

Search for the Decay KL to $\pi^0 e^+ e^-$ and Study
of the Decay KL to $e^+ e^- \gamma \gamma$

by

Peter L. Mikelsons

B.A., Macalester College, 1991

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics

1999

This thesis entitled:
Search for the Decay KL to $\pi^0 e^+ e^-$ and Study of the Decay KL to $e^+ e^- \gamma$
gamma
written by Peter L. Mikelsons
has been approved for the Department of Physics

Anthony Barker

Uriel Nauenberg

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Mikelsons, Peter L. (Ph.D., Physics)

Search for the Decay $K_L \rightarrow \pi^0 e^+ e^-$ and Study of the Decay $K_L \rightarrow e^+ e^- \gamma \gamma$

Thesis directed by Prof. Anthony Barker

The particle decay $K_L \rightarrow \pi^0 e^+ e^-$ is a probe of direct CP violation, a phenomenon previously only seen in $K_L \rightarrow \pi\pi$ decays. Understanding direct CP violation is an important part of understanding violation of CP symmetry in general. Experimentally, one of the obstacles to studying $K_L \rightarrow \pi^0 e^+ e^-$ is the rare decay $K_L \rightarrow e^+ e^- \gamma \gamma$, which can mimic $K_L \rightarrow \pi^0 e^+ e^-$. A study of $K_L \rightarrow \pi^0 e^+ e^-$ and $K_L \rightarrow e^+ e^- \gamma \gamma$ was made as part of the KTeV E799 experiment. $K \rightarrow \pi^0 \pi^0_{Dalitz}$ decays were used for normalization, and a K_L flux of $(2.65 \pm 0.18) \times 10^{11}$ decays was measured. We observed 1578 $K_L \rightarrow e^+ e^- \gamma \gamma$ candidate events, of which 1516.5 ± 1.8 remain after background subtraction. These events allow measurement of the Bergström, Massó, and Singer $K_L \gamma \gamma$ vertex form-factor parameter, $\alpha_{K^*} = +0.015 \pm 0.12_{stat.} \pm 0.03_{sys.}$, in mild disagreement with the previously fit value of -0.28 ± 0.08 . This form-factor implies a corresponding branching ratio of $\Gamma(K_L \rightarrow e^+ e^- \gamma \gamma, E_\gamma^* > 5 \text{ MeV}) / \Gamma(K_L \rightarrow \text{all}) = (5.82 \pm 0.15_{stat.} \pm 0.31_{sys.} \pm 0.19_{BR}) \times 10^{-7}$, in agreement with the QED prediction.

The search for $K_L \rightarrow \pi^0 e^+ e^-$ found two candidate events. However, 1.06 ± 0.41 events were expected from background processes. Therefore, we do not claim observation of $K_L \rightarrow \pi^0 e^+ e^-$. Instead, with a single-event sensitivity of 1.00×10^{-10} , we set an upper limit on the $K_L \rightarrow \pi^0 e^+ e^-$ branching ratio of 4.86×10^{-10} at the 90% confidence level.

Acknowledgements

For guidance, inspiration, teaching, and motivation, my advisor Anthony Barker.

For their intelligence and industry applied to $K_L \rightarrow \pi^0 e^+ e^-$, the other members of the “pi e e” working group, Leo Bellantoni and Greg Graham.

For the hard work and good times, the other University of Colorado at Boulder workers and learners, including Bruce Broomer, Eric Erdos, David Fillmore, Polly Fordyce, Charles Griffis, Jim Hirschauer, Sumita Jaggar, Jason “Elmarth” LaDue, Uriel Nauenberg, Jose-Maria Ramas, Gerhard Schultz, Pat Toale, and Jin-Yuan Wu.

For the amazing feat of making it all happen, the other once and current members of the KTeV collaboration, including (in random order) Hans Kobrak, Ram Ben-David, Roland Winston, Taku Yamanaka, Marj Corcoran, Valeri Jejer, Amitabh Lath, Rick Tesarek, Sam Childress, Steve Schnetzer, Bruce Winstein, Breese Quinn, Myungyun Pang, Suzanne Averitte, William Slater, Ashkan Alavi-Harati, Angela Bellavance, Yau Wai Wah, Ruth Pordes, Tsuyoshi Nakaya, Doug Bergman, Masashi Hazumi, Hogan Nguyen, Rick Coleman, Satoshi Hidaka, Kazunori Hanagaki, Paul Johnson, John Belz, Changqing Qiao, Albert Erwin, John Jennings, Ron Ray, Peter Shawhan, Earl Swallow, Ed Blucher, Bob Swanson, John Adams, Erik Ramberg, Gordon Thomson, Karla Hagan-Ingram, Yee Bob Hsiung, Mike Arenton, Robert Tschirhart, Sunil Somalwar, Ichiro Suzuki, Arun Tripathi, Jim Graham, Katsumi Senyo, Rick Ford, Arthur McManus, Peter Shanahan, Julie Whitmore, Nickolas Solomey, Sasha Ledovskoy, Greg Bock, Doug Jensen, Stephen Field, Colin Bown, Elliott Cheu, Vivian O'Dell, Eva

Halkiadakis, Brad Cox, Valmiki Prasad, Aaron Roodman, Eric Zimmerman, Katsushi Arisaka, Ivone Albuquerque, Theo Alexopoulos, Emmanuel Monnier, Mike Crisler, Steve Bright, Rick Kessler, Sidney Taegar, Masayoshi Sadamoto, and Herman White.

For supporting me and never doubting my ability to do it, my mother (special thanks for proofreading), father, and brother.

Contents

Chapter

1	Introduction	1
1.1	CP Symmetry	1
1.1.1	Neutral Kaon System	2
1.1.2	Rare Decays	3
1.2	$K_L \rightarrow \pi^0 e^+ e^-$	5
1.2.1	Standard Model	5
1.2.2	Outside Standard Model	8
1.2.3	Previous Searches	8
1.2.4	Backgrounds	9
1.3	$K_L \rightarrow e^+ e^- \gamma \gamma$	11
1.3.1	Previous Measurements	12
1.3.2	Standard Model	12
1.3.3	Backgrounds	15
1.4	Method Overview	16
2	KTeV Experimental Setup	18
2.1	Kaon Beams	18
2.1.1	Proton Beam	18
2.1.2	Target	19

2.1.3	Sweeping and Collimating	21
2.1.4	Beam Description	23
2.2	Detector	24
2.2.1	Spectrometer	24
2.2.2	Calorimeter	30
2.2.3	Transition Radiation Detectors	33
2.2.4	Vetoos	36
2.2.5	Trigger Detectors	40
3	Event Selection and Data Collection	42
3.1	Level 0 Trigger: Digitization	42
3.2	Level 1 Trigger: Prompt Logic	45
3.2.1	Processing L1 Sources	45
3.2.2	Logic	46
3.2.3	Calibration Triggers	49
3.2.4	L1 Output	49
3.3	Level 2 Trigger: Hardware Logic	50
3.3.1	Calorimeter Cluster Counting	50
3.3.2	DC Hit Counter	52
3.3.3	Track Finding	53
3.3.4	“Stiff” Tracks	53
3.3.5	Control Logic	54
3.4	Level 2 Trigger: The Hit Counting System	56
3.4.1	DC Signals	57
3.4.2	FASTBUS Racks	58
3.4.3	Auxiliary Cards: Blossoms	61
3.4.4	Latches: Kumquats	61

3.4.5	Trigger TDC: Bananas	64
3.4.6	Hit Summing	73
3.4.7	Outputs to Level 2	75
3.4.8	FASTBUS interface	76
3.4.9	FERA Interface	77
3.5	Readout	79
3.5.1	Crates	79
3.5.2	Streams	80
3.6	Level 3 Trigger: Software	80
4	Taking Data	82
4.1	Runs	82
4.2	Run History	83
4.3	Data Yield	85
5	Reconstruction	86
5.1	Calorimeter Clusters	86
5.1.1	Channel Energy	86
5.1.2	Clusters	87
5.2	Spectrometer Tracks	91
5.2.1	Track Candidates	91
5.2.2	Vertex	95
5.2.3	Momenta	98
5.2.4	Extra Tracks	100
5.3	Kinematic Quantities	100
5.4	TRD Particle ID	101

6	Simulation	104
6.1	Kaons	104
6.1.1	Kaon Tracing	105
6.1.2	Kaon Decay	110
6.2	Decay Products	112
6.2.1	Tracing	112
6.2.2	Material Interactions	114
6.2.3	Digitization	120
6.3	Accidental Overlay	124
6.3.1	Simulating Run Number	125
6.3.2	Overlaying	126
6.4	Trigger	126
7	Reduction	128
7.1	Split	128
7.2	Crunch	128
7.2.1	Reconstruction	129
7.2.2	Cuts	129
7.2.3	Output	131
8	Normalization	132
8.1	Why $K \rightarrow \pi^0 \pi^0_{Dalitz}$?	132
8.2	Simulated Events	133
8.3	General Cuts	135
8.3.1	“Bad Spills”	136
8.3.2	Reconstruction	136
8.3.3	Fiducial Region	139
8.3.4	Trigger Verification	142

8.3.5	Calorimeter	143
8.3.6	TRD	145
8.3.7	E/p	150
8.3.8	Spectrometer	151
8.3.9	Summary	154
8.4	Mode-Specific Cuts	155
8.4.1	$\gamma\gamma$ Mass	155
8.4.2	$e^+e^-\gamma$ Mass	156
8.4.3	Transverse Momentum	158
8.4.4	$e^+e^-\gamma\gamma\gamma$ Mass	160
8.4.5	Summary	163
8.5	Flux	164
8.6	Simulation Quality	166
9	$K_L \rightarrow e^+e^-\gamma\gamma$	170
9.1	Simulated Events	170
9.2	General Cuts	172
9.3	Mode-Specific Cuts	172
9.3.1	Photon Energy Threshold	172
9.3.2	Transverse Momentum	177
9.3.3	$M_{ee\gamma\gamma}$	177
9.3.4	θ_{min}	181
9.3.5	Summary	186
9.4	Other Backgrounds	186
9.4.1	Charged Pions	187
9.4.2	$K \rightarrow \pi^0\pi^0\pi^0_{Dalitz}$	187
9.5	Branching Ratios	188

9.5.1	$E_\gamma^* > 5 \text{ MeV}$	189
9.5.2	Other Thresholds	190
9.6	Form Factor	193
9.6.1	Measuring Form Factor	193
9.6.2	Statistical Error	196
9.6.3	Systematic Errors	198
9.6.4	Effect on Branching Ratio	199
9.7	Simulation Quality	200
9.8	Conclusions	201
10	$K_L \rightarrow \pi^0 e^+ e^-$	205
10.1	Simulated Sample	205
10.2	General Cuts	205
10.3	Obvious, Mode-Specific Cuts	207
10.3.1	$M_{\pi^+ \pi^- \gamma \gamma}$	207
10.3.2	TRD	208
10.3.3	M_{ee}	208
10.3.4	Transverse Momentum	209
10.3.5	$M_{ee \gamma \gamma}$	209
10.4	Non-Obvious, Mode-Specific Cuts	212
10.4.1	$M_{\gamma \gamma}, M_{ee \gamma \gamma}$ Plane	212
10.4.2	$K_L \rightarrow e^+ e^- \gamma \gamma$ Background Checks	215
10.4.3	Optimizing Kinematic Cuts	219
10.4.4	$M_{\gamma \gamma}$	221
10.5	Cut Summary and Final Background Estimate	223
10.6	Alternative Estimates of Backgrounds	226
10.6.1	$K_L \rightarrow \pi^0 \gamma e^+ e^-$	226

10.6.2	$K \rightarrow \pi^0 \pi_{Dalitz}^0$	226
10.6.3	$K \rightarrow \pi^0 \pi^0 \pi_{Dalitz}^0$	226
10.6.4	“Strip”	226
10.6.5	K_{e3}	228
10.6.6	$K_L \rightarrow e^+ e^- \gamma(\gamma)$	230
10.6.7	Summary	231
10.7	Result	232
10.8	Vector Model Result	233
10.9	Conclusions	236
Bibliography		238
 Appendix		
A	Z Slope and DC Inefficiency	242

Tables

Table

1.1	Previous $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ limits. Who: experiment that produced the result. When: year of publication or announcement. BR: limits are at the 90% C.L. Ref.: reference number in the bibliography.	9
1.2	Previous $\text{BR}(K_L \rightarrow e^+ e^- \gamma \gamma, E_\gamma^* > 5 \text{ MeV})$ measurements. Who: experiment that produced the result. When: year of publication or announcement. Stat.: statistical uncertainty. Sys.: systematic uncertainty. Events: number of events observed, after subtracting background. The preliminary number of events listed for NA48 is without background subtraction. Ref.: reference number in the bibliography.	12
1.3	Values of R_X . All masses are from the PDG [37]. Reference [35] gives R_{K^*} as 0.311, which would be true for $K^*(892)^\pm$ mesons, which have mass $891.7 \text{ MeV}/c^2$	15
2.1	Wires and cables per DC plane. Listed by increasing Z . ‘Primed’ planes refer to downstream plane of plane-pair.	29
6.1	Materials in path of decay products. “DC1U” is material between the air gap and 2 cm upstream of wires in DC1. Helium bag 2 sits between DC2 and DC2.	116

6.2	Materials in path of decay products. See table 6.1 for more information. Prob. is the probability of bremsstrahlung being simulated.	120
7.1	Summary of filter tags and physics. All decays are from kaons, unless noted. Tagged is the fraction of events on 2ENEUT tapes accepted by the cuts.	130
8.1	Numbers of simulated $K \rightarrow \pi^0 \pi_{Dalitz}^0$. % is the number on that line di- vided by the number generated. K_L weight is the sum of $ \eta_{00}\mathcal{A}_L(t) ^2/ \mathcal{A}_S(t)+$ $\eta_{00}\mathcal{A}_L(t) ^2$ in good spills. L1, L2, and L3 trigger and EEGG filter sums are unweighted and in all spills. L3 triggers are 2e, n-clus tags only. . .	135
8.2	Bad spill bits summary. The Winter and Summer columns list the per- centage of events in each data set with that bad-spill bit, regardless of whether a cut is made on that bit. The data used pass loose cuts for $K_L \rightarrow \pi^0 \pi_{Dalitz}^0$. The total fraction of data lost to bad spill bits is 9.10%. 137	137
8.3	Non-database bad spills.	138
8.4	Events successfully reconstructed.	138
8.5	Summary of K_L flux numbers.	165
8.6	Uncertainties on fluxes, in units of $10^9 K_L$ s.	165
8.7	Changes in flux when cuts are changed, in units of $10^9 K_L$ s.	166
9.1	Dalitz and radiative-Dalitz kaon decays generated, with various thresh- olds.	171
9.2	Numbers of $K_L \rightarrow e^+ e^- \gamma(\gamma)$. % is the number on that line divided by the number generated. L1, L2, and L3 trigger and EEGG filter sums are in all spills. L3 triggers are 2e, n-clus tags only.	173

9.3	Numbers of $K \rightarrow \pi^0 \pi_{Dalitz}^0$ simulated as background to $K_L \rightarrow e^+ e^- \gamma \gamma$. % is the number on that line divided by the number Generated. K_L weight is the sum of $ \eta_{00} \mathcal{A}_L(t) ^2 / \mathcal{A}_S(t) + \eta_{00} \mathcal{A}_L(t) ^2$ in good spills. L1, L2, and L3 trigger and EEGG filter sums are unweighted and in all spills. L3 triggers are 2e, n-clus tags only.	173
9.4	Branching ratio numbers for $K_L \rightarrow e^+ e^- \gamma \gamma, E_\gamma^* > 5 \text{ MeV}$. Generated number is in good spills only. Expected backgrounds include TRD ac- ceptance; background uncertainties are from statistics and BRs.	189
9.5	Uncertainties on $\text{BR}(K_L \rightarrow e^+ e^- \gamma \gamma, E_\gamma^* > 5 \text{ MeV})$. All sums of uncer- tainties are quadrature sums.	191
9.6	Branching ratio numbers for $K_L \rightarrow e^+ e^- \gamma \gamma$ with different thresholds. Expected backgrounds include TRD acceptance; background uncertain- ties are from statistics and BRs.	192
10.1	Numbers of $K_L \rightarrow \pi^0 e^+ e^-$ generated. % is the number on that line divided by the number Generated. L1, L2, and L3 trigger and EEGG filter sums are in all spills. L3 triggers are 2e, n-clus tags only.	206
10.2	Numbers of $K_L \rightarrow \pi^0 e^+ e^-$ generated using vector model. % is the num- ber on that line divided by the number Generated. L1, L2, and L3 trigger and EEGG filter sums are in all spills. L3 triggers are 2e, n-clus tags only.	206
10.3	Alternative estimates of background to $K_L \rightarrow \pi^0 e^+ e^-$	231
A.1	ρ_X, ρ_Y , and related numbers for Summer and Winter, with various cuts and weights. Stars (*) indicate ρ s actually used. $K \rightarrow \pi^0 \pi_{Dalitz}^0$ decays with all cuts are used.	244

Figures

Figure

1.1	Directly CP violating $K_L \rightarrow \pi^0 e^+ e^-$ decay penguin.	5
1.2	CP conserving $K_L \rightarrow \pi^0 e^+ e^-$ decay.	7
1.3	$K_L \rightarrow e^+ e^- \gamma \gamma$. The other tree-level Feynmann diagrams have the radiated photon coming from the positron and exchange the radiated and direct photons.	13
1.4	$K_L \gamma \gamma$ diagrams. (a) is the pseudoscalar-pseudoscalar transition, while (b) is the vector-vector transition.	14
2.1	Secondary beam, plan view. X and Z scales are different. Dashed lines indicate nominal beam position. X positions of edges that do not impinge on the beams are represented arbitrarily. Elements which do not affect the beams in E799 are not included.	20
2.2	Kaon momenta, as reconstructed in $K_L \rightarrow \pi^+ \pi^- \pi^0$ decays.	23
2.3	KTeV detector. X and Y scales are blown up, relative to Z scale. . . .	25
2.4	Spectrometer elements. X and Z scales differ. The plan view is shown, but the elevation view differs mainly in that no deflection occurs in Y in the magnet.	26
2.5	DC plane pair schematic. Short segment of plane pair, showing upstream and downstream planes, and hexagonal geometry of wire cells.	28
2.6	CsI calorimeter.	31

2.7	TRD cross section, plan view. Scale is approximate, except for wire spacing and as noted.	34
2.8	TRD chamber, beam view.	36
2.9	Veto detectors, plan view. X and Z scales differ.	38
3.1	Example of energy in calorimeter. The shading indicates energy deposited in that crystal, in steps of 0.01–0.1 GeV, 0.1–1.0 GeV, 1.0–10.0 GeV, and more than 10.0 GeV. The circles indicate where the HCC algorithm found clusters.	51
3.2	DCHC overview.	57
3.3	DCHC FASTBUS rack layout, viewed from front. The Summer configuration of Kumquats and Bananas is shown.	60
3.4	Block diagram of Kumquat module.	62
3.5	Block diagram of receiving/latching circuit	63
3.6	Example of hit counting.	64
3.7	Photograph of Banana modules. In the left module, the hit logic and flash TDC chips are visible behind the stiffener plate. The right module shows chips for all other circuits.	65
3.8	Block diagram of Banana module.	66
3.9	Block diagram of TDC circuit. Only one polarity is shown for differential ECL signals for simplicity.	67
3.10	Banana TDC accuracy. Measured Banana TDC time minus LeCroy 3377 TDC time, in units of Banana TDC counts (1.6 ns), for a single channel.	68
3.11	Banana t_1 vs. t_2 and in-time pair region boundary for hits on tracks.	71
3.12	Banana t_1 vs. t_2 and in-time pair region boundary for isolated pairs rejected by Banana.	74
3.13	Structure of DCHC adder tree.	75

5.1	Event display of calorimeter. The circles are centered on found clusters. The dark squares are on channels with HCC bits set.	88
5.2	Event display of drift chambers. Upper plot is X-view, lower plot is Y-view. Solid lines are reconstructed tracks. Dashed lines are reconstructed photon trajectories.	92
5.3	“Raw” SOD. Y-view hits on tracks in $K \rightarrow \pi^0 \pi_{Dalitz}^0$ decays. The entries at SOD=0 are isolated hits.	94
5.4	“Corrected” SOD. Y-view hits on tracks in $K \rightarrow \pi^0 \pi_{Dalitz}^0$ decays. (B) is the peak in (A), with a gaussian fit.	99
5.5	TRD ADC values for $e^\pm$ and $\pi^\pm$ tracks. All from upstream plane in TRD chamber 3, Ke3 decays in runs 8384, 8387, and 8397.	101
5.6	Π_{TRD} for $e^\pm$ and $\pi^\pm$ tracks. Ke3 decays in runs 8384, 8387, and 8397.	102
6.1	Momentum vs. vertical angle at target. As mentioned in chapter 2, the proton beam angle is -0.48×10^{-2} rad. K^0 and $\overline{K}^0$ distributions are similar and are averaged here.	106
6.2	(A) Vertex Z and (B) momentum for a sample of generated $K_L \rightarrow e^+ e^- \gamma$ MC.	108
6.3	P vs. Z for a sample of generated $K \rightarrow \pi^0 \pi_{Dalitz}^0$ MC.	109
6.4	(A) Vertex Z and (B) momentum for a sample of generated $K \rightarrow \pi^0 \pi_{Dalitz}^0$ MC.	109
6.5	Momentum times multiple-scattering angle projected onto plane, for e^+ and e^- at vacuum window. The fit to a gaussian goes from -1.4 mrad GeV/c to $+1.4$ mrad GeV/c, a region which includes 98% of the distribution.	117
6.6	Log of the bremsstrahlung opening angle θ	121
6.7	Energy of the bremsstrahlung photon.	121

8.1	Number of tracks in $K \rightarrow \pi^0 \pi_{Dalitz}^0$ reconstruction. Dots are data and line is MC, normalized to data. No cut has limited the number of clusters yet.	139
8.2	(Left) Momentum resolution in simulated $K \rightarrow \pi^0 \pi_{Dalitz}^0$. Reconstructed total momentum minus generated total momentum in a sample of MC events, near fiducial limits. (Right) Reconstructed momentum for all data and MC, after all other cuts.	140
8.3	Decay-vertex-Z resolution in simulated $K \rightarrow \pi^0 \pi_{Dalitz}^0$. Reconstructed Z minus generated Z in a sample of MC events, near fiducial limits. (Left) Reconstructed vertex uses DC tracks, (Right) reconstructed vertex uses Z_{π^0}	141
8.4	Vertex Z for all data and MC, after all other cuts. (Left) DC track vertex, (Right) Z_{π^0}	141
8.5	Minimum distance from track to upstream V-bank beam hole. All other cuts are made.	143
8.6	Collar Anti energy. Maximum sum of energy in any counter added to its neighbors, in GeV. All other cuts are made.	144
8.7	Total calorimeter energy (GeV), for ETOT verification. All other cuts are made, except for total momentum cut.	144
8.8	Smallest cluster-hole distance (meters). All other cuts are made.	146
8.9	All calorimeter cluster energies (GeV). All other cuts are made, except for ETOT verification.	146
8.10	Number of TRD hits per DC track. All other, non-TRD cuts are made.	147
8.11	(A) Π_{TRD} for all data after all other cuts. (B) E/p for tracks failing and passing Π_{TRD} cut. The passing spectrum is normalized to the failing spectrum.	148

8.12 Π_{TRD} for each track in all data events, after all other cuts. The box indicates the cut.	149
8.13 E/p for each track in Data and MC events, after all other cuts.	150
8.14 E/p vs. $ \vec{p} $ (GeV/c) for each track in Data and MC events, after all other cuts.	152
8.15 E/p for each track in Data and MC events, after all other cuts. The parameters are results of fitting to sum of two Gaussians.	152
8.16 Track separation at DC1 in X (Top) and Y (Bottom), in meters. All other cuts are made except track separation and opening angle. The y-axis for the ratio plots is from 0.8 to 1.2.	153
8.17 Track opening angle (radians). All other cuts are made.	154
8.18 Two-photon invariant mass, in GeV/c^2 , for all pairings (Left) and for best pairing (Right). Dots are data and solid line is MC. All other cuts are made.	157
8.19 Two-photon invariant mass, in GeV/c^2 , data (Left) and simulated $K \rightarrow \pi^0 \pi_{\text{Dalitz}}^0$ (Right). All other cuts are made.	157
8.20 $e^+e^-\gamma$ invariant mass, in GeV/c^2 , data (Left) and simulated $K \rightarrow \pi^0 \pi_{\text{Dalitz}}^0$ (Right). All other cuts are made.	158
8.21 Definition of transverse momentum, $\vec{P}_\perp$	159
8.22 $\log_{10} \vec{P}_T^2$ using track-vertex Z (Dots) and Z_{π^0} (Line), measuring $\vec{P}_T^2$ in $(\text{GeV}/c)^2$. All other cuts are made.	159
8.23 $\vec{P}_T^2$ (using Z_{π^0}), in GeV^2/c^2 . All non- $\vec{P}_T^2$ cuts are made.	161
8.24 $\log_{10} \vec{P}_T^2$ (using Z_{π^0}), where the dots are data. The smooth line is a fit to the data. The MC is shown as a histogram, normalized to make the peak heights match below the cut. $\vec{P}_T^2$ is measured in $(\text{GeV}/c)^2$. All non- $\vec{P}_T^2$ cuts are made.	161

8.25	$e^+e^-\gamma\gamma\gamma$ invariant mass, in GeV/c^2 . The straight line is the fit to the indicated region, extended to higher masses. All other cuts are made.	162
8.26	Fits to $e^+e^-\gamma\gamma\gamma$ invariant mass, in GeV/c^2 . All other cuts are made.	163
8.27	$\log_{10} \vec{P}_T^2$ in $(\text{GeV}/c)^2$ (y-axis) vs. $M_{e\gamma\gamma\gamma}$ in GeV/c^2 (x-axis). The box indicates the cuts. All other cuts are made.	164
8.28	Kaon beam angles, equal to vertex X or Y divided by Z_{π^0} , in radians. All cuts are made.	167
8.29	Run number distribution for Winter (Left) and Summer (Right). The MC is normalized so that its integral over Winter+Summer is equal to the integral of data over Winter+Summer. All cuts are made.	168
8.30	e^+e^- invariant mass, in GeV/c^2 . All cuts are made.	169
9.1	E_γ^* in GeV after all other cuts. The data/MC ratio is made using the sum of MC distributions.	175
9.2	E_γ^* as generated in simulation and as reconstructed, for $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC, in GeV . (Right) Difference between E_γ^* s when both are less than 25 MeV. All other cuts are made.	176
9.3	Efficiency vs. E_γ^* , in GeV . All other cuts are made.	176
9.4	$\gamma\gamma$ invariant mass in GeV/c^2 after all other cuts. The data/MC ratio is made using the sum of MC distributions.	178
9.5	$\gamma\gamma$ invariant mass as generated in simulation and as reconstructed, for $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC, in GeV/c^2 . (Right) Difference between $M_{\gamma\gamma}$ s when both are less than 60 MeV/c^2 . All other cuts are made.	179
9.6	Efficiency vs. $M_{\gamma\gamma}$, in GeV/c^2 . All other cuts are made.	179
9.7	$\log_{10} \vec{P}_T^2$ in $(\text{GeV}/c)^2$, after all other cuts. The data/MC ratio is made using the $K_L \rightarrow e^+e^-\gamma\gamma$ MC distribution.	180

9.8	$e^+e^-\gamma\gamma$ invariant mass in GeV/c^2 after all other cuts. The data/MC ratio is made using the summed MC distributions.	182
9.9	$e^+e^-\gamma\gamma$ invariant mass in GeV/c^2 after all other cuts, for data (Left) and simulated $K_L \rightarrow e^+e^-\gamma\gamma$ (Right).	183
9.10	$\log_{10} \theta_{min}$, where θ_{min} is measured in radians, after all other cuts. The data/MC ratio is made using the summed MC distributions.	184
9.11	$\log_{10} \text{MINTPD}$, where MINTPD is measured in meters, after all cuts except θ_{min} are made. The data/MC ratio is made using the summed MC distributions.	186
9.12	$M_{ee\gamma\gamma}$ for simulated $K \rightarrow \pi^0\pi^0\pi^0_{Dalitz}$, in GeV/c^2 . The MC sample size corresponds to $\sim 31.6\%$ of the flux. All other cuts are applied.	188
9.13	Square of BMS form factor, vs α_{K^*} and M_{ee} in GeV/c^2 . Note that equation 1.3 gives $f(x)$, while $f(M_{ee}, \alpha_{K^*})^2$ is plotted here.	194
9.14	e^+e^- invariant mass, in GeV/c^2 . Note that the summed MC is normalized to have the same area as the data, while the other MC distributions are normalized by flux and BR. The data/MC ratio plot uses summed MC. All cuts are applied.	195
9.15	e^+e^- invariant mass for data and MC with several α_{K^*} , in GeV/c^2 . The error bars on the MC levels are smaller than the symbols at the levels.	197
9.16	χ^2 for data fit to summed MC with varied α_{K^*} . The arrow indicates the minimum of the fit at $\alpha_{K^*} = +0.0145$ and $\chi^2 = 4.30$	197
9.17	Ratio of data to MC distributions of e^+e^- invariant mass for $K \rightarrow \pi^0\pi^0\pi^0_{Dalitz}$, in GeV/c^2 . Note only the range 0-0.15 GeV/c^2 is shown.	199
9.18	$\gamma\gamma$ invariant mass in GeV/c^2 after all other cuts. The data/MC ratio is made using the sum of MC distributions.	202
9.19	Phase-space variable y_γ	203
9.20	Phase-space angles y_γ and ϕ . All cuts are made.	203

10.1	$\pi^+\pi^-\gamma\gamma$ invariant mass, in GeV/c^2 . Only general, non-TRD cuts are made. The solid line is simulated $K_L \rightarrow \pi^0 e^+ e^-$ with arbitrary normalization.	207
10.2	TRD Π_{TRD} for both tracks in data events.	208
10.3	e^+e^- invariant mass, in GeV/c^2 , after general, $M_{\pi^+\pi^-\gamma\gamma}$, and TRD cuts. The $K_L \rightarrow \pi^0 e^+ e^-$ distribution has arbitrary normalization.	210
10.4	Transverse momentum, $\vec{P}_T^2$, using Z_{π^0} , in $(\text{GeV}/c^2)^2$, after general, $M_{\pi^+\pi^-\gamma\gamma}$, TRD, and M_{ee} cuts. The $K_L \rightarrow \pi^0 e^+ e^-$ distribution has arbitrary normalization.	211
10.5	$e^+e^-\gamma\gamma$ invariant mass, using Z_{π^0} , in $(\text{GeV}/c^2)^2$, after general, $M_{\pi^+\pi^-\gamma\gamma}$, TRD, M_{ee} , and $\vec{P}_T^2$ cuts. The $K_L \rightarrow \pi^0 e^+ e^-$ distribution has arbitrary normalization. The fit “box” has been excluded in the data and $K_L \rightarrow e^+e^-\gamma(\gamma)$ samples. The excess data below $420 \text{ MeV}/c^2$ are mostly from K_{e4} decays.	211
10.6	$M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , for data, before kinematic cuts.	213
10.7	$M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , for simulated $K_L \rightarrow e^+e^-\gamma(\gamma)$, before kinematic cuts.	214
10.8	$M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , before kinematic cuts. The shaded boxes are the results of the background fit. The empty boxes are data, plotted on a linear Z scale. Data events inside the “box” are not shown.	215
10.9	$M_{\gamma\gamma}$ in GeV/c^2 , before kinematic cuts, in $M_{\gamma\gamma}$, $M_{ee\gamma\gamma}$ “swath.” Crosses are data and solid line is fit.	216
10.10	$M_{ee\gamma\gamma}$ in GeV/c^2 , before kinematic cuts, in $M_{\gamma\gamma}$, $M_{ee\gamma\gamma}$ “swath.” Crosses are data and solid line is fit.	217
10.11	$\theta_{\min}$ for events in swath but outside box. Simulated $K_L \rightarrow \pi^0 e^+ e^-$ events are from inside box, and have arbitrary normalization.	217

10.12	Absolute value of y_γ for events in swath but outside box. Simulated $K_L \rightarrow \pi^0 e^+ e^-$ events are from inside box, and have arbitrary normalization.	218
10.13	Absolute value of $ \phi $ for events in swath but outside box. Simulated $K_L \rightarrow \pi^0 e^+ e^-$ events are from inside box, and have arbitrary normalization.	218
10.14	Estimated branching ratio limit, as function of cuts on θ_{min} and $ y_\gamma $	220
10.15	$M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , after all other cuts. The shaded boxes are the results of the background fit. The empty boxes are data, plotted on a linear Z scale. Data events inside the “box” are not shown.	222
10.16	$e^+ e^-$ invariant mass, in $(\text{GeV}/c^2)^2$. All non- $M_{\gamma\gamma}$ cuts have been applied to MC. The M_{ee} cut shown is only used for the $M_{\gamma\gamma}$ resolution study.	224
10.17	$\gamma\gamma$ invariant mass, in $(\text{GeV}/c^2)^2$, for $K_L \rightarrow \pi^+ \pi^- \pi^0$ data (Left) and simulated $K_L \rightarrow \pi^0 e^+ e^-$ (Right).	224
10.18	$M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , after all other cuts, for simulated $K \rightarrow \pi^0 \pi^0_{Dalitz}$	227
10.19	$M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , after all other cuts, for simulated $K \rightarrow \pi^0 \pi^0 \pi^0_{Dalitz}$. Note that the $M_{ee\gamma\gamma}$ scale is expanded in this plot.	227
10.20	TRD Π_{TRD} for data events with $130 < M_{\gamma\gamma} < 140 \text{ MeV}/c^2$ and $400 < M_{ee\gamma\gamma} < 465 \text{ MeV}/c^2$, after all other cuts. The only region accepted by the TRD cuts is the bin that has 18 events.	229
10.21	$e^+ e^- \gamma\gamma$ invariant mass after all other cuts, in GeV/c^2 . The squares are the background fit, which excludes the strip, and the fit error bars are not shown. The fit and the data use the same $5 \text{ MeV}/c^2$ bins. No data from inside the box are shown. The $M_{\gamma\gamma}$ cut is applied.	229
10.22	$M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , after all other cuts, for data <i>failing</i> Π_{TRD} cut. Thus all events shown here, including those in the box, are rejected.	230

10.23	$M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , after all other cuts. The box is open. The shaded bins are $K_L \rightarrow \pi^0 e^+ e^-$ MC, normalized to 1.	232
10.24	Estimated branching ratio limit, as function of cuts on θ_{min} and $ y_\gamma $, assuming vector model for $K_L \rightarrow \pi^0 e^+ e^-$	234
10.25	$e^+ e^-$ mass, in GeV/c^2 , for $K_L \rightarrow \pi^0 e^+ e^-$ MC samples after all other cuts.	235
A.1	Ratios of data to MC vs. vertex Z , in meters, for $K \rightarrow \pi^0 \pi^0_{Dalitz}$ decays. Winter is on top, Summer on the bottom, charged-track Z on the left, and Z_{π^0} on the right. All cuts are made.	243
A.2	Ratios of data to MC vs. vertex Z , with ρ weighting on MC. Winter is on top, Summer on the bottom, charged-track Z on the left, and Z_{π^0} on the right. All cuts are made, and $K \rightarrow \pi^0 \pi^0_{Dalitz}$ decays are used. . . .	246

Chapter 1

Introduction

1.1 CP Symmetry

In physics, symmetry exists when some change can be imposed on a system but the properties of the system remain constant. The conservation law that results from the the existence of a symmetry can be crucial in understanding the universe; i.e., symmetry under a shift in time coordinate gives conservation of energy. On the other hand, a symmetry which is sometimes broken or violated can also provide insights; i.e. when beta decay of Co^{60} breaks parity conjugation symmetry — that is, $(x, y, z) \rightarrow (-x, -y, -z)$ — it is because of the existence of parity symmetry violation at a much larger energy scale (the W boson). Another such broken symmetry, which has received much attention for the last 40 years, is that of charge and parity conjugation (CP). Parity (P), as mentioned above, refers to reversing the spatial coordinates of a system. Charge (C) refers to charge conjugation, exchanging particles with antiparticles. Prior to 1956, it was thought that C and P were both exact symmetries. That is, the decay $a \rightarrow bcd$ proceeds identically to $\bar{a} \rightarrow \bar{b}\bar{c}\bar{d}$ if it obeys C symmetry. If it obeys P symmetry, then $a \rightarrow bcd$ would have the same properties as $a \rightarrow bcd$ with reverse coordinates. Indeed, many processes, including all electromagnetic interactions, conserve both C and P. However, P symmetry violation was seen in the beta decay mentioned above in 1957 [1]; observation of P violation in

$\pi \rightarrow \mu\nu$ decay soon followed [2][3]. These decays simultaneously violate C symmetry. However, these decays are symmetric under the combined operations of C and P. Thus it was thought that that CP might be a good symmetry. But that belief persisted for only a few years, until CP violation was seen in $K \rightarrow \pi\pi$ decays, as discussed below.

There are several reasons for the fascination with CP violation. One reason stems from the invariance of CPT (CP + time reversal conjugation), which holds in any local field theory with Lorentz invariance. If CPT is not violated but CP is, then T violation is implied, meaning that microscopic physical reactions see an arrow of time. A second reason is that CP violation points towards the existence of a third family of quarks and leptons, somewhat like how P violation points towards the existence of the W, and CP violation was observed well before any members of the third family were observed. A third reason relates to the preponderance of matter over antimatter in the universe. Assuming that matter and antimatter were present in equal amounts in the early universe, one of the conditions for the observed matter dominance is that there is CP violation [4]. (The other conditions require thermal disequilibrium and baryon-number non-conservation.)

1.1.1 Neutral Kaon System

The breaking of CP symmetry was found by studying neutral kaon decays. The existence of two types of neutral kaons has been known since 1956 [5]. The short-lived kaon, K_S , can decay into two pions ($\pi^+\pi^-$ or $\pi^0\pi^0$), while the long-lived kaon, K_L , can decay into three pions ($\pi^+\pi^-\pi^0$ or $\pi^0\pi^0\pi^0$). The two-pion final state has CP eigenvalue +1 (CP even) while the three-pion final state has CP eigenvalue -1 (CP odd). Therefore, if CP were an exact symmetry, then K_L decays into two pions would be forbidden. CP violation (CPV) was first observed in 1964 in the decay of K_L particles into two pions [6].

CP was demolished as an exact symmetry, although it could still be considered

a good symmetry because of the rarity of $K_L \rightarrow \pi\pi$ decays. The small size of CPV, as compared to CV and PV, made a detailed understanding of the processes of CPV desirable. The superweak model proposed that CPV occurs entirely because the physical kaons are mixtures of CP eigenstates [7]. This mixture would occur through a $\Delta S = 2$ interaction that allowed $K^0 \leftrightarrow \bar{K}^0$ mixing. If K_1 is the CP-even eigenstate ($|K_1\rangle \sim |K^0\rangle + |\bar{K}^0\rangle$) and K_2 is the CP-odd eigenstate ($|K_2\rangle \sim |K^0\rangle - |\bar{K}^0\rangle$) then

$$|K_S\rangle = |K_1\rangle + \epsilon|K_2\rangle$$

$$|K_L\rangle = \epsilon|K_1\rangle + |K_2\rangle$$

The parameter ϵ determines the magnitude of this “indirect” CPV. Other models, including the standard model, allow the existence of “direct” CPV. In direct CPV, the asymmetry occurs not in the parent particle but in the decay process. The magnitude of direct CPV for a decay mode is parameterized by ϵ' . By identifying and counting different $K^0 \rightarrow \pi\pi$ decays, the ratio $\text{Re}(\epsilon'/\epsilon)$ can be measured for the two-pion mode. The E832 experiment by the KTeV collaboration has found $\text{Re}(\epsilon'/\epsilon) = (28.0 \pm 4.1) \times 10^{-4}$ [8], seeming to confirm the existence of direct CPV and thus disproving superweak models.

1.1.2 Rare Decays

CPV in $K_L \rightarrow \pi\pi$ occurs mostly via indirect CPV, as shown by the small value of $\text{Re}(\epsilon'/\epsilon)$. Other K_L decays that have larger degrees of direct CPV would provide additional tests of the standard model’s ability to predict direct CPV. Even if the standard model can accommodate this value of $\text{Re}(\epsilon'/\epsilon)$ (and [14] suggests that may be difficult), it is still only one value. Measurements of other manifestations of direct CPV would provide discriminating tests of a model’s accuracy. The class of decays $K_L \rightarrow \pi^0 l \bar{l}$ have been identified as having CPV components, and possibly significant directly CP violating contributions. Briefly summarizing each mode:

- $K_L \rightarrow \pi^0 e^+ e^-$: This decay is discussed further below. To summarize, it can proceed via direct CPV, indirect CPV, or CP conserving channels. The exact contribution of each channel is uncertain, although they are all expected to be within an order of magnitude of each other. The overall branching ratio (BR) is expected to be in the 10^{-11} neighborhood. This decay has not been observed, but an upper limit on the BR of 4.3×10^{-9} at the 90% C.L. has been measured [9].
- $K_L \rightarrow \pi^0 \mu^+ \mu^-$: This mode is much like $K_L \rightarrow \pi^0 e^+ e^-$. Phase space considerations alone would indicate $\text{BR}(K_L \rightarrow \pi^0 \mu^+ \mu^-) \simeq 0.2 \times \text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ [10], though the lepton masses might enhance form factors that could significantly enhance the BR of the muon mode [11]. The BR of this decay has been measured to be less than 3.4×10^{-10} at the 90% C.L. [12].
- $K_L \rightarrow \pi^0 \nu \bar{\nu}$: This mode proceeds through direct CPV primarily [13]. It is expected to have a BR of $(1.6 \text{ to } 3.9) \times 10^{-11}$ [14]. Because the neutrinos are not observed, it is difficult to search for this mode. Searching for the decay with a secondary Dalitz decay by the pion, $\pi^0 \rightarrow e^+ e^- \gamma$, has given the best limit on the BR, 5.9×10^{-7} at the 90% C.L. [16].

An additional potential benefit of searching for any rare decay mode is the sensitivity to new physics. Processes outside the standard model may allow rare decays to occur much more often. Observing the rare decay and measuring a BR well above the standard model predictions would clearly prove the existence of such a process. Alternatively, not observing the rare decay and establishing a new upper limit on the BR disproves the existence of some processes. Thus the search for rare decays can be fruitful even when the sensitivity of an experiment is too poor to detect the standard model prediction.

1.2 $K_L \rightarrow \pi^0 e^+ e^-$

The primary goal of the analysis described in this thesis is to search for the decay $K_L \rightarrow \pi^0 e^+ e^-$. This section details some of the theoretical predictions for the decay and gives an overview of experimental issues.

1.2.1 Standard Model

The decay $K_L \rightarrow \pi^0 e^+ e^-$ can violate CP directly through the $\Delta S = 1$ process $K_L \rightarrow \pi^0 \gamma^*$. The directly CP violating part of $K_L \rightarrow \pi^0 e^+ e^-$ is fairly straightforward under the standard model. The decay occurs through penguin diagrams shown in Figure 1.1. Similar penguins allow direct CPV in the kaon decays discussed above. In $K_L \rightarrow \pi^0 \mu^+ \mu^-$, the electron lines are simply replaced with muons. In $K_L \rightarrow \pi^0 \nu \bar{\nu}$, neutrinos replace the electrons and no γ diagram appears. In $K_L \rightarrow \pi \pi$, a $u\bar{u}$ or $d\bar{d}$ pair replaces the electrons and the γ, Z can be replaced by a gluon. It is destructive interference between the gluon diagrams and the γ, Z diagrams that suppresses direct CPV in two-pion decays. It is the lack of such interference that provides hope of seeing large direct CPV effects in $K_L \rightarrow \pi^0 l l$ decays.




Figure 1.1: Directly CP violating $K_L \rightarrow \pi^0 e^+ e^-$ decay penguin.

Because of the $s \rightarrow u, c, t \rightarrow d$ transitions in Figure 1.1, measuring direct CPV constrains the Cabibbo-Kobayashi-Maskawa quark-mixing matrix in the standard model. The direct CPV part of $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ is related to CKM matrix element parameters by [22]:

$$\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)_{\text{direct}} = 4.16(y_{7A}^2 + y_{7V}^2)(\text{Im}\lambda_t)^2, \quad (1.1)$$

$$\text{where } \text{Im}\lambda_t = \text{Im}V_{td}V_{ts}^* = |V_{ub}||V_{cb}|\sin\delta = A^2\lambda^5\eta = |V_{cb}|^2|V_{us}|\eta. \quad (1.2)$$

The V s are standard CKM parameters, while A , λ , and η are from the Wolfenstein parameterization. The factor 4.16 comes from the BR for the related decay, $K^+ \rightarrow \pi^0 e^+ \nu$. The parameters y_{7A} and y_{7V} are QCD corrections of order unity times α_{EM} that can be extracted from tables in [22] using the top-quark mass ($173.8 \pm 5.2 \text{ GeV}/c^2$); respectively they are 0.005410 and -0.005314 . The role of η in the Wolfenstein parameterization is the magnitude of the imaginary component of the CKM matrix; if the CKM matrix is real, then the standard model forbids CPV. A global standard-model fit, including the new measurements of $\text{Re}(\epsilon'/\epsilon)$, gives $\text{Im}\lambda_t = (1.38 \pm 0.14) \times 10^{-4}$ [14]. The same authors predict the $K_L \rightarrow \pi^0 e^+ e^-$ branching ratio, scanning over standard-model parameters, to be

$$\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)_{\text{direct}} = (2.8 \text{ to } 6.5) \times 10^{-12}.$$

However, to observe direct CPV in $K_L \rightarrow \pi^0 e^+ e^-$ it will have to be disentangled from the contributions that are indirectly CP violating and CP conserving.

