

KAON IDENTIFICATION AND THE SEARCH FOR HEAVY NEUTRINOS AT NA62

Francis Newson

*A thesis submitted to
the University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY*



Supervisors:
Prof. Cristina Lazzeroni
Dr. Evgueni Goudzovski

Particle Physics Group,
School of Physics and Astronomy,
University of Birmingham
September 2015

Abstract

A search for heavy neutrinos in the decay $K^+ \rightarrow \mu^+ \nu_h$ has been performed using the data collected by the NA62- R_K experiment in 2007. The signal event selection is presented, along with an estimate of the expected background in the signal region, obtained by Monte Carlo simulation as well as data driven techniques. This estimate is supported by studies of the detector performance in muon identification efficiency and momentum resolution. Using a sample corresponding to 6×10^7 kaon decays in the fiducial volume, limits of order 10^{-5} have been set on the branching ratio, $\mathcal{B}(K^+ \rightarrow \mu^+ \nu_h)$, for neutrino masses in the range $275 \text{ MeV}/c^2 - 375 \text{ MeV}/c^2$, under the assumption that the heavy neutrino does not decay into visible particles inside the detector. In the range $330 \text{ MeV}/c^2 - 375 \text{ MeV}/c^2$, there are no previous limits from peak searches. Over the rest of the range, the new limits do not improve upon existing results.

Improved limits might come from the NA62 experiment designed to measure the branching ratio of the ultra rare kaon decay $K^+ \rightarrow \pi^+ \bar{\nu} \nu$, which began taking data in 2015. Kaon identification is performed by a differential Cherenkov detector. The timing resolution of the photo-detectors was measured in a pilot run in 2014 and used to improve the simulation of the detector. An algorithm was developed to align the detector during live data taking, and this was used at the start of the 2015 run.

Declaration of author's contribution

The search for heavy neutrinos presented in Chapter 3 is the result of my own work. The data sample analysed was collect by the NA62- R_K experiment in 2007. I was not involved in the design, construction or operation of the detector used in 2007 and the description provided in Chapter 2 is only included to support the heavy neutrino analysis.

I was not involved in the design of the NA62 detector described in Chapter 4 but I took part in the installation of the KTAG detector described in Chapter 5. The simulation studies in Chapter 5 and alignment studies presented in Chapter 6 are my own work except where explicitly stated.

No other qualifications have been achieved with the work reported in this thesis.

Acknowledgements

I would like to thank all the people who have helped me over the last four years.

I have learnt a huge amount from my supervisors, Cristina Lazzeroni and Evgueni Goudzovski, about physics, research, collaborations, teaching, speaking, writing and more. I am very grateful for all their support and advice when work was going smoothly, and when it wasn't.

I thank Tonino Sergi, Angela Romano, Karim Massri and Chris Parkinson for their patient and friendly guidance about every aspect of NA62, physics, computers and the rest. Thanks also to my fellow students Jimmy, Nathaniel, Andy, Tim and Nicolas, for helping/putting up with/encouraging/distracting me as required.

As always, I am grateful to my family for their love, support and irrepressible confidence in me and for proofreading several chapters of my thesis. Finally, I thank my wife Vicky for her unfailing love and support and her patience during the last eight years of physics.

Contents

Introduction	1
1. Theory	3
1.1. Mass in the Standard Model	3
1.2. The CKM framework	9
1.2.1. Parameter count	10
1.2.2. CP Violation	10
1.2.3. Parametrization	12
1.2.4. Unitarity triangles	13
1.3. $K^+ \rightarrow \pi^+ \nu \bar{\nu}$	15
1.3.1. Standard Model prediction	16
1.3.2. Impact on the unitarity triangle	18
1.4. Neutrinos in the Standard Model	19
1.4.1. Massless Origins	19
1.4.2. Neutrino Oscillations	21
1.5. Theories of neutrino mass	22
1.5.1. Dirac Mass	23
1.5.2. See-saw mechanisms	24
1.5.3. ν MSM	28
1.6. Experimental searches for heavy neutrinos	28
1.6.1. Peak searches in kaon decay	30
1.7. Heavy neutrino kinematics	31
1.7.1. Production	31
1.7.2. Decay	34
2. The NA62 R_K detector	36
2.1. Fixed target kaon experiments at CERN	36

2.2. NA62- R_K : Lepton Universality test in K^\pm decays	38
2.3. The beam and detector	40
2.3.1. The charged kaon beam line	40
2.3.2. The decay volume	41
2.3.3. The magnetic spectrometer	42
2.3.4. The hodoscope	43
2.3.5. The liquid krypton electromagnetic calorimeter	44
2.3.6. The neutral hodoscope	45
2.3.7. The hadronic calorimeter	46
2.3.8. The muon veto system	46
2.3.9. Anti counters	47
2.3.10. Trigger	48
2.3.11. Data acquisition	49
3. Search for heavy neutrinos at NA62-R_K	50
3.1. Analysis strategy	50
3.2. Selection of data taking periods	52
3.3. Pre-selection	53
3.3.1. Single Track Pre-selection	53
3.3.2. Signal shapes	54
3.3.3. Background composition	54
3.3.4. Background suppression	59
3.4. Simulation and corrections	60
3.4.1. Physics simulation	60
3.4.2. Beam	60
3.4.3. Liquid Krypton Calorimeter	62
3.4.4. Spectrometer corrections	64
3.5. Spectrometer resolution tails	64
3.5.1. LKr selection	66
3.5.2. Kinematic fit	68
3.5.3. Charged track selection	71
3.5.4. Kinematic distributions	72
3.5.5. Spectrometer resolution correction	73
3.5.6. Scattering parameters	76

3.6. Muon veto efficiency	79
3.6.1. Muon reconstruction	79
3.6.2. Simulation of scattering applied to existing MC samples	81
3.6.3. Muon sample for MUV efficiency measurement	83
3.6.4. MUV Bad bursts	85
3.6.5. MUV Efficiency measurement	85
3.7. Beam Halo	88
3.7.1. Halo Simulation	90
3.7.2. Data driven approach	91
3.7.3. Halo Pre-selection	92
3.7.4. Upper region	94
3.7.5. Lower region	97
3.7.6. Halo missing mass distribution	99
3.7.7. Punch-through	99
3.7.8. $K^- \rightarrow \pi^- \pi^- \pi^+$ cross-talk	102
3.7.9. Scale factor between K^+ and K^- data	103
3.8. The final signal selection	109
3.8.1. Track requirements	109
3.8.2. Photon veto	109
3.8.3. Muon Identification	110
3.8.4. Kinematic cuts	110
3.8.5. Signal and background in the final selection	110
3.9. Statistical framework	111
3.10. Systematic Uncertainties	114
3.10.1. Spectrometer tails	114
3.10.2. Total number of kaon decays	114
3.10.3. Muon veto	116
3.10.4. Trigger efficiency	119
3.10.5. Halo	121
3.11. Limits on the branching ratio to heavy neutrinos	125
3.11.1. Signal window width	125
3.11.2. Final sensitivity	126
3.11.3. Results	128

3.12. Conclusion	132
4. $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ at NA62	133
4.1. Experimental requirements	133
4.2. Previous experiments	134
4.3. NA62 Analysis Strategy	136
4.4. The NA62 Detector	139
4.4.1. Beam	141
4.4.2. Tracking and momentum measurement systems	142
4.4.2.1. The GigaTracker (GTK) spectrometer	142
4.4.2.2. The Charged ANTI (CHANTI) detector	143
4.4.2.3. The Straw Tracker	144
4.4.3. Particle identification	145
4.4.3.1. Muon Veto System	145
4.4.3.2. Ring Imaging Cherenkov Detector (RICH)	146
4.4.3.3. The charged hodoscope	147
4.4.4. Photon vetoes	147
4.4.4.1. Large Angle Vetoes	148
4.4.4.2. The LKr Calorimeter	149
4.4.4.3. Small Angle Vetoes	149
4.4.5. Trigger and data acquisition	150
5. KTAG at NA62	152
5.1. Principle of operation	153
5.1.1. Ring broadening	154
5.1.2. The number of photons per kaon	156
5.2. The CEDAR detector	158
5.3. The KTAG upgrade	159
5.4. Photo-multiplier tubes	161
5.5. Front-end and read-out system	161
5.6. Simulation	162
5.6.1. Cherenkov light propagation	163
5.6.2. Lens and mirror geometry	166
5.6.3. PMT Time Response	167

5.7. Conclusion	170
6. KTAG Alignment	172
6.1. Introduction	172
6.1.1. Alignment mechanism	172
6.1.2. Alignment parameters	173
6.1.3. Required alignment precision	174
6.1.4. Alignment procedures	174
6.2. Quadrant Asymmetries	176
6.2.1. Simplified models	177
6.2.2. Complete NA62MC simulation	181
6.2.3. Calibration of asymmetry vs misalignment	182
6.2.4. Performance of the asymmetry method	185
6.3. Monte Carlo Templates	187
6.3.1. χ^2 fits	187
6.3.2. Event selection and PMT grouping	188
6.3.3. PMT variation	189
6.3.4. Pressure variation	191
6.4. 2014 Data	194
6.4.1. χ^2 alignment test data sets	194
6.4.2. Consistency checks	195
6.4.3. Comparison of χ^2 and asymmetry precision	196
6.4.4. Fine tuning	200
6.5. Conclusions for the Online Monitor	201
7. Conclusions	204
Bibliography	206
A. Notation	215
B. PMT Properties	218

List of Figures

1.1. Experimental constraints on the unitarity triangle of the CKM matrix.	15
1.2. Feynman diagrams for the 1 loop processes by which the decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ can proceed.	16
1.3. The expected constraints on the $\bar{\rho}, \bar{\eta}$ plane from the NA62 and KOTO experiments. .	20
1.4. Limits on $ U_{eh} $ versus m_h in the mass range $10 \text{ MeV}/c^2 - 100 \text{ GeV}/c^2$ from peak searches and ν_h decays.	30
1.5. Limits on $ U_{\mu h} $ versus m_h in the mass range $100 \text{ MeV}/c^2 - 100 \text{ GeV}/c^2$ from peak searches and ν_h decays.	31
1.6. Current limits on $ U_{\mu h} ^2$ from peak searches in kaon decays.	32
1.7. Ratios of kinematic quantities for the heavy and massless neutrinos as a function of neutrino mass (helicity and phase space suppression).	33
1.8. Partial widths for heavy neutrino decays into SM particles as a function of heavy neutrino mass	35
1.9. Heavy neutrino lifetime, assuming decays to SM particles, as a function of heavy neutrino mass	35
2.1. The beamline and detector for NA62 data taking in 2007.	41
2.2. The NA62 subdetectors downstream of the decay volume.	42
2.3. The charged hodoscope detector	43
2.4. The LKr and neutral hodoscope detectors	45
2.5. The hadronic calorimeter	46
2.6. The three planes of the muon veto system	47
2.7. A schematic diagram of the lead wall used to study muon energy deposition in the LKr	49
3.1. Heavy neutrino signal distributions	55
3.2. Kinematic distributions of the muon produced in $K^+ \rightarrow \mu \nu_h$ decay.	55

3.3. The squared missing mass spectra of the six most common decay modes of K^+ when reconstructed assuming the muon mass for the charged particle detected. The MC samples are normalized according to the decay branching ratios, while the halo distribution is data driven and scaled to match the data. The distributions shown are the predicted spectra when only the single track pre-selection is applied.	56
3.4. The weights applied to MC events in run 20410 as a function of simulated kaon momentum in order to improve the width of the kaon momentum distribution. . . .	61
3.5. The effect of weighting the MC samples according to the kaon momentum.	62
3.6. The projective structure of the LKr calorimeter cell towers, showing how the cluster position must be corrected. x_p and x_r in the figure correspond to x_0 and x in equation 3.19.	63
3.7. Components of the residual magnetic field along the z -axis of the decay region. . . .	65
3.8. MC spectrum of $(m_{\text{miss}}^{\pi^0})^2 = (p_K - p_{\pi^0})^2$ for $K^+ \rightarrow \pi^+ \pi^0$ decays, using the LKr calorimeter for selection and reconstruction.	67
3.9. The geometry of $K^+ \rightarrow \pi^+ \pi^0$ decay showing the variables used in the kinematic fit. . . .	69
3.10. MC distributions of fit χ^2 and E/p for events reconstructed in the LKr calorimeter . . .	70
3.11. Distributions of $(m_{\text{miss}}^{\pi^+})^2 - m_{\pi^0}^2$ and track p_T with respect to the kaon momentum, reconstructed using the magnetic spectrometer, for $K^+ \rightarrow \pi^+ \pi^0$ candidates selected using the LKr calorimeter	73
3.12. Angular deflection, ϑ , determined from upstream and downstream track directions . . .	74
3.13. The distribution used to generate random numbers for populating the far tails of the multiple Coulomb scattering distribution. Here the effect of adding the extra scattering distribution $P(\zeta)$ to a Gaussian distribution is shown. For a given Gaussian distribution and value of ζ_0 , the function is integrated numerically and normalized to an area of 1. . . .	76
3.14. The effect of adding extra scattering to the MC simulation of the spectrometer.	77
3.15. Comparison of (a) the CDA distribution and (b) the p_T distribution of $K^+ \rightarrow \pi^+ \pi^0$ events reconstructed using the spectrometer, with and without extra scattering simulated.	78
3.16. The distribution of $(m_{\text{miss}}^{\pi^+})^2 - m_{\pi^0}^2$ in $K^+ \rightarrow \pi^+ \pi^0$ events	79
3.17. The reconstructed m_{miss}^2 spectrum, $(p_K - p_\mu)^2$, for the sum of all the background contributions showing the effect of applying scattering to the MC simulation.	79

3.18. The effect of the muon veto multiplier ξ in equation 3.45 on muon identification efficiency, shown for a selection including all track momenta, and for a more restricted selection.	82
3.19. The parameters used to describe multiple Coulomb scattering across a single slice in z	82
3.20. The momentum and squared missing mass distribution of events passing the single track pre-selection.	85
3.21. The distribution of muon veto efficiencies measured on individual bursts.	86
3.22. Muon detection efficiency for the bursts in run 20438. All of the bursts in the region near the start of the run are excluded.	86
3.23. Muon veto efficiency distributions observed in data as function of momentum and xy coordinate at muon veto	87
3.24. Track xy positions at the muon veto. (a) all tracks with $30 \text{ GeV}/c < \vec{p} _{\mu} < 60 \text{ GeV}/c$ (b) the selection applied in the xy plane for further study.	88
3.25. Muon detection efficiency as a function of momentum, for tracks in the region shown in figure 3.24.	89
3.26. b	89
3.27. Simulation of the beam halo background with the HALO and CMC programs.	91
3.28. Kinematic distributions of halo events reconstructed with a single track pre-selection. showing how the events are divided into <i>upper</i> and <i>lower</i> regions for further analysis.	93
3.29. Kinematic distributions of halo events in the ‘upper’ and ‘lower’ kinematic regions Cuts are shown as red lines.	95
3.30. Kinematic distributions of signal events for neutrinos with masse $250 \text{ MeV}/c^2$ (blue) and $350 \text{ MeV}/c^2$ (green).	98
3.31. The squared missing mass distribution for positively charged muon tracks obtained from the K^- data taking period.	99
3.32. The effect of kaon punchthrough on the halo measurement.	100
3.33. The distribution of reconstructed kaon mass (the invariant mass of the three pions) in $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ decays, shown here for K^+ data. The cut applied for the beam momentum measurement is shown by red arrows.	101

3.34. The squared missing mass distribution of the negatively charged (μ^-) beam halo in the K^+ data taking period for the signal upper kinematic selection (with negative track requirement), compared with the K_{less} period and the simulated contribution from $K^+ \rightarrow \pi^+ \pi^+ \pi^-$, using a relaxed selection described in section 3.7.8.	102
3.35. The squared missing mass distribution of the positively charged (μ^+) halo for the K^- data taking period, using a relaxed selection described in section 3.7.8, compared with the K_{less} data sample.	104
3.36. The distribution of the positively charged (μ^+) halo for the K^- data taking period, as described in figure 3.35, but here using the full signal selection.	104
3.37. The halo distribution measured in the K^- period, before and after MC $K^- \rightarrow \pi^- \pi^- \pi^+$ subtraction (statistical errors only).	105
3.38. (a) The negative region of the squared missing mass distributions for K^+ (points with error bars) and K^- (solid line) data periods, after applying the full selection in the upper kinematic region. (b) The distribution of reconstructed squared missing mass and muon opening angle w.r.t the kaon trajectory for all halo events in the K^- period.	106
3.39. The fraction of halo events which passed the final selection, relative to those which passed the R_K selection.	107
3.40. Squared missing mass distributions for the control sample used to scale the halo estimate from K^- data to match the K^+ data. All plots show $\#events/(0.005 \text{ GeV}^2/c^4)$	108
3.41. The final estimate of the halo contribution to the final signal selection. (The signal region begins at $m = 275 \text{ MeV}/c^2$, $m^2 = 0.75 \text{ GeV}^2/c^4$.)	108
3.42. The simulated distribution of background events for the final heavy neutrino selection	112
3.43. The signal acceptance in the final heavy neutrino selection as a function of neutrino mass. Discontinuities are due to the discrete signal window sizes (see section 3.11.1).	112
3.44. The fractional error on the background estimate coming from the simulation of the far tails of the spectrometer resolution. The two bins with higher content at high squared missing mass are the result of statistical fluctuation in the MC estimate of the spectrometer contribution. The contribution at these masses is negligible compared to other systematic errors (see figure 3.50).	115
3.45. Extraction of the integrated kaon flux and integrated uncertainty.	117

3.46. Muon veto efficiency in various kinematic projections: (a) squared missing mass, (b) momentum, (c) muon - kaon opening angle, θ , (d) vertex z -coordinate, (e) track radius at MUV 1, (f) track azimuthal angle ϕ	118
3.47. The $Q_1 \times 1\text{TRKL}$ trigger efficiency, measured using the Q_1 data sample.	120
3.48. Simulated kinematic distributions used for the halo model validity check described on page 3.10.5. The quantity 'errscale' is $1 - f$, where f is the scale factor required to achieve a good χ^2 value.	124
3.49. The optimum half-width, σ_h for the heavy neutrino signal window as a function of neutrino mass. (see section 3.11.1 and equation 3.82).	127
3.50. The squared uncertainties which contribute to the final upper limit on the number of heavy neutrinos observed at each mass. The coloured bands show the contributions to the uncertainty on the background estimate (added together). The blue line shows the statistical contribution from the expected fluctuation of the data about the true background value. The contribution labelled 'halo statistical' is a separate contribution which comes from the statistical uncertainty on the estimate of the true halo contribution.	127
3.51. The expected sensitivity, at 90% C.L, to $ U_{\mu 4} ^2$ and the single event sensitivity for this analysis, shown along with the existing limits from production searches in $K \rightarrow \mu \nu_h$ decays at the E949 and the experiment at KEK.	129
3.52. The squared missing mass distribution of the predicted backgrounds, along with examples of the signals that would be expected for $\mathcal{B}(K^+ \rightarrow \mu \nu_h) = 1 \times 10^{-4}$, shown for heavy neutrino masses from $270 \text{ MeV}/c^2$ to $370 \text{ MeV}/c^2$ in $20 \text{ MeV}^2/c$ intervals.	129
3.53. Data and background estimate comparison for the missing mass spectrum (a) on a log-scale and (b) on a linear scale for the signal region. The statistical errors for data are shown by error bars. The width of the filled region indicates the uncertainty on the background prediction.	130
3.54. (a) The distributions of pulls for each mass. (b) The upper limit (at 90% C.L.) on the number of heavy neutrinos detected at each mass, together with bands showing the expected 1σ and 3σ fluctuation around the expected sensitivity.	130

3.55. The upper limits using the 2007 data set from the NA62 experiment, as a function of m_h , the neutrino mass for (a) the branching ratio $K^+ \rightarrow \mu^+ \nu_h$, and (b) the matrix element $ U_{\mu 4} ^2$, shown with the existing limits from peak searches in $K^+ \rightarrow \mu^+ \nu_\mu$ decay at the experiment at KEK and the E949 experiment E949	131
4.1. The E949 detector at BNL	135
4.2. The $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ events recorded by the E787 and E949 collaborations . ‘This analysis’ refers to E949-PNN2 events.	137
4.3. The evolution of limits on $BR(K^+ \rightarrow \pi^+ \nu \bar{\nu})$	137
4.4. Distribution of m_{miss}^2 for the main background channels. The signal distribution is multiplied by 10^{10}	139
4.5. The NA62 Experimental Setup	140
4.6. Schematic of the NA62 beam line close to the beam spectrometer, showing the nominal kaon path (purple), and the beam envelope (blue)	142
4.7. A CHANTI station in exploded and assembled views.	143
4.8. The four views of a straw chamber	144
4.9. The three MUV subdetectors following the Liquid Krypton Calorimeter	145
4.10. Schematic design of MUV3 tiles	146
4.11. The RICH subdetector	147
4.12. The layout of the downstream NA62 sub-detectors, showing the nominal beam trajectory in red and the beam envelope in blue.	148
4.13. A single LAV station	149
4.14. The trigger system system for NA62	150
5.1. A photograph of the light diaphragm showing two of the eight sectors	154
5.2. The effect of various broadening mechanisms on the light ring seen by the CEDAR detector. a) Good alignment. b) Broadening due to multiple coulomb scattering. c) Misalignment (see chapter 6). d) Effect of beam divergence	155
5.3. Inefficiency for kaon tagging as a function of the number of photons expected per kaon and the number of sectors required in coincidence (called an n-Fold coincidence). The dashed line shows the target inefficiency for KTAG at 5%.	157
5.4. The layout of the CEDAR detector	158
5.5. The optical system the CEDAR detector	159

5.6. The optical components of the KTAG upgrade to CEDAR as simulated in Geant4. . . .	160
5.7. Left: Simulation of path taken by Chrenkov photons from the CEDAR quartz windows (blue), to the new PMT planes (green), via the spherical mirrors (red). Right: The mechanical framework for KTAG.	160
5.8. PMTs used in KTAG. The 16 R7400U-03 PMTs are coloured red. The 32 R9880U-210 PMTs are coloured yellow.	161
5.9. Wavelength distributions in nitrogen. The refractive index is always greater than $1/\beta$ so Cherenkov radiation is emitted over the whole range of wavelengths shown. . . .	164
5.10. Transmission and reflectivity spectra of CEDAR optical components. The absorption properties of all the internal optical components are combined into a single efficiency.	165
5.11. Distribution of photons at the CEDAR diaphragm	165
5.12. Cherenkov radiation from kaons (red) and pions (blue) at the CEDAR diaphragm. In the top plot, wavelength is plotted against radius, showing the long tails which overlap at low wavelengths (before the effect transmission and quantum efficiency). In the bottom plot, the number of photons is normalized taking into account the expected flux of kaons and pions in the beam composition.	166
5.13. PMT quantum efficiencies and the source and filter spectra used to measure them. The three spectra: Corning Blue, Standard Eye, and Tungsten source ($T=2856$ K) are described in Appendix B	167
5.14. Light profile on KTAG light-guides for various combinations of spherical mirror radius, R_M , and optical cap lens focal length, F_L . White circles indicate the locations of the PMTs. The red dot indicates the location of the new PMT introduced in the chosen configuration.	168
5.15. Timing distributions for Hamamatsu R7400U-03 PMTs. For each PMT hit ΔT is the difference in time between the individual hit and the average of all the other PMT hits in the reconstructed kaon candidate.	169
5.16. Timing distributions for Hamamatsu R9880U-210 PMTs	170
6.1. KTAG inefficiency for 6 fold coincidences as a function of CEDAR misalignments in x , compared with the maximum inefficiency allowed in the NA62 specification.	175
6.2. The numbering scheme for KTAG sectors. The kaon beam travels in the positive z direction, which is <i>into</i> the page in this figure.	176

6.3. A toy MC simulation of the Cherenkov light from kaons passing through the diaphragm aperture. The dashed lines represent the diaphragm aperture, the dark fill represents photons which pass through the aperture and the light fill represents photons which miss the aperture.	177
6.4. The effect of misalignment for the toy model described in section 6.2.1. The left hand plot shows N_U and N_D as a function of simulated y misalignment. The right hand plot shows both $A_{U/D}$ and $A_{L/R}$ as a function of simulated y misalignment.	179
6.5. The second toy model from section 6.2.1 with a narrower diaphragm aperture.	179
6.6. The distribution of photons as a function of radius at the diaphragm aperture with a “pencil” beam when the CEDAR is aligned (left), and when the simulated y misalignment is 0.5 mm (right). The diaphragm aperture is 1 mm so no light is seen outside the range 99.5 mm – 100.5 mm. Since the detector is approximately left/right symmetrical, the distributions for sectors 5,6,7 and 8 are qualitatively similar to those for sectors 4,3,2 and 1.	180
6.7. The same distributions as for figure 6.6 but using a Gaussian angular beam profile instead of a pencil beam.	180
6.8. The dependence of $A_{U/D}$ on the simulated y misalignment, with 1 mm diaphragm aperture, for two toy models of the beam angular profile.	180
6.9. The angular distribution of the beam in the NA62MC simulation compared with a Gaussian curve of the same rms for both x and y components.	181
6.10. NA62 MC simulated response of asymmetries at various misaligned positions, with diaphragm apertures between 1 and 5 mm. No corrections are applied to the asymmetries so the lines do not pass through the origin.	183
6.11. NA62MC simulated response of <i>corrected</i> asymmetries at various misaligned positions, with diaphragm apertures between 1 and 5 mm. The individual sector counts are re-weighted before computing the asymmetries.	184
6.12. Calibration curves for estimating CEDAR misalignment in x and y from the asymmetries $A_{L/R}$ and $A_{U/D}$, with diaphragm aperture 1.5 mm. In each plot, the points mark the asymmetries found in simulations of known misalignment. The line is the calibration used subsequently to determine the misalignment from asymmetries.	185
6.13. Estimated misalignment in x (left) and y (right) using the asymmetry method indicated by colour, as a function of simulated misalignment in x and y	186

6.14. Residual misalignment: simulated misalignment – estimated misalignment, indicated by colour.	186
6.15. (left) $A_{L/R}$ as a function of simulated x misalignment, showing the breakdown at large misalignment. (right) Re-weighted quadrant counts as a function of simulated x misalignment.	186
6.16. $\ln \chi^2$ distributions obtained for a simulated misalignment of $x = 0.8$ mm, $y = 0.8$ mm	189
6.17. Visualisations of PMT layout.	190
6.18. Correlation between measures of PMT efficiency.	192
6.19. The effect of variation in efficiency between PMTs. (left) the residual misalignment when using the χ^2 method with known PMT efficiencies. (right) the residual misalignment when the PMT efficiencies in the simulated templates vary with respect to the PMT efficiencies in the test samples (see section 6.3.3)	192
6.20. A pressure scan from a test beam in 2011 showing the number of coincidences per trigger as a function of the CEDAR radiator pressure. The peaks correspond (from left to right) to the pions, kaons and protons.	193
6.21. MC Alignment plots when the radiator gas pressure is not the nominal one. The plots show χ^2 vs misalignment in x and y using the sector grouping method. In the top row, the data simulation is aligned. In the bottom row the misalignment is [+1 mm, +1 mm].	193
6.22. CEDAR motor positions during a 1 hour period of data taking on 27th November 2014. 5 periods are identified with distinct configurations. The table lists the recorded motor positions for each period.	195
6.23. Plots of $\ln \chi^2$ vs template misalignment for the first burst of alignment period 1 (defined in figure 6.22). (a) sector grouping. (b) sixths grouping. Minima are circled in red.	196
6.24. Plots of $\ln \chi^2$ vs template misalignment for a MC sample produced with similar misalignment to that seen in figure 6.23. (a) sector grouping. (b) sixths grouping.	197
6.25. Estimates, from the χ^2 fits to all PMTs, of the motor positions at which the CEDAR should be perfectly aligned. (a) All estimates. (b) Estimate most consistent with other values.	198
6.26. Estimates, from the χ^2 fits to only the new PMTs, of the motor positions at which the CEDAR should be perfectly aligned. (a) All estimates. (b) Estimate most consistent with other values.	198

6.27. The asymmetries $A_{L/R}$ and $A_{U/D}$ for each burst in the 5 periods under consideration plotted against the estimated misalignment in x and y determined from the χ^2 method.	200
6.28. The ratio, in MC simulation, of the two expressions for the number of photo-electrons per kaon, ϕ_{68} and ϕ_{78} , as a function of misalignment in x .	201
6.29. The alignment tab of the KTAG online monitor	203

List of Tables

1.1. Dimensions and charges of the Standard Model fields.	4
3.1. Relevant statuses of the NA62- R_K detector configuration for each of the data taking periods in 2007.	52
3.2. The muon status definitions depending which MUV planes have associated hits.	80
3.3. Numerical values for constants used in determining the scattering contribution in equation 3.42	81
3.4. The cuts in the plane of track momentum (p) and vertex z coordinate (z_{vertex}), used in the R_K analysis and adopted for the MUV efficiency study.	84
4.1. The most common K^+ decay modes along with the strategies for suppressing them at NA62.	138
6.1. The correction factors used to re-weight the number of counts in each sector when calculating the asymmetries $A_{U/D}$ and $A_{L/R}$	182

Introduction

The study of kaons has been instrumental in the development of the Standard Model of Particle Physics in the second half of the last century; now kaon physics plays a crucial role in the search for New Physics, beyond the Standard Model.

The kaon was first discovered as two separate particles, θ^+ and τ^+ , distinguished by their decays into states of different parities, but inexplicably similar in mass and lifetime. The $\tau - \theta$ puzzle was resolved by the realization that parity is not a conserved symmetry of the weak interaction and that the θ^+ and τ^+ are in fact the same particle, now called the K^+ . Less than a decade later, it was discovered that the weak interaction also violates the combined symmetry of parity and charge conjugation. Again the revelation came from kaons, this time in the decays of the neutral kaons, K^0 and \bar{K}^0 . As the Standard Model was developed, the kaon sector provided evidence for many of the necessary insights such as ‘strangeness’ to explain the long lifetime of the kaon and the GIM mechanism to explain the absence of flavour changing neutral currents in kaon decays, leading to the prediction of the charm quark.

The results in this thesis come from the NA62- R_K and NA62 experiments, the latest of a long history of fixed target kaon experiments at the CERN North Area (NA) served by the Super Proton Synchrotron (SPS) accelerator. Today, NA62 aims to study the ultra-rare decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ along with a wide selection of searches for other rare decays as well as making precision tests of the Standard Model. The NA62 experiment began taking data this year (2015) but a data sample was also collected with the NA62- R_K detector which was used for precise measurements of lepton universality, the fact that the W boson couples to the electron and the muon with the same strength.

This thesis is divided into two parts: a search for heavy neutrinos using the 2007 NA62- R_K dataset, and a description of my contribution to the development of the kaon identification subdetector for the NA62 experiment.

The first chapter outlines the theoretical motivation for the search for heavy neutrinos and for the measurement of $K^+ \rightarrow \pi^+ \nu \bar{\nu}$: both are related to the origins of mass in the Standard Model.

Neutrinos are unlike the other fundamental fermions in the Standard Model, in that their masses are so small that they were once considered massless. Natural explanations for the smallness of the neutrino masses lead to the prediction of new particles, ‘heavy neutrinos’, which could be discovered in the data from NA62. In the quark sector, the relationship between mass and flavour in the Standard Model is encoded in the CKM matrix. The properties of this matrix have deep implications for particle physics and the decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ is a precision tool for examining them. In chapter 2, the 2007 detector is described and then in chapter 3, an analysis is reported in which the 2007 data set is used to place limits on the coupling between heavy neutrinos and muons. The limiting factors discovered in the analysis suggest that improved measurement could be achieved the NA62 detector. In chapter 4, the NA62 detector is described, followed, in chapter 5, by a more detailed description of the KTAG detector for kaon identification. A study of alignment methods for the KTAG is described in chapter 6.

Chapter 1

Theory

In this chapter, I introduce the theoretical background and motivation for the experimental work described in the following chapters. A brief description of electroweak symmetry breaking is presented, focussing on the generation of fermion masses. The CKM matrix and the question of neutrino mass arise as natural consequences of the Higgs mechanism and are discussed in subsequent sections. Some additional definitions are presented in appendix A.

1.1. Mass in the Standard Model

The Standard Model (SM) is constructed from fields which transform under representations of the symmetry group

$$SU(3) \otimes SU(2)_L \otimes U(1)_Y \quad (1.1)$$

where the $SU(3)$ symmetry corresponds to *colour* in strong interactions, $SU(2)_L$ corresponds to *weak isospin* and $U(1)_Y$ corresponds to *hypercharge* [1]. Table 1.1 shows the dimension of the representation and the charges of the SM particle fields under the various symmetries. Any term entering the Lagrangian of the SM must be invariant under local $SU(3) \otimes SU(2)_L \otimes U(1)_Y$ gauge transformations.

Mass generation in the SM is driven by the $SU(2)_L \otimes U(1)_Y$ subgroup. In the lepton sector, the

Field	$SU(3)$	$SU(2)_L$	T^3	$Y/2$	$Q = T^3 + Y/2$
g_μ^a	8	1	0	0	0
(W_μ^\pm, W_μ^0)	1	3	$(\pm 1, 0)$	0	$(\pm 1, 0)$
B_μ^0	1	1	0	0	0
$q_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}$	3	2	$\begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}$	1/6	$\begin{pmatrix} 2/3 \\ -1/3 \end{pmatrix}$
u_R	3	1	0	2/3	2/3
d_R	3	1	0	-1/3	-1/3
$\chi_L = \begin{pmatrix} \nu_L \\ e_L \end{pmatrix}$	1	2	$\begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}$	-1/2	$\begin{pmatrix} 0 \\ -1 \end{pmatrix}$
e_R	1	1	0	-1	-1
$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}$	1	2	$\begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}$	1/2	$\begin{pmatrix} 1 \\ 0 \end{pmatrix}$
$\widehat{\Phi} = \begin{pmatrix} \phi^0 \\ \phi^- \end{pmatrix}$	1	2	$\begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}$	-1/2	$\begin{pmatrix} 0 \\ -1 \end{pmatrix}$

Table 1.1.: Dimensions and charges of the Standard Model fields.[2]

For example, the left-chiral up-type quarks, u_L , are triplet representations of $SU(3)$ (they have 3 colors), they form part of a doublet of $SU(2)_L$ along with the down-type quarks. They have a third component of weak isospin $T^3 = +1/2$ and weak hypercharge $Y/2 = 1/6$.

g_μ^a are the gluon fields. W_μ^\pm and W_μ^0 are the gauge bosons of the weak isospin $SU(2)_L$ symmetry. B_μ is the gauge boson of the hypercharge $U(1)_Y$ symmetry. W_μ^0 and B_μ^0 mix to form the Z_μ and A_μ fields corresponding to the Z boson and the photon. q_L is the electroweak doublet of left-chiral quark fields, combining up-type, u_L and down-type, d_L . The right-chiral fields u_R and d_R form electroweak singlets. The lepton fields follow the same pattern but ν_R is absent in the SM. Φ is the Higgs doublet of complex scalar fields and $\widehat{\Phi}$ is the Lorentz conjugate state $\widehat{\Phi} = i\sigma^2\Phi^*$, where σ^2 is the second Pauli matrix.

fields, χ_L and e_R transform under this group as follows:

$$\chi_L \rightarrow \chi' = \exp\left[-ig\frac{\vec{\sigma}}{2} \cdot \vec{\Delta} + i\frac{1}{2}g'\Lambda\right]\chi_L \quad (1.2)$$

$$e_R \rightarrow e' = \exp(ig'\lambda)e_R \quad (1.3)$$

where g and g' are the $SU(2)$ and $U(1)_Y$ couplings, $\Delta(\vec{x}) = (\Delta^1(x), \Delta^2(x), \Delta^3(x))$ and $\Lambda(x)$ specify arbitrary local $SU(2)$ and $SU(1)$ gauge transformations, and $\vec{\sigma} = (\sigma^1, \sigma^2, \sigma^3)$ are the Pauli matrices. Kinetic terms such as $\overline{\chi}_L\gamma^\mu\partial_\mu\chi_L$ are not invariant under these local transformations. Local gauge invariance can be restored by introducing the four vector fields $\vec{W} = (W^1, W^2, W^3)$ and B_μ , which

have the gauge transformations:

$$\vec{W}_\mu \rightarrow \vec{W}'_\mu = W_\mu - g\vec{\Delta} \times \vec{W}_\mu + \partial_\mu \vec{\Delta} \quad (1.4)$$

$$B_\mu \rightarrow B'_\mu = B_\mu + \partial_\mu \Lambda \quad (1.5)$$

and defining the covariant derivative

$$D_\mu = \partial_\mu + \frac{i}{2} g \vec{\tau} \cdot \vec{W}_\mu + i g' \frac{Y}{2} B_\mu. \quad (1.6)$$

With these definitions, one can construct the term $\bar{\chi}_L \gamma^\mu D_\mu \chi_L$, which is invariant under local $SU(2)_L \otimes U(1)_Y$ transformations.

The sub-group $SU(2)_L \otimes U(1)_Y$ describes a *chiral* symmetry because the left- and right-chiral fields (e.g. u_L and u_R) have different weak isospin and hypercharge numbers. The chirality of spinor fields is defined by the projection operators:

$$P_L = \frac{1}{2}(1 - \gamma_5), \quad P_R = \frac{1}{2}(1 + \gamma_5), \quad P_L + P_R = 1 \quad (1.7)$$

where $\gamma_5 = i\gamma_0\gamma_1\gamma_2\gamma_3$ is formed from the product of the Dirac matrices. Hence, a spinor field, ψ , can be decomposed into its chiral components:

$$\psi = P_L \psi + P_R \psi = \psi_L + \psi_R, \quad (1.8)$$

and it is these chiral components which participate in the weak interaction. The asymmetry between left- and right-chiral fields is seen in the weak charged current fields W^\pm , which couple only to the left-chiral fermions (χ_L and q_L in table 1.1). For example in the lepton sector the charged current interaction for the first generation is described by:

$$\mathcal{L}_{CC} = -\frac{g}{\sqrt{2}} \bar{\chi}_L \gamma^\mu \left[\sigma^+ W_\mu^+ + \sigma^- W_\mu^- \right] \chi_L, \quad (1.9)$$

where $\chi_L = \begin{pmatrix} \nu_e \\ e \end{pmatrix}^T$ is an $SU(2)_L$ doublet and σ^\pm are combinations of the Pauli matrices, $\sigma^\pm = \sigma^1 \pm \sigma^2$. The right-chiral electron field, e_R , plays no part in the interaction. In the SM there is no ν_R field, so it cannot play any part.

The different gauge interactions of left- and right-chiral fermions prevent the introduction of Dirac

type fermion masses into the SM Lagrangian. A fermion mass term would look like

$$-\mathcal{L}_{\text{mass}} = m\bar{\psi}\psi = m(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L) \quad (1.10)$$

which is not invariant under $SU(2)_L \otimes U(1)_Y$ transformations. For example, the down-quark mass is described by:

$$-\mathcal{L}_{\text{down mass}} = m\bar{d}d = m(\bar{d}_Ld_R + \bar{d}_Rd_L), \quad (1.11)$$

and, referring to table 1.1, it can be seen that the term $m\bar{d}_Rd_L$ has hypercharge $+1/3 + 1/6 = 1/2$, so it is clearly not $U(1)_Y$ invariant. Masses for the vector gauge bosons are also forbidden by their respective symmetries. A mass term for \vec{W}_μ would violate $SU(2)_L$ and a mass term for B_μ would violate $U(1)_Y$.

Both gauge boson and fermion masses can be generated spontaneously with the introduction of the Higgs boson. The charged scalar Higgs doublet $\Phi = \begin{pmatrix} \phi^+ & \phi^0 \end{pmatrix}^T$ can be expanded in terms of four real scalar fields

$$\phi^+ = \frac{\phi_1 + i\phi_2}{\sqrt{2}}, \quad \phi^0 = \frac{\phi_3 + i\phi_4}{\sqrt{2}}. \quad (1.12)$$

Using the covariant derivative defined in equation 1.6, the scalar part of the Lagrangian can be written

$$\mathcal{L}_\Phi = (D_\mu\Phi)^\dagger D^\mu\Phi - V(\Phi), \quad V(\Phi) = \mu^2(\Phi^\dagger\Phi) - \lambda(\Phi^\dagger\Phi)^2. \quad (1.13)$$

If $\lambda < 0$ and $\mu^2 < 0$ then the potential $V(\Phi)$ has a minimum when

$$\Phi^\dagger\Phi|_{\text{min}} = \frac{\mu^2}{2\lambda} \quad (1.14)$$

Equation 1.14 describes a degenerate circle of minima in the phase space of Φ . The $SU(2)_L \times U(1)_Y$ symmetry is broken when the vacuum ground state spontaneously chooses a particular direction in this basis. Without loss of generality, we can choose the direction where all but one of the fields are zero:

$$\langle 0 | \phi_1 | 0 \rangle = \langle 0 | \phi_2 | 0 \rangle = \langle 0 | \phi_4 | 0 \rangle = 0 \quad \langle 0 | \phi_3 | 0 \rangle^2 = v^2 = \frac{\mu^2}{\lambda} \quad (1.15)$$

That is, the field ϕ_3 has a *vacuum expectation value* (vev) of v , while the other fields have zero vev. Furthermore, it is possible to expand the Higgs field in a particular gauge, the *unitary gauge*, such that the fields with zero vev are eliminated:

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H \end{pmatrix}, \quad (1.16)$$

where v is the Higgs vev and H represents the excitation of the Higgs field above its vev (corresponding to the Higgs Boson particle).

After spontaneous symmetry breaking, the vector gauge bosons W^\pm and Z^0 obtain masses through interaction with the Higgs via the covariant derivative in equations 1.6 and 1.13. Fermion masses arise from *Yukawa couplings*.

The addition of the Higgs doublet allows for additional gauge invariant terms which can generate fermion masses. In particular, the following term for the electron (described by the field e) is $SU(2) \times U(1)_Y$ invariant:

$$\mathcal{L}_{\text{Yukawa}}(\text{electron}) = Y_e [\bar{\chi}_L \Phi e_R + \bar{e}_R \Phi^\dagger \chi_L], \quad (1.17)$$

because the relevant charges of the lepton and Higgs fields sum to zero. Here e is electron spinor field operator, with chiral projections $e_{L/R} = P_{L/R}e$, and Y_e is the strength of the *Yukawa coupling* of the electron to the Higgs field. After spontaneous symmetry breaking, the Lagrangian can be rewritten (using equation 1.16) as

$$\mathcal{L}_{\text{Yukawa}}(e) = -\frac{Y_e v}{\sqrt{2}} \bar{e}e - \frac{Y_e}{\sqrt{2}} \bar{e}eH, \quad (1.18)$$

so the electron obtains a mass as well as a coupling to the Higgs boson. Since there is no available ν_R term, the neutrino remains exactly massless in the SM, by construction. The muon and tau masses are generated by exact copies of the Yukawa coupling shown for electrons.

The generation of mass for the down type quarks follows a similar pattern but here we explicitly allow for the possibility that the basis of quark states in which couplings to the Higgs field are diagonal is not necessarily the basis in which the weak current couplings are diagonal.

$$\mathcal{L}_{\text{Yukawa}}(d) = \sum_{f=1,2,3} -\left[\bar{q}_L^f Y_{ff}^d \Phi d_R^{f'} + h.c. \right] \quad (1.19)$$

Here f runs over the 3 quark generations and Y^d is the down-type matrix of Yukawa couplings. The physical states can be obtained by diagonalizing Y^d with a bi-unitary transformation to produce the diagonal mass matrix:

$$M_{\text{diag}}^d = V_L^d Y^d V_R^d, \quad (1.20)$$

where V_L^d and V_R^d , are the unitary matrices describing how the left- and right-chiral states, (d_L, d_R) are related to the physical ones $(d_L^{(m)}, d_R^{(m)})$:

$$d_L = V_L^d d_L^{(m)}, \quad d_R = V_R^d d_R^{(m)}. \quad (1.21)$$

It is not possible to generate mass for the up-type quarks using the field Φ , since the term $q_L^\dagger \Phi u_R$ is not invariant. Instead, one can construct the Lorentz conjugate state:

$$\widehat{\Phi} = i\sigma^2 \Phi^* = \begin{pmatrix} H + v \\ 0 \end{pmatrix}, \quad (1.22)$$

which is an $SU(2)_L$ doublet but has the opposite hypercharge to Φ , $Y(\widehat{\Phi}) = -1/2$. Then it is possible to obtain invariant Yukawa terms for the up-type quarks:

$$\mathcal{L}_Y(u) = \sum_{f=1,2,3} - \left[\bar{q}_L^f Y_{ff'}^u \widehat{\Phi} u_R^{f'} + h.c. \right]. \quad (1.23)$$

Again the physical states can be obtained by diagonalizing the Yukawa mass matrix:

$$M_{\text{diag}}^u = V_L^u Y^u V_R^u \quad (1.24)$$

In general, it is not possible to simultaneously diagonalize Y^u and Y^d . This leads to an irreducible mixing between the physical states and the weak eigenstates, parametrized by the CKM matrix:

$$V_{\text{CKM}} = V_L^d V_L^{u\dagger} \quad (1.25)$$

The Higgs mechanism of mass generation by spontaneous symmetry breaking raises many questions and opens the door to many further avenues of investigation. In this thesis, two of those avenues are explored. Firstly, what is the structure of the CKM matrix and what can we learn from it? Secondly, how can we account for the now irrefutable experimental evidence that, in fact, neutrinos are not

massless but have masses many orders of magnitude below those of the other fermions?

1.2. The CKM framework

As discussed above, the CKM matrix arises in the SM because the Yukawa mass matrices for the up-type and down-type quarks ($Y^{u,d}$ in equations 1.19 and 1.23) cannot be simultaneously diagonalized. The non-trivial quark mixing can be seen in the expression for the charged weak currents, which couple to the W^\pm bosons in the Lagrangian:

$$J_{CC}^\mu = \bar{u}_L \gamma^\mu d_L \quad (1.26)$$

From equations 1.20 and 1.24, this interaction can be rewritten in terms of the mass eigenstates.

$$J_{CC}^\mu = \bar{u}_L^{(m)} \gamma^\mu (V_L^{u\dagger} V_L^d) d_L^{(m)} \quad (1.27)$$

Here the CKM matrix ($V_{CKM} = V_L^d V_L^{u\dagger}$) enters as the coupling of quarks of different generations to the W^\pm bosons.

Without loss of generality, the up type quarks can be considered in a basis where $u_L^{(m)} = u_L$ implying that Y^u is diagonal, $V_L^u = V_R^u = \mathbb{1}$ and $V_{CKM} = V_L^d$. Now V_{CKM} can be interpreted as the mixing between the down type physical mass states and the down type states for which the charged current interactions are diagonal. Thus, it is notated:

$$\begin{pmatrix} d \\ s \\ b \end{pmatrix} = V_{CKM} \begin{pmatrix} d^{(m)} \\ s^{(m)} \\ b^{(m)} \end{pmatrix}, \quad V_{CKM} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \quad (1.28)$$

where d, s, b are the weak eigenstates which interact with u, c, d in the charged current weak interactions and $d^{(m)}, s^{(m)}, b^{(m)}$ are the physical mass states. For example, the ratio V_{cd}/V_{cs} determines the ratio of amplitudes for a charm quark to decay into an strange or a down quark mass state. Since experiments show that $V_{cd} \ll V_{cs}$, decays involving $c \rightarrow d$ are said to be *Cabbibo suppressed*. Neutral currents which couple to the Z^0 boson do not mix up- and down-type quarks in the weak basis:

$$J_{NC}^\mu = \frac{1}{2} (\bar{u}_L^f \gamma^\mu u_L^f) + \dots \quad (1.29)$$

By the unitarity of the four transformation matrices $V_{L,R}^{u,d}$, these terms remain diagonal in the mass

basis:

$$J_{NC}^\mu = \frac{1}{2} \left(\bar{u}_L^{(m)f} \gamma^\mu u_L^{(m)f} \right) + \dots \quad (1.30)$$

This explains the absence of *flavour changing neutral currents* (FCNC) at tree level (since there are no cross terms coupling different flavours). The unitarity of the CKM matrix also ensures the suppression of FCNC's in loop diagrams, through the GIM mechanism which was, historically, the motivation for the prediction of the charm quark. [3].

1.2.1. Parameter count

A general $n \times n$ matrix of complex numbers has $2n^2$ free parameters. However, the number of parameters which the CKM matrix introduces into the SM is less than this. The unitarity condition,

$$V_{ji}^* V_{jk} = \delta_{ik} \quad \forall i, k \quad (1.31)$$

provides n^2 constraints, reducing the number of free parameters to n^2 . In addition, there are a number of residual arbitrary phases in the definitions of the quark mass states. The following transformations:

$$u^{(m)f} \rightarrow e^{i\phi_f^u} u^{(m)f} \quad d^{(m)f} \rightarrow e^{i\phi_f^d} d^{(m)f} \quad (1.32)$$

would result in transforming the CKM matrix:

$$V_{\text{CKM}} \rightarrow U_{\phi^u} V_{\text{CKM}} U_{\phi^d} \quad (1.33)$$

where U_{ϕ^u} and U_{ϕ^d} are diagonal matrices of complex phases ($2n$ in total), but this transformation has no physical implication [4]. This invariance under reparameterization can be used to eliminate $2n - 1$ relative phases from V_{CKM} , leaving one global phase which does matter. The result is $n^2 - (2n - 1) = (n - 1)^2$ free real parameters. Since a unitary matrix is also an orthogonal matrix, $n(n - 1)/2$ of these can be identified as rotation angles between n independent bases, leaving $\frac{1}{2}(n - 1)(n - 2)$ complex phases. For the three dimensional CKM matrix, this means 3 rotation angles and 1 complex phase.

1.2.2. CP Violation

The presence of a complex phase (not possible with 2 generations) permits the possibility of \mathcal{CP} violation in the quark sector. This can be seen from the effect of parity, \mathcal{P} , and charge conjugation,

\mathcal{C} , transformations on fermion fields.

Parity, \mathcal{P} is the transformation which inverts the sign of spatial coordinates:

$$\vec{x} \rightarrow \mathcal{P} \vec{x} = -\vec{x} \quad (1.34)$$

The parity operator acting on a quantum state should produce a new state which is identical to the original apart from the reversal of spatial components. For a state of definite momentum \vec{p} , spin value s and charge q , we expect:

$$\mathcal{P} |\vec{p}, s, q\rangle = \eta_P |-\vec{p}, s, q\rangle. \quad (1.35)$$

where $\eta_P = \pm 1$. If the Dirac equation (see appendix A) is to be invariant under parity transformations, then Dirac spinors, which describe fermion fields, must transform as

$$\mathcal{P}^\dagger \psi(t, \vec{x}) \mathcal{P} = \eta_P \gamma_0 \psi(t, -\vec{x}), \quad (1.36)$$

i.e. some of the spinor components are rearranged by the γ_0 matrix.

Charge conjugation, \mathcal{C} , is the transformation which changes the sign of all quantum charges associated with a field. For a fermion field

$$\mathcal{C} \psi(t, \vec{x}) \mathcal{C}^{-1} = \eta_C \gamma_0 C \psi^*(t, \vec{x}) = \eta_C \widehat{\psi}(t, \vec{x}), \quad (1.37)$$

where η_C is a complex phase, and C is an antisymmetric matrix defined in terms of the Dirac gamma matrices by:

$$C^{-1} \gamma_\mu C = -\gamma_\mu^T \quad (1.38)$$

Its exact form depends on the chosen representation of the γ matrices.

Both \mathcal{C} and \mathcal{P} are maximally violated by the weak interaction. For example, in the weak decay $\pi^+ \rightarrow \mu^+ \nu_L$ the muons are always detected with left-handed helicity. The parity-reversed process $\pi^+ \rightarrow \mu^+ \nu_R$, which would produce muons with right-handed helicity, is not observed (because the W^\pm does not couple to right-chiral neutrinos). Similarly the charge conjugated process $\pi^- \rightarrow \mu^- \bar{\nu}_L$ is not observed (because the W^\pm does not couple to left-chiral anti-neutrinos). On the other hand, the $\mathcal{C}\mathcal{P}$ conjugated process $\pi^- \rightarrow \mu^- \bar{\nu}_R$ decay does occur. Moreover the $V - A$ structure of the weak interaction which leads to \mathcal{C} and \mathcal{P} violation naturally conserves $\mathcal{C}\mathcal{P}$.

Historically, the conservation of $\mathcal{C}\mathcal{P}$ appeared to restore a kind of natural left/right symmetry

to the Standard Model, after the discovery of separate charge and parity violation. However, CP violation was subsequently discovered in the neutral kaon mixing [5], and later in B mesons in the interference between mixing and decay[6]. CP violation is not an automatic consequence of the structure of the weak interaction but it can arise through the CKM matrix.

Consider a single Yukawa term (and its Hermitian conjugate) in the SM Lagrangian

$$\mathcal{L}_{\text{Yukawa}} = Y \bar{\psi}_L^A \Phi \psi_R^B + Y^* \bar{\psi}_R^B \Phi^\dagger \psi_L^A, \quad (1.39)$$

where the Yukawa matrix, Y, has been reduced to a single (possibly complex) number. The only way for \mathcal{L}_{Yuk} to be invariant under the operation of CP is if the first term transforms into the second and vice versa. Applying the definition of CP to the first term in 1.39 gives:

$$-Y \eta_P(\Phi) \eta_P^*(A) \eta_P(B) \eta_C(\Phi) \eta_C^*(A) \eta_C(B) \bar{\psi}_R^B \Phi^\dagger \psi_L^A \quad (1.40)$$

An irreducible complex phase in Y (meaning $Y^* \neq Y$), implies CP-violation because the term in 1.40 cannot be made to equal the second term in 1.39. This explains the significance of the fact that the three-dimensional CKM matrix does indeed contain an irreducible complex phase.[7]

1.2.3. Parametrization

The standard parameterization of the CKM matrix [8] defines it in terms of the angles θ_{12} , θ_{23} and θ_{13} and the complex phase δ_{CP} .

$$V_{CKM} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{pmatrix}, \quad (1.41)$$

where $c_x \equiv \cos \theta_x$ and $s_x \equiv \sin \theta_x$.

In nature, the absolute magnitudes of these numbers are seen to follow a hierarchical structure, in which off-diagonal elements are suppressed by the number of generations they skip [8]:

$$|V_{CKM}| = \begin{pmatrix} 0.97425 \pm 0.00022 & 0.2253 \pm 0.0008 & 0.00413 \pm 0.00049 \\ 0.225 \pm 0.008 & 0.986 \pm 0.016 & 0.0410 \pm 0.0013 \\ 0.0084 \pm 0.0006 & 0.040 \pm 0.0027 & 1.021 \pm 0.032 \end{pmatrix}. \quad (1.42)$$

It can be seen, for example, that $|\theta_{12}| \gg |\theta_{23}| \gg |\theta_{13}|$. Inspired by this, there is an alternative approximation given by the Wolfenstein parametrization, which makes the hierarchy more explicit:

$$V_{CKM} = \begin{pmatrix} 1 - \frac{1}{2}\lambda^2 & \lambda & A\lambda^3 \left[\rho - i\eta \left(1 - \frac{1}{2}\lambda^2 \right) \right] \\ -\lambda & 1 - \frac{1}{2}\lambda^2 - iA^2\lambda^4\eta & A\lambda^2 (1 + i\lambda^2\eta) \\ A\lambda^3 (1 - \rho - i\eta) & -A\lambda^2 & 1 \end{pmatrix} + \mathcal{O}(\lambda^4). \quad (1.43)$$

The parametrizations are connected by the relations:

$$s_{12} = \lambda = \frac{|V_{us}|}{\sqrt{|V_{ud}|^2 + |V_{us}|^2}} \quad (1.44)$$

$$s_{23} = A\lambda^2 = \lambda \left| \frac{V_{cb}}{V_{us}} \right| \quad (1.45)$$

$$s_{13} e^{i\delta_{CP}} = V_{ub}^* = A\lambda^3 (\rho + i\eta) \quad (1.46)$$

The location of the complex phase is not physically significant. Other parametrizations could be constructed with the complex terms in different locations without affecting the physics of the CKM matrix. The existence of CP violation only depends on there being a complex phase which cannot be removed by any rephasing. This can be expressed by writing out quantities which are invariant under quark rephasing. One such quantity is the Jarlskog invariant [9]:

$$J_{CP} \equiv |\mathcal{I}(V_{ij}V_{kl}V_{il}^*V_{kj}^*)| = s_{12}s_{13}s_{23}c_{12}c_{13}^2c_{23} \sin \delta_{CP} \simeq \lambda^6 A^2 \eta \quad (1.47)$$

where there is no sum over i, j, k, l . Since there is only one CP violating phase in the CKM matrix, any CP violation in the SM must be proportional to J_{CP} .

1.2.4. Unitarity triangles

An interesting example of a CP violating quantity comes from considering the unitarity condition on the CKM matrix:

$$(V^\dagger V)_{ij} = \sum_k V_{ik}^\dagger V_{kj} = \sum_k V_{ki}^* V_{kj} = \delta_{ij} \quad (1.48)$$

The three diagonal terms result in relations between real numbers

$$\sum_i |V_{ij}|^2 = 1 \quad (1.49)$$

which have no impact on CP violation. On the other hand, each of the off-diagonal elements produces a sum of three complex numbers which must total zero if the matrix is to be unitary.

$$\sum_k V_{ki}^* V_{kj} = 0 \quad i \neq j \quad (1.50)$$

These six relations can be viewed as triangles in the complex plane and are found to have area $|J_{CP}|/2$. Several of the triangles have sides with very different lengths so their angles are close to 0, $\pi/2$ or π . Two of the triangles have sides with similar lengths and correspondingly large angles:

$$V_{td}^* V_{ud} + V_{ts}^* V_{us} + V_{tb}^* V_{ub} = 0 \quad (1.51)$$

$$V_{ub}^* V_{ud} + V_{cb}^* V_{cd} + V_{td}^* V_{td} = 0 \quad (1.52)$$

In both of these triangles, the side lengths are all $\mathcal{O}(\lambda^3)$. The second triangle is commonly used to present the constraints on the CKM matrix coming from experiment. Normalizing so that one side has unit length, the equation becomes:

$$\frac{V_{ub}^* V_{ud}}{V_{cb}^* V_{cd}} + 1 + \frac{V_{td}^* V_{td}}{V_{cb}^* V_{cd}} = 0 \quad (1.53)$$

Slightly modified parameters $\bar{\rho}$ and $\bar{\eta}$ can be defined such that the relation

$$\bar{\rho} + i\bar{\eta} = -\frac{V_{ud} V_{ub}^*}{V_{cb}^* V_{cd}} \quad (1.54)$$

is independent of phase convention. They are related to the original parameters ρ and η by the relation

$$A\lambda^3(\rho + i\eta) = \frac{A\lambda^3(\bar{\rho} + i\bar{\eta})\sqrt{1 - A^2\lambda^4}}{\sqrt{1 - \lambda^2}[1 - A^2\lambda^4(\bar{\rho} + i\bar{\eta})]}. \quad (1.55)$$

Studying this triangle is a major part of experimental flavour physics; deviation from unitarity or discrepancies in the measured values of $\bar{\rho}$ and $\bar{\eta}$, when determined from different processes, would be a clear indication of New Physics, beyond the Standard Model.

The current experimental status of the unitarity triangle is summarised by the limits in the $\bar{\rho}$, $\bar{\eta}$

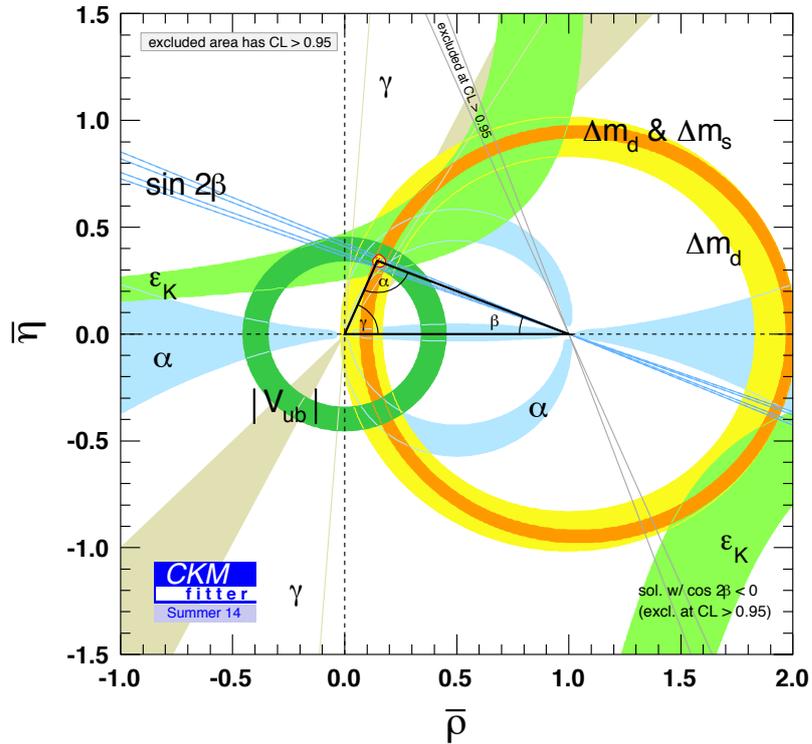


Figure 1.1.: Experimental constraints on the unitarity triangle of the CKM matrix. [10]

plane coming from measurements of both the side lengths and the angles, defined by:

$$\alpha = \arg\left(-\frac{V_{td}V_{tb}^*}{V_{ud}V_{ub}^*}\right) \quad (1.56)$$

$$\beta = \arg\left(-\frac{V_{cd}V_{cb}^*}{V_{td}V_{tb}^*}\right) \quad (1.57)$$

$$\gamma = \arg\left(-\frac{V_{ud}V_{ub}^*}{V_{cd}V_{cb}^*}\right) \quad (1.58)$$

These are shown in figure 1.1.

1.3. $K^+ \rightarrow \pi^+ \nu \bar{\nu}$

In the kaon sector, two rare decay modes, $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ and $K_L \rightarrow \pi^0 \nu \bar{\nu}$, play an especially important role in studying the CKM matrix. Their branching ratios are very low ($\mathcal{O}(10^{-11})$) and, crucially, they can be predicted very accurately in the Standard Model, providing clean measurements of the

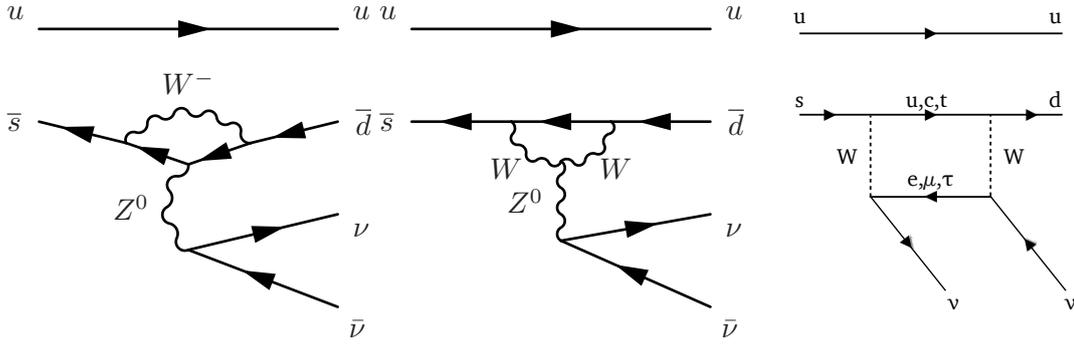


Figure 1.2.: Feynman diagrams for the 1 loop processes by which the decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ can proceed.

parameters $\bar{\rho}$ and $\bar{\eta}$.

1.3.1. Standard Model prediction

Both decays involve flavour changing neutral currents so, as shown in equation 1.30, they cannot proceed by tree level Feynman diagrams with a Z^0 boson. Figure 1.2 shows the mechanisms by which they can proceed: two ‘penguin’ diagrams and one ‘box’ diagram [11].

The GIM mechanism suppresses these decays beyond the factor of G_F^2 which would be expected for a loop diagram (the Fermi constant $G_F = \frac{\sqrt{2}}{8} \frac{g^2}{m_W^2}$, where g , is the weak coupling constant in equation 1.9). Summing over the quarks in the loop (as well as the flavour of the final state neutrinos since these are not distinguished) the amplitudes must have the form [12]

$$A(K \rightarrow \pi \nu \bar{\nu}) \sim \sum_{e,\mu,\tau} \sum_{f=u,c,t} X_f V_{fd} V_{fs}^*, \quad (1.59)$$

where V_{ij} are CKM matrix elements and X_f is a function of quark mass. By the unitarity of the CKM matrix, this amplitude vanishes in the case that there is complete degeneracy between the up and down type quarks. The amplitude is only non-zero because X_f is a function of quark mass. Moreover it can be shown that $X_f \propto m_f^2/m_W^2$ (this is a so-called power-like GIM mechanism), so the process is dominated by the top quark contribution. This makes these decays suitable for studying the V_{td} element which is hard to probe directly.

Studies of long distance contributions from hadronic degrees of freedom and photon exchange have shown that both contributions are small [13–15]¹, so the low-energy effective Hamiltonian for

¹small for the charm quark loops, negligible for the top quark.

the K^+ decay can be written in terms of a single operator: [16, 17]

$$\mathcal{H}_{\text{eff}} = \frac{4G_F}{\sqrt{2}} \frac{\alpha}{2\pi \sin^2 \theta_W} \sum_{l=e,\mu,\tau} (\lambda_c X_c^l + \lambda_t X_t^l) (\bar{s}_L \gamma_\mu d_L) \times (\bar{\nu}_{lL} \gamma^\mu \nu_{lL}) + h.c. \quad (1.60)$$

Here $\lambda_x = V_{is}^* V_{id}$ are the CKM factors, and X_c^l, X_t^l are the contributions from the charm and top loops, which are expansions in powers of $x_q = m_q^2/m_W^2$. The dependence on the lepton mass comes from the box digram contribution and is negligible for the top quark loop so one can write $X_t^l = X_t$. The branching ratio can be extracted from this effective Hamiltonian and is found to be: [18]

$$\mathcal{B}(K^+ \rightarrow \pi^+ \nu \bar{\nu}) = \kappa_+ (1 + \Delta_{\text{EM}}) \left[\left(\frac{\Im \lambda_t}{\lambda^5} X_t \right)^2 + \left(\frac{\Re \lambda_c}{\lambda} (P_c + \delta P_{c,u}) + \frac{\Re \lambda_t}{\lambda^5} X_t \right)^2 \right]. \quad (1.61)$$

Here λ is the Wolfenstein parameter ($\approx |V_{us}|$) and P_c contains the short distance charm contribution from both the electron and tau loops:

$$P_c = \frac{1}{\lambda^4} \left(\frac{2}{3} X_c^e + \frac{1}{3} X_c^\tau \right). \quad (1.62)$$

P_c has been computed at NLO order [17] and partially at NNLO [19] and with 1 loop electroweak corrections [20]. $\delta P_{c,u}$ is the correction for long distance contributions to the charm quark loops, found to be $\delta P_{c,u} = 0.04 \pm 0.02$ [21] and $\Delta_{\text{EM}} = -0.003$ is the long distance QED correction [22]. The function X_t has been evaluated with QCD NLO corrections [16] and two-loop electroweak corrections [18].

κ_+ contains the hadronic matrix element along with electroweak corrections. The hadronic matrix element can be extracted from the well measured semi-leptonic decays $K^+ \rightarrow \pi^0 l^+ \nu$,

$$\kappa_+ = r_{K^+} \frac{3\alpha^2 \mathcal{B}(K^+ \rightarrow \pi^0 e^+ \nu)}{2\pi^2 \sin^4 \theta_W} \lambda^8 \quad (1.63)$$

taking into account isospin breaking corrections (given by $r_{K^+} = 0.901$) [23].

The predicted value in the Standard Model is [18]

$$\mathcal{B}(K^+ \rightarrow \pi^+ \nu \bar{\nu}) = (7.81_{-0.71}^{+0.80} \pm 0.29) \times 10^{-11}, \quad (1.64)$$

where the first error comes from uncertainties in the input parameters (dominated by V_{cb} and $\bar{\rho}$) and the second comes from the theoretical error (dominated by $\delta P_{c,u}$ and QCD corrections to X_t).

The neutral mode is CP-violating so only the purely imaginary components of the CKM matrix contribute:

$$\mathcal{B}(K_L \rightarrow \pi^0 \nu \bar{\nu}) = \kappa_L \left(\frac{\text{Im} \lambda_t}{\lambda^5} X_t \right)^2, \quad (1.65)$$

where

$$\kappa_L = r_{K_L} \frac{3\alpha^2 \mathcal{B}(K_L \rightarrow \pi^\pm e^\mp \nu)}{2\pi^2 \sin^4 \theta_W} \lambda^8. \quad (1.66)$$

The Standard Model prediction is

$$\mathcal{B}(K_L \rightarrow \pi^0 \nu \bar{\nu}) = (2.43_{-0.37}^{+0.80} \pm 0.06) \times 10^{-11}. \quad (1.67)$$

Here the parametric error is dominated by V_{cb} and the theoretical uncertainty is dominated by the QCD contribution to X_t .

1.3.2. Impact on the unitarity triangle

Measured values of branching ratios can be put back into the Standard Model in order to help constrain the CKM matrix and unitarity triangle. In the Wolfenstein parametrization, the charged decay mode branching ratio can be expressed:

$$\mathcal{B}(K^+ \rightarrow \pi^+ \nu \bar{\nu}) = \bar{\kappa}_+ |V_{cb}|^4 X_t^2 \frac{1}{\sigma} [(\sigma \bar{\eta})^2 + (\rho_c - \bar{\rho})^2], \quad (1.68)$$

where

$$\bar{\kappa} = \frac{\kappa}{\lambda^8}, \quad \sigma = \left(\frac{1}{1 - (\lambda^2/2)} \right)^2 \quad \text{and} \quad \rho_c = 1 + \frac{\lambda^4 P_c}{|V_{cb}|^2 X_t}. \quad (1.69)$$

This means that a measured value of $\mathcal{B}(K^+ \rightarrow \pi^+ \nu \bar{\nu})$ determines an ellipse in the $(\bar{\rho}, \bar{\eta})$ plane, centred at $(\rho_c, 0)$. The fact that $\rho_c \neq 1$ is due to the charm contribution to the decay.

Finally, a combination of measurements of the two modes can lead to a measurement of $\sin 2\beta$, in which many of the parametric uncertainties cancel. Defining the reduced branching ratios

$$B_1 = \frac{\mathcal{B}(K^+ \rightarrow \pi^+ \nu \bar{\nu})}{\kappa_+} \quad B_2 = \frac{\mathcal{B}(K_L \rightarrow \pi^0 \nu \bar{\nu})}{\kappa_L} \quad (1.70)$$

it can be shown that

$$\cot \beta = \sqrt{\sigma} \frac{\sqrt{\sigma(B_1 - B_2)} - P_c(X)}{\sqrt{B_2}}, \quad (1.71)$$

which depends only on the two measurable branching ratios and the parameter P_c , so it is free from

the hadronic uncertainties in X_t . A measurement of $\sin 2\beta$, independent of the measurements from B physics, provides a strong test of the SM and minimal flavour violation (MFV) models, in which the structure of CP-violating transitions is still governed by the CKM matrix. [24]. The value of $\sin 2\beta$ has been measured with $\sim 5\%$ precision at LCHb [25], while the kaon sector measurement could yield $\sim 10\%$ precision with the current generation of experiments and $\sim 5\%$ precision in the longer term [11].

The only observations of $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ so far come from the E787 and E949 experiments at the Brookhaven National Laboratory. The experiments are described in chapter 4 and the final result was [26]

$$\mathcal{B}(K^+ \rightarrow \pi^+ \nu \bar{\nu}) = 1.73^{+1.15}_{-1.05} \times 10^{-10} \quad (1.72)$$

So far there have been no observations of the decay $K^0 \rightarrow \pi^0 \nu \bar{\nu}$. An upper limit was set by the E931a experiment at the KEK 12 GeV proton synchrotron [27]

$$\mathcal{B}(K^0 \rightarrow \pi^0 \nu \bar{\nu}) < 2.6 \times 10^{-8} \quad (90\% \text{ C.L.}) \quad (1.73)$$

The current generation of rare kaon decay experiments is expected to set limits on the branching ratios of $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ and $K^0 \rightarrow \pi^0 \nu \bar{\nu}$ which are comparable with the theoretical uncertainty. NA62 expects to measure $\mathcal{B}(K^+ \rightarrow \pi^+ \nu \bar{\nu})$ with 10% precision [28] while *KOTO* aims to measure $\mathcal{B}(K^0 \rightarrow \pi^0 \nu \bar{\nu})$ to within a few percent [29]. Figure 1.3 shows how the $\bar{\rho}, \bar{\eta}$ plane will be constrained if each of these experiments measures the expected SM value with 10% precision [30].

1.4. Neutrinos in the Standard Model

1.4.1. Massless Origins

When neutrinos were first proposed as a “desperate remedy” for the problem of missing energy in β decay [31] it was not essential that they be massless; their masses needed only to be much smaller than that of the electron in order to be compatible with the observed spectrum of electrons energies. By the time the Standard Model was formalized in the 1970s, there was no evidence for non-zero neutrino mass. Anti-neutrinos from beta-decay were always observed in states of positive helicity, consistent with massless right-chiral particles (and vice-versa for anti-neutrinos) and there was no evidence for the decay $\mu^- \rightarrow e^- \gamma$, which implied conservation of lepton flavour number.

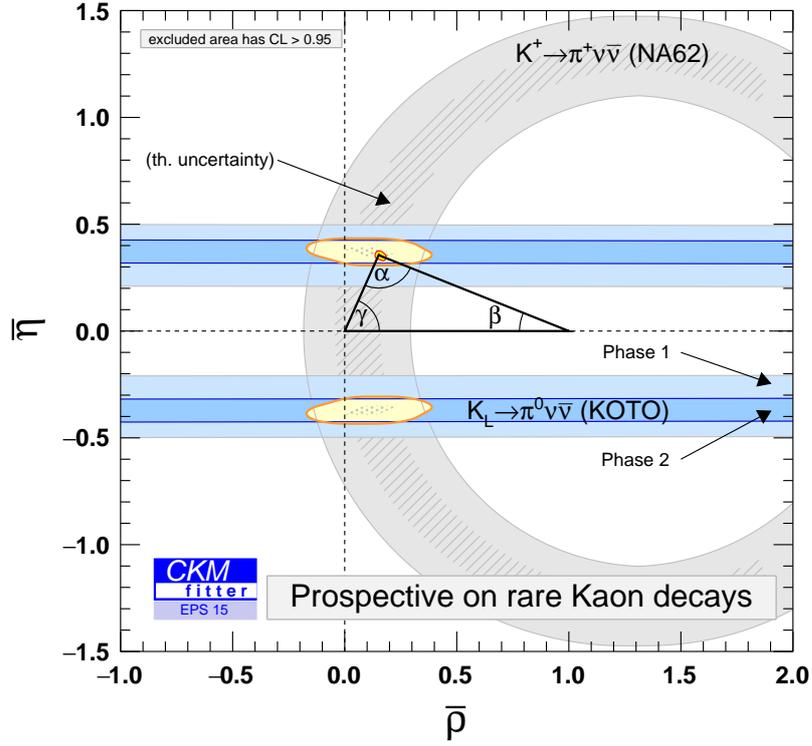


Figure 1.3.: The expected constraints on the $\bar{\rho}$, $\bar{\eta}$ plane from the NA62 and KOTO experiments.[10]

In light of this, the Standard Model is defined without introducing right-chiral neutrino fields. As a result, there is no opportunity for the Higgs mechanism to imbue neutrinos with mass, and they remain massless to all orders. In this limit, chirality and helicity coincide and are both Lorentz invariant and eigenstates of the SM Hamiltonian. This means neutrinos have half as many degrees of freedom as other SM particles.

It remains true that neutrino masses have never been directly measured kinematically. The most sensitive limits come from studying the endpoint region of β -decay spectra. By studying tritium β -decay, the Mainz experiment set an upper limit $m_{\nu_e} < 2.8$ eV, 95% CL in 1999 [32], and in 2011 the Troitsk experiment improved this limit to $m_{\nu_e} < 2.05$ eV, 95% CL [33]. The KATRIN experiment is expected to improve this limit by about an order of magnitude [34], and will begin data taking in 2016 [35]. However, there is now a wealth of evidence for the phenomenon of neutrino oscillations, in which neutrinos produced in one flavour eigenstate (ν_e, ν_μ, ν_τ), may later be detected in another flavour eigenstate [36].

1.4.2. Neutrino Oscillations

As early as the 1960s, there was evidence that the description of neutrinos in the Standard Model was incomplete. The Homestake experiment [37] was designed to measure the flux of neutrinos coming from the sun using inverse β decay:



It measured an (electron) neutrino flux which was only $(27 \pm 4)\%$ of that predicted by models of solar nuclear processes. This was called the *Solar Neutrino problem* [38]. The solution, proposed by Pontecorvo in 1968 [39], was that the electron neutrinos, produced by nuclear reactions in the sun, were oscillating into other flavours on their journey to earth, so they were not detectable by the Homestake experiment. Later experiments confirmed this picture by detecting muon neutrinos coming from the sun. Super-Kamiokande was sensitive to the combined flux of ν_e , ν_μ and ν_τ neutrinos [40], while the SNO experiment was sensitive separately to the ν_e flux and the total neutrino flux [41] and the two experiments produced results consistent with ν_e oscillations into other flavours.

The theory of neutrino oscillations is built on the assumption of neutrino mass so that there is a mixing matrix between neutrino mass states (ν_1, ν_2, ν_3) and flavour states (ν_e, ν_μ, ν_τ): the leptonic equivalent of the CKM matrix. This matrix is called the *Pontecorvo-Maki-Nakagawa-Sakata (PMNS)* matrix. It is defined

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (1.75)$$

and can be parameterized by three angles and a complex phase in the same way as the CKM matrix [8]

$$U_{\text{PMNS}} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{pmatrix}. \quad (1.76)$$

Neutrino oscillation occurs because, although neutrinos are produced by the weak interaction in flavour eigenstates, it is the mass states which propagate as physical particles through space. The rate at which the phase of a neutrino wavefunction progresses as it propagates depends on the mass

of the neutrino so the phases of the different mass states progress at different rates. As a result, after a period of propagation, there is a non-zero probability for a neutrino to be detected in a flavour state other than its original one. In the general case of n neutrino mass states and flavour states, the probability for a neutrino with energy, E , and flavour, α , to be detected with flavour β after travelling a distance L , is given by:

$$\begin{aligned}
 P(\nu_\alpha \rightarrow \nu_\beta)(t) = & \delta_{\alpha\beta} - 4 \sum_{i>j} \Re(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E} \right) \\
 & + 2 \sum_{i>j} \Im(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E} \right)
 \end{aligned} \tag{1.77}$$

With three neutrinos, there are three independent angles θ_{12} , θ_{23} , θ_{13} and two independent squared mass differences, m_{12}^2 and m_{13}^2 . Since the combination $(\Delta m_{ij}^2 L/E)$ enters the probability for mixing, the three mass differences and corresponding mixing angles can be measured in experiments with different combinations of neutrino energy and propagation distance.

θ_{13} is determined by studying the disappearance of anti electron neutrinos produced in nuclear reactors, after they have travelled distances around 1 km (short baseline experiments) [42]. Long baseline experiments in which muon neutrinos produced in accelerators travel hundreds of km are sensitive both to θ_{23} through ν_μ disappearance (oscillating into ν_τ) and θ_{13} through ν_e appearance. [43]. Atmospheric neutrinos can be used to study θ_{23} via $\nu_\mu \rightarrow \nu_\tau$ (and c.c.) oscillations measured as function of zenith angle. The angle θ_{12} is determined from solar oscillations. Matter effects complicate the pattern of neutrino oscillations, since the three flavours ν_e , ν_μ and ν_τ have different interactions with electrons (the MSW effect [44]): ν_e can interact via both the charged current and the neutral current, while ν_μ and ν_τ only interact via the neutral current. Nonetheless, experiments so far have successfully placed constraints on all the parameters describing neutrino oscillation [45]. Some tension remains when attempting to fit all the experimental observations into a 3 flavour model of neutrino mixing [46] but the existence of neutrino mass is now beyond doubt.

