

A search for $\nu_\mu \rightarrow \nu_\tau$ oscillations using
the three-prong tau decay signature
in the NOMAD experiment

B.D. Yabsley



*A thesis
submitted for the degree of
Doctor of Philosophy
at the
University of Sydney*

August, 1999

Abstract

The NOMAD experiment is a short baseline ν_τ appearance search, using a kinematic method to identify the decay of tau leptons produced in $\nu_\tau + N \rightarrow \tau^- + X$ interactions, above a background due to neutral- and charged-current interactions of other neutrino species. A search for “three prong” tau decays in the NOMAD data is developed, based on a fundamental analysis of the kinematics of such decays, and the neutral current background.

It is found that the traditional “transverse plane” description of candidate tau decay events, when employed together with the isolation method more recently developed, introduces a pronounced and complicated correlation between the descriptive kinematic variables. Beginning with the isolation cut, a consistent set of kinematic variables is developed, describing candidate tau decays in a more transparent manner. After taking into account a kinematic effect of the invariant mass of the tau decay product, it is found that the most important kinematic features may be described by just three variables.

Variables sensitive to the “internal structure” of tau lepton production and decay are also studied. It is found that only the invariant mass variables contribute significant discrimination power.

After implementing additional measures to reject charged-current events, a search for three-prong tau decays is carried out for the 1995-96 data, based on simple cuts in the kinematic variables. It is found that at present, a background process—most probably a charged-current background—contributes surviving events at a higher level than predicted by the Monte Carlo simulation. After correcting for this background we find no evidence for $\nu_\mu \rightarrow \nu_\tau$ oscillations, setting a limit in the mixing parameter

$$\sin^2 2\theta < 0.020$$

for large mass-squared-differences, at 90% confidence, based on the three-prong tau decays alone. The present sensitivity of the analysis to the mixing parameter is found to be 0.032.

An independent body of work on particle identification techniques is also described, and its potential application is discussed.

Acknowledgements

Over $5\frac{1}{2}$ years in the NOMAD collaboration I have received invaluable assistance from many people. In the first place my associate supervisor, Dr Kevin Varvell, has been a constant example, a humble and insightful critic, and a good friend throughout our association. Much of my work would have been impossible without his selfless dedication to facilitating physics analysis within the NOMAD collaboration; I will not embarrass him any further by listing my remaining debts to his teaching, work and good humour.

Drs Paul Soler and Burkhard Schmidt have carried much of the technical burden of the Trigger-Veto group, and have encouraged and taught me in my halting efforts on the technical side. They have also given me support and friendship over many years: to Paul in particular I am indebted for many small matters of physics, lots of good advice, and his companionship.

I have learnt a great deal of physics, and received much specific help, from my occasional collaborators within NOMAD: Drs Dario Autiero, Achim Geiser and Peter Hurst. In particular, it seems that I learn something every time I speak with Achim. Drs Martin Sevier and Fred Weber have also been invaluable sources of ideas and information. Peter and Fred, furthermore, have been most companionable neighbours in our offices at CERN. Among the senior staff I am in debt to the vigorous example and clear presentations of our spokesman, Dr Luigi Di Lella, and to a tutorial course and many discussions on experimental statistics with Professor Bob Cousins of UCLA.

From my fellow students I must single out my work- and office-mate of five years, Mr (now Dr) Steven Boyd, for his companionship, and his tolerance in the face of our strongly opposed work-styles and temperaments. I will, although he may not believe this, miss him. In their very different ways now-Drs Josh Long and David Steele of Johns Hopkins have been good friends and interesting workmates. My friends and fellow-students from the University of Melbourne, Mr Lyle Winton and now-Drs Fawzi Fares, Nikki Hyett and Caroline Poulsen, helped to keep me sane at CERN, were considerate flatmates in St Genis, and have always made me welcome in Melbourne.

My supervisor, Dr Juris Ulrichs, has signed my forms, kept me up-to-date with paperwork and other duties, encouraged me in my work, read my thesis chapters, and done all of those other things one needs done, with a uniform good humour. I am grateful to him and to all the other members of the Falkiner High Energy Physics Department for making it such an easy-going place to work. Last but not least of all I must thank our boss, A/Prof. Lawrence Peak: for his tutoring in the Third Year Lab (many years ago now), for keeping a place for me in the department over many uncertain years, and for his loyal and unaffected leadership.

Special acknowledgement: Particle identification studies

Much of the work described in chapter 5 was “published” within the NOMAD Collaboration in 1997 and 1998 with help and encouragement from various people: Steven Boyd gave help and advice on a sample of leakage tracks from the FCAL for a study mentioned in section 5.3.7; Achim Geiser contributed the idea for the discussion of sensitivity to distribution “slope” in section 5.3.8; Mr (now Dr) Kyan Schahmaneche supplied the sample of Λ^0 from the 1995 data, used throughout section 5.3, and gave advice on V^0 cuts; and Dr Slava Valuev provided advice on a number of features of the standard TRD code.

Particular thanks are due to Slava and all of the TRD group within NOMAD, whose TRD simulation and (accessible and documented!) analysis code I freely used and adapted; and to Peter Hurst, who provided advice, analysis code and a filtered sample of data from an isolated-hadron study, on which all of this work was originally tested. The technique for enriching a particle sample in kaons (section 5.3.7), and several other matters, arose in the course of many discussions with Peter.

Author's contribution

The historical and theoretical discussion of chapter 1 was prepared from the references cited, using the bibliographies and commentaries of Ezhela et al. (1996) and Cahn and Goldhaber (1989) as resources. Theoretical discussions of Halzen and Martin (1984), Kim and Pevsner (1993) and contributors to the Review of Particle Properties (Caso et al., 1998) were also repeatedly consulted. The arrangement and presentation of the material is my own.

The description of the West Area Neutrino Facility, the detector and data acquisition (sections 2.2, 2.4 and 2.5 respectively) draw on various written sources in addition to the author's "working knowledge" of the experiment: these have all been acknowledged *in situ*. The "slow control" system for the veto counters (sections 2.4.2 and 2.5.2) is the work of the author; the remaining software and the experimental equipment described in chapter 2 are the work of other members of the NOMAD collaboration. The presentation and the discussion of the strategy of the tau search (section 2.3) and the experimental design (section 2.4) are due to the author.

The discussion of variables previously introduced in the analysis of three-prong tau decays (section 3.3) is a summary of the work of others, as noted in the references, although the presentation is my own. The rest of chapter 3 is original work.

The analysis of the physical principles underlying the 3π structure variables presented in chapter 4 is original work. Most of the *quantities* discussed are variables in common use, or advocated by other members of the NOMAD collaboration: an exception is the "required mass" variable M_R introduced in section 4.3, which is due entirely to the author. In the discussion of polarisation-sensitive variables (section 4.5), the simple argument to establish the complete polarisation of the tau lepton in $\nu_\tau + N \rightarrow \tau^- + X$ is due to Luigi Di Lella and the author; the remainder of the discussion of polarisation draws on the theoretical literature and experimental (*i.e.* LEP) papers on the question, as acknowledged.

The TRD simulation code, and the electron/pion discrimination algorithm described in chapter 5 are the work of the TRD group within NOMAD, while the interpretation of the acceptance variables for electron veto purposes (section 5.2.4) is original. Both the idea of applying this method to particle discrimination by dE/dx , and the implementation and application of the method described in that chapter (section 5.3), are due to the author. Assistance from other members of the NOMAD Collaboration with the material of that chapter has been noted separately above.

The quality cuts presented in section 6.2 are in common use within the NOMAD collaboration, while the neutral-secondary and hadronic calorime-

ter studies described in section 6.3 are my own independent work. Of the elements of the muon identification and veto algorithms (section 6.6), the acceptance calculations (6.6.4) are the work of the author, while the remaining elements are the work of other collaboration members. The discussion of the principles, and the final choice of cuts (section 6.7) are my own.

The application of the analysis to the data (chapter 7) is my own work, using an internal NOMAD paper (Feldman and Geiser, 1999) as a technical reference for the oscillation limit calculation.

The solution for the longitudinal momentum of the final-state neutrino, used in the derivation of the required mass variable M_R in appendix B.1, follows an argument used by Luigi Di Lella in a $\tau^- \rightarrow \nu_\tau \pi^-$ analysis. I am grateful to Luigi for our discussion of this question. The remainder of the appendices is original.

Except where explicitly acknowledged, figures and tables are the work of the author.

Preface

This thesis is a report of work carried out for the degree of Doctor of Philosophy at the University of Sydney. It describes the NOMAD experiment and the neutrino oscillation search carried out there, in particular presenting techniques for observing three-prong tau decays in the NOMAD data, as a signal of tau appearance. A search for such decays in the 1995-96 data is performed, and both the results and their implications for the work presented here are discussed. Independent studies of particle identification techniques for the NOMAD experiment are also presented.

The first chapter describes our developing understanding of neutrinos, neutrino mass and mixing, and the phenomenon of neutrino oscillations. The motivation for the present neutrino oscillation search, and the present evidence for oscillations due to other experiments, are also discussed. The NOMAD experiment is then described in chapter 2, paying attention both to the strategy of the neutrino oscillation search and the design of the detector.

The principles of the three-prong decay analysis are then developed in chapters 3 and 4. Chapter 3 presents an analysis of three-prong tau decays, as distinguished from neutral current interactions, in the approximation where particle masses and the invariant masses of systems of particles are ignored. This analysis is then extended in chapter 4 where the “structure” of tau production and the $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ decay are taken into account.

Techniques for discriminating between particle species, based on information from the NOMAD Transition Radiation Detector (TRD), are then discussed in chapter 5. The standard algorithm for electron/pion discrimination is presented, and its application to the rejection of electrons in hadronic decay and other analyses is discussed (section 5.2). Tools for discrimination between other particle species, based on ionisation energy deposited in the TRD straw tubes, are then developed, and their performance and prospects for their use are assessed (section 5.3).

Returning to the tau decay search (chapter 6), the reconstruction of the total event momentum, quality cuts on the data, and other technical details are discussed in sections 6.2 through 6.4. The background due to charged-current events, and techniques to suppress it, are then discussed in sections 6.5 through 6.7.

The search for three-prong tau decays itself is presented in chapter 7, where its results are also discussed. A summary of the findings of this analysis is then given in chapter 8.

Contents

Abstract	i
Acknowledgements	ii
Author's contribution	iv
Preface	vi
Contents	vii
List of Figures	xiii
List of Tables	xix
1 Neutrinos and neutrino oscillation	1
1.1 Introduction	1
1.2 The neutrino as a light, weakly interacting particle	2
1.3 Neutrinos and the Standard Model	5
1.3.1 Key experimental results	5
1.3.2 Neutrinos as members of weak isospin doublets	10
1.3.3 Neutrino mass	15
1.4 Neutrino oscillations	20
1.4.1 The solar neutrino problem	20
1.4.2 Solar neutrino oscillations	25
1.4.3 Neutrino oscillation experiments	32
1.4.4 The motivation for a short baseline $\nu_\mu \rightarrow \nu_\tau$ appearance search	35
1.4.5 Recent evidence for neutrino oscillation	40
1.4.6 The current experimental situation	43
2 The NOMAD experiment	45
2.1 Introduction	45
2.2 The West Area Neutrino Facility	46
2.2.1 The elements of the WANF beamline	46
2.2.2 The SPS proton beam on the T9 target	49
2.2.3 The expected neutrino spectrum at the NOMAD	51
2.3 The strategy of the τ search in the NOMAD experiment	55
2.3.1 The kinematic method for detection of τ decay	55

2.3.2	Tau decay modes	56
2.3.3	Experimental backgrounds	59
2.3.4	Choice of decay modes for analysis	61
2.4	Detector design and performance	63
2.4.1	The NOMAD magnet and subdetector layout	65
2.4.2	The veto counters	66
2.4.3	The front calorimeter	68
2.4.4	The drift chambers	70
2.4.5	The trigger counters	73
2.4.6	The transition radiation detector	74
2.4.7	The preshower detector	77
2.4.8	The electromagnetic calorimeter	78
2.4.9	The hadronic calorimeter	82
2.4.10	The muon detection system	84
2.4.11	The silicon target	86
2.5	Data recording and handling	88
2.5.1	Triggering	88
2.5.2	Data acquisition	89
2.5.3	The event reconstruction code in outline	92
3	Variables to distinguish τ decay from NC events	95
3.1	Introduction	95
3.2	The construction and classification of analysis variables	96
3.2.1	The kinematic variables	96
3.2.2	Definition of coordinate axes	97
3.2.3	Internal structure variables	98
3.2.4	Quality cuts	99
3.2.5	Three-pion and three-prong events	99
3.3	Some variables previously used in three-prong decay analysis	100
3.3.1	Angles in the transverse plane: $\phi_{3\pi-h}$ and ϕ_{m-h}	101
3.3.2	Ratios of transverse momenta: $\rho^{3\pi}$ and ρ^{had}	103
3.3.3	The isolation variable q_T	105
3.4	The choice of five kinematic variables	110
3.4.1	The transverse plane	110
3.4.2	The plane of momenta	111
3.4.3	The correlation between q_T and $(\rho^{3\pi}, \rho^{had})$	115
3.4.4	The distribution of events in $(\rho^{3\pi}, \rho^{had}, M)$	118
3.4.5	Correlations between q_T and cuts in the transverse plane	123
3.4.6	A unified description of the kinematic variables	127
3.4.7	Choice of a fifth kinematic variable	133
3.5	Summary: The five kinematic variables	138

