

STUDY OF THE RARE DECAY $K_L \longrightarrow e^+e^-\mu^+\mu^-$ AT KTEV

by

Jason C. Hamm

A Dissertation Submitted to the Faculty of the
DEPARTMENT OF PHYSICS
In Partial Fulfillment of the Requirements
For the Degree of
DOCTOR OF PHILOSOPHY
In the Graduate College
THE UNIVERSITY OF ARIZONA

2 0 0 2

Insert correct approval form here AND PAGE NUMBER!!!

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at The University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgment of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the head of the major department or the Dean of the Graduate College when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

SIGNED: _____

ACKNOWLEDGEMENTS

After five years, I'm leaving Tucson with a Ph.D. and a lot less hair. There are people responsible for both, but for now I'd like to thank the people who made my graduate career a rewarding and successful one. First mention goes to Elliott Cheu, my graduate adviser. Without his advice, knowledge, humor, and patience, this could never have come to fruition.

KTeV is a big experiment, and my work was made possible only with the assistance of many collaborators. While I feel privileged to have worked with each of these talented, respected people, there are a few who deserve special mention. Sydney showed me the ropes, helped me get started, and introduced me to Chicago. Thanks to him, Julie, Peter, and Patty, I was able to feel at home during my time at Fermilab. Additionally, I want to acknowledge the members of the Colorado group, Tony, Pat, Jason, and Hanwen, both for the radiative corrections MC and for letting me hang out with them after collaboration meetings. I also had a lot of help along the way from Bob T., Rick K., both Sasha G.s, Val, Elizabeth, Breese, Amit, and Jianbo.

Returning to Tucson...I want to thank the members of the EEPP group for their assistance and instruction over the years. Thanks also go to Mike and Phil for all their computer help. Of course, the LSS crowd has been a big part of my life during my time here. Special recognition goes to Mikey, Tommy, Chuckie, Rob, and Todd, who have been there for me from start to finish, and Sarah, who came along near the end but whose friendship is just as valuable. A debt of gratitude also goes my friends from OC and Berkeley who have stuck by me over the years: Reeve, Holly, Brians C. and K., Jay, Andrew, and Jo, among many others.

Most of all I want to thank my family, especially my parents Jerry and Barbara, for their incredible support. During the toughest times, they have always provided words of encouragement. Even though my life is taking a pretty sharp turn right now, I know they will always be there for me, no matter what happens. Thank you.

Jason C. Hamm
Tucson, AZ
5 July, 2002

DEDICATION

For my parents.



“Things are gonna slide...

slide in all directions.

There won't be nothing...

nothing you can measure anymore.”

- Leonard Cohen
“The Future”

TABLE OF CONTENTS

LIST OF FIGURES	9
LIST OF TABLES	18
ABSTRACT	19
1 Theoretical Motivation for the Study of $K_L \rightarrow e^+e^-\mu^+\mu^-$	20
1.1 Quark Mixing and the CKM Matrix	20
1.2 Review of the Neutral Kaon System	22
1.3 The Physics of $K_L \rightarrow \mu^+\mu^-$	25
1.3.1 Connecting $K_L \rightarrow \mu^+\mu^-$ to the CKM matrix	26
1.3.2 Calculating the long-distance contribution to $K_L \rightarrow \mu^+\mu^-$	29
1.3.3 Models of the $K_L \gamma^* \gamma^*$ Vertex	31
1.3.4 Current experimental status of α and β	36
1.4 Introduction to The Rare Decay $K_L \rightarrow e^+e^-\mu^+\mu^-$	38
1.4.1 Branching ratio predictions for $K_L \rightarrow e^+e^-\mu^+\mu^-$	39
1.4.2 Experimental status of $K_L \rightarrow e^+e^-\mu^+\mu^-$	43
1.5 Dissertation Overview	44
2 The KTeV Experiment	45
2.1 Creation of the K_L beams	46
2.1.1 The proton beam	46
2.1.2 The NM2 enclosure	48
2.2 The KTeV Detector	50
2.2.1 Vacuum region	52
2.2.2 Spectrometer	52
2.2.3 Transition radiation detectors	56
2.2.4 Trigger hodoscopes	56
2.2.5 Electromagnetic calorimeter	56
2.2.6 Hadron anti	62
2.2.7 Muon identification system	63
2.2.8 Photon vetos	65
3 The KTeV Trigger and Event Selection	67
3.1 Level 1	67
3.1.1 Beam timing sources	68
3.1.2 Etotal	68

3.1.3	DC-ORs	69
3.2	Level 2	72
3.2.1	Hardware Cluster Counter	72
3.2.2	DC Hit Counters (the fruit)	82
3.2.3	Y-Track Finder	83
3.3	Physics Triggers	83
3.3.1	Physics triggers in 1997	83
3.3.2	Physics triggers in 1999	85
3.4	Calibration Triggers	86
3.5	Digitization and Readout	86
3.6	Level 3	87
3.7	Online Split	88
4	Event Reconstruction	90
4.1	Tracking	90
4.1.1	Good SOD, Bad SOD	90
4.1.2	y track finding	93
4.1.3	x track finding	93
4.2	Clustering	94
4.2.1	Hardware clustering	94
4.2.2	Software clustering	94
4.2.3	Cluster positioning	95
4.2.4	Clustering corrections	95
4.3	Vertexing	102
5	KTeV Monte Carlo	107
5.1	Simulation of K_L Production	108
5.2	Simulation of K_L Decay	109
5.2.1	$K_L \longrightarrow e^+e^-\mu^+\mu^-$ generator	109
5.2.2	$K_L \longrightarrow \pi^+\pi^-\pi^0$ generator	111
5.2.3	$\pi^0 \longrightarrow e^+e^-\gamma$ generator	111
5.2.4	$K_L \longrightarrow \mu^+\mu^-\gamma$ generator	112
5.2.5	$K_L \longrightarrow \pi^+\pi^-e^+e^-$ generator	112
5.3	Particle Tracing	113
5.4	Particle Interactions	113
5.4.1	Photon conversion	114
5.4.2	Multiple scattering	114
5.4.3	δ -rays	116
5.4.4	Bremsstrahlung	116
5.5	Particle Detection	117
5.5.1	Photon veto simulation	117

5.5.2	Spectrometer simulation	118
5.5.3	VV', MU2, and MU3 simulation	122
5.5.4	Calorimeter simulation	122
5.5.5	Etotal and HCC simulation	124
5.6	Accidental Overlays	127
5.7	Trigger Simulation	128
5.8	Monte Carlo Samples	129
6	Branching Ratio Analysis	130
6.1	Crunch	130
6.2	Bad Spill Cut	132
6.3	Trigger Verification	135
6.4	Particle Identification	136
6.5	Initial Cuts	138
6.6	Normalization Mode	141
6.6.1	Further cuts	141
6.6.2	Persistent acceptance effects	145
6.6.3	Charged pion loss correction	148
6.6.4	Kaon flux calculation	152
6.7	Signal Mode	152
6.7.1	Further cuts	152
6.7.2	Backgrounds	153
6.7.3	Background and cut summary	163
6.8	Branching Ratio Calculation	169
6.8.1	Systematic errors	170
6.8.2	Tabulation of systematic errors	185
6.9	Limit on $K_L \longrightarrow e^\pm e^\pm \mu^\mp \mu^\mp$ Branching Ratio	185
7	$K_L \gamma^* \gamma^*$ Form Factor Analysis	187
7.1	Expected Sensitivity	187
7.2	Measuring α Using a Shape Analysis	189
7.3	Measuring β Using a Shape Analysis	196
7.4	Measuring α Using the $K_L \longrightarrow e^+ e^- \mu^+ \mu^-$ Branching Ratio	196
7.5	Search for CP -violation in the Decay $K_L \longrightarrow e^+ e^- \mu^+ \mu^-$	199
8	Conclusions	202
	REFERENCES	206

LIST OF FIGURES

1.1	A flavor-changing transition mediated by the weak force. The top quark (t) changes into a down quark (d) by coupling to a W gauge boson. The probabilities of such transitions are described in the Standard Model by the CKM matrix (in this example, the matrix element V_{td}).	20
1.2	One of several diagrams contributing to oscillations between K^0 and its antiparticle, \bar{K}^0 . The oscillations are made possible by second order, $\Delta S = 2$ weak interactions.	23
1.3	Short-distance contributions to the $K_L \rightarrow \mu^+ \mu^-$ amplitude. The quark loops in these diagrams are dominated by top, making this mode sensitive to the CKM matrix element V_{td} (or ρ , in the Wolfenstein parameterization).	26
1.4	Long-distance contribution to the $K_L \rightarrow \mu^+ \mu^-$ amplitude. This diagram must be subtracted from the total $K_L \rightarrow \mu^+ \mu^-$ amplitude in order to extract the short-distance contribution, and with it, ρ . Unfortunately, the form of the $K_L \gamma^* \gamma^*$ vertex is not well understood, making this calculation difficult.	27
1.5	Vector Meson Dominance model for the $K_L \gamma^* \gamma^*$ vertex. Interactions between the pseudoscalar K_L and the virtual photons γ^* are mediated by vector mesons, V .	32
1.6	Diagram for the rare decay $K_L \rightarrow e^+ e^- \mu^+ \mu^-$. This decay serves as the best available probe of the quadratic β term of the DIP form factor.	38
1.7	Graphical representation of ϕ , the angle between the normals to the electron and muon decay planes in the K_L rest frame.	40
2.1	Calendar depicting the three run periods of the KTeV E799 experiment: Winter, Summer, and 1999.	46
2.2	y view of the beamline elements located in the NM2 enclosure. The common absorber and jaw collimators were not used during E799 running.	49
2.3	y view schematic of the KTeV detector in the E799 configuration.	51
2.4	y view schematic of the field and sense wire arrangement within a drift chamber.	54
2.5	V and V' trigger hodoscope planes.	57
2.6	Schematic of the KTeV electromagnetic calorimeter. 3100 pure CsI crystals were used in the array.	58
2.7	E/p of calibration K_{e3} electrons in all 3100 channels after five iterations.	62
2.8	y view of the KTeV detector downstream of the EM calorimeter.	63

2.9	Schematic of muon trigger planes. Dashed lines indicate the 1 cm counter overlaps in MU2.	64
2.10	Schematics of typical photon veto detectors: on the left, RC10; on the right, SA4.	65
2.11	Schematic of the CA on the calorimeter face. The blocks represent the small crystals of the inner region of the calorimeter.	66
3.1	Top: The five curves shown represent (from left to right) the total energy distribution from all calibration events, and the total energy distributions of calibration events that fired the ET1, ET2, ET3, and ET4 Level 1 trigger sources. Energies are measured in GeV. Bottom: Threshold curves for L1 sources ET1–4, obtained by dividing each of the L1 source distributions by the distribution for all events. These plots were created from the online monitoring system during the 1999 run. The smiley faces indicate that the Etot total thresholds, widths, and efficiencies are within tolerances.	70
3.2	ET3 threshold as a function of time during the 1999 run. The line depicts the nominal value of 25 GeV. Upward drift away from the nominal value was caused mainly by radiation damage to the calorimeter. The threshold could be returned to the nominal value by adjusting the discrimination level on the NIM module, or by raising the high voltage on the calorimeter PMTs. This variation was modeled in the KTeV monte carlo.	71
3.3	Top: The open histogram is the distribution of in-time energy deposited in calorimeter channel number 1492 from all calibration electrons. The solid histogram is the in-time energy distribution for events in which the HCC bit for this channel fired. Energies are measured in GeV. Bottom: HCC threshold curve for calorimeter channel 1492, obtained by dividing the solid histogram by the open one. The turn-on region has a narrow width and is centered at approximately 1 GeV.	74
3.4	Distribution of all 3100 HCC bit thresholds before (top) and after (bottom) gain matching. Individually tuning the high voltages of the calorimeter PMTs has moved the mean of the distribution closer to the desired HCC threshold value of 1 GeV, and has also reduced the spread.	75
3.5	(a) Higher (solid) and lower (dashed) energy dynode pulses, as seen by a comparator on a first level Etot sum board. In this example, both pulses exceed the comparator threshold. Note that the higher energy pulse crosses threshold earlier. (b) Comparator output. Because the leading edge of the output from the higher energy pulse occurs during the in-time window, an HCC bit would be set for the channel. The lower energy output misses the in-time window, meaning no HCC bit would be set. The overall HCC threshold for this channel is therefore biased towards a higher value. The relative location of the in-time window can be shifted either by changing the global RF phase to L1, or by shortening the length of the dynode cable for the channel.	77

3.6	HCC thresholds and turn-on widths vs. RF phase as measured at L1. The nominal phase is 0 ns. Channel 844 (left) nominally sits in an unstable region – a small negative shift in the phase would lead to large shifts in HCC threshold and width. This channel had its dynode cable trimmed by 3 ns in order to move its nominal position onto the plateau. A normal channel, 1042, is shown for comparison (right). The nominal position of this channel is on the stable plateau.	78
3.7	Examples of clusters from EM shower energy deposited in the calorimeter. Blocks with their HCC bit on are denoted by heavy black boxes. The two clusters found by the HCC are circled. The other activity in the calorimeter is out-of-time; therefore, these blocks do not have their HCC bits on, and no clusters are found there. The energy scale in this event display spans almost three orders of magnitude.	80
3.8	Possible HCC bit patterns and corresponding weights for a 2×2 group of blocks. Filled blocks represent those with their HCC bit on. Note that the +2 patterns could also be viewed as having no right turns – the degeneracy was broken in favor of +2 in order to correctly count events in which two clusters touch at the corner.	81
3.9	Possible wire hit patterns in a drift chamber plane pair, and the corresponding hit counts returned by the fruit.	83
4.1	Categories of SODs from a charged track. Open circles represent sense wires of the plane pair, while the dashed lines indicate the inferred trajectories of the drift electrons.	91
4.2	Distribution of SODs from all plane pairs in 1999 $\pi^+\pi^-\pi_D^0$ data. Not shown are isolated hits (which were assigned a SOD of 0 m). Note that the mean of the distribution is consistent with the drift chamber sense wire spacing of 6.35 mm. SODs within 1 mm of this nominal value were considered “good”. Tracks were allowed to contain at most two bad SODs, or one bad SOD and one isolated hit.	92
4.3	A situation in which the overlap clustering correction is necessary. Two clusters (their calculated positions circled) share blocks between them. Block energies are shown in GeV. Blocks with their HCC bit on are denoted by heavy black boxes.	96
4.4	A situation requiring the neighbor clustering correction. The clusters (calculated positions circled) do not share any blocks, so the overlap correction does not apply. However, the bottom cluster spills into the 7×7 grid used to calculate the raw energy of the top cluster. Block energies are shown in GeV. Blocks with their HCC bit on are denoted by heavy black boxes.	98
4.5	An example of a cluster requiring the missing block correction, due to energy lost in a calorimeter beam hole (the white, vacant region). Block energies are shown in GeV.	99
4.6	A situation requiring the sneaky energy cluster correction. Energy from the left cluster sneaks across the beam hole to appear underneath the right cluster, and vice-versa. Block energies are shown in GeV.	100

- 4.7 Mean E/p vs. p of calibration K_{e3} electrons in small blocks, after all clustering corrections except the linearity fudge. Residual non-linearities cause the mean to drift over the range of p . The linearity fudge uses this distribution to cancel out these residual effects. . . . 102
- 4.8 Calorimeter view of a $K_L \rightarrow e^+e^-\mu^+\mu^-$ event after reconstruction. Four clusters are circled: two hardware clusters from the electrons, and two software clusters from the muons. The vertex position (diamond within the right beam hole) and upstream/downstream track segments (lines) are projected to the calorimeter face. Note that in this particular event, all tracks are matched to clusters. 105
- 4.9 y view of the same $K_L \rightarrow e^+e^-\mu^+\mu^-$ event after reconstruction. Scales are given in meters. Detector elements are labeled in Figure 2.3. The upstream and downstream track segments pass through hit pairs in the drift chambers, and clearly meet at the plane of the analysis magnet. The four-track vertex is located at 122 m in z . Note that the muon tracks (left) point to hits in the muon banks at the far downstream end of the detector. The electron tracks (right) end at the calorimeter, where the electrons lose all of their energy to EM showers in the CsI. 106
- 5.1 Examples of $\mathcal{O}(\alpha^3)$ radiative corrections to $K_L \rightarrow e^+e^-\mu^+\mu^-$ that are included in the KTeV monte carlo. 110
- 5.2 Examples of DC maps, used to correct residual tracking discrepancies between data and monte carlo. These particular maps were applied to the x view of DC1 during the generation of 1999 monte carlo. The top map served to tune the high SOD probability distribution, while the bottom corrected the isolated hit probability distribution. Note the need for additional high SOD correction in the beam region, caused mainly from radiation damage to the sense wires in that area of the chamber. 120
- 5.3 1997 $\pi^+\pi^-\pi_D^0$ track illumination at the x view of DC1. x position is measured in meters, and the dots represent data. In the plots on the left, the monte carlo (histogram) was generated without the use of DC maps, which correct for residual data/MC discrepancies in high SOD and isolated hit inefficiency distributions. A 1.5σ slope is evident in the ratio. This slope disappears when the maps are turned on (plots on the right). 121
- 5.4 Ratio of energy in the in-time (top), in-time + 1 (middle), and in-time + 2 (bottom) slices to the total energy in all slices for small seed blocks in the “medium” energy range of the calibration K_{e3} dataset. Most of the energy from real K_{e3} electrons is expected in the in-time slice, meaning electrons outside the vertical lines are from out-of-time accidental activity. Removing these events guarantees a pure calibration sample for E_{total} and HCC threshold extraction. . . . 126

- 6.1 Fast vertexing results for data on a single split tape. Plotted is the z position of the found vertex, in meters downstream of the target. Note the dominant peak from interactions at the vacuum window. The bumps at 138 m and 152 m correspond to interactions within RC7 and RC9, respectively. 131
- 6.2 Severe bad spill bits for all 22.7 million events passing the crunch stage. Bits from the 1997 data are plotted in the open histogram, while bits from the 1999 data are plotted in the cross-hatched histogram. Bit meanings are listed in Table 6.1. Arrows point to the cut bits, corresponding to about 7% of the dataset. 134
- 6.3 MU2 counter illumination (number of TDC hits per scintillator paddle) for fully reconstructed 1999 $\pi^+\pi^-\pi_D^0$ events after all analysis cuts. The counter map was shown in Figure 2.9. Despite the fact that this plane was in veto, over 1% of events have in-time hits in the MU2 TDCs, illustrating the need for additional trigger verification. 135
- 6.4 Distribution of the ratio E/p for all tracks in the 1999 data. The peak at $E/p = 1$ comes mostly from electrons, while the peak at 0 is from MIPs (muons and pions), as well as from particles that traveled down a beam hole, depositing no energy in the calorimeter. The hump in between is caused by hadronic interactions from pions in the CsI. Tracks with $E/p > 1$ result from calorimeter and spectrometer resolution effects, or from accidental activity on top of a real electron cluster. 137
- 6.5 (a) $\pi^+\pi^-\pi_D^0$ and (b) $e^+e^-\mu^+\mu^-$ invariant mass distributions for all data after track/cluster reconstruction and particle identification. Additional cuts have not yet been made. 139
- 6.6 Definition of the variable P_t^2 . P_t is the component of the reconstructed momentum perpendicular to a line drawn from the target to the vertex. For a well reconstructed event, this variable should be near 0. 140
- 6.7 Definition of the variable $pp0kine$. This kinematic quantity is defined as the square of the longitudinal momentum of the reconstructed π^0 in the frame in which the longitudinal momentum component of the $\pi^+\pi^-$ pair is 0. $pp0kine$ should be positive for true $K_L \rightarrow \pi^+\pi^-\pi_D^0$ decays. 142
- 6.8 Data (dots) vs. monte carlo (histogram) comparisons for 1999 $K_L \rightarrow \pi^+\pi^-\pi_D^0$ events. The histograms are normalized to each other by area. (a) and (b) Overlays in the kinematic parameters P_t^2 and $pp0kine$. The data excesses outside the cuts at $0.00025 \text{ GeV}^2/c^2$ and $-0.005 \text{ GeV}^2/c^2$, respectively, result from background and misreconstructed events. (c) Overlay in e^+e^- hit separation at DC1. The large disagreement at low separation is caused by a combination of photon conversion background and resolution effects. This region is removed from the dataset by requiring the separation to be greater than 2 mm. (d) Photon energy, as measured by the CsI calorimeter. The disagreement at low energies is caused by a combination of calibration, resolution, accidental activity, and bremsstrahlung effects. The cut at 2 GeV removes this region from the dataset. 143

- 6.9 Reconstructed M_K and M_{π^0} mass distributions for 1997 ((a) & (b)) and 1999 ((c) & (d)) after all cuts. Histogram is monte carlo, dots are data. 145
- 6.10 1997 $K_L \rightarrow \pi^+\pi^-\pi_D^0$ vertex z (left) and kaon momentum (right) distributions after all cuts. Histogram is monte carlo, dots are data. While there is excellent agreement in vertex z position, a small slope is seen in the data/monte carlo ratio in kaon momentum. This slope will be discussed further in Section 6.8.1. 146
- 6.11 1999 $K_L \rightarrow \pi^+\pi^-\pi_D^0$ vertex z (left) and kaon momentum (right) distributions after all cuts. Histogram is monte carlo, dots are data. While there is good agreement in kaon momentum, an excess of data at the downstream end of vertex z leads to a slope in the data/monte carlo ratio. This slope will be discussed further in Section 6.8.1. . . 147
- 6.12 1999 $K_L \rightarrow \pi^+\pi^-\pi_D^0$ P_t^2 (a) and vertex χ^2 (b) distributions after all other cuts have been made. Histogram is monte carlo, dots are data. The lines on the overlays show the positions of the cuts on these variables. The disagreement seen here will be discussed further in Section 6.8.1. 149
- 6.13 $K_L \rightarrow \pi^+\pi^-\pi_D^0$ electron E/p for (a) Winter, (b) Summer, and (c) 1999 after all other cuts. Dots are data, histogram is monte carlo. The disagreement seen here will be discussed further in Section 6.8.1. 150
- 6.14 Track-cluster matching inefficiency for charged pions from $K_L \rightarrow \pi^+\pi^-\pi^0$ events, as a function of track momentum in GeV/c. The distribution is approximated by a third-order polynomial. 151
- 6.15 (a) Vertex χ^2 distributions of signal monte carlo (histogram) and data (dots). The line shows the cut at 40. (b) Signal invariant mass distribution after the vertex χ^2 and extra hardware cluster cuts. Compare to Figure 6.5(b). 154
- 6.16 $e^+e^-\mu^+\mu^-$ monte carlo distributions for (a) muon track offset at the analysis magnet, (b) muon cluster energy in the CsI, and (c) muon track momentum. Entries in the 0 bin in (b) are from beam hole muons. The lines show the cuts at 2 mm, 1.5 GeV, and 7 GeV/c, respectively. 155
- 6.17 Electron hit separation at DC1 (in millimeters) vs. M_{ee} (in GeV/c²) for (a) conversion monte carlo, (b) signal monte carlo, and (c) signal data. Plots are logarithmic in z . Cutting out the box rejects 99% of the conversions, leaving approximately 0.7 expected background events from this source. 157
- 6.18 Invariant mass distributions of 1997 $\pi^+\pi^-\pi_D^0$ data (dots) reconstructed in the signal trigger. Overlaid are monte carlo histograms generated with various pion decay/punchthrough ratios. Upper left plot shows 0% decay/100% punchthrough MC, upper right is 25/75, middle left is 60/40, middle right is 65/35, lower left is 70/30, and lower right shows 100% decay/0% punchthrough MC. The best agreement seems to come from 60% decay/40% punchthrough (middle left). 159

- 6.19 Invariant mass distributions of 1999 $\pi^+\pi^-\pi_D^0$ data (dots) reconstructed in the signal trigger. Overlaid are monte carlo histograms generated with various pion decay/punchthrough ratios. Upper left shows 50% decay/50% punchthrough, upper right is 55/45, middle left is 60/40, middle right is 65/35, lower left is 70/30, and lower right is 75/25. The best agreement seems to come from 55% decay/45% punchthrough (upper right). 160
- 6.20 Distributions of pion track momenta from 1997 (top) and 1999 (bottom) $K_L \rightarrow \pi^+\pi^-\pi_D^0$ events. All cuts have been made except for $p_\pi > 7$ GeV/c. Note that the 1999 spectrum is slightly softer than in 1997, due to the lower P_t kick from the analysis magnet. This would be expected to contribute to a larger pion decay probability, which is not seen in the fits. 161
- 6.21 Invariant mass comparison between 80% pion decay, 20% pion punchthrough (dots) and 60% decay, 40% punchthrough (histogram) monte carlo after reconstruction as signal. The shape of the reconstructed invariant mass distribution, and hence the background estimate, is quite insensitive to the choice of decay/punchthrough ratio. 162
- 6.22 Invariant mass distributions of (a) $\pi^+\pi^-e^+e^-$ and $\pi^+\pi^-\gamma$, and (b) $\pi^+\pi^-\pi_D^0$ pion decay/punchthrough monte carlo, when reconstructed as signal. Overlay of the sum of these datasets with data is shown in (c). The signal region has been blacked out to properly normalize the distributions. Approximately 0.06 background events are expected from decay/punchthrough events. 164
- 6.23 P_t^2 vs. invariant mass of all $e^\pm e^\pm \mu^\mp \mu^\mp$ events. The box shows the location of the signal region in the plane. Extrapolating into the signal region yields an estimate of 0.08 background events from double $K_{\mu 3}$ decays. 165
- 6.24 (a) Signal P_t^2 vs. invariant mass for all signal data after all other cuts. 132 events are in the box, which denotes the signal region. (b) Projection onto the invariant mass axis after the cut on P_t^2 . Events from 1997 are plotted in the open histogram, while events from 1999 are plotted in the cross-hatched histogram. The lines mark the boundaries of the signal region. 167
- 6.25 (a) Invariant mass comparison between the 132 signal data events (dots) and signal monte carlo (histogram). (b) Comparison in P_t^2 . The data/monte carlo agreement is good in both variables. 168
- 6.26 1999 $K_L \rightarrow \pi^+\pi^-\pi_D^0$ vertex z (left) and kaon momentum (right) distributions after all cuts, with the cut on electron separation at DC1 moved from 2 to 25 mm. Histogram is monte carlo, dots are data. Note that the disagreement at the downstream end in vertex z is much smaller than in Figure 6.11, and the slope of the data/monte carlo ratio is now consistent with 0. 171

6.27	$\pi^+\pi^-\pi_D^0$ electron E/p for (a) Winter, (b) Summer, and (c) 1999 after “nudging” monte carlo cluster energies during clustering at the crunch stage. Data is dots, histogram is monte carlo. Compare to Figure 6.13. Using the new monte carlo cluster energies led to a total shift in the branching ratio of only 0.1%.	174
6.28	Shifts in percent of the branching ratio as the cuts on P_t^2 , vertex χ^2 , and p_{MIP} are varied. Note the walk in vertex χ^2 . Deviations in the other two variables appear to be statistical.	175
6.29	Shift in the total kaon flux as the cut on vertex χ^2 is varied. This provides a better estimate of the systematic effect than the walk in the branching ratio, due to the greater statistical power. An error of 1.10% is used from the deviation at a vertex χ^2 cut at 25.	176
6.30	Data (dots) and monte carlo (histogram) comparisons for (a) number of hardware clusters, and (b) number of y view tracks for reconstructed 1999 $\pi^+\pi^-\pi_D^0$ events. Both of these variables are dependent on the level of accidental activity. The agreement seen in these overlays demonstrates that the level of simulated accidental activity in the monte carlo is fairly well understood.	178
6.31	Distributions of the variable $xdiff$ (in meters) for 1997 long muon run data (dots) and monte carlo (histogram), plotted in 5 GeV/c muon momentum bins.	180
6.32	Fits to $xdiff$ distributions of 1997 long muon run data. The distributions are well approximated by the sum of two gaussians. Reconstructing signal monte carlo with these fit parameters led to a total shift in the branching ratio of 0.55%.	181
6.33	Results of the four-track trigger systematic study. The top plot shows the calculated kaon flux (in arbitrary units), totaled over the three time periods used, as four-track trigger sources were individually imposed on a two-track normalization sample. The rightmost point represents a complete simulation of the four-track trigger. The bottom plot shows the difference, in percent, between each of these fluxes and the flux calculated from the default two-track normalization, taking correlations into account. There is a difference of 1.81% between the fluxes measured from the two-track and simulated four-track triggers.	183
7.1	M_{ee} and $M_{\mu\mu}$ distributions for the 132 accepted $K_L \longrightarrow e^+e^-\mu^+\mu^-$ signal events (dots) overlaid on MC (histogram) generated with $\alpha = \beta = 0$. The disagreement seen in the $M_{\mu\mu}$ overlay suggests the existence of a non-trivial form factor.	188
7.2	Top: Magnitude of the DIP form factor (Equation 1.32) as a function of M_{ee} and $M_{\mu\mu}$ for $\alpha = -1.54$, $\beta = +2$. Bottom: Projections of the two-dimensional distribution onto each mass axis, as well as the relative contributions to the total from each term. Note that the total is almost completely insensitive to the β term – a fact even further magnified by the fact that the form factor is squared in the calculation of the $K_L \longrightarrow e^+e^-\mu^+\mu^-$ rate (Equation 1.53).	190

7.3	Comparison between monte carlo M_{ee} and $M_{\mu\mu}$ distributions at generation for $\beta = +2$ (dots) and $\beta = +10$ (histogram). $\alpha = -1.54$ for both MC sets. Note that the minimal amount of sensitivity to β in these plots arises only in the high end of the M_{ee} spectrum, where the statistics are the lowest. There is no significant $M_{\mu\mu}$ sensitivity to β	191
7.4	Comparison between monte carlo M_{ee} and $M_{\mu\mu}$ distributions at generation for $\alpha = -1.54$ (dots) and $\alpha = -5$ (histogram). $\beta = 0$ for both MC sets. Note the sensitivity of the $M_{\mu\mu}$ spectrum to α . The shape of the M_{ee} distribution is essentially unaffected by changes in α	192
7.5	(a) Log-likelihood as a function of α_i in the $K_L \rightarrow e^+e^-\mu^+\mu^-$ shape analysis. The parameter β is fixed at 0. (b) Close-up of the maximum of the log-likelihood distribution, which corresponds to a fit value of $\alpha_{shape} = -4.53^{+1.81}_{-2.70}$	194
7.6	Log-likelihood as a function of β_i in the $K_L \rightarrow e^+e^-\mu^+\mu^-$ shape analysis. The parameter α is fixed at -4.53 . The maximum of the log-likelihood distribution corresponds to a fit value of $\beta_{shape} = +12.3 \pm 77.1$	197
7.7	Theoretical prediction of the ratio $\mathcal{B}(K_L \rightarrow e^+e^-\mu^+\mu^-) / \mathcal{B}(K_L \rightarrow \gamma\gamma)$ as a function of the DIP parameter α . Also shown is the extrapolated value of α_{BR} , based on the measured $K_L \rightarrow e^+e^-\mu^+\mu^-$ branching ratio of $(2.63 \pm 0.23_{stat} \pm 0.13_{syst}) \times 10^{-9}$, and assuming $\beta = 0$ (dots). The error bars on the experimental result (dotted lines) include the uncertainty on $\mathcal{B}(K_L \rightarrow \gamma\gamma)$. This method provides an estimate for α_{BR} of -1.43 ± 0.39 , consistent with earlier measurements. The triangles show the effect of setting β to $+2$	198
7.8	Distribution of 132 signal data events (dots) and monte carlo with no CP -violation (histogram) in $\sin\phi\cos\phi$. The measured asymmetry in the data is $\mathcal{A} = -5.3 \pm 12.3\%$, providing a 90% C.L. limit of $ \mathcal{A} < 25.5\%$	200
8.1	Summary plot of theoretical calculations for the $K_L \rightarrow e^+e^-\mu^+\mu^-$ branching ratio. This result is consistent with the prediction from Vector Meson Dominance. If the DIP form factor is assumed, the measured branching ratio is most consistent with the value $\alpha = -1.43 \pm 0.39$, if β is near 0.	203
8.2	Summary plot of all measurements made of the DIP parameter α to date. The world average comes to $\overline{\alpha} = -1.53 \pm 0.10$, almost identical to the total result from the $K_L \rightarrow \mu^+\mu^-\gamma$ analysis.	204

LIST OF TABLES

2.1	z positions and dimensions of KTeV spectrometer elements.	52
3.1	Trigger accept rates for a typical beam intensity (8×10^{12} protons on target per spill) during the 1999 run. At this intensity the trigger was live around 65% of the time.	89
5.1	Type, purpose, and number (by run) of monte carlo events generated for this analysis.	129
6.1	Bad spill cut summary. Only “severe” bits were cut on.	133
6.2	Kaon flux results based on $K_L \rightarrow \pi^+ \pi^- \pi_D^0$ normalization events. N , P , A , and \mathcal{F} are defined in the text. Recall that the normalization mode prescale was changed from 10 to 20 about 5% into the 1999 run. Errors are purely statistical.	152
6.3	Summary of all cuts applied to normalization and signal samples.	166
6.4	Backgrounds to $K_L \rightarrow e^+ e^- \mu^+ \mu^-$ for 1997, 1999, and total.	168
6.5	$K_L \rightarrow e^+ e^- \mu^+ \mu^-$ branching ratios for 1997, 1999, and all data. A and \mathcal{B} are calculated using a radiative cutoff at $x_{e^+ e^- \mu^+ \mu^-} > 0.95$. The branching ratio for the complete dataset was calculated by taking the weighted average of the 1997 and 1999 branching ratios. Errors at this stage are purely statistical.	170
6.6	Total kaon flux, calculated in equally populated bins of vertex z . As might be expected from Figure 6.11, the largest variation from the nominal flux occurs at the most downstream bin.	170
6.7	Total kaon flux, calculated in equally populated bins of kaon momentum. No significant effect is evident.	173
6.8	Changes in the branching ratio due to various populations of accidental overlays.	177
6.9	Systematic errors on the $K_L \rightarrow e^+ e^- \mu^+ \mu^-$ branching ratio.	185
7.1	Fitted value of α_{shape} versus the number of bins used in the M_{ee} and $M_{\mu\mu}$ plots. The average of these entries, -4.50 , is almost identical to the default value of -4.53	195

ABSTRACT

The rare decay $K_L \longrightarrow e^+e^-\mu^+\mu^-$ offers the best opportunity to study the dynamics behind the $K_L \gamma^* \gamma^*$ vertex. This information is needed to subtract the long-distance contribution to the decay $K_L \longrightarrow \mu^+\mu^-$, allowing for a calculation of ρ , the real part of the CKM matrix element V_{td} . Analysis of data from the 1997 and 1999 runs of the KTeV experiment has isolated a sample of 132 $K_L \longrightarrow e^+e^-\mu^+\mu^-$ events, with an estimated background of 0.8 events. Based on this data, the branching ratio for this mode was calculated to be $\mathcal{B}(K_L \longrightarrow e^+e^-\mu^+\mu^-) = (2.63 \pm 0.23_{\text{stat}} \pm 0.13_{\text{syst}}) \times 10^{-9}$, where a radiative cutoff has been imposed at $x_{e^+e^-\mu^+\mu^-} > 0.95$. The D'Ambrosio, Isidori, and Portolés $K_L \gamma^* \gamma^*$ linear form factor parameter α was measured to be -1.52 ± 0.38 , consistent with previous experimental results. Little sensitivity was found to the quadratic parameter β , resulting in a measured value of $+12.3 \pm 77.1$. The asymmetry in the angular distribution of the $K_L \longrightarrow e^+e^-\mu^+\mu^-$ decay products was measured to be $|\mathcal{A}| < 25.5\%$ with 90% confidence, constraining the size of any CP -violating contribution to the $K_L \gamma^* \gamma^*$ interaction. Finally, a 90% C.L. upper limit of 4.12×10^{-11} was placed on the branching ratio of the lepton flavor-violating mode $K_L \longrightarrow e^\pm e^\pm \mu^\mp \mu^\mp$, a full factor of 3 improvement over the current limit.

CHAPTER 1

Theoretical Motivation for the Study of $K_L \longrightarrow e^+e^-\mu^+\mu^-$

1.1 Quark Mixing and the CKM Matrix

A distinguishing characteristic of the weak interaction is that quark flavor eigenstates are not equivalent to weak eigenstates. Under the weak force, transitions are allowed both within and between quark generations. An example of such a transition is diagrammed in Figure 1.1. To accommodate this phenomenon in the Standard Model of particle physics, a matrix was constructed to contain the coupling strength associated with every possible transition. The unitary *CKM matrix* (named for Cabibbo, Kobayashi, and Maskawa) assumes three quark generations,




Figure 1.1: A flavor-changing transition mediated by the weak force. The top quark (t) changes into a down quark (d) by coupling to a W gauge boson. The probabilities of such transitions are described in the Standard Model by the CKM matrix (in this example, the matrix element V_{td}).

and is defined as follows (Kobayashi and Maskawa 1973):

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = V \times \begin{pmatrix} d \\ s \\ b \end{pmatrix}, \quad V = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix}. \quad (1.1)$$

Here d , s , and b represent the flavor eigenstates of the down, strange, and bottom quarks. These states take part in strong interactions. d' , s' , and b' are the corresponding weak eigenstates. The CKM matrix V acts to connect the two representations, describing the probability of weak transitions between flavor eigenstates. For example, the matrix element V_{td} gives the relative likelihood of the $t \rightarrow d + W$ vertex shown in Figure 1.1.

Because transitions within generations are the most likely to occur, the largest elements in the CKM matrix are expected to be close to unity and on the diagonal, while the smallest are expected to be in the opposite corners (describing the most distant transitions between the first and third generations). Additionally, the unitarity condition implies that the CKM matrix elements are not completely independent. As an example, requiring unitarity provides the relation

$$V_{ud}V_{ub}^* + V_{cd}V_{cb}^* + V_{td}V_{tb}^* = 0. \quad (1.2)$$

An alternate parameterization for the CKM matrix, proposed by Wolfenstein (Wolfenstein 1983), reveals that the matrix can be reduced to four independent parameters: three real numbers, and one complex phase:

$$V = \begin{pmatrix} 1 - \lambda^2/2 & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1 - \lambda^2/2 & A\lambda^2 \\ A\lambda^3(1 - \rho - i\eta) & -A\lambda^3 & 1 \end{pmatrix} + \mathcal{O}(\lambda^4) + \dots \quad (1.3)$$

The elements are expanded in powers of λ , which is defined to be equivalent to the matrix element V_{us} . The complex phase has magnitude η , and is confined to the

corners of the matrix in this parameterization. This phase empirically describes the phenomenon of CP -violation in the Standard Model, which will be discussed in the following section.

The parameters λ and A have been well measured experimentally; their values are known to within a few percent (Groom et al. 2000a). This means that the majority of elements in the CKM matrix are well determined. However, values for ρ and η are still relatively unknown, translating into large uncertainties on V_{td} and V_{ub} . Estimates vary, but a conservative range for $|V_{td}|$ is $(6.7-9.3) \times 10^{-3}$ (Buras 2001). This uncertainty is expected to decline over the next few years, as new results from the B sector (notably $B_d^0 - \bar{B}_d^0$, $\frac{\Delta M_{B_s}}{\Delta M_{B_d}}$, and $b \rightarrow s\gamma$) and the charged kaon sector ($K^+ \rightarrow \pi^+ \nu \bar{\nu}$) become available.

