

THE UNIVERSITY OF CHICAGO

SEARCH FOR THE RARE DECAY $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$

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

DEPARTMENT OF PHYSICS

BY
GABRIEL N. PERDUE

CHICAGO, ILLINOIS

JUNE 2008

Copyright © 2008 by Gabriel N. Perdue

All rights reserved

Dedicated to everyone who refuses to stop searching for the “impossible” to find.

ABSTRACT

The E391a collaboration searched for the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ decay at the KEK 12 GeV proton synchrotron in Tsukuba, Japan. We performed a blind analysis on a sample of $(5.13 \pm 0.4) \times 10^9$ Kaon decays. In a blind analysis, the signal region is kept hidden until all signal definition criteria are fixed in order to minimize the experimenter's bias in the process of cut selection. We established a background event estimate of 0.45 ± 0.13 events through the combined use of simulation and data. After finalizing our signal criteria and background estimation, we opened our signal region and observed zero events. Therefore, we have set a new upper limit on the branching ratio for the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ decay at 6.7×10^{-8} at the 90% confidence level using Poisson statistics.

ACKNOWLEDGMENTS

First and foremost I must thank my parents, James and Gayle Perdue, for their love and support. Everything that I am today and could be tomorrow is because of them. They gave me a thirst for knowledge and never once let me get away with thinking I couldn't accomplish something if I really wanted to do it. I can only wish every child were so lucky as to have the kind of love and support I grew up with.

Second, I would like to extend my gratitude to my grandmothers, Louise Perdue and Marian Greenwood for all the love and encouragement they have given me over the course of my life, and to my Uncle Mark Perdue for always asking me so many tough questions. I would also like to thank the rest of my family for all the help they have offered me over the years.

Next I must thank Yau Wah for teaching me how to think critically about nature and how to divine her secrets. $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ is a difficult measurement. It is not so much a search for a needle in a haystack so much as it is search for a specific grain of sand in the Sahara. I have been most fortunate to have as good a relationship with my thesis advisor as I have. When people ask me for advice about choosing a graduate school, my answer is always that it doesn't matter where you go to school; what matters is your advisor. I have been luckier in that regard than I had any business being. Regardless of what happens in my future, it has been a true honor to be able work with Dr. Wah here at Chicago.

I would like to thank Jon Nix, my office mate for many years. Jon was my best reference for just about any subject and he is definitely in the starting line-up of my trivial pursuit team should I ever find myself playing for a serious stake. I have profited immeasurably from our discussions.

It is critical I thank Toshi Sumida for his many contributions to this work. The

$K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ search in Run-II was conducted as a team effort between the two of us. Toshi's willingness to work at odd hours of the night in Japan so that we could communicate in real time via instant messenger is a sacrifice I will never forget. He was also a teacher for me about Japan and the Japanese. My trips to that country were made infinitely richer through his perspective and willingness to share.

I would like to extend my thanks to Dr. Hiroaki Watanabe for his guidance in the design and construction of the Beam-Anti. His experience and depth of knowledge, in addition to his creative spirit, were critical to the success of that project. I should also thank him for everything he taught me about how to work and how to push solutions through impossible problems with sheer force of will.

I must thank Dr. Takao Inagaki for teaching me so much about about the large scale problems behind doing an experiment and for always taking the time to see both the forest and the trees with equal clarity. No matter how simple or complex the plots I showed him were, he always asked a good question.

I would like to thank Dr. Hideki Okuno for his masterful design and construction efforts in our experiment. Everything would literally have been impossible without him.

There are many students from the E391a experiment that should be acknowledged. I must thank Ken Sakashita for his tireless dedication to our early software and analysis efforts. Jiasen Ma was a good office mate who always kept me honest by refusing to accept half-answers. This document is undoubtedly better thanks to him. I would also like to thank Hiroaki Ishii, Hyosang Lee, Misha Doroshenko, and Sergei Podolski for their insight and enthusiasm.

An experiment like this one relies critically on good people in the institutions that run it. From Chicago, I am indebted to Elizabeth Pod for her mechanical and design expertise, to Nobuko McNeil for showing me how to smoothly operate the

inner workings of the University, and Helmut Krebs for teaching me how to shape metal with my bare hands (and a lathe). I also should thank Aspasia Sotir-Plutus and Veronica McClain for making day-to-day life in HEP run so smoothly, and Marty Dippel and Mary Heintz for being the best kind of system administrators - invisible ones (think about it)! From JINR I need to thank Anatoli Lednev for his great help with the mathematics in the five cluster kinematic fit routines. Finally, from KEK, I must thank Gei Youb Lim for not only helping me navigate the paperwork in Japan, but also for never missing a detail on any plot I consulted with him about. Of course, I am also indebted to crew of the KEK 12-GeV proton synchrotron and to all the staff and administrators at the lab.

I would like to thank Dan Reich, my freshman physics instructor from my days at Johns Hopkins. I entered JHU as a pre-med majoring in mechanical engineering and enrolled in the physics course for majors only because the course was offered at 11 am, while engineers went to class at 8. Dr. Reich opened my eyes though to the world of possibilities that existed in physics research. By the end of his class, I had changed majors and my career track for good. As a Chicago alumnus, he also provided a great deal of encouragement in my decisions to apply to and attend Chicago. I should also thank Adam Falk and Jon Bagger for being truly stand-out instructors in my undergraduate education.

I would also like to thank Gerben Wierda, who maintains the TeXShop version of L^AT_EX for the Macintosh [1] [2]. The open source community is an important contributor to the success of science today. It is impossible to thank everyone involved, so I will simply ask Gerben to accept my thanks on behalf of the community at large.

Finally, I suppose I must thank Bungie Studios and Microsoft for the game *Halo*, and the role it played in keeping me sane while at Chicago. But there is no question I would have made it out of school quicker without it. ;)

TABLE OF CONTENTS

ABSTRACT	iv
ACKNOWLEDGMENTS	v
LIST OF FIGURES	xi
LIST OF TABLES	xxiii
1 INTRODUCTION	1
1.1 Kaon Phenomenology	4
1.1.1 Indirect and Direct CP Violation in the Kaon System	8
1.2 The Standard Model and CP Violation.	10
1.2.1 The CKM Matrix	12
1.2.2 The Unitarity Triangle	14
1.3 $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ and the CKM Matrix	16
1.3.1 Theoretical Framework	16
1.3.2 SM Branching Ratio	19
1.4 Beyond the Standard Model	20
1.5 Guide to this Dissertation	22
2 METHOD OF THE EXPERIMENT	24
2.1 Critical Summary	24
2.2 Philosophy of the Measurement	24
2.3 KEK	28
2.3.1 Beam Line	29
2.3.2 Operation of the Synchrotron	33
2.3.3 Running Conditions	33
2.4 Brief Survey of the Detector	34
2.5 DAQ Electronics	36
2.6 First Level Event Selection - The Trigger	37
2.6.1 Physics Trigger	37
2.6.2 Other Triggers	38
2.7 Event Reconstruction	39
2.7.1 Cluster Finding	41
2.7.2 Energy and Position Corrections	42
2.7.3 Pion and Kaon Reconstruction	45

3	DETECTOR AND CALIBRATION	48
3.1	Critical Summary	48
3.2	The Full Detector	48
3.3	CsI	49
3.3.1	CsI Calibration	53
3.3.2	Cosmic Calibration	57
3.3.3	$K_L^0 \rightarrow 3\pi^0$ Calibration	57
3.3.4	Al-Target Run	59
3.4	Main Barrel	62
3.4.1	Main Barrel Calibration	63
3.4.2	Barrel Charged Veto	65
3.5	Front Barrel	65
3.6	Charged Veto	66
3.7	The Back-Anti (BA)	66
3.7.1	Beam-Hole Charged Veto	71
3.8	Collar Counters	72
3.8.1	CC00	72
3.8.2	CC02	73
3.8.3	CC03	73
3.8.4	CC04 and CC05	74
3.8.5	CC06 and CC07	75
3.9	The Vacuum System	76
4	MONTE CARLO	78
4.1	Critical Summary	78
4.2	Introduction	78
4.2.1	Data and MC Processing	79
4.3	Accidental Overlay	80
4.3.1	Accidental Trigger Timing	81
4.3.2	Accidental Loss Estimation - Pure MC and Overlay	82
4.4	MC Tuning	84
4.5	Event Weights	85
4.6	Decay Probability	88
4.7	Acceptance for the Calibration Modes	89
4.8	The Invariant Mass Spectrum and Energy Resolution	91
4.9	The Four Cluster Invariant Mass Spectrum	94
4.10	Veto Distributions	97
4.10.1	Veto Timing	98
4.10.2	CC02	99
4.10.3	Main Barrel	99
4.10.4	Beam-Anti	101
4.10.5	CsI	103

4.11	Kinematic Distributions	107
4.11.1	Decay z Vertex	107
4.11.2	Total Momentum	112
4.11.3	Transverse Momentum and Beam Shape	112
4.11.4	Beam Angles	116
4.11.5	Pairing χ^2	116
4.12	Flux and Systematic Errors	118
4.12.1	Systematic Error	118
4.12.2	The CsI Veto and the Systematic Error	120
4.12.3	Flux Values	128
5	PHOTON VETO INEFFICIENCY	133
5.1	Critical Summary	133
5.2	Photon Veto Inefficiency	133
5.2.1	Five Cluster Events	135
5.2.2	The 1-C Fit	140
5.2.3	Five Cluster Event Reconstruction with the E391a Calorimeter	143
5.2.4	Opportunities at Future Experiments	149
6	BACKGROUNDS AND EVENT SELECTION	152
6.1	Critical Summary	152
6.2	Event Selection	152
6.2.1	Kinematic Cuts	153
6.2.2	Veto Cuts	159
6.3	Neutron Backgrounds	159
6.3.1	Collar Counter 02 Events	160
6.3.2	Charged Veto Events	164
6.4	Kaon Backgrounds	169
6.4.1	$K_L^0 \rightarrow \gamma\gamma$	169
6.4.2	$K_L^0 \rightarrow \pi^0\pi^0$	171
6.4.3	$K_L^0 \rightarrow \pi l\nu$ and $K_L^0 \rightarrow \pi^+\pi^-\pi^0$	180
6.5	Total Background Check	187
7	CONCLUSIONS	189
7.1	Critical Summary	189
7.2	$K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ Event Generator	189
7.3	Acceptance and Sensitivity	190
7.4	Final Results	194
7.5	How an E391a-style Experiment Can Discover $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$	197
	REFERENCES	199

LIST OF FIGURES

1.1	Diagrams mediating K^0 and \bar{K}^0 mixing. The charge of the W bosons is suppressed.	5
1.2	Direct and indirect CP violation in the decay of the $K_L^0 \rightarrow \pi\pi$	9
1.3	The Unitarity Triangle with the impacts of $K \rightarrow \pi\nu\bar{\nu}$ on the parameters visualized.	15
1.4	The penguin and box Feynman diagrams for $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$. The up quark contributions are only required for Glashow-Iliopoulos-Maiani suppression and computationally the process is dominated by the top and charm quarks. Figure adapted from [9].	18
2.1	The general concept of the E391 experiment in diagram form. We search for $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ by identifying a single $\pi^0 \rightarrow \gamma\gamma$ in our fiducial volume with transverse momentum consistent with the decay and inconsistent with backgrounds. We further require all veto detectors to be “quiet.” We do not detect the neutrinos from the decay, and the fiducial decay region is kept at a low vacuum level to minimize interactions with air (or other gas) molecules. Signal is defined by its location in (z, P_T) parameter space with kinematic cuts and photon vetoes imposed. Events in the space are considered signal candidates and events outside are not, making E391a a counting experiment. . .	27
2.2	The lay out of the East Hall at KEK. The E391a experimental area is the shaded region in the center of the figure and the PI0 beam-test area is below and to the left on the diagram.	29
2.3	The K^0 beam-line peak-to-halo comparisons for different beam products, reproduced from [36].	30
2.4	The K^0 beam-line for E391a, reproduced from [36]. The top of the figure shows how the geometrical sizes and positions of the components (note the different vertical and horizontal scales), and the bottom of the figure details the collimation scheme.	31
2.5	The E391a Detector shown in plan view with a numeric length scale. Note that our coordinate system identifies the origin at the beginning of the Front Barrel. The z -axis points downstream and the y -axis points vertically upwards (the coordinate system is “right-handed.”) .	34
2.6	The full E391a Detector, shown with human figures for scale. Not shown are the Collimator Vetoes installed inside the beam-line itself.	35
2.7	Hardware cluster assignments for the physics trigger.	37
2.8	A flow-chart representation of the clustering algorithm. A cluster must contain at least two crystals. Note that clusters do not grow diagonally except through neighbors of neighbors paths.	43

2.9	A diagrammatic representation of the cluster correction table application. Here, “COG” is an acronym for center-of-gravity (energy weighted mean), and “Cor” is an abbreviation for <i>corrected</i>	44
2.10	Before (blue) and after (red) the energy and position correction routines as applied to (a) the photon hit position (in x), (b) the photon energy, and (c) the pion z -vertex. “True - Rec.” means the known true MC value minus the reconstructed value. The associated means and widths (RMS) for the corrected clusters are displayed on the left side of each sub-figure, while the original center-of-gravity calculations are provided on the right side of each sub-figure.	45
2.11	A diagrammatic representation of π^0 reconstruction according to Equation 2.1. An additional relation required for reconstruction is the law of cosines, here written as $r_{12}^2 = d_1^2 + d_2^2 - 2d_1d_2 \times \cos \theta$	46
2.12	A diagrammatic representation of $K_L^0 \rightarrow 3\pi^0$ reconstruction, where each pion is first reconstructed according to Figure 2.11.	47
3.1	CsI pedestal stability before and after an Earthquake on February 16, 2005, shown as pedestal count (channel number) for CsI block 62 versus run number. (The earthquake on February 23, corresponding to about run 4120 above did not have as dramatic an impact on the electronics.) The low range (0.05 pc/Channel) is shown in the top half and the high range (04 pc/Channel) is shown in the lower half. The error bars displayed are five times the gaussian width of the fit (for the low range) or five times the RMS (for the high range). The fit to the pedestal evolution is linear: Pedestal = $P_0 + P_1 \times$ Run Number.	50
3.2	The components of the outer edge of the CsI array. The Sandwich Counter modules are the triangular blocks. They were built from lead and plastic scintillator plates, oriented parallel to the beam. There were a total of 24 such counters, grouped into eight groups of three for readout purposes. The deformed CsI blocks were each read-out individually and used in photon clustering and as veto counters in the exact same way as the rest of the CsI.	51
3.3	The CsI-PMT attachment scheme employed for the KEK CsI (a) and the KTeV CsI (b).	52
3.4	The sandwich counter module design. See Figure 3.2 for placement of the different types in the array.	53

3.5	CsI Xenon Monitor propagation. The drift in the gain is likely due to PMT fatigue and Xe lamp dimming. Uncertainty about this drift helped to motivate the $K_L^0 \rightarrow 3\pi^0$ kinematic fit calibration described below. “Day-by-day” cosmic ray calibration employed all the cosmic tracks found on the given calendar day during Run II with both the dedicated cosmic trigger (using the Main Barrel) and the X_{cls} , or X-cluster trigger (usually four or more hardware clusters).	54
3.6	CsI On/Off Spill Xenon. The χ^2 values from the individual channel fits are used as errors, but are not shown in the plot. As the boxed parameter indicates, there was very little in the way of beam-loading for the CsI, and the gain values computed using off-spill cosmic rays were compatible with gain values computed from on-spill $K_L^0 \rightarrow 3\pi^0$ decays.	55
3.7	The CsI energy resolution as measured by a 5x5 block array. The fit is $\sigma = P_1 \times \sqrt{E} + P_2 \times E$	56
3.8	A cosmic ray track in the CsI array.	58
3.9	CsI gain stability as computed with cosmic rays.	58
3.10	A clean $K_L^0 \rightarrow 3\pi^0$ event from Run II data that passed all cuts. The CsI array is in the upper right hand corner. On the left-top are the upstream channels of the Main Barrel (inner and outer), the Front Barrel (inner and outer layers, with CC02 inside the inner layers), and on the left-bottom are the Main Barrel downstream channels (inner and outer layers, with the Barrel Charged veto as well). Note that there is some soft activity in the BA (bottom-right hand corner) scintillator layers (separated by diagonal lines), but no activity in the quartz layers (separated by horizontal lines) - a likely candidate for a low-energy neutron interaction either in-time or slightly early. The channel marked “APC” was reserved for a detector element not employed in Run II. .	60
3.11	An individual cluster from the event shown in Figure 3.10 with $E = 395$ MeV and $(x, y) = (43.5, -20.75)$. For reference, the center of the most energetic block has (x, y) coordinates of $(42.0, -21.0)$, with x increasing towards the right hand side of the figure and y increasing with the vertical. Recall that each of the “KEK CsI” blocks has face dimensions of 7×7 cm ²	61
3.12	The Main Barrel and vacuum tank.	64
3.13	A single Main Barrel module.	64
3.14	The Outer Charged Veto (OCV).	67
3.15	The Beam-Hole Charged Veto (BHCV, left) and Beam-Anti (BA, right). The channel numbering scheme is displayed above, but the figure is not precisely to scale. For the BA, note that the scintillator layers are separated by vertical lines and the quartz layers are separated into vertically stacked blocks to provide transverse segmentation.	68

3.16	The BA scintillator energy resolution as measured at the T552 beam test. The fitting function is $\sigma/\mu = P_1/\sqrt{E} + P_2$	69
3.17	The BA scintillator energy linearity as measured at the T552 beam test. The fitting function is $E_{Dep} = P_1 + P_2 \times E_{Beam}$	70
3.18	CC02, shown here looking downstream (the beam would pass through the center gap region).	74
3.19	CC03 as seen looking down the beam axis.	75
3.20	CC06/CC07 - both detectors were constructed identically out of lead-glass blocks.	76
3.21	The vacuum system.	77
4.1	BA multi-hit timing comparison for $K_L^0 \rightarrow 3\pi^0$ (left) and accidentals (right) data.	83
4.2	Shown above is the reconstructed radial spectrum of the Kaon at C6 at the end of the Test 10 tuning process before (left) and after (right) a weight factor was applied for six cluster ($K_L^0 \rightarrow 3\pi^0$) data and MC. All kinematic cuts (except the cut on radial position, see Table 4.2) are applied. Here and throughout for overlay plots, data are shown in the upper half as black dots with error bars and the MC events are shown in solid red. The lower half shows the bin-by-bin ratio of data and MC, normalized by total event count. Here, the fit curve is $y = A_0 + A_1 \times r$ and is plotted in red. Errors on both distributions are counting errors (\sqrt{N}). The $\chi^2/d.o.f.$ variable in the upper left-hand corner is a measure of bin-by-bin discrepancies, here with 100 bins (one degree of freedom is used for normalization). The scale variable in the upper right-hand corner is the number used to scale the MC so that the integrated event number for both distributions agree.	87
4.3	Shown above is the collection of event weights for six cluster events (left) and the radial generation position at C6 (right) for a special version of the MC with interactions turned off (pure decay kinematics). On the right, the re-weighted distribution is in red and the un-weighted is in black (with error bars = \sqrt{N}). Overlay plot conventions are the same as in Figure 4.2 (and throughout this document). The event at very high r is due to a non-zero Kaon occupancy in the beam halo.	88
4.4	The six cluster invariant mass spectrum with all cuts applied. On the left we have the $K_L^0 \rightarrow 3\pi^0$ mass plot for data (top, $\sim 1/4$ of the data set) and for MC (bottom). Data are represented by black dots and the MC by the solid red line in the upper half of the right-hand figure. The lower half shows the bin-by-bin ratio with counting errors in dashed black and the fit line in solid red (again, normalized by total event count). Note the data's peak is slightly wider (see Table 4.7).	93

4.5	Reconstructed photon energy (with energy leakage corrections applied) minus the sum of the constituent crystal energies for the native GEANT3 simulation and the crystal by crystal smeared simulation.	94
4.6	On the left is the four cluster invariant mass overlay with analysis cuts imposed, normalized by signal events. The sum MC is in red and the data points are black with error bars. On the right is the four cluster invariant mass final plot, again signal normalized.	96
4.7	On the left is the $K_L^0 \rightarrow 3\pi^0$ four cluster z -vertex versus invariant mass plot, with the same distribution for data on the right. The high-mass contribution from the $K_L^0 \rightarrow 3\pi^0$ MC appears to be a very small fluctuation in z away from participating in the mass plot shown in Figure 4.6.	97
4.8	CsI time versus energy plots for two cluster data and $K_L^0 \rightarrow \gamma\gamma$ and (two cluster) $K_L^0 \rightarrow \pi^0\pi^0$ MC. The horizontal bars depict the time windows used for veto.	100
4.9	CC02 timing plots for four cluster data and $K_L^0 \rightarrow \pi^0\pi^0$ and (four cluster) $K_L^0 \rightarrow 3\pi^0$ MC. In the MC plot, the early time hits from $K_L^0 \rightarrow 3\pi^0$ are clearly visible while in the data no such peak separation exists.	100
4.10	Energy overlays for CC02 using four cluster data and MC. No timing cut is imposed on CC02, but all other four cluster analysis cuts are applied.	101
4.11	Energy overlays for the Main Barrel using four cluster data and MC. The energy value reported is equal to $\sqrt{E_{up} \times E_{down}}$, where <i>up</i> and <i>down</i> refer to the up and downstream ends of the Barrel. The energy produced removes the dependency on the attenuation length (the accompanying factor of e is swallowed by the calibration database). Obviously, if either the up or downstream end has zero energy, the product is zero. To keep positive energy values, we require the channels in the product to have energy deposition over the TDC threshold. The mismatch below 0.001 GeV is due to a TDC threshold of ~ 0.001 GeV in data, and ~ 0 GeV in the MC. The difference above 0.006 GeV is due to a tighter online veto in the data than in the simulation (where the online veto was kept deliberately loose at the generation and skim stages to avoid edge bias).	102
4.12	Energy overlays for the BA Scintillator channels using four cluster data and MC. All the channels are collected together in this plot.	103
4.13	“Energy” overlays for the BA Quartz channels using four cluster data and MC. All the channels are collected together in this plot.	104

4.14	Shown here is a close-up of the CsI array from Figure 3.10. The “single-hit” crystals are outlined in thick black. The featured crystal close to the beam-hole is likely due to a beam-associated halo particle. The featured crystal near the bottom of the figure may be due to an accidental hit as well, but is more likely actually associated with the cluster next to it.	105
4.15	Shown here is the relevant two dimensional plot for the CsI veto, with $K_L^0 \rightarrow \gamma\gamma + K_L^0 \rightarrow \pi^0\pi^0$ two cluster MC on the left and data on the right. The distributions are signal normalized, and the apparent deficit at high distance to cluster and energy in the data is a function of the z-scale and not indicative of a mismatch in that region.	106
4.16	Here is a slice formed from the Figure 4.15 on the left made by integrating the distribution over minimum distance to cluster from 17 to 25 cm. On the left we show all-time hits (no timing window cut is imposed) and on the right only in-time hits . We shall have significantly more to say about the slope in these distributions.	106
4.17	The four-cluster decay- z spectrum normalized by total events (left) and signal events (right). All analysis cuts are applied except for the fiducial (z -vertex) cut.	108
4.18	The six-cluster decay- z spectrum normalized by total events. All analysis cuts are applied except for the fiducial (z -vertex) cut to the figure on the left. The CsI veto is additionally removed from the figure on the right.	109
4.19	On the left is the overlay comparison for photon pass-through points in the x -direction at the end of CC02. On the right is the six-cluster decay- z spectrum with an additional cut imposed on this pass-through region.	109
4.20	If the CC02 aperture is not defined correctly, upstream events are differentially removed.	110
4.21	The geometry of the Outer Charged Veto in the MC, looking down the beam axis.	111
4.22	The $K_L^0 \rightarrow \gamma\gamma$ decay- z spectrum with analysis cuts imposed. The figure on the left is normalized by total events and the figure on the right is normalized using the signal region only. Note that the $K_L^0 \rightarrow \pi^0\pi^0$ contribution appears only at the level of about 10^{-4} relative to $K_L^0 \rightarrow \gamma\gamma$	111
4.23	The Kaon total momentum spectrum derived from $K_L^0 \rightarrow \pi^0\pi^0\pi^0$. All analysis cuts are imposed for the figure on the left while the CsI veto is removed from the figure on the right.	113
4.24	The Kaon total momentum spectrum derived from $K_L^0 \rightarrow \pi^0\pi^0$. All analysis cuts are imposed.	113

4.25	Shown above on the left is the four cluster transverse momentum, <i>normalized by signal</i> . The figure on the right is the six cluster transverse momentum. All analysis cuts are applied except for P_T	114
4.26	Shown above on the left is the four cluster Kaon radial position at C6, <i>normalized by signal</i> . The figure on the right is the six cluster radial position at C6. All analysis cuts are applied except for the radial cut.	115
4.27	The reconstructed photon x and y positions for $K_L^0 \rightarrow 3\pi^0$ decays. The “jagged” nature of the distributions is due to a tendency in the clustering algorithm to pull reconstructed hit positions towards the center of crystals.	115
4.28	The polar and azimuthal beam angles at C6 with analysis cuts imposed for $K_L^0 \rightarrow 3\pi^0$	117
4.29	The polar and azimuthal beam angles at C6 with analysis cuts imposed for $K_L^0 \rightarrow \pi^0\pi^0$	117
4.30	The lowest (best) χ_z^2 value (left) and the second-lowest minus the lowest χ_z^2 value (right) with all analysis cuts applied but those on the χ_z^2 values.	118
4.31	The lowest (best) χ_z^2 value with all analysis cuts applied but those on the χ_z^2 values for $K_L^0 \rightarrow \pi^0\pi^0$	119
4.32	The exclusive acceptance fractional differences, (Data - MC)/Data, for the $K_L^0 \rightarrow \gamma\gamma$ Add-BG MC. A <i>positive</i> value indicates the acceptance <i>loss</i> is <i>higher</i> in the MC than it is in data. The acceptance-weighted sum in quadrature is 7.6%.	120
4.33	The exclusive acceptance fractional differences, (Data - MC)/Data, for the $K_L^0 \rightarrow \gamma\gamma$ Pure MC. A <i>positive</i> value indicates the acceptance <i>loss</i> is <i>higher</i> in the MC than it is in data. Note that these values do not include the effects of accidental activity. The acceptance-weighted sum in quadrature is 6.8%.	121
4.34	The exclusive acceptance fractional differences, (Data - MC)/Data, for the $K_L^0 \rightarrow \pi^0\pi^0$ Add-BG MC. A <i>positive</i> value indicates the acceptance <i>loss</i> is <i>higher</i> in the MC than it is in data. The acceptance-weighted sum in quadrature is 7.8%.	121
4.35	The exclusive acceptance fractional differences, (Data - MC)/Data, for the Pure MC. A <i>positive</i> value indicates the acceptance <i>loss</i> is <i>higher</i> in the MC than it is in data. Note that these values do not include the effects of accidental activity. The acceptance-weighted sum in quadrature is 7.3%.	122
4.36	Shown above is the exclusive acceptance fractional differences, (Data - MC)/Data, for the Add-BG MC. A <i>positive</i> value indicates the acceptance <i>loss</i> is <i>higher</i> in the MC than it is in data. The acceptance-weighted sum in quadrature is 7.7%.	122

4.37	Shown above is the exclusive acceptance fractional differences, (Data - MC)/Data, for the Pure MC. A <i>positive</i> value indicates the acceptance <i>loss</i> is <i>higher</i> in the MC than it is in data. Note that these values do not include the effects of accidental activity. The acceptance-weighted sum in quadrature is 6.9%.	123
4.38	The “bugged” MC geometry employed in Test 10. CC03 is highlighted with an arrow.	124
4.39	The $ P $ and z -decay spectra with the momentum restricted such that $2 \text{ GeV}/c < P < 2.8 \text{ GeV}/c$. The positive slope for the z -decay spectrum overlay supports a low energy scale in data.	128
5.1	A candidate five cluster event for kinematic fitting.	137
5.2	Angle cosine distributions for the two solutions to the Kaon energy (recall that Equation 5.15 is quadratic). In these plots, the azimuthal ϕ angle has been integrated out.	142
5.3	Reconstructed (fitted) energy versus true MC energy for the missing photon with no event selection cuts for perfect and actual experiment quality cluster information resolution. No set of selection cuts dramatically improves the energy reconstruction accuracy of the routine for our level of position and energy resolution (see Table 5.1). This means that with E391a data, the best we can do is fit for the reconstructed angle of the missing photon - the energy resolution of that photon is too poor to allow for a small energy bin analysis.	144
5.4	1-C Fit χ^2 for the “minus” solution (see Equation 5.15) with a cut on the plus/minus solution flag imposed (the discriminant for only either the plus or minus solution may be real valued).	145
5.5	The reconstructed versus true hit positions in the Main Barrel and azimuthal angles for the missing photons in the full Run II MC with all analysis cuts imposed. The MC sample used here represents about 10% of the real number of $K_L^0 \rightarrow 3\pi^0$ events in Run II.	146
5.6	The reconstructed mass using the special five cluster fusion finder. On the left we compare $K_L^0 \rightarrow 3\pi^0$ MC (solid) and data (dashed). On the right we compare data (top) to $K \rightarrow \pi^0\pi^0\gamma$ MC (bottom) to show how the cut operates on true five photon final state events. Those events are pushed to high mass while true fusion events are reconstructed at the Kaon mass.	147
5.7	The observed energy spectra in the MB for the five cluster MC (left) and data (right). Note that the low energy spectra differs due to the behavior of the TDC threshold: $\sim 1 \text{ MeV}$ in data and 0 MeV in the MC.	148

- 5.8 Here we compare the performance of the 1-C Fit (minus solution) for a hypothetical higher resolution experiment (here, employing the KTeV CsI array) and E391a in radial bins. Each radial bin represents evenly distributed randomly generated $K_L^0 \rightarrow 3\pi^0$ decays inside rings defined by the upper and lower bounds of the bin. So, the first bin is a circle with a radius of 1 cm. The second bin is a ring with an outer radius of 2 cm and an inner radius of 1 cm, etc. “Pass Percentage” is the number of events passing the C-Fit cuts divided by the total number of events (photon veto cuts are not imposed). “Angle Percentage” is the number of events reconstructed with angle cosines within 2% of the correct value (i.e., greater than 0.96) divided by the total number of events. “Tolerance Percentage” is the number of events that were reconstructed with angle cosines correct to within 2% and energy reconstructed correctly to within 10 MeV divided by the total number of events. 151
- 6.1 Shown here is the γ -RMS comparison for $K_L^0 \rightarrow \gamma\gamma$ MC and two cluster data (left) and the reconstructed angle versus the angle computed from the vertex to the cluster for even-pairing $K_L^0 \rightarrow \pi^0\pi^0$ MC. 156
- 6.2 Shown here is the $\sum_\gamma \Delta\theta\chi^2$ for even (left) and odd (right) pairing $K_L^0 \rightarrow \pi^0\pi^0$ two cluster MC with all other kinematic cuts applied. The $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ analysis cut is at 1.0. 157
- 6.3 Shown here is π^0 -projection cut parameters (in solid white) with all other cuts applied to $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ MC. Events outside the bounds are rejected. The thin, vertical red lines at $z = 300, 500$ are to guide the eye only. 158
- 6.4 A cartoon illustration of the Al Target Run (π^0 -Run) set-up. The Al plate is located about 6.5 cm downstream of the edge of CC02 and is 5 mm thick. 161
- 6.5 The invariant mass spectrum from the π^0 -Run under special cuts is shown on the left, and the results of applying the $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ physics analysis cuts to the π^0 -Run is on the right. In the mass plot, the π^0 peak is dominant, with a second peak at high mass from $\eta \rightarrow \gamma\gamma$. The intermediate region is due to $K \rightarrow \gamma\gamma$ and the very low-mass region is dominated by soft photons from other neutron interactions. The P_T vs. z -vertex plot with $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ analysis cuts is compared to similar distributions for the physics data in Figures 6.6 and 6.7. 162
- 6.6 The z -vertex spectrum for the physics data (dotted) and π^0 -Runs (solid) with $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ analysis cuts imposed. The P_T vs. z -vertex plot for the π^0 -Runs is shown in Figure 6.5 and for the data in Figure 6.26 162

6.7	The $ P $ (left) and P_T (right) spectra for the physics data (dotted) and π^0 -Runs (solid) with analysis cuts imposed. The P_T vs. z -vertex plot for the π^0 -Runs is shown in Figure 6.5 and for the data in Figure 6.26.	163
6.8	Shown here is the calculated momentum versus the true momentum for the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC (left) and the π^0 -Run data with and without a cut on the “missing” momentum (right).	164
6.9	The CV- π^0 z -vertex distribution with all cuts (left) and kinematic cuts (right). With all cuts, there are 17 events in the data and 18.2 ± 6.1 with P.O.T. scaling in the MC (we must scale the MC up by roughly a factor of two).	165
6.10	The geometrical relationship behind the idea of a bifurcation analysis.	166
6.11	The mass spectrum from the π^0 -Run compared with the composite MC. Normalization is via protons on target (P.O.T.) from the E391a beam-line simulation [45]. There is a 19.8% scaling loss applied to the MC to account for accidental activity and an additional factor of ~ 1.3 used to scale the contribution from $K_L^0 \rightarrow \gamma\gamma$. The most important feature though is the agreement in relative peak height between η and π^0 decays with only P.O.T. normalization. The very low mass discrepancy is rooted in an energy cut applied in the MC to remove very low energy neutron interactions that were incapable of producing interesting secondaries.	167
6.12	The total momentum (left) and transverse momentum (right) spectra from the π^0 -Run compared with the composite MC with all analysis cuts (no requirements on P_T). Both plots consider only the region in Figure 6.11 with mass greater than $520 \text{ MeV}/c^2$.	168
6.13	Shown here is the sum of the angle χ^2 values (left) and the γ -RMS distributions (center and right) for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC, the π^0 -Run data, and the η -MC with all other cuts applied. The π^0 -Run and the η -MC are normalized to each other by proton on target, but the scale of the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC is arbitrary.	169
6.14	Shown here is a set of vectors comparing the exclusive “acceptance” of the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC and the various neutron background modes. Here, “eta” refers to the CV- η events and “cv” refers to the CV- π^0 events. Some entries for the CV- π^0 acceptance are larger than one because low statistics in the MC forced us to employ some ratio arguments in a “bifurcation” analysis. Some of these ratios involved small numbers that due to rounding error or some other problem caused the acceptance estimate to fail. The other entries are to be considered reliable though. The entry “gnid” refers to an unused cut.	170
6.15	The acoplanarity angle as defined in Equation 6.5.	171

- 6.16 Shown here is the $K_L^0 \rightarrow \gamma\gamma$ background plot with loose cuts (left) and the corresponding acoplanarity angle (see Equation 6.5) distribution (right) for a MC sample roughly four times the size of the data. The cut line for the acoplanarity angle is also displayed with the distribution (all $K_L^0 \rightarrow \gamma\gamma$ events are removed by the cut). 172
- 6.17 Fusion takes place when (at least) two true photon hits are reconstructed as a single hit. Here, the true photon hit positions are illustrated with circles with radius roughly proportional to their energies. 174
- 6.18 $K_L^0 \rightarrow \pi^0\pi^0$ backgrounds with only veto cuts applied (left) or only kinematic cuts applied (right). The two circled events in the plot with vetoes applied only are the two events that survive *all* cuts. Low P_T events from the plot on the right are removed by the acoplanarity cut (below 0.05 GeV/c) and by the π^0 -projection cuts. 175
- 6.19 Here we plot P_T versus $\Delta z = z_{True} - z_{Rec}$ for two cluster $K_L^0 \rightarrow \pi^0\pi^0$ events with photon vetoes only applied (left) and kinematics only applied (right). Even background events are those with small values of Δz 176
- 6.20 $K_L^0 \rightarrow \pi^0\pi^0$ background levels with various *exclusive* cut sets. The axis label denotes which cut (or group of cuts) is removed, with all other cuts applied. (See Table 6.3.) 177
- 6.21 The number of hits in the Main Barrel with all cuts applied *except* the MB are shown versus P_T on the left and against reconstructed vertex on the right. All hits come from events inside the signal region. . . . 178
- 6.22 The number of hits in the CsI with all cuts applied *except* the CsI are shown versus P_T on the left and against reconstructed vertex on the right. All hits come from events inside the signal region. 179
- 6.23 Shown here is a cartoon depicting two ways in which particles may “exchange” their electric charge in a Ke3 event: 1) by $\pi^- + p \rightarrow n + \pi^0$ and 2) by $e^+ + e^- \rightarrow \gamma\gamma$. Figure adapted from [56]. 182
- 6.24 Shown here is the distribution of background events from the Ke3 MC under loose cuts. The signal event counts are not scaled for statistics. 184
- 6.25 Shown here is the radial position versus z -vertex for the charged pion decays in the Ke3 MC with loose cuts imposed (all the charged veto subsystems are removed in addition to some of the kinematic cuts). Almost all of the events with (z,r) positions beyond ~ 550 cm in z and ~ 15 cm in r were counted as charged veto hits. 186
- 6.26 Shown here is the two cluster data final plot with the signal region still blind. Note the four numbered regions. In (P_T, z) space Region 1 stretches from (300 cm, 0.12 GeV/c) to (340 cm, 0.24 GeV/c); Region 2 stretches from (340 cm, 0.12 GeV/c) to (400 cm, 0.24 GeV/c); Region 3 stretches from (400 cm, 0.12 GeV/c) to (500 cm, 0.24 GeV/c); Region 4 stretches from (300 cm, 0.0 GeV/c) to (500 cm, 0.12 GeV/c). . . . 188

7.1	Shown above is the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC exclusive acceptance for the kinematic cuts. The exclusive acceptance for a cut is the acceptance for that cut only with all other cuts applied.	192
7.2	Shown above is the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC exclusive acceptance for the photon veto cuts. The exclusive acceptance for a cut is the acceptance for that cut only with all other cuts applied.	193
7.3	Shown above is the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC final plot (all analysis cuts are imposed).	194
7.4	Shown here is the two-cluster data final plot (all $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ analysis cuts are imposed). There are 141 events in the plot with a single overflow event at $z \sim 550$ with very high P_T	195

LIST OF TABLES

1.1	Some key parameters of the Unitarity Triangle. [22]	16
2.1	A representative set of parameters for the online cuts in Run-II (they were occasionally changed in response to ongoing analysis and electronics problems).	38
4.1	Accidental losses for the main calibration modes. Here, errors are purely statistical.	84
4.2	Kinematic Cuts for Six Cluster Events. These cuts are used in the analysis along with photon veto cuts to compute the Kaon flux. Additionally, these cuts served as the primary cut set for the MC tuning process.	86
4.3	Decay probability. This quantity applies to events generated at C6 in the MC, and has no other meaningful interpretation. This value is important for flux computation because we define our flux as the number of Kaons decaying inside the fiducial, and our acceptance as the fraction of those decays that are visible (reconstructed). We chose this definition for the flux for ease of understanding and naturalness.	90
4.4	Acceptance estimates. Here the quoted error is the <i>statistical error</i> . The discrepancy for $K_L^0 \rightarrow 3\pi^0$ is likely rooted in the fact that most of the acceptance loss due to photon vetoes is from accidental losses, and we do not capture accidental activity perfectly in our accidental trigger. The other modes have real photons from other Kaon decays that dominate the acceptance loss in the vetoes.	90
4.5	MC Sample Sizes. Here, K is short-hand for K_L^0 and a flux value of 5.13×10^9 is generically used to normalize to Run II. The Kmu3 and Ke3 samples are one half in the positive pion channel and one half in the negative channel.	91
4.6	Crystal-by-crystal resolution function parameters (see Equation 4.3).	92
4.7	$K \rightarrow 3\pi^0$ invariant mass characteristics in data and MC with all cuts applied. The mass is biased slightly high due to a feature of the reconstruction routine.	93
4.8	Integral sums over the invariant mass spectrum of the MC and Data.	95
4.9	The veto cuts. Recall that $1 \text{ GeV} = 10^3 \text{ MeV} = 10^9 \text{ eV}$	130
4.10	Kinematic Cuts for Four Cluster Events.	131
4.11	Kinematic Cuts for $K_L^0 \rightarrow \gamma\gamma$ Events.	131
4.12	Components of the systematic error for $K_L^0 \rightarrow \pi^0\pi^0$	132
4.13	Flux estimates. Here the quoted error is the <i>total error</i> and is dominated by the systematic error. The “Geo-Fix” samples refer to estimates made with reduced size samples that were free of the CC03/CC04 geometry bug.	132

5.1	Energy and position resolution values applied to MC cluster reconstruction during the 1-C Fit study. These values do not align precisely with the true resolution of the two experiments, but are quite close [48].	143
6.1	Kinematic Cuts for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ Events.	153
6.2	Normalization and control regions for the CC02-background prediction through the π^0 -Runs.	163
6.3	Background levels for $K_L^0 \rightarrow \pi^0 \pi^0$ exclusive cut sets (see Figure 6.20). The scale factor is dictated by the size of the MC sample relative to data (MC:data = 9.6:1).	177
6.4	The correlation is here defined according to Equation 6.6.	180
6.5	The raw hits and kinematic suppression factors for the different event topologies (single and double charged system hits). The last column is a product of the first three and represents the background estimate prior to scaling (up) for MC statistics (by a factor of $1/0.021 \sim 50$ for Ke3). The assumed veto suppression for the CV is based on an upper limit estimated through beam-line muon studies [57]. The other suppressions are speculative.	185
6.6	Predicted and observed backgrounds in the region around the signal box. The predicted and observed background levels are statistically consistent. See Figure 6.26 for the associated plot.	187
7.1	Summary of the Run II results.	196
7.2	Comparison of the Run I and Run II results.	197

CHAPTER 1

INTRODUCTION

The study of symmetries, and their violation, is the heart and essence of particle physics. Perhaps the most fundamental symmetry in all of physics is CPT - a symmetry conglomerated from three others: Charge Conjugation symmetry (C), a symmetry under multiplicative quantum number sign exchange; Parity symmetry (P), an inversion of space; and finally Time reversal symmetry (T). Originally something of a curiosity, CPT has been elevated to the highest importance. It is deeply related to Lorentz invariance and therefore causality, making it a cornerstone of any “rational” system of physics.

Initially, each of the components of CPT were all thought to be good symmetries. All of them are known to hold in electrodynamics and they were all originally assumed to hold in the strong and weak interactions as well. P was a sacred geometrical symmetry of spacetime prior to the suggestion of Lee and Yang [3] that there was no actual *evidence* the symmetry was conserved in weak interactions. P violation was shortly thereafter observed experimentally in 1957 [4], beginning a revolution in thinking about symmetry violation in physics. C symmetry was originally thought of as the correct formulation of a matter-antimatter exchange symmetry, and the early equations of quantum field theory gave no basis for discrimination between the two. However, not long after P violation was shown to be true, C violation was established as well [5]. Lev Landau suggested that CP instead of C might actually be the appropriate symmetry for matter-antimatter exchange - under CP the weak interactions seemed to make sense again. The weak interaction only “sees” left-handed neutrinos. Under C , $\nu \rightarrow \bar{\nu}$, but chirality is unchanged. The interaction couples only to left-handed ν 's but not at all to left-handed $\bar{\nu}$'s. But with P also invoked, $\nu_L \rightarrow \bar{\nu}_R$

and the traditional concept of an anti-particle is restored. It appeared at the time that CP and T were both good symmetries in the weak interactions, and the vector-axial ($V - A$) theory built in 1957 and 1958 to describe weak interactions has these symmetries embedded within it.

However, CP violation was shown to exist in the Kaon system in 1964 [6]. Since we believe CPT to be a good symmetry, the violation of CP implies the violation of T . The surprising fashion in which CP fell has left an indelible mark on physics. When faced with a new symmetry, the *assumption* now is that it is violated. At least the possibility it is violated is *always* held open. While we firmly believe, for instance, that CPT is a good symmetry (at least to very high order), it does not raise an eyebrow when an experimental test is announced. Nothing is truly “sacred” anymore. For quite some time thereafter, the neutral Kaon system was the only system observable in nature that exhibited CP violation. In fascinating contrast, while C and P violation in weak interactions are individually maximal - the symmetries are expressly *not* symmetries of our universe - CP is *almost* conserved. Deviations are small and rare. While we eventually observed CP violation in other systems, notably the B system (mesons composed of bottom quarks), Kaons continue to maintain an important place in particle physics in the study of fundamental symmetry violation.

Understanding the source of CP violation is quite important. Not only is it interesting purely in its own right, it has the potential to shed light on matter-antimatter asymmetries in the universe and it offers a powerful crucible to test extensions to the Standard Model (SM) of particle physics. The decay $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ is particularly special for studying CP violation owing to its exceptional *theoretical* cleanness. Most other prominent meson decays (especially in the B system) face theoretical uncertainties in their predicted SM branching ratios as high as $\pm 10\%$, but in the case of $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$, the uncertainty is an order smaller, at the level of $1 - 2\%$. The

reason for this is the high suppression of long-distance and high order electroweak effects while the required hadronic matrix elements may be extracted from ratios using $K^+ \rightarrow \pi^0 e^+ \nu$ [7], [8]. $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ is almost purely CP violating in the SM, making it an ideal testing ground for the process of CP violation and excellent probe of any new physics which affect the way the symmetry is broken in the SM [9, And references within.]

$K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ is a very challenging decay to measure. The predicted SM branching ratio is quite small:

$$Br(K_L^0 \rightarrow \pi^0 \nu \bar{\nu})_{\text{SM}} \simeq (2.5 \pm 0.4) \times 10^{-11}, \quad (1.1)$$

where the error is dominated by experimental uncertainties in the inputs (a breakdown of the errors is provided after Equation 1.42.) This branching ratio is almost four orders of magnitude lower than the best measurements prior to E391a [10]. The decay is particularly difficult because the final state neutrinos are not detectable, leaving an experimental signature of a single π^0 + missing transverse momentum (P_T). Not only can a number of Kaon decays fake this signature, $K_L^0 \rightarrow \pi^0 \pi^0$ with one missing π^0 for example, but non-Kaon related processes are quite capable of faking this signature as well. Therefore, any attempt to measure the branching ratio of $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ is required to have exceptionally good signal-to-noise and background suppression characteristics. Indeed, because of the low branching ratio, the entire measurement is essentially a pure exercise in background estimation and rejection.

As mentioned above, $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ is almost purely CP violating in the Standard Model (SM), with the indirect (mixing) component highly suppressed. This means that CP violation manifests in the decay amplitude itself and the decay process depends very strongly on the mechanism of CP violation in the SM. A positive

signal at the level of 10^{-11} is an indicator of indirect CP violation contributing to the branching ratio at the level of only one percent, making this mode an excellent candidate to study direct CP violation in a theoretically clean fashion. Because the decay is almost entirely CP violating in the SM, only the top quark contribution is important and uncertainties associated with the charm sector are eliminated.

1.1 Kaon Phenomenology

This discussion mostly follows [11], [12], and [13].

Kaons are mesons formed by the combination of a strange quark and a first generation (up or down) quark. The neutral Kaon is distinguished from its anti-particle by strangeness only: $S(K^0) = 1$, $S(\bar{K}^0) = -1$. Strangeness is a good quantum number for strong and electromagnetic interactions, but $|\Delta S| = 1$ transitions like $K^0 \rightarrow \pi^+\pi^-$ and $\bar{K}^0 \rightarrow \pi^+\pi^-$ prove that strangeness is not conserved by weak interactions. Neutral Kaons are produced in strong interactions, but because strangeness is conserved in those interactions, they must decay through the weak force. What's more, because K^0 and \bar{K}^0 share final states, they are coupled through $|\Delta S| = 2$ transitions - the neutral Kaon and it's anti-particle can mix - see Figure 1.1. By definition, $S(K^+) = 1$, and the neutral Kaons are defined as:

$$\begin{aligned} K^0 &= \begin{pmatrix} d \\ \bar{s} \end{pmatrix} & (S = +1), \\ \bar{K}^0 &= \begin{pmatrix} \bar{d} \\ s \end{pmatrix} & (S = -1). \end{aligned} \tag{1.2}$$

The similarity between \bar{K}^0 and K^0 turns out to yield some very crucial interference effects. CP violation in the Standard Model exists as a complex phase. We can only access this phase in a quantum mechanical world by examining interference effects

like those we see in the neutral Kaon (and B) system.




Figure 1.1: Diagrams mediating K^0 and \bar{K}^0 mixing. The charge of the W bosons is suppressed.

In particular, we note the following:

$$\begin{aligned} C|K^0\rangle &= |\bar{K}^0\rangle, \\ C|\bar{K}^0\rangle &= |K^0\rangle, \end{aligned} \tag{1.3}$$

$$\begin{aligned} P|K^0\rangle &= -|K^0\rangle, \\ P|\bar{K}^0\rangle &= -|\bar{K}^0\rangle, \end{aligned} \tag{1.4}$$

$$\begin{aligned} CP|K^0\rangle &= -|\bar{K}^0\rangle, \\ CP|\bar{K}^0\rangle &= -|K^0\rangle, \end{aligned} \tag{1.5}$$

where these equations fix our phase conventions. The neutral Kaon is not an eigenstate of CP . In order to have a CP eigenstate, we must define new, linear combinations of the two, motivated physically by the $|\Delta S| = 2$ transitions that couple them:

$$\begin{aligned} |K_1^0\rangle &= \frac{1}{\sqrt{2}} \left(|K^0\rangle + |\bar{K}^0\rangle \right), \\ |K_2^0\rangle &= \frac{1}{\sqrt{2}} \left(|K^0\rangle - |\bar{K}^0\rangle \right), \end{aligned} \tag{1.6}$$

where $|K_1^0\rangle$ is CP even and $|K_2^0\rangle$ is CP odd. Originally, these combinations were

proposed as eigenstates to the complete Hamiltonian for the neutral Kaon [14]. While close, this turns out not to be true.

