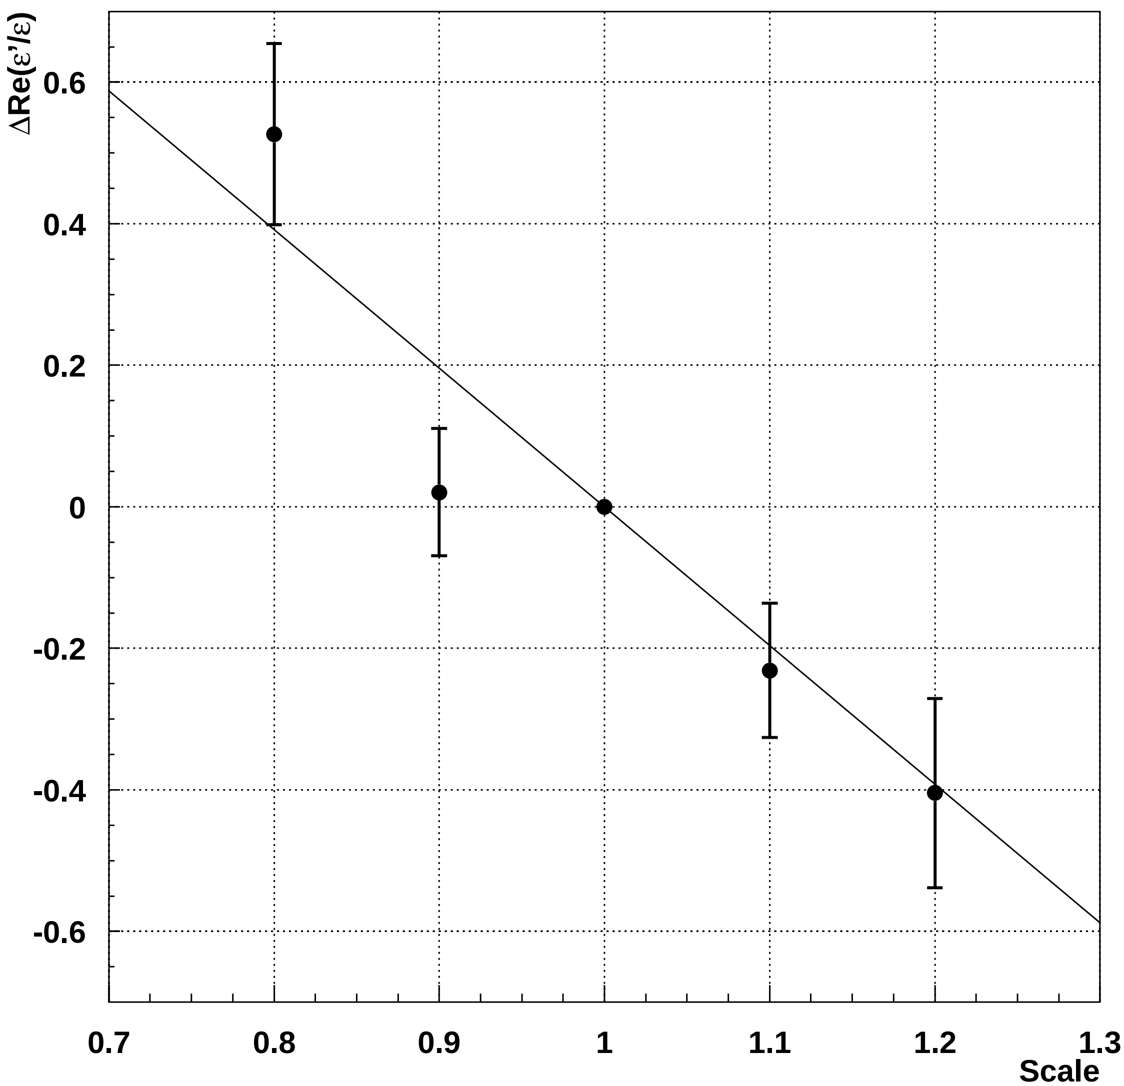

Figure 10.7: The predicted bias on $\mathrm{Re}(\epsilon'/\epsilon)$ as a function of the scale factor applied to the DC-map inefficiency and high-SOD probabilities, relative to the reference point which has a scale factor of 1.0. The error bars show the uncertainty in each case due to MC statistics.

Table 10.8: Effect of various early accidental models on Re(ϵ'/ϵ).

Accidental model	$\Delta Re(\epsilon'/\epsilon)(\times 10^{-4})$	Full Bias	z-slope ($\times 10^{-4}$)
nominal	—	—	—
1/2 timing window	-0.286 ± 0.07	0.32	0.20 ± 0.29
Digital	-0.28 ± 0.1	0.33	0.42 ± 0.04
Digital + 0.4 Geom	0.075 ± 0.07	0.13	0.15 ± 0.03
2× Geom	-0.136 ± 0.06	0.17	-0.24 ± 0.03
2× exp. length	0.0012 ± 0.06	0.06	0.12 ± 0.03
No accidentals	0.93 ± 0.42	1.13	0.91 ± 0.2

in Re(ϵ'/ϵ). However, this large change is reflected in the global acceptance check (the "z-slope"), which will be evaluated separately in a later section. We quote a systematic of 0.30×10^{-4} based on the change we see when we reduce the early timing window.