1.5. Theories of neutrino mass

In order to accommodate neutrino mass, new terms must be added to the SM lagrangian. In contrast to the quark sector, there is more freedom in how the SM should be modified for neutrinos, so there is not yet a single theory of neutrino mass.

1.5.1. Dirac Mass

The simplest approach to generating neutrino masses is to introduce right-handed neutral fields, N_{lR} , analogous to the right handed quarks [47–49]². Since they are right-handed, N_{lR} are $SU(2)_L$ singlet states, and they have no charge so hypercharge $Y = 0$. This means they have no interaction with any of the gauge-bosons. One can now generate neutrino mass by introducing Yukawa couplings:

$$-\mathcal{L}_{\text{Yuk}} = \sum_{l,l'} Y_{ll'} \bar{\chi}_{lL} \hat{\Phi} N_{l'R} + h.c. \quad (1.78)$$

where $Y_{ll'}$ is the mixing matrix which permits the physical mass states to be misaligned with respect to the weak states. After spontaneous symmetry breaking, the part of the Lagrangian responsible for neutrino mass looks like

$$-\mathcal{L}_{\text{mass}} = \sum_{l,l'} Y_{ll'} \frac{v}{\sqrt{2}} \bar{\nu}_{lL} N_{l'R} + h.c., \quad (1.79)$$

which contains the mass matrix:

$$M_{ll'} = \frac{v}{\sqrt{2}} Y_{ll'} \quad (1.80)$$

As in the quark sector, the neutrino masses can be obtained by diagonalizing the mass matrix with a biunitary transformation and defining the physical neutrino mass states ν_α .

$$U^\dagger M V = m, \quad \nu_{lL} \equiv \sum_\alpha U_{l\alpha} \nu_{\alpha L}, \quad N_{lR} \equiv \sum_\alpha V_{l\alpha} \nu_{\alpha R} \quad (1.81)$$

In terms of the propagating mass states, the Lagrangian becomes

$$-\mathcal{L}_{\text{mass}} = \sum_\alpha \bar{\nu}_{\alpha L} m_\alpha \nu_{\alpha R} + h.c. \quad (1.82)$$

where m_α are the three neutrino masses. This model successfully introduces neutrino mass into the SM Lagrangian but at the cost of another matrix with apparently arbitrary parameters. These must be fixed from experimental data as there is no prediction from this theory. In particular, to be consistent with the observed limits on neutrino mass, the coefficients of $Y_{ll'}$ must be several orders of magnitude lower than the corresponding coefficients of the quark Yukawa couplings and there is no explanation for this new scale.

²The following sections are largely drawn from these sources.

1.5.2. See-saw mechanisms

More interesting models can be built by using the new singlet fields to introduce additional terms into the Lagrangian. Defining the Lorentz covariate conjugate field

$$\widehat{N}_L = \widehat{(N_R)} = \gamma_0 C N_R^*, \quad (1.83)$$

the bare mass term

$$-\mathcal{L}_{\text{bare}} = \frac{1}{2} \sum_{l,l'} B_{ll'} \widehat{N}_{lL} N_{l'R} + h.c. \quad (1.84)$$

is invariant under $SU(2)_L \times U(1)_Y$ ³. It does violate the $B - L$ (baryon number – lepton number) symmetry, which has been a global symmetry of the SM so far. However, conservation of $B - L$ is an accidental symmetry of the Standard Model - it does not arise from any gauge symmetry so there is no difficulty in introducing this kind of term. It is called a Majorana mass term.

Using the identity

$$\bar{\nu}_{lL} N_{l'R} = \widehat{N}_{l'L} \widehat{\nu}_{lR} \quad (1.85)$$

to rewrite the first term in equation 1.79, the combination of Dirac and Majorana mass terms can be written as the sum of equations 1.79 and 1.84 in matrix form

$$-\mathcal{L}_{\text{mass}} = \frac{1}{2} \begin{pmatrix} \bar{\nu}_L & \widehat{N}_L \end{pmatrix} \begin{pmatrix} 0 & M \\ M^T & B \end{pmatrix} \begin{pmatrix} \widehat{\nu}_R \\ N_R \end{pmatrix} + h.c. \quad (1.86)$$

$$= \frac{1}{2} \begin{pmatrix} \bar{\nu}_L & \widehat{N}_L \end{pmatrix} \mathcal{M}_{\text{mass}} \begin{pmatrix} \widehat{\nu}_R \\ N_R \end{pmatrix} \quad (1.87)$$

where M is the mass matrix from equation 1.79 and B is the matrix from equation 1.84. Both are $n \times n$ matrices, where n is the number of lepton generations.

The effect of the new mass term can be seen in the case $n = 1$, where now M and B are simply numbers. The usual biunitary transformation is required to find the mass eigenstates. In this case,

³ No corresponding terms could be written for any of the other fermions in the SM. For example $\widehat{\nu}_{lR} \nu_{l'L}$ would violate the $U(1)_Y$ symmetry.

choosing the orthogonal matrix:

$$O = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}, \quad \tan 2\theta = 2M/B \quad (1.88)$$

and an additional diagonal matrix

$$K^2 = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \quad (1.89)$$

the mixing matrix can be written as

$$\mathcal{M}_{\text{mass}} = O^T m K^2 O, \quad (1.90)$$

where $m = \text{diag}(m_1, m_2)$ contains the masses of physical states. These are given by

$$m_{1,2} = \frac{1}{2} \left| \sqrt{B^2 + 4M^2} \mp B \right|. \quad (1.91)$$

The mass eigenstates are then:

$$\begin{pmatrix} n_{1L} \\ n_{2L} \end{pmatrix} = O \begin{pmatrix} \nu_L \\ \widehat{N}_L \end{pmatrix} \quad \begin{pmatrix} n_{1R} \\ n_{2R} \end{pmatrix} = OK^2 \begin{pmatrix} \widehat{\nu}_R \\ N_R \end{pmatrix}. \quad (1.92)$$

and the mass part of the Lagrangian becomes

$$-\mathcal{L}_{\text{mass}} = m_1 \bar{n}_{1L} n_{1R} + m_2 \bar{n}_{2L} n_{2R} + h.c. \quad (1.93)$$

constructed from the two physical particles

$$n_1 = n_{1L} + n_{1R} = \cos \theta (\nu_L - \widehat{\nu}_R) - \sin \theta (\widehat{N}_L - N_R) \quad (1.94)$$

$$n_2 = n_{12} + n_{2R} = \sin \theta (\nu_L + \widehat{\nu}_R) + \cos \theta (\widehat{N}_L + N_R). \quad (1.95)$$

These expansions in terms of the weak states show that

$$n_1 = -\widehat{n}_1, \quad n_2 = \widehat{n}_2. \quad (1.96)$$

The expression $\psi = \lambda \widehat{\psi}$, where λ is some complex phase, is the definition of a *Majorana particle*.

A Majorana particle is its own anti-particle so a Majorana spinor field has half as many degrees of freedom as a Dirac spinor. In the *Majorana basis* (in which all the Dirac matrices are purely imaginary), the additional constraint on a Majorana spinor is expressed by the condition $\psi = \psi^*$, i.e. a Majorana particle is represented by a real field.

In summary, there were originally four weak fields (2-component Weyl spinors), ν_L and $\widehat{\nu}_L$ which couple to the weak currents W^\pm and Z^0 , along with the N_R and \widehat{N}_R , which have no gauge couplings (sometimes called *sterile* neutrino fields). If the Lagrangian had contained only Dirac terms, the neutrino would have been a Dirac particle, whose particle and anti-particle states ($\nu = \nu_L + \nu_R$, $\widehat{\nu} = \widehat{\nu}_L + \widehat{\nu}_R$) had identical masses. Introducing the Majorana mass terms produces Majorana neutrinos, in which there are two states with different masses (n_1 and n_2 in equation 1.95), and each mass state is its own anti-particle.

If the mixing matrix had produced eigenvalues with equal magnitudes (i.e. by setting $B = 0$), n_1 and n_2 could have been combined into a single Dirac particle, which requires twice as many degrees of freedom as a Majorana particle. This is also a feature of more complex models. In general, diagonalization of a mass matrix in n generations produces $2n$ Majorana fermions, but for every pair of equal and opposite eigenvalues, a single Dirac particle is produced instead of two Majorana particles.

The term M in the mass matrix ($\mathcal{M}_{\text{mass}}$ in equation 1.87) comes from the Yukawa coupling so it is expected to be of the same order as the other fermions in the same generation. The term B does not have to be tied to the same scale. If $B \gg M$, then equation 1.91 becomes approximately

$$m_1 \approx \frac{M^2}{B} \quad , \quad m_2 \approx B. \quad (1.97)$$

In this way, the existence of a very heavy neutrino can explain the lightness of the SM neutrino. This is called the *see-saw mechanism*. The problem of explaining the smallness of the light neutrino masses has now been shifted to the problem of explaining the largeness of the mass scale generated by B in the mixing matrix. However, in the context of Grand Unified Theory in which the SM model is the low energy effective theory, the see-saw mechanism can arise naturally to explain the small neutrino masses as a result of new physics at a higher energy scale.

The mechanism described above is the *Type I* see-saw mechanism [50]. The Type I see-saw mechanism can be generalized by replacing the zero matrix in equation 1.87 to obtain a more general

symmetric mixing matrix:

$$-\mathcal{L}_{\text{mass}} = \frac{1}{2} \begin{pmatrix} \overline{\nu}_L & \widehat{\overline{N}}_L \end{pmatrix} \begin{pmatrix} M_L & M_D \\ M_D^T & B \end{pmatrix} \begin{pmatrix} \widehat{\nu}_R \\ N_R \end{pmatrix} + h.c. \quad (1.98)$$

Here M_L is the new matrix connecting $\overline{\nu}_L$ and $\widehat{\nu}_R$ and M has been renamed M_D . If the contribution from M_L dominates that of B , the model is said to involve a *Type II* see-saw mechanism [51]. (If neither type dominates, the model is *hybrid* or *mixed*.) Type II see-saw mechanisms can arise when, for example, the model includes an additional right-chiral symmetry at high energies: $SU(2)_L \times SU(2)_R \times U(1)_Y$. Beyond right-handed fermions, the see-saw effect can appear in other scenarios, such as with an expanded Higgs sector or in models with spontaneous $B - L$ symmetry breaking (which predict the associated Majoron particle) [47].

Returning to a model with 3 lepton generations, the complete Lagrangian must allow for mixing of the neutrino flavour states. The Lagrangian describing neutrino mass (shown here after spontaneous symmetry breaking) is then [49]:

$$-\mathcal{L}_{\text{mass}} = M_D^{\alpha\beta} \overline{\nu}_L^\alpha \nu_R^\beta + \frac{1}{2} B^{\alpha\beta} \widehat{\overline{\nu}}_R^\alpha \nu_R^\beta + \frac{1}{2} M_L^{\alpha\beta} \overline{\nu}_L^\alpha \widehat{\nu}_L^\beta + h.c., \quad (1.99)$$

where α and β label the flavour eigenstates e , μ and τ . The matrix formulation, analogous to equation 1.87 is

$$\mathcal{L}_{\text{mass}} = \frac{1}{2} \begin{pmatrix} \overline{\nu}_L & \widehat{\overline{\nu}}_R \end{pmatrix} \begin{pmatrix} M_L & M_D \\ M_D^T & B \end{pmatrix} \begin{pmatrix} \widehat{\nu}_L \\ \nu_R \end{pmatrix} + h.c. \quad (1.100)$$

$$= \frac{1}{2} \begin{pmatrix} \overline{\nu}_L & \widehat{\overline{\nu}}_R \end{pmatrix} \mathcal{M}_{\text{mass}} \begin{pmatrix} \widehat{\nu}_L \\ \nu_R \end{pmatrix} + h.c., \quad (1.101)$$

where $\mathcal{M}_{\text{mass}}$ is now a $2n \times 2n$ matrix and

$$\nu_L = \begin{pmatrix} \nu_L^e \\ \nu_L^\mu \\ \nu_L^\tau \end{pmatrix}, \quad \widehat{\nu}_R = \begin{pmatrix} \widehat{\nu}_R^e \\ \widehat{\nu}_R^\mu \\ \widehat{\nu}_R^\tau \end{pmatrix} \quad (1.102)$$

Diagonalizing $\mathcal{M}_{\text{mass}}$ as in the 1 generation case, the following relation between the neutrino

flavour states (ν_l) and mass states (n_k), via the $2n \times 2n$ mixing matrix U , can be obtained:

$$\nu_{lL} = \sum_{k=1}^{2n} U_{lk} n_{kL}, \quad \widehat{\nu}_{lR} = \sum_{k=1}^{2n} U_{\bar{l}k} n_{kL} \quad (1.103)$$

Here, l labels the n flavour states (ν_e, ν_μ, ν_τ) and k labels $2n$ mass states (ν_1, ν_2, \dots). The index \bar{l} labels the n lower rows of the mixing matrix U . In this model, each of the three flavour states, ν_e, ν_μ and ν_τ , is a superposition of the six Majorana mass states.

1.5.3. ν MSM

The ν MSM (Neutrino Minimal Standard Model) is the simplest possible renormalizable extension to the Standard Model [52, 53], including only the 3 right-handed neutrinos needed to introduce the observed neutrino masses via the type I see-saw mechanism. The complete Lagrangian is given by

$$\mathcal{L}_{\nu\text{MSM}} = \mathcal{L}_{\text{SM}} + \overline{N}_R^\alpha i \partial_\mu \gamma^\mu N_R^\alpha + -Y^{\alpha\beta} \overline{\chi}_L^\alpha N_R^\beta \Phi - \frac{B^\alpha}{2} \overline{N}_R^\alpha N_R^\alpha + h.c. \quad (1.104)$$

which contains 18 new parameters beyond those in SM. It has been shown that suitable choices of those parameters can explain dark matter and the baryon asymmetry of the universe while remaining consistent with observed neutrino oscillations. The heavy neutrino masses are chosen to be at or below the electroweak scale, which is a motivation for the search for heavy neutrinos in kaon decay [54].

1.6. Experimental searches for heavy neutrinos

If the masses of the SM neutrinos are determined through a mechanism like the ones described above then it might be possible to find evidence of the heavy neutrino mass states in addition to the three light SM states.

A weak decay with one or more neutrinos in the final state will produce them in a state with definite flavour ($\nu_l = \nu_e, \nu_\mu, \nu_\tau$) corresponding to the associated charged lepton produced. The flavour is uniquely determined by the decay, e.g. the neutrino produced in $K^+ \rightarrow \mu^+ \nu$ must always be ν_μ . A subsequent measurement of the neutrino's mass will produce one of the SM neutrino masses, $m_{1,2,3}$, with probabilities determined by $U_{1,2,3}$ in the PMNS mixing matrix (equation 1.75). If there is an additional heavy neutrino mass state, ν_h , which is kinematically accessible in the decay under

consideration, then there will be a possibility to measure m_h with probability determined by the element U_{lh} in the extended mixing matrix defined in equation 1.104.

For heavy neutrinos with masses in the range $10 \text{ MeV}/c^2 \lesssim m_h \lesssim m_K$, limits can be placed on the mixing of heavy neutrinos with $\nu_e \nu_\mu$ by looking for peaks in the missing mass spectrum two-body decays of pions and kaons [55]. Few assumptions need to be made on the decay properties of the heavy neutrinos so the bounds from “peak searches” are typically robust.

An alternative approach is to look for the decay products of heavy neutrinos [56]. In the framework described above, heavy neutrinos would decay via charged current and neutral current interactions into neutrinos and other particles such as pions, muons and electrons which could be detected. “Decay searches” are typically performed in beam dump experiments, where heavy neutrinos are potentially produced in meson decays. The detector can be placed a long way from the point of production with a large amount of shielding in between so that the search for neutrino decay products is effectively background-free. In turn this means that very large fluxes can be used because the rate is not limited by the detector. Limits from decay searches are typically stronger than those from peak searches but less robust because they depend on assumptions about how heavy neutrinos decay. If some fraction of decays are actually into invisible particles, then the bounds are weakened.

Another form of decay search involves lepton number violating decays ($\Delta L = 2$), mediated by an on-shell or virtual Majorana neutrino. As in the beam dump experiments, limits depend on the assumed lifetime of the intermediate neutrino. A search for the decay $B^- \rightarrow \pi^+ \mu^- \mu^-$ was performed at LHCb and set limits on $|U_{\mu 4}|^2$ for lifetimes up to $\approx 1000 \text{ ps}$ [57] and for heavy neutrino masses between $250 \text{ MeV}/c^2$ and $5 \text{ GeV}/c^2$. The lowest upper limit was around 10^{-4} at $2 \text{ GeV}/c^2$. Similar searches for the decay $K^+ \rightarrow \pi^- \mu^+ \mu^+$ lead to limits of around 10^{-5} for neutrino masses in the range $250 \text{ MeV}/c^2 - 380 \text{ MeV}/c^2$.

Figures 1.4 and 1.5 show the current limits on $|U_{eh}|^2$ and $|U_{\mu h}|^2$ for heavy neutrino masses between $100 \text{ MeV}/c^2$ and $100 \text{ GeV}/c^2$ for decay and peak searches. The goal of the analysis presented in this thesis will be to extend the region of masses for which limits can be placed on $|U_{\mu h}|^2$ using peak searches in the kaon decay $K^+ \rightarrow \mu^+ \nu_\mu$. A similar analysis of the decay $K^+ \rightarrow e^+ \nu_e$ could be used to set limits on $|U_{e2}^2|$ but this decay is not studied in this thesis.

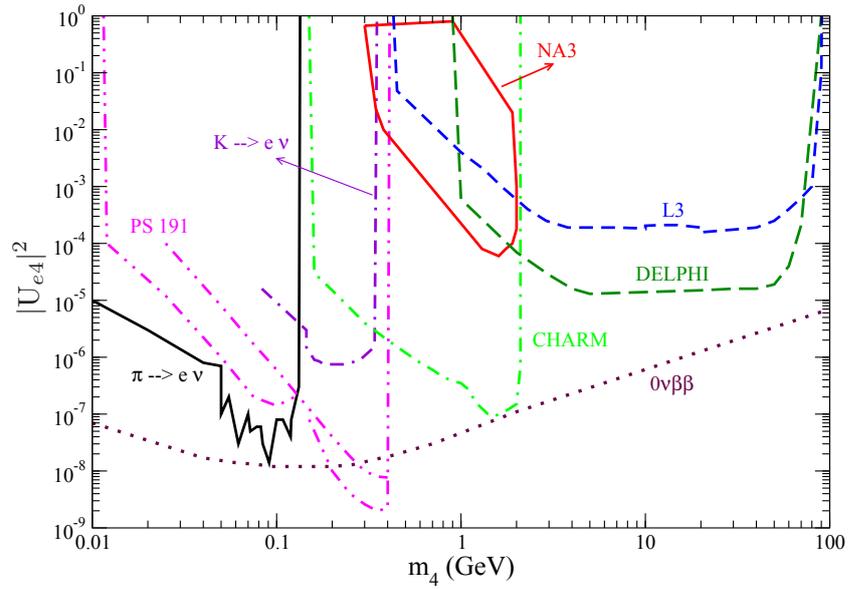


Figure 1.4.: Limits on $|U_{eh}|$ versus m_h in the mass range $10 \text{ MeV}/c^2 - 100 \text{ GeV}/c^2$ from peak searches and ν_h decays [56]. The limits labelled $\pi \rightarrow e \nu$ and $K \rightarrow e \nu$ are from peak searches at TRIUMPH [58] and KEK [59]. The beam dump experiments PS191 [60], CHARM [61], and NA3 [62] set limits by looking for neutrino decay products. The L3 [63], and DELPHI [64] experiments at LEP, looked for heavy neutrino production and subsequent decay in Z^0 decays.

1.6.1. Peak searches in kaon decay

Limits on $|U_{\mu h}|^2$ from peak searches in kaon decay come from the two-body decay

$$K^+ \rightarrow \mu^+ \nu_\mu. \quad (1.105)$$

The existing peak search limits on heavy neutrinos come from experiments using stopped kaons. In 1982, a dedicated experiment was performed at the KEK laboratory, searching for heavy neutrinos in the decays of stopped kaons [65]. The momentum spectrum of muons from the decay $K^+ \rightarrow \mu^+ \nu_\mu$ was measured with a high-resolution magnetic spectrograph. In total, 2.54×10^6 candidate decays were collected and the experiment was able to set limits for neutrino mass between $70 \text{ MeV}/c^2$ and $320 \text{ MeV}/c^2$.

In early 2015, a new limit was set with data from the E949 experiment at BNL [70]. This experiment was built to search for the decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ using stopped kaons and took data in 2002. A reanalysis of these data from 1.70×10^{12} stopped kaons was performed in order to set limits for neutrino masses

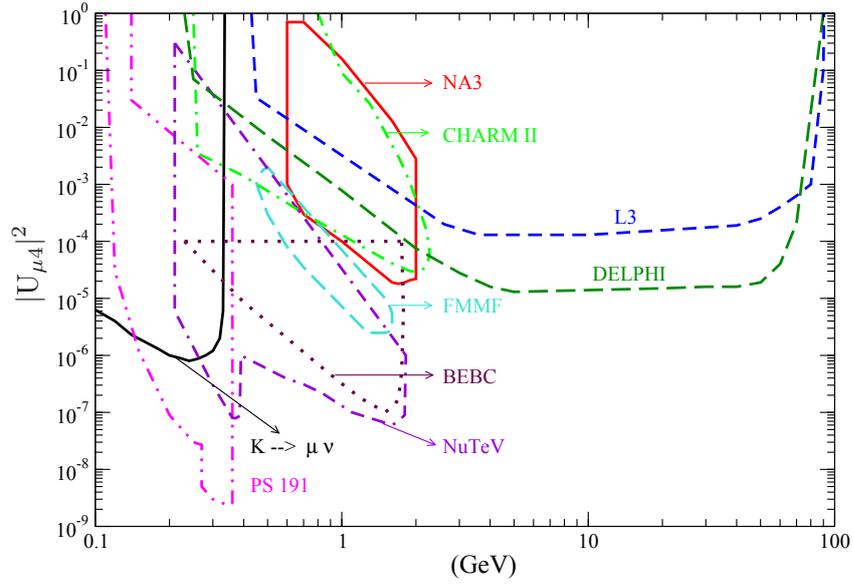


Figure 1.5.: Limits on $|U_{\mu h}|$ versus m_h in the mass range $100 \text{ MeV}/c^2 - 100 \text{ GeV}/c^2$ from peak searches and ν_h decays [56]. The limits marked $K \rightarrow \mu \nu$ come from peak searches at KEK [65], PS191 [60], NuTeV [66], BEBC [67], FMMF [68] and CHARM II [69] are all beam dump experiments which set limits on heavy neutrino decay. An additional result from E949, published after this plot, is shown in figure 1.6.

between $150 \text{ MeV}/c^2$ and $300 \text{ MeV}/c^2$.

The current limits on $|U_{\mu h}|^2$ from peak searches are shown in figure 1.6. The limits become weaker for lower neutrino mass because of the background from $K^+ \rightarrow \mu^+ \nu_\mu \gamma$ which mimics a signal if the photon is not detected; and for high neutrino mass because the slower muons, associated with heavier neutrinos, do not travel far enough in the experiment to be detected.

1.7. Heavy neutrino kinematics

1.7.1. Production

The mixing parameter, $|U_{\mu h}|^2$, can be extracted from $K^+ \rightarrow \mu^+ \nu_\mu$ studies by measuring or setting limits on the ratio of kaon decay width into heavy neutrinos to the decay width into Standard Model light neutrinos [55]:

$$R_{h0} \equiv \frac{\Gamma(K^+ \rightarrow \mu^+ \nu_\mu)_h}{\Gamma(K^+ \rightarrow \mu^+ \nu_\mu)_0}. \quad (1.106)$$

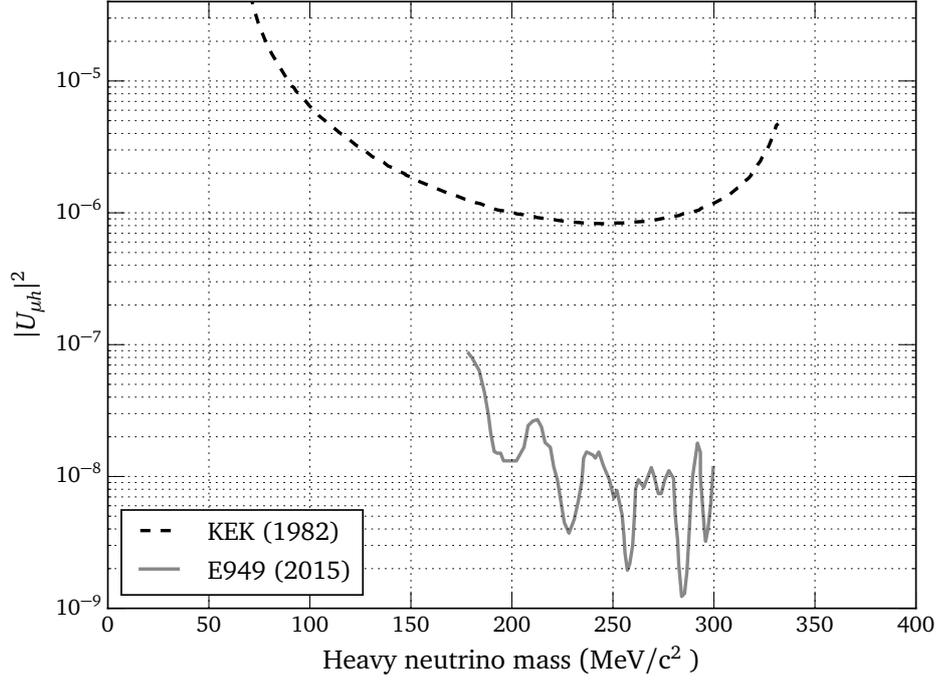


Figure 1.6.: Current limits on $|U_{\mu h}|^2$ from peak searches in kaon decays. The KEK result [65] was a dedicated experiment, while the E949 result [70] comes from a reanalysis of the existing $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ data set.

Here h labels the heavy neutrino mass state and 0 labels the Standard Model mass states. Summing over the three SM mass states, the SM decay width is:

$$\Gamma(K^+ \rightarrow \mu^+ \nu_\mu)_0 = \sum_{i=1,2,3} U_{\mu i} \Gamma(K^+ \rightarrow \mu^+ \nu_i), \quad (1.107)$$

Since the mixing to heavy neutrino mass states is known to be small, we have

$$\sum_{i=1,2,3} U_{\mu i} \simeq 1, \quad (1.108)$$

so the ratio R_{h0} can be written in terms of the total decay width to muon flavoured neutrinos:

$$R_{h0} \simeq \frac{\Gamma(K^+ \rightarrow \mu^+ \nu_h)}{\Gamma(K^+ \rightarrow \mu^+ \nu_\mu)}. \quad (1.109)$$

The kinematics of a two body decay to a massive neutrino are modified with respect to the decay to SM (approximately massless) neutrinos. There is a phase space contribution because the Q value of

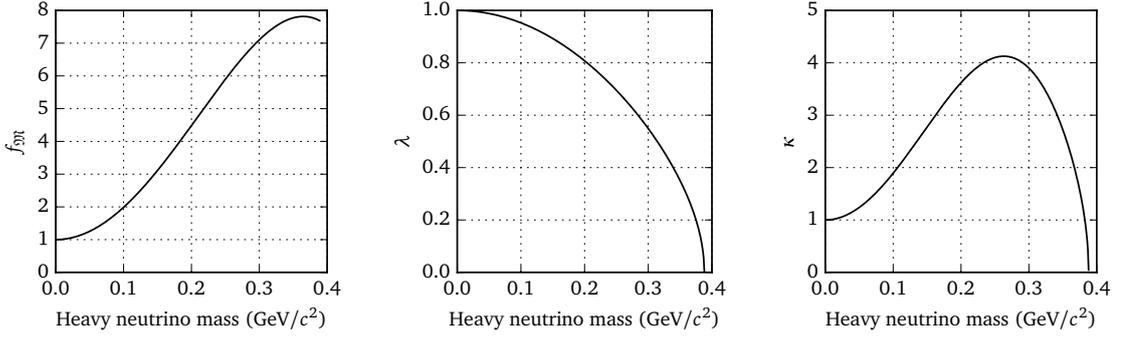


Figure 1.7.: Ratios of kinematic quantities for the heavy and massless neutrinos as a function of neutrino mass. $f_{\mathfrak{M}}$ (left) is the contribution from helicity suppression. λ (centre) is the contribution from phase space suppression. κ (right) is the product of the two, which enters in equation 1.110

the decay is lower, and a matrix element factor because helicity suppression coming from conservation of angular momentum is lifted in the case of massive neutrinos. These kinematic effects mean that R_{h0} depends on the mass of the heavy neutrino as well as the strength of the coupling given by the matrix element $|U_{\mu h}|^2$.

For a heavy neutrino of mass m_h , the ratio can be written:

$$R_{0h} = |U_{\mu h}|^2 \lambda(m_h) f_{\mathfrak{M}}(m_h), \quad (1.110)$$

where $\lambda(m_h)$ is the phase space factor:

$$\lambda(m_h) = \frac{\sqrt{m_K^4 + m_\mu^4 + m_h^4 - 2(m_\mu^2 m_K^2 + m_\mu^2 m_h^2 + m_h^2 m_K^2)}}{m_K^2 - m_\mu^2} \quad (1.111)$$

and $f_{\mathfrak{M}}(m_h)$ accounts for the helicity suppression:

$$f_{\mathfrak{M}} = \frac{\left(\frac{m_\mu}{m_K}\right)^2 + \left(\frac{m_h}{m_K}\right)^2 - \left(\left(\frac{m_\mu}{m_K}\right)^2 - \left(\frac{m_h}{m_K}\right)^2\right)^2}{\left(\frac{m_\mu}{m_K}\right)^2 \left(1 - \left(\frac{m_\mu}{m_K}\right)^2\right)^2} \quad (1.112)$$

Figure 1.7 shows the two ratios λ and $f_{\mathfrak{M}}$ as a function of heavy neutrino mass. The phase space function λ goes to zero at the kinematic end point $m_h = m_K - m_\mu$, when the muon and neutrino are produced at rest in the kaon frame.

Rearranging equation 1.106 and using equation 1.110, the magnitude of the mixing matrix element can be determined (or excluded above a certain value) from a measurement of relative branching ratios

$$|U_{\mu h}|^2 = R_{0h} \frac{1}{\lambda(m_h) f_{\text{SM}}(m_h)} = \frac{\Gamma(K^+ \rightarrow \mu^+ \nu_h)}{\Gamma(K^+ \rightarrow \mu^+ \nu_\mu)} \frac{1}{\lambda(m_h) f_{\text{SM}}(m_h)} \quad (1.113)$$

It can be seen from figure 1.7 that, for the same limits on the branching ratio, stricter limits could be placed on mixing with neutrinos of mass $\sim 250 \text{ MeV}/c^2$ than for lighter neutrinos, to which decays are more helicity suppressed, or to heavier neutrinos, to which decays are more phase space suppressed.

1.7.2. Decay

Peak searches can in principle be weakened if the heavy neutrinos produced decay into SM particle within the experiment so the event does not have the expected signature in the detector. Under the assumption that heavy neutrinos decay only to SM particles, the lifetime of a heavy neutrino is determined by its mass, since no new unknown coupling constants are introduced. For the heavy neutrino masses which can be studied at NA62, the dominant decay modes are [54]

$$\nu_h \rightarrow \pi^0 \nu_l \quad (l = e, \mu, \tau), \quad \nu_h \rightarrow \pi^+ l^- \quad (l = e, \mu) \quad (1.114)$$

Below the threshold for pion production, the dominant decay mode is

$$\nu_h \rightarrow \nu_\alpha \bar{\nu}_\beta \nu_\beta, \quad (1.115)$$

where α and β are any of e, μ, τ . Figure 1.8 shows the partial widths of all the principle decays of heavy neutrinos into SM particles, along with the total width as a function of mass, all at a fixed value for μh mixing, $|U_{\mu h}|^2 = 1 \times 10^{-5}$. The resulting mean lifetime for heavy neutrinos as a function of mass is shown in figure 1.9 and the consequences for a peak search at NA62 are discussed in chapter 3.

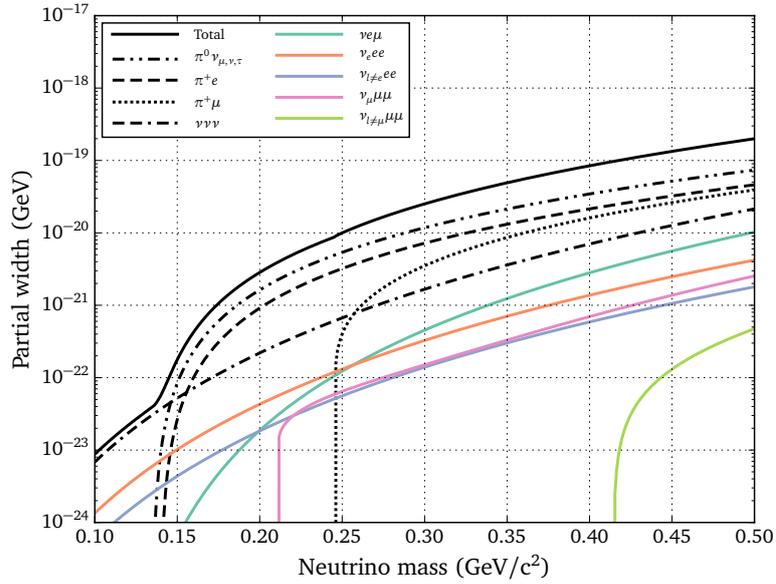


Figure 1.8.: Partial widths for heavy neutrino decays into SM particles as a function of heavy neutrino mass

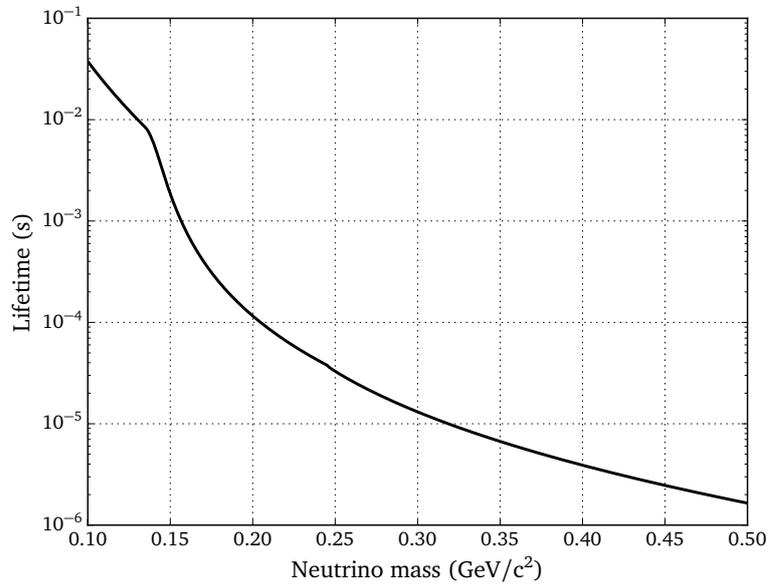


Figure 1.9.: Heavy neutrino lifetime, assuming decays to SM particles, as a function of heavy neutrino mass

The NA62 R_K detector

2.1. Fixed target kaon experiments at CERN

The NA62 experiment is the latest in a series of fixed target kaon decay in flight experiments at CERN.

In 1979, the source of CP violation in the kaon sector was not yet completely understood. In particular, all experimental evidence to date was consistent with the superweak model, in which all CP violation came from the mixing between the K_S and K_L (parameterized by ε) and there was no contribution from direct CP violation (parameterized by ε'). Between 1979 and 1991, the NA31 experiment [71] made a measurement of the double ratio of neutral kaon decay rates:

$$R = \frac{|\eta_{00}|}{|\eta_{+-}|} = \frac{\Gamma(K_L \rightarrow \pi^0 \pi^0)/\Gamma(K_L \rightarrow \pi^+ \pi^-)}{\Gamma(K_S \rightarrow \pi^0 \pi^0)/\Gamma(K_S \rightarrow \pi^+ \pi^-)}. \quad (2.1)$$

A departure from $R = 1$ would indicate a non-zero value of ε' (rejecting the superweak model) via the relation:

$$\Re(\varepsilon'/\varepsilon) = \frac{1}{6} \left(1 - \left| \frac{\eta_{00}}{\eta_{+-}} \right|^2 \right) \quad (2.2)$$

Neutral kaons, both K_S and K_L , were produced by impinging high energy protons (with momentum 450GeV/c) on a beryllium target and allowed to decay in flight along a common decay region. Neutral decay products were detected in a lead/liquid argon calorimeter. Charged tracks were observed with wire-chambers and their energies measured in the liquid argon calorimeter and iron/scintillator hadron calorimeter. The final result was [72]

$$R = 0.9862 \pm 0.0039, \quad (2.3)$$

implying $\Re\epsilon(\epsilon'/\epsilon) = (2.30 \pm 0.65) \times 10^{-3}$. The result, which was statistics limited, was more than three standard deviations from zero. However, other results at the time were not in complete agreement; the E731 experiment at Fermilab reported a value of $\Re\epsilon(\epsilon'/\epsilon) = (0.74 \pm 0.59) \times 10^{-3}$ [73], which, while not completely inconsistent with the NA31 result, was also consistent with zero.

In 1990, the experiment NA48 was proposed [74] to measure $\Re\epsilon(\epsilon'/\epsilon)$ with an accuracy of 2×10^{-4} . It would handle a tenfold increase in beam intensity relative to NA31 in order to reduce the statistical error on the measurement. Like NA31, the K_L and K_S beams were passed into a common decay volume but at NA48, in order to reduce possible sources of systematic error, all four decay modes were recorded simultaneously.

The layout of the downstream detectors would remain almost unchanged for the rest of the NA48 experiments. Charged particle tracks were measured in a magnetic spectrometer, housed in a helium tank and separated from the evacuated decay volume by a Kevlar window. In order to reconstruct neutral particles, electromagnetic showers were measured in the liquid Krypton calorimeter, a quasi-homogeneous detector with an active volume of $\sim 10 \text{ m}^3$ of liquid Krypton. Hadronic energy was measured in a calorimeter made of iron and scintillators. Finally, scintillator planes shielded by 80 cm thick iron walls were used to detect muons from background decays such as $K_L \rightarrow \pi\mu\nu$.

After taking data during 1997, 1998 and 1999, the NA48 measurement of the double ratio was [75]

$$R = 0.99098 \pm 0.00101 \pm 0.00126, \quad (2.4)$$

giving $\Re\epsilon(\epsilon'/\epsilon) = (15.3 \pm 2.6) \times 10^{-4}$, 5.9 standard deviations away from zero. An additional, low-intensity, data-taking period in 2001 allowed for checks that the result was not sensitive to beam intensity and resulted in an improved combined result of $\Re\epsilon(\epsilon'/\epsilon) = (14.7 \pm 2.2) \times 10^{-4}$ [76].

In 1999, the experiment NA48/1 was proposed [77] to study K_S and neutral hyperon decays with a modified K_S beam. Removing the K_L target made it possible to increase the intensity of the K_S beam by a factor $\sim 200 - 500$. The experiment proceeded in two phases. During Phase I in 2000, the wire chambers of the spectrometer were damaged and could not be used. Instead NA48/1 was able to exploit the energy resolution of the liquid krypton calorimeter to study K_S decays with only photons in the final state (such as $K_S \rightarrow \gamma\gamma$ [78]). In phase II, in 2002 with the repaired and upgraded spectrometer, a range of K_S decays (such as $K_S \rightarrow \pi^0\mu^+\mu^-$ [79]) and hyperon decays (such as $\Xi^0 \rightarrow \Sigma^+\mu^-\bar{\nu}_\mu$ [80]) could be studied.

Having established direct CP violation in the kaon sector, it was important to measure it in other

processes, especially as the world average of values for $\Re(\epsilon'/\epsilon)$ from neutral kaons was above most theoretical predictions. The NA48/2 experiment [81] was proposed to look for evidence of direct CP violation in the decays of charged kaons to pions. It was important to maintain symmetry in acceptance for K^+ and K^- decays. The beam line was redesigned to allow the simultaneous transport of positively and negatively charged kaons of the same momentum (60 GeV/c) and allow them to decay in the same volume. The system of magnets was designed so that paths of positively and negatively charged particles could be swapped in order to study systematic detector effects. Additionally it was possible to block one beam to study K^+ or K^- decays separately – this possibility was later exploited in the study of lepton universality.

The experiment looked for evidence of direct CP violation in the kinematic dependence of the matrix element for kaon decays to three pions. For the charged mode ($\pi^\pm\pi^+\pi^-$) and the neutral mode ($\pi^\pm\pi^0\pi^0$), the charge asymmetries of the Dalitz plot linear slopes were measured to be [82]

$$A_g^c = (-1.5 \pm 2.2) \times 10^{-4} \quad , \quad A_g^n = (1.8 \pm 1.8) \times 10^{-4}. \quad (2.5)$$

These limits were an order of magnitude better than the previous results and could be used to constrain physics beyond the Standard Model which might predict enhancements to direct CP violation.

2.2. NA62- R_K : Lepton Universality test in K^\pm decays

In 2006, the new NA62 experiment to measure ultra-rare kaon decays had already been proposed [28], requiring a new detector which would include some of the existing NA48 subdetectors. However, following a theoretical study [83] of lepton flavour violating effects, it became clear that another period of data taking with the existing detector could have a significant impact.

The leptonic decay rates of pseudoscalar mesons are predicted to a few % precision in the Standard Model. To lowest order, the decay width is

$$\Gamma(P \rightarrow l \nu) = \frac{G_F^2}{8\pi} f_P^2 m_l^2 m_P \left(1 - \frac{m_l^2}{m_P^2}\right)^2 |V_{q_1} V_{q_2}|^2, \quad (2.6)$$

where G_F is the Fermi coupling constant, m_P is the mass of the meson P , and $|V_{q_1} V_{q_2}|^2$ is the CKM matrix element between the two quarks flavours in P . The parameter f_P is the meson decay constant, which is related to the overlap of the quark and antiquark, and is difficult to predict from theory [8]. However,

ratios such as $R_K = \Gamma(K \rightarrow e \nu)/\Gamma(K \rightarrow \mu \nu)$ can be predicted with much better precision because the decay constants cancel. R_K is highly suppressed in the SM by helicity constraints (responsible for the term m_l^2 in the decay width), and this makes it particularly sensitive to New Physics in which the universality of the weak coupling between different lepton flavours is violated [83].

In 2007, the PDG precision of the measured value of R_K was $\sim 5\%$, worse than the error on the SM prediction [84]

$$R_K^{SM} = (2.472 \pm 0.001) \times 10^{-5}. \quad (2.7)$$

Furthermore, New Physics theories such as the Minimal Supersymmetric Standard Model predicted enhancements of $\mathcal{O}(1\%)$. It was proposed [85] that a few months of data taking, dedicated to measuring R_K , could achieve sub-percent level precision.

The subdetectors used in the measurement were the magnetic spectrometer, for tracking charged particles, the scintillator hodoscope for triggering on charged particles and the liquid krypton calorimeter for photon vetoing and particle identification. The number of decays $K^+ \rightarrow e^+ \nu_e$ and $K^+ \rightarrow \mu^+ \nu_\mu$ were counted simultaneously so that several systematic effects, such as the absolute kaon flux, cancelled at first order. Data was taken over a period of about 4 months in 2007 with a strategy optimized to measure the $K^+ \rightarrow e^+ \nu_e$ decay and its two main backgrounds.

The dominant background to $K^+ \rightarrow e^+ \nu_e$ came from $K^+ \rightarrow \mu^+ \nu_\mu$ decays, with the muon misidentified as an electron; several strategies were employed to reduce this background. Kinematic separation was improved by increasing the beam momentum (from 60 GeV/c to 74 GeV/c) as well as the momentum kick in the spectrometer magnet (from 120 MeV/c to 265 MeV/c) so that $K^+ \rightarrow e^+ \nu_e$ and $K^+ \rightarrow \mu^+ \nu_\mu$ decays could be distinguished by their reconstructed missing mass over a greater range of momenta.

Muons and electrons could also be distinguished by their energy deposition in the liquid krypton calorimeter. Electrons typically deposit nearly all of their energy whereas muons are normally minimally ionising particles. By cutting on the ratio of energy deposition to particle momentum, E_{LKr}/p , (with a momentum dependent cut-off ranging from 0.9 to 0.95) it was possible to identify electrons with 99% efficiency with muon contamination around 4×10^{-6} (non-negligible in comparison with the value of R_K). The probability for a muon to deposit a large fraction of its energy in the LKr was measured with a pure muon sample, obtained by installing a lead bar to shield part of the LKr from electrons during some of the runs.

The second background to $K^+ \rightarrow e^+ \nu_e$ (which also affected $K^+ \rightarrow \mu^+ \nu_\mu$) came from the beam

muon halo. Since there was no kaon tagging in the beam, there was a background from the decay of beam particles into muons, upstream of the fiducial decay volume. This had a direct impact on the $K^+ \rightarrow \mu^+ \nu_\mu$ measurement as well as on $K^+ \rightarrow e^+ \nu_e$ via muon decays to electrons in the fiducial region. Muon scraping magnets were installed on the beam line, optimized for halo suppression in the K^+ data samples. In addition, some data samples were taken with both the K^+ and K^- beams blocked off in order to measure the halo directly.

The final measurement of R_K by NA62- R_K with the 2007 data set was [86]:

$$R_K = (2.488 \pm 0.007_{\text{stat.}} \pm 0.007_{\text{syst.}}) \times 10^{-5}, \quad (2.8)$$

which dominates the world average. The result is consistent with the SM prediction but the experimental uncertainty is still an order of magnitude larger than the precision of the theoretical prediction so future measurements (for example by NA62) will continue to be of interest.

The minimum bias sample of kaon decays collected for the R_K measurement has made possible a number of other measurements, including the search for heavy neutrinos described here. The complete detector in its 2007 form is described below.

2.3. The beam and detector

2.3.1. The charged kaon beam line

The measurement of CP violation in kaon decays to three pions required a novel beam design [82], allowing for simultaneous beams of oppositely charged particles. Protons with momentum 400 GeV/c from the CERN SPS were impinged on a beryllium target producing a secondary beam containing protons, pions and kaons. There is an additional electron component from the production and decay of π^0 in the target, followed by photon conversion. This component is measured to be $\sim 10\%$.

The secondary beam momentum of 74 GeV/c was determined by passing the beam through an achromat comprising four dipole magnets of alternating sign (see figure 2.1). Particles of opposite signs were given opposite deflections in the vertical direction and were recombined after the achromat with no net deflection. A pair of beam dump/collimators, TAX1 and TAX2 between the second and third dipole magnets ensured that only particles of the required momentum were selected (TAX stands for ‘Target Attenuator eXperimental areas’). Each TAX was made from 1.6 m (along the beam axis) of copper and iron, with slits to allow particles of the correct charge and momentum to pass through.

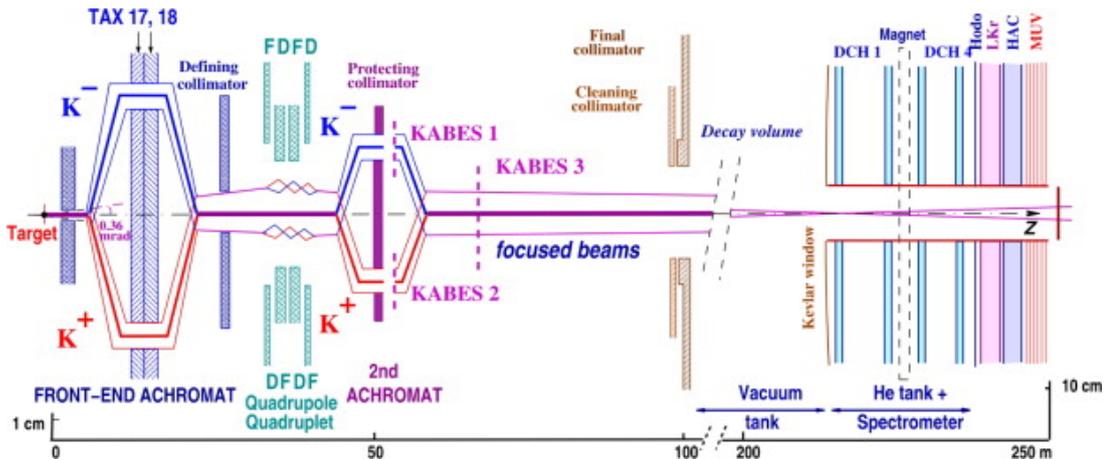


Figure 2.1.: The beamline and detector for NA62 data taking in 2007. [87]

The width of the slits determines the momentum spread of the selected particles, which was $1.4 \text{ GeV}/c$ (rms). By closing one slit, one polarity could be blocked completely to study K^+ or K^- separately.

After the achromat, the beam passed through a second collimator, defining the spatial acceptance followed by four quadrupole magnets which focussed the beams (in a charge symmetric manner) in both the horizontal and vertical planes. A second achromat, placed after the quadrupoles, made it possible in principle to measure the beam particle momenta with the KABES (TPC type) detectors [88] but these were not installed in the R_K measurement. Before entering the decay volume, the beam passed through a final pair of collimators.

The SPS delivered primary protons in spills lasting 4.8 s and containing 1.8×10^{12} protons. The process repeated every 14.4 or 16.8 s so the secondary beam arrived in bursts, each containing about 2.5×10^7 particles. The secondary beam had a transverse size at the entrance to the decay volume described by $\delta x = \delta y = 4 \text{ mm}$ (rms) and horizontal and vertical angular divergences of $20 \mu\text{rad}$ (rms). At this point, the fractions of K^+ , π^+ , p^+ , e^+ and μ^+ were 0.05 , 0.63 , 0.21 , 0.10 and 0.01 respectively. 18% of kaons decayed in the decay volume.

2.3.2. The decay volume

In order to avoid interactions before decaying, the fiducial volume for kaon decay was enclosed in a vacuum tank kept at $< 10^{-4} \text{ mbar}$ [89]. The tank started at the final collimator, 102 m from the proton target, and extended up to the start of the spectrometer, 114 m downstream. A convex Kevlar window separated the vacuum from the helium gas at atmospheric pressure which filled the spectrometer; the window contributed $3 \times 10^{-3} X_0$ to the material budget. Undecayed beam particles continued

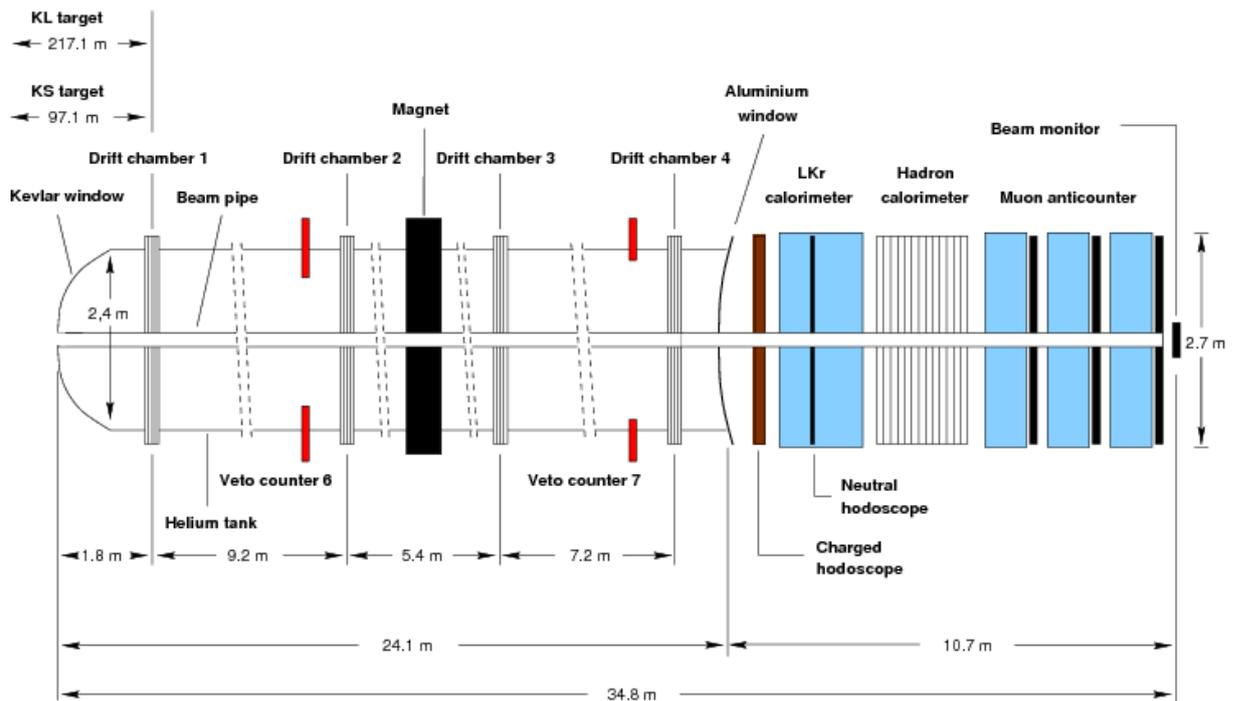


Figure 2.2.: The NA62 subdetectors downstream of the decay volume. [75]

in vacuum, contained in an aluminium pipe of inner diameter 152 mm and 1.2 mm thickness. The vacuum tank had cylindrical cross-section with a diameter of 1.92 m for the first 65 m and 2.4 m for the rest.

2.3.3. The magnetic spectrometer

The detection and momentum measurement of charged particles was performed by a magnetic spectrometer, consisting of four high-resolution Drift Chambers (DCH), two either side of a dipole magnet (see figure 2.2). In order to minimize scattering, the spectrometer was housed in a tank, about 23 m long and 2.8 m in internal diameter, filled with helium at atmospheric pressure. The helium purity was $(92 \pm 4)\%$ so all the gas in the tank contributed $(2.6 \pm 0.3) \times 10^{-3} X_0$ to the material budget [86].

The magnet [90] produces a vertical magnetic field over an active region with aperture 2.4 m and, in 2007, the field strength was set to give a horizontal momentum kick of 265 MeV/c. A map of the magnetic field was measured using a set of Hall probes in order to correct for variations in the field. During the 2007 data taking, the polarity of the field was reversed three times in order to cancel any systematic effects due to asymmetry in the detector acceptance.

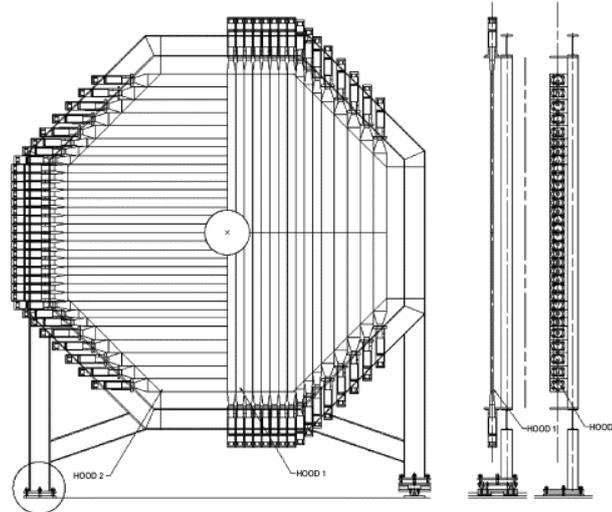


Figure 2.3.: The charged hodoscope detector [87]

The four drift chambers had an octagonal shape with a fiducial area of about 4.5 m^2 and a central hole of diameter 16 cm to allow the beam to pass through. Each chamber was divided into eight planes of sense wires, in four different orientations (views) in the xy plane, $0^\circ(X)$, $90^\circ(Y)$, $-45^\circ(U)$ and $+45^\circ(V)$. Each view contained two planes of wires staggered by 5 mm (half the inter-wire separation) to resolve left-right ambiguity.

The drift chambers operated with a 50:50 mixture of Ar and C_2H_6 gas, resulting in a gas gain of about 8×10^4 and an average drift velocity of the electrons of $45 \mu\text{m}/\text{ns}$. The hit time resolution depends on the hit location along the sense wire so the spatial resolution of individual hits varied between $90 \mu\text{m}$ and $170 \mu\text{m}$ as a function of the hit position. When combining information from hits in all eight sense wire planes, the spatial resolution was found to be better than $100 \mu\text{m}$ in the X and Y directions. The momentum resolution (measured with electron beams of momentum 25 and $100 \text{ GeV}/c$) was found to be $\sigma_p/p = 0.48\% \oplus 0.009p[\text{GeV}/c]\%$, where the first term is due to multiple scattering in the He tank and the second term is due to the position resolution of the chambers. The time resolution for reconstructed tracks was 700 ps.

2.3.4. The hodoscope

Triggering on charged particles was provided by the hodoscope, a system of plastic scintillators. The detector comprised two planes of 64 scintillator strips, aligned vertically in the upstream plane and horizontally in the downstream plane.

Each strip was 2 cm thick and the strips close to the beam pipe were 6.5 cm wide, while those further out were 9.9 cm wide. Scintillation light was collected by a Plexiglas light guide and recorded by a photomultiplier, one per strip. Strips were painted black on the side opposite the PM to prevent spurious signals from reflected light. The two planes were separated by 74 cm in order to distinguish prompt signals from calorimeter backscattering.

Each plane was divided into four quadrants of 16 counters so that specific triggers could require various combinations of hit locations in each plane. The time resolution was measured to better than 200 ps per counter.

2.3.5. The liquid krypton electromagnetic calorimeter

Electromagnetic calorimetry and photon detection were provided by the liquid krypton calorimeter (LKr). The detector is a quasi-homogeneous ionization chamber with a cylindrical volume of surface area 5.3 m² and depth 1.27 m, filled with liquid krypton. It allowed for full development of electromagnetic showers (with > 99% containment), providing a stable signal which was well correlated with the energy of the incoming photon or electron.

The volume was divided into 13 248 ionization cells by thin Cu-Be ribbons, used to construct longitudinal towers with square 2 cm × 2 cm cross section. Since the Molière radius of liquid krypton is 6.1 cm and its radiation length is 4.7 cm, the shower from individual photons encompasses a number of cells, forming a *cluster*. Energies were calculated by summing the energies recorded by all the cells within a cluster (typically about 100), while time and position information were calculated from only the most energetic cells.

The ratio between the current measured in an ionization cell and the total energy deposited depends strongly on the size of the cell so it was important to maintain strict tolerances on individual cell widths. The Cu-Be ribbons were held under 2 N of tension each to hold them in place against spacer-plates which were positioned at 21 cm intervals in *z*. In order to facilitate calibration measurements, the pitch of the cells was increased by 0.19% between successive spacers to give LKr a projective structure, diverging from a point 110 m upstream. This allowed the energy scale of the calorimeter to be measured with π^0 decay, independently of any tracking information.

The measured energy resolution of the LKr was (all energies in GeV):

$$\frac{\sigma_E}{E} = \frac{0.031}{\sqrt{E}} \oplus \frac{0.09}{E} \oplus 0.0042. \quad (2.9)$$

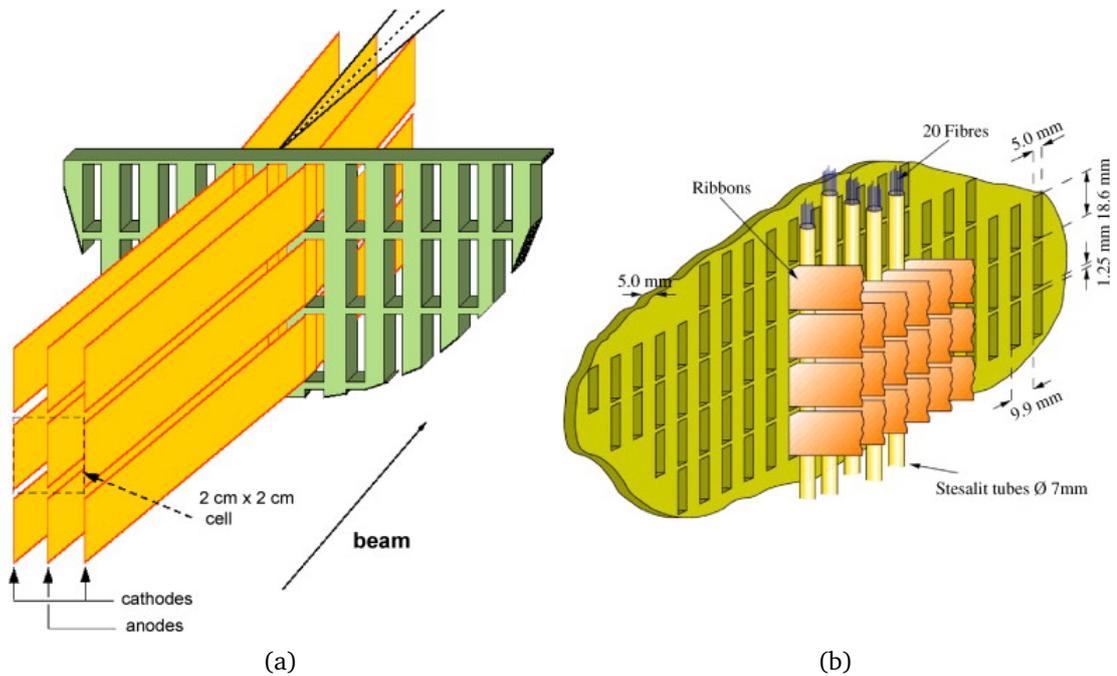


Figure 2.4.: (a) An example of the LKr calorimeter cell structure. (b) The neutral hodoscope embedded in the LKr calorimeter [87].

The first term is a sampling effect coming from the intrinsic (Poisson) fluctuations in energy deposition in electromagnetic showers as well as the finite size of clusters used to reconstruct shower energies. The second term is the contribution from instrumental noise (σ_E independent of E). The constant term comes from the accuracy of the calibration between the reconstructed energy of the shower and the real energy deposition in the calorimeter.

2.3.6. The neutral hodoscope

In order to improve the time resolution for neutral events, which were not detected by the CHOD, an independent time measurement was made using the neutral hodoscope embedded in the LKr. Scintillating fibres were installed vertically at a depth of about $9.5X_0$ inside the LKr, where the electromagnetic shower had its maximum (figure 2.4b). The fibres were grouped into bundles of about 20, each 5 mm in diameter, and placed inside epoxy-fibreglass tubes at ~ 1 cm intervals. Light signals were detected by photomultipliers located inside the LKr but outside of the active volume of the calorimeter. The timing resolution for neutral triggers was better than 200 ps.

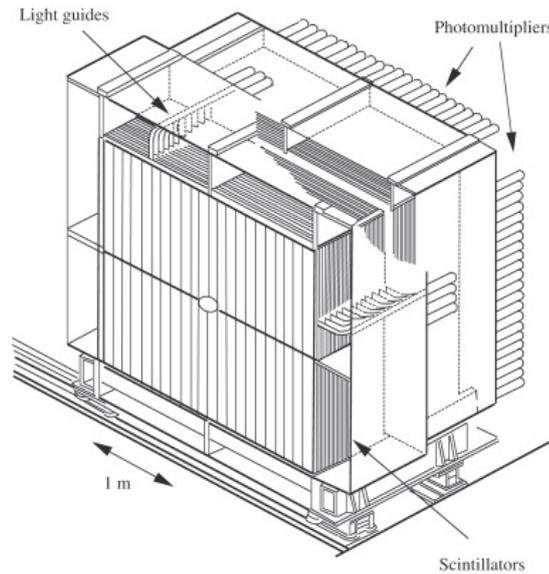


Figure 2.5.: The hadronic calorimeter [87]

2.3.7. The hadronic calorimeter

The hadronic calorimeter (HAC, not used in this analysis) was an iron-scintillator calorimeter, with cross-section $2.6\text{ m} \times 2.6\text{ m}$, mainly used in triggering. It was divided longitudinally into two modules, front and back, each module containing 24 iron plates, 25 mm thick, resulting in a total of 1.2 m of iron (or 6.7 nuclear interaction lengths). Between the iron plates were layers of plastic scintillator, divided into 22 strips in one direction and split in half in the other direction, making 44 strips per layer. The orientation of strips was alternated between consecutive layers so the location of the energy deposition could be determined in both x and y and the HAC could be used to trigger on the energy sum in each quadrant.

Signals were read out by photomultipliers, coupled to the scintillator strips by a Plexiglas light guide. The signal shapes were modified in order to share the same digitizer used for the LKr calorimeter. In 1995 the energy resolution was determined to be $\sigma_E/E = 0.69\sqrt{E}$ for hadronic showers (E measured in GeV).

2.3.8. The muon veto system

The muon veto was originally used in the NA31 experiment [71]. Its primary purpose in both the NA31 and NA48 experiments was to reduce trigger rates by vetoing $K_L \rightarrow \mu^+ \pi^- \nu_\mu$ decays, as well as allowing for positively triggering on muons. However, it was not switched on for all of the data

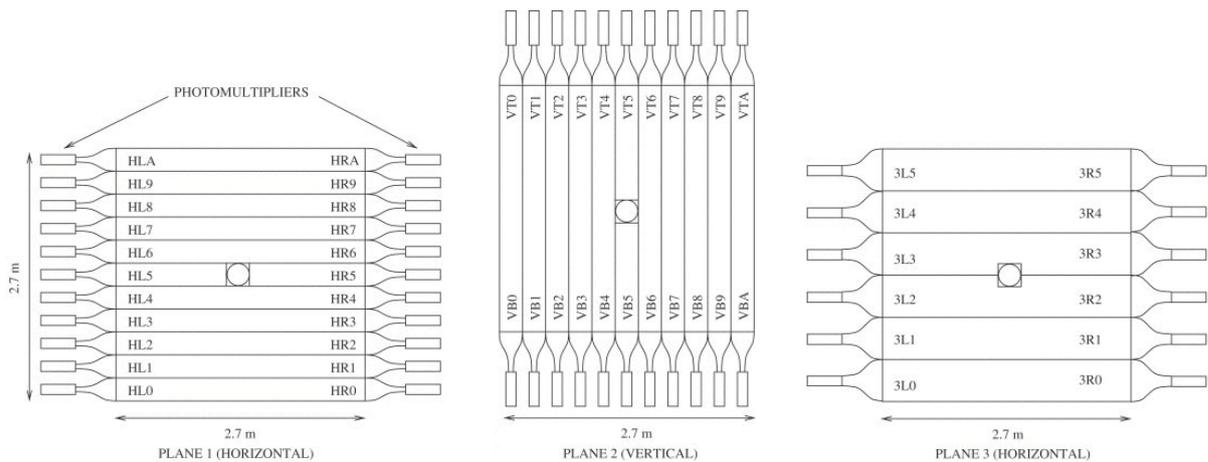


Figure 2.6.: The three planes of the muon veto system [87]

taking in 2007.

The detector was made from three planes of scintillating plastic strips (the first two for efficient muon vetoing and the third plane for extra efficiency and absorption studies [91]). In front of each plane was an iron wall, 0.8 m thick, to act as an absorber for hadrons. Planes 1 and 2 were made from strips 2.5 m wide and 1 cm thick, oriented horizontally (plane 1) and vertically (plane 2). Each strip was read out via two EMI 9813 KB photomultiplier tubes, one at each end, coupled to the strip by a semiadiabatic light-guide. The one exception was the central strip in each plane which was split into two halves to allow the beam to pass through, each half read out by a single PM. The third plane comprised 6 strips, 45 cm wide, with a Philips 56A VP photomultiplier tube at each end.

In order to avoid inefficiencies in muon detection due to gaps between the scintillator strips, the strips were overlapped by an amount varying from 3 mm at the outer edges of the detector to 21 mm near the beam pipe where the rate of muons was highest. The measured efficiency for single muon events (with $p_{\text{muon}} > 10 \text{ GeV}/c$) was found to be $> 99\%$ with a time resolution of $\sim 350 \text{ ps}$.

2.3.9. Anti counters

In order to detect photons outside of the acceptance of LKr calorimeter, seven large-angle Anticounters (AKL) were installed along the beam pipe. Each Anticounter station (or pocket) comprised two layers of scintillation counters, covering an annular shape around the LKr acceptance, with a thick steel plate in front to act as a photon converter. For decays in the fiducial region, the AKL provided coverage for photon angles from 10 to 50 mrad, with total photon conversion efficiency $\approx 95.5\%$, and time

resolution ≈ 550 ps. The AKL is not used in the final selection for the search for heavy neutrinos, but it provides a useful check of the composition of the remaining backgrounds after the selection.

2.3.10. Trigger

Although the rate of kaon decays in the fiducial region was $\mathcal{O}(100\text{ kHz})$, the detector was subject to a total rate of particles $\mathcal{O}(1\text{ MHz})$. A multi-level trigger system was employed to reduce the rate of events to $\sim 10\text{ kHz}$ which could be written to disk. Data from each sub-detector were read-out and stored in circular buffers $204.8\ \mu\text{s}$ deep so the trigger system was required to process events within this time. The trigger system contained a Charged Chain and a Neutral Chain, both fed into the Trigger Supervisor, which produced the final decision to accept or reject each event.

In 2007, the Charged Chain was a fast logic trigger, which received signals from the drift chambers and the hodoscope. The minimum bias trigger was based on the Q_1 signal: a coincidence of at least one hit in each of the hodoscope scintillator planes, both hits belonging to the same quadrant. Single track events were selected by setting upper and lower limits on the number of hits allowed in the drift chambers. In particular the 1TRKLM trigger (1 track, loose modified), required >0 hits in >1 views and <15 hits in any view in three out of the four drift chambers: DCH1, DCH2 and DCH4. $K^+ \rightarrow \mu^+ \nu_\mu$ events were selected by requiring both Q_1 and 1TRKLM triggers (written $Q_1 \times 1\text{TRKLM}$).

The Neutral Chain processed data from the liquid krypton calorimeter and neutral hodoscope. The $E_{\text{LKr}}(10\text{ GeV})$ signal indicated a total energy deposition of $> 10\text{ GeV}$ in the calorimeter. An independent trigger from the neutral hodoscope (written NHOD) was used for efficiency studies of the E_{LKr} trigger.

The signals from the charged and neutral triggers were sent to the Trigger Supervisor, which combined all the information for a single event in order to make a final trigger decision. The high beam rate was required to obtain a large number $K^+ \rightarrow e^+ \nu_e$ decays, which were collected with a dedicated trigger requiring a minimum energy deposit in the LKr calorimeter. Other trigger streams could be downscaled, allowing a fraction of minimum bias triggers such as Q_1 to be recorded without overloading the data acquisition system. The exact downscaling factors varied over the course of the data taking (see the following section) but were typically $\mathcal{O}(1000)$ for the Q_1 trigger and $\mathcal{O}(100)$ for the $Q_1 \times 1\text{TRKLM}$ and NHOD triggers. The status of the various trigger streams was encoded in a 16-bit trigger word, which was output by the Trigger Supervisor along with a timestamp for the event. Whenever an event was recorded on disk, the complete set of trigger signals were stored in

Pattern Units, which could later be used for efficiency studies.

2.3.11. Data acquisition

During data taking, about 3 Tb of raw data were produced every day. The raw data were processed at CERN, producing reconstructed events which were stored in the COMPACT data format.

Data are divided into ‘bursts’, each corresponding to a single spill of protons from the SPS. A spill lasted about 4.8 s and typically contained $\sim 10^{12}$ protons, resulting in $\sim 5 \times 10^4$ recorded kaon decays in the vacuum tank. In the 2007 R_K analysis, with the 1TRKLM trigger downscaled by 150 and signal acceptance of $\sim 40\%$, there were about 300 reconstructed $K^+ \rightarrow \mu^+ \nu_\mu$ events per burst. Bursts are grouped into ‘runs’ containing $\mathcal{O}(1000)$ bursts. A run corresponds to a few hours of data taking in which the detector configuration was unchanged.

Data were collected over a period of ~ 120 days, divided into 5 ‘periods’, with different data taking conditions (including beam intensity, trigger definitions and downscaling, kaon charge etc.). For the R_K measurement, it was important to understand the probability for a muon to deposit a large fraction of its energy in the LKr calorimeter, producing an electron-like signal. This probability was measured with a pure muon sample, obtained by installing a lead bar to shield part of the LKr from electrons during periods 1 - 4 (see figure 2.7), reducing the LKr acceptance for these periods, so there are much higher backgrounds from events with photons in the final state.

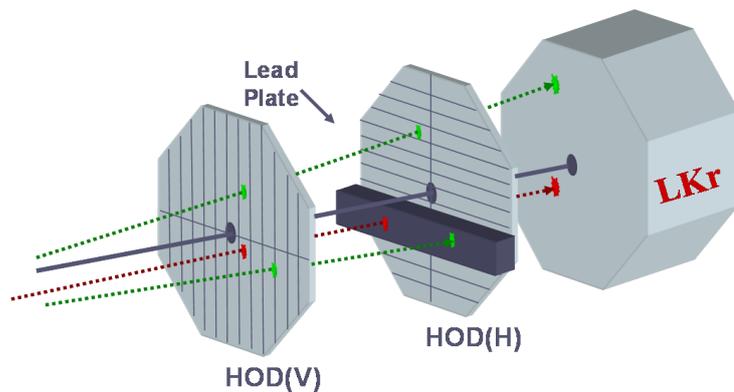


Figure 2.7.: A schematic diagram of the lead wall used to study muon energy deposition in the LKr

Search for heavy neutrinos at NA62- R_K

3.1. Analysis strategy

As part of a measurement of Lepton Universality in 2007, the NA62- R_K experiment recorded a large sample of $K^+ \rightarrow \mu^+ \nu_\mu$ decays with an unbiased trigger using the detector setup described in chapter 2 (different from final NA62 detector used for the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ measurement and described in chapter 4). Reanalysis of this sample provides an opportunity to look for evidence of heavy neutrinos in the kinematic spectrum of the muons produced in association with them.

In the peak search approach taken in this analysis, the experimental signature for heavy neutrinos is a single charged track from the muon. If heavy neutrinos decay into SM particles before leaving the detector, they can produce other signatures in the detector so the acceptance for heavy neutrinos is reduced. There is no danger of background from the decays of pions, which make up the largest fraction of the beam, because the Q-value for the decay $\pi^+ \rightarrow \mu^+ \nu_\mu$ is small enough that if a pion decays in the fiducial decay volume, the muon will not be emitted at a large enough angle to reach the drift chambers of the spectrometer.

Figure 1.9, on page 35, shows that for the relevant masses of heavy neutrino the lifetime is always greater than 1×10^{-6} s, assuming the neutrino decays to SM particles. Furthermore, MC simulation of events within the NA62 acceptance shows that in the relevant decays, heavy neutrinos are produced with momenta between 20 GeV/c and 65 GeV/c, so the relativistic γ factor is always greater than ~ 50 . Using these minimum values, the mean flight distance for heavy neutrinos produced at NA62- R_K ,

must be at least

$$\lambda_{\min} = \gamma\beta c\tau \quad (3.1)$$

$$= 50 \times 1 \times (3 \times 10^8 \text{ ms}^{-1}) \times (1 \times 10^{-6} \text{ s}) \quad (3.2)$$

$$= 1.5 \times 10^4 \text{ m}. \quad (3.3)$$

A rough upper limit on the fraction of neutrinos which decay inside the detector can be obtained by assuming that on average, a neutrino will travel $l_{\text{det}} \sim 100 \text{ m}$ before escaping:

$$f_{\text{decay}} = 1 - \exp(-l_{\text{det}}/\gamma\beta c\tau) = 1 - \exp(-100 \text{ m}/15000 \text{ m}) = 0.66\%, \quad (3.4)$$

This has a negligible impact on the sensitivity of NA62- R_K .

In the decay $K^+ \rightarrow \mu^+ \nu_\mu$, the neutrino mass can be reconstructed from the four momenta of the kaon and muon:

$$m_\nu^2 = m_{\text{miss}}^2 = (p_K - p_\mu)^2. \quad (3.5)$$

At NA62- R_K , the average kaon momentum was measured during data taking, with the decay $K^+ \rightarrow \pi^+ \pi^+ \pi^-$, (the momentum spread was $\sigma_p = 1.4 \text{ GeV}/c$) and the muon momentum is measured separately for each event, using the magnetic spectrometer. $K^+ \rightarrow \mu^+ \nu_\mu$ events can be selected by a single charged track, positively identified as a muon, with no photons detected in the electromagnetic calorimeter. The decay with SM light neutrinos is simulated in Monte Carlo, along with the other kaon decay modes which act as backgrounds, in order to determine the expected spectrum of reconstructed m_{miss}^2 . A comparison between the expected and observed spectra is used to set limits on the branching ratio to heavy neutrinos as a function of mass. This can be translated into a limit on the mixing between the muon flavour state and a heavy mass state for neutrinos.

The following sections describe this analysis in detail. A loose single track selection is used to determine the potential backgrounds to the signal decay. The additional selections required to suppress these backgrounds are presented, along with supporting studies to quantify their performance.

The complete selection is analysed in order to determine the contributions to the mixing limits from systematic errors. Finally the expected sensitivity to kaon decay to heavy neutrinos is presented as a function of mass and limits are set using the observed data.

Period	1	2	3	4	5	6
Beam	K^+/K^-	K^+/K^-	K^+/K^-	$K^+/K^-/$	K^+	K^-
MUV performance	bad	bad	bad	good	good	bad
Lead bar	yes	yes	yes	yes	no	no

Table 3.1.: Relevant statuses of the NA62- R_K detector configuration for each of the data taking periods in 2007.

3.2. Selection of data taking periods

As described in section 2.3.11, data taking in 2007 was divided into several ‘periods’, with different experiment configurations. The design of the present analysis restricts the data taking periods which could be used in the final sample.

The experiment was designed to make it possible to run with both positive and negative beams simultaneously (*c.f.* section 2.3.1) but the beam sweeping magnets designed to minimize the background from the beam halo were optimized for K^+ decays. Since the beam halo constitutes a significant background to heavy neutrinos, the search has been confined to K^+ data. The Muon Veto is critical for the suppression of backgrounds due to kaon decays into pions, so only periods in which it was working well can be used for analysis. An exception is for data used to study the beam halo, which is dominated by muons so can be studied without requiring particle identification. For periods with the lead bar in place, the efficiency for vetoing background producing photons is reduced, so these periods are excluded from the present analysis. Table 3.1 shows how these constraints apply to each of the data taking periods.

Period 5 is the only one suitable for analysis of K^+ decays. It contains 50 runs and 88051 bursts. As discussed in section 3.7, the background from beam halo muons can be analysed using data taken without a K^+ beam. Period 6 data is suitable for this purpose as it contains 19 runs (24,000 bursts) with a K^- beam only and 7 runs (7000 bursts) in which both beams were blocked. In the following sections, these data taking periods are referred to as K^+ (period 5), and K^- and $Kless$ periods (both period 6).

3.3. Pre-selection

3.3.1. Single Track Pre-selection

Single track events were selected using the $Q1 \times 1TRKLM$ trigger, which places restrictions on the number of hits in the spectrometer drift chambers. The limits on hit multiplicity do not exclude multi-track events. For events which pass the trigger, an additional pre-selection is applied to select a cleaner sample of single track events which are suitable for further analysis. This pre-selection is similar to the $K^+ \rightarrow \mu^+ \nu_\mu$ selection used for the R_K analysis [92]. In general, no additional cuts are imposed by the choice of trigger since the single track pre-selection is more restrictive than the $Q1 \times 1TRKL$ requirement. An inefficiency due to electromagnetic showers in the spectrometer is discussed in section 3.10.4.

