

THE UNIVERSITY OF CHICAGO

OBSERVATION OF DIRECT CP VIOLATION IN $K_{S,L} \rightarrow \pi\pi$ DECAYS

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

PETER S. SHAWHAN

CHICAGO, ILLINOIS

DECEMBER 1999

Copyright © 1999 by Peter S. Shawhan
All rights reserved

TABLE OF CONTENTS

LIST OF FIGURES	vii
LIST OF TABLES	ix
ABSTRACT	x
ACKNOWLEDGEMENTS	xi
1 INTRODUCTION	1
2 <i>K</i> MESON PHENOMENOLOGY	5
2.1 <i>K</i> Meson Quantum States and Mixing	5
2.2 Origin of Off-Diagonal Elements	10
2.3 <i>K</i> Meson Decay Amplitudes	10
2.4 Past Measurements of $\text{Re}(\epsilon'/\epsilon)$	13
2.5 Theoretical Predictions for $\text{Re}(\epsilon'/\epsilon)$	14
2.6 Other Manifestations of <i>CP</i> Violation	16
3 THE EXPERIMENTAL TECHNIQUE	18
3.1 Overview of the KTeV Apparatus	18
3.2 Acceptance Considerations	22
3.3 Comments on Other Experiments	22
4 THE KTEV APPARATUS	25
4.1 The Beamline	25
4.2 The Regenerator	28
4.3 The Drift Chambers	30
4.4 The Analysis Magnet	31
4.5 The Calorimeter	32
4.6 Trigger and Veto Counters	39
4.7 The Trigger System	41
4.7.1 Level 1 Trigger	42
4.7.2 Level 2 Trigger	43
4.7.3 Level 1 and Level 2 Trigger Definitions	44
4.7.4 Level 3 Trigger	45

5	THE DATA	47
5.1	Data Collection	47
5.2	Problems Encountered	49
5.2.1	Calorimeter Readout Problems	49
5.2.2	Calorimeter Trigger Problems	50
5.2.3	Drift Chamber Problems	51
5.2.4	Miscellaneous Other Problems	56
5.3	Data Reduction	57
5.4	Data Samples Used for This Analysis	58
6	SELECTION OF THE $\pi^+\pi^-$ SAMPLES	60
6.1	Trigger Requirements	60
6.2	Track Finding	61
6.3	Event Reconstruction	64
6.4	Analysis Cuts	65
6.4.1	Data Quality	67
6.4.2	Trigger Verification	67
6.4.3	Veto Counters	68
6.4.4	Extra-Particle Cuts	69
6.4.5	Track Quality Cuts	69
6.4.6	Aperture Cuts	71
6.4.7	Pion Identification Cuts	72
6.4.8	Kinematic Cuts	73
6.4.9	Fiducial Cuts	74
6.5	Background Subtraction	76
6.5.1	Background Processes	79
6.5.2	Normalization of Background Contributions	82
6.6	Summary	87
7	SELECTION OF THE $\pi^0\pi^0$ SAMPLES	88
7.1	Trigger Requirements	88
7.2	Cluster Finding	89
7.3	Event Reconstruction	92
7.4	Analysis Cuts	96
7.4.1	Data Quality	96
7.4.2	Trigger Verification	96
7.4.3	Veto Counters	97
7.4.4	Extra-Particle Cuts	97
7.4.5	Event Quality Cuts	98
7.4.6	Aperture Cuts	99
7.4.7	Kinematic Cuts	100
7.4.8	Fiducial Cuts	100
7.5	Adjustment of the Photon Energy Scale	101

7.6	Background Subtraction	106
7.6.1	Background Processes	106
7.6.2	Normalization of Background Contributions	109
7.7	Summary	110
8	THE MONTE CARLO SIMULATION	113
8.1	Kaon Production, Evolution and Decay	114
8.1.1	Kaon production	114
8.1.2	Kaon transport	115
8.1.3	Evolution of the kaon quantum state	115
8.1.4	Kaon decay	117
8.2	Tracing of Decay Products	117
8.3	Simulation of the Drift Chambers	118
8.4	Simulation of the Calorimeter	124
8.5	Simulation of the Veto Counters	126
8.6	Accidental Overlays	127
8.7	Simulation of the Trigger System	128
8.8	Generation and Analysis of the Monte Carlo Samples	129
9	EXTRACTION OF $\text{Re}(\epsilon'/\epsilon)$	133
9.1	The Fitting Software	134
9.1.1	Inputs and Binning	134
9.1.2	Calculation of Decay Distributions	135
9.1.3	Calculation and Minimization of the Fit χ^2	139
9.1.4	Other Fit Types	141
9.2	The Fit for $\text{Re}(\epsilon'/\epsilon)$	141
9.3	Fitting Cross Checks	143
9.4	Summary	146
10	SYSTEMATIC ERRORS	147
10.1	Trigger Inefficiencies	148
10.1.1	Charged-Mode Trigger	148
10.1.2	Neutral-Mode Trigger	149
10.2	Event Reconstruction Biases	150
10.2.1	Calorimeter Cluster Energy Measurement	150
10.2.2	Calorimeter Cluster Position Measurement	152
10.2.3	Spectrometer Alignment and Calibration	153
10.3	Analysis Cuts	154
10.4	Backgrounds	155
10.5	Detector Acceptance	156
10.5.1	Limiting Apertures	156
10.5.2	Effective Regenerator Edge	158
10.5.3	Detector Resolutions	158

10.5.4	Drift Chamber Modeling	159
10.5.5	Accidental Activity	160
10.5.6	Z Distributions as a Global Check of the Acceptance	160
10.6	K_S/K_L Flux Ratio	163
10.7	Dependence on Other Physics Parameters	163
10.8	Cross Checks	164
10.9	Summary of Systematic Errors	165
11	CONCLUSION	167
11.1	Implications for Theory	168
11.2	Prospects for Additional Measurements of $\text{Re}(\epsilon'/\epsilon)$	169
	REFERENCES	171

LIST OF FIGURES

2.1	The two ways that CP violation can lead to $K_L \rightarrow \pi\pi$	12
2.2	Past measurements of $\text{Re}(\epsilon'/\epsilon)$	13
3.1	Plan view of the KTeV apparatus as configured to measure $\text{Re}(\epsilon'/\epsilon)$	19
3.2	Three-dimensional cutaway view of the KTeV apparatus	21
3.3	Decay vertex distributions for the $K \rightarrow \pi^+\pi^-$ and $K \rightarrow \pi^0\pi^0$ decay modes	23
4.1	The KTeV kaon beam collimation system	27
4.2	Interference in kaon decays downstream of the Regenerator	29
4.3	The hexagonal drift cell geometry of the drift chambers	30
4.4	The arrangement of CsI blocks in the calorimeter	33
4.5	Measured uniformity of light output for one particular CsI crystal, with and without reflective wrapping	34
4.6	The ratio of calorimeter energy to track momentum (" E/p ") for electrons from K_{e3} decays	38
4.7	The resolution of the calorimeter as a function of energy	39
4.8	Geometry of the trigger hodoscope planes	40
4.9	Geometry of the Collar Anti counters around the beam holes of the calorimeter	41
5.1	Division of running time during the 1996-1997 fixed-target run . . .	48
5.2	Drift chamber sum-of-distances distributions for the E773 and KTeV experiments	53
6.1	Event display for a $K_L \rightarrow \pi^+\pi^-$ event	66
6.2	Definition of angle used to calculate p_T^2	67
6.3	Vertex χ^2 distributions for the data	70
6.4	$\pi^+\pi^-$ invariant mass distributions	74
6.5	p_T^2 distributions for the data	75
6.6	Vertex Z distributions for the charged-mode data	77
6.7	Kaon energy distributions for the charged-mode data	78
6.8	Kaon scattering from the edges of the final collimator	81
6.9	Cartoon of kaon scattering in the final collimator	81
6.10	Distribution of p_T^2 vs. $m_{\pi\pi}$ for vacuum-beam data events	83

6.11	Distributions used to normalize background contributions to vacuum-beam $\pi^+\pi^-$ sample	84
6.12	p_T^2 distributions for data and background components	86
6.13	Vacuum-beam $m_{\pi\pi}$ distribution for data and semileptonic background components	87
7.1	The three ways of pairing four photons to make two π^0 's	94
7.2	Event display for a $K_L \rightarrow \pi^0\pi^0$ event	95
7.3	Pairing χ^2 distributions after all other analysis cuts have been applied	99
7.4	Four-photon invariant mass distributions	101
7.5	Neutral-mode ring number distributions	102
7.6	Vertex Z distributions for the neutral-mode data	103
7.7	Kaon energy distributions for the neutral-mode data	104
7.8	Using the Regenerator edge to adjust the photon energy scale	105
7.9	Position of energy centroid at CsI, showing Regenerator scattering extending under the vacuum beam	107
7.10	Vacuum-beam mass distribution for data and $K_L \rightarrow 3\pi^0$ background	109
7.11	Ring number distributions for data and background components	111
8.1	Spatial dependence of high-sod and missing-hit probabilities in drift chamber 1X	120
8.2	Time dependence of high-sod probability in drift chambers 1X and 3Y	121
8.3	Effect of the high-sod/missing-hit simulation on the p_T^2 distribution	123
8.4	Drift chamber retriggering model used in KTEVMC	124
8.5	Acceptance for $K \rightarrow \pi^+\pi^-$ and $K \rightarrow \pi^0\pi^0$ decays as a function of E_K and Z	131
9.1	Binning used for the fit	136
9.2	Power-law fit to $(f(0) - \bar{f}(0))/k$ measured from KTeV $\pi^+\pi^-$ data	138
9.3	Fitted values of $\text{Re}(\epsilon'/\epsilon)$ in 10-GeV bins of kaon energy	144
9.4	Sample fit from alternative technique to fit for $\text{Re}(\epsilon'/\epsilon)$	145
10.1	Comparison of different determinations of calorimeter energy scale correction	152
10.2	Data versus Monte Carlo comparisons of vacuum-beam Z distributions	161
11.1	Measurements of $\text{Re}(\epsilon'/\epsilon)$, by publication date	168

LIST OF TABLES

2.1	Some recent standard-model calculations of $\text{Re}(\epsilon'/\epsilon)$	15
6.1	Backgrounds in the $\pi^+\pi^-$ samples	85
7.1	Backgrounds in the $\pi^0\pi^0$ samples	110
8.1	Average acceptance for each of the $\pi\pi$ samples	130
9.1	Event totals and average acceptances for “grand ratio” calculation .	146
10.1	Cut variations which lead to significant changes in $\text{Re}(\epsilon'/\epsilon)$	154
10.2	Systematic uncertainties on $\text{Re}(\epsilon'/\epsilon)$	166

ABSTRACT

This thesis concerns the fundamental symmetry properties of elementary particle interactions. Specifically, it is an experimental investigation of the slight violation of CP symmetry (the combination of charge conjugation and parity inversion) by the weak force. We have used data from the KTeV experiment, located at the Fermi National Accelerator Laboratory (Fermilab), to compare the decay rates of the short- and long-lived neutral K mesons (K_S and K_L) to $\pi^+\pi^-$ and $\pi^0\pi^0$ final states. From this comparison, we find that the “direct” CP -violation parameter $\text{Re}(\epsilon'/\epsilon)$ is equal to $(28.0 \pm 3.0 \pm 2.8) \times 10^{-4}$, where the first error is statistical and the second is systematic. This nonzero result definitively establishes the existence of CP violation in a decay process.

ACKNOWLEDGEMENTS

I arrived at the University of Chicago almost nine years ago, thinking that I would most likely be a theorist—or, if an experimentalist, then certainly not in particle physics. By the end of my first year, I had found that I really enjoyed doing experiments after all, and was attracted to the fundamental questions which particle physics can address. I am grateful to my advisor, Bruce Winstein, for giving me the opportunity to play an important role in an experiment that I have found very satisfying. His wisdom and guidance have been eminently helpful over the years, and he has been a good friend.

The material presented in this thesis represents the hard work of many people. The core $\text{Re}(\epsilon'/\epsilon)$ analysis group included Ed Blucher, Colin Bown, Jim Graham, Bob Hsiung, Rick Kessler, Myungyun Pang, Val Prasad, Aaron Roodman, Peter Shanahan, and Bruce Winstein. Of course, none of this would have been possible at all without the efforts of all of the members of the KTeV Collaboration to build and commission the detector, collect the data, and contribute software and ideas to the general analysis endeavor. They are a fine bunch, and I have enjoyed working with all of them. (Even on owl shifts.)

The contributions of technical and support staff tend, unfortunately, to be taken for granted, despite the fact that they play a vital role. I am grateful for the outstanding engineering and technical assistance of Dick Armstrong, Elizabeth Pod, and Larry Fiscelli, as well as the many highly talented specialists at Fermilab.

The success of the cesium iodide calorimeter project is due in no small part to the skill and dedication of many University of Chicago undergraduates, who did most of the crystal preparation and testing. I would like to thank all of them, including Jesse Stone, Aaron Helsinger, Ashwin Tirodkar, Ben Howard, Andy Poon, Leo Chan, and especially Damon Diehl, Greg Gbur, Matt Valenta and Joe Nikolaus. Elizabeth Pod deserves special mention for her extraordinary crystal-wrapping skills.

Thanks go to Gordon Holmblad of Argonne National Laboratory for making it possible to do accurate tests of radiation damage in CsI crystals, as well as for pleasant conversation and root beer.

I have benefitted enormously from the accumulated wisdom, advice, and well-written theses of the previous generations of $K \rightarrow \pi\pi$ graduate students: Ritchie Patterson, Lawrence Gibbons, Roy Briere, and Bernhard Schwingenheuer. I have also enjoyed the company of the other Chicago graduate students in the HEP third-floor terminal room over the years. Colin Bown deserves extra credit for discovering that EPSILON PRIME has the anagram PI MESON PERIL.

I am grateful for the financial support of the National Science Foundation, and for Grainger and Harper Fellowships from the University of Chicago.

The support of my family over the years has meant very much to me. I regret that my father, who was my first physics teacher, did not live to see this thesis. Finally, I am especially grateful for the love and support of my wife, Julia, who has had to put up with my long hours and long years of graduate school.

Peter S. Shawhan

August 1999

CHAPTER 1

INTRODUCTION

Fundamental symmetries and conservation laws have played, and continue to play, a central role in helping us to understand the properties and interactions of subatomic particles. The study of K mesons (or “kaons”) has provided several crucial pieces of evidence, beginning shortly after the first kaon decay was observed in 1946 [1].¹ The peculiar features of the associated production of kaons and hyperons, and the long lifetime of the charged kaon, led M. Gell-Mann, and independently T. Nakano and K. Nishijima, to propose a new quantum number called “strangeness” which is conserved by the strong interaction, by which these particles are produced, but not by the weak interaction which is responsible for their decay [3]. In modern terms, we would say that the K mesons are the lightest particles containing a strange quark (or antiquark).

The *neutral* K mesons, in particular, provide a singularly beautiful demonstration of fundamental quantum mechanics; for a comprehensive review, see reference [4]. In 1955, M. Gell-Mann and A. Pais showed that because the two neutral kaons (called θ^0 and $\bar{\theta}^0$ at that time) can decay to the same $\pi^+\pi^-$ final state, they must mix, yielding a short-lived state that decays to $\pi\pi$ and a long-lived state that must decay to three (or more) particles [5].² The long-lived state was, in fact, discovered

¹The initial observation was referred to as a “V particle” after the shape made by the diverging tracks of the charged particles produced in the decay; a few years later, this would be called a “theta” particle. “Theta”, “tau”, “kappa”, “chi”, and “K” were names given initially to specific observed decay modes, which would eventually be recognized as all coming from what we now call K mesons [2].

²Jim Cronin recalls that the idea of mixing was stimulated by a comment by Enrico Fermi while Gell-Mann was working as an instructor at the University of Chicago [6]. Richard Dalitz recalls that, prior to the publication of the paper, Pais gave a seminar at Brookhaven National Laboratory in which he mentioned, almost in passing, the possible existence of a long-lived state, and that audience members essentially ran out of the room to set up experiments to look for it [7].

very quickly [8]. As we shall describe later, the dynamics of K^0 - \bar{K}^0 mixing and the superposition of decay amplitudes permit precise studies of the three fundamental discrete symmetries which may exist in nature:

C : “Charge conjugation”, the replacement of a particle by its antiparticle;

P : “Parity”, the inversion of the spatial coordinates (changing the coordinate system from left-handed to right-handed); and

T : “Time reversal”, the inversion of the time coordinate.

It was understood early on that, with only very general assumptions, the combination of all three operations, CPT , should be a true symmetry of nature [9, 10, 11]. This implies, for instance, that a particle must have the same mass and lifetime as its antiparticle. It initially seemed natural to assume, as well, that the laws of nature should be symmetric with respect to each operation individually.

One early mystery was the so-called “tau-theta puzzle”, since the “tau” particle (which decayed to $\pi^+\pi^+\pi^-$) and the “theta” particle (which decayed to $\pi^+\pi^-$) seemed to have the same mass, despite the fact that an analysis by R. H. Dalitz had shown that they must come from parent particles with different parity [12]. T. D. Lee and C. N. Yang suggested that the puzzle could be resolved if parity were violated by the weak interaction [13], a proposal which, while contrary to the prevailing opinion at the time, was not experimentally excluded. The observation of maximal parity violation within the following year [14, 15, 16] was a breakthrough for understanding the nature of the weak interaction, and the first demonstration of a fundamental *asymmetry* of nature.

The parity-violation experiments showed that C was violated as well [17]. However, L. Landau immediately proposed that symmetry could be restored if one considered the combined operation CP [18]. The explanation of the short- and long-lived kaon states as being the CP eigenstates of the neutral kaon system, following the proposal of Gell-Mann and Pais,³ seemed a great success, and so it was widely accepted that CP was a good symmetry of nature. Therefore, the 1964 observation by

³The original paper was written in terms of the C operation only, but is easily recast in terms of CP .

Christenson, Cronin, Fitch and Turlay of a few dozen decays of the *long*-lived kaon into two pions [19]—revealing that CP symmetry is violated to a small degree—was a complete surprise, and was initially greeted with considerable skepticism [20, 21]. 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 they were indeed the two eigenstates of the K meson system [22]. Later experiments showed that the $\pi^0\pi^0$ decays behaved similarly, indicating that the dominant effect is a CP -violating asymmetry in K^0 - \bar{K}^0 mixing [23, 24].

Other experiments over the next few years verified that the data was consistent with CPT conservation [25]. This continues to be studied, with increasing precision, both using kaons and in other physical systems, but will not be discussed further in this thesis. We may conclude that nature appears to obey CPT but no other discrete symmetry.

Besides being a cornerstone of elementary-particle physics, CP violation has important implications for understanding the predominance of matter over antimatter in the universe. In a famous paper, A. D. Sakharov showed that CP violation is a necessary ingredient for producing this predominance after the (matter-antimatter-symmetric) Big Bang [26], although at present it does not appear that the CP violation observed in the K meson system is sufficiently large to be responsible for the observed abundance of matter. This suggests that CP violation may occur in other systems, inaccessible to today’s experiments but possibly detectable at higher energies.

Whereas P (and C) violation is a natural consequence of the left-handed nature of the W boson which mediates the weak interaction, the origin of CP violation is unclear. The Standard Model of particle physics, which has been remarkably successful in so many ways, can accommodate CP violation in a natural way, but only through a parameter whose value is not known *a priori*. There are also alternative models which could be responsible for CP violation. To learn more about CP violation, we must study it in different ways. A question of fundamental importance is whether CP violation is manifested only in K^0 - \bar{K}^0 mixing, or whether it also occurs in the decay process itself; the latter is referred to as “direct CP violation”. A com-

parison of the $\pi^+\pi^-$ and $\pi^0\pi^0$ decays can be used to answer this question, in a way to be described in the next chapter. By the end of the 1960's, it was already known that direct CP violation, if it exists, is a much smaller effect than CP violation in $K^0-\bar{K}^0$ mixing. Since then, a long series of experiments at various laboratories have attempted to observe direct CP violation [27]. As of 1993, there was some evidence for a nonzero effect, but it was not conclusive, as will be explained in Section 2.4.

A new experiment, called “KTeV” (for “Kaons at the Tevatron”), has been constructed at the Fermi National Accelerator Laboratory (“Fermilab”) to continue the search for direct CP violation with substantially higher precision than previous measurements. This thesis presents an analysis of some of the data collected by KTeV which makes a new measurement of the direct- CP -violation parameter $\text{Re}(\epsilon'/\epsilon)$; a brief account of this analysis has already been published [28]. We will show how CP violation arises, and how the $\pi^+\pi^-$ and $\pi^0\pi^0$ decays can be used to distinguish between “indirect” and “direct” CP violation, in Chapter 2. Chapter 3 will present a “big picture” overview of the experimental technique, while Chapters 4 and 5 will describe the apparatus and data collection in detail. The analyses of the $\pi^+\pi^-$ and $\pi^0\pi^0$ samples will be described in Chapters 6 and 7, respectively. Accurate simulation of the detector, described in detail in Chapter 8, is essential for the extraction of $\text{Re}(\epsilon'/\epsilon)$ from the data, which is described in Chapter 9. Systematic errors are very important for this measurement, and much work went into evaluating them, as explained (in considerable detail) in Chapter 10. We will conclude with a statement of our result, which is the most precise measurement to date and sheds new light on direct CP violation, and will make some final remarks about its implications for particle theory and about the prospects for future measurements.

CHAPTER 2

K MESON PHENOMENOLOGY

In this chapter we will discuss the properties of the K meson system and define the main parameters used to describe it. In particular, we will find that there are *two* types of CP violation which may be observed in kaon decays, the second of which, “direct” CP violation, is the subject of this thesis. We will summarize the current experimental knowledge of direct CP violation, as well as the status of theoretical predictions.

2.1 K Meson Quantum States and Mixing

The formalism of K meson mixing and CP violation may be found in many reviews; see, for example, [29, 30, 31, 27, 4]. We will cover the important points without going into too much mathematical detail.

There are two neutral kaons, which are distinguished by their different quark content:

$$|K^0\rangle = d\bar{s} \quad , \quad |\bar{K}^0\rangle = s\bar{d} \quad . \quad (2.1)$$

These states are related by the operation CP as follows (using a particular phase convention):

$$CP |K^0\rangle = |\bar{K}^0\rangle \quad , \quad CP |\bar{K}^0\rangle = |K^0\rangle \quad . \quad (2.2)$$

We know, experimentally, that CP is “almost” conserved. Therefore, it will be advantageous to define the CP eigenstates

$$\begin{aligned}
|K_1\rangle &\equiv \frac{1}{\sqrt{2}} (|K^0\rangle + |\bar{K}^0\rangle) & (CP= +1) \\
|K_2\rangle &\equiv \frac{1}{\sqrt{2}} (|K^0\rangle - |\bar{K}^0\rangle) & (CP= -1)
\end{aligned}
\tag{2.3}$$

and to express the true eigenstates of the system in terms of $|K_1\rangle$ and $|K_2\rangle$. As Gell-Mann and Pais pointed out [5], the fact that the K^0 and \bar{K}^0 can decay to a common final state, $\pi\pi$, allows mixing via the weak interaction; these linear combinations then take on physical meaning rather than merely being mathematical constructs. The CP eigenstates are expected to have very different lifetimes, since the $|K_1\rangle$ can decay to $\pi^+\pi^-$ or $\pi^0\pi^0$, both of which are states of definite $CP=+1$, while the $|K_2\rangle$ must decay to a three-pion state with $CP=-1$, or else semileptonically (*i.e.* $K \rightarrow \pi^\pm e^\mp \nu$ or $K \rightarrow \pi^\pm \mu^\mp \nu$), either of which is disfavored by phase space. Thus it is natural to identify the short- and long-lived kaon states with the CP eigenstates, in the absence of evidence to the contrary.

Any general kaon quantum state $|K\rangle$ can be expressed as a linear combination of $|K_1\rangle$ and $|K_2\rangle$. The time evolution of the state is governed by the Schrödinger equation,

$$i \frac{d}{dt} |K(t)\rangle = \mathbf{H}_{\text{eff}} |K(t)\rangle , \tag{2.4}$$

where the matrix \mathbf{H}_{eff} is the effective Hamiltonian. The physically observed kaon states $|K_S\rangle$ and $|K_L\rangle$ have definite masses and lifetimes, so they must be solutions to the Schrödinger equation of the form

$$|K_{S,L}(t)\rangle = e^{-ims_{S,L}t} e^{-\Gamma_{S,L}t/2} |K_{S,L}(0)\rangle . \tag{2.5}$$

Substituting this into Equation 2.4, we must find the solutions to

$$(m_{S,L} - i\Gamma_{S,L}/2) |K_{S,L}\rangle = \mathbf{H}_{\text{eff}} |K_{S,L}\rangle . \tag{2.6}$$

We will re-express this as a matrix equation in a vector space with $|K_1\rangle$ and $|K_2\rangle$ as basis vectors. First, however, we note that \mathbf{H}_{eff} may be separated into two physically meaningful components,

$$\mathbf{H}_{\text{eff}} = \mathbf{M} - \frac{i}{2} \mathbf{\Gamma} \quad (2.7)$$

where the “mass matrix” \mathbf{M} and the “decay matrix” $\mathbf{\Gamma}$ are individually Hermitian, even though \mathbf{H}_{eff} is not. Thus we can write the matrix equation explicitly as

$$(m_{S,L} - i\Gamma_{S,L}/2) \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} = \left(\begin{bmatrix} M_1 & M_{12} \\ M_{12}^* & M_2 \end{bmatrix} - \frac{i}{2} \begin{bmatrix} \Gamma_1 & \Gamma_{12} \\ \Gamma_{12}^* & \Gamma_2 \end{bmatrix} \right) \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} \quad (2.8)$$

where M_1 , M_2 , Γ_1 , and Γ_2 are real.

So far we have made no assumption about *CPT* symmetry. The formalism can be developed further in this general case, but for simplicity we will assume at this point that \mathbf{H}_{eff} is invariant under the operation *CPT*. This requires that

$$\begin{aligned} \langle K^0 | \mathbf{H}_{\text{eff}} | K^0 \rangle &= \langle \bar{K}^0 | \mathbf{H}_{\text{eff}} | \bar{K}^0 \rangle \\ \text{and} \quad \langle K^0 | \mathbf{\Gamma} | K^0 \rangle &= \langle \bar{K}^0 | \mathbf{\Gamma} | \bar{K}^0 \rangle, \end{aligned} \quad (2.9)$$

where the second constraint arises from its connection with physical decay amplitudes [30]. Using Equation 2.8 and the fact that K^0 and \bar{K}^0 are proportional to $(1, 1)$ and $(1, -1)$ in our vector space, these imply

$$\begin{aligned} (M_{12} + M_{12}^*) - \frac{i}{2} (\Gamma_{12} + \Gamma_{12}^*) &= 0 \\ \text{and} \quad (\Gamma_{12} + \Gamma_{12}^*) &= 0 \end{aligned} \quad (2.10)$$

from which we conclude that both M_{12} and Γ_{12} must be pure imaginary numbers.

If the off-diagonal elements M_{12} and Γ_{12} were both equal to zero, then the eigenstates would simply be $|K_1\rangle$ and $|K_2\rangle$ with eigenvalues $M_1 - i\Gamma_1/2$ and $M_2 - i\Gamma_2/2$, respectively, so that we could identify $M_1 = M_S$, $\Gamma_1 = \Gamma_S$, etc. More generally, allowing the off-diagonal elements to be nonzero, we find that the eigenstates are

$$\begin{aligned} |K_S\rangle &= \frac{1}{\sqrt{1+|\epsilon|^2}} (|K_1\rangle + \epsilon|K_2\rangle) \\ &= \frac{1}{\sqrt{2(1+|\epsilon|^2)}} [(1+\epsilon)|K^0\rangle + (1-\epsilon)|\bar{K}^0\rangle] \end{aligned} \quad (2.11)$$

and

$$\begin{aligned} |K_L\rangle &= \frac{1}{\sqrt{1+|\epsilon|^2}} (|K_2\rangle + \epsilon|K_1\rangle) \\ &= \frac{1}{\sqrt{2(1+|\epsilon|^2)}} \left[(1+\epsilon)|K^0\rangle - (1-\epsilon)|\bar{K}^0\rangle \right] \end{aligned} \quad (2.12)$$

where

$$\epsilon = \frac{\text{Im}(M_{12}) - i \text{Im}(\Gamma_{12})/2}{i \Delta m - \Delta\Gamma/2} \quad (2.13)$$

with $\Delta m \equiv m_L - m_S$ and $\Delta\Gamma \equiv \Gamma_S - \Gamma_L$ (defined in this way so that both are positive quantities). It can be shown [30] that $\text{Im}(\Gamma_{12})$ is small compared to $\text{Im}(M_{12})$,¹ so that ϵ is often said to describe “mixing in the mass matrix”. Note also that

$$\langle K_L | K_S \rangle = \frac{2 \text{Re}(\epsilon)}{1 + |\epsilon|^2}, \quad (2.14)$$

i.e. the two eigenstates are not orthogonal.

The discussion above has shown that, in general, the physical kaon states need not be the CP eigenstates. This provides a very attractive explanation for the experimental fact that the K_L can decay into a $\pi\pi$ final state, as first observed in 1964 [19]. Simply put, the K_L consists of mostly $|K_2\rangle$, but with a little bit of $|K_1\rangle$; it is this $|K_1\rangle$ component which decays 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 $\bar{K}^0 \rightarrow K^0$ is faster than the rate for $K^0 \rightarrow \bar{K}^0$; this is, perhaps, the clearest way to see that the mixing is CP -violating. It also means that the K_L contains slightly more K^0 than \bar{K}^0 . (And so does the K_S !)

The value of ϵ can be measured in a few different ways. In light of the discussion above, the most obvious method is to measure the ratio of the K_L and K_S decay amplitudes to either of the $\pi\pi$ final states,

$$\eta_{+-} \equiv \frac{A(K_L \rightarrow \pi^+\pi^-)}{A(K_S \rightarrow \pi^+\pi^-)} \quad \text{and} \quad \eta_{00} \equiv \frac{A(K_L \rightarrow \pi^0\pi^0)}{A(K_S \rightarrow \pi^0\pi^0)}. \quad (2.15)$$

both of which should be equal to ϵ in the absence of direct CP violation (as will be

¹In fact, $\text{Im}(\Gamma_{12})$ is proportional to ϵ' .

discussed later), which in any case is known to be smaller than the experimental errors on these measurements. The magnitudes of η_{+-} and η_{00} may be *independently* determined from branching ratios and the K_S and K_L lifetimes, while measurement of their complex phases ϕ_{+-} and ϕ_{00} requires interference between the K_S and K_L decay amplitudes. The magnitudes and phases thus determined agree within the experimental errors [32]. Experiments which collect both charged and neutral decays have allowed the *ratio* $|\eta_{00}/\eta_{+-}|$ and the phase *difference* $\Delta\phi \equiv \phi_{00} - \phi_{+-}$ to be determined more precisely. From a fit to all information available as of 1998 [32],

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

Note that the errors on $|\eta_{+-}|$ and $|\eta_{00}|$ are correlated. We may take the phase of ϵ to be the average of ϕ_{+-} and ϕ_{00} , $(43.5 \pm 0.5)^\circ$.

Another way to determine ϵ is to measure the charge asymmetry in the semileptonic decays $K_L \rightarrow \pi^\pm e^\mp \bar{\nu}_e$ (called “ K_{e3} ”) and $K_L \rightarrow \pi^\pm \mu^\mp \bar{\nu}_\mu$ (called “ $K_{\mu3}$ ”), defined as

$$\delta \equiv \frac{\Gamma(K_L \rightarrow \pi^- \ell^+ \nu) - \Gamma(K_L \rightarrow \pi^+ \ell^- \bar{\nu})}{\Gamma(K_L \rightarrow \pi^- \ell^+ \nu) + \Gamma(K_L \rightarrow \pi^+ \ell^- \bar{\nu})} . \quad (2.16)$$

These decay amplitudes probe the $|K^0\rangle$ and $|\bar{K}^0\rangle$ composition of the K_L directly, since an ℓ^+ can come only from a K^0 decay, and likewise for ℓ^- and \bar{K}^0 , according to the “ $\Delta S = \Delta Q$ rule”. Recent results from the CPLEAR experiment at CERN constrain possible deviations from the $\Delta S = \Delta Q$ rule to a very small level [33], and so $\delta = 2 \operatorname{Re}(\epsilon)$ to an excellent approximation. Averaging published measurements from both $K_L \rightarrow \pi e \nu$ and $K_L \rightarrow \pi \mu \nu$, $\delta = (3.27 \pm 0.12) \times 10^{-3}$ [32]; using the phase of ϵ given above, this corresponds to $|\epsilon| = (2.25 \pm 0.09) \times 10^{-3}$. The fact that this is consistent with the measured values of $|\eta_{+-}|$ and $|\eta_{00}|$ gives us confidence that the same underlying mixing phenomenon is (at least mostly) responsible for all three experimental signatures.

From these measurements, we can conclude that $|\epsilon|$ is about 2.28×10^{-3} and is quite accurately known. To explore *why* it has this nonzero value, we need to consider the origin of the nonzero values of M_{12} and Γ_{12} . This will involve kaon

decay amplitudes and will lead us to consider whether CP violation may also occur in the decay process itself.

2.2 Origin of Off-Diagonal Elements

Using perturbation theory [29], we can relate the off-diagonal elements of the mass and decay matrices to specific transition amplitudes. Switching back to the K^0 - \bar{K}^0 basis, we can write

$$\Gamma_{12} = 2\pi \sum_f \langle K^0 | \mathbf{H}_W | f \rangle \langle f | \mathbf{H}_W | \bar{K}^0 \rangle \quad (2.17)$$

$$M_{12} = \langle K^0 | \mathbf{H}_{\text{SW}} | \bar{K}^0 \rangle + \sum_v \frac{\langle K^0 | \mathbf{H}_W | v \rangle \langle v | \mathbf{H}_W | \bar{K}^0 \rangle}{m_{K^0} - E_v} \quad (2.18)$$

where f represents a (physical) final state and v represents a (virtual) intermediate state. We have written \mathbf{H}_W for the ordinary weak-interaction part of \mathbf{H}_{eff} , which changes the strangeness quantum number by one unit ($\Delta S = 1$), and have written \mathbf{H}_{SW} to allow for the possibility of an additional $\Delta S = 2$ “superweak” interaction.

From these expressions, we see that the decay amplitudes contribute to the off-diagonal terms in both the mass matrix and the decay matrix. Note that if $\langle f | \mathbf{H}_W | K^0 \rangle$ and $\langle f | \mathbf{H}_W | \bar{K}^0 \rangle$ were to have the same complex phase, then $\langle K^0 | \mathbf{H}_W | f \rangle \langle f | \mathbf{H}_W | \bar{K}^0 \rangle$ would be a real number. Thus, a nonzero imaginary component for the off-diagonal elements requires the existence of a *phase difference* between $\langle f | \mathbf{H}_W | K^0 \rangle$ and $\langle f | \mathbf{H}_W | \bar{K}^0 \rangle$. In the next section, we will explicitly examine the effect of phase differences on K_S and K_L decay amplitudes.

2.3 K Meson Decay Amplitudes

We wish to consider the decay amplitudes for K^0 and \bar{K}^0 to two pions; it will be most profitable to do so in terms of isospin states, so that we will be able to relate $\pi^+\pi^-$ and $\pi^0\pi^0$. To satisfy Bose statistics, the two-pion state must have either $I = 0$ or $I = 2$. We can write the K^0 decay amplitude to the state with isospin I as

$$\langle I | \mathbf{H}_W | K^0 \rangle \equiv A_I e^{i\delta_I} , \quad (2.19)$$

where δ_I is the phase shift from final-state interactions. *CPT* invariance then implies

$$\langle I | \mathbf{H}_W | \bar{K}^0 \rangle = A_I^* e^{i\delta_I} . \quad (2.20)$$

Using Clebsch-Gordan coefficients to express the $|\pi^+\pi^- \rangle$ state in terms of the isospin states, we have

$$\begin{aligned} \langle \pi^+\pi^- | \mathbf{H}_W | K_S \rangle &= \sqrt{\frac{2}{3}} \langle I=0 | \mathbf{H}_W | K_S \rangle + \sqrt{\frac{1}{3}} \langle I=2 | \mathbf{H}_W | K_S \rangle \\ &= \frac{1}{\sqrt{2(1+|\epsilon|^2)}} \sqrt{\frac{1}{3}} \left\{ \sqrt{2} [(1+\epsilon)A_0 + (1-\epsilon)A_0^*] e^{i\delta_0} \right. \\ &\quad \left. + [(1+\epsilon)A_2 + (1-\epsilon)A_2^*] e^{i\delta_2} \right\} \\ &= \frac{1}{\sqrt{2(1+|\epsilon|^2)}} \sqrt{\frac{4}{3}} \left\{ \sqrt{2} [\text{Re}(A_0) + i\epsilon \text{Im}(A_0)] e^{i\delta_0} \right. \\ &\quad \left. + [\text{Re}(A_2) + i\epsilon \text{Im}(A_2)] e^{i\delta_2} \right\} . \end{aligned} \quad (2.21)$$

The expressions for $\langle \pi^+\pi^- | \mathbf{H}_W | K_L \rangle$, $\langle \pi^0\pi^0 | \mathbf{H}_W | K_S \rangle$, and $\langle \pi^0\pi^0 | \mathbf{H}_W | K_L \rangle$ are similar, but with different placement of factors of ϵ , $\sqrt{2}$, and (-1) . We are free to choose the phase for *one* of the A_I ; for convenience, we follow Wu and Yang [34] and take A_0 to be real. The relative phase between A_0 and A_2 is small, so that by dropping terms of order $\epsilon \text{Im}(A_2)$, we find

$$\eta_{+-} = \epsilon + \frac{\epsilon'}{1 + \omega/\sqrt{2}} \quad \eta_{00} = \epsilon - \frac{2\epsilon'}{1 - \sqrt{2}\omega} \quad (2.22)$$

where we have used

$$\epsilon' \equiv \frac{i}{\sqrt{2}} \frac{\text{Im}(A_2)}{A_0} e^{i(\delta_2 - \delta_0)} \quad \omega \equiv \frac{\text{Re}(A_2)}{A_0} e^{i(\delta_2 - \delta_0)} . \quad (2.23)$$

ω represents the ratio of the $\Delta I = 3/2$ transition amplitude to the $\Delta I = 1/2$ transition amplitude, and is experimentally found to have a magnitude of $\sim 1/22$; this dominance of the $\Delta I = 1/2$ transition is referred to as the “ $\Delta I = 1/2$ rule”. The difference between the phase shifts is determined from other data [35] to be

$\delta_2 - \delta_0 = (-42 \pm 4)^\circ$, so that ϵ' has phase $\frac{\pi}{2} + \delta_2 - \delta_0 = (48 \pm 4)^\circ$, *i.e.* essentially the same phase as ϵ . This is a fortunate coincidence which turns out to maximize the influence of ϵ' on the $K \rightarrow \pi\pi$ decay rates.

It is worth pausing to reflect on what we have just shown. First, η_{+-} and η_{00} can be different! Second, this difference depends on A_0 and A_2 having a relative phase (*i.e.* $\text{Im}(A_2)$ being nonzero in our phase convention). Third, a derivation analogous to Equation 2.21 shows that

$$\begin{aligned} \langle \pi^+ \pi^- | \mathbf{H}_W | K_2 \rangle &= \sqrt{\frac{2}{3}} i \text{Im}(A_2) e^{i\delta_0} \\ &\propto \epsilon', \end{aligned} \quad (2.24)$$

so we have just shown that the K_2 (which has $CP = -1$) can decay to a $\pi\pi$ final state (which has $CP = +1$). That is, ϵ' is a parameter describing CP violation in the decay process itself, or “direct” CP violation. So there are *two* types of CP violation which can allow the decay $K_L \rightarrow \pi\pi$, as illustrated in Figure 2.1.

Equation 2.22 tells us that direct CP violation contributes differently to $K_L \rightarrow \pi^+\pi^-$ and $K_L \rightarrow \pi^0\pi^0$, so we can measure ϵ' by comparing $K \rightarrow \pi\pi$ decay rates. Keeping terms up to order $\omega \epsilon'/\epsilon$,

$$\frac{\Gamma(K_L \rightarrow \pi^+\pi^-) / \Gamma(K_S \rightarrow \pi^+\pi^-)}{\Gamma(K_L \rightarrow \pi^0\pi^0) / \Gamma(K_S \rightarrow \pi^0\pi^0)} = \frac{|\eta_{+-}|^2}{|\eta_{00}|^2} \approx 1 + 6 [1 + \sqrt{2} \text{Re}(\omega)] \text{Re}(\epsilon'/\epsilon) \quad (2.25)$$

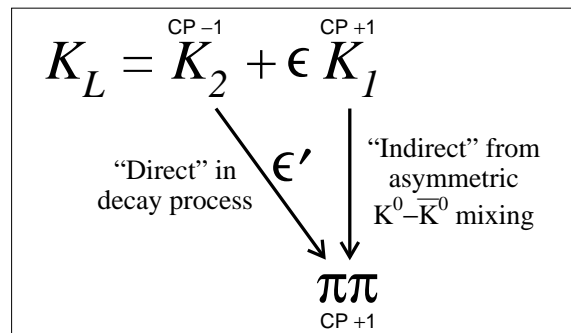


Figure 2.1: The two ways that CP violation can lead to $K_L \rightarrow \pi\pi$.

Using the measured value of $\text{Re}(\omega)$, this becomes $1 + (6.28) \text{Re}(\epsilon'/\epsilon)$. However, it is traditional to ignore ω altogether, so that the coefficient is simply 6. Thus, to be consistent with other usage, the operational definition of $\text{Re}(\epsilon'/\epsilon)$ used in this analysis is

$$\text{Re}(\epsilon'/\epsilon) \approx \frac{1}{6} \left[\frac{\Gamma(K_L \rightarrow \pi^+\pi^-) / \Gamma(K_S \rightarrow \pi^+\pi^-)}{\Gamma(K_L \rightarrow \pi^0\pi^0) / \Gamma(K_S \rightarrow \pi^0\pi^0)} - 1 \right] \quad (2.26)$$

2.4 Past Measurements of $\text{Re}(\epsilon'/\epsilon)$

The search for direct CP violation, manifested as a nonzero value of $\text{Re}(\epsilon'/\epsilon)$, has been carried out with ever-increasing precision for over 30 years. The most recent round of experiments consists of Fermilab experiment E731, which ran in 1985 and 1987-88, and CERN experiment NA31, which ran from 1986 through 1989. The history of the published measurements from these experiments is shown in Figure 2.2. In 1988, NA31 claimed the first evidence for direct CP violation, with a measured value three standard deviations from zero, based on their first year of data

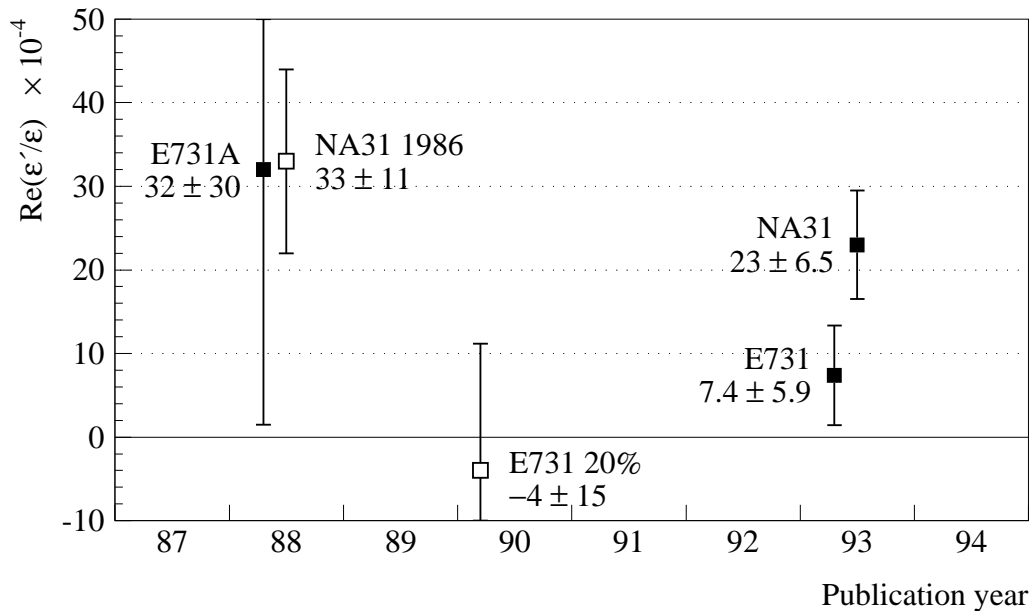


Figure 2.2: Past measurements of $\text{Re}(\epsilon'/\epsilon)$, by publication date. The measurements marked with open boxes are superseded by later measurements which include the same data.

collection [36]. Their final result is $(23 \pm 6.5) \times 10^{-4}$ [37], or about 3.5σ from zero. On the other hand, all values measured by E731 have been consistent with zero, with a final result of $(7.4 \pm 5.9) \times 10^{-4}$ [38, 39, 40]. The final results from the two experiments are not entirely inconsistent, but certainly have different qualitative implications for CP violation. The Particle Data Group, for instance, has chosen to quote the average with the error inflated to cover the spread between the two central values, $(15 \pm 8) \times 10^{-4}$ [41]. Thus, while there is some evidence for a positive value of $\text{Re}(\epsilon'/\epsilon)$, the general consensus has been that it is inconclusive and that further experiments would be needed to resolve the issue.