10.9.5 z Distributions as a Global Check of the Acceptance

As mentioned in Section 8.3, we use the global z vertex distribution as a final measure of our understanding of the global acceptance of detector. Many systematic effects can give rise to an acceptance slope. For example, if we were to not cut on track separation, we would see a z slope of -1.2×10^{-4} m^{-1}; and generating Monte Carlo without accidentals overlaid induces a 0.97×10^{-4} m^{-1} z-slope. Once we have a good handle on these types of effects, we use the remaining acceptance z-slope to measure our understanding of the acceptance.

The acceptance slope in $K_L \to \pi^+\pi^-$, as shown in Figure 10.8, is $(-0.41 \pm 0.34) \times 10^{-4}$ m^{-1} leading to a 0.53×10^{-4} systematic. Confirming this systematic is the value of the acceptance K_{e3} z-slope, $(0.05 \pm 0.2) \times 10^{-4}$, which gives us confidence in the $K \to \pi^+\pi^-$ slope. Note that the K_{e3} slope in the vacuum beam does not depend on the amount of primary K_S in the vacuum beam. We choose to use the $K \to \pi^+\pi^-$ z-slope instead of the K_{e3} slope because we feel that there is not a deep enough understanding of the differences between $K \to \pi^+\pi^-$ and K_{e3} behavior in the detector. For example, differences in the illuminations in the drift chambers have not been studied in detail.

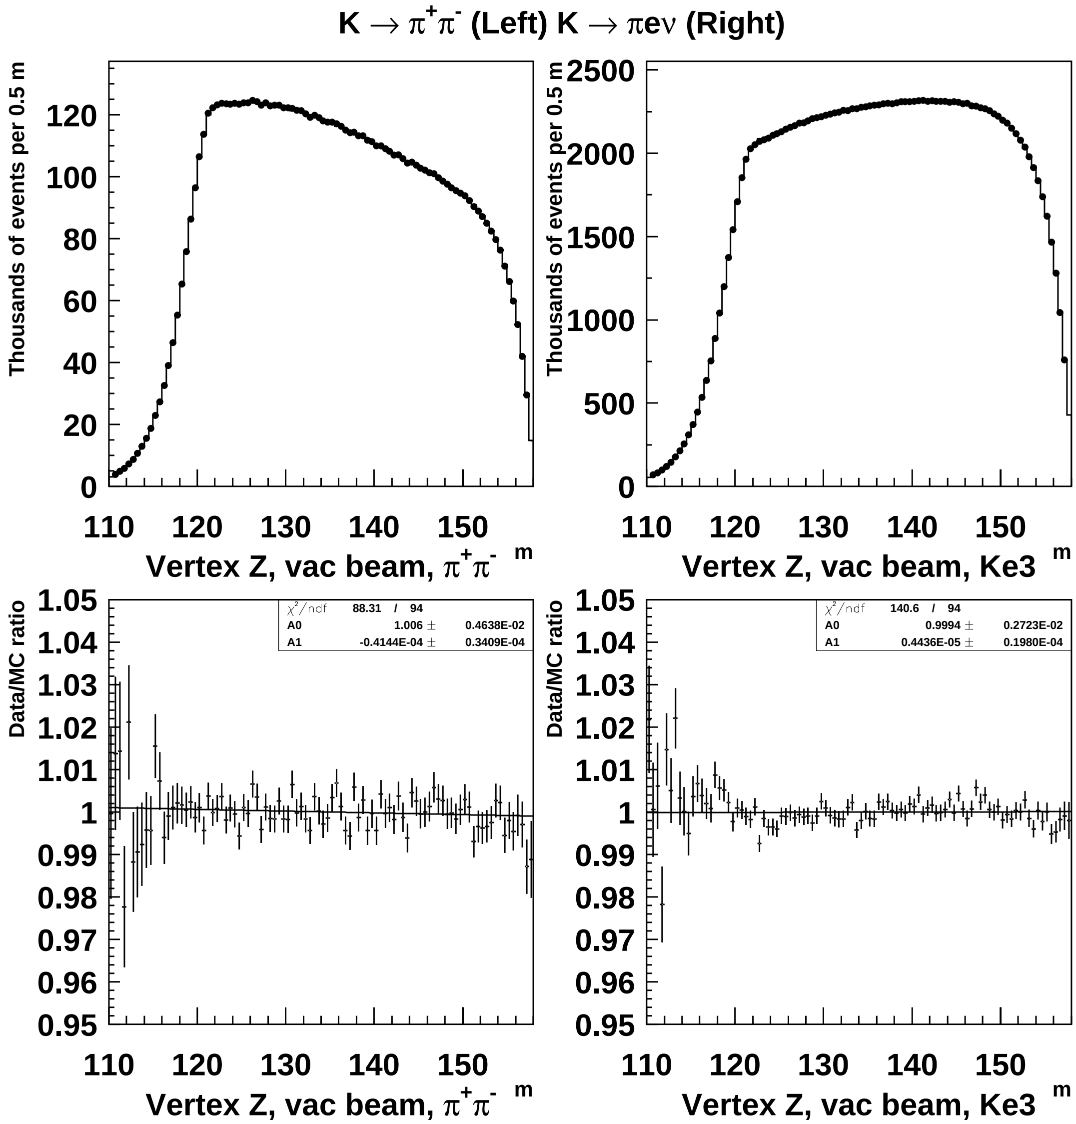

Figure 10.8: Comparison of the vacuum beam z distributions for data (dots) and MC (histogram), and the normalized ratio, for $K \to \pi^+\pi^-$ (left column) and K_{e3} (right column). The MC has been reweighted as a function of energy in the range $40 - 160$ GeV for $K \to \pi^+\pi^-$, and in the range $20 - 100$ GeV in observed energy for K_{e3}.