Of course, the experimental motivation is not simply to determine the CKM parameters, but to overconstrain them. Measuring these parameters using a variety of decay modes from different sectors provides an excellent test of the Standard Model. An estimate of ρ from the neutral kaon sector would therefore be useful for this purpose. The “golden mode” for this measurement is $K_L \rightarrow \mu^+ \mu^-$. In order to understand exactly what a “ K_L ” is, however, a review of the neutral kaon system is in order.

1.2 Review of the Neutral Kaon System

The strong eigenstates of the neutral kaon, K^0 (quark content $d\bar{s}$) and \bar{K}^0 ($\bar{d}s$), are not independent under all interactions. Because strangeness S is not necessarily a conserved quantum number under the weak force, a K^0 ($S = +1$) can be transformed into its $S = -1$ antiparticle (and vice versa) via a second-order, $\Delta S = 2$




Figure 1.2: One of several diagrams contributing to oscillations between K^0 and its antiparticle, \overline{K}^0 . The oscillations are made possible by second order, $\Delta S = 2$ weak interactions.

weak transition. One diagram that contributes to these oscillations is shown in Figure 1.2.

The presence of such oscillations suggests that the strong eigenstates are not the best way to describe neutral kaons within a weak interaction Hamiltonian. As an alternative, Gell-Mann and Pais suggested constructing eigenstates of the charge conjugation operator C (Gell-Mann and Pais 1955):

$$|K_1\rangle = \frac{1}{\sqrt{2}} (|K^0\rangle + |\overline{K}^0\rangle) \quad (1.4)$$

$$|K_2\rangle = \frac{1}{\sqrt{2}} (|K^0\rangle - |\overline{K}^0\rangle). \quad (1.5)$$

Since $C|K^0\rangle = |\overline{K}^0\rangle$ and $C|\overline{K}^0\rangle = |K^0\rangle$ it can easily be verified that

$$C|K_1\rangle = |K_1\rangle, \quad C|K_2\rangle = -|K_2\rangle. \quad (1.6)$$

Furthermore, K_1 and K_2 are eigenstates of the combination of the charge conjugation and parity operators, CP :

$$CP|K_1\rangle = |K_1\rangle, \quad CP|K_2\rangle = -|K_2\rangle. \quad (1.7)$$

Because of these properties, K_1 is known as the *CP-even* state, while K_2 is called the *CP-odd* state. It can be shown (Griffiths 1987) that when the pseudoscalar

kaon decays to two pions, the final state is CP -even, while the decay to three pions produces a CP -odd final state. Therefore, if CP were a perfect symmetry, the decays

$$K_1 \rightarrow 2\pi, \quad K_2 \rightarrow 3\pi \quad (1.8)$$

would be allowed, while

$$K_1 \rightarrow 3\pi, \quad K_2 \rightarrow 2\pi \quad (1.9)$$

would be forbidden. Because there is much more phase space available to the two pion final state of the K_1 , its lifetime is expected to be much shorter than that of the K_2 .

It turns out, however, that the forbidden decay $K_2 \rightarrow 2\pi$ is observed, at a level of about 1 part in 500 (Christenson et al. 1964). This phenomenon of CP -violation means that not even the K_1 and K_2 states are perfect descriptions of the neutral kaon system in weak interactions. Instead, the weak eigenstates are written as

$$\begin{aligned} |K_S\rangle &= \frac{1}{\sqrt{1+|\epsilon|^2}} (|K_1\rangle + \epsilon|K_2\rangle) \\ |K_L\rangle &= \frac{1}{\sqrt{1+|\epsilon|^2}} (|K_2\rangle + \epsilon|K_1\rangle). \end{aligned} \quad (1.10)$$

K_L represents the long-lived state of the neutral kaon: mostly K_2 , with a small admixture of K_1 that provides for the CP -violating decays observed experimentally. Conversely, K_S is the short-lived state: mostly K_1 with a small K_2 component. The states are aptly named – with a mean lifetime of approximately 50 ns, the K_L takes almost 600 times longer to decay than its counterpart (Groom et al. 2000b). K_S and K_L appear to be the true eigenstates of the weak interaction.

The size of the “contamination” leading to CP -violation is given in Equation 1.10 by the parameter ϵ , approximately 2×10^{-3} . Calculation of the $\Delta S = 2$ weak transition amplitude reveals that ϵ is proportional to the imaginary phase η from Equation 1.3 (Belusevic 1999).

The astounding significance of CP -violation can be revealed by rewriting Equation 1.10. Doing so for the K_L provides

$$|K_L\rangle \propto (1 + \epsilon)|K^0\rangle - (1 - \epsilon)|\overline{K^0}\rangle. \quad (1.11)$$

In other words, matter and anti-matter are treated inherently differently! Note that there is nothing in the Standard Model to describe *why* CP -violation exists in the first place – the phenomenon is only *empirically* described by an imaginary phase in the CKM matrix. The true source of CP -violation remains one of the biggest questions in the field of high energy physics.

1.3 The Physics of $K_L \longrightarrow \mu^+\mu^-$

As mentioned in Section 1.1, the decay $K_L \longrightarrow \mu^+\mu^-$ offers the best hope of measuring ρ within the neutral kaon sector. The decay can proceed via several “short-distance” weak interactions, pictured in Figure 1.3. Because of its disproportionately large mass, the top quark dominates the loops shown in these diagrams (Buchalla and Buras 1994). This provides a coupling between top and down quarks, making the $K_L \longrightarrow \mu^+\mu^-$ decay amplitude sensitive to ρ (the real part of V_{td}). The diagrams shown in Figure 1.3 are easily calculable, implying that the magnitude of ρ can be linked directly to the $K_L \longrightarrow \mu^+\mu^-$ amplitude (Wolfenstein 1983).

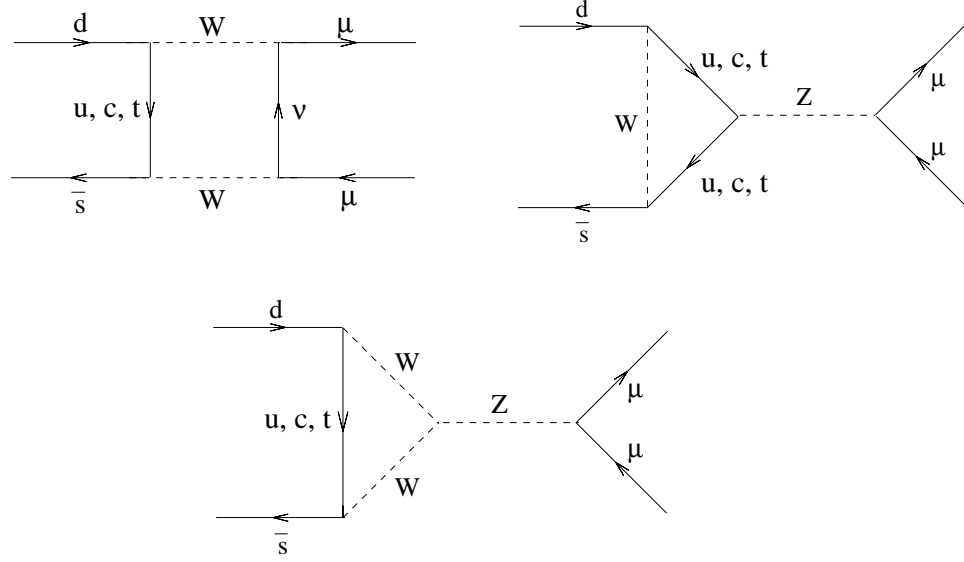


Figure 1.3: Short-distance contributions to the $K_L \longrightarrow \mu^+ \mu^-$ amplitude. The quark loops in these diagrams are dominated by top, making this mode sensitive to the CKM matrix element V_{td} (or ρ , in the Wolfenstein parameterization).

Unfortunately, this measurement is complicated by a “long-distance” contribution to the amplitude. This contribution (pictured in Figure 1.4) is mediated by low-energy electromagnetic coupling between the K_L and virtual photons, and must be subtracted off from the total amplitude before any CKM information can be obtained.

1.3.1 Connecting $K_L \longrightarrow \mu^+ \mu^-$ to the CKM matrix

To perform the calculation described above, the $K_L \longrightarrow \mu^+ \mu^-$ branching ratio is decomposed into two terms, as follows:

$$\mathcal{B}(K_L \longrightarrow \mu^+ \mu^-) = |A|^2 = |\Re A|^2 + |\Im A|^2. \quad (1.12)$$




Figure 1.4: Long-distance contribution to the $K_L \rightarrow \mu^+ \mu^-$ amplitude. This diagram must be subtracted from the total $K_L \rightarrow \mu^+ \mu^-$ amplitude in order to extract the short-distance contribution, and with it, ρ . Unfortunately, the form of the $K_L \gamma^* \gamma^*$ vertex is not well understood, making this calculation difficult.

where \Re and \Im respectively designate the real and imaginary parts of the amplitude A . The magnitude of the imaginary part, which corresponds to the case where the two photons are real (on-shell), can be calculated in closed form using QED (Sehgal 1969):

$$|\Im A|^2 = \frac{\alpha_{EM}^2 M_\mu^2}{2M_K^2 \beta_\mu} \left[\ln \frac{1 - \beta_\mu}{1 + \beta_\mu} \right]^2 \times \mathcal{B}(K_L \rightarrow \gamma\gamma), \quad (1.13)$$

where M_μ is the muon mass, M_K is the kaon mass, and

$$\beta_\mu = \sqrt{1 - \frac{4M_\mu^2}{M_K^2}}. \quad (1.14)$$

The most recent measurement of the $K_L \rightarrow \gamma\gamma$ branching ratio places $|\Im A|^2$ at a value of $(7.07 \pm 0.18) \times 10^{-9}$ (Groom et al. 2000b). This result almost completely saturates the total measured $K_L \rightarrow \mu^+ \mu^-$ branching ratio of $(7.18 \pm 0.17) \times 10^{-9}$ (Ambrose et al. 2000). The small difference between the two numbers must be due to the real part of the total amplitude. Taking correlations into account, the difference was calculated to be $(0.11 \pm 0.18) \times 10^{-9}$. This result for $|\Re A|^2$ can be

translated into a 90% confidence level limit, excluding the unphysical region below 0. Doing so yields an upper limit for the square of the real part of the $K_L \longrightarrow \mu^+ \mu^-$ amplitude of 0.37×10^{-9} (Ambrose et al. 2000).

This tiny, residual contribution can be further split up into long- and short-distance terms:

$$\Re A = \Re A_{long} + \Re A_{short}. \quad (1.15)$$

The short-distance term has been calculated by summing amplitudes from Z-penguin and box diagrams, some of which are shown in Figure 1.3. This calculation results in the expression

$$|\Re A_{short}|^2 = \frac{\alpha_{EM}^2 \mathcal{B}(K^+ \longrightarrow \mu^+ \nu) \tau(K_L)}{V_{us}^2 \pi^2 \sin^4 \theta_W} \frac{\tau(K_L)}{\tau(K^+)} [\Re V_{cs}^* V_{cd} Y_{NL} + \Re V_{ts}^* V_{td} Y(x_t)]^2. \quad (1.16)$$

The variable x_t is defined as M_t^2/M_W^2 (M_t is the mass of the top quark), τ is the mean lifetime of the particle at rest, and θ_W is the weak mixing angle (the Standard Model parameter that quantifies the breaking of the unified electroweak field into its electromagnetic and weak components). The function $Y(x)$ describes the effects of both single-loop electroweak and double-loop QCD top quark corrections, while Y_{NL} accounts for corrections due to non-negligible contributions from charm (Buchalla and Buras 1994), (Buras and Fleischer 1998). Note the appearance of V_{td} in the above equation.

In order to proceed, it is necessary to expand Equation 1.3 to $\mathcal{O}(\lambda^5)$. This expansion yields (Buras and Fleischer 1998)

$$V_{td} = A\lambda^3 (1 - \bar{\rho} - i\bar{\eta}), \quad (1.17)$$

with the definitions

$$\bar{\rho} = \rho \left(1 - \frac{\lambda^2}{2}\right), \quad \bar{\eta} = \eta \left(1 - \frac{\lambda^2}{2}\right). \quad (1.18)$$

After a bit of algebraic manipulation, along with evaluation of some of the physical constants, the following expression can be obtained:

$$|\Re A_{short}|^2 = (1.68 \times 10^{-9}) A^4 Y^2(x_t) \frac{(\bar{\rho}_0 - \rho)^2}{\sigma}, \quad (1.19)$$

where

$$\bar{\rho}_0 = 1 + \frac{Y_{NL}}{\lambda^4 A^2 Y(x_t)}, \quad \sigma = \left(\frac{1}{1 - \frac{\lambda^2}{2}}\right)^2. \quad (1.20)$$

For $150 \text{ GeV}/c^2 < M_t < 190 \text{ GeV}/c^2$, Equation 1.19 can be rewritten as the approximate relation (Buras and Fleischer 1998)

$$|\Re A_{short}|^2 = 0.9 \times 10^{-9} (1.2 - \bar{\rho})^2 \left[\frac{M_t}{170 \text{ GeV}/c^2} \right]^{3.1} \left[\frac{|V_{cb}|}{0.040} \right]^4. \quad (1.21)$$

Solving for $\bar{\rho}$, and making a substitution from Equation 1.15, this finally yields

$$\bar{\rho} > 1.2 - \max \left(\frac{|\Re A| \pm |\Re A_{long}|}{3 \times 10^{-5}} \left[\frac{M_t}{170 \text{ GeV}/c^2} \right]^{-1.55} \left[\frac{|V_{cb}|}{0.040} \right]^{-2} \right), \quad (1.22)$$

where the sign choice is made to set the stricter lower limit. The only remaining unknown in Equation 1.22 is $|\Re A_{long}|$, the size of the long-distance contribution to $K_L \longrightarrow \mu^+ \mu^-$ shown in Figure 1.4.

1.3.2 Calculating the long-distance contribution to $K_L \longrightarrow \mu^+ \mu^-$

Calculation of the long-distance contribution to $K_L \longrightarrow \mu^+ \mu^-$ starts by writing down a general expression for the vertex amplitude (Ecker et al. 1988):

$$A(K_L \longrightarrow \gamma^*(q_1, \epsilon_1) \gamma^*(q_2, \epsilon_2)) = i \epsilon_{\mu\nu\rho\sigma} \epsilon_1^\mu \epsilon_2^\nu q_1^\rho q_2^\sigma F(q_1^2, q_2^2). \quad (1.23)$$

Here, q_1^ρ and q_2^σ represent the components of the momentum four-vectors q_1 and q_2 for the two virtual photons, while ϵ_1^μ and ϵ_2^ν are the components of the corresponding polarization four-vectors ϵ_1 and ϵ_2 . $\varepsilon_{\mu\nu\rho\sigma}$ signifies the anti-symmetric Levi-Civita tensor operator. $F(q_1^2, q_2^2)$ is the long-distance $K_L \gamma^* \gamma^*$ form factor, required to be symmetric under exchange of q_1^2 and q_2^2 . A non-trivial calculation leads to the result

$$|\Re A_{long}|^2 = \frac{2\alpha_{EM}^2 M_\mu^2 \beta_\mu}{\pi^2 M_K^2} \mathcal{B}(K_L \rightarrow \gamma\gamma) \left| \Re R(M_K^2) \right|^2, \quad (1.24)$$

where

$$R(q^2) = \frac{2i}{\pi^2 q^2} \int d^4\ell \frac{q^2 \ell^2 - (q \cdot \ell)^2}{\ell^2 (\ell - q)^2 [(\ell - p)^2 - M_\mu^2]} \frac{F(\ell^2, (\ell - p)^2)}{F(0, 0)}, \quad (1.25)$$

and $p^2 = M_\mu^2$ (Ametller et al. 1984). The normalization factor of $F(0, 0)$ in the denominator of the integrand corresponds to the case where both photons are on-shell. This term can be related to the experimentally accessible $K_L \rightarrow \gamma\gamma$ width:

$$|F(0, 0)| = \left[\frac{64\pi\Gamma(K_L \rightarrow \gamma\gamma)}{M_K^3} \right]^{1/2}. \quad (1.26)$$

The problem has now been reduced to evaluating the integral in Equation 1.25. Before this is done, however, a decision must be made regarding the form of F . Many models have been suggested over the past 35 years; however, most suffer from theoretical uncertainties that prevent a useful calculation of $|\Re A_{long}|^2$ (and therefore ρ). Some of the most commonly cited models will be described in the following section.

1.3.3 Models of the $K_L \gamma^* \gamma^*$ Vertex

Kroll–Wada type (phase space) interactions

One of the earliest studies of pseudoscalar meson–photon–photon interactions was published by Kroll and Wada in 1955. The Kroll–Wada formula describes the dynamics responsible for internal conversion of the π^0 , providing QED predictions for the amplitudes of the decays $\pi^0 \longrightarrow e^+e^-\gamma$ and $\pi^0 \longrightarrow e^+e^-e^+e^-$ (Kroll and Wada 1955). The original theory included no momentum–dependent form factors; only phase space contributes to the total amplitude. Additionally, the authors neglected the effects of additional interference terms caused by the exchange of identical particles.

In extending the Kroll–Wada calculation to the $K_L \gamma^* \gamma^*$ vertex, Miyazaki and Takasugi included these interference effects (Miyazaki and Takasugi 1973). They assumed that the form factor carries no momentum dependence, and that no CP –violating effects exist in the interaction (in other words, K_L is equivalent to the CP –odd eigenstate K_2).

In her calculation done almost twenty years later, Uy allowed for a CP –violating contribution to the $K_L \gamma^* \gamma^*$ vertex (Uy 1991). The strengths of the CP –violating and CP –conserving terms were described by two form factors, g and h , which were assumed to be momentum–independent over the range of consideration. The Uy calculation is of interest because it is the most general of the Kroll–Wada type models.




Figure 1.5: Vector Meson Dominance model for the $K_L \gamma^* \gamma^*$ vertex. Interactions between the pseudoscalar K_L and the virtual photons γ^* are mediated by vector mesons, V .

Vector Meson Dominance

In the Vector Meson Dominance approximation, there is no direct coupling between the pseudoscalar mesons and the virtual photons. Instead, the interaction is mediated by vector mesons (the ρ , ω , and ϕ), as shown in Figure 1.5. These particles are theoretically allowed to “fluctuate” into photons because vector mesons and photons share identical quantum numbers.

Because the vector–pseudoscalar and vector–photon couplings are theoretically well understood, application of Vector Meson Dominance greatly simplifies the description of the $K_L \gamma^* \gamma^*$ interaction. The approximation leads to a momentum–dependent form factor proportional to

$$f(q_1^2, q_2^2) = \sum_i \frac{1}{(q_1^2 + M_i^2)(q_2^2 + M_i^2)}, \quad (1.27)$$

where q_1 and q_2 are the invariant masses of the two virtual photons, and M_i is the mass of vector meson propagator i (Quigg and Jackson 1968). Note that the form factor can be expanded to an arbitrary number of terms in order to allow for multiple vector meson couplings. However, because the ρ is so much lighter than the other vector mesons, it is often sufficient to include only a single term with $M_1 = M_\rho$.

Chiral Perturbation

Chiral Perturbation Theory is often used to describe long-distance (low-energy) interactions that are not calculable in closed form. The framework is based on the breaking of the chiral symmetry of the QCD Lagrangian at low energies (below the QCD scale, approximately 300 MeV). Pions are treated as the Goldstone bosons resulting from the spontaneous breaking of the symmetry. The Chiral Lagrangian can be expanded in terms of momentum to accommodate higher-order effects (so long as the energies involved remain below the QCD scale). The theory has been quite successful in predicting several parameters in the neutral kaon sector, such as those describing the decays $K_L \rightarrow \pi\pi l\nu$ (Holstein 1995).

A two-term form factor has been proposed to describe the $K_L \gamma^* \gamma^*$ vertex using Chiral Perturbation Theory (Zhang and Goity 1998). The first term is momentum-independent, and describes contributions from π^0 , η , and η' poles:

$$F_1 = -0.89 \frac{\alpha_{EM} C_8}{2\pi F_\pi}. \quad (1.28)$$

In this expression, F_π is the pion decay constant (93 MeV), and C_8 is the octet coupling in the non-leptonic Chiral Lagrangian of order p^2 ($C_8 \approx 3.12 \times 10^{-7}$).

The second term of the Zhang–Goity form factor accounts for single chiral loop corrections (similar to the loop corrections in QED, except with pions, the lowest mass chiral state, instead of electrons) and counterterms:

$$F_2(q_1^2, q_2^2) = \frac{\alpha_{EM} C_8}{192\pi^3 F_\pi^3} \left[-(a_2 + 2a_4) D(q_1^2, q_2^2, \mu) + C(\mu) (q_1^2 + q_2^2) \right], \quad (1.29)$$

where

$$\begin{aligned} D(q_1^2, q_2^2, \mu) &= (q_1^2 + q_2^2) \left[\frac{10}{3} - \left(\ln \frac{M_K^2}{\mu^2} + \ln \frac{M_\pi^2}{\mu^2} \right) \right] \\ &+ 4 \left[F(M_\pi^2, q_1^2) + F(M_K^2, q_1^2) + F(M_\pi^2, q_2^2) + F(M_K^2, q_2^2) \right], \end{aligned} \quad (1.30)$$

and

$$F(m^2, q^2) = \left(\left(1 - \frac{y}{4} \right) \sqrt{\frac{y-4}{y}} \ln \frac{\sqrt{y} + \sqrt{y-4}}{-\sqrt{y} + \sqrt{y-4}} - 2 \right) m^2, \quad y = \frac{q^2}{m^2}. \quad (1.31)$$

μ is a scaling factor commonly taken to be the ρ mass, $(a_2 + 2a_4) = 1.5 \pm 0.3$, and $C(\mu = M_\rho) = -10.3 \pm 7.3$.

DIP

A general expression for the $K_L \gamma^* \gamma^*$ form factor proposed by D’Ambrosio, Isidori, and Portolés reads

$$f(q_1^2, q_2^2) = \frac{F(q_1^2, q_2^2)}{F(0, 0)} = 1 + \alpha \left(\frac{q_1^2}{q_1^2 - M_V^2} + \frac{q_2^2}{q_2^2 - M_V^2} \right) + \beta \frac{q_1^2 q_2^2}{(q_1^2 - M_V^2)(q_2^2 - M_V^2)}, \quad (1.32)$$

where M_V is conventionally evaluated at the ρ meson pole ($M_V = M_\rho$) (D’Ambrosio et al. 1998). This “DIP” form factor incorporates several of the features from the models described earlier in this section: it is the most general expression compatible with the Lagrangian from $\mathcal{O}(p^6)$ Chiral Perturbation Theory, and contains the vector meson poles indicative of the VMD approximation. α and β are real parameters indicating the relative strengths of the linear and quadratic terms, respectively.

By inserting Equation 1.32 into the integrand of Equation 1.25, it can be shown that

$$\begin{aligned}
\Re R \left(M_K^2 \right) &= -3 \left[\ln (\Lambda / M_0) + 2\alpha \ln (\Lambda / M_\alpha) + \beta \ln (\Lambda / M_\beta) \right] \\
&= -3 \left[\ln (M_\beta / M_0) + 2\alpha \ln (M_\beta / M_\alpha) \right] - 3(1 + 2\alpha + \beta) \ln (\Lambda / M_\beta) \\
&= -5.25 - 3.47\alpha - 3(1 + 2\alpha + \beta) \ln (\Lambda / M_\beta), \tag{1.33}
\end{aligned}$$

where the known values of M_0 , M_α , and M_β (140, 452, and 806 MeV/c², respectively) have been inserted in the final step, and Λ is an ultraviolet cutoff, included to prevent divergence of the integral (D'Ambrosio et al. 1998).

Finally, inserting Equation 1.33 into Equation 1.24 provides a result for the total long-distance contribution to the $K_L \longrightarrow \mu^+ \mu^-$ amplitude:

$$\begin{aligned}
|\Re A_{long}| &= \left[\frac{2\alpha_{EM}^2 M_\mu^2 \beta_\mu \mathcal{B}(K_L \longrightarrow \gamma\gamma)}{\pi^2 M_K^2} \right]^{1/2} \\
&\times |5.25 + 3.47\alpha + 3(1 + 2\alpha + \beta) \ln (\Lambda / M_\beta)|. \tag{1.34}
\end{aligned}$$

The only remaining unknowns in Equation 1.34 are α and β . The values of these parameters must be determined before any further progress can be made.

The authors attempt to theoretically constrain α and β by evaluating the $K_L \gamma^* \gamma^*$ form factor using perturbative QCD, in the realm where $q_1^2 = q_2^2 = q^2 \gg M_K^2$. By matching this high-energy behavior with Equation 1.32 in the realm $\Lambda_{QCD}^2 \ll q^2 \ll 4M_C^2$ (where Λ_{QCD} is the QCD scale and M_C is the charm mass) they obtain the condition (D'Ambrosio et al. 1998)

$$|1 + 2\alpha + \beta| \approx 0.3. \tag{1.35}$$

The authors also propose constraining the form factor parameters by applying the factorization hadronization model (D'Ambrosio and Portoles 1997) to the

pseudoscalar–vector–vector weak vertex (assuming Vector Meson Dominance). This calculation provides the relation

$$1 + 2\alpha + \beta \approx -0.01. \quad (1.36)$$

Along with the experimental results for α from the following section, Equations 1.35 and 1.36 can provide an estimate of the magnitude of the quadratic parameter β .

1.3.4 Current experimental status of α and β

To date, only two direct measurements of the DIP parameter α have been made, both by the KTeV experiment. The first measurement was performed by studying the decay mode $K_L \longrightarrow \mu^+ \mu^- \gamma$ (Alavi-Harati et al. 2001a). Because the second photon in this mode is real, q_2^2 in Equation 1.32 is 0. This mode is therefore sensitive only to α .

The $K_L \longrightarrow \mu^+ \mu^- \gamma$ study was done using two different methods: a shape analysis and a branching ratio analysis. Because the parameter α connects to the momentum of the virtual photon that forms the $\mu^+ \mu^-$ pair, the shape of the invariant mass distribution $M_{\mu\mu}$ is sensitive to the value of α . Monte carlo was generated over a range of α , and a fit was done to the data in $M_{\mu\mu}$ to extract the most likely value of the DIP parameter:

$$\alpha_{\mu\mu\gamma}^{shape} = -1.73_{-0.18}^{+0.14}. \quad (1.37)$$

Meanwhile, the differential decay rate for $K_L \longrightarrow \mu^+ \mu^- \gamma$ is given by (D'Ambrosio et al. 1998)

$$\frac{1}{\Gamma(K_L \longrightarrow \gamma\gamma)} \frac{d\Gamma(K_L \longrightarrow \mu^+ \mu^- \gamma)}{dq^2} = \frac{2}{q^2} \left(\frac{\alpha_{EM}}{3\pi} \right) |f(q^2, 0)|^2 \lambda^{3/2} \left(1, \frac{q^2}{M_K^2}, 0 \right) G_\mu(q^2), \quad (1.38)$$

where

$$\lambda(a, b, c) = a^2 + b^2 + c^2 - 2(ab + bc + ac), \quad (1.39)$$

and

$$G_\ell(q^2) = \left(1 - \frac{4M_\ell^2}{q^2}\right)^{1/2} \left(1 + \frac{2M_\ell^2}{q^2}\right). \quad (1.40)$$

By inserting Equation 1.32 (with $q_2^2 = 0$) into Equation 1.38 and integrating over the kinematically allowed range of q , the $K_L \rightarrow \mu^+ \mu^- \gamma$ branching ratio can be expressed as a function of only one unknown: α . Using the measured branching ratio of $(3.62 \pm 0.09) \times 10^{-7}$ (Alavi-Harati et al. 2001a), α was determined to be

$$\alpha_{\mu\mu\gamma}^{BR} = -1.38 \pm 0.13. \quad (1.41)$$

A weighted average of the shape and branching ratio results provided the published measurement of α from this mode:

$$\alpha_{\mu\mu\gamma} = -1.54 \pm 0.10. \quad (1.42)$$

Using this value of α in Equations 1.35 and 1.36 implies that the value of β is approximately equal to +2.

The second KTeV measurement was performed using the decay mode $K_L \rightarrow e^+ e^- e^+ e^-$. Only a shape analysis was done (Alavi-Harati et al. 2001b), providing the result:

$$\alpha_{4e}^{shape} = -1.1 \pm 0.6. \quad (1.43)$$

Note that this value is consistent with the result from $K_L \rightarrow \mu^+ \mu^- \gamma$. Because the second photon in this mode is virtual, q_2^2 is non-zero. In principle, this makes $K_L \rightarrow e^+ e^- e^+ e^-$ sensitive to the quadratic DIP parameter β . However, due to the kinematic limits imposed by the small electron mass, in practice the sensitivity




Figure 1.6: Diagram for the rare decay $K_L \rightarrow e^+e^-\mu^+\mu^-$. This decay serves as the best available probe of the quadratic β term of the DIP form factor.

is marginal. A shape analysis fit for β yielded the value (Halkiadakis 2001):

$$\beta_{4e}^{shape} = +11.9 \pm 37.8, \quad (1.44)$$

which is certainly consistent with the theoretical prediction, but does nothing to advance understanding of the $K_L \gamma^* \gamma^*$ vertex. A different mode must be used to attempt a measurement of β , and the most obvious choice is $K_L \rightarrow e^+e^-\mu^+\mu^-$.

1.4 Introduction to The Rare Decay $K_L \rightarrow e^+e^-\mu^+\mu^-$

To lowest order, the decay $K_L \rightarrow e^+e^-\mu^+\mu^-$ proceeds via the virtual two-photon state shown in Figure 1.6. Because the invariant mass of the muon pair is much larger than that of the electron pair from $K_L \rightarrow e^+e^-e^+e^-$, coupling to the quadratic DIP parameter β is theoretically much stronger. Additionally, since the

decay $K_L \longrightarrow \mu^+ \mu^- \mu^+ \mu^-$ remains unobserved (estimates of the branching ratio are at the 10^{-13} level (Uy 2002)), $K_L \longrightarrow e^+ e^- \mu^+ \mu^-$ is the only mode to have a real chance to accurately determine β .

While the outlook seems good from a form factor perspective, it must be remembered that $K_L \longrightarrow e^+ e^- \mu^+ \mu^-$ is a *rare* decay mode. This is mainly because the interaction requires two conversions, meaning it occurs at an order α_{EM}^2 lower than $K_L \longrightarrow \gamma\gamma$. Additionally, the (relatively) large invariant mass needed to create a muon pair from one of the virtual photons requires a substantial amount of phase space available for the decay, suppressing the rate even further. Several of the models for the $K_L \gamma^* \gamma^*$ vertex described in Section 1.3.3 offer predictions for the $K_L \longrightarrow e^+ e^- \mu^+ \mu^-$ branching ratio – these predictions will be the subject of the following section.

1.4.1 Branching ratio predictions for $K_L \longrightarrow e^+ e^- \mu^+ \mu^-$

Kroll–Wada type (phase space) predictions

The QED calculation of Miyazaki and Takasugi, which ignores CP -violating effects and assumes a momentum-independent form factor, leads to the prediction (Miyazaki and Takasugi 1973)¹:

$$\frac{\Gamma(K_L \longrightarrow e^+ e^- \mu^+ \mu^-)}{\Gamma(K_L \longrightarrow \gamma\gamma)} = 2.84 \times 10^{-6}. \quad (1.45)$$

Using the most recent measurement of the $K_L \longrightarrow \gamma\gamma$ branching ratio of $(5.86 \pm 0.15) \times 10^{-4}$ (Groom et al. 2000b),

$$\mathcal{B}(K_L \longrightarrow e^+ e^- \mu^+ \mu^-) = (1.66 \pm 0.04) \times 10^{-9}. \quad (1.46)$$

¹The published rate for $\Gamma(K_L \longrightarrow e^+ e^- \mu^+ \mu^-) / \Gamma(K_L \longrightarrow \gamma\gamma)$ in Table II of this paper is actually 1.42×10^{-6} . This is an error; many of the entries in this table are off by a factor of 2.




Figure 1.7: Graphical representation of ϕ , the angle between the normals to the electron and muon decay planes in the K_L rest frame.

The extension of this model by Uy allowed for the presence of a CP -violating contribution to the vertex. The sizes of CP -violating and CP -conserving contributions are respectively parameterized by the momentum-independent form factor coefficients g and h . The branching ratio, of course, depends on the relative sizes of these coefficients. The differential rate is given by (Uy 1991)

$$\begin{aligned}
 \frac{1}{\Gamma(K_L \rightarrow \gamma\gamma)} \frac{d\Gamma(K_L \rightarrow e^+e^- \mu^+\mu^-)}{d\phi} = & \\
 & \frac{F\sigma_1 \left[l_1 \sin^2 \phi + (g/h)^2 l_2 \cos^2 \phi - (g/h) \sin \sigma l_3 \sin \phi \cos \phi + l_4 + (g/h)^2 l_5 \right]}{\left[1 + 2(g/h)^2 \right]} \\
 = & \frac{F\sigma_2 \left[(h/g)^2 s_1 \sin^2 \phi + s_2 \cos^2 \phi - (h/g) \sin \delta s_3 \sin \phi \cos \phi + (h/g)^2 s_4 + s_5 \right]}{\left[(h/g)^2 + 2 \right]}.
 \end{aligned} \tag{1.47}$$

Here ϕ is the angle between the normals to the electron and muon decay planes in the K_L rest frame, shown in Figure 1.7. The coefficients F , σ , δ , σ_i , s_i , and l_i are known constants (calculated during numerical integration of the QED matrix element). For the extreme case where $h = 1$, $g = 0$ (pure CP -conserving), it is

found that

$$\begin{aligned}\frac{\Gamma(K_L \longrightarrow e^+e^-\mu^+\mu^-)}{\Gamma(K_L \longrightarrow \gamma\gamma)} &= 2.86 \times 10^{-6} \\ \mathcal{B}(K_L \longrightarrow e^+e^-\mu^+\mu^-) &= (1.68 \pm 0.04) \times 10^{-9},\end{aligned}\quad (1.48)$$

where the most recent value for the $K_L \longrightarrow \gamma\gamma$ branching ratio has again been used. Note that this result is consistent with the prediction from Miyazaki and Takasugi, as one would expect. Meanwhile, for the $h = 0$, $g = 1$ (pure CP -violating) case,

$$\begin{aligned}\frac{\Gamma(K_L \longrightarrow e^+e^-\mu^+\mu^-)}{\Gamma(K_L \longrightarrow \gamma\gamma)} &= 0.644 \times 10^{-2} \\ \mathcal{B}(K_L \longrightarrow e^+e^-\mu^+\mu^-) &= (3.78 \pm 0.10) \times 10^{-6}.\end{aligned}\quad (1.49)$$

The large sensitivity to g comes about because of the last term in the numerator of the middle line of Equation 1.47. The coefficient l_5 turns out to be four orders of magnitude larger than the other l_i s, meaning that turning on even a small, non-zero value of g leads to a significant increase in the rate. Since the differential rate is sensitive to the amount of CP -violation present, the $K_L \longrightarrow e^+e^-\mu^+\mu^-$ branching ratio can theoretically be used to measure the quantity $(g/h)^2$.

In addition to its effect on the branching ratio, any CP -violating contribution to the form factor (non-zero g) will create an asymmetry in the angular distribution of the decay products. Due to interference between certain contributions to the matrix element, such an effect is observed in the decay $K_L \longrightarrow \pi^+\pi^-e^+e^-$. An asymmetry of approximately 13% is present in the distribution $\sin\phi\cos\phi$, where ϕ is the angle between the electron and pion decay planes in the K_L rest frame (Alavi-Harati et al. 2000). The same distribution (with muons substituting for pions, of course) can be studied to determine the presence of CP -violating contributions to $K_L \longrightarrow e^+e^-\mu^+\mu^-$.

Vector Meson Dominance prediction

A Vector Meson Dominance calculation by Quigg and Jackson assumes that the $K_L \gamma^* \gamma^*$ interaction is mediated by a linear combination of vector meson

states (Quigg and Jackson 1968). Based on the nonet model of pseudoscalar–vector–vector coupling, the interaction Lagrangian contains the term

$$V = (\rho\rho + \omega\omega - 2\phi\phi)\eta. \quad (1.50)$$

This model predicts

$$\begin{aligned} \frac{\Gamma(K_L \longrightarrow e^+e^-\mu^+\mu^-)}{\Gamma(K_L \longrightarrow \gamma\gamma)} &= 4 \times 10^{-6} \\ \mathcal{B}(K_L \longrightarrow e^+e^-\mu^+\mu^-) &= (2.34 \pm 0.06) \times 10^{-9}, \end{aligned} \quad (1.51)$$

where CP -violating effects are ignored.

Chiral Perturbation prediction

The Chiral Perturbation calculation of the $K_L \gamma^* \gamma^*$ vertex by Zhang and Goity was carried out to order p^6 , using the form factor terms given in Equations 1.28 and 1.29. CP -violating effects enter into the Lagrangian at $\mathcal{O}(p^4)$, but are suppressed by the size of the parameter ϵ (recall that $\epsilon \approx 2 \times 10^{-3}$), meaning their contribution could be safely ignored. This led to the prediction (Zhang and Goity 1998)

$$\begin{aligned} \frac{\Gamma(K_L \longrightarrow e^+e^-\mu^+\mu^-)}{\Gamma(K_L \longrightarrow \gamma\gamma)} &= (2.20 \pm 0.25) \times 10^{-6} \\ \mathcal{B}(K_L \longrightarrow e^+e^-\mu^+\mu^-) &= (1.29 \pm 0.15) \times 10^{-9}. \end{aligned} \quad (1.52)$$

Note that this result is smaller than the CP -conserving prediction by Uy given in Equation 1.48.

DIP prediction

The DIP form factor (Equation 1.32) cannot be used to predict the $K_L \longrightarrow e^+e^-\mu^+\mu^-$ branching ratio *a priori*, since it depends on the unknown parameters α and β . However, by integrating the differential rate over kinematically allowed values of q_e^2 and q_μ^2 , the $K_L \longrightarrow e^+e^-\mu^+\mu^-$ branching ratio can be used to extract

values for the DIP parameters. In terms of photon momentum, this differential rate can be expressed as

$$\frac{1}{\Gamma_L^{\gamma\gamma}} \frac{d\Gamma_L^{e^+e^-\mu^+\mu^-}}{dq_e^2 dq_\mu^2} = \frac{2}{q_e^2 q_\mu^2} \left(\frac{\alpha_{EM}}{3\pi} \right)^2 |f(q_e^2, q_\mu^2)|^2 \lambda^{3/2} \left(1, \frac{q_e^2}{M_K^2}, \frac{q_\mu^2}{M_K^2} \right) G_e(q_e^2) G_\mu(q_\mu^2), \quad (1.53)$$

where λ and G_ℓ were defined in Equations 1.39 and 1.40, and $f(q_e^2, q_\mu^2)$ comes from Equation 1.32. This analysis will be done in Chapter 7.

1.4.2 Experimental status of $K_L \longrightarrow e^+e^-\mu^+\mu^-$

Because the decay $K_L \longrightarrow e^+e^-\mu^+\mu^-$ is so rare, experimental measurements have become possible only recently. The first observation of the decay was announced by the E799-I collaboration (the predecessor to KTeV) in 1996. Based on a single event, a branching ratio of

$$\mathcal{B}(K_L \longrightarrow e^+e^-\mu^+\mu^-) = (2.9_{-2.4}^{+6.7}) \times 10^{-9} \quad (1.54)$$

was published (Gu et al. 1996).

