Abstract

The interference between the long-lived and short-lived neutral kaons (K_L and K_S) provides an arena for very sensitive tests of CPT symmetry. An observation of CPT violation would have a significant impact on physics, as it has been shown that CPT is an exact symmetry in almost all local quantum field theories.

In a fixed target experiment at Fermilab, the E731 collaboration has collected large samples of K_L and K_S decays into the $\pi^+\pi^-$ and $\pi^0\pi^0$ final states. This experiment uses a regenerator technique, in which one of two K_L beams passes through a block of material in order to provide a K_S component that is fully coherent with the K_L state. By measuring the interference between the K_L and K_S decays, the phases ϕ_{+-} and ϕ_{00} of the two amplitude ratios $\mathrm{amp}(K_L \to \pi^+\pi^-)/\mathrm{amp}(K_S \to \pi^+\pi^-)$ and $\mathrm{amp}(K_L \to \pi^0\pi^0)/\mathrm{amp}(K_S \to \pi^0\pi^0)$ can be determined. CPT symmetry predicts that the difference between ϕ_{00} and ϕ_{+-} should be much less than one degree and that their average should be close to 43.7 degrees.

Using data from a 1987–88 run, we find that $\phi_{00} - \phi_{+-} = -0.3° \pm 2.4°(\mathrm{stat}) \pm 1.2°(\mathrm{syst})$. The values obtained for ϕ_{+-} and ϕ_{00} — $47.7° \pm 2.0°(\mathrm{stat})$ and $47.4° \pm 1.4°(\mathrm{stat})$ — have larger systematic errors, mostly due to the limited knowledge about the phase associated with the K_S regeneration. (This "regeneration phase" cancels in the phase difference.) Using the world average value of ϕ_{+-}, we find that $2\phi_{+-}/3 + \phi_{00}/3 = 44.5° \pm 1.5°$. The results are in agreement with the predictions of CPT symmetry.

FERMILAB
LIBRARY

Table of Contents

Abstract i

List of Tables vi

List of Figures vii

Acknowledgements xiii

1 INTRODUCTION 1

1.1 Neutral Kaons 3

1.2 The CPT Predictions 4

1.3 Where's the Phase? 5

1.4 The Experiment 10

1.5 The Data 12

2 PHENOMENOLOGY 15

2.1 Mixing of Neutral Kaons 15

2.2 Decays of Neutral Kaons 17

2.3 The Bell-Steinberger Relation 20

2.4 Predictions of CPT Symmetry 22

2.5 Regeneration 22

3 APPARATUS 25

3.1 Beamline 25

3.2 Regenerator 28

3.3 Trigger Hodoscopes 29

3.4 Spectrometer 32

3.5 Calorimeter 33

3.6 Background Rejection 36

4 DATA COLLECTION 39

4.1 Beam 39

4.2 Rates 41

4.3 Triggers 44

4.3.1 Charged Trigger 44

4.3.2 Neutral Trigger 46

4.3.3 Other Triggers 48

4.3.4 Trigger Timing 49

4.4 Data Acquisition 49

4.5 Monitoring 51

5 CALIBRATION 53

5.1 Coordinate System 53

5.2 Analysis Magnet 54

5.3 Drift Chambers 54

5.3.1 Time to Distance Conversion 55

5.3.2 Resolution 58

5.3.3 Track Finding 60

5.3.4 Track Matching 61

5.3.5 Momentum Determination 62

5.4 Calorimeter 62

5.4.1 Cluster Finding 64

5.4.2 Cluster Energy and Position 67

5.4.3 Energy Calibration 68

5.5 Photon Vetoes 69

6 DATA REDUCTION **75**

6.1 Charged Mode Analysis 76

6.1.1 Reconstruction of Two-Track Events 77

6.1.2 Identification of the $\pi^+\pi^-$ Signal 80

6.1.3 Miscellaneous Cuts 87

6.1.4 Final Data Sample and Background Subtraction 89

6.2 Neutral Mode Analysis 99

6.2.1 Reconstruction of Four-Photon Events 100

6.2.2 Kinematic Cuts 101

6.2.3 Other Cuts 104

6.2.4 Final Data Sample and Background Subtraction 109

7 DETECTOR SIMULATION **117**

7.1 Event Generation 117

7.2 Comparison with Data 120

8 EXTRACTION OF THE PHASES **141**

8.1 Predicted Decay Rates 141

8.2 Fitting Method 143

8.3 The Result 147

8.4 Systematics 150

9 CONCLUSIONS **173**

Bibliography **177**

List of Tables

4.1 Rates. 43

4.2 Rates for the charged trigger. 45

4.3 Rates for the neutral trigger. 47

8.1 Fixed parameters (see text). 142

8.2 Result from the fit described in the text. The errors are one standard deviation statistical errors. 147

8.3 Systematic errors for ϕ_{+-}, ϕ_{00} and $\Delta\phi$. 150

List of Figures

1.1 Wu-Yang diagram. 4

1.2 Interference in the $\pi\pi$ decay rate using target K_S. 7

1.3 Interference in the $\pi\pi$ decay rate using regenerated K_S. 8

1.4 The E731 Basic Design. 10

1.5 Decay vertex of $\pi^0\pi^0$ events in the two beams. 12

3.1 The E731 decay volume and detector. 27

3.2 Schematic view of the regenerator. 28

3.3 Beam's view of the B (*top*) and C (*bottom*) hodoscopes. 31

3.4 Configuration of lead glass blocks in the calorimeter. 34

4.1 Event distribution within the spill, summed over G set data. 40

4.2 Beam intensity as measured by the muon hodoscope. 42

5.1 A track passing through two complementary planes. 55

5.2 Drift time distribution. 56

5.3 Inefficiency versus distance from the wire. 57

5.4 Sum of drift distances for one plane pair. 58

5.5 Average track residual (*top*) and rms (*bottom*) versus position in the cell. 59

5.6 Typical $\pi^0\pi^0$ event in the lead glass. 63

5.7 Typical $\pi^+\pi^-$ event in the lead glass. 65

5.8 Event from Figure 5.6 seen by the Hardware Cluster Finder. 66

5.9 Energy of photons missing the calorimeter. (Monte Carlo) 71

5.10 Ratio of reconstructed and generated energies. (Monte Carlo) . . . 72

5.11 Invariant photon-photon mass. 72

5.12 Observed energy versus the predicted energy. 73

6.1 Number of drift chamber hits for charged triggers. 76

6.2 Track momentum for two-track events (two entries per event). . . . 77

6.3 Transverse momentum squared for all two-track events on one raw data tape (G set data). 79

6.4 Invariant $\pi^+\pi^-$ mass for all two-track events on one raw data tape. Regenerated beam (*top*) and vacuum beam (*bottom*). 81

6.5 Energy from the calorimeter divided by the track momentum. One entry for each track that matches a cluster. 82

6.6 Geometric track cut at the lead glass for two-track events. 83

6.7 Invariant two-track mass for the lambda hypothesis. 85

6.8 Invariant $\pi^+\pi^-$ mass for two-track events (one raw data tape) after the $\pi e\nu$, $\pi\mu\nu$ and and Λ cuts. Regenerated beam (*top*) and vacuum beam (*bottom*). 86

6.9 Track quality cuts. 88

6.10 Track projection ($x^2 + y^2$) at the vacuum window (z=159 m). 89

6.11 Track projection (x and y) at the T and V hodoscopes (z=138 m). . 90

6.12 Track projection (x and y) at the analysis magnet (z=169 m). . . . 91

6.13 Track projection (x and y) at the muon hodoscope (z=190 m). . . . 92

6.14 Track x projection at $z = 162$ m. 93

6.15 Time distribution of the RA trigger signal. 93

8.3 $\pi^+\pi^-$ vertex distribution (regenerated beam, 30–40 GeV). 154

8.4 $\pi^+\pi^-$ vertex distribution (regenerated beam, 40–50 GeV). 155

8.5 $\pi^+\pi^-$ vertex distribution (regenerated beam, 50–60 GeV). 156

8.6 $\pi^+\pi^-$ vertex distribution (regenerated beam, 60–70 GeV). 157

8.7 $\pi^+\pi^-$ vertex distribution (regenerated beam, 70–80 GeV). 158

8.8 $\pi^+\pi^-$ vertex distribution (regenerated beam, 80–90 GeV). 159

8.9 $\pi^+\pi^-$ vertex distribution (regenerated beam, 90–100 GeV). 160

8.10 $\pi^+\pi^-$ vertex distribution (regenerated beam, 100–110 GeV). 161

8.11 $\pi^+\pi^-$ vertex distribution (regenerated beam, 110–120 GeV). 162

8.12 $\pi^+\pi^-$ vertex distribution (regenerated beam, 120–130 GeV). 163

8.13 $\pi^0\pi^0$ vertex distribution (regenerated beam, 40–50 GeV). 164

8.14 $\pi^0\pi^0$ vertex distribution (regenerated beam, 50–60 GeV). 165

8.15 $\pi^0\pi^0$ vertex distribution (regenerated beam, 60–70 GeV). 166

8.16 $\pi^0\pi^0$ vertex distribution (regenerated beam, 70–80 GeV). 167

8.17 $\pi^0\pi^0$ vertex distribution (regenerated beam, 80–90 GeV). 168

8.18 $\pi^0\pi^0$ vertex distribution (regenerated beam, 90–100 GeV). 169

8.19 $\pi^0\pi^0$ vertex distribution (regenerated beam, 100–110 GeV). 170

8.20 $\pi^0\pi^0$ vertex distribution (regenerated beam, 110–120 GeV). 171

Acknowledgements

I would first like to thank George Gollin for being my thesis advisor during these last four years, for introducing me to the odd (or is it even?) world of CPT symmetry, and for guiding this work to its completion. George is a good advisor, and I have learned a lot from him.

Research in high energy physics requires the combined effort of a large group of people. I enjoyed the company of my fellow graduate students — Lawrence Gibbons, Vaia Papadimitriou and Ritchie Patterson — and I wish them success in their future careers. I have learned much from Bob Hsiung, Yau Wah, Bruce Winstein, Hitoshi Yamamoto and Taku Yamanaka. It was a pleasure to meet and to work with Tony Barker, Grahame Blair, Greg Bock, Roy Briere, Jean-Claude Brisson, Rick Coleman, Pascal Debu, Janice Enagonio, Gayle Grazer, Greg Makoff, Jeff Okamitsu, Bernard Peyaud, Sunil Somalwar, Ken Stanfield, Ray Stefanski, Earl Swallow, Bob Tschirhart, Rene Turlay, Bertrand Vallage, Roland Winston and Mike Woods. I look forward to meeting you again.

Six years ago, I didn't know the difference between hodoscopes and kaleidoscopes. Ten years ago, I couldn't tell a torque from a quark, and twenty-nine years ago, I didn't know how to hold a spoon. I still drop my spoon sometimes, but now they tell me that I'm a trained physicist. This work is dedicated to the people who made it possible: my parents Rolf and Marianne and all my teachers.

Chapter 1
INTRODUCTION

Symmetries of Nature play a central role in physics. We find them to be very useful tools when constructing models of our universe, and we also appreciate the elegance and simplicity of a symmetric theory. On the other hand, any deviation from exact symmetry is always very exciting, since it usually leads us towards new phenomena and sometimes even a new interaction. In fact, physics would be rather dull if there were no strange effects to investigate. With recent experiments in high energy physics agreeing so well with the standard model, would it not be interesting if some boring old symmetry could be replaced by a new exciting *asymmetry*?

There are three symmetry operations in particular that have been of much interest in the study of elementary particles and their interactions. These are time reversal (T), parity reversal (P) and charge conjugation (C). Initially, these were all thought to be exact, but then it was discovered in 1957 that the weak interaction violates both P and C.[1] Since the C and P violations are closely related, the combined operation of charge conjugation and parity reversal (CP) remained a good symmetry. For example, consider the $\pi^+ \rightarrow \mu^+\nu_L$ decay, where the neutrino is always left-handed. Parity reversal changes the direction of momentum vectors, but leaves spin vectors unchanged. Consequently, a reflection of the decay in the P mirror would lead to a *right*-handed neutrino, which is not allowed. A reflection in the C mirror, on the other hand, results in the decay $\pi^- \rightarrow \mu^-\bar{\nu}_L$ with a left-handed *anti*-neutrino, which is not allowed either. Only the combined reflection CP results in the allowed decay $\pi^- \rightarrow \mu^-\bar{\nu}_R$, with a *right*-handed *anti*-neutrino. In the standard model, the weak interaction couples only to left-handed neutrinos and right-handed anti-neutrinos, thus explaining the "maximal" violation of the P and C symmetries.

It did not take very long, however, until a crack in the CP mirror appeared.

If CP symmetry were exact, the long-lived neutral kaons would never decay into two pions, but in 1964 that particular decay mode was in fact observed.[2] Unlike the P and C violations, the deviation from perfect CP symmetry is a very small effect (10^{-3} level), and its origin remains uncertain. In the standard model, the effect is assumed to arise from a non-zero phase in the weak interaction quark-quark couplings (Kobayashi-Maskawa matrix).[3] It has also been proposed that the effect is due to an entirely new interaction.[4]

Meanwhile, the stakes have been raised by the CPT theorem[5–7], which shows that the combination of all three operations is an exact symmetry in almost all local quantum field theories. This result does not imply that CPT violations are impossible, only that the consequences of a deviation from exact symmetry would be more far-reaching. Any experimentalist who is thirsting for "new physics" should certainly continue to test the predictions of CPT symmetry. One such prediction is the equality between particle-antiparticle masses and lifetimes, and direct measurements of mass and lifetime differences ($\delta M/M$ and $\delta\Gamma/\Gamma$) typically provide limits at the 10^{-3}–10^{-4} level. Another prediction is that the magnitudes of the magnetic moments of particle and anti-particle should be equal. A comparison of the "g minus 2" measurements for the electron and the positron gives a limit at the 10^{-11} level. For a compilation of CPT tests, see page 46 of reference 8. Since the mass difference between K^0 and $\overline{K^0}$ is limited by the difference (Δm) between K_L and K_S, and with $\Delta m/m_K = 7 \times 10^{-15}$, the kaons are expected to set a particularly strong limit. Recalling that CP violation has only been observed in the neutral kaon system, one might even argue that this is the best place to look for a CPT violation.

So here you are, reading the introduction to a recent experimental test of CPT symmetry in the neutral kaon system. The remainder of this first chapter will tell you about the basic ideas of the experiment, while the following chapters describe the various aspects — apparatus, analysis, etc. — in more detail.

1.1 Neutral Kaons

Excellent reviews of the neutral kaon system together with the CP and CPT phenomenology can be found in the literature.[9–14] Here, a very brief outline will be presented in order to introduce the predictions of CPT symmetry that can be tested experimentally. For simplicity, all the relations given in this introduction assume CPT invariance. See chapter 2 for a somewhat more detailed review of the algebra.

Because the strong interaction conserves strangeness (S), the K^0 ($S = +1$) and $\overline{K^0}$ ($S = -1$) are the neutral kaon eigenstates during production. The kaons decay thanks to the weak interaction which allows $\Delta S = \pm 1$ transitions. They also mix with each other through intermediate states with $S = 0$ (*e.g.*, 2π and 3π). Taking into account the small ($\epsilon \simeq 2 \times 10^{-3}$) CP violation in the K^0-$\overline{K^0}$ mixing, the vacuum eigenstates are written

$$K_S = [(1+\epsilon)K^0 + (1-\epsilon)\overline{K^0}]/\sqrt{2(1+|\epsilon|^2)}$$

$$K_L = [(1+\epsilon)K^0 - (1-\epsilon)\overline{K^0}]/\sqrt{2(1+|\epsilon|^2)}$$

The K_S state is mostly CP-even, allowing it to decay into the CP-even 2π final states. In fact, this is the reason why K_S is the short-lived state ($\tau_S \simeq 0.89 \times 10^{-10}$ sec). The long-lived K_L state ($\tau_L \simeq 5.2 \times 10^{-8}$ sec) is mostly CP-odd, but the small CP-even impurity introduced by ϵ permits it to decay occasionally into the 2π states. This "indirect" CP violation was first observed in 1964.[2]

There are two different 2π final states ($\pi^+\pi^-$ and $\pi^0\pi^0$), and we define two complex parameters η_{+-} and η_{00} as the ratios of the K_L and K_S decay amplitudes:

$$\eta \equiv \mathrm{amp}(K_L \to \pi\pi)/\mathrm{amp}(K_S \to \pi\pi)$$

Neglecting the small ($\simeq 5\%$) violation of the $\Delta I = 1/2$ rule, we have the two relationships

$$\eta_{+-} \equiv |\eta_{+-}|e^{i\phi_{+-}} = \epsilon + \epsilon'$$

$$\eta_{00} \equiv |\eta_{00}|e^{i\phi_{00}} = \epsilon - 2\epsilon'$$

which are often visualized in the Wu-Yang diagram (Figure 1.1). The param-

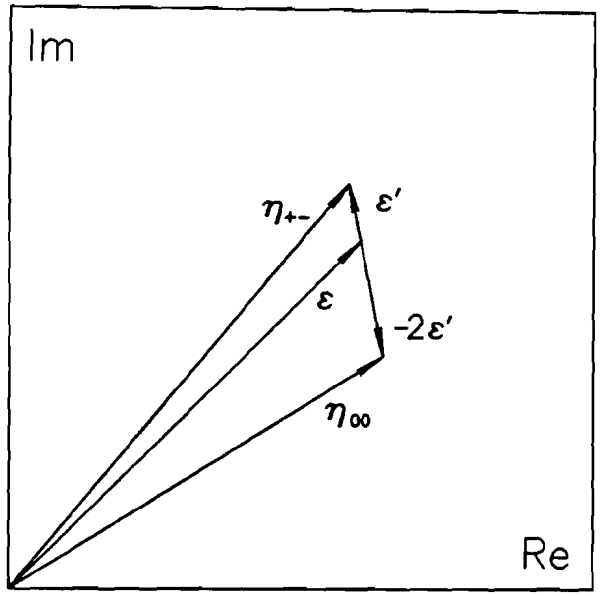

Figure 1.1: Wu-Yang diagram.

eter ϵ' represents any additional CP violation arising directly from the 2π decay amplitudes (known as "direct" CP violation). By measuring the ratio of $|\eta_{+-}|$ and $|\eta_{00}|$, the quantity $\mathrm{Re}(\epsilon'/\epsilon)$ can be determined through the relationship $|\eta_{+-}|^2/|\eta_{00}|^2 \simeq 1 + 6\,\mathrm{Re}(\epsilon'/\epsilon)$. The most recent result[15] — which was obtained from the same data set that is discussed in this thesis — is consistent with zero: $\mathrm{Re}(\epsilon'/\epsilon) = -0.0004 \pm 0.0014(\mathrm{stat}) \pm 0.0006(\mathrm{syst})$. Another recent measurement[16] by the NA31 collaboration at CERN shows a three sigma deviation from zero: $\mathrm{Re}(\epsilon'/\epsilon) = 0.0033 \pm 0.0007(\mathrm{stat}) \pm 0.0008(\mathrm{syst})$. Since it is also known that the argument of ϵ' is close to the argument of ϵ, it follows that $|\epsilon'|$ is much smaller than $|\epsilon|$. The magnitude and relative angle of ϵ' are greatly exaggerated in Figure 1.1 for clarity.

1.2 The CPT Predictions

CPT symmetry results in predictions for the phases of ϵ and ϵ', which can be tested by measuring the phases of η_{+-} and η_{00}. It should be stressed that these phases are meaningful quantities only because certain other phases have been fixed

by convention. The CPT requirements can, of course, be formulated in a phase-independent way. The expression for ϵ is

$$\epsilon = \frac{1}{2}\frac{\langle\overline{K^0}|H|K^0\rangle - \langle K^0|H|\overline{K^0}\rangle}{(m_L - m_S) + i(\Gamma_S - \Gamma_L)/2}$$

Here, $H = M - i\Gamma/2$ is the Hamiltonian of the K^0-$\overline{K^0}$ system, and m_L, m_S, Γ_L and Γ_S are the K_L and K_S masses and decay rates. The mass and decay matrices (M and Γ) are both hermitean, making the numerator equal to $2i\,\mathrm{Im}\,M_{12} + \mathrm{Im}\,\Gamma_{12}$. Using results from the 3π and semileptonic decay modes, one can show that the contribution to ϵ from the decay matrix is very small compared to the contribution from the mass matrix. Consequently, the phase of ϵ — which is measured as $2\phi_{+-}/3+\phi_{00}/3$ — should be close (within a degree) to the argument of $\Gamma_S/2+i\Delta m$. This "natural angle" is equal to $43.7 \pm 0.2°$.

The other prediction is obtained by considering the phase of ϵ'. This parameter is written as

$$\epsilon' = \frac{i}{\sqrt{2}}\frac{\mathrm{Im}\,A_2}{\mathrm{Re}\,A_0}e^{i(\delta_2-\delta_0)}$$

where A_I is the decay amplitude of K^0 into a 2π final state of isospin I with a phase shift δ_I from final state interactions. Using the experimental value[17] $\delta_2 - \delta_0 = -45 \pm 10°$, we see that ϵ' is nearly parallel with ϵ. Since $\mathrm{Re}(\epsilon'/\epsilon)$ is known to be very small, it then follows that the difference between ϕ_{+-} and ϕ_{00} should be much less than one degree. However, the world average values[8] $\phi_{+-} = 44.6 \pm 1.2°$ and $\phi_{00} = 54 \pm 5°$ violate both predictions at about two standard deviations. Considering the importance of CPT symmetry, it is desirable to make more precise measurements of the two phases, particularly ϕ_{00}.

1.3 Where's the Phase?

Thanks to the remarkably small mass difference, $\Delta m \simeq 3.5 \times 10^{-6}$ eV, between the K_L and K_S states, it is possible to observe interference between the $K_S \to \pi\pi$ and $K_L \to \pi\pi$ decay amplitudes. In fact, with $\Delta m \simeq 0.5$ in units of K_S lifetimes, the mass difference is very well adjusted for an interference experiment. Recalling the definition of η as the ratio between the K_L and K_S amplitudes, and assuming

that there is at least some coherence between the K_S and K_L components, it follows that the phase of η becomes observable in the $\pi\pi$ decay rate via the interference term. The amount of coherence depends on whether the K_S component comes directly from the kaon production target or is being regenerated from a K_L state. This experiment uses K_S from a regenerator, but let us first consider the interference between K_S and K_L from the target.

Kaons are produced in a target as K^0 or $\overline{K^0}$, and with $K^0 \sim K_S + K_L$ and $\overline{K^0} \sim K_S - K_L$, the $\pi\pi$ decay rate as a function of proper time t from the target is proportional to

$$e^{-\Gamma_S t} + |\eta|^2 e^{-\Gamma_L t} + 2D|\eta| e^{-(\Gamma_S+\Gamma_L)t/2} \cos(\Delta m\, t - \phi_\eta)$$

where $\Gamma_{S,L}$ are the $K_{S,L}$ decay rates and ϕ_η stands for either ϕ_{+-} or ϕ_{00}. The "dilution factor" $D = (N - \overline{N})/(N + \overline{N})$, where N and $\overline{N}$ are the K^0 and $\overline{K^0}$ fluxes, enters because the interference terms from K^0 and $\overline{K^0}$ have opposite signs. In other words, the amount of coherence between K_S and K_L is determined by $|D|$. The decay rate is plotted in Figure 1.2 for the case of a K^0 at the target (*i.e.*, $D = +1$). The pure K_S and K_L terms and their sum are indicated by dashed lines. The factor D is always positive — if the target is made of ordinary matter — since it is easier to produce a K^0 than a $\overline{K^0}$. This production asymmetry increases with the ratio of kaon energy to incident proton energy and also with the targeting angle. The NA31 experiment at CERN uses this method, and for them D varies from about 0.1 at low kaon energies to about 0.4 at higher energies.

In our experiment, we wait until most of the target K_S have decayed and let the resulting (almost pure) K_L beam pass through a regenerator. This recreates a K_S component which is fully coherent with the K_L. The regeneration phenomenon is discussed in more detail at the end of chapter 2, and it is sufficient to note for now that the kaon state emerging from the regenerator in the forward direction is described by $K_L + \rho K_S$, where ρ is called the coherent regeneration amplitude. The $\pi\pi$ decay rate as a function of proper time from the regenerator is proportional to

$$|\rho|^2 e^{-\Gamma_S t} + |\eta|^2 e^{-\Gamma_L t} + 2|\rho||\eta| e^{-(\Gamma_S+\Gamma_L)t/2} \cos(\Delta m\, t + \phi_\rho - \phi_\eta)$$

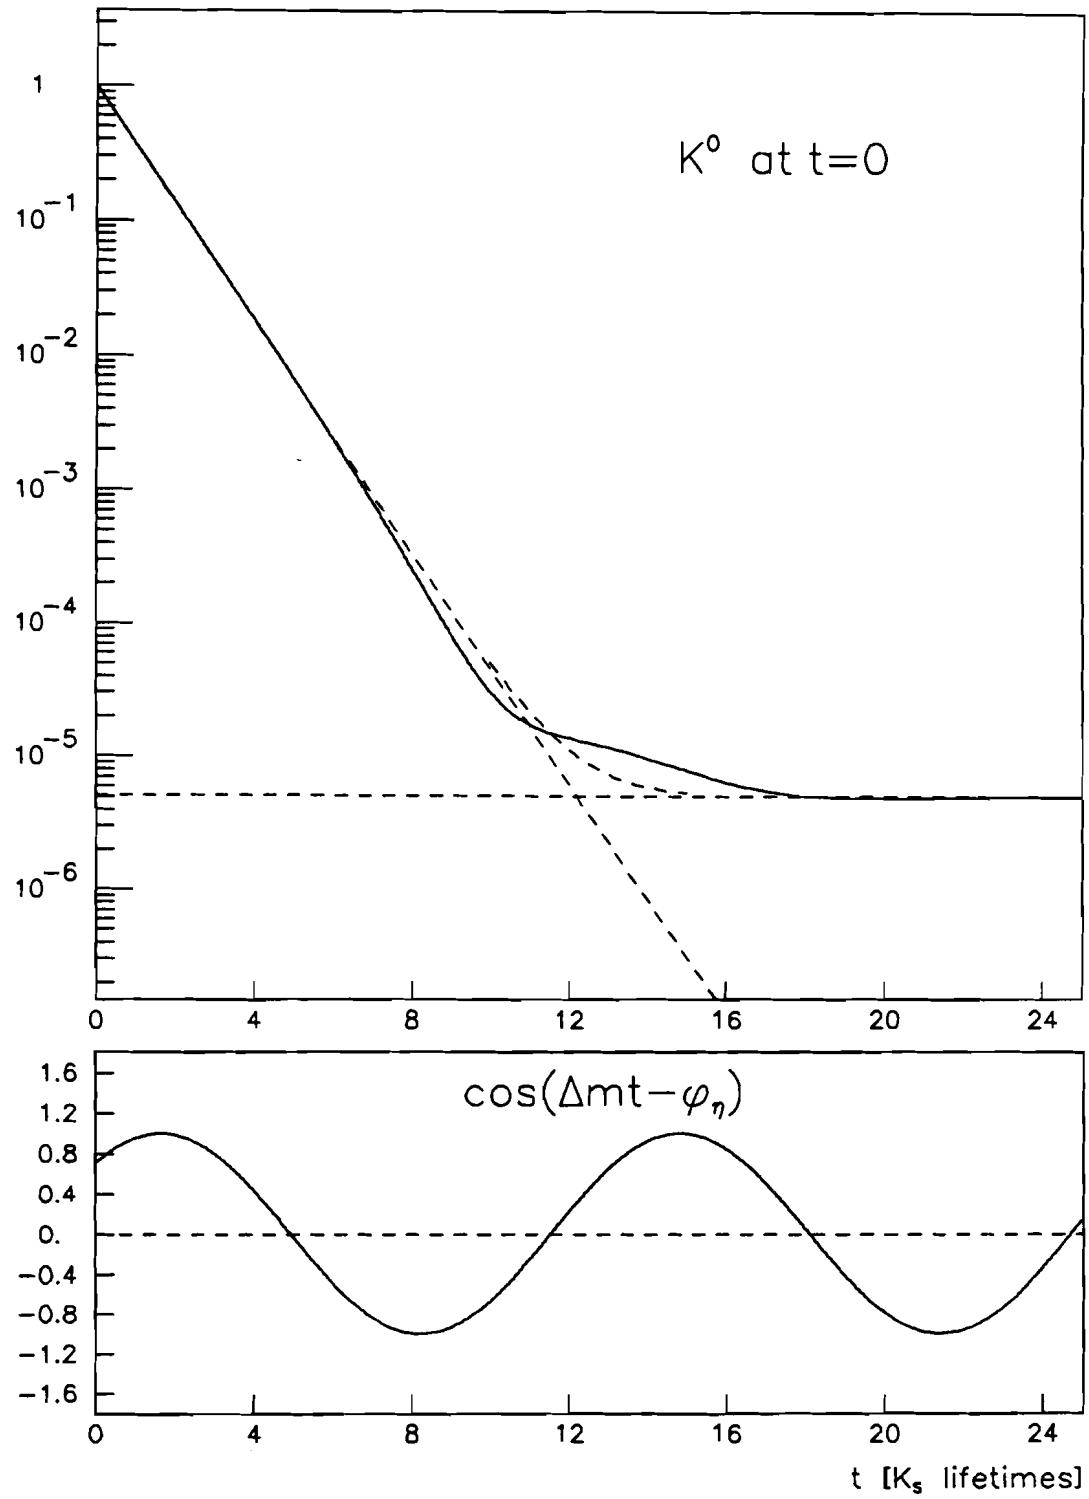

Figure 1.2: Interference in the $\pi\pi$ decay rate using target K_S.

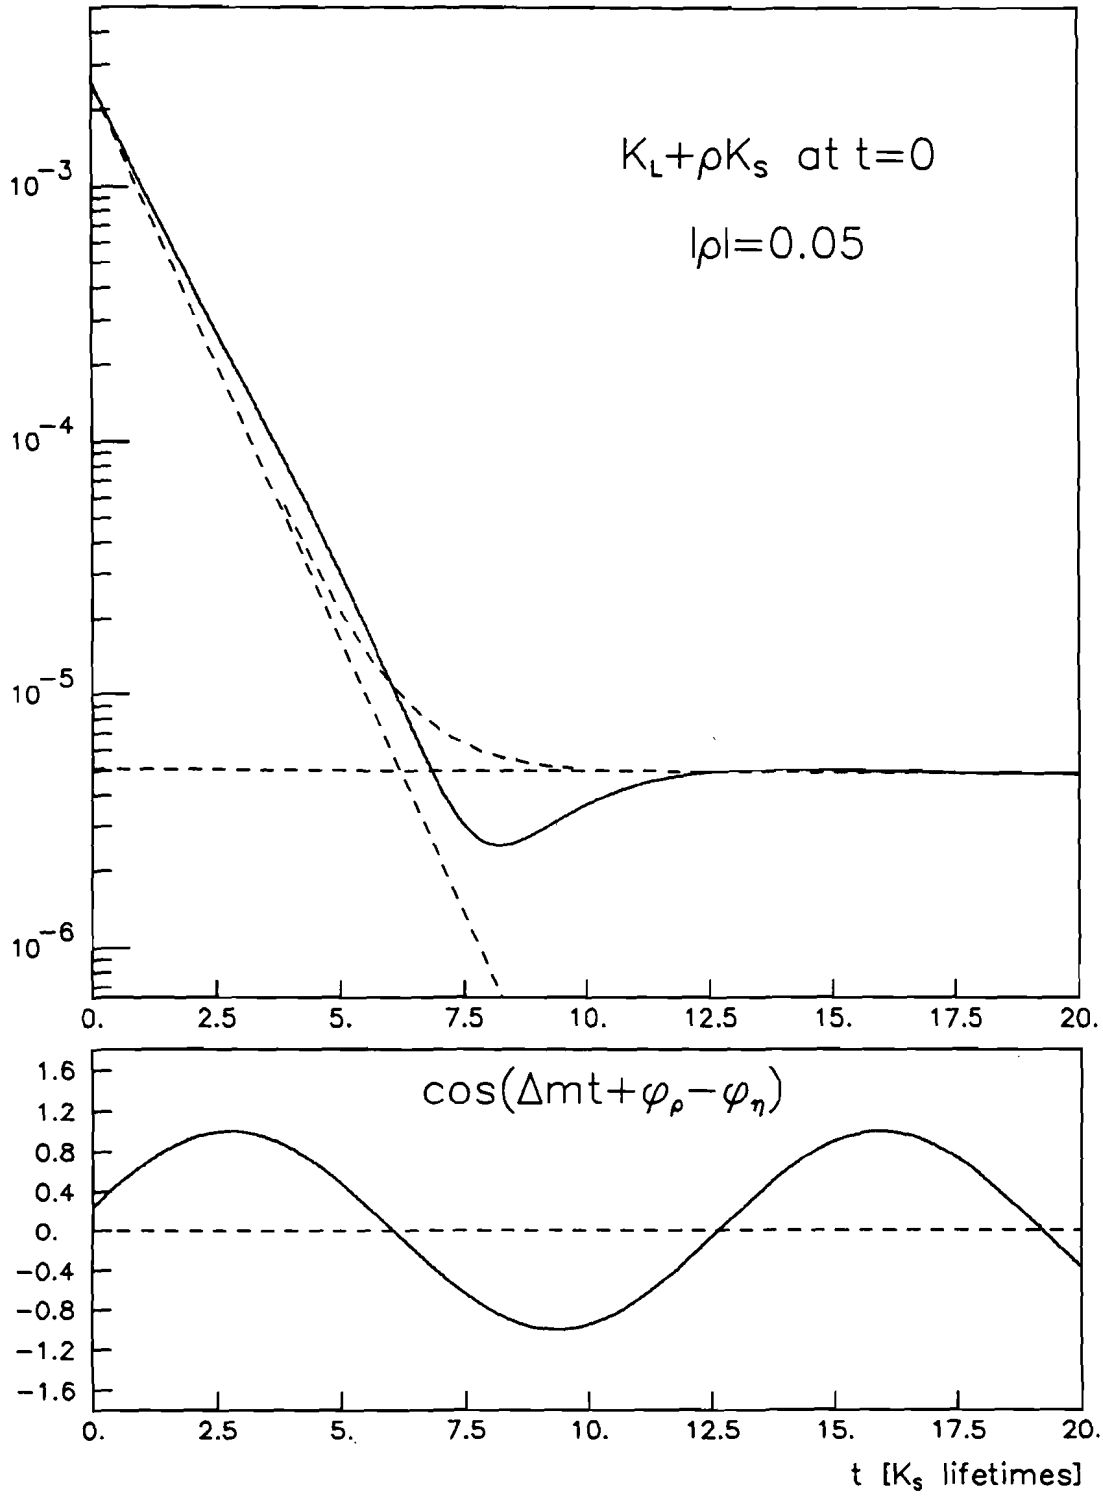

Figure 1.3: Interference in the $\pi\pi$ decay rate using regenerated K_S.

which is plotted in Figure 1.3. The magnitude of ρ depends on the thickness of the regenerator and on the kaon energy. In our case, $|\rho|$ varies from about 5% at a kaon energy of 40 GeV to about 2% at 160 GeV. The decay rate plots show just how fortunate we are with the size of Δm.

In order to compare the statistical power of the two methods for extracting ϕ_η, we need to look at the magnitude of the interference term ($\sim e^{-\Gamma_S t/2}$) relative the magnitudes of the pure K_S term ($\sim e^{-\Gamma_S t}$) and the pure K_L term ($\sim$ constant). For the proper times that we are interested in here, we can assume that $\Gamma_L \approx \Gamma_S/580$ is zero. Ignoring for a moment the cosine factor, let us consider the ratio of the interference term to the sum of the other two terms. This ratio has its maximum at the proper time t_0 defined as the time where the K_S and K_L terms are equal. In the case of K_S from the target, t_0 is determined by $e^{-\Gamma_S t_0} = |\eta|^2$, and we find after some simple algebra that the ratio is equal to D. In the case of K_S from a regenerator, t_0 is given by $|\rho|^2 e^{-\Gamma_S t_0} = |\eta|^2$ and we find that the same ratio is now equal to one. This shows explicitly how the factor D dilutes the sensitivity to the interference term for target K_S, while there is no such dilution for regenerated K_S because of the full coherence between K_S and K_L. (The two ratios can be compared directly with each other only for a given K_L flux, and it should be remembered that the regenerator attenuates this flux through absorption.)

For times larger than t_0, the statistical power quickly decreases as the K_L term stays constant and the interference term dies out. In the other direction, the K_S decays dominate the rate, but it is very important to note that the magnitude of the interference term is proportional to the square root of the K_S term, implying that the statistical power *per proper time interval* remains constant all the way to $t = 0$. However, since the total decay rate increases, the statistical power *per recorded event* decreases exponentially towards $t = 0$. For the target K_S method, it is best to define the start of the decay volume by a veto counter and avoid the huge number of K_S decays at the earliest proper times. (Such a veto counter will also eliminate junk from the target.) For the regenerator method, there is no problem with too many K_S, as $|\rho|^2$ is very small compared to one. The magnitude of ρ is essentially proportional to the length of the regenerator up to a certain limit. A

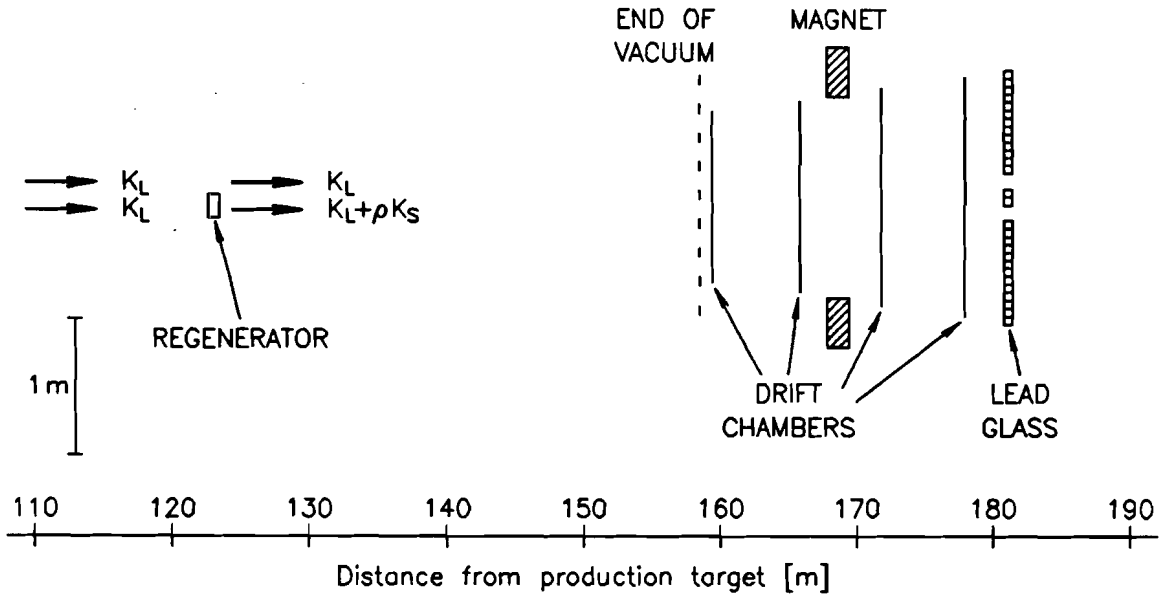

Figure 1.4: The E731 Basic Design.

thinner regenerator is better, not because $|\rho|$ is smaller, but because there is less absorption. However, if $|\rho|$ becomes too small (comparable too $|\eta|$), the interference term drowns in the K_L flux. The important point is that once the decay volume is defined and the data recorded, the statistical power is constant from the start of the decay volume until $t \approx t_0$. The effect of the cosine factor is to modulate the sensitivity to the phase. For example, when the the cosine changes slowly (near ± 1), there is less statistical sensitivity to ϕ_η.

1.4 The Experiment

A very schematic view of our experiment is shown in Figure 1.4. The beamline and the detector will be described in more detail in chapter 3, and only the major features are introduced here. Two neutral kaon beams — with an average kaon energy around 70 GeV — are produced at the target by the 800 GeV protons from the Fermilab Tevatron. By the time the kaons reach the regenerator, almost all the K_S have decayed. Being this far away from the target also means that most lambdas have decayed and that the accidental activity is quite low. After the regenerator we have a "regenerated" beam which is a coherent mixture of K_L and

K_S, and a "vacuum" beam which is pure K_L. It is in the regenerated beam that the phase of η can be measured from the interference between K_L and K_S. The vacuum beam has no such phase information, but it can be used to normalize the incident kaon flux. This normalization is possible because the events in the regenerated and vacuum beams are recorded simultaneously and are affected equally by dead-time and losses due to accidental activity. Changes in detector performance or in the electronics over a period of time also have the same effect on regenerated and vacuum events. However, since there are small differences between the top and bottom beams, it is necessary for the regenerator to alternate frequently between the two positions. For example, the incident top and bottom beams have slightly different intensities, momentum spectra, profiles and angular divergence, and there are most likely small differences in detector acceptance for events from the top and bottom beams. But because the regenerator alternates, there are no such differences between what we call the regenerated and vacuum beams, since they are averaged over the top and bottom positions. Having just one beam and moving the regenerator in and out would certainly avoid the differences arising between two separate beams, but the dead-time and accidental losses would be different for regenerated and vacuum events. Only the "Double Deluxe" recipe, with two beams and alternating regenerator, cancels virtually all systematic differences between events from the regenerated and vacuum beams.

When the kaons pass through the regenerator, there is both coherent and incoherent scattering. The coherently scattered kaons, which peak sharply in the forward direction, emerge with the same beam profile and angular divergence as the incident beam, while the incoherently scattered kaons (elastic and inelastic) emerge with finite transverse momentum and form a background under the forward peak. It is the *coherently* scattered kaons that are described by the expression $K_L + \rho K_S$, where ρ is the coherent regeneration amplitude and where the relative phase between K_S and K_L is simply given by the phase of ρ. The relative phase between the incoherently scattered K_S and K_L is not as simple, because of the more complex processes involved in the incoherent scattering. For the phase measurement, it is therefore necessary to extract the coherent "signal" by subtracting the incoherent

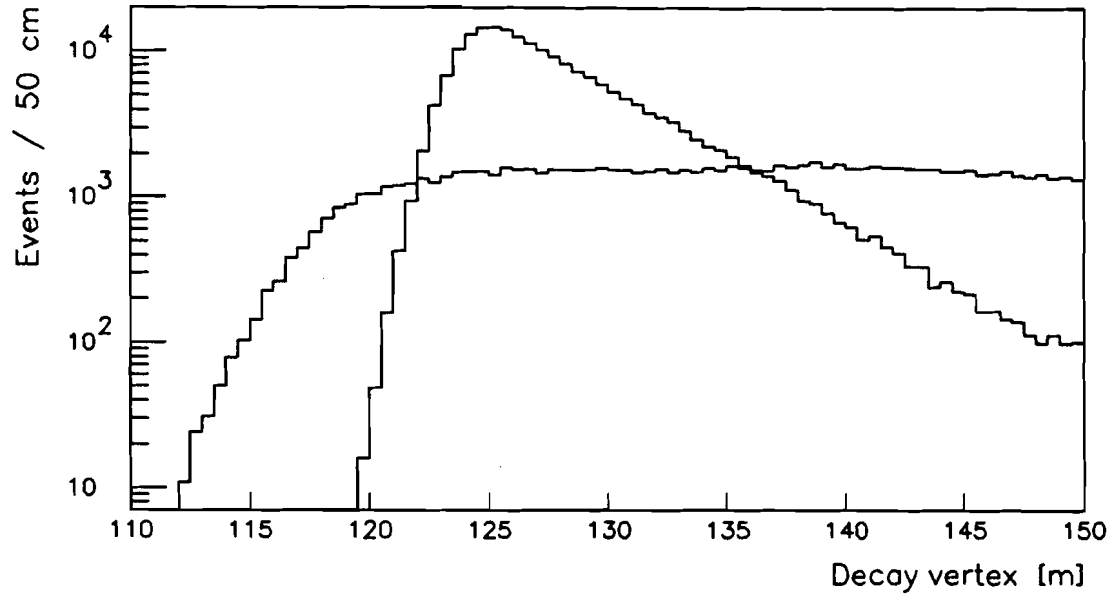

Figure 1.5: Decay vertex of $\pi^0\pi^0$ events in the two beams.

"background." By instrumenting the regenerator with scintillator and phototubes, we reject a good fraction of the inelastically scattered kaons at the trigger level.

1.5 The Data

The decay vertex distributions for reconstructed $\pi^0\pi^0$ events are shown in Figure 1.5. (The longitudinal vertex position is measured from the production target, which is located 123 meters upstream of the regenerator.) The flat K_L shape in the vacuum beam and the falling K_S shape in the regenerated beam are quite visible, even though the distributions have not been corrected for detector acceptance. (The decay length is 3.8 m for a 70 GeV K_S and about two kilometers for a K_L.) For the regenerated beam, there is an edge around 123 m due to the veto counters at the end of the regenerator. A little before the regenerator, there is a lead mask with two holes for the beams, which causes the acceptance for upstream decays in the vacuum beam to fall off quickly. Since the decay distributions are integrated over kaon energies from 40 GeV to 120 GeV, the interference effects towards the downstream end of the regenerated beam are not visible in this plot. Note that the kaon *flux* is mostly K_L even in the regenerated beam, and that this

flux is about 6% of the vacuum beam flux because of absorption in the regenerator and an upstream absorber. Only a small fraction ($|\rho|^2 \simeq 10^{-3}$) of the K_L are regenerated into K_S, but this is enough to make the $\pi\pi$ decay rates in the two beams comparable, since the $K_L \to \pi\pi$ rate is suppressed by $|\eta|^2 \simeq 5 \times 10^{-6}$. There are, of course, large numbers of K_L decays into the 3π states ($\pi^+\pi^-\pi^0$ and $\pi^0\pi^0\pi^0$) and the semileptonic states ($\pi^\pm e^\mp \nu$ and $\pi^\pm \mu^\mp \nu$), and these will have to be dealt with by the triggers and during the event reconstruction.

The two charged pions from the $\pi^+\pi^-$ decays are tracked and momentum analyzed with a spectrometer consisting of an analysis magnet and four drift chambers. The "charged" trigger is essentially a two-track trigger based on scintillator hodoscopes with a muon veto to reject the $\pi\mu\nu$ decays. The decay vertex is obtained directly from the tracking, and the invariant mass and the kaon energy are calculated by assuming the (charged) pion mass for both tracks. Both the $\pi^+\pi^-\pi^0$ and the $\pi e \nu$ are cut in the analysis, the former by the invariant mass and the latter by electron identification in the lead glass. The transverse kaon momentum is calculated and used for the subtraction of the incoherent kaons and the small remaining semileptonic background.

The four photons that result from a $\pi^0\pi^0$ decay are detected with a calorimeter consisting of an array of lead glass blocks. Each photon shower deposits energy in a "cluster" of blocks, and the number of such clusters is determined by a trigger processor. The "neutral" trigger relies on the total energy in the lead glass and on the number of clusters. From the energies and positions of the four photon clusters, it is possible (assuming the π^0 mass) to calculate the distance from the lead glass to the decay vertex and also the invariant kaon mass. The kaon energy is just the sum of the four photon energies. There is a background from $\pi^0\pi^0\pi^0$ decays with two photons missing the calorimeter or merging with other photons. This background is suppressed by using veto counters to catch escaping photons and by trying to distinguish the single-photon clusters from those that result when photons merge. For the subtraction of incoherently scattered kaons, we must use the center of energy at the lead glass instead of the transverse kaon momentum. The remaining $\pi^0\pi^0\pi^0$ background under the invariant mass peak is also subtracted.

The triggers and the data taking are described in more detail in chapter 4, and all calibration issues are discussed in chapter 5. The detailed event analysis is presented in chapter 6, together with the background subtractions. We do not calculate the proper time of the decay for each event. Instead we bin the data according to decay vertex and kaon energy in 1 m by 10 GeV bins. Corresponding Monte Carlo distributions are obtained in order to determine the acceptances. The simulation of the beam and the detector is described in chapter 7 together with some comparisons between the data and the Monte Carlo.

There are a couple of different ways to extract the phases of η_{+-} and η_{00} from the data distributions. One way would be to divide the number of regenerated and vacuum events in each vertex and energy bin, and extract the phases from a fit. This would eliminate most of the acceptance uncertainty, but the division would also result in a loss of statistical power. We choose instead to fit the data distributions directly against the Monte Carlo distributions. This gives better statistical precision, but on the other hand requires the acceptances to be well understood. The details of the fitting procedure can be found in chapter 8, together with a discussion of systematic uncertainties. The final results with conclusions are presented in chapter 9, and that's it!

Chapter 2
PHENOMENOLOGY

This chapter is intended to be a summary of the relevant neutral kaon phenomenology — which according to reference 11 "boils down to multiplication and addition of complex numbers." Most of the algebra can be found in the earlier work[18–20], but in order to get the most up-to-date information on the status of CP and CPT symmetry, one should start by consulting the more recent reviews[9–14]. The notation used in this chapter is a mix of that in references 10–12.

2.1 Mixing of Neutral Kaons

Because of the mixing introduced by the weak interaction, the neutral kaon eigenstates in vacuum will be linear combinations of the strong interaction eigenstates K^0 and $\overline{K^0}$. In the absence of CP violation, and with the convention $\mathrm{CP}|K^0\rangle = |\overline{K^0}\rangle$, the vacuum eigenstates would be the CP-even state $(K^0+\overline{K^0})/\sqrt{2}$ and the CP-odd state $(K^0-\overline{K^0})/\sqrt{2}$. Not assuming anything, we can write the observed eigenstates, K_S and K_L, as perturbations of the odd and even states: $K_S \propto |\mathrm{even}\rangle + \epsilon_S|\mathrm{odd}\rangle$ and $K_L \propto |\mathrm{odd}\rangle + \epsilon_L|\mathrm{even}\rangle$. Including normalization, we write

$$K_S = [(1+\epsilon_S)K^0 + (1-\epsilon_S)\overline{K^0}]/\sqrt{2(1+|\epsilon_S|^2)}$$
$$K_L = [(1+\epsilon_L)K^0 - (1-\epsilon_L)\overline{K^0}]/\sqrt{2(1+|\epsilon_L|^2)}$$

Using the Wigner-Weisskopf approximation[21], which is valid for times larger than $\hbar/m_K \simeq 10^{-24}$ sec, the time evolution of a neutral kaon is described by $dA/dt = -iHA$, where $A = (a, \bar{a})$ is the time dependent amplitude vector with K^0 and $\overline{K^0}$ components. H is the Hamiltonian of the K^0-$\overline{K^0}$ system and can be written in terms of the mass and decay matrices as $H = M - i\Gamma/2$. (We generally

work in the rest frame of the kaon, so that the elements of M and Γ can be associated with invariant masses and decay rates.) Since the kaon has zero spin, H can be thought of as a two-by-two matrix operating on the two-component vector A. The eigenstates K_S and K_L have corresponding eigenvalues $\lambda_{S,L} = m_{S,L} - i\Gamma_{S,L}/2$ where $m_{S,L}$ and $\Gamma_{S,L}$ are the $K_{S,L}$ masses and decay rates. The four complex numbers ϵ_S, ϵ_L, λ_S and λ_L are determined by the equations $H|K_{S,L}\rangle = \lambda_{S,L}|K_{S,L}\rangle$, and it turns out to be convenient to write

$$\epsilon_S = \epsilon + \Delta$$

$$\epsilon_L = \epsilon - \Delta$$

Assuming that the magnitudes of ϵ_S and ϵ_L are small, we find that (to first order in ϵ and Δ)

$$\lambda_S + \lambda_L = H_{11} + H_{22} = M_{11} + M_{22} - i(\Gamma_{11} + \Gamma_{22})/2$$

$$\lambda_S - \lambda_L = H_{12} + H_{21} = M_{12} + M_{21} - i(\Gamma_{12} + \Gamma_{21})/2$$

and

$$\epsilon = \frac{1}{2}\frac{\langle\overline{K^0}|H|K^0\rangle - \langle K^0|H|\overline{K^0}\rangle}{\lambda_L - \lambda_S} = \frac{1}{2}\frac{i(M_{12} - M_{21}) + (\Gamma_{12} - \Gamma_{21})/2}{(\Gamma_S - \Gamma_L)/2 - i(m_L - m_S)}$$

$$\Delta = \frac{1}{2}\frac{\langle\overline{K^0}|H|\overline{K^0}\rangle - \langle K^0|H|K^0\rangle}{\lambda_L - \lambda_S} = \frac{1}{2}\frac{i(M_{11} - M_{22}) + (\Gamma_{11} - \Gamma_{22})/2}{(\Gamma_S - \Gamma_L)/2 - i(m_L - m_S)}$$

CPT symmetry requires the K^0 and $\overline{K^0}$ masses and lifetimes to be equal, which means that Δ is a CPT-violating parameter.

It is useful to introduce the angle

$$\Phi = \arctan 2\Delta m/\Gamma_S = 43.7 \pm 0.2^\circ$$

where $\Delta m \equiv m_L - m_S$ and $\Delta\Gamma \equiv \Gamma_S - \Gamma_L \simeq \Gamma_S$. Decomposing ϵ and Δ into parallel and perpendicular components with respect to Φ and assuming that M and Γ are both hermitean, we see that

$$\epsilon_{\parallel} = \frac{-\mathrm{Im}\, M_{12}}{\sqrt{2}\,\Delta m}$$

$$\epsilon_{\perp} = \frac{\mathrm{Im}\,\Gamma_{12}}{\sqrt{2}\,\Gamma_S}$$

$$\Delta_\perp = \frac{M_{11} - M_{22}}{2\sqrt{2}\,\Delta m}$$

$$\Delta_\parallel = \frac{\Gamma_{11} - \Gamma_{22}}{2\sqrt{2}\,\Gamma_S}$$

and

$$m_S + m_L = M_{11} + M_{22}$$

$$\Gamma_S + \Gamma_L = \Gamma_{11} + \Gamma_{22}$$

$$m_S - m_L = 2\,\mathrm{Re}\,M_{12}$$

$$\Gamma_S - \Gamma_L = 2\,\mathrm{Re}\,\Gamma_{12}$$

The CPT-violating components of Δ are proportional to the K^0-$\overline{K^0}$ mass and decay rate differences. Since the K^0 and $\overline{K^0}$ are not vacuum eigenstates, their masses and lifetimes are not directly measurable quantities. The components of the CP-violating parameter ϵ depend on the relative phase between the K^0 and $\overline{K^0}$ states, which has not yet been explicitly fixed. This relative phase has been constrained implicitly, however, since the assumption that ϵ is small forces the arguments of M_{12} and Γ_{12} to be very close to zero. Finally, we observe that the eigenstates K_S and K_L are in general not orthogonal:

$$\langle K_S | K_L \rangle = \epsilon_S^* + \epsilon_L = 2\,\mathrm{Re}\,\epsilon - 2i\,\mathrm{Im}\,\Delta$$

This section has mostly been an exercise in linear algebra. In the next two sections, the mixing parameters ϵ and Δ will be related to observable quantities.

2.2 Decays of Neutral Kaons

CP violation has been observed in the 2π decay modes ($\pi^+\pi^-$ and $\pi^0\pi^0$) and in the semileptonic modes ($\pi e \nu$ and $\pi \mu \nu$). Since the 2π states are CP-even, the $K_L \to 2\pi$ amplitudes are suppressed compared to the $K_S \to 2\pi$ amplitudes, and we define

$$\eta_{+-} \equiv \frac{\mathrm{amp}(K_L \to \pi^+\pi^-)}{\mathrm{amp}(K_S \to \pi^+\pi^-)}$$

$$\eta_{00} \equiv \frac{\mathrm{amp}(K_L \to \pi^0\pi^0)}{\mathrm{amp}(K_S \to \pi^0\pi^0)}$$

both of which have been experimentally determined. It is useful to write the 2π final states in terms of states with definite isospin:

$$|\pi^+\pi^-\rangle = \sqrt{\frac{2}{3}}|2\pi, I{=}0\rangle + \sqrt{\frac{1}{3}}|2\pi, I{=}2\rangle$$

$$|\pi^0\pi^0\rangle = \sqrt{\frac{1}{3}}|2\pi, I{=}0\rangle - \sqrt{\frac{2}{3}}|2\pi, I{=}2\rangle$$

We can now define a few more amplitude ratios:

$$\epsilon_0 \equiv \frac{\mathrm{amp}(K_L \to 2\pi, I{=}0)}{\mathrm{amp}(K_S \to 2\pi, I{=}0)}$$

$$\epsilon_2 \equiv \frac{1}{\sqrt{2}}\frac{\mathrm{amp}(K_L \to 2\pi, I{=}2)}{\mathrm{amp}(K_S \to 2\pi, I{=}0)}$$

$$\omega \equiv \frac{1}{\sqrt{2}}\frac{\mathrm{amp}(K_S \to 2\pi, I{=}2)}{\mathrm{amp}(K_S \to 2\pi, I{=}0)}$$

The ϵ's are related to the η's by

$$\epsilon_0 = \frac{2}{3}\eta_{+-}(1+\omega) + \frac{1}{3}\eta_{00}(1-2\omega)$$

$$\epsilon_2 = \frac{1}{3}\eta_{+-}(1+\omega) - \frac{1}{3}\eta_{00}(1-2\omega)$$

or, the other way around,

$$\eta_{+-} = \frac{\epsilon_0 + \epsilon_2}{1+\omega}$$

$$\eta_{00} = \frac{\epsilon_0 - 2\epsilon_2}{1-2\omega}$$

Everything will be expressed in terms of the K^0 and $\overline{K^0}$ amplitudes

$$\mathrm{amp}(K^0 \to 2\pi, I) \equiv A_I e^{i\delta_I}$$

$$\mathrm{amp}(\overline{K^0} \to 2\pi, I) \equiv \overline{A}_I e^{i\delta_I}$$

where the final state phase shifts δ_I have been separated from the more fundamental weak interaction amplitudes A_I. We now find that

$$\epsilon_0 = \epsilon_L + r_0$$

$$\epsilon_2 = \epsilon' + \epsilon_L\omega$$

where

$$r_0 = \frac{A_0 - \overline{A}_0}{A_0 + \overline{A}_0}$$

$$\epsilon' = \frac{1}{\sqrt{2}} \frac{A_2 - \overline{A}_2}{A_0 + \overline{A}_0} e^{i(\delta_2 - \delta_0)}$$

and

$$\omega = \frac{1}{\sqrt{2}} \frac{A_2 + \overline{A}_2}{A_0 + \overline{A}_0} e^{i(\delta_2 - \delta_0)}$$

We now explicitly fix the relative phase between the K^0 and $\overline{K^0}$ states by requiring that the phases of A_0 and $\overline{A}_0$ are equal. As a result, r_0 is a real quantity, and since CPT symmetry requires that $\overline{A}_0 = A_0^*$, it follows that r_0 parametrizes CPT violation in the 2π (I=0) decay amplitudes. The parameter ω represents the small violation of the $\Delta I = 1/2$ rule. If one assumes that r_0 is zero, the relationships between the η's and the ϵ's can also be written

$$\eta_{+-} = \epsilon_0 + \frac{\epsilon'}{1+\omega}$$

$$\eta_{00} = \epsilon_0 - \frac{2\epsilon'}{1-2\omega}$$

$$\epsilon' = \frac{1}{3}(\eta_{+-} - \eta_{00})(1 - \omega - 2\omega^2)$$

The equation $\epsilon_0 = \epsilon_L + r_0$ relates $\epsilon - \Delta$ to the experimentally determined ϵ_0, but for $\mathrm{Re}(\epsilon - \Delta)$ there is a complication due to r_0.

We now turn to the 3π decay modes ($\pi^+\pi^-\pi^0$ and $\pi^0\pi^0\pi^0$). Since the 3π final states are mainly CP-odd, it is natural to define

$$\eta_{+-0} \equiv \frac{\mathrm{amp}(K_S \to \pi^+\pi^-\pi^0)}{\mathrm{amp}(K_L \to \pi^+\pi^-\pi^0)}$$

$$\eta_{000} \equiv \frac{\mathrm{amp}(K_S \to \pi^0\pi^0\pi^0)}{\mathrm{amp}(K_L \to \pi^0\pi^0\pi^0)}$$

These parameters are expected to be comparable in magnitude to η_{+-} and η_{00}, but only upper limits exist so far.

In the semileptonic decay modes, CP violation has been observed by measuring the charge asymmetry δ defined by

$$\delta_l \equiv \frac{\Gamma(K_L \to \pi^- l^+ \nu) - \Gamma(K_L \to \pi^+ l^- \bar{\nu})}{\Gamma(K_L \to \pi^- l^+ \nu) + \Gamma(K_L \to \pi^+ l^- \bar{\nu})}$$

where l is either e or μ. One can write

$$\delta_l = 2\,\mathrm{Re}\,\epsilon_L - \mathrm{Re}(2y_l + x_l - \bar{x}_l)$$

where

$$y_l = \frac{\mathrm{amp}(\overline{K^0} \to \pi^+ l^- \bar{\nu})^* - \mathrm{amp}(K^0 \to \pi^- l^+ \nu)}{\mathrm{amp}(\overline{K^0} \to \pi^+ l^- \bar{\nu})^* + \mathrm{amp}(K^0 \to \pi^- l^+ \nu)}$$

parametrizes CPT violation in the semileptonic decay amplitudes, and where

$$x_l \equiv \frac{\mathrm{amp}(\overline{K^0} \to \pi^- l^+ \nu)}{\mathrm{amp}(K^0 \to \pi^- l^+ \nu)}$$

$$\bar{x}_l \equiv \frac{\mathrm{amp}(K^0 \to \pi^+ l^- \bar{\nu})}{\mathrm{amp}(\overline{K^0} \to \pi^+ l^- \bar{\nu})}$$

represent violations of the $\Delta S = \Delta Q$ rule. CPT symmetry implies $y_l = 0$ and $\bar{x}_l = x_l^*$. By assuming that the $\Delta S = \Delta Q$ rule is exact and that there is no CPT violation in the semileptonic decay amplitudes, one gets the familiar relationship $\delta_l = 2\,\mathrm{Re}\,\epsilon_L$. We have here another relationship between $\mathrm{Re}(\epsilon - \Delta)$ and an experimentally determined quantity (δ_l), but now there is a complication due to y_l.

The K^0 and $\overline{K^0}$ states can be expressed in terms of the K_S and K_L:

$$K^0 \simeq [(1-\epsilon_L)K_S + (1-\epsilon_S)K_L]/\sqrt{2}$$

$$\overline{K^0} \simeq [(1+\epsilon_L)K_S - (1+\epsilon_S)K_L]/\sqrt{2}$$

where quadratic and higher terms in $\epsilon_{S,L}$ have been dropped. In $p\bar{p}$ annihilations, it is possible to tag initial K^0 and $\overline{K^0}$ states. By comparing their decay rates as functions of proper time, relationships for ϵ and Δ can be obtained without any complications from the parameters r_0 and y_l. For example, the decay rate of an initial K^0 state at large proper times into any final state is proportional to $|1-\epsilon_S|^2$, while the corresponding rate for an initial $\overline{K^0}$ state is proportional to $|1+\epsilon_S|^2$. The ratio of these rates is thus equal to $1 - 4\,\mathrm{Re}(\epsilon + \Delta)$.

2.3 The Bell-Steinberger Relation

Consider an arbitrary kaon state $a_S|K_S\rangle + a_L|K_L\rangle$ at $t = 0$. Since K_S and K_L are the vacuum eigenstates, the time evolution in vacuum of this state is expressed

by $a_S e^{-i\lambda_S t}|K_S\rangle + a_L e^{-i\lambda_L t}|K_L\rangle$. By invoking conservation of probability, one finds that

$$i(\lambda_L - \lambda_S^*)\langle K_S|K_L\rangle = \sum_f [\mathrm{amp}(K_S \to f)]^*[\mathrm{amp}(K_L \to f)]$$

where f runs over all decay channels. This equation is usually called the Bell-Steinberger relation[22], and it can also be derived from the eigenvalue equations $H|K_{S,L}\rangle = \lambda_{S,L}|K_{S,L}\rangle$ by assuming that M and Γ are hermitean. Neglecting Γ_L compared to Γ_S, we can write this relation as

$$(1 + i\frac{2\Delta m}{\Gamma_S})(\mathrm{Re}\,\epsilon - i\,\mathrm{Im}\,\Delta) = \sum_f \alpha(f)$$

where

$$\alpha(f) \equiv \frac{1}{\Gamma_S}[\mathrm{amp}(K_S \to f)]^*[\mathrm{amp}(K_L \to f)].$$

Note that the argument of $(1 + 2i\Delta m/\Gamma_S)$ is equal to Φ. The most important α's can be evaluated as

$$\alpha(2\pi, I{=}0) = \epsilon_0$$

$$\alpha(2\pi, I{=}2) = 2\omega^*\epsilon_2 = 2\omega^*\epsilon' + 2|\omega|^2\epsilon_L$$

$$\alpha(\pi^0\pi^0\pi^0) = \frac{\Gamma_L}{\Gamma_S}\frac{\Gamma(K_L \to \pi^0\pi^0\pi^0)}{\Gamma(K_L \to \mathrm{all})}\eta^*_{000}$$

$$\alpha(\pi^+\pi^-\pi^0) = \frac{\Gamma_L}{\Gamma_S}\frac{\Gamma(K_L \to \pi^+\pi^-\pi^0)}{\Gamma(K_L \to \mathrm{all})}\eta^*_{+-0}$$

$$\alpha(\pi l\nu) = \frac{\Gamma_L}{\Gamma_S}\frac{\Gamma(K_L \to \pi l\nu)}{\Gamma(K_L \to \mathrm{all})}(\delta_l + 2\Delta^* + x_l^* - \bar{x}_l^*)$$

The sum is dominated by $\alpha(2\pi, I{=}0)$, and we write

$$(1 + i\frac{2\Delta m}{\Gamma_S})(\mathrm{Re}\,\epsilon - i\,\mathrm{Im}\,\Delta) = \epsilon_0 + \tilde{\alpha}$$

where $\tilde{\alpha}$ is the sum of all $\alpha(f)$ except $\alpha(2\pi, I{=}0)$. Together with the equation $\epsilon_0 = \epsilon_L + r_0 = \epsilon - \Delta + r_0$ from the previous section, we can solve for the components of ϵ and Δ.

$$\epsilon_{||} = \epsilon_{0||} + \mathrm{Re}\,\tilde{\alpha}\cos\Phi$$

$$\epsilon_{\perp} = -\mathrm{Im}\,\tilde{\alpha}\cos\Phi$$

$$\Delta_{||} = \mathrm{Re}\,\tilde{\alpha}\cos\Phi + r_0\cos\Phi$$

$$\Delta_{\perp} = -\epsilon_{0\perp} - \mathrm{Im}\,\tilde{\alpha}\cos\Phi - r_0\sin\Phi$$

2.4 Predictions of CPT Symmetry

If CPT symmetry is exact, Δ and r_0 vanish, and the equations for the components of ϵ and Δ simplify to

$$\epsilon_{\|} = \epsilon_{0\|}$$

$$\epsilon_{\perp} = \epsilon_{0\perp} = -\mathrm{Im}\,\tilde{\alpha}\cos\Phi$$

The experimental limit for $\mathrm{Im}\,\tilde{\alpha}$ constrains the perpendicular component of ϵ_0. Since $\epsilon_0 = 2\eta_{+-}/3 + \eta_{00}/3$ (neglecting ω), CPT symmetry predicts

$$\frac{2}{3}\phi_{+-} + \frac{1}{3}\phi_{00} \simeq \Phi$$

Another consequence of CPT invariance is that $\overline{A}_I = A_I^*$, and from the expression for ϵ', we have $\arg\epsilon' = \delta_2 - \delta_0 \pm \pi/2$. Together with the experimentally determined $\delta_2 - \delta_0$ and $\mathrm{Re}(\epsilon'/\epsilon)$, this leads to the prediction

$$\phi_{+-} \simeq \phi_{00}$$

Finally, we note that since CPT invariance requires r_0 and y_l to be zero, the relationship $\delta = 2\,\mathrm{Re}\,\epsilon_0$ is also a prediction of CPT symmetry.