10.9.6 The 1997A z-Slope

One of the largest systematic issues in our first analysis was the large z acceptance bias seen in the $K \to \pi^+\pi^-$ data to Monte Carlo comparison [2]. The slope of this comparison was $(1.67 \pm 0.67) \times 10^{-4}$ m^{-1}. We have not been able to find the cause of this large acceptance slope. However, we do know that the bulk of the effect appears in the left beam, and seems to be an effect in the data, because data-to-data comparisons between 1997A and the first parts of 1997B show a large z slope. We have investigated many effects, and can eliminate the following as candidates:

- Alignment and calibration of the drift chamber system.
- Resolution of the drift chamber system.
- Magnet field corrections.

We have generated a new Monte Carlo acceptance correction for the 1997A dataset. Currently the z-slope is $(-2.1 \pm 0.59) \times 10^{-4}$. This slope leads to a systematic of 2.2×10^{-4}. This change in the acceptance slope is consistent with the change in the Monte Carlo statistics.

10.10 Dependence on Other Physics Parameters

The measured value of $\mathrm{Re}(\epsilon'/\epsilon)$ depends on the values of Δm and τ_S through the prediction of the shape of the regenerator beam. The dependence is

$$\begin{aligned}\Delta\mathrm{Re}(\epsilon'/\epsilon) \quad &= \quad \left(+0.05 \times 10^{-4}\right) \times \frac{\Delta m - 0.5267}{0.0015} \\ &\quad + \left(-0.06 \times 10^{-4}\right) \times \frac{\tau_S - 0.8965}{0.0005},\end{aligned} \tag{10.11}$$

where Δm and τ_S are in units of 10^{10} $\hbar$s^{-1} and 10^{-10} s respectively. The numerators are our best fit values for the two parameters, and the denominators are the total errors. For the systematic uncertainty on $\mathrm{Re}(\epsilon'/\epsilon)$ due to variations of these parameters, we will use one-sigma variations, leading to a 0.05×10^{-4} uncertainty from Δm and 0.06×10^{-4} uncertainty from τ_S.

$\mathrm{Re}(\epsilon'/\epsilon)$ has almost no dependence of the value of ϵ or τ_L to the precision to which they are measured. The $K \to \pi^+\pi^-$ vacuum beam is sensitive to primary K_S from the target, especially at high energy. Floating the amount of primary K_S in the vacuum beam increases the K_S component by 6.7% and reduces the z-slope for the 1997B data sample to $(-0.17 \pm 0.33) \times 10^{-4}$, and has a -0.13×10^{-4} effect on $\mathrm{Re}(\epsilon'/\epsilon)$. Hence our z-slope systematic uncertainty is conservative.

We vary the screening models amongst the three models discussed in Section 9.2.2. This leads to a 0.16×10^{-4} systematic uncertainty. We also tested our assumption of analyticity by allowing the regeneration phase ϕ_ρ to vary by $\pm 0.25°$ from the value predicted by the analyticity constraint (Equation 9.5). This results in a change of 0.07×10^{-4} on $\mathrm{Re}(\epsilon'/\epsilon)$.

10.11 Neutral-Mode Systematic Uncertainties

We do not detail all the hard work undertaken to obtain the neutral-mode systematic uncertainty, and instead refer the reader to Reference [1]. We comment on some of the larger systematic issues.

10.11.1 Linear Energy Scale

The absolute energy scale for the calorimeter is set by matching the z vertex distribution at the regenerator edge. At the downstream end of the decay region, the vacuum window serves as a production point for hadronic "junk" and for ηs that decay $\eta \to 3\pi^0$. We match the vacuum window in a manner similar to matching the regenerator edge, and see a shift of 2.1 cm. Interpreting this as a linear energy scale of the form $E = E_0(1 + \alpha)$ along the decay region results in a 1.37×10^{-4} systematic on $\mathrm{Re}(\epsilon'/\epsilon)$.

10.11.2 CsI Energy Non-Linearities

In addition to a linear energy scale in the calorimeter, we have considered a number of models of non-linearities in the calorimeter. These models include low-energy pedestal shifts and non-linearities based on corrections for neighboring clusters, to list two. The non-linearity models are applied to the data and compared to the Monte Carlo prediction (see Reference [1]). We set a systematic uncertainty of 0.66×10^{-4} on Re(ϵ'/ϵ) due to CsI non-linearities.