4	Variables sensitive to the $a_1^- \rightarrow \pi^- \pi^+ \pi^-$ structure	139
4.1	Introduction	139
4.2	Invariant masses of the 3π system	140
4.2.1	The three-pion mass $M_{inv}^{3\pi}$	141
4.2.2	Two-pion masses: 3π “substructure”	146
4.3	The choice of transverse mass variable	150
4.3.1	The transverse mass M_{\perp}	150
4.3.2	The “required mass” M_R	151
4.3.3	Relative performance of M_R and M_{\perp}	154
4.3.4	A kinematic effect of the three-pion mass	157
4.3.5	Summary: M_R as a transverse mass variable	158
4.4	Opening angles of the 3π system	159
4.4.1	The opening angles and their distributions	159
4.4.2	The $1/\gamma^{3\pi}$ scale of the opening angle	161
4.4.3	The scale-free part of the opening angle	162
4.4.4	The isolation angle	165
4.5	Sensitivity to the τ polarisation?	169
4.5.1	Polarisation of the τ lepton in $\nu_{\tau} + N \rightarrow \tau^- + X$	169
4.5.2	Measurement of τ polarisation at e^-e^+ colliders	171
4.5.3	Polarisation measurement in the NOMAD experiment	172
4.5.4	Prospects for the use of polarisation-sensitive variables	178
4.6	Selection of the 3π -decay candidate	179
4.7	Summary: Variables sensitive to the 3π structure	183
5	Particle identification using TRD information	185
5.1	Introduction	185
5.2	Electron/pion discrimination	186
5.2.1	Transition radiation in the NOMAD	186
5.2.2	A likelihood ratio for e/π discrimination	189
5.2.3	Construction of acceptance functions	190
5.2.4	Interpretation of acceptance functions in analysis	193
5.2.5	Electron/proton discrimination	197
5.2.6	The effect of track angle	199
5.2.7	Treatment of overlapping tracks	200
5.3	Discrimination between two particle species: the general case	202
5.3.1	Energy loss due to ionisation	202
5.3.2	Adaption of the likelihood ratio method	206
5.3.3	A correction for track angle	208
5.3.4	Treatment of overlapping tracks	209
5.3.5	Five cases of discrimination between species	210
5.3.6	A test of p/π discrimination using Λ^0 decays	219

5.3.7	A test of $K/\pi(\mu)$ discrimination using the calorimeters	230
5.3.8	A test of the π/μ discrimination variables	235
5.3.9	A search for $\phi \rightarrow K^+K^-$ decays	241
5.3.10	Discrimination between two particle species using TRD information: a summary	252
6	Event-quality and charged current rejection	253
6.1	Introduction	253
6.2	Cuts on the primary vertex	254
6.3	Choice of objects for estimation of \vec{p}^{tot}	259
6.3.1	Treatment of secondary vertices	260
6.3.2	Cuts on V^0 vertices	260
6.3.3	Treatment of hanging tracks	262
6.3.4	Treatment of neutral secondary “interactions”	264
6.3.5	Use of hadronic calorimeter information	269
6.3.6	A note on the use of charged tracks	273
6.3.7	Summary: reconstruction of event momentum \vec{p}^{tot}	274
6.4	Treatment of detector malfunctions and software problems	276
6.4.1	Events with known subdetector problems	276
6.4.2	Events with reconstruction or simulation problems	277
6.4.3	A special cut on the Fermi momentum in the Monte Carlo	277
6.5	The charged-current backgrounds	281
6.6	Elements of muon identification and veto algorithms	283
6.6.1	Reconstructed track segments in the muon system	283
6.6.2	Association of drift-chamber and muon tracks	285
6.6.3	The muon veto scintillators	286
6.6.4	Acceptance of the muon detection system	287
6.6.5	Energy deposition in the calorimeters	288
6.6.6	Muon decay, large angle scattering, <i>etc.</i>	290
6.6.7	Event structure: The transverse size ratio R_T	291
6.7	Choice of lepton rejection cuts	292
6.7.1	Cuts against identified muons	293
6.7.2	The kinematic properties of surviving events	294
6.7.3	The muon veto for high- p_T tracks	297
6.7.4	Cuts against identified electrons	301
6.7.5	The electron veto for high- p_T tracks	302
6.7.6	Choice of cuts on the transverse size ratio	307
6.7.7	Summary: The lepton-veto cuts	319

7	The search for three-prong tau decays	321
7.1	Introduction	321
7.2	Definition of cuts on the normalised Monte Carlo events . . .	322
7.2.1	Normalisation of the simulated events	322
7.2.2	A problem in the estimation of efficiency of cuts	326
7.2.3	The choice of variables and cuts against the total back- ground	327
7.3	A search for three-prong τ^+ decays using an equivalent analysis	330
7.3.1	Distributions for Monte Carlo and data positives . . .	331
7.3.2	Events surviving cuts in the data positives	337
7.3.3	Adjustment of cuts and construction of a correction factor	341
7.3.4	“Data simulator” corrections for the neutral current background	343
7.4	Results of the analysis for the 1995-1996 data	344
7.4.1	Choice of final q_T cut and blind region	345
7.4.2	Comparison of distributions for Monte Carlo and data events	347
7.4.3	Results in the signal region	352
7.5	Estimation of neutrino oscillation parameters	353
7.5.1	Quantities determining the oscillation probability . . .	354
7.5.2	Limit and exclusion region for $\nu_\mu \rightarrow \nu_\tau$ oscillation . . .	355
7.6	A discussion: The present state and prospects of this analysis	359
8	Conclusion	363
	Appendices:	365
A	Formulae for the relation between planes of the event	367
A.1	Representation of the rotation using Euler angles	367
A.2	Some notes on expressions for $p^{3\pi}$, p_T^{had} and p_T^{miss}	369
B	Some results concerning transverse mass variables	371
B.1	Derivation of M_R	371
B.2	Comparison of M_R , M_T and M_\perp	373
B.3	The interpretation of M_R as the “required mass” of the tau .	375
B.4	Bounds on the transverse mass given ideal measurement . . .	377
B.5	A transverse-plane approximation to the required mass	379

C	Constraints on the opening angle of $\tau \rightarrow \nu 3\pi$ pions	381
C.1	The constraint on an individual pion	381
C.1.1	Motivation	381
C.1.2	The condition to escape a cone of given opening angle	383
C.1.3	Variation of the factor k with momentum	384
C.2	Decay angles in the backward hemisphere	386
C.2.1	Escape to large angles $\theta > 1/\gamma$	386
C.2.2	Lower bounds on θ for the backward hemisphere	387
C.3	Decay angles in the forward hemisphere	389
C.3.1	Lower bounds on θ near $\theta^* = 90^\circ$	389
C.4	Constraints on the distribution of $\gamma^{3\pi}\theta_{\max}$	390
C.4.1	The upper bound on $\gamma^{3\pi}\theta_{\max}$	390
C.4.2	The shape of the $\gamma^{3\pi}\theta_{\max}$ distribution	391
C.4.3	Approximate lower bounds for $\gamma^{3\pi}\theta_{\max}$	391
C.4.4	Summary	393
	Bibliography	395

List of Figures

1.1	Hadronic cross-section near the Z^0 mass.	9
1.2	The solar neutrino spectrum predicted by the standard solar model.	21
1.3	Neutrino oscillations through the MSW effect.	29
1.4	The $(\sin^2 2\theta, \Delta m^2)$ plane for the MSW effect.	30
1.5	Allowed regions in $(\sin^2 2\theta, \Delta m^2)$ for the MSW effect.	31
1.6	Exclusion regions in $(\sin^2 2\theta, \Delta m^2)$ for appearance and disappearance searches.	33
1.7	Exclusion regions in $(\sin^2 2\theta, \Delta m^2)$ for $\nu_\mu \rightarrow \nu_\tau$	36
1.8	Confidence intervals for the LSND neutrino oscillation claim.	41
1.9	Confidence intervals for the Super-Kamiokande neutrino oscillation claim.	42
2.1	Schematic layout of the WANF beam line.	47
2.2	The principle of the horn and reflector magnetic lens system.	48
2.3	The SPS cycle for the West Area Neutrino Facility.	50
2.4	Parent particles for ν_μ at the NOMAD.	51
2.5	The predicted energy spectra and transverse position distributions of neutrinos at the NOMAD.	53
2.6	The Neutrino Oscillation Magnetic Detector (NOMAD).	64
2.7	Front view of the veto counters.	67
2.8	The NOMAD front calorimeter.	69
2.9	Overview of a NOMAD drift chamber.	71
2.10	Layout of a trigger plane.	73
2.11	Top view of TRD modules.	75
2.12	Response of a TRD module to pions, electrons, and ^{55}Fe	76
2.13	An exploded view of the preshower detector.	78
2.14	An ECAL lead-glass counter.	79
2.15	Preshower and ECAL response to 5 GeV/c electrons and pions.	81
2.16	A front view of the HCAL.	83
2.17	The NOMAD muon detection system.	85
2.18	Schematic side view of the silicon target.	86

2.19	Layout of an individual silicon layer.	87
2.20	The NOMAD data acquisition system.	90
3.1	A three-prong candidate event, ignoring “internal structure”	97
3.2	The transverse plane for $\tau^- \rightarrow \nu_\tau \bar{\nu}_\mu \mu^-$ analysis	101
3.3	The transverse plane for $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$	102
3.4	The momentum triangle in the transverse plane	104
3.5	Distribution of track p_T relative to the hadronic system.	106
3.6	Distributions of q_T for NC background and signal.	107
3.7	Physical origin of q_T for signal and NC background.	108
3.8	The transverse plane for $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$	110
3.9	The plane of momenta	112
3.10	Momentum $p^{3\pi}$ for signal and background in bins of q_T	113
3.11	Momentum $p^{3\pi}$ for signal and background in bins of q_T	114
3.12	$(\rho^{3\pi}, \rho^{had})$ for signal and background in bins of q_T	116
3.13	$(\rho^{3\pi}, \rho^{had})$ in bins of q_T (<i>cont.</i>)	117
3.14	Summary plot of $(\rho^{3\pi}, \rho^{had})$ for $q_T > 0.7\text{GeV}/c$	118
3.15	$(\rho^{3\pi}, \rho^{had})$ for signal and background in bins of M	119
3.16	$(\rho^{3\pi}, \rho^{had})$ in bins of M (<i>cont.</i>)	120
3.17	Transverse momenta for ideal signal and NC events.	121
3.18	Transverse mass M_\perp in bins of q_T	124
3.19	s_T for signal and background in bins of q_T	125
3.20	s_T in bins of q_T (<i>cont.</i>)	126
3.21	A generic event as a rotation of the plane-of-momenta.	129
3.22	M for signal and background in bins of q_T	131
3.23	M in bins of q_T (<i>cont.</i>)	132
3.24	Mean values of the p_T variables in bins of q_T	134
3.25	Distributions of the quantity $(p_T^{had} + p_T^{miss} - p_T^{3\pi})$ in bins of q_T	136
3.26	$\rho^{3\pi}$ in bins of q_T , after the transverse mass cut.	137
4.1	Three-pion mass for a_1^- and other three-prong decays.	141
4.2	The quantity q_{lep}/p^{had} for ranking 3π combinations.	142
4.3	Three-pion mass in tau decay and neutral current events.	143
4.4	Three-pion mass under kinematic cuts.	144
4.5	Three-pion mass with kinematic (including energy) cuts.	145
4.6	Candidate selection including a mass cut.	145
4.7	Two-pion (rho) mass under kinematic cuts.	147
4.8	Two-pion (rho) mass with an $M_{inv}^{3\pi}$ cut imposed.	147
4.9	Comparison between two-pion mass definitions.	148
4.10	Squared masses s_1 and s_2 in tau decay and NC events.	149
4.11	Momentum vectors for calculation of the transverse mass.	152