2.5 Theoretical Predictions for $\text{Re}(\epsilon'/\epsilon)$

At the beginning of the 1970's, the prevailing model for the weak interaction consisted of four quarks² (in two doublets) with a single parameter, the Cabibbo angle. However, M. Kobayashi and T. Maskawa showed that this model could not explain CP violation because all complex phases could be chosen to be zero [42]. They considered a few alternative models, one of which extended the system to *six* quarks (three doublets) and allowed CP violation due to the fact that the 3×3 unitary matrix describing the system had one non-trivial complex phase. Thus, in a sense, CP violation hinted at the existence of the bottom and top quarks long before their discovery. In any case, the so-called Cabibbo-Kobayashi-Maskawa (CKM) model has become a cornerstone of the Standard Model of particle physics.

The first realistic standard-model calculations of $\text{Re}(\epsilon'/\epsilon)$ were done in the late 1970's, with the dominant contribution coming from “QCD penguin” operators [43, 44]. They predicted values of order $\sim 10^{-3}$ – 10^{-2} and helped to motivate the experiments E731 and NA31 [45]. Later, it was recognized that the electroweak penguin operators contribute and tend to make $\text{Re}(\epsilon'/\epsilon)$ smaller [46], especially if the mass of the top quark is large—which eventually proved to be the case. By

²The charm quark had not actually been observed yet, but was believed to exist since it was necessary to explain the absence of strangeness-changing neutral currents via the GIM mechanism.

the early 1990's, it was even found to be possible for $\text{Re}(\epsilon'/\epsilon)$ to be zero due to a cancellation between the QCD and electroweak penguin contributions.

Further refinements over the past several years have made progress toward calculational precisions of better than 10^{-3} . Bertolini, Fabbrichesi, and Eeg have written an accessible review of modern calculation methods [47]; we will only make a few brief comments. The main uncertainties are in the determination of the hadronic matrix elements, especially those related to the QCD and electroweak penguins. Different groups have estimated these matrix elements in different ways, all expressed relative to the “vacuum saturation approximation” (VSA). Some recent calculations are listed in Table 2.1. These calculations differ substantially, and at present it is not clear which is the most likely to be reliable. The first two agree better with the E731 measurement, but given the spread, it is fair to say that we cannot conclude much more than that the calculations are of the same order as the experimental values.

Further efforts are under way to refine the calculations above, and to try new approaches. One study, recently released, uses lattice gauge theory methods to attempt to calculate $\text{Re}(\epsilon'/\epsilon)$ from first principles [51]. In contrast to the calcu-

Table 2.1: Some recent standard-model calculations of $\text{Re}(\epsilon'/\epsilon)$. All of these calculations were published *before* the announcement of the experimental result presented in this thesis.

“Home” of group	Method used to calculate hadronic matrix elements	Value of $\text{Re}(\epsilon'/\epsilon)$ ($\times 10^{-4}$)	Reference
Rome	Lattice gauge calculations	4.6 ± 3.0	Ciuchini [48]
Munich	“Phenomenological” — mostly determined from experimental data, but penguin diagram matrix elements based on VSA	8.5 ± 5.9	Buras [49]
Trieste	Chiral quark model	17_{-10}^{+14}	Bertolini <i>et al.</i> [50]

lations above, it concludes that the QCD penguin contribution greatly dominates and has a sign opposite that of conventional phenomenology, suggesting a value of $(-122 \pm 68) \times 10^{-4}$ for $\text{Re}(\epsilon'/\epsilon)$. Clearly, it will be interesting to see how the various theoretical calculations converge in the future.

The most prominent alternative to the CKM model, in terms of explaining CP violation, is the “superweak” model proposed by L. Wolfenstein shortly after CP violation was first discovered [52]. This hypothesizes the existence of a very weak interaction which changes strangeness by two units (“ $\Delta S = 2$ ”), and may be CP -violating. As shown in Equation 2.18, this interaction would connect the K^0 and \bar{K}^0 states directly, but would not play any role in kaon decays. Thus it, and not the normal weak interaction, could be responsible for all of the CP violation in the mass matrix, but would not cause any direct CP violation, so that $\text{Re}(\epsilon'/\epsilon)$ would be zero (except for some tiny higher-order effects). Therefore, conclusive experimental evidence that $\text{Re}(\epsilon'/\epsilon)$ is nonzero would rule out the possibility that a superweak interaction is the sole source of CP violation.

2.6 Other Manifestations of CP Violation

We have already described how CP violation is observable in $K_L \rightarrow \pi\pi$ decays and as a charge asymmetry in the semileptonic decays of the K_L . It has also been measured by E731 in the radiative decays $K_{S,L} \rightarrow \pi^+\pi^-\gamma$, for which $\eta_{+-\gamma}$ has been found to be consistent in magnitude and phase with η_{+-} and η_{00} , within the experimental error [53].

CP violation is expected to contribute significantly to certain rare decay modes of the K_L . Because the K_L contains both $|K_1\rangle$ and $|K_2\rangle$ components, there can be two (or more) decay amplitudes which interfere, affecting the decay rate and/or angular distributions. One particularly nice effect, recently measured using rare-decay data from KTeV [54], is seen in the decay $K_L \rightarrow \pi^+\pi^-e^+e^-$ as a large asymmetry in the distribution of the angle between the $\pi^+\pi^-$ and e^+e^- planes in the center-of-mass. This is a consequence of CP violation in mixing, and agrees well with the prediction from a calculation using the known value of ϵ [55]. Other, rarer K_L decay

modes ($K_L \rightarrow \pi^0 e^+ e^-$, $K_L \rightarrow \pi^0 \mu^+ \mu^-$, $K_L \rightarrow \pi^0 \nu \bar{\nu}$) are expected to include large contributions from *direct CP* violation, but these decays have not yet been detected.

The B meson system promises to provide much more information about *CP* violation over the coming years. Mixing occurs in neutral B mesons, as it does in neutral kaons, and *CP*-violating contributions to certain decay modes are expected to be rather large. From the multitude of B decay modes, several will yield complementary measurements of the parameters of the CKM model, with the goal of ultimately testing the self-consistency of the model. At present, the only experimental evidence for *CP* violation in the B meson system is a preliminary measurement of the quantity $\sin(2\beta)$ by the CDF collaboration which is slightly less than two standard deviations away from zero [56]. This is the first evidence of *CP* violation anywhere other than in the K meson system.

CHAPTER 3

THE EXPERIMENTAL TECHNIQUE

The smallness of $\text{Re}(\epsilon'/\epsilon)$ means that a great deal of care is required when designing an experiment to measure it. The experiment must be able to collect enough $K \rightarrow \pi\pi$ decays so that the statistical error is small. More importantly, the design of the experiment should minimize possible biases.

The KTeV experiment was designed [57] with the goal of ultimately measuring $\text{Re}(\epsilon'/\epsilon)$ with a precision of order 10^{-4} , sufficient to detect a nonzero value even at the low end of the range of the standard-model predictions. It uses the same basic technique as E731 [58], but with a new beamline and with higher-precision detector systems. The basic design of the KTeV apparatus is summarized below; a more detailed description of the beam and of each detector system will be given in Chapter 4. We will make some additional comments about the importance of the detector acceptance in the KTeV technique, and will close with some remarks about other experimental techniques.

3.1 Overview of the KTeV Apparatus

The first requirement for measuring $\text{Re}(\epsilon'/\epsilon)$ is, of course, a source of K mesons. Protons accelerated to an energy of 800 GeV by Fermilab's Tevatron are directed onto a target, and a series of sweeping magnets and collimators are used to produce two side-by-side beams of neutral particles. The beams travel about 90 m before entering the KTeV decay region, allowing short-lived particles (K_S and hyperons) to decay away and leaving mostly K_L and neutrons in roughly equal numbers.

Figure 3.1 shows the two beams entering the evacuated decay volume which extends to a large Kevlar-Mylar vacuum window at $Z = 159$ m. One of the K_L




Figure 3.1: Plan view of the KTeV apparatus as configured to measure $\text{Re}(\epsilon'/\epsilon)$. The label “CsI” indicates the electromagnetic calorimeter.

beams passes through a “Regenerator” which coherently converts some K_L to K_S (to be explained in Section 4.2). This allows K_L and K_S decays to be collected *simultaneously* by the KTeV detector.¹ This is the most important feature of the E731/KTeV technique, since it makes the measurement of the *ratio* of K_L and K_S decay rates insensitive to the inevitable time variation of beam intensity and trigger/readout deadtime, and causes detector inefficiencies and the effects of extra activity in the detector to nearly cancel as well. The Regenerator alternates sides between Tevatron beam-extraction cycles (about once per minute) to minimize the effect of any left-right beam or detector asymmetry.

The KTeV detector, beginning just downstream of the decay volume, has a geometry which is optimized to catch the highly-boosted particles resulting from

¹KTeV also collects $\pi^+\pi^-$ and $\pi^0\pi^0$ decays simultaneously, but this analysis happens to use $\pi^+\pi^-$ and $\pi^0\pi^0$ data from different running periods, for reasons to be described later.

decays of the high-energy (average ~ 70 GeV) kaons in the two beams. The main components are a charged-particle spectrometer and an electromagnetic calorimeter.

The spectrometer consists of four rectangular drift chambers, each with horizontal and vertical wires to measure the positions of charged particles passing through them, and a large dipole magnet. The full trajectory of a particle is determined by extrapolating from the upstream and downstream pairs of chambers, while its momentum is measured from the bending angle in the magnetic field. The spaces between the drift chambers are filled with helium gas to reduce scattering.

The calorimeter, located downstream of the spectrometer, uses cesium iodide (CsI) crystals as an interaction and scintillation medium to measure particle energies and positions. It is designed to have excellent resolution for the photons from $K \rightarrow \pi^0\pi^0$ decays (with each π^0 decaying immediately to $\gamma\gamma$), for which it is the only detector used to reconstruct the event. It also is used to distinguish between charged pions and electrons based on fractional energy deposit.

Because the calorimeter has finite area, it is necessary to have other detectors to catch photons which miss the calorimeter in order to reject $K_L \rightarrow \pi^0\pi^0\pi^0$ decays, which otherwise would be a large source of background in the $\pi^0\pi^0$ samples. These “photon veto” detectors are installed at various points along the length of the decay volume, at the outer edges of the drift chambers and calorimeter, and behind the beam holes in the calorimeter. The arrangement and geometry of these detectors is more clearly seen in a three-dimensional view (Figure 3.2). An additional detector called the “Mask Anti” (MA) determines the beginning of the decay volume, while the “Collar Anti” (CA) sharply defines the active area around the beam holes of the calorimeter.

A scintillator hodoscope upstream of the calorimeter is used to trigger on charged particles. An additional hodoscope, located downstream behind 4 m of steel, is used to detect muons for veto purposes. (The transition radiation detectors [TRDs] shown in Figure 3.2 were not needed for the data used in this thesis (experiment E832), and therefore were removed from the beamline during data collection. A hodoscope called the “Hadron Anti” and a second muon hodoscope behind additional steel were physically present but are not used for this analysis, so they will not be discussed.)

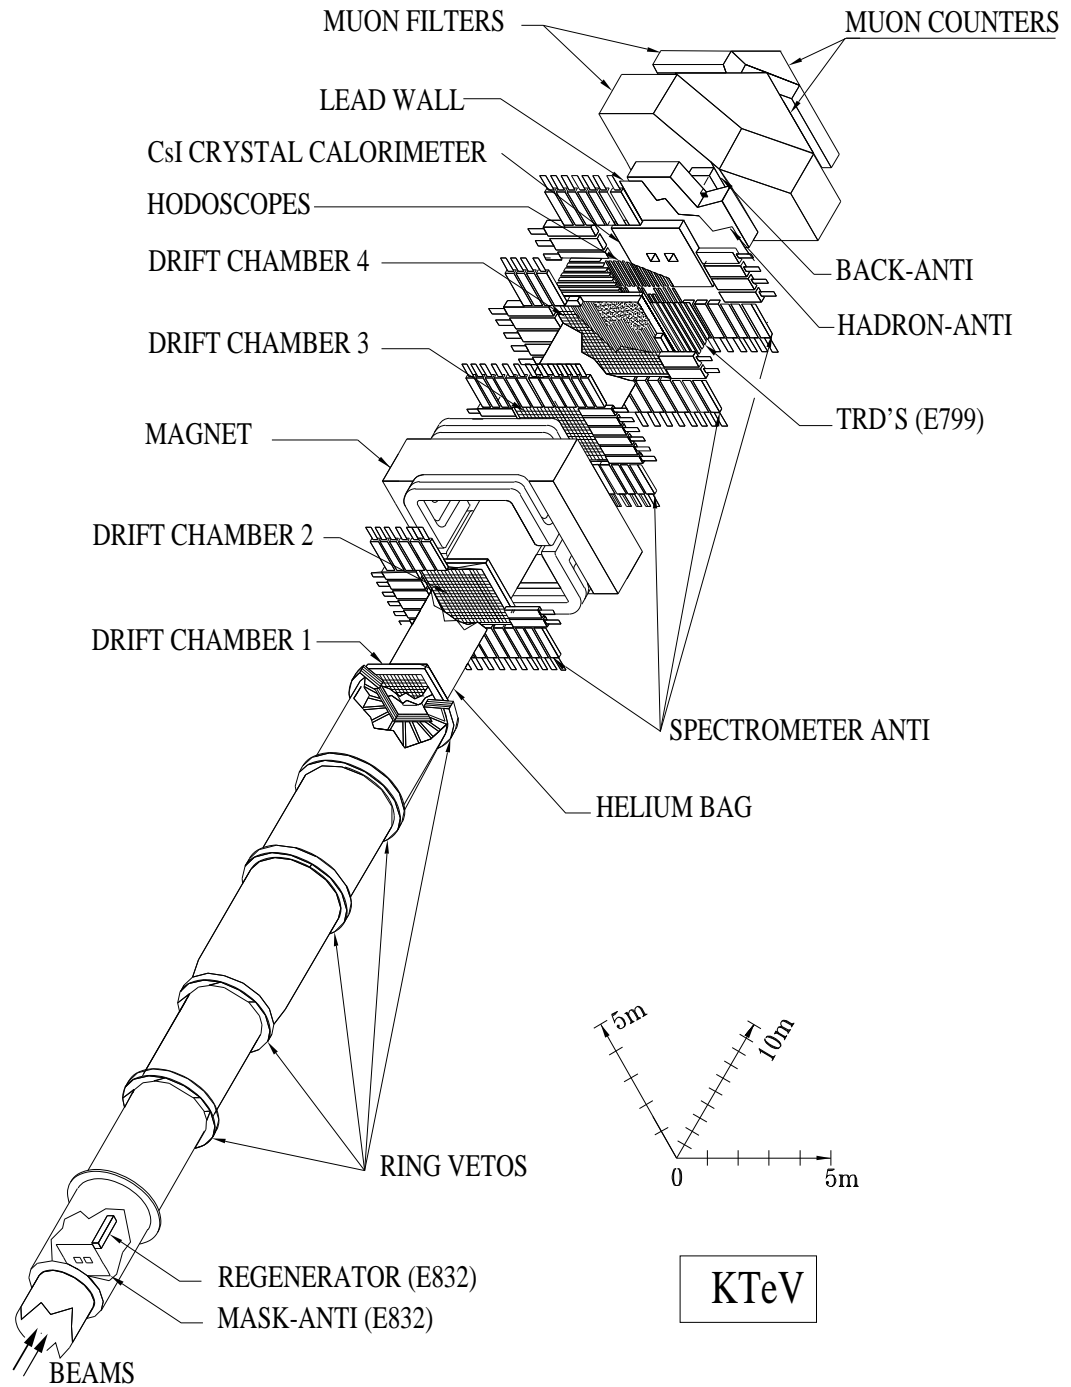


Figure 3.2: Three-dimensional cutaway view of the KTeV apparatus. (Adapted from a drawing by Elizabeth Pod.)

3.2 Acceptance Considerations

To calculate a decay rate from a number of observed events, one must know the flux of parent particles and the acceptance (*i.e.* the fraction of decays which are successfully reconstructed and identified by the detector). However, to measure $\text{Re}(\epsilon'/\epsilon)$ we only need to know the double ratio of decay rates. We have already pointed out that the double-beam technique causes the integrated beam intensity to cancel; thus the K_S/K_L flux ratio is determined solely by the material of the Regenerator (and an extra absorber in the regenerated beam, which will be discussed later), and will be the same for the $\pi^+\pi^-$ and $\pi^0\pi^0$ modes.

Triggering, event reconstruction, and event selection are done with identical criteria for decays in either beam, so deadtime and veto-counter occupancy affect the K_L and K_S samples identically. However, the detector acceptance varies substantially as a function of kaon decay position, and although the *local* acceptance should be the same for the two beams (except for small intensity-dependent effects which will be discussed later), the *overall* acceptance will be manifestly different for the K_L and K_S samples due to their different lifetimes. This is evident in the comparison of decay vertex distributions, shown in Figure 3.3 as a function of Z , the distance from the kaon production target. There is no particular cancellation between the charged and neutral modes in this case since they are reconstructed with different detector systems. Therefore, the most crucial requirement for measuring $\text{Re}(\epsilon'/\epsilon)$ with this technique is a precise understanding of the Z dependence of the acceptance for each mode. This will require a detailed “Monte Carlo” simulation, to be described in Chapter 8.

3.3 Comments on Other Experiments

As previously mentioned, the experimental technique used in KTeV is essentially the same as was used in E731. The main differences are that, for much of the E731 run, charged and neutral decays were collected separately in alternating time periods, and one photon from each $\pi^0\pi^0$ decay was required to convert in a lead sheet placed in the middle of the decay region. The largest sources of systematic uncertainty for

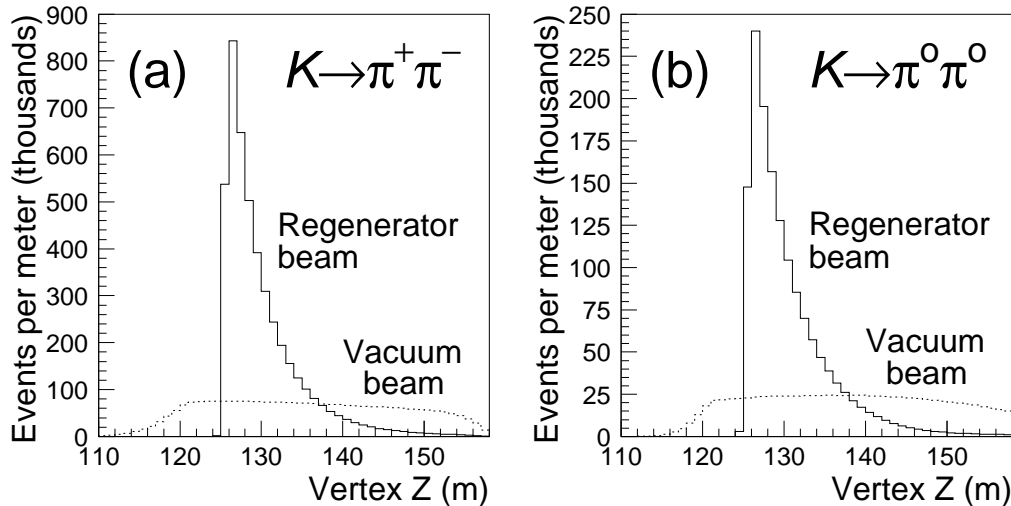


Figure 3.3: Decay vertex distributions for the (a) $K \rightarrow \pi^+ \pi^-$ and (b) $K \rightarrow \pi^0 \pi^0$ decay modes, showing the difference between the “regenerated” (K_S) and “vacuum” (K_L) beams.

E731 were due to calorimeter energy scale, beamline material, detector acceptance, and accidental activity in the detector.

The NA31 experiment at CERN [37] used quite a different technique to measure $\text{Re}(\epsilon'/\epsilon)$. They collected K_L and K_S decays in alternating periods, though in each case they collected the charged and neutral decays simultaneously so that the integrated beam flux canceled in the ratio. They produced K_S using protons incident on a target located within their decay volume, and moved the target in steps along the length of the decay volume so that the Z distribution for the K_S decays was similar to that for the K_L decays; thus the detector acceptance nearly canceled in the ratio. The detector did not include a magnetic spectrometer, and instead used a hadron calorimeter to reconstruct the $\pi^+ \pi^-$ decays. The largest systematic uncertainties were due to backgrounds, accidental activity, and calorimeter energy scale.

The NA48 experiment [59], now under way at CERN, is nominally the successor to NA31, but uses a technique much more similar to E731/KTeV. K_L and K_S decays (both charged and neutral) are collected simultaneously, with K_S 's being produced at 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. $\text{Re}(\epsilon'/\epsilon)$ is extracted using a weighting technique so that the double ratio can be directly calculated with only a small acceptance correction. After four years of data collection, NA48 expects to achieve an error on $\text{Re}(\epsilon'/\epsilon)$ of order 2×10^{-4} .

The KLOE experiment, now beginning to collect data, is unique in that it uses the DAΦNE “phi factory” in Frascati, Italy as its source of kaons. DAΦNE collides electrons and positrons of the proper energy to create ϕ mesons at rest, which decay 34% of the time to a nearly back-to-back $K_S K_L$ pair. The major challenges for the measurement of $\text{Re}(\epsilon'/\epsilon)$ will be to produce a sufficient number of ϕ mesons and to understand reconstruction biases for the low-energy decay products. KLOE will also be able to study many interesting *CPT* effects using the quantum-mechanically entwined kaon states.

CHAPTER 4

THE KTeV APPARATUS

The KTeV beamline and detector encompass many large and complex systems, and a great deal of effort went into their design and production. In this chapter we will describe several of the systems in some detail.

4.1 The Beamline

The earlier series of neutral-kaon experiments at Fermilab (E731, E773, E799-I) were done in the Meson Center beamline. The existing “building” (actually a Quonset-hut enclosure with a few adjacent trailers to house electronics and experimenters) was too small for the KTeV detector, but even with a new building the MC beamline would not have been suitable for KTeV because the beam was located above ground and there was no practical way to add enough shielding to allow the high-intensity beams required by KTeV and a proposed successor experiment, KAMI. Therefore, KTeV was sited in the “Neutrino” area on the underground beamline that previously led to the New Muon Lab, called NM. An existing enclosure was adapted for the KTeV primary target and secondary beam components, while a completely new building was constructed for the KTeV detector.

Much work went into designing clean, well-defined kaon beams for KTeV [57, 60]. They are produced using 800-GeV protons extracted from the Tevatron and steered to the KTeV target station. The Tevatron operates on a cycle about one minute long, taking ~ 40 seconds to inject and accelerate protons and then delivering beam in a “spill” ~ 20 seconds long. The beam has a 53 MHz RF structure, with the protons bunched in “buckets” 18.9 ns apart and less than 1 ns wide.

The KTeV target is a narrow rod of beryllium oxide 30 cm long (1.1 interaction lengths). After final magnetic focusing, the RMS size of the beam spot on the

target is about $250 \mu\text{m}$ in the X and Y directions, smaller than the target itself. During most of the data collection period, the beam position on the target was held steady by an “Autotune” program [61]. The proton beam and target are tilted downward at an angle of 4.8 mr relative to the “secondary” (kaon) beams formed by the collimation system; at this targeting angle, the number of neutrons is suppressed relative to the number of kaons.

The residual proton beam, plus charged particles produced in the target, are magnetically swept to the side and “dumped” in copper and steel. Special care was required to sweep away muons effectively, because of the fact that muons in the magnet return yoke bend the opposite direction. A set of three magnets were configured for this purpose, one of them specially constructed for the task. Some low-energy muons could still be bent inward toward the KTeV detector, so blocks of iron were buried around the beam pipe to stop these muons before they reached the detector.

The beam collimation system, shown in Figure 4.1, was designed to produce beams with sharply defined edges and to minimize “halo” from scattered kaons. The final beam shapes are determined by the “defining” collimator, which is located about 85 m downstream of the target. The “primary” collimator, located at $Z = 20 \text{ m}$, serves to reduce the particle flux on the defining collimator. Each of these collimators has two square holes which are tapered so that particles from the target cannot hit the inner surfaces. The geometry is such that the beam halo cannot hit the CsI calorimeter. In addition, a “slab” collimator keeps particles which scatter in the absorbers or primary collimator on one side from crossing over to the hole on the opposite side of the defining collimator. (A set of “jaw” collimators was also installed, but was used only for special studies.) A magnet just downstream of the defining collimator sweeps away any remaining charged particles from interactions in the collimators or from particle decays upstream of that point. The beams are transported in vacuum beginning at about $Z = 28 \text{ m}$.

After sweeping away charged particles, the neutral beams contain mostly photons and neutrons. Two “common absorbers” are placed in the beams; the first, consisting of 20 inches of beryllium, attenuates neutrons more than kaons because

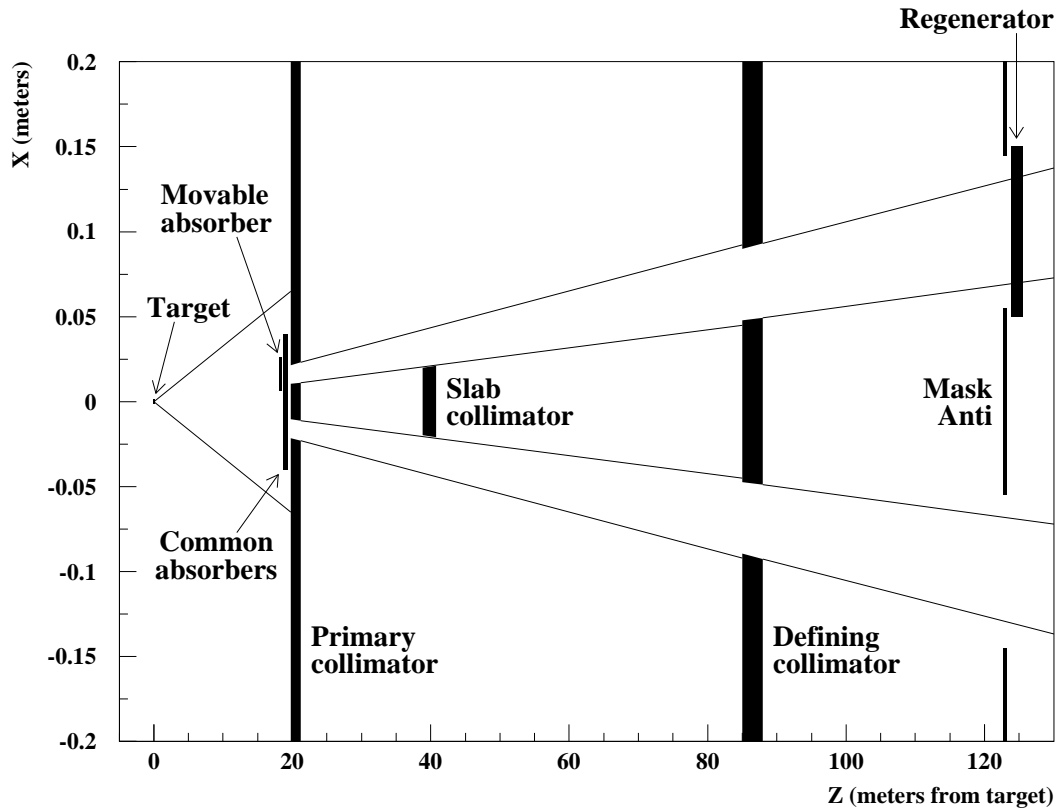


Figure 4.1: The KTeV kaon beam collimation system.

of their different interaction cross-sections, while the second, consisting of 3 inches of lead, converts photons and absorbs the ensuing electromagnetic showers. An additional beryllium absorber moves from side to side to further attenuate the flux incident on the Regenerator; its length, 18 inches, was chosen to reduce the particle flux through the KTeV detector from hadronic interactions in the Regenerator without lowering the number of K_S 's produced to the point where it greatly affects the statistical error on the $\text{Re}(\epsilon'/\epsilon)$ measurement.

By the time the beams reach the beginning of the decay region, they consist mostly of K_L and neutrons in a 1:2 (vacuum beam) or 1:1.5 (regenerated beam) ratio [62]. However, at high energy, a significant K_S component survives.

4.2 The Regenerator

The K^0 and \bar{K}^0 states have different quark content ($d\bar{s}$ vs. $\bar{d}s$), and so they interact differently with ordinary matter, which contains u and d quarks but not \bar{u} and \bar{d} . As a result, a K_L ($\approx (K^0 - \bar{K}^0)/\sqrt{2}$) which passes through matter emerges as a different mixture of K^0 and \bar{K}^0 ; thus, in the K_S - K_L basis, a nonzero K_S component is “regenerated” in the material. The KTeV Regenerator takes advantage of this remarkable property to provide the K_S decays needed for the measurement of $\text{Re}(\epsilon'/\epsilon)$.

The desire is to use kaons which undergo only coherent forward scattering, without any measurable momentum transfer. The state of a kaon emerging from the Regenerator may then be expressed as $K_L + \rho K_S$, where ρ is a complex number.¹ Of course, the vast majority of kaons incident on the Regenerator interact inelastically or quasi-elastically, with substantial momentum transfer; when this happens, the kaon generally emerges with large transverse momentum and in a quantum state with a different effective regeneration parameter ρ than for the coherent sample. To reject such cases, the KTeV Regenerator is made of blocks of plastic scintillator viewed by photomultiplier tubes to detect the recoil products from the inelastic interaction. The signals from the PMTs are integrated and digitized with ADCs; a few of them are also discriminated, latched, and used by the trigger system. There remains a type of scattering, called “diffractive,” in which the kaon acquires transverse momentum but there are no detectable recoil products; this is an unavoidable background which will need to be subtracted. However, the length of the KTeV Regenerator, around two interaction lengths, was chosen to minimize the amount of diffractive scattering relative to coherent regeneration.

The fact that the Regenerator makes a *coherent* mixture of K_L and K_S , both of which can decay to $\pi\pi$, means that interference occurs in the decay process. This is evident in the distribution of decays downstream of the Regenerator, shown in Figure 4.2 for kaons in a limited energy range.

¹The kaon state actually is still more K_L than K_S , but essentially all of the K_S component decays away within the KTeV decay region, compared to only a few percent of the K_L component.




Figure 4.2: Interference in kaon decays downstream of the Regenerator.

Many kaons decay within the Regenerator, but we do not want to include these in our final samples. Therefore, the last module at the downstream end of the Regenerator contains two lead pieces, each 5.6 mm (one radiation length) thick and each followed by a 4-mm-thick piece of plastic scintillator viewed by PMTs. For a $\pi^0\pi^0$ decay within the Regenerator, there is almost certain to be at least one photon conversion in the lead, so the event will not be reconstructed.² A $\pi^+\pi^-$ decay will deposit energy in the scintillator, and will normally be rejected. However, some kaon decays, very close to the downstream end of the Regenerator, will survive and will be included in the final data samples. This must be taken into account when extracting the value of $\text{Re}(\epsilon'/\epsilon)$ and when estimating systematic errors. The lead piece contributes significantly to regeneration and to scattering, so it must be included in the model used in the Monte Carlo simulation.

²It does not matter whether the conversion products are detected; the mere fact that a photon was destroyed guarantees that the event will not be reconstructed.

4.3 The Drift Chambers

The only parts of the E731 detector to be re-used for KTeV were the frames of the four drift chambers. All of the wires were re-strung for KTeV, with gold-plated aluminum used for the field-shaping wires (which are thicker and more numerous) to reduce particle scattering. Each of the four drift chambers has two vertical and two horizontal planes of sense wires to measure charged-particle positions in the X (horizontal) and Y (vertical) directions, respectively. The system contains a total of 1972 sense wires. The field-shaping wires are arranged in a hexagonal cell pattern around each sense wire, as shown in Figure 4.3. In most cases, a particle leaves a “hit” on one wire in each plane, and the staggered configuration allows the particle position to be determined unambiguously. It also provides a distinct signature for an in-time particle: the sum of the drift distances (“sum-of-distances”, or “sod”) should equal the cell spacing, 6.35 mm. In cases where one wire is “inefficient”, *i.e.* it fails to record a hit, the sod signature is not available and the particle position




Figure 4.3: The hexagonal drift cell geometry of the drift chambers. A typical charged-particle trajectory through the drift chamber is shown; the thin lines represent the drift paths of ionization electrons toward the sense wires.

has a twofold ambiguity which must be resolved using information from the other drift chambers.

The drift chambers were operated with a 50%-50% mixture of argon and ethane, bubbled through isopropyl alcohol. For the data analyzed in this thesis, the voltages were set to -2450 or -2500 , determined by a compromise between high efficiency and high current draw. Under these conditions, the drift speed is nearly constant over most of the drift cell at about $50 \mu\text{m}/\text{ns}$. The electrical pulse on the sense wire is amplified in two stages and discriminated, with the time measured by a multi-hit TDC with a precision of 0.5 ns .

The calibration of the drift chamber system required several steps and will only be summarized here. First, a timing offset was determined for each wire to align all of the raw hit-time distributions in a given sense plane. Next, the relationship between drift distance and drift time was inferred from the raw hit-time distribution (summed over all wires in a plane), using the assumption that the illumination is uniform (on average) over a drift cell. This stage was iterated a few times. Once the chambers had been individually calibrated, they were aligned to each other and to the target and calorimeter by projecting tracks from samples of muons (taken with the analysis magnet turned off, so that the tracks were straight), $K_L \rightarrow \pi e \nu$ and $K \rightarrow \pi^+ \pi^-$ decays.

After final calibration and alignment, the drift chambers measure horizontal and vertical position with a typical resolution of $110 \mu\text{m}$ and momentum with a resolution of 0.4% at $36 \text{ GeV}/c$, which is the mean momentum of pions from $\pi^+ \pi^-$ decays.

4.4 The Analysis Magnet

A new analysis magnet was constructed for KTeV. It was operated at a current of 1600 A and imparted a horizontal momentum kick of $0.412 \text{ GeV}/c$ to charged particles; the exact field integral was determined (in each of several time periods) from $\pi^+ \pi^-$ data by constraining the reconstructed invariant mass to match the known K mass. The polarity of the field was flipped about once per day to avoid

biases for $\pi^\pm e^\mp \bar{\nu}_e$ asymmetry studies, etc., but this is irrelevant for $\pi^+\pi^-$ decays because of their inherent charge symmetry. After the magnet was assembled, the field integral was mapped using a “zip track” system and found to be uniform to better than 0.25% over most of the clear aperture, as expected; corrections were applied based on this map when determining track momentum.

There was a small but non-negligible “fringe field” from the analysis magnet which extended past chambers 2 and 3 and slightly displaced the hit positions in those chambers. The magnitude of the field was determined from the data, and a correction was applied during track reconstruction to undo this displacement.

4.5 The Calorimeter

The calorimeter was the most ambitious (and most expensive) new detector for KTeV. The lead-glass calorimeter used for E731 (as well as for E773 and E799-I) had been a major limitation, although after years of study it was very well characterized. KTeV needed a new calorimeter which would have excellent energy resolution and linearity, good position resolution, fast signals, and which could withstand a fairly high radiation dose without serious degradation. It was decided to use scintillating crystals, and after building and testing some prototypes using various candidate materials (including barium fluoride and lead fluoride), pure cesium iodide (CsI) was chosen.³

CsI has two scintillation components, one “fast” (decay time ~ 25 ns, peak wavelength ~ 305 nm) and one “slow” (decay time ~ 1 μ s, peak wavelength ~ 480 nm). The fast component has sufficient light so that, at KTeV energies, the resolution term from photostatistics is acceptably small. Studies of radiation damage in small CsI pieces suggested that it could be made sufficiently radiation hard, although this had not been shown for large pieces. However, CsI also has some properties which make handling inconvenient: it is fairly soft, which makes polishing difficult

³While other experiments, notably CLEO, had used CsI doped with tellurium (for increased light production), KTeV and KEK experiment 162 are the first to build large calorimeters from pure CsI.

and makes the crystals susceptible to scratches and bending, and it is somewhat hygroscopic, requiring a dry environment for storage and handling.

The KTeV calorimeter [63] was designed as a square array, 1.9 m on a side, as shown in Figure 4.4. There are a total of 3100 blocks in two sizes: 2.5×2.5 cm in the central region (where position resolution and the ability to distinguish two close photons are more important) and 5×5 cm in the outer region. All of the blocks are 50 cm long (27 radiation lengths) so that the electromagnetic shower from a high-energy photon or electron is almost fully contained, minimizing the resolution term from fluctuations in the amount of energy lost out the back of the CsI array. Two 15 cm square beam holes (supported by square carbon-fiber tubes) allow passage of the neutral beams through the calorimeter.

The crystals were purchased from three vendors: Solon Technologies (later a division of Bicron) in Ohio, Quartz & Silice (later called Crismatec) in France, and




Figure 4.4: The arrangement of CsI blocks in the calorimeter.

Horiba in Japan. Most of the 50-cm-long blocks from the first two vendors actually consisted of two crystals glued together before final machining and polishing.

All of the crystals were received and processed at the University of Chicago in a dry room built for that purpose. The main task was to wrap each crystal for uniform light collection, in order to make the energy measurement insensitive to fluctuations in the shape of the electromagnetic shower. The wrapping was mylar, 13 μm thick, printed black on one side and aluminized on one half of the other side, and secured with small pieces of thin, backless tape. By choosing where along the length of the crystal to put reflector and where to leave black, we were usually able to make the light collection uniform to within 5%, as determined using an automated test setup which measured the light output from a Cs-137 source moved to ten places along the length of the crystal. Figure 4.5 shows the effect of wrapping for one particular crystal, along with the resulting improvement in the predicted energy resolution.

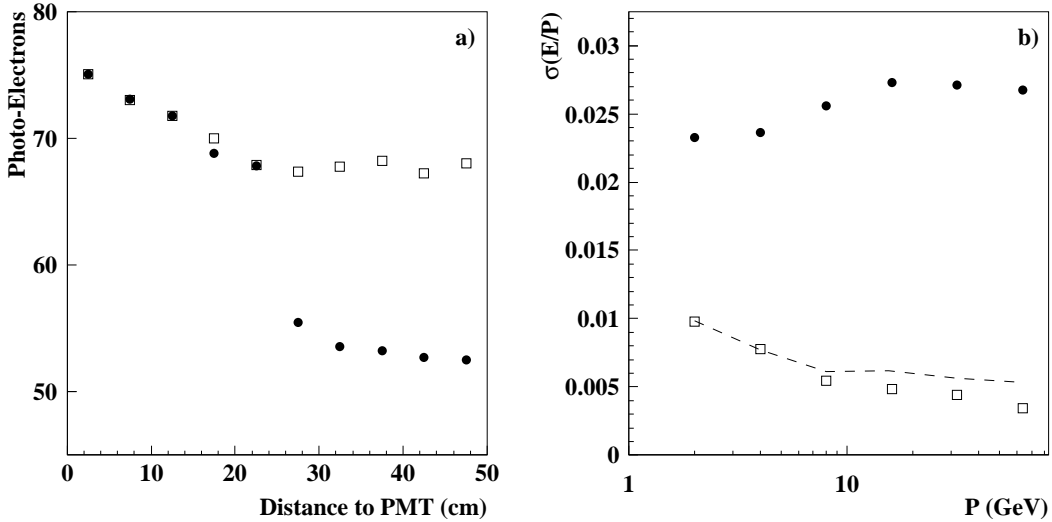


Figure 4.5: Measured uniformity of light output for one particular CsI crystal, with and without reflective wrapping. In (a), the solid circles show the uniformity with a non-reflective wrapping, while the open squares show the uniformity when aluminized mylar is added over the region from 25 to 50 cm from the PMT. Panel (b) shows the predicted energy resolution, as a function of photon energy, in each case. The dashed line indicates the energy resolution expected for a crystal with perfectly uniform light collection.

Other procedures done at Chicago included measuring the precise physical dimensions of each crystal at 20 points along its length and gluing on a plastic “flange” with a magnetic shield and a coupling for an optical fiber. A final scintillation-uniformity test, using cosmic-ray muons, was then done to certify that the crystal was ready to be used.

Simulations indicated that the crystals nearest the beams would receive a radiation dose of several kilorads over the course of the experiment. (Such a dose would have been fatal for the E731 lead glass, which degraded seriously after doses of only a few hundred rads.) Existing information about the radiation hardness of CsI was based mostly on small samples, so we began a program of radiation damage tests on full-size crystals. A few initial tests we done by placing crystals in the E799-I neutral beam, just in front of the beamstop; later on, crystals were irradiated with a powerful Co-60 source at Argonne National Laboratory. In general, the radiation induced a slope in the uniformity along the crystal,⁴ but with significant variations from sample to sample. Some crystals were observed to “recover” substantially on a time scale of days to weeks. We found reasonable correlations among crystals from the same manufacturer and/or batch, with crystals from Horiba and Solon generally faring the best. We did not find correlations with other quantities, such as the relative amounts of fast- and slow-component light output, which could have allowed us to predict the radiation hardness of a crystal *a priori*.

Crystals were assigned locations in the array subject to several constraints. First, the crystals in a given row all had to have nearly the same height, to make a flat base for the row above. Second, each pair of small-crystal rows had to have nearly the same total width, to meet the large crystals at the ends. (The widths of the rows vary, with the widest row being 6.5 mm wider than the narrowest row, so the array is not truly square. Thus the exact position of each crystal must be used in the analysis software.) Third, we required that the crystals near the center of the array all be made by Horiba or Solon and *not* have reflective wrapping on the upstream end, to be more resistant to radiation damage and to leave us a means of correcting

⁴There was also an overall drop in light output, but this is much less of a concern for KTeV.

for any degradation that might occur.⁵ Finally, we generally placed the crystals with the best uniformities closer to the center of the array. The crystals were stacked one row per day, and each row was tested with cosmic-ray muons overnight before stacking the next row.⁶

The calorimeter was assembled in a sealed metal enclosure, affectionately known as “the blockhouse,” in the KTeV hall. The blockhouse has its own dehumidifying and air conditioning systems to protect the crystals and to cool the electronics. Vertical “ribs”, just downstream of the crystals, support the readout electronics and convey power, digital and analog signals, and optical fibers. The density of electronics and cables pushes the limit of the available space; it is still possible to replace a PMT or an optical fiber, but it takes special tools and skilled hands! The front of the blockhouse has a “window” of several layers of aluminized mylar to allow particles to enter through minimal material while keeping the interior dark.

The scintillation light from each CsI block is collected by a photomultiplier tube (PMT) coupled to the crystal with a soft silicone “cookie” and a disk of Schott UG-11 filter glass, which blocks a significant fraction of the slow-component scintillation light. The PMTs were custom-designed for KTeV by Hamamatsu to meet the stringent specification on linearity at the rather low gain required by KTeV. For many crystals, we used a cookie with an embedded black annular “mask” to reduce the amount of light seen by the PMT, so that we can operate the PMT at a higher gain where the linearity is better.

The CsI signal is read out by a novel “digital photomultiplier tube” (DPMT) device [64], in which the PMT anode is connected by a very short cable to a circuit which digitizes the signal, buffers the digitized value, and transmits it on demand. The heart of the device is an ASIC (Application Specific Integrated Circuit) developed at Fermilab called the QIE (for “charge integrator and encoder”) which divides the input signal into several ranges (with $1/2$, $1/4$, $1/8$, ... of the input signal), accumulates charge on capacitors over a given time interval (“slice”), selects the range

⁵In practice, upstream-face wrappings were modified *in situ* only for a few crystals.

⁶Long-term changes in crystal uniformities are monitored *in situ* using cosmic-ray muons. Some degradation is seen on average, but it correlates only slightly with radiation dose.

for which the voltage on the capacitor lies within a certain window, and outputs that voltage to be digitized by a separate flash ADC chip. The range selected is indicated by a binary “exponent” code, which is recorded with the digital “mantissa” from the flash ADC. Each range actually uses four capacitors in a rotating sequence, so that the device continuously produces outputs for each time slice. A second ASIC called the DBC (for “driver, buffer, clock”) contains a FIFO to store the digitized information from the QIE and flash ADC on each clock cycle, and transfers the data for several consecutive time slices to a “pipeline” module when an event is accepted by the trigger. Thus, this device provides a wide dynamic range with very low noise (pedestal width ~ 0.5 MeV) and with pulse shape information.

The multitude of DPMT ranges and capacitors makes calibration rather complex. In addition, during 1996 there were three classes of events with different timing of the DPMT slices relative to when the energy was deposited. Cross-calibration is done using a laser calibration system which distributes shaped pulses of UV light (produced in a dye by the primary laser pulse) to each crystal along a quartz fiber connected to a hole in the flange; the light intensity is scanned up and down over the full dynamic range, and referenced to PiN photodiodes with excellent intrinsic linearity. This system is also useful for monitoring short-term drifts in the PMT gains during normal running.

The channel gains are then calibrated using electrons, by matching the energy measured by the calorimeter to the momentum measured by the spectrometer, in an iterative procedure. For the 1996 running period, the calibration process used a total of 1.9×10^8 electrons from $K_L \rightarrow \pi e \nu$ decays collected during normal running. Figure 4.6 shows the ratio of the calorimeter energy to the track momentum (“ E/p ”) for these electrons after all calorimeter corrections. Figure 4.7 shows the energy resolution of the calorimeter itself (inferred by subtracting the component of the E/p resolution due to the spectrometer) as a function of energy. This indicates that the calorimeter has an average energy resolution of 0.7% for photons from $\pi^0 \pi^0$ decays, which have a mean energy of 19 GeV. The X and Y positions of electrons and photons are measured with an average resolution of roughly 1 mm in the small crystals and 1.8 mm in the large crystals.

Electrons from $K \rightarrow \pi e \nu$




Figure 4.6: The ratio of calorimeter energy to track momentum (“ E/p ”) for electrons from K_{e3} decays. The cutoff at the left edge of the plot is due to an online cut used to select candidates K_{e3} decays.