With K^0 and \bar{K}^0 mixing, we write the time evolution as:

$$i\frac{d\psi(t)}{dt} = \mathcal{H}\psi(t), \quad \psi(t) = \begin{pmatrix} |K^0(t)\rangle \\ |\bar{K}^0(t)\rangle \end{pmatrix}. \quad (1.7)$$

This is a beautiful example of a two-state quantum system. Without mixing, the Hamiltonian is $\mathcal{H} = M - i\Gamma/2$, where M is the mass and Γ is the lifetime. We can easily extend this to the case with mixing as:

$$\mathcal{H} = \hat{M} - i\frac{\hat{\Gamma}}{2} = \begin{pmatrix} M_{11} - i\Gamma_{11}/2 & M_{12} - i\Gamma_{12}/2 \\ M_{21} - i\Gamma_{21}/2 & M_{22} - i\Gamma_{22}/2 \end{pmatrix}, \quad (1.8)$$

where the subscripts label the K_1 (CP -even) and K_2 (CP -odd) states. We can simplify this expression by invoking hermiticity ($M_{21} = M_{12}^*$) and CPT ($M_{11} = M_{22} = M$ and $\Gamma_{11} = \Gamma_{22} = \Gamma$), leaving us with:

$$\mathcal{H} = \begin{pmatrix} M - i\Gamma/2 & M_{12} - i\Gamma_{12}/2 \\ M_{12}^* - i\Gamma_{12}^*/2 & M - i\Gamma/2 \end{pmatrix}. \quad (1.9)$$

If we then diagonalize this matrix, we find as eigenstates:

$$\begin{aligned} K_L^0 &= \frac{(1 + \tilde{\epsilon})K^0 + (1 - \tilde{\epsilon})\bar{K}^0}{\sqrt{2(1 + \tilde{\epsilon}^2)}} = \frac{K_2 + \tilde{\epsilon}K_1}{\sqrt{1 + \tilde{\epsilon}^2}}, \\ K_S^0 &= \frac{(1 + \tilde{\epsilon})K^0 - (1 - \tilde{\epsilon})\bar{K}^0}{\sqrt{2(1 + \tilde{\epsilon}^2)}} = \frac{K_1 + \tilde{\epsilon}K_2}{\sqrt{1 + \tilde{\epsilon}^2}}, \end{aligned} \quad (1.10)$$

where $\tilde{\epsilon}$ is a small, complex parameter defined through the relation:

$$\frac{1 - \tilde{\epsilon}}{1 + \tilde{\epsilon}} = \sqrt{\frac{M_{12}^* - (i/2)\Gamma_{12}^*}{M_{12} - (i/2)\Gamma_{12}}} \equiv r e^{i\kappa}. \quad (1.11)$$

The eigenvalues of the above Hamiltonian may be written as:

$$\begin{aligned} M_{L,S} &= M \pm \text{Re}(Q), \\ \Gamma_{L,S} &= \Gamma \mp 2 \text{Im}(Q), \end{aligned} \quad (1.12)$$

with Q defined as:

$$Q = \sqrt{(M_{12} - (i/2)\Gamma_{12})(M_{12}^* - (i/2)\Gamma_{12}^*)}. \quad (1.13)$$

This allows us to write $\Delta M = M_L - M_S = 2 \text{Re}(Q)$ and $\Delta \Gamma = \Gamma_L - \Gamma_S = -4 \text{Im}(Q)$. We should note that $\tilde{\epsilon}$ is dependent on phase convention, and so not a true physical observable. However, $\text{Re}(\tilde{\epsilon})$ and r are independent of phase convention. The deviation of r from unity, in fact, is a measure of CP violation in neutral Kaon mixing:

$$r = 1 + \frac{2|\Gamma_{12}|^2}{4|M_{12}|^2 + |\Gamma_{12}|^2} \text{Im}(M_{12}/\Gamma_{12}). \quad (1.14)$$

We pause to collect the consequences of the derivations above. The eigenstates of the weak Hamiltonian we have written down are *not* equal to K_1 and K_2 . They are quantum superpositions of those states, which were themselves superpositions of the strong eigenstates K^0 and \bar{K}^0 . So neutral Kaons formed as strong eigenstates mix into CP eigenstates, which in turn mix into weak eigenstates in order to decay! The particles we observe in the lab and call Kaons are superpositions of superpositions of quantum states - fascinating beasts indeed.

1.1.1 Indirect and Direct CP Violation in the Kaon System

This subsection is of primarily historical interest. It is only within the last decade that the existence of *direct* CP violation was established. The subject has a rich history at the University of Chicago and forged many new doctorates.

Two pion final states are CP even and three pion final states are CP odd. Therefore, K_S^0 and K_L^0 decay preferentially into two and three pion states because the parameter $\tilde{\epsilon}$ that describes the admixture of states above is small. In other words, since $K_S^0 \sim K_1$ and $K_L^0 \sim K_2$, the lifetime of K_S^0 is shorter than K_L^0 by a factor of almost six hundred due to the enlarged phase space for a two pion decay versus three.

However, K_S^0 is partially K_2 and K_L^0 contains a small piece of K_1 . Therefore, we can see decays like $K_L^0 \rightarrow \pi\pi$ through the K_1 piece of K_L^0 . This is what is meant by *indirect* CP violation. The measure for this form of violation is defined by the relation:

$$\begin{aligned} \epsilon &= \frac{A(K_L^0 \rightarrow (\pi\pi)_{I=0})}{A(K_S^0 \rightarrow (\pi\pi)_{I=0})}, \\ &= \frac{e^{i\pi/4}}{\sqrt{2}\Delta M_K} (\text{Im}(M_{12}) + 2\zeta \text{Re}(M_{12})), \end{aligned} \tag{1.15}$$

where $\zeta = \text{Im}(A_0)/\text{Re}(A_0)$ and I is an isospin state label. Note that the piece of ϵ involving $\text{Im}(M_{12})$ represents $\tilde{\epsilon}$ from above, but the phase dependence of ζ cancels the phase dependence from $\tilde{\epsilon}$, leaving ϵ free of any such dependencies.

While indirect CP violation is manifested here a consequence of the superposition of different CP eigenstates in the weak eigenstate, direct CP violation appears through the direct transition of a CP even to odd state (or vice versa) - see Figure 1.2.

Direct CP violation is characterized in two pion decays as the complex parameter ϵ' :

$$\epsilon' = \frac{1}{\sqrt{2}} \text{Im} \left(\frac{A_2}{A_0} \right) e^{i\Phi_{\epsilon'}}, \quad (1.16)$$

where $\Phi_{\epsilon'} = \pi/2 + \delta_2 - \delta_0$ and the A_I and δ_i parameters are the isospin and amplitudes and strong phases respectively from K^0 to two pion amplitudes:

$$\begin{aligned} A(K^0 \rightarrow \pi^+\pi^-) &= \sqrt{2/3}A_0e^{i\delta_0} + \sqrt{1/3}A_2e^{i\delta_2} \\ A(K^0 \rightarrow \pi^0\pi^0) &= \sqrt{2/3}A_0e^{i\delta_0} - 2\sqrt{1/3}A_2e^{i\delta_2} \end{aligned} \quad (1.17)$$

While the isospin amplitudes A_I are complex and dependent on phase conventions, ϵ' measures the difference between the phases and is a physical quantity. The strong phases can be pulled out of $\pi\pi$ scattering, leaving $\Phi_{\epsilon'} \approx \pi/4$.




Figure 1.2: Direct and indirect CP violation in the decay of the $K_L^0 \rightarrow \pi\pi$.

Experimentally, ϵ and ϵ' can be measured through the ratios:

$$\begin{aligned} \eta_{+-} &= \frac{A(K_L \rightarrow \pi^+\pi^-)}{A(K_S \rightarrow \pi^+\pi^-)} = |\eta_{+-}| e^{i\phi}, \\ \eta_{00} &= \frac{A(K_L \rightarrow \pi^0\pi^0)}{A(K_S \rightarrow \pi^0\pi^0)} = |\eta_{00}| e^{i\phi}. \end{aligned} \quad (1.18)$$

If we assume ϵ and ϵ' are small, we can write:

$$\begin{aligned}\eta_{00} &= \epsilon - \frac{2\epsilon'}{1 - \sqrt{\omega}} \approx \epsilon - 2\epsilon', \\ \eta_{+-} &= \epsilon + \frac{\epsilon'}{1 + \omega/\sqrt{2}} \approx \epsilon + \epsilon',\end{aligned}\tag{1.19}$$

with $\omega = \text{Re}(A_2)/\text{Re}(A_0)$. If direct CP violation does not exist, $\eta_{00} = \eta_{+-}$. Any deviation then can be compactly expressed:

$$|\eta_{00}/\eta_{+-}|^2 \approx 1 - 6 \text{Re}(\epsilon'/\epsilon).\tag{1.20}$$

The measured value is $\text{Re}(\epsilon'/\epsilon) = (20.7 \pm 1.48_{stat} \pm 2.39_{syst}) \times 10^{-4} = (20.7 \pm 2.8) \times 10^{-4}$ [15] (with further details in [16].) Significantly, it is *not* zero, thus demonstrating the existence of direct CP violation.

1.2 The Standard Model and CP Violation.

With a basic understanding of the Kaon in hand, let us turn to how it fits into the larger picture of the Standard Model (SM). It is possible to say a great deal about the SM, but we will offer only a few scandalously brief details before turning further questions over to the references and honing in on the ingredients key to understanding the importance of $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$. This discussion mostly follows [13].

The SM is defined by three generations of quarks and leptons with interactions defined by the gauge group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$, broken spontaneously into $SU(3)_C \otimes U(1)_Q$, with the following features:

- Eight gluons, G_α mediate the strong interactions, while the W^\pm , Z^0 , and γ are the mediators of the electroweak interaction.

- Left-handed quarks and leptons are put into $SU(2)_L$ doublets according to the electroweak force:

$$\begin{pmatrix} \nu_e \\ e^- \end{pmatrix} \quad \begin{pmatrix} \nu_\mu \\ \mu^- \end{pmatrix} \quad \begin{pmatrix} \nu_\tau \\ \tau^- \end{pmatrix} \quad (1.21)$$

$$\begin{pmatrix} u \\ d' \end{pmatrix} \quad \begin{pmatrix} c \\ s' \end{pmatrix} \quad \begin{pmatrix} t \\ b' \end{pmatrix} \quad (1.22)$$

The corresponding right-handed fields transform as $SU(2)_L$ singlets and the primes will be discussed shortly.

- Flavor violating charged-current processes are mediated by the W^\pm with strength given by the gauge coupling and a 3×3 unitary matrix:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix} = V_{CKM} \begin{pmatrix} d \\ s \\ b \end{pmatrix} \quad (1.23)$$

This matrix, the *Cabibbo-Kobayashi-Maskawa* matrix [17] [18], connects the weak eigenstates (denoted with primes above) and the mass eigenstates. The leptonic sector shares an analogous mixing matrix between neutrino flavors.

- Because the CKM matrix is unitary, flavor-changing neutral current transitions are forbidden at the tree level (loop diagrams are required). This suppression is called the GIM mechanism [19].
- Because the CKM matrix is unitary, the V_{ij} parameters may be complex - this allows “built-in” CP violation in the quark sector of the SM because the

charged-current Lagrangian is given by:

$$\mathcal{L}_{CC} = \frac{g}{\sqrt{2}} \left[\bar{u}_i V_{ij} d_j W^- + \bar{d}_j V_{ij}^* u_i W^+ \right], \quad (1.24)$$

where u and d are quark fields and the W 's are mediators of the weak force.

This Lagrangian transforms under CP according to:

$$\mathcal{L}_{CC} \xrightarrow{CP} \frac{g}{\sqrt{2}} \left[\bar{d}_j V_{ij} u_i W^+ + \bar{u}_i V_{ij}^* d_j W^- \right]. \quad (1.25)$$

So, if $V_{ij} = V_{ij}^*$, CP is conserved, but complex entries would necessarily violate the symmetry.

1.2.1 The CKM Matrix

The CKM matrix initially has nine free real parameters, but this number is reduced to four by the definition of the quark phase. There are several popular parameterizations of V . One involves three angles (θ_{12} , θ_{23} , and θ_{31}) and a phase (δ):

$$V = \begin{pmatrix} c_{12} c_{13} & s_{12} c_{13} & s_{13} e^{i\delta} \\ -s_{12} c_{23} - c_{12} s_{23} s_{13} e^{i\delta} & c_{12} c_{23} - s_{12} s_{23} s_{13} e^{i\delta} & s_{23} c_{13} \\ s_{12} c_{23} - c_{12} s_{23} s_{13} e^{i\delta} & -c_{12} c_{23} - s_{12} s_{23} s_{13} e^{i\delta} & c_{23} c_{13} \end{pmatrix} \quad (1.26)$$

with c_{ij} defined as $\cos(\theta_{ij})$ and s_{ij} defined as $\sin(\theta_{ij})$. We know that s_{13} and s_{23} are small numbers ($\sim 10^{-3}$), so to a good approximation $c_{13} = c_{23} = 1$ and the four independent parameters are: $s_{12} = |V_{us}|$, $s_{13} = |V_{ub}|$, $s_{23} = |V_{cb}|$, and δ . This is the so-called ‘‘standard’’ parameterization and has the advantages of both being free of any further conventions and quite useful for numerical computations. It has the disadvantage though of being somewhat opaque with regards to any hierarchy or

structure in the parameters.

Perhaps the most famous parameterization is due to Wolfenstein [20], who wrote down the matrix as an expansion of $\lambda = |V_{us}| \approx 0.22$:

$$V = \begin{pmatrix} 1 - \lambda^2/2 & \lambda & A \lambda^3 (\rho - i \eta) \\ -\lambda & 1 - \lambda^2/2 & A \lambda^2 \\ A \lambda^3 (1 - \rho - i \eta) & -A \lambda^2 & 1 \end{pmatrix} + \mathcal{O}(\lambda^4) \quad (1.27)$$

This parameterization wonderfully reveals the internal structure between the parameters, but has the serious drawback of not being unitary - higher order terms in λ must be included to provide consistent calculations.

The most efficient way to remove these problems while keeping a transparent structure is return to the standard parameterization and simply define:

$$\begin{aligned} s_{12} &= \lambda \\ s_{23} &= A \lambda^2 \\ s_{13} e^{-i \delta} &= A \lambda^3 (\rho - i \eta) \end{aligned} \quad (1.28)$$

to all orders in λ . Therefore, we have:

$$\rho = \frac{s_{13}}{s_{12} s_{23}} \cos \delta, \quad \eta = \frac{s_{13}}{s_{12} s_{23}} \sin \delta. \quad (1.29)$$

By making this change of variables we have a matrix that is a function of (λ, A, ρ, η) that satisfies unitarity exactly. If we expand in powers of λ again, we can recover Equation 1.27 and find explicit corrections for higher order terms in λ if we need them

(see [13] for a listing). To excellent accuracy though, we may write:

$$\begin{aligned}
 V_{us} &= \lambda \\
 V_{cb} &= A \lambda^2 \\
 V_{ub} &= A \lambda^3 (\rho - i \eta) \\
 V_{td} &= A \lambda^3 (1 - \bar{\rho} - i \bar{\eta})
 \end{aligned}
 \tag{1.30}$$

where

$$\bar{\rho} = \rho \left(1 - \frac{\lambda^2}{2} \right), \quad \bar{\eta} = \eta \left(1 - \frac{\lambda^2}{2} \right).
 \tag{1.31}$$

1.2.2 The Unitarity Triangle

The CKM matrix is unitary - therefore the dot of a column or row with its complex conjugate is equal to one, and the dot with the complex conjugate of other rows or columns must be zero. If we apply this condition to the first and third columns of the CKM matrix we obtain a relation expressible as a sum of three vectors in the complex plane that sum to zero - in other words, a triangle in the complex plane. This is one formulation of the famous Unitarity Triangle:

$$V_{ub}^* V_{ud} + V_{cb}^* V_{cd} + V_{tb}^* V_{td} = 0.
 \tag{1.32}$$

The area of this triangle is proportional to a fundamental quantity, J , that is independent of phase parameterization convention:

$$J = \text{Im} \left[V_{ij} V_{kl} V_{kj}^* V_{il}^* \right] \quad (i \neq l, j \neq k),
 \tag{1.33}$$

the so-called Jarlskog parameter [21]. In the Wolfenstein parameterization, $J \approx \lambda^6 A^2 \eta$.

We can construct the triangle according to the following steps:

1. The product $V_{cd}V_{cb}^* = -A\lambda^3 + \mathcal{O}(\lambda^7)$ is to an excellent approximation real. Therefore, $1/(A\lambda^3) \times |V_{cd}V_{cb}^*| \approx 1$ forms a vector between the origin and (1,0) in the complex plane.
2. If we re-scale by that same factor, $A\lambda^3$, we find the following:

$$\frac{1}{A\lambda^3}V_{ud}V_{ub}^* = \bar{\rho} + i\bar{\eta} \quad \frac{1}{A\lambda^3}V_{td}V_{tb}^* = 1 - (\bar{\rho} + i\bar{\eta}). \quad (1.34)$$

3. The sum of these three complex vectors closes (as it should by construction) as illustrated in Figure 1.3.




Figure 1.3: The Unitarity Triangle with the impacts of $K \rightarrow \pi\nu\bar{\nu}$ on the parameters visualized.

A comprehensive review of the parameters of the Unitarity Triangle and the methods of their estimation may be found in [22]. For a brief summary of the most im-

portant parameters, see Table 1.1.

Quantity	1- σ CL (+/- denote uncertainties)
λ	$0.2265 + 0.0025 - 0.0023$
A	$0.801 + 0.029 - 0.020$
$\bar{\rho}$	$0.189 + 0.088 - 0.070$
$\bar{\eta}$	$0.358 + 0.046 - 0.042$

Table 1.1: Some key parameters of the Unitarity Triangle. [22]

1.3 $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ and the CKM Matrix

1.3.1 Theoretical Framework

The most important calculations for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ rely on a technique called Operator Product Expansion (OPE) [13] [23]. OPE is an organizational scheme for the phenomenology of hadronic weak decays in which the Hamiltonian is written as:

$$\mathcal{H}_{eff} = \frac{G_F}{\sqrt{2}} \sum_i V_{CKM}^i C_i(\mu) Q_i, \quad (1.35)$$

with G_F equal to the Fermi constant and the Q_i are built from quark and lepton fields and function as local operators governing the decays in question. The V_{CKM}^i are parameters of the CKM matrix and the C_i are *Wilson coefficients* governing the strength for a given operator in the Hamiltonian. The amplitude for a decay like $K \rightarrow F$ for some final state F is written in this language like:

$$A(K \rightarrow F) = \langle F | \mathcal{H}_{eff} | K \rangle = \frac{G_F}{\sqrt{2}} \sum_i V_{CKM}^i C_i(\mu) \langle F | Q_i | K \rangle, \quad (1.36)$$

where $\langle F | Q_i(\mu) | K \rangle$ are the hadronic matrix elements between K and F .

The largest advantage of the OPE scheme is that it separates the calculation of the amplitude into two pieces - a perturbative (short range) piece built out of the Wilson coefficients, and a non-perturbative (long distance) piece built from the hadronic matrix elements.

The Wilson coefficients depend on the masses of the top quark and any other new particles and can be found through box and penguin diagrams. These diagrams must include all mediator (W^\pm and Z^0) and particle exchange and short distance quantum chromodynamics (QCD) effects.

As for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$, the decay proceeds through penguin and box diagrams with internal top and charm quark exchanges. See Figure 1.4 for an illustration. Evaluating these diagrams (and similar diagrams for other K and even B decays), one finds a set of universal functions of the top quark mass, $F_r(x_t)$, where $x_t = m_t^2/M_W^2$. We may express the OPE Hamiltonian as a function of these:

$$A(K \rightarrow F) = P_0(K \rightarrow F) + \sum_r P_r(K \rightarrow F) F_r(x_t). \quad (1.37)$$

This equation is known as a Penguin-Box Expansion (PBE) [13]. This expansion was originally intended to highlight the dependence of flavor-changing neutral current (FCNC) processes on the mass of the top, but with the top mass fairly well measured, this is no longer a fashionable application. However, the PBE is excellent for studying extensions of the SM where new particles appear in these loops. The $F_r(x_t)$ functions acquire a dependency on the mass of new particles (e.g. supersymmetric particles, the Higgs boson, etc.), but the forms of the P_0 and P_r functions end up unchanged, making the effects of new physics explicit.




Figure 1.4: The penguin and box Feynman diagrams for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$. The up quark contributions are only required for Glashow-Iliopoulos-Maiani suppression and computationally the process is dominated by the top and charm quarks. Figure adapted from [9].

1.3.2 SM Branching Ratio

This section mostly follows [9] (more recent) and [13] (more readable). The effective Hamiltonian for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ (and $K^+ \rightarrow \pi^+ \nu \bar{\nu}$) is computed using the relevant Z^0 box and penguin diagrams (see Figure 1.4) and written in the SM as:

$$\mathcal{H}_{\text{eff}} = \frac{G_{\text{F}}}{\sqrt{2}} \frac{\alpha}{2\pi \sin^2 \Theta_{\text{W}}} V_{ts}^* V_{td} X(x_t) (\bar{s}d)_{V-A} (\bar{\nu}\nu)_{V-A} + h.c., \quad (1.38)$$

where $X(x_t)$ is a function dependent mostly on m_t and α_s that collects a lot of the complicated guts of the calculation. The dependence on V_{td} is plain and due to flavor-changing vertices in the diagrams in Figure 1.4. The decay is dominated by short-distance loops with top quark exchanges and the charm contribution is essentially negligible, removing some significant theoretical uncertainties relating to m_c and some associated scales. Employing the OPE approach, after some significant algebra we can use the above expression to produce:

$$A(K_L \rightarrow \pi^0 \nu \bar{\nu}) = \frac{1}{\sqrt{2}} [F(1 + \bar{\epsilon}) - F^*(1 - \bar{\epsilon})] \langle \pi^0 | (\bar{s}d)_{V-A} | K^0 \rangle (\nu \bar{\nu})_{V-A}, \quad (1.39)$$

where ϵ is the neutral Kaon mixing parameter and F is a complex function which looks like:

$$F = \frac{G_{\text{F}}}{\sqrt{2}} \frac{\alpha}{2\pi \sin^2 \Theta_{\text{W}}} V_{ts}^* V_{td} X(x_t), \quad (1.40)$$

for one neutrino flavor.

This formalism is somewhat intimidating, but the key feature is that the expression is broken into two pieces: one piece is a function of F and F^* that is calculable directly and a matrix element that is not. However, we may invoke a useful isospin relation, $\langle \pi^0 | (\bar{d}s)_{V-A} | \bar{K}^0 \rangle = \langle \pi^0 | (\bar{s}u)_{V-A} | K^+ \rangle$ and use a similar equation for the branching ratio of charged Ke3 and the *measurement* of that decay to re-write the

matrix element as a ratio of known quantities. Together with our expression for the effective Hamiltonian and the difference in the lifetimes between the K^+ and K_L^0 , we can use this information to derive a prediction for the branching ratio. The critical remaining piece is calculating the isospin breaking corrections. This calculation has recently been extended beyond leading order and we may now extract the matrix element with an accuracy of a few parts per mil [8]. The dominant errors on this calculation are due to Ke3 experimental errors.

In some senses though, the errors on those corrections are unimportant because they are far from dominant in the overall calculation. If we massage the expressions above, we find:

$$Br(K_L^0 \rightarrow \pi^0 \nu \bar{\nu}) \propto r_{K_L} Br(K^+ \rightarrow \pi^0 e^+ \nu) \times \frac{\alpha^2}{\sin^4 \Theta_W} \times \left[\frac{\text{Im} V_{ts}^* V_{td}}{|V_{us}|} \right]^2 \times X(m_t, \alpha_s)^2, \quad (1.41)$$

where r_{K_L} is the isospin breaking correction and Θ_W is the weak mixing angle. If we plug in all the appropriate values we find:

$$Br(K_L^0 \rightarrow \pi^0 \nu \bar{\nu}) = (2.49 \pm 0.39) \times 10^{-11}, \quad (1.42)$$

where the error budget is 25% due to X_t (a function of m_t and α_s), only 4% due to long-distance uncertainties related to the matrix element, and 71% due to uncertainties in the CKM parameters (see Table 1.1) [8].

1.4 Beyond the Standard Model

All of the parameters of the CKM matrix may be extracted from tree-level processes that, while if not as theoretically clean as $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$, are easier to measure. If the

only purpose of measuring $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ is to over-constrain the CKM matrix and provide a clean measurement of one of the parameters, it is fair to ask whether the experiment is worth the trouble. The answer is yes for several reasons:

- Because the decay is so theoretically clean, any deviation from the SM prediction is a clear signature of new physics.
- Since $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ is sensitive to loop processes and not the tree-level, shifts in the branching fraction are due to in the introduction of new particles in the loop that must be summed over. Any new theories must be consistent with the measurement, giving it extraordinary discriminating power when it comes to favoring or ruling out extensions to the SM.
- The sensitivity of the decay through loop-processes to new particles makes the measurement highly complementary to energy-frontier measurements which may be able to find signatures of new physics and still struggle with characterization.

This three piece approach has wide application. For example, suppose a new particle is discovered at the Tevatron [24] or the Large Hadron Collider (LHC) [25]. New models will have to be developed to explain it. Each model will have its own unique $X(\nu)$ function in analogue to the X function above, where $\nu = \{m_t, \alpha_s, \alpha_N, m_N\}$ with α_N and m_N the new coupling and mass. This function can be solved for with the branching ratio for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ and favored or disfavored. If X is generalized, the result can be compared to predictions for $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ and either favored or outright killed. With more new particles (and more masses and couplings), these constraints only get stronger. In this sense, $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ can really be thought of as a very precise way of measuring short distance loop functions that will constrain almost any new physics discovered.

Numerous papers on the impact of $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ on new physics scenarios exist. See [26], [27], [28], or [29] for a representative cross-section.

1.5 Guide to this Dissertation

This document is organized around the results of a counting experiment. In the final analysis, we observe some number of events with a prediction in hand for the number of background events. Signal analysis at this stage is simple. However, justification of the background estimate and the overall sensitivity of the search is not so easy.

We begin in Chapter 2 with a high-level overview of the experimental method employed by E391a in the search for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$. Most of the material in that Chapter is derived from work done by the whole E391a group, and especially from those students who analyzed the first data Run of E391a.

Chapter 3 also deals with material best attributed to the group as a whole as it discusses the overall mechanical design of the detector. However, some of the material in that Chapter, namely, discussions of the Collar Counter 02 (CC02) veto and the Beam-Anti (BA, or beam-stop veto) are of more direct importance insofar as the author's involvement. We constructed CC02 in Chicago and the author was heavily involved in the mechanical fabrication down to the level of fashioning pieces in the Chicago machine shop. The author was primarily responsible for the design of the BA and went back and forth between simple prototypes, simulation, advanced prototypes, more simulation, and final construction.

Chapter 4 describes the E391a Monte Carlo (MC) simulation in detail. The author wrote the first fully functional MC from the ground up. Although this MC was not ultimately adopted by the group, many of the techniques and insights gained from the process were extremely helpful when the author took over the task of updating the

Run I MC to Run II and validating it for the unique circumstances of the Run. The MC was critical for acceptance and sensitivity estimation, as well as for background prediction. As a consequence, the process of validation was crucial for the success of the experiment. In a sense, the extent to which the MC has been validated is the extent to which the final results can be trusted.

In Chapter 5 we present a method developed in E391a for estimating the photon veto inefficiency through neutral product decay modes only (in particular, through $K_L^0 \rightarrow 3\pi^0$). Although the success of the method is not critical to the final result for this data sample, it could well prove critical for any future attempts to measure the branching fraction for the $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ decay. Therefore, it is included as a full Chapter rather than as an Appendix.

Background processes and our strategies for defeating them are outlined in Chapter 6. The Run II analysis was largely conducted by the author and another student working closely together [30]. Our division of labor was such that the author was primarily responsible for “all things Kaon” and the other student was responsible for “all things neutron.” Even though there was some significant collaboration across these boundaries, that Chapter will focus more on the Kaon background estimation, which was the focus of the author.

Finally, in Chapter 7, we present the final results of the E391a Run II analysis and provide a discussion of lessons learned from the experiment and possible directions for the future.

Branching ratios used hereafter for calibration and normalization purposes are taken from [31].

CHAPTER 2

METHOD OF THE EXPERIMENT

2.1 Critical Summary

E391a operates by searching for a single $\pi^0 \rightarrow \gamma\gamma$ decay with transverse momentum (P_T) consistent with a $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ decay and inconsistent with any background processes. The critical innovation of the experiment is full 4π hermetic photon veto coverage and a tightly collimated beam to secure the P_T resolution and allow us to forgo any charged tracking apparatus. We use a single level event selection trigger with a rate of approximately 200 Hz (~ 400 triggers in a 2 second on, 2 second off beam spill). We do not pre-scale the number of accepted events. Photon identification is handled via a cesium-iodide (CsI) inorganic crystal scintillator array with single photon position resolution on the order of 1 cm and energy resolution (σ/μ) of $\sim 1\% \oplus 1\%/\sqrt{E}$, where E is measured in GeV.

2.2 Philosophy of the Measurement

The E391a collaboration is small for a particle physics experiment. There are eleven institutes representing five countries. Japan is the host country, but The United States, South Korea, Russia, and Taiwan all have provided critical contributions.

There were three long running periods in the E391a experiment. Run I took place between the beginning of February and the middle of July of 2004. Run II began in February of 2005 and finished in April of that year. Finally, Run III took place in the fall of 2005, from October through early December. This dissertation will present results from Run II, in which the author devoted the majority of his time in the preparation and data-taking phases. Results from Run I have been published [32]

where we set a new limit of 2.1×10^{-7} at the 90% C.L. Further details are available in [33] and [34].

The signal for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ is one π^0 plus missing transverse momentum (P_T) kinematically consistent with the decay. Attempting a full final state reconstruction by detecting the neutrinos is unfeasible due to the very low interaction probability of those particles. We chose to use only the two photon decay mode for the π^0 ; therefore our basic approach is to look for two photons in our calorimeter while requiring no signal in any veto detectors. We did this for two reasons: 1) to exploit the high branching fraction ($\sim 98.8\%$), and 2) in order to avoid the materials in our decay volume that would be required for charged particle tracking. Here and throughout, vetoes are made distinct from the calorimeter by purpose. Our calorimeter must have good position and energy resolution, but since the vetoes are used to reject events with any sort of hit outside the calorimeter, their performance requirements are somewhat looser. Position resolution is generally not critical, but because we apply tight veto thresholds, the low-energy (\sim MeV) behavior of those detectors should be well understood. Assuming two photon clusters in our calorimeter are produced by a single π^0 , we can reconstruct a vertex because we use a tightly collimated beam that restricts the (x, y) position. We can also compute a transverse momentum and our signal is defined by requiring the resulting P_T to sit inside a specific range and the reconstructed z -vertex to be in our primary decay volume.

We require a P_T greater than 0.120 GeV/c in order to keep the signal region clear of backgrounds from $\Lambda \rightarrow \pi^0 n$ decays, which have a maximum P_T of 0.109 GeV/c. We further required the P_T to be below 0.240 GeV/c since the maximum allowed P_T in $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ decays is 0.231 GeV/c. We also require the reconstructed Z vertex to have a value between 340 and 500 cm from the entrance to the detector itself. It is typical practice to place the origin of the coordinate system at the target (here,

in an experiment like this, so be careful to note the origin of the coordinate system here. This range was chosen to isolate the signal from background sources and will be justified in Chapter 6. (Our target sits 11 meters upstream of the origin.)

The general concept of the experiment is very simple. Prior to E391, the best limit for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ came from the KTeV collaboration. [10] The lessons of that measurement were the following:

- To suppress backgrounds from Kaons with decay products escaping the calorimeter, 4π steradian hermetic photon veto coverage is required.
- To minimize backgrounds from secondary beam particles interacting with detector materials, the beam “halo” (the tail of the radial distribution of beam density) must be highly suppressed and the decay region kept as “clear” as possible (no hodoscopes, wire chambers, or gas in the decay region).
- Because the decay region is to be kept clean, we cannot use charged particle tracking. Therefore, P_T resolution is quite important and we need as many constraints as possible on the decay vertex. This necessitates a very thin “pencil” beam. There is a trade-off here with the overall flux, but the general idea is to fix the decay vertex in $X - Y$ space as best as possible.

See Figure 2.1 for a diagrammatic illustration of these principles.

The challenge for E391a is not a complicated or ambiguous signal definition - it is a counting experiment. The problem is almost purely one of background rejection. Despite the care taken to address likely background sources, when striving for a measurement with a sensitivity high enough to search for a one part in one hundred billion level process, even rare background processes can be extremely problematic. It turns out that there are a number of ways to produce two clusters in our calorimeter

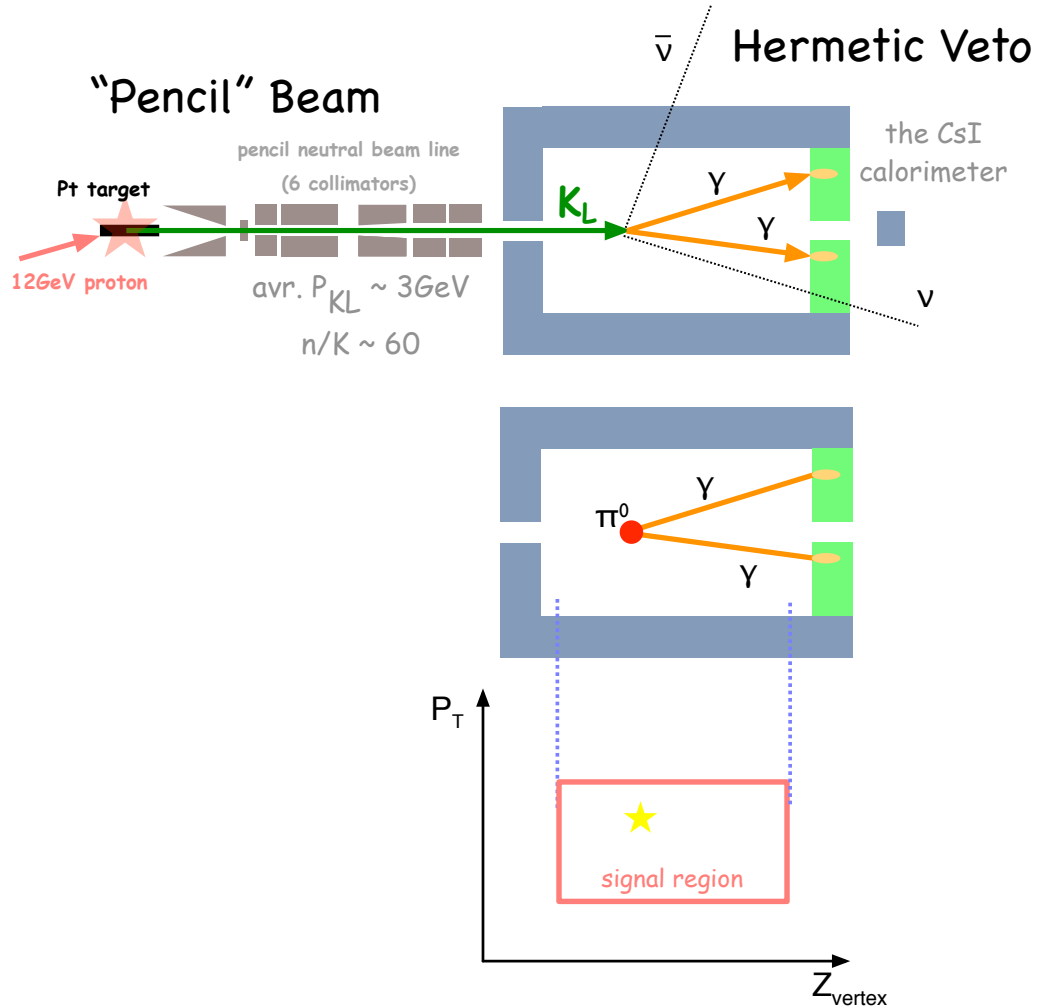


Figure 2.1: The general concept of the E391 experiment in diagram form. We search for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ by identifying a single $\pi^0 \rightarrow \gamma\gamma$ in our fiducial volume with transverse momentum consistent with the decay and inconsistent with backgrounds. We further require all veto detectors to be “quiet.” We do not detect the neutrinos from the decay, and the fiducial decay region is kept at a low vacuum level to minimize interactions with air (or other gas) molecules. Signal is defined by its location in (z, P_T) parameter space with kinematic cuts and photon vetoes imposed. Events in the space are considered signal candidates and events outside are not, making E391a a counting experiment.

kinematically consistent with a single π^0 through background processes. We need effective photon veto and kinematic selection criteria to remove these backgrounds.

This challenge is compounded by the fact that this analysis was done *blind* in order to minimize human bias. See, for example, [35] for a discussion of the blind analysis technique. There is a tendency in a rare process search to continue to tighten cuts until the data matches the experimenter's expectations. This sort of bias can make discovery of new physics impossible. Operating blind in this case means that we did not investigate the signal region itself while choosing event selection criteria. Instead we used a variety of methods to study the region of parameter space around our signal definition. Once we felt we understood those regions and could reproduce them, we used the tools developed in that process to predict a background level in the signal box. Significant deviation over that background prediction is to be taken as signal in a blind analysis, so great care must be taken to produce a truthful and precise background estimate.

2.3 KEK

E391a was conducted in the East Hall at the KEK (Kou-Enerugi-Kenkyujyo - literally, High Energy Research Area) 12 GeV proton-synchrotron in Tsukuba, Japan. See Figure 2.2 for a diagram of the counter hall. During the running periods for E391a, the East Hall hosted a number of beam tests at the PI0 test area immediately next to the K0 beam-line. (E391a conducted several beam tests in this area prior to the experimental runs.) Because the PI0 beam-line was well off-set from the K0 area and the two beam-lines were well shielded, cross-talk between experimental areas was completely negligible.

KO BEAM LINE LAYOUT IN EAST HALL AT KEK



Figure 2.2: The lay out of the East Hall at KEK. The E391a experimental area is the shaded region in the center of the figure and the PI0 beam-test area is below and to the left on the diagram.

2.3.1 Beam Line

The E391a beam-line is examined in detail in [36]. However, features critical to understanding this document will be discussed here.

The experimental method employed by E391a placed two stringent requirements on the beam. First, we required a “pencil” beam - i.e., a very narrow, well-collimated beam. Such a beam was necessary in order to calculate with sufficient certainty the transverse momentum of single π^0 's from the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ decay. Decays from a substantial beam halo would have smeared our P_T resolution to an unacceptable degree. Furthermore, halo particles interact with detector components and produce primary false veto signals directly, or secondary false veto signals from their interaction products. From peak-to-halo, our beam intensity dropped by almost five orders of magnitude, with the bulk of our beam within 1-2 cm of the nominal beam axis, as shown in Figure 2.3. We will see that even with a beam this sharply defined, interactions of halo particles with detector subsystems formed a formidable background.

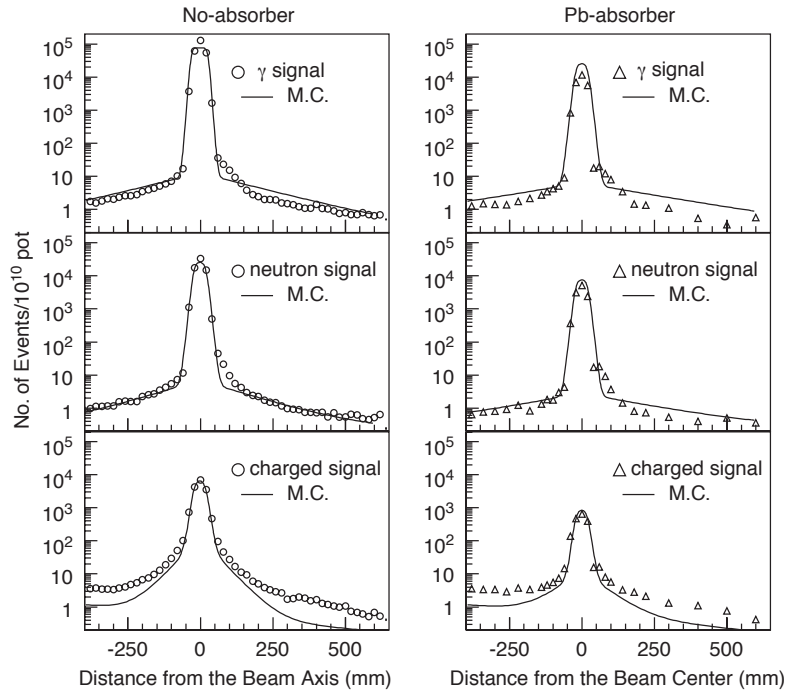


Figure 2.3: The K^0 beam-line peak-to-halo comparisons for different beam products, reproduced from [36].

Second, we required a very “clean” beam in the sense that we sought to maximize the fraction of our beam composed of actual K_L^0 ’s. Stray neutrons and photons, since they are not electrically charged, are difficult to remove from the beam but pose serious problems, by either creating accidental false veto events and driving down acceptance or by faking the signature for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$.

The beam line itself consisted of a set of six collimators, labeled C1 through C6 in order away from the target, a pair of sweeping magnets, and two in-beam absorbers (one lead (Pb) and one beryllium (Be)). These components are shown in Figure 2.4.

The principle components of the beam and their functions are (from [36]):

1. The primary proton beam is incident on a platinum target 60 mm thick ($0.68 \lambda_I, 20 X_0$) and 8 mm in diameter.
2. The beam-line is at a 4° angle with respect to the primary proton beam, pro-

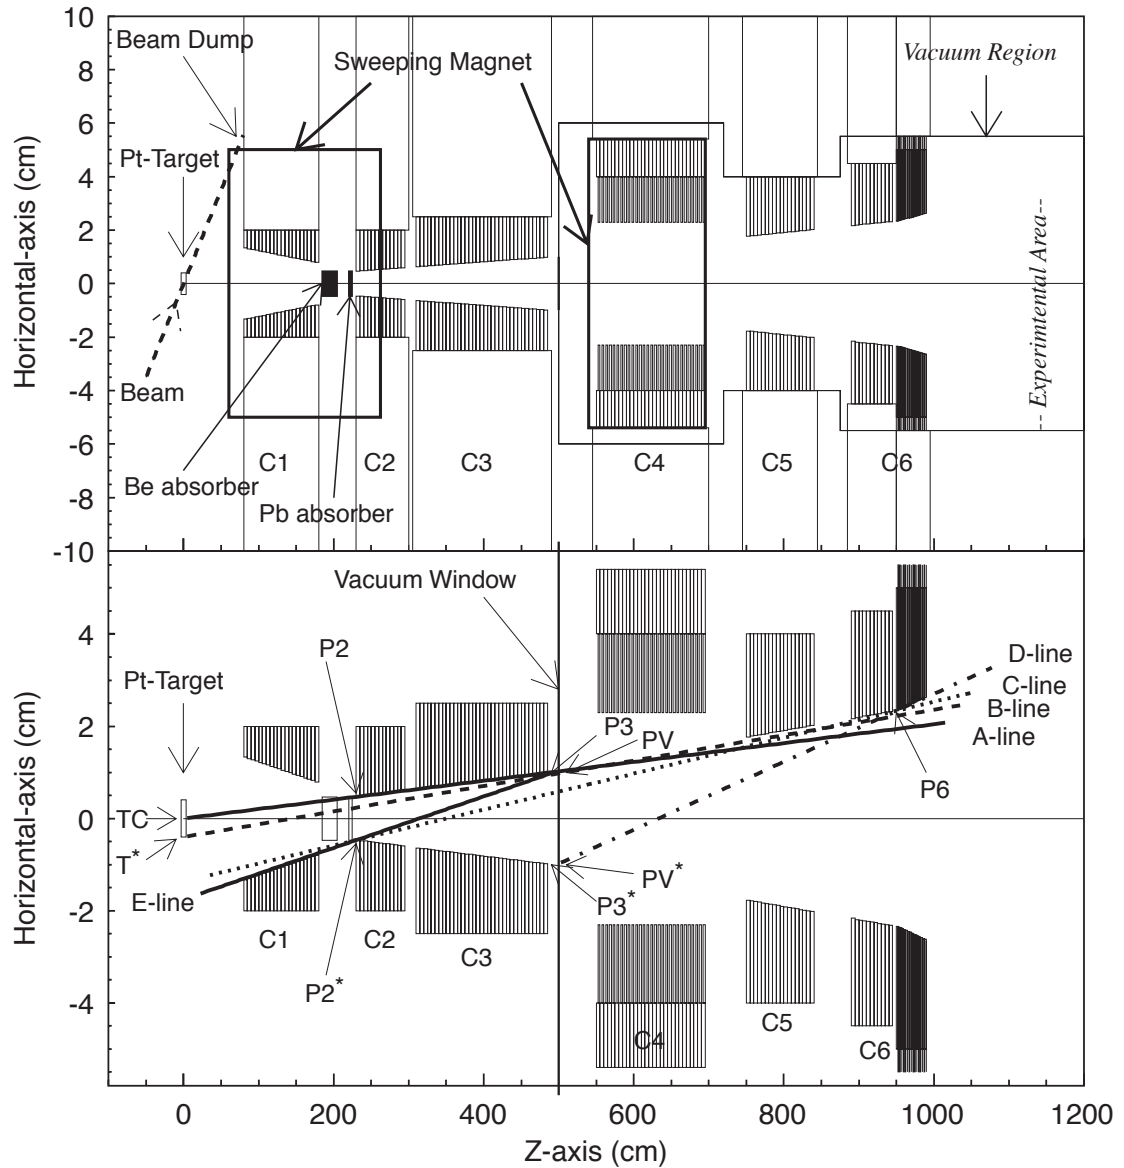


Figure 2.4: The K^0 beam-line for E391a, reproduced from [36]. The top of the figure shows how the geometrical sizes and positions of the components (note the different vertical and horizontal scales), and the bottom of the figure details the collimation scheme.

ducing Kaons with a mean momentum of 3.5 GeV/c at the target.

3. The beam-line is 10 m in length to reduce the hyperon and K_S content of the beam to negligible levels. With the full detector installed, the target was 11 m from the beginning of the detector and 14.4 m from the start of our fiducial volume.
4. Three collimators (C1, C2, and C3) are used to define the beam profile (half-cone angle of 2 mrad).
5. Two collimators (C5 and C6) are used to trim the beam halo. The most downstream collimator (C6) has an active component (plastic scintillator) to veto background events generated in collimator interactions (this veto turned out to be unnecessary).
6. Two movable absorbers are available to reduce photon and neutron components of the beam. During Run II physics data-taking, we employed both absorbers. One absorber was a 7 cm thick block of lead (high Z to reduce the photon content of the beam) and the other was a 30 cm block of beryllium (low Z to reduce the neutron component). Reducing the photon content of the beam was particularly important for the functioning of the very last veto, the Beam-Anti (discussed in Chapter 3). After analyzing our results from Run I (which was conducted with the lead absorber only), we decided to also use the beryllium absorber to reduce the neutron backgrounds in exchange for a small sacrifice in flux (the absorber also reduces the Kaon content of the beam, but it preferentially selects against neutrons).
7. Two dipole magnets are deployed to sweep charged particles from the beam.

8. One of the middle collimators (C4) contained thin Gd_2O_3 sheets to reduce the low energy thermal neutron ($E \sim kT \sim 1/40$ eV) content of the beam.
9. The air pressure in the second half of the beam-line was lowered to the single pascal level, and separated from the fiducial decay region by a thin membrane. (The fiducial decay region was held at 10^{-5} Pa during physics data taking.)

2.3.2 Operation of the Synchrotron

The total extraction cycle was four seconds long. Over the course of the first two seconds, protons were accelerated in the proton synchrotron (PS) to a kinetic energy of 12 GeV. During the next two seconds, the protons were extracted to the KEK East Counter Hall where they were incident on a platinum (Pt) production target at angle of four degrees with respect to the E391a neutral beam line. Typical proton intensities ranged from 2×10^{12} to 3×10^{12} protons per spill. A four degree extraction angle was chosen to minimize the neutron to kaon ratio since beam neutrons can, through several mechanisms, produce serious backgrounds for the flagship $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ decay mode. We achieved a n/K ratio of roughly 60 at the target. Further details on the operation of the synchrotron are available in [34].

2.3.3 Running Conditions

In order to use the cleanest possible data, we employed a number of beam quality cuts. These quality cuts included: Secondary Emission Counter (SEC) values, accelerator stability, online veto plot stability, and pedestal stability. We required the SEC count to be within a range around the empirical values observed during ideal beam conditions wide enough that data loss was only a few percent. Individual runs and spills taken when the accelerator failed were removed “by hand,” along with runs

that showed poor veto system behavior for various reasons in the online monitor plots (confirmed off-line). In total, we collected 1.4×10^{18} protons on target in the physics runs that passed all basic quality cuts during Run II.

2.4 Brief Survey of the Detector

Individual detector sub-systems will be discussed in more depth in Chapter 3, but some salient features will be discussed immediately. In particular, note that the origin of the coordinate system used throughout is displayed relative to other detector elements in Figure 2.5, while Figure 2.6 provides a sense of scale relative to human-sized figures for the detector.

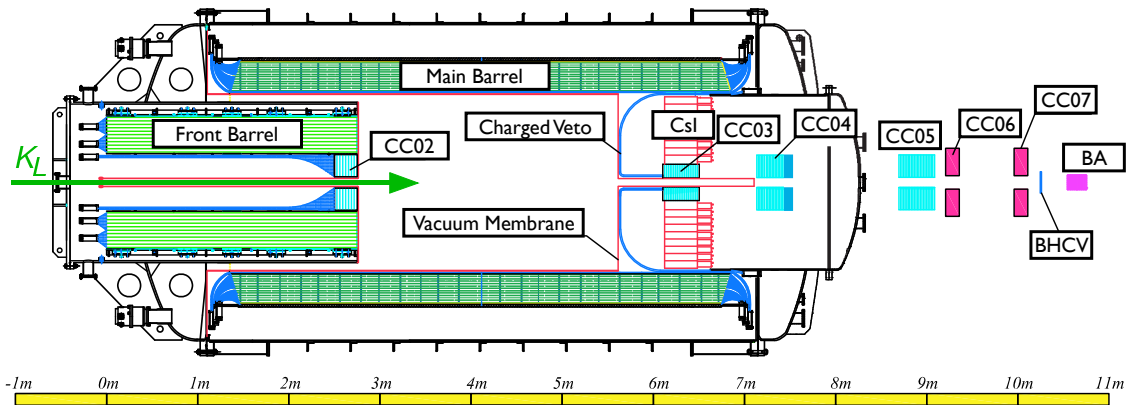


Figure 2.5: The E391a Detector shown in plan view with a numeric length scale. Note that our coordinate system identifies the origin at the beginning of the Front Barrel. The z -axis points downstream and the y -axis points vertically upwards (the coordinate system is "right-handed.")