The 1997 run of the KTeV experiment significantly increased the world sample of events (KTeV run periods will be described in detail in the following chapter). Based on a total of 43 observed events, the branching ratio was measured to be (Alavi-Harati et al. 2001c)

$$\mathcal{B}(K_L \longrightarrow e^+e^-\mu^+\mu^-) = (2.62 \pm 0.40_{\text{stat}} \pm 0.17_{\text{syst}}) \times 10^{-9}. \quad (1.55)$$

However, because of the limited size of the dataset, no attempt was made to extract any DIP form factor information.

A final run of the KTeV experiment was approved for 1999, providing the chance to collect even more events from this important decay mode.

1.5 Dissertation Overview

The analysis described in this dissertation is based on the entire KTeV dataset. In addition to reexamining the data leading to the 1997 result given earlier, an analysis of data from the 1999 run will be presented for the first time. Using the total dataset, a new value for the $K_L \rightarrow e^+e^-\mu^+\mu^-$ branching ratio is extracted. A new limit on the branching ratio for the lepton flavor-violating counterpart $K_L \rightarrow e^\pm e^\pm \mu^\mp \mu^\mp$ will also be calculated. Measurements of the DIP $K_L \gamma^* \gamma^*$ form factor parameters α and β , as well as a limit on CP -violating contributions to the $K_L \gamma^* \gamma^*$ vertex, will also be presented.

The following chapter will describe the KTeV beamline and detector, while the KTeV trigger system will be outlined in Chapter 3. In Chapter 4, the software algorithms used to reconstruct events are discussed. The KTeV monte carlo, used to simulate events in the detector, is the topic of Chapter 5. The branching ratio analysis, along with detailed descriptions of associated systematic errors, is presented in Chapter 6. Finally, Chapter 7 is devoted to the extraction of the $K_L \gamma^* \gamma^*$ form factor parameters.

CHAPTER 2

The KTeV Experiment

The name “KTeV” stands for “kaons at the Tevatron”, and actually represents two distinct experiments. E832 was designed to measure the direct CP -violating parameter¹ $\Re(\epsilon'/\epsilon)$ to a precision on the order of 10^{-4} . E799, while primarily approved to search for the direct CP -violating decay $K_L \longrightarrow \pi^0 e^+ e^-$, also provided an opportunity to study a wide range of rare kaon and hyperon decays.

The configuration of the KTeV beamline and detector varied slightly between the two experiments. The main difference was the use of an active regenerator in E832 in order to deliver simultaneous, parallel K_L and K_S beams to the experiment. This regenerator was removed for E799 running, leaving two parallel K_L beams. Absorber and collimator settings also differed somewhat between the two experiments. The following analysis is based on data collected while running in the E799 configuration, which will be discussed in detail over the following pages.

The E799 experiment was broken into three run periods, depicted in Figure 2.1. The first period, known as the *Winter* run, spanned from late January to late March of 1997. The *Summer* run started in late July and lasted until early September of 1997. These two periods are known collectively as the *1997* run. This was followed by a two year downtime, which allowed for significant detector upgrades. The final run started in mid-September of 1999 and lasted through mid-January of 2000 — this period is referred to as the *1999* run.

¹*Direct* CP -violation occurs when the pure CP -odd K_2 decays directly into the CP -even two pion final state. This is to be compared to the phenomenon of *indirect* CP -violation, described in Section 1.2, in which the *mostly* CP -odd K_L decays to two pions.

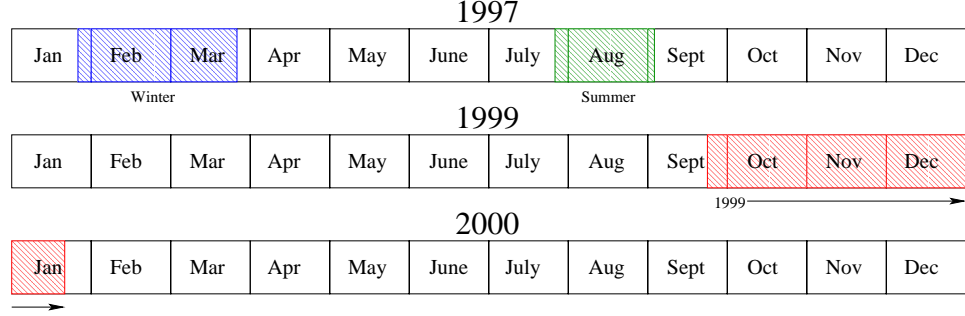


Figure 2.1: Calendar depicting the three run periods of the KTeV E799 experiment: Winter, Summer, and 1999.

2.1 Creation of the K_L beams

The KTeV experiment was located on the Neutrino–Muon (NM) fixed–target beamline at Fermi National Accelerator Laboratory, better known as Fermilab. At the heart of Fermilab is the Tevatron, a circular proton/anti–proton accelerator capable of reaching energies of up to 1 TeV per beam. It is here that the description of the KTeV beam begins.

2.1.1 The proton beam

Protons were accelerated in the Tevatron to an energy of 800 GeV before being delivered to the NM beamline. The time required for this acceleration was 40 seconds, referred to as the *off-spill* period. After acceleration, the protons were delivered continuously for a period referred to as *on-spill*, or simply *the spill*. During the 1997 run, the spill lasted for 20 seconds. Improvements to the beamline magnets prior to the 1999 run allowed for stretching this period to 40 seconds, leading to a large improvement in duty cycle for the experiment.

Protons were delivered from the Tevatron in 1 – 2 ns pulses, spaced by 19 ns. A copy of this 56 MHz beam RF signal was sent to KTeV in order to synchronize

the timing of the various detectors. Beam intensity varied over the course of the experiment. Typically, between 2 and 4×10^{12} protons were delivered per spill for 1997; during the 1999 run the intensity was increased to $6 - 10 \times 10^{12}$ protons per spill. Note that due to the longer on-spill period used in the later run, the *instantaneous* intensities were roughly equal between 1997 and 1999.

After being focused to a width of less than $250 \mu\text{m}$, the proton beam was directed onto a BeO target housed in a beamline enclosure known as NM2. The composition of the target was chosen for resistance to thermal stresses. The center of the target served as the origin of the KTeV coordinate system, with the $+z$ axis parallel to the downstream beam direction, and $+y$ pointing upward towards the ceiling. The $+x$ direction was chosen to complete a right-handed system.

The target presented a 3 mm square cross-section in the $x - y$ plane and extended 30 cm in z , corresponding to 1.1 proton interaction lengths; optimal for kaon production. Furthermore, the incident proton beam was directed downwards with respect to the $+z$ axis at an angle of 4.8 mrad. This angle was chosen to maximize the kaon to neutron production ratio.

A target monitor was constructed to detect accidental beam interactions unrelated to kaon decays. Located 1.8 m away from the target and aligned perpendicular to the beam direction, this monitor consisted of three sequential scintillator counters that viewed the target through a hole in its shielding. A coincidence in the three counters fired an accidental trigger, and the event was recorded for later use in monte carlo simulations. These *accidental overlays* are discussed in detail in Section 5.6. The amount of accidental activity per event was roughly proportional to the intensity of the incident beam.

2.1.2 The NM2 enclosure

The beam exiting the target contained some neutral kaons, but consisted mostly of undesired particles such as protons, pions, muons, and photons. Most of the beamline elements in the NM2 enclosure were designed to rid the beam of these particles and collimate the remaining beam into the desired shape. A schematic of the enclosure is shown in Figure 2.2.

The first element, known as the *target sweeper*, was used to deflect remaining protons into a beam dump. The target sweeper extended from 0.6 to 4.4 m in z , and imparted a transverse momentum (P_t) kick of 475 MeV/c to the protons in the $-y$ direction. The *primary proton dump*, a 4.5 m block of water-cooled copper, was offset in y to receive the deflected protons.

Downstream of the primary proton dump was *μ sweep1*. This magnet spanned from 12.3 to 17.8 m in z , and imparted a P_t kick of 3806 MeV/c in the $+x$ direction to muons in order to sweep them out of the beam.

The *Pb absorber*, a 3 inch (14 radiation length) lead wall, was placed at $z = 19$ m in order to convert photons in the beam to electron-positron pairs. Immediately downstream of this wall was the *primary collimator*, used to initially shape the K_L parallel beams. The collimator consisted of a 2 m long brass and steel block with two tapered, rectangular holes. The center-lines of these holes lay in the $x - z$ plane and pointed back to the target, separated by an angle of 1.6 mrad. During the Winter run, the openings at the downstream face of the primary collimator measured 1.18×1.29 cm. They were enlarged to 1.62×1.73 cm for the Summer and 1999 runs.

At $z = 21.8$ m the beams entered a vacuum region, capped at the upstream end by a 0.005 inch thick titanium window. Just inside of this region was *μ sweep2*, a 6 m long magnet used to sweep away charged particles created from interactions in the Pb absorber, primary collimator, and the titanium window. For the Winter

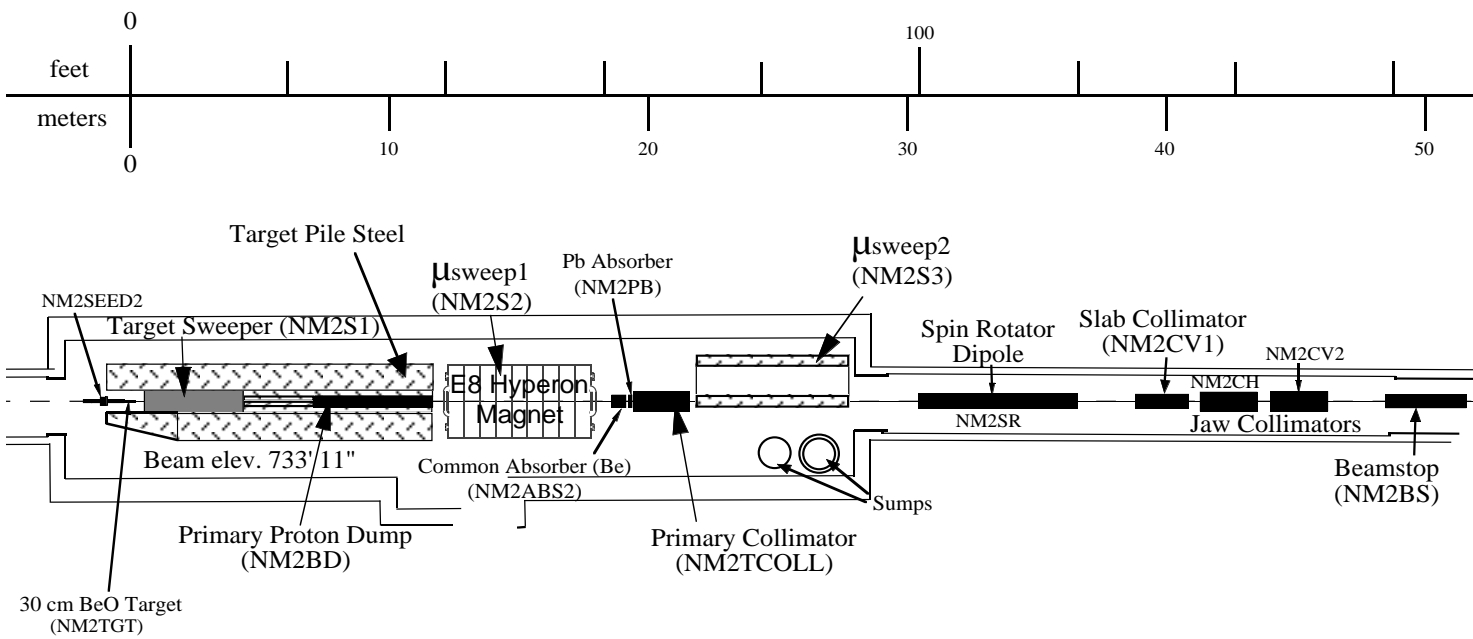


Figure 2.2: y view of the beamline elements located in the NM2 enclosure. The common absorber and jaw collimators were not used during E799 running.

run, *μsweep2* was set to a P_t kick of 3135 MeV/c. The kick was lowered to 1854 MeV/c for the Summer and 1999 runs.

The *spin rotator dipole* stretched from 30.5 to 36.5 m in z . This magnet was used to select the polarization state of Ξ^0 and Λ^0 particles for hyperon studies. The spin rotator had no effect on the spinless K_L s.

The *slab collimator* was a 2 m long, tapered, stainless steel block starting at $z = 38.8$ m. This collimator was initially used in the Winter run to prevent crossover between the two beams, though it was found to be unnecessary and was removed before the start of the Summer run.

A 6 m thick iron beam stop was located at 46.4 m. The beam stop was inserted as a safety measure during accesses to the experiment, and was also used to create muon beams for special detector calibration runs.

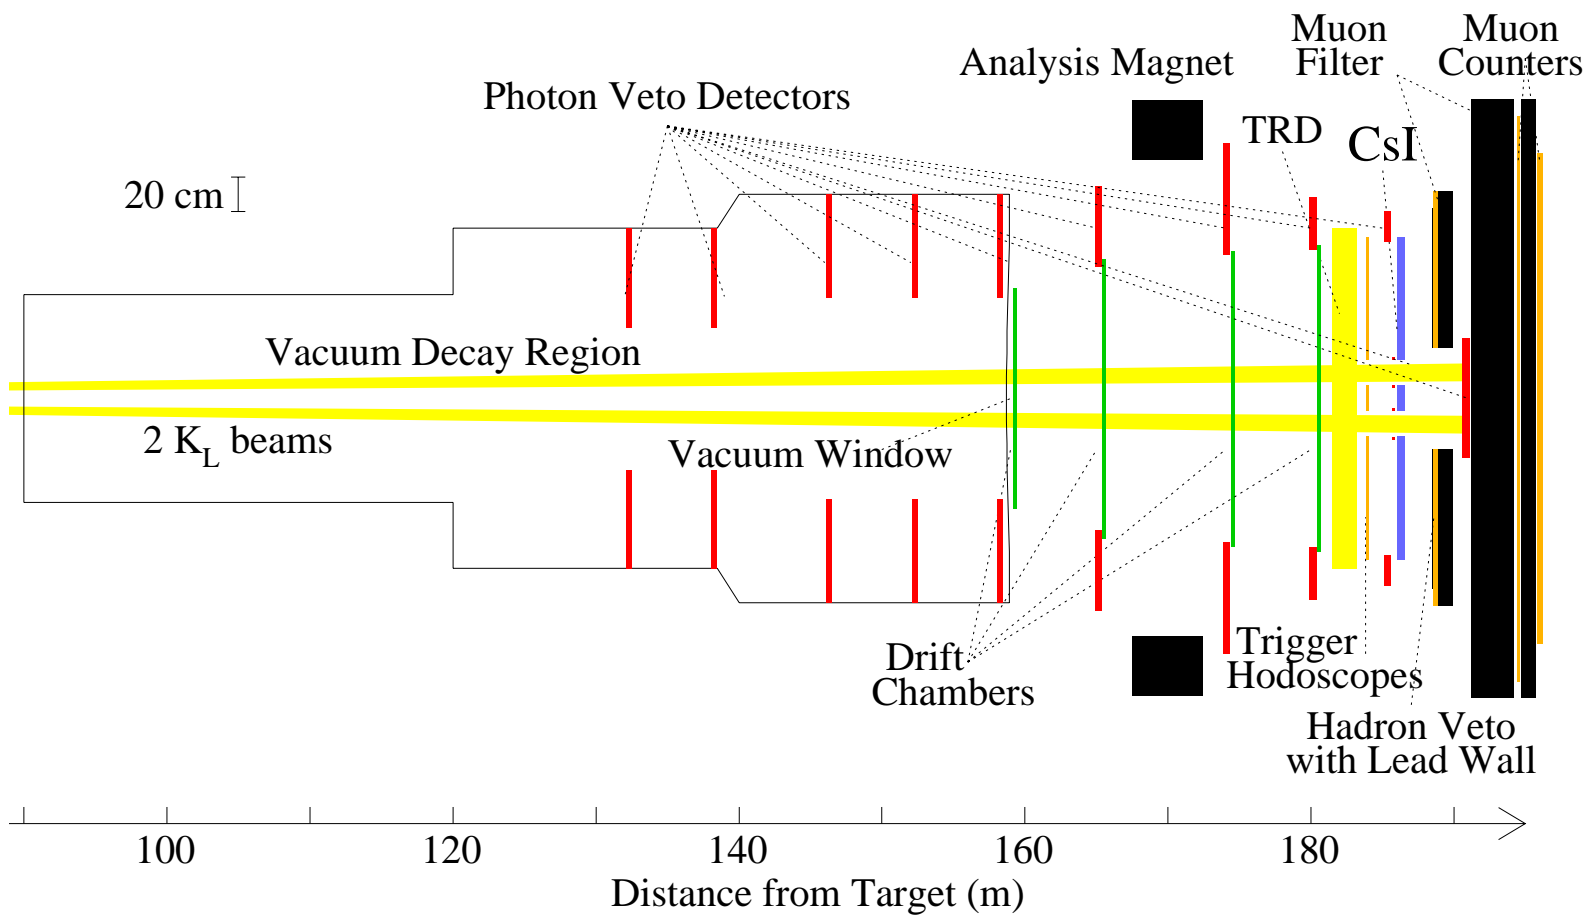
Final shaping of the neutral beams was done by the *defining collimator*, a 3 m long tungsten and iron block starting at $z = 85$ m. Tapered holes provided openings at the downstream face measuring 4.4 cm square for the Winter run, and 5.2 cm square for the Summer and 1999 runs. After exiting the block, the beams passed through the *final sweeper*, a 3 m long magnet designed to remove charged particles from interactions within the defining collimator or from upstream decays.

At this point, formation of the neutral beams was complete. The beam composition was about 3:1 neutron to K_L , with small fractions of remaining K_S , Λ^0 , and Ξ^0 (Bocean et al. 1998). The neutral hadron rate was measured to be between 25 and 50 MHz.

2.2 The KTeV Detector

A schematic of the KTeV detector is given in Figure 2.3. The main components of the detector are outlined in the following sections.

Figure 2.3: y view schematic of the KTeV detector in the E799 configuration.



Spectrometer element	z (m)	Dimensions, $x \times y$ (m)
DC1	159.4	1.30×1.30
DC2	165.6	1.64×1.44
Analysis magnet	170.0	2.90×2.00
DC3	174.6	1.74×1.64
DC4	180.5	1.90×1.90

Table 2.1: z positions and dimensions of KTeV spectrometer elements.

2.2.1 Vacuum region

The vacuum region was a 65 m long tank that defined the fiducial volume for K_L decays. This tank was maintained at a pressure of 1 μ Torr. Veto counters ringed the interior of the tank in order to detect particles escaping the fiducial volume (more details on the KTeV veto systems are provided in Section 2.2.8).

The downstream end of the vacuum region was sealed by a 1.8 m diameter Mylar-laminated, Kevlar window. Kevlar was needed to withstand the enormous force exerted on the window by atmospheric pressure, while Mylar was chosen to reduce permeability of the window to air. This combination led to a total window thickness of only 0.0015 radiation lengths.

2.2.2 Spectrometer

To measure the momenta and charges of charged particles from K_L decays, KTeV used a spectrometer consisting of four drift chambers and an analysis magnet. The z positions and dimensions of these spectrometer elements are given in Table 2.1. Helium bags were hung between the spectrometer elements to reduce the probabilities of multiple scattering, photon conversions, and beam interactions in air.

Analysis magnet

The analysis magnet produced a field along the y axis with a strength on the order of 0.5 T. In 1997, the field was tuned such that relativistic charged particles traveling perpendicular to the plane of the magnet received a P_t kick of 205 MeV/c in the $x - z$ plane². In order to increase the detector acceptance for four-track decay modes, the field was lowered to provide a kick of 150 MeV/c for the 1999 run. This cut down on the number of “soft” particles that were kicked outside of the fiducial volume of the experiment by the analysis magnet. The polarity of the field was flipped regularly to cancel out any possible systematic effects.

Drift chambers

The drift chambers were strung with 25 μm diameter gold-plated tungsten field (cathode) wires and 100 μm diameter gold-plated aluminum sense (anode) wires. The sense wires were arranged to form *plane pairs*: two parallel planes, offset by a half-wire in order to resolve left-right ambiguity. This arrangement formed hexagonal cells measuring 12.7 mm across (see Figure 2.4). Each drift chamber contained two plane pairs. The wires of the upstream pair, known as the x *view*, were aligned parallel to the y axis, while the y *view* wires were aligned parallel to the x axis.

Mylar windows were mounted on the drift chambers to seal in the drift gas. The gas mixture consisted of equal parts argon and ethane, along with approximately 1% isopropyl alcohol. The alcohol absorbed ultraviolet light, which was found to cause deposits on the chamber wires. During the 1997 runs, the field wires were held at voltages between 2450 and 2600 V with respect to the sense wires, corresponding to an electron drift speed of approximately 50 $\mu\text{m}/\text{ns}$ and a drift time of less than 200 ns. During the two-year downtime, the sense wires were cleaned, and components on the chamber-mounted electronics cards were upgraded. These

²These values for the P_t kicks are approximate. Exact values were calibrated on a regular basis by measuring the reconstructed invariant mass distributions from $K_L \rightarrow \pi^+ \pi^-$ decays.




Figure 2.4: y view schematic of the field and sense wire arrangement within a drift chamber.

improvements allowed for running the chambers at higher gain with lower voltages during the 1999 run, typically between 2350 and 2450 V (Corcoran). Running at lower voltages is advantageous because noise levels are significantly reduced. Signals collected on the chamber sense wires were amplified and discriminated by custom chamber-mounted electronics before being digitized by time-to-digital converter (TDC) modules operating in common-stop mode.

Calibration

In order to reconstruct particle tracks, it was necessary to map the time information from the TDC modules into distances within the chamber cells. Special runs were taken to generate these *XT maps*. During these runs, the beam stop was inserted to block neutral particles, the analysis magnet was turned off, and the settings for the sweeper magnets in NM2 were adjusted to allow charged particles to enter the detector. This resulted in a beam of straight-through muons that were used for online calibration and alignment purposes; these runs were therefore known as *muon runs*. For the more precise offline calibration, $K \rightarrow \pi e \nu$ (also called K_{e3}) events were used to generate the XT maps.

Resolution

The overall momentum resolution of the spectrometer was found to be

$$\frac{\sigma(P)}{P} = 0.38\% \oplus 0.016\% P, \quad (2.1)$$

with P measured in GeV/c (Hanagaki 1998). The first term in this expression results from multiple scattering effects, and the second is from the finite resolution of measured hit positions (about 100 μm) within the chambers.

2.2.3 Transition radiation detectors

Pion–electron discrimination is very important in certain analyses (the search for $K_L \rightarrow \pi^0 e^+ e^-$, for example). A transition radiation detector (TRD) was used to assist in particle identification. The KTeV TRD system consisted of eight planes, the first located at $z = 181.1$ m with the remaining planes staggered at 31 cm intervals. Each plane contained a polypropylene felt radiator followed by a multiwire proportional counter (MWPC). Discrimination was based on the X-ray signature detected by the MWPCs, providing pion rejection ratios of up to 300 : 1 with 90% electron efficiency (Solomey 1998). The TRD system was not used in this analysis.

2.2.4 Trigger hodoscopes

The Level 1 trigger (discussed in the following chapter) required information regarding the existence of charged particles on a much faster timescale than the drift chambers could provide (recall that the maximum electron drift time in the chambers was approximately 200 ns). For this reason, two sequential trigger hodoscope planes, designated V (upstream plane) and V' (downstream plane), were placed at $z = 184$ m. These planes measured 1.9 m square and consisted of 1 cm thick scintillator paddles, each paddle being viewed by a photomultiplier (PMT) tube. The two planes contained different numbers and sizes of paddles, designed to prevent cracks between paddles that would be visible to incident particles. 14 cm square beam holes were cut in each plane to decrease the number of beam interactions, reduce trigger rates, and prevent radiation damage to the calorimeter. Schematics of the V and V' planes are given in Figure 2.5.

2.2.5 Electromagnetic calorimeter

The KTeV electromagnetic (EM) calorimeter was located two meters downstream of the trigger hodoscopes, at $z = 186$ m. It consisted of 3100 0.5 m long blocks of




Figure 2.5: V and V' trigger hodoscope planes.

pure CsI crystal stacked in a 1.9 m square array. 868 larger blocks (5 cm square) were used to form the outside of the array, while 2232 smaller blocks (2.5 cm square) were used for the interior region, where greater position resolution was needed. Two 15 cm square holes allowed the neutral beams to pass through without damaging the crystals (see Figure 2.6 (Alavi-Harati 1999)).

The depth of the array is equivalent to 27 radiation lengths, ensuring that most EM showers were completely contained within the calorimeter. The array depth also translates to 1.4 nuclear interaction lengths, giving a probability of approximately 30% that a charged pion would hadronically interact and shower within the calorimeter.

CsI crystals

The light yield in a CsI crystal was measured to be about 20 photoelectrons per MeV deposited (Roodman 1997). There are two components to this scintillation light. The “fast” contribution peaks at a wavelength of 305 nm and falls off with




Figure 2.6: Schematic of the KTeV electromagnetic calorimeter. 3100 pure CsI crystals were used in the array.

a decay time of 25 ns. The origin of the fast component has only recently been discovered: a radiative decay of self-trapped excitons (STEs), dominated by the contribution from the off-center 4.3 eV STE state (Nishimura et al. 1995). The “slow” component peaks at 480 nm and decays on the μ s scale. Minimally ionizing particles (MIPs) traveling parallel to the crystal axis deposit approximately 320 MeV of energy in the crystal (Prasad 2002).

Scintillation response along the length of the crystals was found to be somewhat non-uniform, the magnitude varying from crystal to crystal. To compensate for this effect, the crystals were wrapped in 13 μ m thick black and/or reflective mylar coverings. The type and placement of covering was determined by studying the longitudinal scintillation response of each crystal individually. With the wrappings, these responses were measured to be uniform to within 5% (Shawhan 2000).

Optical readout

A RTV rubber cookie was used to optically couple the crystals to their respective PMTs. In certain instances, an opaque mask was used to compensate for variation of light output from crystal to crystal. In addition, a filter was placed on the cookie to block the slow component of the scintillation light, in order to prevent pileup.

The large crystals were coupled to 6-stage Hamamatsu 1.5 inch R5330 PMTs, while 5-stage 0.75 inch R5364 Hamamatsu tubes were used to view the small crystals. The operating voltages of these PMTs were individually chosen to tune the thresholds for the Hardware Cluster Counter (discussed in Section 3.2.1), but typically varied between 900 and 1500 V. The signal at the final dynode stage of the PMT was picked off for use by the Etot system of the Level 1 trigger, discussed in Section 3.1.2. The anode signal was sent to a DPMT board for digitization.

DPMTs and digitization

The DPMT (digital PMT) board was designed to digitize and buffer the anode signal from the calorimeter PMTs. The board contained an 8-bit flash analog-to-digital converter (ADC), and two custom chips: the charge integrator and encoder (QIE) and the driver-buffer-clock (DBC).

The QIE divided the signal current from the PMT among eight capacitors, each receiving a different fraction of the current ($I/2, I/4, \dots, I/256$). The voltage across each capacitor was integrated and compared to a reference voltage to determine which value was in range. This voltage was then digitized by the flash ADC and stored as a mantissa. A 3-bit exponent containing the selected range was also stored.

Four such circuits were present in each QIE. They were activated in round-robin fashion each RF bucket. The circuit number, along with the mantissa and exponent, formed a data word that was written to a FIFO (first-in, first-out) buffer in the DBC. This buffer had a depth of 32 words, and was read out only if certain Level 1 trigger conditions were satisfied. In this case, the data were sent to a custom VME pipeline and sparsification system. The sparsification process removed a channel from the readout list if the energy in that channel was less than a given threshold, greatly reducing the amount of memory required to describe an event. Data were read out from the VME pipeline only if the event passed the Level 2 trigger.

Calibration

Light pulses from a dye laser were distributed to each crystal via optical fiber. During normal running, the laser flashed continuously at constant intensity. These *quiescent laser flashes* were used to flag time-dependent calorimeter problems (drifting PMT gain, for example). During special runs, known as *laser scans*, the entire

DPMT dynamic range was scanned by slowly turning a variable filter wheel in front of the laser source. The laser flashes were also sent to several highly linear PIN diodes, providing a reference. The relationship between the charge as seen by the PIN diode and the flash ADC output from the DPMT was found to be linear; the slope and offset for each DPMT range were extracted as calibration constants.

The next step was to find the conversion factor from charge to energy for each channel (known as the Q/E constant). This was done by studying the ratio E/p of calibration K_{e3} electrons. Here E is the electron energy deposited in the calorimeter, and p is the corresponding momentum as measured by the spectrometer. Since these relativistic electrons were expected to deposit all of their energy in the EM calorimeter, their E/p should be 1.

The Q/E calibration was an iterative process. An initial guess was made at the value of the Q/E constant for a given channel. The E/p for every calibration electron that hit that channel was calculated and histogrammed. Based on the mean of the resulting distribution, another guess was made at the constant and the process was repeated. Figure 2.7 shows the E/p for all calibration electrons in all channels after five such iterations. Note that the mean of the distribution has converged to the expected value of 1.

Resolution

The overall energy resolution of the spectrometer was found to be

$$\frac{\sigma(E)}{E} = 0.45\% \oplus \frac{2\%}{\sqrt{E}}, \quad (2.2)$$

with E measured in GeV (Shawhan 2000). The constant term is caused by an amalgam of effects, such as light leakage, noise, and uncorrected non-uniformities. The second term results from the photostatistics of the scintillation light. For reference, the mean electron energy from $K_L \rightarrow e^+e^-\mu^+\mu^-$ decays is about 16 GeV.

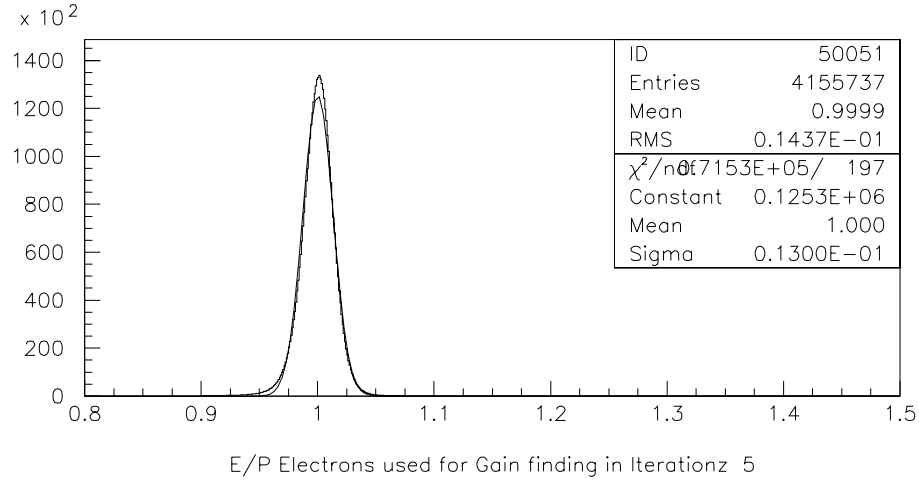


Figure 2.7: E/p of calibration K_{e3} electrons in all 3100 channels after five iterations.

The position resolution was measured to be approximately 1 mm in the small crystals, and about 1.8 mm in the large crystals. The method used to determine position in the calorimeter will be discussed in Section 4.2.3.

2.2.6 Hadron anti

A 10 cm thick lead wall was located just downstream of the calorimeter (see Figure 2.8). This wall served two purposes: it absorbed any EM shower leakage out of the back of the calorimeter, and it induced hadronic showers for detection by the *hadron anti* (HA). The HA was a single plane of 28 scintillator paddles, measuring 2.24 m square. The analog sum of HA channels was discriminated and sent to the trigger, providing a veto on hadronic activity. The HA was not used in this analysis.

Directly behind the HA was a 1 m thick steel wall. Known as MUF1 (muon filter 1), this wall protected the HA from backslash off of the neutral beam dump located further downstream. Holes were cut in the lead wall and the HA to allow

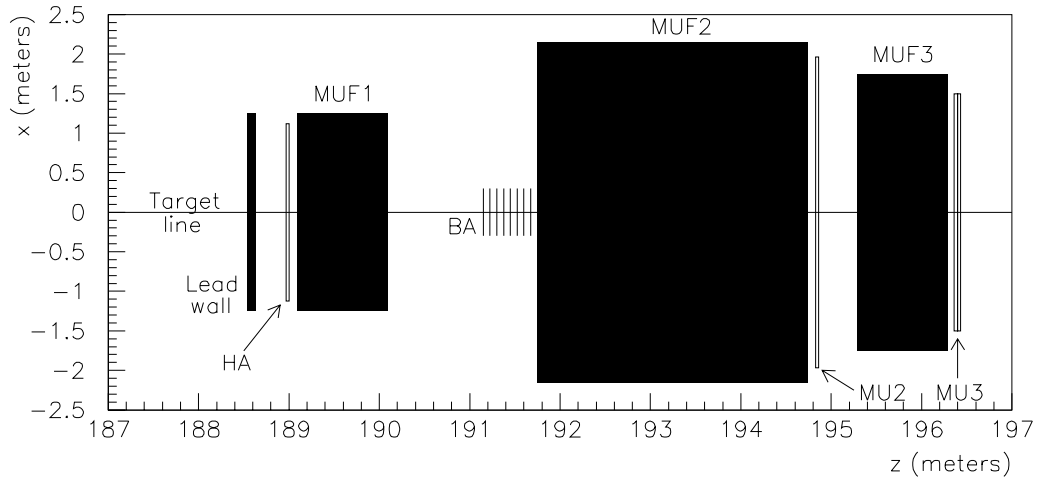


Figure 2.8: y view of the KTeV detector downstream of the EM calorimeter.

for passage of the neutral beams.

2.2.7 Muon identification system

A 3 m deep steel wall known as MUF2 served as the neutral beam dump. Directly downstream of MUF2 was a bank of 56 scintillator paddles. Called MU2, this bank served as a muon veto. The paddles in MU2 overlapped by 1 cm in order to prevent muons from passing undetected through any cracks (see Figure 2.9).

Following MU2 was another filter, a 1 m thick steel wall known as MUF3. Finally, at the tail end of the experiment ($z = 196.4$ m) were two muon identification planes, MU3Y and MU3X (known collectively as MU3). These planes each contained 40 non-overlapping, 15 cm wide scintillator paddles. The paddles of MU3Y were aligned parallel to the x axis, and vice versa. These planes are also depicted in Figure 2.9.

Taken together, the lead wall and the three muon filters add up to a total of 31 nuclear interaction lengths, translating to a 0.5% probability that a hadronic

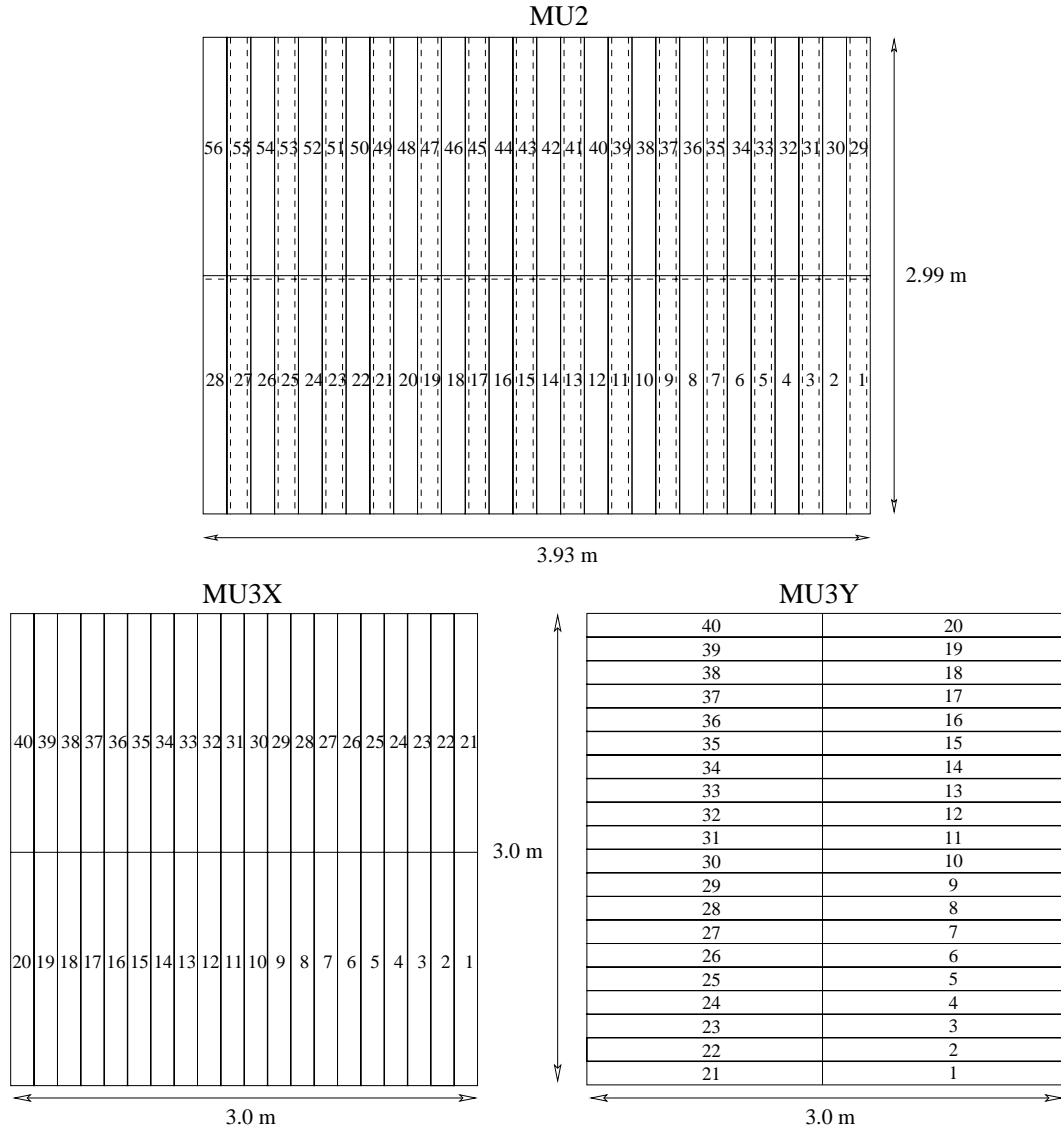


Figure 2.9: Schematic of muon trigger planes. Dashed lines indicate the 1 cm counter overlaps in MU2.




Figure 2.10: Schematics of typical photon veto detectors: on the left, RC10; on the right, SA4.

shower from a 20 GeV/c pion could leak through to MU3. The momentum threshold for a muon reaching MU3 was found to be approximately 7 GeV/c (Quinn 2000).

2.2.8 Photon vetos

Several different types of detectors were employed to veto on particles escaping the fiducial volume of the experiment. Within the vacuum region were five *ring counters*, labeled RC6–10. These had a round outer perimeter to fit flush against the inner wall of the vacuum tank, and a rectangular inner perimeter (see Figure 2.10). The RCs were segmented azimuthally into 16 overlapping paddles. Each paddle contained 24 lead–scintillator layers for a total of 16 radiation lengths. The paddles were individually instrumented with PMTs. The *spectrometer anti*s (SA2–4) surrounded drift chambers 2 through 4, while the *CsI anti* (CIA) framed the face of the calorimeter. These rectangular detectors were also segmented into 16 radiation length thick paddles. A typical SA is also shown in Figure 2.10. If at least 0.5 GeV of in–time energy was deposited in any paddle of the RCs, SAs, or CIA, a signal was sent to the Level 1 trigger to be used as a veto.