The decay can conserve CP when the K_L first decays into $\pi^0 \gamma^* \gamma^*$, and the photons then convert into electrons, as shown in Figure 1.2. This component of $K_L \rightarrow \pi^0 e^+ e^-$ can be understood by studying $K_L \rightarrow \pi^0 \gamma \gamma$. A measurement of $\text{BR}(K_L \rightarrow \pi^0 \gamma \gamma)$ was first published only in 1990 [18], but it was found to be a factor of three more common than the $\mathcal{O}(p^4)$ chiral perturbation theory (ChPT) prediction. By including $\mathcal{O}(p^6)$ effects, much better agreement between experiment and theory was obtained

for the BR [19]. A similar discrepancy between $\mathcal{O}(p^4)$ and $\mathcal{O}(p^6)$ ChPT predictions occurs for the decay $K_L \rightarrow \pi^0 \gamma e^+ e^-$, with $\mathcal{O}(p^4)$ giving a BR of 1.0×10^{-8} while $\mathcal{O}(p^6)$ gives a BR of 2.3×10^{-8} [20]. KTeV recently measured $\text{BR}(K_L \rightarrow \pi^0 \gamma e^+ e^-)$ to be $(2.20 \pm 0.48 \pm 0.11) \times 10^{-8}$ [21] [46] (preliminary), supporting the validity of the $\mathcal{O}(p^6)$ calculations. Using $\mathcal{O}(p^6)$, Donoghue and Gabbiani predicted CP-conserving $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ as a function of the (then-uncertain) vector-meson exchange parameter, a_V [22]. KTeV has recently measured $a_V = -0.72 \pm 0.05 \pm 0.06$ [23]. This a_V suggests that

$$\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)_{\text{conserving}} = (1 \text{ to } 3) \times 10^{-12}.$$




Figure 1.2: CP conserving $K_L \rightarrow \pi^0 e^+ e^-$ decay.

The indirectly CP violating part of $K_L \rightarrow \pi^0 e^+ e^-$ occurs because of the same mass mixing that allows most $K_L \rightarrow \pi\pi$ decays. That is, the indirectly CP violating decay is $K_L \rightarrow K_S \rightarrow \pi^0 e^+ e^-$. Thus the indirect amplitude of $K_L \rightarrow \pi^0 e^+ e^-$ is just ϵ times the amplitude of $K_S \rightarrow \pi^0 e^+ e^-$. However, there is considerable theoretical uncertainty in $\mathcal{A}(K_S \rightarrow \pi^0 e^+ e^-)$. Standard model predictions for $\text{BR}(K_S \rightarrow \pi^0 e^+ e^-)$ are about $\mathcal{O}(10^{-9})$, although they could well be as much as an order of magnitude higher or lower [24]. Current measurements give $\text{BR}(K_S \rightarrow \pi^0 e^+ e^-) < 1.1 \times 10^{-6}$ at the 90% C.L. [25], so considerable progress remains to be made in understanding this

mode. Multiplying the prediction by $|\epsilon|^2$ and the ratio of the lifetimes gives

$$\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)_{\text{indirect}} \sim 3 \times 10^{-12}.$$

To determine the magnitude of direct CPV after $K_L \rightarrow \pi^0 e^+ e^-$ has been observed, two effects must have been measured: $K_S \rightarrow \pi^0 e^+ e^-$ by itself and the interference of the CP violating and CP conserving amplitudes [24]. The CPV-CPC interference may cause the electrons to have more energy than the positrons, or vice-versa, depending on the exact value of parameters [10]. The presence of direct CPV may substantially change the size of this asymmetry, enough that measurements of both $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ and electron-positron asymmetry can unambiguously signal the existence of direct CPV [22].

1.2.2 Outside Standard Model

Some nonstandard models predict enhancements to $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ higher than those of standard model predictions. Flynn and Randall examined several non-standard models for enhancements, but found few possibilities [28]. A process that has received more recent attention is an enhancement of the effective $\bar{s}dZ$ vertex predicted in a supersymmetric extension of the standard model [26]. An enhanced $\bar{s}dZ$ vertex would increase $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$, but it would also enhance the direct CPV part of $K_L \rightarrow \pi\pi$, so measuring $\text{Re}(\epsilon'/\epsilon)$ constrains such models. With the (now-obsolete) constraint $0 < \epsilon'/\epsilon < 20 \times 10^{-4}$, $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ could be as large as 2.2×10^{-10} [27], although this would be an extreme case. The measurement $\text{Re}(\epsilon'/\epsilon) = (28.0 \pm 4.1) \times 10^{-4}$ could still allow $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-) < \sim 3 \times 10^{-10}$, an enhancement of two orders of magnitude above the standard model.

1.2.3 Previous Searches

Previous searches for $K_L \rightarrow \pi^0 e^+ e^-$ decays have found none. The results for the most recent round of experiments are summarized in Table 1.1. The best limit

comes from the E799 phase-i experiment, the direct predecessor to KTeV's E799 phase-ii experiment (herein referred to as simply E799). Clearly, these limits are not even close to standard model predictions, nor even to the limits of nonstandard models.

Table 1.1: Previous $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ limits. Who: experiment that produced the result. When: year of publication or announcement. BR: limits are at the 90% C.L. Ref.: reference number in the bibliography.

Who	When	$\text{BR} \times 10^9 <$	Ref.
E-731	1990	7.5	[29]
E845	1990	5.5	[30]
E799-i	1993	4.3	[9]

1.2.4 Backgrounds

Several kaon decays can mimic $K_L \rightarrow \pi^0 e^+ e^-$ decay, and these need to be understood and rejected in any search for $K_L \rightarrow \pi^0 e^+ e^-$.

- $K_L \rightarrow e^+ e^- \gamma \gamma$: This is the primary background to $K_L \rightarrow \pi^0 e^+ e^-$ in this search. It shares the same final state particles with the signal mode, and thus can only be rejected by kinematic and phase space cuts. The decay $K_L \rightarrow e^+ e^- \gamma \gamma$ itself is very rare, with fewer than 221 events having been seen in previous experiments. Section 1.3 discusses this mode.
- $K_L \rightarrow e^+ e^- \gamma$: The Dalitz decay of the K_L can mimic $K_L \rightarrow \pi^0 e^+ e^-$ if an extra photon is introduced. This can happen if the e^+ or e^- radiates a photon while passing through matter in the detector. In such cases there tends to be a very small angle between the trajectories of the e and the photon, which can be cut on. Also, an extra photon can come from accidental activity in the detector. In that case, the invariant mass of the $e^+ e^- \gamma \gamma$ system tends to be higher than a K^0 mass, so a cut on $M_{ee\gamma\gamma}$ suppresses the background.

- $K_L \rightarrow \pi^0 \gamma e^+ e^-$: This decay is potentially a background if the photon from the kaon decay is lost. This will cause $M_{ee\gamma\gamma}$ to be less than the K^0 mass, and because this photon tends to be energetic, $M_{ee\gamma\gamma}$ tends to be quite far from M_{K^0} [46]. In addition, the BR of $K_L \rightarrow \pi^0 \gamma e^+ e^-$ is small, $\sim 2.2 \times 10^{-8}$.
- $K_L \rightarrow \pi^0 \gamma \gamma$: If one or two photons convert into a electron-positron pair in matter in the detector, this decay can be a background. If one photon converts, another photon must be lost, causing $M_{ee\gamma\gamma}$ to be low. If the lost or converted photons come from the decaying pion, then the pion mass will not be reconstructed. Also, M_{ee} in photon conversions tends to be very small. If two photons convert, the decay can be a background if an electron from one pair and a positron from the other pair are lost. In this case, M_{ee} can be large. However, this mode is still suppressed by kinematics (two lost particles) and by the low probability that a single photon will convert ($\sim 1.5\%$). Yet another background mode for this decay would have a Dalitz decay of the pion. This mode would be just like $K \rightarrow \pi^0 \pi_{Dalitz}^0$ (below), but with no reconstructed pion mass and suppressed by the small BR of $K_L \rightarrow \pi^0 \gamma \gamma$.
- $K \rightarrow \pi^0 \pi_{Dalitz}^0$: Although CP violating, $K_L \rightarrow \pi^0 \pi_{Dalitz}^0$ is still common enough to be a potential background. The K_S impurity in our K_L beam is small, but $K_S \rightarrow \pi^0 \pi_{Dalitz}^0$ is also a potential background. They are both suppressed because one photon must be lost, so $M_{ee\gamma\gamma}$ tends to be low. Also, M_{ee} is restricted to being less than the π^0 mass, while there is no such kinematic limit on the signal mode.
- $K_L \rightarrow \pi^0 \pi^0 \pi_{Dalitz}^0$: This mode is very common, but three of the photons must be lost, making kinematic cuts very effective. Also, like $K \rightarrow \pi^0 \pi_{Dalitz}^0$, M_{ee} must be less than the π^0 mass.

- $K_L \rightarrow \pi^+\pi^-\pi^0$: This common mode is identical to the signal mode, save that the electrons are replaced with charged pions. This causes $M_{ee\gamma\gamma}$ to be low. Also, the charged pions are susceptible to pion/electron discrimination techniques, such as those provided by the TRD and the calorimeter.
- $K_L \rightarrow \pi e \nu$ (Ke3): This was the limiting background in E799-i. This mode has to accumulate two extra photons to be background, either from accidental activity or from the electron radiating in matter. Kinematic cuts can then suppress this background. The charged pion also makes the decay susceptible to pion/electron discrimination.

1.3 $K_L \rightarrow e^+e^-\gamma\gamma$

To understand $K_L \rightarrow e^+e^-\gamma\gamma$ as a background to $K_L \rightarrow \pi^0 e^+e^-$, better measurements of $K_L \rightarrow e^+e^-\gamma\gamma$ are needed, and $K_L \rightarrow e^+e^-\gamma\gamma$ is of interest in its own right. Therefore, I study this mode and measure its BR before searching for $K_L \rightarrow \pi^0 e^+e^-$.

As a radiative decay, one must specify an infrared cutoff when discussing $K_L \rightarrow e^+e^-\gamma\gamma$. For comparison with previous measurements and predictions, I use the cutoff $E_\gamma^* > 5$ MeV. This means that the energy of each photon in the kaon rest frame, E_γ^* , must be greater than 5 MeV. Another natural cutoff to consider would be on the invariant mass of the two photons, $M_{\gamma\gamma}$. This variable is very important in rejecting $K_L \rightarrow e^+e^-\gamma\gamma$ background in the search for $K_L \rightarrow \pi^0 e^+e^-$, because the latter decay has $M_{\gamma\gamma} = M_{\pi^0} \pm$ resolution effects, while the former has a broad spectrum in $M_{\gamma\gamma}$. The Monte Carlo (MC) simulation of the detector, which measures detector acceptance, uses a cutoff of $M_{\gamma\gamma} > 1$ MeV when simulating $K_L \rightarrow e^+e^-\gamma\gamma$.

1.3.1 Previous Measurements

Previous experimental results are given in Table 1.2. Only 221 $K_L \rightarrow e^+e^-\gamma\gamma$ events have ever been observed. So while the existence of this decay mode has been firmly established, there has been no precise study of its BR or kinematics.

Table 1.2: Previous BR($K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV}$) measurements. Who: experiment that produced the result. When: year of publication or announcement. Stat.: statistical uncertainty. Sys.: systematic uncertainty. Events: number of events observed, after subtracting background. The preliminary number of events listed for NA48 is without background subtraction. Ref.: reference number in the bibliography.

Who	When	BR $\times 10^7$	Stat. $\times 10^7$	Sys. $\times 10^7$	Events	Ref.
BNL845	1992	6.6	± 3.2	—	17	[31]
E799i	1994	6.5	± 1.2	± 0.6	58	[32]
NA31	1998	8.0	± 1.5	$^{+1.4}_{-1.2}$	40	[33]
NA48-prelim	1998	4.6	± 0.7	± 1.4	(106)	[34]

1.3.2 Standard Model

The decay $K_L \rightarrow e^+e^-\gamma\gamma$ occurs mainly when a photon is radiated internally in kaon Dalitz decay as shown in Figure 1.3. This radiative part of the decay is calculated using quantum electrodynamics. The same QED calculations are used to determine the radiative corrections to Dalitz decays, so study of $K_L \rightarrow e^+e^-\gamma\gamma$ is useful as a check of these calculations.

$K_L \rightarrow e^+e^-\gamma$, and thus $K_L \rightarrow e^+e^-\gamma\gamma$, proceed through an effective $K_L\gamma\gamma$ vertex. The vertex dynamics are parameterized by Bergström, Massó, and Singer with a form factor $f(x)$ [35] [36]:

$$f(x) = \frac{1}{1 - R_\rho x} + \frac{C\alpha_{K^*}}{1 - R_{K^*}x} \left[\frac{4}{3} - \frac{1}{(1 - R_\rho x)} - \frac{1}{9(1 - R_\omega x)} - \frac{2}{9(1 - R_\phi x)} \right] \quad (1.3)$$

Here $x = (M_{ee}/M_{K^0})^2$ and $R_X = (M_{K^0}/M_X)^2$. Table 1.3 lists values of R_X . C is equal to 2.5, and is a combination of known coupling constants. The variable α_{K^*} measures the relative strength of the two terms in the form factor. The first term represents




Figure 1.3: $K_L \rightarrow e^+e^-\gamma\gamma$. The other tree-level Feynmann diagrams have the radiated photon coming from the positron and exchange the radiated and direct photons.

the pseudoscalar-pseudoscalar process in Figure 1.4a, while the second term represents the vector-vector process in Figure 1.4b. The value of α_{K^*} has been measured in $K_L \rightarrow e^+e^-\gamma$ decays to be -0.28 ± 0.08 [37]. NA48 recently measured $-0.36 \pm 0.06_{stat.} \pm 0.02_{sys.}$ [38]. Measuring the M_{ee} spectrum in $K_L \rightarrow e^+e^-\gamma\gamma$ could provide a check of the form factor and α_{K^*} .




Figure 1.4: $K_L\gamma\gamma$ diagrams. (a) is the pseudoscalar-pseudoscalar transition, while (b) is the vector-vector transition.

Using this form factor, Greenlee predicted the BR of $K_L \rightarrow e^+e^-\gamma\gamma$ as a fraction of the BR of $K_L \rightarrow \gamma\gamma$ [35]. Using the Particle Data Group (PDG) fit [37], $\text{BR}(K_L \rightarrow \gamma\gamma) = (5.92 \pm 0.15) \times 10^{-4}$. Multiplying by Greenlee’s fraction, $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV})$ is expected to be $(6.02 \pm 0.15) \times 10^{-7}$. The uncertainty comes only from uncertainty in $\text{BR}(K_L \rightarrow \gamma\gamma)$.

Although the radiative-Dalitz decay is expected to be the dominant channel for $K_L \rightarrow e^+e^-\gamma\gamma$, “direct-emission” decay might make a significant contribution. Direct-emission $K_L \rightarrow e^+e^-\gamma\gamma$ would proceed through a $K_L\gamma\gamma\gamma^*$ vertex. This vertex is unlikely to be observed through on-shell photons, because $K_L \rightarrow \gamma\gamma\gamma$ is very highly suppressed by

Table 1.3: Values of R_X . All masses are from the PDG [37]. Reference [35] gives R_{K^*} as 0.311, which would be true for $K^*(892)^\pm$ mesons, which have mass $891.7 \text{ MeV}/c^2$.

X	meson	$M_X(\text{MeV}/c^2)$	R_X
ρ	$\rho(770)$	770.0	0.418
ω	$\omega(782)$	781.9	0.405
K^*	$K^*(892)^0$	896.1	0.308*
ϕ	$\phi(1020)$	1019.4	0.283

the requirement that each $\gamma\gamma$ pair have two units of angular momentum. Zero angular momentum for a $\gamma\gamma$ pair is forbidden by gauge invariance, just as $0 \rightarrow 0$ radiative transitions are forbidden. Angular momentum 1 for a $\gamma\gamma$ pair is not allowed by Bose statistics. As a result, the BR of $K_L \rightarrow \gamma\gamma\gamma$ is expected to be $\sim 3 \times 10^{-19}$ [39]. However, if one of the photons is massive, then the above restrictions apply only to the on-shell $\gamma\gamma$ pair. A naive estimate for the BR of $K_L \rightarrow \gamma\gamma\gamma^*$ would be the BR of $K_L \rightarrow \gamma\gamma$ times α_{EM} , or 4×10^{-6} . The direct portion of the BR of $K_L \rightarrow e^+e^-\gamma\gamma$ would then be suppressed by another factor of α_{EM} , giving a BR contribution of about 0.3×10^{-7} .

Although the naive BR is substantial, various models suggest direct emission is suppressed. ChPT estimates a suppression of at least 10^{-3} due to $\mathcal{O}(p^8)$ effects or 10^{-7} in a CP violating channel [40]. Another possible direct emission process would be $K_L \rightarrow K_S\gamma^*$, then $K_S \rightarrow \gamma\gamma$ and $\gamma^* \rightarrow e^+e^-$. $K_L \rightarrow \pi^0\pi^0e^+e^-$ is expected through a similar charge-radius radiation. By scaling a prediction for $\text{BR}(K_L \rightarrow \pi^0\pi^0e^+e^-)$ [41], one gets a contribution of only about 10^{-15} to direct emission.

1.3.3 Backgrounds

The same backgrounds plague $K_L \rightarrow e^+e^-\gamma\gamma$ as $K_L \rightarrow \pi^0e^+e^-$ (see section 1.2.4). Of course, $K_L \rightarrow e^+e^-\gamma\gamma$ is a background to $K_L \rightarrow \pi^0e^+e^-$, while the opposite is not true. Also, the backgrounds with charged pions, K_{e3} and $K_L \rightarrow \pi^+\pi^-\pi^0$, are very strongly suppressed by kinematic cuts and electron identification (E/p and TRD). Differences in the background spectrum from $K_L \rightarrow \pi^0e^+e^-$ are discussed below.

- $K \rightarrow \pi^0 \pi_{Dalitz}^0$: The angle between any electron and any photon in the kaon center-of-mass is largely uncorrelated in this mode. In $K_L \rightarrow e^+ e^- \gamma \gamma$, an electron tends to point in the same direction as the photon it radiated (internally). This correlation can be used to reject $K \rightarrow \pi^0 \pi_{Dalitz}^0$ background. However, using M_{ee} to reject this background, as in the $K_L \rightarrow \pi^0 e^+ e^-$ analysis, would reject too much $K_L \rightarrow e^+ e^- \gamma \gamma$ signal.
- $K_L \rightarrow e^+ e^- \gamma$: When the extra photon comes from an electron radiating in matter, the angle between the radiating electron and its photon tends to be even smaller than in $K_L \rightarrow e^+ e^- \gamma \gamma$ decays. Thus, selecting a cut on this angle is a matter of degree and must be done with care. $K_L \rightarrow e^+ e^- \gamma$ plus an accidental photon is a small but largely unavoidable background.

1.4 Method Overview

This thesis describes how I study these decay modes. After a description of the detector apparatus and data analysis methods, I discuss my measurement of the K_L flux using $K \rightarrow \pi^0 \pi_{Dalitz}^0$ decays. While a flux estimate is needed for any branching ratio measurement, studying the relatively plentiful $K \rightarrow \pi^0 \pi_{Dalitz}^0$ decays also allows evaluation of data quality.

Next, I discuss my study of the $K_L \rightarrow e^+ e^- \gamma \gamma$ decay mode. First, while keeping the analysis similar to that used for $K_L \rightarrow \pi^0 e^+ e^-$, I measure $\text{BR}(K_L \rightarrow e^+ e^- \gamma \gamma)$ and α_{K^*} as quantities of interest in their own right. Second, I consider how well we understand the $K_L \rightarrow e^+ e^- \gamma \gamma$ decay, as expressed by the Monte Carlo simulation, as a potential background to $K_L \rightarrow \pi^0 e^+ e^-$.

Finally I move on to searching for $K_L \rightarrow \pi^0 e^+ e^-$. The essential step in this search is to make a background estimate. This is done by simulating the primary background ($K_L \rightarrow e^+ e^- \gamma \gamma$) and forming a distribution in some variable using the simulation.

The simulated distribution is then fit to the data, *without* fitting in any region of the distribution where signal might be accepted (blind analysis). The fit distribution is used to estimate the expected number of background events. Cuts on other variables are selected that may improve the BR measurement. The fit is repeated with the new cuts, and a new background estimate is made. Using the background estimate and the number of data events that pass all cuts, I measure a confidence interval for $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$.

Chapter 2

KTeV Experimental Setup

This chapter describes the “hardware,” or physical apparatus, of the KTeV experiment. First, the process of creating kaon beams is described. Second, the detector itself is described, with descriptions for detector components used for hermeticity, charged-particle tracking, calorimetry, and electron/pion discrimination.

2.1 Kaon Beams

The KTeV E799-ii experiment requires a pure, high-intensity beam of neutral kaons. A narrow cross section is helpful, but not a primary requirement. As it turns out, neither is a narrow range of kaon momenta, because the detector can accept kaons that differ in momentum by an order of magnitude.

2.1.1 Proton Beam

The process of generating our kaon beam begins with the Fermi National Accelerator Laboratory (Fermilab) Tevatron. The Tevatron accelerates protons in a cyclotron to energies of 800 GeV. When fixed target experiments are taking data, the Tevatron then distributes proton beam to the various fixed targets.

The protons arrive at the targets in a specified time structure. On the nano-scale, protons arrive in 1–2 ns long pulses every 19 ns. Each 19 ns long period is known as a *RF bucket*. These buckets arrive at 53 MHz for about 19 seconds. Each 19 second period,

known as a *spill*, is followed by an off-spill period of 41 seconds when no protons arrive. (Spills were extended to 23 seconds of the 60 second cycle late in E799 data taking.) Over an entire spill 2.5×10^{12} to 5.0×10^{12} protons arrive at KTeV's target in the Neutrino Muon beam line.

2.1.2 Target

When protons reach KTeV, they are focussed onto the target (see Figure 2.1). The center of the target forms the origin of the KTeV coordinate system. The positive Z axis extends from the target towards the detector; $+Z$ is referred to as “downstream” while $-Z$ is “upstream.” The positive Y direction points vertically up, and the positive X direction is defined in the conventional way ($\hat{y} \times \hat{z} = \hat{x}$).

The proton beam strikes the target parallel to Z in the X plane, but with an angle of 4.8mrad below the $+Z$ direction in the Y plane. This angle is to inhibit neutrons (which tend to be produced more forward than kaons) entering the detector. The proton beam is less than $250 \mu\text{m}$ across in the transverse dimensions. The target is made of BeO, $\sim 30\text{cm}$ long in the Z direction, and 3 mm by 3 mm in X and Y . Thirty centimeters of BeO is ~ 1.1 interaction lengths, which optimizes kaon production efficiency. BeO is used (instead of, say, copper or tungsten) because of its resilience to the thermal stresses of the proton beam environment.




Figure 2.1: Secondary beam, plan view. X and Z scales are different. Dashed lines indicate nominal beam position. X positions of edges that do not impinge on the beams are represented arbitrarily. Elements which do not affect the beams in E799 are not included.

2.1.3 Sweeping and Collimating

Proton interactions in the target produce a zoo of particle types, of which only the K_L s are wanted. “Sweeper” magnetic fields are used to deflect charged particles out of the beam. The first sweeper magnet is the Target Sweeper, which operates between about 0.6 and 4.4 meters in Z , and gives protons a momentum kick of 475 MeV/c in the $+X$ direction. Downstream of the Target Sweeper is the Primary Proton Dump, which absorbs primary-beam protons. It is located between 7.2 and 11.8 m in Z , offset in the $-Y$ direction, water cooled, and made of copper. This offset allows neutral particles to travel unimpeded above the Dump, while the Target Sweeper deflects protons into the Dump. That is followed by another magnet, μ sweep1, between 12.3 and 17.8 m in Z , which gives positive particles a 3806 MeV/c kick in the $+X$ direction. Photons from the target are absorbed by a layer of lead 3 inches ($14 X_0$) thick at $Z = 19$ m, the “Pb Absorber.”

Just downstream of the Pb Absorber is the Primary Collimator. Made of brass, it has two tapered rectangular cylinders cut through it that begin to shape the kaon beams. Each cylinder points back towards the target, as shown in Figure 2.1. The center lines of each cylinder lie in the X - Z plane and make an angle of 0.8 mrad with the Z axis in the X - Z plane. The transverse dimensions of the cylinder were made larger for the “Summer” data taking period of E799 than for the “Winter” period: from 1.18×1.29 cm at the downstream face in Winter to 1.62×1.73 cm in Summer. This change is discussed further in section 4.2. The reason for having two kaon beams is related to KTeV’s study of ϵ'/ϵ , but the split beam has little effect on E799’s rare decay searches.

After the primary collimator, the beams begin to travel in a vacuum chamber, contained at the upstream end ($Z = 21.8$ m) by a 5 mil titanium window. Another sweeper magnet, μ sweep2, operates between 21.9 and 27.7 meters in Z . It deflects

positive particles in the $+X$ direction, with a field strength of 3135 MeV/c for Winter and 1854 MeV/c for Summer. This sweeper removes particles created by interactions in the Pb Absorber, Primary Collimator, and vacuum window. Downstream of μ sweep2 is the “Spin Rotator Dipole,” operating between 30.5 and 36.5 meters in Z with its field aligned with the X axis. This magnet is used to change the polarization of neutral hyperons (Ξ^0 and Λ^0), enhancing KTeV’s capability to search for rare hyperon decays. As kaons are spinless, this magnet has no effect on kaon studies. Its momentum kick was 2407 MeV/c for Winter and 1180 MeV/c for Summer. When in use, the Slab Collimator is downstream of the Spin Rotator. The Slab Collimator was not used during Summer. It is a wedge of stainless steel, located between 38.8 m and 40.8 m in Z . It is tapered towards the target in the X - Z plane but is not tapered vertically. It prevents particles from crossing from one beam to the other.

Downstream of the Slab Collimator are the Beam Stops, a pair of iron blocks starting at $Z = 46.4$ m, with a combined thickness of 5.95 m. The Beam Stops can be moved to block the beams. When the Beam Stops are in place, personnel can access the detector for short periods to make repairs to the detector. Alternatively, with the Beam Stops in place and suitable adjustments to the sweeper magnets, beams of muons can be sent to the detector for calibration purposes.

There then follows a long (~ 32 m) vacuum region. It is ended by the Defining Collimator at $Z = 85$ – 88 m. It has much the same geometry as the Primary Collimator, but is made of iron. The Defining Collimator determines the transverse edges of the kaon beams in the detector. Like the Primary Collimator, its apertures were larger for Summer than for Winter. Before the change, they were 4.4×4.4 cm square at the downstream face. After, they were 5.2×5.2 cm square.

The Defining Collimator is followed by the Final Sweeper magnet. This sweeper removes charged particles from decays upstream of the defining collimator and from interactions therein. The Final Sweeper extends from 90 to 93 m, and gives an integrated

transverse momentum kick of 1.1 GeV/c. The vacuum in this region is reduced to two 8.9 cm square steel pipes enclosing the beams, the downstream ends of which are limiting apertures for the beams.

2.1.4 Beam Description

At this point, the beams themselves are composed mostly of neutrons and long-lived kaons. Neutrons are difficult to avoid in an apparatus like this, and they outnumber the K_L s by about 3:1. However, the neutrons' long lifetime inhibits their decaying in the detector, so they tend to only appear as neutral, accidental activity. There are a small number of other particles present. These include photons, but also neutral hadrons with short lifetimes that have not yet decayed upstream of $Z = 93$ m: K_S ($c\tau = 2.67$ cm), Λ^0 ($c\tau = 7.89$ cm), and Ξ^0 ($c\tau = 8.71$ cm). The total rate from neutral hadrons that decay (or, like neutrons sometimes do, deposit energy in the detector) is 25 to 50 MHz.




Figure 2.2: Kaon momenta, as reconstructed in $K_L \rightarrow \pi^+\pi^-\pi^0$ decays.

The kaon beams are square in cross section. Their size is 0.50 mrad by 0.50 mrad for the Winter data, and 0.59 mrad by 0.59 mrad for the Summer data. The beam

centers are separated by 1.6 mrad in the X - Z plane. The momentum distribution of kaons is shown in Figure 2.2. The kaons that decay in a 65 m long vacuum region downstream of the Final Sweeper are used for the experiment.

This vacuum region is kept at a pressure of 1.0×10^{-6} Torr. Its downstream end at $Z = 158.89$ m is sealed with a window made of Kevlar laminated with Mylar. It is circular, with a radius of 0.90 m, and is 0.0156% of a radiation length thick in Z .

2.2 Detector

An overview of the detector is shown in Figure 2.3. Beginning upstream, there is the decay region and ring vetoes. Next comes the charged-particle spectrometer and its accompanying vetoes. They are followed by the TRD system (electron/pion discrimination), which is in turn followed by a trigger counter. The Cesium-Iodide calorimeter comes next. Downstream of that are more vetoes, including the muon counters.

2.2.1 Spectrometer

The purpose of the spectrometer is to measure the momentum and position of charged particles, as part of reconstructing kaon decays. This is done with a set of four drift chambers (DCs) and one analysis magnet. The location of these elements is shown in Figure 2.4. In order to minimize multiple scattering by particles in air, plastic bags filled with Helium are placed between each drift chamber.

2.2.1.1 Analysis Magnet

The analysis magnet is used to reconstruct the momentum of charged particles. It is an electromagnetic dipole that produces a vertical field of ~ 2000 gauss centered around $Z = 170$ m. The field strength can be mapped using a Hall effect probe moved around inside the magnet; this was done before KTeV starting running to obtain the




Figure 2.3: KTeV detector. X and Y scales are blown up, relative to Z scale.




Figure 2.4: Spectrometer elements. X and Z scales differ. The plan view is shown, but the elevation view differs mainly in that no deflection occurs in Y in the magnet.

field strength within about $\pm 1\%$ before other means could be used (sections 2.2.1.3 and 2.2.2.3). The field direction was flipped between up and down daily. A relativistic, charged particle crossing this field receives a transverse momentum “kick” of $205 \text{ MeV}/c$ in the $\pm X$ direction. By measuring the change in angle of the trajectory of the particle in the X - Z plane, the momentum of the particle can be easily calculated.

2.2.1.2 Drift Chambers

The drift chambers are basically volumes of gas across which anode and cathode wires apply an electric potential. Charged particles pass through the gas, ionizing gas atoms and freeing electrons. The freed electrons are accelerated by the electric field towards the nearest anode wire. Gaining energy, the electrons ionize yet more gas atoms, freeing more electrons. The number of freed electrons rapidly increases, forming an avalanche. When the electron avalanche reaches an anode wire, it creates a detectable current. From the position of the current-bearing wires and the time that the current arrives, the position where the particle passed through the drift chamber can be reconstructed.

There are two types of drift chamber wires in KTeV: the field (cathode) wires, made of $25 \mu\text{m}$ diameter gold-plated tungsten; and the sense (anode) wires, made of $100 \mu\text{m}$ diameter gold-plated aluminum. The voltage difference between the sense and field wires was between 2450 and 2600 V during E799. The wires are strung in parallel and organized into “plane pairs;” so called because each contains two planes of sense wires. The sense wire planes within each pair are offset, removing left/right ambiguity in which cell a track comes from. The wire structure within each plane pair is shown in Figure 2.5.

Each drift chamber contains two plane pairs. The upstream plane pair has wires parallel to the Y axis, while the downstream plane pair has wires parallel to the X axis. The number of wires in each plane is listed in Table 2.1. A pair of Mylar windows

contain the gas for each drift chamber. Each window has a plane of field wires next to it, which ensure that the field line contours for the planes near the windows are similar to those away from the windows. The gas itself is equal parts argon and ethane, with 0.5% to 1.0% isopropyl alcohol. The electron drift speed is of order $50 \mu\text{m}/\text{ns}$, and drift times are typically less than 200 ns.




Figure 2.5: DC plane pair schematic. Short segment of plane pair, showing upstream and downstream planes, and hexagonal geometry of wire cells.

Each sense wire is connected to an electronic amplifier. The amplifiers are located in “pre-amp cards,” electronics modules mounted on the DC frames. Each pre-amp card amplifies signals from 16 adjacent wires in a given sense-wire plane, and sends its output on a 17-pair (16 differential ECL pairs plus a “ground” pair) ribbon cable to a LeCroy 4413 discriminator module in an adjacent electronics station. The number of pre-amp cards and cables is listed in Table 2.1. The discriminator produces a 40 ns long pulse for each wire when there is a signal above threshold. This logic pulse then goes to various trigger logic circuits and to a LeCroy 3377 TDC (Time to Digital Converter), operating with 500 picosecond resolution and in common stop mode. The common stop is provided by trigger logic (L1TRIG, see chapter 3) for potentially good events. The stop is timed such that the DC signal arrives at the TDC between 115 and 350 ns before the stop. These TDCs can record multiple DC hits on a single wire for a single event,

Table 2.1: Wires and cables per DC plane. Listed by increasing Z . ‘Primed’ planes refer to downstream plane of plane-pair.

DC View	Wires	Cables
1Y	101	7
1Y’	101	7
1X	101	7
1X’	101	7
2Y	112	7
2Y’	112	7
2X	128	8
2X’	128	8
3Y	128	8
3Y’	128	8
3X	136	9
3X’	136	9
4Y	140	9
4Y’	140	9
4X	140	9
4X’	140	9
total	1972	128

with a multiple-pulse resolution limited by the discriminator width (40 ns). The times measured by the 3377s can be used to reconstruct the positions of tracks as they pass through each plane pair. This reconstruction is discussed further in chapter 5.

2.2.1.3 DC Calibration

As mentioned in section 2.1.3, KTeV’s neutral beams can be changed into muon beams. Using the beam stops, neutral particles are filtered from the beam. The sweeping magnets are adjusted to allow penetrating charged particles (primarily muons) to enter the detector. The analysis magnet is turned off, so the muon tracks are straight in the spectrometer. These straight tracks can then be used to determine the mapping of TDC time to track position (time-to-distance maps) for each cell, as well as the relative positions of the planes. Such “muon runs” were performed every few days during E799.

Another calibration, made using normal kaon data, uses the ability of the calorimeter (section 2.2.2) to measure the energy of electrons. Because electrons deposit all of their energy in the calorimeter, their ratio of measured energy to measured momentum (E/p) should be close to 1. The momentum measurement is more sensitive to the assumed analysis magnet momentum kick, so the kick is estimated such that $E/p \cong 1$. This kick calibration is necessary because the field is changed as often as daily (see section 2.2.1.1), and precise magnet currents may not be reproduced with each cycle.

2.2.2 Calorimeter

The purpose of the calorimeter is to measure the position and quantity of energy deposited in it by kaon decay products. It is the only tool KTeV has for measuring photons. Also, even though charged particles are seen by the spectrometer, the calorimeter is crucial in completing charged particle observations. The spectrometer separately measures X and Y positions, but the calorimeter measures both together, allowing correlation of spectrometer data (see Track-Cluster Matching, section 5.2.2.2).

The structure of the calorimeter is shown in Figure 2.6, a stack of CsI crystals, each separately instrumented and read out. The upstream face is at $Z = 186.0$ m.




Figure 2.6: CsI calorimeter.

2.2.2.1 Crystals

The calorimeter uses 3100 scintillating crystals made of pure Cesium-Iodide, each 50 cm long. Of these, 2232 are “small” crystals, measuring 2.5 cm wide and 2.5 cm high. Around these are arranged 868 “large” crystals, measuring 5.0 cm wide and 5.0 cm high. The small crystals are removed in regions where the kaon beams would hit, creating two beam holes. The scintillation light from a single particle typically makes a pulse 8 to 10 ns long with wavelength 305 nm, along with a longer wavelength, slower response over some microseconds. Each crystal is individually wrapped with reflective and dark materials to make each light tight and to provide uniform light transmission efficiency over its length. The total light yield is at least 10 photo-electrons per MeV deposited per crystal [42]. Virtually all energy from incident electromagnetically-interacting particles

(electrons and photons) is deposited in the calorimeter, because 50 cm of CsI provides 27 radiation lengths. Pions and muons typically deposit only minimal ionizing energy in this quantity of material (1.36 nuclear interaction lengths), about 280 MeV for tracks parallel to Z .

2.2.2.2 Calorimeter Instrumentation

Each CsI crystal is instrumented by a PMT (Photo-Multiplier Tube) and a “DPMT” (Digital PMT base). The PMTs are 1.5 inch Hamamatsu R5330s and 0.75 inch Hamamatsu R5364s for large and small crystals, respectively. The optical link from the crystal to the tube (a transparent RTV rubber cookie) includes a filter to remove the “slow” scintillation light. The tubes operate at 1200 V with a typical gain of 5000. The dynode outputs of the PMTs are used for trigger purposes, while the anode outputs go to the DPMTs.

A DPMT is a circuit board device which digitizes the signal from a PMT and stores the result for readout. The first stage in the DPMT is the custom QIE (Charge into Energy) chip, which integrates the current output of the PMT. The current is divided by powers of two (2, 4, 8, 16, 32, 64, and 128) into smaller currents. Each smaller current charges a different capacitor. Voltage comparators select which capacitor has the highest, non-overflow voltage. The QIE outputs the voltage of this capacitor and a number corresponding to the capacitor (the “range”). There are four such circuits in the QIE, and each is charged in round-robin fashion every bucket. The QIE also outputs a number (the “cap ID”) corresponding to which circuit is active.

The voltage output of the QIE is the input to a FADC (Flash Analog to Digital Converter) on the DPMT. Every bucket, the FADC converts this voltage into an 8-bit mantissa. The mantissa, the range, and the cap ID are combined to make a data word. This word is input to the DBC (Driver-Buffer-Clock), another custom chip. Up to 32 of such words are buffered in the DBC in FIFOs and read out in response to signals from

the trigger logic.

2.2.2.3 Calorimeter Calibration

An important tool in calibrating the calorimeter is a laser light distribution system. A single laser source distributes light pulses to all crystals via optical fibers. Fibers also deliver light to several PIN diodes with highly linear response that determine the total light in each pulse. The response of each crystal-channel can then be compared to the PIN diode. A major component of each channel's response is the capacitance of the capacitor used for that pulse. This is measured by varying the light pulse size to see the relationship of ADC counts to photo-electrons in each range. In practice, this relationship is assumed to be linear, so that only a slope and offset need be calculated for each range during calibration for later use in event reconstruction. Enough pulses of different sizes must be provided to determine this relationship for all of the 32 capacitors in the QIE. This is done during "laser scans," periods when the detector only collects this calibration data. Laser scans were performed every 2.2 days, on average, during E799.

The other major calibration tool was the requirement that $E/p = 1$ for the electrons in $K_L \rightarrow \pi e \nu$ decays. This is used to determine the relationship between ADC counts and energy deposited in a crystal. This was done before the analysis magnet variations were determined using E/p (section 2.2.1.3), when the magnet strength was known from direct field measurements.

2.2.3 Transition Radiation Detectors

The purpose of the Transition Radiation Detectors (TRDs) is to discriminate between electrons and pions, which is important in rejecting $K_L \rightarrow \pi e \nu$ background. They use the transition radiation produced when charged particles cross a boundary between materials with different indices of refraction. TR energy is inversely proportional to

the mass of the charged particle. For the relativistic particles found in KTeV (tens of GeV), the electron TR takes the form of x rays (tens of KeV). These x rays can be detected when they ionize the gas in a Multi-Wire Proportional Chamber (MWPC). More information on the KTeV TRD system is available in [43].

2.2.3.1 TRD Chambers

TRD components can be divided into the radiator material and the x-ray detector.

Figure 2.7 shows a section of a TRD chamber.




Figure 2.7: TRD cross section, plan view. Scale is approximate, except for wire spacing and as noted.

For charged particles with tens of GeV of energy, there is $\sim 1\%$ probability of useful TR per interface crossed by the particle [37]. For this reason, the radiation material in KTeV is a mat of polypropylene fiber. A low-Z material, polypropylene, is used to minimize absorption of x rays in the radiator. Dry gas flows through the radiator volume to minimize contamination.

The x rays are detected when they ionize molecules in a gas and the freed electrons are collected to measure the x-ray energy. The ionization in the KTeV TRDs is enhanced by having a gas volume downstream of the radiator with a high x-ray cross section (80% Xenon and 20% CO_2). The MWPC electric fields cause the ionized electrons to

avalanche, and the avalanches are collected by the MWPC sense wires. The CO_2 is added to the Xenon to speed electron drift. One technical problem with the Xenon volume is the tendency for bulges to form in the windows containing it, so that the gain varies over the face of the detector. The bulges are caused by the gas pressure stretching the window and, towards the bottom, by gravity pooling the dense gas, and can even change in time as the barometric pressure changes. KTeV corrects for these deformations by surrounding the conversion gas with buffer gas volumes of 80% C_2F_6 and CO_2 , which has low x-ray cross section but high density. The buffer thickness is kept small to minimize x-ray absorption.

Next, the electrons are gathered by the MWPC. The inner gas windows are made of aluminized mylar which is kept at a voltage lower than the cathode wires, so that drifting electrons produced between the cathode planes and windows will make their way to the anode wires. Cathode wires are at +200 V, the windows are at +250 V, and the anode wires are at about +2400 V. Electron drift time is reduced by using MWPCs, which have two directions of electron drift, and by using two MWPCs that filled the conversion gas volume. Charge collection time is ~ 225 ns. The anode wires connect to pre-amplification cards mounted on the chamber frames; the preamps send their outputs to postamps in electronics racks near the chambers. The postamp outputs go to 10-bit LeCroy ADCs in the counting room, operating with a 300 ns gate and 0.25 pC per count resolution.

See Figure 2.8 for the transverse profile and position of the TRD chambers. Each TRD chamber has a 2.1 m square window. Data is collected only from the central 1.8 m of wires. To reduce interactions in the beam regions, 15 cm square holes are cut in the radiator mat. In the same regions, the anode wires are deadened by thickening them with electroplating. To save electronics overhead, the anode wires in a plane, upstream or downstream, are ganged together. They are ganged in groups of two for the central 64 cm of each chamber and in groups of four for the outer wings.

There are eight such TRD chambers in the TRD system. The upstream-most of them is at $Z = 181.1$ m and the remainder follow at 31 cm intervals. The downstream anode plane of the last one is at $Z = 183.4$ m.




Figure 2.8: TRD chamber, beam view.

2.2.3.2 TRD Calibration

In order to track the gain of each chamber, a ^{53}Fe 5.9 KeV x-ray source is placed in a corner of each TRD chamber. The high voltage is adjusted to keep the gain constant during the experiment. The next step in calibration is determining wire-by-wire gain. This is done using the minimum-ionizing peak from pions in $K_L \rightarrow \pi e \nu$ decays, with the pions identified by requiring E/p significantly less than 1. The last step is to determine $\phi(x; Y)$, the probability that particle Y produces x ADC counts. This is done, again, with a $K_L \rightarrow \pi e \nu$ sample of high purity where E/p separates pions from electrons. How $\phi(x; Y)$ is used to reject pions is described in chapter 5.

2.2.4 Vetoes

The purpose of the vetoes is to detect particles in a decay which are not observed by the spectrometer-calorimeter system. Such decays with missing particles can be

a background to signal decay modes, such as when a kaon decay to $\pi^0\pi_{Dalitz}^0$ with a missing photon mimics $K_L \rightarrow \pi^0 e^+ e^-$ or $K_L \rightarrow e^+ e^- \gamma\gamma$. “Veto” or “anti” detectors are a means of rejecting such events. Figure 2.9 shows the location of veto detectors.

2.2.4.1 Ring Vetoes

One class of KTeV vetoes is intended to detect particles which leave the fiducial region at high angle. These are the Ring Counters (RCs), Spectrometer Antis (SAs), and the Cesium-Iodide Anti (CIA). These are all ring shaped detectors, with large holes in their centers. Each is composed of alternating sheets of lead and 2.5 mm thick scintillator, stacked in Z . The detectors are segmented azimuthally into modules. Module edges are angled to provide overlap, preventing particles from escaping down cracks between modules. The scintillation light from a single module is fed to a PMT. Each PMT signal is discriminated to make a logic level available to the trigger system. The PMT signals are also digitized by 10-bit LeCroy 4300 ADC modules.

The RCs are located in the vacuum decay region. The two farthest upstream, “RC6” and “RC7,” are somewhat smaller than the downstream three: “RC8,” “RC9,” and “RC10.” The outer edge of a RC is circular, located flush with the vacuum vessel walls. The inner aperture is square. There are 16 modules per RC, with boundaries that point towards the Z axis. The 16 upstream-most lead layers are 0.5 radiation lengths thick each and the eight downstream lead layers are 1.0 radiation length thick, for a total of $16 X_0$ per RC.




Figure 2.9: Veto detectors, plan view. X and Z scales differ.

Downstream of the RCs are the SAs and the CIA. “SA2” surrounds DC2, “SA3” surrounds DC3, “SA4” surrounds DC4, while the CIA is next to the CsI calorimeter. The SAs and CIA have similar construction to the RCs, except that their modules are rectangular. Thus, the inner and outer edges of the SAs and CIA are rectangles. Like the RCs, they have $16 X_0$ of lead, but segmented into 32 layers of 0.5 radiation lengths each.

2.2.4.2 Downstream Vetoes

Just downstream of the calorimeter is a set of detectors for detecting decay products which pass through the spectrometer-calorimeter system. Such particles can travel in or near the neutral beam or can pass through the calorimeter leaving only the minimum ionization energy (muons and charged pions).

The Collar Antis (CAs) surround the beam holes on the upstream face of the calorimeter. Their purpose is to reject events with electromagnetic particles that hit the calorimeter close to these holes. Because much of the shower energy can be lost out of the beam holes, the energy and position is difficult to measure for such particles. The CAs are 1.5 cm-wide square rings. Their inner edges are flush with the calorimeter’s 15 cm-square beam holes. Each has three layers of tungsten and scintillator. The total thickness of the tungsten is $9.7 X_0$. Each CA has four modules (top, bottom, left, right) with their own PMT instrumentation. Signals from each PMT are discriminated and sent to ADCs, just as for the ring vetoes.

Downstream of the calorimeter is a 15 cm (0.88 nuclear interaction lengths) thick lead wall. Its purposes are to stop the tails of any showers escaping the calorimeter and to increase the probability that hadrons passing through the calorimeter will shower. These hadron showers are detected by the Hadron Anti (HA), located on the downstream face of the lead wall. The HA has 28 scintillator paddles, 14 above the beams and 14 below the beams. Their long axes are vertical. Rectangular holes in both the wall and

the HA allow the neutral beams to pass through this detector. The hole is 60 cm wide and 30 cm high in the wall and 64 cm wide and 34 cm high in the HA. HA paddle signals are sent to ADCs. An analog sum of HA signals is made for trigger logic.

Downstream of the HA, the beams pass through another 60×30 cm hole in a 1 m thick wall of steel. This steel protects the HA from backslash caused when the beams hit the Back Anti (BA). The BA is a $30 X_0$ (about one nuclear interaction length) thick lead and scintillator sandwich detector. It is 60 cm wide and 30 cm high, and sits squarely in the path of the beams. It is intended to reject events where a particle is lost down a calorimeter hole, but beam neutrons cause so much activity that the BA is of no use for this analysis.