2.5 Regeneration

An interesting and very peculiar phenomenon occurs when a neutral kaon beam passes through matter. When the effects of both the strong and the weak interactions are taken into account, the eigenstates are neither the production states K^0 and $\overline{K^0}$ nor the vacuum states K_S and K_L. The purpose of this section is to summarize the quantitative description of how neutral kaons evolve in matter. In this context, the effect of CP violation is negligible, and one can use the approximate relationships $K_S \simeq (K^0 + \overline{K^0})/\sqrt{2}$ and $K_L \simeq (K^0 - \overline{K^0})/\sqrt{2}$.

Consider first an arbitrary initial kaon state $aK^0 + b\overline{K^0}$. When passing through matter, the K^0 and $\overline{K^0}$ components interact differently with the nucleons. For example, the reaction $\overline{K^0} + n \rightarrow \Lambda + \pi^0$ has no counterpart for the K^0. It follows that the initial state will turn into a different linear combination of K^0 and $\overline{K^0}$. In

particular, an incident pure K_L state will turn into a mixture of K_L and K_S, and this is called regeneration of K_S.[23]

The scattering amplitudes of the K^0 and the $\overline{K^0}$ are usually denoted $f(\theta)/k$ and $\overline{f}(\theta)/k$, where k is the wave number of the kaon. It is the difference between f and $\overline{f}$ that causes regeneration. At energies of interest, the kaon wavelength is very small compared to the typical distance between nuclei in matter, with the result that different scattering centers act coherently only in the forward direction. This means, in particular, that there is a forward peak of regenerated K_S, and this peak is known as coherent regeneration.[24]

Because there is coherence only in the forward direction, it is convenient to think of the kaon scattering as a combination of transmission (described by an index of refraction) and incoherent scattering off individual nuclei (elastic or inelastic). In this experiment, we are only interested in the coherent transmission, and we will treat the incoherent scattering as a background. The optical theorem relates the indices of refraction for K^0 and $\overline{K^0}$ to the forward scattering amplitudes $f(0)$ and $\overline{f}(0)$. The general description of the time evolution of the K^0-$\overline{K^0}$ system in matter is then obtained by combining the effect of the two indices of refraction with the weak interaction Hamiltonian. Here, we are interested in what happens to a pure K_L state passing through a block of matter (usually called the regenerator). The result of the general treatment is that an initial K_L state turns into the state $K_L + \rho K_S$, ignoring an attenuation factor common to K_L and K_S. The coefficient ρ is known as the coherent regeneration amplitude and is given by

$$\rho = \pi i N L g[f(0) - \overline{f}(0)]/k$$

where N is the density of scattering centers and L the length of the regenerator. The geometric factor $g = (1-e^{-x})/x$, where $x = (1/2 - i\Delta m/\Gamma_S)L/\Lambda_S$ with Λ_S the K_S decay length, results from the integration over the length of the regenerator. (At the energies that we are concerned with here, g is close to unity.) From the expression for ρ, it is seen that the phase of ρ (known as the regeneration phase) can be written $\phi_\rho = \pi/2 + \phi_f + \phi_g$. Here, ϕ_f is the argument of $[f(0) - \overline{f}(0)]/k$, and ϕ_g is the phase of the geometric factor g (which is calculated from well known

parameters).

The coherent regeneration of K_S can be studied with Regge-pole theory.[25] The result is two predictions for ρ, one for the magnitude and one for the phase. The magnitude of $[f(0) - \overline{f}(0)]/k$ is predicted to have a power law dependence on the kaon momentum, which means that we can write $|\rho| \propto p_K^{-\alpha}$. The other prediction states that $\phi_f = -(2-\alpha)\pi/2$, and this is also a prediction for ϕ_ρ (since the phase of g is well known). These predictions are supported by previous experiments[26] and, as will be seen later, by our own data. The value of α is around 0.6, from which it follows that ϕ_f is approximately -126°.

Chapter 3
APPARATUS

Why are high energy detectors so big? Maybe it is because there are so many people in each collaboration. If they all build something, then pretty soon there is a lot of stuff. Or maybe it is the other way around. Anyway, this chapter describes the E731 stuff, all of which can be found in the Meson Center experimental area at Fermilab. Additional information can be found in references 27–30. The various signals that are used in the triggers are introduced here, but how they combine into actual triggers is discussed in the next chapter, together with rates and other aspects of the data collection.

3.1 Beamline

The primary beam of 800 GeV protons from the Tevatron strikes a beryllium target, with a targeting angle of about five milliradians in the horizontal plane. Beryllium is chosen because the absorption of neutrons relative to the absorption of kaons tends to increase with lower atomic number. The choice of targeting angle further improves the kaon to neutron ratio (which increases with the angle) without losing too much kaon flux. The K^0 to $\overline{K^0}$ ratio also depends on the targeting angle, but since we use a regenerator to produce K_S from K_L, that ratio is not important.

After the target, there are sweeping magnets to deflect charged particles, and a six meter copper collimator which is also the beam dump. The two holes in the collimator measure one by one centimeter and are separated by 1.8 centimeters at the downstream end, about fifteen meters from the target. Right after the collimator, both beams pass through absorbers made of twenty inches of beryllium and three inches of lead. The beryllium once again improves the kaon to neutron ratio, while the lead converts photons into electron-positron pairs which are removed by

additional sweeping magnets. In the beam that will go through the regenerator, there is also a "moveable absorber" consisting of another eighteen inches of beryllium. This absorber and the regenerator both alternate between the two beams, and together they reduce the kaon flux by a factor of fifteen. This is necessary in order to make the two beams have roughly comparable $\pi\pi$ decay rates after the regenerator. Further downstream, additional collimators define the final beam size. The slab collimators between the beams also prevent cross-over from one beam to the other.

Around 122 meters from the target, there is a lead mask with two 10-by-10 centimeter holes for the beams. This mask can be thought of as defining the start of the decay volume, since it causes the acceptance for upstream decays to fall off quite rapidly (see Figure 1.5). A side view of the decay volume and all detector elements is shown in Figure 3.1, where the longitudinal scale is compressed by a factor of 10 in order to make the figure fit on one page. (But remember that in the kaon rest frame, the detector looks even more compressed.) The vacuum ends at the thin vacuum window, but in order to prevent multiple scattering and maintain good tracking resolution, the spaces between the drift chambers are filled by three helium bags.

There is actually a total of three magnets, but only the analysis magnet appears in the figure. Right before the lead mask (at about 120 m), there is a sweeping magnet which is used when only $\pi^0\pi^0$ decays are recorded. Since we prefer to collect $\pi^+\pi^-$ and $\pi^0\pi^0$ decays at the same time, this magnet is usually off. After the first trigger plane (at about 139 m), there is a "separator" magnet which can be used in a particular type of $\pi^0\pi^0$ trigger where one of the photons is converted in a lead sheet. The resulting electron-positron pair is then separated by the magnet and triggered on by the scintillator hodoscopes. This one is also off most of the time, since the standard trigger for the $\pi^0\pi^0$ decays relies on the total energy and number of clusters in the lead glass calorimeter. All the magnets are on, however, during the special lead glass calibration runs when photons from a copper target are converted into electron pairs by a thin copper foil. (For these runs, the lead in the absorbers is removed.) The analysis magnet and the separator magnet deflect

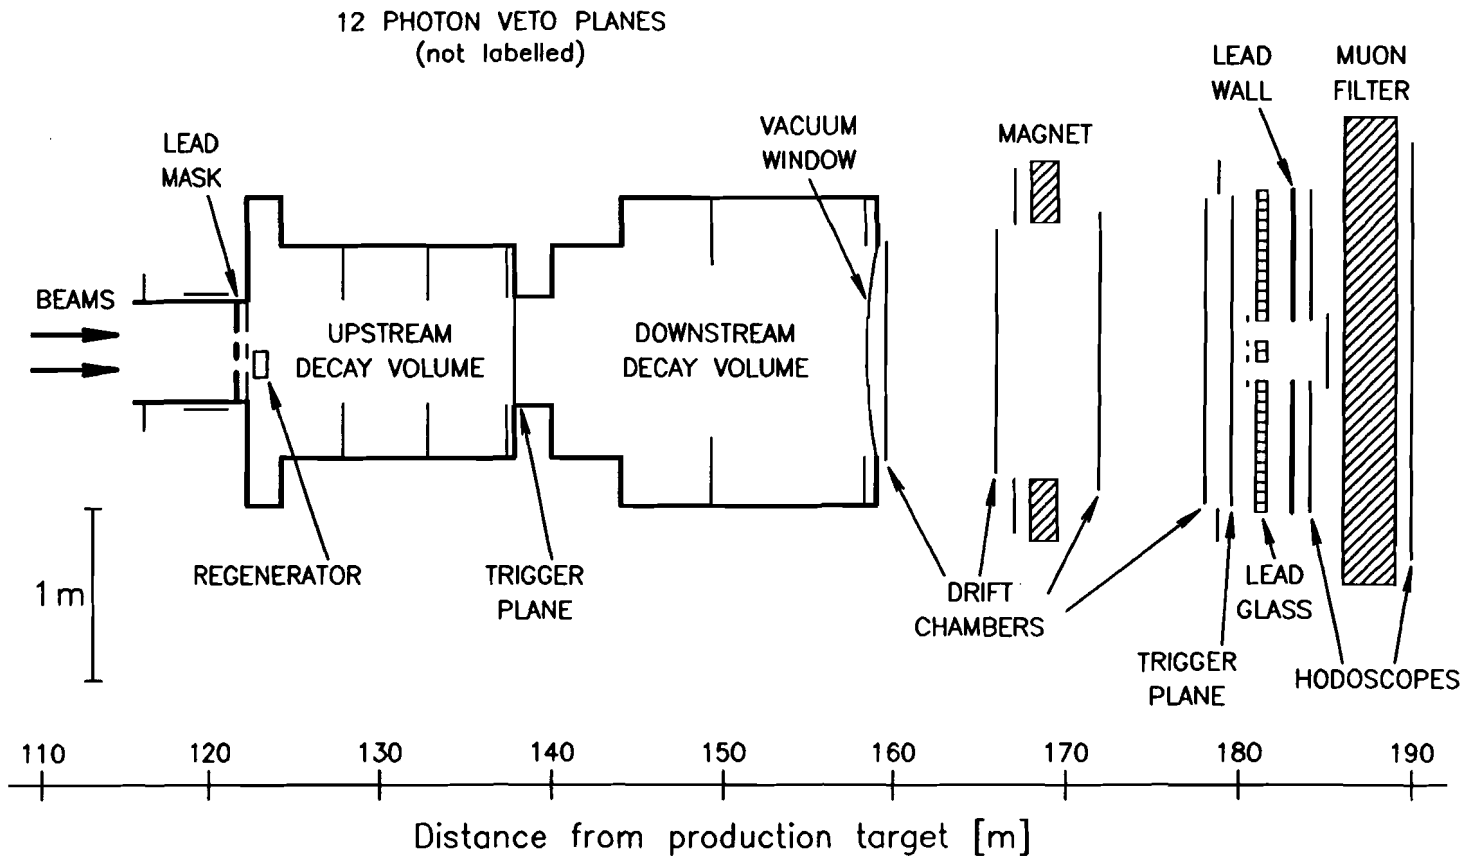

Figure 3.1: The E731 decay volume and detector.

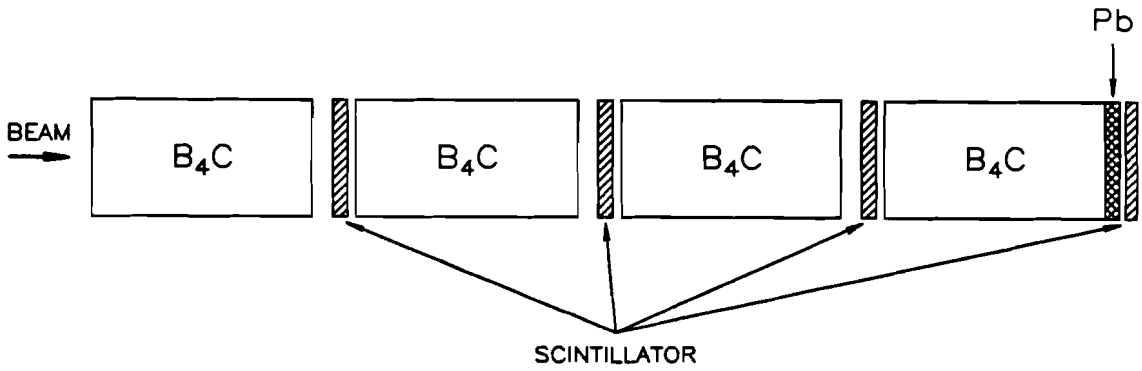

Figure 3.2: Schematic view of the regenerator.

the electrons in the horizontal direction, while the sweeping magnet deflects them in the vertical direction. By varying the magnet currents, all lead glass blocks are illuminated by electrons.

3.2 Regenerator

Right at the beginning of the decay volume, one of the beams passes through the regenerator. It is here that the (almost) pure K_L beam turns into a coherent mixture of K_L and K_S. Figure 3.2 shows a schematic view of the regenerator, which consists of four 19 centimeter blocks of boron carbide (B_4C). The regeneration phenomenon was discussed in chapter 2, and we recall that an incident K_L state turns into the amplitude $e^{-L/2\Lambda}(K_L+\rho K_S)$, where L is the length of the regenerator and Λ the kaon interaction length. The coherent regeneration amplitude ρ is given by $\rho = \pi i N L g(f-\overline{f})/k$, where N is the density of scattering centers, g a geometric factor close to unity, and $(f-\overline{f})/k$ the difference between the K^0 and $\overline{K^0}$ forward scattering amplitudes. Our regenerator has a total length equal to 2Λ in order to maximize the absolute size of the K_S flux, which is proportional to $L^2 e^{-L/\Lambda}$. This choice is good for the $\mathrm{Re}(\epsilon'/\epsilon)$ determination for which the experiment is optimized, because in that measurement the total number of K_S decays in the regenerated beam is compared to the total number of K_L decays in the vacuum beam. For the

interference measurements, it would be better to have a thinner regenerator and not attenuate the flux as much. The magnitude of ρ is determined from the data, and with a B_4C density of 2.5 g/cm^3, the magnitude of $(f - \overline{f})/k$ is found to be about 6 millibarn per B_4C molecule (for 70 GeV kaons).

In addition to the coherent signal, the regenerator also produces incoherently scattered kaons which are not wanted in the analysis. The inelastic part of this background can be suppressed by rejecting events with charged particles emerging from the regenerator, and for this purpose, scintillator counters are installed after each boron carbide block. The phototube signals are added and discriminated to form the Regenerator Anti signal (RA), which is used to veto triggers (see chapter 4). The fourth and last scintillator layer also serves as a hardware-defined edge for the decay distribution in the regenerated beam. The piece of lead after the fourth block converts photons, so that there is an edge for the $\pi^0\pi^0$ decays as well as for the $\pi^+\pi^-$ ones. (However, the edge is shifted slightly upstream for the neutral decays.)

As explained in the introduction, it is necessary to have the regenerator alternate between the two beams if one wants to compare the decay rates with and without the regenerator. At Fermilab, the beam is on for twenty seconds every minute, and the regenerator and the upstream absorber move from one position to the other during the forty seconds between beam pulses.

3.3 Trigger Hodoscopes

There are two sets of scintillator hodoscopes, both of which are used in the $\pi^+\pi^-$ trigger. The first set, which consists of two planes of one millimeter thick scintillator, is located at about 138 m where it separates the upstream and downstream parts of the decay volume. The scintillator has to be very thin in order to maintain a good resolution for the decay vertex, which is determined by extrapolating the two charged tracks to their intersection point. The two planes are known as the T and V hodoscopes and have transverse dimensions of 50 cm (horizontal) by 62 cm (vertical). Both planes are made out of six scintillator pieces (vertical

for T and horizontal for V), and each piece is equipped with phototubes at both ends. The twelve output signals for each plane are summed together and discriminated to form the T and V trigger signals which are used in the $\pi^+\pi^-$ trigger. The discriminator level is set to require two minimum ionizing particles.

Since the T and V counters are required in the $\pi^+\pi^-$ trigger, the decay distribution is cut off at 138 m and this causes a certain loss of statistical power in the determination of ϕ_{+-}. If the experiment were optimized for the phase measurement, the trigger counters would be located much further downstream. For the $\pi^0\pi^0$ decays, both the upstream and the downstream decay volumes can be used, but it will be necessary to take into account regeneration and absorption in the trigger counters.

The second set of trigger counters also consists of two scintillator planes, which are located at around 180 m. These planes, which are both about one centimeter thick, are known as the B and C hodoscopes. (Since they are located after the spectrometer, their thickness does not affect the tracking resolution.) Figure 3.3 shows the layout of the two hodoscopes, whose transverse dimensions are about 1.8 by 1.8 meters. Most of the B counters are 10 centimeters wide, except for a few 17 centimeter ones on the sides. The C counters are all 15 centimeters wide. The phototube signals are discriminated at a level requiring one minimum ionizing particle, and the discriminated signals are then added together in different ways to make trigger signals. The 1B (and 1C) signals are logical sums over all 30 B (or 24 C) counters, thus requiring at least one minimum ionizing particle. The 2B and 2C signals require at least two counters to have a minimum ionizing particle. The B_W and B_E signals are logical sums over the west half (B1–B8 and B16–B23) and east half (B8–B15 and B23–B30) of the B hodoscope. Note that the two middle counters belong to both halves. The C_U signal sums the upper two thirds of the C hodoscope (C1–C8, C13–C20), and C_D is a sum of the lower two thirds (C5–C12, C17–C24). The B_W, B_E, C_U and C_D signals are used in the trigger to require a symmetric two-track topology (see chapter 4.3).

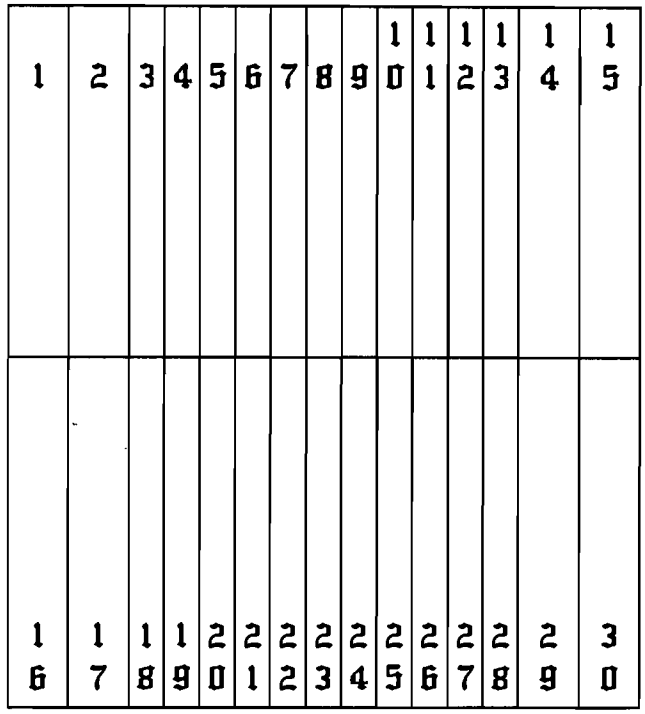

1	13
2	14
3	15
4	16
5	17
6	18
7	19
8	20
9	21
10	22
11	23
12	24

Figure 3.3: Beam's view of the B (*top*) and C (*bottom*) hodoscopes.

3.4 Spectrometer

The trajectories of charged particles are measured with four drift chambers, two on each side of the large aperture analysis magnet. With this configuration, the direction of a track is known both before and after the magnet, and the particle momentum is then calculated from the track deflection at the magnet. The main component of the magnetic field is in the vertical direction, and all particles with unit electric charge receive a transverse momentum kick of about 200 MeV/c in the horizontal plane. The polarity of the field can be changed by turning the magnet upside down or — more easily — by reversing the current. Of course, we have never tried the first method, but we do reverse the current periodically during the data taking.

The drift chambers all have two x planes (vertical wires) and two y planes (horizontal wires). The spacing between sense wires is 12.7 millimeters (half an inch), and the two planes in each view — also called complementary planes — are offset by half of that amount, so that the left-right ambiguities can be resolved. The field wires are arranged in a hexagonal pattern around the sense wires, with the two complementary planes sharing the field wires in between them. This way, the two planes are only 11 millimeters apart, making the chambers particularly compact (about three inches from window to window). Together with thin windows (25 micrometers) and the helium bags between the chambers, this reduces the multiple scattering of charged particles and improves the resolution. Except for the field wires, there is nothing else to separate the x and y views from each other. The transverse size of the chambers increases from about 100 sense wires per plane for the first chamber to about 150 sense wires per plane for the fourth one. The chamber gas is a 50/50 argon-ethane mixture at atmospheric pressure, with a small amount of ethanol.

The wire pulses propagate to the amplifier and discriminator cards that are mounted on the chamber edges. Twisted pair cables then take the signals to 32-channel TDC modules which reside in special purpose Camac crates. These TDC's have a 10-bit range with a resolution of one nanosecond, but we only need the range

0–250 ns. A nice feature of this TDC system is the ability to "autotrim" all channels and to test the calibration with a built-in reference clock. By autotrimming the TDC's regularly during the data taking, we keep them calibrated to within 1-2 ns over the entire 0–250 ns range. Almost all drift times are less than 150 ns (see Figure 5.2), but there is a tail extending above 200 ns. The TDC's operate in common stop mode, with start signals from the wires and a stop signal derived from the trigger, and with the timing adjusted so that hits close to the wire will result in a digitized time around 240 ns. The TDC's do not restart on a second pulse unless it arrives more than 250 ns later, but even so there is a risk that late pulses will kill the real signal and increase the inefficiency. Fortunately, most of the late pulses can be suppressed by inhibit signals derived from the trigger.

The special crate controllers do a sparse read-out of the TDC modules and transfer the data to interface modules in regular Camac crates. The transfer time depends on the number of words (*i.e.*, the number of wire hits), but a typical total time for digitization and data transfer is around 400 microseconds. The Camac read-out of the TDC data then takes another 300 microseconds on average. This total acquisition time for the drift chamber data is comparable to the read-out time for other systems which is mostly done via Fastbus.

The TDC modules also provide a "fast OR" of the wire signals, and this is useful for the trigger. The Chamber 2 x planes are divided into two halves, east and west. Since good $\pi^+\pi^-$ decays should have zero transverse momentum and therefore one track on each side, we form a coincidence between the east sum and the west sum. Allowing for some inefficiency, a logical sum of the two coincidences (one per plane) is taken. This is the Chamber 2 east-west topology condition that is used in the charged trigger (see chapter 4.3).

3.5 Calorimeter

The electromagnetic calorimeter consists of 804 lead glass blocks, which are stacked according to Figure 3.4. Each block is $5.8 \times 5.8\,\mathrm{cm}^2$ across and 61.0 cm deep. The depth is equal to 20 radiation lengths. There are two holes for the beams, since

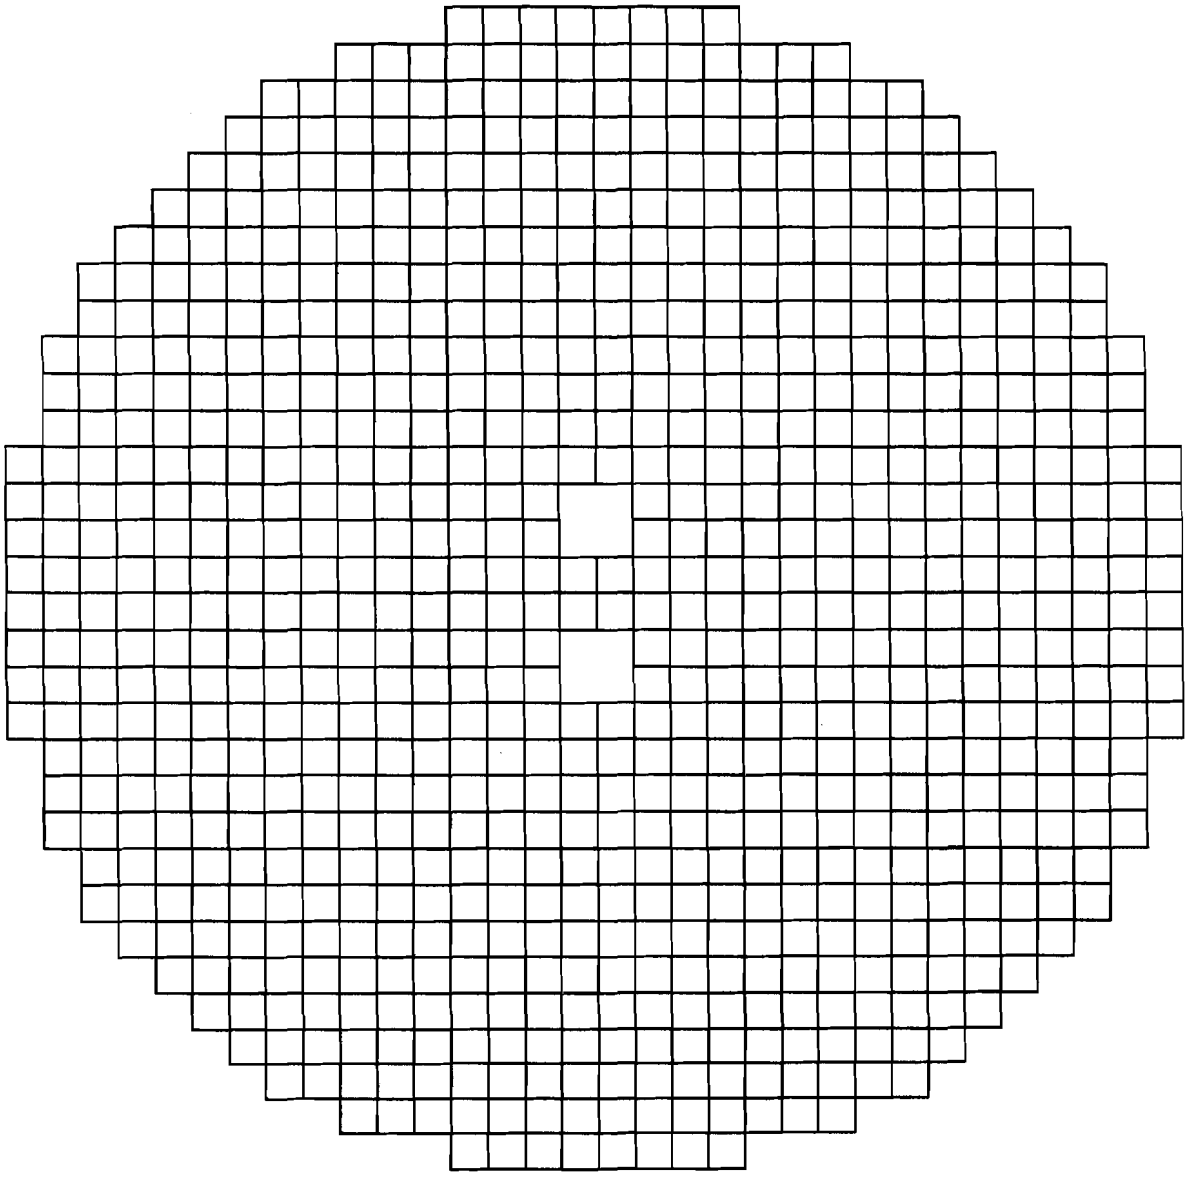

Figure 3.4: Configuration of lead glass blocks in the calorimeter.

the lead glass is not able to withstand the high rates that would result if the beams were to pass through it. The blocks closest to the beam holes suffer some radiation damage, which increases the light attenuation in the glass and changes the energy response. Fortunately, it is possible to reverse almost all of the damage by shining ultraviolet light at the lead glass. This "curing" was performed periodically during beam down time.

The energy from an electromagnetic shower is typically contained in a three-by-three "cluster" of blocks. The light from each block is collected by a phototube which is attached to the back of the block, and the 804 phototube signals are sent to 96-channel Fastbus ADC modules with 15-bit resolution. The 150 ns ADC gate is derived from the trigger, and the gains are set so that 1 GeV corresponds to about 200 counts. The long gate reduces the sensitivity of the gains to timing drifts, but it increases the risk of having out-of-time accidental photons contaminate the signal. However, the phototube signals are also summed in groups of nine and digitized with a 30 ns gate. These "adder" signals can be used to tell whether a cluster of energy was in time with the trigger or not. All adder signals are summed together to form a total energy signal, which is discriminated for use in the trigger. The ET_{HI} signal corresponds to an energy of about 28 GeV, while the ET_{LO} signal has a much lower threshold.

Special pedestal events are recorded in between beam pulses. (The beam is on for twenty seconds each minute.) Average pedestals are calculated for all ADC channels and down-loaded into the Fastbus data acquisition system. During real events, these pedestals are subtracted from the digitized signals, and only channels above a certain threshold are read out. (Pedestal events are also taken while the beam is on, and these can be used for studying intensity related pedestal shifts and similar effects.)

The 804 phototube signals are also fed into a trigger processor known as the Hardware Cluster Finder, where they are digitized by flash ADC's. A threshold corresponding to an energy of 1 GeV is applied to each signal, and the processor determines how many energy clusters there are in the lead glass by looking for "islands" above threshold. This number of clusters is used in the trigger (see

chapter 4.3). Because of the very short ADC gate, there is very little sensitivity to accidental clusters.

3.6 Background Rejection

One important feature of the detector is the large number of photon veto planes that are used for suppressing the background from $K_L \rightarrow \pi^0\pi^0\pi^0$ decays. These veto counters can be seen in Figure 3.1 as the unlabelled detector planes which are distributed along the decay volume and the detector in order to cover as much solid angle as possible. Each set of veto counters provides a trigger signal, which is usually referred to by the initials given below. Some of them are used in the triggers to veto photons and/or charged particles that would otherwise escape the detector. The phototube signals are digitized by ADC's (or at least latched) and can be used in the analysis for background suppression.

The first three (counting from upstream) are the Pinching Anti (PA), Sweeping Anti (SA) and Mask Anti (AM). These are all made of scintillator and detect photons thanks to the metal — beam pipe or lead mask — that is in front of them. The next five are all inside the decay volume and are known as VA1, VA2, DRAC/N, VA3 and VA4 (for Vacuum Anti and Decay Region Anti). The four VA's consist of a scintillator layer followed by two lead-lucite sandwiches. DRAC/N is made of a scintillator layer followed by a lead-scintillator sandwich. (When you get hungry on shift, you go out to the beam-line for a bite.) All five are able to distinguish between photons and charged particles thanks to the first layer of scintillator. The next two veto planes are the Magnet Anti (MA, in front of the analysis magnet) and the Lead Glass Anti (LGA, right after the last drift chamber). Like the VA's, they consist of a scintillator layer followed by two lead-lucite sandwiches.

Right in front of the lead glass is the Collar Anti (CA), which is made of copper and scintillator. CA forms a collar around the beam holes with a width equal to one half of a lead glass block. A couple of meters behind the lead glass is the Back Anti (BA) which covers both beams and is made of lead and lucite in several layers. It is thick enough so that a photon shower should not have a

signal in the last layer. This way, the photons can be distinguished from neutron interactions, and the trigger signal that identifies photons is known as BAγ.

At the end of the detector there is the traditional muon filter — three meters of steel — followed by a scintillator hodoscope known as MU2, with transverse dimensions (width and height) of 2.5 m × 2.4 m. The signal from MU2 is used for vetoing $K_L \rightarrow \pi\mu\nu$ decays. Before the muon filter, there is another scintillator hodoscope (roughly 2.1 m × 2.1 m) which is named MU1, even though it is not used for muon identification. It is used instead as a hadron veto together with the 12 cm lead wall. The lead glass (61 cm) and the lead wall contribute about twenty radiation lengths each, which is enough to contain almost all photon showers. Hadrons, on the other hand, have their shower maximum at MU1. Both the lead wall and the hodoscope have a single large hole (16 cm × 40 cm) for the two beams.

Chapter 4
DATA COLLECTION

Our experiment took data during a fixed target run at Fermilab which lasted from July 1987 to February 1988. This chapter describes the typical running conditions and how the data were collected.

We started out alternating between charged mode ($K \rightarrow \pi^+\pi^-$) and neutral mode ($K \rightarrow \pi^0\pi^0$) data taking, but towards the end of the run we collected all $K \rightarrow \pi\pi$ decays simultaneously. The data sets from these different running periods are labelled alphabetically from A to H. The analysis presented in this thesis is based on data from the G set which contains about 20% of our total data sample. It was during the G set that most of the simultaneous data taking occurred, which means that the study of systematic errors is easier for this data set than for the other sets.

4.1 Beam

During fixed target running at Fermilab, the 800 GeV protons are extracted from the Tevatron in "spills" which occur once per minute and last about twenty seconds. Each one-minute machine cycle starts with the injection of protons into the Tevatron followed by acceleration to the maximum energy. This batch of protons is then extracted at a nearly uniform rate during the twenty second spill. The extracted beam is split into many beams which are steered onto the various targets in the experimental areas. Our target typically receives 10^{12} protons per spill. The time distribution within the spill of reconstructed $K \rightarrow \pi^0\pi^0$ decays is shown in Figure 4.1. (This is summed over all spills in the G set.) The dips in the distribution are present whenever the neutrino experiments are running. Since they prefer short

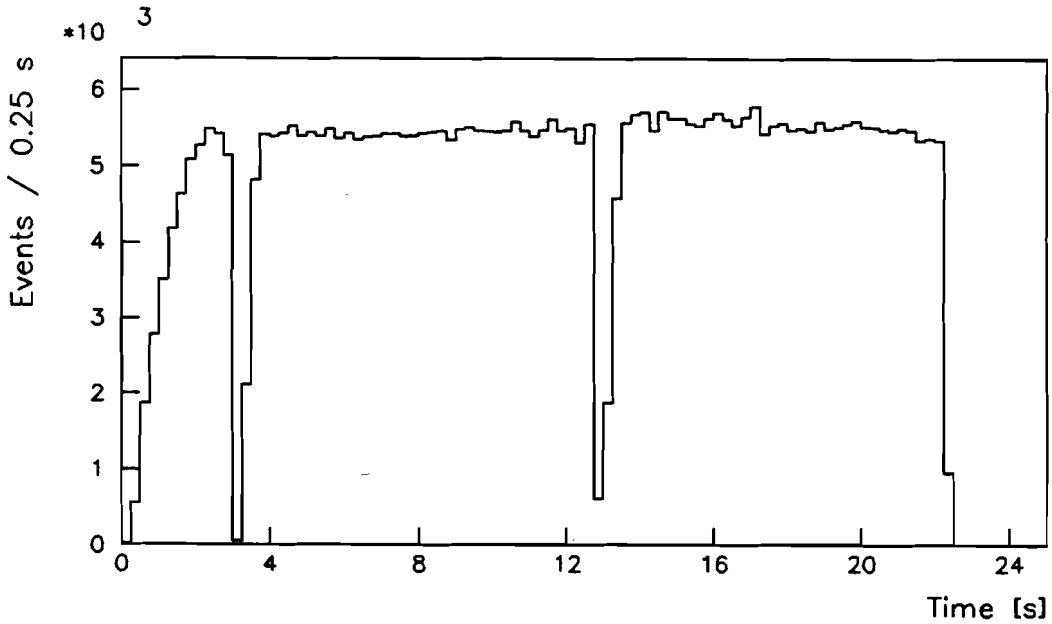

Figure 4.1: Event distribution within the spill, summed over G set data.

bursts of high intensity, all available beam is steered onto their targets twice each spill.

There is a very fine structure to the beam within each spill since the protons circulate in the Tevatron in bunches as opposed to a continuous stream. This is called the RF (radio frequency) structure. As a bunch of protons passes the extraction point, a small fraction is shaved off and sent to the experimental areas where it becomes known as a "bucket". The buckets are separated by about twenty nanoseconds and each one is roughly two nanoseconds wide. With an accelerator radius of one kilometer, there are about one thousand bunches in the ring and each bunch circulates a million times during a spill. As a result, we have 10^9 buckets per spill hitting the target, with an average of one thousand protons per bucket.

The question now arises whether all buckets are equally populated. To study this issue, we looked at the number of out-of-time coincidences between two separate groups of counters in the MU2 hodoscope, one group at the top and one at the bottom of the counter bank. Let N_{top} and N_{bot} be the scaler counts for the two signals, and let N_{coin} be the number of coincidences between top and bottom, where one of them is delayed by an amount equivalent to several buckets. If all

buckets have the same beam intensity, the quantity $N_{top}N_{bot}/N_{coin}$ is equal to the total number of buckets in the spill. Any intensity variation on time scales larger than a few buckets would introduce a correlation between the two signals, and the value of $N_{top}N_{bot}/N_{coin}$ would decrease and become equal to the effective number of populated buckets. The average value measured during the data taking is 0.93×10^9. This does not rule out large intensity variations from bucket to bucket, but it shows that any such variation is random.

There is also a very large scale structure to the beam intensity due to such things as magnet failures and other accelerator problems that cause the beam to be off. This structure is too irregular and unpredictable to be discussed here. An event census shows that there were 17,103 spills in the analyzed data (G set) with at least one good $\pi^0\pi^0$ event. The G set data taking lasted 23 days or 33,000 minutes, from which it follows that 52% of the time was used taking analysis quality data. The remaining 48% was spent calibrating the detector, taking special study runs, fixing problems, and occasionally waiting for beam.

4.2 Rates

Before going into the details of the triggers, we should look at the hardware signals that are available. These signals come from the various detector elements and can be used in the triggers either in coincidence with other signals or in veto. Descriptions of how they are formed can be found in the relevant sections of chapter 3. In order to measure the rates in the detector, each signal is connected to a scaler channel. The gate for the scalers is open during the spill and the total counts are read out at the end of each spill. Table 4.1 is a list of these hardware signals and their average rates during the data taking (G set). The averages are taken over events and not over time, which ensures that low intensity spills with no events do not contribute. In calculating the rates, an effective spill time of 20 seconds is assumed (see Figure 4.1). MC6SEM is a beam monitor that tells us the total number of protons on target during a spill. The average beam intensity for the G set data is 8×10^{11} protons on target per spill, which translates to 40 GHz.

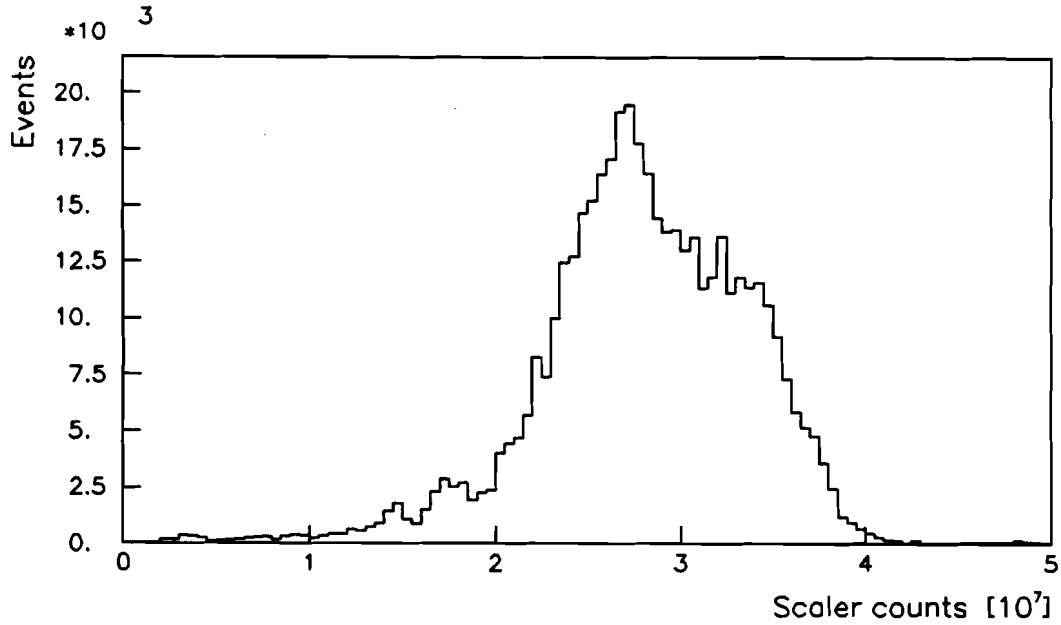

Figure 4.2: Beam intensity as measured by the muon hodoscope.

A few other hardware signals are also available but not scaled. In addition to 1B and 1C, the B and C hodoscopes provide 2B, 2C (at least two particles in B and C respectively), B_E, B_W (one particle in the right or left half of B), C_U, C_D (one particle in the upper two thirds or lower two thirds of C), and B5 (five minimum ionizing particles in at least one B counter).

Not surprisingly, the beam intensity varies from spill to spill. Figure 4.2 shows the total scaler count per spill from the MU2 hodoscope with one entry for each reconstructed $K \rightarrow \pi^0\pi^0$ event. During spills with average intensity, we find, on the average, fifteen good $\pi^0\pi^0$ events from the regenerated beam (coherent kaons only) and about eight from the vacuum beam. (This is for kaon energies between 40 and 160 GeV and a decay vertex between 110 and 150 meters from the target.) With a typical dead-time of 45% and an acceptance of 15% , this implies that there is an average of 90 $K_L \rightarrow \pi^0\pi^0$ decays per spill in the vacuum beam (energy and vertex range as above). Using the known K_L decay length and the branching ratio into $\pi^0\pi^0$, one calculates that there is a flux of six million K_L (40–160 GeV) per spill in the vacuum beam through our decay volume. The kaon flux in the regenerated beam is about 6% of the vacuum beam flux, but since a small fraction of those

Table 4.1: Rates.

Signal	Threshold (MIP = minimum ionizing particle)	Rate	Ref.
V	2 MIP in V hodoscope	980 kHz	3.3
T	2 MIP in T hodoscope	1020 kHz	3.3
1B	1 MIP in B hodoscope	1560 kHz	3.3
1C	1 MIP in C hodoscope	1890 kHz	3.3
ET_{HI}	28 GeV in the lead glass	61 kHz	3.5
ET_{LO}	6 GeV in the lead glass	180 kHz	3.5
MU1	2 MIP in MU1 hodoscope	590 kHz	3.6
MU2	1 MIP in MU2 hodoscope	1390 kHz	3.6
CA	6 GeV ($\sim$ 25 MIP) in CA	28 kHz	3.6
BAγ	6 GeV in BA	240 kHz	3.6
PA	1 MIP in PA scintillator	210 kHz	3.6
SA	1 MIP in SA scintillator	250 kHz	3.6
AM	1 MIP in AM scintillator	190 kHz	3.6
RA	1 MIP in RA scintillator	600 kHz	3.2
VA1	1 MIP in VA1 scintillator	1850 kHz	3.6
VA2	1 MIP in VA2 scintillator	1150 kHz	3.6
DRAC	1 MIP in DRAC scintillator	1250 kHz	3.6
VA3	1 MIP in VA3 scintillator	1040 kHz	3.6
VA4	1 MIP in VA4 scintillator	960 kHz	3.6
MA	1 MIP in MA scintillator	550 kHz	3.6
LGA	1 MIP in LGA scintillator	890 kHz	3.6
VA4γ	$\sim$ 0.3 GeV in VA4 lead-lucite	590 kHz	3.6
LGAγ	$\sim$ 1.0 GeV in LGA lead-lucite	460 kHz	3.6
MC6SEM	Number of protons on target	$4 \times 10^{10}\,\mathrm{sec}^{-1}$	

kaons are K_S with much higher 2π decay rates, the number of observed decays is larger than in the vacuum beam.

4.3 Triggers

We will now look at how the triggers are formed. The charged trigger (for $K \to \pi^+\pi^-$) and the neutral trigger (for $K \to \pi^0\pi^0$) are the ones that the analysis is based on. We also took some data with a variety of other triggers in order to study and calibrate the detector. The subtriggers are referred to here by their abbreviations (like MU2 for the muon hodoscope). Please see chapter 3 for descriptions of the relevant detector components.

4.3.1 Charged Trigger

The charged decays result in one positive and one negative pion which both produce a track in the spectrometer and minimum ionizing signals in the four trigger hodoscopes. (The V and T hodoscopes at the end of the upstream decay volume, and the B and C hodoscopes after the spectrometer.) The pions often shower in the lead glass or the lead wall. The chance of anything making it through the muon filter to the muon veto counters (MU2) is very small, and the probability that a pion will decay into a muon before the muon filter is also small.

First-level charged trigger:

$$(\mathrm{V}+\mathrm{T})\cdot \mathrm{B}_{+-}\cdot \mathrm{C}_{+-}\cdot \overline{\mathrm{MU2}}\cdot \overline{\mathrm{RA}}\cdot \overline{\mathrm{PA}}\cdot \overline{\mathrm{AM}}\cdot \overline{\mathrm{VA2}}\cdot \overline{\mathrm{VA3}}\cdot \overline{\mathrm{VA4}}\cdot \overline{\mathrm{LGA}\gamma}$$

where $\mathrm{B}_{+-} = 2\mathrm{B}\cdot \mathrm{B}_E\cdot \mathrm{B}_W$ (at least two B counters hit, at least one in the east half and at least one in the west half) and $\mathrm{C}_{+-} = \mathrm{C}_U\cdot \mathrm{C}_D$ (at least one hit in the upper two thirds and one hit in the lower two thirds).

Second-level charged trigger: Chamber 2 east-west topology condition.

The timing is defined by the accelerator RF signal (see section 4.3.4). By taking the logical sum of the V and T hodoscopes, we are much less sensitive to inefficiencies in either one. The total transverse momentum of the two pions should be zero, both before and after the Analysis Magnet (assuming both pions

Table 4.2: Rates for the charged trigger.

Trigger Stage	Rate
$(V+T)\cdot\overline{LGA\gamma}$	1350 kHz
after $B_{+-}\cdot C_{+-}\cdot\overline{PA}\cdot\overline{AM}\cdot\overline{RA}$	8.3 kHz
after $\overline{VA2}\cdot\overline{VA3}\cdot\overline{VA4}$	4.2 kHz
after $\overline{MU2}$	2.6 kHz

experience the same field). We therefore add the topological requirements in the B and C hodoscopes. The Chamber 2 topology trigger is not formed in time for a coincidence, and is therefore applied at the second level. The first-level trigger starts the data acquisition, and if the Chamber 2 condition is not satisfied, the read-out is aborted and a fast clear is issued to all systems.

The $K_L \to \pi\mu\nu$ decays are rejected by the MU2 veto, and inelastically produced kaons in the regenerator by the RA veto. The remaining vetoes reject events with particles leaving the detector at relatively large angles. (PA, AM, VA2, VA3 and VA4 are scintillator vetoes, while LGAγ is a lead-lucite veto.) Table 4.2 shows the average rates at various stages of the trigger as measured by scalers during the data taking. The Chamber 2 topology requirement vetoes 29% of the events accepted by the first-level trigger.

In addition to the $\pi^+\pi^-$ decays, the trigger also accepts a large fraction of the $\pi e\nu$ and $\pi^+\pi^-\pi^0$ decays. In previous runs of the experiment, these modes were rejected by inserting a lead curtain in front of the B hodoscope during charged mode running. The electron from a $\pi e\nu$ decay and the photons from a $\pi^+\pi^-\pi^0$ decay would shower in the lead, and the events were vetoed with the B5 signal. It would not be good to use the lead curtain during the simultaneous running of charged and neutral modes, since the photons from the neutral decays would start to shower in the curtain instead of in the calorimeter. Furthermore, the $\pi e\nu$ and $\pi^+\pi^-\pi^0$ decays are useful for calibration. Therefore, we did not use this technique at all, except for a few special runs.

4.3.2 Neutral Trigger

The neutral decays result in two neutral pions which immediately decay into two photons each. (A 20 GeV π^0 has a decay length of about four microns.) The four photons shower in the lead glass calorimeter, giving a large total energy (ET) to trigger on.

First-level neutral trigger:

$$\mathrm{ET}_{HI} \cdot \overline{\mathrm{V}} \cdot \overline{\mathrm{BA}\gamma} \cdot \overline{\mathrm{VA4}\gamma} \cdot \overline{\mathrm{LGA}\gamma} \cdot \overline{\mathrm{CA}} \cdot \overline{\mathrm{MU1}} \cdot \overline{\mathrm{RA}} \cdot \overline{\mathrm{PA}} \cdot \overline{\mathrm{AM}} \cdot \overline{\mathrm{VA2}} \cdot \overline{\mathrm{VA3}} \cdot \overline{\mathrm{VA4}}$$

Second-level neutral trigger: Four or six clusters found in the lead glass by the Hardware Cluster Finder.

The timing is defined by the accelerator RF signal (see section 4.3.4). The threshold for the ET_{HI} signal corresponds to a total energy of about 28 GeV in the lead glass. (In the analysis we use events with kaon energies above 40 GeV.) The main problem for this trigger is what to do with the large numbers of $K_L \rightarrow \pi^+\pi^-\pi^0$ and $K_L \rightarrow \pi^0\pi^0\pi^0$ decays that also deposit a large energy in the calorimeter. The V veto rejects $\pi^+\pi^-\pi^0$ decays from the upstream decay volume, while the downstream events are mostly accepted. (V is one of the two trigger hodoscopes at the end of the upstream decay volume.) The $\pi^+\pi^-\pi^0$ decay mode is very valuable for checking the lead glass calibration, and many of the upstream events are accepted by the charged trigger. The decays accepted by the neutral trigger, however, turn out to be less useful, since this trigger favors events where the charged pions shower and therefore interfere with the electromagnetic showers.

The $\pi^0\pi^0\pi^0$ decays with all six clusters in the calorimeter can be reconstructed, and because of the large number of these events, they provide more information about the detector acceptance than the $\pi^0\pi^0$ decays. For this reason, we accept events with four or six clusters found by the Hardware Cluster Finder. The number of clusters found is not available for the first-level trigger, so this requirement is added at the second level. If at least one photon from a $\pi^0\pi^0\pi^0$ decay escapes the lead glass or if two or more photon showers merge, the event is of no use. In the case where the resulting number of clusters in the lead glass is four, these events form a background to the $\pi^0\pi^0$ signal. To reduce this background at the trigger level, we

Table 4.3: Rates for the neutral trigger.

Trigger Stage	Rate
$ET_{HI} \cdot \overline{V}$	29 kHz
after $\overline{PA} \cdot \overline{AM} \cdot \overline{RA} \cdot \overline{VA2} \cdot \overline{VA3} \cdot \overline{VA4}$	23 kHz
after $\overline{MU1}$	4.0 kHz
after $\overline{BA\gamma}$	3.0 kHz
after $\overline{CA} \cdot \overline{VA4\gamma} \cdot \overline{LGA\gamma}$	1.6 kHz

add the lead-lucite vetoes BAγ, VA4γ and LGAγ in order to reject events with one or more photons missing the lead glass. (BAγ identifies photons that go through the beam holes. VA4γ and LGAγ are positioned to cover solid angle outside the lead glass.)

The CA veto rejects events where a photon hits the calorimeter near the beam holes. If this edge were not hardware defined, it would be difficult to understand the acceptance in the region around the beam holes. The MU1 veto kills events with large hadronic showers in the lead glass, such as neutron interactions. (The lead glass is followed by a lead wall and the MU1 scintillator hodoscope. Electromagnetic showers are absorbed in the glass and the wall, while hadron showers reach their maximum at MU1.) Inelastically produced kaons in the regenerator are vetoed by RA. The remaining scintillator vetoes (PA, AM, VA2, VA3 and VA4) reject events with charged particles leaving the detector at relatively large angles. Table 4.3 shows the average rates at various stages of the trigger as measured by scalers during the data taking. 10% of the first-level triggers have four clusters, and 7% have six.

In previous runs of the experiment, a different kind of trigger was used for the neutral decays. One of the four photons was converted into an e^+e^- pair by inserting a thin lead sheet between the T and V trigger hodoscopes. The separator magnet opened up the two tracks, and the analysis magnet made them reconverge at the calorimeter. The trigger would look for two charged tracks (very much like the $\pi^+\pi^-$ trigger), with the T(rigger) counters in coincidence and the V(eto) counters in veto. The requirement of exactly one conversion cuts the signal significantly, and

it is not possible to use the downstream decay volume. It would be difficult to take charged and neutral mode data simultaneously, since the charged mode mass and vertex resolutions are worsened by scattering of the pions in the lead sheet. Also, the required magnet currents for the conversion technique are very different from the standard charged mode magnet settings.

4.3.3 Other Triggers

There were also a few special purpose triggers, which were usually prescaled and taken together with the main triggers (charged and neutral). Their timing was also defined by the accelerator RF signal, except for the flasher, pedestal and end-of-spill triggers.

The $\pi\mu\nu$ ("Kμ3") trigger was essentially the same as the charged trigger, with MU2 in coincidence instead of in veto, and without the Chamber 2 topology condition. The rate was about 1.6 kHz.

There was a muon trigger (B · MU2) with an average rate of 600 kHz.

In order to study the accidental activity, an accidental trigger was formed (on purpose, not accidentally) by taking a coincidence between two small scintillator counters. These were located downstream of the target, but at an angle from the beam direction, so that a particle going through the counters would not go through any part of the detector. This way, the accidental trigger was correlated with beam activity in general, but uncorrelated with activity in any particular detector element.

Pedestal events were taken about once per second, both during the spill and in between spills. The timing was defined by a radioactive source to avoid correlations with any other frequency (*e.g.*, 60 Hz).

Flasher events were also taken about once per second, with a light source illuminating each block in the lead glass array. These events provided short-term gain tracking for the lead glass blocks.

There were end-of-spill events in order to read out the scalers and some miscellaneous information such as magnet currents and readings from the beam intensity and beam profile monitors.

4.3.4 Trigger Timing

As mentioned previously, the timing of the triggers is defined by the accelerator RF, which is a signal with the same frequency as the bucket structure in the beam. This method avoids the timing jitter in the trigger signals due to the size of the trigger counters (usually long scintillator counters). However, since the RF signal is generated far away from our experiment (and is sent to us through long cables), it becomes crucial to monitor the relative timing between the RF and events in our detector. Any change in the electronics generating or transmitting the RF, for example, could easily introduce a phase shift in the signal. Temperature changes could change the propagation speed in the cables and have a significant effect on the timing.

To monitor the RF timing we measure the interval between a particular event in the detector and the RF signal using a TDC. The event chosen is a coincidence between one of the B counters and one of the C counters, because it defines a small area of scintillator and therefore minimizes the jitter. These measurements are accumulated independently of the data taking, and if the RF drifts relative to the beam events, we adjust its timing (typically by 0.5–1.0 ns) by adding or removing short delay cables.

4.4 Data Acquisition

The charged and neutral triggers were taken simultaneously during the G set. In order to get roughly equal numbers of reconstructed $\pi^+\pi^-$ and $\pi^0\pi^0$ events, the charged trigger was prescaled by a factor of eight. (The charged mode branching ratio is about a factor of two higher, and the average final acceptance for $\pi^+\pi^-$ decays is almost three times higher than for $\pi^0\pi^0$ decays.) The neutral mode second-level trigger accepted all four and six cluster events during the G set data taking. In earlier data sets, the six cluster events were usually prescaled by a factor of eight. The muon trigger (prescaled by 2^{16}), the accidental trigger (prescaled by 20) and the Kμ3 trigger (prescaled by 50) were also connected. The flasher, pedestal and end-of-spill events were included in all data sets.

All triggers are combined into one unified first-level trigger. This signal starts the data acquisition system and defines the gates for the ADC's and the latches and also the stop signal for the drift chamber TDC's. Since it is necessary in the analysis to know the trigger type for the event, each trigger is also connected to one channel in a latch module for later reference by the off-line software. The neutral mode second-level trigger (from the Hardware Cluster Finder) comes about twenty microseconds after the first-level trigger, while the charged mode second-level trigger (Chamber 2 east-west logic) is ready within one microsecond. In either case, if the second-level condition is not satisfied, the read-out is aborted and all systems cleared.

The ADC's reside in two Fastbus crates, one for the lead glass channels and one for all other channels. Each crate is read out by a crate controller, which transfers the data to a Fastbus memory buffer. All other data (including the drift chamber TDC's) come from Camac crates and are read out by a microprocessor-based Camac module acting as a branch driver. The data is transferred from the memory of this module through a Fastbus/Camac interface and a Fastbus crate controller into another memory buffer. The three memory buffers are read out by a PDP 11/45 computer and combined into one single event record which is written to tape. The average event size is about six hundred 16-bit words. The transfer of data between Camac and Fastbus and from the memory buffers to the tape is asynchronous with the basic read-out and does not contribute to the dead-time.

The data acquisition was dead for about one millisecond after each trigger, which resulted in an average dead-time of 45% during the G set. Triggers that arrived randomly (*e.g.*, the neutral triggers) had a 55% chance of being accepted, while the prescaled charged triggers were accepted 67% of the time. (The effective prescale factor was about six and a half instead of eight.) The delay time for the second-level triggers (not greater than 20 μs) made a negligible contribution to the overall dead-time.

4.5 Monitoring

The scaler rates were displayed by the microprocessor-based module that also did the Camac read-out. A fraction of the events written to tape were transferred to a μVAX, which accumulated histograms of basic variables from most of the detector systems. No event reconstruction was done at this stage. One tape from each eight-hour shift was analyzed off-line by a standard monitoring job. Checking the output from this job kept the next shift busy.

Chapter 5
CALIBRATION

As described in chapter 4, the data acquisition system collects all event information from the read-out electronics and writes event records to tape. These records consist entirely of digital information, such as ADC counts, TDC counts and various other bits and bytes. Before doing any analysis, it is necessary to convert this stream of ones and zeroes into useful quantities like the momenta and directions of charged tracks in the spectrometer and the energies of electromagnetic showers in the calorimeter. How to accomplish this is the subject of the following sections.

5.1 Coordinate System

We use a Cartesian coordinate system with the z axis along the beam direction, which happens to point almost directly north. The plane perpendicular to the beam is spanned by the horizontal x axis and the vertical y axis. Positive x is west and positive y is up, which makes our coordinate system a right-handed one. (Surprisingly, this has no effect on the measurements of CP and CPT violation.) The origin is chosen to be at the center of the target, and the z axis passes by definition through the center of the lead glass calorimeter. Positive and negative x are also referred to as beam left and beam right, respectively. The positions of all beam and detector elements in this coordinate system are determined through surveys. Muon tracks are also used to locate many parts of the detector.

5.2 Analysis Magnet

The main field component of our analysis magnet is B_y (vertical), which bends the tracks in the horizontal plane. The transverse momentum kick is about 200 MeV/c with the standard magnet current of about 1700 amperes.

The magnet was calibrated with the ZIPTRACK apparatus which is a Fermilab tool for mapping magnetic fields. All three field components were measured at points in a three dimensional grid with a step size of two inches in x and y and one inch in z. The field strength was also measured with an NMR probe and found to be 3.8 kilogauss at the center.

In the data analysis, the small B_x and B_z components are neglected, while B_y is represented by a look-up table giving the effective transverse momentum kick as a function of x and y. This table is obtained directly from the field measurements, but with the overall scale adjusted to make the average invariant kaon mass in the $K \rightarrow \pi^+\pi^-$ decays come out right. Because the absolute field strength changes slightly after each field reversal, a separate fudge factor is derived for each period between reversals. (The polarity of the field was changed regularly during the data taking to reduce the systematic uncertainty for the measurement of the charge asymmetry δ in $\pi e \nu$ decays.)

5.3 Drift Chambers

There are four drift chambers, each with two x planes (vertical wires) and two y planes (horizontal wires). The two planes in each view are offset by one half of the distance between sense wires, so that the left-right ambiguity can be resolved. The wire spacing is half an inch in all the planes, which together with the offset between complementary planes gives a fundamental cell width d of 6,350 micrometers.

Figure 5.1 illustrates schematically a track passing through the two complementary planes of a view. (Only sense wires are shown.) One wire in each plane is hit, and the drift times (t_1, t_2) are measured by the TDC's. If the track fitting algorithm were to use only the positions of the hit wires, the single plane resolution would be $2d/\sqrt{12} = 3.7$ mm. By using the drift times to calculate the drift

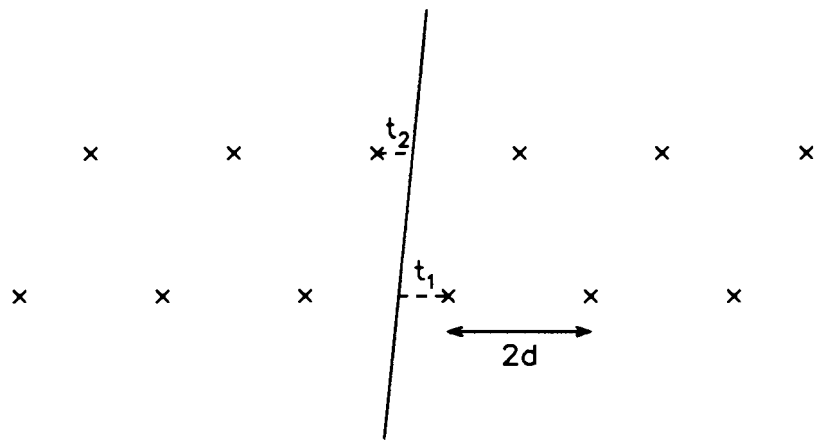

Figure 5.1: A track passing through two complementary planes.

distances between the wires and the track, the resolution is greatly improved.

5.3.1 Time to Distance Conversion

Figure 5.2 shows the drift time distribution for hits in one of the sixteen planes. (The figures illustrating the drift chamber calibration are made from all two-track events on one raw data tape.) Since the TDC's operate in common stop mode, the drift time is obtained by inverting the digitized time, using a somewhat arbitrary origin. The sharp edge of the distribution corresponds to zero drift distance (*i.e.*, hits near the wire), and the largest drift times correspond to the maximum drift distance ($= d$). If only we knew the drift velocity everywhere, we would also be able to convert any other drift time into a drift distance without any difficulty. However, since the velocity is not very well known, another approach is necessary.

The solution comes from observing that the track illumination becomes uniform when the wire plane is collapsed onto the fundamental cell. Knowing then that the space distribution of the hits in Figure 5.2 is uniform, the mapping from time to space is obtained by integrating and normalizing the time distribution. More specifically, if a fraction x ($0 \leq x \leq 1$) of the hits have a drift time less than t, then t corresponds to a drift distance of $x \cdot d$, where d is the cell width.

To guarantee a uniform distribution in the space variable, only hits from good tracks should be entered in the time distribution plot. This means that the track

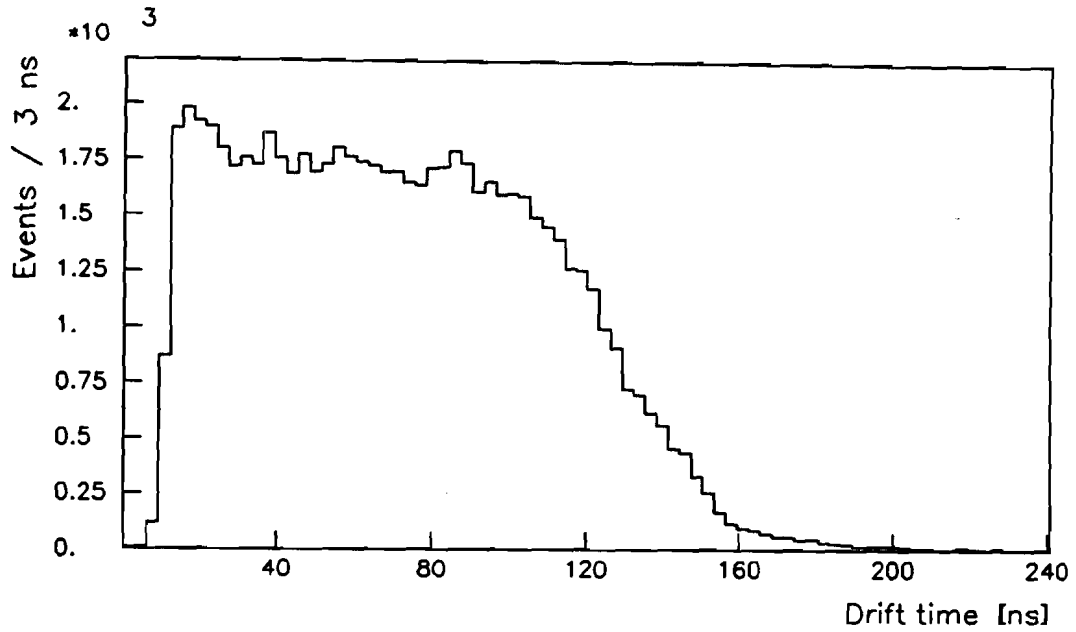

Figure 5.2: Drift time distribution.

finder has to be running when the time to distance calibration is done. On the other hand, the track finder usually requires time to distance conversions, so how does one ever get started? — The easiest way is probably to find a preliminary time to distance conversion by using all hits (not just track hits). This would be good enough to run the track finder, and then the final calibration could be obtained.