Figure 2.11: Schematic of the CA on the calorimeter face. The blocks represent the small crystals of the inner region of the calorimeter.

The *collar anti* (CA) was designed to veto on particles that hit the face of the calorimeter near the beam holes. This was a concern because leakage of EM shower energy into a hole could lead to a mismeasurement of the shower energy. The CA consisted of two 1.5 cm wide square rings that framed the beam holes on the face of the calorimeter (see Figure 2.11). The rings were constructed from 3 tungsten–scintillator layers, for a total of 9.7 radiation lengths. Each segment of the rings was individually instrumented with a PMT and read out. If at least 14 GeV of in–time energy was deposited in any segment, a veto signal was sent to the Level 1 trigger.

The *back anti* (BA) was located on the face of MUF2, in the beam region. The purpose of the BA was to veto on particles that escaped down the beam holes. Thirty lead–scintillator layers (equivalent to about 30 radiation lengths or one nuclear interaction length) were divided longitudinally into three modules in order to distinguish between electromagnetic and hadronic particles based on the observed shower shape. The BA was not used in this analysis.

CHAPTER 3

The KTeV Trigger and Event Selection

As mentioned in Section 2.1.2, the neutral hadron rate in the KTeV beams was found to be as high as 50 MHz, with a 3:1 neutron to kaon ratio. Due to the long lifetime of the K_L , only about 5% decayed within the fiducial region of the detector (Quinn 2000). This still translates to a K_L decay rate on the order of 1 MHz – far too high for the data acquisition system (DAQ) to handle. For this reason, a three-tiered trigger was used to select events of interest. This system was able to reduce the event rate by three orders of magnitude, down to the order of 1 kHz.

3.1 Level 1

The Level 1 (L1) trigger system processed the fastest signals, usually PMT pulses, from the KTeV detector elements. L1 was clocked by the Tevatron RF, meaning that the system was ready to make a trigger decision every 19 ns bucket. For this reason, L1 is classified as a *deadtimeless* trigger.

The detectors used by L1 sent their trigger decision signals (also known as *sources*) to an array of programmable logic modules. Cable delays were employed to ensure that all of the sources arrived at the modules simultaneously. These modules held memory lookup tables, downloaded through a CAMAC interface at the start of each data collection run. Logical combinations of trigger sources were compared to the tables to determine if an event passed L1.

A total of 80 L1 sources were used, most coming from the detectors described in Chapter 2. The remaining sources require some extra description.

3.1.1 Beam timing sources

Two beam timing sources were transmitted from the Tevatron to L1. The first, known as *GATE*, was sent when the beam was on-spill. The second signal, *NEUTRINO PING* (NP), was turned on when beam was being delivered to E815, an experiment on an adjacent beamline. Due to occasional beam leakage into the NM beamline, the NP source was used as a veto.

3.1.2 Etototal

The *Etotal* system accepted signals from the final dynode stage of the calorimeter PMTs. The system had three functions:

- formation of the total analog sum of the dynode signals for all 3100 CsI channels. This sum represented the total amount of in-time energy deposited in the calorimeter during an event. The sum was sent to a multi-channel NIM discriminator for comparison to four preset thresholds. The logical outputs from the discriminator, named ET1–4, were used as L1 sources.
- generation of an *HCC bit* for a channel if the amount of in-time energy deposited in the channel exceeded a preset threshold. These bits were used as inputs to the Hardware Cluster Counter, a Level 2 processor described in Section 3.2.1.
- monitoring circuitry mounted on the Etototal boards allowed experimenters in the KTeV control room to view the dynode signal from any calorimeter PMT on a dedicated oscilloscope. This was used for diagnosing calorimeter

problems (bad PMTs, PMT bases, or poor connections to Etototal) without having to access the experiment.

There were three tiers to the Etototal system. Sixty *first level sum boards* were split into four custom, water-cooled VME crates. Each of these boards received up to 56 dynode inputs. The outputs included the HCC bits, a monitoring signal (if activated), and the analog sum over the input channels. These sums were sent to four *quadrant sum boards*, which then passed their output sums to the *full sum board*.

During the 1997 run, the four Etototal thresholds were set to 10, 18, 25, and 38 GeV. For 1999, ET1 was bumped up to 11 GeV, and ET2 was lowered to 16 GeV. Figure 3.1 shows a plot of these Etototal thresholds taken from the online monitoring system during the 1999 run. These thresholds tended to drift somewhat over time, as radiation damage to the calorimeter accumulated over the course of the run. Since this damage led to less scintillation light from the CsI crystals, over time the PMT pulse height dropped for a given energy, resulting in a gradual rise of the Etototal thresholds. Figure 3.2 shows the drift of the ET3 threshold through the 1999 run. The thresholds were periodically returned to their nominal values by either tuning the discrimination levels on the NIM module, or by increasing the high voltage on the calorimeter PMTs. The latter procedure was known as *gain matching*, and will be discussed further in Section 3.2.1. Variations in the values of the Etototal thresholds were modeled in the KTeV monte carlo, the subject of Chapter 5.

3.1.3 DC-ORs

The L1 trigger was able to use information from the two upstream spectrometer drift chambers, despite the lengthy 200 ns maximum drift time. This was done by recognizing that since a particle track always passes between two sense wires, one

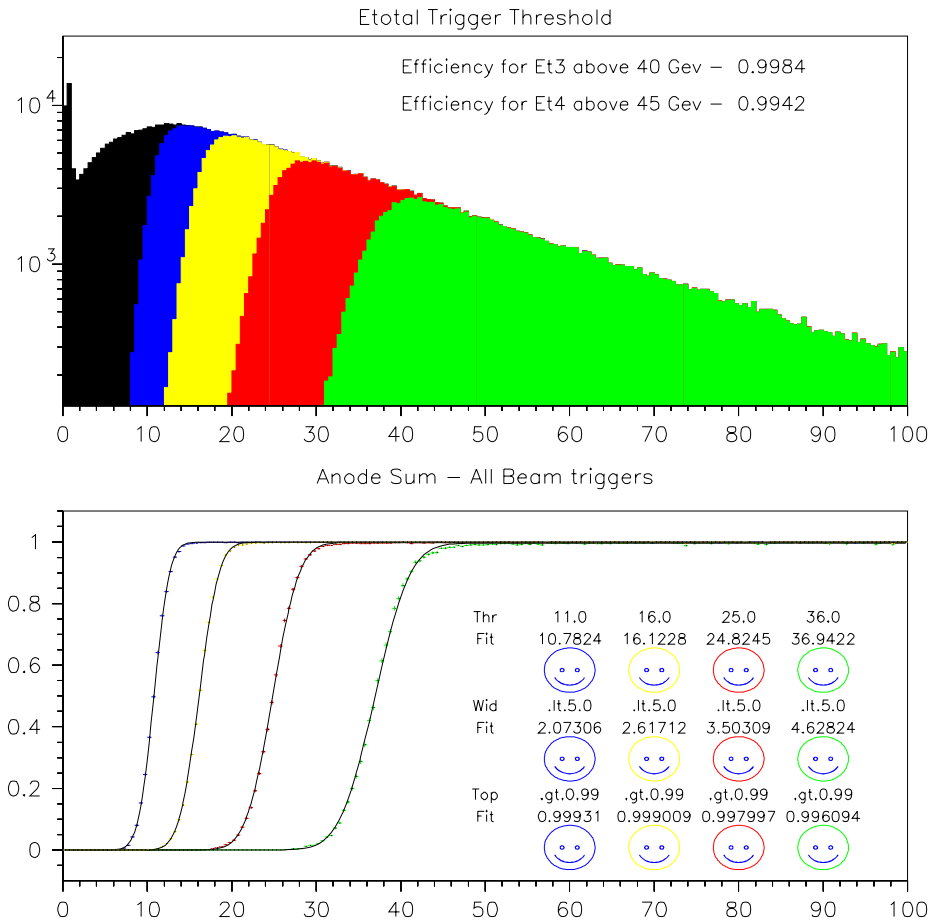


Figure 3.1: Top: The five curves shown represent (from left to right) the total energy distribution from all calibration events, and the total energy distributions of calibration events that fired the ET1, ET2, ET3, and ET4 Level 1 trigger sources. Energies are measured in GeV. Bottom: Threshold curves for L1 sources ET1–4, obtained by dividing each of the L1 source distributions by the distribution for all events. These plots were created from the online monitoring system during the 1999 run. The smiley faces indicate that the Ettotal thresholds, widths, and efficiencies are within tolerances.




Figure 3.2: ET3 threshold as a function of time during the 1999 run. The line depicts the nominal value of 25 GeV. Upward drift away from the nominal value was caused mainly by radiation damage to the calorimeter. The threshold could be returned to the nominal value by adjusting the discrimination level on the NIM module, or by raising the high voltage on the calorimeter PMTs. This variation was modeled in the KTeV monte carlo.

of the two resulting hits should arrive in less than half the maximum drift time. By taking the logical OR of the sense wires, hit information could then be passed along to L1 on a sufficiently fast timescale. These sources were therefore known as DC-ORs (drift chamber ORs).

The sense wires in the x and y views were grouped into 16 paddles. The logical OR of each paddle was taken and sent to a central controller, which counted the number of hit paddles in each view of a chamber. These counts were then passed along to L1. The DC-ORs were credited with cutting the L1 output rate by over a factor of two (Alexopoulos et al. 1994).

3.2 Level 2

The output rate from L1 averaged about 75 kHz. Elements of the Level 2 trigger (L2) utilized hardware processors to analyze the topology of an event and reduce the event rate to approximately 10 kHz. These processors required between 0.8 and $2.5 \mu\text{s}$ to analyze an event. While L2 was processing, new events accepted by L1 were inhibited, meaning that L2 caused the trigger system to be dead for a fraction of the time. Outputs from the processors were sent to the L2 state machine, which issued a decision on whether to abort (requiring another $0.5 \mu\text{s}$ to reset the L2 processors) or accept the event. In the case of an accept, the event was fully digitized and buffered for readout by Level 3 of the trigger.

The three L2 processors used in this analysis, the Hardware Cluster Counter, the DC Hit Counters, and the Y-Track Finder, are described below.

3.2.1 Hardware Cluster Counter

A count of the number of electromagnetic showers found in the calorimeter can provide important information about an event. For example, the ultra-rare direct

CP-violating mode $K_L \longrightarrow \pi^0 e^+ e^-$ contains four EM particles in the final state. Requiring at least¹ four showers in the trigger for this mode would significantly cut down the rate by excluding many other less interesting two-track decays.

The *Hardware Cluster Counter* (HCC) employed an innovative pattern recognition algorithm and parallel processing to calculate the number of EM showers in under 2 μ s (Bown et al. 1996). 3100 HCC bits were used as inputs. These bits and the clustering algorithm that uses them will now be discussed in detail.

HCC bits

As described in Section 3.1.2, an HCC bit was generated by the Etototal system for a calorimeter channel if the in-time energy deposited in the channel (measured at the final dynode stage of the channel's PMT) exceeded a preset threshold, nominally 1 GeV. Threshold decisions were made by comparator chips located on the first level Etototal sum boards. A typical HCC threshold profile is shown in Figure 3.3.

Thresholds could be adjusted on a channel-by-channel basis by changing the high voltage on the individual calorimeter PMTs. This procedure was called *gain matching*, and was performed every few months to counter gain loss from radiation damage to the calorimeter. Figure 3.4 shows the obvious improvement in the HCC threshold distribution for all 3100 calorimeter channels before and after gain matching.

Because the Etototal system set HCC bits based on the amount of *in-time* energy in the calorimeter channels, drift in the phase of the Tevatron RF signal had an effect on the measured HCC thresholds and turn-on widths. This is because lower energy dynode pulses cross the comparator threshold later than higher energy pulses. If the in-time window for a given channel is in the wrong place, low energy pulses wouldn't generate an HCC bit, effectively raising the value of the channel's

¹This flexibility allows for extra showers from hard bremsstrahlung photons or accidental activity.




Figure 3.3: Top: The open histogram is the distribution of in-time energy deposited in calorimeter channel number 1492 from all calibration electrons. The solid histogram is the in-time energy distribution for events in which the HCC bit for this channel fired. Energies are measured in GeV. Bottom: HCC threshold curve for calorimeter channel 1492, obtained by dividing the solid histogram by the open one. The turn-on region has a narrow width and is centered at approximately 1 GeV.

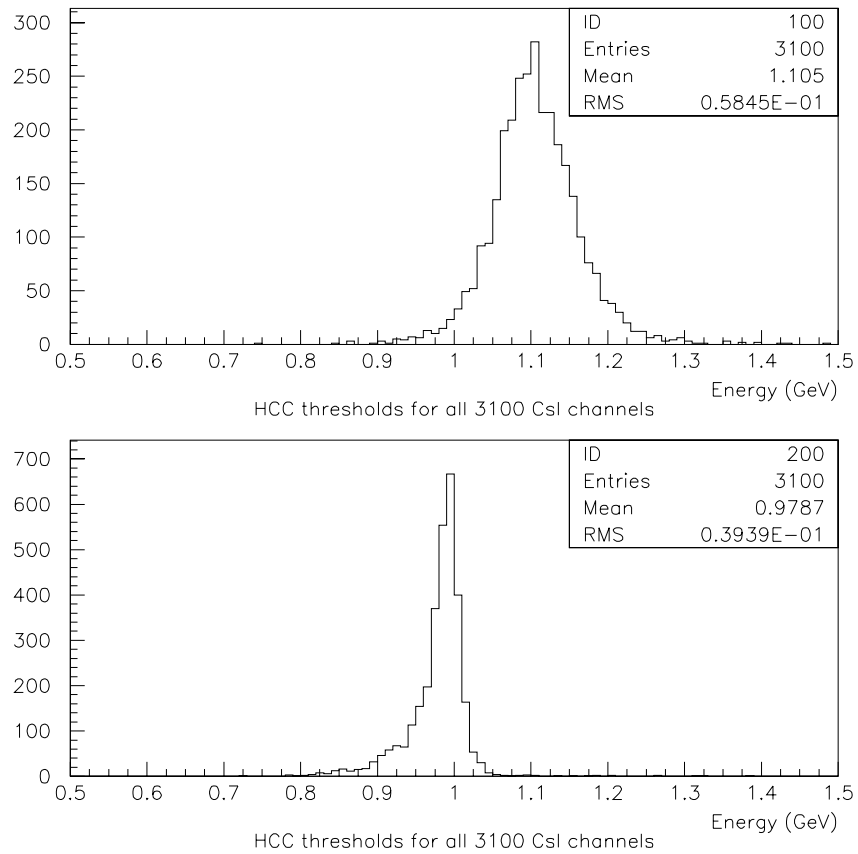


Figure 3.4: Distribution of all 3100 HCC bit thresholds before (top) and after (bottom) gain matching. Individually tuning the high voltages of the calorimeter PMTs has moved the mean of the distribution closer to the desired HCC threshold value of 1 GeV, and has also reduced the spread.

HCC threshold (see Figure 3.5). The window widths were dependent on individual channel characteristics, but were typically between 6 and 10 ns. The width was less than the full RF period of 19 ns to allow for setup and hold time requirements and signal propagation delays from electronic components on the first level sum boards. The relative location of the in-time window was dictated by the phase of the global Tevatron RF, which clocks L1, as well as by individual characteristics (PMT speed, dynode signal cable length, etc.) that varied from channel to channel. In addition, since the intrinsic width of the HCC turn-on curve for a given channel was caused by jitter in pulse arrival times, late dynode pulses also led to larger turn-on widths.

Unfortunately, the phase of the Tevatron RF, as measured at L1, drifted considerably over the course of a run, as much as ± 3 ns from its nominal value. Certain channels were very susceptible to these variations in terms of their HCC threshold and width, as shown in Figure 3.6. A list of the 160 most susceptible channels was compiled at the end of the 1997 run (Bown et al. 1997). These channels had their dynode signal cables trimmed between 1.5 to 4 ns, depending on the severity of the instability. This had the effect of shifting the in-time windows for these channels onto a plateau. Because of this, HCC threshold and width distributions were more stable in 1999.

Retriggering capability was also a concern. In some channels, very large dynode pulses would not return to ground by the start of the next RF bucket. This effect was due to stray capacitances that varied from channel to channel. If the large pulse were followed in the next bucket by a pulse barely above threshold, there was a chance that an HCC bit would not be set for the smaller signal. The HCC threshold values for these channels were therefore biased upwards: a 50 GeV pulse effectively increased the threshold of the following bucket by 1 GeV (Kabassian et al. 1994). The susceptible channels were located using a laser scan, and a simulation of this retriggering inefficiency was added to the monte carlo.

Shortly after the beginning of the Winter run, it was discovered that two

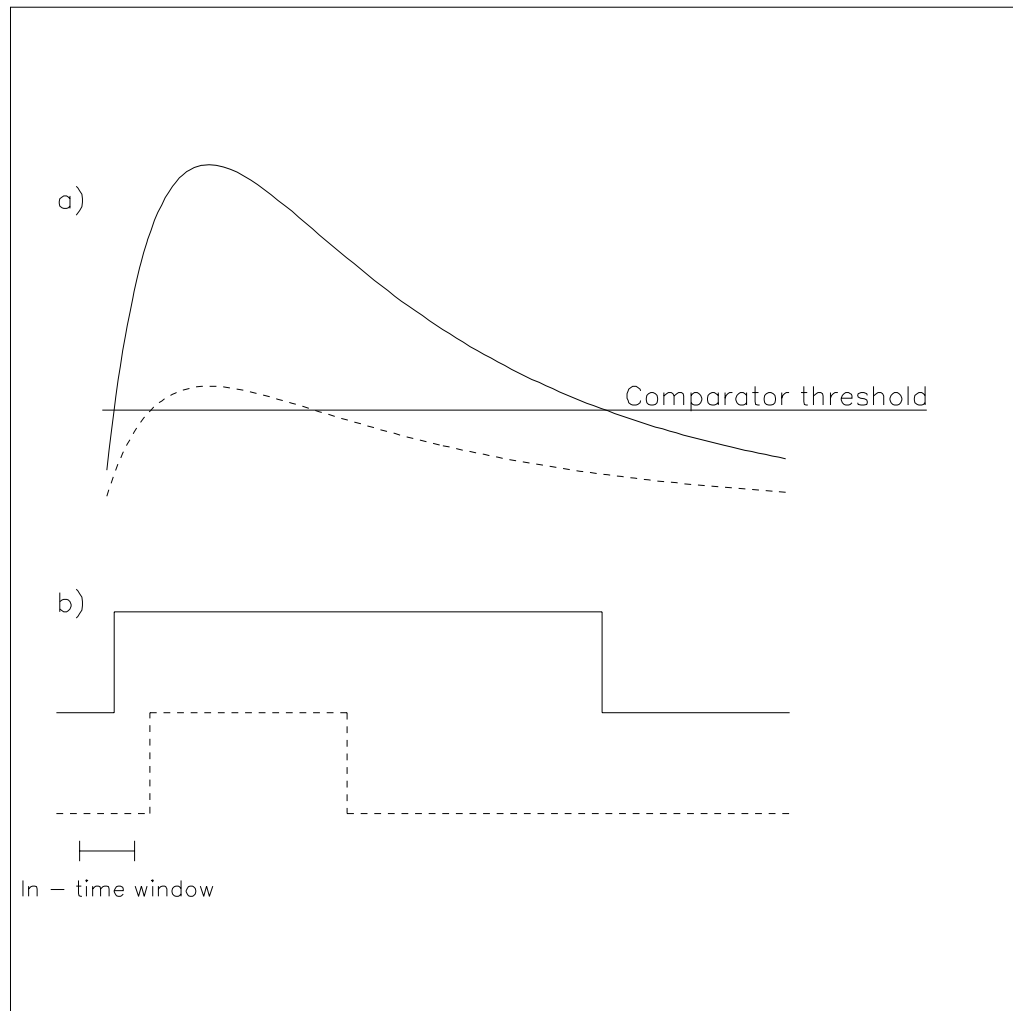


Figure 3.5: (a) Higher (solid) and lower (dashed) energy dynode pulses, as seen by a comparator on a first level Etot sum board. In this example, both pulses exceed the comparator threshold. Note that the higher energy pulse crosses threshold earlier. (b) Comparator output. Because the leading edge of the output from the higher energy pulse occurs during the in-time window, an HCC bit would be set for the channel. The lower energy output misses the in-time window, meaning no HCC bit would be set. The overall HCC threshold for this channel is therefore biased towards a higher value. The relative location of the in-time window can be shifted either by changing the global RF phase to L1, or by shortening the length of the dynode cable for the channel.




Figure 3.6: HCC thresholds and turn-on widths vs. RF phase as measured at L1. The nominal phase is 0 ns. Channel 844 (left) nominally sits in an unstable region – a small negative shift in the phase would lead to large shifts in HCC threshold and width. This channel had its dynode cable trimmed by 3 ns in order to move its nominal position onto the plateau. A normal channel, 1042, is shown for comparison (right). The nominal position of this channel is on the stable plateau.

cables carrying HCC bits from Etotol to the HCC were swapped. This resulted in erroneous HCC output for events in which one of the sixteen affected channels were part of a cluster. While this mistake was simulated in the monte carlo, data from this period (which amounted to less than 0.5% of the total 1997 dataset) were generally not used in E799 analyses.

During the Summer run, the water cooling system servicing the Etotol crates failed. This caused the first level sum boards to overheat, damaging many of the comparator chips. The damaged chips would output spurious signals, setting an HCC bit even if there was insufficient energy in the affected channel. These *hot bits* therefore biased the cluster count². About 5% of the channels were affected at any one time. These hot bits were simulated in the Summer monte carlo. The affected comparators were replaced before the start of the 1999 run, and a temperature interlock system was installed to prevent the Etotol system from overheating again. Because of these changes, the rate of hot bits was effectively zero during the 1999 run.

HCC clustering

The HCC bits from a shower formed clusters, as shown in Figure 3.7. By counting these clusters, the HCC was able to determine the number of EM showers. The counting was done by recognizing that four right turns are needed to traverse the perimeter of any cluster of rectangular blocks. The total number of right turns in the entire array is therefore equal to four times the total number of clusters. Right turns were quantified by breaking the array into 2×2 groups of blocks. All possible combinations of HCC bits in these groups are shown in Figure 3.8. The HCC processed all 2×2 groups in the calorimeter array, assigning each group a weight. The sum of these weights was then divided by four, and passed to the L2

²Since the cluster count was biased *upwards*, the hot bits problem did not lead to a loss of events. Because of this, the HCC was left in the trigger for the remainder of the run.




Figure 3.7: Examples of clusters from EM shower energy deposited in the calorimeter. Blocks with their HCC bit on are denoted by heavy black boxes. The two clusters found by the HCC are circled. The other activity in the calorimeter is out-of-time; therefore, these blocks do not have their HCC bits on, and no clusters are found there. The energy scale in this event display spans almost three orders of magnitude.


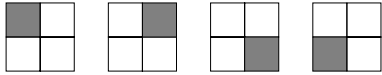


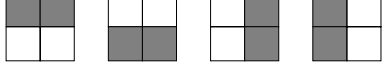
Hit block patterns	Pattern weight
	0: No turns
	+1: One right turn
	-1: One left turn
	+2: Two right turns
	0: No turns

Figure 3.8: Possible HCC bit patterns and corresponding weights for a 2×2 group of blocks. Filled blocks represent those with their HCC bit on. Note that the +2 patterns could also be viewed as having no right turns – the degeneracy was broken in favor of +2 in order to correctly count events in which two clusters touch at the corner.

trigger as a four-bit source (allowing for 8 clusters and an overflow bit).

Some regions of the array required special consideration. Groups straddling the large–small block boundary were handled by treating a large block as four small ones. The beam holes were considered to be filled with “ghost” blocks, needed to find clusters around the hole edges. The HCC bits for the ghost blocks were, of course, always off. A single row of ghost blocks also bordered the outside edge of the array, to allow for cluster finding along the outer perimeter of the calorimeter.

One problem with the HCC algorithm is that if a block with no HCC bit is surrounded by hit blocks, no cluster will be found. These “doughnut” patterns yield no clusters because the four right turns from the outer perimeter cancel out the four left turns from the inner perimeter. However, monte carlo studies of $K_L \rightarrow \pi^0 \pi^0$ decays have shown that the formation of doughnuts is extraordinarily rare. This inefficiency was also accounted for in the monte carlo.

3.2.2 DC Hit Counters (the fruit)

The DC Hit Counter system allowed the L2 trigger to quickly reject events with an insufficient number of charged hits. The system consisted of two types of processors: *Kumquats* and *Bananas* (known collectively as “the fruit”). The Kumquats collected wire hit information from the drift chamber plane pairs over a 205 ns gate (electron drift time plus cable delay), started by L1. The Bananas were much more complicated, using flash TDCs and memory lookup tables to determine whether wire hits were in–time. This was done by measuring the sum of drift distances (SODs) from hit pairs. SODs will be discussed in detail in Section 4.1. Only in–time wire hits were counted by the Banana modules.

The hit counting algorithm was the same for both processors. N sequential in–time hit wires in a plane pair were counted as $N - 1$ hits. Isolated hit wires were counted as 1 hit (see Figure 3.9). Taken together, the fruit required approximately

900 ns to process an event. Hit counts from both views of each chamber were sent to the L2 trigger as three-bit sources.

The Y-Track Finder (YTF) was a programmable logic unit used to identify hit patterns in the Y-views of the drift chambers. Inputs to the YTF came from the fruit. The processor searched for hits consistent with tracks in the upper half and lower half of the chambers. The YTF returned a four bit word containing the number and location of found tracks.

The architecture of the KTeV trigger system allowed for the definition of 16 higher-level *physics triggers*. These were logical combinations of L1 and L2 sources designed to select events from particular decay modes. The triggers used for this analysis were changed between the 1997 and 1999 runs; these changes are described below.

The *dimuon* trigger was a physics trigger constructed to search for $K_L \rightarrow e^+e^-\mu^+\mu^-$, the direct-CP violating decay $K_L \rightarrow \pi^0\mu^+\mu^-$, and the muonic Dalitz

decay $K_L \longrightarrow \mu^+ \mu^- \gamma$. In 1997, this trigger was written as:

$$GATE * \overline{NP} * 2V * DC12 * 2MU3 * \overline{PHV1} * 2HCY_LOOSE * HCCGE1.$$

The level 1 trigger elements required:

- the event occurs during the on-spill period (GATE)
- veto on NEUTRINO PING signal (\overline{NP})
- ≥ 2 hits in V and ≥ 1 hit in V' , or vice versa (2V)
- ≥ 1 hit from the DC-ORs in both views of drift chambers 1 and 2 (DC12)
- ≥ 2 hits in both views of MU3 (2MU3)
- veto on hits in RC, SA, or CIA ($\overline{PHV1}$)

while level 2 required:

- ≥ 2 hits in y view of chambers 3 and 4, with ≥ 1 y hit in chamber 1 and ≥ 2 y hits in chamber 2 or vice versa, from the fruit (2HCY_LOOSE)
- ≥ 1 HCC cluster (HCCGE1)

Normalization events were collected for the branching ratio calculation and systematic checks. The decay mode chosen for normalization in this analysis is $K_L \longrightarrow \pi^+ \pi^- \pi_D^0$,³ due to the similarity in final state topology to $K_L \longrightarrow e^+ e^- \mu^+ \mu^-$. During the 1997 runs, these events were collected on the *two-track* trigger:

$$GATE * \overline{NP} * 2V * DC12 * 2HCY_LOOSE.$$

The two-track trigger had a prescale factor of 500 applied at L2 to control the rate.

³Here π_D^0 denotes the Dalitz decay $\pi_D^0 \longrightarrow e^+ e^- \gamma$

3.3.2 Physics triggers in 1999

The most restrictive element of the dimuon trigger in 1997 was the $2MU3$ requirement, which rejected events in which both muons hit the same paddle in a single view of the MU3 bank (even if they were well separated in the other view). In order to increase the acceptance for $K_L \rightarrow e^+e^-\mu^+\mu^-$ in 1999, the $2MU3$ requirement was changed to $2MU3_LOOSE$, which allowed for a single missing hit in MU3. Of course, this change alone would significantly increase the rate of the dimuon trigger. To counter this effect, the number of required HCC clusters was raised⁴ to ≥ 2 . The dimuon trigger in 1999 therefore read⁵:

$$GATE * 2V * DC12 * 2MU3_LOOSE * \overline{PHV1} * 2HCY_LOOSE * HCCGE2.$$

Taken together with the lower P_t kick from the analysis magnet described in Section 2.2.2, this trigger change resulted in an acceptance increase for $K_L \rightarrow e^+e^-\mu^+\mu^-$ of about 30% (Hamm 1999).

Due to an unfortunate error, $K_L \rightarrow \pi^+\pi^-\pi_D^0$ normalization events were taken from the *four-track* trigger in 1999, instead of the two-track trigger. The four-track trigger was designed to collect $K_L \rightarrow \pi^+\pi^-e^+e^-$ events, and contained several more restrictive elements:

$$GATE * 3V_TIGHT * 2DC12_MED * ET1 * \overline{MU2} * \overline{PHV1} * \overline{CA} * 34_HCY * \\ HCC_GE2 * YTF_UDO * 3HC2X.$$

The new elements are:

- ≥ 3 hits in V and ≥ 3 hits in V' ($3V_TIGHT$)
- ≥ 2 hits from the DC-ORs in 3 of 4 views of drift chambers 1 and 2, ≥ 1 hit in the other view ($2DC12_MED$)

⁴This change was possible because $K_L \rightarrow \mu^+\mu^-\gamma$ was no longer of interest in 1999.

⁵The NP source was also removed in 1999, as E815 wasn't running at this time.

- total energy deposited in calorimeter ≥ 11 GeV (ET1)
- veto on hits in MU2 ($\overline{MU2}$)
- veto on hits in CA (\overline{CA})
- ≥ 3 y view hits in drift chambers 1 and 2, with ≥ 4 y hits in chambers 3 and 4, from the fruit (34_HCY)
- ≥ 3 x view hits in drift chamber 2, from the fruit (3HC2X)
- ≥ 1 track in upper and lower halves of drift chambers, from the YTF (YTF_UDO)

The strict nature of this trigger will necessitate an extra systematic study. This will be presented in Chapter 6.

3.4 Calibration Triggers

Another set of 14 higher-level triggers was reserved for collecting events used to calibrate the various KTeV detectors. The quiescent laser flashes, laser scan flashes, and muon run events described earlier were all accepted by calibration triggers. Additional triggers collected DPMT and ADC pedestal events (to record noise levels for pedestal subtraction during event reconstruction), as well as cosmic ray tracks through the calorimeter (to measure uniformity of the scintillation response along the crystal axes).

3.5 Digitization and Readout

Upon receiving an accept signal from L2, analog signals from the TRDs and photon vetos were digitized by fast encode and readout (FERA) ADC modules. During the 1997 run, 12-bit ADCs were used. These were replaced by 8-bit modules in

1999 to cut the required readout time. Latch modules were used to record bitmasks of which L1 and L2 sources, beam triggers, and calibration triggers fired during an event. The Fermilab-designed “Deep Latches” used for this purpose during the 1997 run were capable of latching sixteen time slices of four 8-bit data words in a FIFO buffer, and therefore required a significant amount of time for readout. Before the 1999 run they were replaced by University of Arizona-designed “AZLatches”, which were specifically built for the KTeV readout system. These modules latched only one time slice of three 8-bit data words (Cheu et al. 1999). This resulted in a drop of approximately 25% in the total trigger readout time. Sixteen-bit scaler modules assigned spill and event numbers to the event at this stage.

A fully digitized event, consisting of ADC, TDC, scaler, latch, and pipelined CsI information, was then sent to one in a bank of four VME memory modules. These modules contained enough memory to buffer all of the events from an entire spill: a total of 4.6 Gb in 1997, upgraded to almost 4.9 Gb for 1999 (LaDue). Each module belonged to a *DAQ plane*. Three of the four planes were devoted to physics triggers, while the fourth was used for calibration triggers and monitoring events⁶. The modules were read out over the course of a spill by four SGI Challenge computers, one per plane. These computers ran software that constituted the third level of the KTeV trigger.

3.6 Level 3

The Level 3 (L3) trigger consisted of a library of optimized software routines used to find tracks, clusters, and vertices in order to fully reconstruct an event. These reconstruction algorithms will be described in detail in Chapter 4. Optimization of the code was crucial – if L3 failed to finish analyzing all of the events from a spill, new triggers were inhibited, leading to trigger deadtime at the start of the following

⁶Monitoring events were simply prescaled copies of physics events, used for online monitoring of the detectors during a run.

spill. CPUs were selected for the Challenge machines to help ensure that all the events could be processed in time, meaning virtually no deadtime resulted from L3 at nominal beam intensities.

If an event satisfied a predefined set of conditions, it was assigned a *L3 tag*. The tags used in this analysis are:

- $e^+e^-\mu^+\mu^-$, assigned to events in the dimuon trigger that passed a “fast” four-track vertexing algorithm, which searched for at least three tracks that form a common vertex.
- *four-track minimum bias*, assigned to events in the two-track trigger that satisfied the fast four-track vertexing algorithm. These were usually $K_L \rightarrow \pi^+\pi^-\pi_D^0$ decays. Events with this tag were used for normalization in 1997.
- *four-track prescaled*, assigned to events in the four-track trigger that passed the fast four-track vertexing algorithm. As mentioned earlier, these tags were accidentally selected for normalization during the 1999 run. These events were initially prescaled by a factor of 10 at this stage; in order to cut the output rate, the prescale was raised to 20 about 5% of the way into the run.

Approximately 15% of input events passed L3 and were written to DLT storage tapes for analysis. Table 3.1 summarizes the trigger accept rates at L1, L2, and L3 for a typical 1999 beam intensity of 8×10^{12} protons on target per spill. The trigger system was live approximately 65% of the time at this intensity. Most of the deadtime was attributed to the HCC (the slowest L2 processor) and readout.

3.7 Online Split

During the 1997 run, a total of ten DLT tape drives were used for recording events. Physics events were written in round-robin fashion to nine of the drives (three per

	Dimuon trigger only (kHz)	All physics triggers (kHz)
L1 accept	8	74
L2 accept	0.4	8
L3 accept	0.1	1

Table 3.1: Trigger accept rates for a typical beam intensity (8×10^{12} protons on target per spill) during the 1999 run. At this intensity the trigger was live around 65% of the time.

plane). All calibration events were sent to the remaining drive. While this system was simple, it required a significant amount of work (about six months!) to go back through all the tapes after the run and split off events into data sets for calibration and analysis. The offline split produced a total of 20 split tapes for 1997 dimuon analyses.

In order to eliminate this lag time for the 1999 data, an innovative *online split* system was implemented. Copies of events were sent to different disk caches on the Challenge machines, based on their L3 tags. When the size of a cache approached the capacity of a DLT tape (about 20 Gb), the online split system would alert the shift crew to mount a pre-labeled blank tape into a free drive. After the cache had been successfully written to the split tape, it was deleted from disk. Data was also written to the ten drives, as in 1997, in case the online split failed (fortunately, an unnecessary precaution).

Thanks to the online split system, data sets were available for calibration and analysis immediately after the 1999 run. A total of 59 split tapes were written from the dimuon cache, which included $e^+e^-\mu^+\mu^-$, $\pi^0\mu^+\mu^-$, and normalization tags.

CHAPTER 4

Event Reconstruction

Each event contained about 6.5 kB of data. In its raw form (ADC, TDC, and DPMT counts), this data was of little use. In order to gain any material information about an event, the data had to be reconstructed into physically significant quantities, such as charged track momenta and calorimeter cluster energies. This reconstruction process consisted of three stages of software algorithms: tracking, clustering, and vertexing. The routines described in this chapter are from Version 5.04 of the KTeVANA code library.

4.1 Tracking

The first stage of event reconstruction was the determination of tracks from charged particles. The tracking algorithm used TDC information from hits in the spectrometer drift chambers. As mentioned in Section 2.2.2, maps were used to translate TDC counts from hit times into drift distances. Hits were classified as being in-time if they fell within a 235 ns wide window. If there were multiple hits on a wire, only the first in-time hit was used.

4.1.1 Good SOD, Bad SOD

If the charged particle also left a hit in the complementary plane of the plane pair, the two hits were grouped together as a *hit pair*. The sum of the electron drift distances from the two hits of the hit pair should equal the 6.35 mm cell spacing of

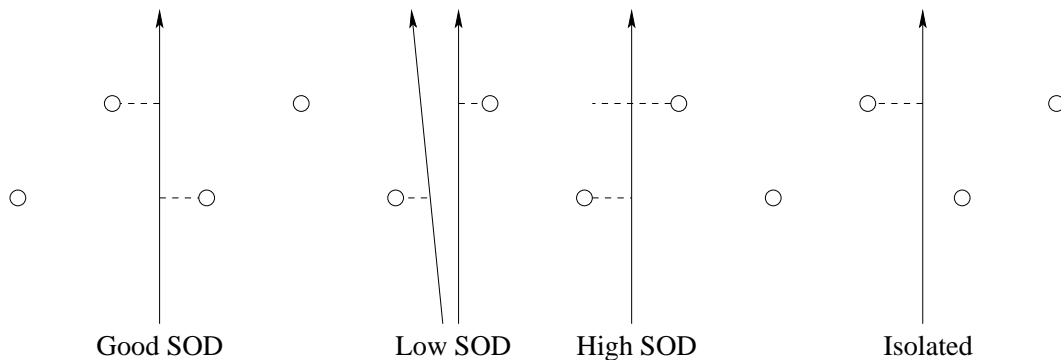


Figure 4.1: Categories of SODs from a charged track. Open circles represent sense wires of the plane pair, while the dashed lines indicate the inferred trajectories of the drift electrons.

the drift chamber sense wires (see Figure 2.4). This *sum of distances* (SOD) was used to classify the quality of the hit pair. Hit pairs with SODs within 1 mm of the nominal value were classified as having “good” SODs.

Unfortunately, it was possible for perfectly good tracks to leave bad SODs. Because this led directly to tracking inefficiencies, it is important to understand the pathologies of bad SODs, pictured in Figure 4.1. Low SODs were typically the result of two tracks occupying the same cell. In this case, both sense wires would register two in-time hits. However, since only the *first* in-time hit was used, the sum of the drift distances would be too small. The second track could be from the decay, a δ -ray, or from accidental activity. High SODs were caused by inefficiencies in detecting the earliest electrons from the drift cloud. These inefficiencies resulted from radiation damage to the chambers, insufficient signal gain, and noise from the chamber electronics. Isolated hits, by definition, had no measurable SOD. These were most probably caused by defects on the sense wires. Despite having no SOD, isolated hits were accepted by the tracking algorithm. Figure 4.2 shows the distribution of SODs from all plane pairs in the 1999 normalization data.




Figure 4.2: Distribution of SODs from all plane pairs in 1999 $\pi^+\pi^-\pi^0$ data. Not shown are isolated hits (which were assigned a SOD of 0 m). Note that the mean of the distribution is consistent with the drift chamber sense wire spacing of 6.35 mm. SODs within 1 mm of this nominal value were considered “good”. Tracks were allowed to contain at most two bad SODs, or one bad SOD and one isolated hit.