Unless otherwise specified a single reference frame is used throughout this analysis, in which the final collimator at the entrance of the decay volume ends at $z = -18$ m and the first drift chamber of the spectrometer is at $z = 97.08$ m.

In order to avoid rejecting too many events unnecessarily, the pre-selection allows for the presence of “ghost tracks”, where a single charged particle is reconstructed more than once and “bad tracks” which can be ignored. A selected event must have no more than one “good track”, but may have additional ghost tracks. Initially, any tracks with momentum outside of the range 3 - 65 GeV/c are ignored. The remaining tracks are all defined as “good” initially but can be relabelled as “ghost” or “bad” tracks. As part of the standard selection to reduce backgrounds from photon conversion, if two tracks are within 0.5 cm as they enter the first drift chamber, the one with lower quality is labelled a ghost track and the other is labelled a good track. This cut is applied applied for consistency with other NA62- R_K analyses but has negligible impact on the selection of $K^+ \rightarrow \mu^+ \nu_\mu$ events.

A good track is relabelled “bad” if the recorded track time is more than 62.5 ns away from the trigger time recorded by the hodoscope, since an out of time track is likely to come from a different event and can be safely ignored. After reconstructing the vertex between the kaon and the charged particle (without applying any corrections), the track is labelled “bad” if the *closest distance of approach* (CDA) between the kaon and the charged particle is more than 10 cm or if the reconstructed vertex is outside of the decay region. Here the decay region is defined as between -20 m and 90 m,

The single track pre-selection is often used as part of other selections in later sections, sometimes with modifications to the initial ranges of allowed momentum and vertex location.

3.3.2. Signal shapes

Figure 3.1 shows the spectrum of reconstructed squared missing mass for simulated heavy neutrino events which pass the single track pre-selection, along with some of their other kinematic properties. The number of heavy neutrinos simulated is the same in each case.

The resolution on reconstructed squared missing mass, $\sigma_{m_{\text{miss}}^2}$, is defined as half the width which contains 68.2% of the signal events. The resolution on the mass peaks improves at higher mass because the associated muon is more often produced with lower momentum in the laboratory frame when the neutrino mass is greater.

Signal acceptance is defined as the number of simulated events passing the pre-selection, divided by the number of events with the z coordinate of the simulated vertex between -18 m and 80 m. The acceptance decreases at higher mass because lower momentum muons are more likely to be emitted at small angles to the kaon in the laboratory frame and fail to reach the inner radius of the first drift chamber. This can be seen in figure 3.2.

There is an irreducible background to the search for heavy neutrinos from the Standard Model decay $K^+ \rightarrow \mu^+ \nu_\mu$ (i.e the $m = 0$ case shown in the plots in figure 3.2). Since the $K^+ \rightarrow \mu^+ \nu_\mu$ branching ratio is large (63%), it is only viable to look for heavy neutrinos with masses sufficiently large that the signal peak is well separated from the $K^+ \rightarrow \mu^+ \nu_\mu$ peak. The expected sensitivity to heavy neutrinos with the NA62 2007 data set is only an improvement on previous experimental limits for neutrino masses $> 300 \text{ MeV}/c^2$. For this analysis, in order to avoid significant contributions from the far tails of the $K^+ \rightarrow \mu^+ \nu_\mu$ peak (see section 3.5) and contributions from punch-through kaons (see section 3.7.7), the *signal region* covers the range of heavy neutrino masses $275 \text{ MeV}/c^2 - 375 \text{ MeV}/c^2$, corresponding to squared missing mass, m_{miss}^2 , in the range $0.075 \text{ GeV}^2/c^4 - 0.14 \text{ GeV}^2/c^4$

3.3.3. Background composition

Figure 3.3 shows the reconstructed squared missing mass spectrum for simulated background events passing the pre-selection. (The halo distribution comes from data, see section 3.7.) There are clear peaks from the $K^+ \rightarrow \mu^+ \nu_\mu$ and $K^+ \rightarrow \pi^+ \pi^0$ decays and contributions from three body decays at higher m_{miss}^2 . The properties of each background are discussed below in order to motivate the selection strategy to improve the ratio of signal to background.

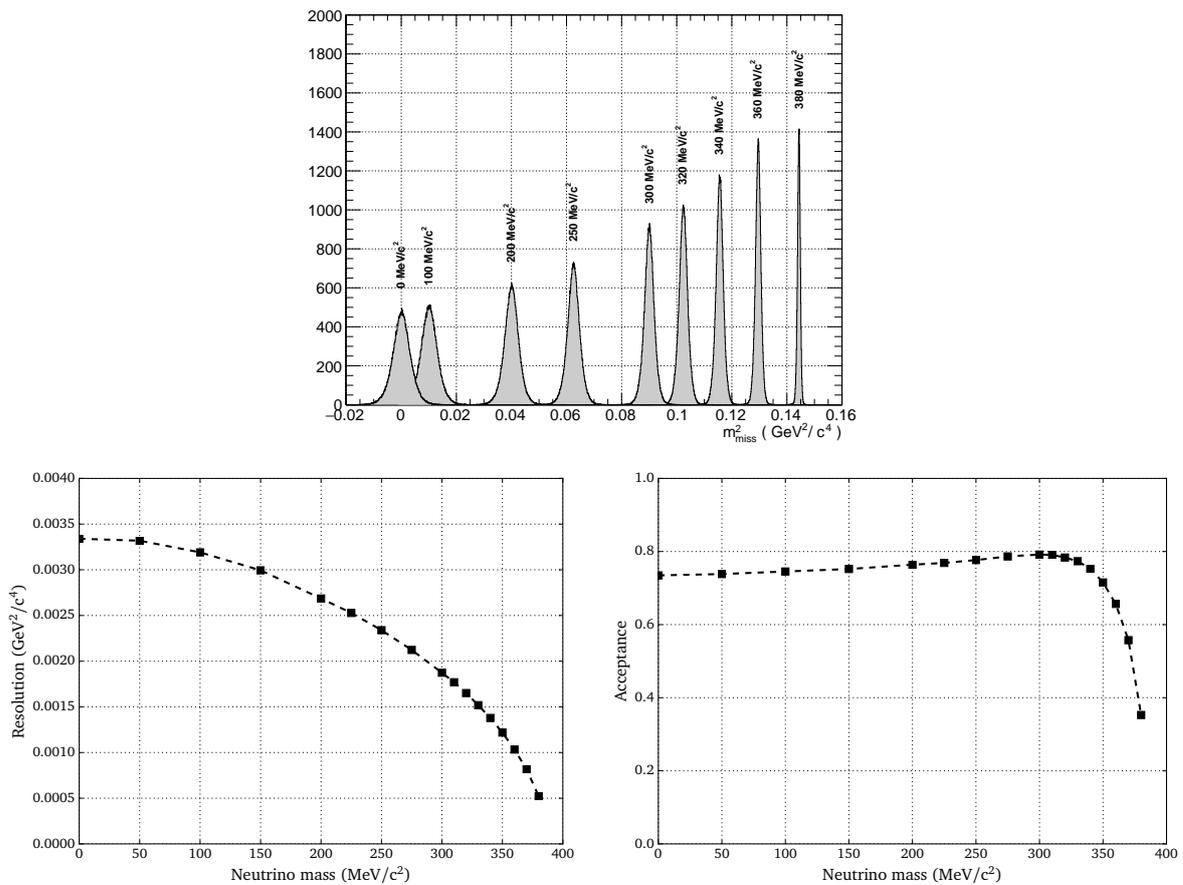


Figure 3.1.: Heavy neutrino signals. *top*: the expected distribution of reconstructed squared missing mass, for several different heavy neutrino masses (indicated by the label above each peak). *left*: resolution on reconstructed m_{miss}^2 as a function of neutrino mass. *right*: signal acceptance as defined in section 3.3.2

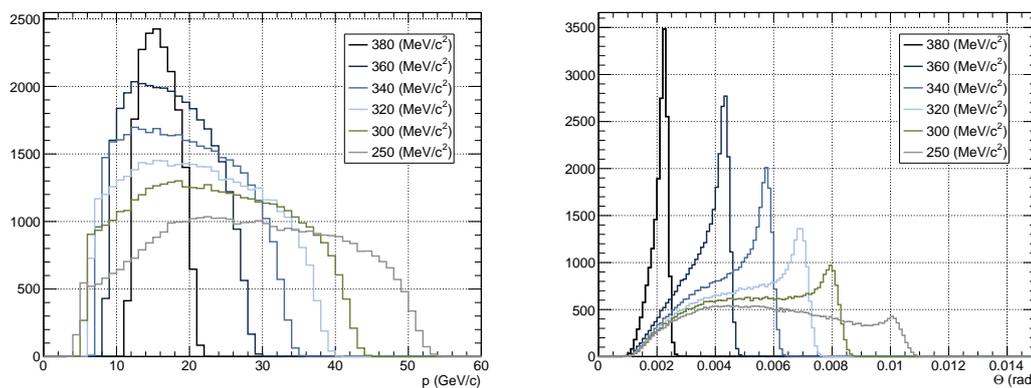


Figure 3.2.: Kinematic distributions of the muon produced in $K^+ \rightarrow \mu \nu_h$ decay. *left*: momentum, *right*: angle between track and nominal kaon trajectory.

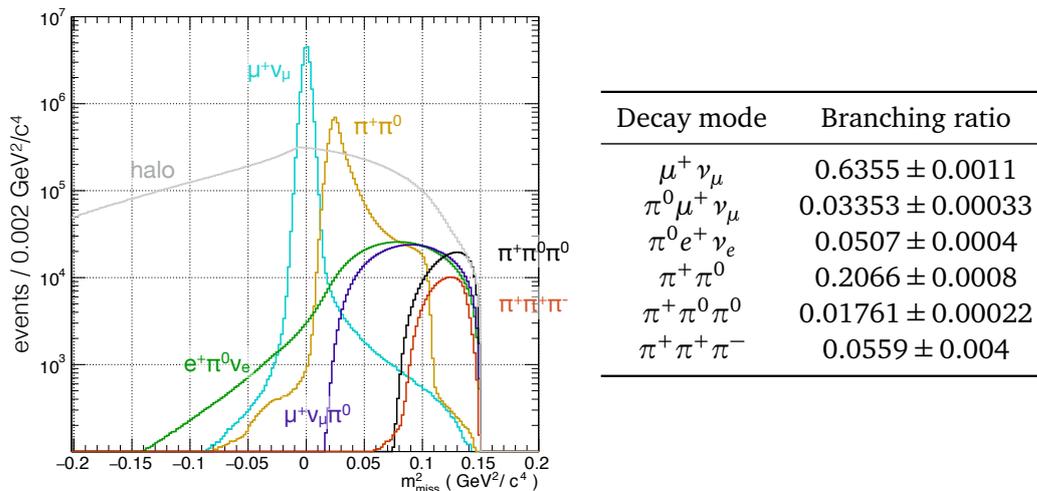


Figure 3.3.: The squared missing mass spectra of the six most common decay modes of K^+ when reconstructed assuming the muon mass for the charged particle detected. The MC samples are normalized according to the decay branching ratios [8], while the halo distribution is data driven and scaled to match the data. The distributions shown are the predicted spectra when only the single track pre-selection is applied.

$K^+ \rightarrow \mu^+ \nu_\mu (\gamma)$

Decay to a muon and muon neutrino is the most common decay mode of charged kaons, with branching ratio 0.6355 ± 0.0011 . Within the resolution of the experiment, the three Standard Model neutrinos all appear massless, so there is a single peak at zero missing mass with a width ($\Delta m^2 \approx 3 \times 10^{-3} \text{ GeV}^2/c^4$), determined by the width of the kaon momentum spectrum and the resolution of the spectrometer. The contribution from heavy neutrinos can only be resolved if their mass lies outside this peak, so its width determines the minimum heavy neutrino mass which could be observed.

Kaon decay to a muon and neutrino can be accompanied by photon emission in one of two ways. The first is Inner Bremsstrahlung (IB), in which the photon is emitted from the incoming kaon or the outgoing muon; this represents a radiative correction which can be calculated in QED. The second contribution is the Structure Dependent part, in which the photon is emitted from internal hadronic states, so it is sensitive to the structure of the kaon. This contribution is divided into two parts (SD^+ and SD^-), associated with linear combinations of the vector and axial kaon form factors ($F_V \pm F_A$). In addition, there can be constructive and destructive interference (INT^\pm) between the IB and SD

components. The decay width of each of the contributions can be calculated from the phase space integral [93]:

$$\Gamma_I = \int dx dy \rho_I(x, y, F_A, F_V), \quad I = SD^\pm, INT^\pm \quad (3.6)$$

where, ρ_I are the Dalitz plot densities, x and y are related to the energy of the photon and the muon in the kaon rest frame:

$$x = 2E_\gamma/m_K, \quad y = 2E_\mu/m_k \quad (3.7)$$

and F_A, F_V are form factors which can be calculated in chiral perturbation theory or the Light Front Quark Model [94]. The predictions on the branching ratios vary by 5% between the various models. For this analysis, the branching ratios used are derived from the $O(p^6)$ chiral perturbation theory described in [94]. The results are:

$$\mathcal{B}(SD^+) = 8.73 \times 10^{-6}, \quad \mathcal{B}(SD^-) = 1.92 \times 10^{-6}, \quad (3.8)$$

$$\mathcal{B}(INT^+) = 1.47 \times 10^{-5}, \quad \mathcal{B}(INT^-) = -4.61 \times 10^{-5}, \quad (3.9)$$

where the negative sign for $\mathcal{B}(INT^-)$ indicates that there is the destructive interference between the IB and SD contributions so the simulated events should be *subtracted* from the total background distribution.

$K^+ \rightarrow \mu^+ \pi^0 \nu_\mu$

The π^0 from $K^+ \rightarrow \mu^+ \pi^0 \nu_\mu$ decays immediately ($\tau_{\pi^0} \sim 10^{-16}$ s) almost always (99%) into two photons. If the photons are not detected, then the experimental signature is the same as for the signal decay. The reconstructed neutrino mass under the assumption of the signal decay now depends on the pion momentum:

$$m_{\text{miss}}^2 = (p_K - p_\mu)^2 = (p_\nu + p_\pi)^2. \quad (3.10)$$

The range of pion momenta allowed by three body kinematics constrains this background to the region:

$$m_{\pi^0}^2 < m_{\text{miss}}^2 < (m_K - m_\mu)^2, \quad (3.11)$$

$$0.018 \text{ GeV}^2/c^4 < m_{\text{miss}}^2 < 0.154 \text{ GeV}^2/c^4. \quad (3.12)$$

$K^+ \rightarrow \pi^+ \pi^0$

The hadronic decay $K^+ \rightarrow \pi^+ \pi^0$ is only reconstructed as a signal event if both photons from the π^0 decay are missed and the π^+ is misidentified as a muon (or decays into a muon). If \vec{p}_{π^+} is measured in the spectrometer before it decays, but the pion is reconstructed under the muon hypothesis, then the reconstructed missing mass depends on the pion momentum in the laboratory frame (determined by the angle of emission in the kaon rest frame)

$$m_{\text{miss}}^2 = m_{\pi^0}^2 + m_{\pi^+}^2 - m_{\mu}^2 + 2E_K |\vec{p}_K| \left[\left(1 + \frac{m_{\pi^+}^2}{\vec{p}_{\pi^+}^2} \right)^{\frac{1}{2}} - \left(1 + \frac{m_{\mu}^2}{\vec{p}_{\pi^+}^2} \right)^{\frac{1}{2}} \right]. \quad (3.13)$$

Alternatively, if the pion decays into a muon before it reaches the spectrometer, the missing mass evaluated is:

$$m_{\text{miss}}^2 = (p_K - p_{\mu})^2 = (p_{\pi^0} - p_{\nu})^2 = m_{\pi^0}^2 - E_{\nu}(E_{\pi^0} - p_{\pi^0} \cos \theta_{\pi\nu}). \quad (3.14)$$

The result depends on both the initial direction of the π^0 and the subsequent direction of the ν_{μ} so the range of reconstructed masses is broader than for the case in which the π^+ momentum is measured before it decays.

 $K^+ \rightarrow \pi^+ \pi \pi$

The backgrounds due to decays to three pions are naturally suppressed because they involve either three tracks ($K^+ \rightarrow \pi^+ \pi^+ \pi^-$) or four photons ($K^+ \rightarrow \pi^+ \pi^0 \pi^0$) so they can be easily rejected. The events which do pass the tracking and veto selections typically appear at the upper end of the missing mass spectrum because m_{miss}^2 really corresponds to two pions.

 $K^+ \rightarrow e^+ \nu_e (\pi^0)$

Decays with electrons in the final state are typically easy to reject with particle identification. $K^+ \rightarrow e^+ \nu_e$ would appear at negative m_{miss}^2 and its branching ratio is very small due to helicity suppression. $K^+ \rightarrow e^+ \pi^0 \nu_e$ would appear at higher values of m_{miss}^2 but can also be rejected by detecting the photons from the π^0 , so it is not a significant background contribution.

Beam halo

There is an additional background contribution which does not come from kaon decays in the decay volume but instead comes from the beam halo: muons produced by kaon and pion decay upstream of the decay volume. These muons pass through the final collimator and into the decay volume, where they can only be distinguished from muons produced by fiducial kaon decays by their kinematic distributions.

3.3.4. Background suppression

The backgrounds described in the previous section motivate three approaches to distinguishing signal and background. They are outlined briefly here, and explored further in the following sections.

Firstly, a photon veto is required to identify and reject events with π^0 s in the final state. The LKr calorimeter is used to veto photons. Its performance has been evaluated in previous NA48 and NA62 (2007 data) analyses (for example in the reconstruction of $K^+ \rightarrow \pi^+ \pi^0 \pi^0$ [82]) and the MC simulation of the detector response has been tuned to the 2007 data. The reconstruction of photons is discussed in section 3.4.3.

Secondly, particle identification is required in order to distinguish muons, which are expected in the signal decay, from electrons and pions which appear in background decay products. The LKr could also be used to distinguish electrons from muons, by the amount of energy deposited in the LKr relative to the momentum measured in the spectrometer: $E_{\text{LKr}}/|\vec{p}_{\text{DCH}}|$ (shortened to E/p where there is no risk of ambiguity). Electrons, having low mass, will deposit most of their energy in an electromagnetic shower ($E/p \approx 1$) while muons act like Minimum Ionising Particles (MIPs) and typically deposit only a small fraction of their energy in the calorimeter through ionisation ($\sim 99\%$ of muons have $E/p < 0.1$). However, the showering of charged pions in the LKr is not well controlled in MC, so the background from $K^+ \rightarrow \pi^+ \pi^0$ is difficult to simulate reliably when cuts are applied on the value of E/p . Instead, muons are identified using the muon veto, and no requirements are placed on E/p in the final selection. The muon veto is studied in more detail in section 3.6.

Finally, for the beam halo which cannot be distinguished from the signal mode by particle identification, a kinematic selection is required to take advantage of the different distributions of muons associated with heavy neutrinos and muons from the beam halo. The modelling and suppression of the halo background is described in section 3.7.

3.4. Simulation and corrections

In this section and the following three sections, the main components of the analysis are presented along with studies to demonstrate that the signal and backgrounds are well understood and can be modelled accurately. There is a brief description of the MC simulation software used for NA48/2 and NA62 in 2007, and how it is tuned to match the measured properties of the beam, the LKr calorimeter and the magnetic spectrometer. New studies of the far tails of the spectrometer resolution and the muon veto efficiency have been performed and these are described in detail. Finally a study is made of the beam halo so that its contribution to the background in the signal region can be minimized.

3.4.1. Physics simulation

MC simulations of the backgrounds from kaon decay modes were produced by the CMC software package [95], designed for NA48/2 and the R_K phase of NA62. It uses GEANT3 [96] for the simulation of particle propagation and reconstruction in the detector, with additional simulation of the detector response, tuned to data. Separate simulations were performed for each run to account for time varying quantities such as the beam momentum and bending magnet polarities, with sample sizes proportional to the number of kaon decays recorded in each run. The simulated detector responses for MC were stored in the same format as the experimental data (*COMPACT*) so the same analysis could be applied to both data and MC.

During the analysis phase a number of corrections to the MC are applied to improve the agreement between MC and data. Conversely, some corrections applied to the data are not required in MC.

3.4.2. Beam

During data taking, the average kaon 3-momentum was measured on a run by run basis, using the decay $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ which can be reconstructed entirely from the decay products. The typical beam momentum is $|\vec{p}| = 74 \text{ GeV}/c$ with an event by event standard deviation of $2 \text{ GeV}/c$. The beam is deflected in the x direction on entry into the decay region in order to counter the effect of the spectrometer magnet. When the polarity of the spectrometer magnet is positive, the initial kaon direction is $dx/dz = 2.4 \times 10^{-4}$. When the polarity is reversed, the initial kaon direction is $dx/dz = -2.0 \times 10^{-4}$.

When reconstructing the squared missing mass $(p_K - p_\mu)^2$ for a single event, the kaon 3-momentum is taken to have the average value obtained for the run under consideration.

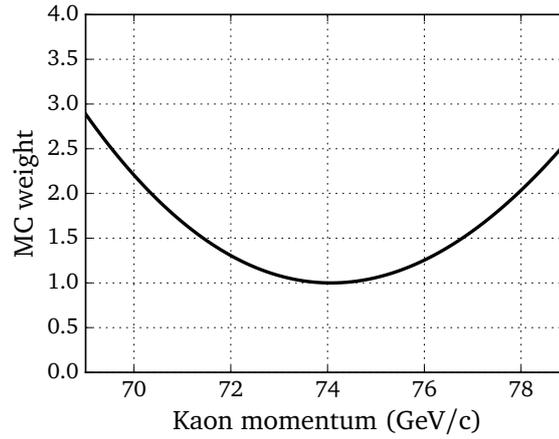


Figure 3.4.: The weights applied to MC events in run 20410 as a function of simulated kaon momentum in order to improve the width of the kaon momentum distribution.

In MC, the appropriate central momentum, $|\vec{p}_K|$ as well as the direction ($dx/dz, dy/dz$) was simulated for each run. The width of the kaon momentum spectrum is incorrectly reproduced in MC. To compensate for this, weights are applied to individual events as a function of the MC Truth kaon momentum:

$$w = 1 + \lambda_r (|\vec{p}_K| - 74 \text{ GeV}/c)^2, \quad (3.15)$$

where λ_r is a run dependent slope which varies over time from -0.08 to $+0.08$. Additional corrections are made to fine tune the spectrum in certain momentum ranges. The additional weights have the form:

$$w = 1 + \mu_r (|\vec{p}_K| - p_{0r}), \quad (3.16)$$

where the values of μ_r and p_{0r} are a function of run and momentum range.

Figure 3.6 shows, as an example, the weights applied in MC for run 20410, and figure 3.5 shows the effect on the reconstructed missing mass spectrum of applying these additional weights. The improvement is not significant but the impact of the poorly controlled shape on the present analysis is negligible as the signal region in the squared missing mass spectrum is well separated from the peak at zero mass.

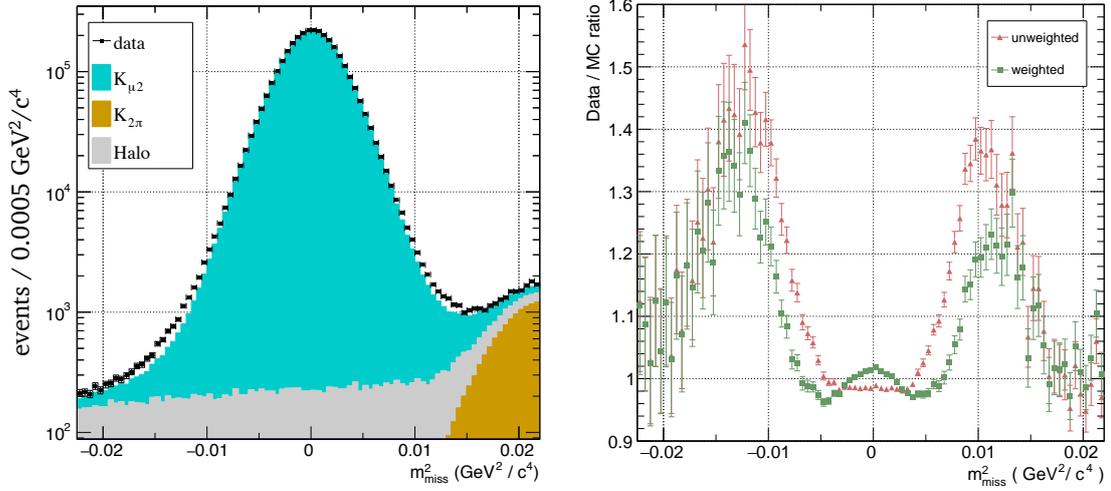


Figure 3.5.: The effect of weighting the MC samples according to the kaon momentum. Figure (a) shows the data and MC backgrounds in the region around the peak after applying the weighting, while (b) shows the data/MC ratios with and without weighting.

3.4.3. Liquid Krypton Calorimeter

Several corrections are applied to the cluster energies obtained from the LKr calorimeter. For data (not MC), there is a run dependent scale correction $E_{\text{corr}} = E_0(1 + C_{\text{run}})$, where C_{run} ranges from -0.4% to 0.9%) as well as a non-linearity correction $E_{\text{corr}} = E_0(1 + f(E_0))$ for clusters with energies less than 11 GeV ($f(11 \text{ GeV}) = 0$, $f(0 \text{ GeV}) \simeq 2\%$). Additionally, a cell by cell correction (rms 0.004) is applied to every cluster.

In both data and MC, cluster positions must be corrected to account for the projectivity of the LKr detector, which was introduced in the design to facilitate measurement of the calorimeter's relative energy scale. The cell towers diverge from a point $\Delta z_{\text{LKr}} = 110 \text{ m}$ in front of the calorimeter so that the angle between a pair of photons originating from a known distance upstream could be determined independently of longitudinal fluctuations. In order to determine the x, y position of a cluster from the IDs of the cells in which energy was deposited, it is necessary to know the depth at which the shower was at its maximum. The cluster depth (distance in z from front face of the LKr to the shower maximum) is estimated from the total energy deposited:

$$z_{\text{depth}} = k_z + k_E \ln \frac{E_{\text{clus}}}{E_0}, \quad (3.17)$$

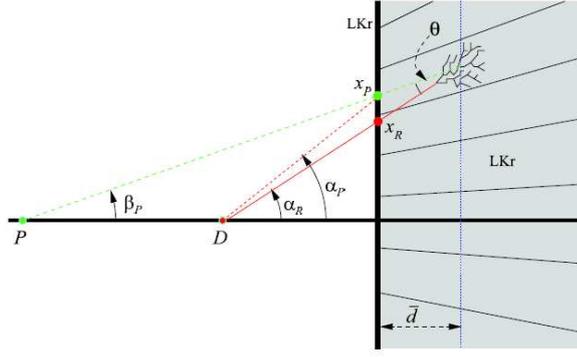


Figure 3.6.: The projective structure of the LKr calorimeter cell towers, showing how the cluster position must be corrected [97]. x_p and x_r in the figure correspond to x_0 and x in equation 3.19.

where, $E_0 = 1$ GeV, $k_E = 4.3$ cm and

$$k_z = \begin{cases} 16.5 \text{ cm} & \text{for electrons} \\ 20.8 \text{ cm} & \text{for photons} \end{cases} \quad (3.18)$$

The x , y and z coordinates of the cluster, used for future calculations are then given by:

$$z = z_{\text{LKr}} + z_{\text{depth}}, \quad x = x_0 \left(1 + \frac{z_{\text{depth}}}{\Delta z_{\text{LKr}}} \right), \quad y = y_0 \left(1 + \frac{z_{\text{depth}}}{\Delta z_{\text{LKr}}} \right) \quad (3.19)$$

where x_0 and y_0 are the coordinates of the LKr cell at the front face of the detector.

A final coordinate transformation is required because the LKr coordinate system $x_{\text{LKr}}, y_{\text{LKr}}$, is not exactly aligned with the coordinate system used in the spectrometer, $x_{\text{DCH}}, y_{\text{DCH}}$.

For data, the transformation is:

$$\begin{pmatrix} x_{\text{DCH}} \\ y_{\text{DCH}} \end{pmatrix} = \begin{pmatrix} 0.136 \text{ cm} \\ 0.300 \text{ cm} \end{pmatrix} + \begin{pmatrix} 1 & 0.87 \times 10^{-3} \\ -0.87 \times 10^{-3} & 1 \end{pmatrix} \begin{pmatrix} x_{\text{LKr}} \\ y_{\text{LKr}} \end{pmatrix} \quad (3.20)$$

In MC, the LKr and spectrometer coordinate systems are already aligned except for a small shift in x , so the transformation is:

$$\begin{pmatrix} x_{\text{DCH}} \\ y_{\text{DCH}} \end{pmatrix} = \begin{pmatrix} -0.013 \text{ cm} \\ 0 \end{pmatrix} + \begin{pmatrix} x_{\text{LKr}} \\ y_{\text{LKr}} \end{pmatrix} \quad (3.21)$$

3.4.4. Spectrometer corrections

Measurements of track momentum in the spectrometer were corrected offline in order to account for the change in the field integral of the dipole bending magnet when the polarity is switched, and for misalignment of the drift chambers. The magnetic field correction (parameterized by β) has the same sign for all tracks, while the correction required for the chamber misalignment (parameterized by α) depends on the charge of the track so the correction has the form:

$$|p_{\text{corr}}| = |p_0|(1 + \beta)(1 + \alpha|p_0|q). \quad (3.22)$$

The required values of α and β for each data taking run were determined by reconstructing the invariant mass in the decay $K \rightarrow \pi^+ \pi^+ \pi^-$ and comparing it with the nominal kaon mass from the PDG. The signs of both α and β depend on the spectrometer polarity. α was $\pm 3.9 \times 10^{-5}$ for all runs while β varied between runs. β had a magnitude $\mathcal{O}(10^{-3})$ and rms variation $\mathcal{O}(10^{-4})$. The tabulated run by run values were available in a database for use in the present analysis.

Once the momentum at the spectrometer has been correctly determined, the track must be extrapolated back to the decay vertex, taking into account the residual magnetic field in the decay volume, in order to determine the momentum at the decay point. The magnetic field along the z -axis is shown in figure 3.7.

First, a zero order approximation to the decay vertex is computed as the point of closest approach between the kaon track (defined by database parameters for the current run) and the raw track definition from the spectrometer. The residual magnetic field is then integrated over the length of the track, in order to determine an improved estimate of the vertex location and track direction.

3.5. Spectrometer resolution tails

Since the heavy neutrino analysis involves the far tails of the $K^+ \rightarrow \mu^+ \nu_\mu$ spectrum, it is important to understand the spectrometer resolution, especially the effect of multiple Coulomb scattering on the far tails, which is not necessarily well reproduced in GEANT3. The LKr calorimeter can be used to cross check the spectrometer resolution using the decay $K^+ \rightarrow \pi^+ \pi^0$.

In this section, a selection is described to obtain a control sample of $K^+ \rightarrow \pi^+ \pi^0$ events and several kinematic variables are plotted to compare the data and MC. An empirical approach to correcting the MC is proposed, based on the expected distribution of the tails of multiple Coulomb scattering.

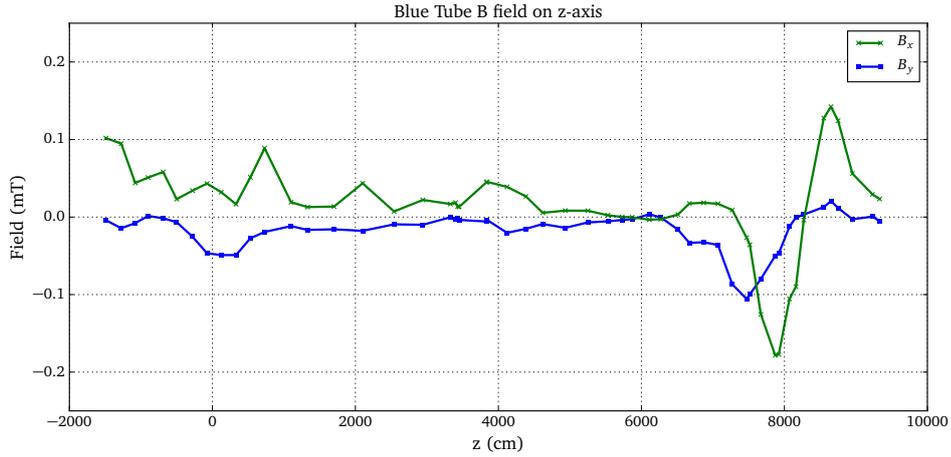


Figure 3.7.: Components of the residual magnetic field at $x, y = 0$ along the z -axis of the decay region

The modified MC is shown to overestimate the tails of the missing mass spectrum, so the size of the modification gives an upper limit on the systematic error which should be included in the final analysis.

The decay $K^+ \rightarrow \pi^+ \pi^0$ is kinematically closed so the four-momentum of the π^+ can be reconstructed once the K^+ and π^0 momenta are determined. The π^0 can be reconstructed by detecting a pair of photons in the LKr calorimeter and the kaon momentum can be estimated from the run by run average as described in section 3.4.2. This means it is possible to select a sample of $K^+ \rightarrow \pi^+ \pi^0$ decays without using the spectrometer at all, and use it to improve the MC simulation of the spectrometer performance.

The selection is divided into three parts. First, the LKr is used to perform a cut based selection to find events with a topology consistent with $K^+ \rightarrow \pi^+ \pi^0$. Then, a kinematic fit is constructed from the photon clusters and the average kaon momentum, and a cut on the fit χ^2 is used to reject events which do not match the topology and kinematics of $K^+ \rightarrow \pi^+ \pi^0$. Finally, a standard single track selection is performed on the spectrometer data, independently of the fit. The resulting sample can be used to study the drift chamber resolution either by comparing the measured π^+ momentum with the kinematic fit value or by studying the distribution of variables reconstructed only using the spectrometer, such as the transverse momentum or the missing mass of the π^0 .

The selection of $K^+ \rightarrow \pi^+ \pi^0$ events is not entirely independent of the spectrometer as the track time is used as a reference and the track direction is used to determine which cluster in the LKr, if any, should be attributed to the π^+ . However, the measured π^+ momentum is not used in the kinematic

fit so using the pion track does not have a large biasing effect on the selection.

In this section, several ‘squared missing mass’ quantities are referred to. $(m_{\text{miss}}^{\pi^+})^2 = (p_K - p_{\pi^+})^2$ is computed from the measured π^+ momentum and for $K^+ \rightarrow \pi^+ \pi^0$ decays it peaks at $m_{\pi^0}^2$. $(m_{\text{miss}}^{\pi^0})^2 = (p_K - p_{\pi^0})^2$ is computed from the measured π^0 momentum (reconstructed from photons) and for $K^+ \rightarrow \pi^+ \pi^0$ decays it peaks at $m_{\pi^+}^2$.

3.5.1. LKr selection

The following cuts are applied before performing a kinematic fit:

- Bad bursts for Physics, DCH, LKr, and MUV are removed. These are bursts during which one or more of the various subdetector performances have been identified as sub-optimal due to, for example, malfunctioning components or problems in the readout chain.
- The NHOD (neutral hodoscope) trigger is used, instead of the $Q_1 \times 1\text{TRLK}$ trigger (see section 2.3.10), since it is independent of the spectrometer. It was downscaled online by a factor of 150.
- Tracks with associated muons in the MUV are rejected in order to suppress background from $K^+ \rightarrow \mu^+ \pi^0 \nu_\mu$ and secondary $\pi^+ \rightarrow \mu^+ \nu_\mu$ decays.
- Clusters in the LKr are ignored if they are more than 12 ns out of time with respect to the track time given by the spectrometer.
- There should be two or three in-time clusters with energies greater than $10 \text{ GeV}/c^2$: two corresponding to the photons from the π^0 decay and, optionally, one corresponding to the π^+ track.
- The track cluster, if any, is defined as the cluster which is closest to the extrapolated track position at the front face of the LKr and within 40 cm of it. The cut on the track-cluster separation is required to be loose because the pion can produce a shower with a large shower and this process is poorly described in MC.
- The photon clusters are the remaining two clusters with energies greater than $10 \text{ GeV}/c^2$ (or the only two if no cluster was associated with the pion).
- The photon clusters should be separated from the track cluster (if there is one) by at least 30 cm and from each other by at least 20 cm.

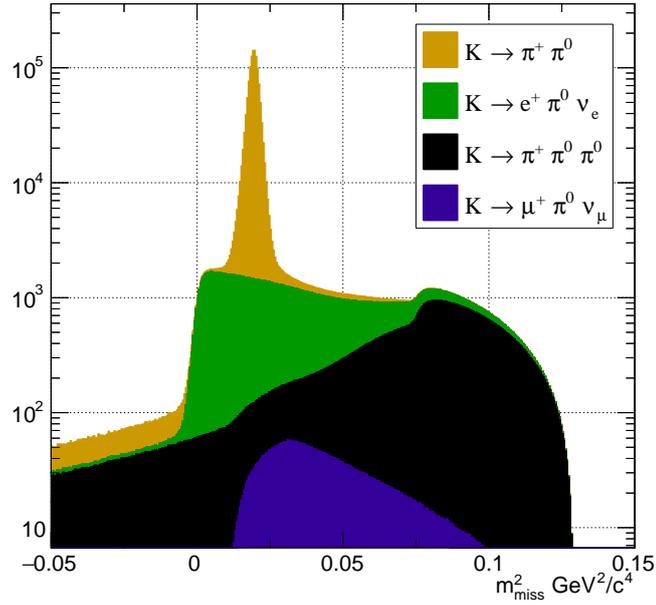


Figure 3.8.: MC spectrum of $(m_{\text{miss}}^{\pi^0})^2 = (p_K - p_{\pi^0})^2$ for $K^+ \rightarrow \pi^+ \pi^0$ decays, using the LKr calorimeter for selection and reconstruction.

- After the kinematic fit (see section 3.5.2), the reconstructed photon trajectories should have a large enough transverse component that they are more than 15 cm away from the beam axis when they are level with the first drift chamber. This is to minimize the effect of photons interacting with the beam pipe.

Clusters with energies less than $10 \text{ GeV}/c^2$ are ignored when determining which clusters should be associated with photons, as described above. However, after the photon clusters have been assigned, if no pion cluster has been assigned but there is a low energy ($< 10 \text{ GeV}/c^2$) cluster within 30 cm of the extrapolated track position, this low energy cluster is defined as the pion cluster for the purposes of E/p studies.

Figure 3.8 shows the MC spectrum of the reconstructed squared missing mass, $(m_{\text{miss}}^{\pi^0})^2 = (p_K - p_{\pi^0})^2$, for MC simulation of the events which pass the selection outlined above. The $K^+ \rightarrow \pi^+ \pi^0$ signal appears as a peak at the nominal value of $m_{\pi^+}^2$, while the backgrounds come from $K^+ \rightarrow e^+ \pi^0 \nu_e$ with the positron misidentified as the π^+ ; $K^+ \rightarrow \pi^+ \pi^0 \pi^0$ with two missed photons; and $K^+ \rightarrow \mu^+ \pi^0 \nu_\mu$ events in which the muon was not detected in the MUV. (The procedure for including the muon veto inefficiency in the MC simulation is described in section 3.6.)

3.5.2. Kinematic fit

A kinematic fit is performed on the kaon and two photon clusters, in order to reconstruct the momenta of all the particles in the decay. The inputs to the fit are: the photon cluster energies and positions, the kaon three-momentum (p_x, p_y, p_z) , and the kaon position (x_0, y_0) in the xy plane at $z = 0$.

The kaon parameters are taken from the database of run by run averages, measured with $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ decays. The errors are taken as the standard deviation of the event by event variation, corresponding to the beam spot size and the angular divergence:

$$\sigma_{p_x, p_y} = 1 \times 10^{-3} \text{ GeV}/c, \quad \sigma_{p_z} = 1.5 \text{ GeV}/c, \quad \sigma_{x_0, y_0} = 0.3 \text{ cm}, \quad (3.23)$$

where p_x, p_y, p_z are the three components of the kaon momentum and x_0, y_0 are the kaon position at $z = 0$.

The resolution on the cluster positions is energy dependent:

$$\sigma_{x,y} = \frac{0.42}{\sqrt{E}} \oplus 0.06, \quad (3.24)$$

where x, y are the cluster coordinates at the front face of the LKr calorimeter, and the energy resolution is given by (in units of GeV and cm) [87]

$$\frac{\sigma_E}{E} = \frac{0.032}{\sqrt{E}} \oplus \frac{0.09}{E} \oplus 0.0042. \quad (3.25)$$

In total, there are 29 unknown quantities used in the reconstruction of the $K^+ \rightarrow \pi^+ \pi^0$ event. These are: the four-momenta (E, \vec{p}) of the K, π^+, π^0 and two photons ($5 \times 4 = 20$ unknowns); the coordinates, (x, y, z) , of the kaon decay vertex (3 unknowns); the xy coordinates of the photon clusters in the LKr ($2 \times 2 = 4$ unknowns); and the kaon position at $z = 0$, denoted x_0, y_0 (2 unknowns).

The measured cluster positions and energies, along with the kaon momentum and position constitute 11 constraints. There are further kinematic constraints from the 5 particle masses (5 constraints) and the requirement of momentum conservation at both the K^+ and π^0 decay vertices (6 constraints). The photon momenta are fixed by the z -coordinate of the decay vertex and the cluster positions at the LKr (6 constraints) The x and y coordinates of the decay vertex are fixed by the z -coordinate of the decay vertex and the kaon track parameters taken from the database (2 constraints). The decay geometry is shown in figure 3.9.

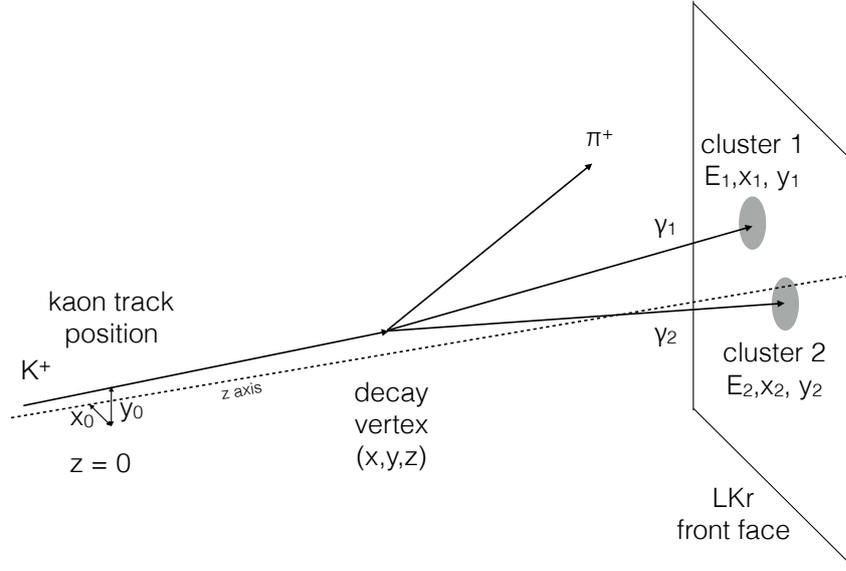


Figure 3.9.: The geometry of $K^+ \rightarrow \pi^+ \pi^0$ decay showing the variables used in the kinematic fit.

Since there are 11 measured quantities and 19 constraints but only 29 unknowns, every event can be reconstructed by a fit with 1 degree of freedom. The resulting χ^2 can be used in a cut to remove background events.

For a given set of fit parameters, χ^2 is computed as follows:

- The angle, θ , between the two photons is determined from their measured energies (E_1 and E_2), using the π^0 mass with:

$$\cos \theta = 1 - \frac{m_{\pi^0}^2}{2E_1 E_2}, \quad (3.26)$$

which follows from conservation of four-momentum:

$$p_{\pi^0} = p_1 + p_2. \quad (3.27)$$

- The z coordinate of the vertex is computed numerically such that the vectors joining the computed vertex position to the measured cluster positions have an angle θ between them which satisfies equation 3.26.
- The photon four momenta are constructed from their energies and directions. The π^0 four momentum is the sum of the two photon four momenta.
- The π^+ four momentum is computed from $p_{\pi^+} = p_{K^+} - p_{\pi^0}$. (The four components of $p_{\pi^+}^{\text{fit}}$ are

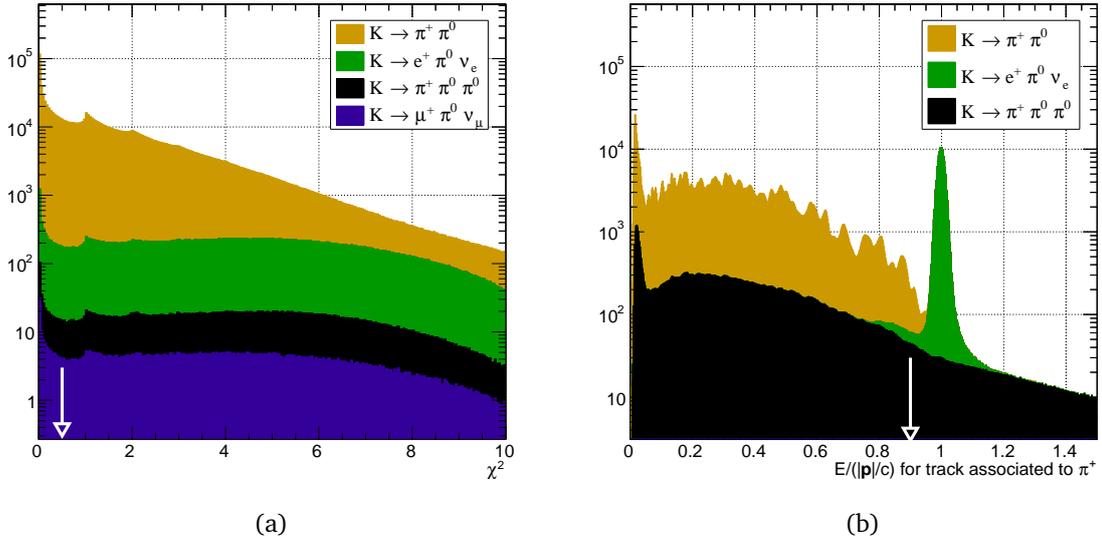


Figure 3.10.: MC distributions of (a) fit χ^2 and (b) E/p for events reconstructed in the LKr calorimeter. (The cuts applied are shown by white arrows.)

independent at this point, so in general $m_{\pi^+}^{\text{fit}} = (p_{\pi^+}^{\text{fit}})^2 \neq m_{\pi^+}$.

- The χ^2 value for this set of fit parameters is defined by

$$\chi^2 = \sum_i \frac{(X_i - X_i^0)^2}{\sigma_{X_i}^2} + 10^8 (m_{\pi^+}^{\text{fit}} - m_{\pi^+})^2, \quad (3.28)$$

where i runs over the 11 measurements, X_i^0 is the measured value and X_i is the fit value. The final term is a penalty factor to impose the π^+ mass constraint - the large value is chosen ad hoc to ensure that the constraint is enforced without causing problems related to machine precision and the precision on the known value of m_{π^+} . Add this term to the chi-squared.

For each event, the fit parameters are determined by minimizing the value of χ^2 defined in equation 3.28. The minimization is performed by the ROOT *Minimizer* library, which implements the MINUIT *MIGRAD* algorithm. After the fit has been performed, events which are not a good match for the $K^+ \rightarrow \pi^+ \pi^0$ hypothesis are rejected with a tight cut requiring $\chi^2 < 0.5$ to minimize background in the kinematic distributions produced with the selected events.

3.5.3. Charged track selection

Once a sample of $K^+ \rightarrow \pi^+ \pi^0$ events has been selected with the LKr calorimeter as described above, a standard single track selection is performed on the data from the spectrometer:

- A single track pre-selection is performed as described in section 3.3.1, but initially a looser track definition is used: $3 < |\vec{p}_{\pi^+}| < 80 \text{ GeV}/c$; $-50 \text{ m} < z_{\text{vertex}} < 100 \text{ m}$; $\text{CDA} < 15 \text{ cm}$. This decreases the acceptance for multi-track events because an extra track is less likely to be marked as ‘bad’ and ignored.
- Track vertex z coordinate must be in the range $-15 \text{ m} < z_{\text{vertex}} < 65 \text{ m}$
- Track charge must be +1.
- Track quality must be > 0.7 , where ‘quality’ corresponds to the fraction of hits in the DCH chambers which are close in time to the average for the event.
- Track radius at the first and last spectrometer drift chambers should satisfy:
 $12 \text{ cm} < r_{\text{DCH1}} < 115 \text{ cm}$; $14 \text{ cm} < r_{\text{DCH4}} < 115 \text{ cm}$ to ensure that the track is in the spectrometer acceptance and does not pass through the beam pipe at too shallow an angle.
- Track time should be within 20 ns of the hodoscope trigger time.
- Track CDA (the Closest Distance of Approach between the reconstructed and the nominal kaon track) should be $< 3.5 \text{ cm}$.
- Track momentum must be in the range $10 \text{ GeV}/c < p < 60 \text{ GeV}/c$, corresponding to the range used in the final signal selection.
- Clusters associated with the track are ignored if they have energy $E < 9 \text{ GeV}$, since the distribution of low energy clusters is not well simulated in the MC.
- If there is a cluster associated with the track, then the cluster energy, E , and the track momentum $|\vec{p}|$, should be such that the ratio $E/(|\vec{p}|c) < 0.9$, in order to reject electron tracks which have $E/(|\vec{p}|c) \sim 1$.

Figure 3.10b shows the simulated contributions from each kaon decay channel to the $E/(|\vec{p}|c)$ spectrum for reconstructed track clusters. The irregular peaks in the spectrum for $K^+ \rightarrow \pi^+ \pi^0$ arise because the electromagnetic shower in the LKr is not fully simulated for each event but is instead taken from

a ‘library’ of a limited number of pre-computed showers. The probability for a pion to have an energy deposition such that $E/(|\vec{p}|c) > 0.9$ has been determined in previous measurements so the effect of the $E/(|\vec{p}|c)$ cut on $K^+ \rightarrow \pi^+ \pi^0$ can be simulated separately, using the values from data.

3.5.4. Kinematic distributions

Once a sample of $K^+ \rightarrow \pi^+ \pi^0$ events has been selected using the LKr calorimeter, it can be used to study the spectrometer performance, in one of two ways.

For quantities such as the π^+ momentum \vec{p}_π^+ , the value obtained in the fit can be compared directly with the value measured in the spectrometer. This approach is only applicable when the fit resolution is well known and does not dominate the spectrometer resolution.

Alternatively, the spectrometer can be studied by reconstructing quantities which have a known true value in $K^+ \rightarrow \pi^+ \pi^0$ decays. In this approach the spectrometer resolution can be studied independently of kinematic fit values. The reconstructed missing mass $(m_{\text{miss}}^{\pi^+})^2 = (p_K - p_{\pi^+})^2$ should always be given by the π^0 mass so this is a good choice of variable to study. The transverse momentum of the π^+ with respect to the kaon momentum depends on the emission angle in the kaon rest frame but it has a maximum value when the pion momenta are perpendicular to the kaon direction. In this case

$$p_T = |\vec{p}| = \frac{1}{2m_K} [(m_K^2 - (m_{\pi^+} + m_{\pi^0})^2)(m_K^2 - (m_{\pi^+} - m_{\pi^0})^2)]^{\frac{1}{2}} \quad (3.29)$$

$$= 0.205 \text{ GeV}/c. \quad (3.30)$$

The π^+ transverse momentum spectrum has a sharp end point at this value, so tails in the distribution beyond this point can be used to study the spectrometer resolution.

Figure 3.11 shows the two distributions (missing mass w.r.t. the π^+ and track transverse momentum) reconstructed with the spectrometer. Each has a region, $(m_{\text{miss}}^{\pi^+})^2 < -0.03 \text{ GeV}/c^2$ and $p_T > 0.22 \text{ GeV}$, in which MC simulation shows that no background (non- $K^+ \rightarrow \pi^+ \pi^0$) events are expected, so these are good regions to study the spectrometer resolution. In both cases, the data points lie consistently above the MC simulation, an indication that the far tails of the multiple Coulomb scattering distribution of charged tracks in the spectrometer could be underestimated in MC.

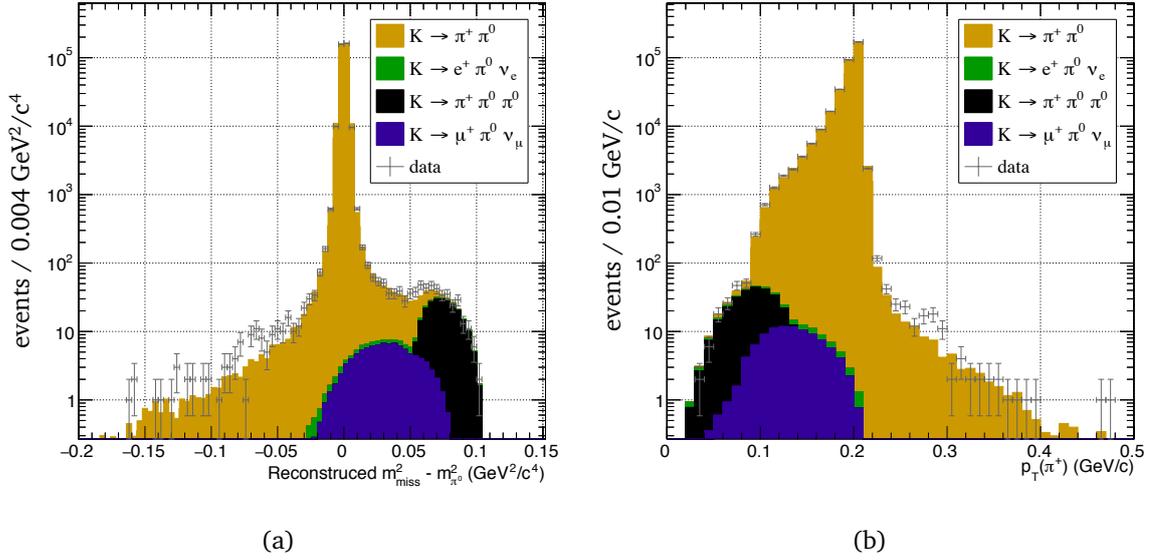


Figure 3.11.: Distributions of (a) $(m_{\text{miss}}^{\pi^+})^2 - m_{\pi^0}^2$ and (b) track p_T with respect to kaon momentum, reconstructed using the magnetic spectrometer, for $K^+ \rightarrow \pi^+ \pi^0$ candidates selected using the LKr calorimeter

3.5.5. Spectrometer resolution correction

Additional scattering can be introduced into the MC simulation of the magnetic spectrometer, in order to improve the data/MC agreement in the kinematic distributions described in the previous section. Scattering of charged particles as they pass through the spectrometer affects both the measurement of the direction of the upstream track and the measurement of the track momentum. Although the two effects have the same source, they are treated separately in this empirical approach.

The direction of the charged track upstream of the bending magnet is defined by two directions, $\theta_x = p_x/p_z$ and $\theta_y = p_y/p_z$, where $(p_x, p_y$ and p_z are the components of the track momentum). Multiple Coulomb scattering introduces a momentum dependent smearing of the direction measurement

$$\sigma_{\theta_{x,y}}^{\text{MCS}} = \frac{\text{const}}{|\vec{p}|}. \quad (3.31)$$

The magnitude of the track momentum, $|\vec{p}|$, is determined from the angular deflection, ϑ , of the track in the xz plane by the dipole bending magnet:

$$|\vec{p}| = \frac{qBL}{\vartheta}, \quad (3.32)$$

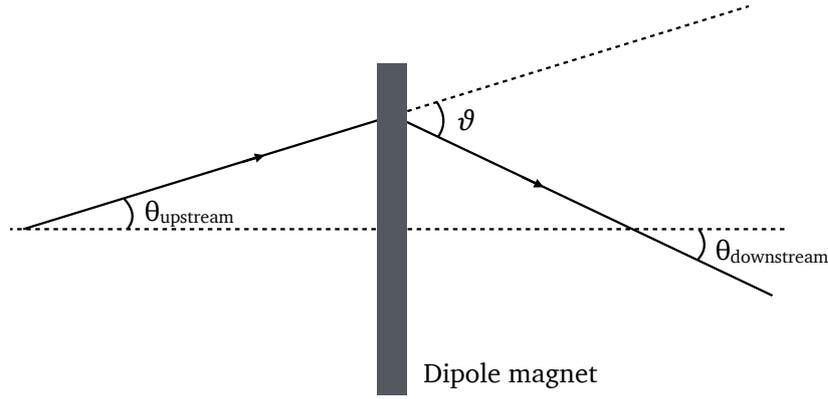


Figure 3.12.: Angular deflection, ϑ , determined from upstream and downstream track directions

where q is the particle charge, B is magnitude of the magnetic field and L is the length of magnet and $\vartheta = (p_x^{\text{downstream}} - p_x^{\text{upstream}})/p_z$ (see figure 3.12). Hence the multiple Coulomb scattering contribution to the momentum resolution is proportional to the contribution to the angular resolution of the spectrometer:

$$\frac{\sigma_{|\vec{p}|}^{\text{MCS}}}{|\vec{p}|} = \frac{\sigma_{\vartheta}^{\text{MCS}}}{\vartheta} = \sigma_{\vartheta}^{\text{MCS}} \cdot \frac{|\vec{p}|}{qBL} \quad (3.33)$$

Since the smearing of the angular resolution due to multiple Coulomb scattering, $\sigma_{\vartheta}^{\text{MCS}}$ is proportional to $1/|\vec{p}|$, the result is that the contribution from multiple Coulomb scattering to the fractional momentum resolution is independent of momentum:

$$\frac{\sigma_{|\vec{p}|}^{\text{MCS}}}{|\vec{p}|} = \text{const.} \quad (3.34)$$

In summary, additional multiple Coulomb scattering should be introduced into the simulation to affect the resolution of the track direction, θ_x, θ_y , with standard deviation $\sigma_{\theta_{x,y}} \propto 1/|\vec{p}|$ and to affect the resolution of the track momentum $|\vec{p}|$ with an standard deviation $\sigma_{\theta_{|\vec{p}|}} \propto |\vec{p}|$.

The central 98% of the angular distribution from multiple Coulomb scattering is well described by a Gaussian approximation [8], so the pdf at fixed momentum for scattering by an amount $\Delta\theta$ is given by:

$$P(\Delta\theta) = \frac{1}{\sigma_{\theta}^{\text{MCS}} \sqrt{2\pi}} \exp \left\{ -\frac{1}{2} \left(\frac{\Delta\theta}{\sigma_{\theta}^{\text{MCS}}} \right)^2 \right\} \quad (3.35)$$

and for the momentum

$$P(\Delta p) = \frac{1}{\sigma_p^{\text{MCS}} \sqrt{2\pi}} \exp \left\{ -\frac{1}{2} \left(\frac{\Delta p}{\sigma_p^{\text{MCS}}} \right)^2 \right\}, \quad (3.36)$$

where both $\sigma_\theta^{\text{MCS}}$ and σ_p^{MCS} are momentum dependent (from equations 3.31 and 3.34). However, infrequent ‘‘hard’’ scatters produce non-Gaussian tails, tending to the $1/(\Delta\theta)^4$ distribution of Rutherford scattering. In order to populate the tails of the kinematic distributions which are underestimated in MC, additional scattering is applied to a small fraction of events to improve the data/MC agreement in the p , θ_x and θ_y distributions. Empirically, it was found that using a $1/x^3$ distribution produced a better fit to the data, in all three distributions, than a $1/x^4$ distribution.

The same mechanism is used to generate ‘kicks’ according to a $1/x^3$ distribution for all three variables. A fraction, f , of events is selected and for each of these events, an additional ‘kick’ is applied to populate the far tails of the scattering distribution. The kick is determined by drawing a random number ζ from the pdf:

$$P(\zeta) = \tanh \left[\left(0.983 \frac{\zeta}{\zeta_0} \right)^5 \right] \cdot \frac{1}{\zeta^3}, \quad (3.37)$$

$P(\zeta)$ tends towards $1/\zeta^3$ for large values as desired. The tanh factor ensures that $P(\zeta)$ tends to zero at small values so events are preferentially scattered into the far tails. ζ_0 controls the cut-off point at which the distribution starts tending to zero; the factor 0.983 is included for convenience so that the function peaks at ζ_0 . Figure 3.13 shows how using this distribution increases the number of events in the far tails of the distribution without significant impact on the central Gaussian part.

The extra scattering contributions are described by two parameters for the magnitude of the momentum, f_p and ζ_{0p} , and two values for the track direction, $f_{x,y}$ and $\zeta_{0x,y}$. The four parameters were chosen to ensure good agreement between data and MC in the $|\vec{p}|$ and p_x, p_y distributions as described in the next section.

For the momentum correction, a fraction, f_p , of events was selected. For each event, a random number, ζ , was sampled from the distribution in equation 3.37 and then multiplied by the original momentum (so that $\sigma_{|\vec{p}|} \propto |\vec{p}|$) before being added to the original momentum:

$$|\vec{p}| \rightarrow |\vec{p}| + \zeta |\vec{p}|. \quad (3.38)$$

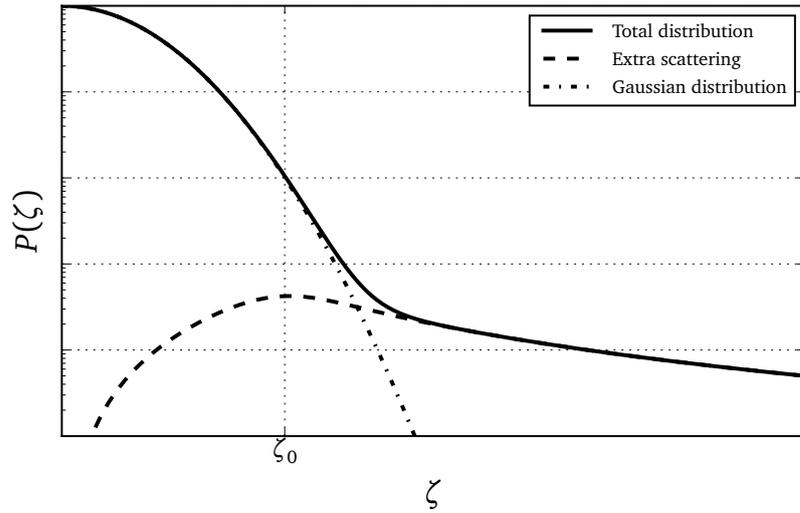


Figure 3.13.: The distribution used to generate random numbers for populating the far tails of the multiple Coulomb scattering distribution. Here the effect of adding the extra scattering distribution $P(\zeta)$ to a Gaussian distribution is shown. For a given Gaussian distribution and value of ζ_0 , the function is integrated numerically and normalized to an area of 1.

A similar procedure is used to modify the distributions of θ_x and θ_y . Two more independent subsets of the events are randomly selected, each containing the same fraction of events $f_x = f_y \neq f_{|\vec{p}|}$. Values of ζ are drawn from the pdf in equation 3.37, but with a different value of ζ_0 . The original angles are then modified by the transformation:

$$\theta_{x,y} \rightarrow \theta_{x,y} \pm \zeta \frac{1 \text{ GeV}/c}{|\vec{p}|}, \quad (3.39)$$

which ensures that $\sigma_\theta \propto 1/|\vec{p}|$.

3.5.6. Scattering parameters

Figure 3.14a shows the distribution of $p_x^{\text{DCH}} - p_x^{\text{fit}}$, for the $K^+ \rightarrow \pi^+ \pi^0$ events selected as described in sections 3.5.1 to 3.5.3, where p_x^{DCH} is the x component of the track momentum measured by the spectrometer and p_x^{fit} is the value obtained from the kinematic fit to the kaon momentum and photon clusters in the LKr calorimeter. The data points lie above the simulated spectrum so additional scattering has been introduced by modifying the values of θ_x and θ_y as described above, using the parameters:

$$f_{x,y} = 0.0080, \quad \zeta_{0x,y} = 0.0011. \quad (3.40)$$

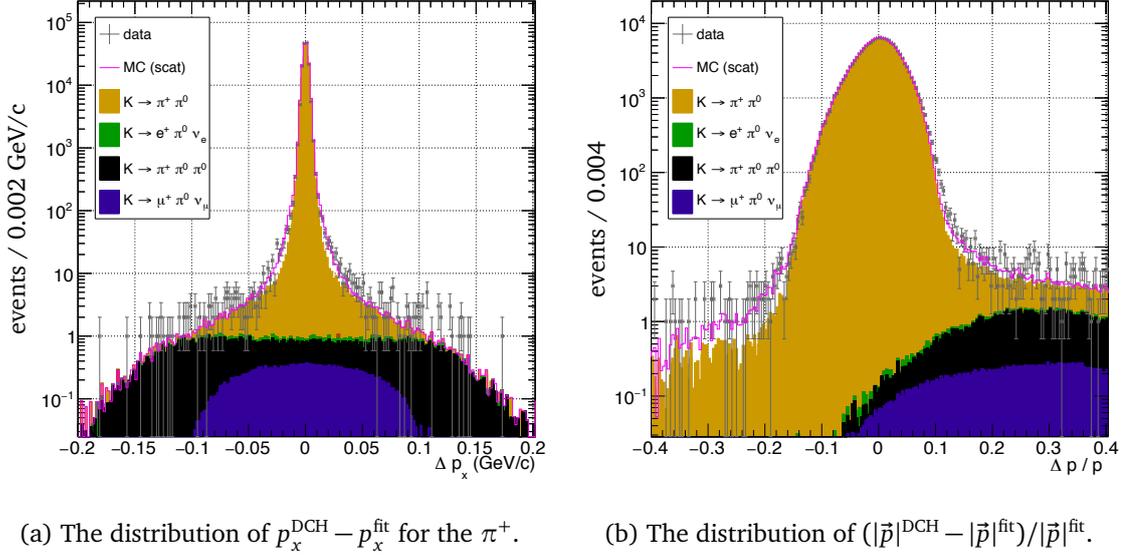


Figure 3.14.: The effect of adding extra scattering to the MC. Quantities labelled DCH are measured in the spectrometer and those labelled fit come from the kinematic fit without using the spectrometer. The original MC distribution is shown as a stack of histograms, the data points are shown with error bars and the MC distribution after applying extra scattering is shown as a solid (magenta) line.

This produces a better fit to the data, shown by a solid line. A similar result is obtained for p_y (not shown) using the same parameters.

The same procedure is applied to the magnitude of the momentum measurement, shown in figure 3.14b. In this case the scattering parameters are:

$$f_p = 0.0037, \quad \zeta_{0p} = 0.0080. \quad (3.41)$$

Figures 3.15 and 3.16 show the effect of introducing this additional scattering on the CDA, p_T and $(m_{\text{miss}}^{\pi^+})^2$ distributions, which are obtained from the π^+ momentum measured in the spectrometer, without using any information from the LKr calorimeter. In these cases, the modified MC spectrum lies above the data points, suggesting that too much scattering has been introduced, and that the discrepancy seen in the p_x and $|\vec{p}|$ distributions must in part have come from the imperfectly modelled LKr resolution.

Nevertheless, the overestimated additional scattering correction can be applied to the events passing the heavy neutrino signal selection, in order to obtain an upper limit on the possible background

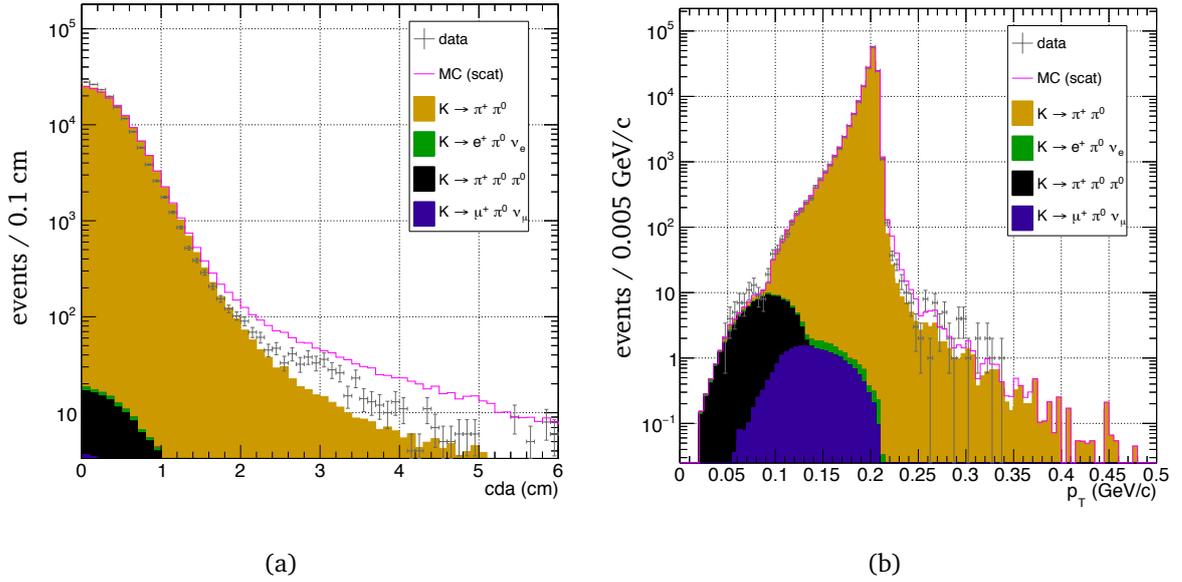


Figure 3.15.: Comparison of (a) the CDA distribution and (b) the p_T distribution of $K^+ \rightarrow \pi^+ \pi^0$ events reconstructed using the spectrometer, with and without extra scattering simulated.

contribution in the signal region from the far tails of the spectrometer resolution. The results are shown in figure 3.17. While there is a significant excess of events close to the SM $K^+ \rightarrow \mu^+ \nu_\mu$ peak, the extra scattering introduced in the MC simulation does not have a large impact on the signal region ($> 0.09 \text{ GeV}/c^2$) so the upper limit estimate of the scattering contribution can be used (see section 3.10.1), without severely worsening the final limit obtained for heavy neutrinos. The residual discrepancy in the upper shoulder of the momentum distribution in figure 3.14b is attributed to poor simulation of the LKr calorimeter, since in the distributions not relying on the LKr fit values, the points from the corrected MC simulation now lie above the data.

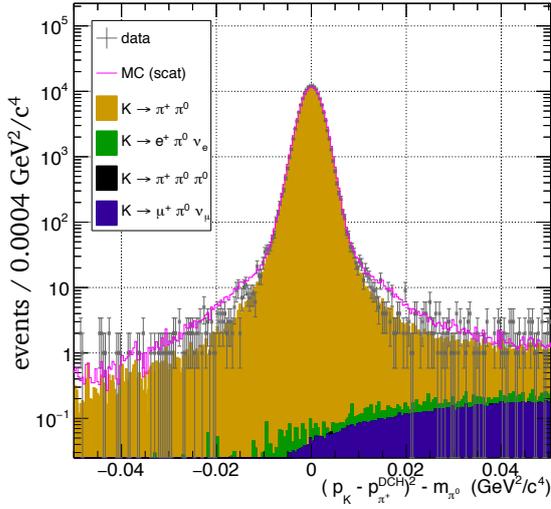


Figure 3.16.: The distribution of $(m_{\text{miss}}^{\pi^+})^2 - m_{\pi^0}^2$ in $K^+ \rightarrow \pi^+ \pi^0$ events

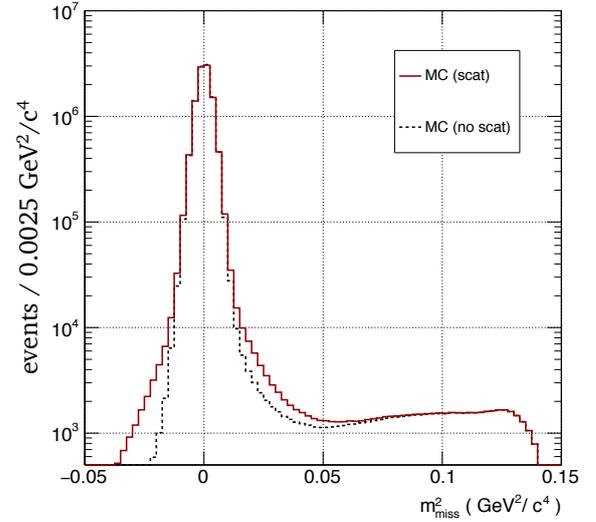


Figure 3.17.: The reconstructed m_{miss}^2 spectrum, $(p_K - p_\mu)^2$, for the sum of all the background contributions showing the effect of applying scattering to the MC simulation.

3.6. Muon veto efficiency

The Muon veto (MUV) is used in the final selection to suppress backgrounds with pions in the final state, so its efficiency must be included in the simulation of background distributions. For the MC samples used in this analysis, the interaction of particles with the detector (including scattering, showering etc.) is only simulated upstream of the LKr calorimeter. Downstream of the LKr calorimeter, only a simple extrapolation of the track is performed, so the probability of a charged track producing a signal in the MUV cannot be directly taken from the MC. Instead, the MUV efficiency for the K^+ data sample is measured and parameterized using a pure sample of muons. This parameterization is then used to apply weights to the MC samples, to reproduce the effect of the MUV inefficiency.

3.6.1. Muon reconstruction

The final selection for heavy neutrinos uses a definition of a reconstructed muon which relies on only the first two MUV planes (MUV1 and MUV2). (The third MUV plane was used only for monitoring studies during data taking.) For this analysis, a muon is defined by PMT hits in both planes 1 and 2, which are consistent in position with the extrapolated track from the spectrometer. There are no

Muon status	MUV1	MUV2	MUV3
1	✓	✓	✓
2	✓	✓	
3		✓	✓
4	✓		✓

Table 3.2.: The muon status definitions depending which MUV planes have associated hits.

requirements placed on signals in plane 3. Table 3.2 shows that this corresponds to reconstructed muons with status 1 or 2 in the terminology of the reconstruction software. The x (y) coordinate of the reconstructed muon is determined from PMT hits in MUV2 (MUV1), which has vertically (horizontally) orientated scintillator strips. In order to determine if a PMT hit is consistent with a track from the spectrometer, the two coordinates are considered separately. In each case, the track is extrapolated to the relevant plane and a momentum dependent cut is applied on the distance between the track and the reconstructed muon.

The cut on separation between track and the reconstructed muon as a function of momentum is calculated by considering the scattering of the muon as it passes through the LKr and hadronic calorimeters, as well as the resolution of the spectrometer. In the standard reconstruction for NA62 2007 data, the standard deviation for the scattering contribution in the x or y plane (either $\sigma_x^{\text{scatter}}$ or $\sigma_y^{\text{scatter}}$) is modelled by:

$$\sigma_{x,y}^{\text{scatter}} = (z_{\text{MUV}n} - z_{\text{LKr}}) \frac{13.6 \text{ MeV}}{pc} \sqrt{\frac{N^{\text{rad}}}{3}} \quad (3.42)$$

where $z_{\text{MUV}n}$ ($n = 1, 2$) is the position of the MUV plane considered, z_{LKr} is the position of the LKr calorimeter ($z_{\text{MUV}1} - z_{\text{LKr}} = 6.7 \text{ m}$, $z_{\text{MUV}2} - z_{\text{LKr}} = 7.9 \text{ m}$) and N^{rad} is the number of radiation lengths of material the muon has passed through. The number of radiation lengths is determined by the length of the detector component and the radiation length of the material it is made from:

$$N^{\text{rad}} = \frac{l_{\text{LKr}}}{X_{\text{LKr}}^0} + \frac{l_{\text{HAC}}}{X_{\text{iron}}^0} + \frac{l_{\text{iron}}}{X_{\text{iron}}^0} \times n_{\text{MUV}}, \quad (3.43)$$

where l_{LKr} and l_{HAC} are the thicknesses of the LKr and HAC detectors, l_{iron} is the thickness of iron placed in front of each muon veto, and $n_{\text{MUV}} = 1, 2$ is the MUV station. The numerical values are given in table 3.3 and the total number of radiation lengths to the first MUV plane is 112.4, so for a

constant:	l_{LKr}	l_{HAC}	l_{iron}	X_{LKr}^0	X_{iron}^0
value (cm):	123.0	120.0	80.0	1.76	4.70

Table 3.3.: Numerical values for constants used in determining the scattering contribution in equation 3.42

muon with a momentum of 20 GeV/c, $\sigma_x = 2.7$ cm.

The contribution from the spectrometer resolution as a function of momentum is:

$$\sigma_{x,y}^{\text{DCH}} = \left(1 \times 10^{-4} + \frac{0.085 \text{ MeV}}{pc} \right) (z_{\text{MUVn}} - z_{\text{DCH}}) \quad (3.44)$$

where z_{DCH} is the z -coordinate of the front face of the last drift chamber ($z_{\text{MUV1}} - z_{\text{DCH}} = 10.2$ m, $z_{\text{MUV2}} - z_{\text{DCH}} = 11.4$ m). The first term comes from the error on the direction measurement and the second term comes from scattering in the helium gas and the kevlar window at the entrance to the helium tank. This contribution is at least an order of magnitude smaller than the contribution from $\sigma_{x,y}^{\text{scatter}}$, for all relevant momenta.

Cut on the separation between the extrapolated track position and the reconstructed muon position is calculated as:

$$s_{x,y}^{\text{cut}} = \frac{1}{2} w_{\text{strip}} + \xi \sqrt{(\sigma_{x,y}^{\text{DCH}})^2 + (\sigma_{x,y}^{\text{scatter}})^2} \quad (3.45)$$

where $w_{\text{strip}} = 25$ cm is the width of a single scintillator strip and ξ is a dimensionless parameter which can be tuned to ensure negligible inefficiency. Figure 3.18 shows how the muon detection efficiency varies as a function of ξ (using the selection described in section 3.6.3). The default value for analysis 2007 data at NA62 is 4, and the figure shows that at this value the efficiency has plateaued so it is a suitable value for the present analysis.

3.6.2. Simulation of scattering applied to existing MC samples

The MC simulation of the muon veto must account for detection inefficiencies which are intrinsic to the detector, as well as inefficiencies due to muon scattering when the muon is scattered by more than the distance allowed for in equation 3.45, or outside of the detector acceptance. Since the available MC samples do not simulate scattering downstream of the LKr calorimeter, a simple model of multiple Coulomb scattering is used to simulate the distribution of muons detected in the x y plane with respect to the extrapolated track position, starting from the true muon 4-momentum in the

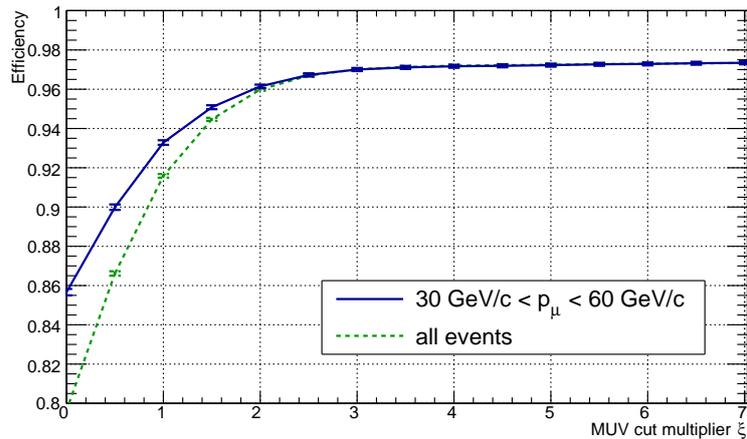


Figure 3.18.: The effect of the muon veto multiplier ξ in equation 3.45 on muon identification efficiency, shown for a selection including all track momenta, and for a more restricted selection.