The most important piece of the detector is a cesium iodide (CsI) inorganic crystal scintillator array. This crystal array forms our calorimeter and is used for event reconstruction according to methods described later in this Chapter.

The CsI array was placed at the downstream end of a lead-scintillator sandwich




Figure 2.6: The full E391a Detector, shown with human figures for scale. Not shown are the Collimator Vetoes installed inside the beam-line itself.

barrel veto (with planes parallel to the beam direction) called the Main Barrel (MB). In front of this large barrel veto we placed a smaller barrel veto of similar structure, the Front Barrel (FB). This formed a “double decay region.” The concept behind this structure was to capture $K_L^0 \rightarrow \pi^0\pi^0$ decays outside the Main Barrel where two photons from the subsequent $\pi^0 \rightarrow \gamma\gamma$ decays traveled downstream to the calorimeter and two photons escaped in the transverse direction. By placing an additional barrel veto in front of the Main Barrel, the acceptance for the two downstream photons from this event topology was reduced essentially to zero.

At the end of the beam-line we placed a Beam-Anti veto (BA) to reject $K_L^0 \rightarrow \pi^0\pi^0$ events where photons escaped down the beam pipe and $K_L^0 \rightarrow \pi^+\pi^-\pi^0$ events where both charged pions escaped down the beam pipe. Because the BA was exposed directly to the beam, we had to be sensitive to accidental losses due to interactions with beam neutrons.

Finally, the beam axis was surrounded with a set of collar vetoes (referred to usually as collar “counters”) to reject beam halo interactions, and very backwards and forward going daughter particles of Kaon decays. Some of these collar counters were located to cover any remaining uncovered escape phase space for Kaon decay products from the fiducial region and ensure 4π hermetic coverage once the calorimeter, barrel vetoes, and beam-anti were designed and in place.

2.5 DAQ Electronics

The E391a experiment uses photo-multiplier tubes (PMTs) to translate scintillation light in detectors into electrical pulses, recorded by analog-to-digital converters (ADC’s). ADC signals are fed into special NIM modules constructed for E391a called Amplifier-Discriminator (AD) modules. The AD’s bundle the analog sum of eight channels and send the sum along with the individual channel information and a logic signal for the timing, handled through a time-to-digital converter (TDC). The signals were transmitted via coaxial delay cables as long as 90 m to the NIM electronics station. Sum signals were transferred to the decision logic modules for trigger decisions using shorter, 30 m cables. The timing system employed a common start and individual stop scheme, with the timing window opened by the trigger signal and closed by the stop logic from the AD modules. In Run II, DAQ live-time was approximately 90%. Further details on the DAQ are available in [34] and [37].

2.6 First Level Event Selection - The Trigger

2.6.1 Physics Trigger

The physics trigger system employed just one decision level. Trigger decisions were made by summing the signal output from individual crystals into bundles of eight, forming “hardware clusters” (HWC). See Figure 2.7 for the lay-out of the crystals corresponding to the different HWC. Our physics trigger required two or more HWC’s, with an energy threshold of approximately 80 MeV per cluster. This threshold was chosen to discriminate against soft beam secondaries and was comfortably below the ultimate analysis level cuts on minimum photon energy (at 150 MeV). The physics trigger was not pre-scaled (every trigger request was accepted) with a rate of roughly 200 Hz (roughly 400 trigger requests per 2 second beam spill).




Figure 2.7: Hardware cluster assignments for the physics trigger.

We additionally enforced a series of “online” vetoes for the physics trigger in order

to avoid storing data that would fail even the loosest cursory cuts. See Table 2.6.1 for a summary of the online cuts.

Detector	Electronics Threshold (mV)	Energy Equivalent (MeV)
CC02	-48.6	~ 25
CC03	-29	~ 15
CC04	-57.1	~ 45
CC05	-34.3	~ 25
Main Barrel (Upstream)	-31.2	~ 15
Front Barrel	-100	~ 30
Inner Charged Veto	-28.9	~ 1
Outer Charged Veto	-25.7	~ 1

Table 2.1: A representative set of parameters for the online cuts in Run-II (they were occasionally changed in response to ongoing analysis and electronics problems).

2.6.2 Other Triggers

We also employed several calibration triggers, a pair of minimum bias triggers, and several accidental triggers. Most of these triggers were pre-scaled to various levels throughout the experiment. On average, the total trigger rate was roughly 500 Hz (1000 trigger requests per 2 second spill). For calibration we used a cosmic ray trigger based on energy deposition in opposing groups of channels in the main barrel (here, opposing implies separation by roughly 180 degrees around the barrel). Additionally, we used a muon trigger based on the coincidence of Collar Counters 02 and 04 (the detector geometry is described in detail in Chapter 3).

The minimum bias triggers were designed to allow us to study the effects of our online veto selections by mimicking the cluster requirements of the physics trigger at interesting thresholds, but not generally enforcing the online vetoes. One minimum bias trigger required only one or more HWC, while the other required two or more.

We used three different accidental triggers at E391a. The first was a trigger

based on an energy sum threshold in the Beam-Anti (BA). The second was based on an energy sum threshold in the final part of the collimator system (C6). The final and most important trigger was based on an energy threshold in a scintillator plane positioned adjacent to the target, the “target monitor trigger.” The target monitor trigger was effectively proportional to beam intensity and proved most effective in replicating data (as seen through MC/Data overlays). Therefore, throughout this document, the “accidental trigger” is assumed to mean the target monitor trigger. Unfortunately, an accidental trigger was not originally included in the design of the electronics and cabling. As a consequence, the timing of the accidental trigger was slightly out of synch with respect to other triggers. Because the beam intensity was relatively flat and this timing off-set was small, the effect of this was almost invisible for every detector except the BA (which was directly exposed to the beam and so sensitive to the beam microstructure). The consequences of this timing offset in the BA will be discussed in Chapter 4.

2.7 Event Reconstruction

The $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ signal in the E391a detector is based on classifying a $\pi^0 \rightarrow \gamma\gamma$ decay. When dealing with π^0 's, we generally only consider the decay $\pi^0 \rightarrow \gamma\gamma$, which has a branching fraction of almost 99%. Of course, we take into account the deviation in the branching ratio from unity where appropriate, but as a convention throughout this analysis, any decay of the π^0 is assumed to be in the $\gamma\gamma$ channel unless explicitly noted otherwise.

Our most important calibration modes are $K_L^0 \rightarrow \gamma\gamma$, $K_L^0 \rightarrow \pi^0\pi^0$, and $K_L^0 \rightarrow 3\pi^0$. As a consequence, our reconstruction routines are based on first finding and pairing photons in our calorimeter. Because we have no independent way to calculate

the decay vertex, we must assume the decay originated at $(0, 0)$ in the (x, y) plane. We then assume the parent mass and compute the z -vertex. In decays with multiple pions, we must examine all photon pairing combinations. Once we have selected the best pairing combination, we shift the vertex to sit on a line between the target and center of energy in the calorimeter and assign the z -vertex as the error-weighted average of the component pion vertices.

Because the particles we wish to detect are electrically neutral, we cannot constrain the decay vertex using charged particle tracking. Furthermore, in this experiment we do not have sufficient position, energy, and timing resolution in our calorimeter to compute the incident angle of a photon to better than roughly ten degrees. This angular resolution turns out to be poor enough to introduce a completely unacceptable uncertainty in the vertex, although it will still prove useful for some forms of background rejection.

Event reconstruction proceeds then as follows:

1. Events are classified by the number of clusters in the CsI array, and separated into streams.
2. Events from the two cluster stream are checked for single π^0 decays and $K_L^0 \rightarrow \gamma\gamma$ decays.
3. Events from the four cluster stream are checked for $K_L^0 \rightarrow \pi^0\pi^0$ events and those from the six cluster stream are checked for $K_L^0 \rightarrow 3\pi^0$ decays.
4. All events with successful reconstructions subject to an extremely loose set of cuts become candidates for analysis, and all others are rejected.

The other streams are generally not used, but the five cluster stream serves some important functions with regards to veto inefficiency estimation, discussed in Chapter

5. There is, of course, some acceptance loss due to the clustering approach taken, but we believe our MC models this loss well. Of course, we must also take care to include the effects of accidental clusters in our MC. We do this through overlays in some cases and through the use of a simple numerical factor in others. The probability of an accidental event shifting the number of clusters in an event is quite small though - $99.33 \pm 0.04_{Stat}\%$ of all events do not see a shift in the number of clusters due to accidental activity. Accidentals will be discussed in depth in Chapter 4.

2.7.1 Cluster Finding

The first step in particle reconstruction is cluster finding in the CsI array. This is a multi-step process:

- First we find all possible cluster seeds in the array, where a cluster seed is any crystal with 5 MeV or more energy deposited. We additionally compile a list of all crystals with energy over 1 MeV.
- Second, beginning with the seed crystal with the largest energy deposited, we add crystals to the cluster by including every neighbor from the list of possible cluster seeds. Note that neighbors are defined here to be crystals that share an edge. Crystals with more than 1 MeV that did not qualify as seeds are added to clusters where appropriate at this stage as well. We then check each neighbor's neighbor and so on.
- Once we have exhausted the list of possible crystals for the first cluster, we take the crystal from the list of remaining seed candidates with the highest energy and repeat the neighbor search process.
- When all the seeds have been used, we count the number of local maxima in

each cluster, where a local maximum occurs when a crystal has energy higher than all four of its neighbors. If any cluster has more than one local maximum, the event is rejected at this stage as a fusion event candidate. Signal acceptance loss for this clustering stage cut was estimated to be about 8%.

- Crystals with energy below the seed threshold and seeds without neighbors are classified as “single-hit” crystals, and not as clusters for cluster-counting purposes. Many single-hit crystals are found on the diagonals of good clusters and are an artifact of the clustering algorithm. Others are due to very soft photon interactions. The Moliere radius in CsI is about 3.5 cm and the width of most of the crystals is 7 cm. As a consequence, some electromagnetic showers are completely contained by single crystals. Finally, some single-hit crystals are due to other particles.
- Clusters are classified by their *size*, or number of crystals with energy deposition over 5 MeV, and by their *csize*, or number of crystals with energy deposition over 1 MeV. These sizes are used in the photon quality cuts for event selection.

See Figure 2.8 for a flow-chart representation of this process.

2.7.2 *Energy and Position Corrections*

Once all the clusters are formed, we perform a set of corrections to calculate the *incident* energy and position of the photon. Because of our shallow and broad CsI crystals, there can be significant shifts from the center-of-energy position and total deposited energy to the true incident position and energy.

The correction routines employed are described in detail in [38], but the essential details will be summarized here. Two tables were constructed using a GEANT4 MC

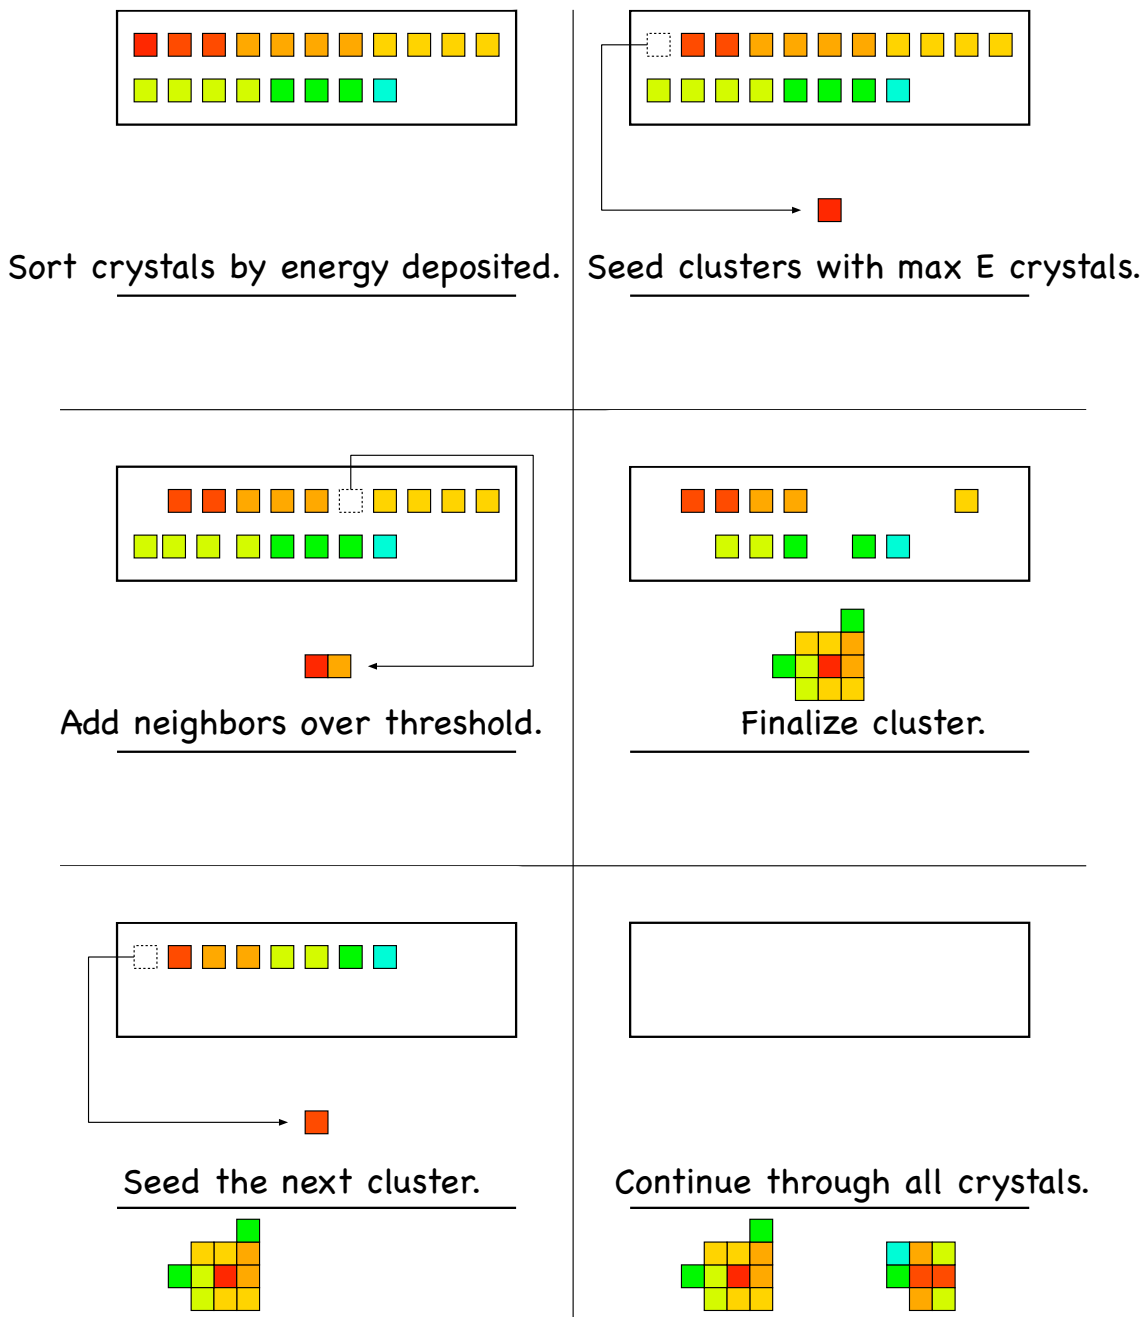


Figure 2.8: A flow-chart representation of the clustering algorithm. A cluster must contain at least two crystals. Note that clusters do not grow diagonally except through neighbors of neighbors paths.

[39] by injecting photons into an 11×11 crystal array at a variety of energies and angles. One table was designed to correct for energy “leakage” out of a cluster, either transversally or longitudinally. The other table was designed to correct the incident photon angle and position.

The correction is iterative. Once we have the clusters, we must begin event reconstruction, as described below. With a vertex in hand, we can return to the question of the correction, applying first the energy correction table, and the angular correction table on that output. Of course, with a new cluster energy and position, we must again reconstruct the event and check all possible photon pairing combinations. Typically, events converge within three iterations. See Figure 2.9 for an illustration of this process.

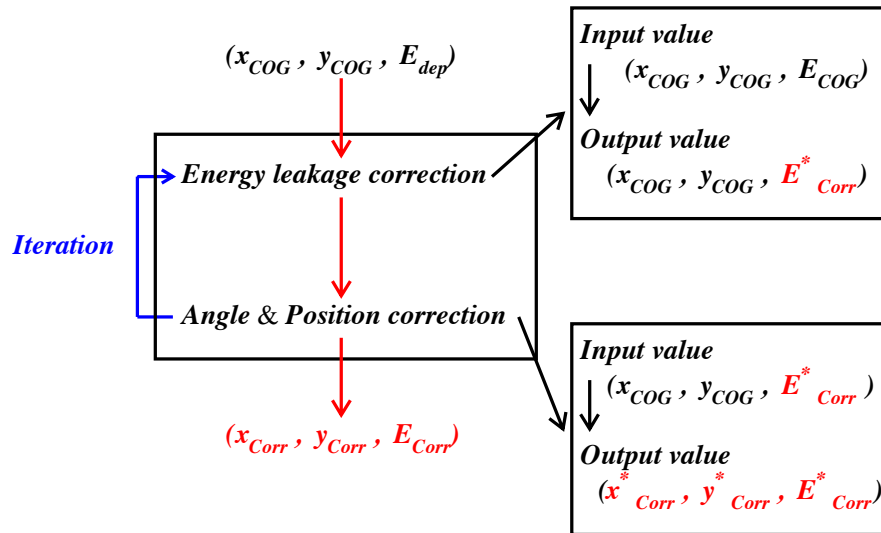


Figure 2.9: A diagrammatic representation of the cluster correction table application. Here, “COG” is an acronym for center-of-gravity (energy weighted mean), and “Cor” is an abbreviation for *corrected*.

The correction tables improved the photon hit position resolution (the gaussian width of the distribution of true position minus reconstructed position) from 2.6 cm

to 1.1 cm, the energy bias from 20.7 MeV to -5.7 MeV, and the pion vertex resolution from 8.9 cm to 6.6 cm, as displayed in Figure 2.10.

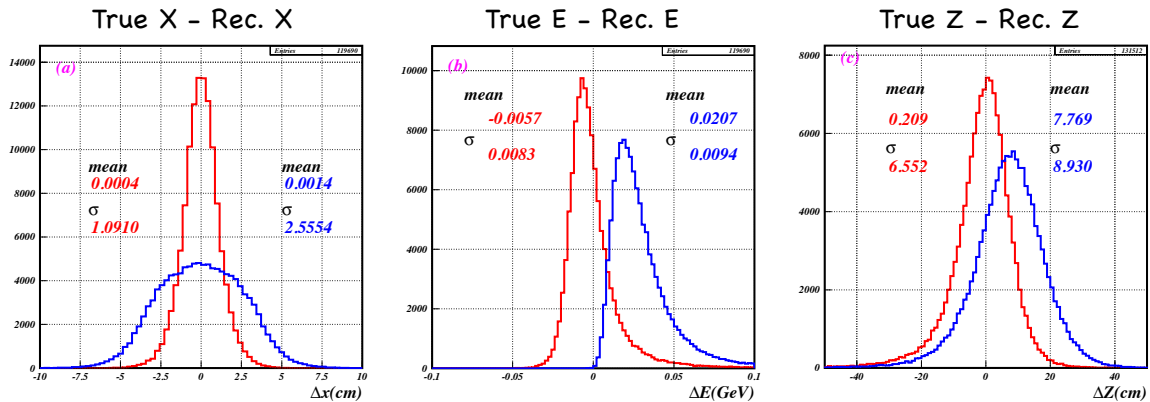


Figure 2.10: Before (blue) and after (red) the energy and position correction routines as applied to (a) the photon hit position (in x), (b) the photon energy, and (c) the pion z -vertex. “True - Rec.” means the known true MC value minus the reconstructed value. The associated means and widths (RMS) for the corrected clusters are displayed on the left side of each sub-figure, while the original center-of-gravity calculations are provided on the right side of each sub-figure.

2.7.3 Pion and Kaon Reconstruction

With cluster candidates in hand, we reconstruct π^0 's by requiring the four momenta of a cluster pairing to be consistent with a pion:

$$\begin{aligned}
 m_\pi^2 &= (p_1 + p_2)^2 \\
 &= p_1 p_1 + p_2 p_2 + 2 \times p_1 p_2 \\
 &= 2 \times (E_1 E_2 - \mathbf{p}_1 \cdot \mathbf{p}_2) \\
 &= 2E_1 E_2 (1 - \cos \theta)
 \end{aligned}
 \tag{2.1}$$

where p_i is the four momentum for the i -th photon (covariant notation is suppressed) and \mathbf{p}_i is the three momentum for the i -th photon. Because we are generally working

with large angles, we cannot make a small-angle approximation for the cosine function, however the “rule of thumb” $m^2 \sim E_1 E_2 \times \theta^2$ will be a useful relation to keep in mind for estimating the impact of slight changes to the energy, angle, or mass in the process of a reconstruction. See Figure 2.11 for an illustration of the reconstruction. Note that for $K_L^0 \rightarrow \gamma\gamma$ events, we simply replace the pion mass with the Kaon mass in this equation.




Figure 2.11: A diagrammatic representation of π^0 reconstruction according to Equation 2.1. An additional relation required for reconstruction is the law of cosines, here written as $r_{12}^2 = d_1^2 + d_2^2 - 2d_1d_2 \times \cos \theta$.

To reconstruct Kaons from multiple π^0 's, we first build pions from all the possible unique pairings of photons. Then, for each combination we calculate a *pairing* χ^2 according to:

$$\chi^2 = \sum_{i=1}^n \frac{(z - \bar{z})^2}{\sigma_i^2}, \quad (2.2)$$

$$\bar{z} = \frac{\sum_{i=1}^n z_i / \sigma_i^2}{\sum_{i=1}^n 1 / \sigma_i^2},$$

where $n = 2$ for $K_L^0 \rightarrow \pi^0\pi^0$ and 3 for $K_L^0 \rightarrow 3\pi^0$. Here, σ is the result of error

propagation on the uncertainty in the energy and position through Equation 2.1. See Figure 2.12 for a cartoon of this procedure as applied to $K_L^0 \rightarrow 3\pi^0$. We then order solutions by pairing χ^2 and keep the two combinations with the lowest χ^2 values. The lowest- χ^2 solution is our preferred solution. In order to discriminate against cases where two solutions are close together and the correct solution might have fluctuated to the higher value of χ^2 , we cut on the second solution and force its χ^2 value to be relatively large (kinematic cut points are detailed later).




Figure 2.12: A diagrammatic representation of $K_L^0 \rightarrow 3\pi^0$ reconstruction, where each pion is first reconstructed according to Figure 2.11.

CHAPTER 3

DETECTOR AND CALIBRATION

3.1 Critical Summary

In this Chapter we discuss the E391a detector in moderate detail with heavy use of references to more complete discussions when available. We discuss both the mechanical structure of the detector and outline the methods used to calibrate all the critical components. Because the author was responsible for the design of the Back-Anti (BA) and heavily involved in its fabrication, it is discussed in somewhat disproportionate depth to the other detector subsystems.

3.2 The Full Detector

The detector is discussed at length in [33] and [34], but a selection of the important details will also be covered here. Furthermore, the calibration of each detector subsystem is discussed in detail in [40], but some essential details will be discussed here for ease of understanding.

Photomultiplier (PMT) pedestals were recorded via a clock-trigger. The E391a ADC can automatically switch resolutions between a low (0.05 pC/channel) and high range (0.4 pC/channel) to accommodate incoming signals. We recorded pedestal signal for both modes and fit and subtract these pedestals dynamically for each data run prior to ADC-to-energy conversion, taking care to avoid distortions in the pedestal value due to noise or accidental activity. In general, because statistics for the high range were quite low, the high range pedestal was chosen as the peak channel for each run instead of the Gaussian mean. On-Off spill pedestal comparison showed no beam loading effects for the pedestals. Because the range was almost exclusively low

when using the automatic mode, we forced the first three spills of each run to be in the high range mode in order to have some pedestal information in those cases it was needed.

Most of the pedestals were stable in peak and width over the course of Run-II, with only a few exceptions. During Run-II there were two earthquakes, on February 16 and 23 of 2005, large enough to shut down the synchrotron and perturb the position of the detector. Due to the way the electronics were grounded, this introduced a drift into some channels as the multi-ton detector “settled” back into position. See, for example, Figure 3.1. Because the pedestal database was updated run by run though, this effect had negligible impact on our calibration.

3.3 CsI

We used a large array of Cesium-Iodide (CsI) crystals to construct our primary electromagnetic calorimeter. There were 576 crystals in the array of varying sizes and shapes. The bulk of the array was built from crystals that were $7\text{ cm} \times 7\text{ cm} \times 30\text{ cm}$ (approximately $16 X_0$, or “radiation lengths”) deep. There were 496 such crystals, borrowed from a previous experiment at KEK and referred to internally as the “KEK CsI.” The innermost ring of the array, consisting of 24 crystals, was built from crystals that were both narrower and longer at $5\text{ cm} \times 5\text{ cm} \times 50\text{ cm}$ (approximately $27 X_0$) deep. The inner crystals were borrowed from the spare stock used to construct the calorimeter for the KTeV experiment at Fermi National Accelerator Laboratory (FNAL) and so were referred to as the “KTeV CsI.” [41] The outer edge of the array consisted of a blend of “deformed” CsI cut and melted to fit inside the cylindrical drum built to hold the array and lead-scintillator sandwich counters, used to fill particularly vexing gaps around the outer edge. See Figure 3.2 for a schematic

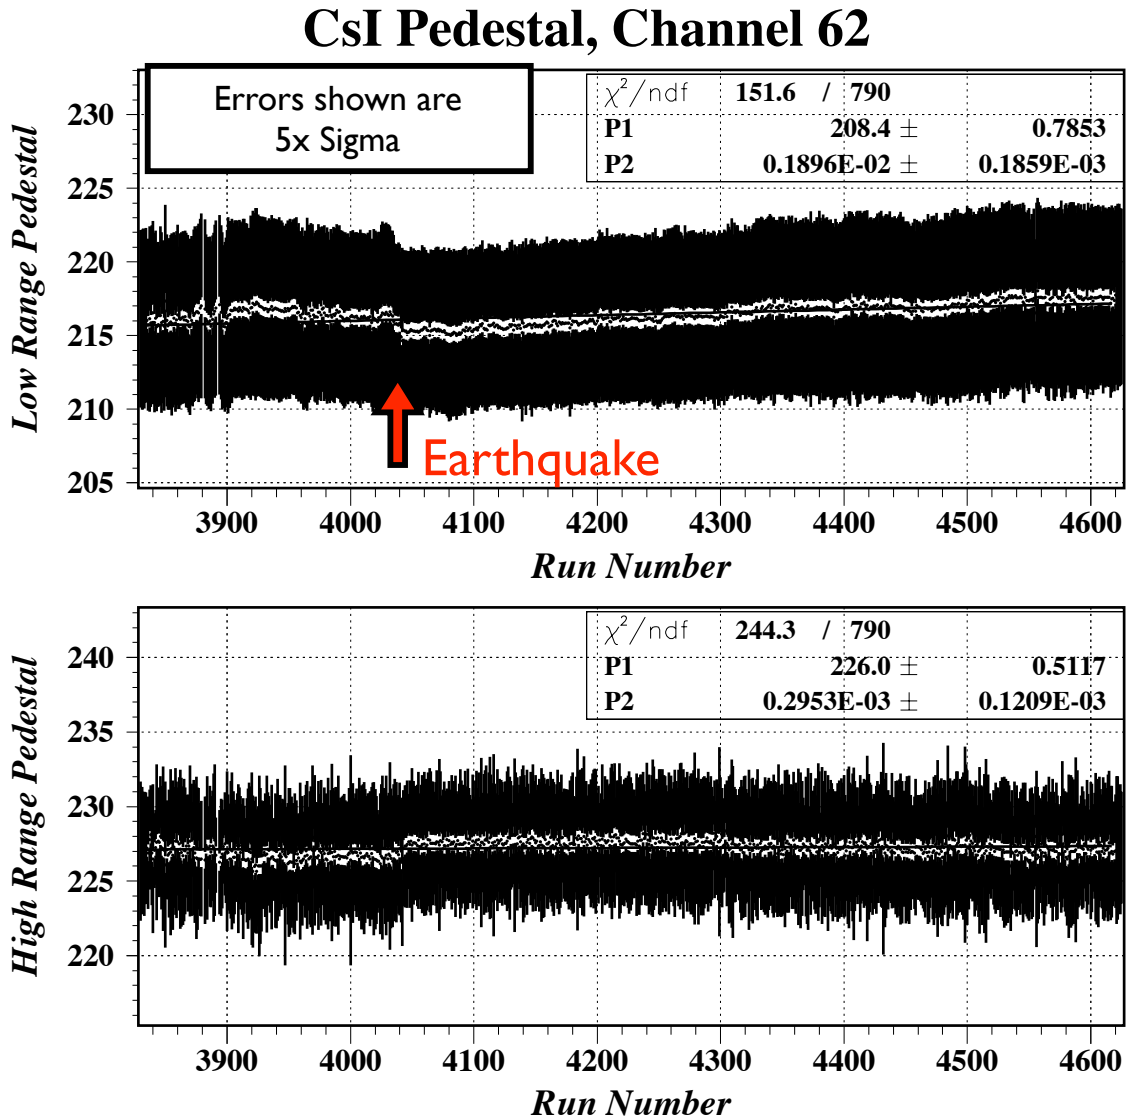


Figure 3.1: CsI pedestal stability before and after an Earthquake on February 16, 2005, shown as pedestal count (channel number) for CsI block 62 versus run number. (The earthquake on February 23, corresponding to about run 4120 above did not have as dramatic an impact on the electronics.) The low range (0.05 pc/Channel) is shown in the top half and the high range (04 pc/Channel) is shown in the lower half. The error bars displayed are five times the gaussian width of the fit (for the low range) or five times the RMS (for the high range). The fit to the pedestal evolution is linear: Pedestal = $P_0 + P_1 \times \text{Run Number}$.

of the deformed CsI and Sandwich Counters.

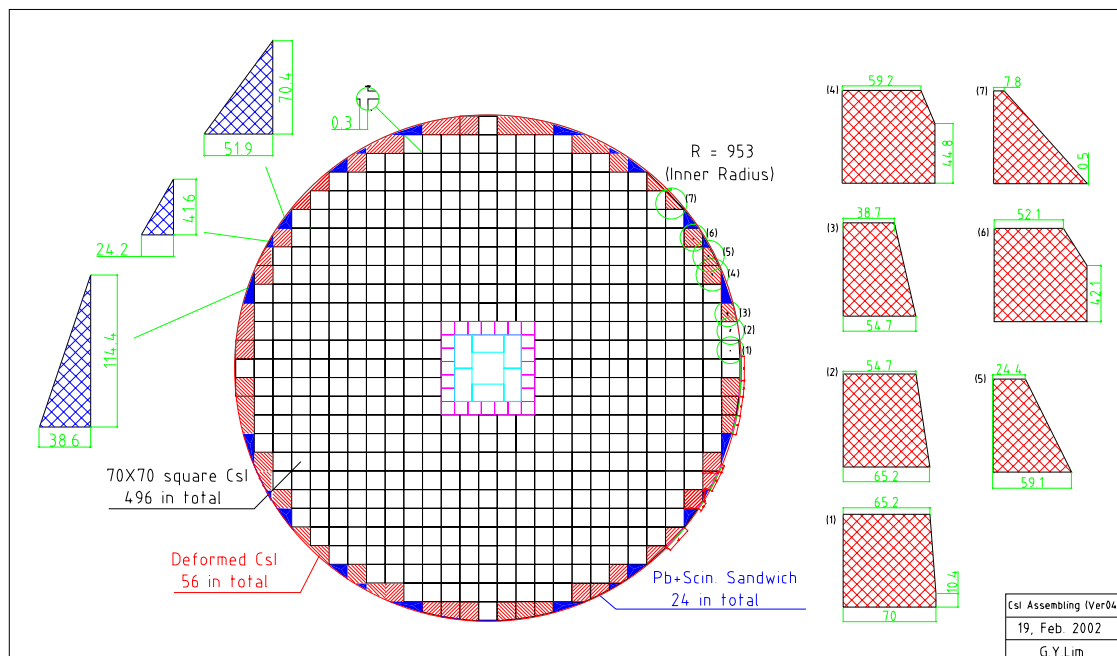


Figure 3.2: The components of the outer edge of the CsI array. The Sandwich Counter modules are the triangular blocks. They were built from lead and plastic scintillator plates, oriented parallel to the beam. There were a total of 24 such counters, grouped into eight groups of three for readout purposes. The deformed CsI blocks were each read-out individually and used in photon clustering and as veto counters in the exact same way as the rest of the CsI.

The KEK CsI crystals were individually wrapped in 100 micron thick teflon sheets and 20 micron thick aluminized mylar sheets for optical isolation. The average light yield for the KEK CsI was 15 photo-electrons per MeV. Attached to each crystal was a 2-inch Hamamatsu R4275-02 PMT, optically coupled using a 3 mm silicone sheet and a UV transmission filter designed to remove the slow component of the scintillation light emitted by the CsI. Because the PMTs were operated in vacuum, special care was taken to design a cooling and temperature stability system. Details of this system are reported elsewhere [34], [42].

The KTeV CsI crystals were wrapped using 13 micron thick mylar sheets, with

reflective coating placed to smooth the light yield across each crystal. Each crystal was attached to a 1.5 inch Hamamatsu R580-UV PMT, optically coupled with a 5 mm silicone sheet and UV filter. See Figure 3.3 for a schematic detailing the crystal-PMT attachments employed for the KEK and KTeV CsI.

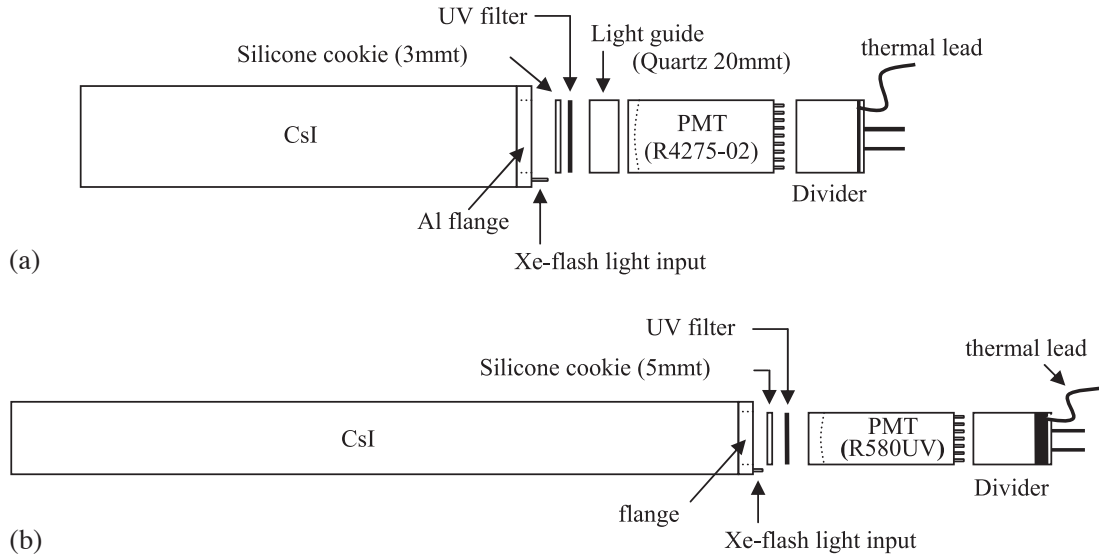


Figure 3.3: The CsI-PMT attachment scheme employed for the KEK CsI (a) and the KTeV CsI (b).

In order to fill the more difficult shaped gaps and to provide support, we additionally employed lead-scintillator “Sandwich Counters” in the CsI array around the edges. These counters were not used in clustering, but were used in veto. Each module was built using alternating, vertically stacked (parallel to the beam) layers of 1 mm thick lead and 5 mm thick scintillator, read out using wavelength-shifting fibers glued into groves in the scintillator. The fibers were silverized on the outer end and read-out terminated in 1.125 inch Hamamatsu H1398 PMTs. Light yields ranged from 10 to 20 photoelectrons per MeV, dependent on the module type. See Figure 3.4 and [34] for additional information.

The gain of the CsI crystals was monitored on a spill-by-spill basis using a Xenon




Figure 3.4: The sandwich counter module design. See Figure 3.2 for placement of the different types in the array.

flasher monitoring system. See references [34] and [42] for detailed information on this system. See Figure 3.5 for an illustration of Xenon data used to track PMT gain values and Figure 3.6 for a comparison of on and off-spill Xenon flasher data.

3.3.1 CsI Calibration

Prior to installation, 25 KEK CsI crystals were tested extensively with the KEK 3 GeV electron-hadron beam. Additionally, each crystal was installed in a cosmic test bench and analyzed for response uniformity and to check the PMT attachments. See [42] for a detailed discussion. During this beam test, we measured a resolution of $\sim 1\% + 1\%/\sqrt{E}$ (see Figure 3.7).

Once the CsI was installed, we additionally performed three other forms of calibration. First, we used cosmic rays in an attempt to set a baseline absolute energy scale and refine relative timing. Second, we used $K_L^0 \rightarrow 3\pi^0$ decays, along with a constrained kinematic fit, to refine the relative energy scale between crystals. Finally, we also used a special “Al-Target Run” to cross-check both the absolute scale and relative calibration between crystals. In this Target Run, we inserted a thin (5 mm) aluminum target in the beam directly in front of CC02 (discussed later). This target forced reactions like $n \rightarrow \pi^0 + X$ and provided a known vertex for $\pi^0 \rightarrow \gamma\gamma$ decays. We additionally used a xenon flasher to monitor PMT gain drift and kept detailed

CsI Xe/Cosmic Day-by-Day Comp, Chan 19

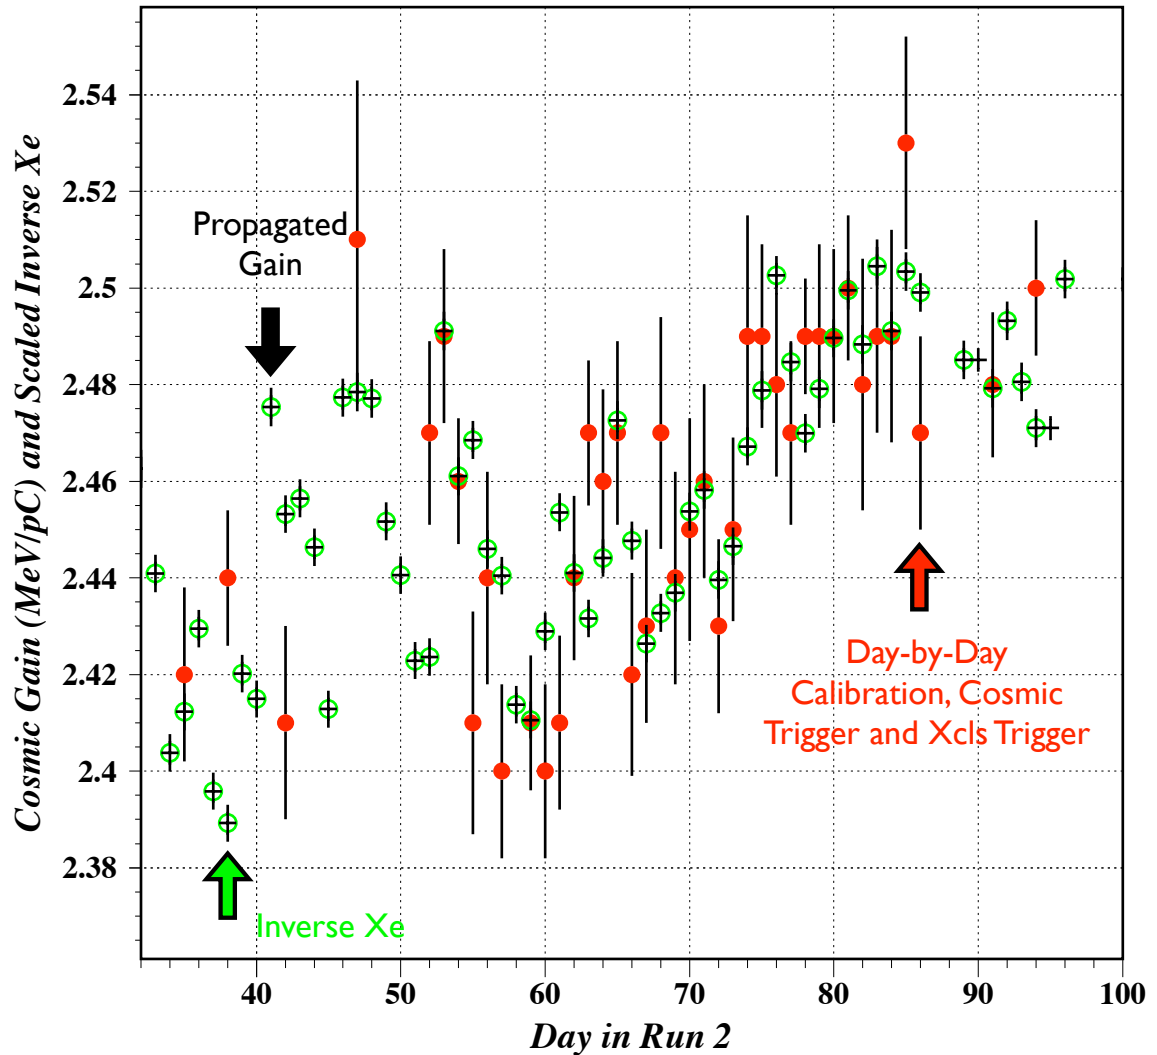


Figure 3.5: CsI Xenon Monitor propagation. The drift in the gain is likely due to PMT fatigue and Xe lamp dimming. Uncertainty about this drift helped to motivate the $K_L^0 \rightarrow 3\pi^0$ kinematic fit calibration described below. “Day-by-day” cosmic ray calibration employed all the cosmic tracks found on the given calendar day during Run II with both the dedicated cosmic trigger (using the Main Barrel) and the X_{cls} , or X-cluster trigger (usually four or more hardware clusters).

On/Off-Spill Ratios for Day-Averaged Xe Signal in Run II




Figure 3.6: CsI On/Off Spill Xenon. The χ^2 values from the individual channel fits are used as errors, but are not shown in the plot. As the boxed parameter indicates, there was very little in the way of beam-loading for the CsI, and the gain values computed using off-spill cosmic rays were compatible with gain values computed from on-spill $K_L^0 \rightarrow 3\pi^0$ decays.




Figure 3.7: The CsI energy resolution as measured by a 5x5 block array. The fit is $\sigma = P_1 \times \sqrt{E} + P_2 \times E$.

track of the temperature of the CsI and counter hall in order to enforce gain stability. Some additional comments on these features follow.

3.3.2 *Cosmic Calibration*

During Run II, we employed a cosmic-ray trigger that took the coincidence of opposite modules of the Main Barrel (the Main Barrel is discussed below). We searched that sample of events for cases where we could identify a track in the CsI. Then we fit a path through the crystals in (x, y) space using a least-squares fit and rejected all tracks shorter than 70 cm in length or with poor fit χ^2 values. See Figure 3.8. Deposited charge was normalized to path-length through the crystal (with ambiguity in the path-length on the z-axis being uncorrectable) and we assumed a minimum ionizing energy loss of 5.63 MeV/cm. We connected tracks passing through CC03 when possible. This method performed most poorly near the edges of the calorimeter in the x -direction.

One positive aspect of the cosmic calibration routines was our ability to take cosmic data off-spill and during accelerator down-time. Figure 3.9 shows the stability of the gain over time for one channel. Most of the timing calibration work was done with cosmic rays [40].

3.3.3 $K_L^0 \rightarrow 3\pi^0$ Calibration

We used a kinematic fit to constrain the relative energy between clusters in well-reconstructed $K_L^0 \rightarrow 3\pi^0$ decays [43]. See Figures 3.10 and 3.11 for examples of clean events and clusters as seen in our event display. Each photon cluster in a reconstructed $K_L^0 \rightarrow 3\pi^0$ decay provides three measured variables as calculated by our clustering routine - the (x, y) position and energy, E . We have three unknowns to solve for -

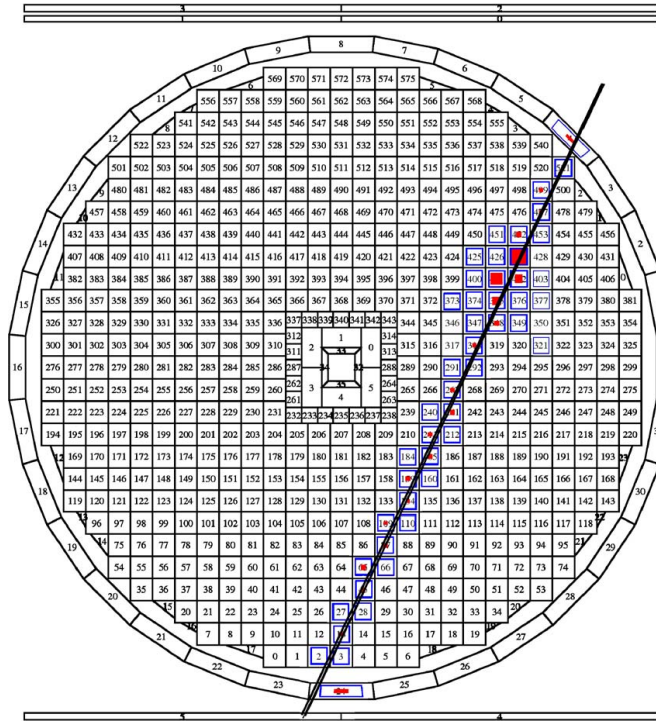


Figure 3.8: A cosmic ray track in the CsI array.

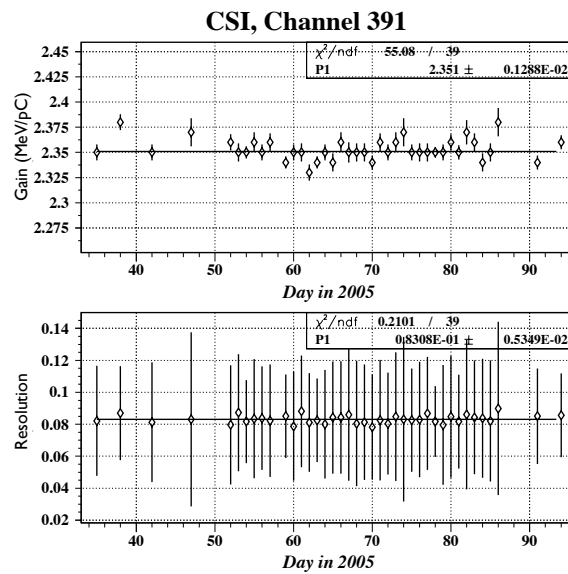


Figure 3.9: CsI gain stability as computed with cosmic rays.

the vertex of the Kaon decay, and six constraints:

- $M_{6\gamma} = M_{K_L}$
- $M_{\gamma\gamma} = M_{\pi^0} (\times 3)$
- $\sum_i (x_i \times E_i) = v_x \times \sum_i E_i$ (The Kaon vertex is at the corrected center of energy.)
- $\sum_i (y_i \times E_i) = v_y \times \sum_i E_i$ (The Kaon vertex is at the corrected center of energy.)

This system can be solved by least-squares with Lagrange multipliers with three degrees of freedom, shifting the energy and positions of the clusters to best satisfy the constraints.

For calibration we can also assume the energy of one cluster is unknown, leaving us with two degrees of freedom. We can then solve for the energy of that photon and use the ratio between the measured and calculated energies to shift the gain of the CsI crystals (we only shift the gain for each crystal in a cluster using an energy weighted formula that makes smaller shifts for crystals less involved in a given cluster). Since we have six clusters in an event, we can apply this correction to all six photons.

This calibration method can only adjust relative gains, and cannot be used to set the absolute energy scale. However, the average value of the Kaon decay vertex can be monitored with this method, and that value is sensitive to the absolute scale.

