Fermilab FERMILAB-THESIS-1994-40

THE UNIVERSITY OF CHICAGO

A MEASUREMENT OF THE BRANCHING RATIO OF $\pi^0 \to e^+e^-$

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
KEVIN SCOTT MCFARLAND

CHICAGO, ILLINOIS
MARCH 1994

ACKNOWLEDGMENTS

Obviously the work represented here in this thesis was not all done by me. I owe tremendous debts to both my collaborators in this effort and to all those who have worked on the precursor experiments to FNAL-E799. I cannot hope to do justice to all of their hard work here, and so in the spirit of the many who have reached this conclusion before I did, I will not try to do so. Nevertheless there are many people who I worked closely with in my years at Chicago and Fermilab to whom I must explicitly give my deepest thanks:

To Deborah Appel Harris, who was the only other person with whom I shared all the day-to-day joys and frustrations [1] of analyzing this data. Thank you for your helping hand, for your willingness to listen when I needed to talk, for your ready advice when I was at a loss, and for being a friend through all of our long hours together. It most certainly does not suck to be us.

To Yau Wah and Bruce Winstein, for their guidance and advice through all aspects of this work: from deciding to join the group, through the analysis, to preparing to leave the group and move on. I am intellectually richer for having had this opportunity to come to Chicago and work with you.

To my other colleagues at the University of Chicago: Roy Briere, Bernhard Schwingenheuer, Sunil Somalwar, Elliott Cheu, Aaron Roodman, Peter Shawhan, Lawrence Gibbons, Tony Barker, Pat Krolak, Greg Makoff and Roland Winston. They have truly made Chicago a more productive and interesting (though not always both at the same time) environment in which to work.

Finally all my collaborators and I owe a great debt to the tremendous work done on our experiment by the Fermilab support staff. In particular, Herculean efforts by the Computer Department, the constant vigilance of the Accelerator and Operations staff, the precision work of the surveying crews, and long hours put in by Gene Beck and John Krider were key to the success of our experiment.

Of course, I do not only owe thanks to my collaborators. There are those who helped me with this work in more less direct, but possibly more important ways. In particular I must give my thanks:

To the physics lunch crowd (and occasional study group), Jim, Arthur, Monica, Mark and Debbie, who not only proved the old cliche about learning more from your classmates than from your classes in graduate school, but who also provided hours of diversion, laughter, companionship, and some fantastic bowling experiences. How did we all turn into professional physicists anyway?

To my family, particularly my mother Judy, for all their support and love through the long years of gainless unemployment that I have called my education.

To the memory of my father, Dr. James McFarland, who got me interested in experimental science in the first place. You may never have recovered from my scoffing at chemistry in favor of physics, but I hope you'd be pleased that I have followed in your footsteps in a career in academic science.

Finally, to the woman who became my wife in the middle of this work. Theresa, you've put up with my late hours and constant obsession with all of this; you've lived at Fermilab for a week at a time while I stayed out all night pulling cables; and you've listened to me talk on and on about work. Thank you for your loving support and patience through these last four years. Now that you've seen how much fun this is, I suppose you'll be wanting your own PhD?

THE FNAL-E799 COLLABORATION

Katsushi Arisaka, Doug Roberts, William Slater, Matt Spencer, Matt Weaver
University of California at Los Angeles

Roy Briere, Elliott Cheu, Deborah Harris, Pat Krolak, Greg Makoff, Kevin McFarland, Aaaron Roodman, Bernhard Schwingenheuer, Yau Wah, Bruce Winstein, Roland Winston
University of Chicago and The Enrico Fermi Institute

Tony Barker
University of Colorado

Earl Swallow
Elmhurst College and The Enrico Fermi Institute

Greg Bock, Rick Coleman, Mike Crisler, Janice Enagonio, Rick Ford, Yee Bob Hsiung, Doug Jensen, Erik Ramberg, Bob Tschirhart
Fermi National Accelerator Laboratory

Eric Collins, George Gollin
University of Illinois

Tsuyoshi Nakaya, Taku Yamanaka
Osaka University

Peter Haas, William P. Hogan, Sunkee Kim, John N. Matthews, Seong Sook Myung, Ping Gu, Steve Schnetzer, Sunil Somalwar, Gordon Thomson, Yu Zou
Rutgers University

TABLE OF CONTENTS

ACKNOWLEDGMENTS ii

THE FNAL-E799 COLLABORATION iv

LIST OF TABLES ix

LIST OF ILLUSTRATIONS xi

ABSTRACT xviii

CHAPTER

1 INTRODUCTION 1
- 1.1 Decays of the π^0 3
- 1.2 The Decay $\pi^0 \to e^+e^-$ 8
  - 1.2.1 Non-Standard Model Contributions to $\pi^0 \to e^+e^-$ 11
- 1.3 Radiative Corrections to $\pi^0 \to e^+e^-$ 12
- 1.4 Previous Measurements of $\pi^0 \to e^+e^-$ 16
- 1.5 Measurement Technique 18

2 THE DETECTOR 19
- 2.1 The K_L Beam 19
  - 2.1.1 K_L Production 20
  - 2.1.2 Formation of The Neutral Beam 23
  - 2.1.3 The Decay Volume 27
- 2.2 The Detector 29
  - 2.2.1 The Charged Spectrometer 29
    - 2.2.1.1 Material In the Spectrometer 35
  - 2.2.2 The Electromagnetic Calorimeter 39
  - 2.2.3 Other Detector Elements 43
    - 2.2.3.1 Photon Vetoes 44
    - 2.2.3.2 Trigger Planes 52
    - 2.2.3.3 The Accidental Trigger 56
- 2.3 The Experimental Trigger 58
  - 2.3.1 The Level 1 Trigger 58
  - 2.3.2 The Level 2 Trigger 60
    - 2.3.2.1 The Track Processor 61
    - 2.3.2.2 The Hardware Cluster Finder 67
  - 2.3.3 Trigger Performance 68

2.4 Data Acquisition 69
2.4.1 The Event Content 70

3 THE DATA 73
3.1 The Data Set 73
3.2 The Triggers of E799 74
3.2.1 Two Electron Triggers 77
3.2.2 The Dimuon Trigger 78
3.2.3 The Four-Track Trigger 79
3.2.4 The Neutral Trigger 79
3.2.5 The $e\mu$ Trigger 80
3.2.6 The Λ^0 Triggers 80
3.2.7 Minimum Bias Trigger 81
3.2.8 Other Triggers 81
3.2.9 Performance of E799 Triggers 82
3.3 Special Runs 82
3.4 Online Monitoring 84
3.5 The Split 85

4 CLUSTER AND TRACK RECONSTRUCTION 89
4.1 Cluster-Finding Algorithm 89
4.2 Tracking Algorithm 91
4.3 Track-Cluster Matching 101

5 CALIBRATION OF THE CHARGED SPECTROMETER
AND DETECTOR ALIGNMENT 103
5.1 Survey Information 104
5.2 Calibration of the Drift Chamber TDCs 105
5.3 Translation of Drift Times to Distances 107
5.3.1 Variation of $X(t)$ Relation with Time 111
5.3.2 Evaluation of Chamber Resolution 112
5.3.3 Time Dependence in TDC Calibration 115
5.4 Internal Chamber Alignment 116
5.4.1 Alignment with Straight-Through Tracks 116
5.4.2 Corkscrew Rotations 118
5.4.3 Quality of Internal Chamber Alignment 119
5.4.4 Changes in Chamber Alignment with Time 120
5.5 Transverse Alignment of The Detector 124
5.5.1 Target 124
5.5.2 Lead-Glass 126
5.5.3 Scintillator Banks and Photon Vetoes 128
5.5.4 Iron Ring 128
5.5.5 Beam Collimators 129
5.5.6 Upstream Pipes 130
5.6 Alignment and Field Strength of the Magnet 133

5.6.1 Field Strength 134
5.6.2 Location of the "Bend Plane" 137
5.7 Determination of Chamber Efficiencies 137
5.8 Theoretical Calculation of the Chamber Resolution 140
5.9 Conclusion 143

6 CALIBRATION AND MODELING OF THE LEAD-GLASS CALORIMETER 145
6.1 Model of Lead-Glass Response 146
6.1.1 Corrections for Non-Linearity in the Response 149
6.1.2 Corrections for Photon Response 154
6.1.3 Photostatistics 156
6.1.4 Energy Resolution 156
6.2 Determining the Energy of a Lead Glass Cluster 160
6.2.1 Leakage Off the Edge of the Array 160
6.2.2 Rate Dependent Pedestal Shift 161
6.2.3 Threshold Correction 162
6.2.4 Leakage Outside of the 3×3 163
6.2.5 The Non-Linearity Correction 163
6.2.6 Electron-Photon Difference 163
6.2.7 Bremsstrahlung Correction 163
6.3 Electron Calibration Procedure 164
6.4 Cluster Position Measurements 169
6.5 Fusion of Clusters 172
6.6 Conclusion 177

7 MONTE CARLO SIMULATION OF THE DETECTOR 179
7.1 Production and Decay 180
7.2 Tracing of Particles 182
7.3 Detector Simulations 184
7.3.1 Simulation of the Drift Chambers 184
7.3.2 Simulation of the Lead-Glass Response 186
7.3.3 Simulation of Other Detector Elements 196
7.4 Simulation of Accidental Activity 196
7.4.1 Accidental Activity in the Detector 201
7.5 Trigger Simulation 209
7.6 Conclusion 211

8 ANALYSIS OF THE $\pi^0 \to e^+e^-$ SAMPLE 215
8.1 The Initial Data Reduction 215
8.1.1 Monitoring the Crunch 222
8.2 The Analysis 224
8.2.1 Analysis of Tracks 224
8.2.2 Analysis of Calorimeter Clusters 235
8.2.3 Photon Veto Cuts 241

8.2.4 Kinematic Cuts 244
8.3 Backgrounds 252
8.4 Removal of the Four Electron Background 259
8.5 Extraction of the Branching Ratio for $\pi^0 \to e^+e^-$ 265
8.5.1 Monte Carlo Samples 271
8.5.2 Estimation of Background to the $\pi^0 \to e^+e^-$ data sample . . . 273
8.5.3 Radiative Corrections to $\pi^0 \to e^+e^-$ 274
8.6 Systematic Errors 276

9 CONCLUSIONS 279
9.1 This Result 279
9.2 All Experimental Results 281
9.3 Implications for Models of $\pi^0 \to \gamma\gamma \to e^+e^-$ 283
9.4 Constraints on non-Standard Model Processes 287
9.5 Summary 287

REFERENCES 289

LIST OF TABLES

1 Decays of the π^0 3

2 Different theoretical predictions for the real part of the amplitude for $\pi^0 \to e^+e^-$ 10

3 Effects of Be absorbers in the beam on hadron and K_L flux 27

4 Drift chamber sizes, measured in wires per sense plane 33

5 Distance from the target to elements of the charged spectrometer 35

6 Radiation lengths of material upstream of or inside the spectrometer 36

7 Radiation lengths of material downstream of the spectrometer 37

8 Summary of trigger and particle identification scintillator banks 52

9 Summary of Level 1 trigger sources 59

10 Typical rates at each stage of the trigger 68

11 Contributions to dead time 69

12 Summary of detector information written into the data stream 71

13 Rare K_L decays and the associated experimental triggers 76

14 Triggers per spill written from each trigger 77

15 Rate at each level of the trigger for each experimental trigger 83

16 The streams of the split 87

17 The tracking algorithm's classification of chamber hits 93

18 Locations along the beam direction of "immovable" detector elements from different surveys 105

19 A study of run-to-run $X(t)$ variations 113

20 Glass Calibrations in E799 164

21 Level 1 trigger veto acceptance losses as measured in accidental events . . . 202

22 Level 1 trigger rates in Accidental events 203

23 The number of chamber hits in Accidental events 204

24 Fraction of events passing each stage of the crunch 220
25 Effects of fusion cuts on data and Monte Carlo 239
26 The fraction of four-electron background events passing cuts on extra hits in the upstream drift chambers 263
27 The numbers of events in data and Monte Carlo $\pi^0 \to e^+e^-$ and $\pi^0 \to e^+e^-\gamma$ normalization samples after each cut 268
28 Monte Carlo event samples used in the analysis 272
29 The systematic errors 278

LIST OF ILLUSTRATIONS

1 Feynman diagram for $\pi^0 \to e^+e^-$ through the $\pi^0 \to \gamma\gamma$ intermediate state . 2

2 An illustration of the effect of conservation of helicity on the final-state electron spins. 2

3 The Feynman diagram illustrating the lowest order QED contribution to π^0 Dalitz decay 5

4 The Feynman diagram illustrating the vector meson contribution to Dalitz decay 5

5 π^0 form-factor slope measurements from Dalitz decays 7

6 Diagrams considered in calculating radiative corrections to $\pi^0 \to e^+e^-$ . . . 13

7 BR($\pi^0 \to e^+e^- +$ photons) as a function of x from Dalitz decay, $\pi^0 \to e^+e^- + \gamma_{\text{brem}}$, and their interference 15

8 Previous measurements of BR($\pi^0 \to e^+e^-$) 17

9 An illustration of the beam bucket occupancy autocorrelation 21

10 The energy spectrum of produced K_L's 22

11 Diagram of the beam production and collimation portion of the experimental apparatus 24

12 Formation of the beam by the collimation in horizontal ($\hat{x}$) view 25

13 Formation of the beam by the collimation in vertical ($\hat{y}$) view 26

14 Diagram of the decay volume of the detector 28

15 Schematic drawing of the detector 30

16 The sense and field wire geometry in the drift chambers 31

17 The lead-glass array 40

18 Measured radiation dose in the lead-glass 42

19 Geometry of the "square" photon vetoes 45

20 Geometry of the "round" photon vetoes 46

21 Schematic of the Back Anti counter 47

22 Cross-Section of a lead-scintillator photon veto counter 48

23 The location of the iron ring 49

24 Photon veto coverage as a function of decay vertex and photon angle 50

25 Photon veto coverage as a function of decay vertex and photon angle (continued) 51
26 Geometry of the B scintillator bank 53
27 Geometry of the C scintillator bank 54
28 Geometry of the MU1 scintillator bank 55
29 The accidental trigger counter 57
30 The use of the memories on the Zuma board 61
31 The original track processor design 63
32 The hit-counting portion of the track processor as used in E799 64
33 Illustration of track processor front-end hit location assignment 65

34 Illustration of an HCF cluster fusion 90
35 An illustration of pair finding with groups of adjacent hits 94
36 Track-Finding: $\hat{x}$-View Segment Candidates 95
37 Track-Finding: $\hat{x}$-View Track Candidates 96
38 Track-Finding: $\hat{y}$-View Segment Candidates 99
39 Track-Finding: $\hat{y}$-View Track Candidates 100

40 The raw distribution of TDC values for all wires in plane 16 of the drift chamber 106
41 An example of the measurement of wire timing offsets in the drift chamber TDCs 108
42 The missing hit correction applied to the raw drift time distribution for the $X(t)$ determination 110
43 A typical drift time to distance relation 111
44 The sum-of-distances distribution in a drift chamber 114
45 Time variation of a TDC timing offset for one wire 115
46 Chamber alignment fits using muon data 117
47 Illustration of the effect of corkscrew rotations at the decay vertex 119
48 Fit to measure the chamber corkscrew rotation 120
49 Time dependence of Chamber 2 alignment constants 121
50 Rate of change of Chamber 2 alignment constants 122
51 Time dependence of the measured corkscrew rotation 123
52 $m_{p\pi}$ for the candidate $\Lambda \to p\pi^-$ decays used for target location 125

53 The fit to measure the location of the lead-glass array in the chamber coordinate system 127

54 Illustration of the loss of large radius photons from the Iron Ring in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, sample 129

55 The upper beam intensity as a function of $\hat{x}$ and $\hat{y}$ position compared in data and Monte Carlo 131

56 The lower beam intensity as a function of $\hat{x}$ and $\hat{y}$ position compared in data and Monte Carlo 132

57 Schematic of the upstream vacuum pipes in the decay volume 133

58 $m_{\pi\pi}$ for candidate $K_L \to \pi^+\pi^-$ decays in data and Monte Carlo 135

59 Variation of the momentum scale with time as measured in $K_L \to \pi^+\pi^-$ decays 136

60 The fit to locate the z-position of the magnet "bend plane" 138

61 The chamber wire inefficiencies from one run in the upstream plane of Chamber 2Y. 139

62 Corrected and uncorrected chamber plane efficiencies in the downstream plane of Chamber 4X 139

63 Contributions to the calculated chamber resolution 141

64 Illustration of Shower Fluctuations 149

65 Shower Depth as a Function of Energy 150

66 $\log E$ vs. $\log p$ Without the Absorption Correction 151

67 The non-linearity correction for electrons as a function of energy for different absorptions 153

68 The electron-photon response difference as a function of energy for different absorptions 155

69 σ_E/E from shower fluctuations and from photostatistical smearing as a function of energy for different absorptions 157

70 σ_E/E from shower fluctuations and from photostatistical smearing as a function of absorption for different energies 158

71 Electron energy resolution versus p for electrons from $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays 159

72 Definition of the center, edge and corner blocks of the central 3×3 blocks of a lead-glass cluster 161

73 $m_{\gamma\gamma}$ in $K_L \to \pi^+\pi^-\pi^0$ decays before and after the energy correction in the pipe blocks 162

74 Deviations of the non-linearity correction from its power law parameterization 166

75 Mean absorptions of lead-glass blocks as a function of distance from a beam hole . . . 167
76 The rate of increase of absorption for each less-glass block . . . 168
77 Energy spectrum of calibration electrons in different regions of the lead-glass, compared with the spectrum of final state particles from $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$ . . . 170
78 Energy ratios used in the lead-glass cluster position algorithm . . . 171
79 A 5×5 Fusion Event . . . 174
80 A "Column" Fusion Event . . . 175
81 A "Corner" Fusion Event . . . 176

82 Multiple-scattering angular distribution in the Monte Carlo . . . 183
83 Distance from the target to the decay vertex in data and Monte Carlo for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$. . . . 185
84 "Asymmetrical Gaussian" functional form . . . 186
85 Sample fits to the lead-glass response for Monte Carlo parameterizations . . 188
86 Fit parameters used in the Monte Carlo lead glass response as a function of the absorption and initial conversion depth . . . 189
87 First pass data and Monte Carlo comparison of E/p Mean and RMS for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays . . . 190
88 Second-pass data and Monte Carlo comparison of E/p Mean and RMS versus absorption . . . 193
89 Second-pass data and Monte Carlo comparison of E/p Mean and RMS versus momentum . . . 194
90 E/p shapes by ring of blocks in the lead glass . . . 195
91 Secondary versus primary beam intensity . . . 198
92 Accidental Trigger rate versus primary beam intensity . . . 199
93 Accidental Trigger rate versus measures of secondary beam intensity 200
94 The sum-of-distances for chamber hit pairs in Accidental Triggers 204
95 The number of tracks found in Accidental Triggers . . . 205
96 The mean energy deposition in each lead-glass block for Accidental Triggers 206
97 The radiation dosage per block as measured in Accidental Triggers 207
98 Minimum Cluster Energy in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$. . . . 210
99 Total Energy in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$. . . . 212

100 Effect of No Chamber Sharing Requirement in Tracking . . . 216

101 Track Vertex Distribution from the Two-Electron Stream 217

102 Tracks Found in the Two-Electron Stream 218

103 E/p in the Two-Electron Stream 219

104 E/p vs. E/p in the Two-Electron Stream 221

105 The Yield of $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, by Run 222

106 The K_L Mass Resolution in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, by Run 223

107 Data and Monte Carlo comparison of wire illumination in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays 226

108 The electron momentum spectra for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, decays and for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays in Monte Carlo 227

109 The track segment χ^2 in data and Monte Carlo 229

110 The track offset at the magnet bend plane in data and Monte Carlo 230

111 The population of hits in each wire of the upstream plane of the $\hat{y}$-view of drift chamber 2 for $\pi^0 \to e^+e^-$ and $\pi^0 \to e^+e^-\gamma$ normalization Monte Carlo 231

112 The vertex fit χ^2 in data and Monte Carlo 233

113 The e^+e^- invariant mass of the data sample after trigger, crunch and tracking requirements. Obviously, the full reconstruction of the parent K_L decay is essential for this analysis! 234

114 The E/p for electrons in data and Monte Carlo 237

115 E/p for pions from $K_L \to \pi^+\pi^-\pi^0$ decays 239

116 The maximum number of ADC counts in blocks containing candidate soft photons in data and Monte Carlo 240

117 The photon veto response in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays, and the analysis cuts used to remove events with extra photons 242

118 Energy deposited in the $\hat{x}$-views of the first two longitudinal layers of the Back Anti 243

119 The π^0 pairing χ^2 in data and Monte Carlo 245

120 The transverse momentum with respect to the kaon direction in data and Monte Carlo for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays 247

121 The total invariant mass in data and Monte Carlo for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays 248

122 The e^+e^- invariant mass in data and Monte Carlo for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays 249

123 The total invariant mass and e^+e^- invariant mass in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, Monte Carlo 251

124 The total invariant mass vs. the e^+e^- invariant mass in the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, data sample 253

125 The the e^+e^- invariant mass in the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, data sample and $\pi^0 \to e^+e^-\gamma$ background Monte Carlo 254

126 Position of electrons at VA0 in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, candidates . . . 255

127 Misreconstruction of a four-electron final state as a high e^+e^- invariant mass two-electron final state 257

128 The π^0 pairing χ^2 for the low e^+e^- mass and high e^+e^- background compared 258

129 The total invariant mass versus the e^+e^- invariant mass for four-electron backgrounds in Monte Carlo 260

130 The the e^+e^- invariant mass in data and a background Monte Carlo including $\pi^0 \to e^+e^-\gamma$ and four electron backgrounds 261

131 The total invariant mass versus the e^+e^- invariant mass for four-electron backgrounds in Monte Carlo after cuts to remove the four-electron background 264

132 The total invariant mass vs. the e^+e^- invariant mass in the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, data sample after all cuts 266

133 The e^+e^- invariant mass in the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, data sample and $\pi^0 \to e^+e^-\gamma$ background Monte Carlo after all cuts. 267

134 The acceptance and number of decays as a function of run number for the $\pi^0 \to e^+e^-\gamma$ normalization sample 270

135 The predicted e^+e^- invariant mass spectrum from radiative corrections to the $\pi^0 \to e^+e^-$ process 275

136 The branching ratio for $\pi^0 \to e^+e^-$ for $x > x_{\rm cut}$ as a function of $x_{\rm cut}$ 280

137 All experimental measurements of $\mathrm{BR}(\pi^0 \to e^+e^-)$ 282

138 Probability distributions for $,{}^{0}_{ee}/,{}_{\rm all}$ from the three most significant measurements of $\pi^0 \to e^+e^-$ and the combined probability distribution 284

139 Probability distributions for the real part of the amplitude for $\pi^0 \to e^+e^-$ from this experiment alone and from the best fit from all experimental data 284

140 The implications of the measurement of $\eta \to \mu^+\mu^-$ and $\pi^0 \to e^+e^-$ for the sign of the real part of the amplitude for $\eta \to \mu^+\mu^-$ 286

For Tess and Baby

ABSTRACT

This thesis describes a measurement of the rare decay $\pi^0 \to e^+e^-$. The π^0's studied were produced in a neutral kaon beam via the decay in flight of $K_L \to \pi^0\pi^0\pi^0$. We have observed nine events consistent with $\pi^0 \to e^+e^-$, and we measure

$$\mathrm{BR}(\pi^0 \to e^+e^-, (\frac{m_{ee}}{m_{\pi^0}})^2 > 0.95) = (7.6^{+3.9}_{-2.8}(\mathrm{stat}) \pm 0.5(\mathrm{syst})) \times 10^{-8},$$

over and above the contribution to this rate from Dalitz decays ($\pi^0 \to e^+e^-\gamma$). This measurements probes the $\pi^0\gamma\gamma$ decay vertex and limits non-Standard-Model processes that can contribute to this decay.

CHAPTER 1

INTRODUCTION

The neutral pion is the lightest known hadron, and electrons are the lightest charged leptons. No fundamental symmetry prohibits the direct decay of the π^0, which is after all a quark-antiquark pair, into e^+e^-, a lepton-antilepton pair. However, models for this decay predict the rate for $\pi^0 \to e^+e^-$ to be more than a factor of 10^7 smaller than that for $\pi^0 \to \gamma\gamma$, the most common π^0 decay.

Why is such a rare process an interesting subject of study? In general, rare processes are forbidden or suppressed because they violate exact or approximate conservation laws. When this suppression is large, the rare decay can become an effective laboratory for detecting small violations of the symmetries from which these conservation laws arise. For example, the unexpected violation of the discrete symmetry CP was discovered with the observation of the rare decay $K_L \to \pi^+\pi^-$ [2]. Measurements of rare processes can also offer an opportunity to test accepted models in new ways. In the rare decay of interest here, $\pi^0 \to e^+e^-$ can be used to measure the structure of the $\pi^0\gamma\gamma$ vertex in different kinematic regions from those probed in previous measurements.

The $\pi^0 \to e^+e^-$ decay is expected to be suppressed because it is presumed to proceed through the $\pi^0 \to \gamma\gamma$ intermediate state. The Feynman diagram representing this process is shown in Figure 1. Because of the vector nature of electromagnetic interactions, helicity is conserved at each $ee\gamma$ vertex. This has an effect on the final state electron spins, as shown in Figure 2. If the electron were massless, the helicity conservation would be equivalent to a conservation of S_z, the electron spin along the decay axis. This would result in an $S = 1$ two-electron final-state, which is forbidden by conservation of angular momentum and parity. However, because the electron is massive, there is some overlap between a given helicity state and the opposite spin state. This results in a factor of $(1-\beta)$ at one of the $ee\gamma$

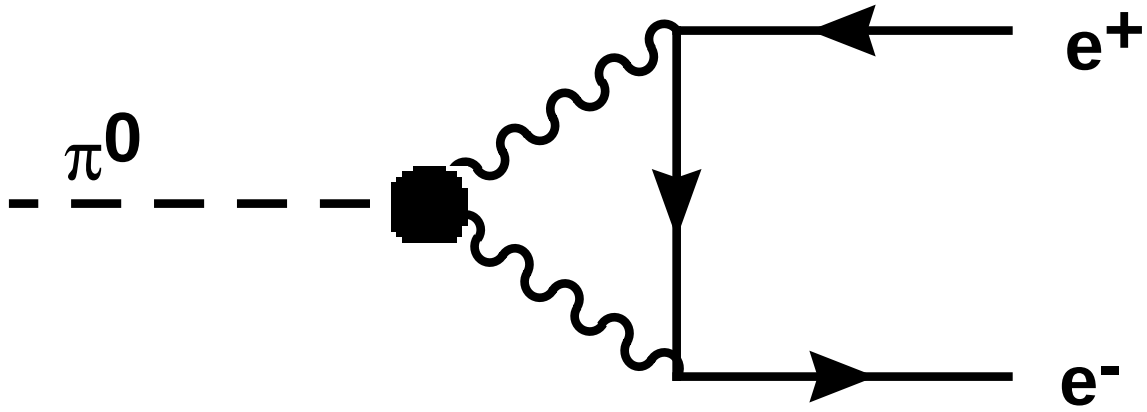

Figure 1. Feynman diagram for $\pi^0 \to e^+e^-$ through the $\pi^0 \to \gamma\gamma$ intermediate state

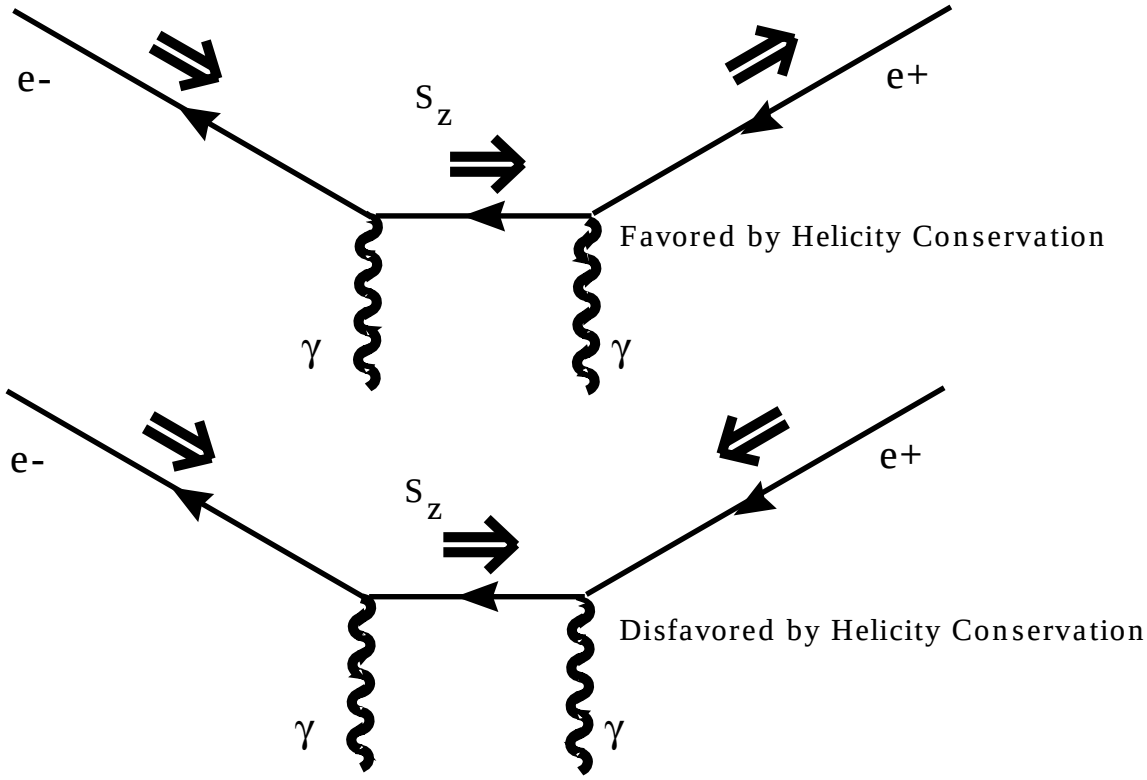

Figure 2. An illustration of the effect of conservation of helicity on the final-state electron spins.

Table 1. Decays of the π^0

Final State	Branching Fraction	Reference
$\gamma\gamma$	$(98.798 \pm 0.032)\%$	[3]
$e^+e^-\gamma$	$(1.198 \pm 0.032)\%$	[3]
$e^+e^-e^+e^-$	$(3.14 \pm 0.30) \times 10^{-5}$	[4]
$\nu\nu$	$< 8.3 \times 10^{-7}$, 90% C.L.	[5]
$\gamma\gamma\gamma$	$< 3.1 \times 10^{-8}$, 90% C.L.	[6]
$\gamma\gamma\gamma\gamma$	$< 2 \times 10^{-8}$, 90% C.L	[6]
$e^+\mu^- + e^-\mu^+$	$< 1.7 \times 10^{-8}$, 90% C.L.	[7]

vertices, where β is v/c of the electron in the π^0 rest frame. This gives a helicity suppression factor of

$$\left(\frac{2m_e}{m_{\pi^0}}\right)^2 .$$

Also, relative to , $(\pi^0 \to \gamma\gamma)$, , $(\pi^0 \to e^+e^-)$ is reduced by two factors of $\alpha_{\rm EM}$, one from each $ee\gamma$ vertex. However, if other types of interactions exist which contribute to the rate for $\pi^0 \to e^+e^-$, these suppression factors will not necessarily be present. Therefore, an accurate measurement of the branching ratio can serve to detect or to limit other such possible interactions.

1.1 Decays of the π^0

A list of observed π^0 decays, aside from $\pi^0 \to e^+e^-$, along with limits on other decay channels are shown in Table 1. By far the most common decay of the π^0 is its decay into two photons. In some sense, this is the *only* π^0 decay, since all other known decays of the π^0 are consistent with proceeding through a two-photon intermediate state. If this picture of the π^0 decay were correct, the only things to measure would be the π^0 lifetime and the structure of the $\pi^0\gamma\gamma$ vertex. Unfortunately, the decay $\pi^0 \to \gamma\gamma$ is not a good place to observe this vertex structure since it occurs at a fixed kinematic point, namely the point where the two photons are on mass shell.

The only observed decays of the π^0, other than $\pi^0 \to \gamma\gamma$, are the so-called "Dalitz" decays which involve the internal conversion of photons into e^+e^- pairs. The single-Dalitz decay ($\pi^0 \to e^+e^-\gamma$) rate and e^+e^- invariant mass spectrum can be evaluated in a straightforward QED calculation [8][9], assuming no structure in the $\pi^0\gamma\gamma$ vertex and including only the lowest order diagram shown in Figure 3. Define

$$\rho \equiv \frac{,(\pi^0 \to e^+e^-\gamma)}{,(\pi^0 \to \gamma\gamma)}.$$

In calculations neglecting radiative corrections, the calculated ratio of rates obtained is denoted here by ρ_0, and

$$\rho \equiv \rho_0 + \rho_{\rm rad}.$$

The differential rate for Dalitz decays with respect to the e^+e^- invariant mass is calculated to be [9]

$$\frac{d\rho_0}{dx} = \frac{2\alpha_{\rm EM}}{3\pi}\frac{1}{x}(1-x)^3\left(1+\frac{r}{2x}\right)\sqrt{1-\frac{r}{x}},$$

where x is $(m_{ee}/m_{\pi^0})^2$, and r is the minimum value of x, $(2m_e/m_{\pi^0})^2$. Note that the e^+e^- invariant mass spectrum is sharply peaked at low invariant mass. The total rate calculated in this manner is given by [8]

$$\rho_0 = 0.01185,$$

which is in agreement with branching ratio measurements [3].

The effects of including higher-order (virtual and radiative) QED processes in this total rate have been calculated analytically [10][11][12] and numerically [13] with the result that the effect on the total branching ratio is extremely small,

$$\rho^{\rm rad} \approx 1.04 \times 10^{-4}.$$

Effects in the differential rate, $d\rho_{\rm rad}/dx$, however, are significantly larger, and distort the observed e^+e^- invariant mass spectrum.

Another possible deviation from the above calculation can result from structure in the $\pi^0\gamma\gamma$ vertex. Such non-pointlike behavior can result from diagrams like the one shown in Figure 4 where there are intermediate vector mesons. However, because the mass-scale for these effects is much larger than masses of the virtual photons in Dalitz decays, these effects

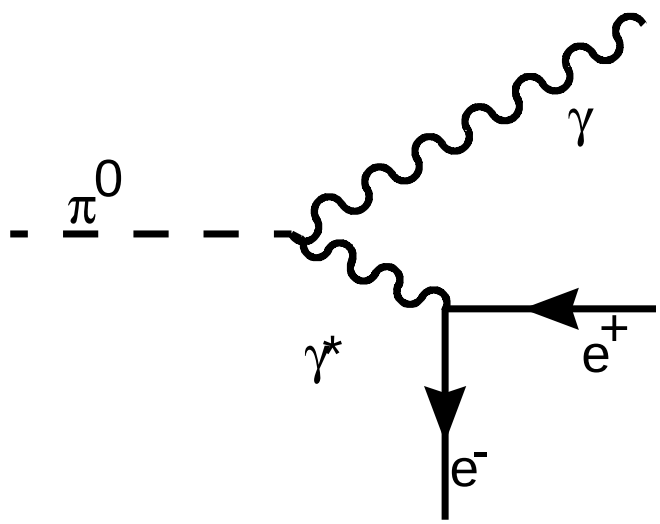

Figure 3. The Feynman diagram illustrating the lowest order QED contribution to π^0 Dalitz decay

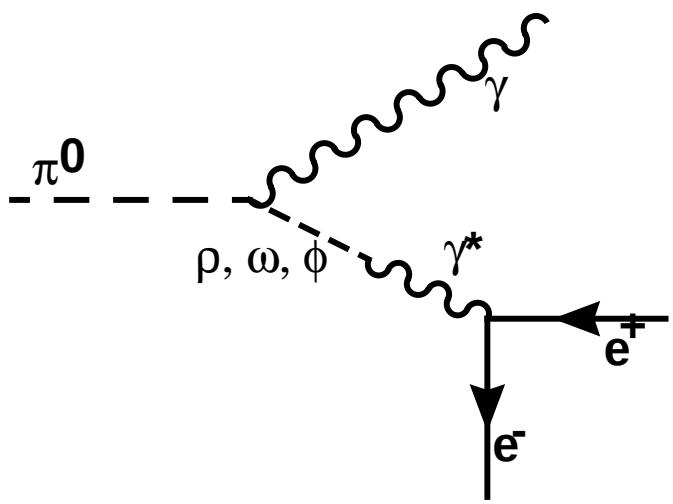

Figure 4. The Feynman diagram illustrating the vector meson contribution to Dalitz decay

are expected to be small. Independent of models, one parameterizes the effects of structure in the $\pi^0\gamma\gamma$ vertex by a form-factor, which is a function of the square of the four-momenta of the three particles at the vertex. This form-factor is denoted by

$$F(k_1^2, k_2^2; m_{\pi^0}^2),$$

where k_1 and k_2 are the four-vectors of the two emitted photons. By convention, this form-factor is normalized so that

$$F(0, 0; m_{\pi^0}^2) \equiv 1.$$

In a Dalitz decay with one real photon, $k_2^2 = 0$, so one can write

$$F(k_1^2, 0; m_{\pi^0}^2) \equiv F(x).$$

As first pointed out by Berman and Geffen [14], the analyticity of the form-factor, coupled with its expected small size from strong interaction effects, allows the form-factor effect in the differential Dalitz decay rate to be approximated by

$$F(x) \approx (1 + ax).$$

In the Dalitz decay rate, this form factor enters into the differential rate as

$$\frac{d\rho}{dx} = \frac{d\rho_{\mathrm{QED}}}{dx} \times |F(x)|^2 \,. \approx \frac{d\rho_{\mathrm{QED}}}{dx}(1 + ax)^2$$

Theoretical predictions for the form-factor slope, a, in different models are summarized in a recent review by Landsberg [15]. The most successful model in other similar systems appears to be a Vector Meson Dominance (VMD) model, which predicts a to be approximately the square of the ratio of the masses of the π^0 and the intermediate vector meson, a ρ^0 for example. Thus, VMD predicts $a = (m_{\pi^0}/m_\rho)^2 \approx 0.03$. Baryon Loop Models [16] predict very low form-factor slopes, in contradiction with experiments measuring form-factor slopes in other pseudoscalar meson decays. Potential models [17] and Non-local Quark Models (NQM) [18] still are viable competing candidate models. The latter model makes a prediction for the form-factor effect in π^0 decays which differs somewhat from the VMD model.

Slope (a)	Notes	Date	Reference
-0.24 ± 0.16	rate analysis	1961	[4]
-0.15 ± 0.10	stat. error only	1961	[19]
$+0.01 \pm 0.11$		1969	[20]
$+0.02 \pm 0.10$		1972	[21]
$+0.10 \pm 0.03$	stat. error only	1978	[22]
-0.01 ± 0.08	rate analysis	1987	[23]
-0.11 ± 0.08		1989	[24]
$+0.025 \pm 0.030$		1992	[25]
$+0.026 \pm 0.054$	rate analysis	1992	[26]

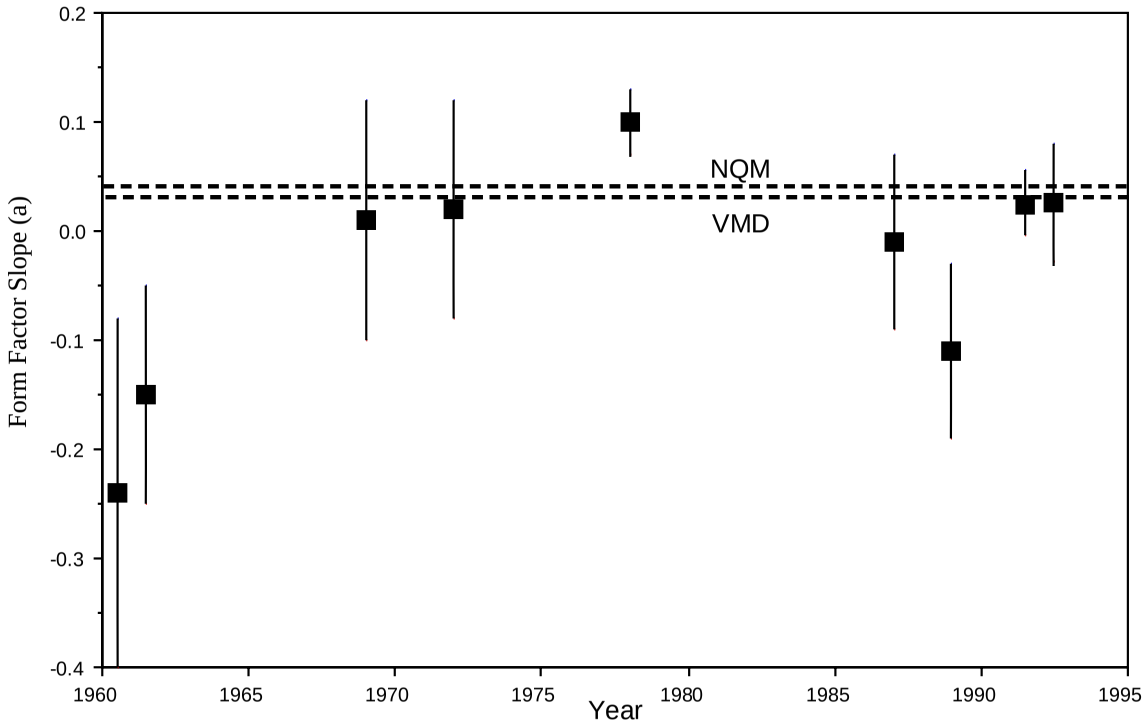

Figure 5. π^0 form-factor slope measurements from Dalitz decays

The form-factor does not produce measurable changes in the predicted $\pi^0 \to e^+e^-\gamma$ branching ratio. However, the effect of a form-factor slope can be observed by looking at the differential e^+e^- invariant mass spectrum in these decays. The history of such measurements is shown in Figure 5. A distinction is made between experiments which use a "rate analysis", which measures the rate of decays in high mass part of the e^+e^- mass spectrum, normalized to the number of π^0's produced, and those experiments which use "shape" analyses which attempt to fit the shape of the e^+e^- invariant mass spectrum from Dalitz decay. The best single experiment [25] obtains $a = 0.025 \pm 0.014(\text{stat}) \pm 0.026(\text{syst})$. Measurements of this same form factor were performed by the CELLO collaboration in a different kinematic region using the reaction $e^+e^- \to e^+e^-\pi^0$ [27] which translates to a form-factor slope in Dalitz decays of $a = 0.0326 \pm 0.0026$, much more precise than any measurement from Dalitz decays.

The branching ratio for π^0 double-Dalitz decay, $\pi^0 \to e^+e^-e^+e^-$, can also be calculated and compared with experiment. However, more interesting in this decay is the examination of final-state angular distributions. Because the e^+e^- are emitted preferentially in the plane of the electric-field vector of the parent photon, correlations between the planes of e^+e^- pairs in $\pi^0 \to e^+e^-e^+e^-$ decay allow a determination of the relative polarities of final-state photons in $\pi^0 \to \gamma\gamma$ decay. This decay provides the only direct measurement of the π^0 parity (which is negative) in π^0 decay [28].

1.2 The Decay $\pi^0 \to e^+e^-$

The role that the rare decay $\pi^0 \to e^+e^-$ might play in measurements of the π^0 form factor was first discussed by Drell in 1959 [29]. The calculation of the decay rate for the process shown in Figure 1 in terms of the π^0 form-factor, $F(k_1^2, k_2^2; m_{\pi^0}^2)$, is completely straightforward. Following the notation of Bergström, Massó *et al.* [30],

$$\frac{,(\pi^0 \to \gamma^*\gamma^* \to e^+e^-)}{,(\pi^0 \to \gamma\gamma)} = 2\sqrt{1 - \frac{4m_e^2}{m_{\pi^0}^2}} \left(\frac{\alpha_{\text{EM}}}{\pi}\frac{m_e}{m_{\pi^0}}\right)^2 |R|^2 ,$$

where R, which contains the dynamical aspects of the decay, is given by

$$R \equiv -\frac{2\imath}{\pi^2 m_{\pi^0}^2} \int d^4k_1 \frac{q^2k_1^2 - (q\cdot k_1)^2 F(k_1^2, (q-k^1)^2; m_{\pi^0}^2)}{[k_1^2 + \imath\varepsilon]\,[(q-k_1)^2 + \imath\varepsilon]\,[(k_1-p)^2 - m_e^2 + \imath\varepsilon]},$$

where q and p are the four-vectors of the π^0 and the outgoing electron, respectively. The imaginary part of R, which is the contribution to the amplitude corresponding to two real photons, is independent of F (which is normalized to 1 at the point where the photons are real), and is given by

$$\mathrm{Im}\{R\} = -\frac{\pi}{2\beta_\pi} \log \frac{1+\beta_\pi}{1-\beta_\pi},$$

where $\beta_\pi^2 = 1 - 4m_e^2/m_{\pi^0}^2$. This implies a bound on $|R|^2$,

$$|R|^2 \geq (\mathrm{Im}\{R\})^2 ,$$

which is known as the "unitary limit" for this decay mode. Therefore,

$$\frac{,(\pi^0 \to \gamma^*\gamma^* \to e^+e^-)}{,(\pi^0 \to \gamma\gamma)} > 4.75 \times 10^{-8}.$$

However, in the limit of a pointlike interaction at the $\pi^0\gamma\gamma$ vertex, (i.e. constant F), the real part of R diverges logarithmically. Therefore, there must be a "cut-off" point beyond which the form-factor becomes small for the real part of R to not diverge. Drell interpreted this cut-off in F as corresponding to a "pion size", and wrote

$$F(0,0;Q^2) = \begin{cases} 1 & Q^2 < \Lambda^2 \\ 0 & Q^2 > \Lambda^2 \end{cases} .$$

This leads to a prediction for the real part of the amplitude of

$$|\mathrm{Re}\{R\}| = \left[\log \frac{\Lambda}{m_{\pi^0}} \log \frac{\Lambda}{m_e}\right] .$$

This cut-off corresponds to a "pion size" in the sense that one might expect contributions to this vertex to fall off at sizes greater than, say, the nuclear Compton wavelength, $1/M_N$, where M_N is the mass of a nucleon.

More "realistic" models for the form-factor can, of course, produce different results for the rate of $\pi^0 \to e^+e^-$. A summary of results from various models is shown in Table 2. Most theories predict a smaller contribution from the real portion of the amplitude than from the imaginary part. These predictions suggest

$$5.2 \times 10^{-8} < \frac{,(\pi^0 \to \gamma^*\gamma^* \to e^+e^-)}{,(\pi^0 \to \gamma\gamma)} < 7.2 \times 10^{-8},$$

Table 2. Different theoretical predictions for the real part of the amplitude for $\pi^0 \to e^+e^-$

$\frac{\lvert\mathrm{Re}\{R\}\rvert}{\lvert\mathrm{Im}\{R\}\rvert}$	Model	Ref.
0.81 if $\Lambda = M_N$	Unsubtracted Dispersion Relation with Phenomenological Cutoff $F(0,0;Q^2) = \theta(\Lambda^2 - Q^2)$	[29]
0.53 if $\Lambda = M_N$	Phenomenological Cutoff $F(k_1^2, k_2^2; m_{\pi^0}^2) = \Lambda^2/(\Lambda^2 - k_1^2 - k_2^2)$	[14]
0.63	Vector Meson Dominance (ρ, ω, ϕ)	[31]
0.53	VMD (Dominance of Quark Triangle Anomaly)	[32]
0.56	VMD with asymptotic behavior from QCD	[33]
0.59 ± 0.03	VMD	[34]
1.40	Nucleon Loop Model	[16]
0.28	Nucleon Loop Model (correction)	[35]
0.36	Non-Local Quark Model	[36]
3.2	A Once-Subtracted Dispersion Relation (Criticized because the F obtained is unphysical)	[37] [38]
0.48	Quark Triangle Loop (QTL) Model in Soft Meson Limit	[35]
0.72	QTL, $m_q = 100$ MeV	[39]
0.47	$m_q = 350$ MeV	
0.53	"Q^2 Duality": results are relatively insensitive to details of model, and only depend on cut-off. Thus, same results for VMD and QTL	[30]
0.5-0.8	Chiral Perturbation Theory Calculation based on BR$(\eta \to \mu\mu)$ [41]	[40]

with the VMD and Quark Triangle Loop (QTL) models favoring values of

$$\frac{,(\pi^0 \to \gamma^*\gamma^* \to e^+e^-)}{,(\pi^0 \to \gamma\gamma)} \approx 6.5 \times 10^{-8}.$$

It has been pointed out [30] that one expects many of these seemingly different approaches to produce similar results. In particular in the case of VMD models and constituent quark models, the results are quantitatively similar because the real part of the amplitude is mostly sensitive to the value of the cut-off in the form-factor, and not to the details of the model.

A recent calculation [40] has used the measured process $\eta \to \mu^+\mu^-$ as input to a form factor model-independent calculation using Chiral Perturbation Theory. This results in a prediction of

$$\frac{,(\pi^0 \to \gamma^*\gamma^* \to e^+e^-)}{,(\pi^0 \to \gamma\gamma)} = (7 \pm 1) \times 10^{-8},$$

which is consistent with the model-dependent results quoted above.

1.2.1 Non-Standard Model Contributions to $\pi^0 \to e^+e^-$

Of course, this need not be the entire story. Because of its large suppression in the Standard Model, the decay $\pi^0 \to e^+e^-$ has been considered as a possible candidate system for viewing non-Standard effects. In particular, this decay could have large components from certain types of direct lepton-hadron coupling, or from π^0 mixing with some exotic pseudoscalar capable of coupling to leptons.

The first suggestions that $\pi^0 \to e^+e^-$ might be a good place to look for non-Standard Model interactions came with the discovery of the increase with increasing energy of

$$R = \frac{\sigma(e^+e^- \to \text{hadrons})}{\sigma(e^+e^- \to \mu^+\mu^-)}$$

at CEA in 1973 [42]. One possible explanation for such an effect is a direct electron-hadron coupling. With the hindsight of history, of course, we now know that this rising cross-section was due to the effects of the as yet undiscovered charm quark. Nevertheless, it was suggested at the time by Mel Schwartz [43] that if such a direct coupling were responsible for this effect, it might also be observable in decays of pseudoscalars into electron pairs.

This idea was further developed in the context of the CEA and similar SPEAR results [44] and later generalized [45].

Suggestions have also been made that the decay $\pi^0 \to e^+e^-$ could be sensitive to effects from light pseudoscalar Higgs particles [46] or from certain types of axions [47]. Limits that can be placed on exotic contributions will be discussed in the conclusions drawn in Chapter 9.

1.3 Radiative Corrections to $\pi^0 \to e^+e^-$

An important complication for any measurement of $\pi^0 \to e^+e^-$ are the effects of final-state radiation. These effects can not only change the total observed amplitude, but also will introduce a radiative tail in the e^+e^- invariant mass spectrum from internal bremsstrahlung. Since distinguishing $\pi^0 \to e^+e^-$ from $\pi^0 \to e^+e^-\gamma$ requires cutting tightly on the e^+e^- invariant mass and removing events with extra photons, the presence of this tail will reduce the number of identified $\pi^0 \to e^+e^-$ events. Thus, it is important to account for radiative effects in extracting an experimental measurement of the rate for this process.

The model chosen for calculating radiative effects [48] assumes a pointlike pseudoscalar constant coupling at a $\pi^0e^+e^-$ vertex. This results in a model-independent treatment of the radiative effects. One might ask if this is a good approximation for contributions from the intermediate $\pi^0 \to \gamma\gamma$ state because it neglects photon emission from the internal electron line. This assumption is argued [48] to be reasonable since this electron will be far off the mass shell for the loop momenta which dominate the electromagnetic amplitude for $\pi^0 \to e^+e^-$. Because bremsstrahlung is only prominent for particles near mass shell, the radiation from this internal electron can safely be neglected. Also note that the assumption of a constant coupling independent of the final state e^+e^- invariant mass is probably not such a good approximation for very hard bremsstrahlung emission. The bulk of the bremsstrahlung is, however, soft, so any error from this approximation is probably quite small in any observable rate. This model results in three diagrams which contribute at order $\alpha_{\rm EM}$ which are shown in Figure 6.

The matrix element for the bremsstrahlung processes shown in Figure 6 is

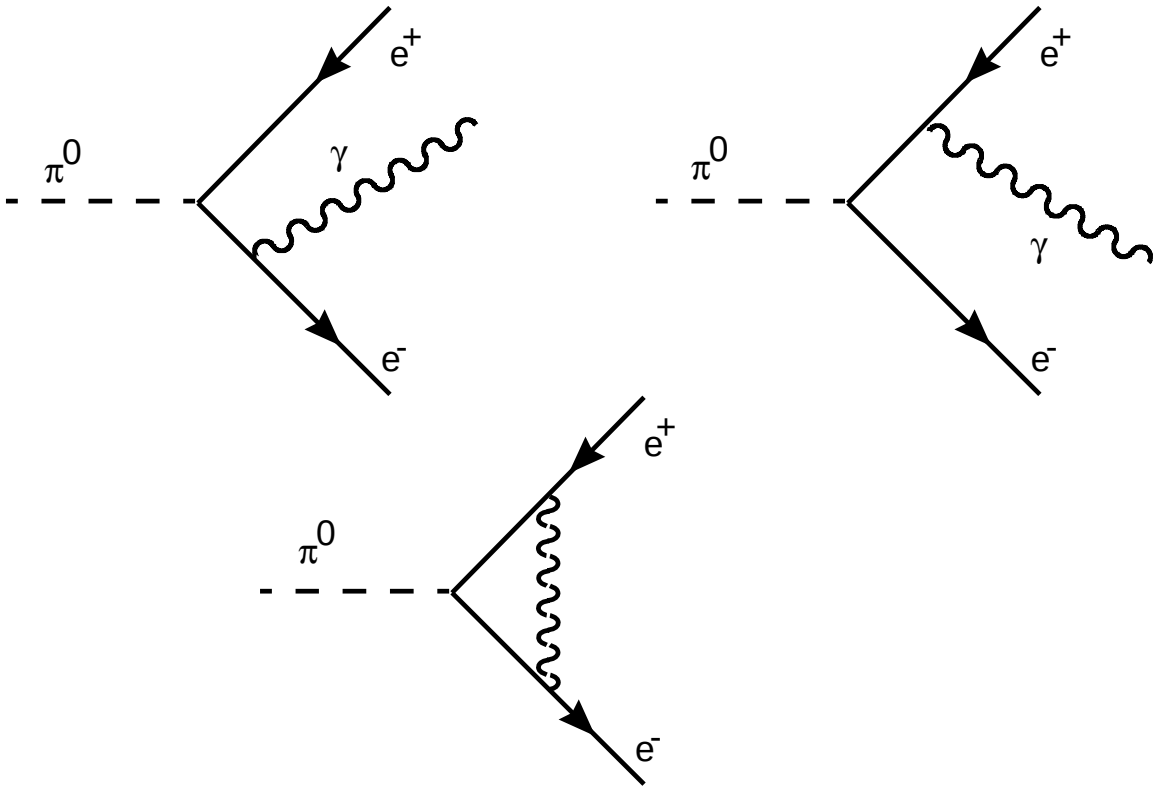

Figure 6. Diagrams considered in calculating radiative corrections to $\pi^0 \to e^+e^-$

$$M = -eg\overline{u}(p_-)\left[\not{\epsilon}\frac{\imath}{\not{p}_- + \not{k} - m_e + \imath\varepsilon}\gamma_5 + \gamma_5\frac{\imath}{-\not{p}_+ + \not{k} - m_e + \imath\varepsilon}\not{\epsilon}\right]v(p_+),$$

where g is the strength of the pseudoscalar coupling, p_+ and p_- are the four-momenta of the positron and electron and k and ϵ are the four-momentum and polarization of the bremsstrahlung photon. This matrix element was then evaluated in a straightforward manner to calculate the bremsstrahlung diagrams, and the vertex correction diagram was calculated using a subtracted dispersion relation [48]. The results are that the total rate for the radiative diagrams relative to the tree-level diagram is given by

$$\frac{\Gamma^{rad}}{\Gamma^{0}_{ee}} = \frac{\alpha_{\rm EM}}{\pi}\left[\frac{3}{2}\log\frac{1-\beta_\pi}{1+\beta_\pi} + \frac{9}{4} + \mathcal{O}\left(\frac{m_e^2}{m_{\pi^0}^2}\right)\right],$$

which is a fairly small negative correction, approximately 3%. The effects in the differential rate from bremsstrahlung are larger, however, and are given by

$$\frac{1}{\Gamma^{0}_{ee}}\frac{d\Gamma^{rad}}{dx} = \frac{\alpha_{\rm EM}}{\pi}\left[\frac{x^2+1}{1-x}\log\frac{1-\beta(x)}{1+\beta(x)} - \frac{2x\beta(x)}{1-x}\right],$$

where $\beta(x) = \sqrt{1 - 4m_e^2/x m_{\pi^0}^2}$.

Once one includes radiative corrections, the distinction between Dalitz decays and $\pi^0 \to e^+e^-$ begin to blur. The addition of a bremsstrahlung photon to the latter can, after all, make it have the same final state as the former. To be clear then, by $\pi^0 \to e^+e^-\gamma$ or Dalitz decay, I mean the process shown in Figure 3 which involves an internal photon conversion. By $\pi^0 \to e^+e^-$, I mean diagrams which can procedure a photon-free final state if the electrons do not radiate, such as the one shown in Figure 1 or a direct $\pi^0 e^+e^-$ coupling.

The differential rate as a function of x is shown in Figure 7 for three processes that contribute to $\pi^0 \to e^+e^-\gamma$ at $\mathcal{O}(\alpha_{\rm EM})$, π^0 Dalitz decay, $\pi^0 \to e^+e^-$ with radiation and the interference between the two [48]. $\pi^0 \to e^+e^-$ appears as a small bump at the endpoint of the Dalitz e^+e^- invariant mass spectrum. The magnitude of the interference between the two processes cannot be determined because the phase of $\pi^0 \to e^+e^-$ is unknown. However, even assuming the maximum possible magnitude, the interference between the two is negligible in any region of x because the two processes populate different regions of phase space. However, it is worth noting that integrated over all x the total rate for interference can actually greater that the total rate from $\pi^0 \to e^+e^-$ itself!

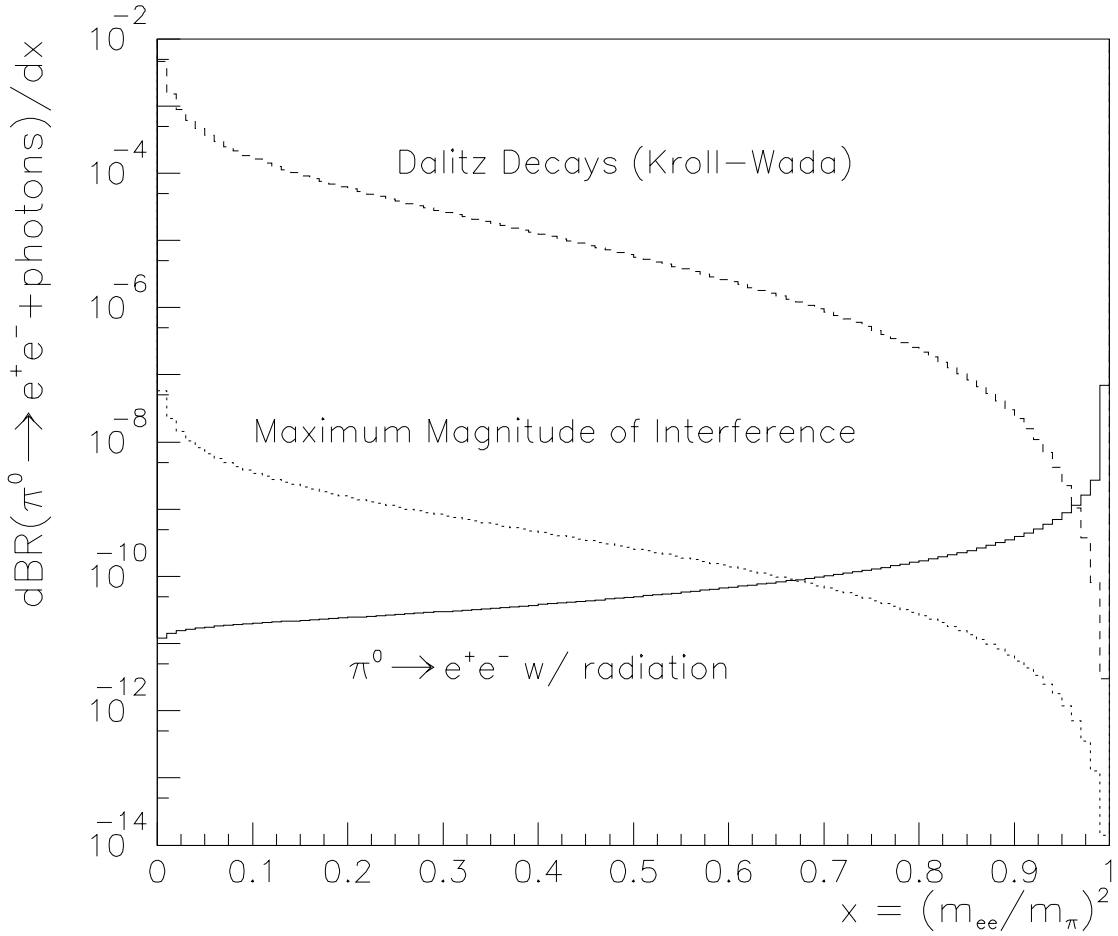

Figure 7. BR($\pi^0 \to e^+e^-$ + photons) as a function of x from Dalitz decay, $\pi^0 \to e^+e^- + \gamma_{\rm brem}$, and their interference

1.4 Previous Measurements of $\pi^0 \to e^+e^-$

Before this search, there were several experimental efforts to look for $\pi^0 \to e^+e^-$. The results of these past experiments are shown in Figure 8. The most striking feature of these results is the apparent discrepancy between the two early observations of $\pi^0 \to e^+e^-$from the Geneva-Saclay and LAMPF groups with a combined determination of BR($\pi^0 \to e^+e^-$)= $(19^{+6}_{-5}) \times 10^{-8}$, and the limit from the SINDRUM collaboration of BR($\pi^0 \to e^+e^-$)$< 13 \times 10^{-8}$ at 90% confidence level. The minimum χ^2 for the consistency of these three results is 4, and therefore the three experiments are consistent at only the 5% confidence level. Thus it makes no sense to combine all of the available data to obtain a "best-fit" result.

The two production techniques used in previous experiments were the decay in flight of $K^+ \to \pi^+\pi^0$ and interactions of π^-'s in a target, producing π^0's through the reaction $p\pi^- \to n\pi^0$. Since the π^0 has a very short lifetime, $c\tau \approx 25$ nm, the π^0 decay is effectively instantaneous, and thus the e^+e^- pair in $\pi^0 \to e^+e^-$ is created at the point of π^0 production.

There are advantages of each of these two production techniques. Both methods allow for a tagged production of the π^0, by detecting the π^+ in K^+ decay and by detecting the neutron from the target in charge-exchange. However, in decay in flight of the K^+, there is less material through which the final-state electrons must pass than in the case of charge-exchange production in a target. This results in a better measurement of the initial e^+e^- invariant mass with less smearing from multiple scattering and bremsstrahlung, which then allows a better rejection of backgrounds from $\pi^0 \to e^+e^-X$. On the other hand, if the incident π^- beam is stopped in the production target as was done in the SINDRUM measurement, the kinematics of the decay are completely determined, and the neutron time of flight can be used to discriminate between π^0 production and background processes not involving π^0's. The biggest problem, however, with both techniques is a potentially severe background problem from non-resonant e^+e^- production in the processes $K^+ \to \pi^+e^+e^-$ and $p\pi^- \to ne^+e^-$ which are indistinguishable from the signal if $m_{e^+e^-} = m_{\pi^0}$.

Branching Ratio (units of 10^{-8})	Technique	Date	Experiment
22.3^{+24}_{-11}	K^+ decay in flight (90% error bars)	1978	Geneva-Saclay [49]
$17 \pm 6(\text{stat}) \pm 3(\text{syst})$	$\pi^- p \to \pi^0 n$ in flight	1983	LAMPF [50]
< 53	$\pi^- p \to \pi^0 n$ in flight (90% confidence)	1987	OMICRON [51]
< 13	$\pi^- p \to \pi^0 n$ at rest (90% confidence)	1989	SINDRUM [52]

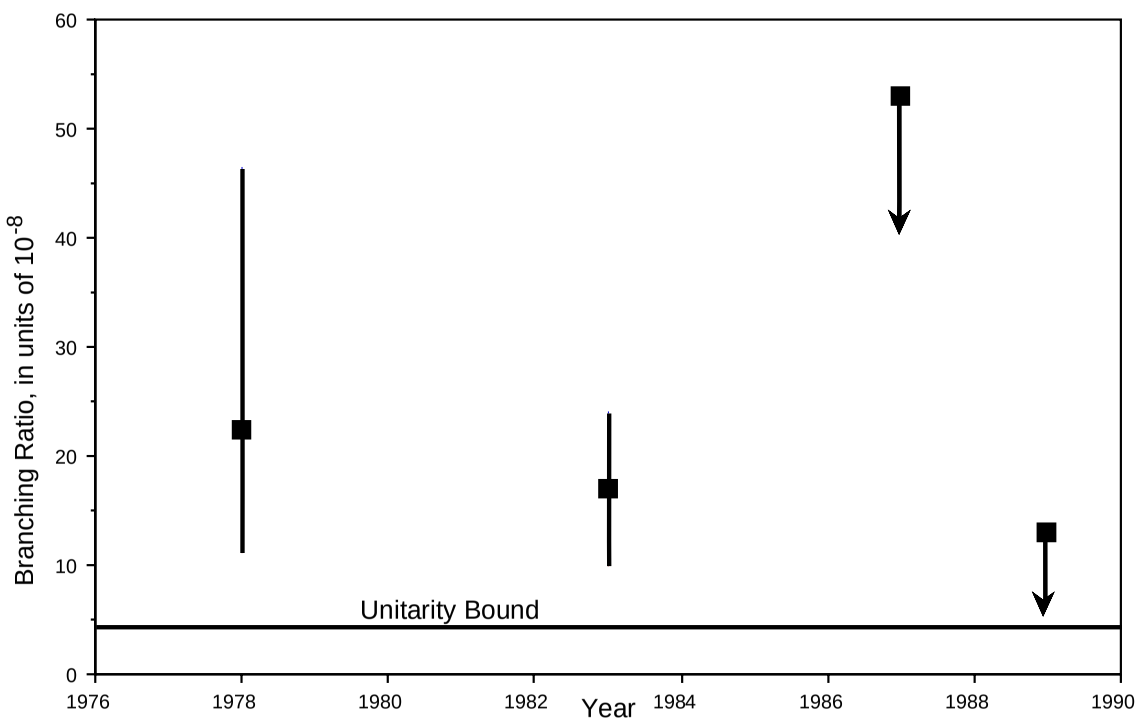

Figure 8. Previous measurements of BR($\pi^0 \to e^+e^-$). The first two experiments observed these decays and reported a branching ratio measurement. The last two did not and reported upper limits on the branching ratio at 90% confidence.

1.5 Measurement Technique

This new measurement of $\pi^0 \to e^+e^-$ was undertaken as part of the E799 experiment at Fermilab. This experiment was one of a series of experiments using essentially the same detector and beamline, all concerned with high statistics measurements of neutral kaon decays. Previous experiments using this apparatus, E731 and E773, were designed to measure CP and CPT violating parameters in the $K_{L,S} \to \pi\pi$ system. E799, by contrast, was a dedicated experiment to search for and study rare decays of the K_L.

π^0's in E799 were produced through the decay in flight $K_L \to \pi^0\pi^0\pi^0$. In addition to having a large branching fraction, this decay also produces three π^0's, any of which can decay to e^+e^-, thus tripling the π^0 production. This decay has the advantages of the similar K^+ decay in flight production in that the decays occur in a vacuum and in that the π^0 decay can be tagged by the full reconstruction of the parent K_L decay. One disadvantage relative to the K^+ production technique is the six body final state of the K_L decay, with two π^0's each decaying to $\gamma\gamma$ and one to e^+e^-. This high multiplicity final state results in a lower geometrical acceptance for full reconstruction of the decay. However, a huge advantage of the K_L production is that there are no common K_L decays with two electrons in the final state. The branching ratio for the continuum e^+e^- production process, $K_L \to \pi^0\pi^0e^+e^-$, is expected to be 2×10^{-8} with an e^+e^- invariant mass spectrum mostly well below the π^0 mass [53]. This means that backgrounds to $\pi^0 \to e^+e^-$ will be entirely from π^0 decays which makes the experimental situation much more clean than in previous measurements.

The full reconstruction of $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, requires good multi-body acceptance, photon calorimetry, good e^+e^- invariant mass resolution, and of course a large number of K_L decays. The production of K_L's and the detector will be described in Chapter 2, and then the data set acquired in E799 and experimental triggers will be discussed in Chapter 3. Chapter 4 will then describe the reconstruction of photons and electrons in the detector. In Chapters 5 and 6 the calibration of the spectrometer and calorimeter, respectively, are described, and the detector simulation is discussed in Chapter 7. Finally, the search for $\pi^0 \to e^+e^-$ is described in Chapter 8, and conclusions based on the extracted measurement of $\pi^0 \to e^+e^-$ are drawn in Chapter 9.

CHAPTER 2

THE DETECTOR

Like following life thro' creatures you dissect,
You lose it in the moment you detect.
—*Alexander Pope*

In the previous chapter, the principle of measuring π^0 decays by fully reconstructing $K_L \to \pi^0\pi^0\pi^0$ decays was introduced. The details of the production of the neutral K_L beam and of the apparatus used to observe these decay products are given in this chapter.

2.1 The K_L Beam

E799-I was run in the Meson Center beamline at Fermilab during the 1991 fixed target run. In the fixed target operation, the Fermilab Tevatron produced a beam of 800 GeV protons. This beam was delivered to the fixed target experiments in a 22 second "spill" during each 58 second machine cycle, resulting in a beam delivery duty-cycle of 38%. In each spill,protons arrived grouped in radio frequency (RF) "buckets", approximately 2 ns wide and separated by 18.9 ns.

Although in principle the Tevatron should have produced uniformly populated buckets of protons, this was, in fact, not the case. During the run, we monitored this effect by observing coincidences between our accidental counter (described in Section 2.2.3.3) and the MU2 muon trigger plane (see Section 2.2.3.2). Since these two detectors viewed different angular regions from the target, it was extremely unlikely that the same particle could interact in both the accidental counter and MU2. Coincidences between these two counters then should occur randomly. Therefore, if the bucket population is uniform, the rate of

coincidences between these two counters in the same bucket should be the same as the rate of coincidences between, for example, the accidental counter being hit several buckets later than the MU2 trigger. The ratio of in-time to out-of-time coincidences,

$$R = \frac{\langle I(t)I(t)\rangle}{\langle I(t)I(t+6 \text{ buckets})\rangle},$$

where $I(t)$ is the bucket occupancy as a function of time, t, is a measure of the "Bucket Occupancy Autocorrelation" (B.O.A.),

$$\text{B.O.A.} = \frac{\langle I(t)I(t)\rangle}{(\langle I(t)\rangle)^2}.$$

R and the B.O.A. are only strictly equal if there is no correlation in the occupancy of a bucket with the occupancy in the bucket 6 RF cycles later. The B.O.A. provided a measure of the increase in average accidental activity that occurred because of the non-uniformity of the proton population in the buckets. Because accidental activity was a major source of loss of acceptance in the detector, it was important to keep the B.O.A. as low as possible during data-taking. Typically during E799, the measured R ranged from 2 to 3. An illustration of the effect of a non-trivial B.O.A. is shown in Figure 9. When it rose significantly above this range, the accelerator operators were able to make adjustments which reduced R.

During the spill, the protons extracted from the Tevatron were fanned out to many experiments, including E799. During the running of the experiment, the average primary beam intensity for good spills was 1.39×10^{12} protons per spill on target.

2.1.1 K_L Production

Kaons were produced from interactions of the protons in a one nuclear interaction length Be target, a square rod measuring approximately 3.2 mm on a side and 36.2 cm long mounted on a ceramic target stand. As the proton beam profile was approximately Gaussian with a 0.4 mm width, it was typically well contained within the target. The primary beam was monitored during the data-taking by a small wire chamber (SWIC) located approximately 3 meters upstream of the target, and final corrections to the beam position at the target were made using two steering dipole magnets located approximately 30 m upstream of the target. The targeting angle with respect to the secondary beam-line, 4.8 mrad, was chosen

as a compromise between the maximization of kaon flux, which occurs at small angles, and the minimization of the neutron to kaon ratio, which is lowest at large angles. Downstream of the target were several sweeping magnets which swept charged particles away from the secondary beam path, leaving neutral particles like K_L's, K_S's, neutrons, photons, Λ^0's and Ξ^0's.

The energy spectrum of kaons produced at the target was measured precisely by experiment E731, which ran in the Meson Center beamline in an earlier fixed target run with an identical production geometry and a similar detector. The spectrum was similar to what one predicted by relating the well-measured K^+ and K^- production spectra, parameterized by Malensek [54], to K^0 and $\overline{K^0}$ production based on quark content. Corrections on the order of 40% to this spectrum were determined from the K_L spectrum of the vacuum beam of E731 [55]. The energy spectra used in the Monte Carlo for K^0 and $\overline{K^0}$ production were

$$\frac{d\sigma_{K^0}}{dpd\theta} = \frac{sin\theta}{800}\left(B_1\frac{X(1-X)^{A_1}(1+5e^{-D_1X})}{1+q^2/{M_1}^2)^4} + B_2\frac{X(1-X)^{A_2}(1+5e^{-D_2X})}{(1+q^2/{M_2}^2)^4}\right)F(p)$$

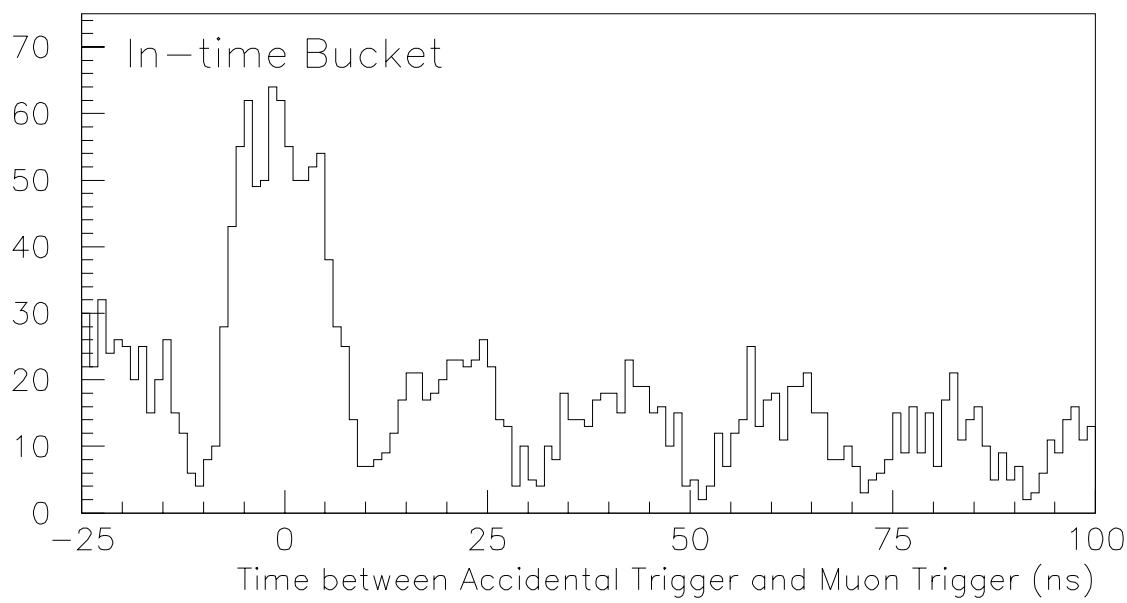

Figure 9. An illustration of the beam Bucket Occupancy Autocorrelation using accidental triggers. The MU2 trigger source TDC is plotted in these triggers; note the excess of MU2 triggers in the same bucket as the accidental trigger.

and

$$\frac{d\sigma_{\overline{K^0}}}{dpd\theta} = \frac{sin\theta}{400}\left(B_2\frac{X(1-X)^{A_2}(1+5e^{-D_2X})}{(1+q^2/{M_2}^2)^4}\right)F(p),$$

with

$$A_1 = 2.924 \qquad A_2 = 6.107$$

$$B_1 = 14.15 \qquad B_2 = 12.33$$

$$D_1 = 19.89 \qquad D_2 = 12.78$$

$$M_1 = 1.079(\mathrm{Gev/c^2}) \qquad M_2 = 1.048(\mathrm{Gev/c^2}),$$

where X is the ratio of kaon to proton energies, p is the kaon momentum, q is its transverse momentum, θ is the production angle, and with

$$F(p) = 1 + w_1p + w_2p^2 + w_3p^3 + w_4p^4,$$

$$w_1 = 6.033 \times 10^{-3}(\mathrm{GeV/c})^{-1} \qquad w_2 = -4.283 \times 10^{-6}(\mathrm{GeV/c})^{-2}$$

$$w_3 = -1.106 \times 10^{-7}(\mathrm{GeV/c})^{-3} \qquad w_4 = 1.802 \times 10^{-10}(\mathrm{GeV/c})^{-4}.$$

The resulting K_L production spectrum is shown in Figure 10.

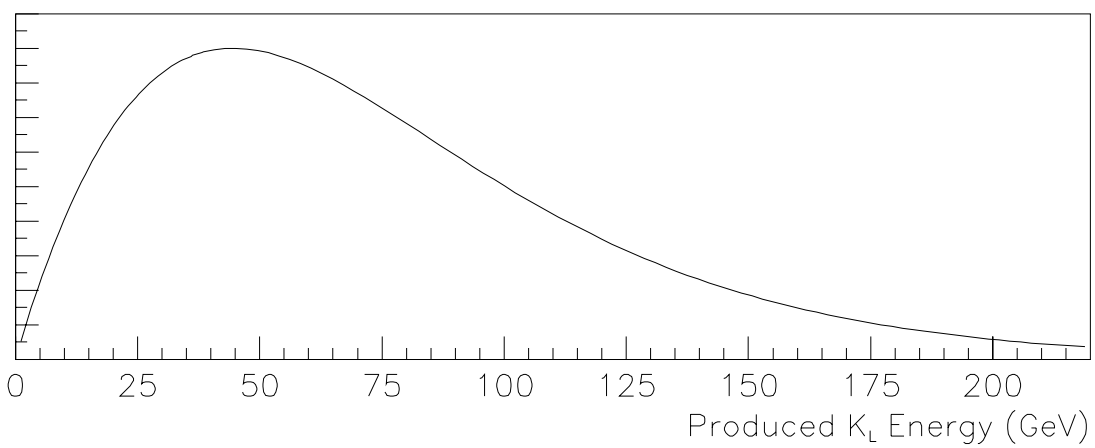

Figure 10. The energy spectrum of produced K_L's

2.1.2 Formation of The Neutral Beam

Secondary particles leaving the target were formed into a neutral beam by a series of collimators and sweeping magnets. The entire upstream beam-shaping apparatus is shown in Figure 11. Immediately after the target, were sweeping magnets to bend the path of charged particles produced in the target. Nine meters downstream of the target, a 5.8 m thick copper collimator with two tapered beam channels performed the initial formation of the secondary beams. Because of the sweeping magnets, the bulk of charged particles in the beam were not able to follow the secondary beam channel and were dumped into the collimator. Downstream of this initial collimator was a 3 inch thick (0.44 interaction lengths or 14 radiation lengths) absorber made of lead, which was designed to convert the photons in the beam so the resulting e^+e^- pairs could be removed by the downstream sweeping magnets. Following the absorber, there was a slab collimator, for defining the edges between the two beams and eliminating cross-over from one beam into another, and then a sweeping magnet to remove charged particles produced in interactions with the collimators or in the lead absorber. The next group of collimators was located 20 meters downstream of this magnet: a final slab collimator, two sets of jaw collimators, one oriented horizontally and one vertically, which defined the outer edges of the two beams, and a final sweeping magnet. Thirty meters downstream of these collimators were the last two sets of jaw collimators for the final outer edge definition of the neutral beam. The effective decay volume in E799 began just a few meters downstream of these last collimators. Figures 12 and 13 show the positions of each collimator along with the final beam edges in order to illustrate where each edge was defined.

Short-lived particles in the neutral beam, like the K_S ($c\tau \sim 2.7$ cm) and to a lesser extent Λ^0's and Ξ^0's ($c\tau \sim 7.9$ cm and 8.7 cm, respectively), decayed before reaching the K_L decay volume about 90 meters from the target. The very long-lived neutrons ($c\tau \sim 2.7 \times 10^{10}$ m) practically never decayed in flight. They passed through the detector, and were dumped into the Back Anti veto counter and the muon filter. Most of the K_L's ($c\tau \sim 15.5$ m) did so also, but a small fraction, about 3% at 70 GeV, did decay in the decay volume.

In the precursor experiments to E799, there was a 20 inch Be absorber (1.25 interaction

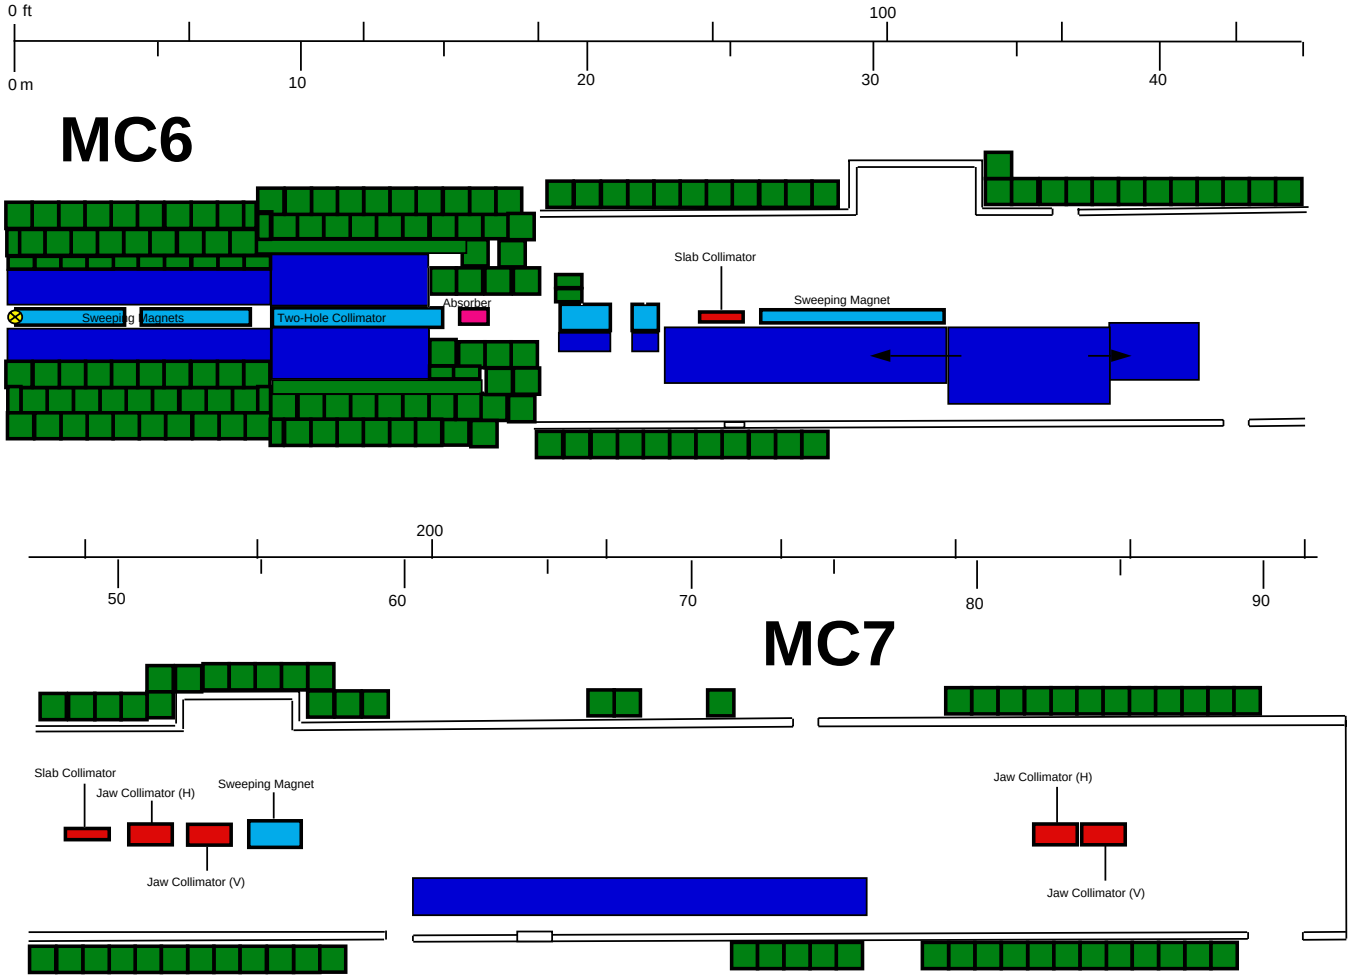

Figure 11. Diagram of the beam production and collimation portion of the experimental apparatus

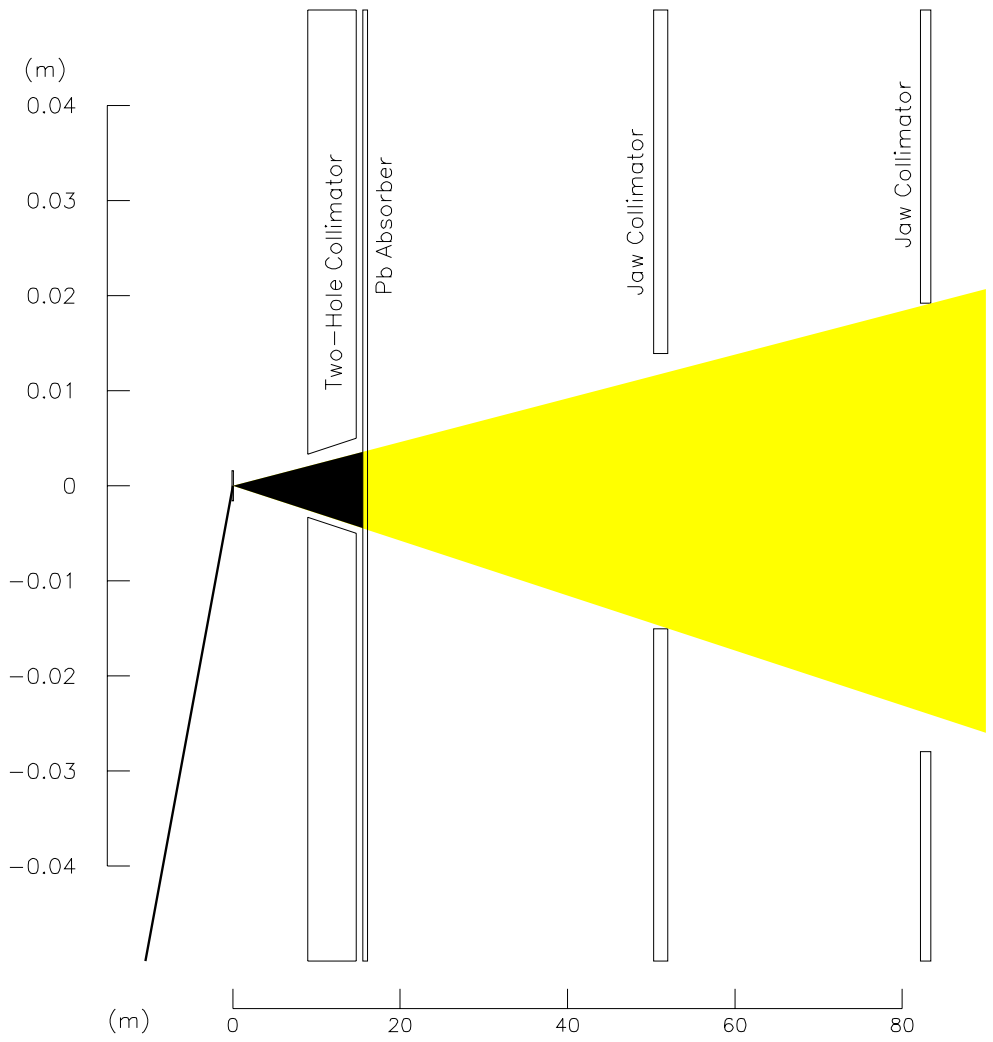

Figure 12. Formation of the beam by the collimation in horizontal ($\hat{x}$) view. The collimator positions shown are those used in the Monte Carlo simulation of the collimation (see Sections 5.5.5 and 7.1).

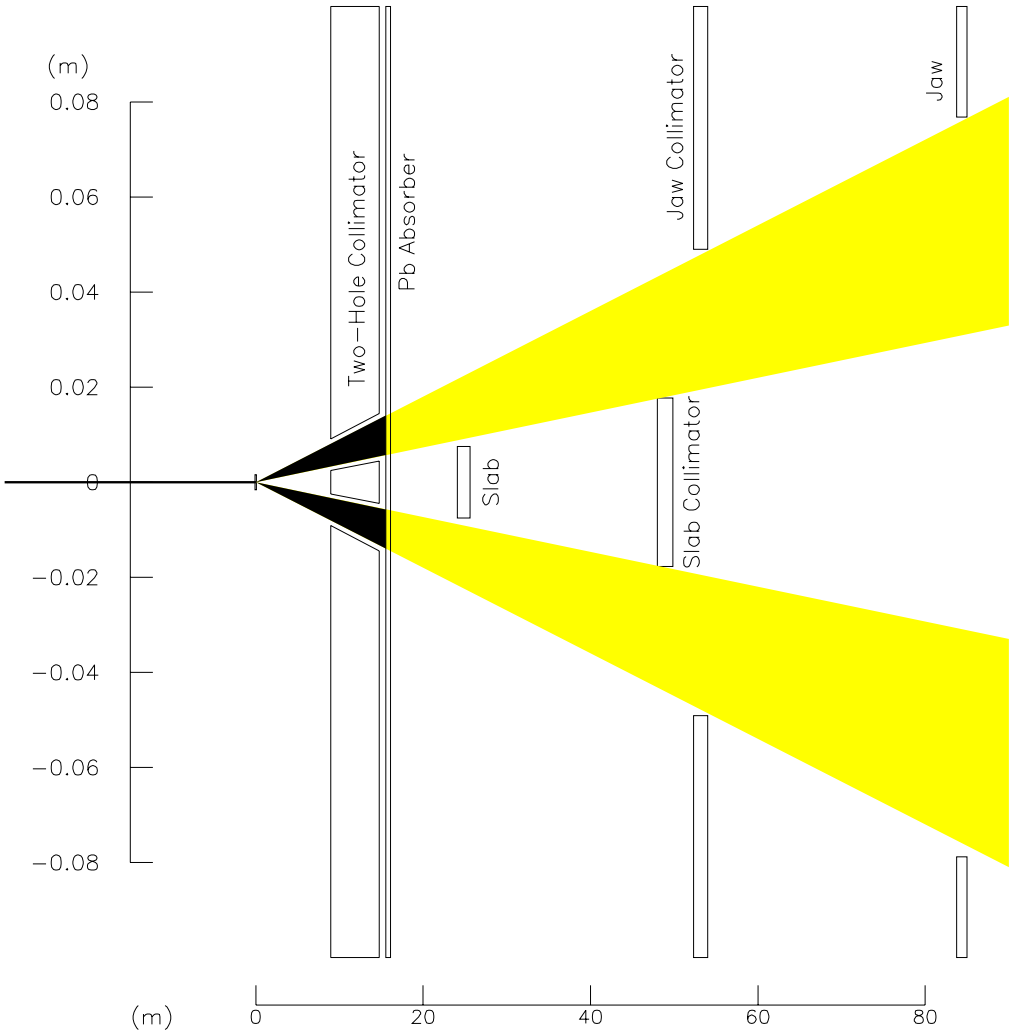

Figure 13. Formation of the beam by the collimation in vertical ($\hat{y}$) view

Table 3. Effects of Be absorbers in the beam on hadron and K_L flux

Inches of Be	Hadron Flux (relative units)	K_L Flux (relative units)	Average Calorimeter Energy Deposition (GeV)
20	1	1	0.83
8	2.1	1.9	n/a
none	3.6	3.0	1.44

The hadron flux was measured using rates in the last longitudinal layer of the Back Anti
The kaon flux was measured from rates of reconstructed $K_L \to \pi^0\pi^0\pi^0$ and
$K_L \to \pi^+\pi^-\pi^0$ decays

lengths) in the neutral beam located just upstream of the lead absorber. This absorber was in place to reduce preferentially the neutron content of the beam relative to the kaon content. However, because the primary beam intensity allocated to our beamline was limited, these absorbers had to be removed in order to obtain increases in the kaon flux. Before the final decision was made to remove these absorbers, the effect of these absorbers on kaon flux and total hadron flux was studied. The results are shown in Table 3. The total hadron flux was studied by looking at rates in the most downstream portion of the Back Anti photon veto counter (Section 2.2.3.1) into which the hadron beam was dumped, and the kaon flux was studied by looking at rates in triggers designed to find all-neutral and two-track kaon final states. These results indicate that the neutron to K_L ratio did not increase drastically as a result of removing the Be absorbers. The neutron and K_L beam rates could be calculated approximately from scaler rates in the Back Anti photon veto, which measured the total hadron flux into the beam dump, and from the total flux of kaons which was measured by reconstructing of kaon decays in the data. The average bucket occupancy of hadrons triggering the last layer of the Back Anti was 0.09; the average number of K_L which either decayed in flight or reached the Back Anti was approximately 0.05 per RF bucket.

2.1.3 The Decay Volume

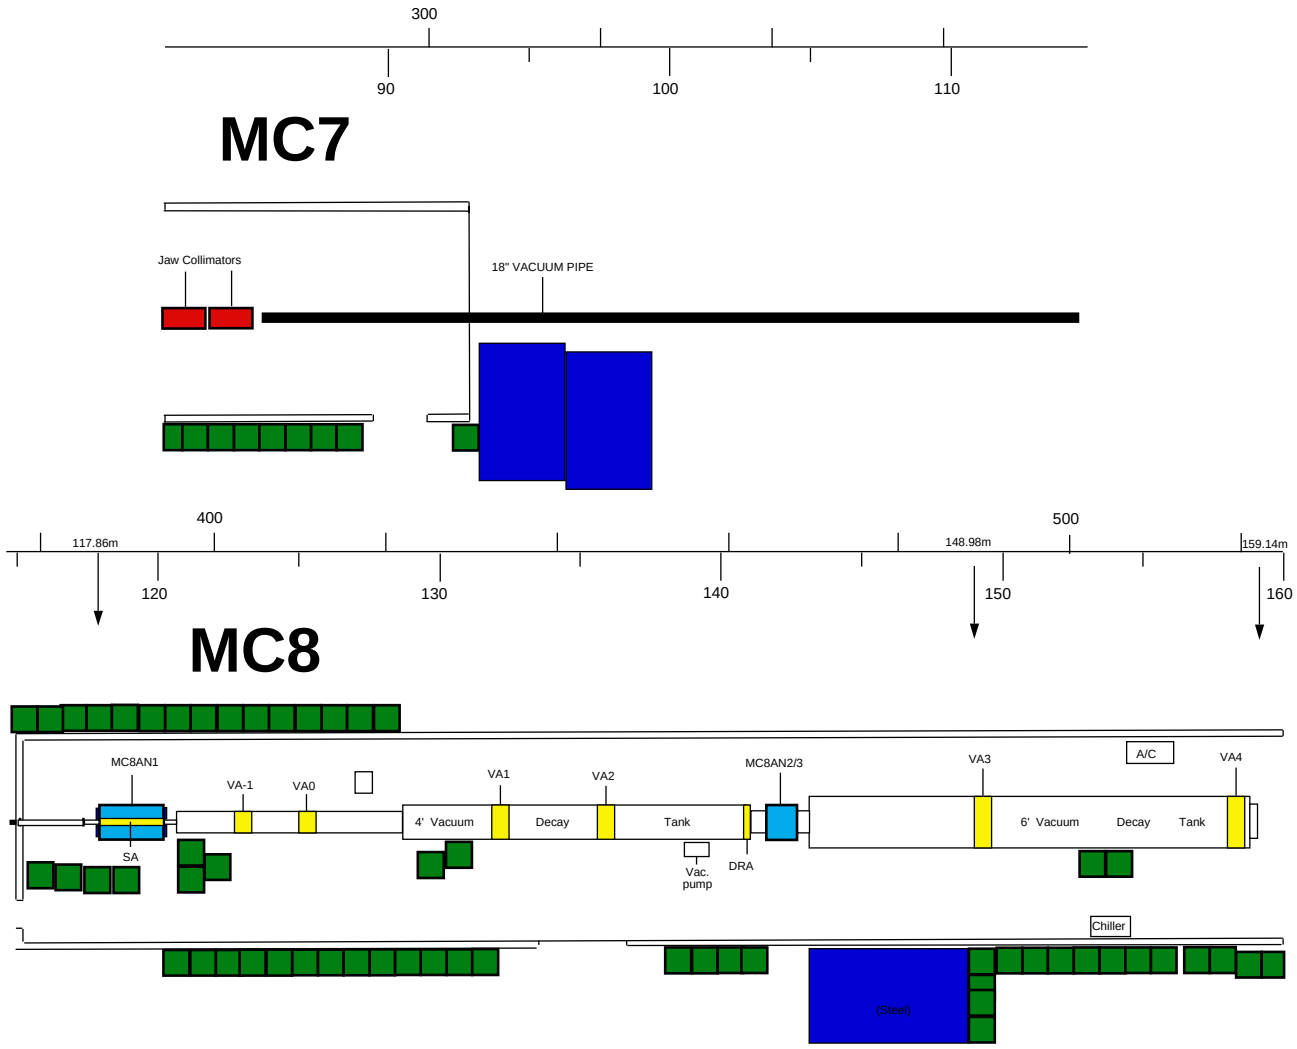

Figure 14. Diagram of the decay volume of the detector

The decay volume is shown in Figure 14. By the time the neutral kaons in the beam had reached the final collimators, 85 meters downstream of the production target, the K_S component had effectively decayed away. The fractional contamination of K_S's from the target after 85 meters was less than 10^{-3} for even the highest energy kaons (220 GeV) observed in our detector. The vacuum in the decay region was maintained below 15 mTorr throughout the run. An interaction length in the residual gas was greater than 3.7×10^6 meters, thus making regeneration of K_L's into K_S's in the vacuum system a negligible effect. The end of the decay volume was defined by a 48 inch diameter vacuum window just before the first drift chamber.