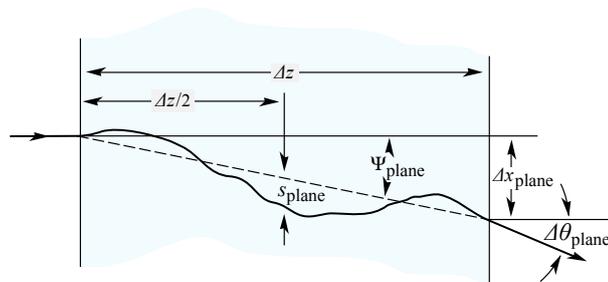


Figure 3.19.: The parameters used to describe multiple Coulomb scattering across a single slice in z . [8]

existing samples. This is particularly important in the case of decays with pions in the final state, such as $K^+ \rightarrow \pi^+ \pi^0$, which then decay into muons. The muon will not in general travel in the same direction as the measured pion track, so the probability to match the muon to the track is sensitive to the tails of the scattering distribution.

The simulation of multiple Coulomb scattering follows the approach described in the PDG [8]. Scattering effects in the x and y planes are treated independently. The path from the muon production point to the muon veto is divided into slices in z , with a new slice for each detector or gap between detectors, as well as for the iron walls used as absorbers in front of the muon vetoes. For each slice, the direction and position at the downstream end of the slice are determined from their values at the upstream end of the slice (see figure 3.19) by generating two independent random numbers, ζ_1, ζ_2 , from Gaussian distributions with mean zero and variance one. These are transformed into changes in

position and angle over the plane by

$$\Delta x_{\text{plane}} = \frac{\zeta_1 z \theta_0}{\sqrt{12}} + \frac{\zeta_2 z \theta_0}{2} \quad (3.46)$$

$$\Delta \theta_{\text{plane}}^x = \zeta_2 \theta_0, \quad (3.47)$$

which ensures the correct correlation between Δx_{plane} and $\Delta \theta_{\text{plane}}^x$. Here, θ_0 is the rms width of the Gaussian approximation for the angular distribution given by:

$$\theta_0 = \frac{13.6 \text{ MeV}}{\beta c p} \sqrt{\frac{z}{X_0}} \left[1 + 0.038 \ln \frac{z}{X_0} \right] \quad (3.48)$$

An equivalent procedure is applied for the y component. As an example, for a muon with momentum 20 GeV/c, produced in pion decay at DCH3, the rms scattering in x by the time it reaches MUV2 is 4.7 cm. At 40 GeV/c, the rms is only 2.3 cm. For the y -component, which is measured at MUV1, the rms values at 20 GeV/c and 40 GeV/c are 3.7 cm and 1.8 cm.

For each MC event in both signal and background simulations, the MUV response is determined by starting with the MC truth 4-momentum of the muon and generating a position at the MUV1 and MUV2 planes using the scattering prescription described above. The cut in equation 3.45 is then applied to determine whether the event passes the muon veto requirement. Additionally, a weight is applied to events which pass the cut, to simulate the intrinsic inefficiency of the muon veto and to ensure that the final muon detection efficiency matches that measured in the data, as described in the following section.

3.6.3. Muon sample for MUV efficiency measurement

The following selection is used to obtain a pure sample of muons from the K^+ data:

- A single track pre-selection is performed as described in section 3.3.1, with initial track parameters $3 < |\vec{p}_{\pi^+}| < 75 \text{ GeV/c}$; $-20 < z_{\text{vertex}} < 90 \text{ m}$; $\text{CDA} < 10 \text{ cm}$.
- Track charge must be +1.
- Track quality must be > 0.7 .
- Track radius at first and last spectrometer drift chambers should satisfy: $12 \text{ cm} < r_{\text{DCH1}} < 115 \text{ cm}$; $14 \text{ cm} < r_{\text{DCH4}} < 115 \text{ cm}$.

Min $ \vec{p} _\mu$ GeV/c	Max $ \vec{p} _\mu$ GeV/c	Min z_{vertex} m	Max z_{vertex} m
5	25	-10	90
25	30	0	90
30	45	5	90
45	55	10	90
55	65	15	90

Table 3.4.: The cuts in the plane of track momentum (p) and vertex z coordinate (z_{vertex}), used in the R_K analysis and adopted for the MUV efficiency study.

- The track should be within the LKr acceptance, and at least 8 cm away from areas corresponding to any missing or malfunctioning digitizers.
- Clusters in the LKr calorimeter are ignored if they have energy < 2 GeV or are out of time by > 12 ns. Clusters consistent with bremsstrahlung from the track upstream of the spectrometer (within 6 cm of the extrapolated upstream track) are also ignored. Any remaining cluster should be consistent with the track downstream of the spectrometer (cluster position within 10 cm of the extrapolated track impact point).
- The track cluster may not be in either of the ‘hot’ cells 56 and 57 in CPD (Calorimeter Pipeline Digitizer) 134, and must be at least 2 cm away from any dead cell.
- Cuts are applied in the plane of track momentum (p) and the z coordinate of the track vertex. The ‘standard’ pz cuts from the 2007 R_K analysis [86] are applied. These are shown in table 3.4.
- Electron contamination is suppressed by a cut on the energy deposited in the LKr cluster (if there is one): $E_{\text{cluster}}/|\vec{p}|_{\text{track}} < 0.85$.
- A pure sample of $K^+ \rightarrow \mu^+ \nu_\mu$ events is selected by cutting on the reconstructed missing mass, $|m_{\text{miss}}^2| \equiv |(p_K - p_\mu)^2| < 0.005 \text{ GeV}^2/c^4$, comparable with the resolution on the mass peak $\sigma_{m^2} = 0.003 \text{ GeV}^2/c^4$.

Figure 3.20 shows the distribution in track momentum and reconstructed m_{miss}^2 of the events passing the single track pre-selection. The main background for muon veto efficiency measurement comes from $K^+ \rightarrow \pi^+ \pi^0$ decays, but by cutting at $|m_{\text{miss}}^2| < 0.005 \text{ GeV}^2/c^4$, these events can be rejected without requiring the LKr for further pion/muon separation.

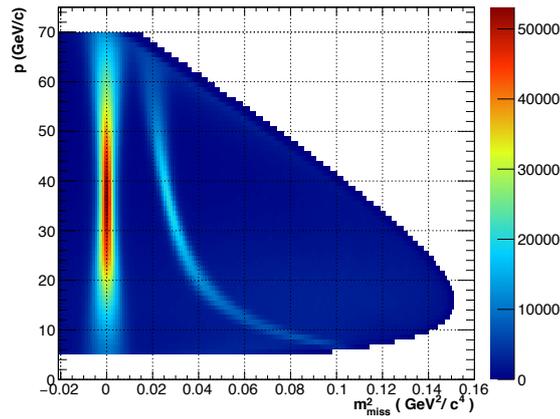


Figure 3.20.: The momentum and squared missing mass distribution of events passing the single track pre-selection. The curved band comes from $K^+ \rightarrow \pi^+ \pi^0$, while the vertical band comes from $K^+ \rightarrow \mu^+ \nu_\mu$.

The only remaining tracks besides those from $K^+ \rightarrow \mu^+ \nu_\mu$ come from the beam halo (contributing a fraction 2×10^{-4} of events). Since the beam halo consists of muons, these events can safely be included in the sample.

3.6.4. MUV Bad bursts

The sample of muons is first used to identify ‘Bad MUV’ bursts, in which the MUV is not performing as expected. Figure 3.21 shows the distribution of muon detection efficiency by the MUV by the MUV measured for individual bursts. A ‘Bad MUV’ burst is defined as one in which the efficiency is measured to be less than 90% at 95% confidence level (taking only statistical errors into account). 52 bad bursts were identified in this way and removed from the data sample. In a small fraction of runs, there are periods with many bad bursts close together, interspersed with a few good ones. Figure 3.22 shows an example. All the bursts within these periods are excluded, even if they have efficiency greater than 90%. In total 2440 (2.8%) bursts were removed from the original set of 88051 (2095 of them are from the run 20438 shown in figure 3.22).

3.6.5. MUV Efficiency measurement

The muon detection efficiency has been investigated as a function of track momentum and the track position when extrapolated to the muon veto. The position is expressed in terms of

x : the x coordinate of the track when extrapolated to MUV2

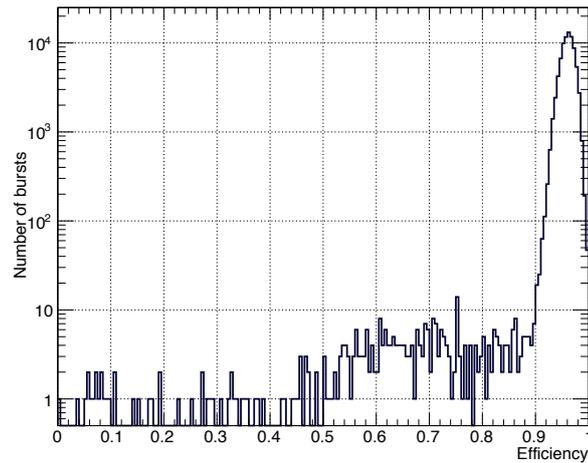


Figure 3.21.: The distribution of muon veto efficiencies measured on individual bursts.

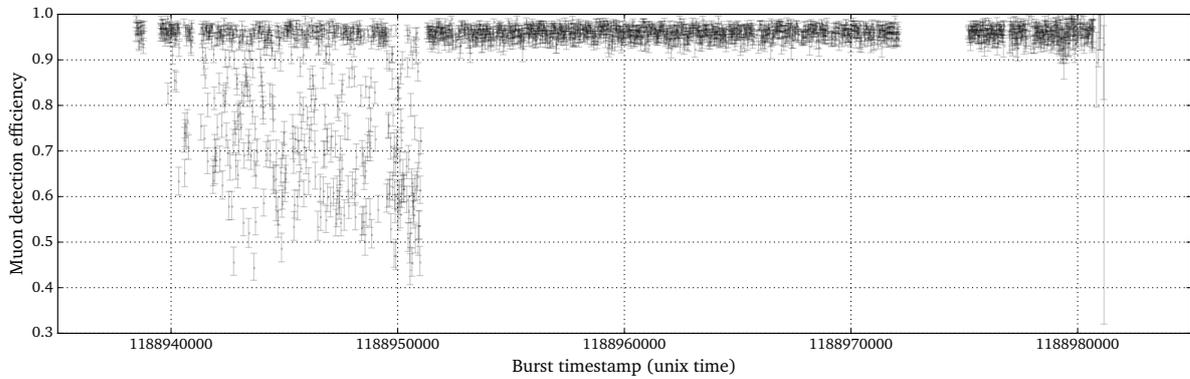


Figure 3.22.: Muon detection efficiency for the bursts in run 20438. All of the bursts in the region near the start of the run are excluded.

y : the y coordinate of the track when extrapolated to MUV1

since when reconstructing muons, the x and y coordinates are taken from MUV2 and MUV1, respectively.

Figure 3.23 shows the efficiency of the muon veto as a function of momentum and as a function of position in the xy plane, revealing structure both in distributions. The drop in efficiency at low momentum is expected from increased multiple Coulomb scattering, while the visible horizontal and vertical bands in the xy plane come from variation in efficiency between individual scintillator strips. The drop in efficiency at high momentum comes from the correlation between momentum and xy position in $K^+ \rightarrow \mu^+ \nu_\mu$ decays.

If all signal and backgrounds had the same distribution in xy and p as $K^+ \rightarrow \mu^+ \nu_\mu$ decays, then it

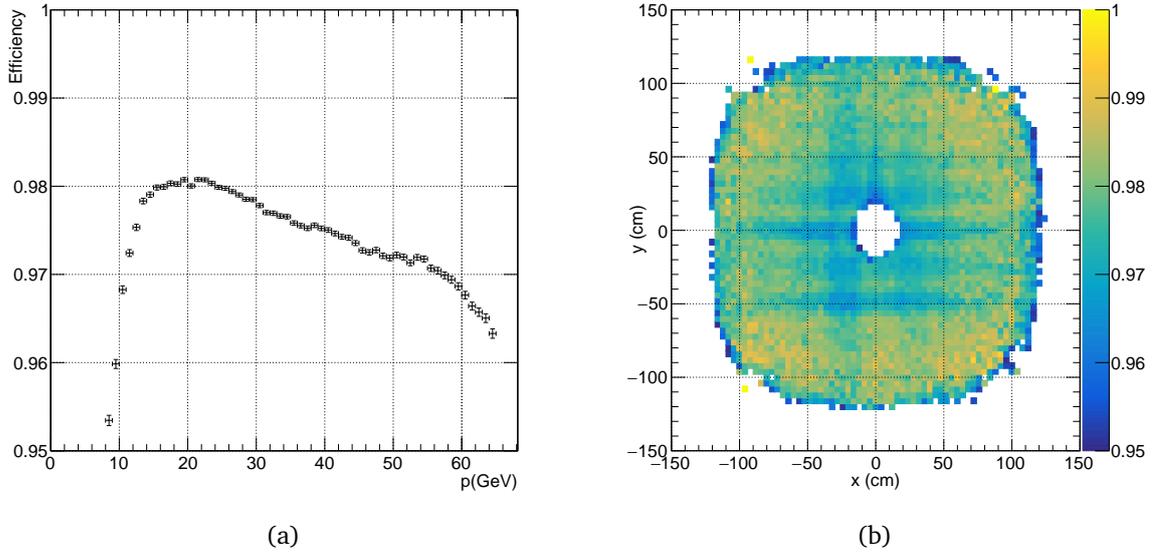


Figure 3.23.: Muon veto efficiency distributions observed in data as function of (a) momentum and (b) xy coordinate at muon veto

would be enough to measure the muon detection efficiency as a function of just one variable and apply the same correction to all MC productions. Since, in reality, the distributions vary between decays, the accuracy of the simulation can be improved by measuring efficiency as a function of both momentum and xy position at the muon veto.

Since the purely momentum dependent contribution to the muon detection efficiency from multiple Coulomb scattering occurs primarily at low momentum ($p < 10 \text{ GeV}/c$), a first approximation to the efficiency as a function of xy (independent of momentum) can be obtained by studying high momentum muons:

$$30 \text{ GeV}/c < |\vec{p}|_{\mu} < 60 \text{ GeV}/c. \quad (3.49)$$

Figure 3.24a shows that only a fraction of the MUV plane can be studied in this manner because the outer region of the detector is only illuminated by low momentum muons. This region was then directly selected by cuts in the xy plane so that the edges of the region were aligned with the 2D histogram bin edges, and the efficiency as a function of x and y was measured as shown in figure 3.24b. This efficiency xy -dependence can then be used to re-weight the MC distribution for all momenta.

Figure 3.25 shows the momentum spectrum of muons passing through the xy region shown in

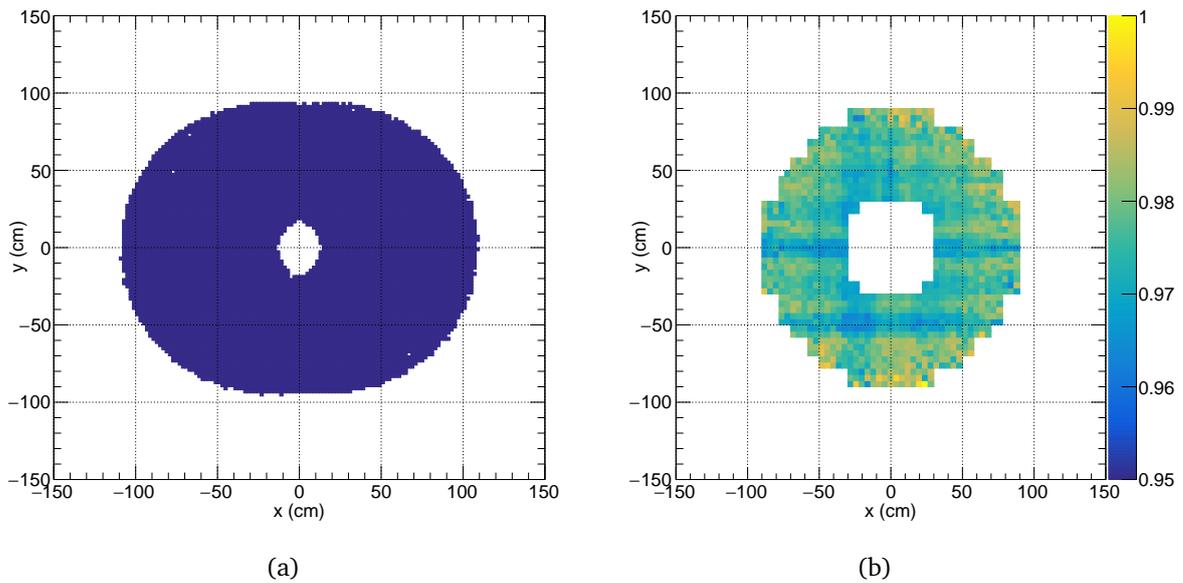


Figure 3.24.: Track xy positions at the muon veto. (a) all tracks with $30 \text{ GeV}/c < |\vec{p}|_\mu < 60 \text{ GeV}/c$ (b) the selection applied in the xy plane for further study.

figure 3.24b. The MC events have been re-weighted according to the measured xy efficiency, and the negative slope at high momentum is partially reproduced.

The remaining discrepancy is absorbed by a momentum dependent weight, given by the ratio of the two measured efficiency curves in figure 3.25, normalized to have a maximum value of 1.0. The resulting correction is shown in figure 3.26a. Since the momentum dependence is now assumed to be entirely described by the curve shown in figure 3.26a (independent of xy position), it can be introduced as a weight in the MC, and the xy dependence can be remeasured for the whole xy plane. The result is shown in figure 3.26b.

The momentum dependent weights in figure 3.26a and the xy dependent weights are used in the MC simulation for the final heavy neutrino selection. The accuracy of the model is assessed as part of the systematic error estimation in section 3.10.3.

3.7. Beam Halo

When using a single-track selection, there is background from beam halo events: those in which a muon is produced outside of the fiducial volume but leaves a track in the spectrometer which is consistent with a signal-like kaon decay inside the fiducial region. Since the signature of the

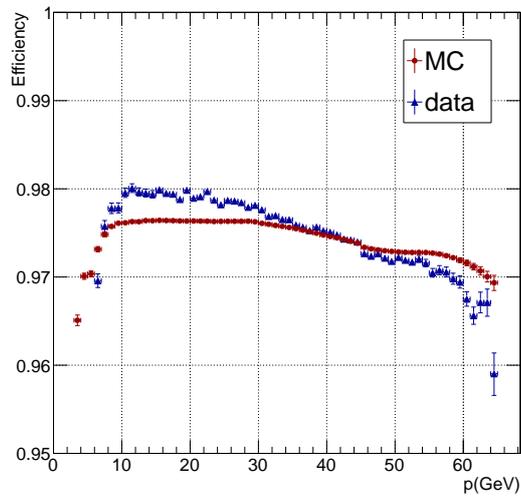


Figure 3.25.: Muon detection efficiency as a function of momentum, for tracks in the region shown in figure 3.24, after correcting for xy dependent MUV efficiency. Data in blue, MC in red.

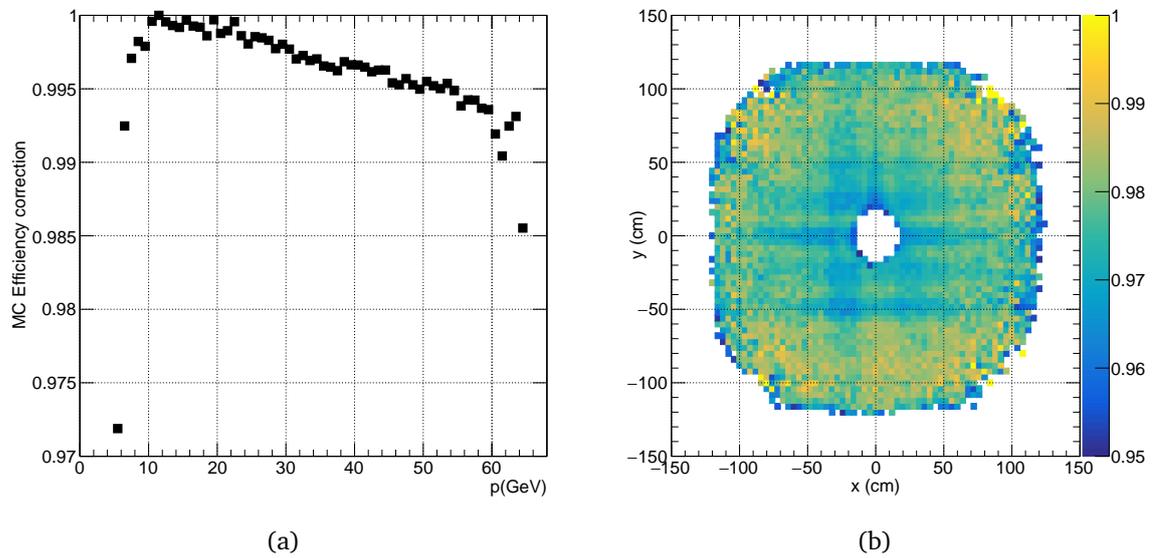


Figure 3.26.: Final MUV efficiency corrections applied to MC. (a) momentum dependent. (b) xy dependent

heavy neutrino decay mode is also a single muon, the halo background cannot be reduced using particle identification. Instead, signal and halo events are separated by their different kinematic and geometrical distributions.

The halo distribution cannot be accurately simulated (although important features of its kinematic distribution can be reproduced) so a data driven approach is used in modelling this background component. There are significant systematic errors (of order 20% – 50%) associated with this approach so despite only being a minor component of the total background, the halo background is responsible for a large fraction of the final uncertainty on the background estimate.

In the following sections, a simulation is used to demonstrate the main features of the halo distribution. Then a data driven approach is introduced and used to design a selection which minimizes the halo background while preserving signal acceptance. The deficiencies in the data driven approach are discussed and a control sample is used to estimate the uncertainty on the final background estimate.

3.7.1. Halo Simulation

Figure 3.27 shows the expected distribution of halo events from a simulation of the beam halo. The distribution was produced using the HALO program [98], which simulates the secondary beam, from its production at the beryllium target, through the beam optics system and into the decay volume, using a transfer matrix approach. Muons from decays of kaons and pions in the beam are generated with the appropriate kinematic distributions and tracked through the magnetic fields along the beam line, simulating energy loss and scattering as they pass through material elements. The standard CMC simulation can then be used to predict the detector response to these muons and produce the expected distribution of reconstructed events.

In figure 3.27b there are two distinct contributions to the halo background. One contribution comes from events with high momentum muons appearing to come from a vertex at low z . These events are from muons which were produced in decays far upstream of the decay volume and travelled along with the pions and kaons all the way through the beam optics before scattering in the collimator near the beam axis. The other contribution comes from low momentum muons which have been swept away from the beam line by the muon scraper magnets (designed for this purpose) and reach the collimator some distance away from the beam axis. The effect of the first of these beam scraping magnets can be seen in figure 3.27a in the deflection of tracks towards positive x at $z = -9000$ cm

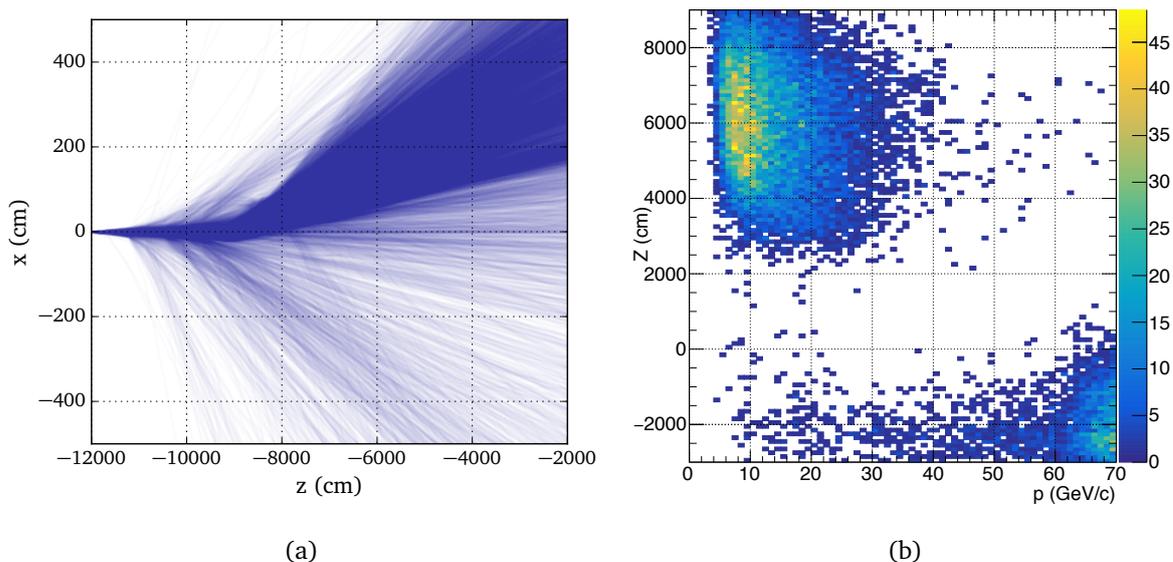


Figure 3.27.: Simulation of the beam halo background with the HALO and CMC programs. (a) the xz projection of the paths taken by the muons in the HALO simulation. (b) the distribution of reconstructed halo events, in terms of the track momentum and the z coordinate of the reconstructed vertex, when the muon tracks from the output of HALO program are used as input for the CMC program. The decay volume starts at $z = -1800$ cm.

and other magnets further along the beamline produce similar effects. The steel collimator is about 100 radiation lengths thick resulting in scattering of muons with rms 0.1 rad at 10 GeV/c. A few of the muons will be scattered back across the beam line at high z creating an apparent vertex when the muon is detected in the spectrometer.

The following sections show that the HALO program successfully reproduces the qualitative features of the halo distribution seen in data, but the precision (limited by the uncertainties on the fringe fields of the beam line magnets) is not sufficient for background subtraction. Instead the halo background is estimated using data driven methods.

3.7.2. Data driven approach

The distribution of halo background events can be estimated from data by using data taking periods in which no K^+ entered the decay region (either K^- periods, or $Kless$ periods in which both K^+ and K^- beams were blocked by the TAX absorbers). Since no K^+ decays are expected to occur in the fiducial volume, the majority of reconstructed positively charged muons must come from beam halo muons. (Two exceptions, due to K^+ particles which pass through the absorbers, and the mis-identified π^+ s

from $K^- \rightarrow \pi^- \pi^- \pi^+$ decays are considered later.) Therefore, applying a K^+ selection to K^- or $Kless$ data reproduces approximately the beam halo spectrum expected in K^+ data. Since there are only a few runs in the $Kless$ period, and only one spectrometer polarity, the K^- period is used for estimating the halo distribution, and the $Kless$ period is used as an additional check.

3.7.3. Halo Pre-selection

A loose single (positive) track selection is applied to K^- periods in order to see the general structure of the halo kinematic distribution:

- The single track pre-selection is applied with parameters $3 \text{ GeV}/c < p_\mu < 80 \text{ GeV}/c$; $CDA < 15 \text{ cm}$; and no restriction on z initially
- Track charge must be +1.
- Track quality must be > 0.7 .
- The track momentum should be in the range $3 \text{ GeV}/c < p_\mu < 70 \text{ GeV}/c$.
- The reconstructed vertex should have a z coordinate in the range $-30 \text{ m} < z_{\text{vertex}} < 90 \text{ m}$.
- The distance of the track from the z axis at first and last spectrometer drift chambers should satisfy: $12 \text{ cm} < r_{\text{DCH1}} < 115 \text{ cm}$; $14 \text{ cm} < r_{\text{DCH4}} < 155 \text{ cm}$.
- The track should be within the LKr acceptance, and the LKr photon veto is applied as described in the muon sample selection on page 83.

The reconstruction makes use of a nominal kaon momentum and trajectory in order to reconstruct the decay vertex. The procedure for determining the nominal kaon parameters, described in section 3.4.2, is not applicable when reconstructing halo events, since there were no positively charged kaons entering the decay volume whose average momenta could be measured. Instead when reconstructing halo events from K^- and $Kless$ data samples, the reconstruction was performed using the K^+ momenta measured in the K^+ data sample. For each halo event, a set of kaon track parameters was chosen at random from all the measured K^+ parameter sets, weighted by the number of $K^+ \rightarrow \mu^+ \nu_\mu$ events recorded with that set.

The distributions of the events passing this pre-selection are shown in figure 3.28, both in the pz plane (shown for the MC simulation in figure 3.27b) and in the $z\theta$ plane, where θ is the angle

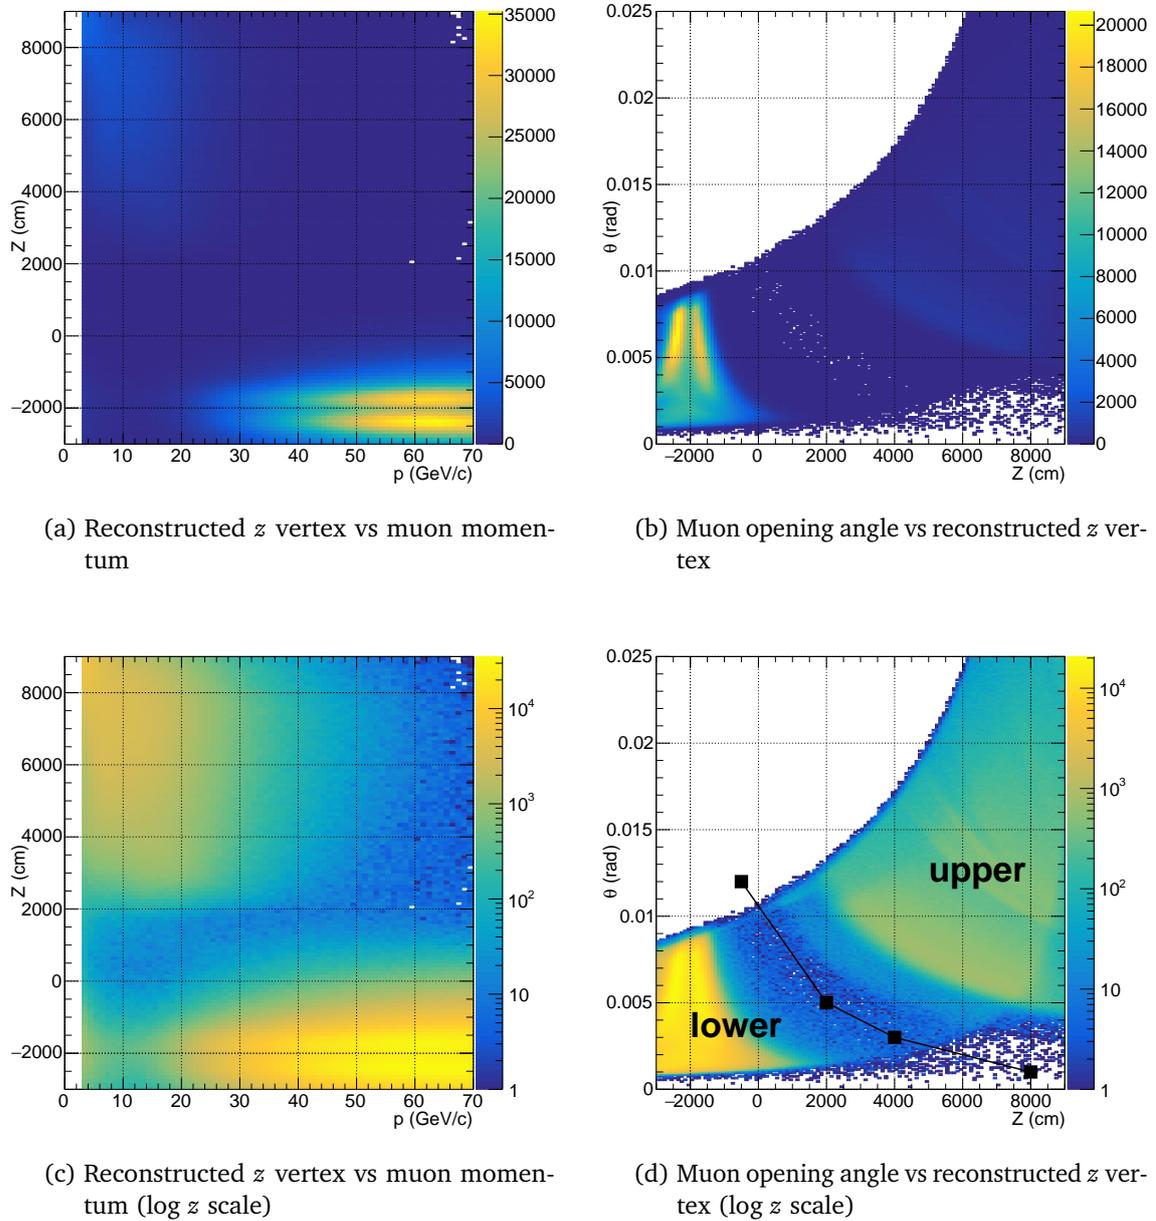


Figure 3.28.: Kinematic distributions of halo events reconstructed with a single track pre-selection. Figure (d) shows how the events are divided into *upper* and *lower* regions for further analysis.

between the nominal K^+ momentum and the muon momentum in the laboratory frame, z is the z coordinate of the reconstructed vertex and p is the reconstructed track momentum. As in the MC simulation, there is a clear separation of the events into two kinematic regions. Figure 3.28d shows a line in the plane of z and θ which is used to define the upper and lower regions. The coordinates of the marked points are:

$$\begin{aligned} &(-500 \text{ cm}, \quad 0.012 \text{ rad}), \\ &(2000 \text{ cm}, \quad 0.005 \text{ rad}), \\ &(4000 \text{ cm}, \quad 0.003 \text{ rad}), \\ &(8000 \text{ cm}, \quad 0.001 \text{ rad}). \end{aligned}$$

As has already been discussed, halo events in the upper and lower regions have different physical origins and, furthermore, the relative frequencies of the two classes of event are not necessarily the same in the K^+ data sample as in the K^- and $Kless$ sample because in the K^- and $Kless$ samples, the K^+ beam is absorbed 105 m upstream of the decay volume. For this reason, the two regions are treated separately, with independent weights in the final background estimate. Since they have different kinematic distributions, different cuts can be applied in each region to maximize the ratio of signal to background.

The kinematic distributions of both signal and backgrounds depend on the polarities of the beam magnets and the spectrometer magnet. In MC, the correct ratio of events with positive and negative magnet polarities is simulated, so any asymmetries are automatically accounted for. For the data driven approach to the halo, the K^- data samples with positive and negative magnet polarities are weighted to match the ratio in the K^+ data taking period.

In the following sections, the upper and lower regions are studied separately in order to motivate the final kinematic selection, which is applied in addition to the pre-selection outlined above.

3.7.4. Upper region

Figure 3.29 shows the distribution of halo events in the upper kinematic region and the selections used to suppress them. For reference, the signal distributions for neutrinos with masses $250 \text{ MeV}/c^2$ and $350 \text{ MeV}/c^2$ are shown in figure 3.30. Figure 3.29a shows the distribution of the z coordinate of the reconstructed vertex and the opening angle, θ between the nominal kaon direction and the measured muon direction in the lab frame. There is a concentration of events at high z and θ so this

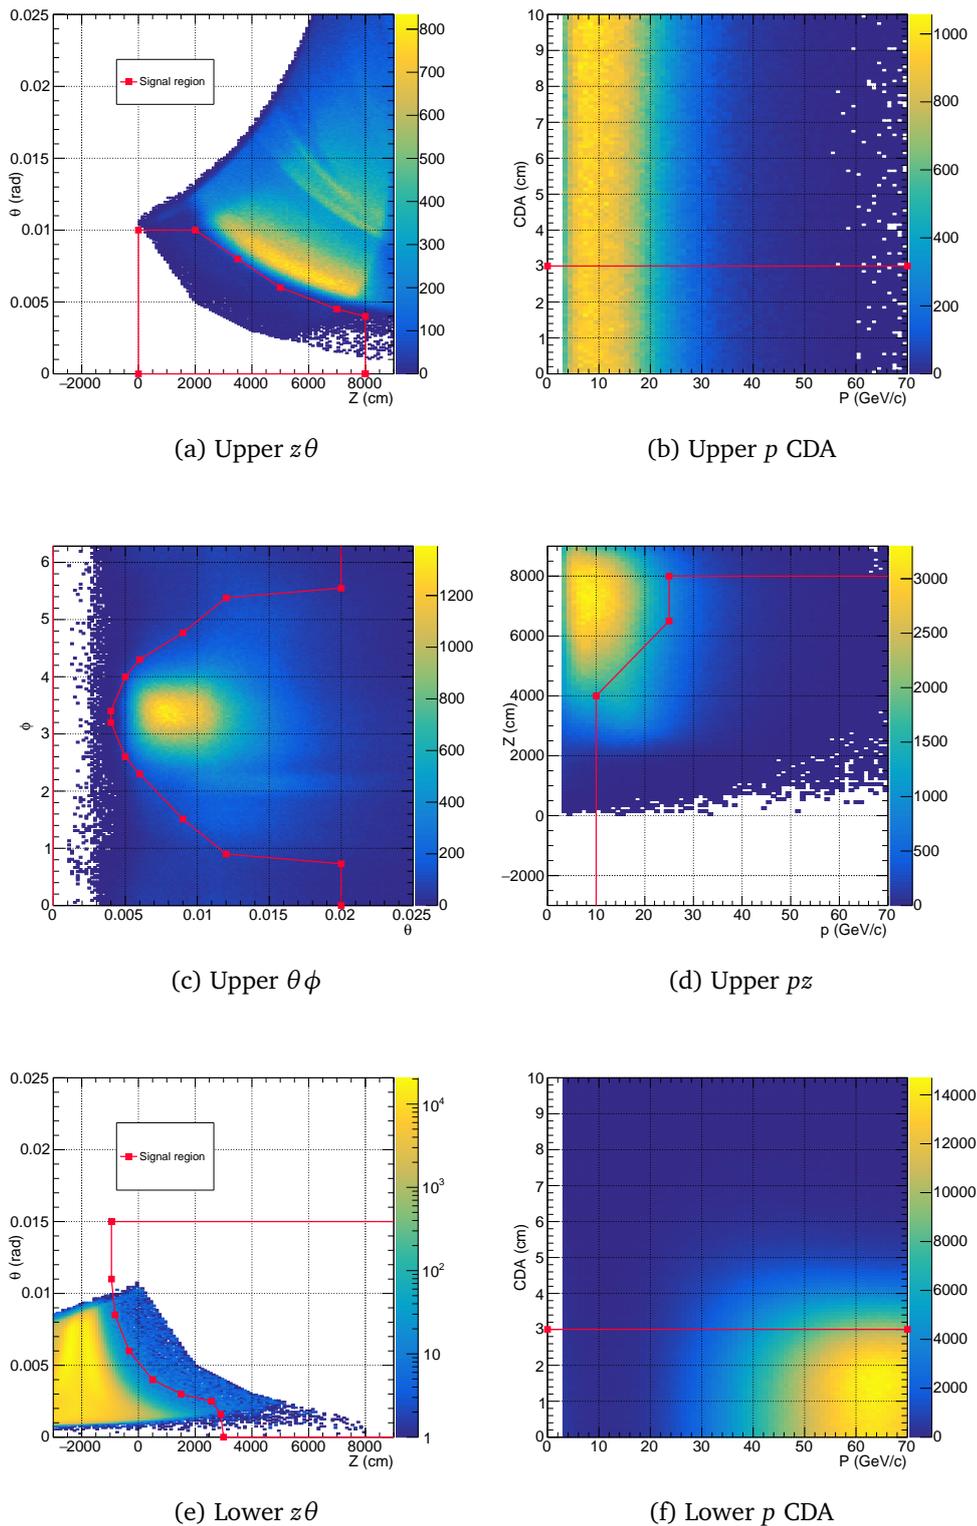


Figure 3.29.: Kinematic distributions of halo events in the ‘upper’ kinematic region ((a) to (d)) and in the ‘lower’ kinematic region ((e) and (f)). The kinematic variables are: p , the track momentum; z , the z -coordinate of the reconstructed track vertex; CDA, the track CDA; θ , the angle between the track and the nominal kaon trajectory and ϕ , the track azimuthal angle. Cuts are shown as red lines.

region is excluded by the selection shown. The marked points are:

(8000 cm, 0.000 rad)	(8000 cm, 0.004 rad)	(7000 cm, 0.0045 rad)	(5000 cm, 0.006 rad)
(3500 cm, 0.008 rad)	(2000 cm, 0.010 rad)	(0 cm, 0.010 rad)	(0 cm, 0.000 rad)

These high- z events are further suppressed by cutting in the plane of z -vertex and track momentum as shown in figure 3.29d. The marked points are:

(10 GeV/c, 4000 cm)	(25 GeV/c, 6500 cm)	(25 GeV/c, 8000 cm)
---------------------	---------------------	---------------------

and all events with $p < 10$ GeV/c or $z > 8000$ cm are excluded.

Figure 3.29b shows that the events are distributed uniformly in the CDA of the reconstructed vertex, which is expected since the muon tracks have not come directly from the kaon beam line but have been scattered before entering the decay volume. Since signal events have a CDA distribution which peaks at zero, the signal to background ratio can always be improved by decreasing the maximum allowed CDA. However, the exact shape of the kaon beam profile in x and y , and hence the CDA distribution, is not well described in MC, so there is a significant systematic error if too tight a cut is used. For the final selection, the cut is left at $CDA < 3.0$ cm.

Figure 3.29c shows the direction of halo muons in the plane of the opening angle θ and the azimuthal angle, ϕ of the muon track, given by $\phi = \arctan(p_y/p_x)$. There is a peak around $\theta = 0.08$ rad, $\phi = 3.4$ rad. This is consistent with the action of the beam sweeping magnets which shift muons towards positive x , so the ones which are later scattered back across the beam line have negative dx/dz . A large fraction of halo events can be rejected by cutting in the $\phi\theta$ plane as shown in the figure. The points defining the allowed signal region are (in radians):

(0.000, 0.00)	(0.020, 0.00)	(0.020, 0.73)	(0.012, 0.90)
(0.009, 1.51)	(0.006, 2.30)	(0.005, 2.60)	(0.004, 3.20)
(0.004, 3.40)	(0.005, 4.00)	(0.006, 4.30)	(0.009, 4.77)
(0.012, 5.38)	(0.020, 5.55)	(0.020, 6.40)	(0.000, 6.40)

3.7.5. Lower region

The lower region is shown in the $z\theta$ plane in figure 3.29e, along with the definition of the low z boundary, given by the points:

(−950 cm, 0.0200 rad)	(−950 cm, 0.0110 rad)	(−825 cm, 0.0085 rad)	(−325 cm, 0.0060 cm)
(500 cm, 0.0040 rad)	(1500 cm, 0.0030 rad)	(2575 cm, 0.0025 rad)	(2900 cm, 0.0016 cm)
(3000 cm, 0.0000 rad)	(10000 cm, 0.0000 rad)	(10000 cm, 0.0200 rad)	(−950 cm, 0.0200 cm)

Figure 3.29f shows that events in the lower region have high momentum and low CDA, consistent with muons which have travelled through the beam optics along the kaon beam line. As in the upper region a cut is placed at $CDA < 3.0$ cm.

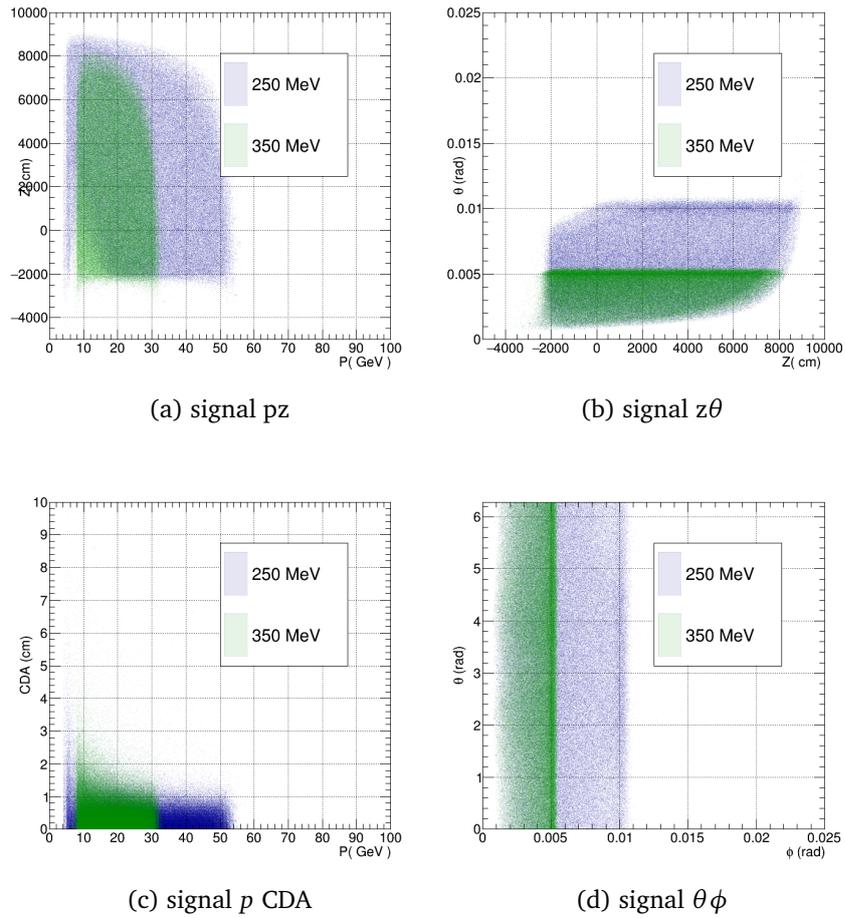


Figure 3.30.: Kinematic distributions of signal events for neutrinos with masse $250\text{MeV}/c^2$ (blue) and $350\text{MeV}/c^2$ (green).

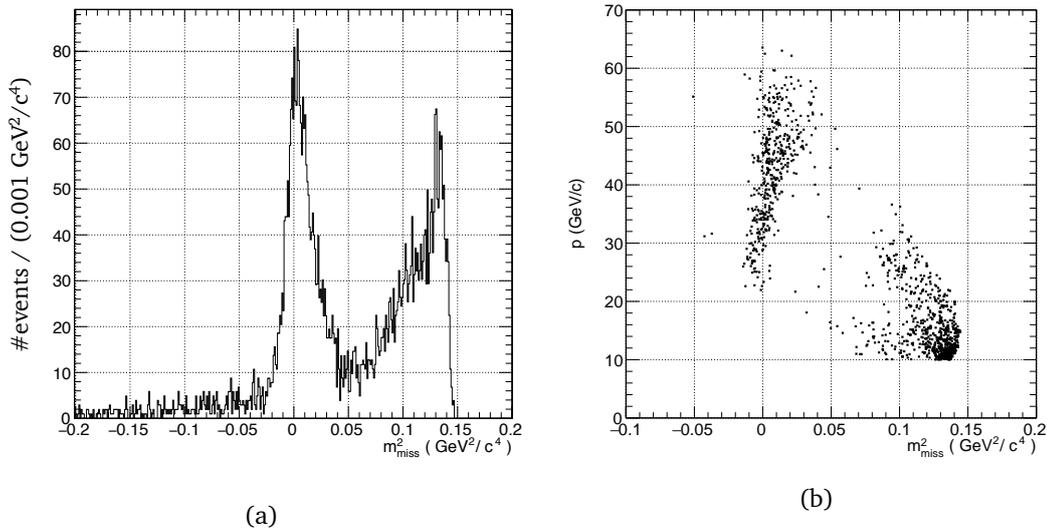


Figure 3.31.: The squared missing mass distribution for positively charged muon tracks obtained from the K^- data taking period.

3.7.6. Halo missing mass distribution

After applying the cuts defined in the previous sections in order to reject halo events, the remaining positive track events in the K^- sample have the distribution shown in figure 3.31. The two peaks in the squared missing mass distribution around $0 \text{ GeV}^2/c^4$ and $0.13 \text{ GeV}^2/c^4$ do not come from the beam halo itself, but have separate sources: kaon punch-through (the peak around zero), and $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ cross-talk (the peak at higher squared missing mass). These are described in the following sections.

3.7.7. Punch-through

During K^+ and K^- data taking periods, positively charged kaons are blocked upstream of the decay volume by two layers of copper and iron absorbers (TAX, each 1.6 m thick) in between the second and third magnets in the first achromat. For K^+ data taking periods, positively charged kaons are allowed to pass through holes in the absorbers. K^+ particles are blocked by shifting the absorber layers relative to each other so that the holes no longer overlap. However, this leaves areas where the absorber is only half its nominal thickness (around 7 nuclear interaction lengths instead of 14) and a small fraction of kaons (5×10^{-4}) can pass through and into the decay volume. Since the holes

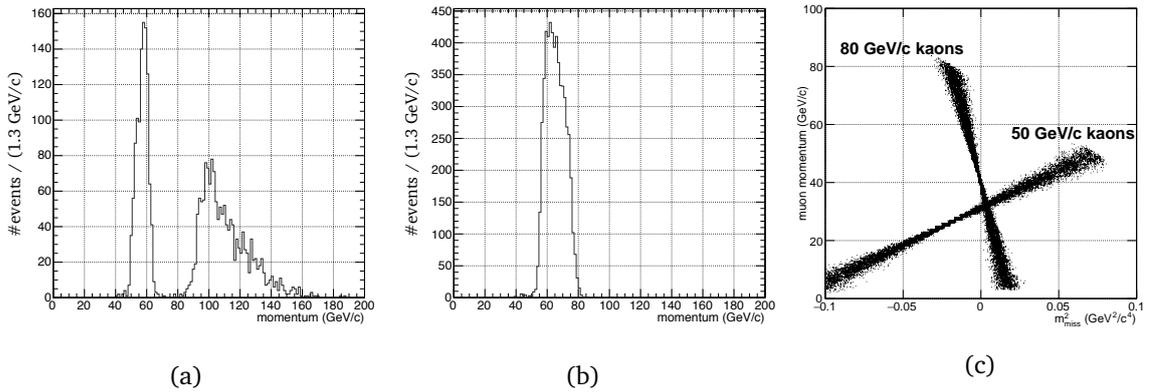


Figure 3.32.: The effect of kaon punchthrough on the halo measurement. (a) the momentum spectrum of K^+ particles reconstructed from $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ decays during the K less data taking period. (b) the momentum spectrum of K^+ particles reconstructed in the K^- data taking period. (c) a toy MC demonstration of the distribution of momentum and reconstructed squared missing mass for $K^+ \rightarrow \mu^+ \nu_\mu$ decays by kaons of different momenta, assuming the nominal kaon momentum (74 GeV/c) in the reconstruction.

are displaced from their nominal positions, the kaons which pass through do not have the nominal momentum which is obtained when the holes are aligned.

This phenomenon can be studied using $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ decays to unambiguously reconstruct K^+ particles, in the K less and K^- data taking periods.

The selection for $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ events is as follows:

- One three-track vertex should be reconstructed with $\chi^2 < 40$.
- The total charge of the three tracks should be +1.
- No two tracks should be separated by less than 0.5 cm when extrapolated to the plane of the first drift chamber (in order to suppress the contribution from photon conversion in the kevlar window at the upstream end of the spectrometer).
- The deflection in y between drift chambers 1 and 4 should be < 0.6 cm for all tracks.
- The total transverse momentum of the three tracks should be less than 0.001 GeV/c.
- The z -coordinate of the reconstructed vertex should be between -18 m and 70 m.
- The reconstructed kaon mass: $m_K = \sqrt{(p_{\pi^+} + p_{\pi^+} + p_{\pi^-})^2}$ should be within 0.003 GeV/c 2 of the true kaon mass.

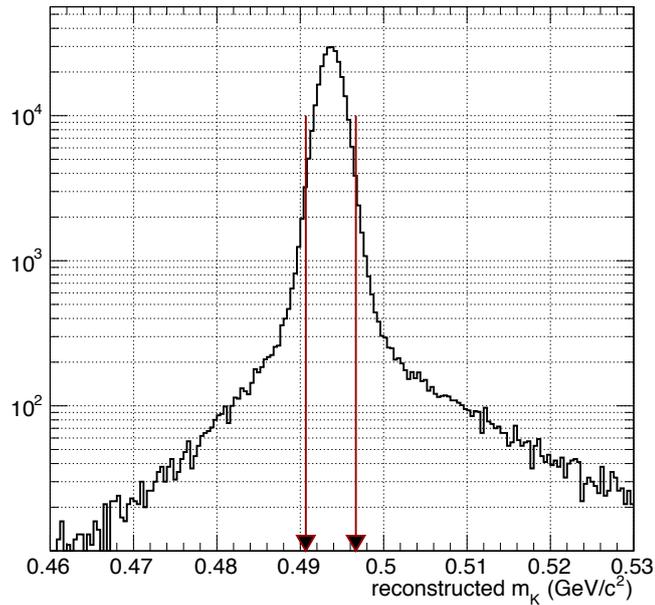


Figure 3.33.: The distribution of reconstructed kaon mass (the invariant mass of the three pions) in $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ decays, shown here for K^+ data. The cut applied for the beam momentum measurement is shown by red arrows.

Since all three charged pions are detected, $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ decays can be used to reconstruct the four momentum of the kaon. The invariant mass is used as a cut to select the decay and the kaon momentum is then plotted for the selected events. Figure 3.32 shows the momentum of positively charged kaons reconstructed in the K_{less} and K^- periods. The distributions are different because the alignment of the absorbers required to block both beams for K_{less} periods is different from the alignment required to block just the K^+ beam for K^- periods.

Figure 3.32c shows a toy MC simulation of the squared missing mass (as a function of track momentum) reconstructed under the assumption that the kaon momentum is $74 \text{ GeV}/c$, when the true value is larger or smaller ($50 \text{ GeV}/c$ or $80 \text{ GeV}/c$, corresponding to the limits of the distribution in figure 3.32b). The results show correlation between reconstructed squared missing mass and track momentum, which is also seen in the data in figure 3.31b, however the detector acceptance has not been simulated so no quantitative results can be extracted. Since the contribution from kaon punchthrough is all at small squared missing mass and does not affect the signal region ($m_{\text{miss}}^2 > 0.0075 \text{ GeV}^2/c^4$), no further simulation is required.

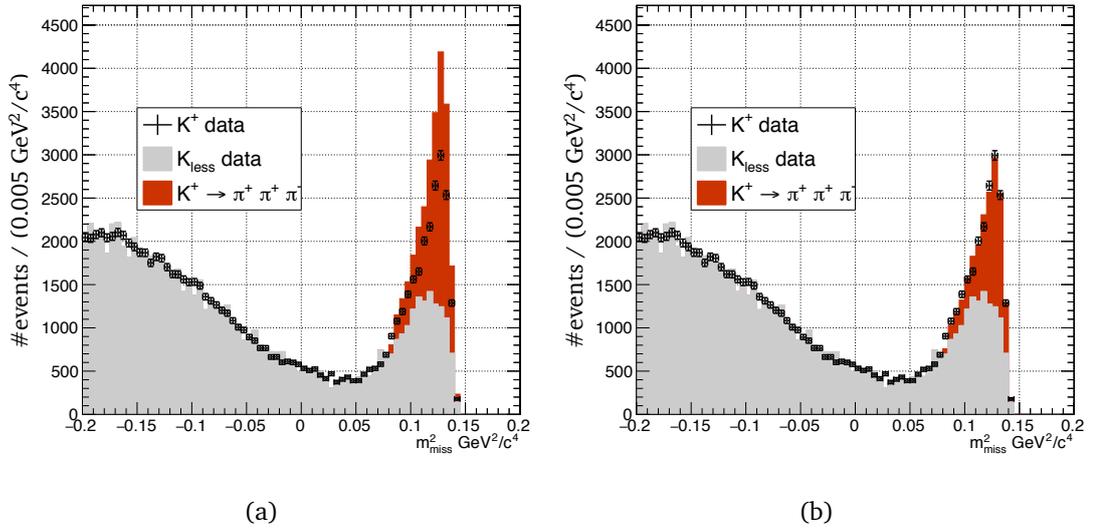


Figure 3.34.: The squared missing mass distribution of the negatively charged (μ^-) beam halo in the K^+ data taking period for the signal upper kinematic selection (with negative track requirement), compared with the K_{less} period and the simulated contribution from $K^+ \rightarrow \pi^+ \pi^+ \pi^-$, using a relaxed selection described in section 3.7.8. In figure (a), the $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ contribution is scaled using the kaon flux estimated from $K^+ \rightarrow \mu^+ \nu_\mu$ decays. In figure (b), the $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ contribution is scaled to fit the data.

3.7.8. $K^- \rightarrow \pi^- \pi^- \pi^+$ cross-talk

The peak at high squared missing mass in figure 3.31a comes from K^- decays to three charged pions, in which the two negatively charged pions are not detected, leaving a single positively charged track which cannot be distinguished from a muon because the muon veto was not available in the K^- and K_{less} data taking periods. These events must be simulated with MC and subtracted from the measured squared missing mass spectrum in order to obtain an estimate for the underlying halo distribution.

This cross-talk effect is studied here with larger statistics by looking at the halo distribution measured when negative tracks are reconstructed in K^+ data taking periods (i.e. the reverse of the normal approach) since the muon sweeping magnets were not optimized to remove halo background in this configuration. In addition, to increase the statistics for this study for this study, the cuts in the p_z plane (figure 3.29d) are removed, the CDA cut is relaxed to 3.5 cm and the muon veto is not used.

Figure 3.34a shows the distribution of squared missing mass for negative tracks in K^+ data compared with K_{less} data. The difference is attributed to the cross talk from $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ decays in which

only the π^- is detected. The number of $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ decays to simulate was determined from the estimated total kaon flux for the K^+ period, measured using $K^+ \rightarrow \mu^+ \nu_\mu$ decays (see section 3.10.2) and the $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ branching ratio. However, this overestimates the $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ contribution, partly because the $Q_1 \times 1\text{TRKL}$ trigger is not 100% efficient for these decays with two lost pions (see section 3.10.4) and partly because the acceptance for these events is not well reproduced in MC: the K^+ and K^- beamlines diverge in the decay volume so events due to this cross-talk effect are more likely to come from the poorly understood tails of the CDA spectrum.

Figure 3.34b shows the simulated $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ distribution rescaled to match the observed distribution in data by minimizing the χ^2 between the two distributions for the high mass region $0.11 \text{ GeV}^2/c^4 < m_{\text{miss}}^2 < 0.14 \text{ GeV}^2/c^4$. The required correction factor is:

$$r_{K3\pi} \equiv \frac{N_K^{K3\pi}}{N_K^{K\mu2}} = 0.583 \pm 0.017 (\text{stat.}), \quad \chi^2/n\text{DOF} = 21.3/24, \quad (3.50)$$

where $N_K^{K\mu2}$ is the number of kaons expected from the $K^+ \rightarrow \mu^+ \nu_\mu$ peak measurement, and $N_K^{K3\pi}$ is the effective number of kaons determined by the measurement of $K^+ \rightarrow \pi^+ \pi^+ \pi^-$.

The same correction factor can be used to subtract $K^- \rightarrow \pi^- \pi^- \pi^+$ events from the distribution observed in K^- periods. Figure 3.35 shows that there is good agreement at high squared missing mass using the same correction factor. (The disagreement at low squared missing mass is due to the different punchthrough kaon momentum spectra which is not simulated, as discussed previously). Finally the procedure is applied to events passing the full signal selection. The total halo distribution after $K^- \rightarrow \pi^- \pi^- \pi^+$ subtraction is shown in figure 3.37.

3.7.9. Scale factor between K^+ and K^- data

Having determined the shape of the expected squared missing mass distribution of the halo from K^- data taking periods, the distribution must be multiplied by a scale factor in order to match the distribution in K^+ periods. In total, there are four scale factors which must be determined since for each spectrometer polarity, a scale factor is required for both the upper and lower kinematic regions. To determine the scale factor, a selection is defined which is dominated by the halo contribution so that the scale is given by the ratio of events detected in the K^- and K^+ data taking periods.

Figure 3.3 (page 56) shows that in the distribution of squared missing mass, the region $m_{\text{miss}}^2 < -0.05 \text{ GeV}^2/c^4$ is dominated by the halo because backgrounds from kaon decays lead to positive

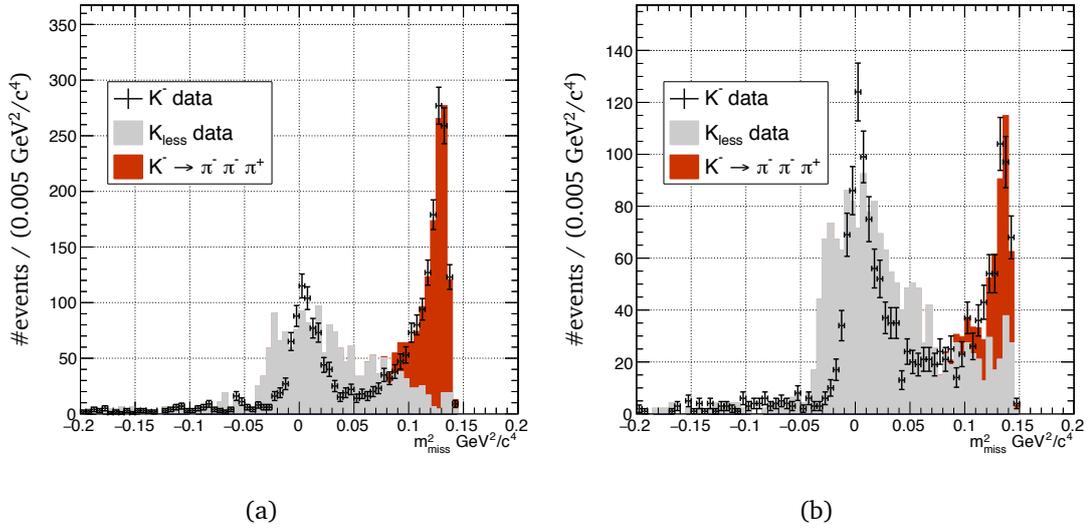


Figure 3.35.: The squared missing mass distribution of the positively charged (μ^+) halo for the K^- data taking period, using a relaxed selection described in section 3.7.8, for (a) the upper kinematic region and (b) the lower kinematic region, compared with the K_{less} data sample. The simulated $K^- \rightarrow \pi^- \pi^- \pi^+$ contribution is scaled using the estimated K^- flux and the correction factor in equation 3.50. The disagreement near $m_{\text{miss}}^2 = 0$ comes from the kaon punch-through contribution which is not simulated, see section 3.7.7.

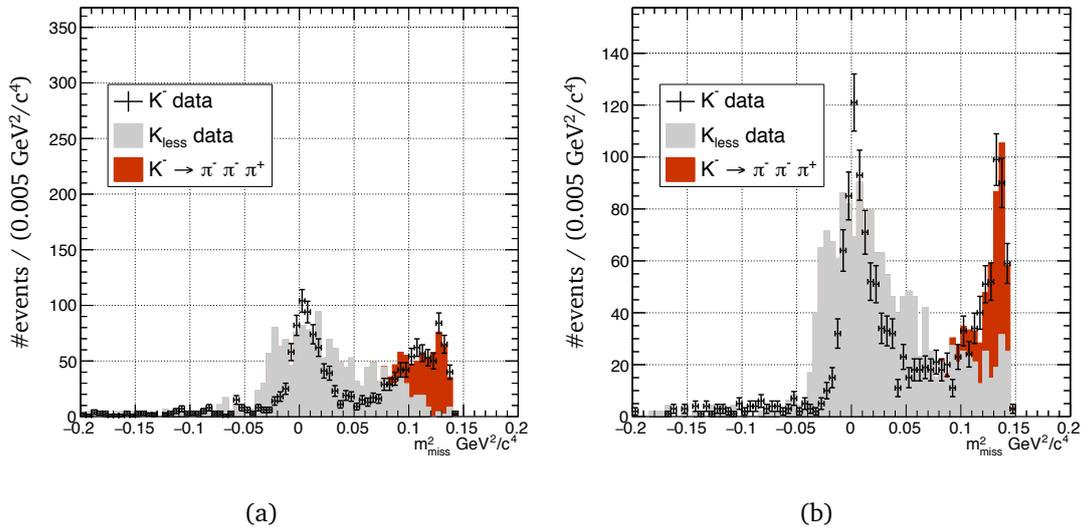


Figure 3.36.: The distribution of the positively charged (μ^+) halo for the K^- data taking period, as described in figure 3.35, but here using the full signal selection.

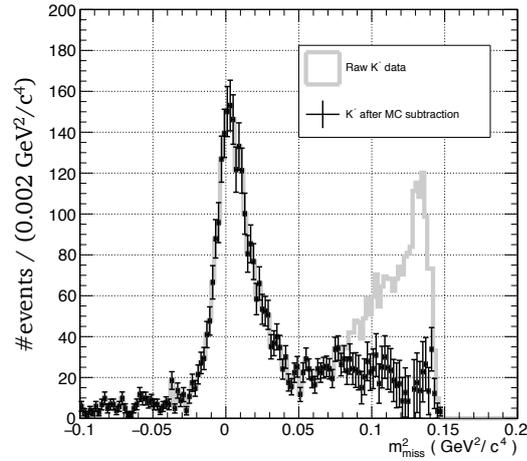


Figure 3.37.: The halo distribution measured in the K^- period, before and after MC $K^- \rightarrow \pi^- \pi^- \pi^+$ subtraction (statistical errors only).

reconstructed squared missing mass¹. In the 2007 R_K analysis [86], this region was used to determine the scale factor for the halo by integrating the number of events in the region $-0.3 \text{ GeV}^2/c^4 < m_{\text{miss}}^2 < -0.1 \text{ GeV}^2/c^4$. Figure 3.38a shows the distributions from the K^+ and K^- data samples after applying the signal selection, from which the halo scale factor could be extracted. However, this approach is less suitable for the present analysis. The negative squared missing mass control region is not directly adjacent to the signal region and the events have a different kinematic distribution. For example, as figure 3.38b shows, events with negative squared missing mass typically have large opening angles, while in the final signal selection, the opening angle is restricted to $\theta < 0.01$ rad. The measured scale factors were:

$$s_{\text{halo}}^{\text{neg}m^2}(\text{+ve upper}) = 2.90 \pm 0.25 \quad (\text{stat.}) \quad (3.51)$$

$$s_{\text{halo}}^{\text{neg}m^2}(\text{-ve upper}) = 4.09 \pm 0.37 \quad (\text{stat.}) \quad (3.52)$$

$$s_{\text{halo}}^{\text{neg}m^2}(\text{+ve lower}) = 5.26 \pm 0.70 \quad (\text{stat.}) \quad (3.53)$$

$$s_{\text{halo}}^{\text{neg}m^2}(\text{-ve lower}) = 8.59 \pm 1.41 \quad (\text{stat.}) \quad (3.54)$$

where +ve and -ve refer to data samples with positive and negative spectrometer polarity and upper and lower refer to the upper and lower kinematic regions, but these values are not used in the final

¹ The background from $K^+ \rightarrow e^+ \pi^0 \nu_e$ can generate negative reconstructed squared missing mass but with two photons and an electron in the final state it is easily rejected so its contribution is negligible.

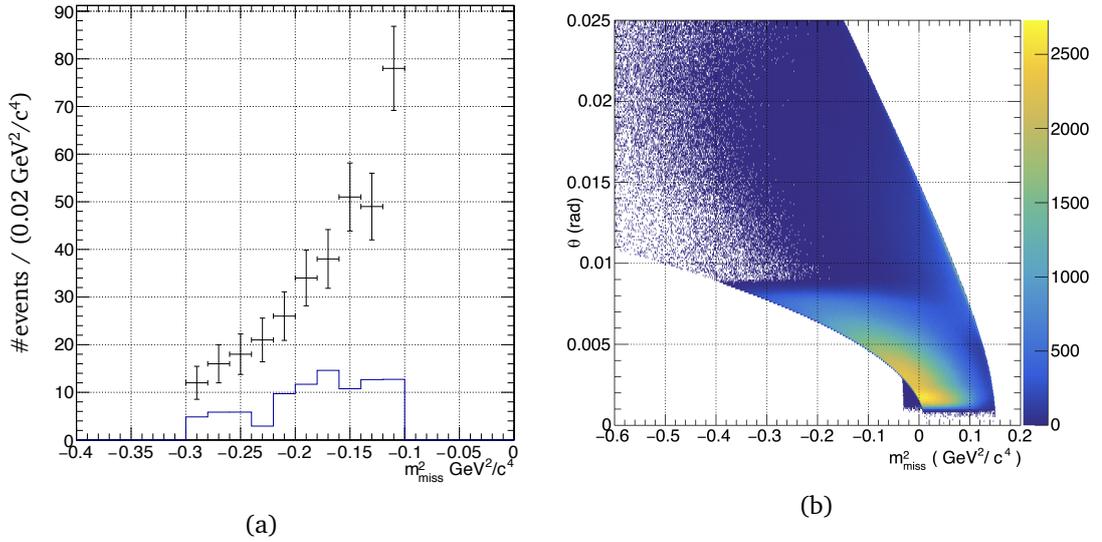


Figure 3.38.: (a) The negative region of the squared missing mass distributions for K^+ (points with error bars) and K^- (solid line) data periods, after applying the full selection in the upper kinematic region. (b) The distribution of reconstructed squared missing mass and muon opening angle w.r.t the kaon trajectory for all halo events in the K^- period.

analysis, only as a check on the scale factors obtained below.

A more precise and reliable estimate of the halo scale factors can be made by choosing a different selection which is still dominated by the halo contribution but is more similar to the signal selection. A control sample was selected by using the final signal selection but inverting the CDA cut to select only events with high CDA. This means that events from real kaon decays in the fiducial region were excluded, since these events typically have low CDA, leaving a sample dominated by the halo. For the lower region, the accepted range was $2 \text{ cm} < \text{CDA} < 5 \text{ cm}$ while for the upper region, the range was $5 \text{ cm} < \text{CDA} < 8 \text{ cm}$. In addition, the selection was restricted to high squared missing mass: $0.05 \text{ GeV}^2/c^4 < m_{\text{miss}}^2 < 0.14 \text{ GeV}^2/c^4$, covering the signal region for heavy neutrinos. MC simulation showed that for this selection, the contribution from kaon decays was dominated by the contribution from the halo. The MC simulation of the small kaon decay contribution was subtracted from the K^+ data distribution before it was used to scale the K^- distribution.

Figure 3.40 shows the squared missing mass distributions for the control selection, in which the K^- distributions have been scaled to match the K^+ distributions. The four scale factors applied to the K^- data were determined with a χ^2 fit to the ratio of the number of K^+ and K^- in each bin, and

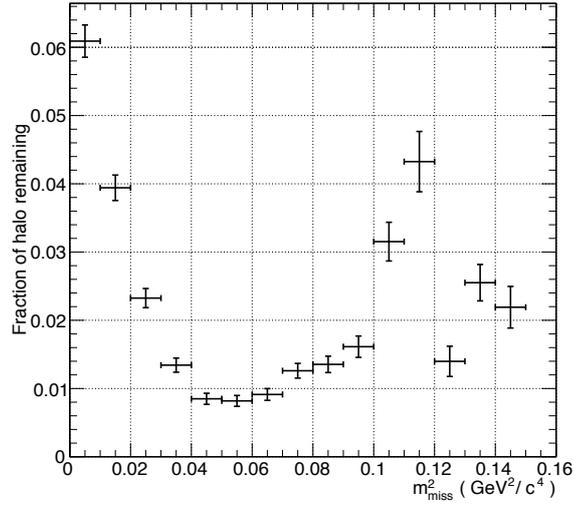


Figure 3.39.: The fraction of halo events which passed the final selection, relative to those which passed the R_K selection.

were found to be:

$$s_{\text{halo}}(+\text{ve upper}) = 3.25 \pm 0.20 \quad (\text{stat.}) \quad (3.55)$$

$$s_{\text{halo}}(-\text{ve upper}) = 2.22 \pm 0.14 \quad (\text{stat.}) \quad (3.56)$$

$$s_{\text{halo}}(+\text{ve lower}) = 5.58 \pm 0.35 \quad (\text{stat.}) \quad (3.57)$$

$$s_{\text{halo}}(-\text{ve lower}) = 4.04 \pm 0.29 \quad (\text{stat.}) \quad (3.58)$$

The scale factors obtained from the inverted-CDA-cut control sample can be compared with the scale factors obtained from the negative squared missing mass region in equations 3.51 to 3.54. For the positive spectrometer polarity data periods, the two scale factors are in agreement, within errors. For negative polarity, the scale factors from the inverted CDA control sample are significantly smaller than those obtained from the negative squared missing mass sample. Since the control sample is more similar to the final signal selection, the smaller factors are used and an appropriate error is attributed to the scale factor (described in section 3.10.5).

The final estimate of the halo contribution to the final signal selection (see section 3.8) is shown in figure 3.41. Figure 3.39 shows that relative to the R_K selection, between 1% and 5% of the halo background remains in the signal region. The systematic errors associated with the remaining halo contribution are discussed in section 3.10.5. The signal acceptance after applying the cuts described here is shown in section 3.8.

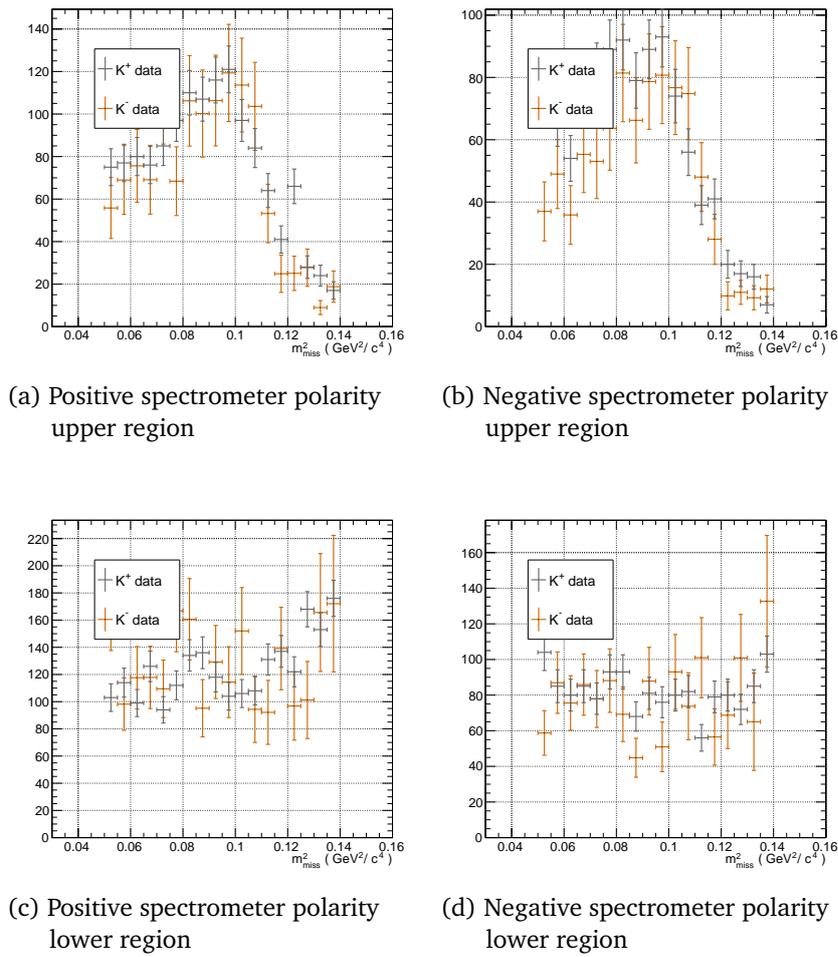


Figure 3.40.: Squared missing mass distributions for the control sample used to scale the halo estimate from K^- data to match the K^+ data. All plots show $\#events/(0.005 \text{ GeV}^2/c^4)$

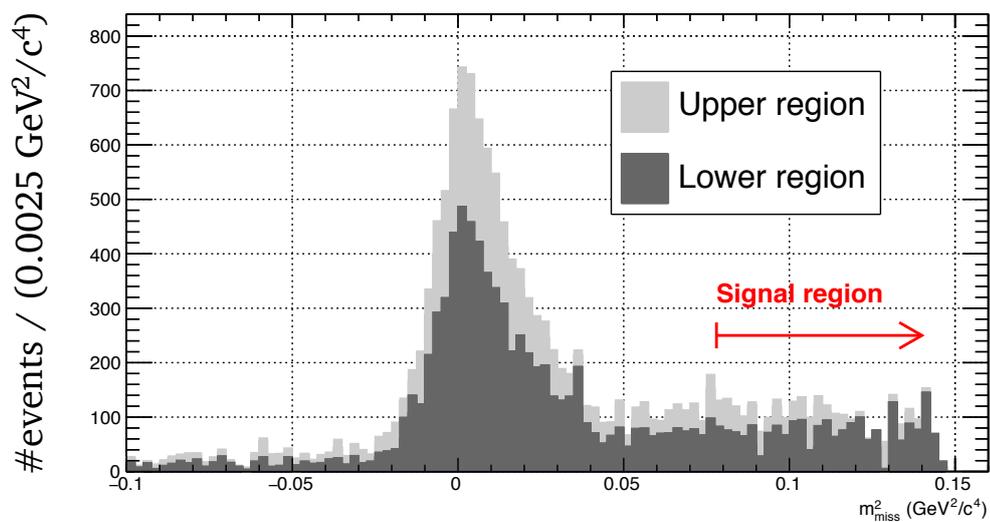


Figure 3.41.: The final estimate of the halo contribution to the final signal selection. (The signal region begins at $m = 275 \text{ MeV}/c^2$, $m^2 = 0.75 \text{ GeV}^2/c^4$.)

3.8. The final signal selection

The final selection is designed to achieve the lowest possible expected limit on the branching ratio to heavy neutrinos with masses in the range 300 to 350 MeV/c². Ideally, the selection should minimize the number of background events in the signal region while maximizing the acceptance for signal events. The following sections describe the selections which are applied in addition to the single track pre-selection described in section 3.3.1

3.8.1. Track requirements

- Track charge must be +1.
- Track quality must be > 0.7.
- Track CDA with respect to the kaon beam axis should be < 3.0 cm.
- Track should be within the spectrometer acceptance defined by $12 \text{ cm} < r_{\text{DCH1}} < 115 \text{ cm}$; $14 \text{ cm} < r_{\text{DCH4}} < 115 \text{ cm}$, in order to ensure that the charged muon does not pass through the beam pipe at too shallow an angle.
- Track should be in time with the trigger: $|t_{\text{DCH}} - t_{\text{CHOD}}| < 20 \text{ ns}$
- The track should be within the LKr acceptance, and at least 8 cm away from any missing or malfunctioning CPDs (Calorimeter Pipeline Digitizers)
- Track should be in the MUV acceptance, defined by an octagonal set of cuts when the track is extrapolated to the front face of MUV1, in xy and $u = 1/\sqrt{2}(x + y)$ and $v = 1/\sqrt{2}(x - y)$. All the values $|x|$, $|y|$, $|u|$, and $|v|$ should be < 135 cm. Additionally at this point, the track radius, $\rho = \sqrt{x^2 + y^2}$ should satisfy $\rho > 15 \text{ cm}$, since the hole for the beam pipe reduces the MUV efficiency inside this radius.

3.8.2. Photon veto

- Clusters in the LKr calorimeter are ignored if they have energy < 2 GeV or are out of time w.r.t to the track time by > 12 ns. Clusters consistent with Bremsstrahlung from the track upstream of the spectrometer (within 6 cm of the extrapolated track) are also ignored. Any remaining cluster should be consistent with the track downstream of the spectrometer, i.e. the cluster

position should be within 40 cm of the extrapolated track impact point. This loose cut on cluster position is used because the distribution of track-cluster separation is not well reproduced in MC so a tight cut would lead to an unreliable background estimate.

- The track cluster may not be in either of the ‘hot’ cells 56 and 57 in CPD 134, and must be at least 2 cm away from any dead cell.

3.8.3. Muon Identification

- There should be a reconstructed muon in the MUV with status 1 or 2 (hits in both planes 1 and 2) close to the extrapolate track position (see section 3.6.1).
- The separation between reconstructed muon position and the extrapolated track position should be less than the distance given by equation 3.45, with $\xi = 4$ (the standard cut for NA62 analyses of 2007 data) .