A more serious concern about the assumption of uniform distribution, is the possibility that the wire inefficiency depends on the distance of the track from the wire. If a track has no hit in a particular plane, it is not represented in the time distribution plot for that plane, possibly causing the corresponding space distribution to be non-uniform. Fortunately, the time distribution of these missing hits can be created by using the drift times of the hits in the complementary plane. (The chance of missing both hits is very small.) For each missing hit, one takes the drift time of the complementary hit and translates it back into a drift time for the missing hit. The mapping from a time in one plane to a time in the other plane is obtained from the large number of events with hits in both planes. The missing hits are then added to the existing hits before the time distribution is integrated and normalized. The final time to distance calibration is given as a look-up table

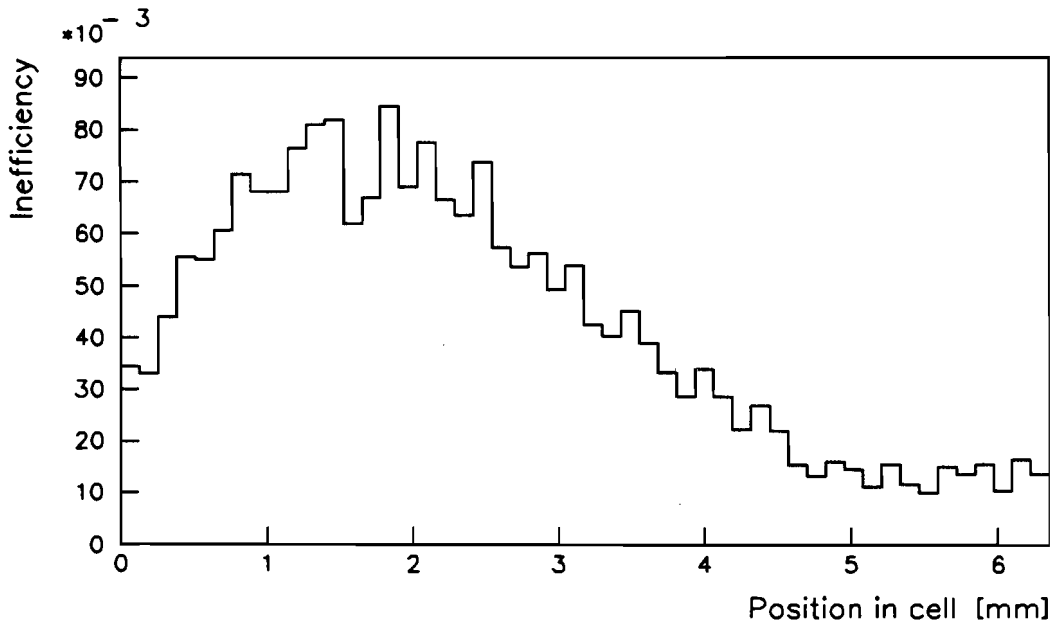

Figure 5.3: Inefficiency versus distance from the wire.

with one x value for each nanosecond of drift time. (The step size of the TDC's is one nanosecond.)

It is worth pointing out that the time distribution (if properly normalized) is also a plot of the drift velocity as a function of position in the cell (as given by the drift time). Since most of the area is under the flat portion of the curve, it follows that the drift velocity is constant throughout most of the cell. The velocity is typically around 50 μm/ns and normally does not vary more than a few percent during the run. In order to track the variations, the time to distance calibration is updated periodically. To avoid any problems from possible timing differences between triggers, the calibration is always done using events from the real data triggers (charged and neutral).

The procedure for finding the distribution of missing hits also gives information about the plane inefficiency. Figure 5.3 shows the fraction of hits that are missing (*i.e.*, the inefficiency) as a function of the distance from the wire. The plot is for the first (*i.e.*, most upstream) y plane in Chamber 1. As can be seen in the figure, this plane has an inefficiency of about five percent. All other planes average around one percent.

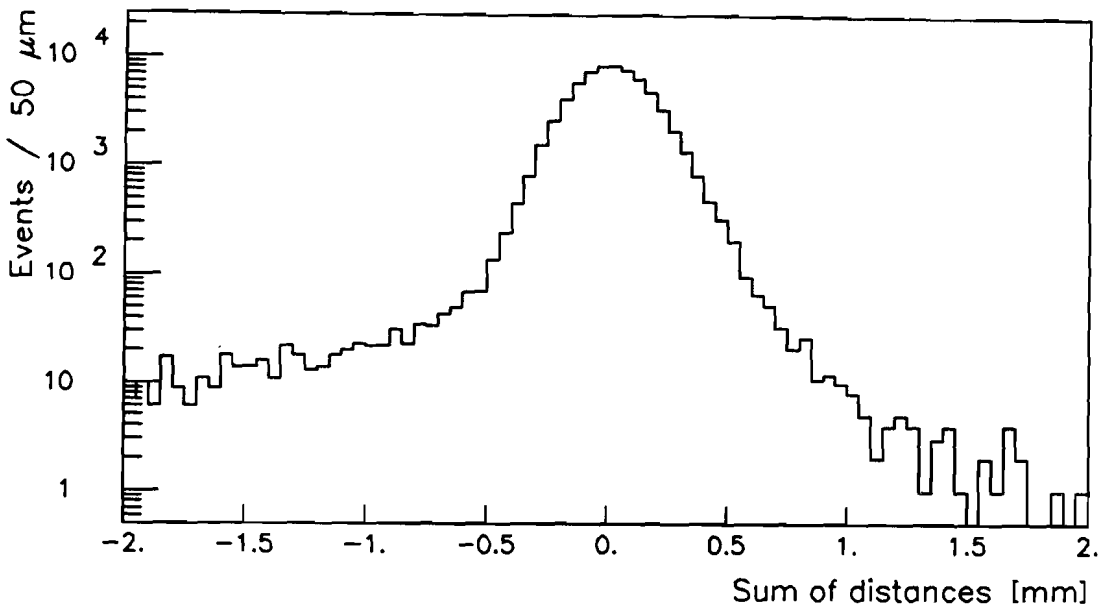

Figure 5.4: Sum of drift distances for one plane pair.

5.3.2 Resolution

By comparing the two track position values obtained from the two complementary planes, we can calculate the resolution. It is easy to see in Figure 5.1 that the sum of the two drift distances (with a small correction for the track angle) should be equal to the cell width d. This "sum of distances" for one pair of wire planes is plotted in Figure 5.4 (after subtracting d). Changes in the drift velocity makes the peak move away from zero, indicating a need for a new calibration. The resolution is now obtained directly from the width of the peak. Cutting at ± 1000 μm, the "rms" is typically 150–160 μm, which is equivalent (through a $\sqrt{2}$) to a single plane resolution of 110–120 μm.

A change in the relative timing between the tracks and the TDC stop signal (*i.e.*, the trigger) would also affect the sum of distances. This can be used to eliminate out-of-time tracks. If a track is not in time with the event, all drift times (and therefore all sum of distances) are shifted. We remove a track if at least two of the sum of distance values are more than 1000 μm off.

A final check is to plot the average track residual as a function of the drift distance. This is done in Figure 5.5 for one of the wire planes. It appears very

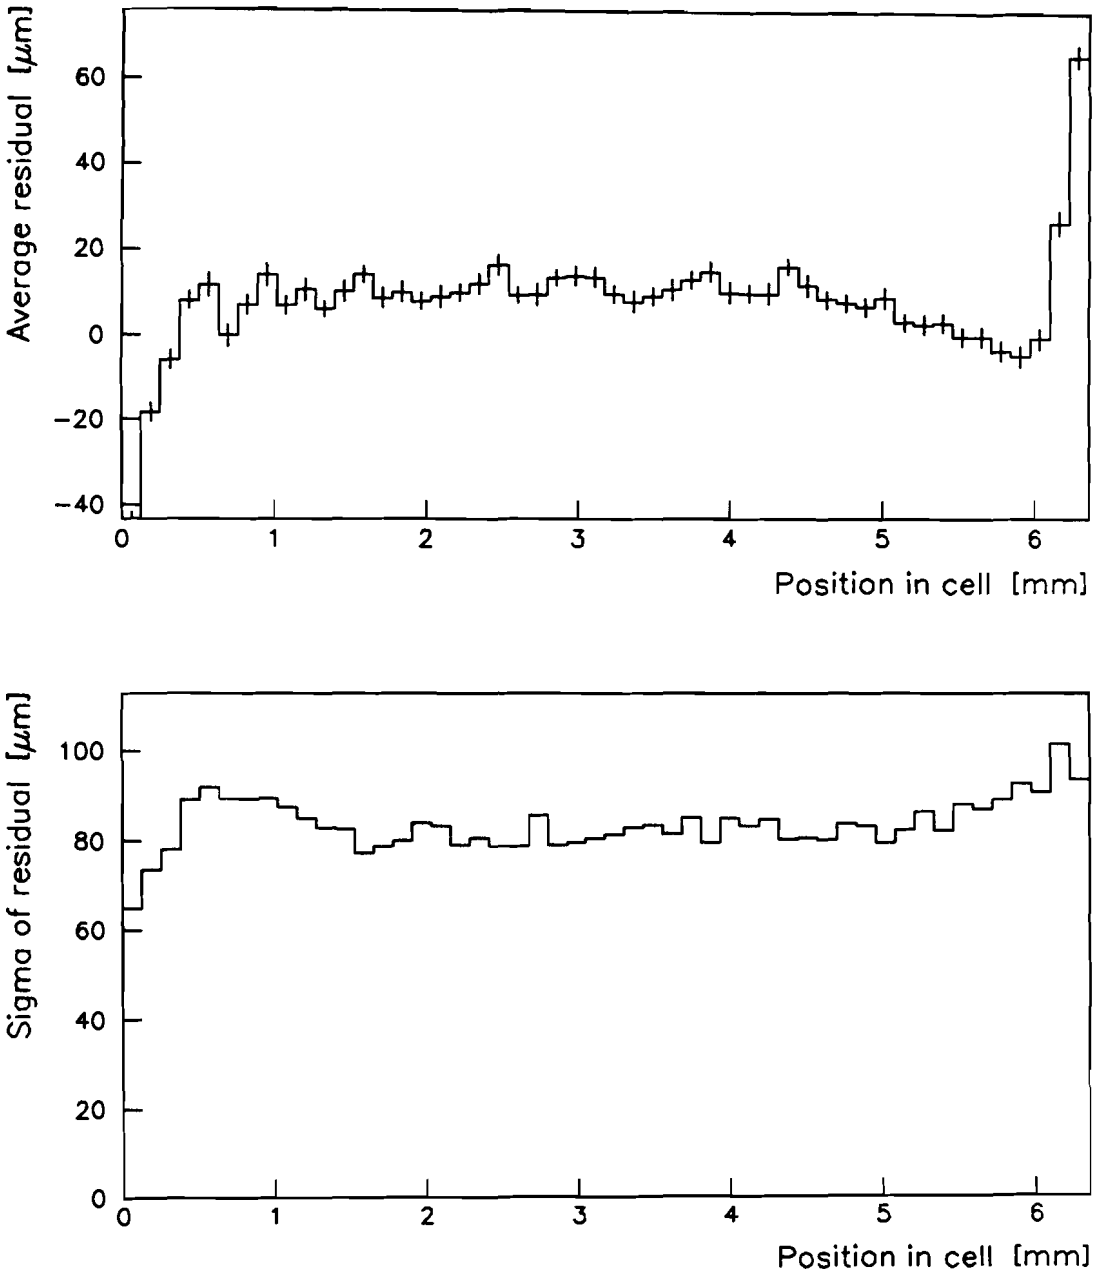

Figure 5.5: Average track residual (*top*) and rms (*bottom*) versus position in the cell.

difficult to make the average residual be exactly zero without fudging the time to distance calibration. However, the discrepancy is small compared to the effects caused by normal variations in drift velocity over time. The large deviations from zero at the end-points are caused by the track fitting algorithm. Since the fitted track position is essentially the average of the hit positions from the two planes, the tracks tend to be pushed away from the cell boundaries (by an amount related to the resolution). For example, if the drift distance in one plane is very close to zero, the fitted track position is almost always further away from the wire, and this makes the average residual non-zero. (The effect can be hidden by plotting the average residual against the fitted track position instead of against the drift distance.) The "rms" of the track residual is equal to the plane resolution divided by $\sqrt{2}$.

5.3.3 Track Finding

The track finding consists of identifying which hits belong to a track, and then fitting a line to those hits in order to obtain the track parameters. Since the vertical wires (x view) provide information only in the xz plane and the horizontal wires (y view) only in the yz plane, we have essentially two completely independent tracking systems. From the x view, we get lines in the xz plane (called x tracks), and from the y view, lines in the yz plane (called y tracks). Both x and y tracks are lines in two-dimensional spaces, specified by two parameters each (usually a slope and an intercept). Combining an x track with a y track, we then get the four parameters necessary to completely determine a track in our three-dimensional space. In addition to the separation between the x and y views, the x tracks are split into two segments, one upstream and one downstream of the analysis magnet (which bends the tracks in the horizontal plane).

For example, in the upstream x view, there are four wire planes, two in chamber 1 and two in chamber 2. All combinations of one hit from a chamber 1 plane and one hit from a chamber 2 plane are tried, with a straight line drawn through the wires that were hit. The two remaining planes are then searched for hits close enough to this line, and the track is accepted if at least three out of the four planes

have good hits. After resolving (wherever possible) the left-right ambiguities, the drift times are converted into drift distances and added to the wire positions to get the final x values of the hits. The z values are simply the z positions of the wire planes (measured in the survey), and the track parameters are then obtained from a least square fit to the (x, z) points.

The downstream x view is treated exactly the same way, and complete x tracks are found by matching the upstream and downstream segments at the analysis magnet. The transverse separation between the segments at the magnet is required to be below a certain value, and any remaining left-right ambiguity is resolved by minimizing this separation.

In the y view, where there is no bending, we look for tracks going through all four chambers. Combinations of hits from two different wire planes are tried, and the remaining six planes are searched for consistent hits. A track is accepted if five out of the eight planes have good hits. Here, we can resolve the left-right ambiguity for "single" hits (hits without a hit in the complementary plane) by choosing the sign that agrees best with a line determined by the hits in the other chambers. After all ambiguities are resolved, the drift distances are added to the wire positions to get the final y values of the hits, and the track parameters are obtained from a least square fit to the (y, z) points.

For all hits belonging to a track (x or y), a residual can be calculated. This residual is simply the difference between the projection of the track in the wire plane and the actual position of the hit as given by the sum of the wire position and the drift distance. Track "chi-squares" are calculated by summing the squares of the residuals and dividing by the number of hits minus two. (The tracks are obtained from a two parameter fit.) To get real chi-squares, one should also divide by the square of the single plane resolution.

5.3.4 Track Matching

So far, we have only discussed how to find the x and y tracks separately. In order to find the complete tracks (in three-dimensional space), it is necessary to match the x and y views with each other. If there is more than one track in the

spectrometer, it is impossible to tell, based on chamber information alone, which x track belongs to which y track. By correlating the tracks with hits in the B and C hodoscopes, we get a fairly good matching of x and y tracks. The most reliable match is obtained by using the cluster positions in the lead glass calorimeter. It should be pointed out that this method does not require that all tracks hit the calorimeter.

Once the tracks are matched, it is possible to correct for two small effects. One of these comes from the signal propagation time along the sense wires, and the other effect arises if, for example, the x wires are not exactly vertical. In both cases, corrections are made by using the known track position from the other view. The final track parameters are obtained from a new fit to the corrected hit positions. If there are enough hits in the y view, we also fit for upstream and downstream y segments. The bending at the magnet of y tracks is small but not exactly zero, so the resolution is improved slightly by using separate track parameters upstream and downstream.

5.3.5 Momentum Determination

Finally, we determine the magnitude of the momentum vector from the deflection at the magnet of the x track and the known transverse momentum kick. The direction of the momentum vector is, of course, the same as the direction of the track. The kick depends on the transverse track position at the magnet according to the look-up table obtained from the magnet calibration. This transverse position is taken from the intersection point of the upstream and downstream segments with a correction to account for the fact that the track bends along a smooth curve.

5.4 Calorimeter

When the calorimeter data is analyzed, the first step is to translate the ADC counts into energy values. This is done using a gain file which specifies the conversion factor from counts to GeV for each lead glass block. Figure 5.6 shows a typical example of the energy deposit from the four electromagnetic showers of a

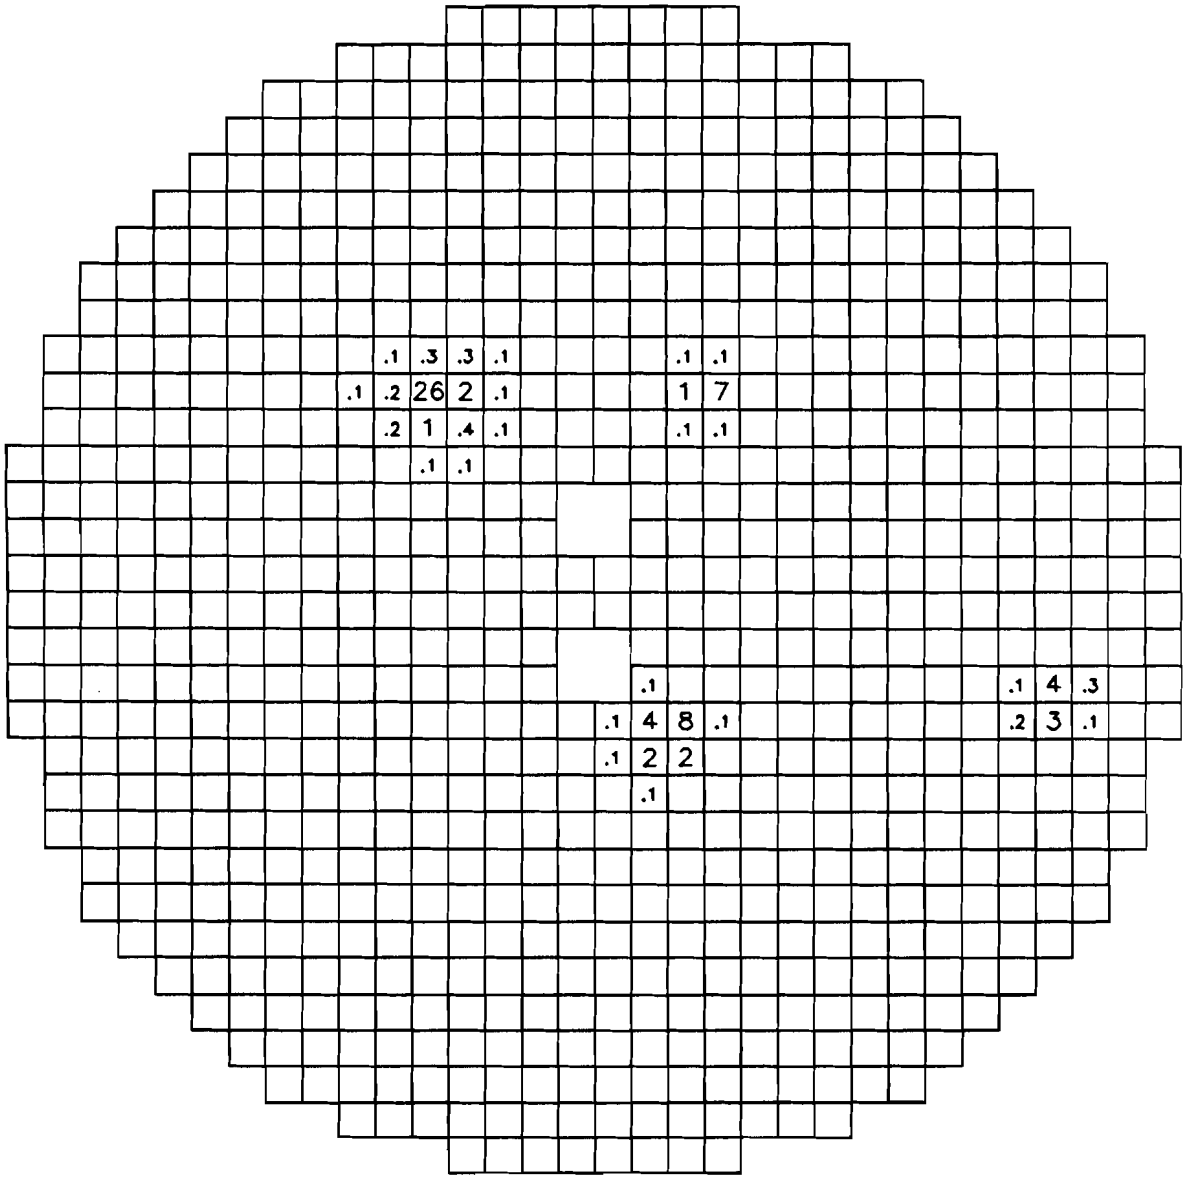

Figure 5.6: Typical $\pi^0\pi^0$ event in the lead glass.

$\pi^0\pi^0$ decay. Energies above 1 GeV have been rounded to the nearest GeV and energies below 1 GeV to the nearest tenth of a GeV. (An empty block has an energy less than 0.05 GeV.) Figure 5.7 shows a fairly typical $\pi^+\pi^-$ decay where one of the pions showered while the other one went straight through (resulting in a minimum ionizing signal). Hadron showers are usually more spread out than the electromagnetic ones.

Assuming for the moment that a good gain file already exists, the second step consists of identifying the clusters and calculating their energies and positions. The energy of a cluster is obtained by summing over the blocks belonging to it and correcting for things like the non-linear response of the lead glass and the "missing" energy for clusters close to the edge. The cluster energy is then checked using calibration events in order to determine the final gain file.

5.4.1 Cluster Finding

The clusters are found by searching for local energy maxima in the lead glass array. If a block has a larger energy than each of its eight neighbors (and if this energy is above 0.2 GeV), it is taken to be the center block of a cluster. For the determination of the cluster energy, we then use the nine blocks in a three-by-three array centered on this particular block. In order to be counted as a genuine cluster, the total energy in the nine blocks is usually required to be at least 0.5 GeV.

Sometimes two photons are so close to each other in the lead glass that there is only one energy maximum, and in this case we say that the photons have fused. This is of particular concern for the analysis of the $\pi^0\pi^0$ decay mode, where $3\pi^0$ decays with fusions form a background. In order to reduce this background, a number of criteria have been developed to identify fusions. These empirically derived conditions are based on quantities like the ratio of corner block energy to center block energy, the ratio of side column (row) energy to middle column (row) energy and the amount of energy outside the three-by-three array.

For certain triggers (including the $\pi^0\pi^0$ but not the $\pi^+\pi^-$), the event record contains information from the Hardware Cluster Finder (HCF). Figure 5.8 shows the event from Figure 5.6 as seen by the HCF. A threshold of 1 GeV is applied to all

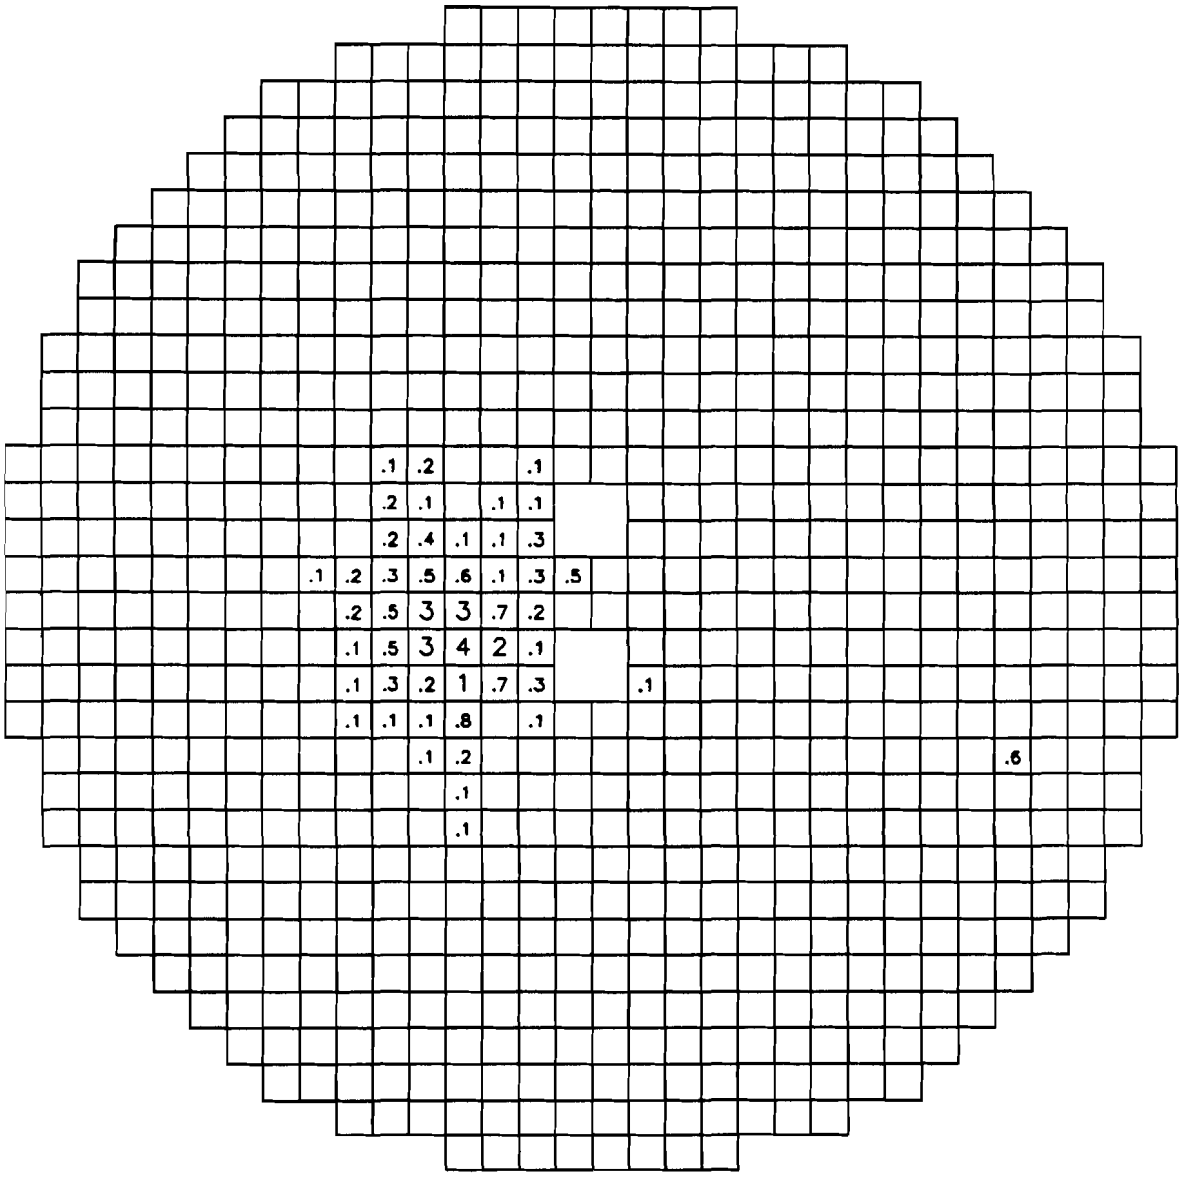

Figure 5.7: Typical $\pi^+\pi^-$ event in the lead glass.

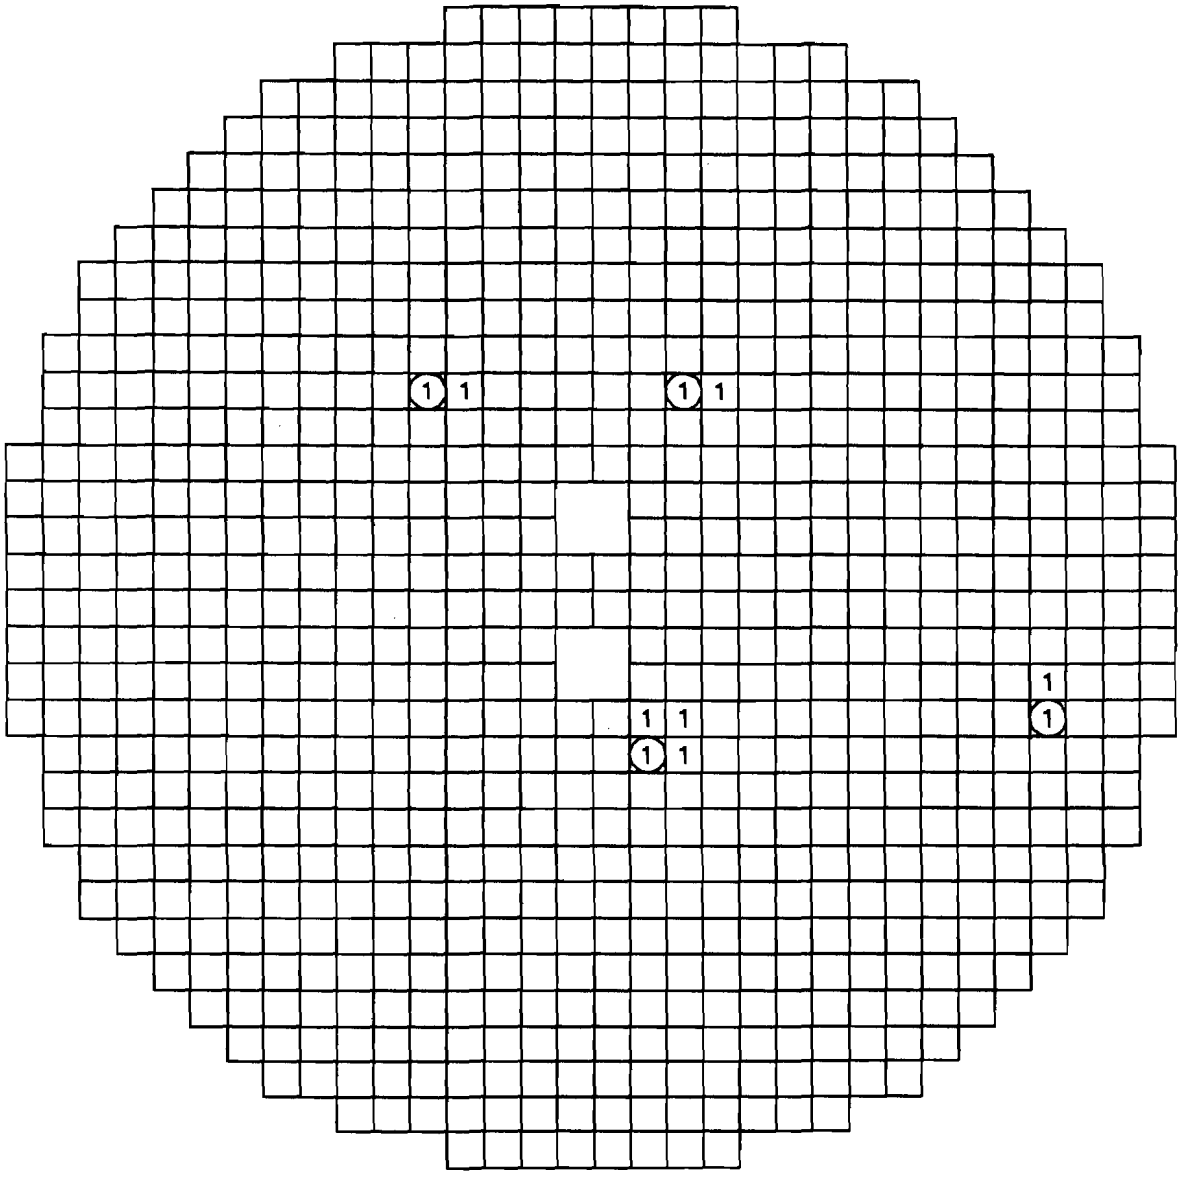

Figure 5.8: Event from Figure 5.6 seen by the Hardware Cluster Finder.

blocks, and the HCF determines how many islands there are above this threshold. The first block encountered in each cluster is considered to be the "seed" block and is marked with a circle in the figure. Since the HCF uses a much more narrow gate than the regular ADC's, it is much less likely to find accidental clusters, and the off-line search is therefore improved by looking only near the seed blocks.

It might be of some interest to note that the local maximum definition of a cluster makes it possible to use a parallel architecture for a cluster finding trigger processor. For each block, eight energy comparisons would be made, and an output signal formed to indicate whether a local maximum exists or not. The comparisons could be made either between the digital output from fast ADC's or between analog signals obtained by integrating the charge from the phototubes. Some straightforward logical circuitry would then sum all the individual block signals to get the number of clusters. This procedure would be quite fast, especially if an analog energy comparison was chosen.

5.4.2 Cluster Energy and Position

The get the final cluster energy, it is necessary to make a number of corrections to the basic three-by-three sum. Most of the corrections are based on detailed studies of electromagnetic showers using the EGS Monte Carlo. First of all, there are several ways in which the three-by-three sum does not include all the energy. If a cluster is close to the beam holes or the outside edge of the array, one or more of the nine blocks is missing. Block energies below the sparse read-out threshold are lost, and so is the small amount of energy deposited outside the three-by-three array. In all of these cases, the missing energy is predicted from the EGS studies and added to the cluster energy. Secondly, if two clusters overlap, the energy in the shared blocks has to be divided between them. This can also be done using detailed knowledge of the shower shape. The most important correction is due to the non-linear response caused by the light attenuation inside the block and the shower leakage at the end of the block. A photon with higher energy has a shower maximum deeper inside the block, resulting in less attenuation but more leakage. There is also a difference between electrons and photons since the latter need an

initial conversion before showering and therefore has a deeper shower maximum. The sizes of these effects are determined from the EGS studies. There is also a correction for the loss of energy if the shower starts in the B or C hodoscopes, which are located right in front of the calorimeter. If a cluster position is found to correlate well with hits in the B and C counters, the cluster energy is increased by three percent.

The cluster position is determined by comparing the energy sums in the three columns (for x) and the three rows (for y). The position is always taken to be inside the center block (which, after all, is the maximum), but pulled towards the side with the most energy. How much the position is shifted from the middle of the center block towards the edge depends on the ratio of energy in the side column (row) to the energy in the middle column (row). This algorithm is designed to make the distribution of cluster positions uniform across the block. There is also a small final correction to the cluster position due to slight imperfections in the stacking of the lead glass blocks. For electron clusters, the position can be compared with the more precise value obtained by extrapolating the track to the lead glass, and the resolution (x or y) is found to be about three millimeters. Note that the resolution would be much worse if all the energy was contained in one block, since with a block width of 5.8 cm one would get $5.8/\sqrt{12} = 1.7$ cm.

5.4.3 Energy Calibration

The absolute energy determination is done with special calibration data taken periodically throughout the run. During these calibration runs, a thin copper foil inserted in the beam converts photons into electron-positron pairs which are then deflected by the magnets in such a way that all the lead glass blocks are illuminated. The observed cluster energy is compared with the known electron momentum, which is measured in the spectrometer with a precision of about one percent. The energy of a cluster is shared between several blocks, but since a particular block can be combined with other blocks in different ways, the individual block gains are sorted out from an average over all events. The energy resolution is found to be given roughly by $2.5\% + 5\%/\sqrt{E}$ for photons and $1.5\% + 5\%/\sqrt{E}$ for electrons (E in

GeV). The resolution is worse for photons because of the variation in conversion depth (*i.e.*, how far into the block the photon converts) which results in an extra uncertainty in the light output.

Between the calibration runs, the gains are tracked using the flasher events which are taken once a second with a common light source illuminating all the blocks. Gain variations in individual blocks are monitored by assuming that the average gain of all blocks stays constant over time. The average excludes blocks in the inner part of the array near the beams, since the gain of these blocks is known to decrease with time due to radiation damage. During the data taking, all blocks are gain matched to within 5% of each other, which is necessary in order to have the Hardware Cluster Finder work properly. (The gains are held at about 200 ADC counts per GeV.)

Additional checks of the calibration can be made with the electrons from $\pi e \nu$ decays and with the photons from $\pi^+\pi^-\pi^0$ decays. In the latter case, the decay vertex is known from the tracking, and we can calculate the invariant π^0 mass from the observed photon energies in the lead glass. By comparing with the nominal π^0 mass, the calibration can be adjusted. The final check of the absolute energy scale comes from the edge at the regenerator of the $\pi^0\pi^0$ decay vertex distribution. (This edge occurs because of the veto counters at the downstream end of the regenerator.) Since the distance of the $\pi^0\pi^0$ decay vertex from the lead glass is proportional to the overall energy scale, the location of the decay distribution edge provides a direct check of the absolute calibration. It is found necessary to decrease all energies by about 0.5% in order to get a good agreement.

5.5 Photon Vetoes

There are twelve planes of photon veto counters, whose purpose is to detect photons from $\pi^0\pi^0\pi^0$ decays going outside the lead glass calorimeter. See chapter 3.6 for a list of these counters. Some of the vetoes are included in the trigger (chapter 4), but most of them are used in the off-line analysis for background suppression. One might ask, do they really need to be calibrated? The ones that are

in the trigger need some gain matching for the phototubes within each group of counters, but this can be done well enough with muons. Since veto counters are only used to cut events, it is sufficient for the analysis to optimize these cuts, and it does not seem necessary to know exactly how many counts result from a 1 GeV photon. The main reason why we like to have an absolute calibration of the photon vetoes is that we want to study the background from $\pi^0\pi^0\pi^0$ decays with the Monte Carlo. In order to simulate this background correctly, the same cuts have to be applied to data and Monte Carlo events, which means that the absolute energy scale has to be known when the Monte Carlo events are generated.

Most of the photon veto counters are positioned along the side of the detector (see Figure 3.1), in order to catch photons that escape with a relatively large angle. Such photons have quite low energies, typically around a few GeV or less. The Collar Anti counters (which form a collar around the beam holes in front of the calorimeter) and the Back Anti counters (behind the beam holes) see much more energetic photons, since they cover the small angle region. Photons hitting the CA or BA counters have energies similar to or higher than the ones that hit the calorimeter, and the average energy of the latter ones is around 20 GeV. The absolute calibration of the CA and BA counters is done with electrons — either from $\pi e \nu$ decays or from the special electron beam data that is also used for the lead glass calibration. The electron energy is known from its track momentum, which is measured in the spectrometer.

Unfortunately, it is not possible to calibrate all of the photon vetoes with electrons (or muons), since the drift chambers do not cover enough angle to provide the necessary position and momentum determinations. We use instead the large number of $\pi^+\pi^-\pi^0$ decays that are accepted by the charged two-track trigger. If one of the two photons from the neutral pion hits the calorimeter and the other one goes outside to one of the photon veto counters, the decay can be used to calibrate that veto counter. The track information from the spectrometer gives the momenta of the two charged pions and also the decay vertex. Together with the energy of the photon that is in the calorimeter, this is enough information to calculate the energy and direction of the "missing" photon. Requiring the total transverse momentum to

be zero, we get the transverse momentum of the missing photon, and by assuming that the total invariant mass is equal to the kaon mass, we find a solution for the longitudinal momentum. (There are usually two solutions, but one can pick the one that makes the photon-photon mass be closest to the neutral pion mass.)

This somewhat complicated reconstruction scheme can be checked with Monte Carlo $\pi^+\pi^-\pi^0$ events. Figure 5.9 shows the energy distribution of the photons that go outside the calorimeter. (This is the energy generated by the Monte Carlo.) Using the algorithm specified above, the missing photon is reconstructed from the other three particles, and the generated and reconstructed energies are compared. The ratio is plotted in Figure 5.10. The reconstruction seems to work reasonably well, even though the average is shifted slightly from one. The resolution of the reconstructed energy is not very good, because the energy of the missing photon is small and is calculated as the difference between two larger numbers.

For the real data events that are used in the calibration, the invariant photon-photon mass is shown in Figure 5.11. This plot is made after a few other cuts have been applied in order to clean up the event sample. The mass is required to be within 10 MeV of the true π^0 mass. For the remaining events, we plot the

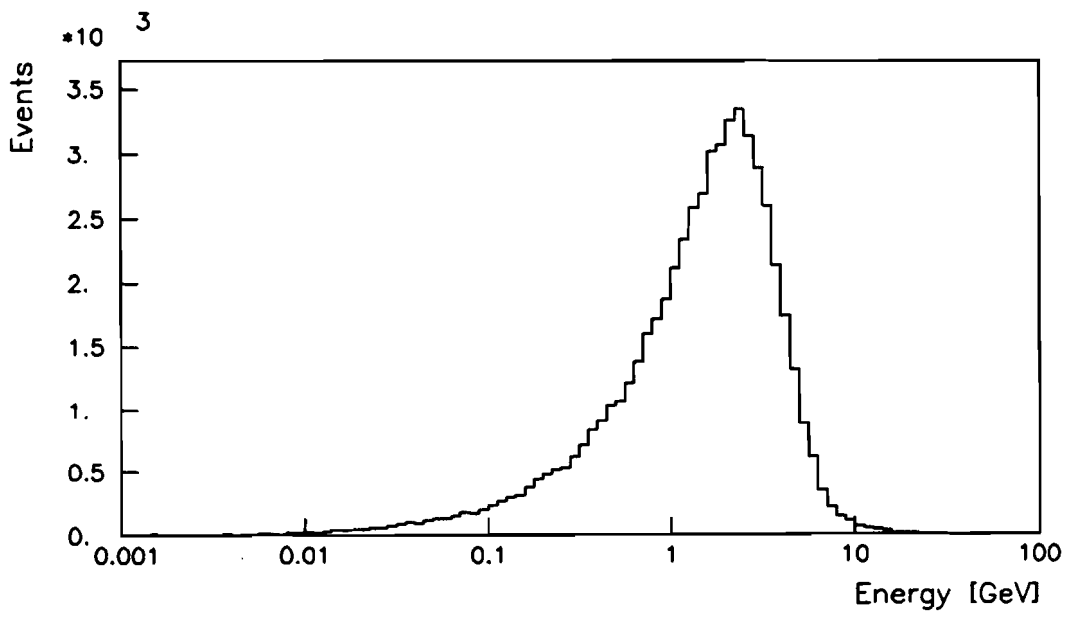

Figure 5.9: Energy of photons missing the calorimeter. (Monte Carlo)

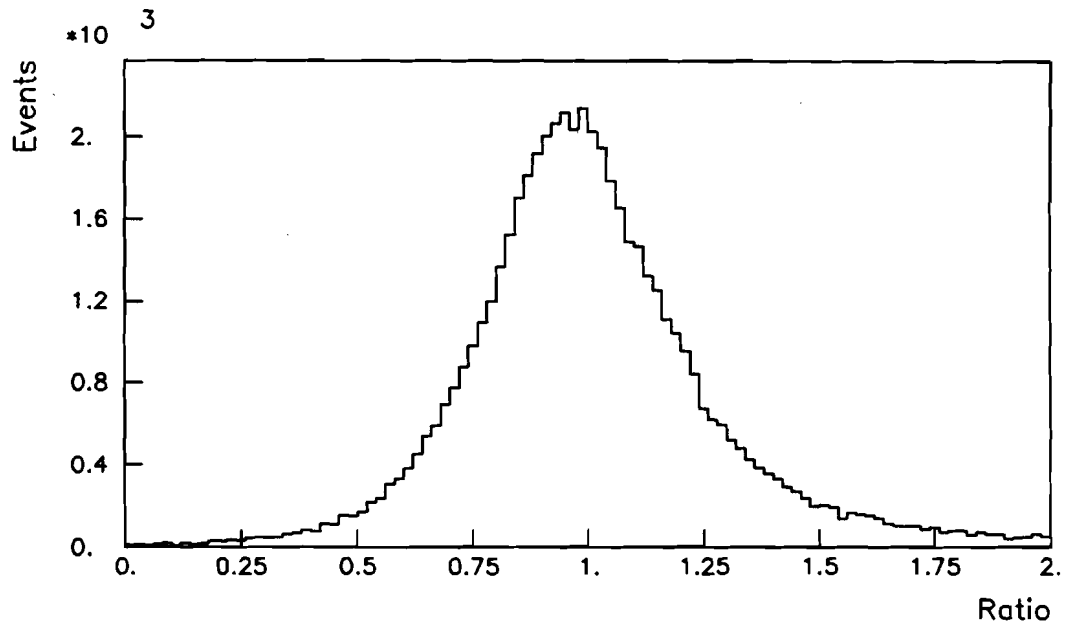

Figure 5.10: Ratio of reconstructed and generated energies. (Monte Carlo)

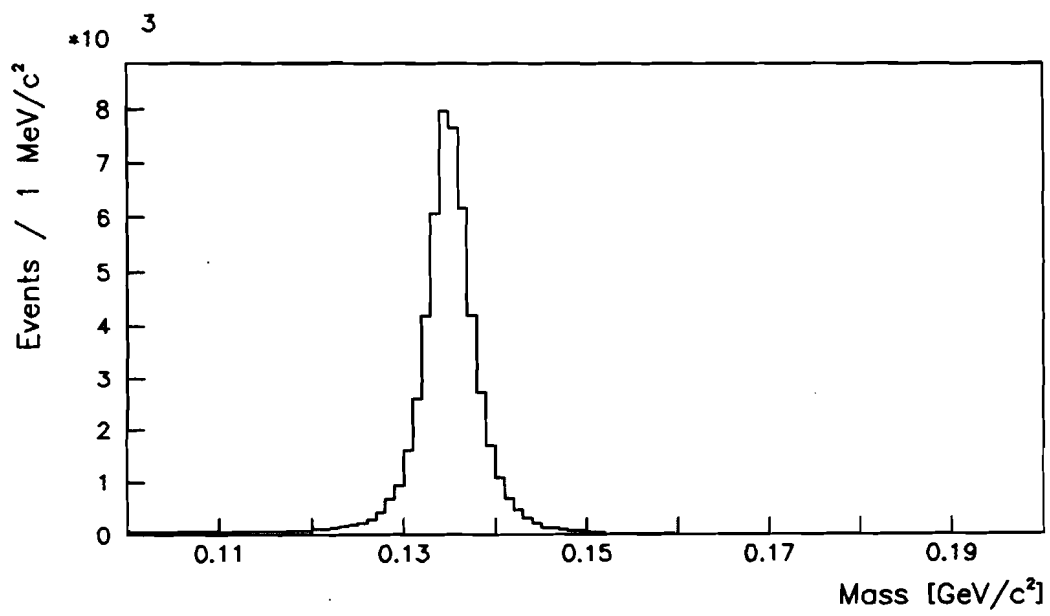

Figure 5.11: Invariant photon-photon mass.

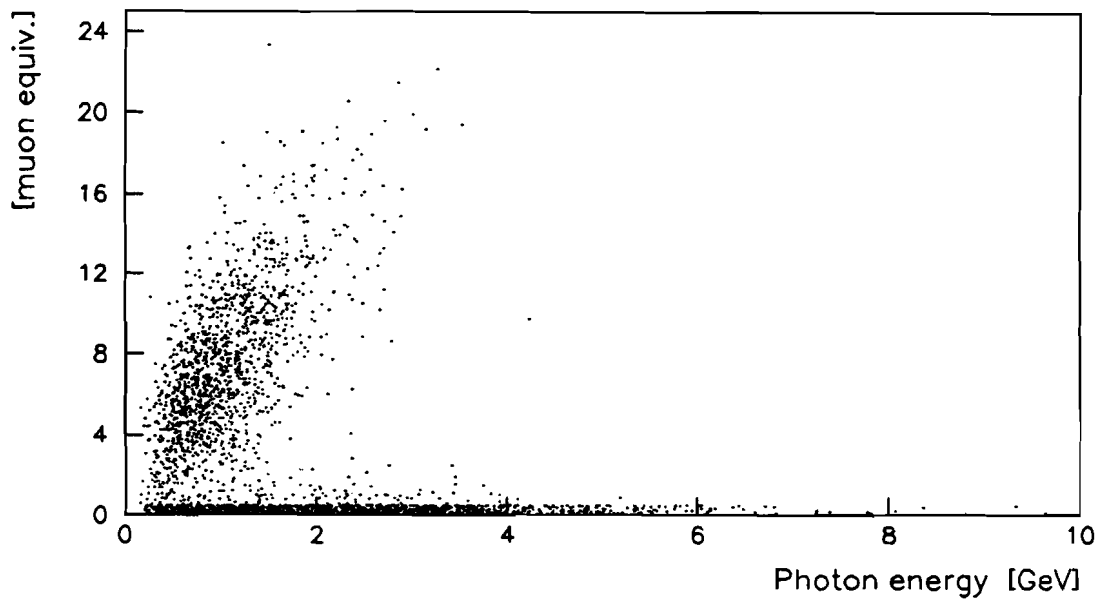

Figure 5.12: Observed energy versus the predicted energy.

observed number of counts in the photon veto counters against the predicted energy of the missing photon. Figure 5.12 shows such a plot for the VA3 counters. The observed energy is expressed in "muon equivalents," which roughly correspond to the number of counts that result from a minimum ionizing particle. The events with almost zero observed energy are events where the missing photon did not actually hit the counter that was predicted from the reconstructed momentum vector. These events are eliminated by requiring the observed energy to be above a certain minimum. The counter gains are then determined from the ratio of observed energy to predicted energy.

Chapter 6
DATA REDUCTION

As in most high energy experiments, there is a huge amount of raw data, and it becomes necessary to do the data reduction in several steps. During the data taking (from 7/87 to 2/88), we wrote 5,000 tapes of raw data. These are standard 9-track open-reel magnetic tapes with a density of 6,250 bpi (bits per inch per track) and a length of about 2,400 feet. As mentioned at the beginning of chapter 4, the analysis and results presented in this thesis are based on 20% of the total data sample (*i.e.*, about 1,000 tapes). For this particular data set, all triggers were recorded simultaneously. Each of these 1,000 tapes has about 50,000 raw data events from the charged trigger (chapter 4.3.1) and another 50,000 events from the neutral trigger (chapter 4.3.2). There is also a small number of events from the other miscellaneous triggers such as the muon and accidental triggers (chapter 4.3.3). One tape corresponds to 25–30 minutes of data taking (*i.e.*, 25–30 spills).

Only a small fraction of the raw data events are $\pi^+\pi^-$ and $\pi^0\pi^0$ decays. Most of the charged triggers are $\pi e\nu$ and $\pi^+\pi^-\pi^0$ decays, and most of the neutral triggers are $\pi^0\pi^0\pi^0$ and $\pi^+\pi^-\pi^0$ decays. In order to separate the different event categories from each other and to eliminate events that do not reconstruct well, the raw data tapes go through a first-pass analysis which is also known as the "crunch." This took place during the summer of 1988, and the total size of the data sample was reduced to about 200 tapes. The reduction in the number of tapes is partially due to the elimination of bad events that are of no interest, but it is also due to a compression of the data. This compression (by a factor of 2–3) is possible because the data acquisition system — which has to work fast — does not write the data to tape in the most compressed form possible. The main benefit of the crunch is the separation of the various event categories. After the separation, there are

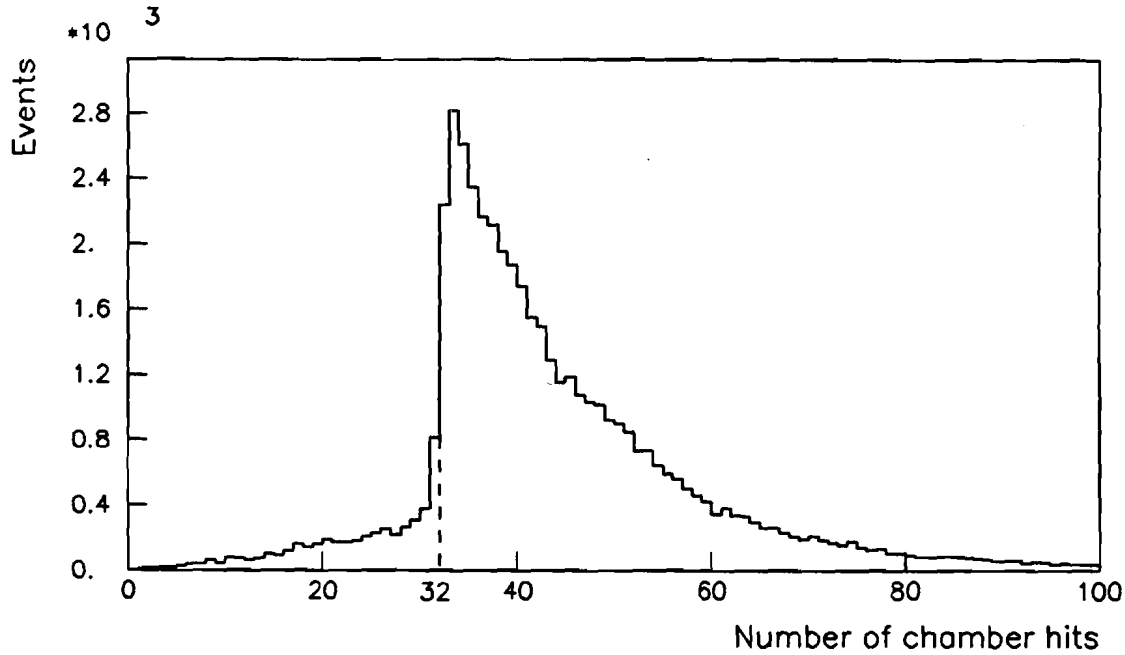

Figure 6.1: Number of drift chamber hits for charged triggers.

only fourteen tapes each for the $\pi^+\pi^-$ and $\pi^0\pi^0$ categories, and these tapes are then processed further in order to extract the final signal. The other tapes, which contain events like $\pi e\nu$ and $\pi^0\pi^0\pi^0$, are also analyzed further in order to study the detector and to check the calibration.

This chapter describes in detail the reconstruction and analysis of the $\pi^+\pi^-$ and $\pi^0\pi^0$ decays (charged and neutral mode, respectively), from which the phases of η_{+-} and η_{00} are extracted. Only the final event selection is described here.

6.1 Charged Mode Analysis

The charged trigger is designed to find events with two charged particles (see chapter 4.3.1). Figure 6.1 shows the total number of hits in the drift chambers for raw data events from this trigger. Since there are sixteen wire planes in the spectrometer, an ideal two-particle event will have 32 hits, and this is indicated by the dashed line. (The bin immediately to the right of the dashed line has exactly 32 hits.) It follows from the figure that most of the triggers do have two tracks and also that there are not very many extra hits from accidental activity or

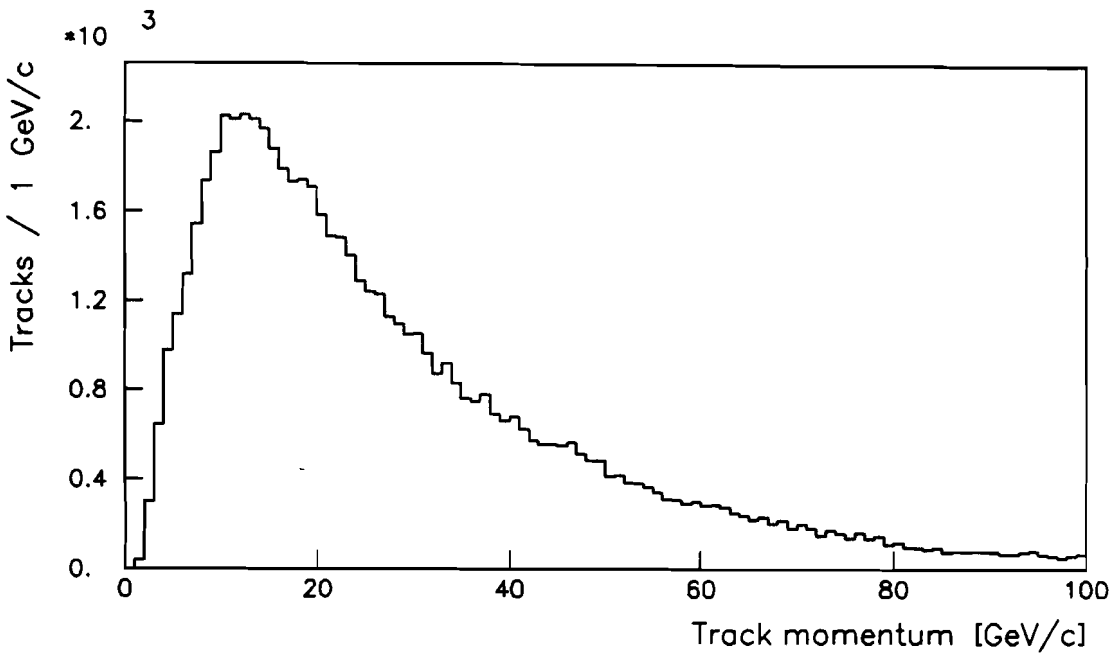

Figure 6.2: Track momentum for two-track events (two entries per event).

noise. How to determine the particle trajectories from the wire hits and how to calculate the momenta of the tracks were discussed in chapter 5.3. For the $\pi^+\pi^-$ analysis, we require that the events have exactly two tracks and that these two tracks have opposite charge (*i.e.*, that they are bent in opposite directions by the analysis magnet). About 67% of the raw triggers satisfy these conditions and will be referred to as two-track events. For about 8% of the triggers, the number of chamber hits is too large for the reconstruction code. Around 23% have less than two complete tracks, and only 2% have more than two tracks. The sum-of-distance cut (see 5.3.2) eliminates out-of-time tracks and reduces the chance of finding extra accidental tracks. Out of the events that have exactly two tracks, only half a percent do not have oppositely charged tracks.

6.1.1 Reconstruction of Two-Track Events

We will first look at how the two-track events are reconstructed. The figures that illustrate this reconstruction are all made from one raw data tape. (The full data sample is about 1,000 raw data tapes.) Figure 6.2 shows the distribution

of track momenta for two-track events. The momentum resolution is about one percent, which is mostly due to multiple scattering. The internal chamber resolution (100 micrometers per plane) leads to an uncertainty that is proportional to the momentum and is about 0.3% for a 20 GeV/c track.

The first step in the reconstruction is to find the decay vertex by projecting the two tracks upstream to their intersection point. Due to the finite tracking resolution, the tracks do not actually intersect, and we find instead the point of closest approach. The decay vertex is taken as the midpoint of the shortest line segment connecting the two tracks. Since the tracks are almost parallel with the z axis (*i.e.*, the beam direction), this line segment is nearly perpendicular to the z axis. The length of the segment (also called the distance of closest approach) is typically 1–2 millimeters, which means that the transverse vertex resolution is of the order of a millimeter. The z component of the vertex, however, has a resolution of about 20 centimeters, due to the small angle between the tracks and the z axis. It is possible to do a more sophisticated vertex fitting. Tracks with low momentum, who are more likely to scatter, and tracks with missing hits could be assigned lower weights. Such an algorithm would pick another point than the midpoint on the shortest line segment between the tracks, and this might improve the transverse resolution somewhat. However, in the case of two tracks, the z component of the vertex would not change.

Assuming now that the two tracks are the only final state particles, we obtain the momentum of the particle that decayed by summing the two track momenta: $\mathbf{p} = \mathbf{p}_1 + \mathbf{p}_2$. The total momentum vector $\mathbf{p}$ is projected backwards in order to find the transverse position at the regenerator. By comparing this position with the known location of the regenerator (up or down), the event is assigned either to the regenerated beam or to the vacuum beam. For example, if the sign of the y component is positive and the regenerator is in the lower position, the event comes from the vacuum beam. (The regenerator moves between the top and bottom beams every spill.) It should be stressed that exactly the same cuts are applied to all events regardless of their origin, and that the transverse position is only used for separating the regenerated and vacuum event samples from each other. For good

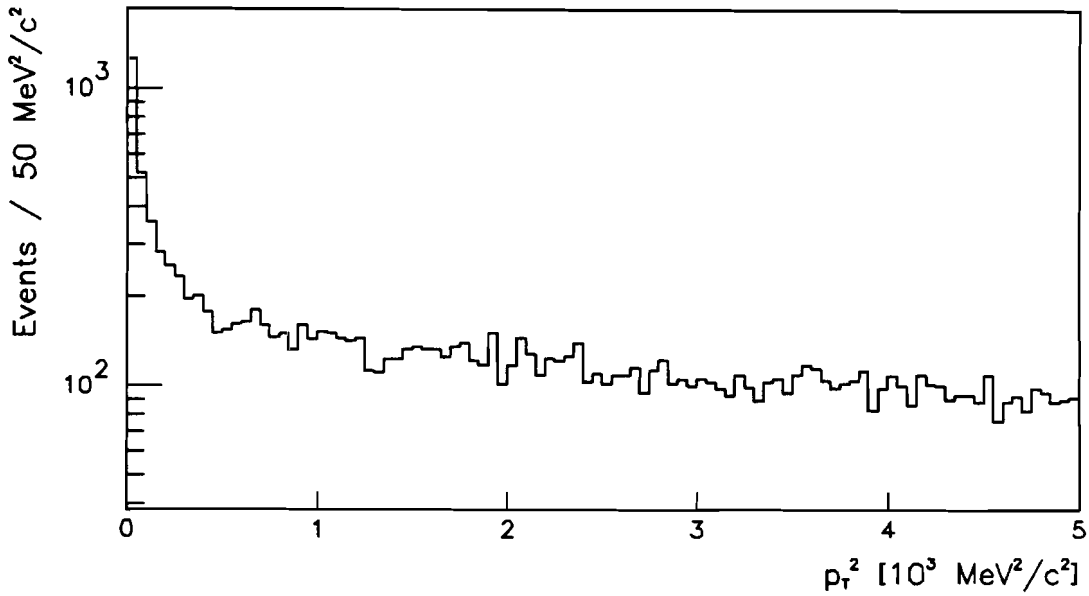

Figure 6.3: Transverse momentum squared for all two-track events on one raw data tape (G set data).

$\pi^+\pi^-$ events, the two beams are very well separated (see Figure 6.19), and there is no ambiguity in selecting which beam the event comes from. But events with an extra neutral particle (*e.g.*, $\pi^+\pi^-\pi^0$ and $\pi e\nu$ decays) will often project outside the beams, unless the transverse momentum of the neutral particle is taken into account. For these event categories, the transverse vertex position can be used instead.

The transverse momentum of the two-track system is calculated with respect to a line from the target to the projected position at the regenerator. The transverse momentum is equal to the transverse component of $\mathbf{p}$ with respect to that line. Figure 6.3 shows p_T^2 for all two-track events. Most of the events with non-zero transverse momentum are $\pi^+\pi^-\pi^0$ and $\pi e\nu$ decays, but there is also a contribution from $\pi^+\pi^-$ decays of kaons that scattered incoherently at the regenerator.

We now introduce the hypothesis that the two-track event is a $K \to \pi^+\pi^-$ decay. The kaon energy is then given by $E = \sqrt{p_1^2 + m_\pi^2} + \sqrt{p_2^2 + m_\pi^2}$, where m_π is the $\pi^\pm$ mass and $p_i = |\mathbf{p}_i|$, and the invariant mass is obtained directly from $m^2 = E^2 - \mathbf{p}^2$ (double precision recommended). Since m_π is much smaller than p_1

and p_2, it is more instructive to write m^2 to first order in m_π^2/p_i^2:

$$m^2 = m_\pi^2(2 + \frac{p_1}{p_2} + \frac{p_2}{p_1}) + 2p_1p_2(1 - \cos\theta)$$

Here, θ is the angle between $\mathbf{p}_1$ and $\mathbf{p}_2$. The invariant mass is shown in Figure 6.4. The number of $\pi^+\pi^-$ decays is comparable in the two beams, but due to the big difference in K_L flux, there are about fifteen times more $\pi^+\pi^-\pi^0$ and $\pi e\nu$ decays in the vacuum beam than in the regenerated beam. Of these two decay modes, it is only the $\pi e\nu$ mode that is really a background to the $\pi^+\pi^-$ signal. The two-track invariant mass for $\pi^+\pi^-\pi^0$ decays is at least a π^0 mass below the K mass. There is also a small contribution to the background from $\Lambda \rightarrow p\pi^-$ and $\overline{\Lambda} \rightarrow \overline{p}\pi^+$ decays. The following section will describe the cuts that are used to separate the $\pi^+\pi^-$ signal from the other decay modes.

6.1.2 Identification of the $\pi^+\pi^-$ Signal

To suppress the background from $\pi e\nu$ decays, it is necessary to distinguish electrons from pions, and this is done by using the lead glass. For tracks that match a cluster in the lead glass, we calculate the ratio of the cluster energy to the track momentum, and this ratio is known as the track E/p ("E over p"). That a track matches a cluster means that the projected track position at the lead glass is within 7 cm of the cluster position. If a track does not match a cluster, E/p is considered to be zero. Figure 6.5 shows the E/p distribution for all two-track events (or rather, for all tracks in two-track events that match a cluster). The electrons from $\pi e\nu$ decays deposit all their energy in the lead glass and show up as the peak at $E/p = 1$. Pions, on the other hand, do not always shower in the lead glass. The broad distribution is due to the pions that do shower, while the peak near zero corresponds to pions that only deposit a minimum ionizing signal. The electrons are rejected by requiring both tracks to have an E/p less than 0.8.

Because the trigger hodoscopes and the drift chambers are rectangular in shape (essentially squares), some tracks project into one of the four corners not covered by the circular lead glass. There are also a number of tracks that go through the beam holes in the lead glass. Since the E/p cut does not work for

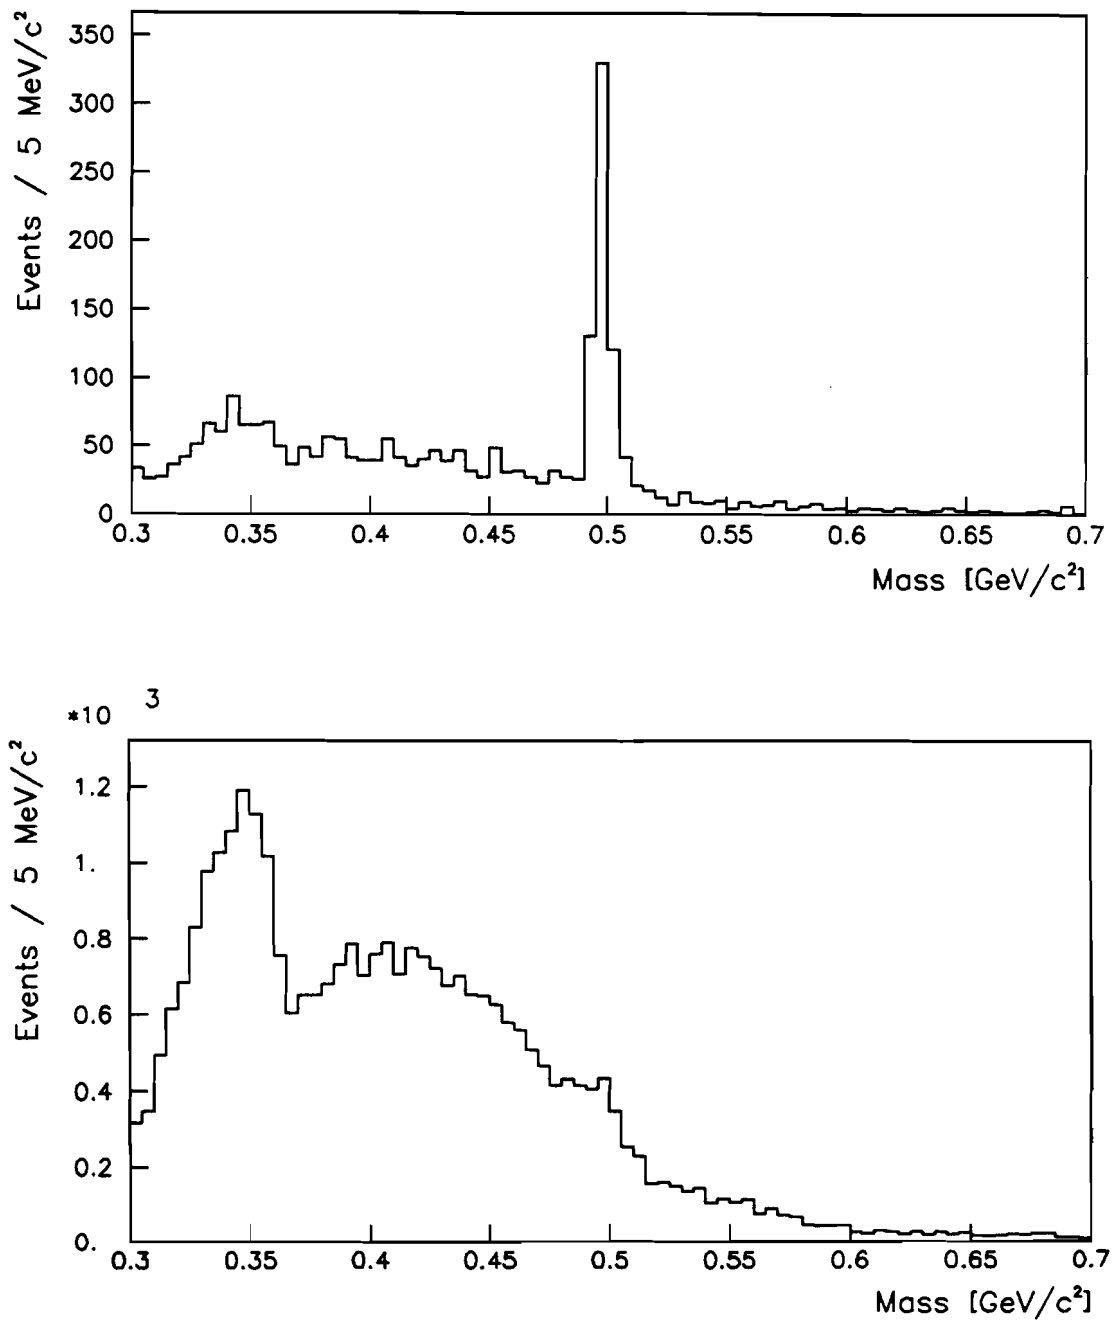

Figure 6.4: Invariant $\pi^+\pi^-$ mass for all two-track events on one raw data tape. Regenerated beam (*top*) and vacuum beam (*bottom*).