4.12	$\rho^{3\pi}$ in bins of q_T , for different transverse mass cuts.	155
4.13	The effect of M_R and M_\perp cuts on the transverse plane.	156
4.14	Tau decays with high- and low- $p_T^{3\pi}$	157
4.15	Opening angles of the 3π system.	159
4.16	The opening angle $\theta_{\pi\pi}$ in tau decay and neutral current events.	160
4.17	The opening angle θ_{\max} in tau decay and neutral current events.	160
4.18	The opening angle for a pion in the 3π system.	162
4.19	The scale-free opening angle $\gamma^{3\pi}\theta_{\max}$ under an energy cut.	163
4.20	The scale-free angle $\gamma^{3\pi}\theta_{\max}$ under mass and energy cuts.	164
4.21	The isolation angle θ_{isol}	166
4.22	The isolation angle θ_{isol} for tau decay and neutral current events.	167
4.23	The isolation angle θ_{isol} under cuts.	168
4.24	Particle spins in the quark-lepton CMS.	170
4.25	The polar angle of the tau decay θ_a in the Monte Carlo.	173
4.26	Angle between the tau direction in the quark-lepton CMS and laboratory frames.	175
4.27	The polarisation variable $\cos^2\beta_a$ under a strong energy cut.	176
4.28	The polarisation variable $\cos 2\gamma_a$ under a strong energy cut.	177
4.29	The polarisation variable $\sin 2\gamma_a$ under a strong energy cut.	177
4.30	Comparison of 3π selection efficiency for two simple schemes.	179
4.31	Improvement in 3π selection efficiency under kinematic cuts.	180
4.32	The 3π selection efficiency under tighter kinematic cuts.	181
5.1	Energy deposition by 10 GeV/ c muons in the TRD.	187
5.2	Energy deposition by delta rays in the TRD	187
5.3	The number of TR photons emitted by a TRD radiator.	188
5.4	Normalised energy deposition spectra for 10 GeV/ c π and e	189
5.5	Likelihood ratio distributions for 10 GeV/ c pions and electrons	191
5.6	Expected pion acceptance ϵ_π for fixed values of electron identification efficiency ϵ_e	192
5.7	Pion acceptance $\tilde{\epsilon}_\pi$ for primary tracks, with an electron rejection cut $\tilde{\epsilon}_e < 0.10$ shown.	194
5.8	Pion acceptance $\tilde{\epsilon}_\pi$ for primary tracks, under progressively tighter cuts in $\tilde{\epsilon}_e$	195
5.9	Pion acceptance $\tilde{\epsilon}_\pi$ for primary tracks in the ν_e CC Monte Carlo.	196
5.10	Positive primaries in the ν_e CC Monte Carlo, by species.	197
5.11	Pion acceptance $\tilde{\epsilon}_\pi$ for positive primaries in the ν_e CC Monte Carlo, with protons suppressed.	198
5.12	Pion acceptance $\tilde{\epsilon}_\pi$ for pions from Λ^0 decays in the data.	199
5.13	Simulated TRD response for isolated and overlapping tracks.	201
5.14	Mean rate of energy loss due to ionization in various materials.	203

5.15	Energy deposition in the TRD by pion and proton tracks.	204
5.16	Likelihood ratio distribution for low momentum p, π	206
5.17	Pion acceptance $\tilde{\epsilon}_\pi^{p/\pi}$, with and without track angle correction.	209
5.18	Distributions of p/π likelihood ratio at three momenta.	211
5.19	Expected pion acceptance for different proton identification efficiencies at $p = 0.5, 1.0, 10.0$ GeV/ c	211
5.20	Distribution of p/π likelihood ratio at 2.0 GeV/ c	212
5.21	Distributions of p/K likelihood ratio at three momenta.	213
5.22	Expected kaon acceptance for different proton identification efficiencies at $p = 0.5, 1.0, 10.0$ GeV/ c	213
5.23	Distributions of K/π likelihood ratio at three momenta.	215
5.24	Expected pion acceptance for different kaon identification efficiencies at $p = 0.5, 10.0$ GeV/ c	215
5.25	Distribution of K/π likelihood ratio at $p = 0.25, 0.35$ GeV/ c	216
5.26	Distributions of π/μ likelihood ratio at low momentum.	217
5.27	Distributions of π/μ likelihood ratio at intermediate momentum.	217
5.28	Armenteros plot from the 1995-1997 data.	220
5.29	Detail of the Armenteros plot from the 1995-1997 data.	221
5.30	Armenteros plot for the 1995 data Λ^0 sample.	223
5.31	Distributions of p/π discrimination variables for the positive arm of $\Lambda^0 \rightarrow p\pi^-$	224
5.32	Pion acceptance for the positive arm of $\Lambda^0 \rightarrow p\pi^-$, in three momentum bins.	225
5.33	Armenteros plot for the Λ^0 sample, with pion rejection cuts.	226
5.34	Proton acceptance $\tilde{\epsilon}_p$ for the negative arm of the Λ^0 s, compared to the positive arm.	227
5.35	Proton acceptance $\tilde{\epsilon}_p$ for the negative arm of the Λ^0 s below 1.0 GeV/ c , expanded scale.	228
5.36	Pion acceptance $\tilde{\epsilon}_\pi$ for the negative arm of the Λ^0 s, compared to the positive arm (protons) for $p^+ < 1.0$ GeV/ c	229
5.37	Momentum spectrum for both negative and positive arms of the Λ^0 below 1.5 GeV/ c	229
5.38	K/π variables for tracks isolated in the calorimeters.	231
5.39	K/μ variables for matched μ s isolated in the calorimeters.	232
5.40	Interaction length for K^\pm, π^\pm as a function of momentum.	233
5.41	K/μ discrimination in the data using an $\tilde{\epsilon}_\mu^{K/\mu}$ cut.	234
5.42	Distribution of $\tilde{\epsilon}_\mu^{\pi/\mu}$ for muons and muon-veto tracks.	235
5.43	Distribution of $\tilde{\epsilon}_\pi^{\pi/\mu}$ for muons, and muon-veto tracks.	237
5.44	The variable $\tilde{\epsilon}_\pi^{\pi/\mu}$ for 10,000 simulated tracks at 2.0 GeV/ c	238
5.45	Distribution of $M(h^+h^-)$ in the 1995/96 data.	242

5.46	$M(h^+h^-)$ at high-momentum in the data and Monte Carlo.	243
5.47	$M(h^+h^-)$ at high-momentum under cuts (data).	244
5.48	$M(h^+h^-)$ at high-momentum under cuts (Monte Carlo)	246
5.49	$M(h^+h^-)$ at low-momentum in the data and Monte Carlo	247
5.50	The K/π discrimination variable $\tilde{\epsilon}_\pi$ for 0.70 GeV/ c kaons.	248
5.51	$M(h^+h^-)$ at low-momentum under cuts (data)	249
5.52	$M(h^+h^-)$ at low-momentum under cuts (Monte Carlo)	250
6.1	Distribution of primary vertices in (x, y) in the NOMAD data.	255
6.2	Distribution of primary vertices in z in the NOMAD data.	256
6.3	Fiducial volume cuts in x and y	257
6.4	Total charge at the primary vertex.	258
6.5	Maximum momentum error among primary vertex tracks.	259
6.6	Comparison of V^0 momentum and presumed line-of-flight.	261
6.7	Closest approach of hanging-track momentum to the primary.	263
6.8	Comparison of neutral interaction momentum and presumed line-of-flight.	264
6.9	Angle between neutral interaction momentum and line-of-flight.	265
6.10	Event p_T^{tot} for differing treatment of neutral interactions.	267
6.11	Error $ \Delta\vec{p}_T $ for differing treatment of neutral interactions.	268
6.12	Event p_T^{tot} for differing treatment of calorimeter neutrals.	271
6.13	Event p_T^{tot} where no heavy neutrals strike the calorimeters.	272
6.14	Fermi momentum of the struck nucleon in Monte Carlo events.	278
6.15	Isolation variable q_T for events in the Fermi momentum tail.	279
6.16	Required mass M_R for events in the Fermi momentum tail.	280
6.17	The 3π energy $E^{3\pi}$ for events in the Fermi momentum tail.	280
6.18	A charged current event interpreted as tau decay.	282
6.19	Reconstruction of tracks in the muon chambers.	284
6.20	Distribution of q_T for ν_μ CC events surviving cuts.	294
6.21	Distribution of M_R for ν_μ CC events.	295
6.22	Distribution of $E^{3\pi}$ for ν_μ CC events surviving cuts.	296
6.23	The effect of the muon track veto.	300
6.24	The electronic backgrounds after the muon track veto.	303
6.25	Electron acceptance $\tilde{\epsilon}_e$ for leading- p_T tracks	304
6.26	The effect of the electron track veto.	305
6.27	Effect of the electron track veto on the ν_μ CC sample.	307
6.28	The transverse size ratio R_T^- for three-prong decays.	309
6.29	Transverse size ratio R_T^- for the charged-current background.	310
6.30	Transverse size ratio R_T^- under changing q_T cuts.	311
6.31	The transverse size ratio R_T^{all} for three-prong decays.	312
6.32	Transverse size ratio R_T^{all} for the charged-current background.	314

6.33	Effect of transverse size ratio cuts on the muonic backgrounds.	315
6.34	Effect of transverse size ratio cuts on the electronic backgrounds.	316
6.35	Comparison of the effect of electron veto and R_T^- cuts.	317
6.36	Comparison of two veto schemes on the electronic backgrounds.	318
7.1	Expected background, and the potential signal at high q_T . . .	329
7.2	The isolation variable q_T in the Monte Carlo and data positives.	331
7.3	The kinematic variables in the Monte Carlo and data positives.	333
7.4	Invariant mass $M_{inv}^{3\pi}$ in the Monte Carlo and data positives. . .	334
7.5	The 3π structure variables in the Monte Carlo and data positives.	335
7.6	The isolation angle θ_{isol} in the Monte Carlo and data positives.	336
7.7	The transverse size ratio variables in the Monte Carlo and data positives.	336
7.8	Isolation q_T in the Monte Carlo and data positives after cuts. .	337
7.9	Kinematic variables for surviving events in the data positives .	338
7.10	Invariant mass $M_{inv}^{3\pi}$ for surviving events in the data positives.	339
7.11	The transverse size ratio variables for surviving events in the data positives.	340
7.12	Events surviving all cuts in the Monte Carlo and data positives.	341
7.13	Correction to the number of survivors predicted by the Monte Carlo	342
7.14	Sensitivity to neutrino oscillations as a function of q_T^{cut}	346
7.15	Events surviving cuts in the data, outside the blind region. . .	347
7.16	The isolation variable q_T in the Monte Carlo and data.	348
7.17	The kinematic variables in the Monte Carlo and data.	349
7.18	Invariant mass $M_{inv}^{3\pi}$ in the Monte Carlo and data.	350
7.19	The 3π structure variables in the Monte Carlo and data. . . .	351
7.20	The transverse size ratio variables in the Monte Carlo and data.	352
7.21	Events surviving cuts in the data, for the full range in q_T^{cut} . . .	353
7.22	The exclusion plot for this analysis in $(\sin^2 2\theta, \Delta m^2)$	357
A.1	The plane of momenta, with cartesian coordinates shown. . . .	368
C.1	Pion momentum components in 3π decay.	382

List of Tables

1.1	Predicted and observed rates for solar neutrino experiments. . .	23
2.1	Relative abundance of neutrino species at the NOMAD.	52
5.1	Number of tracks required for muon contamination to be visible in a sample of pion tracks at 2.0 GeV/c.	239
7.1	Normalisation of the ν_μ CC sample to the 1995-96 data. . . .	323
7.2	Relative normalisation of charged-current MC samples	324
7.3	Correction for the minor neutral current processes.	325
7.4	Relative normalisation of the neutral-current MC sample. . . .	325
C.1	Range in θ^* for which fast decay pions escape to $\theta > a/\gamma$. . .	387

Chapter 1

Neutrinos and neutrino oscillation

1.1 Introduction

Experimental results on neutrinos are of very wide interest. Their importance is due in part to our imperfect understanding of neutrinos—questions which have long been answered for quarks remain unresolved for neutrinos—so that even rather basic information can make a significant contribution to the state of our knowledge. Neutrinos are also intrinsically interesting because they hold the privileged position of participating in the weak interaction alone among the fundamental processes.

This present work describes a search for *neutrino oscillation*, the transformation of one neutrino species into another as it propagates, analagous to the well-known phenomenon of *mixing* in the quark sector. The observation of neutrino oscillation is one of the principal goals of contemporary neutrino research as (*inter alia*) it would establish that at least one of the species of neutrino is massive. Such a result, although widely anticipated, would be the first definite signature of physics beyond the present Standard Model and, depending on the neutrino mass involved, could have cosmological implications.