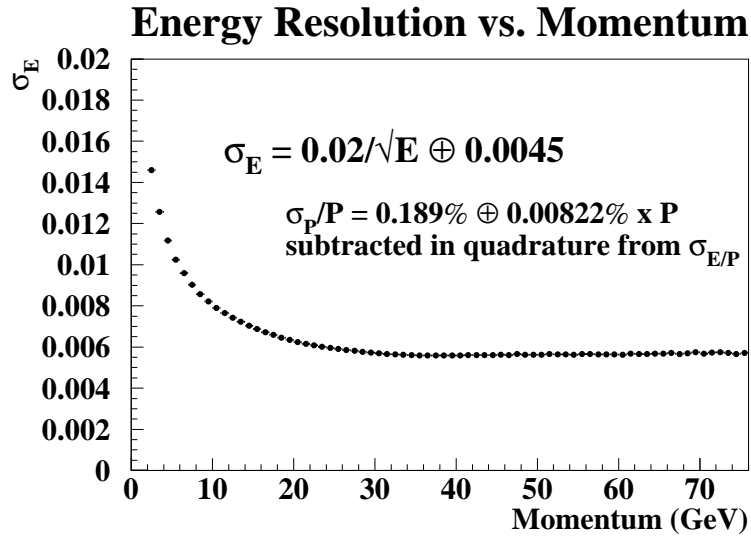


Figure 4.7: The resolution of the calorimeter as a function of energy, measured using electrons. The momentum resolution due to the spectrometer has been subtracted.

4.6 Trigger and Veto Counters

The trigger hodoscope is the primary component for the charged-mode trigger. It consists of two planes of plastic scintillator, called “V0” and “V1”, each 0.95 m square and divided into individual counters as shown in Figure 4.8. The counters are viewed from the top and bottom ends by PMTs attached to lucite light guides. The cracks between counters are in different locations for the two planes, so that a given particle cannot pass through cracks in both planes (except at discrete locations along the horizontal cracks). The beam holes, each 14 cm square, were added to the design after it was determined that a significant fraction of the radiation dose in the E731 lead glass was due to hadron interactions in the E731 trigger hodoscopes, which had no beam holes. The output pulses from the PMTs are fast enough to allow single-bucket timing.

There are several types of veto counters collectively referred to as “photon vetoes”. These include the “Ring Counters” (installed along the vacuum tank), “Spectrometer Antis” (at the outer edges of the drift chamber and the calorimeter), and




Figure 4.8: Geometry of the trigger hodoscope planes. The black boxes represent the beam holes.

the “Back Anti” (BA) (downstream of the beam holes in the calorimeter, just in front of the muon filter steel). All of these counters are made of alternating layers of lead and scintillator and are divided into segments with individual PMTs. The BA is also segmented longitudinally to help distinguish between electromagnetic and hadronic showers. For each photon veto counter, a signal above a certain threshold generates a digital pulse which is sent to the trigger system and may be used in the Level 1 trigger decision.

The Collar Anti is also intended to detect photons, but has a different construction and also has the important task of defining the inner aperture for photons landing near the beam holes of the calorimeter. The CA consists of two identical detectors around the two beam holes of the calorimeter. Each one is 1.5 cm wide to occlude the inner part of the crystals immediately surrounding the beam holes, as shown in Figure 4.9. Each half consists of layers of tungsten and scintillator, read out by long fibers going to 4 PMTs mounted outside the edges of the calorimeter.

4.7 The Trigger System

The beam buckets delivered by the Tevatron to the KTeV target every 19 ns contain, on average, about 3000 protons. However, only about one proton in 10^5 produces




Figure 4.9: Geometry of the Collar Anti counters around the beam holes of the calorimeter. The grid indicates the CsI crystals (2.5 cm square) while the black areas are the two parts of the Collar Anti.

a kaon which succeeds in passing through the collimation system and reaching the decay region, and only a fraction of these kaons decay in useful ways. The task of the trigger system is to detect these useful decays when they occur and to cause the detector information to be read out and written to tape.

The KTeV trigger system is described in detail in Chapter 3 of Ref. [65]. It is built in three “levels” with different strategies and deadtime implications. Level 1 uses fast synchronous signals from various detector elements and simple logic to evaluate the state of the detector for every RF bucket. Level 2 consists of a set of specialized electronic processors which evaluate the number and pattern of signals on individual drift chamber wires and calorimeter channels. For events which pass the Level 2 trigger, the detector information is transferred to computer memories and a “Level 3 trigger” process, running in parallel on 24 CPUs, reconstructs events and applies some loose kinematic cuts to select $\pi^+\pi^-$ and $\pi^0\pi^0$ candidates plus prescaled samples of other decay modes.

4.7.1 Level 1 Trigger

Several dozen boolean “source” signals are available for the Level 1 trigger decision. For example, one source indicates that at least two counters in the V0 plane of the trigger hodoscope were hit, while another indicates that the total energy deposit in the calorimeter (based on an analog sum of PMT dynode outputs) is above 26 GeV. These sources are formed and combined using mostly commercial electronics and simple logic. To keep the trigger rate at a manageable level, triggers may be inhibited by fast veto signals from the Regenerator, the MA, the photon vetoes, and the muon veto hodoscope.⁷ An “accept” decision by the Level 1 trigger initiates the digitization of various detector signals (ADCs, TDCs, and latches) and buffers the continuously-digitized calorimeter information in FIFOs. The raw Level 1 trigger rate is about 60 kHz under normal running conditions.

⁷Thus, while the Level 1 trigger decision is deadtimeless, accidental activity in various detectors can veto otherwise good events.

4.7.2 Level 2 Trigger

Only the Y views of the drift chambers are used in the Level 2 trigger decision; thus, the trigger is evaluated using the same drift chamber wires for decays in either beam. The hit-counting system, consisting of “Banana” and “Kumquat” modules, takes signals from each wire and counts the number of pairs of complementary wires (or isolated wires) in each chamber which were hit within an appropriate time interval.⁸ The hit-counting system takes around 800 ns to provide a result.

An additional Level 2 processor called the “Y Track Finder” (YTF) was added to evaluate whether the drift chamber hits found by the hit-counting system are consistent with two straight tracks. It is designed to be used for $\pi^+\pi^-$, and so it looks for one upward-going track and one downward-going track, with an overlap region in the central region of the drift chamber system which satisfies both the “up” and “down” requirements. It was built from commercial programmable-logic and memory-lookup modules. It uses coarse-grained outputs from the hit-counting system (each one an OR of typically 16 contiguous wires) as its inputs, and generally returns a decision a few hundred nanoseconds before the hit-counting system.

A “Hardware Cluster Counter” (HCC) [66] was built to count the number of “clusters” in the calorimeter, where a cluster is defined as a set of contiguous channels each with energy deposit over some threshold, nominally 1 GeV. Channels touching only at a corner are not considered contiguous. This processor is very useful at reducing the rate of the neutral-mode trigger, which at Level 1 just requires a certain amount of energy in the calorimeter and therefore is often satisfied by the copious K_{e3} and $3\pi^0$ decays. The HCC has the longest decision time of all the Level 2 processors, about 1.5 μ s.

If an event passes the Level 2 trigger, then digitization is allowed to finish and the front-end modules are read out, a process which takes around 18 μ s. If the event fails, then the front-end modules are cleared and the trigger system is re-enabled

⁸The Banana modules, used for chambers 1 and 2, were designed to be able to make a more sophisticated decision based on the correlated hit times on each pair of complementary wires—in essence, to trigger on good-sod pairs. However, that capability was not used when collecting data to measure $\text{Re}(\epsilon'/\epsilon)$ because of concerns about inefficiencies, so they behaved essentially the same as the Kumquat modules on chambers 3 and 4, which simply latched any hit within a gate.

after a few microseconds. Front-end readout and clearing together cause a fractional deadtime of about 35% under normal running conditions. The total Level 2 trigger rate is about 10 kHz.

4.7.3 Level 1 and Level 2 Trigger Definitions

Both the Level 1 and Level 2 control logic use programmable memory lookup units (MLUs) to make the trigger decision based on the inputs. A text configuration file is used to define symbols for logical combinations of low-level trigger inputs and to specify the trigger requirements for 16 trigger types; the file is parsed to generate the necessary memory maps to be loaded into the MLUs. We have adopted a philosophy of trying to make the triggers essentially 100% efficient for the $\pi\pi$ modes, to minimize possible biases on $\text{Re}(\epsilon'/\epsilon)$. Thus, for example, we define loose trigger-hodoscope, hit-counting and YTF trigger requirements to be insensitive to a single crack or inefficiency, and allow either 4 or 5 HCC clusters in the $\pi^0\pi^0$ to be insensitive to “hot” channels.

Besides the main $\pi^+\pi^-$ and $\pi^0\pi^0$ triggers, there are special triggers with looser requirements to permit study of inefficiencies in the nominal triggers. There are also extra triggers to collect $3\pi^0$ decays for detector studies and to select candidates for several rare kaon and hyperon decay modes. Finally, an “accidental” trigger is formed, using scintillation counters near the kaon production target, to randomly record the underlying activity in the KTeV detector with the same instantaneous-intensity distribution as the physics data.

In addition to the beam-related triggers, there is a separate set of hard-wired “calibration” triggers used mostly to monitor the response of the calorimeter to light from the laser calibration system and to cosmic-ray muons, and to measure “pedestals” (baseline digitized values reported by the ADCs and DPMTs when no energy has been deposited).

4.7.4 Level 3 Trigger

The detector data is read out into VME memory modules with sufficient capacity to hold the data from an entire 20-second spill. Therefore, the Level 3 computers have a whole minute to process the data without introducing any deadtime. A minimal set of kinematic and particle-identification cuts are applied to select $\pi\pi$ candidates, and the cuts are loose since event reconstruction had to rely on preliminary detector calibrations.

For the $\pi^+\pi^-$ trigger, track and vertex candidates are reconstructed first, and the approximate invariant mass is calculated for each vertex candidate. If any vertex candidate has an invariant mass of over $450 \text{ MeV}/c^2$, the calorimeter information is unpacked and clusters are found and matched to the tracks.⁹ The quantity E/p is calculated for each track and required to be less than 0.9 in order for the event to be tagged as a $\pi^+\pi^-$ candidate. Note that an “inclusive” approach is used: rather than trying to select the *best* vertex candidate, the event is accepted if *any* vertex candidate passes the cuts. Determination of the *best* vertex candidate is left for the off-line analysis, after final calibration of the detector.

Level 3 processing is somewhat more straightforward for the $\pi^0\pi^0$ trigger. The cluster-finding software must find four clusters corresponding to those found by the HCC. The $\pi^0\pi^0$ decay vertex is found by pairing the photon clusters to make two π^0 's (see Section 7.3), and the invariant mass is required to be at least $450 \text{ MeV}/c^2$. If the second-best pairing is also reasonably consistent with being from two π^0 's, then the invariant mass is calculated for that case as well, and the event is kept if *either* pairing yields a mass of at least $450 \text{ MeV}/c^2$.

Candidate $\pi^+\pi^-$ and $\pi^0\pi^0$ events are tagged as such and marked to be written to tape. In addition, the Level 3 software tags prescaled samples of $K_L \rightarrow \pi e \nu$, $K_L \rightarrow \pi^+\pi^-\pi^0$, and other decay modes from these triggers. Also, a fraction of all input events are tagged as “random accepts” regardless of the result of Level 3 processing. Any event tagged as part of one or more samples is written to tape. In

⁹There is not sufficient CPU time available to do clustering for all events, since clustering is much more time-consuming than tracking.

46

all, about 40,000 events are written to tape from each accelerator spill at nominal intensity.

CHAPTER 5

THE DATA

In this chapter, we will make some comments about data collection. In particular, we will spend some time discussing a number of serious and not-so-serious problems which were encountered. We will then make a few general remarks about the data-reduction process and identify the data samples used for the analysis presented in this thesis.

5.1 Data Collection

KTeV collected data during the 1996-1997 Fermilab fixed-target run. It took a great effort to install and commission the KTeV detector before the arrival of beam, with around-the-clock activity beginning months in advance. Stacking of the cesium iodide crystals finished on July 21, 1996. Beam was first delivered to the KTeV target on July 31, with the beamstop kept closed so that only muons were seen in the detector. After a month of occasional muon beam, the beamstop was opened on August 31, and kaon decays were observed for the first time with the KTeV detector. The first high-intensity data run with full-size beams was on September 5; studies of beam characteristics and trigger rates verified that the livetime of the detector was acceptable and that the yield of reconstructed kaon decays was good. There followed a 7-week period of intense detector debugging, with major overhauls to some systems, as described in Section 5.2. On October 24, all systems were judged to be working well, and the first physics-quality data was collected.

The remaining running time was divided between the two operational modes of KTeV: experiment E832, which is designed primarily to measure $\text{Re}(\epsilon'/\epsilon)$, and experiment E799, which studies a variety of rare kaon, π^0 and hyperon decays.

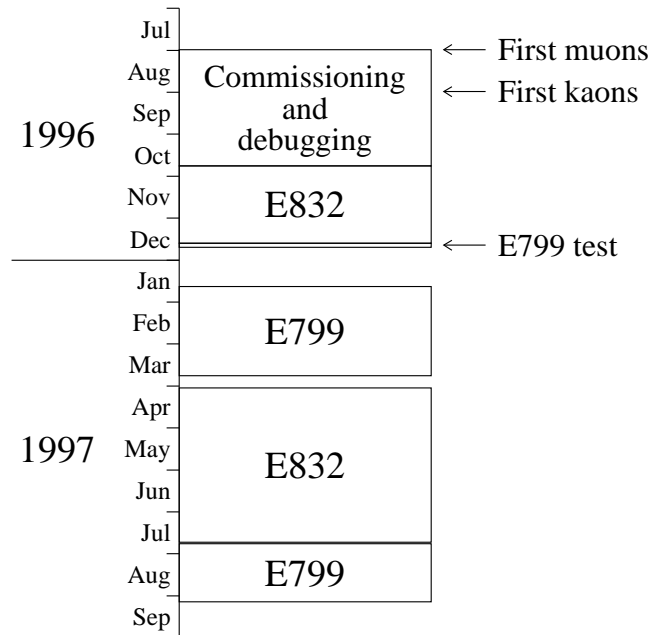


Figure 5.1: Division of running time during the 1996-1997 fixed-target run.

The division of running time is shown in Figure 5.1. The detector configurations differed somewhat for the two experiments: for E799 running, the beryllium part of the common absorber was taken out, the Regenerator and MA were swung out of the beams, and the eight transition radiation detectors (TRDs) were installed downstream of drift chamber 4. Also, new collimator inserts were installed for the second E799 running period to increase the size of the beams.

Due to a number of improvements, the data-taking efficiency was significantly higher for the second E832 running period than for the first. In fact, about 3.5 times as many good $\pi\pi$ decays were recorded even though the running time, in days, was only twice as long.

Data was collected around the clock, with three people on shift at all times. Shift activities included starting new “runs” (which could last up to ~ 10 hours, if the accelerator operated smoothly), monitoring the detector for various problems, resetting the trigger or readout electronics when there was a glitch, and calling in experts to handle any serious problems which occurred.

5.2 Problems Encountered

Any experiment as large and complex as KTeV is certain to have problems, from outright equipment failures to electronic glitches to mysterious effects which come and go without apparent cause. KTeV turned out to be especially prone to failures of the calorimeter readout electronics, although other custom electronics modules had their share of failures as well. On the whole, about 15–20% of otherwise good beam time was lost due to various problems with the KTeV beamline and detector. The Tevatron itself was down at least 10% of the time in a typical week, with occasional longer downtimes.

It would be difficult to try to catalogue all of the problems encountered; we will just mention the most important ones.

5.2.1 *Calorimeter Readout Problems*

After a long development process and a number of prototypes, the DPMT readout system succeeded in meeting the stringent resolution, linearity, and low-noise requirements for the KTeV calorimeter. However, once a large number of channels had been installed and operated for an extended period of time, problems of a different sort surfaced. DPMT boards began to fail in several different ways; the average lifetime for a single channel was a few years, but with a few thousand channels in the array, failures were happening at the rate of a few per day. One class of failures was determined to be due to a production problem with the DBC ASIC, and as a result the DBCs were remade (with design changes for added robustness against any repeat of the production problem) and replaced on the DPMT cards during the month of September 1996. Early in October, it was found that a large fraction of the QIEs (one production batch) occasionally indicated an incorrect range when operated at the desired frequency of 53 MHz. Therefore, during 1996 data collection the DPMTs were operated at “RF/3” (~ 18 MHz), so that each DPMT clock cycle integrated the signal over three beam buckets. The QIEs also failed outright in a number of different ways, most apparently due to the same production problem which affected the DBCs, necessitating frequent replacements during data collec-

tion. The QIEs were remade and replaced over the December-January shutdown, allowing the DPMTs to operate at 53 MHz during 1997 data collection. However, the new batch of QIEs still had similar production problems (recently demonstrated to be intrinsic to the particular chip fabrication process used), and chips continued to fail at the rate of about one per day throughout 1997.

Any serious DPMT failure was generally noticed within a few minutes by the monitoring software and triggered a recorded audible alarm. We generally replaced the DPMT immediately, which involved steering away most of the primary proton beam and putting in the beam stop to allow people to go into the CsI blockhouse. Thankfully, the actual replacement was simple and only took a few minutes. We followed this policy of quick replacement for the whole run, even though Monte Carlo studies indicated that the bias on $\text{Re}(\epsilon'/\epsilon)$ from a broken DPMT would be small, even if present for a long time period and ignored in the analysis.

5.2.2 *Calorimeter Trigger Problems*

While the calorimeter readout was the biggest crisis during detector commissioning, there were also major problems with the “E-Total” system, which uses fast analog signals from the calorimeter (taken from the last dynode stage of each PMT) to generate information for the trigger. As its name suggests, it performed an analog sum of the total energy deposit in the calorimeter, which was used in the Level 1 trigger. However, it also discriminated the signal from each of the 3100 calorimeter channels to produce 3100 bits of input to the HCC. In mid-September 1996, all of the E-Total boards had to be removed so that a set of leaky capacitors could be replaced; this was done at the same time as the DBC chips on the DPMTs were replaced. Then, in early October, it was discovered that the E-Total system was occasionally setting bits for channels with no energy deposit, causing the HCC to find extra clusters. The cause of these “hot bits” was determined to be due to comparators on the E-Total boards which had been damaged by high temperature when the air conditioning system failed on a few occasions. The worst comparators were identified and replaced over the next few weeks, and the system was declared operational on October 24; this defined the beginning of good physics running.

However, there were still some hot bits at a very low level throughout the rest of 1996, so the neutral-mode trigger was set up to allow an extra HCC cluster.

5.2.3 *Drift Chamber Problems*

The drift chambers had a number of problems. After they were installed in the KTeV Hall, it was found that the wires picked up noise which, amplified by the electronics, was large enough to be recorded as fake hits. In particular, there was substantial noise during the asynchronous readout of the digitized TDC information, when the detector should have been ready to record another event. A number of modifications were made to the preamplifier gains and to the grounding network which eventually reduced the noise to below the discriminator threshold.

During commissioning, there were a number of instances in which excessive current was drawn from the high-voltage supplies. A few such cases turned out to be due to discharges from wire stubs or other metal objects; cleaning the interiors of the chambers also seemed to help at one point. When high-intensity beam was delivered for the first time, there was a concern about how much current was being drawn, and the chamber voltages were reduced. Even so, there were several instances of unexplained current draw during data collection which were cured by cycling the high voltage. This current draw was generally at a low level and did not seem to adversely affect the drift chamber efficiencies.

On several occasions, one of the drift chambers would begin to “oscillate”, producing a continuous stream of fake hits on many of its wires. The chambers remained in this state until the high voltage was cycled. The oscillations seemed to be related to a feedback loop with a high-impedance pickoff of the discriminated signals, used by the trigger system, although it was never understood what caused the oscillations to begin.

The most insidious problem with the drift chamber system was not discovered until after the 1996 E832 running period. During studies of the Level 3 software for the E799 test period in December, several events were observed in which the track reconstruction code was unable to find two tracks even though the drift chambers appeared to have all the necessary hits. Closer inspection showed that in these cases,

one of the hit-pairs had a surprisingly high sum-of-distances, larger than the nominal 6.35 mm by more than the 1-mm window imposed by the track reconstruction code.¹ This window had been chosen based on the sod distributions seen in E773, which had a very small tail at high sod values, as shown in Figure 5.2; the equivalent KTeV distribution turns out to have a much larger tail, and as a result many events were being rejected online by the Level 3 filter code. The essence of the problem is that one of the wires records a hit time which is substantially delayed (by 20 ns or more) from the expected time; this causes the drift distance to that wire to be overestimated.

Once this “high-sod effect” was discovered, the tracking software was quickly modified to allow high-sod pairs to be used to reconstruct tracks. Fortunately, there was sufficient online computing power to handle the modest increase in CPU time needed for Level 3 processing. All E799 data, plus the 1997 E832 data, was collected with Level 3 using the modified tracking code, which was sensitive to high-sods and other track imperfections only at second or third order (depending on the positions and types of the imperfections).

Given the obvious difference in the figure, it may seem surprising that this problem was not discovered earlier. The primary reason is that sod plots like this are always made for hits on tracks, because the drift chambers naturally contain extra hit-pairs from out-of-time tracks which appear to have a low or high sod and must be excluded. Thus, any track which had a high sod was excluded from the sod plot until after the tracking code was modified. We could also have stumbled across the problem earlier by scanning events by hand, but it turns out that the magnitude of the high-sod problem was fairly small at the beginning of data collection, when the most attention was being paid to the performance of the detector and the reconstruction software; it only became large later on, after the chamber voltages were lowered (without lowering the discriminator thresholds to compensate for the reduced gain).

¹Low-sod pairs can easily result from a delta ray or an extra charged particle, so these were used in track reconstruction from the beginning; see Section 6.2 for more information on the tracking code.

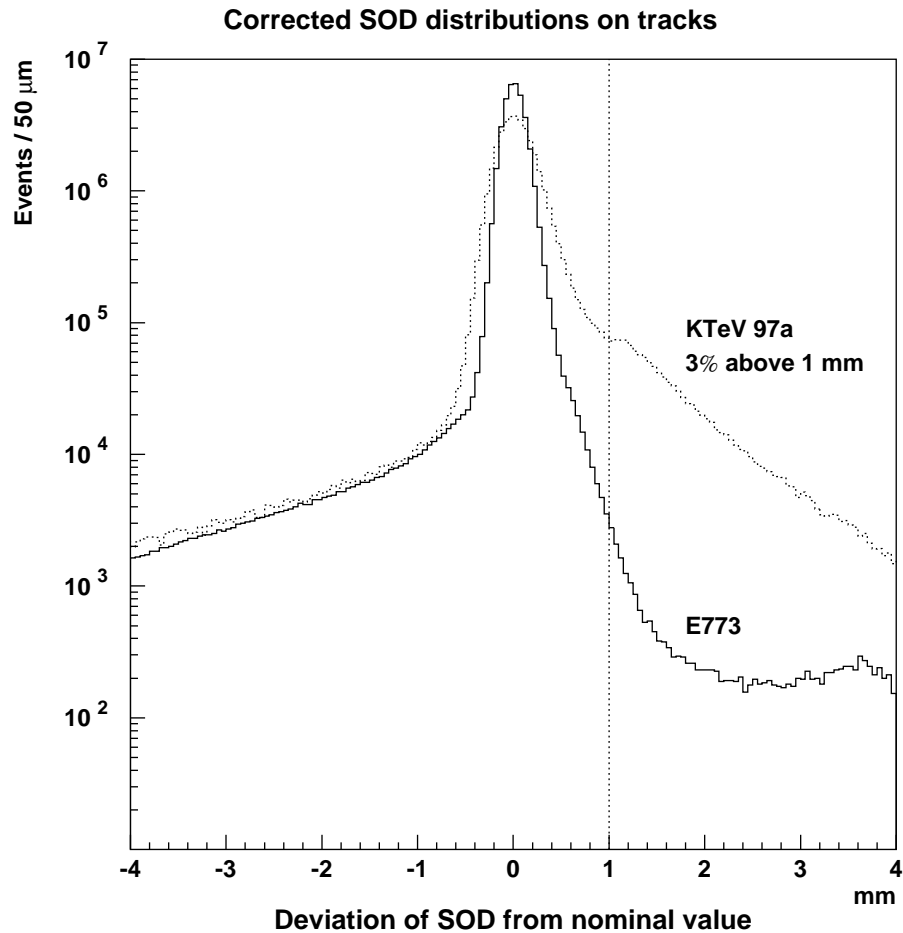


Figure 5.2: Drift chamber sum-of-distances distributions for the E773 and KTeV experiments. The Level 3 filter code used during 1996 data collection ignored all hit pairs with a sum-of-distances high by more than 1 mm (indicated by the vertical dotted line).

Because of the Level 3 filtering, the $\pi^+\pi^-$ samples written to tape during 1996 had already been subjected to the requirement that no chamber had a high sod. A study of $\pi^+\pi^-$ events in the random-accept sample revealed that, on average, the chance of any given hit time being mismeasured was less than 1%—but with 32 hits in each $\pi^+\pi^-$ event, a total of 22% of otherwise good events had been discarded! It also indicated (from the differential loss between the vacuum and regenerated beams) that the resulting bias on $\text{Re}(\epsilon'/\epsilon)$ was $(10 \pm 5) \times 10^{-4}$, where the error is

due to the limited number of $\pi^+\pi^-$ events in the random-accept sample. Because the random-accept sample also contained a much larger number of K_{e3} and $\pi^+\pi^-\pi^0$ decays, it was reasonable to suppose that we could use these decays to understand the high-sod effect in detail and simulate it; there were a sufficient number of tracks so that the bias on $\text{Re}(\epsilon'/\epsilon)$ could be understood, at least in principle, to the 10^{-4} level, which would be an acceptable systematic error contribution for the $\text{Re}(\epsilon'/\epsilon)$ measurement.

With this in mind, we embarked on a detailed study of the high-sod effect in the 1996 data, using the updated tracking code to find tracks whether or not they had a high sod. A number of basic characteristics soon became evident. First, the $\pi^+\pi^-$ event loss probability varied dramatically as a function of time, from 7% at the beginning of physics data collection, to a high of 31% in the middle of the 1996 running period, and back down to 21% at the end of the 1996 running period. The vacuum- and regenerated-beam loss probabilities tracked each other very well. The biggest change in the loss probability occurred when the voltages of the drift chambers were lowered from 2500 to 2450 V. The remaining time dependence showed some correlation with atmospheric pressure. There was also some dependence on beam intensity, though the loss probability changed by only $\sim 30\%$ of itself between nominal intensity and very low intensity, so the bulk of the effect is independent of intensity. The intensity dependence is greatest in the regions of the drift chambers where the beams pass through, which is a major concern because the ambient particle flux is higher in the vacuum-beam region than in the regenerated-beam region. In fact, this effect seems to be responsible for much of the observed $\text{Re}(\epsilon'/\epsilon)$ bias.

It was established early on that most instances of the high-sod effect occur when a particle passes close to a wire (within a few hundred microns), and it is the *closer* wire which mismeasures the drift time. This suggested a mechanism which we now believe to be the correct explanation for the high-sod effect: the total gain of the chamber gas and electronics was too low for a single drift electron to reliably cross the discriminator threshold. When a particle passes far from a wire, several drift electrons reach the wire at about the same time, giving a high, short pulse; whereas, when a particle passes close to a wire, the drift electrons tend to dribble in, giving

a low, long pulse. (Figure 4.3 shows these two cases.) In particular, the first few drift electrons may arrive one at a time and stay below the discriminator threshold, and then two or more may arrive simultaneously and trigger it. The result is that a hit is recorded, but not from the closest drift electrons.

This interpretation explains why the high-sod effect got worse when the chamber high voltages were lowered: the gain dropped as a result, so the chambers became even less efficient for a single drift electron.

The probability of mismeasuring a hit distance varies significantly from wire to wire, presumably due to small gain or threshold variations from channel to channel. It also varies along the length of a given wire. Intriguingly, many wires have one or more spots where the high-sod probability is *very* high, which we have named “freckles.” The freckles are highly localized, with sizes of order a few millimeters, and generally persist throughout the whole 1996-1997 run. They are thought to be related to tiny black fibers which have been found on the wires. Usually the wire also has a significant inefficiency (*i.e.* a chance that no hit at all will be recorded) at the same spot, supporting the idea that a freckle is a place where the wire has a very low effective gain.

We used the K_{e3} and $\pi^+\pi^-\pi^0$ decays in the random-accept sample to map out the time and position dependences of the high-sod probability (with separate maps for the two Regenerator positions) and added these effects to the drift chamber simulation in the Monte Carlo. The simulation reproduced the basic characteristics of the 1996 high-sod loss, including the vacuum-vs.-regenerated-beam difference, but we could not be certain that the agreement was quantitative and did not see how to evaluate the systematic error on $\text{Re}(\epsilon'/\epsilon)$. Therefore, we decided *not* to use the 1996 $\pi^+\pi^-$ data for the $\text{Re}(\epsilon'/\epsilon)$ measurement, but to use $\pi^+\pi^-$ data from the early part of the 1997 running period.

The high-sod effect still occurred during the 1997 running period; in fact, it was substantially worse in the beam regions of the drift chambers due to “damage” from the high kaon and neutron fluxes experienced during E799 running. However, the tracking code was no longer sensitive to this effect at first order, and so the potential bias on $\text{Re}(\epsilon'/\epsilon)$ was smaller to begin with. We still need to understand higher-order

losses, and so the mapping procedure mentioned above was refined further and made into a standard part of the Monte Carlo simulation, as described in Section 8.3.

Even though the 1996 $\pi^+\pi^-$ data is no longer used directly in the $\text{Re}(\epsilon'/\epsilon)$ measurement, we still use it to characterize kaon scattering from the Regenerator during the 1996 running period, which is essential to subtract backgrounds from the $\pi^0\pi^0$ samples. The Monte Carlo includes the full simulation of position-dependent inefficiencies and the high-sod effect, and is believed to be sufficiently accurate for this purpose. Even so, the systematic error assigned due to the $\pi^0\pi^0$ background subtraction will include a component from the uncertainty in the 1996 $\pi^+\pi^-$ acceptance.

5.2.4 *Miscellaneous Other Problems*

A wide variety of problems were encountered and dealt with while commissioning the detector, and we will not try to list them here. We will, however, mention a few things which came up during otherwise good running.

The digitized information from the various parts of the detector was transmitted to the computer memories along a rather convoluted path. The calorimeter information was assembled asynchronously by a “Crate-Trigger Interface/Readout Controller” (CTI/RC) in each crate of pipeline modules, using a custom backplane. For all other systems, readout within a crate was based on the “FERA” readout protocol developed by LeCroy for the ADC modules of the same name; it relied on front-panel cables which tended to be fragile, especially in crates where modules of different types were in the same readout chain. In all cases, the data from each crate was transmitted to the computer memories along ribbon cables using the RS-485 protocol, but because total distances of a few hundred feet were involved, it was necessary to add “repeaters” at several places to retransmit the digital information. All of these systems were subject to occasional glitches, which generally either froze the data-acquisition system or else caused errors when events were unpacked, and required the system to be reset. Monitoring software was developed and improved (*e.g.* by adding audible messages) over time to alert the people on shift to problems as they occurred. These readout hang-ups, like hang-ups of the trigger system

(which also occurred sometimes), are not a concern for the $\text{Re}(\epsilon'/\epsilon)$ measurement because they cause identical losses for decays in the vacuum and regenerated beams.

The machinery which moved the Regenerator back and forth between the two beams became increasingly sticky during the 1996 running period. For a while, the frequency of motion was changed from once per minute to once every three minutes in order to prolong its life. Even so, there were a few times at which it got stuck. This could potentially have biased $\text{Re}(\epsilon'/\epsilon)$, by causing us to assign events to the wrong beams. However, the two positions of the Regenerator (and of the movable absorber, which must remain synchronized) were independently sensed by separate microswitches and read out with every event. Thus it was easy to discard all events during the periods when the Regenerator mover was stuck.² The mechanism was redesigned over the winter shutdown to be more reliable.

The decay volume was kept evacuated by two large diffusion pumps and associated equipment, which should have been capable of keeping the pressure at 10^{-5} torr or better. However, during the 1996 running period, the pressure was around 2×10^{-4} . While higher than expected, this was still too low to be an issue for the $\text{Re}(\epsilon'/\epsilon)$ measurement. During the winter shutdown, the vacuum window was found to have a small tear near the edge and was replaced (as planned in any case). After that, the pumps were able to keep the pressure below 10^{-5} torr, except for a few occasions when the vacuum system controller failed and had to be restarted.

5.3 Data Reduction

KTeV generated a very large amount of data. For E832 alone, over five billion events were written out onto about 3000 DLT tapes (nominal capacity of 10 GB before compression), with a total data stream size of about 35 terabytes. Accessing the data presented a major challenge, since obviously it was not practical to read through all 3000 tapes multiple times for the various analyses to be performed. It was decided to first “split” the data into separate datasets according to the sample

²It was also easy to identify these periods by comparing the number of K_{e3} decays in each beam during a given spill.

tags determined by the Level 3 processing; for example, all events tagged as $\pi^+\pi^-$ candidates were copied to a set of 277 tapes, while the $\pi^0\pi^0$ candidates occupied 116 tapes. The large samples of K_{e3} , $\pi^+\pi^-\pi^0$ and $3\pi^0$ candidates were also separated out. During the copying procedure, the calorimeter and HCC information was “squeezed” into a more compact format, reducing the average event size by about 40% without any loss of information.

These datasets were still fairly unwieldy, and so the next step was to “crunch” the data, applying some loose analysis cuts and reducing the number of candidate events by a factor of a few. (The cuts used will be described later.) Finally, we processed the crunch output and recorded high-level reconstructed quantities (track trajectories, cluster energies and positions, veto-counter energies, etc.) in a flexible “data summary tape” (DST) format. This multi-tiered approach provided optimal datasets for different aspects of the analysis. For example, one could modify the analysis cuts and reprocess the DST sample in less than a day; this was done many dozens of times during the course of the analysis. If the event reconstruction software was modified, one could go back and process the crunch output in a matter of a few days. The crunch itself was redone twice for the $\pi^+\pi^-$ data and once for the $\pi^0\pi^0$ data, in each case to loosen some cuts and to take advantage of improvements in the event reconstruction software.

It is important to note that vacuum- and regenerated-beam events were not separated at any time during the data reduction process. This removed the possibility of a $\text{Re}(\epsilon'/\epsilon)$ bias due to a damaged tape or disk file.

5.4 Data Samples Used for This Analysis

It was decided, long before data collection had even begun, to first measure $\text{Re}(\epsilon'/\epsilon)$ from a fraction of the data, nominally 20%. This would provide a measurement with higher precision than either E731 or NA31 in a (hopefully) timely fashion. (The time required is not determined by the processing of the $\pi\pi$ signal modes, but by the amount of time and effort required to calibrate the detector and to reduce/understand all the systematic issues to a level comparable to the statistical

error.) It seemed natural to use the data collected in 1996 for this measurement, corresponding to 23% of the total. However, after detailed study of the high-sod problem in the 1996 $\pi^+\pi^-$ data, we decided to use $\pi^+\pi^-$ data from 1997 for the $\text{Re}(\epsilon'/\epsilon)$ measurement. The first 18 days of the 1997 running period (ending with a downtime due to a failure of the vacuum pump control system) had nearly as much good data as all of 1996, so we decided to use the $\pi^+\pi^-$ data from that period, giving it the name “97a”.

The fact that we use $\pi^+\pi^-$ and $\pi^0\pi^0$ data from different running periods opens up some additional systematic issues, but we will see in Chapter 10 that they do not contribute much to the total systematic error.

CHAPTER 6

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

This chapter will describe all aspects of the selection of the $K \rightarrow \pi^+\pi^-$ samples, starting with the hardware trigger requirements. We will then describe how the pion “tracks” through the spectrometer are reconstructed, and will describe in some detail the event-selection criteria (“cuts”) which define the final sample of $\pi^+\pi^-$ candidates. We will end with a description of the (small) backgrounds in the samples, and the procedure for subtracting them.

6.1 Trigger Requirements

The goal of the charged-mode trigger is to select events with two charged particles. The main requirement is that there be at least two hit counters in the V0 plane and at least one in the V1 plane, or vice versa; this loose requirement makes the trigger largely insensitive to the cracks between the counters. We also impose a simple topological requirement: among the hit counters, one (or more) must be in the upper set of counters, one in the lower set, one in the “east” region, and one in the “west” region. The east and west regions overlap between the beam holes; a counter in the overlap region satisfies both requirements. The upper and lower regions also overlap as a natural result of the counter geometries in the two planes (see Figure 4.8). Because the desired $\pi^+\pi^-$ events have no net transverse momentum, they always satisfy the east-west-up-down trigger requirement. Decays from scattered kaons, as well as some K_{e3} and $\pi^+\pi^-\pi^0$ decays, tend to fail this requirement, reducing the trigger rate.

A kaon may decay downstream of drift chamber 1, leaving signals in the trigger hodoscope but not allowing us to reconstruct complete tracks. The frequency of this

happening is actually rather large. Therefore, we instrumented the X and Y views of chambers 1 and 2 with fast “DC-OR” electronics, which indicate whether any wire in a given chamber view was hit in a 90-ns time interval. We require that at least three out of the four instrumented views have at least one such hit. This reduces the Level 1 trigger rate, lessening the deadtime incurred from making Level 2 trigger decisions, without any significant loss of signal.

The charged-mode trigger was vetoed if there was significant energy in the Regenerator, the Spectrometer Antis, or the muon hodoscope.

At Level 2, a hit-count requirement was imposed only on the Y views of the drift chambers, using the counts returned by the Banana and Kumquat modules. The event was required to have at least two hits in any three out of the four chambers, and at least one hit in the remaining chamber. This loose requirement made us insensitive to inefficiencies at first order. The trigger also required that the Y Track Finder find one track in the upper half of the chambers, and one track in the lower half (with an overlap region somewhat larger than the size of the beams). One of the tracks was allowed to be missing a hit in chamber 1 or 2 (but not in chamber 3 or 4).

As described in Section 4.7.4, the Level 3 filter process reconstructed events and required that the $\pi^+\pi^-$ invariant mass be greater than $450 \text{ MeV}/c^2$ and that the E/p be less than 0.9 for both tracks.

6.2 Track Finding

The offline analysis of $\pi^+\pi^-$ candidates begins by finding tracks in the spectrometer.¹ The track-finding algorithm consists of several steps, which we will now summarize.

The first step is to unpack the drift chamber TDC information and convert each hit time to a drift distance according to a relationship measured from the data. We then classify each hit according to whether there is a hit on an adjacent wire in the complementary sense plane, and if so, according to the sum of the drift distances for

¹This was done during Level 3 processing, but is repeated from the beginning during the offline analysis after final calibration of the detector.

the two wires (sum-of-distances, or “sod”). For any given wire, only the earliest hit within the in-time window is considered for forming tracks.² Ideally, a track would always leave one hit per plane, and pairs of hits in complementary planes would always have a good sum-of-distances. This does, in fact, happen most of the time, but the tracking code needs to have high efficiency for finding tracks even when there are imperfections, and since an ordinary two-track event leaves a total of 32 hits, even rare effects can show up in a significant fraction of the events. Examples of imperfections include:

- A track emits a “delta ray” (an energetic electron knocked out of an atom of the drift chamber gas) which travels toward the wire, ionizing the gas as it goes. The ionization from the delta ray can reach the wire before the ionization from the original track, causing us to underestimate the drift distance. This leads to a low sod for the pair.
- Two particles (both from the in-time bucket, or else one from an “accidental” particle in another bucket) pass through the same drift cell or adjacent drift cells. This too will lead to a low sod.
- As described in Section 5.2.3, the drift time was occasionally mismeasured due to the chambers having some inefficiency for registering a single drift electron. This leads to a high sod. (During the 1996 running period, the Level 3 software discarded all such hit pairs at this stage.)
- The wires had some absolute inefficiency, generally a fraction of a percent. Thus a track might leave a hit in only one plane in a given pair.

“Good” sods are defined to be those within ± 1 mm of the cell size (± 1.5 mm in chambers 3X and 4X, for which larger track angles smear out the sods more). As the hits are classified, the software makes a list of space points to be used to form

²There could, in principle, be useful information from additional hits on the wire. However, extra “afterpulse” hits are often produced from the ionization trail from a single particle, and may fool us. This process is not simulated in the Monte Carlo, so we prefer to ignore all additional pulses in the analysis for consistency.

track candidates. Low- and high-sod pairs have *two* entries, one for each hit, since any given track will only use one of the two hits. Isolated single hits also have two entries, one for each side of the wire.

Most events are fairly clean, with only a handful of extra in-time hits in the drift chambers. However, some events have a great deal of extra activity, occasionally with 100 or more extra in-time hits in the drift chambers. In such cases, the tracking code may have to consider a very large number of track candidates. To keep the combinatorics under control, the tracking code checks how many space points are found for each view of each chamber; if this number is more than 32, then it discards one or more classes of low-quality space points (*e.g.* single hits with out-of-time partners) to bring the number back below 32. Studies have shown that only a small number of events are lost due to this limitation and that it does not lead to a bias on $\text{Re}(\epsilon'/\epsilon)$.

The next step is to find track candidates in the X and Y views. This involves looping over the list of space points in each chamber. In the X view, a bend at the mid-plane of the magnet is permitted, as long as the upstream and downstream segments meet at the magnet within 6 mm. In the Y view, all four space points must lie near a single straight line, as determined from a fit. Certain requirements are made on the classifications of the hits in order to reduce the number of fake track candidates found; for examples, a Y-view track candidate must have at least two good-sod pairs. All track candidates which satisfy the requirements are kept at this point, without requiring them to be mutually exclusive. Therefore, there will often be more than two track candidates in one or both views.

Because the KTeV drift chambers have only horizontal and vertical wires, there is no way to associate a given X track candidate with a Y track candidate without using additional information, even if there are only two X track candidates and two Y track candidates. We resolve this ambiguity using clusters in the calorimeter; specifically, we require a given X-Y track candidate combination to point within 7 cm of a cluster. We then make a list of two-track vertex candidates for which the tracks are consistent with originating from a common point within the vacuum

decay volume, allowing a “slop” of 2 mm in the position of each space point. For a given vertex candidate, the two tracks are not allowed to share any hits.

Now that we have associated track candidates in the X and Y views, we can apply corrections to the hit positions due to drift chamber rotations and to the finite propagation speed of signals on the wires. We also recalculate the sum-of-distance for each pair, correcting for the track angle. For tracks which have a pair with corrected sod more than 0.6 mm low, we interpret this as a delta ray and choose the hit which is more consistent with the extrapolation from the other chambers.

In the X view, we adjust the track positions at chambers 2 and 3 to correct for the (momentum-dependent) horizontal displacement resulting from the fringe field of the analysis magnet. In the Y view, we re-fit the upstream and downstream segments separately, no longer assuming the track to be straight in the Y view. This is important because the analysis magnet causes “vertical focusing” (bending and displacement in the Y view) due to the three-dimensional shape of the magnetic field, so the tracks are not expected to be exactly straight. We use the offset of the upstream and downstream segments at the mid-plane of the magnet to calculate an “offmag χ^2 ” for each track. The momentum of each track is recalculated using the field integral measured at the place where the track passes through the analysis magnet.

6.3 Event Reconstruction

We are finally ready to select the *best* vertex from among the candidates we have found. For each candidate, we determine the vertex position (*i.e.* the position of the kaon decay) by minimizing the “vertex χ^2 ”, which is calculated from the hit positions in drift chambers 1 and 2 using the position resolutions and the expected amount of multiple scattering in the vacuum window and chamber 1. The “figure of merit” used to compare vertex candidates is a combination of the vertex χ^2 with the offmag χ^2 ’s for the two tracks, plus a penalty if there are hits missing on either of the tracks. Besides the two tracks which make up the vertex, we see if the leftover hits in the drift chambers can be used to form any additional tracks.

Figure 6.1 shows a reconstructed $\pi^+\pi^-$ event. The tracks originate from a vertex in the vacuum beam, so this is a K_L decay. This particular event has one pion which is minimum-ionizing in the calorimeter (depositing 0.36 GeV) and one which makes a shower with $E/p = 0.277$.

The two most important kinematic variables are the invariant mass $m_{\pi\pi}$ and the “ p_T^2 ”, defined as the square of the magnitude of the component of the momentum of the $\pi^+\pi^-$ system which is perpendicular to the original kaon direction. For a real $K \rightarrow \pi^+\pi^-$ decay, whether from a K_L or from a coherently-regenerated K_S , the invariant mass should be close to the known K mass ($497.67 \pm 0.03 \text{ MeV}/c^2$) while the transverse momentum should be close to zero. For technical reasons, we always calculate the transverse momentum assuming that the kaon scattered at the Z position of the downstream end of the Regenerator.³ That is, referring to Figure 6.2, we calculate $|\vec{p}|^2 \sin^2 \phi_R$ rather than $|\vec{p}|^2 \sin^2 \phi_V$. The difference is fairly minor, but it is important to note that we calculate p_T^2 in the same way for decays in the vacuum beam as well.

There is a slight subtlety in the way we assign an event to one beam or the other. It is possible for a kaon to scatter in the Regenerator but then cross over to the vacuum-beam side of the decay volume before decaying. We would like to assign such events to the regenerated beam so that the measured p_T^2 distribution reflects the full range of scattering in the Regenerator. Therefore, for an event with a decay vertex downstream of the Regenerator, we project the total momentum vector back to the Z position of the Regenerator and assign the event to one beam or the other based on the projected X position.

6.4 Analysis Cuts

A wide variety of analysis cuts are applied to select the final samples. Most of these cuts are intended to reduce backgrounds or to more cleanly define the acceptance.