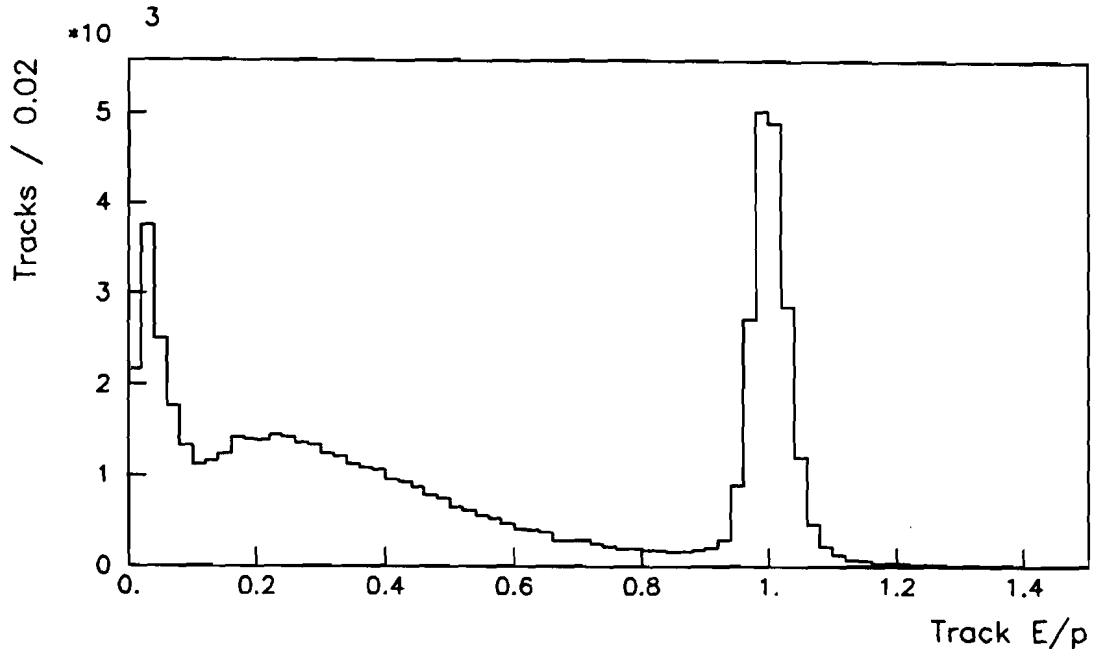

Figure 6.5: Energy from the calorimeter divided by the track momentum. One entry for each track that matches a cluster.

those tracks, we require that all tracks hit the lead glass. This geometric cut is illustrated by the dashed line in Figure 6.6. The tracks are required to be half a block away from the edge, in order to ensure a good energy resolution. One might wonder why we do not keep all tracks and accept a higher background. In fact, since most of the tracks hit the lead glass, the background would not be intolerably high. The problem is that the E/p cut also removes the good pion tracks that are in the tail of the pion E/p distribution, and this means that the acceptance for good $\pi^+\pi^-$ decays would depend on whether the tracks hit the lead glass or not. To include this effect in the Monte Carlo, it would be necessary either to simulate the pion showers or to find an empirical constant describing what fraction of good events are killed by the E/p cut. We prefer to use the geometric cut and loose some signal.

Another complication with the E/p cut occurs if the two tracks have almost the same x position at the lead glass. In this case, the matching of x and y tracks is not very reliable, and it is not clear which track momentum (p) belongs to which cluster energy (E). For this reason, we cut an event if the track x positions at the

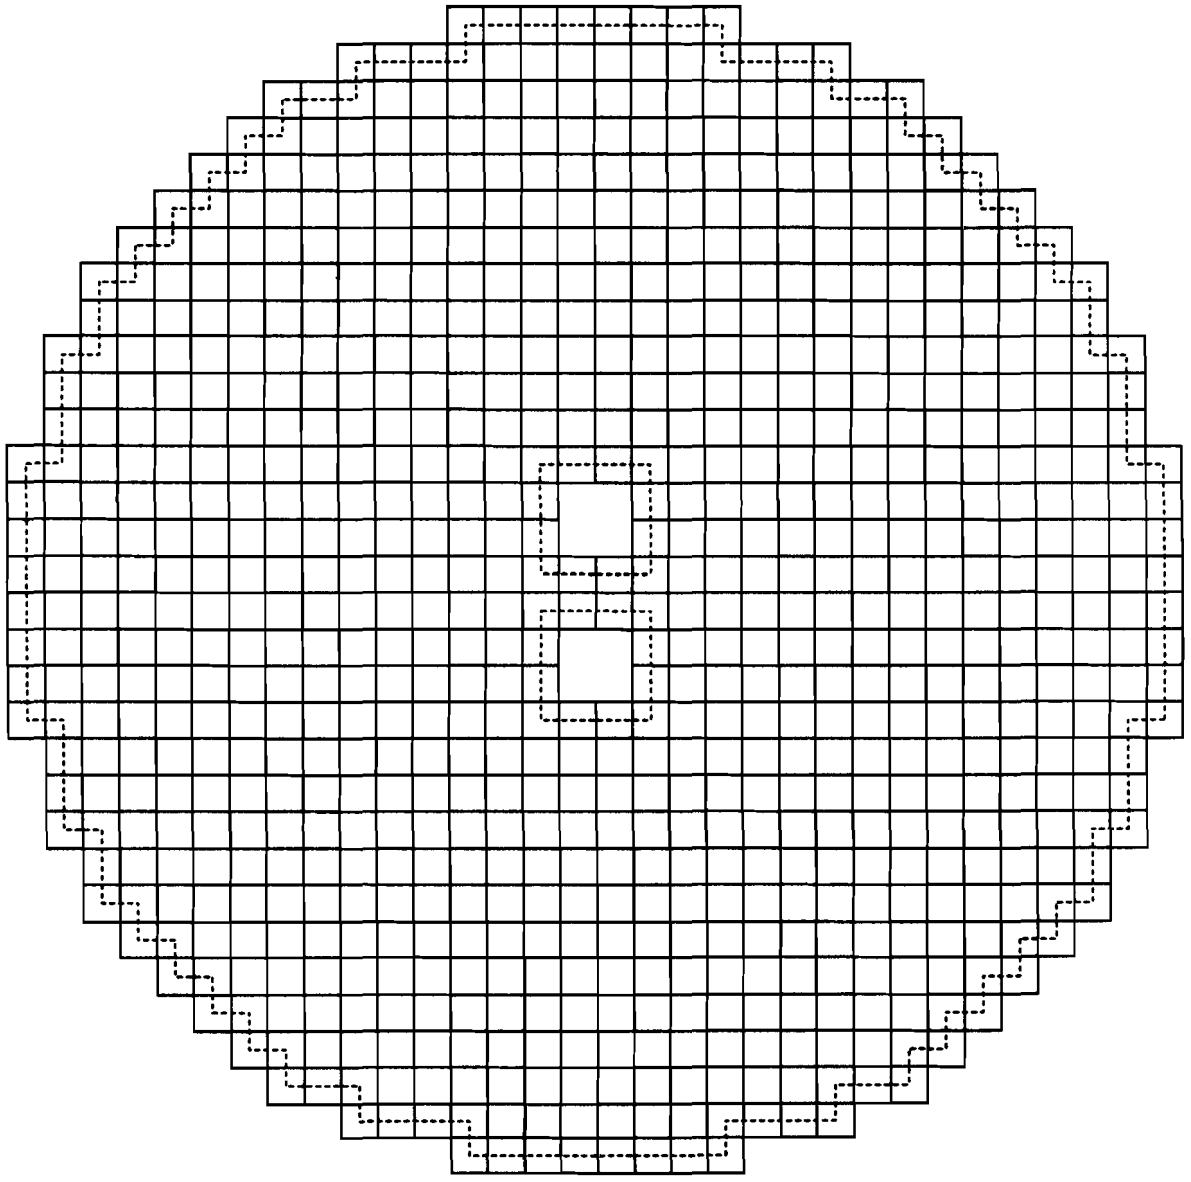

Figure 6.6: Geometric track cut at the lead glass for two-track events.

lead glass are within 2 cm of each other. Note that there is no problem if the y positions are close, since the track momenta are determined from the x tracks.

Most of the $\pi\mu\nu$ decays are vetoed at the trigger level by the muon hodoscope behind the muon filter. The discriminated signals from the muon counters are also latched individually, and just in case there is some inefficiency in the trigger veto, an event is cut if any latch is on. Since low energy muons might not make it through the muon filter, we also require that both track momenta are above 7 GeV/c (see Figure 6.2). This cut removes some signal as well, but once again it is easier to cut than to simulate the probability that a muon is not vetoed.

In order to reject the lambda decays, we also calculate the invariant mass assuming that the two-track event is a $\Lambda \rightarrow p\pi^-$ (or $\overline{\Lambda} \rightarrow \bar{p}\pi^+$) decay. Since the proton and pion masses are different, it is necessary to choose which track is the proton. It turns out that it is safe to assume that the track with the highest momentum is the proton. This is because there is not much energy available in the center of mass frame for the proton and the pion, and their momenta in the lab frame are therefore dominated by the boost from the lambda frame to the lab frame. Since the proton has a much higher mass, it also has the higher momentum. It is then possible to tell the difference between a Λ and a $\overline{\Lambda}$ by checking which way the proton is bent by the magnetic field. Denoting the higher track momentum by p_p and the lower by p_π, the energy of the lambda is given by $E = \sqrt{p_p^2 + m_p^2} + \sqrt{p_\pi^2 + m_\pi^2}$. The invariant mass is calculated from $m^2 = E^2 - \mathbf{p}^2$, but once again it is more instructive to write m^2 to first order in m_p^2/p_p^2 and m_π^2/p_π^2:

$$m^2 = m_p^2(1 + \frac{p_\pi}{p_p}) + m_\pi^2(1 + \frac{p_p}{p_\pi}) + 2p_p p_\pi(1 - \cos\theta)$$

This relation sets a lower limit of 3.1 (and an upper limit of 14.6) for the ratio p_p/p_π. Figure 6.7 shows the invariant $p\pi$ mass for two-track events with an energy above 100 GeV and a track momentum ratio greater than 3. (Only high-energy lambdas will survive until the start of the decay volume.) The lambda peak in the figure agrees very well with the nominal mass of 1.1156 GeV/c^2. Recall from chapter 5 that the absolute scale of the momentum determination is adjusted to make the invariant mass of $K \rightarrow \pi^+\pi^-$ decays be equal to the kaon mass. The fact that the

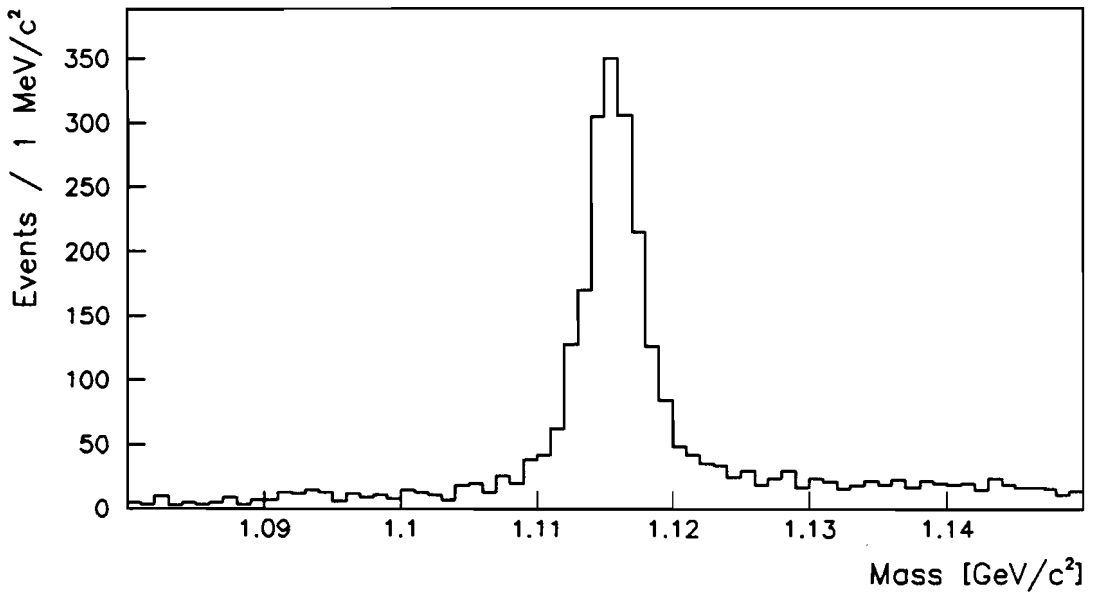

Figure 6.7: Invariant two-track mass for the lambda hypothesis.

invariant mass of the much more asymmetric $\Lambda \rightarrow p\pi$ decays is also right, gives us some confidence in the calibration. The lambda decays are rejected by cutting events with a $p\pi$ mass between 1.10 and 1.13 GeV/c^2 (but only if the momentum ratio is above 3 and the energy above 100 GeV). It should be noted that the two calculated values for the two-track energy (kaon hypothesis and lambda hypothesis) are essentially equal, since the energy is dominated by the track momenta. This means that the lambda cut only affects the data in energy bins above 100 GeV.

After the cuts described in this section, the invariant mass plot looks like Figure 6.8. The $\pi^+\pi^-\pi^0$ decays are easily eliminated by a mass cut. Comparing with Figure 6.4, it is seen that about 40% of the $\pi^+\pi^-$ signal has been lost. This is mostly due to the geometric track cut related to the E/p cut and also to the track momentum cut. After a mass cut, the background in the p_T^2 distribution will be due to incoherently scattered kaons in the regenerated beam and the few remaining $\pi e\nu$ decays in the vacuum beam.

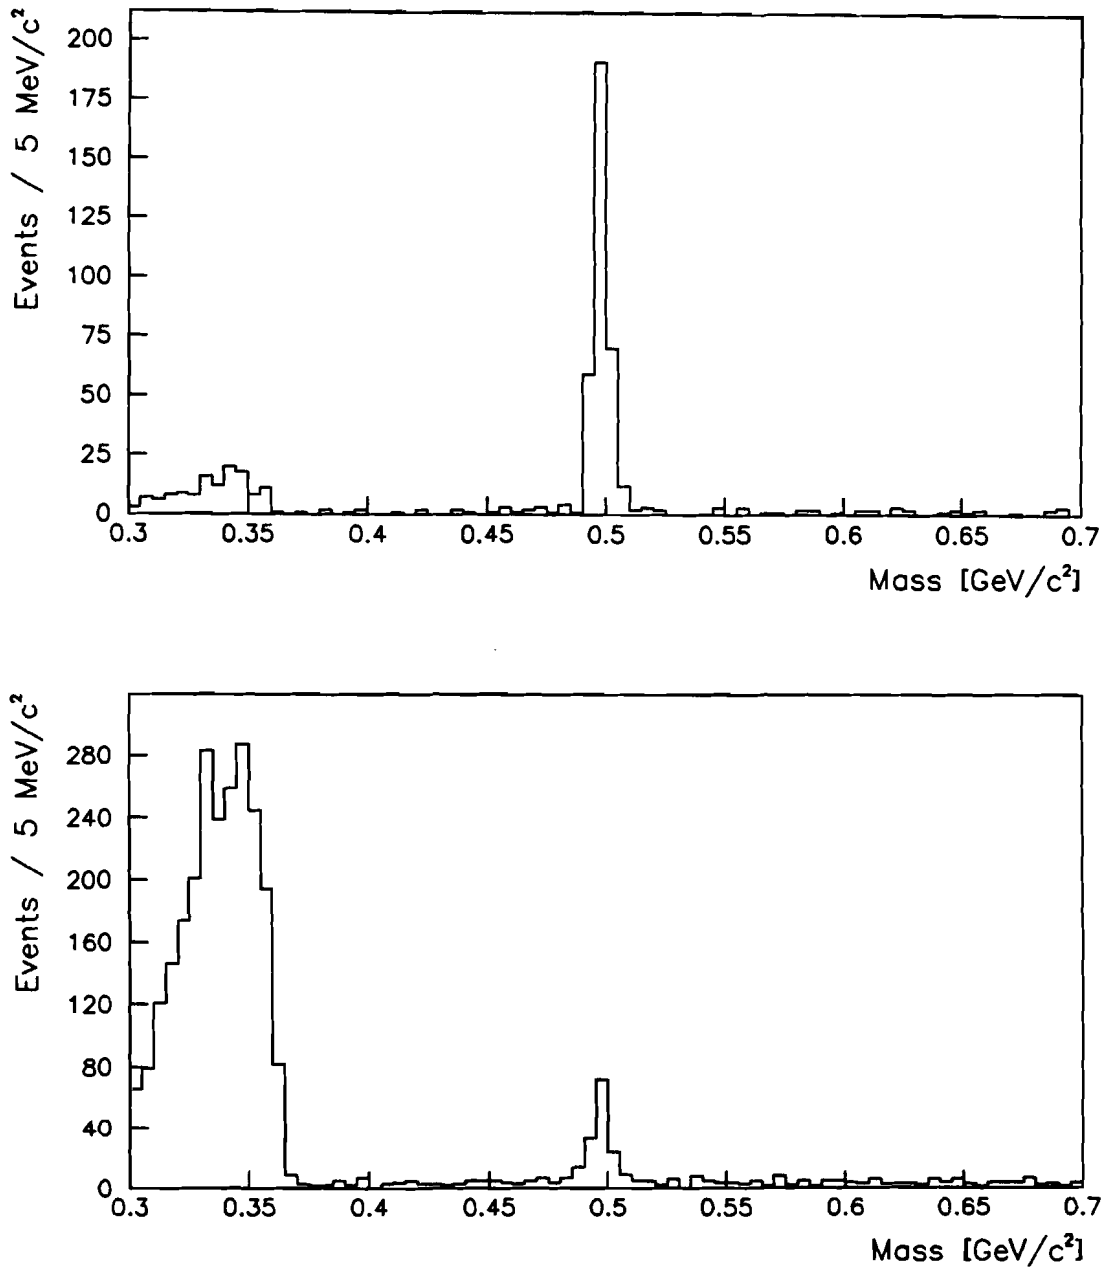

Figure 6.8: Invariant $\pi^+\pi^-$ mass for two-track events (one raw data tape) after the $\pi e\nu$, $\pi\mu\nu$ and and Λ cuts. Regenerated beam (*top*) and vacuum beam (*bottom*).

6.1.3 Miscellaneous Cuts

A number of cuts are used to clean up the $K \rightarrow \pi^+\pi^-$ sample. The figures represent the full data sample, and the cuts are usually indicated by arrows. Figure 6.9 shows the distributions of the various track quality variables after all other cuts. For each upstream and downstream track segment, there is a "chi-square" which is the sum of the squared residuals divided by the number of degrees of freedom for the track fit (see 5.3.3). With a single hit resolution of 100 μm, one should divide by 10^{-8} to get real chi-squares. The magnet offset variable is calculated by projecting the upstream and downstream x tracks to the magnet. Ideally, they should intersect in the middle of the magnet, but in general there is an offset between them due to multiple scattering and the chamber resolution. In order to find the number of standard deviations, the offset is divided by the empirically derived amount $220 + 2700/p$ μm, where p is the track momentum in GeV/c. The vertex quality variable is the distance of closest approach divided by the distance between the vertex and the first drift chamber and by the empirically derived quantity $(24.5 + 8500/p_1^2 + 8500/p_2^2) \times 10^{-6}$, where p_1 and p_2 are the two track momenta in GeV/c.

There are a number of aperture cuts, which are illustrated in Figures 6.10–14 (after all non-aperture cuts). The projections of the two tracks are required to be inside the edges of the circular vacuum window (6.10), the rectangular T and V hodoscopes (6.11), the rectangular aperture of the analysis magnet (6.12) and the rectangular muon hodoscope (6.13). The x projection cut at $z = 162\,\mathrm{m}$ (6.14) ensures that the tracks do not hit one of the frames holding the helium bags.

As mentioned in chapter 4, the charged trigger includes an east-west topology requirement that is based on the hits in the Chamber 2 x planes. In principle, this requires that there is one track on the east side and one on the west side. However, it is possible for an event with both tracks on the same side to pass the trigger, if there is noise or accidental activity on the other side. Extra hits may actually save events that would otherwise not be accepted. Since this never happens in the Monte Carlo, where exactly two tracks are generated, it is necessary to re-apply the east-west condition in the data analysis. By requiring that the condition is satisfied

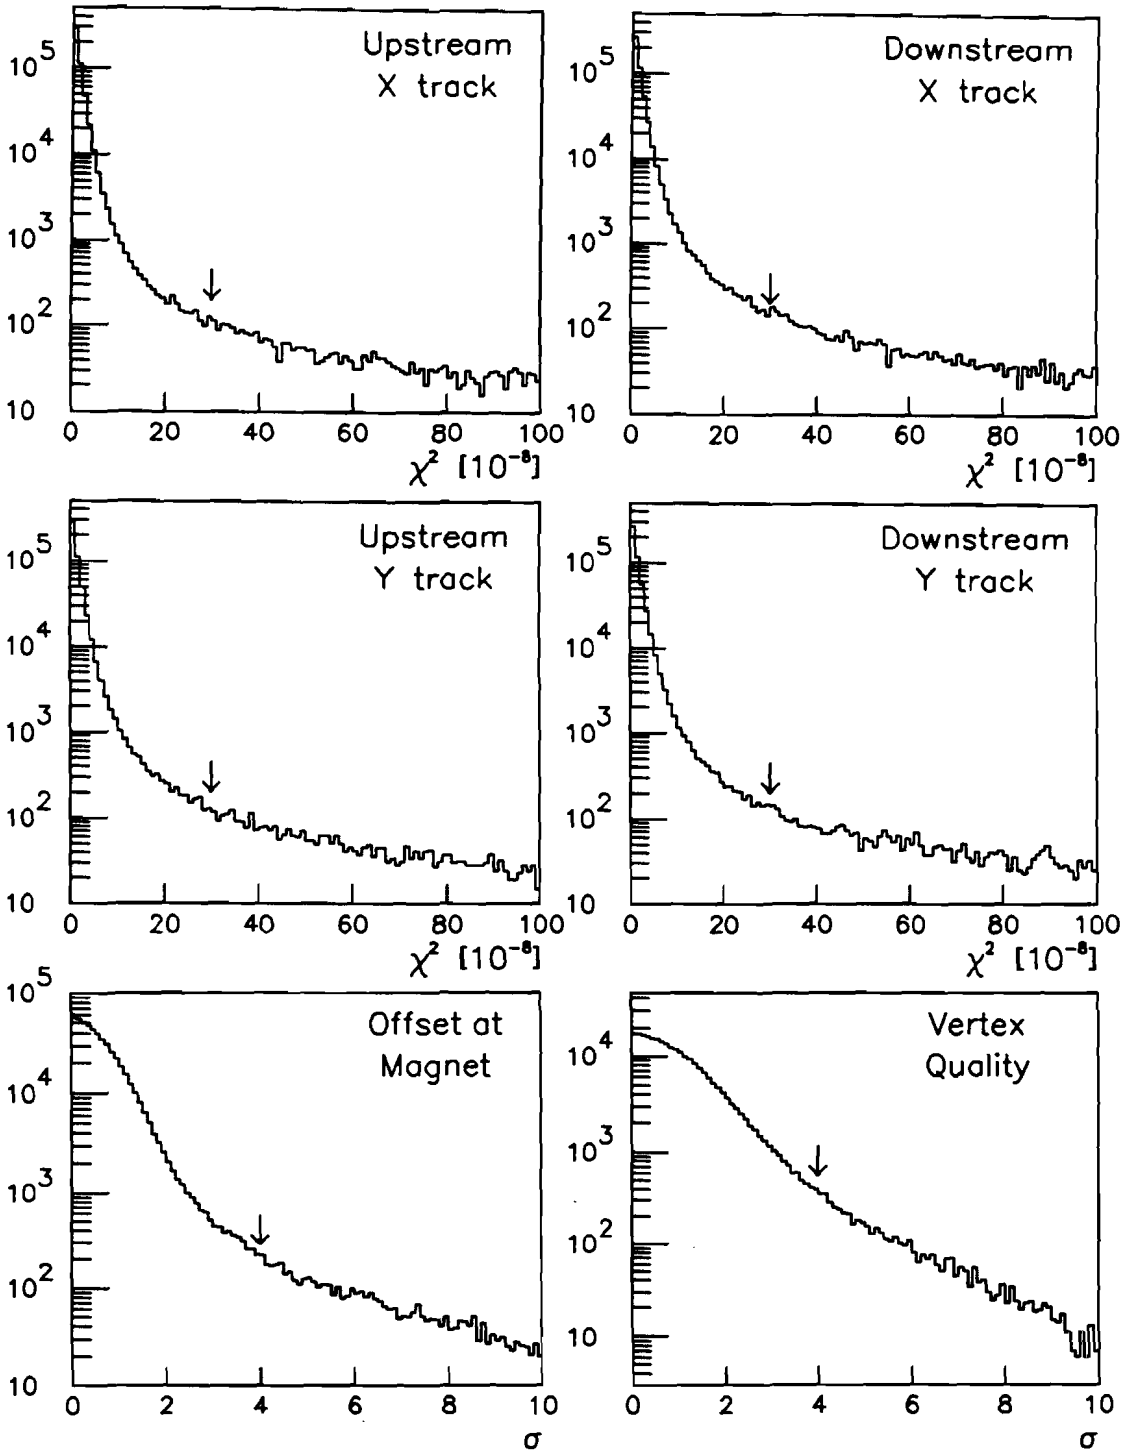

Figure 6.9: Track quality cuts.

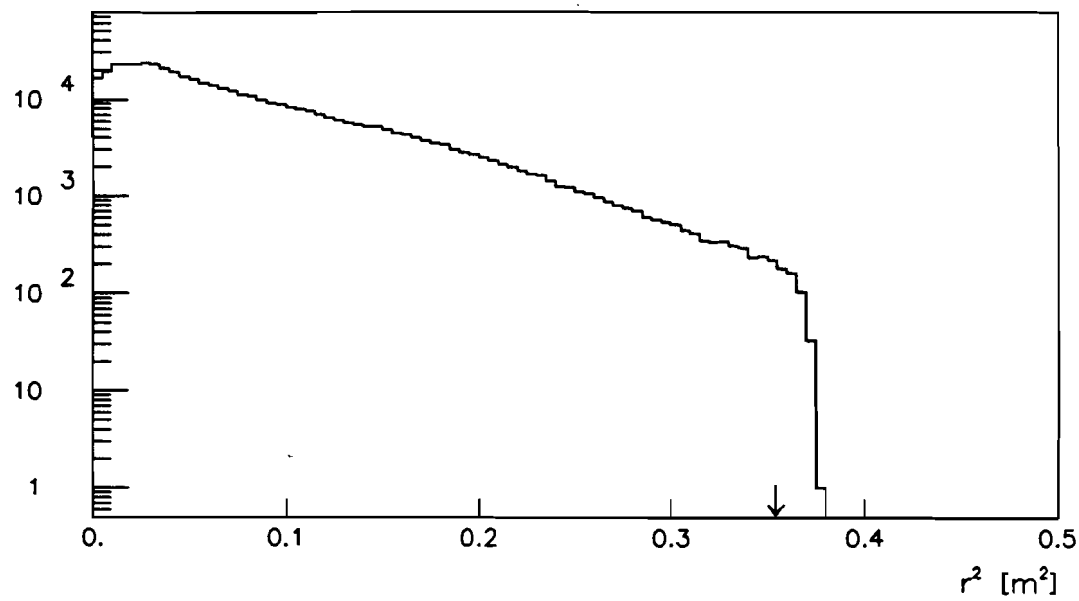

Figure 6.10: Track projection $(x^2 + y^2)$ at the vacuum window (z=159 m).

by the hits that are associated with the two tracks, we make sure that data and Monte Carlo events are treated equally.

The charged trigger has the RA trigger signal in veto, in order to reject kaons that scatter inelastically in the regenerator. The timing of the RA signal is measured by a TDC with a 0.1 ns resolution and shown in Figure 6.15 after all other cuts. The TDC is started by the event trigger and stopped by the RA signal, with a quite arbitrary origin for the time scale. Each peak corresponds to a bucket in the beam structure (see chapter 4.1), and the bucket which is in time with the trigger is empty because RA is in veto. Since the rate of the RA signal is around 600 kHz, most of the entries are overflows. In fact, out of 226,000 entries in the figure, there are about 215,000 overflows. Just in case there is some inefficiency in the RA veto, we cut an event if the time is inside the interval indicated by the arrows.

6.1.4 Final Data Sample and Background Subtraction

After all cuts, we use the invariant $\pi^+\pi^-$ mass (m) and the transverse momentum squared (p_T^2) for identifying the final $K \to \pi^+\pi^-$ sample and subtracting the background. The signal region is defined by $484 < m < 512\,\mathrm{MeV/c^2}$ and

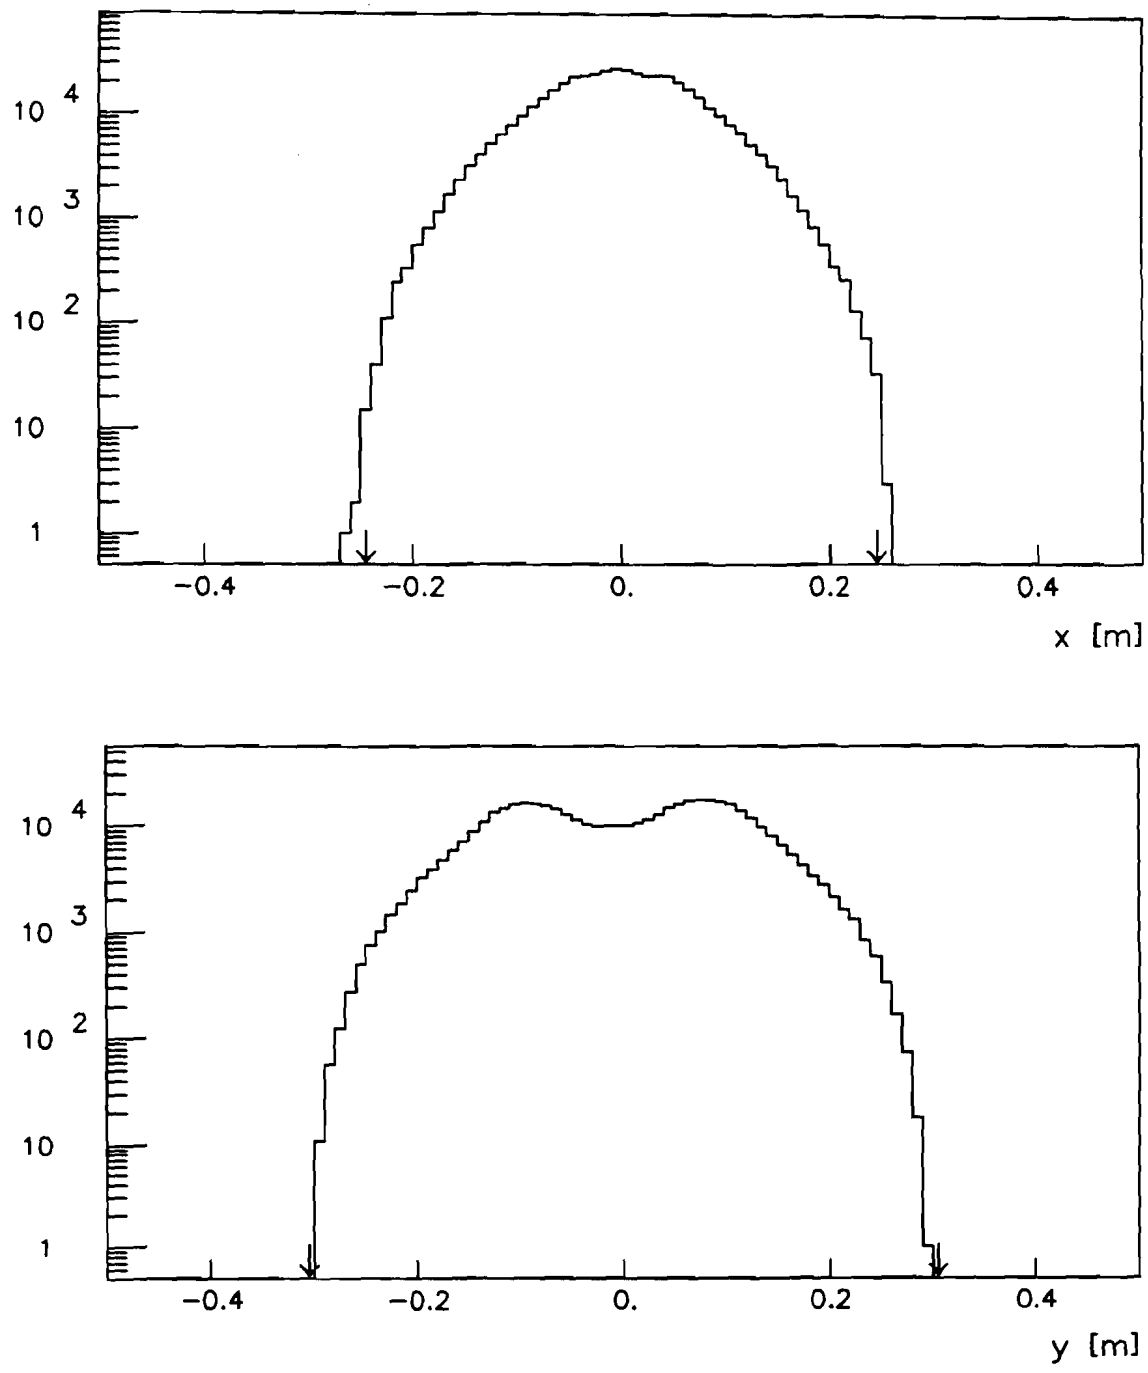

Figure 6.11: Track projection (x and y) at the T and V hodoscopes (z=138 m).

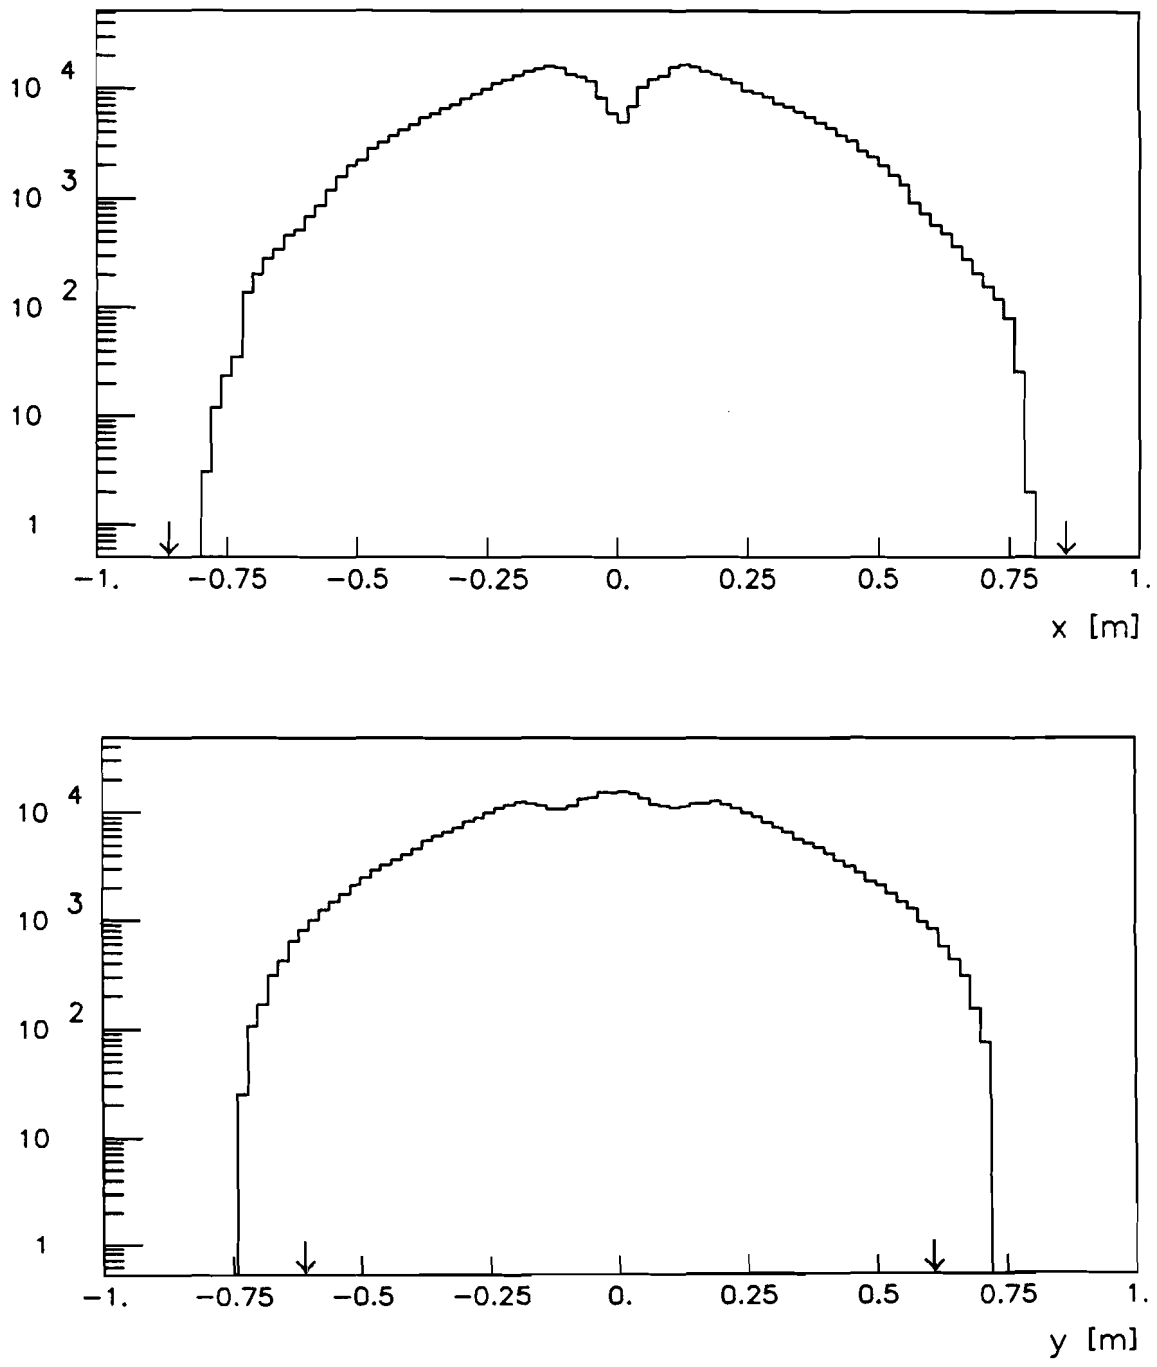

Figure 6.12: Track projection (x and y) at the analysis magnet (z=169 m).

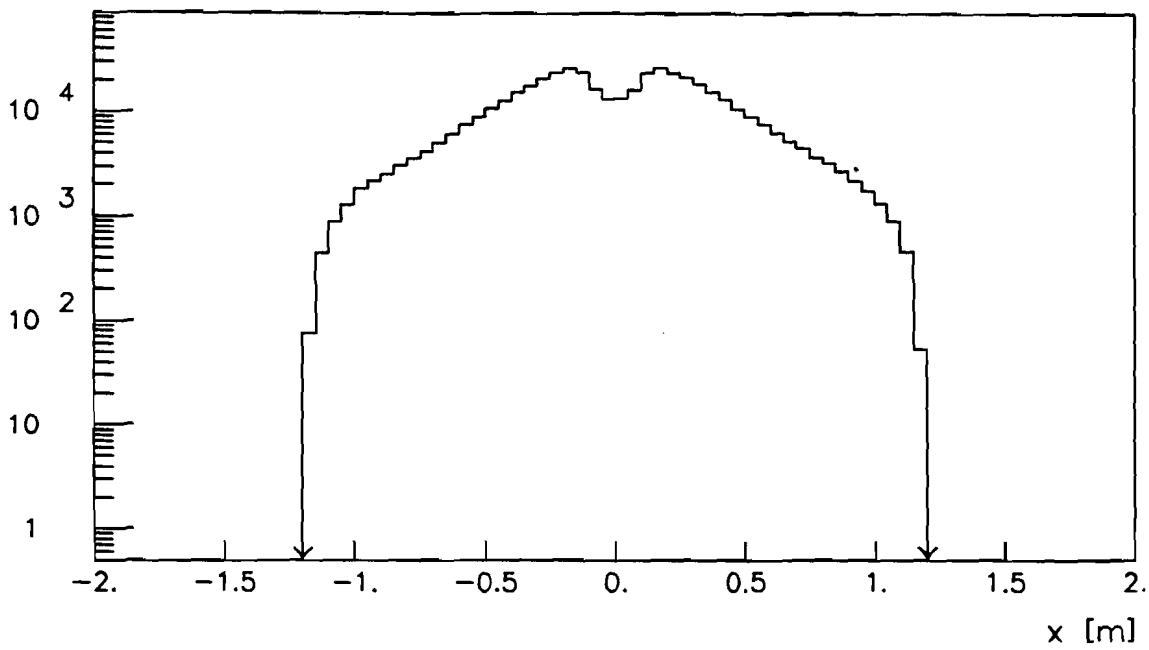

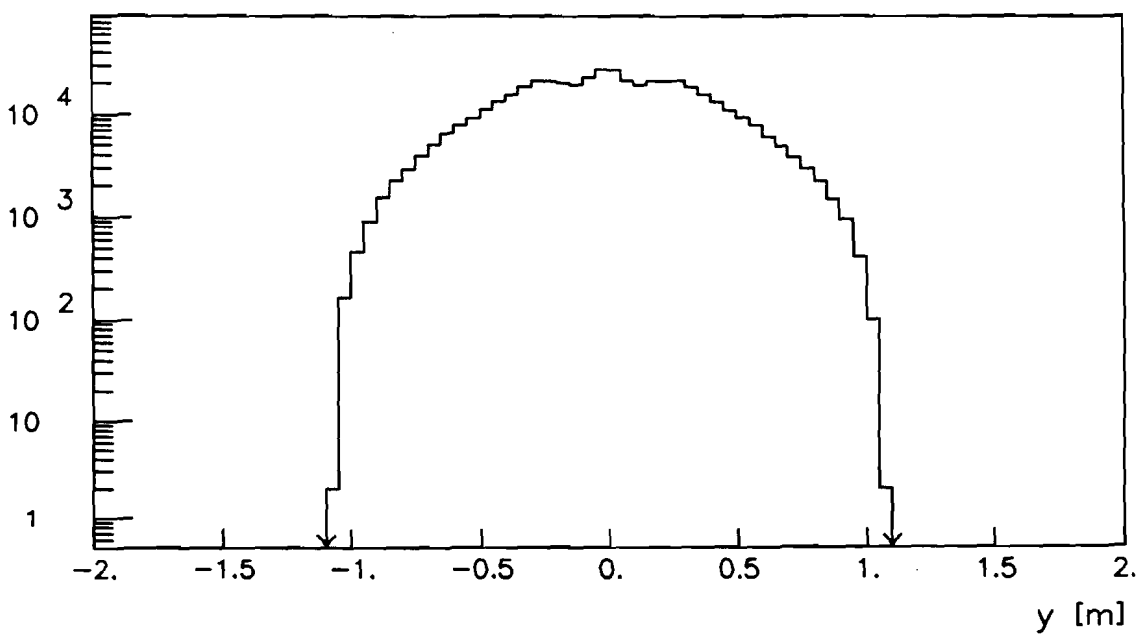

Figure 6.13: Track projection (x and y) at the muon hodoscope (z=190 m).

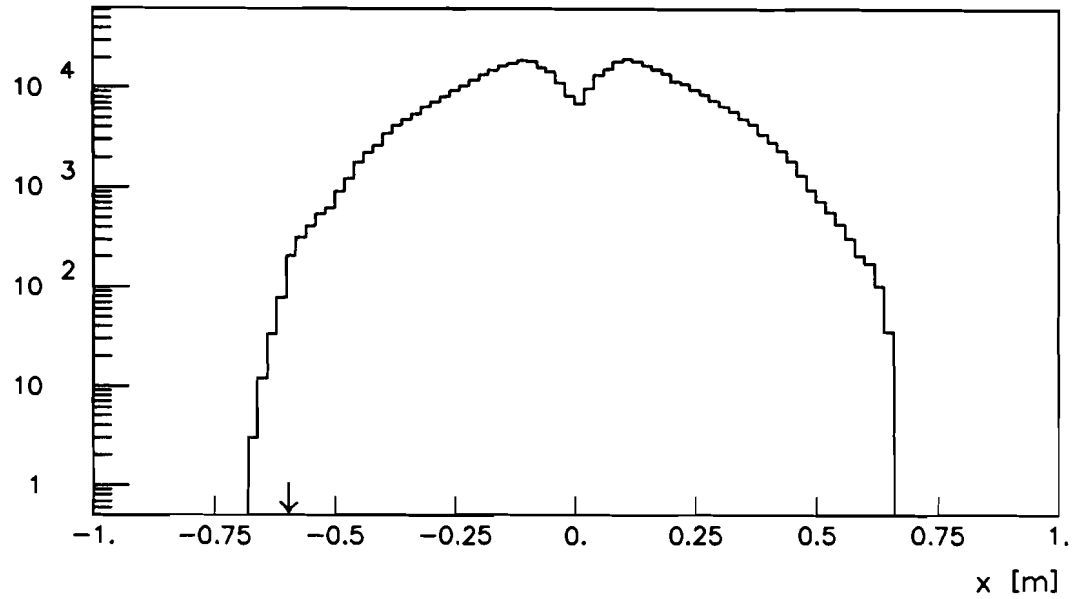

Figure 6.14: Track x projection at $z = 162\,\mathrm{m}$.

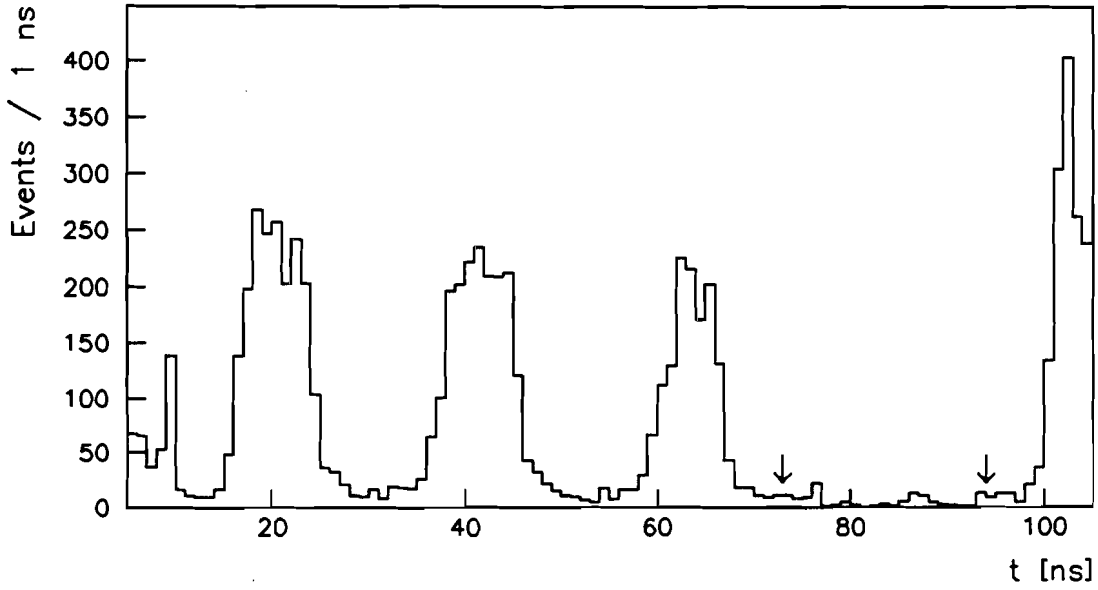

Figure 6.15: Time distribution of the RA trigger signal.

$p_T^2 < 250\,\mathrm{MeV}^2/\mathrm{c}^2$. The signal events are binned according to their energy (E) and longitudinal vertex position (z), with a bin size of 10 GeV$\times$1 m. Only events in the range $30 < E < 130\,\mathrm{GeV}$ and $123 < z < 137\,\mathrm{m}$ are used in the final analysis. Figure 6.16 shows the invariant mass for both beams after all other cuts have been applied (including the p_T^2, E and z cuts). The events in the low-side tail are radiative $\pi^+\pi^-\gamma$, and there is a very small remaining background from $\pi e\nu$ decays.

The transverse momentum squared is plotted in Figure 6.17 after all other cuts (including m, E and z). In the regenerated beam, the background is dominated by incoherently scattered kaons decaying into $\pi^+\pi^-$. For the p_T^2 range in the figure, the elastic scattering dominates over the inelastic. In the vacuum beam, the background comes from the remaining $\pi e\nu$ decays. The background level is determined from a fit, with exponential terms describing the different contributions. The amount of background does not seem to depend on the kaon energy or the decay vertex, and it is found to be 0.13% in the regenerated beam and 0.32% in the vacuum beam. Since the $\pi^+\pi^-\gamma$ decays form a tail on the signal, they do not get subtracted. The radiative part of this decay mode is the same for K_S and K_L, so it does not matter if it is subtracted or not. The non-radiative part ("direct emission") is small enough so that the difference between K_S and K_L is negligible.

Figure 6.18 shows the distribution of kaon energies after all other cuts. There is a small difference between the regenerated beam and the vacuum beam, since the regeneration amplitude depends on the kaon energy. Figure 6.19, finally, shows the transverse position of the kaon at the regenerator after all cuts. This is the variable that is used for deciding which beam an event comes from. As seen in the figure, there is no overlap between the two beams. (This plot is not corrected for the regenerator motion, and the two peaks correspond to the top and bottom beams.) After all cuts (but before the background subtraction), there are 198,264 signal events in the regenerated beam and 40,925 in the vacuum beam.

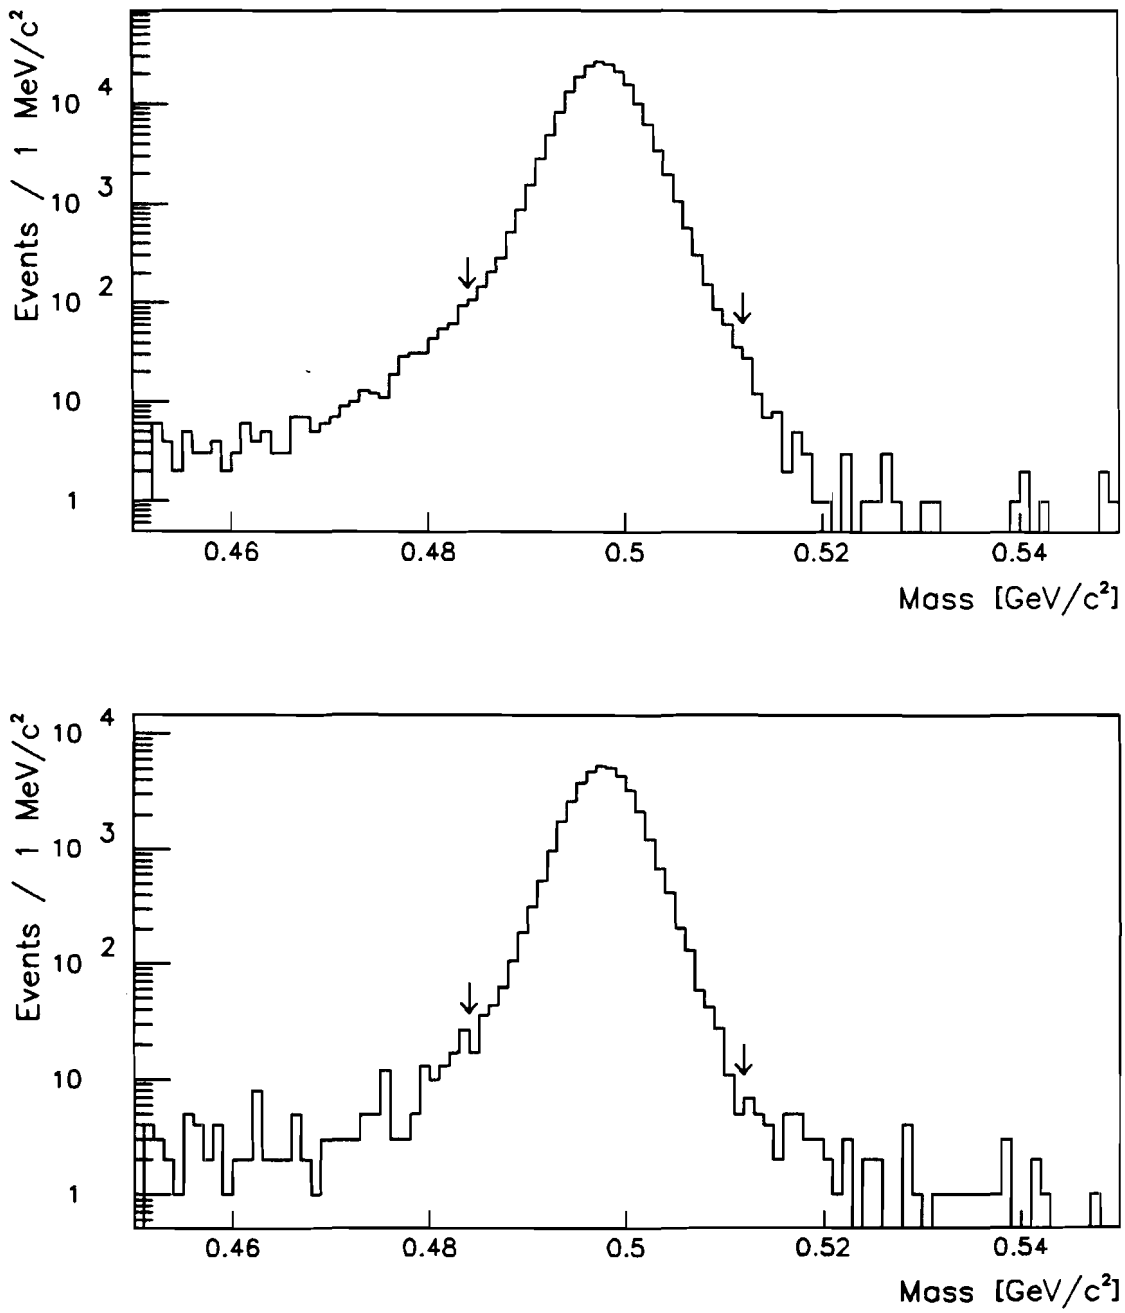

Figure 6.16: Invariant $\pi^+\pi^-$ mass after all cuts, with arrows indicating the signal region. Regenerated beam (*top*) and vacuum beam (*bottom*).

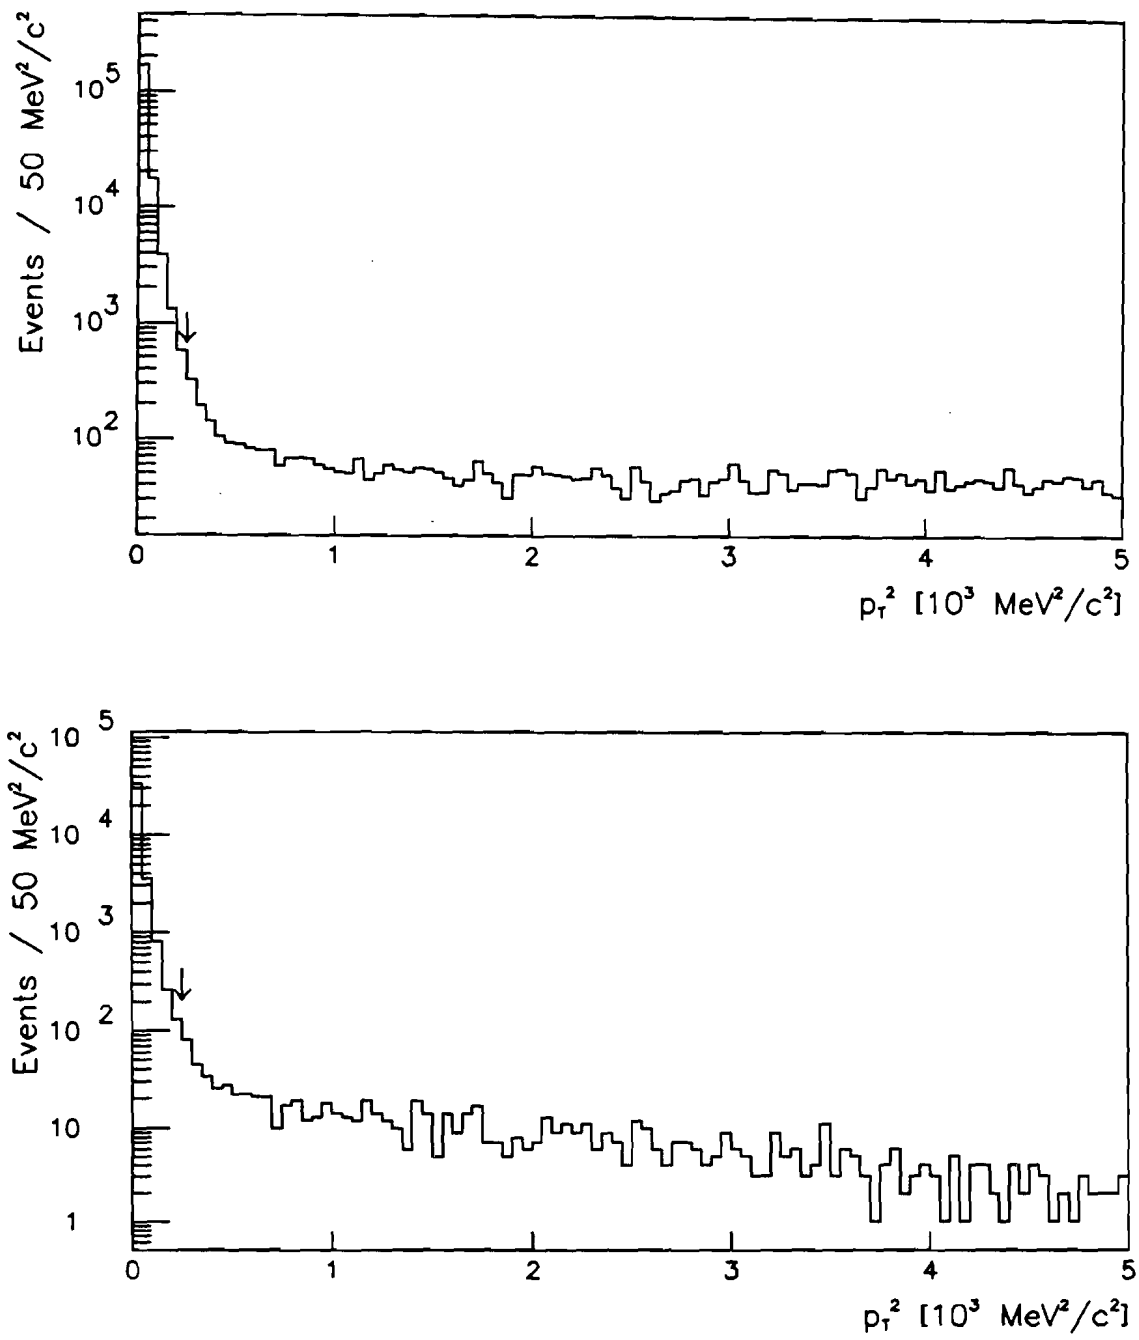

Figure 6.17: Transverse momentum squared after all cuts, with arrows indicating the signal region. Regenerated beam (*top*) and vacuum beam (*bottom*).

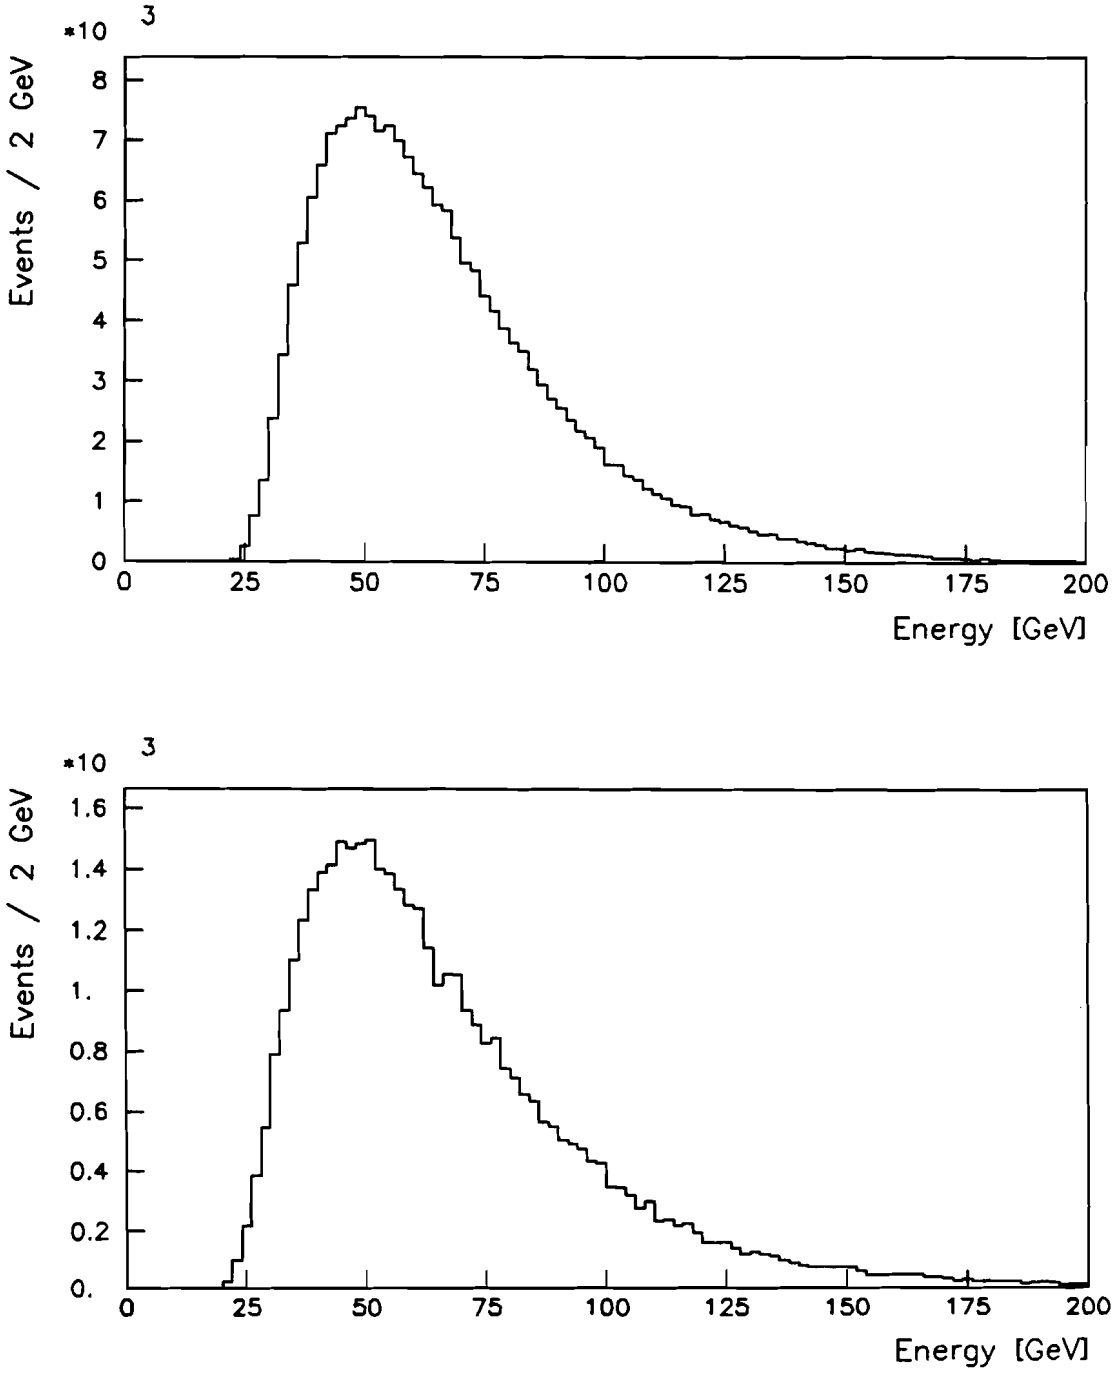

Figure 6.18: Kaon energy after all cuts. The range 30–130 GeV is used in the analysis. Regenerated beam (*top*) and vacuum beam (*bottom*).