A very light, neutral particle with a small cross-section for interactions was originally proposed in order to explain the spectrum of electrons in beta decay. The *neutrino* was thus introduced in entirely negative terms, and without any clear hope that it might be observed in its own right. Nonetheless, the successful theory of *weak interactions* was built on this foundation and neutrinos themselves were eventually observed in 1953. This early work and the basic properties of the neutrino are reviewed in section 1.2, while the

place of neutrinos in the now-Standard Model of the fundamental particles and interactions is briefly described in section 1.3, where a framework for understanding neutrino mass is also presented. The long-standing discrepancy in the observed neutrino flux from the sun, the so-called *solar neutrino problem*, is discussed in section 1.4.1, together with the concept of neutrino oscillations which provides the most appealing way to explain the phenomenon (section 1.4.2).

The interpretation of other experiments searching for neutrino oscillation is discussed in section 1.4.3. The experimental and cosmological motivation of the current round of accelerator-based $\nu_\mu \rightarrow \nu_\tau$ appearance searches is then briefly reviewed (section 1.4.4) and the implications of the recent strong claim to have observed oscillation through the disappearance of “atmospheric” (cosmic-ray induced) ν_μ are discussed (sections 1.4.5 and 1.4.6).

The West Area Neutrino Facility at CERN and the NOMAD experiment, where the present study was conducted, are described in some detail in chapter 2. The method of analysis and results are then discussed in the third and following chapters.

1.2 The neutrino as a light, weakly interacting particle

The particle later known as the *neutrino* was first proposed by Wolfgang Pauli in 1930 to explain the continuous β decay spectrum (Pauli, 1934): decays of the form

$$A \longrightarrow A' + e^- + \nu$$

provided for a range of electron energies, rather than the unique energy which would be required by energy and momentum conservation if the reaction had the form $A \rightarrow A' + e^-$. This hypothesis, together with ideas from the then newly-developed quantum electrodynamics, was used by Fermi (1934) to construct a theory of β decay by a new *weak (nuclear) interaction*. Even at this early stage the key properties of the hypothetical neutrino were already established:


1. it was known to be light ($m_\nu \lesssim m_e$) from the shape of the β decay spectrum;
2. it was known to be weakly interacting in the colloquial sense, *i.e.* possessing a small cross-section for interactions with matter, from the lack of direct observation.

This second property was soon underlined when Bethe and Peierls (1934) used the Fermi theory, and observed rates of nuclear β decay, to make a quantitative estimate of the cross-section for the inverse process

$$\nu + A' \longrightarrow e^{\mp} + A$$

and found $\sigma < 10^{-44} \text{ cm}^2$, yielding a mean free path of $\gtrsim 10^{19}$ metres in matter.

Despite the embarrassment of there being “no practically possible way of observing the neutrino” (so Bethe and Peierls) the Fermi theory enjoyed continued experimental success, including in particular the discovery of the K-capture process in gallium



by Alvarez (1938), exactly the type of inverse process which occurs naturally in the Fermi theory but would by no means follow from other explanations of the beta decay process (energy/momentum non-conservation at the event-by-event level, and the like). After 23 years the neutrino itself (strictly speaking, the antineutrino) was observed by Reines and Cowan (1953) through the interaction

$$\bar{\nu} + p \longrightarrow e^{+} + n$$

in the unprecedented antineutrino flux made available by nuclear reactors, first at Hanford, and later at Savannah River (Cowan et al., 1956; Reines and Cowan, 1956).

By this stage the nuclear capture of the μ^{-} (at this point known alternatively as the “ μ meson” or the “mesotron”) had already been interpreted as a weak interaction process

$$\mu^{-} + p \longrightarrow \nu + n$$

(Pontecorvo, 1947; Lee et al., 1949) with the same coupling constant as that in beta decay. Together with the continuous spectrum of electrons produced in the decay of the μ^{-} (Hincks and Pontecorvo, 1949; Leighton et al., 1949), interpreted as due to final state neutrinos

$$\mu^{-} \longrightarrow e^{-} \nu \nu$$

(Lee et al., 1949) the most basic feature of the modern understanding of leptons—that a universal weak interaction couples the charged leptons e^{\mp}, μ^{\mp} and the much-later discovered τ^{\mp} to the neutrinos—had been put in place.

While neutrino interactions are now routinely observed in large numbers, using tertiary beams of neutrinos produced at particle accelerators, all attempts to date at direct measurement of a neutrino *mass* have failed. Experiments instead have set rather stringent upper limits:

- A sustained effort at measurement of the electron neutrino mass by fitting the tritium beta decay spectrum near its endpoint has proved to be limited by systematic errors not presently understood, including the two recent careful analyses of Stoeffl and Decman (1995) and Belesev (1995). While this precludes setting a formal limit based on these experiments, the Particle Data Group conclude that a mass $m_{\nu_e} > 10 \sim 15 \text{ eV}$ would distort the spectrum in a manner not observed in the data (Caso et al., 1998).
- Meson decays $\pi^+ \rightarrow \mu^+ \nu_\mu$ allow a measurement of the mass of the “muon neutrino” (see the next section), with the best current result setting an upper limit of $m_{\nu_\mu} < 0.17 \text{ MeV}/c^2$ at 90% confidence (Assamagan et al., 1996).
- The most recent measurements of the mass of the tau neutrino ν_τ in three- and five-prong tau decays sets an upper limit $m_{\nu_\tau} < 18.2 \text{ MeV}/c^2$ at 95% confidence (Barate et al., 1998).

The experiments described to this point all rely on the effect of neutrino mass on the kinematics of a particle decay, although the experimental method varies from case to case. The observation of interactions from neutrinos produced in the recent supernova 1987A (Hirata et al., 1987; Bionta et al., 1987) introduced neutrino mass measurement based on a completely different effect: the joint measurement of the arrival times and energies of the neutrinos allows for the detection of any relationship between energy and apparent time-of-flight. Based on these data Loredo and Lamb (1989), using a likelihood method and carefully taking into account effects due to neutron star formation, neutrino mass and the acceptance of both the Kamiokande-II and IMB detectors, set a conservative upper limit of $23 \text{ eV}/c^2$ on the electron neutrino mass.

So although the current limits on muon- and tau-neutrino mass allow for electron-like masses, the *electron* neutrino has a mass several orders of magnitude smaller, if its mass is indeed non-zero. The muon- and tau-neutrino are similarly limited to masses orders of magnitude smaller than their corresponding charged leptons. This raises the question of whether neutrinos have a mass at all, and on what grounds one might judge the probability of neutrinos possessing mass, short of a direct measurement. Such questions, and an understanding of the different types of neutrino which we have so far introduced without discussion, requires that neutrinos be set within a broader theoretical framework.

1.3 Neutrinos and the Standard Model

As a unique probe into the weak interaction, neutrinos have played an important role in establishing the Standard Model of the fundamental interactions and particles. This role, together with the place of neutrinos in the Standard Model and in possible extensions of that model, is reviewed in the following sections.

The experimental results establishing the existence of different neutrino species and their behaviour, leading to the idea of successive families or *generations* of fundamental particles, are discussed in section 1.3.1. The current understanding of the weak interaction, based on the electroweak theory, is then briefly reviewed in section 1.3.2. The existence of *mixing* between quark eigenstates of the weak interaction, and the quark mass eigenstates, naturally raises the question of whether mixing also occurs among the leptons. This brings us once again to the question of neutrino mass: the theoretical apparatus for describing and understanding possible neutrino masses, in extensions of the Standard Model, is discussed in section 1.3.3.

1.3.1 Key experimental results

Parity, and the ν_μ as a separate or “sequential” neutrino

As noted in the previous section, the continuous electron energy spectrum in muon decay had been firmly established in 1949: together with the non-observation of other particles, or gamma rays in these decays (Hincks and Pontecorvo, 1948; Sard and Althaus, 1948; Piccioni, 1948), this made the interpretation that these decays possessed two neutrinos in the final state

$$\mu^- \longrightarrow e^- \nu \nu$$

an increasingly natural one. The long muon lifetime also suggested that the process was due to the weak interaction. Pontecorvo (1947) had earlier suggested that the capture of muons on nuclei could also be understood by analogy with the β -decay process, and this idea was cast in quantitative and wide-ranging terms by Lee, Rosenbluth, and Yang (1949) who noted that the three processes of β decay of nucleons, the β -like decay of muons shown above, and muon capture, which they interpreted as

$$\mu^- + p \longrightarrow n + \nu,$$

could all be interpreted as Fermi-type interactions with coupling constants within an order of magnitude of each other. Such a “universal” Fermi interaction would perhaps (they speculated) be mediated by a field with respect to which all particles had the same “charge”.

Since Fermi interaction in its original application to β decay coupled the electron to the neutrino, the application to muon decay raised a question as to the identity of the neutrinos produced in muon decay. Konopinski and Mahmoud (1953) argued that both particles were alike, *i.e.* that a neutrino pair $\nu\nu$ rather than a $\nu\bar{\nu}$ combination was produced. Within a framework where a universal Fermi interaction directly coupled combinations of *four* fermions to each other, it was necessary to explain why many new processes such as $\mu^- \rightarrow e^-e^+e^-$ were not observed. These authors proposed that only interactions where two particles (as opposed to a particle-*antiparticle* combination) were annihilated and two particles produced, and the processes related by crossing, were allowed: identifying the μ^+ as the particle and μ^- as the antiparticle then immediately allowed for $\mu^- \rightarrow e^-\nu\nu$ but not the muon-decay-to-electrons or many other processes which were not observed in nature. This argument is the antecedent of the modern idea of *lepton number* conservation.

The identification of the particles in muon decay was made more difficult by the confusion then-current concerning the Lorentz structure of the weak interaction. This was subsequently determined to be an equal combination of vector and axial vector couplings (Sudarshan and Marshak, 1958; Feynman and Gell-Mann, 1958), with opposite signs (“V - A theory”) to account for the non-conservation of *parity* (invariance under spatial inversion) in the weak interaction, suggested by Lee and Yang (1957a) and established in the β decay of ^{60}Co by Wu et al. (1957). This allowed for the interpretation of the neutrino as a massless particle of fixed helicity (Lee and Yang, 1957a; Landau, 1957), determined to be *negative* or *left-handed* by Goldhaber et al. (1958). Experimental spectra then required that the neutrinos produced in muon decay had opposite helicities

$$\mu^- \longrightarrow e^- \nu_R \nu_L$$

(Nishijima, 1957) so that the interpretation opposite to the one given above, namely that a neutrino and *anti*-neutrino were produced,

$$\mu^- \longrightarrow e^- \bar{\nu} \nu$$

became the natural one.

The non-observation of processes such as $\mu^- \rightarrow e^-e^+e^-$ and $\mu \rightarrow e + \gamma$ was then most readily understood if the weak interaction coupled electrons and muons to *different* neutrinos, ν_e and ν_μ respectively, with *lepton numbers* for both electron- and muon-type particles being separately conserved. Encouraged by the detection of $\bar{\nu}$ interactions—now interpreted as interactions of the state $\bar{\nu}_e$ —Pontecorvo (1959) proposed that this idea be tested by

producing a beam of $\bar{\nu}_\mu$ in $\pi \rightarrow \mu$ decay at an accelerator, and determining if the reaction

$$\bar{\nu}_\mu + p \longrightarrow e^+ + n$$

took place. An experiment according to principles suggested independently by Schwartz (1960) was subsequently carried out at the Brookhaven AGS, where electron-producing reactions were not observed but the reactions

$$\begin{aligned}\bar{\nu}_\mu + p &\longrightarrow \mu^+ + n \\ \nu_\mu + n &\longrightarrow \mu^- + p\end{aligned}$$

were seen (Danby et al., 1962). The modern understanding of different families or *generations* of particles, within which a charged lepton and its corresponding neutrino are coupled by the weak interaction, was thereby established.

The τ lepton and the third generation of fermions

A similar scheme of generations of fundamental particles was subsequently developed to describe the hadrons. Soon after the introduction of the mixing angle θ_C by Cabibbo (1963), describing the weak interactions of the hadrons then-known, *quarks* were proposed as fundamental building-blocks for the hadrons (Gell-Mann, 1964; Zweig, 1964a,b). At this stage three quarks were envisaged, a doublet (subsequently the *up* and *down* quarks) and a singlet (*strange*) to account for hadrons with “strangeness” $S \neq 0$. A fourth or *charm* quark was suggested as early as 1964 by Bjorken and Glashow and was given strong phenomenological motivation by Glashow, Iliopoulos, and Maiani (1970) who showed that such a particle would cause various higher-order weak processes which violated observed selection rules (*e.g.* strangeness-changing neutral currents) to be suppressed. In the resulting scheme there were two generations each of leptons and quarks, with the weak interaction coupling the particles within each doublet (up to the effects of Cabibbo mixing, to be discussed further below).