10.11.3 Backgrounds

The backgrounds to the $\pi^0\pi^0$ data samples are large, and ignoring them would cause a shift of 13×10^{-4} to Re(ϵ'/ϵ). The accuracy of our understanding of the background determination is limited by our understanding of kaon scattering in the regenerator, specifically the shape of the ring number distributions for these scattered events in the vacuum and regenerator beams.

As stated in Section 6.7.1, we discovered a $+1.7 \times 10^{-4}$ error in our determination of the regenerator scattering background in our published result [55], so we have invested additional effort to understand the scattering and the systematic uncertainties associated with it. The studies used to determine the systematic are detailed in Reference [1]. The neutral mode background is 1.06×10^{-4}.

10.12 Summary of Systematic Uncertainties on Re(ϵ'/ϵ)

Table 10.9 summarizes the estimates of the systematic uncertainties on Re(ϵ'/ϵ), for the 1997B dataset. Table 10.10 summarizes the estimates of the systematic uncertainties for the entire 1997 dataset dataset. Adding all the contributions from the charged-mode in quadrature, the total charged systematic is 1.24×10^{-4}. In turn, combining with the neutral mode systematic uncertainties and the common uncertainties such as the dependence on Δm and τ_S, the total systematic uncertainty on Re(ϵ'/ϵ) is 2.26×10^{-4}. This is larger than the statistical error on the measurement, indicating that we are approaching the limits of this technique to measure Re(ϵ'/ϵ). A large component of the systematic uncertainty arises from the understanding of the neutral-mode energy scale between data and Monte Carlo, and there is some hope that improvements in the neutral mode reconstruction and MC can be made in the future to reduce this systematic error. With an estimate of the systematic uncertainty, the analysis of $\mathrm{Re}(\epsilon'/\epsilon)$ is complete.

Table 10.9: Systematic uncertainties on $\mathrm{Re}(\epsilon'/\epsilon)$ for 1997B.

	Uncertainty ($\times 10^{-4}$)	
Source of uncertainty	from $\pi^+\pi^-$	from $\pi^0\pi^0$
Trigger and level 3 filter	0.62	0.16
Energy/Resolution scale	0.16	1.37
Calorimeter nonlinearity	—	0.66
Detector calib, align	0.18	0.38
Analysis cuts	0.36	0.37
Backgrounds	0.20	1.06
Detector acceptance		
Limiting apertures	0.33	0.48
Detector resolutions	0.15	0.08
Drift chamber inefficiencies	0.21	—
Accidental activity	0.30	—
Global check of z distributions	0.53	0.26
K_S/K_L flux ratio (Transmission)		
Energy dependence	0.24	
Dependence on other physics parameters		
Δm	0.05	
τ_S	0.06	
Screening assumption	0.16	
Relax analyticity assumption	0.07	
TOTAL	2.30	

Table 10.10: Systematic uncertainties on $\mathrm{Re}(\epsilon'/\epsilon)$ for 1997.

Source of uncertainty	Uncertainty ($\times 10^{-4}$) from $\pi^+\pi^-$	from $\pi^0\pi^0$
Trigger and level 3 filter	0.56	0.18
Energy/Resolution scale	0.16	1.27
Calorimeter nonlinearity	—	0.66
Detector calib, align	0.18	0.38
Analysis cuts	0.32	0.37
Backgrounds	0.20	1.07
Detector acceptance		
Limiting apertures	0.30	0.48
Detector resolutions	0.15	0.08
Drift chamber inefficiencies	0.21	—
Accidental activity	0.30	—
Global check of z distributions	0.89	0.32
K_S/K_L flux ratio (Transmission)		
Energy dependence	0.24	
Dependence on other physics parameters		
Δm	0.05	
τ_S	0.06	
Screening assumption	0.16	
Relax analyticity assumption	0.07	
TOTAL	2.26	

CHAPTER 11
CONCLUSION

Using the 1996-1997 dataset from the KTeV experiment, we have measured the kaon sector parameters Δm, τ_S, and ϕ_{+-}, and we have measured the direct CP-violating parameter $\mathrm{Re}(\epsilon'/\epsilon)$. We summarize the measurements below.

11.1 Measurements of the Kaon Sector Parameters

11.1.1 The Measurements of Δm and τ_S

The KTeV measurements of Δm and τ_S are the most precise single measurements to date, and are competitive with the world averages on these quantities. In general, our data are consistent with the later measurements of both these quantities, favoring a lower Δm and a higher value for τ_S. These measurements will affect the new PDG averages.

From the charged-mode measurements, with ϕ_{+-} fixed to be ϕ_{SW}, our results are:

$$\begin{aligned}
\Delta m &= [0.5267 \pm 0.0006\ (\mathrm{stat}) \pm 0.0014\ (\mathrm{syst})] \times 10^{10}\ \hbar s^{-1} \\
&= (0.5267 \pm 0.0015) \times 10^{10}\ \hbar s^{-1}, \\
\tau_S &= [0.8965 \pm 0.0003\ (\mathrm{stat}) \pm 0.0004\ (\mathrm{syst})] \times 10^{-10}\ s \\
&= (0.8965 \pm 0.0005) \times 10^{-10}\ s.
\end{aligned}$$