3.8.4. Kinematic cuts

The kinematic phase space is separated into two regions in the $z\theta$ plane as described in section 3.7 and kinematic selections are applied as described in sections 3.7.4 and 3.7.5. In the upper region, the cuts are in the plane of $z\theta$, p_z and $\theta\phi$, while in the lower region cuts are only applied in the $z\theta$ plane.

3.8.5. Signal and background in the final selection

Figure 3.42 shows the squared missing mass distribution of background events passing the final heavy neutrino selection. The largest contribution to the total background in the signal region comes from $K^+ \rightarrow \mu^+ \pi^0 \nu_\mu$ decays, in which both of the photons from π^0 decay are outside of the LKr calorimeter acceptance. The halo contribution varies between 5% and 10% of the total background composition, but in the limit calculation discussed in the following sections it is still the largest contribution to the uncertainty.

Figure 3.43 shows the acceptance, $a_\nu(m)$ for heavy neutrinos as a function of the assumed mass when the final selection is applied, and a signal window of 1σ (containing 68% of the signal events, see section 3.11.1) is used.

$$a_\nu(m) = \frac{n^{\text{pass}}(m)}{n^{\text{fid}}}, \quad (3.59)$$

where n^{pass} is the number of events which pass the full selection and $n^{\text{fid}} = 7.0 \times 10^5$ is the number of events simulated in the fiducial region $-1800 \text{ cm} < z_{\text{vertex}}^{\text{MC}} < 8000 \text{ cm}$, for each m_ν .

3.9. Statistical framework

In this section, the statistical methods for setting limits on the branching ratio to heavy neutrinos are outlined. In the following section, contributions to the systematic uncertainties are estimated in order to arrive at a final sensitivity for the analysis.

For each neutrino mass in the range under consideration, the calculation of the limit on the branching ratio $\mathcal{B}(K^+ \rightarrow \mu \nu_h(m))$, can be reduced to a simple counting experiment. The following arguments apply separately to each neutrino mass.

For a neutrino mass, m , a signal window can be defined in the squared missing mass spectrum, centred on m^2 . The background estimate, b , is the total number of events predicted to lie in this region, in the absence of any heavy neutrino signal. It has an associated uncertainty, σ_b . The Monte Carlo simulation also gives an estimate of the signal acceptance, ϵ : the number of signal events detected in the window divided by the total number of signal decays simulated in the fiducial region. Performing the experiment provides one more number, x , the number of events actually observed in the signal region. We require a prescription to produce a confidence interval for μ , the number of heavy neutrino decays in the fiducial region, given values of b , σ_b , ϵ and x .

For a measurement, μ , of a quantity whose true value is μ_t , there is more than one interpretation of what it means to construct a confidence interval for μ_t . In the frequentist approach, introduced by Neyman [99], a confidence interval, $[\mu_1, \mu_2]$, at confidence level $1 - \alpha$ has the property that, for a large ensemble of independent experiments, the confidence intervals will contain the true value, μ_t , in a fraction $1 - \alpha$ of experiments.

The Neyman definition of a confidence interval does not uniquely specify the values μ_1 and μ_2 . For a given experiment, there can be infinitely many possible intervals satisfying the frequentist definition, so a prescription is required to choose an interval according to some other constraints which lead to intervals which are intuitively reasonable (for example, one may wish to avoid returning empty intervals, or intervals containing non-physical parameters). Once a prescription has been defined, its coverage can be studied to ensure that the intervals produced really do contain the true value in a fraction $1 - \alpha$ of cases.

In the case of looking for a small signal, μ , on top of a small background b , the Feldman-Cousins

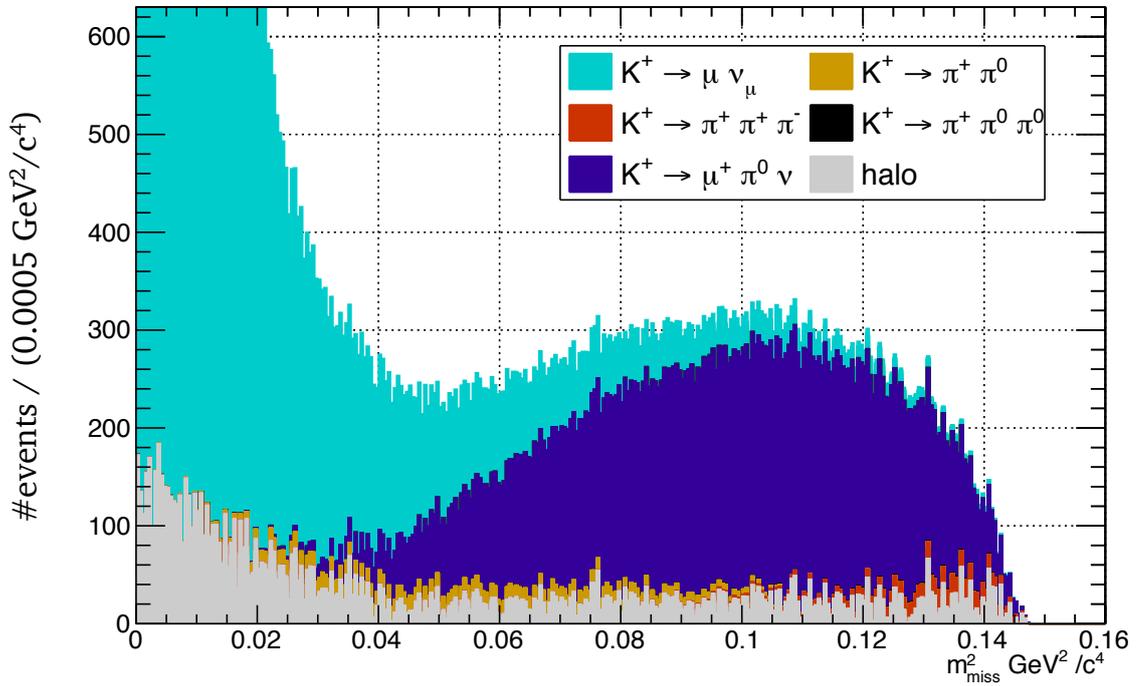


Figure 3.42.: The simulated distribution of background events for the final heavy neutrino selection

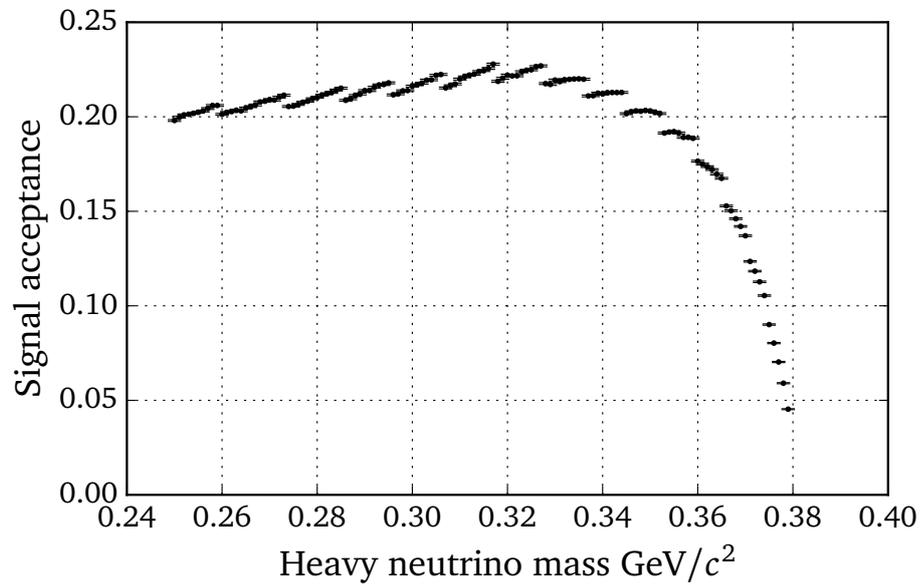


Figure 3.43.: The signal acceptance in the final heavy neutrino selection as a function of neutrino mass. Discontinuities are due to the discrete signal window sizes (see section 3.11.1).

approach [100] gives a prescription for constructing confidence intervals using likelihood ratios. For a given experimental measurement, x , one can construct the likelihood function

$$\mathcal{L}(x|\mu) \equiv P(x|\mu), \quad (3.60)$$

which is the probability of obtaining the measurement x , given some possible signal, μ . Feldman and Cousins define the ratio:

$$R = \frac{\mathcal{L}(x|\mu)}{\mathcal{L}(x|\hat{\mu})}, \quad (3.61)$$

where $\hat{\mu}$ is the particular choice of μ which maximises $\mathcal{L}(x|\mu)$ while still being in the physically allowed region. In the present example of a counting experiment with background, $\hat{\mu} = \max(0, x - b)$. A confidence interval can be constructed by taking values of μ_i ordered (descending) by the corresponding value of R , until the total probability reaches the desired confidence level: $\sum_i P(x|\mu) = \alpha$. The resulting confidence intervals have good coverage and naturally produce either upper limits or two-sided limits depending on the value of x .

The Feldman-Cousins approach assumes that the background, b , is known exactly, so some extension is required for cases in which the true value of the background, b_t is also unknown. The true value of b_t is called a nuisance parameter since it must be included in the model but is not of direct interest.

The approach of Rolke and Lopez [101] makes use of the joint likelihood $\mathcal{L}(x, y|\mu, b)$, which is the probability of observing x events in the signal region and y events in the background estimate given real signal μ and background b . For example, in the present analysis, the number of background events observed is assumed to follow a Gaussian distribution, with mean b and variance σ_b so the likelihood function for a given hypothesis of signal and background is:

$$\mathcal{L}(x, y|\mu, b) = \frac{(\mu + b)^x}{x!} e^{-(\mu+b)} \frac{1}{\sqrt{2\pi}} e^{-(y-b)^2/(2\sigma_b^2)} \quad (3.62)$$

The profile likelihood function, λ , is given by the ratio:

$$\lambda(x, y|\mu) = \frac{\mathcal{L}(x, y|\mu, \hat{b}(\mu))}{\mathcal{L}(x, y|\hat{\mu}, \hat{b})}. \quad (3.63)$$

In the numerator, $\hat{b}(\mu)$ is the value of b which maximises $\mathcal{L}(x, y|\mu, b)$ for the measured values x, y and the value of μ under consideration. In the denominator, \hat{b} and $\hat{\mu}$ are the values of b and μ which jointly maximize $\mathcal{L}(x, y|\mu, b)$ for the measured values of x and y . Values of $\lambda(x, y|\mu)$ lie in the

interval $[0, 1]$.

It can be shown that in the large sample limit, $-2 \ln \lambda$ follows a χ^2 distribution with n degrees of freedom, where n is the number of parameters of interest (in this case, just 1, the value of μ). This means that the confidence interval for μ at C.L. $(1 - \alpha)$ can be extracted by starting at $\mu_1 = \mu_2 = \hat{\mu}$ (i.e. where $-2 \ln \lambda$ is minimum) and expanding the interval $[\mu_1, \mu_2]$ to the points where $-2 \ln \lambda$ has risen by an amount corresponding to $(1 - \alpha)$ in the appropriate cumulative χ^2 distribution. For example, for a 90% confidence interval, $-2 \ln \lambda$ should rise by 2.706 because $\int_0^{2.706} \chi_1^2(x) dx = 0.90$.

It will sometimes be the case that, with this approach, μ_1 or both μ_1 and μ_2 are < 0 , i.e. outside the physically allowed region. In these cases, the lower limit μ_1 is set to 0 and the upper limit μ_2 is determined using the method described above but starting from 0 rather than $\hat{\mu}$. This approach is called the method of bounded likelihood.

The previous sections have shown the various studies used to determine the b , the estimate of the background in each mass window. The following sections will outline the main contributions to σ_b , the error on this estimate. The final uncertainty on the background estimate σ_b is taken as the sum in quadrature of the contributing uncertainties, and, by the Central Limit Theorem, is assumed to follow a Gaussian distribution.

3.10. Systematic Uncertainties

3.10.1. Spectrometer tails

The far tails of the spectrometer resolution were studied in section 3.5, and an upper limit on the contribution from multiple Coulomb scattering was obtained as shown in figure 3.17. For the final analysis, no additional scattering is applied to the MC simulation of the signal region. Instead the whole contribution which would be produced by simulating scattering with the maximum parameters obtained in section 3.5 is treated as a systematic error. Figure 3.44 shows the relative contribution in the signal region as a function of squared missing mass.

3.10.2. Total number of kaon decays

The total number of kaon decays in the fiducial region, N_K^{fid} , cancels in the calculation of the branching ratio of decay to heavy neutrinos because the limits are normalized w.r.t. the $K^+ \rightarrow \mu^+ \nu_\mu$ decay. However, N_K^{fid} must be estimated so that the MC simulation of the expected background can be scaled

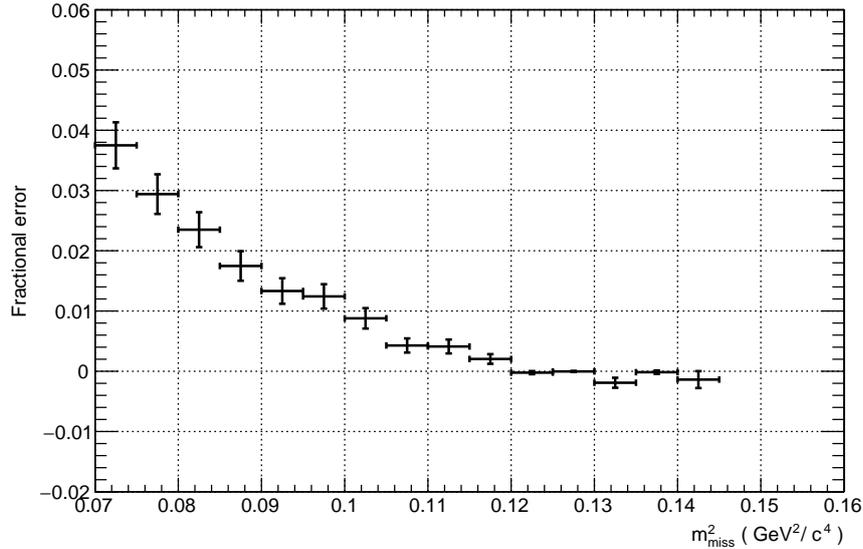


Figure 3.44.: The fractional error on the background estimate coming from the simulation of the far tails of the spectrometer resolution. The two bins with higher content at high squared missing mass are the result of statistical fluctuation in the MC estimate of the spectrometer contribution. The contribution at these masses is negligible compared to other systematic errors (see figure 3.50).

appropriately. The flux is estimated from the $K^+ \rightarrow \mu^+ \nu_\mu$ peak using a ‘loose selection’ with looser cuts than the final signal selection in order to reduce the systematic error from cutting on the tails of poorly simulated kinematic distributions.

The selection is the same as the final signal selection with the exceptions:

- muon veto requirement is removed
- all cuts in momentum (p), z -vertex, and opening angle (θ) are removed and replaced with the cuts in the p - z plane shown in table 3.4 (page 84).

The estimated kaon flux in the fiducial region is defined as the number of kaon decays between $z = -1800$ cm and $z = 8000$ cm. The number of $K^+ \rightarrow \mu^+ \nu_\mu$ decays in data is determined by integrating the number of events in the peak at zero in the squared missing mass distribution, after subtracting the contribution from the beam halo, and the acceptance is determined from the MC simulation of $K^+ \rightarrow \mu^+ \nu_\mu$. The total number of kaon decays is obtained by dividing by the $K^+ \rightarrow \mu^+ \nu_\mu$ branching ratio (0.6355).

Figure 3.45 shows the squared missing mass distribution (3.45a) and the data/MC ratio (3.45b) for the loose selection described above. It has a width (containing 68% of the peak) $\sigma_{m^2} =$

$0.0032 \text{ GeV}^2/c^4$. The cuts on the squared missing mass are chosen to include the tails of the $K^+ \rightarrow \mu^+ \nu_\mu$ distribution so that the final estimate is insensitive to the exact shape of the peak. In figures 3.45c and 3.45d the relative variation in the estimated kaon flux is shown as a function of the cut in squared missing mass, for positive and negative spectrometer polarities. There is little variation after $0.015 \text{ GeV}^2/c^4$ (about $5\sigma_{m^2}$) so this value is chosen for the cut: $|m_{\text{miss}}^2| < 0.015 \text{ GeV}^2/c^4$.

An estimate on the systematic uncertainty associated with the integrated kaon flux estimate was obtained by varying the cut on $|m_{\text{miss}}^2|$ by $1\sigma_{m^2}$. The variation in estimated integrated kaon flux was 0.2%. Another estimate comes from comparing the integrated flux measured using the loose selection described in this section with the integrated flux determined using the final signal selection. The difference in estimates was also 0.2%. This uncertainty is comparable with the precision on $\mathcal{B}(K^+ \rightarrow \mu^+ \nu_\mu)$ in the PDG [8] which is 0.15%. Therefore the systematic uncertainty on the estimated integrated kaon flux is taken to be $\sqrt{0.2^2 + 0.15^2} = 0.25\%$. The estimated values, separated by polarity are:

$$N_K^{\text{fid}}(+\text{ve polarity}) = (3.630 \pm 0.009) \times 10^7 \quad (3.64)$$

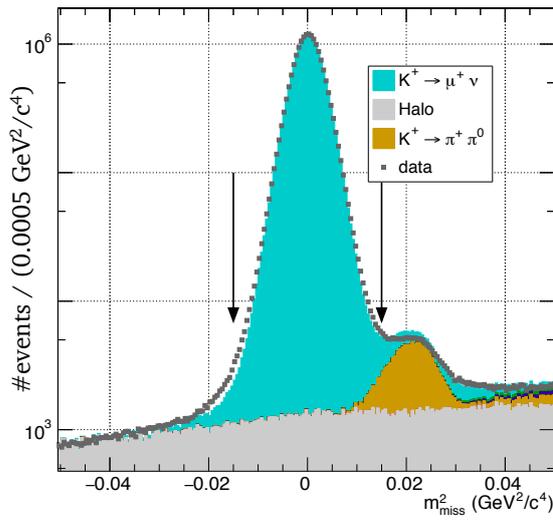
$$N_K^{\text{fid}}(-\text{ve polarity}) = (2.347 \pm 0.006) \times 10^7 \quad (3.65)$$

3.10.3. Muon veto

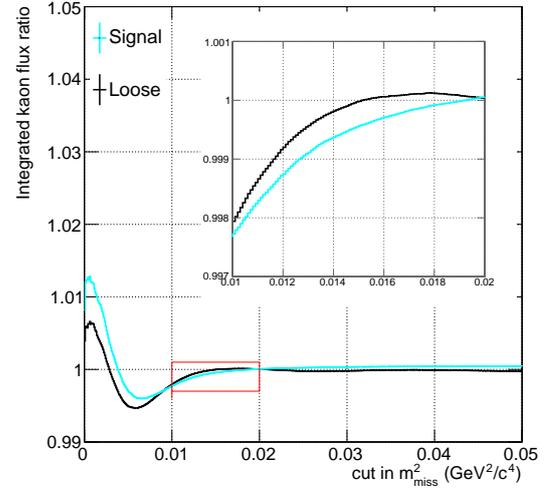
Section 3.6 described how the muon veto efficiency is measured and modelled in MC as a function of track momentum and impact point on the MUV planes. The residual disagreement between MC and data can be estimated by comparing MUV efficiency as a function of other kinematic variables.

Figure 3.46 shows the comparison between data and MC for the muon selection applied to K^+ data as described in section 3.6, after applying the measured MUV efficiency as a weight in MC. The residual plots show that the difference between data and MC is typically $< 0.5\%$ but at the extremes of the kinematic distributions it can rise to 1%. In the final analysis, the residual difference included as a systematic error on the background estimate. For a given signal window, denoting the number of expected background events N_{BG} , the uncertainty, σ_{MUV} , associated with the MUV efficiency measurement is given by:

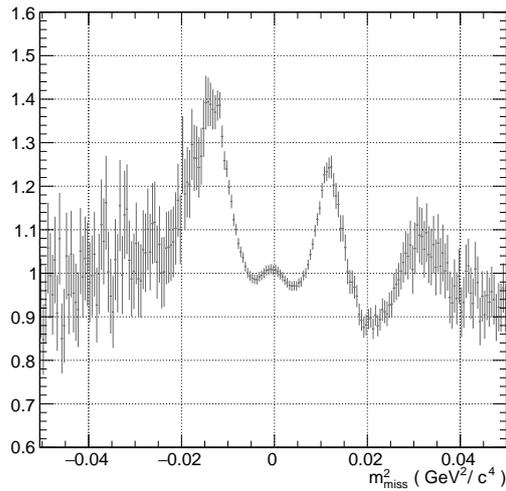
$$\sigma_{\text{MUV}} = 1\% \times N_{\text{BG}}. \quad (3.66)$$



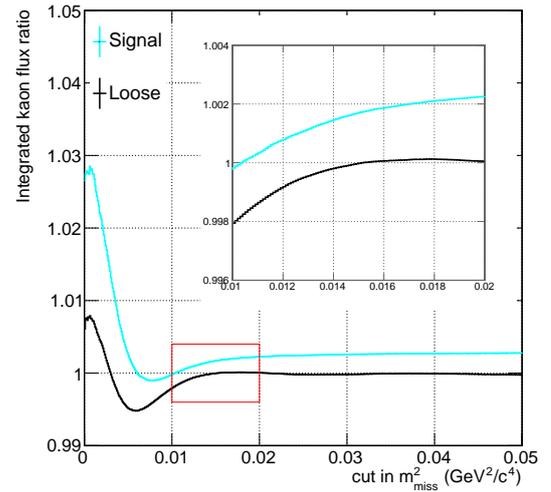
(a)



(c)



(b)



(d)

Figure 3.45.: Extraction of the integrated kaon flux and integrated uncertainty. Figure (a) shows the $K^+ \rightarrow \mu^+ \nu_\mu$ peak obtained with a loose selection used to measure the kaon flux (arrows indicate the cuts applied in m^2_{miss}). The bin by bin data/MC ratio for this selection is shown in figure (b). The data/MC ratio for the integral of the peak at $m^2_{\text{miss}} = 0$ as a function of the peak width is shown in figures (c), for positive parity, and (d), for negative parity, and the data/MC ratios for the final signal selection are superimposed (in light blue). The normalization is chosen so that the ratio for the loose selection is 1 when the half-width is $0.02 \text{ GeV}^2/c^4$. The region around the final cut at $0.015 \text{ GeV}^2/c^4$ is shown, enlarged and inset.

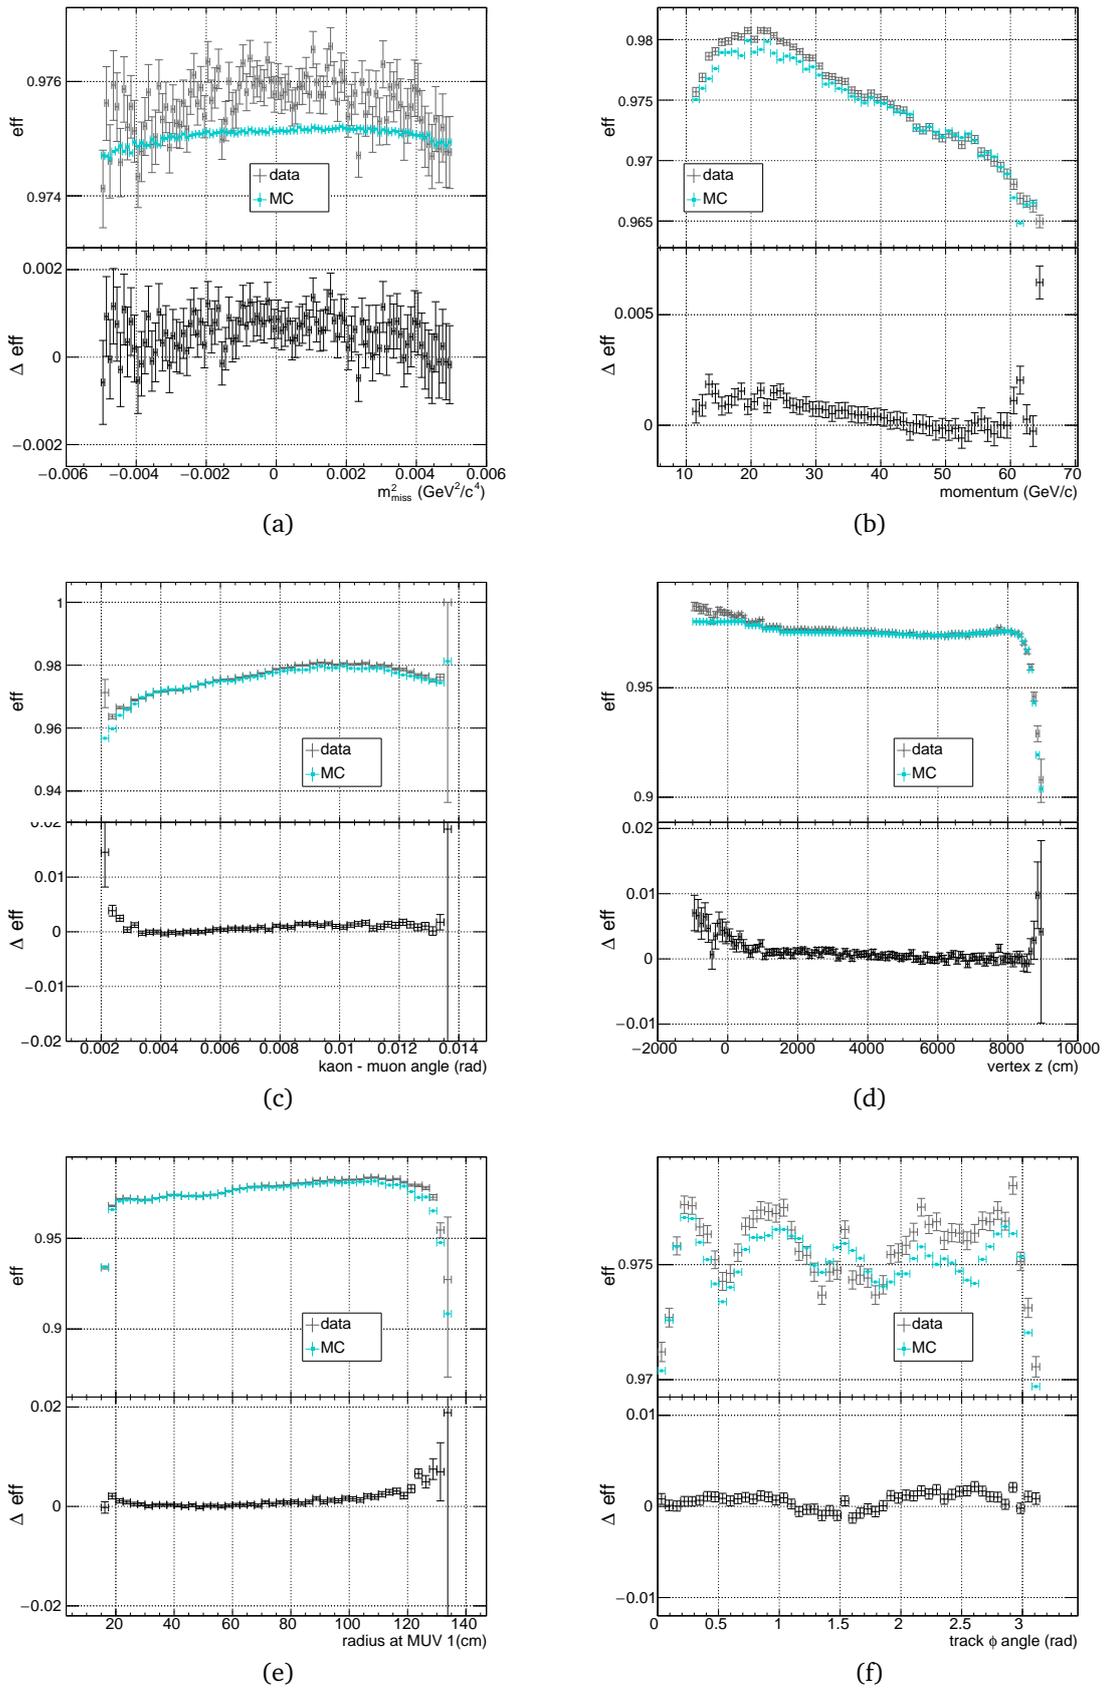


Figure 3.46.: Muon veto efficiency in various kinematic projections: (a) squared missing mass, (b) momentum, (c) muon - kaon opening angle, θ , (d) vertex z -coordinate, (e) track radius at MUV 1, (f) track azimuthal angle ϕ

3.10.4. Trigger efficiency

There were two triggers running in the K^+ data taking period which could be used for selecting heavy neutrinos. The minimum bias trigger Q_1 was downscaled by a factor of 600, while the $Q_1 \times 1\text{TRKL}$ trigger had a downscaling factor of 150. Using the $Q_1 \times 1\text{TRKL}$ trigger increases the available sample size by a factor of 4 with respect to the minimum bias sample, but it also introduces a new source of error in the final limit, coming from the trigger efficiency.

It has already been seen (section 3.3.3) that decays such as $K^+ \rightarrow \pi^+\pi^0$ and $K^+ \rightarrow \pi^+\pi^+\pi^-$ can be reconstructed as signal-like events if some of the particles in the final state are outside of the detector acceptance. These events can fail the 1TRKL trigger if a π^+ or a photon from a π^0 interacts in the drift chambers and creates a shower resulting in more than 15 hits in a single view. This showering is not well simulated in MC so the trigger inefficiency must be measured separately in order to have a reliable estimate of these backgrounds.

Relative trigger efficiencies can be extracted from data because, when an event passes a trigger, it is recorded along with information about any other triggers it passed. Thus, the inefficiency of the $Q_1 \times 1\text{TRKL}$ trigger with respect to the Q_1 trigger can be measured by applying the final heavy neutrino selection to events passing the Q_1 trigger and counting how many would also have passed the 1TRKL trigger. Since the downscaling factors are large, there is little overlap between the Q_1 sample the $Q_1 \times 1\text{TRKL}$ sample, so the estimates of the trigger efficiency and the expected background in the final sample are uncorrelated. The Q_1 inefficiency is small ($<1\%$) and does not vary significantly between the kaon decay channels so it has a negligible contribution.

The relative trigger efficiency is described as a function of reconstructed missing mass:

$$\eta(m) = \frac{N_{\text{passed}}(m)}{N_{\text{all}}(m)}, \quad (3.67)$$

where $N_{\text{all}}(m)$ is the number of events which would be collected by an unbiased trigger after applying the full selection and $N_{\text{passed}}(m)$ is the number of events which would be collected by the biased trigger under consideration, again after applying the full selection. (The m is suppressed for the remainder of this section - the logic applies separately to each mass window).

In the final analysis, we require an estimate of the number of background events expected when using the $Q_1 \times \text{TRKL}$ trigger, $B_{Q_1\text{T}}^{\text{est}}$. This number is estimated from the raw number of background

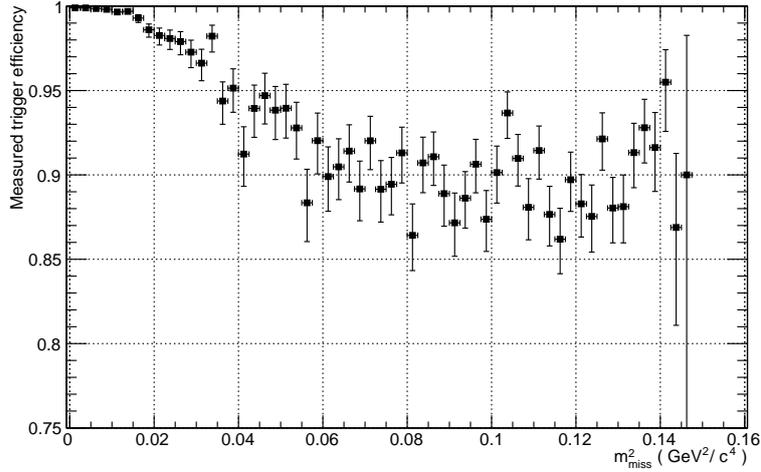


Figure 3.47.: The $Q_1 \times 1\text{TRKL}$ trigger efficiency, measured using the Q_1 data sample.

events predicted by MC and data driven methods, B_{raw} , and the measured trigger efficiency, η_{meas} :

$$B_{Q11T}^{\text{est}} = \eta_{\text{meas}} B_{\text{raw}} \quad (3.68)$$

The measurement of η introduces an additional uncertainty on the value of B_{Q11T}^{est} , which must be included in the final limit calculation. A measurement of $\eta = N_{\text{passed}}/N_{\text{all}}$, follows a binomial distribution: N_{passed} is the number of successes from N_{all} trials. The uncertainty on η is given by (in one prescription) the Clopper-Pearson interval [102], which provides fully frequentist estimates for the upper and lower limits at a confidence level $1 - \alpha$, given the values of N_{passed} and N_{all} :

$$\beta\left(\frac{\alpha}{2}; N_{\text{passed}}, N_{\text{all}} - N_{\text{passed}} + 1\right) < \eta < \beta\left(1 - \frac{\alpha}{2}; N_{\text{passed}} + 1, N_{\text{all}} - N_{\text{passed}}\right), \quad (3.69)$$

where $\beta(p; \nu, w)$ is the p th quantile from a beta distribution with shape parameters ν and w . In the limit calculation for a given neutrino mass, the Clopper-Pearson interval is calculated using $\alpha = 0.68$ and half the interval is taken as an uncertainty, σ_{η} , which is included in the uncertainty on the background estimate. The measured $Q_1 \times 1\text{TRKL}$ trigger efficiency using the Q_1 data sample is shown in figure 3.47.

3.10.5. Halo

The uncertainty on the estimated halo distribution can be split into statistical and systematic uncertainties. The statistical uncertainty comes from the limited sample size of the K^- data taking period, while the systematic uncertainties come from limitations in the model: i.e. the assumption that the halo distribution in K^- data taking periods accurately reproduces that of the K^+ periods.

For a given mass window, the total number of halo events, n_{halo} is determined from the raw number of events detected in the K^- period, n_{K^-} , for each spectrometer polarity and kinematic region, after subtracting the simulated number of accepted $K^- \rightarrow \pi^- \pi^- \pi^+$ decays, $n_{K3\pi}$:

$$n_{\text{halo}} = \sum_{\text{polarity}} \sum_{\text{region}} s_{\text{halo}} (n_{K^-} - s_{K3\pi} n_{K3\pi}). \quad (3.70)$$

Here, s_{halo} is the scale factor used to scale the halo estimate in order to match the K^+ data (see section 3.7.9) and $s_{K3\pi}$ is the scale factor applied to the MC sample of $K^- \rightarrow \pi^- \pi^- \pi^+$ decays in order to match the K^- data sample (see section 3.7.8). The scale factor $s_{K3\pi}$ is calculated from the branching ratio of $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ and the size of the MC sample:

$$s_{K3\pi} = BR(K^- \rightarrow \pi^- \pi^- \pi^+) \frac{N_{K^-}^{\text{data}}}{N_{K3\pi}^{\text{MC}}} r_{K3\pi}, \quad (3.71)$$

where, $N_{K^-}^{\text{data}}$ is the number of K^- decays in K^- data taking period, $N_{K3\pi}^{\text{MC}}$ is the number of simulated $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ decays and $r_{K3\pi}$ is the correction factor defined in equation 3.50 (on page 103).

Statistical uncertainties: Considering only the statistical uncertainties for n_{K^-} and $n_{K3\pi}$, the contribution to the uncertainty on the halo distribution is:

$$(\sigma_{\text{halo}}^{\text{stat}})^2 = \sum_{\text{polarity}} \sum_{\text{region}} s_{\text{halo}}^2 (n_{K^-} + s_{K3\pi}^2 n_{K3\pi}) \quad (3.72)$$

$K^+ \rightarrow \pi^+ \pi^+ \pi^-$ scale factor: The scale factor, $r_{K3\pi}$, used for scaling the simulation of $K^\pm \rightarrow \pi^\pm \pi^\pm \pi^\mp$ (to account for un-simulated detection inefficiency) was estimated using the negatively charged beam halo in the K^+ data taking period, as shown in figure 3.34 (on page 102). The value of $r_{K3\pi}$ was chosen as the one which minimizes the χ^2 between the K^+ distribution and the distribution

from the sum of K less data and the $K^- \rightarrow \pi^- \pi^- \pi^+$ simulation . The value obtained was:

$$r_{K3\pi}(K^+ \text{ period, upper region}) = 0.583 \pm 0.017, \quad \chi^2/\text{nDOF} = 21.3/24. \quad (3.73)$$

(Only the upper kinematic region was used as the contribution from $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ in the K^+ data sample was negligible in the lower region). The same scale factor was then shown to produce consistent results when used for K^- data taking periods in figure 3.35 (on page 104) and with the full signal selection applied in figure 3.36.

In order to quantify the consistency, the scale factor was determined independently for the K^- data taking periods. For the K^- data samples there was a non-negligible contribution from $K^- \rightarrow \pi^- \pi^- \pi^+$ in both the upper and lower kinematic regions so two measurements could be made. The values obtained were:

$$r_{K3\pi}(K^- \text{ period, upper region}) = 0.561 \pm 0.020, \quad \chi^2/\text{nDOF} = 16.3/19, \quad (3.74)$$

$$r_{K3\pi}(K^- \text{ period, lower region}) = 0.578 \pm 0.060, \quad \chi^2/\text{nDOF} = 16.6/19. \quad (3.75)$$

These values are in agreement with the value measured in the K^+ data taking period, and in all cases the fit χ^2 is less than 1 per DOF so it is concluded that it is valid to assume a single scale factor for the $K^\pm \rightarrow \pi^\pm \pi^\pm \pi^\mp$ inefficiency.

The original value from the K^+ period is used as the final input to the measurement along with the associated error from the χ^2 fit. The contribution to the final uncertainty on the halo background is:

$$(\sigma_{\text{halo}}^{r_{K3\pi}})^2 = \sum_{\text{polarity}} \sum_{\text{region}} \left(s_{\text{halo}} n_{K3\pi} \frac{\sigma_{r_{K3\pi}}}{r_{K3\pi}} \right)^2 \quad (3.76)$$

The source of this detection inefficiency has not been fully accounted for. In particular it has not been conclusively demonstrated that it can be described by a single scale factor so, although the fit χ^2 is good, the uncertainty associated with $r_{K3\pi}$ may be underestimated in this approach.

Halo Scale factor between K^+ and K^- data: The final scale factors applied to the K^- data for the halo estimate in order to match the distribution in K^+ data were determined from the control sample and listed in equations 3.55 to 3.58 (page 107). An alternative approach was also considered, using events with negative reconstructed squared missing mass (equations 3.51 to 3.54).

A systematic error is assigned to each of the four scale factors, by comparing the value obtained

from the control sample with the value obtained from the negative missing mass region. The relative error is determined by finding the scale factor which each value would have to be multiplied or divided by in order to make the values consistent within the statistical error bars:

$$\sigma_s^{\text{rel}} = \sqrt{\frac{s^{\text{neg}} - \sigma_{s^{\text{neg}}}}{s^{\text{ctrl}} + \sigma_{s^{\text{ctrl}}}} - 1}, \quad (3.77)$$

where s^{neg} is the scale factor obtained using events in the negative squared mass region selected with the final signal selection and s^{ctrl} is the scale factor obtained using the control selection.

The final values for each region are:

$$s_{\text{halo}}(+\text{ve, upper}) = 3.25 \pm 0.50 \quad (3.78)$$

$$s_{\text{halo}}(-\text{ve, upper}) = 2.22 \pm 0.59 \quad (3.79)$$

$$s_{\text{halo}}(+\text{ve, lower}) = 5.58 \pm 0.86 \quad (3.80)$$

$$s_{\text{halo}}(-\text{ve, lower}) = 4.04 \pm 1.20, \quad (3.81)$$

where +ve and -ve refer to data samples with positive and negative spectrometer polarity and upper and lower refer to the upper and lower kinematic regions.

Halo model validity check: The validity of using K^- data to model the halo background in K^+ data taking periods is assessed using the control sample described in section 3.7.9. After applying the control selection and scaling the K^{less} data by the appropriate factor, the K^{less} and K^+ data samples are compared in figure 3.48. The left-hand column shows several kinematic distributions and in each case, the two samples are superimposed, while the centre column shows the ratio of the two samples, with statistical errors. The ratio is fit to a constant value shown with a horizontal line. In the right hand column, the errors on each point are scaled up until the χ^2 per degree of freedom for the fit is < 1 . For the CDA, m_{miss}^2 and momentum distributions, no scaling has to be applied to the error bars to achieve $\chi^2/DOF < 1$. For the opening angle and z -vertex distributions, a scale factor of 1.15 is required. This value is used as an additional relative error on the halo estimate.

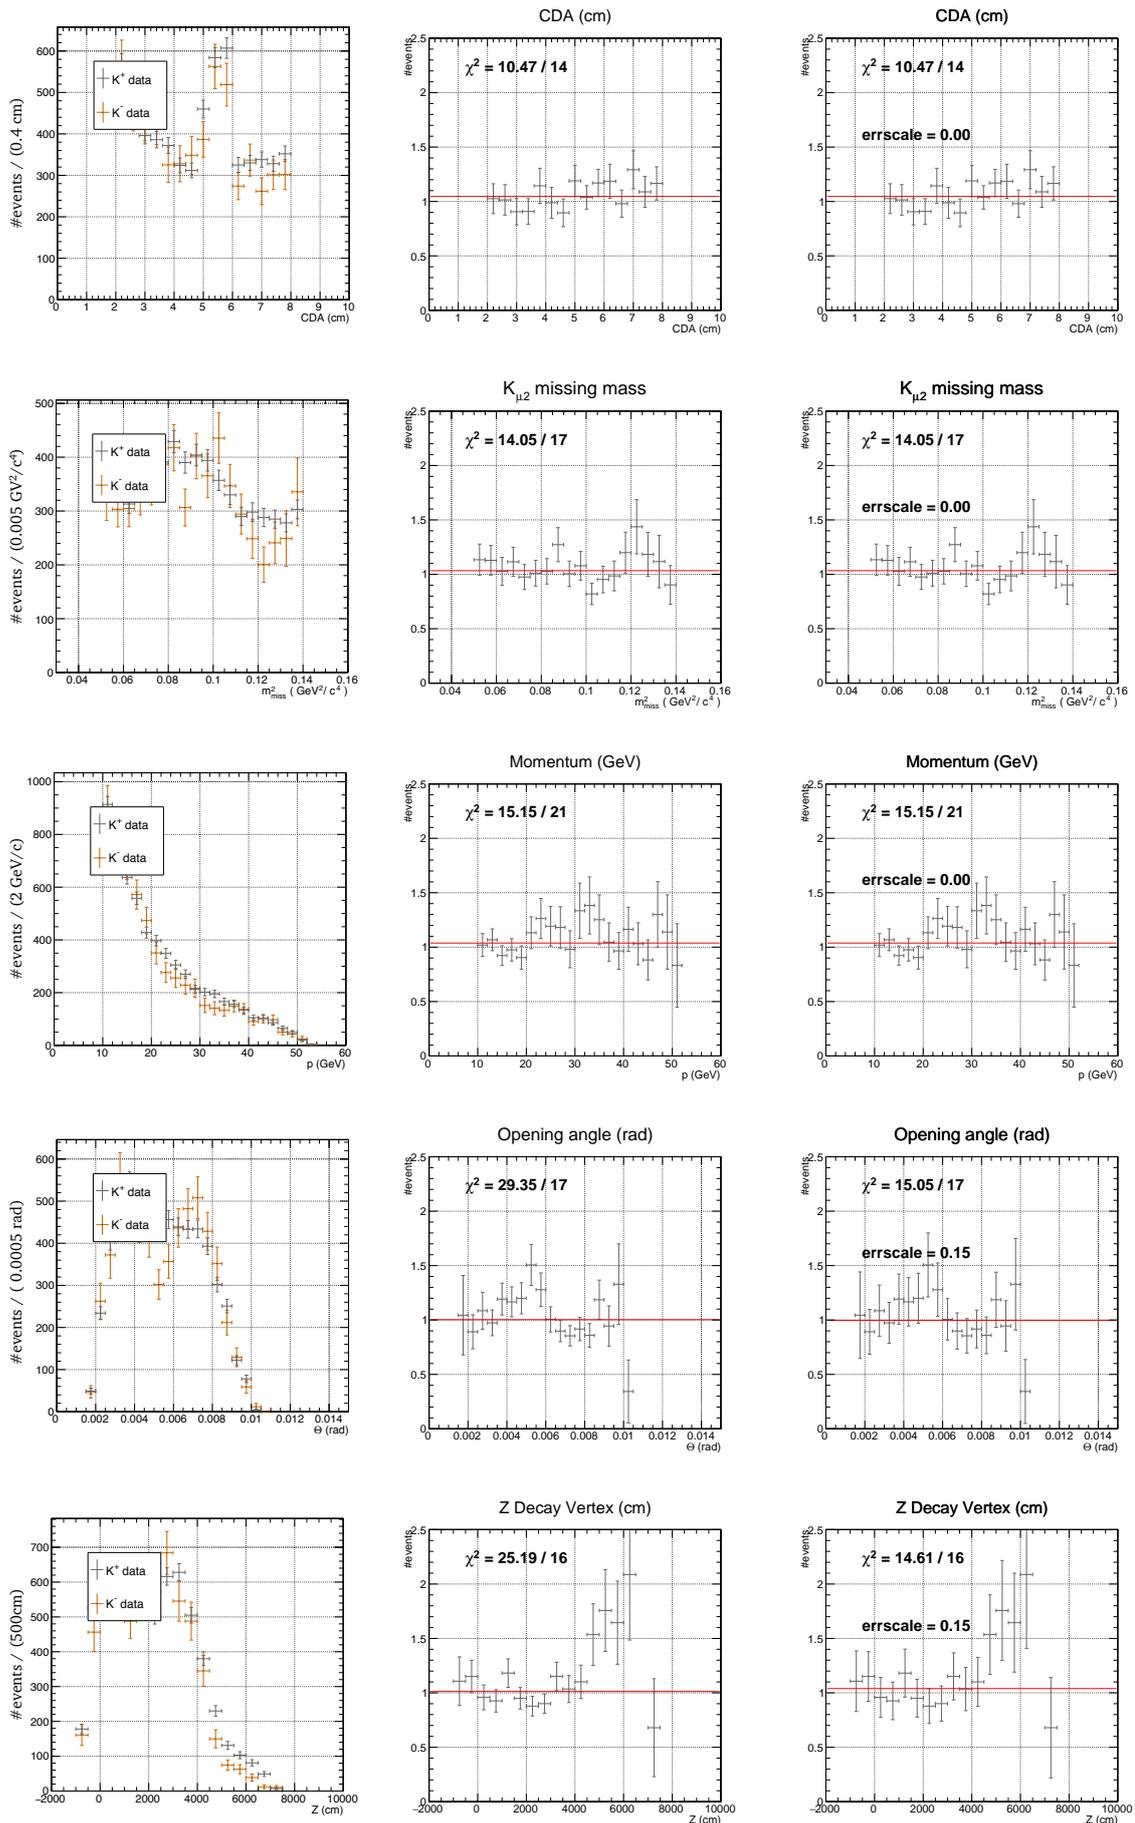


Figure 3.48.: Simulated kinematic distributions used for the halo model validity check described on page 3.10.5. The quantity 'errscale' is $1 - f$, where f is the scale factor required to achieve a good χ^2 value.

3.11. Limits on the branching ratio to heavy neutrinos

In this section, the width of the signal window for each neutrino mass is determined by maximizing the *expected sensitivity*: the upper limit which would be set if the number of events observed in data was exactly equal to the number of expected background events, (the *Asimov data set*). The expected sensitivity is presented along with relative contributions of the systematic errors described in the previous sections. Finally, the data observed in the signal region are compared with the expected backgrounds in order to set limits on the branching ratio to heavy neutrinos.

3.11.1. Signal window width

A scan for a heavy neutrino signal is performed by generating heavy neutrino data samples with MC simulation with neutrino masses at $1 \text{ MeV}/c^2$ intervals, from $275 \text{ MeV}/c^2$ to $375 \text{ MeV}/c^2$. Once the background distribution and the signal acceptance have been modelled, the expected sensitivity depends on the choice of the width of the signal window in reconstructed squared missing mass. Increasing the width of the window increases the final acceptance for heavy neutrinos, but also increases the number of background events expected, so there is an optimum width which balances the two effects.

The signal window in m^2 for a given neutrino mass, m_h , was defined in terms of the signal width:

$$(m_h^2 - n\sigma_h) < m^2 < (m_h^2 + n\sigma_h), \quad (3.82)$$

where σ_h is defined such that the range $(m_h^2 - \sigma_h)$ to $(m_h^2 + \sigma_h)$ contains 68.3% of the expected signal events, and n is the parameter to be optimized. For each window, the width was rounded to the nearest multiple of $10^{-4} \text{ GeV}^2/c^4$ (leading to the discontinuities in acceptance shown in figure 3.43). In order to determine the optimum width for a given heavy neutrino mass the expected sensitivity was calculated as function of m_h in steps of 0.05. The optimum width found for each neutrino mass is shown in figure 3.49. For masses below $280 \text{ MeV}/c^2$, where there is a large background uncertainty from the spectrometer resolution, the optimum signal width tends to lower values, but throughout the signal region above $300 \text{ MeV}/c^2$, the optimum signal width fluctuates around a stable value of $1 \sigma_h$. In order to avoid over-tuning the selection to fluctuations in the background estimate, the signal window is defined as $\pm 1 \sigma_h$ for all neutrino masses. This signal window definition was used for the plot of acceptance as a function of simulated heavy neutrino mass (figure 3.43, page 112).

3.11.2. Final sensitivity

Figure 3.50 shows the squares of the uncertainties which contribute to the final upper limit on the number of heavy neutrinos observed for each of the masses considered. The total squared error on the background estimate is greater than the square of the expected fluctuation of the data around the background value (the blue line). However, the largest contribution to the background uncertainty comes from the statistical uncertainty in the halo measurement, so the present result is ultimately limited by the available statistics in the K^+ and K^- data samples.

The expected sensitivity to heavy neutrinos in NA62 2007 data is determined as a function of neutrino mass using the Asimov data set in which the number of events observed is taken to be exactly equal to the number of background events predicted. Limits on the number of heavy neutrinos observed are determined using the profile likelihood approach described in section 3.9.

The upper limits on the number of heavy neutrinos observed, n_{UL} , are converted to upper limits on the branching ratio $\mathcal{B}(K^+ \rightarrow \mu^+ \nu_h)$ using the relation:

$$n_{\text{UL}} = B_{\text{UL}}(K^+ \rightarrow \mu^+ \nu_h) a(m_h) N_K, \quad (3.83)$$

where $a(m_h)$ is the signal acceptance, shown in figure 3.43, N_K is the total number of kaon decays in the fiducial region (see section 3.10.2), and $B_{\text{UL}}(K^+ \rightarrow \mu^+ \nu_h)$ is the upper limit on the branching ratio for the decay $K^+ \rightarrow \mu^+ \nu_h$.

The branching ratio is related to the neutrino mixing matrix element $|U_{\mu 4}|^2$ by equation 1.110 (on page 33) so the limit on $|U_{\mu 4}|^2$ is given by:

$$|U_{\mu 4}|^2 = \frac{n_{\text{UL}}}{N_K \mathcal{B}(K^+ \rightarrow \mu^+ \nu_h) a(m_h) \lambda(m_h) f_{\text{M}}(m_h)} \quad (3.84)$$

where the kinematic factors $\lambda(m_h)$ and $f_{\text{M}}(m_h)$ are described in section 1.7.

The single event sensitivity, \mathcal{B}_{SES} , for the branching ratio to heavy neutrinos is defined by:

$$\mathcal{B}_{\text{SES}} = \frac{1}{a(m_h) N_K}, \quad (3.85)$$

i.e., setting $n_{\text{UL}} = 1$ in equation 3.83. The corresponding limit on $|U_{\mu 4}|^2$ is shown in figure 3.51. The limits on $\mathcal{B}(K^+ \rightarrow \mu^+ \nu_h)$ are of order 10^{-5} (see figure 3.55a). For illustration, the MC signals for $\mathcal{B}(K^+ \rightarrow \mu^+ \nu_h) = 10^{-4}$ are shown on top of the estimated background in figure 3.52 (page 129), for

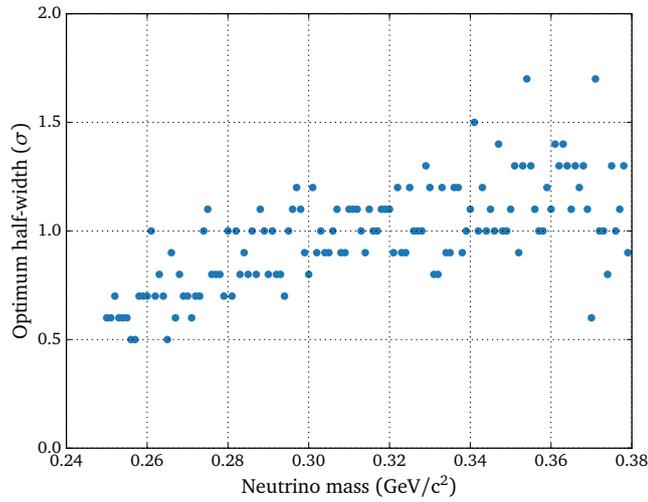


Figure 3.49.: The optimum half-width, σ_h for the heavy neutrino signal window as a function of neutrino mass. (see section 3.11.1 and equation 3.82).

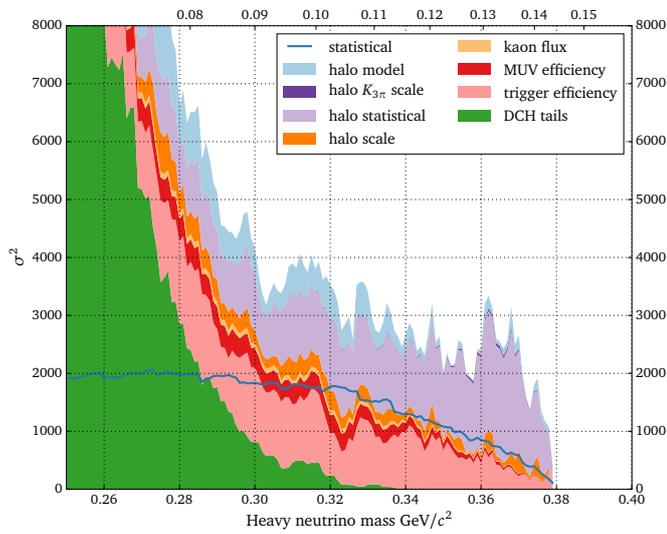


Figure 3.50.: The squared uncertainties which contribute to the final upper limit on the number of heavy neutrinos observed at each mass. The coloured bands show the contributions to the uncertainty on the background estimate (added together). The blue line shows the statistical contribution from the expected fluctuation of the data about the true background value. The contribution labelled 'halo statistical' is a separate contribution which comes from the statistical uncertainty on the estimate of the true halo contribution.

heavy neutrino masses from $270 \text{ MeV}/c^2$ to $370 \text{ MeV}/c^2$ in $20 \text{ MeV}/c^2$ intervals.

3.11.3. Results

In figure 3.53, the spectrum of reconstructed squared missing mass for the full data set is compared with the estimated background spectrum, showing qualitatively good agreement across the whole range of missing mass. The fluctuation of the data with respect to the expected background is compared with the combined uncertainty, σ , coming from the estimate of the expected background, σ_{BG} , and the expected fluctuation of the data about this background, $\sqrt{N_{\text{BG}}}$:

$$\sigma = \sqrt{\sigma_{\text{BG}}^2 + N_{\text{BG}}}. \quad (3.86)$$

The distribution of pulls, $(N_{\text{data}} - N_{\text{BG}})/\sigma$, is shown in figure 3.54a, with mean = 0.03 and rms = 0.95. There are no signal windows in which a 5σ excess is observed so the profile likelihood procedure described in section 3.9 is used to set upper limits at 90% confidence level. These are shown in figure 3.54b.

The upper limits on $\mathcal{B}(K^+ \rightarrow \mu^+ \nu_h)$ and $|U_{\mu 4}|^2$ were calculated as described in the previous section and the results are shown in figure 3.55.

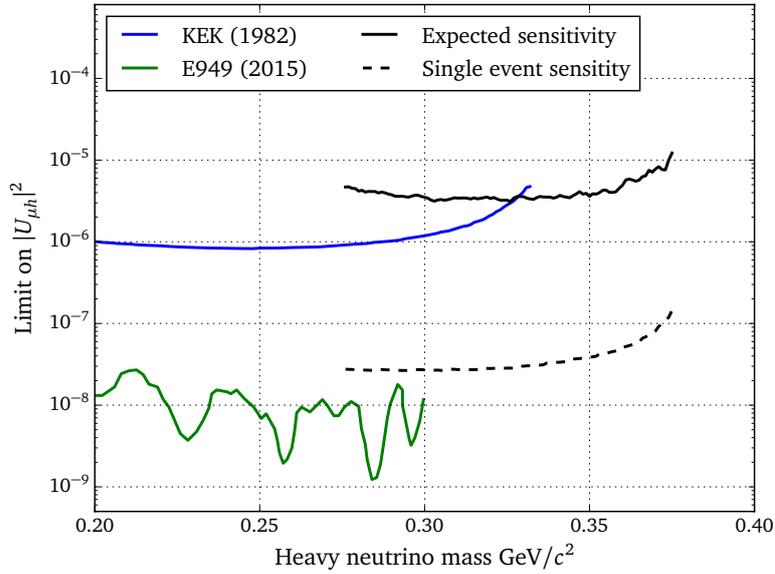


Figure 3.51.: The expected sensitivity, at 90% C.L, to $|U_{\mu 4}|^2$ and the single event sensitivity for this analysis, shown along with the existing limits from production searches in $K \rightarrow \mu \nu_h$ decays at the E949 [70] and the experiment at KEK [59].

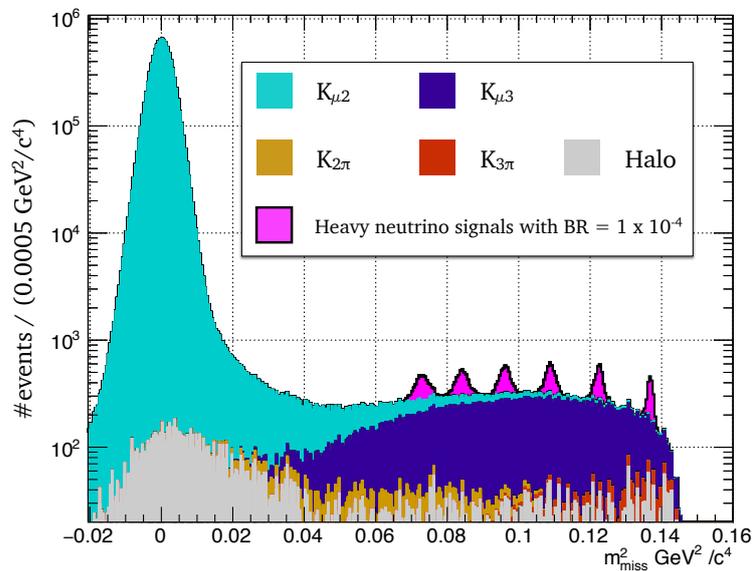


Figure 3.52.: The squared missing mass distribution of the predicted backgrounds, along with examples of the signals that would be expected for $\mathcal{B}(K^+ \rightarrow \mu \nu_h) = 1 \times 10^{-4}$, shown for heavy neutrino masses from $270 \text{ MeV}/c^2$ to $370 \text{ MeV}/c^2$ in $20 \text{ MeV}^2/c^2$ intervals.

3.11. Limits on the branching ratio to heavy neutrinos

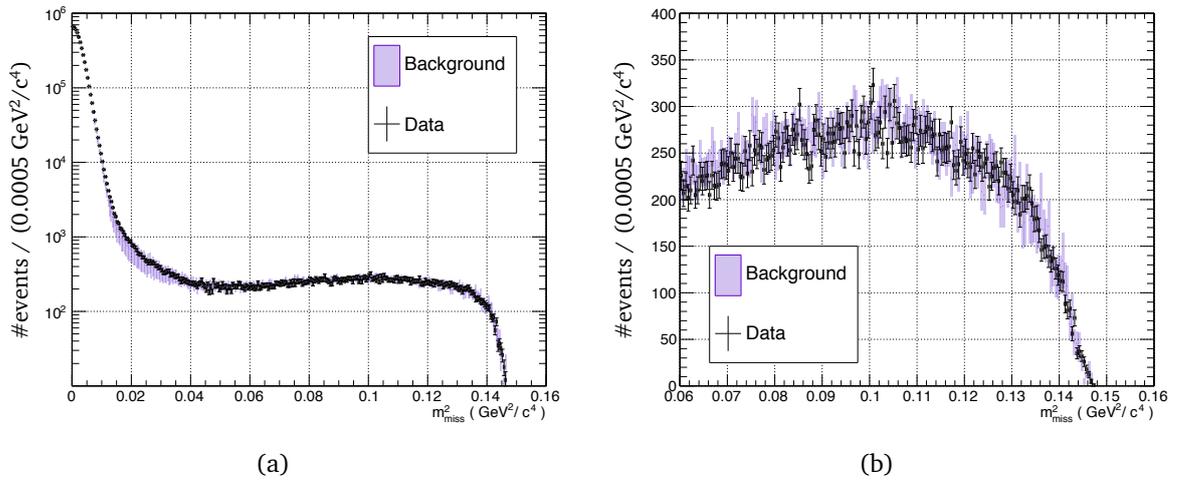


Figure 3.53.: Data and background estimate comparison for the missing mass spectrum (a) on a log-scale and (b) on a linear scale for the signal region. The statistical errors for data are shown by error bars. The width of the filled region indicates the uncertainty on the background prediction.

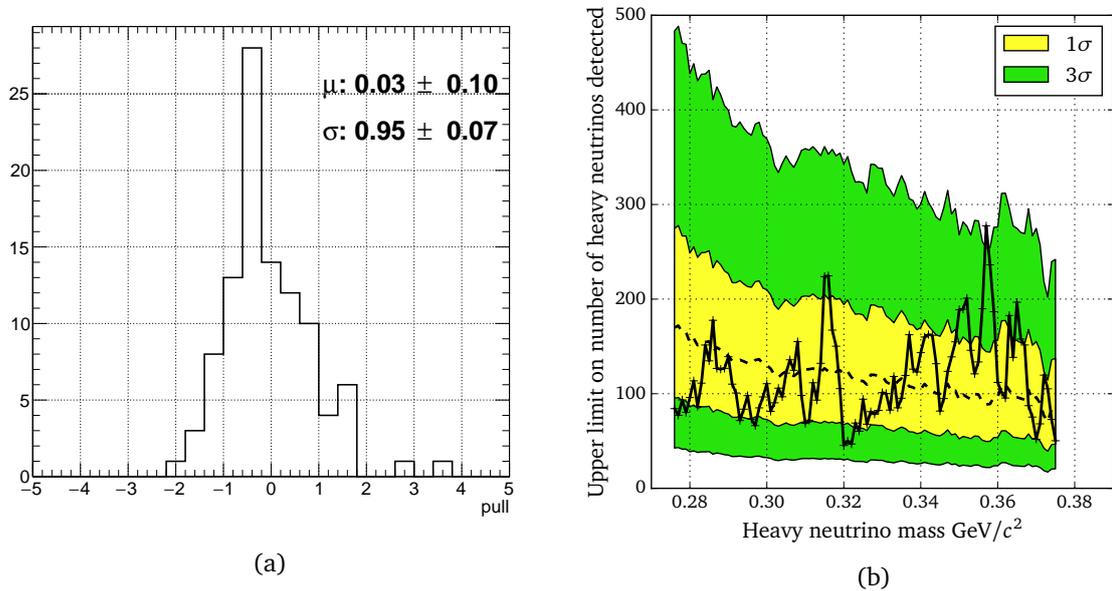
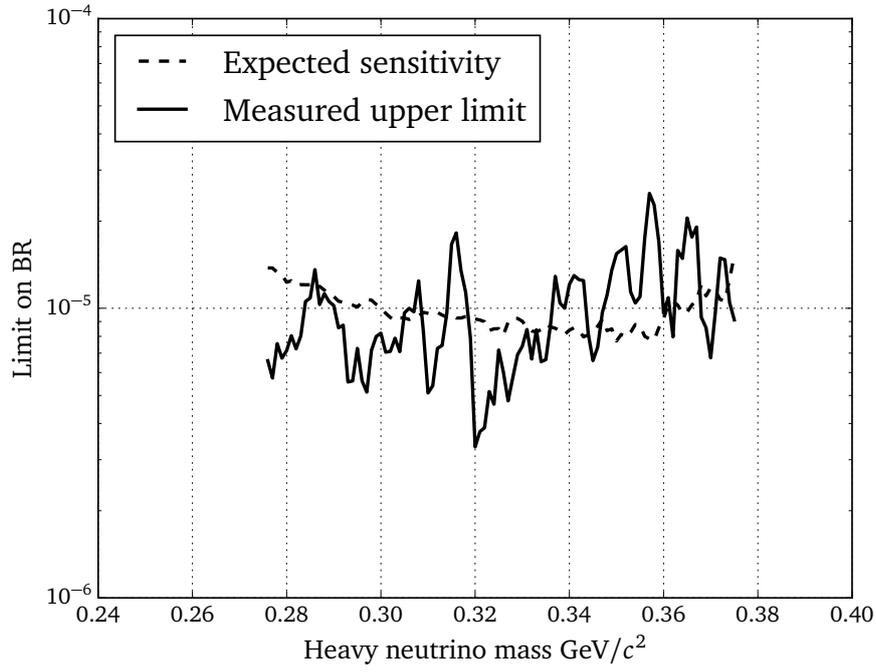
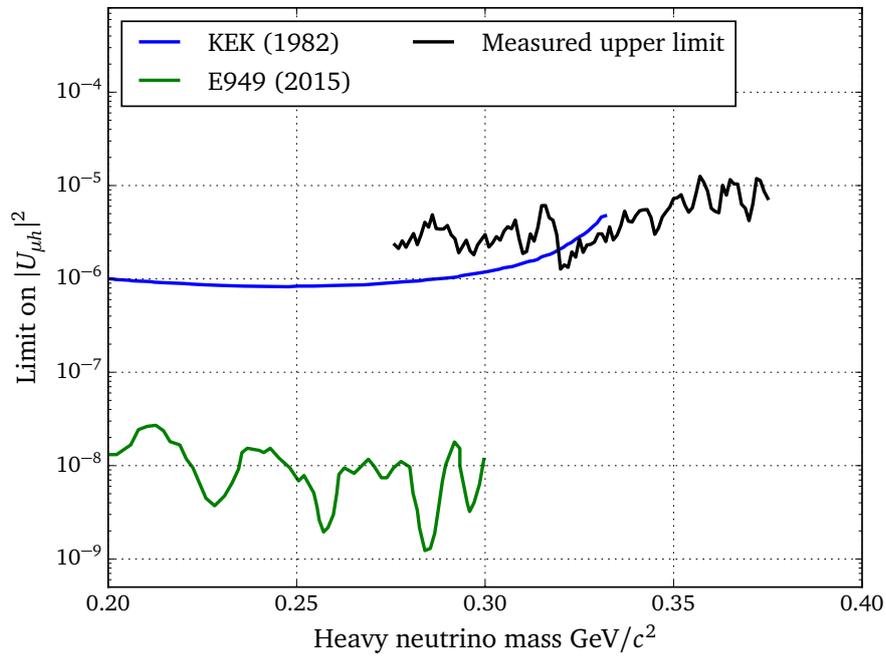


Figure 3.54.: (a) The distributions of pulls for each mass. (b) The upper limit (at 90% C.L.) on the number of heavy neutrinos detected at each mass, together with bands showing the expected 1 σ and 3 σ fluctuation around the expected sensitivity.

3.11. Limits on the branching ratio to heavy neutrinos



(a)



(b)

Figure 3.55.: The upper limits using the 2007 data set from the NA62 experiment, as a function of m_h , the neutrino mass for (a) the branching ratio $K^+ \rightarrow \mu^+ \nu_h$, and (b) the matrix element $|U_{\mu 4}|^2$, shown with the existing limits from peak searches in $K^+ \rightarrow \mu^+ \nu_\mu$ decay at the experiment at KEK [59] and the E949 experiment E949 [70].

3.12. Conclusion

A peak search has been performed on the squared missing mass spectrum already observed in $K^+ \rightarrow \mu^+ \nu_\mu$ decays in part of 2007 dataset from the NA62- R_K experiment. Limits have been set on the mixing between muon and heavy neutrino states, for neutrino masses in the range $275 \text{ MeV}/c^2$ to $375 \text{ MeV}/c^2$. The results extend the range of masses for which upper limits have been set on the value of $|U_{\mu 4}|^2$ from peak searches.

The dominant contribution to the background uncertainty comes from the statistical error on the estimate of the halo distribution, while the largest number of background events comes from the $K^+ \rightarrow \mu^+ \pi^0 \nu_\mu$ decay.

The source of the of the extra scale factor applied to the MC simulation of $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ crosstalk has not been identified. The assumption of charge asymmetry as been used to extra the required scale factor from the data but it would be interesting to find the real source. One explanation, which has been suggested in this note, is that the different paths of the K^+ and K^- beams in the decay volume make the contribution from cross talk strongly dependent on the CDA distribution of $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ events, which is poorly described in MC. Further study could either confirm that this is the cause or determine the true cause of the discrepancy.

The NA62 detector designed for studying the decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ (see following chapters) would be well equipped to perform a similar analysis. The background from the beam halo would be significantly reduced by the KTAG and GigaTracker detectors which can be used to ensure that muons originate from kaon decays in the decay volume, and not from further upstream. In addition the hermetic photon veto provided by the Large Angle Vetoes, LKr calorimeter, Inner Ring Calorimeter and Small Angle Calorimeter would significantly reduce the contribution from $K^+ \rightarrow \mu^+ \pi^0 \nu_\mu$ decays. Finally, the improved momentum resolution from the STRAW chambers means that the signal regions in the squared missing mass distribution could be made narrower, decreasing the number of expected background events without reducing signal efficiency. As well as reduced background, a much larger kaon flux will be available. Assuming a similar downscaling of the minimum bias trigger at NA62 as in 2007, the kaon sample available will be $\sim 10^3$ times larger than the one used for the analysis in this note.

Chapter 4

$K^+ \rightarrow \pi^+ \nu \bar{\nu}$ at NA62

The NA62 experiment was designed to measure the ultra-rare decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ described in section 1. Here, the current experimental status is described before outlining the NA62 measurement strategy and detector design. In the subsequent chapters, I focus on the KTAG detector and my contributions to its development.

4.1. Experimental requirements

The design of experiments to measure $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ is driven by the signature of the decay, as well as the associated background distributions. The signal signature is a single π^+ associated with a K^+ , and no other particles detected (since neutrinos will escape any detector, unimpeded). Since there is freedom in the distribution of energy between the daughter particles, the momentum of the π^+ in the kaon rest frame can take a range of values up to the kinematic limit of 227 MeV/c. Most backgrounds to the signal decay come from other kaon decays with similar decay signatures when one or more of the decay products is misidentified or not detected at all. In particular, the decay $K^+ \rightarrow \pi^+ \pi^0$ mimics the signal mode when the two photons from π^0 decay are lost. Since $K^+ \rightarrow \pi^+ \pi^0$ is a two body decay, the π^+ momentum has a sharp peak at 205 MeV/c. This naturally divides the kinematic distribution of $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ into two regions: region 1, where the π^+ momentum is above 205 MeV/c and region 2, below it.

Decay kinematics can also be expressed in terms of squared missing mass, defined by

$$m_{\text{miss}}^2 = (p_K - p_\pi)^2 \quad (4.1)$$

where p_K and p_{π^+} are the particle four momenta. In this case, the reconstructed missing mass for $K^+ \rightarrow \pi^+ \pi^0$ is that of the π^0 , $0.021 \text{ GeV}^2/c^4$. In region 1, m_{miss}^2 is below $m_{\pi^0}^2$ and in region 2, above it.

Many backgrounds to $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ can be rejected purely kinematically. For example, the charged track associated with the muon from $K^+ \rightarrow \mu^+ \nu_\mu$ has a momentum of $232 \text{ MeV}/c$, which is above the kinematic end point of the signal decay. Experiments should therefore have good momentum resolution in order to ensure kinematic rejection of backgrounds. For backgrounds which cannot be rejected kinematically, experiments must rely on particle identification, to positively select events containing pions rather than muons or electrons; and on photon vetoing, to reject decays such as $K^+ \rightarrow \mu^+ \pi^0 \nu_\mu$.

4.2. Previous experiments

The earliest experimental results for $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ were obtained at the Zero-Gradient Synchrotron, Argonne in 1969, using stopping K^+ mesons in a heavy-liquid bubble chamber [103]. The technique relied on pion identification through the $\pi \rightarrow \mu \rightarrow e$ decay chain, photon vetoing by looking for gamma-ray conversion into an electron-positron pair, and kinematic selection to reject events from the two body decay $K^+ \rightarrow \pi^+ \pi^0$. An upper limit on the branching ratio $\text{BR}(K^+ \rightarrow \pi^+ \nu \bar{\nu})$ of 5.7×10^{-5} was set at 90% confidence level [104].

This result was superseded by an experiment at the Berkely Bevatron, which tracked the decay products of stopped kaons in a stack of spark chambers. The original experiment was sensitive to pions with momenta near the kinematic limit of $227 \text{ MeV}/c^2$ (region 1) and set an upper limit of 1.4×10^{-6} [105]. Subsequently, the addition of a nearly hermetic lead-glass photon detector allowed the study of pions with energies below that for $K^+ \rightarrow \pi^+ \pi^0$ (viz. 205 MeV). By extending the search region to $141 < p_{\pi^+} < 200 \text{ MeV}/c^2$ (region 2), the upper limit was improved to 5.6×10^{-7} [106].

In 1981, an experiment at the KEK proton synchrotron improved the upper limit using the high energy region $214 < p_{\pi^+} < 227 \text{ MeV}$ [107]. Particle trajectories were measured in multi-wire proportional chambers and lead glass Cherenkov detectors were used to veto photons from $K^+ \rightarrow \pi^+ \pi^0$. The analogue signals of the $\pi \rightarrow \mu \rightarrow e$ chain in plastic scintillator were digitized in order to identify pions. The final result was an upper limit on the branching ratio of 1.4×10^{-7} .

The current experimental value for the branching ratio of $K^+ \rightarrow \pi^+ \nu \bar{\nu}$, comes from the E787 [108] and E949 [109] experiments. These were a series of stopped kaon experiments at the Brookhaven

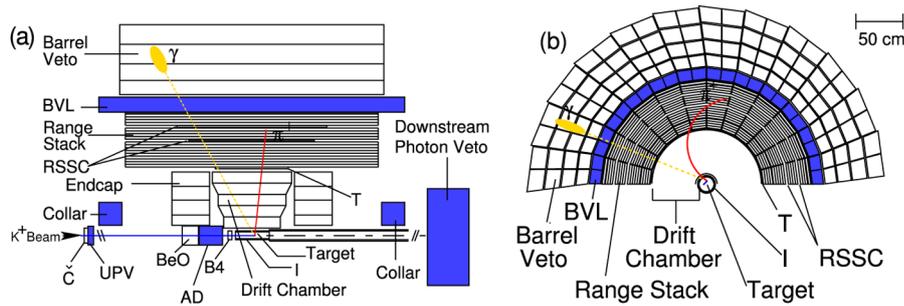


Figure 4.1.: The E949 detector at BNL [110]

National Laboratory (BNL), beginning in the early 1980s. The E949 experiment is shown in its final form in figure 4.1.

For the E787 experiment, an 800 MeV/c kaon beam (also containing pions and protons) was extracted from Alternating Gradient Synchrotron (AGS), slowed in a BeO degrader and stopped in an active target made from scintillating fibres. Kaons were tagged in time by scintillating counters and positively identified by Cherenkov counters. Wire proportional chambers were used to monitor the beam profile and identify multiple incoming particles.

The entire detector was contained and supported by a solenoid magnet producing the 1 T axial field required for momentum determination from charged tracks. A momentum resolution of 2.5% was achieved with a cylindrical drift chamber surrounding the target. It occupied the radial region between 95 and 432 mm, subtending a solid angle of 2π sr, and was filled with a mixture of argon and ethane at atmospheric pressure. The resolution was limited by multiple coulomb scattering to between 130 and 250 μm in the azimuthal direction and between 2.2 and 4.2 mm in the axial direction [111].

Surrounding the drift chamber, a range stack was used to measure the energy, range and decay sequence of charged particles. It consisted of 20 layers of plastic scintillation counters, interspersed with two multi-wire proportional chambers for tracking information. At the trigger level, the deepest scintillator layer reached by a particle could be used to reject muons from $K^+ \rightarrow \mu^+ \nu_\mu$ because their expected range was greater than that of pions from the signal decay. Offline, the digitized signals from the scintillating layers could be used to positively identify a pion by its decay into a mono-energetic muon, which subsequently decays into an electron. The final muon suppression was at the level of about 5.5×10^{-6} .

The outer layer of the detector was a photon veto, key to suppressing backgrounds from $K^+ \rightarrow \pi^+ \pi^0$

and other decays with photons in the final state. The photon veto was split into the barrel assembly, surrounding the range stack, and two end caps, upstream and downstream of the drift chamber. Photons were detected by their electromagnetic showers in alternating layers of lead and scintillator. Detection inefficiencies for single photons varied between 10^{-2} and 10^{-4} for photon energies between 20 and 225 MeV. Taking into account the correlation between photons and kinematic selection of the π^+ , the π^0 detection inefficiency was determined to be around 1×10^{-6} .

In the first phase, E787 set a 90% CL. limit on the branching ratio of 2.4×10^{-9} , using data from pions with energies greater than those from $K^+ \rightarrow \pi^+ \pi^0$ [112]. A separate limit of 1.7×10^{-8} was extracted from data with pions below the peak from $K^+ \rightarrow \pi^+ \pi^0$ [113].

Between 1992 and 1994, E787 underwent a significant upgrade[114]. The drift chamber was completely replaced with a new ultra-thin chamber, improving momentum and axial coordinate resolutions by a factor of 2. The lead scintillator end-caps were replaced with CsI crystals improving the timing resolution by a factor 2 and were placed closer to the tracking chamber in order to increase coverage.

A run in 1995 using the new detector yielded the first ever recorded $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ candidate [115], in region 1. With a total of 1.49×10^{12} kaons on target and a total acceptance of 1.6×10^{-3} , the expected background was 0.08 events. The corresponding branching ratio was $(4.2_{-3.5}^{+9.7}) \times 10^{-10}$. A subsequent data set with similar sensitivity resulted in the second ever recorded event, and a new branching ratio measurement of $(1.57_{-0.82}^{+1.75}) \times 10^{-10}$ [116]. Later analysis extended the kinematic region to include region 2, producing one more candidate event which had little impact on the combined estimate of the branching ratio but did improve sensitivity to non-standard scalar and tensor interactions. [117].

In 2001, the experiment was taken over by the E949 collaboration and was upgraded again to increase the rate of kaons on target and improve detector efficiency. The new collaboration obtained a further 1.7×10^{12} stopped K^+ decays, including 4 more signal events. Now the combined measurement of the branching ratio stands at $1.73_{-1.05}^{+1.15} \times 10^{-10}$ [26], consistent with the Standard Model prediction (equation 1.64). The signal events recorded by E787 and E949 are shown in figure 4.2.

4.3. NA62 Analysis Strategy

The NA62 experiment aims to measure the branching ratio $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ at the CERN SPS with 10% precision, requiring $\mathcal{O}(100)$ events, while controlling systematic errors at the percent level [28].