The BA is followed by a 3 meter (17.9 interaction lengths) thick steel beam dump. The only particles likely to penetrate this are muons, which can be rejected by “MU2,” a plane of overlapping scintillator paddles downstream of the 3 m steel. Because scattering in the steel can deflect the muons by a large angle, MU2 is relatively large in X and Y : 3.78 m wide and 3 m high. MU2 was used for a short period early in E799, but was removed because only accidental activity occurs there for this analysis; there are no significant background modes that generate muons.

MU2 is followed by yet another meter of steel, which is followed by a scintillator hodoscope, “MU3.” This detector is used to identify particles for analyses of decay modes with muons, and is not used in this analysis.

2.2.5 Trigger Detectors

Several detectors exist to produce simple logic levels indicating that an event should receive further attention. These detectors are used either for collecting special calibration information or for reducing trigger rates in kaon data collection.

2.2.5.1 V-Bank Counters

Although the drift chambers detect charged particles, they impose a delay of ~ 200 ns, the maximum drift time, while doing so. The V-Bank counters are used to make prompt trigger decisions on the presence of charged particles. There are two V-Banks, and each is a plane of scintillator paddles. The upstream bank (known as V) has its upstream face at $Z = 183.9$ m (between TRD and calorimeter), while the downstream bank (V') is adjacent to V. They are 1.90 m square, with 14 cm square beam holes. Each has 16 paddles above the beam and 16 below, with the long axes of the paddles aligned vertically. The paddles in V and V' are arranged such that the gaps between paddles in one plane do not overlap gaps in the other plane, reducing the risk of particles slipping through the V-Banks undetected. LeCroy 4413 modules discriminate the PMT signals from each paddle. The analog multiplicity currents of the 4413s are summed and discriminated in order to determine a digital hit multiplicity for each bank. These hit multiplicities are used by the trigger.

2.2.5.2 Accidental Counters

A fair amount of activity in an event can occur in the detectors which is unrelated to the decay in the event. This “accidental” activity is mostly from other beam particle interactions, and it must be measured to accurately simulate the detector. This is done by recording detector activity when the accidental counters fire. Both accidental counters are located near the target. The “accidental muon” counter is a two-scintillator muon telescope downstream from the target, at a large angle from Z , and outside the steel pile surrounding the target region. The “90 degree target monitor” is a three-scintillator telescope that looks at the target from a 90° angle to Z . When the three counters fire in coincidence, “ACC90” fires. The accidental muon counter is more correlated with non-accidental activity, so only ACC90 is used to collect accidental events.

Chapter 3

Event Selection and Data Collection

This chapter describes the method E799 uses to select data to be stored digitally for later study. Because of the high rate of the neutral beams, a trigger system is used to reject uninteresting events. Events must be accepted by three successive levels of the trigger, each more sophisticated but slower than the last, before being saved. Full readout of the various detector systems by the data acquisition system (DAQ) does not occur until after the second stage of trigger processing.

3.1 Level 0 Trigger: Digitization

The Level Zero (L0) trigger system prepares detector signals for Level One (L1), described in the next section, but does not reject or accept events. L0 produces a set of bits known as *L1 Source* bits. The electronics for producing them is located in racks next or near to their detector elements. The bits relevant to this analysis are described below.

Each veto counter produces a L1 Source. Each counter's PMT signal in the RCs, SAs, and CIA is discriminated, and the discriminator signals are ORed together for each veto. The PMT gains and discriminator levels are set to fire when more than 500 MeV is deposited in an RC module or 400 MeV is deposited in a SA or CIA module. The L1 Sources produced by these ORs are known as RC6, RC7, RC8, RC9, RC10, SA2, SA3, SA4, and CIA.

The Collar Antis use a similar scheme. The four PMT signals from each module are discriminated and ORed together. The energy threshold for CA modules is 13 GeV. The L1 Sources are called CA_LEFT and CA_RIGHT.

The L1 Sources for the Hadron Anti are formed by performing an analog sum of PMT outputs. The L1 Source HA_AC turns on when an AC-coupled sum goes above a level corresponding to the energy deposited by about 2.5 minimum-ionizing particles. Another HA related L1 Source, HA_DC, used a threshold of 7 MIPs for the first part of E799 and used a DC-coupled sum with a 2.5 MIPs for the remainder of E799.

L1 Sources for MU2 are made by discriminating the PMT outputs with LeCroy 4413s and making a sum of the analog multiplicity outputs. Discriminating the sums with two different thresholds gives two L1 Source bits: 1MU2, meaning one paddle hit in MU2; and 2MU2, meaning two paddles hit in MU2.

As mentioned in section 2.2.5.1, an analog multiplicity is formed for the V-Bank Counters by discriminating the paddle signals and summing the analog multiplicity currents. These are then fanned out by LeCroy 428F modules to discriminators with thresholds set to discriminate on one or more, two or more, or three or more counters hit. The L1 Sources are then V0_GE1S, V0_GE2S, V0_GE3S, V1_GE1S, V1_GE2S, and V1_GE3S. V0 refers to the upstream plane (V) and V1 refers to the downstream plane (V').

To check at an early trigger stage that there is sufficient DC activity in an event, the DC-OR logic is used. The DC-ORs stretch each discriminated wire signal to 90 ns long. The drift time for a track passing through the center of a cell is about 60 ns, so at least one of the two stretched signals from a track will be on during the time interval 60 to 90 ns after the track occurs. By keeping the signal width short, the DC-ORs can be used in prompt trigger logic. An OR is then taken of groups of 16 wires: 8 from the upstream plane and 8 from the downstream plane. The number of such paddles with DC activity is then summed by plane-pair. Three plane-pairs had 14 paddles each

(DC1X, DC1Y, and DC2Y), while a fourth plane-pair had 16 paddles (DC2X). There are two L1 Sources for each plane-pair: 1 or more paddles hit and 2 or more paddles hit. These Sources are named 1DC1X, 2DC1X, 1DC1Y, 2DC1Y, 1DC2X, 2DC2X, 1DC2Y, and 2DC2Y.

Another trigger element that checks a detector element for a minimum level of activity is the ETOT (Energy TOTal) system. It estimates the total energy deposited in the calorimeter during an event by summing the PMT dynode outputs. This AC-coupled analog sum goes to four LeCroy 821 discriminator channels, each set at a different threshold. The L1 Sources are called ET_THR1, ET_THR2, ET_THR3, and ET_THR4, corresponding to nominal energy thresholds of 10 GeV, 18 GeV, 25 GeV, and 38 GeV.

There are two additional L1 Sources related to beam operations. The SPILL bit indicates that the Tevatron is in its 19-second long spill period and could be providing protons to the target. The NC_PING (neutrino ping) bit indicates that very high intensity beam is being provided to another fixed-target experiment. There are several neutrino pings at the beginning of each spill during the “fast-spill” phase, lasting less than a second. However, some of this intensity could spill over into KTeV’s beam causing very high detector activity, so NC_PING is used as a veto.

Once the L1 Source bits are formed at the detector trigger stations, they must be transmitted to the L1 Trigger system. The L1 Trigger is located in a counting room upstairs and adjacent to the detector hall, connected by feed-through holes near $Z = 188$ m. Because trigger stations can be as much as 62 m away from the feed-throughs, L1 Sources are sent to L1 on “hard-line” RF waveguides to reduce trigger processing time. The hard-lines carry signals at about 0.88 times the speed of light, while standard coaxial cables’ speed is about 0.64c.

3.2 Level 1 Trigger: Prompt Logic

The Level 1 Trigger makes a decision to accept or reject an event for every RF bucket. This decision is represented by the L1TRIG bit. L1TRIG will be on if any one of up to 32 different logical conditions is true. These conditions are known as triggers, and are divided into 16 “beam” triggers that use L1 Sources as input and 16 “calibration” triggers that are satisfied when other auxiliary inputs are on. The L1 trigger system is implemented, for the most part, with LeCroy modules mounted in NIM and CAMAC crates in five adjacent electronics racks in the KTeV counting room.

3.2.1 Processing L1 Sources

The L1 Sources associated with a single event arrive at the L1 system with a spread of ~ 80 ns. All but ~ 10 ns of this spread is removed by using cable delays on the earlier signals. These delayed signals are fanned out to TDCs and to the L1 logic. The TDCs are LeCroy 3377 modules, and are used to check the timing of all L1 Sources with respect to the RF phase for jitter or shifting.

The L1 Sources, after delay cables and fan-out, are then synchronized with each other. Each goes to a LeCroy 622 coincidence module, along with a copy of the RF signal. The global phase of the RF signals is adjusted so that the RF will come on after all L1 Sources become valid for a particular event. The leading edges of the outputs of the 622s are then synchronized. In addition, if any L1 Source is on for longer than one bucket, then the 622 output for that source will have a falling and rising edge for all such buckets. Each 622 channel fans out five “synched” L1 Source outputs. One output continues in the L1 logic chain and some or all of the remainder go to some form of monitoring. All go to FERA pipeline latches built by KTeV that are read out for all events accepted by the Level 2 trigger. All also go to scalers for general rate monitoring. The scaler inputs for the DC-ORs are taken from the L1 sources *before*

they are synchronized with RF, because the DC-ORs are always on for at least 90 ns (~ 4.7 buckets).

The scalers for the L1 Sources are LeCroy 2551 modules. Each module has 12 channels that can count up to 24 bits each. A wire patch allows pairs of channels to be joined, creating a channel with a 48-bit counter. This modification is needed for most L1 Sources because of the high rates possible in most channels; the maximum average rate that can be scaled with a single channel is $2^{24}/23 \text{ sec} = 729 \text{ KHz}$. This modification was made to 12 modules, while two modules were left unmodified for low rate Sources. The scalers are read out over a CAMAC bus several times during a spill and are cleared at the end of the spill. The results are displayed in the control room for experimenters to view, and the end of spill scaler totals are stored in a database for later use.

The synched L1 Sources that continue in the L1 system go into LeCroy 4413 discriminators. These modules do several things for the signals: give them identical widths in time, translate them from NIM levels to ECL levels, group them onto ribbon cables, and fan them out. The signals are grouped into *Source Groups* made up of 8 signals each. Each 4413 puts two Source Groups on a 16-pair ribbon cable, and fans out an identical cable. These cables serve as inputs to LeCroy 4508 programmable lookup units.

3.2.2 Logic

The 4508s have two memory lookup tables each. Each memory is addressed by eight bits and has eight bits of output. Thus each Source Group cable addresses two lookup tables that are in separate modules. The outputs of the 4508s are ANDed together and recombined to form a mask of trigger bits.

To put it another way, each Source Group can be designated S_i , where $i = 1, 2, \dots, 12$. S_i represents a mask of eight bits. The Groups are fanned out and combined

onto 12 cables:

$$(S_1, S_2), (S_1, S_2), (S_3, S_4), (S_3, S_4), \dots, (S_{11}, S_{12}), (S_{11}, S_{12}).$$

Each lookup table can perform an arbitrary functions F on its input S , with a different F_j possible for every lookup table. F_j represents a mask of eight bits. Thus the cable outputs of the 4508s can be designated:

$$\begin{aligned} & \left(F_1[S_1], F_2[S_2] \right), \left(F_3[S_1], F_4[S_2] \right), \left(F_5[S_3], F_6[S_4] \right), \left(F_7[S_3], F_8[S_4] \right), \\ & \dots \left(F_{21}[S_{11}], F_{22}[S_{12}] \right), \left(F_{23}[S_{11}], F_{24}[S_{12}] \right). \end{aligned}$$

Two stages of LeCroy 4516 modules AND together the bits on these cables; i.e. all of the bit-0s are ANDed together with each other to form a new bit-0. The two output cables of the 4516s are then

$$\begin{aligned} & \left(F_1[S_1] \cdot F_5[S_3] \cdot F_9[S_5] \cdot F_{13}[S_7] \cdot F_{17}[S_9] \cdot F_{21}[S_{11}], \right. \\ & \left. F_2[S_2] \cdot F_6[S_4] \cdot F_{10}[S_6] \cdot F_{14}[S_8] \cdot F_{18}[S_{10}] \cdot F_{22}[S_{12}] \right), \\ & \left(F_3[S_1] \cdot F_7[S_3] \cdot F_{11}[S_5] \cdot F_{15}[S_7] \cdot F_{19}[S_9] \cdot F_{23}[S_{11}], \right. \\ & \left. F_4[S_2] \cdot F_8[S_4] \cdot F_{12}[S_6] \cdot F_{16}[S_8] \cdot F_{20}[S_{10}] \cdot F_{24}[S_{12}] \right) \end{aligned}$$

These are input to a pair of LeCroy 2365 modules, which can AND bits within input cables. Bit 0 is ANDed with bit 8, bit 1 with bit 9, and so on. The output of each 2365 would then be an 8 bit mask, but these 8 bits are repeated twice for fanout. The two output cables are represented

$$\begin{aligned} & \left(F_1[S_1] \cdot F_2[S_2] \cdot F_5[S_3] \cdot F_6[S_4] \cdot F_9[S_5] \cdot F_{10}[S_6] \right. \\ & \left. \cdot F_{13}[S_7] \cdot F_{14}[S_8] \cdot F_{17}[S_9] \cdot F_{18}[S_{10}] \cdot F_{21}[S_{11}] \cdot F_{22}[S_{12}], \text{ same thing} \right), \\ & \left(F_3[S_1] \cdot F_4[S_2] \cdot F_7[S_3] \cdot F_8[S_4] \cdot F_{11}[S_5] \cdot F_{12}[S_6] \right. \\ & \left. \cdot F_{15}[S_7] \cdot F_{16}[S_8] \cdot F_{19}[S_9] \cdot F_{20}[S_{10}] \cdot F_{23}[S_{11}] \cdot F_{24}[S_{12}], \text{ same thing} \right) \end{aligned}$$

These ECL cables are translated onto NIM cables using LeCroy 4616 modules. One copy of the masks is used for monitoring, but the other copy is recombined to form the

16-bit *L1 Beam Trigger Mask*, which can be represented:

$$\begin{aligned} & \left(F_1[S_1] \cdot F_2[S_2] \cdot F_5[S_3] \cdot F_6[S_4] \cdot F_9[S_5] \cdot F_{10}[S_6] \right. \\ & \cdot F_{13}[S_7] \cdot F_{14}[S_8] \cdot F_{17}[S_9] \cdot F_{18}[S_{10}] \cdot F_{21}[S_{11}] \cdot F_{22}[S_{12}], \\ & F_3[S_1] \cdot F_4[S_2] \cdot F_7[S_3] \cdot F_8[S_4] \cdot F_{11}[S_5] \cdot F_{12}[S_6] \\ & \left. \cdot F_{15}[S_7] \cdot F_{16}[S_8] \cdot F_{19}[S_9] \cdot F_{20}[S_{10}] \cdot F_{23}[S_{11}] \cdot F_{24}[S_{12}] \right) \end{aligned}$$

Thus, all Trigger Mask bits “see” all of the L1 Source bits. Also, the logic that can be performed for each Beam Trigger Mask bit is arbitrary within a Source Group, but can only be an AND between processed Source Groups.

The data for this analysis uses the first bit in the Beam Trigger Mask, known as 2E-NCLUS. Although there were small changes in the requirements of 2E-NCLUS during E799, a typical configuration is shown below. These changes in the 2E-NCLUS definition are discussed in section 4.2.

$$\begin{aligned} 2E-NCLUS_{L1} = & \text{SPILL} \cdot \overline{\text{NC_PING}} \cdot \overline{\text{RC6}} \cdot \overline{\text{RC7}} \cdot \overline{\text{RC9}} \cdot \overline{\text{RC10}} \cdot \overline{\text{SA2}} \cdot \overline{\text{SA4}} \cdot \overline{\text{CIA}} \\ & \cdot \overline{(\text{CA_LEFT} + \text{CA_RIGHT})} \cdot \overline{\text{HA_AC}} \cdot \overline{\text{MU2}} \\ & \cdot [(\text{V0_GE2S} \cdot \text{V1_GE1S}) + (\text{V0_GE1S} \cdot \text{V1_GE2S})] \\ & \cdot [(1\text{DC1X} \cdot 1\text{DC1Y} \cdot (1\text{DC2X} + 1\text{DC2Y})) \\ & + [(1\text{DC1X} + 1\text{DC1Y}) \cdot 1\text{DC2X} \cdot 1\text{DC2Y}]] \\ & \cdot \text{ET_THR3} \end{aligned}$$

In other words, this trigger requires that a normal spill be in progress and that no veto counters are above threshold. In the trigger counter, at least two paddles must be hit in one plane and at least one paddle in the other plane. In the drift chambers, at least one 16-wire paddle must be hit in DC1X, DC1Y, DC2X, and DC2Y, although one chamber with no hits is allowed. There must be more than about 25 GeV deposited in the calorimeter. There are additional Level 2 requirements for 2E-NCLUS, described in section 3.3.5.

The L1 Beam Trigger Mask can then be prescaled by a Fermilab built CAMAC prescaler, so high rate triggers can be restricted. An OR is taken of the *Prescaled Beam Trigger Mask* bits to form the overall L1 Beam Trigger bit. The Prescaled Beam Trigger Mask also goes to latches and scalers for monitoring.

3.2.3 Calibration Triggers

Another 16-bit mask is formed in parallel to the Beam Trigger Mask, called the *Calibration Trigger Mask*. Calibration triggers are used to collect information about the detector, but not about physics events. The first 14 bits of the Calibration Trigger Mask are essentially equal to Calibration Source signal bits. One examples of a calibration source is a signal indicating that the calorimeter laser has flashed. The OR of these 14 bits is known as the *Calibration Trigger*. The last two bits of the Calibration Trigger Mask are both set equal to the Calibration Trigger. This serves as a flag later in the trigger that an event is calibration data, but it also means that the beam trigger definitions can never allow bits 14 and 15 to be on in the same event.

3.2.4 L1 Output

The OR of the L1 Beam Trigger and the Calibration Trigger is input to a LeCroy 821 discriminator. The output of the discriminator is *L1TRIG*. However, L1TRIG is restricted by the L1VETO signal going into the veto input of the 821, which prevents L1TRIG from occurring at inappropriate times. L1VETO is an OR of L1TRIG itself, the Calibration Trigger, “DYC-CTIRC Full”, and “L1 Inhibit”. DYC-CTIRC Full is an OR of the Full outputs of DYC and CTIRC modules, which are intermediate devices in the readout process, described in section 3.5.1. L1 Inhibit is a signal issued by L2 Control, described in section 3.3.5. While L1VETO is on, no new L1TRIGs are issued. L1TRIG is fanned out to various locations, including L2 processors and control logic. L1TRIG also latches the Trigger Masks, Beam and Calibration, which are then sent to

the Level 2 trigger.

A typical, average, instantaneous L1TRIG rate is 57 KHz with 4×10^{12} protons hitting the target per spill. The corresponding 2E-NCLUS trigger rate by itself is 22 KHz. These rates include dead time of about 34% from DAQ.

3.3 Level 2 Trigger: Hardware Logic

The Level 2 trigger system decides whether to accept or reject events when a L1TRIG occurs. Several *L2 Processors*, trigger subsystems implemented with custom-built or programmable electronics, begin processing an event when they get L1TRIG. This processing time can take up to several microseconds. The *L2 Control* system monitors the L2 Processor outputs with a state machine. Judging these outputs, L2 Control then issues an ABORT, to reset the detector electronics, or a L2TRIG, allowing the event to proceed to readout and Level 3.

A typical, average, instantaneous L2TRIG rate is 11 KHz with 4×10^{12} protons hitting the target per spill. The corresponding 2E-NCLUS trigger rate is 2.1 KHz. As above, these rates include dead time of about 34% from DAQ.

The L2 Processors are discussed below. All produce some data bits and two control signals, DONE and BUSY. The DONE indicates that the data bits are valid and can be used by L2 Control to decide on ABORT/L2TRIG. The BUSY comes on when the L2 Processor receives L1TRIG and goes off after an abort or after readout of the processor's data is complete.

3.3.1 Calorimeter Cluster Counting

The L2 Processor that instruments the calorimeter is a custom-built electronic subsystem known as the Hardware Cluster Counter (HCC). When electromagnetic particles shower in the calorimeter, they leave a characteristic cluster of crystals containing energy above some threshold. Figure 3.1 shows an example of clusters. The HCC counts




Figure 3.1: Example of energy in calorimeter. The shading indicates energy deposited in that crystal, in steps of 0.01–0.1 GeV, 0.1–1.0 GeV, 1.0–10.0 GeV, and more than 10.0 GeV. The circles indicate where the HCC algorithm found clusters.

the number of these clusters and outputs the result to L2 Control. The inputs for the HCC are produced by the ETOT system. There is one HCC input bit for each of the 3100 CsI channels; ETOT turns the bit on if the PMT dynode signal is above threshold when an RF clock is received. The threshold is reached when ~ 1 GeV is deposited in the crystal. All 3100 bits are stored in a FIFO. When a L1TRIG is received, the FIFO word associated with the L1TRIG is transmitted to the HCC subsystem in the counting room.

The HCC algorithm works by counting corners. Consider a rectangular grid of bits, in which there is a cluster of ON bits. Start at any point on the line bounding the cluster and follow it in clockwise direction. The corner count increments by $+1$ for every 90° right turn and -1 for every 90° left turn. Upon returning to the starting point, the corner count is $+4$. The cluster count for the grid is then the corner count divided by four. This is implemented in the HCC by having the bits from every 2×2 array in the calorimeter address a lookup table. The tables output a corner count, which can then be summed to produce a cluster count. The tables are addressed in parallel, keeping the overall processing time low. From L1TRIG to valid HCC output takes $2.5 \mu\text{sec}$. More information about the algorithm and details about the HCC design can be found in [44].

The output of the HCC is a three-bit cluster multiplicity, with a fourth bit indicating eight or more clusters. The 2E-NCLUS trigger requires that the HCC count four or more clusters.

3.3.2 DC Hit Counter

Primary instrumentation in L2 of the drift chambers is provided by the DC Hit Counter (DCHC) subsystems. Although these are described in considerably more detail in section 3.4, a brief summary is given here. There is one DCHC L2 Processor for each drift chamber plane-pair. There are two types of DCHCs: *Bananas* and *Kumquats*. The

Bananas use flash TDCs to identify DC hits associated with a particular RF bucket. Kumquats simply identify DC hits occurring within a fixed gate period. Both types output a three-bit hit count for a plane-pair. They also have outputs for other L2 Processors: masks of DC hits and ORs of DC hits in wire-paddles.

The 2E-NCLUS trigger requires that there be at least two hits in every Y-view, with one missing hit allowed in either DC1Y or DC2Y. For most of E799, at least one hit in DC2X was required (see section 4.2).

3.3.3 Track Finding

The Y-Track Finder (YTF) is a L2 Processor that looks for patterns of DC hits consistent with straight tracks in the upper and lower halves of the chambers. For input, it uses the ORs of hits from the Y-view DCHCs. These are available in paddles of 4, 8, 16, or 32 wires; the YTF uses the wider paddles towards the edges of the chambers. The YTF algorithm is implemented on a commercial programmable logic unit. It has four bits of output describing the position and quality of tracks found.

The 2E-NCLUS trigger makes no requirements on the YTF.

3.3.4 “Stiff” Tracks

The Stiff Track Trigger (STT) L2 Processor looks for patterns of DC hits consistent with very high momentum tracks. “Stiff,” because high momentum tracks will not bend much in the analysis magnet. The STT is used for hyperon physics modes in E799. For input, it uses masks of latched DC hits from the beam regions of the X-view DCHCs. Its logic is implemented on commercial modules. The output is a single flag indicating whether a stiff track was detected or not.

No requirements on the STT are made by the 2E-NCLUS trigger.

3.3.5 Control Logic

The Level 2 Control logic uses the L2 Processors to decide if an event should be rejected or passed on to the Level 3 trigger. In the latter case, it also controls readout of detector electronics. L2 Control is implemented using LeCroy modules, located in four electronics racks adjacent to the L1 Trigger.

The first stages of L2 Control prepares signals from the L2 Processors and the L1 Trigger for the L2 Control state machine. The data and done bits from the L2 Processors go to three LeCroy 2373 programmable lookup tables, one for X-view DCHCs, one for Y-view DCHCs, and one for the other processors (HCC, YTF, STT, etc.) The modules produce 16-bit mask outputs, NAMA, NAMB, and NAMC, short for Not(Abort Mask A/B/C). Each bit corresponds to a trigger condition. A bit in a mask is OFF if the processor data for that mask fails to satisfy the trigger associated with the bit. For example, bit 0 of NAMB, associated with the 2E-NCLUS trigger, would be OFF if any Y-view DCHC issued a DONE and had a hit count of zero. The other Y-view DCHC would not have to be DONE for bit 0 to go low, because the hit count of 0 means the event cannot be accepted. An AND is taken of NAMA, NAMB, and NAMC to produce another 16-bit mask, NAM.

$$\text{NAM}_i = \text{NAMA}_i \cdot \text{NAMB}_i \cdot \text{NAMC}_i$$

This is in turn ANDed with the latched Beam Trigger Mask (LBTM) from L1, and ORed with the latched Calibration Trigger Mask (LCTM) from L1 to form a mask called L2MASK.

$$\text{L2MASK}_i = \text{LCTM}_i + \text{LBTM}_i \cdot \text{NAM}_i.$$

Bit 15 of the LCTM (equal to the OR of bits 0–13 of LCTM) is used to veto the LBTM, so if there are any calibration triggers present, only they make it into L2MASK. An OR is taken of the bits in L2MASK to form the signal notL2ABORT.

This same calibration-vetoed-LBTM addresses another 2373 lookup table. The output, `Use_L2_Proc`, is a mask of bits corresponding to L2 Processors. If ON, it means that data from the processor is needed to satisfy the current trigger conditions. This is then ANDed with the mask of inverted DONE signals from the processors. An OR is taken of the resulting mask to form the signal L2WAIT;

$$\text{L2WAIT} = \sum_{i=0}^{15} \text{Use_L2_Proc}_i \cdot \overline{\text{DONE}_i}.$$

L2WAIT means that there are L2 Processors that are not finished that are needed to satisfy the L2 requirements of triggers that satisfy L1 requirements.

The signals `notL2ABORT` and L2WAIT are inputs to the L2 Control state machine. This is implemented with a 2373 lookup table, is strobed with a 10 MHz clock, and has a six-bit state. When the state machine receives L1TRIG as an input, it prepares for a new event. It turns on L1 Inhibit, to prevent more L1TRIG signals. If `notL2ABORT` goes OFF, meaning that L2 Processors have rejected all triggers accepted at L1, the state machine immediately issues an ABORT (and turns off L1 Inhibit shortly thereafter.) The L2 trigger then resets for a new L1TRIG. If `notL2ABORT` remains ON and L2WAIT goes off, it means that no L2 Processors have rejected the event and all L2 Processors needed for the event's trigger are DONE. Then, the state machine issues a L2TRIG, the readout process commences, and the event goes to Level 3.

The L2 Control system controls the readout, the transfer of data from detector electronics into RAM. It does this by issuing L2TRIG signals, which prepare the detector electronics for readout. It also sends the Read Out Control (ROC) mask to be inserted in readout data streams. The ROC mask includes a four-bit event number incremented by L2TRIG, a four-bit sparsification mask, and an eight-bit "DAQ destination" mask indicating how the event should be processed by Level 3. The state machine monitors the progress of readout with two BUSY inputs. One is the OR of all L2 Processor BUSYs, which are described above. The other is the OR of all DYC BUSYs; DYCs

are intermediate devices in the readout process, discussed in section 3.5.1. When the BUSYs go off, L2 Control issues a CLEAR signal (and turns off L1 Inhibit) to prepare the trigger for a new event.

The state machine monitors if certain inputs (including L2WAIT, the L2 Processor BUSY, and the DYC BUSY) are taking too long to arrive with a pair of time-out inputs produced by gate generators when L1TRIG occurs. There are four such time-out gates, multiplexed onto two inputs to the state machine. A single bit output from the state machine controls the multiplexing. If a time-out condition occurs, the state machine outputs an ERROR bit and stops all trigger activity. Other illegal conditions can cause this error state as well.

Two inputs to the state machine, Reset and End_of_Run, are produced by a CAMAC pulser. This allows these signals to be produced by experimenters using software in the control room. Reset clears error conditions and prepares the trigger to operate, while End_of_Run places it in a stopped state.

3.4 Level 2 Trigger: The Hit Counting System

The purpose of the DCHC system is, obviously, to count hits. A “hit” is defined as activity in a DC plane-pair produced by a single track. Hits can be *in-time*, caused by a track occurring during the RF bucket that the trigger is responding to, or *out-of-time*, caused by an earlier or later track or accidental activity. The DCHCs attempt to ignore out-of-time hits; Kumquat DCHCs do so with a simple gate, while Banana DCHCs use a more sophisticated method involving flash TDCs and memory lookups.

A simplified view of the flow of signals in the DCHC system is shown in Figure 3.2.




Figure 3.2: DCHC overview.

3.4.1 DC Signals

The first stage in transporting signals to the hit counting logic occurs at the *DC Repeater* boards. The inputs to these boards are the amplified and discriminated DC wire signals. They are picked off of the 17-pair twist-and-flat cables between the discriminator outputs and the inputs to the DC TDCs. The custom-built Repeater board circuits are fairly simple: a Motorola 100314 differential ECL receiver/repeater chip, with a 300 Ohm pull-down resistor to -5.2 V on each output trace. The Repeater board is implemented as a single circuit board mounted on an electronics rack. On one side are input and output pins and on the other side are surface-mounted components. The chips are socketed for easy replacement. Each Repeater board can handle 36 cables of input and 36 cables of output. There is one Repeater board for each DC, so each has a number of spare cable connections (except for DC4, where all Repeater connections are used.) See Table 2.1 for a summary of how many wires and cables come from each DC.

The outputs of the DC Repeaters go onto 17-pair twist-and-flat cables. These cables carry the signals from the detector hall to the counting room. They also delay the signals so that they are available to the DCHC system after L1TRIG has been issued for the corresponding event. The total length of the cables ranges from ~ 600 ns for DC4

cables to ~ 700 ns for DC1 cables.

Because the cables are so long, the signals suffer considerable attenuation. The signals are sensed and restored to normal ECL levels by the *DC Receiver* circuits. These circuits consist of a 100314 chip, a 300 Ohm pull-down resistor to -5.2 V on the negative input, a 750 Ohm pull-down resistor to -5.2 V on the positive input, with a 33 Ohm resistor and $3.3 \mu\text{H}$ inductor in series between the negative and positive leads. Both polarities of the output are pulled down with 510 Ohm resistors. These circuits are implemented in two custom-built crates of 16 modules each. 64 input cables connect to the backplane of each crate, as do the 64 output cables from each crate. This design allows Receiver modules to be swapped without any uncabing. All input and output cable connections on both crates are used; however, not all Receiver channels are used because some cables have unused wire pairs. Another function of the Receivers is to reverse the order of signals on one half of the cables. In the cables from upstream Y planes and from downstream X planes, the highest numbered wire signals are on the ribbon cable pair adjacent to the ground pair. The opposite is true for downstream Y and upstream X cables. The Receivers reverse the order for downstream Y and upstream X on the output cables. This function occurs in the cables that connect the crate backplane to the modules. The DC Receiver outputs are carried on ribbon cables into the adjacent FASTBUS electronics racks, where the DCHC subsystems are located.

3.4.2 FASTBUS Racks

The arrangement of the DCHC FASTBUS racks is shown in Figure 3.3. There are four racks, and each contains two FASTBUS crates and one CAMAC crate. Each FASTBUS crate can hold the modules for a single DCHC subsystem. As shown in the Figure, the X-view DCHCs are above the Y-view DCHCs, and the DC1 crates are on the left and the DC4 crates are on the right (as viewed from the front). The CAMAC crates hold modules for FERA (Fast Encoding and Readout ADC) readout of DCHCs.

DC Receiver output cables enter each racks from under the counting room's raised floor through a hole in the bottom of each rack. Inside each rack (and not shown in the Figure) are 6 power supplies, which provide $+5\text{ V}$, -2 V , and -5.2 V to each crate.

Because the electronics in each rack consume up to 5 KW of electricity, some attention must paid to cooling. The cooling system consists of forced air and chilled water circulation. The air is forced by a blower, which draws air from beneath the raised floor. The sides of the racks are sealed to keep air flowing through the crates. As it rises, the air passes through three heat exchangers — below, between, and above the two crates. Water is chilled in a separate area of the KTeV hall and pumped to the racks and through the heat exchangers.

Several systems monitor the FASTBUS racks. Each rack has a smoke detector, an air flow sensor, a water flow sensor, and several drip detectors (sensitive to leaks or to condensation dripping from chilled-water pipes). If any sensor detects an unsafe condition, the rack is “tripped”: water flow and electricity are shut off. There is also monitoring of the voltage and current output of the power supplies, performed by the “Mango” module in each rack. The Mangos displays these values on their front panels, and also output the values to the KTeV Slow Data Acquisition system for monitoring in the control room.




Figure 3.3: DCHC FASTBUS rack layout, viewed from front. The Summer configuration of Kumquats and Bananas is shown.

3.4.3 Auxiliary Cards: Blossoms

On the back of FASTBUS crates are slots for auxiliary cards that connect to the the main modules. Each DCHC module has such a card, known as a *Blossom* board. Most of the I/O cables for the DCHC system connect through the Blossoms, allowing the main modules to be accessed with a minimum of uncabing. The Blossom designs are fairly simple, containing only receiver/driver chips for some cables. After the DC Repeater output cables enter the racks, they connect to the Blossoms. Other cables connected to Blossoms carry inter-module signals for performing the hit count and allow FERA readout of the modules.

3.4.4 Latches: Kumquats

A Kumquat DCHC subsystem/crate contains one Kumquat module for every two DC signal cables from the DC view. A Kumquat module contains a single circuit board and spans two FASTBUS slots. Figure 3.4 is a simplified block diagram of a Kumquat module. The GATE input is produced by the L1 system whenever L1TRIG occurs. GATE lasts for 220 ns, the maximum drift time plus some time to account for differing Kumquat-input-cable delays. While GATE is on, DC signals set latches. The hit count logic counts the number of DC hits consistent with the wires with latched hits, and adds this number to hit counts that come from other Kumquat modules.

3.4.4.1 Input Latching

Thirty-two signals from adjacent DC wires are each latched by a circuit show in Figure 3.5. This circuit uses ECL (Emitter-Coupled Logic) exclusively. One-fifth of a National Semiconductor 100314 differential ECL receiver buffers each input channel. When a DC signal arrives at this circuit, the D flip-flop (1/3 of a 100331) is inactive (Q output is off). The HIT signal goes to the hit counting circuit and also goes the the




Figure 3.4: Block diagram of Kumquat module.

clock input of the flip-flop. If the GATE signal is on, then the leading edge of the HIT turns on the flip-flop output. When the flip-flop is on, the AND gate has an input that is off, and inputs from the DC wire will not change the state of HIT.




Figure 3.5: Block diagram of receiving/latching circuit

The flip-flop is reset after every event by CLEAR to prepare the circuit for the next event. FASTBUS commands can also set, reset, and read out the state of the latch to allow monitoring and testing.

3.4.4.2 Hit Counting

The latches pass a mask of 32 bits to the hit counting circuit, implemented in two Altera EPM7096 EPLDs (Erasable, Programmable, Logic Devices). The hit counting algorithm counts N continuous hits in the input mask as $N - 1$ track hits. A single, isolated hit is counted as one track hit. In order to use this algorithm correctly for an entire drift chamber, each module needs a latched hit from its neighboring modules. These data are transmitted through the Kumquat Blossoms. See Figure 3.6 for an example of hit counting. The local hit count is available on the Kumquat front panel as NIM signals for monitoring. It is also sampled several times per second and used to

light a number of LEDs on the front panel, providing a visual “temperature gauge” of DC activity.




Figure 3.6: Example of hit counting.

The hit counting circuit also receives inputs from zero, one, or two other Kumquat modules via the Blossoms. The input from each module includes a bit indicating that that hit-count cable has been connected to another Kumquat, a count-ready bit, and a 3-bit count. If an input is present, the hit-counting logic waits for the input-count-ready bit to come on before producing output. When the input readies are on, the logic adds the input counts to its own local count, and turns on its own count-ready bit. If all three bits are on in the count, it means seven or more hits. Section 3.4.6 describes how these counts produce a hit count for each DCHC subsystem.

3.4.5 Trigger TDC: Bananas

The Banana DCHC subsystems operate in a very similar manner to the Kumquats, with additional circuits for rejecting out-of-time hits. These circuits require that Banana modules use more circuit board “real estate” than Kumquat modules. Bananas span three FASTBUS slots and use two of the three slots to support circuit boards: a mother board and a daughter board. Figure 3.7 shows two views of Banana modules.

Figure 3.8 is a simplified block diagram of a Banana module. It is similar to

the Kumquat diagram (Figure 3.4) with the addition of the TDC and memory lookup blocks. Also, CLEAR in Kumquats is replaced with ABORT in Bananas.




Figure 3.7: Photograph of Banana modules. In the left module, the hit logic and flash TDC chips are visible behind the stiffener plate. The right module shows chips for all other circuits.

3.4.5.1 Input Latching

The Bananas and Kumquats use the same circuit for latching input signals, described in section 3.4.4.1. There are some differences in the inputs and outputs. The GATE input for Bananas is much longer, at about 464 ns. This is because Bananas can use very-early or very-late DC wire hits to determine that an isolated in-time hit belongs to an out-of-time pair. In this way, extending the Banana GATE can actually reduce trigger rates. Another difference is that the outputs from the Banana input latch circuit are not called HIT, but START. In normal running, an input signal will turn on START even when GATE is off. After GATE goes on, only the first input during GATE will turn on START. START then goes to a TDC.




Figure 3.8: Block diagram of Banana module.

3.4.5.2 TDC Circuit

The core of the Banana module is the flash TDC circuit. Figure 3.9 is a simplified schematic of this circuit. The circuit uses fast ECL devices driven at a clock (CLK) frequency of 625 MHz.




Figure 3.9: Block diagram of TDC circuit. Only one polarity is shown for differential ECL signals for simplicity.

We carefully designed the physical properties of the TDC circuit to allow high rate performance. The clock and many key signals are differential ECL to ensure performance. The lengths of the trace pairs of the differential signals are matched to reduce the time skew. The circuits are laid out on a $100\ \Omega$ impedance-controlled printed-circuit board. All components are surface mounted. The pull-down resistors are single piece devices, rather than multiple packages, to avoid possible coupling in the package. The 32 channels in a single board are arranged in 16 identical channel-pair blocks to ensure even performance.

The circuit measures the time between the arrival of the START signal and the end of the GATE signal, digitizing the time in nine bits plus an overflow bit. The

control signal RUN is on at the beginning of an event, and goes off when GATE goes off. While RUN is on, the counting circuit counts. Each counting circuit is made up of a Sony CXB1136Q counter, 1/3 of a Motorola MC100E431 flip-flop (providing a ninth data bit), and 1/2 of a Motorola MC100E131 flip-flop (providing an overflow bit). If a START signal comes from the input latching circuit while RUN is on, then the counting circuit will be reset to zero. When RUN goes off, the count freezes and START is blocked to preserve the count. RUN later goes on again when the module receives L2TRIG or ABORT.

The parameters of this design satisfy the requirements of Banana subsystems. From Figure 3.9, one can calculate that the worst case TDC resolution is two clock cycles, or 3.2 ns. The double-pulse resolution is four clock cycles, or 6.4 ns. The minimum time between the leading edge of START and the trailing edge of RUN (start-stop offset) is five clock cycles, or 8.0 ns. In other words, if the START comes between zero and five clock cycles before RUN turns off, then the TDC will record $T = 0$.




Figure 3.10: Banana TDC accuracy. Measured Banana TDC time minus LeCroy 3377 TDC time, in units of Banana TDC counts (1.6 ns), for a single channel.

The performance of these TDCs is evaluated by comparing their measurements to those made by the DC TDCs, the LeCroy 3377 modules used to measure DC hit positions. Figure 3.10 shows $t_{ban} - t_{3377}$ for 3,392 hits on a drift chamber wire, where t_{ban} is the time read out for that channel’s Banana TDC and t_{3377} is the time read out from a 3377 channel. The offset of $t_{ban} - t_{3377}$ is arbitrary and is chosen to make the mean close to zero. The time in Figure 3.10 is measured in units of Banana TDC counts (1.6 ns). The 3377 TDC counts time intervals of 0.5 ns, so some rounding errors are inevitable in measuring $t_{ban} - t_{3377}$. These errors would cause an RMS of 0.4 counts, even if both TDCs agree perfectly. The actual RMS in Figure 3.10 is 0.89 counts, over the range -2 to $+7$ counts. The rest of this error comes from time jitter introduced in the DC-Repeater/delay-cable/RC-Receiver system. However, we found this inaccuracy acceptable because the Banana TDCs only had to be able to resolve that events arrived within a 19 ns (~ 12 Banana count) bucket.

3.4.5.3 Lookup Tables

The RAM-lookup tables and hit-logic circuits correlate the raw TDC values to decide if the DC signals come from out-of-time events. First, raw data bits from the TDC circuits are converted from ECL to TTL (Transistor-Transistor Logic) levels and placed on the “time-address bus” traces. These data then address RAM lookup tables. Every RAM gets TDC values from two adjacent DC wires. Each memory location in a RAM represents a pair of possible arrival times from its two input channels. Thus each lookup table can be thought of as a two-dimensional map, with each input time as an axis. Every Banana module has 32 RAMs, one for each pair of inputs; the 32nd RAM receives a TDC value from the the lowest channel in the next-highest board via the Blossom boards.

Besides going to a RAM, every nine-bit TDC value goes via the time-address bus to an address control circuit. These circuits are implemented in eight Altera EPM7064

EPLDs. Each of these EPLDs is attached to four TDC circuits. The address control circuits allow FASTBUS readout of TDC values. FASTBUS commands can cause the address-control circuit to take control of the time-address buses to load and verify the contents of the RAMs.

Each RAM has 18 address bits as input and four data bits as output. One output bit (the “C” bit) is on if the two input times (t_1, t_2) correspond to an in-time pair. A definition of an in-time pair is drawn on Figure 3.11. The dots are 10,903 Banana time pairs from DC hits that are reconstructed as being part of good tracks. There are several features to note in Figure 3.11. One is that, because START signals reset TDC counters, larger t_1 and t_2 correspond to earlier times. Another is the nonlinearity of t_1 vs. t_2 . If DC drift speeds were constant, $t_1 + t_2$ would be a constant. Factors such as acceleration of the electron avalanche or non-uniform electric field shape curve the in-time t_1 vs. t_2 plot into a characteristic shape. This shape, which resembles a boomerang or banana, is the source of the name of this system. Another feature of Figure 3.11 comes from delta rays, knock-on electrons sometimes released by charged particles passing through the DCs. Because the delta rays move at large angles to the incident track and can ionize the DC gas, they can reduce measured drift time. Delta rays fill in the the region above and to the right of the in-time “banana.” Because we can reconstruct good tracks that have delta-ray hits, the in-time region is drawn to accept delta rays.

There are three other RAM outputs, of which two are used (“T” and “D” bits) and one is unused. The T bit is on if t_1 is an in-time hit: it could have come from an in-time pair. The T bit ignores t_2 . If the memory map for the T bit were drawn on Figure 3.11, it would look like a vertical stripe running from top to bottom and overlapping the in-time pair region. The D bit is on if the time pair could have come from a single track, whether in-time or out-of-time. In principle, the D bit map would be formed by shifting the C bit map by multiples of one bucket in both t_1 and t_2 . In




Figure 3.11: Banana t_1 vs. t_2 and in-time pair region boundary for hits on tracks.

practice, because the C bit map is so wide, the D bit map forms a diagonal stripe from lower left to upper right, overlapping the C map (in-time pair region).

3.4.5.4 Hit Logic

Next, the 32×3 data bits from the RAMs go to an array of eight EPM7064 EPLDs that make up the hit logic circuit. The overflow bits from all TDCs also go to the hit-logic array. For each DC wire channel, the hit logic decides if there is an acceptable hit by looking at the data from that wire, its neighbors, and its next-to-neighbors. If the overflow bit for any channel is on, then the RAM data bits from that channel are ignored, because the TDC values that produced them are invalid. Then, only channels that are in-time hits (T bit is on) are considered as good-hit candidates. The hit logic then considers the RAM data bits for each channel produced by correlations with its neighbors. For example, suppose the hit logic is looking at channel 3. The hit logic will flag a channel as a good hit if any of the following conditions are true:

$$\begin{array}{ll}
 \text{Isolated hit:} & T_3 \cdot \overline{D}_{23} \cdot \overline{D}_{34} \\
 \text{In-time-pair:} & T_3 \cdot (C_{23} + C_{34}) \quad \text{(The simplest case)} \\
 \text{Triple hit:} & T_3 \cdot [(C_{12} \cdot \overline{D}_{34}) + (C_{45} \cdot \overline{D}_{23})] \\
 \text{Quintuple hit:} & T_3 \cdot C_{12} \cdot C_{45}
 \end{array}$$

X_{ij} is the X bit from the RAM that had channels i and j as inputs.

To make these multi-channel correlations across module boundaries, a number of bits have to be passed through the Banana Blossoms. From the upper neighbor comes the two lowest overflow bits and the lowest C bit. From the lower neighbor comes the two highest overflow bits, the highest two C bits, and the highest D bit.

The hit-logic array is also a FASTBUS interface for the RAMs. The hit-logic array can load one bit of the RAMs while preserving the contents of the other three

bits. A register in the hit-logic chips can mask out any RAMs during writing operations.

The performance of the hit logic can be seen in Figure 3.12. The data in this plot comes from 112,646 occasions when any two adjacent DC wires were hit and the wires adjacent to those two were not hit. So if there were four wires in a row, then the first was not hit, the second and third were hit, and the fourth was not hit. By looking only at these isolated pairs, there is no confusion from isolated, triple, and quintuple hits. The events selected in Figure 3.12 are those where the Bananas identified no good hits on either wire of the pair. The dark diagonal stripe comes from out-of-time tracks. The bands in the stripe correspond to early or late buckets. These are events that the Banana is correctly identifying as being out-of-time (D bit is on), and rejecting.

The sparsely populated rectangles in the upper left and lower right are events where t_1 and t_2 are uncorrelated and out-of-time (T bits are off), and the Banana rejected the hits. The unpopulated regions in Figure 3.12 are where good pairs could occur and the Banana accepted one or both of the hits.