4.1.2 y track finding

Track candidates were found separately in the x and y views. Tracks in the y view were not affected by the momentum kick from the analysis magnet, so their hit pairs should lie in a straight line through all four drift chambers (although chamber alignment and fringe field effects necessitated some corrections, to be discussed in Section 4.3). Track candidates in the y view were allowed to contain at most two hit pairs with bad SODs, or one bad SOD and one isolated hit. A track χ^2 was calculated for each candidate to describe how well a track fit the hit pairs (recall from Section 2.2.2 that the hit position resolution was about $100\ \mu\text{m}$). Only candidates with a track $\chi^2 < 4$ were used. The algorithm fit all possible y tracks in this manner. If fewer than two y track candidates were found that didn't share any hits, the event was rejected.

4.1.3 x track finding

Because of the field orientation from the analysis magnet, tracking in the x view was more complicated. Segments were found separately upstream and downstream of the magnet plane. Upstream segments (found using drift chambers 1 and 2) were allowed to contain two bad SOD pairs, or one isolated hit. Downstream segments (DC3 and 4) were allowed only one bad SOD pair or one isolated hit. The tracking algorithm looped over all possible pairs of upstream and downstream segments, scanning for pairs that intersected at the magnet plane. A 6 mm “miss distance” was allowed in order to accomodate resolution effects. Joined segments were allowed to have at most two bad SOD pairs between them, or one bad SOD and one isolated hit. The event was rejected if less than two x track candidates were found that didn't share any hits.

4.2 Clustering

The clustering algorithm was used to reconstruct energy deposits in the calorimeter. The process consisted of four stages: hardware clustering, software clustering, cluster positioning, and corrections.

4.2.1 Hardware clustering

The first stage was known as *hardware* clustering because it made use of the HCC bits generated by the Etototal trigger. The algorithm began by looping over all crystals with their HCC bit turned on, searching for local energy maxima. Because the Etototal system set HCC bits based on the amount of in-time energy in the calorimeter channels, out-of-time clusters from accidental activity were suppressed. If a local energy maximum was more than 0.1 GeV^1 , the channel was designated as a hardware cluster *seed*. The raw cluster energy was calculated by summing up the energies from channels in a grid centered on the seed block. The grid size depended on the size of the crystals: 7×7 channels for small blocks, and 3×3 for large blocks. At the large-small block boundary the 3×3 grid size was used, and four small crystals were considered as a single large crystal.

4.2.2 Software clustering

Minimally ionizing particles (MIPs) such as muons and pions also left energy deposits in the calorimeter. Typically, the amount of energy would be too small to turn on any HCC bits, so no MIP signatures would be found during the hardware clustering process. The *software* clustering algorithm searched for seeds containing in-time energy over a certain preset threshold, nominally 50 MeV. Seeds already

¹This seemingly redundant requirement was necessary to ensure that hot HCC bits were not found to be cluster seeds.

found during hardware clustering were excluded. Raw cluster energies were calculated using the same grids described above. Software clusters were retained only if their raw in-time energies exceeded a threshold of 100 MeV.²

4.2.3 Cluster positioning

In most cases, the energy of a cluster was spread out over several blocks. It was therefore not a trivial task to determine exactly where the incoming particle struck the calorimeter. Cluster positions were located by calculating the ratios of total energy in the 1×3 block row and 3×1 column containing the seed block to the energies of adjacent rows and columns. These ratios were compared to values in look-up tables, which returned the x and y coordinates of the cluster position on the calorimeter face. The look-up tables were generated by studying $K_L \rightarrow \pi^0 \pi^0$ data and monte carlo, and were cross-checked using calibration K_{e3} events. The tables were binned by cluster energy, and separate tables were used for large and small block regions of the calorimeter. This method resulted in a position resolution of approximately 1 mm for clusters in small blocks, and approximately 1.8 mm in large blocks (Shawhan 2000).

4.2.4 Clustering corrections

At this stage, the raw cluster energies were corrected to account for a myriad of second-order effects. Two types of corrections were applied: block-by-block corrections modified the energies of individual channels within a cluster, while multiplicative corrections modified the energy of the cluster as a whole.

²Recall that MIPs traveling parallel to the crystal axis deposited approximately 320 MeV of energy in the crystal.




Figure 4.3: A situation in which the overlap clustering correction is necessary. Two clusters (their calculated positions circled) share blocks between them. Block energies are shown in GeV. Blocks with their HCC bit on are denoted by heavy black boxes.

Block-by-block corrections

The *overlap* correction was needed to separate the energies of clusters that shared crystals (see Figure 4.3). Energies of all common crystals were divided between the two clusters, and the cluster energies and positions were recalculated. This process was iterated either until these quantities converged (cluster energies constant to within 5 MeV, x and y positions to within 100 μm), or the maximum number of iterations (20) was reached.

It was also possible for neighboring clusters to bias raw cluster energy measurements without sharing blocks. This would happen if energy from the neighbor spilled into the grid used for the raw energy calculation for the cluster of interest

(see Figure 4.4). This bias would become quite significant if a very high energy cluster neighbored a low energy cluster. The *neighbor* correction was used to subtract off any energy leakage. Studies of shower shapes from GEANT simulations were used to calculate a mean transverse shower profile, covering a 13×13 array of small blocks. The algorithm looped over all pairs of clusters, using the GEANT profile to predict the amount of energy leakage from one cluster into its neighbor’s grid. This excess was then subtracted from the neighbor’s raw energy.

Clusters bordering a calorimeter edge or a beam hole required a correction to account for lost energy (see Figure 4.5). The *missing block* correction used the same transverse shower distribution as the neighbor correction to estimate the amount of lost energy in the “ghost blocks” that would have otherwise participated in the cluster’s energy measurement grid. Resumming over block energies in the grid then provided a more accurate value for the cluster energy.

Oftentimes, some of this “lost” energy reached the crystals on the far side of the beam hole. If another cluster was located on the far side, its energy measurement would be biased upwards by this “sneaky” energy (see Figure 4.6). This effect was accounted for by applying the *sneaky energy* correction. Correction factors were tabulated by analyzing calibration K_{e3} events in which the electron cluster was situated near a beam hole. Sneaky energy fractions were measured in 140 blocks in a C-shaped region around the opposite side of the hole. These fractions could then be used to predict the amount of sneaky energy residing in each channel of a cluster.

Recall from Section 2.2.5 that as data from the DPMTs were pipelined, channels were removed from the readout list if the channel energies were less than a given threshold, a process known as *sparsification*. Because some of these channels might have been on the outer edges of a cluster, removing them would have the effect of skewing the cluster energy downwards. The *threshold* correction was used to counteract this bias by adding E_{thr} amount of energy to channels in a cluster,




Figure 4.4: A situation requiring the neighbor clustering correction. The clusters (calculated positions circled) do not share any blocks, so the overlap correction does not apply. However, the bottom cluster spills into the 7×7 grid used to calculate the raw energy of the top cluster. Block energies are shown in GeV. Blocks with their HCC bit on are denoted by heavy black boxes.




Figure 4.5: An example of a cluster requiring the missing block correction, due to energy lost in a calorimeter beam hole (the white, vacant region). Block energies are shown in GeV.



Figure 4.6: A situation requiring the sneaky energy cluster correction. Energy from the left cluster sneaks across the beam hole to appear underneath the right cluster, and vice-versa. Block energies are shown in GeV.

where E_{thr} is given by

$$E_{thr} = P_1(r) + P_2(r) \times \log(E). \quad (4.1)$$

P_1 and P_2 depend on the distance of the block from the cluster's center, r , and were measured during a special calibration run with no sparsification requirement.

Multiplicative corrections

While mylar wrappings were successful in reducing non-uniformities in the longitudinal scintillation response along the axes of the CsI crystals, transverse response of the crystals varied significantly. Because of this effect, the raw energy of a cluster depended upon where the particle struck the face of the cluster's seed block. This was rectified by applying the *intra-block* correction. In order to generate this correction, every block in the calorimeter array was divided into 25 transverse position bins. The ratio E/p was measured in every bin from calibration K_{e3} electrons. This data was used to generate a map of transverse scintillation response for each crystal. The map provided a multiplicative factor to correct a cluster's energy for variations in the response over the face of its seed block.

The *non-linearity* correction was applied to remove residual variations in longitudinal scintillation response. Correction factors were estimated by convoluting GEANT generated showers with the measured longitudinal response in each channel. Because shower depth depended on the type and energy of incident particle, different corrections were generated for electrons and photons, in energy bins of 2, 4, 8, 16, 32, and 64 GeV.

Even after applying all of the corrections described above, residual non-linearity was apparent in the calorimeter response to electrons. In a perfectly linear calorimeter, the mean value of the ratio E/p would be expected to remain constant over all ranges of p ; however, variations of a few tenths of a percent were observed (see Figure 4.7 (Roodman 1998)). These E/p vs. p distributions were used to




Figure 4.7: Mean E/p vs. p of calibration K_{e3} electrons in small blocks, after all clustering corrections except the linearity fudge. Residual non-linearities cause the mean to drift over the range of p . The linearity fudge uses this distribution to cancel out these residual effects.

generate a *linearity fudge*, applied after all other corrections. This correction had the effect of canceling out the residual non-linearities. Separate fudges were used for large and small blocks. The final energy resolution of the calorimeter was given in Equation 2.2.

4.3 Vertexing

At this stage, the charged tracks were extrapolated upstream to see if they formed a common vertex, i.e., the K_L decay point. The vertexing algorithm first looped over all combinations of y tracks, searching for common intersections (y view vertex candidates). A 2 mm miss distance was allowed for tracks at the intersection points, which were required to be within the fiducial decay volume. Minimum and maximum possible z values for the y view vertex candidates were also calculated, using the uncertainties on the fitted track parameters (slope and intercept). Sharing of two y tracks was allowed, meaning only two intersecting y tracks were required for each candidate. The event was thrown out if no y vertex candidates were found.

The algorithm then looped over all combinations of upstream x track segments, searching for x view vertex candidates containing at least four tracks. As was the case in the y view, a 2 mm tolerance was allowed for tracks at the intersections. Ranges of possible z positions for the x view vertex candidates were calculated based on the uncertainties in the fitted upstream track segment parameters. If no x view vertex candidates were found, the event was discarded.

At this point, y view tracks were associated with those in the x view. This process was complicated by ambiguity: two tracks might be well separated in the x view but overlap in the y view, and vice versa. This ambiguity was resolved using a procedure known as *track-cluster matching*. For each track, the matching algorithm looped over the list of all calorimeter clusters, measuring the distance between a cluster’s position and the projected track position at the calorimeter face. If the minimum separation was less than 7 cm, the track was said to be “matched” to the cluster that minimized the separation distance. x view tracks could then be associated with y view tracks if both were matched to the same cluster, and if their projections were within 1.5 cm of each other in x and y on the calorimeter face. Tracks that traveled down a beam hole, leaving no cluster, were “matched” to the beam hole for the purposes of track association. In this case, the maximum allowable track-cluster separation distance of 7 cm was assigned to the match. Therefore, for tracks just barely inside the beam hole, any match to a cluster near the hole took precedence over the beam hole match. Clusters not matched to tracks were listed as “extra” clusters, and were considered as photon candidates during particle identification (Section 6.4).

A list of final vertex candidates was compiled by comparing the possible z ranges between x view and y view candidates, searching for overlaps. If an overlap was found, several corrections were applied to the hit pairs used by the candidate. First, the hit positions were recalculated to account for signal transit time from the drift chamber sense wire to its discriminator. Also, because the two planes of a plane pair were offset in z , a correction was needed to remove the dependence of

the SOD on the angle of the incoming track relative to the alignment of the drift chamber cell. Another SOD correction was needed to account for the alignment of the drift chambers with respect to each other. A final correction was used to cancel the effect of fringe fields from the analysis magnet, which extended out to drift chambers 2 and 3.

After these corrections were applied, track and vertex parameters were recalculated. Two variables that described the quality of a vertex candidate were also calculated at this point. The first, vertex χ^2 , is based on how well the upstream track segments intersect at the candidate vertex position. The second, offmag χ^2 , describes how well the upstream and downstream track segments meet at the plane of the analysis magnet. A comprehensive quality variable for vertex candidates was formed by combining the vertex χ^2 , the offmag χ^2 , and the total number of bad SODs used. The candidate with the smallest value of this variable was selected as the final vertex.

The last step was the momentum calculation for the tracks used in the final vertex. Using the fitted x view track parameters, the bend angle at the plane of the analysis magnet was computed. The track momentum is proportional to the magnitude of the magnet kick (already known from calibration) divided by the bend angle. The momentum resolution of the KTeV spectrometer was given in Equation 2.1.

Figures 4.8 and 4.9 show event displays of a fully reconstructed $K_L \rightarrow e^+e^-\mu^+\mu^-$ event from the data. Hardware and software clusters, tracks, and a vertex are all pictured.




Figure 4.8: Calorimeter view of a $K_L \rightarrow e^+e^-\mu^+\mu^-$ event after reconstruction. Four clusters are circled: two hardware clusters from the electrons, and two software clusters from the muons. The vertex position (diamond within the right beam hole) and upstream/downstream track segments (lines) are projected to the calorimeter face. Note that in this particular event, all tracks are matched to clusters.




Figure 4.9: y view of the same $K_L \rightarrow e^+ e^- \mu^+ \mu^-$ event after reconstruction. Scales are given in meters. Detector elements are labeled in Figure 2.3. The upstream and downstream track segments pass through hit pairs in the drift chambers, and clearly meet at the plane of the analysis magnet. The four-track vertex is located at 122 m in z . Note that the muon tracks (left) point to hits in the muon banks at the far downstream end of the detector. The electron tracks (right) end at the calorimeter, where the electrons lose all of their energy to EM showers in the CsI.

CHAPTER 5

KTeV Monte Carlo

Simulated datasets generated using monte carlo statistical techniques (MC for short) were essential for the successful completion of this analysis. Monte carlo served three functions:

- Because of the geometry, efficiency, and resolution of the detectors, many kaon decays were lost during data collection. Additionally, cuts used to reduce background invariably tossed out a number of otherwise good events. Monte carlo was used to determine the overall *acceptance*¹ of the normalization and signal modes in this analysis. These acceptance numbers were needed to measure the $K_L \longrightarrow e^+e^-\mu^+\mu^-$ branching ratio and form factor.
- Monte carlo was used to estimate the number of background events remaining in the signal data after all cuts. These events must be subtracted off before the calculation of the branching ratio.
- Comparisons between data and MC were used to ensure that the geometry, efficiency, and resolution of the KTeV detector systems were well understood. In the case of a discrepancy, the monte carlo was used to estimate the magnitude of the systematic effect.

This chapter describes the process of KTeV monte carlo generation.

¹The acceptance is defined as the number of events reconstructed in a given decay mode after all cuts, divided by the total number of decays that actually occurred in that mode. Acceptances at KTeV are typically on the order of a few percent, depending on the mode in question.

5.1 Simulation of K_L Production

Generation of monte carlo events began with $K^0 - \overline{K}^0$ production at the BeO target. The particle production cross section for protons on a beryllium target was parameterized by Malensek, and reads

$$\frac{d^2N}{dP d\Omega} = K P \frac{(1-x)^A (1+5e^{-Dx})}{(1+P_t^2/M^2)^4}, \quad (5.1)$$

where P is the lab frame momentum of the produced particles, x is the ratio of P to the momentum of the protons incident on the target, and P_t is the transverse momentum of the produced particles with respect to the incoming protons (Malensek 1982). K , A , D , and M are constants determined by fitting to data from 400 GeV/c protons. Only measurements of charged kaon production have been performed. However, by counting valence and sea quark combinations, a reasonable guess could be made at the magnitude of the neutral kaon production probability, σ :

$$\sigma(K^0) \approx \frac{\sigma(K^+) + \sigma(K^-)}{2} \quad (5.2)$$

$$\sigma(\overline{K}^0) \approx \sigma(K^-). \quad (5.3)$$

These relations predict a mix of approximately 55% K^0 , 45% \overline{K}^0 from the target. The parameters in Equation 5.1 were tuned to match the momentum spectrum from actual $K_L \rightarrow \pi^+\pi^-$ decays measured at KTeV.

Once generated, the kaon was projected to its decay position. If the kaon trajectory intersected the primary beam collimator, tracing was terminated, and a new kaon was generated. Kaons striking the defining collimator were either terminated or allowed to scatter. Scattering was also simulated throughout the beam absorbers in the NM2 enclosure.

The initial K^0/\overline{K}^0 state was evolved up to the point of decay. The monte carlo kaons used in this analysis were generated over a momentum range of 20 to

220 GeV/c, and were forced to decay between 90 and 160 m downstream of the target. Therefore, except for very high energy kaons, the final state was almost completely K_L within the allowed decay region. Detector acceptance for decays outside of these ranges was negligible.

5.2 Simulation of K_L Decay

Different decay generators were used for different kaon decay modes. For some modes, the monte carlo even offered a choice of generators from competing theories. The generators used in this analysis are described below.

5.2.1 $K_L \longrightarrow e^+e^-\mu^+\mu^-$ generator

Signal events were generated using the DIP form factor (Equation 1.32). For the branching ratio analysis presented in Chapter 6, the value of α was fixed at -4.53 , the central value of the result from the form factor analysis to be described in Section 7.2, while β was fixed at 0. Of course, there is some uncertainty associated with the choice of these values, leading to a systematic error on the signal acceptance. This will be discussed in the following chapter.

The DIP monte carlo contains radiative corrections of $\mathcal{O}(\alpha^3)$, requiring calculation of quite a few additional diagrams (some of which are shown in Figure 5.1). These corrections include contributions from inner bremsstrahlung, vacuum polarization, and pentagon graphs. An infrared cutoff on the energy of the radiated photon is imposed at 400 keV to prevent any divergences (Barker), (Toale et al. 2002).




Figure 5.1: Examples of $\mathcal{O}(\alpha^3)$ radiative corrections to $K_L \rightarrow e^+e^-\mu^+\mu^-$ that are included in the KTeV monte carlo.

5.2.2 $K_L \longrightarrow \pi^+\pi^-\pi^0$ generator

The decay to three pions was the first step in the generation of normalization mode monte carlo. The matrix element is proportional to

$$|M|^2 \propto 1 + gu + hu^2 + jv + kv^2, \quad (5.4)$$

where

$$u = \frac{(s_3 - s_0)}{M_\pi^2}, \quad v = \frac{(s_1 - s_2)}{M_\pi^2}, \quad (5.5)$$

and

$$s_i = (P_K - P_i)^2, \quad s_0 = \frac{1}{3} \sum_{i=1}^3 s_i. \quad (5.6)$$

P_K and P_i are the four-momentum vectors of the kaon and the three pions, respectively (Groom et al. 2000b). The values used for the Dalitz coefficients g , h , and k in this analysis were taken from the most recent KTeV analysis of $K_L \longrightarrow \pi^+\pi^-\pi^0$ decays (Barrio 2001). The coefficient j is related to CP-violating effects, and was fixed at zero.

Because the lifetime of the neutral pion is so short, it is assumed that the π^0 promptly decays at the same position as the parent kaon decay.

5.2.3 $\pi^0 \longrightarrow e^+e^-\gamma$ generator

The second stage in the generation of normalization MC is the Dalitz decay $\pi^0 \longrightarrow e^+e^-\gamma$. As mentioned in Chapter 1, the amplitude for this decay was initially calculated assuming a form factor equal to unity (Kroll and Wada 1955). Experimental evidence, however, shows that a form factor of the type

$$f(x) = 1 + ax, \quad x = \frac{M_{ee}^2}{M_{\pi^0}^2} \quad (5.7)$$

is required for the theory to match $\pi^0 \longrightarrow e^+e^-\gamma$ data, with the empirical coefficient a measured to be 0.032 ± 0.004 (Groom et al. 2000b). This is the generator used by the KTeV monte carlo to simulate Dalitz decays.

Radiative corrections to this mode occur at order α^2 , and are therefore fairly significant (Mikaelian and Smith 1972). A cutoff of $M_{\gamma\gamma} > 1 \text{ MeV}/c^2$ is imposed to prevent divergences from the inner bremsstrahlung contribution.

5.2.4 $K_L \longrightarrow \mu^+\mu^-\gamma$ generator

The decay $K_L \longrightarrow \mu^+\mu^-\gamma$, coupled with conversion of the photon to a e^+e^- pair, proved to be the most significant background to the signal. Monte carlo for these decays was needed to obtain an estimate of remaining background events for subtraction before calculation of the $K_L \longrightarrow e^+e^-\mu^+\mu^-$ branching ratio. Because this mode is simply the muonic Dalitz decay of the K_L , the Kroll–Wada formula with a form factor of unity was used as a decay generator. Radiative corrections to $\mathcal{O}(\alpha^2)$ were included, again with a cutoff at $M_{\gamma\gamma} > 1 \text{ MeV}/c^2$.

The monte carlo code used to generate these background events was slightly modified in order to force the conversion of the photon within the vacuum window. A discussion of how photon conversion is handled by the KTeV MC is given later in this chapter.

5.2.5 $K_L \longrightarrow \pi^+\pi^-e^+e^-$ generator

This mode appeared as a background if the charged pions decayed to muons, or punched through the muon filter, firing counters in MU3. The matrix element used for $K_L \longrightarrow \pi^+\pi^-e^+e^-$ in the KTeV monte carlo is quite complex (Sehgal and Wanning 1992). The amplitude includes contributions from bremsstrahlung, direct emission, and K^0 charge radius diagrams. Interference between the first two terms, which leads to a large CP-violating effect in the angular distribution of the decay products, is fully simulated.

5.3 Particle Tracing

After generation, the daughter particles were rotated by a random angle about the direction axis of the parent kaon, and were boosted into the laboratory reference frame. Each of the particles was then traced through the KTeV detector. The tracing was done in sequential steps of z , from the upstream to the downstream edge of each detector element and air gap. Several types of particle interactions were allowed to occur at each step; these interactions will be described in detail in the following section.

Photons and electrons were traced until they escaped the detector geometry or were absorbed by an element (a veto or the EM calorimeter). Charged pions were followed until they escaped or struck the upstream face of the calorimeter (or, if they passed through a calorimeter beam hole, the hadron anti), while muons were allowed to pass through every element without being stopped. At each step, pions were allowed to decay to muons, with a probability in accordance with their exponential lifetime distribution. In the event of pion decay, the daughter muon was added to the particle list and traced through the detector. Neutrinos were not traced by the monte carlo.

Particles that escaped from the KTeV detector geometry were considered “lost”. Tracing of these particles was stopped, and depending on the user-requested conditions for the final state, the event could be rejected at this stage in order to conserve computer processing time.

5.4 Particle Interactions

As mentioned in the previous section, several types of interactions were allowed within each detector element. These include photon conversion, multiple scattering, the creation of δ -rays, and bremsstrahlung.

5.4.1 Photon conversion

Photons of energy greater than 100 MeV were allowed to convert into a $e^+ e^-$ pair. The conversion probability in each detector element was given by

$$P_{conv} = 1 - e^{-\frac{7}{9}(X/X_0)}, \quad (5.8)$$

where X is the total path length through the element, and X_0 is the characteristic radiation length of the element. The z position of conversion within the element was selected randomly.

The energy distribution between the electron and positron was determined by the Bethe–Heitler spectrum (Mikelsons 1999):

$$P(\epsilon_+, \epsilon_-) \approx \epsilon_+^2 + \epsilon_-^2 + \left(\frac{2}{3} - \frac{1}{9\sigma}\right) \epsilon_+ \epsilon_-, \quad (5.9)$$

where $\epsilon_{\pm} = E_{e\pm}/E_{\gamma}$, and σ is a known constant (≈ 3.7). The electron and positron directions were offset by a small angle, rotated by a random amount about the axis of the parent photon direction. This offset angle was a function of E_{γ} , ϵ_{\pm} , and X_0 , and was calculated using a distribution given by the EGS4 code library from SLAC (Nelson et al. 1985).

5.4.2 Multiple scattering

Charged particles passing through matter are continuously deflected by Coulombic interactions with nuclei. In order to simulate the effect of multiple scattering, the KTeV monte carlo randomly chose a scattering angle θ immediately after a charged particle exited a detector element. For all charged particles upstream of the calorimeter, the probability spectrum for θ was a gaussian, centered on 0, with a width θ_0 taken from the Molière distribution (Groom et al. 2000c), (Bethe 1953):

$$\theta_0 = \frac{13.6 \text{ MeV}}{\beta c p} z \sqrt{X/X_0} [1 + 0.038 \ln(X/X_0)], \quad (5.10)$$

where p , βc , and z are the momentum, velocity, and charge of the scattered particle. The direction of scatter was chosen by rotating θ about the previous particle direction by a random amount.

As mentioned earlier, muons were the only particles traced past the upstream face of the calorimeter. The scattering simulation after this point was more complicated – much work was done to tune the monte carlo simulation of the muon system for the purpose of matching the data distributions observed during the analysis of $K_L \rightarrow \mu^+ \mu^- \gamma$ (Quinn 2000). This extra effort was required in order to simulate the effects of energy loss within the muon filters. Because the normalization modes used in the $K_L \rightarrow \mu^+ \mu^- \gamma$ and $K_L \rightarrow e^+ e^- \mu^+ \mu^-$ analyses contain no muons, the final branching ratio results are sensitive to the momentum threshold of the muon system. An inaccurate simulation of the system would lead to a mismeasurement of the threshold, the signal acceptance, and the branching ratio.

Starting at the CsI calorimeter face, muon energy loss was simulated using the Bethe–Bloch relation (Groom et al. 2000c):

$$-\frac{dE}{dX} = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \left(\frac{2 M_e c^2 \beta^2 \gamma^2 T_{max}}{I^2} \right) - \beta^2 - \frac{\delta}{2} \right], \quad (5.11)$$

where Z and A are the effective atomic number and mass of the detector element, γ is the common relativistic parameter $(1 - \beta^2)^{1/2}$, T_{max} is the maximum possible kinetic energy transfer to a free electron in a collision, I is the mean excitation energy, and $K \approx 0.31 \text{ g}^{-1} \text{ cm}^2$. δ is a small correction made to account for density effects (from polarization of the surrounding media).

Random fluctuations were superimposed on the Bethe–Bloch energy loss result. It was found (Quinn 2000) that in order to properly simulate the muon threshold, different fluctuation distributions were required for different ranges of $\kappa = \xi/T_{max}$, where

$$\xi = 153.4 \rho X \frac{Z}{A} \left(\frac{z}{\beta} \right)^2, \quad (5.12)$$

and ρ is the density of the element (Patrick 1986). Landau fluctuations (Landau

1944) were used for $\kappa < 0.01$, Vavilov (Vavilov 1957) for $0.01 \leq \kappa \leq 10$, and Gaussian fluctuations were used for $\kappa > 10$.

5.4.3 δ -rays

δ -rays are electrons that are “knocked on” by incoming particles. Because they could leave extra hits on sense wires (which contributes to drift chamber inefficiency), δ -ray production was simulated within the drift cells. Assuming that the incoming particle with energy E is highly relativistic, the probability of generating a δ -ray (with energy T above a fixed cutoff T_0) in the KTeV monte carlo is given by

$$P_{T>T_0} = (15400 \text{ keVcm}^3/\text{gm}) \frac{Z}{A} \frac{1}{E} \rho X. \quad (5.13)$$

Only one such interaction was allowed in each drift cell. The direction of the resulting δ -ray was assumed to be perpendicular to that of the incoming particle, and the energy was chosen to conform to a T^{-2} spectrum (Groom et al. 2000c). It was assumed that the transverse momentum kick to the parent particle was small enough to be ignored; its trajectory after the interaction was therefore left unchanged. Further discussion regarding the behavior of δ -rays within a drift cell will be presented later in this chapter.

5.4.4 Bremsstrahlung

Electrons (and positrons) were allowed to radiate a single bremsstrahlung photon as they traversed a detector element. Assuming that the particles were ultra-relativistic, and that the Born approximation is valid, the probability of emission of a photon of momentum $k > k_0$ (where k_0 is a cutoff value) was

$$P_{k>k_0} = \frac{X}{X_0} \frac{1}{18\zeta + 1} \left[2(12\zeta + 1) (k_0/E_0 - \ln(k_0/E_0) - 1) + 9\zeta \left(1 - (k_0/E_0)^2 \right) \right], \quad (5.14)$$

where $\zeta = \ln(183/Z^{1/3})$, and E_0 is the initial energy of the electron (Mikelsons 1999). The default cutoff in the KTeV monte carlo was $k_0 = 0.001E_0$.

The expression for the differential momentum spectrum and angular distribution of the radiated photon is quite complex (a full description can be found in (Mikelsons 1999)). To conserve computation time, the full angular simulation was performed only if the radiation occurred upstream of the analysis magnet (meaning the electron would be well separated from the photon before striking the calorimeter) *and* $E_0 > 1.5$ GeV (approximately the efficiency threshold for electron cluster detection by the HCC trigger). If these conditions were not satisfied, the momentum spectrum for the radiated photon was given by

$$\frac{d\sigma}{dk} = 4\alpha_{EM} Z^2 r_e^2 \frac{1}{k} \left[\left(1 + \left(\frac{E}{E_0} \right)^2 - \frac{2}{3} \frac{E}{E_0} \right) \zeta + \frac{1}{9} \frac{E}{E_0} \right], \quad (5.15)$$

obtained by integrating out the angular dependence of the differential cross section (Mikelsons 1999).

5.5 Particle Detection

Detector response was simulated for each of the detector elements traversed by particles. The simulation procedure was different for each element, and will be described in detail below.

5.5.1 Photon veto simulation

Electrons and photons traced to a photon veto detector were terminated at the veto plane. The energy of the particle, after Gaussian smearing to account for resolution effects, was compared to the threshold of the hit veto segment (the thresholds were calibrated offline using data from special runs (Hanagaki 1997)). An amount of energy exceeding threshold in any segment caused a L1 trigger source to be set. If the source was used as a veto condition in the physics trigger being evaluated

during generation, the event was rejected at this stage in order to improve processing efficiency.

5.5.2 Spectrometer simulation

Hit simulation

Each of the 16 drift chamber planes (1 plane pair per view per chamber) was simulated individually. Hit locations were determined by selecting the sense wire closest to the charged particle within the plane. The distance between the particle and the wire was Gaussian smeared to simulate the position resolution of the drift chamber (approximately $100\ \mu\text{m}$), and the SOD (sum of distances – see Section 4.1) was calculated. Drift distances were correlated to drift times using the XT maps described in Section 2.2.2. Inefficiency related to discriminator response was modeled by considering a wire “dead” for 45 ns after a hit (the approximate width of a sense wire signal pulse entering the discriminator), meaning additional hits on the wire during that period were ignored.

Two corrections were made to the monte carlo SODs at this stage. Effects from discrete ionization of the drift gas near the sense wires led to a high-side tail on the SOD distribution. To account for this, the calculated SOD value was skewed slightly upwards (although this was done in such a way as to maintain the mean of the SOD distribution at the sense wire spacing of 6.35 mm). Secondly, due to a mass dependence of the ionization efficiency, there was a small difference (approximately $15\ \mu\text{m}$) observed between the mean SOD for electrons and heavier particles (pions and muons). A scaling factor, calculated from K_{e3} electron and pion SOD distributions, was used to simulate this effect.

δ -ray simulation

As mentioned in Section 5.4.3, knock-on electrons (δ -rays) could be produced by incoming particles. Only one such interaction was allowed in each drift cell traversed by the incoming particle. If a δ -ray was produced within a drift cell, its effective range was calculated; a hit left on a sense wire within the effective range of the δ -ray was treated identically to a hit left by any other “real” particle (i.e., the sense wire was considered “dead” for 45 ns after the hit).

DC maps

Unfortunately, even after these corrections, significant tracking discrepancies remained between data and monte carlo. Inaccurate simulation of isolated hit inefficiencies and high SODs could lead to misreconstruction of track segments, causing mismeasurement of track momenta, vertex position, and vertex χ^2 . To account for the residual discrepancies, maps were generated for each drift chamber plane. These *DC maps* describe isolated hit and high SOD probabilities as a function of position within the plane and time during the run, and were applied after all other corrections. Some example maps are shown in Figure 5.2 (LaDue). After application of the DC maps, improvement in the data/MC agreement is clear (see Figure 5.3).

Analysis magnet simulation

A transverse momentum kick in the $\pm x$ direction was imparted to charged particles as they crossed the plane of the analysis magnet. As mentioned in Section 2.2.2, the magnitude of this kick was approximately 205 MeV/c in 1997, and 150 MeV/c in 1999. Exact values were calibrated on a regular basis by measuring the mean invariant mass of reconstructed calibration $K_L \rightarrow \pi^+\pi^-$ events, and tuning the P_t kick in order to match the kaon mass. These kick values were installed into the KTeV database for use during monte carlo generation and event reconstruction.

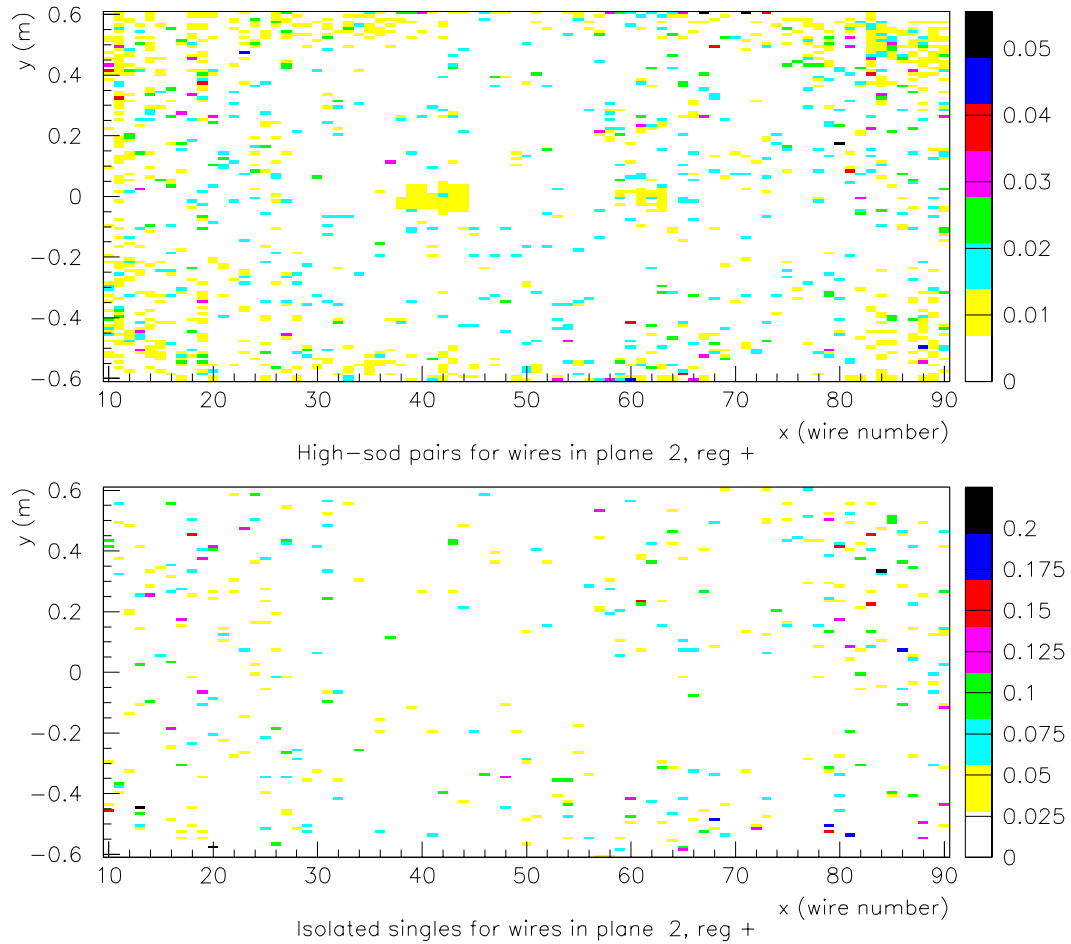


Figure 5.2: Examples of DC maps, used to correct residual tracking discrepancies between data and monte carlo. These particular maps were applied to the x view of DC1 during the generation of 1999 monte carlo. The top map served to tune the high SOD probability distribution, while the bottom corrected the isolated hit probability distribution. Note the need for additional high SOD correction in the beam region, caused mainly from radiation damage to the sense wires in that area of the chamber.




Figure 5.3: 1997 $\pi^+\pi^-\pi_D^0$ track illumination at the x view of DC1. x position is measured in meters, and the dots represent data. In the plots on the left, the monte carlo (histogram) was generated without the use of DC maps, which correct for residual data/MC discrepancies in high SOD and isolated hit inefficiency distributions. A 1.5σ slope is evident in the ratio. This slope disappears when the maps are turned on (plots on the right).

5.5.3 VV', MU2, and MU3 simulation

When charged particles in the monte carlo were traced to a counter in the VV', MU2, or MU3 banks, a random number was generated and compared to the counter efficiency to determine if the counter registered a hit. The counter efficiencies were individually measured using calibration data. Cracks between scintillator paddles were also simulated, as well as the light propagation time from the hit location to the attached phototube. In-time hits were used as sources for the L1 trigger.

Although there existed a small probability for charged pions to hadronically interact and shower within the VV' hodoscope, this effect was not simulated within the default monte carlo. This systematic difference between the data and monte carlo necessitated a correction to the calculated normalization acceptance, which will be discussed in Section 6.6.3.

5.5.4 Calorimeter simulation

EM showers

As mentioned earlier, the tracing of most types of particles was stopped at the upstream face of the CsI calorimeter. At this stage, a detailed electromagnetic shower simulation was initiated for photons and electrons. The mean longitudinal shower position within the crystal, \bar{z} , was calculated. The \bar{z} distribution was found to be different for electrons and photons:

$$\begin{aligned}\bar{z}_e &= 0.11 + 0.18 \ln E_e \\ \bar{z}_\gamma &= 0.12 + 0.18 \ln E_\gamma\end{aligned}\tag{5.16}$$

where E is measured in GeV, and \bar{z} in meters (Prasad 2002).

At this point, an electromagnetic shower was selected from the shower library. Entries in the library were generated using the GEANT simulation package,

and spanned a transverse area of 32.5×32.5 cm (corresponding to a 13×13 array of small blocks; large blocks were considered to be four small blocks). Showers were binned according to the energy of the parent particle (6 bins, with upper boundaries of 2, 4, 8, 16, 32, and 64 GeV) and the particle's x and y position within the block when projected to \bar{z} ($325\,500 \times 500\,\mu\text{m}$ bins to cover an octant of the crystal). Each bin contained five entries, for a total library size of 9,750 showers (Prasad 2002). In order to account for variation in longitudinal scintillation response along the length of a crystal, the showers were segmented into 25 bins in z . This allowed for a convolution of the shower's longitudinal energy distribution with the previously measured longitudinal response of the crystal.

The energy of the incoming particle was smeared to account for calorimeter resolution effects. The amount of smearing depended on the type (large or small) and location of the hit block, and was tuned to match the resolutions observed in the data. After a shower was selected from the library, its total energy was scaled to match the smeared energy of the incoming particle. Energy from showers bordering a beam hole was allowed to “sneak” across to crystals on the opposite side of the hole. This effect was simulated using the distributions from the sneaky energy correction, discussed in Section 4.2.4.