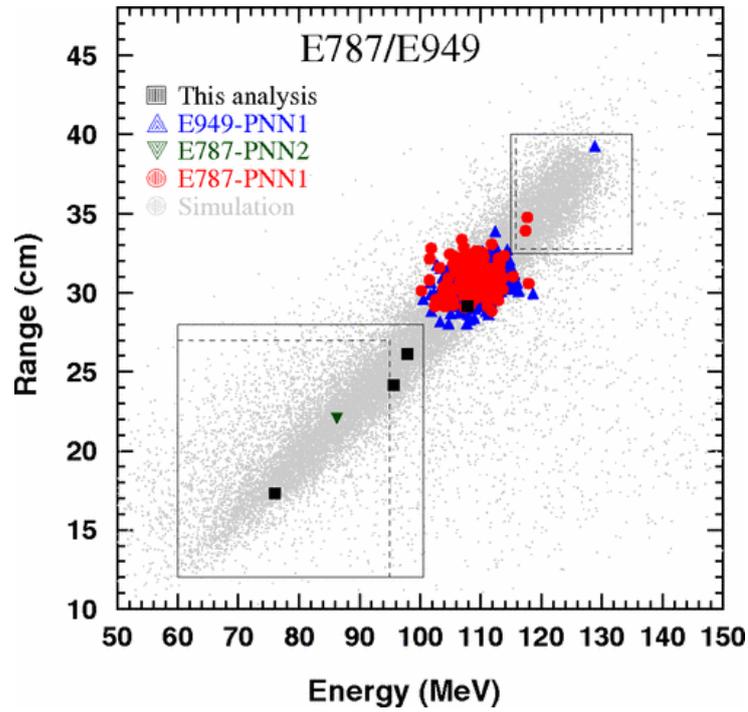


Figure 4.2.: The $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ events recorded by the E787 and E949 collaborations [26]. ‘This analysis’ refers to E949-PNN2 events.

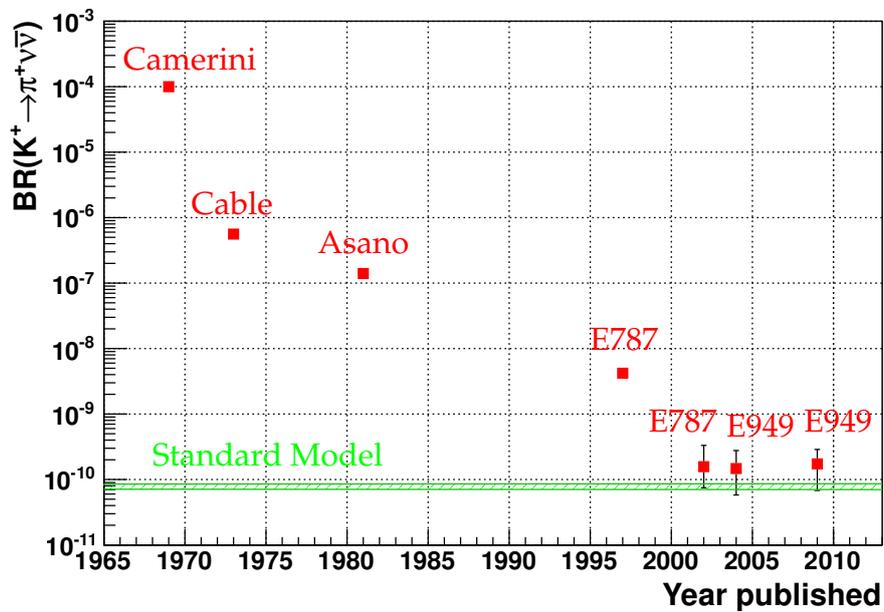


Figure 4.3.: The evolution of limits on $BR(K^+ \rightarrow \pi^+ \nu \bar{\nu})$

The approach is to study the decay of kaons in flight (the first time this approach has been used for $K^+ \rightarrow \pi^+ \nu \bar{\nu}$) using an unseparated 75 GeV/c beam of kaons, protons and pions. Large statistics are achieved by using a high intensity beam and maximising signal acceptance. Systematic errors are controlled with large background rejection and high redundancy between subdetectors. The final signal-background ratio expected is $S/B \approx 10$.

Table 4.1.: The most common K^+ decay modes along with the strategies for suppressing them at NA62. [118]

Decay Channel	Branching ratio (%)	Suppression Strategy
$K^+ \rightarrow \mu^+ \nu$	63.55 ± 0.11	μ veto + 2-body kinematics
$K^+ \rightarrow \pi^+ \pi^0$	20.66 ± 0.08	Photon veto + 2-body kinematics
$K^+ \rightarrow \pi^+ \pi^+ \pi^-$	5.59 ± 0.04	Charged particle veto + kinematics
$K^+ \rightarrow \pi^0 e^+ \nu$	5.07 ± 0.04	E/p + photon veto
$K^+ \rightarrow \pi^0 \mu^+ \nu$	3.353 ± 0.034	μ veto + photon veto
$K^+ \rightarrow \pi^+ \pi^0 \pi^0$	1.761 ± 0.022	Photon veto + kinematics

The most common kaon decay modes are shown in table 4.1, along with the strategies used to suppress them. As in previous experiments, background suppression requires high resolution particle tracking and momentum measurement combined with particle identification and vetoing. Another background associated with the decay in flight technique comes from beam induced accidentals, since a scattered pion from the beam has the same experimental signature as the signal decay. Accidentals of this nature can be suppressed by positively identifying (tagging) the kaon before it decays. Precise timing resolution is required to match the tagged kaon to its decay products downstream.

Figure 4.4 shows the distribution of squared missing mass m_{miss}^2 for the main backgrounds and the signal. As already discussed, the continuous signal distribution is split by the peak from $K^+ \rightarrow \pi^+ \pi^0$. Then, at higher values of m_{miss}^2 , there is a sharp turn on for the decay $K^+ \rightarrow \pi^+ \pi^+ \pi^-$, while at negative missing mass, there is a large contribution from $K^+ \rightarrow \mu^+ \nu_\mu$, when the muon is reconstructed under the assumption that it is a pion. The two m_{miss}^2 signal regions for NA62 are therefore:

Region I: between 0 and the $K^+ \rightarrow \pi^+ \pi^0$ peak

Region II: between the $K^+ \rightarrow \pi^+ \pi^0$ peak and the $K^+ \rightarrow \pi^+ \pi^+ \pi^-$ threshold $\approx 4m_{\pi^\pm}^2$

The decays which can be kinematically constrained in this way constitute 92% of the kaon decay width. The remaining unconstrained backgrounds (and tails of the constrained backgrounds from resolution effects and radiative decays) must be suppressed by using photon and muon vetos to detect

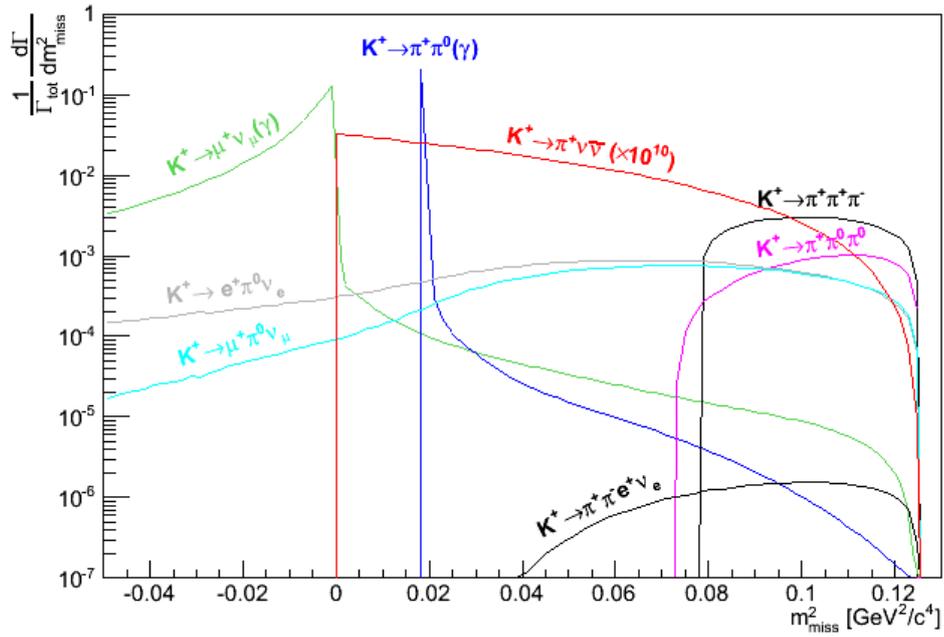


Figure 4.4.: Distribution of m_{miss}^2 for the main background channels. The signal distribution is multiplied by 10^{10} . [118]

extra particles in the final state and using particle identification to distinguish between pions and other particles (i.e. muons and electrons).

The main background in the signal kinematic regions comes from $K^+ \rightarrow \pi^+ \pi^0$ decays with an additional radiated photon. It is important to efficiently detect the π^0 in order to reject these events. Here the decay in flight technique has an advantage with respect to techniques using stopped kaons since the initial kaon momentum of 75 GeV/c means that the decay products are boosted in the lab frame so they can be detected more efficiently. In particular, requiring $|\vec{p}_{\pi^+}| < 35$ GeV/c ensures that the photons from the associated π^0 in $K^+ \rightarrow \pi^+ \pi^0$ have a total energy of at least 40 GeV.

4.4. The NA62 Detector

Figure 4.5 shows the layout of the NA62 detector, which was placed on the same CERN-SPS extraction line used by NA62- R_K in 2007 for the measurement of R_K [86]. Some of the detector is reused, notably the liquid krypton electro-magnetic calorimeter but there are also significant upgrades and most of the sub-detectors are completely new.

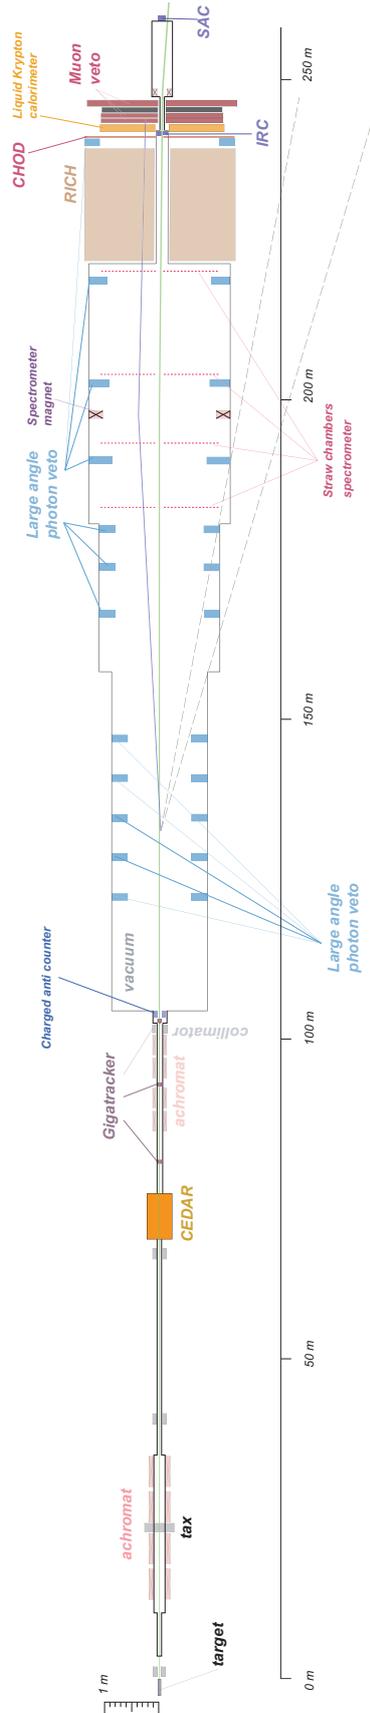


Figure 4.5.: The NA62 Experimental Setup [118]

4.4.1. Beam

The NA62 experiment uses an unseparated beam, obtained by impinging 400 GeV/c protons from the SPS onto a beryllium target (T10). The secondary beam is selected to have a momentum of 75 GeV/c, with 1% standard deviation. It is advantageous to use a K^+ rather than K^- beam since with the chosen momentum, the ratio K^+/K^- per 400 GeV/c proton is ≈ 2.1 , while the fraction of the kaons is the same in the positive and negative beams [119]. The exact choice of kaon momentum was motivated by various factors. A momentum of 75 GeV/c is high enough to ensure that if a particle (other than a neutrino) is produced in association with an accepted π^+ , it must have an energy greater than 40 GeV so it can be easily detected. Conversely, a much higher momentum would reduce the number of kaon decays in the fiducial region.

After the T10 target, the beam passes through a series of quadrupole magnets which define the acceptance angles in the horizontal and vertical planes as well as focussing the beam at the collimators, TAX1 and TAX2. The collimators are blocks made from copper and iron which are water cooled to allow them to absorb the remaining primary proton beam. They are situated in the middle of the ‘front end achromat’, a series of four bending magnets which ensures that only particles with the chosen momentum are able to pass through the holes in the collimators. After the front end achromat, the beam has narrow momentum bite of $\Delta p/p = 1.0\%$ (r.m.s.) about the central momentum of 75 GeV/c.

Another series of quadrupoles refocuses the beam before it passes through three muon sweeping magnets. These dipole magnets, surrounding the beam, are designed to sweep aside the halo of muons which come from the decays of particles in the beam. The beam passes through the almost field-free, 40 mm diameter, bore and any deviation due to stray fields is corrected by steering (TRIM) magnets before and after the sweeping magnets.

The beam then passes through a pair of quadrupoles which ensure that the beam is parallel before it passes through the CEDAR detector. The CEDAR detector is tuned to tag the K^+ component of the beam and is described in chapter 5.

In the final section before the decay volume, the beam momentum is measured by the GigaTracker (GTK). The three GigaTracker stations are placed before, after and in the middle of a ‘second achromat’ composed of four vertically-deflecting dipole magnets. Just before the second GigaTracker station, a ‘scraper’ magnet deflects low momentum muons away as well as blocking any remaining negative or neutral particles in the beam. After the final GigaTracker station, the CHANTI detector surrounds the

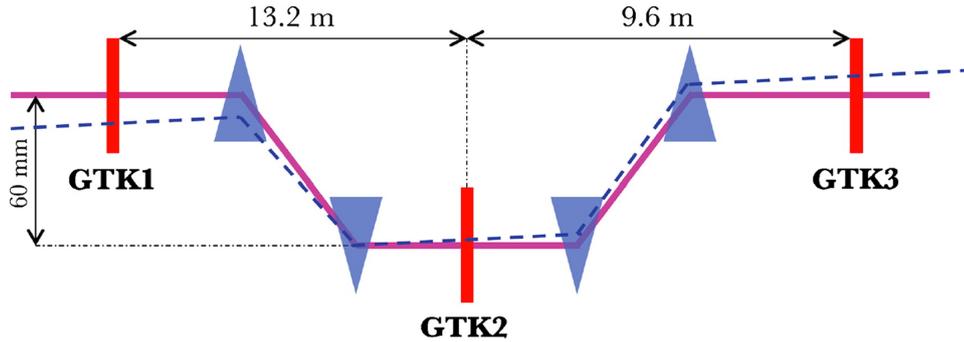


Figure 4.6.: Schematic of the NA62 beam line close to the beam spectrometer, showing the nominal kaon path (purple), and the beam envelope (blue) [118].

beam in order to veto interactions in the exit layer of GTK 3.

Finally, before GTK3, a steering magnet deflects the beam horizontally away from the straight axis by 1.2 mrad, in preparation for the deflection in the opposite direction from the spectrometer magnet, MNP-33.

4.4.2. Tracking and momentum measurement systems

It is vital to have accurate measurements of both the kaon and pion momenta in order to use the reconstructed missing mass as a selection variable. In order to suppress the tails of backgrounds on the edge of the allowed kinematic region, the resolution on m_{miss}^2 should be less than $10^{-3} \text{ GeV}^2/c^4$ [118].

4.4.2.1. The GigaTracker (GTK) spectrometer

The kaon spectrometer, called the GigaTrackEr (GTK) provides measurements of the momentum, time and direction of the $75 \text{ GeV}/c$ beam [120]. Figure 4.6 shows the layout of the three GTK stations within the second achromat. Simulation has shown that in order to achieve the required missing mass resolution, the GTK must measure the kaon momentum with a relative resolution of $\sigma(p_K)/p_K \sim 0.2\%$ and the direction with a resolution of around $16 \mu\text{rad}$. Furthermore, it must achieve these resolutions under strict conditions. Each tracking station is subject to a high and non-uniform beam rate: 750 MHz spread over the 1620 mm^2 area with a peak value of $1.3 \text{ MHz}/\text{mm}^2$. In order to avoid mismatching between kaons and their decay products in such a high rate environment, the time resolution for a reconstructed kaon must be better than 150 ps . The total amount of material from all three stations must constitute less than 5×10^{-3} radiation lengths (X_0) per station in order to prevent scattering

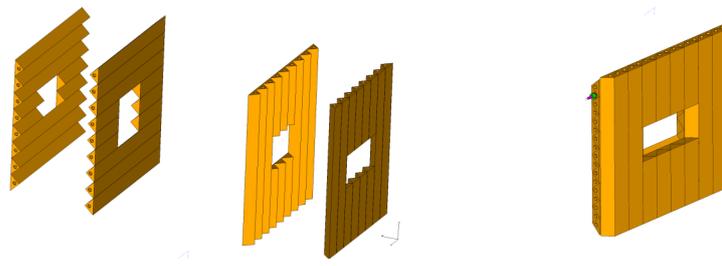


Figure 4.7.: A CHANTI station in exploded and assembled views.

particles which might interact with other sub-detectors downstream.

To meet these requirements, each station is comprised of a hybrid silicon detector containing 18,000 pixels ($300\ \mu\text{m} \times 300\ \mu\text{m}$) in a 90×200 grid which matches the expected beam dimensions of about $60\ \text{mm} \times 27\ \text{mm}$. The pixels are $200\ \mu\text{m}$ thick. The time resolution for each pixel hit is about 200 ps. In order to cool the chip with minimal additional material, a micro-cooling system is used [121], in which liquid C_6F_{14} flows through micro-channels $70\ \mu\text{m}$ deep and $200\ \mu\text{m}$ wide embedded in a silicon membrane $150\ \mu\text{m}$ deep.

4.4.2.2. The Charged ANTI (CHANTI) detector

Despite the small material budget of the GTK, there is a potential background due to inelastic interactions of the beam with the final GTK station. If a pion is emitted at low angle, it can reach the STRAW tracker and mimic a signal event. GEANT4 simulation has shown that the probability of scattering in GTK is 1×10^{-3} so the combination of analysis cuts and inelastic collision veto must reject these events with an inefficiency less than 10^{-8} . The Charged ANTI (CHANTI) detector [122] is designed to veto inelastic interactions in the GTK by detecting particles emitted between 34 mrad and 1.38 rad with respect to the beam axis. It consists of 6 stations, placed downstream of the GTK, each consisting of two layers of scintillating bars placed at 90° to each other. An example is shown in figure 4.7. Light is collected by a wavelength shifting fibre placed inside each bar and read out by a silicon photomultiplier. The expected rate of events vetoed by the CHANTI due to inelastic interactions as well as muons from the beam halo is 2 MHz. In order not to introduce an unacceptable inefficiency through random vetoing, the CHANTI must have time resolution $< 2\ \text{ns}$. For signal-like events, the CHANTI vetoes events with inelastic collisions with almost 99% efficiency, and prototype measurements have shown a time resolution better than 1 ns.

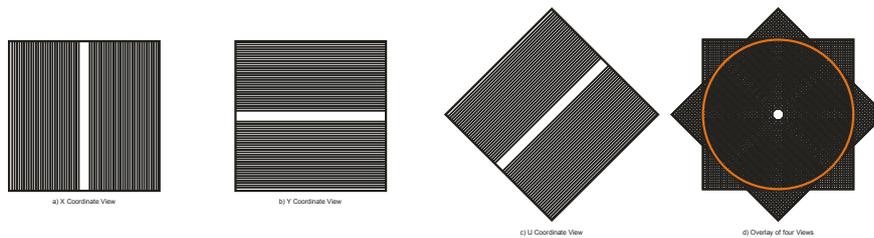


Figure 4.8.: The four views of a straw chamber

4.4.2.3. The Straw Tracker

Kaon decay product momenta are measured by a magnetic spectrometer, the Straw Tracker [123], which consists of four chambers in vacuum, two upstream and two downstream of a dipole magnet producing a vertical B-field of 0.36 T and corresponding transverse momentum kick $\Delta p_T = 270 \text{ MeV}/c$. Secondary particles must be reconstructed with momentum resolution $\sigma(p)/p \leq 1\%$ and $\sigma(\theta_{K\pi}) \leq 60 \mu\text{rad}$ with minimal Coulomb scattering, particularly in the first chamber. These requirements imply a spatial resolution $\leq 130 \mu\text{m}$ per coordinate and material budget of $\leq 0.5\%$ of a radiation length for each chamber. The detector must operate with an average particle rate of 40 MHz, and up to 500 kHz per straw for straws close to the beam. In order to minimize multiple scattering, the spectrometer is operated in vacuum with no physical separation from the upstream decay volume.

Each chamber consists of two modules, each containing two views, x,y or u,v, (shown in figure 4.8), so there are four coordinates measured for every charged particle passing through a chamber. There are 448 straws per view, arranged in four staggered layers to resolve the left-right ambiguity and ensure that a charged particle must cross at least two straws. A central strip is left empty in each view, creating an octagonal hole in the chamber for the beam to pass through. The straws are made from $36 \mu\text{m}$ thick Mylar, coated on the inside with a layer of copper and a layer of gold to create an electrically conducting cathode. The volume inside the straws is filled with argon and CO_2 at atmospheric pressure, and the anode sense wires in the centre are made from gold-plated tungsten. In order to prevent the straws sagging under the force of gravity, they are held under tension and supported by spacer rings to hold them in position.

The expected performance is $\sigma(|\vec{p}|)/|\vec{p}| \approx 0.3\% \pm 0.007\%$ and $\sigma(dX,Y/dZ) \approx 15 - 45 \mu\text{rad}$, depending on the track momentum.

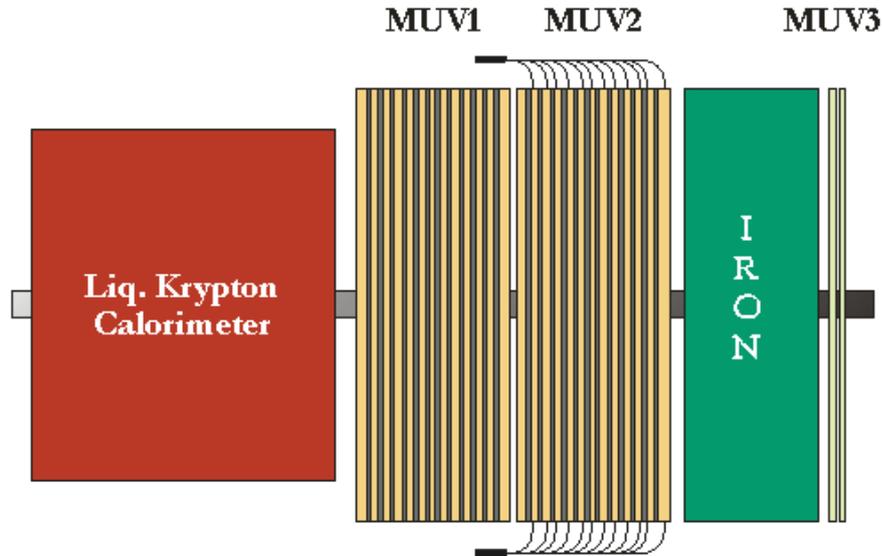


Figure 4.9.: The three MUV subdetectors following the Liquid Krypton Calorimeter [118].

4.4.3. Particle identification

Particle identification (PID) is important both for beam particles and their decay products. Kaon tagging in the beam is presented in subsequent chapters, while identification of decay products is discussed here. The main purpose of identifying decay products is to reject the background $K^+ \rightarrow \mu^+ \nu_\mu$ by distinguishing pions and muons. The branching fraction of $K^+ \rightarrow \mu^+ \nu_\mu$ is of order 1 so a suppression factor of 10^{-12} is required in order to reach the required sensitivity. The kinematic cut on m_{miss}^2 provides a suppression factor of 10^{-5} so the Muon Veto system and RICH detector must between them provide another factor of 10^{-7} . Since the expected rate of muons is so high (10 MHz), a muon veto is needed at the lowest trigger level to reduce the rate below the maximum which can be handled by the readout system.

4.4.3.1. Muon Veto System

The Muon Veto System consists of three subdetectors shown in figure 4.9 [118]. MUV1 and MUV2 are classic iron scintillator sandwich calorimeters consisting of alternating layers of iron and plastic scintillator. They are placed immediately downstream of the LKr calorimeter and work as hadronic calorimeters, measuring the deposited energies and shower shapes of incident particles. MUV1 is a new detector. It has 24 layers of iron and scintillator strips (48 strips per layer), with orientation alternating between vertical and horizontal. The scintillators are read out via wavelength shifting

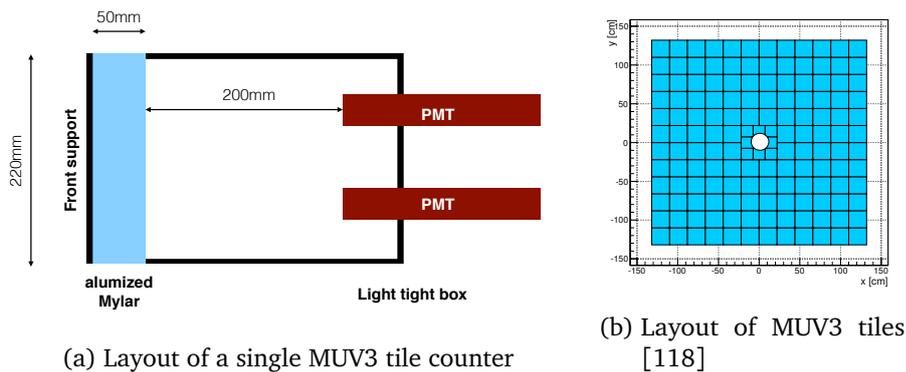


Figure 4.10.: Schematic design of MUV3 tiles

fibres (two per strip), with each longitudinal row of scintillators (i.e one from each layer) read out by a single PMT.

The MUV2 is a reused part of the NA48 hadronic calorimeter. Again the scintillators are arranged in layers of alternating orientation but in MUV2, each strip spans only half the calorimeter so each plane is made of two half-planes. Scintillator strips are connected to PMTs via plexiglass lightguides.

MUV3, located behind an 80 cm thick iron wall, is used as a muon veto at the lowest trigger level. The scintillator is divided into 148 tiles, arranged as shown in figure 4.10. To ensure a good timing resolution, no wavelength shifting fibres are used; light travels directly from each scintillator tile to a pair of PMTs. There is a possible error in timing measurement if a particle passing through a PMT window emits Cherenkov photons, which would arrive 2 ns earlier than those emitted by the scintillator. This source of error is suppressed by requiring a coincidence of the two PMTs in each tile to define the time of the particle passing through. In the NA62 Technical Run, the measured time resolution of MUV3 was $\sigma_t \approx 500$ ps. The total rejection factor of the MUV system for $K^+ \rightarrow \mu^+ \nu_\mu$ decays is expected to be 10^{-5} .

4.4.3.2. Ring Imaging Cherenkov Detector (RICH)

The remaining factor 10^{-2} of muon suppression is provided by the RICH [124]. This is a Ring Imaging Cherenkov designed to distinguish between pions and muons as well as helping to suppress decays containing electrons such as $K^+ \rightarrow e^+ \pi^0 \nu_e$. According to the NA62 analysis strategy, the RICH must be able to distinguish muons and pions with momenta between 15 and 35 GeV/c. A time resolution of better than 100 ps is required in order to match the identified pion with the kaon upstream.

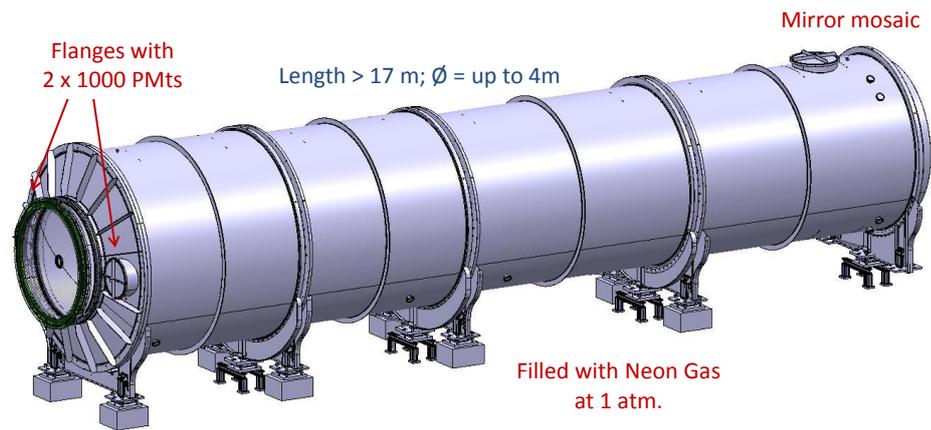


Figure 4.11.: The RICH subdetector [118]

Neon is a suitable choice of radiator gas as, at standard temperature and pressure, the Cherenkov threshold for pions is 12.5 GeV/c so the detector is efficient for momenta above 15 GeV/c. The low index of refraction corresponds to a smaller number of Cherenkov photons per kaon (see section 5.1.2) so a long (17 m) radiator is required in order to produce enough radiation to be detected. A mosaic of 18 hexagonal mirrors at the downstream end of the detector reflects Cherenkov light back towards the upstream end of the detector, where it is focused onto two flanges, each housing 1000 PMTs (see figure 4.11). The PMTs have a maximum separation of 18 mm, chosen to optimise the angular resolution of the detector. Within the momentum range 15 to 35 GeV/c, test beam results have shown that π^+/μ^+ separation can be achieved with the required purity [125]. The expected time resolution is below 100 ps.

4.4.3.3. The charged hodoscope

The plastic scintillator hodoscope is reused from the NA48 experiment (see section 2.3.4), to provide a charged particle trigger signal or to act as a veto for suppressing background signals.

4.4.4. Photon vetoes

Kaon decays can imitate the signal decay if one or more photons escape the detector undetected. In particular, suppressing $K^+ \rightarrow \pi^+\pi^0$ requires π^0 rejection with an inefficiency $< 10^{-8}$ and this has driven the design of the photon veto system. Hermetic coverage for photons with angles up to 50 mrad is divided into three regions. The Large Angle Vetoes (LAVs) cover the region from 8.5 to

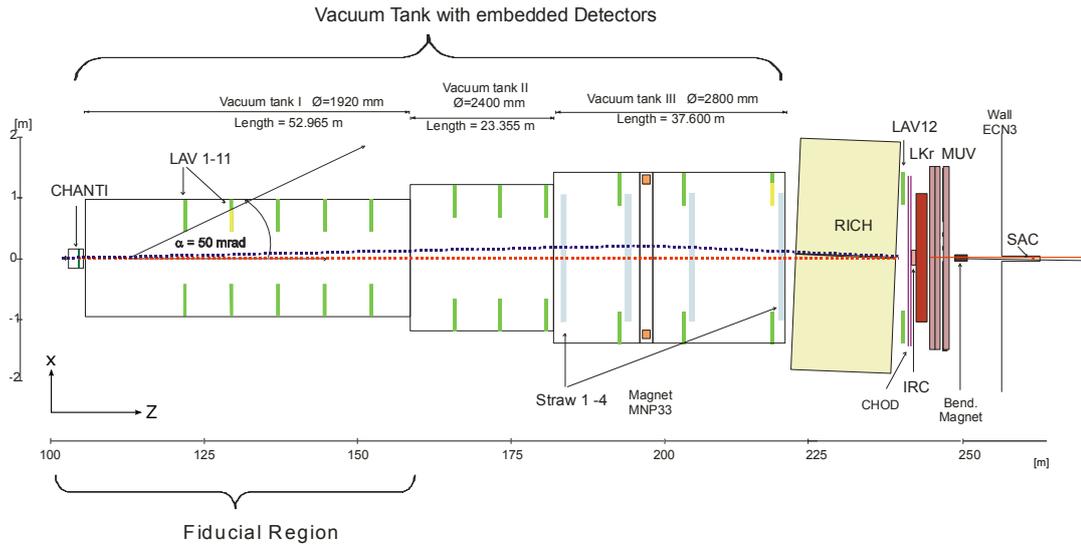


Figure 4.12.: The layout of the downstream NA62 sub-detectors, showing the nominal beam trajectory in red and the beam envelope in blue.

50 mrad, the Liquid Krypton Electromagnetic calorimeter (LKr) covers the region from 1.5 to 8.5 mrad and the Small Angle Vetoes (SAV) cover angles less than 1.5 mrad. Figure 4.12 shows the layout of the downstream detectors including the photon vetoes. Requiring that the π^+ momentum is between 15 and 35 GeV/c ensures that the two photons from the π^0 have a total energy of at least 40 GeV. A photon is lost out of acceptance in only 0.2% of $K^+ \rightarrow \pi^+ \pi^0$ decays and there are no configurations in which both photons are out of acceptance.

4.4.4.1. Large Angle Vetoes

In order to achieve 10^{-8} inefficiency for π^0 detection, the LAVs must detect photons with energies down to 200 MeV with an inefficiency less than 10^{-4} , and with a time resolution around 1 ns in order to match detected photons to the right event [126].

There are 12 LAV stations. The first 11 intersect the vacuum decay tube so they must operate in a vacuum of 10^{-6} mbar, while the last is downstream of the RICH detector, exposed to the atmosphere. They make use of lead-glass blocks, recycled from the OPAL electromagnetic calorimeter barrel [127]. Electromagnetic showers in the lead glass produce Cherenkov light which is detected by PMTs, via a glass (SF47) light guide. In each LAV station, the blocks are arranged around the inside of a segment of vacuum tank in rings facing inwards, with the PMTs on the outside (see figure 4.13). There are 4 or 5 rings in each station, staggered azimuthally to provide complete hermeticity.

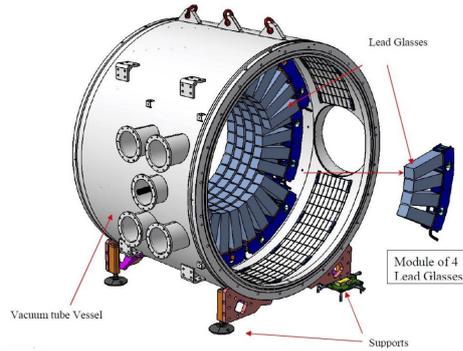


Figure 4.13.: A single LAV station [118]

In 2007, the LAV photon detection inefficiency was measured at the Frascati Beam Test Facility and found to be $(1.2^{+0.9}_{-0.8}) \times 10^{-4}$ at 200 MeV and $(1.1^{+1.8}_{-0.7}) \times 10^{-5}$ at 500 MeV [128].

4.4.4.2. The LKr Calorimeter

The LKr calorimeter is reused from the NA48 experiment as described in chapter 2. However it is equipped with a new readout system, based on the CREAM module [129], in order to cope with high expected rate. Analysis of data taken in 2004 set an upper limit of 0.9×10^{-5} at 90% C.L. for the inefficiency of detecting photons (with energies > 10 GeV) in the LKr [130].

4.4.4.3. Small Angle Vetoes

Photons at small angles (< 1.5 mrad) are detected by the small angle vetoes: the Inner Ring Calorimeter (IRC), and the Small Angle Calorimeter (SAC). Both sub-detectors are based on “Shashlyk” technology: alternating lead and scintillator plates, read out using wavelength shifting fibres and PMTs. In both calorimeters, there are 70 iron-scintillator planes of thickness $1.5 \text{ mm} + 1.5 \text{ mm}$ corresponding to about $16X_0$.

The IRC has a cylindrical shape, segmented into 4 parts longitudinally and 4 parts azimuthally so it is read out by 4 PMTs. It is located just before the LKr calorimeter, surrounding the beam pipe. Although its purpose is to veto photons from $K^+ \rightarrow \pi^+ \pi^0$, the majority of events in the IRC will come from the muon halo.

The SAC is made from a single block of alternating lead and scintillator plates. Light is extracted by 480 wavelength shifting fibres and detected by 4 PMTs. A magnet, 248 m from the target, deflects the kaon beam into an off axis beam dump. The SAC is situated in line with the undeflected beam to

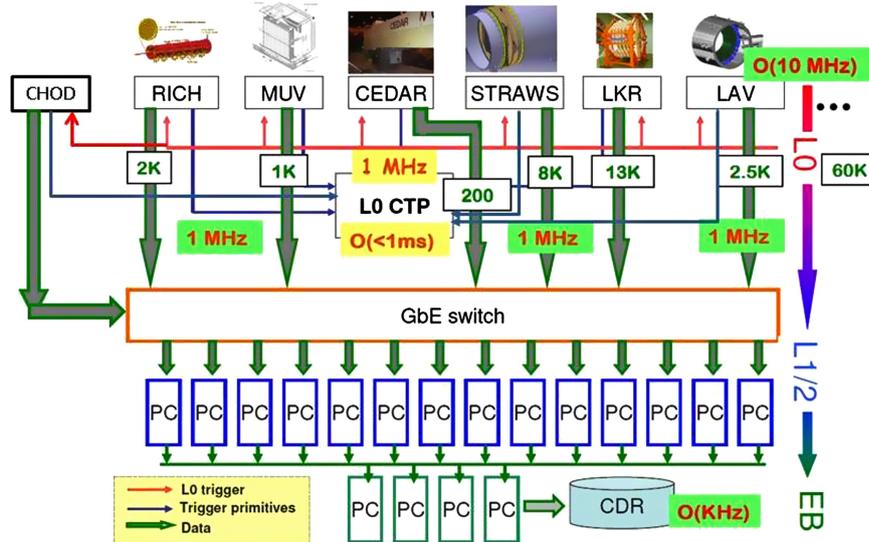


Figure 4.14.: The trigger system system for NA62

detect photons from $K^+ \rightarrow \pi^+ \pi^0$ which dominate the rate of hits it sees.

During a test run in 2006, the inefficiency of the SAC was evaluated using a 25 GeV electron beam. It was found to be $(2.9 \pm 0.3) \times 10^{-5}$, surpassing the 10^{-4} level of inefficiency required by NA62 [131].

4.4.5. Trigger and data acquisition

The expected event rate on the downstream detectors is $\mathcal{O}(10)$ MHz, so a high-performance trigger and data acquisition system (TDAQ) is required in order to achieve a final rate of selected events $\mathcal{O}(\text{kHz})$ which can be written disk, while maintaining good signal efficiency, both for $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ and for other decays of interest. NA62 operates a 3-level trigger, with a hardware L0 trigger and software L1 and L2 triggers (see figure 4.14) [132].

The L0 trigger is designed to reduce the event rate from $\mathcal{O}(10)$ MHz to $\mathcal{O}(1)$ MHz, based on ‘trigger primitives’, encoding a particular set of conditions that have been matched, from a subset of the sub-detectors. The maximum trigger latency at L0 is 1 ms, determined by the size of the buffers where sub-detector data is stored before the trigger decision is available. The L0 trigger processor is implemented on field-programmable gate arrays (FPGAs), which have firmware implementations of the trigger algorithms. A PC based alternative for the LOTP is also under consideration [133].

In response to a L0 trigger signal, data from sub-detectors are sent to the read-out PC farm, where the L1 and L2 trigger decisions are made. The L1 trigger makes decisions based on conditions met

by individual sub-detectors, such as the number of tracks seen in the STRAW detector. If the L1 trigger conditions are met, then the event is reconstructed at L2, combining information from all sub-detectors, to produce the final trigger decision. Events which satisfy the L2 trigger are written to tape.

Chapter 5

KTAG at NA62

A consequence of using a high momentum beam for NA62 is that kaons cannot be efficiently separated from other charged hadrons (mainly pions and protons) in the beam. As a result, when it reaches the upstream detectors, the beam contains only 6% kaons, with the rest made up of pions ($\sim 70\%$) and protons ($\sim 20\%$). The total rate is 750 MHz, 17 times higher than the rate of kaons.

The presence of particles other than kaons in the beam leads to a potential source of background. If a beam pion scatters in residual gas in the vacuum tank, it can pass through the downstream detectors producing exactly the same signature as the signal decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$. In order to keep the number of background events of this nature below 1 per year, the pressure in the first 80 m of the vacuum tank should be below 10^{-6} mbar. However, this requirement can be relaxed by an order of magnitude if kaons are tagged in the beam. This is the role of the KTAG/CEDAR sub-detector.

The detector is based on the Cherenkov Detector with Achromatic Ring focus (CEDAR). This is a gas-filled differential Cherenkov counter, which has been used at CERN since the early 1980s. Using nitrogen gas as a radiator, it is able to separate kaons and pions with momenta up to 150 GeV/c, while still detecting protons with momenta as low as 12 GeV/c.

The specifications of KTAG are driven by the properties of the beam and the background rejection requirements of the NA62 experimental strategy. KTAG must reject pions with a mistagging rate below 10^{-4} , in order to keep the rate of accidentally reconstructed events below 1% of the expected signal, while maintaining a kaon tagging efficiency above 95%. Since the total rate of beam particles is 750 MHz, the kaon crossing time must be measured with a resolution better than 100 ps in order to match kaons with their decay products downstream. The components close to the beam line must be able to withstand a high radiation level while maintaining stable and reliable operations.

In order to meet these requirements, and in particular to work at a high rate, the light detection and read-out systems of the CEDAR have been completely replaced with a new system, KTAG (Kaon TAGger). In this chapter, the CEDAR principle of operation is described. This is followed by a description of the CEDAR/KTAG detector and its readout system. Finally, I describe my work on simulating the detector and its performance.

5.1. Principle of operation

The CEDAR detector exploits the Cherenkov effect. When particles travel through a gas at a velocity greater than the local speed of light, they emit Cherenkov radiation at an angle, θ , which is a function only of the particle velocity, β , and the refractive index of the medium, n :

$$\cos \theta = \frac{1}{n\beta}. \quad (5.1)$$

A Cherenkov detector works by focussing this light to give the image of a ring in its focal plane, so the particles with different velocities can be distinguished by the diameter of their light rings. In a beam with fixed momentum, particle velocity is a function of mass:

$$\beta = \frac{1}{\sqrt{1 + m^2/p^2}}, \quad (5.2)$$

so ring radius can be used to distinguish particles with different masses. Two particles with different masses will have ring radii, R , which differ by:

$$\Delta R = f \Delta \theta = \frac{f}{\theta} \frac{\Delta(m^2)}{p^2}, \quad (5.3)$$

where f is the focal length of the system. The CEDAR detector selects particles of a single type by using a mechanical diaphragm of fixed radius and variable annular aperture, w_A to select rings with a particular radius. In order to distinguish between two particles, we require $\Delta R > w_A$.

Since the radius of the aperture is fixed, the selected particle mass is changed by varying the refractive index of the radiator gas. In terms of electron oscillators, the index of refraction of a gas for a particular frequency of light ω is given by a sum over the resonance frequencies ω_k with associated

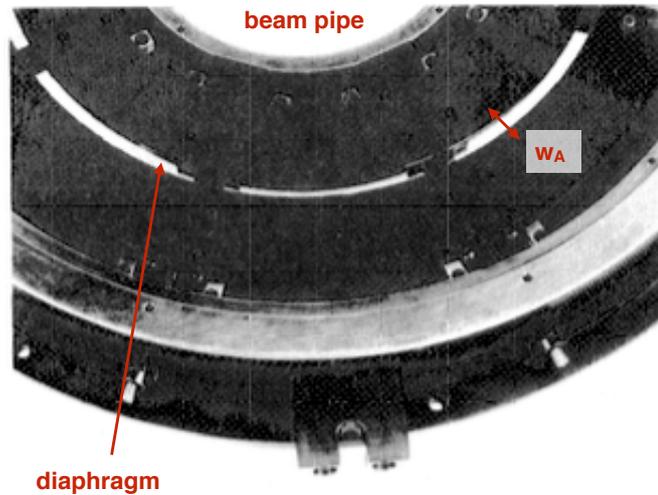


Figure 5.1.: A photograph of the light diaphragm showing two of the eight sectors [134].

damping factors γ_k and weighted by oscillator strengths f_k :

$$n - 1 = \frac{Nq_e^2}{2\epsilon_0 m} \sum_k \frac{f_k}{\omega_k^2 - \omega^2 + i\gamma_k \omega}. \quad (5.4)$$

Here, N is the number of atoms per unit volume, so for a particular gas, $n - 1$ is proportional to gas density. Using the perfect gas equation, we find:

$$n - 1 = (n_0 - 1) \left(\frac{T_0}{P_0} \right) \frac{P}{T}, \quad (5.5)$$

where n_0 is the refractive index at standard temperature and pressure (T_0, P_0), and T and P are the current temperature and pressure.

5.1.1. Ring broadening

The best mass resolution and hence rejection of unwanted particles is achieved in the limit of setting the diaphragm aperture, w_A , to zero. However, the light ring due to a single kind of particle is broadened by a number of mechanisms so closing the light diaphragm too far reduces the efficiency of detecting the desired particle. These effects are visualized in figure 5.2.

The largest effect is *chromatic dispersion*. The variation of refractive index with the wavelength of

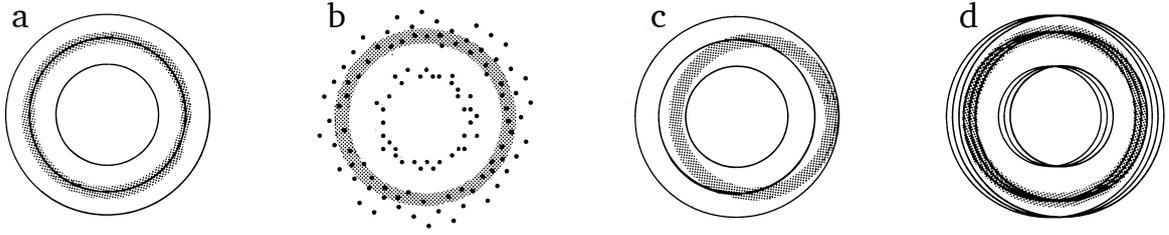


Figure 5.2.: The effect of various broadening mechanisms on the light ring seen by the CEDAR detector. a) Good alignment. b) Broadening due to multiple coulomb scattering. c) Misalignment (see chapter 6). d) Effect of beam divergence

light spreads the ring over a width:

$$\Delta R = f\theta \frac{1}{2\nu} \left(1 + \frac{1}{\theta^2 \gamma^2} \right), \quad (5.6)$$

where ν is the Abbé number [135], which characterizes the dispersion of the gas (a small Abbé number means high dispersion), and the effect depends on the speed of the particle through $\gamma = 1/\sqrt{1-\beta^2}$. This effect can be mitigated by using a chromatic corrector made from two lenses to compensate for the dispersion of radiator gas. When optimized for a particular regime (i.e. separating pions and kaons), the γ dependence can be ignored and the correction is limited only by the difference in curvature of the dispersion functions, $n(\lambda)$, for the radiator gas and lens material.

Once chromatic dispersion has been corrected for, the most important remaining effect is due to *multiple scattering* of the beam as it traverses the gas radiator. If the direction of a particle changes as it passes through the detector, the photons are spread around non-concentric circles, broadening the ring. For a particle with momentum p , the contribution is:

$$\sigma_R = f \frac{13.6 \text{ MeV}}{\beta c p} \sqrt{\frac{LP\rho_0}{3X_0P_0}}, \quad (5.7)$$

where L is the length of the gas radiator which is at pressure P and X_0 and ρ_0 are the radiation length and density of the gas at atmospheric pressure, P_0 . For the NA62 CEDAR, $\sigma_R = 0.06$ mm.

There can also be small contributions from the inhomogeneity of the index of refraction of the radiator so it is important to maintain uniform pressure and temperature throughout the CEDAR vessel.

The largest broadening effect is not intrinsic to the CEDAR detector but comes from the divergence

of the beam. Unlike in the other effects, there is no broadening for an individual particle but only when averaging over many particles. The beam divergence is $70 \mu\text{rad}$ in x and y , so the contribution to ring broadening is:

$$\sigma_R = f \sqrt{\sigma_{\theta_x}^2 + \sigma_{\theta_y}^2} = 0.45 \text{ mm} \quad (5.8)$$

5.1.2. The number of photons per kaon

The number of Cherenkov photons emitted by a charged particle per unit length travelled, dx , and per unit wavelength, $d\lambda$, is given by the Frank-Tamm equation [136]:

$$\frac{dN}{dx d\lambda} = 2\pi\alpha \frac{1}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2} \right), \quad (5.9)$$

where α is the fine structure constant. The emitted energy is strongly peaked at short wavelength, but is saved from ultraviolet divergence because the condition $\beta n > 1$ is typically only satisfied above a particular wavelength:

$$\frac{dN}{dx} = 2\pi\alpha \int_{\beta n > 1} \left(1 - \frac{1}{\beta^2 n^2(\lambda)} \right) \frac{d\lambda}{\lambda^2}. \quad (5.10)$$

The number of photons detected per kaon is important both for detection efficiency and timing resolution.

Kaon tagging efficiency can be modelled by Poisson statistics. Typically, a kaon is tagged by the coincidence of photons detected in a specified number of the 8 sectors. If the mean number of detected photons expected per sector per kaon is ϕ , then the probability of detecting no photons in a particular sector is $e^{-\phi}$ and so the efficiency of a single sector is:

$$\eta_{\text{sector}} = 1 - e^{-\phi}. \quad (5.11)$$

Assuming identical sectors, the probabilities of n -fold coincidences are given by the binomial coefficients:

$$P(n\text{-fold}) = \binom{8}{n} \eta^n (1 - \eta)^{8-n} \quad (5.12)$$

and the efficiency when requiring n -fold coincidence is given by the cumulative sum:

$$\eta_n = \sum_{x=n}^8 P(x\text{-fold}) \quad (5.13)$$

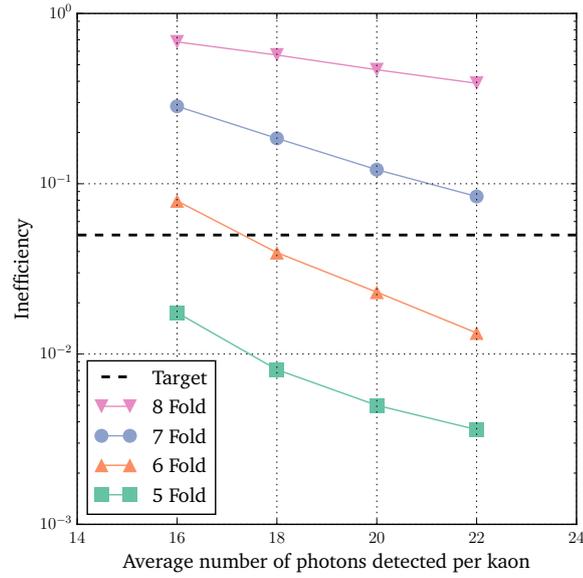


Figure 5.3.: Inefficiency for kaon tagging as a function of the number of photons expected per kaon and the number of sectors required in coincidence (called an n -Fold coincidence). The dashed line shows the target inefficiency for KTAG at 5%.

If we require all 8 sectors in coincidence, then the inefficiency for kaon detection is given by:

$$(1 - \eta)_8 = 1 - \eta^8 = 1 - (1 - e^{-\phi})^8. \quad (5.14)$$

The number of photons expected per kaon at NA62 is about 18, or 2.25 per sector. This would correspond to an inefficiency of about 60% if all 8 sectors were required in coincidence. Figure 5.3 shows that if the coincidence of 6 sectors is required, the inefficiency is less than the 5% specified in the NA62 design.

The relations given by equation 5.12 can be rearranged to determine the number of photons per kaon, from ratios of n -fold coincidences. In particular, one can derive two expressions for ϕ (the number of photons per kaon) in terms of the 6, 7 and 8-fold efficiencies, ϕ_{78} which depends only on η_7 and η_8 and ϕ_{68} which depends only on η_6 and η_8 :

$$\phi_{78} = \ln \left(1 + \frac{8}{\eta_7/\eta_8 - 1} \right) \quad (5.15)$$

$$\phi_{68} = \ln \left(1 + \frac{14}{\sqrt{4 - 7(1 - \eta_6/\eta_8)} - 2} \right) \quad (5.16)$$

These expressions have the advantage that only the ratios of efficiencies enter, so they can be evaluated

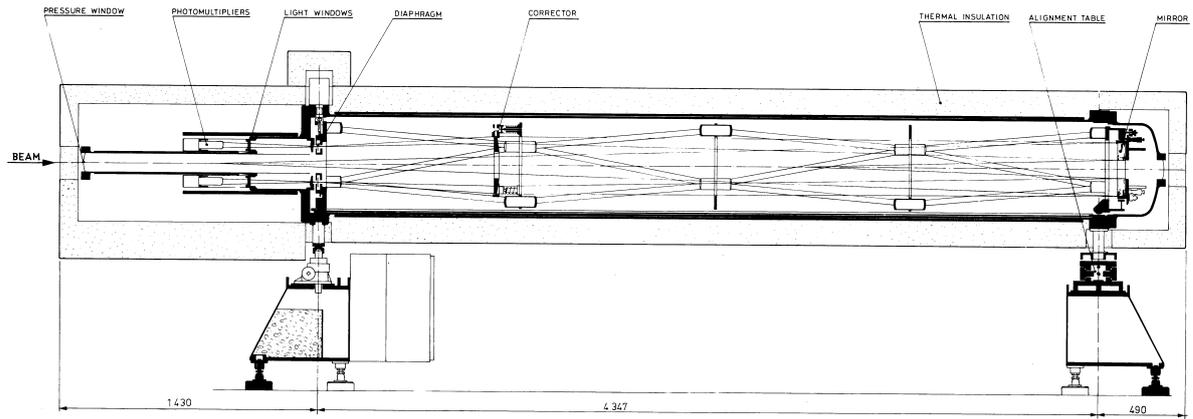


Figure 5.4.: The layout of the CEDAR detector

from the number of n -fold coincidences without normalization.

The timing of Cherenkov photons is important for the definition of a kaon crossing time. In a CEDAR detector, Cherenkov radiation is reflected by a mirror at the downstream end of the detector and is detected at the upstream end of the detector. By construction, photons emitted at the upstream end of the vessel travel for a greater distance than those emitted further downstream with the result that the arrival times of photons due to a single particle are almost independent of the point of emission. There is a spread of ~ 10 ps because the kaon velocity is greater than the photon velocity. If the arrival times of individual photons are measured with resolution $\sigma_{T,\gamma}$, then the resolution on the kaon crossing time is given by:

$$\sigma_{T,K} = \frac{\sigma_{T,\gamma}}{\sqrt{N_\gamma}}, \quad (5.17)$$

where N_γ is the number of photons detected.

5.2. The CEDAR detector

Figure 5.4 shows the general layout of the CEDAR detector. A steel vessel, 4.5 m long and 55.8 cm in external diameter is filled with N_2 gas at a nominal pressure of 1.7 bar. It is sealed at either end with aluminium windows of thickness $150 \mu\text{m}$ (upstream) and $200 \mu\text{m}$ (downstream) separating the nitrogen from the beam pipe. The whole vessel is encased in 100 mm polyurethane foam for thermal insulation.

Figure 5.5 shows the optical system. Photons are first reflected by the back surface of the Mangin

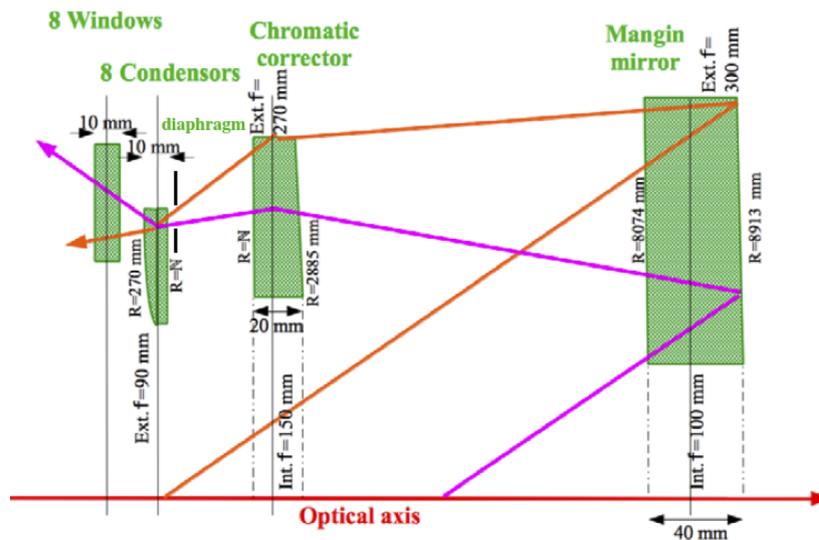


Figure 5.5.: The optical system the CEDAR detector

mirror which acts as the first lens of the chromatic corrector. The mirror is 300 mm in diameter with a 100 mm diameter hole in the centre to allow the beam to pass through. The chromatic corrector lens shares the same axis as the Mangin mirror and has a wider hole of diameter 150 mm in order to allow the propagation of Cherenkov light emitted upstream.

After chromatic correction, Cherenkov light reaches the diaphragm designed to select only rings of a certain radius. The diaphragm consists of 8 elongated apertures made from 8 outer and 8 inner segments which are moved simultaneously to vary the width of diaphragm from 0.03 mm to 20 mm in steps of 0.01 mm.

Light passing through the diaphragm is focussed by condenser lenses to form 8 light spots at the quartz window plane. In the original CEDAR design, one PMT was placed at each quartz window but these were replaced by the KTAG photon detector for NA62.

5.3. The KTAG upgrade

The NA62 nominal rate of kaons expected at CEDAR is 45 MHz. Since a single kaon typically yields about 200 photons at the upstream exit windows of the CEDAR, this would correspond to a rate of greater than 250 MHz of photons per PMT in the original 8 PMT configuration. The original PMTs cannot be used at such a high rate so a new system has been designed in order to spread the light out over more PMTs.

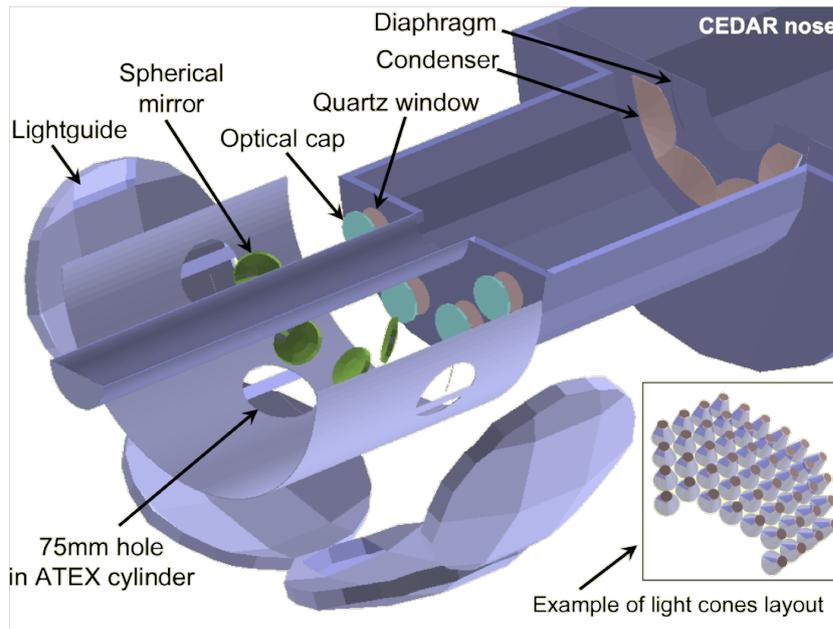


Figure 5.6.: The optical components of the KTAG upgrade to CEDAR as simulated in Geant4.

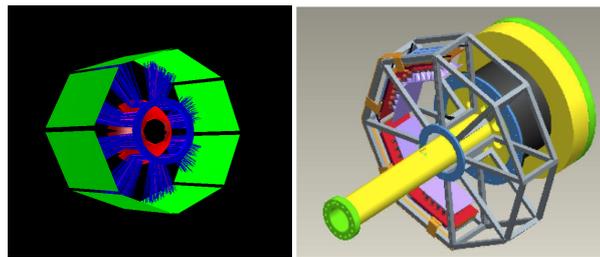


Figure 5.7.: Left: Simulation of path taken by Cherenkov photons from the CEDAR quartz windows (blue), to the new PMT planes (green), via the spherical mirrors (red). Right: The mechanical framework for KTAG.

Figure 5.6 shows the layout of the optical components in the KTAG upgrade. Instead of immediately entering a PMT, Cherenkov light passes through an ‘Optical Cap’ lens before reflecting off a spherical mirror, travelling out radially. The spherical mirror spreads the light out over a ‘light-guide’, a piece of machined aluminium, indented with Mylar coated ‘light-cones’ which guide the light towards PMTs at their narrower ends. Figure 5.7 shows the path taken by Cherenkov light. The light-guides are built to house up to 64 PMTs but simulation shows that KTAG can meet the required specifications with 48 PMTs per octant, making 384 channels in total. With this set-up the rate per PMT is ~ 5 MHz.

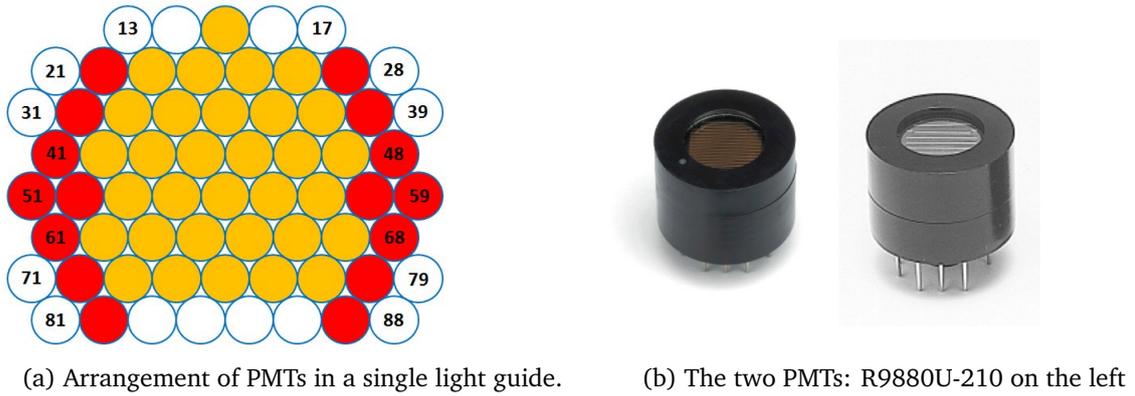


Figure 5.8.: PMTs used in KTAG. The 16 R7400U-03 PMTs are coloured red. The 32 R9880U-210 PMTs are coloured yellow.

5.4. Photo-multiplier tubes

The quantities used to describe the PMT properties are defined in Appendix B. KTAG uses two kinds of PMTs with properties suitable for single photon counting in a high rate environment, the Hamamatsu R7400U-03 and R9880U-210 modules. The former are reused from the Technical Run in 2012 while the latter are new for data-taking in 2014. Their arrangement is shown in figure 5.8. Both are single-anode metal package PMTs with a compact form; the active area is 8 mm in diameter, while the whole tube has a 16 mm \times 12 mm cylindrical shape. They have ~ 300 ps time resolution for single photons and a low dark count rate of a few Hz.

The PMT types differ in their quantum efficiency which typically peaks at $\sim 20\%$ for the R7400U-03 and $\sim 40\%$ for the newer R9880-210. They also differ in the number of dynodes in the electron multiplier system. The R7400U-03 has 8 dynodes, resulting in a typical gain of 1×10^6 while the R9880-210 has 10, resulting in a higher typical gain of 2×10^6 .

The PMTs operate at 900 V, split across the 8 or 10 dynodes by a custom-printed circuit board (PCB) which also converts the output response to a differential signal suitable for further processing [137].

5.5. Front-end and read-out system

When a photon is detected by a PMT, the response is flow of charge at the anode of about 240 fC. The corresponding voltage signal is a roughly triangular peak with $\mathcal{O}(1)$ ns rise time and 10 mV amplitude. Eventually this signal must be recorded digitally as a time corresponding to the arrival of the photon.

The first stage of processing is performed by the NINO ASIC [138]. The NINO operates in *Time-Over-Threshold* (TOT) mode, converting the differential analogue signal from the PMT anode into a *Low Voltage Differential Signal* (LVDS) output, whose width depends on the length of time the anode signal was above threshold. The threshold can be varied to maximize efficiency for real signals while minimizing the contribution from noise. Each light-box is equipped with a single board housing 8 NINO mezzanine chips. The board is also fitted with two Control Area Network (CAN) bus connectors used for communications with the Detector Control System (DCS) through the Embedded Local Monitoring Board (ELMB).

After the NINO, the LVDS signals are passed through passive splitter boards. These distribute the signals over non-consecutive readout channels in order not to overflow the latency buffers which are shared by groups of channels. The signals are processed by a *Time to Digital Converter Module* (TDC), which uses four *High Performance Time to Digital Converter* (HPTDC) chips [139] to convert the leading and trailing edges into precise timestamps. Recording both leading and trailing timestamps allows for slewing corrections. These corrections are required because anode pulses of different heights will have different rise-times leading to variation in the measured crossing time (slewing). The width of the pulse is correlated with its height so it can be used to correct for this variation.

Timestamps are processed by the TEL62 (Trigger Electronics for NA62) board [140], which is an upgraded version of the TELL1 board used by LHCb [141]. Four FPGAs, each with a 2 GigaByte memory buffer, handle incoming data from the TDC boards (one board per TDC). Another FPGA handles triggers and data synchronisation. Data and trigger primitives are passed on via four 1 Gigabit Ethernet channels, while slow control is implemented through a commercial Credit-Card PC (CCPC).

A single TEL62 with four HPTDC boards can support $4 \times 128 = 512$ channels but the expected rate is so high that if every channel was used, the readout would become inefficient. This is because, on the TDC boards, groups of 8 channels share a single latency buffer, which cannot support a total rate greater than 5 MHz. In order not to overflow the TDC buffers, only one channel is used out of each group of 8. The 384 PMT signals are therefore distributed over 3084 channels provided by 6 TEL62s.

5.6. Simulation

Simulation is vital to NA62, both to inform the construction of the detector and to interpret the experimental results. The NA62 Framework [142] is a suite of software which models kaon decays, the propagation of decay products and the interaction of decay products with the detector. All aspects

of simulation rely on Monte Carlo methods: random physics processes, such as particle decay or optical transmission, are simulated by repeatedly sampling the expected distribution of each process and aggregating the results to study output distributions.

The kaon beam (including pions and protons) is simulated using the TURTLE program which implements a transfer matrix approach to evaluate the effect of aberrations in beams with small phase-space volume [143]. Kaon decays are simulated using FORTRAN code, tuned with NA48/2 data before the decay products are passed to GEANT4 [144], which simulates particles as they travel through the detector. In the case of the CEDAR/KTAG detector, important aspects include the simulation of scattering of beam particles in the radiator gas as well as the transmission/absorption of Cherenkov photons by the optical components of the CEDAR detector. These three components are combined in the NA62MC program.

If a particle simulated by GEANT4 reaches a sensitive volume (such as a PMT), it is *digitized* by the NA62Reconstruction program. This converts MonteCarlo *truth* into data which has the same format as real experimental data by simulating aspects of detector response such as efficiencies, resolutions and hardware failures.

The CEDAR/KTAG simulation describes the propagation of Cherenkov radiation from production in the nitrogen radiator to the PMTs in the light-guides. The CEDAR digitization software then models the quantum efficiency and time resolution of the PMTs as well as any other sources of inefficiency and time resolution broadening not accounted for elsewhere. Finally, for both MC and real data, the CEDAR reconstruction software handles the conversion from a collection of digital PMT *hits* into a reconstructed *kaon candidate*. The rest of this section highlights the parts of CEDAR/KTAG simulation, which I have been involved in.

5.6.1. Cherenkov light propagation

Equation 5.1 implies that the Cherenkov photons can be produced whenever the refractive index, n , is greater than the reciprocal of the fractional velocity of the kaon, $1/\beta$. Figure 5.9a shows that this is the case for all wavelengths between 100 nm and 1000 nm. Figure 5.9b shows the Cherenkov spectrum predicted by equation 5.9 and, superimposed, the spectrum of photons produced by GEANT4. It is not necessary to simulate photons with wavelengths below 180 nm or above 700 nm because the detector is not sensitive in those regions.

The photon spectrum seen by the PMTs depends on the transmittance of the optical components

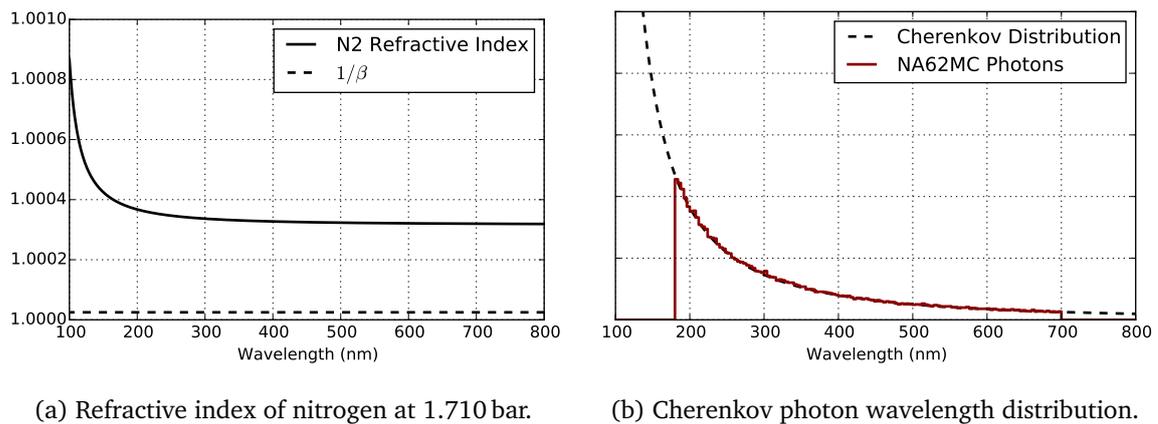


Figure 5.9.: Wavelength distributions in nitrogen. The refractive index is always greater than $1/\beta$ so Cherenkov radiation is emitted over the whole range of wavelengths shown.

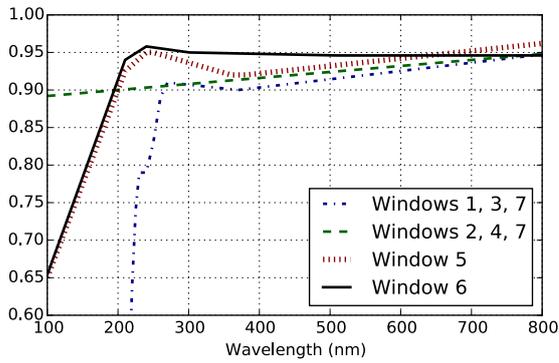
in CEDAR. Figure 5.10 shows the wavelength dependence of the transmittance. There is significant variation between the 8 quartz windows, which has an impact on alignment procedures (see chapter 6).

Figure 5.11a shows the spectrum of light reaching the CEDAR diaphragm. The tail at low wavelengths is misleading because high energy photons are not efficiently transported from the diaphragm to the PMTs. In figure 5.11b the spectrum is re-weighted by the transmission/reflection efficiency of the quartz windows, optical cap lens and spherical mirrors.

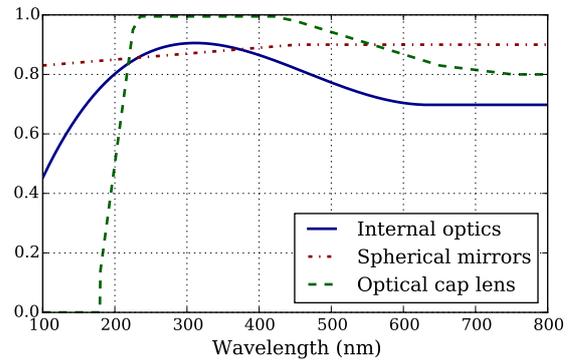
The choice of the width of the diaphragm aperture is a compromise between maximizing efficiency for detecting kaons and minimizing contamination from pions. Figure 5.12 shows the light from both pions and kaons, at the standard diaphragm width of 1.5 mm. For wider diaphragms, pion contamination increases quickly, particularly if the CEDAR is misaligned. This is discussed further in chapter 6.

The last contribution to the wavelength dependence comes from the quantum efficiency of the PMTs, plotted in figure 5.13. Also shown are the spectra of the sources and filters used to parameterize PMT response (as described in appendix B). PMTs are illuminated with light from a standard light source (a tungsten filament lamp) which has passed through one of two standard filters: a ‘Standard Eye’ designed to reproduce the response of the human eye, or the Corning Blue filter, which transmits more light at lower wavelengths than the Standard Eye.

It is clear that a spectrum determined using the Corning Blue filter is more similar to the expected distribution on the PMTs than a spectrum weighted by the ‘Standard Eye’, but it is still the case that most of the photons hitting the PMTs are expected to have wavelengths outside the range measured

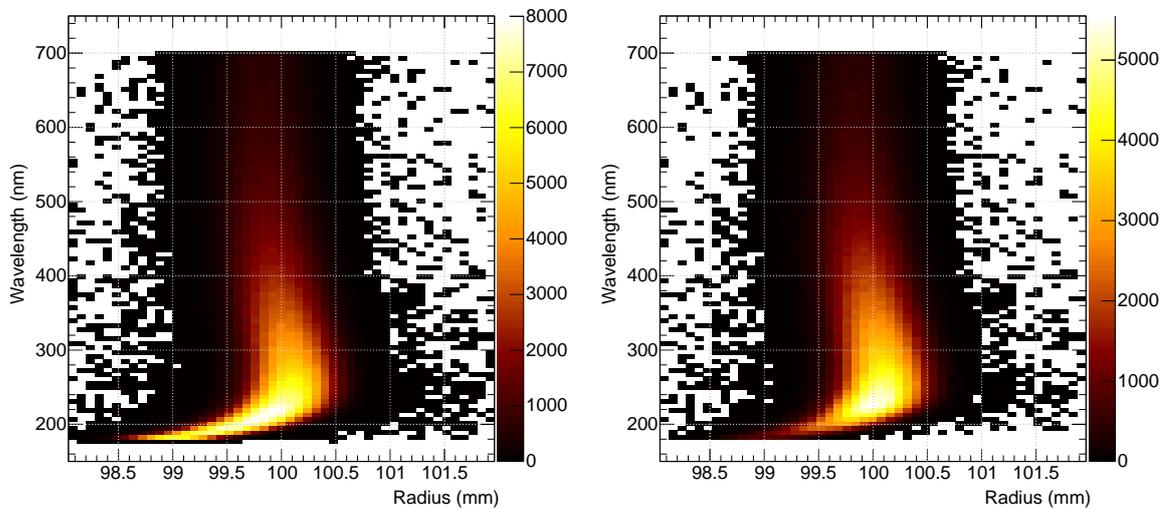


(a) The transmittance of the 8 quartz windows



(b) Mirrors, lenses and internal optics

Figure 5.10.: Transmission and reflectivity spectra of CEDAR optical components. The absorption properties of all the internal optical components are combined into a single efficiency.



(a) The radial position at which Cherenkov particles from kaons reach diaphragm plane as a function of wavelength.

(b) Photons re-weighted by transmission efficiency of the optics and quantum efficiency of PMTs in order to reproduce the distribution of light corresponding to PMT hits.

Figure 5.11.: Distribution of photons at the CEDAR diaphragm

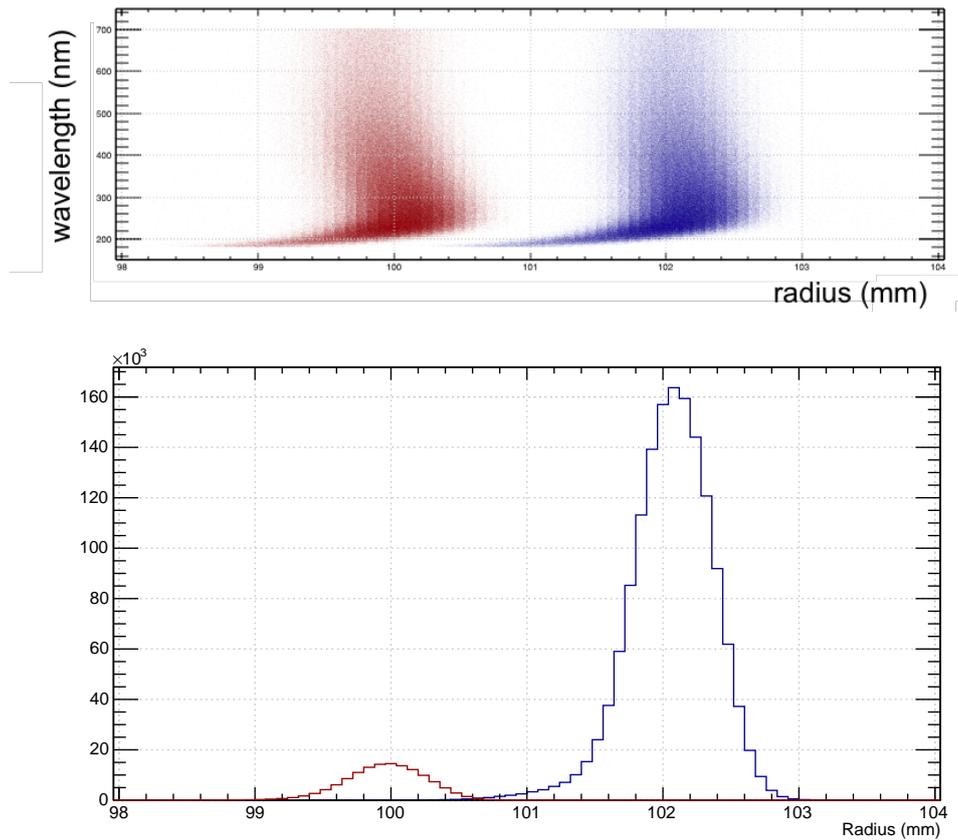


Figure 5.12.: Cherenkov radiation from kaons (red) and pions (blue) at the CEDAR diaphragm. In the top plot, wavelength is plotted against radius, showing the long tails which overlap at low wavelengths (before the effect transmission and quantum efficiency). In the bottom plot, the number of photons is normalized taking into account the expected flux of kaons and pions in the beam composition.

using the Corning Blue filter. This means that the Blue Sensitivity index is a reliable measure of quantum efficiency of individual PMTs only as far as the quantum efficiency varies uniformly, independent of wavelength.

5.6.2. Lens and mirror geometry

The purpose of the optical cap lenses and spherical mirrors is to distribute light across the array of PMTs in each lightguide. Both the focal length of lenses and the radius of the spherical mirrors impact the shape of the distribution; the best combination was chosen by simulating the available options. The optical cap lenses are plano-convex lenses made from fused silica, available with focal length 200 mm or 250 mm, and the possibility of omitting the lens completely was also considered.