Although a charmed state had not yet been observed, additional fifth and sixth quarks were suggested when Kobayashi and Maskawa (1973) found that at least three generations of quarks would be required if violation of the CP symmetry, observed in the neutral kaon system, was to be accommodated within the electroweak theory, then recently developed (see the next section). As it happened, the third generation of fermions was first found for the leptons with the discovery by Perl et al. (1975, 1976) of the tau lepton, through pair production $e^+e^- \rightarrow \tau^+\tau^-$ at the SPEAR electron-positron collider, followed by decays $\tau^\pm \rightarrow \mu^\pm$ and $\tau^\pm \rightarrow e^\pm$ with significant missing

energy. Subsequent studies soon found that the decays of the tau involved a final-state neutral of significantly smaller mass, and that a weak interaction coupling of the $V - A$ form, just as for (ν_e, e) and (ν_μ, μ) , was favoured (Perl et al., 1977). In other words the new tau lepton could consistently be interpreted as a successor to the electron and muon, coupled by the weak interaction to a neutrino ν_τ , which was itself presumably to be distinguished from the other neutrinos.

By contrast, although bound states of the fifth (*beauty*) quark were found later in the same year, the third generation of quarks was not formally completed until 1995.

The number of “light” neutrino species

Thus the leptons, and in particular the neutrinos, played a key role in determining the picture of successive generations of particles. They were also to provide clear evidence that the sequence of generations does not continue beyond the three described so far. With the development of electron-positron colliders able to reach a CMS energy equal to the mass of the neutral weak intermediate vector boson Z^0 , a copious and clean source of decays to fermion-antifermion pairs

$$\begin{aligned} Z^0 &\longrightarrow q\bar{q} \\ &\longrightarrow e^-e^+ \\ &\longrightarrow \mu^-\mu^+ \\ &\longrightarrow \tau^-\tau^+ \end{aligned}$$

and in particular

$$\longrightarrow \nu\bar{\nu}$$

became available. One of the earliest results was a determination that the number of light neutrino species was 2.8 ± 0.6 (and < 3.9 at 95% confidence), in a measurement of the Z^0 width by the MARK-II collaboration (Abrams et al., 1989). This measurement was subsequently confirmed and refined by the four experiments at the Large Electron Positron collider (LEP), using striking fits to the hadronic cross-section for e^+e^- centre-of-mass energies close to the Z^0 mass (figure 1.1).

More recent results use the “invisible width” Γ_{inv} of the Z^0 —calculated by subtracting the width due to Z^0 decays to charged leptons and jets (quarks) from the total width—and form its ratio with the partial width to massless charged leptons Γ_l (very close to the width to electrons Γ_e). Since the Standard Model predicts the ratio of neutrino to charged lepton widths to high




Figure 1.1: Cross-section for $e^+e^- \rightarrow \text{hadrons}$ near the Z^0 mass in the ALEPH data at LEP, with the expectation for 2, 3 and 4 light neutrino species superimposed. (ALEPH (1993); *cf.* Decamp et al. (1992), figure 12.)

precision $((\Gamma_\nu/\Gamma_l)_{\text{SM}} = 1.991 \pm 0.001)$, the number of light neutrino species may be determined from the double ratio

$$N_\nu = \frac{\Gamma_{\text{inv}}/\Gamma_l}{(\Gamma_\nu/\Gamma_l)_{\text{SM}}} \quad (1.1)$$

$$= 2.993 \pm 0.011, \quad (1.2)$$

fitting to all of the LEP data (Caso et al., 1998).

Here “light” means significantly below $\frac{1}{2}m_{Z^0} = 45.6 \text{ GeV}$ so that rather heavy neutrinos are also encompassed by the result. In any of these measurements the assumption is that the invisible width of the Z^0 is due entirely to particles which have the Standard Model neutrino couplings to the Z^0 . They thus carry out the double function of testing the Standard Model itself, in that the number N_ν is consistent with being an integer, and (if this first test succeeds!) counting the number of species with standard neutrino couplings.

1.3.2 Neutrinos as members of weak isospin doublets

In the preceding sections we have reviewed the role of neutrinos in establishing the now-Standard Model of the fundamental particles and their interactions. We now turn to a brief description of the Standard Model itself, as it applies to the neutrinos.

The electroweak theory

The weak interaction as originally introduced by Fermi was built on an analogy with quantum electrodynamics. The *electroweak* theory at the heart of the Standard Model completes this analogy by describing the electromagnetic and weak interactions as manifestations of a single underlying interaction.

The lepton states of this model are

$$\begin{array}{ccc} \begin{pmatrix} \nu_e \\ e \end{pmatrix}_L & \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L & \begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L \\ \epsilon_R & \mu_R & \tau_R \end{array} \quad (1.3)$$

where the subscripts indicate left- and right-handed states of the chirality operator γ_5 . The left-handed states in each doublet are coupled by the exchange of intermediate vector bosons W^\pm , while the right-handed singlets do not engage in these charged-current interactions: this reflects the observed $V - A$ nature of the weak interaction, discussed above.

The possibility that the weak interaction was mediated by a field with heavy quanta of this kind, along the lines of the theory of Yukawa (1935), had been apparent from earliest times. The electromagnetic interaction is described by a Lagrangian density

$$\begin{aligned} \mathcal{L} &= -j_\mu^{(e)}(x)A^\mu(x) \\ &= -(-ie)\bar{\psi}_e(x)\gamma_\mu\psi_e(x)A^\mu(x) \end{aligned} \quad (1.4)$$

where an electron current $j_\mu^{(e)}(x)$ interacts with the radiation field $A^\mu(x)$ describing the photon. The weak interaction responsible for (say) beta decay, by comparison, is described by the interaction of *two currents*

$$\begin{aligned} \mathcal{L} &= -\frac{G_F}{\sqrt{2}}j^\mu{}^{(N)}(x)j_\mu^{(l)}(x) \\ &= -\frac{G_F}{\sqrt{2}}\bar{\psi}_p\gamma^\mu(1 - g_A\gamma_5)\psi_n\bar{\psi}_e\gamma_\mu(1 - \gamma_5)\psi_\nu \end{aligned} \quad (1.5)$$

(dropping the spatial dependence (x) for clarity): a nucleon current $j_\mu^{(N)}(x)$ coupling the proton and neutron, and a lepton current $j_\mu^{(l)}(x)$ coupling the

electron and neutrino. In the heavy-intermediate-boson interpretation, each current couples to a new field $W_\mu(x)$ by analogy with the electromagnetic case (1.4): its quanta are *charged*, W^\pm , accounting for the charge-changing currents; they are also *massive*, accounting for the Fermi coupling constant which is both dimensional, and small ($G_F = 1.166 \times 10^{-5} \text{ GeV}^{-2}$). Lee and Yang, who had much earlier canvassed the idea of heavy field quanta when proposing the universal Fermi interaction (Lee, Rosenbluth, and Yang, 1949), also noted that the nonlocal effects produced by a boson W^\pm propagating for a finite distance—rather than supposing the four fermions to interact at a point—might be observed in the muon decay process (Lee and Yang, 1957b).

The development of this basic idea into a workable theory proved surprisingly difficult. The $V - A$ character of the weak interaction was a crucial element (Lee and Yang, 1957b), as was the assignment of unit spin to the intermediate boson (Schwinger (1957) and Sudarshan and Marshak (1958), *contra* Tanikawa (1957)). The theory of *massless* vector field quanta had already been developed by Yang and Mills (1954) but the full theory of *spontaneous symmetry breaking*, allowing the field quanta to acquire a mass (without the cost of generating additional unknown massless bosons) was not developed for another decade (Higgs, 1964b,a, 1966). It also turned out to be difficult to explain the weak interactions then-known in isolation: based on a study of symmetries Glashow (1961) proposed instead that there were *four* related intermediate vector bosons, namely the photon γ and the weak intermediaries W^+ and W^- , and an additional *neutral* weak intermediary, the particle we call the Z^0 .

The now-accepted synthesis of electromagnetic and weak processes into a single electroweak interaction, due to Weinberg (1967, 1971) and Salam (1968), incorporates all of these features (note also the “speculation” of Glashow et al. (1970)). It posits underlying interactions coupling to *weak isospin*, obeying the $SU(2)$ symmetry group; and *weak hypercharge*, obeying the $U(1)$ symmetry group. Massless gauge bosons W^+ , W^0 and W^- , corresponding to the generators of the $SU(2)$ group, couple the fermions to each other according to their weak isospin T : the charged bosons W^\pm couple the left-handed leptons in the doublets ($T = \frac{1}{2}$) of (1.3) to each other, while the neutral boson W^0 couples each lepton to itself; the right-handed, weak isospin singlet states e_R, μ_R and τ_R , with $T = 0$, do not participate in this interaction. The massless gauge boson B^0 , corresponding to the single generator of the $U(1)$ group, couples each of the particles, the right-handed states included, to itself according to its weak hypercharge Y .

The underlying symmetry

$$SU(2)_L \times U(1)_Y$$

is then spontaneously broken in such a way that a new, derivative $U(1)$ symmetry is still obeyed. The new boson corresponding to its generator is the (massless) photon γ , and this group governs the electromagnetic interaction, coupling to *electric charge*

$$Q = T^3 + \frac{Y}{2}, \quad (1.6)$$

where T^3 is the third component of weak isospin. The intermediate vector bosons W^\pm acquire mass, as does a new boson Z^0 , a linear combination of the isospin and hypercharge intermediaries

$$Z^0 = \cos \theta_W W^0 - \sin \theta_W B^0;$$

the orthogonal linear combination yielding the photon

$$A = \sin \theta_W W^0 + \cos \theta_W B^0$$

The observed neutral interactions of the theory, due to the massless γ and the massive Z^0 , thus “cut across” the underlying isospin/hypercharge structure.

A *neutral current* weak interaction, due to the Z^0 , is thus a straightforward prediction of this model. Strangeness-changing neutral currents were known not to occur (or at least to be highly suppressed) amongst hadrons, but general weak neutral current processes are experimentally difficult to disentangle from the far more numerous electromagnetic interactions, which also involve neutral currents. This difficulty is not present, however, for interactions of neutrinos (with charge $Q = 0$) and the prediction of weak neutral currents was vindicated by the discovery of the processes

$$\begin{aligned} \bar{\nu}_\mu + e^- &\longrightarrow \bar{\nu}_\mu + e^- \quad \text{and} \\ \nu_\mu + N &\longrightarrow \mu^- + X \end{aligned}$$

by the GARGAMELLE collaboration at CERN (Hasert et al., 1973b,a), later confirmed by results at Fermilab (Benvenuti et al., 1974).

The CKM matrix: mixing in the quark sector

The quark states of the Standard Model are

$$\begin{array}{ccc}
 \begin{pmatrix} u \\ d' \end{pmatrix}_L & \begin{pmatrix} c \\ s' \end{pmatrix}_L & \begin{pmatrix} t \\ b' \end{pmatrix}_L \\
 u_R & c_R & t_R \\
 d_R & s_R & b_R
 \end{array} \tag{1.7}$$

in a close analogy with the leptons, the key differences being the charge assignments ($Q = +\frac{2}{3}, -\frac{1}{3}$) for the members of each doublet, rather than ($Q = 0, -1$), and the additional right-handed singlets d_R, s_R and b_R . A very important additional feature is the *mixing* of the weak eigenstates. Weak interactions of hadrons with nonzero *strangeness* were long ago taken into account by Cabibbo (1963), with the introduction of the *mixing angle* or *Cabibbo angle* θ_C , which when cast in terms of quark states posits that the weak interaction couples the *up* quark u , not to the *down* quark d , but to the linear combination

$$d' = d \cos \theta_C + s \sin \theta_C \tag{1.8}$$

including the *strange* quark s . This accounts for the relationship between decays of the kaon, with quark content $u\bar{s}$, such as

$$K^+ \longrightarrow \nu_\mu + \mu^+,$$

and decays of the pion, with quark content $u\bar{d}$,

$$\pi^+ \longrightarrow \nu_\mu + \mu^+,$$

together with many other decay modes, in a consistent way. The generalisation to three generations of quarks leads to the mixing matrix

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix} \tag{1.9}$$

called the *CKM matrix* after Cabibbo, and Kobayashi and Maskawa (1973) who first noted that a third generation of quarks was needed to allow for violation of the combined CP symmetry within the electroweak model. To a good approximation the up quark is coupled to the state shown at (1.8)

above—mixing taking place between two generations only, with the single parameter $\sin \theta_C = 0.220 \pm 0.002$ —while full three-generation mixing has the V_{ij} determined by four parameters, which can be chosen as (for example) three mixing angles and a complex phase, the latter being the as-yet-unconfirmed source of the CP symmetry violation.

Mixing among the leptons?

The well-established mixing among quarks suggests that the same phenomenon might also occur among the leptons. It is a matter of convention whether mixing is ascribed to the isospin-up fermions (u, c, t quarks, or the neutrinos), or the isospin-down fermions (d, s, b quarks, or the charged leptons): we will follow the near-universal convention of describing mixing in terms of the isospin-down quarks, but in terms of the isospin-up particles—the neutrinos—in the lepton sector.