Simulation of non-EM particles

A separate library contained a collection of hadronic GEANT showers, used to simulate the $\approx 30\%$ of charged pions that hadronically interact within the CsI. Aside from their larger size ($102.5 \times 102.5\text{ cm}^2$, corresponding to an array of 41×41 small blocks) and coarser transverse position binning (25 $x \times y$ bins on the crystal face), hadronic showers were handled in the same manner as the EM showers described above. Non-interacting pions were treated as MIPs, leaving 320 MeV in the calorimeter. As mentioned in Section 5.4.2, muons deposited energy in the CsI according to the Bethe-Bloch distribution with fluctuations. The energy from

simulated MIPs was confined to a single crystal within the calorimeter array.

Digitization and readout simulation

After simulating energy deposits from all particles that struck the calorimeter, the Q/E constants (described in Section 2.2.5) were used to convert the energy in each block to a corresponding amount of charge. This charge was spread out over a series of time slices, using pulse shape distributions measured from the data. Smearing of the total charge in each slice was included to account for the effects of photostatistics. Digitization of the monte carlo calorimeter information included a detailed simulation of the DPMT board for each channel, using the calibration constants described in Section 2.2.5. Channels were “read out” for further use within the monte carlo only if they exceeded the VME pipeline energy threshold, just as in the data.

5.5.5 Etot and HCC simulation

As discussed in Chapter 3, the thresholds for the L1 Etot and L2 Hardware Cluster Counter (HCC) trigger sources varied with time as radiation damage to the calorimeter crystals accumulated, and as the RF timing drifted with respect to the internal L1 clock. This section describes how the thresholds were extracted in order to model this time dependency in the KTeV monte carlo.

Cuts on accidental activity

Etot and HCC thresholds were measured using calibration K_{e3} electrons collected on a special minimum-bias trigger (in which the Etot and the HCC sources were read out, but not required by the trigger logic). However, these data events contained a substantial amount of accidental activity. Extra energy in the calorimeter

would have biased the calibration results, and therefore had to be removed before the thresholds could be extracted. This was done by applying a series of cuts to eliminate out-of-time energy. The cut locations were determined by studying the distribution of in-time energy over a series of time slices.

The algorithm to remove accidental energy first isolated the seed blocks from hardware clusters in reconstructed calibration K_{e3} events. Histogrammed in the top plot of Figure 5.4 is the ratio of the energy within the in-time slice to the sum of energy over all time slices ($E(\text{in-time})/E(\text{all slices})$) for these seed blocks. The detector timing was synchronized such that the energy from real K_{e3} electrons was mostly contained within the in-time slice (leading to the large peak bounded by the vertical lines). Anything outside the lines is probably residual energy from accidental activity occurring during earlier time slices, and the corresponding events were removed from the calibration dataset.

The middle plot in Figure 5.4 shows the ratio $E(\text{in-time}+1)/E(\text{all slices})$, corresponding to the slice immediately following the in-time slice. Because of the rapidly falling tail of dynode pulses from the calorimeter PMTs (Figure 3.5), less energy would be expected from real K_{e3} electrons in this slice, meaning the peak (bounded by the vertical lines) should shift downwards. Entries below the lower line are still from residual early accidental activity, while entries above the upper line come from new accidental energy arriving in this late slice. Cutting out all electrons outside the lines removed even more out-of-time events from the calibration dataset.

The bottom plot shows the energy fraction distribution two slices after the in-time slice, ($E(\text{in-time}+2)/E(\text{all slices})$). Even less of the in-time energy remains in this slice, and the location of the in-time peak (bounded by the vertical lines) has again shifted downwards. As with the two previous plots, everything outside the vertical lines was removed. The process was repeated a total of 6 times (up through 5 slices after the in-time slice), leaving an essentially pure calibration sample with no accidental activity remaining. Since PMT pulse shapes vary with energy and




Figure 5.4: Ratio of energy in the in-time (top), in-time + 1 (middle), and in-time + 2 (bottom) slices to the total energy in all slices for small seed blocks in the “medium” energy range of the calibration K_{e3} dataset. Most of the energy from real K_{e3} electrons is expected in the in-time slice, meaning electrons outside the vertical lines are from out-of-time accidental activity. Removing these events guarantees a pure calibration sample for E_{total} and HCC threshold extraction.

crystal size, different cuts were made for low (< 3 GeV), medium (between 3 and 15 GeV) and high (> 15 GeV) energy pulses, as well as for small and large blocks.

Threshold extraction

Examples of threshold curves for Etot sources ET1–4 were shown in Figure 3.1. The curves were fit with error functions, using logistic regression. The turn-on threshold of a source was defined as the energy at which the fit value was 0.5 (corresponding to the source being 50% efficient). The turn-on width was defined as the difference in energy between 90% and 10% efficiency. Threshold curves were fit to each of the 3100 channels of the HCC in the same fashion. An example of such a fit was given in Figure 3.3.

Thresholds and widths for the Etot sources and HCC bits were measured as a function of time (the 1999 run, for example, was broken into 20 calibration ranges) and were then installed into the KTeV database. The monte carlo used these parameters to create probability distributions describing whether a particular Etot source or HCC bit fired during generation of an event. As in the data, the Etot sources (ET1–4) were used as direct inputs to the L1 trigger. The HCC bits that fired were stored in an array, which was passed to a simulated HCC processor. The processor determined how many hardware clusters were present in the monte carlo event, and passed on this information to the L2 state machine in the form of a L2 source.

5.6 Accidental Overlays

At this point, the monte carlo event is still fairly “clean”. While some interactions (the creation of δ -rays, for example) might have taken place, the only activity present in the event is the direct by-product of particles from the initial decay. In real life, however, this is far from the case. Beam interactions, cosmic rays, and junk

from the target and vacuum window all contributed to an ever-present background of accidental activity in the detector. This activity could lead to tracking inefficiencies (by leaving an early hit on a drift chamber sense wire), mismeasurement of cluster energies (if an accidental photon landed on a real electron cluster in the calorimeter), or could even fire trigger sources (if a cosmic ray happened to cross MU3 during the in-time window).

As mentioned in Section 2.1.1, a special trigger was set up to record events containing only accidental activity. These events were stored on disk to be used as *overlays* during monte carlo generation. After pedestal subtraction, ADC values (such as photon veto and calorimeter crystal energies) from the overlay event were added directly to those from the generated event. Calorimeter energies were added on a slice-by-slice basis. TDC entries (such as DC hits) from the overlay event were added to the list of those in the generated event. Latch information from the overlay and the generated event was combined using a logical OR. Since accidental activity could lead to inefficient or spurious triggers, the overlay addition was performed *before* evaluation of any trigger sources.

5.7 Trigger Simulation

Physics trigger definitions used by the monte carlo were identical to those used during data taking. As mentioned earlier, trigger sources for the photon vetos were evaluated during particle tracing in order to conserve processing time. The remainder of the L1 trigger was evaluated upon successful generation of the event. Algorithms used by the L2 processors (HCC, YTF, and the fruit) were replicated in the monte carlo code, as were the software routines constituting L3. Event generation was stopped as soon as the event failed any level of the trigger. Events that passed L3 were written to disk in exactly the same format as the data.

Type	Purpose	N_{Winter}	N_{Summer}	N_{1999}
$K_L \longrightarrow e^+e^-\mu^+\mu^-$	signal acceptance	150K	150K	150K
$K_L \longrightarrow \pi^+\pi^-\pi_D^0$	normalization acceptance, systematic studies	2M	2M	5M
$K_L \longrightarrow \mu^+\mu^-\gamma,$ $\gamma \longrightarrow e^+e^-$	background	1M	1M	1M
$K_L \longrightarrow \pi^+\pi^-\pi_D^0,$ π^\pm decay/punchthrough	background	100M	100M	150M
$K_L \longrightarrow \pi^+\pi^-e^+e^-,$ π^\pm decay/punchthrough	background	5M	5M	5M

Table 5.1: Type, purpose, and number (by run) of monte carlo events generated for this analysis.

5.8 Monte Carlo Samples

Table 5.1 lists the type, purpose, and number (by run) of monte carlo events generated for this analysis. Events were generated using Version 5.04 of the KTeVMC library. The relatively small number of generated signal monte carlo events is due to the lengthy processing time required for the radiative corrections. However, because $K_L \longrightarrow e^+e^-\mu^+\mu^-$ is such a rare decay mode, the monte carlo sample statistics will have only a marginal effect on the precision of the branching ratio measurement. This measurement, including discussion of the three background modes listed in the table, is the focus of the following chapter.

CHAPTER 6

Branching Ratio Analysis

As mentioned at the end of Chapter 3, almost 80 DLT tapes of dimuon and four-track data were written during the three run periods, corresponding to a total data sample of over 200 million events. Clearly, searching for interesting events was not a trivial task (the proverbial needle in the haystack comes to mind). This chapter describes cuts made to the dataset to isolate signal $K_L \longrightarrow e^+e^-\mu^+\mu^-$ and normalization $K_L \longrightarrow \pi^+\pi^-\pi_D^0$ events, and how these events were used to calculate the $K_L \longrightarrow e^+e^-\mu^+\mu^-$ branching ratio.

6.1 Crunch

The first stage of analysis was known as the *crunch*. The crunch served three purposes:

1. isolating $e^+e^-\mu^+\mu^-$ and $\pi^+\pi^-\pi_D^0$ L3 tags on the split tapes
2. cutting out vacuum window junk events (described below)
3. writing crunched events to DLT tapes in a compressed format.

Only events with the appropriate L3 tags were crunched; all other events were discarded. After running the tracking and clustering routines outlined in Chapter 4,




Figure 6.1: Fast vertexing results for data on a single split tape. Plotted is the z position of the found vertex, in meters downstream of the target. Note the dominant peak from interactions at the vacuum window. The bumps at 138 m and 152 m correspond to interactions within RC7 and RC9, respectively.

the crunch called the fast vertexing algorithm, first described in Section 3.6¹. Because fast vertexing required only three tracks from a common vertex instead of four, much less time was needed to crunch an event at this stage than with the full four-track vertexing routine. The vertex z distribution of all successfully vertexed events from a single split tape is shown in Figure 6.1. The distribution is dominated

¹Reconstructing the event again at the crunch stage allowed for the use of the more precise offline calibration constants, which were unavailable to L3 during data collection.

by a peak at $z = 159$ m, the location of the kevlar/mylar window at the downstream end of the vacuum region. These “vacuum window junk” events are not real kaon decays; instead, they resulted from particle interactions within the kevlar. (Interactions within other aperatures are also evident in Figure 6.1, although to a much lesser degree.) The primary aim of the crunch was to remove these vacuum window interactions from the dataset.

Events were discarded if the fast vertexing routine returned a good vertex within 1 m upstream of the vacuum window, where “good” was defined to mean that the uncertainty on the vertex z position (calculated from the uncertainties in the slopes and intercepts of the fitted track segments used to form the vertex) was less than 2 m. Events that survived this cut then underwent the full four-track vertexing. If a four-track vertex was found, the event was compressed and written out to tape. The crunch output filled a total of 12 DLT tapes with 22.7 million events, a full order of magnitude reduction of the dataset. Monte carlo studies show that the crunch caused a loss of only 0.8% of signal events and 1.5% of normalization events. All monte carlo events used in this analysis were crunched in exactly the same manner as the data.

6.2 Bad Spill Cut

Unfortunately, a variety of problems during the run rendered some of the data unusable. Hardware failures affected almost every detector at some point in time. Some of these failures (of the PMTs, DPMTs, or drift chambers, for example) were severe, as they led to corrupted measurements and faulty triggers. On the other hand, problems with the TRDs or the hyperon trigger could safely be ignored, as these systems were not used in this analysis.

A list of problems during data collection was compiled on a spill-by-spill basis, using calibration data and the KTeV logbooks. These problems were classified

Bit number	Problem	Cut?
1	trigger	yes
2	DPMT pedestal exponent > 0 (bad QIE)	yes
3	bad DPMT capacitor (out of rotation)	no
4	blown QIE comparator	yes
5	misc dead DPMT	yes
6	DPMT pedestal drift	no
7	PMT gain drift	yes
8	broken PMT dynode	yes
9	CsI pipeline problems	yes
10	global CsI problems	yes
11	Etotal	yes
12	FERA ADC	no
13	drift chamber	yes
14	photon veto	yes
15	V-bank	yes
16	muon system	yes
17	HCC	yes
18	fruit	yes
19	TRD trigger	no
20	hyperon trigger	no
21	DAQ, L3	yes
22	not 799 run	yes
23	short run	yes
24–28	TRD	no
29	beam	yes
30–32	unused	no

Table 6.1: Bad spill cut summary. Only “severe” bits were cut on.

according to cause and severity, and were installed in a “bad spill” database. This database could be accessed by calling a special routine that returned a bitmask of problems for any given spill. The mask contained 64 bits: 32 “warning” bits, and 32 “severe” bits. Table 6.1 lists the meaning of each bit, as well as which bits were cut on for this analysis. Only events from spills with “severe” bits were cut. Figure 6.2 shows the distribution of bad spill bits for all events passing the crunch stage, with




Figure 6.2: Severe bad spill bits for all 22.7 million events passing the crunch stage. Bits from the 1997 data are plotted in the open histogram, while bits from the 1999 data are plotted in the cross-hatched histogram. Bit meanings are listed in Table 6.1. Arrows point to the cut bits, corresponding to about 7% of the dataset.

arrows pointing to the bits that were cut on for this analysis. Applying the bad spill cut led to a reduction of approximately 7% of the dataset.

It is clear from Figure 6.2 that the largest bad spill losses in this analysis came from Level 3 trigger problems (bit 21) in 1999. At seemingly random times during the 1999 run, arrays of calorimeter information became corrupted for certain physics triggers during L3 processing. When this happened, the abrupt change in




Figure 6.3: MU2 counter illumination (number of TDC hits per scintillator paddle) for fully reconstructed 1999 $\pi^+\pi^-\pi_D^0$ events after all analysis cuts. The counter map was shown in Figure 2.9. Despite the fact that this plane was in veto, over 1% of events have in-time hits in the MU2 TDCs, illustrating the need for additional trigger verification.

L3 acceptance set off an audio alarm, and the trigger had to be manually reset by the shift crew. Unfortunately, the source of this problem was never found, and the data from these periods had to be thrown out. HCC hardware problems (damaged comparator chips and blown input buffers, see Section 3.2.1) caused the largest losses in 1997 (bit 17).

6.3 Trigger Verification

Because of accidental activity, timing jitter, and inefficiencies in the trigger electronics, it was possible for some unwanted events to slip through the KTeV trigger system. For example, Figure 6.3 shows the MU2 counter illumination for fully reconstructed 1999 normalization events after all analysis cuts (these cuts will be described later). Despite that fact that MU2 was used as a veto in the four-track

trigger, over 1% of normalization events contained an in-time hit in one of the MU2 TDCs. This illustrates the need for *trigger verification* – cuts used to ensure that the reconstructed event could indeed have satisfied the trigger requirements.

For the 1997 data, only the VV' trigger source was verified for the signal and normalization modes. This is because the other elements (HCC, MU3 banks, the fruit, etc.) were verified either implicitly or explicitly during the process of event reconstruction. The 1997 verification was done by projecting the reconstructed downstream track segments to the VV' hodoscope planes. If a sufficient number of tracks matched paddles with in-time hits, the event was considered verified.

While verification for the signal mode remained unchanged between 1997 and 1999, the process was slightly more complicated for 1999 normalization data (due to the more restrictive nature of the four-track trigger; see Section 3.3.2). These events were thrown out if any in-time hits were reconstructed in the MU2 veto TDCs, or if the total calorimeter energy was measured to be less than 11 GeV (verifying the ET1 requirement), in addition to the usual VV' verification.

Trigger verification at this stage reduced the total size of the dataset by another 6%.

6.4 Particle Identification

The ratio E/p was used to associate tracks with particle types. As mentioned earlier, relativistic electrons were expected to deposit all of their energy in the EM calorimeter, giving an E/p of 1. MIPs such as muons and pions typically deposited very little energy in the calorimeter, leading to an E/p near 0. Particles traveling down a beam hole have an E/p of exactly 0, since they deposited no energy in the calorimeter. About 30% of the time, a pion would interact hadronically within the CsI, leading to a pion shower and an intermediate value of E/p . Tracks with $E/p > 1$ are also possible, due to calorimeter and spectrometer resolution effects,




Figure 6.4: Distribution of the ratio E/p for all tracks in the 1999 data. The peak at $E/p = 1$ comes mostly from electrons, while the peak at 0 is from MIPs (muons and pions), as well as from particles that traveled down a beam hole, depositing no energy in the calorimeter. The hump in between is caused by hadronic interactions from pions in the CsI. Tracks with $E/p > 1$ result from calorimeter and spectrometer resolution effects, or from accidental activity on top of a real electron cluster.

or from accidental activity on top of a real electron cluster. The E/p distribution for all tracks in the 1999 data is shown in Figure 6.4.

Tracks with E/p between 0.8 and 1.2 were initially identified as electrons. Charged pions in the normalization mode were required to have an E/p less than 0.8 *and* hit the calorimeter at least 7 cm from the outside edges and 5 cm from either

beam hole. The photon from the Dalitz decay in normalization events was identified by looping over all clusters not matched to tracks and taking the combination that minimized the value $|M_{\pi^0} - M_{e+e-\gamma}|$. Muons in the signal mode were required to have E/p less than 0.8 *and* hit at least two out of the three muon trigger planes (MU2, MU3X, and MU3Y).

Signal events were accepted if they contained exactly two electrons and two muons. Both electron tracks and at least one muon track were required to match to clusters in the calorimeter, allowing for one muon down a beam hole. Normalization events were accepted if there were exactly two oppositely charged electrons, two oppositely charged pions, and the mass of the best π^0 fell between 120 and 150 MeV/c². All tracks in the normalization mode were required to match to clusters.

Approximately 560,000 signal and 1.4 million normalization events remained after the particle ID stage. The invariant mass distributions of reconstructed signal and normalization events from all runs are shown in Figure 6.5.

6.5 Initial Cuts

Cuts that were then applied to both samples are listed below:

- $20 \text{ GeV}/c < p_K < 220 \text{ GeV}/c$
- $90 \text{ m} < \text{vertex } z < 158 \text{ m}$
- $0.95 < E/p < 1.05$ for electrons
- $P_t^2 < 0.00025 \text{ GeV}^2/c^2$
- $0.482 \text{ GeV}/c^2 < M_K < 0.512 \text{ GeV}/c^2$
- vertex $\chi^2 < 40$

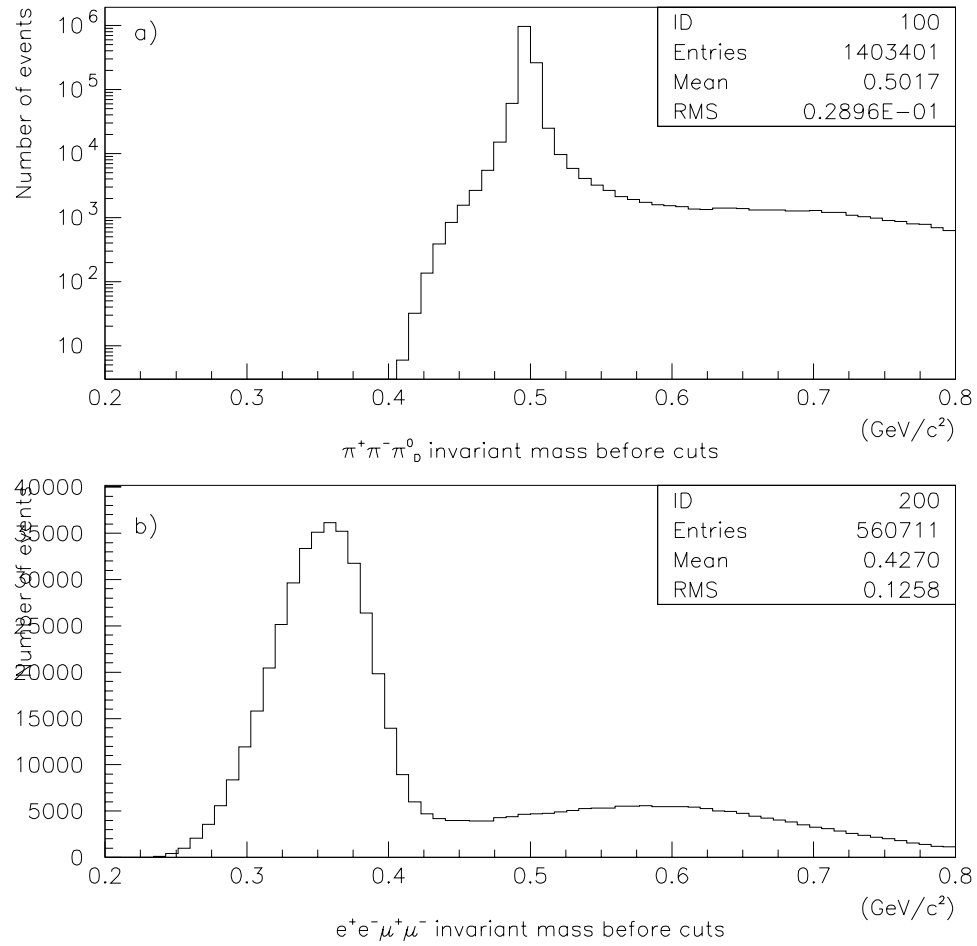


Figure 6.5: (a) $\pi^+\pi^-\pi_D^0$ and (b) $e^+e^-\mu^+\mu^-$ invariant mass distributions for all data after track/cluster reconstruction and particle identification. Additional cuts have not yet been made.




Figure 6.6: Definition of the variable P_t^2 . P_t is the component of the reconstructed momentum perpendicular to a line drawn from the target to the vertex. For a well reconstructed event, this variable should be near 0.

- vertex required to be in beam hole when projected to calorimeter face

The cuts on reconstructed kaon momentum (p_K) and vertex z position were chosen because the acceptance for K_L decays drops off very quickly outside of these ranges. Any events observed outside of these boundaries were most likely misreconstructed K_L decays or background (probably K_S s or hyperons).

The parameter vertex χ^2 was first described in Section 4.3. The placement of the cut at a value of 40 was chosen to eliminate background from two two-track decays in the same RF bucket. This will be discussed in Section 6.7.

P_t^2 is the square of the transverse component of the reconstructed kaon momentum perpendicular to a line drawn from the target to the vertex (see Figure 6.6). If all the particles from the decay were accounted for, and the event was reconstructed properly, this variable should be 0 (except for resolution effects). The P_t^2 distributions for 1999 normalization data and monte carlo are shown in Figure 6.8(a). A clear excess of poorly reconstructed and background events is seen in the data above the cut.

The cut on M_K , the reconstructed invariant mass, was determined by fitting a gaussian to reconstructed signal monte carlo, and allowing a region 3σ wide on either side of the fitted mean. This mass region was also known as the *signal region*.

Further cuts were specific to each mode, and will be described separately below.

6.6 Normalization Mode

As mentioned earlier, the normalization mode served two main purposes. The first was to provide a measurement of the *kaon flux* – the total number of K_L decays that occurred, both observed and unobserved, in the KTeV detector. The flux is needed to calculate the branching ratio of the signal mode $K_L \longrightarrow e^+e^-\mu^+\mu^-$. Secondly, high-statistics normalization mode data was compared to monte carlo to search for discrepancies in the distributions of measured parameters. Such discrepancies could point to the presence of background events, or could signify that the monte carlo is not modeling the detector correctly, leading to possible mismeasurement of the detector acceptance². For these reasons, it was important to study the normalization mode very carefully.

6.6.1 Further cuts

$pp0kine$ is a kinematic quantity used to detect any missing momentum in a reconstructed event. It is defined as the square of the longitudinal momentum of the reconstructed π^0 in the frame in which the longitudinal momentum component of the $\pi^+\pi^-$ pair is 0 (see Figure 6.7). Mathematically, it can be expressed as

²Due to the rarity of $K_L \longrightarrow e^+e^-\mu^+\mu^-$, signal events cannot be used for these kind of studies. However, the topological similarities between the normalization and signal modes imply that any systematic effects are probably common between the two datasets, leading to their cancellation in the branching ratio calculation.




Figure 6.7: Definition of the variable *pp0kine*. This kinematic quantity is defined as the square of the longitudinal momentum of the reconstructed π^0 in the frame in which the longitudinal momentum component of the $\pi^+\pi^-$ pair is 0. *pp0kine* should be positive for true $K_L \longrightarrow \pi^+\pi^-\pi_D^0$ decays.

$$pp0kine = \frac{(M_K^2 - M_{\pi^0}^2 - M_{\pi\pi}^2)^2 - 4M_{\pi^0}^2 M_{\pi\pi}^2 - 4M_K^2 (P_t^2)_{\pi\pi}}{4[(P_t^2)_{\pi\pi} + M_{\pi\pi}^2]}, \quad (6.1)$$

where $M_{\pi\pi}$ is the invariant mass of the $\pi^+\pi^-$ pair, $(P_t^2)_{\pi\pi}$ is the square of the transverse $\pi^+\pi^-$ momentum with respect to the kaon line of flight, and M_K is the kaon mass (Senyo 1999). For properly reconstructed $K_L \longrightarrow \pi^+\pi^-\pi_D^0$ decays, *pp0kine* should be positive. A cut was made at *pp0kine* $> -0.005 \text{ GeV}^2/c^2$ to allow for resolution effects. *pp0kine* distributions for 1999 data and monte carlo are shown in Figure 6.8(b).

A large background contribution to the normalization mode dataset was $K_L \longrightarrow \pi^+\pi^-\pi^0$ in which one of the photons from the π^0 decay converted into a e^+e^- pair in the material upstream of the first drift chamber (DC1). This led to a final state identical to that of $K_L \longrightarrow \pi^+\pi^-\pi_D^0$. In order to discriminate between the two, the e^+e^- hit separation distance at the first drift chamber was calculated. Because DC1 was located immediately downstream of the vacuum tank, the separation distance of conversion events tended to be very small. Furthermore, drift chamber performance was not very well understood at small track separations. These resolution effects also helped to contribute to a significant data/monte carlo




Figure 6.8: Data (dots) vs. monte carlo (histogram) comparisons for 1999 $K_L \rightarrow \pi^+ \pi^- \pi_D^0$ events. The histograms are normalized to each other by area. (a) and (b) Overlays in the kinematic parameters P_t^2 and $pp0kine$. The data excesses outside the cuts at $0.00025 \text{ GeV}^2/c^2$ and $-0.005 \text{ GeV}^2/c^2$, respectively, result from background and misreconstructed events. (c) Overlay in $e^+ e^-$ hit separation at DC1. The large disagreement at low separation is caused by a combination of photon conversion background and resolution effects. This region is removed from the dataset by requiring the separation to be greater than 2 mm. (d) Photon energy, as measured by the CsI calorimeter. The disagreement at low energies is caused by a combination of calibration, resolution, accidental activity, and bremsstrahlung effects. The cut at 2 GeV removes this region from the dataset.

discrepancy at small separation distances, seen clearly in Figure 6.8(c). A cut on the e^+e^- hit separation was made at 2 mm in order to remove this region from the dataset.

A comparison of Dalitz photon energy distributions between data and monte carlo shows strong disagreement at low values of E_γ (see Figure 6.8(d)). Because very low energy calibration electrons were kicked outside the fiducial volume of the experiment by the analysis magnet, it was difficult to accurately measure the calorimeter response in this energy regime. Soft bremsstrahlung and accidental photons also complicated understanding of the photon response at low energies. A cut on photon energy was therefore made at 2 GeV in order to avoid these systematic effects.

Several other cuts were applied to guarantee the quality of the remaining normalization sample. Pion showers from hadronic interactions in the calorimeter were not understood as well as those from EM interactions. This led to the concern that pion-photon shower overlap would not be modeled correctly by the monte carlo. For this reason, events in which the photon cluster was located within 20 cm of either charged pion cluster were thrown out. Events were also thrown out if they contained less than three or more than five hardware clusters. This wide range allowed for the possibility of one or both pions showering in the calorimeter. Cuts on pion track momentum (> 7 GeV/c) and pion track offset at the analysis magnet (< 2 mm) were made to the normalization mode in order to cancel any possible systematic effects from similar cuts on the muons in the signal mode. Finally, the cut on the mass of the reconstructed π^0 was tightened to $125 \text{ MeV}/c^2 < M_{e^+e^-\gamma} < 145 \text{ MeV}/c^2$.




Figure 6.9: Reconstructed M_K and M_{π^0} mass distributions for 1997 ((a) & (b)) and 1999 ((c) & (d)) after all cuts. Histogram is monte carlo, dots are data.

6.6.2 Persistent acceptance effects

12,490 normalization events survived all cuts in the 1997 data, while 458,448 remained from the 1999 data³. Reconstructed M_K and M_{π^0} mass distributions for these events are shown in Figure 6.9. Comparison between data and monte carlo in vertex z position shows excellent agreement in 1997. However, there is a small

³Recall from Section 3.6 that the L3 tags used in 1997 were prescaled by a factor of 500 at L1, while the tags used in 1999 were only prescaled by 20 (10 for the first 5% of the run) at the online split stage.




Figure 6.10: 1997 $K_L \rightarrow \pi^+ \pi^- \pi_D^0$ vertex z (left) and kaon momentum (right) distributions after all cuts. Histogram is monte carlo, dots are data. While there is excellent agreement in vertex z position, a small slope is seen in the data/monte carlo ratio in kaon momentum. This slope will be discussed further in Section 6.8.1.

slope in kaon momentum when the ratio is taken between data and monte carlo (see Figure 6.10). Fortunately, this disagreement appears to be insignificant, as shown in Section 6.8.1. The disagreement between 1999 data and monte carlo in vertex z appears to be more serious (shown in Figure 6.11). The excess of data at the downstream end of the decay volume leads to a 3.5σ slope in the data/monte carlo ratio. This effect is still not understood, and will be discussed further in Section 6.8.1.




Figure 6.11: 1999 $K_L \rightarrow \pi^+\pi^-\pi^0$ vertex z (left) and kaon momentum (right) distributions after all cuts. Histogram is monte carlo, dots are data. While there is good agreement in kaon momentum, an excess of data at the downstream end of vertex z leads to a slope in the data/monte carlo ratio. This slope will be discussed further in Section 6.8.1.

Other cut variables that exhibit severe data/monte carlo disagreement in the normalization mode are vertex χ^2 and P_t^2 (problems common to most KTeV analysis modes). These overlays are shown in Figure 6.12. The source of these discrepancies has only recently been identified. In the default monte carlo simulation of δ -rays (Section 5.4.3), it was assumed that only one interaction could occur in each drift chamber cell, that the δ -ray was knocked in a direction perpendicular to the trajectory of the parent particle, and that no transverse momentum was imparted to the parent particle during the interaction. A new monte carlo simulation of δ -rays, in which these assumptions are avoided, provides much better agreement in χ^2 and P_t^2 . The new simulation, however, was not used for this analysis.

Finally, there is also a discrepancy between data and monte carlo in the electron E/p distributions for all three runs: Winter, Summer, and 1999. The monte carlo seems to overestimate the E/p resolution in all three datasets (see Figure 6.13).

Systematic errors due to all these persistent acceptance effects will be discussed further in Section 6.8.1.

6.6.3 Charged pion loss correction

In the later stages of this analysis, it was discovered (Kessler 2001) that charged pions could be lost through several mechanisms that were not simulated in the default monte carlo. Charged pion clusters in the calorimeter could be lost if the pion showered in the VV' trigger hodoscopes, or if the pion traveled down the crack between two CsI crystals. This would lead to an inefficiency in track-cluster matching for pions. Since charged pions are present in the normalization mode but not in the signal mode, this loss might be expected to bias the kaon flux estimate, and therefore, the measurement of the branching ratio.

The track-cluster matching inefficiency was measured by reconstructing

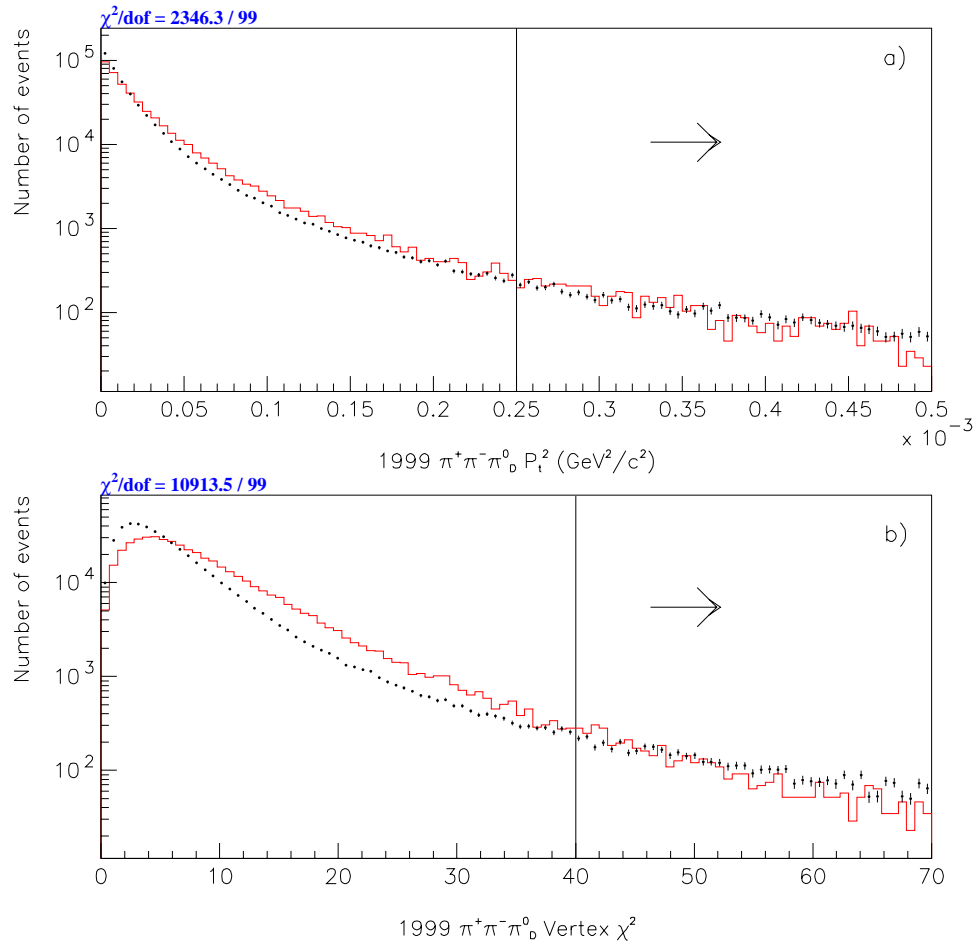


Figure 6.12: 1999 $K_L \rightarrow \pi^+ \pi^- \pi_D^0$ P_t^2 (a) and vertex χ^2 (b) distributions after all other cuts have been made. Histogram is monte carlo, dots are data. The lines on the overlays show the positions of the cuts on these variables. The disagreement seen here will be discussed further in Section 6.8.1.




Figure 6.13: $K_L \rightarrow \pi^+ \pi^- \pi_D^0$ electron E/p for (a) Winter, (b) Summer, and (c) 1999 after all other cuts. Dots are data, histogram is monte carlo. The disagreement seen here will be discussed further in Section 6.8.1.




Figure 6.14: Track-cluster matching inefficiency for charged pions from $K_L \rightarrow \pi^+\pi^-\pi^0$ events, as a function of track momentum in GeV/c. The distribution is approximated by a third-order polynomial.

$K_L \rightarrow \pi^+\pi^-\pi^0$ decays and searching for charged pion tracks that failed to match to clusters in the calorimeter (Kessler 2001). When plotted as a function of track momentum, the inefficiency spectrum is well fit by a third-order polynomial, as shown in Figure 6.14. This fit to the inefficiency data was used to tag “lost” pions during monte carlo generation. These events were thrown out during reconstruction in order to account for pion loss mechanisms in the data. This resulted in approximately a 1% correction to the normalization acceptance.

	1997	1999	Total
N	12,490	458,448	—
P	500	10/20	—
A	$1.52 \pm 0.01\%$	$1.60 \pm 0.01\%$	—
\mathcal{F}	$2.74 \pm 0.03 \times 10^{11}$	$3.64 \pm 0.02 \times 10^{11}$	$6.39 \pm 0.04 \times 10^{11}$

Table 6.2: Kaon flux results based on $K_L \longrightarrow \pi^+\pi^-\pi_D^0$ normalization events. N , P , A , and \mathcal{F} are defined in the text. Recall that the normalization mode prescale was changed from 10 to 20 about 5% into the 1999 run. Errors are purely statistical.

6.6.4 Kaon flux calculation

The total number of kaons decaying in the detector during a given run, \mathcal{F} , is given by:

$$\mathcal{F} = \frac{N \times P}{A \times \mathcal{B}} \quad (6.2)$$

where N is the number of data events after all cuts, P is the prescale used during the run, A is the acceptance as measured from the monte carlo, and \mathcal{B} is the branching ratio of the normalization mode: $\mathcal{B}(K_L \longrightarrow \pi^+\pi^-\pi_D^0) = (1.504 \pm 0.047) \times 10^{-3}$ (Groom et al. 2000b). Results for \mathcal{F} are shown for 1997 and 1999 in Table 6.2.

6.7 Signal Mode

6.7.1 Further cuts

Two main sources of background are evident in Figure 6.5(b). The low-mass peak arises mostly from $K_L \longrightarrow \pi^+\pi^-\pi_D^0$ decays in which the pions either decay to muons in flight or punch through the muon steel and fire the muon trigger. Most of this background was rejected by requiring exactly two hardware clusters in order to cut on the extra photon from the Dalitz decay. The broad high-mass feature is mostly due to two $K_L \longrightarrow \pi\mu\nu$ ($K_{\mu 3}$) decays occurring in the same RF bucket.

If both pions showered in the calorimeter (which happened about 9% of the time), the pions would be identified as electrons due to their large values of E/p , and the final state would resemble that of the signal mode. However, since two separate two-track decays in the same RF bucket are unlikely to yield a good four-track vertex during reconstruction, the cut on vertex χ^2 (described in Section 6.5) was effective in reducing background from this source. Figure 6.15 demonstrates the power of these two cuts in eliminating the visible background. Note that a peak is now visible at the kaon mass.

Additional cuts specific to the signal mode made at this point were:

- muon track offset at analysis magnet < 2 mm
- muon cluster energy < 1.5 GeV
- muon track momentum > 7 GeV/c
- particle charges consistent with $e^\pm e^\mp \mu^\pm \mu^\mp$.

These cuts are illustrated in Figure 6.16. The cut on the muon track offset at the analysis magnet helped to reject muons that resulted from pion decay inside the spectrometer. The cut on muon cluster energy ensured that the cluster was left by a MIP. Since the momentum threshold of the muon system is approximately 7 GeV/c, the cut on track momentum served to verify that the track could really have come from a muon.

6.7.2 Backgrounds

Photon conversion background

The most significant source of background was $K_L \rightarrow \mu^+ \mu^- \gamma$ when the photon converted to a $e^+ e^-$ pair in the material upstream of DC1. This material consisted of the vacuum window, the first helium bag, a small air gap (to allow for the




Figure 6.15: (a) Vertex χ^2 distributions of signal monte carlo (histogram) and data (dots). The line shows the cut at 40. (b) Signal invariant mass distribution after the vertex χ^2 and extra hardware cluster cuts. Compare to Figure 6.5(b).