3.4.5.5 Hit Counting

Hit counting in Bananas is the same as in Kumquats, as described in section 3.4.4.2. The only difference is that the mask of good hits comes from the hit logic, described in the last section, instead of from the input latches. The time delay between the end of GATE and a Banana's local hit-count-ready signal is about 300 ns.

3.4.6 Hit Summing

As described above, each Kumquat/Banana module adds its local hit count to the hit counts from up to two other modules. This allows a DCHC subsystem to quickly total the hits from an entire DC view by using an adder tree. Figure 3.13 gives a block diagram of an adder tree with nine modules. For DCHC subsystems with fewer than nine modules, upper nodes of the tree in Figure 3.13 are pruned first. The adder tree




Figure 3.12: Banana t_1 vs. t_2 and in-time pair region boundary for isolated pairs rejected by Banana.

for Kumquat crates does not differ in principle from that of Banana crates. The total-count bits and count-ready bit are available as NIM outputs on the front panel of all DCHC modules. The count and ready of the lowest module in the tree go to the L2 Control logic, with the ready signal as the DONE signal for its DCHC subsystem. The DONE is produced within 540 ns of the end of the GATE/L1TRIG signal.




Figure 3.13: Structure of DCHC adder tree.

3.4.7 Outputs to Level 2

In addition to the hit-count outputs, two other types of data are output. One is the mask of DC hits: latched signals from Kumquats and in-time hits from Bananas. This is output on two 17-pair ECL connections on the front panel of each module. Each connection has hits for 16 adjacent DC wires. The STT processor uses four such sets of outputs per DC X-view. The other output is the OR of paddles of hits. There are ORs of 4, 8, 16, and all 32 channels, for a total of 15 OR bits. These are output on NIM connections on the front panel of each module, and are inputs for the YTF processor.

3.4.8 FASTBUS interface

The FASTBUS interface in DCHC subsystems allows initialization and testing. No on-line readout of DCHC data occurs through FASTBUS. The FASTBUS interface in DCHC modules is implemented in five MAX 7000 EPLDs. After a DCHC crate powers up, various registers and flip-flops must be initialized to well-defined states and, for Bananas, RAM lookup tables must be loaded. We perform these tasks using the BiRa FSCC (FASTBUS Smart Crate Controller) module in each FASTBUS crate. The FSCCs are single board computers connected to the local ethernet network. Thus we can log in to each crate's FSCC and execute software programs that perform FASTBUS operations on Banana or Kumquat modules, such as the aforementioned register initialization. Also, all data and many relevant signals can be read out via FASTBUS. This provides a handy tool for debugging and testing.

The other major use of FASTBUS, loading the Banana RAMs, requires a non-trivial procedure. Each module has three RAMs per channel and 32 channels, and each RAM has 2^{18} bit-addresses, for a total of 786,432 bits per module. Passing this much data via FASTBUS is relatively slow: loading one Banana takes about one minute. Because only one module can be addressed by the FSCC at a time, the loading of a crate's modules (up to eight) must be done in series. Transferring these maps from hard drives to the FSCCs via ethernet would add even more time to this loading procedure. Instead, the FSCCs read only a minimal amount of data from disk: a time offset for each channel in the crate and a set of time-pairs which describe the outline of the in-time-pair region. From these data, the FSCCs can construct the maps to be loaded for each channel.

3.4.9 FERA Interface

The FERA readout circuit allows the DAQ to read data from Banana and Kumquat modules. The FERA data stream is made up of 16-bit “words.” The 16th bit is on only for the first word, and so it indicates where the data from a module begins. The remaining 15 bits are available for data. The first four words from each module carry header information: module identification number, mask of hits, local hit count, and mask of ORs. For Kumquat modules, these are all of the words. For Banana modules, each following word carries data for a TDC channel. Any TDC channel for which the overflow bit is off contributes a word containing its TDC count and a number indicating which TDC it is. Therefore, the number of words in the data stream from a Banana for one event ranges from 4 to 36, depending on DC activity.

A readout process is started by L2TRIG. In Bananas, the FERA control circuit starts composing the data stream in the first of two event-buffer FIFO (First In, First Out) chips. The control circuit is implemented in an Altera EPM7096, called BFCTRL. BFCTRL sends control signals to store the header words in the first FIFO. Then another EPM7096 EPLD called GEN_CH1 takes control. It searches for non-overflow TDC channels and sends control signals so that TDC words for those channels are stored in that first FIFO. Next, BFCTRL promptly pushes the contents of the first FIFO into the second FIFO. BFCTRL uses the 17th and 18th bits of the second FIFO to flag the first and last words of an event. This FIFO buffering allows a Banana module to process an event even if the DAQ system has not yet read out previous events. This feature increases the instantaneous rates at which the Bananas may function, but it was not used during the 1996–1997 run of KTeV. It is being used during KTeV’s 1999–2000 run [45].

In Kumquats, the FIFO buffering is not present. The Kumquat FERA control circuit multiplexes the FERA data words from the Kumquat circuits directly to the

FERA data bus.

In both module types, the FERA control circuit then communicates with the crate's DYC, in a process that pops out the data event by event. The FERA readout circuit can transport 20 mega-words per second. By setting a jumper, it can transport at half speed: 10 mega-words/sec; we use this feature to ensure accurate transmission. The readout procedure follows the FERA standard, but there are several features unique to the DCHC readout. These arise from the unusual distance between the source of FERA data (Banana/Kumquat modules) and their destination (DYC module). The DCHC modules are read out from connections on the Blossom boards on the back of the FASTBUS crates, while the DYC's FERA connections are on the front of CAMAC crates in the FASTBUS racks (see Figure 3.3). The connecting cables can be as long as 8 feet. To preserve signal integrity over these distances, several measures are taken:

- Differential ECL is used for all signals between the FASTBUS and CAMAC crates. This requires the addition of one or more CAMAC modules to translate to the single-ended ECL used by the DYC.
- All FERA data words are carried on 17-pair ribbon cables wrapped in copper foil.
- The FERA data bus is split into two cables for crates that have eight or more modules. This reduces the maximum cable length for such crates.
- The WST signal is not transmitted on a bus from the DCHC modules to the CAMAC crate, as is customary in FERA readout. Instead, the WST from each module is transmitted as differential ECL to the CAMAC crate, where an active circuit ORs the WSTs to form a single-ended WST.
- The REQ signal for the DCHC modules is generated using a separate CAMAC module. This module turns on REQ when it receives L2TRIG. It also receives

the REN/PASS token signal from the last DCHC module; when the token arrives, the module turns off REQ and passes the token back to the DYC.

3.5 Readout

After an event is accepted by the L2 Trigger system, the detector is read out; that is, all digital data that has been recorded for that event is collected into one central location. This process can be divided into two stages: the crate stage where data moves from individual modules to the DYCs, and the stream stage where data moves from the DYCs to the online computers.

3.5.1 Crates

Data from most detector elements uses FERA readout to DYC modules. The exception is readout of the calorimeter, discussed below. The DYC3 (Damn Yankee Controller) is a CAMAC module made by Fermilab. After the modules connected to a DYC have received L2TRIG and have sent a REQ to the DYC, the DYC reads out its modules. During this readout, a DYC sends a DYC BUSY signal to L2 Control. If the DYC's internal buffer becomes half-full, the DYC produces a DYC FULL signal, also used by L2 Control. The DYC puts the contents of its buffer onto a RS485 bus when it receives a permit-in signal, and turns on a permit-out signal when this transfer is complete. These permit-in/permit-out allow the DYCs to belong to token rings that control access to the RS485 buses.

The DPMTs of the calorimeter are read out into a buffer known as the *Pipeline* system following the L2TRIG. The Pipeline, a custom VME system, sparsifies the DPMT data by flagging channels above an energy threshold and cutting some of the 32 bucket-words read from each DPMT. The Pipeline is then read out by CTIRC modules, devices made by Fermilab and similar in function to DYCs. The CTIRCs transfer the calorimeter data onto RS485 buses.

3.5.2 Streams

There are six RS485 buses, known as “streams.” Four streams come from the calorimeter and two from other detectors. Data flows in the streams from the DYCs and CTIRCs to VME buffers. These buffers are known as DDDs, for the three devices that make up each: a DM115 input controller, a DC2 FIFO, and three DPM (6390 VSB/VME Dual Port Memory) modules. There are 24 DDDs, arranged into a grid of six streams and four “planes.” Each stream inputs its data to four DDDs, one for each plane. The DM115s in a plane decide whether to accept or reject an event’s data, based on the Read Out Control mask present in every stream. Each plane sends the events it has collected to a computers. There are four such computers, one for each plane, which perform the last online event selection.

3.6 Level 3 Trigger: Software

The final stage of online event selection is Level 3. In L3, a streamlined reconstruction of decays occurs, and events passing relatively loose cuts are accepted and have their data written to magnetic tape.

L3 software runs on four SGI Challenge computers. Each Challenge has eight 200 MHz R4400 CPUs together providing 1400 MIPS of processing power. Events from one plane are input to each Challenge. One Challenge is only used for monitoring and for calibration events, and writes events to one DLT (Digital Linear Tape) drive. The other three Challenges filter events and write them to one of three or four DLT drives attached to each Challenge.

The filter plane Challenges have several software processes running in parallel. The reconstruction of events uses code similar to that described in chapter 5. For the 2E-NCLUS trigger, the requirements are as follows.

- An event must have DC hits consistent with at least two, oppositely-charged

particle tracks in both the X and Y views.

- The DC tracks must be consistent with having come from a single decay vertex.
- Energy clusters in the calorimeter are identified using an algorithm similar to that used by the HCC. Events with fewer than four such clusters are rejected.
- Each of the two DC tracks must point towards one of the clusters in both X and Y .

Of these events, those with $E/p > 0.75$ on both tracks are tagged as “2e” events. Those with $E/p < 0.75$ on both tracks are tagged as “2pi” events and those with $E/p > 0.75$ for one track and $E/p < 0.75$ for the other are tagged as “epi”. In addition, a random prescale accepts 1 in 250 2E-NCLUS events; these are tagged as “random accepts.”

Although the average L2TRIG rate can be 22 KHz, that only persists for the 19 second spill. The DDDs can buffer the data for an entire spill, allowing Level 3 to use the entire 60 second cycle for event processing. The Level 3 trigger requirements accept only 12% of events overall, so the average rate of writing to tapes is about 840 events per second. For the 2E-NCLUS trigger, the L3 requirements save 7.0% of L2 2E-NCLUS events. Of the L2 2E-NCLUS, L3 tags 6.4% as 2e events.

There are eight filter processes running on each filter-plane Challenge. The filter processes pass the data for accepted events to a single process which handles DLT I/O for each plane. This process writes each event to one of the three or four DLT drives on each plane.

In the monitoring/calibration-plane Challenge, all events with calibration triggers are written to a single DLT. Monitoring processes, similar to filter processes, also run on this Challenge. The monitoring processes make histograms of various quantities based on a small fraction of events sent to L3. These histograms are stored in shared memory where they can be viewed by experimenters during data taking.

Chapter 4

Taking Data

This chapter describes some procedures used in KTeV data taking. It also describes changes made to the detector during data taking.

4.1 Runs

The “run” is a basic unit of time in KTeV data taking. A run is a period during which the detector and beam conditions are uniform. Trigger requirements, detector element high voltages, and position of E799 beam elements remain the same during a run. If any significant changes in the primary proton beam occur, the current run is stopped. Runs are subdivided into 60-second spills. Good E799 runs typically last several hours: a few hundred spills. Ideally, runs are stopped after eight hours, but some were allowed to continue for 13 hours, until the output DLTs reached capacity. Runs are numbered consecutively.

Besides standard E799 data-taking runs, there are several other types of runs used to get calibration data. One of these is the muon run, described in section 2.2.1.3. This type of run is performed daily to study the drift chambers. During a muon run, the beam stop is moved into place and the sweeper magnets are adjusted to allow muons produced by the proton beam to enter the detector (see section 2.1.3). The primary beam intensity is reduced to about 4×10^{11} protons on target per spill. The analysis magnet field is turned off, so the muons are not deflected. After some tens of spills the

muon run is ended, and the beam configuration is returned to normal. The analysis magnet is turned on again, but with the opposite polarity to what it had before the muon run.

The other major type of calibration run is the laser scan run (see section 2.2.2.3). This run type is preferably taken during the inevitable periods when the Tevatron is unable to provide proton beam to the detector. This is because laser scan do not require beam, and intense beam could conceivably contaminate the scan data. Laser scans were performed every 2.2 days, on average, during E799. Laser scans last long enough for the CsI laser system to fill all DPMT ranges with sufficient statistics. With a laser flash rate of about 5.5 Hz, this takes about 45 minutes.

4.2 Run History

E799 data taking was divided into two periods, both during 1997. The first was *Winter*, which started with run number 8088 on January 24 and ended with run 8913 on March 24. After run 8913 the detector took data for the E832 (ϵ'/ϵ) experiment. The second E799 period was *Summer*, which started with run 10463 on July 24 and ended with run 10970 on September 3.

The Winter started with the 0.5×0.5 mrad neutral beams described in section 2.1. The L1 trigger requirement for the 2E-NCLUS trigger was as described in section 3.2.2, with the exception that the CA was not in veto. The L2 requirements were as described in sections 3.3.1 and 3.3.2: four or more clusters in the CsI and four hits in all DC Y-views with one missing hit allowed in either 1Y or 2Y. Kumquats were used for all Y-view DCHC systems.

This trigger configuration changed a number of times during E799, either because of hardware problems with trigger electronics, repairs of previously broken components, commissioning of new components, or realizing that the trigger could be improved. These changes, as they pertain to the 2E-NCLUS trigger, are described below.

The first major trigger change was in run 8262, when a Banana DCHC system was added for DC2X. A requirement of at least one hit in this system was added to the 2E-NCLUS trigger. In run 8280, the CA, SA3, and RC8 L1 sources were allowed to veto 2E-NCLUS. In addition, the DC-OR requirement was tightened to require an OR “hit” in all four views (1X, 1Y, 2X, 2Y). Soon thereafter, in run 8284, the RC8 veto was removed. In run 8518, the DC-OR requirement was loosened to again allow one view to have no hits. In run 8536, because of hardware problems with the DC-ORs for DC2Y, the DC-OR requirement was changed to $1DC1X \cdot 1DC1Y \cdot 1DC2X$. In run 8577, the MU2 veto was dropped. In run 8761, the DC-OR requirement was changed to require $1DC2Y$ but not $1DC1Y$.

The next major changes were for the Summer period. The big change was the increase in the neutral beam sizes, to 0.59×0.59 mrad. This change also entailed changing some sweeper magnet field strengths and removing the slab collimator. These changes were motivated by a planned reduction in primary beam intensity from the Tevatron, brought on by the number of other experiments requiring beam that summer. The net result was an increase in kaon yield per spill for the Summer.

In the trigger, the DC-ORs went back to allowing one missing view. The DCHC configuration was changed to have Bananas for all DC1 and DC2 views and Kumquats for all DC3 and DC4 views. The requirement of a hit in the 2X Bananas was removed, but the Y-view requirements were retained. In run 10464, the DC-OR requirement was tightened again to require activity in all four views. In run 10482, the 2X Banana requirement returned. In run 10540, the HA veto was changed from HA_AC to HA_DC (see section 3.1).

In run 10934, the Tevatron began providing 23 seconds of beam per spill instead of 19 seconds, increasing the kaon yield per spill.

4.3 Data Yield

All told, there were 312 runs in E799, lasting on average 213.0 spills each. The data taking period lasted 105 days. (The overall live-time of data taking was thus 44 %, including accelerator down times.) The data from these runs occupies 950 DLT III output tapes with about 12 Gb of data per tape. Of these events, there are 129.4 million events with the 2E-NCLUS trigger. Chapter 8 addresses the question of how many kaon decays occurred to produce these data.

Chapter 5

Reconstruction

This chapter describes how a kaon decay is reconstructed using the data recorded and digitized when L2TRIG occurs. Reconstruction includes determining the identities of and measuring the momenta or energies of particles produced by kaon decays. All reconstruction is done using computer programs, written in Fortran, using a code package known as ktevana. In some sense, the purpose of this chapter is to describe the algorithms used by ktevana.

5.1 Calorimeter Clusters

One stage in reconstruction is measuring the location and quantity of energy deposited in the CsI calorimeter by various decay products. First, the energies of individual channels are determined. Second, channels are combined into clusters, and the energy and position of the clusters is determined.

5.1.1 Channel Energy

To determine the energy in a CsI crystal channel, first a “charge” (Q) is calculated for each of four RF buckets. Buckets are also known as “slices” in the context of DPMTs. The first, or in-time, slice corresponds to the bucket when the largest amount of energy is seen in a crystal. Then,

$$Q(i_{chan}, i_{slice}) = M(i_{chan}, i_{range}, i_{cap})N_{adc}(i_{chan}, i_{slice}) + C(i_{chan}, i_{range}, i_{cap}).$$

Here, i_{chan} is the index of the crystal, i_{slice} is the index of the slice, i_{range} is the ADC value's exponent read out for that slice, and i_{cap} is the capacitor number read out. M and C are constants calculated during laser-scan runs. N_{adc} is the ADC value's mantissa. The charge $Q(i_{chan}, i_{slice})$ is then converted into an energy $E(i_{chan}, i_{slice})$ by multiplying by a constant $F(i_{chan})$. This constant is calculated by a calibration procedure that requires electrons in K_{e3} decays to have $E/p = 1$. The energies from the in-time slice and the following three slices are summed together to get $E(i_{chan})$.

5.1.2 Clusters

The next step is to find CsI clusters. (Clusters are also discussed in section 3.3.1.) See Figure 5.1 for an example of completed cluster reconstruction.

5.1.2.1 Seeds

Cluster finding begins by finding cluster “seeds”: the calorimeter channels that are local maxima. Only channels where $E(i_{chan}) > 0.1 \text{ GeV}$ and for which the HCC input bit is on are considered as candidates for local maxima. The setting of HCC input bits is described in section 3.3.1. A local maximum candidate is compared with its neighbors that also have their HCC input bits on. If the candidate channel's $E(i_{chan})$ is greater than its neighbors', it is flagged as a local maximum. The requirement that the HCC input bits were on has the effect of suppressing out-of-time clusters, because the ETOT hardware only flags channels with more than about 1 GeV seen during the in-time slice.

After the cluster seeds are found, the “raw” energy associated with each cluster is estimated. This is done by summing $E(i_{chan})$ for all channels in a 7×7 grid centered on the cluster seed. The grid is only 3×3 for large crystals.

KTEV Event Display

/crypt/mikelson/summer.dat

Run Number: 10463

Spill Number: 26

Event Number: 3112343

Trigger Mask: 1

All Slices

Track and Cluster Info

HCC cluster count: 5

ID Xcsi Ycsi P or E

[T 1: -0.2885 0.0220 -14.42

C 2: -0.2913 0.0228 14.97

[T 2: 0.6248 0.3915 +3.01

C 5: 0.6393 0.3877 2.91

C 1: -0.4375 0.0489 49.76

C 3: -0.2638 0.2018 56.06

C 4: 0.0091 -0.1177 72.52

Vertex: 2 tracks, 2 clusters

X Y Z

-0.0963 0.0148 96.195

Mass=1.0034 (assuming pions)

Chisq=0.03 Pt2v=0.048167





Figure 5.1: Event display of calorimeter. The circles are centered on found clusters. The dark squares are on channels with HCC bits set.

5.1.2.2 Position

The next step is to determine the precise position of each cluster. To do this, $E(i_{chan})$ is summed by row and column in a 3×3 grid around the seed channel. To get the X position, two ratios are calculated: the ratio of energy in $+X$ column to the energy in the seed column (R_+), and the ratio of energy in $-X$ column to the energy in the seed column (R_-). R_+ and R_- then address a lookup table that returns the X position of the cluster relative to the seed crystal position. The positions returned by using R_+ and R_- are interpolated to get the final position. A similar procedure returns the relative Y position.

These lookup tables are created by looking at photon clusters in $K \rightarrow \pi^0 \pi^0 \rightarrow \gamma\gamma\gamma\gamma$ decays and assuming the distribution of cluster positions within each crystal is flat. The distribution of R_+ or R_- is made for such decays then integrated. The y-axis of the integrated function is then rescaled to go from one edge of a crystal to the other edge; this is the position lookup function. There are separate tables for small and large crystals, and also for different bins of raw cluster energies. The lower limits of the energy bins are 0, 2, 4, 8, 16, and 32 GeV.

5.1.2.3 Corrections

What follows are a series of modifications to the calorimeter energies. These are made to correct for various subtle effects occurring in the calorimeter. The first set adjust the energies of individual crystal channels and recalculates overall cluster energies, while the second set only scales the cluster energies.

The first of the first set is the overlap correction. Overlapping clusters are those for which the 7×7 crystal arrays around the cluster seeds share crystals. The overlap array for large crystals is 3×3 . The correction divides the energy in the overlap crystals between the clusters. The second correction adjusts the energy in individual crystals

that belong to only one cluster but are neighbors with a different cluster. In such cases, the other cluster may contribute significant energy to crystals outside of its cluster boundary. This energy is estimated and subtracted. The third correction occurs when cluster boundaries fall outside of the calorimeter or overlap the beam holes. In this case, a table of the typical transverse energy profile of a cluster is used to estimate the energy in the “missing” crystals. The fourth correction is for “sneaky energy,” the component of an electromagnetic shower that leaves the CsI, crosses a beam hole, and is deposited on the other side. Sneaky energy is a problem when two clusters are near the same beam hole. The correction subtracts the estimated sneaky energy from crystals near the beam holes. After the sneaky energy correction, the first three corrections (overlap, neighbor, and missing crystal) are repeated. The fifth correction is for crystals in clusters for which no there was no read out because the energy in the crystal was below threshold. The energy in the crystal is estimated using its position with respect to the cluster’s center and the cluster’s energy.

After these corrections are applied to individual crystals, a series of scale factors are applied to the cluster energies. The first of these is for the fraction of energy in a shower which goes out of the sides (or front or back) of the crystals used in a cluster. This factor is only a function of crystal size and size of the cluster array. The second factor corrects for the non-uniform response of the crystals across their faces. The scintillation light from a particle hitting the corner of a crystal is different from the light of a particle hitting the center. This factor is determined using the cluster position information. (This correction is not made for events simulated by Monte Carlo.) The third factor corrects for a nonlinear response of the calorimeter with respect to cluster energy. This factor is a function of cluster energy and calorimeter region. The fourth scale factor and final correction is for small time variations in the response of the calorimeter. This factor is a function of run number and calorimeter region.

5.2 Spectrometer Tracks

The other major part of reconstruction is finding charged-particle tracks in the drift-chamber system. This proceeds by finding acceptable DC hits, and correlating them to find track segments, track candidates. These candidates are matched to each other in X and Y using calorimeter clusters. The tracks can then be evaluated to see if they point towards a common decay vertex. The momenta of the tracks can be estimated using the analysis magnet.

The tracking differs from the clustering in that there are places where the tracking algorithm will abort if it looks like certain basic requirements are absent. In this analysis, these requirements are that at least two distinct tracks be present in both the X and Y views. For an example of completed track reconstruction, see Figure 5.2.

5.2.1 Track Candidates

Track candidates are possible components of tracks, as found by looking at the X-view or Y-view of the DCs independently. Thus track candidates are called either X-tracks or Y-tracks.

5.2.1.1 Evaluating Hits

If there are at least two DC wire hits of some kind in every plane-pair, then track reconstruction begins with identifying and classifying DC hits. The TDC values for the DC wire hits are corrected for the time offset of each wire. If this corrected time lies within the in-time hit window (115–350 ns), then it is converted into a distance using a lookup table. This time-to-distance table is created using data from muon runs. If there are multiple hits on a wire, only the earliest in-time hit is used. If two adjacent wires are hit, one in the upstream plane and one in the downstream plane, then they form a hit pair. The classification of each pair is based its sum-of-distances (SOD): the distance

KTEV Event Display

/crypt/mikelson/summer.dat

Run Number: 10463

Spill Number: 38

Event Number: 5352428

Trigger Mask: 1

All Slices

Track and Cluster Info

HCC cluster count: 5

ID Xcsi Ycsi P or E

[T 1: 0.4799 -0.0023 +7.63

C 5: 0.4848 -0.0005 7.67

[T 2: 0.1653 -0.1355 -47.49

C 3: 0.1648 -0.1363 47.74

C 1: -0.4038 0.1414 5.54

C 4: 0.1400 0.2680 73.18

C 2: 0.2572 -0.4483 17.08

Vertex: 2 tracks, 2 clusters

X Y Z

0.0816 0.0312 98.733

Mass=0.7201 (assuming pions)

Chisq=2.84 Pt2v=0.004758




Figure 5.2: Event display of drift chambers. Upper plot is X-view, lower plot is Y-view. Solid lines are reconstructed tracks. Dashed lines are reconstructed photon trajectories.

for one wire in the pair added to the distance for the other. The Banana system works on the same principle (section 3.4.5.3). Figure 5.3 shows the distribution of SODs at this point. If the SOD is within 1.00 mm of the wire separation, (6.35 mm), then the pair has a good SOD and is considered to be in-time. (Good SODs in DC3X and DC4X are 6.35 ± 1.5 mm because of the large angles of tracks deflected in the analysis magnet.) If the SOD is low, possibly caused by a delta ray, then the hits are not treated as a pair. Low SODs can also occur when two tracks are so close that they share a cell. If the SOD is high, from resolution effects or DC pathologies, then the hits are not treated as a pair, and both cannot be used as part of tracks.

5.2.1.2 Y Tracks

Next, Y-tracks are sought. This is done with four nested loops, one for each DC. The first loop is over pairs/hits in DC1 and the second is over DC4. The next loop is over pairs/hits in DC2 that lie within 5 mm of a line drawn between the pairs/hits being considered in DC1 and DC4 (the “road”). This is then repeated for DC3. If a pair/hit was found within the road in both DC2 and DC3, then the track candidate is fit to a straight line. The candidate is only used if the χ^2 per degree-of-freedom of this fit is less than $(2 \text{ mm})^2$. Furthermore, the track candidate may only use up to two hits from low SOD pairs, *or* one low SOD hit and one isolated hit. This continues until all possible Y-tracks have been considered. If there are at least two, they are checked to see whether they can coexist. Generally, this means that two track candidates may not share any hits. The exception is that they may share, in one plane-pair, a hit which is shared by two in-time pairs; this could happen if the tracks pass through adjacent cells and the SODs happen to be good for both pairs.




Figure 5.3: “Raw” SOD. Y-view hits on tracks in $K \rightarrow \pi^0 \pi_{Dalitz}^0$ decays. The entries at SOD=0 are isolated hits.

5.2.1.3 X Tracks

If there are at least two exclusive Y-tracks, then X-tracks are sought. The first step here is to form X-track segments by pairing hits in DC1 to hits in DC2. These hits can include only in-time pairs or hits from low-SOD pairs. Furthermore, the angle the segments make with the Z axis must be less than 100 mrad. If there are two upstream candidates, the code goes on to try to add downstream segments. These are found the same way as upstream segments, but the angle can be as much as 150 mrad and must include at least one in-time pair. Downstream and upstream segments are combined into X-tracks if they intercept the magnet plane (at $Z = 170.008$ m) within 6 mm of each other. Like Y-tracks, the X-tracks may only use one or two hits from low SOD pairs, *or* one low SOD hit and one isolated hit. Also like Y-tracks, the X-tracks are checked to see if they can coexist, with the same criteria. At least two non-coexisting X-tracks must be found to satisfy this stage of tracking requirements.

5.2.2 Vertex

The next stage in track reconstruction is to determine if the DC activity is consistent with two tracks intersecting at a point upstream where a parent particle decayed (a vertex). One use for this vertex position is for reconstructing photon momenta, in combination with the photons' cluster energy. If there are more than two X-tracks or Y-tracks, then there is more than one possible vertex, and the best one must be picked.

5.2.2.1 Identifying Vertex Candidates

The next step is to look for possible vertices using only Y-tracks or only X-tracks. First, all pairs of Y-tracks are looped over. For each pair, the largest Z from which both tracks could come is calculated. This is done using the tracks' slopes and intercepts as well as the maximum uncertainty in slopes, and only using the upstream segment (DC1-DC2) of each track. If this Z is sensible (> 0), then the minimum Z , nominal Z , and

vertex weight for the pair are calculated. The vertex weight is used in averaging X and Y vertices together, and is equal to the difference in slopes squared. If a Y-pair vertex candidate is found, then the procedure is repeated for the upstream segments of the X-tracks. Next, each Y-vertex is compared to each X-vertex. If the Z range overlaps for an X-vertex and a Y-vertex, then the vertex pair is considered a vertex candidate. The nominal Z for the candidate is the weighted mean of the nominal Z s of the X and Y vertices.

5.2.2.2 Track-Cluster Matching

If at least one vertex candidate is found, then X-tracks are matched to Y-tracks. Because of the independence of the DC X and Y views, this matching requires a hodoscope of some sort. In KTeV, the calorimeter fills this requirement nicely. Therefore, cluster finding must be performed before track matching. In practice, tracks candidates are found before clusters, allowing some events with insufficient tracks to be rejected before that CPU intensive stage.

The downstream segments of both X and Y tracks are projected to the calorimeter. Tracks for which the X or Y positions fall outside the calorimeter are not matched. Then each pairing of an X-track with a Y-track is considered. If the X, Y position of a pair places the pair in a beam hole, the pair is not considered. The calorimeter cluster which is closest to the pair is found. Next, the vertex candidates are reconsidered. For the tracks in each vertex candidate, there are two ways to match the X-tracks to Y-tracks. The way that gives the smallest sum of squares of track-cluster separations is selected for each vertex candidate. This X-Y matching is then used for that vertex candidate.

5.2.2.3 Selecting a Vertex

Several corrections are applied for each vertex candidate. These corrections are for subtle detector effects, analogous to those made for cluster energies (see section 5.1.2.3). Then a figure-of-merit is calculated for each vertex candidates, and the best one is selected.

The first set of corrections is made to the SODs and positions of DC hits used by the vertex candidate. One is made for the transit time of DC signals between the track and the preamp; i.e. the X/Y position of a DC hit is corrected using the Y/X position of the Y/X -track that matches for that vertex candidate. The SODs are recalculated using these new positions. The SODs are then modified to account for the track angle; because the two planes which measure SOD are offset in Z , a tracks that makes an angle with the Z axis will give a SOD not equal to the wire cell size. Finally, X positions in DC2 and DC3 are corrected for fringe fields from the analysis magnet upstream of DC2 and downstream of DC3.

Next, DC hits that are shared are considered. These can occur if tracks pass through adjacent cells. The shared hit is then made a hit pair with the adjacent hit that gives the best SOD, and the hit pair joins that hit's track. The remaining hit is treated as an isolated hit, its SOD is set equal to zero, and it is assigned to the other track.

The list of corrected SODs for hit pairs is then considered. Each view of each track must have at least one hit pair with a SOD of 6.35 ± 0.6 mm. Any pairs with SOD less than 5.75 mm are considered to be low SOD pairs. (Uncorrected SODs less than 5.35 mm have already been moved to the isolated hits category.) If a track has exactly one such low SOD pair, then the pair is split into two isolated hits. Only one hit is used: the one that gives the smallest offset between the upstream and downstream track segments when both are projected to the analysis magnet plane. The other hit

is disregarded. This procedure allows us to use DC hits where there may have been a delta ray or accidental activity. Figure 5.4 shows the SOD distribution at this point.

The next step is fit the position of the current vertex candidate using the corrected hit data. The inputs to the fit are both track positions in X , Y , and Z at DC1 and DC2 and the position resolutions in X and Y at DC1 and DC2. For hit pairs, the position values are the means of the the positions of the individual hits. The resolution for hit pairs is the quadratic mean of the resolutions of the individual hits ($\sqrt{\sigma_u^2 + \sigma_d^2}/4$). Individual hit resolutions are a function of run number and the DC plane to which the hit belongs. These resolution are generally within $3\mu\text{m}$ of $90\mu\text{m}$. In addition, the resolution values are increased for DC2 because of momentum-dependent multiple scattering. Once the positions and resolutions are decided upon, an iterative procedure fits the position of the vertex candidate. For the first iteration, the tracks are weighted equally. For following iterations, the track slopes are recalculated using the new vertex, and higher momentum tracks — with less multiple scattering — are given more weight. If the new vertex Z is within 1 mm of the previous iteration’s Z , then the new vertex is accepted and the loop stops. This generally occurs within only a few iterations. This vertex position is used to get vertex-corrected track slopes and intercepts.

The best vertex candidate is chosen by considering several things. These include the χ^2 for the vertex fit, the χ^2 for the hypothesis that the upstream and downstream segments of each track intersect at the magnet plane, and the number of poor-SOD or isolated hits used in finding the vertex. A sum is formed using these quantities, and the vertex that minimizes the sum is used for the event.

5.2.3 Momenta

The momentum of each track is then reconstructed. The track X and Y at the analysis magnet plane are found by taking the mean of the projected upstream and downstream segments. These X and Y are then used to look up the “kick”, or magnet




Figure 5.4: “Corrected” SOD. Y-view hits on tracks in $K \rightarrow \pi^0 \pi^0_{Dalitz}$ decays. (B) is the peak in (A), with a gaussian fit.

strength: the magnitude of the transverse momentum given to a particle passing through the magnet at that position. The momentum's magnitude is equal to the kick divided by the track's change in slope in the magnet.

5.2.4 Extra Tracks

In the final stage of DC reconstruction, we search for additional tracks not used for the vertex. The unused track candidates are matched. SOD corrections are applied for each matching as for vertex candidates: DC wire transit times, track angle, and fringe fields. Low SOD pairs are split as for vertex candidates. Extra tracks must have at least one good SOD pair. The “best” extra track is the one that minimizes the sum of several quantities: the χ^2 for the upstream-downstream segments meeting at the magnet, the χ^2 for the track pointing to the cluster, and the number of poor SODs used. The process then continues for the remaining track candidates, which cannot have been used by the vertex tracks or by previously accepted extra tracks.

For this analysis, no events with extra tracks are used.

5.3 Kinematic Quantities

After cluster and track reconstruction, enough information is available to reconstruct the four-momenta of the particles in the event. These four-momenta can then be used to calculate various decay-related kinematics, such as invariant masses for combinations of particles and angles between particles in various reference frames.

For a charged particle, the three-momentum is the vertex-corrected direction of its track (upstream of the magnet), multiplied by the magnitude of its momentum. The energy of the track depends on the rest mass assumed for the particle, by $E^2 = \vec{p}^2 + m^2$.

For a photon, the direction of the three-momentum comes from the difference between the photon's cluster position and the position of the vertex. The cluster Z is within the calorimeter, and is reconstructed as a logarithmic function of the cluster

energy. The three-momentum is the direction times the cluster energy.

5.4 TRD Particle ID

The reconstructed tracks are also used in electron/pion identification using the TRDs. As discussed in section 2.2.3, electrons (and positrons) tend to deposit more energy in TRD chambers than charged pions. This energy is summed into groups of two or four adjacent wires in each plane. Figure 5.5 shows this energy as an ADC value. Electrons in the plot are tracks with $0.95 < E/p < 1.05$ and pions are tracks with $E/p < 0.8$.




Figure 5.5: TRD ADC values for $e^\pm$ and $\pi^\pm$ tracks. All from upstream plane in TRD chamber 3, Ke3 decays in runs 8384, 8387, and 8397.

The distributions in Figure 5.5 could make a crude e/π discrimination. The pion ADC distribution provides the probability $\phi(x)$ of a pion giving ADC value x . The confidence level $\psi(x)$ for a given measured x coming from a pion is

$$\psi(x) = \int_x^\infty \phi(y) dy.$$

It can be seen from Figure 5.5 that a single $\psi(x)$ would not provide much discrim-

ination. Therefore, ADC values from multiple TRD are considered. To identify which TRD wire groups to use, the downstream track segments of each track to be identified are projected through the TRD chambers. A track will use the ADC info from a TRD wire group if it passes through the plane of that group, if the track does not hit inside the beam region or outside the chamber, and if no other track passes within 5 mm of the group. Thus the track has a set of confidence levels $\psi(x_i)$ associated with it, where the minimum i is 1 and the maximum is 16, the number of TRD wire planes. Events with tracks that have no usable TRD ADC values are rejected in this analysis.




Figure 5.6: Π_{TRD} for $e^\pm$ and $\pi^\pm$ tracks. Ke3 decays in runs 8384, 8387, and 8397.

The combined TRD “pion probability”, Π_{TRD} , for a track with N ADC values is

$$\Pi_{\text{TRD}} = \left(\prod_{i=1}^N \psi(x_i) \right) \sum_{j=0}^N \frac{\left[-\ln \left(\prod_{i=1}^N \psi(x_i) \right) \right]^j}{j!}.$$

The product $\prod_{i=1}^N \psi(x_i)$ is the probability that a pion will generate a set of ADC values x_i or higher. The quantity Π_{TRD} can be thought of as a confidence level for the hypothesis that the ADC values are consistent with those of a $\pi^\pm$ track. Figure 5.6 shows Π_{TRD} for electron and pion tracks (with same E/p definitions as in the previous

plot). Using a cut on Π_{TRD} , pion rejection of $276 \pm 38 : 1$ can be attained while retaining 90% of electrons. The overall single pion rejection power varies somewhat with run due to changes in TRD gain and calibration, but was better than 200:1 with 90% electron efficiency in E799 [46].

Chapter 6

Simulation

The rare decay studies in this analysis require an estimate of the acceptance of the detector, the probability that a kaon decay in a particular mode will be identified in a particular way. This includes the acceptance for $K_L \rightarrow \pi^0 e^+ e^-$ and $K_L \rightarrow e^+ e^- \gamma \gamma$ being correctly identified, but also the acceptance for other, background modes being identified as $K_L \rightarrow \pi^0 e^+ e^-$ or $K_L \rightarrow e^+ e^- \gamma \gamma$. Due to the complexity of the detector and the analysis, acceptance is calculated using a Monte Carlo (MC) method. Kaon decays are simulated by computer one at a time, beginning with production of the kaon and ending with the trigger response after all decay products have been accounted for. The simulation decides any random (quantum-mechanical or statistical) events using pseudo-random numbers. Decays are simulated for each decay mode until a sufficient sample is accepted or a sufficient lack of acceptance is demonstrated. The validity of the simulation is judged by comparing spectra of reconstructed variables with data; this is done in chapters 8, 9, and 10.

The KTeV MC is done using a package of Fortran code called `ktevmc`. Much as the last chapter described algorithms in `ktevana`, this chapter is a description of `ktevmc`.

6.1 Kaons

The first part of the simulation handles a kaon. This can be divided into two sub-parts: tracing the particle to its decay point in the detector, and determining the

direction and momentum of its decay products.

6.1.1 Kaon Tracing

One of the first things done when generating a particle is to determine what time period (run) is being simulated. For example, the run number controls whether small (Winter) or large (Summer) neutral beams will be simulated. Section 6.3 discusses the run-number determination in more detail. Another early step is to determine the type of particle to generate. The probability of starting with a pure K^0 is 55.25%, with the remaining 44.75% being pure $\bar{K}^0$. In this analysis, the relative $K^0/\bar{K}^0$ amplitudes are only important for simulating $K \rightarrow \pi^0\pi^0_{Dalitz}$.

The initial trajectory of the neutral kaon is parameterized by the magnitude of its momentum P_0 , its vertical angle θ_0 , and its horizontal angle ϕ_0 . Both angles are measured with respect to the primary proton beam. P_0 and θ_0 are determined randomly using the distribution in Figure 6.1. The angle ϕ_0 is distributed uniformly between -0.3 and $+0.3$ radians.

The particle’s flight to the decay region is then simulated. Scattering can occur in the Pb Absorber, changing the direction of flight. This direction is projected to the Z coordinates of the primary, slab (if present), and defining collimators. If the particle hits matter in any collimator, the event is rejected and a new particle is simulated. Some kaons that scatter in the defining collimator material enter the decay region, and some of these even undergo $K_L \rightarrow K_S$ regeneration, but this effect is not simulated in this analysis. Most such events can be excluded by their large transverse momentum.

If the the kaon passes through the apertures safely, then it will be “generated” — it will decay within the fiducial region of this analysis. The fiducial region is 20–220 GeV/c in momentum (limited by the distribution in Figure 6.1) and 90–160 meters in Z (limited below). A running total of the number of generated kaons is kept for calculating acceptance.




Figure 6.1: Momentum vs. vertical angle at target. As mentioned in chapter 2, the proton beam angle is -0.48×10^{-2} rad. K^0 and $\bar{K}^0$ distributions are similar and are averaged here.

Next, the kaon's amplitude is evolved to the upstream end of the decay region. The kaon state can be expressed as a superposition of K_S and K_L states:

$$\begin{pmatrix} \mathcal{A}_S(t) \\ \mathcal{A}_L(t) \end{pmatrix}.$$

The initial states in the target are:

$$K^0 = \begin{pmatrix} +1 \\ +1 \end{pmatrix} \text{ and } \bar{K}^0 = \begin{pmatrix} +1 \\ -1 \end{pmatrix}.$$

In vacuum, the kaon state is evolved in time using the mass matrix:

$$\begin{pmatrix} e^{-(im_S + \Gamma_S/2)t} & 0 \\ 0 & e^{-(im_L + \Gamma_L/2)t} \end{pmatrix}. \quad (6.1)$$

The time t is the particle's proper time,

$$t = \frac{m\Delta z}{|\vec{p}|},$$

where m is the mass, Δz is the distance traveled, and $\vec{p}$ is the momentum. In matter, off-diagonal terms appear in Matrix 6.1, allowing regeneration to occur. This is simulated only in the primary target and the Pb Absorber. Otherwise, the vacuum matrix is applied up to $Z = 90$ m.

Following this, a location is selected between 90 m and 160 m for the kaon to decay. For most decays, the location is randomly chosen to produce an exponential decay distribution consistent with the K_L lifetime. The relatively flat distribution in Figure 6.2(A) shows that a fairly large fraction of K_L decay upstream or downstream of the fiducial region. The K_S amplitude has nearly all decayed away by $Z = 90$ m, and is neglected for most decays.

However, $K_S \rightarrow \pi^0\pi^0_{Dalitz}$ has a much greater branching ratio than $K_L \rightarrow \pi^0\pi^0_{Dalitz}$ (by a factor of $|\eta_{00}|^{-2} = 1.932 \times 10^5$), so the K_S component becomes important. The probability of a $K \rightarrow \pi^0\pi^0$ decay over some period dt is $|\mathcal{A}_S(t) + \eta_{00}\mathcal{A}_L(t)|^2$.

If the kaon is in vacuum, then using Matrix 6.1 this works out to

$$P(t) = |\mathcal{A}_S(t_0)|^2 e^{-\Gamma_S t} + |\eta_{00}|^2 |\mathcal{A}_L(t_0)|^2 e^{-\Gamma_L t} \\ + 2\text{Re}(\mathcal{A}_S^*(t_0)\mathcal{A}_L(t_0)\eta_{00}) \cos[(m_S - m_L)t + \phi_L - \phi_S + \phi_\eta] e^{-(\Gamma_L + \Gamma_S)t/2}.$$

The angles ϕ_L and ϕ_S are the initial phases of the L and S amplitudes, while ϕ_η is the phase of η . Figure 6.3 shows the resulting momentum and vertex distribution in KTeV MC. The peak on the right side is predominantly high momentum K_S that decay upstream, while the remainder is from the K_L amplitude. Figures 6.4 are projections of 6.3 for comparison to Figures 6.2.




Figure 6.2: (A) Vertex Z and (B) momentum for a sample of generated $K_L \rightarrow e^+e^-\gamma$ MC.

After randomly determining the decay Z , the kaon amplitudes are evolved in vacuum from $Z = 90$ m to the vertex. A running total of the normalized K_L amplitudes,

$$\frac{|\eta_{00}|^2 |\mathcal{A}_L|^2}{|\mathcal{A}_S + \eta_{00}\mathcal{A}_L|^2},$$

is kept. This total is used to determine the “number generated” for calculating $K \rightarrow \pi^0 \pi^0_{Dalitz}$ acceptance. This is discussed further in chapter 8.




Figure 6.3: P vs. Z for a sample of generated $K \rightarrow \pi^0 \pi^0_{Dalitz}$ MC.




Figure 6.4: (A) Vertex Z and (B) momentum for a sample of generated $K \rightarrow \pi^0 \pi^0_{Dalitz}$ MC.

6.1.2 Kaon Decay

This is handled somewhat differently for each of the three modes being considered in this analysis.

6.1.2.1 $K \rightarrow \pi^0 \pi^0_{Dalitz}$

The kaon decay to $\pi^0 \pi^0$ is a simple two-body decay. A random direction is chosen for one pion, the other pion gets the opposite direction, the magnitudes of their momenta are both set equal to $\sqrt{(m_K/2)^2 - m_{\pi^0}^2}$, and both momenta are boosted into the laboratory reference frame.

The decay of the pions is handled next. The π^0 lifetime ($c\tau = 25$ nm) is short enough on the energy scale of the experiment ($\gamma \sim 100$) that the π^0 decay position can be safely treated as being identical to the K decay position. The pion decay to two photons is another simple two-body decay, simulated in the same way as $K \rightarrow \pi^0 \pi^0$.

The simulation of the Dalitz pion (where $\pi^0 \rightarrow e^+ e^- \gamma$) is somewhat more involved. The algorithms used are the same as for kaon Dalitz decay, discussed in section 6.1.2.2. The ratio of radiative Dalitz decay to non-radiative Dalitz decay for pions is 0.1618, somewhat less than the same ratio for kaon decays. Like the kaon decays, the infrared threshold is $M_{\gamma\gamma} > 1$ MeV/c². The $\pi^0 \gamma \gamma^*$ vertex form-factor simulated is simply $f(x) = 1 + 0.032x$, where x is the square of the ratio of virtual-photon invariant mass to pion mass.

6.1.2.2 $K_L \rightarrow e^+ e^- \gamma$

The Dalitz decay code simulates the effects of all QED processes to second order in α_{EM} . In the non-radiative decay, this means that loop diagrams appear through interference with the tree-level diagram. When calculating the partial width, these loop terms diverge to negative infinity. To remove this divergence, the tree-level radiative process needs to be included. The partial width for the radiative decay diverges to

positive infinity because of infrared divergence. The sum of the partial widths for non-radiative and radiative decay is finite. The radiative part of the sum can be divided into two parts based on whether the radiated photon energy is above some threshold (hard radiation) or below threshold (soft radiation). The soft-radiative term still cancels out the non-radiative term's infinity, but the hard-radiative term is finite. The threshold is set so that experiments studying the decay will not be able to distinguish the soft-radiative process from the non-radiative process.