We combine the above results with the neutral-mode measurements [1]:

$$\begin{aligned}
\Delta m &= [0.5237 \pm 0.0011 \text{ (stat)} \pm 0.0018 \text{ (syst)}] \times 10^{10}\ \hbar s^{-1} \\
&= (0.5237 \pm 0.0021) \times 10^{10}\ \hbar s^{-1}, \qquad \text{(Neutral-mode measurement)} \\
\tau_S &= [0.8964 \pm 0.0005 \text{ (stat)} \pm 0.0014 \text{ (syst)}] \times 10^{-10}\ s \\
&= (0.8964 \pm 0.0015) \times 10^{-10}\ s. \qquad \text{(Neutral-mode measurement)}
\end{aligned}$$

We weight these values by the statistical and systematic uncertainties. The common external uncertainties, such as the systematic uncertainties on the regenerator, are kept separate in the weighting. The combined values are:

$$\begin{aligned}
\Delta m &= (0.5262 \pm 0.0015) \times 10^{10}\ \hbar s^{-1}, \qquad \text{(Combined charge and neutral)} \\
\tau_S &= (0.8964 \pm 0.0005) \times 10^{-10}\ s. \qquad \text{(Combined charge and neutral)}
\end{aligned}$$

11.1.2 The Measurement of ϕ_{+-}

KTeV has enough data to make the most statistically precise measurement of ϕ_{+-} to date, as shown in Equation 9.12. We have elected to float all three of Δm, τ_S, and ϕ_{+-} in our nominal fit for ϕ_{+-}, to avoid any possible reliance on the superweak assumption through a value of τ_S determined under that assumption. References [57] and [58] fixed τ_S to the world average at the time. The value of τ_S for the latter measurement of ϕ_{+-} includes measurements of τ_S that used the superweak assumption.

While this fit strategy frees us from any assumption of CPT conservation, the decision has exposed us to larger statistical uncertainties, and also larger systematic uncertainties, due to the regenerator screening correction discussed in Section 10.2.2. However, even with these issues, it is clear that we have a large dataset from which we can extract a precise value of ϕ_{+-}. Furthermore, an independent measurement of τ_S that does not assume the superweak phase will allow us to immediately improve our value of ϕ_{+-}.

Our result is:

$$\begin{aligned}\phi_{+-} &= [44.12 \pm 0.72 \text{ (stat)} \pm 1.14 \text{ (syst)}]^\circ \\ &= (44.12 \pm 1.35)^\circ\,.\end{aligned}$$

As stated in Chapter 1, the deviation of ϕ_{+-} from ϕ_{SW} is a test of CPT violation. It can be shown [62] that

$$\tan(\phi_{+-} - \phi_{SW}) = \frac{1}{\epsilon}\frac{m_{K^0} - m_{\overline{K^0}}}{2\sqrt{2}\kappa\Delta m}. \tag{11.1}$$

With our new measurements of ϕ_{+-}, ϕ_{SW}, and Δm, and using the known value of ϵ_K, we can limit

$$\left|\frac{m_{K^0} - m_{\overline{K^0}}}{m_{K^0}}\right| \lesssim 2.4 \times 10^{-18} \tag{11.2}$$

at the 95% confidence level.

In all, the combined kaon sector parameter measurements presented in this thesis further improve limits on CPT violation and further improve our understanding of the kaon sector.

11.2 Measurement of Re(ϵ'/ϵ)

From the data-set statistically independent from our published preliminary result [55], we find

$$\begin{aligned}\mathrm{Re}(\epsilon'/\epsilon) &= [19.8 \pm 1.73 \text{ (stat)} \pm 2.26 \text{ (syst)} \pm 0.80 \text{ (MC stat)}] \times 10^{-4} \\ &= (19.8 \pm 2.9) \times 10^{-4}\,, \qquad \text{(this result)}\end{aligned}$$

where we have combined all errors in quadrature in the second expression. This new result is consistent with a reanalysis of our published result,

$$\begin{aligned}\mathrm{Re}(\epsilon'/\epsilon) &= [23.2 \pm 3.0\ (\mathrm{stat}) \pm 3.2\ (\mathrm{syst}) \pm 0.8\ (\mathrm{MC\ stat})] \times 10^{-4} \\ &= (23.2 \pm 4.4) \times 10^{-4}\ . \qquad \text{(reanalysis of [55])}\end{aligned}$$

Combining these two results and accounting for the common systematic uncertainties, we obtain our full result

$$\begin{aligned}\mathrm{Re}(\epsilon'/\epsilon) &= [20.7 \pm 1.50\ (\mathrm{stat}) \pm 2.26\ (\mathrm{syst}) \pm 0.56\ (\mathrm{MC\ stat})] \times 10^{-4} \\ &= (20.7 \pm 2.7) \times 10^{-4}\ . \qquad \text{(full result)}\end{aligned}$$

This result is shown graphically in Figure 11.1, along with the recent results from the NA48 collaboration [64]. For the first time in the history of direct CP-violation, we have a good understanding of the value of $\mathrm{Re}(\epsilon'/\epsilon)$. Combining this result with the three most recent independent measurements [53, 54, 64], the new world average is

$$\mathrm{Re}(\epsilon'/\epsilon) = (17.4 \pm 1.7) \times 10^{-4}\ . \qquad \text{(new world average)}$$

The probability for this average is 13%. The fundamental conclusion for this thesis is that $\mathrm{Re}(\epsilon'/\epsilon)$ is now a well-determined parameter.