3.3.4 *Al-Target Run*

We suspended a 5 mm thick plate of aluminum 6.5 cm downstream from the end of CC02 (281.5 cm downstream from the beginning of the detector and 333.3 cm upstream from the face of the CsI array) during a special series of runs at the end of Run II. It was not possible to quickly move the target in and out of position to allow for periodic “ π^0 -Runs” as we called them because the apparatus used to suspend the

Run 3950 Spill 11 Event 356
 Tue Feb 8 06:52:00 2005
 -- Trig = 0x0011 : OnSpill, Nclus.,
 -- Nclus=6

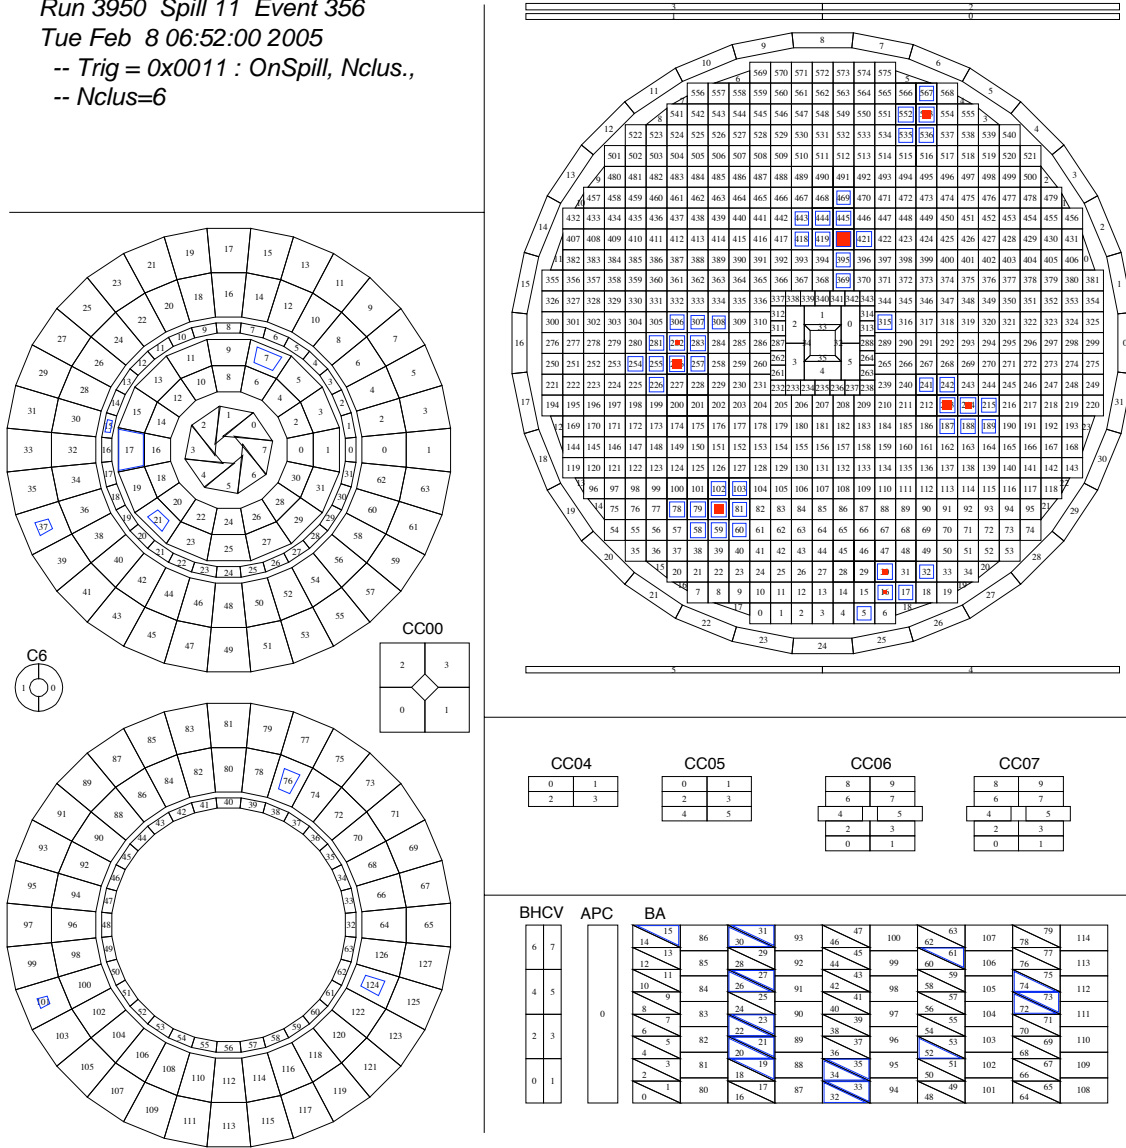


Figure 3.10: A clean $K_L^0 \rightarrow 3\pi^0$ event from Run II data that passed all cuts. The CsI array is in the upper right hand corner. On the left-top are the upstream channels of the Main Barrel (inner and outer), the Front Barrel (inner and outer layers, with CC02 inside the inner layers), and on the left-bottom are the Main Barrel downstream channels (inner and outer layers, with the Barrel Charged veto as well). Note that there is some soft activity in the BA (bottom-right hand corner) scintillator layers (separated by diagonal lines), but no activity in the quartz layers (separated by horizontal lines) - a likely candidate for a low-energy neutron interaction either in-time or slightly early. The channel marked “APC” was reserved for a detector element not employed in Run II.

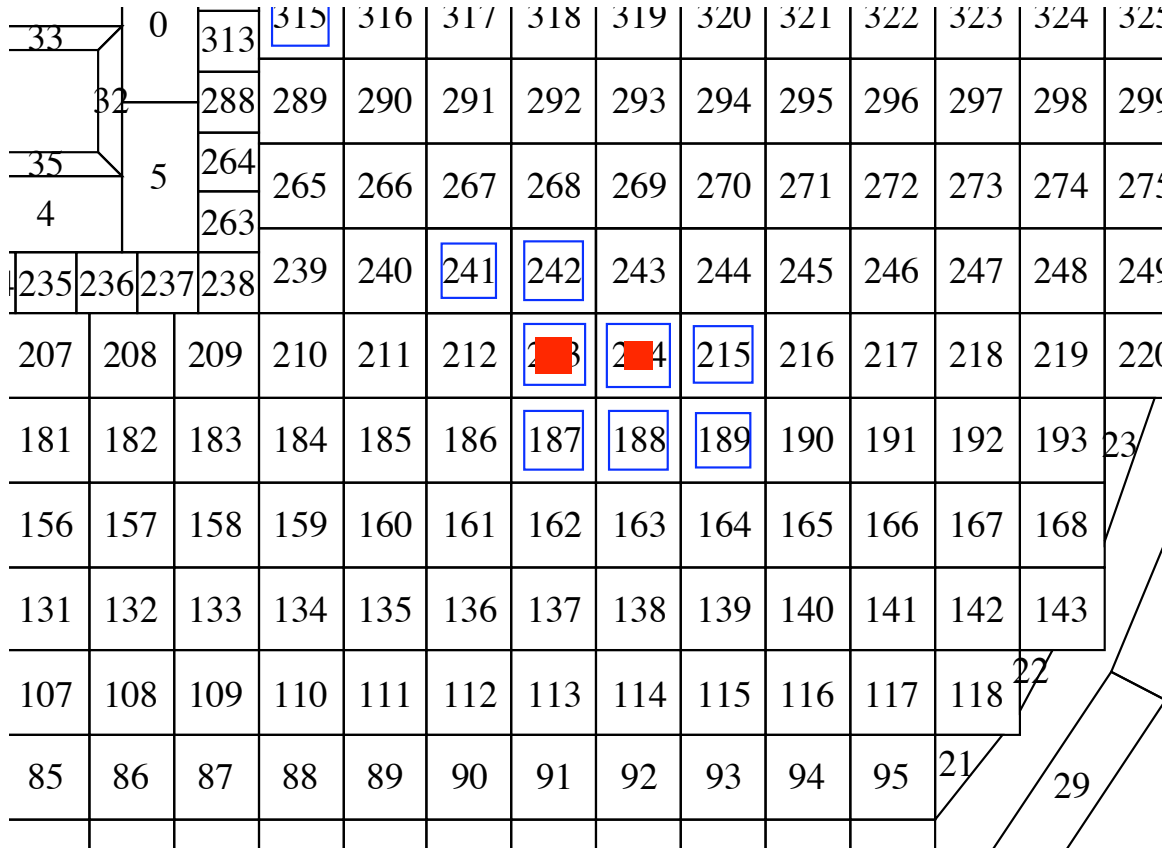


Figure 3.11: An individual cluster from the event shown in Figure 3.10 with $E = 395$ MeV and $(x, y) = (43.5, -20.75)$. For reference, the center of the most energetic block has (x, y) coordinates of $(42.0, -21.0)$, with x increasing towards the right hand side of the figure and y increasing with the vertical. Recall that each of the “KEK CsI” blocks has face dimensions of 7×7 cm².

plate was installed as an afterthought during Run I. As a consequence, we had to break vacuum in order to install the plate. (It took roughly two weeks to pump down to 10^{-5} Pa.)

We performed the Target Runs in order to provide a point with known vertex for π^0 production. In our event reconstruction process, we assume the π^0 mass. If there is a mistake in our overall energy scale, this will systematically shift the z -vertex we reconstruct. We would be blind to a shift of this nature and given the non-linear distribution of the decay vertex, this would lead to a systematic error on our acceptance and flux.

In order to correct for this error, we used the aluminum target and performed π^0 reconstruction using a fixed vertex rather than a fixed mass. Errors in the energy calibration of the CsI would now manifest themselves as a shift in the π^0 mass that we could correct by updating our calibration constants in an iterative fashion. Furthermore, the target provided a source of $\eta \rightarrow \gamma\gamma$ events that gave us a second point (the η mass) to check our calibration with. The Target Runs proved useful not only for calibration, but also for background estimation and background cross-checks, as will be detailed in Chapter 6.

3.4 Main Barrel

The Main Barrel (MB) was a lead-scintillator sandwich veto designed to reject Kaon backgrounds containing spare photons that missed the CsI calorimeter in the transverse direction. In total, the MB was $13.5 X_0$ thick. It surrounded the decay region, overlapping with the end of the Front Barrel and the beginning of the CsI array. See Figure 3.12 for a schematic of how the MB fit inside its container vessel and Figures 2.5 and 2.6 for information about placement relative to the other detector subcompo-

nents. Due to its sheer size and the area it covered, the MB was the most important subcomponent after the CsI itself. The MB was 5.5 m long on the beam axis, with an inner diameter of 2 m and an outer diameter of 2.76 m.

The MB was built from 32 trapezoidal modules, arranged around the beam axis in such a fashion as to present no “holes” or “cracks” to rays drawn from the beam axis outwards. See Figure 3.13 for a diagram of a single module, viewed from the end. Each module was read-out by four PMTs - two from each end. Internally, each module was divided into two pieces - an inner and outer piece - with different lead-scintillator layering. The inner and outer segments were read-out by different PMTs (one at each end of the segment). The inner module was built from fifteen pairs of alternating 1 mm lead and 5 mm scintillator layers. The outer module was built from thirty pairs of alternating 2 mm lead and 5 mm lead layers. See [33] and [34] for further details on construction and mounting for the MB.

3.4.1 Main Barrel Calibration

The MB was calibrated primarily using cosmic rays. Both the gain of each part of each module and all the timing information were computed by applying coincidence triggers to opposing modules (actually, opposing group sums.) Because high statistics were required for this, especially for the modules mounted on the side of the MB, we took long cosmic runs whenever the beam was not being delivered as well as during off-spill times. Additionally, long cosmic runs were held before and after each of the three main Runs for E391a.




Figure 3.12: The Main Barrel and vacuum tank.




Figure 3.13: A single Main Barrel module.

3.4.2 Barrel Charged Veto

Because of the high probability of a Kaon decay producing a charged particle, we placed a charged veto system, the Barrel Charged Veto (BCV) within the inner radius of the MB. Like the MB, the BCV was constructed in 32 segments, with placement offset just slightly from direct alignment with each corresponding MB module. Each panel of the BCV was made from two 5 mm thick pieces of scintillator, each 550 cm long. Wavelength-shifting fibers were glued into grooves between each pair of plates with optical cement and read into one inch Hamamatsu R7899-EGP PMTs, with one PMT placed at both the up and downstream ends of each module of the BCV.

3.5 Front Barrel

The Front Barrel (FB) was a lead-scintillator sandwich detector similar in structure and purpose to the Main Barrel. It was located upstream of the main decay volume, forming the “double-decay region” upstream of Collar Counter 02 (CC02). The FB was $17.5 X_0$ thick transverse to the beam and divided into sixteen modules, read-out from the upstream end using two PMT’s (one inner-layer bundle and one outer-layer bundle) for each module, with the downstream ends of the optical fibers covered with aluminized mylar to minimize light loss.

The FB was calibrated using tracks built off a trigger in the MB, with coincidence with CC02 required. Because we could only collect calibration data from the downstream end of the barrel, energy deposition in the barrel was defined using that normalization. Attenuation was handled using values measured prior to installation.

3.6 Charged Veto

The most common decay modes for K_L^0 and the bulk of its summed branching fractions include charged particles (e.g., $K_L^0 \rightarrow \pi^\pm e^\mp \nu$, $K_L^0 \rightarrow \pi^\pm \mu^\mp \nu$, and $K_L^0 \rightarrow \pi^+ \pi^- \pi^0$). Therefore, it was important our calorimeter be protected by a charged veto to protect against backgrounds associated with charged decays by K_L^0 . Our main Charged Veto (CV) consisted of two basic subcomponents, distinguished by read-out method. The larger portion was defined as the “Outer Charged Veto,” or OCV or simply CV, and the smaller, inner portion surrounding the beam was the “Inner Charged Veto,” or ICV.

The OCV consisted of 32 panels, placed in overlapping positions and fanning away from the beam. The panels extended as far as 50 cm from the face of the CsI before bending to read-out through the far edges of the CsI array. The ICV consisted of four panels placed inside the third collar counter veto, CC03, and read out through the back of the array. Each of the CV panels was read into a two inch Hamamatsu R329 PMT. See Figure 3.14 for details on panel configuration. Because of its “bent” shape, the light yield from the OCV was a non-linear function that we had to model in our simulation. Details of the light yield are discussed in [34].

3.7 The Back-Anti (BA)

Chicago was heavily involved in the design and fabrication of the Back-Anti (also called the Beam-Anti, or BA). The prototype model employed in our engineering run was entirely designed and built in Chicago, with most of the superstructure actually built by students in the Chicago Student Machine Shop. The Chicago group designed the final version used in Runs I and II and the author and J. Nix performed the bulk of the fabrication on-site at KEK. A different version of the BA was used for Run III




Figure 3.14: The Outer Charged Veto (OCV).

with an eye towards prototype testing for the next generation $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ experiment and to leverage lessons learned during the first two Runs. A side view of the BA used for Run II can be seen in Figure 3.15.

On February 21st and 22nd, 2004, the BA was tested at the PI0 area at KEK during the T552 beam test (see Figure 2.2). Two different triggers were employed in that test. The first was an “electron” trigger which utilized a gas Cerenkov counter to identify high speed electrons. The second was a “hadron” trigger that required a hodoscope hit and a null signal from the Cerenkov counter. The hadron trigger could identify not just pions and protons, but also muons. The momentum resolution of the test beam was quite good with a worst-case resolution for higher energy particles of roughly three percent. See Figures , 3.16, and 3.17 for the results of the test.

The BA, along with the Beam-Hole Charged Veto (discussed in the next section) was calibrated using muons produced by closing the beam shutter. “Muon-runs” were used to calibrate the Collar Counters as well (discussed below), but different triggers were employed for the different detectors. Because there were no other detectors that could be used to create a coincidence with the BA, the BA-BHCV muon trigger was




Figure 3.15: The Beam-Hole Charged Veto (BHCV, left) and Beam-Anti (BA, right). The channel numbering scheme is displayed above, but the figure is not precisely to scale. For the BA, note that the scintillator layers are separated by vertical lines and the quartz layers are separated into vertically stacked blocks to provide transverse segmentation.




Figure 3.16: The BA scintillator energy resolution as measured at the T552 beam test. The fitting function is $\sigma/\mu = P_1/\sqrt{E} + P_2$.




Figure 3.17: The BA scintillator energy linearity as measured at the T552 beam test. The fitting function is $E_{Dep} = P_1 + P_2 \times E_{Beam}$.

a self-trigger made between the first and last scintillator layers.

Calibration was complicated by the fact that there was measurable PMT fatigue between each muon-run. Muon runs were separated by up to a couple of days at a time because with the shutter closed, we could not take physics data. Furthermore, there was significant evidence of beam-loading - different effective gains when on-spill versus off-spill. The low intensity muon-run environment well mimicked the off-spill gain, but it is the on-spill gain we need for physics analysis. As a consequence, the BA gain values were calibrated for each muon run, and values were extrapolated between runs by scaling with LED flasher data, and on/off-spill corrections were also calculated by making comparisons between on/off-spill LED flasher signals.

The scintillator and quartz layers were calibrated differently. The scintillator layers were calibrated using minimum ionizing energy deposits, estimated through GEANT3 to be 1.8 MeV/cm for perpendicular transversal. The quartz layers though, were Cerenkov counters and their calibration via MC scaling was less straightforward. We simulated the quartz channel light yield by using a “look-up table” built from a special study that computed the number Cerenkov photons entering the PMTs as a function of energy and angle of charged particles passing through the quartz blocks - providing a base MC output in numbers of photoelectrons from the PMTs. We elected to calibrate the quartz layers from data through simple counting of minimum ionizing particles (MIPs), and then scaled the MC using a muon-run simulation to match that energy scale definition of photoelectrons to MIPs.

3.7.1 *Beam-Hole Charged Veto*

The Beam-Hole Charged Veto (BHCV) was constructed from eight scintillator plates, each 0.5 cm thick. The BHCV was designed to reject $K_L^0 \rightarrow \pi^+ \pi^- \pi^0$ decays where

both charged pions travelled down the beam-pipe and somehow failed to trigger the BA veto. It is not clear the BHCV was actually *needed* for this purpose, but because it sits forward of the BA, it helps to expand the solid area useful for vetoing decays downstream of the CsI and because it was an extremely easy detector to build and install (it sits completely outside the vacuum region) it certainly did not hurt the experiment to have it.

3.8 Collar Counters

The general purpose of the Collar Counters (CC's) was to veto particles escaping from the main fiducial volume with trajectories near the beam axis. There were seven such counters constructed in several different ways. It is worth noting that the number sequence for the CC's skips CC01 for historical reasons (the detector was originally included in the design, but later removed [44].) Most of the CC's were calibrated using beam muons (produced and kept "clean" by closing the beam shutter.) In particular, it was possible to form coincidence triggers with CC00, CC02, and CC0{4-7}. Because of the orientation of its scintillator planes, CC03 was calibrated using cosmic rays.

3.8.1 CC00

CC00 was a tungsten-scintillator sandwich detector installed around the beam-pipe and outside the vacuum region after Run-I in an attempt to further reduce the neutron content in the beam halo with an active detector element. Unfortunately, it is difficult to isolate the impact of CC00 since at the same time it was installed, the vacuum membrane that caused problems in Run I [34] was also repaired.

3.8.2 *CC02*

Chicago was entirely responsible for the design and construction of CC02. It is a *shashlik* style lead-scintillator sandwich veto. In a typical lead-scintillator sandwich detector, light produced inside the scintillator is extracted by optical fibers placed in grooves running parallel to the plane of the scintillator sheet. Due to its placement inside the Front Barrel, this was not an option for CC02. However, the detector would perform best if the lead and scintillator plates were kept perpendicular to the beam direction in order to minimize the possibility of a particle passing through a scintillator plate only (drastically lowering the probability of an interaction). To solve this problem, light was extracted from CC02 using optical fibers that were inserted in holes drilled into the lead and scintillator sheets such that the fibers passed perpendicular to the material planes.

CC02 is not the world's first shashlik style veto, but it is undoubtedly one of the most ambitious. A series of MC studies and test bench prototypes fixed the fiber density, but the task of drilling and aligning hundreds of thousands of holes in the plates and still fitting everything together with sufficient tolerance for the detector to play a role in the alignment and support of the Front Barrel was non-trivial. See Figure 3.18 for an illustration of the module lay-out.

3.8.3 *CC03*

CC03 is a “traditional” sandwich detector (as compared to a shashlik detector) composed of tungsten-scintillator modules with fiber readout parallel to the plates. CC03's modules are oriented such that the planes of the material are parallel to the beam, rather than perpendicular. See Figure 3.19.

CC03 was calibrated using cosmic ray muons in the same procedure as used for




Figure 3.18: CC02, shown here looking downstream (the beam would pass through the center gap region).

the CsI. The procedure is complicated only slightly by the presence of “horizontal” (long axis parallel to the x -direction) and “vertical” (long axis parallel to the y -direction) modules. In general, the same track fitting procedure used for the CsI can be employed in CC03 (typically, CC03 tracks are also used simultaneously for CsI calibration), here assuming 2 MeV/cm for minimum ionizing particles.

3.8.4 CC04 and CC05

CC’s 04 and 05 are lead-scintillator sandwich detectors positioned behind the CsI array. CC04 sits inside the vacuum vessel with the CsI itself, while CC05 sits outside and the vessel and even further downstream. They are both $40 \times 40 \text{ cm}^2$ in area on the front face, with $6.2 \times 6.2 \text{ cm}^2$ square apertures in their centers for the vacuum beam pipe (which stretches almost the BHCV). Both detectors feature “charged veto” and “calorimetric” sections. The charged-veto sections are only a few layers thin and read




Figure 3.19: CC03 as seen looking down the beam axis.

out by two PMT's (one from each transverse side), while the calorimetric sections are much thicker, and read out by their own pair of PMT's. CC's 04 and 05 differ in that CC04 has only one charged layer on the upstream face of the detector (and is thus read out in four total channels), while 05 has two charged-veto layers (one on the upstream face, and one on the downstream for a total of six channels). The detectors are described in greater detail in [33].

3.8.5 CC06 and CC07

While CC's 04 and 05 are very similar, CC's 06 and 07 are essentially identical. They were each constructed from ten blocks of lead-glass, sized $15 \times 15 \times 30 \text{ cm}^3$ (with the long axis perpendicular to the beam). See Figure 3.20 for a schematic of their design. CC06 was initially intended to protect a very narrow sliver of the 4π hermetic veto coverage through which a photon might escape and CC07 was placed to guard

against splash-back from the BA.

Both detectors served their purposes, although CC06 was not a very important veto insofar as background rejection from photons was concerned. However, both additionally provided protection from charged pion decays that took place downstream of the detector for catching transverse-going (and forward boosted) muons.

CC06/CC07 (front view)




Figure 3.20: CC06/CC07 - both detectors were constructed identically out of lead-glass blocks.

3.9 The Vacuum System

The decay region was evacuated to 10^{-5} Pa. At 10^{-4} Pa, we expect the background contribution for the E391a beam design to be negligible (≤ 0.1 events) for a S.E.S. of $\mathcal{O}(10^{-10})$ [45]. Out-gassing from detector components make this vacuum level

difficult to achieve, so the decay region was separated from the detector components by a membrane, and the “low vacuum region” we held the detectors in was kept at 0.1 Pa. The details of the vacuum system can be found elsewhere [34]. See Figure 3.21 for an illustration of the basic lay-out.




Figure 3.21: The vacuum system.

CHAPTER 4

MONTE CARLO

4.1 Critical Summary

Here we examine the E391a primary simulation in significant depth. The author was particularly involved in the design, coding, tuning, and validation of the simulation. Two critical results presented here are the consistency of our flux estimate between $K_L^0 \rightarrow \gamma\gamma$, $K_L^0 \rightarrow \pi^0\pi^0$, and $K_L^0 \rightarrow 3\pi^0$ and the overall systematic uncertainty attached to our flux (and therefore to our sensitivity). The first is important for establishing an understanding of our detector in the case of three very different event topologies and the second summarizes our overall precision. Our flux is $(5.13 \pm 0.40) \times 10^9$ and we achieve consistency in our flux estimates between decay modes at the level of 6%. The systematic uncertainty on our normalization is 7.3%. This Chapter is plot and table dense, but this information ultimately characterizes our understanding of the data.

4.2 Introduction

The E391a Monte Carlo (MC) is a GEANT3 (FORTRAN) program [46] that incorporates every physical detector element in precise detail. We did not employ a “Fast” MC for any major purpose once we finished the detector design stage. The primary reason for this was that we felt it would be difficult to trust a Fast-MC for background estimation in a rare-decay search. Unfortunately, this means we are statistics limited for a number of decay modes of interest (especially $K_L^0 \rightarrow \pi^0\pi^0\pi^0$ and the high branching ratio charged modes like Ke3, etc.). However, limited man-power and the problem of tuning a Fast MC correctly left us with few viable options. The E391a

MC, “Gsim,” went through a number of iterations, each labeled by “Test” number. In the early stages of the experiment, there was an additional, completely separate GEANT3 MC written entirely by the author called the “Chicago MC.” After Gsim Test 5, we ceased maintaining the Chicago MC and folded some of its superior features (most notably an accidental overlay procedure) into Gsim. Two MC Tests were used for our large statistics sets; Test 7 for Run I and Test 10 for Run II. The author was primarily responsible for tuning, preparing, and producing Test 10.

4.2.1 *Data and MC Processing*

MC events are produced and translated in the exact same file format used for real data. Information contained in these files differs from real data in the sense that the information is GEANT3 output, and so already coded in terms of energy and time. In the first stage, real data is coded in terms of ADC and TDC. MC and data files are fed into our analysis software to produce “skimmed” files. At this stage, we transform the information coded in data from ADC and TDC to energy and time with our calibration database and count the number of clusters reconstructed in the CsI calorimeter. We apply additional smearing and correction functions to the MC, but do not change the way the information is coded (it remains in an energy and time basis). We do not “digitize” the MC at any stage. Digitization involves transferring energy and time basis information into ADC and TDC basis information. This process allows MC files to be handled *exactly* in the same way as data and has many advantages. Studies conducted very early in the experiment lead us to believe the advantages did not outweigh the burden of additional processing time and the additional data storage that would be required. However, a “digitization-like” smearing procedure is applied to the CsI energies and the BA quartz crystals.

The MC comes in two basic forms, with differentiation at the *skim stage* - where raw data and/or MC files are separated by the number of clusters reconstructed in the CsI calorimeter. One is the so-called “Pure” MC and the other is the “Add-BG” MC. The Pure MC is derived from the GEANT3 output directly and makes no attempt to incorporate accidental activity except in the form of a post-analysis numerical factor to account for accidental acceptance loss (discussed later in this report). This post-analysis factor was calculated by studying events from the Run II data set taken with the target-monitor accidental trigger. The Add-BG MC takes events from the target-monitor accidental trigger and applies them additively to the raw MC events on a one-for-one basis. Both versions of the MC have pluses and minuses - the Add-BG MC is plagued by a timing offset in the accidental trigger and the Pure MC misses interesting topological effects created by accidental hits.

Ultimately, it was the consensus of the group that the (accidental-loss corrected) Pure MC estimates for the *flux* are appropriate. However, in order to account for the effects of accidental activity in direct data-MC distribution comparisons, we use the Add-BG MC, except where explicitly noted. Furthermore, for background estimation it is important to include the topological changes in events introduced by ambient detector activity, so the Add-BG MC is used in those cases as well. Of course, the predictions for the flux are cross-checked between the two methods and their agreement helps provide confidence in our estimate.

4.3 Accidental Overlay

“Accidentals” refer to the ambient activity in the detector. When protons strike our target, they produce a very large number of particles, relatively few of which are Kaons. Many of these particles can be removed via sweeping magnets and absorbers,

but there is a large flux of photons and neutrons that enter our detector in-time with the Kaons we are interested in. These additional particles interact with detector materials to create false vetoes (accidental losses) or, worse, false signal events. Unfortunately, simulating ambient beam activity correctly is very difficult and time consuming.

We employed two solutions to this problem. The first was to analyze the so-called “Pure” MC - the simulation output prior to overlay - and accidental data separately. With this method, we developed separate groups of timing cuts and combined the results from the accidental study in a multiplicative fashion. This is discussed further in the next section. The second solution was to use a so-called “Add-BG” (Add-Background) MC that employed a direct event-for-event overlay. To combine timing information for each channel, we kept the earliest time between the native MC output (GEANT-based time-of-flight timing with some corrections for time-walk, or energy dependence, and electronics) and the accidental event. This MC had the difficulty of more ambiguous timing windows. Our solution was to center each timing window using the Pure MC, but set the width using accidental and physics data. As a consequence of the out-of-synch zeroing, there is some bias present in the Add-BG MC. However, the Add-BG MC was required to create matching overlay plots, especially of the veto counters, and for background studies.

4.3.1 Accidental Trigger Timing

The initial trigger scheme for E391a did not include an accidental trigger. Shortly before the first experimental Run, we installed a scintillator panel beside the target and set signal cable to feed information from this target monitor into the DAQ system. Because the detector was installed late, the relative timing between signals from the

target monitor and the physics trigger was not well established. In principle, one would like to synch the delay between all detector elements such that they share a common zero in time. Because of the finite transmission time through signal cables and electronics, it is possible to create differentials in the timing information from different detector elements.

Unfortunately, this ended up the case for our accidental trigger, as can be seen in Figure 4.1. In the Figure we compare the mutli-hit information for the BA using the physics trigger (specifically, using $K_L^0 \rightarrow 3\pi^0$ events which have roughly zero photon content in the beam hole from Kaon decays) and using the accidental trigger. Because the BA timing is multi-hit, it can show the beam microstructure. (Multi-hit TDC is capable of measuring more than just one stop signal for each channel after the start signal.) Notice in the figure that the peak of activity corresponds roughly with zero for the physics trigger, but is out of synch for the accidental trigger. This means that if we look at accidental overlays of the MC with the same definition of zero time as used in data, we are significantly under-sampling the ambient beam activity.

4.3.2 Accidental Loss Estimation - Pure MC and Overlay

For the Pure MC, the accidental loss formula is $B \times C_1 \times C_2$, where:

- B = the inclusive loss due to all photon vetoes except the CsI,
- C_1 = the loss due to additional clusters,
- C_2 = the mode specific loss in the CsI due to accidental activity that does not produce an extra cluster.

B was estimated by examining accidental data and counting the pure number of events vetoed regardless of the number of clusters in the CsI - which is to say, “signal” was




Figure 4.1: BA multi-hit timing comparison for $K_L^0 \rightarrow 3\pi^0$ (left) and accidentals (right) data.

defined simply as passing all veto cuts except the CsI. C_1 was estimated by simply counting the fraction of accidental triggers that contained a reconstructed cluster in the calorimeter. No photon quality cuts were imposed on these events because they initiate losses by shifting the number of clusters we reconstruct. In the case of a real $K_L^0 \rightarrow \pi^0\pi^0$ event for example, any number of extra clusters from accidental activity shifts the event out of the four-photon “stream” and so it is not analyzed.

The CsI is separated from the other pieces because the veto is a function of the distance between “single-hit” crystals and clusters (it is discussed in more detail later in this Chapter). As a consequence, it has an effective dependence on the *number* of clusters, and so cannot be estimated from accidental data directly (only $\sim 0.7\%$ of which even contains any clusters in the CsI). In this sense, because it is effectively a multi-particle cut, the CsI veto should be classified as a kinematic cut. However, for historical reasons, it has been kept as part of the “photon-veto” group of cuts.

C_2 is estimated by looking at the ratio of the exclusive acceptance of the Pure and Add-BG MC's, where the exclusive acceptance is the acceptance of one cut with all others applied.

Mode	$K_L^0 \rightarrow \pi^0\pi^0$	$K_L^0 \rightarrow 3\pi^0$	$K_L^0 \rightarrow \gamma\gamma$
Inclusive Acceptance		$84.76 \pm 0.05\%$	
1 - Loss due to extra clusters		$99.33 \pm 0.04\%$	
Mode-specific CsI acc.	$97.99 \pm 1.35\%$	$98.45 \pm 1.43\%$	$97.95 \pm 0.57\%$
Total	$82.50 \pm 1.35\%$	$82.89 \pm 1.43\%$	$82.47 \pm 0.57\%$
Weighted average		82.57%	

Table 4.1: Accidental losses for the main calibration modes. Here, errors are purely statistical.

4.4 MC Tuning

Using the Test 10 MC, we produced K_L^0 events at C6, the exit of the E391a beam collimator system (see Chapter 2). Our MC does not produce K 's at the target. In order to save CPU time, we studied production at the target and used it as a basis to model “production” at a later point in the detector geometry.

Gsim employs an event generator routine that creates particles with initial positions and momenta drawn randomly from user-defined distributions and allows the particles to evolve subsequently under the control of GEANT. Typically, the branching fraction of the particle is modified to be unity in a mode of interest - e.g., for $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ studies, the branching ratio for the mode is set to unity for the MC sample. One important exception is the branching ratio of the $\pi^0 \rightarrow \gamma\gamma$, which is held fixed at 98.802%.

The original user-defined distributions that govern the initial particle distributions were, for the most part, taken initially from beam-line simulation results and empirical observations from E391a Run I. The initial distributions for the Run II MC did not

produce *final* distributions of key variables (momentum, decay spectrum, etc.) that matched those of the K 's reconstructed in data well at all. This initial mismatch was due in part to detector simulation issues (e.g., correctly reproducing the energy and timing resolution for each veto counter), but also due to an improper set of choices for those initial distributions. Both the detector response functions and beam profile distributions needed to be tuned by hand.

Therefore, in order to produce a matching set of spectra, we tuned the seed distributions in Gsim using an iterative correction process. The tuning was initially done by using a set of kinematic and veto cuts looser than the analysis cuts (looser to boost event statistics and reduce the time between iterations). Later the tuning process shifted to use only the kinematic cuts, but at analysis-level tightness. This shift was performed to minimize the impact of problems in simulating the behavior of individual veto counters, and to isolate beam replication from the veto response corrections and changes in the photon veto cuts. It was possible to make a shift like this because the mode used for tuning was $K_L^0 \rightarrow \pi^0\pi^0\pi^0$ - the veto cuts were not necessary to achieve a very high order of background reduction and high quality signal-to-noise ratios. See Table 4.2 for a listing of the cuts used in the tuning process.

4.5 Event Weights

At the end of the MC tuning process, a significant mismatch remained in the radial shape of the Kaon beam (see Figure 4.2). Some piece of this was due to the use of a cylindrically symmetric generator - it cannot perfectly match an ovalar beam. However, some significant component was correctable. Owing to time constraints, we could not adjust the generator function as many times as we would have liked. Instead, we chose to begin mass-production with a faulty beam shape and then *re-*

$3\pi^0$ Cut	Values	Comments
Photon Energy Min.	150 MeV	
Photon CsI Hit Position	$17.5 < r < 88$ cm	The inner dimension forms a square around the beam pipe, the outer is radial.
Photon Hit Separation Min.	17.5 cm	
Photon Fusion Neural Network Min.	0.5	
Pion z -difference Max.	10 cm	
Pion Mass-difference Max.	$5.125 \text{ MeV}/c^2$	
Best z - χ^2 Max.	3	
2nd Best Minus Best z - χ^2 Min.	4	
r^2 at C6 Max.	4.5 cm^2	Also cut if $r(z_{decay})$ is less than R(C6).
Decay z	$340 < z < 500$ cm	
Transverse Momentum Max.	$1.25 \times 10^{-4} (\text{GeV}/c)^2$	

Table 4.2: Kinematic Cuts for Six Cluster Events. These cuts are used in the analysis along with photon veto cuts to compute the Kaon flux. Additionally, these cuts served as the primary cut set for the MC tuning process.

weight the MC. In order to simplify the analysis, we decided to assign only a single weight to each event.

The radial distribution of the Kaon beam at the exit of our collimator system was our worst matching variable, so we chose to assign an event weight based on that position, using the real, “generator” information in the MC. This was done to make the weight function portable between decay modes. The empirical weight function is specified by:

$$\begin{aligned}
 R_{Gen}^2 &= gx(1)^2 + gy(1)^2, \\
 W &= 1.305 - 0.192559 \times R_{Gen}^2 + 0.0029822 \times \left(R_{Gen}^2\right)^2.
 \end{aligned}
 \tag{4.1}$$

Here “gx” and “gy” are specific variables from the common “gnana” routines - they

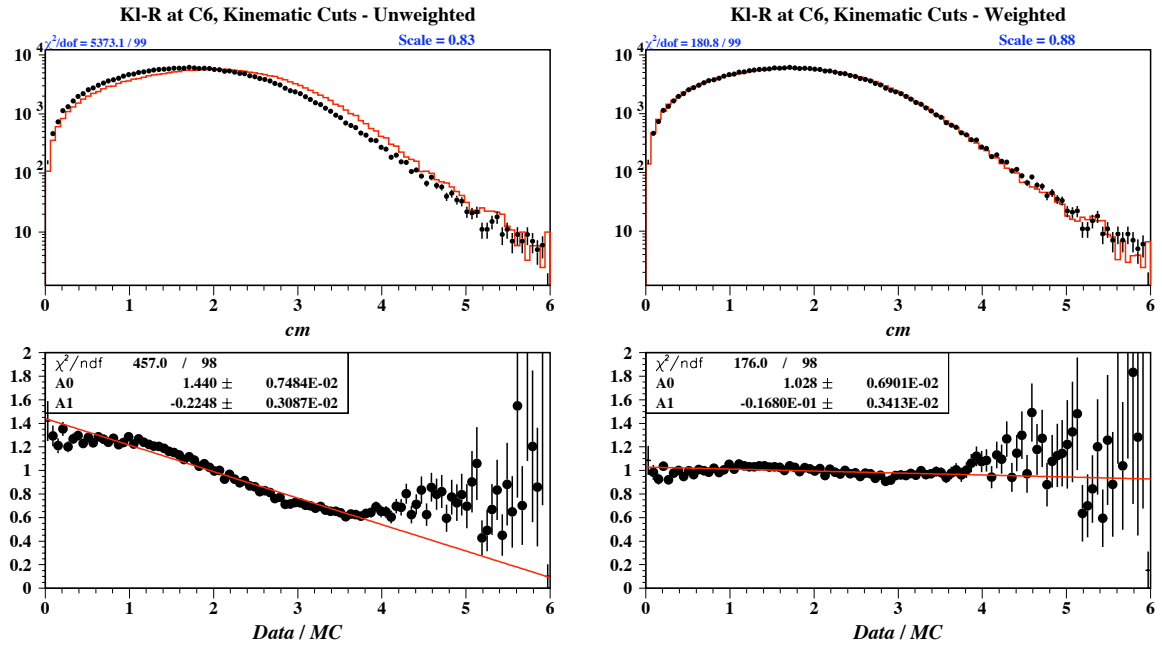


Figure 4.2: Shown above is the reconstructed radial spectrum of the Kaon at C6 at the end of the Test 10 tuning process before (left) and after (right) a weight factor was applied for six cluster ($K_L^0 \rightarrow 3\pi^0$) data and MC. All kinematic cuts (except the cut on radial position, see Table 4.2) are applied. Here and throughout for overlay plots, data are shown in the upper half as black dots with error bars and the MC events are shown in solid red. The lower half shows the bin-by-bin ratio of data and MC, normalized by total event count. Here, the fit curve is $y = A_0 + A_1 \times r$ and is plotted in red. Errors on both distributions are counting errors (\sqrt{N}). The $\chi^2/d.o.f.$ variable in the upper left-hand corner is a measure of bin-by-bin discrepancies, here with 100 bins (one degree of freedom is used for normalization). The scale variable in the upper right-hand corner is the number used to scale the *MC* so that the integrated event number for both distributions agree.

are the *generated* x and y positions of the Kaon at C6. The MC events are produced with weights of unity by Gsim and re-weighted at the *analysis stage* to better match the radial spectrum at C6 (the MC tends to be widely distributed relative to the data in Test 10). See Figure 4.3 for a sample distribution of the weights (the mean value of the distribution is approximately 0.95).

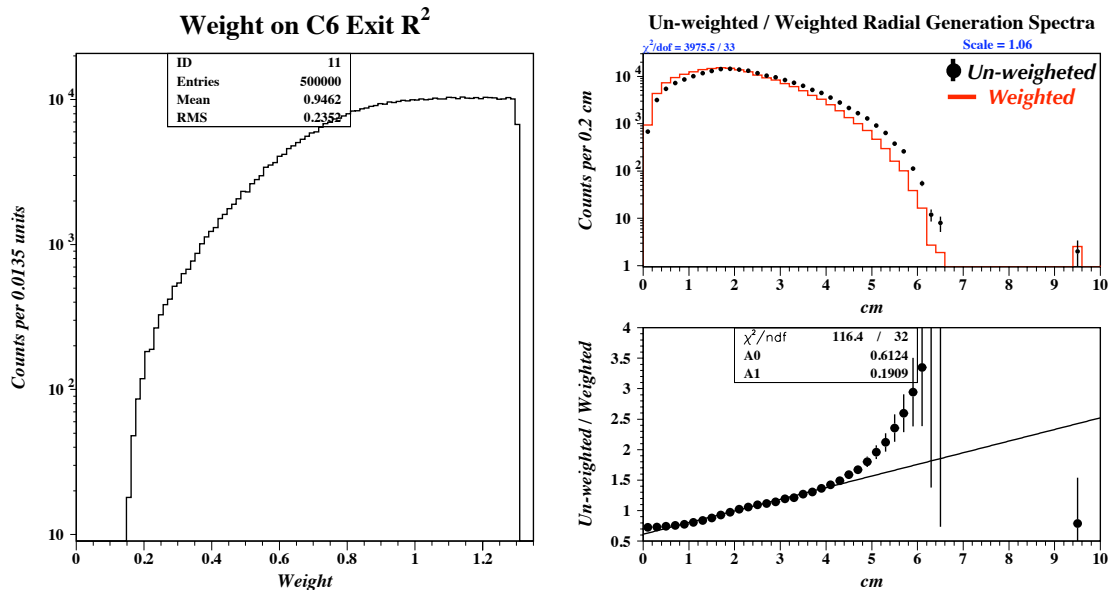


Figure 4.3: Shown above is the collection of event weights for six cluster events (left) and the radial generation position at C6 (right) for a special version of the MC with interactions turned off (pure decay kinematics). On the right, the re-weighted distribution is in red and the un-weighted is in black (with error bars = \sqrt{N}). Overlay plot conventions are the same as in Figure 4.2 (and throughout this document). The event at very high r is due to a non-zero Kaon occupancy in the beam halo.

4.6 Decay Probability

Gsim generates particles at C6 and lets them propagate from that point. It would certainly be possible to compute the Kaon flux for E391a Run II in terms of Kaons at the exit of the collimator system, but that measure is somewhat unnatural and arbitrary. What we are truly interested in is the number of Kaons that *decay* inside

our main detection chamber (inside the Main Barrel, before the CsI face). Our standard fiducial volume is taken to be 340 - 500 cm, where 0 cm is defined to be the beginning point of the Front Barrel veto.

In this coordinate system, the exit of C6 is located at -199.5 cm in Run II. However, due to a software bug, the actual generation point in the Test 10 MC was -192.5 cm (7 cm downstream of C6). In principle, this has no major effects on the MC, but it is important to remember that when we refer to the “exit of C6” in this report, we are actually referring to a point 7 cm downstream of that physical point in the detector system.

Events produced at C6 are not guaranteed to decay inside our fiducial volume and indeed most do not decay inside the detector at all. The decay probability inside the fiducial volume is $0.0214 \pm 0.000207_{Stat}$. See Table 4.3 for a list of the important numbers required for this calculation. We used a special 500,000 event run with all interactions turned off (essentially, the special run was a brute-force numerical integral over the momentum spectrum that appears in Test 10). It is important to remember we use the *weighted* value when working with the radial-spectrum weighted MC. Note also that this particular table is not applicable in the case of a different choice of fiducial volume.

4.7 Acceptance for the Calibration Modes

The MC serves three primary purposes: acceptance calculation for normalization and signal modes, background level estimation, and as a detector calibration aid. The acceptance calculation allows us to compute the *flux*, or the number of Kaon’s that decayed inside our fiducial region. See Table 4.4 for a listing of the acceptances for the three main normalization modes and Table 4.5 for a list of the Kaon MC samples

Events Generated at C6 - Raw	500,000
Total Decays in the MC Mother Volume - Raw	96,854
Total Decays in the MC Mother Volume - Weighted	91,555.6
Decays in the Fiducial [340,500] - Raw	11,296
Decays in the Fiducial [340,500] - Weighted	10,713.8
Raw Decay Probability	$(2.26 \pm 0.0213_{Stat}) \times 10^{-2}$
Weighted Decay Probability	$(2.14 \pm 0.0207_{Stat}) \times 10^{-2}$

Table 4.3: Decay probability. This quantity applies to events generated at C6 in the MC, and has no other meaningful interpretation. This value is important for flux computation because we define our flux as the number of Kaons decaying inside the fiducial, and our acceptance as the fraction of those decays that are visible (reconstructed). We chose this definition for the flux for ease of understanding and naturalness.

produced and their relative size. Of course, to compute the relative size of each sample, we first need to know the Kaon flux. That result was computed first but will be discussed at the end of the Chapter.

Mode	Pure Accid. Loss	Pure MC Accept.	Add-BG Accept.
$K \rightarrow \gamma\gamma$	17.53%	$(6.93 \pm 0.03) \times 10^{-3}$	$(6.95 \pm 0.03) \times 10^{-3}$
$K \rightarrow \pi^0\pi^0$	17.50%	$(3.35 \pm 0.03) \times 10^{-4}$	$(3.35 \pm 0.03) \times 10^{-4}$
$K \rightarrow \pi^0\pi^0\pi^0$	17.11%	$(7.13 \pm 0.06) \times 10^{-5}$	$(7.02 \pm 0.06) \times 10^{-5}$

Table 4.4: Acceptance estimates. Here the quoted error is the *statistical error*. The discrepancy for $K_L^0 \rightarrow 3\pi^0$ is likely rooted in the fact that most of the acceptance loss due to photon vetoes is from accidental losses, and we do not capture accidental activity perfectly in our accidental trigger. The other modes have real photons from other Kaon decays that dominate the acceptance loss in the vetoes.

Broadly speaking, our $2\pi^0$ sample is significantly larger than our data set, and our $3\pi^0$ sample is significantly smaller. The reason for this disparity is two-fold: 1) the $2\pi^0$ MC usually involves only four photons in the final state (neglecting the Dalitz decay of the pion), while the $3\pi^0$ MC must contend with 50% more particles, and 2) the relative branching fraction for $3\pi^0$ is much higher. The result of this disparity is that we must either assign a large numerical weight to each $3\pi^0$ event, or examine

Mode	K_L^0 's at C6	K_L^0 's in the Fiducial	Sample Size
$K \rightarrow \pi^0 \nu \bar{\nu}$	2×10^8	4.28×10^6	NA
$K \rightarrow \pi^\pm \mu^\mp \nu$ (Kmu3)	7×10^7	1.498×10^6	$\sim 0.11\%$ of Run II
$K \rightarrow \pi^\pm e^\mp \nu$ (Ke3)	2×10^9	42.8×10^6	$\sim 2.1\%$ of Run II
$K \rightarrow \pi^+ \pi^- \pi^0$	2×10^9	42.8×10^6	$\sim 6.6\%$ of Run II
$K \rightarrow \gamma \gamma$	5×10^8	10.7×10^6	$\sim 3.81 \times$ Run II
$K \rightarrow \pi^0 \pi^0$	2×10^9	42.8×10^6	$\sim 9.6 \times$ Run II
$K \rightarrow \pi^0 \pi^0 \pi^0$	10^{10}	214×10^6	$\sim 21\%$ of Run II

Table 4.5: MC Sample Sizes. Here, K is short-hand for K_L^0 and a flux value of 5.13×10^9 is generically used to normalize to Run II. The Kmu3 and Ke3 samples are one half in the positive pion channel and one half in the negative channel.

only sub-sections of the data when we compare our four-cluster MC to the data. For four cluster events, the weighting factor for $3\pi^0$ events is computed as:

$$W_{3\pi^0} = \frac{N_{2\pi^0} \times \mathcal{BR}(3\pi^0)}{N_{3\pi^0} \times \mathcal{BR}(2\pi^0)} = \frac{2 \times 10^9 \times 0.1956}{10^{10} \times 8.69 \times 10^{-4}} \approx 45.0 \quad (4.2)$$

This weight dominates the statistical error in the combined $2\pi^0 + 3\pi^0$ MC. We compute an analogous weight for the combined $\gamma\gamma + 2\pi^0$ two cluster MC, but the statistics are rich enough in that case that the weights are small.

4.8 The Invariant Mass Spectrum and Energy Resolution

The electromagnetic package of GEANT3 automatically includes statistical fluctuation in electromagnetic showers. However, the effective energy resolution in our data is lower than this theoretical level due to crystal imperfections, temperature gradients, small calibration errors, electronics noise, etc. Rather than try to model each of those effects separately, we include an *additional* energy smearing term to account for these factors. This method is incapable of reproducing the sort of variability between individual crystals in our CsI array that probably exists, but is vastly simpler than a

piece-wise function for every part of the calorimeter.

At the skim stage, prior to clustering (and prior to accidental overlay), we shift the energy in each CsI crystal by a normally distributed amount with a width defined according to:

$$\frac{\sigma}{E_{Crystal}} = A \oplus \frac{B}{\sqrt{E_{Crystal}}}. \quad (4.3)$$

The parameters A and B for the different crystal types can be found in Table 4.6. Since this additive shift can be negative, it is possible to shift a very low energy to an unphysical, negative value. In this case, we assign zero to the energy value in that crystal. In principle, this could potentially bias the energy in the MC to the high-side (the average energy shift will be greater than zero if negative tails are set to zero). However, since the size of the shift is a function of energy, this bias ends up being very small - on the order of hundreds of eV for sub-MeV crystals, and essentially zero for crystals with more than 1 MeV of energy deposited. No special additional terms are applied to the deformed crystals and no attempt is made to adjust the resolution of the sandwich counters (which are lead-scintillator).

Note that the parameters in Table 4.6 should not be interpreted as the true *total* parameter set for the energy resolution. We measured the resolution using a five by five array as discussed in Chapter 3 (see Figure 3.7). The true intrinsic resolution is difficult to measure in the course of physics data taking because we have no way to tag the energy of incident particles on the CsI.

Crystal Type	Constant Term (A)	Square-root Term (B)
KEK	0.4%	0.8%
KTeV	0.4%	0.6%

Table 4.6: Crystal-by-crystal resolution function parameters (see Equation 4.3).

The crystal-by-crystal parameters were tuned using the *Pure* MC by scanning over

different values for the constant and square-root terms while watching the invariant mass overlay plot in an attempt to match the shape of the mass peak. Interestingly, when comparing the *Add-BG* mass peak to the data, as we do here (in Figure 4.4), the MC mass peak is slightly narrower than the data (and also therefore narrower than the Pure MC).