In KTeV, the threshold is set at a two-photon invariant mass of $1\text{MeV}/c^2$. For kaon decay, this means there is a 27.87% chance of simulating $K_L \rightarrow e^+e^-\gamma\gamma$ instead of $K_L \rightarrow e^+e^-\gamma$. The sample of simulated $K_L \rightarrow e^+e^-\gamma$ therefore includes events where $K_L \rightarrow e^+e^-\gamma\gamma$ was simulated. I keep track of the number generated of each type by making histograms of $M_{\gamma\gamma}$ and E_γ^* during simulation.

Besides the QED processes of Dalitz decay, the $K_L\gamma\gamma$ form factor is simulated. The form factor used is the same as in Equation 1.3. using the PDG value $\alpha_{K^*} = -0.28$ [37].

After these methods determine relative momenta for the three or four decay products, the entire system is rotated by a random angle, and the momenta are boosted into the lab frame.

6.1.2.3 $K_L \rightarrow \pi^0 e^+ e^-$

The decay $K_L \rightarrow \pi^0 e^+ e^-$ is randomly generated using only uniform phase space constraints. This is appropriate because the sensitivity of this search for $K_L \rightarrow \pi^0 e^+ e^-$ is such that only non-standard-model physics could allow events to be found, and the phase-space model offers the least bias in such a search.

However, in order to see what implications can be drawn about the standard model from our search, an alternative simulation of $K_L \rightarrow \pi^0 e^+ e^-$ is made using the vector model for the direct CP violating part of the decay. This is done assuming the

decay proceeds through a $K\pi^0\gamma^*$ vertex. This changes the form of the partial width from

$$\frac{d\Gamma}{dx dy} \sim \lambda^{1/2}(1, x_\pi, x)$$

for phase space to

$$\frac{d\Gamma}{dx dy} \sim \lambda^{3/2}(1, x_\pi, x)(1 - y^2)f^2(x).$$

Here, $\lambda(a, b, c) = a^2 + b^2 + c^2 - 2(ab + bc + ca)$, $x = (p_{e+} + p_{e-})^2/m_K^2$, $x_\pi = m_{\pi^0}^2/m_K^2$,

$$y = \frac{2p_K \cdot (p_{e+} - p_{e-})}{M_K^2 \lambda^{1/2}(1, x, x_\pi)},$$

and $f(x) = (1 + 0.03x/x_\pi)$. The form factor $f(x)$ is obtained by relating the K_{e3} form factor f_+ [37][47].

6.2 Decay Products

After the kaon decay has been fully generated, the next step is to trace the decay products through the detector. The flight of each such particle is traced from the decay vertex until the particle either exits the detector region or is absorbed in the detector material. Because the only decay products being simulated in this analysis are electrons, positrons, and photons, the decay of these particles is not a concern. As they travel through matter, these particles can “create” other electrons or photons that move through the detector. Finally, the energy deposited and digitized in the detector is simulated for each particle.

6.2.1 Tracing

A decay product is traced from its point of creation along its direction of travel. Usually it begins in the vacuum region and moves downstream. Then it is traced to the Z of each RC. If X and Y indicate that an electron or photon hit an RC, then the

energy of that particle (maximum 1 GeV) is added to the energy deposited in that RC. That particle is no longer traced. If the RC is a veto condition for the trigger, and the energy deposited is above the veto threshold, then the event is rejected. Tracing for the particle ends if X and Y indicate that it passed outside the RC.

The particle is then traced to the vacuum window. If it hits outside the window, it is no longer traced. A photon that hits the window can undergo e^+e^- pair conversion (see section 6.2.2.2), while an electron undergoes multiple scattering (section 6.2.2.3) and possibly bremsstrahlung (section 6.2.2.4). The particle is then traced through the air gap downstream of the window, where the same material interactions may occur.

The particle is then traced through DC1 and DC2. Material interactions can occur in DC windows or buffer gas. The probability of hitting a sense or field wire, per DC, is 0.7% and 9.9%, respectively. Material interactions in field wires is either in aluminum (59%) or in gold (41%). The X and Y position of the particle is recorded at each DC plane. Material interactions can occur in the helium bags; each bag is divided into four “slices” for simulating where material interactions occurs. Tracing the particle through SA2 occurs just as with the RCs. After tracing to the nominal magnet “kick” plane, the particle can be lost if it goes outside the magnet aperture. If it is an electron, then it is deflected in the $\pm X$ direction by an amount based on the X, Y map of field strength.

Tracing through DC3 and DC4 is similar to tracing through DC1 and DC2. The particle can be lost if it goes outside SA3 or SA4, and the event can be rejected if sufficient energy is deposited in a veto counter.

When the particle is traced through the TRD system, there can be material interactions at each TRD chamber. The radiation length of each chamber is $0.0176X_0$, or $0.0051X_0$ in the beam-hole regions.

Next, the particle is traced through the VV' trigger bank. If the particle is an electron, or a photon that converts in VV', then the counter's digital response is

simulated. The location of the particle determines exactly which counters are hit in each plane. The particle may hit a crack in either plane and so only hit one plane, in which case the radiation length is halved for purposes of material interactions. A TDC value for each struck counter is simulated using the particle's position within the counter, and a trigger latch bit is set for the counter.

The particle is then traced to the CIA, where the simulation is the same as for the SAs and the RCs. Next, the particle is traced to the CA. Electrons and photons that hit any part of the CA deposit all of their energy, smeared with a gaussian with a sigma of 25%. Such particles are no longer traced. The particle is then traced to the front face of the CsI calorimeter. If it is not outside the array or inside a beam hole, then its energy, X , and Y are noted for later use. Electron and photons are stopped if they hit the calorimeter.

A particles that misses the CsI can be traced to the Back Anti but this analysis does not use the BA, so its simulation is not discussed here. The particle could in principle be traced to the muon counters, but electrons and photons are stopped at the iron wall first. The only way the muon veto can be on in MC is through accidental overlay (section 6.3.1). The Hadron Anti is the other detector downstream of the CsI that is used in the 2ENCLUS trigger. Its response is simulated for pion showers in the CsI, but not for electron or photon showers.

6.2.2 Material Interactions

When a particle is in vacuum, tracing its path is a simple matter of projecting its direction vector to the next detector element. But when it moves through matter, several effects may have to be simulated. Because the particles being simulated here are high-energy photons and electrons, the effects that need to be simulated are pair production, multiple scattering, and bremsstrahlung.

6.2.2.1 Measuring Material

However, all three of these effects are sensitive to the thickness in radiation lengths of detector elements. The radiation lengths of most elements are estimated by determining the composition and measuring the thickness, either by surveying or using a ruler.

A novel method was employed to measure the material between the vacuum vessel and DC1. The method is described in full in [50], but can be summarized as follows. This material includes the vacuum window, four mylar windows, an air gap, and two enclosed gas buffers. This region is important because a large amount of material has to exist here (the vacuum window), and because pairs here may mimic decay products, without the opportunity to use lack of DC1 activity to reject them. The measurement technique was to take a special run with the analysis magnet turned off. $K \rightarrow \pi^0 \pi^0 \pi^0$ decays were selected by reconstructing them using only calorimeter data. If a photon from a pion decay produced an e^+e^- pair, then the pair would stay close together and probably produce only one spectrometer track and one calorimeter cluster. Then the probability of conversion upstream of DC1 could be measured by dividing the number of tracks by the number of photons observed. By doing so, the radiation thickness was measured to be $(3.55 \pm 0.17) \times 10^{-3} X_0$.

Table 6.1 gives the number of radiation lengths used in the simulation for some of the parts of the detector.

6.2.2.2 Pair Production

A photon passing through matter may convert into a positron-electron pair. In the simulation, this can occur if the photon energy is more than 100 MeV. The probability of conversion for high-energy photons in a thin sheet of material is the radiation length multiplied by 7/9 [37]. The conversion point is evenly and randomly distributed through

Table 6.1: Materials in path of decay products. “DC1U” is material between the air gap and 2 cm upstream of wires in DC1. Helium bag 2 sits between DC2 and DC2.

Part	Composition	Radiation Length ($X/X_0 \times 10^{-3}$)
Vacuum Window	mylar, kevlar	1.56
Air Gap	air, mylar	1.47
“DC1U”	mostly mylar	0.50
DC field wire	aluminum	0.885
DC field wire	gold	0.615
DC sense wire	tungsten	5.700
Helium Bag 2	helium	1.62
TRD chamber	(sect. 2.2.3.1)	17.6
TRD beam hole	(sect. 2.2.3.1)	5.1

the Z of the material. The photon's energy is distributed between the two new particles according to the Bethe-Heitler spectrum:

$$P(e_+, e_-) \sim e_+^2 + \left(\frac{2}{3} - \frac{1}{9Z}\right)e_+e_- + e_-^2,$$

where $Z = 3.74059$ and $e_{\pm} = E_{\pm}/E_{\gamma}$. Each new particle has the same direction as the photon, but offset by a small angle (θ) and rotated by a random angle around the photon's direction. The angle θ is a function of the photon's energy, the new particle's energy, and the elemental composition of the surrounding matter, the function to simulate θ comes from the EGS4 code [48].




Figure 6.5: Momentum times multiple-scattering angle projected onto plane, for e^+ and e^- at vacuum window. The fit to a gaussian goes from -1.4 mrad GeV/c to $+1.4$ mrad GeV/c, a region which includes 98% of the distribution.

6.2.2.3 Multiple Scattering

Multiple scattering is simulated for every charged particle as it passes through matter. The particle's direction of travel is altered at the point where it exits the material, deflected by a small angle θ and rotated by a random angle around the former direction. The space angle θ is inversely proportional to momentum; Figure 6.5 plots the product of plane angle and momentum. The fit is made for comparison to the PDG's

expression for the width of this distribution [37]:

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

which predicts a width of 0.405 mrad GeV/c for the vacuum window ($X/X_0 = 1.56 \times 10^{-3}$).

6.2.2.4 Bremsstrahlung

Bremsstrahlung by electrons or positrons in matter is of particular importance in this analysis because it can cause $K_L \rightarrow e^+ e^- \gamma$ decays to mimic and be background to $K_L \rightarrow e^+ e^- \gamma \gamma$ decays. Anytime a simulated electron (or positron) passes through a matter segment of the detector, there is a probability that it will radiate one photon. Multiple scattering is also simulated for the electron, upstream and downstream of the radiation point.

For the energy scales under consideration, the differential cross section is [49]:

$$d\sigma_{k,\theta_0,\theta,\phi} = \frac{Z_{\text{eff}}^2}{137} \left(\frac{r_0}{2\pi}\right)^2 [1 - F(q, Z_{\text{eff}})] \frac{dk}{k} \frac{p}{p_0} \frac{d\Omega_k d\Omega_p}{q^4} \left\{ \frac{p^2 \sin^2 \theta (4E_0^2 - q^2)}{(E - p \cos \theta)^2} + \frac{p_0^2 \sin^2 \theta_0 (4E^2 - q^2)}{(E_0 - p_0 \cos \theta_0)^2} \right. \\ \left. - \frac{2pp_0 \sin \theta \sin \theta_0 \cos \phi (4EE_0 - q^2)}{(E - p \cos \theta)(E_0 - p_0 \cos \theta_0)} + \frac{2k^2 (p^2 \sin^2 \theta + p_0^2 \sin^2 \theta_0 - 2pp_0 \sin \theta \sin \theta_0 \cos \phi)}{(E - p \cos \theta)(E_0 - p_0 \cos \theta_0)} \right\}.$$

The variables used here are as follows:

- E_0, p_0 are the initial energy, momentum of the electron.
- E, p are the final energy, momentum of the electron.
- k is the energy of the photon.
- θ_0 is the angle between the initial electron and the photon.
- θ is the angle between the final electron and the photon, a.k.a. opening angle.
- ϕ is the angle between the initial-electron/photon plane and the final-electron/photon plane.

- $q^2 = p^2 + p_0^2 + k^2 - 2p_0k \cos \theta_0 + 2pk \cos \theta - 2p_0p(\cos \theta \cos \theta_0 + \sin \theta \sin \theta_0 \cos \phi)$.
- $F(q, Z_{\text{eff}}) = 1 / \left[1 + \left(111q/Z_{\text{eff}}^{1/3} \right)^2 \right]$, the atomic form factor, assuming complete screening.
- $Z_{\text{eff}} = (\sum_i P_i Z_i (Z_i + 1))^{1/2}$, the effective atomic number. P_i is the relative fraction of atom i in the material.
- r_0 is the classical electron radius: $e^2/m_e c^2 = 2.82 \times 10^{-15} \text{ m}$.

This formula for $d\sigma_{k,\theta_0,\theta,\phi}$ assumes the Born approximation holds: that $(2\pi Z_{\text{eff}}/137\beta_0)$ and $(2\pi Z_{\text{eff}}/137\beta) \ll 1$. If one assumes all particles involved are ultra-relativistic ($E, E_0, k \gg m_e$), and uses the form factor given, and integrates over the angles, the differential cross section becomes:

$$d\sigma_k = \frac{4Z^2 r_0^2}{137} \frac{dk}{k} \left\{ \left[1 + \left(\frac{E}{E_0} \right)^2 - \frac{2}{3} \frac{E}{E_0} \right] \zeta + \frac{1}{9} \frac{E}{E_0} \right\},$$

where $\zeta = \ln(183/Z_{\text{eff}}^{1/3})$. From $d\sigma_k$, one can integrate over k from some threshold energy, k_0 , to infinity, and over some thickness X to get the total probability of radiation:

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

This $P(k > k_0)$ is used for each simulated electron with the radiated photon energy threshold at 1/1000 of the incident electron's energy. Table 6.2 gives the probability of radiation for some detector elements.

If this probability indicates that a bremsstrahlung event is to be simulated, then one of the above partial cross sections is used to generate k , θ , θ_0 , and ϕ . E is obtained from energy conservation, $E_0 = k + E$. The simpler $d\sigma_k$ is used for bremsstrahlungs where the opening angle is expected to be unimportant. There are three such cases. One is for decays which are not $K_L \rightarrow e^+ e^- \gamma$. The second is when the initial electron has less than 1.5 GeV of energy, as occurs in 7.66% of electrons in the vacuum window in $K_L \rightarrow e^+ e^- \gamma$. When the energy is so low, the efficiency to detect the electron or its

Table 6.2: Materials in path of decay products. See table 6.1 for more information. Prob. is the probability of bremsstrahlung being simulated.

Part	Z_{eff}	ζ	Prob.
Vacuum Window	5.556	4.639	0.01313
Air Gap	7.799	4.525	0.01238
“DC1U”	5.737	4.627	0.00421
DC field wire	13.501	4.342	0.00745
DC field wire	79.555	3.751	0.00518
DC sense wire	74.554	3.772	0.04804
Helium Bag 2	2.452	4.910	0.01364
TRD chamber	3.915	4.755	0.14816
TRD beam hole	3.915	4.755	0.04293

radiated photon is very small. The third case is when the radiation occurs downstream of the analysis magnet. In this case, the photon and electron are not split apart by the magnet and their calorimeter clusters will almost entirely overlap. In addition, the majority of bremsstrahlung occurs in the relatively thick TRD chambers, and using the simpler $d\sigma_k$ here saves a significant amount of computing time.

Thus the more complex $d\sigma_{k,\theta_0,\theta,\phi}$ is used for generating radiation after $K_L \rightarrow e^+e^-\gamma$ decays, for energetic electrons in matter upstream of the magnet. Figures 6.6 and 6.7 show distributions of θ and k for all electron energies for bremsstrahlung in material upstream of DC1 sense wires (thus including the vacuum window, the air gap, DC1 window membranes, and DC1 sense wires). Note that both θ and k tend to be quite small. Indeed, 69.9% of the photons here have $k < 1.0$ GeV, and are thus unlikely to produce clusters that the HCC will count.

6.2.3 Digitization

After the simulation has traced the decay products through the detector, it goes back and determines in detail what the particles did in detector elements and what the digital response of the detector was. The main areas of concern are the spectrometer (DCs) and the calorimeter (CsI). Digitizing other detector components is limited. The




Figure 6.6: Log of the bremsstrahlung opening angle θ .




Figure 6.7: Energy of the bremsstrahlung photon.

VV' system was already mentioned in section 6.2.1. The TRDs are essentially not simulated at all. Instead, TRD-related acceptance issues are studied using decay modes with very low pion contamination. The remaining, veto detectors are simulated with a simple threshold on the energy deposited.

6.2.3.1 Drift Chambers

To simulate the DC response, each plane is considered in turn. The X and Y where each simulated particle intercepts the DC plane is determined. If X and Y put the particle outside the DC, then the particle is ignored for this plane. The closest wire in the plane is found. The distance to that wire is then randomly smeared by a Gaussian function with one standard deviation equal to the position resolution of that plane. There is a global resolution value for each plane, broken down by about 40 run ranges. This value is multiplied by a correction factor that is based on hit region — in the neutral beam, in same X/Y band as the beam, or away from the beam — and is broken down into about 20 run ranges.

Next, this smeared position is converted into a TDC value in several steps. A small distance is added to account for discrete ionization of DC gas, which can add to SOD. Then the distance is multiplied by a correction factor that is broken down like the previous resolution correction. The distance is converted to time linearly, i.e. $t = t_0 + md$, where t_0 and m are functions of wire and plane number. Additional time is added to reflect propagation time from the hit position along the sense wire to the pre-amp electronics. Some hits have additional time added to reflect a high-side tail that appears in the SOD distribution; the probability of a hit receiving this treatment is different for regions on each DC wire and is broken into Winter and Summer run periods.

The digitized hit can then be eliminated to simulate single-hit inefficiency. The map of inefficiencies is broken down by region, DC wire, and Winter/Summer period.

The hit can produce a delta ray. The probability is proportional to how much of the hexagonal sense cell is traversed by the track, but can be as large as 54.7%. If a delta ray occurs, its range and direction are simulated. The position where it ranges out is treated as an additional DC hit, with position smearing and time adjustments.

After all DC hits have been found, the finite pulse width of the inputs to the TDCs is simulated. Pulses widths of about 45 ns are given to each hit, modeled on the observed widths. Any additional hits on that wire within the pulse width after a hit are removed from the list of hits.

6.2.3.2 Calorimeter

The simulation looks at each particle that hit the CsI calorimeter in turn, and distributes its energy to CsI blocks. The energy of the particle that goes into the simulated cluster is smeared to simulate detector resolution. The sigma of the smearing depends on CsI block size, whether the block is on the outer edge of the array, whether the block is near a beam hole, and whether the energy is more or less than 12 GeV. Some energy is subtracted to simulate an observed low-end tail.

Next, a shower is selected for the particle. In this context, a shower is a profile of energy distribution to the 13×13 array of small CsI blocks around where a particle hits. Large blocks are treated as the sum of four small blocks. Showers are generated ahead of time with a GEANT simulation and stored on disk. The showers are binned by energy (lower edges at 0, 2, 4, 8, 16 and 32 GeV), particle type (electron or photon), and position of the particle within the seed block. There are 25 bins in X and 25 bins in Y, spread over one octant of a block. There are a number of showers in each bin, which are used in turn by particles requiring them. The intra-block position is determined by projecting the particle along its direction of travel to a depth within the CsI that depends on energy, block size, and intra-block position at the CsI face.

Next, the selected shower is used to add energy to CsI blocks. The shower is

rotated or reflected for the octant of the block hit. The shower energies are scaled by the particle energy. The shower also contains the longitudinal (Z) distribution of energy, divided into 25 bins. The energy assigned to blocks in the shower is the sum of the product of the energy in each bin and a measured non-uniformity factor at that Z . Parts of the shower for which no blocks exist are lost.

There is an additional shower library for “sneaky” energy — energy from a cluster near a beam hole that crosses the beam hole. The sneaky showers are divided into 22 position bins, one for each block within four blocks of a hole and in one octant extending from the hole. The sneaky showers add energy to crystals on the opposite side of the beam hole to the cluster being simulated.

However, in reality clusters deposit energy beyond a 13×13 block region. This is simulated by adding energy to blocks within a 27×27 block region (but outside the 13×13 region) using an exponential transverse distribution. If sneaky energy has already been added to a block, this modification is not made to that block.

The digital response of the DPMT to the energy in each crystal is then simulated. A lookup table of pulse shapes is used to determine the energy in each DPMT “slice.” Some smearing due to photostatistics is allowed in each slice.

6.3 Accidental Overlay

In any given bucket, there can be some amount of accidental activity in the detector, activity that is not caused by the kaon decay products. This is simulated by superimposing, or overlaying, a real event from the detector on every MC event. The real event is one where an accidental counter fired, as described in section 2.2.5.2. The accidental events are stored during data taking, and then used by the simulation afterwards.

6.3.1 Simulating Run Number

The accidental events also determine which run and spill numbers are simulated. Even before the kaon is traced, an accidental event is selected, thus allowing run-dependent calibration constants to be chosen. Accidental runs are binned by run number in disk files. A “multiple run” (MRN) file determines which accidental file to access. The MRN file gives the fraction of generated events that should use each accidental file. Each MRN fraction is proportional to an estimated K_L flux for each run, as found using a separate analysis of $K \rightarrow \pi^0 \pi^0 \pi^0_{Dalitz}$ decays.

The accidental files created during data taking have some problems. The number of events in each is proportional to the accidental rate for a given run, which is not necessarily proportional to the kaon rate. Also, the events in an accidental file are time-ordered. A given simulation might then end up using only the first part of the file and thus not simulating conditions at the end of the run. This problem could be fixed if the order of the events in each accidental file could be scrambled. However, then accidental files from runs where the accidental rate was proportionally higher than the kaon rate would still contain a fraction of unused events, a waste of computing resources.

To get around these problems, we take several steps. From each “raw” accidental file we take some events selected by a random prescale factor r_i , where

$$r_i = \frac{A_{max}}{N_{max}} \frac{N_i}{A_i}.$$

The run number is i , A_i is the number of “raw” accidentals for run i , and N_i is the number of K_L (same as used for the MRN file). A_{max} is the number of accidentals on tape for the run with the maximum ratio A_i/N_i , while N_{max} is the K_L flux for that run.

This procedure does not re-order the accidentals, so the events remain ordered in time. To prevent biasing the accidentals used towards early spills or away from late spills the number of events generated for each decay mode is such that all accidental events

are used the same number of times. This is done by making the number generated an integer multiple of the number of accidentals on disk.

6.3.2 Overlaying

The data overlayed is of two types, ADC and TDC.

For an ADC channel, overlay is essentially a matter of adding the simulated ADC value to the accidental event's ADC value. The threshold for readout of DPMTs is only two counts for accidental events, so virtually any accidental energy present gets added to simulated activity in the calorimeter. The readout thresholds of ADCs for veto counters are not lower for accidental events. This means that the scenario where a low energy decay product is traced to a veto counter and low level accidental activity pushes the counter's response above veto threshold may not be well simulated.

Overlay on most TDC channels is done by adding the accidental TDC hits to the list of simulated TDC hits. Any hit coming within 10 ns after a previous hit on the same channel is removed. If the TDC channel is used for a DC wire, a random pulse-width is assigned to simulated hits as mentioned in section 6.2.3.1, and in the same manner to accidental hits. If two overlapping hits are from accidentals, then the later hit is not removed. However, an accidental hit may cancel a later, overlapping, simulated DC hit.

6.4 Trigger

The MC imposes trigger logic conditions on the event identical to those used on data. Levels 1, 2, and 3 are all simulated. Level 1 source signals from veto counters are turned on if energy in the counters is above threshold. Simulations of the DC OR, ETOT, and HCC subsystems reproduce their hardware logic using the TDC or DPMT data. In the case of ETOT, its efficiency near and above its energy threshold is simulated. These inefficient outputs of ETOT are passed to the HCC logic, thus simulating inefficiency in the HCC.

To simulate the Kumquats, DC TDC values are used to find hits in the Kumquat in-time window and to set the Kumquat latch bits. These bits are then used to get a hit count in the same way as the hardware. This Kumquat logic simulation is the same code used in online monitoring of Kumquat performance.

The online monitoring code is also used to simulate Banana response. A Banana TDC is simulated by converting the first DC TDC hit after the start of the banana gate, or the last hit before the gate, into Banana TDC counts. These times are then correlated to find in-time hits and counted in the same way as the Banana hardware. One discrepancy in the simulation is that the Banana TDCs see a wider time window (819 ns) than the DC TDCs (512 ns). Because the MC Bananas only get their hits from the smaller window, they cannot simulate the case of an accidental, out-of-time pair where one hit is in-time and the other hit is well out-of-time and the pair is rejected. This corresponds to a pair of triangles of data in Figure 3.12 above and to the right of the in-time-pair region. Because the out-of-time hit is not simulated, the “pair”, which looks like an isolated hit, is accepted. This means that the acceptance of the simulated trigger is a bit larger than the real trigger, at the L2 stage. However, in general the hit-count requirements on Bananas are loose enough that the overall effect is minimal.

The Level 3 logic is simulated using the same code and requirements as during data taking. The only calibration constants used are those that were available during data taking.

Chapter 7

Reduction

As mentioned in section 4.3, the data for this analysis is spread over 950 raw-data DLTs. To reduce this to a manageable size, there were two filtering stages. The first primarily used trigger type and the second used loose, relatively simple analysis cuts.

7.1 Split

The E799 data split was a project that divided the raw data DLTs by trigger type or simple analysis cuts. For this analysis, the split output DLTs are from the “2ENEUT” set. There are 98 such DLTs, and they contain data events with beam triggers 01 (2ENCLUS) and 08 (NEUTRAL). The NEUTRAL trigger’s main requirement is six or more clusters in the HCC at L2, and is not used in this analysis. The 2ENEUT tapes contain 129,360,401 2ENCLUS triggers.

No split is necessary for MC data, because I only simulate the 2ENCLUS trigger.

7.2 Crunch

The 2E-NCLUS “crunch” project reduced the data volume even more. The crunch worked in two ways: writing output at higher tape density and filtering using loose analysis cuts. Tape density was increased by writing the output with DLT7000 drives with DLT IV media. The output tapes could hold as much as 42.8 Gb (minus some safety/inefficiency factor), as opposed to about 12 Gb for the input tapes, which were

written with DLT4000 drives on DLT III media. The filter aspect of the crunch is described below.

MC data was crunched as well. The output generally went to disk files, but the filtering was identical.

7.2.1 Reconstruction

Reconstruction of decays worked the same way as described in chapter 5, with some minor differences. The calibration constants used were those that were available on April 2, 1998, whereas the final reconstruction uses calibration constants available on September 1, 1999. New constants developed over those 18 months include improved distance-time maps for the DCs and the cluster energy-nonlinearity corrections.

The crunch used version 4.08 of ktevana, whereas chapter 5 describes version 5.00. The crunch thus had some minor code differences. Some code differences were bug fixes, such as for the track reconstruction not correctly handling shared hits between tracks. Another code difference was in cluster energy reconstruction, where the crunch did not have correction factors for sneaky energy, energy nonlinearity, or run-to-run variations.

7.2.2 Cuts

The crunch identified and tagged ten different physics modes within the 2EN-CLUS triggers. These are summarized in Table 7.1

The cuts listed in Table 7.1 have the following definitions:

- L3*, Level 3 tagged the event as 2e-nclus or 3pi0dal.
- NTRK, number of tracks.
- T3FVTX4, a four-track vertex can be reconstructed.
- T3FVTX, a two-track vertex can be reconstructed.

Table 7.1: Summary of filter tags and physics. All decays are from kaons, unless noted. Tagged is the fraction of events on 2ENEUT tapes accepted by the cuts.

Name	Cuts	Tagged	Physics
L3RAND	L3 Random Accept Tag	0.33%	systematic studies
PI0TEE	$L3^* \cdot (NCLUS \geq 6)$ $\cdot EOP^* \cdot (M_{ee} \geq 0.07 \text{ GeV})$	4.20%	$\pi^0 \rightarrow e^+e^-$
3T6SC	$L3^* \cdot (NCLUS \geq 6)$ $\cdot EOP^* \cdot (NTRK \geq 3)$ $\cdot (E/p(3) \geq 0.9)$	0.81%	$\pi^0 \rightarrow e^+e^-$ BG
4TRACK	$L3^* \cdot EOP^* \cdot T3FVTX4$	3.73%	$e^+e^-e^+e^-$ $\pi^0 \rightarrow e^+e^-e^+e^-$ 2 or $3\pi^0$, with 2 or $3\pi^0 \rightarrow e^+e^-\gamma$
2T8C	$L3^* \cdot T3FVTX \cdot EOP^*$ $\cdot \overline{T3FVTX4}$ $\cdot (NHCLUS \geq 8)$	1.21%	$\pi^0\pi^0\pi^0, \pi^0 \rightarrow e^+e^-\gamma\gamma$ $\pi^0\pi^0\pi^0\gamma, \pi^0 \rightarrow e^+e^-\gamma$
2PI0EE	$L3^* \cdot T3FVTX \cdot EOP^*$ $\cdot (NHCLUS = 6)$ $\cdot (P_T^2 < 0.001 \text{ GeV}^2)$ $\cdot (M > 0.44 \text{ GeV})$	3.01%	$\pi^0\pi^0e^+e^-$
3T7C	$L3^* \cdot T3FVTX \cdot EOP^*$ $\cdot (NTRK = 3)$ $\cdot (NHCLUS \geq 7)$	0.64%	Background to $3\pi_{DD}^0$ with a missing track.
3PI0D	$L3^* \cdot T3FVTX \cdot EOP^*$ $\cdot (NHCLUS = 7)$	20.6%	$\pi^0\pi^0\pi^0, \pi^0 \rightarrow e^+e^-\gamma$
EEGGG	$L3^* \cdot T3FVTX \cdot EOP^*$ $\cdot (NHCLUS = 5)$ $\cdot (Vtx \subset \text{CsIhole} \pm 5 \text{ cm})$ $\cdot (M > 0.380 \text{ GeV})$	3.19 %	$\pi^0e^+e^-\gamma$ $\pi^0\pi^0, \pi^0 \rightarrow e^+e^-\gamma$ $\pi^0\gamma\gamma, \pi^0 \rightarrow e^+e^-\gamma$
EEGG	$L3^* \cdot T3FVTX \cdot EOP^*$ $\cdot (NHCLUS = 4)$ $\cdot (Vtx \subset \text{CsIhole} \pm 5 \text{ cm})$ $\cdot (P_T^2 < 0.001 \text{ GeV}^2)$	1.22%	$\pi^0e^+e^-$ $e^+e^-\gamma\gamma$

- EOP*, at least two tracks must have $E/p > 0.9$.
- NCLUS, the software cluster count, which was not discussed in chapter 5. It is rather like the clustering described there, but with lower thresholds.
- NHCLUS, number of clusters, as described in chapter 5.
- P_T^2 , transverse momentum squared.
- $(Vtx \subset Cshole \pm 5 \text{ cm})$, a line drawn from the target through the vertex position intersects the calorimeter face within 5 cm of a beam hole.

In addition to the cuts listed there, a number of 2ENEUT data were rejected because of questionable data quality. Only data from runs where the detector was in a normal E799 configuration were used, and only from runs that had more than 5,000 $K \rightarrow \pi^0 \pi^0 \pi_{Dalitz}^0$ events identified at Level 3. Runs with less than 5,000 $K \rightarrow \pi^0 \pi^0 \pi_{Dalitz}^0$ make up about 1% of the full data set, and were typically cut short because of severe detector or beam problems.

7.2.3 Output

The crunch output was written to 12 DLTs. These were then split by crunch tag onto separate DLTs. There are 1,723,166 events identified as EEGG (1,015,703 in Winter and 707,463 in Summer), with a data volume of 10,998,847 Kb, which fit onto one DLT IV or two DLT III-XT. There are 4,775,474 events identified as EEGGG (2,763,666 in Winter and 2,011,808 in Summer), with a data volume of 31,824,901 Kb, which fit onto two DLT IV or three DLT III-XT.

Chapter 8

Normalization

This chapter describes the measurement of the flux, the number of K_L decays in the fiducial region during E799. This flux provides normalization for both the study of $K_L \rightarrow e^+e^-\gamma\gamma$ and the search for $K_L \rightarrow \pi^0e^+e^-$. In addition, the normalization mode ($K \rightarrow \pi^0\pi_{Dalitz}^0$) decays are relatively abundant and are used for evaluating the accuracy of the Monte Carlo simulation with respect to the data. A number of cuts were developed by looking at $K \rightarrow \pi^0\pi_{Dalitz}^0$; these cuts were later applied to signal candidate events.

8.1 Why $K \rightarrow \pi^0\pi_{Dalitz}^0$?

This decay is used for normalization because its branching ratio is relatively large and well known, and because it shares a number of characteristics with the signal modes. The BR is the $K_L \rightarrow \pi^0\pi^0$ branching ratio multiplied by the probability of *either* pion undergoing Dalitz decay:

$$\begin{aligned}\text{BR}(K_L \rightarrow \pi^0\pi_{Dalitz}^0) &= \text{BR}(K_L \rightarrow \pi^0\pi^0) \times 2 \times \text{BR}(\pi \rightarrow e^+e^-\gamma) \times \text{BR}(\pi \rightarrow \gamma\gamma) \\ &= (9.36 \pm 0.02) \times 10^{-4} \times 2 \times (1.198 \pm 0.032) \times 10^{-2} \times 0.98798 \\ &= (2.216 \pm 0.076) \times 10^{-5}\end{aligned}$$

The PDG values for the BRs are used [37]. The BR uncertainties are added in quadrature to get $\pm 3.42\%$. The final state of $K \rightarrow \pi^0\pi_{Dalitz}^0$ is identical to the final states of

$K_L \rightarrow e^+e^-\gamma\gamma$ and $K_L \rightarrow \pi^0e^+e^-$, with one extra photon,. Also, two of the photons come from one pion, just as in $K_L \rightarrow \pi^0e^+e^-$.

There are some other K_L decays that could be used for normalization. One is $K_L \rightarrow \pi^+\pi^-\pi^0$, which has similar decay kinematics to $K_L \rightarrow \pi^0e^+e^-$ and a large BR known to 1.6% [37]. However, the interactions of charged pions in the calorimeter are different than electron interactions. The acceptance calculation for $K_L \rightarrow \pi^+\pi^-\pi^0$ in the B01 trigger, which has a high calorimeter-energy threshold, is highly dependent on the quality of the pion shower simulation. As this quality was in doubt, this mode was rejected.

Another possible normalization mode is $K_L \rightarrow e^+e^-\gamma$. This is especially appropriate for normalizing $K_L \rightarrow e^+e^-\gamma\gamma$, because they both go through a $K_L\gamma\gamma^*$ vertex and the latter mode is predicted as a fraction of the former. However, $K_L \rightarrow e^+e^-\gamma$ is not accepted by the B01 trigger without the addition of an externally-radiated or accidental photon. In addition, while the BR (9.1×10^{-6}) is comparable to that of $K_L \rightarrow \pi^0\pi_{Dalitz}^0$, it is only known to 5.5% of itself [37].

8.2 Simulated Events

A total of 28,122,680 $K \rightarrow \pi^0\pi_{Dalitz}^0$ decays were generated between $Z = 90$ m and $Z = 160$ m with momentum between 20 and 220 GeV/c, as described in chapter 6. I make no attempt to measure flux in periods when the detector had known problems. Therefore, a “bad-spill cut” is made at the generation stage. This cut is discussed in section 8.3.1. Of the generated decays, 90.72% are in “good” (=not-bad) spills. Table 8.1 gives the MC numbers broken down by Winter and Summer periods.

The goal of normalization for $K_L \rightarrow \pi^0e^+e^-$ or $K_L \rightarrow e^+e^-\gamma\gamma$ is to measure the flux of K_L decays, N_L . However, when a neutral kaon is observed to decay to two pions, there is no way to know whether the K_L or K_S component “caused” the decay. To measure N_L , all kaon decays are counted, but only the generated K_L s are used. How

this is done is described below.

Consider some period in proper time dt . The probability of a kaon decaying to $\pi^0\pi^0$ during this time is proportional to $|\mathcal{A}_S(t) + \eta_{00}\mathcal{A}_L(t)|^2$. Thus the branching ratio B_K of the mixed kaon state to two pions can be thought of as a function of t and not a constant. However, the BR of the K_L , B_L , is a constant, independent of t . B_L is a fraction of the full kaon branching ratio:

$$B_L = \frac{|\eta_{00}\mathcal{A}_L(t)|^2}{|\mathcal{A}_S(t) + \eta_{00}\mathcal{A}_L(t)|^2} B_K(t)$$

Now consider extending dt into a t bin over which $\mathcal{A}_S(t)$ and $\mathcal{A}_L(t)$ are nearly constant.

The number of decays in this bin is then proportional to the branching ratio:

$$N_L^t = \frac{|\eta_{00}\mathcal{A}_L(t)|^2}{|\mathcal{A}_S(t) + \eta_{00}\mathcal{A}_L(t)|^2} N_K^t \quad (8.1)$$

Next, suppose that one could observe and count K_L decays: K decays from the K_L component only. The number of $K_L \rightarrow \pi^0\pi^0$ decays observed in a t bin, n_L^t , would be the product of the number of K_L decays, the BR, and the acceptance, $\epsilon(t)$: $n_L^t = N_L^t B_L \epsilon(t)$. (Acceptance ϵ in a t bin should be the same for K and K_L .) The K_L flux is obtained by summing over all t bins:

$$N_L = \sum_{t \text{ bins}} \frac{n_L^t}{B_L \epsilon(t)}$$

But n_L^t is not observable, only n^t , the number of $K \rightarrow \pi^0\pi^0$ decays. However, using Equation 8.1 the flux can be written as:

$$N_L = \sum_{t \text{ bins}} \frac{n^t}{B_L \epsilon(t)} \frac{|\eta_{00}\mathcal{A}_L(t)|^2}{|\mathcal{A}_S(t) + \eta_{00}\mathcal{A}_L(t)|^2}$$

Expressing the acceptance as $\epsilon = A/G$, the number of MC events passing all cuts divided by the number of MC events generated, the flux is:

$$N_L = \sum_{t \text{ bins}} \frac{n^t G^t}{B_L A^t} \frac{|\eta_{00}\mathcal{A}_L(t)|^2}{|\mathcal{A}_S(t) + \eta_{00}\mathcal{A}_L(t)|^2}$$

For this analysis, we assume n^t/A^t is a constant with respect to t . Taking all the constants out of the sum, the flux is:

$$N_L = \frac{n}{A B_L} \sum_t G^t \frac{|\eta_{00}\mathcal{A}_L(t)|^2}{|\mathcal{A}_S(t) + \eta_{00}\mathcal{A}_L(t)|^2}$$

By considering each generated event to be in its own t bin, the sum can then be made over all generated events.

$$N_L = \frac{n}{A \text{ BR}_L} \sum_{\text{generated}} \frac{|\eta_{00}\mathcal{A}_L(t)|^2}{|\mathcal{A}_S(t) + \eta_{00}\mathcal{A}_L(t)|^2}$$

In effect, the number generated is weighted by the K_L fraction. B_L above needs to be multiplied by twice the pion-Dalitz BR times the pion- $\gamma\gamma$ BR to get the correct formula for flux using $K \rightarrow \pi^0\pi^0_{Dalitz}$, but that does not change the argument. The sum of these weights in good spills is 23,507,662.

Table 8.1: Numbers of simulated $K \rightarrow \pi^0\pi^0_{Dalitz}$. % is the number on that line divided by the number generated. K_L weight is the sum of $|\eta_{00}\mathcal{A}_L(t)|^2/|\mathcal{A}_S(t) + \eta_{00}\mathcal{A}_L(t)|^2$ in good spills. L1, L2, and L3 trigger and EEGGG filter sums are unweighted and in all spills. L3 triggers are 2e, n-clus tags only.

Stage	Winter		Summer	
	Number	%	Number	%
Generated	15998576	100.00	12124104	100.00
Good Spills	14809014	92.56	10703911	88.29
K_L weight	13645544	85.29	9862118	81.34
L1 trigger	2318134	14.49	1824625	15.05
L2 trigger	1459213	9.12	1208303	9.97
L3 trigger	1043358	6.52	820074	6.76
EEGGG Filter	513163	3.21	399943	3.30

8.3 General Cuts

This section describes cuts made to remove data of questionable quality. For the most part, the same event requirements are made on $K_L \rightarrow \pi^0 e^+ e^-$ and $K_L \rightarrow e^+ e^- \gamma\gamma$ candidates.

Any plot in this chapter that compares a spectrum in data to a spectrum in MC has the MC spectrum normalized to have the same integrated area as the data spectrum, unless noted.

8.3.1 “Bad Spills”

Events occurring during runs or spills when there were known hardware failures in the detector are rejected. Most such bad spills are identified by referring to a database of bad-spill bits. Table 8.2 lists these bits, as well as the approximate fraction of data with those bits set. The data used for these fractions are events where track reconstruction finds a good, two-track vertex, there are three photon clusters, and $M_{ee\gamma}$ is within $15 \text{ MeV}/c^2$ of the π^0 mass. There are 736,096 such events passing the EEGGG filter.

Another set of bad spills was identified but which does not appear in the database. These types of spills are listed in Table 8.3. Of the events tagged as EEGGG in the crunch, 0.54% are in this kind of bad spill in the Winter, and 1.22% are in Summer, for a total of $\sim 1.76\%$ of all data lost to this second type of bad spill cuts.

8.3.2 Reconstruction

Several cuts are made to get rid of events that can not be reconstructed as described in chapter 5. These cuts occur after the non-database bad-spill cuts, but before the database bad-spill bit cuts. The reconstruction cut requirements are, in sequence: minimum number of track candidates, ok vertex, minimum number of tracks (2), and minimum number of clusters. The events present at each stage in the $K \rightarrow \pi^0 \pi^0_{Dalitz}$ analysis are listed in Table 8.4.

It should be noted that these events undergo reconstruction twice before reaching these stages: once in the L3 trigger and once in the Crunch. Thus events are lost here only because of changed calibration constants or differences in the analysis. Events are lost at this track candidate stage because there must be at least two mutually

Table 8.2: Bad spill bits summary. The Winter and Summer columns list the percentage of events in each data set with that bad-spill bit, regardless of whether a cut is made on that bit. The data used pass loose cuts for $K_L \rightarrow \pi^0 \pi_{Dalitz}^0$. The total fraction of data lost to bad spill bits is 9.10%.

Bit	Description	Cut?	% lost of Winter	% lost of Summer
1	trigger	yes	0.00	0.00
2	DPMT pedestal exponent > 0	yes	0.82	1.04
3	Bad DPMT capacitor	only in Winter	0.60	16.90
4	Blown QIE comparator	yes	0.00	0.00
5	Misc. dead DPMT	yes	1.16	0.23
6	DPMT pedestal drifting	no	4.75	0.51
7	DPMT gain drifting	yes	0.00	0.00
8	Broken DPMT dynode	yes	0.00	0.00
9	DPMT pipeline problem	yes	0.01	0.00
10	global Csl problem	yes	0.00	0.00
11	ETOT trigger	yes	0.00	0.00
12	FERA ADC	yes	0.05	0.02
13	Drift chambers	yes	1.83	0.00
14	Misc. Vetoes	yes	1.13	0.00
15	VV'	yes	0.00	0.00
16	Muon Veto/Counter	only runs < 8577	0.00	0.00
17	HCC trigger	only runs < 8245	8.60	0.51
18	Kumquat/Banana	yes	0.00	0.00
19	TRD trigger	no	0.00	0.00
20	Hyperon trigger	no	4.28	0.00
21	DAQ/L3	yes	0.00	0.00
22	non-E799/E832 run	yes	0.00	0.00
23	short run	yes	0.00	9.61
24	non-standard TRD HV	no	14.87	0.00
25	one dead TRD plane	no	0.36	13.89
26	> 1 dead TRD planes	yes	2.29	0.05
27	TRD voltage sag	no	0.08	2.82
28	severe TRD problem	yes	0.03	0.06
29	beam problem	yes	0.00	0.00
30	unused	no	0.00	0.00
31	unused	no	0.00	0.00
32	miscellaneous	no	0.11	0.06
% of each period cut			7.32	11.53
% of total in each period			57.55	42.45
% of total cut in each period			4.21	4.89

Table 8.3: Non-database bad spills.

Run	Spills	Problem
8428	0–148	CA HV
8453	120–200	CsI gain drift
8913	all	TRD chambers out
10599	55–end	one stuck DPMT mantissa, one DPMT pedestal exponent > 0
10765	all	special target-angle scan run
10790	245–end	ETOT controller failure
10904	all	TRD chambers out
10906	all	TRD chambers out
10909	all	special high intensity run
10914	all	special high intensity run

Table 8.4: Events successfully reconstructed.

Stage	Data		MC	
	Winter	Summer	Winter	Summer
EEGGG crunch tag	2763666	2226242	513163	399943
non-db bad spill	2736844	2165366	475276	351863
track candidate	2289657	1802698	425866	309844
vertex	2288236	1801195	425792	309785
two tracks	2009107	1553815	422591	306930
five clusters	1995334	1532508	422526	306885

exclusive Y track candidates. The L3 and Crunch reconstructions allowed overlap in the Y view. Also, previously only two or more tracks were required. At this stage, exactly two tracks are required. This removes a significant number of four-track background from data events, as shown in Figure 8.1. The four-track data events are dominated by $K_L \rightarrow \pi^0 \pi^0 \pi^0$ decays where four electrons appear, either through two photon pair conversions, two pion-Dalitz decays, one pair conversion and one Dalitz, or one $\pi^0 \rightarrow e^+ e^- e^+ e^-$. MC events with more than two tracks come from accidentals or simulated photon conversions.




Figure 8.1: Number of tracks in $K \rightarrow \pi^0 \pi^0_{Dalitz}$ reconstruction. Dots are data and line is MC, normalized to data. No cut has limited the number of clusters yet.

The five-cluster requirement is intended to identify $K \rightarrow \pi^0 \pi^0_{Dalitz}$ decays, which deposit two electron plus three photon clusters. In the $K_L \rightarrow \pi^0 e^+ e^-$ and $K_L \rightarrow e^+ e^- \gamma \gamma$ analyses, this stage of reconstruction requires exactly four calorimeter clusters.

8.3.3 Fiducial Region

Cuts on total momentum and vertex Z are made to insure that only data events inside the simulated region are studied. Figures 8.2 and 8.3 shows the momentum

and position resolutions in MC near the fiducial region boundaries. The reconstructed momentum is required to be between 20.3 and 216 GeV/c, several sigma from the 20–220 generation limits. Figure 8.2 also shows the ratio of reconstructed momentum in data to simulated events. The sharp drop off at low P is due to the ETOT cut (section 8.3.4). The high-P data are mostly K_S decays in the upstream part of the decay region.