11.3 Implications for Theory

This measurement shows that ϵ' is non-zero, which precludes the possibility of a "superweak" ($\Delta S = 2$) interaction being the sole cause of CP violation in the K meson system. The measurement of $\mathrm{Re}(\epsilon'/\epsilon)$ does not mean a superweak interaction cannot exist, but it does remove a major motivation for hypothesizing its existence.

The world average value of $\mathrm{Re}(\epsilon'/\epsilon)$ is compatible with CKM-model predictions, which supports the notion of a non-zero phase in the CKM matrix. Coupled with

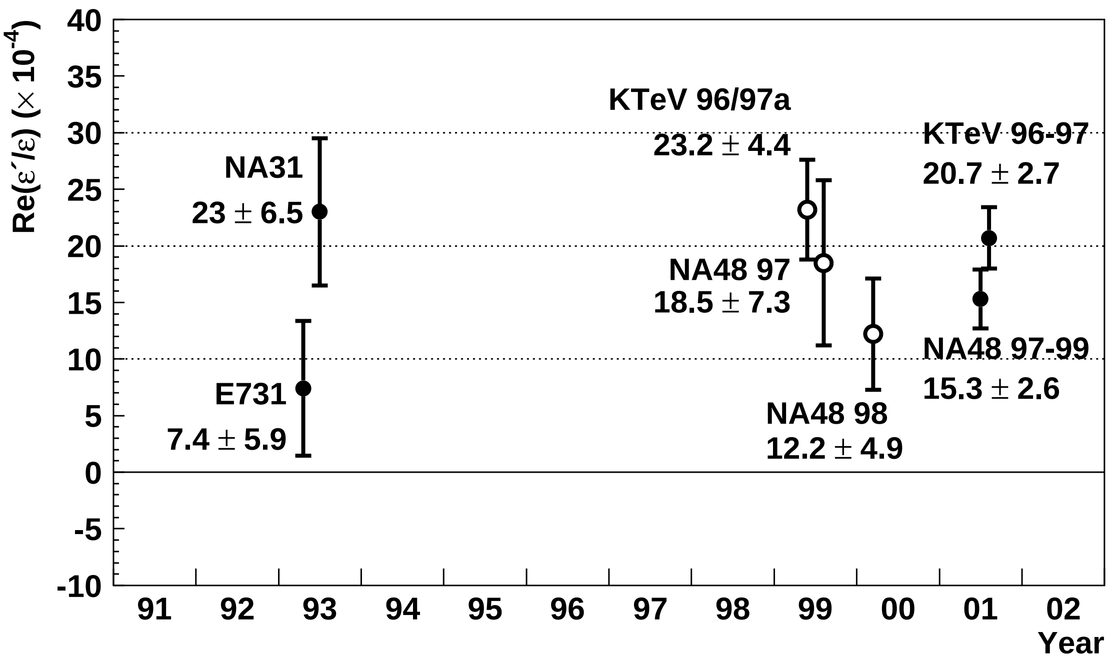

Figure 11.1: Measurements of $\mathrm{Re}(\epsilon'/\epsilon)$ by date. The measurements marked with open boxes are superseded by later measurements which include the same data.

the recent measurements of CP violation in the B system [51, 52], it seems that the Standard Model continues to work well.

The large value of $\mathrm{Re}(\epsilon'/\epsilon)$ presented in our first paper [55] had stimulated interest that there was a real discrepancy between experiment and the Standard Model. This interest led to re-examining the input parameters of the $\mathrm{Re}(\epsilon'/\epsilon)$ calculations, such as the strange quark mass, and interesting non-Standard Model contributions, such as super-symmetric processes [93]. However, with the convergence of the experimental numbers, the general consensus seems to be that there is no large problem between experiment and theoretical calculations.

To extract CKM matrix parameters from measurements of $\mathrm{Re}(\epsilon'/\epsilon)$ is difficult. The hope is that lattice gauge calculations will give us insight into this problem.

11.4 Prospects for Additional Measurements of Re(ϵ'/ϵ)

The kaon programs at both FNAL and CERN had hoped to reach a sensitivity to Re(ϵ'/ϵ) of $\sim 1 \times 10^{-4}$. The NA48 collaboration has finished running and has published their result for Re(ϵ'/ϵ), reaching a statistical error of 1.7×10^{-4}. KTeV has an equal dataset from the 1999 fixed target run on tape that is currently being analyzed. We have collected enough $K_L \to \pi^0\pi^0$ events to lower the statistical error to $\sim 1 \times 10^{-4}$ on the combined sample; however, as presented in this thesis and Reference [1], considerable effort must be expended to reduce the systematic uncertainty to this level. This is a formidable challenge, indeed.

11.5 Summary

The experimental work covered in this thesis composes one half of the necessary effort to extract the CP violating parameter Re(ϵ'/ϵ). In doing this work, we have also been able to measure the kaon sector parameters Δm, τ_S, and ϕ_{+-} to a precision competitive with the world averages of all of these quantities, which reveals the depth of the KTeV experiment.