2.2 The Detector

The detector is shown in Figure 15. It consisted of a decay volume, surrounded by veto counters, a charged-particle spectrometer, a lead-glass calorimeter, and muon identification counters.

For future reference, it is useful to define a global coordinate system when discussing the elements of the detector. The origin of the detector coordinate system is the production target, and the line from the production target to the center of the lead-glass calorimeter is defined as the $\hat{z}$-axis, which is approximately parallel to the beam direction. "Upstream" and "downstream" will be used to denote the negative and positive $\hat{z}$-directions, respectively. The positive $\hat{y}$ direction is approximately "up". Strictly speaking, it was defined as the direction parallel to the vertical edges of the calorimeter blocks. Requiring a right-handed coordinate system, the positive $\hat{x}$ direction is then towards the left side of the calorimeter when viewed from upstream.

The individual elements of the detector are discussed below.

2.2.1 The Charged Spectrometer

The primary measurement apparatus for charged particles was a charged spectrometer which consisted of four drift chambers, two on each side of an analyzing magnet. The spectrometer allowed measurement of a charged particle's path through the detector and

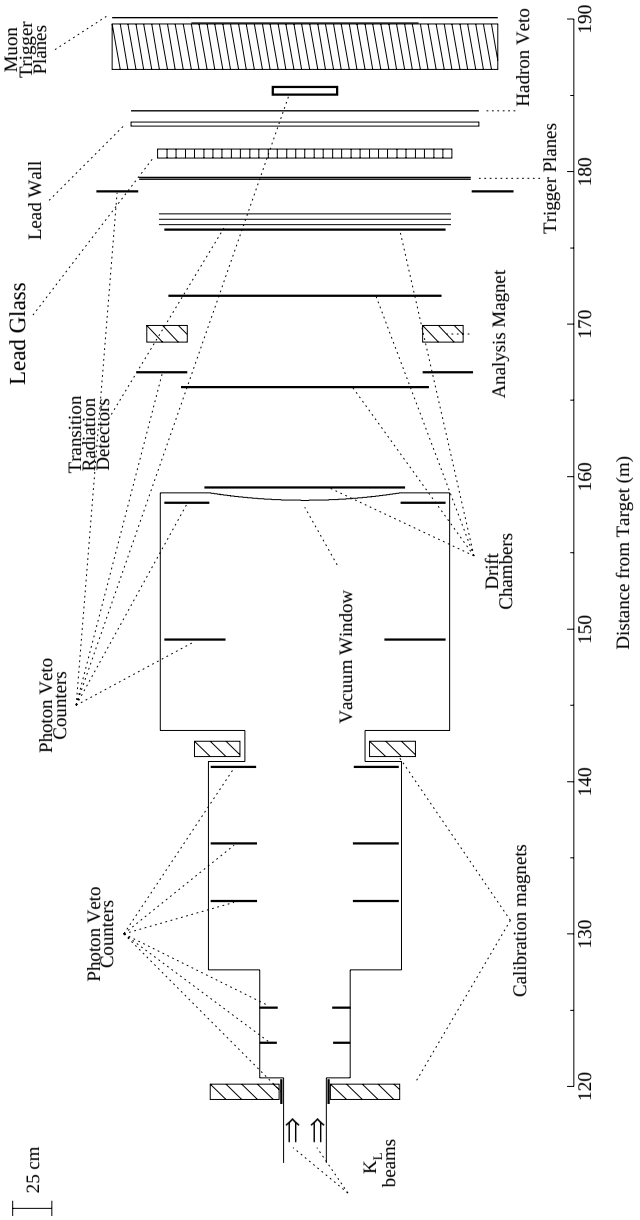

Figure 15. Schematic drawing of the detector. The scale in the direction transverse to the beams has been expanded for clarity.

a measurement of the momentum from the bend angle at the analyzing magnet. With multiple charged particles observed in the final state, the location of the decay vertex could be determined by extrapolating tracks upstream to a common origin. Thus information from the spectrometer played many important roles in the analysis.

The drift chambers consisted of a gas filled volume, bounded by two wire-supported Aclar windows and strung with field-shaping and sense wires. When charged particles passed through the gas, they produced a trail of electron-ion pairs with a mean separation of 300 μm. The high electric field between the field-shaping and sense wires caused the electrons to drift towards the sense wires, while the ions drifted towards the field wires. Because of the competing effects of acceleration of the electrons from the field and the collisions of the electrons with the gas, the electrons drifted with a velocity that was roughly

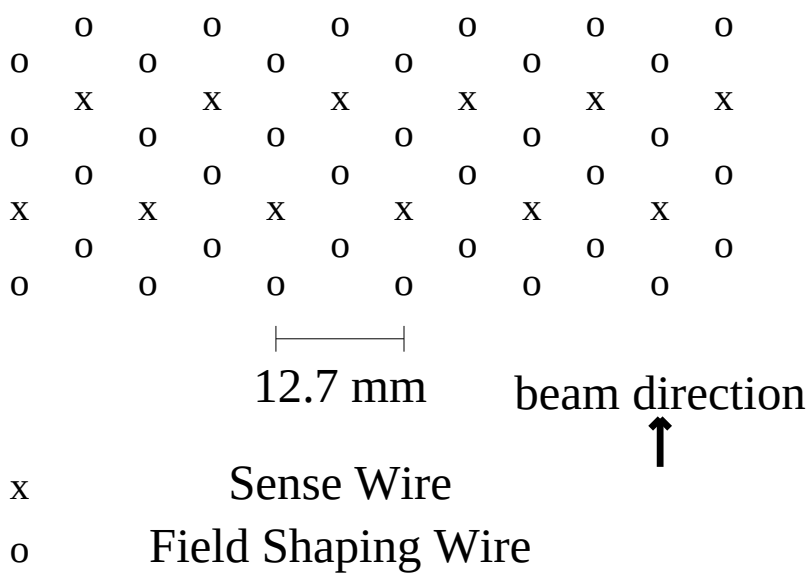

Figure 16. The sense and field wire geometry in half of a drift chamber (one view) is shown. Each sense wire was surrounded by a hexagonal group of six field-shaping wires.

constant over much of the active chamber volume. The field-shaping and window support wires were 100 μm thick gold-plated Cu-Be, and the sense wires were 25 μm gold-plated Tungsten. When the electrons came very close to the thin sense wires, the very high electric fields accelerated electrons in between collisions to sufficient energies that they could create more electron-ion pairs. This avalanche effect resulted in an amplification of the electron signal reaching the sense wire.

Each of the four drift chambers consisted of four planes of sense wires, two each in the $\hat{x}$ and $\hat{y}$ views. The sense wire spacing was 12.7 mm in each plane, and the two planes in the same view were offset by half of a wire spacing. This allows a dual measurement in each chamber-view to resolve the left-right ambiguity that results from ionization drifting towards the nearest wire, rather than in a specific direction. These two planes of sense wires in the same chamber-view are often referred to as "complimentary planes". The field-shaping geometry was hexagonal (see Figure 16), with equal spacing between every plane of wires in the beam-direction. The four chambers were of different sizes (see Table 4), with the largest downstream, to increase acceptance for the high angle tracks. A high voltage of -2650 Volts was maintained on the field wires during data-taking. To reduce dark currents and potential aging effects in the chambers, this voltage was reduced by 20% in between spills.

The drift chamber gas used was a mixture of 50% Argon, 50% Ethane. With the operating voltage of -2650 Volts, the drift velocity was about 50 μm/ns. The drift velocity was almost independent of the local field in the range of fields across the drift region of the cell, varying by less than 15% in fields from 50 V/mm to 500 V/mm. A small admixture (about 1%) of ethyl alcohol was added to the chamber gas for its quenching properties in order to reduce aging in the chambers. At one point during the run, several gas cylinders contaminated with Hydrogen Sulfide were delivered to our beamline and this gas was flowed through our chambers. Because H_2S acts to quench ionization, the efficiencies of our drift chambers decreased substantially, an effect that was apparent from online monitoring. Mixtures containing some H_2S were flowed through the chambers over a 13 hour period. After the bad gas was discovered, these cylinders were removed and we resumed running with no apparent permanent damage to the drift chambers.

Table 4. Drift chamber sizes, measured in wires per sense plane

Chamber Number	View	Number of Sense Wires Per Plane
1	X	101
1	Y	101
2	X	128
2	Y	112
3	X	136
3	Y	128
4	X	140
4	Y	140

The signal from each sense wire was discriminated at the chambers and sent to a repeater card from which copies of the signal were sent to both the chamber TDCs (Time-to-Digital Converters) and the track processor front-end boards. The chamber TDCs used were commercially available LeCroy 4291B TDCs with a common stop provided by the Level 1 trigger signal. Once the TDC received a signal, it was dead for approximately 250 ns, the double-pulse resolution time of the TDCs. This long resolution time had the advantage of ensuring that secondary pulses in the analog signal shape from the sense wire, such as reflections on the chamber wire, would not cause the TDCs to restart.

By converting the drift time measured at the TDCs to a drift distance, a position measurement can be made from each chamber wire hit with an accuracy of about 100 μm. The details of how the chambers were aligned and calibrated, as well as the chamber performance, are discussed in Chapter 5.

The 100D40 analysis magnet was a dipole magnet with a vertical field of approximately 4 KGauss. The transverse dimensions of the magnet aperture measured 2.53 m $\times$ 1.46 m, and the pole faces extended 3.37 m in the $\hat{z}$-direction. The current to the magnet was monitored and written to tape each spill. Past experience [56][57] showed that problems with the magnet power supplies could be correlated with changes in the magnet field strengths; however, during E799, the power supply system was stable and required no repair. The determination of the field-strength and mapping of the magnet is described in Section 5.6. The magnetic field imparted a momentum kick, approximately 200 MeV, in the $\hat{x}$-direction. The magnet's field was negligible at the second and third drift chambers, and was very small in the horizontal direction; therefore, to a good approximation, the entire effect of the magnetic field was to simply bend the straight x-view tracks between chambers 2 and 3. This bend angle was used to measure the momentum of charged particle. The locations along the beam axis of the spectrometer elements are given in Table 5. Because of the addition of Transition Radiation Detectors (TRDs) for the E799 run, the distance between Chamber 3 and Chamber 4 is approximately 2 meters shorter than that between Chambers 1 and 2. This had the effect of degrading the momentum resolution somewhat because of the reduced lever arm in measuring downstream tracks. However, the resolution was still dominated by multiple scattering over most of the relevant momentum range.

Table 5. Distance from the target to elements of the charged spectrometer

Spectrometer Element	Distance from Target to Element Midpoint (m)
(Vacuum Window)	158.946
Chamber 1	159.287
Chamber 2	165.861
Magnet Bend Plane	168.864
Chamber 3	171.856
Chamber 4	176.200
(Lead-Glass Face)	180.886

2.2.1.1 Material In the Spectrometer

For Monte Carlo simulation of multiple scattering, bremsstrahlung and photon conversions, it was important to understand the amount of material through which decay products pass. In particular, material upstream of the spectrometer contributed significantly to the invariant mass resolution, and the multiple-scattering in the material internal to the spectrometer largely determined the momentum resolution. Therefore, a good deal of attention was paid to reducing and then measuring the material contributed by each detector element. For example, to reduce multiple scattering in the air between drift chambers, helium-filled bags were placed between drift chambers, and then the sizes of the remaining air gaps and bag material were carefully measured. The total material, in radiation lengths, for every element in the beam between the end of the decay volume and the lead-glass is listed in Tables 6 and 7. Special cases where calculation of the radiation lengths is not straightforward are discussed below.

The vacuum window at the end of the decay volume was the most significant single piece of material upstream of or inside the spectrometer. It consisted of 23 mils of a Kevlar-29 woven fiber fabric (in a "basket weave" according to the manufacturer) and 5 mils of Mylar. The radiation length of Kevlar-29 ($C_{14}H_{10}N_2O_2$) was calculated to be 41.29 g/cm^2 by mass weighting of the elemental radiation lengths, and the Kevlar fabric weave used had a sheet

Table 6. Radiation lengths of material upstream of or inside the spectrometer

Detector Element	Material	Thickness(cm)	Radiation Lengths
Vacuum Window	Kevlar-29 Fibers	0.058	0.00114
	Mylar	0.013	0.00045
Air Gap	Air	39.1	0.00129
Chamber 1	Windows, Gas	8.2	0.00084
	Wires (average)	n/a	0.00060
Air Gap	Air	35.6	0.00117
Helium Box	Mylar Windows (2)	0.005 ea.	0.00019 ea.
	Polyethylene Bag (2)	0.014 ea.	0.00029 ea.
	He	587	0.00103
Air Gap	Air	26.4	0.00087
Chamber 2	Windows, Gas	8.2	0.00084
	Wires (average)	n/a	0.00060
Air Gap	Air	24.1	0.00079
Helium Bag	Polyethylene Bag (2)	0.014 ea.	0.00029 ea.
	He	537	0.00095
Air Gap	Air	30.0	0.00099
Chamber 3	Windows, Gas	8.2	0.00084
	Wires (average)	n/a	0.00060
Air Gap	Air	21.8	0.00071
Helium Bag	Polyethylene Bag (2)	0.014 ea.	0.00029 ea.
	He	379	0.00067
Air Gap	Air	25.4	0.00083
Chamber 4	Windows, Gas	8.2	0.00084
	Wires (average)	n/a	0.00060

Table 7. Radiation lengths of material downstream of the spectrometer

Detector Element	Material	Thickness(cm)	Radiation Lengths
Air Gap	Air	80	0.0026
TRD 1	Windows, Gas, Wires, Radiators	13	0.012
Air Gap	Air	30	0.0010
TRD 2	Windows, Gas, Wires, Radiators	13	0.012
Air Gap	Air	30	0.0010
TRD 3	Windows, Gas, Wires, Radiators	13	0.012
Air Gap	Air	150	0.0049
C-Bank	Scintillator	1.0	0.024
B-Bank	Scintillator	1.0	0.024
Air Gap	Air	120	0.0039

density of 0.0469 g/cm^2. Therefore, the entire vacuum window represented $1.6 \times 10^{-3} X_0$ of material.

A second significant contribution to the material inside the spectrometer was the presence of air gaps in between the vacuum window and the chambers and in between the chambers and the helium bags. Since the radiation length of air is relatively short, only 304 meters, a small amount of air could contribute significantly to the total radiation length. The gaps between chambers and helium bags were measured at five points: the center of the chamber and at the four corners. Each helium bag face has a roughly spherical bulge, but a single "average" gap length was actually used in the Monte Carlo. This average was defined as 3/4 of the center gap plus 1/4 times the average of the four corner gaps. For the distance between the vacuum window and the first chamber, 3/4 of the maximum bulge (10.9 cm at 1 atmosphere pressure difference) was used to obtain an "average" gap length. The air gaps

ranged from 20 to 50 cm and accounted for a large fraction of the material internal to the spectrometer. It is also worth noting that the Helium bag radiation lengths were calculated assuming pure Helium in the bag without considering any possible admixture of air.

Finally, the chambers themselves contained significant material. Each chamber had two outer windows, consisting of 1 mil Aclar, $0.00019X_0$ each, and 82 mm of Ar-Ethane, $0.00046X_0$ total. Each chamber also has window wires, field wires and sense wires. In the Monte Carlo simulations of bremsstrahlung and multiple-scattering, these wires were treated as discrete points of material, with a certain interaction probability depending on their thickness. For photon conversion however, the wires were treated as a single thin sheet of material in the chambers. In the approximation of the wires as thin sheets of material, the wires contributed $0.00060X_0$ to the total radiation length.

Given these radiation lengths, one can calculate the contribution to momentum resolution from multiple scattering. The primary effect of multiple scattering is to change the bend angle measured at the magnet. For the purposes of this approximation, only material located between the midpoint between Chambers 1 and 2 and the midpoint between Chambers 3 and 4 was considered to contribute to the change in the bend angle. (Material between Chambers 2 and 3 will contribute fully to this effect; material between the other chambers will contribute less because some of the change in angle is measured by the drift chambers.) The total considered material was 0.94% of a radiation length. The RMS multiple scattering angle from a thin sheet of material is given by [3],

$$\theta_{ms} = \frac{0.0136}{p}\sqrt{t}\,[1 + 0.038\log_e t]\,,$$

where p is the track momentum in GeV/c, and t is the material thickness is radiation lengths. The RMS "momentum kick" from multiple scattering is then 1.1 MeV in this approximation. This then contributes a momentum-independent term in the momentum resolution of 0.55%. This is to be compared with the error from the 100 μm position resolution of the chambers of $0.011p\%$, where p is the momentum in GeV/c. For many-body final-state decays such as those investigated in this analysis, $p < 50$ GeV typically, so the multiple-scattering term will dominate the momentum resolution.

2.2.2 The Electromagnetic Calorimeter

The electromagnetic calorimeter consisted of 804 blocks of Schott F-2 lead-glass, $5.8 \times 5.8 \times 60.2$ cm^3 each. The blocks were oriented with the long axis parallel to the beam direction and stacked in a roughly circular array with two holes in the center through which the hadron beams could pass (see Figure 17). The radiation length of Schott F-2 lead glass is 3.21 cm, so each block was approximately 18.8 radiation lengths in depth.

The Čerenkov light produced in electromagnetic showers in the lead-glass was observed using phototubes mounted on the back of each block. The phototubes used were Amperex 2202 ten-stage tubes with bialkali cathodes which were operated at approximately 1200 V, providing for gains on the order of 10^5. The tubes made optical contact with the blocks using a silicon gel whose index of refraction (1.45) was in between that of the lead-glass (1.6) and the phototube window (1.4). Embedded in the gel was a Wratten 2A filter which cut off the part of the Čerenkov spectrum where the absorption in the lead-glass was large (below 400 nm). As will be discussed in Chapter 6, this absorption introduced a significant non-linearity to the glass response and also degraded the resolution. This filter then played an important role by minimizing this absorption. Each block was individually wrapped in 0.5 mil aluminized mylar to increase light collection and optically isolate each block from its neighbors.

In order to monitor phototube gains, each lead-glass block viewed light from a xenon flash lamp (the "flasher") with a quartz window through a system of optical fibers, one of which was mounted on the front face of each block. The lamp provided light in a spectrum similar to that of Čerenkov radiation with each tube viewing light approximately equivalent to the amount of light created in a 40 GeV shower. By monitoring each block during these flashes, changes in tube gain could be monitored and then corrected by changing the high voltage on the tube. Because previous experiments had found that the phototube gains changed by as much as 10% as the incident particle flux increased at the beginning of each spill, these same fibers were used to transmit a constant low level of LED-produced light to the face of each block. This low level of ambient light was sufficient to eliminate this effect.

The entire array was enclosed in a special "glass house". The glass house was kept

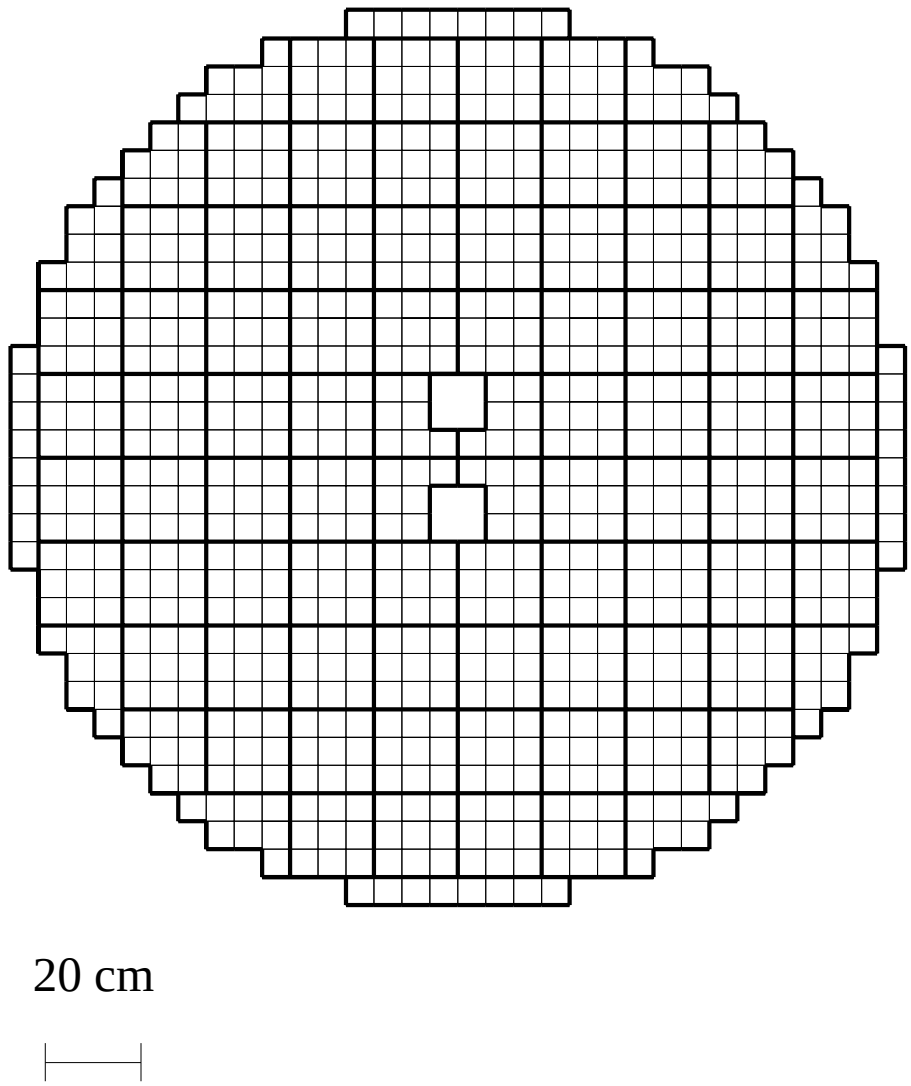

Figure 17. The lead-glass array. Outlined in the heavy lines are the boundaries of the lead-glass adders.

light tight, in order to avoid damage to the phototubes or noise from ambient light, and climate controlled by a system of air conditioners, fans and a household thermostat. During E799, this climate control system was able to keep the temperature at the phototube bases constant to within 1.5°C, which effectively limited gain fluctuations due to temperature variations.

The signal from each phototube was carried by an 80 meter RG-58 cable (which provided sufficient delay for formation of the Level 1 trigger) to a FASTBUS Le Croy 1885 ADC (Analog-to-Digital Converter) module. Although Čerenkov light itself was produced promptly, the signal at the ADC has a substantial time-spread due to broadening in the cables and because of some scintillation in the lead-glass. The ADC integrated the phototube signal over a 100 ns gate, and then converted the integrated pulse height to a digital readout using a 12-bit ADC in the module which was operated in two different ranges. The first of these ranges, "low range", was used for smaller pulse heights; for pulse heights corresponding to more than approximately 3700 ADC counts in low range, a second "high range" was used instead with an input signal attenuation of a factor of 8 before digitization. This effectively extended the dynamic range of the ADC module to that of a 15 bit ADC. With our gate and phototube gains, each low-range ADC count corresponded to between 6 and 7 MeV for most lead-glass blocks. However, several of the blocks closest to the hadron beam-holes suffered severe radiation damage during the run, and for these blocks the low-range gains were as high as 15 MeV/count.

A small portion of the lead-glass analog signal was diverted for use in fast lead-glass energy sums. Using electronics located just outside the glass house, the analog signals from the blocks were added together into 92 groups of approximately nine blocks each (see Figure 17). These signals were then sent on cables similar to the lead-glass block signal cables to ADC modules which used a 30 nsec gate. This narrower gate allowed for a comparison of the adder energies with the energies in the lead-glass blocks to identify events with out-of-time energy in the lead-glass. A small part of each of the adder signals was tapped to form analog energy sums for the lead-glass quadrants, and the quadrant signals were added to form a total lead-glass energy sum, E_t. Three copies of the E_t signal were also sent to the same ADC module as the adder signals with three different cable lengths in

order to form a digitized E_t-early, E_t-late and E_t-in-time quantities. The analog quadrant sums and E_t were also discriminated near their point of formation to make trigger signals which were sent to the Level 1 trigger electronics.

During the experiment, a significant radiation dosage was absorbed by the blocks around the pipes through which the hadron beams passed (the "pipe blocks"). Figure 18 shows estimates of the dosage received per block for various parts of the lead-glass based on radiation monitors located immediately behind blocks of the lead-glass array during two weeks of typical running. The effect of this radiation damage on the Lead-Glass was to

Block Location	Estimated Dosage During the Run
A	23 krad
B	5.6 krad

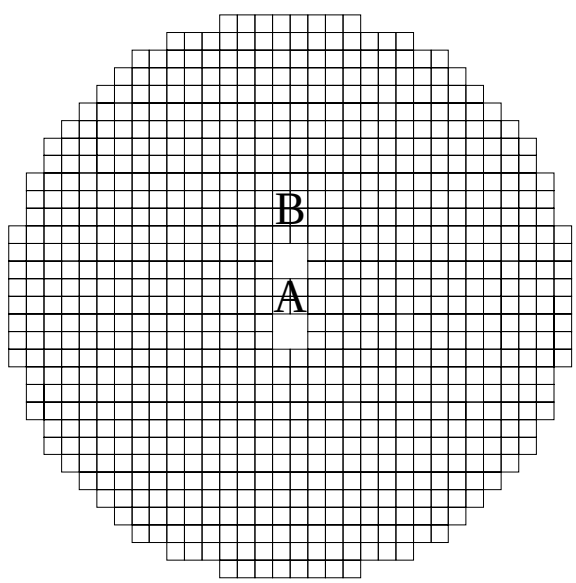

Figure 18. The radiation dosage per block in the lead-glass calorimeter is shown as estimated using radiation monitors during a typical two-week running period.

increase the absorption of Čerenkov light produced in electromagnetic showers. To attempt to "cure" this damage in the pipe blocks, ultraviolet light from two 400 Watt mercury vapor lamps was shone upon the front of the central blocks of the Lead-Glass array for about twelve hours after approximately every two weeks of beam-time. During a one week accelerator downtime in the middle of the run, the UV lamps were used for a three-day "mega-cure", in order to correct the rapidly increasing absorption. Past performance in lower intensity experiments using the Lead-Glass array, such as E731 [55] and E773 [58], showed that this curing procedure did reverse approximately 90% of the damage to the lead-glass. In E799, the rate of damage in the pipe blocks was so large that the damage steadily worsened from calibration to calibration despite any temporary improvements from curing.

2.2.3 Other Detector Elements

Aside from the spectrometer and calorimeter, there exist a number of other detector elements which fall into two major categories. There are several photon vetoes which are designed to detect photons or charged particles escaping the fiducial volume of the detector and a number of scintillator trigger planes used for forming fast triggers and for particle identification. These two categories of detectors are discussed below.

There were also several test detectors installed in E799 which did not play a role in the event reconstruction. At the beginning of the run, two beamhole calorimeters, one made of PbF_2 and one of BaF_2 were placed behind the lead-glass beamholes in an attempt to increase acceptance. These were not used during the bulk of the run, however, because they caused neutral hadrons from the beam to interact in trigger counters downstream of the lead-glass and also because the BaF_2 calorimeter was severely radiation-damaged within a very short time [1]. There were also three transition radiation detectors (TRDs) installed for E799 which were designed to lessen the probability of misidentifying charged pions as electrons. These, unfortunately, were not operational for much of the run, thus limiting their effectiveness in analysis.

2.2.3.1 Photon Vetoes

The photon vetoes come in many varieties, but all have the basic goal of detecting charged particles and photons which escape the fiducial volume of the calorimeter or the spectrometer. Several—the Vacuum Antis (VA's), numbered -1 through 4, and the Decay Region Anti (DRA)—were located inside the evacuated decay volume, bolted to the vacuum pipe wall. The Magnet Anti (MA) shadowed the magnet aperture. Two—the Lead Glass Anti (LGA) and the Iron Ring—surrounded the outer edge of the lead-glass array, and the Back Anti (BA) was positioned to detect photons escaping down beamholes in the center of the lead-glass. The diverse positioning of the vetoes also implied a geometric diversity as shown in Figures 19, 20 and 21.

There was also a diversity of composition in the different vetoes. All of the Vacuum Antis, the Magnet Anti and the Lead Glass Anti consisted of a scintillator layer for detection of charged particles, followed by two modules of three radiation-length lead-lucite sandwiches for detection of photons (see Figure 22). The Decay Region Anti was two planes of scintillator counters with a lead sheet between them in order to convert photons. The Iron Ring (shown in Figure 23) consisted of a 2.9 radiation length thick ring of rusty iron hung about 30 cm in front of the B and C scintillator banks, designed to convert photons that would otherwise have missed both the lead-glass and the LGA so the resulting shower could be observed in the scintillators. Finally the Back Anti counter was a segmented 48 layer lead-lucite sandwich, 28.1 radiation lengths deep. UV transmitting light pipes combined the light from eight different layers in the same view together and delivered it to the 72 readout phototubes. The segmentation divided the counter into three regions of depth, and into 1 inch strips in alternating horizontal and vertical lucite planes. Because of the coarse longitudinal segmentation, the BA could distinguish electromagnetic from hadronic showers by checking for significant energy deposition in the most downstream portion of the counter which should not be present in electromagnetic showers.

The total coverage for escaping photons given by the Photon Veto system is illustrated in Figures 24 and 25.

Photon Veto	x_i (m)	y_i (m)	x_o (m)	y_o (m)	Distance from Target (m)
VA -1 *	0.0947	0.1782	0.2050	0.2863	122.873
VA 0 *	0.0947	0.1721	0.2050	0.2889	125.175
DRA	0.249	0.310	radius 0.605 m		140.948
MA	0.908	0.743	1.057	1.067	166.836

*Had a 3.5° tilt in the $\hat{x}$-$\hat{y}$ plane with respect to the axes

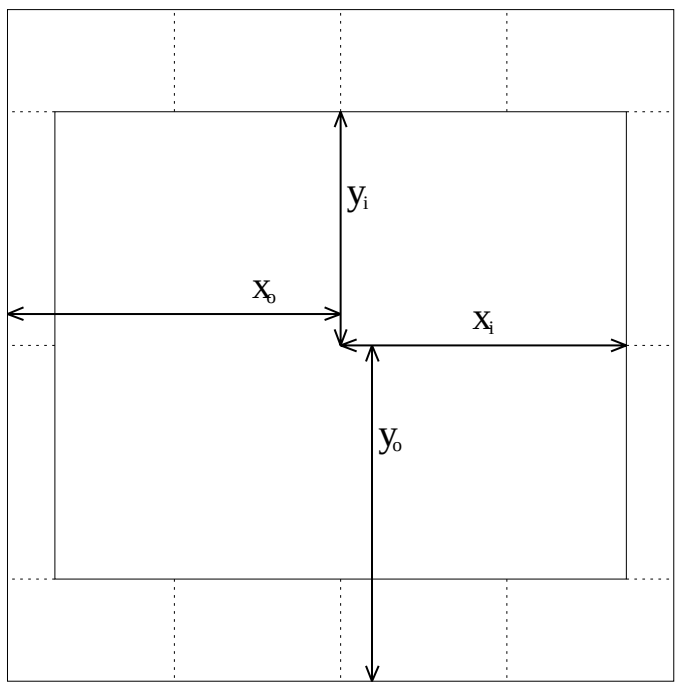

Figure 19. The geometry of the "square" photon vetoes. The segmentation of counters varies. The MA segmentation is shown.

Photon Veto	r_i (m)	r_o (m)	Distance from Target (m)
VA 1	0.303	0.595	132.172
VA 2	0.303	0.595	135.933
VA 3	0.502	0.885	149.291
VA 4	0.606	0.889	158.273
LGA	0.908	1.403	178.710
Iron Ring	0.848	0.931	179.266

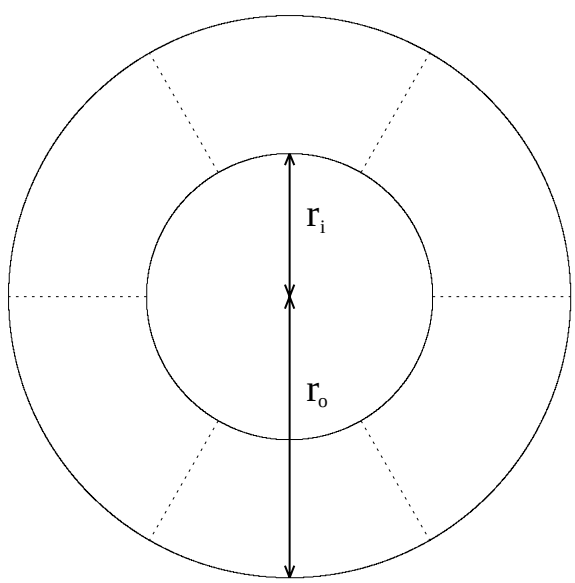

Figure 20. The geometry of the "round" photon vetoes. The segmentation of counters varies. The VA0 segmentation is shown.

Photon Veto	x (m)	y (m)	Distance from Target (m)
BA	.298	.502	185.047

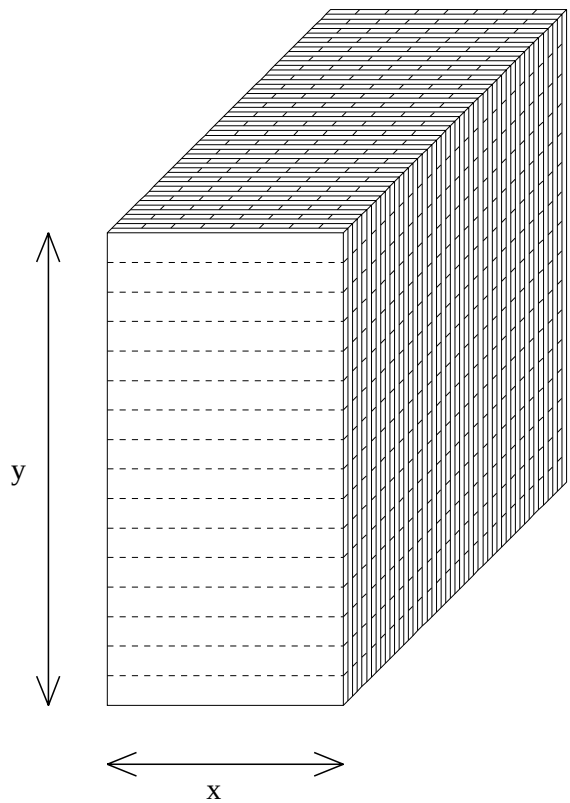

Figure 21. Schematic of the Back Anti counter

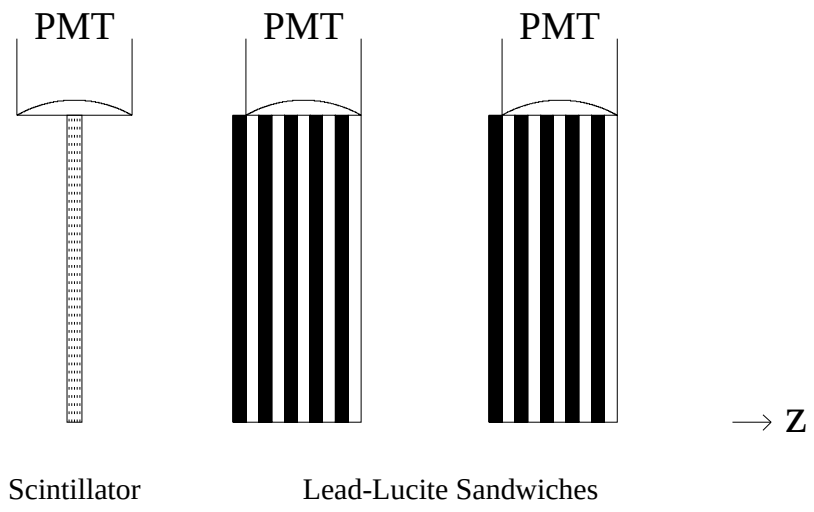

Figure 22. Cross-Section of a lead-scintillator photon veto counter

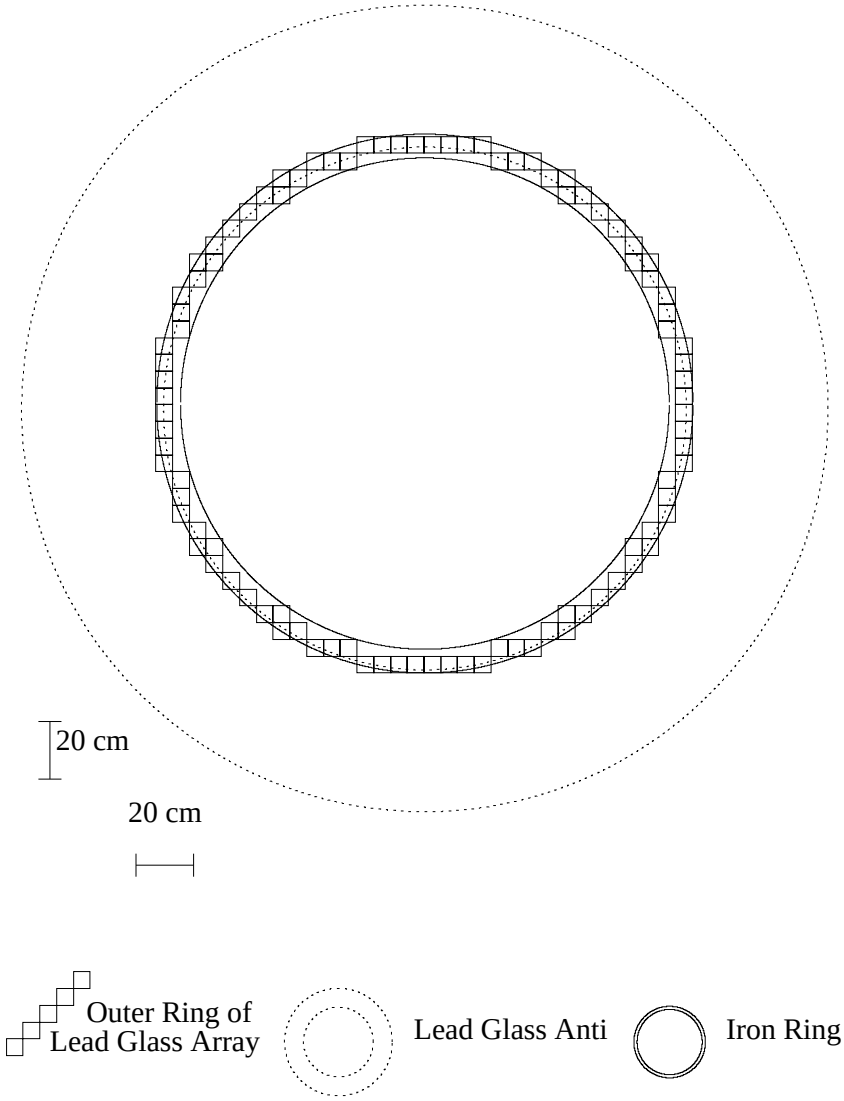

Figure 23. The location of the iron ring is shown with the positions of the outer lead-glass blocks and the LGA superimposed. The iron ring was designed to convert high-angle photons in the gap between the lead-glass and the LGA.

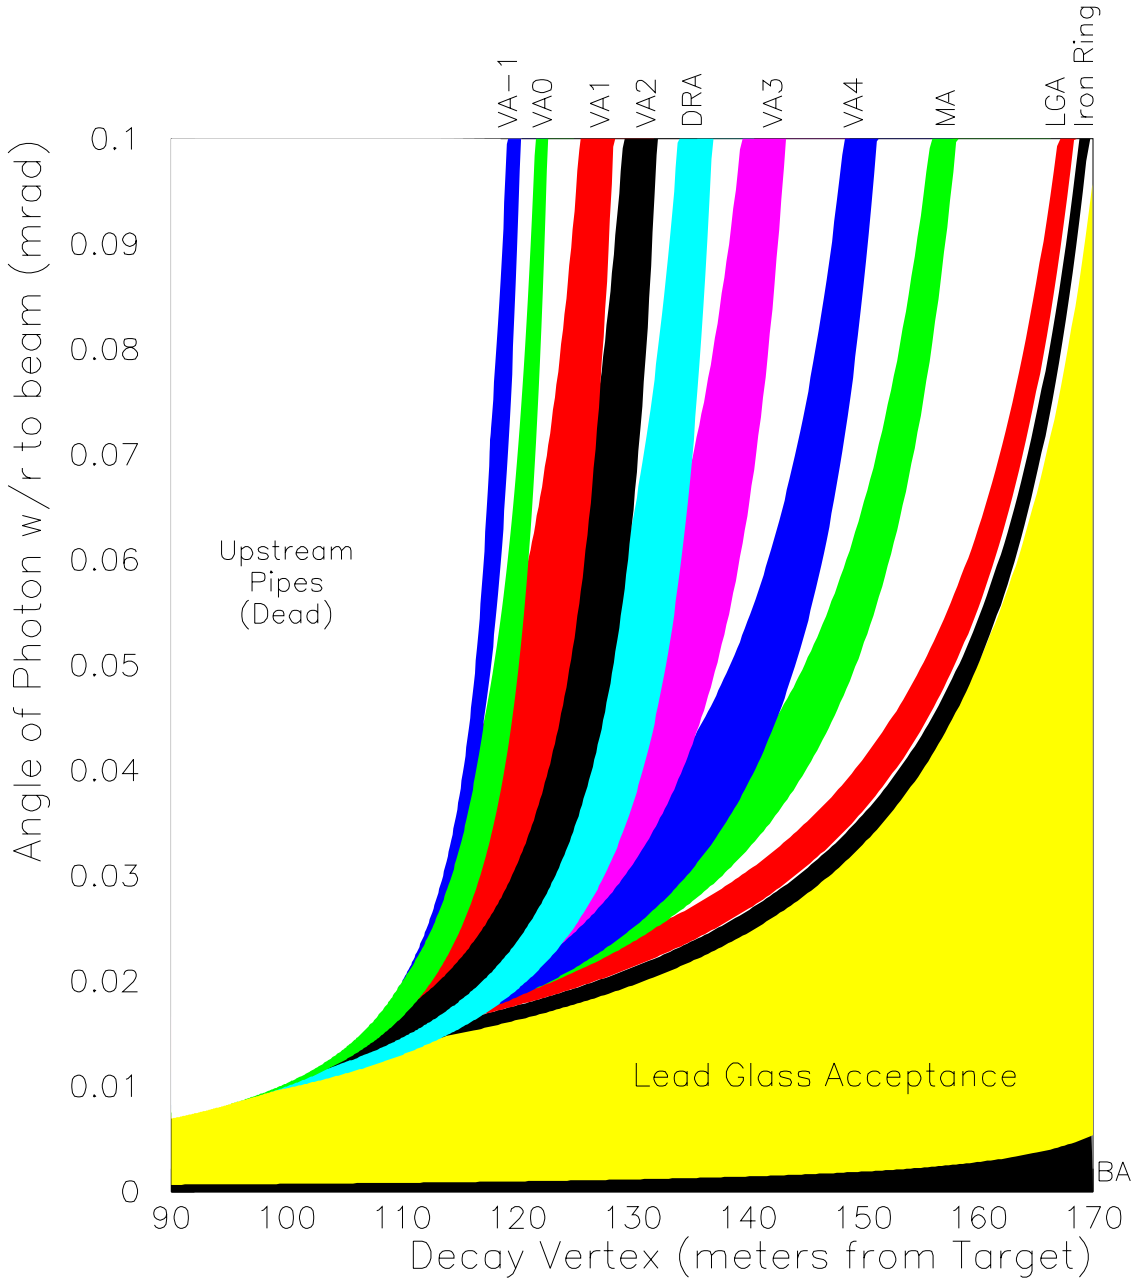

Figure 24. The coverage of the photon vetoes as a function of decay vertex and photon angle is shown for photons created at the center of the *upper* beam, traveling in the $-\hat{y}$ direction.

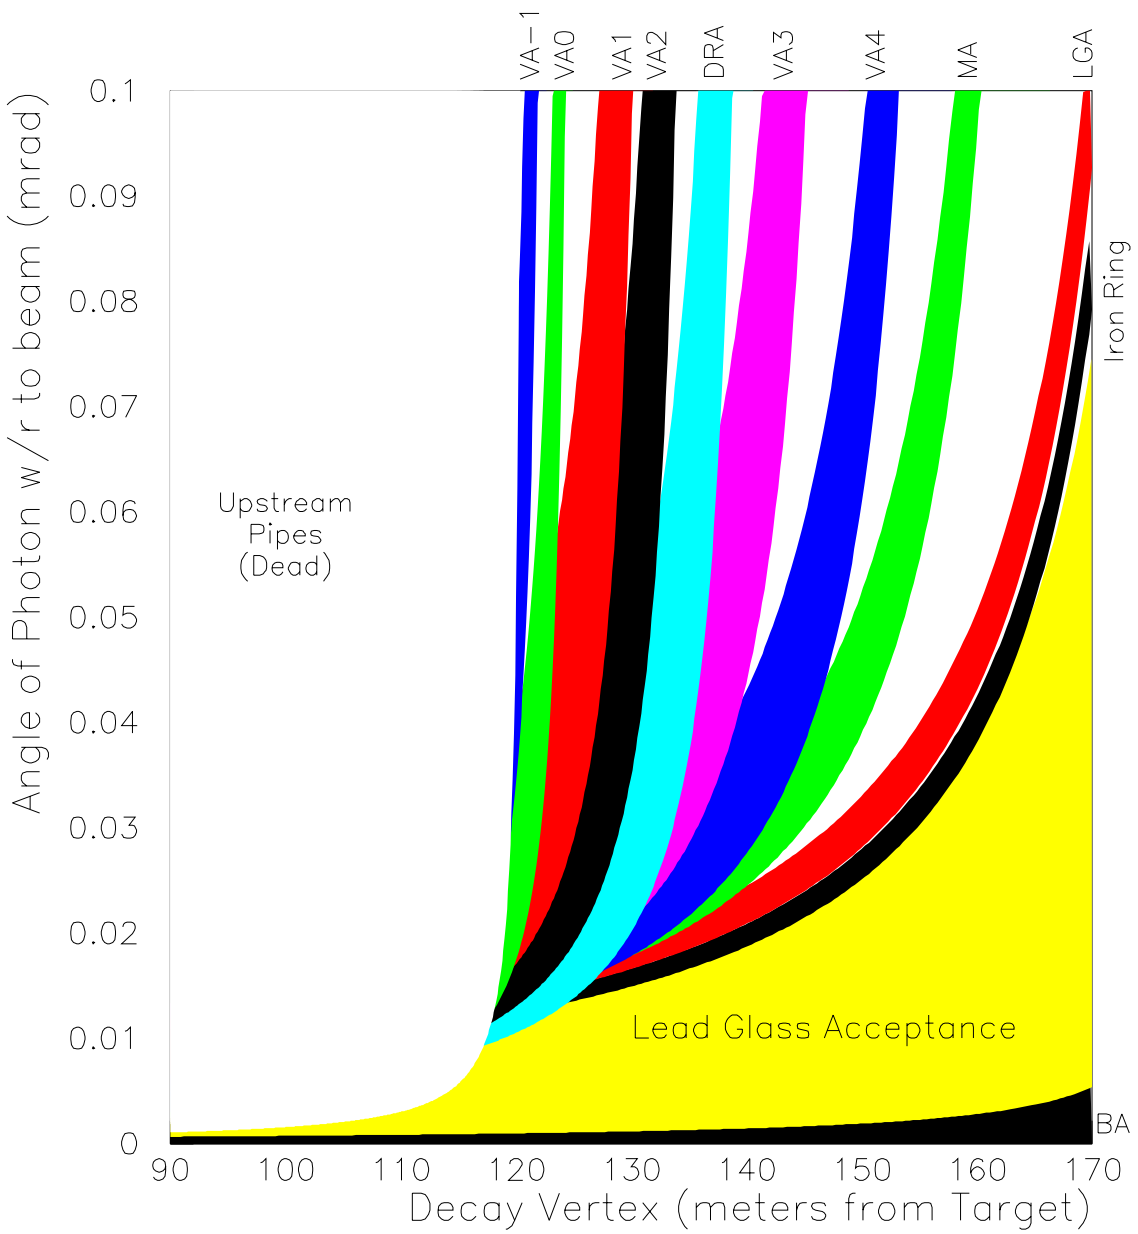

Figure 25. The coverage of the photon vetoes as a function of decay vertex and photon angle is shown for photons created at the center of the *lower* beam, traveling in the $-\hat{y}$ direction.

Table 8. Summary of trigger and particle identification scintillator banks

Bank	Material	Nuclear Interaction Lengths	Distance from the Target (m) (upstream face of element)
C			179.536
B			179.635
	Lead-Glass Array	1.7	180.886
	Lead Wall	0.7	183.2
MU1			184.0
	Muon Steel	20.0	186.0
MU3			189.55
MU2			189.91

2.2.3.2 Trigger Planes

Planes of scintillator counters were very important in the E799 trigger because of their efficient detection of charged particles and fast response. By placing scintillator planes downstream of absorbing material, they can also be used to distinguish electrons from charged pions from muons at trigger level. A list of the positions of the scintillator banks and the intervening material is shown in Table 8.

For triggers selecting final states with charged particles, two banks were placed in front of the lead-glass, the C bank with horizontally-oriented counters and the B bank with vertical counters. The geometries of the B and C banks are shown in Figures 26 and 27, respectively. These banks were finely segmented in order to allow the trigger to require hits in separate counters without a large loss in acceptance. The two-track triggers, for example, required hits in two different scintillators of both the B and C banks. There were also two small trigger scintillators placed behind the beamholes of the lead-glass array referred to as the Beamhole Counters, which were used primarily for triggers on Λ^0 decays to detect high-momentum protons going down the beamhole.

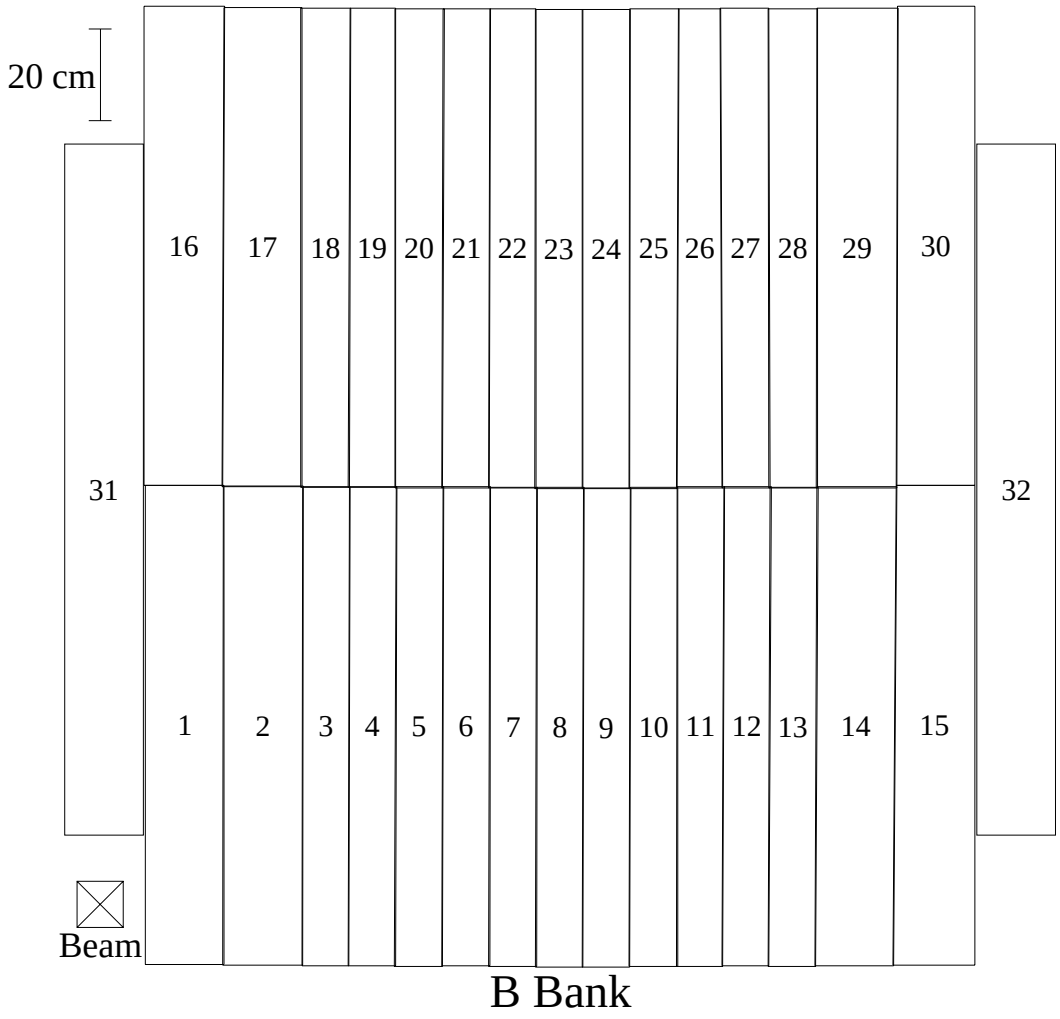

Figure 26. Geometry of the B scintillator bank

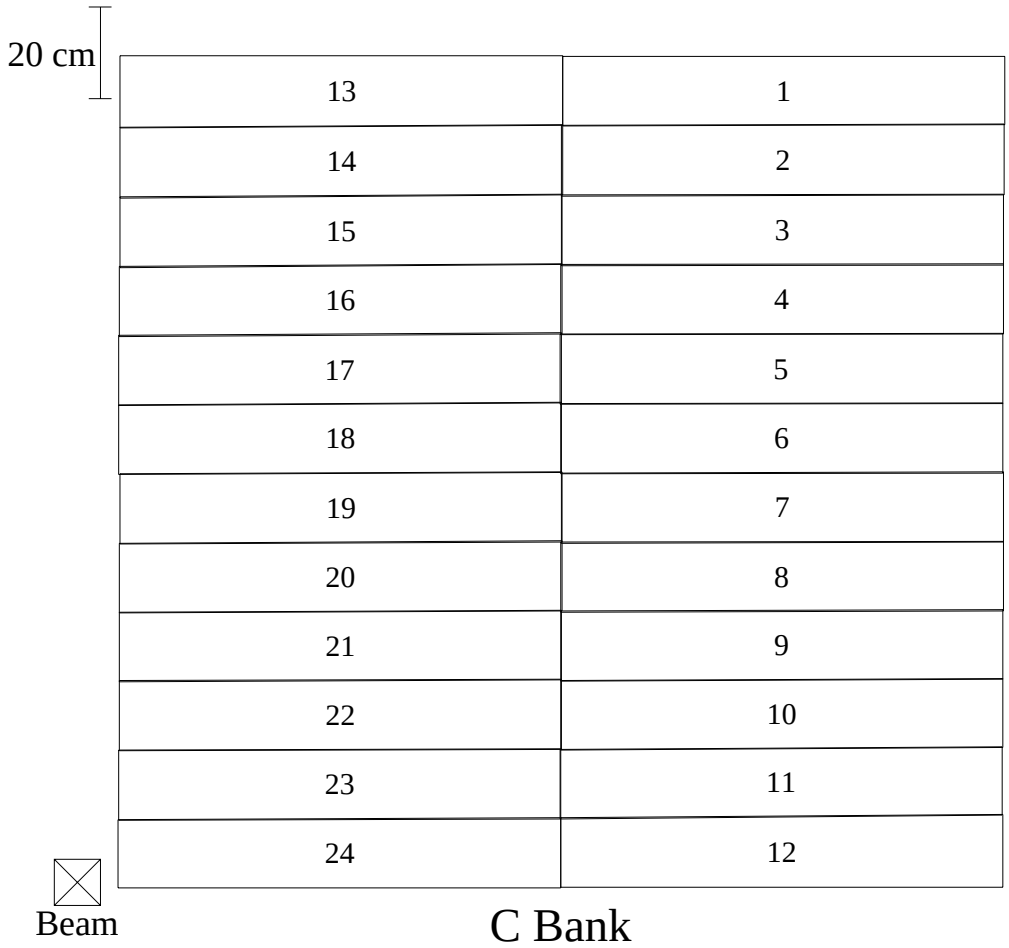

Figure 27. Geometry of the C scintillator bank

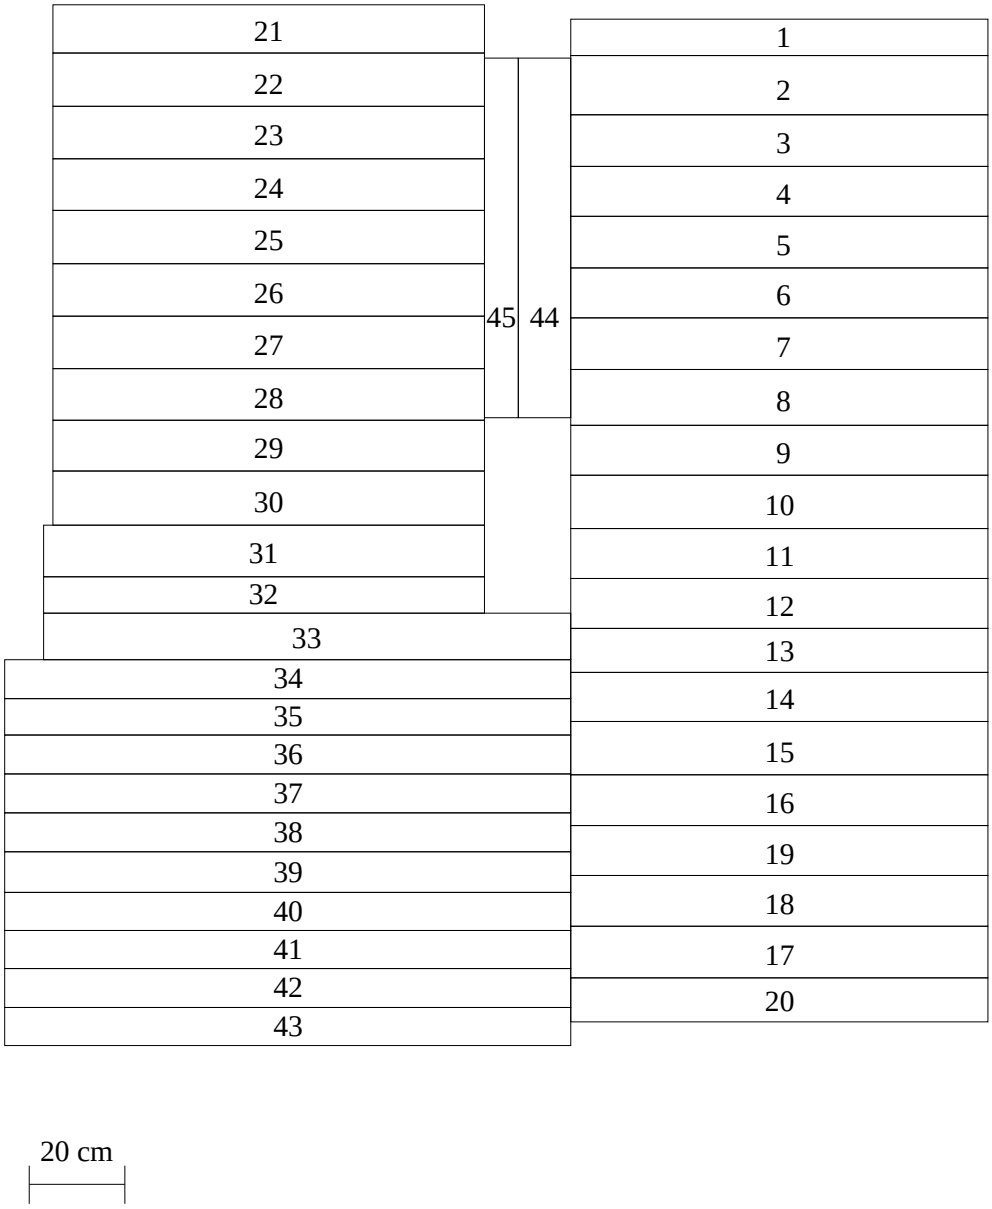

Figure 28. Geometry of the MU1 scintillator bank

For particle identification, the three banks MU1, MU2 and MU3 were used both at trigger level and in the final analysis cuts. MU1 (see Figure 28) was located downstream of the lead-glass and a wall of lead bricks which together comprised a 41 radiation length shield from any electromagnetic showers starting in the lead-glass. Leakage of electromagnetic energy into this bank was extremely small; showers from charged pions, however, are much less well-contained in the lead and will often deposit large amounts of energy in the MU1 scintillator bank. In the trigger, the analog sum of the MU1 counters was required to be below a signal roughly equivalent to that of three minimum ionizing particles. The probability of a single pion initiating a hadronic shower before the end of the lead wall is about 60%. Studies show a 75% rejection of $K_L \to \pi^+\pi^-\pi^0$ decays at trigger level using the MU1 bank; therefore, one concludes that the probability of rejecting a charged pion which showers before the MU1 bank is approximately 80%. Downstream of a 3 meter thick steel muon filter are two muon identification trigger planes, MU2 and MU3. MU2 is a large bank of overlapping scintillator counters designed primarily for the positive detection of single muons passing through the spectrometer. MU3 is a smaller bank of non-overlapping vertically-oriented scintillator counters, flush against the muon filter, equipped with special electronics to trigger on the presence of muons in two non-adjacent counters. Further details of the use and calibration of these banks can be found elsewhere [59][1].

2.2.3.3 The Accidental Trigger

It is important in Monte Carlo simulations of the detector response to include some model of the activity in the detector not caused by the actual decay of interest. This activity occurs in two varieties: activity caused by protons interacting in the target (called "accidentals"), and activity uncorrelated with any real particles incident on the detector (called "noise").

The method used to simulate both types of activity is to imbed all events generated in our Monte Carlo inside real detector events take with some low-bias trigger. If one only wished to include noise in this simulation, the trigger could be essentially random; however, the desire to include "accidentals" as well forces a trigger to be chosen which selects events with protons interacting in the target. Such a trigger must be totally independent of the rest of the detector to avoid biasing this sample. Therefore, a special accidental counter was

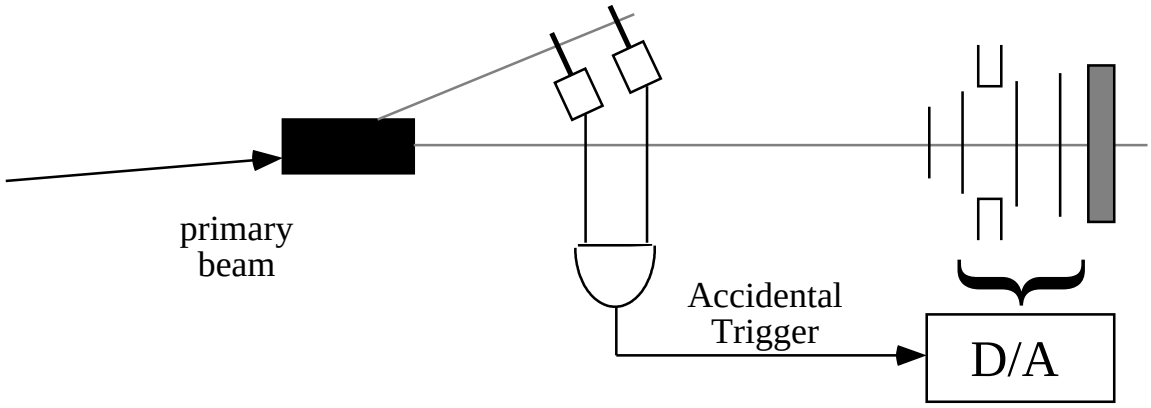

Figure 29. The accidental trigger counter. When the counter detects a muon from the target, the data acquisition system records the activity in the detector.

installed to select events with muons originating in our target, but which are well outside the fiducial volume of the detector. The muon telescope used to form this trigger, shown in Figure 29, was placed about 30 m from the target at an angle of 25 mrad away from any active element of the detector. It was very unlikely that a single muon could interact in both this accidental trigger muon telescope and in any other part of the detector, thus eliminating bias in the selection of these events.

2.3 The Experimental Trigger

Central to the success of the experiment was the triggering apparatus used to select in real time which events would be written to tape. The design of the trigger was complicated to allow for many different triggers which reflected the diversity of physics topics, and because of the need for the maximum possible rejection over undesirable events. The basic structure of the trigger was two-tiered: the Level 1 trigger provided a decision for each RF bucket using information available from fast detectors, and then the Level 2 trigger, started by a positive Level 1 decision, made decisions based on information from the lead-glass and from the drift chambers before selecting the event to be written to tape.

2.3.1 The Level 1 Trigger

The Level 1 Trigger consisted of a set of trigger "sources" which were fed into a series of Memory Lookup Units (MLU's) that produced trigger decisions based on their programming. The trigger sources were produced by a wide variety of detector elements, as shown in the list in Table 9. The trigger sources were strobed by an RF signal synchronous with the 53 MHz bucket structure of the beam. Among the sources were two prescaled RF signals, which allowed for high prescales to be imposed on frequently satisfied trigger conditions (such as those of the Accidental or Single Muon triggers) in order to reduce their rates to acceptable levels. However because these prescales operated by reducing the frequency of the RF signal instead of by counting events, they were potentially susceptible to bias because of variations in RF bucket occupancy. Therefore, these prescaled RF sources were not used to provide the precision prescales necessary for normalization or signal modes.

Table 9. Summary of Level 1 trigger sources

Detector Element	Source Name	Description
B Bank	1B	one or more counters in the bank hit
	2B	two or more counters in the bank hit
	3B	three or more counters in the bank hit
	BE	one or more counters in the east half of bank hit
	BW	one or more counters in the west half of bank hit
C Bank	1C	one or more counters in the bank hit
	2C	two or more counters in the bank hit
	3C	three or more counters in the bank hit
	CU	one or more counters in the upper half of bank hit
	CD	one or more counters in the lower half of bank hit
B and C Banks	BCPH	requires a hit in a central counter of each bank
Lead-Glass	E_t Low	Rising edge above low ($\approx$ 6 GeV) threshold
	E_t High	Rising edge above high ($\approx$ 50 GeV) threshold
	Q1,Q2,Q3,Q4	one quadrant of glass above low threshold
Photon Vetoes	VA -1, VA0	a hit in the veto scintillator
	VA1-4, DRAC	a hit in the veto scintillator
	LGA	signal above threshold in the LGA lucite counters
	BA12	signal above $\approx$ 2.5 GeV in X-view of first 2/3 of BA
	BA3	signal above $\approx$ 5 GeV in X-view of last 1/3 of BA
MU1	MU1	signal above threshold in analog sum of counters
MU2	MU2	a hit in any counter
MU3	MU3	hits in two non-adjacent MU3 counters
Accidental Trigger	ACCID	a coincidence in the Accidental muon telescope
Beamhole Counters	UPBM	a hit in the counter in the upper beamhole
	DNBM	a hit in the counter in the lower beamhole

The MLUs which made the trigger decisions were programmed remotely using an environment which was able to interpret FORTRAN code containing the logic for various triggers. Eight trigger sources were placed into a single MLU, and a wide variety of logical operations were performed on these sources to produce 16 outputs for each MLU, corresponding to the sixteen available Level 1 trigger sources. The outputs from the six MLU's were then combined with a logical AND to form the final Level 1 trigger decisions. This structure did place some limits on possible trigger combinations, but by appropriate grouping of trigger sources in the first layer of MLUs, all needed trigger logic could be implemented.

The first level trigger also contained a precision prescaling capability in addition to the rough prescaling provided by the prescaled RF trigger sources. An additional module was able to provide integral prescaling factors between 2 and 16 by counting unprescaled Level 1 decisions.

The Level 1 trigger system returned a trigger decision each bucket for each of up to 16 pre-programmed triggers, 12 of which were available for physics uses. If one or more of the triggers was satisfied, the Level 2 trigger processors and digitization of the lead-glass and drift chamber information was started, and the Level 2 trigger system was informed of which Level 1 triggers had been satisfied.

2.3.2 The Level 2 Trigger

The Level 2 trigger consists of two sets of hardware trigger devices, the track processor and the hardware cluster finder (HCF). The final Level 2 decision typically took about 30 μsec to reach, limited by the speed of the HCF which was triggered for each event even if its decision was not used in the trigger. However, if one of the elements of the track processor was able to reject an event before the HCF information was available, a Level 2 abort signal was generated in order to reduce deadtime from the Level 2 trigger processing. Triggers could use information from either the Zuma or Plotin track processor boards, from the hit-counting module of the track processor, and an HCF count of clusters. During the Level 2 processing, the trigger system was dead, or unable to process any more incoming triggers. The elements of the track processor and the HCF are described below.

2.3.2.1 The Track Processor

The "track processor" was the name given to a loose confederation of different systems which all used information from the drift chambers as their inputs. These three different triggering devices, the Zuma board, the Plotin board, and the hit-counting system, were electronics-based trigger processors. All of the electronics shared some common physical characteristics: all boards were housed in a FASTBUS crate and could communicate via FASTBUS protocol with a host PDP-11 which was used to test and initialize boards. The input information for all boards was also supplied via special repeater boards which intercepted the digital drift-chamber signals before the TDCs and made an extra copy of these signals for the track processor. The repeater boards also rearranged the signals for convenient cabling and produced a "wire-or" signal which was the logical OR of sixteen adjacent wires in the two planes of one view of a drift chamber, dividing each chamber-view into between 14 and 18 such wire-ors.

The Zuma board [56], which was used as the primary track-processor triggering device in E773, used the wire-ors produced by the repeater boards as inputs and was designed primarily for simple pattern recognition. At the heart of the Zuma board, was a series of of programmable memories which used the inputs of the wire-ors as an address at which was stored a trigger decision. The Zuma memories were divided into nine lookup tables, the first eight of which contained a first level decision based on a subset of the wire-ors,

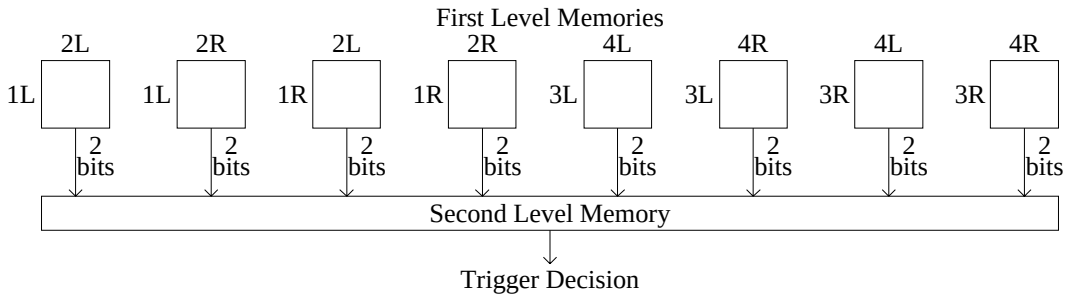

Figure 30. The use of the memories on the Zuma boards. The letters outside the first level inputs signify the inputs used in that memory. In "3L" for example, the "3" indicates the inputs come from chamber 3 and the "L" from the left half of the chamber wire-ors.

and the last of which was a second level memory. The details of this structure are shown in Figure 30. Two identical Zuma boards were used in trigger decision, one for the X view and one for the Y view. As a matter of practice, the boards were programmed such that the first level memories returned the number of observed segments in each subset of wire-ors, and the second level memories checked for a sufficient total number of segments. In E799, the Zuma boards were used in the four-track trigger. In the X view, one segment on each side in the upstream chambers were required, and the downstream chambers were required to contain at least three segments which were compatible with the upstream pattern. In the Y view, at least two upstream segments anywhere in the chambers were required, along with at least three compatible downstream segments. The decision from the Zuma board was returned quite quickly, approximately 100 ns after the gate for the chamber information ended.

The Plotin board was designed specifically for the Λ triggers and used drift chamber information to select high-momentum tracks down the center of the chamber. The board used the 20 center wires from the X view of each chamber as its inputs, where the 10 wires from each of the two sense planes were correctly interleaved with those of the complimentary plane. The board then required a "road" in the chambers with at least one wire hit per chamber. The tracks of interest for this trigger were high-momentum protons from the Λ decays, and these were typically straight through the chamber system with only a very small bend at the magnet. Track segments upstream and downstream of the analyzing magnet were found using the requirement that the wire number hit in one chamber differ by no more than one from the wire number hit in the neighboring chamber. The upstream segment was then matched to the downstream segment by requiring that the same wire number be hit in Chambers 2 and 3. Careful initial chamber alignment was thus key to the operation of the Plotin board. In order to select protons preferentially over anti-protons, a small offset (about 4 mm, the bend distance of a positively-charged 200 GeV track) was added during the chamber-positioning process to the position of the fourth drift chamber. The minimal processing requirement on this board made it quite fast as well, comparable in speed to the Zuma boards.

The final, and most complicated portion of the track processor was the hit-counting

system. The electronics used to perform the hit-counting was originally designed to be part of a more grandiose track-processor with the capability of finding tracks in the drift chamber at trigger level. Figure 31 shows the flow of information among the boards in the originally-planned track processor. However, because the immense number of person-years required to construct such a device, only the controller and front-end sparse readout portions of the total system were functional at the start of the E799 run. These completed portions of the track processor were useful in the Level 2 trigger, however, because they allowed trigger requirements to be placed on the number of hits in each chamber and because of the precision prescaling boards attached to the controller modules. This hit-counting configuration of the track processor, as used in the E799 run, is shown in Figure 32. The electronics used to perform the hit-counting, prescaling and control functions occupied 31 full-sized FASTBUS boards, 24 for the front-end readout, and 7 for the prescaling and controller functions. The boards were physically located next to the drift chamber TDCs, in three FASTBUS crates

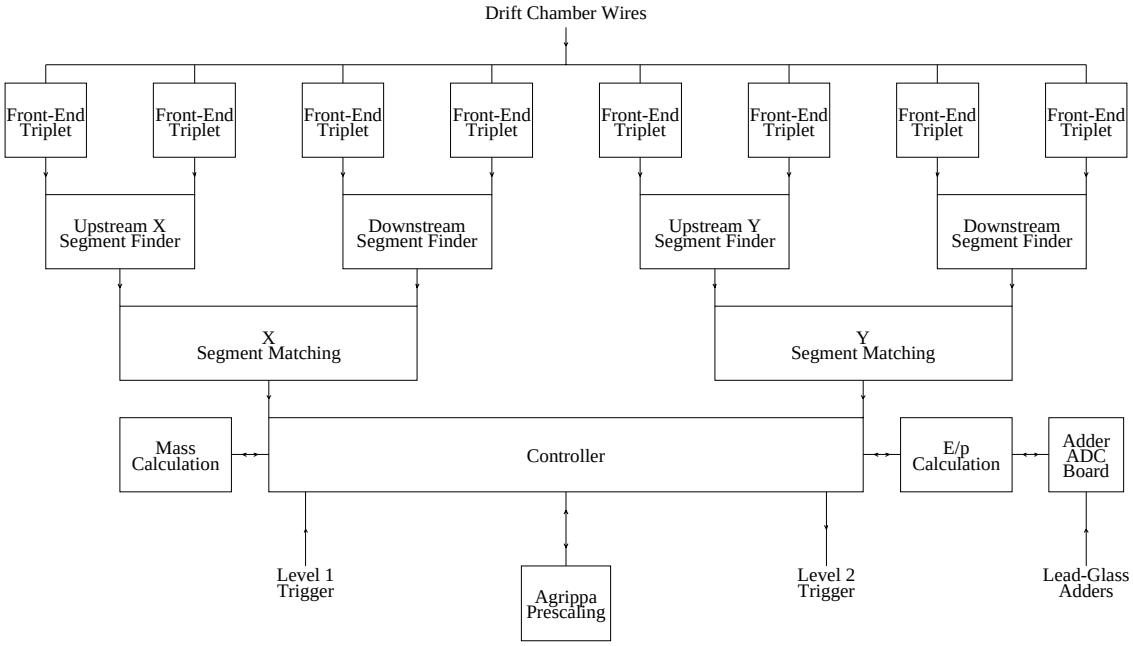

Figure 31. The original track processor design

equipped with water cooling to disperse the nearly 5 kWatts of power consumed by the fast ECL-standard electronics.

The front-end boards which carried out the hit-counting function were originally designed to perform a sparse readout of the drift chamber hits for a set of track-finding hardware processors; therefore, the architecture of these boards was somewhat more intricate than what would have been actually required to simply count hits. A triplet of front-end boards served each pair of sense planes in a single chamber-view. The functions of these triplets could be divided into two broad categories: the "slave" functions of reading drift chamber hits into latches and then reading back that same information on demand, and the "master" functions that controlled the flow of information from the latches onto an internal readout bus. Of the three boards in each front-end triplet, two of them—named "Spartacus" boards—contained the logic necessary only to perform slave functions; the third—called the "Crassus" board—contained a small block of slave logic, plus all the logic necessary to carry out the master functions.

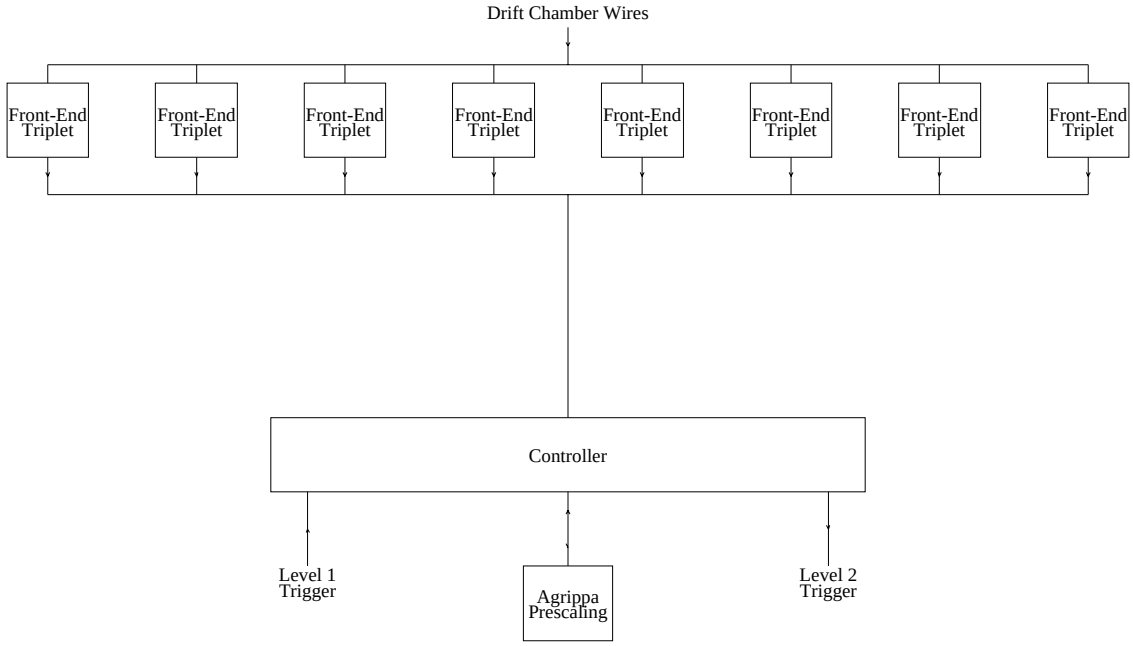

Figure 32. The hit-counting portion of the track processor as used in E799

The slave logic was simple and mostly local. Signals from the drift chamber wires were fanned out to blocks of latch chips. Also distributed to these latches was a gate signal, generated by dedicated hardware triggered by a positive Level 1 decision. If a signal appeared on a drift chamber wire during this gate, the latch at the corresponding location was set. The drift chamber wires from the two complimentary planes in each chamber-view were interleaved before determining the hit location, and a pair of hits in the same chamber cell were considered as a single hit; hits without a corresponding hit in the complimentary chamber plane were assigned to locations in between the locations of the two cells they bordered (see Figure 33 for examples). After the gate was closed, the slave logic then enabled a sparse readout of the hits. The lowest numbered hit in each block of latches would broadcast its location on an internal bus; a similar system tied all the different latch blocks from a triplet of boards together, thus enabling a readout of the first hit location in the chamber-view. Once this address was read out, the latch corresponding to this hit would reset itself, and this readout sequence would repeat until no more latches were set. A special interface to the FASTBUS backplane allowed for the manual loading of hits via the PDP-11 interface for board testing, and also allowed the disabling of latches not corresponding to any physical wire. The boards were tested rigorously before the run and during downtimes in the middle of the run to ensure proper functioning of the latches and readout logic.

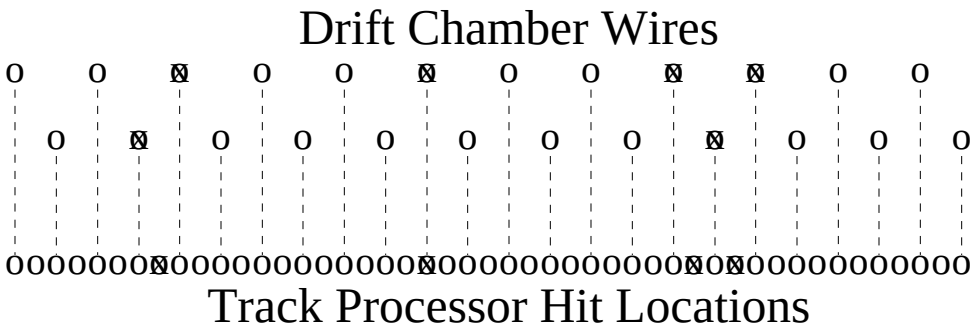

Figure 33. Illustration of track processor front-end hit location assignment

The master logic was somewhat more complicated, and used a hardwired sequencer driven by a 70 MHz clock to orchestrate the readout process. Basically, the role of master logic was to facilitate and direct the readout described above. In addition to the directing the normal flow of the readout logic, the Crassus board was also capable of aborting the readout sequence, resetting the front-end latches, and broadcasting a hit-count and status at the end of the readout sequence to higher level processors. Although most of the logic required for readout was local to each slave board, the master logic was responsible for controlling communications among the different slave boards and for issuing commands from the master to the slaves that were required to initiate or abort readout. A special bus was installed in the FASTBUS backplane to allow this communication between the boards without the use of the comparatively slow FASTBUS protocol. To facilitate board testing, the Crassus sequencer could be run in a special single-step mode, so that logical signals could be checked on the board during the board's readout cycle. Because the higher-level processors to which the Crassus board fed the sparse readout of hits did not exist, the signal from the next level board to proceed with readout was provided by a cable which connected the send and acknowledge signals of the hard-wired hand-shaking protocol on the Crassus boards.

The controller module communicated with the rest of the Level 2 trigger system, performed the readout of the hit-counts from the Crassus boards, and implemented the precision prescaling available with the Agrippa boards. The logic of the controller boards was implemented primarily using a programmable sequencer which allowed for the use of the board in a more limited role than the role of controlling the Level 2 trigger system for which it was originally designed. The operation of the controller module in E799 is best described sequentially. The controller first received a list of valid Level 1 triggers from the trigger system. At this point, it sent this list to the Agrippa prescaling board, which returned a new set of the Level 1 triggers satisfied after taking their prescales into account. This prescaled set of triggers was then sent to the Level 2 logic for use in decisions involving the faster Zuma or Plotin boards. The controller then polled the Crassus boards as they finished their readout to get the hit-count in each chamber view. Cuts on these hit-counts were then applied, based on the requirements of each Level 1 trigger type, and then those

trigger types whose hit-counting requirements were satisfied, if any, were passed back to the Level 2 trigger for use with the HCF decision. If at any time no valid Level 1 triggers remained after cuts, the Level 2 trigger could be aborted by the controller. Similarly, if a decision from the Zuma or Plotin boards was negative, or if Level 2 trigger was aborted for any other reason, the controller order the Crassus boards to abort the readout and reset its latches.

As actually used in the trigger, the hit-counting implemented cuts requiring at least two hits per chamber-view in most of the triggers designed for modes with two final state particles, as well as a cut requiring four hits per chamber-view (except for Chamber 1, Y-view, where three hits were required) which was used in a four-prong trigger. The precision prescaling of the Agrippa board was also used in several of the normalization triggers in order to obtain necessary prescales beyond the factor of 16 available in the Level 1 trigger. The trigger decision, for an unaborted full cycle of the hit-counting, was typically returned within a few μsec of the Level 2 trigger.

2.3.2.2 The Hardware Cluster Finder

The Hardware Cluster Finder (HCF) has been described in great detail elsewhere [60][61], and so will be only be described briefly here.

The HCF was designed to find and count clusters in the lead-glass calorimeter at trigger level. As its input, it took a small portion of the charge from lead-glass phototubes which it then amplified, integrated over a 30 ns gate and digitized with flash ADCs. This digitized value was then required to be over a certain threshold, and all such blocks above threshold were stored as HCF hits on the processing boards. Clusters were identified by grouping of HCF hits together if they were in neighboring blocks. (Blocks which shared one corner were counted as neighbors.) The logic of the HCF hardware started each cluster with a "seed" block, and this seed block, along with all the other HCF hits were read out into the data stream for accepted events. The most important HCF output, however, was the number of clusters found, which was used in the final Level 2 trigger decisions.. The HCF was the slowest of the hardware Level 2 trigger processors, taking about 25 μsec to process a typical four-cluster event.

Table 10. Typical rates at each stage of the trigger

Stage	Rate
RF	53 MHz
After Level 1 (includes some prescaling) (includes deadtime)	32 kHZ
After Track Processor (includes some prescaling)	2.3 kHZ
After HCF (rate to tape)	450 Hz
47% Deadtime (all sources) 1.43×10^{12} protons per spill	

Previous experiments had found that the HCF injected a small amount of charge into the lead-glass ADCs when it was triggered. In E799, the HCF was triggered every event, whether or not it was used in the trigger decision. This avoided the problem of having the HCF affect different triggers differently and ensured that this charge injection would cancel out in calibration.

2.3.3 Trigger Performance

Table 10 shows typical trigger rates during E799 running. At the nominal primary beam intensity, the system was unable to accept new trigger approximately 50% of the time. It is difficult to estimate the deadtime from each individual element of the trigger and data acquisition systems, partly because of the widely varying time required to make track-processor decisions, depending on whether the event was rejected because of a prescale, because of a fast decision from the Zuma or Plotin boards, or because of the hit-counting. It is known from runs without the track processor portion of the trigger that the average dead time from digitization and event readout was approximately 950 μs, dominating the

Table 11. Contributions to dead time

Source	Fraction of Time Dead Due to Source
Level 1 Trigger	none
Track Processor	.04
HCF	.07
Data Acquisition	.42

observed deadtime. Using this to estimate the contribution to dead time from the data acquisition system, one reaches the conclusions regarding the sources of deadtime shown in Table 11. More details of the performance of each individual trigger are given in Chapter 3.

2.4 Data Acquisition

Once an event was accepted by the trigger, all the information in an event was gathered and written to tape. Once all digitization was completed, data was read via FASTBUS and CAMAC interfaces, and a FASTBUS crate which contained crate controllers, the event builder, and memories for storing the event data. From this fastbus crate, the event builder shipped completed events to a VME controller crate which controlled the writing of the data to tape. All data was written onto 8mm tapes using a system of four Exabyte drives running in parallel. All four drives wrote almost continuously during data-taking, and spent the time between spills emptying the buffers which had been filled up during the spill; however, since the event-builder was busy building events during the spill, the rate at which it sent events to the VME crate slowed substantially. The data writing rate of this four drive system, approximately 150 kbytes/sec per drive, placed an upper limit of 12000 on the maximum number of events which could be written during 40 usable seconds of the 58 second spill cycle.

2.4.1 The Event Content

The bulk of the information in an event that was written to tape were latched, discriminated signals from scintillator banks or digitized data from various detector elements. The digitized data from ADCs or TDCs were written in a sparse readout format since the vast majority of detector channels were empty in any given event. This had the disadvantage of requiring pedestals to be subtracted in hardware and requiring a threshold to be imposed on all channels to reduce noise entering the data stream. However, such a sparsification was absolutely necessary to limit event size and the time required to read out the event. The data written out from each detector element is shown in Table 12.

Table 12. Summary of detector information written into the data stream

Detector	Information Written Out
VA -1, VA 0 VA 1-4 MA, LGA	The Scintillator Latches Lucite ADCs
DRA	The Upstream Scintillator Latches and ADCs The Downstream Scintillator ADCs
Drift Chambers	TDCs from each wire
B-Bank	Counters Latches(B1-B30) and ADCs(all)
C-Bank	Counter Latches
Lead-Glass	Each Block ADC Each Adder ADC E_t ADC
MU1	Counter ADCs
BA	ADCs
MU2	Counter Latches
MU3	Counter Latches, ADCs and TDCs
TRDs	Each Wire ADC
Level 1 Trigger Sources	Latches TDCs
Track Processor	Hit Counts and Diagnostic Information
HCF	Cluster Counts, Seed Blocks, and All Blocks Above Threshold

CHAPTER 3

THE DATA

Where is the Life we have lost in living?
Where is the wisdom we have lost in knowledge?
Where is the knowledge we have lost in information?
—*T.S. Eliot*

This chapter describes the data written to tape during the running of E799. The running conditions, experimental triggers used, special runs, and the data handling are described below.

3.1 The Data Set

The E799 data set was accumulated over a 14 week running period at the end of 1991. The first twenty days of the run were devoted to the switchover from the previous experiment that ran in the Meson Center beamline, FNAL-E773, and startup for E799. During this startup period, many beam studies were performed; the effects of installing beamhole calorimeters were studied; the experimental triggers were finalized, and the final fixes were made to new equipment installed for E799. After this period, there were approximately 9 weeks of running with the "standard" E799 triggers. The accelerator was down for one week in the middle of the run for magnet repair, and eight days of running time were devoted to $K_L \to \pi^0\gamma\gamma$ special runs.

During the run, approximately 400 Million physics triggers were written onto 1330 8mm tapes. Under optimal running conditions, approximately 6.5×10^5 triggers could be written to tape per hour. Over the 9 weeks of running, the accelerator delivered 60% of

the requested primary beam intensity. Our experiment spent about 70% of this available accelerator time writing physics events. The remainder of the beam time was spent on special runs, equipment repair, lead-glass cures and calibrations, and normal operational downtimes.

3.2 The Triggers of E799

E799, though conceived originally as an experiment to search for the rare decay $K_L \rightarrow \pi^0 e^+ e^-$, actually had quite general capabilities for the detection of rare K_L decays. Because of the large angle acceptance of the calorimeter and spectrometer, acceptance for multi-body final states is quite good compared to experiments making dedicated searches for two-body K_L final states like $K_L \rightarrow \mu^\pm e^\mp$. This allows for a large variety of different decays to be studied, and final states with up to eight particles were triggered on and studied. The high multi-body acceptance was of course key to the idea of π^0 tagging via $K_L \rightarrow \pi^0\pi^0\pi^0$ discussed in Chapter 1.

Because of its high flux and good multi-body acceptance, E799 was in principle in a position to make improved measurements on all non-two-body rare K_L decays. However, the realities of the trigger rates in a high-flux environment and the bandwidth of the data acquisition system forced certain compromises along the way. For example, in modes with charged pions the trigger rates were too high in order to enable effective triggering. This was because of two problems that result from charged pions. First, the rates from the semi-leptonic decays $K_L \rightarrow \pi^\pm e^\mp \nu$ and $K_L \rightarrow \pi^\pm \mu^\mp \nu$ provided a formidable background, at least at trigger level, to any such search. In modes without charged pions, the trigger rate was kept manageable by the use of the MU1 veto; however, such a veto was obviously unusable in modes involving charged pion final states. Second, a charged pion in the lead-glass would frequently leave no clusters above the threshold of the hardware cluster finder, but could occasionally leave one or more such clusters due to the large transverse spread typical of hadronic showers. This made using the HCF to count multiplicity of final state particles in the decay very inefficient. These two factors caused triggers to search for decay modes such as $K_L \rightarrow \pi^0\pi^\pm e^\mp \nu$ and $K_L \rightarrow \pi^0\pi^\pm \mu^\mp \nu$ to have prohibitively high trigger rates for

implementation in E799. In fact, the only rare modes with charged pions which could be triggered on were four-track modes where the hit-counting capability of the track processor was able to provide a large rejection at trigger level over background from semi-leptonic decays.

All-electromagnetic final states, particularly those with high particle multiplicity in the glass, and two-muon final states were significantly less costly to trigger on, and such modes made up the bulk of the physics agenda of E799. A secondary line of physics investigations attempted in E799 were studies of Λ^0 beta decays. Although these decays were typically low multiplicity and contained charged hadrons in the final state, the trigger rate was kept manageable because of the special geometry of these decays. Since the Λ^0 lifetime is fairly short, $c\tau \approx 7.9$ cm, only very energetic Λ^0's, typically between 200 and 400 GeV, survive to decay in a region of high acceptance. Because most of the decaying Λ^0's momentum is carried off by the daughter proton, the presence of a high momentum proton was a distinctive Λ^0 decay signature. Therefore the trigger required a hit in the scintillators behind the beamholes of the lead-glass, and a stiff track detected by the Plotin board of the track processor. This tracking requirement in conjunction with the requirement that the proton not hit the lead-glass not only greatly reduced the backgrounds from sources other than Λ^0 decays, but also allowed the use of the HCF to count multiplicity of the remaining decay products since there was no proton shower in the glass.

The list of "physics" triggers and the associated decay modes is shown in Table 13. Because the data acquisition system typically ran 50% dead, triggers were at a premium. The prescales listed in Table 13 reflect a desire to balance the rate of various triggers against their perceived importance. During typical running conditions, the number of triggers of each type written per spill is shown in Table 14. These numbers were highly variable depending on running conditions such as the primary beam intensity on target and the Bucket Occupancy Autocorrelation, discussed in Section 2.1.

The requirements implemented in the various triggers will be discussed in detail in the following sections.