Figure 6.16: $e^+e^-\mu^+\mu^-$ monte carlo distributions for (a) muon track offset at the analysis magnet, (b) muon cluster energy in the CsI, and (c) muon track momentum. Entries in the 0 bin in (b) are from beam hole muons. The lines show the cuts at 2 mm, 1.5 GeV, and 7 GeV/c, respectively.

lowering of a safety shutter in front of the window during accesses to the experiment), and the upstream mylar window of DC1. The conversion probability was measured for each run by searching for charged tracks in reconstructed $K_L \rightarrow 3\pi^0$ decays (Zimmerman 1999), and was found to be 0.00241 in the Winter run, 0.00246 during the Summer run, and 0.00218 in 1999 (LaDue and Toale). Conversions were more likely in 1997 due to a leak in the first helium bag (effectively turning it into an air bag). The leak was fixed in preparation for the 1999 run.

The $K_L \rightarrow \mu^+\mu^-\gamma$ conversion background is particularly dangerous, as it reconstructs in the invariant mass signal region. Based on monte carlo studies, about 47 background events could be expected in the signal region at this stage of cuts. Since conversions occurred just upstream of DC1, the electron separation measured at DC1 is very small. Another characteristic is low reconstructed M_{ee} , the invariant mass of the e^+e^- pair. Figure 6.17 shows these variables in two-dimensional space for conversion monte carlo, signal monte carlo, and all signal data. Monte carlo predicts that cutting out the boxed region (hit separation > 2 mm or $M_{ee} > 2.75$ MeV/c²) rejects 99% of conversion events, while retaining 85% of the signal. This cut removed 57 events from the data (67 events were predicted by monte carlo, a 1.2σ discrepancy) leaving an expected background of 0.680 ± 0.013 events from $K_L \rightarrow \mu^+\mu^-\gamma$ conversions.

Pion decay/punchthrough background

As mentioned earlier, charged pion decay and/or punchthrough could fire the muon trigger. Although the extra hardware cluster cut significantly reduced the background from $K_L \rightarrow \pi^+\pi^-\pi_D^0$, events in which the photon has low energy or is lost survived the cut. The decay modes $K_L \rightarrow \pi^+\pi^-e^+e^-$ and $K_L \rightarrow \pi^+\pi^-\gamma$ might also be expected to contribute to this background (although a photon conversion would also be required in the case of $\pi^+\pi^-\gamma$).

The ratio of decay probability to punchthrough probability was measured by

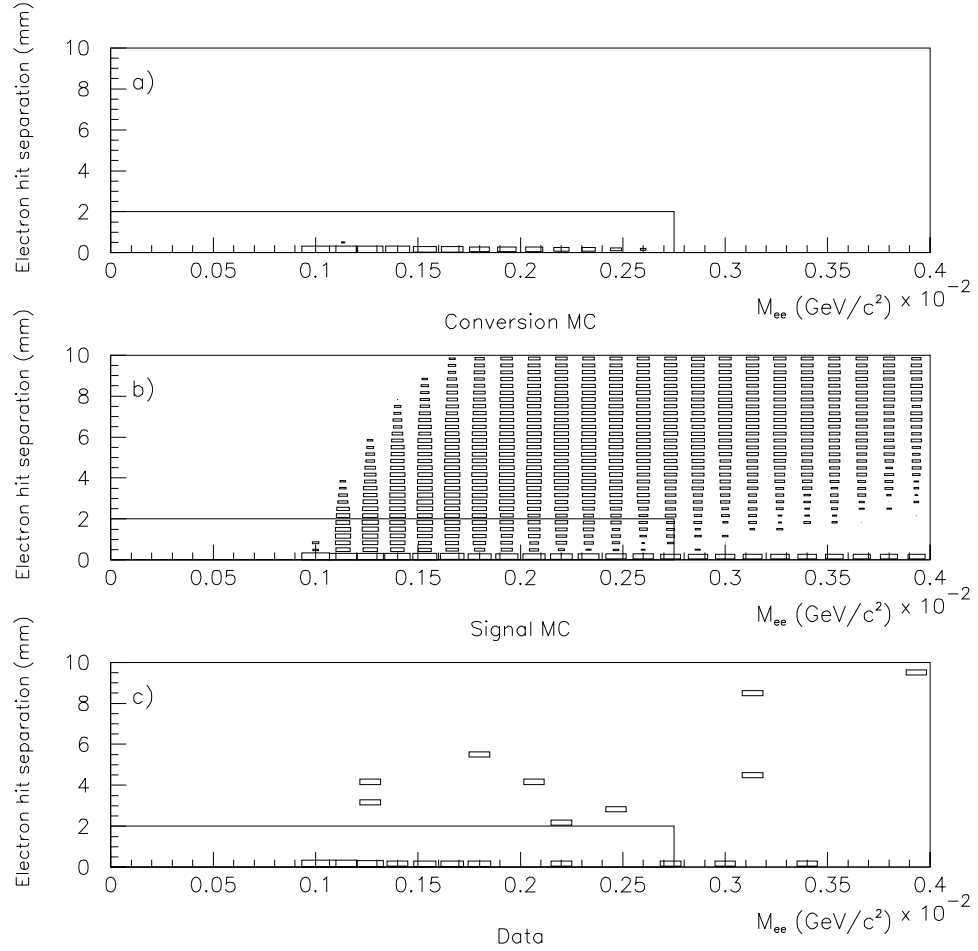


Figure 6.17: Electron hit separation at DC1 (in millimeters) vs. M_{ee} (in GeV/c^2) for (a) conversion monte carlo, (b) signal monte carlo, and (c) signal data. Plots are logarithmic in z . Cutting out the box rejects 99% of the conversions, leaving approximately 0.7 expected background events from this source.

reconstructing signal tags in the data as $K_L \longrightarrow \pi^+\pi^-\pi_D^0$ and applying all signal cuts, except the cut on extra hardware clusters. $\pi^+\pi^-\pi_D^0$ monte carlo sets were generated with different pion decay/punchthrough ratios, accepted on the dimuon trigger, and reconstructed in the same manner. Events with pion decays tended to reconstruct poorly due to the smeared kinematics, while punchthroughs tended to reconstruct in a much narrower peak. Comparison of the subsequent data/monte carlo overlays in $M_{\pi^+\pi^-\pi_D^0}$ and $E_{\pi^+\pi^-\pi_D^0}$ show that a pion decay/punchthrough ratio of 60/40 produces the best agreement for 1997 data, while a 55/45 ratio appears to match better for the 1999 data (see Figures 6.18 and 6.19).

Because the pion momentum spectrum was slightly softer in 1999 (due to the lower P_t kick from the analysis magnet – see Figure 6.20), one would actually expect the decay probability to be slightly *higher* than during 1997. Within the errors of the overlays, however, no real significant conclusion can be made regarding the relative values of the fit results for the 1997 and 1999 decay/punchthrough probabilities. Additionally, when reconstructed as signal, the shape of the low-mass background (and hence the background estimate) turns out to be quite insensitive to the choice of decay/punchthrough ratio. This is demonstrated in Figure 6.21.

Several hundred million background $K_L \longrightarrow \pi^+\pi^-\pi_D^0$ and $K_L \longrightarrow \pi^+\pi^-\pi^+\pi^-$ monte carlo events were generated in which each charged pion was forced to either decay or punch through the muon steel in the appropriate ratio. (Because the shape of $\pi^+\pi^-\gamma$ background distribution is virtually identical to that of $\pi^+\pi^-\pi^+\pi^-$, no monte carlo was generated for the former mode. The background estimate from $\pi^+\pi^-\pi^+\pi^-$ was scaled appropriately to account for the presence of $\pi^+\pi^-\gamma$ events.) The invariant mass distributions of these backgrounds are plotted in Figure 6.22(a) and (b).

In order to obtain an estimate of the total background from pion decay/punchthrough events, the background monte carlo datasets were added together in a proportion that was allowed to float. The invariant mass distribution of the

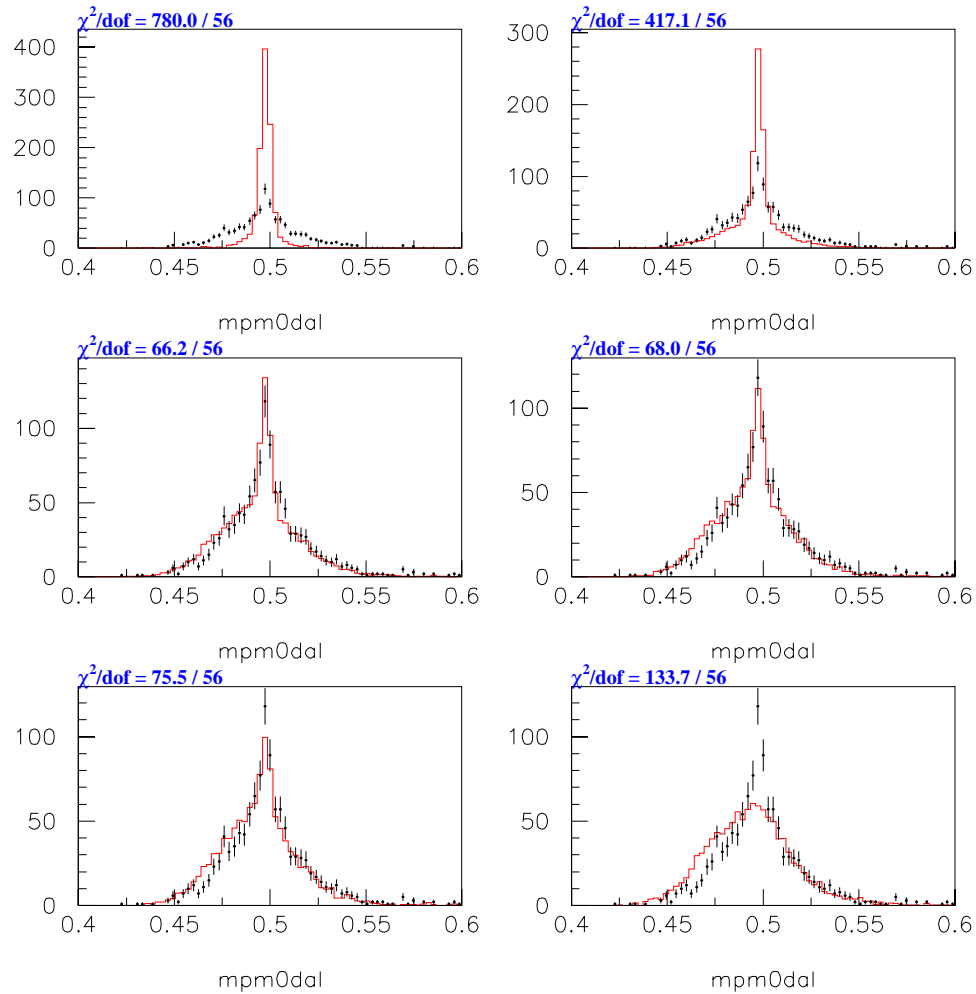


Figure 6.18: Invariant mass distributions of 1997 $\pi^+\pi^-\pi_D^0$ data (dots) reconstructed in the signal trigger. Overlaid are monte carlo histograms generated with various pion decay/punchthrough ratios. Upper left plot shows 0% decay/100% punchthrough MC, upper right is 25/75, middle left is 60/40, middle right is 65/35, lower left is 70/30, and lower right shows 100% decay/0% punchthrough MC. The best agreement seems to come from 60% decay/40% punchthrough (middle left).

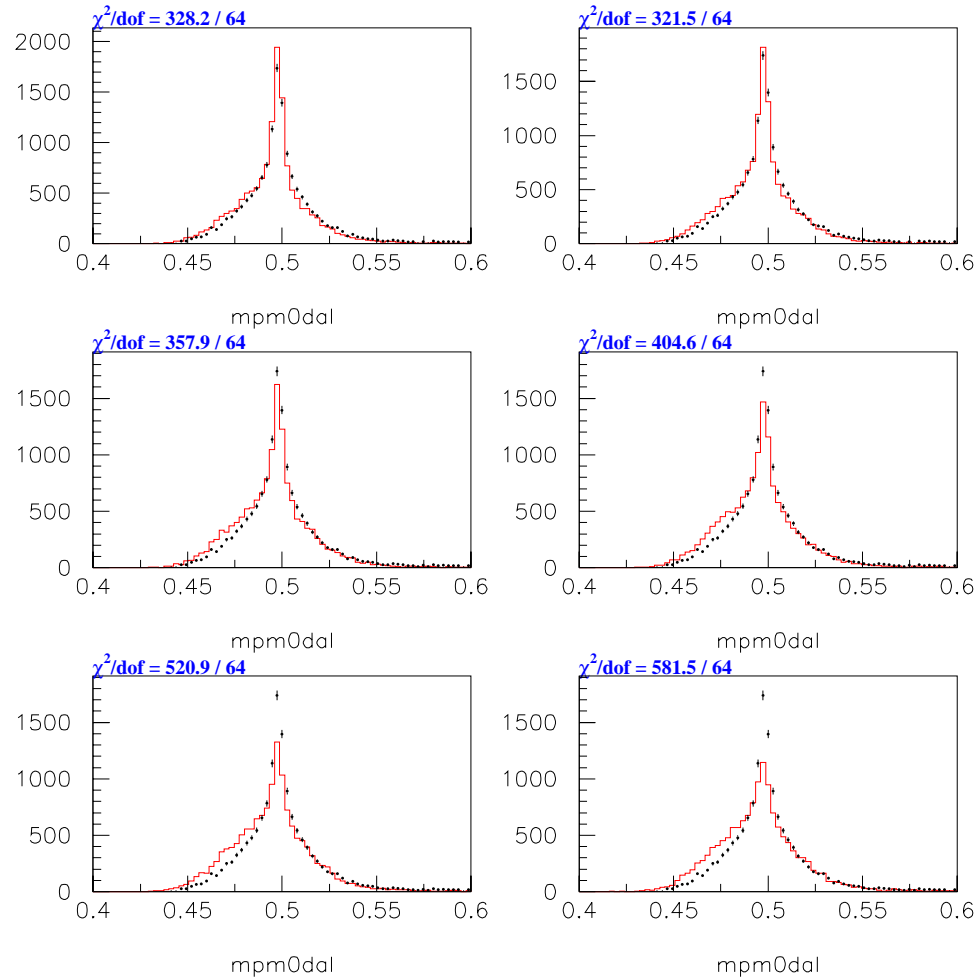


Figure 6.19: Invariant mass distributions of 1999 $\pi^+\pi^-\pi_D^0$ data (dots) reconstructed in the signal trigger. Overlaid are Monte Carlo histograms generated with various pion decay/punchthrough ratios. Upper left shows 50% decay/50% punchthrough, upper right is 55/45, middle left is 60/40, middle right is 65/35, lower left is 70/30, and lower right is 75/25. The best agreement seems to come from 55% decay/45% punchthrough (upper right).




Figure 6.20: Distributions of pion track momenta from 1997 (top) and 1999 (bottom) $K_L \rightarrow \pi^+\pi^-\pi_D^0$ events. All cuts have been made except for $p_\pi > 7$ GeV/c. Note that the 1999 spectrum is slightly softer than in 1997, due to the lower P_t kick from the analysis magnet. This would be expected to contribute to a larger pion decay probability, which is not seen in the fits.




Figure 6.21: Invariant mass comparison between 80% pion decay, 20% pion punchthrough (dots) and 60% decay, 40% punchthrough (histogram) monte carlo after reconstruction as signal. The shape of the reconstructed invariant mass distribution, and hence the background estimate, is quite insensitive to the choice of decay/punchthrough ratio.

monte carlo sum was then normalized by area to the low mass region in the data⁴. The proportion between $K_L \longrightarrow \pi^+\pi^-\pi_D^0$ and $K_L \longrightarrow \pi^+\pi^-e^+e^-$ monte carlo was varied until the χ^2 between the data/monte carlo distributions was minimized. The final overlay with signal data is shown in Figure 6.22(c). Using the constants found from this relative normalization, a total of 0.064 ± 0.033 background events were estimated, dominated by the contribution from $K_L \longrightarrow \pi^+\pi^-\pi_D^0$ ⁵.

Double decay background

Remaining background from double $K_{\mu 3}$ decays (two decays occurring in the same RF bucket) was estimated without the use of monte carlo by assuming that wrong-signed ($e^\pm e^\pm \mu^\mp \mu^\mp$) events are double decays. Double decays yielding these wrong-signed events are expected to occur just as often as those mimicking the correctly-signed $e^+e^-\mu^+\mu^-$. Figure 6.23 shows the distribution of all $e^\pm e^\pm \mu^\mp \mu^\mp$ events in the P_t^2 -invariant mass plane. Assuming that the events are evenly distributed throughout the plane, a linear extrapolation into the signal region gives a background estimate of 0.080 ± 0.010 events from double $K_{\mu 3}$ decays.

6.7.3 Background and cut summary

Table 6.3 summarizes the cuts applied to the normalization and signal samples. After all cuts, 132 $K_L \longrightarrow e^+e^-\mu^+\mu^-$ events remain in the signal region – 46 from the 1997 run and 86 from 1999 (see Figures 6.24 and 6.25). The expected background contributions for 1997 and 1999 are summarized in Table 6.4.

⁴An absolute normalization would require knowing the probability of punchthrough as a function of pion momentum, as well as the probability of photon loss as a function of energy – distributions that are both very model-dependent.

⁵The dominance of the contribution from $K_L \longrightarrow \pi^+\pi^-\pi_D^0$ is due to the relatively tiny branching ratios of $K_L \longrightarrow \pi^+\pi^-e^+e^-$ (3.5×10^{-7}) and $K_L \longrightarrow \pi^+\pi^-\gamma$ combined with a vacuum window photon conversion ($\approx 1.3 \times 10^{-7}$).




Figure 6.22: Invariant mass distributions of (a) $\pi^+\pi^-e^+e^-$ and $\pi^+\pi^-\gamma$, and (b) $\pi^+\pi^-\pi^0_D$ pion decay/punchthrough monte carlo, when reconstructed as signal. Overlay of the sum of these datasets with data is shown in (c). The signal region has been blacked out to properly normalize the distributions. Approximately 0.06 background events are expected from decay/punchthrough events.




Figure 6.23: P_t^2 vs. invariant mass of all $e^\pm e^\pm \mu^\mp \mu^\mp$ events. The box shows the location of the signal region in the plane. Extrapolating into the signal region yields an estimate of 0.08 background events from double $K_{\mu 3}$ decays.

$K_L \rightarrow \pi^+ \pi^- \pi_D^0$	$K_L \rightarrow e^+ e^- \mu^+ \mu^-$
$20 \text{ GeV}/c < p_K < 220 \text{ GeV}/c$	
$90 \text{ m} < \text{vertex } z < 158 \text{ m}$	
vertex in beam hole	
$0.95 < E/p < 1.05 \text{ for electrons}$	
vertex $\chi^2 < 40$	
$P_t^2 < 0.00025 \text{ GeV}^2/c^2$	
$0.482 \text{ GeV}/c^2 < M_K < 0.512 \text{ GeV}/c^2$	
$p_\pi > 7 \text{ GeV}/c$	$p_\mu > 7 \text{ GeV}/c$
π track magnet offset $< 2 \text{ mm}$	μ track magnet offset $< 2 \text{ mm}$
DC1 $e^+ e^-$ hit sep $> 2 \text{ mm}$	DC1 $e^+ e^-$ hit sep $2 > \text{mm}$ OR $M_{ee} > 2.75 \text{ MeV}/c^2$
$3 \leq N_{hclus} \leq 5$	$N_{hclus} = 2$
all tracks matched to clusters	1 muon allowed down beam hole
$pp0kine > -0.005 \text{ GeV}^2/c^2$	
$125 \text{ MeV}/c^2 < M_{e^+e^-\gamma} < 145 \text{ MeV}/c^2$	
$E_\gamma > 2 \text{ GeV}$	
$\pi - \gamma$ cluster separation $> 20 \text{ cm}$	
	$E_\mu^{Csl} < 1.5 \text{ GeV}$

Table 6.3: Summary of all cuts applied to normalization and signal samples.

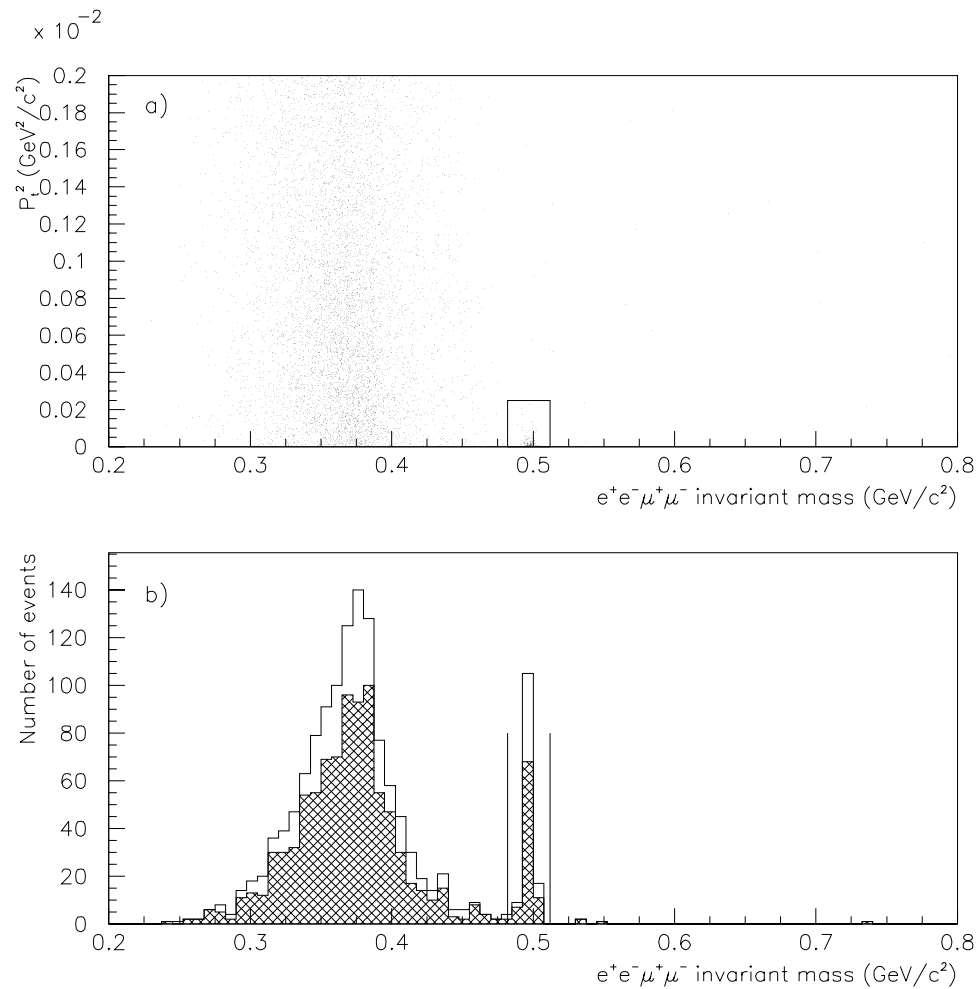


Figure 6.24: (a) Signal P_t^2 vs. invariant mass for all signal data after all other cuts. 132 events are in the box, which denotes the signal region. (b) Projection onto the invariant mass axis after the cut on P_t^2 . Events from 1997 are plotted in the open histogram, while events from 1999 are plotted in the cross-hatched histogram. The lines mark the boundaries of the signal region.




Figure 6.25: (a) Invariant mass comparison between the 132 signal data events (dots) and signal monte carlo (histogram). (b) Comparison in P_t^2 . The data/monte carlo agreement is good in both variables.

	1997	1999	All data
$K_L \longrightarrow \mu^+\mu^-\gamma$ conversions	0.261 ± 0.006	0.419 ± 0.012	0.680 ± 0.013
$K_{\mu 3}$ double decays	0.030 ± 0.006	0.050 ± 0.008	0.080 ± 0.010
Pion decay/punchthrough	0.017 ± 0.013	0.047 ± 0.030	0.064 ± 0.033
Total background	0.308 ± 0.016	0.516 ± 0.033	0.824 ± 0.037

Table 6.4: Backgrounds to $K_L \longrightarrow e^+e^-\mu^+\mu^-$ for 1997, 1999, and total.

6.8 Branching Ratio Calculation

The branching ratio \mathcal{B} is calculated using the following formula:

$$\mathcal{B}(K_L \longrightarrow e^+ e^- \mu^+ \mu^-) = \frac{N}{A \times \mathcal{F}} \quad (6.3)$$

where N is the background-subtracted number of events, A is the acceptance calculated from the monte carlo, and \mathcal{F} is the kaon flux, calculated in Section 6.6.

The radiative corrections included in the signal monte carlo have a fairly large impact on the signal acceptance. If the radiated photon carries a sufficient amount of energy, the invariant mass of the reconstructed signal event will drop below the cut at $0.482 \text{ GeV}/c^2$, and the event would be lost. Therefore, including events from the whole spectrum of radiated photon energies would lead to an extrapolation into a region with no signal acceptance at all. To avoid this effect, the acceptance calculation was limited to events passing an infrared cutoff at $x_{e^+e^-\mu^+\mu^-} > 0.95$, where

$$x_{e^+e^-\mu^+\mu^-} = \frac{M_{e^+e^-\mu^+\mu^-}^2}{M_K^2}. \quad (6.4)$$

This ratio is related to E^* , the energy of the radiated photon in the kaon rest frame, by

$$E^* = \frac{M_K}{2} (1 - x_{e^+e^-\mu^+\mu^-}). \quad (6.5)$$

The cutoff at $x_{e^+e^-\mu^+\mu^-} > 0.95$ therefore corresponds to a E^* value of $0.025M_K$, or 12.4 MeV (Barker). This restricts the calculation of the signal acceptance to approximately 65% of the E^* spectrum.

The branching ratio results for 1997, 1999, and the complete dataset are shown in Table 6.5. The χ^2 between the 1997 and 1999 values comes to $0.28/1 \text{ dof}$, which translates to a probability of agreement of 60%.

	1997	1999	All data
N	45.692	85.484	—
A	$6.77 \pm 0.05\%$	$8.85 \pm 0.08\%$	—
\mathcal{B}	$2.48 \pm 0.37 \times 10^{-9}$	$2.73 \pm 0.29 \times 10^{-9}$	$2.63 \pm 0.23 \times 10^{-9}$

Table 6.5: $K_L \rightarrow e^+e^-\mu^+\mu^-$ branching ratios for 1997, 1999, and all data. A and \mathcal{B} are calculated using a radiative cutoff at $x_{e^+e^-\mu^+\mu^-} > 0.95$. The branching ratio for the complete dataset was calculated by taking the weighted average of the 1997 and 1999 branching ratios. Errors at this stage are purely statistical.

Vertex z range (m)	Total kaon flux ($\times 10^{11}$)	% change in total flux
90–158 (nominal)	6.39 ± 0.04	—
90–112	6.36 ± 0.06	$-0.40 \pm 0.81\%$
112–129	6.34 ± 0.06	$-0.80 \pm 0.81\%$
129–158	6.47 ± 0.07	$+1.24 \pm 0.83\%$

Table 6.6: Total kaon flux, calculated in equally populated bins of vertex z . As might be expected from Figure 6.11, the largest variation from the nominal flux occurs at the most downstream bin.

6.8.1 Systematic errors

1999 z slope

As mentioned in Section 6.6, discrepancies between data and monte carlo exist in several cut variables. One of the most obvious of these is the disagreement of the 1999 $\pi^+\pi^-\pi_D^0$ vertex z distribution at large z (Figure 6.11). In order to obtain a conservative estimate for the systematic effect from this mismeasurement on the branching ratio, the normalization z slope was simply integrated over the range of interest, 90 – 158 m, giving an error of $1.64 \pm 0.56\%$.

Three different methods were used as crosschecks. In the first, the total (1997 + 1999) kaon flux was calculated in equally populated bins of vertex z . The results are shown in Table 6.6. While the flux should be independent of z , there




Figure 6.26: 1999 $K_L \rightarrow \pi^+ \pi^- \pi_D^0$ vertex z (left) and kaon momentum (right) distributions after all cuts, with the cut on electron separation at DC1 moved from 2 to 25 mm. Histogram is monte carlo, dots are data. Note that the disagreement at the downstream end in vertex z is much smaller than in Figure 6.11, and the slope of the data/monte carlo ratio is now consistent with 0.

is a significant shift ($1.24 \pm 0.83\%$) in the most downstream bin. This should not be surprising, as this is where the data/monte carlo disagreement is located in Figure 6.11. The second crosscheck makes use of the fact that the disagreement in 1999 vertex z seems to be correlated with the cut on electron separation at DC1. When the cut is moved from 2 to 25 mm, the downstream disagreement shrinks, and the slope becomes consistent with 0 (see Figure 6.26). Tightening the cut to

25 mm in both 1997 and 1999 changes the total kaon flux by $0.97 \pm 0.37\%$. In the third crosscheck, 1999 normalization and signal data were reweighted to match the monte carlo in vertex z . This shifted the total branching ratio by $0.66 \pm 0.90\%$.

All three crosschecks produced answers consistent with the estimate obtained by integrating the vertex z slope over the fiducial region. This systematic error due to misunderstanding of the vertex z distribution in 1999, $1.64 \pm 0.56\%$, translates to a total effect of 1.91% when 68% (1σ) of the error gaussian area is integrated from 0.⁶ It should be noted that this *is* a rather conservative estimate, as it assumes the data/monte carlo discrepancy is limited only to the normalization mode, and does not similarly affect the signal mode. Since acceptance effects common to the two modes cancel in the calculation of the branching ratio, any systematic estimate that depends solely on the normalization mode is probably an overestimate of the actual error.

1997 p slope

Similar methods were used to determine the size of the systematic effect from the slope in 1997 kaon momentum, shown in Figure 6.10. Simply integrating the 1997 p slope from 20 to 220 GeV/c provides an error estimate of $2.38 \pm 1.80\%$. However, there is ample reason to believe that the effect is not statistically significant. Table 6.7 shows how the total kaon flux varies when calculated in equally populated p bins. No p dependence is evident in the flux measurements. Furthermore, when the 1997 normalization and signal data were reweighted to account for the observed slope in p , the total branching ratio shifted by only $0.56 \pm 1.18\%$. Because there is no compelling evidence for an effect due to the 1997 p slope, no systematic is quoted.

⁶This method is used to account for the statistical errors on all systematic effects in this analysis.

p range (GeV/c)	Total kaon flux ($\times 10^{11}$)	% change in total flux
20–220 (nominal)	6.39 ± 0.04	—
20–60	6.38 ± 0.06	$-0.11 \pm 0.89\%$
60–83	6.39 ± 0.06	$0.0 \pm 0.79\%$
83–220	6.40 ± 0.06	$+0.14 \pm 0.78\%$

Table 6.7: Total kaon flux, calculated in equally populated bins of kaon momentum. No significant effect is evident.

E/p disagreement

Disagreement in E/p (shown in Figure 6.13) was isolated to the mismeasurement of monte carlo cluster energies. In order to determine a systematic error from this effect, these energies were “nudged” during clustering at the crunch stage to obtain better data/MC agreement. For the Winter monte carlo, this involved gaussian-smearing the energies by approximately 15%. For the Summer monte carlo, the energies were smeared by 17% and then shifted up by 0.27%. The 1999 monte carlo required a double-gaussian smear to match the data. Reconstructing the monte carlo with the new cluster energies led to much better looking overlays (see Figure 6.27). The branching ratio, when calculated with the new monte carlo cluster energies, shifted by 0.09%. This value was taken to be the systematic error from E/p misunderstanding.

Vertex χ^2 , P_t^2 disagreement

Vertex χ^2 and P_t^2 are complicated variables that are difficult to nudge in this fashion. As mentioned in Section 6.6.2, disagreement in these distributions was traced to the simulation of δ -rays within the drift chamber cells. Systematic errors from these variables were estimated by shifting the cuts made on them and measuring the effect on the branching ratio. Figure 6.28 shows how the branching ratio moved when these cuts were varied. Also shown is the effect on the branching ratio from shifting the




Figure 6.27: $\pi^+\pi^-\pi_D^0$ electron E/p for (a) Winter, (b) Summer, and (c) 1999 after “nudging” monte carlo cluster energies during clustering at the crunch stage. Data is dots, histogram is monte carlo. Compare to Figure 6.13. Using the new monte carlo cluster energies led to a total shift in the branching ratio of only 0.1%.




Figure 6.28: Shifts in percent of the branching ratio as the cuts on P_t^2 , vertex χ^2 , and p_{MIP} are varied. Note the walk in vertex χ^2 . Deviations in the other two variables appear to be statistical.

momentum cut on the tracks associated with minimally ionizing particles (pions in the normalization mode, muons in the signal mode) – recall that the nominal value of the muon momentum cut was loosened from 10 to 7 GeV/c since the previous study of $K_L \rightarrow e^+e^-\mu^+\mu^-$ was published (Alavi-Harati et al. 2001c).

There appears to be a small walk in the branching ratio as the cut on vertex χ^2 is tightened. It can be argued on the basis of Figure 6.15(a) that the cut on vertex χ^2 is reasonable down to the value of 25, where the data and monte carlo




Figure 6.29: Shift in the total kaon flux as the cut on vertex χ^2 is varied. This provides a better estimate of the systematic effect than the walk in the branching ratio, due to the greater statistical power. An error of 1.10% is used from the deviation at a vertex χ^2 cut at 25.

distributions cross. The full systematic error at this point would correspond to a value of 3.6%. However, this number is dominated by the statistical error on the branching ratio shift. In order to get a more realistic estimate of the effect from this misunderstanding, a plot was made showing how the total kaon *flux* varies as the cut on vertex χ^2 is tightened. Because of the greater statistical power, the walk in the kaon flux gives a better estimate for the systematic error from the cut on vertex χ^2 . The deviation of 1.10%, taken at the vertex χ^2 of 25, is used.

% of events with accidental overlays	% change in branching ratio
100 (nominal)	—
75	$-0.07 \pm 0.19\%$
50	$+0.61 \pm 0.20\%$
25	$-1.55 \pm 0.17\%$
0	$-1.53 \pm 0.29\%$

Table 6.8: Changes in the branching ratio due to various populations of accidental overlays.

Accidental overlay effects

Accidental activity was simulated during the generation of monte carlo by overlaying events taken in a special out-of-time trigger, described in Sections 2.1.1 and 5.6. A cut on extra clusters has been shown to be extremely useful in the elimination of background; it was therefore important to understand how accidental activity affected the acceptance of both modes.

Monte carlo sets were generated with different fractions of events receiving accidental overlays, in order to simulate various levels of accidental activity. Table 6.8 shows that the branching ratio is somewhat dependent on the amount of overlays used. Although it is believed the simulation of accidental activity in the monte carlo is understood fairly well (see Figure 6.30), a conservative error of 1.67% was chosen (corresponding to the largest deviation in Table 6.8).

Charged pion loss uncertainties

Recall that the correction for track-cluster matching inefficiency from charged pion losses was calculated using a third-order polynomial fit to $K_L \rightarrow \pi^+\pi^-\pi^0$ data (Figure 6.14). To account for the effects of uncertainties in the fit parameters, new normalization mode monte carlo was generated in which the fit parameters were each shifted by 1σ , with the sign of the shifts chosen to maximize the effect on the

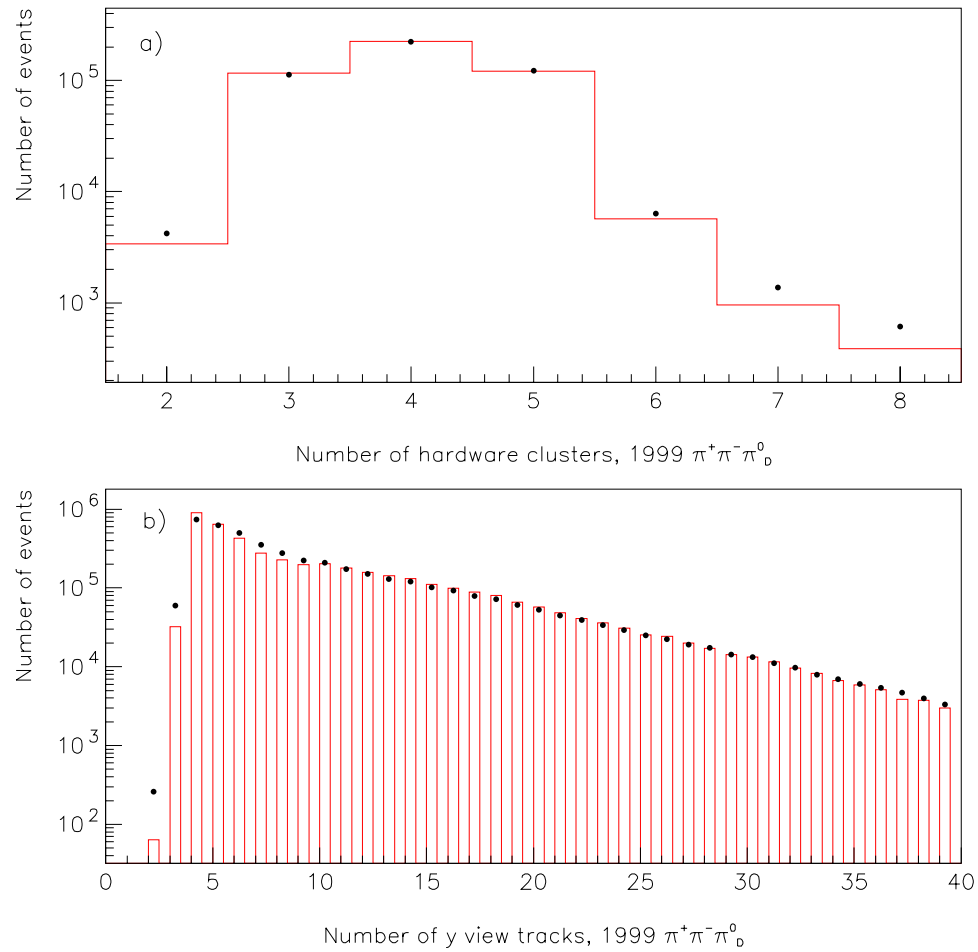


Figure 6.30: Data (dots) and monte carlo (histogram) comparisons for (a) number of hardware clusters, and (b) number of y view tracks for reconstructed 1999 $\pi^+\pi^-\pi^0$ events. Both of these variables are dependent on the level of accidental activity. The agreement seen in these overlays demonstrates that the level of simulated accidental activity in the monte carlo is fairly well understood.

overall correction. Using this new monte carlo to calculate the normalization mode acceptance led to a change of 0.33% in the branching ratio.

Muon scattering

The most difficult systematic to measure was the effect due to misunderstanding of multiple scattering through the muon filter. Much work has been done to improve the monte carlo simulation of muon scattering, details of which can be found in (Quinn 2000). However, discrepancies between data and monte carlo still remain, leading to the concern that the signal mode acceptance might be wrong.

The systematic effects of scattering were parameterized by a variable called *xdiff*, defined as the position difference in x between the muon's downstream track segment when projected to MU2 and the actual x position at MU2. The *xdiff* spectrum is, of course, strongly dependent on muon momentum.

This study made use of data from two special long muon runs taken with the analysis magnet on; one during the Winter run, and the other during the 1999 run. *xdiff* distributions were plotted in 5 GeV/c momentum bins for muons that hit the overlap region between adjacent counters in MU2, so that the x position was known to within a centimeter. Distributions were also plotted for muon run monte carlo containing the default muon scattering simulation. Data/monte carlo overlays for the first nine momentum bins of the 1997 study are shown in Figure 6.31. For reference, the mean muon momentum in generated $e^+e^-\mu^+\mu^-$ events is approximately 18 GeV/c.