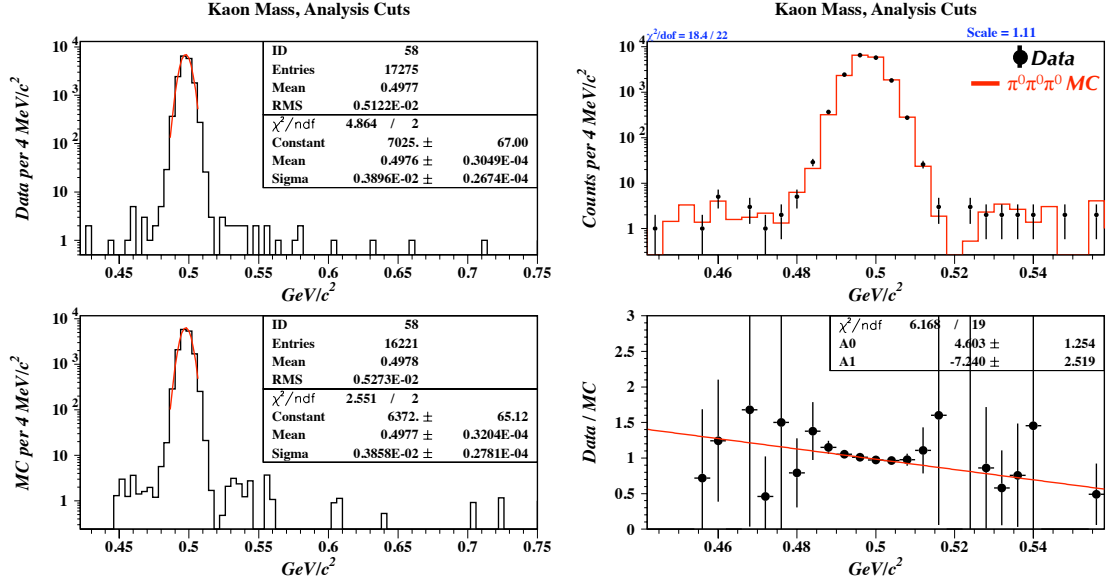


Figure 4.4: The six cluster invariant mass spectrum with all cuts applied. On the left we have the $K_L^0 \rightarrow 3\pi^0$ mass plot for data (top, $\sim 1/4$ of the data set) and for MC (bottom). Data are represented by black dots and the MC by the solid red line in the upper half of the right-hand figure. The lower half shows the bin-by-bin ratio with counting errors in dashed black and the fit line in solid red (again, normalized by total event count). Note the data's peak is slightly wider (see Table 4.7).

	Gaussian Mean	Gaussian Sigma
Data	497.6 \pm 0.03 MeV	3.90 \pm 0.03 MeV
Add-BG MC	497.7 \pm 0.03 MeV	3.86 \pm 0.03 MeV

Table 4.7: $K \rightarrow 3\pi^0$ invariant mass characteristics in data and MC with all cuts applied. The mass is biased slightly high due to a feature of the reconstruction routine.

Perhaps the clearest improvement from the crystal energy smearing can be seen in

the ability of the MC to actually replicate the energy distributions in a cluster with smearing applied. See Figure 4.5 for a dramatic demonstration of the improvement in our ability to replicate the behavior of the CsI.




Figure 4.5: Reconstructed photon energy (with energy leakage corrections applied) minus the sum of the constituent crystal energies for the native GEANT3 simulation and the crystal by crystal smeared simulation.

4.9 The Four Cluster Invariant Mass Spectrum

In the following plots, we do not attempt to perform an absolute normalization. By this we mean that we do not attempt to scale our MC to our data using our calculations for the flux. There are two reasons for this. First some discrepancies remain in the flux estimates between the Pure and Add-BG MC's. Second some discrepancies remain between the flux estimates for $2\pi^0$ and $3\pi^0$. What we are interested in is quantifying the importance of these discrepancies. Therefore, we normalize by *signal events*, where signal events are those with invariant mass values between 481.4 and 512.6 MeV after all cuts are applied. The mass values chosen to define the signal region were chosen based upon a fit to the $3\pi^0$ mass peak (the window is 3σ wide) from about 2005. Our calibration has changed since then, but the signal window has remained static. By then examining the agreement between

Data and MC in the sidebands of the invariant mass spectrum, we can see how close our MC is coming to replicating second order efficiency effects.

We scale our combined MC to the data by forcing the sum of signal events to be the same in both:

$$\begin{aligned}
 \text{Scale}_{MC} &= \frac{\text{Signal}_{2\pi^0} + \text{Signal}_{3\pi^0} \times W_{3\pi^0}}{\text{Signal}_{Data}} \\
 &= \frac{14365 + 1.0868 \times 45}{1500} \\
 &= 9.609
 \end{aligned}
 \tag{4.4}$$

Note that this value is almost equal to the relative fractional size of the $2\pi^0$ sample as quoted in Table 4.5.

In Figure 4.6 we show the combined MC invariant mass spectrum compared with that found in our data. The large weight attached to $3\pi^0$ events is clearly visible in these plots - especially note the “jagged” low mass tail region. Typically, we prefer to make comparisons between data and MC based on overlay ratio plots, but those are difficult to evaluate in this case. Instead, consider the weighted integral sums over the different regions in the mass plot shown in Table 4.8. We are almost within 1σ (errors are statistical and scaled for sample size) for sideband event reproduction. Unfortunately, our MC under-predicts the number of events in the sidebands.

$\pi^0\pi^0$ MC Signal	1494.9 ± 12.4	$\pi^0\pi^0$ MC Sideband	101 ± 3.2
$\pi^0\pi^0\pi^0$ MC Signal	5.1 ± 4.9	$\pi^0\pi^0\pi^0$ MC Sideband	75.6 ± 18.8
Total MC Signal	1500 ± 13.3	Total MC Sideband	176 ± 19.1
Data Signal	1500	Data Sideband	206

Table 4.8: Integral sums over the invariant mass spectrum of the MC and Data.

Generally speaking, events below the signal peak are due to $K_L^0 \rightarrow 3\pi^0$ events with two photons missing through a combination of various channels. Roughly one

third of the low-mass tail (30%) is due two photons lost in other vetoes, one third (40%) is due to five photon hits in the CsI, and one third is due to six photon hits in the CsI (30%). Failed clusters are effectively removed by the CsI veto (discussed in more depth for the analogous two cluster case in Chapter 6.) Events above the peak are due typically to combinatorics (pairing) mistakes and are highly suppressed by the pairing χ^2 cut described below.

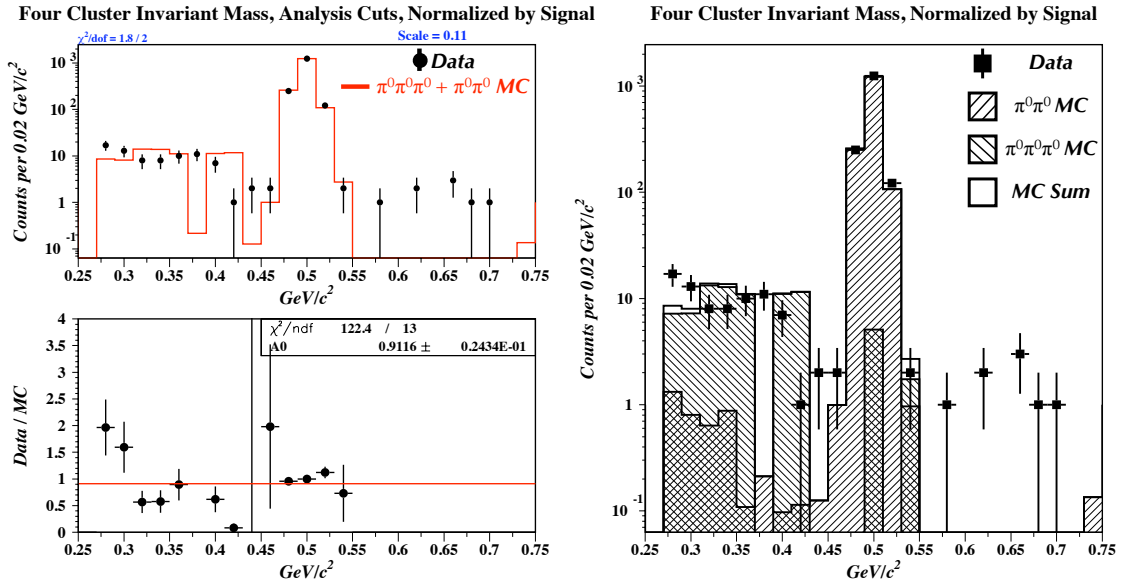


Figure 4.6: On the left is the four cluster invariant mass overlay with analysis cuts imposed, normalized by signal events. The sum MC is in red and the data points are black with error bars. On the right is the four cluster invariant mass final plot, again signal normalized.

The shape of the distribution and overall statistical consistency below the peak is well reproduced by the simulation. However, above the peak, the disagreement is more stark, with nine events in data and ~ 0 in the MC. We believe some of this is due to an imperfect simulation of the resolution of the CsI, but have little conclusive evidence because our $K_L^0 \rightarrow 3\pi^0$ MC sample is, relative to the data, so small. At this stage, we cannot effectively rule out the idea that the high-mass tail mismatch is due simply to statistical fluctuation. As can be seen in Figure 4.7, the $K_L^0 \rightarrow 3\pi^0$

MC is a small fluctuation away from participating in the high-mass tail. Given an approximate event weight of $45/9.6 \sim 4.7$, we very nearly have ~ 4.8 high-mass events in the MC versus 9 events in data, a situation where we would have statistical consistency.




Figure 4.7: On the left is the $K_L^0 \rightarrow 3\pi^0$ four cluster z -vertex versus invariant mass plot, with the same distribution for data on the right. The high-mass contribution from the $K_L^0 \rightarrow 3\pi^0$ MC appears to be a very small fluctuation in z away from participating in the mass plot shown in Figure 4.6.

4.10 Veto Distributions

The cut-points for the different vetoes are listed in Table 4.9. These cuts were chosen in two stages. First, we selected cuts that reduced the sidebands of the $K_L^0 \rightarrow \pi^0\pi^0$ mass distribution while preserving events in the signal peak using a signal-to-noise figure of merit:

$$SN = \frac{(N_S/N_N)_{\text{After}}}{(N_S/N_N)_{\text{Before}}}. \quad (4.5)$$

Second, we tightened the vetoes to handle anticipated backgrounds for the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ search. One and only one set of veto cuts was used for all the analyses presented here. This was done to minimize systematic errors between signal and normalization modes with regards to the veto cuts (i.e., we take discrepancies between the data and MC for the normalization modes to be the same for the signal mode since all the cuts are identical).

4.10.1 Veto Timing

Time zero (t_0) was established for each event as the average time of the two highest energy photon clusters in the CsI. This t_0 was subsequently subtracted from every time channel in each event. We required hits in the veto counters to fall inside specific time “windows” (lower and upper bound pairs) in order to satisfy our veto function. This was done to minimize acceptance loss due to hits well out-of-time with the activity in the CsI array. In general, reproduction of the timing distributions in the MC is extremely difficult. This is because we cannot just take into account kinematic time-of-flight variables, but must also consider physical effects in the detectors, electronics effects, and even energy threshold effects. GEANT natively tracks the time using a clock that begins with zero when the Kaon enters the detector and progresses with the time of flight for all particles. Time stamps are set for each detector volume by the first interaction that deposits energy. For most detectors, this is simply the time-of-flight from C6 (where Kaons enter the MC) to the decay vertex plus the time-of-flight from the vertex to the detector in question.

We chose not to try to reproduce these distributions in the MC, and instead employ different time windows for the MC and the data. The basic approach was to fit the timing distributions seen in two and four cluster data (where we had a good

sample of photons incident on the veto systems) and the appropriate combined MC with a Gaussian to get a mean and a width (sigma) for the distribution. We then built windows around the mean values by extending the windows to the low and high sides by some number of sigma's. We chose the number of sigma's the window extended to the high and low sides by eye, but generally used five or six. For the Pure MC, we built the time windows around the fitted mean and used a matching number of sigma's from the MC fit to the high and low ends as used in data to build the time windows. In all cases, these widths were substantially smaller than the corresponding widths in the data. For the Add-BG MC, in order to hopefully capture the correct number of accidental hits, we used the same center for the windows as used for the Pure MC, but now used the exact same sigma's as used in data to match exactly the size of the timing windows.

These methods had varying degrees of success. See Figures 4.8 and 4.9 for examples of the timing distributions in the CsI and CC02 respectively.

4.10.2 CC02

CC02 was an extremely important veto for the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ search because it surrounds the beam at the entrance to the fiducial decay volume. In Figure 4.10, we show how well CC02's energy spectrum is reproduced. Above 4 MeV, the MC reproduces the data well. We believe the discrepancy between about 1 MeV and 4 MeV is due to the bias in the accidental trigger.

4.10.3 Main Barrel

The Main Barrel (MB) was likely the most important veto subsystem. We effectively did not apply a timing window for the MB (although we did employ an effective

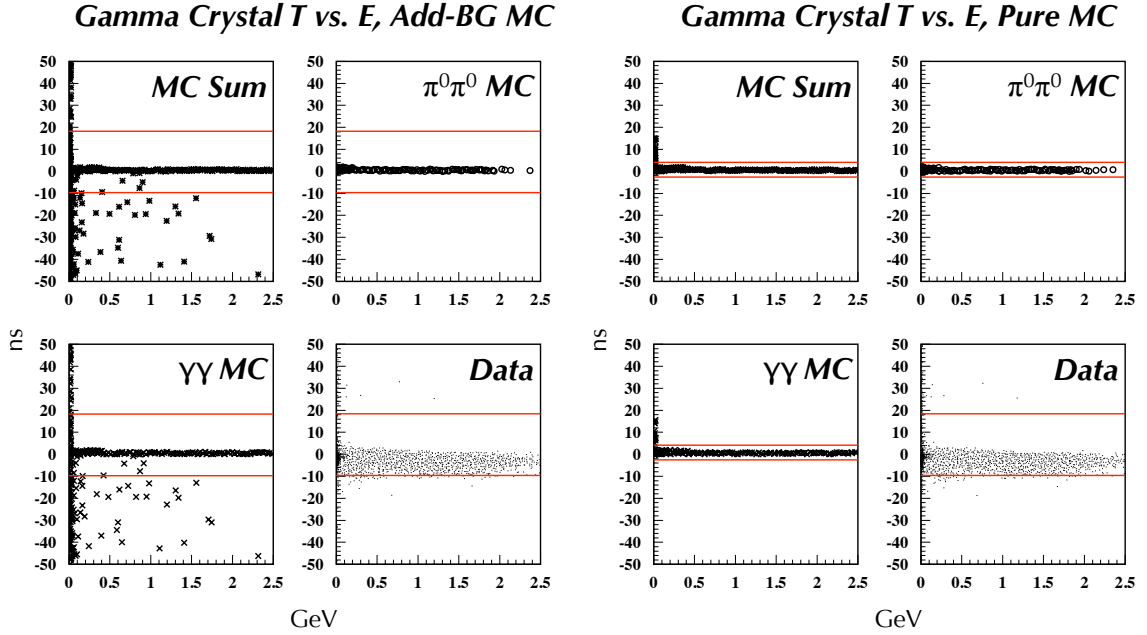


Figure 4.8: CsI time versus energy plots for two cluster data and $K_L^0 \rightarrow \gamma\gamma$ and (two cluster) $K_L^0 \rightarrow \pi^0\pi^0$ MC. The horizontal bars depict the time windows used for veto.

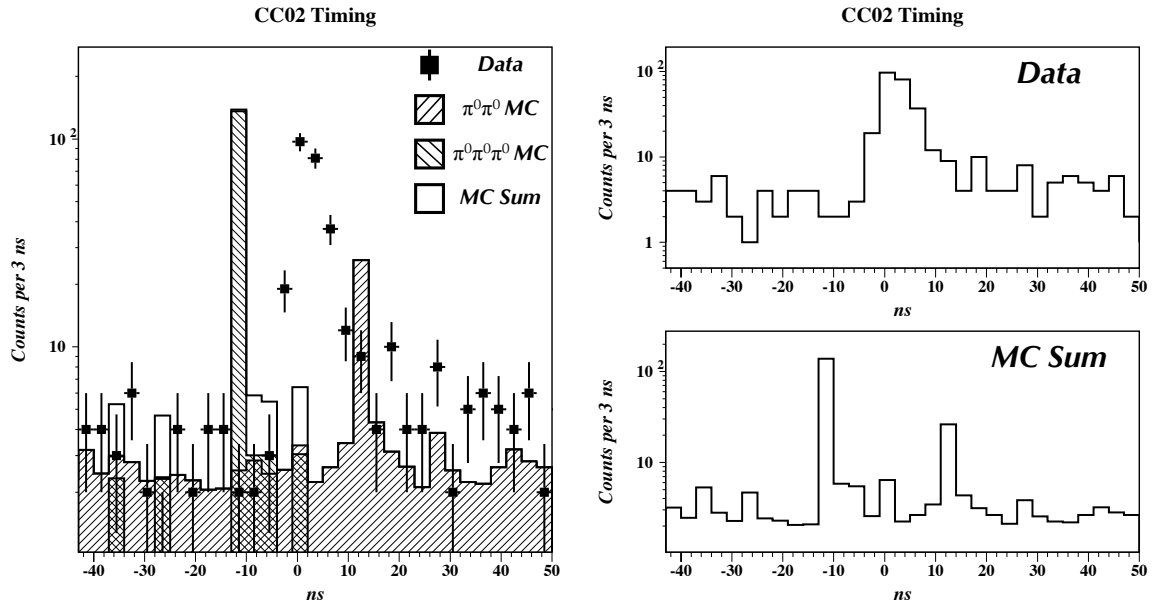


Figure 4.9: CC02 timing plots for four cluster data and $K_L^0 \rightarrow \pi^0\pi^0$ and (four cluster) $K_L^0 \rightarrow 3\pi^0$ MC. In the MC plot, the early time hits from $K_L^0 \rightarrow 3\pi^0$ are clearly visible while in the data no such peak separation exists.

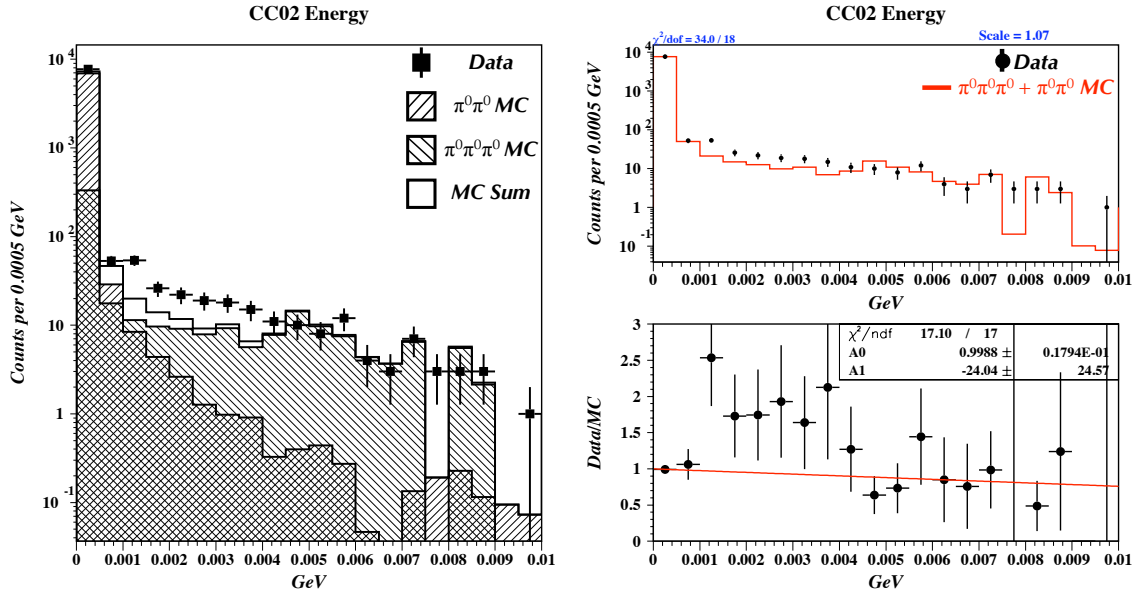


Figure 4.10: Energy overlays for CC02 using four cluster data and MC. No timing cut is imposed on CC02, but all other four cluster analysis cuts are applied.

energy “floor” by requiring hits above the TDC threshold for the veto function). A timing cut could reduce the acceptance loss due to “splash-back” from electromagnetic showers in the CsI array. However, halo neutrons interacting in CC02 or the front barrel could scatter into the Main Barrel late in time. In order to veto these events, we extended the coverage of the MB veto to late times. See Figure 4.11 for an energy overlay plot for four cluster data and MC.

4.10.4 Beam-Anti

The Beam-Anti (BA) contained one easy to simulate component (a lead-scintillator sandwich calorimeter) and one difficult component to simulate (a quartz crystal Cerenkov counter). While GEANT can produce and track Cerenkov photons, the process is prohibitively slow. For this reason, the BA was not included in the Run I MC at all. For Run II though, we included a look-up table built using a homemade

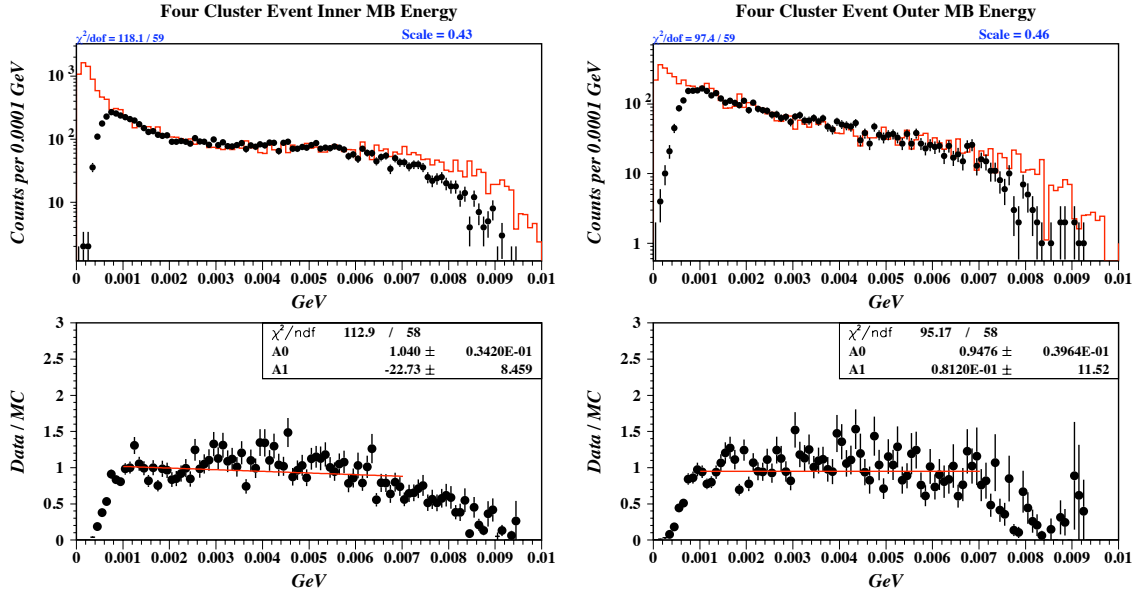


Figure 4.11: Energy overlays for the Main Barrel using four cluster data and MC. The energy value reported is equal to $\sqrt{E_{up} \times E_{down}}$, where *up* and *down* refer to the up and downstream ends of the Barrel. The energy produced removes the dependency on the attenuation length (the accompanying factor of e is swallowed by the calibration database). Obviously, if either the up or downstream end has zero energy, the product is zero. To keep positive energy values, we require the channels in the product to have energy deposition over the TDC threshold. The mismatch below 0.001 GeV is due to a TDC threshold of ~ 0.001 GeV in data, and ~ 0 GeV in the MC. The difference above 0.006 GeV is due to a tighter online veto in the data than in the simulation (where the online veto was kept deliberately loose at the generation and skim stages to avoid edge bias).

numerical simulation and the standard equations for Cerenkov radiation. We translated the number of Cerenkov photons into photoelectrons and then “calibrated” the energy scale in the MC using muons to classify minimum ionizing particles (MIPs) as the base energy scale. See Figure 4.12 for the performance of the scintillator layers and Figure 4.13 for the performance of the quartz simulation. It is important to note though, that activity in the BA is dominated by accidental hits. Since the MC contains accidental overlay, basic agreement is expected, and deviations are more likely an indictment of the accidental trigger than of the simulation itself.

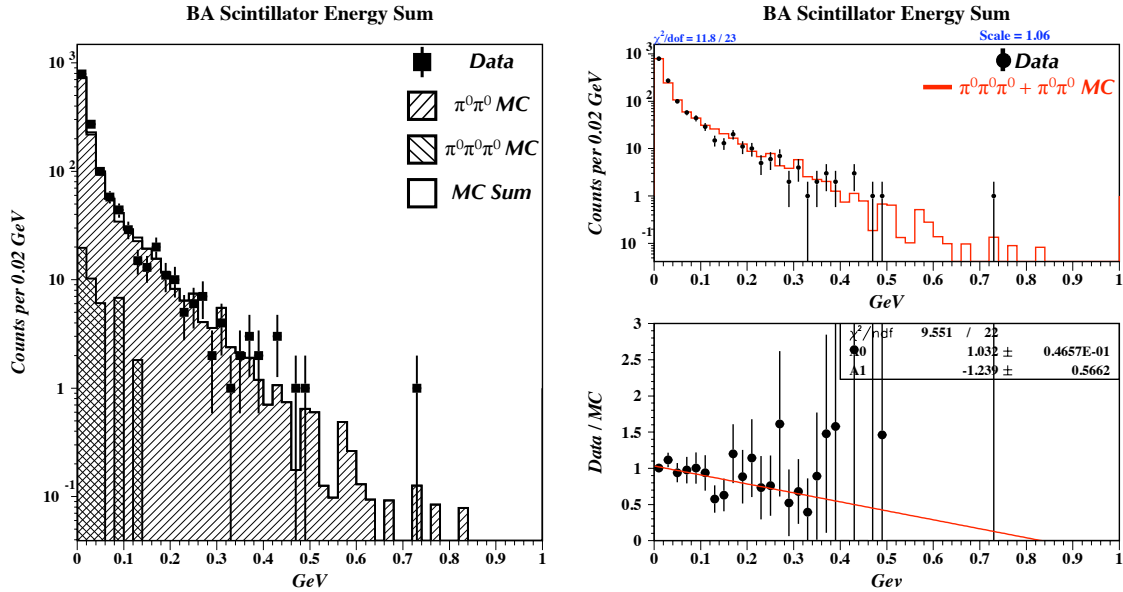


Figure 4.12: Energy overlays for the BA Scintillator channels using four cluster data and MC. All the channels are collected together in this plot.

4.10.5 CsI

The CsI veto is a two-dimensional cut that considers the relationship between *single-hit* (or, as they are unfortunately referred to within the group: s-hit) CsI crystals and their minimum distance to a cluster. We discussed the cluster formation process

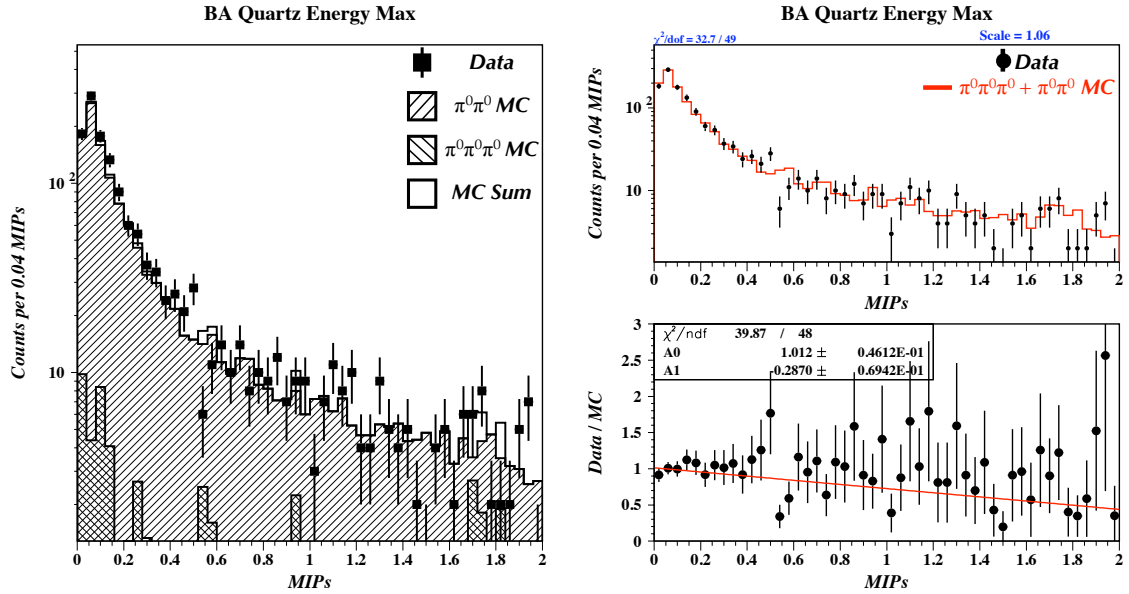


Figure 4.13: “Energy” overlays for the BA Quartz channels using four cluster data and MC. All the channels are collected together in this plot.

in Chapter 2. The critical points are that at least two crystals are required for a “cluster” and that clusters are grown by adding blocks to a seed block using only crystals that share an edge (clusters do not grow diagonally).

Because the bulk of the crystals in the CsI array are $7 \times 7 \text{ cm}^2$ across the face, and the Moliere radius, R_M , in CsI is about 3.5 cm, it is very easy for photons to deposit the entirety of their energy in a shower inside one crystal (about 90% of the energy of a 1 GeV electromagnetic cascade is contained within $2R_M$). Because our clustering routine does not count these single-crystal hits as true clusters, these events provide an avenue for decays like $K_L^0 \rightarrow \pi^0 \pi^0$ to “hide” a photon in the CsI. As a consequence, we cut tightly on the energy in these crystals.

However, because our clusters do not grow diagonally, we must be careful to avoid excessive acceptance loss by misidentifying fluctuations in the shower shapes as additional hits according to the algorithm above. See Figure 4.14 for an illustration of both kinds of single-hit crystals in one event in data.

We collect the in-time single-hit crystals for each event and cut based on the relationship between energy in the crystal and the *minimum* distance to a cluster. We do not consider the energy of the closest cluster. The cut relationship is somewhat complicated, as can be seen in Figures 4.15 and 4.16 which explore the CsI veto distributions for two cluster data and the $K_L^0 \rightarrow \gamma\gamma$ MC. The functional form of the cut is:

1. For $d_{min} < 17$ cm, $E < 10$ MeV.
2. For $17 \text{ cm} \leq d_{min} \leq 25$ cm, $E < 5 - 3/8 \times (d_{min} - 8)$ MeV.
3. For $d_{min} > 25$ cm, $E < 2$ MeV.




Figure 4.14: Shown here is a close-up of the CsI array from Figure 3.10. The “single-hit” crystals are outlined in thick black. The featured crystal close to the beam-hole is likely due to a beam-associated halo particle. The featured crystal near the bottom of the figure may be due to an accidental hit as well, but is more likely actually associated with the cluster next to it.

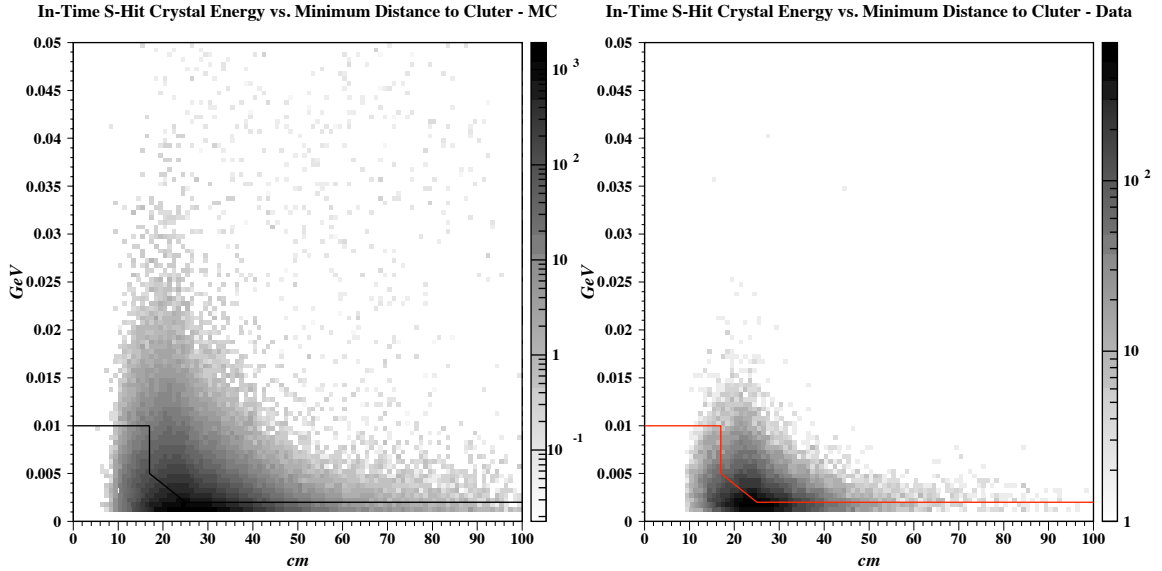


Figure 4.15: Shown here is the relevant two dimensional plot for the CsI veto, with $K_L^0 \rightarrow \gamma\gamma + K_L^0 \rightarrow \pi^0\pi^0$ two cluster MC on the left and data on the right. The distributions are signal normalized, and the apparent deficit at high distance to cluster and energy in the data is a function of the z-scale and not indicative of a mismatch in that region.

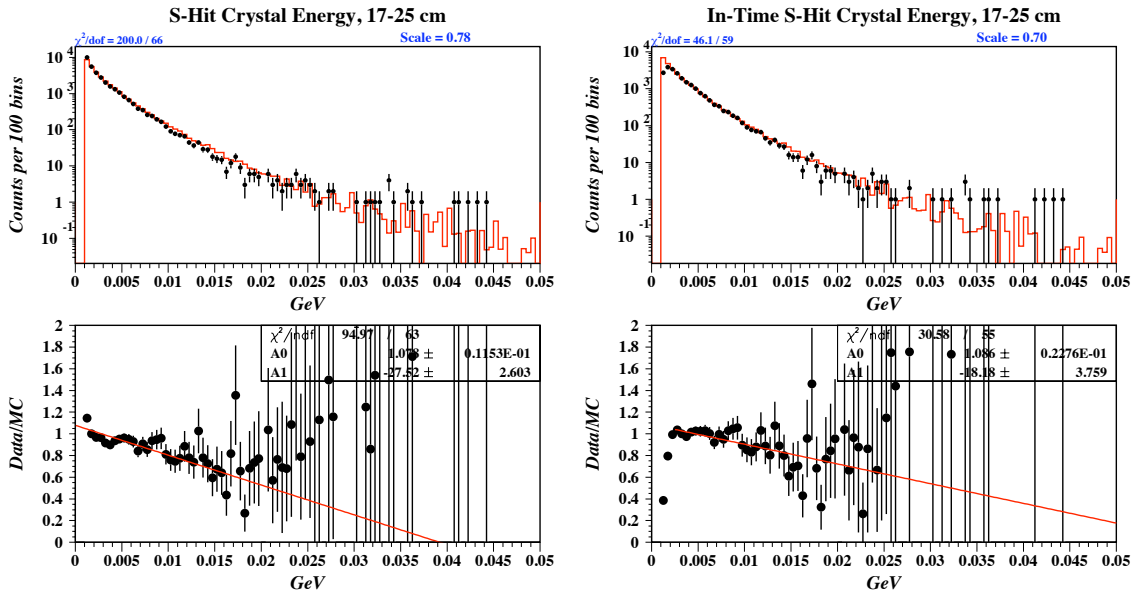


Figure 4.16: Here is a slice formed from the Figure 4.15 on the left made by integrating the distribution over minimum distance to cluster from 17 to 25 cm. On the left we show all-time hits (no timing window cut is imposed) and on the right only in-time hits . We shall have significantly more to say about the slope in these distributions.

4.11 Kinematic Distributions

The usual suite of veto cuts (Table 4.9) are applied first. The same veto cuts are universally applied to all two, four, and six photon mode analyses. Generally, the veto cuts are defined as single particle cuts, or cuts that rely only on a single veto counter. The CsI array serves as both calorimeter and veto and several of the CsI veto components ultimately rely on more than one particle (for instance, there are cuts that compare the distance between photon clusters, and these distances are a function of the cluster occupancy level). Furthermore, several kinematic cuts employed could be classified as single particle cuts (for example, the photon cluster RMS value). However, these “special” kinematic cuts are not typically universal between decay modes.

As for the cuts which do vary from mode to mode, those specific to six cluster events are listed in Table 4.2. Four cluster cuts are shown in Table 4.10 and two cluster $K_L^0 \rightarrow \gamma\gamma$ cuts are shown in Table 4.11. In general, these cuts were chosen to optimize signal-to-noise ratios for the invariant mass under the assumption the parent particle was a K .

4.11.1 Decay z Vertex

The first and probably most important distribution to consider is the decay z -vertex. Recall that the z -vertex coincides with the nominal beam axis. For $K_L^0 \rightarrow \gamma\gamma$, signal is defined by the number of events remaining between $z = [340, 500]$, while for the other modes, a cut is made on z .

See Figures 4.17 and 4.18 for the z spectra for $K_L^0 \rightarrow \pi^0\pi^0$ and $K_L^0 \rightarrow 3\pi^0$ modes respectively. Note that in the case of $K_L^0 \rightarrow \pi^0\pi^0$, most of the $K_L^0 \rightarrow 3\pi^0$ background events seen in Figure 4.17 fall outside the mass signal region. For both these modes,

there is a slope in the overlay distribution that can be removed almost entirely by removing the CsI veto from the analysis. This discrepancy will be explored in more detail shortly. For now, note that the MC tuning was conducted without the CsI veto because the function was still changing at that time.

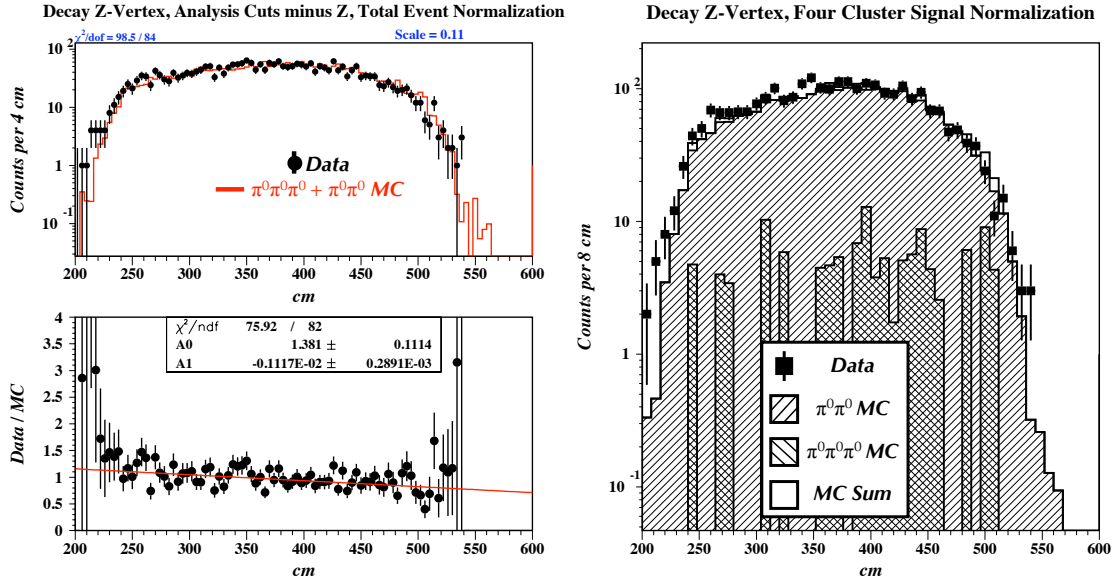


Figure 4.17: The four-cluster decay- z spectrum normalized by total events (left) and signal events (right). All analysis cuts are applied except for the fiducial (z -vertex) cut.

There are two other discrepancies visible at high and low z in those plots. The low z -vertex is explained in Figures 4.19 and 4.20. When the geometry for CC02 was coded into the MC, values differing from those measured during construction were used. As a consequence, the aperture defined in the MC is slightly tighter than in data. This constricts the available space photons from upstream decays may travel through in a differential fashion and preferentially kills upstream events in the MC.

The second discrepancy is also due to a geometry error in the MC. Figure 4.21 displays the geometry of the Outer Charged Veto in the MC, looking down the beam axis. The inner boundary in the MC is circular. However, in the actual detector,

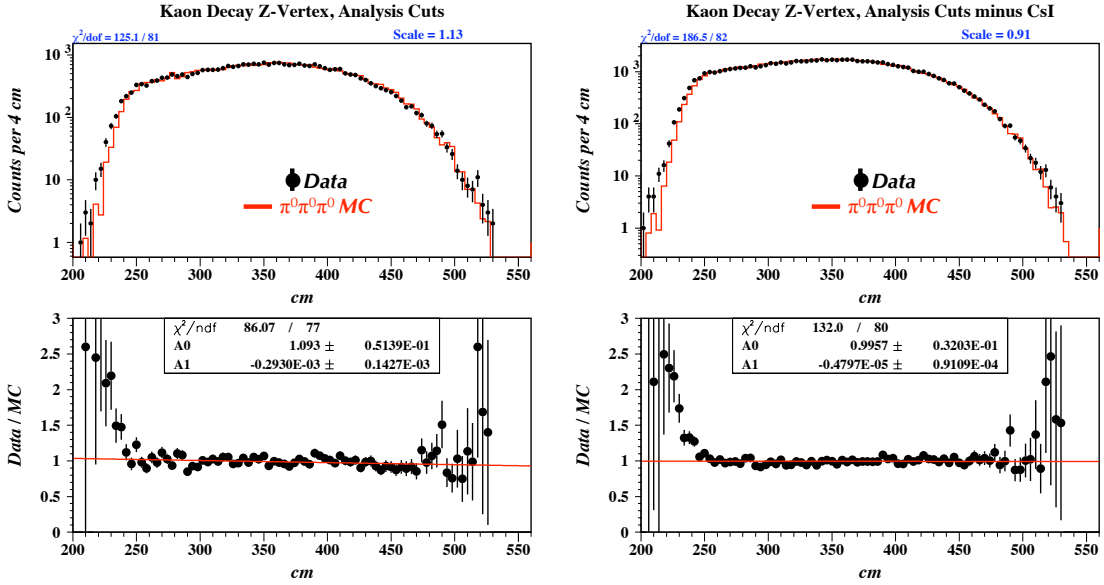


Figure 4.18: The six-cluster decay- z spectrum normalized by total events. All analysis cuts are applied except for the fiducial (z -vertex) cut to the figure on the left. The CsI veto is additionally removed from the figure on the right.

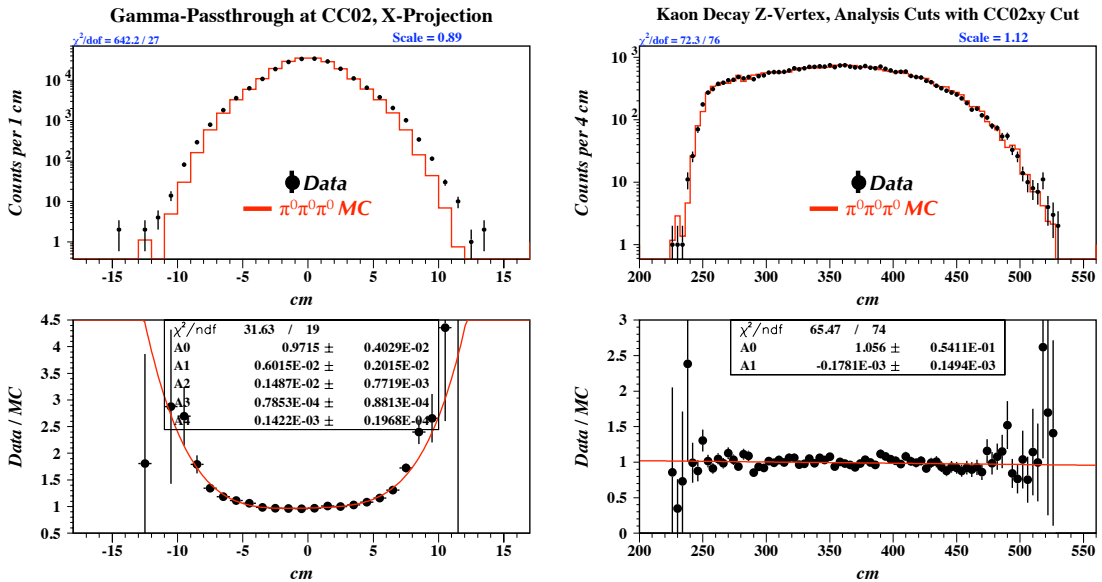


Figure 4.19: On the left is the overlay comparison for photon pass-through points in the x -direction at the end of CC02. On the right is the six-cluster decay- z spectrum with an additional cut imposed on this pass-through region.




Figure 4.20: If the CC02 aperture is not defined correctly, upstream events are differentially removed.

the inner pieces were cut away so that the boundary was a square matching the size of the beam point. The extra material near the beam in the MC provided a region where interaction rates were raised, thus distorting the z spectrum. Of course, this extra material in the beam also affected our background predictions. However, since the effect in the signal region ends up being small, and since it *raises* the background prediction, making our estimate more conservative in some senses, we judged this acceptable.

The final z spectrum plot of interest is Figure 4.22. This is the spectrum for the $K_L^0 \rightarrow \gamma\gamma$ events. The cleanness of this spectrum with regards to contamination from $K_L^0 \rightarrow \pi^0\pi^0$ is important for the $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ analysis. For $K_L^0 \rightarrow \gamma\gamma$, the $\pi^0\pi^0$ contribution can only come from low P_T events. Given that the P_T spectrum falls off exponentially, a contamination of only one part in 10^4 here provides hope that $\pi^0\pi^0$ backgrounds will be very manageable for the flagship mode.




Figure 4.21: The geometry of the Outer Charged Veto in the MC, looking down the beam axis.

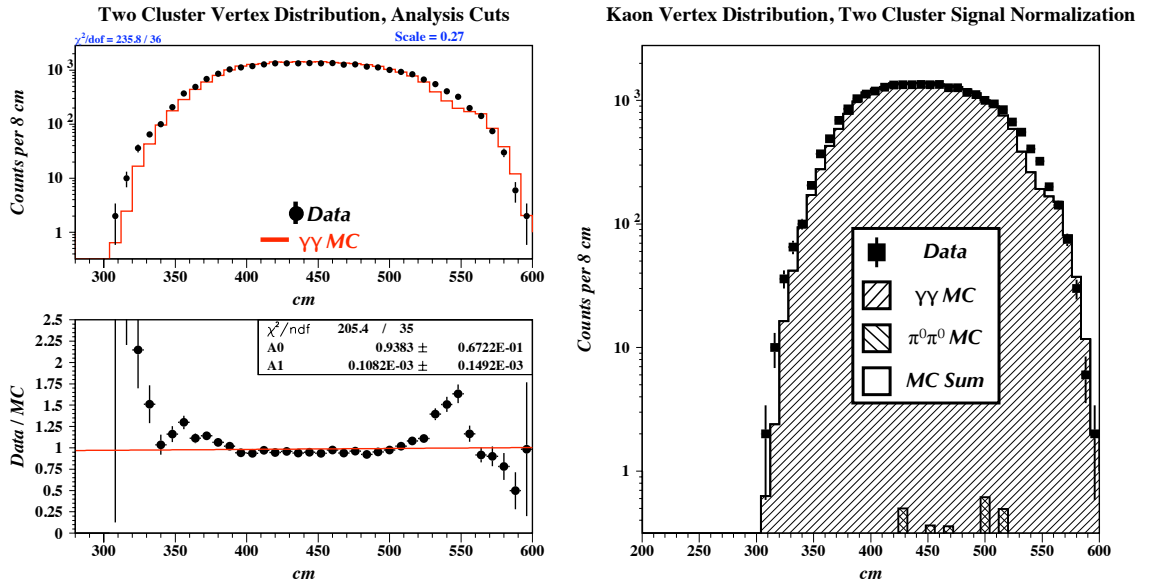


Figure 4.22: The $K_L^0 \rightarrow \gamma\gamma$ decay- z spectrum with analysis cuts imposed. The figure on the left is normalized by total events and the figure on the right is normalized using the signal region only. Note that the $K_L^0 \rightarrow \pi^0\pi^0$ contribution appears only at the level of about 10^{-4} relative to $K_L^0 \rightarrow \gamma\gamma$.

4.11.2 Total Momentum

The momentum and decay spectra are closely related. Early tuning efforts focused on attempting to match the decay spectrum for a single momentum bin because changes to the MC that affect one typically affect the other. In Figure 4.23 we show the $K_L^0 \rightarrow 3\pi^0$ MC total momentum ($|P|$) spectrum compared with the spectrum found in six cluster data. Note that in exactly the same way as in the decay- z spectrum, a distortion is introduced by the CsI veto function (but the overlay is flat without it). This same behavior is also seen in the $K_L^0 \rightarrow \pi^0\pi^0$ momentum spectrum, as seen in Figure 4.24. The functional form of the generator function for the MC (probability density) is:

$$\begin{aligned} \sigma &= \sigma_0 \times (1 - (A + p \times S) \times (p - \mu)), \\ F &= N \times \exp\left(-\frac{(p - \mu)^2}{2 \times \sigma^2}\right), \end{aligned} \tag{4.6}$$

where A is an asymmetric offset, S is an asymmetric slope, μ is the mean momentum, and N is a normalization factor.

4.11.3 Transverse Momentum and Beam Shape

The transverse momentum (P_T) of Kaons in our beam is defined relative to the nominal beam axis (the z axis). During Run II, the actual beam direction was aiming slightly down (in y) and to the left (positive x), but the error in the P_T introduced by defining the P_T relative to the nominal beam axis versus the true beam axis is negligible. Further information about the beam direction is provided below. See Figure 4.25 for a comparison of the P_T spectrum for both four and six cluster events.