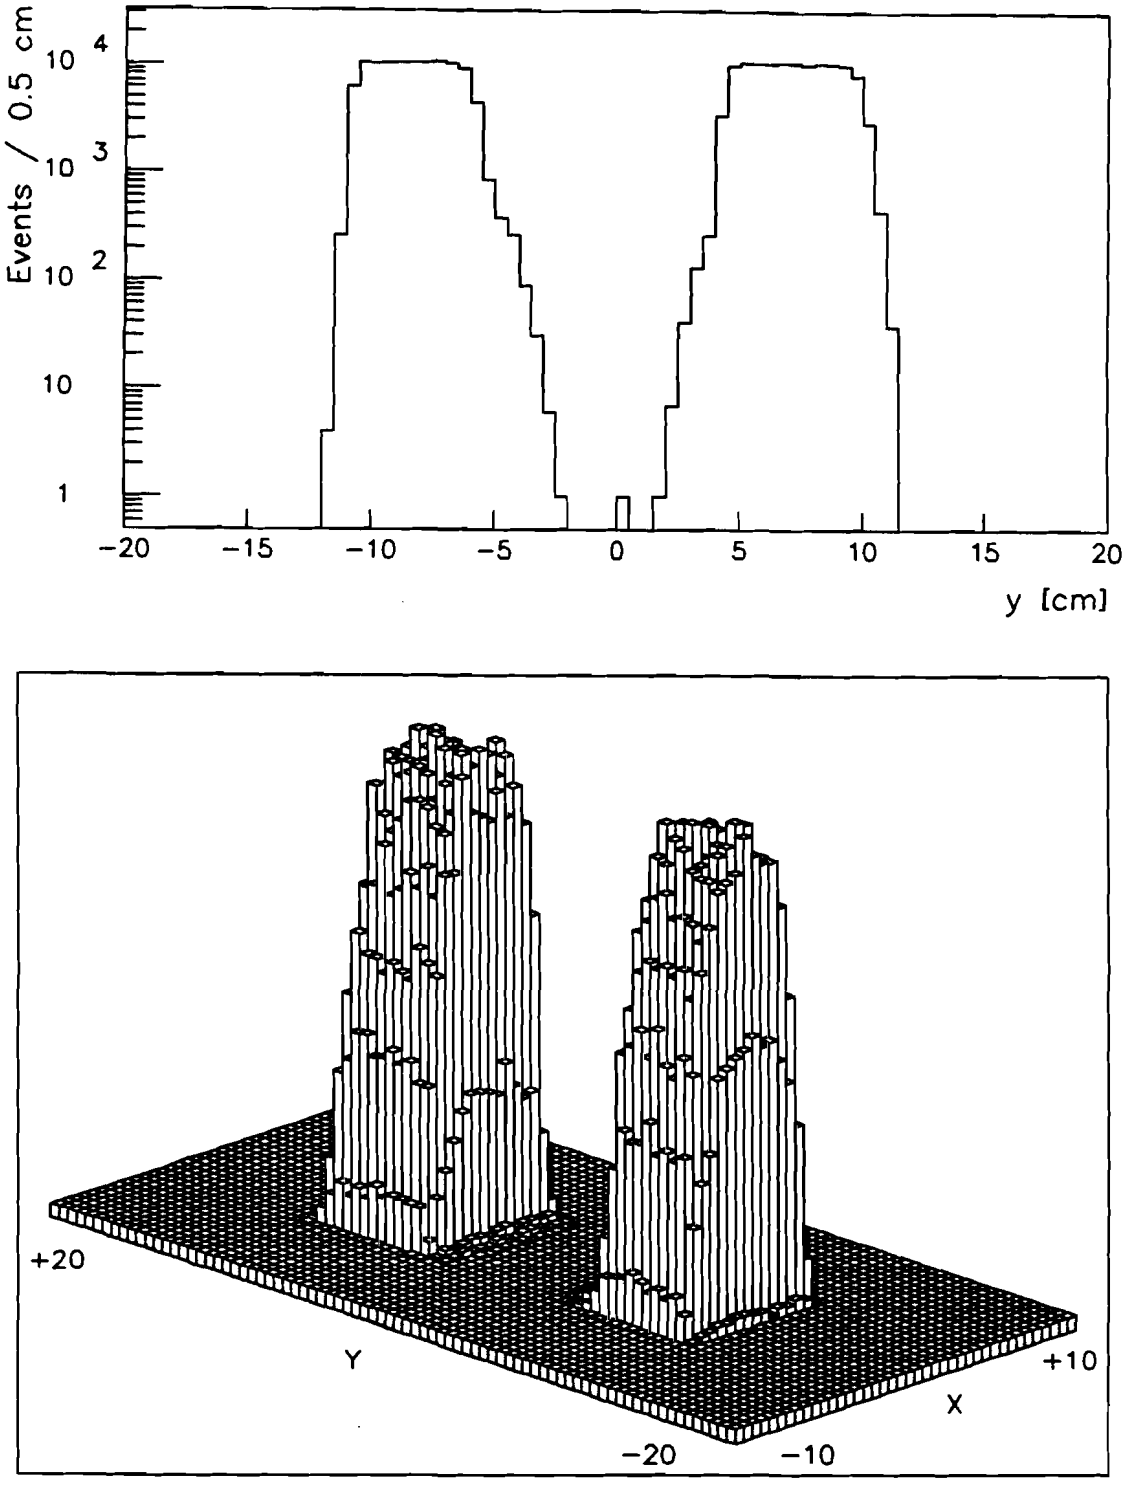

Figure 6.19: Transverse kaon position at the regenerator.

6.2 Neutral Mode Analysis

The neutral mode trigger is based on the total energy in the lead glass and the number of clusters found by the Hardware Cluster Finder (see chapter 4.3.2). The trigger accepts both four and six cluster events, which occur with roughly equal frequency. The six cluster triggers are $\pi^0\pi^0\pi^0$ decays, and most of the four cluster triggers are also $\pi^0\pi^0\pi^0$ decays, but with two photons either missing the lead glass or merging with other photons. It is the goal of the neutral mode analysis to find the $\pi^0\pi^0$ signal among the four cluster triggers.

The Hardware Cluster Finder recognizes clusters as "islands" of lead glass blocks with energies above a threshold of 1 GeV, which means that it is not able to separate overlapping clusters. The software cluster finding algorithm, which looks for energy maxima in the lead glass array, finds that about 20% of the four cluster triggers actually have more than four clusters (*i.e.*, more than four energy maxima). The procedure for finding the clusters and determining their energies and positions is described in more detail in chapter 5.4.1. If two photons are so close to each other that there is only one energy maximum — this is called fusion — not even the software cluster finder can separate them. As mentioned in chapter 5, a number of fusion criteria have been developed to tell whether a cluster is a single photon or not. If these criteria are applied to the events which have four clusters (according to the software algorithm), about half of them are found to have at least one fusion. The remaining events (about 10,000 per raw data tape) might be called four-photon events, even though a significant fraction are still fusions. The next section describes how the energies and positions of the four photons are used in order to reconstruct the decay.

(The fusion criteria are applied here early in the analysis in order to make the point that a lot of the four cluster triggers have more than four photons in the lead glass. If they are used instead at the end of the analysis for suppressing the background, only about 2% of the remaining events are cut. This is because most fusion events are also eliminated by other cuts.)

6.2.1 Reconstruction of Four-Photon Events

The four photons are considered to come from one single vertex since the π^0 decay length is so short (4 μm at 20 GeV). Let us first derive a relationship for the invariant mass of n photons originating at a single point. We will assume that the distance between this common vertex and the calorimeter is much larger than the size of the calorimeter, and in particular that all angles between the photon momenta and the beam direction (z axis) are small. (The decay volume ends more than twenty meters before the lead glass which has a radius less than one meter.) Denoting the photon energies and momenta by E_i and $\mathbf{p}_i$, with $|\mathbf{p}_i| = E_i$, we find that

$$m^2 = (\sum E_i)^2 - (\sum \mathbf{p}_i)^2 = \sum_{i<j} 2E_iE_j(1 - \cos\theta_{ij})$$

where θ_{ij} is the angle between $\mathbf{p}_i$ and $\mathbf{p}_j$. Using the small angle approximation, one has $2(1-\cos\theta_{ij}) = \theta_{ij}^2 = r_{ij}^2/z^2$ where r_{ij} is the separation between photons i and j at the calorimeter and z the distance between the vertex and the calorimeter. The final result for the invariant mass is thus

$$m^2 = \frac{1}{z^2}\sum_{i<j} E_iE_jr_{ij}^2$$

Since all the E_i and r_{ij} are measured by the calorimeter, this is a relationship between the invariant mass and the longitudinal component of the vertex position. Unfortunately, because of the small angles, there is no information about the transverse vertex position.

Applying the above result to all four photons, we have a relationship between z and the kaon mass. Knowing that the four photons are actually two pairs of photons coming from two real pions, and applying the m-z relation to the two photon pairs, we also have two relationships between z and the pion mass. The only slight complication is that there are three different ways of pairing the four photons. The standard approach in our analysis is to — for each pairing — assume the pion mass for both photon pairs and to calculate two z values with errors based on the uncertainties in the cluster energies and positions. A chi-square is calculated for the hypothesis that the two photon pairs come from the same vertex,

and the pairing with the lowest chi-square is the winner. The final vertex position is obtained from a weighted average of the two z values from the best pairing, and the invariant kaon mass is calculated using this vertex position. The resolution for the longitudinal vertex position is about one meter, which is five times worse than for the $\pi^+\pi^-$ analysis.

The invariant mass for all four-photon events on one raw data tape is shown in Figure 6.20. The background is dominated by $K_L \rightarrow \pi^0\pi^0\pi^0$ decays which are fifteen times more common in the vacuum beam because of the big difference in K_L flux. The center of energy at the calorimeter is used to distinguish between the regenerated and vacuum beams. For example, if the center of energy of the four photons is in the upper half of the lead glass and if the regenerator is in the lower position, the event is assumed to be from the vacuum beam. The total kaon energy is simply the sum of the four photon energies.

Note that the vertex position and the invariant kaon mass (via z) depend on which pairing is chosen, and only the kaon energy is pairing independent. An alternative approach is to assume the kaon mass for the four-photon system and obtain a pairing independent vertex position. Given that vertex, two pion masses can be calculated for each pairing, and the pairing with the best combination of pion masses is selected. With this approach, the average of the two pion masses corresponds to the kaon mass of the standard approach and the difference between the pion masses to the chi-square.

6.2.2 Kinematic Cuts

In this section, we describe the cuts that are based on variables derived from the photon energies and positions or, in other words, cuts using only the calorimeter information. Each plot is made after all other cuts and corresponds to the whole data set. Figure 6.21 shows the distribution of photon energies. All four photons are required to have an energy above 1.5 GeV. The position of the photons in the lead glass is plotted in Figure 6.22 relative the outer edge of the lead glass. The distance from the edge is defined in such a way that, for example, all points on the dashed line in Figure 6.6 are half a block from the edge. The bumps occur because

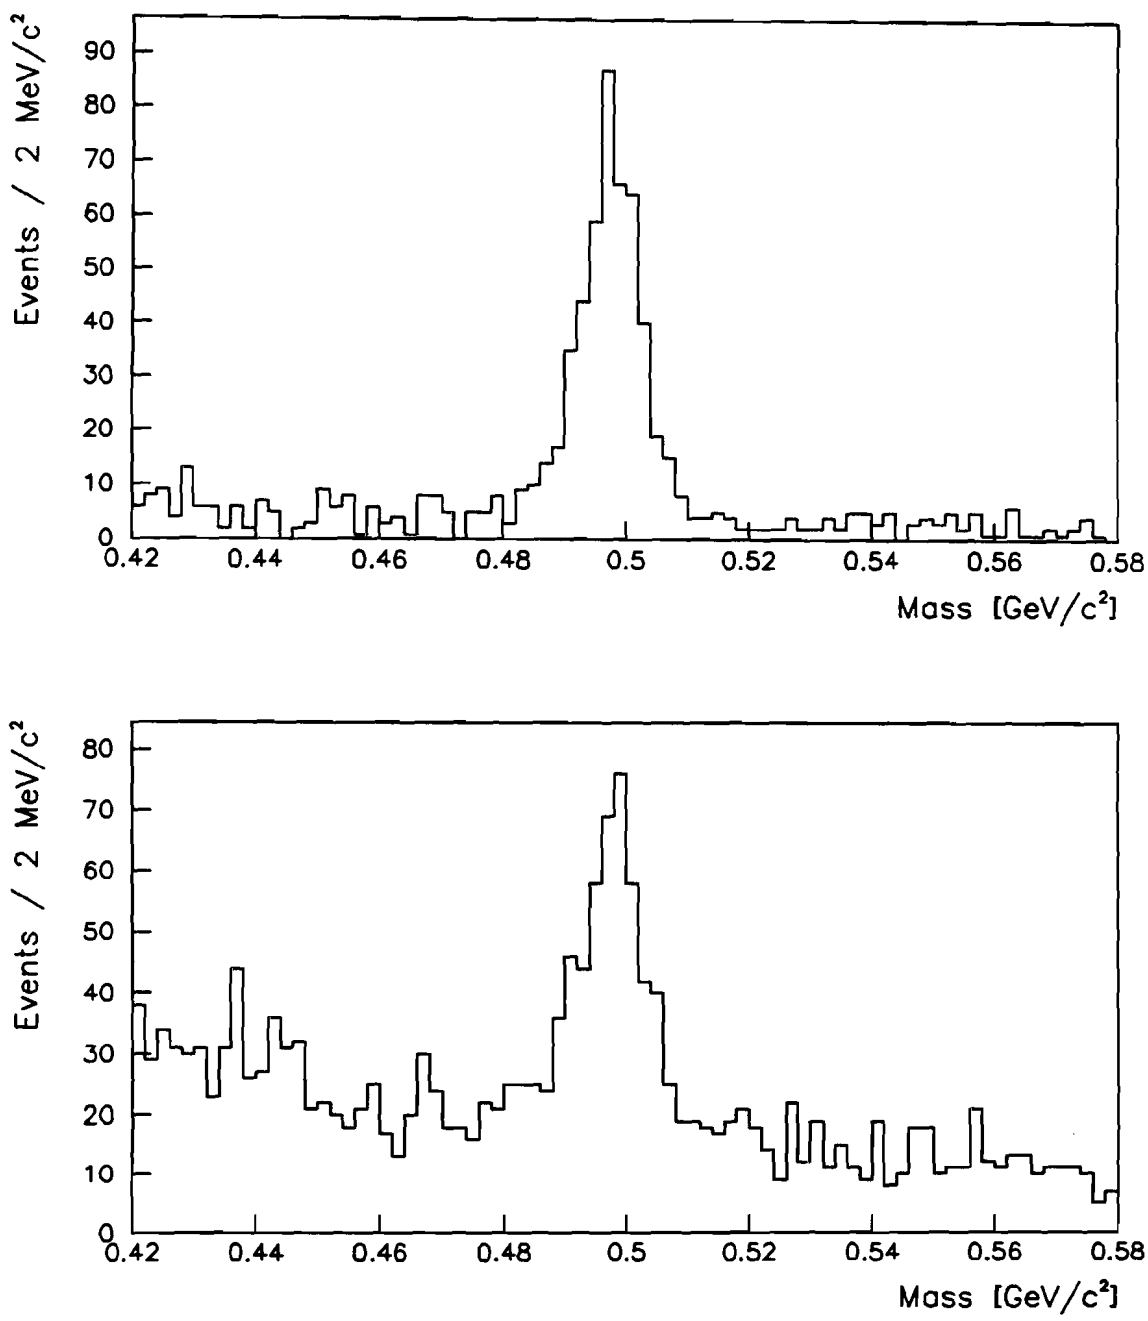

Figure 6.20: Invariant mass for all four-photon events on one raw data tape. Regenerated beam (*top*) and vacuum beam (*bottom*).

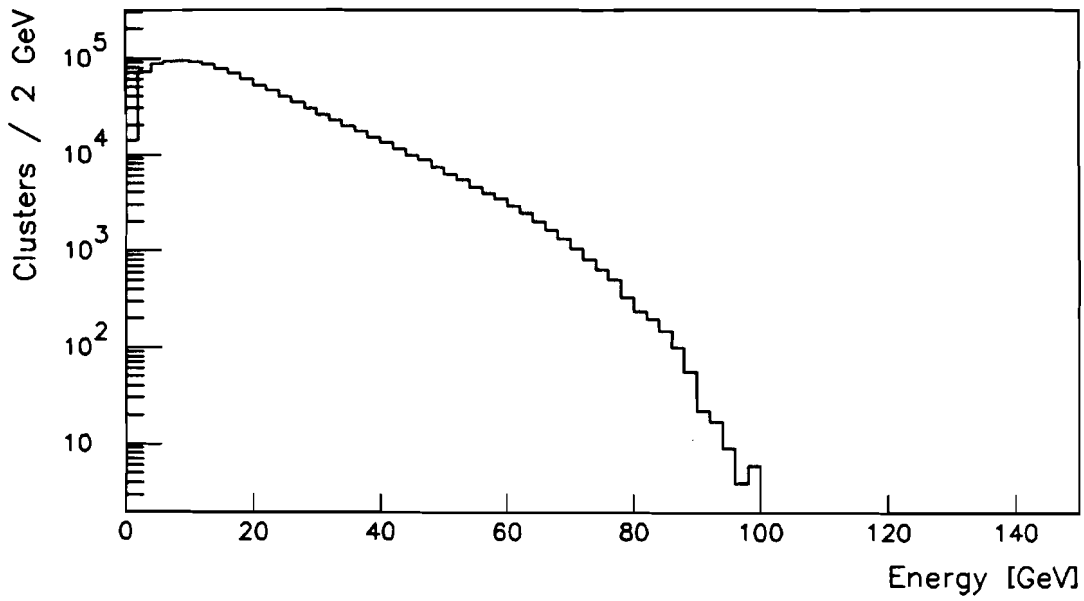

Figure 6.21: Cluster energy distribution.

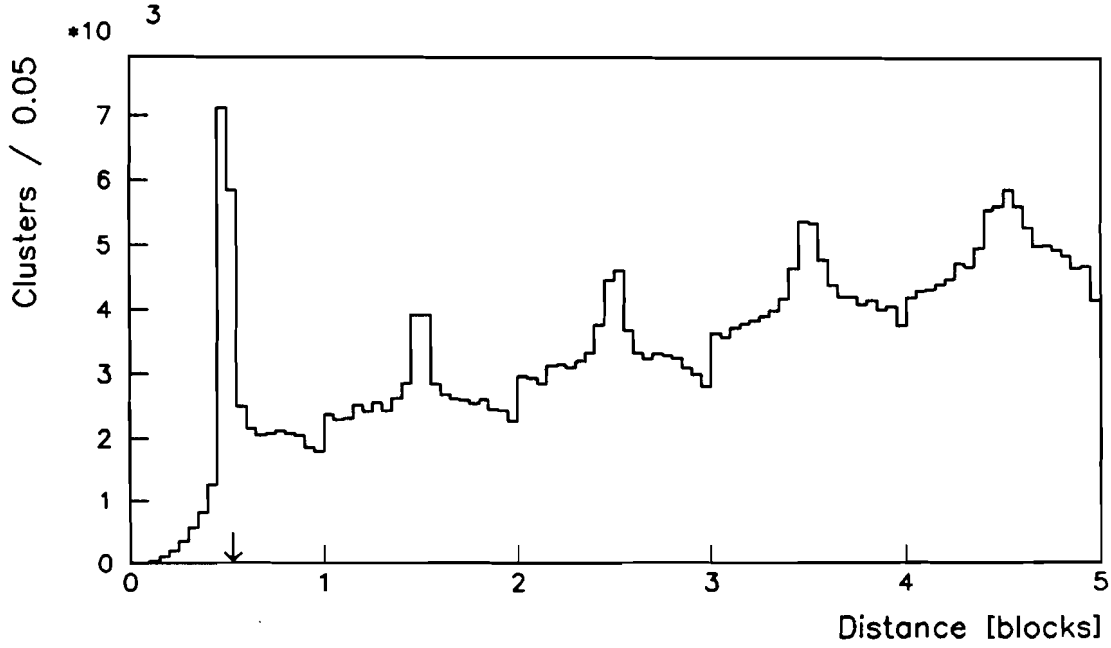

Figure 6.22: Cluster position measured from the lead glass outer edge.

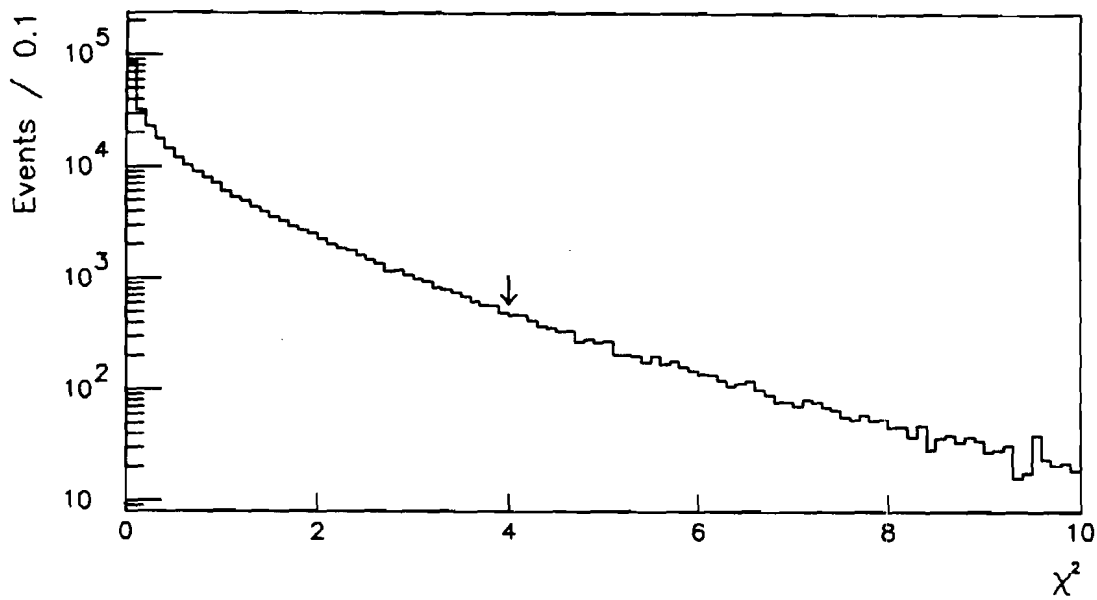

Figure 6.23: Chi-square for the common vertex hypothesis.

the cluster position algorithm (see 5.4.1) is not perfectly tuned. It is particularly difficult to assign positions to clusters near the edge, since there are no blocks on the other side of the edge. All four photons are required to be at least 0.53 blocks from the edge. (A block is 5.8 cm wide.)

As mentioned above, a chi-square is calculated for the hypothesis that both photon pairs originate at the same vertex. The chi-square, which is plotted in Figure 6.23, is required to be less than 4. Since there are three possible photon pairings, the question arises whether the one with the lowest chi-square is really the correct one. If the difference in chi-square between the two best pairings is less than 8, we also calculate the invariant kaon mass for the second best pairing. This mass is plotted in Figure 6.24, and there is indeed a peak at the kaon mass. If the mass is between 470 and 526 MeV/c^2, the event is cut.

6.2.3 Other Cuts

We now turn to the cuts that rely on other detector elements. The photon veto counters are particularly important for the background suppression, since they often catch missing photons from $\pi^0\pi^0\pi^0$ decays. Some of these counters are in veto

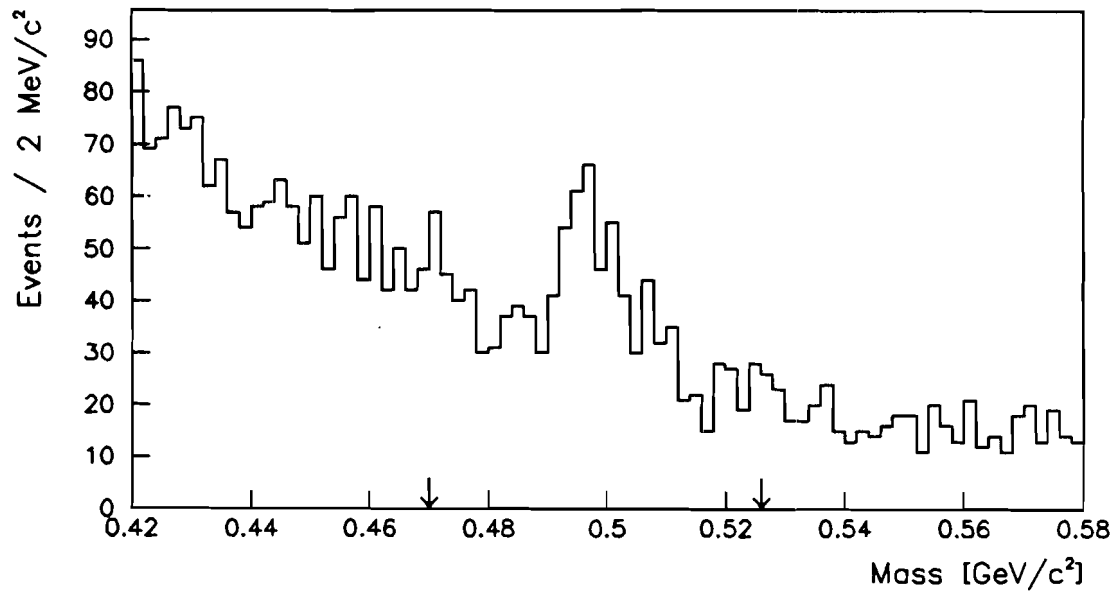

Figure 6.24: Invariant kaon mass for the second best pairing when the chi-square difference between the two best pairings is less than 8.

already in the trigger (chapter 4.3.2). Figures 6.25 and 6.26 show the photon veto cuts. The number of counts is expressed in "muon equivalents," and a photon energy of one GeV typically corresponds to 5–10 muons. The energy distribution in the Collar Anti (CA) counters is shown in Figure 6.27, where the cut is indicated by an arrow. The number of counts in the CA counters is also expressed in muon equivalents, with a photon energy of one GeV roughly equal to four muons. These counters form a collar around the beam holes in front of the lead glass. Because of the holes, the energy and position resolutions for photons become worse near the inside edge of the lead glass, but the CA defines a convenient edge half a block away from the holes. There is also a cut with the SA (Sweeping Anti) counters. The signals from these counters is discriminated and latched, and the cut requires that none of the counters is on.

There is a cut to eliminate events with charged particles in the detector. Figure 6.28 shows the number of drift chamber hits for events with at least one counter hit in both the B and C hodoscopes. (The signals from the B and C counters are latched.) The event is cut if there are fifteen or more chamber hits. Figure 6.29, finally, shows the timing of the RA trigger signal. If the time is between

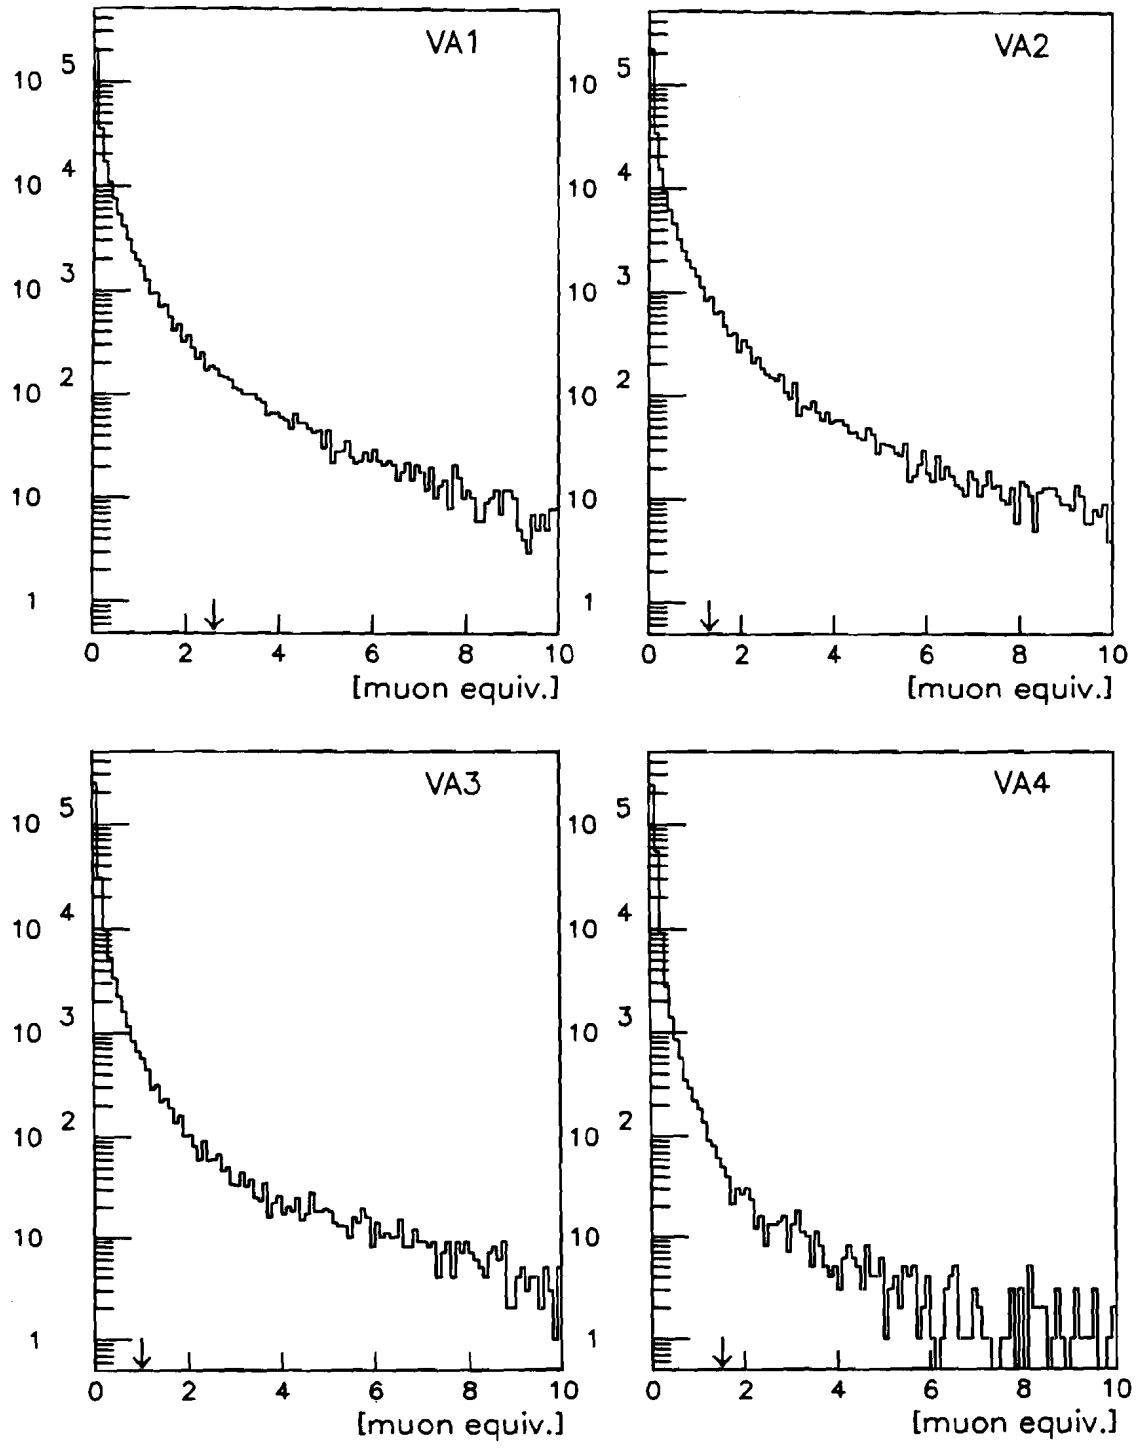

Figure 6.25: Photon veto cuts.

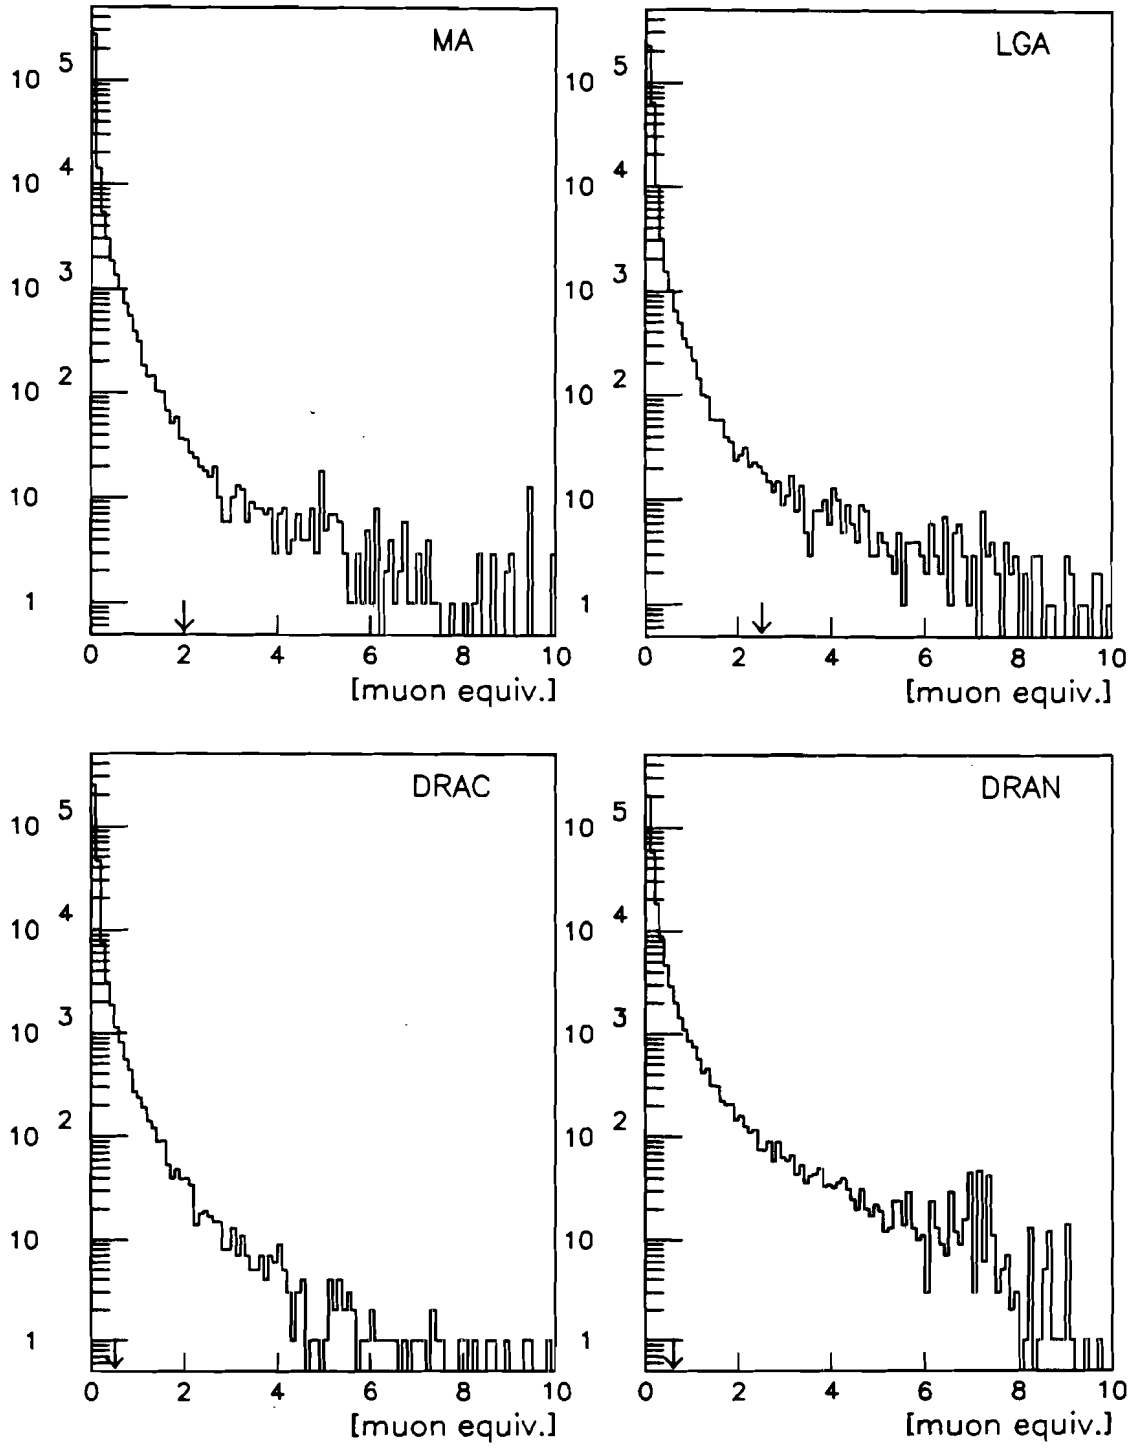

Figure 6.26: Photon veto cuts.

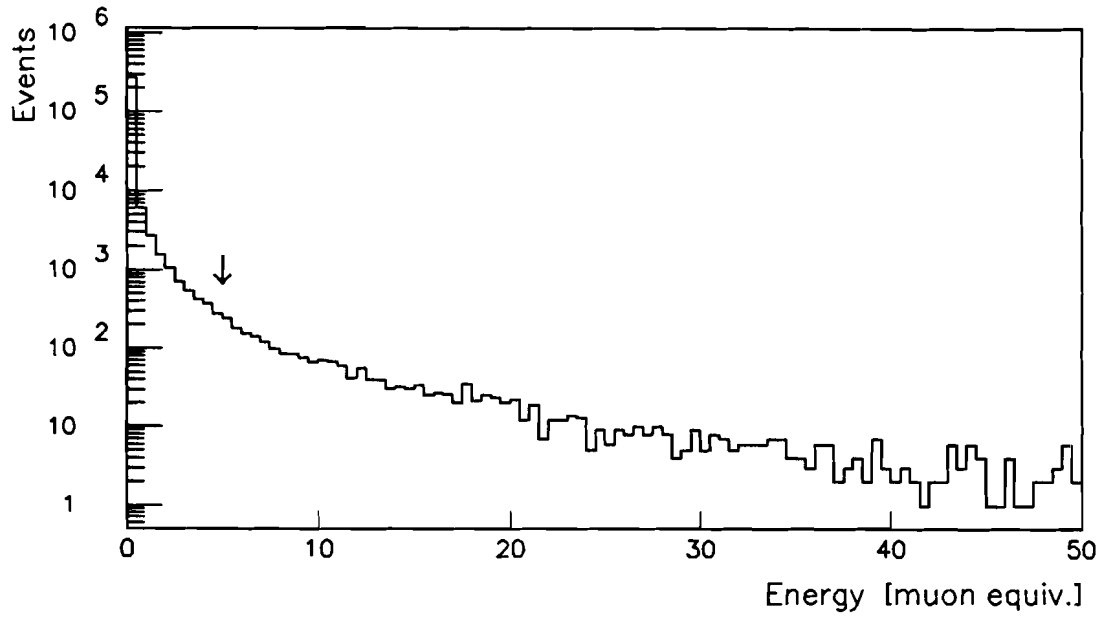

Figure 6.27: Energy in the Collar Anti counters.

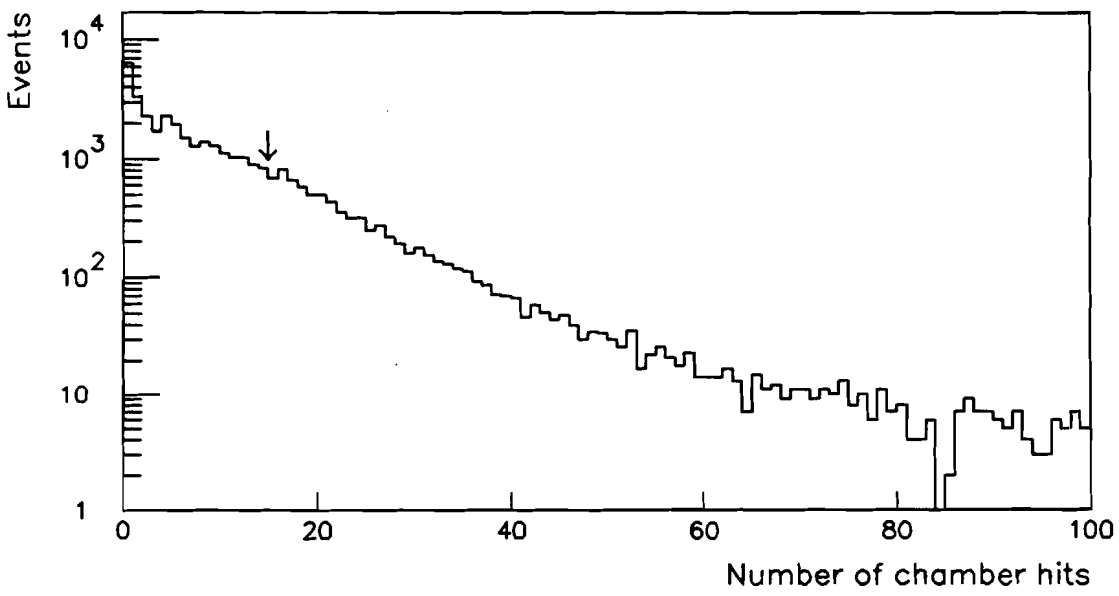

Figure 6.28: Number of drift chamber hits for events with at least one counter hit in both the B and C hodoscopes.

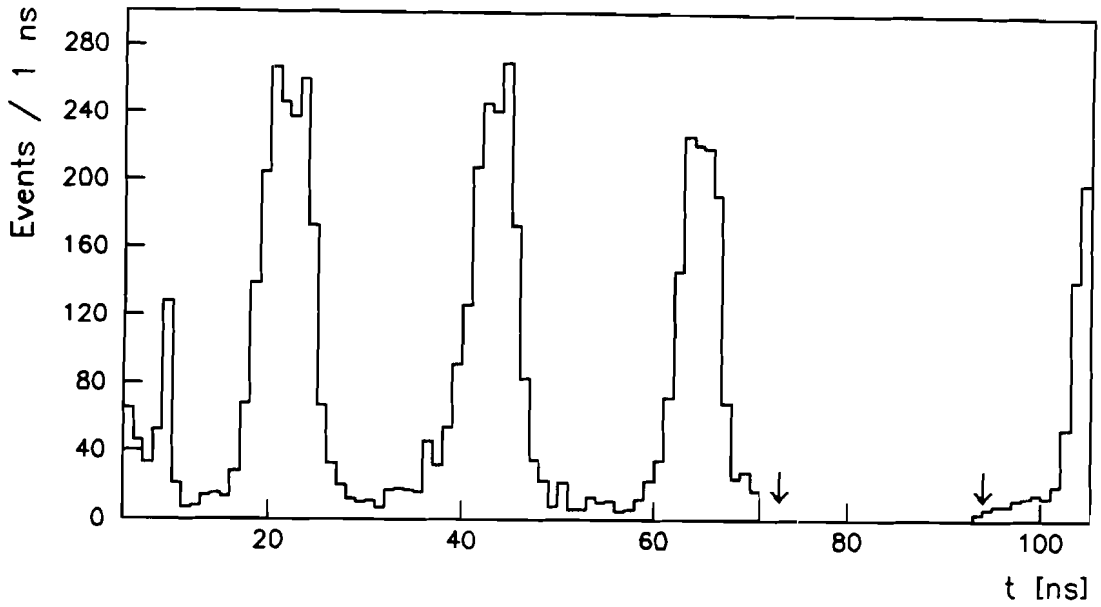

Figure 6.29: Time distribution of the RA trigger signal.

the arrows, the event is cut. This is the same cut as is used in the $\pi^+\pi^-$ analysis (see Figure 6.15).

6.2.4 Final Data Sample and Background Subtraction

As in the charged mode analysis, there is an invariant kaon mass that can be used for background subtraction, but unlike the charged mode, there is no p_T^2 variable available. Instead we use the center-of-energy of the four photons at the lead glass to indicate if the transverse momentum is non-zero. Figure 6.30 shows the center of energy after all other cuts have been applied. This is also the variable that is used for deciding which beam an event comes from. The figure is corrected for the regenerator motion, so that one peak (the taller one) corresponds to the regenerated beam and the other one to the vacuum beam. The two peaks are well separated, but the incoherently scattered kaons, which originate in the regenerated beam, form a background under both peaks. (This does not happen in the charged mode, where the transverse vertex position is available.) By drawing concentric squares around the two beam centers, a "ring number" is determined. This quantity is simply the area of the square on which an event falls in the center of energy plot, and it is

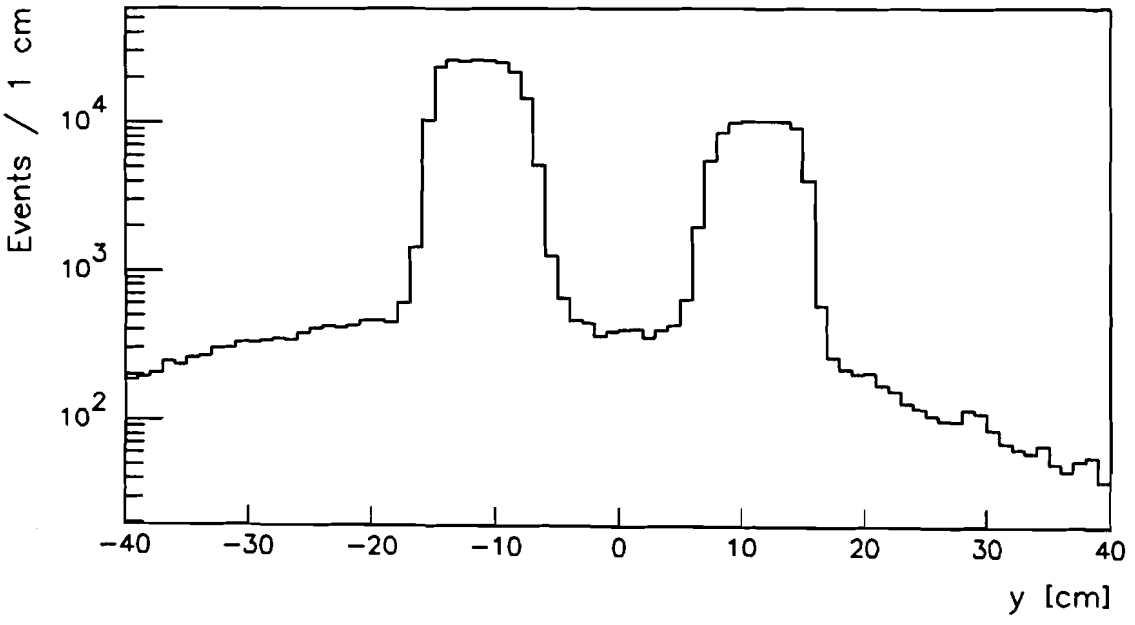

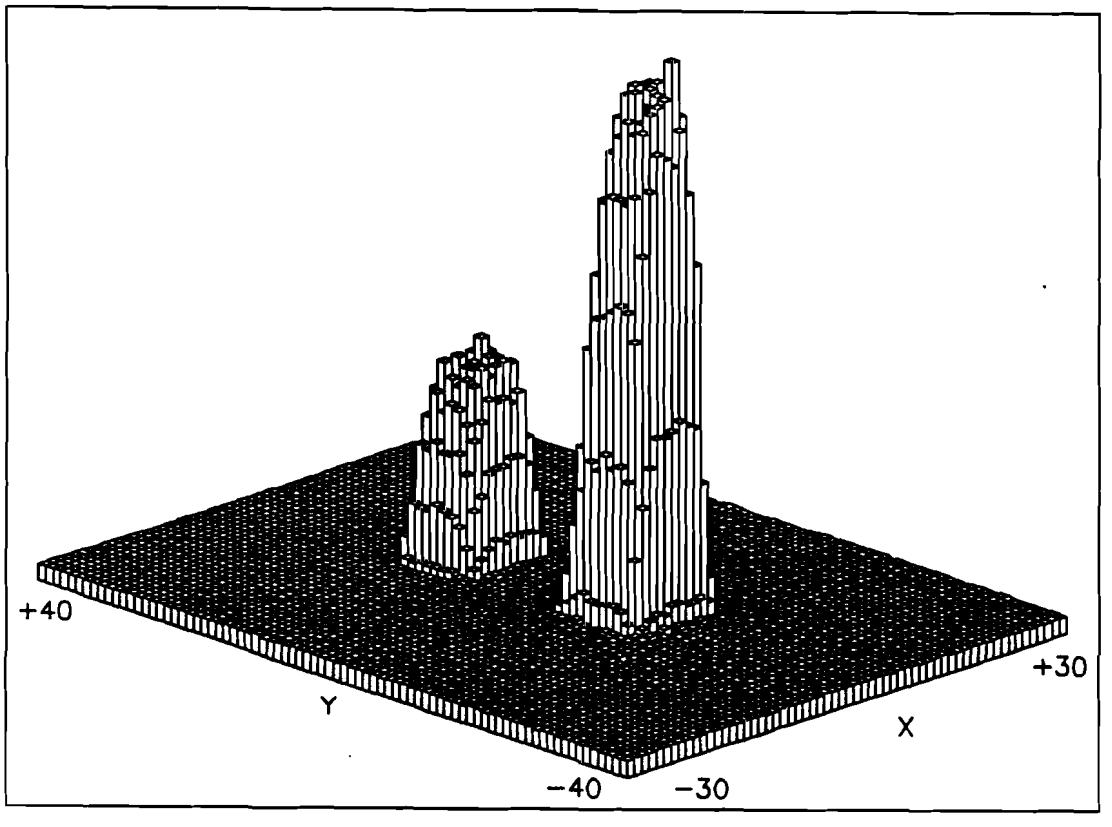

Figure 6.30: Center of energy at the lead glass.

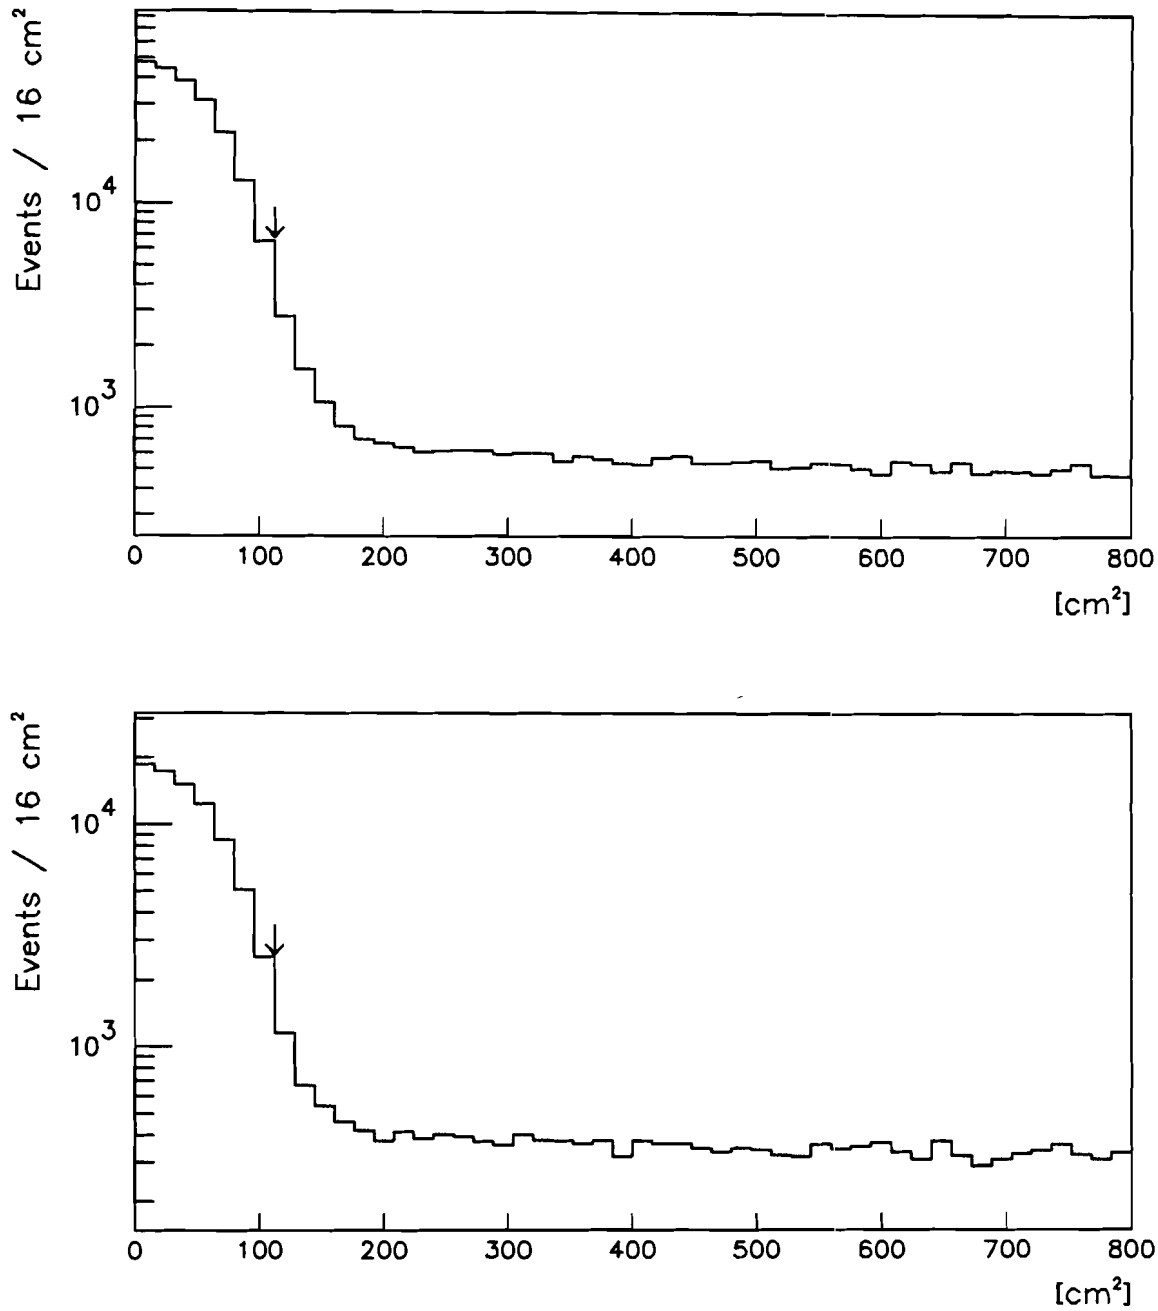

Figure 6.31: Ring number after all cuts, with an arrow indicating the signal region. Regenerated beam (*top*) and vacuum beam (*bottom*).

shown in Figure 6.31 for both beams after all other cuts. The ring number imitates the p_T^2 variable, but since it is not possible to remove the finite width of the beam, there is no sharp peak at zero.

It is this ring number that we use together with the invariant mass (m), in order to identify the final $K \rightarrow \pi^0\pi^0$ sample and subtract the background. The signal region is defined by $480 < m < 516\,\mathrm{MeV/c^2}$ and a ring number less than 112 cm^2. The signal events are binned according to their energy (E) and longitudinal vertex position (z), with a bin size of 10 GeV×1 m. Only events in the range $40 < E < 120\,\mathrm{GeV}$ and $121 < z < 147\,\mathrm{m}$ are used in the final analysis. Figure 6.32 shows the invariant mass for both beams after all other cuts have been applied. The background in the vacuum beam (0.99%) is dominated by $\pi^0\pi^0\pi^0$ decays. For the regenerated beam, the background is so small (0.04%) that there is also a significant contribution from nuclear interactions (producing two neutral pions) in the regenerator and in the trigger hodoscopes at the end of the upstream decay volume.

The background subtraction in the ring number variable is helped by a Monte Carlo simulation of the incoherent scattering. The shape of the observed p_T^2 distribution for $\pi^+\pi^-$ decays is plugged into the Monte Carlo, and a prediction of the background shape in the ring number variable is obtained. Both the size and the shape of this prediction agree very well with the observed background. The predicted shape is then used to extrapolate the observed background under the coherent peak, and the level of the background is found to be 2.7% in the regenerated beam and 2.8% in the vacuum beam. The mass and ring number subtractions are done simultaneously to make sure that the backgrounds are treated correctly. The $\pi^0\pi^0\pi^0$ decays, for example, form a background in both variables and would be over-subtracted if the mass and ring number subtractions were done separately.

The background levels quoted above are averages over all bins, and the actual subtractions are done separately in each energy and vertex bin. The incoherent background increases with the kaon energy, since at higher energies, a larger p_T^2 range stays inside the beam region at the lead glass. The $\pi^0\pi^0\pi^0$ events (when

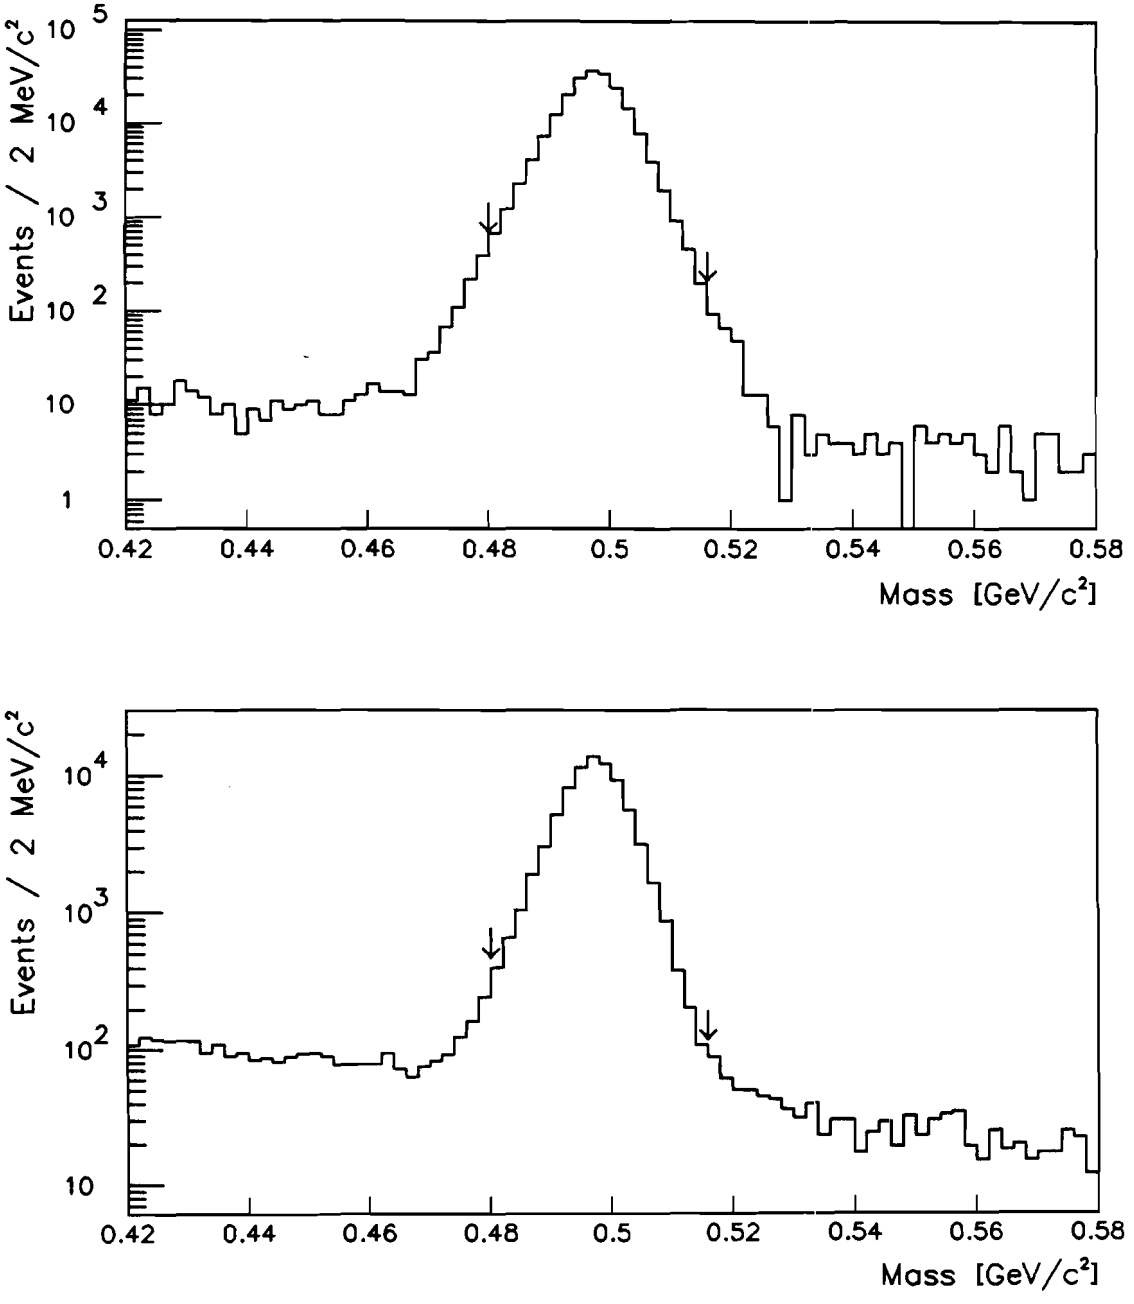

Figure 6.32: Invariant $\pi^0\pi^0$ mass after all cuts, with arrows indicating the signal region. Regenerated beam (*top*) and vacuum beam (*bottom*).

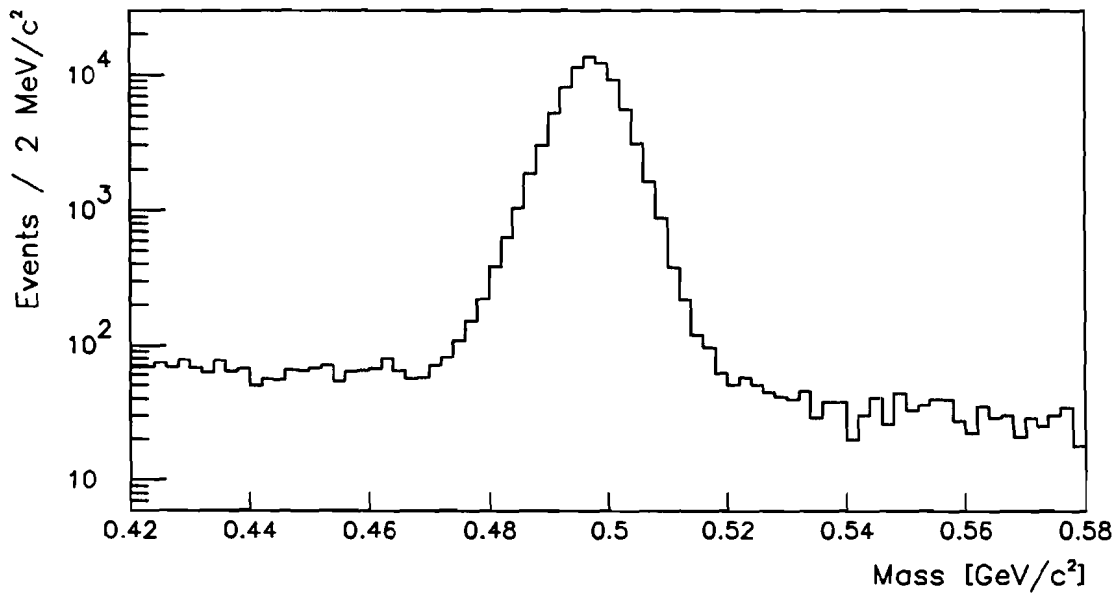

Figure 6.33: Invariant $\pi^0\pi^0$ mass in the vacuum beam with an alternative z cut.

reconstructed as $\pi^0\pi^0$ decays) are pushed towards the lead glass, so that this background increases with the decay vertex. One might wonder what would happen if the decay vertex was calculated by assuming the kaon mass instead of the pion mass, and if this alternative z value was used to bin the signal events. With the same 121–147 m cut, the invariant mass (calculated using the traditional z value) would then appear as in Figure 6.33, where it can be seen that the background is not quite as steep as before. The size of the background is the same, however.

Figure 6.34 shows the distribution of kaon energies after all other cuts. As in the charged mode, there is a small difference between the two beams, since the regeneration amplitude depends on the kaon energy. After all cuts (but before the background subtractions), there are 195,428 signal events in the regenerated beam and 76,177 in the vacuum beam.

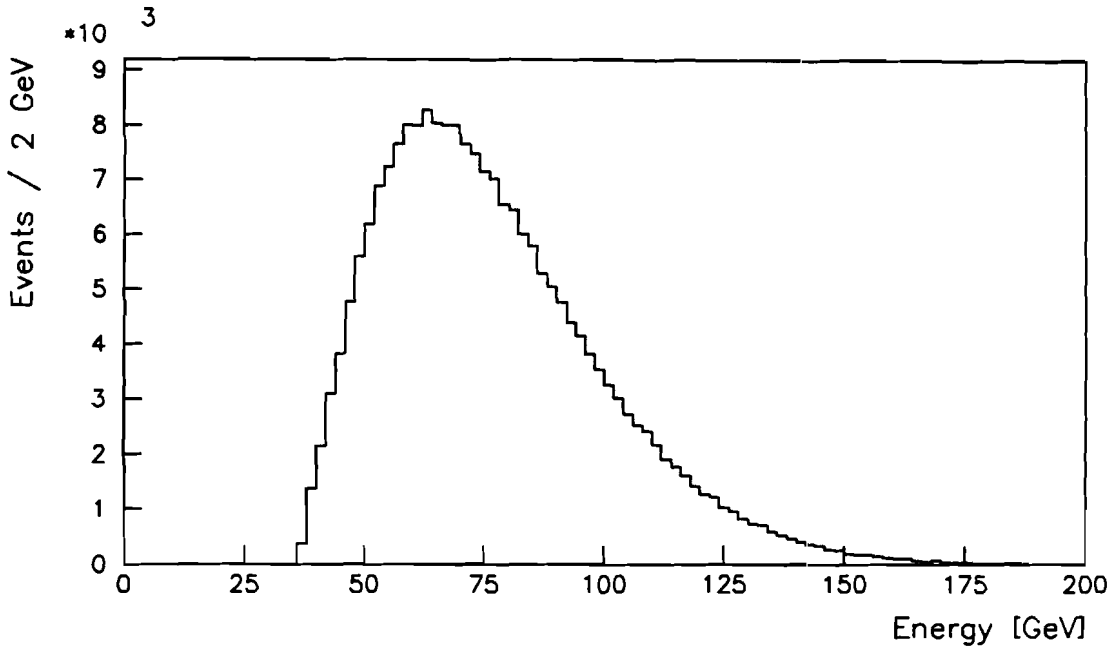

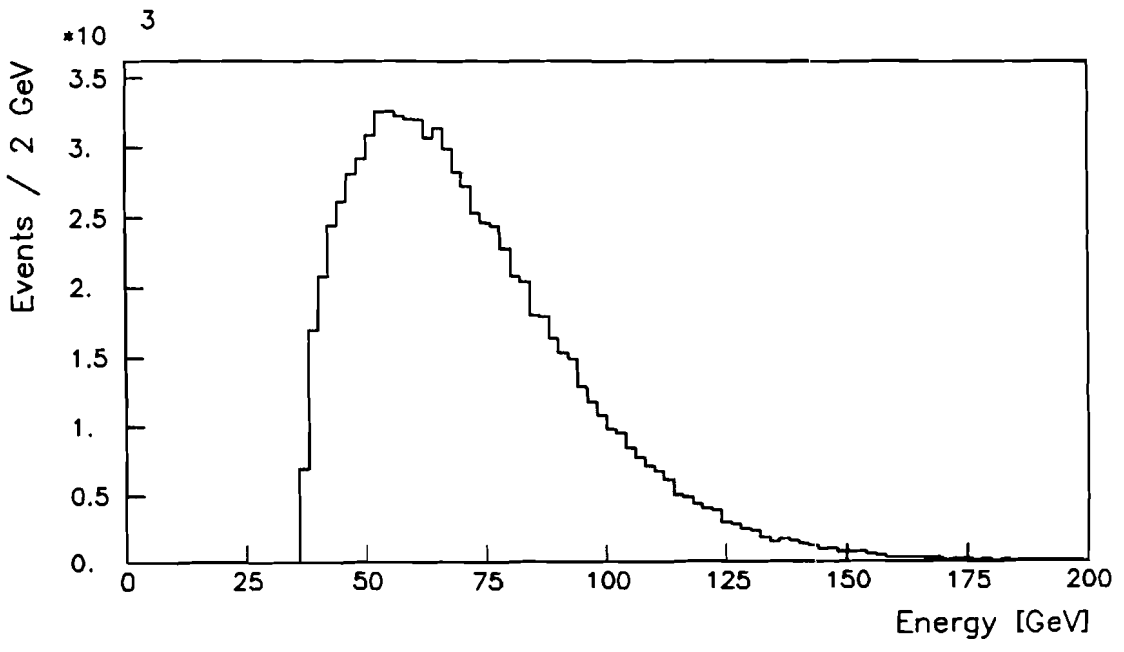

Figure 6.34: Kaon energy after all cuts. The range 40–120 GeV is used in the analysis. Regenerated beam (*top*) and vacuum beam (*bottom*).