Figure 8.2: (Left) Momentum resolution in simulated $K \rightarrow \pi^0 \pi^0_{Dalitz}$. Reconstructed total momentum minus generated total momentum in a sample of MC events, near fiducial limits. (Right) Reconstructed momentum for all data and MC, after all other cuts.

The resolution for two different methods of reconstructing vertex position is shown in Figure 8.3. The first method uses the spectrometer tracks, as described in section 5.2.2. The second method finds the “neutral-pion vertex”, Z_{π^0} , using the constraint on $K \rightarrow \pi^0 \pi^0_{Dalitz}$ decay imposed by the presence of the intermediate pions. The method begins by finding $M_{\gamma\gamma}(= \sqrt{2k_1 \cdot k_2})$ for each of the three possible pairings of photons. The pair that minimizes $|M_{\gamma\gamma} - M_{\pi^0}|$ is considered as coming from the non-Dalitz pion. The distance from the calorimeter to the vertex is scaled to make $M_{\gamma\gamma}$ equal to 134.9764 MeV/c², the PDG π^0 mass [37]. This scaled distance gives Z_{π^0} .

The “charged vertex” Z has much poorer resolution than Z_{π^0} in the upstream




Figure 8.3: Decay-vertex-Z resolution in simulated $K \rightarrow \pi^0 \pi^0_{Dalitz}$. Reconstructed Z minus generated Z in a sample of MC events, near fiducial limits. (Left) Reconstructed vertex uses DC tracks, (Right) reconstructed vertex uses Z_{π^0} .




Figure 8.4: Vertex Z for all data and MC, after all other cuts. (Left) DC track vertex, (Right) Z_{π^0} .

region, 2.6 meters versus 0.83 meters, as shown by Figure 8.3. That is because in that region the opening angle between tracks is quite small and small errors in measuring their directions mean large errors in reconstructing the vertex. In the downstream region, the opening angles can be larger and the resolution is better for Z than for Z_{π^0} , 0.089 meters versus 0.27 meters. The resolution of Z_{π^0} is more constant over the decay region because it uses calorimeter cluster energy information, which is insensitive to opening angle.

The vertex cut requires both Z and Z_{π^0} to be between 96 and 158 meters. These cuts are chosen with concern more for detector elements than for fiducial region boundaries (although the latter were also chosen with the former in mind). At the upstream end, the cut is several sigma from the Final Sweeper magnet and its beam pipes, which end at $Z = 93$ m. The downstream cut eliminates decays at or downstream of the vacuum window, at $Z = 158.89$ m.

Figure 8.4 shows the ratio of data to MC for the Z spectra. The slopes in the fits are discussed in section 8.5.

8.3.4 Trigger Verification

Several cuts are made to reject events that would have failed hardware trigger requirements; i.e., MC events near a poorly simulated trigger threshold. Trigger elements verified in this way include the V-bank counters, the CA, and ETOT.

For the V-bank trigger counters, the DC tracks are extrapolated to each V-bank. The event is cut if the in-time hits corresponding to each track do not meet the trigger requirements: $(V0_GE2S \cdot V1_GE1S) + (V0_GE1S \cdot V1_GE2S)$. After all other cuts, this removes $0.34 \pm 0.03\%$ of data and $0.26 \pm 0.01\%$ of MC. In addition, to check that we know the position of the beam holes in the V-banks a cut requires the minimum distance from each track to the nearest beam hole edge be more than 1 cm. See Figure 8.5 for the distributions of this distance.

For the Collar Anti, the nominal veto threshold was 12 GeV. For an analysis cut, the energy of each CA paddle is added to its neighbor's energy. The maximum of any of these sums must be less than 12 GeV. This has only a minor effect, as can be seen in Figure 8.6.




Figure 8.5: Minimum distance from track to upstream V-bank beam hole. All other cuts are made.

The nominal ETOT threshold was about 25 GeV. To verify this trigger, a sum is made of the energies in all CsI channels read out, as described in section 5.1.1. A cut is made on this sum (“etotsw”) where the Data/MC ratio begins to dip, at 33 GeV; see Figure 8.7.

8.3.5 Calorimeter

There are a couple of additional cuts to insure good quality calorimeter data. One is on the smallest distance from a cluster to a beam hole. The cut requires this distance to be more than 1.25 cm. Below this, the Data/MC ratio becomes relatively large, as Figure 8.8 shows. The outer edge of the CA extends to about 1.25 cm. It may be that




Figure 8.6: Collar Anti energy. Maximum sum of energy in any counter added to its neighbors, in GeV. All other cuts are made.




Figure 8.7: Total calorimeter energy (GeV), for ETOT verification. All other cuts are made, except for total momentum cut.

some particles that hit the CA deposit energy in the CsI behind it, and this energy is then reconstructed as clusters. However, in the simulation any particles that hit the CA stop without depositing energy in the CsI. In addition, Figure 8.8 also shows a dip in the ratio at 2.5, 5.0, and 7.5 cm, multiples of the CsI small-block size. This is caused by a bias in the position reconstruction that pushes MC clusters towards the edges of crystals. This bias may occur because cluster position is sensitive to shower profile. While MC shower profiles are generated by GEANT, both MC and data use the same, data-generated, position lookup tables. At any rate, the cut on minimum-cluster-hole distance is at the mid-point of the block size, so the position bias does not change the acceptance.

The other calorimeter (only) cut is on cluster energy. All cluster energies in an event must be more than 2.5 GeV, the point where the Data/MC ratio starts to look odd. See Figure 8.9.

8.3.6 TRD

TRD cuts are applied only to data events, because no simulation of the TRDs was attempted. The normalization data is used to estimate the acceptance of TRD cuts, which is then used in determining absolute acceptances. Of course, this acceptance correction will cancel out in calculating branching ratios.

The simpler TRD cut requires at least one TRD plane to have a hit corresponding to each DC track. Figure 8.10 shows the number of planes hit in Winter and Summer. There are 83 events in Winter that pass all non-TRD cuts but fail this cut, and 58 events in Summer.

The more interesting TRD cut is on Π_{TRD} (defined in section 5.4). Figure 8.11(A) shows Π_{TRD} for all tracks. The similarity with the distribution for electrons in Figure 5.6 suggests that pion contamination is low. This suggestion is supported by Figure 8.11(B), which shows E/p for tracks failing and passing the requirement $\Pi_{\text{TRD}} < 0.04$;




Figure 8.8: Smallest cluster-hole distance (meters). All other cuts are made.




Figure 8.9: All calorimeter cluster energies (GeV). All other cuts are made, except for ETOT verification.




Figure 8.10: Number of TRD hits per DC track. All other, non-TRD cuts are made.

the E/p for tracks that fail the Π_{TRD} cut is peaked at one, like the tracks that pass, as one would expect for electrons. This Π_{TRD} cut is made fairly loose because the kinematic cuts are very effective in eliminating K_{e3} background to $K_L \rightarrow \pi^0 e^+ e^-$. If Π_{TRD} was uniform for pions, then the single-particle rejection of this cut would be $(1.00 - 0.04)/0.04 = 24$ to 1. Because Π_{TRD} tends to be somewhat biased towards one for pions (see Figure 5.6), the actual rejection ratio is somewhat better.




Figure 8.11: (A) Π_{TRD} for all data after all other cuts. (B) E/p for tracks failing and passing Π_{TRD} cut. The passing spectrum is normalized to the failing spectrum.

Figure 8.12 shows Π_{TRD} for both tracks in all data events. The center region would be expected to contain $K_L \rightarrow \pi^+ \pi^- \pi^0$ background, while the left and lower margins would contain K_{e3} background.

In the Winter, 18,377 events pass both the number of planes and the Π_{TRD} cuts, while 953 events fail one or both cuts. In the Summer, the numbers are 13,110 passing and 880 failing. The acceptances would then be 0.951 ± 0.031 for Winter and 0.937 ± 0.032 for Summer. These are not significantly different, so a combined acceptance of 0.945 ± 0.022 is used for all MC events.




Figure 8.12: Π_{TRD} for each track in all data events, after all other cuts. The box indicates the cut.

8.3.7 E/p

The other major cut for removing non-electron tracks is on E/p , a track's cluster energy divided by the magnitude of its momentum. For relativistic electrons, all of which can be expected to shower in the CsI, E/p should be exactly 1, give or take finite resolution (and neglecting factors of c). Some pions leave only minimum ionizing energy in the CsI, while other pions shower part of the way through the CsI, so pion E/p can be between 0 and 1.

However, the Crunch stage requires $E/p > 0.9$ for two tracks, leaving very little but electrons.

Figure 8.13 compares simulated E/p to the data. The MC peak is somewhat higher than the data peak, and there is some excess data on the low side of the peak. Because the spectra are somewhat different, the E/p cut is kept loose.




Figure 8.13: E/p for each track in Data and MC events, after all other cuts.

A possible source of this difference can be seen in Figure 8.14, which shows mean

and RMS of E/p in 2.5 GeV/c momentum bins. Below 25 GeV/c, the data shows a slight nonlinearity with momentum, as well as an overall scale difference with MC. This is despite the “third” scale factor correction discussed in section 5.1.2.3, made to correct a nonlinear response with respect to energy in the CsI. Nevertheless, the overall effect is small.

It has been found that our E/p spectra can be closely fit by the sum of two Gaussians functions, as shown in figure 8.15. The fit functions can be written as:

$$P(E/p_i)_{data} = 2202 \exp \left[-\frac{1}{2} \left(\frac{E/p - 1.0001}{0.0102} \right)^2 \right] + 113 \exp \left[-\frac{1}{2} \left(\frac{E/p - 0.9958}{0.0297} \right)^2 \right]$$

$$P(E/p_i)_{MC} = 9858 \exp \left[-\frac{1}{2} \left(\frac{E/p - 1.0018}{0.0100} \right)^2 \right] + 329 \exp \left[-\frac{1}{2} \left(\frac{E/p - 1.0001}{0.0331} \right)^2 \right]$$

From these it can be seen that the MC E/p mean is higher by about 0.1%, and that the resolution is a bit worse than 1% for both. Setting the cut at $0.95 < E/p < 1.05$, about five sigma from the central peak, keeps the acceptance insensitive to the data/MC difference.

8.3.8 Spectrometer

To reduce events with ambiguity in track-cluster matching, a cut requires tracks to be more than 2.5 cm apart in X when projected to the calorimeter. However, because the analysis magnet bends the tracks in opposite directions this X separation tends to be fairly large. Therefore this cut is quite minor, removing only 27 data events and 127 MC events.

A more serious concern with spectrometer tracks is our understanding of tracks with small opening angles. These close tracks are expected to be more common in Dalitz decays ($K \rightarrow \pi^0 \pi^0_{Dalitz}$ and $K_L \rightarrow e^+ e^- \gamma \gamma$) than in the signal mode $K_L \rightarrow \pi^0 e^+ e^-$. Figure 8.16 shows the track-separation at DC1, the DC where separation tends to be smallest. An excess in MC below 1 cm can be seen. The cut requires at least 1 cm separation in both X and Y views.




Figure 8.14: E/p vs. $|\vec{p}|$ (GeV/c) for each track in Data and MC events, after all other cuts.




Figure 8.15: E/p for each track in Data and MC events, after all other cuts. The parameters are results of fitting to sum of two Gaussians.




Figure 8.16: Track separation at DC1 in X (Top) and Y (Bottom), in meters. All other cuts are made except track separation and opening angle. The y-axis for the ratio plots is from 0.8 to 1.2.

Even after this cut on relative position at DC1, there remains some misunderstanding of very close tracks. Figure 8.17 shows the track opening angle in the lab reference frame. There is a deficit of MC events at small angles. An additional cut requires this angle to be more than 2.25 mrad. The figure shows that this cut is costly in terms of $K \rightarrow \pi^0 \pi_{Dalitz}^0$ statistics. However, the cut costs much less in $K_L \rightarrow e^+ e^- \gamma \gamma$ acceptance and almost nothing in $K_L \rightarrow \pi^0 e^+ e^-$ acceptance. In addition, the uncertainty of the flux measurement using $K \rightarrow \pi^0 \pi_{Dalitz}^0$ is limited much more by systematic uncertainties than by statistics.




Figure 8.17: Track opening angle (radians). All other cuts are made.

8.3.9 Summary

The cuts described above are summarized here:

- Event in “good” spill.
- At least two separate track candidates in each spectrometer view.
- Two-track vertex can be reconstructed.

- Exactly two tracks present.
- Exactly five clusters present.
- $20.3 \text{ GeV}/c < |\vec{p}_{total}| < 216 \text{ GeV}/c$.
- $96 \text{ m} < Z_{vertex} < 158 \text{ m}$.
- $96 \text{ m} < Z_{\pi^0} < 158 \text{ m}$.
- Tracks intersect V-bank hits.
- Tracks distance to V-bank holes $> 1 \text{ cm}$.
- CA energy $< 12 \text{ GeV}$.
- Reconstructed ETOT $> 33 \text{ GeV}$.
- Cluster distances to CsI beam holes $> 1.25 \text{ cm}$.
- Cluster energies $> 2.5 \text{ GeV}$.
- Tracks match at least one TRD hit.
- $\Pi_{TRD} < 0.04$ for both tracks.
- $0.95 < E/p < 1.05$ for both tracks.
- Track separation in X at CsI $> 2.5 \text{ cm}$.
- Track separation in X at DC1 $> 1.0 \text{ cm}$.
- Track separation in Y at DC1 $> 1.0 \text{ cm}$.
- Opening angle $> 2.25 \text{ mrad}$.

8.4 Mode-Specific Cuts

This section describes cuts which differ significantly in the $K \rightarrow \pi^0 \pi^0_{Dalitz}$ analysis from the cuts in the $K_L \rightarrow \pi^0 e^+ e^-$ and $K_L \rightarrow e^+ e^- \gamma \gamma$ analyses. In general, these cuts are based on kinematic quantities.

8.4.1 $\gamma\gamma$ Mass

As mentioned in section 8.3.3, the three photon clusters and two tracks in this analysis can be put together in three pairings to correspond to $K \rightarrow \pi^0 \pi^0, \pi^0 \rightarrow \gamma\gamma + \pi^0 \rightarrow e^+ e^- \gamma$. The pairing is chosen that makes the invariant mass of the two photons

hypothesized to come from one π^0 closest to the neutral-pion mass. Figure 8.18 shows $M_{\gamma\gamma}$ for all three pairings and for the best pairing.

Figure 8.19 shows the results of fitting Gaussian functions to both data and MC spectra. The MC mean is somewhat lower at $135.41 \text{ MeV}/c^2$ than the data mean of $135.53 \pm 0.02 \text{ MeV}/c^2$. However, both means are higher than the PDG value, $M_{\pi^0} = 134.9764 \pm 0.0006 \text{ MeV}/c^2$ [37]. The reasons for both means being high are not entirely clear. The problem might stem from a downstream stream bias in vertex Z reconstruction (see Figure 8.3) or from differences between electron cluster energy reconstruction (which can be calibrated with E/p) and photon cluster energy reconstruction (which cannot be as directly calibrated). However, the fact that the MC mean $M_{\gamma\gamma}$ is almost as high as the data's suggests that much of the effect(s) is modeled.

The MC width is somewhat wider at $\sigma = 2.225 \pm 0.008 \text{ MeV}/c^2$ than the data width of $\sigma = 2.178 \pm 0.016$. To reduce the impact of these differences, the cut on on $M_{\gamma\gamma}$ is made loose. The cut is centered on the rounded-off data mean of $135.5 \text{ MeV}/c^2$, and requires $M_{\gamma\gamma}$ to be within $8 \text{ MeV}/c^2$: $|M_{\gamma\gamma} - 135.5 \text{ MeV}/c^2| < 8 \text{ MeV}/c^2$.

As discussed in section 8.3.3, $M_{\gamma\gamma}$ is used to find Z_{π^0} . The momentum of all photons is then recalculated using the new vertex position, and these new momenta are used to get the the kinematic cut variables below ($M_{ee\gamma}$, $\vec{P}_T^2$, and $M_{ee\gamma\gamma\gamma}$). This improves the resolution of these variables because the vertex resolution is better. For $M_{ee\gamma\gamma\gamma}$, the resolution is also improved by constraining the mass of one of the two daughter particles.

8.4.2 $e^+e^-\gamma$ Mass

A cut is made on the invariant mass of the Dalitz pion. Figure 8.20 shows $M_{ee\gamma}$ for the photon pairing that gives the best $M_{\gamma\gamma}$, using Z_{π^0} . The fits to Gaussians give pleasantly similar results: a data mean of $134.88 \pm 0.01 \text{ MeV}/c^2$, a MC mean of $134.89 \text{ MeV}/c^2$, a data width of $\sigma = 1.935 \pm 0.012 \text{ MeV}/c^2$, and a MC width of




Figure 8.18: Two-photon invariant mass, in GeV/c^2 , for all pairings (Left) and for best pairing (Right). Dots are data and solid line is MC. All other cuts are made.




Figure 8.19: Two-photon invariant mass, in GeV/c^2 , data (Left) and simulated $K \rightarrow \pi^0 \pi^0_{\text{Dalitz}}$ (Right). All other cuts are made.

$\sigma = 1.945 \pm 0.006 \text{ MeV}/c^2$. The cut requires $M_{ee\gamma}$ to be within 3σ of the PDG value for the pion mass, rounded off: $|M_{ee\gamma} - M_{\pi^0}| < 6 \text{ MeV}/c^2$.

8.4.3 Transverse Momentum

Transverse momentum is the component of total momentum that is perpendicular to a line drawn from the decay vertex to the target, as drawn in Figure 8.21. $\vec{P}_T^2$ is the square of transverse momentum. If the decaying particle did not scatter and if all decay products are observed, then $\vec{P}_T^2$ should be as close to zero as resolution will allow. Because $\vec{P}_T^2$ tends to be small for signal, it is often convenient to look at $\log_{10} \vec{P}_T^2$.

$\vec{P}_T^2$ is calculated using both the charged-track vertex Z and Z_{π^0} , and cuts are made on both. The two distributions are compared for data in Figure 8.22. The Z_{π^0} $\vec{P}_T^2$ has slightly better resolution, and it is used in the remaining plots in this section.




Figure 8.20: $e^+e^-\gamma$ invariant mass, in GeV/c^2 , data (Left) and simulated $K \rightarrow \pi^0\pi_{Dalitz}^0$ (Right). All other cuts are made.

The resolution on $\vec{P}_T^2$ is somewhat better for data than for simulated events. Figure 8.23 shows data peaking higher at $\vec{P}_T^2$ close to zero. Both methods of getting $\vec{P}_T^2$ show




Figure 8.21: Definition of transverse momentum, $\vec{P}_\perp$.




Figure 8.22: $\log_{10} \vec{P}_T^2$ using track-vertex Z (Dots) and Z_{π^0} (Line), measuring $\vec{P}_T^2$ in $(\text{GeV}/c)^2$. All other cuts are made.

the same effect. However, the $\vec{P}_T^2$ cuts are both loose, requiring $\vec{P}_T^2 < 0.001 (\text{GeV}/c)^2$.

Looking at higher $\vec{P}_T^2$, an interesting feature in the data appears. Figure 8.24 shows $\vec{P}_T^2$ up to $1.0 (\text{GeV}/c)^2$. The peak in the data around $0.1 (\text{GeV}/c)^2$ is from kaons that scatter and regenerate in the defining collimator. Both methods of getting $\vec{P}_T^2$ show the same effect. Decays in this peak have the high momentum and upstream decay Z characteristic of K_S decays in E799. However, because they must be deflected in the defining collimator in order to undergo regeneration there and then enter the decay region, their transverse momentum must be non-zero. Therefore, the cut at $0.001 (\text{GeV}/c)^2$ removes most of these events. To estimate the defining collimator background, I fit the data distribution over the range $-2.0 < \log_{10} \vec{P}_T^2 < -1.3$ to an exponential function. The smooth line on Figure 8.24 shows this fit extended to low $\vec{P}_T^2$. Integrating this function from $-\infty$ to -3 , and using the errors from the fit, the background estimate is 128.1 ± 9.4 events, or $0.407 \pm 0.040\%$ of data passing all cuts.

8.4.4 $e^+e^-\gamma\gamma\gamma$ Mass

The last cut is on the invariant mass of the two electron and three photon system. Figure 8.25 shows the spectra in $M_{ee\gamma\gamma\gamma}$ for data and MC. The data distribution shows some background at low masses from $K_L \rightarrow \pi^0\pi^0\pi_{Dalitz}^0$ decays where two photons are not reconstructed. An estimate of this background is made by fitting an exponential function to the data between $420 \text{ MeV}/c^2$ and $470 \text{ MeV}/c^2$. Integrating the fit function over the accepted region and using the errors on the fit gives a background of 10 ± 15 events. Despite the uncertainty, the contamination from this source is obviously negligible.

The data and MC $M_{ee\gamma\gamma\gamma}$ spectra match well near the peak. Figure 8.26 shows the results of fitting a Gaussian function to each peak. The means match well, with the data at $497.37 \pm 0.02 \text{ MeV}/c^2$ and the MC at $497.39 \pm 0.01 \text{ MeV}/c^2$, although both are 0.06% below the PDG K^0 mass of $497.672 \pm 0.031 \text{ MeV}/c^2$ [37]. The widths also match,




Figure 8.23: $\vec{P}_T^2$ (using Z_{π^0}), in GeV^2/c^2 . All non- $\vec{P}_T^2$ cuts are made.




Figure 8.24: $\log_{10} \vec{P}_T^2$ (using Z_{π^0}), where the dots are data. The smooth line is a fit to the data. The MC is shown as a histogram, normalized to make the peak heights match below the cut. $\vec{P}_T^2$ is measured in $(\text{GeV}/c)^2$. All non- $\vec{P}_T^2$ cuts are made.




Figure 8.25: $e^+e^-\gamma\gamma\gamma$ invariant mass, in GeV/c^2 . The straight line is the fit to the indicated region, extended to higher masses. All other cuts are made.

with the data at $\sigma = 3.505 \pm 0.024 \text{ MeV}/c^2$ and the MC at $\sigma = 3.489 \pm 0.011 \text{ MeV}/c^2$.

Like the $M_{ee\gamma}$ cut, the $M_{ee\gamma\gamma\gamma}$ cut requires the mass to be within $\sim 3\sigma$ of the PDG value:

$$|M_{ee\gamma\gamma\gamma} - M_{K^0}| < 10 \text{ MeV}/c^2.$$




Figure 8.26: Fits to $e^+e^-\gamma\gamma\gamma$ invariant mass, in GeV/c^2 . All other cuts are made.

The backgrounds to $K \rightarrow \pi^0\pi^0_{Dalitz}$ can be better understood by considering $M_{ee\gamma\gamma}$ versus $\vec{P}_T^2$, as in Figure 8.27. In the figure, $K \rightarrow \pi^0\pi^0\pi^0_{Dalitz}$ background is to the left and above the signal, because the missing photons carry off mass and transverse momentum. The collimator-scatter background is directly above the signal; no particles are lost so the mass is good, but the parent kaon scatters, adding transverse momentum.

8.4.5 Summary

The cuts described in this section are summarized here:

- $|M_{\gamma\gamma} - 135.5 \text{ MeV}/c^2| < 8 \text{ MeV}/c^2$
- $|M_{ee\gamma} - M_{\pi^0}| < 6 \text{ MeV}/c^2$
- $\vec{P}_T^2 < 0.001 (\text{GeV}/c)^2$ (charged vertex Z)

- $\vec{P}_T^2 < 0.001 \text{ (GeV/c)}^2 \text{ (} Z_{\pi^0} \text{)}$
- $|M_{ee\gamma\gamma\gamma} - M_{K^0}| < 10 \text{ MeV/c}^2$




Figure 8.27: $\log_{10} \vec{P}_T^2$ in $(\text{GeV/c})^2$ (y-axis) vs. $M_{ee\gamma\gamma\gamma}$ in GeV/c^2 (x-axis). The box indicates the cuts. All other cuts are made.

8.5 Flux

After making all above cuts, K_L flux for E799 is calculated. Table 8.5 lists numbers that go into the calculation, as well as fluxes for the Winter and Summer periods. The branching ratio of $K \rightarrow \pi^0 \pi_{Dalitz}^0$ is assumed to be 2.216×10^{-5} (section 8.1) and the acceptances have been multiplied by 0.945 for the missing TRD simulation (section 8.3.6). The final measurement for the number of K_L decays in E799 is 265.07 billion, with uncertainties of $\pm 0.56\%$ from statistics, $\pm 3.4\%$ from the BR, and $\pm 5.8\%$ from systematic errors.

The uncertainties on the flux are listed in Table 8.6. The branching ratio uncertainty comes from the 3.42% combined uncertainty on $\text{BR}(K_L \rightarrow \pi^0 \pi_{Dalitz}^0)$ (section 8.1). Cut variation uncertainties are obtained by changing a cut, seeing how the flux

Table 8.5: Summary of K_L flux numbers.

Quantity	Winter	Summer	All
K_L Generated	13645544.0	9862117.6	23507661.6
MC passing all cuts	77064	56275	133339
% acceptance	0.5337	0.5392	0.5360
Data passing all cuts	18371	13110	31481
$N_{K_L} \times 10^9$	155.36	109.73	265.07

changes, and adding the flux changes from all cut changes in quadrature. Table 8.7 lists the cut variations used to get this systematic error. The ρ weighting uncertainty is motivated by the slope in the ratio of the vertex Z spectra, as seen in section 8.3.3. By assuming there is an unsimulated inefficiency in the beam regions of the DCs, and by weighting MC events with tracks in the beam regions appropriately, this “Z-slope” can be reduced. Taking the sum of these weights as the number of MC accepted and calculating the change in flux gives the systematic error from the “Z-slope.” This procedure is discussed in more detail in appendix A.

Table 8.6: Uncertainties on fluxes, in units of $10^9 K_L$ s.

Uncertainty	Winter	Summer	All
Statistical	1.15	0.96	1.49
Branching Ratio	5.31	3.75	9.07
ρ weighting	+9.95	+3.07	+13.01
Cut Variations	4.64	4.13	8.31
Stat.+BR+ $M_{e e \gamma}$ cut	5.63	3.97	9.48
Stat.+BR+ $M_{e e \gamma}$ cut+ $M_{\gamma \gamma}$ cut	5.64	3.97	9.48

The systematic errors reflected in the ρ weighting and cut variation uncertainties dominate the flux uncertainty. However, the misunderstandings associated with these errors may be reproduced in the signal modes, $K_L \rightarrow e^+ e^- \gamma \gamma$ and $K_L \rightarrow \pi^0 e^+ e^-$. The flux itself is not an interesting physical quantity; $\text{BR}(K_L \rightarrow e^+ e^- \gamma \gamma)$ and $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ are. Because $K \rightarrow \pi^0 \pi^0_{\text{Dalitz}}$ is a similar decay to the signal mode(s), some of these systematic errors may cancel out when calculating cut variation and ρ weighting

systematics for $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma)$ and $\text{BR}(K_L \rightarrow \pi^0e^+e^-)$.

However, the statistical and BR uncertainties on the flux *will* have to be added as uncertainties on $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma)$ and $\text{BR}(K_L \rightarrow \pi^0e^+e^-)$. So will any cut variations that cannot be repeated on the signal mode. That is why Table 8.6 has entries for “Stat.+BR+ $M_{ee\gamma}$ cut” and “Stat.+BR+ $M_{ee\gamma}$ cut+ $M_{\gamma\gamma}$ cut.” The first is the quadratic sum of the flux uncertainties from statistics, BR and varying the $M_{ee\gamma}$ cut; this would be a systematic uncertainty on $\text{BR}(K_L \rightarrow \pi^0e^+e^-)$. The second adds (in quadrature) the uncertainty from varying the $M_{\gamma\gamma}$ cut; this would be a systematic uncertainty on $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma)$.

Table 8.7: Changes in flux when cuts are changed, in units of $10^9 K_L$ s.

Change to Cuts	Change in Flux		
	Winter	Summer	All
$40 < \vec{P} < 160 \text{ GeV}/c$	-0.35	+0.20	-0.15
drop V-bank cuts	-1.68	-1.61	-3.28
CA energy $< 1 \text{ GeV}$	-2.08	-0.74	-2.83
drop ETOT ver. cut	-0.14	-0.00	-0.15
drop cluster-hole dist. cut	+0.12	+0.12	+0.24
drop cluster energy cut	-0.61	-0.02	-0.63
drop TRD cuts and corr.	-0.88	+0.93	+0.04
$0.96 < E/p < 1.04$	+0.25	-0.00	+0.24
drop track sep. at CsI cut	-0.04	-0.02	-0.06
drop track sep. at DC1 cuts	-0.79	+0.01	-0.78
drop opening angle cut	+2.72	+3.38	+6.10
$ M_{\gamma\gamma} - 135.5 < 7 \text{ MeV}/c^2$	+0.24	-0.02	+0.22
$ M_{ee\gamma} - M_{\pi^0} < 5 \text{ MeV}/c^2$	+1.48	+0.86	+2.34
both $\vec{P}_T^2 < 0.0004 (\text{MeV}/c)^2$	+1.65	+0.86	+2.51
$ M_{ee\gamma\gamma} - M_{K^0} < 12 \text{ MeV}/c^2$	-0.11	+0.29	+0.18
quadrature sum	4.64	4.13	8.31

8.6 Simulation Quality

This section contains some additional plots that did not fit in with the discussion of cuts, but which may be of interest.

Figure 8.28 shows the reconstructed transverse distributions of the kaon beam.

These are made using the decay vertex X and Y found with tracking and Z_{π^0} . The dual-beam setup of KTeV appears in the X angle plot, and the slight slope in the distributions in Y is a result of the downward angle of the proton beam. The data-MC agreement shows good understanding of the beam profile.




Figure 8.28: Kaon beam angles, equal to vertex X or Y divided by Z_{π^0} , in radians. All cuts are made.

Figure 8.29 shows the number of data and simulated events as a function of run number. This shows that the time distribution of events is correctly handled by the accidental-overlay procedure described in section 6.3.1.

Finally, Figure 8.30 shows the invariant mass of the e^+e^- pair. This distribution is of interest because M_{ee} is an important cut for eliminating $K \rightarrow \pi^0\pi_{Dalitz}^0$ background to $K_L \rightarrow \pi^0e^+e^-$, and because M_{ee} probes the $K_L\gamma\gamma^*$ form factor in $K_L \rightarrow e^+e^-\gamma\gamma$. The distribution falls sharply at low M_{ee} , where normally Dalitz decays peak, because of the opening angle cut.




Figure 8.29: Run number distribution for Winter (Left) and Summer (Right). The MC is normalized so that its integral over Winter+Summer is equal to the integral of data over Winter+Summer. All cuts are made.




Figure 8.30: e^+e^- invariant mass, in GeV/c^2 . All cuts are made.

Chapter 9

$$K_L \rightarrow e^+e^-\gamma\gamma$$

This chapter studies the decay $K_L \rightarrow e^+e^-\gamma\gamma$, with the primary goals of measuring the BR and the form factor parameter α_{K^*} . The secondary goal will be to evaluate the simulation of $K_L \rightarrow e^+e^-\gamma\gamma$ as a means of predicting background levels in the search for $K_L \rightarrow \pi^0 e^+e^-$, the topic of the next chapter.

As discussed in section 1.3.3, there are two significant backgrounds to $K_L \rightarrow e^+e^-\gamma\gamma$: $K \rightarrow \pi^0\pi_{Dalitz}^0$ and $K_L \rightarrow e^+e^-\gamma$. Their background levels are predicted by simulation in this analysis. To check the prediction, when data spectra are compared to MC the MC spectra for signal and backgrounds are weighted by BR and summed. Unlike the last chapter, MC spectra in this chapter are normalized to the K_L flux.

That backgrounds from charged pion modes are small can and will be shown using TRD data. Likewise, that $K_L \rightarrow \pi^0\pi^0\pi_{Dalitz}^0$ background is small will be shown using MC events.

9.1 Simulated Events

Samples of two different decay modes, $K_L \rightarrow e^+e^-\gamma(\gamma)$ and $K \rightarrow \pi^0\pi_{Dalitz}^0$, are simulated to study the signal and primary backgrounds in this analysis. All $K_L \rightarrow e^+e^-\gamma(\gamma)$ are generated using the BMS form factor with $\alpha_{K^*} = -0.28$.

The first sample, $K_L \rightarrow e^+e^-\gamma(\gamma)$, contains both the signal $K_L \rightarrow e^+e^-\gamma\gamma$ and the background $K_L \rightarrow e^+e^-\gamma$. As discussed in section 6.1.2.2, the simulation can

Table 9.1: Dalitz and radiative-Dalitz kaon decays generated, with various thresholds.

Cut	Number	Fraction generated
$K_L \rightarrow e^+e^-\gamma(\gamma)$ generated	113540415	1
Dalitz ($K_L \rightarrow e^+e^-\gamma$)	81893331	0.72127 ± 0.000151
All radiations ($K_L \rightarrow e^+e^-\gamma\gamma$, $M_{\gamma\gamma} > 1 \text{ MeV}/c^2$)	31647084	0.27873 ± 0.000058
$M_{\gamma\gamma} > 20 \text{ MeV}/c^2$	13674604	0.12044 ± 0.000035
$130 < M_{\gamma\gamma} < 140 \text{ MeV}/c^2$	326438	0.00288 ± 0.000005
$K_L \rightarrow e^+e^-\gamma\gamma$, $E_\gamma^* > 5 \text{ MeV}$	6802810	0.05992 ± 0.000024
$K_L \rightarrow e^+e^-\gamma\gamma$, $E_\gamma^* > 10 \text{ MeV}$	4903726	0.04319 ± 0.000020
$K_L \rightarrow e^+e^-\gamma\gamma$, $E_\gamma^* > 15 \text{ MeV}$	3868350	0.03407 ± 0.000018

generate a $K_L \rightarrow e^+e^-\gamma$ decay with either one or two photons coming from the decay. The probability of one photon (Dalitz decay) is 0.72. The probability of simulating two photons (radiative decay) is 0.28, with the threshold $M_{\gamma\gamma} > 1 \text{ MeV}/c^2$. These two possibilities are the second and third entries in Table 9.1. Table 9.1 also lists the number of events generated above other photon energy thresholds. Depending on which threshold is being used, the number of generated $K_L \rightarrow e^+e^-\gamma\gamma$ signal changes as listed in the table. The number of background generated changes too. The $K_L \rightarrow e^+e^-\gamma$ background sample includes both the decays generated as Dalitz and the radiative decays with photon energy generated below the threshold. These “below-threshold $K_L \rightarrow e^+e^-\gamma\gamma$ ” are regarded as $K_L \rightarrow e^+e^-\gamma$ background in this analysis.

For the branching ratio of $K_L \rightarrow e^+e^-\gamma(\gamma)$, I use the prediction $\text{BR}(K_L \rightarrow e^+e^-\gamma(\gamma))/\text{BR}(K_L \rightarrow \gamma\gamma) = 0.01679$ [35]. This prediction is made assuming the BMS form factor of equation 1.3, with $\alpha_{K^*} = -0.28$. Using the measurement $\text{BR}(K_L \rightarrow \gamma\gamma) = (5.92 \pm 0.15) \times 10^{-4}$ [37], I get $\text{BR}(K_L \rightarrow e^+e^-\gamma(\gamma)) = (9.94 \pm 0.25) \times 10^{-6}$. The branching ratio of $K_L \rightarrow e^+e^-\gamma$ is this 9.94×10^{-6} , minus the fraction of $K_L \rightarrow e^+e^-\gamma(\gamma)$ predicted to be $K_L \rightarrow e^+e^-\gamma\gamma$.

The events listed in Table 9.1 are in all E799 spills and runs. Because the probability of radiative decay is not a function of run number, I get the number of generated MC events by multiplying the fractions in Table 9.1 by the number of events generated

in good spills; see Table 9.2.

The other major sample of simulated events is $K \rightarrow \pi^0 \pi^0_{Dalitz}$, summarized in Table 9.3. This is the same kind of MC sample used in chapter 8, but enlarged to provide good statistics for a background estimate. Also, this $K \rightarrow \pi^0 \pi^0_{Dalitz}$ sample differs from the normalization's MC sample in that these events must have the EEGG crunch filter tag, not EEGGG. The Winter events were generated with an older version of ktevmc; the chief difference was the absence of some DC simulation modifications mentioned in section 6.2.3.1.

9.2 General Cuts

There are a few ways in which this analysis's basic cuts differ from the normalization's. First, the crunch filter tag is EEGG instead of EEGGG. In short, this means four clusters instead of five and $\vec{P}_T^2 < 0.001 (\text{GeV}/c^2)^2$ instead of $M_{ee\gamma\gamma} > 380 \text{ MeV}/c^2$. Also, exactly four clusters must be present in post-filter reconstruction. And of course, there is no π^0 , so there is no vertex Z_{π^0} cut.

9.3 Mode-Specific Cuts

This section is analogous to section 8.4, describing cuts that differ significantly from those presented for the normalization.

9.3.1 Photon Energy Threshold

A cut is made that depends on the infrared threshold for the BR being measured. The resolution for reconstructing the variable in MC is found, and the cut is set several sigma away from the threshold to keep sub-threshold background from contaminating the sample.

Table 9.2: Numbers of $K_L \rightarrow e^+e^-\gamma(\gamma)$. % is the number on that line divided by the number generated. L1, L2, and L3 trigger and EEGG filter sums are in all spills. L3 triggers are 2e, n-clus tags only.

Stage	Winter		Summer	
	Number	%	Number	%
Generate	64590454	100.00	48949961	100.00
Good Spills	59790425	92.57	43212791	88.28
L1 trigger	15217227	23.56	11742975	23.99
L2 trigger	967098	1.50	628650	1.28
L2 trigger	674842	1.04	493097	1.01
EEGG filter	373270	0.58	277490	0.57

Table 9.3: Numbers of $K \rightarrow \pi^0\pi_{Dalitz}^0$ simulated as background to $K_L \rightarrow e^+e^-\gamma\gamma$. % is the number on that line divided by the number Generated. K_L weight is the sum of $|\eta_{00}\mathcal{A}_L(t)|^2/|\mathcal{A}_S(t) + \eta_{00}\mathcal{A}_L(t)|^2$ in good spills. L1, L2, and L3 trigger and EEGG filter sums are unweighted and in all spills. L3 triggers are 2e, n-clus tags only.

Stage	Winter		Summer	
	Number	%	Number	%
Generated	156672281	100.00	121210253	100.00
Good Spills	144840605	92.45	107003845	88.28
K_L weight	133457096	85.18	98578248	81.33
L1 trigger	23269659	14.85	18256452	15.06
L2 trigger	14586034	9.31	12085304	9.97
L3 trigger	10104083	6.45	8199804	6.76
EEGG Filter	1069312	0.68	810186	0.67

9.3.1.1 E_γ^*

E_γ^* is the minimum photon energy in the center-of-mass reference frame. This is reconstructed by boosting the photon lab-frame four-momentum by a velocity $\vec{\beta} = -\vec{P}/E$, where $\vec{P}$ and E momentum and energy summed for all four decay products. Previous measurements of $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma)$ have all used $E_\gamma^* > 5 \text{ MeV}$ [32] [33], and plots in this chapter use this threshold unless otherwise noted.

Figure 9.1 shows the data and MC distributions for small E_γ^* . The simulated signal, $K_L \rightarrow e^+e^-\gamma$, and $K \rightarrow \pi^0\pi_{Dalitz}^0$ distributions are all shown, normalized using the K_L flux and their branching ratios.

The good match between data and MC in Figure 9.1 supports the use of MC to study E_γ^* . The next step is to determine the resolution in reconstructing E_γ^* . Figure 9.2 shows the difference between reconstructed E_γ^* and generated E_γ^* for $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC. The events on the left side of the left plot are from events where the very-soft, internally-radiated photon is lost but is replaced with an accidental or externally-radiated photon. The fit to a Gaussian in the right plot shows that E_γ^* is reconstructed with a resolution of about 0.43 MeV, but also biased upwards by 0.27 MeV. To be on the safe side, cuts on E_γ^* are 3 MeV higher than the thresholds.

Figure 9.3 shows efficiency for low E_γ^* . This plot uses generated E_γ^* only, and is the ratio of accepted to generated $K_L \rightarrow e^+e^-\gamma(\gamma)$. It shows that sensitivity to E_γ^* drops off at small values, but that there is still significant sensitivity at the threshold $E_\gamma^* > 5 \text{ MeV}$.

9.3.1.2 $M_{\gamma\gamma}$

The same arguments can be carried out using $M_{\gamma\gamma}$ instead of E_γ^* . Plots 9.4, 9.5, and 9.6 reproduce the plots in the last section for $M_{\gamma\gamma}$. Note that the signal region for the $K_L \rightarrow \pi^0 e^+ e^-$ search is excluded by the θ_{min} cut (section 9.3.4), so no $K_L \rightarrow \pi^0 e^+ e^-$




Figure 9.1: E_γ^* in GeV after all other cuts. The data/MC ratio is made using the sum of MC distributions.




Figure 9.2: E_{γ}^* as generated in simulation and as reconstructed, for $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC, in GeV. (Right) Difference between E_{γ}^* s when both are less than 25 MeV. All other cuts are made.




Figure 9.3: Efficiency vs. E_{γ}^* , in GeV. All other cuts are made.

candidate events can appear in the data spectrum in Figure 9.4. Figure 9.5 shows the resolution to be $0.69 \text{ MeV}/c^2$, with a mean bias of $+0.29 \text{ MeV}/c^2$; the cut is set at $M_{\gamma\gamma} > 24 \text{ MeV}/c^2$ for the threshold $M_{\gamma\gamma} > 20 \text{ MeV}/c^2$.

9.3.2 Transverse Momentum

Figure 9.7 compares $\vec{P}_T^2$ distributions for data and MC. Because there is no π^0 in the $K_L \rightarrow e^+e^-\gamma(\gamma)$ decays, only $\vec{P}_T^2$ calculated using the charged-track vertex can be used. Only $\vec{P}_T^2 < 0.001 (\text{GeV}/c)^2$ is shown because that is the cut requirement made by the EEGG filter tag. The analysis cut is set at $\vec{P}_T^2 < 0.0003 (\text{GeV}/c)^2$, where the ratio of data to signal MC begins to be not flat. A couple of features to note in Figure 9.7: first, the backgrounds, which come in either through losing or gaining a particle, tend to have higher $\vec{P}_T^2$. Second, the data has slightly better resolution than the MC, just as in the normalization study.

9.3.3 $M_{ee\gamma\gamma}$

Figure 9.8 shows the $e^+e^-\gamma\gamma$ invariant mass distributions for data and MC. A number of features are apparent. At $M_{ee\gamma\gamma} < 425 \text{ MeV}/c^2$ there are a lot of extra data events; these are mostly from $K \rightarrow \pi^0\pi^0\pi_{Dalitz}^0$ decays with three lost photons. Between $425 \text{ MeV}/c^2$ and the kaon mass the main background is $K \rightarrow \pi^0\pi_{Dalitz}^0$, which has a low mass because one photon is missing. Above the kaon mass, $K_L \rightarrow e^+e^-\gamma$ takes over as the primary background. This mode has a high mass because after the θ_{min} cut removes hard external bremsstrahlung photons, it only gets accepted when there is an accidental photon cluster. The accidental tends to give higher energy than the lost soft internal radiation, so the overall mass is high. The ratio of data to summed MC shows that the simulated $K_L \rightarrow e^+e^-\gamma$ does not account for all of the data above the kaon mass. This could be due to an insufficient simulation of accidentals for $K_L \rightarrow e^+e^-\gamma$, or some other unsimulated mode with accidentals. However, the excess is small compared




Figure 9.4: $\gamma\gamma$ invariant mass in GeV/c^2 after all other cuts. The data/MC ratio is made using the sum of MC distributions.




Figure 9.5: $\gamma\gamma$ invariant mass as generated in simulation and as reconstructed, for $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC, in GeV/c^2 . (Right) Difference between $M_{\gamma\gamma}$ s when both are less than $60 \text{ MeV}/c^2$. All other cuts are made.




Figure 9.6: Efficiency vs. $M_{\gamma\gamma}$, in GeV/c^2 . All other cuts are made.




Figure 9.7: $\log_{10} \vec{P}_T^2$ in $(\text{GeV}/c)^2$, after all other cuts. The data/MC ratio is made using the $K_L \rightarrow e^+e^-\gamma\gamma$ MC distribution.

to the signal peak, so the effect is minimal.

Figure 9.9 shows the data and simulated signal $M_{ee\gamma\gamma}$ distributions fitted to Gaussians. The means agree fairly well, with the data at $498.46 \pm 0.17 \text{ MeV}/c^2$ and the MC at $498.24 \pm 0.02 \text{ MeV}/c^2$. The data width is somewhat wider at $\sigma = 4.68 \pm 0.18 \text{ MeV}/c^2$ than the MC at $\sigma = 4.379 \pm 0.023 \text{ MeV}/c^2$. This is consistent with the presence of an evenly distributed background in the data. The cut requires $M_{ee\gamma\gamma}$ within $11 \text{ MeV}/c^2$ of the kaon mass.

9.3.4 θ_{min}

In the $e^+e^-\gamma\gamma$ system, the minimum angle between any electron and any photon in the center-of-mass reference frame, θ_{min} , provides a useful tool for separating $K_L \rightarrow e^+e^-\gamma\gamma$ from its backgrounds. Figure 9.10, which shows $\log_{10}\theta_{min}$ for data and MC samples, demonstrates this. When a high-energy electron radiates a photon through bremsstrahlung while passing through matter, the photon's direction is very closely correlated with the electrons. Thus, $K_L \rightarrow e^+e^-\gamma$ decays with external radiation have much smaller θ_{min} than $K_L \rightarrow e^+e^-\gamma\gamma$. This can be seen in the $K_L \rightarrow e^+e^-\gamma$ triangles in Figure 9.10. On the other hand, the photon directions are relatively uncorrelated with electron directions in $K \rightarrow \pi^0\pi^0_{Dalitz}$ decays. In a $\log_{10}\theta_{min}$ plot, the $K \rightarrow \pi^0\pi^0_{Dalitz}$ distribution is concentrated at high values, as seen in the cross-hatched area of the figure. $K_L \rightarrow e^+e^-\gamma\gamma$ sits between these extremes. The cuts to select it are set where the backgrounds start to rise, at $-1.3 < \log_{10}\theta_{min} < -0.2$.