³The majority of $K \rightarrow \pi^+\pi^-$ candidates in our event sample with $p_T^2 > 0$ are from kaons which scattered in the Regenerator; our definition of p_T^2 allows us to directly relate the observed distribution to the physical scattering processes.

KTEV Event Display

```
/usr/kpasa/data05/shawhan/kz
c074_some.dat
```

```
Run Number: 9097
Spill Number: 210
Event Number: 40474649
Trigger Mask: 1
All Slices
```

Track and Cluster Info

```
HCC cluster count: 4
```

```
ID Xcsi Ycsi P or E
```

```
T 1: -0.1418 0.3870 -20.05
```

```
C 3: -0.1379 0.3855 0.36
```

```
T 2: 0.3079 -0.3072 +28.25
```

```
C 2: 0.3248 -0.2828 7.83
```

```
C 1: 0.2993 -0.4046 1.84
```

```
Vertex: 2 tracks
```

```
X Y Z
0.0952 -0.0146 144.934
```

```
Mass=0.4970 (assuming pions)
```

```
Chisq=0.58 Pt2v=0.000001
```

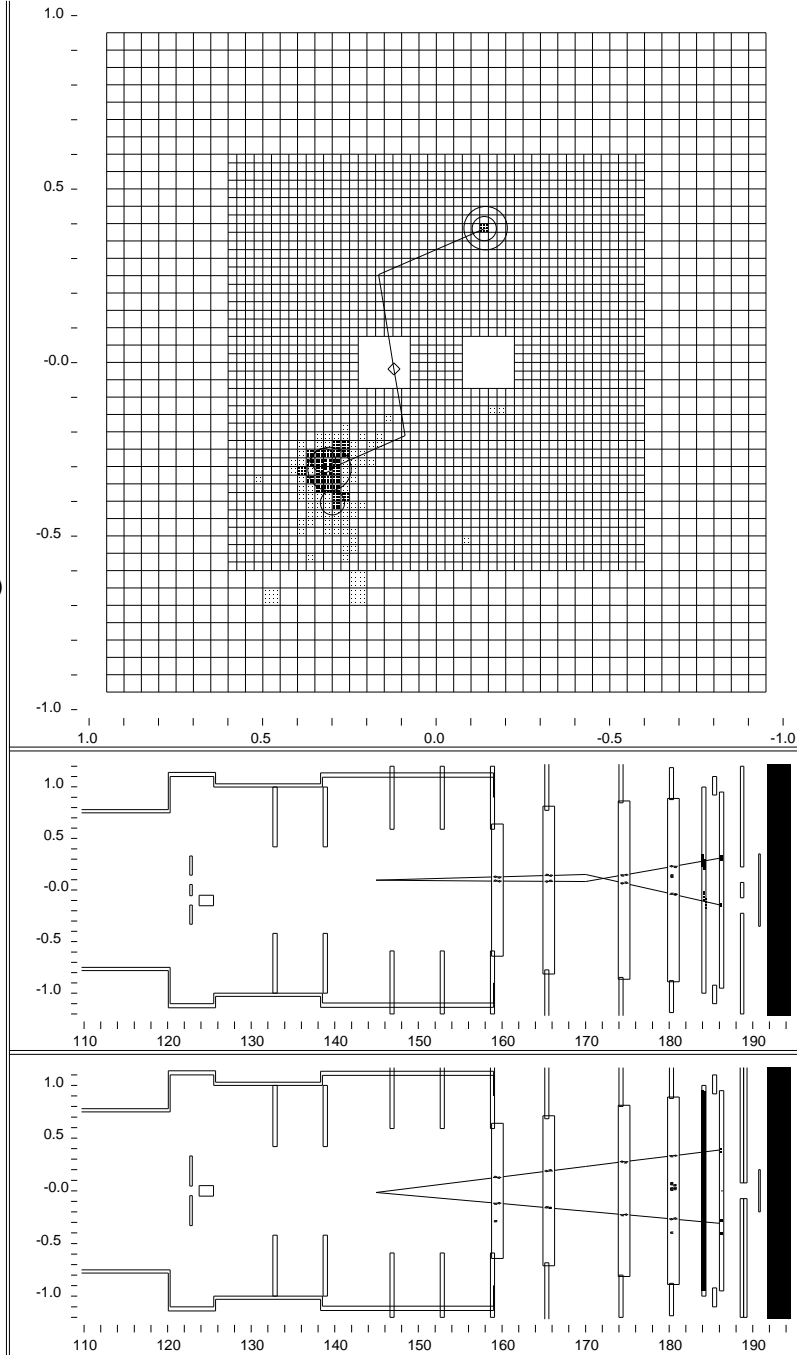


Figure 6.1: Event display for a $K_L \rightarrow \pi^+\pi^-$ event. The bottom two panels show the reconstructed tracks in the X and Y views. The top panel shows the tracks projected onto the face of the calorimeter. The tracks originate from a vertex in the vacuum beam (marked with a diamond in the top panel), pass through the spectrometer (bending in the X view at the magnet), and match up with clusters in the calorimeter.




Figure 6.2: Definition of angle used to calculate p_T^2 . The angle ϕ_R is used instead of ϕ_V .

For instance, some cuts are made to avoid topologies with poor reconstruction efficiency or situations which are not reliably simulated by the Monte Carlo.

6.4.1 Data Quality

A number of data-taking periods were identified as being “bad” based on equipment failures, poor beam quality, etc. Most of these were the result of a DPMT failure or some other problem with the calorimeter readout electronics. All events collected during these “bad” periods (a few percent of the 97a dataset) were discarded in the analysis.

6.4.2 Trigger Verification

It is possible that an event could have passed the trigger due in part to accidental activity in the detector. Therefore, in the offline analysis, we verify that the reconstructed tracks hit enough counters in the trigger hodoscope to satisfy the trigger requirement and that they have the proper east-west-up-down configuration. Only a tiny fraction of otherwise good events fail this cut. We do not attempt to verify the Level 2 requirements (hit-counting and YTF). In Chapter 10 we will evaluate the possible systematic error on $\text{Re}(\epsilon'/\epsilon)$ arising from the trigger.

6.4.3 Veto Counters

We removed events for which one or more counters were hit in the muon veto hodoscope. If the trigger electronics had worked perfectly, we would not have expected any such events, but in fact these make up 0.27% in the vacuum beam, and 0.11% in the regenerated beam, of events which pass all other cuts. The higher percentage in the vacuum beam indicates that real $K_L \rightarrow \pi\mu\nu$ decays are being removed by this cut.

The dominant background in the regenerated beam is from kaons which scatter in the Regenerator before decaying to $\pi^+\pi^-$. To minimize this background, we cut tightly on energy deposit in the Regenerator. The main requirement is that no Regenerator channel can have a measured energy greater than the equivalent of two minimum-ionizing particles (“mip”). The last Regenerator module is a special case, since the lead pieces cause additional scattering and there is less scintillator to detect recoil products. This module also has the highest total rate of energy deposit, so the ADC readout suffers from significant baseline shifts. We correct for this by subtracting the signal integrated by a separate ADC with a gate interval *before* the in-time bucket, and make a cut at 0.7 mip.⁴ Finally, we discard all events for which the Regenerator trigger sources fired in the previous bucket or the following bucket; this ensures that the veto logic works properly and that the ADC readings accurately reflect the energy deposit in the in-time bucket.

We cut on energy deposit in the “outer” veto counters: the Mask Anti, Ring Counters, and Spectrometer Antis. This is not strictly necessary for the $\pi^+\pi^-$ analysis, since none of the major background processes involve extra particles, but is done to be consistent with the $\pi^0\pi^0$ analysis. The particle rate in the MA was so high that the ADC readout had a large amount of common-mode baseline shift, making it impossible to cut cleanly on the ADC values themselves. We got around this problem by subtracting the minimum ADC channel from the maximum channel,

⁴This tight cut also evens out an up-down asymmetry in the online rejection of “Regenerator junk”, a minor background arising from inelastic hadronic interactions in the lead pieces. One specific process which has been observed in the data is the production of a $K^*(892)^0$ particle in the Regenerator, followed by its decay to $K^\pm\pi^\mp$.

canceling the baseline shift to a large extent and allowing us to make a tight cut at the energy deposit equivalent to a 100 MeV photon.⁵ We also discarded events for which the MA trigger source was on. The ADC channels for the Ring Counters also had baseline shifts, so we cut only on the trigger sources, which are immune to this noise. The Spectrometer Antis were quiet enough that we were able to cut on the ADC values, at an energy equivalent to a 300 MeV photon.

It is important to note that cuts on these veto counters, even tight ones, cannot bias $\text{Re}(\epsilon'/\epsilon)$ because any real event leaves no energy in these counters.⁶ This is true even for the Regenerator, since the process of coherent regeneration involves no measurable momentum transfer. Therefore, cuts on veto counters affect the vacuum-beam and regenerated-beam samples identically.

6.4.4 *Extra-Particle Cuts*

Events are rejected if any tracks are found other than the two which come from the vertex. No cut is made on the presence of extra clusters in the calorimeter, since a pion shower often produces one or more extra “satellite” clusters.

6.4.5 *Track Quality Cuts*

The offmag χ^2 's for the two tracks, along with the vertex χ^2 , help to indicate whether “good” tracks have been found. A high value of one or both of these generally indicates that the event has been misreconstructed, for example, that one of the hits used to make a track does not really belong to that track. The data exhibits long tails in these distributions, although these tails are nearly identical

⁵This rather tight cut was chosen early on to have a high efficiency for vetoing events in which a pion passed through the MA as a minimum-ionizing particle, a process which was not reliably simulated in the Monte Carlo. Later on, we tightened the cut on the projected track positions at the MA, eliminating such events from the sample in a more robust way. Therefore, the MA energy cut probably did not have to be this tight after all.

⁶One must be more careful with veto counters near the calorimeter, which may receive some stray particles from pions showering in the CsI; thus no cuts are made on the CA or BA energies in the $\pi^+\pi^-$ analysis. A very small amount of “splashback” energy is seen in the Spectrometer Anti at the outer edge of the CsI, but it is below the level of the cut.

for the vacuum and regenerated beams, as shown in Figure 6.3 for the vertex χ^2 . Studies have shown that these tails are mostly due to a combination of accidental hits and the high-sod effect described in Section 5.2.3. We simulate both effects in the Monte Carlo but do not want to have to rely heavily on our model of the high-sod effect, so we make only *very* loose cuts on the offmag and vertex χ^2 's, at values of 500 and 100, respectively.




Figure 6.3: Vertex χ^2 distributions for the data after all other analysis cuts have been applied. The vacuum-beam (solid) and regenerated-beam (dashed) distributions are nearly identical. However, because of the long tail, we have chosen to make an extremely loose cut, at a χ^2 value of 500.

6.4.6 Aperture Cuts

The KTeV detector has a number of physical apertures which limit the acceptance for $\pi^+\pi^-$ decays. However, there is a limit to the accuracy with which we know the positions of these apertures, and in any case the edges of photon-veto counters are not very sharply defined for pions, which can pass through a significant amount of material without being destroyed. Therefore, we prefer to cut away from these apertures based on projected track positions, so that we do not need to know the exact positions of the physical edges. For convenience, the cuts are defined relative to the measured positions of the apertures, but the only important thing is that the cuts are the same for the data and Monte Carlo samples. This approach does make us somewhat sensitive to our understanding of the spectrometer pointing resolution; we will evaluate the resulting systematic error on $\text{Re}(\epsilon'/\epsilon)$ in Chapter 10.

We require tracks to miss the Collar Anti by at least 2 mm, and to be at least 2.9 cm from the outer edge of the CsI. Tracks must also be at least 2 mm away from the edges of the beam holes in the trigger hodoscope, although most events which fail this cut also fail the CA clearance cut. For a decay upstream of the MA, the tracks must pass through the MA holes at least 3 mm from the edges. (The pointing resolution is poor this far upstream—roughly 2 mm—but it is still better to make this cut than to rely on the physical edge to define the acceptance.) We also discard events if one of the tracks passes “too close” to the edge of a drift chamber, where the wires are less well calibrated and characterized; the exact position of the cut varies from chamber to chamber.

If two tracks are close together in a drift chamber—specifically, in the same drift cell or in adjacent cells—then a hit from one track can obscure a hit from the other track on the same wire. This, by itself, will not cause the tracking code to fail, unless the tracks are close in three or more chambers or are closer than about 1 mm in any one chamber. However, it makes the tracking code more sensitive to additional imperfections in the drift chamber information due to accidental hits, delta rays, or the high-sod effect. We simulate such effects in the Monte Carlo, but do not want to have to rely on our models any more than necessary. Therefore, we require a minimum separation between the tracks, in both the X and Y views,

at every drift chamber. The cut is implemented by assigning each track to the closest wire and requiring these wires to be at least 3 cells apart. Thus the “edge” of the cut is always at the center of a drift cell, where the efficiency is the highest and the geometrical symmetry minimizes any possible systematic bias. This is the most severe cut in the entire $\pi^+\pi^-$ analysis, removing 15.57% in the vacuum beam, and 9.24% in the regenerated beam, of events which pass all other cuts. The large difference between these percentages is because the vacuum beam has many more downstream decays, which tend to have close tracks at the first drift chamber. In fact, the track separation cut has a very significant effect on the shape of the vertex Z distribution at the downstream end of the decay volume.

Finally, we require the the projected track positions at the CsI to be separated by at least 3 cm in the X view and 3 cm in the Y view. This minimizes the chance of the X-view tracks being associated with the wrong Y-view tracks, which relies on matching to clusters in the calorimeter; pion clusters (either hadronic showers or minimum-ionizing energy deposit in a single CsI block) do not give a very accurate position measurement. Mismatching the X and Y tracks generally causes E/p and p_T^2 to be miscalculated, which can lead to loss of $\pi^+\pi^-$ signal events and extra background from $K_L \rightarrow \pi e \nu$.

6.4.7 Pion Identification Cuts

The decays $K_L \rightarrow \pi e \nu$ (“ K_{e3} ”) and $K_L \rightarrow \pi \mu \nu$ (“ $K_{\mu3}$ ”) occur much more frequently than $K_L \rightarrow \pi^+\pi^-$, so they will contribute significant backgrounds unless we make sure that both particles are pions. Electrons deposit essentially all of their energy in the calorimeter, so we calculate the ratio of the energy deposit to the track momentum, E/p , and require it to be less than 0.85. Most muons leave a signal in the muon veto hodoscope, but low-momentum muons may be stopped in the steel. Therefore, we require each particle to have a momentum of at least 8 GeV/ c , so that we can rely on the muon veto hodoscope to be nearly 100% efficient.

6.4.8 Kinematic Cuts

The KTeV beams contain many Λ particles whose dominant decay mode, $\Lambda \rightarrow p\pi^-$, may be mistaken for $K \rightarrow \pi^+\pi^-$. The pion identification cuts listed above do not help in this case because the proton interacts much like a pion. Fortunately, it is simple to reject this background on the basis of kinematics. For each event in our sample, the $p\pi^-$ (or $\bar{p}\pi^+$) invariant mass is calculated under the assumption that the higher-momentum particle is the proton (or antiproton). The event is rejected if this invariant mass lies within the range 1.112 to 1.119 GeV/c^2 . Note that this cut is applied at all energies, even though most of the Λ 's in the beamline actually have rather high energies.

The two remaining kinematic variables are the $\pi^+\pi^-$ invariant mass $m_{\pi\pi}$ and the square of the transverse momentum, p_T^2 (calculated assuming that the kaon scattered at the Z of the downstream edge of the Regenerator, as described in Section 6.3). Figure 6.4 shows the mass distributions for events in the two beams after all other analysis cuts, overlaid to line up the central peaks for the two beams. The shapes of the peaks are nearly the same, with an RMS width of about 1.6 MeV/c^2 . The larger “wings” for the vacuum beam are due to semileptonic K_L decays, as we will show in Section 6.5. The low-mass tail visible in both beams, between about 0.475 and 0.49 GeV/c^2 , is due to the radiative decay $K \rightarrow \pi^+\pi^-\gamma$, where the photon tends to have low center-of-mass energy as expected for this “inner bremsstrahlung” process. We require that $m_{\pi\pi}$ be between 0.488 and 0.508 GeV/c^2 .

Figure 6.5 shows the p_T^2 distributions for events in the two beams after all other analysis cuts, again overlaid to line up the peaks for the two beams. The peaks near $p_T^2 = 0$ are broadened by the spectrometer resolution (including various non-Gaussian processes),⁷ but are very similar for the two beams. The tails extending away from the peaks are from kaons which scatter before decaying to $\pi^+\pi^-$ and from semileptonic K_L decays, as we will see in Section 6.5. We require events to have $p_T^2 < 250 \text{ MeV}^2$ to minimize these backgrounds. This is a fairly tight cut, and so

⁷Radiative $K \rightarrow \pi^+\pi^-\gamma$ decays also broaden the peaks somewhat, but not as much as the spectrometer resolution because the $m_{\pi\pi}$ cut only keeps events with rather small center-of-mass photon energies.

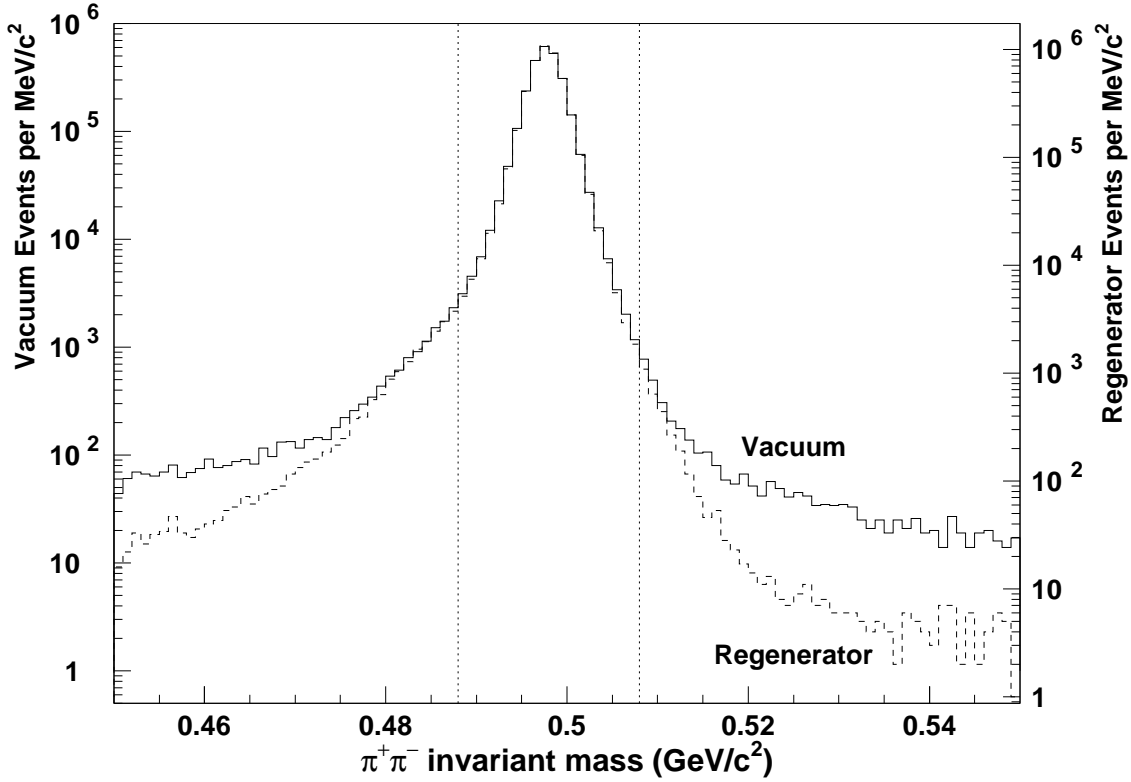


Figure 6.4: $\pi^+\pi^-$ invariant mass distributions after all other analysis cuts have been applied. The dotted lines show the cut made on this quantity.

we will have to assign a significant systematic error on $\text{Re}(\epsilon'/\epsilon)$ due to uncertainty in modeling the p_T^2 tails.

6.4.9 Fiducial Cuts

We observed a small number of events which pass all of the cuts listed above, but which reconstruct well outside of either beam. This is evidence for a slight “halo” of scattered kaons around the beams. The number of events is too small to have a significant effect on $\text{Re}(\epsilon'/\epsilon)$, but we nevertheless make a cut to remove them. Specifically, we project the decay vertex back to the Regenerator along the

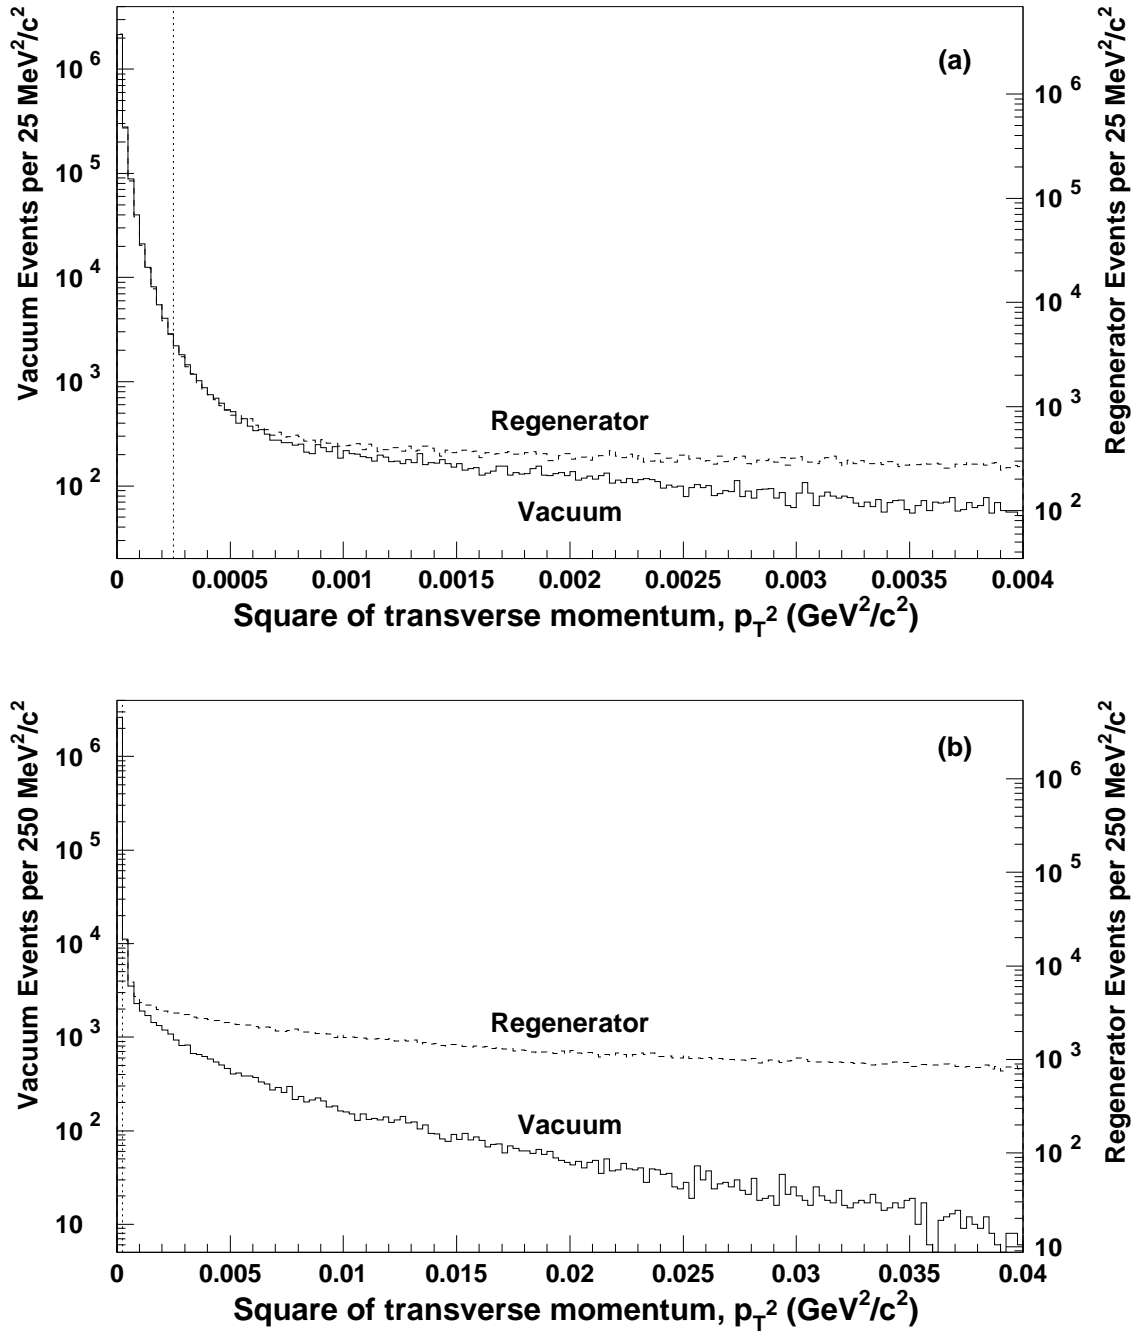


Figure 6.5: p_T^2 distributions for the data after all other analysis cuts have been applied. Panels (a) and (b) show the same data on two different horizontal scales. The dotted lines show the cut made on this quantity.

total momentum vector (or, if the vertex is upstream of the Regenerator, project it forward along a line from the target) and reject the event if the projection falls outside of a square with area 75 cm^2 centered around either beam.

Figure 6.6 shows the vertex Z distributions for the two beams after all other analysis cuts. The geometric cuts described in Section 6.4.6 largely define the accepted range; for example, the vacuum-beam roll-off at the upstream end is due to the cut made on projected track positions at the MA, while the roll-off at the downstream end is due to the track separation cut at drift chamber 1. However, we also explicitly require events to have $110 < Z < 158 \text{ m}$ in the vacuum beam and $122 < Z < 158 \text{ m}$ in the regenerated beam. (There are about a dozen events upstream of $Z = 122 \text{ m}$ in the regenerated beam, from $K_L \rightarrow \pi^+\pi^-$ decays for which the pions manage to pass through the MA beam hole and miss the Regenerator. We remove these events, even though their effect on $\text{Re}(\epsilon'/\epsilon)$ is negligible.)

Figure 6.7 shows the kaon energy distributions for the two beams after all analysis cuts. The distributions are qualitatively similar, though the vacuum beam has more events at both low and high energies (from decays at the downstream and upstream ends of the decay volume, respectively). We will use events in the range $40 < E_K < 160 \text{ GeV}$ for the measurement of $\text{Re}(\epsilon'/\epsilon)$. In this interval, there are a total of 2,609,427 candidate events in the vacuum beam and 4,519,941 candidate events in the regenerated beam. Before we can use these samples to measure $\text{Re}(\epsilon'/\epsilon)$, we need to correct for backgrounds.

6.5 Background Subtraction

The fact that the $K \rightarrow \pi^+\pi^-$ decay has such well-determined kinematics makes it relatively easy to reject backgrounds. The mass and p_T^2 cuts, as well as the pion identification capabilities of the KTeV detector, give us multiple ways to reject other decay modes. Kaons which decay to $\pi^+\pi^-$ after scattering in the detector pass the mass and pion-identification cuts, but the p_T^2 cut alone is sufficient to keep the resulting backgrounds low. Therefore, the total background levels are rather small. Nevertheless, they must be understood before we can determine $\text{Re}(\epsilon'/\epsilon)$ to

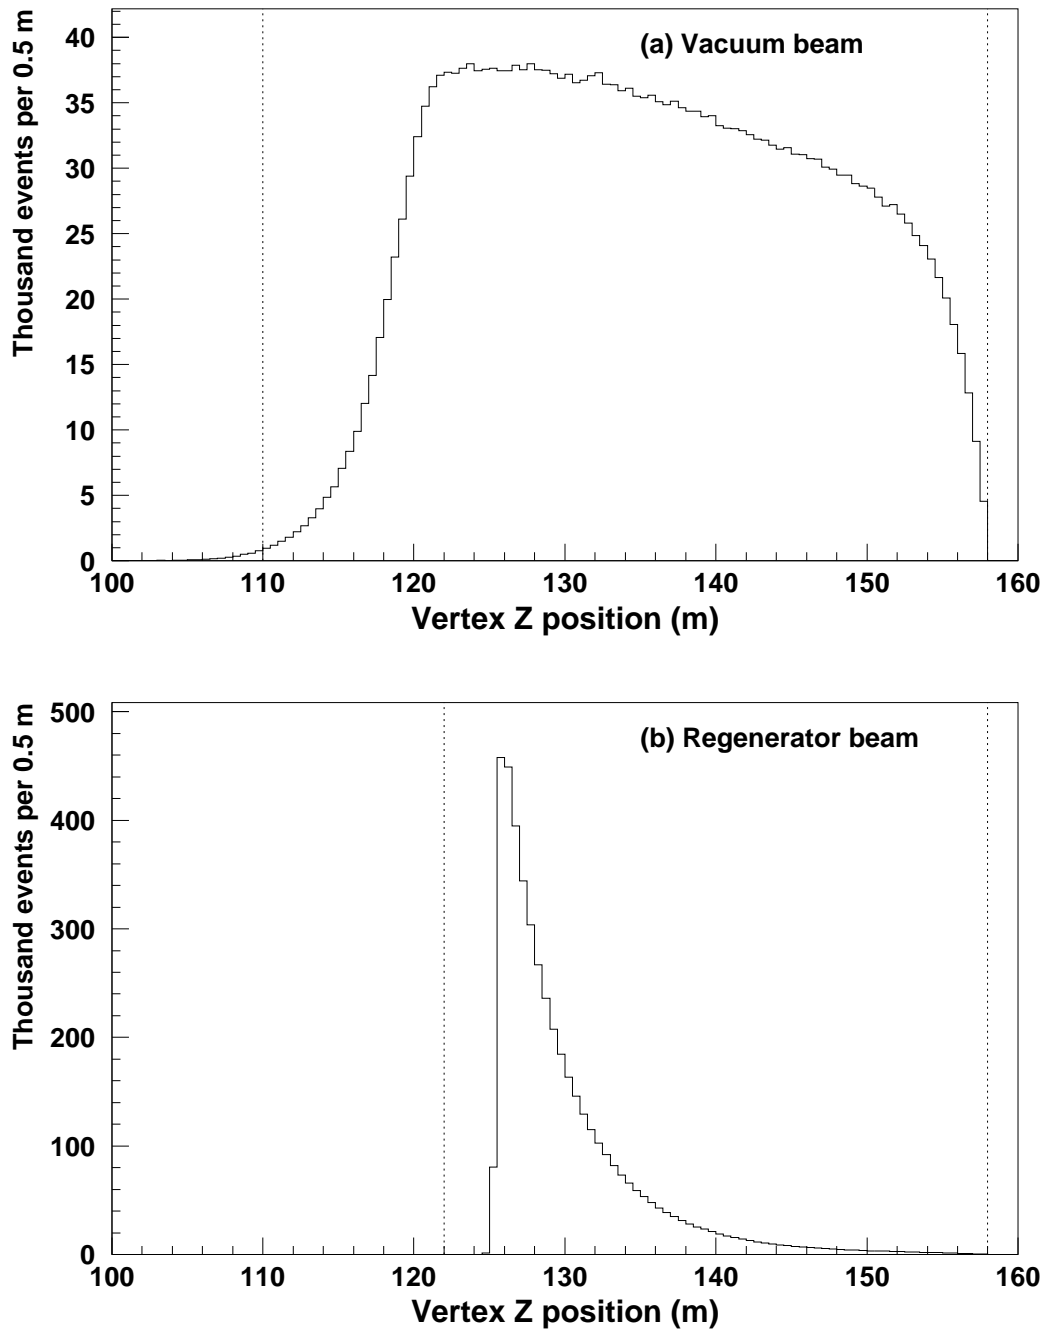


Figure 6.6: Vertex Z distributions for the charged-mode data after all other analysis cuts have been applied. The dotted lines show the cuts made on this quantity.

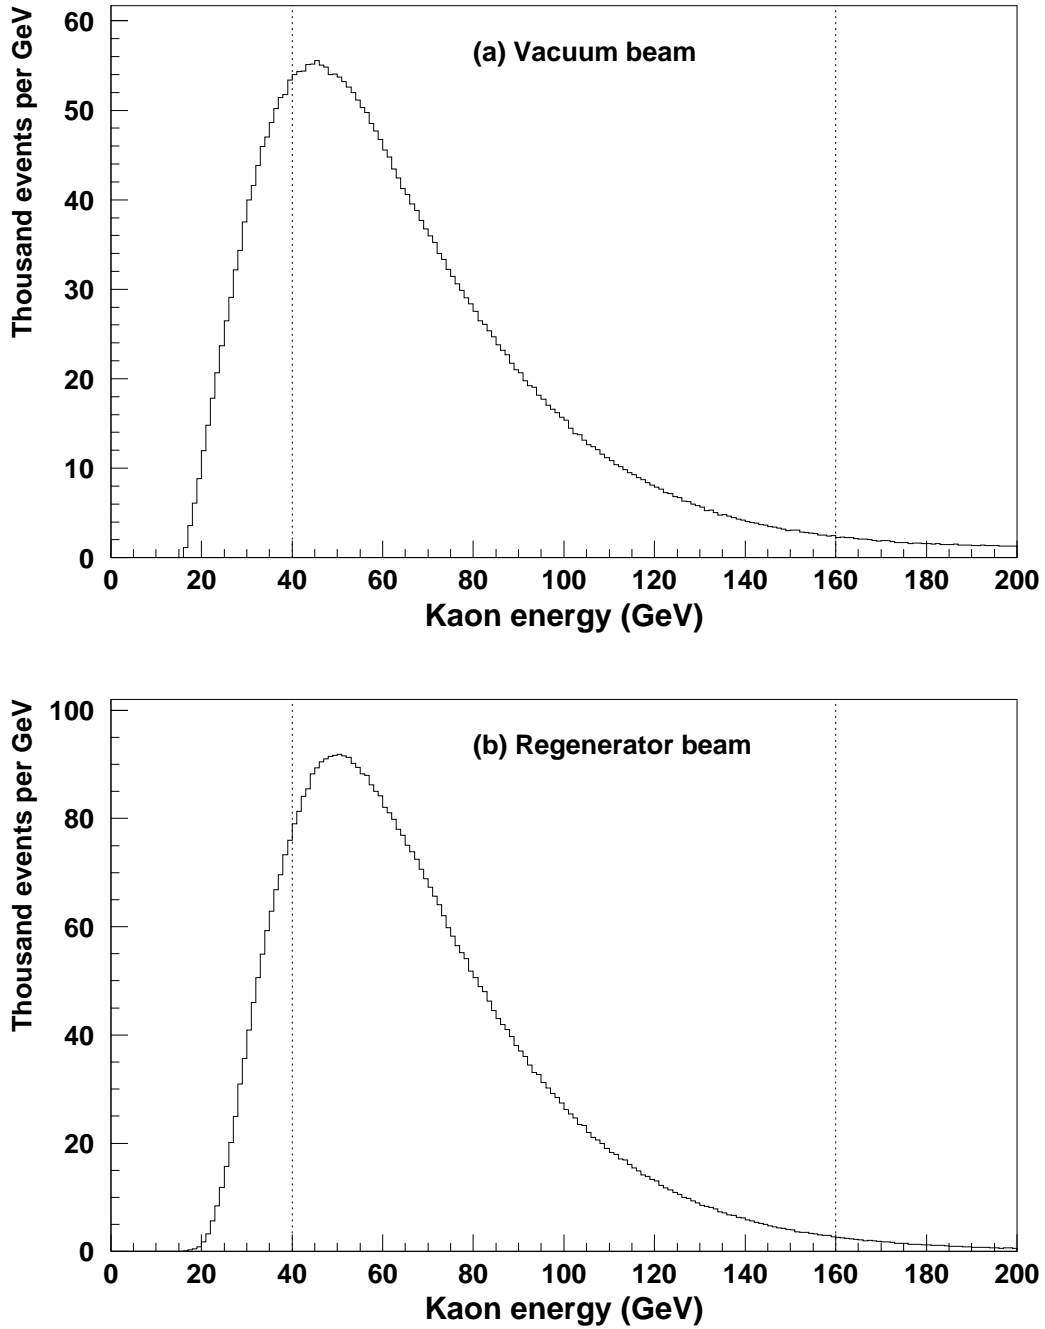


Figure 6.7: Kaon energy distributions for the charged-mode data after all other analysis cuts have been applied. The dotted lines show the interval used to define the final samples for the measurement of $\text{Re}(\epsilon'/\epsilon)$.

the desired precision, and they allow us to characterize the kaon scattering processes which contribute substantial backgrounds to the $\pi^0\pi^0$ samples.

6.5.1 Background Processes

Even after making the pion identification cuts described in Section 6.4.7, the semileptonic decays $K_L \rightarrow \pi e \nu$ and $K_L \rightarrow \pi \mu \nu$ contribute measurable backgrounds to the vacuum-beam $\pi^+\pi^-$ data sample. Both cases require large excursions from the usual interactions of the leptons with matter. An electron normally deposits all of its energy in the CsI, with an RMS of less than a percent; to make a background event, it must deposit less than 85%. In fact, we see a tail which extends down to 50% or perhaps even lower, which is thought to be due to rare electron-nucleus interactions. A muon (with momentum above 8 GeV/c) normally has no difficulty passing through the steel in front of the muon veto hodoscope, but there is some small chance of it stopping in the steel. We cannot expect to make reliable quantitative predictions for the frequency of these rare processes, so we must determine the background level from the data. Semileptonic decays contribute very little background in the regenerated beam because there are far fewer K_L 's (by a factor of ~ 13) and the decay region is somewhat shorter.

The dominant background in the regenerated beam is from kaons which scatter in the Regenerator (without depositing enough energy to be vetoed in the analysis) and then decay to $\pi^+\pi^-$. A number of processes are expected to contribute, such as diffractive scattering off of carbon and lead nuclei, multiple diffractive scattering, and some amount of residual inelastic scattering. Each of these has a different characteristic momentum transfer and effect on the kaon quantum state (which can be thought of as a complex-valued effective regeneration amplitude, $\tilde{\rho}$), resulting in a complicated distribution of kaon decays as a function of kaon energy, p_T^2 , and distance downstream of the Regenerator. We cannot make quantitative predictions for these processes, so we fit the data (in bins of E_K , p_T^2 , and the proper time τ) to a physics-motivated parametrization of the form

$$\frac{d^3N}{dp_T^2 d\tau dE_K} \sim \sum_{j=1}^5 A_j e^{-B_j p_T^2} \left(\frac{E_K}{70 \text{ GeV}}\right)^{\alpha_j} \frac{dN}{d\tau}(\tilde{\rho}_j, \tau) \quad (6.1)$$

where

$$\frac{dN}{d\tau}(\tilde{\rho}_j, \tau) = \left| \eta_{+-} e^{i\Lambda_L \tau} + \tilde{\rho}_j e^{i\Lambda_S \tau} \right|^2 \quad (6.2)$$

is the $K \rightarrow \pi^+ \pi^-$ decay rate at proper time τ from this scattering component. The parameters A_j , B_j , α_j , and $\tilde{\rho}_j$ are floated in the fit. The parametrization was chosen to have five terms to roughly match the number of processes expected to contribute; for one of the terms, $dN/d\tau$ is given a pure K_S time dependence, since inelastic K_L scattering is empirically observed to produce an essentially pure K_S state. This parametrization matches the data rather well.

$K \rightarrow \pi^+ \pi^-$ decays from scattered kaons are also observed in the vacuum beam. The trajectories of the parent kaons, determined for each event from the total momentum vector of the $\pi^+ \pi^-$ system, do not point back to the Regenerator, but to the edges of the final beam-defining collimator. This is clearly seen in Figure 6.8, which shows the projected kaon position for events with $p_T^2 > 5000 \text{ MeV}^2/c^2$ after applying very tight pion identification cuts to eliminate the bulk of the semileptonic kaon decays.

In hindsight, we should have expected this background because there is a significant beam flux incident on the edges of the defining collimator, as is visible in Figure 4.1. Although a kaon entering the front face of the collimator has no chance of passing through the whole 3 meters of steel, it may scatter into the beam hole if it starts out close to the edge, as shown by the cartoon in Figure 6.9.

Many of these collimator-scattered events are clustered at high kaon energy and at the far upstream end of the decay region, indicating that the scattered K_L 's became K_S 's; these events tend to have large values of p_T^2 . There is also a K_L component emerging from the collimator with a much lower characteristic momentum transfer. Note that the collimator-scattered K_L component must also be present in the regenerated beam, but is not easily distinguishable from the Regenerator-scattered events. Therefore, we must properly account for collimator scattering when using the data to characterize Regenerator scattering.

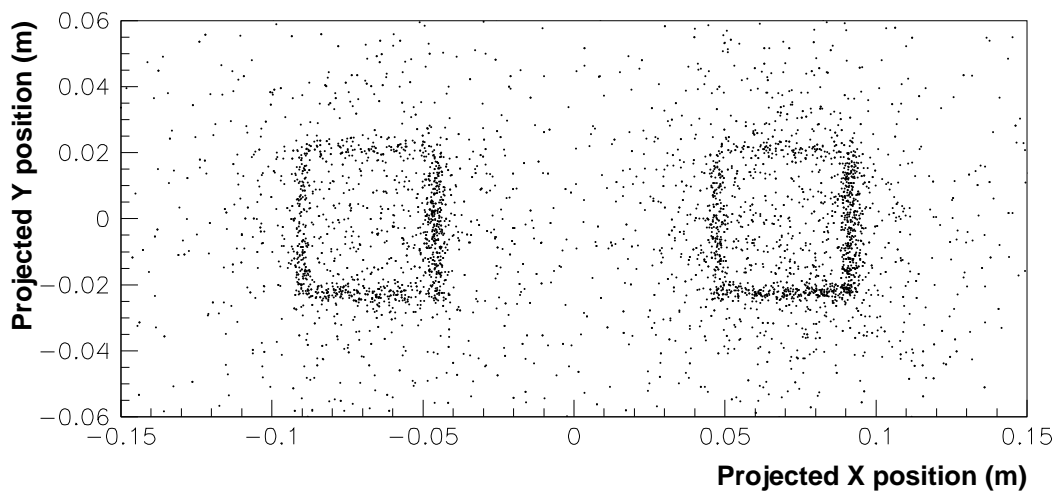


Figure 6.8: Kaon scattering from the edges of the final collimator, demonstrated by selecting vacuum-beam $\pi^+\pi^-$ candidates with $p_T^2 > 5000 \text{ MeV}^2/c^2$, making extra-tight pion identification cuts, and projecting the total momentum vector back to the Z position of the collimator.

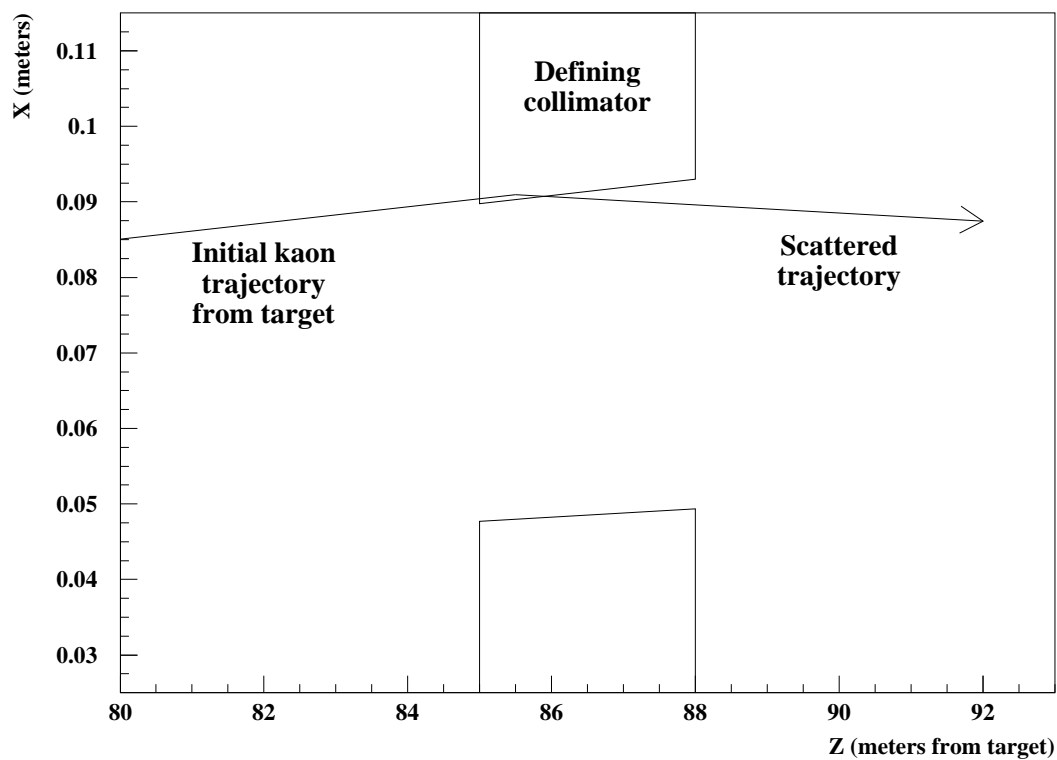


Figure 6.9: Cartoon of kaon scattering in the final collimator.

A fairly simple model, based on the cartoon in Figure 6.9, has been implemented in the Monte Carlo and does a rather good job of reproducing the data distributions. The incident K_L may scatter to a K_S or K_L final state. The p_T^2 imparted to the kaon is chosen from an exponential distribution, with different characteristic scales for the two cases. The probabilities of the two scattering cases are given an energy dependence which is tuned to match the data. Note that smaller-angle scatters are less likely to escape the collimator steel.