It appears that the scattering width is underestimated by the monte carlo at higher values of muon momentum – an effect that is also seen in the 1999 data. While the *xdiff* distributions generated by the default monte carlo are fit well by a single gaussian, those in the data were better approximated by the sum of two gaussians, shown in Figure 6.32. A systematic error quantifying this discrepancy

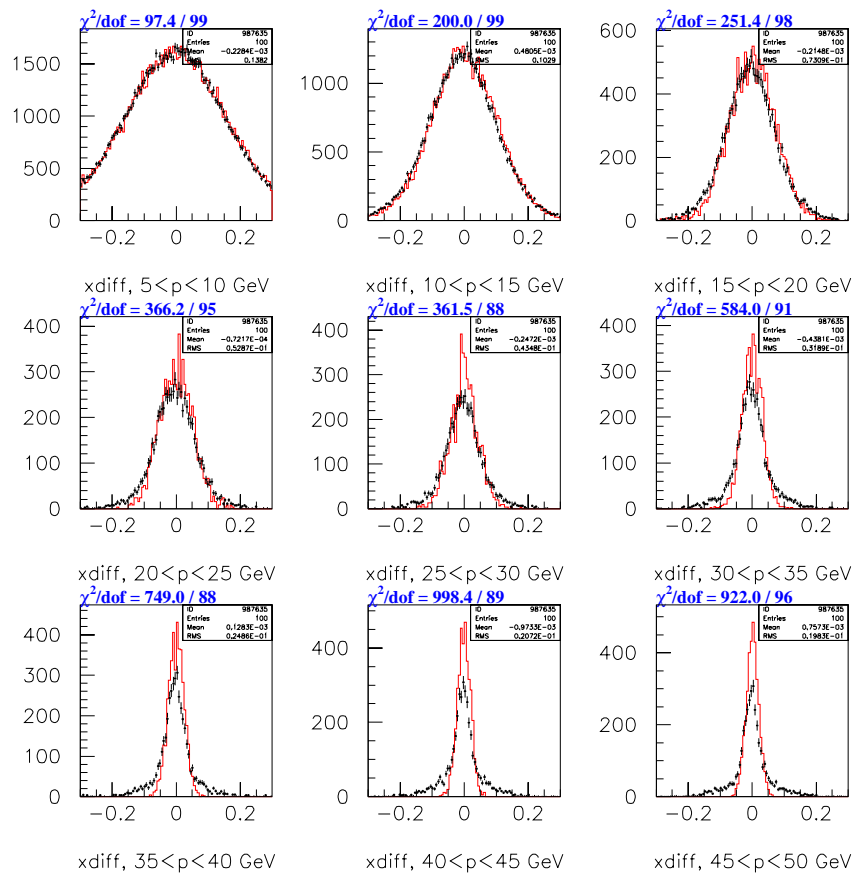


Figure 6.31: Distributions of the variable $xdiff$ (in meters) for 1997 long muon run data (dots) and monte carlo (histogram), plotted in 5 GeV/c muon momentum bins.




Figure 6.32: Fits to x_{diff} distributions of 1997 long muon run data. The distributions are well approximated by the sum of two gaussians. Reconstructing signal monte carlo with these fit parameters led to a total shift in the branching ratio of 0.55%.

was obtained by forcing the x_{diff} distributions of muons from signal monte carlo to match those found in the fits to the muon run data. This was done by manually shifting the x position of a muon the necessary distance just upstream of MU2 during generation. Doing so caused a shift in the branching ratio of 0.55%, implying that the signal acceptance is fairly insensitive to the details of the muon scattering simulation.

Trigger effects

During data collection, the L3 trigger only had access to online calibration constants. However, the monte carlo datasets used in this analysis were generated using the more precise offline constants, calculated at the end of the run. In order to determine if this data/monte carlo procedural discrepancy led to any systematic biases, new monte carlo was generated using the online database. Using this new monte carlo for the acceptance calculations led to a minute change of 0.10% in the branching ratio.

A more significant effect might be expected due to the trigger differences in the normalization mode between 1997 and 1999 (see Section 3.3.2). The four-track trigger used in 1999 required more sources, and was therefore more restrictive, than the two-track trigger used in 1997. A bias introduced by these extra sources might affect the 1999 kaon flux, and therefore the branching ratio. In order to see if this was the case, several thousand two-track events (the kind of tag used for normalization in 1997) were recovered from the raw 1999 data tapes. The tapes were taken from three time periods, spread over the 1999 run. Two-track monte carlo was also generated for these periods.

After reconstruction of the two-track events, the additional L1 and L2 trigger requirements from the four-track trigger were imposed (both individually and collectively) on the data and monte carlo events. Using the acceptance figures from the monte carlo, the kaon flux for each time period was calculated before and




Figure 6.33: Results of the four-track trigger systematic study. The top plot shows the calculated kaon flux (in arbitrary units), totaled over the three time periods used, as four-track trigger sources were individually imposed on a two-track normalization sample. The rightmost point represents a complete simulation of the four-track trigger. The bottom plot shows the difference, in percent, between each of these fluxes and the flux calculated from the default two-track normalization, taking correlations into account. There is a difference of 1.81% between the fluxes measured from the two-track and simulated four-track triggers.

after the four-track trigger simulation. Any difference between the two fluxes was interpreted as a systematic effect from the four-track trigger sources. Figure 6.33 shows the results of the study, summed over the three time periods. Application of

the additional trigger requirements required to transform the two-track trigger into the four-track trigger led to a change in the total flux of $-1.31 \pm 1.03\%$ ⁷, which integrates to a total systematic effect of 1.81%. However, since the four-track trigger was used only for 1999, which makes up 57% of the total KTeV dataset, this error was scaled down to 1.03%.

It should be noted that there appears to be a large discrepancy in the *YTF_UDO* L2 trigger source. While some of the two-track events in the data failed the YTF requirement, all of the monte carlo events had this particular trigger latch bit set. This is not currently understood. However, as a consistency check, the complete simulated four-track fluxes from this study were compared to the actual four-track fluxes from the data, and were found to be equal to within one standard deviation.

Uncertainty on the DIP parameters α and β

As mentioned in Section 5.2.1, the value of the DIP parameter α used to generate monte carlo for the signal acceptance calculation was $\alpha = -4.53$, the central value of the result from the form factor analysis ($\alpha_{shape} = -4.53^{+1.81}_{-2.70}$) to be described in Section 7.2. Of course, the uncertainties associated with this value translate into a systematic effect on the signal acceptance. To obtain an estimate of the size of this effect, new signal MC datasets were generated with $\alpha = -2.72$ and $\alpha = -7.23$, 1σ deviations from the central value in either direction. Recalculating the signal acceptance using the lower value of α led to the more significant shift in the branching ratio, 1.90%. This value was used as a systematic error. Varying the input value for β was found to have an insignificant effect on the overall signal acceptance.

⁷For the curious, flux discrepancies for each of the three individual time periods were $0.4 \pm 2.1\%$, $-3.0 \pm 1.5\%$, and $-0.5 \pm 1.9\%$.

$K_L \longrightarrow \pi^+ \pi^- \pi_D^0$ BR uncertainty	3.12%
1999 vertex z slope	1.91%
DIP α uncertainty	1.90%
Accidentals	1.67%
Monte carlo statistics	1.11%
Vertex χ^2 walk	1.10%
1997/1999 normalization trigger differences	1.03%
Muon scattering	0.55%
Pion loss uncertainties	0.33%
L3 calibration constants	0.10%
E/p disagreement	0.09%
Background uncertainty	0.03%
Total	4.87%

Table 6.9: Systematic errors on the $K_L \longrightarrow e^+ e^- \mu^+ \mu^-$ branching ratio.

6.8.2 Tabulation of systematic errors

Table 6.9 summarizes the systematics described in this section. The total of 4.87% is dominated by the external uncertainty on $\mathcal{B}(K_L \longrightarrow \pi^+ \pi^- \pi_D^0)$, and corresponds to a systematic error in the branching ratio of 0.13×10^{-9} .

The final branching ratio is therefore:

$$\mathcal{B}(K_L \longrightarrow e^+ e^- \mu^+ \mu^- ; x_{e^+ e^- \mu^+ \mu^-} > 0.95) = (2.63 \pm 0.23_{\text{stat}} \pm 0.13_{\text{syst}}) \times 10^{-9}$$

based on 132 events in the signal region, with approximately 0.8 expected background events.

6.9 Limit on $K_L \longrightarrow e^\pm e^\pm \mu^\mp \mu^\mp$ Branching Ratio

The distribution of reconstructed $e^\pm e^\pm \mu^\mp \mu^\mp$ events has already been shown in Figure 6.23. The absence of events in the signal region was used to set a limit on the branching ratio for this lepton flavor-violating decay, which is allowed under

certain extensions to the Standard Model (Technicolor and Supersymmetry, for example (Appel et al. 2000)).

Four-body phase space monte carlo was generated for the three run periods in order to obtain an acceptance estimate of 9.27%. Together with the total kaon flux of 6.39×10^{11} , a single event sensitivity of 1.69×10^{11} was calculated. With no signal events in the absence of background, this corresponds to a 90% confidence level limit⁸ of:

$$\mathcal{B}(K_L \longrightarrow e^\pm e^\pm \mu^\mp \mu^\mp) < 4.12 \times 10^{-11}$$

This is a full factor of 3 improvement over the published limit of 1.23×10^{-10} (Alavi-Harati et al. 2001c).

⁸This limit was calculated using the Feldman–Cousins methodology (Feldman and Cousins 1998).

CHAPTER 7

$K_L \gamma^* \gamma^*$ Form Factor Analysis

Now that signal $K_L \longrightarrow e^+e^-\mu^+\mu^-$ events have been isolated from the dataset, they can be used to study the dynamics of the $K_L \gamma^* \gamma^*$ vertex. Two methods will be used in this chapter to obtain measurements of the DIP form factor parameters α and β . The *shape* method takes advantage of the fact that these parameters connect to the momentum of the virtual photons, and are therefore related to the invariant mass distributions of the electron and muon pairs in the final state. Monte carlo M_{ee} and $M_{\mu\mu}$ distributions, generated with varying values of α and β , can be compared to the corresponding data distributions in order to find the best fit. The *branching ratio* method makes use of the fact that the predicted rate for the decay $K_L \longrightarrow e^+e^-\mu^+\mu^-$ depends strongly on α . Therefore, an independent measurement of α can be obtained using the branching ratio result from the previous chapter. Finally, the angular distribution of the signal decay products will be examined for any asymmetry, which would indicate the presence of a CP -violating contribution to the $K_L \gamma^* \gamma^*$ vertex.

7.1 Expected Sensitivity

As seen in the previous chapter, a total of 132 $K_L \longrightarrow e^+e^-\mu^+\mu^-$ events have been identified in the entire KTeV dataset. That's not a lot of data to work with, especially with the hope of measuring a quadratic form factor term. However, even with the limited statistics, there is evidence for a non-trivial form factor. Figure 7.1 shows the M_{ee} and $M_{\mu\mu}$ distributions for the accepted events, overlaid on monte




Figure 7.1: M_{ee} and $M_{\mu\mu}$ distributions for the 132 accepted $K_L \rightarrow e^+e^-\mu^+\mu^-$ signal events (dots) overlaid on MC (histogram) generated with $\alpha = \beta = 0$. The disagreement seen in the $M_{\mu\mu}$ overlay suggests the existence of a non-trivial form factor.

carlo generated with a flat form factor (Equation 1.32 with $\alpha = \beta = 0$). The almost 3σ disagreement seen in the $M_{\mu\mu}$ overlay suggests the existence of a non-trivial form factor, with $\alpha \neq 0$.

Unfortunately, it appears that this analysis will be insensitive to a β of

the theoretically expected size. Figure 7.2 shows the magnitude of the DIP form factor (Equation 1.32) as a function of M_{ee} and $M_{\mu\mu}$ for $\alpha = -1.54$ and $\beta = +2$ (recall from Chapter 1 that the best experimental measurement, from analysis of $K_L \rightarrow \mu^+ \mu^- \gamma$, is $\alpha = -1.54 \pm 0.10$; and that two theoretical calculations, discussed in Section 1.3.3, place β at a value around +2). Also shown are projections onto both invariant mass axes, as well as the contributions to the total from each term of the form factor. Note that the β term contributes only a very small amount to the total form factor magnitude. This becomes even more of a concern when it is noted that the form factor is *squared* in the expression for the differential $K_L \rightarrow e^+ e^- \mu^+ \mu^-$ rate (Equation 1.53).

The small β dependence is further illustrated in Figure 7.3, which overlays monte carlo M_{ee} and $M_{\mu\mu}$ distributions at generation (before any geometric detector effects, trigger requirements, or analysis cuts) for β values of +2 and +10. It can be seen that the sensitivity to β is confined to the high end of the M_{ee} spectrum, where the statistics are the lowest. Compare this to the shape dependence of $M_{\mu\mu}$ on α , shown in Figure 7.4.

However, while the expected insensitivity to β is somewhat discouraging, it is advantageous in the sense that it allows for fitting the data for α and β independently, instead of dealing with the complexities of a simultaneous two-dimensional fit. Examination of Equation 1.32 and Figures 7.3 & 7.4 reveals that changes in α mostly affect the shape of the $M_{\mu\mu}$ spectrum, while M_{ee} is almost exclusively sensitive to β .

7.2 Measuring α Using a Shape Analysis

A reweighting technique was used to fit for the most likely value of α in the shape analysis. Signal mode monte carlo was generated for known values of α and β ,

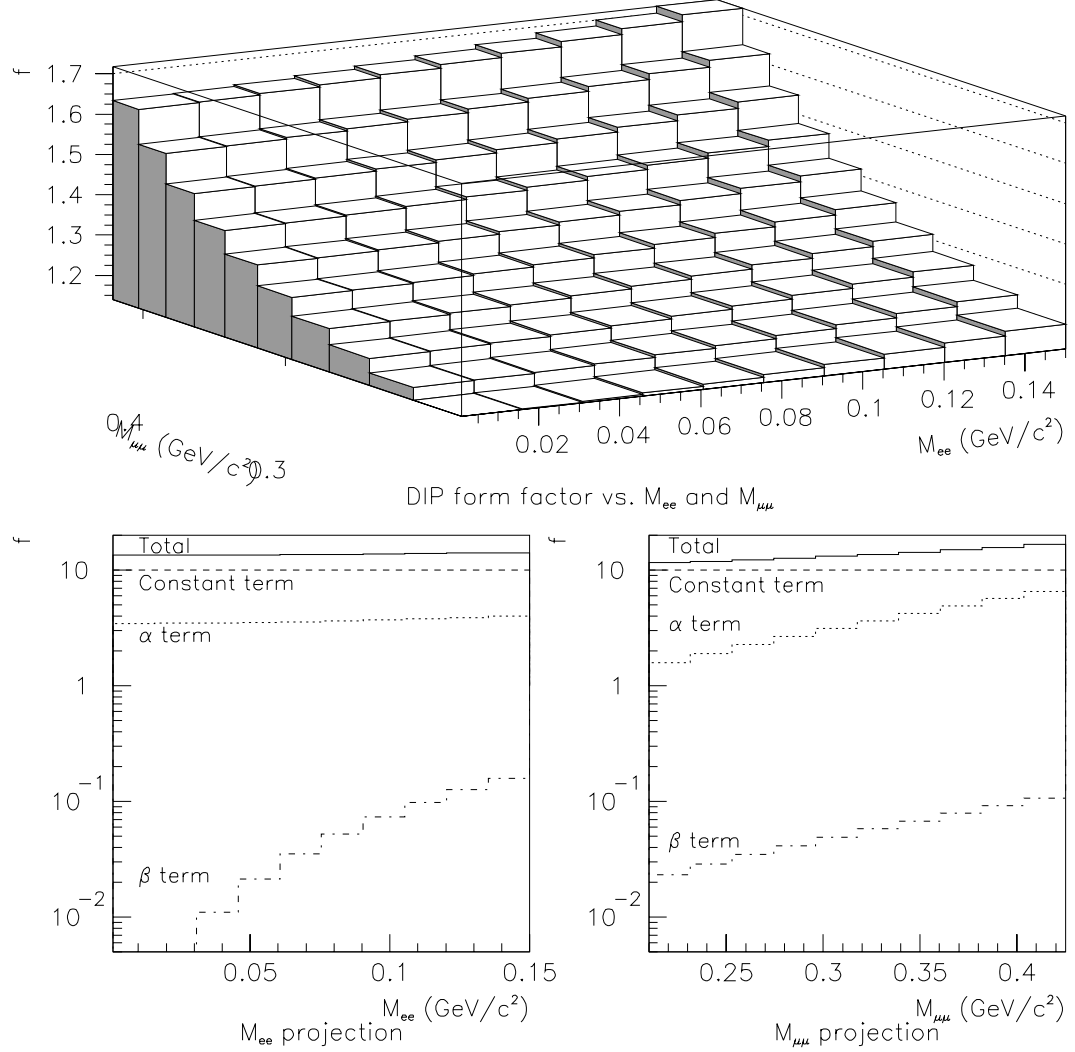


Figure 7.2: Top: Magnitude of the DIP form factor (Equation 1.32) as a function of M_{ee} and $M_{\mu\mu}$ for $\alpha = -1.54$, $\beta = +2$. Bottom: Projections of the two-dimensional distribution onto each mass axis, as well as the relative contributions to the total from each term. Note that the total is almost completely insensitive to the β term – a fact even further magnified by the fact that the form factor is squared in the calculation of the $K_L \rightarrow e^+e^-\mu^+\mu^-$ rate (Equation 1.53).




Figure 7.3: Comparison between monte carlo M_{ee} and $M_{\mu\mu}$ distributions at generation for $\beta = +2$ (dots) and $\beta = +10$ (histogram). $\alpha = -1.54$ for both MC sets. Note that the minimal amount of sensitivity to β in these plots arises only in the high end of the M_{ee} spectrum, where the statistics are the lowest. There is no significant $M_{\mu\mu}$ sensitivity to β .




Figure 7.4: Comparison between monte carlo M_{ee} and $M_{\mu\mu}$ distributions at generation for $\alpha = -1.54$ (dots) and $\alpha = -5$ (histogram). $\beta = 0$ for both MC sets. Note the sensitivity of the $M_{\mu\mu}$ spectrum to α . The shape of the M_{ee} distribution is essentially unaffected by changes in α .

labeled α_0 and β_0 ¹. MC events that passed all analysis cuts after reconstruction were reweighted for many different values of α_i , while β_i was fixed at β_0 . The weights were calculated using the ratio

$$W(\alpha_i, \beta_i) = \frac{f^2(\alpha_i, \beta_0)}{f^2(\alpha_0, \beta_0)}, \quad (7.1)$$

where f is the expression for the DIP form factor (Equation 1.32). These reweighted events were then used to fill histograms in M_{ee} and $M_{\mu\mu}$ for each value of α_i , which were compared to the corresponding distributions in the data. The most likely value of α_i (dubbed α_{shape}) was determined by maximizing the log-likelihood function

$$\mathcal{L}(\alpha_i) = \sum_{j=1}^{nbins} r_j \ln(\mu_j(\alpha_i)) - \mu_j(\alpha_i) - \ln(r_j!), \quad (7.2)$$

where r_j and $\mu_j(\alpha_i)$ are the numbers of data and monte carlo events in each histogram bin, respectively. The summation was carried out over the M_{ee} and $M_{\mu\mu}$ distributions (each partitioned into 10 bins) for every value of α_i . The fit, shown in Figure 7.5, provides a value of

$$\alpha_{shape} = -4.53^{+1.81}_{-2.70}, \quad (7.3)$$

the statistical error being found by measuring the width of the fitted parabola 0.5 units below the maximum.

In the $K_L \rightarrow \mu^+ \mu^- \gamma$ shape analysis, significant shifts in the fitted value of α_{shape} were observed as the analysis cut on track momentum was varied (Quinn 2000). To test if this analysis is similarly affected, the cut on muon momentum was moved from its nominal value of 7 to 10 GeV/c. This led to a shift in α_{shape} of $-13.4 \pm 14.8\%$. Based on data/MC disagreements seen in the branching ratio analysis, other cut variations were also tried. Moving the cut on vertex χ^2 from the nominal value of 40 to 25 moved α_{shape} by $+0.3 \pm 7.5\%$. Tightening the cut on P_t^2 from 250 to 200 MeV/c² changed α_{shape} by $-1.1 \pm 4.7\%$. Moving the downstream

¹ α_0 and β_0 were chosen to be -4.60 and 0, respectively. The choice of initial values ends up having no effect on the final fit results.




Figure 7.5: (a) Log-likelihood as a function of α_i in the $K_L \rightarrow e^+e^-\mu^+\mu^-$ shape analysis. The parameter β is fixed at 0. (b) Close-up of the maximum of the log-likelihood distribution, which corresponds to a fit value of $\alpha_{\text{shape}} = -4.53^{+1.81}_{-2.70}$.

Number of bins	α_{shape} from fit
7	$-4.87^{+2.02}_{-2.90}$
8	$-4.47^{+1.89}_{-2.54}$
9	$-4.55^{+2.18}_{-3.19}$
10 (default)	$-4.53^{+1.81}_{-2.70}$
11	$-4.38^{+1.77}_{-2.63}$
12	$-4.54^{+1.87}_{-2.67}$
13	$-4.11^{+1.71}_{-2.39}$

Table 7.1: Fitted value of α_{shape} versus the number of bins used in the M_{ee} and $M_{\mu\mu}$ plots. The average of these entries, -4.50 , is almost identical to the default value of -4.53 .

cut on vertex z from 158 to 150 m caused α_{shape} to increase by $8.9 \pm 12.0\%$. The time dependence of α_{shape} was studied by separating the total fit result into independent 1997 and 1999 measurements. The values found for these two time periods were $-5.59^{+3.22}_{-7.54}$ and $-4.12^{+2.04}_{-3.19}$, respectively. There appears to be no significant systematic difference between the 1997 and 1999 results.

It was seen during the $K_L \rightarrow e^+e^-\mu^+\mu^-$ branching ratio analysis that the default monte carlo simulation of muon scattering underestimates the scattering width at higher values of muon momentum (see Figure 6.31). This was seen to have a small ($\approx 0.6\%$) effect on the total signal acceptance. However, because the effect is momentum-dependent, it might be expected to distort the shape of the $M_{\mu\mu}$ spectrum, and therefore bias the measurement of α in the shape analysis. Fitting a special set of monte carlo generated with the additional muon scattering led to a shift in α_{shape} of $-5.3 \pm 10.4\%$.

Finally, the value of α_{shape} seems to depend somewhat on the number of bins used in the log-likelihood calculation (see Table 7.1). However, the average of these entries, -4.50 , is almost identical to the default value of -4.53 .

Because none of the systematic effects are clearly statistically significant, no systematic error is quoted on the final value for α_{shape} .

7.3 Measuring β Using a Shape Analysis

The shape analysis for β was carried out in much the same way as that for α . Events passing all analysis cuts were reweighted for many different values of β_i , while α_i was fixed at α_0 (now chosen to be -4.53). The weights are now given by the expression

$$W(\alpha_i, \beta_i) = \frac{f^2(\alpha_0, \beta_i)}{f^2(\alpha_0, \beta_0)}, \quad (7.4)$$

while the log-likelihood function becomes

$$\mathcal{L}(\beta_i) = \sum_{j=1}^{nbins} r_j \ln(\mu_j(\beta_i)) - \mu_j(\beta_i) - \ln(r_j!). \quad (7.5)$$

The fit results are shown in Figure 7.6. The log-likelihood function is maximized at a value of:

$$\beta_{shape} = +12.3 \pm 77.1. \quad (7.6)$$

This value is consistent with the theoretical prediction of $\beta \approx +2$. However, the lack of sensitivity precludes any stronger claim based on this limited dataset. In order to measure β with a statistical error on the order of 1 unit, a $K_L \rightarrow e^+e^-\mu^+\mu^-$ sample of about 750,000 events would be required – a dataset that unfortunately won't be seen for quite some time, if ever.

7.4 Measuring α Using the $K_L \rightarrow e^+e^-\mu^+\mu^-$ Branching Ratio

The differential decay rate of $K_L \rightarrow e^+e^-\mu^+\mu^-$ was given in Equation 1.53. By inserting the DIP form factor (Equation 1.32) into Equation 1.53 and integrating over the kinematically allowed phase space of q_e^2 and q_μ^2 , the branching ratio (normalized to $\mathcal{B}(K_L \rightarrow \gamma\gamma)$) can be directly related to the DIP form factor parameters. Conveniently, this integration was done automatically by the DIP monte carlo, which also accounted for the effects from radiative corrections.




Figure 7.6: Log-likelihood as a function of β_i in the $K_L \rightarrow e^+e^-\mu^+\mu^-$ shape analysis. The parameter α is fixed at -4.53 . The maximum of the log-likelihood distribution corresponds to a fit value of $\beta_{\text{shape}} = +12.3 \pm 77.1$.




Figure 7.7: Theoretical prediction of the ratio $\mathcal{B}(K_L \rightarrow e^+e^-\mu^+\mu^-)/\mathcal{B}(K_L \rightarrow \gamma\gamma)$ as a function of the DIP parameter α . Also shown is the extrapolated value of α_{BR} , based on the measured $K_L \rightarrow e^+e^-\mu^+\mu^-$ branching ratio of $(2.63 \pm 0.23_{stat} \pm 0.13_{syst}) \times 10^{-9}$, and assuming $\beta = 0$ (dots). The error bars on the experimental result (dotted lines) include the uncertainty on $\mathcal{B}(K_L \rightarrow \gamma\gamma)$. This method provides an estimate for α_{BR} of -1.43 ± 0.39 , consistent with earlier measurements. The triangles show the effect of setting β to $+2$.

The predicted dependence of the $K_L \rightarrow e^+e^-\mu^+\mu^-$ branching ratio on α is shown in Figure 7.7. Using the final measurement of $(2.63 \pm 0.23_{stat} \pm 0.13_{syst}) \times 10^{-9}$ ($x_{e^+e^-\mu^+\mu^-} > 0.95$), and assuming $\beta = 0$, this method provides a value of

$$\alpha_{BR} = -1.43 \pm 0.39. \quad (7.7)$$

It should be noted from Figure 7.7 that the shape of the normalized branching ratio vs. α curve does depend on β , but only very weakly. Moving β to a value of +2 changes α_{BR} by less than a percent.

Using a weighted average of the shape and branching ratio measurements of α , a final result from the $K_L \longrightarrow e^+e^-\mu^+\mu^-$ mode can be calculated:

$$\alpha_{e^+e^-\mu^+\mu^-} = -1.52 \pm 0.38.$$

This result is consistent with the earlier measurements from $K_L \longrightarrow \mu^+\mu^-\gamma$ and $K_L \longrightarrow e^+e^-e^+e^-$.

7.5 Search for CP -violation in the Decay $K_L \longrightarrow e^+e^-\mu^+\mu^-$

As mentioned in Chapter 1, a CP -violating contribution to the $K_L \gamma^* \gamma^*$ interaction would result in an asymmetry in the angular distribution of the final state decay products. Figure 7.8 shows the distribution in $\sin \phi \cos \phi$ for the 132 signal events and monte carlo with no CP -violation (the angle ϕ is pictured in Figure 1.7).

The asymmetry \mathcal{A} is defined as

$$\mathcal{A} = \frac{N_+ - N_-}{N_+ + N_-}. \quad (7.8)$$

where N_+ (N_-) is the acceptance corrected number of events in the positive (negative) range of $\sin \phi \cos \phi$. With 63 data events in the positive range and 69 in the negative range, an asymmetry of $\mathcal{A} = -5.3 \pm 12.3\%$ was calculated, consistent with 0. (The value of the asymmetry for the pure CP -conserving monte carlo is $0.4 \pm 0.7\%$).

As with the DIP parameter fits, the time dependence of this asymmetry was studied by separating the total fit result into independent 1997 and 1999 measurements. The values found for these two time periods were -8.5 ± 23.1 and -3.2 ± 13.6 , respectively. The χ^2 of 0.04/1 *dof* corresponds to a probability of




Figure 7.8: Distribution of 132 signal data events (dots) and monte carlo with no CP -violation (histogram) in $\sin\phi\cos\phi$. The measured asymmetry in the data is $\mathcal{A} = -5.3 \pm 12.3\%$, providing a 90% C.L. limit of $|\mathcal{A}| < 25.5\%$

agreement of over 84% (not surprising, since the statistical error bars are so large). Nor were any significant systematic effects observed when analysis cuts (on muon momentum, vertex χ^2 , P_t^2 , and vertex z) or the DIP parameters for the acceptance MC were varied.

Using the Feldman–Cousins methodology (Feldman and Cousins 1998) to eliminate the unphysical region of $|\mathcal{A}| < 0$, a limit of

$$|\mathcal{A}| < 25.5\%, \quad 90\% \text{ C.L.} \quad (7.9)$$

was obtained. Based on this measurement of the asymmetry, it can be concluded that no evidence currently exists for a significant CP -violating contribution to the $K_L \gamma^* \gamma^*$ vertex.

CHAPTER 8

Conclusions

The study of the rare decay mode $K_L \longrightarrow e^+e^-\mu^+\mu^-$ has come a long way since the discovery of a single event by the E799-I collaboration. The 1997 and 1999 runs of the KTeV experiment have increased the world sample by over two orders of magnitude, producing a total sample of 132 signal events, with an expected background of 0.8 events.

Analysis of this signal has produced a myriad of interesting results. The branching ratio was calculated to be

$$\mathcal{B}(K_L \longrightarrow e^+e^-\mu^+\mu^- ; x_{e^+e^-\mu^+\mu^-} > 0.95) = (2.63 \pm 0.23_{\text{stat}} \pm 0.13_{\text{syst}}) \times 10^{-9}, \quad (8.1)$$

which is just over 1σ away from the VMD prediction of $(2.34 \pm 0.06) \times 10^{-9}$ by Quigg and Jackson. If the form factor of D'Ambrosio, Isidori, and Portolés is assumed, then this branching ratio result is most consistent with a value for the linear parameter α_{BR} of -1.43 ± 0.39 , if the quadratic parameter β is near 0 (see Figure 8.1).

A fit to the shape of the $M_{\mu\mu}$ and M_{ee} distributions of the signal data yielded an independent measurement for α_{shape} of $-4.53^{+1.81}_{-2.70}$. Combining the branching ratio and shape results using a weighted average, a final estimate from $K_L \longrightarrow e^+e^-\mu^+\mu^-$ of

$$\alpha_{e^+e^-\mu^+\mu^-} = -1.52 \pm 0.38 \quad (8.2)$$

was obtained, which is consistent with the previous measurements from $K_L \longrightarrow \mu^+\mu^-\gamma$ (-1.54 ± 0.10) and $K_L \longrightarrow e^+e^-e^+e^-$ (-1.1 ± 0.6). These three numbers




Figure 8.1: Summary plot of theoretical calculations for the $K_L \rightarrow e^+e^-\mu^+\mu^-$ branching ratio. This result is consistent with the prediction from Vector Meson Dominance. If the DIP form factor is assumed, the measured branching ratio is most consistent with the value $\alpha = -1.43 \pm 0.39$, if β is near 0.




Figure 8.2: Summary plot of all measurements made of the DIP parameter α to date. The world average comes to $\bar{\alpha} = -1.53 \pm 0.10$, almost identical to the total result from the $K_L \rightarrow \mu^+ \mu^- \gamma$ analysis.

lead to a world average of

$$\bar{\alpha} = -1.53 \pm 0.10 \quad (8.3)$$

(see Figure 8.2).

Little sensitivity was found to the quadratic form factor parameter β , measured to be

$$\beta_{e^+e^-\mu^+\mu^-} = +12.3 \pm 77.1, \quad (8.4)$$

meaning that knowledge of the DIP parameters has not improved since the analysis of $K_L \longrightarrow \mu^+\mu^-\gamma$, when the value $\alpha = -1.54 \pm 0.10$ was used to set the CKM constraint (Quinn 2000)

$$\rho > -0.2. \quad (8.5)$$

This constraint approaches the combined limit of $\rho > 0$ from $|V_{ub}|$, B mixing, and ϵ (Quinn 2000), (Ali and London 1999).

Measurement of the asymmetry in the angular distribution of the $K_L \longrightarrow e^+e^-\mu^+\mu^-$ decay products led to a limit of

$$|\mathcal{A}| < 25.5\%, \quad 90\% \text{ C.L.} \quad (8.6)$$

No evidence exists to support the presence of a significant CP -violating contribution to the $K_L \gamma^* \gamma^*$ vertex.

Finally, the absence of lepton flavor-violating $K_L \longrightarrow e^\pm e^\pm \mu^\mp \mu^\mp$ events in the KTeV dataset allowed a limit to be placed on the branching ratio of

$$\mathcal{B}(K_L \longrightarrow e^\pm e^\pm \mu^\mp \mu^\mp) < 4.12 \times 10^{-11}, 90\% \text{ C.L.}, \quad (8.7)$$

a full factor of 3 lower than the previously published limit (Alavi-Harati et al. 2001c).

No real prospects exist for significantly expanding the world sample of $K_L \longrightarrow e^+e^-\mu^+\mu^-$ events in the near future. New results for the DIP parameter α are expected soon from analysis of the 1999 $K_L \longrightarrow e^+e^-e^+e^-$ and $K_L \longrightarrow e^+e^-\gamma$ KTeV data. However, the outlook for obtaining a significant measurement of β is bleak.

REFERENCES

- Alavi-Harati, A. (1999, April). Ph. D. thesis, University of Wisconsin-Madison.
- Alavi-Harati, A. et al. (2000). *Phys. Rev. Lett.* **84**, 408.
- Alavi-Harati, A. et al. (2001a). *Phys. Rev. Lett.* **87**, 71801.
- Alavi-Harati, A. et al. (2001b). *Phys. Rev. Lett.* **86**, 5425.
- Alavi-Harati, A. et al. (2001c). *Phys. Rev. Lett.* **87**, 111802.
- Alexopolous, T., A. Erwin, and R. S. Kessler (1994, June). KTeV Internal Note #207.
- Ali, A. and D. London (1999). hep-ph/0002167.
- Ambrose, D. et al. (2000). *Phys. Rev. Lett.* **84**, 1389.
- Ametller, L., A. Bramon, and E. Masso (1984). *Phys. Rev.* **D30**, 251.
- Appel, R. et al. (2000). *Phys. Rev. Lett.* **85**, 2450.
- Barker, A. Personal communication.
- Barrio, M. (2001, December). Ph. D. thesis, The University of Chicago.
- Belusevic, R. (1999). *Springer Tracts in Modern Physics* **153**.
- Bethe, H. A. (1953). *Phys. Rev.* **89**, 1256.
- Bocean, V. et al. (1998). Technical Report TM-2046, Fermilab.
- Bown, C. et al. (1996). *Nucl. Instrum. Meth.* **A369**, 248.
- Bown, C., V. Prasad, and S. Taegar (1997, November). KTeV Internal Note #487.
- Buchalla, G. and A. J. Buras (1994). *Nucl. Phys.* **B412**, 106.
- Buras, A. J. (2001). hep-ph/0109197.
- Buras, A. J. and R. Fleischer (1998). *Adv. Ser. Direct. High Energy Phys.* **15**, 65.
- Cheu, E., J. Hamm, and J. Wang (1999, February). KTeV Internal Note #606.
- Christenson, J. H., J. W. Cronin, V. L. Fitch, and R. Turlay (1964). *Phys. Rev. Lett.* **13**, 138.

- Corcoran, M. Personal communication.
- D'Ambrosio, G., G. Isidori, and J. Portoles (1998). *Phys. Lett.* **B423**, 385.
- D'Ambrosio, G. and J. Portoles (1997). *Nucl. Phys.* **B492**, 417.
- Ecker, G., A. Pich, and E. de Rafael (1988). *Nucl. Phys.* **B303**, 665.
- Feldman, G. J. and R. D. Cousins (1998). *Phys. Rev.* **D57**, 3873.
- Gell-Mann, M. and A. Pais (1955). *Phys. Rev.* **97**, 1387.
- Griffiths, D. (1987). *Introduction to Elementary Particles*. John Wiley & Sons.
- Groom, D. E. et al. (2000a). *Eur. Phys. J.* **C15**, 110.
- Groom, D. E. et al. (2000b). *Eur. Phys. J.* **C15**, 395.
- Groom, D. E. et al. (2000c). *Eur. Phys. J.* **C15**, 163.
- Gu, P. et al. (1996). *Phys. Rev. Lett.* **76**, 4312.
- Halkiadakis, E. (2001, May). Ph. D. thesis, Rutgers, The State University of New Jersey.
- Hamm, J. (1999, September). KTeV Internal Note #673.
- Hanagaki, K. (1997, March). KTeV Internal Note #408.
- Hanagaki, K. (1998, September). Ph. D. thesis, Osaka University.
- Holstein, B. R. (1995). hep-ph/9510344.
- Kabassian, G. et al. (1994, April). KTeV Internal Note #196.
- Kessler, R. S. (2001, January). Unpublished KTeV Internal Note.
- Kobayashi, M. and T. Maskawa (1973). *Prog. Theor. Phys.* **49**, 652.
- Kroll, N. M. and W. Wada (1955). *Phys. Rev.* **98**, 1355.
- LaDue, J. Personal communication.
- LaDue, J. and P. Toale. Personal communication.
- Landau, L. (1944). *J. Phys. (USSR)* **8**, 201.
- Malensek, A. J. (1982). Fermilab Report FN-341, Fermilab.
- Mikaelian, K. O. and J. Smith (1972). *Phys. Rev.* **D5**, 1763.
- Mikelsons, P. (1999, December). Ph. D. thesis, University of Colorado.
- Miyazaki, T. and E. Takasugi (1973). *Phys. Rev.* **D8**, 2051.

- Nelson, W. R. et al. (1985). SLAC Report SLAC-265, SLAC-R-265, SLAC.
- Nishimura, H. et al. (1995). *Phys. Rev.* **B51**, 2167.
- Patrick, G. N. (1986). GEANT User's Manual PHYS332, CERN.
- Prasad, V. (2002, June). Ph. D. thesis, The University of Chicago.
- Quigg, C. and J. D. Jackson (1968). UCRL-18487 (unpublished).
- Quinn, G. B. (2000, June). Ph. D. thesis, The University of Chicago.
- Roodman, A. (1997, November). Tucson, Arizona. VII International Conference on Calorimetry.
- Roodman, A. (1998, July). KTeV Internal Note #577.
- Sehgal, L. M. (1969). *Phys. Rev.* **183**, 1511.
- Sehgal, L. M. and M. Wanninger (1992). *Phys. Rev.* **D46**, 1035.
- Senyo, K. (1999, December). Ph. D. thesis, Osaka University.
- Shawhan, P. (2000, June). Ph. D. thesis, The University of Chicago.
- Solomey, N. (1998, February). Vienna, Austria. 1998 Vienna Wire Chamber Conference.
- Toale, P., H. Huang, and A. R. Barker (2002, June). Unpublished KTeV Internal Note.
- Uy, Z. E. S. (1991). *Phys. Rev.* **D43**, 802.
- Uy, Z. E. S. (2002). *Eur. Phys. J.* **C23**, 113.
- Vavilov, P. V. (1957). *Sov. Phys. JTEP* **5**, 749.
- Wolfenstein, L. (1983). *Phys. Rev. Lett.* **51**, 1945.
- Zhang, L. and J. L. Goity (1998). *Phys. Rev.* **D57**, 7031.
- Zimmerman, E. D. (1999). *Nucl. Instrum. Meth.* **A426**, 229.