We define the radial position of Kaon vertices also relative to the nominal beam

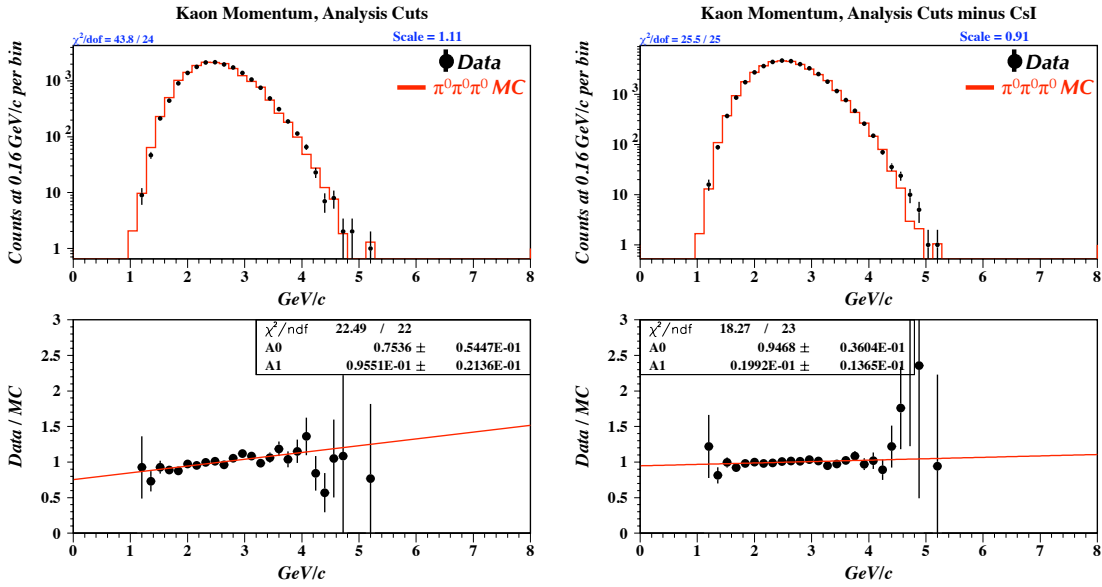


Figure 4.23: The Kaon total momentum spectrum derived from $K_L^0 \rightarrow \pi^0\pi^0\pi^0$. All analysis cuts are imposed for the figure on the left while the CsI veto is removed from the figure on the right.

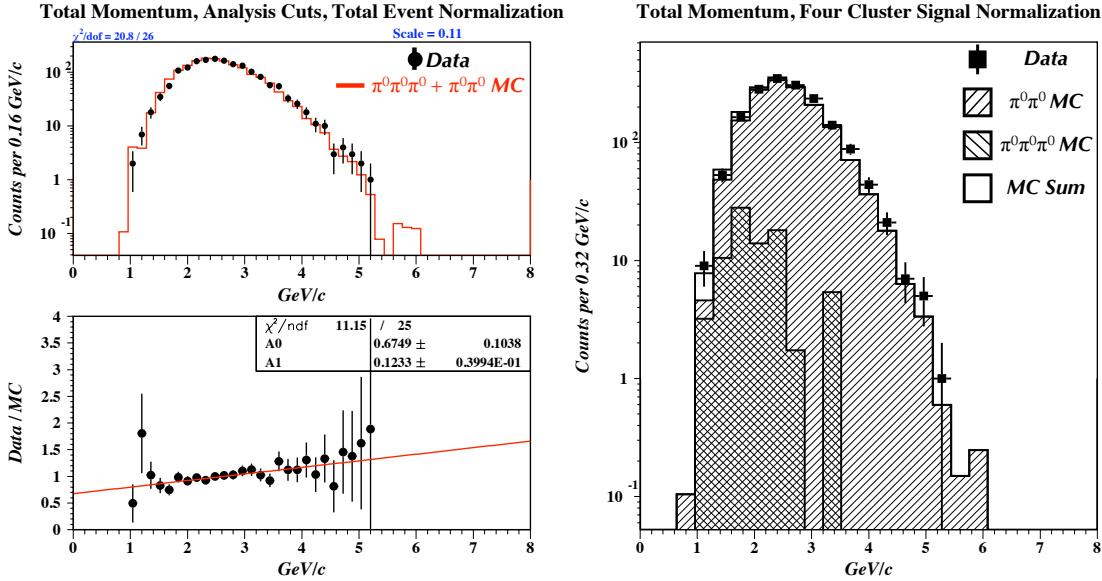


Figure 4.24: The Kaon total momentum spectrum derived from $K_L^0 \rightarrow \pi^0\pi^0$. All analysis cuts are imposed.

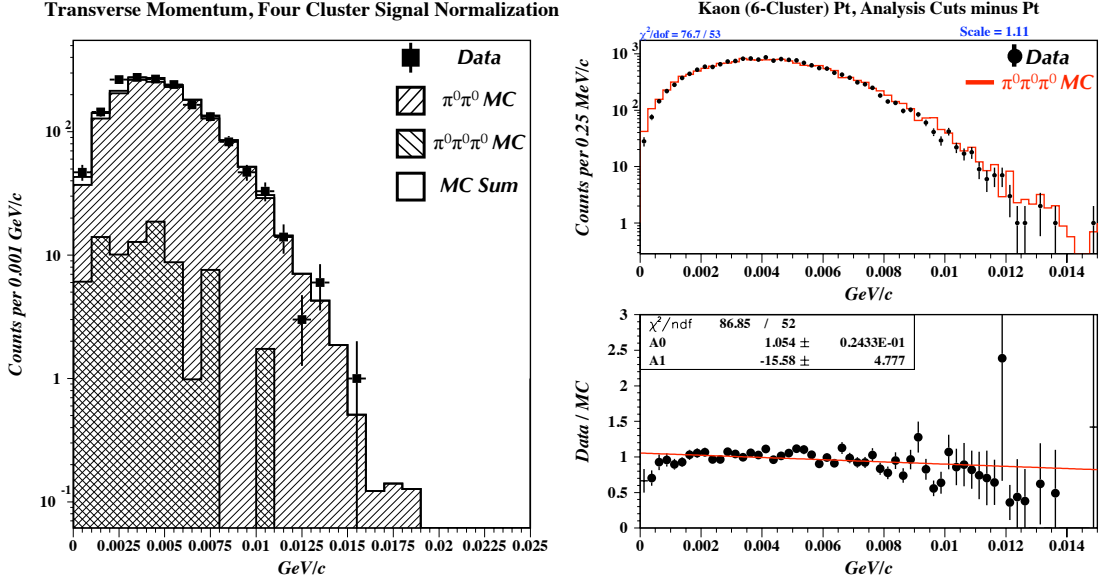


Figure 4.25: Shown above on the left is the four cluster transverse momentum, *normalized by signal*. The figure on the right is the six cluster transverse momentum. All analysis cuts are applied except for P_T .

axis (as opposed to the actual beam axis). Because the radial decay vertex for any given is dependent on both beam divergence and the Kaon momentum, we prefer to use the position at C6, projected back from the reconstructed vertex along the momentum vector. The distribution has the added benefit of being directly related to the MC event generator, which seeds events at C6. The functional form of that generator (probability density) is:

$$F = \frac{N}{1 + \exp((r - x_0)/s) + p_0 + p_1 \times r}, \quad (4.7)$$

where N is a normalization parameter and s represents a slope value. See Figure 4.26 for comparisons of the radial position at C6 for four and six cluster events between data and the MC. Additionally, in Figure 4.27 we show the related distribution of calorimeter hit patterns.

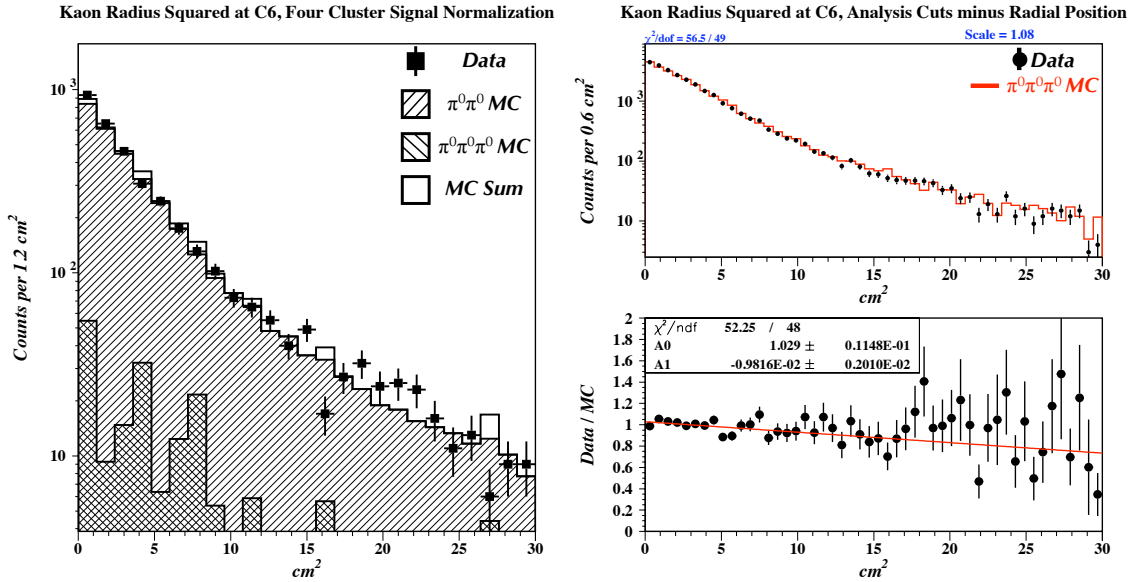


Figure 4.26: Shown above on the left is the four cluster Kaon radial position at C6, *normalized by signal*. The figure on the right is the six cluster radial position at C6. All analysis cuts are applied except for the radial cut.

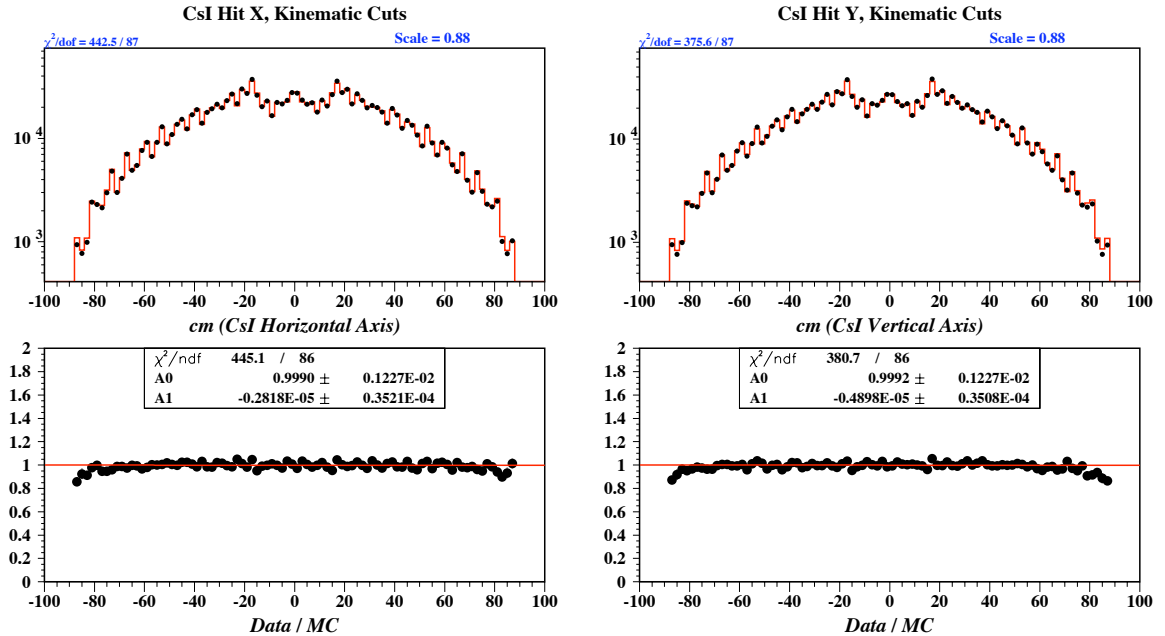


Figure 4.27: The reconstructed photon x and y positions for $K_L^0 \rightarrow 3\pi^0$ decays. The “jagged” nature of the distributions is due to a tendency in the clustering algorithm to pull reconstructed hit positions towards the center of crystals.

4.11.4 Beam Angles

In Figure 4.28 we show the beam direction angles, θ and ϕ , respectively, under analysis cut conditions for $K_L^0 \rightarrow 3\pi^0$ events, while Figure 4.29 shows the same plots for $K_L^0 \rightarrow \pi^0\pi^0$. The beam angle directions are defined in the “usual” way - θ is the *polar* angle between the beam and the z -axis, while ϕ is the *azimuthal* angle and shows the angular distance from the x -axis.

The polar angle is coupled explicitly to the radial position at C6 in the generation stage. As such, the generating function for θ was modified only slightly during the tuning process and it is fortunate that re-weighting to adjust the radial position also seemed to leave θ in decent agreement.

The azimuthal angle is computed separately and underwent major revisions during the tuning process. During Run I, the beam was more or less co-linear with the z axis, but in Run II, the beam points decidedly off axis - down and to the side. This change produced a major shift in the azimuthal angle and the previously almost flat spectrum was replaced with a sinusoidal function (required for continuity across the $\phi = 2\pi$ boundary, although continuity of the first derivative is not explicitly enforced).

4.11.5 Pairing χ^2

The pairing χ^2 , also called the χ_z^2 is defined as:

$$\chi_z^2 = \sum_{i=1}^n \frac{(\bar{z} - z_i)^2}{\sigma_i^2}, \quad (4.8)$$

where $i = 1, \dots, n$ counts the number of π^0 's and σ is computed by straightforward error propagation through Equation 2.1. See Figures 4.30 and 4.31 for data-to-MC comparisons for the pairing χ^2 distributions for $K_L^0 \rightarrow 3\pi^0$ and $K_L^0 \rightarrow \pi^0\pi^0$ respec-

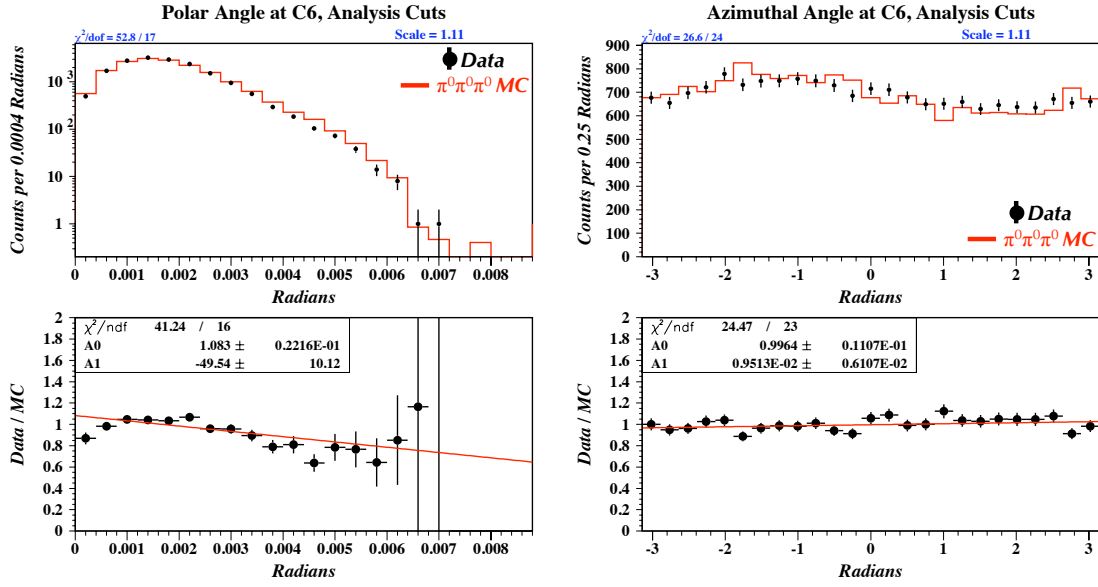


Figure 4.28: The polar and azimuthal beam angles at C6 with analysis cuts imposed for $K_L^0 \rightarrow 3\pi^0$.

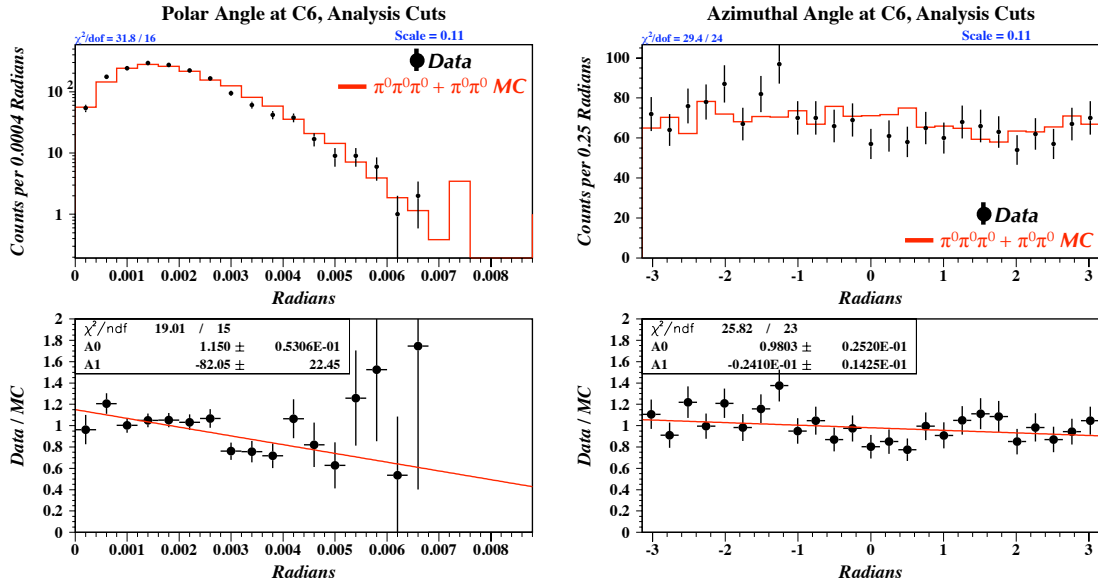


Figure 4.29: The polar and azimuthal beam angles at C6 with analysis cuts imposed for $K_L^0 \rightarrow \pi^0\pi^0$.

tively. The pairing χ^2 distribution is sensitive to a host of resolution effects in the calorimeter and was one of the most difficult distributions to get agreement in until we added the single crystal smearing functions described above. It is one of the few distributions sensitive to the calorimeter that is insensitive to energy scale errors, and this is the primary reason we were ultimately able to find agreement here.

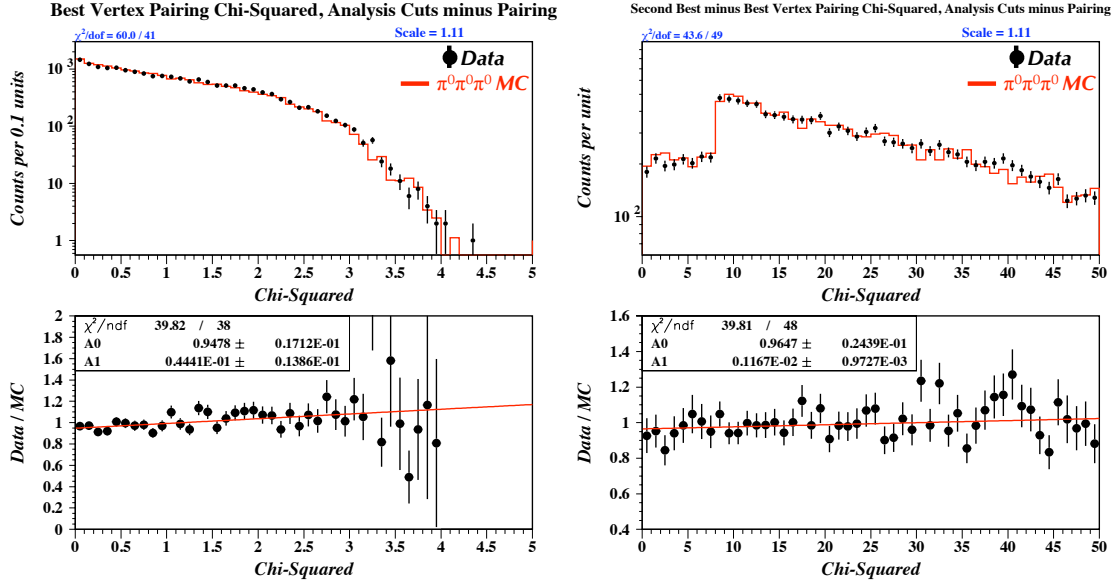


Figure 4.30: The lowest (best) χ_z^2 value (left) and the second-lowest minus the lowest χ_z^2 value (right) with all analysis cuts applied but those on the χ_z^2 values.

4.12 Flux and Systematic Errors

4.12.1 Systematic Error

The systematic error is defined according to the equation:

$$\text{Syst. Err.}^2 = \frac{\sum_{i=\text{All Cuts}} (F_i/A_{Data,i})^2}{\sum_{i=\text{All Cuts}} (1.0/A_{Data,i})^2}, \quad (4.9)$$

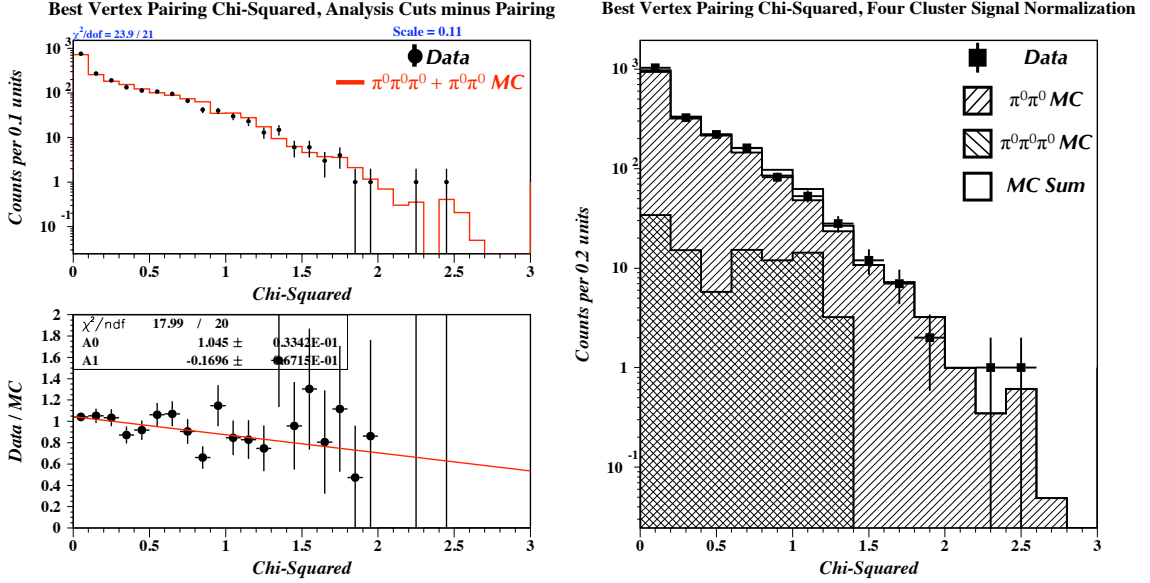


Figure 4.31: The lowest (best) χ_z^2 value with all analysis cuts applied but those on the χ_z^2 values for $K_L^0 \rightarrow \pi^0\pi^0$.

where F is a defined quantity named the *fractional difference*. The *exclusive acceptance*, A , is defined as the acceptance of a given cut with *all other cuts applied*. The fractional difference, F , is the exclusive acceptance of data minus the acceptance of the MC, divided (normalized) by the acceptance of data. The acceptance weighted fractional difference (F/A) is that value divided again by the exclusive acceptance of data. This definition for systematic error was chosen to make explicit the dependence of the accuracy of the flux on the ability of the MC to replicate data.

The fractional differences are summed in quadrature in order to avoid ambiguity in negative contributions. They are further weighted by the exclusive acceptance of the cut in data in order to keep poorly performing, but less relevant cuts (the Back-Anti Scintillator cut for $3\pi^0$, for example) from skewing a meaningful estimate of the error. See Table 4.12 for a summary of the largest contributions to the systematic error for $K_L^0 \rightarrow \pi^0\pi^0$.

See Figures 4.32, 4.33, 4.34, 4.35, 4.36, and 4.37 for vectors showing the fractional

differences in acceptance for $K_L^0 \rightarrow \gamma\gamma$, $K_L^0 \rightarrow \pi^0\pi^0$, and $K_L^0 \rightarrow 3\pi^0$. For each entry in the figures, a value of zero is ideal but realistically indicates the cut had zero impact on the acceptance in both MC and data.




Figure 4.32: The exclusive acceptance fractional differences, $(\text{Data} - \text{MC})/\text{Data}$, for the $K_L^0 \rightarrow \gamma\gamma$ Add-BG MC. A *positive* value indicates the acceptance *loss* is *higher* in the MC than it is in data. The acceptance-weighted sum in quadrature is 7.6%.

4.12.2 The CsI Veto and the Systematic Error

In examining Table 4.12, the total contribution to the systematic error by the CsI veto is striking. The acceptance *loss* in the MC is much higher than in the data. The CsI veto in its current incarnation was introduced relatively late in the analysis. As a consequence, some of the features that lead to the discrepancy escaped early detection.

Fundamentally, there are three basic processes that can contribute to a discrepancy in the CsI veto: 1) an energy scale error, 2) a problem in the simulation of the timing information, and 3) improperly modeled non-linearities in the response. Because the




Figure 4.33: The exclusive acceptance fractional differences, $(\text{Data} - \text{MC})/\text{Data}$, for the $K_L^0 \rightarrow \gamma\gamma$ Pure MC. A *positive* value indicates the acceptance *loss* is *higher* in the MC than it is in data. Note that these values do not include the effects of accidental activity. The acceptance-weighted sum in quadrature is 6.8%.




Figure 4.34: The exclusive acceptance fractional differences, $(\text{Data} - \text{MC})/\text{Data}$, for the $K_L^0 \rightarrow \pi^0\pi^0$ Add-BG MC. A *positive* value indicates the acceptance *loss* is *higher* in the MC than it is in data. The acceptance-weighted sum in quadrature is 7.8%.




Figure 4.35: The exclusive acceptance fractional differences, $(\text{Data} - \text{MC})/\text{Data}$, for the Pure MC. A *positive* value indicates the acceptance *loss* is *higher* in the MC than it is in data. Note that these values do not include the effects of accidental activity. The acceptance-weighted sum in quadrature is 7.3%




Figure 4.36: Shown above is the exclusive acceptance fractional differences, $(\text{Data} - \text{MC})/\text{Data}$, for the Add-BG MC. A *positive* value indicates the acceptance *loss* is *higher* in the MC than it is in data. The acceptance-weighted sum in quadrature is 7.7%.




Figure 4.37: Shown above is the exclusive acceptance fractional differences, $(\text{Data} - \text{MC})/\text{Data}$, for the Pure MC. A *positive* value indicates the acceptance *loss* is *higher* in the MC than it is in data. Note that these values do not include the effects of accidental activity. The acceptance-weighted sum in quadrature is 6.9%.

CsI veto is “absolute” - it does not scale with cluster energy, a scale error can lead to over or under-vetoing. The timing information enters the picture because we use timing windows to minimize losses due to accidental activity. If we are not controlling accidental effects correctly, the time window choice can lead to a discrepancy in the behavior of the veto. Finally, the non-linearity manifests in the resolution as expressed in Equation 4.3. Because we cannot measure our resolution precisely through use of a drift chamber or some similar technology, we have no way to be sure we are matching the behavior of the CsI in our MC. We can only rely on second order plots (particle mass widths, etc.) which are not as reliable. If we are not correctly matching the resolution of the CsI at the single crystal level, this will manifest in the veto behavior, especially for low energy deposits where the constant term is dominant.

Unfortunately, we have one additional contribution. Before beginning mass-production of the Test 10 MC, a member of the E391a collaboration (not the author) modified

the geometry of the steel tank that holds the CsI array to install a curved surface across the downstream face. This was done to more accurately check backgrounds from “backwards-going” pion production downstream of the CsI. (This background turns out to be negligible.) While making this change, the positions of CC03 and CC04 were inadvertently shifted downstream by about 20 cm. This bugged version of the MC was then installed into the general E391a code base without being thoroughly checked and without informing the group of the change. See Figure 4.38 for an illustration. Surprisingly, the energy overlays between MC and Data for CC03 and CC04 matched fairly well even with the bugged geometry. Therefore, the bug was not discovered until it was much too late and many months of CPU time had already been invested in producing a MC sample.




Figure 4.38: The “bugged” MC geometry employed in Test 10. CC03 is highlighted with an arrow.

Because some MC processes are very slow ($K_L^0 \rightarrow 3\pi^0$ and halo neutron events in particular), we did not have the opportunity to reproduce those samples in any meaningful way. $K_L^0 \rightarrow 3\pi^0$ is very slow, for example, because there are at least six particles in the final state. Furthermore, very high statistics are needed for comparison with data owing to the high branching fraction of the mode. In contrast, $K_L^0 \rightarrow \pi^0\pi^0$ is much faster because there are fewer particles to follow in the simulation. What’s

more, the relatively low branching fraction of the decay means we need not run as many events to have a sample competitive in size with our data set. Unfortunately, even relatively quick to simulate modes and those with small branching ratios were impossible to completely reproduce in a reasonable amount of time.

However, we were able to produce MC samples of $K_L^0 \rightarrow \gamma\gamma$ and $K_L^0 \rightarrow \pi^0\pi^0$ that were both roughly the same size as our data set to check the impact of the bug on our analysis. We found that the bug, which affected the geometry around the beam hole, accounted for a significant fraction of the systematic error. With the bug fixed, the contribution from the CsI to the systematic error falls from 6.2% to 4.6% and the overall systematic error falls from 7.3% to 5.4%. Several other cuts were impacted positively by the change, most noticeably, the cut on photon hit position (which is very sensitive to the performance of the simulation near the beam-hole). The contribution to the systematic error for that cut fell from 1.2% to 0.33%. However, even with this geometry fix, the systematic error contribution of the CsI is still dominant and quite large at 4.6%.

Of the three fundamental contributions to the error, we were first able to dismiss the timing simulation discrepancies as a major contribution. This can be understood in two ways. First, this discrepancy exists both between the Pure and the Add-BG MC. Which is to say, the acceptance loss is higher in both forms of the MC. If the problem was only with the CsI time window cut, we would expect the acceptance loss to be lower in the Pure MC because timing cuts are almost irrelevant in that iteration of the MC. The fact the acceptance loss is still higher suggests the problem has to do with the energy behavior of the CsI. The second way to understand this is by studying the Add-BG MC and noting that if the acceptance loss is higher in the MC, that means that if the problem is in the timing, our time window in the MC is too wide. However, we know from the BA that our accidental trigger is likely looking at

events off the intensity peak (in terms of protons on target). This means the energy deposited in the CsI should be, if anything, lower in the MC if the problem is with timing (and therefore accidentals). Finally, even fairly drastic changes (cutting the window in half, for example) to the time window do not change the acceptance of the veto by more than a few percent of itself. So while even though there is clearly a discrepancy in the timing simulation, one that we are not surprised to find, the impact is ultimately too small for the effect to be considered the dominant source of error.

Effects from non-linearities may be contributing. Because we cannot establish our resolution using data directly, it is impossible to rule out. However, as can be seen throughout this document, many higher-order photon quality distributions match well between the MC and data. If there was some sort of problem with our estimate of the resolution, we would expect it to surface in those plots as well.

This means then, that the most likely source of this discrepancy is an energy scale error in the CsI calorimeter. Our relative energy calibration between crystals is quite good due to the $K_L^0 \rightarrow 3\pi^0$ kinematic fitting procedure discussed in Chapter 3. However, the π^0 -Run calibration and cosmic ray calibration procedures were used to establish the absolute energy scale, and the two methods did not agree to better than 7%. Because the CsI veto makes cuts on energy deposition in single-hit crystals (crystals not part of a photon cluster), it is sensitive to the overall energy scale.

When the CsI veto is employed, it introduces a differential effect between data and MC that shifts the average energy of the Kaon between the two. This shifts the momentum spectrum and therefore also the z spectrum. However, it leaves distributions insensitive to the energy scale (like the pairing χ^2 and radial position distributions) unchanged. Energy scale uncertainties are common in particle physics, although they are more often seen in hadronic calorimeters than electromagnetic. The key problem

in E391a is the lack of a resolving magnet and proportional wire chamber.

Without those technologies, we cannot track and measure the momentum of charged particles (e.g. electrons from Ke3 decays) and so lack a common tool for detector calibration. This is unfortunate, but it is simply a price we must pay for a clean fiducial region.

We note that a scale error cannot be established on the basis of the plots shown thus far. This is because with the CsI veto applied, we see distortions in both the momentum and vertex spectra (but good agreement in both without the veto). If our energy scale in data is low (as the current situation would require), this would shift the vertex spectrum closer to the CsI. This would require a higher momentum in the MC to find agreement in the decay vertex spectrum. However, the energy of the Kaon would also be lower, and that would require a lower momentum in the MC to match the spectra. These forces compete and their effects are further convoluted by the geometrical acceptance. Our only hope to establish a scale error is by finding a region of the momentum spectrum that agrees well between the MC and data and compare the vertex spectra for those events.

Fortunately, we are able to do that. See Figure 4.39. If we restrict our attention to the region, $2 \text{ GeV}/c < |P| < 2.8 \text{ GeV}/c$, we find the momentum spectrum overlay slope is consistent with zero. In this region though, the z -decay spectrum is clearly not consistent with zero, and, furthermore, is positive. This is supportive of the decay spectrum for data being shifted closer to the CsI face, and therefore consistent with a low energy scale error in the data.

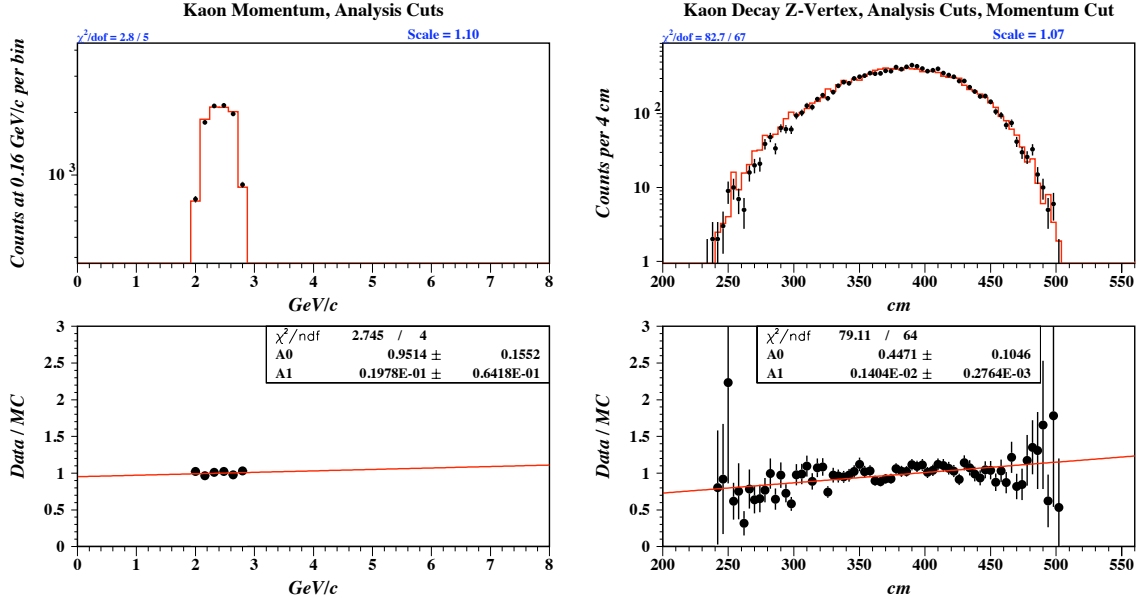


Figure 4.39: The $|P|$ and z -decay spectra with the momentum restricted such that $2 \text{ GeV}/c < |P| < 2.8 \text{ GeV}/c$. The positive slope for the z -decay spectrum overlay supports a low energy scale in data.

4.12.3 Flux Values

In Table 4.13, we provide a summary of our current flux (number of K 's decayed within the fiducial) estimates. We compute our flux according to:

$$\text{Flux} = \frac{N_{\text{Signal}}}{\text{Acceptance} \times \text{Branching Ratio}} \rightarrow F = \frac{N_S}{A \times \mathcal{B}}. \quad (4.10)$$

Our acceptance is defined as:

$$A = \frac{N_{\text{Obs}}}{N_{\text{decay}}} = \frac{N_{\text{Obs}}}{N_{C6} \times p_d}, \quad (4.11)$$

where N_{C6} is the number of Kaons at C6 (the generation point in the MC) and p_d is the probability of decay in the fiducial. Note that for modes containing pions, we must correct for the pion branching fraction, but this correction will cancel out in the

total expression:

$$F_{n\pi} = \frac{N_S}{A_{n\pi} \times \mathcal{B}_{n\pi} \times \mathcal{B}_{\gamma\gamma}^n} = \frac{N_S}{\frac{N_{Obs}}{N_{C6} \times p_d \times \mathcal{B}_{\gamma\gamma}^n} \times \mathcal{B}_{n\pi} \times \mathcal{B}_{\gamma\gamma}^n} = \frac{N_S}{A \times \mathcal{B}}. \quad (4.12)$$

The official choice of E391a for *the* flux is that given by the Pure MC estimate for $K_L^0 \rightarrow \pi^0\pi^0$, which is to say 5.13 ± 0.4 billion. The mode $K_L^0 \rightarrow \pi^0\pi^0$ was chosen as the normalization mode on the basis of relatively similar energy scales for photon clusters with $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$. While a two photon final state normalization, i.e., $K_L^0 \rightarrow \gamma\gamma$, makes sense in many respects, because the energy scale ambiguity ultimately dominates our systematic error, we felt it appropriate to try to match that scale as best as possible in our normalization mode.

Detector	Energy Cut	Comments
CC00	2 MeV	
Front Barrel	1 MeV	Inner & Outer Sum.
CC02	1 MeV	
Barrel CV	0.75 MeV	$\sqrt{Up \times Down}$ (w/TDC).
Inner Main Barrel	1 MeV	$\sqrt{Up \times Down}$ (w/TDC).
Outer Main Barrel	1 MeV	$\sqrt{Up \times Down}$ (w/TDC).
Outer CV	0.3 MeV	
Inner CV	0.7 MeV	
CC03	2 MeV	
CsI: S-Hit Close	10 MeV	Crystal $d < 17$ cm from closest cluster.
CsI: S-Hit Intermediate	$5 - (3/8)(d - 17)$ MeV	Crystal is d cm from closest cluster.
CsI: S-Hit Far	2 MeV	Crystal $d > 25$ cm from closest cluster.
Sandwich Counters	2 MeV	
CC04 Charged Layer	0.7 MeV	
CC04 Calorimeter Layer	2 MeV	
CC05 Charged Layer	0.7 MeV	
CC05 Calorimeter Layer	3 MeV	
CC06	10 MeV	Cerenkov detector with MIP calibration
CC07	10 MeV	Cerenkov detector with MIP calibration
BHCV	0.1 MeV	
BA Scintillator BA Quartz	20 MeV 0.5 MIPs	Sum over layers. Max. layer. (AND logic.)

Table 4.9: The veto cuts. Recall that $1 \text{ GeV} = 10^3 \text{ MeV} = 10^9 \text{ eV}$.

Cut	Values	Comments
Photon Energy Min.	150 MeV	
Photon CsI Hit Position	$17.5 < r < 88$ cm	The inner dimension forms a square around the beam pipe, the outer is radial.
Photon Hit Separation Min.	15 cm	
Photon Fusion Neural Network Min.	0.7	
Pion z -difference Max.	5 cm	
Pion Mass-difference Max.	5 MeV/c ²	
Best z - χ^2 Max.	3	
2nd Best minus Best z - χ^2 Min.	4	
r^2 at C6 Max.	4.0 cm ²	Also cut if $r(z_{decay})$ is less than R(C6).
Decay z	340 $< z <$ 500 cm	
Transverse Momentum Max.	1.25×10^{-4} (GeV/c) ²	

Table 4.10: Kinematic Cuts for Four Cluster Events.

Cut	Values	Comments
γ Energy Min. (Higher Energy γ)	250 MeV	
γ Energy Min. (Lower Energy γ)	150 MeV	
Photon CsI Hit Position	$15 < r < 88$ cm	The inner dimension forms a square around the beam pipe, the outer is radial.
Photon Hit Separation Min.	15 cm	
Photon Fusion Neural Network Min.	0.7	
Photon RMS Max. - Individual	5.2	
Photon RMS Max. - Sum	9.5	
Photon $\Delta\theta$ Min.	-20 degrees	Rec. θ - NN θ
Photon Energy Ratio Min.	0.88	
Acoplanarity Angle Max.	0.1745 radians	
Transverse Momentum Max.	9×10^{-4} (GeV/c) ²	

Table 4.11: Kinematic Cuts for $K_L^0 \rightarrow \gamma\gamma$ Events.

Cut	Pure MC Uncertainty	Add-BG Uncertainty
CsI Veto	6.2%	6.7%
Decay- z	2.3%	2.4%
Radial Position	1.7%	1.8%
Outer CV	1.3%	1.4%
γ Hit Position	1.2%	1.3%
Others (24 Total)	1.9%	1.8%
Total	7.3%	7.8%

Table 4.12: Components of the systematic error for $K_L^0 \rightarrow \pi^0\pi^0$.

Mode	Pure MC Flux	Add-BG Flux
$K \rightarrow \gamma\gamma$	$(5.45 \pm 0.38) \times 10^9$	$(5.43 \pm 0.42) \times 10^9$
$K \rightarrow \pi^0\pi^0$	$(5.13 \pm 0.40) \times 10^9$	$(5.13 \pm 0.43) \times 10^9$
$K \rightarrow \pi^0\pi^0\pi^0$	$(5.02 \pm 0.35) \times 10^9$	$(5.10 \pm 0.40) \times 10^9$
$K \rightarrow \gamma\gamma$ (Geo-Fix)	$(5.05 \pm 0.23) \times 10^9$	
$K \rightarrow \pi^0\pi^0$ (Geo-Fix)	$(5.02 \pm 0.33) \times 10^9$	

Table 4.13: Flux estimates. Here the quoted error is the *total error* and is dominated by the systematic error. The ‘‘Geo-Fix’’ samples refer to estimates made with reduced size samples that were free of the CC03/CC04 geometry bug.

CHAPTER 5

PHOTON VETO INEFFICIENCY

5.1 Critical Summary

In this Chapter we explain a method for veto inefficiency measurement using $K_L^0 \rightarrow 3\pi^0$ with exactly five photons reconstructed in our calorimeter. We then use the method to set an upper limit (un-binned in energy) on the Main Barrel (MB) at 7.66×10^{-4} at the 90% C.L. This result is not strong enough to critically impact the final measurement in E391a, but this method could well prove crucial for any future generation version of this experiment.

5.2 Photon Veto Inefficiency

Photon veto inefficiency poses a potentially crippling threat to a $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ measurement. The reason, as shall be discussed in greater depth in Chapter 6, is that other Kaon decay modes, especially $K_L^0 \rightarrow \pi^0 \pi^0$ can fake a signal event if we lose final state photons. Here, we consider an inclusive definition of photon loss, which is to say that inefficient detection can arise through photonuclear absorption, through too-loose veto thresholds, sampling inefficiency (fluctuation), and photon punch-through.

Photons passing through matter interact primarily through three processes: the photoelectric effect, Compton scattering, and pair production. Below about 1 MeV in lead, for example, the photoelectric effect is the dominant interaction. Right around 1 MeV, Compton scattering dominates, but above that energy, pair production off the nuclear and electron electromagnetic fields begins to dominate, and it is the most important process for energies larger than about 10 MeV. Pair production occurs when a photon interacts with the electromagnetic field of an atom and converts into an

electron-positron pair. These particles then bremsstrahlung and produce photons due to energy losses from interactions with the electromagnetic fields of atomic electrons. Those photons go on to pair-produce and can create an electromagnetic “cascade” of photons and electrons for high enough energy particles. These cascades are what our scintillator detectors observe. The mean free path for pair production is:

$$\lambda_{\text{pair}} = \frac{9}{7}\lambda_{\text{rad}} \quad (5.1)$$

where the radiation length, λ_{rad} , is the distance over which an *electron* loses $1/e$ of its energy through radiation loss (bremsstrahlung) only. The radiation length in grams per centimeter squared is a property of a material, but in units of radiation lengths, we may consider various materials as roughly equivalent. More information is available about these processes in [31].

Because photons are neutral, they are capable of penetration without leaving any ionization signal in our scintillators. This leads to two possible inefficiency sources. First, photons may simply punch directly through a veto detector with probability roughly equal to $\exp(-L_r/\lambda_{\text{pair}})$, where L_r is the thickness of the detector in radiation lengths. Because this quantity is a straightforward calculation, we ensured that all the detector elements of E391a were thick enough to reduce inefficiency from this source to very small levels for a search sensitivity of 10^{-10} . The other inefficiency source though is due to photon conversion deep inside the detector. On average, $\exp(-7/9)$ of the total number of photons incident on a material will convert for each radiation length traversed. This means of course that photons may interact “late” instead of punching through and deposit much less than their total energy. If we have veto thresholds that are loose enough, we will not identify these interactions as photons.

While the interaction is characterized by a small cross-section, direct absorption by the nucleus is a problem for photon detectors. Because excited nuclear states may decay with long half-lives, and may follow decay chains that do not produce particles capable of producing optical photons in our scintillator detectors, it is possible to lose photons completely in this fashion.

All of these background sources are calculable and may be estimated through Monte Carlo (MC). However, it is difficult to completely trust such calculations because they are based on measurements made prior to installation and in some cases rely on cross sections with large uncertainties. Detectors may be damaged in the course of fabrication. Phototubes may underperform, optical fibers can break, and any number of other problems can lead to unanticipated optically “dead” material. It is important to both check for these sorts of problems and verify inefficiency measurements made in lab-bench conditions whenever possible. In this chapter, we present a method for doing just that with kinematic fits of $K_L^0 \rightarrow 3\pi^0$ events with five clusters reconstructed in the CsI. This method has limited power for E391a owing to the energy and position resolution of our calorimeter. However, it is a potentially very important technique for inefficiency estimation verification for future generations of $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ measurements with similar detection strategies.

5.2.1 *Five Cluster Events*

We will use $K_L^0 \rightarrow 3\pi^0$ decays with five clusters in the CsI calorimeter to reconstruct the four momentum of the sixth photon; we then use that information to calculate our photon veto inefficiency. This method was first suggested in E391a for six cluster events, with one cluster “blinded” for CsI calibration purposes [47]. The work in this Chapter is an extension of the results derived there. There are six photons in the final

state under consideration: $\gamma^0, \gamma^1, \gamma^2, \gamma^3, \gamma^4$, and γ^5 , the last of which is missing from the calorimeter and will be referred to as the “missing” photon. The basic strategy is to reconstruct all possible combinations of two π^0 's from the five photons in the calorimeter and keep the pair with the best z -vertex χ^2 , defined as:

$$\chi_z^2 = \sum_{i=1}^2 \frac{(z_{\pi_i^0} - \bar{z})^2}{\sigma_i^2}, \quad (5.2)$$

where σ_i is the error on the z -vertex as computed by the π^0 reconstruction routine (see Chapters 2 and 4.) The remaining fifth photon in the calorimeter shall be referred to as the “spare” photon. We shall employ the following definitions:

- Subsystem 1 \rightarrow The spare photon in the CsI (γ^4). By “spare”, we mean the photon not associated with a pion.
- Subsystem 2 \rightarrow The pair of reconstructed π^0 's.
- Subsystem 3 \rightarrow (1) + (2)
- Subsystem 4 \rightarrow The “missing” photon.

Using the information measured in these subsystems and some reasonable physical assumptions, we shall derive the conditions to employ a one constraint fit to calculate the four-momentum of the missing photon. See Figure 5.1 for a “cartoon” of the topology of the events we use.

We have eight unknowns - the Kaon four-momentum and the missing photon four momentum. We express these as:

$$K_\mu = (K_0, K_1, K_2, K_3) = (P_{Kx}, P_{Ky}, P_{Kz}, E_K) = (K_x, K_y, K_z, K_E), \quad (5.3)$$

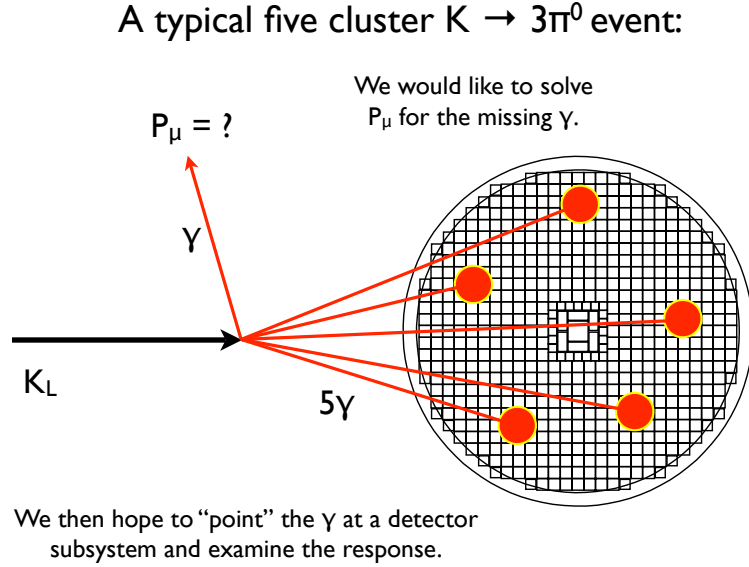


Figure 5.1: A candidate five cluster event for kinematic fitting.

and

$${}^5\gamma_\mu = (P_{5x}, P_{5y}, P_{5z}, E_5) = ({}^5\gamma_x, {}^5\gamma_y, {}^5\gamma_z, {}^5\gamma_E). \quad (5.4)$$

Here the superscript seen in ${}^5\gamma$ references the photon identity. In principle, the kaon vertex is also unknown. However, the z -vertex may be estimated using the best pair of pion vertices found by treating the decay as $K_L^0 \rightarrow 2\pi^0$, with one extra γ . Since the center of energy of the five available clusters is not expected to sit in the same position as the center of energy of all six clusters, the (x, y) -vertex is assumed to sit on the beam axis at $(0, 0)$. This leaves us with a total of eight unknowns.