Chapter 7
DETECTOR SIMULATION

We want to use the decay vertex distributions in the regenerated beam to determine the phases of the interference terms, but before extracting those phases, it is necessary to acceptance-correct the data. This is best done by doing a detailed Monte Carlo simulation of the experiment. The simulation consists of generating a large number of kaon decays and finding out what fraction of these decays reconstruct successfully and survive all the analysis cuts.

7.1 Event Generation

The kaons are generated at the target as either a K^0 or a $\overline{K^0}$ state. For the kaon energy and production angle distributions, we use functions that have been obtained in other kaon experiments. Those experiments looked at charged kaons, while our experiment uses neutral kaons, and this is probably the reason why the predicted kaon energy spectrum does not completely agree with what we observe. To correct for this difference, the Monte Carlo energy spectrum is multiplied by a polynomial which is determined by comparing the data and Monte Carlo energy distributions for $\pi^+\pi^-$ decays. This adjusted energy spectrum is then used for generating both $\pi^+\pi^-$ and $\pi^0\pi^0$ decays.

The kaon that is generated at the target is propagated through the beamline, in order to obtain the correct beam shape. All collimators are simulated, and scattering in the various absorbers is also taken into account. The K_S and K_L amplitudes of each kaon at the start of the decay volume are calculated by including all effects from absorption, regeneration and the probability of the kaon decaying before reaching the decay volume. The magnitudes and relative phase of the K_S and K_L amplitudes are then used to obtain the correct decay vertex distribution,

including the K_S-K_L interference. (In the Monte Carlo, nominal values are used for the parameters that we wish to determine. The phases of η_{+-} and η_{00}, for example, are set to 45 degrees.) The decay vertex is picked randomly for each kaon, in such a way that a large sample of Monte Carlo events will have the correct vertex distribution.

The directions and energies of the decay products are generated in the kaon rest frame. For two-body decays, it is sufficient to pick a random direction for the back-to-back momentum vectors, but other decay modes need to be generated according to some specific phase space factor. The decay products are boosted to the lab frame and are then traced individually from the decay vertex until they decay into other particles or escape the detector or get absorbed in some detector element. The neutral pions, for example, decay immediately into two photons, which are then traced until they escape or get absorbed. The charged pions, on the other hand, do not decay very often before the end of the detector, but the simulation does include this possibility. When the particles pass through the various thin layers of material (trigger hodoscopes, vacuum window, chamber windows), the multiple scattering of charged particles is simulated, and the small probability of photon conversion is taken into account. Electrons that scatter also produce bremsstrahlung photons. The bending of charged particles at the analysis magnet is determined by looking up the transverse momentum kick in the magnet calibration tables.

As the particles are traced through the experiment, the detailed response of the various detector elements is also simulated. For the various trigger hodoscopes, the projected positions of charged particles are used to determine exactly which counters in the hodoscopes are hit. (The positions and efficiencies of the counters are known from muon events.) The inefficiency of the B and C hodoscopes is dominated by the small gaps between the scintillator counters. These gaps can be observed with muon events, and their locations and widths are included in the Monte Carlo. The gaps are typically a millimeter wide, leading to an inefficiency of about one percent.

Charged particles produce wire hits in the drift chambers, with efficiencies that

are typically around 99% and determined from real data. A handful of wires are known to be dead or to have low efficiencies, and this information is also included in the Monte Carlo. After resolution smearing, the transverse positions of the particles are converted into drift times by using the time to distance calibration backwards. The wire hits and the drift times are written to the event record, so that Monte Carlo events can be reconstructed the same way as data events. In particular, the Monte Carlo events are run through the same track finding algorithm as the data. For the trigger simulation, the hits in the second chamber are also checked to see if the east-west topology trigger is satisfied. As the charged particles go through the chambers, the contributions to multiple scattering from the chamber windows and from the field and sense wires are taken into account.

For photons and electrons, energy clusters need to be generated in the lead glass calorimeter. Recall that very detailed shower simulations, studying the non-linear energy response of the glass and the transverse distribution of energy among the blocks in a cluster, were done for the lead glass calibration. This type of simulation requires a lot of computer time and is too slow for the general Monte Carlo. However, the summarized information from the detailed shower studies can be used by the Monte Carlo to simulate the total energy response of the glass, *i.e.*, for a given true shower energy, simulate the observed total energy of the cluster. The transverse energy distribution is then simulated by simply picking a representative cluster (a three-by-three array of block energies) from a large cluster library, which has been compiled using clusters from real $K_L \rightarrow \pi^0\pi^0\pi^0$ events. The block energies in this array are then scaled so that the desired total cluster energy is obtained.

No hadron showers are simulated by the Monte Carlo. In particular, charged pions are always represented by a minimum ionizing signal in a single lead glass block, while in the real world charged pions quite often produce large showers.

Based on the simulated detector response, the various trigger requirements can be evaluated. For example, the calculations made by the cluster finding trigger processor are simulated in the Monte Carlo. If the event passes the trigger, the simulated event information is written to tape using the same format as real data events. After the data and the Monte Carlo have been analyzed using the same cuts,

the two event samples are compared to see if the simulation is accurate enough.

7.2 Comparison with Data

Figures 7.1–18 show how the data and Monte Carlo decay vertex distributions agree in the vacuum beam. In that beam, the vertex distributions do not depend on the parameters that we wish to determine in this experiment (in particular, the phases ϕ_{+-} and ϕ_{00}). The normalization of the two event samples is done only once for each decay mode and not for each energy bin. The chi-square of this data to Monte Carlo comparison is 59.5 for 69 degrees of freedom in the charged mode and 101.3 for 103 degrees of freedom in the neutral mode. (The kaon energy spectrum used in the Monte Carlo was determined from charged mode data, which probably explains the somewhat low chi-square obtained for that mode.)

Since the K_L decay distribution is essentially flat within our decay volume, the plots also show how the acceptance varies with the longitudinal vertex position. For low kaon energies, the particles from upstream decays are more likely to escape outside the detector than those from downstream decays. This results in an acceptance function that increases with vertex position. For higher kaon energies, we get the opposite effect, because the particles from downstream decays are now more likely to go straight down the beam holes in the calorimeter. The average acceptance (after all cuts) is 46% for $\pi^+\pi^-$ decays and 19% for $\pi^0\pi^0$ decays.

Since the data and Monte Carlo agree well in the vacuum beam, we know that the detector acceptance is simulated correctly (within the statistical precision of the vacuum beam data sample). It is then possible to extract the interference information from the vertex distributions in the regenerated beam, after using the Monte Carlo to correct for the acceptance. However, because of the limited event statistics in the vacuum beam, we cannot say that the acceptance correction is perfect, and it is necessary to assign a systematic uncertainty due to this correction. Fortunately, there are large numbers of $K_L \rightarrow \pi e \nu$ and $K_L \rightarrow \pi^0\pi^0\pi^0$ decays collected in the same experiment. By comparing the data and the Monte Carlo for these decay modes, a smaller acceptance uncertainty is obtained. (This assumes

that an acceptance error for the $\pi\pi$ decay modes would also show up in these other decay modes. Since $\pi e\nu$ decays have two tracks like the $\pi^+\pi^-$ and since the $\pi^0\pi^0\pi^0$ decays have even more energy clusters in the lead glass than the $\pi^0\pi^0$, this is not an unreasonable assumption.)

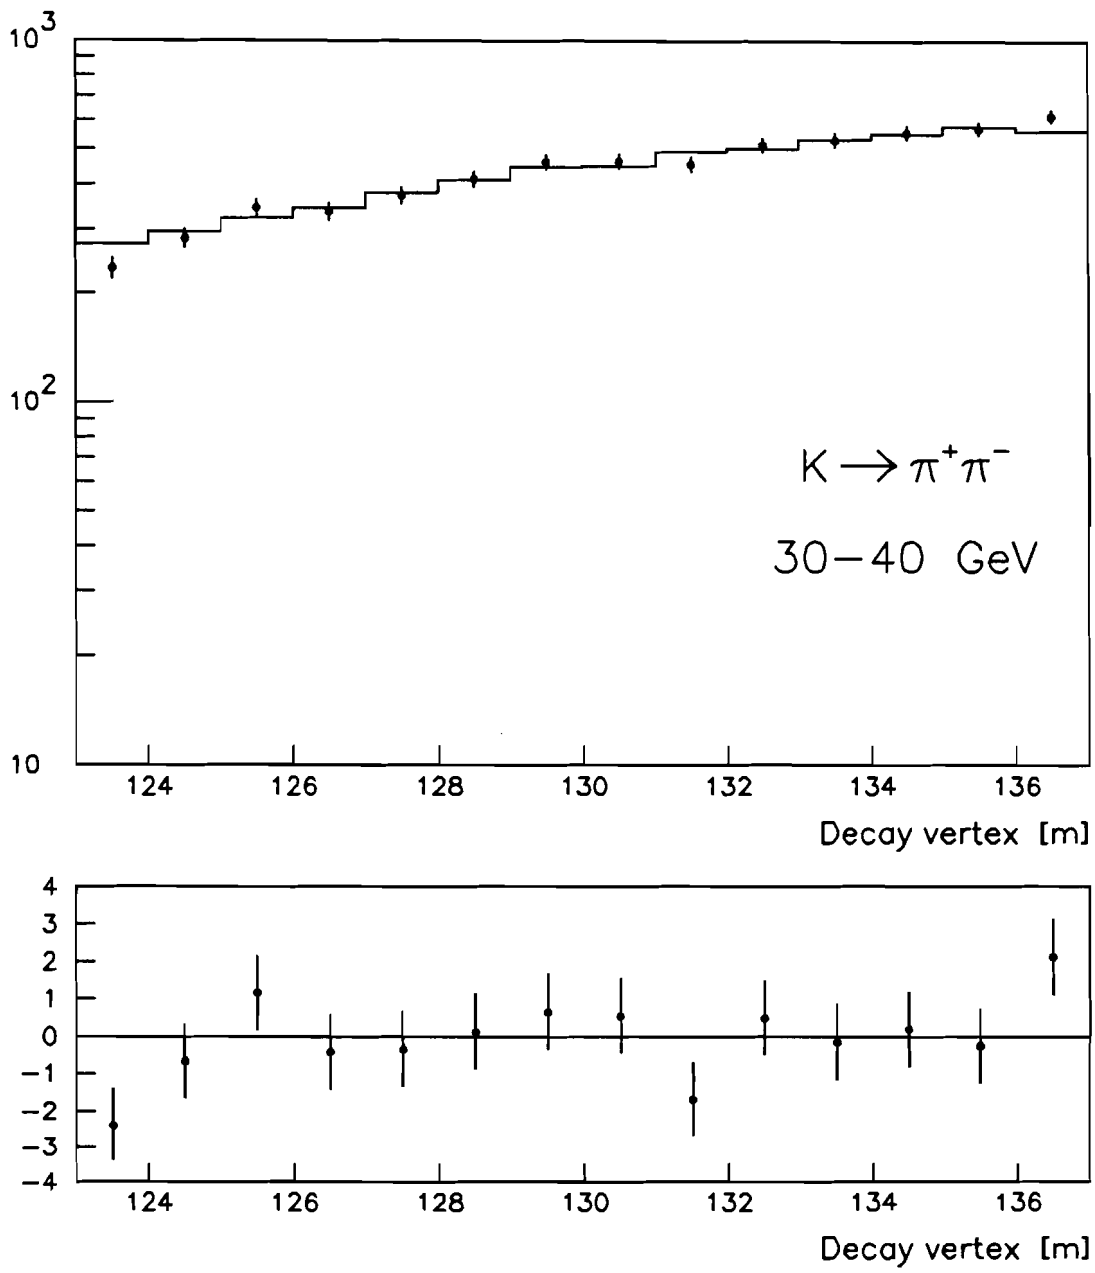

Figure 7.1: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^+\pi^-$ decays between 30 and 40 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

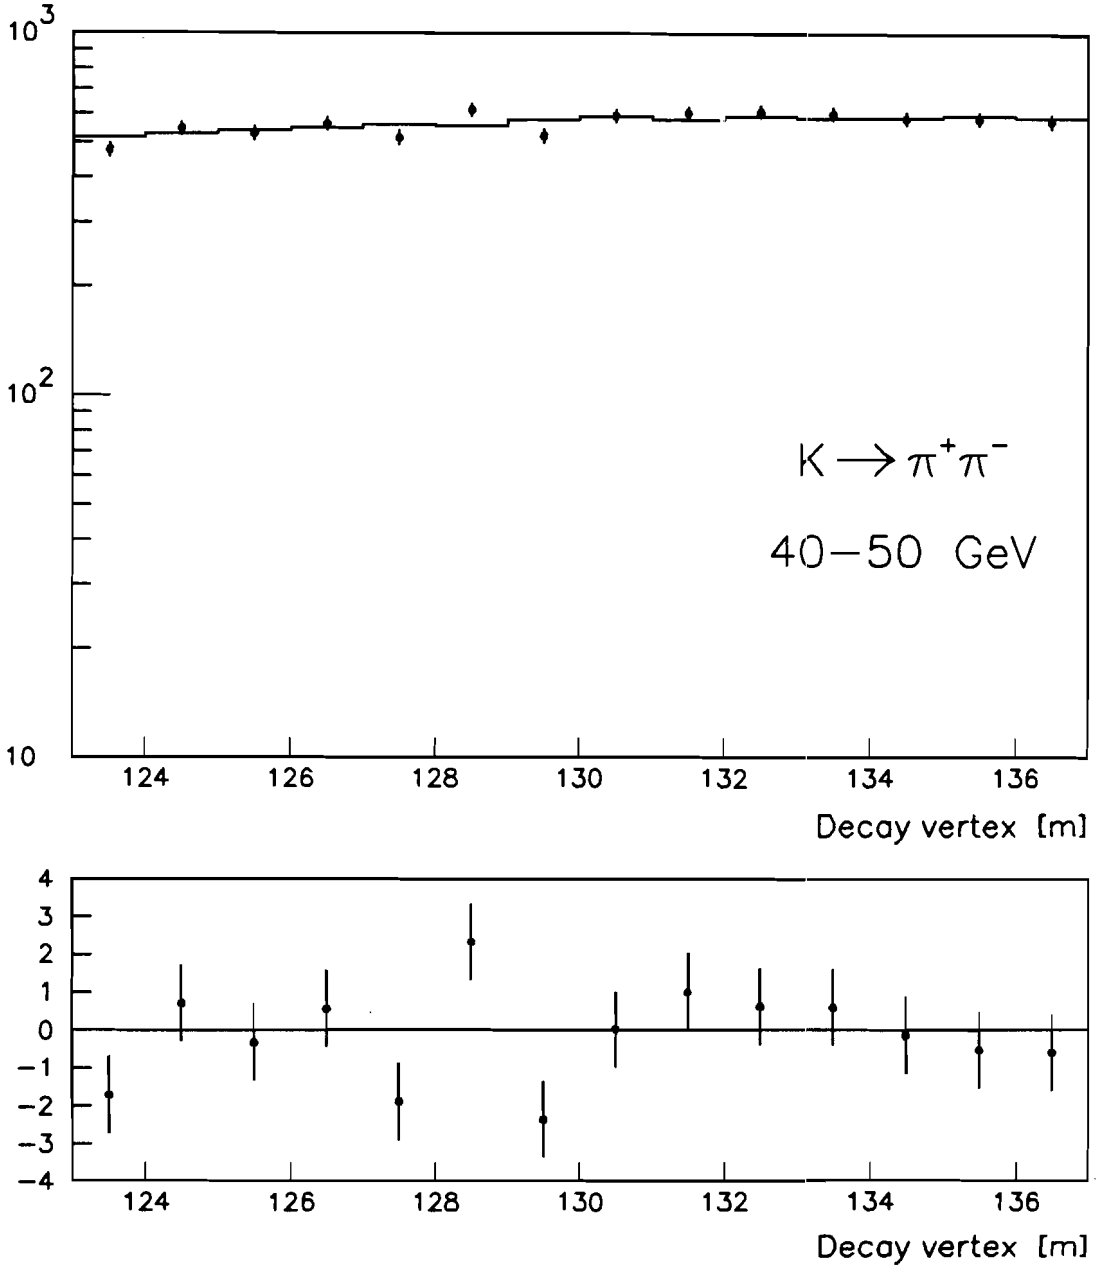

Figure 7.2: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^+\pi^-$ decays between 40 and 50 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

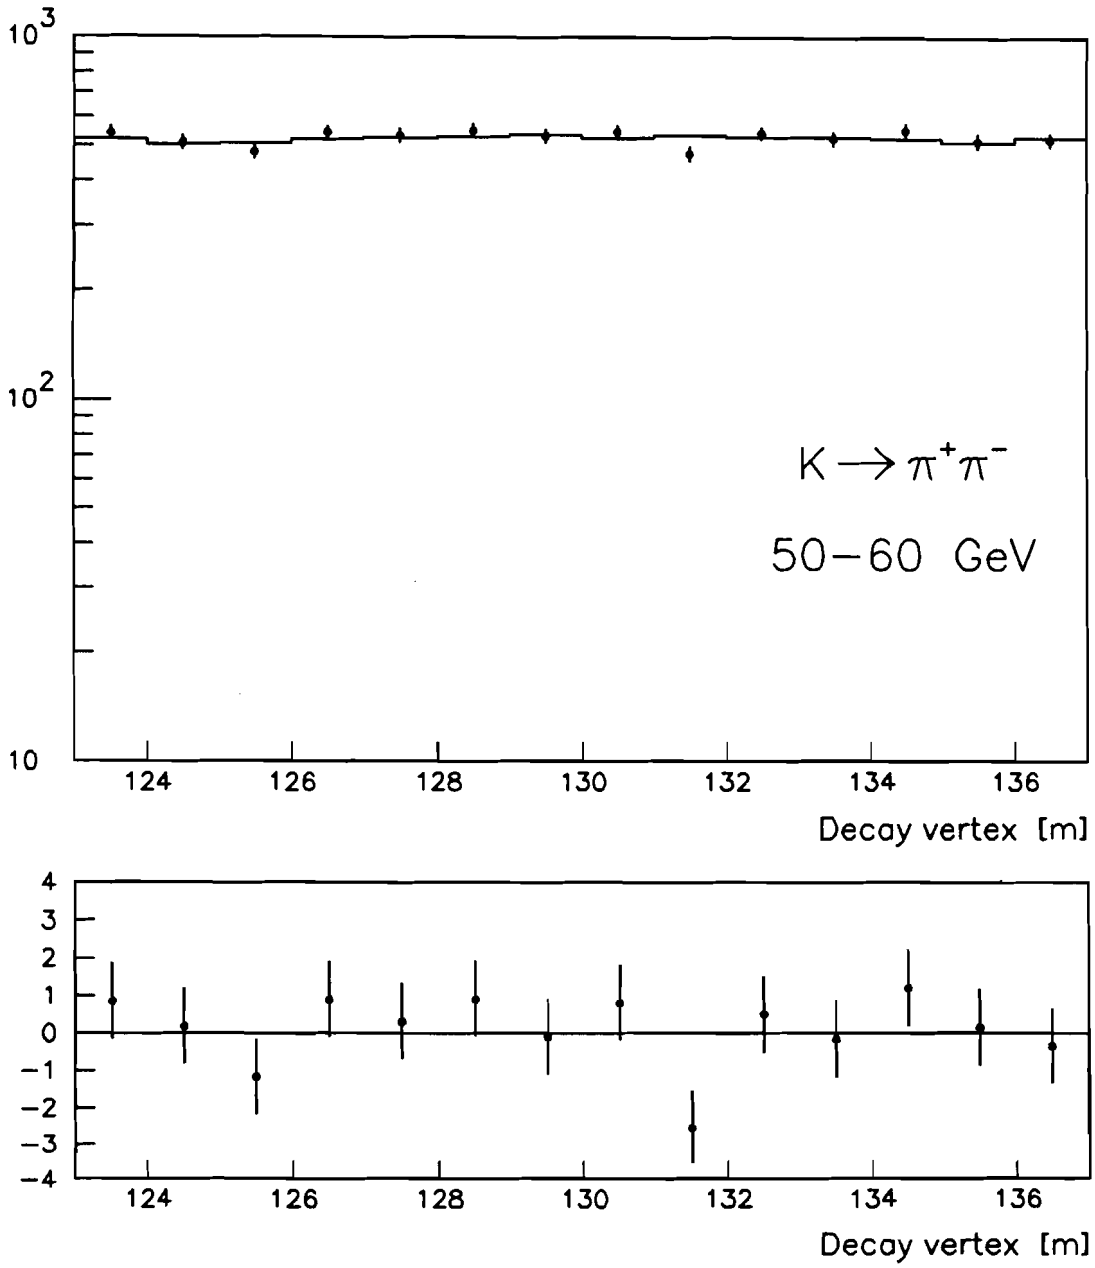

Figure 7.3: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^+\pi^-$ decays between 50 and 60 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

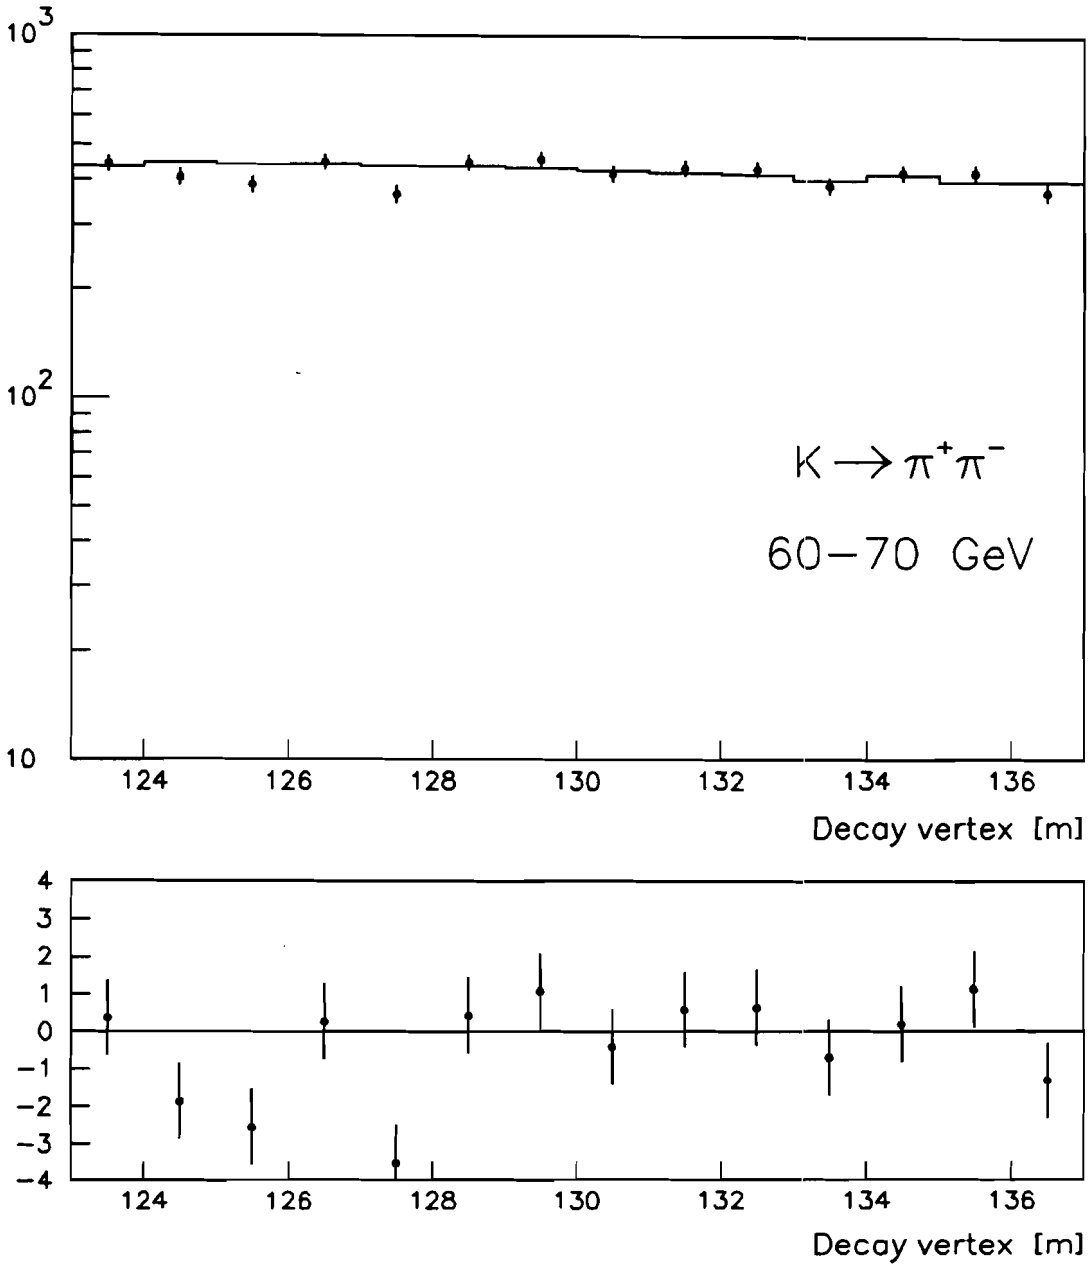

Figure 7.4: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^+\pi^-$ decays between 60 and 70 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

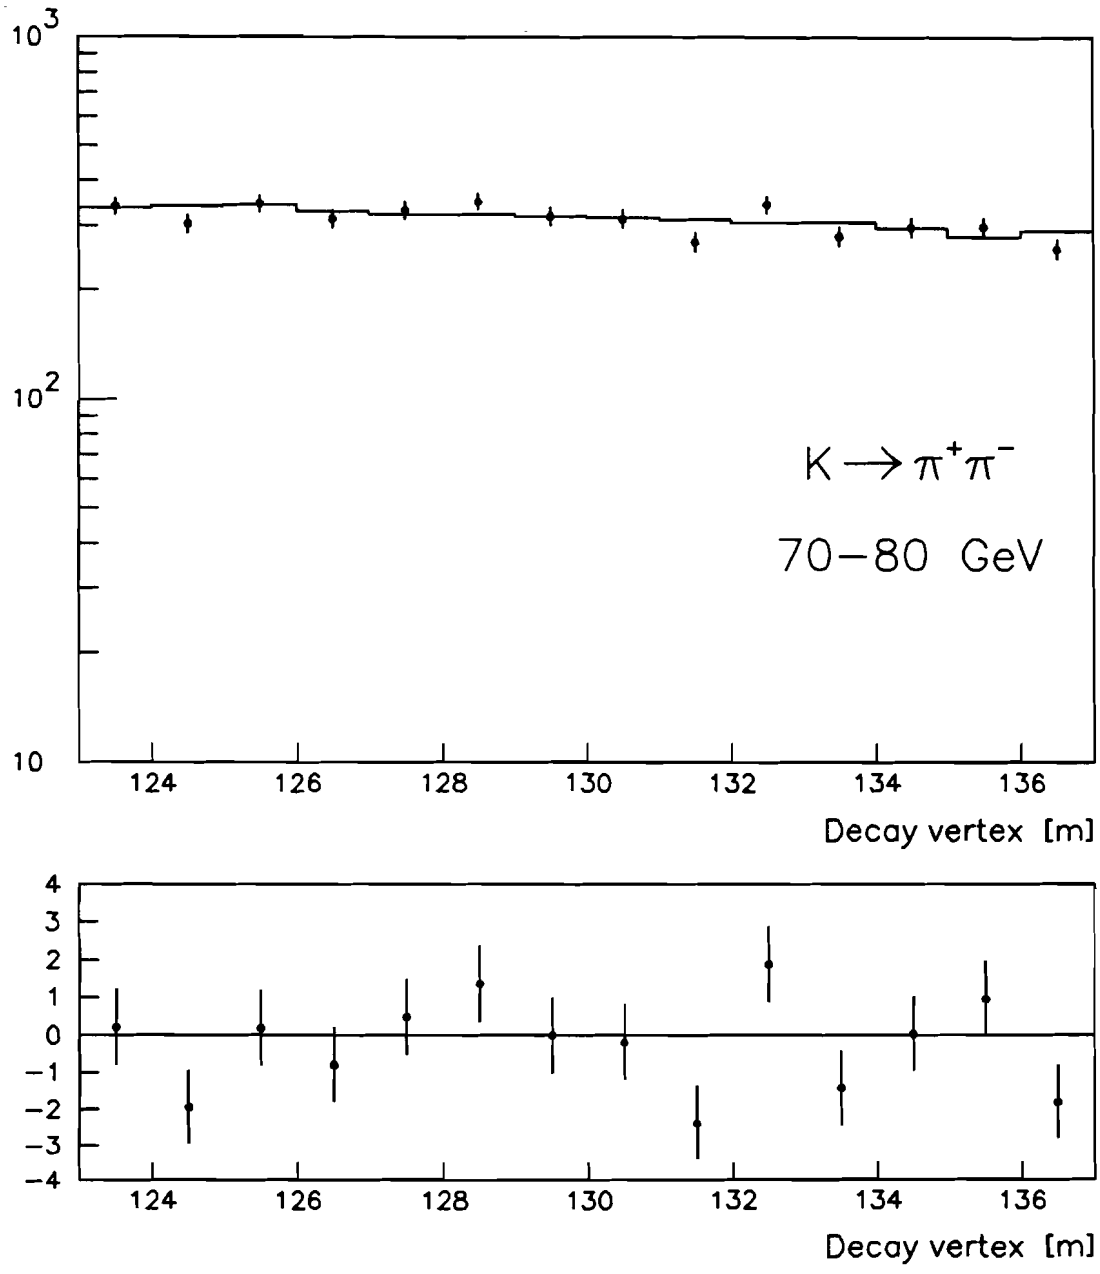

Figure 7.5: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^+\pi^-$ decays between 70 and 80 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

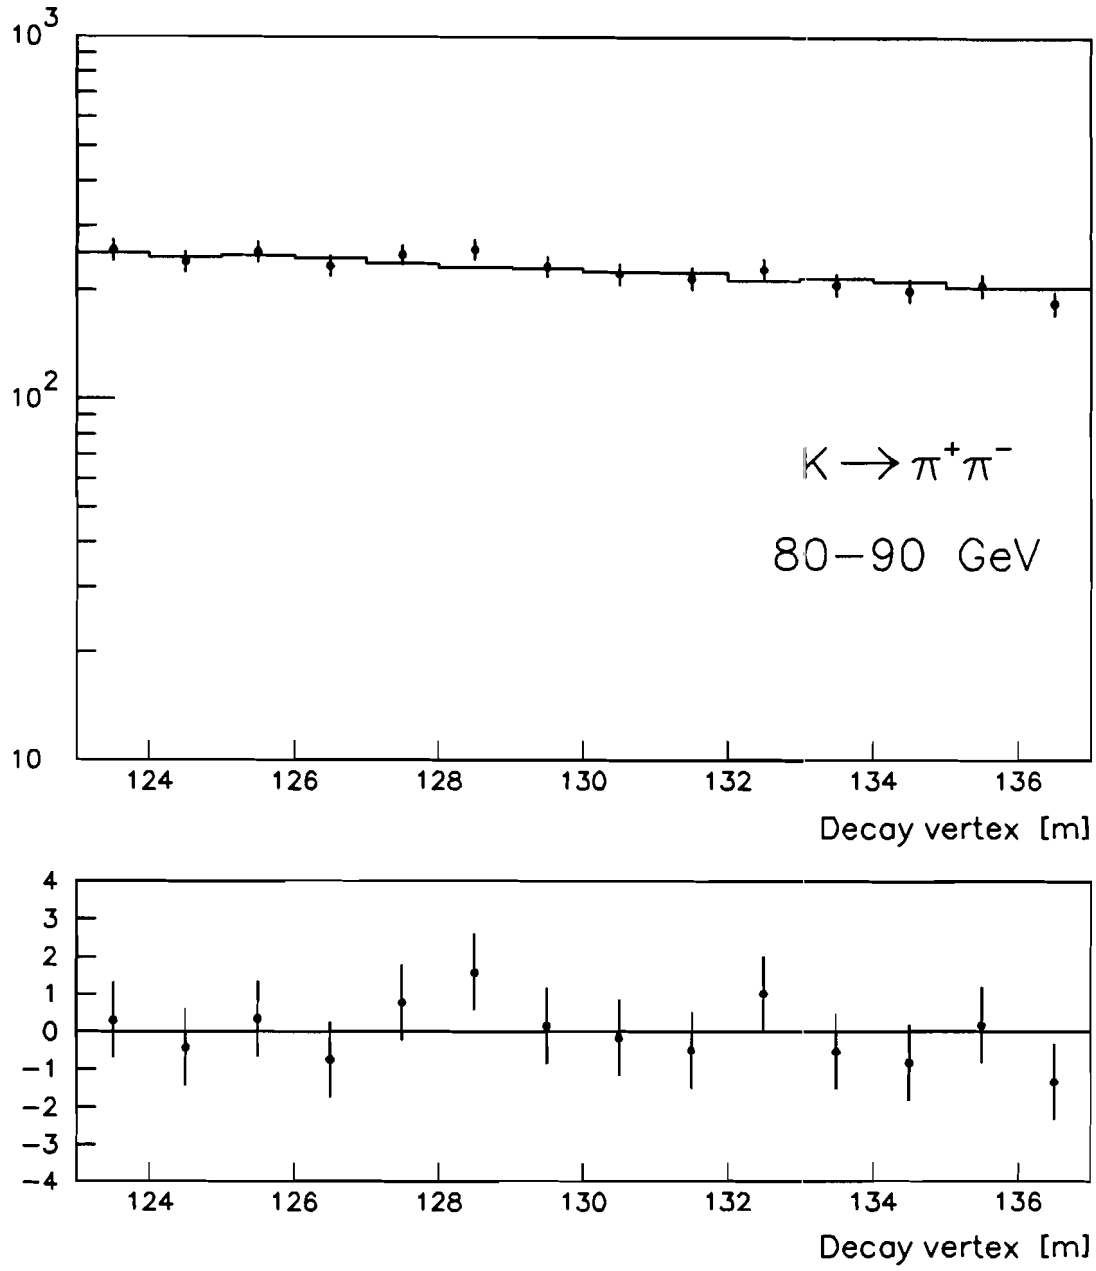

Figure 7.6: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^+\pi^-$ decays between 80 and 90 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

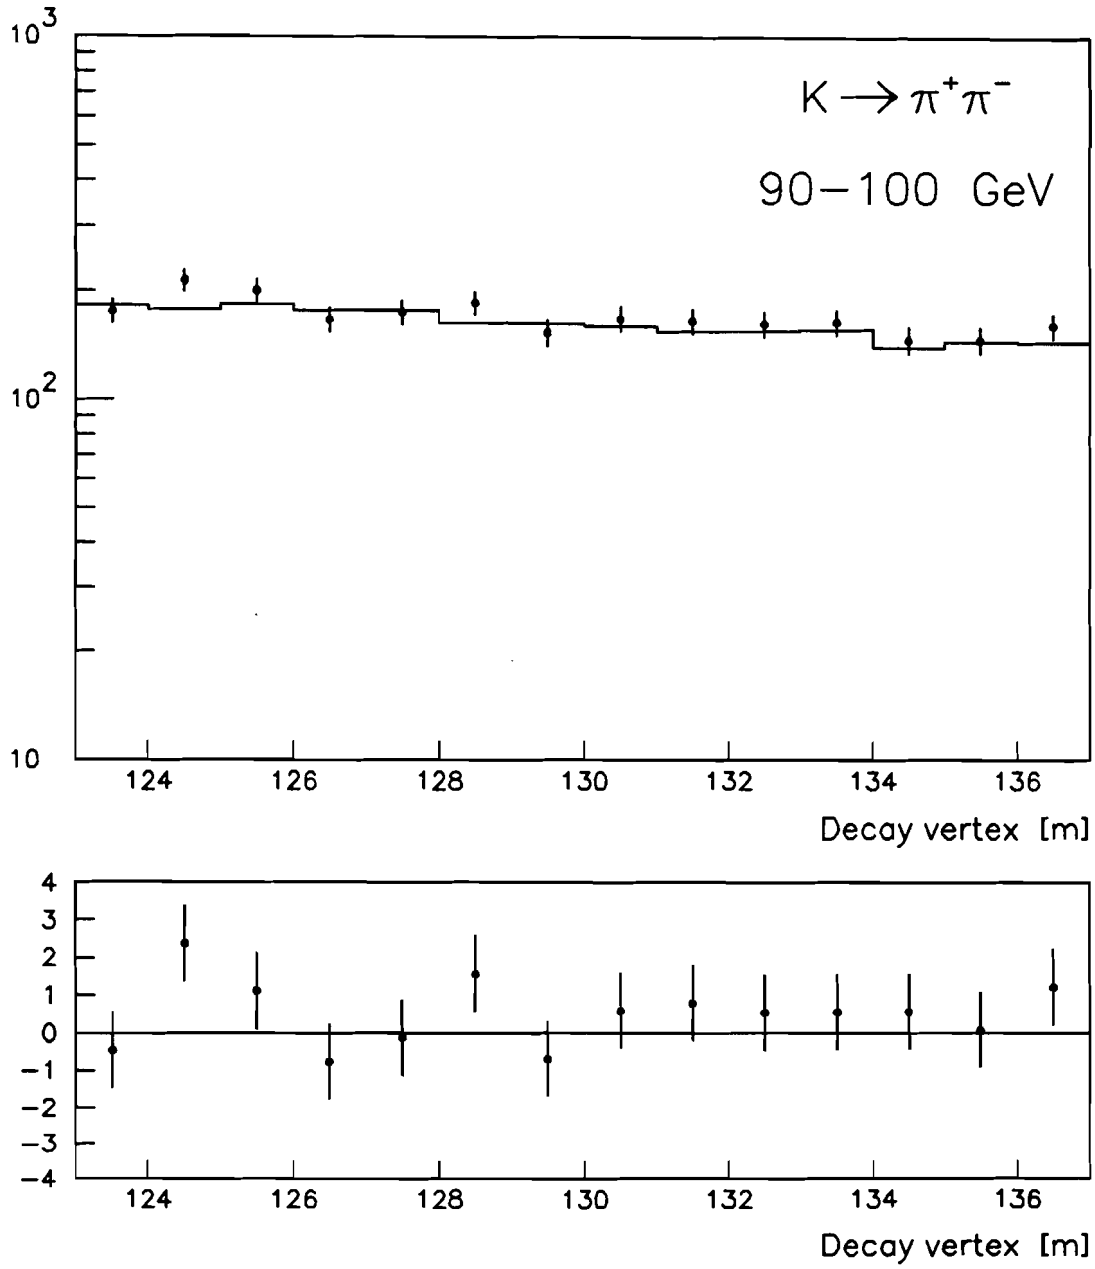

Figure 7.7: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^+\pi^-$ decays between 90 and 100 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

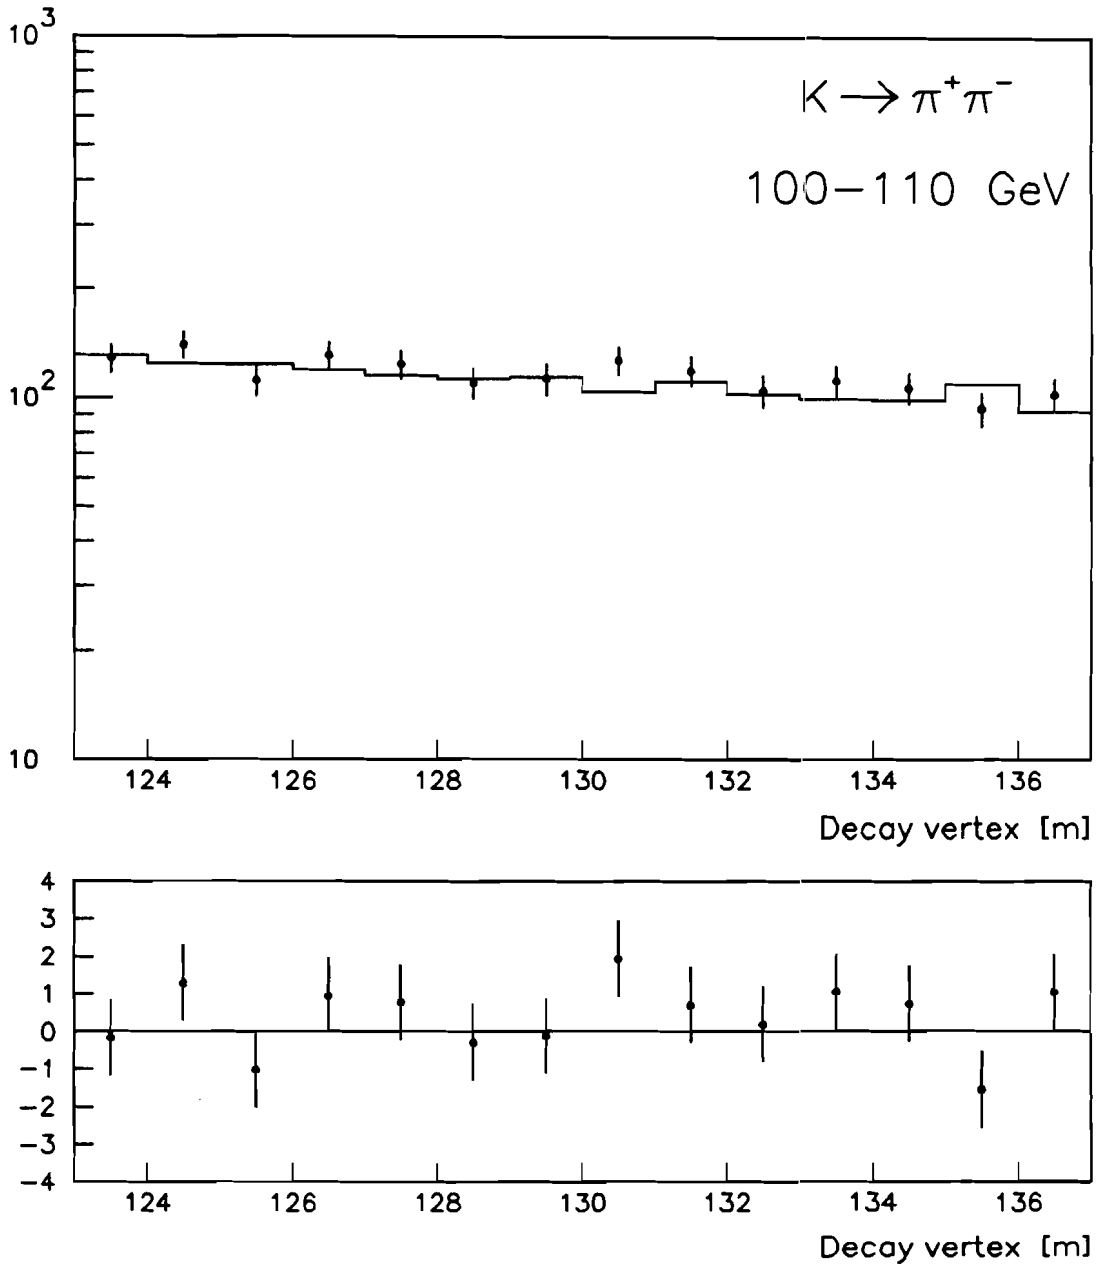

Figure 7.8: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^+\pi^-$ decays between 100 and 110 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

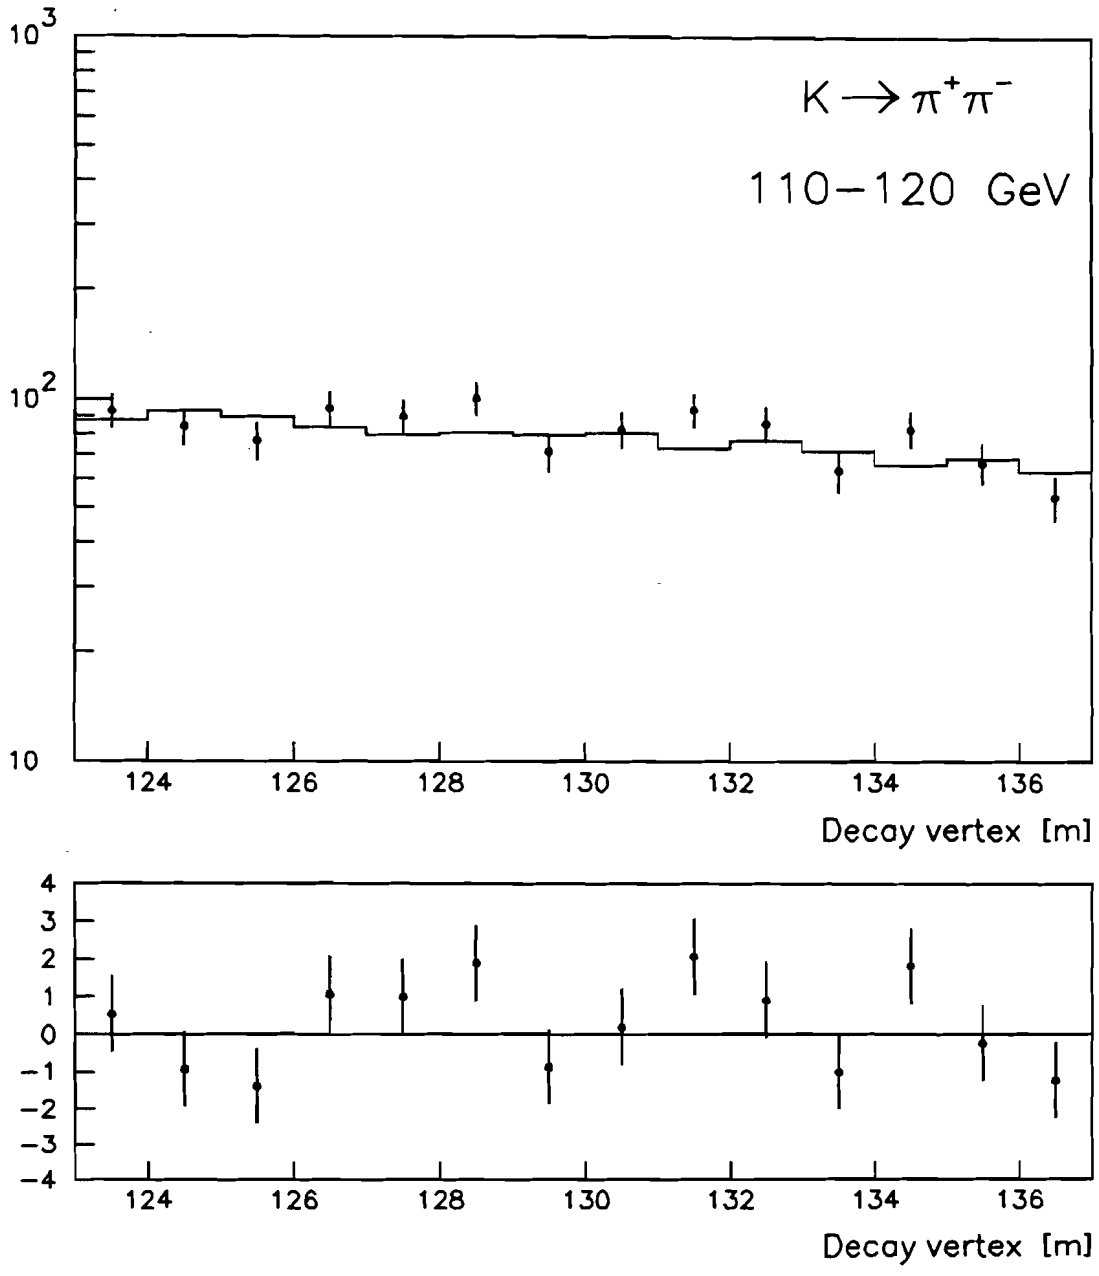

Figure 7.9: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^+\pi^-$ decays between 110 and 120 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

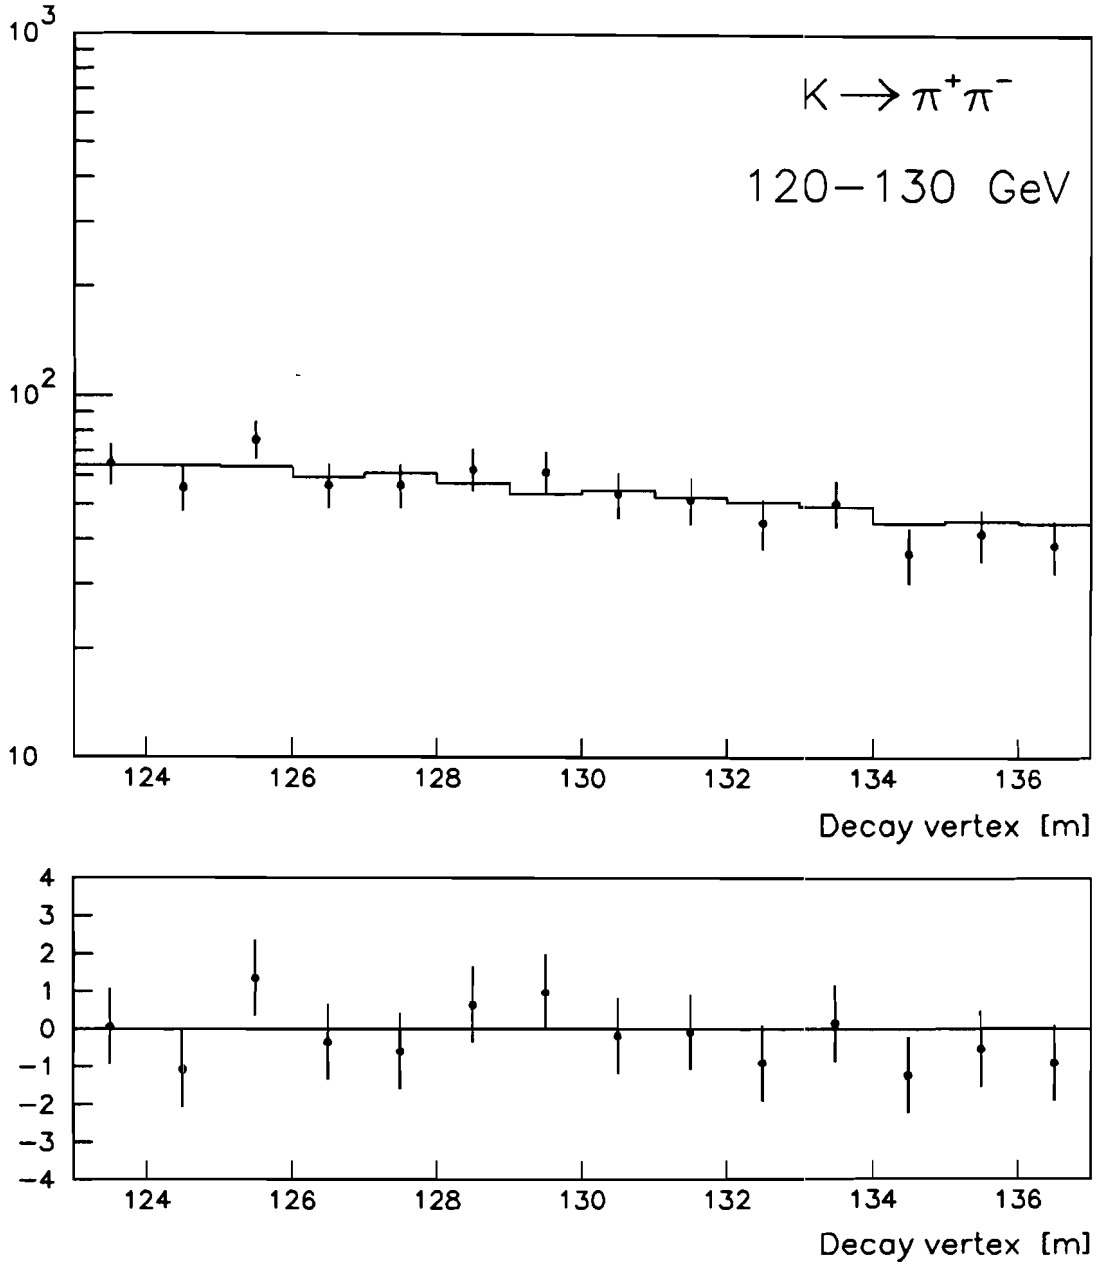

Figure 7.10: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^+\pi^-$ decays between 120 and 130 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

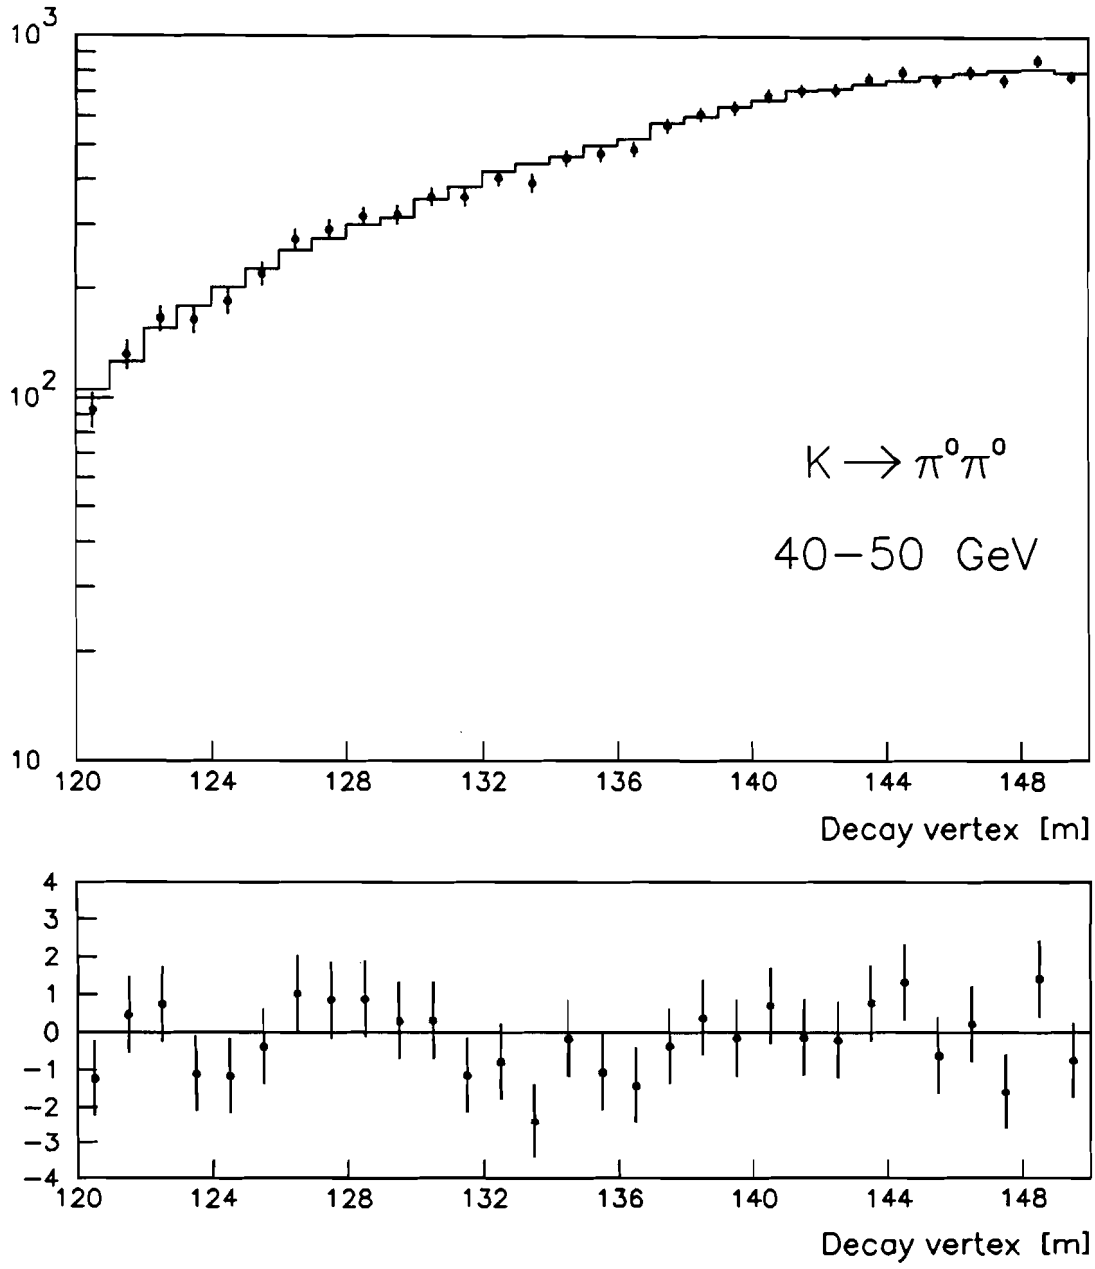

Figure 7.11: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^0\pi^0$ decays between 40 and 50 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

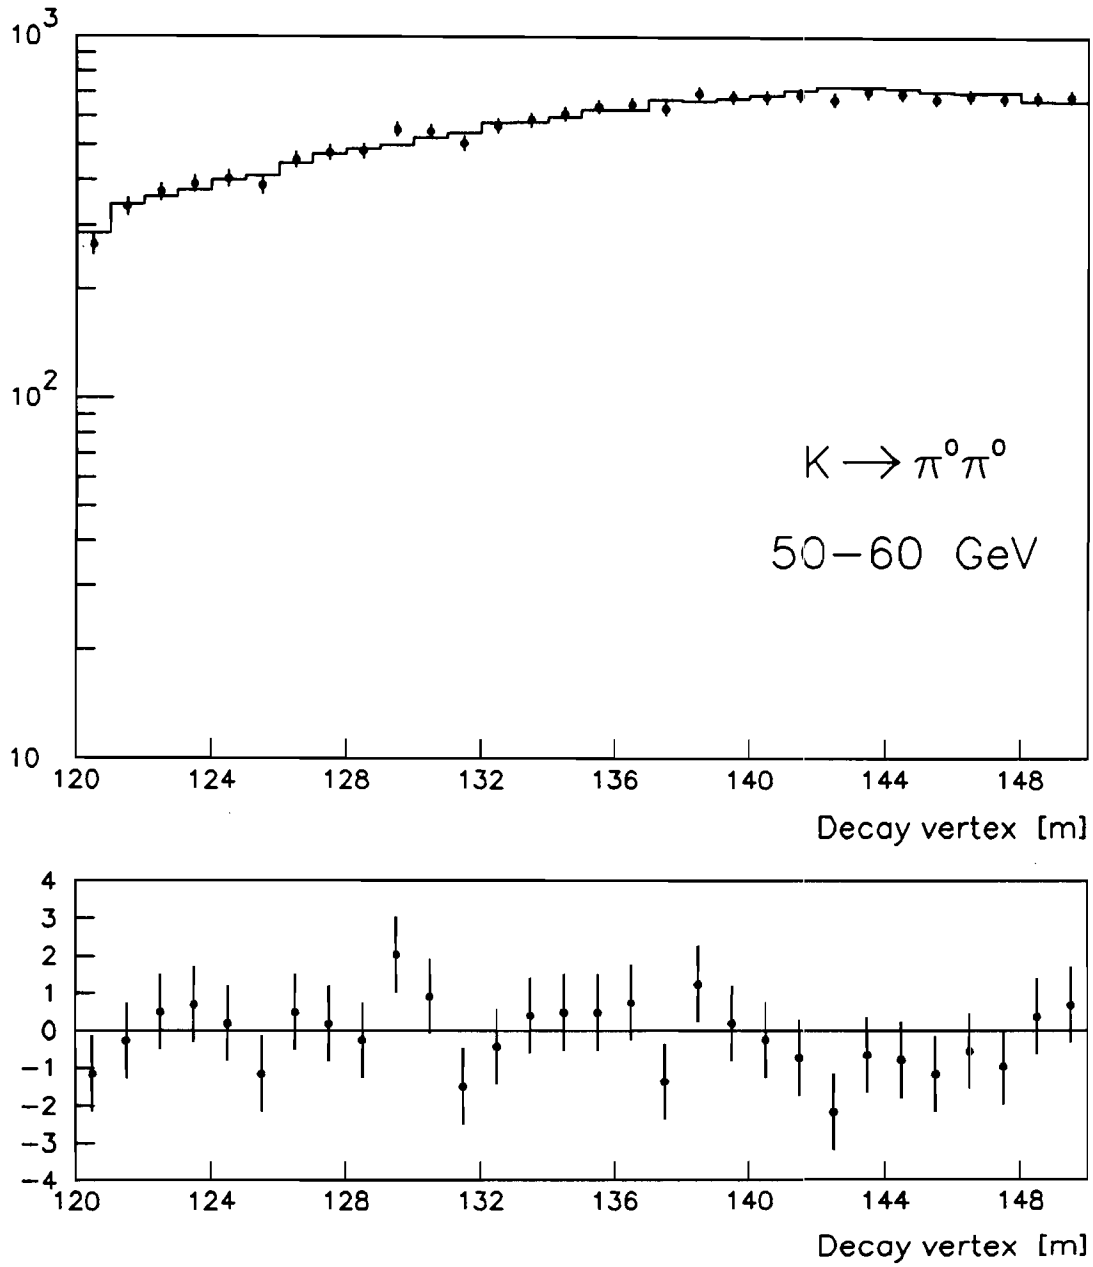

Figure 7.12: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^0\pi^0$ decays between 50 and 60 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

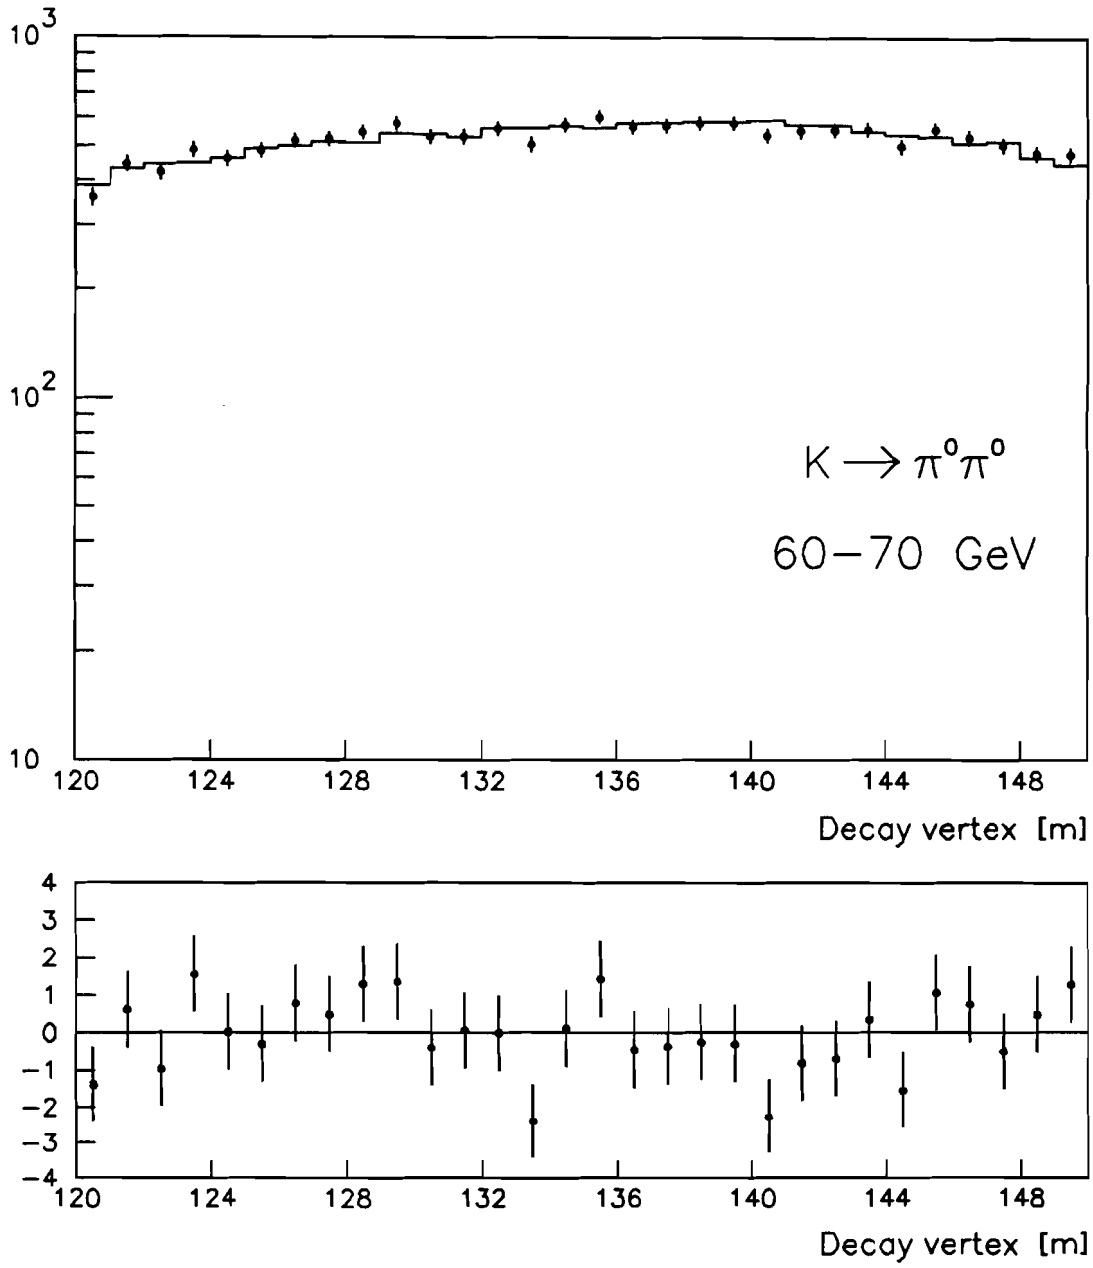

Figure 7.13: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^0\pi^0$ decays between 60 and 70 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