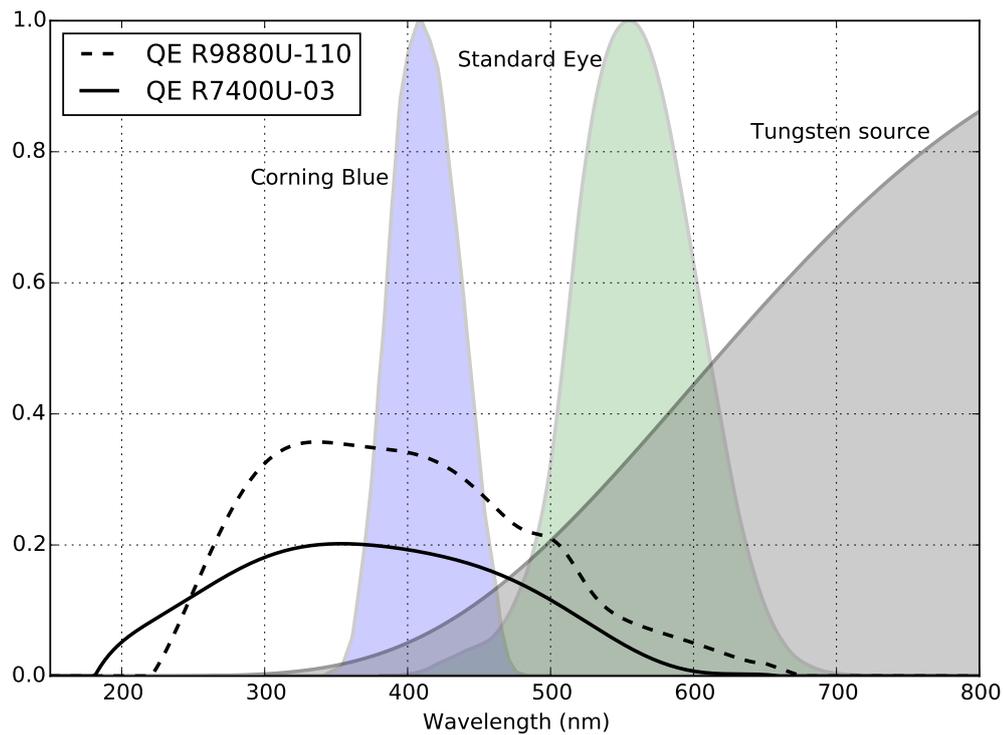


Figure 5.13.: PMT quantum efficiencies and the source and filter spectra used to measure them. The three spectra: Corning Blue, Standard Eye, and Tungsten source ($T=2856$ K) are described in Appendix B

The spherical mirrors were available with a range of radii of curvature. Figure 5.14 shows the simulated light spot on a single lightguide for the various mirror and lens combinations. The optimal configuration is with a lens focal length of 250 mm and a spherical mirror radius of 114 mm. The location of the 48th PMT was chosen based on these plots and is shown as a red circle.

5.6.3. PMT Time Response

The resolution on the kaon crossing time is crucial for matching kaons to their downstream decay products. The final resolution is measured by comparing the reconstructed kaon candidate time (determined from several PMT hits) with the trigger time recorded by the hodoscope (CHOD) downstream. It is also useful, however, to model the time response distributions of individual PMTs in order, for example, to study the time window which should be used to define a single kaon. Since an individual kaon is defined by several hits, constituting a *reconstructed candidate*, the time resolution of individual PMTs can be studied by looking at the hit time of a single PMT with respect to the others

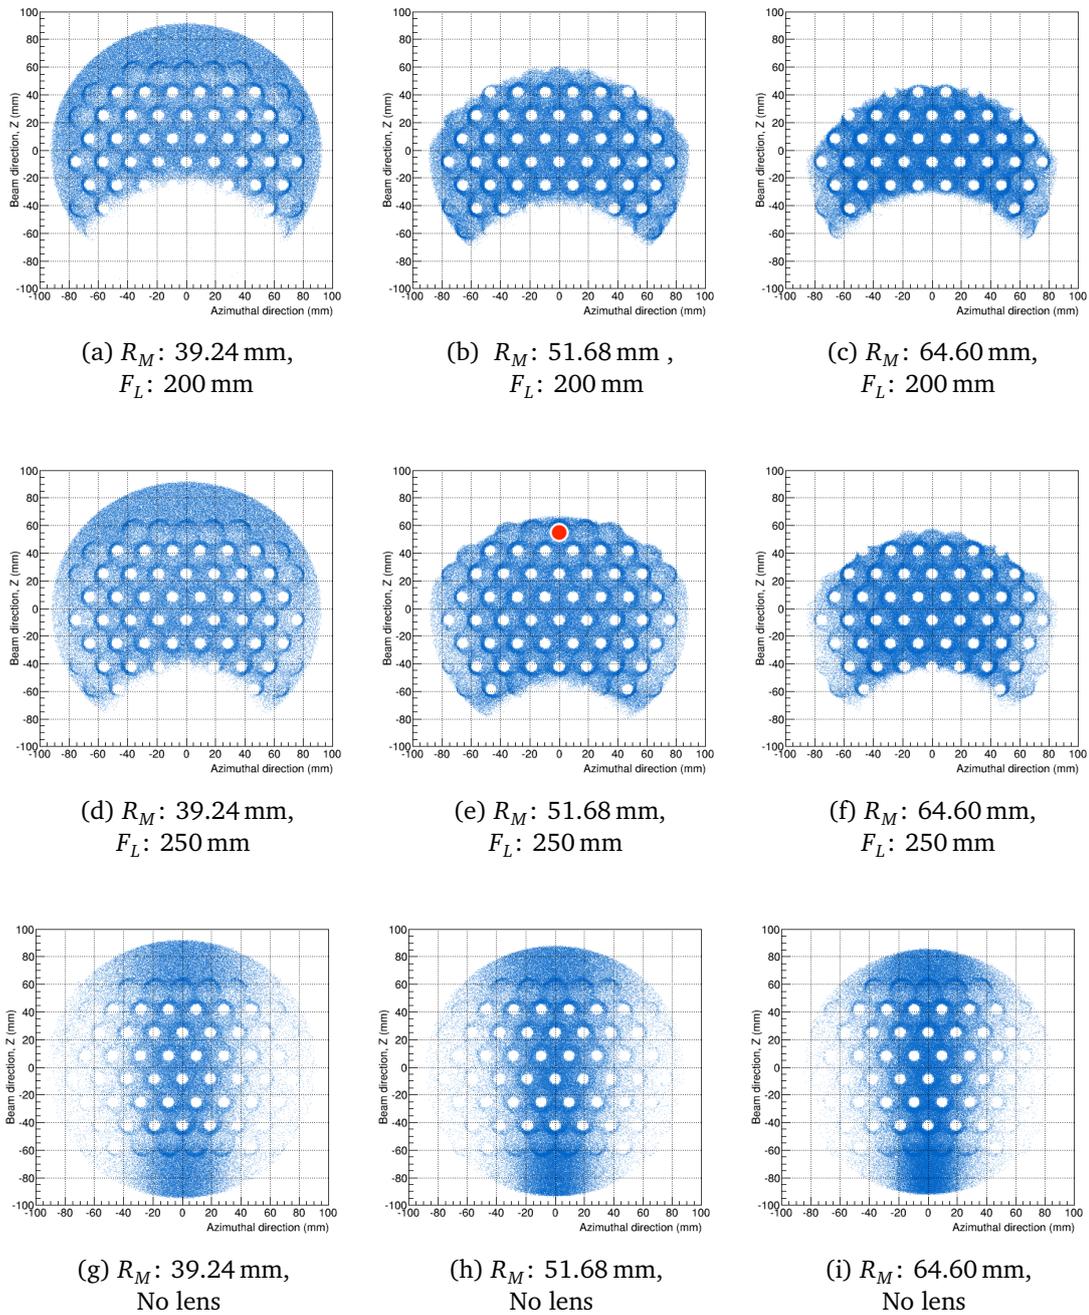


Figure 5.14.: Light profile on KTAG light-guides for various combinations of spherical mirror radius, R_M , and optical cap lens focal length, F_L . White circles indicate the locations of the PMTs. The red dot indicates the location of the new PMT introduced in the chosen configuration.

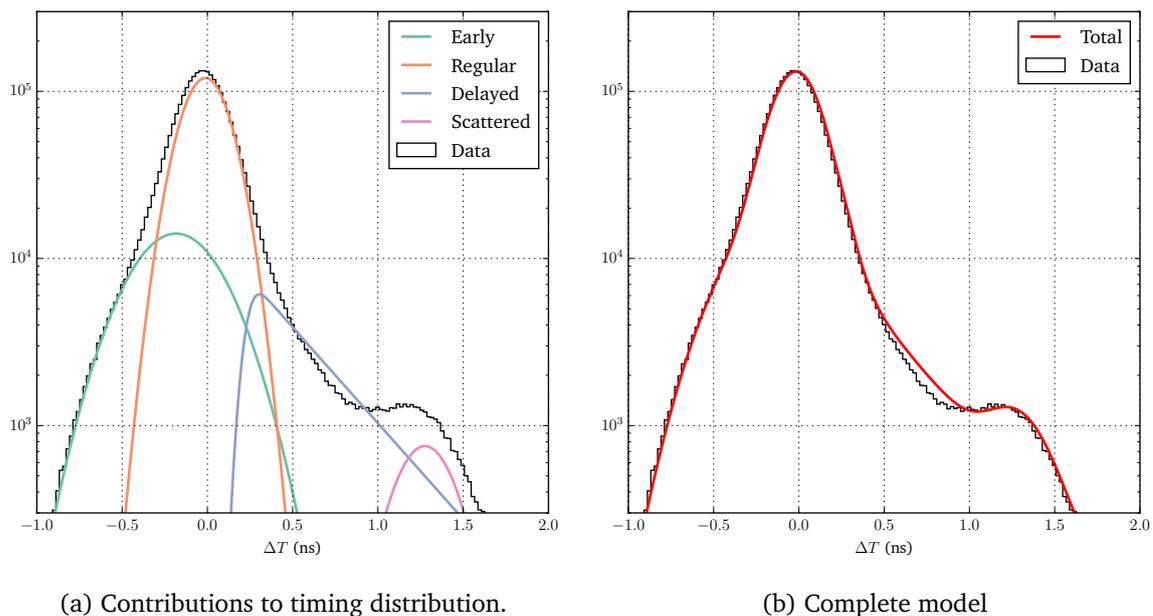


Figure 5.15.: Timing distributions for Hamamatsu R7400U-03 PMTs. For each PMT hit ΔT is the difference in time between the individual hit and the average of all the other PMT hits in the reconstructed kaon candidate.

in the candidate. Figure 5.15 shows the measured timing distribution for PMTs used in a technical run in 2012 in which only the older PMT model, R7400U-03 was used, and only 4 of the 8 sectors were populated. The timing distribution shown is for events in which at least 10 PMTs were hit. The four distributions labelled are a fit to the expected distributions of the mechanisms which contribute to the broadening of the PMT time resolution [145].

The central *regular* shape comes from the ‘normal’ operation of the PMT, in which electron transit times are spread because they take different paths between the dynodes, depending where the original photon hits the cathode and the initial velocity of the photoelectrons. The resulting distribution is Gaussian and accounts for the majority of the events. Electrons which take extreme paths in the inhomogeneous fringes of the electric field will experience suboptimal acceleration, resulting in *delayed* pulses. These follow the distribution of a Gaussian multiplied by an exponential distribution which gives the shape a long tail on the late side. Some electrons undergo elastic (without multiplication) forward scattering at the 1st dynode and arrive at the 2nd dynode early since they do not start from rest. This results in a Gaussian distribution of *early pulses*. These three effects contribute to the main peak of the timing resolution.

The secondary peak is due to electrons which are *back-scattered* at the first dynode. They are

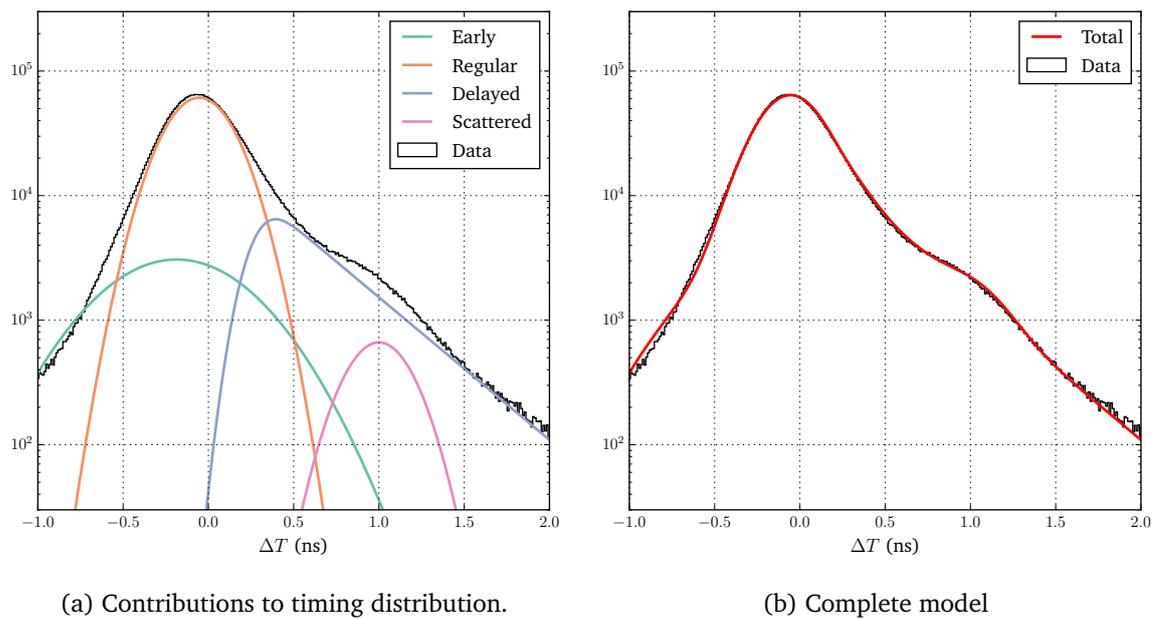


Figure 5.16.: Timing distributions for Hamamatsu R9880U-210 PMTs

re-accelerated by the electric field and produce a cascade as normal on their second impact, resulting in a pulse which is delayed by twice the transit time between two dynodes. The fit value for the centre of the scattered peak is 1.44(3) ns, which is consistent with the quoted transit time for all 8 dynodes of 5.4(3) ns.

Figure 5.16b shows the same fit applied to the new R9880U-210 PMTs in data taken in 2014. The structure is similar but the peak due to scattered photo-electrons now occurs at 1.0 ns, consistent with the shorter quoted transit time of 2.7 ns from cathode to anode.

5.7. Conclusion

The KTAG upgrade to the CEDAR detector enables it to meet the requirements of NA62 for detection efficiency and time resolution; this has been confirmed by MC simulation and experimental data [137]. My studies of the spherical mirror and lens combinations (section 5.6.2) were used to determine the arrangement of the PMTs within the octant and the plots of the photon distribution at the diaphragm aperture (section 5.6.1) are now integrated into the NA62 software validation system to ensure consistent results when new versions of the MC software are produced. I have also introduced realistic simulation of the PMT timing distributions (section 5.6.3) allowing future studies of cross-talk

between PMT channels and more detailed studies of detector timing performance.

Chapter 6

KTAG Alignment

6.1. Introduction

As mentioned in chapter 5, the KTAG detector requires good alignment between the beam and the optical axis of the CEDAR, as misalignment leads to inefficiency in detecting kaons as well as contamination from pions. (Throughout this chapter, alignment will refer to *angular misalignment*, where the optical axis of the CEDAR points in a different direction from the beam axis. KTAG operation is insensitive to *translational misalignment*, where the beam axis and optical axis are parallel but displaced in x and y). In this chapter, two procedures for aligning the CEDAR vessel are studied and their performances are evaluated. The results of these investigations were used to inform the design of the Online Monitor, which displays information to shift operators about the KTAG detector during data taking.

6.1.1. Alignment mechanism

Physically, alignment is controlled by rotating the detector. The CEDAR vessel is supported on three points (see figure 5.4 on page 158). One point is fixed at the upstream end of the detector, under the quartz windows. At the downstream end (under the Mangin mirror), the Cedar vessel is supported by two chariots resting on a V-shaped surface. The two chariots can be moved in the x -direction (horizontally) by motors. Parallel motion of the chariots moves the end of the CEDAR horizontally, controlling alignment in x , while contra motion moves the end of the CEDAR vertically, controlling alignment in y . The user interface to the CEDAR motor control hides the details of parallel and contra motion. Instead the desired CEDAR position is specified as an xy coordinate for the position of the

downstream end of the CEDAR, and the appropriate motor movements are made automatically. The exact origin of this coordinate system is arbitrary: setting the desired CEDAR position to (0, 0) does not align the CEDAR vessel with any special direction.

Since the angle of the CEDAR detector is controlled by moving the motors, the misalignment of the detector is normally expressed in mm, corresponding to the distance that the downstream end of the CEDAR is displaced from its optimum position, rather than quoting the angular misalignment in radians. The conversion between displacement of the downstream end of the CEDAR and angular misalignment is given by the distance, L , between the motors at the upstream end and the fixed point at the downstream end:

$$\Delta x = L \tan \Delta \theta_x \approx 4.347 \text{ m} \times \Delta \theta_x \quad (6.1)$$

$$\Delta y = L \tan \Delta \theta_y \approx 4.347 \text{ m} \times \Delta \theta_y, \quad (6.2)$$

where the small angle approximation is good for the whole range of misalignment expected at NA62. Aligning the CEDAR detector is equivalent to determining the x y coordinates of the motors for which misalignment in x and y is zero. The complete range of motion for the motors is ± 4 mm in x or y , with a step resolution of $10 \mu\text{m}$, corresponding to an angular resolution of $2.3 \mu\text{rad}$.

6.1.2. Alignment parameters

Throughout this chapter, the following terms are used to describe alignment of the detector.

Optimum position: the position of the downstream end of the CEDAR vessel for which the beam axis and the CEDAR optical axis are collinear.

True misalignment: the current displacement of the position of the downstream end of the CEDAR detector with respect to the optimum position (unknown in data, equal to simulated misalignment in MC). The individual coordinates are called the x misalignment and the y misalignment.

Simulated misalignment: the known x and y misalignment specified in a MC production.

Estimated misalignment: an estimate of the current x and y misalignment obtained using one of the methods outlined in the rest of this chapter.

Residual misalignment: when testing an alignment method on MC, the difference between the estimated misalignment obtained from the method and the simulated misalignment specified

in the MC production.

Motor position: the xy coordinate of the downstream end of the CEDAR vessel, in terms of the current motor positions. The individual coordinates are called the motor x position and the motor y position.

Optimum motor position: the motor position corresponding to the optimum position for the CEDAR.

6.1.3. Required alignment precision

There are two requirements in the specifications of the KTAG detector which are relevant for alignment. Firstly, the inefficiency for positively identifying kaons in the beam should not be greater than 5%. Secondly, the probability for misidentifying a pion as a kaon should be less than 10^{-4} . MC studies show that even for large misalignment (2 mm in x and y), the probability of a pion producing a 6-fold coincidence in KTAG is less than 1×10^{-6} so the more stringent requirement comes from kaon detection efficiency.

Kaon detection efficiency as a function of misalignment has been studied using the NA62MC simulation of the KTAG detector, tuned so that the average number of photons per kaon is 22, consistent with data observed in test beams. Figure 6.1 shows the inefficiency for detecting kaons with 6-fold coincidences as a function of the simulated misalignment in x . Assuming no other sources of inefficiency, the CEDAR should be aligned to better than around 0.6 mm in order to keep inefficiencies below the target of 5%. In order to have negligible impact on detection efficiency, the alignment should be better than 0.2 mm.

6.1.4. Alignment procedures

Since the CEDAR must be aligned to the beam, it cannot be positioned in advance. The best position must be determined from data taken with the beam, iterating between measurements and adjustments. NA62 has an Online Monitor, which gives feedback on data just taken by the experiment. In this chapter, various algorithms are discussed, with the purpose of providing alignment information for the Online Monitor.

When evaluating the performance of an alignment procedure, the most important factor is how accurate the final alignment is, i.e. when the estimated misalignment from the procedure is $(0, 0)$, what is the typical residual misalignment? A second important factor is how many iterations it takes

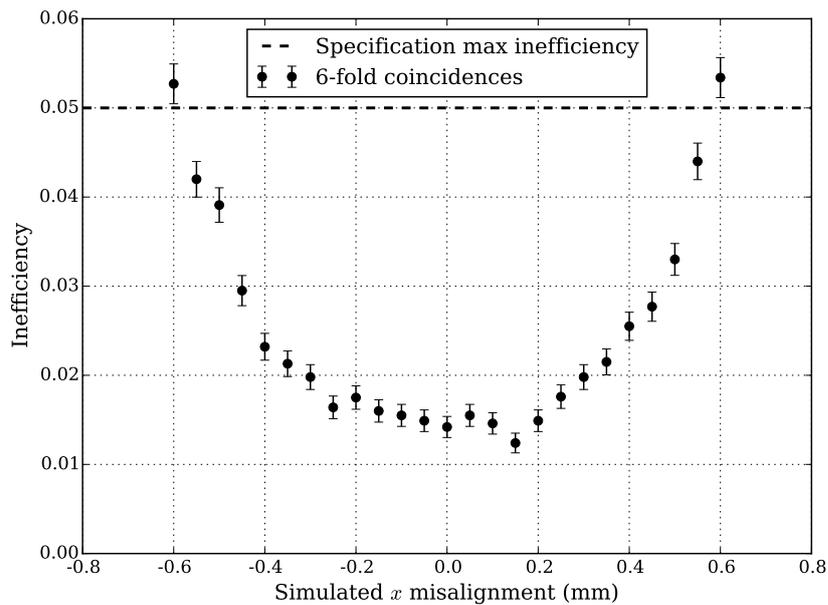


Figure 6.1.: KTAG inefficiency for 6 fold coincidences as a function of CEDAR misalignments in x , compared with the maximum inefficiency allowed in the NA62 specification.

to align the detector when starting from a misaligned position. If the estimated misalignment is always equal to the true misalignment, then the detector can be aligned in a single step, however, if the magnitude of the misalignment is not accurately estimated then several iterations will be required to reach good alignment.

All the methods presented in this chapter are based on taking a burst (corresponding to an SPS spill, see section 2.3.11) or series of bursts of data, and recording the accumulated number of hits in each PMT. The 384 numbers (one for each PMT) are then processed in order to determine how the CEDAR should be moved in order to improve its alignment.

It is often useful to group PMTs into azimuthal regions called *sectors*, each corresponding to a single lightguide. Figure 6.2 shows how they are labelled and their positions in the NA62 coordinate system. The beam direction is conventionally taken to be positive z and it is natural to take positive y to point upwards. KTAG is most easily viewed looking downstream, so it is not obscured by the CEDAR vessel, which results in the positive x direction pointing left.

In the following sections, two approaches to aligning the KTAG detector are presented and their performances are studied, both in MC and with data taken in 2014.

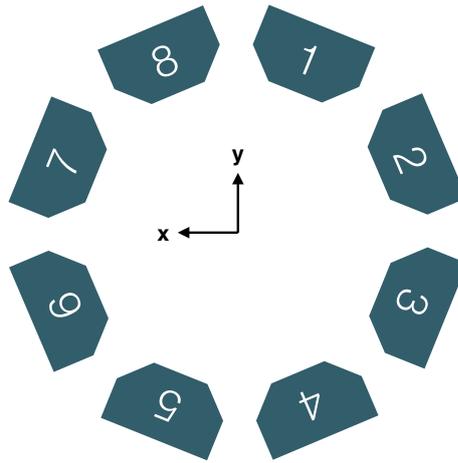


Figure 6.2.: The numbering scheme for KTAG sectors. The kaon beam travels in the positive z direction, which is *into* the page in this figure.

6.2. Quadrant Asymmetries

The traditional method of aligning a CEDAR (defined in the CEDAR operation manual [134]) is based on dividing the 8 sectors into four pairs and comparing the ratios of photon counts between opposite pairs. Four quadrants are defined:

Up (U)	Sectors 8 and 1	Left (L)	Sectors 2 and 3
Down (D)	Sectors 4 and 5	Right (R)	Sectors 6 and 7

The numbers of hits recorded in each quadrant are given by N_U , N_D , N_L , and N_R . Then the left/right and up/down asymmetries are defined:

$$A_{L/R} = \frac{N_L - N_R}{N_L + N_R}, \quad A_{U/D} = \frac{N_U - N_D}{N_U + N_D}. \quad (6.3)$$

If the KTAG detector itself is perfectly symmetric, then when the detector is correctly aligned both asymmetries should be zero. If there are asymmetries in the detector response, due to variation in the transmittance of optical components or in the efficiency of photon detection between different sectors, then the asymmetries $A_{L/R}$ and $A_{U/D}$ will be non-zero even when the detector is aligned in the optimum position. In this case additional corrections can be applied to $A_{L/R}$ and $A_{U/D}$. When using asymmetries to align the detector, it is assumed that misalignment in x (y) generates the asymmetry $A_{L/R}$ ($A_{U/D}$), with no consideration of any other correlations.

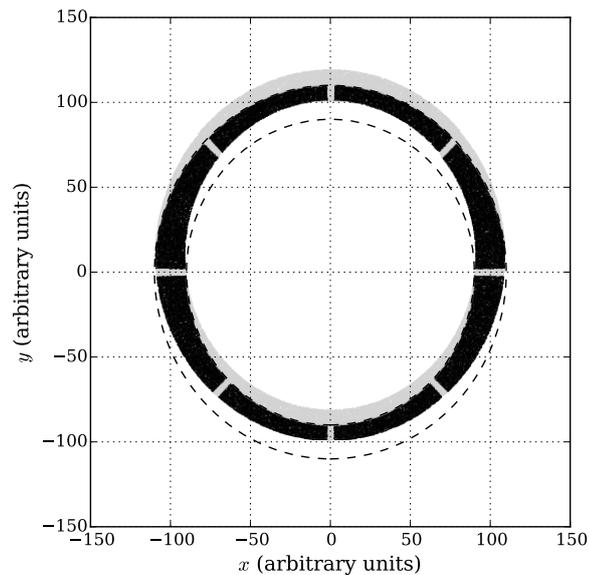


Figure 6.3.: A toy MC simulation of the Cherenkov light from kaons passing through the diaphragm aperture. The dashed lines represent the diaphragm aperture, the dark fill represents photons which pass through the aperture and the light fill represents photons which miss the aperture.

In the following section, a toy model is introduced to show how misalignment generates asymmetries in the first place. Then the full simulation of the KTAG detector in the NA62 MC software is used to study the performance of the method in more detail.

6.2.1. Simplified models

In order to understand how asymmetries can be used to align the CEDAR, it is instructive to look at the effect of misalignment in simplified models, and how it generates asymmetries in the number of hits detected in each quadrant. First a *geometrical toy model* is used to study the overlap of a ring of light with a circular aperture. The *NA62MC package* is then used to produce more realistic photon distributions.

Figure 6.3 shows a *toy model* of the diaphragm aperture, in which the radius of the aperture is 100 units and the width is 10 units. The light from the kaons is modelled as uniformly distributed in the x y plane, with 5° gaps between the sectors. In the example shown, a misalignment in y of 10 units has been simulated to show how this generates an asymmetry in the amount of light passing through the diaphragm aperture.

Figure 6.4 shows that the dominant effect is a reduction of light in both the upper and lower sectors.

The asymmetry is a second order effect arising because more light is lost from the upper sectors than the lower sectors. Asymmetry variation is shown as a function of the misalignment simulated in y . As expected, $A_{L/R}$ is not affected by the misalignment in y while the $A_{U/D}$ asymmetry varies monotonically so it could be used to extract the misalignment. The shape of the asymmetry curve depends on the shape of the diaphragm aperture.

Figure 6.5 shows the asymmetry curve for a model in which the width of both the aperture and the light ring is 4 units. The asymmetry effect is weaker so, for example, at the point where the centre of the light ring has reached the edge of the aperture, the asymmetry is 0.2 for a width of 20 units and 0.05 for a width of 4 units. This makes it more sensitive to fluctuations (both statistical and in the underlying response of the detector).

These toy models can now be compared with the official MC simulation of the KTAG detector in which a more realistic distribution of photons is produced at the diaphragm aperture. The contributions to the intensity profile at the diaphragm can be split into those due to the internal optics of the CEDAR and those due to the divergence of the beam. In order to separate these two effects, two simplified beam profiles were considered.

A simulation was produced using a *pencil beam* (with zero divergence) in order to see just the effects of the CEDAR optics: wavelength dependent transmission (see figure 5.10) and chromatic dispersion (see figure 5.11). Figure 6.6 shows the profile of the light ring at the diaphragm, when the CEDAR is aligned and misaligned in y . The curve of asymmetry vs misalignment with a 1 mm diaphragm aperture is shown in figure 6.8a. The curve has similar features to the result of the toy MC study but in the central region the curve is flat because the light ring is completely contained by the diaphragm aperture for misalignment less than 1 mm.

Figure 6.7 shows the same distributions for a simulation in which the beam had a *Gaussian angular profile* with rms divergence equal to that of the nominal beam: $\sigma_x = \sigma_y = 7 \times 10^{-5}$ rad. The wider shape of the light ring profile at the diaphragm means that a fraction of light is lost as soon as there is any misalignment of the CEDAR so that there is no flat region where misalignment does not create an asymmetry. However, for larger misalignments the asymmetries are smaller than those seen with the pencil beam. Qualitatively similar distributions were obtained for the relationship between $A_{L/R}$ and simulated x misalignment.

From these models it is clear that the relationship between alignment and asymmetry depends strongly on the shape of the beam profile and the diaphragm aperture. In the following section, the complete NA62MC simulation of the beam and the KTAG detector are used for further study.

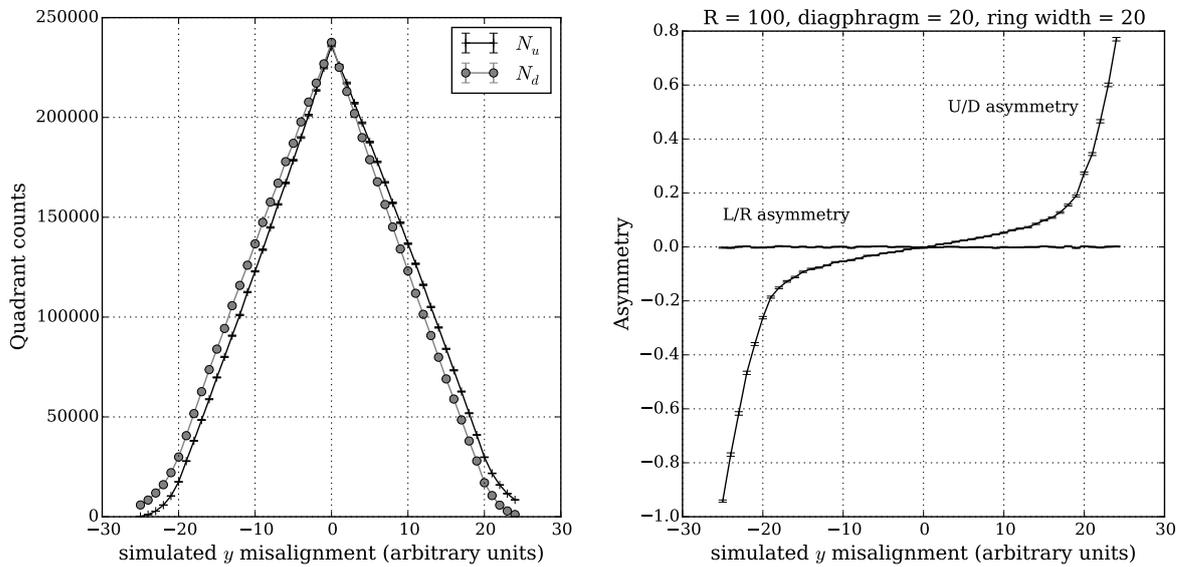


Figure 6.4.: The effect of misalignment for the toy model described in section 6.2.1. The left hand plot shows N_U and N_D as a function of simulated y misalignment. The right hand plot shows both $A_{U/D}$ and $A_{L/R}$ as a function of simulated y misalignment.

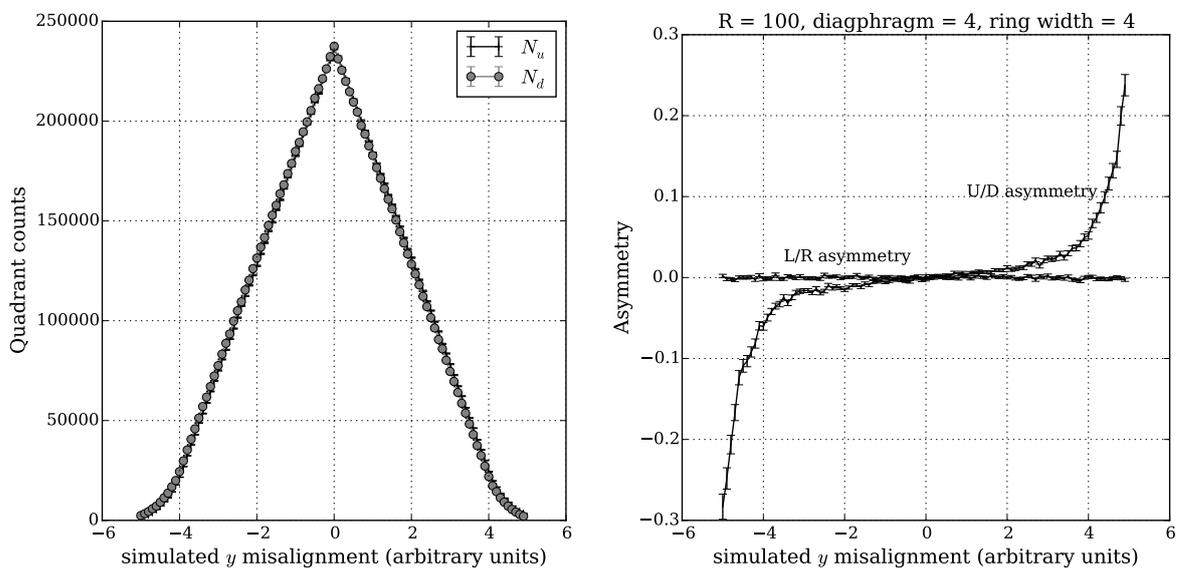


Figure 6.5.: The second toy model from section 6.2.1 with a narrower diaphragm aperture.

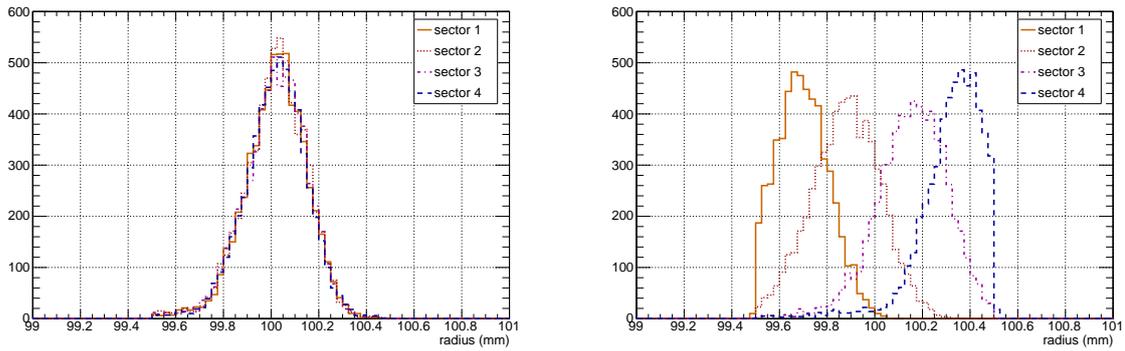


Figure 6.6.: The distribution of photons as a function of radius at the diaphragm aperture with a “pencil” beam when the CEDAR is aligned (left), and when the simulated y misalignment is 0.5 mm (right). The diaphragm aperture is 1 mm so no light is seen outside the range 99.5 mm – 100.5 mm. Since the detector is approximately left/right symmetrical, the distributions for sectors 5,6,7 and 8 are qualitatively similar to those for sectors 4,3,2 and 1.

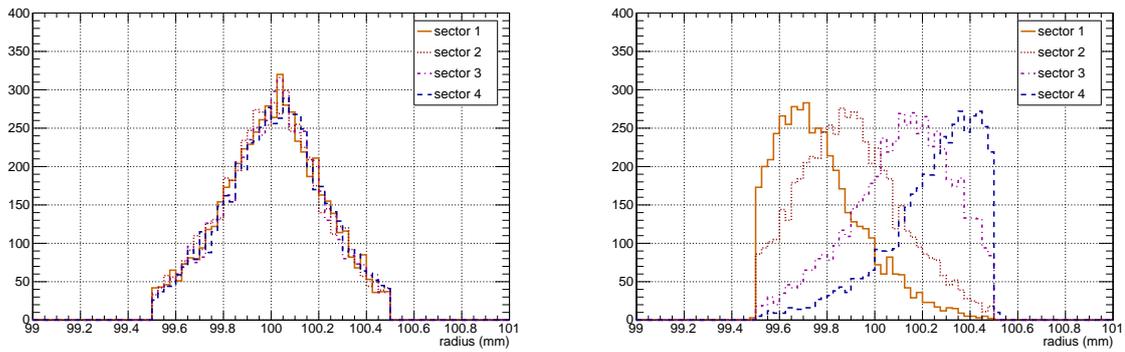


Figure 6.7.: The same distributions as for figure 6.6 but using a Gaussian angular beam profile instead of a pencil beam.

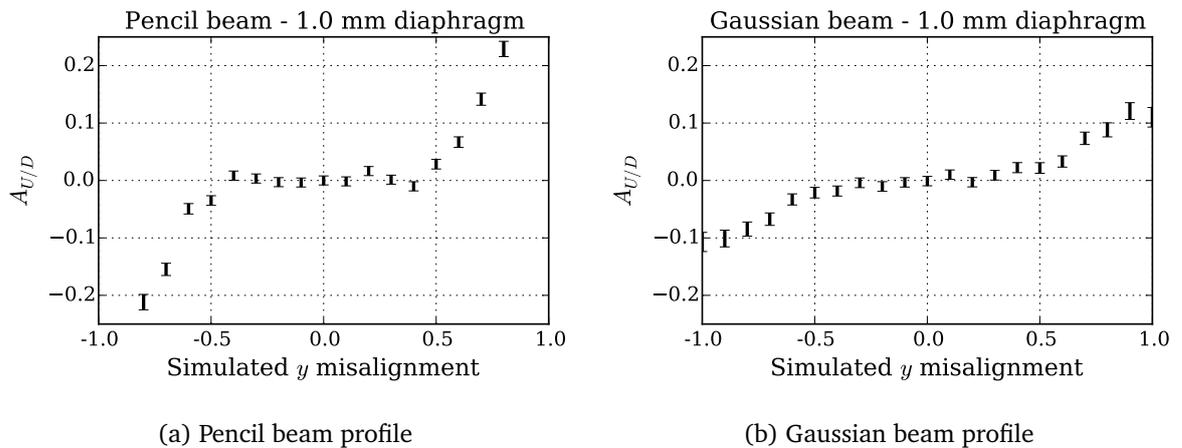


Figure 6.8.: The dependence of $A_{U/D}$ on the simulated y misalignment, with 1 mm diaphragm aperture, for two toy models of the beam angular profile.

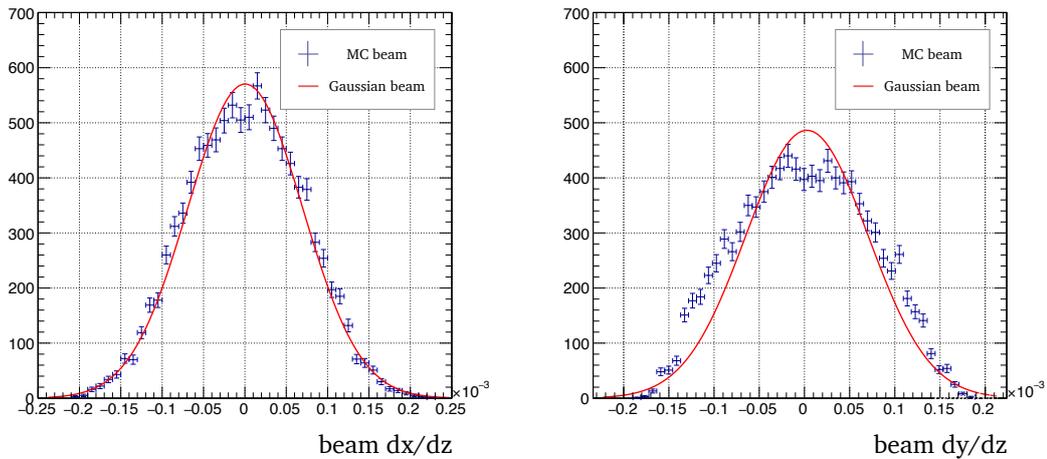


Figure 6.9.: The angular distribution of the beam in the NA62MC simulation compared with a Gaussian curve of the same rms for both x and y components.

6.2.2. Complete NA62MC simulation

The full MC simulation includes two effects not present in the simplified models of the previous section. Firstly, the beam angular profile does not follow a Gaussian distribution, but instead has wider shoulders, particularly in the y component, as shown in figure 6.9. Since the asymmetries $A_{L/R}$ and $A_{U/D}$ are highly sensitive to the tails of distribution, the difference in beam profile can have a dramatic effect. In fact, when the correct beam profile is simulated, the correlation between misalignment and asymmetry is reversed. This can be seen in figure 6.10 where asymmetry is plotted against misalignment for a range of diaphragm apertures.

The second effect arises because the response of the KTAG detector is not the same across all of the 8 sectors. In particular, the transmittances of the quartz windows (measured and implemented as part of the NA62MC software, see figure 5.10) are known to vary between sectors. The effect is visible in figure 6.10 because the measured asymmetries are non-zero even for zero misalignment.

In order to correct for this effect, the photon counts in each of the 8 sectors are re-weighted before calculating $A_{L/R}$ and $A_{U/D}$ to ensure that both asymmetries are equal when the detector is aligned. For real data, this can be achieved by opening the diaphragm aperture wide enough that the Cherenkov light ring is completely contained by the aperture and choosing scale-factors to equalize the counts in all 8 sectors. The approach has been simulated in MC using a diaphragm aperture of 20 mm to re-weight the counts in each sector; the correction factors are shown in table 6.1 and the

sector	1	2	3	4	5	6	7	8
correction	1.032	0.994	1.018	1.007	0.978	0.955	1.026	0.990

Table 6.1.: The correction factors used to re-weight the number of counts in each sector when calculating the asymmetries $A_{U/D}$ and $A_{L/R}$

results are shown in figure 6.11. By construction, the curve for 20 mm passes through the origin; for other diaphragm apertures, the corrected asymmetries are within 1% of zero for zero misalignment. For the rest of this chapter, $A_{L/R}$ and $A_{U/D}$ refer to the corrected asymmetries constructed from the re-weighted sector counts.

6.2.3. Calibration of asymmetry vs misalignment

The sign of the asymmetry is enough to determine in which direction the CEDAR detector should be rotated in order to improve alignment. In order to estimate the magnitude of the required rotation, an asymmetry must be converted into an estimated misalignment. The curves in figure 6.11 can be used to perform this conversion. Figure 6.12 shows splines fitted to simulations of misalignment with a diaphragm aperture of 1.5 mm. As already mentioned, no attempt is made to account for the effect of misalignment in y on the value of $A_{L/R}$ and vice versa. The validity of this approach is assessed in the next section.

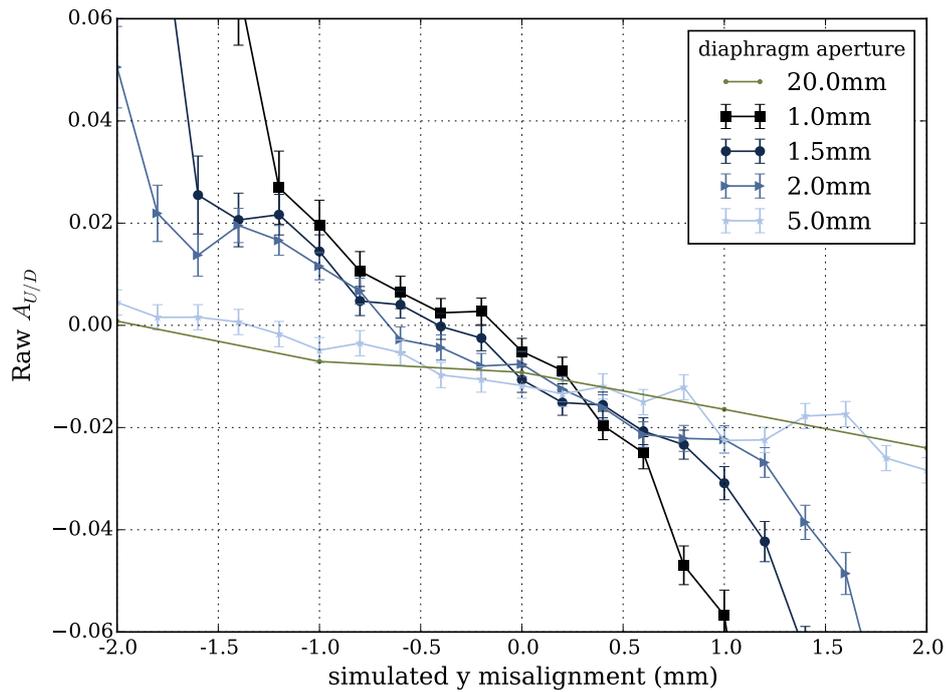
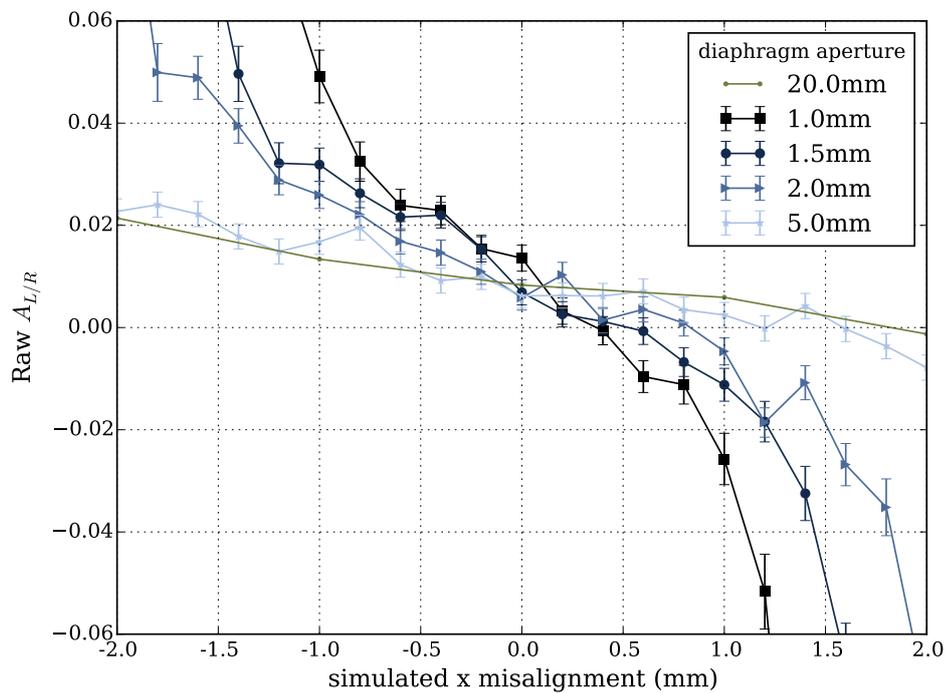
(a) Raw $A_{U/D}$ as a function of y misalignment, with zero x misalignment(b) Raw $A_{L/R}$ as a function of x misalignment, with zero y misalignment

Figure 6.10.: NA62 MC simulated response of asymmetries at various misaligned positions, with diaphragm apertures between 1 and 5 mm. No corrections are applied to the asymmetries so the lines do not pass through the origin.

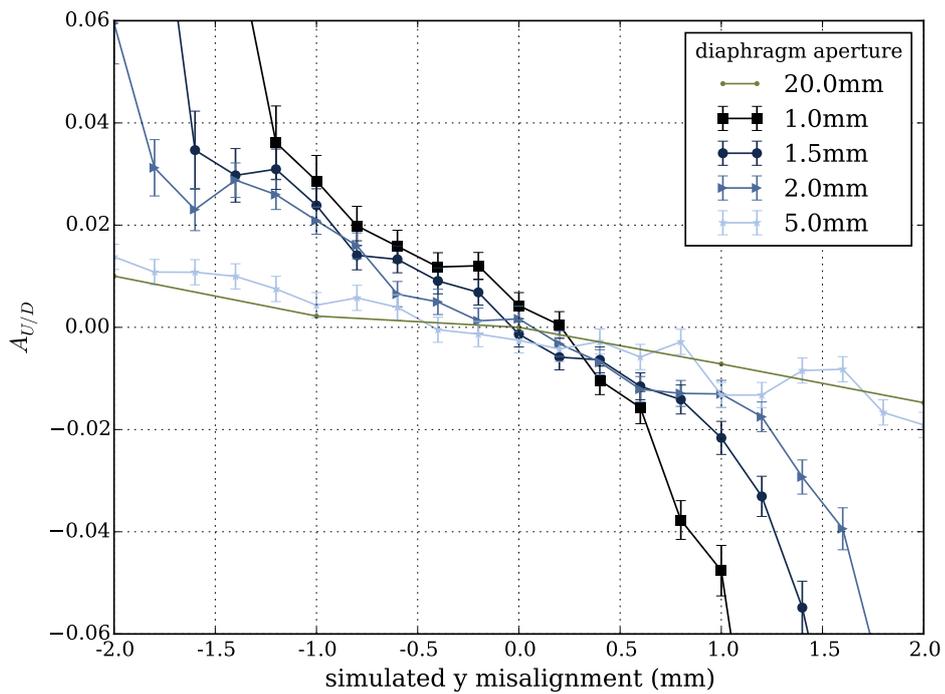
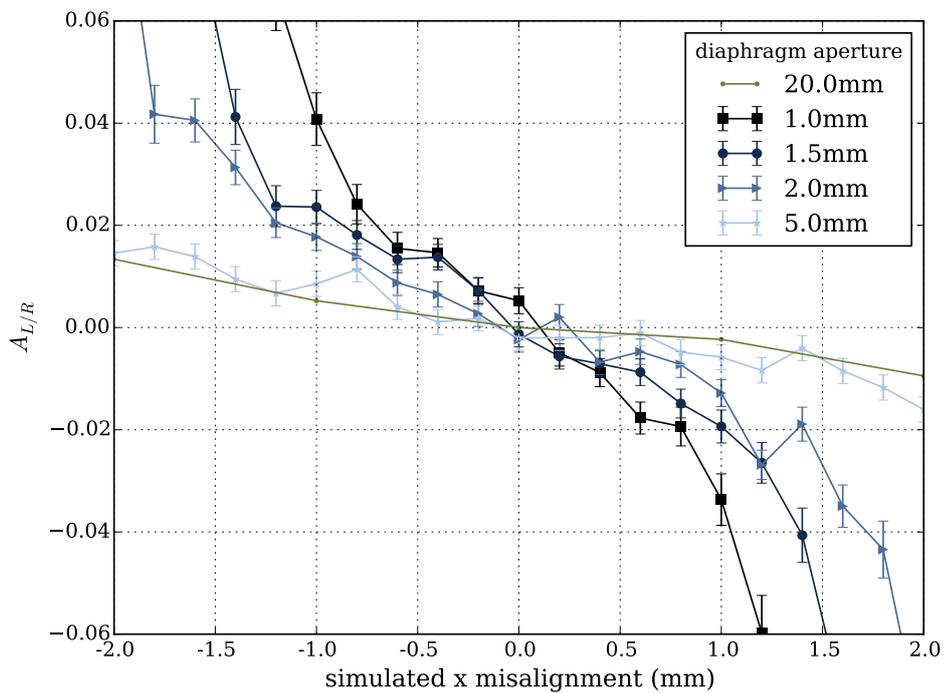
(a) Corrected $A_{U/D}$ as a function of y misalignment, with zero x misalignment(b) Corrected $A_{L/R}$ as a function of x misalignment, with zero y misalignment

Figure 6.11.: NA62MC simulated response of *corrected* asymmetries at various misaligned positions, with diaphragm apertures between 1 and 5 mm. The individual sector counts are re-weighted before computing the asymmetries.

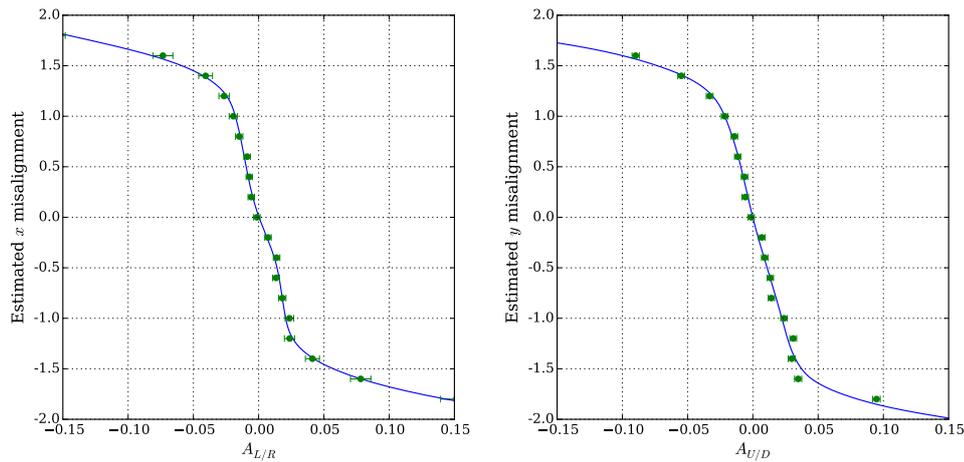


Figure 6.12.: Calibration curves for estimating CEDAR misalignment in x and y from the asymmetries $A_{L/R}$ and $A_{U/D}$, with diaphragm aperture 1.5 mm. In each plot, the points mark the asymmetries found in simulations of known misalignment. The line is the calibration used subsequently to determine the misalignment from asymmetries.

6.2.4. Performance of the asymmetry method

The calibration curves have been tested by applying them to an independent set of Monte Carlo data. Figure 6.13 shows the estimated misalignment in x and y as a function of the simulated misalignment, using the calibrated asymmetry method described in the previous sections. Figure 6.14 shows the residuals: the difference between the simulated misalignment and the estimated misalignment.

In MC simulation, this approach is successful in that it correctly returns zero misalignment when the CEDAR is perfectly aligned and the sign of the estimated misalignment is correct in nearly all cases. Although the sign is correct, figure 6.14 shows that the estimate of the magnitude of misalignment is not always accurate so often several iterations will be required to reach good alignment.

It has been seen in the transition from toy simulations (section 6.2.1) to the full NA62MC simulation (section 6.2.2) that the asymmetries used for alignment are highly sensitive to the shape of the beam angular profile. This is because, as seen in figure 6.4, asymmetries are typically a second order effect generated by the difference between two numbers which are both decreasing rapidly. Figure 6.15 shows the individual counts as a function of simulated x misalignment for a diaphragm aperture of 1.5 mm. In the asymmetry approach described so far, the estimated x misalignment is determined from the two numbers N_L and N_R which are similar in magnitude. A much better estimate of the magnitude of misalignment could come from looking at the ratio of the N_U and N_D counts to the N_L and N_R counts. The next section describes a new approach to alignment making use of all the available information to produce a robust estimate of both the direction and magnitude of the CEDAR misalignment.

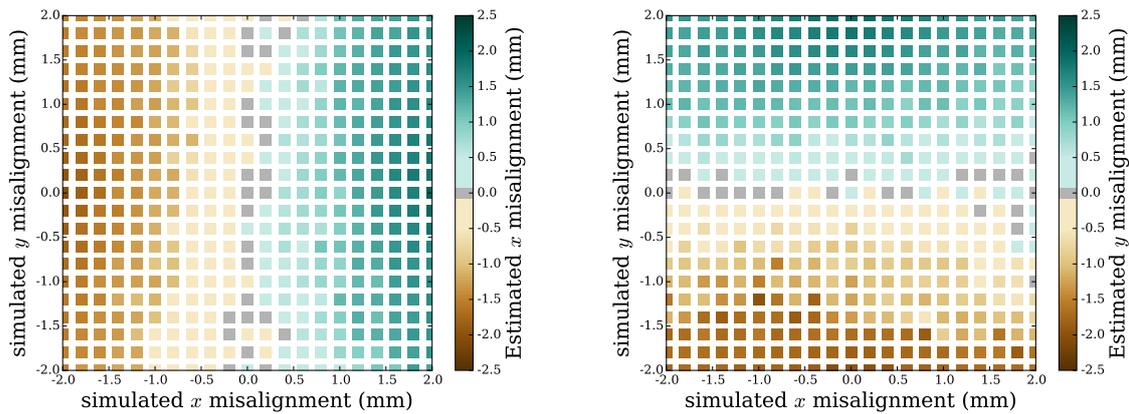


Figure 6.13.: Estimated misalignment in x (left) and y (right) using the asymmetry method indicated by colour, as a function of simulated misalignment in x and y .

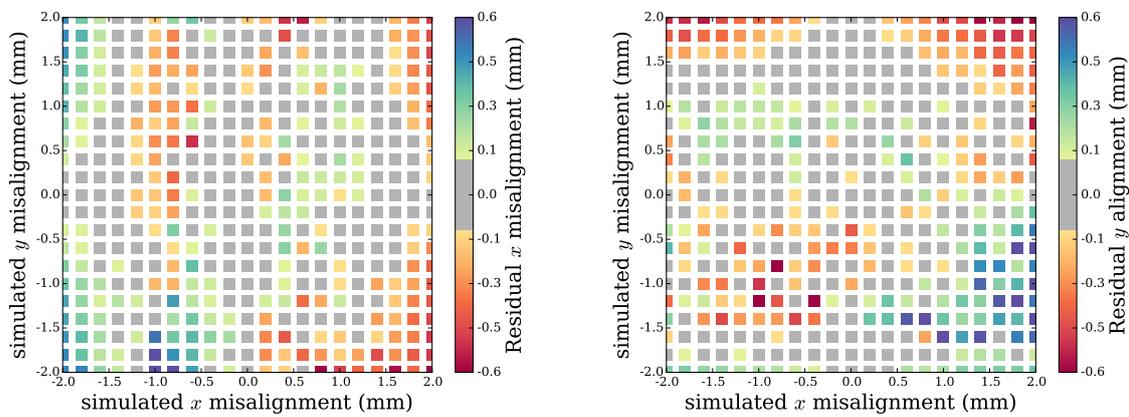


Figure 6.14.: Residual misalignment: simulated misalignment – estimated misalignment, indicated by colour.

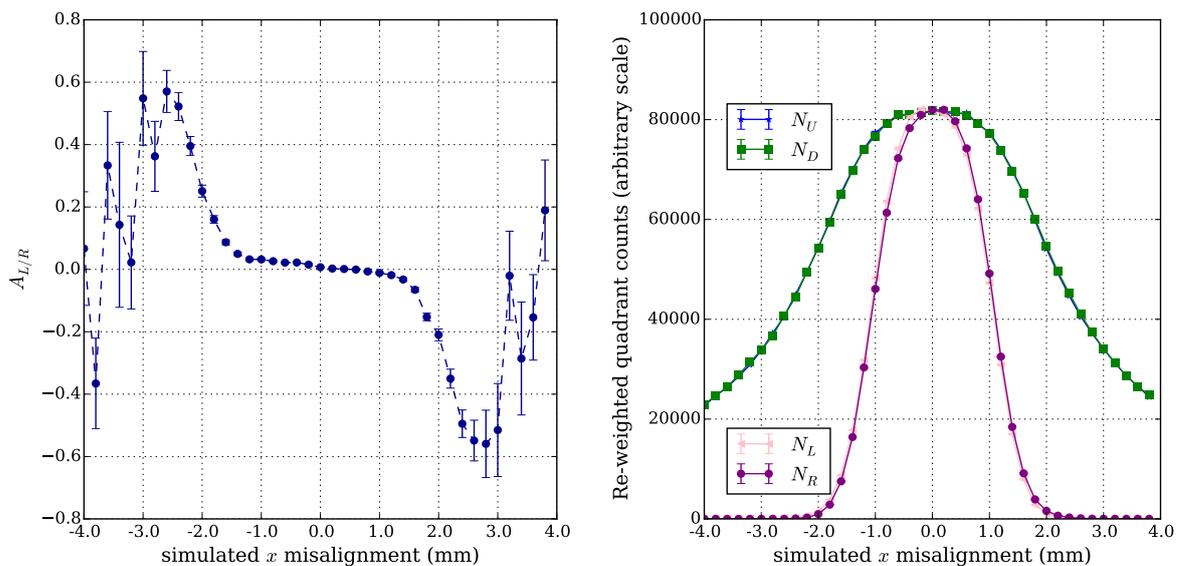


Figure 6.15.: (left) $A_{L/R}$ as a function of simulated x misalignment, showing the breakdown at large misalignment. (right) Re-weighted quadrant counts as a function of simulated x misalignment.

6.3. Monte Carlo Templates

Since it has been seen that some Monte Carlo input will be required in any alignment method, it is natural to consider a more general approach by simulating the response of the CEDAR detector at every point on a grid of possible misalignments. (A simulated data set at particular misalignment is called a *template*.) The current position of the CEDAR is determined by finding the template that best matches the real data.

6.3.1. χ^2 fits

In order to find the MC template which best matches the data, a figure of merit for the agreement between two data sets is required. A suitable quantity is given by

$$\chi^2 = \sum_i \frac{(d_i - r m_i)^2}{d_i + r^2 m_i}. \quad (6.4)$$

Here, d_i and m_i refer to the number of photons detected in data and MC respectively. The index i could run over all 384 PMTs, in which case d_i is simply the number of photons detected in PMT number i in the dataset under consideration. Alternatively, the PMTs could be combined into groups, so d_i would refer to the total number of photons detected by all the PMTs in a single group. The scale factor r is given by

$$r = \frac{\sum_i d_i}{\sum_i m_i}. \quad (6.5)$$

and ensures that the quantity χ^2 is sensitive to the shape of the light distribution across PMTs but not the integrated number of events.

While the χ^2 quantity defined above is clearly related to the χ^2 test statistic, it does not have its usual statistical properties in the present application. The differences between data and MC, even when perfectly aligned, come from real discrepancies in the simulation of the detector, rather than statistical uncertainties. This means the typical values of χ^2 are much greater than 1 per DOF, and usually most conveniently displayed on a logarithmic scale. In a real statistical fit, the value of r would be left as a free parameter, and varied to minimize χ^2 . For the present purpose, the small difference between values of r obtained in this manner and the values obtained from equation 6.5 has a negligible impact.

Once the figure of merit, χ^2 , has been defined, the alignment procedure is as follows:

1. Take a data sample, typically from a single SPS burst, and record the number of hits detected in each PMT
2. For each MC template, compute the χ^2 with respect to the data set just obtained
3. The best estimate of the current misalignment is the MC template with the smallest χ^2 .

6.3.2. Event selection and PMT grouping

When aligning with real data, there are backgrounds to the kaon signal from other particles in the beam (most importantly pions, since the associated Cherenkov ring has a radius close to that of kaons) and from accidentals, including PMT noise, beam halo muons, cosmic rays etc. In order to have a distribution which is well simulated by MC it is useful to apply some selection to the events used for alignment. In the following studies, events were required to have in-time PMT hits in 5 or more sectors.

Deciding on the groupings of PMTs is a tradeoff between maximizing sensitivity to misalignment by increasing the number of data points and minimizing sensitivity to nuisance parameters such as the variation between different PMTs by integrating over a larger number of them. For the largest possible grouping, the hits from all the PMTs in a single sector are summed so in equation 6.4, the index, i , runs over the 8 sectors.

Figure 6.16a shows an example of this approach. The KTAG response was simulated for x and y misalignments at 0.2 mm intervals between -2 mm and $+2$ mm, to create the MC templates. An independent Monte Carlo sample was produced representing data with the CEDAR misaligned to $x = 0.8$ mm, $y = 0.8$ mm. The plot shows the χ^2 between the independent sample and each of the templates. There is a clear minimum at the true value but also a secondary minimum at the opposite point. The secondary minimum arises because of the symmetry of the detector in its response to misalignment. For example, figure 6.17a is a visualisation of the number of hits in each PMT showing that, when the CEDAR is misaligned in y , the number of photons in both the upper and lower pairs of sectors decreases while the counts in the lateral pairs are much less reduced (see also figure 6.15). However, the shape of the light distribution within sectors does not vary symmetrically, so better distinction between the true and secondary minima is obtained if the PMTs are split into smaller groups. Figure 6.17b shows how the PMTs in each sector can be split into 6 smaller groups and the improvement can be seen in figure 6.16b.

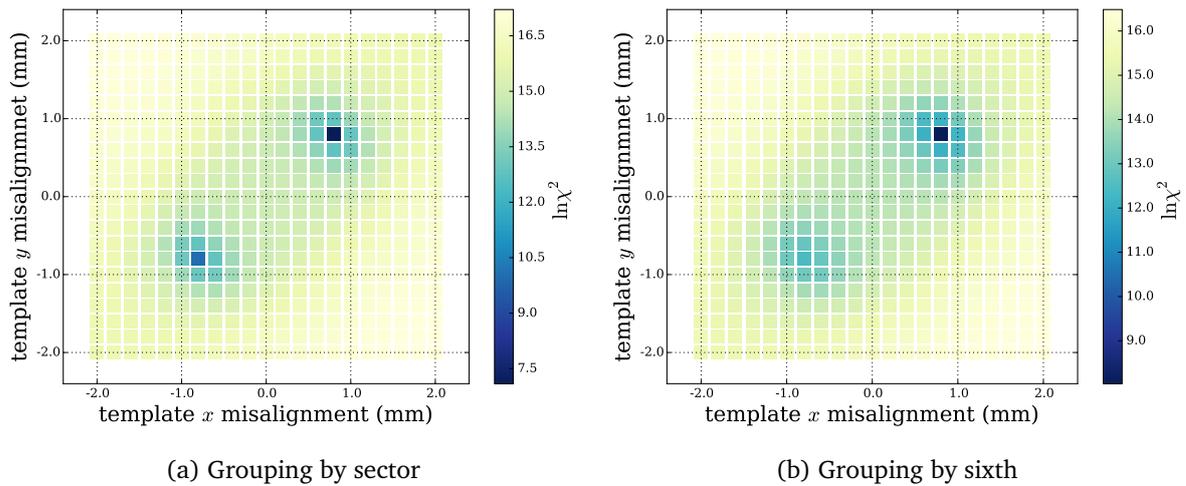


Figure 6.16.: $\ln \chi^2$ distributions obtained for a simulated misalignment of $x = 0.8$ mm, $y = 0.8$ mm

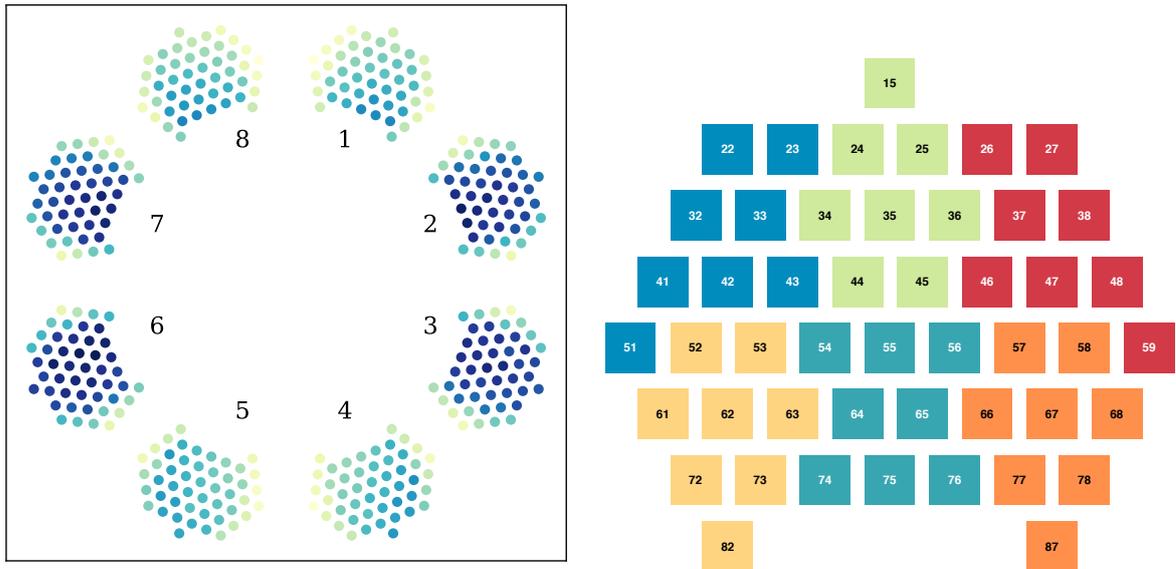
In the following sections, this approach is tested for robustness against variation in PMT efficiency and variation in the CEDAR operating pressure.

6.3.3. PMT variation

When the χ^2 figure of merit is built from smaller groups of PMTs, it becomes more important to consider variation between individual PMTs. In Monte Carlo simulation, the quantum efficiencies of PMTs are taken from the standard curves, shown in figure 5.13. It is assumed that there is no difference between individual PMTs of the same type, so only two efficiency curves are simulated, one for each model of PMT. In reality, there is some variation between PMTs, which could influence asymmetries.

The manufacturer's data sheet [146] provides some information about the relative efficiency of individual PMTs, quoting both the cathode blue sensitivity and cathode luminous sensitivity (defined in appendix B) for every PMT. Both quantities are related to the quantum efficiency, integrated over a range of wavelengths, so either quantity could reasonably be used as a measure of the relative variation in detection efficiency between PMTs.

In addition to the manufacturer's data sheet, there is another set of measurements which were performed by NA62 during the construction of the KTAG detector [137]. PMTs were mounted in the aluminium lightguide (see figure 5.6), which was illuminated with an LED, and the number of hits in each PMT was recorded. The variation in the number of hits between PMTs is a measure of



(a) Visualisation of the hits in each PMT for simulated y misalignment of 1 mm. (Darker spots indicate more hits.)

(b) Division of a sector into 6 groups of PMTs

Figure 6.17.: Visualisations of PMT layout.

the combination of PMT quantum efficiency and the reflectivity of the cones in the lightguide. From these measurements a set of equalization corrections were produced by dividing the number of hits in each PMT by the average.

Figure 6.18a shows the relationship between the cathode blue sensitivity from the data sheets and the measured equalization corrections. The difference between the two models of PMT is clear, but within each group, there is little correlation between the two quantities. This suggests that either the light-cone reflectivity and readout system have a large impact on the overall efficiency, or one or other of the measurements is unreliable. In either case, the conclusion is that there is a variation in overall detection efficiency between PMTs which cannot be completely understood or corrected for, so the alignment method should be robust against this level of variation.

Figure 6.18b shows the distribution of blue sensitivity corrections, which is the correction factor for efficiency derived from the data sheet values of cathode blue sensitivity:

$$\text{blue sensitivity correction} = \frac{\text{data sheet cathode blue sensitivity}}{\text{average cathode blue sensitivity for PMTs of the same type}}. \quad (6.6)$$

The standard deviation for each kind of PMT is less than 10% but there are long tails for the older PMT

model, R7400U-03. The χ^2 alignment method was tested with 15% variation artificially introduced. One set of MC templates was produced with default PMT efficiencies. A second set of MC templates was produced but the number of hits in each PMT was scaled by a number drawn from a Gaussian distribution with mean 1.0 and standard deviation 0.15, to simulate the effect of unknown PMT efficiencies. The effect on the residual misalignments is shown in figure 6.19. In the left hand plot the χ^2 method is tested with independent simulations with no additional variation. Statistical fluctuations sometimes mean that the template with the lowest χ^2 value is not the one corresponding to the simulated misalignment. In this case there is a residual misalignment of, typically, one or two steps in the grid of templates (0.5 mm or 0.1 mm). In the right hand plot the simulations with modified PMT efficiencies are used, and it is more common to see residuals of up to 6 steps (0.3 mm).

A separate study was performed to investigate the effect of correlations between PMT efficiencies, for example if a whole sector of PMTs was less efficient due to a high voltage supply problem or a quartz window with lower transmission. Allowing whole sectors to vary in efficiency by a single scale factor taken from a Gaussian distribution with $\mu = 1$ and $\sigma = 0.05$, the typical residual misalignment was found to be 0.1 mm in x and y .

6.3.4. Pressure variation

During commissioning of the KTAG detector, a pressure scan is performed to determine the correct operating pressure of the radiator gas, at which the Cherenkov ring produced by kaons has exactly the right radius to pass through the diaphragm aperture. The operating point is determined by finding the pressure at which the number of coincidences (6-fold, 7-fold or 8-fold) per kaon is maximum (the ‘kaon peak’ in a plot of coincidences vs pressure, see figure 6.20). If the CEDAR is misaligned during data taking then the kaon peak may be distorted so an incorrect operating pressure is selected. This in turn interferes with the alignment process. Figure 6.21 shows that when the alignment plots are produced with an incorrect pressure, the χ^2 minimum is still centered on the correct value but the shape is blurred, so it is harder to identify the minimum. The correct pressure and alignment are determined by alternating between pressure scans and alignment manoeuvres until consistent results are achieved.

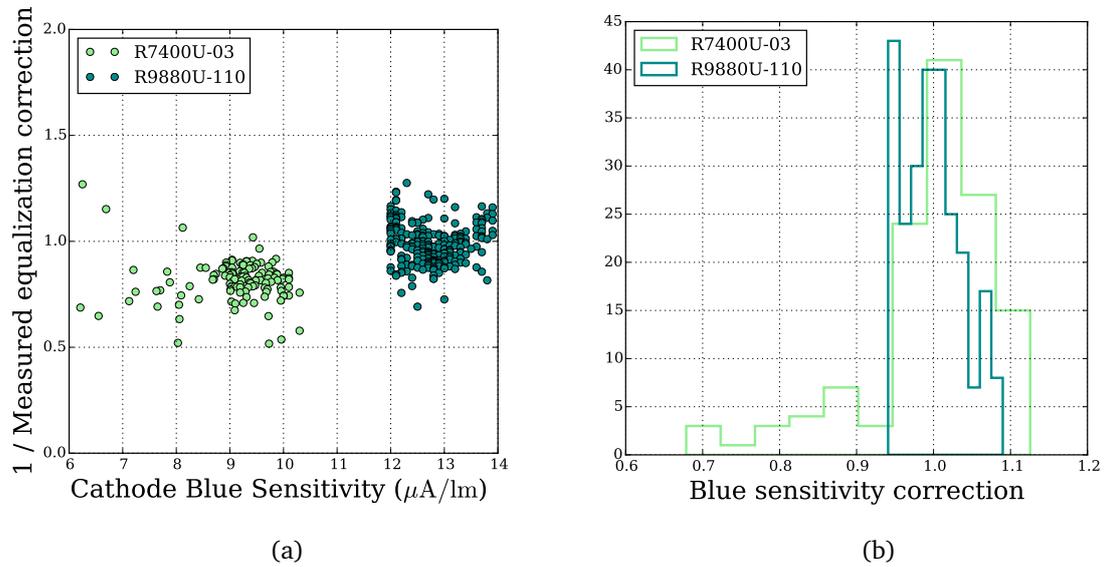
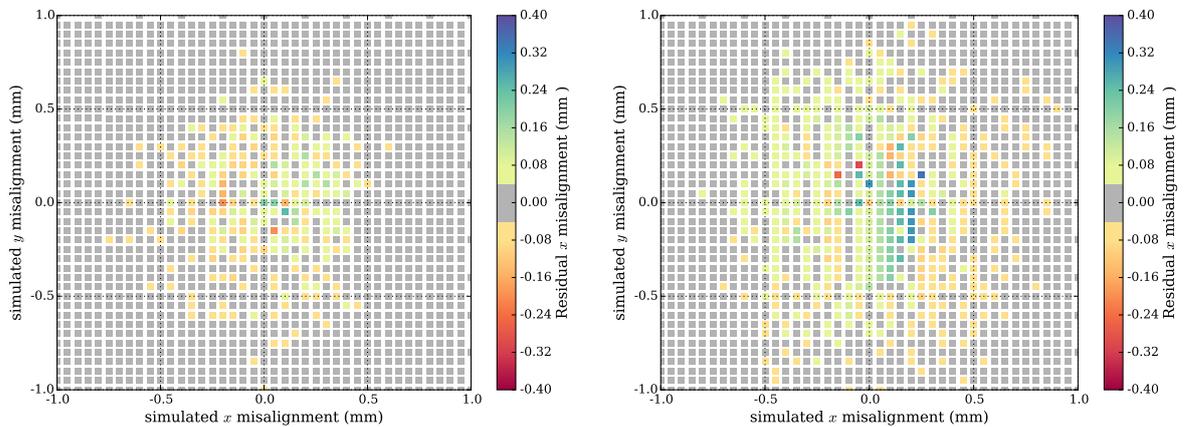


Figure 6.18.: Correlation between measures of PMT efficiency.

Figure 6.19.: The effect of variation in efficiency between PMTs. (left) the residual misalignment when using the χ^2 method with known PMT efficiencies. (right) the residual misalignment when the PMT efficiencies in the simulated templates vary with respect to the PMT efficiencies in the test samples (see section 6.3.3)

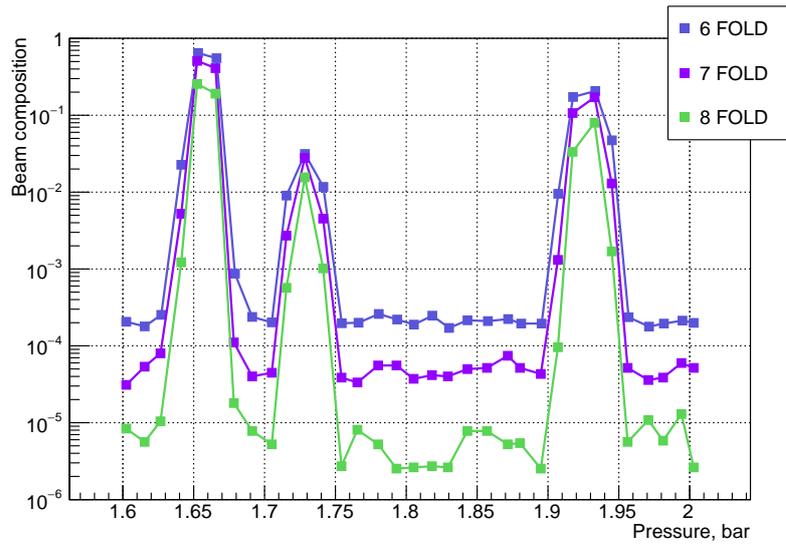


Figure 6.20.: A pressure scan from a test beam in 2011 showing the number of coincidences per trigger as a function of the CEDAR radiator pressure. The peaks correspond (from left to right) to the pions, kaons and protons.