Neutrino mixing would require at least one of the neutrino species to have non-zero mass. This follows from a feature of mixing which we did not previously stress, namely that the *weak eigenstates* d', s' and b' defined at (1.9) are superpositions of different *mass eigenstates* d, s and b . If these mass eigenstates were degenerate, then we would be free to choose new eigenstates by a unitary transformation, without observable consequences: in particular, we could choose the states d', s' and b' . The generations would then be completely decoupled, *i.e.* the quarks would not be mixed.

For neutrino mixing to take place, therefore, at least one of the species must be non-degenerate in mass: *a fortiori* its mass must be non-zero. By contrast, the Standard Model as described in this section assumes that neutrinos are massless. Indeed, the hypothesis that neutrino mass is vanishing was important in the development of the model, since it implied the simple description of neutrinos by a fixed-helicity state, neutrinos being purely left-handed (with both helicity and chirality -1), and anti-neutrinos purely right-handed (helicity and chirality $+1$), the so-called “two-component theory” of the neutrino. This allowed both the violation of parity conservation, and the non-occurrence of various other processes, to be naturally understood, and was the ground for the now-accepted $V - A$ form of the weak interaction (Lee and Yang, 1957a; Landau, 1957; Nishijima, 1957).

Together with the striking failure of attempts at direct measurement of neutrino mass (section 1.2), and the very restrictive upper limits on neutrino mass set by those experiments, it has long been considered natural to assign zero mass to neutrinos in the theory. Electroweak-singlet neutrinos ν_R , which we would otherwise expect by analogy with the quark states (1.7), were therefore omitted from the table of lepton states (1.3), where we listed

neutrinos only as part of the weak isospin doublets $(\nu_l, l)_L$. To assess whether this is reasonable, we must first consider mechanisms which would lead to neutrino masses in a “non-minimal” Standard Model, or in more ambitious extensions to the Standard Model: these are the subject of the following section.

1.3.3 Neutrino mass

The two-component or *Weyl neutrinos* described up to this point are massless by construction: the relevant mass terms are missing from the electroweak Lagrangian, and right-handed neutrino states are absent from the Standard Model by *fiat*. As neutral particles, however, neutrinos could be given mass in a surprising number of ways, of which we will consider only the most important.

Dirac and Majorana masses

The most straightforward source of neutrino mass would be the existence of *Dirac mass terms* by analogy with the other fermions. Given the presence of terms

$$\mathcal{L}_{mass}^D = -(\bar{\nu}_R M_D \nu_L + \bar{\nu}_L M_D \nu_R) \quad (1.10)$$

in the Lagrangian, neutrinos acquire a mass in the same manner as the other states of the model. The fields ν_L and ν_R contain creation and annihilation operators for the weak eigenstate neutrinos; in the single-generation case, M_D is a number giving the Dirac neutrino mass; in the case of several generations, where the $\nu_{L,R}$ terms are vectors

$$\nu_L = \begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix}_L, \quad \nu_R = \begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix}_R,$$

M_D is a complex 3×3 matrix describing both mass and mixing, which we diagonalise to find the appropriate mass eigenstates ν_i and their masses m_i . Since both left- and right-handed states are necessary for the generation of mass in (1.10), and right-handed neutrinos are omitted from the lepton states at (1.3), neutrinos have no Dirac mass in the minimal Standard Model.

There is however an additional mass-generation mechanism made possible by the vanishing charge of the neutrino. Majorana (1937) proposed that the neutrino was its own antiparticle: we would write

$$\begin{aligned} \nu_M &= \mathbf{C} \nu_M \mathbf{C}^{-1} \\ &\stackrel{\text{def}}{=} \nu_M^C \end{aligned} \quad (1.11)$$

where C is the charge-conjugation operator and we use the superscript C to denote a charge-conjugate state. Writing Lagrangian contributions in terms of chiral states as in (1.10) it is possible to add *Majorana mass terms*

$$\mathcal{L}_{mass}^M = -\frac{1}{2} \left(\overline{\nu_L^C} M_L \nu_L + \bar{\nu}_L M_L \nu_L^C \right) \quad (1.12)$$

for such a neutrino, where M_L is the Majorana mass in the single-generation case, and a mass/mixing matrix as above in the case of several generations. We write the mass as M_L since we have not had to introduce any additional right-handed states in order to construct the mass terms (1.12), the conjugate ν_L^C being in fact the right-handed *anti*-neutrino of the minimal Standard Model. If we posit an additional right-handed neutrino state ν_R then further terms

$$\frac{1}{2} \left(\overline{\nu_R^C} M_R \nu_R + \bar{\nu}_R M_R \nu_R^C \right)$$

may be added, with the left- and right-handed Majorana masses different in the general case, $M_L \neq M_R$.

A Majorana neutrino thus introduces neutrino mass very economically into the theory, since no new states are required: by setting the neutrino to be its own antiparticle, however, we have immediately discarded lepton number conservation as an exact property; if we write the neutrino in terms of physical rather than chiral states, then

$$\nu_M = \nu_L + \nu_L^C, \quad (1.13)$$

where the two terms have lepton numbers $+1$ and -1 respectively. Amongst other effects, this allows for the rare second-order weak process of *double beta decay* to proceed through a neutrinoless channel

$${}^Z A \longrightarrow ({}^{Z+2}) A + 2e^- \quad (1.14)$$

in addition to the conventional channel

$${}^Z A \longrightarrow ({}^{Z+2}) A + 2e^- + 2\bar{\nu}_e. \quad (1.15)$$

In (1.14) we may consider the antineutrino from the first of the two beta-decay processes to couple *as a neutrino* to one of the remaining neutrons, inducing the second beta decay. Such a process would add a characteristic line signature to the continuous energy spectrum of the two electrons in (1.15), and has been the subject of many experimental searches. The interpretation of Majorana masses (or mass limits) in such experiments is not straightforward as in the general case more than one mass eigenstate

contributes to the process (1.14), and cancellations occur among the masses due to the complex phases of mixing terms. With this caveat in mind, experiments to date set limits on neutrinoless double beta decay rather than making definite observations of the process: the limit has recently been improved by the Heidelberg-Moscow experiment, requiring an “effective mass” of 0.2eV or below at 90% confidence (Baudis et al., 1999).

The place of massive or massless neutrinos in the theory

Given that neutrinos belong only to left-handed doublets of the electroweak model, and that lepton number is conserved, neutrino mass finds no place in the theory. Together with the role of vanishing neutrino mass in establishing our understanding of the weak interaction (discussed above) and the failure to measure neutrino masses, this has made the assignment $m_\nu = 0$ appear the most natural one. A contributing factor has been the fact that if neutrinos *were* to have mass, the masses involved would be on a much smaller scale than those of the other fermions: the two assertions

$$0 < m_{\nu_e} \ll m_e,$$

taken together, appear unmotivated.

Over time this understanding has changed, in part due to the consideration of effects beyond the Standard Model. Since the Standard Model cannot be extrapolated to indefinitely high energies—certainly not as high as the Planck mass—it has become increasingly natural to interpret it as an effective field theory, valid for energies sufficiently below some scale characterising “new physics”. In a theory describing such new processes, if higher-dimension, non-renormalizable operators occur, then the left-handed neutrinos of the Standard Model in general acquire Majorana masses

$$m_\nu \sim a \frac{v_{EW}^2}{M_X} \quad (1.16)$$

where a is dimensionless, $v_{EW} \cong 250$ GeV is the scale at which the electroweak symmetry is broken, and M_X is the scale of new physics (Shrock, 1994, 1996). In the case $M_X \gg v_{EW}$ neutrino masses are not only generated in a natural way, but are automatically small compared to the other mass terms in the electroweak theory.

If the smallness of supposed neutrino masses can be understood, however, there is still the impediment of lepton number conservation, which even if interpreted as an approximate rather than a fundamental feature—so that Majorana neutrinos (1.13) may occur—is known to hold to very high accuracy. Here too it turns out that a very small, but nonvanishing violation

of lepton number conservation may occur as a matter of course. Lee and Shrock (1977) have shown that in arbitrary $SU(2) \times U(1)$ gauge theories, with states similar to those at (1.3) and (1.7), there is a natural suppression of lepton-number violating effects, so that lepton number conservation holds as a good, but approximate symmetry.

The see-saw mechanism

In the case where we suppose that electroweak singlet neutrinos ν_R are to be added to the states at (1.3), there is a further mechanism by which very small neutrino masses can be generated.¹ Allowing for both Dirac and Majorana mass terms in the Lagrangian, we may write for the general case

$$\mathcal{L}_{mass}^{D-M} = -\frac{1}{2} \left(\bar{\nu}^C \mathcal{M} \nu + \bar{\nu} \mathcal{M} \nu^C \right) \quad (1.17)$$

where we have combined the left-handed states into a vector

$$\nu \stackrel{\text{def}}{=} \begin{pmatrix} \nu_L \\ \nu_R^C \end{pmatrix} \quad (1.18)$$

so that the mass operator is a matrix

$$\mathcal{M} = \begin{pmatrix} M_L & M_D \\ M_D & M_R \end{pmatrix}; \quad (1.19)$$

cf. the expressions (1.10) and (1.12) for Dirac and Majorana mass terms. If we also allow for several generations at (1.18) then the mass terms M_L, M_R and M_D generalise to $N \times N$ matrices as in the preceding cases. Considering one generation only, in the limit

$$M_R \gg M_D, M_L \quad (1.20)$$

¹Despite first appearances there is no restriction on right-handed states due to the SLC/LEP bound on the number of light neutrino species (1.2). Right-handed neutrino states, by analogy with the right-handed charged-lepton states e_R, μ_R, τ_R , would be weak isospin singlets ($T = 0$), so that they would not couple to the bosons W^\pm or W^0 . The physical Z^0 boson does however couple to the l_R due to the weak hypercharge of these states, and it might be supposed that it would couple to ν_R in a similar manner. In fact, due to the construction of electric charge at (1.6), the weak hypercharge of hypothetical right-handed neutrinos is constrained to be

$$\begin{aligned} Y &= 2Q - T^3 \\ &= 0 \end{aligned}$$

so that these states do not couple to the hypercharge intermediary B^0 , and hence have no (standard) electroweak interactions at all.

the matrix (1.19) diagonalises to yield mass eigenvalues

$$m_1 \cong \frac{M_D^2 + \frac{1}{4}M_L^2}{M_R}, \quad m_2 \cong M_R \quad (1.21)$$

describing essentially decoupled states, one of small mass M_D^2/M_R (if we set $M_L = 0$) and another close to the large mass M_R . In extensions of the Standard Model which involve a much larger symmetry group, spontaneously broken to give $SU(3)$ of colour and the electroweak $SU(2)_L \times U(1)_Y$ groups—such as the left-right symmetric model, Pati-Salam $SU(4)_C, SO(10)$ and the like—this mass typically corresponds to the scale at which the higher symmetry is broken. In this case the limit (1.20) is satisfied as a matter of course, and the smaller eigenvalue m_1 is driven to a very low mass, much smaller than the other fermion masses, by the large mass M_R —the *see-saw mechanism*.

Other limits give other mass spectra—if $M_D \gg M_R, M_L$, for example, the eigenstates are conventional ν_L , and corresponding sterile species ν_s nearly degenerate with ν_L in mass—and mixing effects between generations can produce additional complex effects which we will not consider.

Summary

There are thus a number of ways to account for neutrino masses, if right-handed states are allowed or the “law” conserving lepton number is relaxed. Nothing is thereby lost from the Standard Model picture described in the previous section: while *historically* the $V - A$ coupling and other features of the electroweak theory were established (at least in part) by the massless-neutrino hypothesis, they may stand without it. Furthermore, as noted above, it is possible to understand the extremely small masses of the neutrinos (supposing them to be non-zero) without recourse to conspiracy between terms in equations, couplings set to small values by decree, or other unmotivated postulates.

In this case *neutrino mixing*, discussed at the end of section 1.3.2, becomes a general expectation. It is somewhat artificial to argue in this way, however, since the longstanding indirect evidence for neutrino mixing—manifested as *neutrino oscillation*—has been instrumental in rehabilitating the idea of neutrino mass. For the remainder of this chapter, then, we consider the present evidence for neutrino oscillation, and the motivation for the present search and other studies of the phenomenon.


1.4 Neutrino oscillations

There are a number of observational and experimental results which are most naturally understood if the different neutrino species are mixed: if the weak eigenstates are linear combinations of the (different) mass states. The long-standing *solar neutrino problem*, the most important of these phenomena, is described in section 1.4.1; its potential resolution if neutrinos mix, resulting in *neutrino oscillation*, is then discussed in section 1.4.2. The principles and sensitivity of other neutrino oscillation experiments are then briefly described in section 1.4.3, and the motivation for the CHORUS and NOMAD experiments is presented in section 1.4.4. The recent claims to have observed neutrino oscillation, and their implications for the present study, are then discussed in sections 1.4.5 and 1.4.6.