Table 13. Rare K_L decays and the associated experimental triggers

Trigger	Decay Modes Written
Two-Electron (4,6,7,8 Clusters in Glass)	$K_L \to \pi^0 e^+ e^-$ $K_L \to e^+ e^- e^+ e^-$ $K_L \to \pi^0 \pi^0 \pi^0$, $\pi^0 \to e^+ e^-$ $K_L \to \pi^0 \pi^0 e^+ e^-$ $K_L \to \pi^0 \pi^0 \gamma$, $\pi^0 \to e^+ e^- \gamma$ $K_L \to \pi^0 \pi^0 \pi^0$, $\pi^0 \to e^+ e^- \gamma$ $K_L \to \pi^0 \pi^0 \pi^0$, $\pi^0 \to e^+ e^- e^+ e^-$
Dalitz (prescale of 14) (3 or 5 Clusters in Glass)	$K_L \to e^+ e^- \gamma$ $K_L \to \pi^0 \pi^0$, $\pi^0 \to e^+ e^- \gamma$
Dimuon	$K_L \to \pi^0 \mu^+ \mu^-$ $K_L \to e^+ e^- \mu^+ \mu^-$ $K_L \to \mu^+ \mu^- \gamma$
Four-Track (prescale of 10)	$K_L \to \pi^+ \pi^- e^+ e^-$ $K_L \to \pi^+ \pi^- \pi^0$, $\pi^0 \to e^+ e^- \gamma$
Neutral (prescale of 16)	$K_L \to \gamma\gamma$ $K_L \to \pi^0 \pi^0$ $K_L \to \pi^0 \pi^0 \pi^0$
$e\mu$	$K_L \to \pi^0 \mu^\pm e^\mp$ $K_L \to \pi^0 \pi^0 \pi^0$, $\pi^0 \to \mu^\pm e^\mp$
Λ^0 Beta Decay (prescale of 2)	$\Lambda^0 \to p e^- \nu$
Λ^0 Muon Beta Decay (prescale of 16)	$\Lambda^0 \to p \mu^- \nu$
$\Lambda^0 \to p\pi$ (prescale of 600)	$\Lambda^0 \to p \pi^-$
Minimum Bias (prescale 3600)	$K_L \to \pi^+ \pi^- \pi^0$ $K_L \to \pi^+ \pi^-$

Table 14. Triggers per spill written from each trigger

Trigger	Triggers Written per Spill (50% dead, 1.5×10^{12} protons per spill)	% of Total
Two-Electron	4000	33%
Dalitz	950	8%
Dimuon	1400	12%
Four-Track	1500	13%
Neutral	500	4%
$e\mu$	1050	9%
Λ^0 Beta Decay	1900	16%
Λ^0 Muon Beta Decay	150	1.2%
$\Lambda^0 \to p\pi$	100	0.8%
Minimum Bias	300	2%
One-Track Trigger	100	0.8%
Accidental Trigger	100	0.8%

3.2.1 Two Electron Triggers

The Two Electron trigger and the Dalitz trigger were designed to select events with at least two electrons plus extra photons or electrons in the final state. Electrons were selected by requiring hits in the B and C scintillator banks, by counting HCF clusters from the electrons, and by requiring two hits per plane per view in each drift chamber using the hit-counting. In the Dalitz trigger, in which the physics modes of interest had two electrons typically close together, the hit-counting requirement was changed so that only one hit was required in the Y-view of chamber 1. In order to pick out events with all electron or photon final states, the high E_t threshold (about 55 GeV) in the lead-glass was required, and the MU1 hadron veto counter was placed in veto. Since the Dalitz trigger was to allow a search for $K_L \to \pi^0 \nu\overline{\nu}$ in which much of the final state energy is undetected, the Dalitz trigger used the low E_t threshold. In the unprescaled Two Electron trigger, a hardware cluster

finder multiplicity of four, six, seven or eight clusters was required; for the prescaled Dalitz trigger, three and five clusters were selected. Vetoes on the photon veto scintillators, and BA_γ ($\mathrm{BA}_\gamma \equiv \mathrm{BA12} \cdot \overline{\mathrm{BA3}}$) were in place to reject events with charged particles escaping the detector fiducial volume or with photons in the beam hole, respectively. Finally, the MU3 dimuon trigger plane was in veto in the Two Electron triggers for historical reasons involving early problems in the Level 2 trigger handling of events satisfying both the Two Electron and Dimuon Level 1 triggers. The final triggers requirements in the Two Electron and Dalitz triggers were then

$$2\mathrm{B} \cdot 2\mathrm{C} \cdot \mathrm{E_t High} \cdot \overline{\mathrm{MU1}} \cdot \overline{\mathrm{PHV}} \cdot \overline{\mathrm{MU3}} \cdot \overline{\mathrm{BA}_\gamma} \cdot \mathrm{Hitcnt}_{\mathrm{2track}} \cdot \mathrm{HCF}(4,6,7,8),$$

and

$$2\mathrm{B} \cdot 2\mathrm{C} \cdot \mathrm{E_t Low} \cdot \overline{\mathrm{MU1}} \cdot \overline{\mathrm{PHV}} \cdot \overline{\mathrm{MU3}} \cdot \overline{\mathrm{BA}_\gamma} \cdot \mathrm{Hitcnt}_{\mathrm{Dalitz}} \cdot \mathrm{HCF}(3,5),$$

respectively, where

$$\overline{\mathrm{PHV}} = \overline{\mathrm{VA}-1} \cdot \overline{\mathrm{VA0}} \cdot \overline{\mathrm{VA1}} \cdot \overline{\mathrm{VA2}} \cdot \overline{\mathrm{DRAC}} \cdot \overline{\mathrm{VA3}} \cdot \overline{\mathrm{VA4}} \cdot \overline{\mathrm{LGA}}.$$

It is interesting to study the effect of each element of the trigger on the rate. In the Two Electron trigger, after the base trigger of $2\mathrm{B} \cdot 2\mathrm{C} \cdot \mathrm{E_t Hi}$, the MU1 veto reduced the Level 1 trigger rate by approximately a factor of 3, while the photon vetoes cut the rate an additional 30%. The hit-counting requirements added another factor of 3 rejection over the Level 1 trigger rates, and the HCF multiplicity requirement further reduced the rate a factor of 8. These rejection factors are quoted for typical running conditions, and could vary significantly depending on primary beam intensity and Bucket Occupancy Autocorrelation.

3.2.2 The Dimuon Trigger

The Dimuon Trigger was designed to select events with two muons and one or more photons in the calorimeter. Muon selection was done by requiring a dimuon trigger in the MU3 bank, two hits in each of the B and C banks, and by requiring a two track trigger from the hit-counting. Events with photons were selected by requiring E_t in the calorimeter above the low (6 GeV) threshold. To reduce backgrounds from decays with charged pions, the MU1

bank was placed in veto, and the photon veto scintillators and the BA_γ veto were also used to reject backgrounds with extra particles. The trigger is then

$$2B \cdot 2C \cdot E_t Low \cdot MU3 \cdot \overline{MU1} \cdot \overline{PHV} \cdot Hitcnt_{2Track}.$$

3.2.3 The Four-Track Trigger

The Four-Track trigger was intended to select candidates for the four-prong K_L decay mode $K_L \to \pi^+\pi^-e^+e^-$. Those modes with four tracks containing at least two electrons or at least two muons and no charged pions could be found in the Two Electron and Dimuon triggers, respectively. The basic trigger required three hits in each of the B and C banks including a signal from each half of each bank, sufficient calorimeter energy to satisfy the low E_t threshold, and between two and eight HCF clusters in the lead-glass. The core of the trigger requirement, rejecting almost a factor of 100 in trigger rate, were requirements of four hits in each view of each drift chamber with the hit-counting boards, and a four-track topology requirement using the Zuma boards. Since a Zuma board decision was required in this trigger and since Zuma and Plotin decisions could not be considered in the same event due to limitations of the Level 2 trigger logic, it was imperative that no events satisfy both the Level 1 Four-Track and Λ^0 triggers. To avoid any overlap with the Λ^0 Beta Decay and Λ^0 Muon Beta Decay triggers, MU1 was *required* to fire in the Four Track trigger. Because the rate was still prohibitively high for this trigger, it was prescaled by a factor of 10. The Four-Track trigger was then

$$3B \cdot 3C \cdot BE \cdot BW \cdot CU \cdot CD \cdot E_t Low \cdot MU1 \cdot \overline{PHV} \cdot Zuma \cdot Hitcnt_{4Track}.$$

3.2.4 The Neutral Trigger

A Neutral Trigger was taken to study the $K_L \to \gamma\gamma$ decay, as well as to provide a sample of $K_L \to \pi^0\pi^0\pi^0$ decays for glass studies. The base trigger was simply the high E_t requirement, along with a requirement that there be energy in two opposite quadrants of the glass. In veto were the photon veto scintillators, BA_γ, the B and C scintillator banks, the MU1 hadron veto and the MU2 muon veto. Events were also required to have two, four, six,

seven or eight hardware clusters. Because studies of neutral triggers did not require all the available flux, the trigger was prescaled by a factor of 16. The Neutral Trigger was then

$$\mathrm{E_tHi} \cdot [(\mathrm{Q1} \cdot \mathrm{Q3}) + (\mathrm{Q2} \cdot \mathrm{Q4})] \cdot \overline{\mathrm{1B}} \cdot \overline{\mathrm{1C}} \cdot \overline{\mathrm{PHV}} \cdot \overline{\mathrm{MU1}} \cdot \overline{\mathrm{MU2}} \cdot \overline{\mathrm{BA}_\gamma} \cdot \mathrm{HCF}(2,4,6,7,8).$$

3.2.5 The $e\mu$ Trigger

To study the lepton-flavor-number violating decays $K_L \to \pi^0\mu^\pm e^\mp$ and $\pi^0 \to \mu^\pm e^\mp$, tagged by $K_L \to \pi^0\pi^0\pi^0$, an $e\mu$ Trigger was taken. This trigger was similar in all requirements to the Two Electron trigger, except that the muon was not expected to contribute a cluster to the lead glass, and the MU2 trigger bank was required, with the veto on the MU3 dimuon bank dropped. Explicitly, the $e\mu$ Trigger required

$$\mathrm{2B} \cdot \mathrm{2C} \cdot \mathrm{E_tHi} \cdot \mathrm{MU2} \cdot \overline{\mathrm{MU1}} \cdot \overline{\mathrm{PHV}} \cdot \overline{\mathrm{BA}_\gamma} \cdot \mathrm{Hitcnt}_{\mathrm{2track}} \cdot \mathrm{HCF}(3,5).$$

3.2.6 The Λ^0 Triggers

Studies of Λ^0 beta decays $\Lambda^0 \to pl^-\nu$ with electrons or muons required three triggers, one to take electron beta decays, one for muon beta decays and one for $\Lambda^0 \to p\pi^-$ decays for normalization and polarization studies. All three triggers shared the common requirements of two hits in the B and C banks, a stiff track found by the Plotin board, the hit-counting two-track trigger, and a veto on hits in the photon veto scintillators. The $\Lambda^0 \to p\pi^-$ normalization trigger also had a veto on three hits in both the B and C banks, to avoid Level 1 overlap with the Four Track trigger, as discussed in Section 3.2.3. The $\Lambda^0 \to p\mu^-\nu$ trigger required a hit in the MU2 muon trigger plane, zero clusters in the hardware cluster finder, a veto on the low E_t threshold, the MU1 hadron veto, and a hit in one of the two beamhole scintillators. Finally, the $\Lambda^0 \to pe^-\nu$ trigger, which had an electron in the lead glass, required exactly one hardware cluster, enough energy in the lead glass to satisfy the low E_t threshold, and lead-glass energy above this threshold in only one quadrant with a hit in the matching B and C bank quadrant. Like the muon beta decay trigger, the $\Lambda^0 \to pe^-\nu$ trigger also required a hit in one of the two beamhole scintillators. The three triggers, $\Lambda^0 \to p\pi^-$, $\Lambda^0 \to p\mu^-\nu$, and $\Lambda^0 \to pe^-\nu$, were then

$$\mathrm{Lambda} \cdot \overline{\mathrm{PHV}} \cdot \overline{(\mathrm{3B} \cdot \mathrm{3C})},$$

$$\text{Lambda} \cdot (\text{UPBM} + \text{DNBM}) \cdot \overline{\text{PHV}} \cdot \overline{\text{E}_\text{t}\text{Low}} \cdot \overline{\text{MU1}} \cdot \text{MU2} \cdot \text{HCF}(0),$$

and

$$\text{Lambda} \cdot (\text{UPBM} + \text{DNBM}) \cdot \overline{\text{PHV}} \cdot \overline{\text{MU1}} \cdot \text{E}_\text{t}\text{Low} \cdot (\exists!\text{Qn}) \cdot \text{HCF}(1),$$

respectively, where

$$\text{Lambda} = 2\text{B} \cdot 2\text{C} \cdot \text{Hitcnt}_{2\text{track}} \cdot \text{Plotin}.$$

Because of constraints on the trigger rates, the $\Lambda^0 \to p\pi^-$ trigger was prescaled by a factor of 100, the $\Lambda^0 \to p\mu^-\nu$ by 16, and the $\Lambda^0 \to pe^-\nu$ trigger by a factor of 2.

3.2.7 Minimum Bias Trigger

Finally, there was a two-track Minimum Bias trigger, which essentially required two hits in the B and C banks, plus a two-track hit-counting trigger from the track processor, along with the standard veto on photon veto scintillators. This trigger had a precision counted prescale of 3600, of which a factor of 16 was provided by the Level 1 trigger and a factor of 225 provided by the track-processor Agrippa board. It was designed to collect $K_L \to \pi^+\pi^-\pi^0$ decays for normalization of the physics modes in the Dimuon and $e\mu$ triggers. This trigger also accepted $K_L \to \pi^+\pi^-$ decays which were used to study the momentum scale of the charged spectrometer (see Section 5.6.1). Explicitly the Minimum Bias trigger requirement was

$$2\text{B} \cdot 2\text{C} \cdot \overline{\text{PHV}} \cdot \text{Hitcnt}_{2\text{track}}.$$

3.2.8 Other Triggers

There were a plethora of other, "non-physics" triggers also taken during the run for the purpose of monitoring the data. One such trigger was the Accidental trigger, designed to supply a sample of events uncorrelated with kaon decays in the detector in which to embed Monte Carlo events. The hardware which supplied this trigger was described in Section 2.2.3.3, and a prescale of 200, supplied by a prescaled RF trigger source of the Level 1 trigger, was used to suppress the rate to a reasonable level. A One-Track trigger was also taken with a prescale of 7×10^4, again supplied by a prescaled RF source, requiring one hit each in the B and C banks and with MU1 in veto. These One-Track events were written to

tape for studies of track processor efficiency, and therefore no track processor requirement was used in the trigger. For studies of changes in the lead-glass gains, Flasher triggers coincident with flashes of the Xenon monitoring lamp (see section 2.2.2) were written both during and in between spills. Finally, in between spills, the hardware ADC pedestals were written out during a special Beginning-of-Spill event, and scalers from the trigger sources and the track processor prescaling module were written during an End-of-Spill event.

3.2.9 Performance of E799 Triggers

It is of some interest to examine the trigger rates from each of the E799 triggers at each stage of trigger processing. Shown in Table 15 is this information for selected E799 triggers. Note the wide variety of rejection required from the Level 2 trigger system depending on how stringent the Level 1 trigger requirements are for each trigger.

3.3 Special Runs

Aside from the normal data-taking of E799, frequent special runs were performed to provide additional monitoring of the detector. The most frequent of these were the muon alignment runs, taken daily with the analysis magnet off and a loose muon trigger in order to have a sample of straight tracks with which to align the drift chambers. These runs typically lasted only 10 spills, during which approximately 1.5×10^5 events were written to tape. Another class of special muon runs was taken for location and efficiency determinations of the muon trigger bank. Because it was of paramount importance to have a pure muon sample for these studies, the beam was not steered into our target, but rather was dumped into a beam-stop to ensure no hadrons would be present in our beamline. These beam dump muons were written to tape, using a loose trigger, with the analysis magnet on, so they could be momentum-analyzed. Typically this data was taken during periods when a part of the detector other than the chambers was being repaired or during cures of the lead-glass, and a total of approximately 10^7 such "beam-stop-closed" muon triggers were written during the run.

For calibration and monitoring of the lead-glass, special electron calibration runs were

Table 15. Rate at each level of the trigger for each experimental trigger

Trigger	Level 1 Prescale	Level 1 Rate (Hz)	Level 2 Prescale	Rate After TRKP (Hz)	Rate After HCF (Hz)
(50% dead, 1.5×10^{12} protons/spill, 21 sec spill, rates deadtime corrected)					
Two-Electron		9500		3100	400
Dalitz	14	670		220	95
Dimuon		360		140	140
Four-Track	10	15000		150	150
Neutral	16	180		180	38
$e\mu$		1300		420	100
Λ^0 Beta Decay	2	5700		200	170
Λ^0 Muon	16	140		13	11
$\Lambda^0 \rightarrow p\pi$	600	400		9.0	9.0
Minimum Bias	16	29000	225	29	29
One-Track	10^5	9.0		9.0	9.0
Accidental	300	7.1		7.1	7.1
ALL TRIGGERS		62000		4400	1100

performed at six points during the run. Typically, the calorimeter was calibrated with electrons at one point during each curing cycle. To produce electrons for calibration, the beam was run at low intensity with the lead absorber removed to increase the photon content of the beam. Just downstream of the last sweeping magnet, a sheet of 0.09 radiation lengths of Copper and 0.02 radiation lengths of Titanium was inserted in order to convert some fraction of the photons in the beam. The electron-positron pairs were separated by the AN1 and AN2 magnets, located 120 meters and 140 meters downstream of the target, respectively, and different illuminations of electrons at the glass could be achieved by tuning the fields in these magnets. Approximately 1.5 million good electrons in the glass were collected in each calibration run. After calibrations, cures, and voltage changes in the

lead-glass phototubes, special Flasher Trigger runs were also taken in order to have good statistics with the Flashers to determine instantaneous gains for monitoring the change in lead-glass response.

Finally, about ten days of beam time were devoted to an experiment designed to search for the rare decays $K_L \to \pi^0\gamma\gamma$ and $K_L \to \pi^0\pi^0\gamma$. This experiment involved the use of a pre-shower detector which was placed in front of the lead-glass calorimeter in order to reduce backgrounds in the two modes from $K_L \to \pi^0\pi^0\pi^0$with two photons sufficiently close in the calorimeter so that they "fuse" into a single photon in the clustering algorithm. A more complete description of the pre-shower detector and the searches for $K_L \to \pi^0\gamma\gamma$ and $K_L \to \pi^0\pi^0\gamma$ can be found elsewhere [62].

The pre-shower detector would have also had benefits for other rare decay searches, but several complications contraindicated its use in the bulk of E799. First, the presence of the pre-shower detector did degrade the energy resolution of the glass, although the available position resolution was superior. Second, there were no electronics for implementing the pre-shower detector in the trigger, and the presence of the pre-shower detector greatly reduced both the efficiency and background rejection of the calorimeter E_t triggers and the HCF. Third, the presence of the material in the pre-shower detector in front of the lead-glass could potentially have increased the radiation damage in the lead-glass. For these reasons, the pre-shower detector was only used in these special runs, separate from the bulk of the data taking.

3.4 Online Monitoring

During the run, the functioning of the detector and the quality of data written to tape was constantly being monitored. Low-level detector characteristics, such as drift chamber gas flows and voltages, scintillator bank phototube voltages and currents, inflation of the constantly-leaky Helium bags between the drift chambers and residual gas in the vacuum tank were checked three times per day. In addition, scalers which recorded the rates of all trigger sources and primary beam intensity for each spill were checked for any variation which could indicate detector or trigger problems.

The most powerful tool available for monitoring the data being written to tape was an online monitoring package, which was able to analyze events as they were written to tape. In parallel with the process of writing an event to tape, the event was hoisted from a buffer to a VAXStation 3200 which then could run the standard analysis code on each event. The speed of this procedure was limited by both the data transfer rate to the VAXStation and the time required to execute the analysis code. Typically about 100 events per spill could be fully analyzed with this system. This allowed for quick detection of dead channels in detector elements by looking at illuminations in each detector, and for online monitoring of trigger efficiencies, of instantaneous kaon yields and of other relevant quantities.

As mentioned in Section 2.1, there was also a set of scalers installed for the second half of the run to monitor the Beam Occupancy Autocorrelation. These scalers counted coincidences between Accidental Triggers and MU2 triggers in the same bucket, and also counted coincidences with the MU2 trigger signal delayed by 7 buckets. This provided a measurement of the uniformity of the spill, and allowed our experiment to provide feedback to accelerator physicists when they attempted to correct the problem.

3.5 The Split

After the data taking was completed, the first step in data analysis was to separate the different triggers on tape. Because all triggers were written simultaneously, the different triggers were mixed together on the raw data tapes. The meant that in the data reduction procedure for each trigger, each tape would need to be read many times (once for each trigger) as physics data from each trigger was analyzed. Concern about loss of data due to failure of the 8mm tapes after many readings dictated that the data should be initially split off into different streams and written onto separate tapes for subsequent analysis. Fortunately, the observed 8mm tape failure rate, even after many readings, was negligible, so in retrospect this procedure was probably an unnecessary precaution. The split did however simplify the process of initial data set reduction immensely, especially for streams with relatively small amounts of data.

The split was performed on two VAXStation 3200s at Fermilab, each with six 8mm

Exabyte drives. Eight streams of output were created from the original data tapes, and the contents of each stream are listed in Table 16. As the split progressed, tapes were made available for the initial analysis of each data set. The split took approximately 3 months to complete.

After this, the data was ready for physics analysis. However, before the analysis could proceed in earnest, the calibration of the detector has to be completed. The next two chapters describe the calibration of the spectrometer and the lead-glass respectively.

Table 16. The streams of the split

Stream	Triggers Included	Fraction of Data In Split Stream
Two-Electron	Two-Electron Dalitz	39%
Λ^0	Λ^0 Beta Decay Λ^0 Muon Beta Decay $\Lambda^0 \rightarrow p\pi^-$	14%
Muon	Dimuon $e\mu$ Minimum Bias	21%
Four-Track	Four-Track Dimuon* Two-Electron*	34%
Neutral	Neutral	6%
Minimum Bias	Minimum Bias One-Track	3%
Accidental	Accidental	3%
Calibration	Flasher Beginning of Spill End of Spill	2%

* For events in these triggers to be included in Four-Track split stream, the hit-counting was required to have found 2 hits per plane in Chamber 1, 3 hits per plane in Chamber 2, and 4 hits per plane in Chambers 3 and 4.

CHAPTER 4

CLUSTER AND TRACK RECONSTRUCTION

This chapter describes the algorithms used in the basic building-blocks of the data analysis: cluster-finding in the lead-glass, tracking in the drift chambers, and the matching of tracks to clusters. The algorithms described in this chapter were those used in the analysis of data collected in the Two Electron trigger. For other data samples, different algorithms more appropriate to those analyses were often used in place of these, particularly for cluster-finding and the track-cluster matching.

4.1 Cluster-Finding Algorithm

The algorithm for finding clusters in the lead-glass made ample use of information from the Hardware Cluster Finder (HCF). Because the HCF was used in the Two-Electron trigger to select events with exactly the number of final state particles expected in the lead-glass for a given decay mode, it was important to identify each final state particle with a cluster found by the HCF. Fortunately, the readout of the HCF provided a list of the seeds of all clusters found; by using this list of HCF clusters as the starting point for searching for lead-glass clusters, the clustering algorithm was guaranteed to find the clusters selected by the online trigger.

To locate clusters, the local maximum is found using a "walking" search algorithm, starting from each HCF seed block. From a starting block, the search proceeds to the one of the eight neighboring blocks containing the maximum energy. If the starting block contains more energy than the most energetic neighboring block, than that block is identified

as a local maximum. The procedure is iterated until the local maximum corresponding to each HCF seed block is found.

After a cluster center had been found for each HCF seed block, the HCF clusters were then checked for fusions. Because the HCF operated by searching for disjoint clusters of blocks above a fixed threshold, it was possible for the HCF to find only one cluster when more than one cluster above the HCF threshold was present, as illustrated in Figure 34. To identify such candidates, the search for local maxima was repeated, this time using every block found above the threshold of the HCF. If new local maxima were found, a new cluster corresponding to this local maximum was added to the list of clusters. In the Two-Electron trigger, such a fusion occurred in approximately 7% of events. Most of these fusions were due to nearby electrons or photons from K_L decays, although charged pion showers and accidental effects played a role also.

With the complete list of cluster centers, one could find the cluster position using the algorithms described in Section 6.4 and then assign a preliminary energy to the cluster using

	0.2	0.6	1.2	0.5
0.1	1.6	1.9	9.7	0.9
0.3	4.7	2.7	3.2	0.7
	0.1	0.5	0.3	

Figure 34. Illustration of an HCF Cluster Fusion. The numbers shown in each block correspond to the block energies in GeV. The heavily outlined blocks are those above the HCF threshold.

the sum of the energies in the nine center blocks as described in Section 6.2. Corrections to this energy had to wait until the matching is completed, so clusters can be corrected as either candidate electrons or photons.

4.2 Tracking Algorithm

The next step was to find tracks in the drift chambers. The tracking algorithm was very complex, and attempted to use all information available in the drift chambers when selecting tracks. Because of the high-rate environment in E799, the probability of finding extra tracks or extra drift chamber hits in an event was high. The approach of rejecting as many hits, segment candidates and track candidates as possible by using the timing information at early stages in the track-finding helped to reduce backgrounds and losses in the track-finding from accidental activity in the E799 chambers.

The tracking algorithms differed substantially between the $\hat{x}$-view and the $\hat{y}$-view of the drift chambers because of the track bend at the analyzing magnet in the $\hat{x}$-view. Tracking in both views began by identifying and classifying drift chamber hits. In the $\hat{x}$-view, upstream and downstream segments were then identified, and these segments were matched at the magnet to form complete tracks. In the $\hat{y}$-view, more stringent requirements could be imposed on track candidates because these tracks formed a straight line through the four chambers. The $\hat{x}$-view tracks included an *a priori* unknown bend angle at the magnet. The steps of the track-finding algorithm are discussed in detail in the remainder of this section.

First, the hits in the drift chambers were sorted and classified as either single hits or hit pairs. These classifications were useful because the double measurement of hits in the drift chambers by the two planes in each chamber-view (or the isolation of a single hit) provided useful information in addition to the measured position. For example, hits in the chamber without a complimentary hit in the other chamber plane were rare since typical average uncorrected chamber efficiencies were 98% (see Section 5.7). The probability of finding a track with many such missing hits was then very small, and a track-finding routine could discriminate against such tracks. If hits were present in both planes, the position measurements from the two hits could be checked for consistency by calculating the sum of

the distances between the wires and the point where the charged particle passed through the chamber. If this sum-of-distances (SOD) was significantly higher than the known spacing between the two wires in the complimentary planes (6.35 mm), this suggested a track which arrived in an RF bucket before the formation of the trigger.

If the SOD was low, this suggested a late out-of-time track; however, there was another common occurrence which could also cause a low SOD. Passage of charged particles through the material outside or inside the chambers can produce δ-rays, energetic knock-on electrons. If these δ-rays were sufficiently energetic, they could produce their own ionization in the chamber at some distance away from the primary track. If they were soft, they could also disrupt the measurement by traveling some distance in the chamber before being stopped through collisions with the gas, and then drifting like a secondary electron. If both the primary charged particle and the δ-ray passed through the same drift chamber cell, the effect was that each chamber wire saw first the ionization from the closest ionization track, either the primary or the δ-ray. Because of the 250 nsec double-pulse deadtime of the chamber TDCs (Section 2.2.1), only the first observed signal registered in the common-stop TDCs. The net result was a low SOD measurement from δ-rays, an effect which can be easily observed in plots of the sum-of-distances (see Figure 44). Contributions from accidental chamber hits also tended to lower the SOD because only the earlier hit was measured by the TDC. The treatment of chamber hits in the early part of the tracking then attempted to differentiate between all these different categories, based on SOD and isolation.

Hits in each drift chamber were classified according to the following scheme. First, pairs were identified, in order of the difference between their SODs from the nominal chamber cell size. The chamber hits forming each identified pair were then removed from the list of available hits. Pairs with a SOD greater than 1.5 mm from the cell size were identified as being out-of-time. (The typical width of a sum-of-distance distribution was about 150 μm.) However, tracks with a low SOD were treated differently because of the possibility of one of the hits measuring the true position of an in-time track, while the other resulted from a δ-ray or a second track. The pair was stored in a list of hits for tracks, but the individual hits of the pair were stored as well. Finally, once all pairs of hits were removed from the

Table 17. The tracking algorithm's classification of chamber hits

Classification	SOD Requirement	Status Code
In-Time Pair	$\lvert \mathrm{SOD} - d_{\mathrm{cell}} \rvert < 1.5$ mm	0
Late Pair	$(\mathrm{SOD} - d_{\mathrm{cell}}) < -1.5$ mm	individual hits were stored with status code 1; the pair with 3
Early Pair	$(\mathrm{SOD} - d_{\mathrm{cell}}) > 1.5$ mm	3
Isolated Single Hits	n/a	2
Single Hits Next To a Pair	n/a	1
Single Hits Between Two Pairs	n/a	2

list of chamber hits, any remaining chamber hits were classified as individual hits. The list of classifications for chamber hits and the associated "status codes" is shown in Table 17. For hits with a status code of 1 — i.e.: those hits that were part of a late out-of-time pair, or single hits adjacent to a pair of hits — there was no left-right ambiguity because the side with the adjacent hit was overwhelmingly preferred over the side with no neighboring hit. For these hits then, the preferred side of the wire was recorded, and the degree of freedom allowed by the left-right ambiguity was not allowed.

Note that this algorithm included a sensible procedure for large groups of adjacent hits in the drift chambers, as shown in Figure 35. For this topology which was common in low mass electron pairs from photon conversions, the pairs with the best sums-of-distances were chosen. For every classification of chamber hit, the location at the chamber of the hit or hit pair was stored for use in the segment finding. In the subsequent tracking algorithms, no out-of-time pairs were considered in forming segment or track candidates.

Once the chamber hits were classified, they could be used as the building blocks for segments in the $\hat{x}$-view tracking. Using the list of paired hit locations from the two upstream

or the two downstream chambers, all possible segment combinations were considered. Then, for each segment candidate the status codes of the component pairs were summed, and only those segments with a summed status code of less than 3 were considered. In other words, two single hits could only form a valid segment if neither of them had a left-right ambiguity associated with them; otherwise, at least one in-time pair of hits was required.

To reduce backgrounds from segments not likely to have originated from reconstructible decays in the beam, several aperture cuts were made. For the upstream segments, tracks were required to extrapolate within the boundaries of the vacuum window, and the Magnet Anti counters. In addition, a track angle requirement was imposed such that tracks must have been consistent with originating from the beam. This was implemented by requiring that a track more than 10 cm from the beam center at the vacuum window not point upstream away from the beam with an angle greater than 1 mrad. (The beam size at the vacuum window was approximately 9 cm.) Segments in the chambers downstream of the magnet were required to extrapolate to the lead-glass array. For tracking in an electron sample where a cluster in the lead-glass was required for every track, an additional requirement on the downstream segments was imposed requiring that the track extrapolate

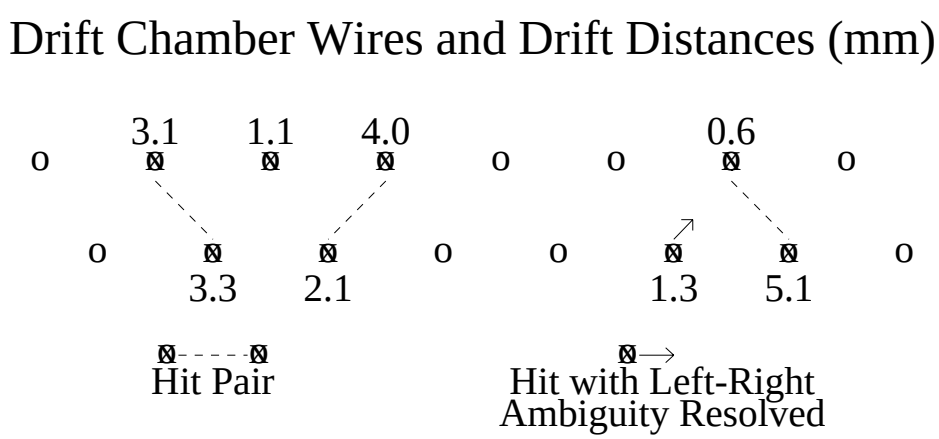

Figure 35. An example of the treatment of pair-finding with large groups of adjacent hits. In the triplet of hits, the pair with the best sum-of-distances is chosen, and the remaining hit is classified as a hit with status code 1 because of the left-right ambiguity is resolved by the presence of neighboring hits. In the quintuplet of hits, however, the remaining hit is between two pairs, so the left-right ambiguity remains.

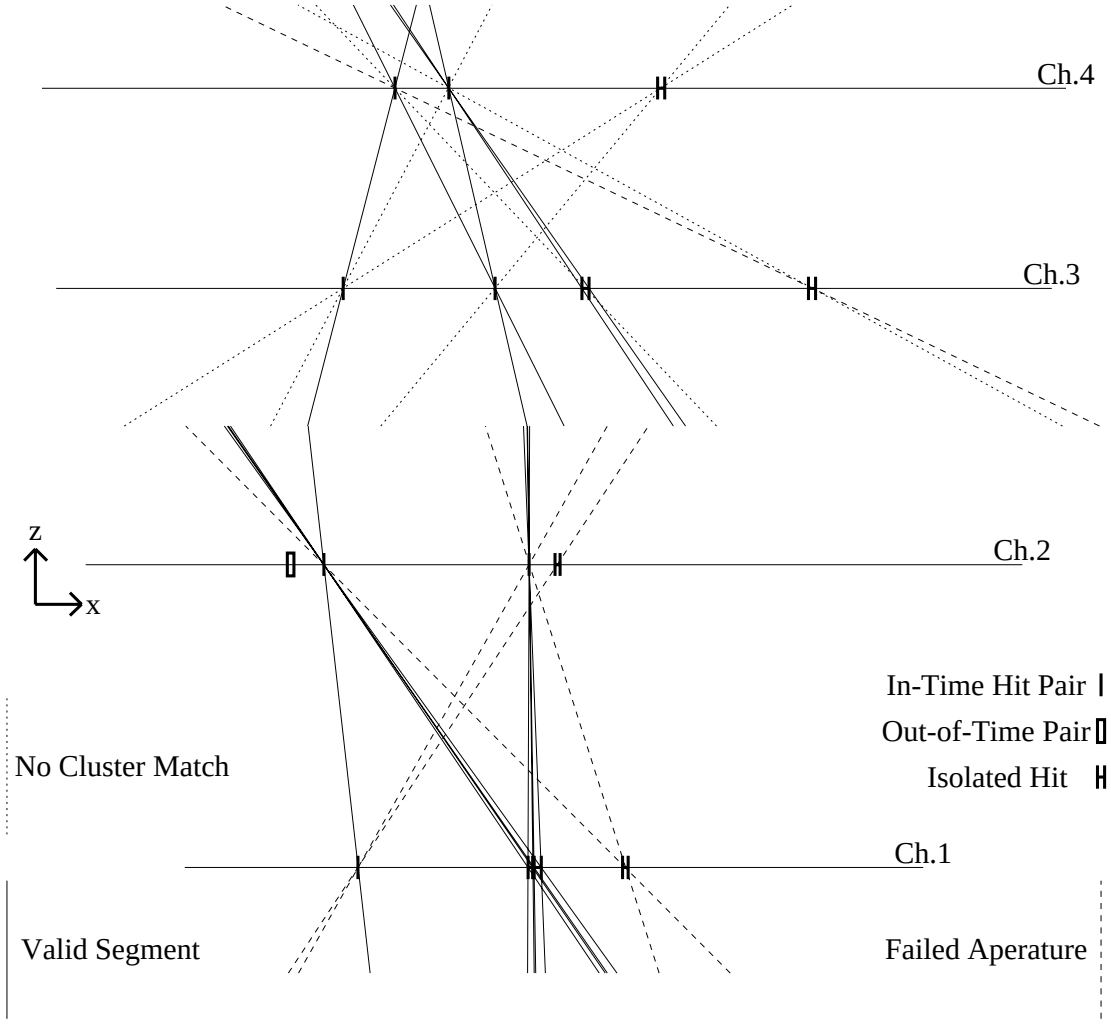

Figure 36. The track-finding algorithm for finding segment candidates is illustrated using an event from the Two-Electron trigger. Note that the segment finding produced a dizzying array of potential segments, many of which were disqualified by aperture cuts or cluster cuts. In the case of isolated single hits, two segments candidates were identified using that hit, one associated with each resolution of the left-right ambiguity. For the record, this amount of accidental activity in this event was significantly below average — typically the segment-finding produced many more candidates than in this case.

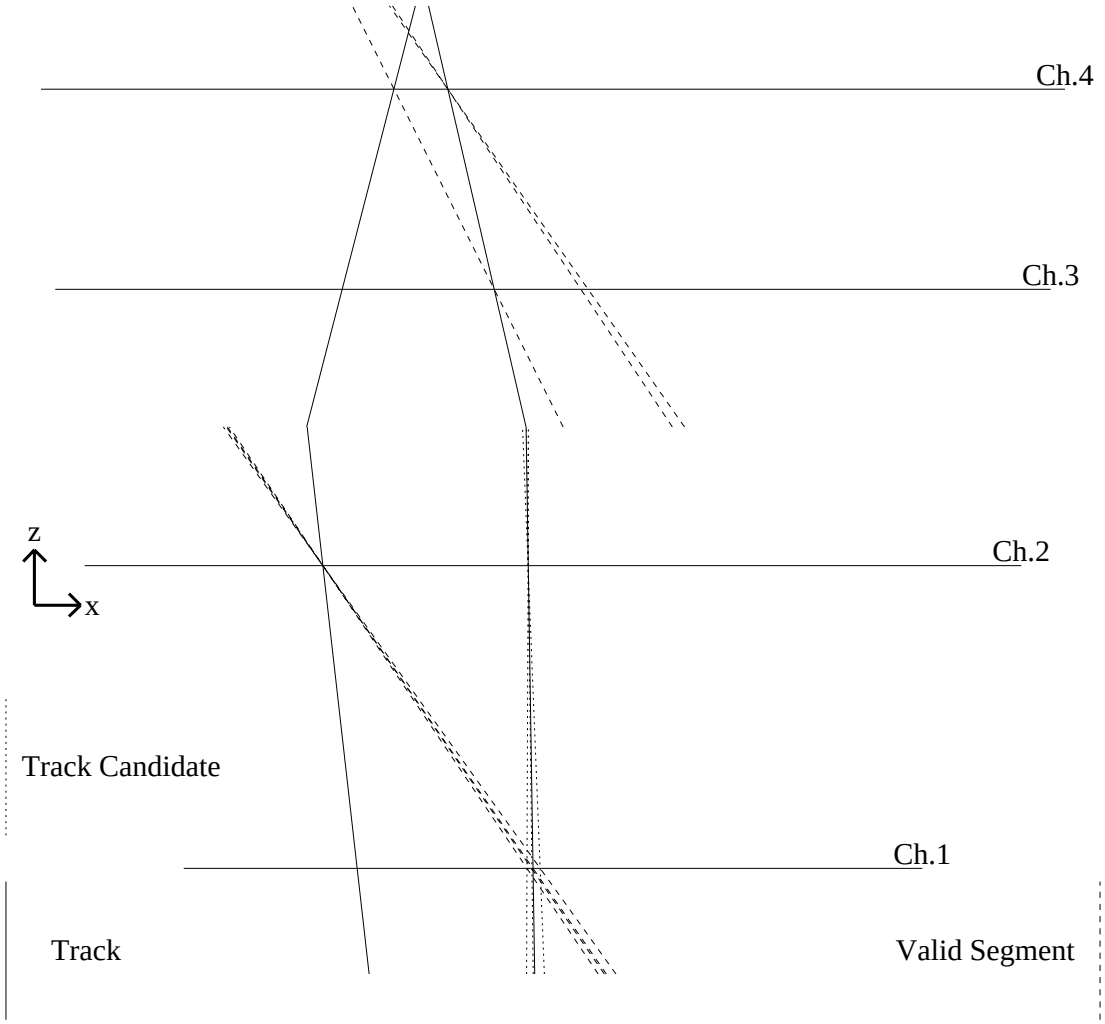

Figure 37. The track-finding algorithm for finding track candidates is illustrated. Of the segments found in Figure 36, only a small number have a sufficiently small offset at the magnet to qualify as track candidates. Of the four closely grouped track candidates sharing the hit near the center of Chamber 2, the ambiguity resolution procedure chose the candidate with the minimum offset at the magnet bend plane. In this case, the magnet offsets of the four candidates ranged from 180 μm to 6 mm.

in the $\hat{x}$-view to within 5 cm of the x-position of a cluster in the glass. This requirement greatly reduced the number of downstream segment candidates in a typical event, and helped to ensure that the real electron tracks could be found, even in an event with many accidental hits. Examples of sample hit locations in the upstream and downstream chambers along with allowable segments are shown in Figure 36.

For each allowable upstream and downstream segment candidate then, a straight-line fit was performed to the hits used in the segment candidate, and the extrapolated position at the magnet bend plane was stored, along with the sum of the status codes (see Table 17) for the chamber hit locations used to form the track. Note that for isolated single hits in the drift chamber, there were two possible segments that can be drawn through each hit because of the left-right ambiguity. In such cases, both segments were identified as possible candidates, as shown in Figure 36.

The next step was to match the upstream and downstream segments to form a complete $\hat{x}$-view track in the drift chambers. The quality variable used in this matching is the offset of the two segments at the magnet. For all successful track candidates, this offset was required to be less than a loose cut of 1 cm (10 times the resolution for a 3 GeV track).

The final step in the track-finding was the "ambiguity resolution" among all the track candidates which was required because of the propensity of track candidates to share hits in the drift chambers. At this point, the total pair status codes for the hits in each track were added to make a summed status code for the track. The minimum such summed status code among all candidates was found, and only track candidates with a track status code within 2 units of this best track candidate were considered for further analysis. For example, in an event with a perfect track — all eight hits present and a perfect sum of distances for each pair — all other track candidates must not have had more than one missing hit, or more than two single hits from low SOD pairs. Then, the remaining track candidates were sorted in order of the magnet offset distance. The lowest offset track candidate was chosen as a "track", and then all other track candidates which shared any hits with this selected track candidate were eliminated. This procedure was repeated until no track candidates remained.

There was one exception to the harsh letter of the law in this ambiguity resolution

procedure. If a selected track candidate shared hits with tracks with a magnet offset within 1 mm of its own, and if there were either two such endangered tracks with an equal summed status code, or one such track with a better summed status code, then the ambiguity resolution algorithm was re-run *forcing* this track to not be considered. If this second pass of the ambiguity resolution found a greater number of tracks, the the result of this second pass was used. This complicated embellishment to the algorithm reduced the likelihood of a single fake track which shared hits with many of the real tracks in an event causing the real tracks to not be found. There were options in the tracking to allow tracks to share hits in the first and second drift chambers without being removed by the ambiguity resolution algorithm. An example of track-finding and ambiguity resolution for the track segments of Figure 36 is shown in Figure 37. With the end of the ambiguity resolution, the $\hat{x}$-view track-finding was complete.

The $\hat{y}$-view tracking used many of the same principles, but was fundamentally different because it looked for one straight track through four chambers. The $\hat{y}$-view tracking first found a "segment" between Chambers 1 and 4. Similar aperture cuts to those described in the $\hat{x}$-view tracking above were imposed on the candidate segments, including cuts on the vacuum window and glass apertures, a looser version of the track angle cut described above, plus the cut requiring consistency with the locations of clusters. An example of this segment finding is shown in Figure 38. Once all the Chamber 1 to Chamber 4 segment candidates were identified, the next step was to search through the list of hit pairs in Chambers 2 and 3. Any hit pair or single hit located within a loose cut of 1.5 cm from the interpolated track position was considered for inclusion in the track. No $\hat{y}$-view track candidate with a summed status code greater than 4 was considered, so only a maximum of two isolated single hits were allowed in any track. For these isolated single hits, there was still a left-right ambiguity to resolve. For such events, the track was fit given each possible resolution of the ambiguity, and the track χ^2 was computed for each scenario, where

$$\chi^2 = \sum_{\text{hits}} (y_{\text{fit}} - y_{\text{hit}})^2.$$

Note that this "χ^2" was not normalized to the expected resolution. For the single-hit resolution of approximately 100 μm, the mean χ^2 should be 6 $\times 10^{-8}$ m^2. The resolution

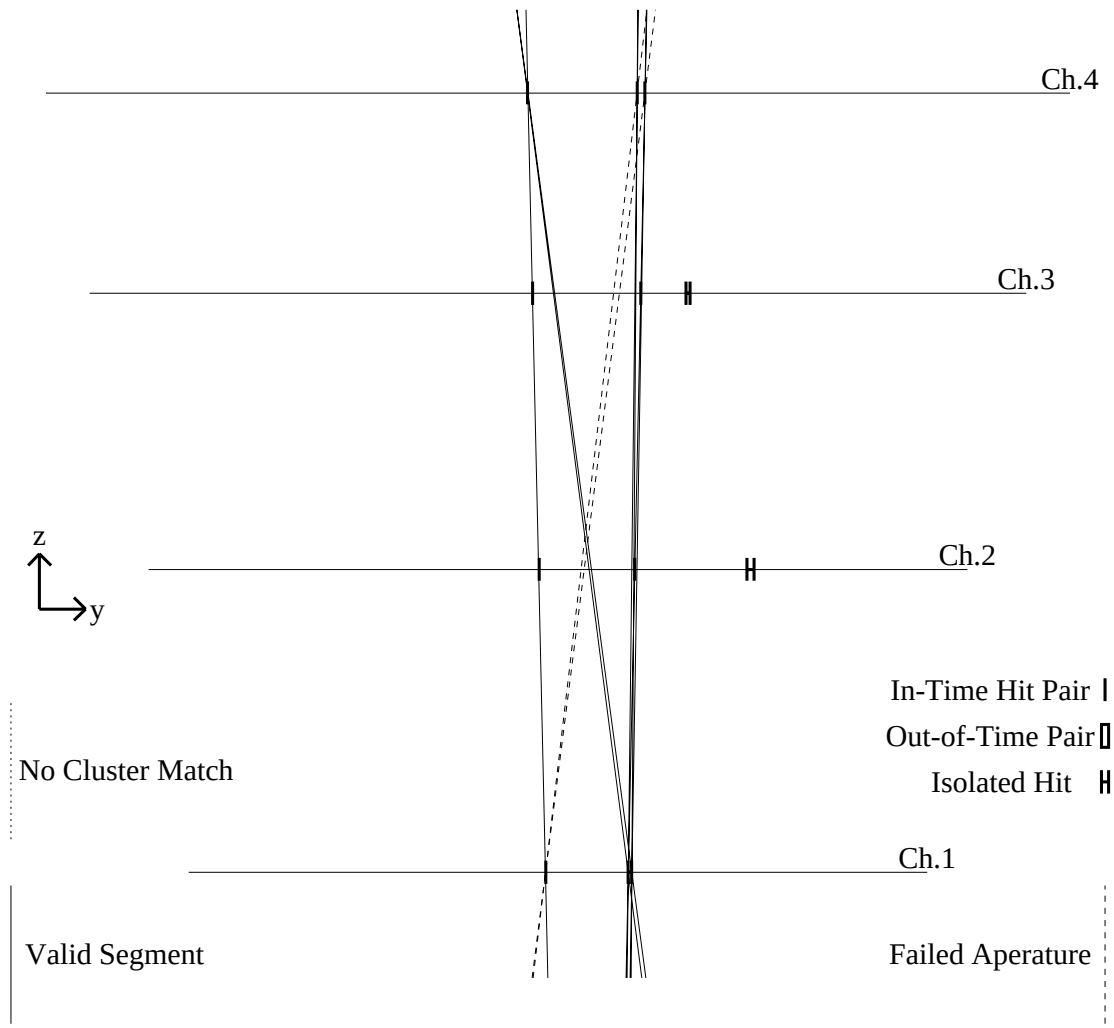

Figure 38. The track-finding algorithm for finding segment candidates is illustrated using an event from the Two-Electron trigger. All segment candidates connecting hits in Chamber 1 and 4 are shown. In the case of isolated single hit in Chamber 1, two segments candidates are identified using that hit, one associated with each resolution of the left-right ambiguity, and the resulting in the best fit χ^2 once all hits are considered will be chosen.

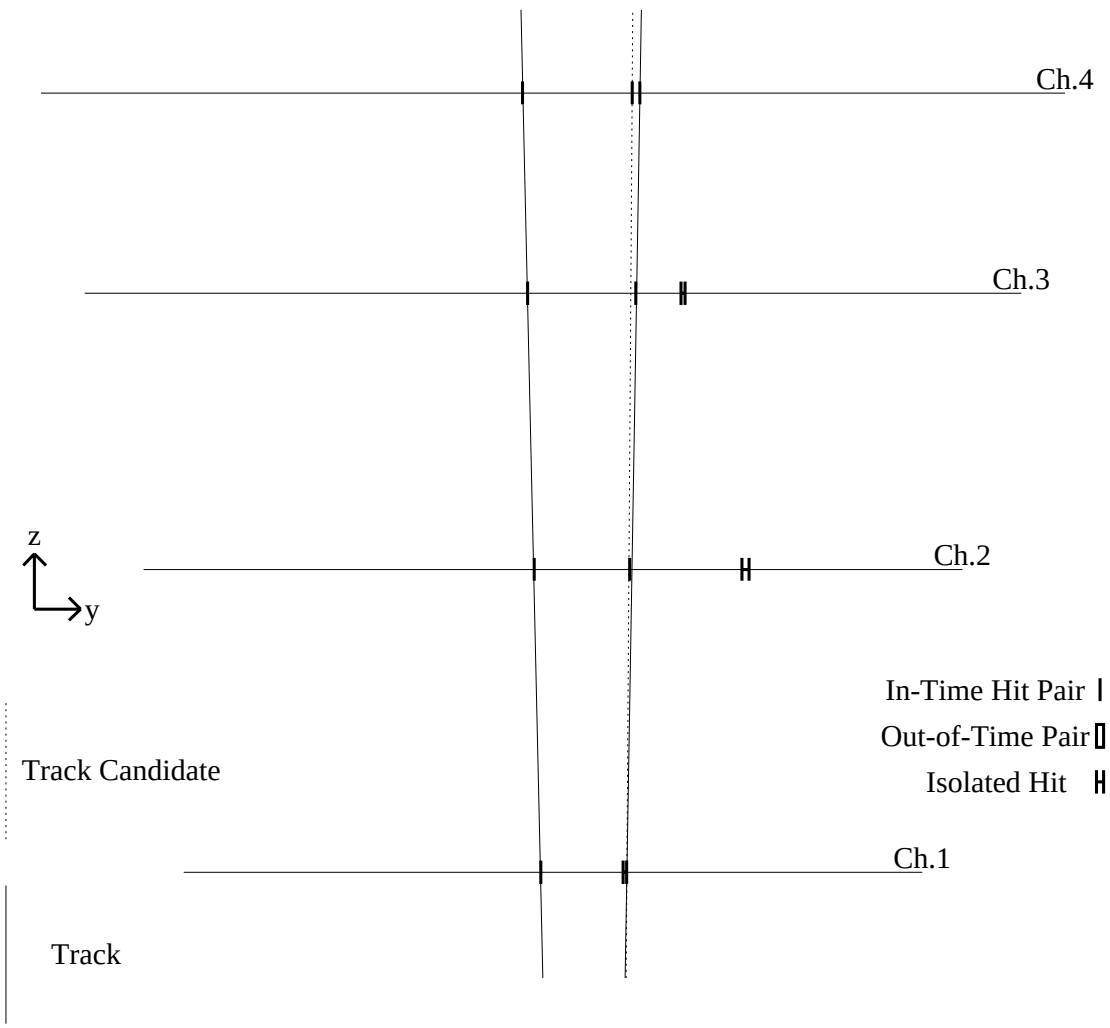

Figure 39. The track-finding algorithm for finding track candidates is illustrated. Of the segments found in Figure 38, only three make reasonable track candidates. The best resolution of the left-right ambiguity in Chamber 1 has been chosen, and the only remaining issue is which of the tracks sharing hits in Chambers 1 through 3 will be chosen. The selected track has a fit χ^2 of 6.41×10^{-8} m^2, much smaller than the 6.77×10^{-6} m^2 of its rival.

of the ambiguity or ambiguities resulting in the lowest χ^2 was chosen for the $\hat{y}$-view track candidate.

The final step in the $\hat{y}$-view tracking was the ambiguity resolution. This proceeded via a similar process to that used in the $\hat{x}$-view tracking, except that the variable used to rank track quality was no longer the offset of upstream and downstream tracks at the magnet, but the χ^2 from the fit to each track candidate. The summed status codes from all track candidates were compared, and those with summed status codes more than 2 above the minimum summed status code of the best track candidate were removed from consideration. Track candidates were then chosen to be "tracks" in order of χ^2, removing all other candidates which share hits with the selected track. This process was then repeated until no track candidates remained. An illustration of the ambiguity resolution for the segments found in Figure 38 is shown in Figure 39.

This algorithm was very efficient in finding tracks, even in the presence of accidentals. In Monte Carlo simulations which included chamber inefficiencies, the efficiency for reconstructing two tracks in each view for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, decays was 99.3% if accidental effects were not included, and 96.7% when accidental activity in the chambers was simulated.

4.3 Track-Cluster Matching

Once tracks and clusters have been found, the next step was to identify the clusters from electrons by matching tracks to clusters in the lead-glass. The optimal track-cluster matching routine varied with the number and types of final-state particles; the one described here was used in final states where the only charged tracks were two electrons.

The procedure was to consider all possible combinations of two $\hat{x}$-view tracks, two $\hat{y}$-view tracks and matching clusters, and to find the combination with the smallest distances between tracks and clusters. Because of the possibility of two partially-overlapping tracks in the $\hat{y}$-view resulting in only one track found by the tracking algorithm, the matching of a single $\hat{y}$-view track to multiple $\hat{x}$-view tracks was permitted. For all combinations of tracks and clusters, the quantity

$$D = (x^{(1)}_{\text{track}} - x^{(1)}_{\text{cluster}})^2 + (x^{(2)}_{\text{track}} - x^{(2)}_{\text{cluster}})^2 + (y^{(1)}_{\text{track}} - x^{(1)}_{\text{cluster}})^2 + (y^{(2)}_{\text{track}} - y^{(2)}_{\text{cluster}})^2$$

was calculated, and the combinations were ranked according to this distance variable. A loose cut on the maximum considered D of $(7\text{ cm})^2$ was applied. (The worst case resolution in the position measurement of a lead-glass cluster was 6 mm.) If no successful match of two tracks to two distinct clusters was found, the event was removed from the sample.

Given a matching of tracks and clusters, the full set of lead-glass energy corrections could be applied since each cluster was now identified as either an electron or a photon. With this done, the ratio of the cluster energy and track momentum for electrons, the E/p, could be calculated and compared with the expected value of 1 for electrons. In the initial data set reduction of the Two-Electron stream, a cut was made requiring $0.65 < E/p < 1.35$ for each electron candidate. If the candidates failed the cut, then the next-best matching of tracks to clusters was considered by re-doing the cluster energy corrections based on the new matching and recalculating E/p. If no acceptable match was found with $D < (7\text{ cm})^2$ which satisfied the E/p cut, the event was rejected.

With these three basic routines, the reconstruction of photons, electrons, and other charged particles can be done in the detector. This allows calibration of the detector and analysis of the data to proceed.

CHAPTER 5

CALIBRATION OF THE CHARGED SPECTROMETER AND DETECTOR ALIGNMENT

Take thou of me smooth pillows, sweetest bed;
A chamber deaf to noise and blind to light,
A rosy garland and a weary head
—*Algernon Sidney*

In this chapter, the process of calibrating the charged spectrometer and the related task of aligning the detector elements are discussed. The essential goal is two-fold. First, one wishes to obtain the best possible momentum and direction measurements from the drift chambers. This requires optimization of the single-hit resolution in the drift chambers, accurate measurement of the field strength of the magnet, and internal alignment of the chamber system. Second, this well-defined chamber coordinate system can then be expanded into a global coordinate system by measuring positions of various detector elements relative to the chamber positions. A consistent global coordinate system is important in order to correctly simulate the effects of the detector geometry, as well as to make possible extrapolation of tracks found in the drift chambers to other detector elements for energy measurement or particle identification. Clearly then, the alignment and calibration of the charged spectrometer are among the most fundamental tasks required to understand the detector, and they have direct impact on every physics measurement in this experiment.

5.1 Survey Information

The starting point for detector alignment was the information available from surveys done of the detector and beamline elements. In the directions transverse to the beam (the $\hat{x}$ and $\hat{y}$-directions), the drift chamber positions must be known *a priori* to approximately one cell-size (6.35 mm) in order to find tracks through the chambers at a crude level. Once this is accomplished, transverse alignment of the detector can proceed through track reconstruction. However, in the direction along the beam, the shallow angles of tracks through the drift chambers (typically less than 50 mrad) severely limit the sensitivity to these positions from track information, and here the surveyors' results for locations in the direction along the beam (the $\hat{z}$-direction) provided the most precise measurement. Thus, the information from these surveys is crucial to the alignment of the detector.

A survey of the key detector elements was done both before and after the E799 run, and a more comprehensive survey was performed approximately three months after the end of the run. Although one might expect some motion of the experimental apparatus to occur during or after the run, in many detector elements this motion is expected to be small if these elements are heavy and not easily moved. Furthermore, the positions of many of these elements are not crucial, and need to be measured only to about 1 mm accuracy. Ironically, the drift chambers were among both the most easily moved detector elements and also the most sensitive to small motions. Therefore, when it was necessary to move the drift chambers for detector work during the course of the run, surveyors were called in the re-position the chambers after the work was completed.

Comparing measurements of immovable detector elements from different surveys gives some measure of the accuracy of such numbers. Table 18 contains a list of such comparisons. From this, one concludes that the individual z-position measurements were accurate to approximately 1 mm. For even the most sensitive detector element, the drift chambers, this determination of the z-position was adequate. As a check that these accuracies are sufficient, consider the error introduced from this uncertainty in measuring the opening angle between two tracks, which is proportional to the two-track invariant mass for light particles like electrons. The opening angle between two tracks is given by

Table 18. Locations along the beam direction of "immovable" detector elements from different surveys in feet. The zero point is arbitrary.

Detector Element	10/91 Survey	1/92 Survey	4/92 Survey
Separator Magnet (Downstream End)	1838.423′	N/A	1838.417′
Vacuum Window	1965.670′	N/A	1965.669′
Analyzing Magnet (Upstream End of Lower Face)	1996.520′	1996.522′	1996.526′
PbG Block 665 (Upstream Face)	2037.654′	2037.652′	2037.658′

$$\theta_{12} \approx \tan\theta_{12} = \frac{r_2 - r_1}{z_2 - z_1}$$

where r_i is the track separation and z_i is the z-position at chamber i. This chamber separation $z_2 - z_1$ was known from survey information to 0.02% of itself. Thus, the surveying inaccuracies introduced a potential invariant mass scale error of 0.02%, which was negligible for measurements in E799.

5.2 Calibration of the Drift Chamber TDCs

Once the survey information has provided location along the beam direction of the drift chambers and preliminary measurements of positions along the directions transverse to the beam, one can perform track finding at the crude level of drawing "roads" through cells in the drift chambers (see Section 4.2). However, for precision tracking with the chambers, one needs to be able to exploit the precise determination of drift distances that one gets from measuring the drift time in the chambers. The drift chamber TDCs were self-calibrating so that it was known that each TDC count corresponded to 1 ns. However, in order to obtain drift distances from the outputs of the chamber TDCs, it is first necessary to know the TDC value for each wire that corresponds to zero drift distance. Then, by knowing

the relationship between the drift distances and drift times (the "$X(t)$ Relation"), one can make precise position measurements in the chambers.

There are several mechanisms that might introduce wire-to-wire variations in the TDC value corresponding to a zero drift-distance. For example, varying cable lengths, different pulse heights or shapes, and channel-to-channel or module-to-module variations in the TDCs can all give timing differences between wires. In order to remove such differences, a very simple technique was used. Because the level one trigger, which provided a timing reference for the TDCs, was correlated with activity in the drift chamber, one could use the absolute position of the timing edge in the raw TDC distribution that comes from the charged particles passing very close to a wire (Figure 40) to measure the zero drift TDC value. By comparing the position of the edge for each individual wire, a set of timing offsets for each wire were generated. This method had the advantage of not even requiring that tracks be found in the chambers since, in general, chamber activity was highly correlated with the experimental triggers. Therefore, this timing offset measurement technique could be used before any chamber alignment had been performed.

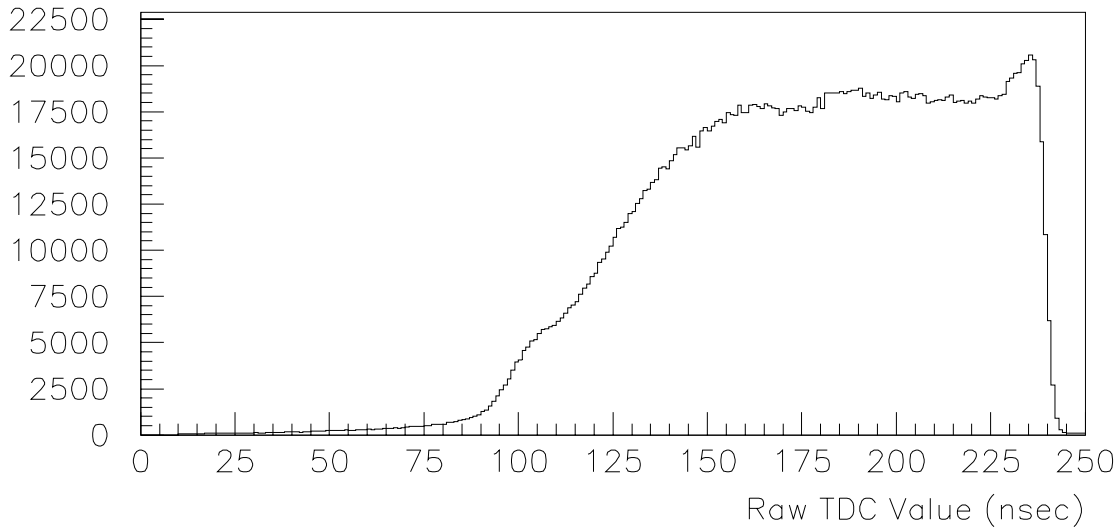

Figure 40. The raw distribution of TDC values for all wires in plane 16. Note the sharp edge from charged particles passing very close to a sense wire at high TDC values. The trigger provided a common stop for the TDCs.

The actual procedure for calculating these timing shifts was to compare the shape near this timing edge for an individual wire with the edge from the sum of the 16 wires in its cable, and similarly to compare each of the 7-9 cables in each sense plane with the overall shape for the whole plane. By sliding these two timing edge distributions past one another and looking for a minimum in the χ^2 for the two distributions to match, one could then find the correct timing offset for that wire or cable (see Figure 41). This offset was determined in 1 ns bins, and no attempt was made to interpolate this offsets to improve this measurement. The offset was determined for each wire and for each cable when there were sufficient statistics to support this technique, using data taken in the Two Electron trigger. Wire timings from two-track triggers in a previous experiment [56] were used for edge wires without sufficient statistics to determine a timing edge in E799 data. This procedure was iterated several times since a change in wire timing constants changed the shape and position of the timing edge for the overall cable or plane.

The wire timing offsets determined here were consistent with timings determined in a previous experiment [56]. Wire timings were computed using this technique at various points during the E799 run, and found to be consistent with each other, except in the case of a few wires. The time dependence of the wire timing offsets will be discussed further in Section 5.3.3. The wire timing offsets used in the analysis were determined using a large block of data from the middle portion of the E799 data set.

5.3 Translation of Drift Times to Distances

Once any wire timing offsets had been removed, the chamber was treated as a unit for the purposes of relating drift times to drift distances (the "$X(t)$ relation") in the chamber. This was a valid approximation because parameters which affect this relation, such as the cell geometry, voltages and gas mixture, were identical for each sense wire in the chamber. Because the $X(t)$ relation was treated as independent of location in the chamber, it is possible to determine this relation without using tracking. Any sample of real hits in the drift chambers from charged particles in-time with the trigger will be approximately uniform in distance from the nearest sense wire, so one can simply use the distribution of drift times

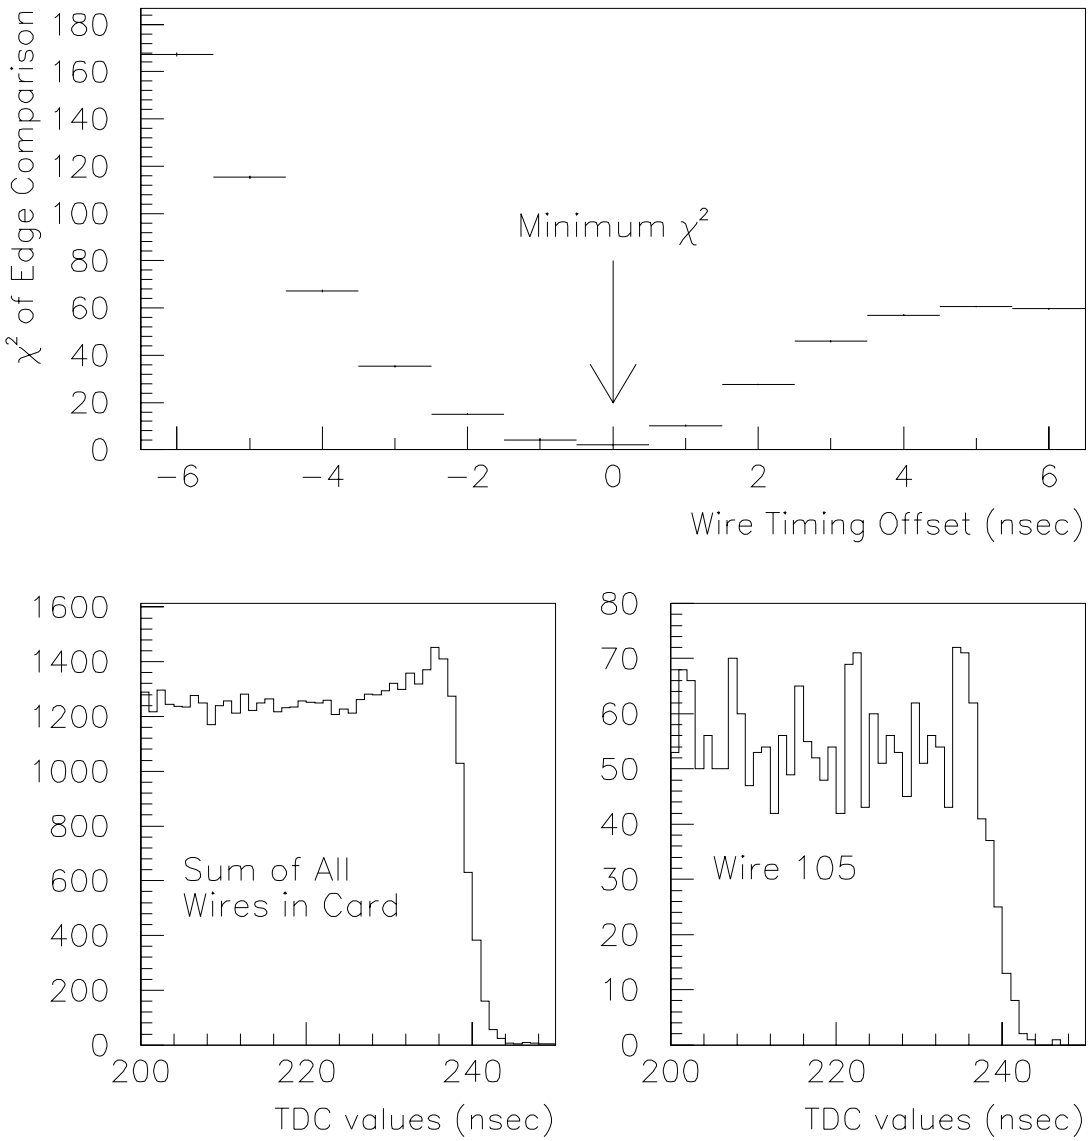

Figure 41. The determination of the wire timing offset wire 105 of plane 16. The χ^2 for each timing shift is shown, along with the timing edge for the card containing this wire and the wire itself. The two timing distributions shown above were compared to produce the plot of χ^2 vs. timing shift.

in this sample to extract the $X(t)$ relation. However, in practice it was necessary to use hits on tracks for the $X(t)$ determination in order to avoid backgrounds from out-of-time hits and from chamber noise.

The sample used for this measurement were events that passed the Two Electron crunch. This sample has several checks built in to ensure that the tracks in the drift chamber were from the same bucket as the trigger. In particular, tracks were required to point to HCF-seeded clusters in the lead-glass with a good E/p in order to pass electron identification cuts. Cuts were also placed on the TDC values of the B-bank and C-bank triggers in order to eliminate any residual out-of-time tracks from the sample. Approximately 50000 events passing all cuts were used to generate the $X(t)$ relation in each run.

Once events were selected for this sample, a distribution of drift times was collected for each plane of the drift chambers. Since there is some correlation between wire efficiency and hit location, a correction had to be made for this effect in order not to bias the drift time distribution. Therefore, when a track did not leave a hit in a given plane, the drift time of the hit was recorded in the complimentary plane of the same chamber in the same view. When another event was found with that same TDC value in the complimentary plane and a hit in the original plane, the drift time in the original plane was stored in a histogram as the reconstructed distribution of drift times for missing hits. This histogram of drift times for the missing hits was then added to the histogram of raw drift times to produce a corrected distribution of drift times (Figure 42). This distribution was then integrated, and by assuming uniform illumination across the cells, one could then assign each TDC value to a location within the cell (Figure 43). Explicitly, the position as a function of drift time is calculated to be

$$X(t) = L_c \frac{\int_0^t N(t')dt'}{\int_0^{t_{max}} N(t')dt'}$$

where L_c is the cell size and $N(t')$ is the number of events with at drift time t'. This procedure was iterated a few times since the track-finding algorithm depends weakly on the $X(t)$ tables implemented in the analysis code.

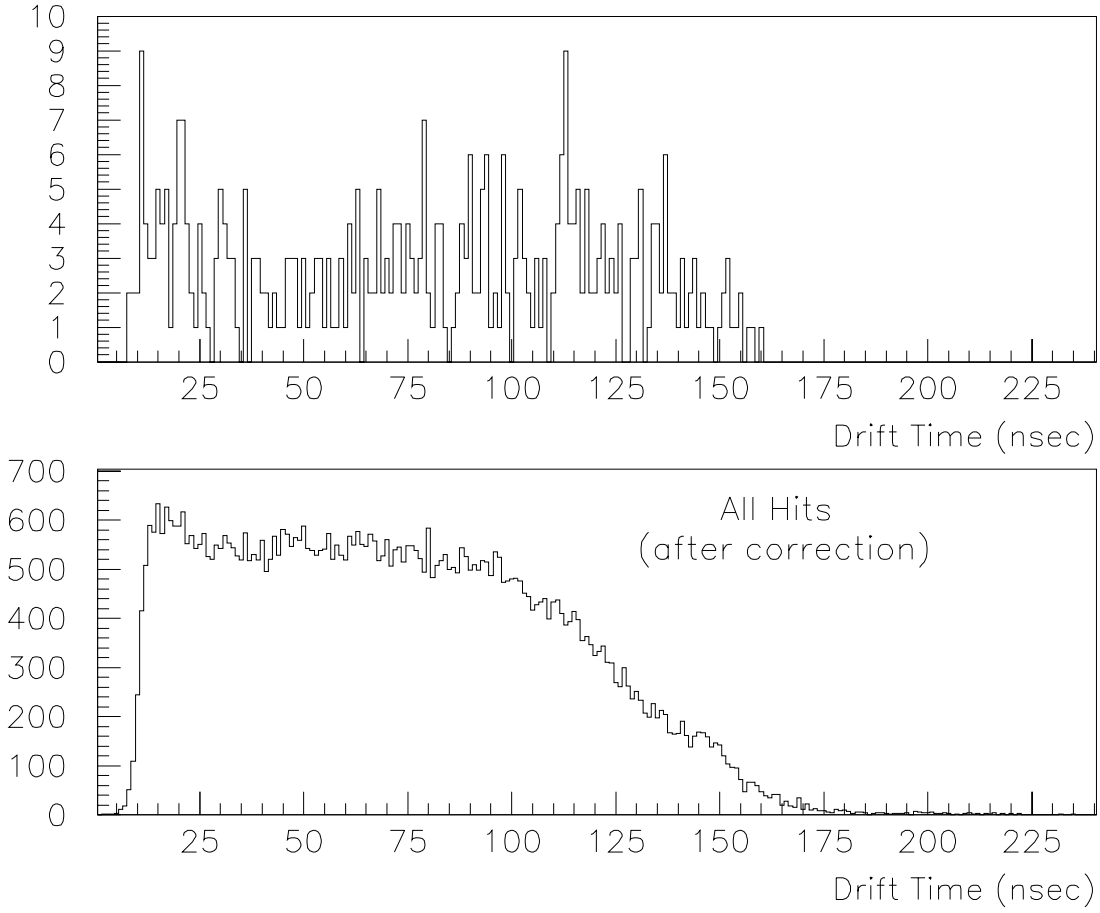

Figure 42. The missing hit correction applied to the raw drift time distribution for the $X(t)$ Determination. Shown is the raw distribution of missing hit "drift times" along with the total corrected distribution of drift times for the downstream plane of Chamber 4Y.

5.3.1 Variation of $X(t)$ Relation with Time

Since it was unknown how the $X(t)$ relation would vary with time, a separate $X(t)$ file was made from data for each run with sufficient statistics, resulting in about 120 separate $X(t)$ tables. It is interesting to study how much the $X(t)$ tables varied from run-to-run, both as a measure of the precision of the determination of the $X(t)$ relation, and as a study of how often $X(t)$ tables need to be generated for future experiments using these chambers. To measure the similarity between two $X(t)$ tables, define

$$\sigma_x = \frac{1}{\sqrt{2}} \left\langle (X_1(t) - X_2(t))^2 \right\rangle ,$$

where $X_1(t)$ and $X_2(t)$ are two $X(t)$ relations, and the average is taken over drift distances. If these two $X(t)$ relations are ideally equal, then σ_x measures the average error in the measurement of x from the statistical and systematic errors in determining the $X(t)$ relation. The contribution of statistical error in the $X(t)$ measurement to σ_x was about 10 μm.

If one believes that the $X(t)$ relation should not vary with time, then the σ_x measured from two runs adjacent in time should be the same as the σ_x between any two runs from the data set, on average. Table 19 shows a test of this hypothesis. In this table, the average σ_x from comparing $X(t)$ measurements from adjacent runs and the average σ_x from comparing

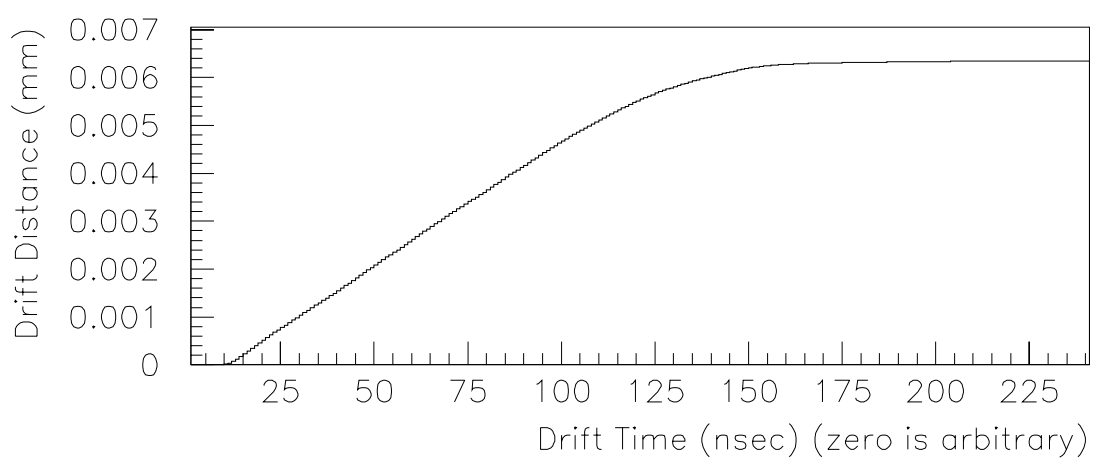

Figure 43. A typical $X(t)$ relation, determined from the drift-time distribution in Figure 42.

each $X(t)$ measurement to a single measurement made from data in the middle of the run are tabulated for each sense plane. One concludes from this comparison that the $X(t)$ relation does vary with time, since the deviations were so much smaller when comparing runs adjacent in time.

It is interesting to note that σ_x between adjacent $X(t)$ relations has large jumps with $\sigma_x \gtrsim 80$ μm at several points in the run. However, in no case could a sudden change in the $X(t)$ relation be associated with any changes in gas mixture, voltages or condition of the chambers as noted in records kept during the run. Records were not kept of temperature and atmospheric pressure variations which may have caused some of this effect. Given the data available to us, one concludes that the run-by-run $X(t)$ relation measurement was necessary to maintain chamber resolution since sudden jumps in the $X(t)$ relation occur that could not be correlated with any recorded change in conditions. This point will be discussed further in Section 5.8.

5.3.2 Evaluation of Chamber Resolution

After determining the $X(t)$ relation, a number of corrections were made on an event-by-event basis to improve the chamber position measurements. The drift time was corrected for propagation time along the sense wire (6 ns along the whole length of the longest sense wires), and a correction was made for sag of wires in the y-view which is as large as 60 μm in the center of the largest chamber. At this point also, a correction was made for a small offset in time between the arrival of decay products from our beam and the RF signal provided by the accelerator which was used to strobe our trigger. Given the chamber resolution and drift velocity of 50 μm/ns, each pair of hits in complimentary chambers provides a better than 3 ns measurement of this timing shift. A perfect two-track event provides 16 independent measurements of this quantity, so typically this correction can be determined to better than 1 ns.

With each iteration of $X(t)$ determination, progress was monitored by plotting the sum of the distances of the two measurements in a single chamber-view of a track. A typical sum-of-distance distributions is shown Figure 44. The low-side tail is due to δ-rays produced when the tracks passed through material inside or outside the chamber which created a

Table 19. A study of run-to-run $X(t)$ variations

Chamber	View	Upstream or Downstream Plane	σ_x (μm) from Adjacent Runs	σ_x (μm) from a Fixed Run
1	X	U	41	70
1	X	D	43	79
1	Y	U	38	60
1	Y	D	42	64
2	X	U	29	67
2	X	D	31	49
2	Y	U	29	61
2	Y	D	29	54
3	X	U	26	50
3	X	D	28	44
3	Y	U	26	53
3	Y	D	27	48
4	X	U	31	59
4	X	D	26	47
4	Y	U	27	59
4	Y	D	29	54

second track through the chambers. This had the effect of lowering the measured sum-of-distances since each wire first observed a signal from the closer of the two tracks to that wire. This sum-of-distance distribution could be fitted to a Gaussian to obtain a measurement of the single-hit chamber resolution. The resolution of each pair of planes in the chamber was measured on a run-by-run basis, and these resolutions were used as input to the Monte Carlo. These resolutions varied between 90 μm and 120 μm and were largely constant with time. The only observed significant variations corresponded to problems with chamber pre-amplifier power supplies and changes in chamber readout thresholds. Chambers 1 and 4 showed, on average, a resolution of approximately 110 μm, while chambers 2 and 3 averaged 105 μm and 95 μm, respectively. It is important to note that the correction described above for the difference between the event time and the strobe from the accelerator RF signal had the effect of artificially narrowing this sum-of-distance distribution since the sum-of-distance itself is used in the correction. However, because the many sum-of-distance measurements in a given event were used to determine this correction, this narrowing effect was only a few percent, and was removed by hand before quoting the resolutions given above.

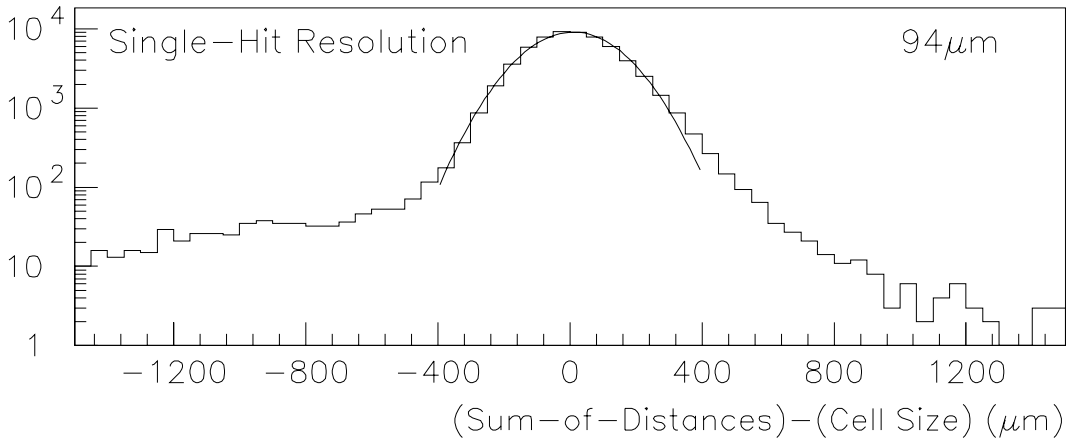

Figure 44. The sum-of-distances distribution in Chamber 4, Y-view

5.3.3 Time Dependence in TDC Calibration

When examining the sum-of-distances distribution for each individual cell in the drift chamber, some significant shifts in the mean as a function of time were observed. In particular, it was often the case that a large shift was observed in two adjacent cells, suggesting that the shift was caused by an offset in the timing for the wire shared by those two cells. Therefore, in the cases where 3σ shifts in the mean sum-of-distance of adjacent cells was observed, a correction was applied to the wire timing offset. Although these corrections to the timing offsets were only implemented based on data from a single run, the offsets were checked using a lower statistics sample from each run. One amusing consequence of this procedure was that six wires were observed to have sudden jumps in the timing offset, like the one shown in Figure 45. One wire (wire 26 in the downstream plane of chamber 1, the $\hat{y}$-view) was observed to cause a double-peaked sum of distance distribution in both of its cells, and hits on this wire were not used in the analysis.

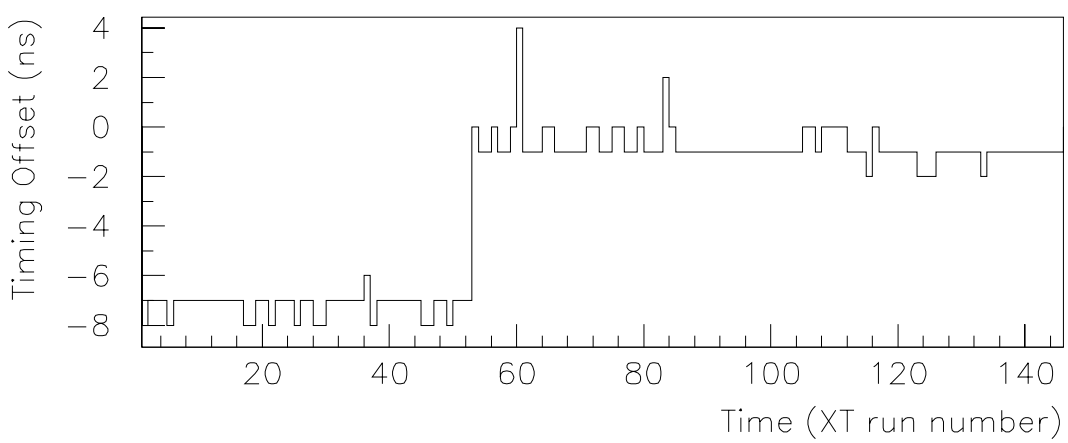

Figure 45. The timing offset for wire 56 of the downstream $\hat{y}$-view plane in chamber 1 as a function of time.

5.4 Internal Chamber Alignment

With both the wire timings and the $X(t)$ relation determined, precision position measurements could be made using the chambers, and one could then begin the task of aligning the chamber system. Two data samples were used to do this alignment. First, approximately once per day during the run, a special run of 10^5 Single-Muon triggers was taken with the analysis magnet off. Using these straight muon tracks, one could determine the offsets in $\hat{x}$ and $\hat{y}$ and rotations in the $\hat{x}$-$\hat{y}$ plane of the inner two drift chambers relative to the positions of the outer two chambers. As a second alignment procedure, a sample of $K_L \rightarrow \pi^+\pi^-\pi^0$ decays from the Two Electron trigger was used in order to find the relative rotation between chamber 1 and chamber 4 in the $\hat{x}$-$\hat{y}$ plane. Such a rotation can not be observed with single muon tracks. However, in events with two tracks originating from a common vertex, this rotation causes an apparent separation of tracks at the vertex. By measuring this separation, one can infer the existence of and correct for this rotation.

Displacements in the $\hat{z}$ direction and rotations in the $\hat{x}$-$\hat{z}$ and $\hat{y}$-$\hat{z}$ planes were not corrected for in the alignment procedures. As was discussed above, the measurements of the chamber z-positions were done through detector surveys, and were accurate to approximately 1 mm. Based on the z-positions of all four corners of the chambers as measured in surveys, any rotations in the $\hat{x}$-$\hat{z}$ and $\hat{y}$-$\hat{z}$ planes were determined to be less than 1 mrad, and were not important in the analysis.

5.4.1 Alignment with Straight-Through Tracks

In order to perform the alignment with straight-through muons, tracks are found using the standard tracking algorithms, but only considering chambers 1 and 4. Cuts are made on the lead-glass adders and the B-bank, C-bank and MU2 bank TDCs to ensure that the candidate muon passed through the detector in the same RF bucket in which the trigger was formed. Typically a sample of 15000 muons passing all cuts was used for the alignment. With each muon in this sample, the muon track is interpolated to each chamber plane in chambers 2 and 3. The offset between the interpolated position and the measured position at the chamber is then calculated, and stored in bins of x- and y-position at the chamber.

Gaussian fits are made for the offsets in each bin, and the mean offset is plotted versus the x- and y-positions of each bin. The resulting graphs of offset versus chamber positions are then fitted to lines, where an overall offset indicates a translation of the chamber, and a slope corresponds to a rotation of the chamber. Sample fits are shown in Figure 46. Typically, about 15000 reconstructed straight-through muons were sufficient to obtain 1 μm statistical errors on the fits.

In principle, this procedure could be performed independently on each sense plane of

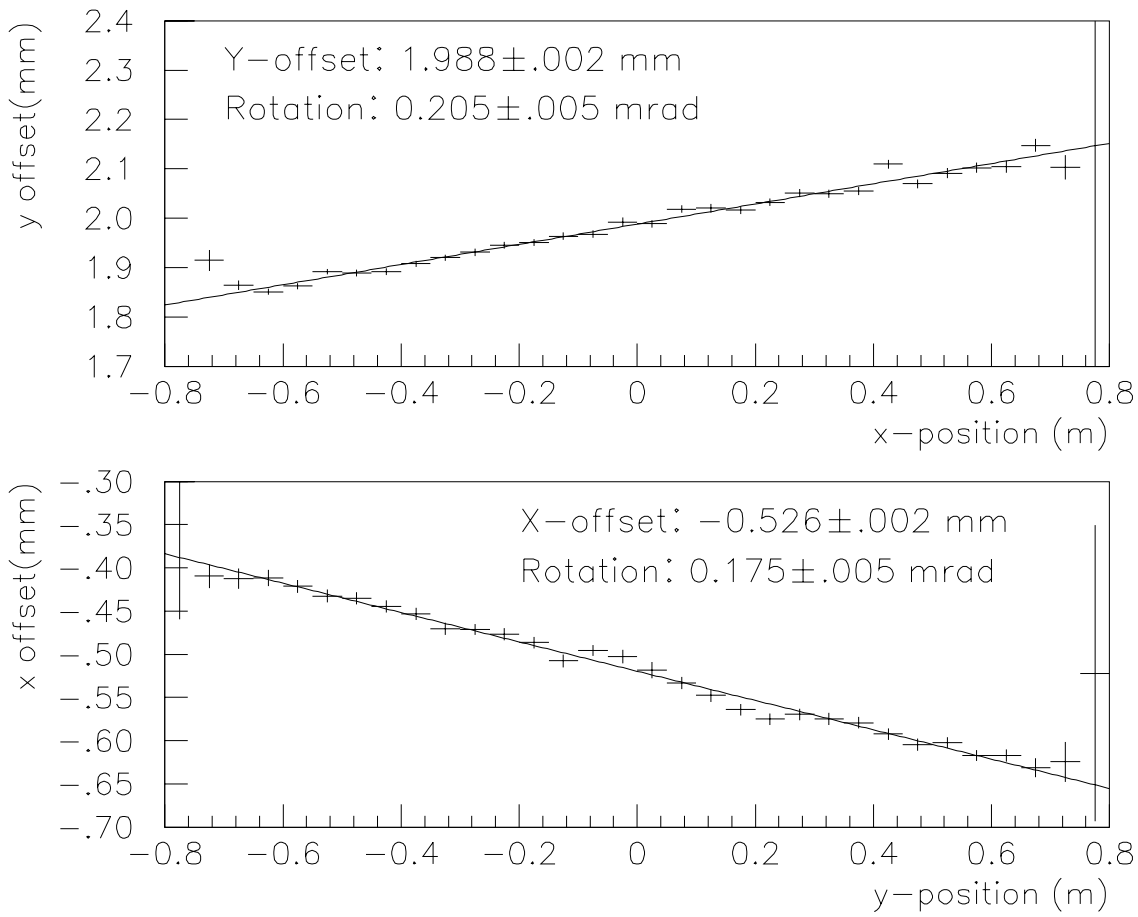

Figure 46. Chamber alignment fits using muon data. Shown is the x-offset versus the y-position and the y-offset versus the x-position. The mean offset from each fit corresponds to a chamber offset, and the slope corresponds to a rotation of the chambers in the $\hat{x}$-$\hat{y}$ plane.

chambers 2 and 3; however, it was observed in the data that all planes in a single drift chamber tended to move together in time as one coherent body. Therefore, only one set of relative offsets and rotations among the four planes in each chamber needed to be determined. To determine the relative offsets and rotations among the planes in chambers 1 and 4, the same muon data was re-analyzed, but this time using tracking information only from chambers 2 and 3 and extrapolating the tracks to chambers 1 and 4. The information extracted from the data was found to be consistent with that measured from the chambers in the lab.

From the set of fits generated in this manner, the offsets in the $\hat{x}$ and $\hat{y}$ directions and the $\hat{x}$-$\hat{y}$ plane rotations for chambers 2 and 3 were extracted. Along with the relative offsets and rotations of individual planes within each chamber, this information is used as a correction in the tracking analysis.

5.4.2 Corkscrew Rotations

There is a further degree of freedom in the internal chamber system that can be removed. While the above procedure does fix the internal rotations of the chamber coordinate system, it does not remove rotations of the four chambers in the $\hat{x}$-$\hat{y}$ plane where the angle of rotations between any two chambers is proportional to the distance between them (thus the name "corkscrew" rotations). The physical effect of this remaining rotation that one can observe in the data is shown in Figure 47. Because there still exists a rotation between chambers 1 and 2, the tracks will not meet at the decay vertex, and the separation between the tracks can be shown to be

$$s\Delta\theta = r_1 r_2 \Delta\phi,$$

where s is the track separation at the vertex, $\Delta\theta$ is the opening angle between the two tracks, r_n the track separation at chamber n, and $\Delta\phi$ the corkscrew rotation between chambers 1 and 2 (measured in radians per meter along the $\hat{z}$-direction). In order to measure this rotation, the vertex separation of a sample of two-track events, $K_L \to \pi^+\pi^-\pi^0$ decays, taken as closely as possible to each muon straight-through alignment run, is studied as a function of $r_1 r_2$. The vertex separation, scaled by $\Delta\theta$, is plotted versus $r_1 r_2$, and a linear

fit is performed to extract the slope, $\Delta\phi$ (Figure 48). A sample of approximately 5000 $K_L \to \pi^+\pi^-\pi^0$ decays was used for each fit.

5.4.3 Quality of Internal Chamber Alignment

Because all four planes in a single drift chamber move as a coherent unit with the frame of the chamber, it is easy to measure the quality of the measurement of each offset and

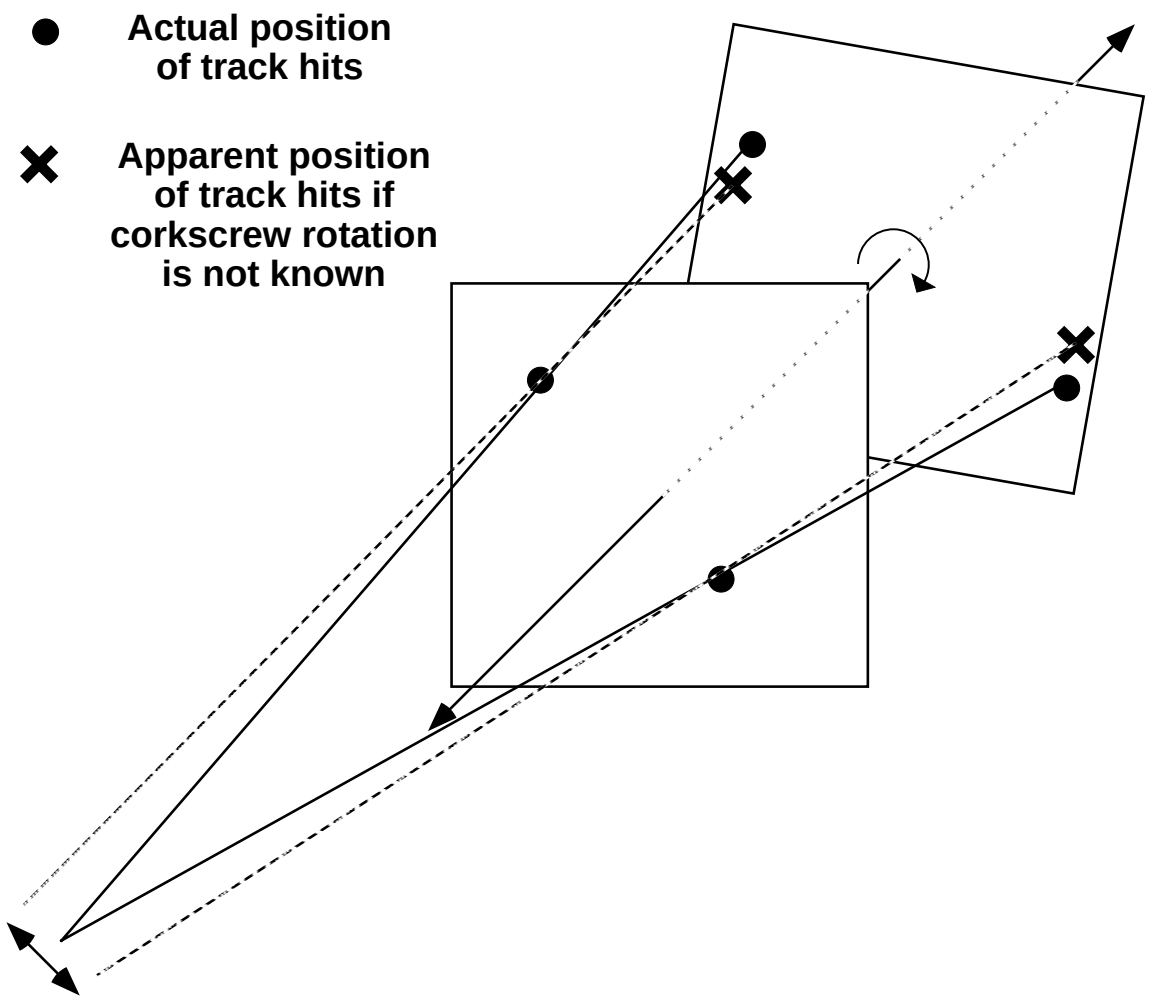

Figure 47. Illustration of the effect of corkscrew rotations. The mismeasurement of the tracks at the second chamber caused by the uncorrected rotation results in a displacement of tracks at the decay vertex.

rotation by comparing independent measurements of the same quantity in different planes. For the chamber offsets, comparing measurements of the two complimentary planes in each chamber-view typically gives an uncertainty of only 1 μm on offsets. Similar comparisons give an error on the determination of the chamber rotation in the $\hat{x}$-$\hat{y}$ plane of 5 μrad. In the case of the corkscrew rotations, the statistical error of the fit result itself is quite large, approximately 6 μrad/m, and probably is an accurate representation of the error in measuring this rotation.

5.4.4 Changes in Chamber Alignment with Time

Although these quantities are measured approximately once per day, they are clearly continuously time-varying quantities. One contribution, then, to the error in determining these rotations and offsets is how rapidly they vary in between alignment runs. The time history of the alignment offsets and rotation of Chamber 2 relative to Chambers 1 and 4 is shown in Figure 49. After having removed the large jumps in alignment offsets and rotations which occurred when chamber 4 was repositioned after work on the TRDs, one could study the

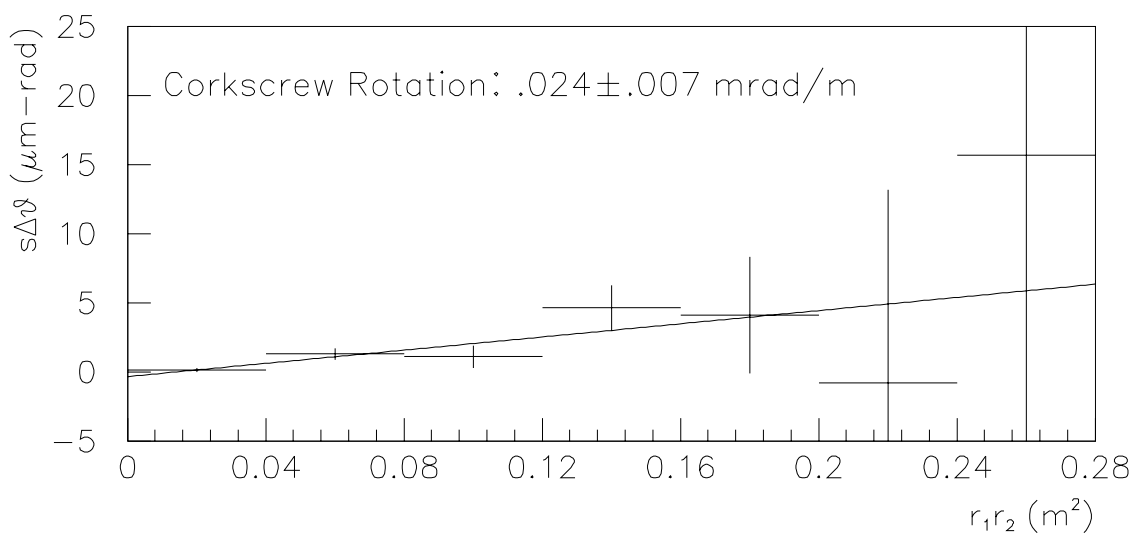

Figure 48. Linear fit to measure the corkscrew rotation. The slope corresponds to the magnitude of the relative rotation between chambers 1 and 2 in radians/meter of chamber separation in the $\hat{z}$-direction.

run-to-run changes in alignment constants to measure the rate at which they change (see Figure 5.4.4). Typically, the chamber offsets changed by 20 μm per day (RMS), and this effect contributed to the effective chamber resolution observed in the experiment. The effects of the 7 μrad per day rotation drift were much smaller. One could also do a similar analysis of the time-dependence of corkscrew rotations, and an RMS drift rate of 8 μrad/m/day was found, comparable with the statistical error on a single measurement (Figure 51).