There is a noticeable dip in the data/MC ratio for $\log_{10}\theta_{min}$, in the region dominated by $K_L \rightarrow e^+e^-\gamma$. In the region below the cut, there are 2,397 data events, 1,222 $K_L \rightarrow e^+e^-\gamma\gamma$ MC, and 1,389 $K_L \rightarrow e^+e^-\gamma$ MC — an excess of $8.9 \pm 2.1\%$ with statistical errors. Some of this excess is from an overall excess in the MC scale. The MC scale comes from the flux and the BRs. The flux uncertainty is at least 3.6% and the uncertainty on $\text{BR}(K_L \rightarrow \gamma\gamma)$ is 2.5%, so some of the excess could be from normalizing the




Figure 9.8: $e^+e^-\gamma\gamma$ invariant mass in GeV/c^2 after all other cuts. The data/MC ratio is made using the summed MC distributions.




Figure 9.9: $e^+e^-\gamma\gamma$ invariant mass in GeV/c^2 after all other cuts, for data (Left) and simulated $K_L \rightarrow e^+e^-\gamma\gamma$ (Right).




Figure 9.10: $\log_{10}\theta_{\text{min}}$, where θ_{min} is measured in radians, after all other cuts. The data/MC ratio is made using the summed MC distributions.

MC. However, the nature of an excess in MC at small θ_{min} still needs to be addressed.

One explanation for an excess in MC at small θ_{min} would be that the simulated acceptance for external radiation is too high, such as if the material in the detector was underestimated. The discrepancy would have to be between 8.2% and 15.4%, depending on what fraction of $K_L \rightarrow e^+e^-\gamma\gamma$ are accepted through external radiation. A deficiency in the simulated material in the detector at this level would have been noticed in other analyses of E799 or E832, and it has not.

Another possibility is that the resolution of θ_{min} is too good in MC. Worsening the resolution would shift $\log_{10}\theta_{min}$ up for the MC at low angles, flattening the data/MC ratio; the overall ratio would still be low by a few percent, but that could be because of branching ratios or flux. Poorly simulated resolution is not unlikely, as can be seen by considering the variable “MINTPD.” MINTPD is the minimum distance between any photon cluster and any track (upstream of the magnet) projected to the calorimeter. MINTPD is closely correlated with θ_{min} , but shows detector effects more clearly. (However, MINTPD is not used as a cut because it offers worse rejection of $K \rightarrow \pi^0\pi^0_{Dalitz}$ than θ_{min} .) Figure 9.11 shows that the photon-electron distance is between 0.1 mm and 10 mm for $K_L \rightarrow e^+e^-\gamma$, peaking around 2.8 mm. Understanding cluster-position resolution at this level is a formidable task in an experiment like KTeV, let alone with track-direction resolution added in.

Instead, I note that increasing all θ_{min} for MC events by 15% flattens out the data/MC ratio for θ_{min} . I use this effect to add a systematic error to $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma)$ in section 9.5. I do not believe that scaling actually models the effect that is causing the excess MC at small θ_{min} . The scaling is merely an ad-hoc technique to estimate a systematic uncertainty from the effect. The scaling is applied only to the $K_L \rightarrow e^+e^-\gamma$ background to get a more conservative (larger) uncertainty.

9.3.5 Summary

The cuts described in this section are summarized here:

- $E_\gamma^* > 8 \text{ MeV}$ (for threshold $E_\gamma^* > 5 \text{ MeV}$)
- $\vec{P}_T^2 < 0.0003 (\text{GeV}/c)^2$
- $|M_{ee\gamma\gamma} - M_{K^0}| < 11 \text{ MeV}/c^2$
- $-1.3 < \log_{10} \theta_{min} < -0.2$ (or $0.05012 < \theta_{min} < 0.6310 \text{ rad}$)




Figure 9.11: $\log_{10}\text{MINTPD}$, where MINTPD is measured in meters, after all cuts except θ_{min} are made. The data/MC ratio is made using the summed MC distributions.

9.4 Other Backgrounds

This section demonstrates the negligible low levels of some other background modes.

9.4.1 Charged Pions

The number of events failing the Π_{TRD} cut (section 8.3.6) can be used to estimate background from charged-pion modes, by assuming that pions have a flat distribution in Π_{TRD} . After all other cuts but the Π_{TRD} cut are made, there are 1,686 data events remaining. These can be divided into three groups:

- 1,578 e^+e^- candidates, with both tracks having $\Pi_{\text{TRD}} < 0.04$.
- 106 K_{e3} candidates, with only one track having $\Pi_{\text{TRD}} > 0.04$.
- 2 $\pi^+\pi^-$ candidates, with both tracks having $\Pi_{\text{TRD}} > 0.04$.

Obviously the $\pi^+\pi^-$ population with two $\Pi_{\text{TRD}} < 0.04$ is tiny. K_{e3} s are another story. In the normalization sample, there are $5.31 \pm 0.13\%$ as many K_{e3} candidates as e^+e^- candidates. If the normalization sample is purely electrons, then this is the fraction of real e^+e^- events that look like K_{e3} . If one assumes that the signal e^+e^- candidates contain few actual pions, then one would expect only $1578 \times 5.38\% = 83.8 \pm 2.0$ K_{e3} candidates. If the excess $106 - 83.8 = 22 \pm 10$ K_{e3} candidates are actual K_{e3} events, then one would expect $22 \times 0.04/(1 - 0.04) = 0.92 \pm 0.42$ K_{e3} to be e^+e^- candidates. In real life, Π_{TRD} is biased towards 1 for pions (see Figure 5.6), so the actual expected background is somewhat lower than 0.92 events.

9.4.2 $K \rightarrow \pi^0\pi^0\pi_{\text{Dalitz}}^0$

An estimate of the background from $K \rightarrow \pi^0\pi^0\pi_{\text{Dalitz}}^0$ can be made using MC events. For convenience, I use a pre-existing sample where two million events were generated for each of the 314 runs in the E799 data set. This is the same sample as used for the run-to-run flux measurement mentioned in section 6.3.1. This sample only has a coarse simulation of time-dependent features of the detector, but it is unlikely that these features could have a major effect on this background estimate. No events

in this sample pass all of the cuts described in this note, so a fit on $M_{ee\gamma\gamma}$ is used to estimate the background. Fitting a decaying exponential function to Figure 9.12 and integrating over the accepted region in $M_{ee\gamma\gamma}$ gives 0.4 ± 2.9 events. Multiplying by the flux $((2.65 \pm 0.18) \times 10^{11})$ and the branching ratio $((7.41 \pm 0.22) \times 10^{-3}$ [37]) and dividing by the number generated (628×10^6) , gives an expected background of 1.3 ± 8.8 events. Even though the errors on the fit leave a lot of uncertainty, by eyeballing Figure 9.12 one can see that the background from $K \rightarrow \pi^0 \pi^0 \pi_{Dalitz}^0$ is negligible.




Figure 9.12: $M_{ee\gamma\gamma}$ for simulated $K \rightarrow \pi^0 \pi^0 \pi_{Dalitz}^0$, in GeV/c^2 . The MC sample size corresponds to $\sim 31.6\%$ of the flux. All other cuts are applied.

9.5 Branching Ratios

This section calculates $\text{BR}(K_L \rightarrow e^+ e^- \gamma \gamma)$ and uncertainties for various thresholds. The formula for BR is as follows:

$$B(K_L \rightarrow e^+ e^- \gamma \gamma) = \frac{1}{N_{K_L}} \frac{G_{\text{sig}}}{A_{\text{sig}} \epsilon_{trd}} \left(n_o - \sum_{\text{backgrounds}} N_{K_L} B(K_L \rightarrow \text{bg}) \frac{A_{\text{bg}} \epsilon_{trd}}{G_{\text{bg}}} \right).$$

G is the number of MC events generated, A is the number of MC events accepted, $\epsilon_{trd} = 0.945$ is the TRD acceptance correction factor (section 8.3.6), and n_o is the

Table 9.4: Branching ratio numbers for $K_L \rightarrow e^+e^-\gamma\gamma$, $E_\gamma^* > 5$ MeV. Generated number is in good spills only. Expected backgrounds include TRD acceptance; background uncertainties are from statistics and BRs.

Quantity	Value
Flux, N_{K_L}	265.07×10^9
Generated $K_L \rightarrow e^+e^-\gamma\gamma$	6.171×10^6
Accepted $K_L \rightarrow e^+e^-\gamma\gamma$	69219
acceptance $A\epsilon_{trd}/G$	1.060%
observed events n_o	1578
Generated $K_L \rightarrow e^+e^-\gamma$	9.683×10^7
Accepted $K_L \rightarrow e^+e^-\gamma$	1437
expected $K_L \rightarrow e^+e^-\gamma$ background	34.74 ± 1.27
Generated $K \rightarrow \pi^0\pi^0_{Dalitz}$	23.20×10^8
Accepted $K \rightarrow \pi^0\pi^0_{Dalitz}$	1117
expected $K \rightarrow \pi^0\pi^0_{Dalitz}$ background	26.72 ± 1.21
n_o minus expected backgrounds	1516.54 ± 1.75
Branching Ratio	5.40×10^{-7}

number of events accepted in the data.

9.5.1 $E_\gamma^* > 5$ MeV

Table 9.4 gives the numbers that go into the branching ratio calculation. Uncertainties on this flux are listed in Table 9.5. Some notes on the errors:

- The statistical error is based on the number of events passing all cuts, 1,578.
- The flux systematic error includes only the change in flux from varying cuts on $M_{ee\gamma}$ and $M_{\gamma\gamma}$.
- The flux branching ratio uncertainty comes from PDG uncertainties on $\text{BR}(K \rightarrow \pi^0\pi^0)$ and $\text{BR}(\pi^0 \rightarrow e^+e^-\gamma)$.
- A systematic from poorly simulated θ_{min} resolution (see section 9.3.4) is calculated by multiplying θ_{min} by 1.15, *only* in the $K_L \rightarrow e^+e^-\gamma$ MC sample. The change in BR resulting from the change in expected background is the uncertainty.

- The ρ weighting systematic (see appendix A) is calculated by applying the ρ weights to MC events and finding the change in the BR, including the change in the flux. Apparently, the DC beam-region inefficiency largely cancels out in the signal: normalization ratio.
- The form-factor systematic uncertainty comes from uncertainty in my calculated value of α_{K^*} . This error is discussed in section 9.6.4.
- Cut variations are done in the same way as for the normalization, where possible. When a cut is changed, the change in flux is included when calculating change in BR. The most costly cut variation is on $\vec{P}_T^2$; this is likely a result of making a tighter cut on the signal than on the normalization.

To summarize, for $K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV}$ the branching ratio is measured to be $(5.40 \pm 0.14_{stat.} \pm 0.29_{sys.} \pm 0.18_{BR}) \times 10^{-7}$. This is only 1.6σ lower than Greenlee's tree-level QED prediction of $(6.02 \pm 0.15) \times 10^{-7}$ (section 1.3.2 and [35]). Using the MC with second-order QED corrections and assuming $\alpha_{K^*} = -0.28$ gives a predicted BR of 6.05×10^{-7} . The combined uncertainty is 5.9%, mostly from measurement systematics.

9.5.2 Other Thresholds

Table 9.6 lists branching ratios for $K_L \rightarrow e^+e^-\gamma\gamma$ with the thresholds $E_\gamma^* > 10 \text{ MeV}$, $E_\gamma^* > 15 \text{ MeV}$, and $M_{\gamma\gamma} > 20 \text{ MeV}/c^2$. The expected background from $K \rightarrow \pi^0\pi_{Dalitz}^0$ is the same for all of these thresholds because none of the $K \rightarrow \pi^0\pi_{Dalitz}^0$ MC events are near any of the infrared thresholds. The same calculations go into the uncertainties as for $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV})$. Again, the BRs do not disagree with predictions.

Table 9.5: Uncertainties on $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV})$. All sums of uncertainties are quadrature sums.

Error	$\sigma_{\text{BR}} \times 10^{-7}$
Statistics	0.136
Flux Branching Ratio	0.183
(internal systematics:)	
Flux statistics	0.030
Flux systematics	0.047
Monte Carlo statistics	0.023
$K_L \rightarrow e^+e^-\gamma$ background	0.005
$K \rightarrow \pi^0\pi^0_{\text{Dalitz}}$ background	0.004
$1.15 \times \theta_{\min}$ in $K_L \rightarrow e^+e^-\gamma$ MC	-0.024
ρ weighting	-0.027
form factor	-0.179
sum of cut variations	0.211
sum of above int. systematics	0.286
Cut Variations:	
$40 < \vec{P} < 160 \text{ GeV}/c$	+0.022
drop V-bank cuts	+0.075
CA energy $< 1 \text{ GeV}$	+0.017
drop ETOT ver. cut	+0.007
drop cluster-hole dist. cut	-0.010
drop cluster energy cut	-0.032
drop TRD cuts and corr.	+0.070
$0.96 < E/p < 1.04$	-0.010
drop track sep. at Csi cut	+0.007
drop track sep. at DC1 cuts	-0.073
drop opening angle cut	-0.086
$\vec{P}_T^2 < 0.00011 (\text{MeV}/c)^2$	-0.114
$ M_{ee\gamma\gamma} - M_{K^0} < 13.2 \text{ MeV}/c^2$	+0.073
drop infrared threshold cut	-0.003
$-1.4 < \log_{10} \theta_{\min} < -0.1$	-0.029

Table 9.6: Branching ratio numbers for $K_L \rightarrow e^+e^-\gamma\gamma$ with different thresholds. Expected backgrounds include TRD acceptance; background uncertainties are from statistics and BRs.

Quantity	Values		
	$E_\gamma^* > 10 \text{ MeV}$	$E_\gamma^* > 15 \text{ MeV}$	$M_{\gamma\gamma} > 20 \text{ MeV}$
Flux, N_{K_L}	265.07×10^9		
Generated $K_L \rightarrow e^+e^-\gamma\gamma$	4.449×10^6	3.509×10^6	12.406×10^6
Accepted $K_L \rightarrow e^+e^-\gamma\gamma$	65991	60251	69417
acceptance $A\epsilon_{trd}/G$	1.40%	1.62%	0.53%
observed events n_o	1502	1379	1580
Generated $K_L \rightarrow e^+e^-\gamma$	9.855×10^7	9.949×10^7	0.9060×10^8
Accepted $K_L \rightarrow e^+e^-\gamma$	1229	1018	1328
expected $K_L \rightarrow e^+e^-\gamma$ bg	29.39 ± 1.12	24.61 ± 0.99	32.10 ± 1.19
Generated $K \rightarrow \pi^0\pi_{Dalitz}^0$	2.320×10^8		
Accepted $K \rightarrow \pi^0\pi_{Dalitz}^0$	1117	1117	1117
expected $K \rightarrow \pi^0\pi_{Dalitz}^0$ bg	26.72 ± 1.21	26.72 ± 1.21	26.72 ± 1.21
n_o minus expected backgrounds	1445.90 ± 1.65	1327.67 ± 1.56	1521.18 ± 1.69
Branching Ratio	3.89×10^{-7}	3.09×10^{-7}	10.85×10^{-7}
BR/BR($K_L \rightarrow e^+e^-\gamma(\gamma)$)	0.0391	0.00310	0.0109
statistical uncertainty	$\pm 0.10 \times 10^{-7}$	$\pm 0.08 \times 10^{-7}$	$\pm 0.27 \times 10^{-7}$
(internal) systematic uncertainty	$\pm 0.24 \times 10^{-7}$	$\pm 0.19 \times 10^{-7}$	$\pm 0.60 \times 10^{-7}$
BR uncertainty (from flux)	$\pm 0.13 \times 10^{-7}$	$\pm 0.10 \times 10^{-7}$	$\pm 0.37 \times 10^{-7}$

9.6 Form Factor

The spectrum of e^+e^- invariant mass informs us about the $K_L\gamma\gamma$ vertex, as mentioned in section 1.3.2. In this section I attempt to measure the BMS model form factor parameter α_{K^*} using the M_{ee} spectrum in $K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV}$ decays. The larger α_{K^*} is, the greater the enhancement of the partial width, especially as $M_{ee} \rightarrow M_{K^0}$; see Figure 9.13.

Figure 9.14 compares the data to MC for M_{ee} . The $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC events are generated assuming $\alpha_{K^*} = -0.28$. A deficit in the simulated events can be seen around $M_{ee} = 70 \text{ MeV}/c^2$, while an excess of simulated events appears above $200 \text{ MeV}/c^2$. Thus, at first glance, it would seem that these data do not support the hypothesis $\alpha_{K^*} = -0.28$.

This measurement of α_{K^*} begins by finding the α_{K^*} that causes the MC M_{ee} distribution to best match the data's M_{ee} distribution. Neither the flux nor the BR is part of the calculation, only the distribution shapes. Next, statistical and systematic errors for this α_{K^*} are quantified. Then, because α_{K^*} affects the acceptance calculation for $K_L \rightarrow e^+e^-\gamma\gamma$, the BR of $K_L \rightarrow e^+e^-\gamma\gamma$ is recalculated using the α_{K^*} observed in these data.

9.6.1 Measuring Form Factor

The procedure used to find a best value for α_{K^*} is to first plot the data as a function of M_{ee} in 10 bins between 0 and $0.5 \text{ GeV}/c^2$. Then MC distributions in reconstructed M_{ee} are found for different values of α_{K^*} . Each $K_L \rightarrow e^+e^-\gamma$ and $K_L \rightarrow e^+e^-\gamma\gamma$ MC event is given a weight $(f(M'_{ee}, \alpha_{K^*})/f(M'_{ee}, -0.28))^2$, where M'_{ee} is M_{ee} as generated in the simulation. All three MC event types ($K_L \rightarrow e^+e^-\gamma\gamma$, $K_L \rightarrow e^+e^-\gamma$, and $K \rightarrow \pi^0\pi^0_{Dalitz}$) are added together, with the $e^+e^-\gamma\gamma/e^+e^-\gamma$ fraction given by the prediction and the fraction of $K \rightarrow \pi^0\pi^0_{Dalitz}$ given by its relative BR and number




Figure 9.13: Square of BMS form factor, vs α_{K^*} and M_{ee} in GeV/c^2 . Note that equation 1.3 gives $f(x)$, while $f(M_{ee}, \alpha_{K^*})^2$ is plotted here.




Figure 9.14: e^+e^- invariant mass, in GeV/c^2 . Note that the summed MC is normalized to have the same area as the data, while the other MC distributions are normalized by flux and BR. The data/MC ratio plot uses summed MC. All cuts are applied.

generated. Only one-half of the $K_L \rightarrow e^+e^-\gamma(\gamma)$ sample is used at this stage, and the other half is saved for a check described in section 9.6.2. The three MC distributions are summed together. The summed MC distribution is normalized to have the same area as the data. Figure 9.15 shows the data and MC for some representative α_{K^*} .

Next, a χ^2 is computed for the data and the MC with each α_{K^*} . Only the first nine bins are used, because there are no data in the tenth bin. These χ^2 are plotted as a function of α_{K^*} in Figure 9.16. The error bar on each χ^2 is calculated assuming that the only uncertainty comes from data statistics in each bin. If

$$\chi^2 = \sum_{\text{bins},i} \frac{(n_i - m_i)^2}{\sigma^2(n_i) + \sigma^2(m_i)} = \sum_{\text{bins},i} \frac{(n_i - m_i)^2}{n_i},$$

where n_i is the number of data in bin i , m_i is the sum of MC weights in bin i , and $\sigma(n_i) = \sqrt{n_i}$, then

$$\sigma^2(\chi^2) = \sum_{\text{bins},i} \frac{(n_i^2 - m_i^2)^2}{n_i^3},$$

Here, $\sigma(\chi^2)$ represents how much spread would be expected to appear in a distribution of χ^2 if E799 were to be repeated very many times. This $\sigma(\chi^2)$ is only used while fitting a parabola to $\chi^2(\alpha_{K^*})$. The minimum of the fit is found at $\alpha_{K^*} = +0.0145$.

9.6.2 Statistical Error

The statistical uncertainty on this measurement is found by using the other half of the $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC sample. This is divided into 22 subsets that, after all cuts are applied, each have about as many events as the data. The above procedure for finding α_{K^*} is then repeated with each of these subsets. The mean α_{K^*} for these 22 measurements is -0.253 , with a RMS of 0.092 . This check confirms that the measurement technique can return the generated α_{K^*} . Repeating this procedure, but with the subsets reweighted to have $\alpha_{K^*} = +0.0145$ gives a mean α_{K^*} from the subsets of $+0.0529$ with an RMS of 0.118 . This RMS is taken as the statistical uncertainty on α_{K^*} .




Figure 9.15: e^+e^- invariant mass for data and MC with several α_{K^*} , in GeV/c^2 . The error bars on the MC levels are smaller than the symbols at the levels.




Figure 9.16: χ^2 for data fit to summed MC with varied α_{K^*} . The arrow indicates the minimum of the fit at $\alpha_{K^*} = +0.0145$ and $\chi^2 = 4.30$.

9.6.3 Systematic Errors

Systematic uncertainties are considered from two sources: DC inefficiency and misunderstanding of acceptance as a function of M_{ee} . These sources are considered to be the most likely causes of a slope in the M_{ee} data/MC ratio, other than form-factor physics.

DC beam-region inefficiency, which might cause a change in the M_{ee} distributions, is tested by applying ρ weighting. After applying ρ weighting to the $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC, the measured α_{K^*} shifts by only +0.0013. Therefore beam-region inefficiency has almost no effect on the shape of M_{ee} , and this error is not used for a systematic uncertainty.

The possibility that the M_{ee} shapes are different between data and MC because of some unsimulated acceptance effect requires attention. Unfortunately there is no decay mode that probes large M_{ee} that does not share the $K_L\gamma\gamma$ form factor with $K_L \rightarrow e^+e^-\gamma\gamma$. KTeV studies of $K_L \rightarrow \pi^+\pi^-(\pi^0)$ have not seen any such acceptance problem in $M_{\pi^+\pi^-}$, but directly comparing pion modes to electron modes is difficult. Dalitz-pion decays are more comparable with $K_L \rightarrow e^+e^-\gamma\gamma$, although they are limited to $M_{ee} < M_{\pi^0}$. However, an acceptance problem is more likely to appear in this region than for high M_{ee} , because of the difficulty of simulating/reconstructing close tracks. A slope appears in the data/MC ratio for M_{ee} in $K \rightarrow \pi^0\pi_{Dalitz}^0$ (see Figure 8.30), but it is not statistically significant. For abundant statistics, I turn to $K \rightarrow \pi^0\pi^0\pi_{Dalitz}^0$ decays. Figure 9.17 shows the ratio of data to MC after a set of cuts designed to select $K \rightarrow \pi^0\pi^0\pi_{Dalitz}^0$ in E799 [51]. No significant slope appears in a line fit to the ratio distribution. Taking the uncertainty on the fit to be the limit of our understanding of acceptance vs. M_{ee} , I weight the MC sums by the fit in Figure 9.17 and measure α_{K^*} again. With a weighting function of $1/(0.994 + (0.038 + 0.096)M_{ee})$, the measured α_{K^*} changes by -0.032 , while a weight $1/(0.994 + (0.038 - 0.096)M_{ee})$ makes the measured

α_{K^*} increase by 0.014. To be conservative, the larger of these changes (-0.032) is taken as a symmetric systematic uncertainty (± 0.032).




Figure 9.17: Ratio of data to MC distributions of e^+e^- invariant mass for $K \rightarrow \pi^0\pi^0\pi^0_{\text{Dalitz}}$, in GeV/c^2 . Note only the range 0-0.15 GeV/c^2 is shown.

9.6.4 Effect on Branching Ratio

The form factor affects the BR measurement by changing the acceptance; i.e., a more negative α_{K^*} puts more events at high M_{ee} where the detector has higher acceptance. (The form factor also effects the actual BR by enhancing or suppressing the vertex amplitude.) To see the acceptance effect, the $K_L \rightarrow e^+e^-\gamma(\gamma)$ events are reweighted by $(f(M'_{ee}, +0.0145)/f(M'_{ee}, -0.28))^2$, as described in section 9.6.1. This reduces the acceptance from 1.06% to 0.984% and the expected $K_L \rightarrow e^+e^-\gamma$ background from 34.75 ± 1.27 to 32.34 ± 1.20 events. The BR for $K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV}$ is then $(5.82 \pm 0.15_{\text{stat.}} \pm 0.31_{\text{sys.}} \pm 0.19_{\text{BR}}) \times 10^{-7}$ with 1518.9 ± 1.7 observed events after background subtraction. The uncertainties are scaled from those in section 9.5.1. This BR measurement is somewhat more self-consistent within the context of this experiment

than the BR presented in section 9.5.1. However, the section 9.5.1 measurement is no less valid, because the difference between them is only the choice of parameter α_{K^*} , and α_{K^*} is in some sense orthogonal to BR. Also, my measurement of α_{K^*} has large uncertainties and differs from the PDG's fit for α_{K^*} by only about two sigma, so the choice of α_{K^*} is somewhat ambiguous.

Greenlee's prediction of $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV}) = (6.02 \pm 0.15) \times 10^{-7}$ assumes $\alpha_{K^*} = -0.28$ to get $\text{BR}(K_L \rightarrow e^+e^-\gamma)$. A larger α_{K^*} would mean a smaller BR (see Figure 9.13), Using our MC, assuming $\alpha_{K^*} = +0.01$, a prediction of $(5.76 \pm 0.14) \times 10^{-7}$ can be obtained, where the uncertainty comes from $\text{BR}(K_L \rightarrow \gamma\gamma)$. This predictions agrees very well with the measured BR.

The statistical uncertainty on my α_{K^*} contributes a systematic uncertainty to the BR measurement. Increasing α_{K^*} by its uncertainty to $+0.134$ increases the BR from 5.82×10^{-7} to 6.00×10^{-7} , while decreasing α_{K^*} by its uncertainty to -0.106 decreases the BR to 5.65×10^{-7} . The larger shift is taken as a symmetric systematic uncertainty, $\pm 0.18 \times 10^{-7}$, on the BR. This uncertainty is also applied to the $\alpha_{K^*} = -0.28$ BR, as seen in the “form factor” entry of Table 9.5, to reflect the error on BR from uncertainty in M_{ee} acceptance.

9.7 Simulation Quality

In the search for $K_L \rightarrow \pi^0 e^+ e^-$ described in the next chapter, the critical feature of the $K_L \rightarrow e^+ e^- \gamma(\gamma)$ simulation is its ability to model spectra in certain variables, chiefly $M_{ee\gamma\gamma}$, θ_{min} , $M_{\gamma\gamma}$, y_γ , and ϕ . The first two of these are addressed in detail in sections 9.3.3 and 9.3.4.

Two-photon mass in the $K_L \rightarrow e^+ e^- \gamma(\gamma)$ MC is discussed for small $M_{\gamma\gamma}$ in section 9.3.1.2. A ratio of data to summed MC is shown for all $M_{\gamma\gamma}$ in Figure 9.18. The data/MC ratio is fairly flat. The region of interest for $K_L \rightarrow \pi^0 e^+ e^-$ is near the pion mass, and the ratio is flat in that neighborhood. The small $K \rightarrow \pi^0 \pi_{Dalitz}^0$ background

near the pion mass might be a concern, but such events also have $M_{ee} < M_{\pi^0}$.

The other two variables listed above, y_γ and ϕ , are the usual phase-space angles, where

$$y_\gamma = 2 \frac{P \cdot (k_1 - k_2)}{M_{K^0}^2 \lambda^{1/2}(1, x, x_\gamma)}.$$

P is the parent particle four momentum, k is photon four momentum, $x = (M_{ee}/M_{K^0})^2$, $x_\gamma = (M_{\gamma\gamma}/M_{K^0})^2$, and $\lambda(a, b, c) = a^2 + b^2 + c^2 - 2(ab + bc + ac)$. In more physical terms, in the photon-pair center-of-mass, y_γ is the cosine of the angle between a photon and the electron-pair vector; see Figure 9.19. The other angle, ϕ , is measured in the parent particle center-of-mass between the plane formed by the two photons and the plane formed by the two electrons. In the $K_L \rightarrow e^+e^-\gamma\gamma$ system, both angles tend to be small because the radiated photon makes a small angle with an electron, and symmetric because the photons are interchangeable. To see the structure in these variables, I plot $\log_{10}(1 - |y_\gamma|)$ and $\log_{10}(1 - |\cos\phi|)$ in Figure 9.20. In general, the MC matches the data well in these plots. There is a deficit of MC for $|y_\gamma|$ very close to 1 ($\log_{10}(1 - |y_\gamma|) < -2.8$), but this is a resolution issue well away from where the cut will be for $K_L \rightarrow \pi^0 e^+ e^-$.

9.8 Conclusions

This chapter presents the measurements $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV}) = (5.40 \pm 0.14_{\text{stat.}} \pm 0.29_{\text{sys.}} \pm 0.18_{\text{BR}}) \times 10^{-7}$ (assuming $\alpha_{K^*} = -0.28$), and $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV}) = (5.82 \pm 0.15_{\text{stat.}} \pm 0.31_{\text{sys.}} \pm 0.19_{\text{BR}}) \times 10^{-7}$ (using the measurement $\alpha_{K^*} = 0.01 \pm 0.12_{\text{stat.}} \pm 0.03_{\text{sys.}}$). The first BR does not disagree with QED predictions or previous measurements, but the second BR better agrees with predictions and is more self consistent. The uncertainty on this BR is smaller than the best previous measurement by a factor of 3.4 [32]. However, unlike previous measurements, this uncertainty is limited by systematics rather than statistics. Future studies of this mode at high flux experiments, such as KTeV99, will have to take care to regard $K_L \rightarrow e^+e^-\gamma\gamma$






Figure 9.18: $\gamma\gamma$ invariant mass in GeV/c^2 after all other cuts. The data/MC ratio is made using the sum of MC distributions.

Figure 9.19: Phase-space variable y_γ .Figure 9.20: Phase-space angles y_γ and ϕ . All cuts are made.

as a precision measurement rather than a rare decay search.

More precise measurements of α_{K^*} have been done with the more abundant $K_L \rightarrow e^+e^-\gamma$. This measurement using $K_L \rightarrow e^+e^-\gamma\gamma$ is done as a check, and does not strongly contradict the $K_L \rightarrow e^+e^-\gamma$ results. Combining uncertainties, this α_{K^*} is 2.0σ larger than the PDG average of -0.28 ± 0.08 [37] and 2.8σ larger than the new NA48 measurement of -0.36 ± 0.06 [38], both of which use $K_L \rightarrow e^+e^-\gamma$. For future studies, it would be of interest to determine if this α_{K^*} is just a statistical fluke or if the form factor is somehow different for $K_L \rightarrow e^+e^-\gamma\gamma$.

Finally, the $K_L \rightarrow e^+e^-\gamma\gamma$ simulation appears able to reproduce the spectra that are in data, for variables important in the $K_L \rightarrow \pi^0 e^+e^-$ search. This check is important because before this experiment there were not enough data events to make such comparisons, and because this simulation will be used to determine how much background is expected to $K_L \rightarrow \pi^0 e^+e^-$.

Chapter 10

$$K_L \rightarrow \pi^0 e^+ e^-$$

This chapter details the final steps in my search for $K_L \rightarrow \pi^0 e^+ e^-$, beyond the basic steps of normalization and studying the background.

10.1 Simulated Sample

The numbers of MC simulated events is listed in Tables 10.1 and 10.2. The simulation used for Table 10.1's sample assumed flat phase space for the $K_L \rightarrow \pi^0 e^+ e^-$ decay, while Table 10.2 used the vector model, which assumes a $K\pi^0\gamma^*$ vertex. The model assumed makes little difference in the acceptance through the EEGG filter tag stage. The following discussion will only refer to the phase-space model sample for $K_L \rightarrow \pi^0 e^+ e^-$, until section 10.8 returns to the vector model.

The simulated $K \rightarrow \pi^0 \pi_{Dalitz}^0$ and $K_L \rightarrow e^+ e^- \gamma \gamma$ event samples described in section 9.1 are re-used in this chapter.

10.2 General Cuts

The basic cuts in this analysis differ from the normalization's cuts described in section 8.3 in a few ways. The EEGG crunch filter tag is used instead of EEGGG. The post-filter reconstruction must find four calorimeter clusters instead of five. Like the normalization, but unlike the $K_L \rightarrow e^+ e^- \gamma \gamma$ study, the signal decay has a neutral pion, so Z_{π^0} can be and is used to reconstruct various cut variables.

Table 10.1: Numbers of $K_L \rightarrow \pi^0 e^+ e^-$ generated. % is the number on that line divided by the number Generated. L1, L2, and L3 trigger and EEGG filter sums are in all spills. L3 triggers are 2e, n-clus tags only.

Stage	Winter		Summer	
	Number	%	Number	%
Generate	615109	100.00	466179	100.00
Good Spills	575325	93.53	411712	88.32
L1 trigger	130080	21.15	98358	21.10
L2 trigger	74785	12.16	56897	12.20
L2 trigger	70613	11.48	53989	11.58
EEGG filter	63606	10.34	44416	10.29

Table 10.2: Numbers of $K_L \rightarrow \pi^0 e^+ e^-$ generated using vector model. % is the number on that line divided by the number Generated. L1, L2, and L3 trigger and EEGG filter sums are in all spills. L3 triggers are 2e, n-clus tags only.

Stage	Winter		Summer	
	Number	%	Number	%
Generate	615109	100.00	466179	100.00
Good Spills	575252	93.52	411651	88.30
L1 trigger	129051	20.98	97846	20.98
L2 trigger	73820	12.00	56429	12.10
L2 trigger	69410	11.28	53321	11.43
EEGG filter	62392	10.14	47107	10.10

10.3 Obvious, Mode-Specific Cuts

The $K_L \rightarrow \pi^0 e^+ e^-$ search uses a blind analysis. Therefore, plots in this section will not use data where a signal peak might be present. Unlike the previous two chapters, plots will generally not show distributions “after all other cuts,” but only after all previously described cuts. This is to illustrate how various backgrounds are suppressed.

10.3.1 $M_{\pi^+ \pi^- \gamma \gamma}$

Figure 10.1 shows the distribution of invariant mass of the particles using the hypothesis that the tracks are pions. Note that *no* TRD cuts have been applied. This shows the large number of $K_L \rightarrow \pi^+ \pi^- \pi^0$ decays present in the EEGG tags. However, a cut requiring $M_{\pi^+ \pi^- \gamma \gamma} > 520 \text{ MeV}/c^2$ strongly suppresses them.




Figure 10.1: $\pi^+ \pi^- \gamma \gamma$ invariant mass, in GeV/c^2 . Only general, non-TRD cuts are made. The solid line is simulated $K_L \rightarrow \pi^0 e^+ e^-$ with arbitrary normalization.

10.3.2 TRD

Figure 10.2 shows Π_{TRD} for the data. The cut on Π_{TRD} is more properly a “general cut” than a “mode-specific” cut, but its effects on K_{e3} backgrounds are more apparent after the mode-specific $M_{\pi^+\pi^-\gamma\gamma}$ cut is applied. After requiring $M_{\pi^+\pi^-\gamma\gamma} > 520 \text{ MeV}/c^2$ and at least one TRD hit per track, there are 12,254 data events where both tracks have $\Pi_{\text{TRD}} < 0.04$, 2,538 events with only one $\Pi_{\text{TRD}} < 0.04$, and 175 events where both tracks have $\Pi_{\text{TRD}} > 0.04$. Following the procedure in section 9.4.1, I would expect 71 of the events passing the Π_{TRD} cut to be K_{e3} s, or 0.58%.




Figure 10.2: TRD Π_{TRD} for both tracks in data events.

10.3.3 M_{ee}

The M_{ee} cut is crucial for getting rid of π_{Dalitz}^0 backgrounds. Figure 10.3 shows M_{ee} for the data and for several MC samples. The $K \rightarrow \pi^0 \pi_{\text{Dalitz}}^0$ and $K_L \rightarrow e^+ e^- \gamma(\gamma)$ samples are normalized by flux and branching ratio, using predicted BR in the case of $K_L \rightarrow e^+ e^- \gamma\gamma$. The low cut requires $M_{ee} > 140 \text{ MeV}/c^2$. This was selected to be $+3\sigma$ from the M_{ee} peak in E799’s $\pi^0 \rightarrow e^+ e^-$ analysis [52]. The data events below

140 MeV/c² are dominated by Dalitz pion decays. This cut costs 18.2% of the signal MC at this stage.

The high cut requires $M_{ee} < 370 \text{ MeV}/c^2$. This is set to be above the hard kinematic bound of $362.7 \text{ MeV}/c^2$, rounded up for resolution.

The $K \rightarrow \pi^0 \pi^0_{Dalitz}$ MC events passing the cut usually have one photon undergoing pair production. The spectrometer magnet then kicks out a soft electron from that pair and a soft electron with opposite charge from the Dalitz pair. Then M_{ee} can be large.

10.3.4 Transverse Momentum

Cuts are made on the $\vec{P}_T^2$ reconstructed using the charged vertex and $\vec{P}_T^2$ using Z_{π^0} , requiring both to be less than $0.001 (\text{GeV}/c)^2$. Because the EEGG crunch tag already made such a requirement, only a handful of events in any sample are lost to the charged vertex $\vec{P}_T^2$ cut. The $Z_{\pi^0} \vec{P}_T^2$ cut removes events without a π^0 , primarily $K_L \rightarrow e^+ e^- \gamma(\gamma)$. Figure 10.4 plots this variable after all previous cuts for some event samples. The normalization is the same as for Figure 10.3.

10.3.5 $M_{ee\gamma\gamma}$

Figure 10.5 shows $M_{ee\gamma\gamma}$, using Z_{π^0} , after all previous cuts. To keep the analysis blind, data with $485 < M_{ee\gamma\gamma} < 510 \text{ MeV}/c^2$ and $130 < M_{\gamma\gamma} < 140 \text{ MeV}/c^2$ are not shown. For comparison, $K_L \rightarrow e^+ e^- \gamma(\gamma)$ MC in the same region are not shown either. Section 10.4.1 discusses the features in this distribution further, as it is essentially a projection of Figure 10.6 onto the $M_{ee\gamma\gamma}$ axis.

Fitting a Gaussian to $M_{ee\gamma\gamma}$ for the $K_L \rightarrow \pi^0 e^+ e^-$ MC give a mean of $497.57 \pm 0.012 \text{ MeV}/c^2$ with a width of $\sigma = 2.527 \pm 0.013 \text{ MeV}/c^2$. The cut is set at the kaon mass $\pm \sim 2\sigma$: $|M_{ee\gamma\gamma} - M_{K^0}| < 5 \text{ MeV}/c^2$.




Figure 10.3: e^+e^- invariant mass, in GeV/c^2 , after general, $M_{\pi+\pi-\gamma\gamma}$, and TRD cuts. The $K_L \rightarrow \pi^0 e^+ e^-$ distribution has arbitrary normalization.




Figure 10.4: Transverse momentum, $\vec{P}_T^2$, using Z_{π^0} , in $(\text{GeV}/c^2)^2$, after general, $M_{\pi^+ \pi^- \gamma \gamma}$, TRD, and M_{ee} cuts. The $K_L \rightarrow \pi^0 e^+ e^-$ distribution has arbitrary normalization.




Figure 10.5: $e^+ e^- \gamma \gamma$ invariant mass, using Z_{π^0} , in $(\text{GeV}/c^2)^2$, after general, $M_{\pi^+ \pi^- \gamma \gamma}$, TRD, M_{ee} , and $\vec{P}_T^2$ cuts. The $K_L \rightarrow \pi^0 e^+ e^-$ distribution has arbitrary normalization. The fit “box” has been excluded in the data and $K_L \rightarrow e^+ e^- \gamma(\gamma)$ samples. The excess data below $420 \text{ MeV}/c^2$ are mostly from K_{e4} decays.

10.4 Non-Obvious, Mode-Specific Cuts

For the variables in the previous section, the optimum cuts can be selected using kinematic and resolution limits. This section describes the method for selecting cuts on kinematic angles and $M_{\gamma\gamma}$, where the optimum cuts are not so obvious. First, a method for estimating background is developed. This is then used to pick angle cuts that would give the lowest upper limit in the absence of signal. Then the $M_{\gamma\gamma}$ cut is selected by looking at $M_{\gamma\gamma}$ resolution in the phase space allowed by the angle cuts.

10.4.1 $M_{\gamma\gamma}, M_{ee\gamma\gamma}$ Plane

This rare decay search uses a background estimate to calculate a confidence interval for the branching ratio. The background estimate is made by fitting the data and interpolating from outside the signal region to inside. This interpolation is done in two dimensions, in the variables $M_{\gamma\gamma}$ and $M_{ee\gamma\gamma}$. In the $M_{\gamma\gamma}, M_{ee\gamma\gamma}$ plane, the signal appears at one point, smeared by resolution. However the backgrounds can appear above or below the signal, in either variable, adding constraints to the interpolation. Furthermore, the area outside the signal region is much larger than the area inside, further constraining the fit.

Figure 10.6 plots $M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ for data events passing all previous cuts. Several regions are outlined and labeled, each highlighting a feature of the distribution.

The “box” is the region where no data events are used while the analysis is blind. The box covers the region $485 < M_{ee\gamma\gamma} < 510 \text{ MeV}/c^2$ and $130 < M_{\gamma\gamma} < 140 \text{ MeV}/c^2$. The “bump” is a region inside the box indicating the cuts that select $K_L \rightarrow \pi^0 e^+ e^-$ candidates. While the analysis is blind, the bump may be placed anywhere within the box. At this point, the bump includes the region $|M_{ee\gamma\gamma} - M_{K^0}| < 11 \text{ MeV}/c^2$ and $|M_{\gamma\gamma} - 135.20| < 2.65 \text{ MeV}/c^2$. (Actually, the bump is the same for the final selections, but it does not have to be.)




Figure 10.6: $M_{\gamma\gamma}$ vs $M_{e\bar{e}\gamma\gamma}$ in GeV/c^2 , for data, before kinematic cuts.

The “strip” covers the region $M_{ee\gamma\gamma} < 465 \text{ MeV}/c^2$ and $130 < M_{\gamma\gamma} < 140 \text{ MeV}/c^2$. Extra backgrounds with $M_{\gamma\gamma}$ near the pion mass appear here, but they fall off quickly in $M_{ee\gamma\gamma}$, so that very little enters the bump. The strip is excluded from fits to interpolate background.

The “swath” is the region where most $K_L \rightarrow e^+e^-\gamma(\gamma)$ appear; see Figure 10.7. The swath is diagonal because $M_{ee\gamma\gamma}$ uses Z_{π^0} . When $M_{\gamma\gamma}$ is higher than than M_{π^0} , then Z_{π^0} is shifted from the charged-vertex Z to reduce $M_{\gamma\gamma}$, which reduces $M_{ee\gamma\gamma}$ proportionally. When $M_{\gamma\gamma}$ is lower than M_{π^0} , the Z_{π^0} shifts to increase $M_{\gamma\gamma}$ and $M_{ee\gamma\gamma}$. If the charged-vertex $M_{ee\gamma\gamma}$ were to be used, then the swath would be vertical because $M_{\gamma\gamma}$ is mostly uncorrelated with $M_{ee\gamma\gamma}$ in $K_L \rightarrow e^+e^-\gamma(\gamma)$ events.




Figure 10.7: $M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , for simulated $K_L \rightarrow e^+e^-\gamma(\gamma)$, before kinematic cuts.

The function that is interpolated inside the bump to determine the background is the sum of a plane and the $K_L \rightarrow e^+e^-\gamma(\gamma)$ sample:

$$f(M_{\gamma\gamma}, M_{ee\gamma\gamma}) = (P_1 + P_2 M_{\gamma\gamma} + P_3 M_{ee\gamma\gamma}) + P_4 A(M_{\gamma\gamma}, M_{ee\gamma\gamma}).$$

P_i is a parameter of the log-likelihood fit and $A(M_{\gamma\gamma}, M_{ee\gamma\gamma})$ is the $K_L \rightarrow e^+e^-\gamma(\gamma)$ sample distribution, which appears in Figure 10.7. Neither the plane nor P_4 are allowed to dip below zero. The function $f(M_{\gamma\gamma}, M_{ee\gamma\gamma})$ is fit to the data outside the strip and the box, with the same binning as in Figure 10.7. Figure 10.8 is the resulting $f(M_{\gamma\gamma}, M_{ee\gamma\gamma})$ and the data. Integrating $f(M_{\gamma\gamma}, M_{ee\gamma\gamma})$ over the bump gives an estimated background of 36.91 ± 1.98 events. Only 0.180 ± 0.043 events come from the plane, and the uncertainty comes from the fit and MC statistics.




Figure 10.8: $M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , before kinematic cuts. The shaded boxes are the results of the background fit. The empty boxes are data, plotted on a linear Z scale. Data events inside the “box” are not shown.

10.4.2 $K_L \rightarrow e^+e^-\gamma\gamma$ Background Checks

The above background estimate shows that $K_L \rightarrow e^+e^-\gamma(\gamma)$ is the most important background remaining. Chapter 9 compared data and MC for $K_L \rightarrow e^+e^-\gamma(\gamma)$, but the analysis is somewhat different here. As a check, some additional comparisons are shown here. The events compared are from outside the box, to keep the analysis

blind, but from inside the swath, to keep the $K_L \rightarrow e^+e^-\gamma(\gamma)$ purity high.

Figures 10.9 and 10.10 project Figure 10.8 onto the $M_{\gamma\gamma}$ and $M_{ee\gamma\gamma}$ axes, for data and fit function. Any bins in Figure 10.8 with centers inside the swath are included.




Figure 10.9: $M_{\gamma\gamma}$ in GeV/c^2 , before kinematic cuts, in $M_{\gamma\gamma}$, $M_{ee\gamma\gamma}$ “swath.” Crosses are data and solid line is fit.

Figures 10.11, 10.12, and 10.13 compare the kinematic angle distributions inside the swath and outside the box. θ_{min} is used extensively in chapter 9, while y_γ and ϕ are introduced in section 9.7. The $K_L \rightarrow e^+e^-\gamma(\gamma)$ samples in all three plots are normalized with the P_4 factor from the background fit. The $K_L \rightarrow \pi^0 e^+e^-$ MC distribution is shown for comparison, and includes events from inside the box.

All three angle variables show good data-MC agreement for small angles, where most $K_L \rightarrow e^+e^-\gamma(\gamma)$ events occur. There tends to be a small excess of data for large angles. These are most likely from accidental backgrounds represented by the plane in the background fit. The plots also show how $K_L \rightarrow \pi^0 e^+e^-$ events are much more uniformly distributed in these variables than the $K_L \rightarrow e^+e^-\gamma(\gamma)$ background, a fact which will be exploited in the next section.