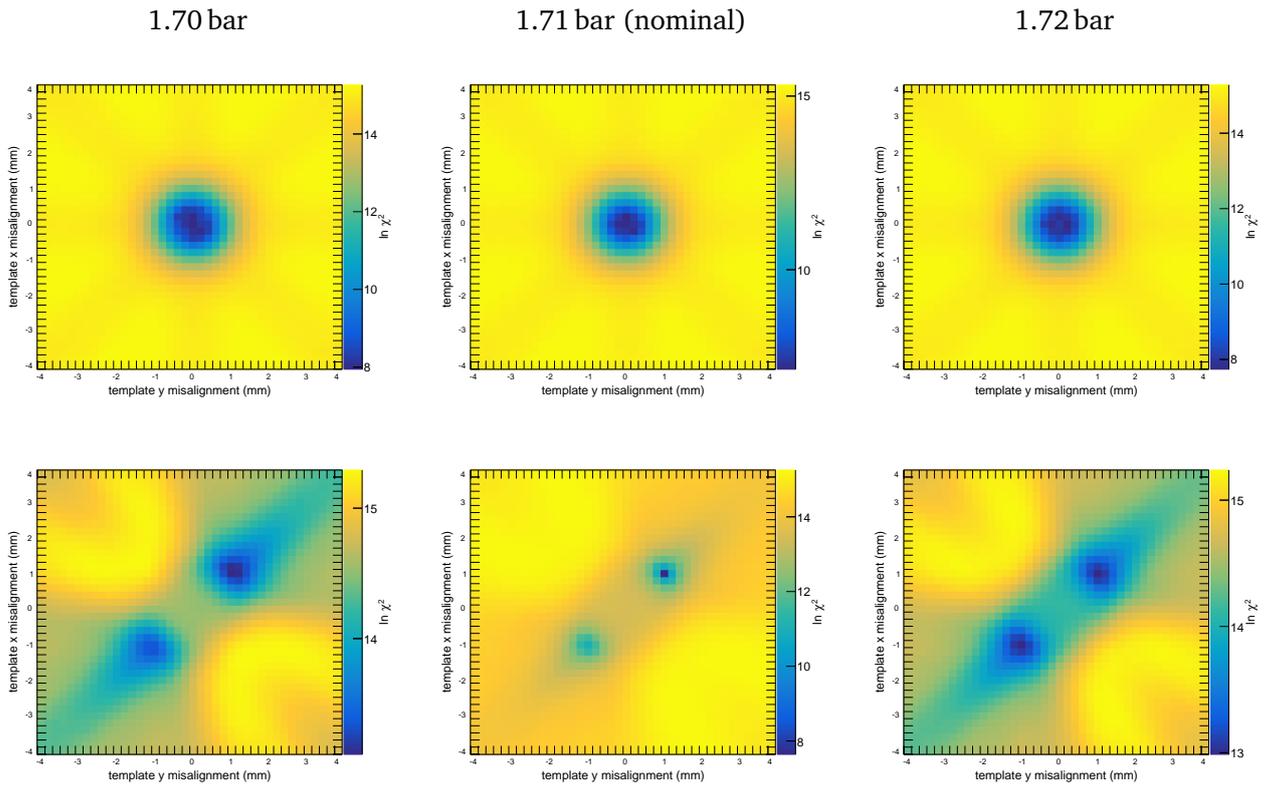


Figure 6.21.: MC Alignment plots when the radiator gas pressure is not the nominal one. The plots show χ^2 vs misalignment in x and y using the sector grouping method. In the top row, the data simulation is aligned. In the bottom row the misalignment is $[+1 \text{ mm}, +1 \text{ mm}]$.

6.4. 2014 Data

In 2014, the NA62 experiment completed a ‘Pilot’ run, with a low intensity beam [147]. The KTAG detector was fully instrumented and data were collected with several different trigger streams. This provided an opportunity to test the alignment methods on real data. In the following sections, the χ^2 approach described in the previous section is used to estimate the misalignment in x and y during data taking. Several different data taking periods are considered, in which the CEDAR motor positions were varied, in order to assess the consistency of the results obtained. The asymmetry method is applied to the same data in order to compare the two methods.

6.4.1. χ^2 alignment test data sets

In order to test the consistency of the χ^2 alignment method, a period of data taking has been analysed, in which the CEDAR motors were moved between a few different positions. For each position, the estimated misalignment was determined using the χ^2 method. Then, the relative change in estimated misalignment was compared with the relative change in motor positions to check that consistent results were obtained in each case.

Figure 6.22 shows the CEDAR motor positions recorded over a period of about 1 hour on 27th November 2014. During this time, the diaphragm aperture was constant at 1.5 mm while the CEDAR vessel alignment motors were moved between several values of x and y . This data set can be used to test the alignment methods since, although the true absolute alignment of the CEDAR was unknown, the relative change in estimated misalignment from one period to the next should be the same as the change in the position of the motors.

Figure 6.23 shows the results of the χ^2 alignment procedure applied to a single burst from the first alignment configuration (called period 1). Figure 6.24 shows the simulated result for a similar misalignment. There is a qualitative difference in the structure of the minima observed. Whereas, in MC simulation, grouping PMTs by sector typically leads to two minima which can be distinguished by sub-grouping into sixths, the opposite pattern is seen in the data. This suggests that while in MC the variation in the light profile within a sector can be used to break the symmetry between two opposite misalignments, the shape of the light profile in data is not well enough described for the same technique to work in all cases.

The most practical solution for performing alignment during data taking is to make the results of both approaches available so that the operator can resolve any ambiguities by comparing the two

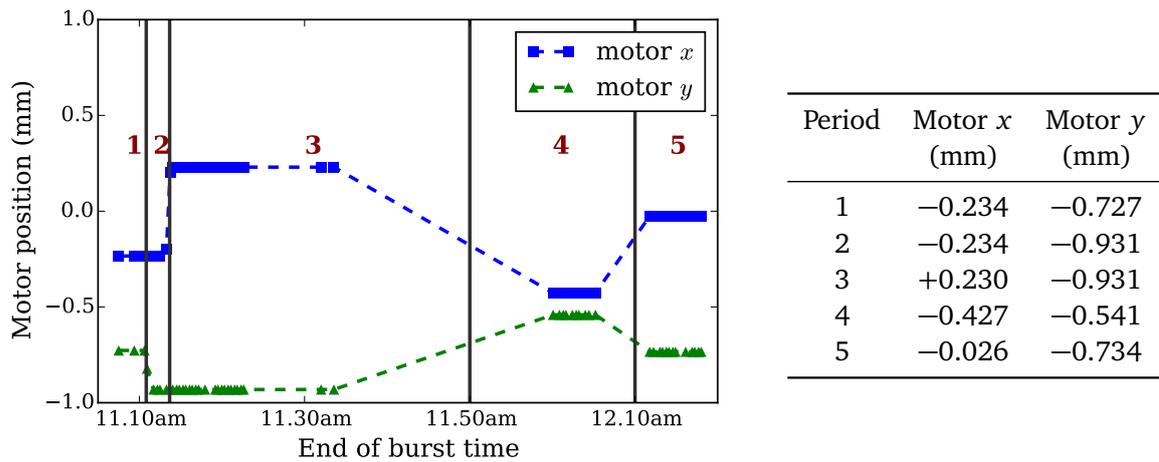


Figure 6.22.: CEDAR motor positions during a 1 hour period of data taking on 27th November 2014. 5 periods are identified with distinct configurations. The table lists the recorded motor positions for each period.

plots. For example, the choice of minima in figure 6.23b is clearly determined by the single minimum in figure 6.23a.

6.4.2. Consistency checks

In this section, the estimated misalignments for each of the 5 data periods defined in figure 6.22 are compared with the recorded CEDAR motor positions to check for consistent results. In order to compare all 5 periods at once, the results are presented in terms of the *estimated optimum motor position*, which is the estimate of the x and y motor positions which the CEDAR should have been moved to in order to obtain perfect alignment:

$$\text{estimated optimum motor position} = \text{current motor position} - \text{estimated misalignment.} \quad (6.7)$$

If the alignment method produces consistent results, then any change in motor position should be matched by an equal change in estimated misalignment, so the estimated optimum motor position should be constant.

Presenting misalignment estimates in terms of the estimated optimum motor position provides a way to remove the ambiguity when there are two minima in the χ^2 distribution resulting in two estimates of misalignment (as seen in figure 6.23b). Each estimate is used separately to determine the optimum motor position and both positions are recorded in a 2D scatter plot of optimum motor x

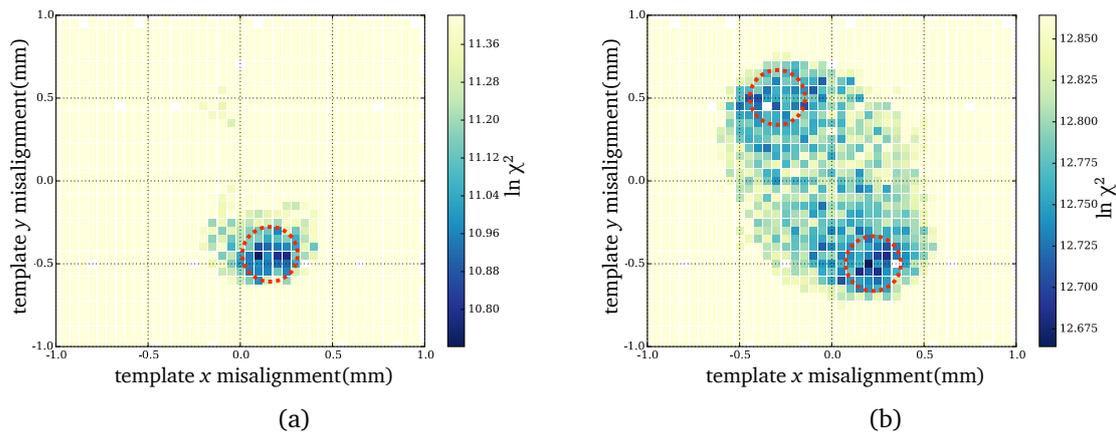


Figure 6.23.: Plots of $\ln \chi^2$ vs template misalignment for the first burst of alignment period 1 (defined in figure 6.22). (a) sector grouping. (b) sixths grouping. Minima are circled in red.

vs optimum motor y . Then the CEDAR motors can be moved to another position and the process is repeated. After several iterations, the estimated optimum motor positions corresponding to the correct choices of minimum should cluster around the same point – the correct value – while those corresponding to incorrect choices will be scattered across different positions.

Figure 6.25 shows the results of this procedure applied to the five data periods. For each burst, there are 2 minima for the fit grouped by sector and 2 minima for the fit grouped by sixths, making 4 separate estimates of the motor position required for perfect alignment. These are all plotted in 6.25a. As expected, there is one position where the estimates cluster, around $(-0.4 \text{ mm}, -0.5 \text{ mm})$.

Assuming that the true position for perfect alignment must be close to this region, the process was repeated but, this time, only the best of the 4 estimates for each burst is selected, defined by proximity to the point $(-0.4 \text{ mm}, -0.5 \text{ mm})$. The results in figure 6.25b show that in all the bursts considered there was at least one estimate close to the assumed best point. The spread in estimates is about $\pm 0.25 \text{ mm}$, consistent with the size of the minima in, for example, figure 6.23a. As a cross-check, the procedure was repeated using only the newer model PMTs (see section 5.4), which occupy the central locations in the lightguides so they are less sensitive to the shape of the tails of the beam profile. The results, shown in 6.26, are consistent with the fits using all the PMTs.

6.4.3. Comparison of χ^2 and asymmetry precision

As a check of the consistency of the two alignment measures discussed in this chapter, the estimated optimum motor position determined from the χ^2 method, $(-0.4 \text{ mm}, -0.5 \text{ mm})$, has been used to

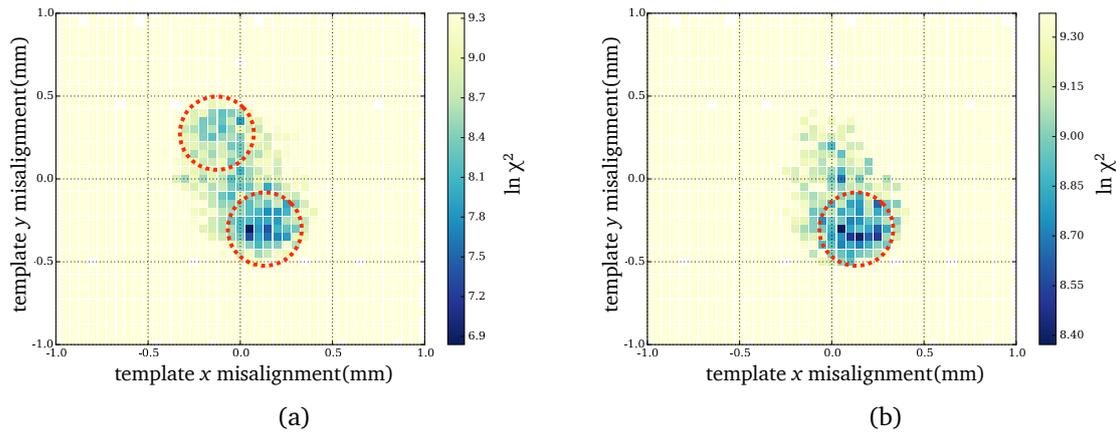


Figure 6.24.: Plots of $\ln \chi^2$ vs template misalignment for a MC sample produced with similar misalignment to that seen in figure 6.23. (a) sector grouping. (b) sixths grouping.

produce a plot of asymmetry vs misalignment for the data. For each point in figures 6.27a and 6.27b, the estimated misalignment on the x axis is calculated from the overall optimum motor position and the known motor position for the burst (burst by burst estimates of the optimum motor position are not used). The asymmetry is calculated from the re-weighted numbers of hits in each sector, using the scale factors from MC shown in table 6.1.

Since the accuracy of misalignment estimates is dominated by systematic uncertainty rather than limited statistics, it is difficult to extract a quantitative measure of the uncertainty on individual measurements. However, rough estimates of the tolerances of each method can be made from the data and used to compare the two methods.

An estimate of the uncertainty in the χ^2 fits can be made from examples such as those in figure 6.23, in which there are several points with values close to the minimum value of $\ln \chi^2$ spread over a circle of radius 0.15 mm. In the χ^2 method, the best estimate of misalignment is determined from the bin with the lowest value of $\ln \chi^2$, with no attempt to smooth out the distribution so inaccuracies of 0.15 mm are expected. In figure 6.26b, the majority of the estimates cluster within a similar range of ± 0.15 mm, although the outlying points suggest that the method can sometimes fail by a larger amount.

An estimate of the uncertainty in asymmetry measurements can be made from the range of asymmetry values obtained for a single motor position, i.e. the difference between the highest and lowest point in each column of figures 6.27a and 6.27b. The typical range of asymmetries measured for bursts with the same motor position (ignoring outliers) is 0.01. Referring to figure 6.11, the

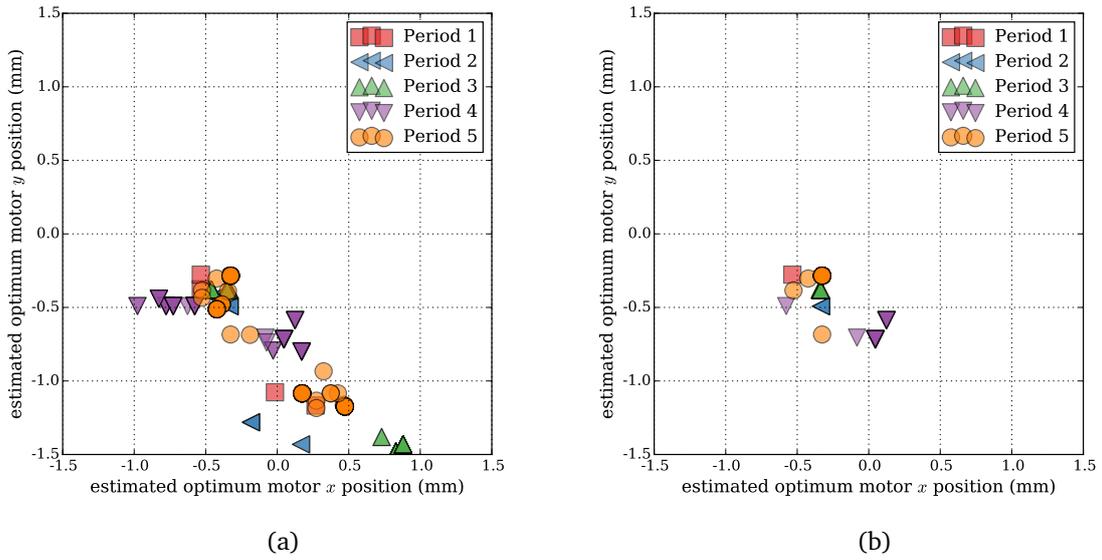


Figure 6.25.: Estimates, from the χ^2 fits to all PMTs, of the motor positions at which the CEDAR should be perfectly aligned. (a) All estimates. (b) Estimate most consistent with other values.

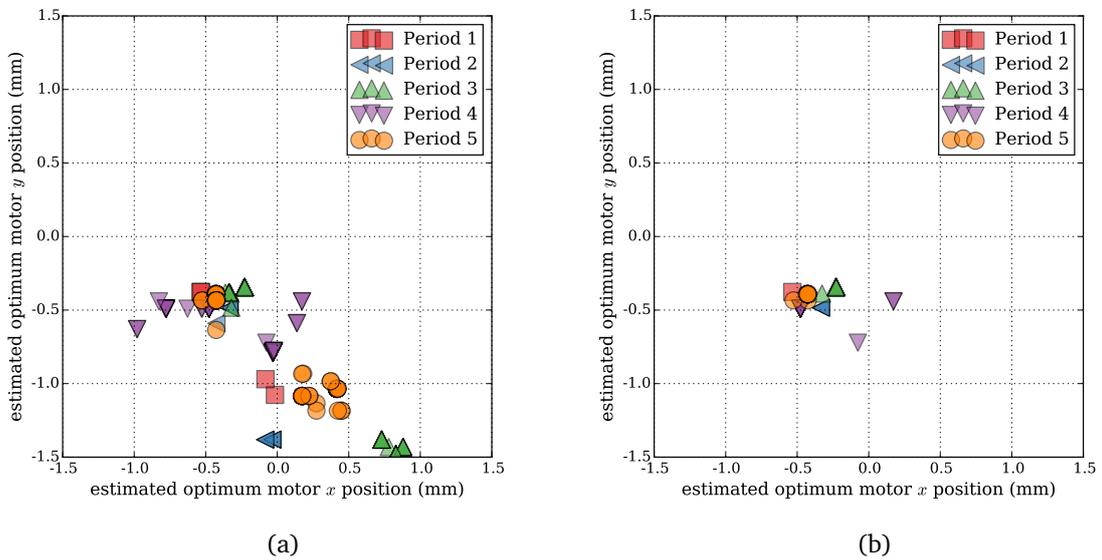


Figure 6.26.: Estimates, from the χ^2 fits to only the new PMTs, of the motor positions at which the CEDAR should be perfectly aligned. (a) All estimates. (b) Estimate most consistent with other values.

relationship between misalignment and asymmetry is approximately linear for small misalignments, with a gradient of around 0.02/mm for all diaphragm apertures. However, in the data, in figure 6.27, the observed gradient is steeper, around 0.1/mm. This is evidence that asymmetries cannot reliably be used to extract the magnitude of misalignment by comparison with MC. Using the data value of 0.1/mm, a shift in asymmetry of 0.01 corresponds to a shift in misalignment of 0.1 mm.

An alternative test of the asymmetry method can be made by constructing an approximate envelope function from the range of asymmetry measurements at each point. This approach is shown in figure 6.27. For each misalignment position, lines are drawn from the maximum and minimum asymmetry measurements (ignoring outliers) at this position to the maximum and minimum asymmetry measurements at other positions. The tolerance of the method is then given by the range of misalignments for which the envelope of all the lines intersects the abscissa corresponding to zero asymmetry. For $A_{U/D}$, this range is -0.21 mm to -0.11 mm giving a tolerance of 0.1 mm y -misalignment. For $A_{L/R}$, the range is 0.05 mm to 0.25 mm corresponding to a large tolerance for x -misalignment of 0.2 mm. However, this value is dominated by the two low lying points from period 4; if these outliers are ignored, the range is significantly reduced. For the rest of this section, the smaller error is used, since the larger would artificially improve the agreement of the two methods. Assuming a tolerance of ± 1.0 mm in x and y , the expected error in the xy plane is 0.14 mm, similar to the error in the χ^2 method.

The x -intercept in figures 6.27a and 6.27b provides an estimate of the agreement of the two methods. For $A_{U/D}$, the intercept (taken as the centre of the range given by the envelope) is at 0.2 mm meaning the asymmetry method would report perfect alignment when the χ^2 method reported misalignment of 0.20 mm. For $A_{L/R}$, the figure is -0.16 mm, so the combined disagreement in the xy plane of misalignment ($\Delta s = \sqrt{(\Delta x)^2 + (\Delta y)^2}$) has magnitude 0.26 mm. Since the errors on the two methods are not statistical, they cannot be reliably combined, but the discrepancy, $\Delta s = 0.26$ mm, between the methods is less than the sum of estimated errors of the two methods, (0.14 mm + 0.15 mm = 0.29 mm), so they can be regarded as consistent within the quoted accuracies.

Even if the entire worst-case discrepancy was attributed entirely to one of the methods, Figure 6.1 shows that a misalignment of 0.26 mm results in a tolerable level of inefficiency for kaon tagging, well below the 5% limit.

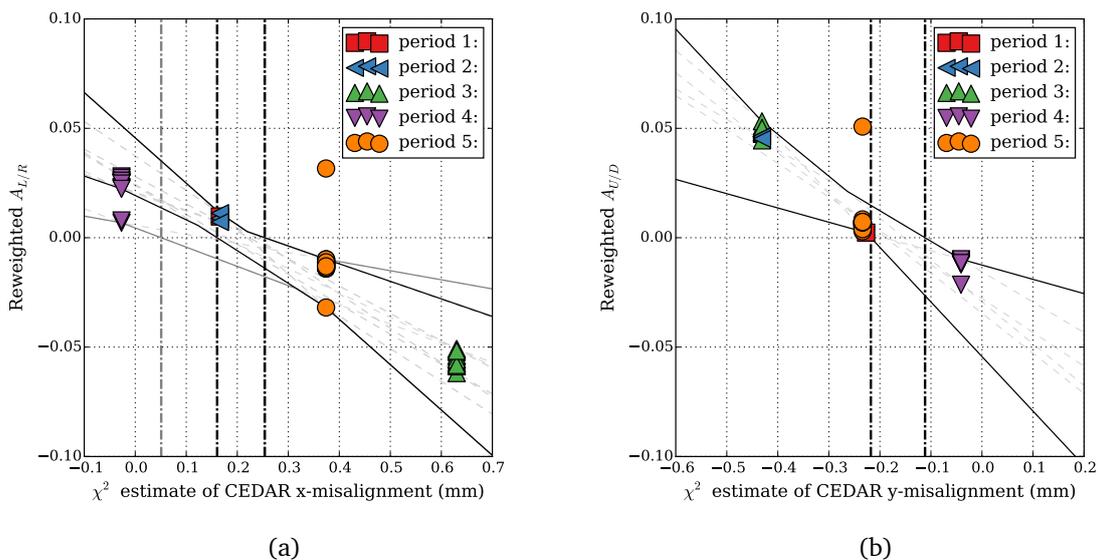


Figure 6.27.: The asymmetries $A_{L/R}$ and $A_{U/D}$ for each burst in the 5 periods under consideration, plotted against the estimated misalignment in x and y determined from the χ^2 method. Solid black lines show the estimated envelope of asymmetry curves described in section 6.4.3, constructed using the dashed gray lines. (In figure (a) a gray line shows the effect of including outlying points). Vertical dash-dotted lines show the intersection of the asymmetry envelope with the line of zero misalignment, used to derive an estimate of the error in the method.

6.4.4. Fine tuning

Once the approximate alignment has been determined by visually inspecting plots like the ones in figure 6.23, the CEDAR position can be fine-tuned to maximize efficiency. One useful measure comes from equations 5.15 and 5.16 for determining the number of photo-electrons per kaon (either ϕ_{78} or ϕ_{68}) from the ratios of 6,7 and 8-fold coincidences¹. If every sector had the same efficiency and the CEDAR was perfectly aligned, ϕ_{78} and ϕ_{68} should be equal, but if the CEDAR is misaligned, then the ratio is no longer expected to be unity. Figure 6.28 shows that for the more realistic case where the sectors have different efficiencies, the ratio is not exactly 1.0 even when perfectly aligned, but it does have a minimum at $x = 0$. This means the ratio ϕ_{68}/ϕ_{78} can be used as a check of fine alignment without knowing the exact individual PMT efficiencies, and without needing any external normalization.

¹ Recall the two expressions $\phi_{78} = \ln\left(1 + \frac{8}{\eta_7/\eta_8 - 1}\right)$ and $\phi_{68} = \ln\left(1 + \frac{14}{\sqrt{4 - 7(1 - \eta_6/\eta_8)^2} - 2}\right)$ are both estimates of the number of photons per kaon.

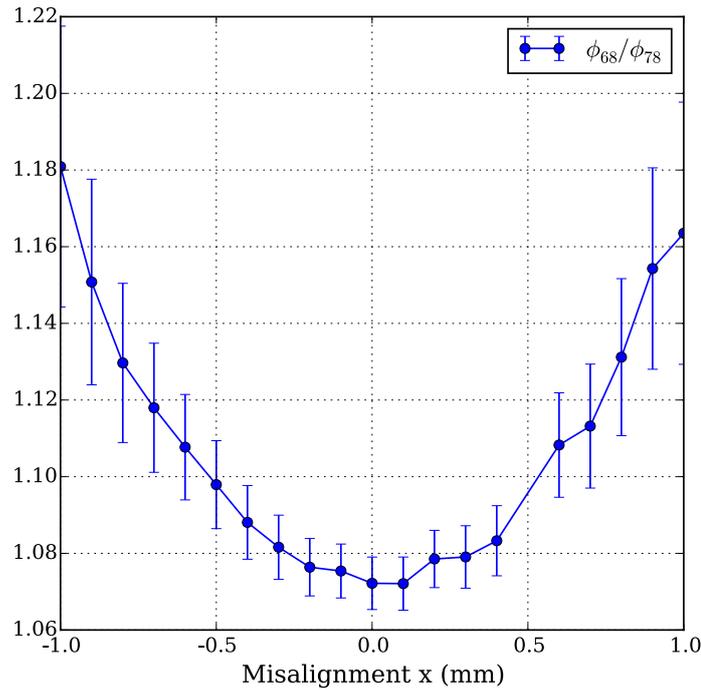


Figure 6.28.: The ratio, in MC simulation, of the two expressions for the number of photo-electrons per kaon, ϕ_{68} and ϕ_{78} , as a function of misalignment in x .

6.5. Conclusions for the Online Monitor

It has been shown that both the alignment methods, quadrant asymmetries and χ^2 fits, can be used to align the CEDAR detector to within 0.3 mm of the optimal position, corresponding to an angular range of 7×10^{-5} rad. In addition, the χ^2 method produces reliable estimates of the magnitude and direction of misalignment, while the asymmetry method performs less well in estimating the magnitude of misalignment, since it is sensitive to the tails of the beam angular distribution which are not well under control. On the other hand the χ^2 method often fails to produce a single minimum and the sign of the asymmetry could be useful in determining the correct minimum in ambiguous cases.

In order to facilitate alignment as quickly and easily as possible, the results of both methods are included in the alignment section of the NA62 KTAG Online Monitor. Figure 6.29 shows the plots as they appear in the Online Monitor, implemented with the ROOT software package. The counts in each PMT are reset every burst and the plots re-generated.

A typical use case might proceed from right to left. The plot on the far right, ' χ^2 Raw sixths', shows

χ^2 vs simulated misalignment for a large range: ± 4 mm in x and y . This can provide information at the very start of the alignment procedure, when the CEDAR might be misaligned beyond the range of the other χ^2 plots and the asymmetries do not provide reliable results. No selection is applied to the number of sectors in coincidence, since for large misalignments the efficiency for 5-fold coincidences is reduced.

The central two plots show the χ^2 plots discussed in this chapter, for simulated misalignment in x and y between -1 mm and $+1$ mm, using 5-fold coincidences. The left hand plot shows the number of counts in each sector along with the re-weighted asymmetries. The three plots, ‘sector counts’, ‘ χ^2 sixths’ and ‘ χ^2 sectors’, can be used together to align the CEDAR with an expected precision of better than 0.3 mm. Finally small adjustments can be made while observing the value of the ‘ ϕ ratio’, which is the value of ϕ_{68}/ϕ_{78} defined above.

The Online Monitor was implemented for the 2015 run of the NA62 experiment and used in the commissioning of the KTAG detector.

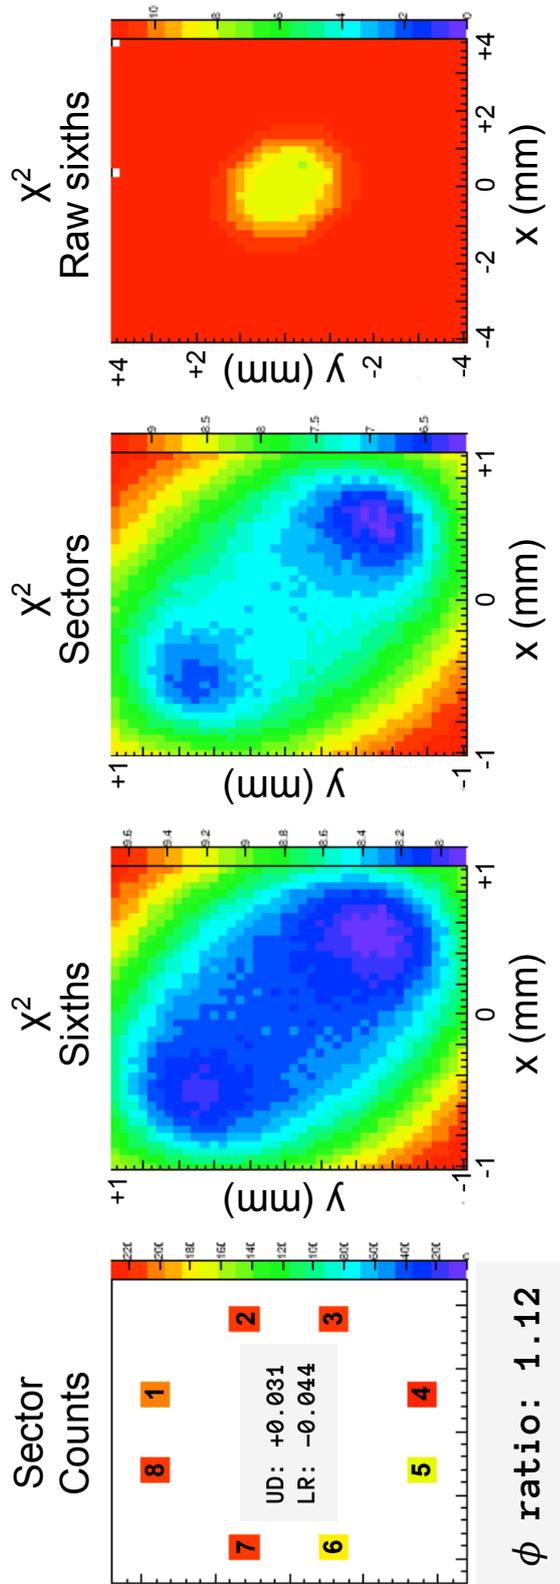


Figure 6.29.: The alignment tab of the KTAG online monitor

Conclusions

The nature of neutrino mass in the Standard Model is an open question and real insight could come from the discovery of heavy neutrinos which could explain the small SM neutrino mass via the seesaw mechanism. In this thesis, the 2007 data set from the NA62- R_K experiment has been used to set limits on the branching ratio $\mathcal{B}(K^+ \rightarrow \mu^+ \nu_h)$ for heavy neutrinos with masses in the range $275 \text{ MeV}/c^2 < m_{\nu_h}^2 < 375 \text{ MeV}/c^2$. The total number of kaon decays (measured with Standard Model $K^+ \rightarrow \mu^+ \nu_\mu$ decays) corresponded to 6×10^7 kaon decays in the fiducial region, and the acceptance for heavy neutrinos varied between 5% and 25%, depending on the mass considered. Backgrounds from decays to charged pions were reduced using the muon veto to require a muon in the final state, and backgrounds with photons in the final state were reduced by vetoing photons using the LKr electromagnetic calorimeter. The final measurement was limited by the background from $K^+ \rightarrow \mu^+ \pi^0 \nu_\mu$ decays in which the photons from π^0 decay escaped detection, and by beam halo muons which could be rejected only by kinematic cuts. The final limits set on the branching ratio of $K^+ \rightarrow \mu^+ \nu_h$ were between 3×10^{-5} and 3×10^{-6} , corresponding to limits on the neutrino mixing matrix element $|U_{\mu h}|^2$ between 1×10^{-6} and 2×10^{-5} . For neutrino masses above $330 \text{ MeV}/c^2$, these represent the strongest experimental limits on $|U_{\mu h}|^2$ from peak searches.

The NA62 experiment will bring precision to the ultra-rare kaon decay sector by measuring the branching ratio of the decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ with 10% uncertainty, helping to constrain the CKM matrix and further our understanding of CP violation in flavour physics. Kaon identification is an essential component in suppressing the backgrounds to the signal mode, and my contribution to the development of the KTAG detector has been outlined, focussing on the procedures for aligning the detector to maximize kaon detection efficiency. Two angular alignment methods have been considered

and shown to produce results consistent to within 7×10^{-5} rad. The results have been used to design the alignment panel of the NA62 Online Monitor for the KTAG detector.

The NA62 experiment has a wide-ranging physics programme, with analysis possibilities including searches for lepton flavour violation and exotic new particles, as well as precision studies of lepton universality and chiral perturbation theory. The limits obtained in the search for heavy neutrinos described in this thesis could be significantly improved with the new detector, taking advantage of the Gigatracker and KTAG detector to suppress the beam halo and the hermetic photon veto to suppress the background from $K^+ \rightarrow \mu^+ \pi^0 \nu_\mu$.

Bibliography

- [1] M. Srednicki. *Quantum Field Theory*. Cambridge University Press, 2007.
- [2] J. D. Wells. *Effective theories in physics*. Springer Briefs in physics. Heidelberg, Germany: Springer, 2012. DOI: 10.1007/978-3-642-34892-1.
- [3] S. L. Glashow, J. Iliopoulos, and L. Maiani. “Weak Interactions with Lepton-Hadron Symmetry”. In: *Phys. Rev. D* 2 (7 1970), pp. 1285–1292. DOI: 10.1103/PhysRevD.2.1285.
- [4] W. N. Cottingham and D. A. Greenwood. *An Introduction to the Standard Model of Particle Physics*. Cambridge University Press, 2007.
- [5] J. H. Christenson et al. “Evidence for the 2π Decay of the K_2^0 Meson”. In: *Phys. Rev. Lett.* 13 (4 1964), pp. 138–140. DOI: 10.1103/PhysRevLett.13.138.
- [6] H. Albrecht et al. “Observation of B_0 - B_0 mixing”. In: *Physics Letters B* 192.1–2 (1987), pp. 245–252. DOI: [http://dx.doi.org/10.1016/0370-2693\(87\)91177-4](http://dx.doi.org/10.1016/0370-2693(87)91177-4).
- [7] M. S. Sozzi. *Discrete symmetries and CP violation: From experiment to theory*. 2008.
- [8] K. Olive et al. “Review of Particle Physics”. In: *Chin.Phys.* C38 (2014), p. 090001. DOI: 10.1088/1674-1137/38/9/090001.
- [9] C. Jarlskog. “Commutator of the Quark Mass Matrices in the Standard Electroweak Model and a Measure of Maximal CP Nonconservation”. In: *Phys. Rev. Lett.* 55 (10 1985), pp. 1039–1042. DOI: 10.1103/PhysRevLett.55.1039.
- [10] J. Charles et al. “CP violation and the CKM matrix: assessing the impact of the asymmetric B factories”. English. In: *The European Physical Journal C - Particles and Fields* 41.1 (2005), pp. 1–131. DOI: 10.1140/epjc/s2005-02169-1.
- [11] A. Buras, S. Uhlig, and F. Schwab. “Waiting for precise measurements of $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ and $K_L \rightarrow \pi^0 \nu \bar{\nu}$ ”. In: *Rev. Mod. Phys.* 80 (3 2008), pp. 965–1007. DOI: 10.1103/RevModPhys.80.965.
- [12] V. Cirigliano et al. “Kaon decays in the standard model”. In: *Rev. Mod. Phys.* 84 (1 2012), pp. 399–447. DOI: 10.1103/RevModPhys.84.399.
- [13] D. Rein and L. M. Sehgal. “Long-distance contributions to the decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ ”. In: *Phys. Rev. D* 39 (11 1989), pp. 3325–3329. DOI: 10.1103/PhysRevD.39.3325.
- [14] J. S. Hagelin and L. S. Littenberg. “Rare kaon decays”. In: *Progress in Particle and Nuclear Physics* 23 (1989), pp. 1–40. DOI: [http://dx.doi.org/10.1016/0146-6410\(89\)90007-0](http://dx.doi.org/10.1016/0146-6410(89)90007-0).
- [15] M. Lu and M. B. Wise. “Long distance contribution to $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ ”. In: *Physics Letters B* 324.3–4 (1994), pp. 461–464. DOI: [http://dx.doi.org/10.1016/0370-2693\(94\)90223-2](http://dx.doi.org/10.1016/0370-2693(94)90223-2).
- [16] G. Buchalla and A. J. Buras. “The rare decays $K^+ \rightarrow \pi^+ \bar{\nu} \nu$ and $K_L \rightarrow \mu^+ \mu^-$ beyond leading logarithms”. In: *Nuclear Physics B* 412.1–2 (1994), pp. 106–142. DOI: [http://dx.doi.org/10.1016/0550-3213\(94\)90496-0](http://dx.doi.org/10.1016/0550-3213(94)90496-0).

- [17] G. Buchalla and A. J. Buras. “The rare decays $K^+ \rightarrow \pi^+ \bar{\nu} \nu$, $B \rightarrow X \nu \bar{\nu}$ and $B \rightarrow l^+ l^-$: an update”. In: *Nuclear Physics B* 548.1–3 (1999), pp. 309–327. DOI: [http://dx.doi.org/10.1016/S0550-3213\(99\)00149-2](http://dx.doi.org/10.1016/S0550-3213(99)00149-2).
- [18] J. Brod, M. Gorbahn, and E. Stamou. “Two-loop electroweak corrections for the $K \rightarrow \pi \nu \bar{\nu}$ decays”. In: *Phys. Rev. D* 83 (3 2011), p. 034030. DOI: 10.1103/PhysRevD.83.034030.
- [19] A. J. Buras et al. “Charm quark contribution to $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ at next-to-next-to-leading order”. In: *Journal of High Energy Physics* 2006.11 (2006), p. 002.
- [20] J. Brod and M. Gorbahn. “Electroweak corrections to the charm quark contribution to $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ ”. In: *Phys. Rev. D* 78 (3 2008), p. 034006. DOI: 10.1103/PhysRevD.78.034006.
- [21] G. Isidori, F. Mescia, and C. Smith. “Light-quark loops in $K^+ \rightarrow \pi^+ \bar{\nu} \nu$ ”. In: *Nuclear Physics B* 718.1–2 (2005), pp. 319–338. DOI: <http://dx.doi.org/10.1016/j.nuclphysb.2005.04.008>.
- [22] F. Mescia and C. Smith. “Improved estimates of rare K decay matrix elements from $K_{\ell 3}$ decays”. In: *Phys. Rev. D* 76 (3 2007), p. 034017. DOI: 10.1103/PhysRevD.76.034017.
- [23] J. Bijnens and K. Ghorbani. “Isospin breaking in $K\pi$ vector form-factors for the weak and rare decays K_{13} , $K \rightarrow \pi \nu \bar{\nu}$ and $K \rightarrow \pi l^+ l^-$ ”. In: (2007). arXiv: 0711.0148 [hep-ph].
- [24] G. Buchalla and A. J. Buras. “ $\sin 2\beta$ from $K \rightarrow \pi \nu \bar{\nu}$ ”. In: *Physics Letters B* 333.1–2 (1994), pp. 221–227. DOI: [http://dx.doi.org/10.1016/0370-2693\(94\)91034-0](http://dx.doi.org/10.1016/0370-2693(94)91034-0).
- [25] R. Aaij et al. “Measurement of CP Violation in $B^0 \rightarrow J/\psi K_S^0$ Decays”. In: *Phys. Rev. Lett.* 115 (3 2015), p. 031601. DOI: 10.1103/PhysRevLett.115.031601.
- [26] A. Artamonov et al. “New Measurement of the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ Branching Ratio”. In: *Phys. Rev. Lett.* 101 (19 2008), p. 191802. DOI: 10.1103/PhysRevLett.101.191802.
- [27] J. K. Ahn et al. “Experimental study of the decay $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$ ”. In: *Phys. Rev. D* 81 (7 2010), p. 072004. DOI: 10.1103/PhysRevD.81.072004.
- [28] G. Anelli et al. “Proposal to measure the rare decay $K^+ \rightarrow \pi^+ \bar{\nu} n \bar{u}$ at the CERN SPS”. In: (2005).
- [29] M. Togawa. “Status of KOTO Experiment at J-PARC”. In: *Journal of Physics: Conference Series* 556.1 (2014), p. 012046.
- [30] J. Charles et al. “CP violation and the CKM matrix: assessing the impact of the asymmetric B factories”. English. In: *The European Physical Journal C - Particles and Fields* 41.1 (2005), pp. 1–131. DOI: 10.1140/epjc/s2005-02169-1.
- [31] W. Pauli. “Pauli letter collection: letter to Lise Meitner”. Typed copy. 1930.
- [32] C. Weinheimer et al. “High precision measurement of the tritium β spectrum near its endpoint and upper limit on the neutrino mass”. In: *Physics Letters B* 460.1–2 (1999), pp. 219–226. DOI: [http://dx.doi.org/10.1016/S0370-2693\(99\)00780-7](http://dx.doi.org/10.1016/S0370-2693(99)00780-7).
- [33] V. N. Aseev et al. “Upper limit on the electron antineutrino mass from the Troitsk experiment”. In: *Phys. Rev. D* 84 (11 2011), p. 112003. DOI: 10.1103/PhysRevD.84.112003.
- [34] A. Osipowicz et al. “KATRIN: A Next generation tritium beta decay experiment with sub-eV sensitivity for the electron neutrino mass. Letter of intent”. In: (2001). arXiv: hep-ex/0109033 [hep-ex].

- [35] K. Valerius. “Status of the Karlsruhe Tritium Neutrino Experiment KATRIN”. In: *Proceedings, 20th International Conference on Particles and Nuclei (PANIC 14)*. 2014, pp. 308–312. DOI: 10.3204/DESY-PROC-2014-04/80.
- [36] G. Bellini et al. “Neutrino oscillations”. In: *Adv.High Energy Phys.* 2014 (2014), p. 191960. DOI: 10.1155/2014/191960. arXiv: 1310.7858 [hep-ph].
- [37] B. T. Cleveland et al. “Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector”. In: *The Astrophysical Journal* 496.1 (1998), p. 505.
- [38] J. N. Bahcall, N. A. Bahcall, and G. Shaviv. “Present Status of the Theoretical Predictions for the ^{37}Cl Solar-Neutrino Experiment”. In: *Phys. Rev. Lett.* 20 (21 1968), pp. 1209–1212. DOI: 10.1103/PhysRevLett.20.1209.
- [39] V. Gribov and B. Pontecorvo. “Neutrino astronomy and lepton charge”. In: *Physics Letters B* 28.7 (1969), pp. 493–496. DOI: [http://dx.doi.org/10.1016/0370-2693\(69\)90525-5](http://dx.doi.org/10.1016/0370-2693(69)90525-5).
- [40] S. Fukuda et al. “Solar ^8B and hep Neutrino Measurements from 1258 Days of Super-Kamiokande Data”. In: *Phys. Rev. Lett.* 86 (25 2001), pp. 5651–5655. DOI: 10.1103/PhysRevLett.86.5651.
- [41] Q. R. Ahmad et al. “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 89 (1 2002), p. 011301. DOI: 10.1103/PhysRevLett.89.011301.
- [42] X. Qian and W. Wang. “Reactor neutrino experiments: θ_{13} and beyond”. In: *Modern Physics Letters A* 29.16 (2014), p. 1430016. DOI: 10.1142/S021773231430016X. eprint: <http://www.worldscientific.com/doi/pdf/10.1142/S021773231430016X>.
- [43] G. J. Feldman, J. Hartnell, and T. Kobayashi. “Long-baseline neutrino oscillation experiments”. In: *Adv. High Energy Phys.* 2013 (2013), p. 475749. DOI: 10.1155/2013/475749. arXiv: 1210.1778 [hep-ex].
- [44] L. Wolfenstein. “Neutrino oscillations in matter”. In: *Phys. Rev. D* 17 (9 1978), pp. 2369–2374. DOI: 10.1103/PhysRevD.17.2369.
- [45] F. Capozzi et al. “Status of three-neutrino oscillation parameters, circa 2013”. In: *Phys. Rev. D* 89 (9 2014), p. 093018. DOI: 10.1103/PhysRevD.89.093018.
- [46] J. Kopp et al. “Sterile neutrino oscillations: the global picture”. English. In: *Journal of High Energy Physics* 2013.5, 50 (2013). DOI: 10.1007/JHEP05(2013)050.
- [47] R. N. Mohapatra and P. B. Pal. “Massive neutrinos in physics and astrophysics. Second edition”. In: *World Sci. Lect. Notes Phys.* 60 (1998). [World Sci. Lect. Notes Phys.72,1(2004)], pp. 1–397.
- [48] S. Petcov. “The Nature of Massive Neutrinos”. In: *Adv.High Energy Phys.* 2013 (2013), p. 852987. DOI: 10.1155/2013/852987. arXiv: 1303.5819 [hep-ph].
- [49] S. M. Bilenky and S. T. Petcov. “Massive neutrinos and neutrino oscillations”. In: *Rev. Mod. Phys.* 59 (3 1987), pp. 671–754. DOI: 10.1103/RevModPhys.59.671.
- [50] P. Minkowski. “ $\mu \rightarrow e\gamma$ at a rate of one out of 10⁹ muon decays?” In: *Physics Letters B* 67.4 (1977), pp. 421–428. DOI: [http://dx.doi.org/10.1016/0370-2693\(77\)90435-X](http://dx.doi.org/10.1016/0370-2693(77)90435-X).
- [51] R. N. Mohapatra et al. “Theory of neutrinos: A White paper”. In: *Rept. Prog. Phys.* 70 (2007), pp. 1757–1867. DOI: 10.1088/0034-4885/70/11/R02. arXiv: hep-ph/0510213 [hep-ph].

- [52] T. Asaka, S. Blanchet, and M. Shaposhnikov. “The ν MSM, dark matter and neutrino masses”. In: *Physics Letters B* 631.4 (2005), pp. 151–156. DOI: <http://dx.doi.org/10.1016/j.physletb.2005.09.070>.
- [53] T. Asaka and M. Shaposhnikov. “The ν MSM, dark matter and baryon asymmetry of the universe”. In: *Physics Letters B* 620.1–2 (2005), pp. 17–26. DOI: <http://dx.doi.org/10.1016/j.physletb.2005.06.020>.
- [54] D. Gorbunov and M. Shaposhnikov. “How to find neutral leptons of the $\hat{\nu}$ MSM?” In: *Journal of High Energy Physics* 2007.10 (2007), p. 015.
- [55] R. Shrock. “New tests for and bounds on neutrino masses and lepton mixing”. In: *Physics Letters B* 96.1–2 (1980), pp. 159–164. DOI: [http://dx.doi.org/10.1016/0370-2693\(80\)90235-X](http://dx.doi.org/10.1016/0370-2693(80)90235-X).
- [56] A. Atre et al. “The search for heavy Majorana neutrinos”. In: *Journal of High Energy Physics* 2009.05 (2009), p. 030.
- [57] R. Aaij et al. “Search for Majorana Neutrinos in $B^- \rightarrow \pi^+ \mu^- \mu^-$ Decays”. In: *Phys. Rev. Lett.* 112 (13 2014), p. 131802. DOI: [10.1103/PhysRevLett.112.131802](https://doi.org/10.1103/PhysRevLett.112.131802).
- [58] D. I. Britton et al. “Measurement of the $\pi^+ \rightarrow e^+ \nu$ branching ratio”. In: *Phys. Rev. Lett.* 68 (20 1992), pp. 3000–3003. DOI: [10.1103/PhysRevLett.68.3000](https://doi.org/10.1103/PhysRevLett.68.3000).
- [59] T. Yamazaki et al. “Search for Heavy Neutrinos in Kaon Decay”. In: (1984). [Conf. Proc.C840719,262(1984)] p. I.262.
- [60] G. Bernardi et al. “Further limits on heavy neutrino couplings”. In: *Physics Letters B* 203.3 (1988), pp. 332–334. DOI: [http://dx.doi.org/10.1016/0370-2693\(88\)90563-1](http://dx.doi.org/10.1016/0370-2693(88)90563-1).
- [61] J. Dorenbosch et al. “A search for decays of heavy neutrinos in the mass range 0.5–2.8 GeV”. In: *Physics Letters B* 166.4 (1986), pp. 473–478. DOI: [http://dx.doi.org/10.1016/0370-2693\(86\)91601-1](http://dx.doi.org/10.1016/0370-2693(86)91601-1).
- [62] J. Badier et al. “Mass and lifetime limits on new long-lived particles in 300 GeV/c π interactions”. English. In: *Zeitschrift für Physik C Particles and Fields* 31.1 (1986), pp. 21–32. DOI: [10.1007/BF01559588](https://doi.org/10.1007/BF01559588).
- [63] O. Adriani et al. “Search for isosinglet neutral heavy leptons in Z^0 decays”. In: *Physics Letters B* 295.3 (1992), pp. 371–382. DOI: [http://dx.doi.org/10.1016/0370-2693\(92\)91579-X](http://dx.doi.org/10.1016/0370-2693(92)91579-X).
- [64] “Search for neutral heavy leptons produced in Z decays”. English. In: *Zeitschrift für Physik C Particles and Fields* 74.1 (1997), pp. 57–71. DOI: [10.1007/s002880050370](https://doi.org/10.1007/s002880050370).
- [65] R. S. Hayano et al. “Heavy-Neutrino Search Using $K_{\mu 2}$ Decay”. In: *Phys. Rev. Lett.* 49 (18 1982), pp. 1305–1309. DOI: [10.1103/PhysRevLett.49.1305](https://doi.org/10.1103/PhysRevLett.49.1305).
- [66] A. Vaitaitis et al. “Search for Neutral Heavy Leptons in a High-Energy Neutrino Beam”. In: *Phys. Rev. Lett.* 83 (24 1999), pp. 4943–4946. DOI: [10.1103/PhysRevLett.83.4943](https://doi.org/10.1103/PhysRevLett.83.4943).
- [67] A. Cooper-Sarkar et al. “Search for heavy neutrino decays in the BEBC beam dump experiment”. In: *Physics Letters B* 160.1 (1985), pp. 207–211. DOI: [http://dx.doi.org/10.1016/0370-2693\(85\)91493-5](http://dx.doi.org/10.1016/0370-2693(85)91493-5).
- [68] E. Gallas et al. “Search for neutral weakly interacting massive particles in the Fermilab Tevatron wideband neutrino beam”. In: *Phys. Rev. D* 52 (1 1995), pp. 6–14. DOI: [10.1103/PhysRevD.52.6](https://doi.org/10.1103/PhysRevD.52.6).

- [69] P. Vilain et al. “Search for heavy isosinglet neutrinos”. In: *Physics Letters B* 343.1–4 (1995), pp. 453–458. DOI: [http://dx.doi.org/10.1016/0370-2693\(94\)01422-9](http://dx.doi.org/10.1016/0370-2693(94)01422-9).
- [70] A. V. Artamonov et al. “Search for heavy neutrinos in $K^+ \rightarrow \mu^+ \nu_H$ decays”. In: *Phys. Rev. D* 91 (5 2015), p. 052001. DOI: 10.1103/PhysRevD.91.052001.
- [71] H. Burkhardt et al. “The beam and detector for a high-precision measurement of {CP} violation in neutral-kaon decays”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 268.1 (1988), pp. 116–143. DOI: [http://dx.doi.org/10.1016/0168-9002\(88\)90598-0](http://dx.doi.org/10.1016/0168-9002(88)90598-0).
- [72] G. Barr et al. “A new measurement of direct {CP} violation in the neutral kaon system”. In: *Physics Letters B* 317.1–2 (1993), pp. 233–242. DOI: [http://dx.doi.org/10.1016/0370-2693\(93\)91599-I](http://dx.doi.org/10.1016/0370-2693(93)91599-I).
- [73] T. Nakada. “Review on CP violation”. In: *AIP Conference Proceedings* 302.1 (1994), pp. 425–463. DOI: <http://dx.doi.org/10.1063/1.45461>.
- [74] G. D. Barr, R. Buchholz, and R. Carosi. *Proposal for a precision measurement of ϵ'/ϵ in CP violating $K^0 \rightarrow \pi$ decays*. Tech. rep. CERN-SPSC-90-022. SPSC-P-253. Geneva: CERN, 1990.
- [75] A. Lai et al. “A precise measurement of the direct CP violation parameter $\text{Re}(\epsilon'/\epsilon)$ ”. English. In: *The European Physical Journal C - Particles and Fields* 22.2 (2001), pp. 231–254. DOI: 10.1007/s100520100822.
- [76] J. Batley et al. “A precision measurement of direct CP violation in the decay of neutral kaons into two pions”. In: *Physics Letters B* 544.1–2 (2002), pp. 97–112. DOI: [http://dx.doi.org/10.1016/S0370-2693\(02\)02476-0](http://dx.doi.org/10.1016/S0370-2693(02)02476-0).
- [77] R. Batley et al. “A high sensitivity investigation of K(S) and neutral hyperon decays using a modified K(S) beam. (Addendum 2 to P253)”. In: (1999).
- [78] A. Lai et al. “Precise measurements of the $K_S \rightarrow \gamma\gamma$ and $K_L \rightarrow \gamma\gamma$ decay rates”. In: *Physics Letters B* 551.1–2 (2003), pp. 7–15. DOI: [http://dx.doi.org/10.1016/S0370-2693\(02\)02955-6](http://dx.doi.org/10.1016/S0370-2693(02)02955-6).
- [79] J. Batley et al. “Observation of the rare decay $K_S \rightarrow \pi^0 \mu^+ \mu^-$ ”. In: *Physics Letters B* 599.3–4 (2004), pp. 197–211. DOI: <http://dx.doi.org/10.1016/j.physletb.2004.08.058>.
- [80] J. Batley et al. “Measurement of the branching ratio of the decay $\Xi^0 \rightarrow \Sigma^+ \mu^- \bar{\nu}_\mu$ ”. In: *Physics Letters B* 720.1–3 (2013), pp. 105–110. DOI: <http://dx.doi.org/10.1016/j.physletb.2013.01.023>.
- [81] J. R. Batley et al. *Addendum 3 (to proposal P253/CERN/SPSC): for a precision measurement of charged kaon decay parameters with an extended NA48 setup*. Tech. rep. CERN-SPSC-2000-003. SPSC-P-253-Add-3. Geneva: CERN, 1999.
- [82] J. Batley et al. “Search for direct CP violating charge asymmetries in $K^\pm \rightarrow \pi^\pm \pi^+ \pi^-$ and $K^\pm \rightarrow \pi^\pm \pi^0 \pi^0$ decays”. English. In: *The European Physical Journal C* 52.4 (2007), pp. 875–891. DOI: 10.1140/epjc/s10052-007-0456-7.
- [83] A. Masiero, P. Paradisi, and R. Petronzio. “Probing new physics through μ - e universality in $K \rightarrow \ell \nu$ ”. In: *Phys. Rev. D* 74 (1 2006), p. 011701. DOI: 10.1103/PhysRevD.74.011701.
- [84] M. Finkemeier. “Radiative corrections to $\pi l 2$ and $K l 2$ decays”. In: *Physics Letters B* 387.2 (1996), pp. 391–394. DOI: [http://dx.doi.org/10.1016/0370-2693\(96\)01030-1](http://dx.doi.org/10.1016/0370-2693(96)01030-1).

- [85] M. Sozzi et al. *Possible Search for New Physics by Measuring the Ratio $\Gamma(K \rightarrow e\nu)/\Gamma(K \rightarrow \mu\nu)$ with the Na48/2 Apparatus*. Tech. rep. CERN-SPSC-2006-015. SPSC-M-745. Geneva: CERN, 2006.
- [86] C. Lazzeroni et al. “Precision Measurement of the Ratio of the Charged Kaon Leptonic Decay Rates”. In: *Phys.Lett.* B719 (2013), pp. 326–336. DOI: 10.1016/j.physletb.2013.01.037. arXiv: 1212.4012 [hep-ex].
- [87] V. Fanti et al. “The beam and detector for the {NA48} neutral kaon {CP} violation experiment at {CERN}”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 574.3 (2007), pp. 433–471. DOI: <http://dx.doi.org/10.1016/j.nima.2007.01.178>.
- [88] B. Peyaud. “KABES: a novel beam spectrometer for {NA48}”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 535.1–2 (2004). Proceedings of the 10th International Vienna Conference on Instrumentation, pp. 247–252. DOI: <http://dx.doi.org/10.1016/j.nima.2004.07.275>.
- [89] C Biino et al. *The simultaneous long- and short-lived neutral kaon beams for experiment NA48*. Tech. rep. CERN-SL-98-033-EA. CERN-NA-48-NOTE-98-16. Geneva: CERN, 1998.
- [90] E. Griesmayer et al. “Comparison of field calculations and measurements of a spectrometer magnet”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 361.3 (1995), pp. 466–471. DOI: [http://dx.doi.org/10.1016/0168-9002\(95\)00142-5](http://dx.doi.org/10.1016/0168-9002(95)00142-5).
- [91] T. W. Robert DeWolf Roger Moore. *Description of the Mu-Veto Counters*. Internal Note NA48-94-1. Cambridge: NA48, 1994.
- [92] C. Lazzeroni et al. “Test of Lepton Flavour Universality in $K^+ \rightarrow l^+ \nu$ Decays”. In: *Phys.Lett.* B698 (2011), pp. 105–114. DOI: 10.1016/j.physletb.2011.02.064. arXiv: 1101.4805 [hep-ex].
- [93] J. Bijnens, G. Ecker, and J. Gasser. “Radiative semileptonic kaon decays”. In: *Nuclear Physics B* 396.1 (1993), pp. 81–118. DOI: [http://dx.doi.org/10.1016/0550-3213\(93\)90259-R](http://dx.doi.org/10.1016/0550-3213(93)90259-R).
- [94] C.-H. Chen, C.-Q. Geng, and C.-C. Lih. “Analysis of $K^+ \rightarrow e^+ \nu_e \gamma$ in light-front quark model and chiral perturbation theory of order p^6 ”. In: *Phys. Rev. D* 77 (1 2008), p. 014004. DOI: 10.1103/PhysRevD.77.014004.
- [95] M De Beer. *NASIM User’s Guide*. Internal Note NA48-00-23. NA48, 2000.
- [96] R. Brun, F. Carminati, and S. Giani. “GEANT Detector Description and Simulation Tool”. In: (1994).
- [97] E. M. Marinova. “Investigation of charged kaon semileptonic decays”. PhD thesis. University of Sofia, Faculty of Physics, 2009.
- [98] F. C. Iselin. *HALO: a computer program to calculate muon halo*. Geneva: CERN, 1974.
- [99] J. Neyman. “Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability”. English. In: *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* 236.767 (1937), pp. 333–380.
- [100] G. J. Feldman and R. D. Cousins. “Unified approach to the classical statistical analysis of small signals”. In: *Phys. Rev. D* 57 (7 1998), pp. 3873–3889. DOI: 10.1103/PhysRevD.57.3873.

- [101] W. A. Rolke, A. M. López, and J. Conrad. “Limits and confidence intervals in the presence of nuisance parameters”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 551.2–3 (2005), pp. 493–503. DOI: <http://dx.doi.org/10.1016/j.nima.2005.05.068>.
- [102] C. J. CLOPPER and E. S. PEARSON. “THE USE OF CONFIDENCE OR FIDUCIAL LIMITS ILLUSTRATED IN THE CASE OF THE BINOMIAL”. In: *Biometrika* 26.4 (1934), pp. 404–413. DOI: 10.1093/biomet/26.4.404. eprint: <http://biomet.oxfordjournals.org/content/26/4/404.full.pdf+html>.
- [103] U. Camerini et al. “Experimental Search for Semileptonic Neutrino Neutral Currents”. In: *Phys. Rev. Lett.* 23 (6 1969), pp. 326–329. DOI: 10.1103/PhysRevLett.23.326.
- [104] D. Ljung and D. Cline. “Experimental Study of the Rare K^+ Decay Modes: $K^+ \rightarrow \pi^+ \pi^0 \gamma$, $K^+ \rightarrow \mu^+ \pi^0 \nu \gamma$, $K^+ \rightarrow \pi^+ \gamma \gamma$, $K^+ \rightarrow \pi^+ \nu \bar{\nu}$, $K^+ \rightarrow \pi^0 \pi^0 e^+ \nu$, and $K^+ \rightarrow e^+ \pi^0 \nu \gamma$ ”. In: *Phys. Rev. D* 8 (5 1973), pp. 1307–1330. DOI: 10.1103/PhysRevD.8.1307.
- [105] J. Klems, R. Hildebrand, and R. Stiening. “Limits on the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ and $K^+ \rightarrow \pi^+(n)\gamma$ decay rates”. In: *Phys.Rev.* D4 (1971), pp. 66–80. DOI: 10.1103/PhysRevD.4.66.
- [106] G. Cable et al. “Search for rare K^+ decays. ii. $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ ”. In: *Phys.Rev.* D8 (1973), pp. 3807–3812. DOI: 10.1103/PhysRevD.8.3807.
- [107] Y. Asano et al. “Search for a Rare Decay Mode $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ and Axion”. In: *Phys.Lett.* B107 (1981), p. 159. DOI: 10.1016/0370-2693(81)91172-2.
- [108] M. Atiya et al. “A detector to search for $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ ”. In: *Nucl.Instrum.Meth.* A321 (1992), pp. 129–151. DOI: 10.1016/0168-9002(92)90382-E.
- [109] Bassalleck et al. *An experiment to measure the branching ratio $B(K^+ \rightarrow \pi^+ \nu \bar{\nu})$* . Tech. rep. BNL-67247 tri-pp-00-06. Brookhaven National Laborator, 1999.
- [110] S. Adler et al. “Measurement of the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ branching ratio”. In: *Phys. Rev. D* 77 (5 2008), p. 052003. DOI: 10.1103/PhysRevD.77.052003.
- [111] J. Cresswell et al. “A Cylindrical Drift Chamber for the Measurement of $K \rightarrow \pi$ Neutrino Anti-neutrino Decay”. In: *IEEE Trans.Nucl.Sci.* 35 (1988), pp. 460–463. DOI: 10.1109/23.12765.
- [112] S. Adler et al. “Search for the Decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ ”. In: *Phys. Rev. Lett.* 76 (9 1996), pp. 1421–1424. DOI: 10.1103/PhysRevLett.76.1421.
- [113] M. Atiya et al. “Search for the decays $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ and $K^+ \rightarrow \pi^+ X^0$ for $150 < M_{X^0} < 250 \text{ MeV}/c^2$ ”. In: *Phys. Rev. D* 48 (1 1993), R1–R4. DOI: 10.1103/PhysRevD.48.R1.
- [114] S. H. Kettell. “A Search for the rare decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ ”. In: (1996). arXiv: hep-ex/9701003 [hep-ex].
- [115] S. Adler et al. “Evidence for the Decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ ”. In: *Phys. Rev. Lett.* 79 (12 1997), pp. 2204–2207. DOI: 10.1103/PhysRevLett.79.2204.
- [116] S. Adler et al. “Further Evidence for the Decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ ”. In: *Phys. Rev. Lett.* 88 (4 2002), p. 041803. DOI: 10.1103/PhysRevLett.88.041803.
- [117] S Adler et al. “Search for the decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ in the momentum region $p_\pi < 195 \text{ MeV}/c$ ”. In: *Physics Letters B* 537.3–4 (2002), pp. 211–216. DOI: [http://dx.doi.org/10.1016/S0370-2693\(02\)01911-1](http://dx.doi.org/10.1016/S0370-2693(02)01911-1).
- [118] F Hahn et al. *NA62: Technical Design Document*. Tech. rep. NA62-10-07. Geneva: CERN, 2010.

- [119] H. Atherton et al. “Precise Measurements of Particle Production by 400-GeV/c Protons on Beryllium Targets”. In: (1980).
- [120] M. Fiorini et al. “High rate particle tracking and ultra-fast timing with a thin hybrid silicon pixel detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 718.0 (2013). Proceedings of the 12th Pisa Meeting on Advanced Detectors La Biodola, Isola d’Elba, Italy, May 20 – 26, 2012, pp. 270 –273. DOI: <http://dx.doi.org/10.1016/j.nima.2012.10.108>.
- [121] G. Romagnoli et al. “Silicon micro-fluidic cooling for {NA62} {GTK} pixel detectors”. In: *Microelectronic Engineering* 145 (2015). Micro/Nano Devices and Systems 2014 An open focused special thematic issue of Microelectronic Engineering, pp. 133 –137. DOI: <http://dx.doi.org/10.1016/j.mee.2015.04.006>.
- [122] F. Ambrosino et al. “The CHarged ANTIcounter for the NA62 Experiment at CERN”. In: *Physics Procedia* 37.0 (2012). Proceedings of the 2nd International Conference on Technology and Instrumentation in Particle Physics (TIPP 2011), pp. 675 –682. DOI: <http://dx.doi.org/10.1016/j.phpro.2012.04.095>.
- [123] H. Danielsson. “Design, construction and testing of the straw tracker for the NA62 experiment”. In: (2013). DOI: 10.1109/NSSMIC.2013.6829462.
- [124] M. Lenti. “Progress of the NA62 RICH detector”. In: *Journal of Instrumentation* 9.01 (2014), p. C01054.
- [125] B. Angelucci et al. “Pion–Muon separation with a {RICH} prototype for the {NA62} experiment”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 621.1–3 (2010), pp. 205 –211. DOI: <http://dx.doi.org/10.1016/j.nima.2010.05.062>.
- [126] F. Ambrosino et al. “The large-angle photon veto system for the NA62 experiment at CERN”. In: *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2011 IEEE*. 2011, pp. 1159–1166. DOI: 10.1109/NSSMIC.2011.6154594.
- [127] “The {OPAL} detector at {LEP}”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 305.2 (1991), pp. 275 –319. DOI: [http://dx.doi.org/10.1016/0168-9002\(91\)90547-4](http://dx.doi.org/10.1016/0168-9002(91)90547-4).
- [128] F. Ambrosino et al. “A prototype large-angle photon veto detector for the P326 experiment at CERN”. In: *Nuclear Science Symposium Conference Record, 2007. NSS ’07. IEEE*. Vol. 1. 2007, pp. 57–64. DOI: 10.1109/NSSMIC.2007.4436288.
- [129] A Ceccucci et al. “The NA62 liquid Krypton calorimeter’s new readout system”. In: *Journal of Instrumentation* 9.01 (2014), p. C01047.
- [130] A. Ceccucci. *NA62/P-326 Status Report*. Tech. rep. SPSC-M-760. CERN-SPSC-2007-035. revised version submitted on 2007-11-27 08:18:51. 2007.
- [131] T. N. P. V. W. Group. *Photon Vetoes - Upgrade to the Proposal*. Note 07-09. NA62, 2007.
- [132] M. Krivda. “The {NA62} trigger system”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 718 (2013). Proceedings of the 12th Pisa Meeting on Advanced Detectors La Biodola, Isola d’Elba, Italy, May 20 – 26, 2012, pp. 164 –167. DOI: <http://dx.doi.org/10.1016/j.nima.2012.11.178>.
- [133] M Pivanti et al. “Implementation of a PC-based Level 0 Trigger Processor for the NA62 Experiment”. In: *Journal of Physics: Conference Series* 513.1 (2014), p. 012008.

- [134] C. Bovet et al. *The CEDAR counters for particle identification in the SPS secondary beams: a description and an operation manual*. Geneva: CERN, 1982.
- [135] H Hovestadt. *Jena glass and its scientific and industrial applications*. Ed. by J. Everett and E. Alice. Macmillan and Co., Limited, 1902.
- [136] R. C. Fernow. *Introduction to Experimental Particle Physics*. Cambridge University Press, 1986.
- [137] E. Goudzovski et al. “Development of the kaon tagging system for the NA62 experiment at CERN”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (2015), pp. –. DOI: <http://dx.doi.org/10.1016/j.nima.2015.08.015>.
- [138] F. Anghinolfi et al. “NINO: an ultra-fast and low-power front-end amplifier/discriminator {ASIC} designed for the multigap resistive plate chamber”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 533.1–2 (2004). Proceedings of the Seventh International Workshop on Resistive Plate Chambers and Related Detectors, pp. 183–187. DOI: <http://dx.doi.org/10.1016/j.nima.2004.07.024>.
- [139] J Christiansen. *HPTDC High Performance Time to Digital Converter*. Tech. rep. Version 2.2 for HPTDC version 1.3. Geneva: CERN, 2004.
- [140] B Angelucci et al. “The FPGA based Trigger and Data Acquisition system for the CERN NA62 experiment”. In: *Journal of Instrumentation* 9.01 (2014), p. C01055.
- [141] G. Haefeli et al. “The {LHCb} {DAQ} interface board {TELL1}”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 560.2 (2006), pp. 494–502. DOI: <http://dx.doi.org/10.1016/j.nima.2005.12.212>.
- [142] A. Sergi. *NA62 Framework*.
- [143] K. Brown and F. Iselin. “Decay TURTLE (Trace Unlimited Rays Through Lumped Elements): A Computer Program for Simulating Charged Particle Beam Transport Systems, Including Decay Calculations”. In: (1974).
- [144] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. DOI: [http://dx.doi.org/10.1016/S0168-9002\(03\)01368-8](http://dx.doi.org/10.1016/S0168-9002(03)01368-8).
- [145] “Overview of Photon Counting Techniques”. English. In: *Advanced Time-Correlated Single Photon Counting Techniques*. Ed. by J. Castleman A.W., J. Toennies, and W. Zinth. Vol. 81. Springer Series in Chemical Physics. Springer Berlin Heidelberg, 2005, pp. 11–25. DOI: [10.1007/3-540-28882-1_2](https://doi.org/10.1007/3-540-28882-1_2).
- [146] H. PHOTONICS. *Photomultiplier Tube FINAL TEST SHEET JNO: 641286, 641287, 726283, 817164, 817165, 817166, 817167, 817168, 817169, 817170, 817171*). Manufacturers data sheet. 2012.
- [147] C. NA62. *2015 NA62 Status Report to the CERN SPSC*. Tech. rep. CERN-SPSC-2015-012. SPSC-SR-157. Geneva: CERN, 2015.
- [148] P. B. Pal. “Dirac, Majorana and Weyl fermions”. In: *Am. J. Phys.* 79 (2011), pp. 485–498. DOI: [10.1119/1.3549729](https://doi.org/10.1119/1.3549729). arXiv: 1006.1718 [hep-ph].

Appendix A

Notation

Where there is variation in notation or terminology in the literature, the following conventions are used.

Vectors

\vec{p} is three-vector. p is Lorentz four-vector. The Einstein summation convention and Minkowski metric are used

$$p^\mu q_\mu \equiv p^\mu q^\nu \eta_{\mu\nu} \equiv \sum_{\mu\nu} p^\mu q^\nu \eta_{\mu\nu}, \quad (\text{A.1})$$

where $\eta = \text{diag}(1, -1, -1, -1)$.

Pauli matrices

The Pauli matrices σ_i , appear in the Quantum Mechanical description of spin:

$$\sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (\text{A.2})$$

Spinors

The Dirac equation is an expression of relativistic energy conservation, $E^2 - |\vec{p}|^2 - m^2 = 0$, via quantum mechanical operators which are first order in time and space derivatives. Writing the four-momentum operator $p = (E, p_x, p_y, p_z)$, the equality

$$p^\mu p_\mu - m^2 = (\gamma^\kappa p_\kappa + m)(\gamma^\lambda p_\lambda - m) \quad (\text{A.3})$$

holds provided the four γ^μ satisfy the anti-commutation relation

$$\{\gamma^\mu, \gamma^\nu\} = 2g^{\mu\nu}. \quad (\text{A.4})$$

This relation can be accommodated if γ^μ are in fact 4×4 matrices. Equation A.4 does not uniquely define the γ^μ matrices, but one solution (the Dirac basis) is:

$$\gamma^0 = \begin{pmatrix} I_2 & 0 \\ 0 & -I_2 \end{pmatrix}, \quad \gamma^k = \begin{pmatrix} 0 & \sigma^k \\ -\sigma^k & 0 \end{pmatrix}, \quad (\text{A.5})$$

where I_2 is identity matrix and σ^k are the Pauli matrices. Using the QM expression for the momentum operator $p = i\partial_\mu$, the right-hand bracket of equation A.3 can be taken to act on a

wave-function, ψ , to produce the Dirac equation:

$$(i\gamma^\mu \partial_\mu - m)\psi = 0 \quad (\text{A.6})$$

Since γ^μ are matrices, the wavefunction ψ must have four-components: it is a bi-spinor (often just called a spinor). The interpretation of the 4 components of a bi-spinor depends on the chosen basis of γ matrices. In the Dirac basis, the first two components correspond to positive and negative helicity particle states, while the last two correspond to negative and positive helicity anti-particle states. In the Weyl basis, the first two components correspond to a right-chiral state, while the second two correspond to a left-chiral state. Helicity and chirality are discussed below.

Helicity and chirality

The *helicity* of a particle state is defined by the operator

$$h_{\vec{p}} = \frac{2\vec{J} \cdot \vec{p}}{|\vec{p}|} \quad (\text{A.7})$$

where \vec{p} and \vec{J} are the momentum and angular momentum of the particle. It has eigenvalues of ± 1 corresponding to ‘positive helicity’ and ‘negative helicity’ eigenstates.

The *chirality* operator is defined as product of the four gamma matrices $\gamma_5 = i\gamma^0\gamma^1\gamma^2\gamma^3$. If the chiral projection operators are defined

$$P_L = \frac{1}{2}(1 - \gamma_5), \quad P_R = \frac{1}{2}(1 + \gamma_5), \quad (\text{A.8})$$

then $\psi_L = P_L\psi$ is the ‘left-chiral’ projection of the state ψ , and ψ_R the ‘right-chiral’ projection. In general, helicity commutes with the Dirac Hamiltonian so it is a conserved quantity, however its value depends on the frame of reference since a boost can always be applied to reverse \vec{p} without affecting \vec{J} . Conversely, chirality does not commute with the Dirac Hamiltonian so it is not conserved, but it is Lorentz invariant. In the massless limit, helicity and chirality are equivalent to each other and are both conserved and Lorentz invariant. Then one can unambiguously describe a particle as ‘right-handed’, meaning it has positive helicity and is right-chiral.

Dirac adjoint

In order to construct Lorentz invariant quantities from Dirac spinors, the definition of an *adjoint spinor*, $\bar{\psi}$ is required.

$$\bar{\psi} = \psi^\dagger \gamma^0, \quad (\text{A.9})$$

where γ^0 is the first of the Dirac gamma matrices, satisfying $\{\gamma^\mu, \gamma^\nu\} = 2g^{\mu\nu}$. With this definition, the quantity $\bar{\psi}\psi$ is a Lorentz scalar.

Anti-particles

If a particle is described by the spinor, ψ , a spinor describing its antiparticle can be constructed from ψ^* . In order to have the required Lorentz transformations, one must construct the *Lorentz covariate conjugate*, $\widehat{\psi}$, defined by:

$$\widehat{\psi}(x) = \gamma_0 C \psi^*(x) = -C \bar{\psi}^T \quad (\text{A.10})$$

where C is an antisymmetric matrix whose components depend on the choice of basis for the Dirac matrices but will always satisfy $C^{-1}\gamma_\mu C = -\gamma_\mu^T$. Some authors use ψ^c to mean the same thing.

Charge conjugation

Charge conjugation replaces all fields by their complex conjugates, thus changing the sign of all quantum charges. In addition, the operation may introduce a complex phase such that for a fermion field, the C matrix is required again to ensure Lorentz invariance:

$$\mathcal{C}\psi(x)\mathcal{C}^{-1} = \eta_C \gamma_0 C \psi^*(x) = \eta_C \widehat{\psi}(x) \quad (\text{A.11})$$

where η_C is a complex phase. Some authors use ψ^c to mean the same thing.

Conjugation and chirality

Let χ be a right-chiral Dirac spinor field, meaning $P_L \chi = 0$ and $P_R \chi = \chi$. It can be shown that $P_L \widehat{\chi} = \widehat{\chi}$ and $P_R \widehat{\chi} = 0$, i.e the Lorentz covariate conjugate of a right-chiral field is left-chiral [148].

On the other hand, charge conjugation does not affect chirality. Continuing with the same definition of χ , we find $P_L(\mathcal{C}\chi\mathcal{C}) = 0$ and $P_R(\mathcal{C}\chi\mathcal{C}) = (\mathcal{C}\chi\mathcal{C})$, i.e. the charge conjugated spinor field of a right-chiral spinor field is still right-chiral.

In fact for a general Dirac spinor, ψ , with $\psi_L(x) = P_L \psi(x)$ and $\psi_R(x) = P_R \psi(x)$:

$$\mathcal{C}\psi_L\mathcal{C} = \eta_c \widehat{\psi}_R \quad (\text{A.12})$$

Since $P_L \widehat{\psi}_R = \widehat{\psi}_R = P_L \widehat{\psi}$, it is natural to define

$$P_L \widehat{\psi} \equiv \widehat{\psi}_L = \widehat{\psi}_R \quad (\text{A.13})$$

Appendix B

PMT Properties

There are several characteristics which drive the choice of PMT for a particular application. *Quantum efficiency*, is the probability that a photon incident on the photo-cathode will result in an electron reaching the first dynode. It is a function of the wavelength of the incident photon:

$$\eta_{QE}(\lambda) = \frac{N_e^{\text{dynode1}}(\lambda)}{N_\gamma^{\text{cathode}}(\lambda)} \quad (\text{B.1})$$

The same property can alternatively be expressed in terms of current and power as the *cathode radiant sensitivity* (units A/W):

$$S_C(\lambda) = \frac{I_C \lambda}{P(\lambda)} \quad (\text{B.2})$$

where $I_C(\lambda)$ is the current per unit wavelength at the first dynode and $P(\lambda)$ is the incident power on the cathode per unit wavelength. Since $I_C(\lambda)$ is determined by the number of electrons and $P(\lambda)$ is determined to the number of photons, quantum efficiency and cathode radiant sensitivity are related by:

$$\eta_{QE}(\lambda) = \frac{1240 \text{ nm}}{\lambda} S_C(\lambda). \quad (\text{B.3})$$

Rather than list the entire spectrum for every PMT, datasheets typically quote a single parameter, the *cathode luminous sensitivity* (Amps per Lumen) as a measure of PMT quantum efficiency:

$$S_C = \frac{I_C}{\phi_V} \quad (\text{B.4})$$

where I_C is the current at the first dynode:

$$I_C = \int_0^\infty S_C(\lambda) W(\lambda) d\lambda \quad (\text{B.5})$$

and ϕ_V is the luminous flux at the cathode:

$$\phi_V = 683 \frac{\text{lm}}{\text{W}} \int_0^\infty \bar{y}(\lambda) W(\lambda) d\lambda \quad (\text{B.6})$$

The standard source spectrum (i.e. the radiant flux), $W(\lambda)$ is that of a tungsten filament lamp

operated at a distribution temperature of 2856 K giving the Black Body spectrum:

$$W(\lambda) = \frac{W_0}{\lambda^5} \frac{1}{\exp\left(\frac{hc}{\lambda k_B T}\right) - 1} \quad (\text{B.7})$$

The standard eye distribution, \bar{y} peaks around 550 nm. The last two equations fix the definition of the lumen.

For PMTs which are not designed to detect light in the visible region, the cathode luminous sensitivity is not a very meaningful quantity since it places the most weight on the region where the PMT is not sensitive. Instead the *blue sensitivity* is often quoted. It is measured using the standard tungsten lamp with an extra blue filter (Corning CS No.5-58) whose transmission peaks around 430 nm so it is more appropriate for lower wavelength applications. The units are A/lmF, where F stands for filtered.

Another important quantity is the gain of the photomultiplier, the number of electrons seen at the anode for each photoelectron at the cathode. It can be expressed as:

$$G = (g\delta)^n \quad (\text{B.8})$$

where δ is the average secondary emission coefficient, g is the transfer efficiency of electrons between dynodes and n is the number of dynodes. The measurable quantity quoted on datasheets is the *anode luminous sensitivity*, defined as:

$$S_A = \frac{I_A}{P} \quad (\text{B.9})$$

which can be related to the gain by

$$G = \frac{S_A}{S_C} \quad (\text{B.10})$$