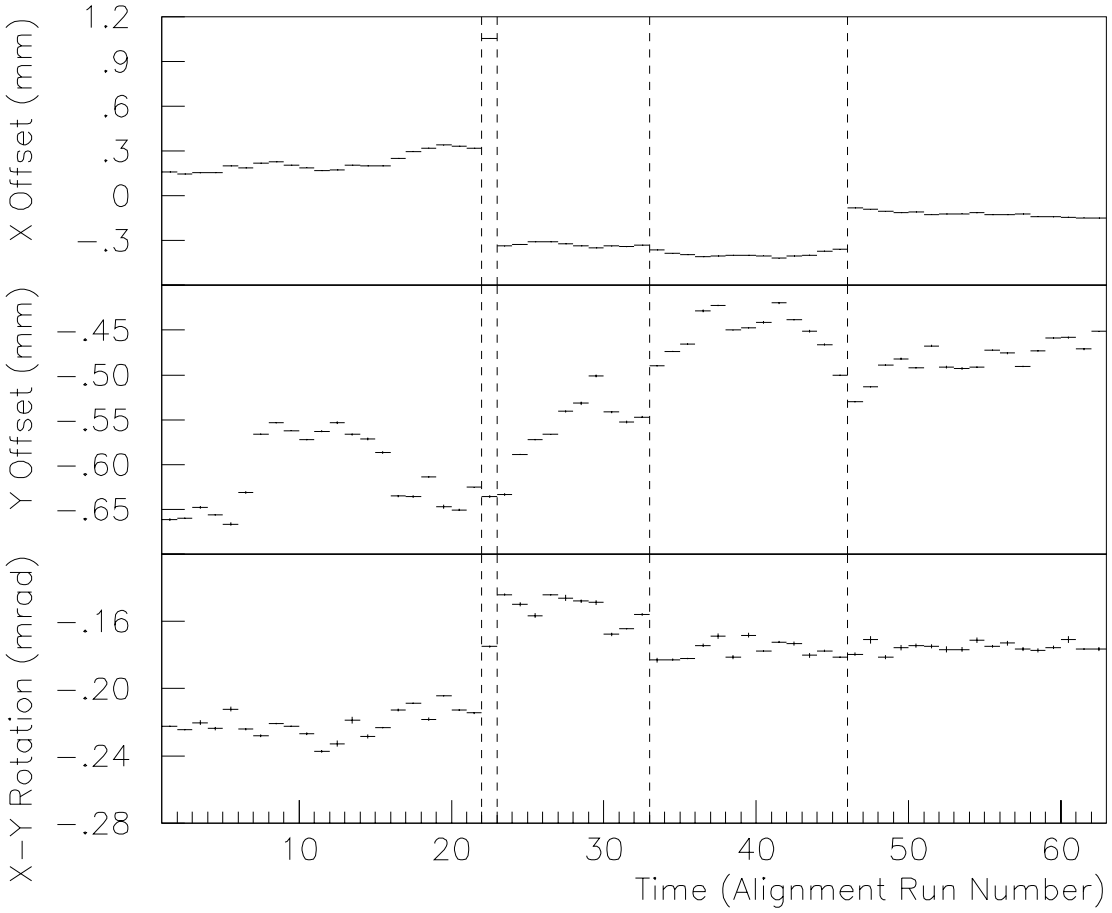

Figure 49. The time dependence of the Chamber 2 alignment constants. The positions and rotation shown are relative to Chambers 1 and 4. The dotted lines indicate times when chamber 4 was moved and subsequently reset into position by surveyors.

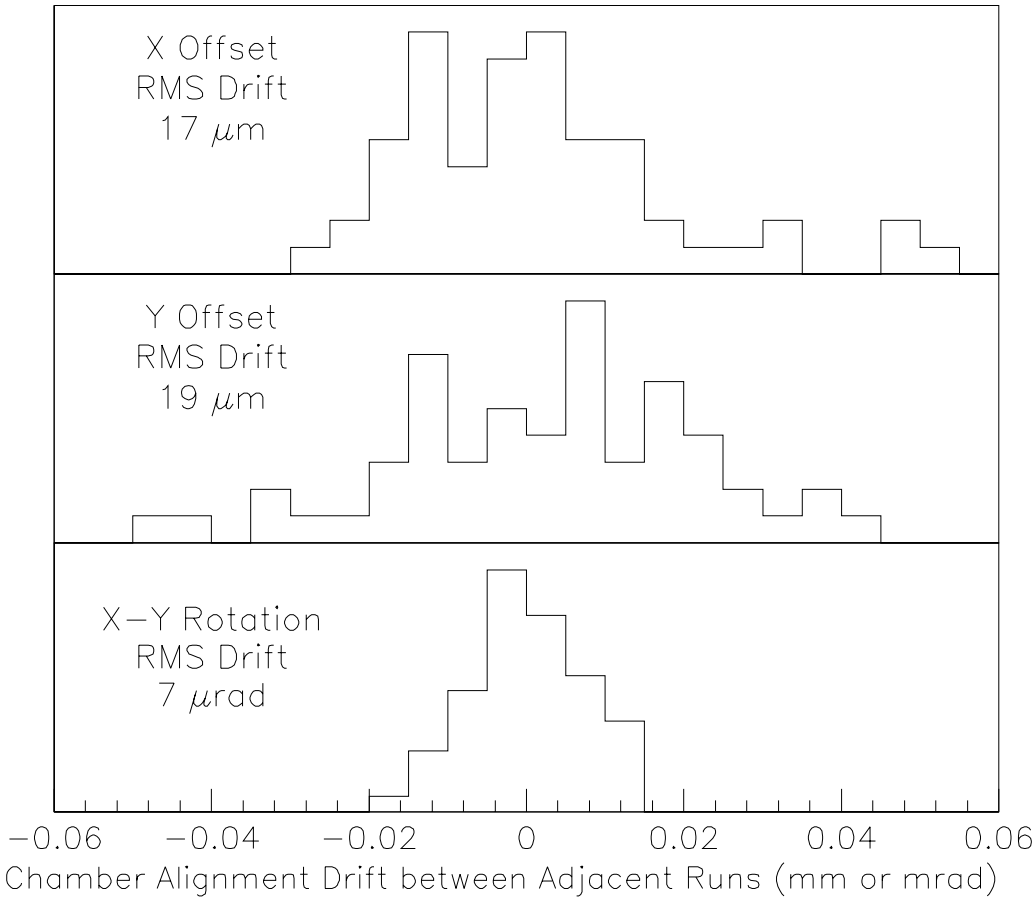

Figure 50. The change in the alignment constants of Figure 49 between adjacent runs. The large changes due to chamber repositioning have been removed. The remaining changes were most likely due to motion of the chambers. From these plots, one measures a second "drift velocity" of about 15 μm/day[a] for each drift chamber!

[a] This RMS velocity is astonishingly small. If the chambers moved at a constant velocity, rather than in a random walk, such a drift velocity would be five times smaller than the typical growth rate of toenails in 50-year-old, male professors at the University of Chicago [63].

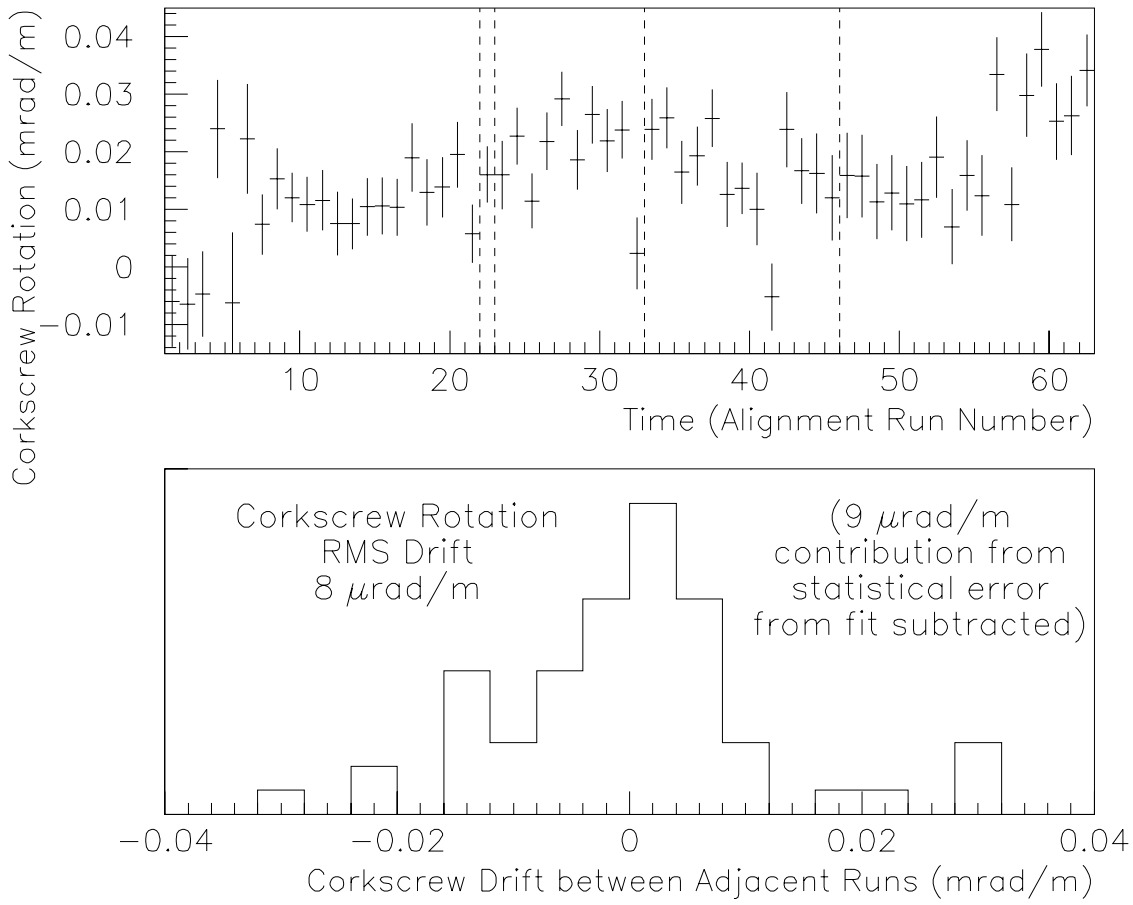

Figure 51. The time dependence of the measured corkscrew rotation. The change in corkscrew rotation between adjacent runs is also shown.

5.5 Transverse Alignment of The Detector

Once the internal alignment of the drift chambers was complete, the task of locating other detector elements in the now well-defined chamber coordinate system could be done. Two points of special interest for the global coordinate system were the target, which defined the origin of the coordinate system, and the lead-glass, which was centered around the $\hat{z}$-axis and whose array was on a grid parallel to the $\hat{x}$- and $\hat{y}$-axes, by definition. In effect then, global offsets and rotations must be applied to the aligned chambers in order to force the target and lead-glass to be at their defined positions. As with the corkscrew rotations, fits to the target and glass positions were performed using data taken near each alignment straight-through muon run, in order to correct for the time-varying chamber motion.

Once this was accomplished, the global-coordinate system was set, and one could begin the process of measuring the positions of other detector elements. Particularly important to align correctly were the scintillator banks used in the trigger, the photon vetoes which served as detector apertures, other pieces of inactive material which acted to limit detector acceptance and finally the collimators which shaped the neutral kaon beam.

5.5.1 Target

The most straightforward detector element to align was the target itself. In principle, since all particles in the neutral beam originate from the target, it is trivial to find. One can fully reconstruct any decay of a particle in the beam, and the momentum vector of the parent particle will point back to the target. In practice, however, certain decays lend themselves to this procedure more easily than others. A particularly good source of decays for this measurement comes from the copious $\Lambda \rightarrow p\pi^-$ decays found in our beam (see Figure 52). Because the Λ has a relatively short life-time ($c\tau \approx 7.89$ cm), the requirement that the decay products be observable in our detector selected particularly high-momentum Λ's (typically $\gtrsim$200 GeV), and the protons in particular typically had very high momentum. This reduced the multiple-scattering contributions to the error in finding the target location on a given event, so typically only about 2500 reconstructed $\Lambda \rightarrow p\pi^-$ decays were used to obtain a target position with a statistical error of about 50 μm. As was the case with chamber

offsets above however, there was significant time variation of this measured target position. The RMS change of the measured target position between adjacent alignment runs was approximately 400 μm per view.

The figure of merit used to evaluate the error on the target position is the effect of this uncertainty on the calculation of p_t^2 of an event. If the target is offset by Δx in the $\hat{x}$-direction and Δy in the $\hat{y}$-direction, the apparent transverse momentum of this event will be

$$|\vec{p_t}| = |\vec{p}| \frac{\sqrt{(\Delta x)^2 + (\Delta y)^2}}{z}$$

where $\vec{p}$ is the parent particle momentum, and z is the distance from the decay vertex to the target. For the highest momentum and most upstream observed kaon decays, 250 GeV and 100 meters from the target, the amount of effective $|\vec{p_t}|$ introduced by uncertainties in target position is less than 2 MeV, which is much less than errors resulting from multiple scattering or lead-glass resolution.

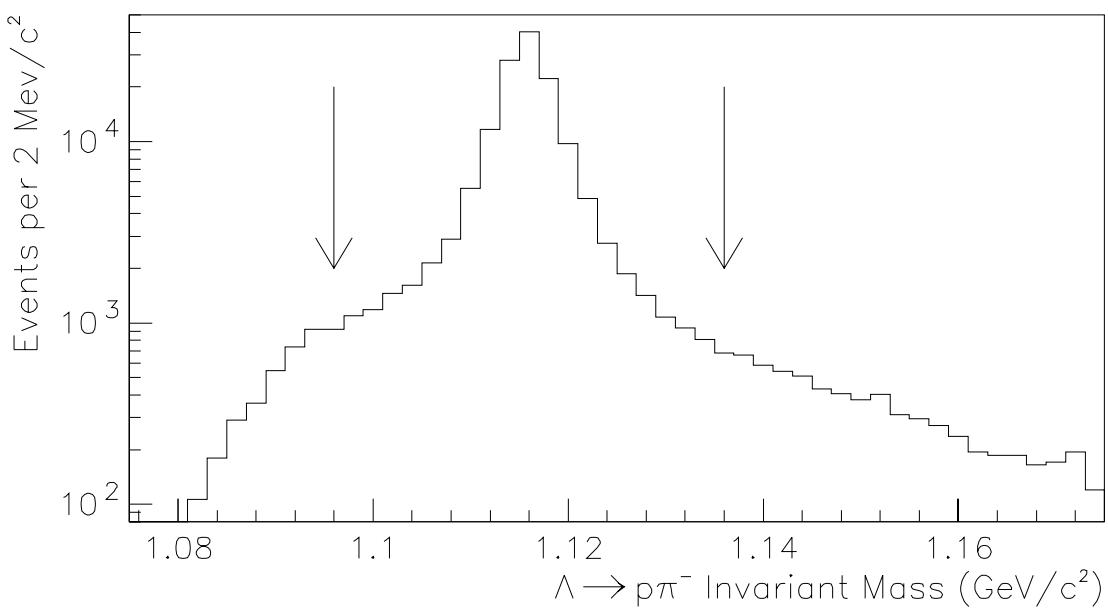

Figure 52. $m_{p\pi}$ for the candidate $\Lambda \rightarrow p\pi^-$ decays used for target location

5.5.2 Lead-Glass

Since the lead-glass could measure positions of electromagnetic clusters with a resolution of only about 4 mm, locating the array in the chamber coordinate system is straightforward in principle. A sample of electrons from the Two Electron trigger is selected by requiring good-quality tracks with $0.8 < E/p < 1.2$ that do not point into the radiation-damaged pipe-blocks. The resulting sample of about 10000 electron candidates used in each fit contained about 90% electrons, while the remaining 10% are charged pions from $K_L \to \pi^+\pi^-\pi^0$ and $K_L \to \pi^\pm e^\mp \nu$ backgrounds to the Two Electron trigger. These electron candidates were then binned by x- and y-position, as was done in the muon chamber alignment, and the offset between track and cluster positions in the $\hat{x}$- and $\hat{y}$-directions were compared as a function of position. Results from a sample fit allowing for offsets and rotations of the glass array in the $\hat{x}$-$\hat{y}$ plane are shown in Figure 53. The offsets were well-measured using this procedure, with a statistical error of about 35 μm.

The issue of relative rotations between the chambers and the glass is somewhat more subtle, however. Because the measured cluster position depends weakly on track angle and because of correlations between track position and track angle, the method used above does not provide a reliable measure of the rotation. It does, however, provide a crude estimate of the maximum rotation present between the glass and the chambers, and limited possible chamber rotations to be $\lesssim$500 μrad. Is this sort of rotation important to correct for in E799? There are two reasons to want to fix the angle between the chamber coordinates and the glass. First, one wants to be able to accurately measure positions at the glass. However, even at the far ends of the chamber, a 500 μrad rotation corresponds to only a 500 μm shift in cluster position, which was only a small fraction of the measurement error in cluster position. Therefore, this did not cause a problem for photon position measurements. Second, one wants a stable chamber coordinate system in the experiment. It is safe to assume that changing chamber-glass rotations are due to moving chambers and not due to rotations in the lead-glass array. In particular, when attempting to find aperture edges, one usually needs to use the highest statistics samples available which are usually accumulated over long periods of time. Therefore, a time-varying overall chamber rotation

relative to other detector elements is potentially problematic. However, in E799 there were no apertures where uncertainties in position at the few hundred μm level could drastically change physics results. (This is in contrast to the situation in an experiment like E731 where precision understanding of acceptances was crucial [57]). Therefore, changes in the overall chamber system rotation of 1 mrad were not important. The chamber-glass relative rotations were then ultimately left uncorrected simply because a sophisticated treatment of the issue was not warranted by the needs of this analysis.

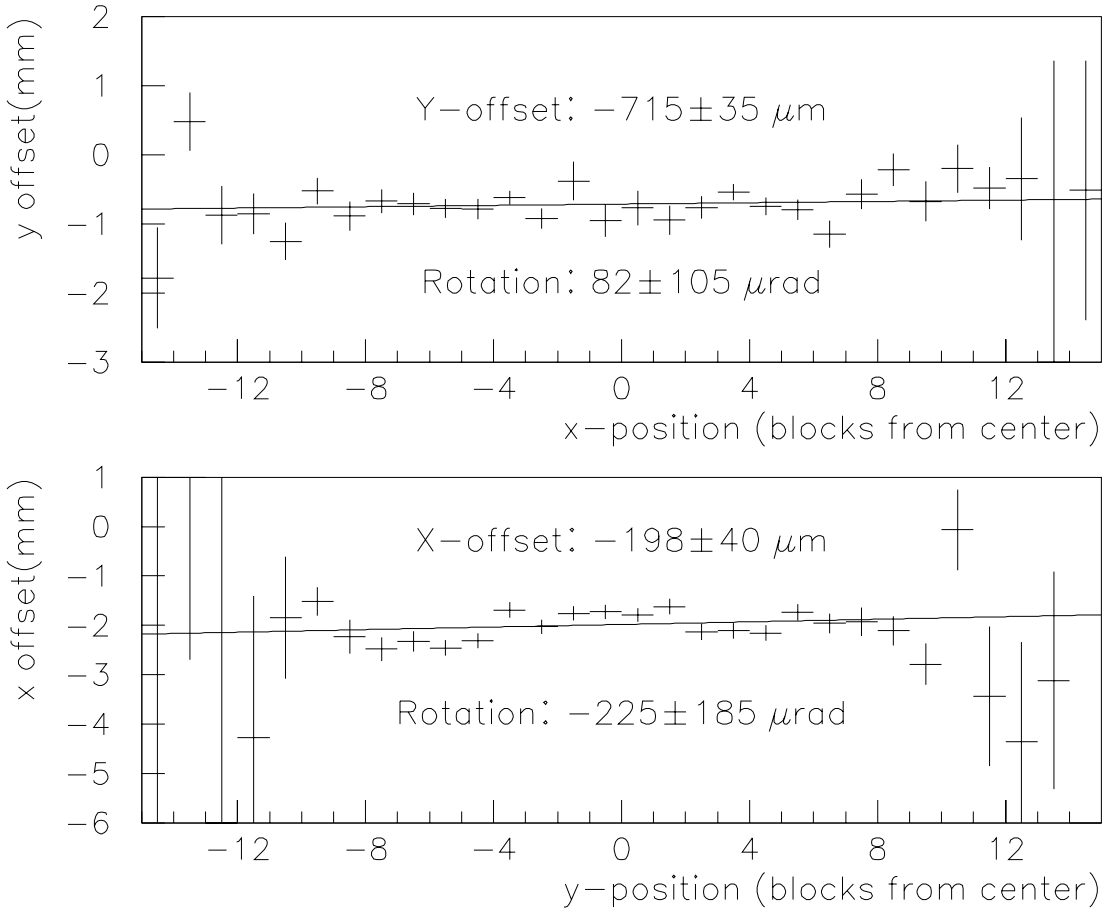

Figure 53. The fit to measure the location of the lead-glass array in the chamber coordinate system

5.5.3 Scintillator Banks and Photon Vetoes

Location of the scintillator banks and photon vetoes was straightforward. Since the banks are active, a large sample of muons written to tape in a special run with the primary beam dumped away from the target was used to point at the scintillator edges in order to locate them. In the case of the photon vetoes, all that was needed was a measurement of the overall position of the counter banks, since the shapes and dimensions of the vetoes were well-measured before installation. For the B and C trigger banks, knowledge of the locations of the dead regions between adjacent counters was important in order to simulate possible position-dependent inefficiencies in the trigger. Therefore, during the E773/E799 switch-over, a large sample of beam-stop-closed muons was written with a $(\overline{1B} \cdot 1C + 1B \cdot \overline{1C}) \cdot MU2$ trigger in order to select events with a muon in an inefficient portion of the counter. The positions determined with this run were be used for both the E773 and E799 Monte Carlo simulations and trigger verifications of the B and C banks. The net effect proved to be fairly small, since less than 1% of the active area of each bank was dead. The problem of locating and simulating the muon trigger banks (MU2, MU3 and MU1) were somewhat different because of the large amounts of material between the chambers. The locations of these banks were not relevant to this analysis however. The details of their alignment are discussed elsewhere [1].

5.5.4 Iron Ring

Finding the iron ring with the data proved to be a somewhat more subtle problem. Unlike the other photon vetoes, there was no active element available with which to locate the ring itself; therefore, it was impossible to find using muons. However, in the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, sample, the effect of the ring is large because in such a seven-body final state, the probability of one of the five photons being at such a large radius is enhanced. In fact, comparing a plot of the maximum photon radius at the iron ring in data and Monte Carlo for this sample, one could see a pronounced difference in these distributions due to the presence of the iron ring in the data and not in the Monte Carlo (Figure 54). The radial position of the ring edge was measured using this data to Monte Carlo comparison

in different angular slices around the ring, and a fit to the ring shape was performed to determine if the ring was off-center in the chamber coordinate-system. The ring center was then measured to be at

$$(x, y) = (+3 \pm 2\ mm, +4 \pm 2\ mm),$$

essentially consistent with zero offset. In the analysis, events containing photons at a radius of greater than 0.835 m, the minimum distance to the ring edge from the $\hat{z}$-axis, were removed.

5.5.5 Beam Collimators

Another set of important detector elements which needed to be aligned were the collimators used to shape the beam. In the Monte Carlo detector simulation for E799, K_L's were generated at the target and then propagated downstream through the collimators in order to produce a simulation of the beam shape seen in the data. While it was not extremely crucial that the fine details of the beam be reproduced for the measurements in E799, the gross features were worth modeling.

The approach taken for measuring collimator positions was to use only information

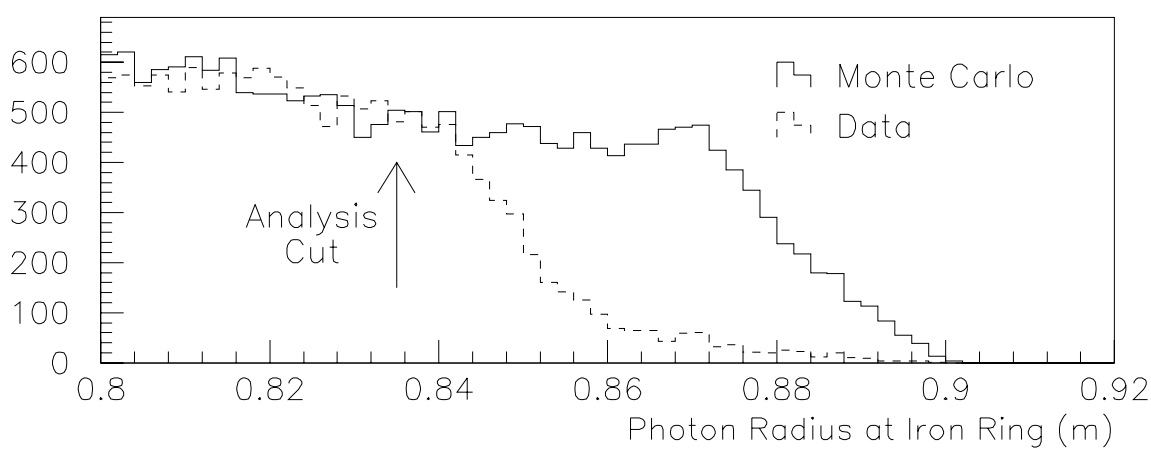

Figure 54. Illustration of the loss of large radius photons from the Iron Ring in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, sample. The cut used in the analysis is shown.

from survey data to measure the change in collimator positions between E773, where the positions were determined using the copious $K^0 \to \pi^+\pi^-$ decays, and E799. During the changeover form E773 to E799, the collimator positions were surveyed as found in E773 and then again as set for the larger beams of E799. By comparing these two surveys, one could simply apply shifts in the collimator positions to the results of the E773 beam shape studies. The resulting beam shape agreement in data and Monte Carlo is shown for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays in E799 in Figures 55 and 56.

5.5.6 Upstream Pipes

A set of potentially important detector apertures which could not be easily observed in studies of the data were those imposed by the small pipes in the upstream portion of the decay volume, particularly the pipe through the BM109 magnet (see Figure 57). These apertures have significant implications for the level of $K_L \to \pi^0\pi^0\pi^0$ background to all-photon final states, such as studies of $K_L \to \pi^0\gamma\gamma$ decays. If one wishes to attempt an absolute normalization of the background, it is important to determine the positions of these pipes in some independent way.

At the 1 mm level, it is possible to measure these pipe positions from survey data collected after the end of the E799 run. The complications, however, were myriad in attempting to tie transverse positions from survey information to the chamber coordinate system. The surveyors worked in a coordinate system where a nominal "beam-centerline" formed the $\hat{z}$-axis and absolute elevations above sea-level defined the y-coordinate. However, it was possible to relate the two coordinate systems by using measurements of the two-hole collimator and lead-glass array positions in surveying coordinates, and then relating those to their known locations in the chamber coordinate system. Amusingly enough, the curvature of the earth results in a 2 mm shift in elevation between the two-hole collimator and the glass, and this effect had to be removed when converting between the two coordinate systems. In the end, this survey information on the positions of the pipes was used in the Monte Carlo, along with measurements of the inside dimensions of the pipes, to simulate any effect from these apertures.

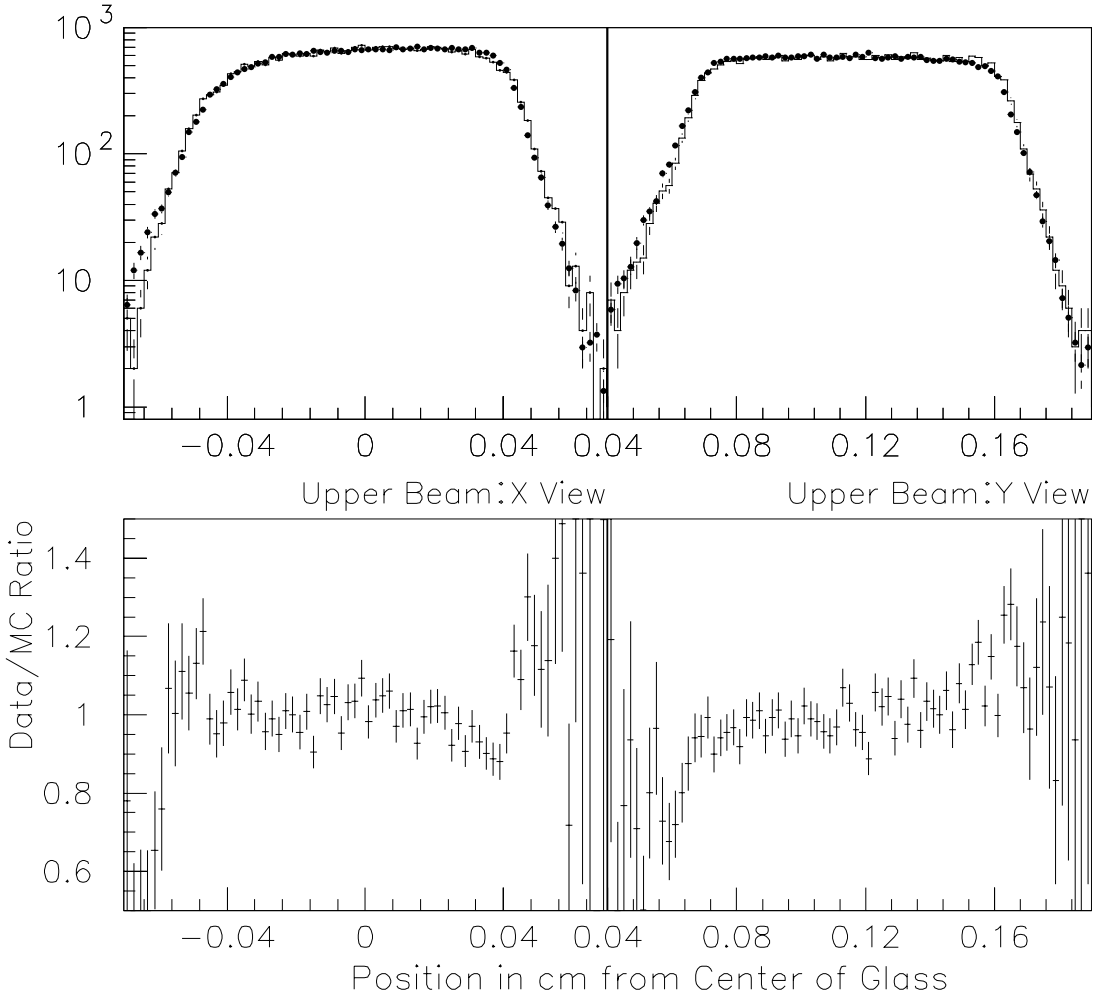

Figure 55. The upper beam intensity as a function of $\hat{x}$ and $\hat{y}$ position compared in data and Monte Carlo

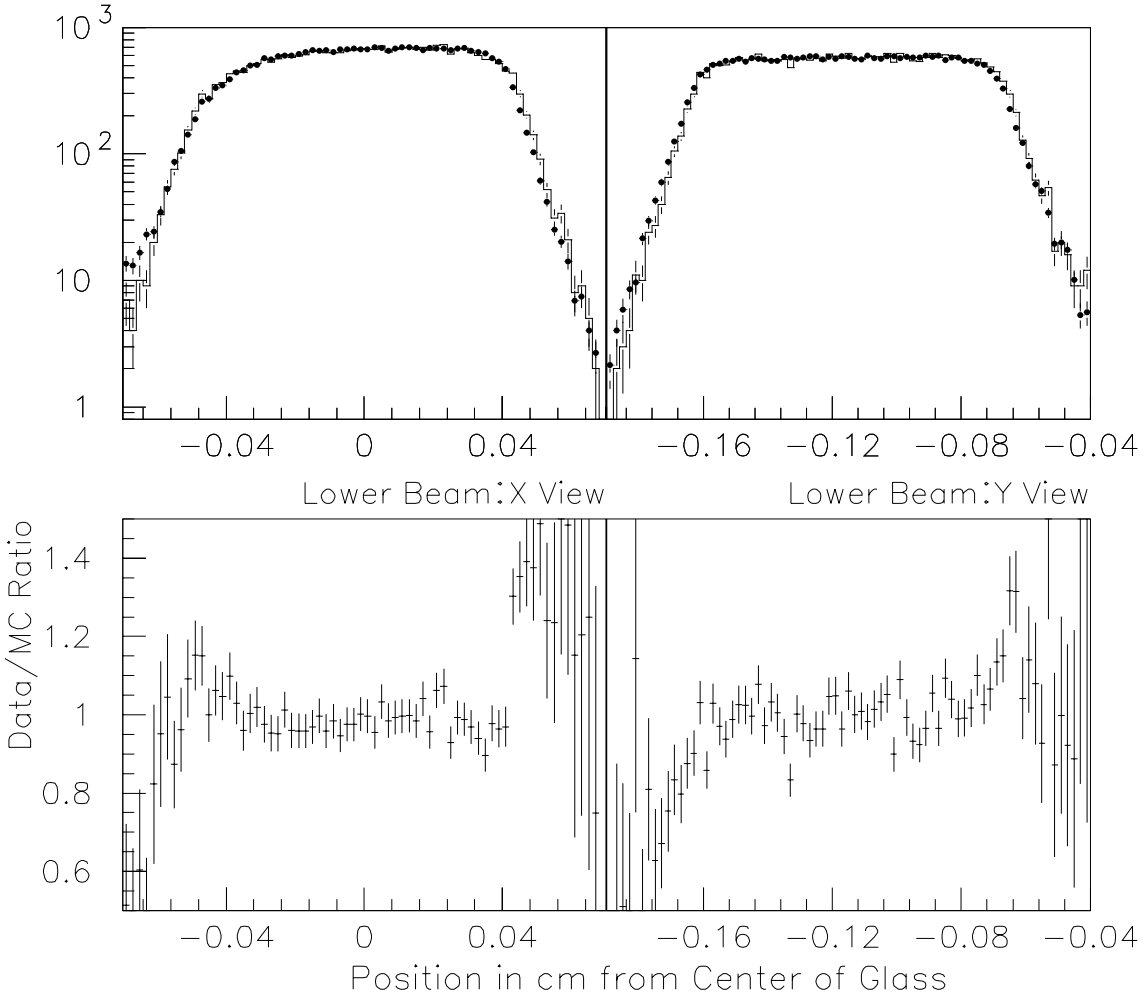

Figure 56. The lower beam intensity as a function of $\hat{x}$ and $\hat{y}$ position compared in data and Monte Carlo

5.6 Alignment and Field Strength of the Magnet

The final element of the spectrometer to be considered is the magnet itself. The transverse momentum kick of the magnet allows the determination of the momentum of each charged track by measuring the bend angle in the magnet aperture. The change in momentum of a charged particle passing through the magnet is given by

$$\begin{aligned} \Delta\vec{p} &\propto \int_{path} \vec{B} \times d\vec{l} \\ &= \hat{x} \int_{path} B_y dl_z - \hat{z} \int_{path} B_y dl_x \end{aligned}$$

where $\vec{B}$ is the magnetic field which is oriented along the $\hat{y}$-direction (other small components are neglected). This results in a momentum kick predominantly along the $\hat{x}$-direction, with a small component in the $\hat{z}$-direction. These components are related to each other by the

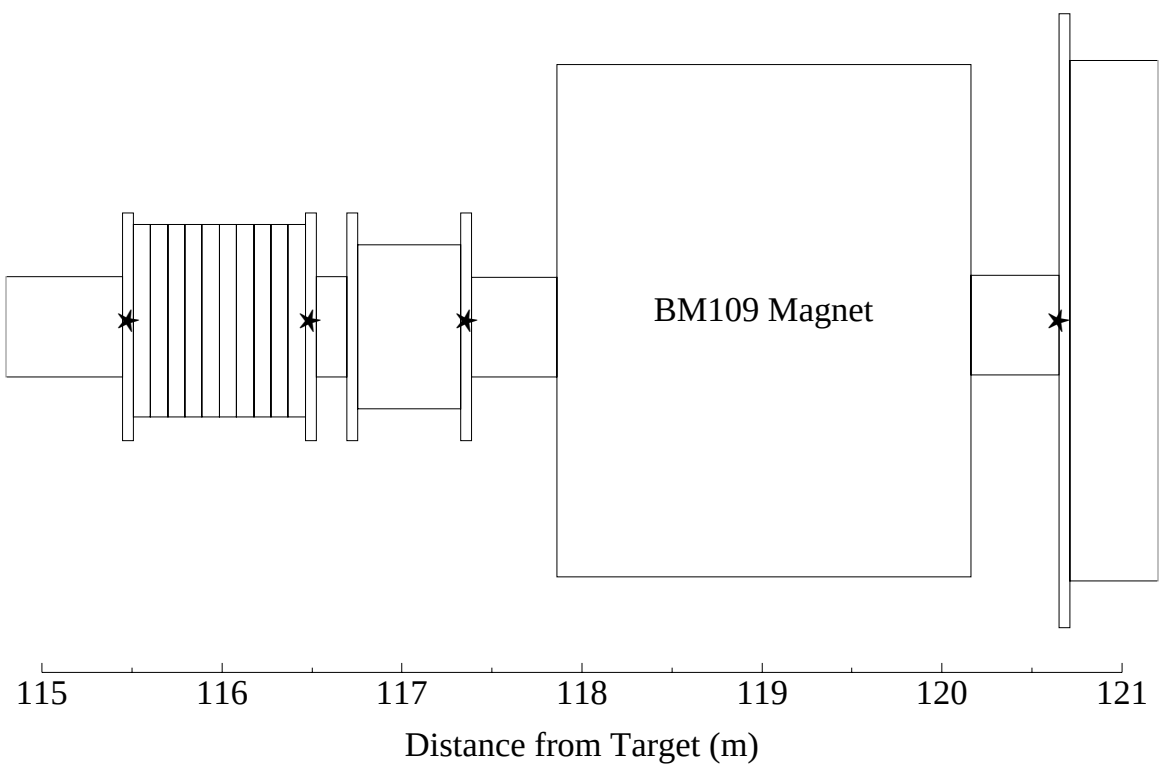

Figure 57. Schematic of the upstream vacuum pipes in the decay volume. The stars indicate the location of the apertures used to simulate the pipes in the Monte Carlo.

average angle in the $\hat{x}$-direction, $(\theta_{ux} + \theta_{dx})/2$, where θ_x is the angle that the projection of the track vector onto the $\hat{x} - \hat{z}$ plane makes with the $\hat{x}$-axis, and where the u and d subscripts denote the upstream and downstream segments of the x track. The bend angle in the $\hat{x}$-direction is related to the momentum kick by

$$\theta_{dx} - \theta_{ux} + \frac{\theta_{dx}^3}{2} - \frac{\theta_{ux}^2\theta_{dx}}{2} + \mathcal{O}(\theta^4) = \frac{\Delta p_x}{p_z}.$$

Thus the total track momentum is given by

$$|\vec{p}| = p_{kick}\frac{(1 + \frac{1}{2}\theta_{ux}^2 + \frac{1}{2}\theta_y^2 - \frac{1}{2}\theta_{dx}^2 - \frac{1}{2}\theta_{dx}\theta_{ux} + \mathcal{O}(\theta^3))}{\theta_{ux} - \theta_{dx}},$$

where p_{kick} is the nominal transverse momentum kick of the magnet, which is defined as the kick given to an infinite momentum particle traveling parallel to the $\hat{z}$-axis. Note that in the analysis code, none of the above corrections to $\mathcal{O}(\theta^2)$ were used.

5.6.1 Field Strength

In order to make a momentum measurement from the above formula, it was necessary to know the nominal p_{kick} of the magnet. There was some variation of the magnetic field as a function of position across the magnet aperture. For previous experiments [57], a map was made of the magnetic field on a 5.1 cm grid which found variations of about 2% across the aperture and measured these to an accuracy of 0.1%. Corrections to the field strength based on this map were applied in the analysis code.

To normalize the overall field strength, the sample of $K_L \to \pi^+\pi^-$ decays taken in E799 was analyzed. Unfortunately because these events could only be found in the Minimum Bias Trigger which was prescaled by 3600, there were only 4000 such events available for study. This made it very difficult study variations in the field strength as a function of time. However, this sample (see Figure 58) was sufficiently large to allow a determination of the average momentum scale to good precision since the $\pi^+\pi^-$ mass is determined to 0.007% of itself. A correction to p_{kick} was then applied based on the shift observed from the value for the $K_L \to \pi^+\pi^-$ mass measured in a sample of Monte Carlo decays generated with a K_L mass equal to the measured K^0 mass given by the Particle Data Group [3] of (497.671 ± 0.031) MeV/c^2. Because the $\pi^+\pi^-$ invariant mass is given by

$$m_{\pi\pi}^2 = 2m_\pi^2\left(1+\frac{p_1}{p_2}+\frac{p_2}{p_1}\right)+2p_1p_2(1-\cos\theta)+\mathcal{O}\left(\frac{m_\pi^4}{p^2}\right),$$

where p_i is the momentum of the ith pion, a momentum scale shift can be related to a K_L mass shift by

$$\begin{aligned}\frac{\delta p_{kick}}{p_{kick}} &\approx \frac{m_K}{\sqrt{m_K^2-2m_\pi^2}}\frac{\delta m_K}{m_K}\\ &\approx 1.09\frac{\delta m_K}{m_K}.\end{aligned}$$

Shown in Figure 59 is a plot of the mean fitted K_L mass from $K_L \to \pi^+\pi^-$ decays for twelve roughly equal data samples during the run. No evidence is seen for a shift in the momentum scale as a function of time, so only one overall correction, with a statistical accuracy of 0.012% was applied for the entire run. It should be noted that this correction

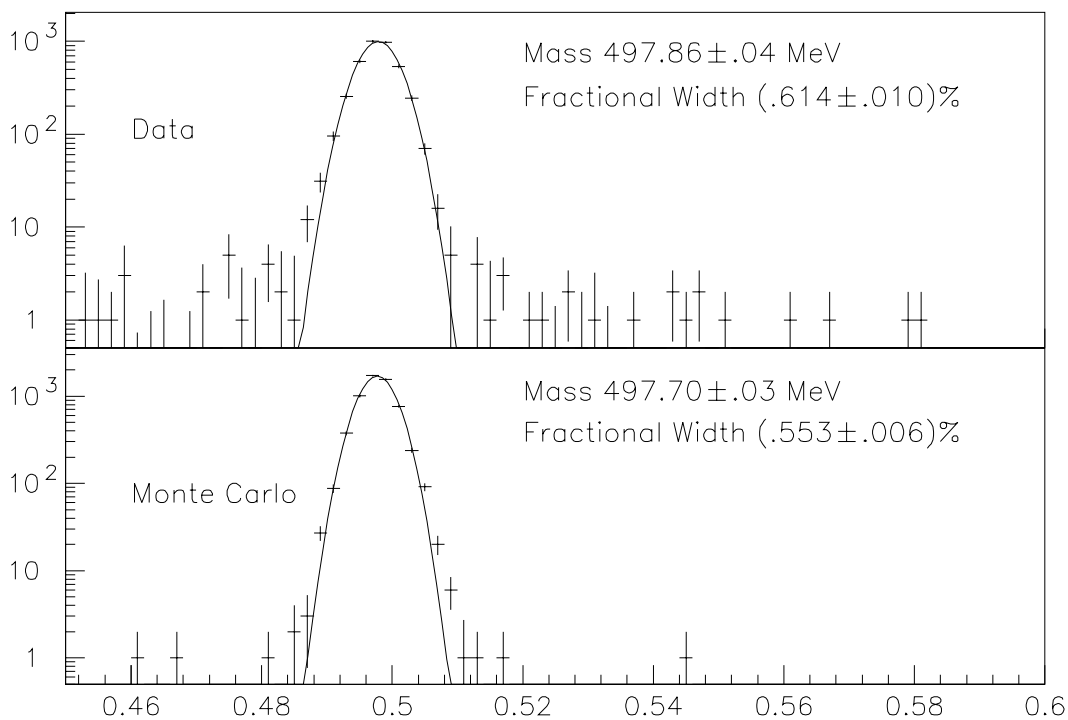

Figure 58. $m_{\pi\pi}$ for candidate $K_L \to \pi^+\pi^-$ decays in data and Monte Carlo. The data and Monte Carlo are both background subtracted using sidebands in p_t^2.

is tangled up with any uncertainty on the measurement of the opening angle that comes with the error on the separation between chambers 1 and 2. Recall from the discussion of survey data (see Section 5.1), that there is an error of approximately 0.02% in the invariant mass of two light ($m \ll m_K$) particles. So, although this study does fix the error on the invariant mass of two light particles to be 0.012%, the error on the momentum scale itself is actually the sum of the uncertainties from the opening angle and from the $K_L \to \pi^+\pi^-$ mass study, or 0.023%.

In addition to measuring the momentum scale, this study with the $K_L \to \pi^+\pi^-$ decays also allows a study of momentum resolution. In Figure 58, the data has an additional (0.27 ± 0.02)% Gaussian smearing above what is seen in the Monte Carlo. The reasons for this smearing are not understood, and the effect of such an unexpected smearing was studied as a source of systematic error. One could also use this plot to place limits on allowed non-Gaussian smearing, for use in systematic studies.

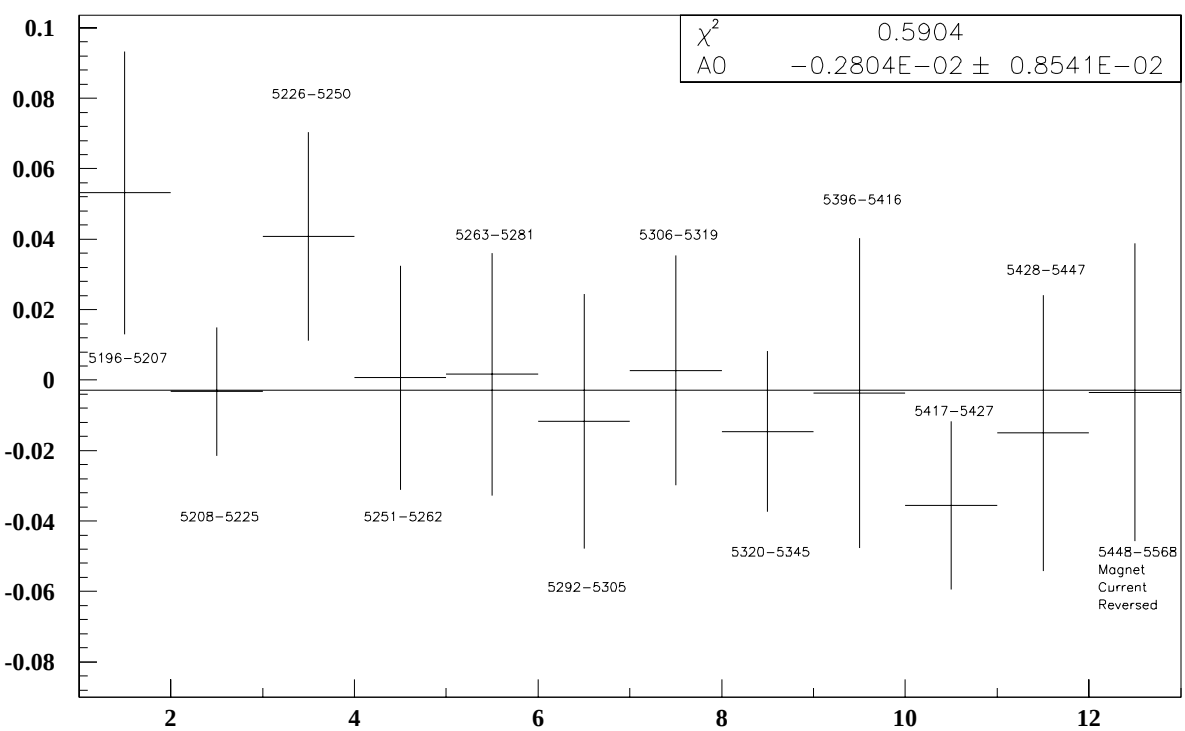

Figure 59. Variation of the momentum scale with time as measured in $K_L \to \pi^+\pi^-$ decays

5.6.2 Location of the "Bend Plane"

One subtle issue which has little effect on the momentum determination of a track, but which may be important in track quality parameters is the position of the "bend plane" of the magnet. In the momentum calculations and in the matching of upstream and downstream $\hat{x}$-view segments, the momentum kick of the magnet was treated as though it were applied all at once and in a constant z-position. This plane of constant z at which magnet kick occurs in this approximation is referred to as the "bend plane". It is possible to make a reasonable determination of this z-position from survey data by assuming it will be at the midpoint of the magnet pole faces. However, since the offset at the magnet is the parameter used in the upstream-downstream $\hat{x}$-view track matching to measure how "good" a track candidate is, it is important to check this approximation independently. An error in the z-position of the bend plane of 2 cm for a 3 GeV track would result in an apparent offset at the magnet on about the same order as the resolution.

To do this, a sample of muons from the beam-dump, taken with the analysis magnet on, was analyzed. For each event, the apparent bend-plane position where the upstream and downstream tracks met was plotted versus location at the magnet and track angles. A 8.4 mm shift from the expected bend plane z-position obtained from the survey data was found, and this shift was incorporated into the analysis code. This offset is shown below as a function of the bend-angle (see Figure 60).

5.7 Determination of Chamber Efficiencies

In order to understand the efficiencies of trigger and reconstruction algorithms using the chambers, it was necessary to measure the efficiencies of wires in the drift chambers. This was done using a very simple algorithm. Events from the output of the Two Electron crunch with two identified electron tracks were checked for tracks with missing hits in the chambers. The tracking algorithm used in this check was slightly different from that used in the Two Electron crunch and the rest of the E799 analysis in that explicit cuts were not placed on the sum of distances in this modified tracking algorithm to avoid a bias from δ-rays or accidental hits in determination of the efficiencies. Every track found with this

new algorithm was checked for missing hits. If a missing hit was found, the wire on which that hit would have been was inferred using the remaining track information. Efficiencies were calculated for each wire by simply calculating the fraction of missing hits on that wire compared to the total number of hits observed. If the observed inefficiency was significantly different at the 4σ level, from the overall plane efficiency, that wire was labeled as inefficient. Once a wire was so selected, its contribution to the overall plane inefficiency was removed, thus giving a corrected inefficiency. Both the overall plane inefficiencies and the individual wire inefficiencies were written out to a special file for use in the Monte Carlo simulation of the detector. This procedure was repeated for each run with sufficient data statistics.

A sample plot of chamber inefficiencies by wire is shown in Figure 61. In general, chamber efficiencies did not vary significantly by run, with the exception of a small number of wires. The plane efficiencies, once corrected for the inefficient wires, were typically above 99%—the worst single plane efficiencies, found in the inner planes of chamber 1, were still above 98.3%. These corrections for inefficient wires could be quite dramatic (Figure 62), especially when the inefficient wire was located in the center of the chamber.

This method of measuring chamber inefficiencies did have limitations. In particular, in-

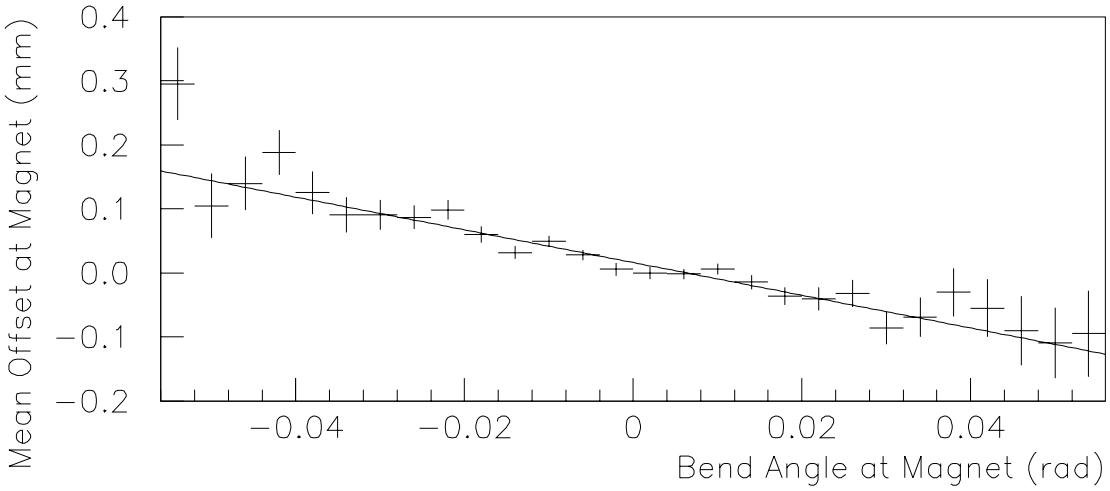

Figure 60. The mean offset of tracks at the magnet shown as a function of the bend angle. The linear dependence shows a small shift in the apparent location of the magnet bend plane.

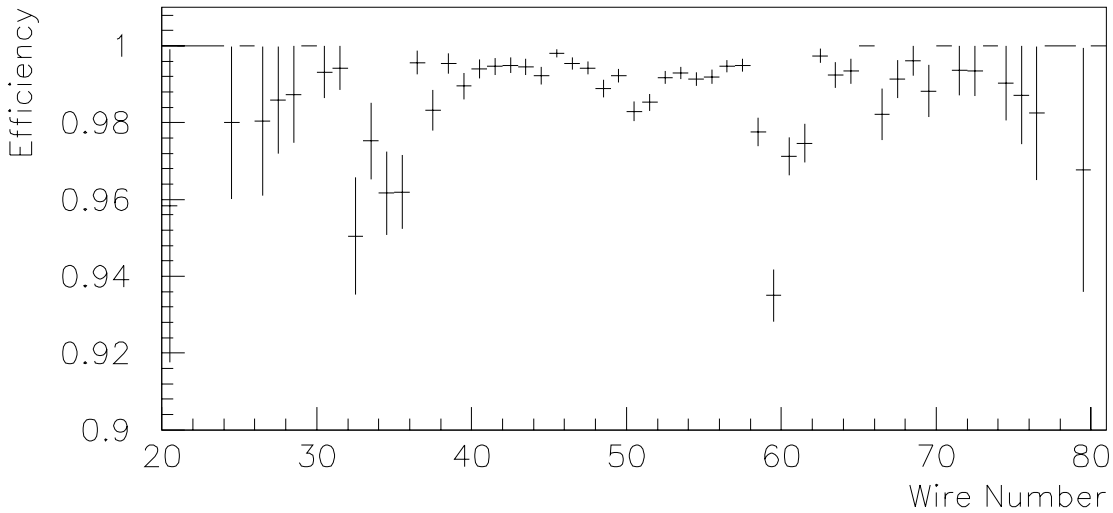

Figure 61. The chamber wire inefficiencies from one run in the upstream plane of Chamber 2Y.

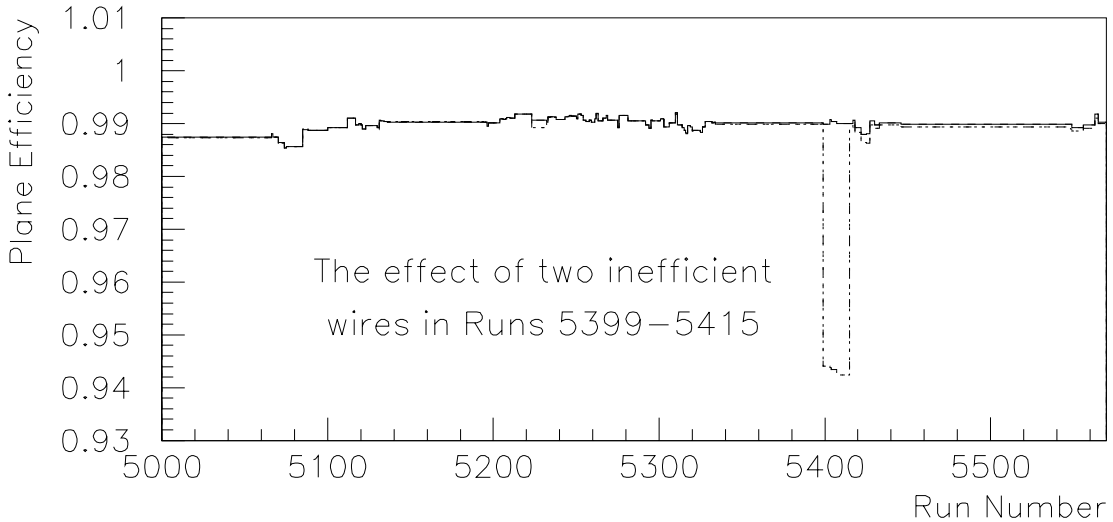

Figure 62. The corrected and uncorrected (shown in the broken line) plane efficiency for the downstream plane of Chamber 4X. Note the large effect of a two wires which were very inefficient over a period of several runs.

efficiencies in the hit-counting trigger of the track processor could have biased this efficiency determination. Also, there was no attempt to study the inefficiencies of the hit-counting with this technique. A more refined discussion of the treatment of these issues can be found in Section 8.2.1.

5.8 Theoretical Calculation of the Chamber Resolution

A great deal effort was put into improving the single-hit resolution of the drift chambers. It is interesting to compare the measured chamber resolutions of 95 to 110 μm with the fundamental limits of resolution in the drift chamber.

One major contribution to the drift chamber resolution was the stochastic nature of the primary ionization process. Because electron-ion pairs are produced in discrete interactions separated on average by some macroscopic distance, the distance drifted will often not be the distance of closest approach of the charged track. The primary interaction rate in many common gasses has been shown to be proportional to their atomic number [64], and by using this proportionality, one estimates the primary interaction density in 50%-50% Argon-Ethane to be 2.9 interactions/mm. This means that the mean distance between the point of closest approach of the track to the wire and the closest primary interaction, denoted by λ, is 170 μm. The RMS contribution to the position resolution introduced by this statistical effect is, in closed form, [1]

$$\sigma_{\mathrm{PI}} = \sqrt{x^2 + 2\lambda^2 - \frac{x^4\pi^2}{4\lambda^2}\left(H_1\left(\frac{x}{\lambda}\right) - N_1\left(\frac{x}{\lambda}\right)\right)^2},$$

where x is the distance of closest approach of the track to the wire, H_1 is the Struve function of order 1, and N_1 is the Neumann function of order 1. In the limit as $x \to 0$, the Struve and Neumann functions diverge; however, σ_{PI} approaches λ in this limit. This contribution to the chamber resolution is shown in Figure 63.

A second major effect contributing to intrinsic chamber resolution was diffusion of electrons during the drift time. As the electrons moved through the drift region of the chamber,

[1] My thanks go to Arthur Kosowsky and Peter Shawhan, former members of the fearsome Washington University Putnam Mathematics Examination team, for their assistance in evaluating the particularly nasty integral involved. This Struve function is for you, Arthur!

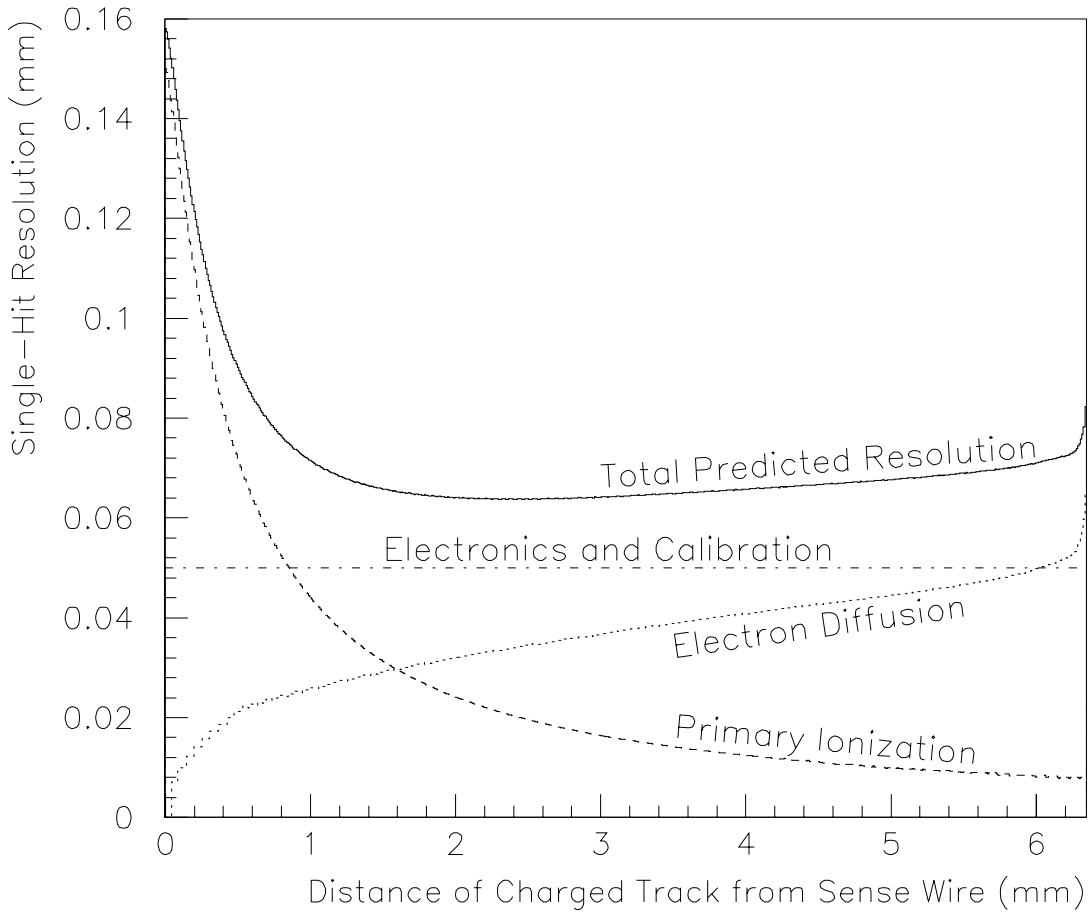

Figure 63. Contributions to the calculated chamber resolution

they diffused at a rate which was largely independent of variations in the electric field. Measurements of the diffusion rate for 50%-Argon, 50%-C_2H_6 gas mixtures using test chambers [65] and calculations from Monte Carlo simulations [66] predict an RMS one-dimensional diffusion rate to be

$$\sigma_D = 8.5\sqrt{t\ \text{(in ns)}}\ \mu\text{m},$$

where t is the drift time, assuming a 50 μm/ns drift velocity. The $\sqrt{t}$ dependence results, of course, from the random-walk nature of the diffusion process. However, this effect was somewhat mitigated by the large number of electron-ion pairs produced by the ionization track. For our gas mixture, the mean number of electron-ion pairs produced at each primary interaction was 3.3 [64], and at large distances from the wire, electrons from several primary interactions arrived essentially synchronously. The effective factor by which the RMS diffusion is reduced for n drifting electrons is [64]

$$\frac{\pi}{2\sqrt{3\log n}},$$

and an approximate expression for the mean number of electrons which arrived within one standard deviation (σ_D) of the arrival time of the closest primary interaction is

$$n \approx \begin{cases} 3.3 & \text{if } x < 300\ \mu\text{m, or} \\ 3.3\left(\frac{\sqrt{2x\sigma_D(x)}}{170\ \mu\text{m}}\right) & \text{otherwise,} \end{cases}$$

where x is the distance of closest approach of the charged track to the sense wire. The contribution to the resolution from diffusion, including this correction is shown as a function of x in Figure 63.

Finally, there were other contributions to the resolution, due to the electronics and remaining calibration uncertainties. The TDCs used have 1 ns granularity in their time measurement, contributing a 15 μm term to the resolution. Similarly, the zero-drift time, as determined in the calibration procedure described in Section 5.2, also has a 1 ns granularity, and contributes another 15 μm. The error on the event-time correction, calculated from the sums of distances in the drift chamber, contributes an additional 26 μm term to the resolution. Systematic and statistical errors in the $X(t)$ determination were evaluated at between 41 μm and 26 μm, depending on the chamber plane. Finally residual alignment

uncertainties were evaluated at 20 μm. All of these contributions except the last contribute to the single-hit resolution in the chamber as measured by the sums of distances, and the net contribution from these effects, between 43 and 53 μm, is shown also in Figure 63.

The contributions to the resolutions discussed here explain much of the observed single-hit resolution as observed in the chambers. Other possible effects contributing to the resolution include changes in macroscopic parameters affecting the chamber drift velocities or gains For example, in our chambers, analytic models of chamber response predict variations of drift velocity with atmospheric pressure of about 0.2 μm/nsec/Hg inch, and variations with temperature of 0.03 μm/nsec/°C [67]. Since average drift times were about 65 nsec in our chambers, rapid changes in temperature or pressure could have significant effects on chamber response. For reasonable changes in atmospheric pressures and temperatures over one day, 0.5 Hg-inch and 5 °C, one estimates an effect that would change position measurements by about 15 μm on average. Similarly, variations in temperature or atmospheric pressure affected the intrinsic gain of the chambers as well. Since the discrimination of the signal from the chamber was done using constant threshold discriminators, and since the analog chamber signals had a finite rise-time, timing variations could result from these gain changes. In addition to these macroscopic effects, small variations in cell geometry, drift field changes from broken field-shaping wires and small variations in gas mixtures could all cause resolution effects in the chambers.

5.9 Conclusion

With the completion of these tasks, the charged spectrometer and chamber-based coordinate system were well-understood. Chamber performance was more than adequate in E799. Chamber single-hit resolution was close to the intrinsic resolution limits. For most multi-body decay modes, multiple-scattering rather than chamber resolution ultimately limited the momentum resolution, and any variations in alignment uncorrected by the calibration procedures were sufficiently small so as to not cause observable effects in the data.

CHAPTER 6

CALIBRATION AND MODELING OF THE LEAD-GLASS CALORIMETER

Get thee glass eyes;
And like a scurvy politician, seem
To see the things that thou dost not.
—*William Shakespeare, King Lear*

The physical characteristics of the lead-glass calorimeter have already been described in Section 2.2.2. This chapter discusses both the framework used to model the lead-glass response, and the calibration procedure for the calorimeter based on this model.

Careful calibration and simulation of the calorimeter was necessary to understand photon or electron response in the lead-glass, and this understanding had a direct impact on physics measurements. In the case of photons, the calorimeter provided the only measurement of energy and position at the calorimeter. This information, along with the decay vertex location, is sufficient to measure the momentum vector of a photon. A typical use of the lead-glass for candidate photons is to measure the invariant mass of a pair of photons and compare it to the mass of a hypothesized parent π^0. For a case such as $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, or $K_L \to \pi^0 e^+e^-$, where the decay vertex has been measured from the charged tracks, the candidate π^0 mass is given by

$$m_{\pi^0}^2 = (E_{\gamma 1} + E_{\gamma 2})^2 - |\vec{p_{\gamma 1}} + \vec{p_{\gamma 2}}|^2$$

$$
\begin{aligned}
&= 2E_{\gamma 1}E_{\gamma 2}\left(1-\cos(\theta_{\gamma 1\gamma 2})\right) \\
&\approx E_{\gamma 1}E_{\gamma 2}\left(\frac{r_{\gamma 1\gamma 2}}{z_{PbG}-z_{decay}}\right)^2,
\end{aligned}
$$

where $r_{\gamma 1\gamma 2}$ is the distance between the two photons at the lead-glass and $z_{PbG} - z_{decay}$ is the distance between the lead-glass showers and the decay vertex. In order to optimize the π^0 mass resolution, then, it is clear that good energy and position information from the lead-glass were needed. Also, understanding the energy and position resolution of the lead-glass will allow the Monte Carlo simulation to accurately predict the efficiencies for kinematic and particle identification cuts.

6.1 Model of Lead-Glass Response

When high-energy electrons pass through material, they lose energy primarily by production of bremsstrahlung photons. High-energy photons in turn interact with matter primarily through pair production or Compton scattering, and both of these processes produce high-energy electrons. These multiplicative processes produce what is referred to as an "electromagnetic shower" from the interaction of a single primary electron or photon. While high-energy particles produce multiple low-energy particles in the shower process, lower energy particles lose energy in non-multiplicative processes. For electrons, these losses occur dominantly in collision processes; for photons, the photoelectric effect is the dominant energy loss mechanism. Thus, as the shower develops deeper into the lead-glass medium, it first grows in particle multiplicity and energy loss, reaching a "shower maximum". Then as the shower particles become lower in energy deep into the medium, multiplicity and energy loss decrease and the shower decays with increasing depth.

As the relativistic electrons in a shower pass through the lead-glass, they emit Čerenkov light. Čerenkov radiation is emitted when a charged particle passes through a medium whose phase velocity is less than the velocity of the charged particle, or

$$\beta > \frac{1}{\sqrt{\epsilon(\omega)}},$$

where β is v/c of the electron, and $\epsilon(\omega)$ is the dielectric constant of the medium at the frequency of the emitted radiation radiation [68]. Čerenkov radiation also has a fixed angle

of emission with respect to the path of the incident charged particle given by

$$\cos\theta_c = \frac{1}{\beta\sqrt{\epsilon(\omega)}}.$$

For a diamagnetic medium like the lead-glass, $\sqrt{\epsilon(\omega)} \approx n$ in the optical frequencies, where n is the index of refraction of the lead-glass. We know then that the Čerenkov angle, θ_c, is always less than the maximum incident angle of total internal reflection, $1/n$. For an electron traveling at an angle in the block with $\beta \approx 1$ then, there was some Čerenkov radiation which did not reflect back into the block. This, however, was a small effect because the block geometry ensured that most of the Čerenkov light from a charged particle traveling at a small angle with respect to the block axis would still be collected. The loss of Čerenkov light because of this effect was consequently neglected in the Monte Carlo simulations. Because the Čerenkov angle is large in lead-glass, 51.3° for $\beta = 1$, the Čerenkov light typically reflected many times in the block, so any absorption of Čerenkov light in the block was averaged over the transverse dimensions. This eliminated the need to treat variations in absorption of Čerenkov light over transverse slices of the lead-glass blocks.

Even before being radiation-damaged, the lead-glass had some absorption in the region of the Čerenkov spectrum selected by the Wratten 2A filter ($\lambda \gtrsim 420$ nm), typically between 1 and 2% per radiation length. When the lead-glass was exposed to large radiation doses (see Figure 18), this absorption increased substantially, so it is important to account for the absorption in the model of the lead-glass response. Denote the effective absorption per radiation length (X_0) for Čerenkov light by α, and note that

$$\alpha = \alpha_c \cos\theta_c,$$

where α_c is the physical absorption length for Čerenkov light in units of $1/X_0$, and θ_c is the Čerenkov emission angle defined above. This rescaling of the absorption length accounts for the extra path length of the Čerenkov light relative to the length of the lead-glass block. Note also that this absorption has been assumed to be independent of the position within the block. This was a reasonable approximation for the dimensions transverse to the beam direction due to the many reflections of Čerenkov light mentioned above; however, the assumption that the effective absorption was independent of depth (the direction parallel to the beam direction), has not yet been justified. This will be discussed further below.

We denote the distribution function of the produced Čerenkov light by $f(E,t)$, where E is the total shower energy and t is the depth into the block in units of radiation lengths, and $f(E,t)$ is normalized such that

$$\int_0^\infty f(E,t)dt = 1.$$

One can then express the total light observed at the phototube by

$$I(E,\alpha,t) = g_0 E \int_0^L f(E,x) e^{-\left(\frac{L-t-x}{\alpha}\right)} dx, \tag{6.1}$$

where g_0 is the gain constant in photoelectrons/GeV, and L is the total length of the lead-glass block. Fluctuations in longitudinal shower shape changed the shower distribution function and thus the light output. Furthermore, these fluctuations worsened the inherent resolution of the calorimeter. Typical fluctuations in shower shape are illustrated in Figure 64.

The average longitudinal shower profile also varied as a function of the shower energy. Higher energy showers typically deposited more energy deeper into the block, as shown in Figure 65. A parameterization, based on fits to EGS Monte Carlo showers (see Section 6.1.1), of the mean depth of the point of maximum shower energy deposition ("shower max") as a function of energy was performed with the result that

$$t_{max}(E) \approx 3.05 + 1.00 \log E + 0.017 (\log E)^2,$$

where $t_{max}(E)$ is the depth of the shower max in radiation lengths (see Figure 65). The dependence of the depth of energy deposition on energy combined with the absorption introduces an inherent non-linearity in the lead-glass energy response. One can estimate the size of this non-linearity by assuming that

$$f(E',t) \approx f(E, t + t_{max}(E') - t_{max}(E)).$$

The non-linearity in light response in this approximation is then

$$\begin{aligned} \frac{I(E',\alpha,t)}{I(E,\alpha,t)} &\approx e^{\alpha(t_{max}(E') - t_{max}(E))} \\ &\approx \left(\frac{E'}{E}\right)^{\left(1+\alpha\left(1.00+0.017\log\frac{E'}{E}\right)\right)}. \end{aligned}$$

The actual non-linearity in the data is shown in Figure 66 for data from special electron

calibration runs. This parameterization of the non-linearity was a reasonable approximation for small absorptions, but began to fail at larger values of α where the shower fluctuations deep in the block became increasingly important in comparison to the location of the shower maximum.

6.1.1 Corrections for Non-Linearity in the Response

Corrections to this non-linearity were implemented using information from showers generated using the EGS4 electromagnetic shower Monte Carlo [69]. EGS simulates the full development of electromagnetic showers, including pair production, bremsstrahlung, and Möller, Compton and Bhabha scattering. For studies of the lead-glass in past experiments, a library of EGS showers in lead-glass was generated, and this same shower library was

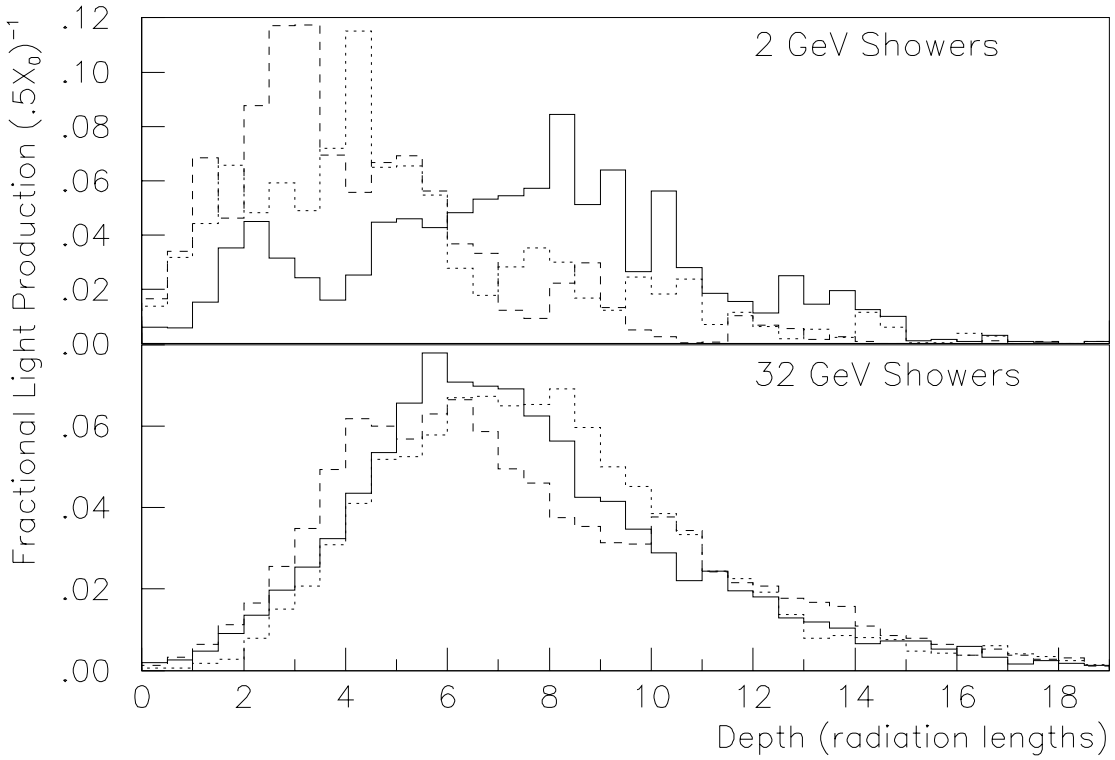

Figure 64. The shower distribution function for three different EGS Monte Carlo generated showers at 2 and 32 GeV.

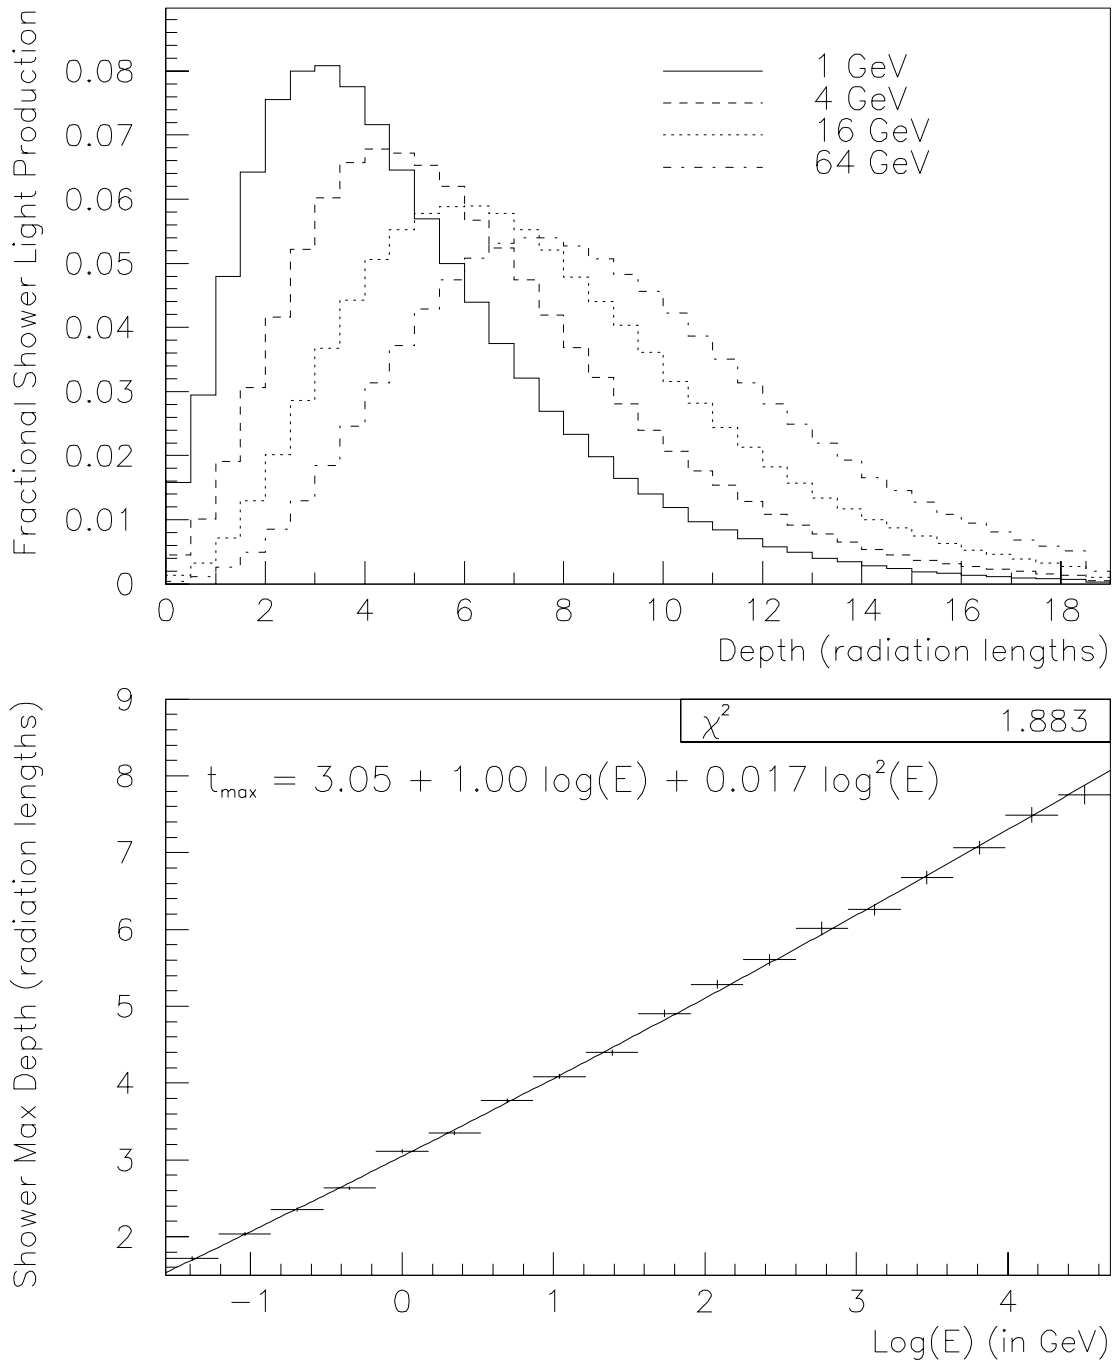

Figure 65. The average longitudinal profile is shown for EGS Monte Carlo showers of 1 GeV, 4 GeV, 16 GeV and 64 GeV. A fit to the location of the shower max as a function of energy is also shown.

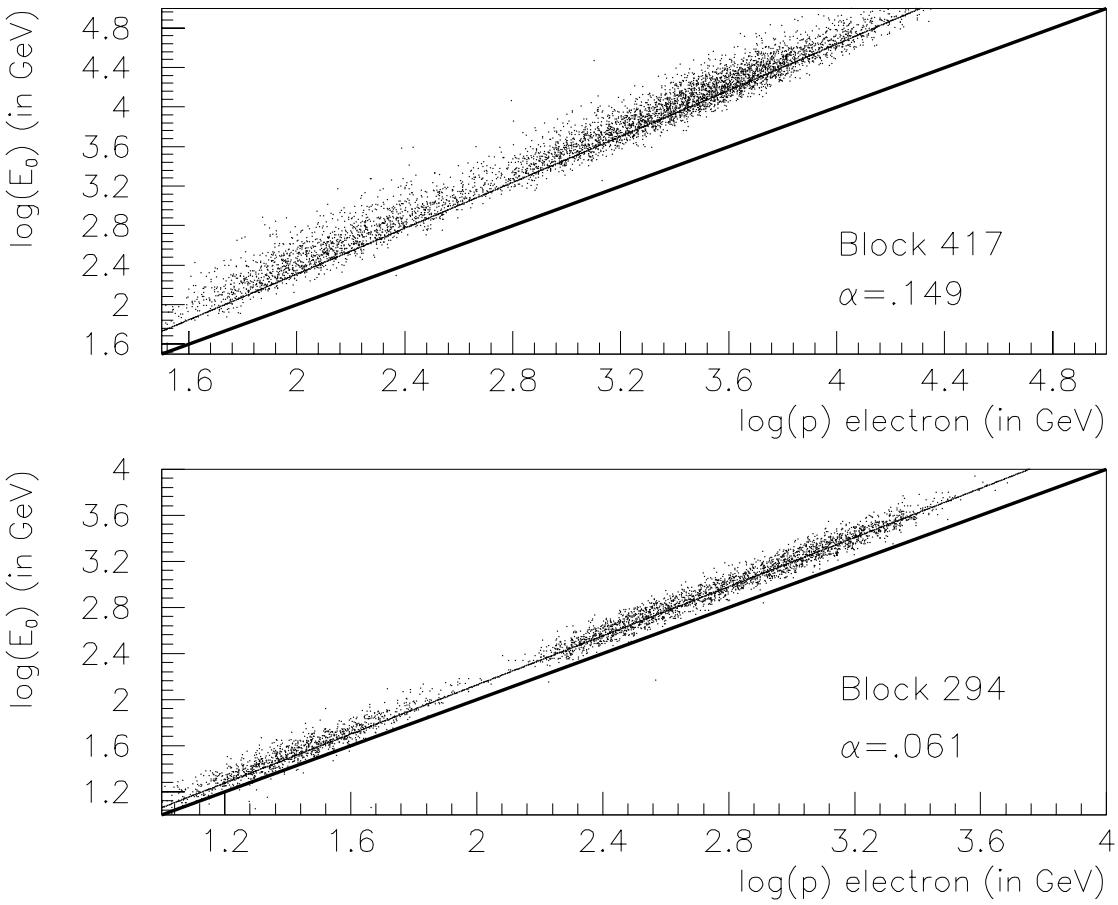

Figure 66. $\log E$ vs. $\log p$ shown without the absorption correction in order to illustrate the inherent non-linearity of the lead-glass. The two lines show the expected result without absorption, and a prediction for the non-linearity based on the fit in Figure 65.

used for simulations of the lead-glass in E799. More detailed descriptions of the EGS shower simulations can be found elsewhere [55]. Each shower was generated in a large block of lead-glass, and once a shower was generated, the block was divided into cells measuring $0.36 \times 0.36\ X_0^2 \times 0.5\ X_0$ in depth (5×5 bins transverse to the beam direction, and 38 longitudinal bins per lead-glass block). The number of charged particles in each cell was stored for each shower, weighted by the probability for Čerenkov emission,

$$1 - \frac{1}{n^2\beta^2}.$$

Showers were generated at 18 energies between 0.25 GeV and 90.51 GeV in evenly spaced increments of $\log E$. The Monte Carlo statistics were poorer for high-energy showers than for low-energy showers because the latter were much more CPU intensive, but at least 210 showers were generated at each energy.

With Monte Carlo shower simulations to provide information about the shower distribution function, $f(E,t)$, one could extract correction factors for showers in the data. The first of these corrections was used to relate the observed light to the incident electron energy. Define

$$C(E,\alpha,0) \equiv \frac{\langle I(E,\alpha,0)\rangle}{g_0 E} e^{\alpha(L-s)},$$

where, as above, $I(E,\alpha,t)$ is the light yield, g_0 is the gain constant, and s is an empirical value, $s = 5.157$, such that $C(E,\alpha,0)$ is approximately 1 for $E = 1$ GeV. From the definition of $I(E,\alpha,t)$ (Equation 6.1), one can write $C(E,\alpha,0)$ more simply as

$$C(E,\alpha,0) = \left\langle \int_0^L f(E,t) e^{-\alpha(s-t)} dt \right\rangle.$$

This equation makes more evident the benefits of using of the EGS Monte Carlo. With potentially large absorptions in the lead-glass, it was important to average over light response rather than to simply average over shower distribution functions ($f(E,t)$). The full shower simulation took into account fluctuations in shower shape at large depths which comprised a large portion of the response for high absorption blocks. Once the $C(E,\alpha,0)$ values had been calculated (see Figure 67), the non-linearity for electron response could be corrected using the relation

$$EC(E,\alpha,0) = \frac{1}{g} I(E,\alpha,0),$$

where g is the gain constant given by

$$g = g_0 e^{\alpha(s-L)},$$

and I is the light response, proportional to the observed signal in the glass ADCs.

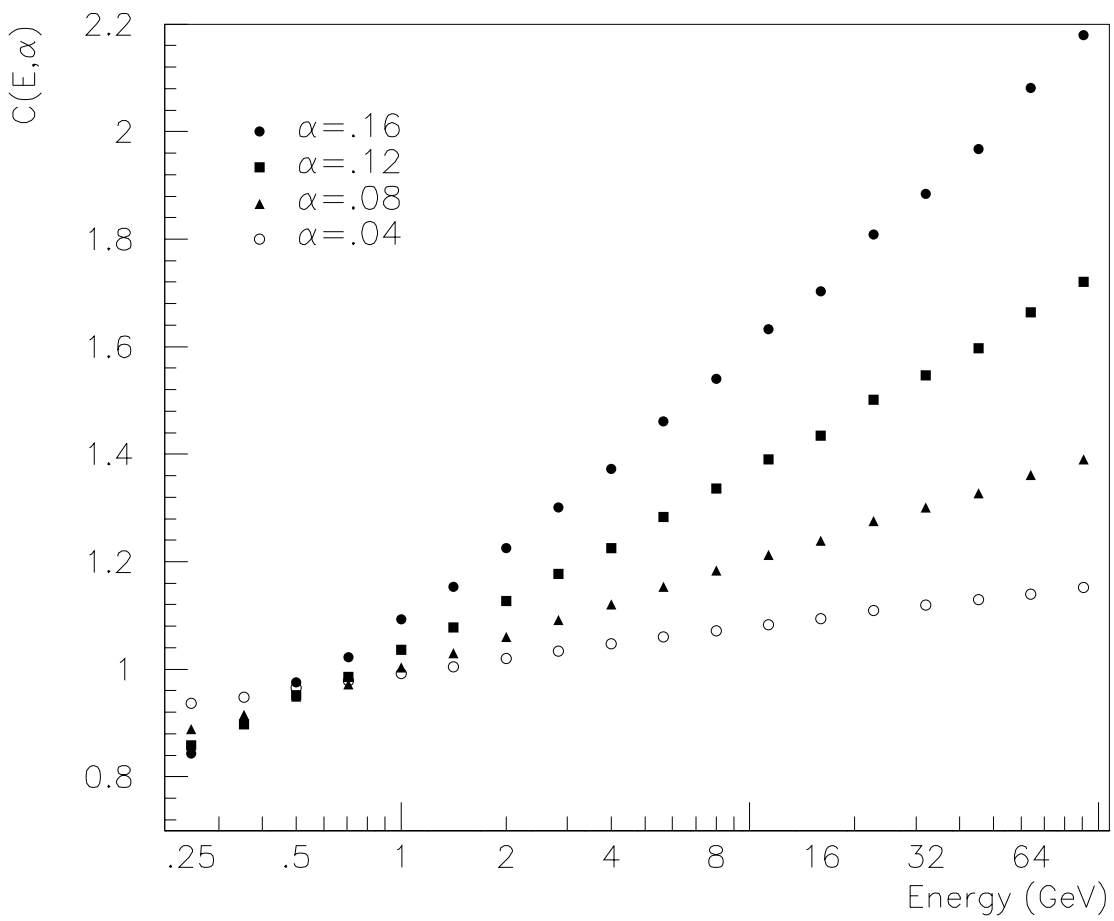

Figure 67. The non-linearity correction ($C(E,\alpha)$) for electrons as a function of energy for different absorptions (α)

6.1.2 Corrections for Photon Response

Once electron response was well-understood, modeling the response of the lead-glass to photons was simple. Photons convert to an electron-positron pair at some depth in the glass governed by the probability distribution,

$$P(t) = \frac{9}{7} e^{-\frac{7}{9}t},$$

where t is the conversion depth in radiation lengths. The spectrum of produced electrons is given by the Coulomb-corrected Bethe-Heitler spectrum [70],

$$P(\epsilon) \approx \epsilon^2 + (1-\epsilon)^2 + \epsilon(1-\epsilon)\left(\frac{2}{3} - a\right),$$

where the electron energies are ϵE and $(1-\epsilon)E$ and a is a material-dependent constant, $a \approx 0.0297$ in the lead-glass. This information is sufficient to allow transformation of the EGS electron shower information into a simulation of photon response also.

Although the showers from photons began deeper in the lead-glass, they had lower-energy primary constituents than showers from electrons of equal energy. Thus, the energy correction for photon showers was close to unity and relatively independent of α. Nevertheless, it was important to include this several percent correction in the data. Consider the response to an average photon shower, $I_\gamma(E, \alpha, t_0)$, in terms of the electron response,

$$I_\gamma(E, \alpha, t_0) = \langle I_e(\epsilon E, \alpha, t_0) + I_e((1-\epsilon)E, \alpha, t_0)\rangle_\epsilon ,$$

where t_0 is the photon conversion depth, and $\langle \cdots \rangle_\epsilon$ denotes an average, weighted by the Bethe-Heitler spectrum, over all ϵ. The average photon response is then given by

$$\langle I_\gamma(E, \alpha)\rangle = \langle\ \langle I_e(\epsilon E, \alpha, t_0)\rangle + \langle I_e((1-\epsilon)E, \alpha, t_0)\rangle\ \rangle_{\epsilon, t_0} ,$$

where the averages without subscripts are understood to be over all shower distribution functions. From the definition of $C(E, \alpha, t)$ above, one can write this as

$$\frac{1}{R(E,\alpha)} \equiv \frac{\langle I_\gamma(E,\alpha)\rangle}{\langle I_e(E,\alpha)\rangle} = \left\langle \frac{\epsilon C(\epsilon E, \alpha, t_0) + (1-\epsilon) C((1-\epsilon)E, \alpha, t_0)}{C(E, \alpha, 0)} \right\rangle_{\epsilon, t_0} ,$$

which can be easily computed from the already tables of $C(E, \alpha, t)$ already extracted from the EGS Monte Carlo showers. This $R(E, \alpha)$ correction function is applied to photon cluster energies in the data. Plots of $R(E)$ for various values of α are shown in Figure 68.

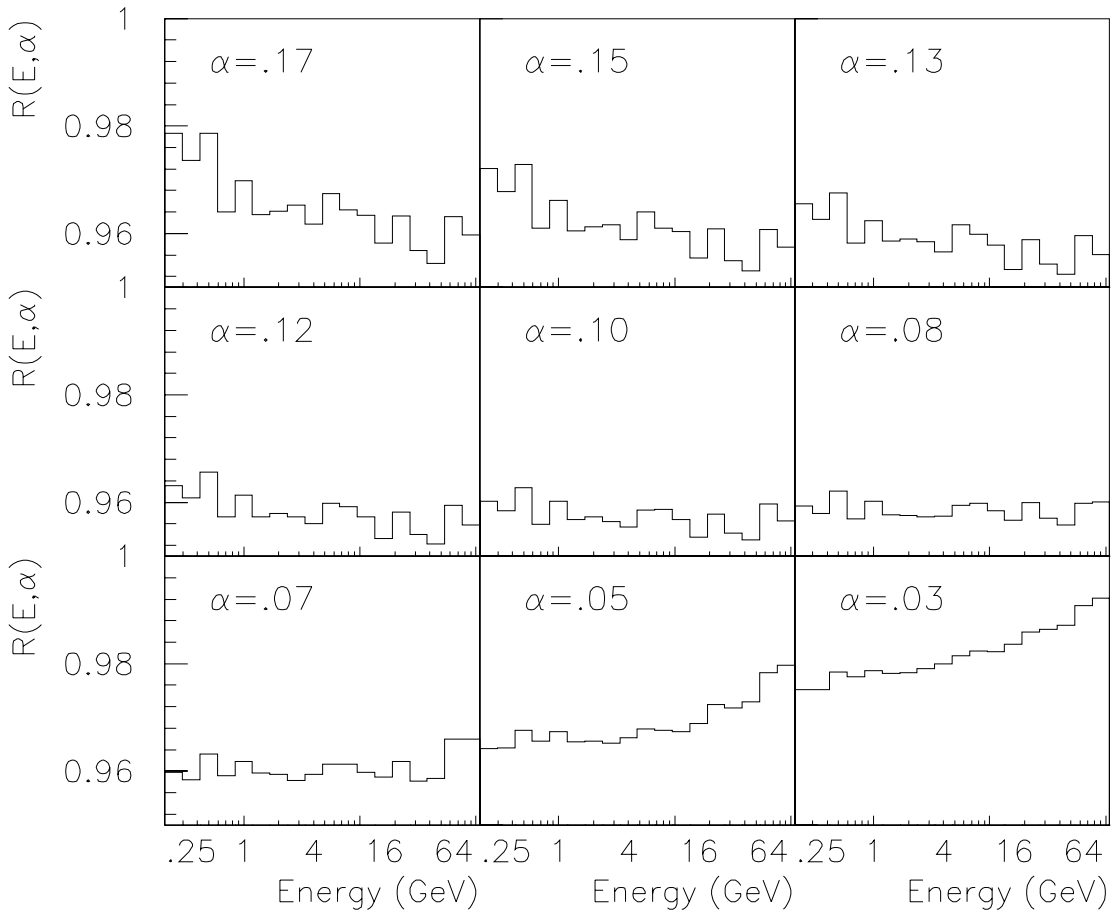

Figure 68. The electron-photon response difference ($R(E,\alpha)$) as a function of energy for different absorptions (α)

6.1.3 Photostatistics

Another contribution to the response of the lead-glass was the smearing of the expected response due to statistical fluctuations in the number of electrons produced at the photocathode of the phototubes. This smearing could be quantified in a straightforward manner using the flasher signal (Section 2.2.2). By measuring the RMS fluctuation observed in the response to this constant signal, it was possible to extract a conversion between the number of photoelectrons and the light output at the phototube. A different g_0 (defined in Equation 6.1) could then be determined for each lead-glass block, and the mean g_0 was 536 photoelectrons/GeV. Note that this constant is independent of α, the phototube gain and the ADC response.

Recall that the number of photoelectrons produced in a shower is related to the energy, E, by

$$gEC(E,\alpha,0) = I(E,\alpha,0).$$

The Gaussian width introduced by photostatistics, σ_p, is then

$$\begin{aligned}\frac{\sigma_p}{E} &= \frac{1}{\sqrt{I(E,\alpha,0)}} \\ &= \frac{1}{\sqrt{gEC(E,\alpha,0)}} \\ &= \frac{4.3\%}{\sqrt{EC(E,\alpha,0)}} e^{\alpha(L-s)/2}.\end{aligned}$$

It is clear that the photostatistical smearing gets significantly worse as the block absorption increases, but note that the quantification of this effect relies heavily on the assumption that the α of the block is uniform with depth. However, the most damaged blocks where this uniform α approximation would be most likely to break down tended to have the highest energy clusters. Because of this, the effect of photostatistics was very difficult to observe in the high α blocks, and therefore it was not possible to use the observed photostatistical smearing as a test of the uniform α model of lead-glass response.