We have the following equations to solve for these unknowns:

$$\begin{aligned} (K_\mu)^2 &= m_K^2 \\ ({}^5\gamma_\mu)^2 &= 0 \\ ({}^4\gamma_\mu + {}^5\gamma_\mu)^2 &= m_{\pi^0}^2 \\ K_\mu &= \sum_{i=0}^{i=5} i\gamma_\mu \end{aligned} \quad (5.5)$$

where the superscript on terms like ${}^i\gamma$ refer to the i -th photon. Since the μ index ranges from 0 to 3, the last of these equations is really a set of four equations.

Thus we have eight unknowns and seven equations, leaving us with an undetermined system. However, since the Kaon x and y momenta are very small, we will assume they are zero to have six unknowns and seven equations - an over-constrained system. We can solve this system by doing a one-constraint, or 1-C, fit.

With $K_x = K_y = 0$, we may reduce the system.

$$K_x = 0 \implies \sum_{i=0}^{i=4} {}^i\gamma_x = {}^5\gamma_x \quad (5.6)$$

$$K_y = 0 \implies \sum_{i=0}^{i=4} {}^i\gamma_y = {}^5\gamma_y \quad (5.7)$$

$$K_z - \sum_{i=0}^{i=4} {}^i\gamma_z = {}^5\gamma_z \quad (5.8)$$

$$K_E - \sum_{i=0}^{i=4} {}^i\gamma_E = {}^5\gamma_E \quad (5.9)$$

Therefore the four-momentum of the missing photon is solved when the Kaon energy and z -momentum are determined. We may rewrite the constraint on the mass of the missing photon, $({}^5\gamma_\mu)^2 = 0$, in a more useful way by noticing that conservation of four-momentum implies

$$(K_\mu - {}^3P_\mu)^2 = 0, \quad (5.10)$$

where the superscript on ${}^3P_\mu$ refers to the subsystem identity. We may also rewrite our constraint on the missing pion mass, $({}^4\gamma_\mu + {}^5\gamma_\mu)^2 = m_{\pi^0}^2$, as

$$(K_\mu - {}^2P_\mu)^2 = m_{\pi^0}^2, \quad (5.11)$$

where the superscript on ${}^2P_\mu$ again refers to the subsystem identity. We reserve $(K_\mu)^2 = (K_E^2 - K_z^2) = m_K^2$ as our constraint equation for the fit.

We shall now reduce these constraints, beginning with the constraint on the missing photon mass:

$$\begin{aligned}
(K_\mu - {}^3P_\mu)^2 &= (K_E - {}^3P_E)^2 - {}^3P_T^2 - (K_z - {}^3P_z)^2 \\
&= K_E^2 + {}^3P_E^2 - 2K_E {}^3P_E - \\
&\quad {}^3P_T^2 - K_z^2 - {}^3P_z^2 + 2K_z {}^3P_z.
\end{aligned} \tag{5.12}$$

We can define a new effective mass using ${}^3m^2 = {}^3P_E^2 - {}^3P_T^2 - {}^3P_z^2$. This allows us to simplify our condition on the photon mass further:

$$K_E^2 - 2K_E {}^3P_E + 2K_z {}^3P_z - K_z^2 + {}^3m^2 = 0. \tag{5.13}$$

Now we apply a similar process to the condition on the missing pion mass:

$$\begin{aligned}
(K_\mu - {}^3P_\mu)^2 &= m_{\pi^0}^2 \\
&= (K_E - {}^2P_E)^2 - {}^2P_T^2 - (K_z - {}^2P_z)^2 \\
&= K_E^2 + {}^2P_E^2 - 2K_E {}^2P_E - {}^2P_T^2 - \\
&\quad K_z^2 - {}^2P_z^2 + 2K_z {}^2P_z.
\end{aligned} \tag{5.14}$$

By introducing another effective mass, we simplify this condition as well:

$$K_E^2 - 2K_E {}^2P_E + 2K_z {}^2P_z - K_z^2 + {}^2m^2 - m_{\pi^0}^2 = 0. \tag{5.15}$$

Since we now have two equations for K_E and K_z , we may solve for the Kaon energy and z -momentum. The solutions are a bit lengthy, but are roughly of the form:

$K_E = A \pm \sqrt{C_1}$ and $K_z = B \pm \sqrt{C_2}$. Note that initially, $C_1 = C_2$.

5.2.2 The 1-C Fit

In practice, the relation $K_E^2 - K_z^2 = m_K^2$ will not be true when using the final formulations for K_E and K_z from the previous section. The reasons for this are two-fold. First, there are errors in the reconstruction of the cluster energies and positions. Second, additional errors are introduced by assuming $(x, y) = (0, 0)$ and $K_x = K_y = 0$. Therefore, we must fit for the best values of K_E and K_z by allowing the values they were derived from to vary, subject to the constraint that $C = 0 = K_E^2 - K_z^2 - m_K^2$.

Since the four-momentum of the missing $2\pi^0$ system and the Kaon vertex are assumed to be correct, the constraint varies as a function of the coordinates x, y , and E of the “spare” cluster in the CsI array. Therefore, we first compute the derivatives: $\partial_x C$, $\partial_y C$, and $\partial_E C$. Since the functions for K_E and K_z are very complicated, these derivatives are computed numerically. For example,

$$\frac{\partial C}{\partial x} = \frac{C(x + \delta x) - C(x - \delta x)}{2\delta x}. \quad (5.16)$$

Here, the error introduced by the numerical approximation is proportional to $(\delta x)^3$ and is negligible for small δx . In this analysis, $2\delta x = 2\delta y = 0.1\text{cm}$ and so $\delta x/x$ was small enough to safely neglect additional corrections. Here δE has been taken to be 1 MeV. The minimum cluster energy used in this study was 70 MeV and the average cluster energy was over 250 MeV. It is not possible to choose any of these δ values to be arbitrarily small without running into problems with machine precision. If we operate on the assumption that neglecting the numerical error is safe, the errors on the derivatives are then simply: $\sigma_x \partial_x C$, $\sigma_y \partial_y C$, and $\sigma_E \partial_E C$, where σ_i is the measurement error on i . We write the total gradient squared as $(\nabla C)^2 = (\partial_x C)^2 + (\partial_y C)^2 + (\partial_E C)^2$.

We use these derivatives to iteratively scan the x, y , and E values for the spare gamma. We shift each of these variables by an amount δq_i for $i = x, y$, and E :

$$\delta q_i = \frac{-C\sigma_i\partial_i C}{(\nabla C)^2}. \quad (5.17)$$

Which is to say $x \rightarrow x + \sigma_x\delta x$, $y \rightarrow y + \sigma_y\delta y$ and $E \rightarrow E + \sigma_E\delta E$, and then the constraint value C is recomputed. This is repeated iteratively until $C < 0.0005 \text{ GeV}^2$ or until twenty iterations pass, whichever comes first. (For clean events with good starting values, convergence is rapid and usually only requires a few iterations.) Recall that Equation 5.15 employed two different constants C_1 and C_2 in the discriminant of the solution. The reason for that now becomes clear. As the iterative fit procedure is executed, C_1 and C_2 do not remain equal and instead evolve in different directions. This feature is how we pick between the plus and minus-solutions of those quadratic equations. In principle, there appears to be two valid solutions. However, it often happens that either C_1 or C_2 eventually become negative in the course of the fit. The partners of these solutions are ultimately the solutions we keep - i.e., when the “plus solution” discriminant is negative, we keep the “minus solution” and vice versa. This costs statistics because we lose all the cases where both the plus and minus solutions have real valued discriminants. However, with no other way to choose between the solutions, this was the procedure employed.

The plus solution tends to be higher energy and the photons in question point either down the beam-hole or into the CsI. The minus solution tends to be lower in energy and the photons are oriented in more transverse directions. They point at the outer rim of the CsI and into the Main Barrel. Because the Main Barrel is both the largest detector (and so the easiest detector to point at) and the most important with regards to a rigorous photon veto inefficiency, we will choose to focus on the minus

solution for the remainder of this document.

We first test this algorithm on MC data where we substitute the true calorimeter hit position and incident energy for the reconstructed values and screen the sample to remove fusion events. It can be difficult to summarize the performance of the fit with low-dimensional figures of merit. However, the angle cosine between the true and reconstructed four momentum of the missing photon provides a representative gauge of the performance of the routine in compact form. See Figure 5.2 for the distribution of angle cosines when no selection cuts are imposed. As can be seen, the ability to reconstruct the direction of the missing photon is good but not perfect. Fits fail even with perfect clustering information because we have assumed the vertex of the Kaon sits at $(0, 0)$ in the (x, y) plane and that the transverse momentum is zero. Neither assumption is strictly true and this sort of baseline plot provides a sense for how much those assumptions distort the functionality of the fit.




Figure 5.2: Angle cosine distributions for the two solutions to the Kaon energy (recall that Equation 5.15 is quadratic). In these plots, the azimuthal ϕ angle has been integrated out.

5.2.3 Five Cluster Event Reconstruction with the E391a Calorimeter

We now quantify how well this technique performs on real E391a data. The central concept behind using this method for measuring inefficiency is pointing a photon and predicting which detector element the photon hits. We then examine the response of the detector to classify whether the photon is detected or not. If the event fails to pass this qualification, there are two possibilities: 1) the detector exhibited a true inefficient event, or 2) we did not correctly point the photon.

To guard against the second case, we employ tight photon vetoes for all subsystems away from the pointed one, but we must still know how accurate the pointing fit is. We first note, as shown in Figure 5.3, that the accuracy of the energy reconstruction with “E391a-level” resolution is simply not good enough to do an analysis with small energy bins. We have put E391a in scare quotes above to indicate that this study was done relatively early in the Run II analysis using a special “fast” MC (with electromagnetic showers turned off). The energy and position were smeared according to Table 5.1. Recall that the real energy and position resolution are best illustrated in Figure 2.10. Also in Table 5.1 we provide a set of resolution parameters labeled as “KTeV.” These parameters were investigated because it is likely the KTeV CsI array will be utilized by the next generation $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ experiment [48] [49].

	σ_x & σ_y (cm)	σ_E (MeV)
“KTeV” CsI	0.1	0.5
“E391a” CsI	1.513	7.3

Table 5.1: Energy and position resolution values applied to MC cluster reconstruction during the 1-C Fit study. These values do not align precisely with the true resolution of the two experiments, but are quite close [48].

In order to gauge the power of this technique, we must use it with the full MC. It is only in this environment that we can apply the same cuts we use in real physics

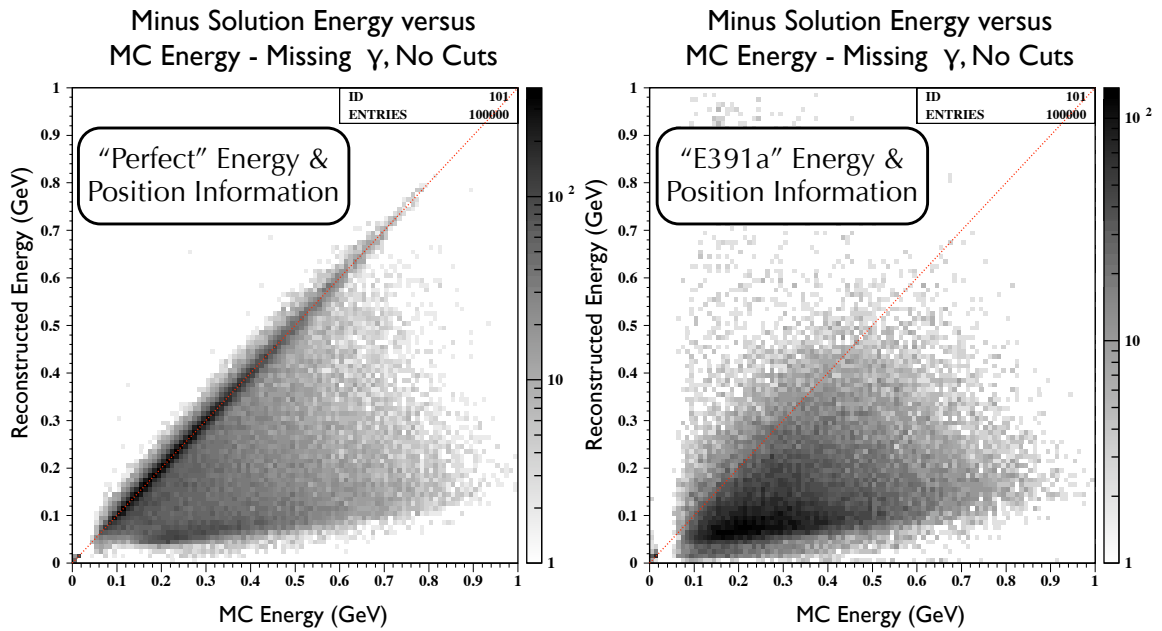


Figure 5.3: Reconstructed (fitted) energy versus true MC energy for the missing photon with no event selection cuts for perfect and actual experiment quality cluster information resolution. No set of selection cuts dramatically improves the energy reconstruction accuracy of the routine for our level of position and energy resolution (see Table 5.1). This means that with E391a data, the best we can do is fit for the reconstructed angle of the missing photon - the energy resolution of that photon is too poor to allow for a small energy bin analysis.

data analysis. We apply a set of veto and kinematic cuts consistent with those listed in Chapter 4 for $K_L^0 \rightarrow 3\pi^0$ decays, with some slight modifications to account for the fact that we only have two pions reconstructed in the final state. Additionally, we cut on the 1-C Fit χ^2 formed from the original spare cluster position and energy and the fitted spare cluster position and energy. See Figure 5.4 for a comparison of this distribution between real five cluster data and our $K_L^0 \rightarrow 3\pi^0$ MC.




Figure 5.4: 1-C Fit χ^2 for the “minus” solution (see Equation 5.15) with a cut on the plus/minus solution flag imposed (the discriminant for only either the plus or minus solution may be real valued).

With all cuts are imposed, we can see the performance of the routine in Figure 5.5 in successfully predicting hits to the Main Barrel. The performance is degraded relative to the ideal case mostly due to photon fusion backgrounds. Photon fusion has been little discussed in this Chapter, but it is a toxic background in five cluster studies. In order to control fusion, we used a method first suggested in the KTeV experiment [50]. For each event we assume the cluster with the highest energy is actually a fused pair of clusters. We then divide the cluster into two with identical (x, y) positions and shared the energy between the two in steps of ten percent from ninety down to fifty. For each sharing scenario, we ran our usual $K_L^0 \rightarrow 3\pi^0$ reconstruction routine and stored the mass and pairing χ^2 (see Equation 4.8). We repeat this procedure for all five clusters and then sort all the solutions by pairing χ^2 .

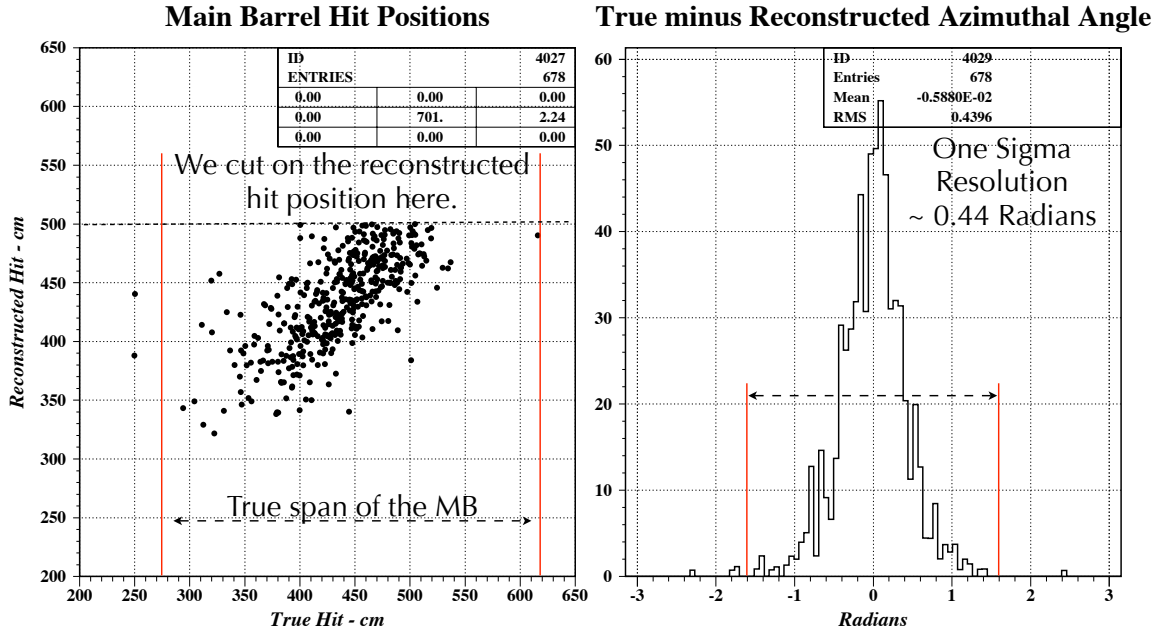


Figure 5.5: The reconstructed versus true hit positions in the Main Barrel and azimuthal angles for the missing photons in the full Run II MC with all analysis cuts imposed. The MC sample used here represents about 10% of the real number of $K_L^0 \rightarrow 3\pi^0$ events in Run II.

In Figure 5.6 we show how this fusion cut operates on $K_L^0 \rightarrow 3\pi^0$ MC and data,

and on $K \rightarrow \pi^0 \pi^0 \gamma$ MC and data. (We studied $K \rightarrow \pi^0 \pi^0 \gamma$ along with the fitting procedure but achieved a sensitivity no better than published limits for the decay.)

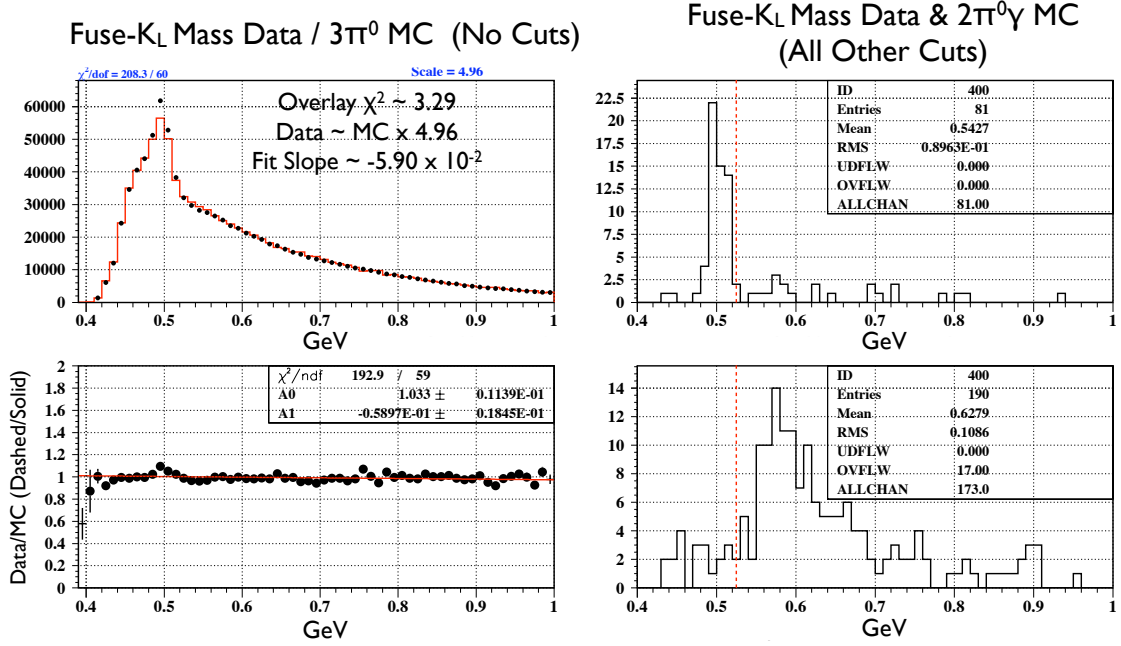


Figure 5.6: The reconstructed mass using the special five cluster fusion finder. On the left we compare $K_L^0 \rightarrow 3\pi^0$ MC (solid) and data (dashed). On the right we compare data (top) to $K \rightarrow \pi^0 \pi^0 \gamma$ MC (bottom) to show how the cut operates on true five photon final state events. Those events are pushed to high mass while true fusion events are reconstructed at the Kaon mass.

Our figure of merit for detection is the integrated energy over the targeted module of the Main Barrel and the neighboring seven to either side (roughly π radians integrated). While this is a wide window, it almost guarantees that should we choose an energy threshold of 1 MeV, that “misses” in the MC are due entirely to mis-pointed photons. We can expect a similar probability of mis-pointing in data and try to establish an inefficiency based on any excess of events below the energy sum threshold. See Figure 5.7 for a comparison of these distributions in the MC and data.

The fitting failure rate in the MC after all cuts with the sample available was $2.31 \pm 0.58_{\text{Stat}}\%$ while the sub-threshold rate is data was 1.95%. In terms of events,




Figure 5.7: The observed energy spectra in the MB for the five cluster MC (left) and data (right). Note that the low energy spectra differs due to the behavior of the TDC threshold: ~ 1 MeV in data and 0 MeV in the MC.

this translates to a mis-pointing expectation of 78.4 ± 19.7 events, but only 64 sub-threshold events observed. If we take this observation as a lack of inefficient “signal,” we can write, using a formula from [51], with $\Delta N = N_{Signal} - N_{Background}$:

$$\begin{aligned}
 2.3 \times \left(1 + \frac{2.3}{2} \times \left(\frac{\sigma_{MC}}{R_{MC}} \right)^2 \right) \times \frac{1}{\Delta N} &= 2.3 \times \left(1 + \frac{2.3}{2} \times \left(\frac{0.0058}{0.0231} \right)^2 \right) \times \frac{1}{3129}, \\
 &= 7.66 \times 10^{-4},
 \end{aligned}
 \tag{5.18}$$

for a limit on our inefficiency of 7.66×10^{-4} at the 90% C.L.

5.2.4 Opportunities at Future Experiments

A limit of essentially 10^{-3} is far from cutting-edge in terms of what can be measured from beam tests and what may ultimately be required for detecting a large sample of $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ events. It is worth examining where this method fails and why the useful event yield is so low.

The largest problem for E391a with regards to these questions is due to the resolution of the CsI array. As can be seen in Figure 5.8, a narrow, well-collimated beam also contributes to a high event yield, but that factor is dominated by the CsI energy and position resolution. In that figure, we show the results of a binned radius study. Rather than speculate on what future beam profiles might look like, we instead use rings of evenly distributed events. These rings are generated by choosing numbers randomly from a uniform distribution between r_1 and r_2 , and then choosing the azimuthal position randomly from a uniform distribution between zero and 2π . The Kaons in this study had zero transverse momentum. When comparing the performance of the E391a CsI (7 cm \times 7 cm face) blocks to the KTeV array (2.5

cm \times 2.5 cm face) blocks, we can see that the event yield for this routine climbs by a full order of magnitude at low beam radii. Fusion events were screened in this study, so any additional advantages in fusion event removal also improve the well-reconstructed event yield. Taken together, these advantages suggest inefficiencies of 10^{-4} or 10^{-5} might be measurable in barrel veto analogs at future $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ experiments. Furthermore, the reconstructed angle resolution might improve enough to allow reliable studies of other detector subsystems that were too small to be reliable pointing targets in E391a.




Figure 5.8: Here we compare the performance of the 1-C Fit (minus solution) for a hypothetical higher resolution experiment (here, employing the KTeV CsI array) and E391a in radial bins. Each radial bin represents evenly distributed randomly generated $K_L^0 \rightarrow 3\pi^0$ decays inside rings defined by the upper and lower bounds of the bin. So, the first bin is a circle with a radius of 1 cm. The second bin is a ring with an outer radius of 2 cm and an inner radius of 1 cm, etc. “Pass Percentage” is the number of events passing the C-Fit cuts divided by the total number of events (photon veto cuts are not imposed). “Angle Percentage” is the number of events reconstructed with angle cosines within 2% of the correct value (i.e., greater than 0.96) divided by the total number of events. “Tolerance Percentage” is the number of events that were reconstructed with angle cosines correct to within 2% and energy reconstructed correctly to within 10 MeV divided by the total number of events.

CHAPTER 6

BACKGROUNDS AND EVENT SELECTION

6.1 Critical Summary

Here we present the total background estimation for E391a. In general, there are two classes of potential background events in E391a: 1) those associated with the interaction of neutrons in the beam halo with detector subsystems, and 2) those associated with other decays of the K_L^0 . As the author's primary focus dealt with backgrounds from Kaon decays, those events will be covered in greater depth. Our total background prediction is 0.45 ± 0.13 events, with 0.17 ± 0.11 events from Kaon decays.

6.2 Event Selection

Event selection in E391a was driven by background rejection. (We, of course, tried to preserve our signal acceptance to the greatest degree possible.) Because our background environment was characterized by sources of widely varying origins, we decided not to try to build a single, total-background simulation. Indeed, some background sources were almost impossible to replicate with simulation anyway. As a consequence, event selection optimization was mostly done by first isolating a background source, devising a cut to remove it, and tightening the cut until the background in question was reduced to a non-threatening level. Because many of the cuts are truly targeted at one background, many of them end up correlated with regards to other backgrounds. This is unfortunate in terms of lost acceptance, but unavoidable in the sense that each cut is needed to defeat at least one background topology. In this Chapter we will describe the event selection criteria and the backgrounds they

are designed to reject. It is difficult to discuss one topic without the other, so the presentation is often intertwined.

6.2.1 Kinematic Cuts

In general, we defined kinematic cuts as those involving two or more particles. Eventually, we expanded this definition to include photon quality criteria that do not always involve multiple particles but are also not calorimeter based threshold vetoes. See Table 6.1 for a listing of the basic features of the kinematic cuts employed for the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ analysis. Some of the cuts are too detailed to fit compactly in the table. These cuts as well as the others will be described in more detail below.

Cut	Values	Comments
Higher Energy γ E-Min.	250 MeV	
Lower Energy γ E-Min.	150 MeV	
γ CsI Hit Position	$18 < r < 88$ cm	The inner dimension forms a square around the beam pipe, the outer is radial.
γ Hit Separation Min.	15 cm	
γ Hit timing = $T_H - T_L$	$-9.6 < \Delta T < 18.4$ ns	Different MC times.
π^0 Energy Max.	2 GeV	
γ RMS Max. - Individual	4.0	
$\sum_{\gamma} \Delta\theta\chi^2$ Max.	1	See text.
γ Energy Ratio Min.	0.88	
Acoplanarity Angle Min.	0.7854 radians (45 degrees)	
TDI Max.	2	
γ Energy Balance Max.	0.75	$(E_H - E_L)/(E_H + E_L)$
Cluster Size Min.	3	5 MeV threshold
Crystal Size Min.	1	1 MeV threshold
π^0 -projection		See text.
Missing momentum		See text.

Table 6.1: Kinematic Cuts for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ Events.

For basic feature selection, we cut on the photon energies (with different cut points for the higher and lower energy clusters) in order to remove low energy hadronic and poorly reconstructed electromagnetic showers. Our cut on the photon hit position used the reconstructed cluster (x, y) with all corrections applied and excluded a box around the nominal beam axis that enclosed the KTeV CsI (the innermost “ring” of crystals in Figure 3.2) as well as a ring 88 cm beyond the nominal beam axis. These dimensions were chosen to exclude $\eta \rightarrow \gamma\gamma$ events generated in the Charged Veto (CV) support structure which had a very high probability of striking the CsI in the inner ring and to exclude hits near the edge of the array which were often poorly reconstructed. We also require the minimum separation between clusters to be at least fifteen centimeters, although in practice this cut turns out to be so loose as to effect almost nothing. In addition to a space-separation requirement, we also require the difference in time between the higher energy photon and lower energy photon (high minus low) to be between roughly -10 and 20 nanoseconds (our timing resolution was ~ 1 nanosecond). This is the same time window applied to the single-hit crystals in the CsI veto. Finally, we also require the reconstructed pion energy to be below 2 GeV to remove very high energy neutron interactions in the collar vetoes.

The γ -RMS for each cluster is computed according to the equation:

$$RMS = \sqrt{\frac{\sum_{crystals} E_i \times (\mathbf{r}_i - \mathbf{r}_0)^2}{\sum_{crystals} E_i}}, \quad (6.1)$$

where \mathbf{r}_0 is the vector from the origin (in the plane of the CsI face) to the center of the cluster and \mathbf{r}_i is the vector from the origin to the center of the i^{th} crystal included in the cluster. In words, it is the energy weighted mean radial dispersion of a cluster. It is sensitive to fusion events, but also sensitive to very high angles of incidence (because it grows quickly for showers spread transversely across several crystals). It

is primarily for its ability to discriminate very high angle gammas (and thus reject particles coming from the CV support structure) that we cut on it so tightly. See Figure 6.1 for the distribution of RMS values in $K_L^0 \rightarrow \gamma\gamma$ MC and data.

The $\Delta\theta\chi^2$ angle cut was done using the reconstructed photon angle between the vertex and CsI hit position and a fit to the shower shape with the fit's accompanying error. The so-called "algorithm" angle is computed using a routine that uses a polynomial expansion of the energy deposition to characterize the angle:

$$r_l = \sqrt{\frac{\sum_{crystals} E_i \times (d_i P_l(\cos \alpha))^2}{\sum_{crystals} E_i}}, \quad (6.2)$$

where P_l are Legendre polynomials [52]. The γ -RMS defined above is immediately recognizable as the zeroth-order term in this expansion, while r_1 is the projected RMS along the center-of-energy direction. Using a single photon MC study, we mapped this r_1 variable to angle as a function of energy and were able to achieve a resolution of nine degrees. Significantly, our map also produced an associated error. In addition to the γ -RMS, Figure 6.1 shows the relationship between the algorithm's prediction for the angle and the angle computed via the vertex for even-pairing two cluster $K_L^0 \rightarrow \pi^0\pi^0$ MC (even-pairing events are those where the reconstructed pion decay sits on top of the true Kaon decay - these occur when the two clusters in the CsI come from the same pion). Cutting on the absolute angle turns out to involve very high acceptance losses, but we can form a χ^2 variable between the angle predicted from the algorithm and the angle calculated using the reconstructed vertex and cluster positions. This turns out to be a very useful variable for rejecting events originating close to the CsI that are reconstructed in a different location (as happens with the $\eta \rightarrow \gamma\gamma$ background.) It is also effective against odd-pairing $K_L^0 \rightarrow \pi^0\pi^0$ background events (where the two clusters in the CsI are each from a different pion) as shown in

Figure 6.2.

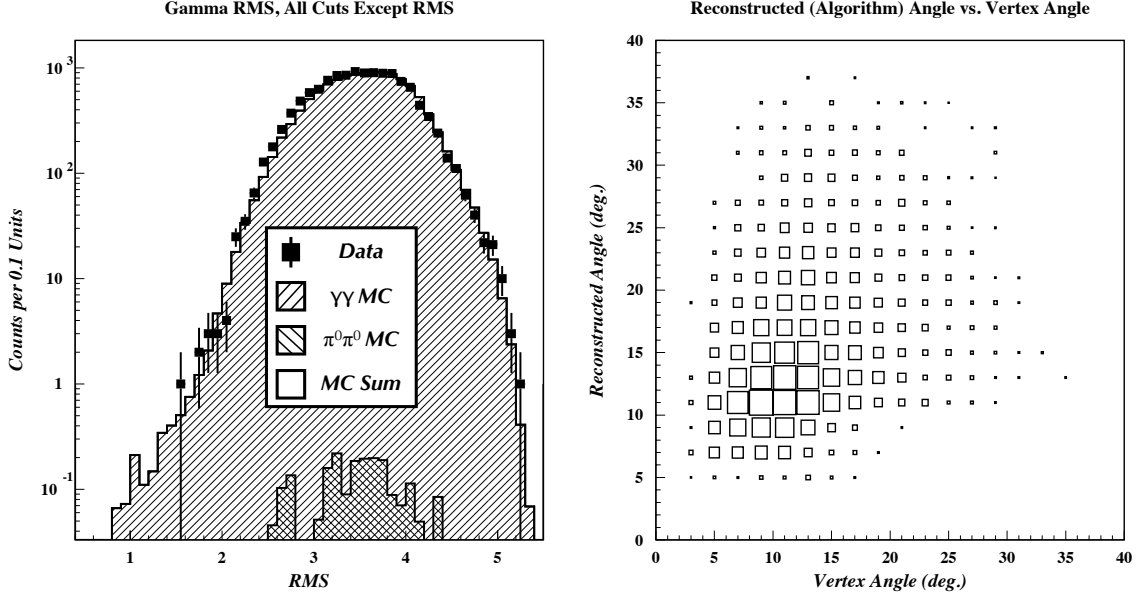


Figure 6.1: Shown here is the γ -RMS comparison for $K_L^0 \rightarrow \gamma\gamma$ MC and two cluster data (left) and the reconstructed angle versus the angle computed from the vertex to the cluster for even-pairing $K_L^0 \rightarrow \pi^0\pi^0$ MC.

The γ Energy Ratio is defined by dividing the energy sum of the three crystals with the highest energy by the total energy sum of all the crystals in the cluster. It functions as a fusion cut and it serves to remove both poorly reconstructed electromagnetic clusters and hadronic showers from direct neutron interactions.

The acoplanarity angle describes the degree to which vectors from the center of the CsI array to the cluster positions deviate from anti-parallel and is described in more detail later in the Chapter (during the discussion of Kaon decay backgrounds).

The TDI is a measure of “time dispersion” in a cluster and is defined by the relation

$$TDI = \frac{\sqrt{\sum_{crystals} (T_i - T_{mean})^2}}{N_{crystals}}, \quad (6.3)$$

where T_{mean} is the average time for all the crystals (we require energy deposition




Figure 6.2: Shown here is the $\sum_{\gamma} \Delta\theta\chi^2$ for even (left) and odd (right) pairing $K_L^0 \rightarrow \pi^0\pi^0$ two cluster MC with all other kinematic cuts applied. The $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ analysis cut is at 1.0.

above the TDC threshold of ~ 1 MeV.) The TDI cut operates as a time-like fusion cut and to remove showers affected by abnormally large fluctuations in the course of development. The TDI distribution is one we could not replicate in the MC, so we estimate the acceptance loss by studying $K_L^0 \rightarrow \gamma\gamma$ events with cluster energies consistent with those seen in the $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ MC.

The γ Energy Balance, $(E_H - E_L)/(E_H + E_L)$ measured the asymmetry in cluster energies and primarily was designed to reject events with one true photon cluster and one soft cluster due to accidental beam activity.

The cluster size cuts on the number of crystals in a cluster over two different thresholds functioned primarily to remove hadronic showers via direct neutron interactions. They also prove to be powerful discriminants against minimum ionizing interactions should a charged pion or muon (from a $K_{\mu 3}$ or $K_{3\pi}$ decay, for example) find their way past the charged veto (discussed later in this Chapter).

The π^0 projection cut is best defined using a plot - see Figure 6.3. However, we will also list the requirements for the cut, with $P_R \equiv P_T/P_Z$ for the π^0 :

- If $Z_\pi < 400$ and $P_R < 0.1$, then cut the event.
- If $Z_\pi > 400$ and the event is such that it is below the line defined by (400 cm, 0.1) and (500 cm, 0.15) in (Z_π, P_R) space, then cut the event.
- If the event is such that it is above the line defined by (300 cm, 0.2) and (500 cm, 0.34) in (Z_π, P_R) space, then cut the event.




Figure 6.3: Shown here is π^0 -projection cut parameters (in solid white) with all other cuts applied to $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC. Events outside the bounds are rejected. The thin, vertical red lines at $z = 300, 500$ are to guide the eye only.

The missing momentum for a $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ event is calculated by assuming the invariant mass of the $\nu \bar{\nu}$ system is zero. With that constraint and the assumption the P_T of the whole systems is balanced, we may solve for the four momentum of the parent Kaon.

$$\begin{aligned}
M_{\nu\bar{\nu}} &\equiv 0 \\
z &= z_{CS1} - z_{\pi^0} \\
M_{\gamma\gamma} &= \sqrt{(E_{\gamma 1} + E_{\gamma 2})^2 - p_{\pi^0, \mu} p_{\pi^0}^\mu} \\
A &= E_{\pi^0}^2 - P_{\pi^0, Z}^2 \\
B &= 0.5 \times (M_K^2 - M_{\pi^0}^2) - P_{\pi^0, T}^2 \\
D &= (B \times P_{\pi^0, Z})^2 - A \times (E_{\pi^0}^2 \times (P_{\pi^0, T}^2 + M_{\nu\bar{\nu}}^2) - B^2) \\
P_{miss} &= (B \times P_{\pi^0, Z} - \sqrt{D})/A \\
P_K &= P_{miss} + P_{\pi^0, Z}
\end{aligned} \tag{6.4}$$

If the reconstructed P_K is greater than 2 GeV/c, the event is rejected.

6.2.2 Veto Cuts

The veto cuts were classified as those operating on the basis of some energy threshold in the calorimetric components of the detector. Many details of the veto performance are discussed in Chapter 4. Our focus in this chapter is on the specific role of the vetoes in handling different background topologies.

6.3 Neutron Backgrounds

The results of this section are explained in more detail in reference [30].

6.3.1 Collar Counter 02 Events

Collar Counter 02 (CC02) was exposed to the beam halo and interactions in the downstream part of the structure like $n \rightarrow \pi^0 + X$ produced $\pi^0 \rightarrow \gamma\gamma$ events that could be reconstructed with large enough transverse momentum (P_T) to enter the signal region. This is because we assume the vertex is on the beam axis for π^0 reconstruction and true vertices far off-axis distort the P_T we calculate.

In principle, the z -vertex reconstructed for CC02 events should be approximately equal to the downstream face plane of CC02 at 275 cm (recall the origin is at the beginning of the detector), or 339.8 cm upstream of the CsI face. This is well separated from the outer edge of our signal region in z , beginning at 340 cm. However, if either photon's energy is measured incorrectly due to shower leakage or a photo-nuclear interaction, this can pull the vertex downstream into the signal region (recall that $m_\pi^2 \sim E_1 E_2 \times \theta^2$). Furthermore, there is a gaussian tail on the position reconstruction accuracy due to normal clustering energy and position resolution and it is possible for a very high event production rate for some high sigma event to occur with a mistaken vertex large enough to threaten the signal region.

This was a difficult background to estimate because our halo neutron MC was not able to reproduce the distribution of events in the neighborhood of CC02. We suspect this was related to the difficulties we had modeling the resolution of the CsI, but it was compounded by the structural complexity of CC02 coupled to the general difficulties faced by low energy hadronic MC's.

Serendipitously, we found a way to do so using a special data run taken with the initial purpose of aiding in the CsI calibration. During a roughly one-week period at the end of Run II, we broke the vacuum to install an aluminum plate in the beam core just downstream of CC02. See Figure 6.4 for an illustration of this set-up. The

purpose of the Al Target Run, or π^0 -Run, was to produce $\pi^0 \rightarrow \gamma\gamma$ decays with a known vertex. This would allow us to reconstruct the mass of the pions and calibrate the gain of the CsI calorimeter to match the true pion mass. In Figure 6.5 we show the invariant two cluster mass calculated from the Target Run data sample with the vertex fixed at the Al target. The π^0 and η peaks are clearly visible.




Figure 6.4: A cartoon illustration of the Al Target Run (π^0 -Run) set-up. The Al plate is located about 6.5 cm downstream of the edge of CC02 and is 5 mm thick.

When we apply cuts designed to enhance our calibration sample, we see a mass plot like what is visible on the left-hand side of Figure 6.5. When we apply our $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ physics analysis cuts, we see a P_T vs z -vertex distribution according to the plot in the right-hand side of that figure. We noticed that the vertex and P_T distributions of the events in Figure 6.5 matched the corresponding distributions seen between 250 and 340 cm in z in our physics data. See Figure 6.6 for a comparison of the z -vertex spectrum and Figure 6.7 for a comparison of the P_T spectrum.

In Figure 6.8 we show the results of our missing momentum calculation with the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC and the results of the cut on the π^0 -Run data. The missing




Figure 6.5: The invariant mass spectrum from the π^0 -Run under special cuts is shown on the left, and the results of applying the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ physics analysis cuts to the π^0 -Run is on the right. In the mass plot, the π^0 peak is dominant, with a second peak at high mass from $\eta \rightarrow \gamma\gamma$. The intermediate region is due to $K \rightarrow \gamma\gamma$ and the very low-mass region is dominated by soft photons from other neutron interactions. The P_T vs. z -vertex plot with $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ analysis cuts is compared to similar distributions for the physics data in Figures 6.6 and 6.7.




Figure 6.6: The z -vertex spectrum for the physics data (dotted) and π^0 -Runs (solid) with $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ analysis cuts imposed. The P_T vs. z -vertex plot for the π^0 -Runs is shown in Figure 6.5 and for the data in Figure 6.26 .




Figure 6.7: The $|P|$ (left) and P_T (right) spectra for the physics data (dotted) and π^0 -Runs (solid) with analysis cuts imposed. The P_T vs. z -vertex plot for the π^0 -Runs is shown in Figure 6.5 and for the data in Figure 6.26.

momentum cut was developed specifically to combat the CC02 background as seen here.

We use event counting and (measured) proton on target (P.O.T.) scaling on the π^0 -Run data in order to make a background prediction for CC02 events. We feel confident this sort of scaling is justified on the basis that the distributions in Figures 6.6 and 6.7 match well. See Table 6.2 for the scaling factors employed. The actual numerical background estimate is provided in Table 6.6 near the end of the Chapter.

	Physics Runs	π^0 -Runs
CC02 Events Inside $z=[200,300]$ cm	120	6824
Control Region $z=[300,340]$ cm	3	$106 \rightarrow 1.9 \pm 0.2$

Table 6.2: Normalization and control regions for the CC02-background prediction through the π^0 -Runs.

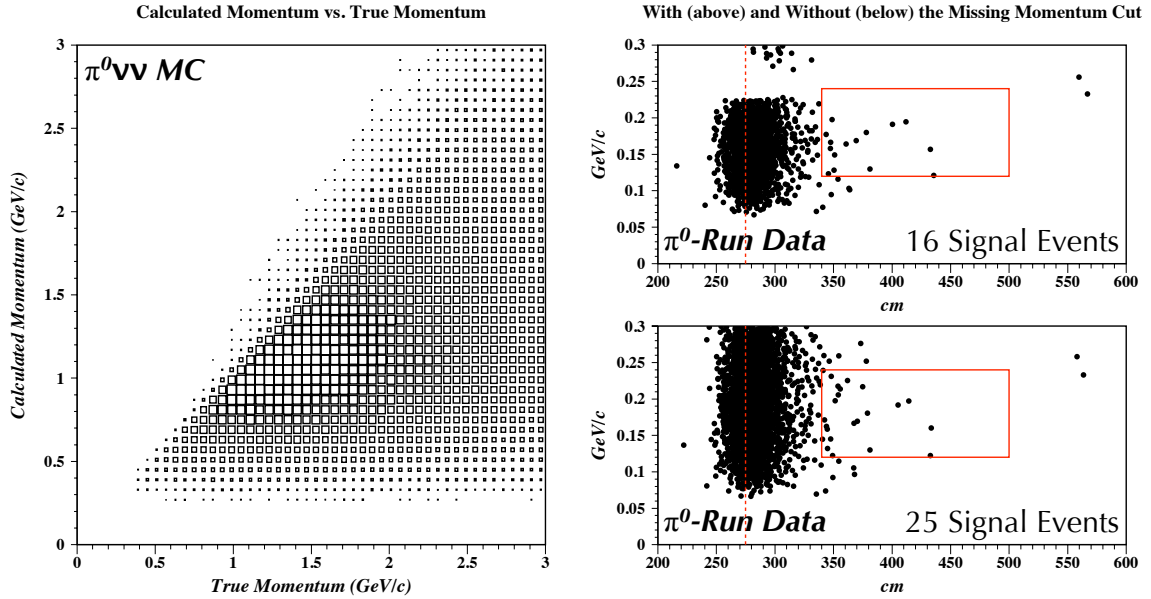


Figure 6.8: Shown here is the calculated momentum versus the true momentum for the $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ MC (left) and the π^0 -Run data with and without a cut on the “missing” momentum (right).

6.3.2 Charged Veto Events

Halo neutron interactions with the Charged Veto (CV) aluminum support structure like $n \rightarrow \pi^0 + X$ and $n \rightarrow \eta + X$ could potentially create background events. $n \rightarrow \pi^0 + X$, called “CV- π^0 ” events could contaminate the signal region if an accidental event overlapped one of the photon clusters, effectively raising the energy, or if multiple- π^0 's were produced and an odd-pairing error was made with the two spare photons lost in the veto system. Due to the interaction source in the misplaced vacuum membrane, this sort of event was the dominant background in Run-I [34]. $n \rightarrow \eta + X$ events could threaten the signal region due to the assumption in the reconstruction process that the parent particle was a π^0 . If the true parent was an η particle with a higher mass than the π^0 , this could push the vertex back upstream into the signal region (recall that $m^2 \sim E_1 E_2 \times \theta^2$). We referred to this background as the “CV- η ” background and treated the two different CV backgrounds with distinct methods, mostly for historical

reasons.

Unlike the CC02-background, we had some success in matching the distribution of events near the CV. See Figure 6.9 for a comparison of the vertex spectra for the halo neutron MC and two cluster data near the CV (the signal region is held blind at this stage). With P.O.T. scaling from our beamline simulation for the halo neutrons, we find statistical consistency in the peak.

The background prediction of the halo neutron MC due to π^0 production at the CV support structure is zero. However, because we must scale our MC up by roughly a factor of two to check the consistency of the peak, it is not clear whether we see zero events in our background MC because of low statistics or because the background truly is zero (or negligible). In order to make our prediction then, we apply a “bifurcation analysis” to the halo neutron MC.

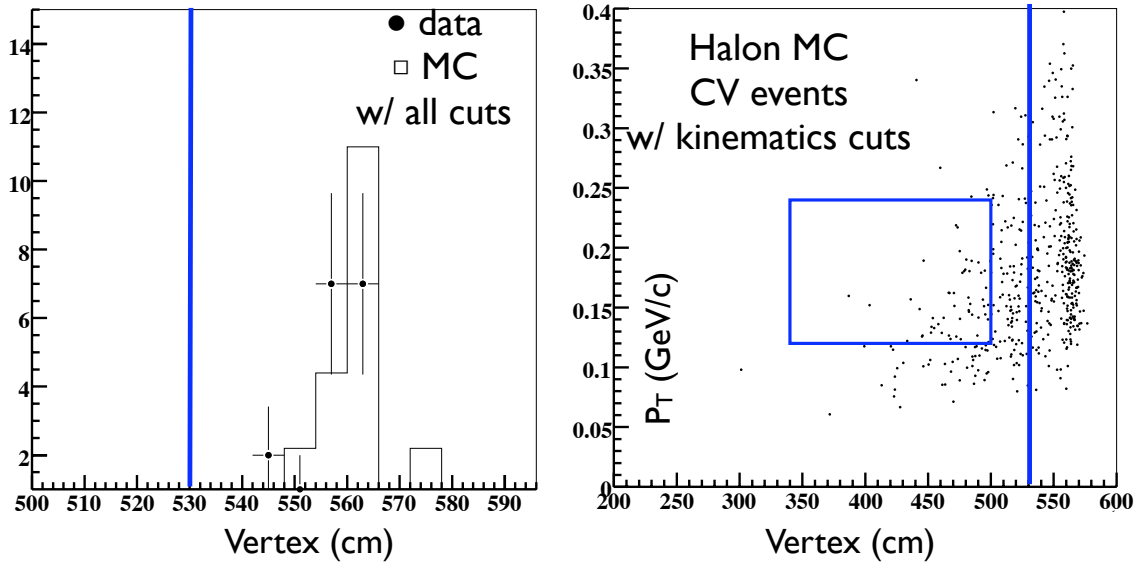


Figure 6.9: The CV- π^0 z -vertex distribution with all cuts (left) and kinematic cuts (right). With all cuts, there are 17 events in the data and 18.2 ± 6.1 with P.O.T. scaling in the MC (we must scale the MC up by roughly a factor of two).

A bifurcation analysis uses a geometrical trick to examine blinded regions of pa-

parameter space. In Figure 6.10 we illustrate this geometry. The basic idea is that groups of uncorrelated cuts form ratios in parameter space that can be used to predict event content in that region when all cuts are applied.

While first applied elsewhere [53], the methods and validity of a bifurcation were greatly extended by E391a. We employed it in a search for $K_L \rightarrow \pi^0 \pi^0 \nu \bar{\nu}$ [54] and in our first $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ analysis as well [32]. See [37] for the most comprehensive discussion available.

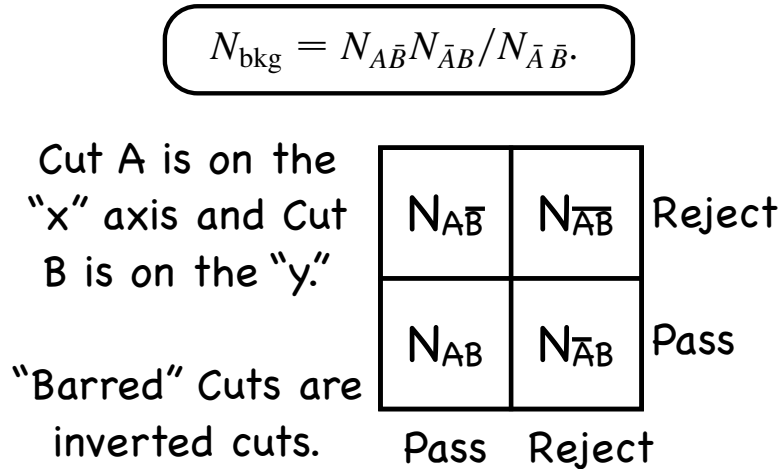


Figure 6.10: The geometrical relationship behind the idea of a bifurcation analysis.