There is a small additional background from “Regenerator junk”, the production of multiple particles by inelastic interactions in the Regenerator, particularly in the lead pieces. This shows up as events reconstructing near the edge of the Regenerator with a broad invariant-mass distribution and a p_T^2 distribution somewhat peaked at zero. The tight cut made on the energy in the last Regenerator module suppresses this background to a very low level, and we have chosen to ignore it in the background-subtraction procedure.

6.5.2 Normalization of Background Contributions

Having identified the sources of backgrounds, we need a procedure to determine how much each one contributes. We must do the vacuum beam first since that will tell us how much collimator scattering is present in regenerated beam, where it cannot be distinguished from Regenerator scattering. The general approach is to use “sideband” regions in the p_T^2 -vs.- $m_{\pi\pi}$ distribution to normalize the different background processes. However, it is tricky to disentangle the different contributions from the basic p_T^2 -vs.- $m_{\pi\pi}$ plot, shown in Figure 6.10. In particular, it is hard to distinguish $K_L \rightarrow \pi\mu\nu$ from $K_L \rightarrow \pi e\nu$ since they populate overlapping regions.

We have developed a more robust procedure to normalize the background contributions. We apply additional cuts to the data to *enhance* each of the backgrounds relative to the others. To enhance the $K_L \rightarrow \pi e\nu$ background, we require that one of the particles has $E/p > 0.75$ and that the other particle deposits at least 1 GeV in the CsI. To enhance the $K_L \rightarrow \pi\mu\nu$ background, we require that one particle be minimum-ionizing while the other has $E/p < 0.5$. To enhance the collimator-scattered $\pi^+\pi^-$ background, we require that both particles have $E/p < 0.5$ but




Figure 6.10: Distribution of p_T^2 vs. $m_{\pi\pi}$ for vacuum-beam data events after all other analysis cuts. Each dot represents about 20 events. The signal region is marked with a box (along the bottom edge, near the center).

neither is minimum-ionizing, and that the incident kaon trajectory projects back near an edge of the defining collimator. Figure 6.11 shows the effects of these cuts on the data (left column), with comparisons to the simulated distributions from the background processes (right column). In each case, an appropriate region (outlined in the plots) is used to normalize the background to the data. We then correct for the efficiencies of the extra particle-identification cuts (measured from the data in a separate study) to deduce the number of background events of each type within the signal region. We also scale the Z distributions for the background samples and subtract them from the data distribution.

The entire process described above is done independently in 10-GeV bins of kaon

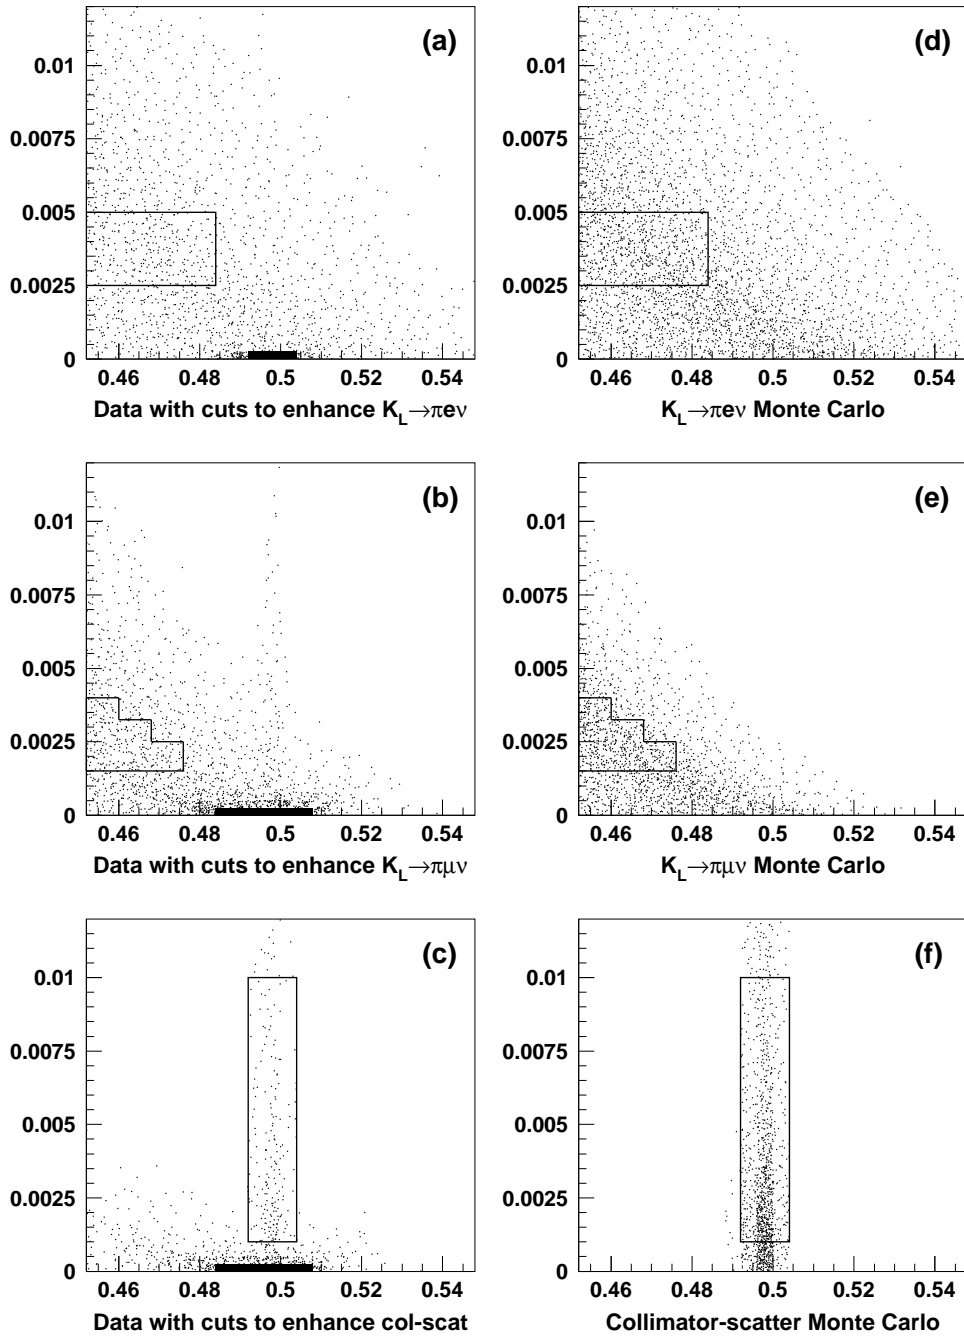


Figure 6.11: Distributions used to normalize background contributions to vacuum-beam $\pi^+\pi^-$ sample. Panels (a), (b), and (c) show the distributions of p_T^2 vs. $m_{\pi\pi}$ for the data, with various extra cuts to enhance each of the expected backgrounds. Panels (d), (e), and (f) show the simulated distributions for the corresponding background processes. The outlines show the regions used for normalization in each case. (The dot-density scales are arbitrary for these plots.)

energy. The $K_L \rightarrow \pi\mu\nu$ normalization scale varies by a factor of three over the range 40–160 GeV, indicating that low-momentum muons are much more likely than high-momentum muons to be stopped in the steel; the normalization scales for the other processes are fairly constant. The normalization determined for collimator scattering in the vacuum beam is used to predict the amount of collimator scattering in the regenerated beam. After subtracting this, the remaining scattered-kaon events in the regenerated beam are used to determine the parameters in the parametrization described in Section 6.5.1, which is incorporated into the Monte Carlo simulation.

The normalized background components, when combined, do a good job of reproducing the p_T^2 distributions observed in the data, as shown in Figure 6.12. The agreement in the vacuum beam is especially encouraging since the semileptonic backgrounds were *not* determined from the events in this distribution, but from events with low invariant mass *and* nonzero p_T^2 . The semileptonic backgrounds also reproduce the tails on the invariant-mass peak in the vacuum beam (after applying the p_T^2 cut), as shown in Figure 6.13.

The fractional backgrounds contributed by each background process are summarized in Table 6.1. The statistical uncertainty on the background normalization is small, but a relative systematic uncertainty of 10% is assigned to the total background level for each beam, based on a judgment about how well the p_T^2 distributions are reproduced.

Table 6.1: Backgrounds in the $\pi^+\pi^-$ samples.

Background process	Vac beam	Reg beam
$K_L \rightarrow \pi e \nu$, e misidentified as π	0.039%	0.002%
$K_L \rightarrow \pi \mu \nu$, μ misidentified as π	0.030%	0.001%
Kaon scattered in collimator	0.014%	0.014%
Kaon scattered in Regenerator	—	0.072%
TOTAL	0.083%	0.089%

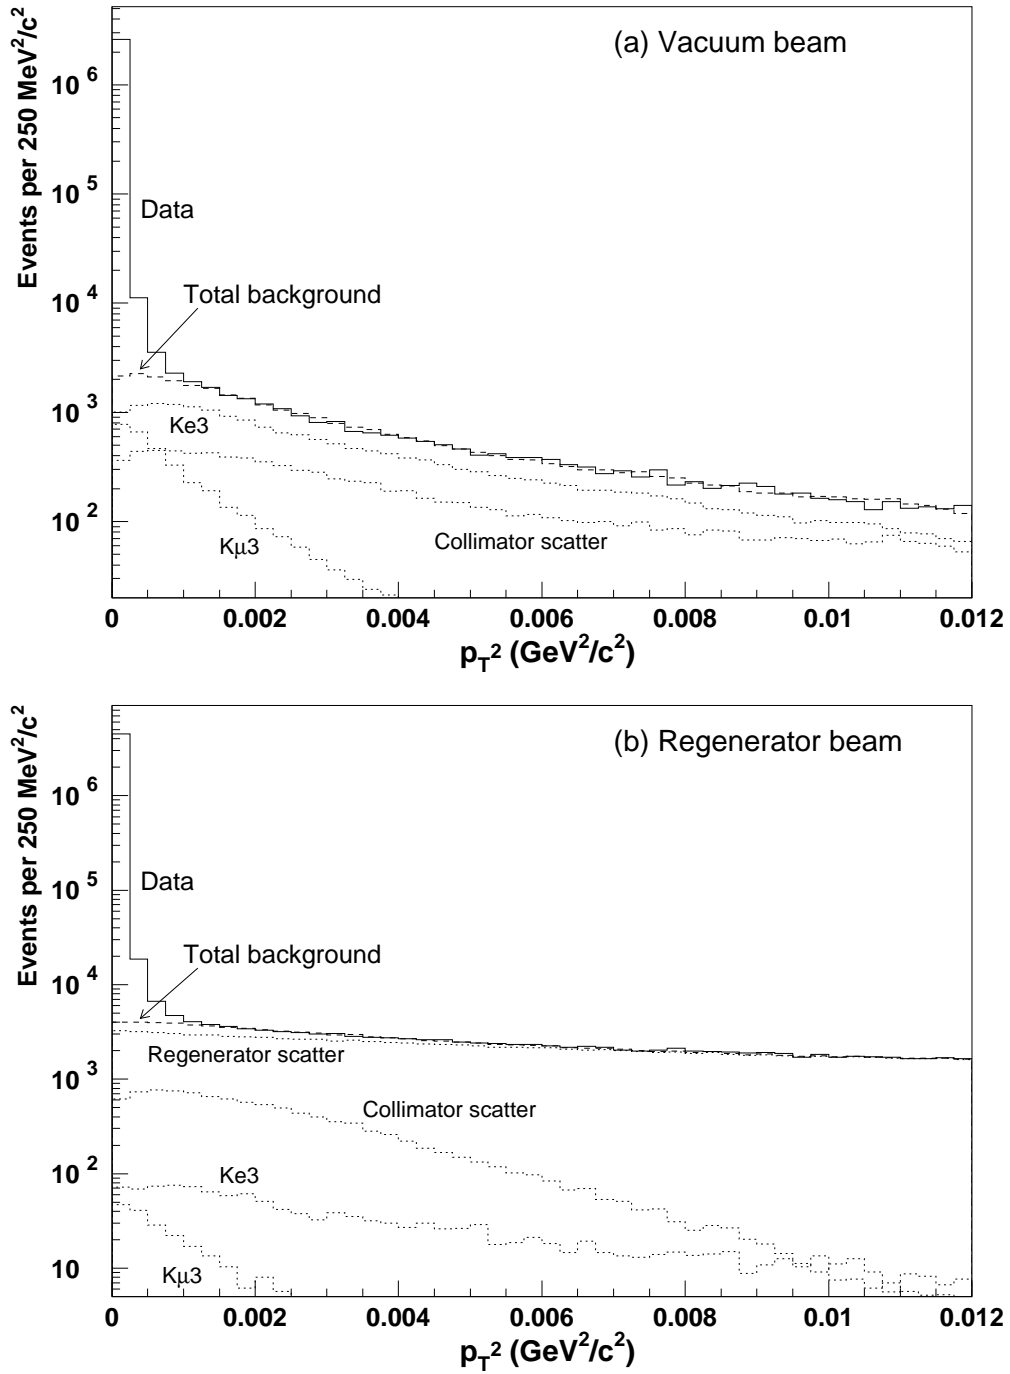


Figure 6.12: p_T^2 distributions for data and background components after all other analysis cuts have been applied, including the invariant mass cut. The p_T^2 cut used to define the final sample keeps just the first bin of each of these plots.

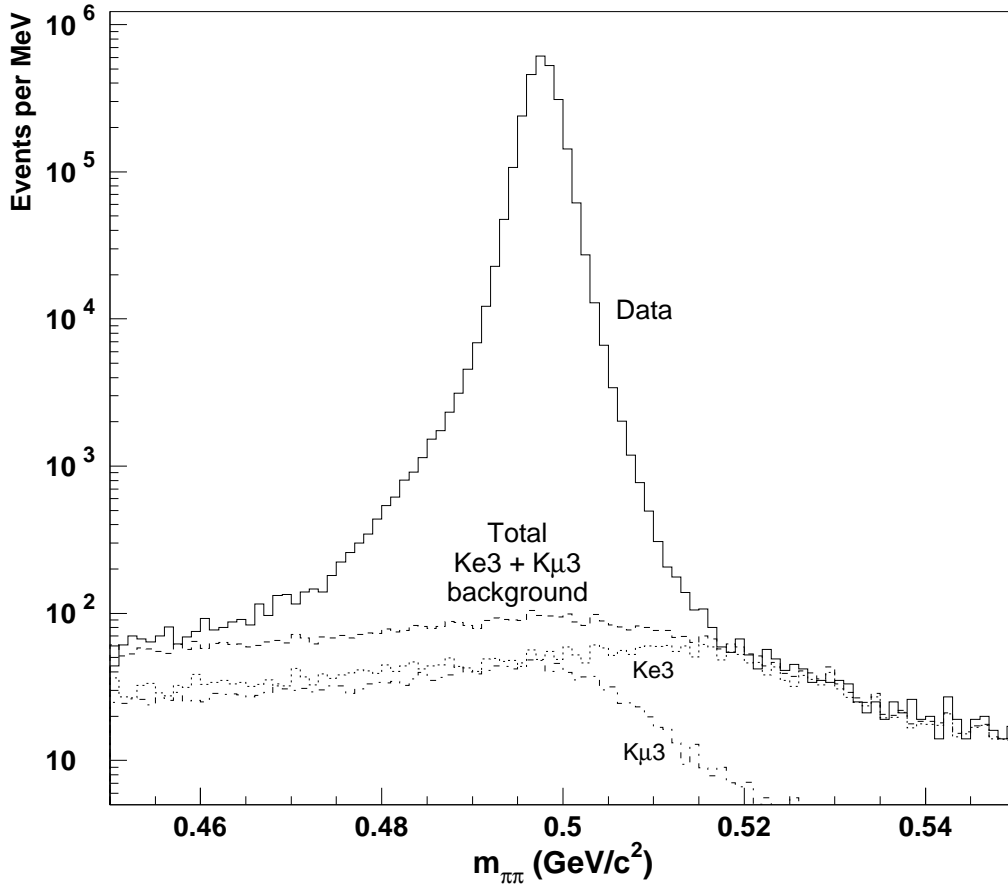


Figure 6.13: Vacuum-beam $m_{\pi\pi}$ distribution for data and semileptonic background components after all other analysis cuts have been applied, including the p_T^2 cut.

6.6 Summary

After background subtraction, the net yields (for the range $40 < E_K < 160$ GeV) are 2,607,274 $\pi^+\pi^-$ in the vacuum beam and 4,515,928 $\pi^+\pi^-$ in the regenerated beam. The background-subtracted (Z, E_K) distributions are ready for input to the fitting program to determine $\text{Re}(\epsilon'/\epsilon)$.

CHAPTER 7

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

This chapter will describe how $K \rightarrow \pi^0\pi^0$ candidates are collected and analyzed. The basic event topology, four photons in the calorimeter (two from each π^0), is initially required by the trigger. Offline analysis consists of refining the measured photon energies and positions and reconstructing the decay by choosing the best “pairing” of the photons. After describing all of the analysis cuts, we will explain how the Z distribution at the downstream edge of the Regenerator is used to make a final adjustment to the calorimeter energy scale. The backgrounds are rather large in this mode, primarily due to kaon scattering, but the scattering distributions are accurately measured from *charged* kaon decays and so the background subtraction procedure is fairly simple. At the end of this chapter, we will have completed all of the $K \rightarrow \pi\pi$ data analysis needed to measure $\text{Re}(\epsilon'/\epsilon)$.

7.1 Trigger Requirements

The goal of the neutral-mode trigger is to select events with four photons in the calorimeter. At Level 1, the trigger uses the fast analog energy sum produced by the E-Total system. The trigger threshold is around 26 GeV. Like the charged-mode trigger, the neutral-mode trigger was vetoed if there was significant energy in the Regenerator, the Spectrometer Antis (except the one at drift chamber 3), or the muon hodoscope. The Mask Anti was also in veto for about the first 30% of the data collected during the 1996 running period.

The common decay modes $K_L \rightarrow 3\pi^0$, $K_L \rightarrow \pi e\nu$, and $K_L \rightarrow \pi^+\pi^-\pi^0$ also generally deposit a large amount of energy in the calorimeter, so the neutral-mode Level 1 trigger rate is quite high. At Level 2, the rate is greatly reduced by checking

the number of clusters reported by the HCC. Because the E-Total “hot bit” problem would occasionally cause a phantom extra cluster, the trigger required either four or five HCC clusters (except for a brief period at the beginning of physics data collection, when it required exactly four).

The HCC relied on the E-Total system to determine which calorimeter channels had energies above a threshold, nominally 1 GeV, based on analog signals from the last dynode of each PMT. The value of the threshold, in millivolts, was the same throughout the E-Total system, so the effective threshold in GeV depended on the product of the light output for a CsI crystal and the gain of its PMT. For about the first quarter of the 1996 running period, the effective thresholds were not very well balanced from channel to channel, and there was a systematic difference between the small- and large-crystal channels. This was not a serious problem, but required more offline calibration work to determine the actual threshold curves for each channel and use them in the Monte Carlo simulation. Partway through the 1996 running period, the PMT high voltages were adjusted to make the thresholds about the same throughout the calorimeter.

As described in Section 4.7.4, the Level 3 filter process calculated cluster energies and positions and evaluated the best and second-best photon pairings to make two π^0 's. At the beginning of the 1996 running period, Level 3 kept events for which the best pairing gave an invariant mass of at least $450 \text{ MeV}/c^2$ and had a pairing χ^2 of less than 50, or for which the second-best pairing gave an invariant mass of at least $450 \text{ MeV}/c^2$ and had a pairing χ^2 of less than 10. However, it was soon decided that these pairing χ^2 cuts were uncomfortably tight, and they were loosened to 500 and 50, respectively. Only about 4% of the data was collected with the tighter cuts.

7.2 Cluster Finding

The neutral-mode analysis begins by identifying clusters in the calorimeter and calculating their energies and positions. This involves many corrections to the raw digitized values, as will be briefly described below. The corrections were developed from many *in situ* studies of the calorimeter, several of them using large samples

of momentum-analyzed electrons from $K_L \rightarrow \pi e \nu$ decays. In the end, the excellent performance of the calorimeter is a result both of its intrinsic properties and of the effort spent on understanding its behavior.

The first step is to unpack the calorimeter data and calculate the energy deposited in each block. For each DPMT time slice, the flash ADC value is converted to an energy with a linear conversion which depends on which QIE range was used. A study of electrons in the data showed that this nominal energy conversion has a residual nonlinearity, probably due to the fact that the laser pulses used to calibrate the DPMTs do not have quite the same shape as CsI scintillation pulses, and so a correction (generally less than 1%) is applied to remove this nonlinearity. The energies for four slices are summed together to give the total energy deposited in the crystal. At this point a correction is made based on the current response of this channel to flashes from the laser calibration system, which tracks short-time-scale drifts in the PMT gains.

The next step is to find “seed” blocks from among those which have their HCC bits set; a block is considered to be a seed if none of its neighbors has an HCC bit set, or if its energy is greater than any neighbor which does have its HCC bit set.¹ Requiring seed blocks to have HCC bits set helps to verify the trigger and to use only the particles which arrived within the in-time bucket, since the reconstructed energy in each channel is a sum over 12 buckets—4 DPMT time slices, each one spanning three buckets. (HCC bits are not required when finding clusters for the charged-mode analysis, since charged pions often deposit less than 1 GeV in the calorimeter.) A seed block is also required to have at least 0.1 GeV of energy, to be insensitive to the occasional “hot” HCC bits.

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 defined the sum of the individual block energies. The cluster X position is determined by summing the block energies in the central column (the one containing the seed block) and in each adjacent column, and comparing the ratios of column energies against look-up

¹Blocks touching diagonally at a corner are considered to be neighbors for purposes of this test.

tables derived from a sample of real photons. Similarly, the Y position is determined from the ratios of row energies.

A number of corrections are applied to the cluster energy based on its position and raw energy. The energy is adjusted to correct for the intrinsic nonlinearity of the CsI crystals due to the nonuniformity of light collection along their lengths, individually measured for each crystal using cosmic-ray muons. If a seed block is located near a beam hole or near the outer edge of the calorimeter, then the square cluster region may include one or more nonexistent blocks; in this case, the energy expected to have been in those blocks (according to the average transverse energy distribution of a photon) is added to the cluster energy. A correction is also made for blocks which are physically present but which were not read out due to the channel readout threshold (nominally 4 MeV).

If two (or more) clusters overlap, then the energy in each shared block is divided between them according to the expectation from the average transverse energy distribution of a photon, and the cluster energies and positions are recalculated; this process is iterated until it converges to a self-consistent division of the energy in the overlap region. A small correction is also made for blocks which are close to another cluster but not close enough to overlap, to correct for “leakage” of energy outside the square cluster region. The leakage problem is more complicated for clusters near a beam hole, for which shower particles can cross the beam hole and deposit energy on the far side. A special correction, derived from the data, is applied to handle such cases.

The final set of corrections are based on studying the response of the calorimeter to electrons as a function of various factors; in each case, the correction to be applied was determined from the electron data. The largest such effect, and the most intriguing, is the observation that the scintillation response of the CsI crystals varies significantly as a function of transverse position. This was not seen during the initial preparation and testing of the crystals since we only measured the response as a function of position along the length of a crystal. In fact, the response is generally lower, by as much as a few percent, near the edges of a given crystal than in the

center. Making this correction dramatically improves the offline energy resolution of the calorimeter.

The calorimeter response was found to have a systematic variation among the three classes of events with different timing of the DPMT slices relative to the in-time bucket. It was also found to vary over time between calorimeter calibrations (typically two days apart), due in part to temperature variations in the CsI block-house and to a slow overall degradation of the CsI crystals due to radiation damage. After applying these and all other corrections, there is a small residual nonlinearity in the calorimeter response. The resulting *ad hoc* correction is a few tenths of a percent in the worst case (very low and very high energies), and generally under one tenth of a percent.

7.3 Event Reconstruction

Once we have established that an event has exactly four in-time clusters, we must determine whether they came from the four photons from a $K \rightarrow \pi^0 \pi^0$ decay. There are three possible ways for the photons to be paired to make two π^0 's, and we must determine which is the correct pairing. The hypothesis that photons a and b come from a π^0 decay allows us to infer, from the π^0 invariant mass constraint, the opening angle θ_{ab} between the photon trajectories:

$$\begin{aligned} m_{\pi^0}^2 &= (E_1 + E_2)^2 - |\vec{p}_a + \vec{p}_b|^2 \\ &= 2E_1 E_2 - \vec{p}_a \cdot \vec{p}_b \\ &= 2E_1 E_2 (1 - \cos \theta_{ab}) \end{aligned} \tag{7.1}$$

Since θ_{ab} is a small angle, we may make the approximations $(1 - \cos \theta_{ab}) \approx \theta_{ab}^2/2$ and $\theta_{ab} \approx r_{ab}/z_{ab}$, where r_{ab} is the distance between the photons at the calorimeter and z_{ab} is the distance from the π^0 decay position to the calorimeter. It follows that

$$z_{ab} \approx \frac{\sqrt{E_1 E_2}}{m_{\pi^0}} r_{ab} \tag{7.2}$$

In general, one of the three pairing combinations will give consistent Z positions for the two π^0 decays (*i.e.* $z_{ab} \approx z_{cd}$), and the other two pairing combinations will not. This is illustrated in Figure 7.1. The actual variable used to choose the best pairing is the “pairing χ^2 ”,

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

where $\sigma_{z_{ab}}$ is the error on z_{ab} due to the finite energy and position resolution of the calorimeter. Because the KTeV calorimeter has excellent resolution, the probability of incorrectly pairing the photons from a $K \rightarrow \pi^0\pi^0$ decay is only about 0.03%.

The Z position of the kaon decay vertex is taken to be the weighted average of z_{ab} and z_{cd} . The (X, Y) position of the decay vertex is not directly measurable from the available information; however, the energy centroid of the four photons at the calorimeter indicates where the kaon was heading when it decayed. Therefore, 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 vertex position, we calculate the momentum vectors for each of the photons and the four-photon invariant mass.

Figure 7.2 shows a reconstructed $\pi^0\pi^0$ event. The energy centroid lies within the vacuum-beam beam-hole in the CsI, so this is most likely a vacuum-beam $K_L \rightarrow \pi^0\pi^0$ decay.

Because the initial kaon trajectory is unknown, the energy centroid is the only available indicator of kaon scattering. If the energy centroid lies outside both of the beams, then we can conclude that the kaon scattered before decaying. To facilitate this check, we calculate a quantity called the “ring number”, defined as

$$4 \times \max(\Delta x^2, \Delta y^2) \quad (7.4)$$

where Δx and Δy are the horizontal and vertical distances between the energy centroid and the center of the closer beam, measured in centimeters. The ring number may also be thought of as follows: if one draws a square which is centered on the closer beam and for which the energy centroid lies on its edge, then the ring number is the area of that square in cm^2 . The beams are nominally 9.3 cm square at the CsI, so all events with ring numbers below 86.49 are within one of the




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

KTEV Event Display

/usr/kpasa/data05/shawhan/kz
n011_some.dat

Run Number: 7095
Spill Number: 221
Event Number: 23751371
Trigger Mask: 8
All Slices

Track and Cluster Info

HCC cluster count: 4


ID	Xcsi	Ycsi	P or E
C 1:	0.2855	-0.4879	8.32
C 2:	-0.1729	0.1387	66.97
C 3:	-0.0323	-0.0253	17.37
C 4:	-0.8546	-0.2040	4.55

Vertex: 4 clusters

X	Y	Z
-0.0997	0.0281	132.091

Mass=0.4965

Pairing chisq=3.62



- - Cluster
- - Track
- - 10.00 GeV
- (dotted) - 1.00 GeV
- (cross-hatched) - 0.10 GeV
- (fine dotted) - 0.01 GeV




Figure 7.2: Event display for a $K_L \rightarrow \pi^0 \pi^0$ event. The thick-lined boxes in the centers of the clusters indicate blocks for which the HCC bit was 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.

beams, neglecting smearing of the beam edges due to the finite size of the target and resolution smearing in the reconstruction.

Note that the ring number only tells us where the kaon ends up, not where it came from. It is possible for a kaon in the regenerated beam to scatter in the Regenerator and head toward the vacuum-beam hole in the CsI, and in this case we would assign the event to the vacuum beam and give it a low ring number. This is a background process that we will have to consider.

7.4 Analysis Cuts

We now describe the cuts applied to the data to select the final samples.

7.4.1 Data Quality

Calorimeter readout problems were significantly more frequent during 1996 than during 1997 (after the QIE chips were replaced). There were also some “special runs” with modifications to the detector configuration. In all, over ten percent of the data collected during the 1996 running period was considered unsuitable for the $\text{Re}(\epsilon'/\epsilon)$ measurement and was discarded in the analysis.

7.4.2 Trigger Verification

The Level 1 trigger hardware required that the energy deposit in the calorimeter exceed a 28 GeV threshold. A TDC was used to record the time at which the threshold was exceeded. Offline, we required this time to lie within an appropriate interval to help verify that the energy really did arrive within the in-time bucket.

As mentioned in Section 7.2, the seed blocks for each of the four clusters are required to have their HCC bits set. This implicitly verifies that the cluster-count trigger requirement should have been satisfied by the four reconstructed photons.

7.4.3 Veto Counters

The primary purpose of most of the veto counters is to suppress the $K_L \rightarrow 3\pi^0$ background by catching one or both of the extra photons. During the neutral-mode analysis, cuts are made on energy deposit in all of the veto counters except the muon veto hodoscope. The cuts made on the Regenerator, Ring Counters, and Spectrometer Antis are essentially the same as those made in the charged-mode analysis, while a looser cut (300 MeV) is made on the Mask Anti. An event is also discarded if any channel in the front section of the Back Anti has more than 5 GeV of energy or if any channel in the Collar Anti has more than 1 GeV.

7.4.4 Extra-Particle Cuts

The veto counters cannot reduce the $K_L \rightarrow 3\pi^0$ background to zero, because it is possible for all six photons to hit the calorimeter but still look like a four-cluster event. The extra photons can be hidden in one of two ways:

- Two photons land at nearly the same place in the calorimeter and are reconstructed as a single cluster. We call this “photon fusion”.
- A photon has energy below about 1 GeV, so it does not cause any HCC bit to be set. Therefore it does not make an HCC-seeded cluster and is not used in reconstructing the event.

We check for photon fusion by calculating a “shape χ^2 ” variable for each cluster, comparing the cluster’s transverse energy distribution with that expected from an electromagnetic shower of that energy. The distribution of shape χ^2 values is not modeled very well by the Monte Carlo simulation, so we make a only a very loose cut, at a value of 48.

We check for low-energy photons in the calorimeter by repeating the cluster-finding algorithm with the “software clustering” option, in which a seed block is simply required to be a local maximum with energy of at least 0.25 GeV, regardless of whether its HCC bit is set. We discard the event if there is any “software” cluster

with energy of at least 0.6 GeV which satisfies the following additional requirements to ensure that it is a real photon:

- There is no track pointing to the software cluster.
- It is not too close to another cluster.
- It is at least 5 cm from the beam holes.
- It has a shape χ^2 of 10 or less.
- The division of energy among the time slices read out by the DPMT is consistent with a photon arriving within the in-time bucket.

Finally, an event is rejected if any track can be reconstructed from the drift chamber information.

7.4.5 Event Quality Cuts

The pairing χ^2 (described in Section 7.3) provides another way of checking whether the four photons came from a real $K \rightarrow \pi^0\pi^0$ decay. As shown in Figure 7.3, the vacuum and regenerated beams have very similar distributions at low values, but at higher values there is an excess in the vacuum beam from $K \rightarrow 3\pi^0$ decays.

The response of the calorimeter is least well understood for energies below a few GeV, because the KTeV analysis magnet sweeps all low-energy electrons away from the calorimeter, removing our primary means of studying the calorimeter response. In addition, the trigger efficiency for low-energy photons depends sensitively on the HCC bit threshold, which varies from channel to channel.² We originally planned to use photons with energies down to 2 GeV, but because of some small differences between data and Monte Carlo in the vicinity of 2 GeV, we decided to require all photons to have an energy of at least 3 GeV.

²These things can, in principle, be studied using low-energy photons from $K_L \rightarrow \pi^+\pi^-\pi^0$ decays, but there are some complications, for example, from disentangling the resolutions for the other particles.

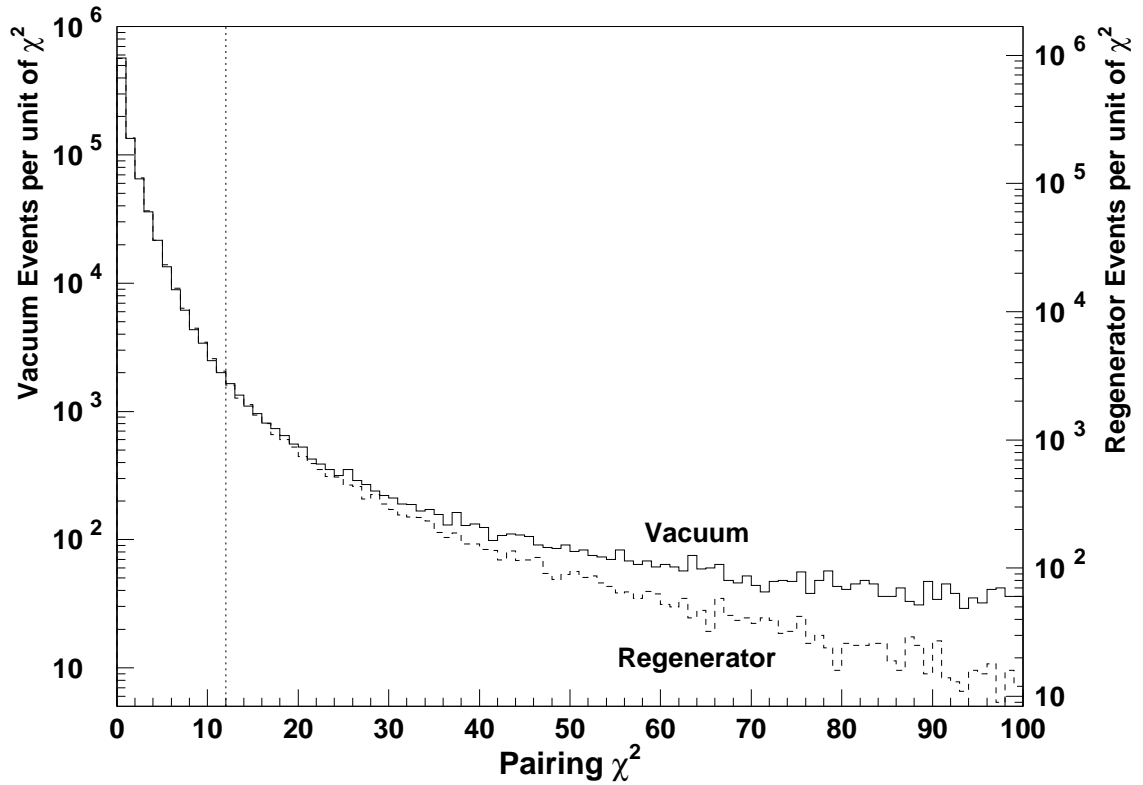


Figure 7.3: Pairing χ^2 distributions after all other analysis cuts have been applied. The dotted line shows the position of the analysis cut on this variable.

7.4.6 Aperture Cuts

As in the charged-mode case, we make some geometrical cuts on particle positions in order to more cleanly define the acceptance. If a photon lands too close to an edge of the CsI array—either a beam hole or the outer edge—then its energy and position are likely to be poorly reconstructed. Therefore, we need to require photons to be some distance from the edge. As has been mentioned before, we use a physical object called the Collar Anti to veto photons closer than 1.5 cm to either of the beam holes. Photons landing just outside the CA (but still within the inner ring

of CsI blocks) are reasonably well constructed, and no cut is made based on the reconstructed position.

The “Cesium Iodide Anti” (the last of the Spectrometer Antis) could, in principle, be used similarly to physically define the usable region at the outer edge of the calorimeter. However, it is located farther in front of the CsI, and many photons there have significant angles, so it would not define a very sharp edge. Therefore, we prefer to use the calorimeter information to define the usable region. Specifically, we require the clusters to *not* have seed blocks in the outermost layer of CsI blocks, *i.e.* the clusters must be at least 5 cm from the outer edge of the array.

Although the cluster-finding software has an algorithm to separate overlapping clusters, we do not want to have to rely on it too heavily. Therefore, we require the clusters to be separated from each other by at least 7.5 cm.

7.4.7 Kinematic Cuts

The four-photon invariant mass distributions, after all other cuts, are shown in Figure 7.4. The mean mass resolution is about $1.5 \text{ MeV}/c^2$ in both beams, slightly *better* than for charged mode. The fairly large number of events at high and low mass in the vacuum beam are from $K_L \rightarrow 3\pi^0$ decays. We require the invariant mass to be between 0.490 and $0.505 \text{ GeV}/c^2$.

The ring number variable does not tell us directly about the kinematics of the event, but it does allow us to identify some fraction of the scattered kaons. The data distributions are shown in Figure 7.5. We require the ring number to be less than 110. This is a rather tight cut, and will lead to a systematic error on $\text{Re}(\epsilon'/\epsilon)$.

7.4.8 Fiducial Cuts

Figure 7.6 shows the vertex Z distributions for the two beams after all analysis cuts (including the cut on Z). The cuts are the same as for charged mode.

Figure 7.7 shows the kaon energy distributions for the two beams after all analysis cuts. In the interval $40 < E_K < 160 \text{ GeV}$, there are a total of 868,572 candidate events in the vacuum beam and 1,451,451 candidate events in the regenerated beam.




Figure 7.4: Four-photon invariant mass distributions after all other analysis cuts have been applied. The dotted lines show the cut made on this quantity.

7.5 Adjustment of the Photon Energy Scale

Reconstruction of $K \rightarrow \pi^0\pi^0$ events depends critically on the correct measurement of photon energies and positions by the calorimeter. As described earlier, the calorimeter is calibrated using electrons, and several corrections are made to smooth out the calorimeter response based on studies with electrons. The final step in the neutral-mode analysis is to ensure that the calorimeter measures *photon* energies correctly.

A small difference ($\sim 0.1\%$) between electrons and photons is expected from the slightly different electromagnetic shower profiles in the CsI blocks, coupled to




Figure 7.5: Neutral-mode ring number distributions after all other analysis cuts have been applied. The dotted line shows the cut made on this quantity.

the nonuniformity of light collection. This is taken into account in the calibration procedure. The overall energy scale is then checked by examining the Z distribution of $\pi^0\pi^0$ events near the downstream edge of the Regenerator. As Equation 7.2 indicates, the reconstructed Z position is directly related to the measured photon energies. Thus, a systematic shift in the energy scale would shift the apparent position of the Regenerator edge. Figure 7.8a shows that there is initially a shift of about 7.6 cm between data and the Monte Carlo simulation, corresponding to an average energy scale error of 0.125%.

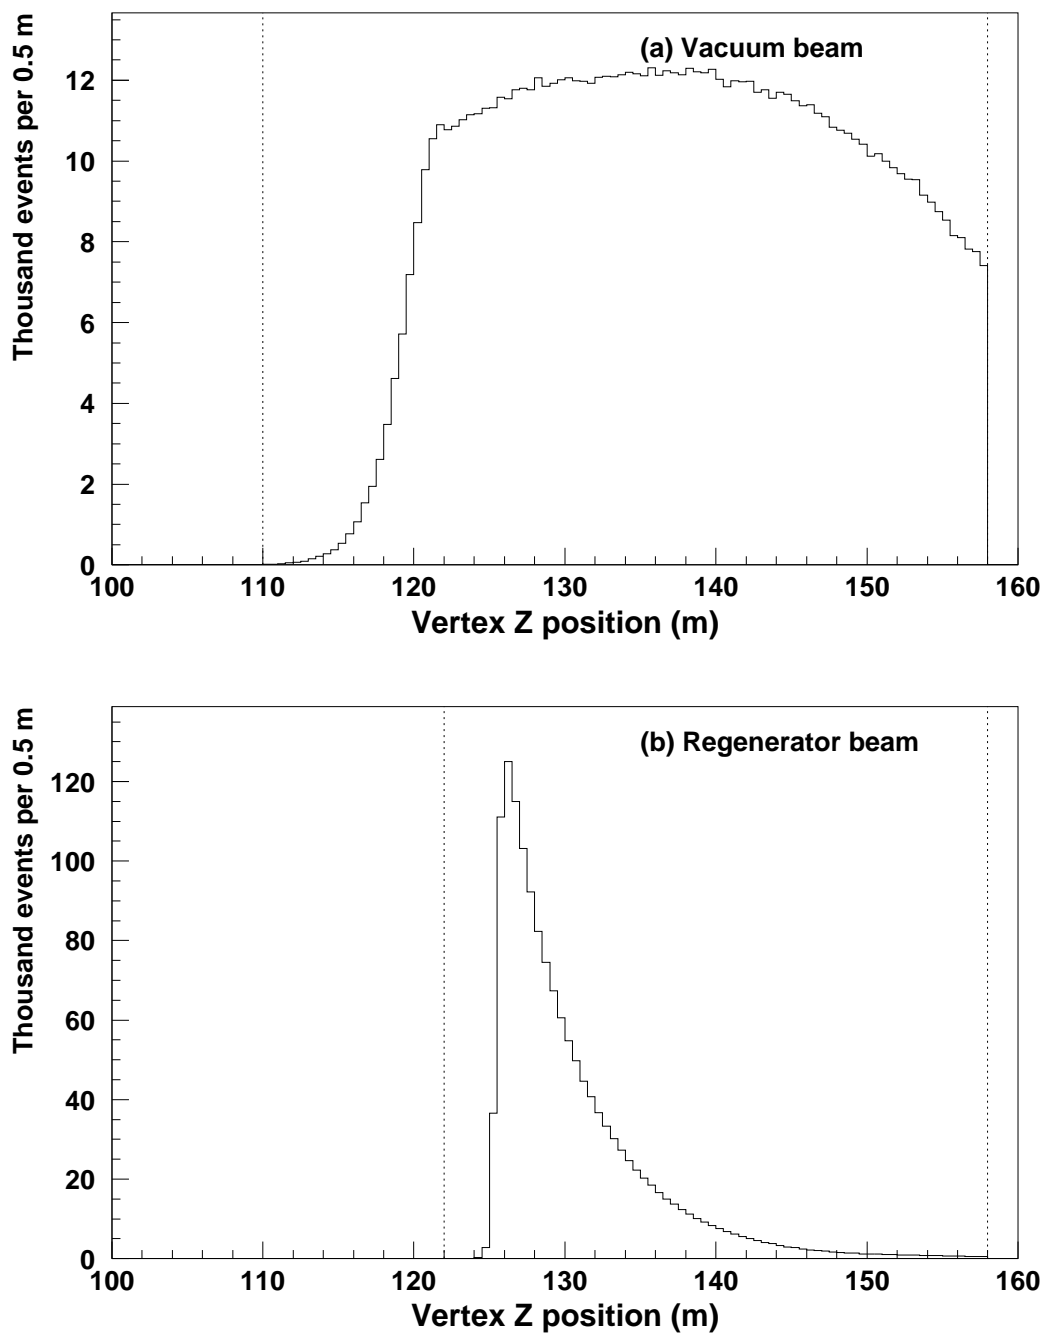


Figure 7.6: Vertex Z distributions for the neutral-mode data after all analysis cuts (including the cut on Z) have been applied. The dotted lines show the cut made.




Figure 7.7: Kaon energy distributions for the neutral-mode data after all other analysis cuts have been applied. The cutoffs at 30 and 180 GeV are from cuts made when the data were “crunched”. The dotted lines show the interval used to define the final samples for the measurement of $\text{Re}(\epsilon'/\epsilon)$.