6.1.4 Energy Resolution

Plots of the RMS energy resolution from shower fluctuation as measured from the EGS Monte Carlo and from photostatistics as parameterized above are shown in Figures 69 and

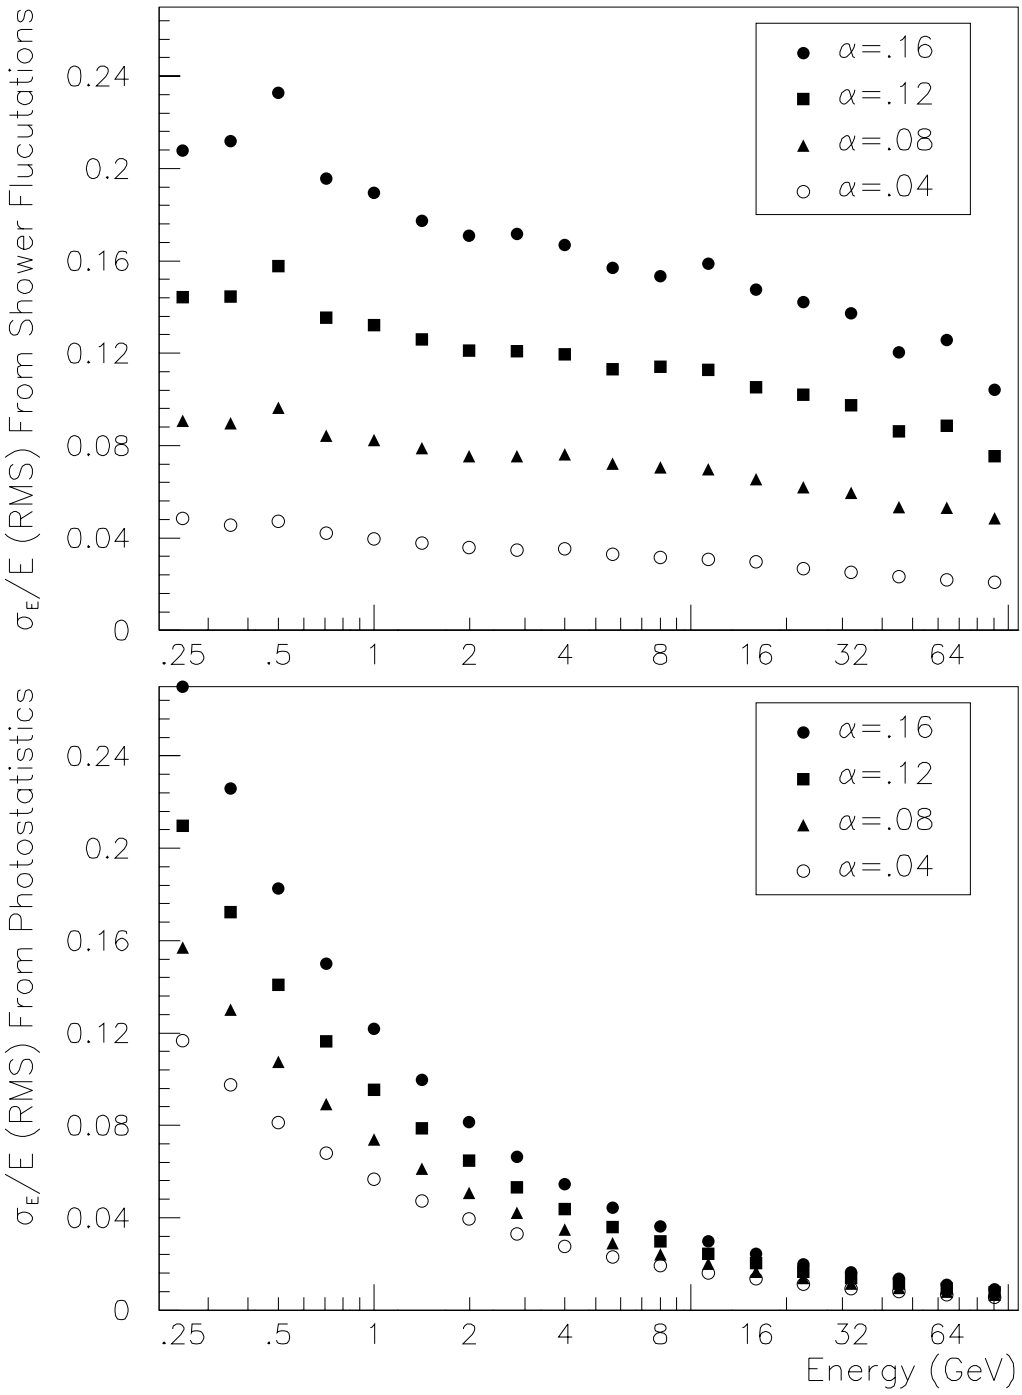

Figure 69. σ_E/E from shower fluctuations and from photostatistical smearing as a function of energy for different absorptions (α)

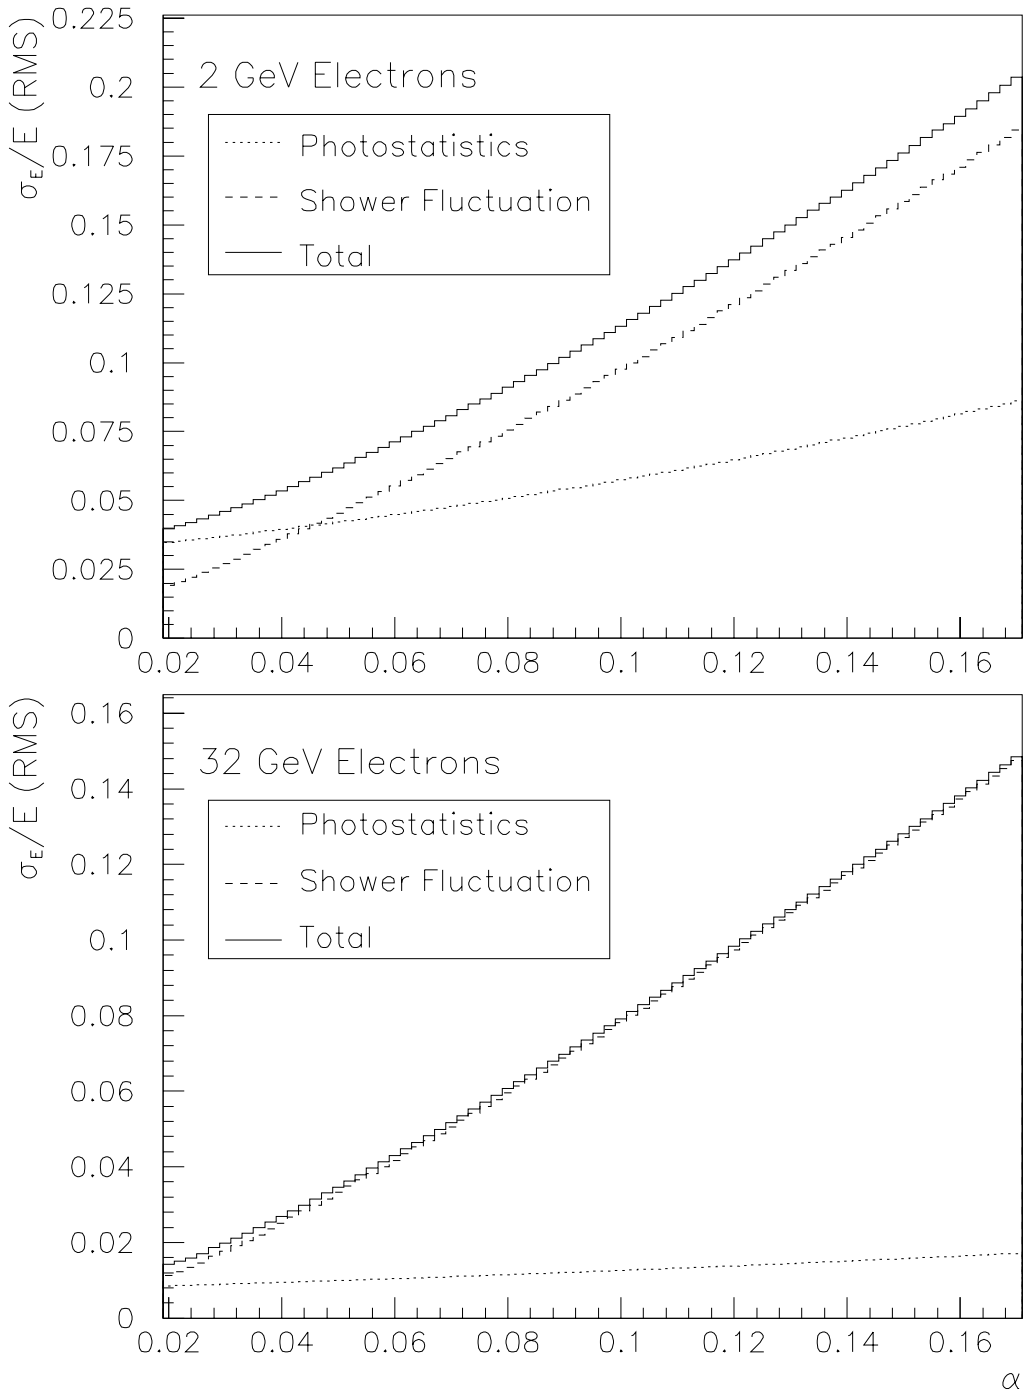

Figure 70. σ_E/E from shower fluctuations and from photostatistical smearing as a function of absorption (α) for different energies

70. Traditionally, calorimeter resolution is parameterized as $\sigma_E/E = a \oplus b/\sqrt{E}$, where the constant term comes from shower fluctuation effects and the second comes from photostatistics. This parameterization, however, makes little sense in E799 for several reasons. First, the "constant" term was not constant as a function of energy, particularly at high absorptions. Second, the photostatistics term did not vary as $1/\sqrt{E}$ but as $1/\sqrt{EC(E,\alpha,0)}$, and $C(E,\alpha,0)$ could vary by almost a factor of 2 over the electron energy range (see Figure 67). Third, there was a significant correlation between electron energy and α in the glass block hit. Since the most damaged blocks were close to the beam, and since the highest energy products of kaon decays tend to travel nearly parallel to the beam because of their high forward Lorentz boost in the lab frame, the highest energy decay products tended to be observed in the blocks with the highest absorptions. This resulted in the average resolution for highest momentum electrons actually being *worse* than that for lower energy electrons,

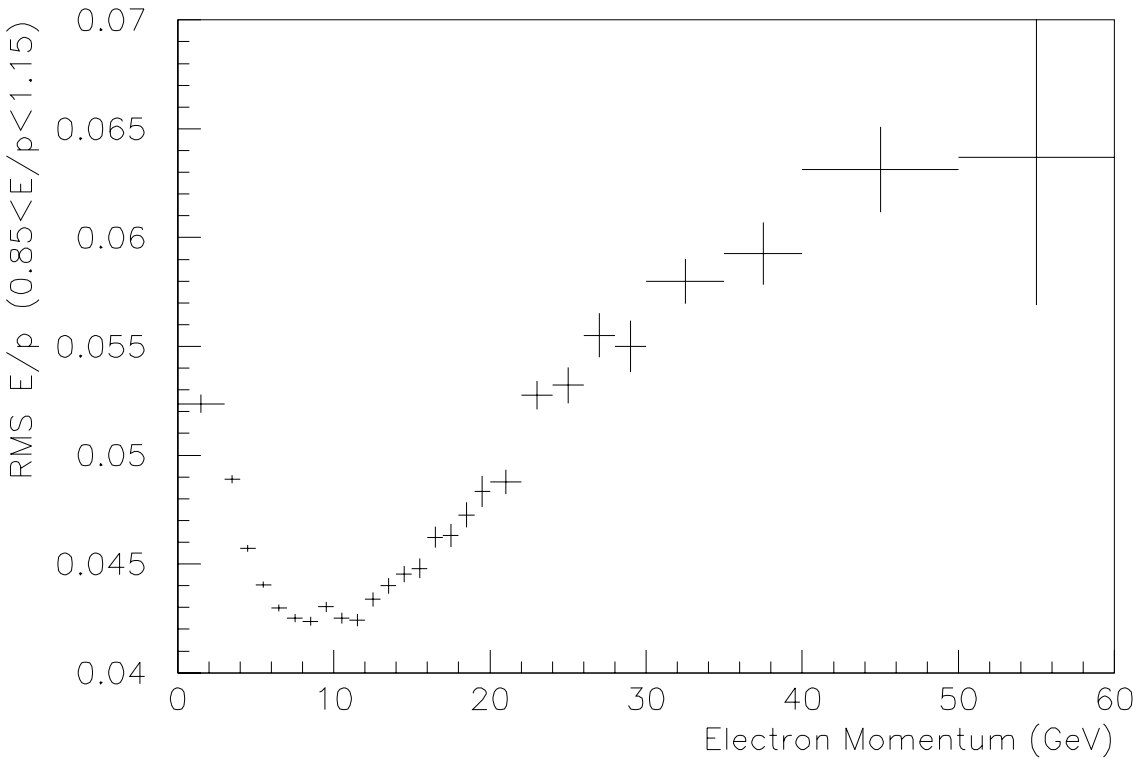

Figure 71. Electron energy resolution versus p for electrons from $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays

as shown in Figure 71. Therefore, I have chosen to quote here an "average" electron energy resolution of 4.4% as measured in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays.

6.2 Determining the Energy of a Lead Glass Cluster

The energy deposited in each lead-glass block was defined to be

$$E_{\text{block}} = g_{\text{block}} n_{\text{block}},$$

where n_{block} is the number of ADC counts in a block and g_{block} is the block gain that was measured in the calibration. For an electromagnetic shower, the bulk of the energy (typically 98%) was contained in a 3×3 array of blocks with the maximum energy block as the center. The baseline energy calculation then assigned the shower energy to be

$$E_{\text{shower}} = \sum_{3\times 3} E_{\text{block}}.$$

However, there were important corrections to be made to this energy calculation. In addition to the non-linearity and electron-photon difference corrections described above, corrections must be made for leakage of shower energy outside of the 3×3 array and leakage off of the edge of the lead-glass array. Corrections were also necessary for the ADC threshold in the glass and for the energy loss due to electron bremsstrahlung in the material upstream of the lead-glass. These corrections will be described in the order in which they were made in the data.

6.2.1 Leakage Off the Edge of the Array

If at least one of the blocks from the 3×3 was missing from the array, either from the inner edge (beamholes) or outer edge of the glass, a correction is made to the energy. If the cluster center was not within one-half block width of any edge of the array, a constant fraction of the observed energy in the cluster was added to correct for the missing block. For a corner block (see Figure 72), the correction was 0.6%, and for an edge block the correction was 1.1%. This additive correction was applied based on a cluster energy with a preliminary correction for threshold and pedestal shift effects.

In the case of a cluster within half a block width of the beamholes, a different approach was used. After all other corrections described in this section, a final scaling factor was applied to the cluster energy, based on the cluster position. The factor was determined using electron calibration data with electrons hitting the pipe blocks [1]. The effect of the this correction for photons is shown in Figure 73.

6.2.2 Rate Dependent Pedestal Shift

Although the pedestals for the lead-glass ADCs were hardware-subtracted, small variations in the actual pedestal were observed as a function of event rate. These could be understood primarily as small variations in the ADC response as a function of the time since the last time the ADC was triggered. These variations caused an effective shift in the pedestal as a function of the time since the last Level 2 trigger. Since this time was recorded for each event, the shift could be corrected for as a function of this time on an event-by-event basis. This shift was studied by examining the variation of the pedestal during full readout

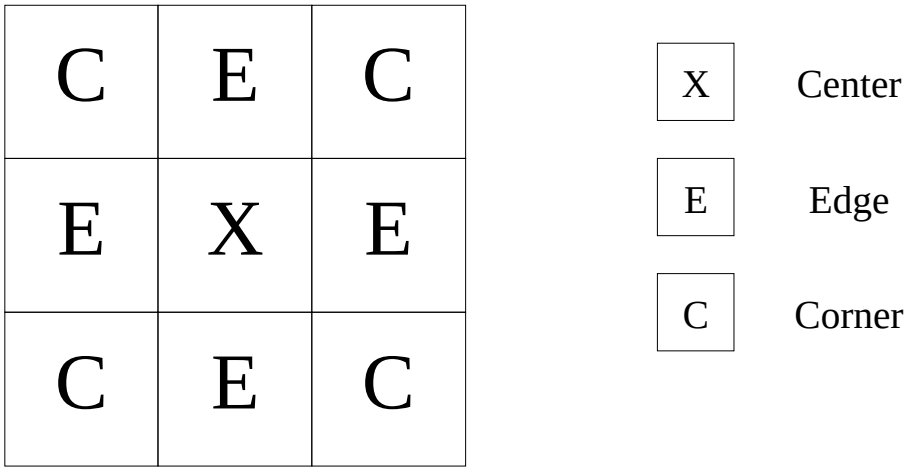

Figure 72. Definition of the center, edge and corner blocks of the central 3×3 blocks of a lead-glass cluster

accidental triggers in four unused ADC channels in each of the ten ADC modules [71]. A correction to the block energies for this pedestal shift was then applied, parameterized as a function of the time since the last Level 2 trigger and the ADC module in which a given block was read out. Energy corrections were typically smaller than 200 MeV per cluster (3 ADC counts per block), and the sign and magnitude of the correction was different for each lead-glass ADC module.

6.2.3 Threshold Correction

Because of the 5 count lead-glass ADC readout threshold, the 3×3 energy was often artificially small when blocks containing small amounts of energy were below threshold. This effect was most noticeable in low energy clusters. A cluster energy correction was applied as a function of energy and the effective lead-glass threshold after accounting for pedestal shifts. The correction was determined using special electron calibration data taken

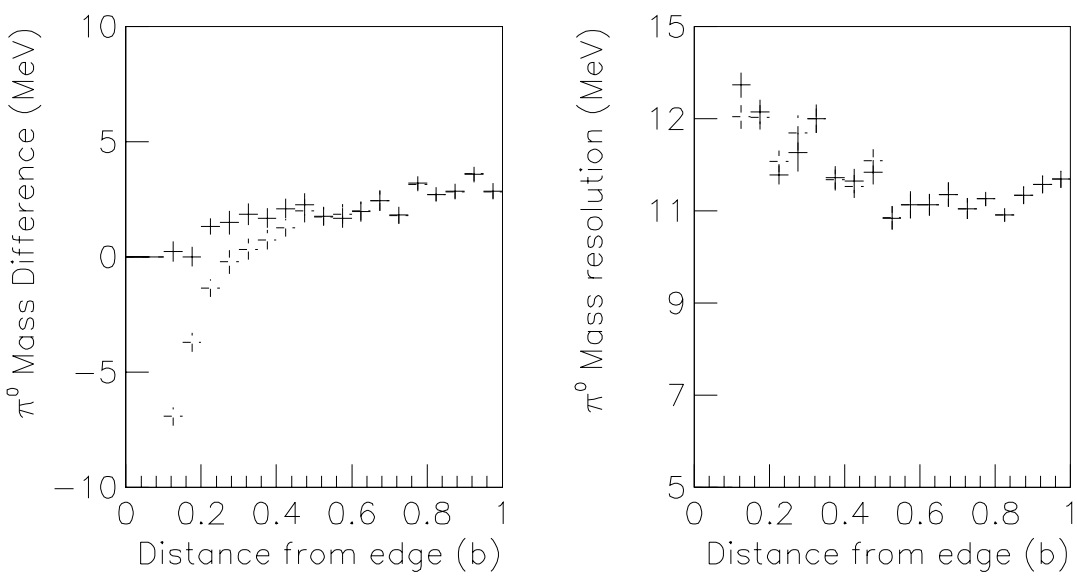

Figure 73. $m_{\gamma\gamma} - m_{\pi^0}$ and the $m_\gamma\gamma$ resolution in $K_L \to \pi^+\pi^-\pi^0$ decays before and after the energy correction in the inner half of the pipe blocks [1]. The original masses and resolutions are shown in dotted lines with the corrected in solid, plotted as a function of distance from the block edge in units of block widths.

with the lead-glass ADCs in full-readout mode [71]. Corrections were typically no greater than 50 MeV per cluster.

6.2.4 Leakage Outside of the 3×3

The EGS Monte Carlo predicts that the fraction of energy of a shower outside the center 3×3 of a lead-glass cluster was about 2.2%, and roughly independent of energy. This was checked in the data with electron clusters, and a significant deviation from this behavior at low energies, below 4 GeV, was observed. For example, the correction for 2 GeV electrons was 2.8%, and for 3 GeV electrons it was 2.5%. This correction was applied to both photon and electron clusters, and was applied independent of cluster position in the lead-glass block.

6.2.5 The Non-Linearity Correction

The non-linearity correction discussed above was applied at this point. Recall that

$$EC(E, \alpha, 0) \propto I(E, \alpha, 0) \propto \sum_{3\times 3} E_{block}.$$

Therefore, one wanted to divide the sum of the block energies by $C(E, \alpha, 0)$ in order to get the non-linearity correction. Because $C(E, \alpha, 0)$ was potentially large, and because it depended on the incident shower energy and not the observed energy, it was necessary to estimate the energy to be used for the correction. This estimation was performed using the power law fit to $C(E, \alpha, 0)$ described below in Section 6.3.

6.2.6 Electron-Photon Difference

At this point, the correction for the response difference of photons from electrons, $R(E, \alpha)$ of Section 6.1.2, was made if the cluster was not matched to an drift chamber track. This correction was fairly small, less than 5%, and relatively independent of energy.

6.2.7 Bremsstrahlung Correction

Clusters matched to tracks were treated as electron candidates. In passing through the material in front of the lead-glass array, primarily the B and C banks, incident electrons

Table 20. Glass Calibrations in E799

Date	Calibration/Cure	Fraction of Two-Electron Data
10/10	Calibration 27	7%
10/15	Cure (13 hours)	
10/21	Calibration 28	16%
10/29	Cure (12 hours)	
11/06	Calibration 29	12%
11/12	Cure (12 hours)	
11/14	Mega-Cure (72 hours)	
12/02	Calibration 30	26%
12/03	Cure (14 hours)	
12/16	Calibration 31	25%
12/18	Cure (12 hours)	
12/23	Calibration 32	14%
01/08	End of Run	

lost some of their energy in the form of low energy bremsstrahlung photons. Although these photons landed on top of the electron cluster, they decreased the effective cluster energy slightly because these low energy photons showered earlier in the glass. An increase in energy of 0.28% was applied to all electron clusters to compensate for this effect.

6.3 Electron Calibration Procedure

The lead-glass calorimeter was calibrated using a sample of electrons produced in special calibration runs which were described in Section 3.3. Using this sample, measurements were made of both the gain and absorption in the lead-glass blocks. Six calibrations were performed during E799; dates, fractions of the data in which a particular calibration was used, and other significant events in the history of the glass are shown in Table 20. In between

calibrations, the variations in gains were tracked using the flashers (see Section 2.2.2); no allowance was made for variations in the absorptions between calibrations.

The calibration procedure was essentially to take the ensemble of all electrons with clusters centered in a given block and to fit the observed energy, after all corrections, to a power law in the measured electron momentum, i.e.

$$E = Gp^{\delta},$$

where G represents a gain factor at 1 GeV, and δ the power-law dependence. The power-law coefficient δ contained any residual non-linearity in the glass response which to be related to uncorrected absorption. Consider the effect of correcting a block with absorption α as though it had an absorption α'. Then the extracted energy, E', would be related to the correctly-extracted electron energy, E, by

$$\frac{E'}{E} = \frac{C(E,\alpha,0)}{C(E,\alpha',0)}$$

If $C(E,\alpha',0)$ can be well-fit by a power law dependence, as suggested by Figure 67, then

$$\delta = \beta(\alpha) - \beta(\alpha'),$$

where

$$C(E,\alpha,0) \approx E^{\beta(\alpha)}.$$

A previous analysis [55] determined the best-fit power law coefficient (for undamaged blocks) to be given by

$$\beta(\alpha) \approx .0058 + .9000\alpha,$$

and this parameterization is used in the calibration procedure.

Although this fit was good to approximately 5% for the reasonable energy range in most blocks, in the high energy response for the most damaged blocks, it could deviate from $C(E,\alpha,0)$ by as much as 30% (Figure 74). However, even such a large deviation from power-law behavior was not fatal to the fitting procedure because the fit procedure was inherently iterative. Recall from above that all clusters that were *centered* in a given block were grouped together for the calibration. Obviously, however, the energy of a cluster

depends on more than just the response for the center block. By iterating this fitting process, both the multi-block nature of clusters and the deviations from a power-law of $C(E, \alpha, 0)$ could be accommodated.

Because of limitations in statistics, not all parameters could be extracted from fits to the available data, and some external input had to be used. In particular, the amount of data in the "high range" of the glass bilinear ADCs (described in Section 2.2.2) was extremely limited in electron calibration data. Therefore the "high range slopes", defined as the ratio between the high range and low range gains, were determined from studies done using

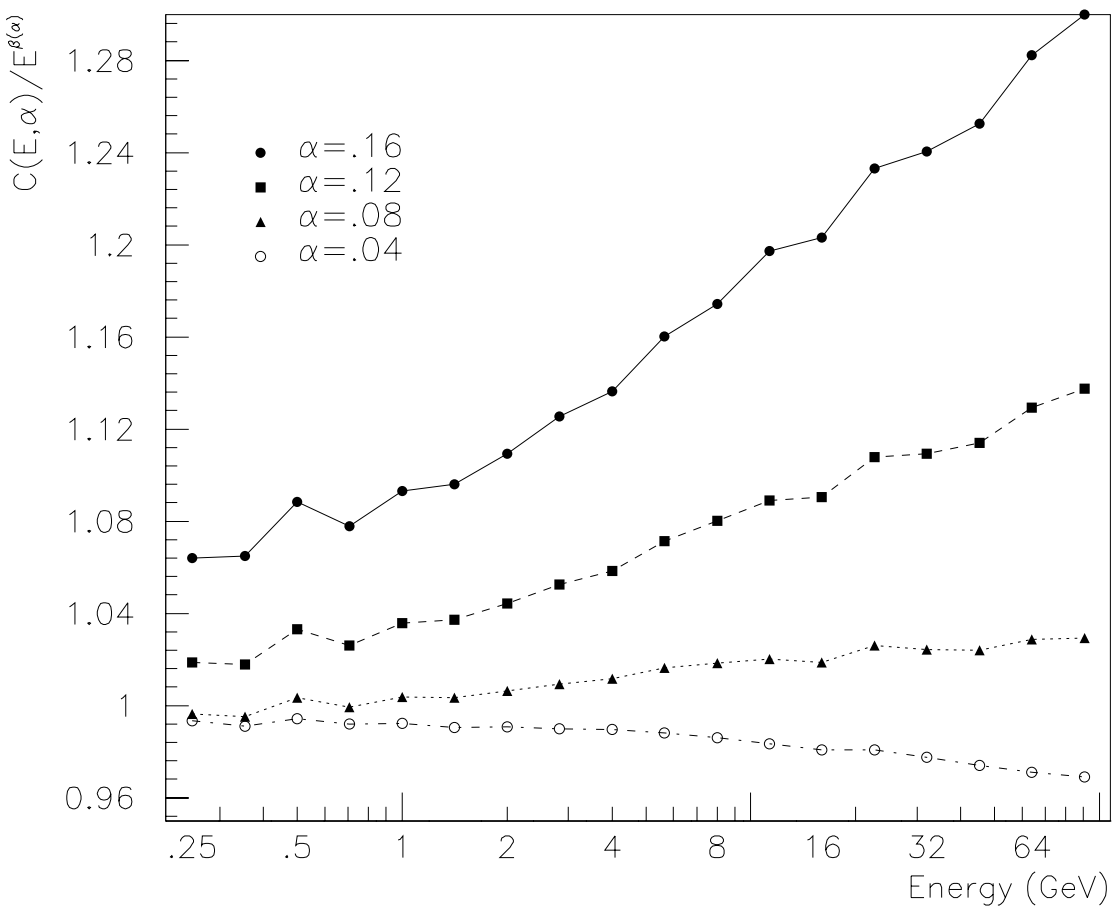

Figure 74. Deviations of the non-linearity correction ($C(E, \alpha, 0)$) from its power law parameterization

$K_L \to \pi^{\pm} e^{\mp} \nu$ data taken in E773. Similarly, the electron calibration data did not provide a sufficient range of energies or statistics for determination of the non-linearity in blocks far from the beam holes; therefore, for any block greater than 50 cm from a beam hole, the absorption was determined from $K_L \to \pi^{\pm} e^{\mp} \nu$ studies in E731. Although, in principle, these blocks could have been damaged after the E731 determinations of the absorptions, any radiation damage in these high radius blocks should have been small. The mean α, by radius in the lead-glass, is shown in Figure 75; the observed discontinuity at the border between determinations of α from the electron data and from E731 $K_L \to \pi^{\pm} e^{\mp} \nu$ data was small for all calibrations.

It is interesting to use the time-evolution of the glass absorption constants to study the radiation damage in the lead-glass. It is clear that the central blocks of the lead-glass degraded steadily and dramatically during the running of E799, despite the attempts to cure the damage with UV light. A plot of the rate of increase in α in each block is shown in Figure 76. The damage suffered by the lead-glass had several important effects in the E799 data. First, the resolution of the most damaged blocks degraded dramatically as the absorption increased and the effects of shower fluctuations were amplified. Second, because the damage also increased the glass non-linearities, the lead-glass trigger thresholds

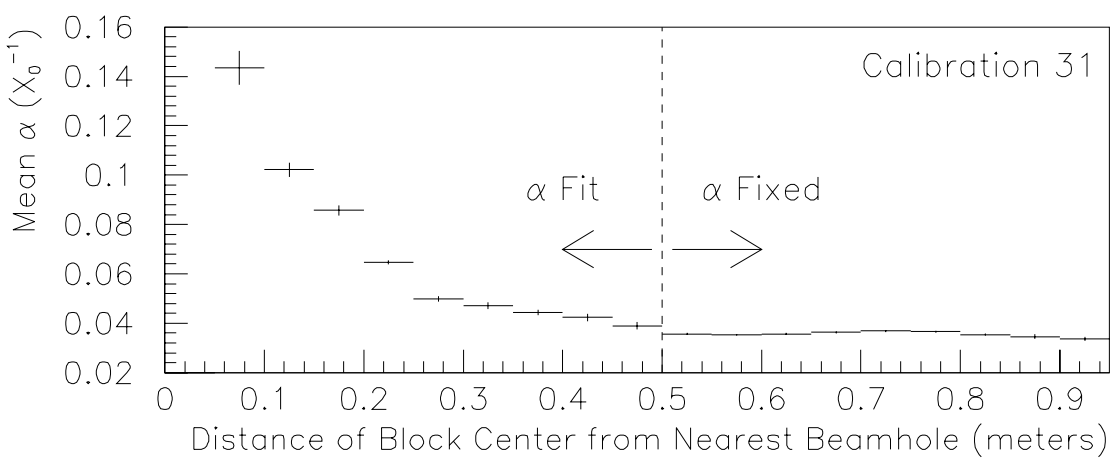

Figure 75. Mean absorptions (α) of lead-glass blocks as a function of distance from a beam hole

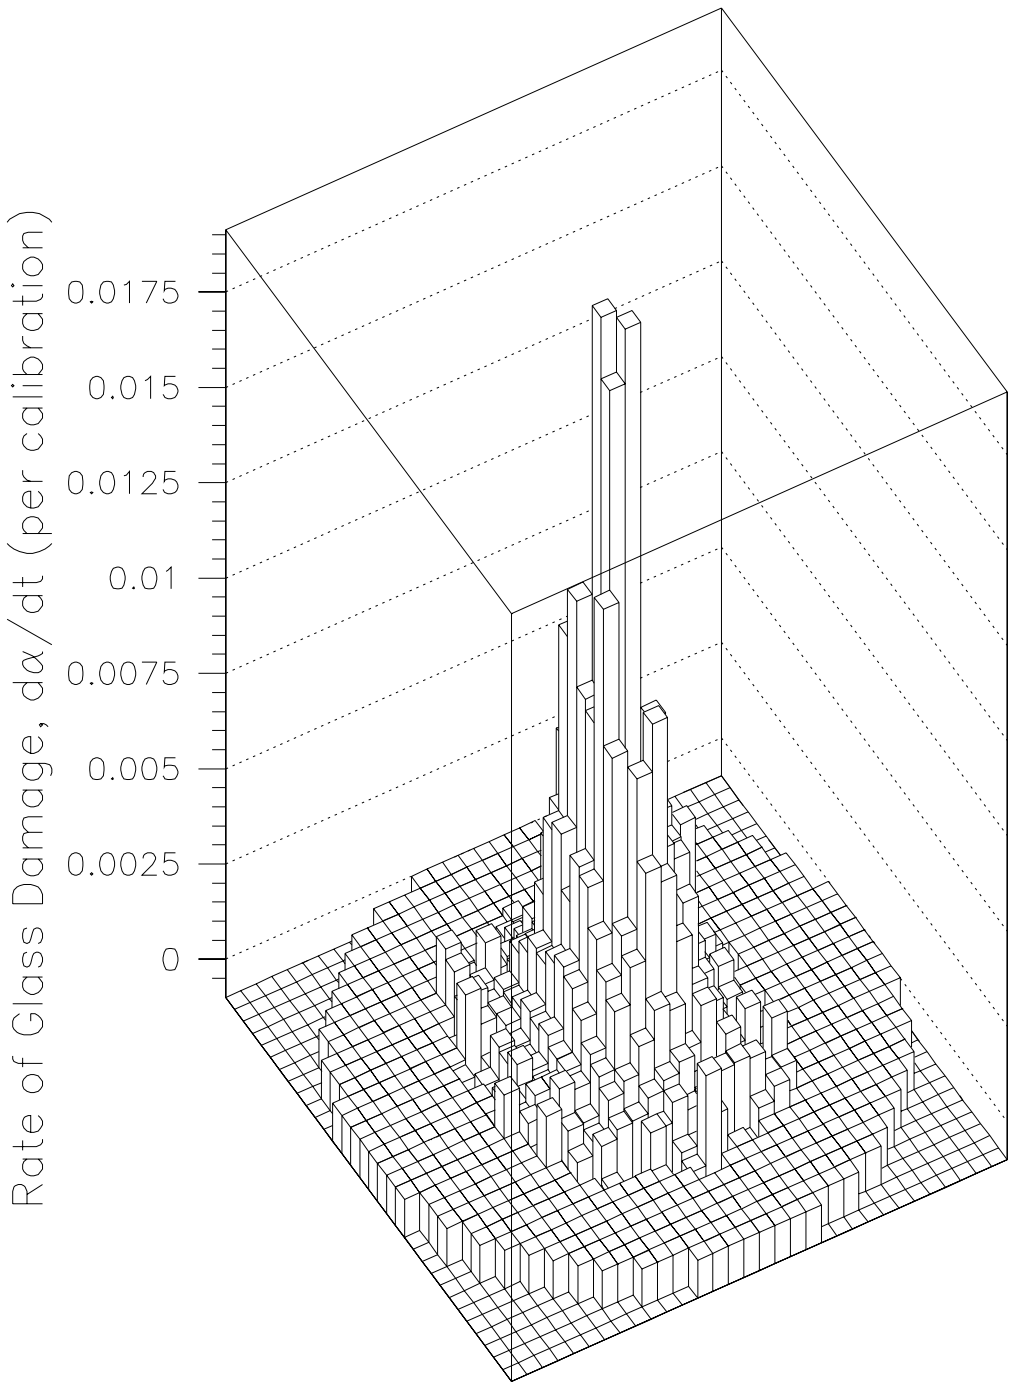

Figure 76. The rate of increase of absorption (α) for each less-glass block

were affected. In particular, the distribution of total kaon energies required to satisfy the E_tHigh Level 1 trigger source became broader and more dependent on the fraction of energy deposited in the most damaged blocks. Third, because of the extreme damage experienced in the pipe blocks, these blocks became less useful for energy measurements, particle identification and position measurements. Furthermore, because these blocks so often contained a photon or electron in high-multiplicity final states, the consequences for acceptance and backgrounds were severe.

One particularly desirable feature of using the electron calibration data for glass calibration is that the electron energy spectrum roughly matches that of high final-state multiplicity Kaon decays, such as $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, as shown in Figure 77. This feature lessens the dependence upon details of the model used to simulate lead-glass, particularly the assumption of uniform absorption as a function of depth. Because the data probes roughly the same region of shower depth as is present in the calibration electron sample, the value of α extracted in the fit will correspond roughly to the effective α at the shower depths of interest.

Comparisons of the results of the glass response after calibration and corrections with the Monte Carlo simulation of this response can be found in Section 7.3.2.

The same data used to calibrate the lead-glass was also used for a calibration of the lead-glass adders. The goal was for the sum of ADC values in the lead-glass blocks comprising a certain adder to be equal to the measured response in that adder. Gain constants for the adders were determined by plotted the ratio of these two quantities for adders with blocks hit by calibration electrons. The adder gains were then adjusted so that this distribution of this ratio for each adder was centered at 1. This calibration allowed this comparison of the adder response to the lead-glass response to be made for each cluster in the lead-glass in order to identify out-of-time clusters.

6.4 Cluster Position Measurements

Two algorithms were used in E799 to measure the position of electromagnetic particles using lead-glass information, using essentially the same method. Once the central block of

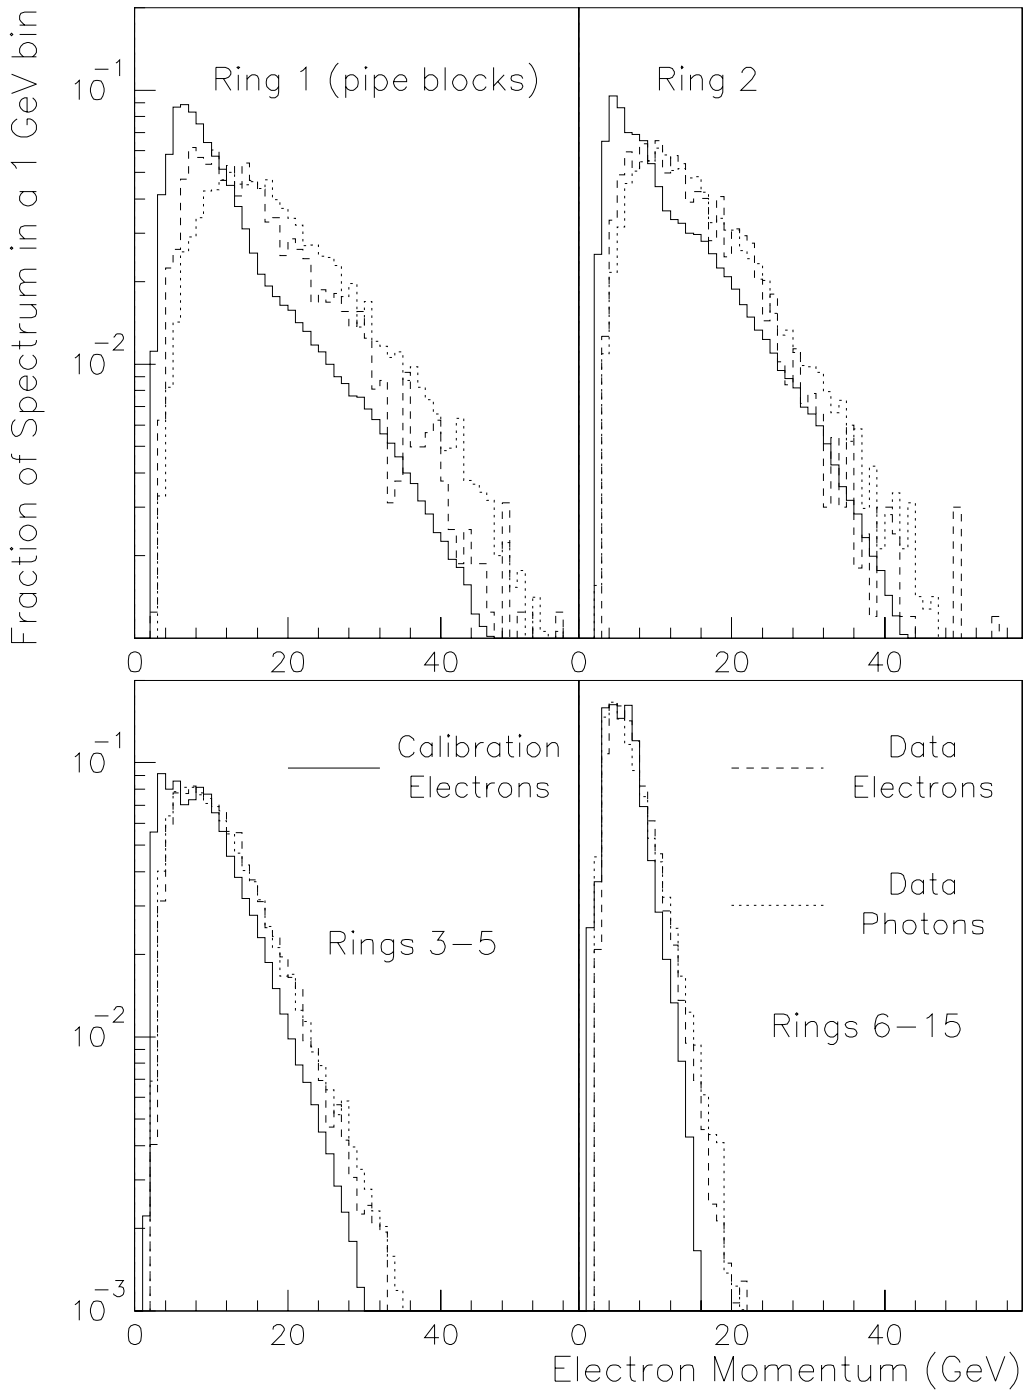

Figure 77. Energy spectrum of calibration electrons in different regions of the lead-glass, compared with the spectrum of final state particles from $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$

a cluster was determined by searching for the maximum block energy in each cluster, the block energies in the surrounding 3×3 of the cluster were used to extract the position of the electron or photon within the central block.

For clusters with all nine blocks of the 3×3 existing (i.e., clusters not centered in the pipe blocks or in outer edge blocks), the cluster position in the $\hat{x}$ and $\hat{y}$ directions were found independently using the following technique. The block energies of the 3×3 were summed in columns and in rows, as shown in Figure 78. The maximum row and the maximum column ratio were then used to determine the cluster position by using this ratio as the index to a lookup table. This table was generated using calibration electrons whose positions were known by extrapolating the downstream track segments in the drift chambers to the lead-

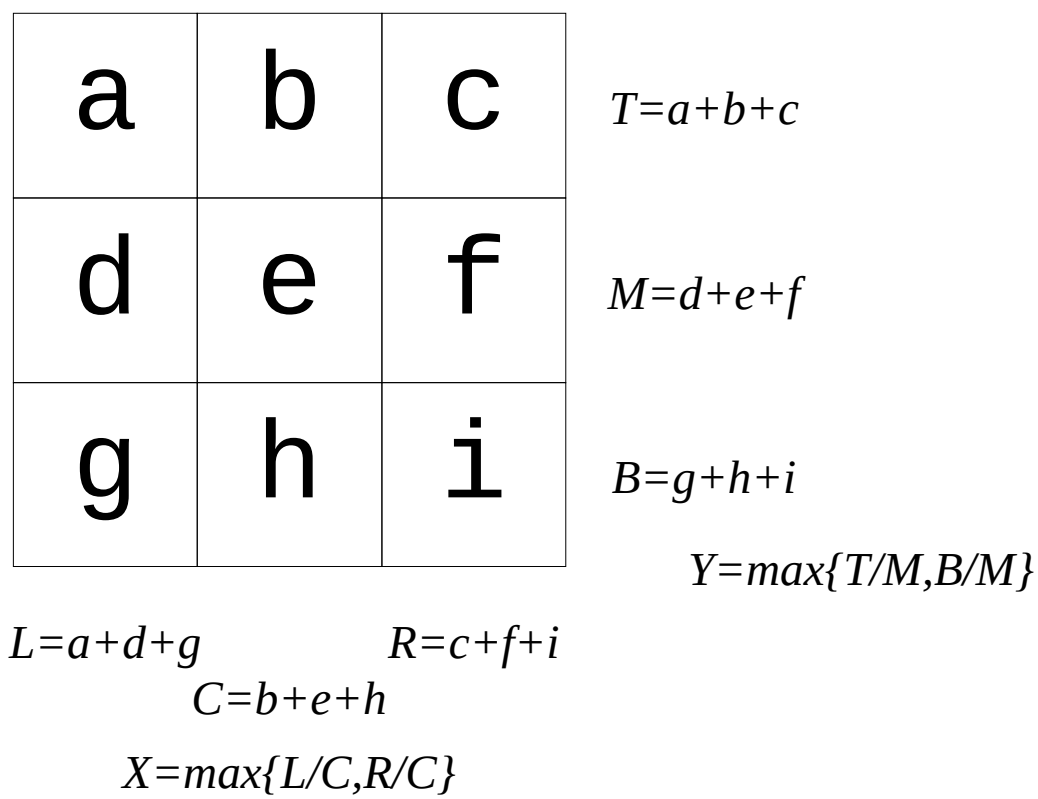

Figure 78. Energy ratios used in the lead-glass cluster position algorithm

glass. The z-position of the lead-glass, for the purposes of this extrapolation, was taken to be 20.3 cm downstream of the face of the block, which is the shower maximum position of a 20 GeV electron shower, since the shower will be centered closer to the shower maximum of the cluster rather than the position on the glass face at which the electron or photon enters the glass. The RMS cluster position resolution is, on average, 2.8 mm, and gets worse as the cluster position moves from the edge of a glass block to the center.

Since the pipe blocks were such an important component of the lead-glass acceptance, separate position algorithms were developed for use in these blocks. A similar method was used, although the missing blocks in the 3×3 of the cluster caused the position of the clusters in the two dimensions to become coupled, thus complicating the problem immensely. Details of this heroic effort are discussed elsewhere [1]. These improved algorithms were developed using calibration electrons within one half-block of the beam hole. (Electrons further from the beam holes had position resolutions comparable to those in the rest of the glass.) For pipe blocks sharing a side with the beamhole, the RMS position resolution in the direction perpendicular to this shared side was 6.0 mm, and the resolution was 4.4 mm in the parallel direction. For pipe blocks sharing only a corner with the beamhole, the RMS resolution was 4.4 mm in both directions.

This position effectively measured the distance of a cluster center from the side of a lead-glass block. However, there were non-trivial variations in the locations of these edges within the array due to a spread in the block sizes. A correction was applied to the cluster position based on both measured block sizes and information from electron data which was used to locate offsets in the block edges introduced during the stacking of the array.

6.5 Fusion of Clusters

One effect which could potentially cause gross mismeasurement of cluster energies and positions was the overlap of multiple showers in the lead-glass. If two electromagnetic showers were located sufficiently close to each other in the glass, it was possible for only a single local maximum to be found in the lead-glass because of the coarse transverse segmentation in the array. This effect was referred to as the "fusion" of clusters. In practice,

two clusters separated by less than one block-width almost always fused in the glass, and often clusters separated by as much as two-block widths could still fuse. Fusion is not a serious problem in electron clusters, since the position and momentum of the electron are measured using the information from the electron's track. Fake electrons, however, could result from a cluster fusion between a charged pion's hadronic shower and an extra cluster which caused the pion shower to appear sufficiently energetic to pass the E/p electron identification cut. For photons, where the lead-glass provides the only measurement of energy and position, fusions could be even more disastrous. It is worthwhile, then, to understand what tools could be used to identify fusions.

The signature of fused clusters in the lead-glass was a transverse profile of energy in the lead-glass that appears inconsistent with that of a single electromagnetic cluster. There were three different algorithms used to detect fusions in E799: a search for extra energy in the 5×5 ring of blocks surrounding the cluster center, comparison among the column and row ratios for inconsistencies, and a comparison between energies of corner and side blocks in the cluster 3×3.

The first and most powerful algorithm looked for anomalously large energy deposits in the 5×5 ring of blocks around the cluster center. First, it discarded any blocks known to belong to the 3×3 of another cluster. Then, all adjacent pairs of the remaining blocks in the 5×5 were checked for a large energy deposition. Any event with greater than 0.8 GeV in a pair of blocks was tagged as a fusion candidate. Figure 79 shows an example of fusion candidate found with this algorithm.

The second check was referred to as the "column" fusion algorithm, and was based on the same technique used to find lead-glass positions. The energies in the cluster 3×3 were summed in rows and columns as in the position algorithm (Figure 78). However, in this case, information from both the minimum and maximum row or column ratios was used. The cluster was identified as a fusion candidate if

$$\begin{aligned} R_{max}(R_{min} - 0.02) &> 0.025 \ \ (\text{if } E_{cluster} > 20 \text{ GeV}) \\ &> 0.015 \ \ (\text{if } 20 \text{ GeV} > E_{cluster} > 4 \text{ GeV}), \end{aligned}$$

where R_{max} is the maximum column or row ratio, and R_{min} is the minimum column or

.	0.1	0.3	.	.	.
.	0.1	1.2	0.6	0.1	.
0.1	0.8	1.9	3.3	0.5	0.2
0.0	0.1	0.1	2.1	0.3	0.2
.	.	0.1	0.3	0.1	.

Figure 79. A cluster fusion found in the $K_L \to \pi^0\pi^0\pi^0$,$\pi^0 \to e^+e^-$ data sample, located because of the excess energy in the cluster 5×5. A subset of the lead-glass array is shown with the energies shown in GeV for each block above threshold. The 3×3 of the cluster is outlined in a heavy box, along with the neighboring blocks in the 5×5 ring which contain more than 0.8 GeV

.	.	0.1	0.1	0.1	.
.	0.1	0.3	0.4	0.4	.
.	0.5	2.5	0.9	0.2	.
0.0	0.1	0.6	0.9	0.1	.
.	.	0.3	0.1	0.1	0.1

Figure 80. A cluster fusion found in the $K_L \to \pi^0\pi^0\pi^0$,$\pi^0 \to e^+e^-$ sample, using the "column" fusion algorithm. The cluster 3×3 and the columns which fail the column ratio cut are outlined in heavy boxes.

row ratio. In the case of separately identified clusters with overlapping center 3×3 blocks, these blocks were excluded from column or row ratios. A fusion candidate found using this technique is shown in Figure 80.

The final algorithm searches for "corner" fusions, events with anomalously large deposits of energy in the corner blocks of the 3×3 (see Figure 72). For this study, the energy in each corner block not part of the central 3×3 of another cluster was compared with the energy of the adjacent edge blocks. If any corner block was found with more than twice the

.	0.1	0.1	0.1	.	.
.	0.0	3.5	3.6	0.0	.
.	0.2	6.5	1.6	0.1	.
.	0.1	0.1	.	0.1	.
.	.	.	.	.	.

Figure 81. A cluster fusion found in the $K_L \to \pi^0\pi^0\pi^0$,$\pi^0 \to e^+e^-$ sample, using the "corner" fusion algorithm. The 3×3 of the fused cluster and the corner and edge block with suspect energy depositions are outlined in heavy boxes.

energy of an adjacent edge block, this event was classified as a fusion candidate. Such an event is shown in Figure 81.

6.6 Conclusion

The model used to describe the lead-glass calorimeter has been discussed, and applications to the glass calibration and data corrections have been shown. Algorithms for extracting cluster positions and rejecting fused clusters were also discussed.

The glass performed reasonably well in E799. However, the blocks near the beamholes of the array suffered serious radiation damage during the run and became increasingly less useful during the course of the data taking as their resolution and electron identification capabilities degraded.

With the calibration of the glass completed, the detector was sufficiently well-understood to allow measurements on the data sample. Information obtained in calibration also allowed for a realistic simulation of the lead-glass response, which is discussed in the next chapter.

CHAPTER 7

MONTE CARLO SIMULATION OF THE DETECTOR

Dans les champs de l'observation le hasard ne favorise
que les esprits préparés
—*Louis Pasteur*

This chapter discusses the Monte Carlo simulation of the E799 detector. Monte Carlo simulations are an important tool commonly used in experimental high-energy physics to increase understanding of physics processes and detector response. In E799, the Monte Carlo was used in the early stages of the experiment to optimize detector design, and to create experimental triggers. The Monte Carlo also provided a measure of the loss in acceptance from experimental cuts, and thus was valuable in helping to choose these cuts. In the extraction of the branching ratio of $\pi^0 \to e^+e^-$, the Monte Carlo was used to relate the number of observed decays to the rate at which the decay occurred. If we observe $N(\pi^0 \to e^+e^-)$ decays in the final event sample, the decay rate is given by

$$\frac{,(\pi^0 \to e^+e^-)}{,(\pi^0 \to \text{all})} = \frac{N(\pi^0 \to e^+e^-)}{N(\pi^0 \to \text{all})} \times \frac{1}{A(\pi^0 \to e^+e^-)}.$$

$A(\pi^0 \to e^+e^-)$ is the "acceptance" for the decay $\pi^0 \to e^+e^-$, or the number of decays observed in the detector divided by the total number of decays which occurred in the decay volume. This acceptance includes effects from both the the response of the detector and also the algorithms used to search for such events. To measure $N(\pi^0 \to \text{all})$, the total number of π^0 decays in the decay volume, a decay of the π^0 with a previously measured branching

ratio was also analyzed in parallel with the final event sample. For the $\pi^0 \to e^+e^-$analysis, the "normalization sample" was $\pi^0 \to e^+e^-\gamma$ decays. The $\pi^0 \to e^+e^-$ rate was then given by

$$\frac{,(\pi^0 \to e^+e^-)}{,(\pi^0 \to \text{all})} = \frac{,(\pi^0 \to e^+e^-\gamma)}{,(\pi^0 \to \text{all})} \times \frac{N(\pi^0 \to e^+e^-)}{N(\pi^0 \to e^+e^-\gamma)} \times \frac{A(\pi^0 \to e^+e^-\gamma)}{A(\pi^0 \to e^+e^-)}.$$

The last term in this expression, the ratio of acceptances, is the the term determined by the Monte Carlo.

An important point to consider is that the Monte Carlo was only used directly in the measurement to predict the *ratio* of acceptances of different decay modes. Therefore, it was desirable to pick a normalization decay which was as similar as possible to the signal decay in order to minimize dependence on the Monte Carlo. For example, in this analysis both the decay and signal mode have two electrons in the final state. Therefore, one expects analysis biases from tracking cuts to be approximately the same for the two decays. Any correction for these biases made in the ratio of acceptances will be small, and thus it is not important for the Monte Carlo to reproduce the details of this bias. This point will be discussed further in Chapter 8.

The Monte Carlo was also used in this analysis to simulate backgrounds to $\pi^0 \to e^+e^-$. This led to a better understanding of the sources of these backgrounds and ultimately gave a prediction of the amount of background to the signal mode in the data sample.

The E799 Monte Carlo incorporated production of the K_L beam, decays of K_L's and daughter particles, the flight of decay products through the detector, and the detector response. This was intended to provide a complete model for what was seen in our detector, including the physics of K_L production and decay. Both data and Monte Carlo generated event samples could be treated with the same analysis procedures for comparison. The details of the mechanisms of each component of the Monte Carlo are supplied below.

7.1 Production and Decay

As an input to the Monte Carlo, the range of generated parent K_L momenta to be used was specified, along with the limits of the decay region in meters from the target. These limits were chosen to be well outside the range observed in the data samples. The shape

of the production spectrum was given by the modified Malensek spectrum discussed in Section 2.1.1. This spectrum depended on the production angle, so the energies chosen varied somewhat as a function of position within the beam. A kaon energy was first chosen in the specified range, using both the K_L spectrum and the decay probability in the specified decay region. Then the K_L was flown from a point within the target through the collimators. If any collimator was hit, the K_L was discarded, and the generation process recommenced. As shown in Figures 55 and 56, this simulation using collimator positions from detector surveys provided an adequate model of the beam shapes in E799. The Monte Carlo energy spectrum of observed kaon decays was also accurate as shown in Figure 99.

If the K_L survived the flight past all the collimators, a decay vertex was chosen, and the initial momenta of the products of the K_L decay were chosen. In the case of the $\pi^0 \to e^+e^-$simulation, this was particularly straightforward since the only degree of freedom for a two-body decay is the decay angle relative to the π^0 boost direction. For $\pi^0 \to e^+e^-\gamma$, the decay was treated as series of two body decays with $\pi^0 \to \gamma\gamma^*$ and then $\gamma^* \to e^+e^-$. The e^+e^- invariant mass and e^+-e^- energy asymmetry for a given decay was chosen according to the Kroll-Wada differential spectrum with a form-factor slope [9][14],

$$\frac{d\rho_0}{dxdy} = \frac{2\alpha_{\mathrm{EM}}}{3\pi}\frac{1}{x}(1-x)^3\left(1+\frac{r}{2x}\right)\sqrt{1-\frac{r}{x}} \times (1+ax)^2(1+y^2),$$

where ρ_0 is the rate, $x = (m_{ee}/m_{\pi^0})^2$, and

$$y = \frac{2(E_{e^+} - E_{e^-})}{m_{\pi^0}(1-x)} \text{ in the } \pi^0 \text{ rest frame.}$$

The form factor slope a was set equal to the CELLO value of (0.0326 ± 0.0026) [27]. Algorithms to implement radiative corrections in both signal and normalization decay modes modified this approach. In the case of the $\pi^0 \to e^+e^-$decay, a routine implementing the analytic calculations discussed in Section 1.3 was used. For the $\pi^0 \to e^+e^-\gamma$ decay, the PHOTOS QED radiative correction Monte Carlo package [72] was used to simulate final-state radiative effects.

7.2 Tracing of Particles

Once the decay products were created the next step was to trace them through the detector. Any π^0's produced in kaon decay were immediately forced to decay because of their very short lifetime ($c\tau \approx 25$ nm). Photons, electrons and charged pions were traced from the decay vertex until they either left the fiducial volume of the detector, hit a photon veto counter, or hit the lead-glass. Muons were traced through the entire detector until the final muon trigger plane after the muon filter.

Tracing through the detector basically consisted of propagating particles in straight lines according to their initial direction. However, there were some possible modifications made in-flight to particle directions and momenta. For example, the decay of a charged pion into a muon and a neutrino between the decay vertex and the lead-glass was simulated. For typical pion momenta and kaon decay vertices, the probability of such a decay was on the order of 3% per charged pion.

Interactions of particles in the detector material were also accounted for in the Monte Carlo. Conversion of photons into e^+e^- pairs according to the Bethe-Heitler spectrum (see Section 6.1.2) was simulated with the final state electrons both produced traveling along the parent photon path. A treatment of multiple scattering was also included for charged particles. For the small amount of material upstream of the lead-glass, the angle of scattering was based on the RMS multiple-scattering angle formula given in Section 2.2.1.1. Although a Gaussian approximation to multiple scattering was sufficient to describe approximately the central 98% of the distribution, the Monte Carlo used a parameterization of the Molière scattering distribution [73] shown in Figure 82 in order to include the effects of the tails in the scattering angle distribution. Bremsstrahlung of electrons was also simulated, again using the Bethe-Heitler cross section.

Some approximations were made to the actual distribution of material in the detector for the simulations of bremsstrahlung, photon conversions and multiple scattering. In the spectrometer simulation, scattering and bremsstrahlung only took place at the drift chambers. So, for example, at Chamber 1, the multiple scattering routine was called using a radiation length equal to the summed radiation length of all the material between the vac-

uum window and the midpoint between Chambers 1 and 2. The wires in the drift chamber itself were treated separately. Each wire was approximated as a square of the same cross-sectional area of the wire, and every charged particle passing through was randomly chosen to either hit one wire or miss all wires, independent of particle position within the cell.

For the simulation of photon conversions, however, material within the spectrometer was grouped differently. Conversions always occurred at the midpoint between chambers, and the amount of material simulated was equal to the total material between chambers. The chamber wires were approximated by a thin sheet of material in this conversion simulation, and only wires upstream of the second sense plane of a given chamber contributed to conversions before that chamber. Material downstream of the spectrometer was grouped at the TRDs and trigger planes for the multiple scattering, bremsstrahlung, and photon conversion simulations. Scattering of muons through the bulk material of the lead-glass,

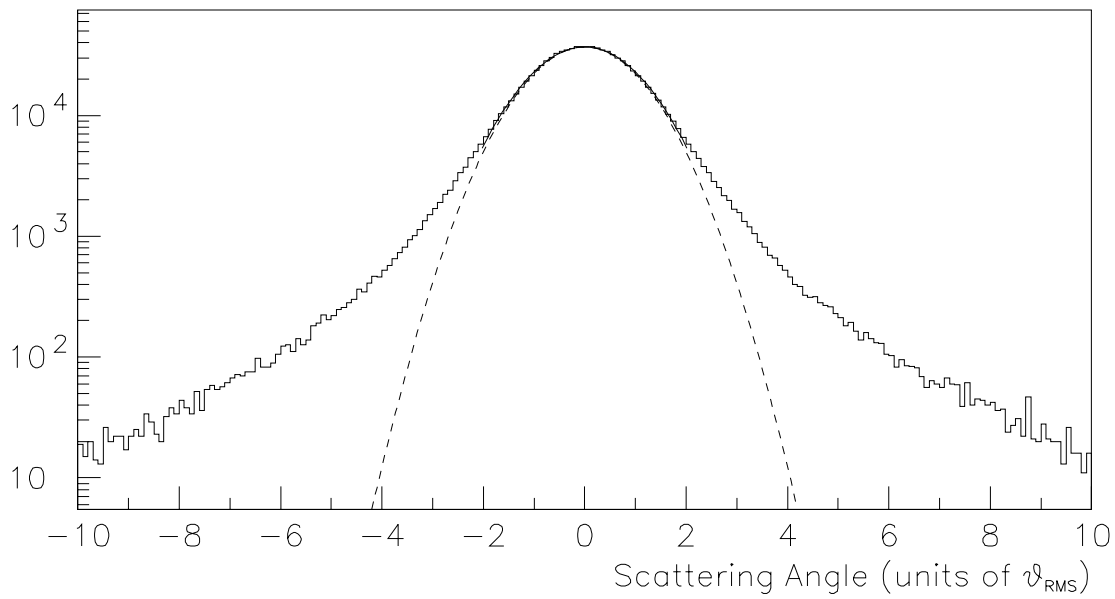

Figure 82. The angular distribution of multiple scattering in the Monte Carlo is shown, compared with the Gaussian approximation of multiple scattering. The parameterization used in the Monte Carlo is based on the scattering theory of Molière as described by Bethe [73].

the lead wall and the muon filter was handled differently, and is discussed in detail elsewhere [1].

A particle was traced only until it hit a detector aperture, such as a photon veto counter or one of the upstream vacuum pipes, or until it escaped the detector volume. The effect of these geometrical acceptance losses varied as a function of the decay vertex location. To test the Monte Carlo, the distribution of decay vertex locations in data and Monte Carlo for $\pi^0 \to e^+e^-\gamma$decays is shown in Figure 83. The reasonable agreement, especially in the upstream region where the dead material in the upstream pipes acts as the limiting aperture, gives some confidence in the Monte Carlo simulation of the geometrical acceptance.

7.3 Detector Simulations

This section describes the detector portion of the Monte Carlo. Of primary importance were the simulations of the drift chambers and of the lead-glass. Also discussed are the photon vetoes and trigger scintillators.

7.3.1 Simulation of the Drift Chambers

As each charged particle was traced through the drift chambers, the position and direction of that particle at the midpoint of the chamber was saved. Using this information, the position of the charged track at each chamber plane was calculated and then smeared according to the chamber resolutions which were input to the Monte Carlo from measurements in the data (see Section 5.3.2). The wire efficiency of the hit wire was then checked (see Section 5.7), and the hit was randomly kept or discarded according to this efficiency. If multiple hits from two different tracks or from accidentals occurred on the same wire, only the earliest of these hits was kept in order to simulate the double-hit response of the chamber TDCs.

The effects of δ-rays in the chamber were not accounted for in the Monte Carlo. This discrepancy resulted in substantial tails in the chamber position measurements observed in data that are not simulated in Monte Carlo. The effects of this disagreement on the analysis will be discussed in Section 8.2.1.

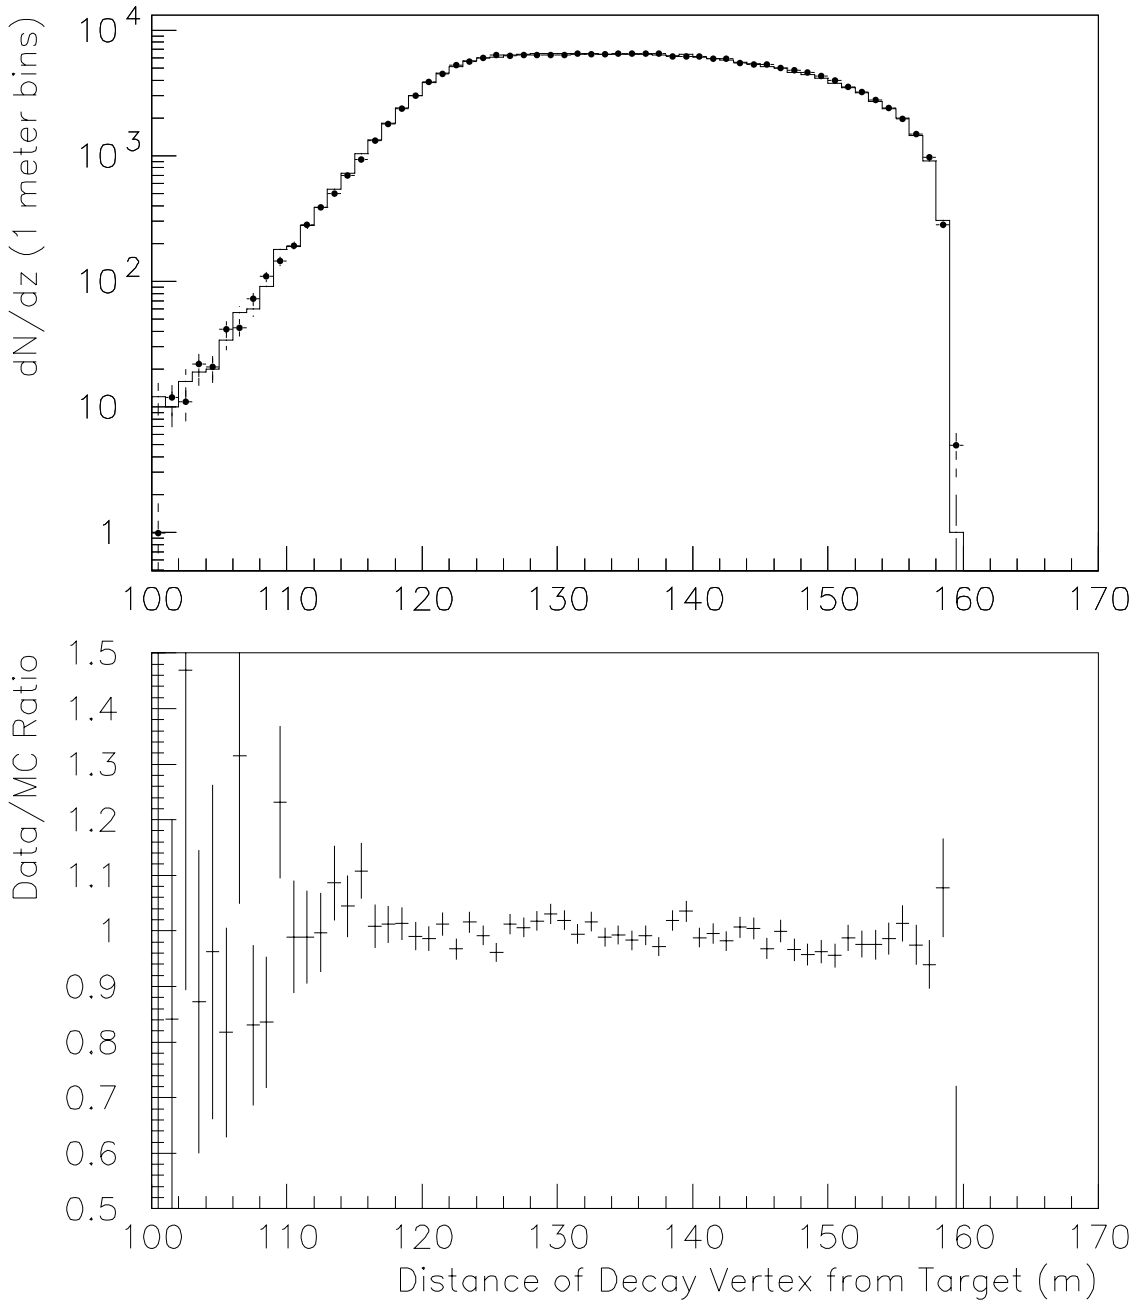

Figure 83. Distance from the target to the decay vertex in data and Monte Carlo for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$.

7.3.2 Simulation of the Lead-Glass Response

The simulation of the lead-glass response to photons and electrons was the most complicated detector simulation in the Monte Carlo. The Monte Carlo lead-glass response was based on parameterizations of showers from EGS Monte Carlo and on information obtained from the data. A parameterization of the EGS showers, rather than the actual EGS showers themselves were used in order to speed up Monte Carlo generation and to smooth out statistical fluctuations from the small shower statistics. The EGS electron shower Monte Carlo used is described in Section 6.1.1. Once showers from electrons were well-understood, shower from photons could be simulated by converting the photon at some depth into the lead-glass and starting two electron showers with fractional energies described by the Bethe-Heitler spectrum as was discussed in Section 6.1.2.

As was done in calculating the non-linearity correction, $C(E, \alpha, t_0)$ (see Section 6.1.1), the distribution of shower light responses,

$$\int_0^{L-t_0} f(E,t)e^{-\alpha(s-t-t_0)}dt$$

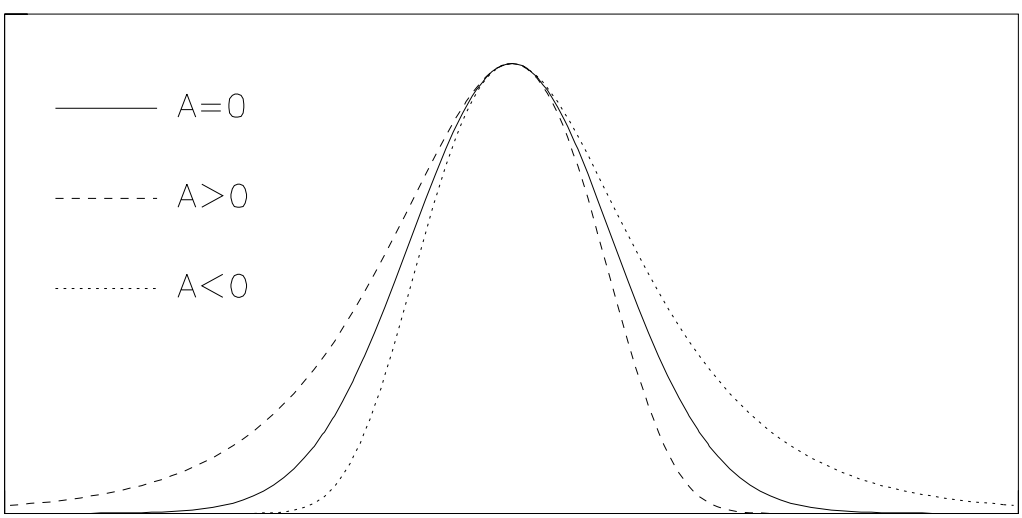

Figure 84. "Asymmetrical Gaussian" functional form

was plotted for 18 energies, logarithmically spaced between 0.25 GeV and $64\sqrt{2}$ GeV, 38 conversion depths between 0.0 X_0 and 18.5 X_0, and 76 α's between 0.02 and 0.17. Each of these 51984 response functions was then fit to the functional form,

$$Ne^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma-A(x-\mu)}\right)},$$

where N is a normalization factor, μ is at the peak of the distribution, and A represents an asymmetry in the distribution. In the limit of $A \rightarrow 0$, this distribution becomes a familiar Gaussian distribution, and for this reason the functional form is referred to as an "asymmetrical Gaussian". As shown in Figure 84, this functional form can accommodate a nearly Gaussian response with either a high-side or low-side tail, depending on the sign of the asymmetry parameter.

This functional form was convenient for a number of reasons. First, it was a good fit to the actual lead-glass response. There are two processes that skew the response away from its nominal Gaussian shape. If energy leaks out the back, this gives a low-side tail, and if the shower fluctuates deep into a high-absorption block, the response has a high-side tail. These two effects tend to cancel somewhat, and so typically only one of the two tails is prominent. The fact that the average χ^2 from these fits was reasonable, 1.4 per degree of freedom, bears out this observation. Second, this form used fewer parameters than if the high and low side tails were fit separately, and this made the fitting procedure more robust. Fewer than 0.2% of the fits failed the first pass of fitting and had to be redone. Sample fits to the glass response are shown in Figure 85, and an example of the variation of the fit parameters as a function of α and t_0 is shown in Figure 86.

Once all fits were complete, the results were smoothed as a function of energy since the Monte Carlo data at each energy were independent. Because one expects smooth behavior as a function of energy, this procedure helps to compensate for the poor EGS Monte Carlo statistics available. The smoothed energy response shapes were then used in the lead-glass Monte Carlo to select the energy of the cluster to be implanted in response to an incident particle. The table of responses as a function of energy, absorption and conversion depth was interpolated to obtain the response for each cluster generated in the Monte Carlo.

To check the Monte Carlo energy response simulation against the data, the sample of

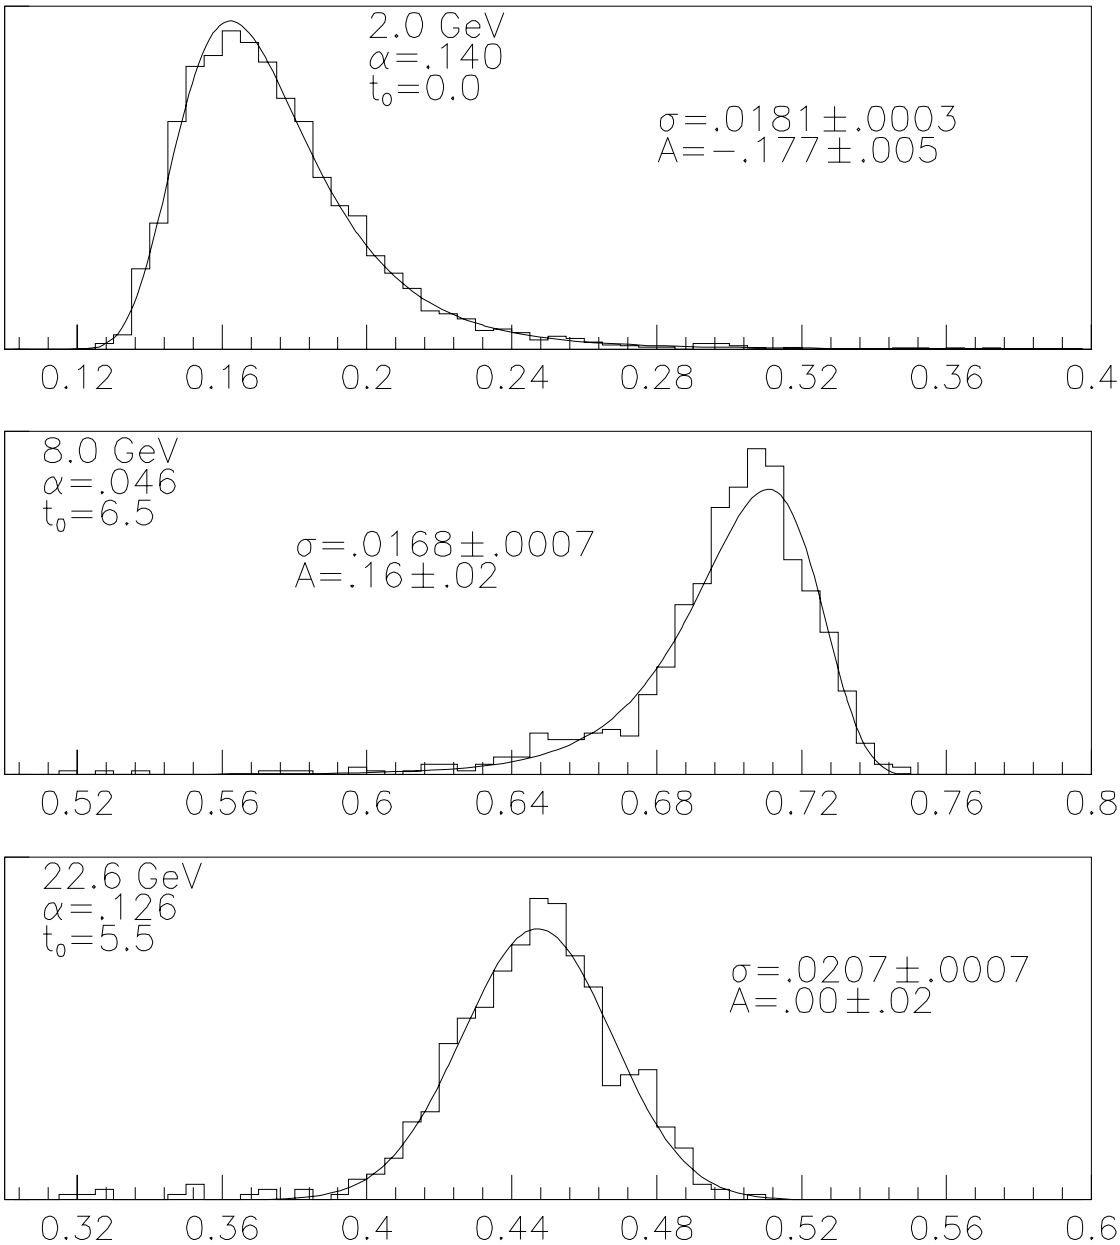

Figure 85. Sample fits to the lead-glass response for Monte Carlo parameterizations. In the third example, note the compensation between leakage and shower fluctuations which leaves a Gaussian response.

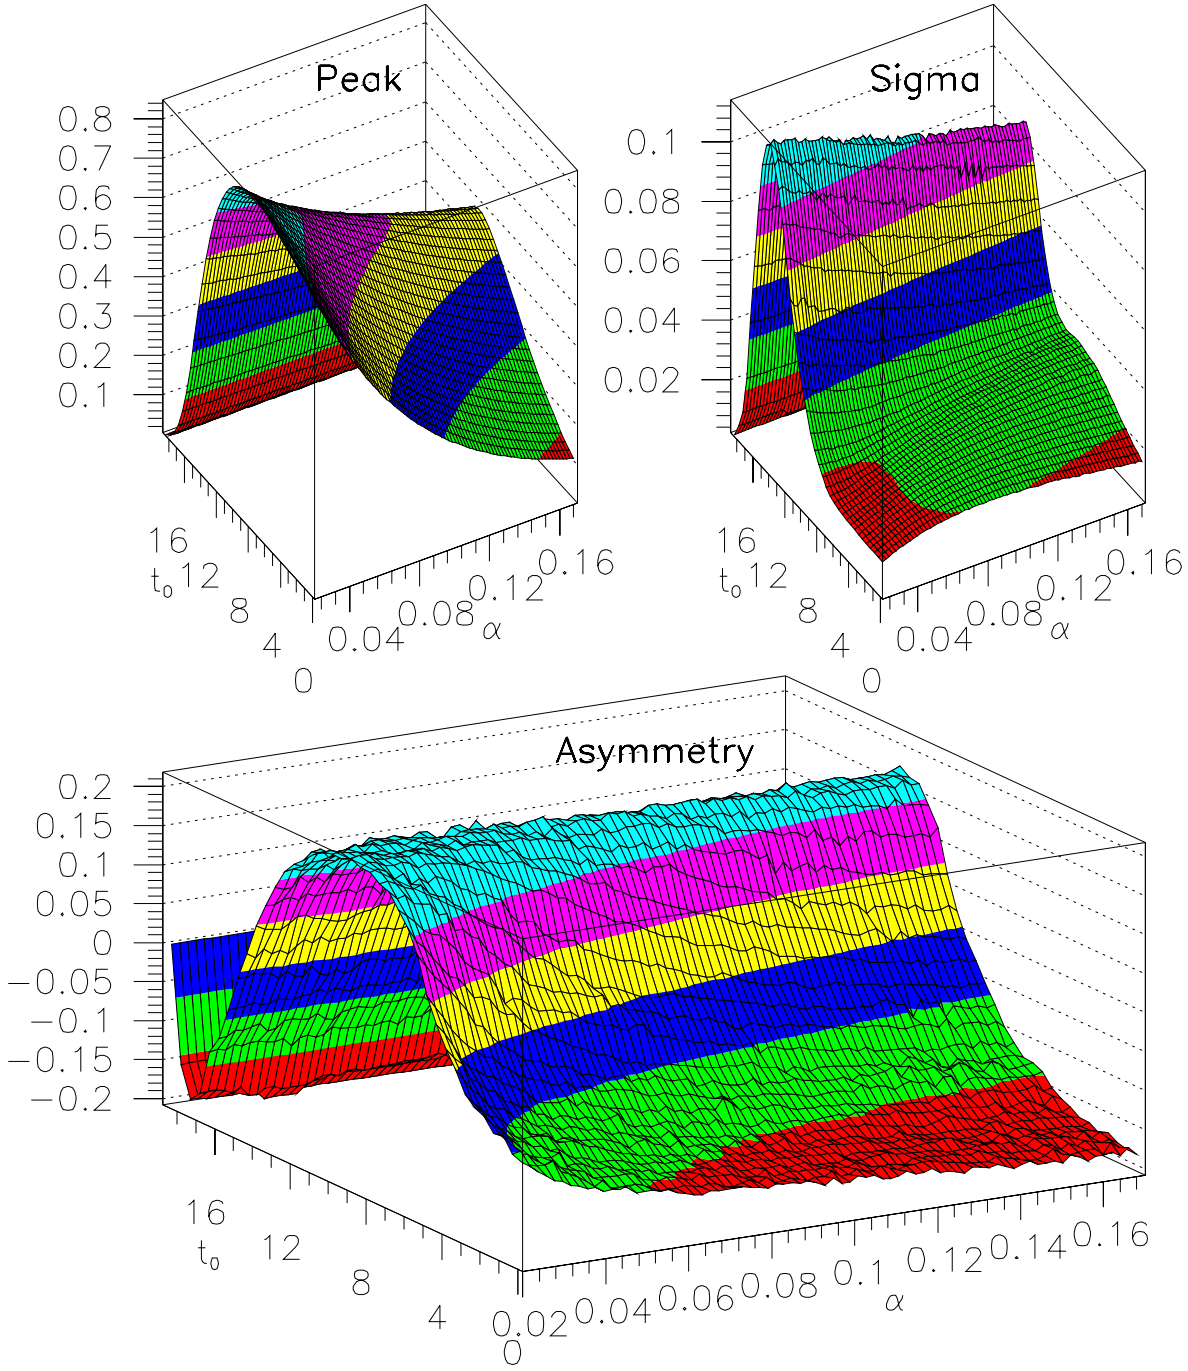

Figure 86. Fit parameters used in the Monte Carlo lead glass response for 2.0 GeV showers as a function of the absorption (α) and initial conversion depth (t_0). The shading indicates bands of constant value.

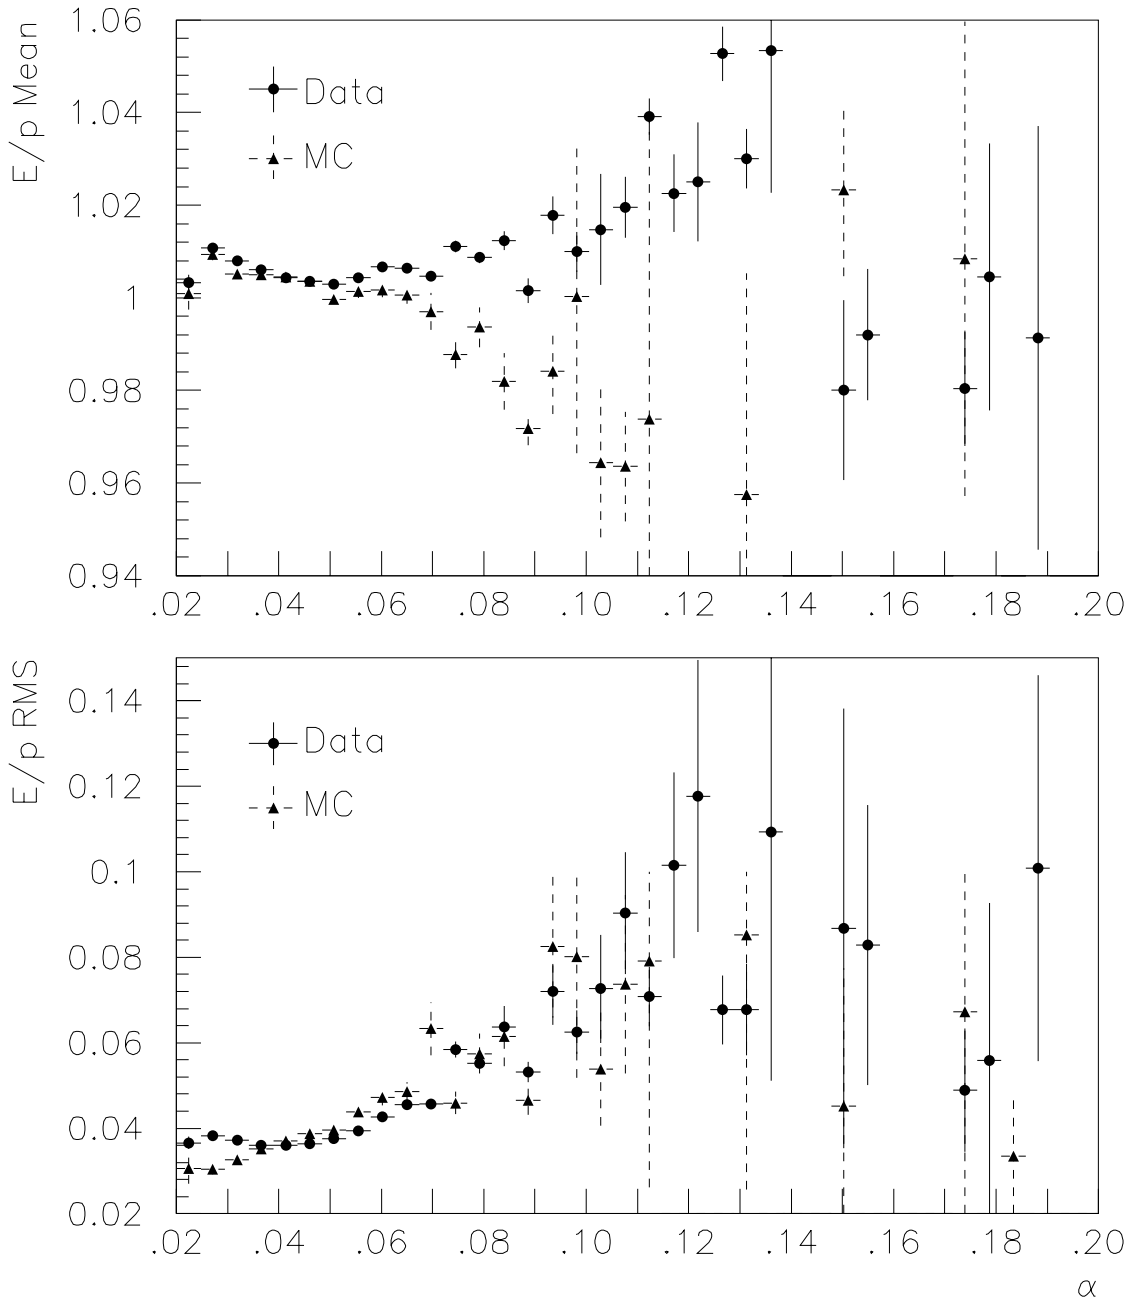

Figure 87. First pass data and Monte Carlo comparison of E/p Mean and RMS for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays

$K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, events was examined both in the data and in Monte Carlo. Electron E/p was plotted as a function of the α in the center block of the electron cluster. This comparison is shown in Figure 87. The variables compared in each bin of α are the mean E/p and the RMS E/p for all electrons in the sample with E/p between 0.85 and 1.15. Despite the low statistics in this first sample, there are some obvious features which can be noted.

First, in the plot of E/p means, note that E/p was approximately 1% high in both data and Monte Carlo. This effect was due to soft accidental energy being deposited on top of clusters in the data, and it was well-simulated in the Monte Carlo with accidental overlay techniques (see Section 7.4).

Second, as α increased the Monte Carlo mean E/p dropped significantly relative to data. This effect resulted from the mean shower response being defined differently in this parameterization and in the determination of $C(E, \alpha, 0)$. Because of the calibration, the mean E/p remained flat as a function of α by definition. Without an analogous calibration procedure in the Monte Carlo, this error caused E/p to decrease with increasing α.

Third, for very low α blocks, $\alpha < 0.045$, the RMS of E/p is significantly underpredicted in the Monte Carlo. This is most likely because the blocks which were identified with very low α actually had real α's that were higher. Recall that because these undamaged blocks were far from the center of the beamhole there were insufficient statistics for a measurement of α in these blocks. In these unmeasured blocks, plus the highest radius blocks for which the absorption was measured, there was a good chance of block absorption being higher than estimated.

Finally, for high α the Monte Carlo overpredicts the RMS of E/p, probably due to a failure of the lead-glass absorption model for very damaged blocks. To attempt to correct for all of these problems, three fudges were introduced into the Monte Carlo, based on fits to the differences between the data and Monte Carlo E/p mean and RMS distributions. The mean, μ, and width, σ, of the asymmetrical Gaussian were fudged as a function of α by

$$\sigma \to \sigma \times (1.37 - 8.3\alpha) \quad \text{if } \alpha < 0.045,$$

$$\sigma \rightarrow \sigma \times (0.86) \quad \text{if } \alpha > 0.09,$$

$$\mu \rightarrow \mu \times (0.989 + 0.25\alpha).$$

With these fudges implemented in the Monte Carlo, a new higher-statistics Monte Carlo sample of $K_L \rightarrow \pi^0\pi^0\pi^0$, $\pi^0 \rightarrow e^+e^-\gamma$, was generated to compare with the data. With somewhat different analysis cuts, these two samples were re-analyzed and compared, and the mean and RMS E/p versus p and α are shown in Figures 88 and 89. Also, shown in Figure 90 are the energy response shapes in data and Monte Carlo for different parts of the glass. The most obvious disagreement between data and Monte Carlo is seen in the mean E/p as a function of α. The low-statistics correction did improve this distribution substantially, but did not correct for the full effect. Some disagreement is also seem in the high-momentum widths, with the Monte Carlo underpredicting the RMS E/p slightly. Although significant disagreements between data and Monte Carlo remain, no further corrections to the lead-glass energy response parameterization were made, and the remaining differences were used to estimate systematic errors.

Once the cluster energy response was simulated, the next step of the lead-glass simulation was the implanting of this energy into different blocks of the lead-glass. For this, a sample of electron clusters from a "cluster library" was used to simulate the transverse shape of the lead-glass energy deposition. Below 4.0 GeV, the clusters from the library were generated by EGS Monte Carlo; above this energy, electron clusters from calibration data were used. Clusters in the library were indexed by position within the central block and by energy. For each electron shower generated in the glass, a random cluster from the library was chosen. Each library cluster consisted of energy information from a 5×5 block array centered around the incident electron, and the energies in this 5×5 were scaled so that the total energy deposition was given by the lead-glass energy response algorithm described above.