In order to estimate the CV- η background, we had to produce a special MC. GEANT3 is not capable of producing η 's, so we imported relevant pieces of the geometry files from Gsim into a GEANT4 framework and ran a special MC with a relatively high low-energy particle production cut-off because we were primarily interested only in neutron interactions that produced η 's. Once the η particle was produced, we stopped the simulation and returned the event to GEANT3 to interact with the whole detector in our primary simulation, Gsim.

To validate this MC we attempted to reproduce the π^0 -Run mass spectrum by placing an aluminum target in the MC geometry at the appropriate position. See

Figure 6.11 for the resulting comparison, where we have normalized the π^0 and η peaks using only P.O.T. scaling from our beamline simulation. It is fair to ask why we were able to reproduce the mass spectrum here but not use the MC for a CC02 background prediction. The answer is it is much easier to match the interactions from a 5 mm piece of aluminum than for the whole of CC02 with its complex geometry and veto behavior.




Figure 6.11: The mass spectrum from the π^0 -Run compared with the composite MC. Normalization is via protons on target (P.O.T.) from the E391a beam-line simulation [45]. There is a 19.8% scaling loss applied to the MC to account for accidental activity and an additional factor of ~ 1.3 used to scale the contribution from $K_L^0 \rightarrow \gamma\gamma$. The most important feature though is the agreement in relative peak height between η and π^0 decays with only P.O.T. normalization. The very low mass discrepancy is rooted in an energy cut applied in the MC to remove very low energy neutron interactions that were incapable of producing interesting secondaries.

In Figure 6.12 we compare the $|P|$ and P_T spectra for the η mass region from Figure 6.11. Matching these spectra and the mass spectrum gave us confidence we could use the GEANT4 special MC to produce η particles with the correct cross section and behavior. Therefore, we based our background estimate on pure P.O.T.

scaling from our beamline simulation.

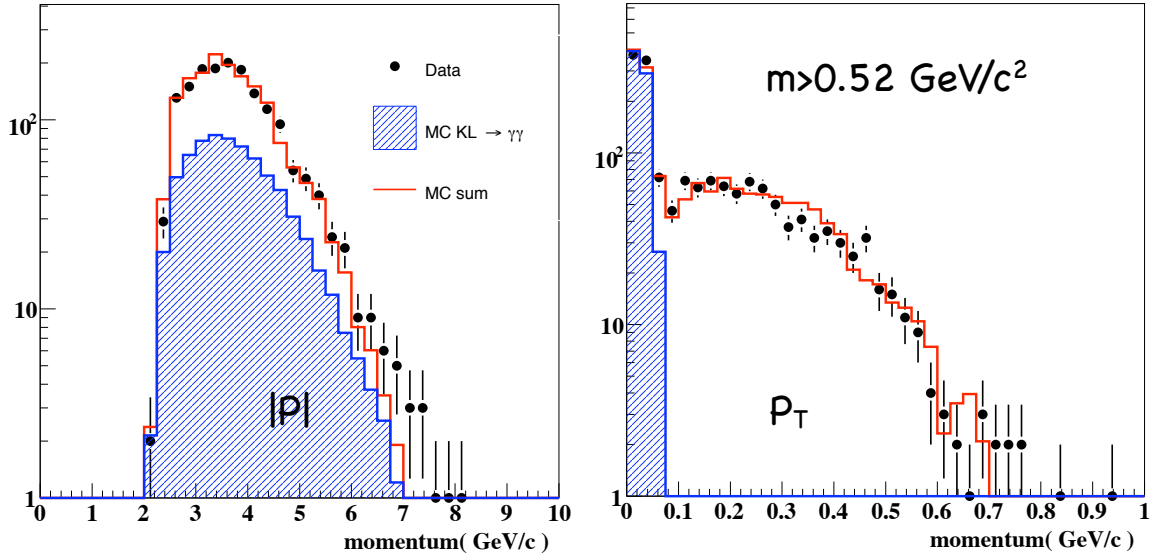


Figure 6.12: The total momentum (left) and transverse momentum (right) spectra from the π^0 -Run compared with the composite MC with all analysis cuts (no requirements on P_T). Both plots consider only the region in Figure 6.11 with mass greater than $520 \text{ MeV}/c^2$.

Reducing the η background to an acceptably low level required some stringent cuts. One of them is the π^0 -projection cut described above. This cut was designed purely to reject events coming from the CV support structure. Another cut which was tightened severely is the γ -RMS cut. This cut is effective against the η background because the true incident angle for photons is high. See Figures 6.13 and 6.14 for plots showing the γ -RMS distributions for the η -MC sample and the relative rejection power of the different cuts for the various neutron backgrounds respectively.




Figure 6.13: Shown here is the sum of the angle χ^2 values (left) and the γ -RMS distributions (center and right) for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC, the π^0 -Run data, and the η -MC with all other cuts applied. The π^0 -Run and the η -MC are normalized to each other by proton on target, but the scale of the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC is arbitrary.

6.4 Kaon Backgrounds

6.4.1 $K_L^0 \rightarrow \gamma\gamma$

In general, well-reconstructed $K_L^0 \rightarrow \gamma\gamma$ decays have P_T far too low to threaten our signal region. However, there are still ways for $K_L^0 \rightarrow \gamma\gamma$ to become a background event:

- Should one cluster energy be poorly measured, due to a photo-nuclear absorption or poor shower containment, it is possible to shift the P_T into the signal region.
- If one real photon is missed completely, an accidental cluster can be formed through a beam halo particle (neutron or γ) interaction or through a halo particle directly.




Figure 6.14: Shown here is a set of vectors comparing the exclusive “acceptance” of the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC and the various neutron background modes. Here, “eta” refers to the CV- η events and “cv” refers to the CV- π^0 events. Some entries for the CV- π^0 acceptance are larger than one because low statistics in the MC forced us to employ some ratio arguments in a “bifurcation” analysis. Some of these ratios involved small numbers that due to rounding error or some other problem caused the acceptance estimate to fail. The other entries are to be considered reliable though. The entry “gnid” refers to an unused cut.

Both of these backgrounds, but the first background especially, are highly suppressed by an *acoplanarity* cut, defined for E391a as:

$$\mathcal{A} = \pi - \arccos \left(\frac{\mathbf{u} \cdot \mathbf{v}}{|\mathbf{u}||\mathbf{v}|} \right), \quad (6.5)$$

where \mathbf{u} and \mathbf{v} are vectors pointing from the center of CsI array to the reconstructed cluster positions in the face-plane of the calorimeter. Essentially, the acoplanarity describes the degree to which the cluster positions deviate from a back-to-back decay on the beam axis (note that the angle is small when the decay is more back-to-back). See Figure 6.15 for a visual definition of the angle and Figure 6.16 for a visualization of the background distribution and its corresponding set of acoplanarity angles. As can be seen in the Figure, the acoplanarity cut completely removes the background.




Figure 6.15: The acoplanarity angle as defined in Equation 6.5.

6.4.2 $K_L^0 \rightarrow \pi^0\pi^0$

The largest background source from Kaon decays came from $K_L^0 \rightarrow \pi^0\pi^0$, where two photons escaped identification in the calorimeter. There are several ways this can

P_T vs. Z and Acoplanarity Angle for $K \rightarrow \gamma\gamma$




Figure 6.16: Shown here is the $K_L^0 \rightarrow \gamma\gamma$ background plot with loose cuts (left) and the corresponding acoplanarity angle (see Equation 6.5) distribution (right) for a MC sample roughly four times the size of the data. The cut line for the acoplanarity angle is also displayed with the distribution (all $K_L^0 \rightarrow \gamma\gamma$ events are removed by the cut).

happen, given that two of the four final state photons are successfully reconstructed:

- The spare two photons strike the veto subsystems (including possibly the CsI) and some combination of the following occurs: photon punch-through without interaction, absorption through photo-nuclear interactions, or energy deposition below the veto threshold (usually due to sampling inefficiency). In the case where a spare photon struck the CsI, it may deposit significant energy, but if all the energy is contained in one crystal, it is not counted as a cluster. (Recall that photon clusters are defined to contain two or more crystals.)
- One photon was lost in the veto systems through one of the processes mentioned above while the other *fused* with another cluster. Recall that fusion events are defined to occur when two real photons in the CsI calorimeter are reconstructed together as a single cluster. See Figure 6.17 for an illustration of a fused cluster in the calorimeter.

- Finally, it is possible for both of the reconstructed clusters to be fused clusters - i.e., four real photons in the CsI are reconstructed as only two clusters.

For the first of these cases we can clearly define two sub-types, referred to as “even” and “odd” events. These labels referring to vertex pairing combinations - for even background events, the two photons reconstructed in the calorimeter both come from the same pion. In odd cases, the assignment is mixed. Obviously, for even cases the reconstructed P_T and z -vertex match correctly with a real pion. In odd cases, the vertex identified does not correspond to the location of the Kaon decay. Because odd events are susceptible to the kinematic cuts, cases with fusion fit a similar scheme conceptually in the sense that some fusions produce clusters that reconstruct into pions quickly killed by kinematic cuts that are not designed to reject fusion at the root of the event. Other fused clusters ultimately produce pions that pass these sorts of cuts. These events must be addressed with specialized cuts. In this sense, we have even and odd backgrounds by classification even for fused cluster events.

We can study the relative proportion of these processes through MC, but it is obviously difficult to cross-check with data. However, we can validate our MC by using fully reconstructed $K_L^0 \rightarrow \pi^0\pi^0$ events where now $K_L^0 \rightarrow 3\pi^0$ with two missing photons becomes a significant background and compare data and MC to check the quality of our simulation. Recall Figure 4.6. While there are some questions about the high-mass tail of the four cluster mass distribution, the low-mass tail is well simulated. This is the portion of the distribution sensitive to losing two photons through the processes outlined above. Our ability to replicate this distribution with our MC gives us confidence to trust the two cluster $K_L^0 \rightarrow \pi^0\pi^0$ simulation with regards to background estimation.




Figure 6.17: Fusion takes place when (at least) two true photon hits are reconstructed as a single hit. Here, the true photon hit positions are illustrated with circles with radius roughly proportional to their energies.

We handle $K_L^0 \rightarrow \pi^0\pi^0$ backgrounds through two basic channels. First, our kinematic cuts are very effective at removing odd-pairing events (gammas from the two pions are mixed in assignment). This leaves even pairing events as our dominant background. The most important event topologies that threaten the signal region involve either four photons in the CsI, three photons in the CsI and one in the MB, or two in the CsI and two in the MB. Events with photons passing down the beam-hole with enough P_T to challenge the signal region are very rare.

With all analysis cuts applied, only one event from a sample roughly ten times the size of the data remains from $K_L^0 \rightarrow \pi^0\pi^0$. While we can make background predictions with such a small event sample, we cannot characterize likely background topologies. Given the importance of $K_L^0 \rightarrow \pi^0\pi^0$ both for E391a and especially for future generations of this measurement, we need to investigate the sample under looser cuts both to better characterize event topology and to check the robustness of

our estimate.

We begin by considering photon vetoes (P.V.) only and kinematics only as shown in Figure 6.18. We can classify background rejection into our even-odd schema by studying the difference between true Kaon decay vertex (known in the simulation) and the reconstructed vertex. We plot this variable against P_T in Figure 6.19. By comparing the occupancy against the P_T , we can see that the $K_L^0 \rightarrow \pi^0\pi^0$ background in the signal region is dominated by “even” events and that background rejection is largely due to the photon vetoes; in other words the $K_L^0 \rightarrow \pi^0\pi^0$ background is largely immune to kinematic cuts in the high P_T region. This is because for even pairing events (one pion vertex correctly identified), the kinematic cuts are ineffective by design - the identified particle is a legitimate pion!




Figure 6.18: $K_L^0 \rightarrow \pi^0\pi^0$ backgrounds with only veto cuts applied (left) or only kinematic cuts applied (right). The two circled events in the plot with vetoes applied only are the two events that survive *all* cuts. Low P_T events from the plot on the right are removed by the acoplanarity cut (below 0.05 GeV/c) and by the π^0 -projection cuts.




Figure 6.19: Here we plot P_T versus $\Delta z = z_{True} - z_{Rec}$ for two cluster $K_L^0 \rightarrow \pi^0 \pi^0$ events with photon vetoes only applied (left) and kinematics only applied (right). Even background events are those with small values of Δz .

We next examine the background levels under various exclusive cut sets in Figure 6.20 and Table 6.3. Here, *exclusive* cuts are understood to mean all cuts except the listed cut (or group of cuts). Rejection is completely dominated by the MB and CsI. This is not to say that the other vetoes are unimportant. Unfortunately, many of the veto cuts are correlated. This is because the fundamental background topology involves (at least) two photons in the CsI, and two photons pointed into other subsystems. Due to the size of the detectors and the forward boost of the decay products, the dominant topologies involve spare photons in the CsI and the MB. As a consequence, with those vetoes applied, if they are efficient, the other vetoes are essentially always paired to them. Interestingly, while the γ -RMS cut is removing events, much of the rejection in this case is “accidental.” There is acceptance loss even for background modes with such a tight cut. (Recall that the RMS cut is intended

to handle the $CV\text{-}\eta$ background.)




Figure 6.20: $K_L^0 \rightarrow \pi^0\pi^0$ background levels with various *exclusive* cut sets. The axis label denotes which cut (or group of cuts) is removed, with all other cuts applied. (See Table 6.3.)

Cut Set	Scaled Background Events	Relative Rejection
All Cuts	0.12 ± 0.11	$\equiv 1$
Minus γ -RMS	0.155 ± 0.13	1.3
Minus CsI	0.38 ± 0.20	3.13
Minus MB	0.81 ± 0.29	6.66
Minus any <i>other</i> single cut	0.12 ± 0.11	1.0
Minus All Kinematic Cuts	0.159 ± 0.13	1.3

Table 6.3: Background levels for $K_L^0 \rightarrow \pi^0\pi^0$ exclusive cut sets (see Figure 6.20). The scale factor is dictated by the size of the MC sample relative to data (MC:data = 9.6:1).

Given the primacy of the MB and CsI, what can we say about the event topologies they are most effective against? We begin with the MB and summarize the effects in Figure 6.21. In 73% of the events in the signal region, there are two hits in the MB. The remaining 27% involve single hits to the MB, and 94% of those events send the other photon in the CsI, with only 6% sending a photon down the beam-hole. Note, the precision here is slightly misleading. Because of low statistics even with this veto set, statistical uncertainties on all these numbers are about ten percent of themselves.

True MB Hits vs. Pion P_T & Z - $2\pi^0$ Background MC
All Photon Vetoes except the Main Barrel




Figure 6.21: The number of hits in the Main Barrel with all cuts applied *except* the MB are shown versus P_T on the left and against reconstructed vertex on the right. All hits come from events inside the signal region.

In Figure 6.22 we summarize the topologies for events that remain if the CsI veto is shut off. This Figure is very illustrative of the reason we have tolerated so many problems in our CsI veto systematics - it is superb at removing additional photon activity. Of the events rejected by the CsI veto, roughly 70% involve four photons in the CsI and 28% involve three photons. There is significant event rejection overlap between the CsI and MB (three photons in the CsI and one in the Main Barrel), but the $\sim 30\%$ is coincidental (the MB rejects more events than the CsI). Most of the additional photon hits in the CsI are failed clusters. Only about 9% of the total events remaining involve clusters with true hit positions in the CsI closer than 14 cm - true single fusion events. All these events come from events with four photons in the CsI (13% of the four photon events are fusion events.) Again, the precision here is somewhat misleading. Statistical uncertainty on each number is about 15% of itself.

True CsI Hits vs. Pion P_T & Z - $2\pi^0$ Background MC
All Photon Vetoes except the CsI

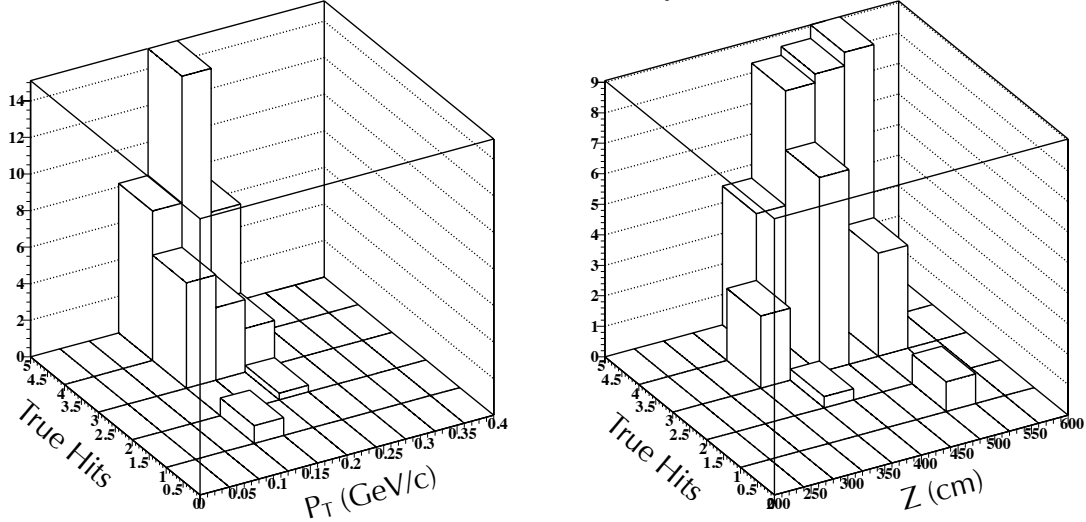


Figure 6.22: The number of hits in the CsI with all cuts applied *except* the CsI are shown versus P_T on the left and against reconstructed vertex on the right. All hits come from events inside the signal region.

One measure of correlation between cuts that is simple to compute with the values at hand is:

$$\rho_{AB} = \frac{N_{All}N_{-AB}}{N_{-A}N_{-B}} - 1, \quad (6.6)$$

where N_{All} is the number of events with all cuts imposed, N_{-AB} is the number of events with both A and B cuts removed, and N_{-A} (N_{-B}) is the number of events with only cut A (B) removed. This measure of correlation is zero for uncorrelated cuts, but is unfortunately not normalized. However, it still provides a useful *relative* measure. See Table 6.4 for a summary of the correlations between the most important vetoes. The high correlation between the CsI and MB is reflective of the most dominant topologies (one spare photon in the MB, one spare photon in the CsI).

Before moving further though, it is worth commenting briefly on “the event.” With all cuts applied, one event with a weight of 1.16 remains in the signal region. Scaled

Removed Cuts	Scaled Background	Implied Correlation	Total Rejection
MB & CsI	29 ± 1.7	11	240
MB & BA	5.1 ± 0.72	5.3	43
CsI & BA	1.3 ± 0.4	2.3	11

Table 6.4: The correlation is here defined according to Equation 6.6.

by MC statistics, the final weight is 0.12 events. The topology of the event is quite interesting. It is an even event, with the kaon vertex coinciding with the reconstructed pion vertex. The missing pion sends a 14 MeV photon to the MB where it is lost due to sampling inefficiency. Interestingly, if our MB cut was single-sided (instead of requiring energy above TDC threshold in both the up and down-stream ends of each module for it to be included in the veto decision) it would remove this event. However, such a tight MB cut would involve significant additional acceptance loss. The other photon strikes the CsI (not a gap between crystals) and punches through cleanly without interacting! For a 15 radiation length crystal, this is a four parts per billion level process. In a back of the envelope estimate, we might suppose that roughly ten million $K_L^0 \rightarrow \pi^0\pi^0$ decays could threaten with background in our data. If we assume an average of two photons in the CsI per decay, we would still expect to see a punch-through only ten percent of the time... and that is *before* suppression due to other selection criteria! This confluence of rare processes highlights the sort of challenges a search like this faces.

6.4.3 $K_L^0 \rightarrow \pi l \nu$ and $K_L^0 \rightarrow \pi^+\pi^-\pi^0$

Roughly four fifth's of the possible decay channels of the K_L is accounted for in the $K_L^0 \rightarrow \pi l \nu$ and $K_L^0 \rightarrow \pi^+\pi^-\pi^0$ modes. $K_L^0 \rightarrow \pi^\pm e^\mp \nu$, called Ke3, has a branching fraction of 40.53%, $K_L^0 \rightarrow \pi^\pm \mu^\mp \nu$, called Kmu3 has a branching fraction of 27.02%, and $K_L^0 \rightarrow \pi^+\pi^-\pi^0$, called K3pi, has a branching fraction of 12.56% [31]. (Ke3 and

Kmu3 together are referred to as K13.) All of these decay channels involve at least two charged particles in the final state. If we assume a very conservative charged veto inefficiency of 10^{-4} [55], and an “acceptance” (probability of the decay products producing two clusters kinematically consistent with a single π^0) one order or so worse than that for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$, we would expect backgrounds due to direct charged particle hits in the CsI to be essentially negligible in a study with sensitivity around 10^{-9} despite their cumulative branching fraction of ~ 1 .

However, that sort of “back-of-the-envelope” calculation is relying on both charged particles interacting with the various charged veto subsystems (the Inner and Outer Charged Vetoes, the Barrel Charged Veto, and the Beam-hole Charged Veto / Beam-Anti complex) and the charged veto subsystems performing to specification. Furthermore, while it seems reasonable to expect the π^0 -like acceptance of these modes to be lower than for real π^0 's, we should quantify that expectation.

The largest problem when dealing with these backgrounds is their very large branching ratios. Owing to the complexity of the E391a MC, it is impossible to generate large (flux times ten) samples to use for background studies. Furthermore, since there is no charged-particle tracking in E391a, we cannot tag and study K13 or K3pi independently in data. The approach here then is to generate as large a sample of MC events as possible in the time available and study those events under loose vetoes and kinematic cuts. Generally, there are two basic background topologies that concern us: 1) charge exchange and 2) charged veto inefficiency.

In the first case, a physical reaction like $\pi^- + p \rightarrow n + \pi^0 \rightarrow n + \gamma\gamma$ occurs very early in one of the charged veto subsystems (or perhaps in a collar counter) and so the charged particle does not leave sufficient visible energy to trip the veto. See Figure 6.23 for a cartoon illustration. The resulting neutron is subsequently lost and the remaining photon pair form a background event. We have two basic forms

of protection against this topology: it is a relatively rare process and the vertex is centered at the charged veto. Unfortunately, because this is a rare process, it is difficult to study with simulation. We observed no evidence of a charge exchange event peak near CC02 or the Outer CV in any of our charged mode simulations. However, as can be seen from Table 4.5, our MC samples are small enough that this lack of observation does not rule the backgrounds out. Nevertheless, we rely on that lack of observation and the ability of our other background estimation processes to match the distributions around the signal box to relegate this background to the status of unimportant. See [56] for further discussion of this background source.

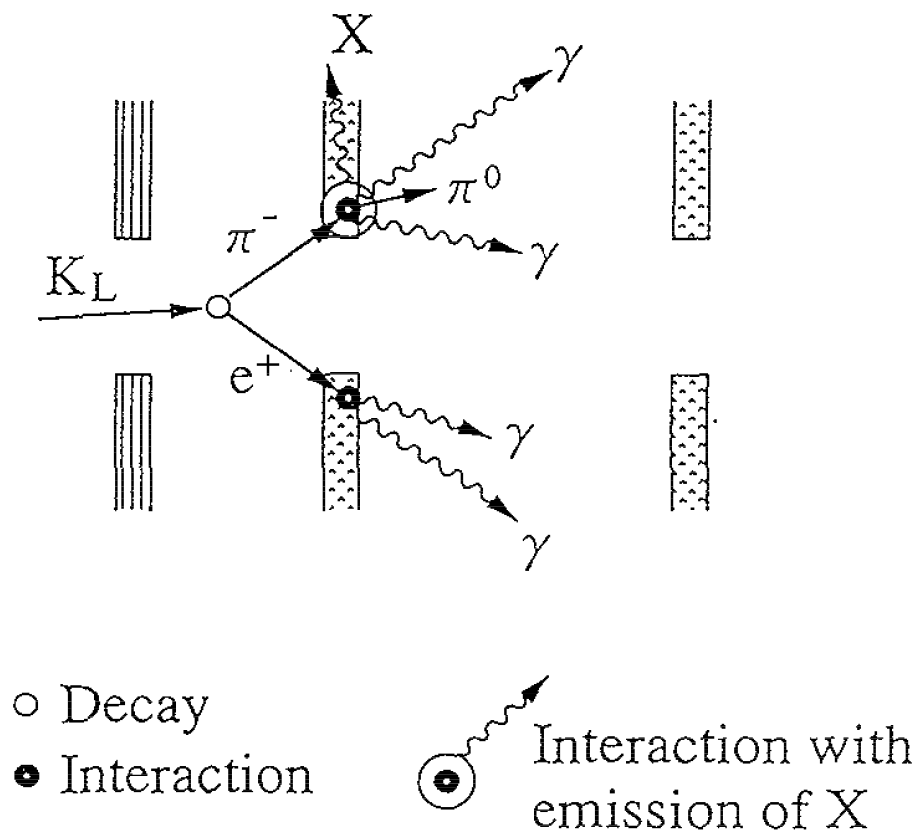


Figure 6.23: Shown here is a cartoon depicting two ways in which particles may “exchange” their electric charge in a Ke3 event: 1) by $\pi^- + p \rightarrow n + \pi^0$ and 2) by $e^+ + e^- \rightarrow \gamma\gamma$. Figure adapted from [56].

A potentially more dangerous background is mechanical failure. If some piece of the charged veto complex was damaged during installation (broken optical fibers, for example) or became inert during the course of the experiment, it could create an avenue for background to leak into the signal region in unanticipated ways. We can explore this background by examining the event topologies of our charged mode simulation and making guesses about the likely *mechanical* inefficiency. The idea here is to examine events reconstructed in the signal region under loose cuts and then factor in suppression from tighter vetoes and kinematics. We then examine the background levels predicted for different levels of performance in the charged veto subsystems. Our primary concern with the charged vetoes is not that the fundamental behavior of plastic scintillator vetoes has somehow degraded seriously from expectation and previous measurements, but rather that some form of unaccounted damage in construction, mechanical or material failure, or misalignment could introduce an inefficiency over what was expected.

The process employed in the estimation was to examine charged mode MC that had passed all of our skim stage vetoes and produced two cluster events. No requirements were made that the clusters originate directly via one or another particle interaction, we only required them to satisfy our clustering algorithm (see Chapter 2). In Figure 6.24 we show an example of the background distributions for Ke3 under some loose cut conditions. In the analysis we first relaxed all the charged veto subsystems. This included the Inner and Outer CV, the Barrel CV, and the entire BHCV and BA complex. Application of only the Outer and Inner CV is enough to remove all events in the MC. The next step was to relax all the kinematic cuts.

We then traced each of the charged daughter particles from the decay in iterative steps of one centimeter from the Kaon decay vertex until they either hit a veto subsystem, decayed, or escaped the mother volume of the MC. Because our MC did




Figure 6.24: Shown here is the distribution of background events from the Ke3 MC under loose cuts. The signal event counts are not scaled for statistics.

not store momentum information for secondary decay daughter products, we could not track the secondary daughter products and did not speculate on their direction. Events were then classified by the number and type of veto subsystem they hit. The following event topologies emerged: single hits to the CV (either Outer or Inner), single hits to the Barrel CV, single hits to the Beam Hole complex (BHCV and BA), and double hits to combinations of those detectors. We found no events where both particles failed to hit at least one of the charged veto subsystems. See Table 6.5 for a summary of these event classes for the Ke3 MC.

We can sum the contributions in Table 6.5 and scale for statistics to form a speculative background prediction. The assumed veto suppression for the CV is based on an upper limit estimated through beam-line muon studies of $(1.1 \pm 0.2) \times 10^{-4}$ [57]. The other suppressions are simple estimates. The sum is clearly dominated by

Veto (& Assumed Suppression)	Raw Hits	Kinematic Suppression Factor	Un-Scaled Background
BH Only (10^{-3})	0	0	0
CV Only (10^{-4})	8.17×10^3	0.00117	9.6×10^{-4}
BCV Only (10^{-3})	0	0	0
BH + BH (10^{-6})	0	0	0
CV + CV (10^{-8})	762×10^3	0.000971	7.4×10^{-6}
BCV + BCV (10^{-6})	0	0	0
BH + CV (10^{-7})	4.21	0	0
BH + BCV (10^{-6})	0	0	0
CV + BCV (10^{-7})	97.5	0.000282	2.7×10^{-9}

Table 6.5: The raw hits and kinematic suppression factors for the different event topologies (single and double charged system hits). The last column is a product of the first three and represents the background estimate prior to scaling (up) for MC statistics (by a factor of $1/0.021 \sim 50$ for Ke3). The assumed veto suppression for the CV is based on an upper limit estimated through beam-line muon studies [57]. The other suppressions are speculative.

single sub-system hits. What is happening in those events? In Figure 6.25 we present a picture of the vertices of the charged pions in those events.

It is highly likely that many of the single subsystem hits are actually double subsystem hits because we cannot track the daughter muons from pion decay and, as can be seen in Figure 6.25, a great many of the single sub-system events are coming from decays where the daughter muon is likely to hit either the Outer CV or Barrel CV, with some others escaping down the beam hole. Those escaping down the beam hole could miss the BHCV and BA, but application of the downstream collar counters makes even these events unlikely.

However, since we cannot actually track these particles and be certain of their hit locations, we will instead simply treat those events as true single subsystem hits and regard the resulting estimate as an upper bound. If we use 10^{-4} as the CV suppression with a 20% error, and 48 ± 3.5 for the statistical scaling factor (where the

Charged Pion Decay Vertex R vs. Z, Loose Kinematics + Vetoes (No Charged Systems)




Figure 6.25: Shown here is the radial position versus z -vertex for the charged pion decays in the Ke3 MC with loose cuts imposed (all the charged veto subsystems are removed in addition to some of the kinematic cuts). Almost all of the events with (z,r) positions beyond ~ 550 cm in z and ~ 15 cm in r were counted as charged veto hits.

error is due to the uncertainty on the flux) we find an estimate of 0.046 ± 0.0097 for Ke3. Repeating this calculation for K3pi and Kmu3 yields background predictions that are completely negligible (on the order of 10^{-4}).

6.5 Total Background Check

In Figure 6.26 and Table 6.6 we provide a final summary of all background predictions and compare predictions for regions around the signal box to our observations from data. We find statistical consistency for all regions and therefore feel confident to open the signal region.

Region	1	2	3	2+3	4
CC02	1.9 ± 0.2	0.11 ± 0.04	0.05 ± 0.03	0.16 ± 0.05	0.26 ± 0.07
CV- π^0	0	0	0.08 ± 0.04	0.08 ± 0.04	0.09 ± 0.04
CV- η	0.02 ± 0.01	0.04 ± 0.02	0.02 ± 0.01	0.06 ± 0.02	0.04 ± 0.01
$K_L^0 \rightarrow \pi^0 \pi^0$	0			0.12 ± 0.11	0
Ke3	< 0.01			0.046 ± 0.0097	< 0.01
Total	1.9 ± 0.2			0.45 ± 0.13	0.39 ± 0.08
Observed	3	NA	NA	NA	2

Table 6.6: Predicted and observed backgrounds in the region around the signal box. The predicted and observed background levels are statistically consistent. See Figure 6.26 for the associated plot.

Two Cluster Data - Analysis Cuts Imposed




Figure 6.26: Shown here is the two cluster data final plot with the signal region still blind. Note the four numbered regions. In (P_T, z) space Region 1 stretches from (300 cm, 0.12 GeV/c) to (340 cm, 0.24 GeV/c); Region 2 stretches from (340 cm, 0.12 GeV/c) to (400 cm, 0.24 GeV/c); Region 3 stretches from (400 cm, 0.12 GeV/c) to (500 cm, 0.24 GeV/c); Region 4 stretches from (300 cm, 0.0 GeV/c) to (500 cm, 0.12 GeV/c).

CHAPTER 7

CONCLUSIONS

7.1 Critical Summary

We establish a value for the acceptance of $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$, where the acceptance is defined as the fraction of events decaying through a given channel inside our fiducial region that we identify as signal, of 0.67%. Taken with our previously computed flux for normalization, we find a single event sensitivity of 2.91×10^{-8} before accounting for error on the flux. Upon examining our signal region we observe zero events. Therefore, we set a new upper limit on the branching ratio for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ at 6.7×10^{-8} at the 90% confidence level using Poisson statistics.

7.2 $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ Event Generator

For the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ event generator we did not produce or track the neutrinos produced by the decay. The pion momentum distribution was produced according to calculations in [58] and [59]. The differential decay rate can be related to the Kaon and pion masses and energies through:

$$\frac{d\Gamma}{dE_\pi} \sim \lambda f_+^2 \left[\left(M_K^2 + M_\pi^2 - q^2 \right)^2 - \frac{2}{q^2} \left(\frac{q^2 \lambda^2}{3} + M_\pi^2 q^4 \right) \right], \quad (7.1)$$

where

$$\begin{aligned} q^2 &= M_K^2 + M_\pi^2 - 2 M_K E_\pi, \\ \lambda^2 &= \left(M_K^2 + M_\pi^2 - q^2 \right)^2 - 4 M_K^2 M_\pi^2. \end{aligned} \quad (7.2)$$

We parameterized the form factor f_+ according to:

$$f_+ = 1 + \lambda_+ \frac{q^2}{M_\pi^2}, \quad (7.3)$$

and we use $\lambda_+ = 0.0284$ as computed by Ke3 and Kmu3 experiments [31].

We begin by computing the π^0 energy in the Kaon rest frame. Initially the pion is assigned an energy equal to its rest mass and assigned a weight of unity. This weight is compared to the function given in Equation 7.1. So long as the weight is higher than the return value, we choose a new energy from a flat distribution between the rest mass and maximum energy (given by $(M_K^2 + M_\pi^2) / (2 M_K)$) and assigned a weight value chosen from a flat distribution between zero and one. With the pion energy in hand we generate the direction of the momentum vector as a random isotropic value, and then boost the resulting four momentum into the lab frame according to the Kaon momentum vector.

7.3 Acceptance and Sensitivity

The single event sensitivity is defined as:

$$\text{S.E.S.}(K_L^0 \rightarrow \pi^0 \nu \bar{\nu}) = \frac{1}{F \times A}, \quad (7.4)$$

where F is the Kaon flux, defined here as the number of decays within our fiducial volume, and A is the acceptance for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$. Note that our use of the word “flux” here is a bit of jargon - a better phrase would be “time integrated flux,” but we simply use the term flux and leave the time integration understood. We formally compute our flux using a well understood normalization mode. In E391a we have chosen $K_L^0 \rightarrow \pi^0 \pi^0$ as our normalization mode on the basis of similarity in photon

energies observed in the final state. The flux is defined according to the equation

$$F = \frac{N_M}{A_M \times Br(M)}, \quad (7.5)$$

where N_M is the number of observed events decaying through the chosen mode, A_M is the acceptance of the normalization mode, and $Br(M)$ is its branching ratio. See Chapter 4 for a more complete discussion of the calculation of the flux for the normalization mode.

The acceptance of $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ is computed through use of the Gsim MC. In Figure 7.1 we show a vector of the exclusive acceptance (the acceptance of a cut with all other cuts imposed) for the kinematic cuts and in Figure 7.2 we show the same vector for the photon veto cuts.

With all cuts applied, the distribution of $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ events reconstructed can be seen in Figure 7.3. We find 90,851.2 events in the signal region (recall that event weights deviate from unity as described in Chapter 4). From a sample of 1.071×10^7 decays in our fiducial, this provides a raw acceptance of $(0.848 \pm 0.003_{stat})\%$. We factor in the acceptance loss due to accidental hits by using the factor computed for $K_L^0 \rightarrow \gamma\gamma$ decays (17.53%) to arrive at a final acceptance value of $(0.670 \pm 0.003_{stat})\%$.

Computing the systematic error on the acceptance (the error due to the differences in the behavior of the cuts between MC and Data) is difficult because we obviously do not have immediate access to a large sample of well-reconstructed $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ events for comparison! Therefore, we choose to simply use the same systematic error as computed for the normalization mode ($K_L^0 \rightarrow \pi^0 \pi^0$) of 7.3%. Because the systematic error is dominated by the CsI veto and the photon counting statistics are lower for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$, this is a conservative estimate even if it does include a somewhat different set of cuts. Therefore, our final expression for the acceptance is $(0.670 \pm$




Figure 7.1: Shown above is the $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ MC exclusive acceptance for the kinematic cuts. The exclusive acceptance for a cut is the acceptance for that cut only with all other cuts applied.




Figure 7.2: Shown above is the $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ MC exclusive acceptance for the photon veto cuts. The exclusive acceptance for a cut is the acceptance for that cut only with all other cuts applied.

$0.003_{stat} \pm 0.049_{syst})\% = (0.670 \pm 0.049)\%$. Using this value, we can compute our single event sensitivity and find $S.E.S. = (2.91 \pm 0.31) \times 10^{-8}$.




Figure 7.3: Shown above is the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ MC final plot (all analysis cuts are imposed).

7.4 Final Results

Finally, in Figure 7.4 we show the results of all analysis cuts applied to the data with the signal region open.

Two Cluster Transverse Momentum vs. Decay-Z, All Analysis Cuts




Figure 7.4: Shown here is the two-cluster data final plot (all $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ analysis cuts are imposed). There are 141 events in the plot with a single overflow event at $z \sim 550$ with very high P_T .

Because we observe zero events, we set an upper limit on the decay. Using Poisson statistics, we find:

$$Br(K_L^0 \rightarrow \pi^0 \nu \bar{\nu}) < 2.3 \times 2.91 \times 10^{-8} = 6.69 \times 10^{-8} \quad (7.6)$$

at the 90% Confidence Level (C.L.). If we take the uncertainty on the flux and acceptance into account, we find (adopting the formula from [51]):

$$\begin{aligned} Br(K_L^0 \rightarrow \pi^0 \nu \bar{\nu}) &< 2.3 \times (1 + 2.3/2 \times (\sigma_{\text{S.E.S.}}/\text{S.E.S.})^2) \times \text{S.E.S.} \\ &< 6.78 \times 10^{-8} \end{aligned} \quad (7.7)$$

at the 90% C.L. At the 95% C.L., our limit is 8.73×10^{-8} without accounting for uncertainties. This result is consistent with Standard Model (SM) expectations. These numbers are all summarized in Table 7.1. This represents roughly a factor of three improvement over our Run I result [32]. See Table 7.2 for a point-by-point comparison between our Run I and Run II results. It is important to emphasize that the Run I result used only $\sim 10\%$ of the data set for that Run. For that analysis, we were background limited due to a mechanical failure in the vacuum membrane (because the analysis was background limited, additional statistics would not have improved the result).

Run II $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ Acceptance	$(0.670 \pm 0.049) \times 10^{-2}$
Run II S.E.S.	2.91×10^{-8}
Run II 90% C.L. Upper Limit	6.69×10^{-8}
Run II 90% C.L. Upper Limit (uncertainties included)	6.78×10^{-8}

Table 7.1: Summary of the Run II results.

Run I 90% C.L. Upper Limit	2.1×10^{-7}
Run I Acceptance	0.657%
Run I Background Estimate	1.9 ± 1
Run I S.E.S.	9.11×10^{-8}
Run II / Run I Signal to Noise ratio	$S.E.S._{II}/S.E.S._I \times BG_I/BG_{II} \sim 12$

Table 7.2: Comparison of the Run I and Run II results.

7.5 How an E391a-style Experiment Can Discover

$$K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$$

E391a was limited by statistics and acceptance, not background. As a consequence, there is hope for success in the future of this difficult measurement using this technique. The largest acceptance loss in the experiment comes from the application of the CsI veto. There are two reasons for the large loss: 1) the large (transverse) size of the crystals makes fusion rejection and even single cluster identification challenging, and 2) the thin (longitudinal) size of the CsI allows for photon punch-through and energy leakage effects that are detrimental to the vertex resolution and lead to a higher than desired fraction of poorly formed clusters. There is no question that the greatest feature of the E391a CsI array was its *availability* - a less than ideal calorimeter still proved much better than no calorimeter at all! However, in terms of recovering acceptance, an improved CsI array is the place to begin.

Ultimately, the largest background source in E391a was due to the interactions of halo neutrons with the detector materials. There are two ways to improve this situation. The first is (obviously) to reduce the neutron halo even further. This may be technically quite challenging. The second way, based on E391a style geometry, is to convert the last collar counter prior to the fiducial (CC02 for E391a) into a purely active detector and move its location upstream. This will help recover acceptance

by shifting the background vertex away from the fiducial, making event separation easier. Additionally, upgrading the CsI calorimeter will improve vertex resolution in this case. Backgrounds associated with the Charged Veto support structure can be addressed in any number of ways: a thinner support structure, an active support structure, or even possibly designing the CV such that there is no support structure!

Even if the beam-line is designed to completely remove the neutron component though, the threat of Kaon backgrounds always loom. $K_L^0 \rightarrow \pi^0\pi^0$ in particular looks irreducible based on simple scaling from the E391a result. However, it is to be emphasized that the $K_L^0 \rightarrow \pi^0\pi^0$ background is again fundamentally rooted to the behavior of the CsI veto. Increasing the longitudinal thickness and shrinking the transverse size will simultaneously reduce punch-through and fusion backgrounds. Of course, the performance of the analogue of the Main Barrel is critical to the rejection of this background as well. Thickening the barrel to reduce punch-through is a necessary step, but so also is taking action to reduce sampling inefficiency.

An experimental proposal has been submitted and approved to extend this search [49] at the J-Parc hadron facility in Tokai, Japan. The E14 experiment hopes to ultimately see 100 events - requiring a S.E.S. of roughly 10^{-13} for a SM observation. This is truly a formidable challenge. However, many of the same collaborators from E391a are participating, as well as many new and excited members. The lessons learned from this analysis will live on in future generations of this experiment!

REFERENCES

- [1] M. Goossens and F. Mittelbach. *The L^AT_EX Companion, Second Edition*. Addison-Wesley, 2004.
- [2] The TeXShop Group. <http://www.uoregon.edu/~koch/texshop/>, 2006.
- [3] T. D. Lee and C. N. Yang. *Physical Review*, 104(254), 1956.
- [4] C. S. Wu et al. Experimental Test of Parity Conservation in Beta Decay. *Physical Review*, 105, 1957.
- [5] T.D. Lee, R. Oehme, and C.N. Yang. *Phys. Rev.*, 106:340–345, 1957.
- [6] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay. Evidence for the 2π decay of the K_2^0 meson. *Physical Review Letters*, 13(4), 1964.
- [7] W. J. Marciano and Z. Parsa. Rare kaon decays with “missing energy”. *Physical Review D*, 53(1), 1996.
- [8] F. Mescia and C. Smith. Improved estimates of rare K decay matrix-elements from K_{l3} decays. arXiv:0705.2025v1, 2007.
- [9] A. J. Buras et al. Waiting for precise measurements of $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ and $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$. hep-ph/0405132, 2004.
- [10] A. Alavi-Harati et al. Search for the decay $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ using $\pi^0 \rightarrow e^+ e^- \gamma$. *Physical Review D*, 61, 2000.
- [11] H. Frauenfelder and E. M. Henley. *Subatomic Physics*. Prentice Hall, 1991.
- [12] R. Sachs. *The Physics of Time Reversal*. The University of Chicago Press, 1987.
- [13] A. J. Buras. CP Violation and Rare Decays of K and B Mesons. hep-ph/9905437, 1999.
- [14] M. Gell-Mann and A. Pais. *Phys. Rev.*, 97(1387), 1955.
- [15] A. Alvi-Harati et al. Measurements of Direct CP Violation, CPT Symmetry, and Other Parameters in the Neutral Kaon System. *Phys. Rev. D*, 67(012005), 2003.
- [16] V. Prasad. *Measurements of the Direct CP-Violation and CPT-Invariance in the Neutral Kaon System*. PhD thesis, The University of Chicago, 2002.
- [17] Nicola Cabibbo. Unitary Symmetry and Leptonic Decays. 10(12), 1963.
- [18] M. Kobayashi and T. Maskawa. CP Violation in the Renormalizable Theory of Weak Interactions. *Progress of Theoretical Physics*, 49(2), 1973.

- [19] S. L. Glashow, J. Iliopoulos, and L. Maiani. Weak Interactions with Lepton-Hadron Symmetry. *Phys. Rev. D*, 2(7):1285–1292, Oct 1970.
- [20] L. Wolfenstein. Parametrization of the Kobayashi-Maskawa Matrix. *Physical Review Letters*, 51(21), 1983.
- [21] C. Jarlskog. Commutator of the Quark Mass Matrices in the Standard Electroweak Model and a Measure of Maximal CP Nonconseration. *Physical Review Letters*, 55(10), 1985.
- [22] J. Charles et al. *Eur. Phys. J. C*, 41:1–131, 2005.
- [23] K. G. Wilson. *Phys. Rev.*, 179, 1969.
- [24] The Fermilab Tevatron. <http://www-bdnew.fnal.gov/tevatron/>.
- [25] The Large Hadron Collider. <http://lhc.web.cern.ch/lhc/>.
- [26] A. J. Buras et al. The Impact of Universal Extra Dimensions on the Unitarity Triangle and Rare K and B Decays. hep-ph/0212143, 2003.
- [27] Christoph Bobeth et al. Upper bounds on rare K and B decays from minimal flavour violation. *Nuclear Physics B*, 726:252–274, 2005.
- [28] G. Isidori and P. Paradisi. Higgs-mediated $K \rightarrow \pi\nu\bar{\nu}$ in the MSSM at large $\tan\beta$. arXiv:hep-ph/0601094v1, 2006.
- [29] M. Blanke et al. Rare and CP-Violating K and B Decays in the Littlest Higgs Model with T-Parity. arXiv:hep-ph/0610298v2, 2006.
- [30] T. Sumida. *Search for the decay $K_L \rightarrow \pi^0\nu\bar{\nu}$* . PhD thesis, Kyoto University, 2008.
- [31] W. M. Yao et al (The Particle Data Group). *J. Phys. G*, 33(1), 2006.
- [32] J. K. Ahn et al. New limit on the $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ decay rate. *Phys. Rev. D (RC)*, 74(051105(R)), 2006.
- [33] M. Doroshenko. *A Measurement of the branching ratio of the $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$ decay*. PhD thesis, SOKENDAI, KEK, 2005.
- [34] K. Sakashita. *Search for the decay $K_L^0 \rightarrow \pi^0\nu\bar{\nu}$* . PhD thesis, Osaka University, 2006.
- [35] A. Roodman. Blind Analysis in Particle Physics. PHYSTAT2003, 2003.

- [36] H. Watanabe et al. Neutral beam line to study $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ decay at the KEK 12-GeV proton synchrotron. *Nuclear Instruments and Methods in Physics Research A*, 545:542–553, 2005.
- [37] J. D. Nix. *First Search for $K \rightarrow \pi^0 \pi^0 \nu \bar{\nu}$ and $K \rightarrow \pi^0 \pi^0 P$* . PhD thesis, The University of Chicago, 2007.
- [38] H. S. Lee. *Search for the Rare Decay $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ at E391a*. PhD thesis, Pusan National University, 2007.
- [39] CERN. GEANT 4.8.3.
- [40] M. Doroshenko. Calibration of E391 detector. Internal Tech-Note, 2007.
- [41] The KTeV Experiment at Fermilab. <http://kpsa.fnal.gov:8080/public/ktev.html>.
- [42] M. Doroshenko et al. Undoped CsI calorimeter for the $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ experiment at KEK-PS. *Nuclear Instruments and Methods in Physics Research A*, 545:278–295, 2005.
- [43] K. F. “Jack” Chen. CsI Calibration Using $K_L^0 \rightarrow 3\pi^0$ Decays. Internal Memo, 2005.
- [44] J. D. Nix. Study of the Necessity of Collar Counters 1 and 4. Internal Memo, 2001.
- [45] H. Watanabe. PhD thesis, Saga University, 2002.
- [46] CERN Program Library Long Writeup W5013. GEANT - Detector Description and Simulation Tool, 1993.
- [47] A. Lednev. Some examples on kinematic fits for E391a. Internal Technote, 2005.
- [48] Y. Wah. KTeV performance parameters. Private conversation.
- [49] T. Yamanaka et al. Proposal for $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ at J-Parc. J-Parc Internal Document, 2006.
- [50] R. Kessler. KTeV analysis. Private conversation.
- [51] R. Cousins and V. Highland. Incorporating systematic uncertainties into an upper limit. *Nuclear Instruments and Methods in Physics Research A*, 320:331–335, 1992.
- [52] K. F. “Jack” Chen. Angle Measurement with Shower Shapes. Internal Memo, 2006.

- [53] S. Adler et al. Evidence for the Decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$. *Physical Review Letters*, 79(2204), 1997.
- [54] J. Nix et al. First search for $K_L \rightarrow \pi^0 \pi^0 \nu \bar{\nu}$. *Phys. Rev. D (RC)*, 76(011101(R)), 2007.
- [55] T. Inagaki et al. Detection inefficiency of a 10 mm thick plastic scintillation counter for 1 GeV/c e^\pm , π^\pm and protons. *Nuclear Instruments and Methods in Physics Research A*, 359:478–484, 1995.
- [56] T. Inagaki et al. Proposal of an Experiment at the KEK 12-GeV Proton Synchrotron: Measurement of the $K_L \rightarrow \pi^0 \nu \bar{\nu}$. KEK Internal 96-13, 1996.
- [57] J. Ma. Performance of the Outer CV according to beam-line muons. Private communication.
- [58] N. G. Deshpande and G. Eilam. Neutrino-Mass Limits from $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ Decay. *Physical Review Letters*, 53(24), 1984.
- [59] L. S. Littenberg. CP -violating decay $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$. *Physical Review D*, 39(11), 1989.