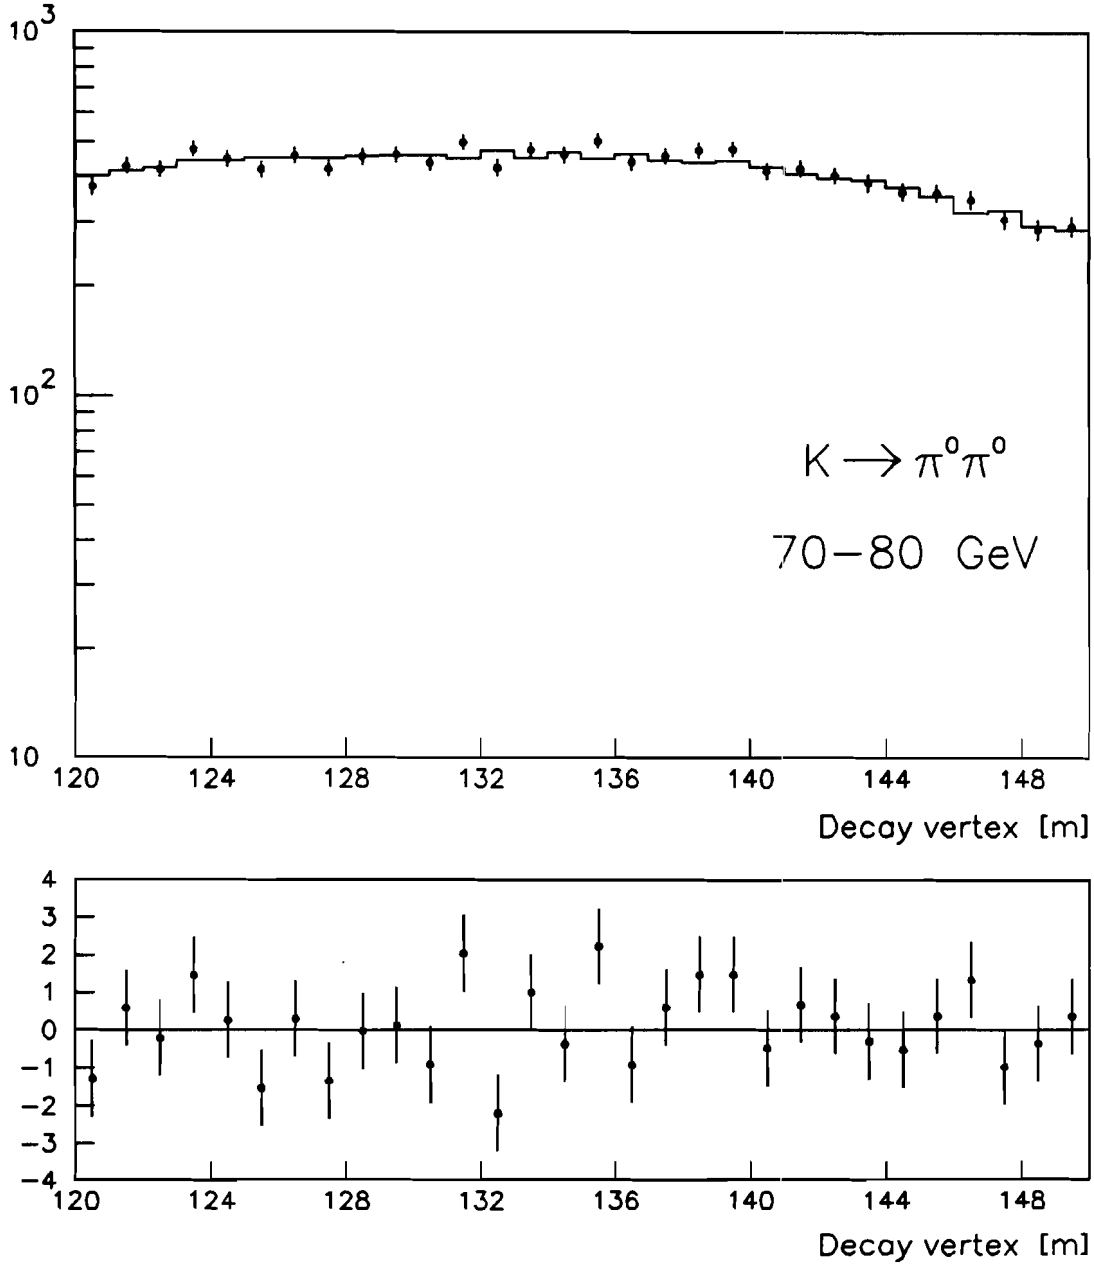

Figure 7.14: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^0\pi^0$ decays between 70 and 80 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

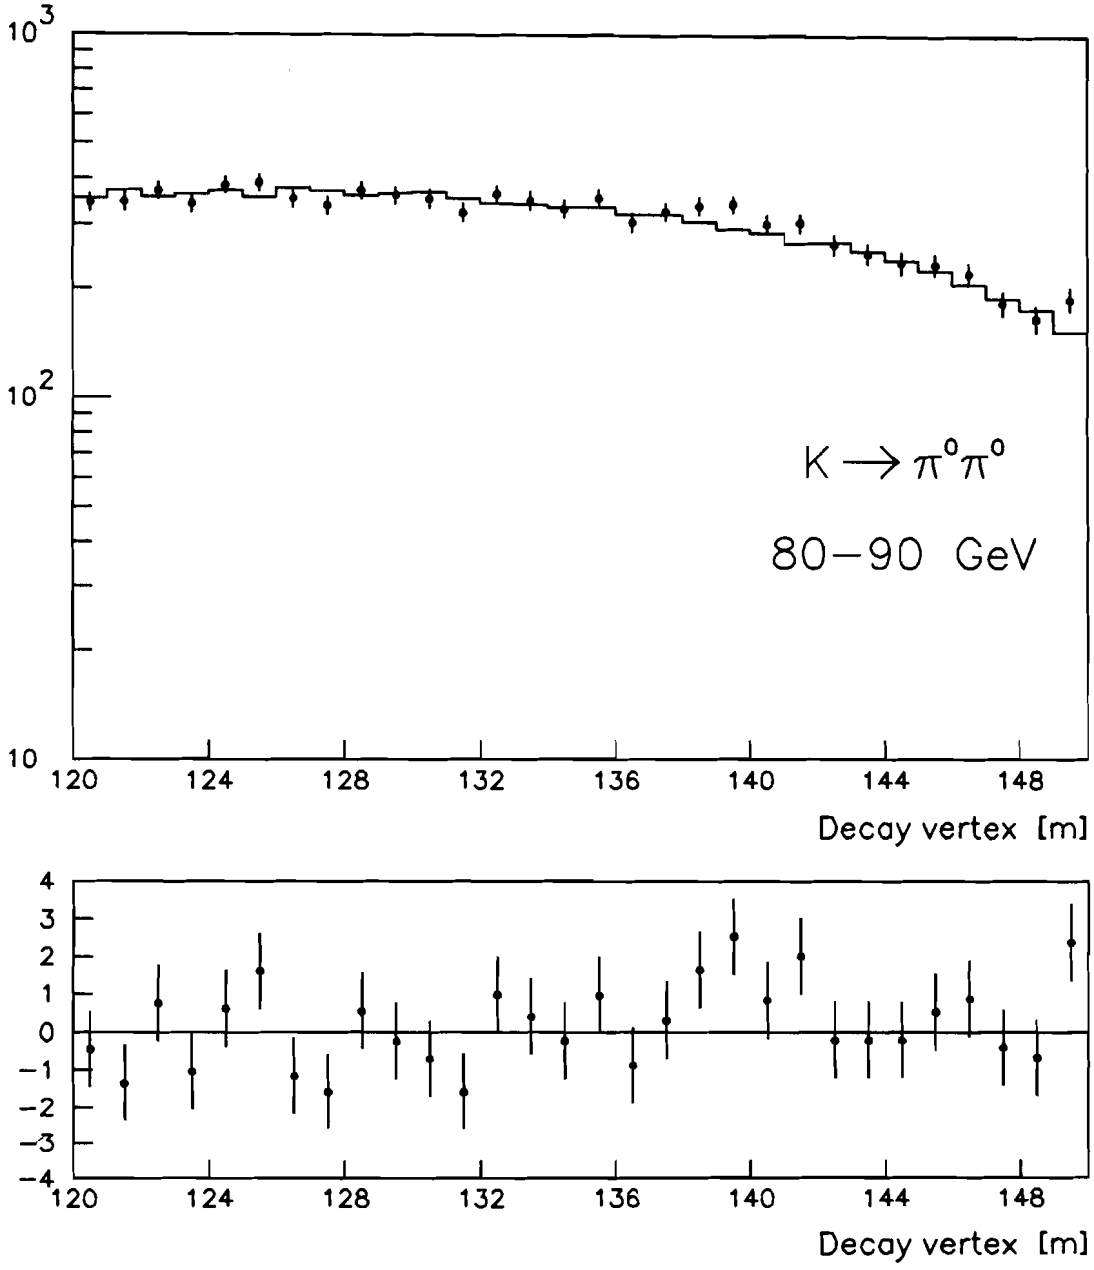

Figure 7.15: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^0\pi^0$ decays between 80 and 90 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

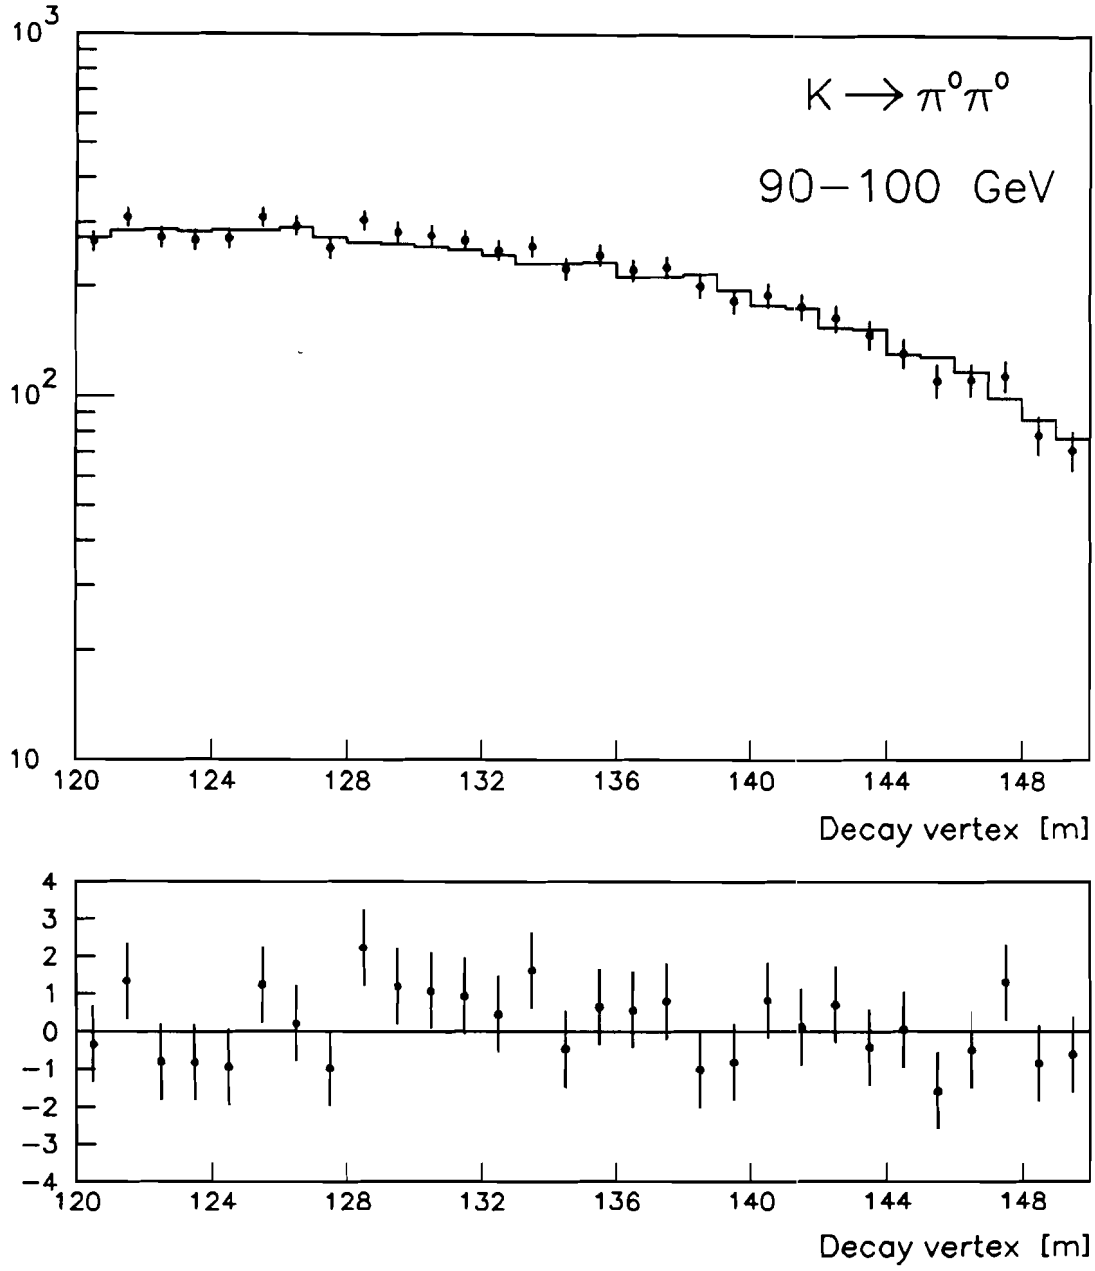

Figure 7.16: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^0\pi^0$ decays between 90 and 100 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

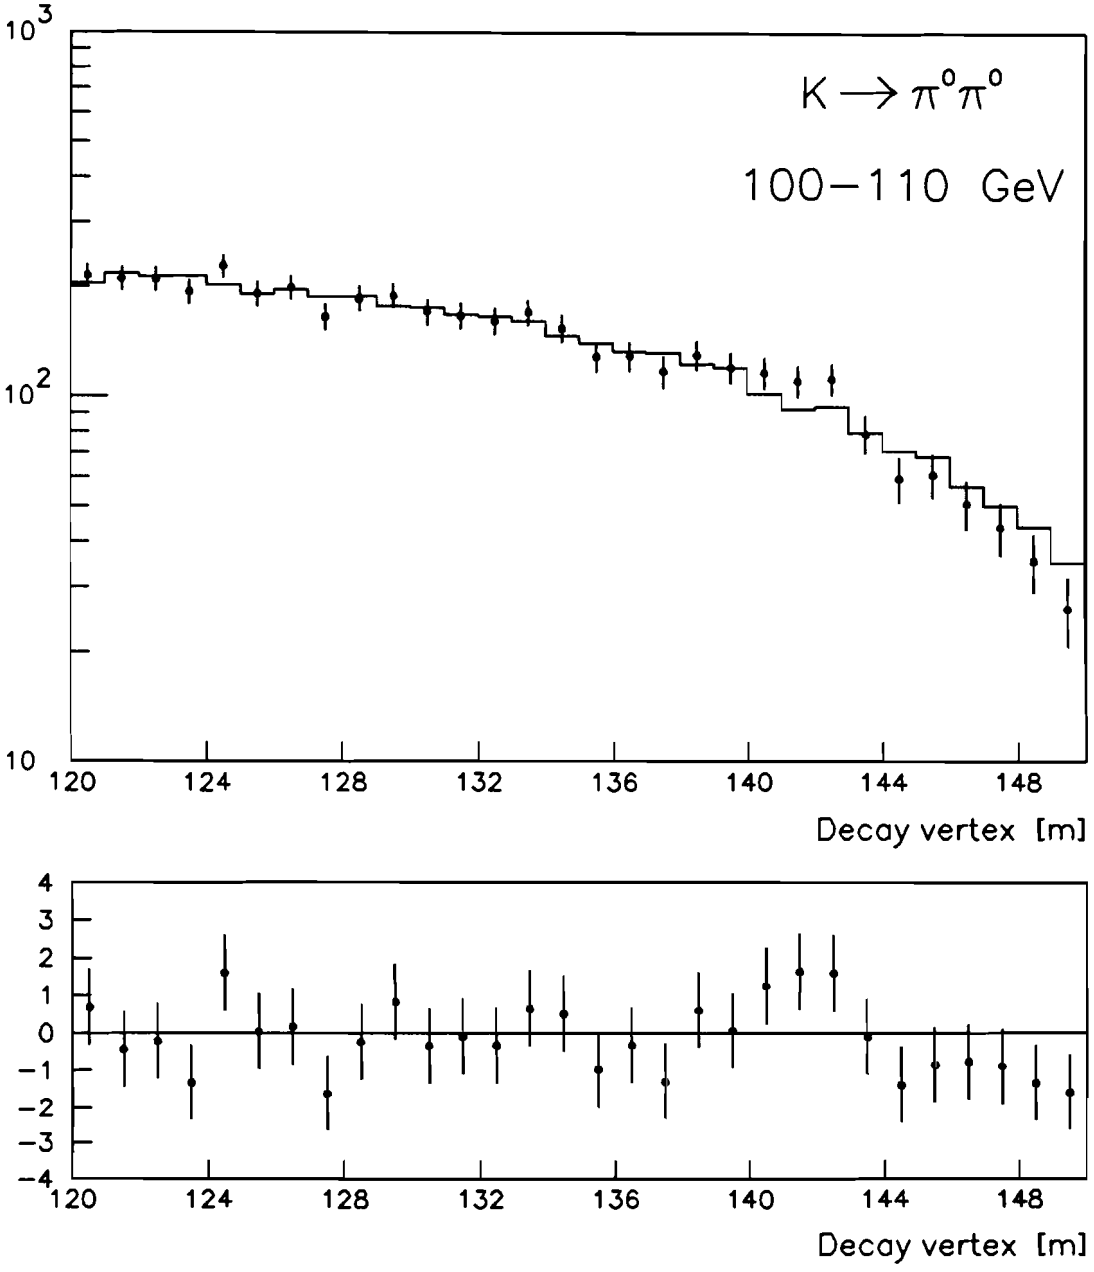

Figure 7.17: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^0\pi^0$ decays between 100 and 110 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

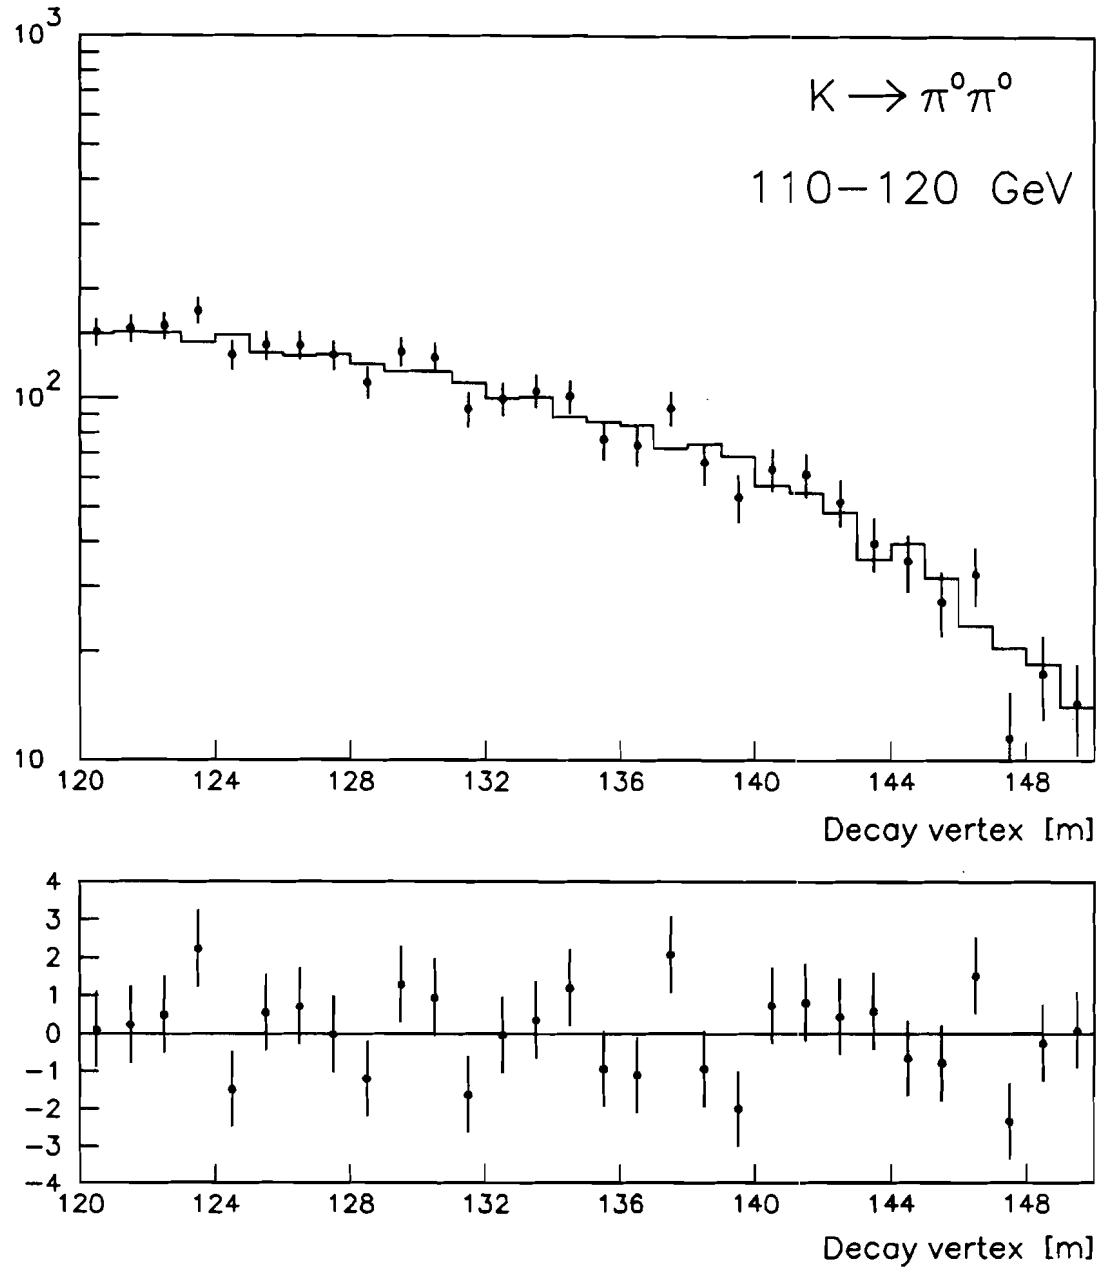

Figure 7.18: *Top:* Data (dots) and Monte Carlo (histogram) decay vertex distributions for $\pi^0\pi^0$ decays between 110 and 120 GeV in the vacuum beam. *Bottom:* Standard deviations for the data-MC difference.

Chapter 8
EXTRACTION OF THE PHASES

In Chapter 6 we obtained the final data samples after background subtractions, and in Chapter 7 corresponding Monte Carlo samples. All of these are binned according to the kaon energy and longitudinal decay vertex with a bin size of 10 GeV by 1 meter. The Monte Carlo samples are intended to imitate the real data as closely as possible, which means that the K_S-K_L interference is included in the event generation, and that all available information about the detector is used for the acceptance simulation. It also means that the Monte Carlo is generated using some nominal values for the parameters that we intend to determine from the data. For example, the phases ϕ_{+-} and ϕ_{00} are both set to 45 degrees. This chapter will describe how to find the parameter values that give the best agreement between the data and the Monte Carlo. Those "best values" will be the result of this experiment.

The first section discusses the various parameters that the predicted decay rates depend on. Parameters that are well known can be fixed, but a few parameters (ϕ_{+-} and ϕ_{00} in particular) will have to be "free." The free parameters are determined from a fit, and the second section describes how that is done. After the result of the fit is presented, a number of systematic uncertainties will be investigated.

8.1 Predicted Decay Rates

The decay rates are most easily written as functions of proper time t and kaon energy E. The proper time is usually measured from the downstream end of the regenerator. Recall that we write the coherent kaon beam coming out of the regenerator as $K_L + \rho K_S$, where $\rho \equiv |\rho| e^{i\phi_\rho}$ is the coherent regeneration amplitude. (Only the coherently scattered kaons are used for the phase determination, for

Table 8.1: Fixed parameters (see text).

Parameter	Value
τ_S	0.8922×10^{-10} sec
τ_L	5.18×10^{-8} sec
Δm	$0.5349 \times 10^{10}\,\hbar\,\mathrm{sec}^{-1}$
e^{-X}	0.0635

reasons that were discussed in the introduction.) The 2π decay rates are

$$|\eta|^2 e^{-t/\tau_L}$$

in the vacuum beam, and

$$e^{-X}[|\rho|^2 e^{-t/\tau_S} + |\eta|^2 e^{-t/\tau_L} + 2|\rho||\eta| e^{-t/2\tau_S} \cos(\Delta m\, t + \phi_\rho - \phi_\eta)]$$

in the regenerated beam. Here, η is either η_{+-} or η_{00}, and ϕ_η is ϕ_{+-} or ϕ_{00}. The factor e^{-X} reflects the difference in flux (about a factor of fifteen) between the two beams due to absorption in the regenerator and the upstream absorber. With the regenerator alternating between the top and bottom positions, the incident kaon momentum spectrum is the same for the regenerated and vacuum beams, except for the flux difference. The energy dependence of the incident kaon flux and the energy dependence of $|\rho|$ have been dropped from the decay rate expressions.

The kaon parameters τ_S, τ_L ($K_{S,L}$ lifetimes) and Δm (K_L-K_S mass difference) are fixed to the values found in reference 8 (see Table 8.1). These are $\tau_S = (0.8922 \pm 0.0020) \times 10^{-10}$ sec and $\tau_L = (5.18 \pm 0.04) \times 10^{-8}$ sec and $\Delta m = (0.5349 \pm 0.0022) \times 10^{10}\,\hbar\,\mathrm{sec}^{-1}$. The factor e^{-X} is determined in the same experiment by comparing the number of $K_L \to \pi^+\pi^-\pi^0$ and $K_L \to \pi^0\pi^0\pi^0$ decays in the two beams. Since the K_S component contributes negligibly to these decay modes, the rate difference between the regenerated and vacuum beams is only due to the difference in K_L flux. The value obtained for e^{-X} is 6.35%, with an uncertainty that is about one percent of the value itself.

Looking at the decay rate expressions, it is seen that all terms are second order in $|\eta|$ and/or $|\rho|$, which means that one of the three parameters $|\eta_{+-}|$, $|\eta_{00}|$

and $|\rho|$ can be fixed to some arbitrary value. We fix $|\eta_{+-}| = 2.275 \times 10^{-3}$. (The actual value chosen is only meaningful if one wants to determine $|\rho|$.) The ratio between $|\eta_{00}|$ and $|\eta_{+-}|$ is related to $\mathrm{Re}(\epsilon'/\epsilon)$ by $|\eta_{00}|/|\eta_{+-}| \simeq 1 - 3\,\mathrm{Re}(\epsilon'/\epsilon)$. The quantity $\mathrm{Re}(\epsilon'/\epsilon)$ is measured in the same experiment, but with the assumption that CPT symmetry is exact (*i.e.*, ϕ_{+-} and ϕ_{00} fixed).[15,29] In the fit described here, it is therefore necessary to let $\mathrm{Re}(\epsilon'/\epsilon)$ be a free parameter. (Since $|\eta_{+-}|$ has been fixed, this means that $|\eta_{00}|$ is free.) The parameter $|\rho|$ depends on the kaon energy E according to the power law $|\rho| = \rho_0 E^{-\alpha}$ (see chapter 2.5). Since ρ_0 and α are not very well known, they are allowed to be free in the fit. (Usually, ρ_0 is taken to be the value of $|\rho|$ at 70 GeV, so that $|\rho| = \rho_0(E/70)^{-\alpha}$, with E in GeV.)

It is the two phases ϕ_{+-} and ϕ_{00} that we wish to determine from the fit, but the decay rate expressions only contain the differences $\phi_{+-} - \phi_\rho$ and $\phi_{00} - \phi_\rho$. The regeneration phase ϕ_ρ can be written $\phi_\rho = \pi/2 + \phi_f + \phi_g$ (chapter 2.5). The geometric phase ϕ_g results from the integration over a thick regenerator and is calculated from well known parameters. The phase ϕ_f is the argument of $(f - \overline{f})/k$ and is not very well known. For this fit, we use the prediction $\phi_f = -(2-\alpha)\pi/2$, with the same α as in the power law for $|\rho|$. This constrains ϕ_ρ to α (which is free).

The description above is in terms of kaon energy and proper time, but the data is binned according to kaon energy and decay vertex. However, for any given kaon energy, the decay vertex position and the proper time are simply proportional to each other, and the decay rates can easily be expressed in terms of energy and vertex position.

8.2 Fitting Method

The previous section introduced the five parameters that will be determined from the fit: ϕ_{+-}, ϕ_{00}, $\mathrm{Re}(\epsilon'/\epsilon)$, ρ_0 and α. We would now like to vary these parameters and find out what set of parameter values gives the best agreement between the predicted rate and the observed data. However, we cannot compare the predicted rate directly with the data, since the observed data sample also depends on the detector acceptance, which is not the same in all bins. There are two ways

to correct for the acceptance. We either remove the acceptance function from the data or include it in the predicted rate. The fit described here uses the second approach, but let us consider briefly how the first one would work.

To remove the acceptance function from the data, we would need to know, for each bin, what the ratio is between the number of observed events and the number of actual decays. This ratio is available directly from the Monte Carlo simulation. Each Monte Carlo event is generated with "true" values for the kaon energy E and the longitudinal vertex position z. Due to finite detector resolution, the reconstructed E and z values will in general be different from the true values. For a particular energy and vertex bin, the acceptance is then defined as the number of Monte Carlo events with reconstructed E and z in that bin divided by the number of events with generated E and z in the same bin. This means that the acceptance also includes effects from event smearing, *i.e.*, events that are generated in one bin but reconstructs in another. To remove the acceptance function from the data, one would simply divide the data by the acceptance in each bin, and the resulting true decay distribution could then be compared with the predicted decay rate. For each set of parameter values, one would integrate the predicted rate over each bin and take the difference between this prediction and the corrected data. By minimizing the chi-square, the best set of parameter values would be obtained.

The procedure appears to be quite straightforward, but there are a couple of complications. One is that the acceptance (as defined above) depends slightly on the parameters that are being determined. This is because the resolution-smearing between bins depends on the shape of the decay distribution. As long as the Monte Carlo is generated with parameter values that are close to the true values, this should not be a big effect. Another question is how to integrate the predicted decay rate over a bin. The integration over decay vertex is easily done analytically, but the integration over kaon energy would have to be done numerically, since the kaon energy spectrum is obtained from the data. One possible simplification is to do the decay vertex integration only for the average kaon energy, but then one would have to worry a little about the low energy bins, where $|\rho|$ changes quite a bit within the bin. A third complication is what to do with bins for which there

are no true decays. This happens at the regenerator, where there is an edge in the decay distribution due to the veto counters attached to the end of the regenerator. Because of resolution, some events will reconstruct in bins upstream of this edge. These events cannot be used in the fit, since the acceptance is not well defined for such bins.

To avoid these complications, we prefer to include the acceptance function in the predicted rate, and then fit against the uncorrected data. One very direct way of doing this would be to generate many Monte Carlo samples, each one with a different set of parameter values. These event samples would all be compared to the observed data and a chi-square determined for each one. When there is a sufficient number of grid points in the free parameter space, an interpolation procedure would find the set of parameters corresponding to the minimum chi-square. The only problem is that an enormous amount of computer time is required to generate enough Monte Carlo samples, especially with several free parameters in the fit. (The fit will of course have to be done many times, since the Monte Carlo is periodically improved, and since there are also systematic effects to investigate.)

A better approach is to simulate the detector only once, and then adjust this standard Monte Carlo sample to reflect changes in the free parameters. This is done by "re-weighting" the Monte Carlo events, which have been generated with a nominal set of parameter values. For each event, we know the true (*i.e.*, generated) energy and decay vertex. For this particular combination of E and z, we calculate the predicted decay rate both with the nominal parameters and with a new set of parameters. The event weight (which is always one in the standard sample) is then multiplied by the ratio of the new rate to the nominal rate. For example, if the new rate is higher than the nominal rate, the event is assigned a weight greater than one. In other words, if this particular event is more likely to occur with the new set of parameters, its weight is increased. By binning the re-weighted events, a new Monte Carlo sample is obtained without having to repeat the detector simulation. (Note that there are no statistical fluctuations between the re-weighted Monte Carlo samples.)

As described, this method is still quite time consuming, since it would be

necessary to run through all events in the standard Monte Carlo sample for each new set of parameter values. However, it is also possible to do the re-weighting *after* the events have been binned. This is done by calculating the new and nominal rates only once per bin, at a point specified by the average generated E and z values. (These average quantities are obtained from the standard Monte Carlo sample.) The weight of the whole bin (*i.e.*, the number of events) is then adjusted to reflect the change in decay rate. (Since the binned sample is stored in a small array, this re-weighting is done quickly for each new parameter set. Note also, that the rates are calculated only at individual points and never integrated over bins.) The minimization then proceeds in a normal fashion, with a chi-square calculated for each set of parameter values. The program MINUIT is used to determine which parameter set minimizes the chi-square.

How good is the bin-by-bin re-weighting compared to the event-by-event one? If the bins are small enough, the difference should be negligible. That this is so can be verified by doing the event-by-event re-weighting once to obtain an alternative standard Monte Carlo sample (preferably with nominal parameters "on the other side" of the minimum). After running the minimization again using the bin-by-bin re-weighting, but starting from the new standard sample, the new result can be compared with the old one. This has been done with our data, and no difference was observed. The test could also be done by generating an entirely new Monte Carlo sample with a different set of nominal parameters, but in this case there would be a statistical fluctuation between the two Monte Carlo samples, and the difference between the two results would not necessarily be zero. With the re-weighting, there is no statistical fluctuation, and the two results should be identical.

It was stated above that there are five free parameters. In the fit, it is also necessary to have an overall normalization parameter for each decay mode (charged and neutral), since the total number of Monte Carlo events is arbitrary relative the total number of data events. (These two extra parameters are labelled N_{+-} and N_{00} below.) The Monte Carlo samples have six to seven times more events than the data, which is enough to ensure that the Monte Carlo does not contribute significantly to the total error. Even so, the statistical error from the Monte Carlo should always

Table 8.2: Result from the fit described in the text. The errors are one standard deviation statistical errors.

Parameter	Central Value	Error
ϕ_{+-}	47.7°	2.0°
ϕ_{00}	47.4°	1.4°
$\mathrm{Re}(\epsilon'/\epsilon)$	-0.2×10^{-3}	1.6×10^{-3}
ρ_0	1.167 mb	0.003 mb
α	0.587	0.003
N_{+-}	0.1900	0.0008
N_{00}	0.1371	0.0008
$\chi^2 = 316$ for 341 degrees of freedom		

be included in the chi-square calculation. (The chi-square is actually calculated in 10 GeV×2 m bins, in order to make sure that no low-statistics bins introduce non-linearities into the minimization. All other calculations are performed with 10 GeV×1 m bins.)

It was mentioned in the introduction that the parameters are determined from a direct comparison between the data and Monte Carlo distributions, and that this requires the Monte Carlo to simulate the detector acceptance very accurately. An alternative approach would be to first take the ratio between the regenerated and vacuum beams in each bin, in order to have the detector acceptance essentially cancel out. However, this would increase the statistical uncertainty, since the total number of events is lower in the vacuum beam.

8.3 The Result

The result of the fit (see Table 8.2) is $\phi_{+-} = 47.7 \pm 2.0°$ and $\phi_{00} = 47.4 \pm 1.4°$, with a χ^2 of 316 for 341 degrees of freedom. (The value 341 is the number of bins minus the number of free parameters. However, since the Monte Carlo energy spectrum is taken from the data, the "correct" number of degrees of freedom is probably somewhat lower.) The charged and neutral modes are almost decoupled in the fit. In fact, the only constraint that connects the two modes is the requirement that the parameter α be the same for both. To get the correct statistical error for

the difference between ϕ_{+-} and ϕ_{00}, the minimization program is also run with $\Delta\phi \equiv \phi_{00} - \phi_{+-}$ replacing ϕ_{00}. The result is $\Delta\phi = -0.3 \pm 2.4°$.

The value for $\mathrm{Re}(\epsilon'/\epsilon)$ from this fit is very close to that obtained when CPT is assumed to be exact. (The latter result[15,29] is -0.4×10^{-3}.) The parameter ρ_0 is given as the magnitude (at 70 GeV) of the regeneration amplitude per atom in the regenerator (averaged over boron and carbon). The value in the table should be decreased by about 2% to correct for the fact that regeneration in the RA veto counters (attached to the regenerator) is not explicitly taken into account by the fitting program. (This is irrelevant for the phase determination.) The parameters N_{+-} and N_{00} are simply the ratios of the total number of data events to the total number of Monte Carlo events for the charged and neutral modes, respectively.

Figures 8.3–20 at the end of this chapter show the data and Monte Carlo decay vertex distributions in each energy bin. The data (solid circles) and the best fit (histogram) are shown together with the best fit without the interference term (dotted histogram). The statistical fluctuations in the best fit are due to the finite number of events in the Monte Carlo sample. The decay vertex distributions are also shown after subtracting the distribution without interference and dividing by the magnitude of the interference term. (The magnitude is calculated from the parameters obtained in the fit.) The resulting distributions correspond to the cosine factors in the decay rates. Due to resolution and smearing effects, the reconstructed decay vertex is not completely proportional to the true proper time of the decay. In particular, the most upstream bins contain events that decayed just downstream of the regenerator edge. This is the reason why the cosine curves level off towards the left in the figures.

Since we know the average true proper time for each bin, it is possible to plot the cosine factor as a function of proper time instead of decay vertex. This is done in Figure 8.1, with several energy bins combined in one plot. The highest energy bins are not included because of their larger error bars due to lower event statistics. (In order to combine different energy bins, it is necessary to correct for the energy dependence of the geometric phase ϕ_g.) A convenient rule of thumb is that one K_S lifetime is equivalent to about 27 degrees.

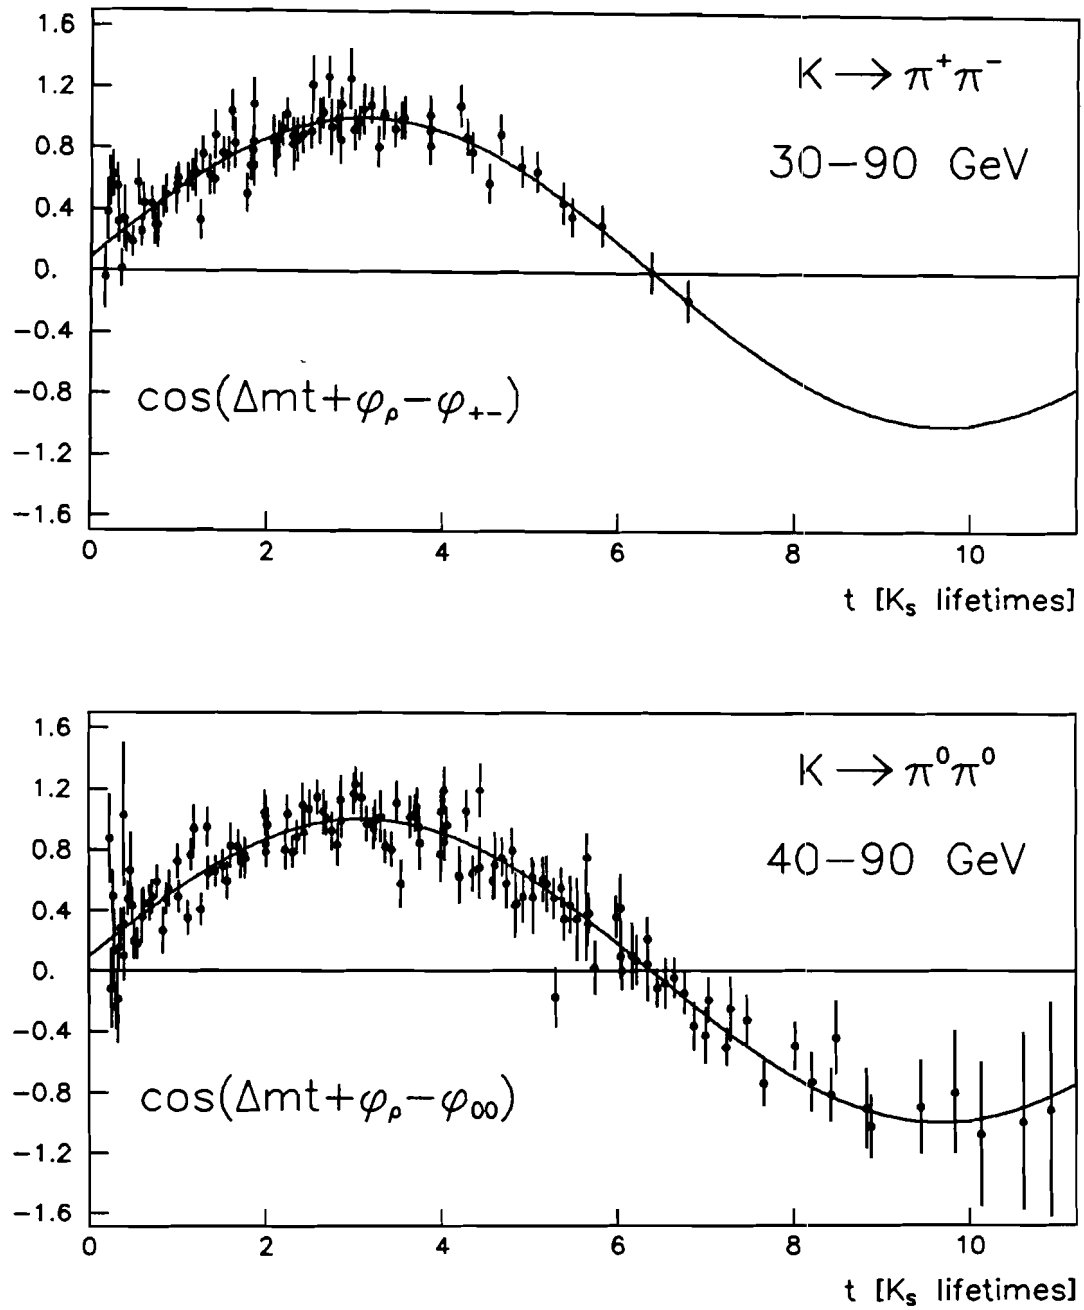

Figure 8.1: Extracted cosine factors from the $\pi^+\pi^-$ decay rate (*top*) and the $\pi^0\pi^0$ decay rate (*bottom*).

Table 8.3: Systematic errors for ϕ_{+-}, ϕ_{00} and $\Delta\phi$.

Source	ϕ_{+-}	ϕ_{00}	$\Delta\phi$
Δm	0.4°	0.5°	0.1°
τ_S	0.8°	0.6°	0.2°
τ_L	0.03°	0.02°	0.01°
e^{-X}	0.4°	0.3°	0.2°
Acceptance	0.8°	0.4°	0.9°
Energy scale	0.2°	0.6°	0.6°
Energy resolution	0.1°	0.3°	0.3°
Background subtraction	0.1°	0.3°	0.3°
Total (excluding ϕ_ρ)	1.3°	1.2°	1.2°

8.4 Systematics

Table 8.3 is a summary of the systematic uncertainties for ϕ_{+-}, ϕ_{00} and the difference $\Delta\phi \equiv \phi_{00} - \phi_{+-}$. The lifetimes τ_S and τ_L and the mass difference Δm are varied one standard deviation around their world average values, and the resulting systematic errors are listed in the table. An increase in Δm (or a decrease in τ_S) leads to an increase in both ϕ_{+-} and ϕ_{00}. There is a coupling between $\Delta\phi$ and $\mathrm{Re}(\epsilon'/\epsilon)$, such that $\Delta\phi$ increases by 0.7° for each increase of 1×10^{-3} in $\mathrm{Re}(\epsilon'/\epsilon)$. This is not a systematic uncertainty, since $\mathrm{Re}(\epsilon'/\epsilon)$ is a free parameter in the fit. The factor e^{-X}, which represents the flux difference between the two beams, is known to about one percent of itself. The corresponding errors in ϕ_{+-} and ϕ_{00} are 0.4° and 0.3°, respectively.

There are a few small effects that are ignored by the decay rate expressions given in the first section. Even though most of the target K_S decay before the start of our decay volume, there is a small contribution to the decay rate at high kaon energies. Another effect is the small amount of regeneration in the trigger hodoscopes at the end of the upstream decay volume. (This is only relevant for the neutral mode.) The Monte Carlo and the fitting program treat both of these effects correctly by following the evolution of the kaon state all the way from the target to the end of the decay volume. The fitting program also corrects for a small beam tail, which has the effect that about 0.1% of the incident K_L flux misses the regenerator.

Because that part of the beam is not attenuated by the regenerator, the K_L decay rate in the regenerated beam is increased by about one percent. (This correction is only for the charged mode, since the neutral mode has a center of energy cut which automatically excludes events from the beam tail.) None of these special decay rate corrections have a big effect on the phases (a few tenths of a degree at the most), and since they are quite well understood, the related systematic uncertainties are negligible.

The main concern is whether the detector acceptance is simulated correctly. We have seen from the chi-square of the fit that the agreement between the data and Monte Carlo distributions is excellent. If we used this agreement to limit the acceptance uncertainty, we would find systematic errors of about two degrees in both modes. Better limits are obtained by checking the Monte Carlo simulation with the 10^7 $\pi e \nu$ and 6×10^6 $\pi^0\pi^0\pi^0$ decays, which were collected at the same time as the 2π decays. These comparisons show that the acceptance function is understood to better than one percent across the full range of decay vertex and kaon energy, and the resulting systematic uncertainties are 0.8° for ϕ_{+-} and 0.4° for ϕ_{00}. (These errors are limited by the statistics of the $\pi e \nu$ and $\pi^0\pi^0\pi^0$ samples.) The charged and neutral mode errors add in quadrature for $\Delta\phi$, since the acceptance functions in the two modes are related to different detector elements.

There is some uncertainty due to the energy calibration in the neutral mode, where both the decay vertex and the kaon energy are calculated from the energy measurements in the lead glass. (In the charged mode, the decay vertex is determined with much higher accuracy from the tracking information. The charged mode energy scale is also quite well known, since the magnetic field is calibrated by using the kaon mass as a constraint.) The neutral mode energy scale and resolution are studied by adding extra scale factors and extra energy smearing in the analysis. The altered data samples are then fit against the Monte Carlo sample, in order to find the effect on the phases. The systematic uncertainty for ϕ_{00} is 0.6° due to the energy scale and 0.3° due to the resolution. There is also a small effect on ϕ_{+-}, since the charged and neutral modes are coupled to each other in the fit.

The effect of the background subtractions can be divided into two parts. The

uncertainties in the average levels of the various backgrounds are small compared to the uncertainty in the factor e^{-X}. (The average background levels and the factor e^{-X} only affect the normalization between the vacuum and regenerated beams.) As discussed above, the e^{-X} uncertainty is taken to be one percent, which leads to errors of 0.3–0.4 degrees for the phases. The average background levels are all small (not above 2–3%) and are understood well enough to make the uncertainties much less than one percent. This means that we only have to worry about the shape of the backgrounds in the regenerated beam. (The background shape in the vacuum beam does not affect any of the parameters in the fit.) Varying the background shapes by reasonable amounts, we find a systematic error of 0.3° for ϕ_{00} and less than 0.1° for ϕ_{+-}.

Accidental activity is found to have a negligible effect on the phases. There is a special accidental trigger (chapter 4.3.3), which records accidental events together with the 2π data. The events are triggered by a coincidence between two small scintillator counters placed in such a way that a particle going through the counters misses the detector by a wide margin. The rate for this trigger is therefore proportional to the $K \to 2\pi$ rates, but uncorrelated with any particular detector element. Possible effects on the relative normalization between the two beams or on the shape of the data distributions are looked for by overlaying the events from the accidental trigger on Monte Carlo 2π events. No effects are seen within the statistical limits of the investigation.

Two assumptions were made for the regeneration amplitude ρ. The first one says that $|\rho|$ has a power law dependence on the kaon energy, and this is clearly supported by our data, since the chi-square of the fit is good. One could relax this assumption and let $|\rho|$ be a free parameter in each energy bin. When this is done, the central values obtained for the phases are consistent with the ones from the standard fit (but with a small increase in the statistical errors). Because no significant deviation from the power law is seen, we do not assign a systematic error to this assumption.

The other assumption is that the phase of ρ is determined by the value of α (the slope parameter in the power law). In particular, we assume that ϕ_ρ is independent

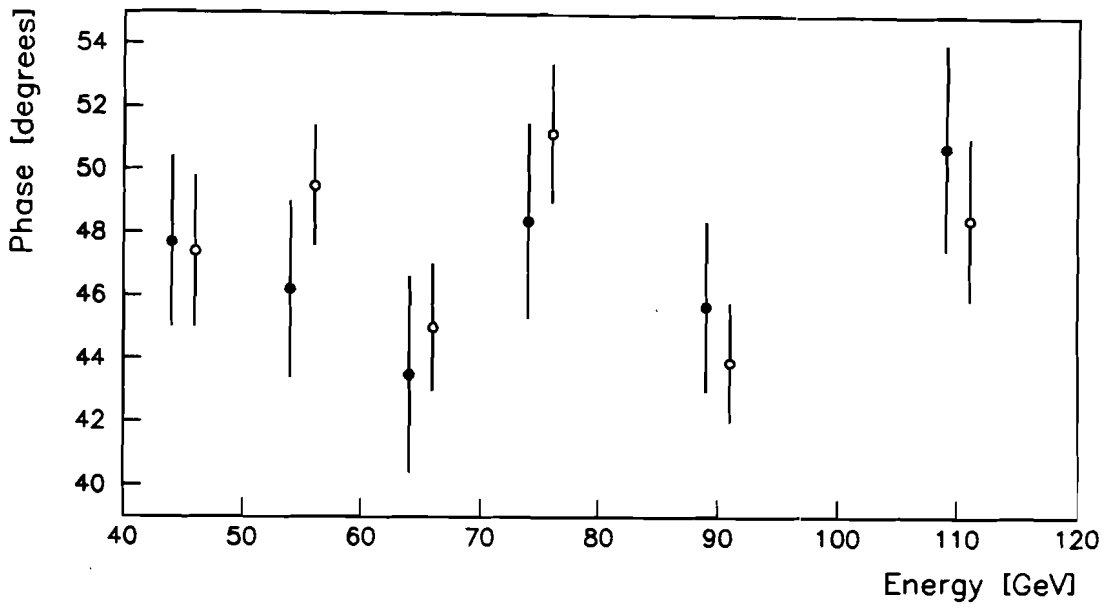

Figure 8.2: ϕ_{+-} (solid circles) and ϕ_{00} (open circles) obtained in separate energy bins.

of the kaon energy. Figure 8.2 shows the phase values that are obtained from fits in separate energy bins. (For these fits, all parameters except the two phases are fixed to the values obtained in the standard fit. Otherwise, there would be no statistical precision in the higher energy bins.) If ϕ_ρ changed with energy, there should be an energy dependence visible in the plot, but no such effect is seen. That ϕ_{+-} and ϕ_{00} appear to track each other is most likely caused by the fact that the kaon energy spectrum used in the Monte Carlo is not perfectly tuned. (Recall that a fourth order polynomial was used to tune the energy spectrum.) Since the phases have no significant energy dependence, we do not assign a systematic error to the assumption that ϕ_ρ is constant. However, some kind of systematic error should be associated with the limited knowledge of the value of ϕ_ρ. Since ϕ_{+-} and ϕ_{00} depend directly on ϕ_ρ, we are not willing to completely trust the value calculated from α. It is difficult to give a precise estimate of this uncertainty, but the range 2–3 degrees seems reasonable. For $\Delta\phi$, fortunately, the ϕ_ρ uncertainty cancels in the difference, and the total error given in Table 8.3 is also the final systematic error for $\Delta\phi$.

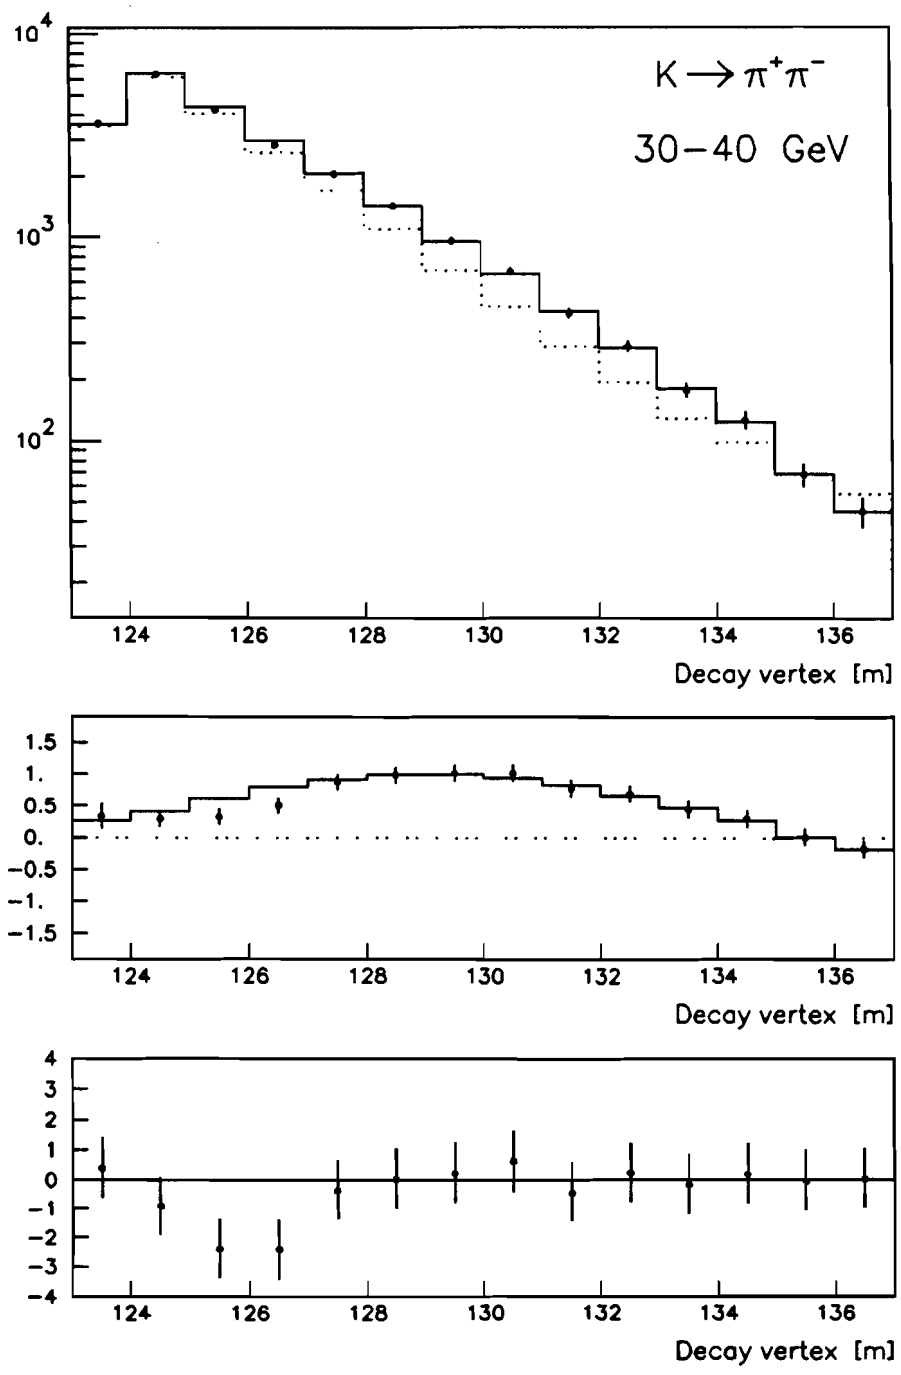

Figure 8.3: *Top:* Decay vertex distributions for $\pi^+\pi^-$ decays (30–40 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

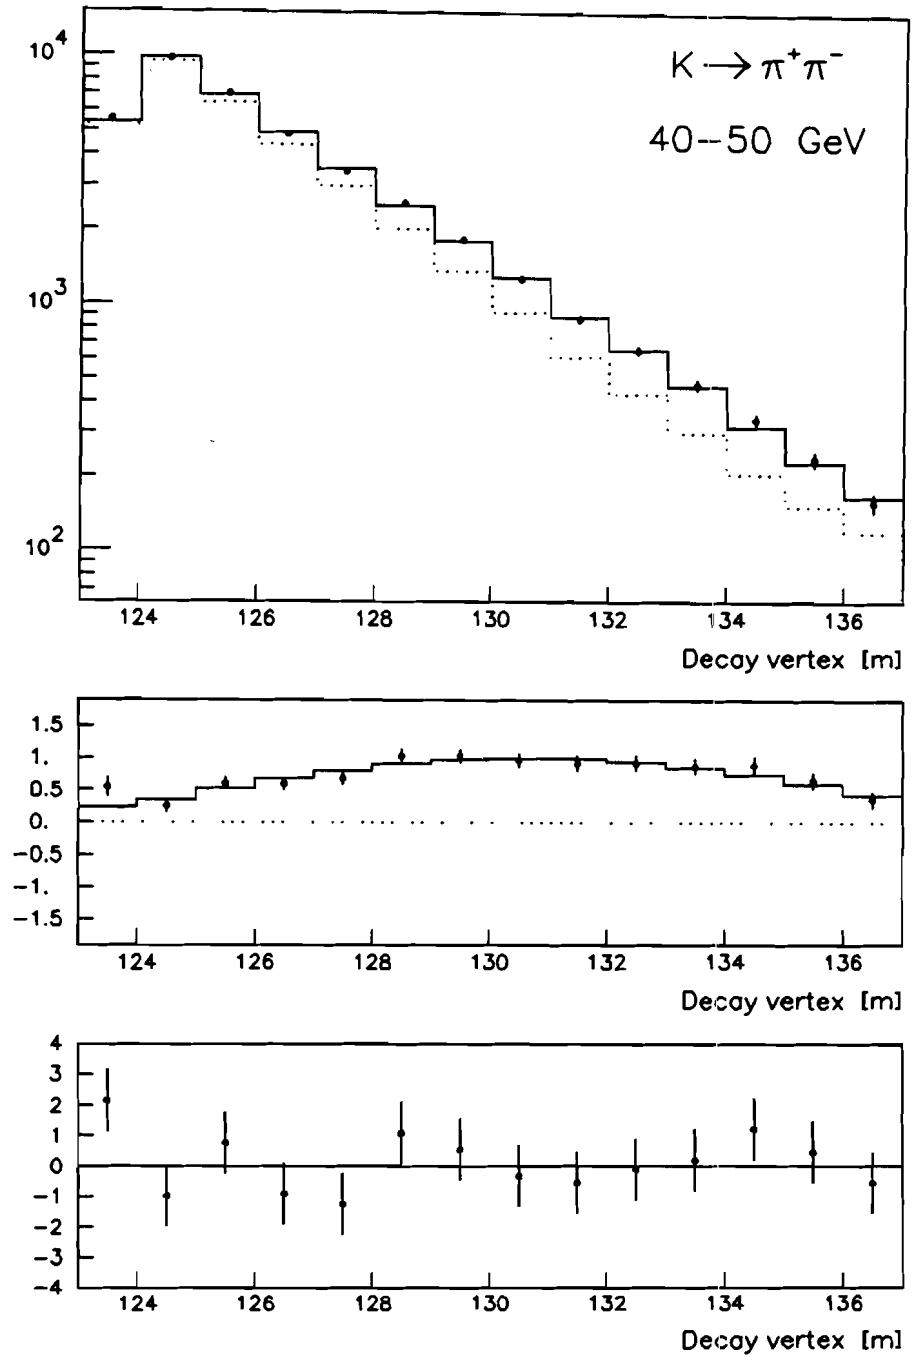

Figure 8.4: *Top:* Decay vertex distributions for $\pi^+\pi^-$ decays (40–50 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

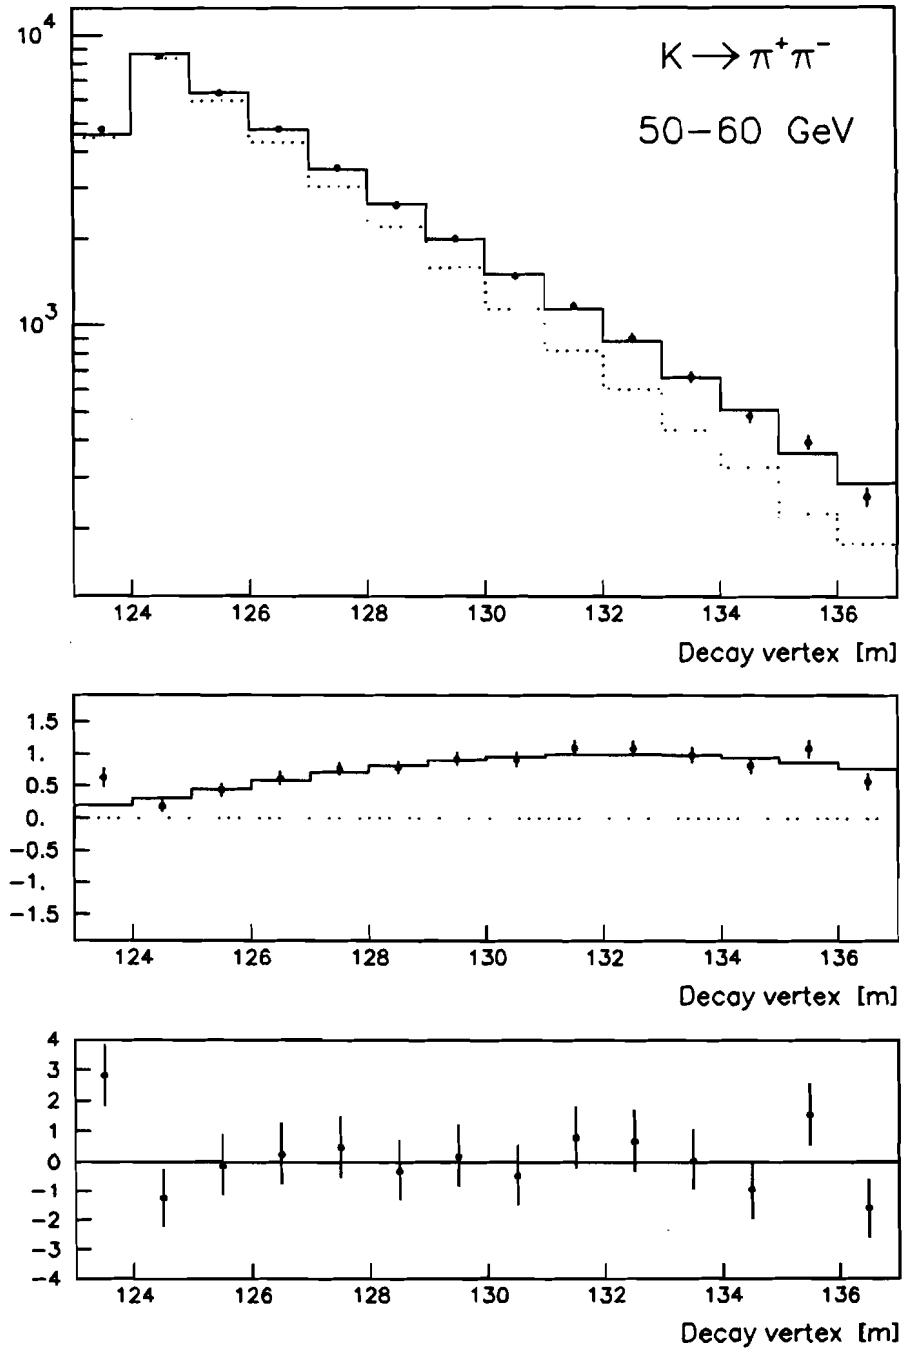

Figure 8.5: *Top:* Decay vertex distributions for $\pi^+\pi^-$ decays (50–60 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

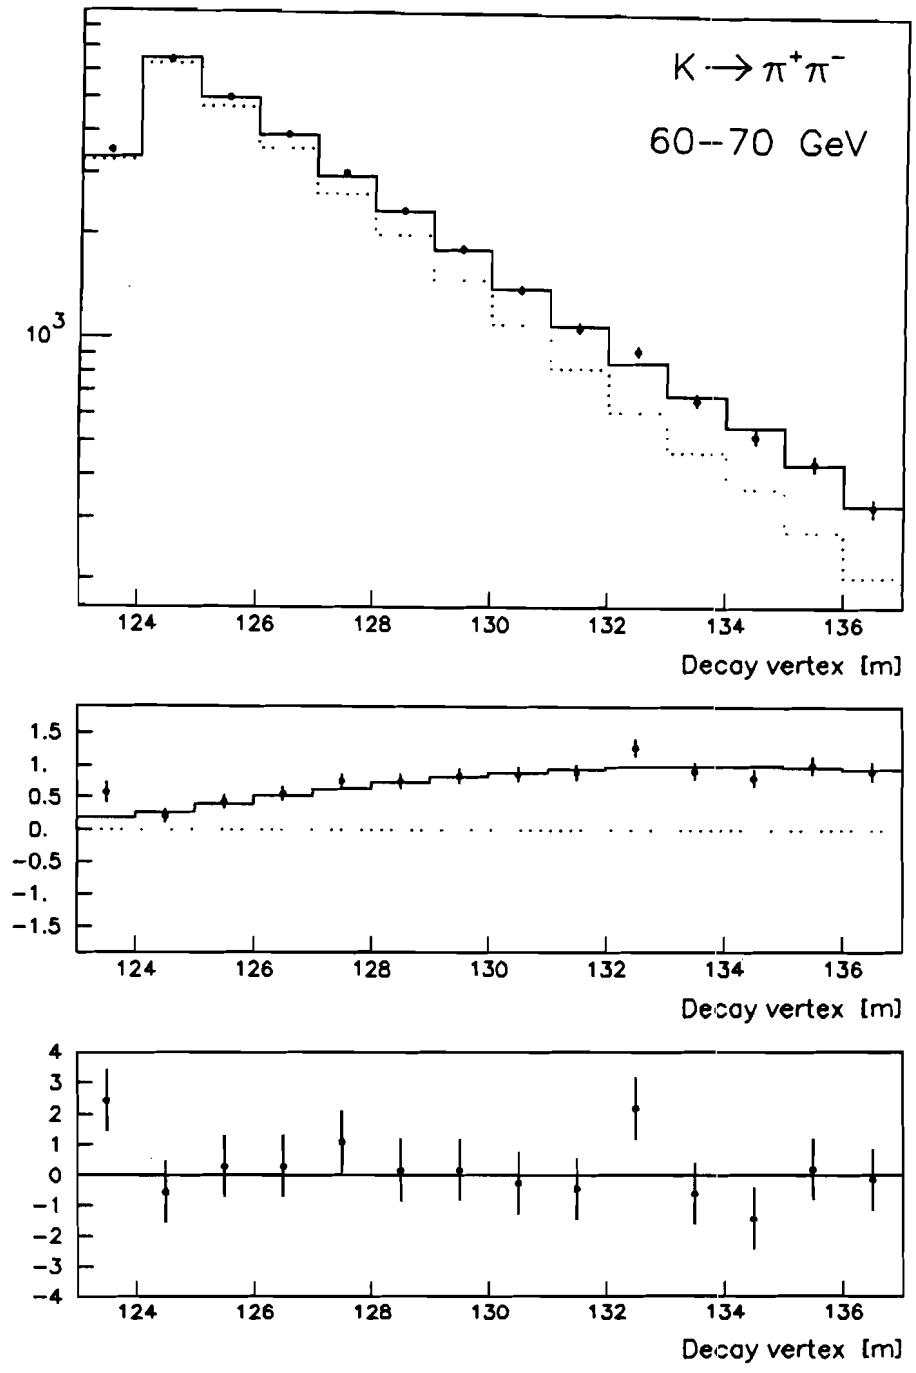

Figure 8.6: *Top:* Decay vertex distributions for $\pi^+\pi^-$ decays (60–70 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