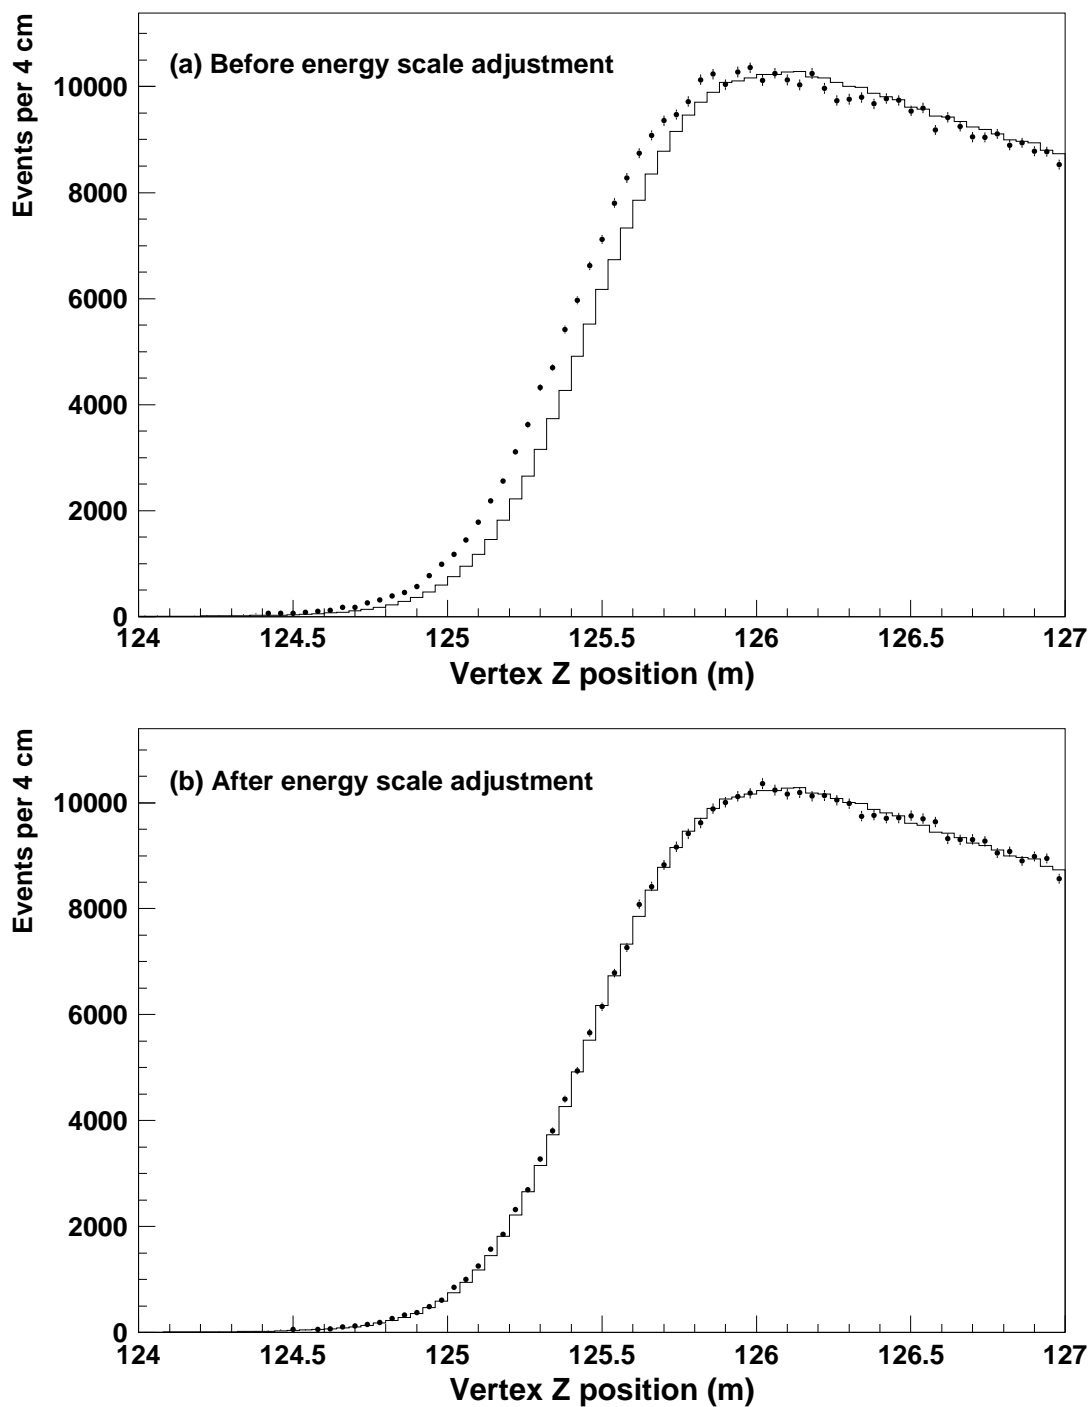


Figure 7.8: Using the Regenerator edge to adjust the photon energy scale. The dots are data, and the histograms are Monte Carlo.

This edge comparison is used to adjust the energy scale,³ and the whole analysis is re-run and cuts re-evaluated. Separate adjustment factors are used for each 10-GeV intervals of kaon energy from 40–160 GeV, ranging from 0.18% (for kaon energies between 40 and 50 GeV) to 0.07%, and the adjustment is applied to all four photon energies in the event.⁴ After the adjustment, the agreement at the Regenerator edge is excellent, as shown in Figure 7.8b.

It is not known why the photon energy scale does not come out right in the first place. The correction needed is about the same size as the expected electron-photon difference, but it is not clear whether this is relevant. In Chapter 10 we will assign a significant uncertainty on $\text{Re}(\epsilon'/\epsilon)$ due to a possible residual misunderstanding of the energy scale.

7.6 Background Subtraction

7.6.1 Background Processes

In Section 7.4 we described several cuts designed to reduce the background from $K_L \rightarrow 3\pi^0$ decays. Even after these cuts, there is a significant background in the vacuum beam, visible in Figure 7.4.

As in charged-mode, the regenerated-beam background is dominated by kaons which scatter in the Regenerator before decaying to $\pi\pi$. However, in this case the background level is much larger, because there is no measurement of p_T^2 and the ring number variable cannot distinguish scattered kaons which stay within the beam. There is also a substantial background in the *vacuum* beam from kaons which scatter in the Regenerator. This is visible in Figure 7.9, which shows the distribution of the position of the energy centroid at the CsI when the Regenerator is on the

³The fact that the Monte Carlo does not simulate kaon decays within the Regenerator is taken into account; this modifies the naive correction by about 0.01%.

⁴In reality, the energy mismeasurement is most likely a function of photon energy rather than kaon energy. It is possible, in principle, to determine a correction as a function of photon energy using $K_L \rightarrow \pi^+\pi^-\pi^0$ decays, but is tricky if one also wants to let the correction depend on position in the calorimeter, etc. Therefore we have stuck with a simple correction.

right side. The broad component, centered on the right (regenerated) beam, is from scattered kaons.

It is important to note that Regenerator scattering is characterized using $K \rightarrow \pi^+\pi^-$ data from the 1996 running period, even though this data is not used to extract $\text{Re}(\epsilon'/\epsilon)$. The cuts made on the Ring Counters and Spectrometer Antis are the same as those made for the $K \rightarrow \pi^0\pi^0$ data, to help ensure that the scattering distribution measured from the charged data accurately describes the scattering




Figure 7.9: Position of energy centroid at CsI, showing Regenerator scattering extending under the vacuum beam. Plot made using $K \rightarrow \pi^0\pi^0$ events for which the Regenerator is on the right side of the detector, but without the final offline cuts on energy in the Regenerator. Thus the scattered background is greater in this plot than it is in the final event sample.

present in the neutral data. The scattering characteristics could, in principle, be different between 1996 and 1997, although we find that they are rather similar.

There is also a significant background from kaons which scatter in the final collimator before decaying to $\pi^0\pi^0$. It is hard to distinguish this from Regenerator scattering in the $\pi^0\pi^0$ data, but we know it must be there because we see it clearly in the vacuum-beam $\pi^+\pi^-$ sample.

Finally, there is a small background from “Regenerator junk”, inelastic interactions in the Regenerator which may produce pairs of π^0 's. These are empirically observed to have a flat $\pi^0\pi^0$ invariant mass distribution and exponential dependences on kaon energy and p_T^2 . A model built along these lines, with parameters measured from the data, was incorporated in the Monte Carlo.

It occasionally happens that the photons from a real $K \rightarrow \pi^0\pi^0$ decay are paired incorrectly, yielding a good ring number but a bad reconstructed mass. We do not consider this a background, but just a part of the signal distribution.

After the background-subtraction procedure was finalized, an additional class of backgrounds was identified: a $K_L \rightarrow \pi^+\pi^-\pi^0$ decay downstream of drift chamber 1 can leave four clusters in the calorimeter which, on occasion, will reconstruct as a $\pi^0\pi^0$ candidate and pass all analysis cuts.⁵ (Any track in the spectrometer normally causes an event to be rejected, but a decay downstream of chamber 1 does not leave enough information to reconstruct tracks.) The resulting background is small and looks rather like the background from $3\pi^0$ decays, so the current background-subtraction procedure essentially counts it as $3\pi^0$ background. The procedure may be modified in the future to treat this background separately.

Finally, there are tiny backgrounds expected from $K_L \rightarrow \pi^0\gamma\gamma$, and from $\Xi^0 \rightarrow \Lambda\pi^0$ with $\Lambda \rightarrow n\pi^0$, in the vacuum beam. These are not subtracted, and thus will contribute slightly to the systematic error on $\text{Re}(\epsilon'/\epsilon)$.

⁵There is also evidence for hadronic interactions in drift chamber material which produce a combination of neutral and charged particles, but no reconstructable tracks.

7.6.2 Normalization of Background Contributions

The various background contributions are normalized and subtracted in a particular order which decouples them as much as possible. First, the $3\pi^0$ background is normalized using sidebands in the mass distribution (separately for the vacuum and regenerated beams, and after accounting for the $\pi^0\pi^0$ mis-pairing contribution to the mass 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 Regenerator junk. The simulation of the $3\pi^0$ background does match the shape of the mass distribution, as shown in Figure 7.10.

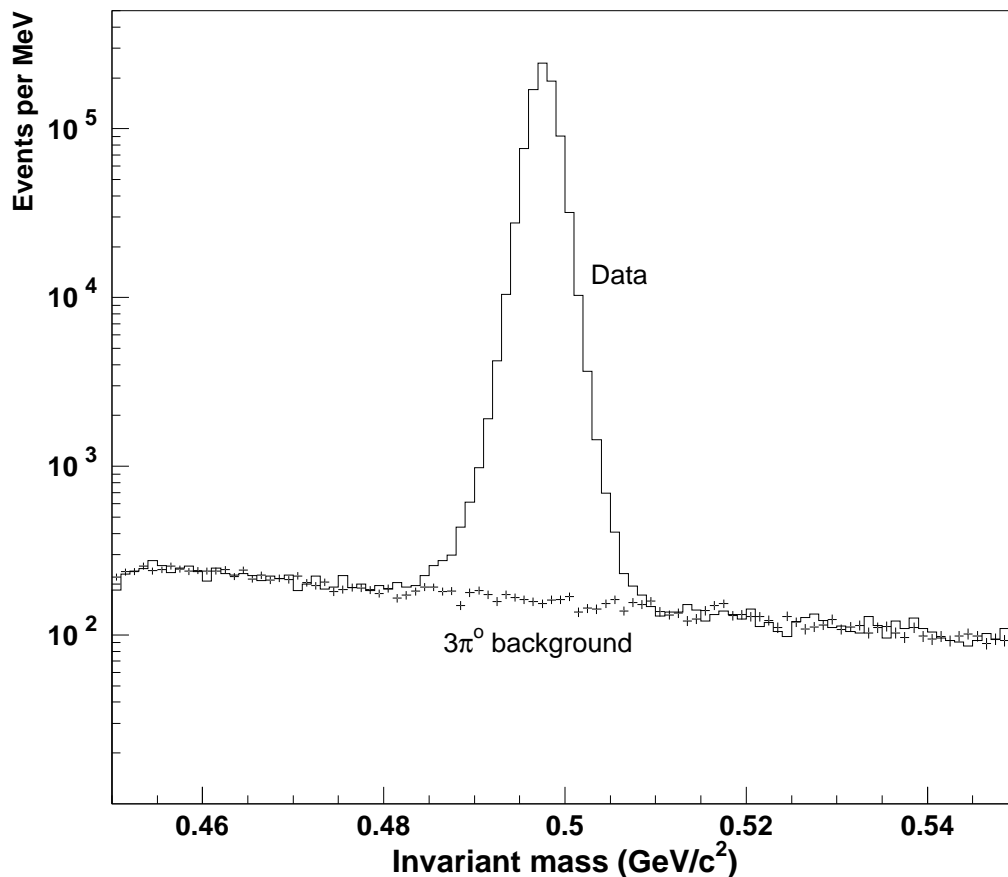


Figure 7.10: Vacuum-beam mass distribution for data and $K_L \rightarrow 3\pi^0$ background after all other analysis cuts have been applied.

The next step is to normalize the Regenerator-junk background. This is done using the tails on the mass distribution for regenerated-beam events with $125 < Z < 126$ m, where essentially all of these background events reconstruct.

Next, the collimator-scattered background is subtracted. This background cannot be normalized cleanly using the $\pi^0\pi^0$ data, so the absolute normalization is taken from the Monte Carlo model, which was tuned with $K \rightarrow \pi^+\pi^-$ data from 1997. We will assign a systematic uncertainty due to the fact that the amount of collimator scattering could be different for 1996 vs. 1997.

Finally, after subtracting all other known backgrounds, we can conclude that the remaining ring-number tail is due to Regenerator scattering. We normalize the MC simulation to the data using the ring-number range from 300 to 800. We do the normalization separately for the vacuum and regenerated beams; the normalization factors agree fairly well.

The sum of all the background contributions does a good job of reproducing the ring-number distributions observed in the data, as shown in Figure 7.11. The fractional backgrounds contributed by each background process are summarized in Table 7.1.

7.7 Summary

After background subtraction, the net yields are 862,254 $\pi^0\pi^0$ in the vacuum beam and 1,433,923 $\pi^0\pi^0$ in the regenerated beam. We have now finished analyzing all the data needed to measure $\text{Re}(\epsilon'/\epsilon)$, but we still need to determine the acceptance of

Table 7.1: Backgrounds in the $\pi^0\pi^0$ samples.

Background process	Vac beam	Reg beam
$K_L \rightarrow 3\pi^0$, lost or fused photons	0.27%	0.01%
Kaon scattered in Regenerator	0.30%	1.07%
Kaon scattered in collimator	0.16%	0.14%
Regenerator junk (π^0 pairs from interactions)	~ 0	0.01%
TOTAL	0.73%	1.22%

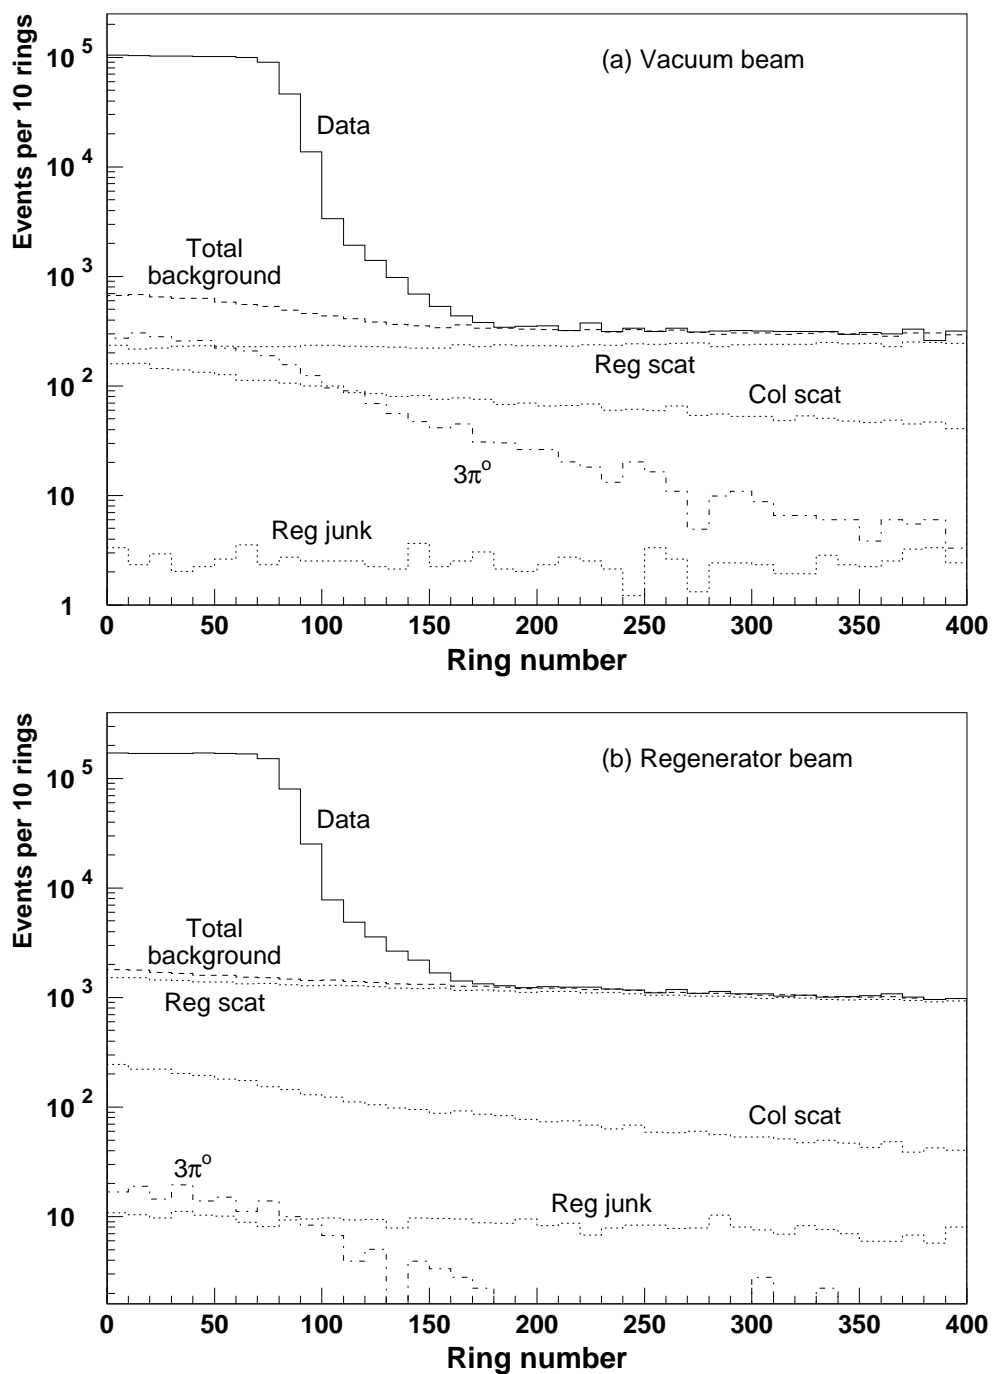


Figure 7.11: Ring number distributions for data and background components after all other analysis cuts have been applied.

the detector and analysis cuts so that we can infer the actual number of kaon decays within the decay volume. We do so using a Monte Carlo simulation, as described in the next chapter.

CHAPTER 8

THE MONTE CARLO SIMULATION

The KTeV Monte Carlo (“KTEVMC”, or simply “the MC”) is a computer program which simulates kaon decays and the response of the KTeV detector in considerable detail. It simulates decays with the same energy and Z distributions as the data and produces output that looks essentially the same as real KTeV data, but with a few extra data banks added to the event format to store information about the simulated event for special studies. Monte Carlo event samples are analyzed as if they were data, with the same reconstruction algorithms (except for some minor differences, described in Section 8.8) and analysis cuts.

In the previous two chapters, we saw that the Monte Carlo is used to simulate various backgrounds to the $\pi^+\pi^-$ and $\pi^0\pi^0$ data samples, often using models derived from studies of the data. However, the most crucial use of the MC for the measurement of $\text{Re}(\epsilon'/\epsilon)$ is to determine the acceptance for the $\pi\pi$ decays, in other words, the fraction of actual decays occurring within the decay volume which are successfully reconstructed and pass all analysis cuts. The acceptance is largely determined by the geometry of the detector and by geometric analysis cuts (*e.g.* apertures and particle-particle separations); however, to understand reconstruction biases it is important to simulate the detector response accurately. We use data whenever possible—often the high-statistics $K_L \rightarrow \pi e \nu$ and $K_L \rightarrow 3\pi^0$ decay modes—to check or tune various aspects of the detector geometry and response, such as the spectrometer and calorimeter resolutions. The MC is also useful for doing systematic studies since it allows us to modify the detector response in a known way and to measure the resulting effect on $\text{Re}(\epsilon'/\epsilon)$.

In this chapter, we will describe the Monte Carlo simulation in some detail. The basic structure, and many of the core routines, are essentially unchanged from the

Monte Carlo programs used for earlier kaon experiments at Fermilab (E731, E773, E799-I). However, the detector simulation code is almost entirely new for KTeV.

8.1 Kaon Production, Evolution and Decay

KTEVMC is capable of generating a wide variety of kaon (and hyperon) decay modes, although ordinarily only one decay mode is generated at a time. For kaons, the code keeps track of the quantum state as the kaon travels along the beamline, and takes it into account when determining the decay distribution.

8.1.1 Kaon production

Simulation of an event begins by choosing a beam (left vs. right, and vacuum vs. regenerated) and an initial kaon eigenstate (K^0 vs. \bar{K}^0). Next, the kaon energy and production angle relative to the primary proton beam are chosen from a combined distribution which is motivated by a parametrization by Malensek [67] of measurements of K^+ and K^- production by 450 GeV protons incident on a beryllium target [68]. Using the valence quark content of the incident proton as a guide, the forward production of neutral kaons is assumed to be related to charged kaon production as follows:

$$K^0 \sim (K^+ + K^-)/2 \qquad \bar{K}^0 \sim K^- \qquad (8.1)$$

The energy distribution given by this parametrization is modified by a polynomial correction in order to match the distributions observed for $K_L \rightarrow \pi^+\pi^-$ decays in KTeV vacuum beam data. The magnitude of this correction differs from unity by $\pm 8\%$ over the energy range of interest (40–160 GeV).

The KTeV neutral beams (after collimation) include a range of production angles, from 4.55 mrad at the bottom to 5.05 mrad at the top. The production model predicts a variation of the flux and energy distribution over this range, and such variations are observed in the data. However, it was found to be necessary to make a small linear correction to the production angle distribution, amounting to +0.8%

at the bottom edge of the beam and -0.8% at the top edge, in order to match the observed beam profile.

8.1.2 *Kaon transport*

The kaon is assigned an initial position within the production target, with the transverse distribution based on the size of the primary proton beam and an exponential distribution along the length of the target based on the proton and kaon interaction cross sections in beryllium. The kaon is then propagated downstream toward the detector. It first passes through the absorbers, including the movable absorber if a regenerated-beam kaon is being simulated. There is a chance of scattering in each absorber, in which case the kaon is given a transverse momentum chosen from an exponential distribution before continuing downstream.

The kaon is propagated through the neutral-beam collimators, whose positions are input from survey information (with cross-checks against the observed beam profiles). The upstream collimators are treated as perfectly absorbing, but a kaon hitting the final collimator is given some probability to scatter and re-enter the neutral beam channel according to the model described in Section 6.5.1. If the scattering is inelastic, then the kaon quantum state is changed to pure K_S .

The attenuation of kaon flux by the Regenerator and movable absorber was tuned by comparing $K_L \rightarrow \pi^+\pi^-\pi^0$ and $K_L \rightarrow 3\pi^0$ data from the vacuum and regenerated beams. In particular, the net transmission was found to decrease with increasing kaon energy, by about 5×10^{-4} per GeV (as a fraction of the average transmission).

8.1.3 *Evolution of the kaon quantum state*

As a kaon travels from the target to the detector, its quantum state evolves from the initial K^0 or \bar{K}^0 into a state which is nearly all K_L , with negligible K_S component except at the highest kaon energies.¹ KTEVMC uses K_S and K_L basis states, and calculates an exact transformation matrix for propagation along each segment

¹This residual K_S component is the only thing which reflects whether a K^0 or a \bar{K}^0 was initially produced.

of the beamline. This matrix is diagonal for the vacuum segments, representing simply the decay of the two states over time. When the kaon traverses material, the transformation matrix acquires additional on-diagonal contributions representing losses due to interactions, plus off-diagonal terms representing coherent regeneration.

Regeneration in the beamline absorbers is simulated but is not particularly important,² while the simulation of the KTeV Regenerator is an important task of KTEVMC. For the extraction of $Re(\epsilon'/\epsilon)$ and other physics parameters, KTEVMC is used only to calculate the detector acceptance in bins of kaon energy and Z , and so it is not critical that KTEVMC produce exactly the right distributions. Nevertheless, we attempt to simulate regeneration as well as possible. The regeneration amplitude for plastic scintillator, which makes up the bulk of the Regenerator, is based on preliminary fits to KTeV data and has a power-law dependence on the kaon energy over the region of interest. The regeneration phase is determined from the power-law coefficient by assuming analyticity of the forward scattering amplitude. Two small lead pieces at the downstream end of the Regenerator (used to convert photons from $K_L \rightarrow \pi^0\pi^0$ decays within the Regenerator) contribute significantly to the total regeneration, and are simulated using regeneration parameters measured in an earlier experiment [69]. The resulting kaon state downstream of the Regenerator is a coherent mixture of K_L and K_S , and KTEVMC reproduces the characteristic interference behavior in $\pi\pi$ decays.

Evolution of the kaon quantum state, as described above, reflects only the results of coherent forward scattering; all other scattering and inelastic interactions are treated as losses. However, we also use KTEVMC to simulate the background from kaons which scatter in the Regenerator. The hardware trigger and the offline analysis cuts reject most inelastic interactions, leaving mostly diffractive scatters from carbon and lead. The scattering probabilities and distributions for events which survive are simulated according to the physics-motivated model described in Section 6.5.1. In particular, the quantum state for a scattered kaon is modified according to the effective regeneration amplitude for scattering of that type.

²The absorbers do affect the observed K_L energy spectrum somewhat due to regeneration of K_S (from the target) into K_L .

8.1.4 Kaon decay

Each kaon generated is forced to decay within a specified region of the detector, slightly larger than the fiducial region to be selected for analysis. This approach makes efficient use of CPU time, but adds some complexity since weighting factors must be calculated at the beginning of the job in order to end up with the correct proportions of vacuum-beam vs. regenerated-beam coherent decays, diffractive and inelastic scatters from the Regenerator, and inelastic collimator scatters, all of which produce different quantum states at the beginning of the decay region. The Z position of the decay is chosen from the expected distribution, which is a function of the kaon energy and quantum state, including the interference term between the K_L and K_S components when appropriate. Decays within the Regenerator are not simulated in KTEVMC.

Each kaon or hyperon decay mode has a separate subroutine which generates a decay in the rest frame of the parent particle. The complexity of these subroutines varies considerably, depending on the physics involved. For example, the simulation of $K \rightarrow \pi^+\pi^-$ includes an appropriate fraction of $\pi^+\pi^-\gamma$ from internal bremsstrahlung. Three-body decays such as $K \rightarrow \pi^+e^-\bar{\nu}$ and $K \rightarrow \pi^+\pi^-\pi^0$ include form factors which affect the resulting kinematic distributions. If a π^0 is produced, it is forced to decay immediately, and options exist to simulate the Dalitz decay $e^+e^-\gamma$ (with full radiative corrections) or various rare decay modes if desired. After decay products have been generated, they are boosted into the lab frame for tracing through the detector.

8.2 Tracing of Decay Products

The products of each kaon decay (except neutrinos) are traced through the KTeV detector in sections. A charged pion is permitted to decay to $\mu\nu$ according to its lifetime, and the muon is traced through the remainder of the detector. Charged particles receive a transverse momentum kick from the analysis magnet, based on a field-integral map which is a function of the transverse position of the particle at the center of the magnet.

Charged particles undergo multiple scattering as they pass through the vacuum window and the material in the spectrometer, with a crude parametrization of a non-Gaussian scattering component. Electrons may undergo bremsstrahlung, losing some energy and producing a photon which is traced through the remainder of the detector. Photons are permitted to convert into e^+e^- pairs.

Synchrotron radiation from electrons in the field of the analysis magnet has been observed in the data (from hits left in the drift chambers by low-energy synchrotron photons) [70], but is not simulated.

Energy deposit in the photon-veto detectors and trigger hodoscope counters is simulated during tracing. Depending on the trigger configuration in use, sufficient energy deposit in a photon-veto detector can cause the event to be rejected immediately, saving CPU time when simulating backgrounds with extra photons.

A particle (other than a muon) which hits a photon-veto detector or some passive material, or any particle which escapes the detector volume, is no longer traced. The user can specify that an event should be rejected immediately in such a case, which can conserve CPU time when a physics signal mode is being generated.

KTEVMC stores the position of each particle at the drift chambers and calorimeter for more detailed simulation later, as described in the sections below. Tracing ends when a particle hits the calorimeter or the Back Anti, except for muons, which are propagated through the downstream material (calorimeter and steel muon filters) in a separate step.

8.3 Simulation of the Drift Chambers

The drift chamber simulation started out fairly simple, but was made much more sophisticated to reproduce the detailed behavior observed from the data. The basic approach is to calculate the distance to the closest wire in a given sense plane, smear the distance with the measured resolution for this plane,³ and convert the modified distance to a drift time. However, there are two additional ways in which the drift

³The position resolution is assumed to be uniform across the drift cell and to be the same for all wires in a sense plane, neither of which is necessarily true.

distance may be modified before converting it to a drift time, which we will now describe.

We account for the discrete ionization of the drift chamber gas, which generally makes the shortest drift distance a bit longer than the normal distance between the wire and the particle track. The true ionization density in argon is known to be about 2.9 electrons per mm; however, the drift chambers were not always sensitive to a single drift electron, especially in the beam regions, so in practice we used the data to measure an *effective* ionization density for a few different regions in each plane and used this in the simulation.

The most elaborate modification to the drift distance is a simulation of the occasional mismeasurement of a drift time which is responsible for the “high-sod effect” described in Section 5.2.3. From studies of the $K \rightarrow \pi^+\pi^-$ data, the probability of such a mismeasurement is found to depend on several things. First, there is a spatial dependence, shown in Figure 8.1a for one particular sense plane in drift chamber 1. The beam regions are clearly visible as regions of enhanced high-sod probability, as are a scattering of point-like “freckles”. We make “maps” such as this one, separately for the two Regenerator positions (because we find that the higher particle flux slightly enhances the probability in the vacuum-beam region), and incorporate them into the Monte Carlo simulation. The probabilities vary as a function of time, as shown in Figure 8.2, and so we apply a multiplicative correction to the high-sod probabilities in each plane depending on the run number.⁴ Note that the greatest change in high-sod probability occurred when the high voltage was lowered for chamber 3. We have also implemented a dependence on the position of the track in the drift cell, since the data indicates that the high-sod probability is much higher close to the wire. When the simulation decides that a hit is to be mismeasured, we increase the effective drift distance by a value chosen from an exponential distribution whose slope is measured from the data.

Sometimes a wire records no hit at all, and we map out the probability of this occurring as a function of position and time in a similar fashion. One such map is

⁴One wire in DC4X behaves differently from the rest of the wires in that plane, so we allow it to have an independent time variation.

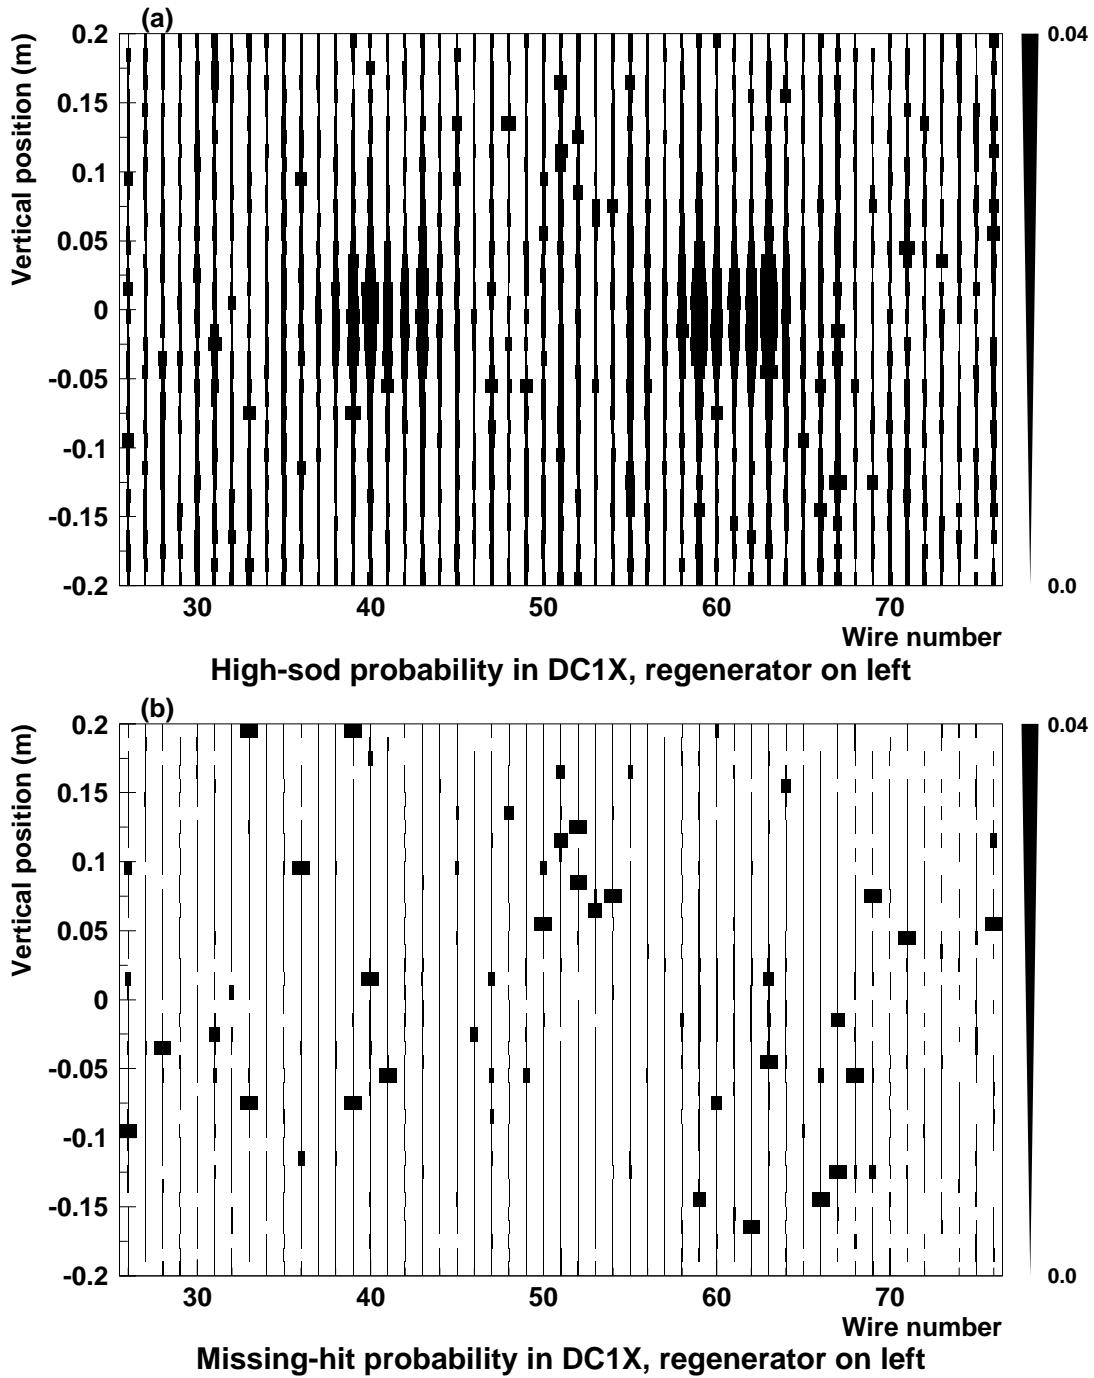


Figure 8.1: Spatial dependence of high-sod (top) and missing-hit (bottom) probabilities in drift chamber 1X. The probabilities are measured in 1 cm position bins along each wire; the line width in the figure above is proportional to the local probability. Only the central region (64×40 cm) of the drift chamber is shown.

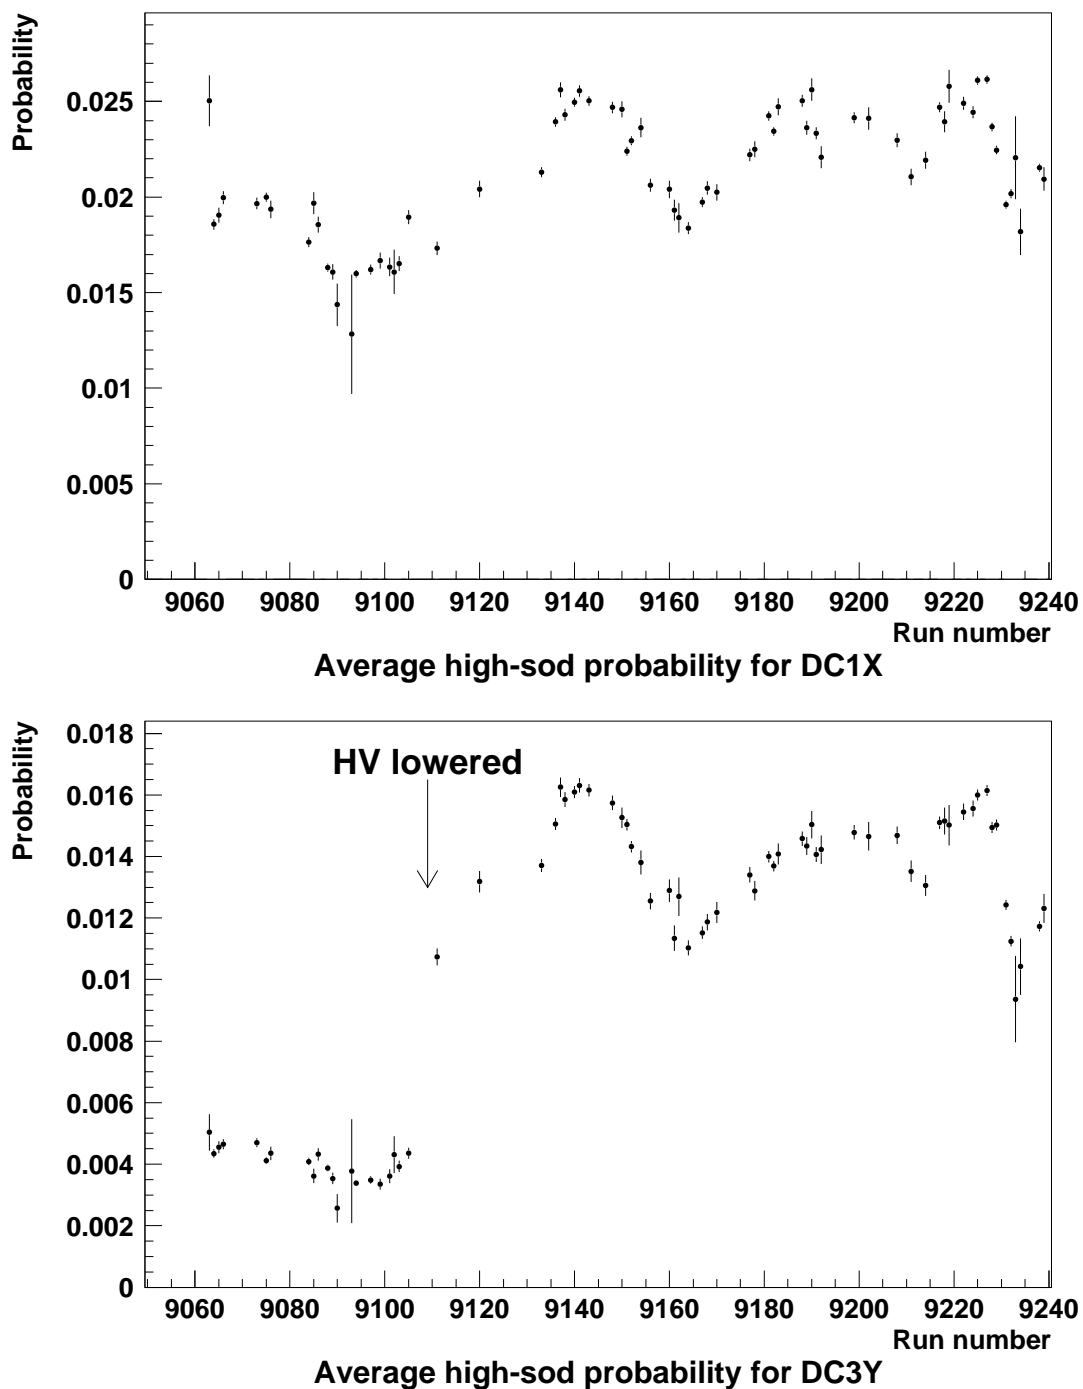


Figure 8.2: Time dependence of high-sod probability in drift chambers 1X and 3Y. The arrow indicates when the DC3Y high voltage was lowered from 2500 V to 2450 V.

shown in Figure 8.1b; note that most of the inefficiency is concentrated in “freckles”, which are in the same places as in the high-sod map and may have missing-hit probabilities of up to $\sim 10\%$. Unlike the high-sod case, the missing-hit probability is greatest when the particle is *far* from the wire.

The high-sod and missing-hit simulation dramatically improves the agreement with the data, especially in the tails of certain distributions. Figure 8.3 shows an early check of this for the p_T^2 variable, before final tuning of the simulation. We find that the agreement becomes even better when we scale up the high-sod and missing-hit probabilities by a factor of 1.25, and so we did this when generating the large MC samples for the $\text{Re}(\epsilon'/\epsilon)$ measurement. In Chapter 10 we will assign a systematic error due to the uncertainty in this *ad hoc* scale factor.

The drift chamber simulation also includes a model for generating a “delta ray” above a certain cutoff energy, which can produce an extra hit on a wire. This hit may arrive before the hit from the original particle, causing the drift distance to be underestimated when the event is reconstructed.

Finally, KTEVMC models the response of the discriminator cards to multiple pulses on a wire. This is important because an “accidental” hit before the beginning of the in-time window can obscure the hit from the simulated particle. The discriminators were designed to have a deadtime of about 50 ns, but we have observed that a hit sometimes can prevent retriggering for an even longer period of time, presumably due to an especially broad analog pulse which stays above the discriminator threshold for a long time. The data is described reasonably well by a non-retriggering probability which is unity up to 46 ns and then falls exponentially. A sample curve, measured for DC1X, is shown in Figure 8.4. Each hit for which the discriminator triggers is digitized after applying a correction due to the signal propagation time along the wire and a wire-dependent time offset (measured from the data, and taken out when reconstructing events).

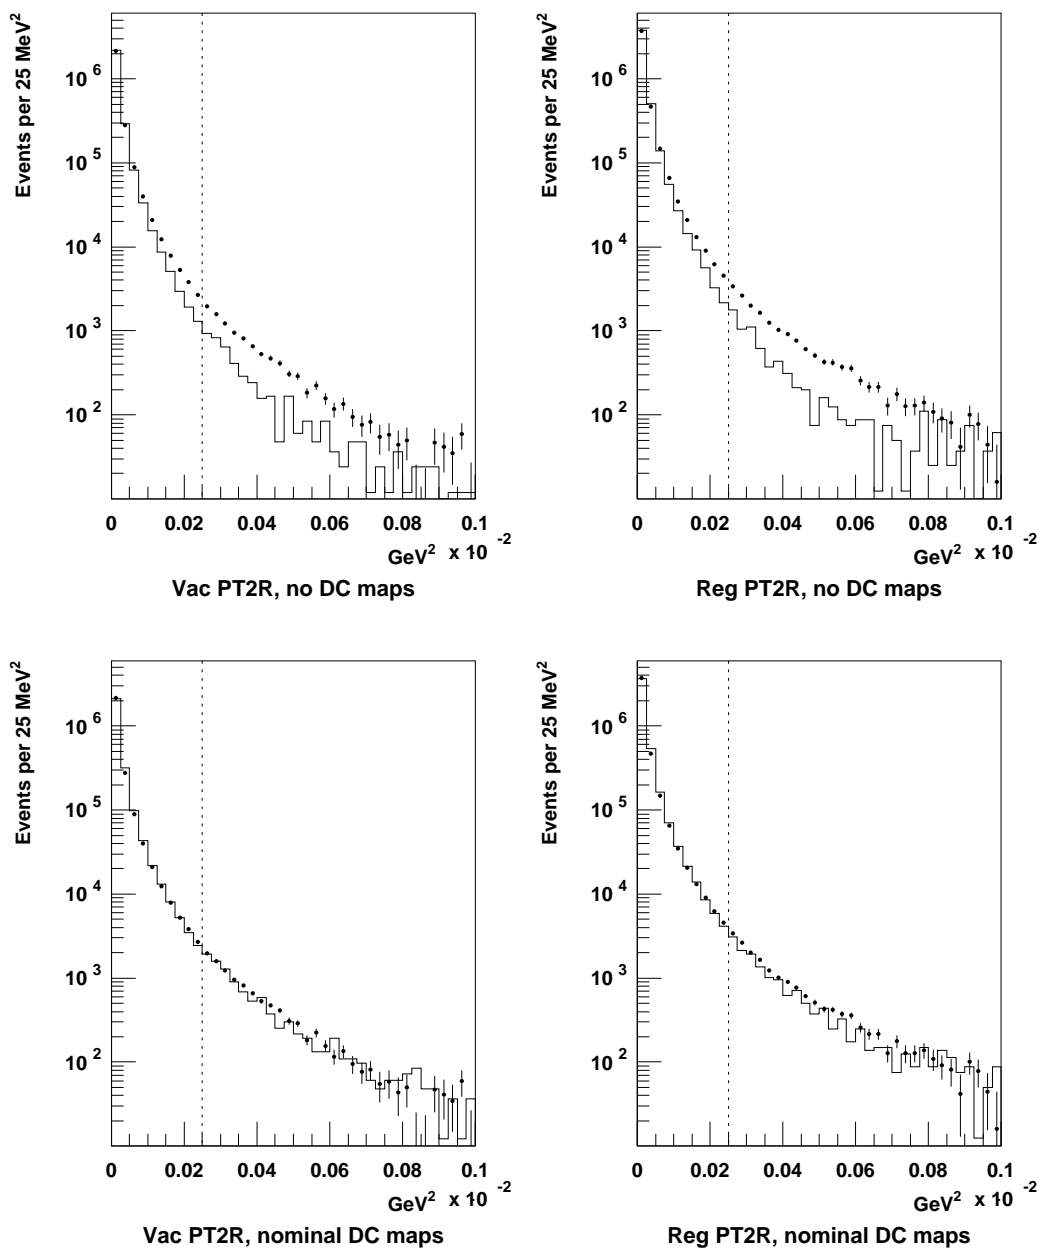


Figure 8.3: Effect of the high-sod/missing-hit simulation on the p_T^2 distribution. The top two plots compare data (dots) and Monte Carlo (histograms) *without* the high-sod/missing-hit simulation, while the bottom plots compare data and Monte Carlo *with* the simulation. The dotted lines show the analysis cut at $250 \text{ MeV}^2/c^2$.