For charged pions and muons, the Monte Carlo of the lead-glass response treated each as minimum ionizing particles, and simply deposited 0.7 GeV in the lead-glass block hit by the pion or muon. With no attempt to simulate hadronic showers, this meant that the glass response of the Monte Carlo was not reliable in modes with charged pions. For this

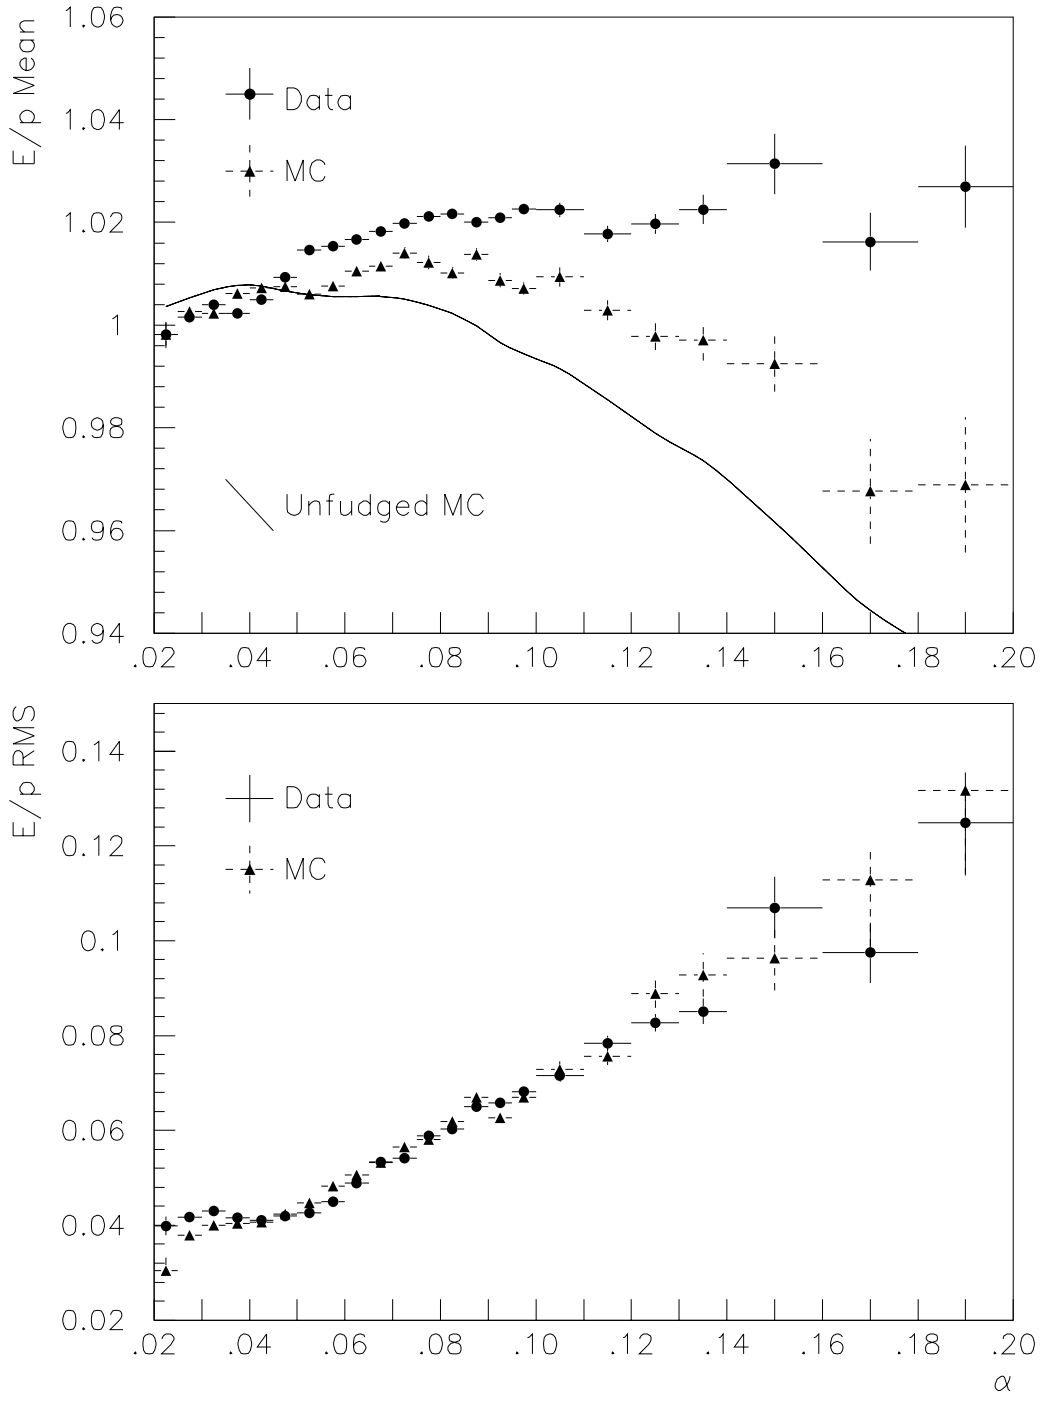

Figure 88. Second-pass data and Monte Carlo comparison of E/p Mean and RMS versus absorption (α)

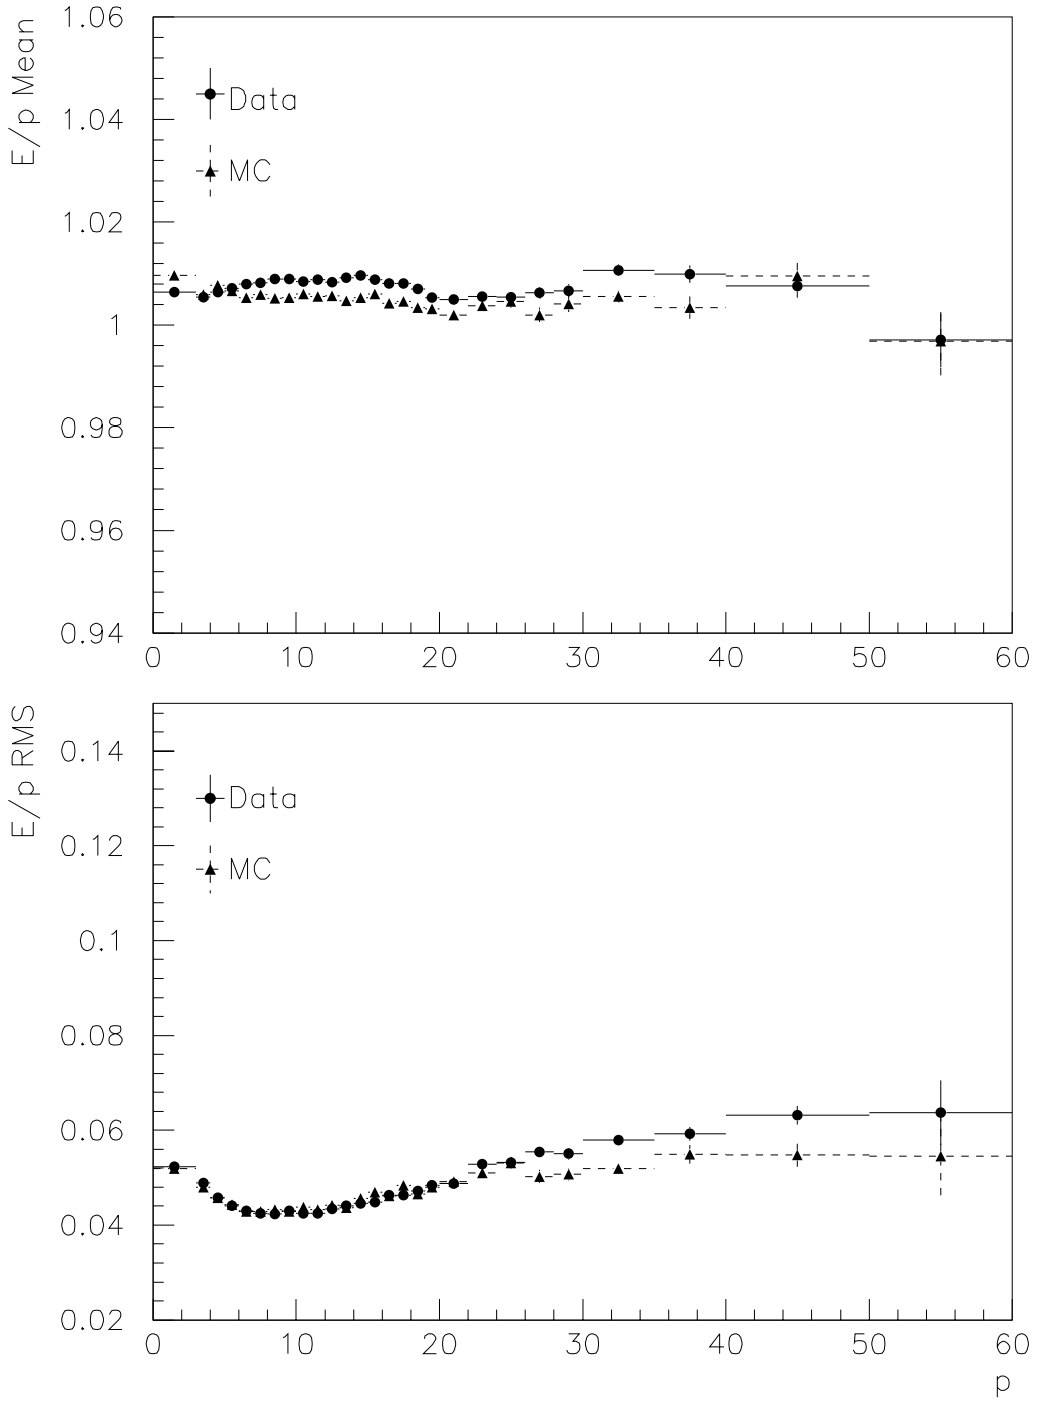

Figure 89. Second-pass data and Monte Carlo comparison of E/p Mean and RMS versus momentum

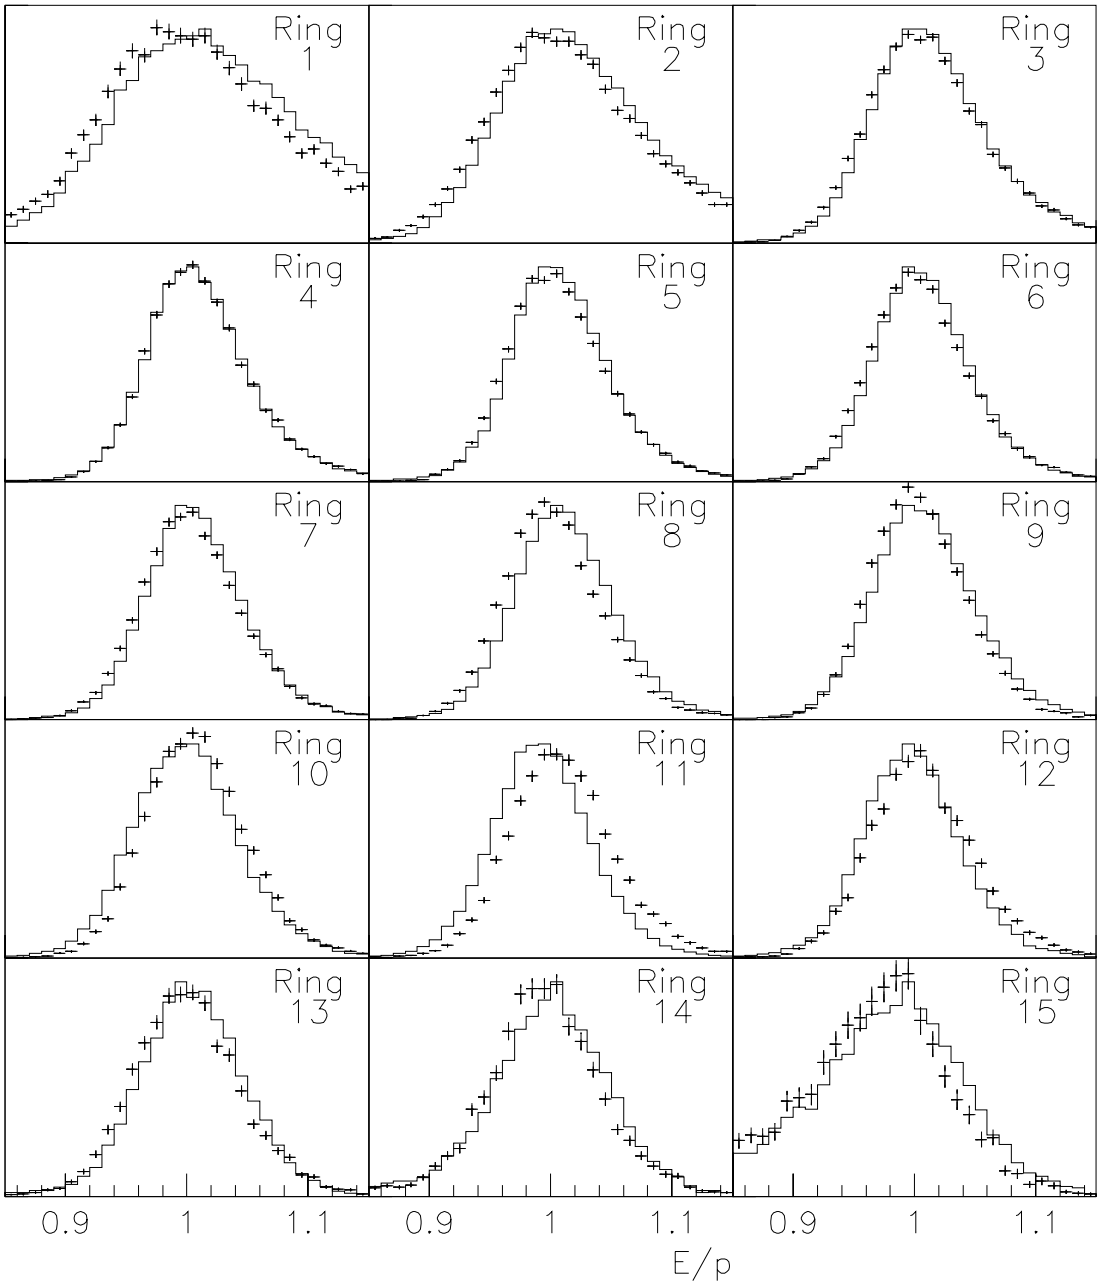

Figure 90. E/p shapes for different rings of blocks in the lead glass. Ring 1 includes all pipe blocks, and subsequent rings comprise concentric rectangles of lead glass blocks. Aside from shifts in the mean of the E/p distribution, the shape agreement between data and Monte Carlo is good.

and other reasons, no attempt was made in these analyses to simulate possible backgrounds with charged pions in the final state.

7.3.3 Simulation of Other Detector Elements

The response of the photon vetoes to charged particles and photons was also simulated in the Monte Carlo. If a charged particle was traced into a photon veto scintillator, the relevant trigger latch bit was set in the Monte Carlo. If a photon was traced into a photon veto, then the energy response in the photon veto ADC was parameterized by a Gaussian whose mean and width were measured from $K_L \to \pi^+\pi^-\pi^0$ data with one photon pointing into a photon veto. The Back Anti was handled differently. For each electron or photon pointing into the BA, its energy was first smeared by the measured resolution, $20\%/\sqrt{E}$ in GeV, and one-half of this energy was added into each view of the energy sum for the first two longitudinal layers of the BA.

The Monte Carlo also included a simulation of the trigger scintillators in the B and C banks, as well as the muon response to the MU1, MU2 and MU3 banks. For the B and C banks, charged particles were extrapolated to a particular counter, and the latch corresponding to this counter was set according to efficiencies and positions measured in the data. Signals from photon conversions in the B and C banks were also modeled. The simulation of the muon trigger banks proceeded in the same way, with the extrapolated or traced track position being used as an input to a position-dependent efficiency or response table for the individual counter hit. The details of this simulation are not important for the analyses described here, and can be found elsewhere [1].

7.4 Simulation of Accidental Activity

A key ingredient of the Monte Carlo was the simulation of accidental activity in the detector. As was discussed in Section 2.2.3.3, data was written to tape during the run using a special Accidental trigger. During Monte Carlo generation, the information from accidental events was used to model background detector activity in a Monte Carlo event.

For each Monte Carlo event generated, an accidental event from the data was read

in from data tapes. Once a Monte Carlo event was successfully generated, the generated event was then embedded into the accidental event. For the most part, the embedding was a straightforward procedure. ADC values and energies from the two events were simply added together. Trigger and latch bits were set if they were on in either event. In the case of chamber hits, the earliest hit from either the accidental event or the Monte Carlo event was used to generate the TDC value for each wire in order to correctly account for the chamber double-pulse response. ADC readout thresholds were always applied to the summed response from the accidental and Monte Carlo events.

One important question is how accurately this procedure reproduced the level of activity in the detector for a real event. The detector activity in accidental events resulted from three sources: noise in the detector, interactions or decays of neutral hadrons in the secondary beam, and proton interactions in the target (aside from the activity resulting from production of the secondary beam) or in the beam dump. The last of these three was what the accidental trigger observed, i.e.: muons from the beam dump or target. Noise was independent of beam parameters by definition; activity from the neutral beam and the primary beam were a function of the primary beam intensity and targeting parameters.

The stability of the targeting parameters presents a potential problem, since the the level of accidental activity from the secondary beam was dependent on the targeting in a way that was not monitored by the accidental trigger. To illustrate this with an extreme example, if the beam missed the target entirely, the accidental trigger would count at approximately the same rate, even though the accidental effects of neutral hadrons on the beam would be missing entirely. The targeting was monitored during the run using a small wire chamber (SWIC) in the primary beam just upstream of the target with a 0.5 mm wire spacing, and adjusted using steering magnets as described in Section 2.1.1. During targeting scans, shifts of 1.2 wires in the mean primary beam position caused 10% decreases in the ratio of secondary to primary beam intensity, and targeting shifts during data taking were usually corrected manually when they exceeded 1 wire. Therefore, one would expect that changes in targeting would cause changes in accidental activity from neutral hadrons of no more than 10%. Studies of trigger rates dominated by the secondary beam rate indicated that this variation was approximately 7% (RMS) during the run. Furthermore, the relationship

between primary and secondary beam intensity as measured in this way, was also linear to about 5% over the whole range of beam intensities observed in E799 as shown in Figure 91.

Because the accidental events were written to tape in the same way as any other trigger, effects such as deadtime which varied as a function of intensity and targeting were the same in data and in the Monte Carlo simulation. Another key assumption in the accidental trigger scheme was that the rate in the accidental trigger counter scaled linearly with the kaon rate in the secondary beam. This was checked using measures of the primary and secondary beam intensities, and, unfortunately deviations from linear behavior were found as shown in

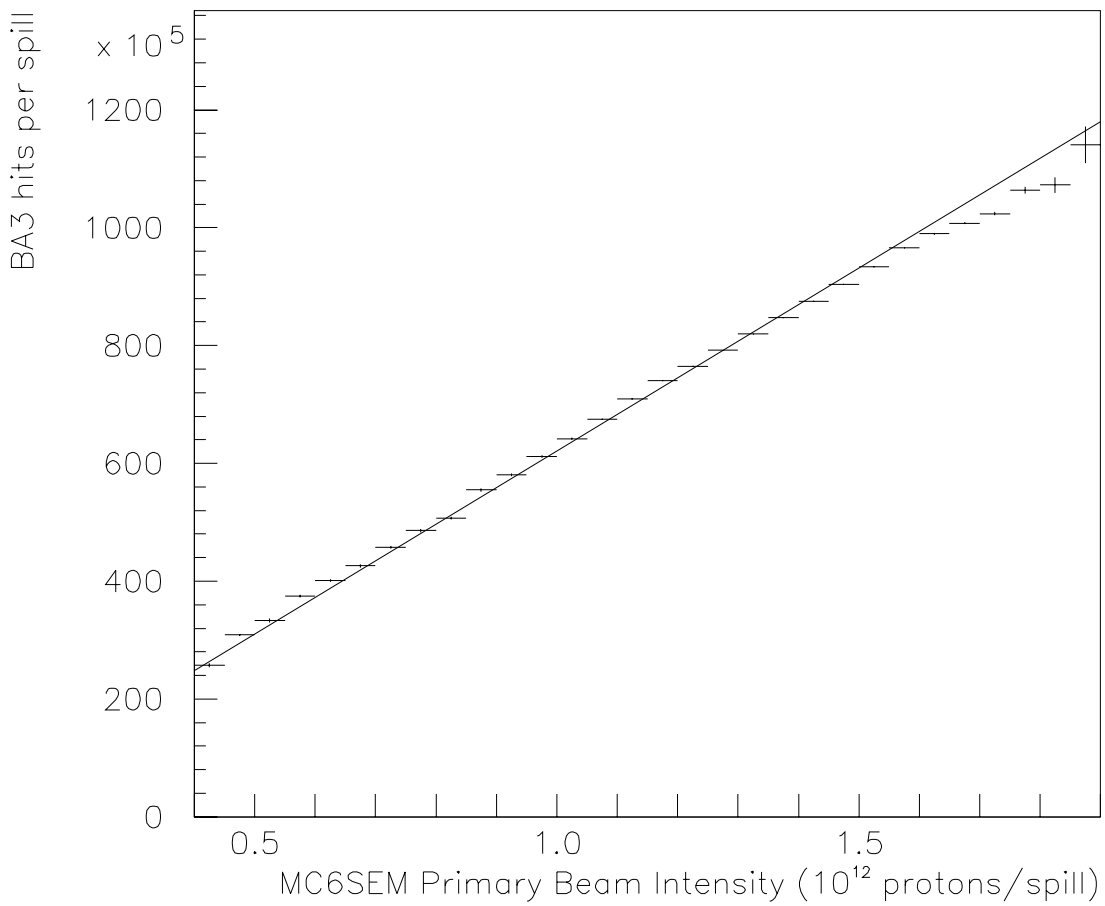

Figure 91. Secondary versus primary beam intensity as measured using the Back Anti counting rate, corrected for multiple occupancy. The relationship is linear to about 5%.

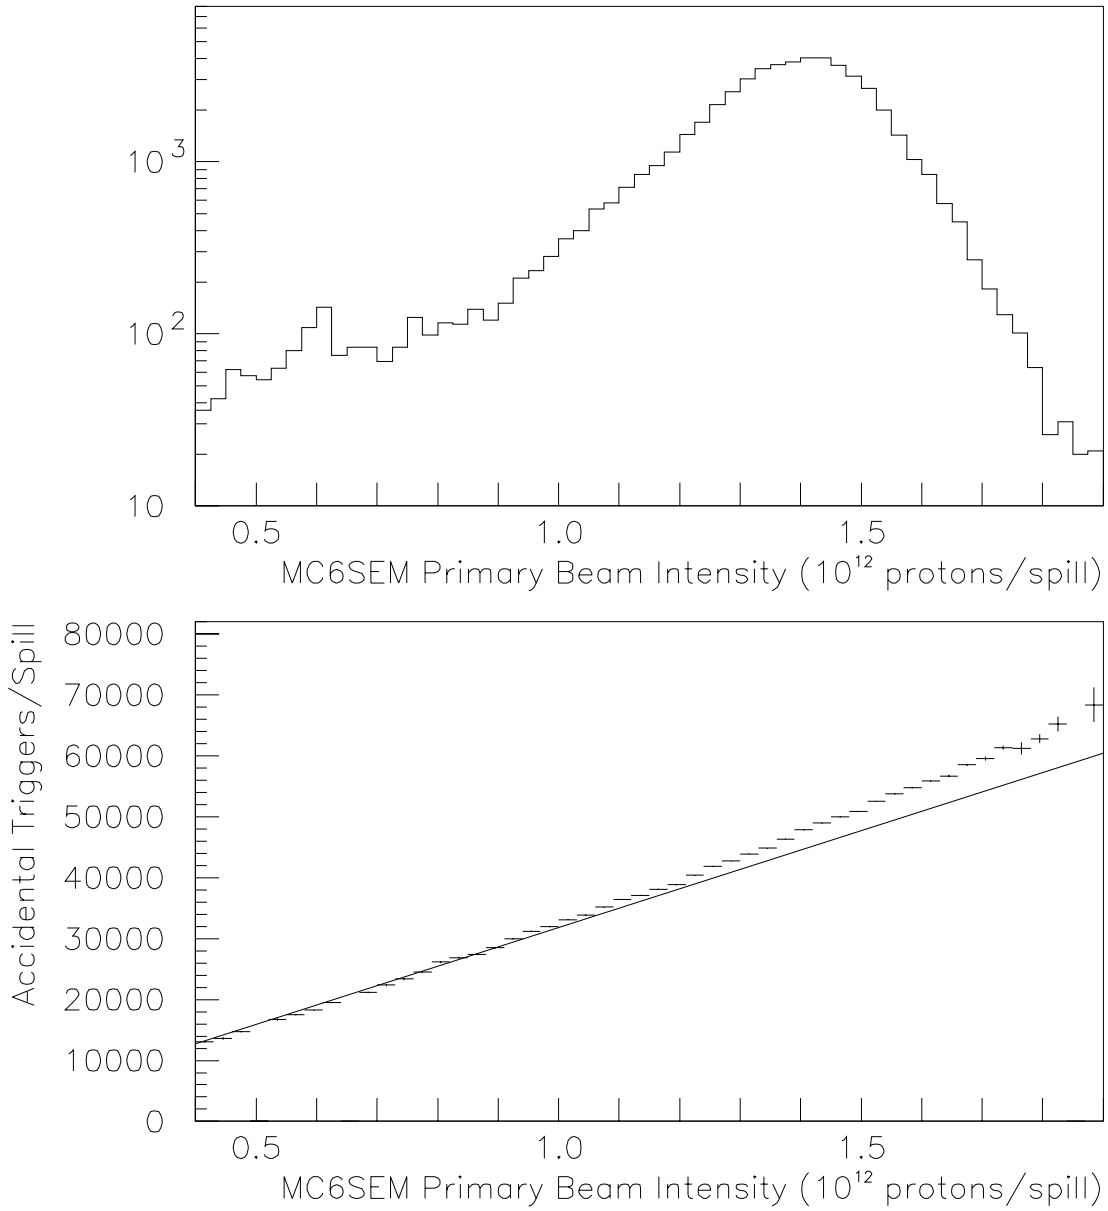

Figure 92. The distribution of primary beam intensities for the E799 run is shown, along with the mean accidental trigger rates in each bin of intensity. The accidental trigger rate deviates substantially from linear above 1.3×10^{12} protons per spill.

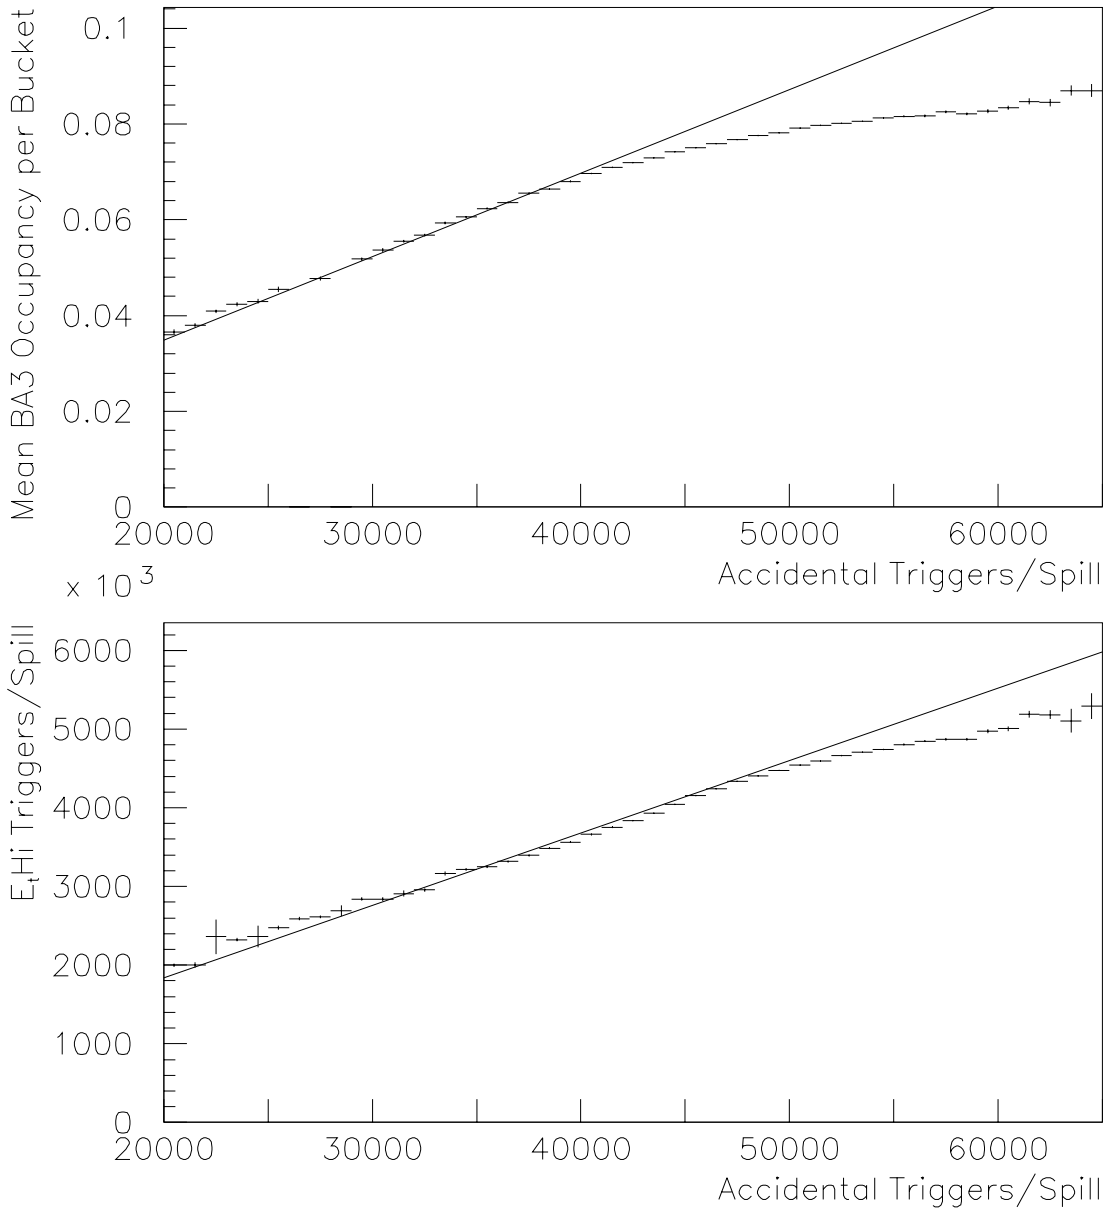

Figure 93. The accidental trigger rate is shown compared with rates in the Back Anti and the E_tHi Lead-Glass trigger. Both of these rates are expected to be linear with the secondary beam intensity, and the accidental trigger rate deviates from a linear response at high intensities.

Figures 92 and 93. These deviations resulted in a somewhat different spectrum of intensities in Monte Carlo and data, with a mean primary beam intensity 1.37×10^{12} protons per spill in data and 1.39×10^{12} protons per spill in the Monte Carlo. No attempt was made to correct for this effect.

7.4.1 Accidental Activity in the Detector

Because the effects of accidental activity were so severe on the trigger and detector acceptance, it is interesting to look at the activity that was present in E799. The sample of accidental events collected for the Monte Carlo simulation provides a perfect opportunity to do this. The effects of accidental activity can be put into several categories. First, accidentals caused direct losses in trigger acceptance by either causing trigger vetoes to fire or by depositing hardware clusters in the lead-glass. Second, accidentals caused trigger rate problems by faking elements of triggers designed to select kaon decays, thus increasing dead-time in the trigger and data acquisition system. Third, accidental activity could make the correct event reconstruction less likely through the presence of extra tracks or lead-glass clusters, or through extra hits in the drift chambers and soft energy deposits in the calorimeter. Finally, by examining accidental events, one can hope to gain insight on what physical processes caused these effects and thus improve the accidental environment in future experiments.

A number of the trigger vetoes caused substantial accidental losses. A list of many of the common vetoes and their rates in accidental triggers is shown in Table 21. In addition, events in the Two Electron, $e\mu$, Neutral and Λ^0 Beta decay triggers required an exact number of clusters from the HCF. Therefore, an extra cluster with sufficient energy to be found by the HCF will also cause an event to not satisfy the trigger. In accidental events, 13.1% of the events had at least one extra hardware cluster. These losses represent a significant loss in sensitivity in many of the triggers; in the Two-Electron trigger in which the data for this analysis was taken, the total loss at trigger level from accidental effects was about 40%.

Accidentals also substantially increased trigger rates in the detector. The rates in accidental events for a number of Level 1 trigger source combinations used in the E799 physics

Table 21. Level 1 trigger veto acceptance losses as measured in accidental events

Trigger Veto	Loss in Accidentals	Triggers Affected
VA-1	1.7%	All Physics Triggers
VA0	1.8%	All Physics Triggers
VA1	5.4%	All Physics Triggers
VA2	5.3%	All Physics Triggers
DRAC	4.2%	All Physics Triggers
VA3	6.8%	All Physics Triggers
VA4	5.4%	All Physics Triggers
LGA	1.7%	All Physics Triggers
All Photon Vetoes	13.7%	All Physics Triggers
BA_γ	6.2%	Two Electron, $e\mu$ and Neutral
MU1	18.6%	Two Electron, Dimuon, $e\mu$ and Λ^0 Beta Triggers
MU2	15.0%	Neutral, $\Lambda^0 \rightarrow pe^-\nu$
MU3	0.6%	Two Electron
3B·3C	3.2%	$\Lambda^0 \rightarrow p\pi^-$

Table 22. Level 1 trigger rates in Accidental events

Trigger	Fraction of Accidental Trigger Satisfying
$1B \cdot 1C$	0.157
$UPBM + DNBM$	0.084
$2B \cdot 2C$	0.070
$3B \cdot 3C$	0.032
$2B \cdot 2C \cdot E_tHi$	0.014
$2B \cdot 2C \cdot E_tHi \cdot \overline{MU1}$	1.7×10^{-3}
$2B \cdot 2B \cdot MU3$	3.4×10^{-3}
$2B \cdot 2B \cdot MU3 \cdot \overline{MU1}$	6.5×10^{-4}
E_tHi	0.020
$E_tHi \cdot \overline{PHV} \cdot \overline{BA_\gamma}$	8.4×10^{-3}
$2B \cdot 2C \cdot (UPBM + DNBM)$	0.030
$2B \cdot 2C \cdot (UPBM + DNBM) \cdot E_tLo \cdot \overline{MU1}$	2.7×10^{-3}

triggers are shown in Table 22. Accidental activity also had a major effect on the hit-counting trigger of the track processor. The mean number of hits or hit pairs found in each drift chamber as counted by the track processor hit-counting in accidental events is shown in Table 23. Only about 27% of each of these chamber hit pairs are consistent with being from tracks in the same bucket as the accidental trigger (Figure 94), but they contributed to the hit-counting trigger rate nevertheless. Finally, as was mentioned above, accidental effects often increased the multiplicity observed in the lead-glass by the hardware cluster finder. The mean number of hardware clusters found in accidental events is 0.22, with 4.8% of accidental events having two or more clusters found by the HCF.

Accidental effects also caused difficulties in event reconstruction. In the drift chambers, the probability of finding extra tracks is fairly large, despite a substantial reduction from the sum-of-distance cuts in the E799 tracking algorithm (see Figure 95). Energy deposited in the

Table 23. The number of chamber hits in Accidental events

Chamber Plane Pair	Mean Number of Hits or Hit Pairs	Mean Number of Hits or Hit Pairs (noisy wires removed)
1Y	0.95	0.85
1X	0.84	0.84
2Y	1.32	1.32
2X	1.36	1.36
3Y	1.41	1.41
3X	1.47	1.42
4Y	1.86	1.67
4X	1.71	1.71

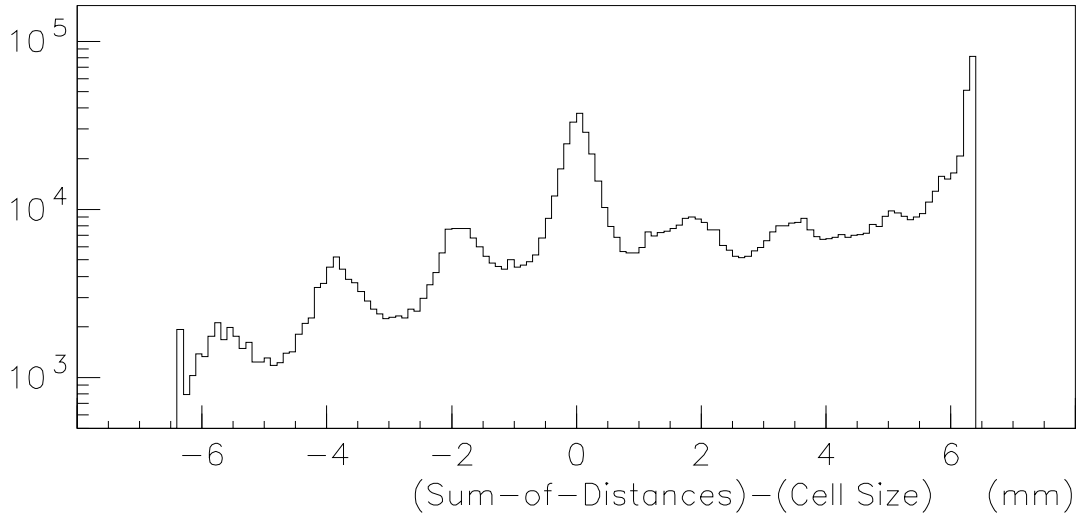

Figure 94. The sum-of-distances for chamber hit pairs in Accidental Triggers

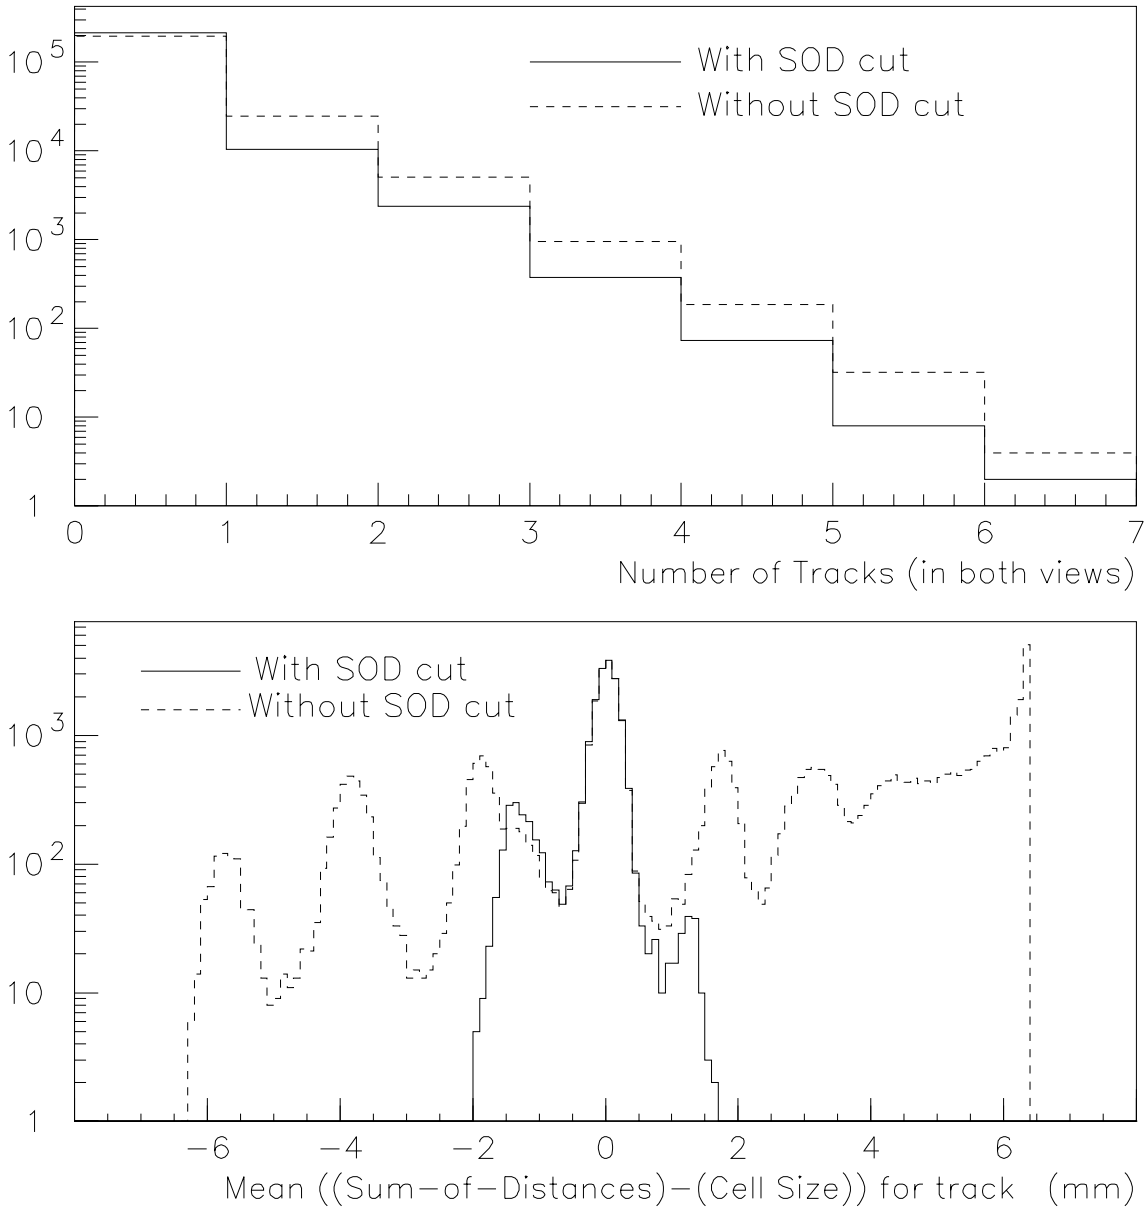

Figure 95. The number of tracks found in both views for Accidental Triggers. Two versions of the tracking described in Section 4.2 are compared: the standard tracking algorithm, and a special version which does not cut on the sum-of-distances for hit pairs.

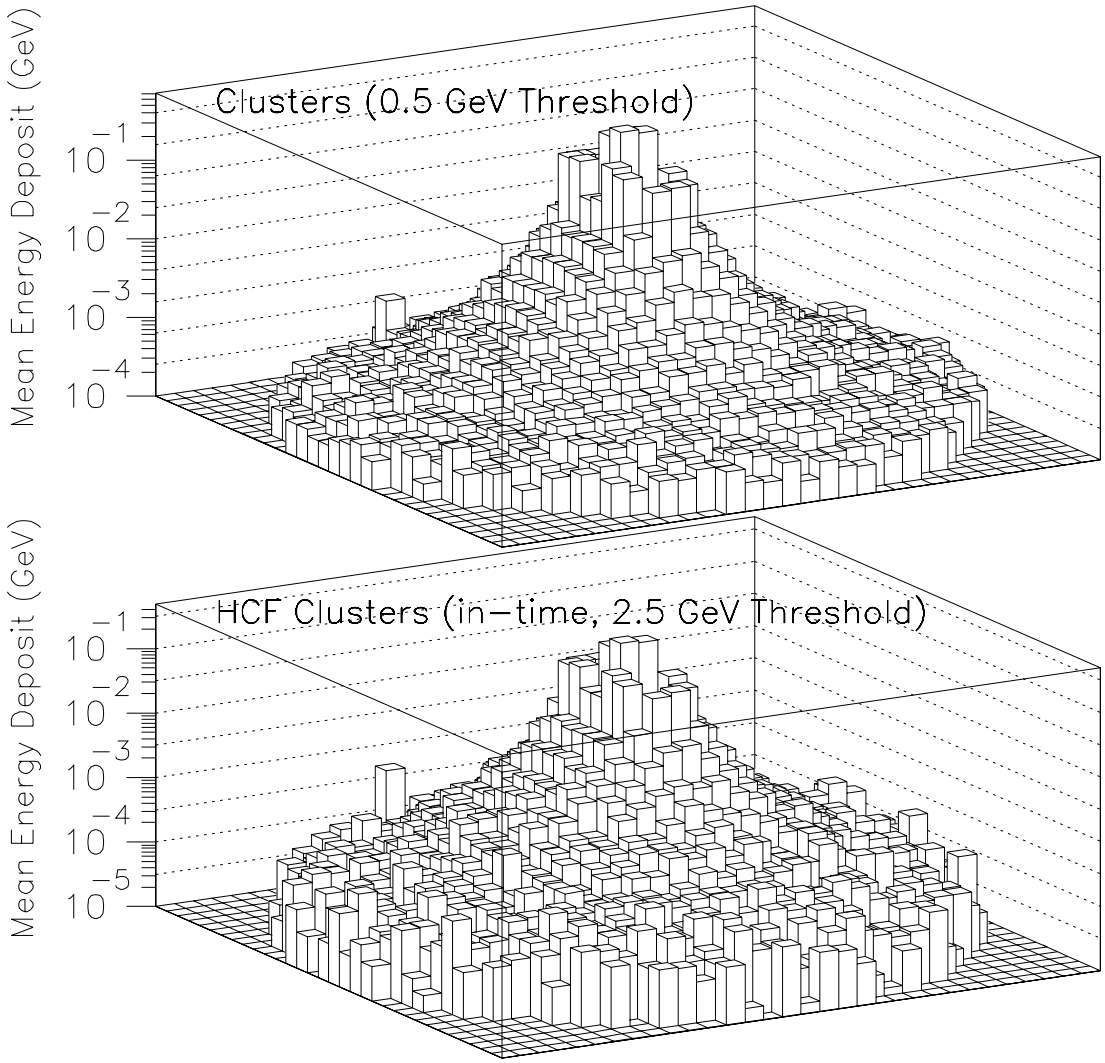

Figure 96. The mean energy deposition in each lead-glass block for Accidental Triggers. The plot depicts a map of the glass with the $\hat{x}$ and $\hat{y}$ axes corresponding to block location and the height in each tower representing the mean energy deposit. The mean energy in each 100 ns glass gate is shown, along with only the energy in clusters found by the Hardware Cluster Finder. The energy deposited is very sharply peaked towards the center of the glass.

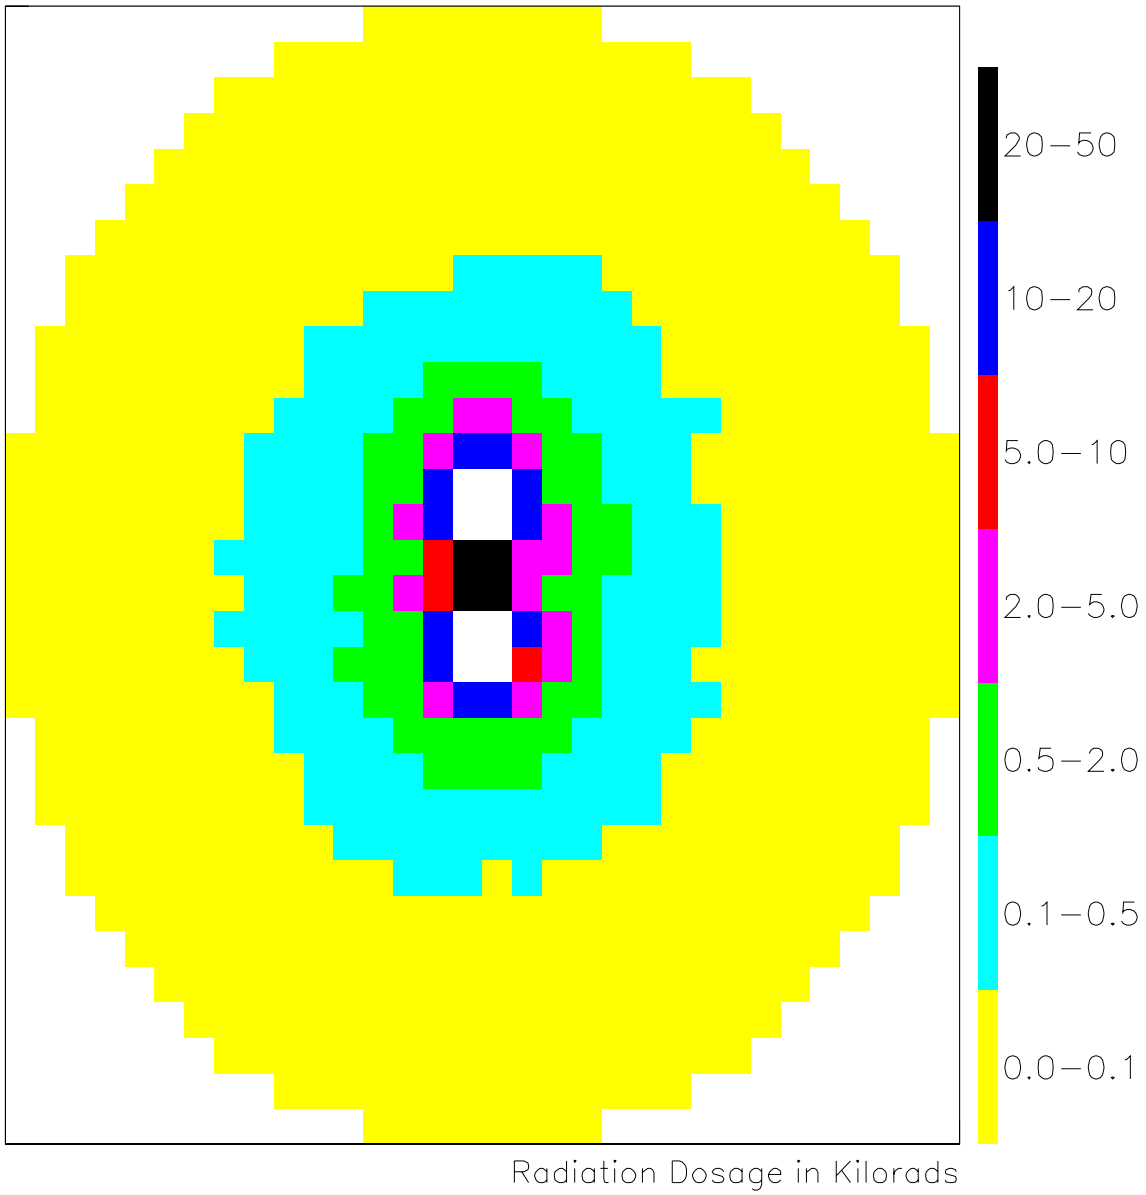

Figure 97. The radiation dosage per block as measured in Accidental Triggers

glass also led to possible confusion in cluster reconstruction. The mean energy deposition, by block, is shown in Figure 96. The energy deposition is sharply peaked towards the beam holes of the glass, which should be no surprise since radiation damage was peaked there also (see Figure 76). In fact, one can use the deposition of energy in accidental events to obtain estimates for radiation dosage in each block, as shown in Figure 97, and these agree roughly with the dosages from radiation monitors given in Figure 18.

Finally, one can use the information contained in Accidental Triggers to attempt to determine the source of the accidental activity observed in the detector. The most damaging effect of accidental activity was clearly the energy deposited into the lead-glass, since this affected the trigger and the event reconstruction, and also caused the radiation damage in the lead-glass. Recent studies both from Monte Carlo simulations and from accidental data [74] have led to the conclusion that a large fraction of the observed energy deposited in the glass is the result of the interaction of hadrons with material in the beam, particularly the B and C-bank scintillators. In particular, for events with no wire hits in the most downstream drift chamber, 39% of energy deposited into the lead-glass adders (which have a narrow 30 nsec gate) were associated with significant activity in the B-bank, indicating an interaction occurring between the last drift chamber and this trigger bank. Even more significantly, 60% of events with hardware clusters in the lead-glass and no hits in the last drift chamber show activity in the B-bank. Thus, the most energetic depositions in the lead-glass appear to be associated with interactions in the trigger scintillator banks which may be responsible for much of the observed radiation damage in E799.

One could also use accidental events to determine the sources of hits and tracks in the drift chambers. Approximately 50% of the single-track events in Accidental Triggers appeared to be muons from the target, and the other half was presumably decay products of neutral hadrons from the beam. The two-track rate, however, appeared to be dominantly from decay products, and contained only a very small component of two muon events, as shown by the low dimuon trigger rate in Table 22. Since single track events were approximately 5 times more prevalent than multi-track events in Accidental Triggers, the target muons still provided a large fraction of the single-hit rate in the chamber however, accounting for about 1/3 of these accidental hits.

7.5 Trigger Simulation

In order to correctly simulate detector response, it was also necessary to reproduce the effect of the trigger on the data. Many of the trigger elements were straightforward to simulate. Trigger sources which used input from scintillator latches could simply use the same information in the Monte Carlo events. Similarly, the track processor triggers were simulated using the drift chamber TDC information generated in the Monte Carlo. However, in cases where the trigger used an analog threshold, the simulation was in general much more complicated because the information used in forming the trigger was not saved or simulated in the data stream. From the list of Level 1 trigger sources in Table 9, one can see that the only trigger sources based on analog inputs were MU1, the LGA, the BA and the Lead-Glass triggers. The response of MU1 to charged pions was not simulated, and the Monte Carlo response of MU1 to muons, described elsewhere [1], was not important for this analysis. In the case of the BA and LGA, these thresholds could safely be ignored by simply cutting more tightly than the trigger did in the offline analysis of both data and Monte Carlo. However, in the case of the lead-glass, the loss in acceptance from cutting above the E_tHi and HCF trigger thresholds would have been too great. Acceptance for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, was reduced approximately 1%/1 GeV increase in the E_tHi threshold and 1.6%/0.1 GeV increase in the HCF cluster threshold. Therefore, it was necessary to simulate the response of these two triggers in Monte Carlo.

In the case of the HCF threshold, the response of the HCF flash ADCs as a function of known lead-glass ADC values had been studied using data from previous experiments using the same ADCs and HCF [60]. The response was parameterized using a Gaussian lineshape, with the mean and RMS determined for each block of the lead-glass array. In E799, the HCF was run at a higher threshold than in these previous experiments and was adjusted in the Monte Carlo accordingly. Also, for the last 14% of the data, the overall effective HCF response dropped by 27% for unknown reasons. This effect was simulated in the Monte Carlo by changing the threshold at this point in the run. This simulation was checked by comparing distribution of the minimum cluster energy in data and Monte Carlo in the highest-multiplicity, highest-statistics decay channel available, $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$

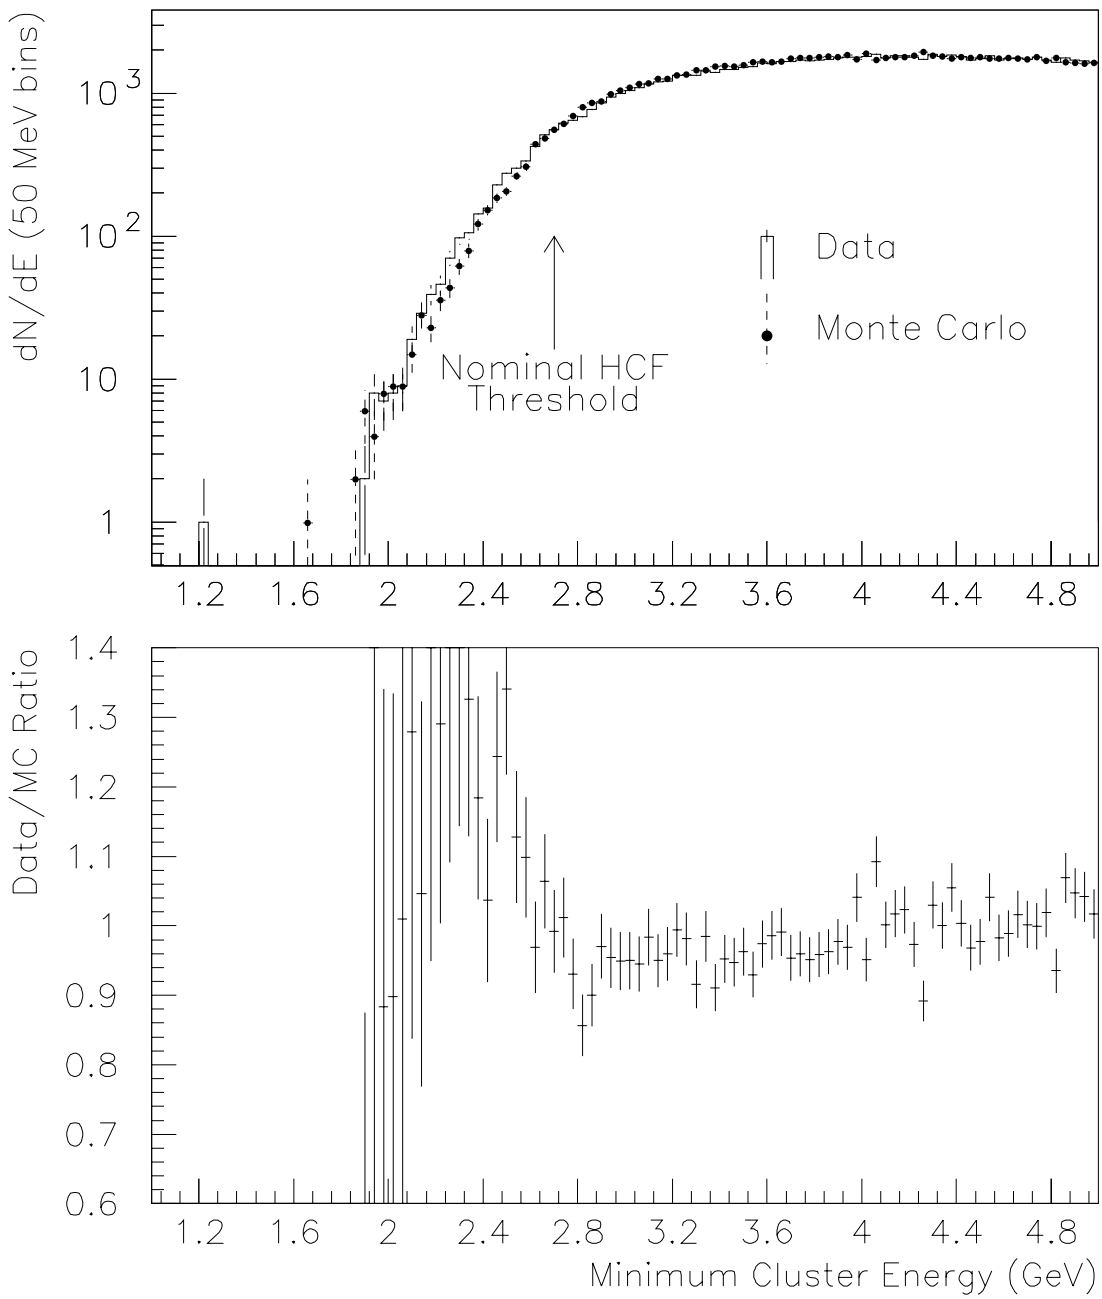

Figure 98. Minimum Cluster Energy in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$.

(Figure 98). Although there is some disagreement on the low tail of the distribution, the agreement is basically good. The threshold offset between data and Monte Carlo appears to be 40 MeV at most, which would correspond to a 0.6% fractional change in acceptance for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$. The Monte Carlo also included the simulation of several dead blocks in the HCF which were found during the initial data reduction.

The simulation of the E_tHi threshold also took advantage of the fact that all of the signals which were combined to form the analog E_t were read out by an ADC. The E_t ADC response was parameterized in the Monte Carlo by

$$E_t\text{ADC} = C + \sum_{n=1}^{92} g_n(\text{AdderADC})_n,$$

where the g_n represents E_t gain constants. These gain constants were determined using events where more than half of the lead-glass adder energy was deposited in a single adder. For such events, the ratio of the E_t ADC as calculated using the relation above was compared to the real E_t ADC value. The mean value of this ratio was then used as a correction to the gain for the most energetic adder in that event. Because the gains of the adders were coupled in this procedure, it was necessary to iterate this several times to get stable gains. The measured gains were significantly different for different adders, varying over a range of 30% from adder to adder. The simulation ultimately was able to reproduce the E_t ADC value on an event-by-event basis with 6% accuracy. In addition to the E_t ADC values, lead-glass quadrant ADC values were also calculated using these same gains. These simulated ADC values were then required to be above thresholds measured from the data in order for the E_tHi, E_tLow or lead-glass quadrant trigger bits to be set in Monte Carlo. The result of this simulation is shown below by comparing total energy from $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays in data and Monte Carlo (Figure 99).

7.6 Conclusion

With a functioning Monte Carlo as a guide, the analysis to search for rare decays can now proceed! Some more subtle Monte Carlo issues beyond the basic ones discussed here arose

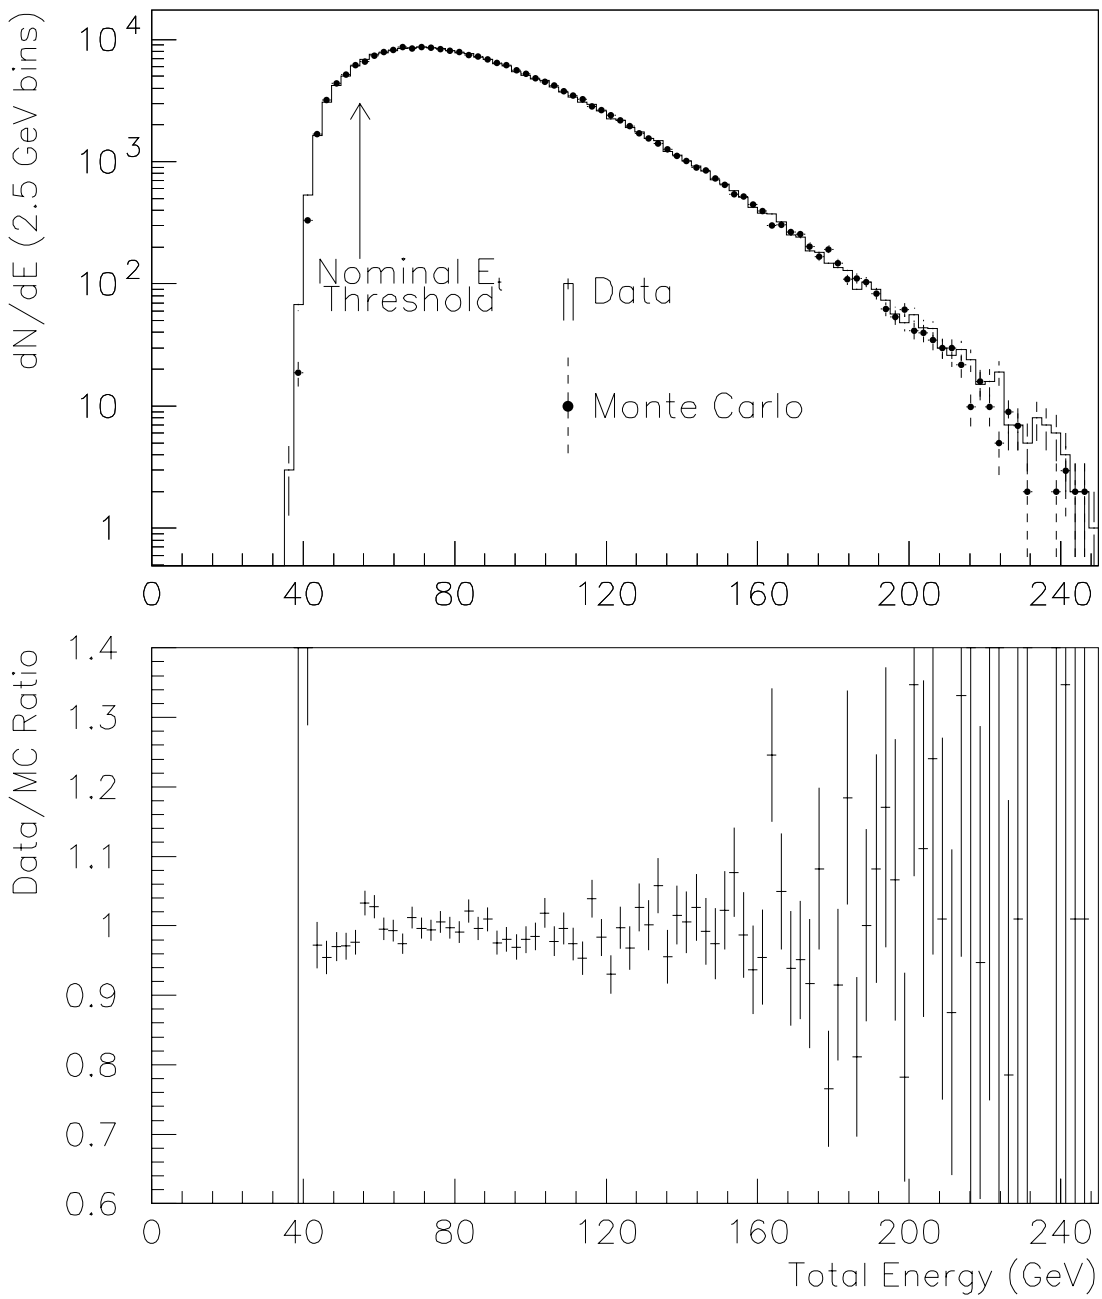

Figure 99. Total Energy in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$.

during the process of background simulation and analysis, and these will be discussed as they arise in the next chapter.

CHAPTER 8

ANALYSIS OF THE $\pi^0 \to e^+e^-$ SAMPLE

This chapter describes the search for $\pi^0 \to e^+e^-$ decays in the data. Because of the large amount of data, over 300 GBytes, it was first necessary to perform an initial reduction of the data set. After this was complete, the search through the remaining events for $\pi^0 \to e^+e^-$ candidates could proceed. The extraction of the $\pi^0 \to e^+e^-$ branching ratio and the evaluation of statistical and systematic errors are also described in this chapter.

8.1 The Initial Data Reduction

After the Two Electron triggers were split from the rest of the data (see Section 3.5), the initial data reduction (often referred to as the "crunch") of each data set could proceed independently of the others. The data described in this thesis was taken entirely in the Two-Electron trigger, so only the crunch of data in the Two-Electron stream will be described here. Recall that this stream included events from both the Two-Electron and the Dalitz triggers.

The goal of the crunch was to reduce the data set size by approximately a factor of ten while making only minimal requirements on the events. The strategy adapted was to simply search for events with at least two electrons hitting the lead-glass calorimeter. In order to do this, clusters of energy were located in the lead-glass; tracks were found in the drift chambers, and each track was then required to match to a lead-glass cluster. The candidate electron cluster energy and track momentum were then compared in order to identify the track as an electron.

The lead-glass cluster finding algorithm used was the Hardware Cluster Finder driven algorithm described in Section 4.1. This guaranteed that each photon and electron candidate found in the crunch had also been found by the HCF. No cut was made on the result of the cluster-finding algorithm in the crunch; however, if the clustering algorithm failed because it found too many clusters or because the HCF information was not present or was corrupted for that event, the event was discarded.

The tracking algorithm is described in Section 4.2. The choice among the available options in the tracking was an important decision to be made before beginning the data analysis. Because the reconstruction of events from the Two Electron trigger would require

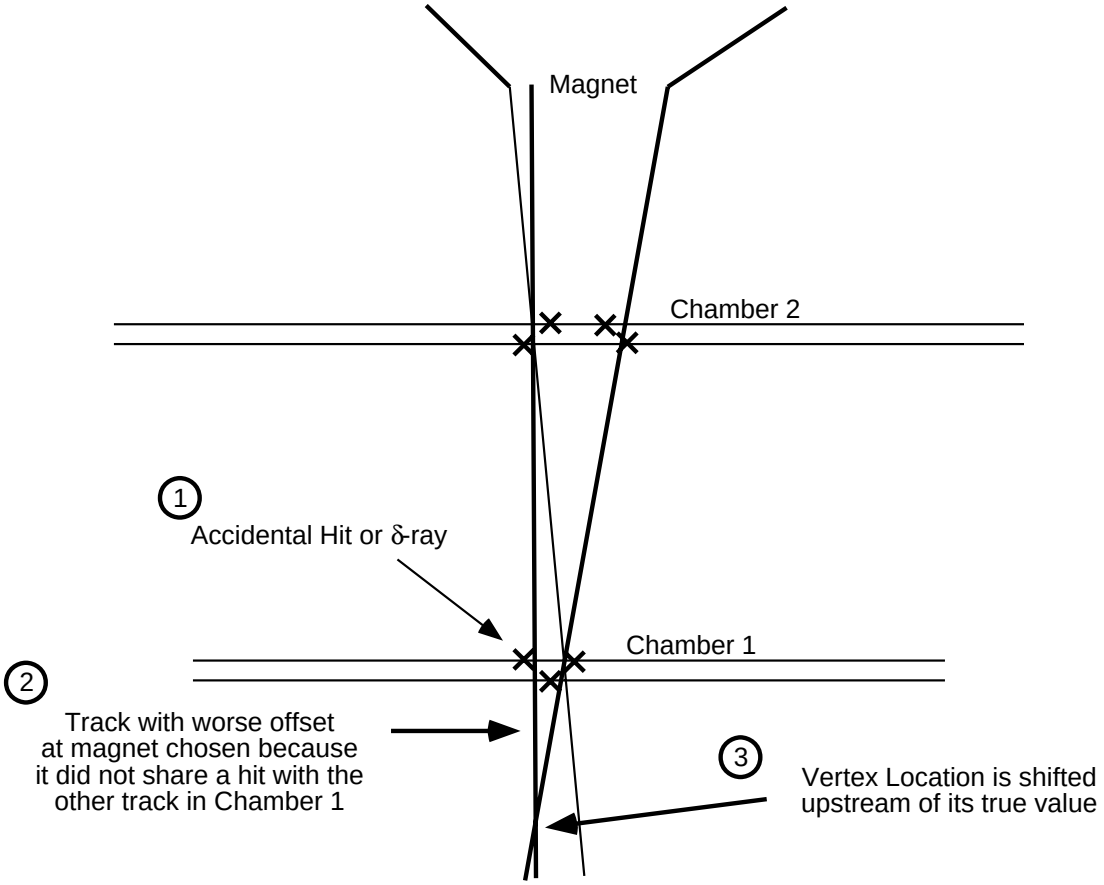

Figure 100. The effect of requiring no chamber sharing in the tracking on conversion pairs from the vacuum window is shown. If the track segments are moved apart by this requirement, the vertex location will shift upstream of its true position.

the electron tracks in the drift chambers to be matched with the electron clusters in the lead-glass, track segment candidates were required to point to locations consistent with cluster positions. A more subtle issue was that of hit-sharing in the drift chambers. Because the track processor hit-counting required distinct hits in each chamber from every track, tracks sharing hits must be eventually discarded from the sample in order to ensure that candidate events passed the trigger. However, not allowing hit-sharing in the track finding could have undesirable effects. A common problem that arose was the following: consider a pair of electrons which were so close together that they shared track hits in the first drift chamber. Because of the large amount of material in front of the first chamber (0.31% of a radiation length, including the chamber window), the probability of one of the electrons producing a secondary electron, i.e. a δ-ray or a conversion product of a bremsstrahlung photon, was large. If this secondary electron hit a different cell in the first drift chamber, and if sharing of hits in Chamber 1 was not allowed, the track algorithm often used the secondary electron's

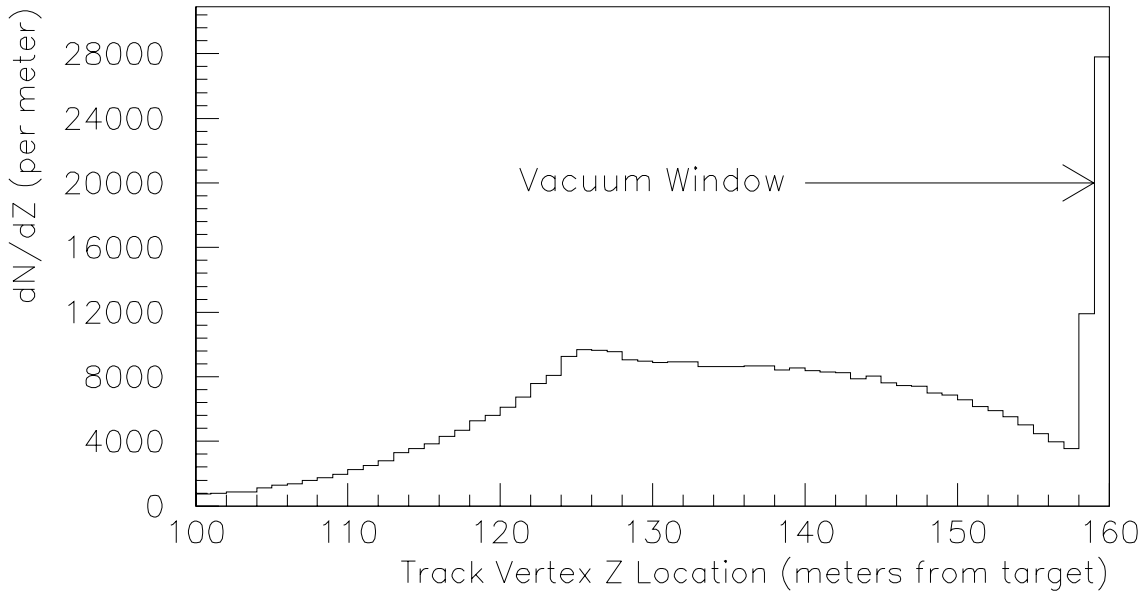

Figure 101. The raw track vertex distribution is shown from a subset of the events gathered in the Two Electron stream. The large downstream peak results from photon conversions to e^+e^- pairs in the vacuum window.

hits to complete one of the tracks. This could have undesirable side-effects, such as shifting of the reconstructed vertex location as shown in Figure 100. The potential magnitude of this problem becomes evident when one looks at a plot of the vertex position for events from the Two Electron trigger (Figure 101). The large observed peak of events with vertices very close to the first drift chamber resulted from conversion of photons in all the material just upstream of the chamber. If such a conversion occurred, it was unlikely that the resulting electrons would be far enough apart at the first drift chamber to produce two separate hits. (This required an opening angle greater than 5 mrad in the electron conversion pair in order to get the required 1.5 mm separation at the chamber, and the RMS opening angle from a thin radiator of $3.1 \times 10{-}3\ X_0$ was less than 0.4 mrad for electrons above 2 GeV.) Therefore, in order to minimize vertex misreconstruction in vacuum-window conversion events, sharing of hits was allowed in the first drift chamber for the Two-Electron crunch track-finding.

With hit sharing allowed in Chamber 1 and with the track-cluster match requirement, the outcome of the tracking in the crunch is shown in Figure 102. In order to be minimally consistent with two electrons in the final state, at least two $\hat{x}$-view tracks and at least one $\hat{y}$-view track were required. Two $\hat{x}$-view tracks were required because the momentum

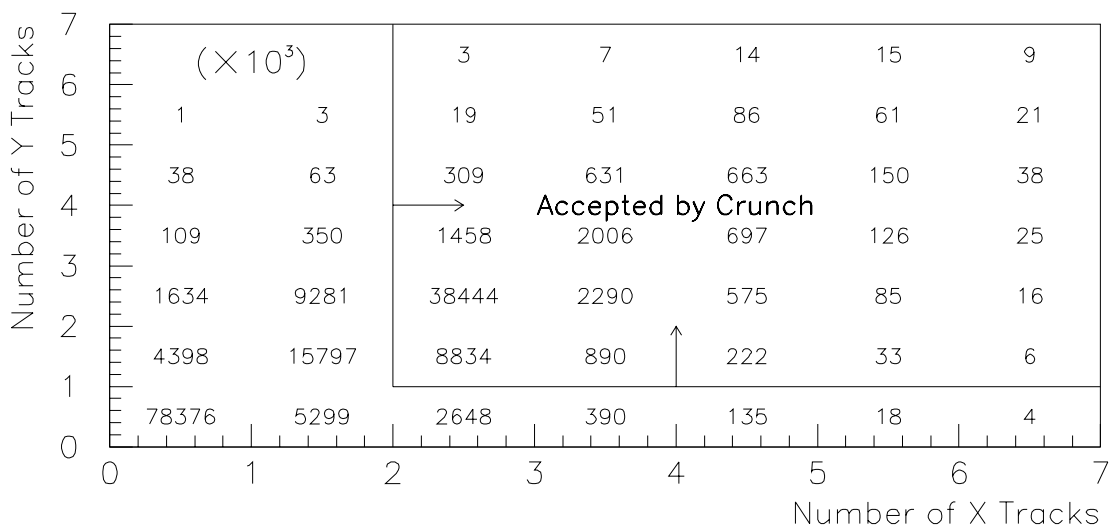

Figure 102. Number of $\hat{y}$-view tracks versus the number of $\hat{x}$-view tracks for all events analyzed in the Two-Electron crunch.

of each electron must be measured in order to use E/p for particle identification. In the case of the $\hat{y}$-view tracking however, only one track was required to allow for the possible, albeit unlikely, case where the $\hat{y}$-view tracks of two electrons overlapped so completely in the chamber that only one track could be found. In this case, the single $\hat{y}$-view track was shared between the two electron candidates. This requirement on the number of tracks found drastically reduced the size of the event sample. Only 1/3 of the events in the Two-Electron stream passed this cut, despite the hit-counting requirement in the trigger.

Once tracks and clusters were identified, the next step was to use the matching algorithm (Section 4.3) to identify the cluster associated with each track. Because of the requirement that track segments be consistent with pointing to lead-glass clusters in the track-finding, most events (78%) contained at least one match of tracks and clusters resulting in two electron candidates.

Once electron candidates were identified by the matching, the ratio of the cluster energy

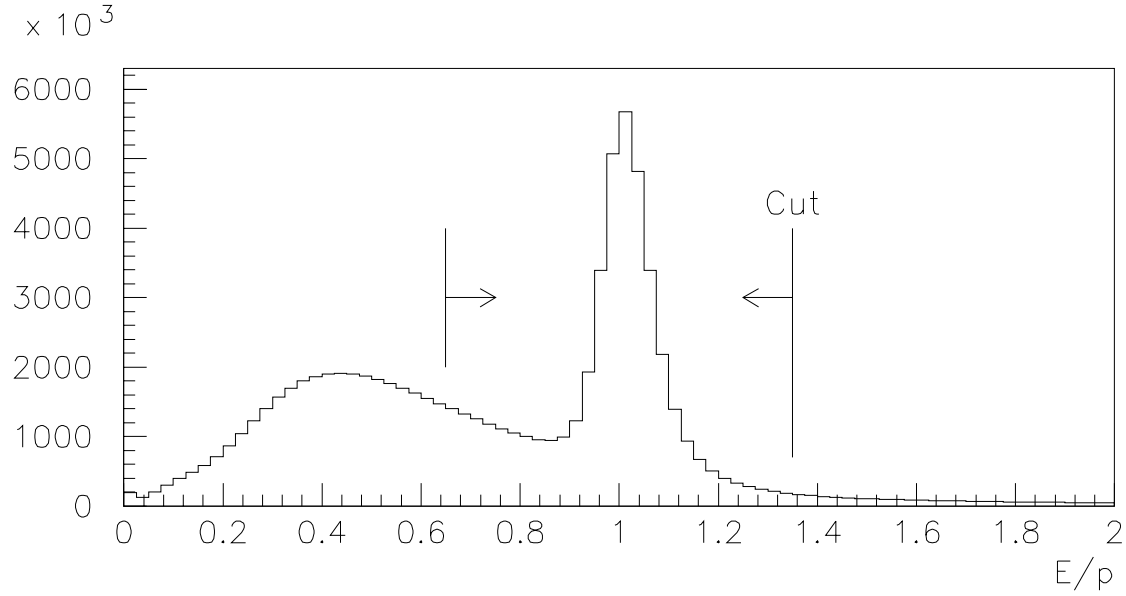

Figure 103. E/p for events in the Two-Electron stream. The electron identification cut used in the crunch is shown on the plot. The peak is at E/p of 1 is from electrons, and the background is from charged pion showers and cluster fusions.

Table 24. Fraction of events passing each stage of the crunch

Cut	Fraction Kept in Data	Fraction Kept in $\pi^0 \to e^+e^-$ Monte Carlo
Track-Finding	0.33	0.975
Track-Cluster Matching	0.78	0.992
E/p Cut	0.42	0.971

to the track momentum, E/p, was calculated. For electrons which pointed into the inner half of a pipe block, E/p was not considered since there was a possibility of large leakage out the edge of the lead-glass array in such showers. The distribution of E/p for all events in the Two-Electron steam is shown in Figure 103. For a track to be identified as an electron candidate, E/p must have been between 0.65 and 1.35.

For the two tracks in the event, one can plot E/p of one track against that as of the other to identify the most likely parent decay of each event as shown in Figure 104. Using such a classification with a tight E/p cut, about 22% of these Two-Electron stream events appear to have two good electrons. 41% of the events were consistent with having two final-state charged pions and were identified as $K_L \to \pi^+\pi^-\pi^0$ candidates. 37% had one good electron and one charged pion candidate, and mostly likely resulted from semi-leptonic K_L decays. Of course, without the trigger requirement of HCF-found clusters for electrons and without the hadron shower veto from the MU1 bank, the fraction of events containing two electrons would have been much lower.

To minimize accidental losses in the data reduction, an event was saved if *any* valid match of tracks to clusters resulted in two electron candidates with good E/p. Recall from Section 4.3 that possible matches of tracks to clusters were ranked by the sum of the squared distances between track and cluster positions. Approximately 38% of the events passed the E/p cut with the best match. Only 3.6% of events passed the E/p cut with an alternate matching of tracks to clusters.

The outcome of the crunch is shown in Table 24. In all, a reduction factor of 9.3 was

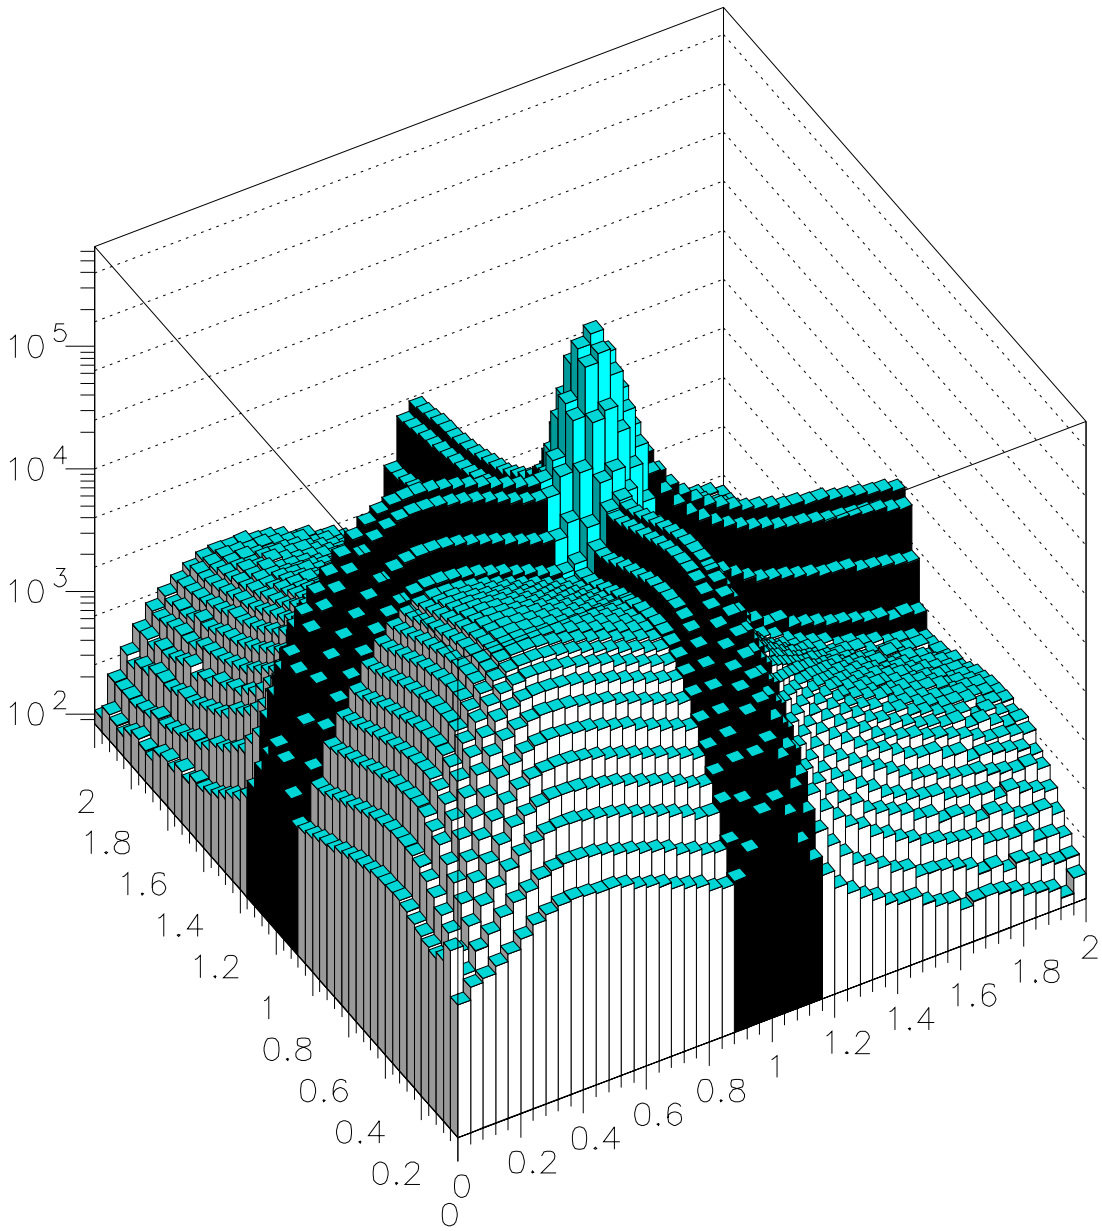

Figure 104. E/p vs. E/p for the two tracks in events in the Two-Electron stream. The events in the center light-colored peak are from events with two real electrons. Events in the dark bands are those with one electron and one charged pion, most likely from semi-leptonic decays. Events in the light-shaded region outside the peak mostly likely result from $K_L \to \pi^+\pi^-\pi^0$ decays.

achieved in the data while only losing approximately 5% of $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, events generated in Monte Carlo which satisfied the trigger requirements.

8.1.1 Monitoring the Crunch

Because of the immense volumes of information produced by the crunch, an automatic monitoring procedure was used to check key information from each data run. Variables related to track quality were checked to verify that the chamber calibration was adequate for each run, and similarly electron E/p was monitored to check the quality of glass calibrations. Because of the extremely high statistics processed in the crunch, it was also possible to check for non-functioning or inefficient HCF seed blocks and trigger counters. This monitoring also

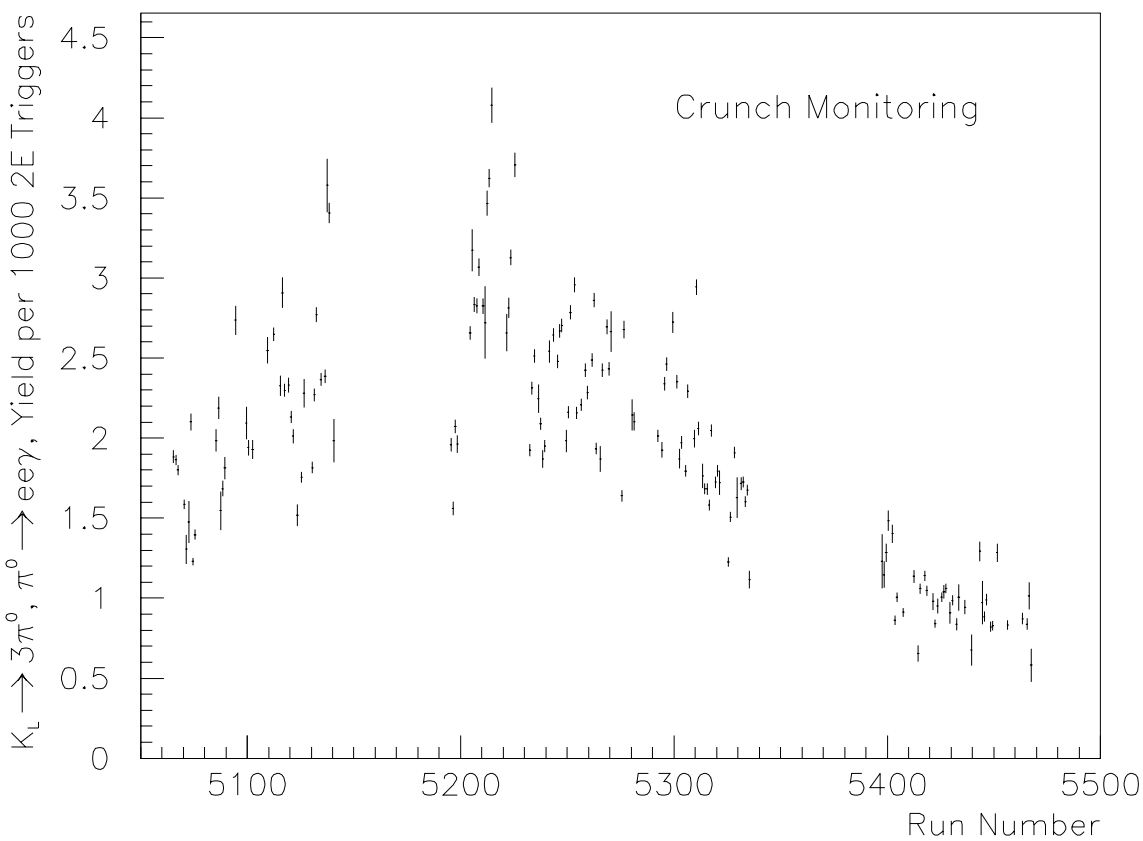

Figure 105. The Yield of $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, by Run

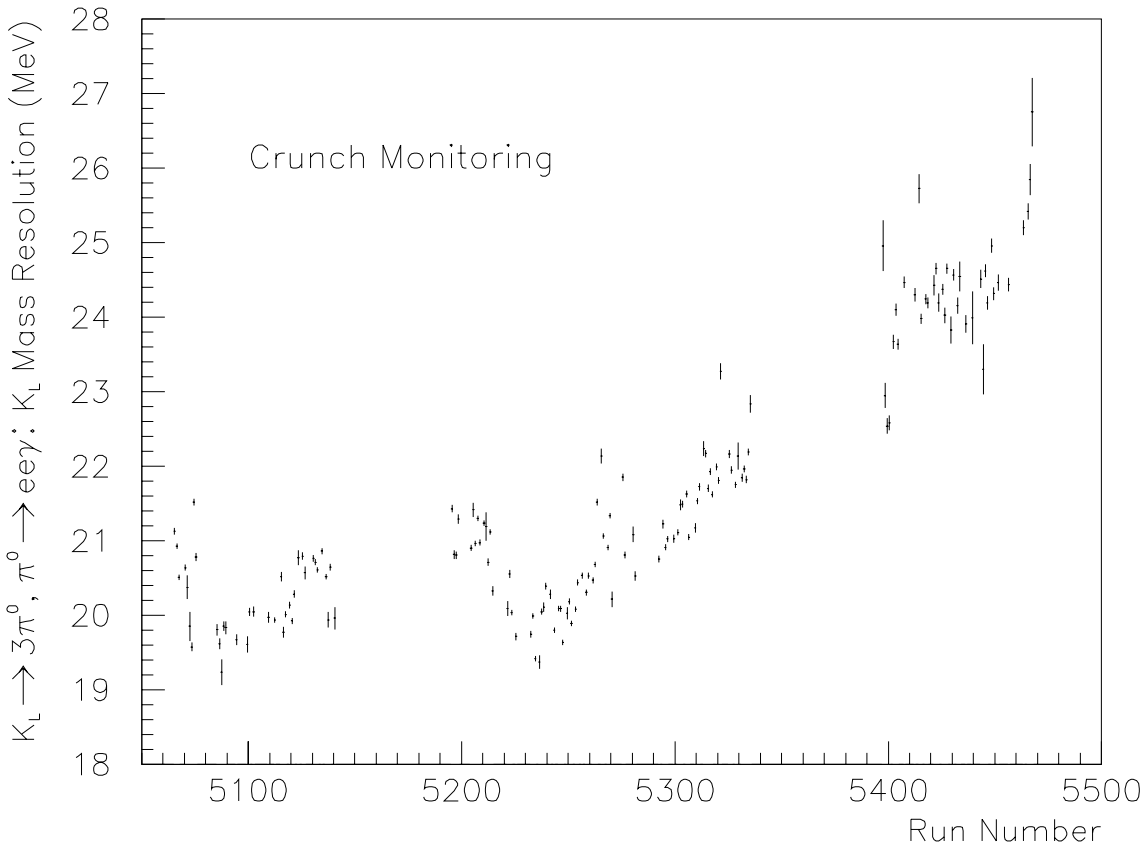

Figure 106. The K_L Mass Resolution in $K_L \rightarrow \pi^0\pi^0\pi^0$, $\pi^0 \rightarrow e^+e^-\gamma$, by Run

provided a graphic demonstration of the effects of radiation damage to the lead-glass on the event yield and kinematic resolution. The yield per trigger of $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays reconstructed using loose cuts, is shown in Figure 105. It is interesting to note the large variations from run-to-run, which were primarily due to variations in beam intensity, and the overall downward trend as the run progressed because of radiation damage in the calorimeter. This damage is further illustrated in Figure 106, where the mass resolution from these reconstructed $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, is shown to have degraded substantially over time.

8.2 The Analysis

Once the crunch was completed and checked by the monitoring procedure, the event sample was sufficiently small (19 million events, 20 Gbytes) that the trial-and-error business of extracting the $\pi^0 \to e^+e^-$ sample could begin in earnest. Also, the high-statistics monitoring of the crunch gave confidence in our understanding and calibration of the detector, especially for time-varying effects. Thus, it was now possible to lump the entire data sample together without being overly worried about diluting important run-dependent effects.

The goal of the analysis was to extract background-free samples of the signal mode, $\pi^0 \to e^+e^-$, and the normalization mode, $\pi^0 \to e^+e^-\gamma$ with an e^+e^- invariant mass greater than 70 MeV/c^2. Each cut imposed on the data sample will be motivated by examining data and Monte Carlo. After all cuts have been described, the data samples will then be examined, providing a natural lead-in to a discussion of backgrounds.

8.2.1 Analysis of Tracks

The first analysis cuts beyond those in the crunch were made on the electron tracks. The initial cut made in the analysis was the removal of events with only one track found in the $\hat{y}$-view. This was done because of the large chance of misreconstructing such events by assigning the wrong $\hat{x}$-view track to a cluster. Such misreconstructed events would have the correct invariant mass, but will have a corrupted e^+e^- momentum measurement in the $\hat{y}$-direction. In addition, although approximately 18% of the events passing the crunch did

share a track in the $\hat{y}$-view, very few of these events were likely to be $\pi^0 \to e^+e^-$ decays, since for typical π^0 energies in $\pi^0 \to e^+e^-$, about 25 GeV, the mean track opening angle in the $\hat{y}$-view was approximately 5 mrad. This cut resulted in only a 0.3% acceptance loss in $\pi^0 \to e^+e^-$ Monte Carlo.

An important step in the tracking analysis was the verification of the trigger. The Level 1 and track processor triggers placed certain requirements on the topology of the electron tracks. In particular, the tracks had to pass through two different counters in each of the B and C trigger banks, and tracks were required to make two hits or hit pairs in every drift chamber to satisfy the Two Electron trigger. Some of the data on tape did not satisfy these triggers because of the electron tracks, but rather because of other activity in the detector, such as accidentals. These triggers were problematic in the analysis, because they relied heavily on the details of the Monte Carlo modeling of accidentals in order to correctly simulate these events. Therefore, all events in both the data and Monte Carlo samples were required in software to satisfy the B and C bank triggers with only drift chamber hits and scintillator latches that resulted from the electrons. This trigger verification resulted in a 2.1% loss from the track processor hit-counting verification, and a 5.0% loss from the B and C trigger bank verification. The loss from the trigger bank verification included the removal of about 0.4% of the data during which a part of the B-bank was not functioning.

Another important class of events to remove from the sample were those events with hits on wires whose efficiencies were not correctly simulated by the Monte Carlo. Recall that the procedure with which chamber efficiencies were determined did not take into account all possible sources of wire inefficiencies, such as those due to the track processor trigger (see Section 5.7). Therefore, it was important to check the Monte Carlo simulation of efficiencies. The complete $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, sample was used for this comparison. A sample comparison in which a disagreement was evident is shown in Figure 107. This disagreement was due to a dead cell in the hit-counting portion of the track processor which was not simulated in the Monte Carlo. Events with track hits in this cell or on another sporadically inefficient wire were removed from the data sample.

The next cuts were designed to tighten the loose quality cuts made in the tracking algorithm. After all corrections to the hit positions, the $\hat{x}$-view and $\hat{y}$-track segments were

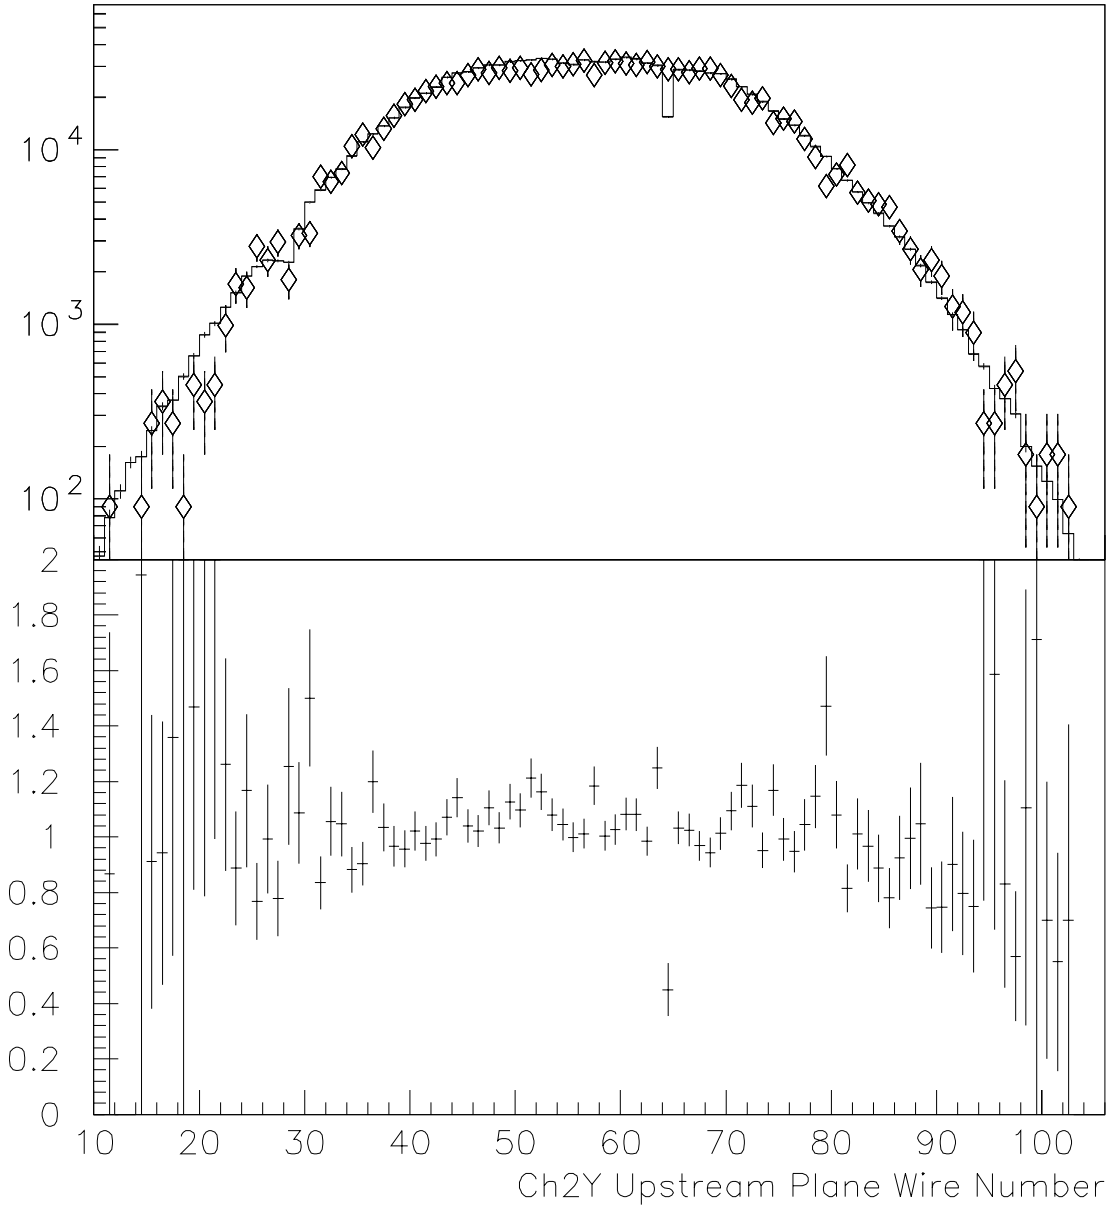

Figure 107. Data and Monte Carlo comparison of wire illumination for the upstream plane of the $\hat{y}$-view of chamber 2 for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays. The disagreement shown was due to an inefficient cell in the track processor trigger which was not modeled in the Monte Carlo.

refit, and the results of these new track fits were checked to ensure good reconstructed track quality. Such a check was important because poorly reconstructed tracks could degrade the resolution of track momentum and direction measurements which were used to calculate the invariant mass. Track quality cuts were only made on the two tracks identified as electron candidates in each event.