REFERENCES

[1] V. Prasad, Ph.D. thesis (The University of Chicago, to be presented March, 2002).

[2] P. S. Shawhan, Ph.D. thesis (The University of Chicago, December, 1999).

[3] J. Schwinger, *Phys. Rev.* **91**, 713 (1953); *Phys. Rev.* **94**, 1362 (1954).

[4] G. Lüders, *Kgl. Danske Videnskab Selskab., Matt-fys. Medd.* **28**, 1 (1954).

[5] W. Pauli, *Niels Bohr and the Development of Physics*, Pergamon Press, Elmsford, NY, 1955.

[6] T. D. Lee and C. S. Wu, *Ann. Rev. Nuc. Sci.* **15**, 381 (1965).

[7] T. D. Lee and C. N. Yang, *Phys. Rev.* **104**, 254 (1956).

[8] C. S. Wu *et al.*, *Phys. Rev.* **105**, 1413 (1957).

[9] R. L. Garwin, L. M. Lederman, and M. Weinrich, *Phys. Rev.* **105**, 1415 (1957).

[10] J. I. Friedman and V. L. Telegdi, *Phys. Rev.* **105**, 1681 (1957).

[11] L. Landau, *JETP* **5**, 336 (1957).

[12] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, *Phys. Rev. Lett.* **13**, 138 (1964).

[13] V. L. Fitch, *Rev. Mod. Phys.* **53**, 367 (1981).

[14] J. W. Cronin, *Rev. Mod. Phys.* **53**, 373 (1981).

[15] V. L. Fitch, R. F. Roth, J. S. Russ, and W. Vernon, *Phys. Rev. Lett.* **15**, 73 (1965).

[16] M. Banner, J. W. Cronin, J. K. Liu, and J. E. Pilcher, *Phys. Rev. Lett.* **21**, 1103 (1968).

[17] I. A. Budagov *et al.*, *Phys. Lett. B* **28**, 215 (1968).

[18] A. D. Sakharov, *JETP Lett.* **5**, 24 (1967).

[19] G. Segre and M. S. Turner, *Phys. Lett.* **B99**, 339 (1981).

[20] J. A. Harvey, E. W. Kolb, D. B. Reiss, and S. Wolfram, *Phys. Rev. Lett.* **47**, 391 (1981).

[21] P. K. Kabir, *The CP Puzzle*, Academic Press, New York, 1968.

[22] R. G. Sachs, *The Physics of Time Reversal*, University of Chicago Press, Chicago, 1987.

[23] K. Kleinknecht, in *CP Violation*, edited by C. Jarlskog, (World Scientific, Singapore, 1989), page 41.

[24] B. Winstein and L. Wolfenstein, *Rev. Mod. Phys.* **65**, 1113 (1993).

[25] R. Belušević, *Neutral Kaons*, Springer-Verlag, Berlin, 1999.

[26] M. Gell-Mann and A. Pais, *Phys. Rev.* **97**, 1387 (1955).

[27] R. G. Sachs, *Phys. Rev.* **129**, 2280 (1963).

[28] D.E. Groom *et al.* (PDG), *The European Physical Journal* **C15**, 1 (2000).

[29] A. Angelopoulos *et al.*, *Phys. Lett.* **B444**, 38 (1998); *Phys. Lett.* **B444**, 43 (1998).

[30] H. Nguyen, to be published in *KAON 2001*, edited by F. Constantini.

[31] T. T. Wu and C. N. Yang, *Phys. Rev. Lett.* **13**, 380 (1964).

[32] W. Ochs, *πN Newsletter* **3**, 25 (1991).

[33] R. A. Briere, Ph.D. thesis (The University of Chicago, June, 1995).

[34] L. Wolfenstein, *Phys. Rev. Lett.* **13**, 562 (1964).

[35] A. J. Buras, in *Probing the Standard Model of Particle Interactions*, edited by F. Davids and R. Gupta, (Elsevier, Holland, 1999).

[36] S. Bertolini, J. O. Eeg, and M. Fabbrichesi, *Rev. Mod. Phys.* **72**, 65-93 (2000).

[37] M. Kobayashi and T. Maskawa, *Prog. Theor. Phys.* **49**, 652 (1973).

[38] L. Wolfenstein, *Phys. Rev. Lett.* **51**, 1945 (1983).

[39] S. Mele, "Indirect Determination of the Vertex and Angles of the Unitarity Triangle", Preprint `hep-ph/0103040`, 2001.

[40] E. Pallante, A. Pich, and I. Scimemi, "The Standard Model Prediction for ϵ'/ϵ", Preprint `hep-ph/0105011`, 2001.

[41] Y.-L. Wu, *Phys. Rev.* **D64**, 016001 (2001).

[42] S. Narison, *Nucl. Phys.* **B593**, 3 (2001).

[43] J. Bijnens and J. Prades, *JHEP* **06**, 035 (2000).

[44] M. Ciuchini *et al.*, "ϵ'/ϵ from Lattice QCD", Preprint `hep-ph/9910237`, 1999.

[45] A. A. Belkov *et al.*, "Phenomenological Analysis of ϵ'/ϵ within an Effective Chiral Lagrangian Approach at $\mathcal{O}(p^6)$", Preprint hep-ph/9907335, 1999.