Figure 8.4: Drift chamber retriggering model used in KTEVMC. The probability of NOT retriggering is unity for times up to 46 ns and falls exponentially after that. Somewhat different exponentials are measured for the different drift chambers.

8.4 Simulation of the Calorimeter

The high precision achieved by the KTeV calorimeter when reconstructing photon and electron energies and positions, and the level to which we need to understand potential systematic effects, demand a very accurate simulation of the calorimeter response. We are concerned not only with first-order performance measures, such as the RMS energy resolution, but also with the full non-Gaussian behavior and tails on the response which result, in part, from fluctuations in the development of electromagnetic showers. Therefore, we make use of the GEANT package [71], which simulates all relevant interaction processes in the CsI crystals (pair production,

bremsstrahlung, etc.) and tracks all shower particles down to energies of a few MeV.⁵

Since an EM shower involves thousands of particles, its simulation is rather time-consuming, and we cannot afford to do the full simulation for the many millions of kaon decays we need to generate with KTEVMC. Therefore, we developed an alternative scheme using a “library” of pre-generated showers which is accessed by KTEVMC as needed. The library contains several thousand showers generated using GEANT from photons and electrons at six different energies (2, 4, 8, 16, 32, and 64 GeV) and incident at 325 distinct positions on the face of the central CsI crystal. The information stored for each shower includes the energy deposited in a 13×13 array of 2.5×2.5 cm crystals, as well as the longitudinal profile of the energy deposit for the central 7×7 region (collectively) and for the four crystals with the most energy (individually). A separate library contains pion “showers” (which sometimes are simply minimum-ionizing energy deposit), generated with GEANT using the FLUKA package for hadronic interactions (rather than the default GHEISHA package, which was found to have some shortcomings). The pion showers were generated at more energies and fewer positions, but otherwise are handled in a manner similar to the photon and electron showers.

For each particle which reaches the CsI, KTEVMC selects a shower from the library according to the particle type, energy, and transverse position in the central CsI crystal. KTEVMC “interpolates” between the discrete shower energies in the library by randomly selecting a shower with lower or higher energy than the actual particle energy, using weights calculated from the logarithms of the energies; thus, the mixture of MC showers used is a continuous function of energy. For purposes of selecting a shower for a particle incident at an angle, the transverse position is calculated at an appropriate depth within the crystal, roughly the energy-weighted mean depth of the shower but with a correction determined from a study of electrons incident at significant angles.

After KTEVMC retrieves a shower from the library, it scales the block-by-block

⁵We could also have used the EGS package [72], which was used to simulate the response of the E731 lead-glass calorimeter.

energies by the ratio of the incident particle energy to the energy with which the shower was generated. A large CsI block is given the sum of the energies of the four small-block units which it encompasses. A small amount of energy is also assigned to blocks outside of the 13×13 region covered by the shower-library info, using a parametrization which depends on the distance from the core of the shower. For crystals near the core of the shower, the longitudinal energy-deposit profile is convolved with the longitudinal light-collection efficiencies measured for those crystals, so that shower fluctuations can couple to light-collection nonuniformities to reproduce nonlinearities or tails on the energy resolution.

The total simulated light output for each crystal is given a time profile and divided into DPMT time slices. The signal in each time slice is smeared according to photostatistics and then digitized using code which simulates the multi-ranging behavior of the DPMT. The hardware readout threshold is applied (*i.e.* the sum of four consecutive DPMT slices must be at least 6 counts above pedestal), so that the digitized calorimeter information looks exactly like real data.⁶

8.5 Simulation of the Veto Counters

The Regenerator scattering parametrization measured from the data and used in KTEVMC is specifically for those events which deposit no significant energy, so that they pass the trigger and all offline cuts. Therefore, we do not need to try to simulate energy deposit in the Regenerator due to kaon scattering.

Energy deposit by photons in the Mask Anti, Ring Counters, and Spectrometer Antis is simulated with an energy-dependent Gaussian resolution parametrized from the data. There is a fixed resolution for minimum-ionizing particles which varies from about 30% to 40%, depending on the counter.

Muons are traced through the steel muon filters and will reach the muon veto hodoscope unless they are stopped in the steel. Energy loss in the steel is modeled

⁶Because the KTEVMC digitization procedure mirrors the hardware so closely, it should be possible to simulate some of the common DPMT failure modes, such as a QIE which is unable to select one particular range. However, we have not attempted to do so, and have chosen instead to discard all data collected while there was a malfunctioning channel in the calorimeter.

as Bethe-Bloch ionization loss, with Landau fluctuations from a parametrization determined from studies with GEANT. Thus we expect the simulation to do a reasonably good job of reproducing the probability of a muon to be stopped as a function of energy, but we do not actually make use of this aspect of the simulation when determining the background from $K_L \rightarrow \pi\mu\nu$.

8.6 Accidental Overlays

Given the high flux of kaons and neutrons through the KTeV apparatus, it is natural that some events with a real kaon decay will also contain a significant amount of underlying activity in the detector. For example, the beam incident on the Regenerator has an interaction rate of order 1.5 MHz, and many of these interactions produce a large number of particles. Such an interaction normally vetoes the trigger for that bucket, but the particles may still affect an event occurring a few buckets earlier or later through extra hits in the drift chambers and/or extra energy in the calorimeter (within the 12-bucket time interval over which block energies are integrated). Stray particles from far-upstream hyperon decays or from beam interactions in the spectrometer material can also contribute to underlying activity, without leaving a signature in any veto counter.

This underlying activity may lead to biases in the event reconstruction, for example by corrupting tracks or by increasing the energies of calorimeter clusters. Since the underlying activity is likely to be different on the two sides of the detector, there can be a bias on $\text{Re}(\epsilon'/\epsilon)$. In fact, interactions in the Regenerator tend to spray particles into the detector fairly uniformly, and so we see a greater concentration of underlying activity near the *vacuum* beam, which has more kaon decays and a higher hadron flux.

We study the effect of the underlying activity, and automatically correct for it in the acceptance determination, by overlaying an event from the “accidental” trigger (see Section 4.7.3) on top of each simulated decay. That is, the energies in each detector element of the accidental event are added to any energy deposit

from the simulated decay,⁷ the lists of drift chamber hits are merged, and trigger-source information is combined. The summed energies are digitized, while the drift chamber retriggering simulation determines which hits from the merged list will be seen by the TDCs. The trigger is evaluated based on the combined information.

Note that the accidental trigger collects events with the proper distribution of instantaneous beam intensity, so we expect it to accurately reproduce the distribution of underlying activity (unlike, for example, a purely random trigger). Of course, if there was a Regenerator interaction in the accidental event, then the Monte Carlo event will end up being rejected by the trigger or in the offline analysis.

The total effect of accidental activity on $\text{Re}(\epsilon'/\epsilon)$ turns out to be rather small. It causes several percent of the events to be lost due to energy deposited in the Regenerator and other veto counters, but these losses do not distinguish between the two beams. On top of this, accidental activity causes $\pi^+\pi^-$ events to be unreconstructable on the order of 1% of the time, but the loss is nearly identical for vacuum- and regenerated-beam decays. It also causes the clusters in $\pi^0\pi^0$ events to have slightly more energy, which should lead to the Z vertex being reconstructed further upstream, but our procedure for adjusting the photon energy scale using the Regenerator edge (described in Section 7.5) corrects for this on average. The net effect of accidental activity on $\text{Re}(\epsilon'/\epsilon)$, determined from a series of Monte Carlo studies, is only of order 10^{-4} .

8.7 Simulation of the Trigger System

KTEVMC simulates the behavior of all of the trigger elements, including the hit-counting system (Banana and Kumquat modules), Y Track Finder, and HCC. The hit-counting and YTF models are rather robust since their behavior depends on digital information from the drift chambers, with only a few minor subtleties having to do with the timing of accidental hits. In contrast, the HCC simulation is sensitive to the modeling of the analog thresholds for setting HCC bits in the E-Total system.

⁷The calorimeter readout threshold is lower for the accidental trigger than for the rest of the triggers, so even energy below the usual threshold can be added.

The positions and widths of these thresholds have been measured from the data for each calorimeter channel and are used in KTEVMC. A number of problematic channels, which are inefficient for very *large* analog signals, are also modeled.

The trigger decision in KTEVMC is made using the same trigger definitions as were used on-line for the data collection. In fact, the same binary trigger-definition maps are used in KTEVMC, with software performing the memory lookups which were originally done in hardware.

8.8 Generation and Analysis of the Monte Carlo Samples

Large Monte Carlo samples were generated to determine the acceptance for the $K \rightarrow \pi\pi$ decay modes. 74 million $K \rightarrow \pi^+\pi^-$ decays were simulated, using ~ 450 CPU hours on a DEC 8400 computer with 440 MHz Alpha processors. About 20 million events (3.3 times the $\pi^+\pi^-$ data) passed all analysis cuts. 164 million $K \rightarrow \pi^0\pi^0$ decays were simulated, using ~ 400 CPU hours and yielding about 20 million events (10 times the $\pi^0\pi^0$ data) after all cuts.

These MC samples were analyzed with the same programs as the data samples, but with a few differences in event reconstruction. First, track finding used a drift-time-to-distance relationship measured from a sample of MC events, in the same way that it was originally measured from the data, to put the data and MC on an equal footing with regard to slight self-inconsistencies in the calibration procedure. Second, cluster reconstruction used transverse energy distributions and position-finding lookups derived from the GEANT showers used by KTEVMC, whereas for reconstructing data these things were derived from the data. Again, this is so data is reconstructed using data-based calibrations while MC is reconstructed using MC-based calibrations. Third, several track and cluster corrections normally applied to the data (*e.g.* to correct for the fringe field of the analysis magnet, DPMT nonlinearities, time dependence of calorimeter channel gains, etc.) are not applied to the MC samples since KTEVMC does not attempt to simulate these effects. Last, a final *ad hoc* calorimeter linearity correction is measured from the MC sample and used in the same way as for the data.

The MC event samples were not required to pass Level 3 filtering. This probably should have been done, but we did not expect it to matter since the Level 3 filtering code should be perfectly efficient for events which satisfy the offline analysis cuts, except for some slight differences due to the fact that the detector calibration was re-done before the offline analysis. We did studies in which we reconstructed data samples with both online and offline calibrations, and found that the differences were very small and showed no sign of a bias on $\text{Re}(\epsilon'/\epsilon)$. This will be quantified as a systematic uncertainty in Chapter 10.

The Monte Carlo jobs produce tables indicating how many events were generated in each beam as a function of E_K and Z . The analysis jobs produce similar tables indicating how many events passed all cuts. The ratio of these is the acceptance. For reference, the average acceptance (for decays with $40 < E_K < 160$ GeV and $110 < Z < 158$ m) is shown in Table 8.1 for the four samples. The fitting software does not simply use these averages, but makes use of the full (E_K, Z) tables to calculate the acceptance in small (E_K, Z) bins, for reasons to be explained in Chapter 9. Contour plots of the vacuum-beam acceptance, for both charged and neutral modes, are shown in Figure 8.5. It is important to note that, for a given E_K and Z , the acceptance is virtually identical for the vacuum and regenerated beams; a difference may arise due to the different underlying activity on the two sides of the detector, but such differences turn out to be quite small.

Background processes are simulated with KTEVMC, but with a few modifications to make efficient use of CPU time. For example, the $K_L \rightarrow \pi e \nu$ background is simulated normally, but rather than trying to simulate the physics process behind

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

Sample	Average Acceptance
Vacuum-beam $\pi^+\pi^-$	23.13%
Regenerated-beam $\pi^+\pi^-$	30.77%
Vacuum-beam $\pi^0\pi^0$	10.36%
Regenerated-beam $\pi^0\pi^0$	13.05%





Figure 8.5: Acceptance for $K \rightarrow \pi^+\pi^-$ (top) and $K \rightarrow \pi^0\pi^0$ (bottom) decays as a function of E_K and Z .

the rare occurrence of the electron being misidentified as a pion, we simply omit the E/p cut when analyzing this sample. Similarly, when generating the $K_L \rightarrow \pi\mu\nu$ background we disable the muon veto hodoscope rather than trying to predict how often a muon will be stopped in the steel. In both cases, we effectively assume that the particle misidentification probability is independent of energy when we generate the samples, and then we let the background-subtraction procedure handle the energy dependence by normalizing in 10-GeV bins of kaon energy.

No similar dramatic simplification is available when simulating the $3\pi^0$ background, since it arises out of the combinatorics of fused clusters and/or missing low-energy photons. We generated a total of 12.7 *billion* $K_L \rightarrow 3\pi^0$ decays, using a total of over 18,000 hours (~ 2 years!) of CPU time on the three SGI Challenge computers originally used for Level 3 processing, which have a total 24 CPUs. It is likely that KTEVMC could predict the absolute background level fairly well, but it is still necessary to use the mass sidebands to normalize this background.

Finally, large samples of $K_L \rightarrow \pi e\nu$ and $K_L \rightarrow 3\pi^0$ decays were generated to compare with the large statistics of these modes in the data. These allowed us to compare data and MC for various quantities, such as the vertex Z distributions, to check our understanding of the detector performance and the acceptance with high precision. This will be revisited when we discuss systematic errors in Chapter 10.

CHAPTER 9

EXTRACTION OF $\text{Re}(\epsilon'/\epsilon)$

We now know the number of observed events in each of the four $\pi\pi$ samples and the acceptance for each. One might think that we could extract $\text{Re}(\epsilon'/\epsilon)$ from a grand ratio of these numbers,

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

There are two important reasons why this approach is not quite correct. The first is that our regenerated-beam $\pi\pi$ samples do not consist of pure K_S decays; there is significant K_S - K_L interference in the decays of the coherently-regenerated kaons, as well as some outright K_L decays. Thus $N(\text{Reg } \pi\pi)$ is not equivalent to $N(K_S \rightarrow \pi\pi)$. If the acceptance were the same for charged and neutral decays, this would merely “dilute” the measured value of $\text{Re}(\epsilon'/\epsilon)$ by a multiplicative factor, but since the charged and neutral acceptances depend on Z in different ways, there can be an additive bias.

The other reason that we do not calculate $\text{Re}(\epsilon'/\epsilon)$ from a simple ratio is because it would require the *average* acceptance to be known extremely well, which would require that our Monte Carlo simulation match the data *exactly*. Any difference in the kaon energy distribution, for example, would cause us to incorrectly weight the contributions from different kaon energies when calculating the average. Similarly, uncertainties in the physics parameters which affect the Z distribution downstream of the Regenerator (Δm , τ_S , and the regeneration amplitude) would couple to the variation of the acceptance with Z and affect the calculated average acceptance. These things were tuned to be very close to the data before the large MC datasets were generated, so we do not expect much of an effect. However, we would not even

be able to evaluate the effect of varying these things without generating entirely new Monte Carlo samples.

The preceding discussion should make it clear that it is necessary to account for the K_L component in the regenerated beam and to make use of the acceptance as a function of (E_K, Z) rather than an overall average. We have developed a program called KFIT¹ as a flexible tool to allow us to extract $\text{Re}(\epsilon'/\epsilon)$, as well as other physics parameters, from the data in accordance with these principles. The basic approach is to use physics parameters ($\text{Re}(\epsilon'/\epsilon)$, Δm , τ_S , etc.) and a regeneration model to calculate kaon decay distributions, then multiply by the acceptance (determined from KTEVMC) to arrive at a prediction for the number of events in (E_K, Z) bins of a certain size. We determine $\text{Re}(\epsilon'/\epsilon)$ by allowing it and certain other parameters to float while KFIT minimizes the χ^2 describing the agreement between the prediction and the data.

We will now give some details about how KFIT bins the data and calculates the prediction in each bin. We will then describe the specific fit done to extract $\text{Re}(\epsilon'/\epsilon)$ and present the result. This will be followed by a summary of some of the cross-checks we have done.

9.1 The Fitting Software

9.1.1 Inputs and Binning

The inputs to KFIT are 2-dimensional (E_K, Z) tables of background-subtracted data, generated MC coherent events, and reconstructed MC coherent events, for $\pi^+\pi^-$ and $\pi^0\pi^0$ decays in each beam and for each Regenerator position. We combine the tables for the two Regenerator positions with a simple sum; this is valid because the intensities of the two KTeV beams are equal to within a fraction of a percent and the detector acceptance is virtually identical for the left and right beams. (In contrast, the intensities of the two E731 beams differed by 8%; they combined the

¹Many parts of the KFIT code were adapted from the fitting software used in E773 to measure the physics parameters ϕ_{+-} (the phase of η_{+-}) and $\phi_{00} - \phi_{+-}$ [73, 74], which in turn descended from the fitting software used in E731.

data from the two Regenerator positions using the geometric mean to minimize their sensitivity to this asymmetry [75, 76]. For KTeV, using the geometric mean changes the measured value of $\text{Re}(\epsilon'/\epsilon)$ by less than 0.2×10^{-4} .)

As previously mentioned, only events with $40 < E_K < 160$ GeV and $110 < Z < 158$ m are used in the fit. At the beginning of the job, the data arrays are rebinned into “fit bins”, each one 10 GeV wide and covering the entire Z range. The arrays of generated and reconstructed MC events are rebinned into smaller “acceptance bins”, each one $2 \text{ GeV} \times 2 \text{ m}$ in size, and are divided to calculate the acceptance and its error in each bin. The two binning schemes are shown graphically in Figure 9.1.

The use of fairly small acceptance bins minimizes our sensitivity to data/MC mismatches in the overall E_K and Z distributions. The important thing is that the data/MC ratio does not vary much *within* an acceptance bin, and the fact that the E_K and Z distributions generated by KTEVMC match the data quite closely means that our acceptance bins do not really have to be very small. In essence, KTEVMC averages the acceptance over the bin with the proper weighting within the bin.

We choose to use large fit bins, without any division in Z , so that we do not have to worry much about smearing of the reconstructed Z position due to finite detector resolutions. In any case the Z distribution of data events gives very little additional information for determining $\text{Re}(\epsilon'/\epsilon)$ (only a very small amount from helping to distinguish the K_S and K_L components in the regenerated beam). However, the acceptance calculation uses smaller bins, so it is more likely for a MC event to reconstruct in a different bin from the one in which it was generated. This is not a problem since the MC E_K and Z resolutions agree quite closely with the data.

9.1.2 Calculation of Decay Distributions

On each iteration of the fit, KFIT propagates K^0 and \bar{K}^0 states from the target to the beginning of the decay region in each beam ($Z = 110$ m in the vacuum beam, and the downstream end of the Regenerator in the regenerated beam) using code which is very similar to KTEVMC. The modeling of kaon production, transmission, and regeneration will be described in this section. The propagation calculation is done for kaons with energies from 40.5 to 159.5 GeV, in 1-GeV steps, and involves




Figure 9.1: Binning used for the fit. The thin lines show the acceptance bins ($2 \text{ m} \times 2 \text{ GeV}$), while the thick lines show the twelve fit bins ($48 \text{ m} \times 10 \text{ GeV}$). The dots show the vacuum-beam $\pi^+\pi^-$ data for reference, with each dot representing about 325 events.

several parameters which may be fixed or floated in the fit, including Δm , τ_S , and the parameters describing the regeneration amplitude as a function of kaon energy. (For the standard $\text{Re}(\epsilon'/\epsilon)$ fit, Δm and τ_S are fixed while the regeneration parameters are floated, as described below.) Note that at high kaon energy, a significant K_S component reaches the decay region and affects the calculated decay distributions.

The energy dependence of kaon production is initially taken to be the “Malensek spectrum” (different for K^0 vs. \bar{K}^0) at the nominal proton-beam targeting angle of

4.8 mrad, modified by the same polynomial correction which is used in KTEVMC (see Section 8.1.1). The energy distribution is further modified to reflect the effect of scattering in the absorbers placed in the beamline, since high-energy kaons are more likely to stay within the beam when they scatter. This is done using a polynomial correction, different for the vacuum and regenerated beams (because of the movable absorber), which is derived from a Monte Carlo study. Even after these corrections, we do not claim to be able to predict the energy distribution exactly; for one thing, the $\pi^+\pi^-$ and $\pi^0\pi^0$ data were collected at different times, and the average targeting angle could have changed. 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.

One additional input to KFIT is the net transmission of the kaon beam through the movable absorber and Regenerator, *i.e.* the ratio of kaon fluxes in the regenerated and vacuum beams. We do not know the kaon-nucleon cross sections well enough to predict this accurately, so we must measure it from the data. Studies of $K_L \rightarrow \pi^+\pi^-\pi^0$ and $K_L \rightarrow 3\pi^0$ data indicate that the average K_L transmission (at 70 GeV) is about 7.70%. More importantly, the transmission is found to depend on energy, with a slope of -0.004% per GeV. KFIT uses these values to adjust the calculated kaon flux in the regenerated beam as a function of kaon energy.

The final important aspect of kaon propagation is the model for coherent regeneration. Regeneration arises from a difference in the forward scattering amplitudes for K^0 and \bar{K}^0 ; the relevant quantity is

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

where $f(0)$ and $\bar{f}(0)$ are the forward scattering amplitudes for K^0 and \bar{K}^0 , respectively, and k is the kaon wave number. Over the range of kaon energies used in this experiment, Regge theory predicts that the magnitude of f_- should vary with kaon energy according to a power law, which we express with respect to our typical kaon energy of 70 GeV:

$$|f_-(E_K)| = |f_-(70 \text{ GeV})| \left(\frac{E_K}{70 \text{ GeV}} \right)^\alpha. \quad (9.3)$$

Figure 9.2 shows a power-law fit to values of f_- measured from KTeV $\pi^+\pi^-$ data, showing that the power-law assumption is a good one.

The complex phase of f_- may be determined from its energy dependence via dispersion relations, with the requirement that the forward scattering amplitudes be analytic. For a power-law energy dependence, this yields a constant phase,




Figure 9.2: Power-law fit to $(f(0) - \bar{f}(0))/k$ measured from KTeV $\pi^+\pi^-$ data. The top plot shows the fit. In the bottom plot, the power-law term has been taken out so that the residuals are visible.

$$\arg(f_-) = -\frac{\pi}{2}(2 + \alpha) . \quad (9.4)$$

For plastic scintillator, which makes up the bulk of the Regenerator, we let the values of $|f_-(70 \text{ GeV})|$ and α float in the fit. For the lead pieces in the Regenerator (and the beryllium and lead absorbers), we use fixed values measured in earlier experiments. Propagation through these objects then involves a matrix transformation of the kaon quantum state which calculates regeneration exactly (*e.g.*, the object does not have to be “thin” relative to $c\tau_S$).

The end result of the calculations described above is the quantum state of a kaon (with K_S and K_L amplitudes a_S and a_L , respectively) at the beginning of the decay region in each beam, as a function of the initial kaon state (K^0 or \bar{K}^0) and energy. This determines the $\pi\pi$ decay rate as a function of z ,

$$r(z) \propto \left| a_S e^{-im_S\tau} e^{-\tau/2\tau_S} + \eta a_L e^{-im_L\tau} e^{-\tau/2\tau_L} \right|^2 \quad (9.5)$$

$$= |a_S|^2 e^{-\tau/\tau_S} + |\eta a_L|^2 e^{-\tau/\tau_L} \\ + 2 |a_S| |\eta a_L| \cos[\Delta m\tau + \arg(a_S/a_L) - \arg(\eta)] e^{-\tau(1/2\tau_S + 1/2\tau_L)} \quad (9.6)$$

where $\tau = (z - z_{beg})/\gamma\beta c$ is the proper time, measured from the beginning of the decay region. Note that a_S and a_L are independent of the decay mode, while η is different for charged and neutral decays:

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

$$\eta_{00} = \epsilon [1 - 2 \text{Re}(\epsilon'/\epsilon)] . \quad (9.8)$$

(These expressions, while not mathematically correct, are valid since ϵ and ϵ' have nearly identical complex phases.) Thus, KFIT varies $\text{Re}(\epsilon'/\epsilon)$ to try to simultaneously match the charged- and neutral-mode data.

9.1.3 Calculation and Minimization of the Fit χ^2

KFIT integrates the decay-rate expression (Equation 9.6) over z analytically, and sums over the 1-GeV E_K steps, to calculate the number of kaon decays in each

acceptance bin. (The bin containing the Regenerator edge is a special case: here, the integral starts a certain distance *within* the Regenerator, reflecting the fact that kaon decays within the last part of the Regenerator can escape. This distance is fixed at 2.2 mm for $\pi^+\pi^-$ decays and 5.8 mm for $\pi^0\pi^0$ decays.) It then multiplies by the acceptance for each bin, and adds together all the acceptance bins within a given fit bin to get the total number of events predicted to be in that bin.

The fit χ^2 is the sum over all fit bins of contributions of the form

$$\Delta\chi^2 = \frac{(N - P)^2}{\sigma_N^2 + \sigma_P^2} \quad (9.9)$$

where N is the number of data events observed and P is the prediction. KFIT calculates the raw kaon decay rate exactly, so σ_P comes only from the statistics of the Monte Carlo used to determine the acceptance.² Since we generate more MC events than there are data events, $\sigma_P < \sqrt{P}$. Because P is subject to smaller fluctuations than N , we take σ_N to be equal to \sqrt{P} rather than \sqrt{N} . (This matters only for fits using small fit bins, which may have low statistics in some bins.)

The χ^2 is minimized using the MINUIT package [77], which is a general-purpose package for minimizing a function of several parameters. The standard $\text{Re}(\epsilon'/\epsilon)$ fit takes about five minutes to run.

One final important aspect of this analysis is that the fit for $\text{Re}(\epsilon'/\epsilon)$ was done “blind”, by hiding the true value until after all data analysis and systematic studies were finalized. This was done within KFIT itself, which used the true value of $\text{Re}(\epsilon'/\epsilon)$ to calculate the decay distributions but reported the value of a fit parameter which was equal to the true value plus an unknown offset, chosen from a roughly Gaussian distribution with a mean of zero and a sigma of 60×10^{-4} .³ The offset was the same every time KFIT was run, so it was straightforward to do systematic

²The error on the acceptance in a given acceptance bin is calculated as the binomial error for the number of MC events generated and reconstructed in that bin.

³KFIT actually multiplied the true value by $(-1)^n$ before adding the offset, where n could have been either 0 or 1, and was initially unknown. However, this turned out to be a hindrance when doing systematic studies since we wanted to be able to check whether the value of $\text{Re}(\epsilon'/\epsilon)$ moved in the expected direction when something was changed. So we “peeked” at the value of n early on and established that it was equal to 1, and took this into account in all later studies.

studies and observe the *change* in $\text{Re}(\epsilon'/\epsilon)$. Although we believed that we would be able to choose analysis cuts, the fitting procedure, etc., without being unduly influenced by any prejudice about the “expected” value of $\text{Re}(\epsilon'/\epsilon)$, it was comforting to not have to worry about it at all.

9.1.4 Other Fit Types

The previous sections described KFIT as it is used for the standard $\text{Re}(\epsilon'/\epsilon)$ fit. In fact, it can be used in a variety of ways. By using small fit bins in the regenerated beam (say, 2 m wide) we can fit for the parameters Δm and τ_S . We can also relax the assumption of *CPT* invariance and fit for the parameters ϕ_{+-} (the phase of η_{+-}) and $\Delta\phi \equiv \phi_{00} - \phi_{+-}$ in order to *test* whether *CPT* is violated. We have, in fact, made preliminary measurements of these quantities using KTeV data, but they will not be presented in this thesis.

KFIT has a number of options and alternative parametrizations which may be used in various parts of the calculation. For example, there are variations on the regeneration model, such as nuclear screening corrections or a relaxation of the analyticity assumption which determines the phase of the regeneration amplitude. These may be used to do robustness checks and systematic studies; a few such variations will be discussed later.

9.2 The Fit for $\text{Re}(\epsilon'/\epsilon)$

As mentioned earlier, there are 12 fit bins for each of the four $\pi\pi$ samples, for a total of 48 fit bins. The kaon flux normalizations for $\pi^+\pi^-$ and $\pi^0\pi^0$ decays are independently floated in each 10-GeV energy interval, making 24 fit parameters. Regeneration is parametrized by $|f_-(70 \text{ GeV})|$ and α , which are allowed to float. The final floating parameter is $\text{Re}(\epsilon'/\epsilon)$. There are 27 floating parameters in all, so the fit has 21 degrees of freedom.

The physics parameter Δm and τ_S are fixed to the values from a fit to all previous experiments, done by the Particle Data Group (PDG) and published in the 1998 Review of Particle Properties [32], which are $0.5301 \times 10^{10} \text{ } \hbar \text{ s}^{-1}$ and $0.8934 \times 10^{-10} \text{ s}$,

respectively. The value of ϵ is fixed to the mean of the PDG values for $|\eta_{+-}|$ and $|\eta_{00}|$, 2.28×10^{-3} . *CPT* symmetry is assumed to be valid, so that the complex phases of ϕ_{+-} and ϕ_{00} are both equal to the so-called “superweak phase”,

$$\begin{aligned}\phi_{SW} &\equiv \arctan\left(\frac{2\Delta m}{(1/\tau_S - 1/\tau_L)}\right) \\ &= 43.5^\circ\end{aligned}\tag{9.10}$$

By mid-February 1999, we had frozen the analysis cuts and Monte Carlo code and had completed the evaluation of the systematic errors. Thus we removed the offset from KFIT and found that the true value of $\text{Re}(\epsilon'/\epsilon)$ from the fit was $(28.0 \pm 3.2) \times 10^{-4}$. We were quite startled to find such a large value, nearly twice the average of E731 and NA31! But we were, and are, confident that the analysis and the fit are valid, based in part on cross-checks to be described in Section 9.3.

The regeneration parameters found by the fit are $|f_-(70 \text{ GeV})| = (1.2984 \pm 0.0006)$ mbarn and $\alpha = -0.5903 \pm 0.0015$. The kaon flux normalization parameters at the twelve different energies vary by about 1% for charged mode and 2% for neutral mode. The fit χ^2 is 30 for 21 degrees of freedom.

The statistical error on $\text{Re}(\epsilon'/\epsilon)$ given above is due to both the data statistics and the error on the acceptance. An alternative fit, with σ_P forced to be zero in Equation 9.9, indicates that the error from the data statistics alone is 3.0×10^{-4} . (The central value changes by less than 0.05×10^{-4} .) The remaining error, 1.0×10^{-4} , is the statistical error on the acceptance due to the finite number of Monte Carlo events generated.

Note that the statistical error is 14% larger than one would naively calculate from the double ratio using \sqrt{N} statistics. This is because the K_L component in the regenerated beam increases the uncertainty on the K_S decay amplitudes, and because the charged and neutral data samples have somewhat different energy distributions.

9.3 Fitting Cross Checks

We have performed a number of checks of the fitting software. In one study, we used KTEVMC to generate a dataset with $\text{Re}(\epsilon'/\epsilon) = 10 \times 10^{-4}$, then fit it (using an independent MC sample for the acceptance calculation). The fitted value of $\text{Re}(\epsilon'/\epsilon)$ was consistent with the value used in generation within the statistical precision of the study, $\sim 2 \times 10^{-4}$. We designed another study to test the internal consistency of KFIT with high statistical precision: we used a simplified acceptance model to calculate an exact parent distribution, then applied Poisson statistical fluctuations to create 100 independent “fake datasets”. We fit these datasets, using the same acceptance model, and found that the mean $\text{Re}(\epsilon'/\epsilon)$ value from the 100 fits differed from the generated value by $(0.11 \pm 0.28) \times 10^{-4}$. We also modified the fit in various ways (size of acceptance bins, screening corrections for regeneration in lead, etc.) with negligible effect on the fitted value of $\text{Re}(\epsilon'/\epsilon)$.

Many possible problems in the detector or the analysis—for example, significant nonlinearities in the photon energy reconstruction, or momentum-dependent tracking inefficiencies in the spectrometer—would tend to cause an energy-dependent bias. However, we find that fits to individual 10-GeV bins of kaon energy, shown in Figure 9.3, give consistent values of $\text{Re}(\epsilon'/\epsilon)$.

We did an alternative fit in which we removed the power-law constraint on the regeneration amplitude, letting the amplitude float in each 10-GeV momentum bin. (*Within* each 10-GeV bin, the regeneration amplitude is still assumed to follow a power law with the α parameter fixed to the value from the standard fit, -0.5903 .) This fit found $\text{Re}(\epsilon'/\epsilon) = 27.5 \times 10^{-4}$, with a χ^2 of 11.3 for 11 degrees of freedom. Interestingly, the statistical error on $\text{Re}(\epsilon'/\epsilon)$ did not increase.

M. Pang has developed an entirely different fitting technique to extract $\text{Re}(\epsilon'/\epsilon)$, along with the other kaon physics parameters, from the KTeV data. This technique starts from the principle that the acceptance for a given value of E_K and Z is the same for the vacuum and regenerated beams. The data is kept in small bins, 4 GeV by 0.5 m (0.1 m near the Regenerator edge), and fits are done to the *ratio* of the




Figure 9.3: Fitted values of $\text{Re}(\epsilon'/\epsilon)$ in 10-GeV bins of kaon energy.

regenerated- and vacuum-beam Z distributions. The acceptance cancels in the ratio, so no Monte Carlo simulation is needed. One such fit is shown in Figure 9.4.

This technique is less statistically powerful than the standard fitting technique. It has the great advantage of not being sensitive to the detector acceptance, although it is more sensitive to other systematic issues such as the non-Gaussian tails on the Z vertex resolution. It also is vulnerable to any slight difference in the acceptance between the vacuum and regenerated beams, which can occur from accidental or rate-dependent effects, and must assign a systematic error to cover such a possibility. The result of the fit (which also was done “blind” until everything was finalized) is

$$\text{Re}(\epsilon'/\epsilon) = \left[20.5 \pm 4.9 \text{ (stat)} \begin{matrix} +3.6 \\ -3.3 \end{matrix} \text{ (syst)} \right] \times 10^{-4} .$$

Thus the central value differs from the standard fit by -7.5×10^{-4} . The correlation




Figure 9.4: Sample fit from alternative technique to fit for $\text{Re}(\epsilon'/\epsilon)$. The dots are the data, and the line is the fit.

between the statistical errors for the two techniques was evaluated using the 100 “fake datasets” mentioned earlier in this section, allowing us to determine that the expected statistical error on the *difference* is 3.5×10^{-4} . The uncorrelated parts of the systematic errors contribute an additional error of 3.8×10^{-4} on the difference, for a total error of 5.2×10^{-4} . Thus, while the difference between the two techniques is almost 1.5σ , we consider this alternative technique to have provided an affirmative check of the standard fit.⁴

We began this chapter by explaining why a grand ratio of the event totals and acceptances was not a reliable way to extract $\text{Re}(\epsilon'/\epsilon)$. However, our Monte Carlo distributions match the data very well, and if we restrict the Z range used in the regenerated beam—say, to $124 < Z < 132$ m—then the K_L component becomes much less important, and the value we arrive at should not be too far off. Table 9.1

⁴One might wonder why we do not take the weighted average of the two fit results. The main reason is that we decided in advance, while the $\text{Re}(\epsilon'/\epsilon)$ values were still hidden, that we would report the result of the standard technique and use the alternative technique as a check, and to change our plan now would violate the principle of the blind analysis. In any case, because the alternative technique has a substantially larger uncorrelated error than the standard technique, the weighted average is only $\sim 1 \times 10^{-4}$ lower than the result of the standard technique alone.

shows the background-subtracted event totals and average acceptances which go into this calculation, which yields a value for $\text{Re}(\epsilon'/\epsilon)$ of $(26.5 \pm 3.0) \times 10^{-4}$. It is nice to know that we can see direct CP violation so clearly without having to do any fit at all.

9.4 Summary

We have succeeded in extracting $\text{Re}(\epsilon'/\epsilon)$ from the data, with alternative techniques giving consistent results. The value is larger than was found in either E731 or NA31, and the difference from zero appears to be highly significant, at least judging from the statistical error. A final conclusion must wait until we have evaluated the systematic error, which is the subject of the next chapter.

Table 9.1: Event totals and average acceptances for “grand ratio” calculation. Note that the values of the individual acceptance-corrected Vac/Reg ratios are not particularly meaningful, since they depend on the Z intervals used; they just happen to be close to one for this calculation. Their ratio, however, is directly related to $\text{Re}(\epsilon'/\epsilon)$.

	Vac $\pi^+\pi^-$ (110–158 m)	Reg $\pi^+\pi^-$ (124–132 m)	Vac $\pi^0\pi^0$ (110–158 m)	Reg $\pi^0\pi^0$ (124–132 m)
Data	2,607,274	3,477,614	862,254	1,048,347
MC generated	37,314,795	36,386,694	83,271,909	81,284,332
MC accepted	8,629,457	11,582,403	8,625,959	10,399,164
Acceptance	0.231261(69)	0.318314(77)	0.103588(33)	0.127936(37)
Acceptance-corrected Vac/Reg ratio	1.0320(09)		1.0158(15)	
Grand ratio $\text{Re}(\epsilon'/\epsilon)$	1.0159(18) $(26.5 \pm 3.0) \times 10^{-4}$			

CHAPTER 10

SYSTEMATIC ERRORS

The struggle to control systematic errors is at the heart of the effort to measure $\text{Re}(\epsilon'/\epsilon)$, and the focus is always on things which may affect K_L and K_S decays differently. It begins with the conceptual design of the experiment: two beams to allow the simultaneous collection of K_L and K_S decays, a high-precision detector with a well-defined geometrical acceptance, and high-efficiency triggers. The analysis software is tuned to reconstruct events accurately and robustly, and cuts are chosen with care to affect the K_L and K_S samples similarly. However, there are unavoidable differences from the different decay distributions, and so a great deal of work goes into making sure that the Monte Carlo simulation accurately reproduces the detector response and acceptance.

Even after all of this effort, we must carefully consider a wide variety of possible remaining biases and constrain them, generally based on studies of the data. The magnitude of the uncertainty is often dictated by the amount of suitable data available to be studied, though in some cases we have simply decided against spending the additional effort required to try to reduce the uncertainty further. We would like, at least, to achieve a total systematic error smaller than the statistical error on $\text{Re}(\epsilon'/\epsilon)$, 3.0×10^{-4} .

This chapter describes our studies of possible sources of systematic errors on our measurement of $\text{Re}(\epsilon'/\epsilon)$. Dozens of effects have been considered, and uncertainties assigned in many cases. The studies have been divided into seven categories which will be covered by the sections below. We will also describe a few additional cross-checks, which do not contribute to the total systematic error but help to give us confidence that we are not missing some significant effect. Finally, we will summarize this chapter with a table listing all contributions to the total systematic error. Note

that all of these studies were completed *before* we knew the true value of $\text{Re}(\epsilon'/\epsilon)$ extracted from the data by the fitting program.

10.1 Trigger Inefficiencies

Besides the normal $K \rightarrow \pi\pi$ triggers, we collected prescaled samples of events with looser trigger requirements. These allowed us to study the small inefficiencies in the normal triggers, both in the hardware and in the Level 3 filtering, and to place limits on the possible biases on $\text{Re}(\epsilon'/\epsilon)$.

10.1.1 Charged-Mode Trigger

We studied the charged-mode Level 1 trigger inefficiency using reconstructed K_{e3} events collected using a special trigger which simply required a certain amount of energy in the calorimeter. The overall inefficiency is measured to be 0.47%, mostly from the cracks in the trigger hodoscope banks, which are well simulated in KTEVMC. Early accidental activity in the trigger hodoscope and in the drift chambers (causing the “DC-OR” electronics to be inefficient) also contribute but are consistent with affecting the vacuum and regenerated beams equally. Based on a study with K_{e3} data, we assign a systematic error on $\text{Re}(\epsilon'/\epsilon)$ of 0.17×10^{-4} from these processes.

A small fraction of the events, 0.04%, fail the Level 1 trigger for unknown reasons; the inefficiency is consistent with being the same for the vacuum and regenerated beams, but to be conservative, we assign the statistical error on this check as a systematic error on $\text{Re}(\epsilon'/\epsilon)$, 0.28×10^{-4} .

We studied the Level 2 trigger inefficiency using a different special trigger with the normal Level 1 requirements but without any Level 2 requirements. The overall Level 2 inefficiency is 0.08%, mostly due to an occasional glitch in the signal path between the hit-counting system and the trigger control hardware, which does not cause a bias because it is not associated with any activity in the detector. The Banana/Kumquat modules had to miss *two* hits in order for an event to fail the hit-count requirement, and this happened very rarely. However, events with a particular

topology (both tracks either above or below the center of drift chamber 4) could fail the YTF requirement if only *one* hit were missed by the Banana/Kumquat system, and this topology is somewhat more common for vacuum-beam decays than for regenerated-beam decays. The data allowed us to limit the bias from this process to 0.07×10^{-4} .

We studied the inefficiency of the Level 3 filter process using events which were collected with the normal hardware trigger but which were not subjected to online filtering. The overall inefficiency is about 0.08% and is consistent with being the same for vacuum- and regenerated-beam events. A similar inefficiency is measured when Level 3 filtering is applied to Monte Carlo events, so we may actually be able to simulate the failure probability correctly. However, to be conservative we combine the uncertainties from data and Monte Carlo in quadrature, and assign a systematic error on $\text{Re}(\epsilon'/\epsilon)$ of 0.38×10^{-4} .

10.1.2 Neutral-Mode Trigger

We studied the inefficiency of the E-Total system using K_{e3} and $\pi^+\pi^-\pi^0$ decays collected with the normal charged-mode trigger. The $\pi^+\pi^-\pi^0$ decays were particularly valuable since they left four clusters of energy in the calorimeter, like $\pi^0\pi^0$ decays. Both samples indicated an overall inefficiency of 0.6%, which is almost entirely due to early accidental energy which keeps the analog total energy above the discriminator threshold, preventing the E-Total source from retriggering; this cannot bias $\text{Re}(\epsilon'/\epsilon)$ since the failure to retrigger applies to the whole calorimeter, not just the blocks with the early energy. The *intrinsic* inefficiency is measured to be 0.04%, and we assign 1/6 of this, or 0.07×10^{-4} , as a systematic error on $\text{Re}(\epsilon'/\epsilon)$.

The HCC only counts clusters for which one of the blocks has an energy above its HCC-bit threshold, which is nominally 1 GeV. However, we allow events in the final $\pi^0\pi^0$ sample to have photons with energies as low as 3 GeV, and a 3-GeV photon which lands near the corner of four CsI blocks will deposit less than 1 GeV in each. Therefore, we expect the HCC requirement to have some inefficiency for otherwise good events. For the data used in this analysis, the average inefficiency

was about 1%.¹ It was slightly different for vacuum- and regenerated-beam decays, and would have biased $\text{Re}(\epsilon'/\epsilon)$ by 2×10^{-4} if this fact had been ignored. But KTEVMC simulates the HCC-bit thresholds, and by comparing the MC and data HCC cluster-count distributions for $K_L \rightarrow 3\pi^0$ decays (collected with a special trigger with no HCC requirement), we limit the possible error to 0.2×10^{-4} .

We studied the inefficiency of the Level 3 filter process using events which were collected with the normal hardware trigger but which were not subjected to online filtering. The inefficiency was rather large (0.66%) at the very beginning of physics data collection, when the online pairing χ^2 cuts were too tight, but only a few percent of the data was collected under these conditions. Averaged over the whole 1996 running period, the inefficiency was 0.04%, and was consistent with being the same for the vacuum and regenerated beams. The statistics of this event sample allow us to limit the possible bias on $\text{Re}(\epsilon'/\epsilon)$ to 0.20×10^{-4} .

10.2 Event Reconstruction Biases

Accurate event reconstruction relies on accurate knowledge of the positions of the detector elements (CsI crystals and drift chamber wires) and calibration of their response. Although we have made an effort to determine these things as well as we can, in this section we consider the possibility that there are systematic shifts. These can affect the measured value of $\text{Re}(\epsilon'/\epsilon)$ by shifting events past the analysis cuts, causing the final $\pi\pi$ samples to either gain or lose events.

10.2.1 Calorimeter Cluster Energy Measurement

Section 7.5 described how we adjust the energy scale of the calorimeter using $\pi^0\pi^0$ events at the downstream edge of the Regenerator. This step is important because a systematic shift in measured photon energies can shift the reconstructed Z vertex positions for $K \rightarrow \pi^0\pi^0$ decays. For example, if the calorimeter overestimates

¹The inefficiency was 2.4% for the first quarter of the data collected during the 1996 running period, when the actual thresholds varied significantly from channel to channel. After the CsI PMT gains were adjusted, the thresholds all were close to 1 GeV and the HCC inefficiency fell to 0.6%.

the photon energies, then the decay vertex will be reconstructed upstream of its actual position. This may cause a decay at $Z = 158.01$ m, say, to reconstruct at $Z = 157.99$ m and thus to be included in the final event sample when it should not have been.² Because there are far more events at the downstream end of the fiducial region in the vacuum beam than in the regenerated beam, this can lead to a significant bias on $\text{Re}(\epsilon'/\epsilon)$. For example, the original edge mismatch of 7.6 cm in Figure 7.8a would bias $\text{Re}(\epsilon'/\epsilon)$ by about 1.2×10^{-4} if not corrected.

The energy scale correction measured with the $\pi^0\pi^0$ decays is statistically very precise, but we can check for a systematic problem by determining the correction in various other ways. A similar method, using $3\pi^0$ decays at the Regenerator edge, yields a similar correction, as does a study of the reconstructed position for Regenerator-junk events produced in the lead pieces in the Regenerator. This agreement is encouraging, but we also want to check that energy scale correction is valid for events throughout the entire decay region, which tend to illuminate the calorimeter differently. The vacuum window at the downstream end of the decay volume provides a serendipitous means to check this, using pairs of π^0 's from “vacuum-window junk” events (hadronic interactions which produce multiple particles). Using the nominal energy scale correction, we find that vacuum-window junk π^0 pairs reconstruct downstream of the true position (determined using charged particles) by about 2 cm on average, suggesting a somewhat smaller correction to the energy scale. An additional check, using a special data run with a tungsten “Z calibrator” inserted at a known position near the downstream end of the decay volume, yields a consistent but statistically poorer result. The energy scale corrections indicated by each of these five methods are shown graphically in Figure 10.1. We continue to use the $K \rightarrow \pi^0\pi^0$ value for our nominal energy scale correction, but assign an uncertainty equal to the difference between this and the vacuum-window-junk value. This corresponds to an uncertainty of 0.7×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$.