1.4.1 The solar neutrino problem

Observation of the solar neutrino flux

In section 1.2 we noted that the existence of the neutrino was finally confirmed using the large antineutrino flux from nuclear (fission) reactors: the dominant source of neutrinos in the vicinity of the earth is the nuclear fusion process taking place near the centre of the sun, with net effect



From 1970 onwards Davis and collaborators endeavoured to measure this neutrino flux in an heroic radiochemical experiment at the Homestake gold mine in South Dakota, based on the inverse beta decay process



in a target of 615 tonnes of liquid C_2Cl_4 . The ${}^{37}\text{Ar}$ atoms produced in these reactions (≈ 0.5 per day) were extracted by purging the system with helium gas, and then counted by observing their slow radioactive decay due to electron capture, in a proportional counter. The rate of reactions (1.23) observed over three decades has been consistently less than half that expected from solar models (Cleveland et al., 1998).

These models of the sun and the fusion reactions taking place at its core have been continuously refined, taking into account measurement and modelling of nuclear reaction rates, measurement of element abundances in the sun, and, more recently, the effects of “heavy” element diffusion. Helioseismological measurements of the speed of sound in the sun as a function of

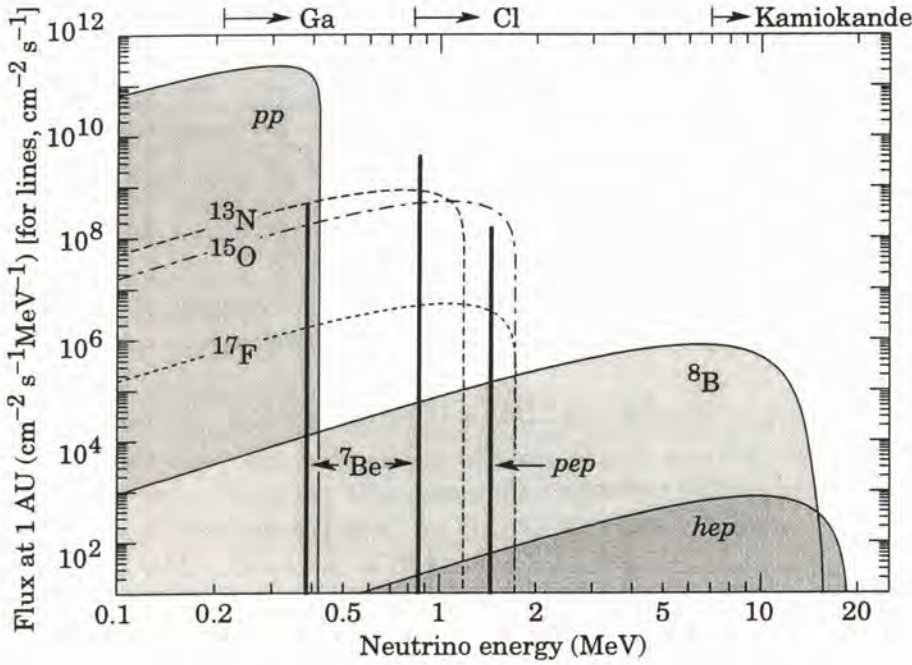


Figure 1.2: The solar neutrino spectrum predicted by the standard solar model. Solid lines show the spectra for the pp chain, and the dashed lines the spectrum for the CNO chain; see the text. The detection thresholds for the gallium and chlorine experiments, and the Kamiokande detector are also shown: note their positions relative to the spectra, in particular the ${}^8\text{B}$ spectrum and the second ${}^7\text{Be}$ line. (Figure due to J.N. Bahcall, taken from Caso et al., 1998).

radius have also recently supplied strong constraints on the temperature and relative element-abundance profiles of solar models: Bahcall et al. (1997) find that the best solar models agree with these constraints to within 2 parts in a thousand. With the possible exception of some of the nuclear reaction rates under solar conditions, the key processes are believed to be well-understood.

The neutrino spectrum expected from the so-called *Standard Solar Model* (SSM) of Bahcall and Pinsonneault (1995) is shown in figure 1.2. The neutrino flux is dominated by the “ pp ” neutrinos produced in the process




at the base of the chain of fusion reactions. The rate of (1.24) largely determines the rate of the net reaction (1.22) so that the observed solar luminosity strongly constrains the pp neutrino flux. Also shown are the monoenergetic

neutrinos produced in electron capture on beryllium,



and the relatively rare *pep* process




which proceeds as an alternative channel to (1.24). The high-energy portion of the spectrum is supplied by neutrinos from the decay



of boron produced in the rare reaction



occurring as an alternative to (1.25) about one time in a thousand. There is also a very small flux due to the “*hep*” reaction



Neutrino rates due to reactions in the competing *CNO cycle*, catalysed by the nuclei of carbon *etc.*, are also shown, dotted.

The reaction (1.23) on chlorine has a threshold of 0.814 MeV (see figure 1.2), so that the Homestake experiment is sensitive to

- the ${}^8\text{B}$ neutrinos, contributing a small fraction of the solar neutrino flux;
- the monoenergetic *pep* neutrinos; and
- the second neutrino line from the electron capture (1.25),

the last of which turns out to be important in the following discussion. More recent experiments detect solar neutrinos via the scattering



in the Kamiokande and Super-Kamiokande water Čerenkov detectors. The thresholds for these experiments are instrumentally determined and have been continually reduced, but even the recent Super-Kamiokande threshold of 6.5 MeV provides sensitivity to the boron neutrino flux alone. However the

Experiment	inverse β on ^{37}Cl (SNU)	inverse β on ^{71}Ga (SNU)	^8B flux ($10^6\text{cm}^{-2}\text{s}^{-1}$)
SSM prediction	$9.3^{+1.2}_{-1.4}$	137^{+8}_{-7}	$6.6(1.00^{+0.14}_{-0.17})$
Homestake	$2.56 \pm 0.16 \pm 0.16$	–	–
GALLEX	–	$69.7 \pm 6.7^{+3.9}_{-4.5}$	–
SAGE	–	73^{+18+5}_{-16-7}	–
Kamiokande	–	–	$2.80 \pm 0.19 \pm 0.33$
Super-Kamiokande	–	–	$2.42 \pm 0.06^{+0.10}_{-0.07}$

Table 1.1: Event rates for the five solar neutrino experiments, as predicted by the SSM of Bahcall and Pinsonneault (1995), and as observed: results are taken from Cleveland et al. (1998) for Homestake, Hampel et al. (1996) for GALLEX, Abdurashitov et al. (1994) for SAGE, Fukuda et al. (1996) for Kamiokande and Fukuda et al. (1998b) for Super-Kamiokande. One Solar Neutrino Unit (SNU) is 10^{-36} captures per atom per second. (SSM prediction taken from Caso et al., 1998, page 329.)

radiochemical experiments GALLEX and SAGE, utilising the inverse beta process on gallium



with a threshold of 0.233 MeV, are sensitive to each of the neutrino sources, in particular the pp neutrinos which dominate the flux.

Taking these differing sensitivities into account, the signal expected for these experiments may be predicted from the Standard Solar Model: these predictions are shown in table 1.1 for the SSM of Bahcall and Pinsonneault (1995), and compared with the observations. In each case a significant deficit is seen, with roughly half the expected rate observed for the gallium experiments, 36% in Super-Kamiokande and a quarter of the expected rate observed in the chlorine experiment. This discrepancy, its persistence over time and presence over a range of experiments with differing techniques, is known as the *solar neutrino problem*.

An astrophysical solution?

The hypothesis that the observed solar neutrino deficit is due to experimental errors, which was economical when only the Homestake experiment was operating, is increasingly untenable. A deficit is seen across five experiments

using three different techniques (four, if the completely different chemical processes of GALLEX and SAGE are taken into account), with agreement on the flux observed between the two gallium experiments, and (separately) the two water-Čerenkov experiments. Furthermore, in addition to the tests of atom extraction efficiencies *etc.* previously carried out, GALLEX and SAGE have calibrated their detectors using bright artificial neutrino sources with neutrino energies close to the two solar ${}^7\text{Be}$ lines (Hampel et al., 1998; Abdurashitov et al., 1999). Both experiments find a response consistent with that quoted for solar neutrinos, with statistical errors of $\approx 10\%$.

This then suggests that the prediction for the solar neutrino flux may be in error. A number of different solar models have been proposed, varying one or other feature or making different assumptions concerning nuclear reaction rates *etc.*. Of those listed in the review of Nakamura (1998) all predict fluxes well above the data shown in table 1.1, with the exception of Dar and Shaviv (1996) who predict a ${}^8\text{B}$ flux consistent with that seen by the water Čerenkov detectors (based on different reaction rates, equation of state *etc.* to Bahcall and Pinsonneault (1995)). The predictions for the signal in the chlorine and gallium experiments from this model, however, while lower than those of Bahcall and Pinsonneault, are still inconsistent with the data.

We have already noted that the sensitivity of the gallium experiments to the neutrinos from the principal reaction (1.24) allows the neutrino flux to be constrained by the observed solar luminosity. The analysis by Hata and Langacker (1997) furthermore finds that even if the relative fluxes of solar neutrinos from the reactions (1.25) and (1.27) are allowed to vary freely, without concern for realistic modelling, the fit to the combined data is extremely poor, excluded at 99% confidence. This situation remains even if one of the three types of experiment (chlorine, gallium or water-Čerenkov) is excluded, or the solar luminosity constraint is omitted, with the fit to the remaining data excluded at $> 98\%$ confidence in each case. This reflects the qualitative problem that the observed deficit varies between the experiments in a manner inconsistent with the neutrino flux originating with the reactions (1.24) through (1.27): applying the ${}^8\text{B}$ flux measurement of the water Čerenkov detectors to the other experiments leads to a best-fit flux of ${}^7\text{Be}$ neutrinos which is significantly *negative* for the chlorine experiment, and even worse for the gallium experiments if the solar luminosity is allowed to constrain the pp neutrino flux.

Even if unlikely fits to the data or the exclusion of several experiments are granted, we should note that the free variation of the ${}^7\text{Be}$ with respect to the ${}^8\text{B}$ flux which seems to be required by the data is extremely difficult to reconcile with the solar origin of the neutrino flux: ${}^8\text{B}$ in the solar core is produced *from* the ${}^7\text{Be}$ by the reaction shown at (1.28). By comparison, the

consistency of the standard solar model with actual solar conditions may be gauged from the helioseismological measurements, which agree at all radii to better than 0.2%; the neglect of even the relatively subtle effect of heavy-element diffusion, on the other hand, introduces disagreements as large as 1%, and 22 times larger in the mean than for the SSM (Bahcall et al., 1997).

For these reasons, even an astrophysical solution to the solar neutrino problem seems increasingly difficult to support.

1.4.2 Solar neutrino oscillations

The concept of neutrino oscillations

By contrast with the difficulty of conventional explanations for the solar neutrino problem, there are a number of elegant solutions based on *neutrino oscillations*. Originally proposed in the form of $\nu \leftrightarrow \bar{\nu}$ oscillations by Pontecorvo (1958), by analogy with the $K^0 - \bar{K}^0$ system, oscillations between different *neutrino* species, $\nu_l \leftrightarrow \nu_{l'}, l' \neq l$, follow generally from the neutrino mixing canvassed in section 1.3.2 above. We set the weak eigenstates ν_l to be linear superpositions of mass eigenstates ν_m

$$|\nu_l\rangle = \sum_m U_{lm} |\nu_m\rangle \quad (1.32)$$

(*c.f.* equation (1.9) for quarks above) where U is a unitary 3×3 matrix, assuming three neutrino species and the theory of section 1.3, including neutrino mass so that mixing is non-trivial. In this case a neutrino emitted as a weak eigenstate ν_l is transformed into a superposition of several weak eigenstates (in general, all three) as it propagates. The “oscillation solutions” to the solar neutrino problem propose that a fraction of the ν_e formed in the core of the sun are transformed into other species, to which the solar neutrino detectors are relatively insensitive, *en route* to the earth.

Each of the mass eigenstates in (1.32) will, after a proper time τ_m , evolve as

$$\begin{aligned} |\nu_m(\tau_m)\rangle &= e^{-im_m\tau_m} |\nu_m(0)\rangle \\ &= e^{-i(E_m t - p_m L)} |\nu_m(0)\rangle \end{aligned}$$

where m_m, p_m and E_m are the mass, momentum and energy of the state, and t, L its propagation time and distance in the laboratory. This yields

$$|\nu_m(L)\rangle \approx e^{-i(m_m^2/2E_m)L} |\nu_m(0)\rangle \quad (1.33)$$

to an excellent approximation for highly relativistic neutrinos with $E_m \gg m_m$, as suggested by the upper limits on the m_m .