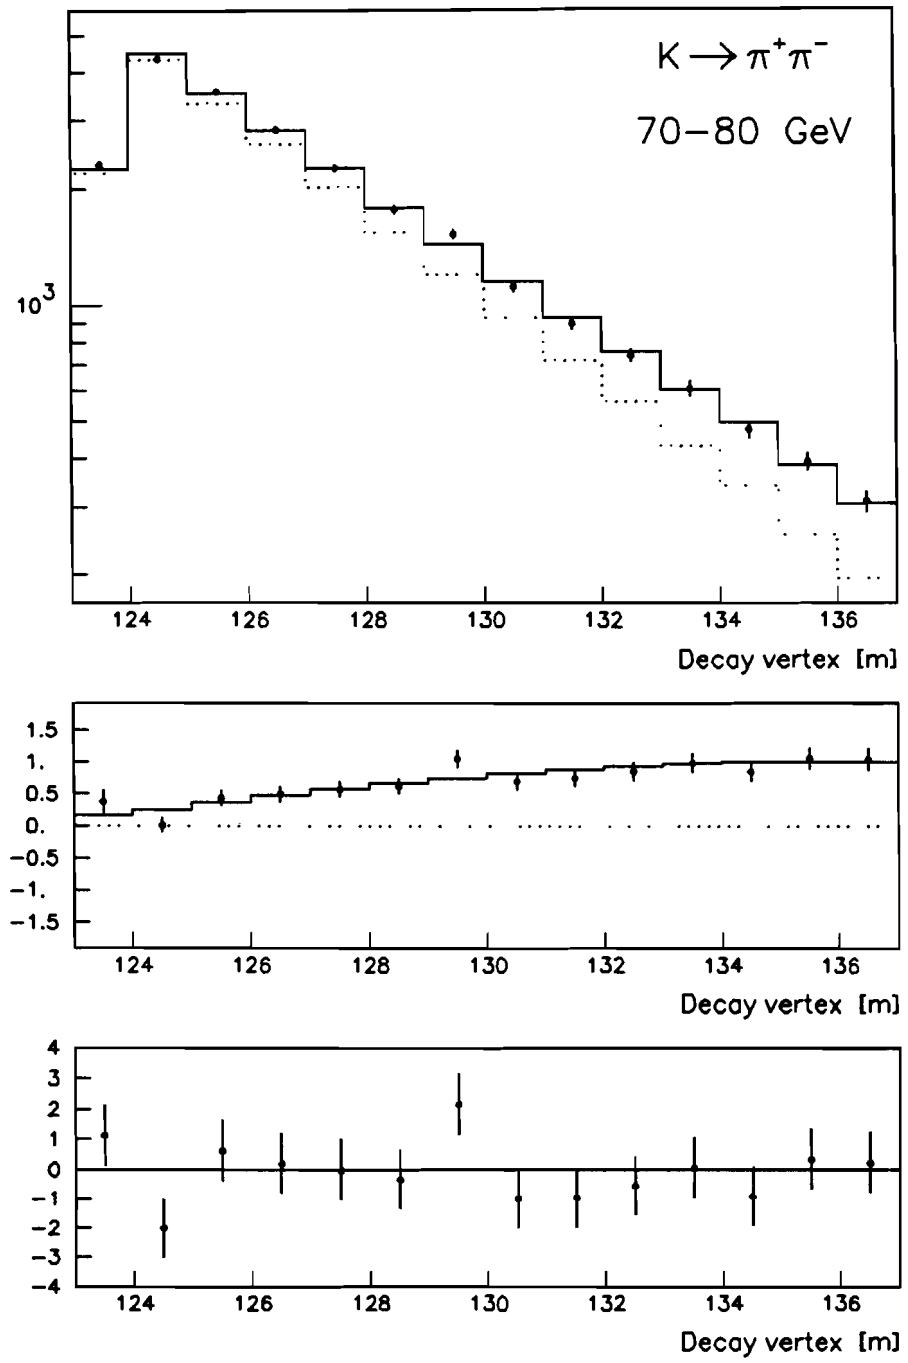

Figure 8.7: *Top:* Decay vertex distributions for $\pi^+\pi^-$ decays (70–80 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

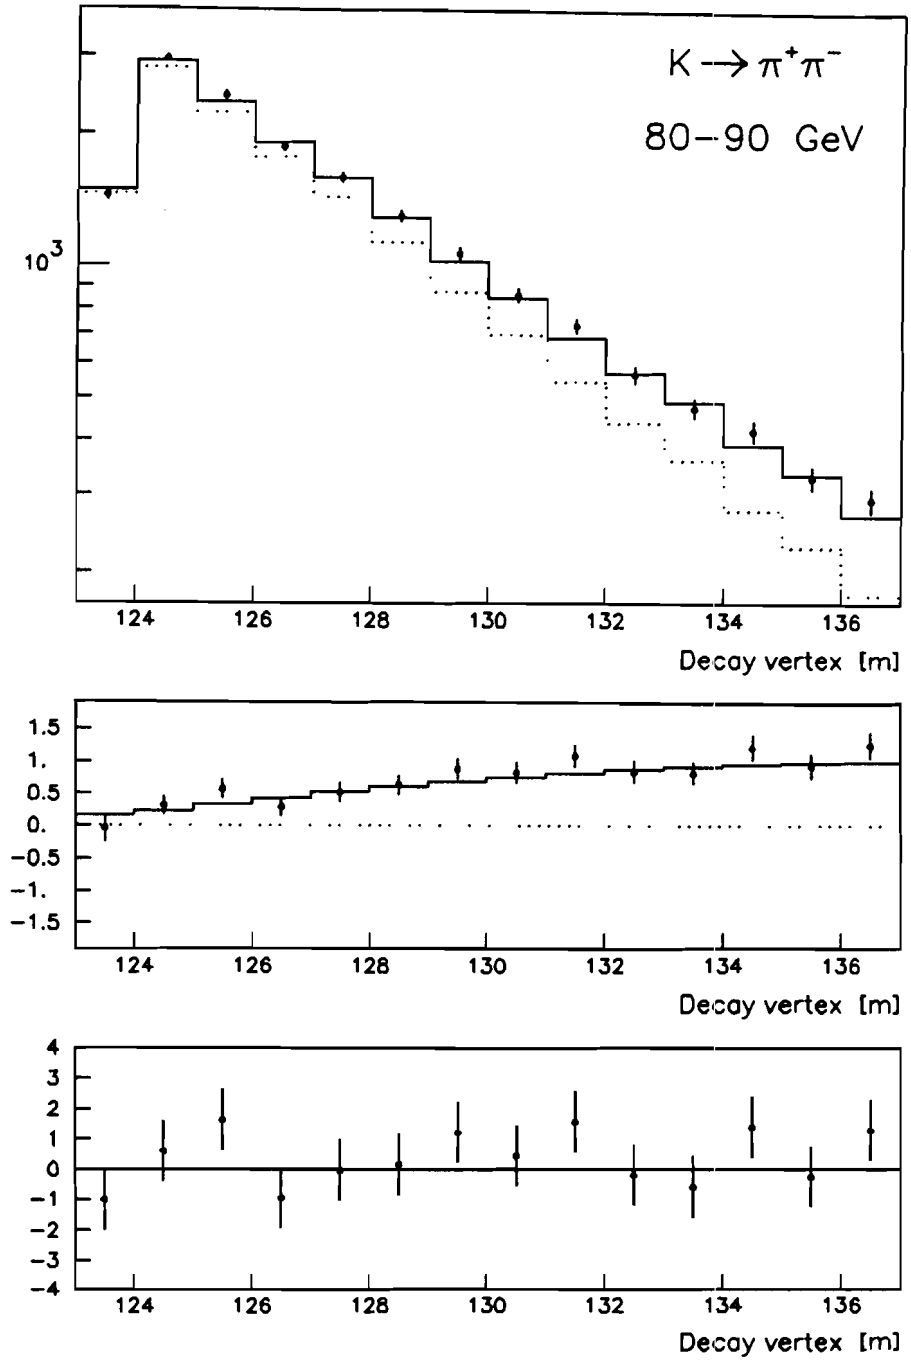

Figure 8.8: *Top:* Decay vertex distributions for $\pi^+\pi^-$ decays (80–90 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

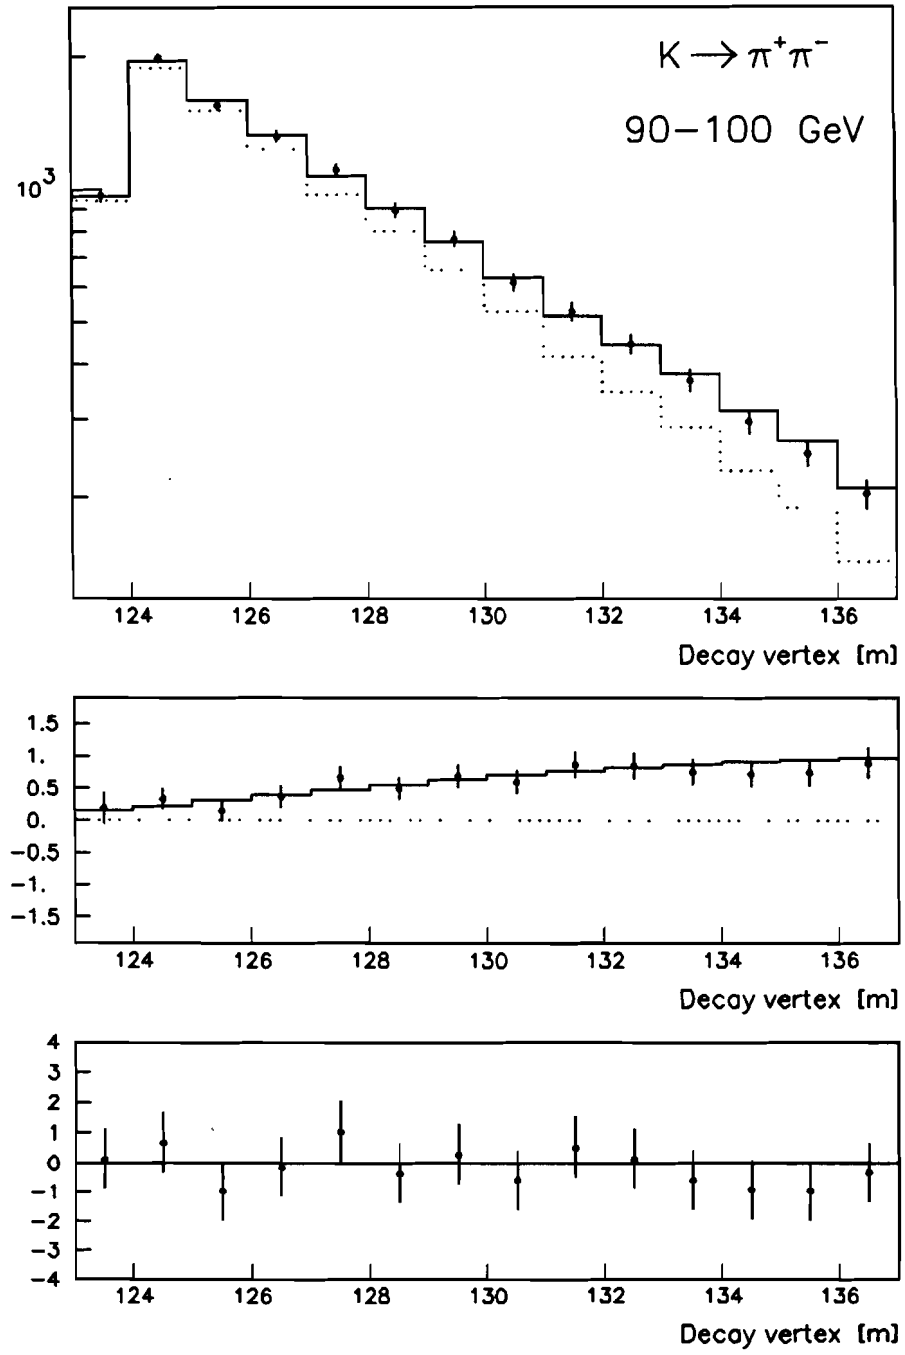

Figure 8.9: *Top:* Decay vertex distributions for $\pi^+\pi^-$ decays (90–100 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

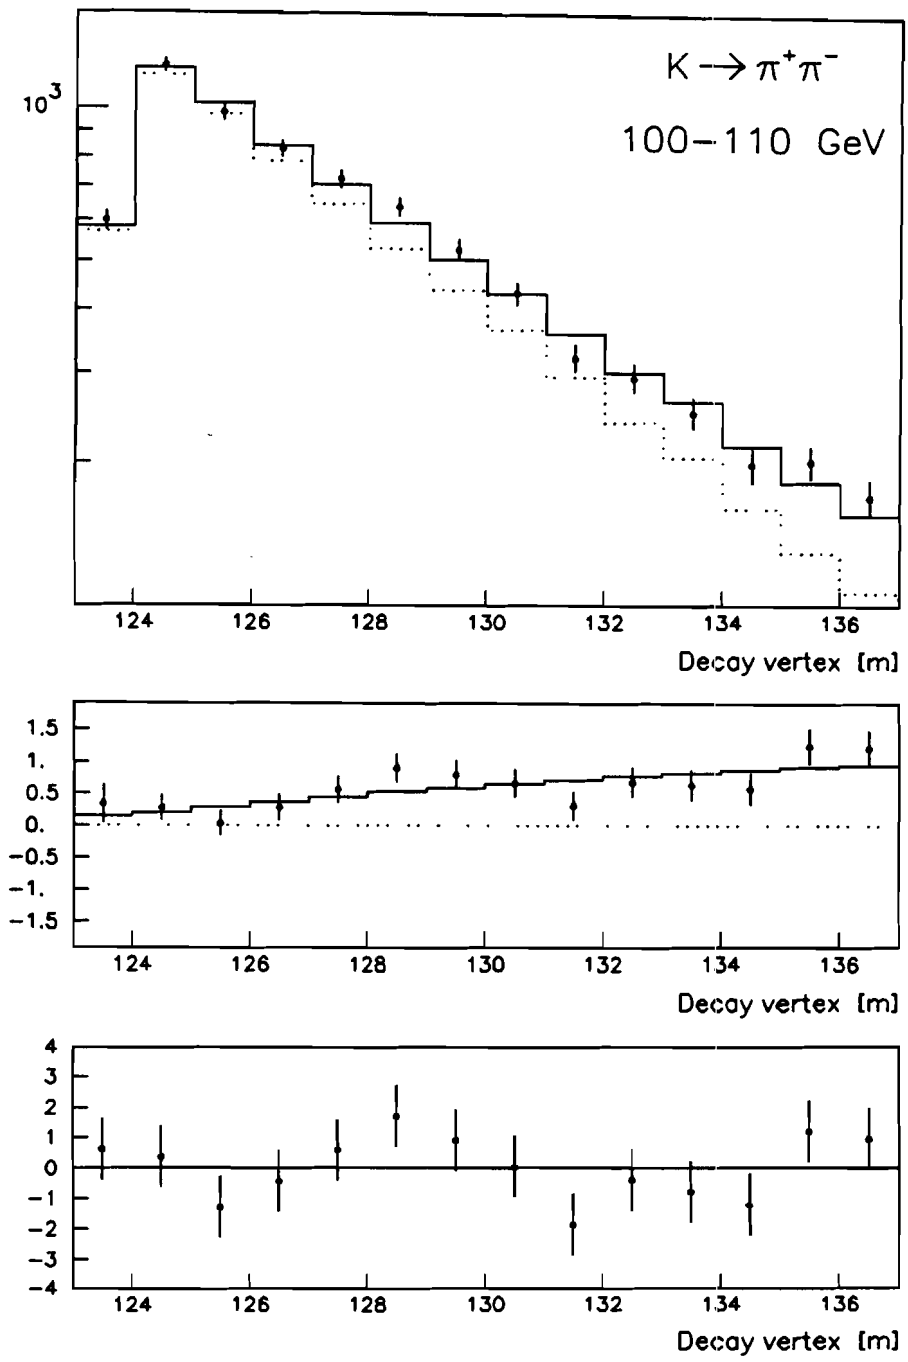

Figure 8.10: *Top:* Decay vertex distributions for $\pi^+\pi^-$ decays (100–110 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

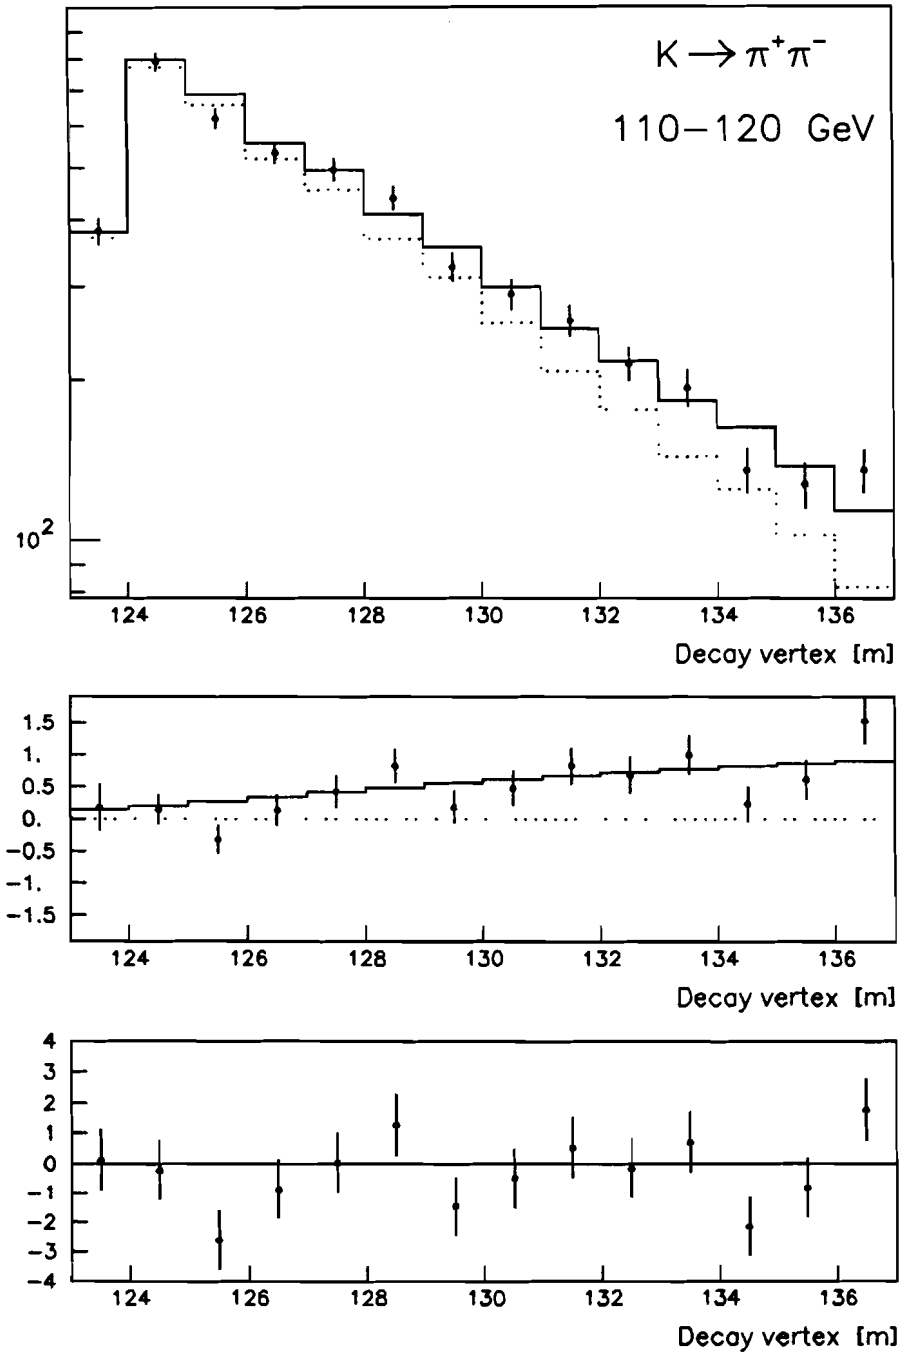

Figure 8.11: *Top:* Decay vertex distributions for $\pi^+\pi^-$ decays (110–120 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

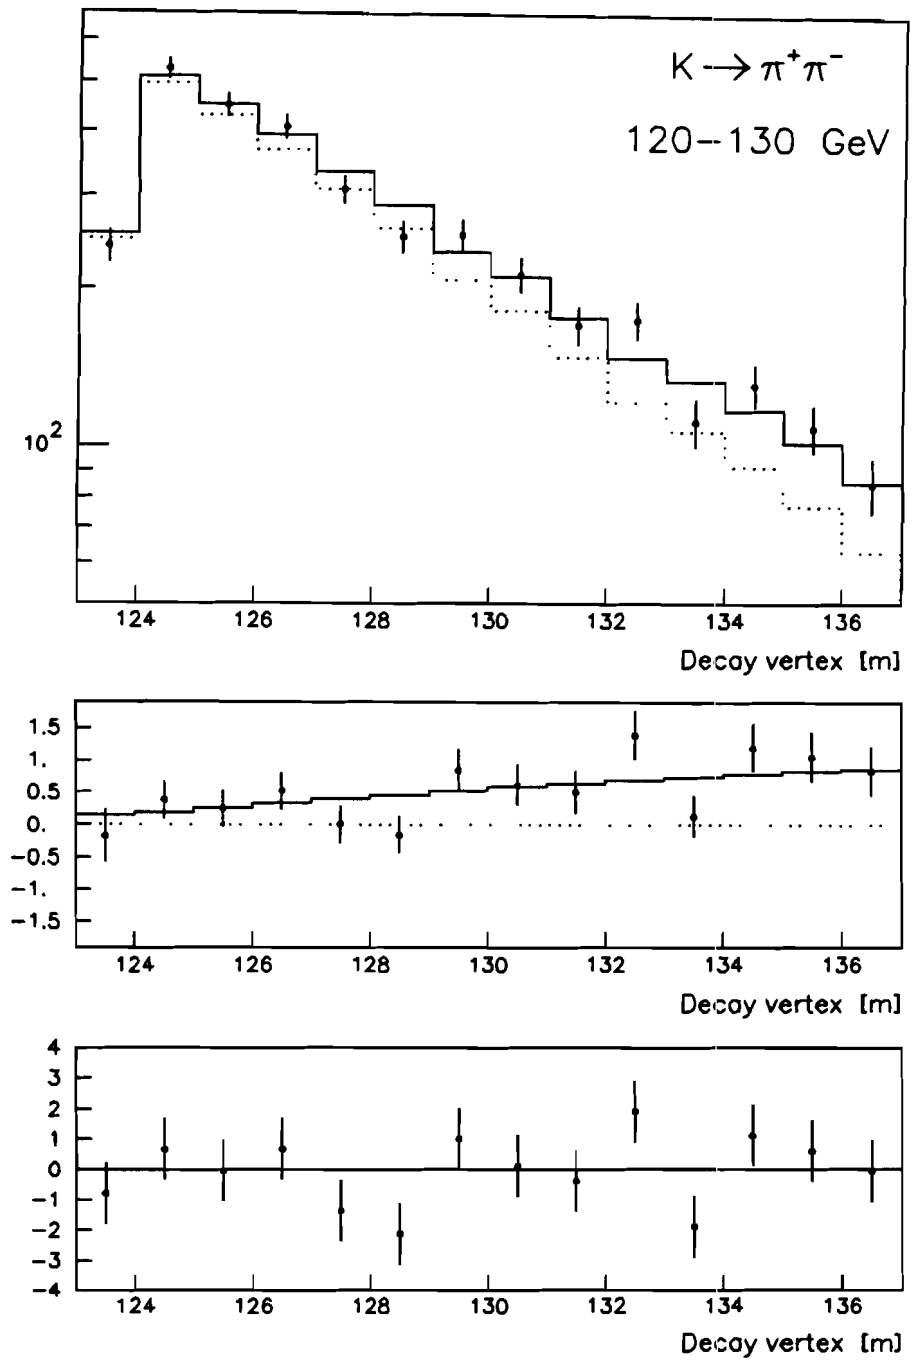

Figure 8.12: *Top:* Decay vertex distributions for $\pi^+\pi^-$ decays (120–130 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

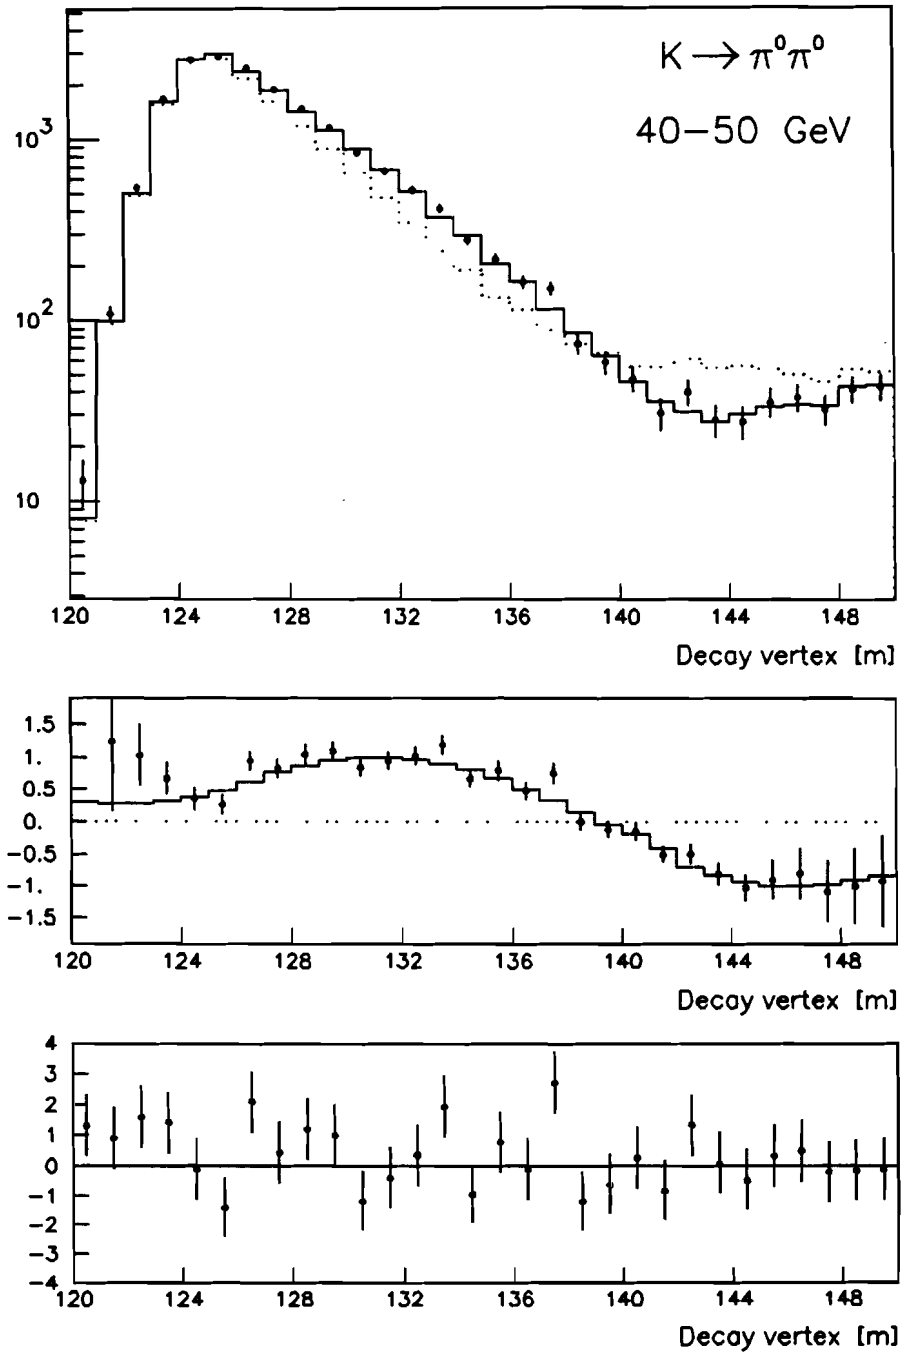

Figure 8.13: *Top:* Decay vertex distributions for $\pi^0\pi^0$ decays (40–50 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

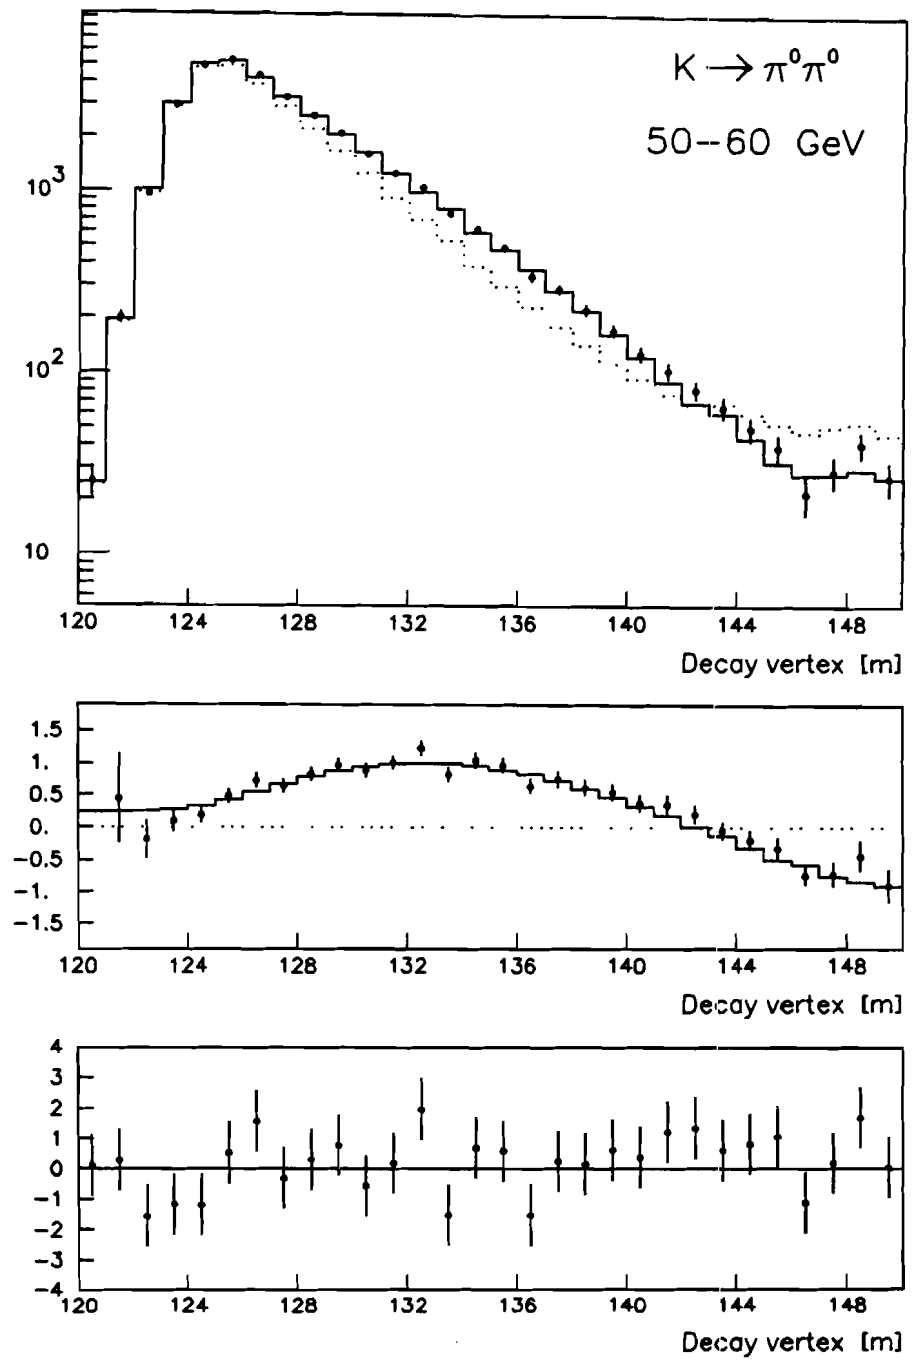

Figure 8.14: *Top:* Decay vertex distributions for $\pi^0\pi^0$ decays (50–60 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

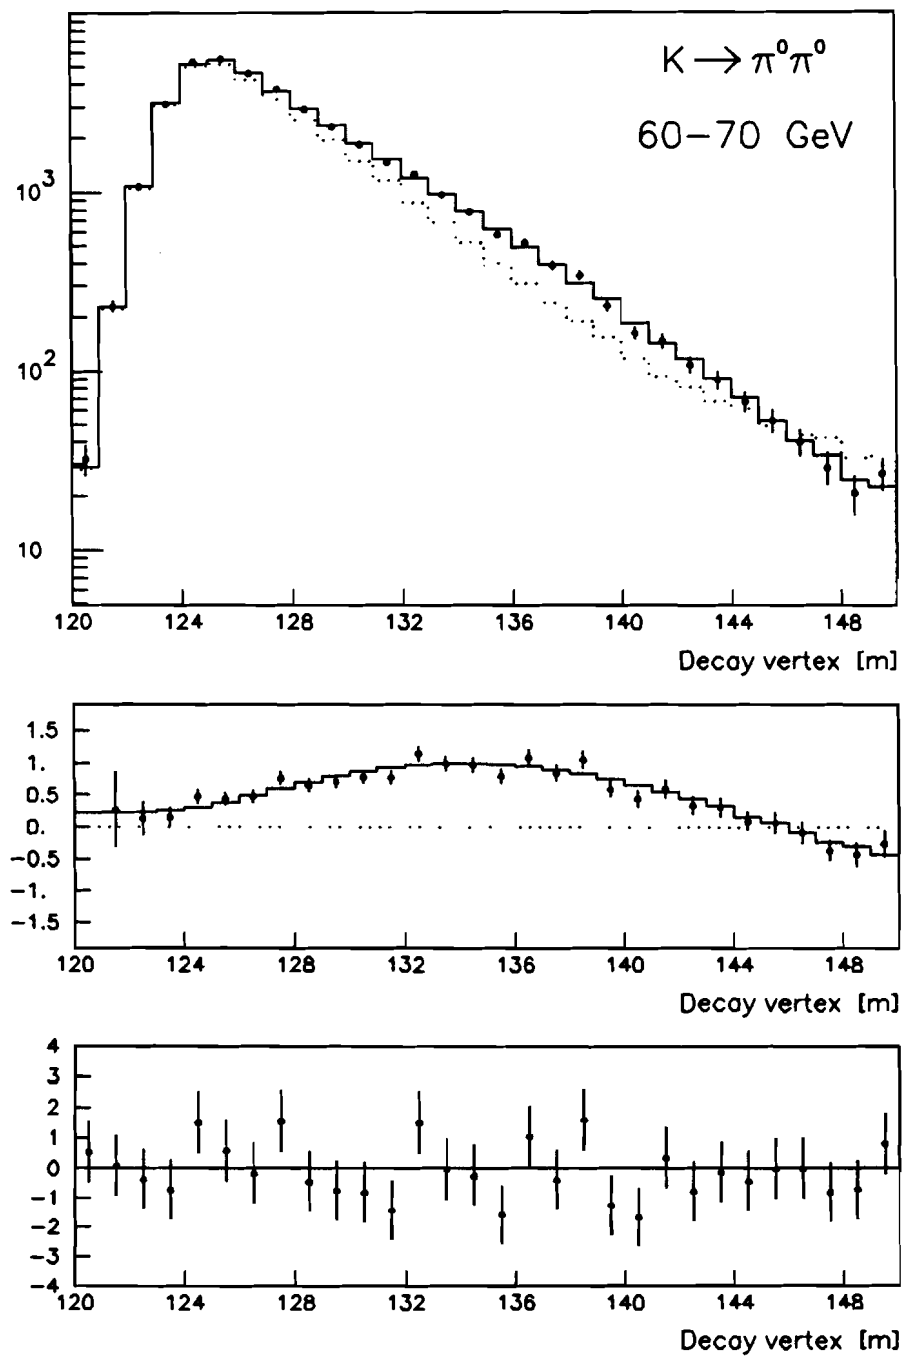

Figure 8.15: *Top:* Decay vertex distributions for $\pi^0\pi^0$ decays (60–70 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

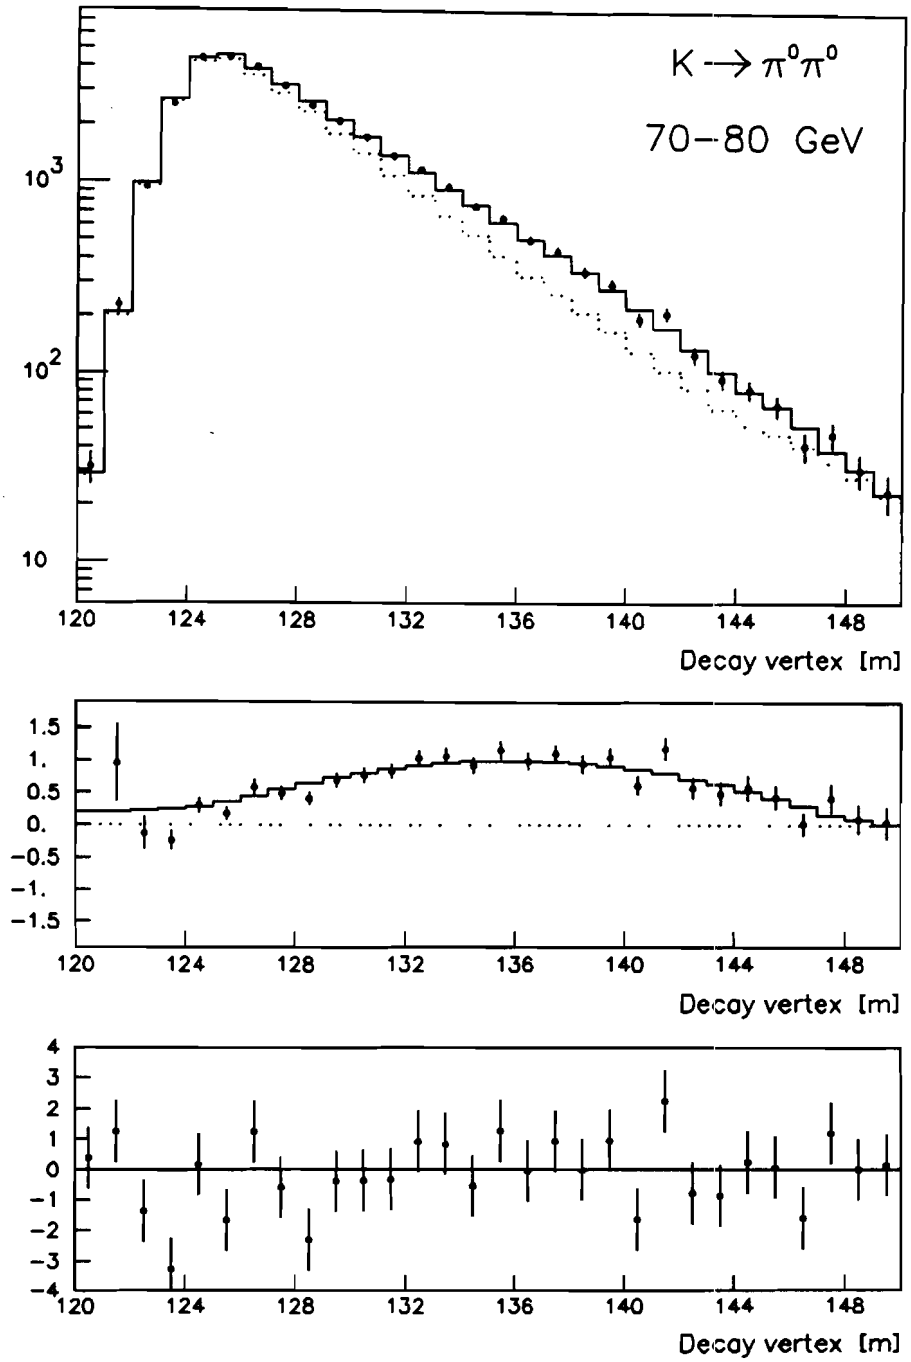

Figure 8.16: *Top:* Decay vertex distributions for $\pi^0\pi^0$ decays (70–80 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

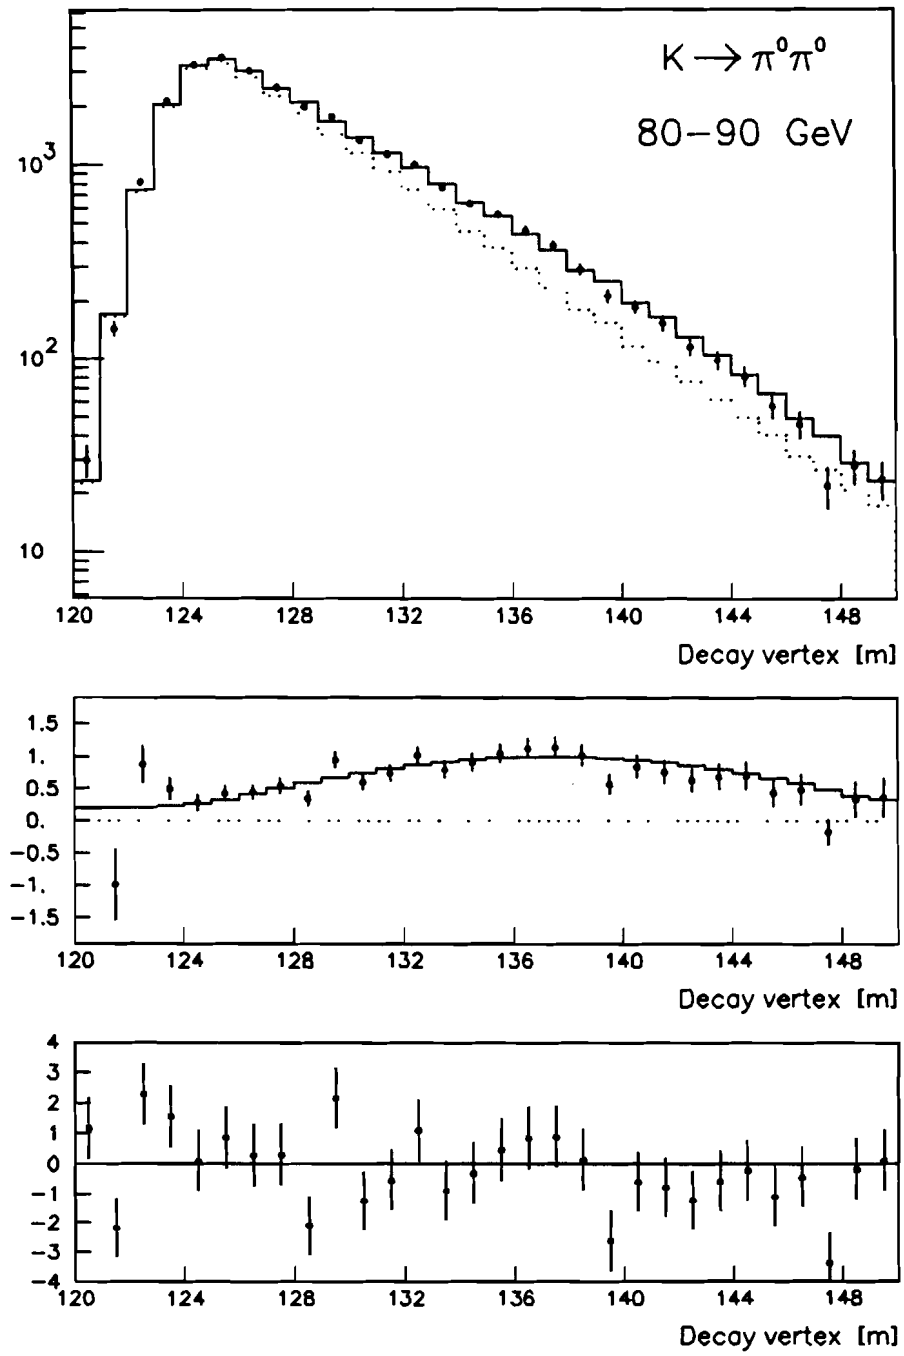

Figure 8.17: *Top:* Decay vertex distributions for $\pi^0\pi^0$ decays (80–90 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

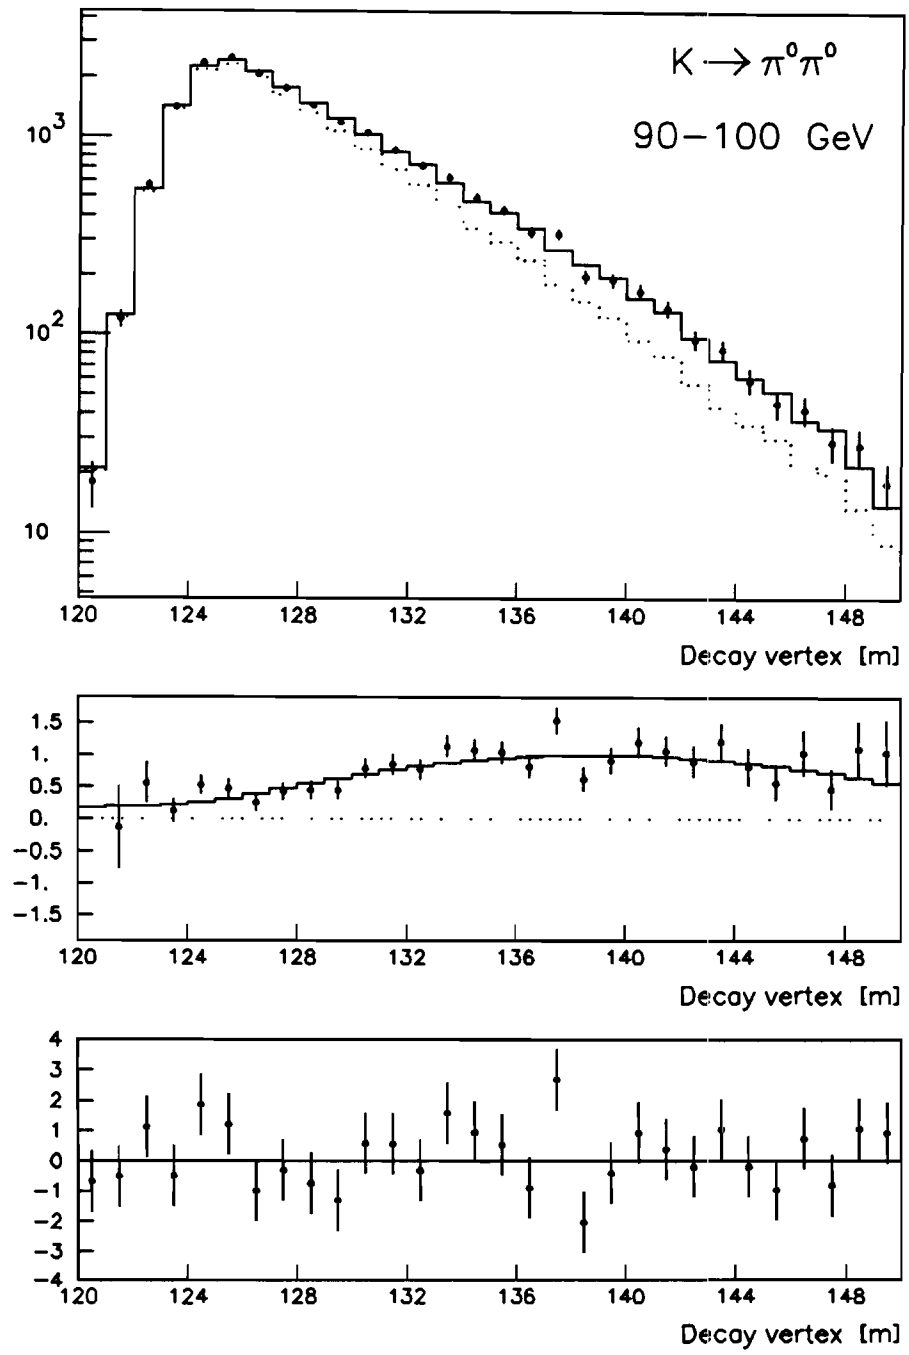

Figure 8.18: *Top:* Decay vertex distributions for $\pi^0\pi^0$ decays (90–100 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

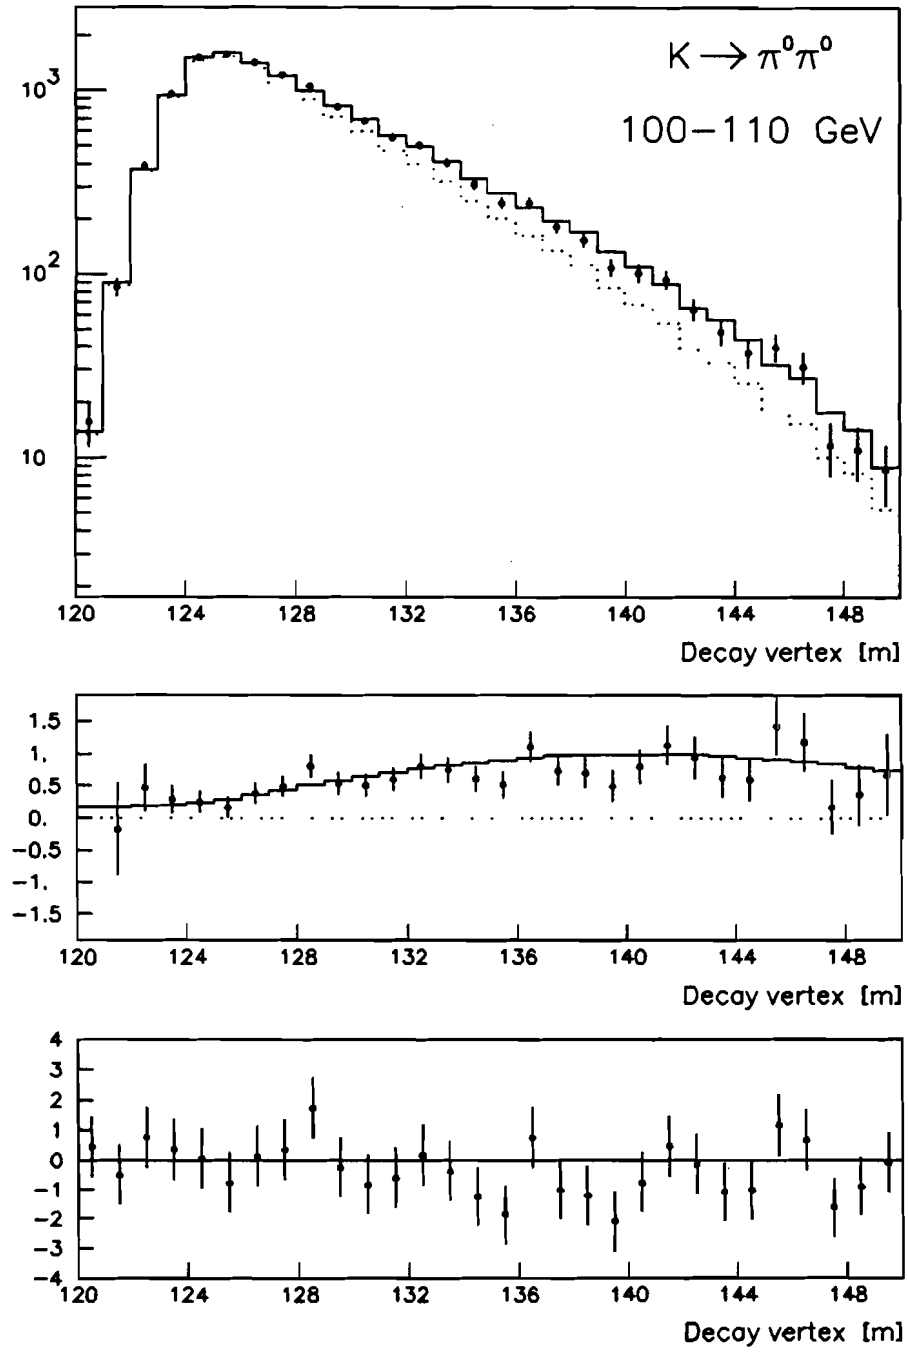

Figure 8.19: *Top:* Decay vertex distributions for $\pi^0\pi^0$ decays (100–110 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

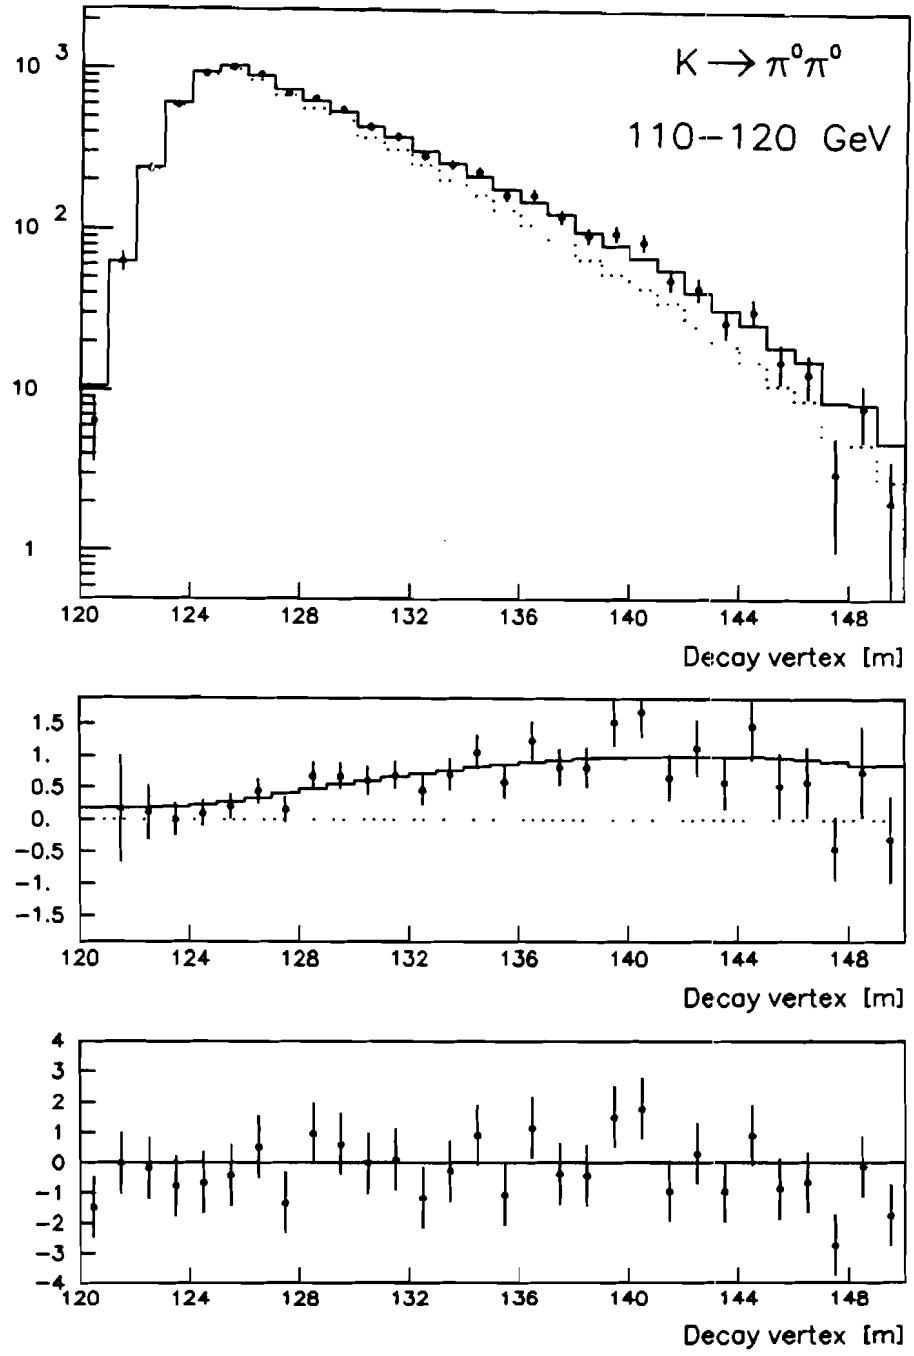

Figure 8.20: *Top:* Decay vertex distributions for $\pi^0\pi^0$ decays (110–120 GeV) in the regenerated beam. Data (solid circles), Monte Carlo (histogram) and Monte Carlo without the interference term (dotted histogram). *Middle:* Extracted "cosine" factor (see the text). *Bottom:* Standard deviations for the data-MC difference.

Chapter 9
CONCLUSIONS

The main result of this analysis is the phase difference $\phi_{00} - \phi_{+-}$, for which several systematic errors essentially cancel. This result[31]

$$\Delta\phi \equiv \phi_{00} - \phi_{+-} = -0.3^\circ \pm 2.4^\circ(\text{stat}) \pm 1.2^\circ(\text{syst})$$

is in good agreement with the prediction of CPT symmetry. Reference 11 has pointed out that (with certain assumptions about isospin) $\Delta\phi$ can be related to the lifetime difference between K^+ and K^-. With that relationship, the $\Delta\phi$ result is comparable in precision to the result from direct lifetime measurements.

The values obtained for the individual phases, $\phi_{+-} = 47.7^\circ \pm 2.0^\circ(\text{stat})$ and $\phi_{00} = 47.4^\circ \pm 1.4^\circ(\text{stat})$, have larger systematic errors. As explained in the previous chapter, this is mostly due to the limited knowledge about the regeneration phase ϕ_ρ. This is somewhat unfortunate, since we would also like to check the CPT prediction for the average of ϕ_{+-} and ϕ_{00}. However, it is possible to calculate a value for the average phase that is independent of the regeneration phase by combining our $\Delta\phi$ measurement with the world average ϕ_{+-} from reference 8, namely $\phi_{+-} = 44.6^\circ \pm 1.2^\circ$. The result

$$2\phi_{+-}/3 + \phi_{00}/3 = 44.5^\circ \pm 1.5^\circ$$

is within one standard deviation of the CPT prediction 43.7°. From this result, one finds the perpendicular component of ϵ_0. (Recall that $\epsilon_0 \simeq 2\eta_{+-}/3 + \eta_{00}/3$.) With $|\epsilon_0| = 2.259 \times 10^{-3}$ (from $|\eta_{+-}|$ and $|\eta_{00}|$ in reference 8), the value calculated above for the average phase leads to

$$\epsilon_{0\perp} = (3.2 \pm 5.9) \times 10^{-5}$$

It is also possible to derive a limit for the K^0-$\overline{K^0}$ mass difference. From chapter 2, we recall the relationship $\Delta_\perp = -\epsilon_{0\perp} - \mathrm{Im}\,\tilde{\alpha}\cos\Phi - r_0 \sin\Phi$, where $\Delta_\perp = \delta m/2\sqrt{2}\,\Delta m$, and Φ is the "natural angle" 43.7°. As usual, Δm is the K_L-K_S mass difference, and here we let δm be the K^0-$\overline{K^0}$ mass difference. (Please refer to chapter 2 for details about the algebra.) With experimental limits from the 3π and semileptonic decay modes, reference 11 has obtained $\mathrm{Im}\,\tilde{\alpha} = (-2.6 \pm 12.0) \times 10^{-5}$. If one assumes that the $\Delta S = \Delta Q$ rule is exact for the semileptonic decay modes, and that the CP violation in the 3π and 2π modes are comparable (*i.e.*, $|\eta_{3\pi}| \simeq |\eta_{2\pi}|$), the limit on $\mathrm{Im}\,\tilde{\alpha}$ is greatly improved and becomes negligible compared to the $\epsilon_{0\perp}$ error. The parameter r_0 represents CPT violation in the 2π decay amplitudes, and it is related to observable parameters by (assuming $\Delta S = \Delta Q$ once again) $r_0 = \mathrm{Re}\,\epsilon_0 - \delta/2 - \mathrm{Re}\,y$, where δ is the semileptonic charge asymmetry (averaged over $\pi e \nu$ and $\pi \mu \nu$) and y represents CPT violation in the semileptonic decay amplitudes. With $|\epsilon_0| = (2.259 \pm 0.018) \times 10^{-3}$ and $\delta = (3.30 \pm 0.12) \times 10^{-3}$ (reference 8) and $\epsilon_{0\perp}$ from above, we obtain $r_0 + \mathrm{Re}\,y = (-8 \pm 8) \times 10^{-5}$. The experimental limit on $\mathrm{Re}\,y$ is not very good — reference 12 estimates $\pm 6 \times 10^{-3}$ — but we will be bold and assume that there is no freak cancellation between r_0 and $\mathrm{Re}\,y$ and pretend that the combined limit can be used for r_0 separately. The result is

$$\Delta_\perp = (2 \pm 8) \times 10^{-5}$$

With $\Delta m/m_K = 7 \times 10^{-15}$, this becomes $\delta m/m_K = (0.4 \pm 1.6) \times 10^{-18}$, which is equivalent to

$$\left|\frac{\delta m}{m_K}\right| < 3 \times 10^{-18} \quad (90\%\ \mathrm{C.L.})$$

A recent experiment by the NA31 collaboration at CERN also finds that $\Delta\phi$ is consistent with zero. They use K_S from the target (no regenerator) to measure the interference, and find[32] that $\phi_{+-} = 46.9° \pm 2.2°$ and $\phi_{00} = 47.1° \pm 2.8°$.

Since both of these experiments have found the average of ϕ_{+-} and ϕ_{00} to be a little above the predicted 43.7°, more precise measurements would be of great interest. Improved measurements of the parameters of the 3π and semileptonic decay modes would also be very interesting, since the limits on these parameters would affect the interpretation of a non-zero $\epsilon_{0\perp}$. (See the discussion above.) Of

the two phases, ϕ_{+-} will probably be the easier one to measure more precisely. However, if the ϕ_{+-} error is reduced significantly, one would also need more precise measurements of $\Delta\phi$ to find the average phase. ($\Delta\phi$ is also of great value by itself, since it is related to CPT violation in the 2π decay amplitudes.) One such experiment (E773 at Fermilab) will take data later this year (1990), and they expect to determine $\Delta\phi$ to a precision better than one degree.

Bibliography

[1] C.S. Wu *et al.*, Phys. Rev. **105**, 1413 (1957).

[2] J.H. Christenson *et al.*, Phys. Rev. Lett. **13**, 138 (1964).

[3] M. Kobayashi and K. Maskawa, Prog. Theor. Phys. **49**, 652 (1973).

[4] L. Wolfenstein, Phys. Rev. Lett. **13**, 569 (1964).

[5] G. Lüders, Dan. Mat. Fys. Medd. **28**, no. 5 (Det Kongelige Danske Videnskabernes Selskab, 1954).

[6] W. Pauli in *Niels Bohr and the Development of Physics*, p 30, edited by W. Pauli (Pergamon Press, London, 1955).

[7] R. Jost, Helv. Phys. Acta **30**, 409 (1957).

[8] Particle Data Group, G.P. Yost *et al.*, Phys. Lett. B **204**, 1 (1988).

[9] K. Kleinknecht, Ann. Rev. Nucl. Sci. **26**, 1 (1976).

[10] J.W. Cronin, Acta Phys. Pol. B **15**, 419, 721 (1984).

[11] V.V. Barmin *et al.*, Nucl. Phys. B **247**, 293 (1984).

[12] N.W. Tanner and R.H. Dalitz, Ann. Phys. **171**, 463 (1986).

[13] L. Wolfenstein, Ann. Rev. Nucl. Sci. **36**, 137 (1986).

[14] R.G. Sachs, *The Physics of Time Reversal* (University of Chicago Press, Chicago, 1987).

[15] J.R. Patterson *et al.*, Phys. Rev. Lett. **64**, 1491 (1990).

[16] H. Burkhardt *et al.*, Phys. Lett. B **206**, 169 (1988).

[17] T.J. Devlin and J.O. Dickey, Rev. Mod. Phys. **51**, 237 (1979).

[18] T.D. Lee, R. Oehme and C.N. Yang, Phys. Rev. **106**, 340 (1957).

[19] T.T. Wu and C.N. Yang, Phys. Rev. Lett. **13**, 380 (1964).

[20] T.D. Lee and C.S. Wu, Ann. Rev. Nucl. Sci. **16**, 511 (1966).

[21] V. Weisskopf and E.P. Wigner, Z. Phys. **63**, 54 (1930).

[22] J.S. Bell and J. Steinberger, Proceedings of the Oxford International Conference on Elementary Particles, 1965, p 195, edited by T.R. Walsh, A.E. Taylor, R.G. Moorhouse, B. Southworth (Rutherford Laboratory, Chilton, England, 1966).

[23] A. Pais and O. Piccioni, Phys. Rev. **100**, 1487 (1955).

[24] M. Good, Phys. Rev. **106**, 591 (1957).

[25] F. Gilman, Phys. Rev. **171**, 1453 (1968).

[26] J. Roehrig *et al.*, Phys. Rev. Lett. **38**, 1116 (1977).

[27] G. Grazer, Ph.D. thesis, Princeton University, 1988.

[28] M. Woods, Ph.D. thesis, University of Chicago, 1988.

[29] J.R. Patterson, Ph.D. thesis, University of Chicago, 1990.

[30] V. Papadimitriou, Ph.D. thesis, University of Chicago, 1990.

[31] M. Karlsson *et al.*, Phys. Rev. Lett. **64**, 2976 (1990).

[32] R. Carosi *et al.*, Phys. Lett. B **237**, 303 (1990).