We have studied residual nonlinearities in the calorimeter energy measurement

²Of course, the finite resolution of the calorimeter guarantees that some events will be reconstructed on the opposite side of the cut, but there will be a comparable number of events moving in each direction since the Z distribution is fairly flat. Uncertainties related to detector resolutions will be considered in Section 10.5.3.




Figure 10.1: Comparison of different determinations of calorimeter energy scale correction. The solid line shows the correction used, 0.99875, which is the value given by $\pi^0\pi^0$ Regenerator-edge matching. The dashed lines show the assigned systematic uncertainty of 0.00073.

by examining the mean $\pi^0\pi^0$ invariant mass as a function of kaon energy. It is observed to drop by about $0.3 \text{ MeV}/c^2$ between 40 GeV and 160 GeV; this may be explained by a nonlinearity with a slope of -5×10^{-5} per GeV along with a small shift of the intercept. Among the cases considered, the largest change in $\text{Re}(\epsilon'/\epsilon)$ was 0.6×10^{-4} , so we assign this as a systematic error.

10.2.2 Calorimeter Cluster Position Measurement

The overall size of the CsI array has been measured in four different ways:

1. The sizes of the individual CsI blocks, measured before stacking, were summed in rows and columns.
2. The total horizontal and vertical dimensions were optically surveyed after stacking.
3. For a large sample of electrons from K_{e3} decays, the reconstructed cluster position was compared to the projected track position (at the mean shower depth in the CsI, which depends logarithmically on the electron energy).

4. Muons, from special runs with the beam-stop closed and the analysis magnet turned off, were projected to the calorimeter. The positions of the *edges* between crystals were determined by finding muons, traveling at an angle, which deposited energy in both crystals. The point at which the muon passed from one crystal to the other was determined from the relative amount of energy in the two crystals.

The nominal dimensions were taken from the optical survey. The other methods yielded dimensions within ± 1 mm, so this was taken as the uncertainty on the overall size. An error in the overall size would affect the reconstructed Z vertex (via Equation 7.2), but our procedure of adjusting the calorimeter energy scale tends to reduce the resulting bias. When this is taken into account, an uncertainty of 1 mm on the overall size of the CsI array corresponds to an uncertainty on $\text{Re}(\epsilon'/\epsilon)$ of 0.15×10^{-4} .

When comparing cluster and track positions for electrons, a systematic difference was observed as a function of the position of the electron within the seed block, indicating an inconsistency in the cluster position calculation. The maximum size of the discrepancy was about 0.25 mm in small blocks and 0.7 mm in large blocks. This effect is probably related to the incident angles of the electrons, but to be conservative we assume that the effect is the same for photons, even though they generally have smaller incident angles. A Monte Carlo study indicates that this effect would shift the measured value of $\text{Re}(\epsilon'/\epsilon)$ by 0.35×10^{-4} , so we assign this as a systematic error.

10.2.3 *Spectrometer Alignment and Calibration*

We have considered a variety of possible errors arising from the alignment and calibration of the drift chambers, including errors in the surveyed Z positions, the drift-time-to-distance calibration, and the determination of the chamber positions and rotations. These are found to have very little effect on $\text{Re}(\epsilon'/\epsilon)$, as expected, and a conservative error of 0.21×10^{-4} is assigned.

The track-finding software assumes that the drift chamber wires are spaced at

intervals of exactly 6.35 mm. In fact, the spacing does vary, sometimes by 0.1 mm or more. This could change the effect of the track separation cut (described in Section 6.4.6), which already cuts more vacuum-beam events than regenerated-beam events. A study of this effect, using a large sample of K_{e3} events to determine the actual wire positions, indicates that the bias on $\text{Re}(\epsilon'/\epsilon)$ is no more than 0.17×10^{-4} .

A mistake in calculating the kaon energy for $\pi^+\pi^-$ events can bias $\text{Re}(\epsilon'/\epsilon)$ by shifting events past the cut on E_K , and by affecting the calculation of the regeneration amplitude when performing the fit. The calculated kaon energy is essentially proportional to the field of the analysis magnet, which is calibrated to ensure that the $\pi^+\pi^-$ invariant mass distribution lines up with the K^0 mass. However, the kaon mass has a fractional uncertainty of 6×10^{-5} ; this is found to lead to an uncertainty of 0.12×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$.

10.3 Analysis Cuts

We have varied or removed many of the analysis cuts, one at a time, and observed the resulting change in the measured value of $\text{Re}(\epsilon'/\epsilon)$ in each case. When the number of events passing all cuts does not change too much (say, of order 1%), the high correlation between the normal and altered cases allows us to detect statistically significant *changes* of less than 10^{-4} in the value of $\text{Re}(\epsilon'/\epsilon)$. In most cases, there is no significant change. However, there *is* a significant change for a few important variables, listed in Table 10.1. In each case, we assign the observed change in $\text{Re}(\epsilon'/\epsilon)$ as a systematic error. Note that the first two variations cause the background levels to roughly double.

Table 10.1: Cut variations which lead to significant changes in $\text{Re}(\epsilon'/\epsilon)$.

Mode	Variable	Normal cut	Altered cut	$ \Delta\text{Re}(\epsilon'/\epsilon) $
Charged	p_T^2	250 MeV ² /c ²	500 MeV ² /c ²	0.57
Neutral	Shape χ^2	48	no cut	0.50
Neutral	Ring number	110	140	0.60

10.4 Backgrounds

The backgrounds to the $\pi^0\pi^0$ samples are quite large; if they were ignored altogether, $\text{Re}(\epsilon'/\epsilon)$ would change by 8×10^{-4} . The accuracy of the background determination is mainly limited by our understanding of kaon scattering in the Regenerator, specifically the shape of the ring-number distribution for these scattered events, since we must normalize this background at large ring-number values and extrapolate under the signal peak. Section 6.5.1 described how we use $\pi^+\pi^-$ data (from the 1996 running period, *i.e.* collected at the same time as the $\pi^0\pi^0$ data) to characterize the distribution of scattered events, then model it in KTEVMC.

We have considered several sources of uncertainty in the scattering model. The most significant is the uncertainty in the acceptance for $K \rightarrow \pi^+\pi^-$ decays, in particular as a function of p_T^2 . Because of the large Level 3 filter inefficiency in the 1996 data, we felt it necessary to use data to check the acceptance rather than relying on MC. We did this by studying the p_T^2 distribution for $\pi^+\pi^-$ pairs from $\pi^+\pi^-\pi^0$ decays, although in the process we found that we had to remeasure the $\pi^+\pi^-\pi^0$ Dalitz parameters, which had last been measured in 1974. In the end, we were able to place a limit on the acceptance uncertainty corresponding to an error of 0.50×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$.

The other major uncertainty concerns the amount of collimator scattering, which affects the $\pi^0\pi^0$ background estimate in two ways: first, as a background in the $\pi^+\pi^-$ sample which must be subtracted before we can determine the Regenerator scattering distribution; and second, as a background in the $\pi^0\pi^0$ sample which cannot be normalized using a sideband, so that we must rely on KTEVMC to predict its magnitude correctly. In reality, an uncertainty in the amount of collimator scattering should tend to cancel between these two mechanisms, but to be conservative we consider them independently and arrive at a systematic error of 0.45×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$.

Variations in the procedure for fitting the Regenerator scattering distribution, and possible dependence on time and kaon energy, lead us to assign an additional error of 0.39×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$. Finally, small uncertainties in the $3\pi^0$ and

Regenerator-junk backgrounds, as well as a few other backgrounds not included in the standard background-subtraction procedure ($K_L \rightarrow \pi^0\gamma\gamma$, and $\Xi^0 \rightarrow \Lambda\pi^0$ with $\Lambda \rightarrow n\pi^0$), contribute an error of 0.17×10^{-4} .

The charged-mode backgrounds are small to begin with and are relatively easy to determine from sidebands. Based on the agreement between the background prediction and the data in Figure 6.12, we assign a conservative error of 10% (relative) on the background level for each beam, corresponding to an error of 0.20×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$.

10.5 Detector Acceptance

Accurate determination of the detector acceptance is the most critical ingredient in the KTeV technique for measuring $\text{Re}(\epsilon'/\epsilon)$. In Chapter 8 we described the Monte Carlo simulation which is used to calculate the acceptance. While the Monte Carlo has been designed to simulate the various parts of the detector in sufficient detail, it can do so only if a multitude of input parameters have the correct values. These include detector positions and sizes, resolutions, inefficiencies, etc. We will consider these specific issues in this section, and will also use a comparison of the Z distributions for data and MC events as a global check of the acceptance.

Even if the Monte Carlo were perfect, there would still be an uncertainty on the detector acceptance due to the fact that it is calculated from a finite number of simulated events. For this analysis, we chose to generate about 3.3 times the data statistics in charged mode and 10 times the data in neutral mode. These event totals contribute uncertainties of 0.52×10^{-4} and 0.90×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$, respectively.³

10.5.1 Limiting Apertures

The three limiting apertures for photons from $K \rightarrow \pi^0\pi^0$ decays are the Mask Anti, the Collar Anti, and the outer ring of CsI crystals in the calorimeter. The effect

³There are different ways to calculate these values which seem to give somewhat different results, but in any case their quadrature sum is consistent with the difference between the standard $\text{Re}(\epsilon'/\epsilon)$ fit and an alternative fit with σ_P forced to be zero in Equation 9.9.

of the outer ring of CsI crystals was included in the study of the overall size of the CsI array (Section 10.2.2) and will not be discussed further. The MA affects only vacuum-beam events, while the CA affects a greater fraction of vacuum-beam events than regenerated-beam events, so an incorrect size for either of these would bias $\text{Re}(\epsilon'/\epsilon)$.

Although the actual physical sizes of the MA and CA are known from their design and construction, it is important to determine their *effective* sizes for purposes of defining the acceptance for photons. These may differ from the actual sizes due to slight misalignment of individual lead (or tungsten) and scintillator pieces, or simply due to the fact that a photon traveling at an angle may encounter only a fraction of the total thickness of the detector.⁴ Therefore, we use data to determine the effective sizes *in situ*. Note that the transverse positions are much less critical due to the symmetry of the detectors.

The positions of the edges of the CA are determined using electrons from K_{e3} decays which point near the CA but do not deposit any energy (or deposit less than the analysis cut made for $\pi^0\pi^0$ events). The effective size of the CA is found to be consistent with its physical size, to within an uncertainty of about 0.1 mm for the average of X and Y for the east and west detectors. This is used to limit the possible bias on $\text{Re}(\epsilon'/\epsilon)$ to 0.45×10^{-4} .

The positions of the edges of the MA have been determined using electrons. However, a consistent and more accurate determination comes from studying the distribution of $K_L \rightarrow 3\pi^0$ decays upstream of the MA, which indicates the effective sizes of the MA beam holes in a more direct manner. Using this method, we limit the possible bias on $\text{Re}(\epsilon'/\epsilon)$ to 0.18×10^{-4} .

The $\pi^+\pi^-$ acceptance does not depend on the positions of these physical aper-

⁴A photon conversion in the Mask Anti is sufficient to “veto” a $\pi^0\pi^0$ event, even if there is no energy deposit in the scintillator, because there are no longer four photons to be reconstructed. On the other hand, the CA vetoes an event only if the photon makes an electromagnetic shower which deposits significant energy in the CA. It is possible, for instance, for a photon traveling at an angle to clip the corner of the CA and convert, with the e^+e^- pair escaping out the side of the CA and entering the CsI. This would not be detected by the CA and would make a perfectly good cluster in the CsI.

tures since we cut away from them based on projected track positions, as explained in Section 6.4.6.

10.5.2 *Effective Regenerator Edge*

The decay region in the regenerated beam effectively begins a bit *before* the downstream edge of the Regenerator, and is different for neutral vs. charged decays. For neutral decays, it is simply a matter of the probability of all four photons passing through the remaining material without converting. It is easy to show that the *average* amount of material that four photons can pass through is 1/4 of the characteristic conversion length (which is, itself, 9/7 of the radiation length). The final 4-mm-thick piece of plastic scintillator is only 1% of a radiation length, so we will neglect it and assume that the preceding lead piece causes all of the photon conversions. The radiation length of lead is 5.6 mm, so the effective Regenerator edge for neutral decays is

$$(4 \text{ mm}) + \frac{1}{4} \times \frac{9}{7} \times (5.6 \text{ mm}) = 5.8 \text{ mm}$$

upstream of the end of the Regenerator. Note that there is no significant uncertainty in this calculation.

The situation is different for charged decays. A $\pi^+\pi^-$ pair will usually exit the Regenerator without being destroyed, but the event may be rejected by the analysis cut on the measured energy deposit in the scintillator of the last Regenerator module. A study of the response of this module to minimum-ionizing particles under high-rate running conditions indicates the the $\pi^+\pi^-$ pair can pass through the last 2.2 ± 0.6 mm of scintillator, on average, without being vetoed. This leads to an uncertainty of 0.26×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$. Note that if we had simply assumed that both charged and neutral decays start at the downstream end of the Regenerator, $\text{Re}(\epsilon'/\epsilon)$ would have been biased by $+1.4 \times 10^{-4}$.

10.5.3 *Detector Resolutions*

The spectrometer resolution affects the widths of kinematic distributions such as the invariant mass and p_T^2 . It also can cause a net “migration” of events past the

analysis cuts on E_K and Z , changing the number of events in the final sample.⁵ We tuned the KTEVMC parameters for the drift chamber resolutions, and the amount of multiple scattering in the material of the spectrometer, to match the data distributions for several variables; one of the most useful was the mean square separation between upstream and downstream track segments at the magnet, whose momentum dependence gives information about both the chamber resolutions and the amount of multiple scattering. Based on these comparisons, we assigned an uncertainty of 10% on the resolution parameters. A Monte Carlo study, which generated the same set of events with varying resolutions, indicated that this corresponds to an uncertainty of 0.35×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$.

One aspect of the spectrometer which is *not* simulated in KTEVMC is “vertical focusing” by the analysis magnet, which bends tracks in the Y direction and can, at some level, affect whether tracks pass the aperture cuts at the calorimeter. This effect was studied and found to have a small effect for low-energy kaons, but no detectable effect for the kaon energies above 40 GeV which are used to measure $\text{Re}(\epsilon'/\epsilon)$.

The energy resolution of the calorimeter was checked by examining the reconstructed Z resolution for Regenerator-junk $\pi^0\pi^0$ events, which are produced in the lead pieces of the Regenerator. In this way, we verified that the resolution in KTEVMC is correct to within 0.2% in quadrature. A resolution change of this size was found to have virtually no effect on the measured value of $\text{Re}(\epsilon'/\epsilon)$.

10.5.4 Drift Chamber Modeling

Section 8.3 described the simulation of the “high-sod” effect, as well as missing hits, in the drift chambers. We found that the simulation matched the data best, *e.g.* the tails on the p_T^2 distribution, when we scaled up the nominal high-sod and missing-hit probabilities by a factor of 1.25. We assign an uncertainty of 0.25 on this

⁵One can show that, for a variable on which a cut is made, the net migration past the cut due to the reconstruction resolution σ is $s\sigma^2/2$, where s is the local slope of the distribution at the position of the cut.

scale factor; a Monte Carlo study indicates that this changes the measured value of $\text{Re}(\epsilon'/\epsilon)$ by 0.62×10^{-4} , so we assign this as a systematic error.

10.5.5 *Accidental Activity*

As described in Section 8.6, KTEVMC overlays an event from the accidental trigger on top of each simulated decay. This procedure should properly reproduce the effect of the underlying activity in the detector. For example, extra hits in the drift chamber system (often in conjunction with a high sod or a missing hit) can cause a track to be misreconstructed or even lost. From a Monte Carlo study, the overall effect of accidental hits on $\text{Re}(\epsilon'/\epsilon)$ is about 1.1×10^{-4} . A comparison of the number of early hits in the drift chambers indicates that KTEVMC agrees with the data at the 10% level, so we assign a systematic error of 0.11×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$.

The most readily apparent effect of accidental activity in the calorimeter is to add extra “software” clusters from particles which arrive early or late. The distribution of the number of software clusters agrees well between the data and MC generated with accidental overlays, giving us confidence that the accidental overlay procedure is doing the proper thing. A less apparent effect of accidental activity is to sometimes add extra energy to one or more of the four photon clusters. This induces a shift in the reconstructed kaon energy and Z vertex, which is corrected for, on average, by our procedure of adjusting the calorimeter energy scale to match the Z distribution at the Regenerator edge. A Monte Carlo study indicates that any residual discrepancy would induce a mismatch in the Z distributions, which is covered in the next section.

10.5.6 *Z Distributions as a Global Check of the Acceptance*

Many potential detector modeling problems would affect the acceptance as a function of Z , so a crucial check of our understanding of the acceptance is to compare the Z distributions for the data against the MC simulation. The statistics are a bit limited in the $K \rightarrow \pi\pi$ samples, so our primary check is to use the higher-statistics $K_L \rightarrow \pi e\nu$ and $K_L \rightarrow 3\pi^0$ samples, which should have similar sensitivities to de-

tector modeling problems. Figure 10.2a shows the vacuum-beam comparisons for all four of these decay modes. In these plots, the data and MC distributions are essentially indistinguishable, and so we plot the *ratio* for each decay mode, normalized to have a mean value of one, in Figure 10.2b. We see that the $\pi e \nu$ comparison may have a bit of a downward slope, while the $3\pi^0$ comparison is consistent with

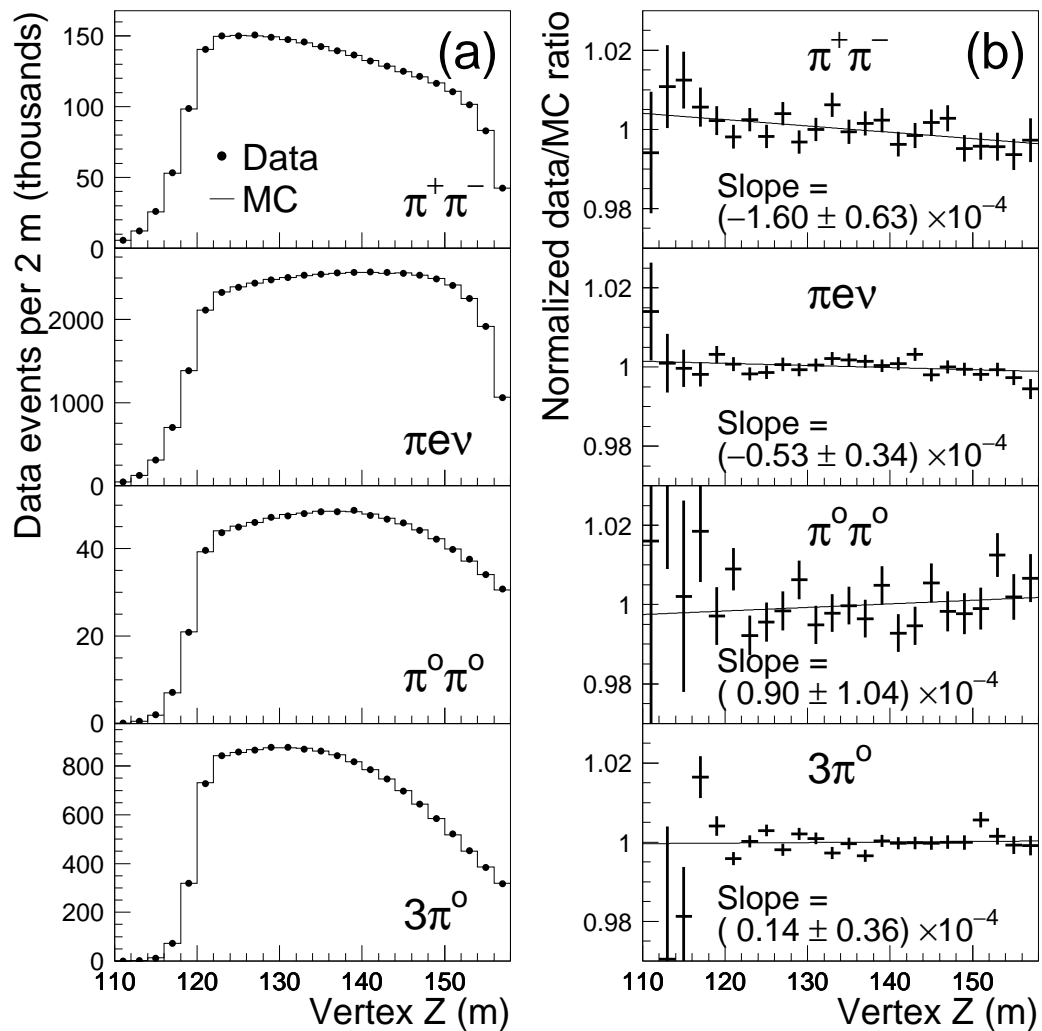


Figure 10.2: (a) Data versus Monte Carlo comparisons of vacuum-beam Z distributions for $\pi^+\pi^-$, $\pi e \nu$, $\pi^0\pi^0$, and $3\pi^0$ decays. (b) Linear fits to the data/MC ratio of Z distributions for each of the four decay modes.

being flat. Adding the central value plus one sigma, we may use these comparisons to limit the magnitudes of the Z slopes to 0.87×10^{-4} and 0.50×10^{-4} .

However, the apparent nonzero slope in the $\pi^+\pi^-$ comparison (2.5σ different from zero) is a cause for concern, and much effort has gone into studying it.⁶ It is natural to suspect the drift chambers, especially since downstream decays pass through the damaged beam regions more often, but the KTEVMC model which reproduces a variety of distributions correctly does not induce a significant Z slope.⁷ Pion “punchthrough” to the muon veto hodoscope, which is not properly simulated in KTEVMC, has also been considered and found to have very little effect on the Z distribution.

The $\pi^+\pi^-$ slope is not so inconsistent with the slope from the $\pi e\nu$ comparison (1.5σ), so there is the possibility that it may be due, at least in part, to a statistical fluctuation. However, in light of the discrepancy, we decided to base our systematic error on the $\pi^+\pi^-$ slope, rather than on the $\pi e\nu$ slope as originally intended. Thus we take the central value of the $\pi^+\pi^-$ slope, 1.6×10^{-4} , as the uncertainty on the charged-mode acceptance slope.

One can show that, to a good approximation, an acceptance slope s affects the measured value of $\text{Re}(\epsilon'/\epsilon)$ according to the simple formula $s \Delta z/6$, where Δz is the difference in the mean z values for vacuum- vs. regenerated-beam events, which is 5.54 m for $\pi^+\pi^-$ and 7.15 m for $\pi^0\pi^0$. Thus, the slope uncertainties assigned above would correspond to systematic errors on $\text{Re}(\epsilon'/\epsilon)$ of 1.48×10^{-4} and 0.60×10^{-4} , respectively. By introducing the slopes directly into acceptance calculation in the fitting program, we arrived at slightly higher values, 1.60×10^{-4} and 0.68×10^{-4} .

Other general checks of the acceptance include data/MC comparisons of track illuminations at the drift chambers and CsI, photon illumination at the CsI, and minimum photon separation distance. These all agree well and indicate no other sources of acceptance misunderstanding.

⁶The regenerated-beam Z slope has a much larger statistical error since the events are clustered together in Z . In any case, it is consistent both with the vacuum-beam slope and with zero.

⁷If the track separation cut is removed, then the model *does* induce a slope, because tracks which are close enough to share wires have incomplete hit information to begin with and are more susceptible to being lost from additional drift chamber inefficiencies and/or delayed hits.

10.6 K_S/K_L Flux Ratio

It is essential that the K_S/K_L flux ratio be the same for the $\pi^+\pi^-$ and $\pi^0\pi^0$ samples. This would automatically be the case if we collected charged and neutral decays at the same time. The fact that we use $\pi^+\pi^-$ and $\pi^0\pi^0$ data from *different* running periods for this analysis is, nevertheless, not a major concern because the Regenerator and movable absorber were the same for both periods.⁸ We must, however, consider the effect of temperature on the densities of the movable absorber and Regenerator. The density drops as the cube of the thermal expansion, but the length increases linearly, so the net result is that the number of interaction lengths drops as the square of the expansion. The length of the Regenerator is such that K_S production is nearly maximal, so only the movable absorber can significantly affect the flux. Temperature records for the target enclosure indicate that the temperature difference between 1996 and 1997 was no more than about 5° C. Using the thermal expansion coefficient for beryllium, we place a conservative limit on the resulting $\text{Re}(\epsilon'/\epsilon)$ bias of 0.20×10^{-4} .

In addition, the $\pi^+\pi^-$ and $\pi^0\pi^0$ samples have somewhat different energy distributions, so the energy dependence of the regenerated-beam flux attenuation makes the average K_S/K_L flux ratio slightly different for the two samples. This energy-dependent attenuation is taken into account when fitting to extract $\text{Re}(\epsilon'/\epsilon)$, but the *uncertainty* on the energy dependence leads to an uncertainty of 0.24×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$.

10.7 Dependence on Other Physics Parameters

The measured value of $\text{Re}(\epsilon'/\epsilon)$ depends on the values of Δm and τ_S in the following way:

⁸We could not have used data from different running periods had we not moved the Regenerator back and forth so that it spent essentially equal amounts of time in the east and west beams. By doing so, any small flux difference between the east and west beams (due, say, to collimator holes of slightly different sizes) averages out for both beams.

$$\begin{aligned} \Delta\text{Re}(\epsilon'/\epsilon) &= (+0.05 \times 10^{-4}) \times \frac{\Delta m - 0.5301}{0.0014} \\ &+ (-0.11 \times 10^{-4}) \times \frac{\tau_S - 0.8934}{0.0008} \end{aligned} \quad (10.1)$$

where Δm and τ_S are in units of $10^{10} \hbar \text{ s}^{-1}$ and 10^{-10} s , respectively. The reference values in the numerators above are the published world averages [32], and were used for the nominal $\text{Re}(\epsilon'/\epsilon)$ fit. The values in the denominators are the errors on the averages. We will report a systematic error on $\text{Re}(\epsilon'/\epsilon)$ which includes one-sigma variations in the values of Δm and τ_S .⁹

The measured value of $\text{Re}(\epsilon'/\epsilon)$ has negligible dependence on the value of ϵ or on the K_L lifetime, given the accuracies to which these are known. There is also very little dependence on the relative numbers of K^0 and \bar{K}^0 produced in the target, and in any case the data verifies, to the level of a few percent, that our model gives the correct relative fluxes.

We allow the complex phase of the regeneration amplitude to deviate from the value predicted by analyticity by $\pm 0.5^\circ$; this leads to a systematic error of 0.15×10^{-4} on $\text{Re}(\epsilon'/\epsilon)$. For comparison, it was argued that the analyticity assumption was good to at least 0.35° for the E773 experiment [79], which included lower kaon energies where more of a deviation is expected.

10.8 Cross Checks

In addition to the studies described above, which assigned systematic errors on $\text{Re}(\epsilon'/\epsilon)$, we have done a number of “cross checks” which give us additional confidence in the result. Some of these have already been described in Section 9.3. Other studies checked for a variation of the result as a function of time or of beam intensity; no significant variation was observed in either case.

One final check deserves special mention. After our initial decision to use $\pi^+\pi^-$

⁹Preliminary fits to KTeV data give $\Delta m = (0.5280 \pm 0.0013) \times 10^{10} \hbar \text{ s}^{-1}$ and $\tau_S = (0.8967 \pm 0.0007) \times 10^{-10} \text{ s}$ [78], the latter of which is significantly different from the published world average. A $\text{Re}(\epsilon'/\epsilon)$ fit done using these values finds $\text{Re}(\epsilon'/\epsilon) = 27.5 \times 10^{-4}$, with a χ^2 of 27.2 for 30 degrees of freedom.

data from 1997 instead of from 1996 because of the high Level 3 inefficiency due to drift chamber problems, we made many improvements to the drift chamber simulation in KTEVMC. As a result, we now expect that the MC should be able to simulate the 1996 acceptance reasonably well. We therefore repeated the full $\pi^+\pi^-$ analysis on the 1996 data and generated a large MC sample to determine the acceptance, so that we could fit for $\text{Re}(\epsilon'/\epsilon)$ from charged and neutral data collected at the same time. This fit found a central value for $\text{Re}(\epsilon'/\epsilon)$ of 24.6×10^{-4} , which is consistent with the standard fit within the systematic uncertainty we estimated from our modeling of the 1996 Level 3 inefficiency, 4×10^{-4} .

10.9 Summary of Systematic Errors

Table 10.2 summarizes all of the estimated contributions.¹⁰ Adding all contributions in quadrature, the total systematic uncertainty on $\text{Re}(\epsilon'/\epsilon)$ is 2.85×10^{-4} —actually slightly less, so that it rounds to 2.8×10^{-4} . This is smaller than the statistical error, but only slightly. It also happens to be almost exactly the same as the systematic error ultimately achieved by E731 (after several years of studying their detector). In any case, with this number in hand, the analysis is finally complete.

¹⁰Some of the categories were combined differently in the published account of this analysis [28].

Table 10.2: Systematic uncertainties on $\text{Re}(\epsilon'/\epsilon)$.

Systematic error source	Uncertainty ($\times 10^{-4}$)	
	from $\pi^+\pi^-$	from $\pi^0\pi^0$
Trigger inefficiencies	0.51	0.29
Event reconstruction biases		
Calorimeter cluster energy measurement	—	0.92
Calorimeter cluster position measurement	—	0.38
Spectrometer alignment and calibration	0.30	—
Analysis cuts	0.57	0.78
Backgrounds	0.20	0.81
Detector acceptance		
Monte Carlo statistics	0.52	0.90
Limiting apertures	—	0.48
Effective Regenerator edge	0.26	—
Detector resolutions	0.35	0.02
Drift chamber inefficiencies	0.62	—
Accidental activity	0.11	—
Global check of Z distributions	1.60	0.68
K_S/K_L flux ratio		
1996 versus 1997		0.20
Energy dependence		0.24
Dependence on other physics parameters		
Δm		0.05
τ_S		0.11
Relax analyticity assumption		0.15
TOTAL		2.85

CHAPTER 11

CONCLUSION

Using a subset of the data collected by the KTeV experiment, we have made a new precision measurement of the direct- CP -violation parameter $\text{Re}(\epsilon'/\epsilon)$. We find that

$$\begin{aligned}\text{Re}(\epsilon'/\epsilon) &= [28.0 \pm 3.0 \text{ (stat)} \pm 2.8 \text{ (syst)}] \times 10^{-4} \\ &= (28.0 \pm 4.1) \times 10^{-4}, \qquad \text{(this result)}\end{aligned}$$

where, in the second expression, the statistical and systematic errors have been added in quadrature. This result is shown graphically in Figure 11.1, along with the other results published since 1988. Our result is consistent with the NA31 result, and substantially higher than E731. We have spent some time reviewing the details of the E731 analysis [76, 58] and have found no reason to question the validity of that result; therefore we will assume it to be an unusually large statistical fluctuation. Combining our result with the final results from E731 and NA31, the new world average is

$$\text{Re}(\epsilon'/\epsilon) = (21.7 \pm 3.0) \times 10^{-4}. \qquad \text{(new world average)}$$

The most basic conclusion is that we have now definitively established the existence of direct CP violation. Finally, 35 years after the initial discovery of $K_L \rightarrow \pi\pi$ decays, there are *two* independent (and well-measured) parameters of CP violation, ϵ and ϵ' .




Figure 11.1: Measurements of $\text{Re}(\epsilon'/\epsilon)$, by publication date. The measurements marked with open boxes are superseded by later measurements which include the same data.

11.1 Implications for Theory

The fact that ϵ' is nonzero rules out the possibility that a “superweak” ($\Delta S = 2$) interaction is the sole source of CP violation in the K meson system. This does not mean that a superweak interaction cannot exist, but it does remove a major motivation for hypothesizing the existence of such an interaction.

The measured value of $\text{Re}(\epsilon'/\epsilon)$ is roughly compatible with CKM-model predictions (see Section 2.5), supporting the notion of a nonzero phase in the CKM matrix. Although this does not constitute proof, it does provide additional circumstantial evidence for the continued success of the Standard Model in describing particle interactions.

However, the new world average is substantially larger than some CKM-model calculations, which has stirred up new interest among theorists in considering the

possibility that there is a real discrepancy. There are two classes of proposed explanations:

1. One (or more) of the parameters used as input to the calculations may be incorrect. The most critical ingredients are the hadronic matrix element parameters for QCD and electroweak penguin diagrams and the mass of the strange quark. For example, a lower value for the strange quark mass leads to a higher calculated value of $\text{Re}(\epsilon'/\epsilon)$, all other things being equal.
2. A new physics process may contribute to direct CP violation. For example, Masiero and Murayama have recently suggested that a supersymmetric model can contribute an amount to $\text{Re}(\epsilon'/\epsilon)$ which is comparable to the measured value [80].

The latter possibility is tantalizing, but we cannot give it much weight until the theoretical uncertainties in the CKM-model calculations are reduced. Efforts are currently under way to determine the input parameters more accurately using lattice gauge calculations, etc.

On the experimental front, future measurements of the neutron electron dipole moment should be able to constrain (or support!) the supersymmetric hypothesis above. Rare K_L decays, especially $K_L \rightarrow \pi^0 \nu \bar{\nu}$, promise to shed new light on direct CP violation once the experimental challenges have been overcome. Ultimately, multiple measurements of CP violation in the B meson system will test the internal consistency of the CKM model.

11.2 Prospects for Additional Measurements of $\text{Re}(\epsilon'/\epsilon)$

The NA48 collaboration has already announced a preliminary measurement of $\text{Re}(\epsilon'/\epsilon)$ based on data collected in 1997 [81, 82]. They find

$$\begin{aligned} \text{Re}(\epsilon'/\epsilon) &= [18.5 \pm 4.5 \text{ (stat)} \pm 5.8 \text{ (syst)}] \times 10^{-4} \\ &= (18.5 \pm 7.3) \times 10^{-4}, \quad \text{(NA48 preliminary result)} \end{aligned}$$

in excellent agreement with the average of the average value presented earlier in this chapter.

KTeV collected roughly three times as much data in 1997 as in 1996, and analysis of the rest of the data is now under way. Also, KTeV is presently collecting additional data with improved drift chamber performance¹ and calorimeter reliability,² and expects to collect enough data to allow an ultimate statistical error of around 1×10^{-4} . Much work will likely be needed to reduce the systematic error to a similar level; in particular, the small data/MC mismatch in the Z distributions will have to be understood. NA48 also has additional data, collected in 1998 after various detector improvements, and is continuing to run. They expect to eventually reach a statistical error of around 2×10^{-4} . The KLOE experiment will also measure $\text{Re}(\epsilon'/\epsilon)$ in the coming years, though it may take many years to reach a precision comparable to KTeV and NA48. The results expected from all of these experiments promise to give us an even more precise knowledge of direct CP violation in the K meson system.

¹In preparation for the 1999 run, the damaged beam-region wires were replaced, the chambers were more tightly sealed, and the gas system was modified to provide better control over the alcohol content. The electronic gain was increased and the noise level reduced, allowing the discriminator threshold to be well below the pulse height from a single drift electron. As a result, the high-sod effect has been virtually eliminated.

²The QIE and DBC chips on all of the DPMTs were replaced once again, and failures are now very rare.

REFERENCES

- [1] G. Rochester and C. Butler, *Nature* **160**, 855 (1947).
- [2] P. Galison, *Image and Logic: A Material Culture of Microphysics*, University of Chicago Press, Chicago, 1997.
- [3] M. Gell-Mann, *Phys. Rev.* **92**, 833 (1953); T. Nakano and K. Nishijima, *Prog. Theor. Phys.* **10**, 581 (1953).
- [4] R. Belušević, *Neutral Kaons*, Springer-Verlag, Berlin, 1999.
- [5] M. Gell-Mann and A. Pais, *Phys. Rev.* **97**, 1387 (1955).
- [6] J. W. Cronin, after-dinner remarks at the KAON99 Conference banquet, Chicago, Illinois, June 25, 1999.
- [7] R. H. Dalitz. “50+ years of kaon physics”. Talk given at the KAON99 Conference, Chicago, Illinois, June 21, 1999.
- [8] K. Lande, E. T. Booth, J. Impeduglia, and L. M. Lederman, *Phys. Rev.* **103**, 1901 (1956).
- [9] J. Schwinger, *Phys. Rev.* **91**, 713 (1953); *Phys. Rev.* **94**, 1362 (1954).
- [10] G. Lüders, *Kgl. Danske Videnskab Selskab., Mat-fys. Medd.* **28**, 1 (1954).
- [11] W. Pauli, *Niels Bohr and the Development of Physics*, Pergamon Press, Elmsford, NY, 1955.
- [12] R. H. Dalitz, “The Modes of Decay of the τ -Meson”. In *Congrès International sur le Rayonnement Cosmique* (informal publication), July, 1953; “Decay of τ -Mesons of Known Charge”, *Phys. Rev.* **94**, 1046 (1954).
- [13] T. D. Lee and C. N. Yang, *Phys. Rev.* **104**, 254 (1956).
- [14] C. S. Wu et al., *Phys. Rev.* **105**, 1413 (1957).
- [15] R. L. Garwin, L. M. Lederman, and M. Weinrich, *Phys. Rev.* **105**, 1415 (1957).
- [16] J. I. Friedman and V. L. Telegdi, *Phys. Rev.* **105**, 1681 (1957).

- [17] T. D. Lee, R. Oehme, and C. N. Yang, *Phys. Rev.* **106**, 340 (1957).
- [18] L. Landau, *JETP* **5**, 336 (1957).
- [19] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, *Phys. Rev. Lett.* **13**, 138 (1964).
- [20] V. L. Fitch, *Rev. Mod. Phys.* **53**, 367 (1981).
- [21] J. W. Cronin, *Rev. Mod. Phys.* **53**, 373 (1981).
- [22] V. L. Fitch, R. F. Roth, J. S. Russ, and W. Vernon, *Phys. Rev. Lett.* **15**, 73 (1965).
- [23] M. Banner, J. W. Cronin, J. K. Liu, and J. E. Pilcher, *Phys. Rev. Lett.* **21**, 1103 (1968).
- [24] I. A. Budagov et al., *Phys. Lett. B* **28**, 215 (1968).
- [25] See, for example, K. R. Schubert *et al.*, *Phys. Lett. B* **31**, 662 (1970).
- [26] A. D. Sakharov, *JETP Lett.* **5**, 24 (1967).
- [27] B. Winstein and L. Wolfenstein, *Rev. Mod. Phys.* **65**, 1113 (1993).
- [28] A. Alavi-Harati *et al.* (KTeV Collaboration), *Phys. Rev. Lett.* **83**, 22 (1999).
- [29] P. K. Kabir, *The CP Puzzle*, Academic Press, New York, 1968.
- [30] R. G. Sachs, *The Physics of Time Reversal*, University of Chicago Press, Chicago, 1987.
- [31] K. Kleinknecht, in *CP Violation*, edited by C. Jarlskog, World Scientific, Singapore, 1989, p. 41.
- [32] Particle Data Group, C. Caso et al., *Eur. Phys. J. C* **3**, 1 (1998).
- [33] A. Angelopoulos et al., *Phys. Lett.* **B444**, 38 (1998); *Phys. Lett.* **B444**, 43 (1998).
- [34] T. T. Wu and C. N. Yang, *Phys. Rev. Lett.* **13**, 380 (1964).
- [35] W. Ochs, *πN Newsletter* **3**, 25 (1991).
- [36] H. Burkhardt et al., *Phys. Lett. B* **206**, 169 (1988).
- [37] G. D. Barr et al., *Phys. Lett.* **B317**, 233 (1993).
- [38] M. Woods et al., *Phys. Rev. Lett.* **60**, 1695 (1988).

- [39] J. R. Patterson et al., *Phys. Rev. Lett.* **64**, 1491 (1990).
- [40] L. K. Gibbons et al., *Phys. Rev. Lett.* **70**, 1203 (1993).
- [41] Particle Data Group, L. Montanet et al., *Phys. Rev. D* **50**, 1173 (1994).
- [42] M. Kobayashi and T. Maskawa, *Prog. Theor. Phys.* **49**, 652 (1973).
- [43] J. Ellis, M. K. Gaillard, and D. V. Nanopoulos, *Nucl. Phys.* **B109**, 213 (1976).
- [44] F. J. Gilman and M. B. Wise, *Phys. Lett.* **83B**, 83 (1979).
- [45] B. Winstein, private communication.
- [46] J. Bijnens and M. B. Wise, *Phys. Lett.* **137B**, 245 (1984).
- [47] S. Bertolini, M. Fabbrichesi, and Jan O. Eeg. “Estimating ϵ'/ϵ : A Review”. SISSA Report 19-98-EP; hep-ph/9802405; to appear in *Rev. Mod. Phys.*
- [48] M. Ciuchini, *Nucl. Phys. Proc. Suppl.* **59**, 149 (1997).
- [49] A. J. Buras, in *Probing the Standard Model of Particle Interactions*, edited by R. Gupta, Elsevier, 1999. This estimate assumes that the mass of the strange quark is $m_s(m_c) = 125 \pm 20 \text{ MeV}/c^2$.
- [50] S. Bertolini et al., *Nucl. Phys.* **B514**, 93 (1998).
- [51] T. Blum et al. “A First Study of ϵ'/ϵ on the Lattice Using Domain Wall Fermions”. RIKEN BNL Research Center preprint BNL-66731; hep-lat/9908025.
- [52] L. Wolfenstein, *Phys. Rev. Lett.* **13**, 562 (1964).
- [53] E. J. Ramberg et al., *Phys. Rev. Lett.* **70**, 2529 (1993).
- [54] A. Alavi-Harati *et al.* (KTeV Collaboration). Submitted to *Phys. Rev. Lett.*; preprint hep-ex/9908020.
- [55] L. M. Sehgal and M. Wanninger. *Phys. Rev.* **D46**, 1035 (1992); *Phys. Rev.* **D46**, 5209(E) (1992).
- [56] CDF Collaboration, “A Measurement of $\sin 2\beta$ from $B \rightarrow J/\psi K_S^0$ with the CDF Detector (preliminary)”, Technical Report CDF/PUB/BOTTOM/CDF/4855, Fermilab, February 1999.
- [57] K. Arisaka et al., *KTeV Design Report*, Technical Report FN-580, Fermilab, 1992.
- [58] L. K. Gibbons et al., *Phys. Rev. D* **55**, 6625 (1997).

- [59] G. D. Barr et al., “Proposal for a Precision Measurement of ϵ'/ϵ in CP Violating $K^0 \rightarrow 2\pi$ Decays”, Technical Report CERN/SPSC/90-22, CERN, 1990.
- [60] V. Bocean et al., “KTeV Beam Systems Design Report”, Technical Report TM-2023, Fermilab, 1997.
- [61] T. Kobilarcik, “Results of Stability Study for the KTeV Beam”, Technical Report TM-2037, Fermilab, 1997.
- [62] V. Bocean et al., “Performance of the KTeV High-Energy Neutral Kaon Beam at Fermilab”, Technical Report TM-2046, Fermilab, 1998.
- [63] A. Roodman. (title?). in *Proceedings of the Seventh International Conference on Calorimetry in High Energy Physics*, edited by E. Cheu et al., page 89. World Scientific, 1998.
- [64] J. Whitmore, *Nucl. Instrum. Methods Phys. Res., Sect. A* **409**, 687 (1998).
- [65] E. D. Zimmerman, Ph.D. thesis (The University of Chicago, March, 1999).
- [66] C. Bown et al., *Nucl. Instrum. Methods Phys. Res., Sect. A* **369**, 248 (1996).
- [67] A. J. Malensek, 1981. Fermilab Reports FN-341, FN-341A (errata).
- [68] H. W. Atherton et al., 1980. CERN preprint 80:07.
- [69] W. R. Molzon, Ph.D. thesis (The University of Chicago, March, 1979).
- [70] G. C. Bown, private communication.
- [71] R. Brun et al., 1994. computer code GEANT 3.21, CERN, Geneva.
- [72] W. R. Nelson, H. Hirayama, and W. O. Rogers, “The EGS4 Code System”, Technical Report SLAC-265, Stanford Linear Accelerator Center, 1985.
- [73] R. A. Briere, Ph.D. thesis (The University of Chicago, June, 1995).
- [74] B. Schwingenheuer, Ph.D. thesis (The University of Chicago, June, 1995).
- [75] J. R. Patterson, Ph.D. thesis (The University of Chicago, December, 1990).
- [76] L. K. Gibbons, Ph.D. thesis (The University of Chicago, August, 1993).
- [77] F. James and M. Roos, Computer code MINUIT, CERN, Geneva, 1994.
- [78] Y. B. Hsiung. “KTeV Result on $\text{Re}(\epsilon'/\epsilon)$ ”. Talk given at the KAON99 Conference, Chicago, Illinois, June 21, 1999.
- [79] R. A. Briere and B. Winstein, *Phys. Rev. Lett.* **75**, 402 (1995).

- [80] A. Masiero and H. Murayama, *Phys. Rev. Lett.* **83**, 907 (1999).
- [81] P. Debu. “First result on direct CP violation in Kaon decays from NA48”. CERN seminar, Geneva, Switzerland, June 18, 1999.
- [82] M. Sozzi. “NA48 Result on ϵ'/ϵ ”. Talk given at the KAON99 Conference, Chicago, Illinois, June 21, 1999.