Figure 10.10: $M_{ee\gamma\gamma}$ in GeV/c^2 , before kinematic cuts, in $M_{\gamma\gamma}$, $M_{ee\gamma\gamma}$ “swath.” Crosses are data and solid line is fit.




Figure 10.11: θ_{min} for events in swath but outside box. Simulated $K_L \rightarrow \pi^0 e^+ e^-$ events are from inside box, and have arbitrary normalization.




Figure 10.12: Absolute value of y_γ for events in swath but outside box. Simulated $K_L \rightarrow \pi^0 e^+ e^-$ events are from inside box, and have arbitrary normalization.




Figure 10.13: Absolute value of $|\phi|$ for events in swath but outside box. Simulated $K_L \rightarrow \pi^0 e^+ e^-$ events are from inside box, and have arbitrary normalization.

10.4.3 Optimizing Kinematic Cuts

Figures 10.11, 10.12, and 10.13 show that some kind of cuts on kinematic variables would be helpful in reducing the ~ 37 expected $K_L \rightarrow e^+e^-\gamma(\gamma)$ background events. The best combination of cuts is found by trying out a number of combinations and seeing how they perform.

The performance figure-of-merit used is a predicted branching-ratio upper limit. The cut combination under consideration is applied to data (outside the box and strip), $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC, and $K_L \rightarrow \pi^0e^+e^-$ MC. The $K_L \rightarrow \pi^0e^+e^-$ MC sample (multiplying by the TRD acceptance) gives the acceptance, ϵ , for those cuts. The $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC sample is used in fitting the data sample to get a background estimate, μ_{bg} , and uncertainty, σ_{bg} . To get a figure to compare with other cut combinations, I calculate what the upper limit on the number of signal events, (μ_u) , would be if the number of events passing the cuts in data, n_o , was the same as the background estimate: $n_o = \mu_{bg}$. This numeric upper limit, divided by the product of acceptance and flux, gives an expected BR upper limit in the absence of signal:

$$B_u = \frac{\mu_u}{\epsilon N_{K_L}}$$

The cut combination that minimizes B_u is optimum.

The numeric upper limit, μ_u , here refers to the upper end of a confidence interval. The confidence interval (μ_l, μ_u) is constructed such that if the experiment were to be repeated many times, the true value, μ_t , would be inside 90% of the experiments' confidence intervals. In other words, if $\mu_t > \mu_u$, then more than 90% of the experiments would see more events than this experiment. The method of interval construction is based on the unified procedure of Feldman and Cousins for a Poisson process with background [53]. This method differs from theirs by allowing the background estimate to have some uncertainty. While reference [53] assumes that the background is known exactly, we assume that the background estimate would vary with repetition of the

experiment, with a Gaussian probability distribution function. Appendix B of reference [46] details the procedure. In brief, the probability of each n is reweighted by the integral of the background pdf.

When estimating an upper limit to evaluate some combination of cuts, I assume $n_o = \mu_{bg}$ and calculate μ_u . However, the method for calculating μ_u assumes that n_o is an integer, so I interpolate $\mu_u(n_o)$ between integer values of n_o . To be more correct, the mean μ_u should be found for a Poisson distributed n_o with mean equal to μ_{bg} . Another approximation I make is to assume $\sigma_{bg} = 0$ when computing μ_u . However, making these corrections has negligible effects for most of the cut combinations considered. The function $\mu_u = 2.44 + 0.834(\mu_{bg})^{0.794}$ fits well for the data points $\mu_{bg} = n_o = 0, 1, 2, 3, 4, 5$. Most cuts considered produced $\mu_{bg} < 5$, so this was the figure-of-merit used for the cuts.




Figure 10.14: Estimated branching ratio limit, as function of cuts on θ_{min} and $|y_\gamma|$.

The optimization considers a grid of cut values for minimum θ_{min} , maximum $|y_\gamma|$, and maximum $|\phi|$ allowed. As it turns out, the optimum $|\phi|$ cut is to make no cut at all. This makes sense if one looks in Figure 10.13. The signal spectrum of $|\phi|$ peaks

(albeit not sharply) in the same place as the background, The signal spectra for $|y_\gamma|$ and θ_{min} have minima where the background peaks. Also, all three variables are somewhat correlated, so having too many cuts only overconstrains the acceptance. In particular, ϕ is highly correlated with y_γ in $K_L \rightarrow e^+e^-\gamma\gamma$, so cutting on both is redundant at best.

So the optimization described here only uses θ_{min} and $|y_\gamma|$. This is nice because two variables are much easier to visualize and monitor than three. Figure 10.14 shows the BR as a function of cut. The cut values are at the bin centers. By eyeballing the plot, one can see that any pair of cuts in the region $\theta_{min} > \sim 0.3-0.5$, $|y_\gamma| < \sim 0.75-0.85$ would be satisfactory. Fitting a two-dimensional quadratic function to the points near the minimum gives the best fit at $\theta_{min} > 0.340 \pm 0.118$ and $|y_\gamma| < 0.788 \pm 0.018$, where the errors come from the fit.

Figure 10.15 shows the effect of these cuts on the $M_{\gamma\gamma}$, $M_{ee\gamma\gamma}$ plane. Integrating the fit inside the bump gives a background estimate of 1.06 ± 0.41 events.

10.4.4 $M_{\gamma\gamma}$

After all other cuts are finalized, the $M_{\gamma\gamma}$ cut is finalized by studying $M_{\gamma\gamma}$ resolution. Previous to this stage a cut of $|M_{\gamma\gamma} - 135.20| < 2.65 \text{ MeV}/c^2$ has been used. Some foresight of this stage was used in selecting this cut, because a substantial change might make other cuts sub-optimal.

A cut on $M_{\gamma\gamma}$ is one of the most obvious ways to get rid of $K_L \rightarrow e^+e^-\gamma\gamma$ background to $K_L \rightarrow \pi^0 e^+ e^-$. The background shows a broad continuum in $M_{\gamma\gamma}$, while the signal is a delta function. The only limit to the rejection power of the cut is detector resolution.

Unlike most other variables in the $K_L \rightarrow \pi^0 e^+ e^-$ search, our understanding of $M_{\gamma\gamma}$ in MC can be checked against “data.” Not actual $K_L \rightarrow \pi^0 e^+ e^-$ data, of course, but the next best thing: $K_L \rightarrow \pi^+ \pi^- \pi^0$. There are many such decays amongst the




Figure 10.15: $M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , after all other cuts. The shaded boxes are the results of the background fit. The empty boxes are data, plotted on a linear Z scale. Data events inside the “box” are not shown.

EEGG filter tags, as shown in Figure 10.1. This mode contains the necessary π^0 , and when the $\pi^\pm$ are assumed to be electrons, it often has large enough M_{ee} to pass the M_{ee} cut.

The cuts made to identify $K_L \rightarrow \pi^+\pi^-\pi^0$ events are identical to those for $K_L \rightarrow \pi^0 e^+ e^-$, with the following exceptions:

- $\Pi_{\text{TRD}} > 0.04$ for both tracks.
- No E/p cut.
- No $M_{ee\gamma\gamma}$ cut.
- Minimum distance between cluster positions more than 15 cm. This is to prevent energy from pions showers from contaminating photon showers.
- $|M_{\pi^+\pi^-\gamma\gamma} - M_{K^0}| < 5 \text{ MeV}/c^2$.

Reconstructing M_{ee} for these events gives the distributions shown in Figure 10.16. Because of concerns that M_{ee} could correlate with the resolution of $M_{\gamma\gamma}$ due to poor resolution at small track angles, the $K_L \rightarrow \pi^+\pi^-\pi^0$ are weighted by the function shown in Figure 10.16. Because M_{ee} is kinematically limited in $K_L \rightarrow \pi^+\pi^-\pi^0$, only events with $140 < M_{ee} < 225 \text{ MeV}/c^2$ are considered.

Figure 10.17 fits $M_{\gamma\gamma}$ for events after these cuts and weights. Fitting a Gaussian gives a mean of $135.17 \pm 0.015 \text{ MeV}/c^2$ with a width of 1.328 ± 0.013 for the data and a mean of 135.24 ± 0.010 with a width of 1.154 ± 0.009 for the MC. The means do not disagree much, but the MC resolution is $13.1 \pm 1.4\%$ too small. Fitting $M_{\gamma\gamma}$ for the MC (not shown), allowing the full range of M_{ee} , gives a mean of $135.27 \pm 0.006 \text{ MeV}/c^2$ with a width of 1.131 ± 0.005 . Based on the (more conservative) data resolution and a desire for a $\pm 2\sigma$ cut, the cut is left unchanged from what was used previously: $|M_{\gamma\gamma} - 135.2| < 2.65 \text{ MeV}/c^2$.

10.5 Cut Summary and Final Background Estimate

The final cuts to identify $K_L \rightarrow \pi^0 e^+ e^-$ candidates are as follows:




Figure 10.16: e^+e^- invariant mass, in $(\text{GeV}/c^2)^2$. All non- $M_{\gamma\gamma}$ cuts have been applied to MC. The M_{ee} cut shown is only used for the $M_{\gamma\gamma}$ resolution study.




Figure 10.17: $\gamma\gamma$ invariant mass, in $(\text{GeV}/c^2)^2$, for $K_L \rightarrow \pi^+\pi^-\pi^0$ data (Left) and simulated $K_L \rightarrow \pi^0 e^+ e^-$ (Right).

- General cuts (section 8.3) for four cluster events.
- $M_{\pi^+\pi^-\gamma\gamma} > 520 \text{ MeV}/c^2$.
- $140 < M_{ee} < 370 \text{ MeV}/c^2$.
- $\vec{P}_T^2 < 0.001 (\text{GeV}/c)^2$, using charged vertex Z .
- $\vec{P}_T^2 < 0.001 (\text{GeV}/c)^2$, using Z_{π^0} .
- $|M_{ee\gamma\gamma} - M_{K^0}| < 5 \text{ MeV}/c^2$.
- $\theta_{min} > 0.349 \text{ radians}$.
- $|y_\gamma| < 0.788$.
- $|M_{\gamma\gamma} - 135.2| < 2.65 \text{ MeV}/c^2$.

Because the $M_{\gamma\gamma}$ cut did not change in section 10.4.4, the expected background is still 1.06 ± 0.41 events. Of these, $0.913 \pm 0.237 \pm 0.112$ are from the $K_L \rightarrow e^+e^-\gamma(\gamma)$ component of the fit; the first uncertainty is from the fit's P_4 uncertainty (which is mostly from data statistics) and the second uncertainty is from the statistics of $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC in the bump. The plane contributes 0.145 ± 0.313 events to the bump. The $K_L \rightarrow \pi^0 e^+e^-$ MC sample has 39,249 events passing all cuts, giving a TRD-corrected acceptance of 0.03758. Using the K_L flux measurement of 265.07×10^9 , the single-event sensitivity is 1.003×10^{-10} .

My upper limit will be calculated using the method outlined in section 10.4.3. If one (1) event is observed, then this method gives $\mu_u = 3.3095$ and $\text{BR}(K_L \rightarrow \pi^0 e^+e^-) < 3.32 \times 10^{-10}$. On the other hand, if four or more events are observed passing all cuts, then the unified approach to confidence intervals dictates, for 1.06 ± 0.41 background, that the lower limit μ_l becomes greater than zero, and an observation of signal must be claimed. In the case of exactly four (4) events observed, $(\mu_l, \mu_u) = (0.5295, 7.7945)$ and $5.31 \times 10^{-11} < \text{BR} < 7.81 \times 10^{-10}$.

10.6 Alternative Estimates of Backgrounds

To check the background fit, a number of background estimates using alternative methods are presented below.

10.6.1 $K_L \rightarrow \pi^0 \gamma e^+ e^-$

As part of E799's study of $K_L \rightarrow \pi^0 \gamma e^+ e^-$, a sample of $K_L \rightarrow \pi^0 \gamma e^+ e^-$ events was simulated. Though this sample is equivalent to 50 times the E799 data set, none of the events pass the $K_L \rightarrow \pi^0 e^+ e^-$ cuts [46]. Therefore, no significant background is expected from this decay mode.

10.6.2 $K \rightarrow \pi^0 \pi_{Dalitz}^0$

The background from $K \rightarrow \pi^0 \pi_{Dalitz}^0$ can be estimated using its MC sample. No events pass all cuts, but 546 pass all cuts except $M_{\gamma\gamma}$ and $M_{ee\gamma\gamma}$, and 340 of these fall into the plot in Figure 10.18. By fitting a Gaussian to the $M_{\gamma\gamma}$ distribution after the $M_{ee\gamma\gamma}$ cut, an estimate of $1.8 \pm \sim 2$ events can be obtained. Scaling by the flux, the number of generated events, and the branching ratio, the background level is about 0.05 events.

10.6.3 $K \rightarrow \pi^0 \pi^0 \pi_{Dalitz}^0$

The background from $K \rightarrow \pi^0 \pi^0 \pi_{Dalitz}^0$ is considered using the same MC sample described in section 9.4.2. Figure 10.19 shows these events. Fitting $M_{ee\gamma\gamma}$ after the $M_{\gamma\gamma}$ cut suggests about 0.025 events in the bump. Using the scale factor from section 9.4.2 gives the background from $K \rightarrow \pi^0 \pi^0 \pi_{Dalitz}^0$ as ~ 0.08 events.

10.6.4 “Strip”

A cluster of data events appears in the strip region ($130 < M_{\gamma\gamma} < 140 \text{ MeV}/c^2$ and $400 < M_{ee\gamma\gamma} < 465 \text{ MeV}/c^2$) of Figures 10.6 and 10.6. These decays contain a π^0 and




Figure 10.18: $M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c², after all other cuts, for simulated $K \rightarrow \pi^0 \pi^0_{Dalitz}$.




Figure 10.19: $M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c², after all other cuts, for simulated $K \rightarrow \pi^0 \pi^0 \pi^0_{Dalitz}$. Note that the $M_{ee\gamma\gamma}$ scale is expanded in this plot.

frequently a $\pi^\pm$; see Figure 10.20. Sections 10.6.2 and 10.6.3 show that $K \rightarrow \pi^0 \pi_{Dalitz}^0$ and $K \rightarrow \pi^0 \pi^0 \pi_{Dalitz}^0$ contribute few events to the strip. Therefore, most strip event are probably K_{e3} decays with a π^0 , possibly K_{e4} ($K_L \rightarrow \pi^0 \pi^\pm e^\mp \nu$).

The concern regarding the strip is that some backgrounds will spill out into the bump. The fit does not use the strip, so whatever background is making events in the strip is not reflected in the fit, and therefore it is not in the background estimate. As a check, there are six data events outside the strip and box after all cuts except $M_{ee\gamma\gamma}$; see Figure 10.21. Figure 10.21 is Figure 10.15, with the $M_{\gamma\gamma}$ cut applied and projected onto the $M_{ee\gamma\gamma}$ axis. Integrating the fit over the same region gives an estimate of 2.1 ± 3.8 events. Most of the uncertainty comes from fit plane parameter errors. Thus the number of data events outside the box and strip in Figure 10.21 is consistent with the background fit. Even if the large uncertainty on the background is neglected, there is still a 1.6% probability that six events could appear when 2.1 are expected. Therefore it appears that the backgrounds in this region are understood to the degree they can be, and there is no significant pollution from the background process that dominates the strip.

10.6.5 K_{e3}

The number of background from K_{e3} where $M_{\gamma\gamma}$ is not near the pion mass can be estimated by removing TRD cuts. Figure 10.22 shows the $M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ for events which have at least one Π_{TRD} of more than 0.04. Most have $\Pi_{TRD} < 0.04$ for one track and $\Pi_{TRD} > 0.04$ for the other. There are 200 events outside the strip region (all but seven have one $\Pi_{TRD} < 0.04$). Scaling into the bump region gives an average of 0.548 ± 0.039 events. There are actually three events inside the bump. This gives a 90% C.L. confidence interval of 0.555 to 6.88 events for the true, average number of events in the bump. From Figure 5.6, the accepted:rejected ratio of single pions with a $\Pi_{TRD} < 0.04$ cut is $1 : 25.29 \pm 0.31$. The K_{e3} background estimate then has a confidence




Figure 10.20: TRD Π_{TRD} for data events with $130 < M_{\gamma\gamma} < 140 \text{ MeV}/c^2$ and $400 < M_{e\gamma\gamma} < 465 \text{ MeV}/c^2$, after all other cuts. The only region accepted by the TRD cuts is the bin that has 18 events.




Figure 10.21: $e^+e^-\gamma\gamma$ invariant mass after all other cuts, in GeV/c^2 . The squares are the background fit, which excludes the strip, and the fit error bars are not shown. The fit and the data use the same $5 \text{ MeV}/c^2$ bins. No data from inside the box are shown. The $M_{\gamma\gamma}$ cut is applied.

interval of 0.0219 to 0.272 events.

The background level here indicates the usefulness of a TRD detector system in this experiment. Those three events, or an expected 0.555–6.88 at 90% C.L., are only rejected by their transition radiation, despite the kinematic cuts on $\vec{P}_T^2$, M_{ee} , $M_{\gamma\gamma}$, $M_{ee\gamma\gamma}$, y_γ , and θ_{min} .




Figure 10.22: $M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , after all other cuts, for data *failing* Π_{TRD} cut. Thus all events shown here, including those in the box, are rejected.

10.6.6 $K_L \rightarrow e^+e^-\gamma(\gamma)$

Finally, background from $K_L \rightarrow e^+e^-\gamma(\gamma)$ can be estimated by scaling MC. The $K_L \rightarrow e^+e^-\gamma(\gamma)$ MC sample contains 67 events that pass all $K_L \rightarrow \pi^0 e^+e^-$ cuts, all of them generated as $K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV}$ decays. Dividing by the number of generated events gives an acceptance of $(1.09 \pm 0.13) \times 10^{-5}$. Multiplying by the flux, the TRD acceptance, and using the branching ratio found in the last chapter (5.40×10^{-7}) gives a background estimate of 1.47 ± 0.21 events. The uncertainty here comes from

Table 10.3: Alternative estimates of background to $K_L \rightarrow \pi^0 e^+ e^-$.

source	method	level
$K_L \rightarrow \pi^0 \gamma e^+ e^-$	MC	none
$K \rightarrow \pi^0 \pi_{Dalitz}^0$	MC	~ 0.05
$K \rightarrow \pi^0 \pi^0 \pi_{Dalitz}^0$	MC	~ 0.08
“strip”	data-fit	$\sim \text{none}$
K_{e3}	data	0.0219–0.272 at 90% C.L.
$K_L \rightarrow e^+ e^- \gamma(\gamma)$	MC	1.34 ± 0.19

MC statistics and branching ratio, which includes flux uncertainties.

As discussed in section 9.6, the $K_L \rightarrow e^+ e^- \gamma(\gamma)$ MC tends to have more events at high M_{ee} than the data, and the $K_L \rightarrow \pi^0 e^+ e^-$ analysis requires $140 < M_{ee} < 370 \text{ MeV}/c^2$. Looking at events that pass all $K_L \rightarrow e^+ e^- \gamma \gamma$ cuts, 38.8% of $K_L \rightarrow e^+ e^- \gamma(\gamma)$ MC pass the $K_L \rightarrow \pi^0 e^+ e^-$ M_{ee} cut, while 35.3% of $K_L \rightarrow e^+ e^- \gamma \gamma$ candidates in data pass the cut. This suggests that predictions from MC of the number of $K_L \rightarrow e^+ e^- \gamma \gamma$ background to $K_L \rightarrow \pi^0 e^+ e^-$ will be high by 9.8%. Including this factor reduces this background estimate to 1.34 ± 0.19 events.

10.6.7 Summary

Table 10.3 lists the background estimates given above. The first five entries are backgrounds that the plane portion of the background fit is intended to address. These five add up to ~ 0.15 – 0.40 . This agrees with the plane background estimate of 0.145 ± 0.313 events.

The $K_L \rightarrow e^+ e^- \gamma \gamma$ -MC background estimate — the last entry in Table 10.3 — is intended to be addressed by the $K_L \rightarrow e^+ e^- \gamma(\gamma)$ portion of the background fit — $0.91 \pm 0.24_{fit} \pm 0.11_{MCstat}$ events. The $K_L \rightarrow e^+ e^- \gamma \gamma$ -MC estimate is 0.43 ± 0.25 events larger than the $K_L \rightarrow e^+ e^- \gamma(\gamma)$ portion of the background fit; the uncertainties come from the fit and the BR, but not the MC statistics because both measurements use

the same 67 $K_L \rightarrow e^+e^-\gamma\gamma$ MC events. This 1.7σ difference casts a bit of doubt on the fit estimate. However, the fit estimate contains our best information on acceptance conditions in the immediate neighborhood in cut space of the signal region. Section 10.4.2 shows that we understand that neighborhood fairly well. Also, the fit estimate is smaller than the $K_L \rightarrow e^+e^-\gamma\gamma$ -MC estimate, so the fit estimate would provide a more conservative, larger BR limit on $K_L \rightarrow \pi^0 e^+ e^-$. So I will use the fit estimate of background for measuring $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$.

10.7 Result

Two events in the data sample pass all cuts. Figure 10.23 shows these events in the “opened” data box. With a background estimate of 1.06 ± 0.41 events the confidence interval is $(0, 4.8515)$ events, and the upper limit on the branching ratio of $K_L \rightarrow \pi^0 e^+ e^-$ is 4.86×10^{-10} , at the 90% confidence level.




Figure 10.23: $M_{\gamma\gamma}$ vs $M_{ee\gamma\gamma}$ in GeV/c^2 , after all other cuts. The box is open. The shaded bins are $K_L \rightarrow \pi^0 e^+ e^-$ MC, normalized to 1.

This confidence interval is calculated using the method outlined in section 10.4.3. If the background uncertainty is neglected, the upper limit drops by a mere 0.0010 events.

Low statistical certainty (two events) dominates the upper limit. However, the studies of flux and $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma)$ showed the presence and size of systematic uncertainties in the sensitivity of this experiment. To check the magnitude of the effect of the systematic uncertainties, I find the change in upper limit by a numeric integration over sensitivity. In this procedure the sensitivity S is varied, μ_{bg} and σ_{bg} are scaled by S , and a new upper limit $\mu_u(S)$ is found. The $\mu_u(S)$ is weighted by a Gaussian function of S with mean of one and sigma equal to the uncertainty. If sigma is 5.5%, approximately the systematic uncertainty found for $\text{BR}(K_L \rightarrow e^+e^-\gamma\gamma, E_\gamma^* > 5 \text{ MeV})$, then the mean μ_u is only 0.0019 events larger than μ_u with no systematic uncertainty. Even with sigma as large as 11%, the upper limit only rises by 0.0099 events. Therefore, no change in the upper limit from sensitivity uncertainties is assigned.

10.8 Vector Model Result

To relate this measurement to the standard model, the analysis is repeated using the vector model of $K_L \rightarrow \pi^0 e^+ e^-$ described in section 6.1.2.3 to get the acceptance. The only difference in the search is that the optimum θ_{min}, y_γ cuts might be different. When the cut optimization procedure is repeated the optimum kinematic cuts are found at $\theta_{min} > 0.560 \pm 0.229$ and $|y_\gamma| < 0.791 \pm 0.018$. Figure 10.24 shows the BR as a function of θ_{min} and y_γ cuts. The optimum cuts do not differ significantly from those found with phase-space $K_L \rightarrow \pi^0 e^+ e^-$, so all cuts are kept the same as for the phase-space model analysis. Then the background estimate and number of events observed do not change, only the signal acceptance.

The vector-model MC sample has 32,754 events that pass all cuts, giving a TRD-corrected acceptance of 3.136%. The single-event sensitivity is 1.201×10^{-10} , and the




Figure 10.24: Estimated branching ratio limit, as function of cuts on θ_{min} and $|y_\gamma|$, assuming vector model for $K_L \rightarrow \pi^0 e^+ e^-$.

branching ratio upper limit is 5.83×10^{-10} .

The 17% drop in acceptance going from the phase-space to vector model can mostly be attributed to M_{ee} . When the M_{ee} cut is dropped, the acceptance is 4.86% using phase-space and 4.78% using vector. As might be expected when the e^+e^- pair is produced by a virtual photon, M_{ee} tends to be lower using a vector model, as seen in Figure 10.25.




Figure 10.25: e^+e^- mass, in GeV/c^2 , for $K_L \rightarrow \pi^0 e^+ e^-$ MC samples after all other cuts.

Next, this branching ratio can be used to limit CKM matrix elements. Using the equation 1.1, $|\text{Im}\lambda_t| < 15.60 \times 10^{-3}$ at the 90% C.L. As mentioned in section 1.2.1, a global fit of standard model parameters including ϵ'/ϵ gives $\text{Im}\lambda_t = (1.38 \pm 0.14) \times 10^{-4}$ [14]. To see what constraint this BR makes on η , the imaginary magnitude of the CKM matrix, I use $\text{Im}\lambda_t = |V_{cb}|^2 |V_{us}| \eta$ with the central values $|V_{cb}| = 0.0395 \pm 0.0017$ and $|V_{us}| = 0.2196 \pm 0.0023$ [37]. This gives $|\eta| < 4.55$ at the 90% confidence level. From the $\text{Im}\lambda_t$ limit, it can be seen that $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ does not yet provide a competitive constraint of the CKM matrix.

10.9 Conclusions

To sum up this experiment: the decay $K_L \rightarrow \pi^0 e^+ e^-$ is sought. The background is expected to be mostly $(86 \pm \%)$ (0.91 ± 0.26 events out of 1.06 ± 0.41) the rare $K_L \rightarrow e^+ e^- \gamma \gamma$ decays. For this reason $K_L \rightarrow e^+ e^- \gamma \gamma$ is studied; 1,578 $K_L \rightarrow e^+ e^- \gamma \gamma, E_\gamma^* > 5$ MeV candidate events are found, of which 1,516.5 remain after background subtraction. These events are used to make the measurements $\text{BR}(K_L \rightarrow e^+ e^- \gamma \gamma, E_\gamma^* > 5 \text{ MeV}) = (5.40 \pm 0.14_{\text{stat.}} \pm 0.29_{\text{sys.}} \pm 0.18_{\text{BR}}) \times 10^{-7}$ (assuming $\alpha_{K^*} = -0.28$), $\alpha_{K^*} = 0.015 \pm 0.12_{\text{stat.}} \pm 0.03_{\text{sys.}}$, and $\text{BR}(K_L \rightarrow e^+ e^- \gamma \gamma, E_\gamma^* > 5 \text{ MeV}) = (5.82 \pm 0.15_{\text{stat.}} \pm 0.31_{\text{sys.}} \pm 0.19_{\text{BR}}) \times 10^{-7}$ (using the measured α_{K^*}).

Furthermore, a large sample of simulated $K_L \rightarrow e^+ e^- \gamma(\gamma)$ is created with similar properties to these 1,578 data events. The simulated $K_L \rightarrow e^+ e^- \gamma(\gamma)$ are used to create a function to fit to data in event-parameter space near where a $K_L \rightarrow \pi^0 e^+ e^-$ is expected ($M_{\gamma\gamma}$ vs. $M_{e\gamma\gamma}$ outside the box). Interpolating this fitted function gives the number of background events expected: 1.06 ± 0.41 total, 0.91 ± 0.26 from $K_L \rightarrow e^+ e^- \gamma(\gamma)$, and 0.145 ± 0.313 from other decays (like K_{e3}). The same cuts accept two events from the E799 data set. The K_L flux is measured to be 2.651×10^{11} decays in the fiducial region, the $K_L \rightarrow \pi^0 e^+ e^-$ acceptance is measured to be 3.76%, and the single-event sensitivity is 1.003×10^{-10} . Based on the number of background and signal, I conclude that there is no evidence for $K_L \rightarrow \pi^0 e^+ e^-$ decay and set an upper limit for the branching ratio of $\Gamma(K_L \rightarrow \pi^0 e^+ e^-) / \Gamma(K_L \rightarrow \text{all}) < 4.86 \times 10^{-10}$ at the 90% confidence level. Assuming a vector model for the $K_L \rightarrow \pi^0 e^+ e^-$ decay mechanics reduces the acceptance somewhat, and the upper limit becomes 5.83×10^{-10} .

The theoretical implications of these analyses are limited. The $K_L \rightarrow e^+ e^- \gamma \gamma$ results do not significantly disagree with theory. The $K_L \rightarrow \pi^0 e^+ e^-$ results are not sensitive enough to impact any predictions. The largest $K_L \rightarrow \pi^0 e^+ e^-$ BR prediction, 3×10^{-10} , from supersymmetric extension of the standard model, is not yet directly

excluded by experimental results. As for the standard model and the CKM matrix, the vector-model BR limits the direct CP-violation parameter $\text{Im}\lambda_t$ to be less than 15.60×10^{-3} ; however, much more precise limits of $\text{Im}\lambda_t$ are possible, including those obtained by measuring direct CP-violation in $K \rightarrow \pi\pi$ decays.

Compared with previous experiments, this experiment makes dramatic improvements. The upper limit on $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ is a factor of 8.7 smaller than the best, previous limit, from experiment E799-i. The combined uncertainty on $\text{BR}(K_L \rightarrow e^+ e^- \gamma\gamma, E_\gamma^* > 5 \text{ MeV})$ of 5.86% is 4.2 times smaller than the combined uncertainty in the best, previous measurement, also E799-i.

As for future experiments, they will need more K_L decays to find $K_L \rightarrow \pi^0 e^+ e^-$. One candidate experiment is KTeV99, nearly identical to E799 but with more integrated luminosity. The method presented here should continue to be viable, if the background sensitivity can be reduced. Experimental resolution of $M_{\gamma\gamma}$, and therefore precise calorimetry, will probably be crucial for this. Particle identification, such as with a TRD system, may continue to be vital. With higher K_L fluxes, future searches for $K_L \rightarrow \pi^0 e^+ e^-$ may well use the relatively abundant $K_L \rightarrow e^+ e^- \gamma\gamma$ as the normalization mode. In the end, sensitivity must improve by a daunting two orders of magnitude to barely reach the standard model BR; even greater advances will be needed to begin the task of unraveling the CP-symmetry physics of $K_L \rightarrow \pi^0 e^+ e^-$.

Bibliography

- [1] C.S. Wu *et al.*, “Experimental Test of Parity Conservation in Beta Decay,” *Phys. Rev.* **105**, 1413 (1957).
- [2] R.L. Garwin, L.M. Lederman, M. Weinrich, “Observations of the Failure of Conservation of Parity and Charge Conjugation in Meson Decays: the Magnetic Moment of the Free Muon,” *Phys. Rev.*, **105**, 1415 (1957).
- [3] J.I. Friedman, V.L. Telegdi, “Nuclear Emulsion Evidence for Parity Non-conservation in the Decay Chain $\pi^+ - \mu^+ - e^+$,” *Phys. Rev.*, **105**, 1681 (1957).
- [4] A.D. Sakharov, “Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe,” *ZhETF Pis'ma* **5**, 32 (1967).
- [5] K. Lande *et al.*, “Observation of Long-Lived Neutral V Particles,” *Phys. Rev.* **103**, 1901 (1956).
- [6] J.H. Christenson, J.W. Cronin, V.L. Fitch, R. Turlay, “Evidence for the 2π Decay of the K_2^0 Meson,” *Phys. Rev. Lett.* **13**, 138 (1964).
- [7] L. Wolfenstein, “Violation of CP Invariance and the Possibility of Very Weak Interactions,” *Phys. Rev. Lett.* **13**, 562 (1964).
- [8] A. Alavi-Harati *et al.* [KTeV Collaboration], “Observation of Direct CP Violation in $K_{S,L} \rightarrow \pi\pi$ Decays,” *Phys. Rev. Lett.* **83**, 22 (1999).
- [9] D.A. Harris *et al.* [E799-i Collaboration], “Limit on the branching ratio of $K_L \rightarrow \pi^0 e^+ e^-$,” *Phys. Rev. Lett.* **71**, 3918 (1993).
- [10] L.M. Sehgal, “CP violation in $K_L \rightarrow \pi^0 e^+ e^-$: Interference of one-photon and two-photon exchange,” *Phys. Rev.* **D38**, 808 (1988).
- [11] P. Heiliger, L.M. Sehgal, “Analysis of the decay $K_L \rightarrow \pi^0 \gamma \gamma$ and expectations for the decays $K_L \rightarrow \pi^0 e^+ e^-$ and $K_L \rightarrow \pi^0 \mu^+ \mu^-$,” *Phys. Rev.* **D47**, 4920 (1993).
- [12] A. Alavi-Harati *et al.* [KTeV Collaboration], “Search for the decay $K_L \rightarrow \pi^0 \mu^+ \mu^-$,” in preparation.
- [13] L.S. Littenburg, “CP-violating decay $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$,” *Phys. Rev.* **D39**, 3322 (1989).

- [14] S. Bosch *et al.*, “Standard model confronting new results for ϵ'/ϵ ,” preprint hep-ph/9904408 (1999).
- [15] J. Adams *et al.* [KTeV Collaboration], “Search for the decay $K_L \rightarrow \pi^0 \nu \bar{\nu}$,” *Phys. Lett.* **B447**, 240 (1999).
- [16] A. Alavi-Harati *et al.* [KTeV Collaboration], “Search for the decay $K_L \rightarrow \pi^0 \nu \bar{\nu}$ using $\pi^0 \rightarrow e^+ e^- \gamma$,” Preprint hep-ex/9907014 (1999).
- [17] A.J. Buras, L. Silvestrini, “Upper bounds on $K \rightarrow \pi \nu \bar{\nu}$ and $K_L \rightarrow \pi^0 e^+ e^-$ from ϵ'/ϵ and $K_L \rightarrow \mu^+ \mu^-$,” *Nucl. Phys.* **546**, 299 (1999).
- [18] G.D. Barr *et al.* [NA31 Collaboration], “Observation Of The Decay $K(L) \rightarrow \pi^0 \gamma \gamma$,” *Phys. Lett.* **B242**, 523 (1990).
- [19] A.G. Cohen, G. Ecker, A. Pich, “Unitarity and $K_L \rightarrow \pi^0 \gamma \gamma$,” *Phys. Lett.* **B304**, 347 (1993).
- [20] J.F. Donoghue, F. Gabbiani, “ $K_L \rightarrow \pi^0 \gamma e^+ e^-$ and its relation to CP and chiral tests,” *Phys. Rev.* **D56**, 1605 (1997).
- [21] E. Cheu, “Status of the direct CP violation measurement from KTeV and other results,” In *Vancouver 1998, High energy physics, vol. 2*, pp. 1001–1005.
- [22] J.F. Donoghue, F. Gabbiani, “Reanalysis of the decay $K_L \rightarrow \pi^0 e^+ e^-$,” *Phys. Rev.* **D51**, 2187 (1995).
- [23] A. Alavi-Harati, *et al.* [KTeV Collaboration], “Measurement of the decay $K_L \rightarrow \pi^0 \gamma \gamma$,” *Phys. Rev. Lett.* **83**, 917 (1999).
- [24] G. D’Ambrosio, G. Ecker, G. Isidori, and J. Portoles, “The decays $K \rightarrow \pi l^+ l^-$ beyond leading order in the chiral expansion,” *Jour. High Ener. Phys.* **9808**, 004 (1998).
- [25] G.D. Barr *et al.* [NA31 Collaboration], “Search for the decay $K(S) \rightarrow \pi^0 e^+ e^-$,” *Phys. Lett.* **B304**, 381 (1993).
- [26] G. Colangelo and G. Isidori, “Supersymmetric contributions to rare kaon decays: beyond the single mass-insertion approximation,” *Jour. High Ener. Phys.* **9809**, 009 (1998).
- [27] A.J. Buras and L. Silvestrini, “Upper bounds on $K \rightarrow \pi \nu \bar{\nu}$ and $K_L \rightarrow \pi^0 e^+ e^-$ from ϵ'/ϵ and $K_L \rightarrow \mu^+ \mu^-$,” *Nucl. Phys.* **B546**, 299 (1999).
- [28] J.M. Flynn and L. Randall, “The CP violating contribution to the decay $K_L \rightarrow \pi^0 e^+ e^-$,” *Nucl. Phys.* **B326**, 31 (1989).
- [29] A. Barker *et al.* [E731 Collaboration], “New limit on $K_L \rightarrow \pi^0 e^+ e^-$,” *Phys. Rev.* **D41**, 3546 (1990).
- [30] K.E. Ohl *et al.* [E845 Collaboration], “Improved experimental limit on $K_L \rightarrow \pi^0 e^+ e^-$,” *Phys. Rev. Lett.*, **64**, 2755 (1990).

- [31] W.M. Morse *et. al.* [E845 Collaboration], “Observation of the decay mode $K_L^0 \rightarrow \gamma\gamma ee$,” *Phys. Rev.* **D45**, 36 (1992).
- [32] T. Nakaya *et. al.* [E799-i Collaboration], “A Measurement of the Branching Ratio of $K_L \rightarrow e^+e^-\gamma\gamma$,” *Phys. Rev. Lett.* **73**, 2169 (1994).
- [33] M.G. Setzu *et. al.* [NA31 Collaboration], “Measurement of the decay $K_L \rightarrow e^+e^-\gamma\gamma$,” *Phys. Lett.* **B420**, 205 (1998).
- [34] H. Blumer, “Measurement of direct CP-violation with the NA48 experiment at the CERN SPS,” In *Vancouver 1998, High energy physics, vol. 2*, pp. 996–1000.
- [35] H.B. Greenlee, “Background to $K_L^0 \rightarrow \pi^0 ee$ from $K_L^0 \rightarrow \gamma\gamma ee$,” *Phys. Rev.* **D42**, 3724 (1990).
- [36] L. Bergström, E. Massó, and P. Singer, “Testing the origin of the $\Delta I = \frac{1}{2}$ rule through K_L Dalitz decays,” *Phys. Lett.* **131B**, 229 (1983).
- [37] C. Caso *et al.* [Particle Data Group], “Review of Particle Physics,” *European Phys. Jour.*, **C3**, 1 (1998).
- [38] V. Fanti *et al.* [NA48 Collaboration], “Measurement of the decay rate and form factor paramter α_{K^*} in the decay $K_L \rightarrow e^+e^-\gamma$,” *Phys. Lett.* **B458**, 553 (1999).
- [39] P. Heiliger, B. McKellar, and L.M. Sehgal, “ K^0 decay into three photons,” *Phys. Lett.* **B327**, 145 (1994).
- [40] Private communication from G. D’Ambrosio and G. Isidori to T. Nakaya, reported in: T. Nakaya [E799-i Collaboration], “Measurement of the Branching Ratio of $K_L \rightarrow e^+e^-\gamma\gamma$,” PhD. Dissertation, Osaka University (1995).
- [41] P. Heiliger and L.M. Sehgal, “Decays $K_{L,S} \rightarrow \pi^0\pi^0\gamma$ and $K_{L,S} \rightarrow \pi^0\pi^0e^+e^-$ as probes of chiral dynamics,” *Phys. Lett.* **B307**, 182 (1993).
- [42] A. Roodman, “The KTeV pure CsI calorimeter,” *Proceedings of the VII International Conference on Calorimetry, Tucson, Arizona*, World Scientific (1998).
- [43] N. Solomey [KTeV Collaboration], “Development and performance of the KTeV transition radiation detector system,” *Nucl. Inst. Meth.*, **A419**, 637 (1998).
- [44] C. Bown, E. Cheu, J. Dusatko, H. Sanders and M. Zeleznik, “The KTeV hardware cluster counter,” *Nucl. Inst. Meth.* **A369**, 248 (1996).
- [45] J. LaDue, private communication (1999).
- [46] G. Graham, “First Observation of the rare decay $K_L \rightarrow \pi^0\gamma e^+e^-$,” PhD. Thesis, The University of Chicago (1999).
- [47] J.F. Donoghue, B.R. Holstein, and G. Valencia, “ $K_L \rightarrow \pi^0 e^+e^-$ as a probe of CP violation,” *Phys. Rev.* **D35**, 2769 (1987).
- [48] A.F. Bielajew, “Improved angular sampling for pair production in the EGS4 code system,” PIRS-0287 (1991).

- [49] H.W. Koch and J.W. Motz, “Bremsstrahlung cross-section formulas and related data,” *Rev. Mod. Phys.* **31**, 920 (1959).
- [50] E.D. Zimmerman, “Use of the $K_L \rightarrow 3\pi^0$ decay as a tagged photon source to measure material thickness in a neutral kaon beam,” *Nucl. Inst. Meth.* **A426**, 229 (1999).
- [51] P. Toale, private communication (1999).
- [52] A. Alavi-Harati *et al.* [KTeV Collaboration], “Measurement of the branching ratio of $\pi^0 \rightarrow e^+e^-$ using $K_L \rightarrow 3\pi^0$ decays in flight,” *Phys. Rev. Lett.* **83**, 922 (1999).
- [53] G.J. Feldman and R.D. Cousins, “Unified approach to the classical statistical analysis of small signals,” *Phys. Rev.* **D57**, 3874 (1998).

Appendix A

Z Slope and DC Inefficiency

The appearance of a bias in the reconstructed vertex Z of simulated $K \rightarrow \pi^0 \pi_{Dalitz}^0$ events (Figure 8.4) is a cause for concern. It suggests that some fundamental aspect of tracking may not be understood. Figure A.1 shows this “Z slope” — the result of fitting a 1st order polynomial to the data:MC ratio in vertex Z .

This appendix proposes that the Z slope is the result of an unsimulated inefficiency in regions of DC1 near the neutral beam. There are proportionally fewer hits in the beam regions of DC1 in MC than in data. One contributing factor to this effect may be that the TRDs are not simulated at all, while the TRD chambers are deadened in their beam regions and data events must have one TRD hit on each track. Another factor may be radiation damage of the DCs by the beams. Beam region inefficiency is of more concern than Z slope itself, because it could bias efficiency measurements in different decay modes that illuminate the beam regions differently. I should note that the statistics of $K \rightarrow \pi^0 \pi_{Dalitz}^0$ decays limit the ability of this analysis to understand the problem. More common charged modes might be better at understanding the root causes of the problem. However, $K \rightarrow \pi^0 \pi_{Dalitz}^0$ is used here because of its similarity to the signal modes.

The DC inefficiency is found by measuring the number of data and MC events with zero or one tracks in a beam region of DC1. The beam regions in X/Y are found




Figure A.1: Ratios of data to MC vs. vertex Z , in meters, for $K \rightarrow \pi^0 \pi^0_{Dalitz}$ decays. Winter is on top, Summer on the bottom, charged-track Z on the left, and Z_{π^0} on the right. All cuts are made.

by drawing lines from the target, through the limiting edges of the defining collimator, to the downstream X/Y plane of DC1. This means that the beam regions are bigger for Summer than for Winter. The X regions for this measurement are vertical strips covering the square beam regions extending to the edges of DC1, The Y region is a horizontal strip. A track is in the X/Y beam region if any hit that it uses in DC1 X/Y is on a wire inside or adjacent to a beam region.

Table A.1 lists the number of events by the number of tracks in the beam regions. Events with 2 tracks in a beam region are not listed. The ρ_i column lists the double ratio:

$$\rho_i = \frac{\left(\frac{\text{data with 1 in } i \text{ beam}}{\text{data with 0 in } i \text{ beam}} \right)}{\left(\frac{\text{MC with 1 in } i \text{ beam}}{\text{MC with 0 in } i \text{ beam}} \right)},$$

where i is either X or Y . This ρ_i is the depletion in data (or excess in MC) of tracks in the beam.

Period	Weight/Cut	i	Data		MC		ρ_i
			ntrk=0	ntrk=1	ntrk=0	ntrk=1	
Winter	slope=0	x	7181	7962	29266	34125	0.951 ± 0.017
		y	11194	5457	45681	23759	0.937 ± 0.017
	slope=-0.0019	x	7181	7962	30333.0	35195.1	$0.956 \pm 0.017^*$
		y	11194	5457	47354.8	24370.1	$0.947 \pm 0.017^*$
	$\sigma_Z < 1$ m	x	4098	4892	17122	21563	0.948 ± 0.022
		y	6318	3740	26467	16466	0.952 ± 0.022
Summer	slope=0	x	5080	5797	21628	24892	0.992 ± 0.021
		y	8039	3830	33787	17012	0.940 ± 0.020
	slope=-0.0020	x	5080	5797	22435.2	25695.6	$0.996 \pm 0.021^*$
		y	8039	3830	35060.7	17461.7	$0.951 \pm 0.020^*$
	$\sigma_Z < 1$ m	x	2925	3595	12746	15942	0.983 ± 0.027
		y	4595	2626	19775	11841	0.954 ± 0.026

Table A.1: ρ_X , ρ_Y , and related numbers for Summer and Winter, with various cuts and weights. Stars (*) indicate ρ s actually used. $K \rightarrow \pi^0 \pi^0_{Dalitz}$ decays with all cuts are used.

The table also lists ρ_X and ρ_Y calculated with different weighting and cuts. This is done because the Z slope may bias measurement of ρ . The farther downstream decays occur, the more likely they are to put a track in a beam region. To compare data and

MC with the same Z distribution, I weight the MC events using their vertex Z s. The weight is equal to Z times the slope in Figure A.1, plus a constant that makes the weight equal to one at $Z = 143$ m, to flatten out the Z slope. For comparison, the ρ s without weighting are listed. As a check, the ρ s with an alternate method of flattening Z slope are listed. Requiring the vertex uncertainty (VTXSIGZ) reported by the reconstruction algorithm to be less than 1 meter removes the Z slope, although at a heavy cost in $K \rightarrow \pi^0 \pi^0_{Dalitz}$ statistics. The ρ s obtained in this way are consistent with the Z slope weighted ρ s.

The ρ s are used to get the systematic error resulting from not simulating the inefficiency. Each MC event used in calculating acceptance is given a weight

$$w = (\rho_X)^{N_x} \times (\rho_Y)^{N_y},$$

where N_x is the number of tracks in the X beam regions and N_y is the number of tracks in the Y beam region. The change in answer using these weighted MC events is taken as a systematic uncertainty.

Figure A.2 shows how the Z slopes change when the ρ weighting is used. For Winter, the slopes are within 1σ of zero. For Summer, significant slopes still remain. However, ρ_X for Summer is conspicuously higher (0.996) than the other ρ s (~ 0.95). It may be that statistical fluctuations pushed Summer's ρ_X up, but the underlying inefficiency remained.




Figure A.2: Ratios of data to MC vs. vertex Z , with ρ weighting on MC. Winter is on top, Summer on the bottom, charged-track Z on the left, and Z_{π^0} on the right. All cuts are made, and $K \rightarrow \pi^0 \pi^0_{\text{Dalitz}}$ decays are used.