The first quality cut was made on the track segment χ^2, which is similar to the track χ^2 discussed in Section 4.2, and is defined as

$$\chi^2 = \sum_{\text{hits}} (r_{\text{fit}} - r\text{hit}), r \in \{x, y\},$$

where r_{fit} denotes the fitted track segment position at the sense plane, and r_{hit} is the position measured from the hit. Note that this is not a "true" χ^2 since it is not normalized by the resolution of this hit position measurement, approximately 100 μm. Because the track segment "fit" was equivalent to drawing a straight line between the average position at the two chambers, this χ^2 was essentially a sophisticated version of the sum-of-distances measurement discussed in Section 5.3.2. In fact, this segment χ^2 can be written as

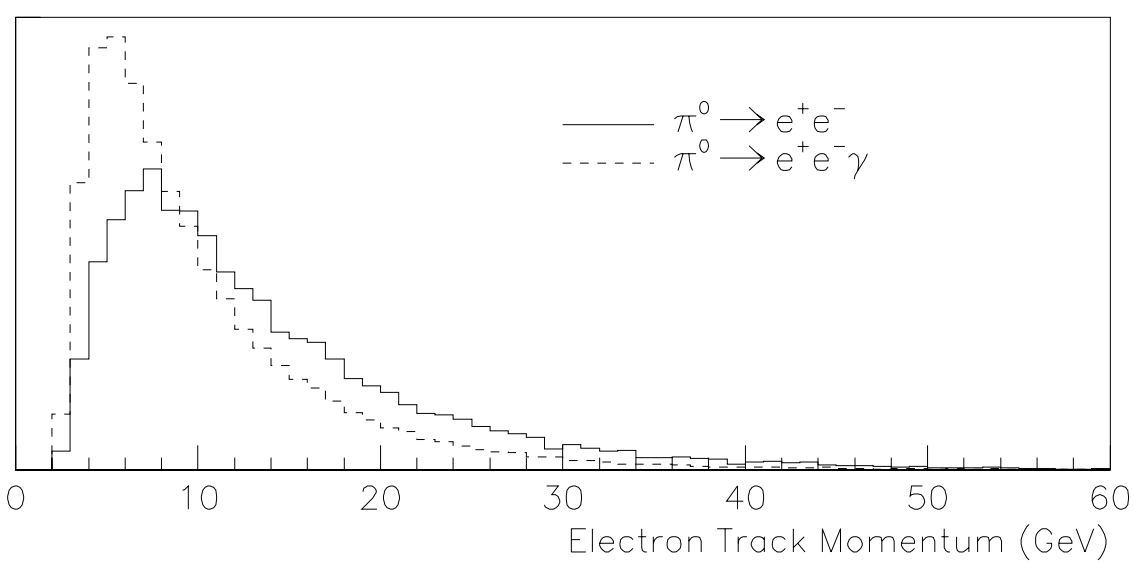

Figure 108. The electron momentum spectra for $K_L \rightarrow \pi^0\pi^0\pi^0$, $\pi^0 \rightarrow e^+e^-$, decays and for $K_L \rightarrow \pi^0\pi^0\pi^0$, $\pi^0 \rightarrow e^+e^-\gamma$, decays in Monte Carlo

$$\chi^2 = \frac{1}{2} \sum_{\text{hit pairs}} (d_1 + d_2 - d_c - \theta\Delta z)^2,$$

where d_i are the drift distances of the two hits, d_c is the 6.35 mm cell size, and $\theta\Delta z$ is a correction for the effect of the finite distance (Δz) between the two sense planes, and the track angle, θ. The mean χ^2 for 100 μm single-hit resolution is thus 2×10^{-8} m^2.

The second quality cut was made on the track magnet offset, which is simply the difference between the upstream and downstream track extrapolations to the magnet "bend plane", as defined in Section 5.6.2. The error in a measurement of this magnet offset was

$$\sigma_{\text{magnet offset}} = \sqrt{(170\ \mu\text{m})^2 + \left(\frac{2.1\ \text{mm}}{p\ (\text{inGeV})}\right)^2},$$

where the first term is from chamber resolution and the second from multiple scattering. The multiple-scattering term clearly dominated as low electron energies; however, as is shown in Figure 108, average electron momenta were sufficiently high that both terms contributed.

Distributions of the segment χ^2's and magnet offsets are shown in Figures 109 and 110 in the normalization data, the normalization Monte Carlo and the signal Monte Carlo. The signal data itself is not shown because of the expected low statistics and high backgrounds. For the segment χ^2, the Monte Carlo and data disagreed significantly. This was, of course, to be expected given that the Monte Carlo did not attempt to simulate the δ-rays which were responsible for much of the non-Gaussian shape of the sum-of-distances distribution. However, because the signal and normalization modes both contained two final state electrons, had similar electron energy spectra (see Figure 108), and hit roughly the same wires in the drift chambers (see Figure 111), effects from δ-rays and other misreconstructions not simulated in the Monte Carlo should cancel in the ratio of acceptances for the two modes. The agreement between normalization data and Monte Carlo in the magnet offsets was much better, probably because the tails of multiple scattering from low momentum tracks which were simulated in the Monte Carlo dominated the tails of the magnet offset distribution.

The analysis cuts were chosen to ensure reasonably well-measured tracks without sacrificing too many events. The track segment χ^2 cut was chosen to be 15×10^{-8} m^2, which is approximately the χ^2 value for a track segment with two sums-of-distances 3σ away from

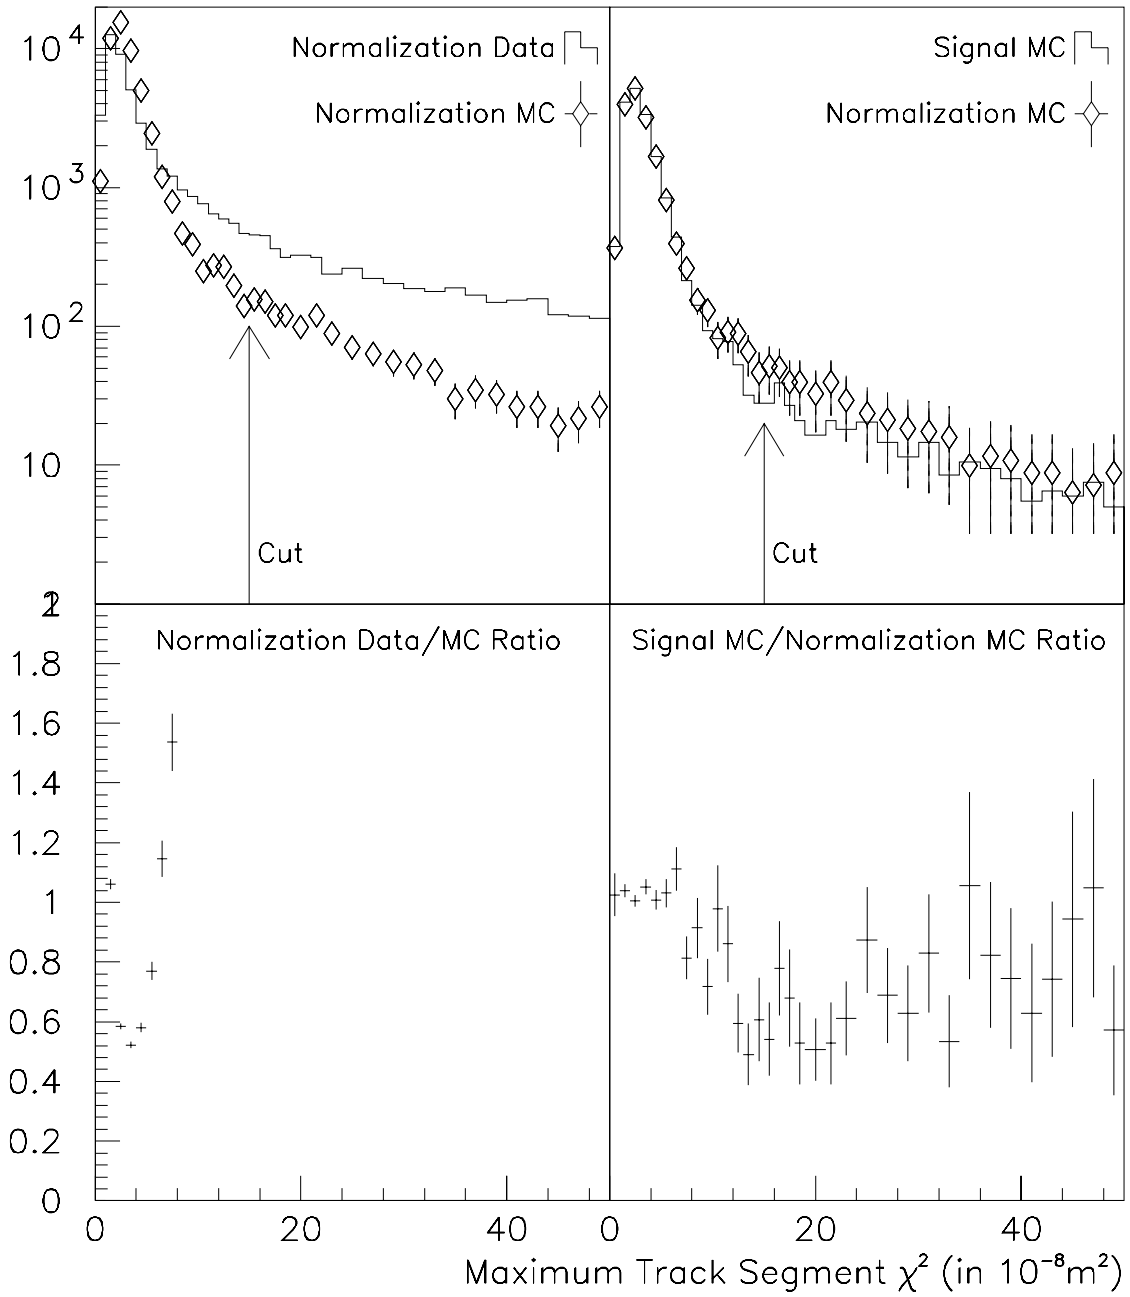

Figure 109. The maximum track segment χ^2 of the electron track segments in an event is shown in data and Monte Carlo. The arrow indicates the location of the analysis cut

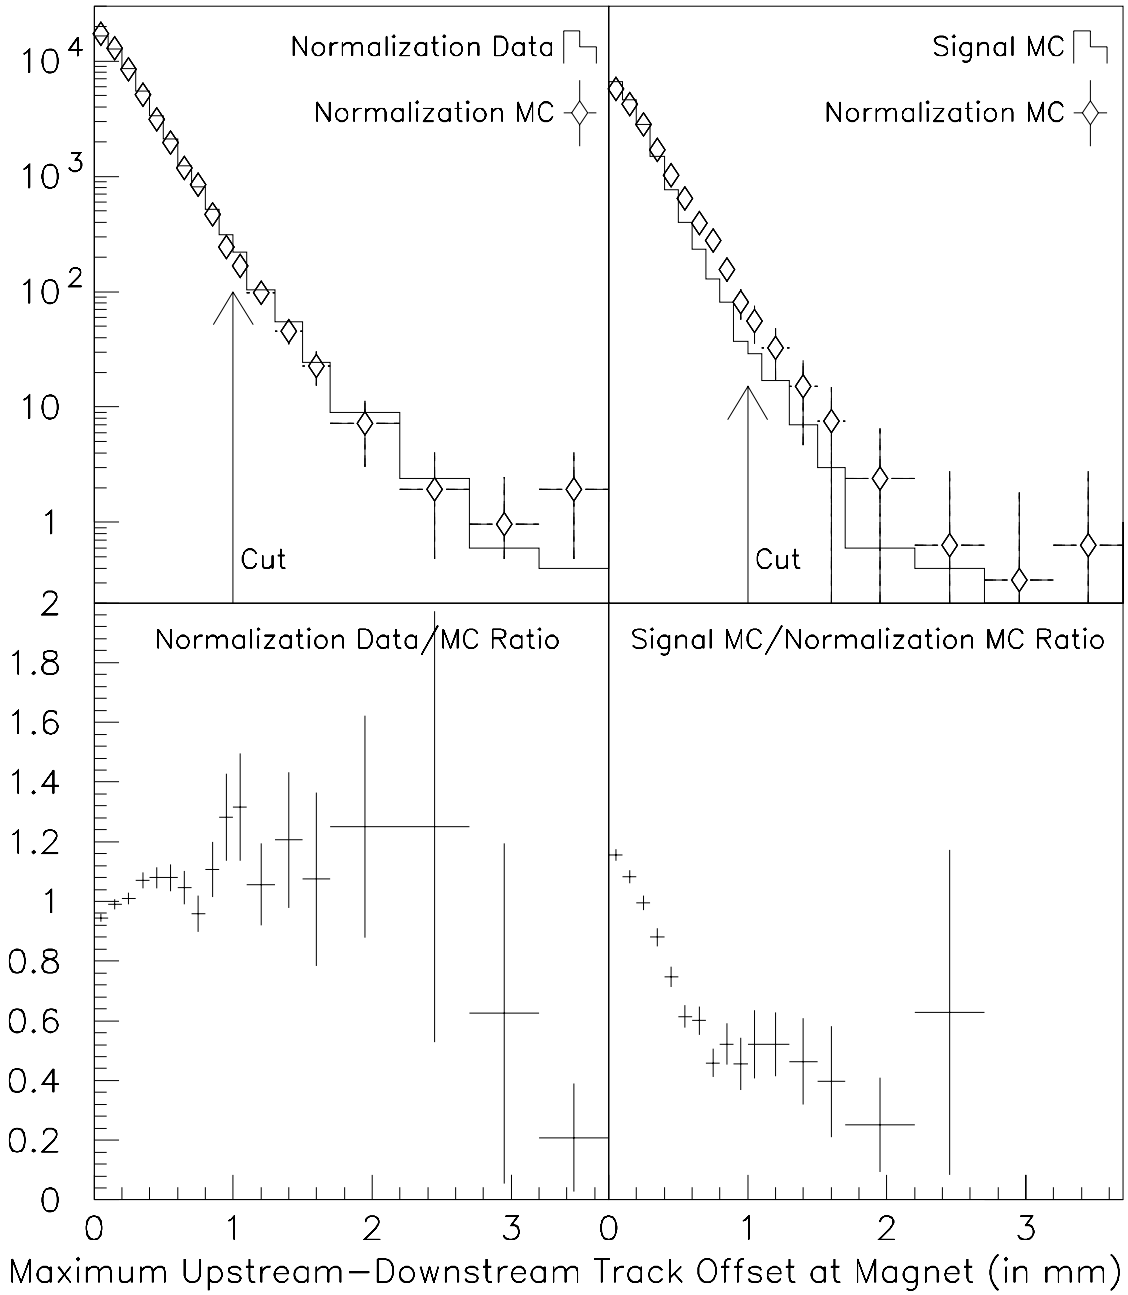

Figure 110. The maximum track offset at the magnet bend plane of all electron tracks in an event is shown in data and Monte Carlo. The arrow indicates the location of the analysis cut

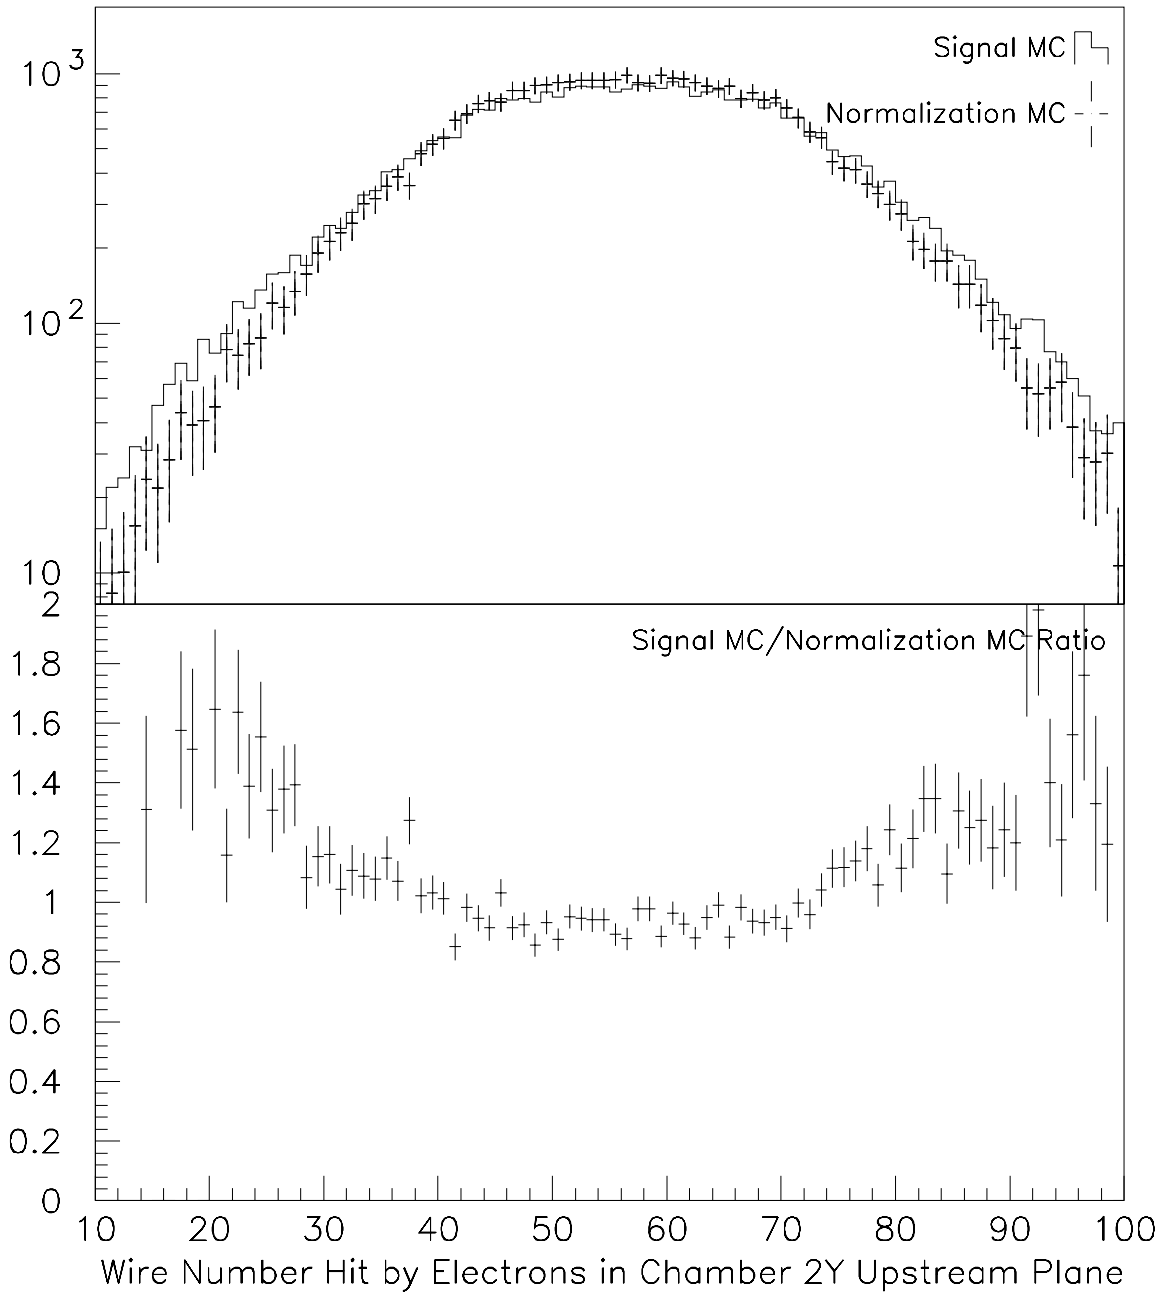

Figure 111. The population of hits in each wire of the upstream plane of the $\hat{y}$-view of drift chamber 2 for $\pi^0 \to e^+e^-$ and $\pi^0 \to e^+e^-\gamma$ normalization Monte Carlo

the cell size. The segment χ^2 cut removed 19% of the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, candidates in the data sample. This cut reduced the acceptance substantially, 10% per electron track, but it played an important role in reducing track position measurement errors. The cut on magnet offset was conceptually more difficult. Because of the strong dependence of resolution on momentum, a momentum dependent cut seemed natural. However, at this point in the analysis, the $\hat{x}$-view tracking algorithm already had ranked track quality using this variable without regard to track momentum. Because the tracking made what was effectively a cut in this variable without regard to track momentum, I felt it was aesthetically cleaner to not make momentum-dependent cut in this variable. A loose cut of 1 mm was chosen, in order to avoid removing too many of the low momentum events. This cut reduced the acceptance by 1.2% in both the normalization data and Monte Carlo samples, and by only 0.5% in the $\pi^0 \to e^+e^-$ Monte Carlo.

Once good quality tracks were chosen, the next step was locating the decay vertex using the electron tracks. This was done in two steps. First, the location of the minimum distance between the two electron tracks was calculated. Then, using this location as a starting point for the fit, the two upstream tracks were refit, using the constraint of a common vertex. This fit took the eight measured positions of the two tracks in the upstream chambers and returned the three space dimensions of the best fit vertex point, along with four new track angles for the two $\hat{x}$-view and two $\hat{y}$-view upstream segments. The fit also returned a χ^2 and errors on the fit parameters. This procedure failed approximately 4.9% of the time in Two-Electron triggers, either because the tracks were too close to being parallel in one view to compute the point of closest approach, or because the vertex fit returned a vertex well outside of the decay volume. The distribution of the vertex χ^2 is shown in Figure 112. The cut chosen on this χ^2 was extremely loose, 15, which represents a 0.01% confidence level cut. This cut had no appreciable loss in data or Monte Carlo for the signal or normalization samples, which only contained events with a high e^+e^- invariant mass. However, for the low e^+e^- invariant mass Dalitz decays which represented the bulk of events with real electrons in the Two Electron trigger, there was a very long tail in the vertex fit χ^2, and this cut would have removed an astonishing 21% of this data sample!

With the tracking cut finalized and the decay vertex measured, the invariant mass of

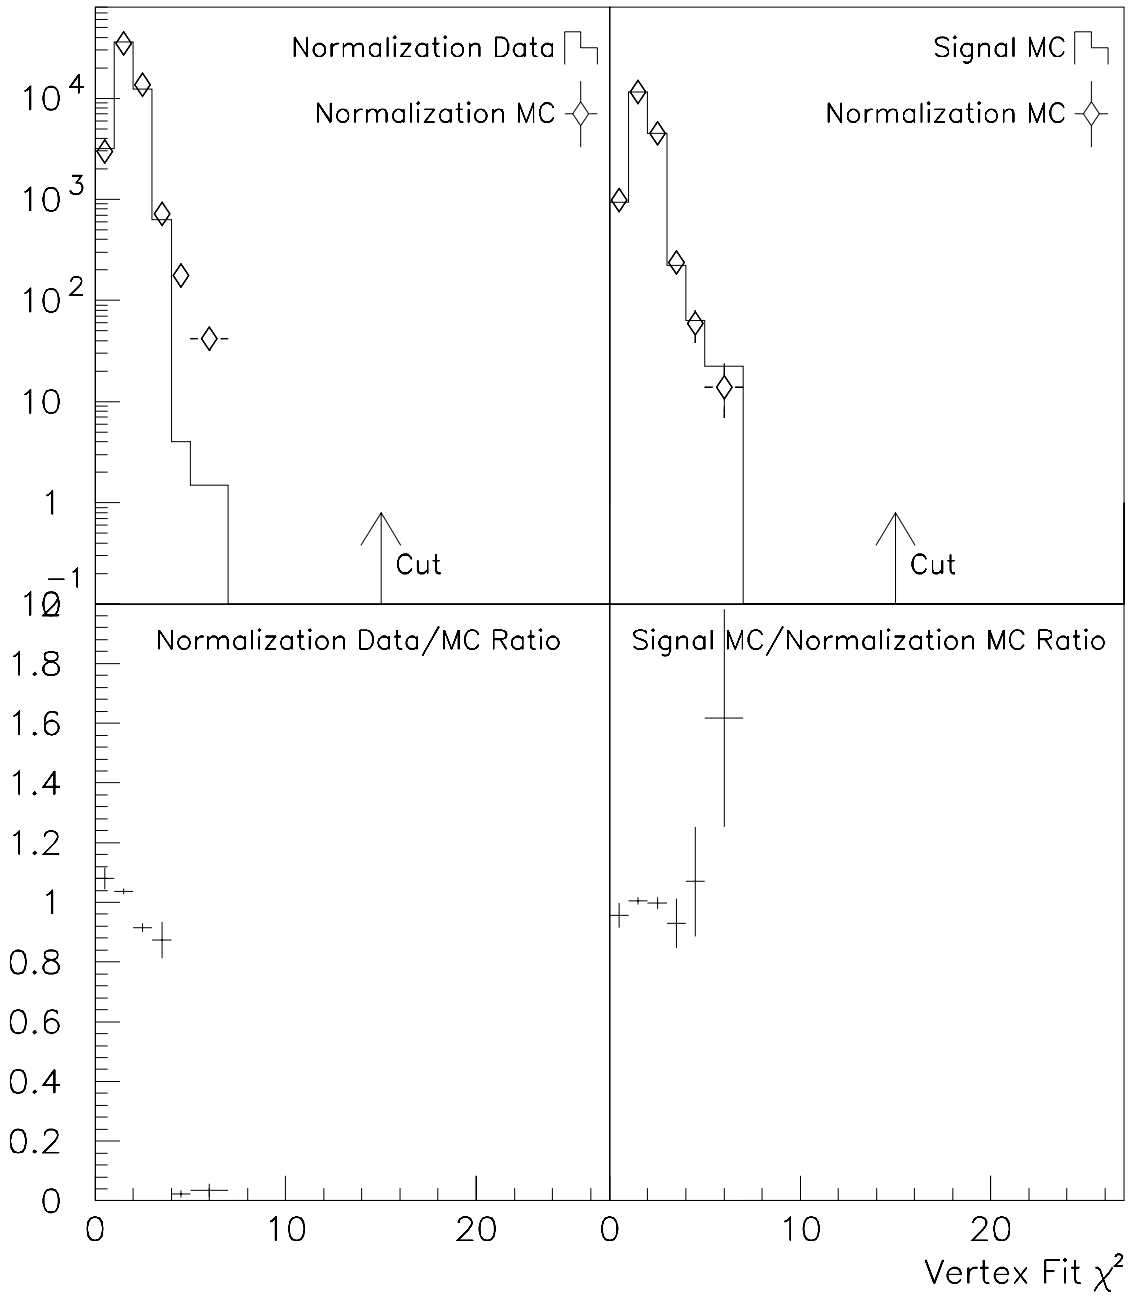

Figure 112. The vertex fit χ^2 is shown in data and Monte Carlo. The arrow indicates the location of the very loose analysis cut

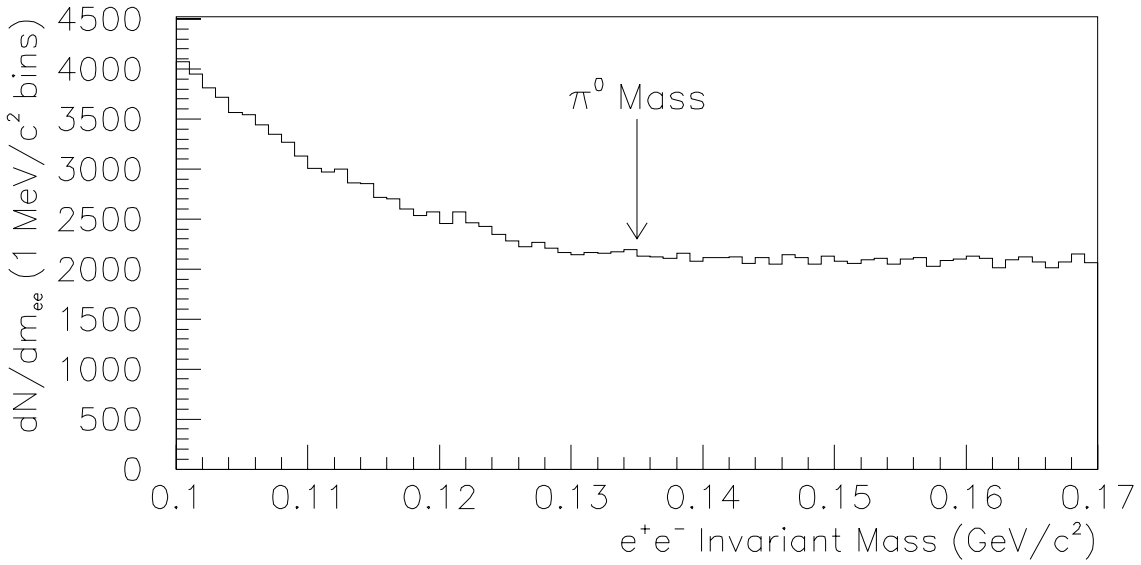

Figure 113. The e^+e^- invariant mass of the data sample after trigger, crunch and tracking requirements. Obviously, the full reconstruction of the parent K_L decay is essential for this analysis!

the two electrons in the event could be calculated. At this point, the size of the data sample was reduced dramatically (about a factor of 100) by requiring an invariant mass between 68 and 240 MeV/c^2 in both samples (the π^0 mass is 135.0 MeV/c^2). All plots shown in this section, and in all subsequent sections use data samples with only these high e^+e^- mass events. One can actually look at the e^+e^- invariant mass in the data at this stage to search for any sign of a peak at the π^0 mass. The only cuts on the events shown were the trigger requirements (including a high lead-glass cluster multiplicity requirement), the cuts in the crunch, including the loose electron identification requirement, and cuts on the track quality. The data are shown in Figure 113. Clearly, there was a formidable background problem remaining. It was the full reconstruction of the parent K_L decay which allowed for real π^0 decays to be distinguished from other background processes.

8.2.2 Analysis of Calorimeter Clusters

Although the tracking was arguably more important to this analysis, the information from the calorimeter was needed for the reconstruction of the two $\pi^0 \to \gamma\gamma$ decays from the parent $K_L \to \pi^0\pi^0\pi^0$ decay. For the normalization mode, the calorimeter was also used to find the photon from the $\pi^0 \to e^+e^-\gamma$ decay. Also, information from the calorimeter was used for electron identification, and for removing backgrounds with extra photons from the $\pi^0 \to e^+e^-$ sample, such as those from $\pi^0 \to e^+e^-\gamma$ decays.

For π^0 reconstruction, the calorimeter information was used in the following way: every cluster in the lead-glass calorimeter found by the HCF that was *not* identified as an electron cluster was classified as a photon candidate. Events with HCF cluster fusions, described in Section 4.1, were removed from the sample. Cuts were also placed on photon positions to ensure that energies were well measured. If the photon position was measured to be within 2.9 mm of a pipe block, the pipe block energy correction routine described in Section 6.2.1 failed, and events with such a photon were discarded. Photon paths were also interpolated between their position at the glass and the charged vertex to the z-position of the Iron Ring. The distance between the photon position as it passed through the Iron Ring and the $\hat{z}$-axis was then required to be less than 83.5 cm to remove events with photons interacting in the ring. The data and Monte Carlo photon radius distributions in the $K_L \to \pi^0\pi^0\pi^0$,

$\pi^0 \to e^+e^-\gamma$, sample before such a cut was made are shown in Figure 54 and described in Section 5.5.4.

In $\pi^0 \to e^+e^-$ events, the four photons were combined to form two π^0's. In $\pi^0 \to e^+e^-\gamma$ decays, the five real photons plus the virtual photon (observed in the lab as an e^+e^- pair) were paired to make three π^0 candidates. For each possible pairing of photons into π^0's, a normalized χ^2 for the hypothesis that the photons originated from π^0 decays was calculated from the expression

$$\chi^2 = \frac{1}{n_{\text{pairs}}} \sum_{\text{photon pairs}} \left(\frac{m_{\gamma\gamma} - m_{\pi^0}}{\sigma_{m_{\gamma\gamma}}} \right)^2 .$$

The two-photon invariant mass was calculated by

$$m^2_{\gamma\gamma} = 2E_{\gamma_1} E_{\gamma_2}(1 - \cos\theta_{12}) \approx E_{\gamma_1} E_{\gamma_2} \left(\frac{r_{\text{PbG}}}{z_{\text{PbG}} - z_{\text{vertex}}} \right)^2 ,$$

or

$$m^2_{\gamma\gamma^*} \approx m^2_{\gamma^*} \left(1 + \frac{E_\gamma}{p_{\gamma^*}} + \frac{E_\gamma m^2_{\gamma^*}}{4p^3_{\gamma^*}} \right) + E_\gamma p_{\gamma^*} \left(\frac{r_{\text{PbG}}}{z_{\text{PbG}} - z_{\text{vertex}}} \right)^2 ,$$

where r_{PbG} is the distance of separation of the photons at the lead-glass calorimeter, $z_{\text{PbG}} - z_{\text{vertex}}$ is the distance between the decay vertex and the calorimeter, as measured from the best fit charged track vertex, and m_{gamma^*} is the invariant mass of the e^+e^- pair observed in the lab. The error estimation for this invariant mass calculation included errors from the photon energy measurement, photon position measurements, track position and momentum measurements if applicable, and the error in the charged vertex measurement. From the three possible pairings of photons into two π^0's for the $\pi^0 \to e^+e^-$ candidates or from the 15 possible pairings into three π^0's in the $\pi^0 \to e^+e^-\gamma$ candidates, the pairing with the lowest χ^2 was chosen. Cuts made on the photon pairing χ^2 will be discussed in Section 8.2.4

For improved electron identification, the cut on E/p of electrons was tightened from the cut used in the crunch. As in the crunch, the candidate electrons were chosen by finding a matching of tracks to clusters with $0.65 < E/p < 1.35$ for both electron tracks. Again, as in the crunch, no cut was placed on electrons pointing into the inner half of a pipe block because of the potentially large shower leakage out of the array. Taking the matching of tracks to clusters chosen with the loose E/p cut, the cut was then tightened to $0.85 < E/p < 1.15$. Electron E/p is shown for the normalization data and for signal and

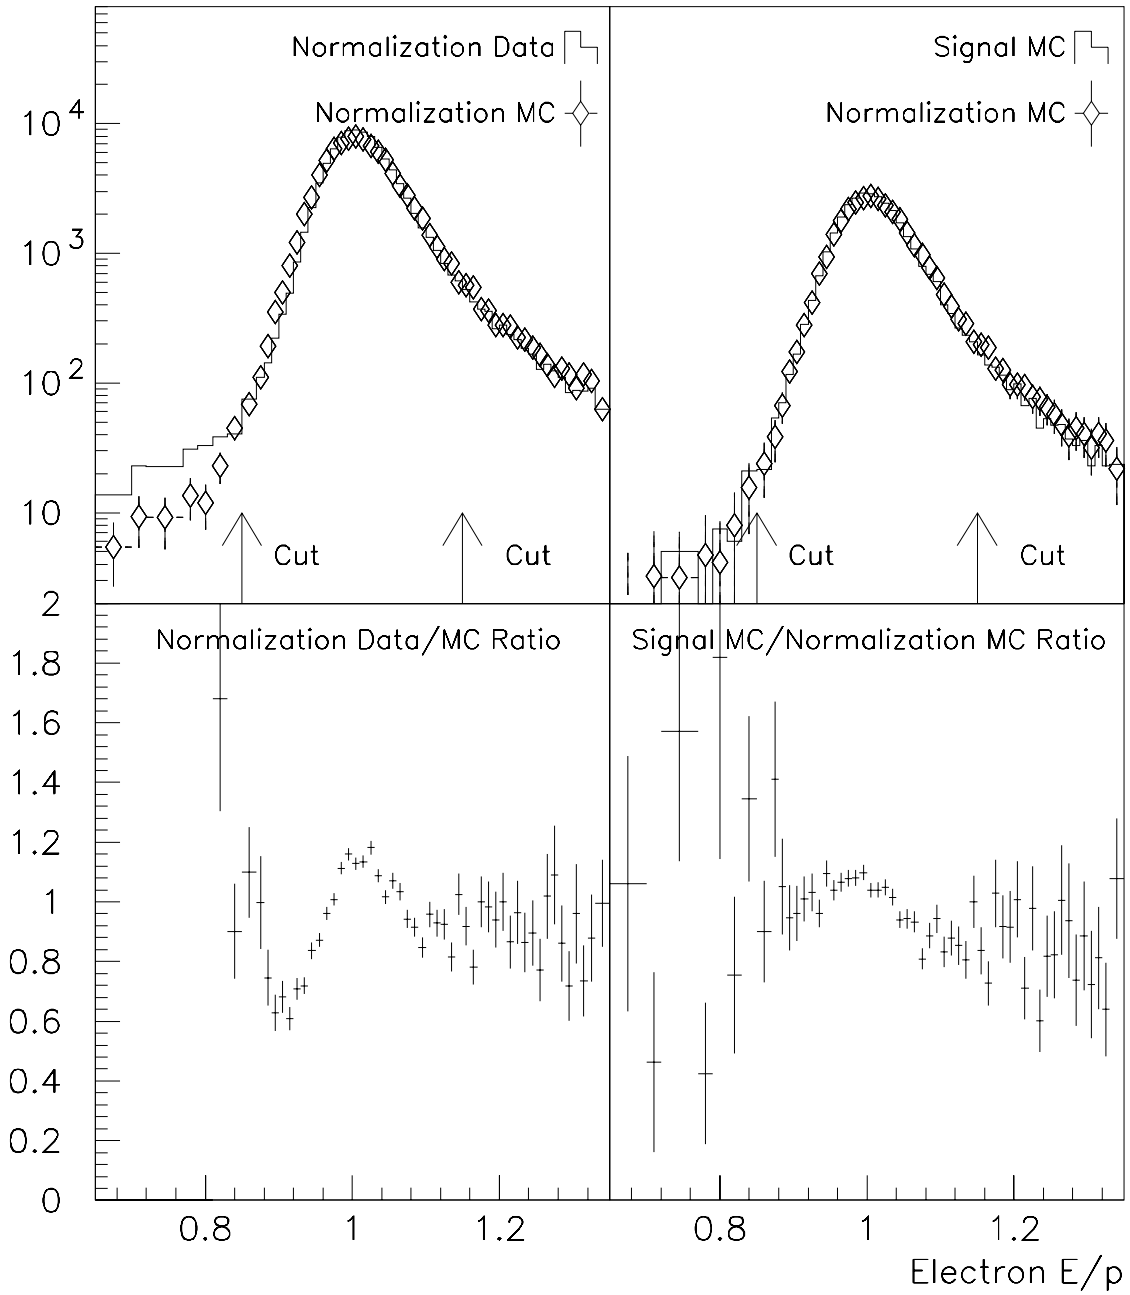

Figure 114. The E/p for electrons in is shown in data and Monte Carlo. Each event has two entries, one for each electron. The arrow indicates the location of the analysis cut. The data and Monte Carlo samples are shown after kinematic cuts designed to select $K_L \to \pi^0\pi^0\pi^0$ decays.

normalization Monte Carlo in Figure 114. The agreement between data and Monte Carlo in the normalization sample was reasonably good. There was though some disagreement at very low E/p, and the peak was sharper in data than in Monte Carlo. However, in the range above the high-side E/p cut, where it was most important for the data and Monte Carlo to agree in order to correctly predict losses from the E/p cut, the agreement was reasonable. The acceptance loss from this cut in the normalization data was 4.7% per electron, and in normalization Monte Carlo it was 4.9%. For this price in acceptance, the probability of a single pion satisfying the E/p cut drops from 10.5% with the crunch cut to 3.6% with this analysis cut (see Figure 115), as measured in $K_L \to \pi^+\pi^-\pi^0$ decays.

A series of cuts were then applied to remove events with extra photons in the lead-glass calorimeter which were not found by the HCF or by the cluster-finding routine. Such events were undesirable because the dominant backgrounds to $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, had extra photons in the final state. Also photon energies and positions could be badly mismeasured if a "photon" actually consisted of two nearby photons fused in the cluster-finding algorithm. The combination of these two occurrences was particularly dangerous, since rejection of backgrounds often depended on kinematic cuts designed to discriminate against missing final state particles. However, if the "missing" photon were fused with another photon in the event, its energy was not lost but simply added to the energy of another photon in the event. Many of these kinematic cuts were ineffective in such a case.

The initial cuts applied to remove such events were the fusion cuts described in Section 6.5. Each cluster in an event was checked first for anomalously large energy deposits in the 5×5 ring of blocks around cluster centers, then for inconsistent column and row energy sums in the central 3×3 blocks of clusters, and finally for excess energy in the corners of the 3×3's. The results of these cuts on data and Monte Carlo samples are shown in Table 25.

An additional cut was made to remove soft photons not located near other clusters. The cut relied on information from the lead-glass adders in order to ensure that the lead-glass energy deposition occurred in the same bucket as the trigger. For every adder which did not contain a block from the center 3×3 array of blocks in a calorimeter cluster, the ratio of the adder signal to the sum of the lead-glass ADC counts in that adder was required to be between 0.6 and 1.4. If this condition was satisfied, it was likely that the bulk of the lead-

Table 25. Effects of fusion cuts on data and Monte Carlo

Sample	Fraction lost in 5×5 cut	Fraction lost in column cut	Fraction lost in corner cut
	for electrons/photons		
Normalization data	0.045/0.285	0.005/0.037	0.0002/0.003
Normalization Monte Carlo	0.036/0.236	0.003/0.033	0.0002/0.002
Signal Monte Carlo	0.044/0.222	0.005/0.031	0.0003/0.002

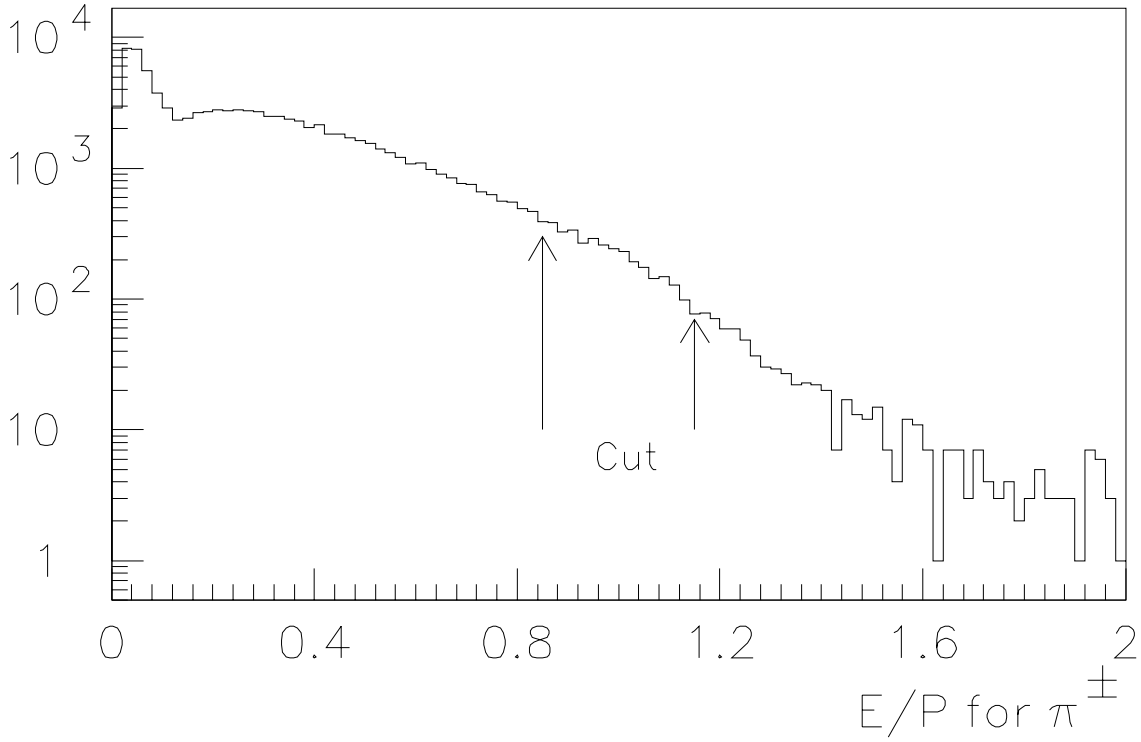

Figure 115. E/p for pions from $K_L \to \pi^+\pi^-\pi^0$ decays

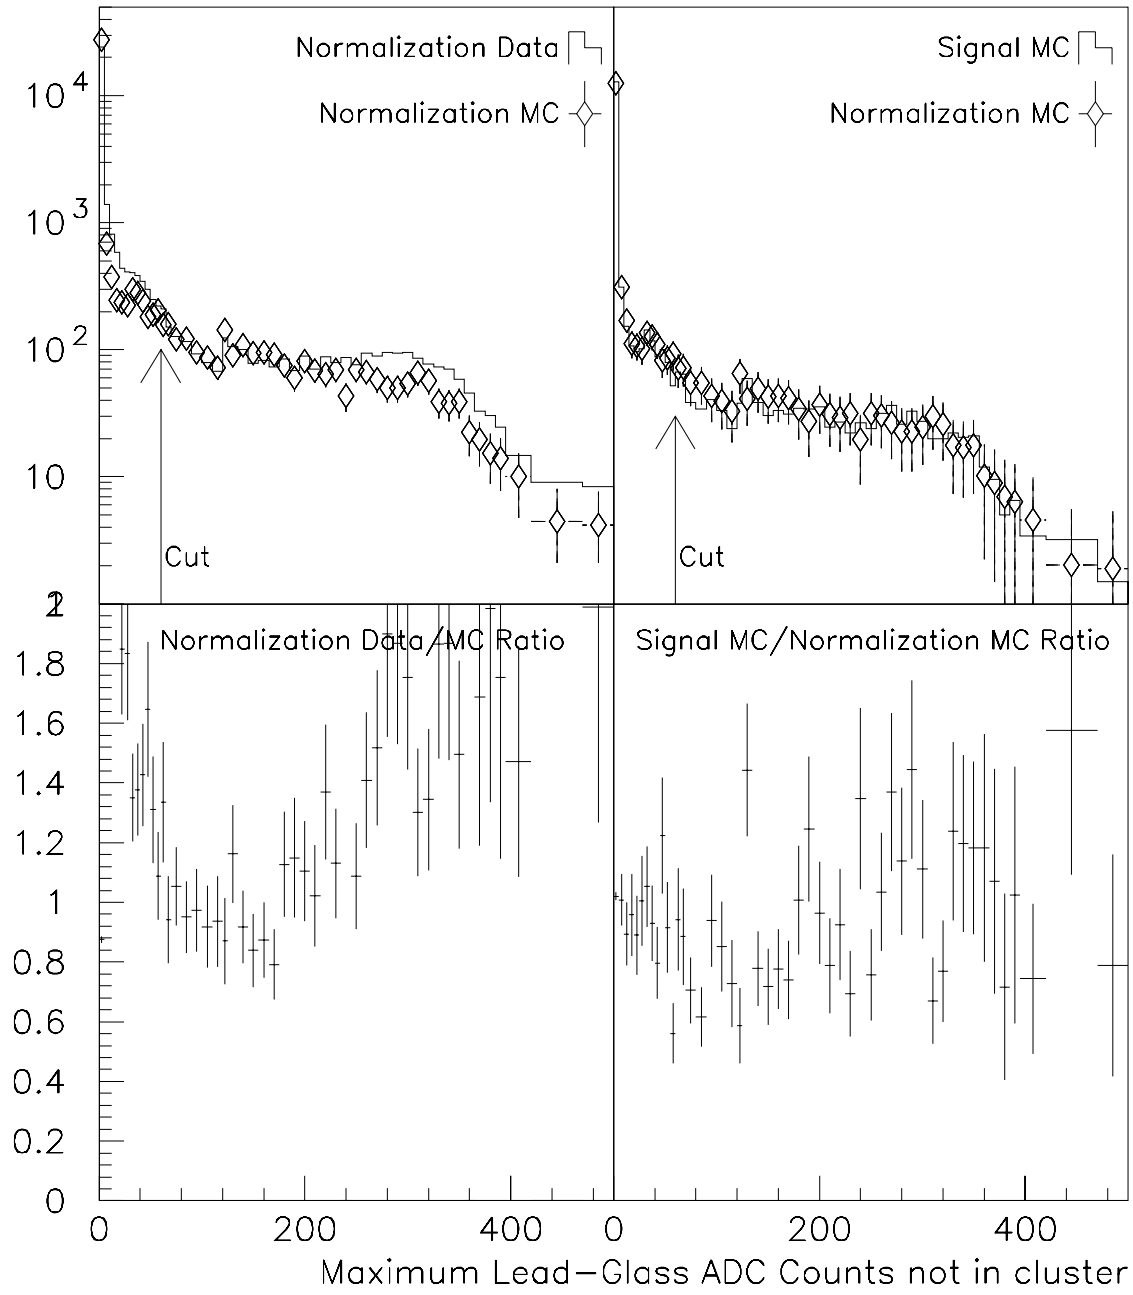

Figure 116. The maximum number of ADC counts in blocks containing candidate soft photons in data and Monte Carlo. The arrow indicates the location of the analysis cut.

glass energy in this adder was deposited in the same RF bucket as the trigger. The maximum number of ADC counts in any single lead-glass block not part of the 5×5 block array of a cluster and located in an adder satisfying this in-time condition was computed. If any such block contained more than 60 ADC counts, the event was rejected. This corresponded to an energy threshold of approximately 500 MeV for detection of these soft photons. The distribution of this maximum number of ADC counts is shown in Figure 116 for data and Monte Carlo. This cut resulted in a loss of acceptance of 16.4% in the normalization data sample, and a loss of 13.5% in the normalization Monte Carlo.

8.2.3 Photon Veto Cuts

To further reduce backgrounds from decays with extra final state photons, events with significant energy deposited into the lead-lucite (lead-scintillator in the case of the DRA) portion of the photon veto counters were removed from the event sample. The cuts were chosen by looking at events from the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, sample where there should be no signal in the photon vetoes from child photons in the reconstructed K_L decay chain. Cuts were then chosen that resulted in a minimal loss in sensitivity from the accidental activity in the photon vetoes. The photon veto response in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays and the chosen cuts are shown in Figure117. The cuts on each veto were between 0.8 to 2.5 minimum ionizing equivalents, which corresponded to a range of 0.11 to 0.37 GeV for photons. The cuts on all nine of these photon vetoes resulted in a 4.7% loss in acceptance in the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, normalization sample.

The Back Anti photon veto was treated separately because it acted as the beam dump for hadrons. Recall that in the trigger, the energy from the first two longitudinal layers of the Back Anti were added together to form an "electromagnetic" part, and the last third of the Back Anti (18.7 radiation lengths deep) was used to detect hadronic showers. If there was significant energy in this last layer, then the information in the first two layers was not used in the trigger. The analysis did not cut on the energy in the first two layers of the Back Anti when the trigger latch for the third layer was on. As in the trigger, only the $\hat{x}$-view layers of the Back Anti were used to form the energy sum on which the analysis cut. This $\hat{x}$-view layer energy sum is shown for normalization data and Monte Carlo in

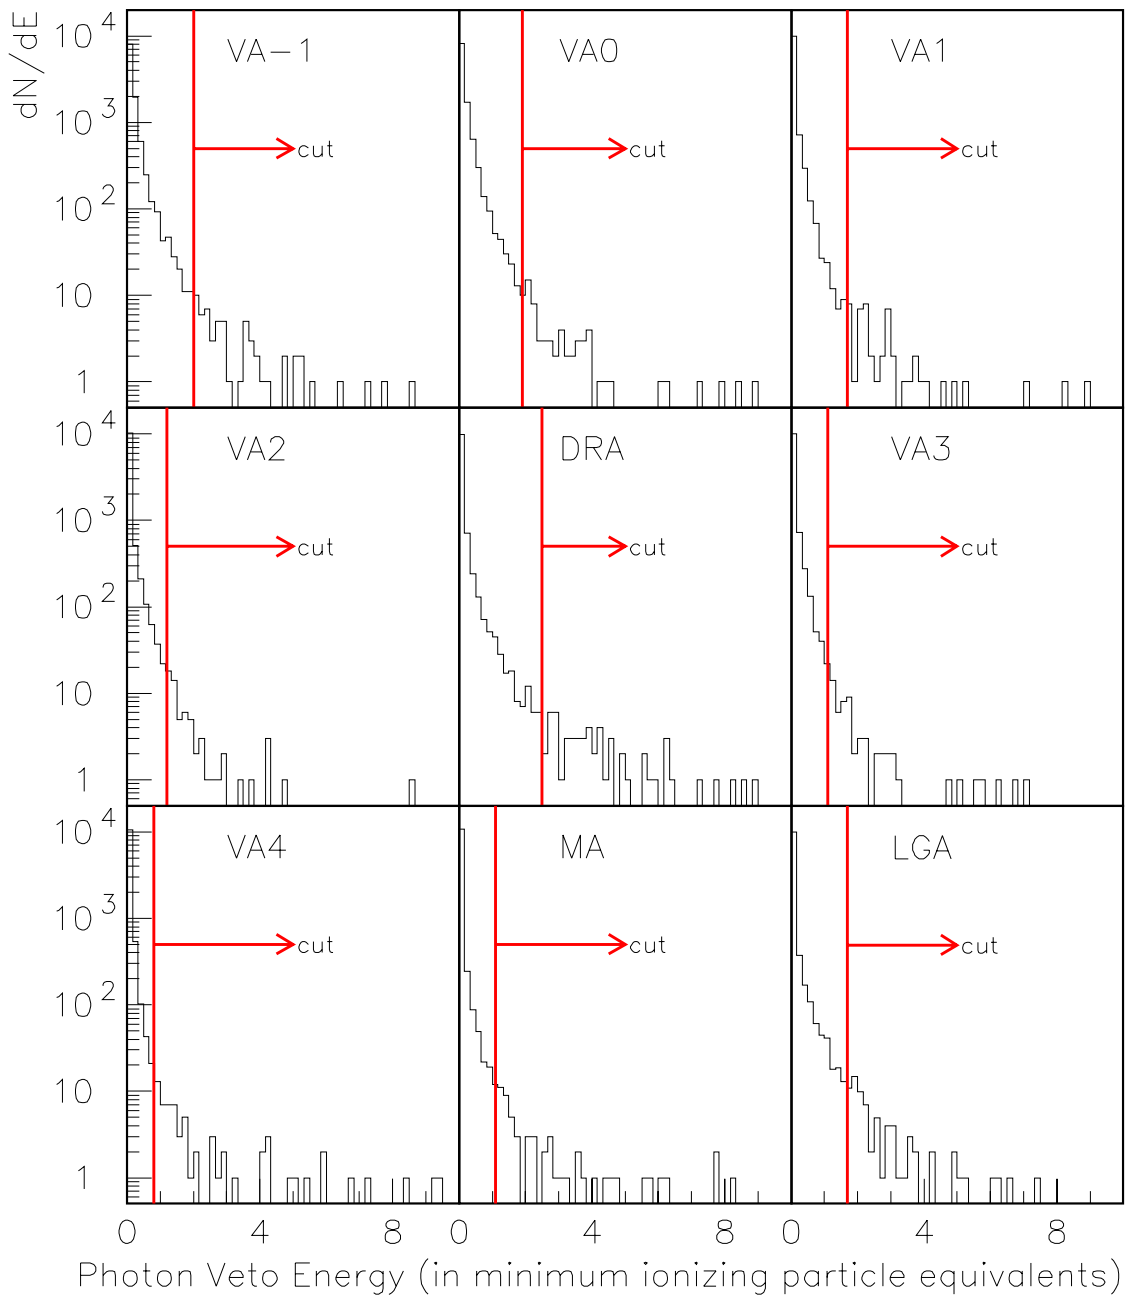

Figure 117. The photon veto response in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays, and the analysis cuts used to remove events with extra photons

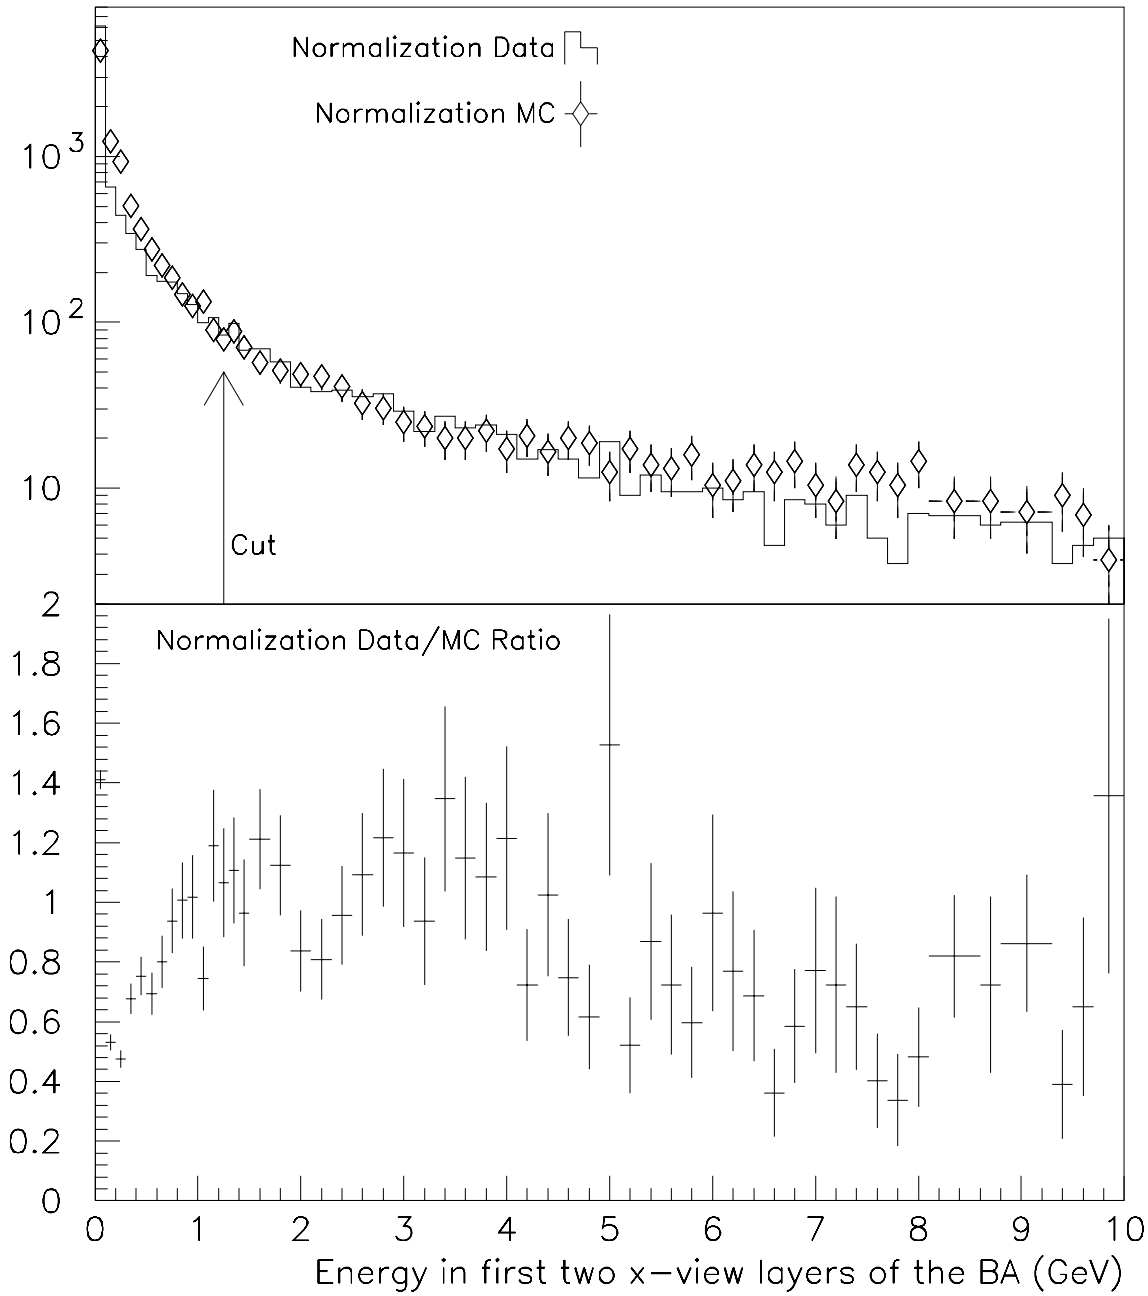

Figure 118. Energy deposited in the $\hat{x}$-views of the first two longitudinal layers of the Back Anti. The arrow indicates the location of the analysis cut.

Figure 118. The BA12 trigger latch was set when this energy was was above a threshold of about 2.2 GeV; however, this threshold was smeared with a Gaussian width of 0.4 GeV. The analysis cut was placed at 1.25 GeV, 2.5σ away from the threshold. Although the data and Monte Carlo did not agree at very low energies, the agreement was good in the region above the cut. This cut resulted in a 17% loss for normalization data and a 19% loss in normalization Monte Carlo.

8.2.4 Kinematic Cuts

The final cuts to be applied in order to extract the $\pi^0 \to e^+e^-$ and $\pi^0 \to e^+e^-\gamma$ signals were the cuts on kinematic quantities: the π^0 pairing χ^2, the transverse momentum of the reconstructed kaon, the total invariant mass, and the e^+e^- invariant mass in the case of $\pi^0 \to e^+e^-$. These cuts were by necessity different in the normalization and signal mode analyses.

The π^0 pairing χ^2 was defined in Section 8.2.2. This kinematic variable was important because it distinguished between events whose parent K_L decay was reconstructed as a $K_L \to \pi^0\pi^0\pi^0$, and events from other sources. It was also effective at rejecting events with undetected photon fusions because such a fusion would affect the measurement of $m_{\gamma\gamma}$ of the parent π^0. The pairing χ^2 distribution is shown for the $\pi^0 \to e^+e^-\gamma$ normalization sample in data and Monte Carlo and for the $\pi^0 \to e^+e^-$ Monte Carlo in Figure 119. The agreement between data and Monte Carlo was less than perfect, with the pairing χ^2 more sharply peaked in data than in Monte Carlo. From this plot, it is evident that the Monte Carlo cannot be expected to correctly predict the loss from a cut in pairing χ^2 because of the 20% disagreement in the magnitude of the high χ^2 tail. A loose cut of $\chi^2 < 15$ was chosen for the normalization sample because backgrounds were not a problem. For the $\pi^0 \to e^+e^-$ analysis, a cut was placed at $\chi^2 < 7.5$, which resulted in a 2.6% loss in $\pi^0 \to e^+e^-$ Monte Carlo acceptance after all other cuts.

Another kinematic variable of interest was the reconstructed transverse momentum of the kaon, or $\vec{p_t}$. Because the target position was by definition the origin of the coordinate system, and because the decay vertex location was measured from the charged tracks, the direction of kaon momentum, $\hat{d_K}$, could be determined with a accuracy of a few tens of

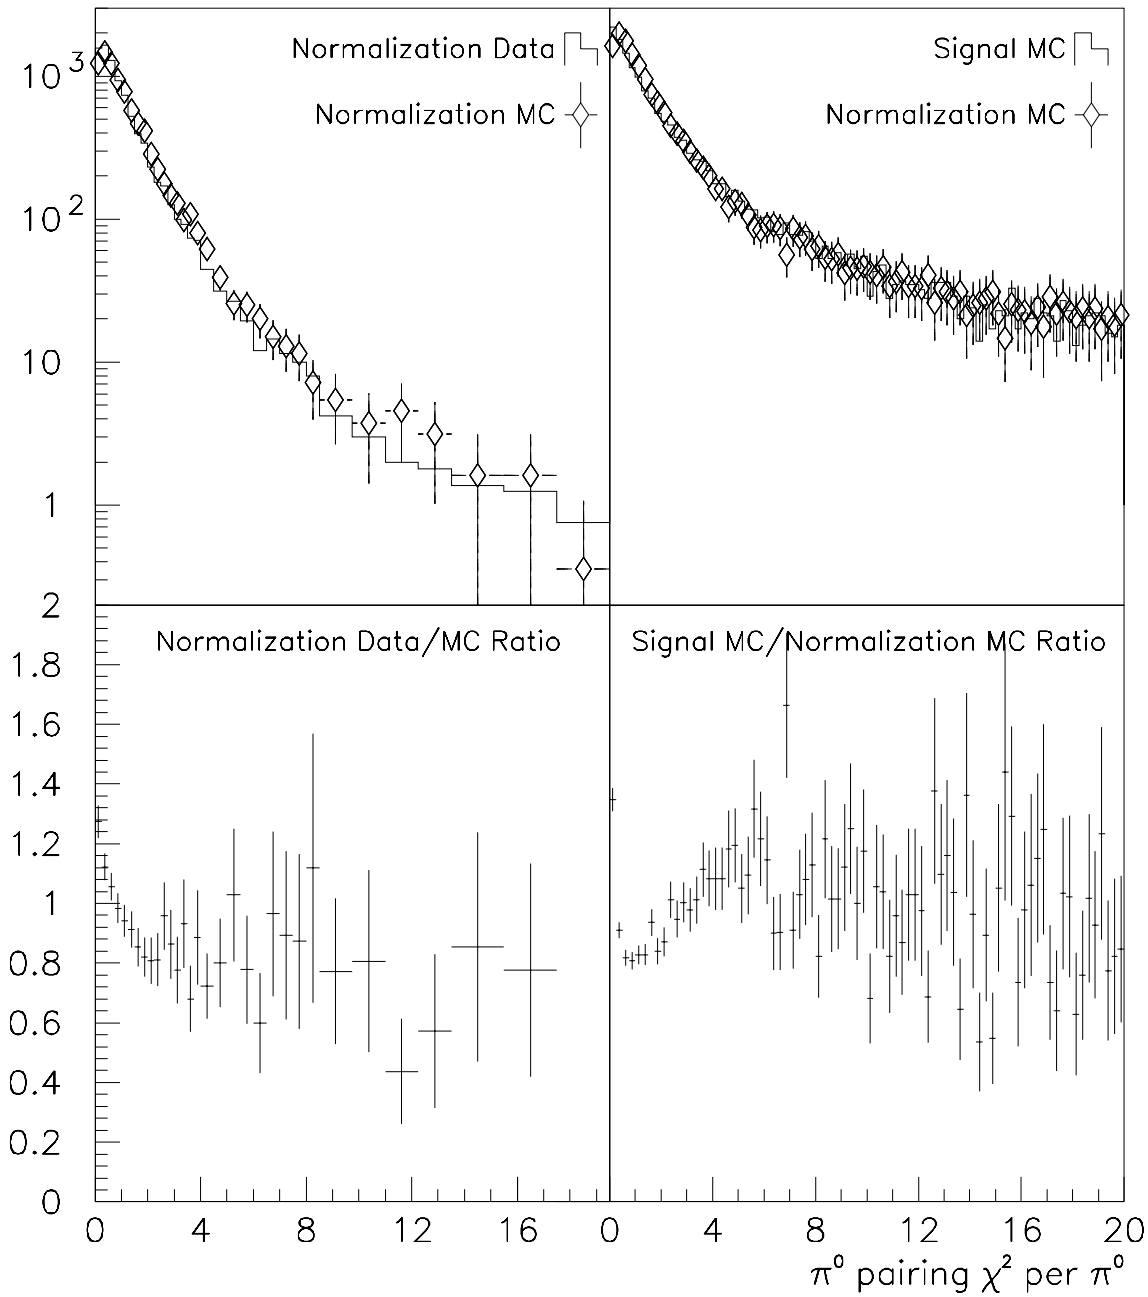

Figure 119. The π^0 pairing χ^2 in data and Monte Carlo. The comparison between normalization data and Monte Carlo uses tighter cuts than the comparison between Monte Carlo samples in order to remove background.

microradians. Also, the kaon momentum, $\vec{p_K}$, could be measured by adding up all the momentum vectors of the observed decay products. The transverse momentum, $\vec{p_t}$, was then defined as

$$\vec{p_t} \equiv \vec{p_K} - (\vec{p_K} \cdot \hat{d_K})\hat{d_K}.$$

The error in the measurement of $\vec{p_t}$ for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$ or $\pi^0 \to e^+e^-\gamma$, final states was typically dominated by the error in the measurement of $\vec{p_K}$ rather than error on $\hat{d_K}$.

The $\vec{p_t}$ was useful for rejecting backgrounds because it was sensitive to the loss of final-state particles. If, for example, a $\pi^0 \to e^+e^-\gamma$ decay occurred in the detector and the photon from the Dalitz decay were not observed, the observed $\vec{p_t}$ would be equal to the transverse momentum of the photon with respect to the K_L direction. On the other hand, for a $\pi^0 \to e^+e^-$ decay, if all of the final state particles were observed, one would expect $\vec{p_t} = 0$. Of course, complications such as electron bremsstrahlung, scattering, and detector resolution introduced errors into the measurement of $\vec{p_t}$.

The distribution of p_t^2 is shown for the $\pi^0 \to e^+e^-\gamma$ normalization sample in data and Monte Carlo in Figure 120. The Monte Carlo overpredicts the height of the large p_t^2 tail by approximately 25%. The mean p_t^2 in data was 280 MeV/c^2 and 330 MeV/c^2 in Monte Carlo. In the normalization sample, a loose cut in p_t^2 was placed at 2500 MeV^2, in order to minimize the effect of this disagreement. For the $\pi^0 \to e^+e^-$ analysis, the cut was placed at $p_t^2 < 1025$ MeV^2, which caused a 2.5% loss in $\pi^0 \to e^+e^-$ Monte Carlo after all other cuts.

The remaining kinematic variables were the invariant masses of the $\pi^0\pi^0\pi^0$ final states for both signal and normalization modes and the e^+e^- invariant mass in the $\pi^0 \to e^+e^-$ signal mode. The total seven-body invariant mass and the e^+e^- invariant mass in data and Monte Carlo for the $\pi^0 \to e^+e^-\gamma$ normalization sample are shown in Figures 121 and 122. As in the case of p_t^2 and pairing χ^2, the Monte Carlo overpredicts the size of the tails in the invariant mass distribution. The RMS of the total invariant mass distribution was 15.8 MeV/c^2 in data, and 17.7 MeV/c^2 in Monte Carlo. To minimize the effect of this disagreement, a very loose cut was placed on the total invariant mass in the normalization sample,

$$|m_{\text{total}} - m_{K_L}| < 70 \text{ MeV/c}^2.$$

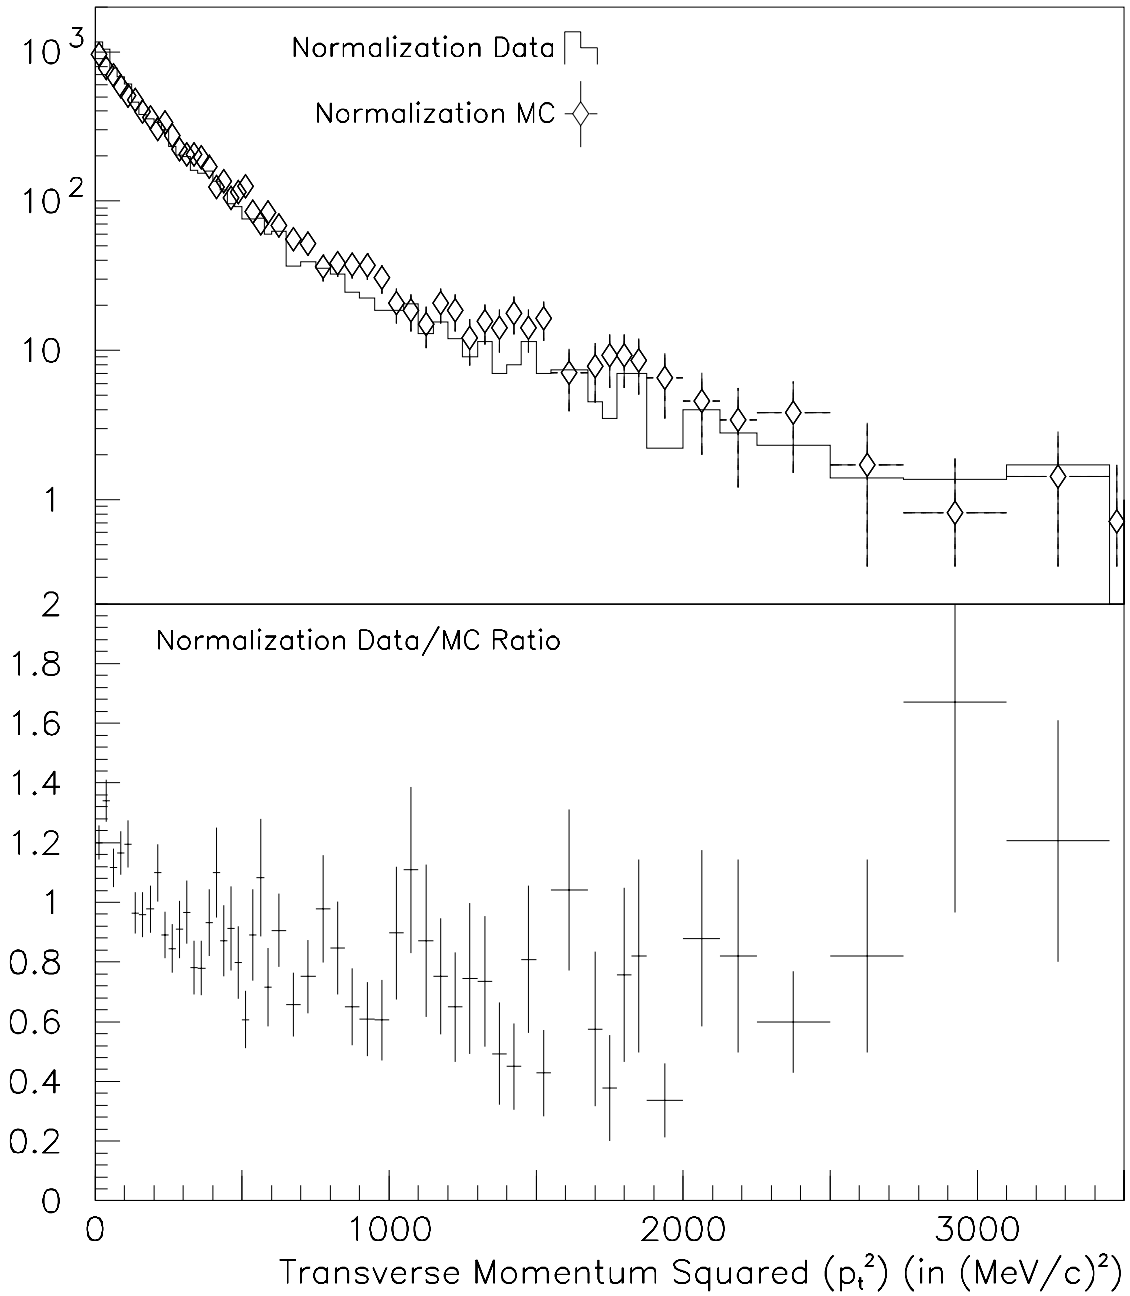

Figure 120. The transverse momentum with respect to the kaon direction in data and Monte Carlo for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays

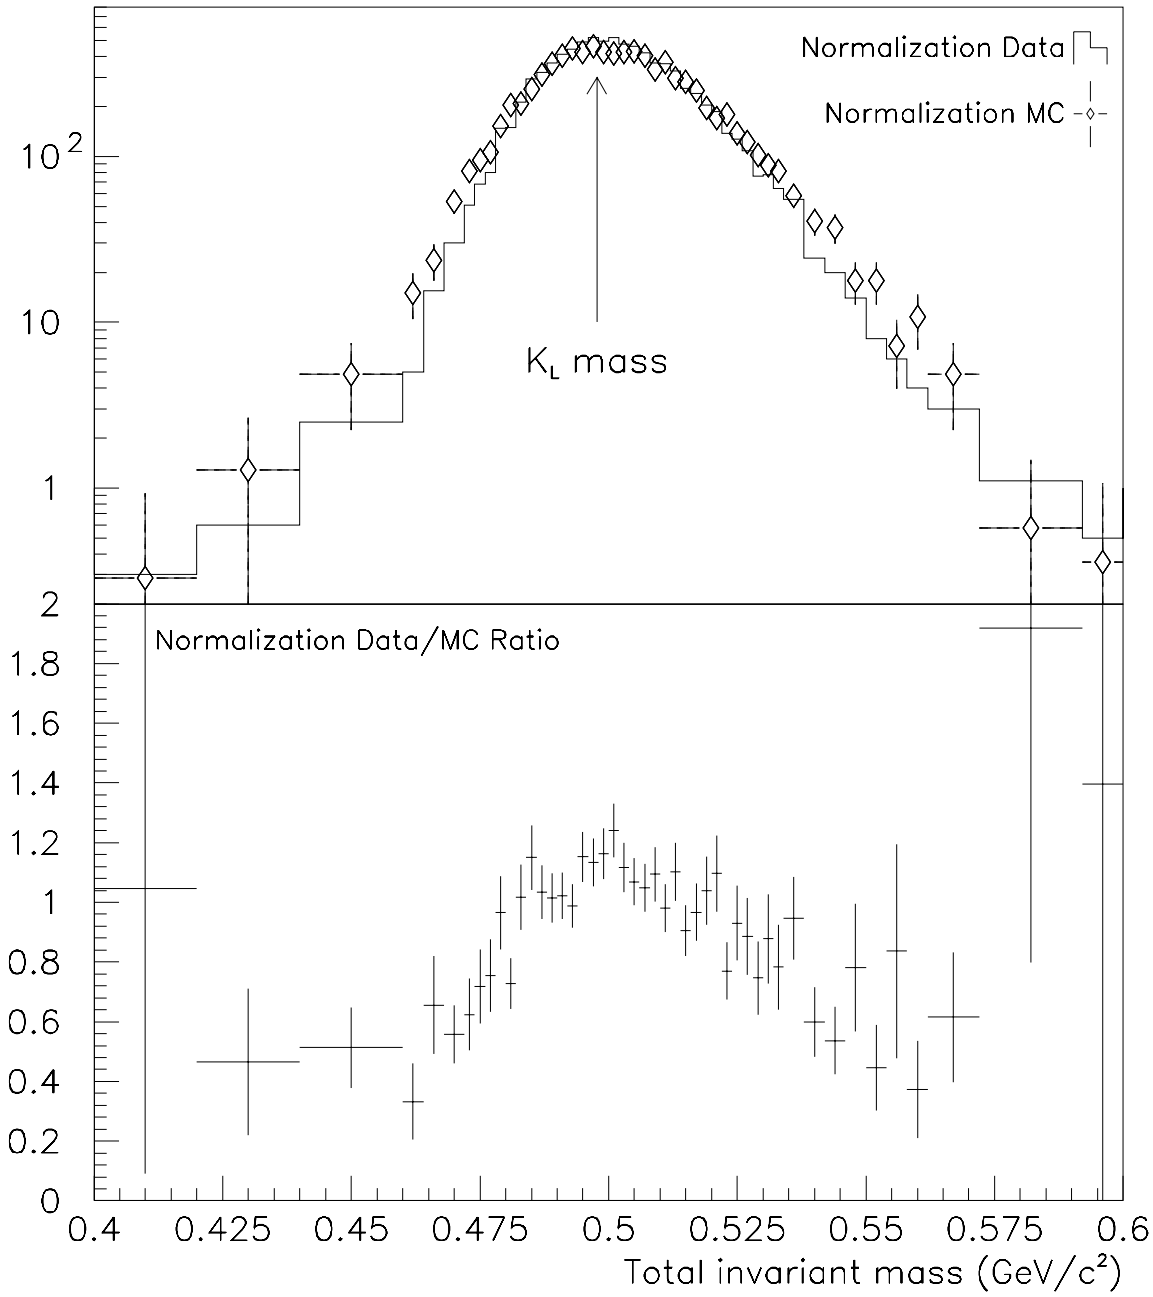

Figure 121. The total invariant mass in data and Monte Carlo for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays

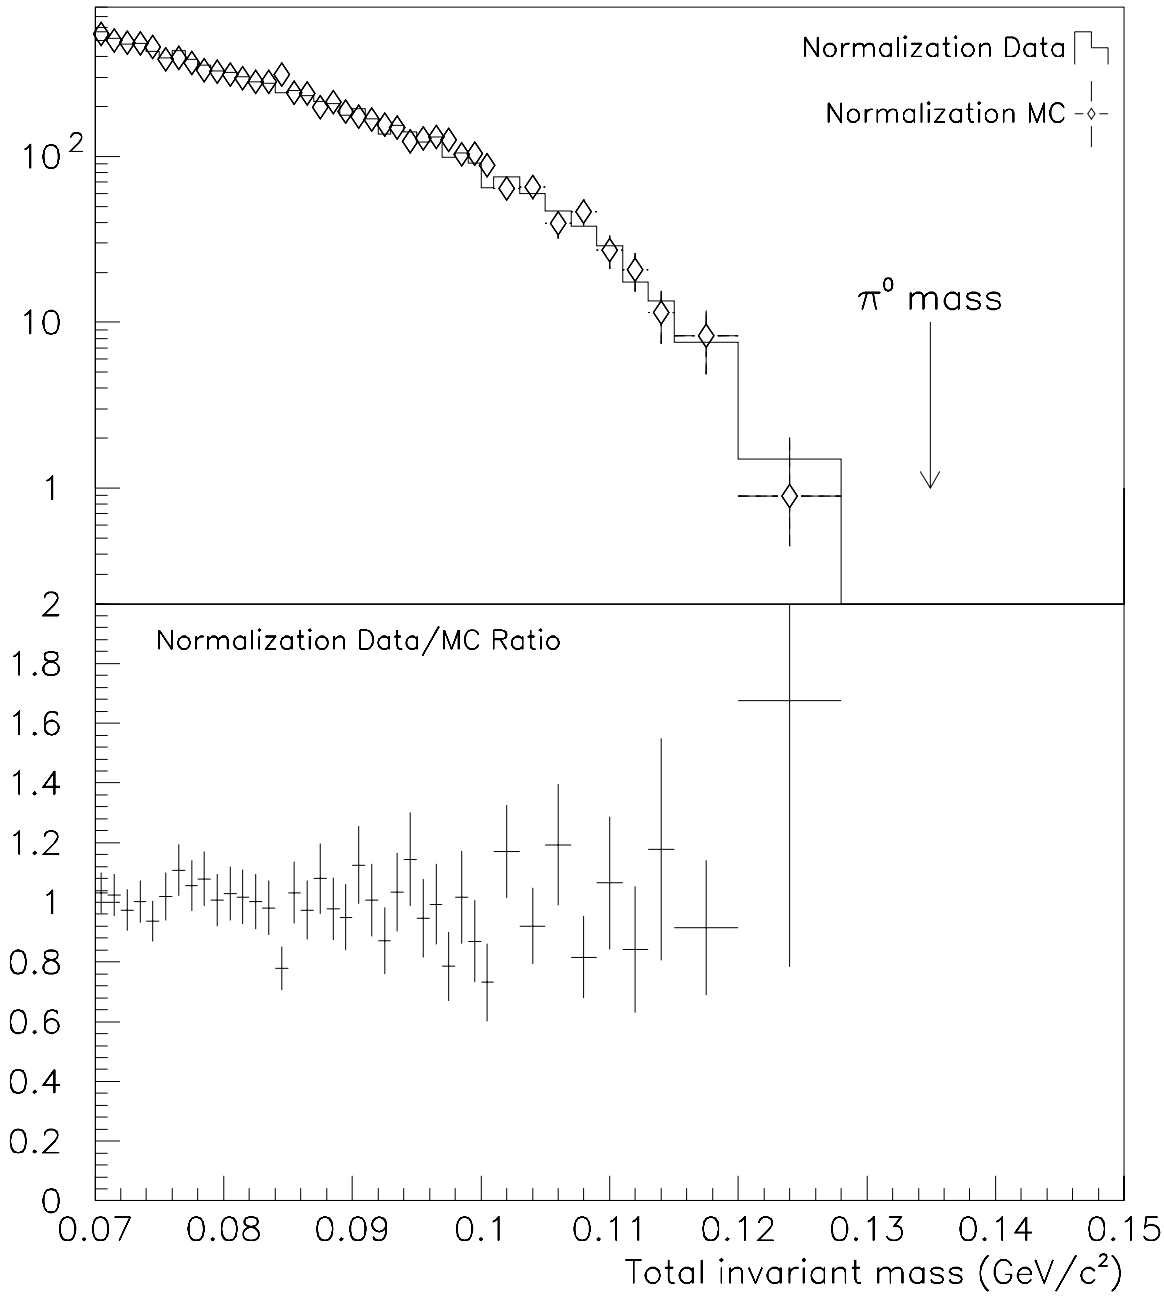

Figure 122. The e^+e^- invariant mass in data and Monte Carlo for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays

The agreement in the e^+e^- invariant mass was very good however, and gave us confidence in the calibration and understanding of the spectrometer.

In the $\pi^0 \to e^+e^-$ mode, cuts in the total and e^+e^- invariant mass were chosen using the Monte Carlo simulation. The Monte Carlo distribution for the total invariant mass and the e^+e^- invariant mass are shown in Figure 123. The cut on the total invariant mass was again very loose,

$$|m_{\text{total}} - m_{K_L}| < 41 \text{ MeV/c}^2,$$

to avoid a problem from the overestimation of the glass resolution tails in the Monte Carlo. This cut corresponds to more than 3 times the fitted Gaussian width of 13.5 MeV/c^2, and resulted in an acceptance loss of only 0.8%. In the e^+e^- invariant mass, though, where the $\pi^0 \to e^+e^-\gamma$ sample shown here and the $K_L \to \pi^+\pi^-$ study discussed in Section 5.6.1 serve as a cross check, this cut could be made tighter. Also, because the chamber resolution was significantly better than that of the glass, the e^+e^- mass cut was the best kinematic quantity for discriminating between $\pi^0 \to e^+e^-$ and $\pi^0 \to e^+e^-\gamma$ decays. Therefore, it was important and acceptable to cut tightly on the e^+e^- invariant mass. Studies of $\pi^0 \to e^+e^-$ decays in Monte Carlo without radiative corrections measured the e^+e^- invariant mass resolution at the π^0 mass to be 1.8 MeV/c^2 (fitted Gaussian width). Therefore, the cut was chosen to be 2.5 times this width, or 4.5 MeV/c^2.

Because the Monte Carlo did not accurately predict losses from kinematic cuts involving measurements from the lead-glass, there was some systematic error associated with these cuts. However, such errors will be relatively small because a total of only 5% of the $\pi^0 \to e^+e^-$ Monte Carlo events were removed by the cuts on pairing χ^2, p_t^2 and total invariant mass. The systematic errors from misunderstanding the lead-glass resolution are discussed and evaluated in Section 8.6.

The cuts discussed in this section represented a best attempt to eliminate all expected sources of background. With these cuts, the normalization mode had no apparent background, so all that remained was to check the $\pi^0 \to e^+e^-$ sample and see what background and signal, if any, were present.

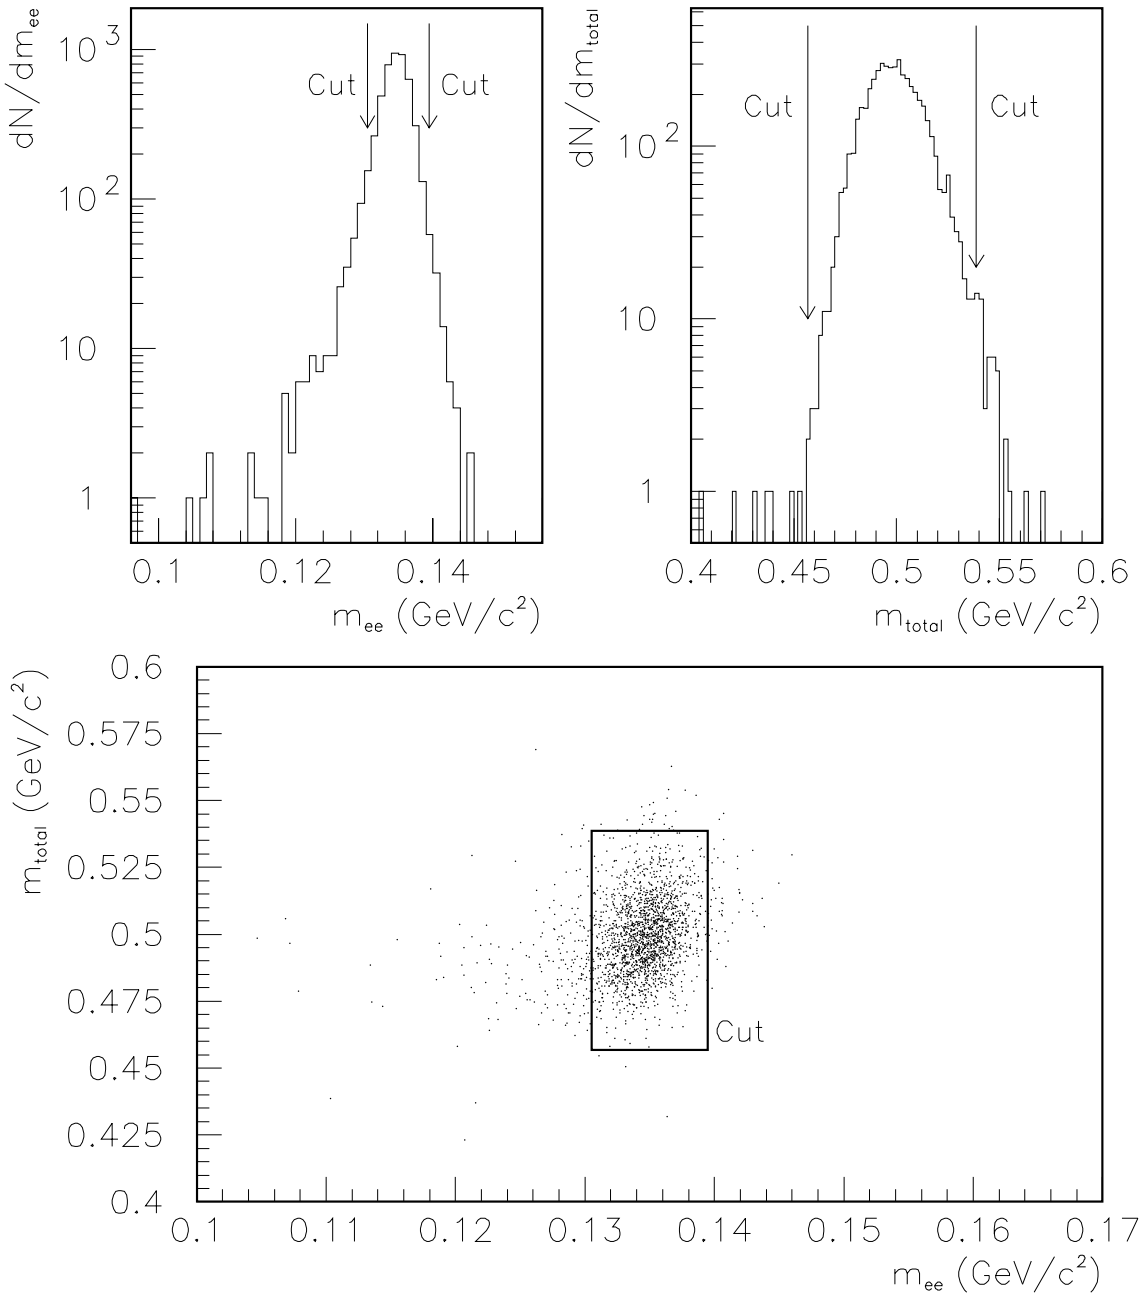

Figure 123. The total invariant mass and e^+e^- invariant mass in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, Monte Carlo. The experimental cuts on these quantities are also shown. Note that much of the radiative tail from $\pi^0 \to e^+e^- + \gamma_{\rm brem}$ was removed by the m_{ee} cut.

8.3 Backgrounds

The title of this section, of course, gives away the answer to the nagging question of what was in the $\pi^0 \to e^+e^-$ data sample. With all the cuts described in the previous sample, the $\pi^0 \to e^+e^-$ data sample is shown in Figure 124. There was no obvious evidence in this sample for an excess of events in the signal region corresponding to a $\pi^0 \to e^+e^-$ contribution; however, there were 50 events in the signal region.

What were these events? A few clues can be found in the event sample itself. First, the total invariant mass of most of these events was near to, but slightly below, the K_L mass. The mean total invariant mass of the events in Figure 124 was 480 MeV/c^2, and 85% were below the K_L mass. This suggests that most of these events came from parent K_L decays with missing particles. The cuts on p_t^2, the photon vetoes, cluster fusions and extra soft photons act to make it more likely that any missing particles would be low energy, and thus not carry away much of the invariant mass of the event. Therefore, the events that remain after all cuts from missing particle backgrounds will be likely to have a total invariant mass near to but below the K_L mass. Second, there appeared to be two components to the background. The first, dominant component had an e^+e^- invariant mass spectrum which fell as it approached the π^0 mass at 135.0 MeV/c^2. A second component, which was only evident above the π^0 mass where the first component was absent, appeared to be roughly flat in $m_{e^+e^-}$.

The dominant structure in e^+e^- invariant mass came from an obvious source, $\pi^0 \to e^+e^- + X$ decays, where X was not observed in the detector. The most obvious such decay was the $\pi^0 \to e^+e^-\gamma$ Dalitz decay, which occurred with a large branching fraction but also with a steeply falling e^+e^- invariant mass spectrum. However, as is shown in Figure 125, $\pi^0 \to e^+e^-\gamma$ only explained a small fraction of this background because the rate for e^+e^- invariant mass pairs near the π^0 mass was so low in Dalitz decay. Therefore, some other π^0 decay must have been responsible for this background.

What other processes can give e^+e^- pairs? The known decays of the π^0 give e^+e^- pairs through Dalitz decays, the internal conversion of a photon. In such decays, the probability of a conversion is approximately 6.0×10^{-3} per photon. There were also external conversions

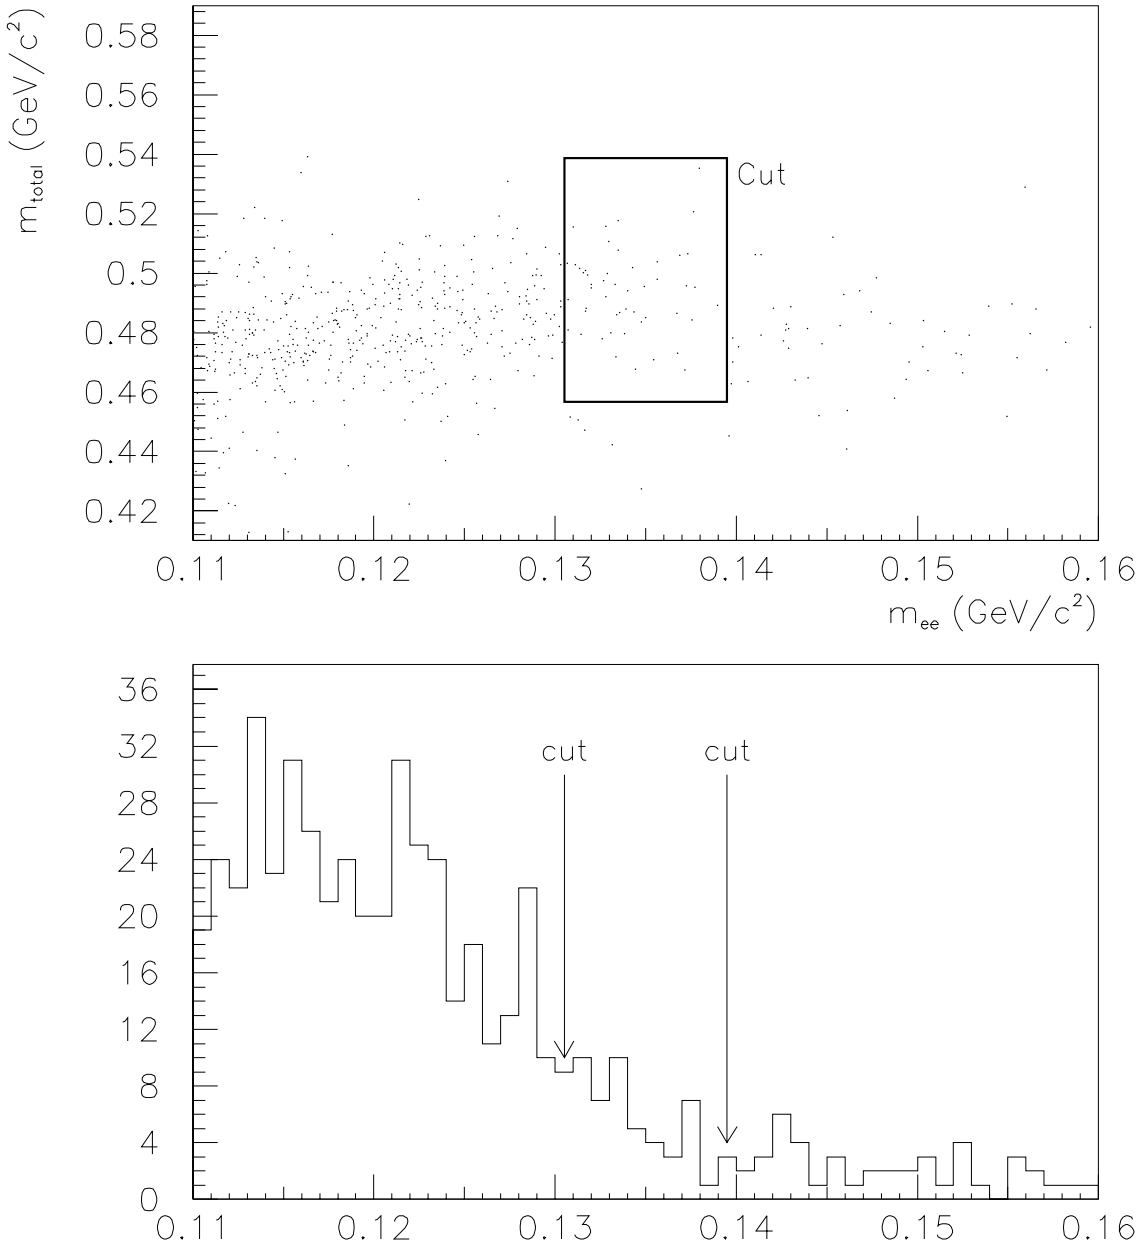

Figure 124. The total invariant mass vs. the e^+e^- invariant mass in the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, data sample. The projection onto the $m_{e^+e^-}$ axis after a cut on the total mass is also shown.