[46] T. Hambye *et al.*, *Nucl. Phys.* **B564**, 391 (2000).

[47] M. Ciuchini, *Nucl. Phys. Proc. Suppl.* **59**, 149 (1997).

[48] S. Bertolini *et al.*, *Nucl. Phys.* **B514**, 93 (1998).

[49] S. Adler *et al.* (E787 Collaboration), *Phys. Rev. Lett.* **84**, 3768 (2000).

[50] A. Affolder *et al.* (CDF collaboration), *prd* **62**, 72005 (2000).

[51] B. Aubert *et al.*, *Phys. Rev. Lett.* **87**, 091801 (2001).

[52] K. Abe *et al.*, *Phys. Rev. Lett.* **87**, 091802 (2001).

[53] L. K. Gibbons *et al.*, *Phys. Rev. Lett.* **70**, 1203 (1993).

[54] G. D. Barr *et al.*, *Phys. Lett.* **B317**, 233 (1993).

[55] A. Alavi-Harati *et al.* (KTeV Collaboration), *Phys. Rev. Lett.* **83**, 22 (1999).

[56] V. Fanti *et al.*, *Phys. Lett.* **B465**, 335-348 (1999).

[57] L. K. Gibbons *et al.*, *Phys. Rev. Lett.* **70**, 1199 (1993).

[58] B. Schwingenheuer *et al.*, *Phys. Rev. Lett.* **74**, 4376 (1995).

[59] L. Bertanza *et al.*, *Z. Phys.* **C73**, 629 (1997).

[60] A. Apostolakis *et al.* (CPLEAR collaboration), *Phys. Lett.* **B458**, 545 (1999).

[61] K. Arisaka *et al.*, *KTeV Design Report*, Technical Report FN-580, Fermilab, 1992.

[62] L. K. Gibbons, Ph.D. thesis (The University of Chicago, August, 1993).

[63] L. K. Gibbons *et al.*, *Phys. Rev. D* **55**, 6625 (1997).

[64] A. Lai *et al.*, to be published in *The European Physical Journal*.

[65] A. R. Barker, in *Les Arcs 1992, Electroweak Interactions and Unified Theories*, edited by J. Tran Thanh Van, (Editions Frontieres, France, 1992).

[66] R. Kessler, private communication.

[67] M. L. Good, *Phys. Rev.* **121**, 311 (1961).

[68] M. Woods, Ph.D. thesis (The University of Chicago, June, 1988).

[69] R. Currier *et al.*, "KTeV Spectrometer Magnet Design Report", KTeV Internal, KTEV-0112, 1993.

[70] C. Bown, Ph.D. thesis (The University of Chicago, to be presented March, 2002).

[71] A. Roodman, in *Proceedings of the Seventh International Conference on Calorimetry in High Energy Physics*, edited by E. Cheu *et al.*, (World Scientific, Singapore, 1998), page 89.

[72] J. Whitmore, *Nucl. Instrum. Methods Phys. Res., Sect. A* **409**, 687 (1998).

[73] E. D. Zimmerman, Ph.D. thesis (The University of Chicago, March, 1999).

[74] P. L. Mikelsons, Ph.D. thesis (The University of Colorado, December, 1999).

[75] C. Bown *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **369**, 248 (1996).

[76] V. Bocean *et al.*, "The Precise Alignment of the Fermilab KTeV Experiment", Technical Report KTEV-0466, Fermilab, 1997.

[77] C. Bown, private communication.

[78] A. J. Malensek, 1981. Fermilab Reports FN-341, FN-341A (errata).

[79] W. R. Molzon, Ph.D. thesis (The University of Chicago, March, 1979).

[80] R. Brun *et al.* Computer code GEANT 3.21, CERN, Geneva, 1994.

[81] F. Sauli. Principles of Operation of Multiwire Proportional and Drift Chambers, 1980. CERN 77-09, Lectures given 1975–1976.

[82] R. Tesarek, private communication.

[83] EECS Department of the University of California at Berkeley. Computer code SPICE 3.

[84] M. Pang, "High SOD in Data and Qualitative Modeling of It", KTeV Internal, KTEV-0449, 1997.

[85] F. Gilman, *Phys. Rev.* **171**, 1453 (1968).

[86] G. J. Bock *et al.*, *Phys. Rev. Lett.* **42**, 350 (1979).

[87] B. Diu and A. Ferrez de Camargo, *Z. Phys.* **C3**, 345 (1980).

[88] A. Gsponer *et al.*, *Phys. Rev. Lett.* **42**, 9 (1979).

[89] L. Bertocchi and D. Treleani, *Nuovo Cim.* **50A**, 338 (1979).

[90] F. James and M. Roos, Computer code MINUIT, CERN, Geneva, 1994.

[91] E. Cheu, X. Qi, P. Shanahan, and S. Taeger, "Analysis of Re(ϵ'/ϵ) Using a Reweighting Method", KTeV Internal, KTEV-0775, 2001.

[92] R. A. Briere and B. Winstein, *Phys. Rev. Lett.* **75**, 402 (1995).

[93] A. Masiero and H. Murayama, *Phys. Rev. Lett.* **83**, 907 (1999).