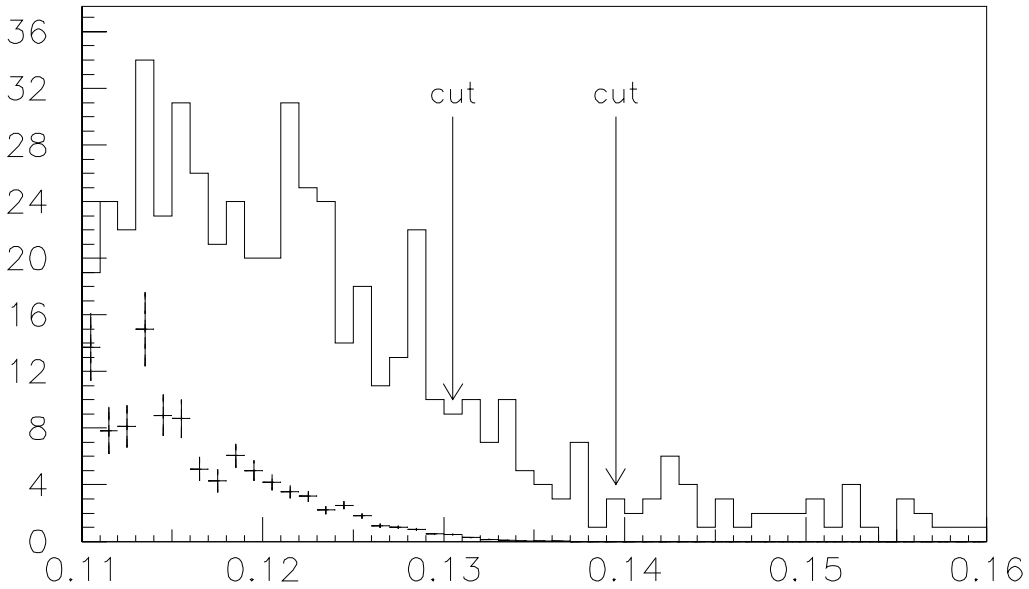

Figure 125. The the e^+e^- invariant mass in the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, data sample and $\pi^0 \to e^+e^-\gamma$ background Monte Carlo. The histogram is the data, and the crosses are Monte Carlo. The background Monte Carlo is absolutely normalized to the observed number of fully reconstructed $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays.

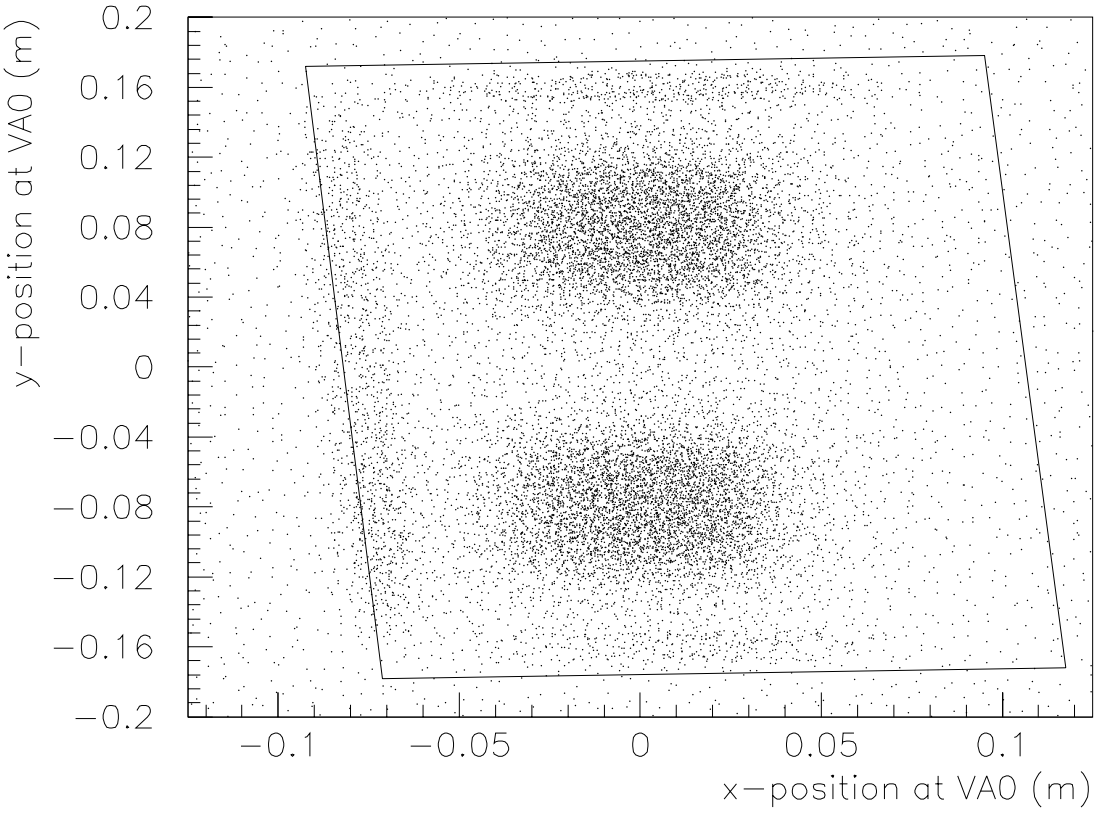

Figure 126. Position of electrons at VA0 in $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, candidates. The events in the two bright spots are from internal photon conversions in decays in the beam. The diffuse background is from decays which occurred far downstream of VA0. The bands and the left, top and bottom sides of the beam result from conversions in the VA0 edge. The position of the VA0 scintillator edge is also shown by the lines.

of final-state photons in the material in the detector, which could also contribute to the rate of observed e^+e^- pairs if the conversion occurred far enough upstream in the detector to be observed in Chamber 1. The material in the detector can be divided into three parts. First, the residual gas in the vacuum tank, at a pressure of less than 15 mTorr, accounted for less than 7×10^{-8} radiation lengths per meter, and was therefore negligible. Second, the vacuum window, air gap, and upstream wires and gas in the first drift chamber, provided a large amount of material in which photons could convert. In order to create a reconstructible track, the photon had to convert before leaving the drift cells of the downstream $\hat{y}$-view plane of Chamber 1. From Table 6, this material comprised $(3.4 \pm 0.2) \times 10^{-3}$ radiation lengths, where the error quoted was primarily from the estimation of the air gap thickness. This gives a $(0.26 \pm 0.02)\%$ conversion probability per photon. A third source of material were the edges of the photon vetoes and vacuum tank walls. Ideally, electrons from photon conversions in this material should have been vetoed in downstream photon vetoes, or should have failed track aperture cuts because they did not originate from the beam. However, in the case of conversions in the upstream photon vetoes, in VA0 in particular, this was not always the case. Shown in Figure 126 is the position at VA0 of electrons from a sample of $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, candidates with very loose cuts. The vertices of such conversions were well-separated from the vertices of internal photon conversions in decays from the beam, and therefore such events were easy to remove from the event sample.

With both internal and external conversions, there were a number of π^0 decays that could contribute to backgrounds. For example, an $e^+e^-\gamma$ final state could be created by either

$$\pi^0 \to e^+e^-\gamma \;\text{ or }\; \pi^0 \to \gamma + (\gamma \to e^+e^-)$$

where $\gamma \to e^+e^-$ is understood to denote external photon conversion. However, this does not help to simulate the background in Figure 125 since the e^+e^- spectrum of internal conversion pairs was even softer than that of Dalitz pairs. However, in a *four* electron final state, such as

$$\pi^0 \to e^+e^-e^+e^-,\; \pi^0 \to e^+e^- + (\gamma \to e^+e^-),\text{ or } \pi^0 \to (\gamma \to e^+e^-) + (\gamma \to e^+e^-),$$

it was more probable for *some* e^+e^- pair to have a high invariant mass. In such a four-

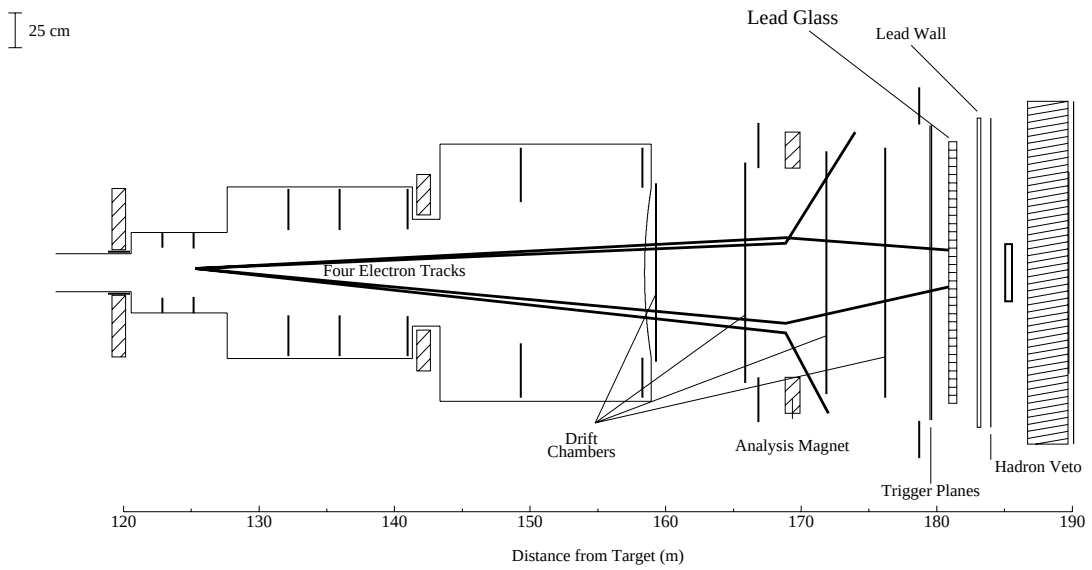

Figure 127. Misreconstruction of a four-electron final state as a high e^+e^- invariant mass two-electron final state. One electron from each pair is bent out of the detector by the magnet, and the remaining e^+e^- pair have a large invariant mass.

electron decay each pair of electrons created by a photon conversion was also likely to have a low invariant mass. However, electron-positron pairs from different parent photons were likely have a fairly large invariant mass. The typical topology of such a decay is shown in Figure 127. Such decays in the detector resulted in a potentially high e^+e^- invariant mass, up to the π^0 mass, and a total invariant mass somewhat below the K_L mass.

A similar decay could account also for the background with an e^+e^- invariant mass above the π^0 mass. Imagine a four-electron final state, but with the two converting photons coming from *different* parent π^0's. In such a decay, it was possible for the reconstructed e^+e^- invariant mass to be above the π^0 mass since the e^+ and the e^- did not originate from the same π^0. Many such four-electron final states could result from $K_L \to \pi^0\pi^0\pi^0$ decays such as

$$\begin{array}{cccc} K_L \to & \pi^0 & \pi^0 & \pi^0 \\ & \downarrow & \downarrow & \downarrow \\ & \gamma\gamma & e^+e^-\gamma & e^+e^-\gamma, \end{array}$$

$$\begin{array}{cccc} K_L \to & \pi^0 & \pi^0 & \pi^0 \\ & \downarrow & \downarrow & \downarrow \\ & \gamma\gamma & \gamma\gamma & e^+e^-\gamma \\ & & \downarrow & \\ & & e^+e^- & \end{array}$$

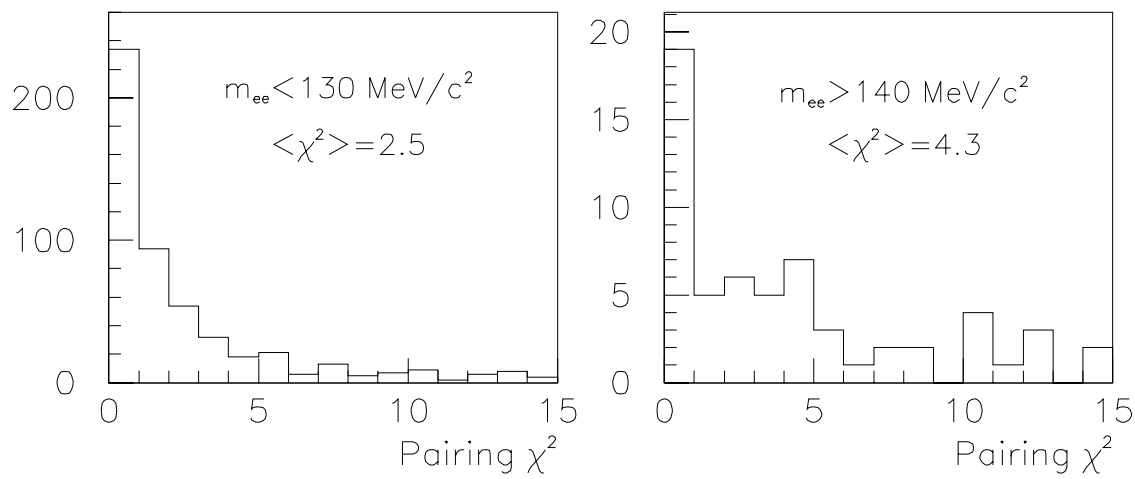

Figure 128. The π^0 pairing χ^2 for the low e^+e^- mass and high e^+e^- background compared

and

$$\begin{array}{cccc} K_L \rightarrow & \pi^0 & \pi^0 & \pi^0 \\ & \downarrow & \downarrow & \downarrow \\ & \gamma\gamma & \gamma\gamma & \gamma\gamma \\ & & \downarrow & \downarrow \\ & & e^+e^- & e^+e^-. \end{array}$$

There was one additional kinematic handle on the rejection of such decays. Because they all required that two photons from different π^0's "fake" a π^0 mass, this meant that the pairing χ^2 cut could be used to discriminate against such decays. A comparison of the π^0 pairing χ^2 for low e^+e^- mass and high e^+e^- mass events in the $\pi^0 \rightarrow e^+e^-$ sample is shown in Figure 128.

To verify that these four electron decays were responsible for this background, Monte Carlo simulations of these backgrounds were generated. The range of e^+e^- background populated by these backgrounds is shown in Figure 129. A comparison between the e^+e^- spectrum observed in data, and that observed in a Monte Carlo including the four electron backgrounds is shown in Figure 130. Clearly, these backgrounds were responsible for the observed events kinematically similar to $\pi^0 \rightarrow e^+e^-$ in the data sample.

8.4 Removal of the Four Electron Background

Because this background was difficult to remove kinematically, another technique had to be found for eliminating background from four electron decays. The method chosen was to search for the extra two electrons in the event. Since each of the electrons in the e^+e^- pair reconstructed in the detector had a sibling electron from the parent $\gamma \rightarrow e^+e^-$ conversion, typically at low mass, these electrons would usually leave hits in the drift chambers upstream of the magnet. (Downstream of the magnet these electrons, which were often soft, could be bent out of the spectrometer by the magnet's momentum kick.) One difficulty with this approach was that the e^+e^- conversion pairs were often very close together in the chambers and could be indistinguishable from a single track.

The cuts chosen to remove these events searched for extra hits in both views of the two upstream drift chambers. The chambers were searched for hits not from electron tracks satisfying one of three categories: (1) in-time hit pairs with a sum-of-distances within 1 mm

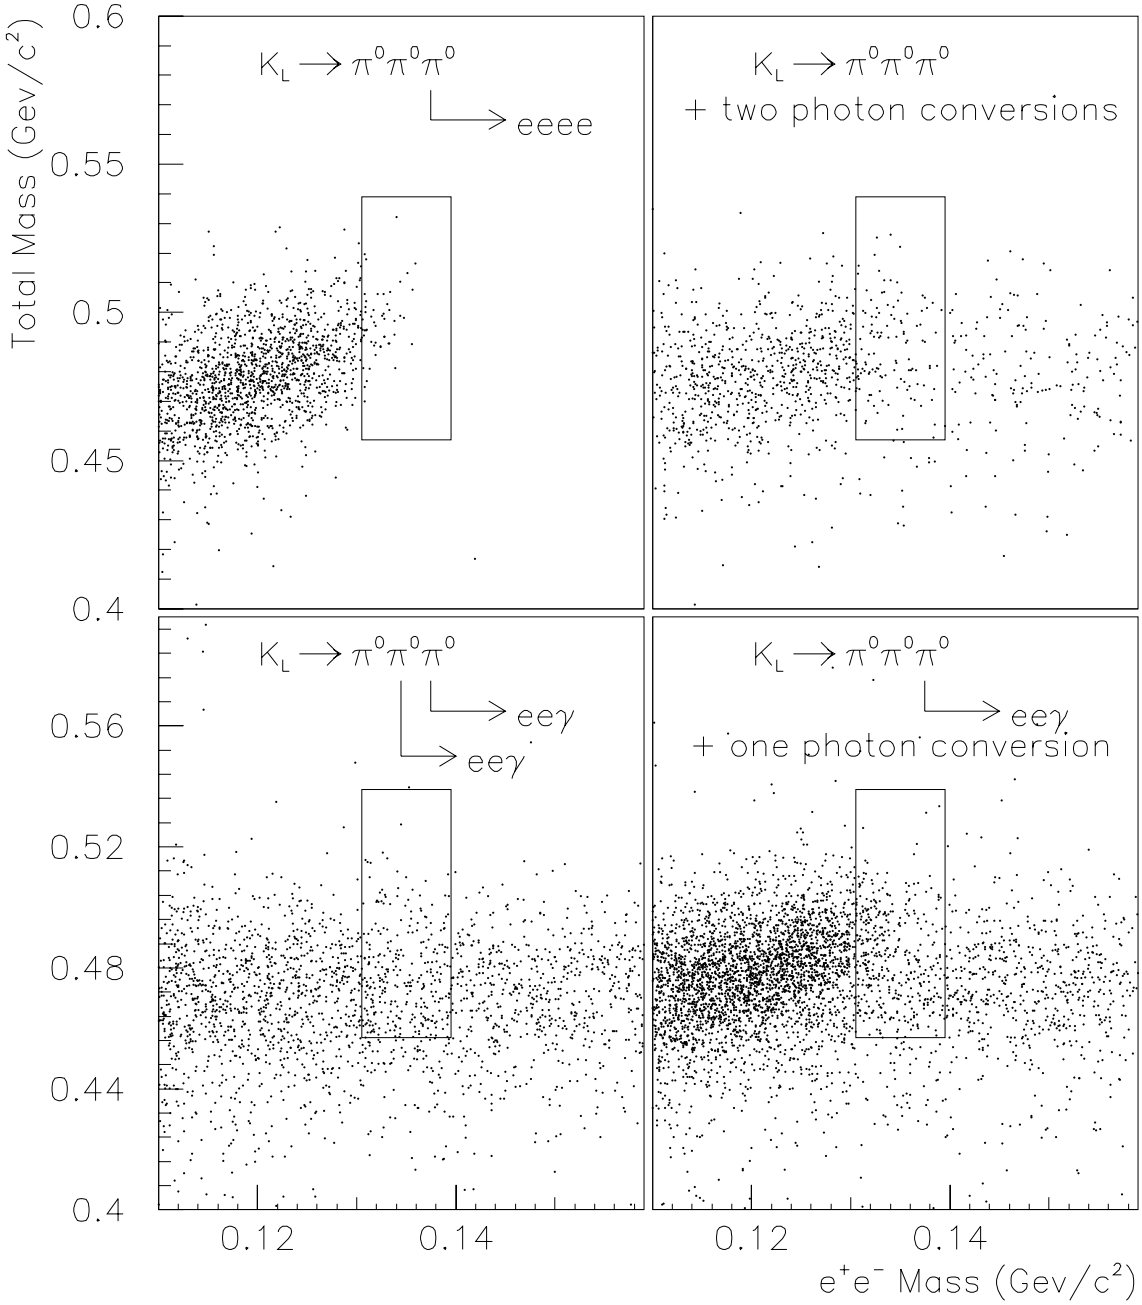

Figure 129. The total invariant mass versus the e^+e^- invariant mass for four-electron backgrounds in an early version of the Monte Carlo simulation. Each simulation represents approximately 19 times the flux of such decays observed in the data.

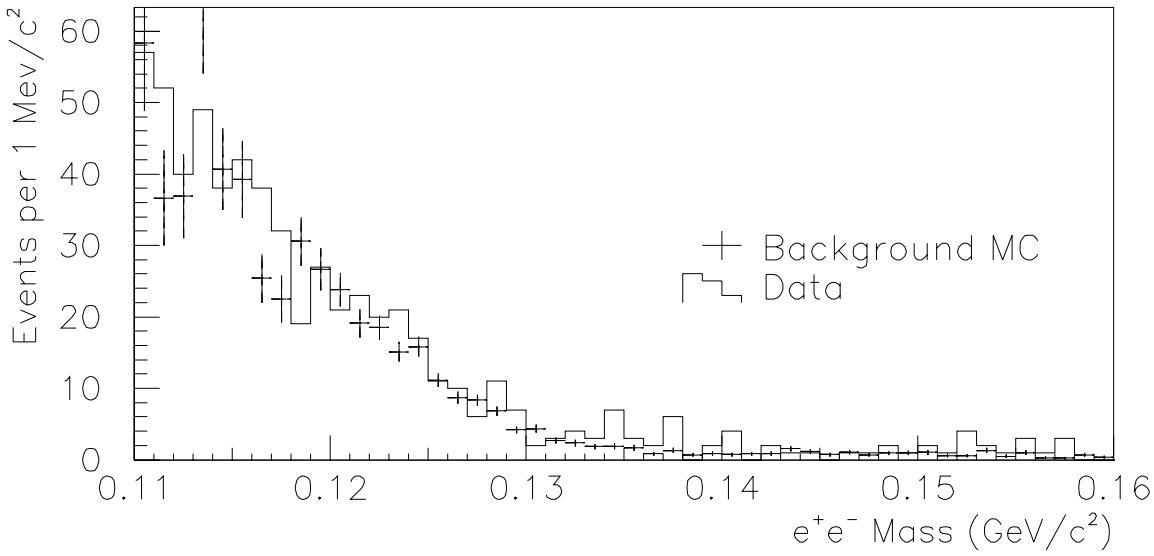

Figure 130. The the e^+e^- invariant mass in data and a background Monte Carlo including $\pi^0 \rightarrow e^+e^-\gamma$ and four electron backgrounds. Because this sample of four electron backgrounds was generated with a weak version of the Chamber 2 extra hit cut discussed in Section 8.4, this cut was also applied to the data and the $\pi^0 \rightarrow e^+e^-\gamma$ background simulation. Nevertheless, this cut was sufficiently loose to still allow for a reasonable comparison, and the four electron processes appear to be sufficient to simulate the background which was not accounted for in Figure 125.

of the cell size, (2) isolated hits adjacent to hits on electron tracks, and (3) isolated hits or out-of-time hit pairs. Note that the second category also selected events with a hit pair from an electron track that had a sum-of-distances not within 1.5 mm of the nominal cell size because only one of the hits was used in the track in such cases. Therefore, removing events with hits in the second category effectively required all hit pairs along the electron tracks to be in-time. This search for extra chamber hits provided a redundant cut criterion since each event had two extra electrons which passed through two views of two chambers. This made the cuts very effective independent of small wire inefficiencies and made it much more likely that the small opening angle between two low invariant mass tracks could be detected.

The exact cuts were chosen using a Monte Carlo simulation of the four electron backgrounds. As shown in Table 26, the most effective cuts were those made on extra hits in Chamber 2. This was not a surprising result, since close together tracks, especially those from photon conversions in the vacuum window, had more chance to separate by the time they reached Chamber 2. In no case was adding a cut on isolated or out-of-time hits shown to be effective. Therefore, the cut to remove the four-electron background discarded any event with either an extra in-time hit pair, or an extra hit adjacent to an electron track in either view of Chamber 2. Such a cut was very effective, removing between 99.4% and 99.9% of four-electron background events, depending on the decay mode.

Because the cuts were chosen using the existing Monte Carlo sample, a completely new sample was generated in order not to bias the background determination. The result of this simulation after all cuts is shown in Figure 131. After the Chamber 2 extra hit cut, less than 1 event background was predicted for the expanded signal region of 100 MeV/c^2 $< m_{ee} <$ 170 MeV/c^2, and 400 MeV/c^2 $< m_{\rm total} <$ 600 MeV/c^2. This cut, however, did result in a 33% loss of acceptance in the normalization data sample. Furthermore, because this cut required that the sum-of-distances of a hit pair in Chamber 2 be within 1.5 mm of the cell size, the cut was affected by the presence of δ-rays in the chambers which were not simulated in the Monte Carlo. The loss from this cut in the $\pi^0 \to e^+e^-\gamma$ normalization Monte Carlo was only 16%. However, because both the $\pi^0 \to e^+e^-$ and $\pi^0 \to e^+e^-\gamma$ normalization decays have

Table 26. The fraction of four-electron background events passing cuts on extra hits in the upstream drift chambers. The fractions shown in boxes are for the cut used in the analysis.

Cut		$K_L \to \pi^0\pi^0\pi^0$ $\pi^0 \to e^+e^-\gamma$ $\pi^0 \to e^+e^-\gamma$	$K_L \to \pi^0\pi^0\pi^0$ $\pi^0 \to e^+e^-\gamma$ $\gamma \to e^+e^-$	$K_L \to \pi^0\pi^0\pi^0$ $\gamma \to e^+e^-$ $\gamma \to e^+e^-$	$K_L \to \pi^0\pi^0\pi^0$ $\pi^0 \to e^+e^-e^+e^-$
No extra in-time pairs in Ch. 1		1.5%	7.3%	89%	0.9%
No extra in-time pairs in Ch. 2		0.8%	1.8%	9.4%	0.1%
No extra in-time pairs in Ch. 2 plus...	No extra in-time pairs in Ch. 1	0.8%	1.0%	8.7%	0.1%
	No extra in-time pairs or adjacent hits in Ch. 1	0.6%	0.6%	7.9%	0.1%
	No extra adjacent hits in Ch. 2	$\boxed{0.6\%}$	$\boxed{0.1\%}$	$\boxed{0.2\%}$	$\boxed{0.1\%}$

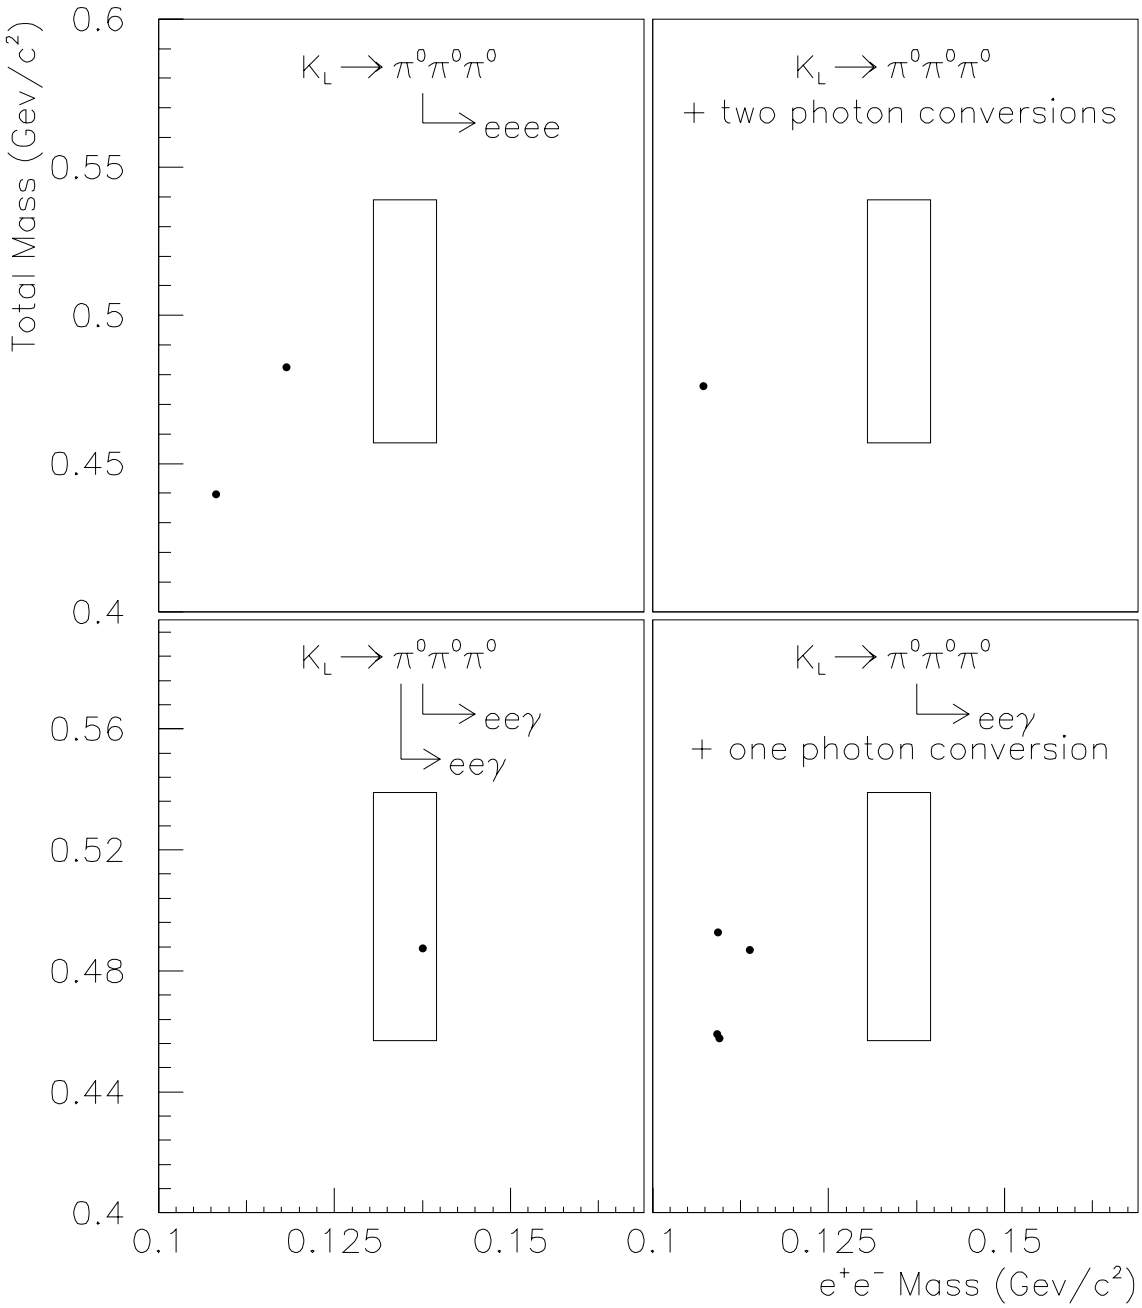

Figure 131. The total invariant mass versus the e^+e^- invariant mass for four-electron backgrounds after cuts to remove the four-electron background. Each simulation represents approximately 8.5 times number of such decays observed in the data.

two electrons with similar energy spectra and chamber illuminations (shown in Figures 108 and 111), this disagreement should cancel in the ratio of Monte Carlo acceptances.

Clearly though, the result was worth the loss in acceptance. The $\pi^0 \to e^+e^-$ data sample, after the addition of the Chamber 2 extra hit cuts, is shown in Figure 132. Nine events were found in the signal region, and the steep fall of the data in the region below the π^0 mass suggested a low background. Note also that there were no events with e^+e^- invariant mass above the high-side π^0 mass cut out to the kinematic limit of $K_L \to \pi^0\pi^0e^+e^-$, 228 MeV/c^2.

The $\pi^0 \to e^+e^-$ data and the $\pi^0 \to e^+e^-\gamma$ background Monte Carlo are shown in Figure 133. The background Monte Carlo prediction was absolutely normalized to the fully-reconstructed $\pi^0 \to e^+e^-\gamma$ normalization sample. The agreement between the two is excellent, and this gives us confidence in using the Monte Carlo to predict the small amount of background expected in the signal region.

With the cuts finalized, losses from each cut in the $\pi^0 \to e^+e^-$ and $\pi^0 \to e^+e^-\gamma$ normalization samples in data and Monte Carlo are shown in Table 27.

8.5 Extraction of the Branching Ratio for $\pi^0 \to e^+e^-$

With the completed analysis of the $\pi^0 \to e^+e^-$ data, the $\pi^0 \to e^+e^-\gamma$ normalization sample, and Monte Carlo predictions for signal, background and normalization modes, the extraction of the branching ratio can proceed. As was discussed in Chapter 7, the branching ratio of $\pi^0 \to e^+e^-$ is measured via the relation

$$\begin{aligned}\frac{,(\pi^0 \to e^+e^-)}{,(\pi^0 \to \text{all})} &\equiv \text{BR}(\pi^0 \to e^+e^-) \\ &= \text{BR}(\text{normalization}) \times \frac{N(\pi^0 \to e^+e^-)}{N(\text{normalization})} \times \frac{A(\text{normalization})}{A(\pi^0 \to e^+e^-)},\end{aligned}$$

where N is the number of observed decays, and A is the acceptance for that decay. Similarly, the number of expected events in the signal region, from a source of background can be measured using the Monte Carlo by

$$B(\text{background}) = N(\text{normalization}) \times \frac{A(\text{background})}{A(\text{normalization})} \times \frac{\text{BR}(\text{background})}{\text{BR}(\text{normalization})},$$

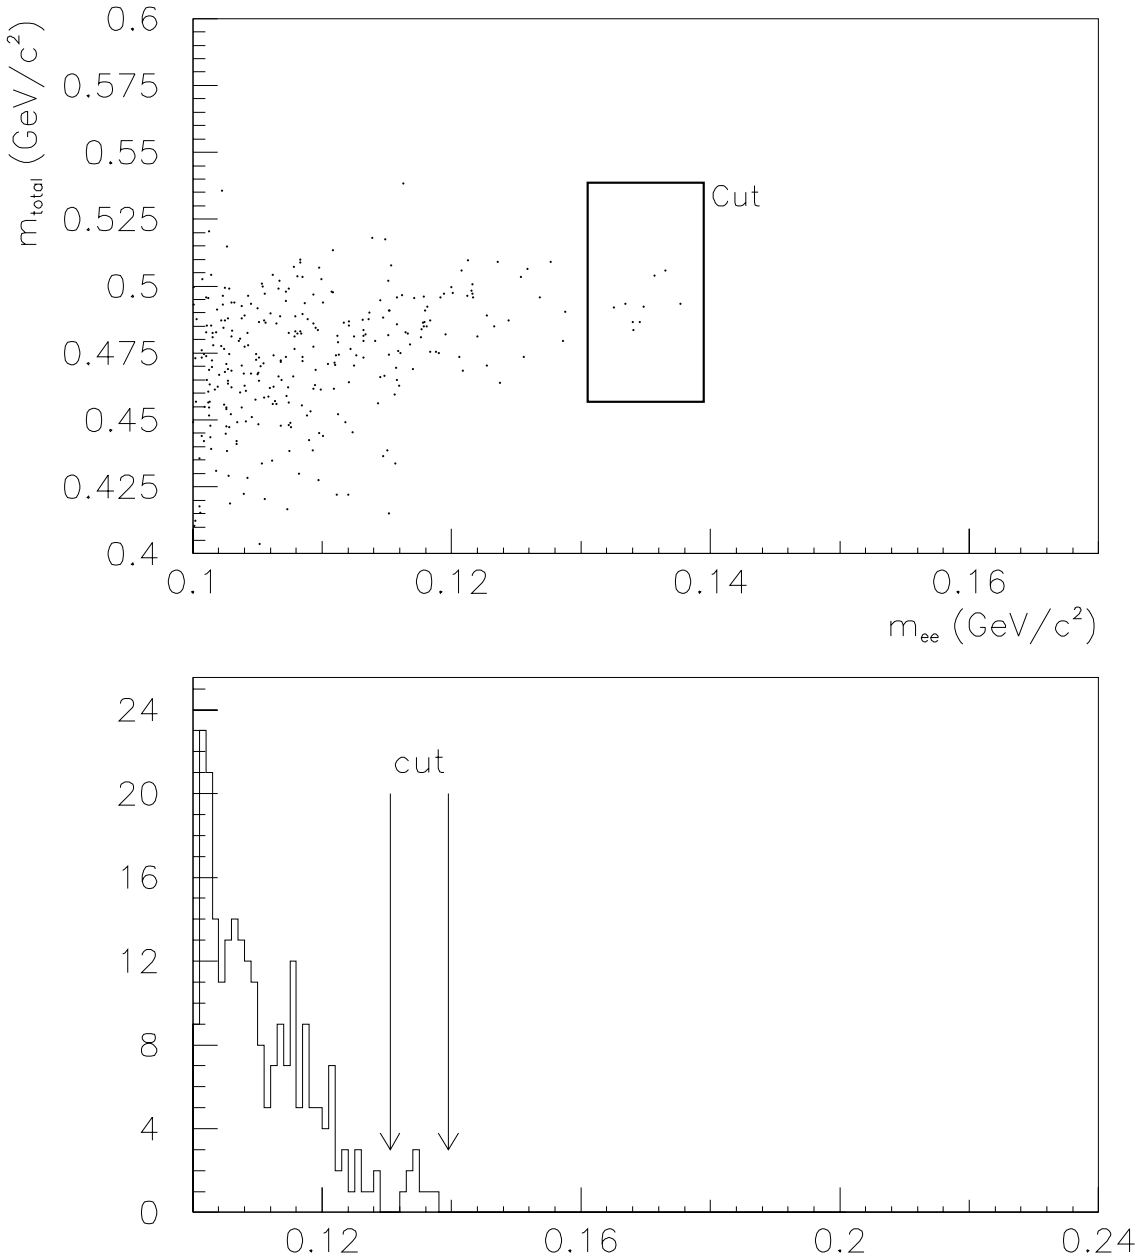

Figure 132. The total invariant mass vs. the e^+e^- invariant mass in the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, data sample after all cuts. The projection onto the $m_{e^+e^-}$ axis after a cut on the total mass is also shown.

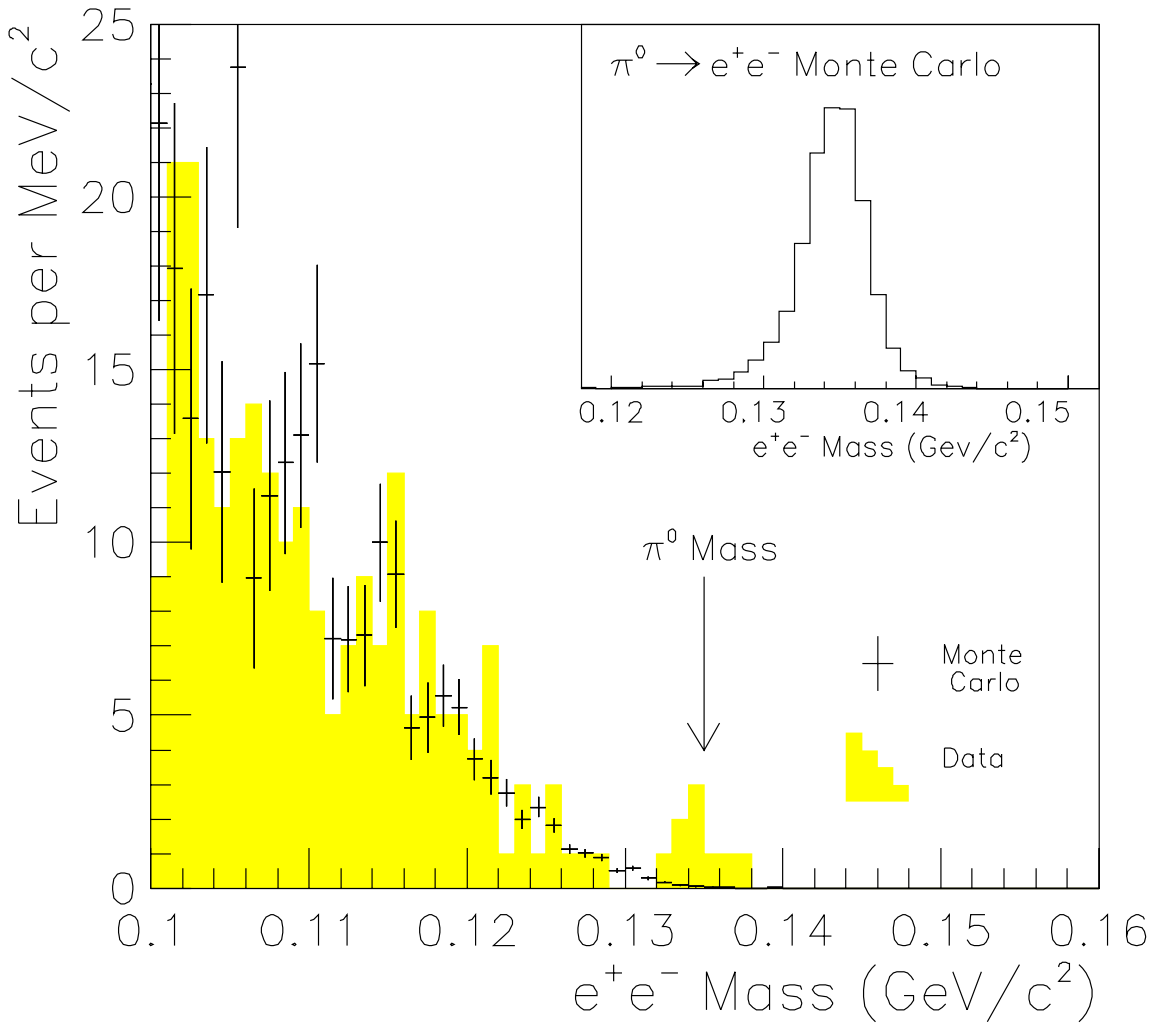

Figure 133. The e^+e^- invariant mass in the $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$, data sample and $\pi^0 \to e^+e^-\gamma$ background Monte Carlo. The histogram is the data, and the crosses are Monte Carlo. The background Monte Carlo is absolutely normalized to the observed number of fully reconstructed $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, decays. The expected distribution for $\pi^0 \to e^+e^-$ is also shown.

Table 27. The numbers of events in data and Monte Carlo $\pi^0 \to e^+e^-$ and $\pi^0 \to e^+e^-\gamma$ normalization samples after each cut

Cut	$\pi^0 \to e^+e^-$ data	$\pi^0 \to e^+e^-\gamma$ data	$\pi^0 \to e^+e^-$ MC	$\pi^0 \to e^+e^-\gamma$ MC
	(# of events)		(fraction of events) (generated)	
Passing Trigger	176,376,610		0.0257	0.00841
Crunch Tracking	57,820,493		0.0251	0.00820
Crunch Matching	45,225,690		0.0249	0.00813
Crunch E/p	18,805,391		0.0242	0.00790
6 or 7 Hardware Clusters	5,486,599		0.0242	0.00790
Two Tracks	5,486,599		0.0241	0.00781
No HCF Fusions	4,782,558		0.0237	0.00764
Vertex Found	4,167,436		0.0236	0.00763
$240 > m_{e^+e^-} > 68$ (MeV/c^2)	301,569	132,421	0.0215	0.00572
Trigger Verification	246,585	107,363	0.0204	0.00539
Track Quality	198,282	84,301	0.0198	0.00512
Vertex χ^2	181,152	77,138	0.0194	0.00502
Chamber Fiducial Cuts	177,424	75,680	0.0188	0.00486
PbG Fiducial Cuts	146,708	63,227	0.0174	0.00441
Fusion Cuts	29,620	24,791	0.0117	0.00300
Electron E/p	26,503	22,862	0.0113	0.00288
Photon Veto Cut	17,923	21,043	0.0108	0.00280
Back Anti Cut	14,105	17,326	0.0085	0.00218
Kinematic Cuts	50	13,834	0.0066	0.00185
Chamber 2 Extra Hits	9	10,039	0.0061	0.00171

where B is number of events in the data that result from the background source, and where A(background) is understood to denote the fraction of generated background events passing all of the cuts in the $\pi^0 \to e^+e^-$ analysis.

For convenience, all Monte Carlo simulations used the same range of generated K_L momenta and decay vertices, with a K_L momentum between 35 and 220 GeV/c and a decay vertex between 90 and 160 meters from the target. In this range of momenta and decay vertex locations, one can measure the number of π^0 decays resulting from $K_L \to \pi^0\pi^0\pi^0$ decays by using the normalization mode. The $\pi^0 \to e^+e^-\gamma$ sample was analyzed with a cut requiring $m_{e^+e^-} > 70$ MeV/c^2; however, in the Monte Carlo, Dalitz decays were generated with an e^+e^- above 65 MeV/c^2 because of the finite resolution on the measurement of $m_{e^+e^-}$. The fraction of Dalitz decays with an e^+e^- invariant mass above 65 MeV/c^2 was obtained by integrating the differential rate for Dalitz decays given in Section 1.1, and was found to be

$$\frac{,(\pi^0 \to e^+e^-\gamma, m_{e^+e^-} > 65 \text{ MeV/c}^2)}{,(\pi^0 \to e^+e^-\gamma)} \approx (3.233 + 2.08a) \times 10^{-2}, \quad \text{for } a \ll 1,$$

where a is the form-factor slope. For a, the measurement from the CELLO experiment [27] of $a = 0.0326 \pm 0.0026$ was used, and for the Dalitz branching ratio itself, the value was taken from the fit of the Particle Data Group, $(1.198 \pm 0.032)\%$ [3]. The Monte Carlo measured acceptance for the $\pi^0 \to e^+e^-\gamma$ normalization sample, including the 70 MeV/c^2 $m_{e^+e^-}$ invariant mass cut, was $(0.171 \pm 0.002)\%$. Radiative corrections to $\pi^0 \to e^+e^-\gamma$ caused a $(5.5 \pm 1.8)\%$ decrease in this acceptance, and were included in the quoted acceptance. 10039 $\pi^0 \to e^+e^-\gamma$ normalization events were observed in the data. Therefore, the total number of π^0 decays from $K_L \to \pi^0\pi^0\pi^0$ was $(14.9 \pm 0.5) \times 10^9$. Such a number must be taken with a large grain of salt, of course, because the Monte Carlo did not always accurately reproduce the effects of cuts, especially those sensitive to δ-rays in the chambers. Calculations involving ratios of acceptances were, however, expected to be much more reliable since effects like this cancel in these ratios.

Nevertheless, this number of decays roughly set the sensitivity scale. If the $\pi^0 \to e^+e^-$ branching ratio were 6×10^{-8} and the acceptance 1%, for example, one would expect about nine $\pi^0 \to e^+e^-$ decays observed. This total number of decays was also useful for

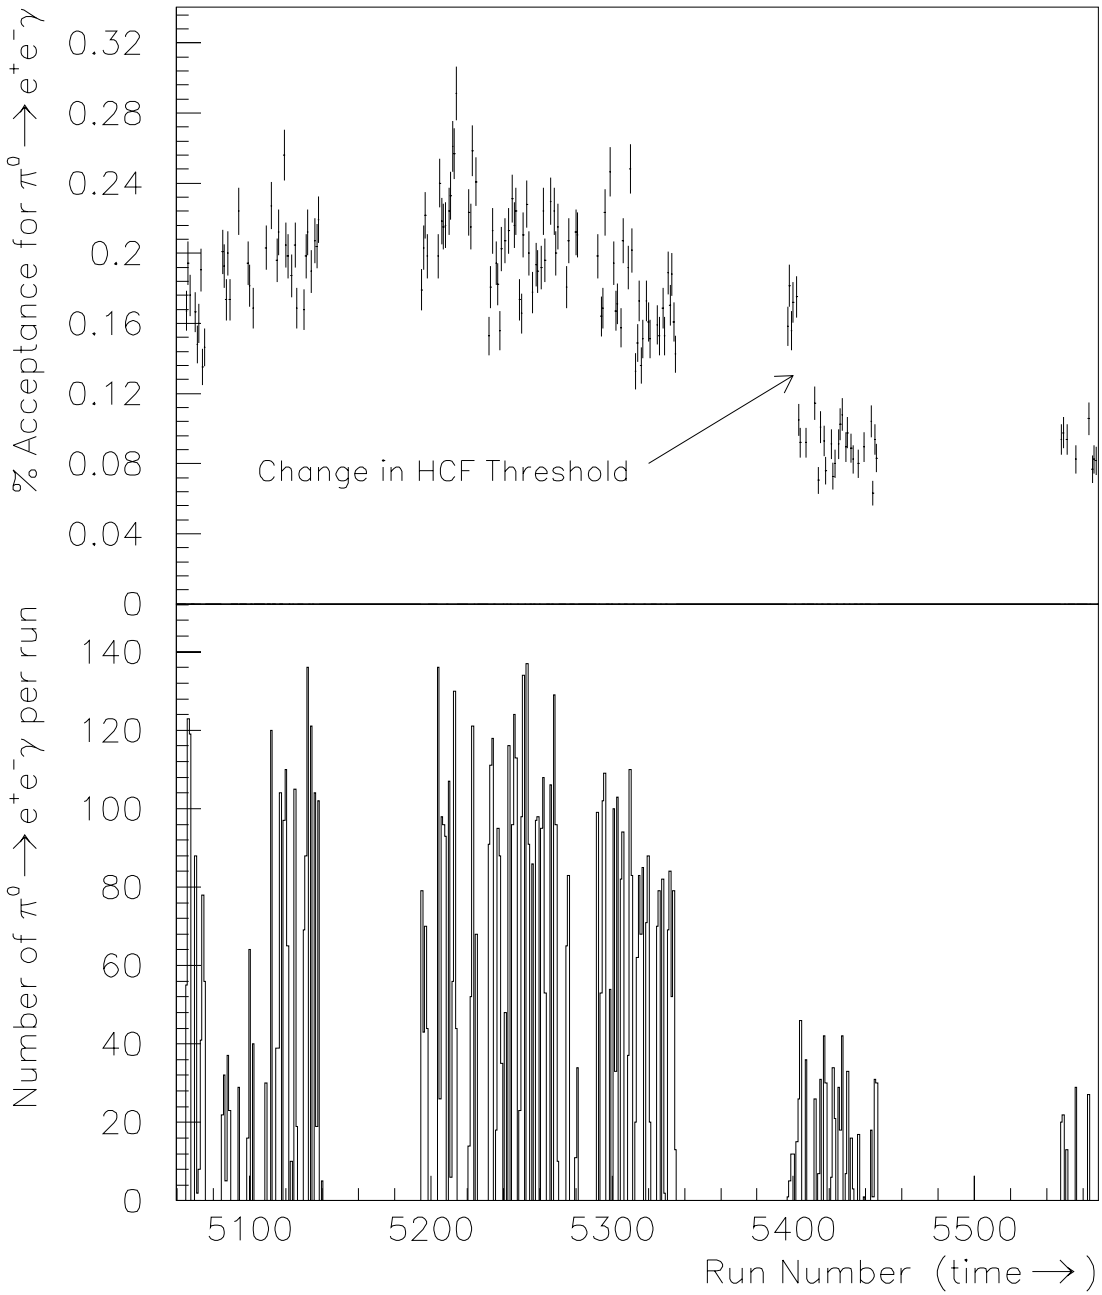

Figure 134. The acceptance and number of decays as a function of run number for the $\pi^0 \to e^+e^-\gamma$ normalization sample. The time-dependence of the acceptance is evident.

determining the number of Monte Carlo events that had to be generated to reproduce the number of decays seen in the data. For a four-electron background, for example, where the rate relative to π^0 decays was about 10^{-4}, approximately 1.5×10^6 events had to be generated in each mode to reproduce the data statistics. This number of decays was also broken down by run number in order to allow the Monte Carlo to simulate changes in acceptance with time. The acceptance for the $\pi^0 \to e^+e^-\gamma$ normalization sample, as a function of run, is shown in Figure 134, along with the number of observed events in the $\pi^0 \to e^+e^-\gamma$ normalization sample.

8.5.1 Monte Carlo Samples

The Monte Carlo samples generated for this analysis are listed in Table 28. The branching ratio given in this table is that for the entire decay chain, including the parent $K_L \to \pi^0\pi^0\pi^0$ decay in each case. The final column gives the ratio of the number of decays generated to the number needed to reproduce the number of decays seen in the data. This ratio, which will be referred to as the "Monte Carlo statistics factor", was equal to

$$\frac{N(\text{background})}{B(\text{background})} = G(\text{background}) \times \frac{A(\text{normalization})}{N(\text{normalization})} \times \frac{\text{BR}(\text{normalization})}{\text{BR}(\text{background})},$$

where G is the number of events generated in the Monte Carlo ($A \times G = N$). In the case of the $\pi^0 \to e^+e^-\gamma$ normalization sample, the error on this ratio came from the statistics of the data and Monte Carlo samples. In the other decays, the error in the Monte Carlo statistics factor came from this same statistical error, plus errors from branching ratio measurements and from the estimation of the external photon conversion probability.

With this Monte Carlo, it was straightforward to calculate the single event sensitivity to $\pi^0 \to e^+e^-$ decays, defined as

$$\frac{\text{BR}(\pi^0 \to e^+e^-)}{N(\pi^0 \to e^+e^-)} = \frac{\text{BR}(\text{normalization})}{N(\text{normalization})} \times \frac{A(\text{normalization})}{A(\pi^0 \to e^+e^-)}.$$

The acceptance of $\pi^0 \to e^+e^-$ was measured by the Monte Carlo to be $(0.614 \pm 0.009)\%$, and the ratio of acceptances was then

$$\frac{A(\text{normalization})}{A(\pi^0 \to e^+e^-)} = 0.279 \pm 0.005,$$

Table 28. Monte Carlo event samples used in the analysis

Sample	Branching Ratio	Events Generated	$\frac{\text{Number Generated}}{\text{Number of Decays in the Data}}$
$K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, $m_{ee} > 65$ MeV/c^2 (normalization)	2.50×10^{-4}	3899547	0.665 ± 0.011(stat)
$K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-$	??	799846	~ 500
$K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, (background)	7.58×10^{-3}	199857† $\sum$ weight $= 1.46 \times 10^9$	252 ± 4
$K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, $m_{ee} > 65$ MeV/c^2 (four electron normalization)‡	2.50×10^{-4}	499927	0.241 ± 0.005(stat)
$K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-e^+e^-$,	2.1×10^{-5}	1399921	8.5 ± 0.9
$K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, $\pi^0 \to e^+e^-\gamma$	9.1×10^{-5}	6449772	8.33 ± 0.44
$K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, $\gamma \to e^+e^-$	9.7×10^{-5}	8699550	10.5 ± 0.9
$K_L \to \pi^0\pi^0\pi^0$, $\gamma \to e^+e^-$, $\gamma \to e^+e^-$	2.1×10^{-5}	2299860	12.9 ± 1.6

† The $\pi^0 \to e^+e^-\gamma$ background Monte Carlo was generated with a flat e^+e^- invariant mass spectrum, but with events weighted according to the Dalitz e^+e^- mass spectrum.

‡ The four-electron background Monte Carlo was generated with a different lead-glass response parameterization, and with no accidentals overlaid. To account for the different definition of acceptance in this Monte Carlo, the same Monte Carlo was used to generate a sample of $\pi^0 \to e^+e^-\gamma$ decays for use in normalizing the four-electron background.

where the quoted error was purely due to Monte Carlo statistics. The single-event sensitivity was then

$$\frac{\mathrm{BR}(\pi^0 \to e^+e^-)}{N(\pi^0 \to e^+e^-)} = \frac{(3.95 \pm 0.10) \times 10^{-4}}{10039} \times (0.279 \pm 0.005) = (1.10 \pm 0.04) \times 10^{-8}.$$

This error was dominated by the error on the measurement of the Dalitz branching ratio. Contributions of other systematic errors will be discussed in Section 8.6.

8.5.2 Estimation of Background to the $\pi^0 \to e^+e^-$ data sample

In the four electron background Monte Carlo, a total of one background event in all four modes was generated that passed all $\pi^0 \to e^+e^-$ requirements. Using the Monte Carlo statistics factor for $K_L \to \pi^0\pi^0\pi^0$, $\pi^0 \to e^+e^-\gamma$, $\pi^0 \to e^+e^-\gamma$, decays (the one background event was observed in this channel), this results in a background prediction of

$$0.12^{+0.28}_{-0.10} \text{ events}$$

from four electron final-states. The asymmetric errors represent a 68% confidence interval from Poisson statistics on the single observed Monte Carlo event.

The background from $\pi^0 \to e^+e^-\gamma$ was higher than this, as can be seen from Figure 133. From the $\pi^0 \to e^+e^-\gamma$ background Monte Carlo,

$$0.85 \pm 0.10 \text{ events}$$

from high e^+e^- mass Dalitz decays were predicted to be in the signal region. The error on this prediction was, again, from Monte Carlo statistics only. The total background was thus predicted to be approximately one event.

Therefore, the number of of observed events attributed to $\pi^0 \to e^+e^-$ was

$$8.0^{+4.1}_{-2.9} \text{ events.}$$

The statistical error on this number represents the 68% confidence interval from the Poisson statistics on nine observed events, plus the small contribution from the statistical error on the background determination. With the number of events above background known, a branching ratio may now be quoted.

8.5.3 Radiative Corrections to $\pi^0 \to e^+e^-$

Before doing extracting this branching ratio, however, some sober reflection on the effects of radiative corrections on the $\pi^0 \to e^+e^-$ measurement is in order. The Monte Carlo generates events with the full predicted e^+e^- spectrum from bremsstrahlung diagrams (Figure 7) and also events which do not produce any final state particles because of the destructive interference between the vertex correction diagram (Figure 6) and the tree-level process. Therefore, the measured branching ratio was actually the branching ratio one would expect in the absence of radiative corrections.

Does it make sense to quote this as a "branching ratio"? I would argue no for two reasons. First this branching ratio actually includes events which "disappear" due to the destructive interference of radiative diagrams! Clearly, this has no business being quoted in a branching ratio. Second, for e^+e^- invariant mass below 110 MeV/c^2, it is possible for the magnitude of the interference of Dalitz decay and $\pi^0 \to e^+e^-$ to be larger than the rate for $\pi^0 \to e^+e^-$ itself. What it means in this case to quote a rate for $\pi^0 \to e^+e^-$ separate from Dalitz decays is less than clear. Therefore I have chosen to quote a branching ratio only for the high e^+e^- invariant mass excess over Dalitz decays from $\pi^0 \to e^+e^-$.

In order to minimize the dependence on the model of radiative corrections, the range of e^+e^- invariant masses included in the branching ratio was chosen to be similar to that selected by the experimental cut, as shown in Figure 135. From the model of radiative corrections, 86.2% of events generated by the Monte Carlo simulation of $\pi^0 \to e^+e^-$ produced an electron-positron pair with an invariant mass in the range $0.95 < x < 1$, where $x = (m_{ee}/m_{\pi^0})^2$. The branching ratio extracted for $\pi^0 \to e^+e^-$ was then,

$$\mathrm{BR}(\pi^0 \to e^+e^-, (\frac{m_{ee}}{m_{\pi^0}})^2 > 0.95) = (0.862) \times N(\pi^0 \to e^+e^-) \times (1.10 \pm 0.04) \times 10^{-8},$$

where the first factor is from the restriction of the e^+e^- invariant mass range, and the last is the single-event sensitivity computed above.

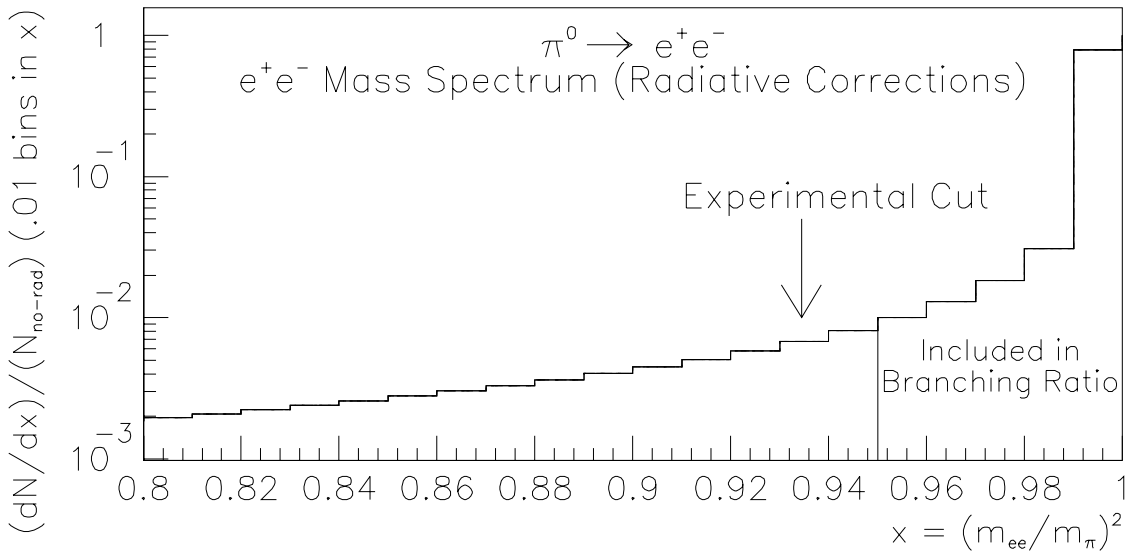

Figure 135. The predicted e^+e^- invariant mass spectrum from radiative corrections to the $\pi^0 \to e^+e^-$ process. The location of the experimental cut and the range included in the quoted branching ratio are also shown.

8.6 Systematic Errors

Although it may seem pointless to discuss systematic errors to a measurement with such a large statistical errors, the studies involved in calculating systematic errors were important to make the final result convincing. Systematic effects were particularly important because the subtraction of the Dalitz background which, after all, was significantly larger than the size of the signal only a few MeV/c^2 in e^+e^- invariant mass away from the signal region.

Systematic errors were measured by making a change in the analysis of Monte Carlo event samples and observing the effect on the ratio of acceptances. For example, to study the effect of residual differences between data and Monte Carlo in the glass energy resolution, energy smearing was added to or removed from cluster energies in Monte Carlo events, and the effect on the acceptance in the normalization mode and the signal or background mode was computed. Since only the ratios of acceptances entered into quantities used to extract the branching ratio, it was the change in these quantities that contributed to systematic errors.

The most important systematic effects were those involving the drift chamber invariant mass resolution and scale. This resolution and scale was checked in studies involving $K_L \to \pi^+\pi^-$ decays (see Section 5.6.1). The invariant mass scale was shifted 2.5σ (0.05%), and two models of additional momentum smearing based on the resolution observed in $K_L \to \pi^+\pi^-$ decays were tested. The total systematic error in the single-event sensitivity was about 1%, and the error on the background prediction was large, 20%, but even this was small compared to the statistical errors.

The effects of misunderstanding the lead-glass were studied by changing the glass energy scale, adding uncorrected pedestal shifts (20 MeV coherent, 60 MeV random), smearing the lead-glass gains, and adding large coherent changes to calibration constants in the most damaged blocks. Because the cuts on quantities affected by the lead-glass energy and position measurement were loose, the effects were fairly small.

Systematic effects also resulted from errors in measurements of branching ratios, the $\pi^0 \to e^+e^-\gamma$ branching ratio in particular, and in the error in the measurement of the π^0 form-factor slope. Because there was some disagreement in past measurements of this slope

in Dalitz decays (see Figure 5), this slope was allowed to vary by ± 0.1. The effect of this systematic error was reduced somewhat by the fact that the the error in single-event sensitivity and background had opposite signs and thus canceled to some extent.

There was some dependence in the number of π^0 decays measured by the normalization mode on the e^+e^- invariant mass cut used in the analysis of this mode. Because this effect was greater than one statistical standard deviation, a systematic error from this change was included to be conservative.

Another systematic study was done to determine the effect inefficient wires in Chamber 2 had on the efficiency of the cut used to remove the four-electron background. Because this cut was so important, it was necessary to know if there was any effect from imperfect knowledge the chamber efficiencies. Four wires, two in the center of each view of the chamber, were turned off in software for the purposes of the Chamber 2 extra hits cut, and no change in the background prediction resulted. Therefore, no systematic error was added to account for uncertainties in wire efficiencies.

Finally, one could ask if there were any backgrounds which were not included in the Monte Carlo simulations. For example, backgrounds involving charged pions, misidentified as electrons, in the final state would fall into this category. Such backgrounds unrelated to $K_L \to \pi^0\pi^0\pi^0$ decays were studied by looking in sidebands in one kinematic variable, while keeping the cuts on all other variables the same. The regions in p_t^2 between 1025 $(\mathrm{MeV/c})^2$ and 10250 $(\mathrm{MeV/c})^2$, and in e^+e^- mass between 139.5 $\mathrm{MeV/c^2}$ and 240 $\mathrm{MeV/c^2}$ were studied. If a background unrelated to $K_L \to \pi^0\pi^0\pi^0$ decays had existed in the sample, it is plausible that such a background would have been flat in p_t^2 and in $m_{e^+e^-}$. No events were observed in these sidebands, and from this fact and the assumption of flat backgrounds, one can derive an upper limit of 0.26 events at the 90% confidence level from such backgrounds.

All the considered sources of systematic error are shown in Table 29. The contribution to the final result from the large fractional errors in backgrounds were of course diminished because the backgrounds themselves were so small. The total systematic error on the branching ratio measurement, combining the effects on the single-event sensitivity and backgrounds in quadrature, was 7.1%.

Table 29. The systematic errors

Source of Error	$\left\lvert\Delta\frac{A(\pi^0 \to e^+e^-)}{A(\text{normalization})}\right\rvert$	$\left\lvert\Delta\frac{A(\text{background})}{A(\text{normalization})}\right\rvert$
Invariant mass scale shift (0.05%)	0.1%	6%
Momentum resolution smear (Gaussian)	0.2%	12%
Momentum resolution smear (non-Gaussian tail)	1.0%	2%
Chamber Total	1.3%	20%
PbG energy scale Shift (1%)	0.6%	1%
Gain Smear (2%)	0.2%	3%
Pedestal Shifts	0.1%	1%
Damaged block study (based on data-MC disagreements in high α blocks)	0.3%	3%
Total PbG	1.2%	8%
Dalitz branching ratio error	2.7%	$< 1\%$
Form-factor slope error (0.1)	5.0%	12%
Normalization e^+e^- mass cut	1.9%	1.9%
Non-simulated backgrounds		27%
Total Error	6.2%	37%

CHAPTER 9

CONCLUSIONS

Even my sweat smells clean.

— *Emily Saliers*

With a measured branching ratio for $\pi^0 \to e^+e^-$, one can compare the result with that expected by theory. The real part of the amplitude for $\pi^0 \to e^+e^-$ is extracted from all available experimental data, and the question of non-Standard Model contributions is considered.

9.1 This Result

With nine events observed, including a predicted background of approximately one event, we have measured

$$\mathrm{BR}(\pi^0 \to e^+e^-, (\frac{m_{ee}}{m_{\pi^0}})^2 > 0.95) = (7.6^{+3.9}_{-2.8}(\mathrm{stat}) \pm 0.5(\mathrm{syst})) \times 10^{-8},$$

over and above the contribution from Dalitz decays [75]. The model of radiative corrections $\pi^0 \to e^+e^-$ developed in Section 1.3 was used to predict an e^+e^- invariant mass spectrum for $\pi^0 \to e^+e^-$. This spectrum can be used to predict the rate for $\pi^0 \to e^+e^-$ with different cuts on $m_{e^+e^-}$, as shown in Figure 136.

The same model of radiative corrections can also be used to predict a rate for $\pi^0 \to e^+e^-$ in the absence of radiative corrections, $,^0_{ee}$. This rate is useful because it is the rate that is predicted by theoretical calculations of $\pi^0 \to e^+e^-$. $,^0_{ee}$ was obtained by integrating over the entire bremsstrahlung spectrum, and then adding back the radiative correction to the total rate. Through this procedure, we obtain

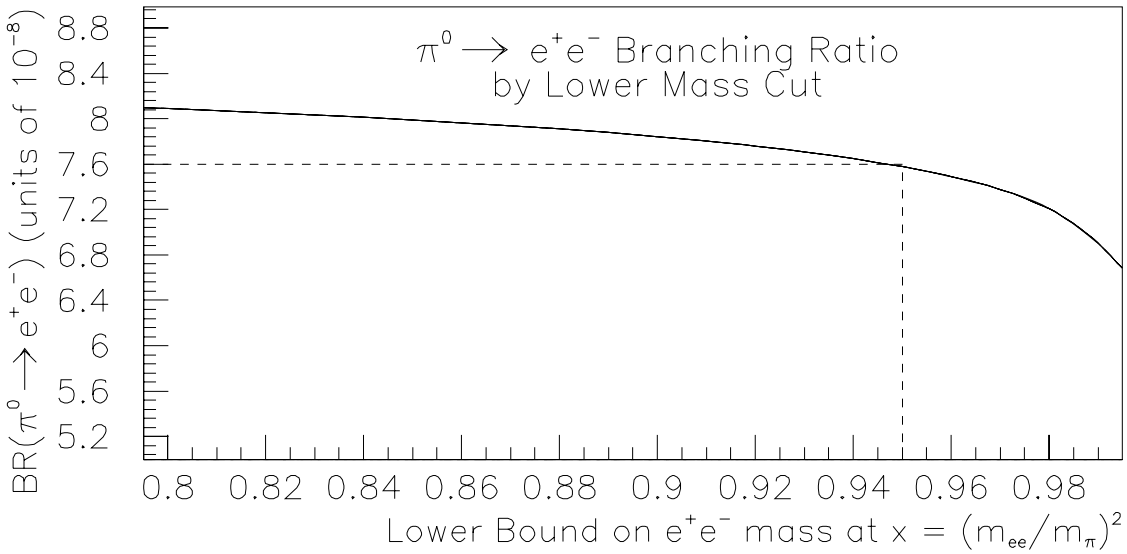

Figure 136. The branching ratio for $\pi^0 \rightarrow e^+e^-$ for $x > x_{\mathrm{cut}}$ as a function of x_{cut} for the measurement described in this thesis

$$\frac{,\ ^{0}_{ee}}{,\ \text{all}} = (8.8^{+4.5}_{-3.2}(\text{stat}) \pm 0.6(\text{syst})) \times 10^{-8},$$

where $,_{\text{all}}$ is the total decay rate for the π^0.

9.2 All Experimental Results

The experimental situation before this measurement was discussed in Section 1.4. The measurement described in this thesis was only slightly more sensitive than the best previous measurements. However, there was a vast improvement in background rejection compared to past efforts. Despite observing only nine events, the background was sufficiently low that the detection of $\pi^0 \to e^+e^-$ had a very high statistical confidence. The probability of the 1 event background in the signal region fluctuating to 9 or more observed events is about 10^{-6}.

There was also a new measurement from the Brookhaven E851 collaboration [76] which was published simultaneously with this result. The E851 measurement was made using the K^+ decay-in-flight technique, and was significantly more sensitive than this measurement, observing 21 ± 7 events attributed to $\pi^0 \to e^+e^-$. However, this measurement had background from both $K^+ \to \pi^+e^+e^-$ and $K^+ \to \pi^+\pi^0$, $\pi^0 \to e^+e^-\gamma$, and only obtained a signal to noise ratio of approximately 1:1 in $m_{e^+e^-}$ at the π^0 mass peak. From this sample, they measure

$$\text{BR}(\pi^0 \to e^+e^-) = (6.9 \pm 2.3(\text{stat}) \pm 0.6(\text{syst})) \times 10^{-8} \text{ [76]}.$$

This branching ratio measurement did not include any radiative corrections or quote any explicit cut-off in the e^+e^- invariant mass. They quote a second branching ratio that results from incorporating a model for radiative corrections [48] into their Monte Carlo.

$$\text{BR}(\pi^0 \to e^+e^-) = (8.0 \pm 2.6(\text{stat}) \pm 0.6(\text{syst})) \times 10^{-8} \text{ [76]}.$$

This second value is equivalent to the $,\ ^{0}_{ee}/,_{all}$ defined here. Branching ratio results from all measurements of $\pi^0 \to e^+e^-$ are shown in Figure 137.

It is interesting to make a best fit to all the available data to extract a branching ratio. Since the early detections of $\pi^0 \to e^+e^-$ from the Geneva-Saclay and LAMPF groups

Branching Ratio (units of 10^{-8})	Technique	Date	Experiment
22.3^{+24}_{-11}	K^+ decay in flight (90% error bars)	1978	Geneva-Saclay [49]
$17 \pm 6(\text{stat}) \pm 3(\text{syst})$	$\pi^- p \to \pi^0 n$ in flight	1983	LAMPF [50]
< 53	$\pi^- p \to \pi^0 n$ in flight (90% confidence)	1987	OMICRON [51]
< 13	$\pi^- p \to \pi^0 n$ at rest (90% confidence)	1989	SINDRUM [52]
$7.6^{+3.9}_{-2.8} \pm 0.5$ ($x > 0.95$)	K_L decay in flight	1993	Fermilab E799 [75]
$6.8 \pm 2.3 \pm 0.6$	K^+ decay in flight	1993	Brookhaven E851 [76]

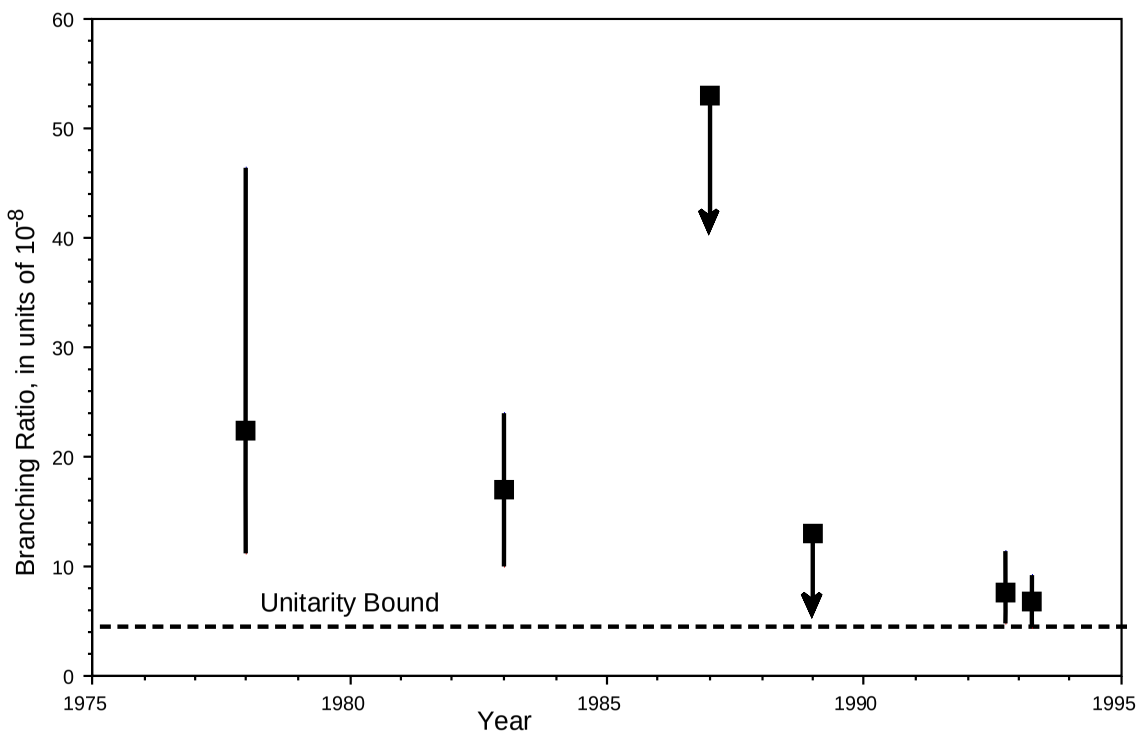

Figure 137. All experimental measurements of $\text{BR}(\pi^0 \to e^+e^-)$

have already been shown to be inconsistent with the SINDRUM measurement, these two measurements were excluded from the fit. The OMICRON measurement was also not included because it did not significantly affect the result.

In order to correctly include the effects of final-state radiation, the values of $,^0_{ee}/,_{\text{all}}$, rather than the measured branching ratios were combined. In the case of the SINDRUM measurement, their quoted branching ratio is equivalent to $,^0_{ee}/,_{\text{all}}$, because their normalization interaction, $\pi^- p \to n e^+ e^-$ included similar final state radiative effects [52]. For this comparison, statistical and systematic errors from each measurement were combined. The probability distributions in $,^0_{ee}/,_{\text{all}}$ extracted from the measurement described here [75], the Brookhaven E851 measurement [76], and the SINDRUM upper limit [52] are shown in Figure 138, along with the combined probability distribution. This combination of the three measurements yields a best fit value of

$$\frac{,^0_{ee}}{,_{\text{all}}} = (7.8 \pm 2.1) \times 10^{-8}.$$

Using the model of radiative corrections described in Section 1.3, this leads to a branching ratio of

$$\text{BR}(\pi^0 \to e^+e^-, (\frac{m_{ee}}{m_{\pi^0}})^2 > 0.95) = (6.7 \pm 1.9) \times 10^{-8}.$$

9.3 Implications for Models of $\pi^0 \to \gamma\gamma \to e^+e^-$

Recall from Section 1.2, that the rate for $\pi^0 \to \gamma\gamma \to e^+e^-$ can be written as

$$\frac{,(\pi^0 \to \gamma^*\gamma^* \to e^+e^-)}{,(\pi^0 \to \gamma\gamma)} = 4.75 \times 10^{-8} \times (1 + r^2),$$

with

$$r = \left|\frac{\text{Re}\{R\}}{\text{Im}\{R\}}\right|,$$

where R is proportional to the amplitude for $\pi^0 \to e^+e^-$. Recall that the imaginary part of this amplitude was known, but the real part was model-dependent. One can use the measurement of the $\pi^0 \to e^+e^-$ branching ratio to place constraints on allowed values for r. Shown in Figure 139 are the probability curves as a function of r from considering this

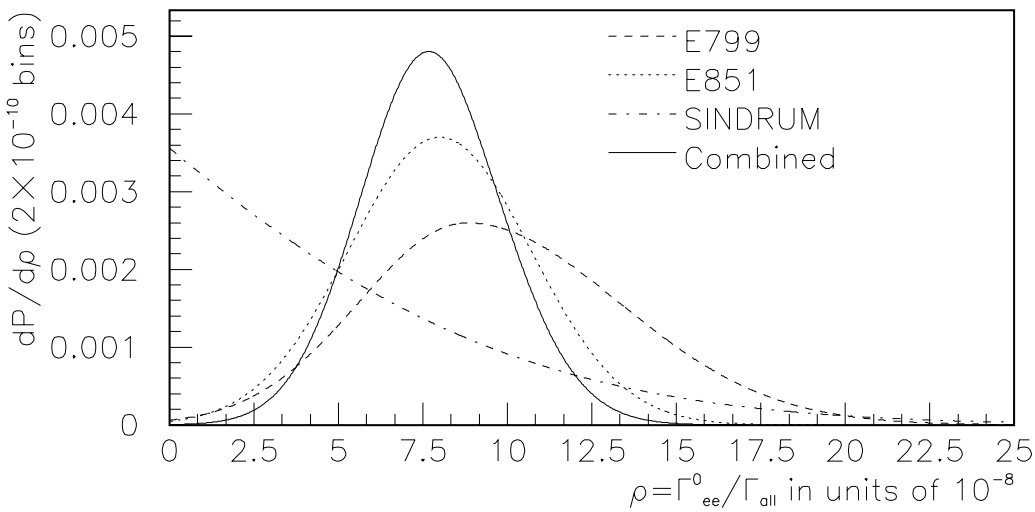

Figure 138. Probability distributions for , $^{0}_{ee}/$, $_{\mathrm{all}}$ from the three most significant measurements of $\pi^0 \to e^+e^-$ and the combined probability distribution

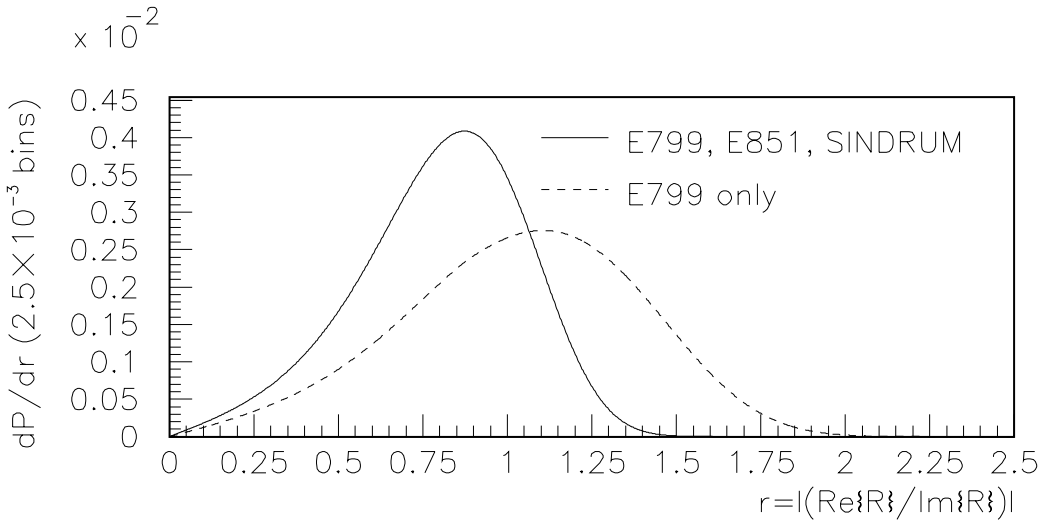

Figure 139. Probability distributions for the real part of the amplitude for $\pi^0 \to e^+e^-$ from this experiment alone and from the best fit from all experimental data

result only, and also from all the experimental data combined. If we consider only the result described in this thesis, one obtains

$$0.68 > r > 1.35 \quad (68\% \text{ confidence})$$
$$0.42 > r > 1.59 \quad (90\% \text{ confidence}).$$

The combined experimental data favors a smaller real portion of the amplitude,

$$0.54 > r > 1.05 \quad (68\% \text{ confidence})$$
$$0.34 > r > 1.17 \quad (90\% \text{ confidence}).$$

Both results are consistent with every model listed in Table 2, except the nucleon loop model [35] and the dispersion relation calculation of Tupper and Samuel [37]. The disagreement with the latter should be no surprise, however, since this calculation reflected an effort to explain the seemingly anomalous early results of the Geneva-Saclay and LAMPF groups.

One can actually use the results from the measurement of $\pi^0 \to e^+e^-$ and the related process $\eta \to \mu^+\mu^-$ to obtain the sign of the real part of the amplitude for $\eta \to \mu^+\mu^-$. From the work of Ametller, Bramon and Massó [34][77],

$$\mathrm{Re}\{R_{\pi ee}\} - \mathrm{Re}\{R_{\eta\mu\mu}\} = 12 \pm 2,$$

where the error includes a estimation of the model-dependence of the result. (If one assumes their Vector Meson Dominance model [34], the error on the difference is much smaller, 11.4 ± 0.4.) Since the measurement of $\eta \to \mu^+\mu^-$ at SATURNE [41] gives

$$\mathrm{Re}\{R_{\eta\mu\mu}\} \approx \pm(3 \pm 1),$$

the sign of the amplitude for $\pi^0 \to e^+e^-$ is clearly positive. All the available data can be combined to attempt to determine the sign of this amplitude for $\eta \to \mu^+\mu^-$, but even with the improvements in the $\pi^0 \to e^+e^-$ measurement, the results are still inconclusive, as shown in Figure 140.

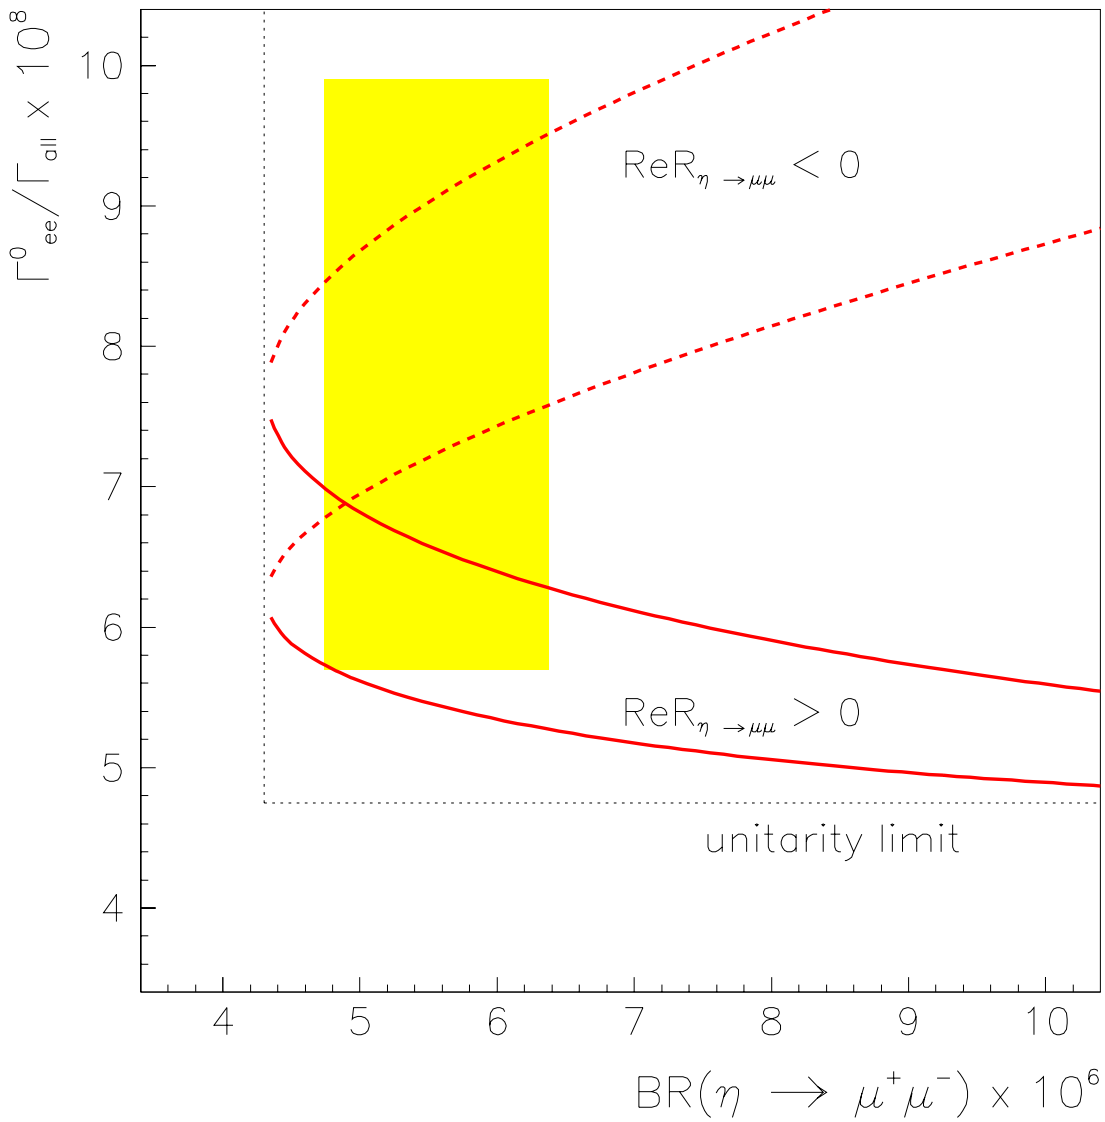

Figure 140. The implications of the measurement of $\eta \to \mu^+\mu^-$ and $\pi^0 \to e^+e^-$ for the sign of the real part of the amplitude for $\eta \to \mu^+\mu^-$. The shaded region reflects one standard deviation errors on each measurement. The area between the curves reflects allowed regions for each sign of the real part of the $\eta \to \mu^+\mu^-$ amplitude, including theoretical errors. *Thanks to R. Kessler for producing this figure.*

9.4 Constraints on non-Standard Model Processes

Unfortunately for those who like the exotic, this new measurement of $\pi^0 \to e^+e^-$ leaves little room for non-Standard Model contributions. The observed rate is now completely consistent with what one expects from models such as Vector Meson Dominance.

Massó has considered the effect of a short-lived axion with a mass of approximately 1.8 MeV/c^2 coupling to first-generation fermions on the decay $\pi^0 \to e^+e^-$ [47]. (Such an axion was briefly thought to have been capable of escaping all previous attempts at experimental detection.) The branching ratio of $\pi^0 \to e^+e^-$ including only the effects of this axion mixing with the π^0 was shown to be greater than 16×10^{-8} for any such axion capable of decaying to electrons ($m_a > 2m_e$). This result thus rules out the existence of such an axion.

The decay $\pi^0 \to e^+e^-$ has also been considered as a place to observe the effects of light pseudoscalar Higgs particles [46]. Unfortunately, the light mass of the π^0 makes any effect very small, and any contribution from such a Higgs would be several orders of magnitude below the unitarity bound. (A similar search in a leptonic decay of a much heavier pseudoscalar, like an η_b, would prove much more fruitful.) Thus, these new results can shed little light on the Higgs sector.

Perhaps the most interesting of the non-Standard Model effects which can be studied in $\pi^0 \to e^+e^-$ are those of leptoquarks. As has been pointed out by several authors [45][78], a pseudoscalar lepton-hadron coupling would appear as a contribution to the real part of the $\pi^0 \to e^+e^-$ amplitude. The limits on the real part of the amplitude presented here could be used to limit such contributions, but better limits on the lepotquarks to which $\pi^0 \to e^+e^-$ is sensitive exist from recent atomic parity violation experiments and from measurements of muon conversion on nuclei [79].

9.5 Summary

The rare decay $\pi^0 \to e^+e^-$ has been observed, and its branching ratio measured. This result represents the only low-background detection of this decay to date, and helps to resolve a long-standing experimental puzzle in the rate for this decay. This result also confirms

theoretical prejudices that predict the rate to be near the unitarity limit for the decay. No non-Standard Model effects are needed to explain the observed rate. Improvements in this measurement could help to distinguish to some extent among various models for structure in the $\pi^0\gamma\gamma$ vertex, although the differences among models are small. The most immediate result from an improvement in the measurement might be a definitive determination of the sign of the real part of the amplitude for $\eta \to \mu^+\mu^-$ which is still uncertain at the present. Planned rare K_L and K^+ decay experiments should be able to make significant improvements in the precision of this measurement in the near future.

REFERENCES

[1] D.A. Harris, PhD thesis, Univerisity of Chicago, 1994.

[2] V.L. Fitch J.H. Christenson, J.W. Cronin and R. Turlay, *Phys. Rev. Lett.* **13**, 138 (1964).

[3] Particle Data Group, *Phys. Rev.* **D45**, 1 (1992).

[4] N. Samios, *Phys. Rev.* **121**, 275 (1961).

[5] Dorenbos *et al.*, *Z. Phys.* **C40**, 497 (1988).

[6] McDonough *et al.*, *Phys. Rev.* **D38**, 2121 (1988).

[7] P. Krolak *et al.* A Limit on the Lepton-Family Number Violating Process $\pi^0 \to \mu^\pm e^\mp$. Preprint EFI-93-63 (accepted by Phys. Lett.), 1993.

[8] R.H. Dalitz, *Proc. Phys. Soc.* **A64**, 667 (1951).

[9] N.M. Kroll and W. Wada, *Phys. Rev.* **98**, 1355 (1955).

[10] B. Lautrup and J. Smith, *Phys. Rev.* **D3**, 1122 (1971).

[11] K. Mikaelian and J. Smith, *Phys. Rev.* **D5**, 1763 (1972).

[12] L. Roberts and J. Smith, *Phys. Rev.* **D33**, 3457 (1986).

[13] D. Joseph, *Nuovo Cim.* **16**, 997 (1960).

[14] S.M. Berman and D.A. Geffen, *Nuovo Cim.* **18**, 1192 (1960).

[15] L.G. Landsberg, *Phys. Rep.* **128**, 302 (1985).

[16] M. Pratap and J. Smith, *Phys. Rev.* **D5**, 2020 (1972).

[17] G.A. Kozlov *et al.*, *Z. Phys.* **C21**, 63 (1983).

[18] M. Dinejhan *et al.*, *Proc. Yad. Fiz.* **34**, 264 (1981).

[19] H. Kobrak, *Nuovo Cim.* **20**, 1115 (1961).

[20] S. Devons *et al.*, *Phys. Rev.* **184**, 1356 (1969).

[21] J. Burger, PhD thesis, Columbia University, 1972.

[22] J. Fischer *et al.*, *Phys. Lett.* **B73**, 359 (1978).

[23] P. Gumplinger, PhD thesis, Oregon State University, 1987.

[24] H. Fonvieille *et al.*, *Phys. Lett.* **B233**, 65 (1989).

[25] R.M. Drees *et al.*, *Phys. Rev.* **D45**, 1439 (1992).

[26] F. Farzanpay *et al.*, *Phys. Lett.* **B278**, 413 (1992).

[27] H.J. Behrend *et al.*, *Z. Phys.* **C49**, 401 (1991).

[28] R. Plano *et al.*, *Phys. Rev. Lett.* **3**, 525 (1959).

[29] S.D. Drell, *Nuovo Cim.* **11**, 693 (1958).

[30] L. Bergström *et al.*, *Phys. Lett.* **B126**, 117 (1983).

[31] C. Quigg and J.D. Jackson. UCRL-18487 (unpublished), 1968.

[32] A.I. Ivanov and V.M. Shekhter, *Sov. J. Nucl. Phys.* **32**, 410 (1980).

[33] K.S. Babu and E. Ma, *Phys. Lett.* **B119**, 449 (1982).

[34] A. Bramon Ll. Ametller and E. Massó. UAB-FT-300-93, 1993.

[35] Ll. Ametller *et al.*, *Nucl. Phys.* **B228**, 301 (1983).

[36] G.V. Efimov *et al.*, *JETP Lett.* **34**, 221 (1981).

[37] G.B. Tupper and M.A. Samuel, *Phys. Rev.* **D26**, 3302 (1982).

[38] L. Bergström and E. Ma, *Phys. Rev.* **D29**, 1029 (1984).

[39] A. Pich and J. Bernabéu, *Z. Phys.* **C22**, 197 (1984).

[40] M. Luke M.J. Savage and M. Wise, *Phys. Lett.* **B291**, 481 (1992).

[41] G.K. Kessler *et al.*, *Phys. Rev. Lett.* **70**, 892 (1993).

[42] A. Litke *et al.*, *Phys. Rev. Lett.* **30**, 1189 (1973).

[43] M. Schwartz, 1973. This idea is credited to M. Schwartz in [44].

[44] J.D. Davies *et al.*, *Nuovo Cim.* **A24**, 324 (1974).

[45] J.C. Pati and A. Salam, *Phys. Rev.* **D11**, 1137 (1975).

[46] G.L. Kane H.E. Haber and T. Sterling, *Nucl. Phys.* **B161**, 493 (1979).

[47] E. Massó, *Phys. Lett.* **B181**, 388 (1986).

[48] L. Bergström, *Z. Phys.* **C20**, 135 (1983).

[49] J. Fischer *et al.*, *Phys. Lett.* **B73**, 364 (1978).

[50] J.S. Frank *et al.*, *Phys. Rev.* **D28**, 423 (1983).

[51] A.G. Zephat *et al.*, *J. Phys.* **G13**, 1375 (1987).

[52] C. Niebuhr *et al.*, *Phys. Rev.* **D40**, 2976 (1989).

[53] D.P. Majumdar and J. Smith, *Phys. Rev.* **187**, 2039 (1969).

[54] A.J. Malensek, Technical Report FN-341, Fermilab, October 1981.

[55] J.R. Patterson, PhD thesis, Univerisity of Chicago, December 1990.

[56] R. Briere, PhD thesis, Univerisity of Chicago, 1994.

[57] L.K. Gibbons, PhD thesis, Univerisity of Chicago, August 1993.

[58] B. Schwingenheuer, PhD thesis, Univerisity of Chicago, 1994.

[59] D.A. Harris *et al.*, *Phys. Rev. Lett.* **71**, 3914 (1993).

[60] V. Papadimitriou, PhD thesis, Univerisity of Chicago, December 1990.

[61] H. Sanders *et al.*, *IEEE Trans. Nucl. Sci.* **36**, 358 (1989).

[62] D. Roberts, PhD thesis, Univerisity of California, Los Angeles, November 1993.

[63] B. Winstein. Private Communication.

[64] F. Sauli. Principles of Operation of Multiwire Proportional and Drift Chambers. CERN 77-09, 1977. Lecutres given 1975-1976.

[65] B. Jean-Marie *et al.*, *Nucl. Inst. Meth.* **159**, 213 (1979).

[66] G.W. Fraser and E. Mathieson, *Nucl. Inst. Meth.* **A257**, 339 (1987).

[67] N. Solomey. Private Communication.

[68] J.D. Jackson, *Classical Electrodynamics.* John Wiley & Sons, 1975.

[69] W.R. Nelson *et al.* The EGS4 Code System. Technical Report SLAC-265, SLAC, December 1985.

[70] R. Ford and W. Nelson. Technical Report UC-32, SLAC, June 1978.

[71] M. Weaver, PhD thesis, Univerisity of California, Los Angeles, 1994.

[72] E. Barberio *et al.*, *Comput. Phys. Commun.* **66**, 115 (1991).

[73] H.A. Bethe, *Phys. Rev.* **89**, 1256 (1953).

[74] R. Tschirhart. Private Communication.

[75] K.S. McFarland *et al.*, *Phys. Rev. Lett.* **71**, 31 (1993).

[76] A. Deshpande *et al.*, *Phys. Rev. Lett.* **71**, 27 (1993).

[77] A. Bramon Ll. Ametller and E. Massó, *Phys. Rev.* **D30**, 251 (1984).

[78] P. Herczeg, *Phys. Rev.* **D16**, 712 (1977).

[79] S. Davidson *et al.* CFPA-93-TH-29, 1993.