A weak eigenstate thus evolves over a distance L to

$$\begin{aligned} |\nu_l(L)\rangle &\approx \sum_m U_{lm} e^{-i(m_m^2/2E_m)L} |\nu_m\rangle \\ &= \sum_{l'} \left(\sum_m U_{lm} e^{-i(m_m^2/2E_m)L} U_{l'm}^* \right) |\nu_{l'}\rangle; \end{aligned} \quad (1.34)$$

where expression (1.32) for the weak eigenstates has been inverted to obtain the second form. That is, a weak eigenstate ν_l becomes in general a superposition of all of the weak eigenstates $\nu_{l'}$. Assuming a superposition sufficiently coherent that subtleties of wave packet analysis *etc.* may be ignored, the probability for a neutrino formed as ν_l to be measured as $\nu_{l'}$ after propagating for a distance L is given by

$$\begin{aligned} P(\nu_l \rightarrow \nu_{l'}; L) &= |\langle \nu_{l'} | \nu_l(L) \rangle|^2 \\ &= \left| \sum_m U_{lm} e^{-i(m_m^2/2E_m)L} U_{l'm}^* \right|^2 \end{aligned} \quad (1.35)$$

The name *neutrino oscillations* is based on the characteristic form of the probability (1.35) as a function of L in the “two-flavour” approximation, where only two weak eigenstates a and b , and two mass eigenstates i and j , are significant. In this case the transition probability is given by

$$P(\nu_a \rightarrow \nu_b; L) = \sin^2 2\theta_{ab} \sin^2 \left(\frac{\Delta m_{ij}^2}{4E} \right) \quad (1.36)$$

in natural units, and

$$= \sin^2 2\theta_{ab} \sin^2 \left(\frac{1.27 (\Delta m_{ij}^2/\text{eV}^2) (L/\text{km})}{E/\text{GeV}} \right) \quad (1.37)$$

in conventional units. The parameter θ_{ab} fully describes the mixing matrix in this approximation,

$$U = \begin{pmatrix} \cos \theta_{ab} & \sin \theta_{ab} \\ -\sin \theta_{ab} & \cos \theta_{ab} \end{pmatrix}, \quad (1.38)$$

and is known as the *mixing angle*. From (1.37) the maximum probability for a transition $\nu_a \rightarrow \nu_b$ is given by $\sin^2 2\theta_{ab}$, so that the mixing angle determines

the strength or degree of mixing. The distance over which one period of the sinusoidal variation of (1.37) occurs is known as the *oscillation length*

$$\begin{aligned} L^{osc} &= \frac{4\pi E}{\Delta m_{ij}^2} \\ &= \frac{2.48 E/\text{GeV}}{(\Delta m_{ij}^2/\text{eV}^2)} (L/\text{km}) \end{aligned} \quad (1.39)$$

The general case of “three-neutrino oscillations” is described by a mixing matrix, the analogue of the CKM matrix, parametrised by three mixing angles and an invariably neglected complex phase; and in addition, two mass-squared differences Δm_{ij}^2 , so that five parameters must in general be considered. Experimental searches for neutrino oscillation typically interpret and report results in terms of two-neutrino oscillations as the phenomenology of (1.37) is straightforward; the general case, furthermore, reduces to the two-neutrino case when all three mixing angles are small.

Vacuum oscillations

The *vacuum oscillation* solution to the solar neutrino problem is the application of this theory to ν_e from the solar core as they propagate to the earth. Given an earth-sun distance of ≈ 150 million kilometres, and a typical neutrino energy of ≈ 1 MeV, we see from (1.39) that the characteristic mass-squared difference for such solutions will be of order 10^{-11} eV². From the large suppression of the solar neutrino flux which must be explained, we see furthermore that the characteristic mixing angles of such solutions will be large. In a fit to all of the experiments Hata and Langacker (1997) find vacuum oscillations to be consistent with the data, even in the two-neutrino approximation, with three allowed regions in the range $\Delta m^2 = (5 - 8) \times 10^{-11}$ eV² and $\sin^2 2\theta = 0.65 - 1$.

Solutions of this kind are also known as “just-so” oscillations, because energy-dependent suppression of the flux by a factor ≈ 2 requires that the oscillation length and the radius of the earth’s orbit be tuned to the level of a few percent. Coincidences of this kind are widely felt to be artificial and vacuum oscillation solutions have been disfavoured as a result. It should nevertheless be stressed that such oscillations are presently consistent with the full data set, and that where mixing between three neutrino states is taken into account, constraints on individual parameters become less, rather than more restrictive.

Matter-enhanced oscillations

Much of the interest in neutrino oscillation solutions to the solar neutrino problem is due to the alternative mechanism of *matter-enhanced oscillations*. In this case, neutrino mixing does occur but the equations (1.33) and (1.37) for neutrino oscillation are modified by an effect due to the medium through which the neutrinos are propagating.

The potential due to interactions with electrons in the solar plasma is significant, and electron neutrinos, which alone among the neutrino species have charged- as well as neutral-current interactions with electrons, may gain an effective mass as large as that of ν_μ or some other (ordinarily) heavier species. That is, the energy levels cross: however, the coupling introduced between ν_e and ν_μ (say) by neutrino mixing has the double effect of forcing the two levels apart, and providing a continuous transition from a high-density region, where the heavier state is essentially ν_e , to the lower-density region, where this same state is essentially ν_μ . Electron neutrinos are thus transformed into another species, if their energy is sufficiently high for the level-crossing to take place. The phenomenon is an example of the well-known quantum-mechanical effect of *anti-crossing*, which occurs quite generally in two-level systems where a perturbation is introduced (see for example Cohen-Tannoudji et al., 1977, pages 405–411). The application to neutrino oscillations in matter was first noted by Wolfenstein (1978) and Mikheyev and Smirnov (1985, 1986), and is known as the *MSW effect*: it is depicted in figure 1.3

The details of this process produce a complicated allowed region in the parameter space ($\sin^2 2\theta, \Delta m^2$) for a given experiment, shown schematically in figure 1.4. Provided the resonant region is large compared to the oscillation length $\pi E/1.27\Delta m^2$, the change in the energy of the levels as the neutrino passes through the sun is too slow to perturb the states: ν_1 and ν_2 are effectively stationary states of the full Hamiltonian and even a very weak coupling ($\sin^2 2\theta \ll 1$) will efficiently transform $\nu_e \rightarrow \nu_\mu$; this “adiabatic” MSW process is thus insensitive to mixing angle but rather fixes a narrow region in Δm^2 (the upper branch in the figure). Where the resonant region is comparable to or smaller than the oscillation length, the mass states ν_1 and ν_2 are no longer stationary states and the transition $\nu_e \rightarrow \nu_\mu$ becomes inefficient (as shown in figure 1.3). A strong vacuum coupling can offset this effect so the solutions corresponding to this “non-adiabatic” process are sensitive to $\sin^2 2\theta$ as well as Δm^2 (the diagonal branch in figure 1.4). Finally, if the coupling in vacuum is sufficiently strong ($\sin^2 2\theta \approx 1$) the transition $\nu_e \rightarrow \nu_\mu$ takes place for a wide range of mass-squared differences Δm^2 (the vertical branch of figure 1.4).

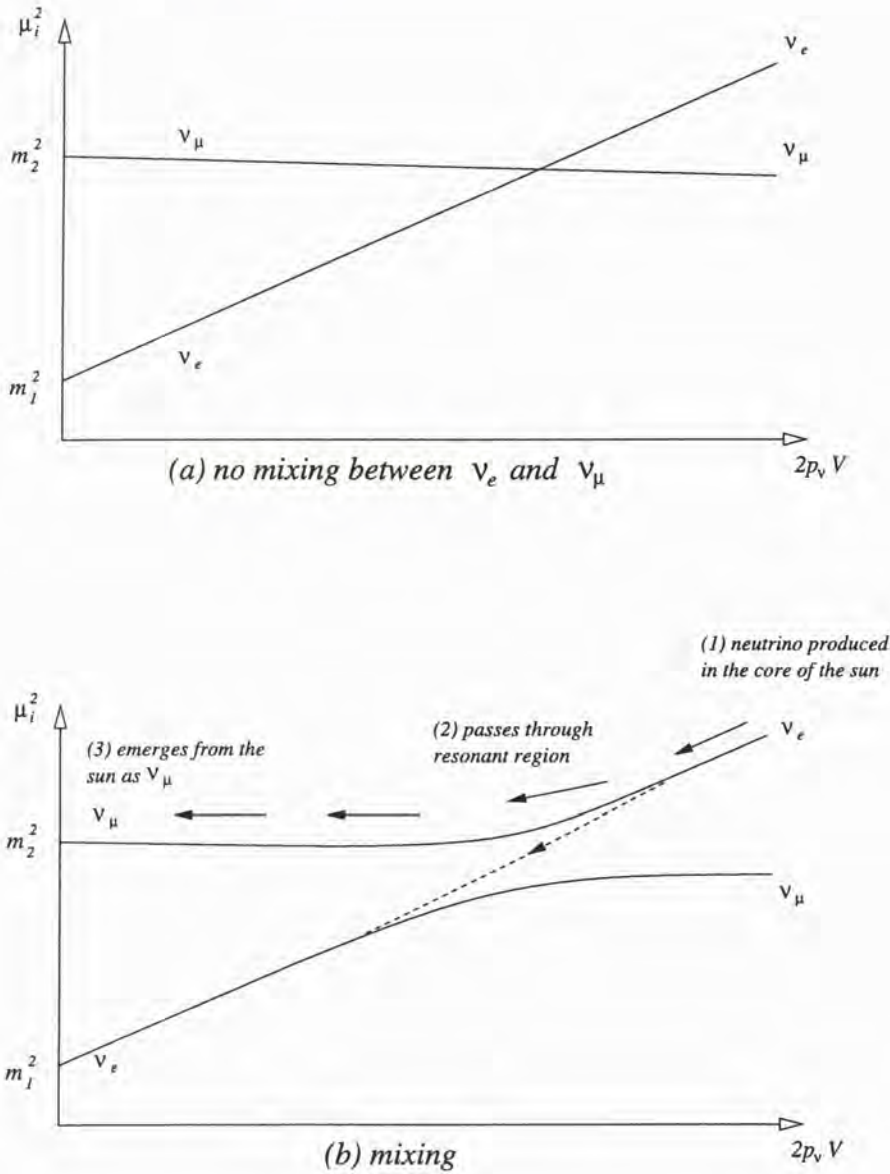


Figure 1.3: Neutrino oscillations $\nu_e \rightarrow \nu_\mu$ through the MSW effect. In the hypothetical case (a) where ν_e and ν_μ are massive (with masses m_1 and m_2) but not coupled, the levels “cross” when the potential V due to interactions with matter becomes sufficiently large, and ν_e becomes the heavier species ($m_1 > m_2$). In the case of coupling (b) the states ν_1 and ν_2 are forced apart in the crossing region, where the mass composition of the weak states also changes: below the crossing region, $\nu_e \approx \nu_1$, while above it, $\nu_e \approx \nu_2$, ($m_2 > m_1$). A sufficiently energetic ν_e formed in the centre of the sun (1) is formed as the heavier state ν_2 , and propagates as ν_2 through the sun where the density ρ and hence the potential V falls through the resonant region (2), and finally emerges from the sun as ν_μ (3). If the $\nu_e \rightarrow \nu_\mu$ oscillation length is comparable to or larger than the size of the resonant region, there is a nonvanishing probability for the transition $\nu_2 \rightarrow \nu_1$ to take place (shown dotted) so that the neutrino emerges from the sun as ν_e .




Figure 1.4: Typical allowed regions in the $(\sin^2 2\theta, \Delta m^2)$ plane for the MSW interpretation of a solar neutrino experiment. See the text.

As the gallium, chlorine and water-Čerenkov experiments sample different parts of the solar neutrino energy spectrum, and this characteristic shape is a function of $\pi E/1.27\Delta m^2$, the sensitivity regions of the experiments are offset from each other in Δm^2 . Requiring consistency between all experiments therefore imposes rather strict constraints on the MSW interpretation of the data. The analysis of Hata and Langacker (1997) however finds that two regions are presently allowed by the full data set, even in the simplest two-neutrino interpretation: both small- and large-mixing angle solutions are possible, with $\Delta m^2 \approx 10^{-5} \text{ eV}^2$ in each case. (See figure 1.5. Regions excluded by an analysis of the E_ν spectrum, to which the water-Čerenkov detectors are sensitive, and also any additional matter-enhanced oscillation effect for neutrinos observed at night, *i.e.* passing through the earth's core, are also taken into account.)

In addition to its consistency with the data, this solution has what may be considered the advantage of requiring no special tuning of the neutrino-mixing and solar-system parameters.

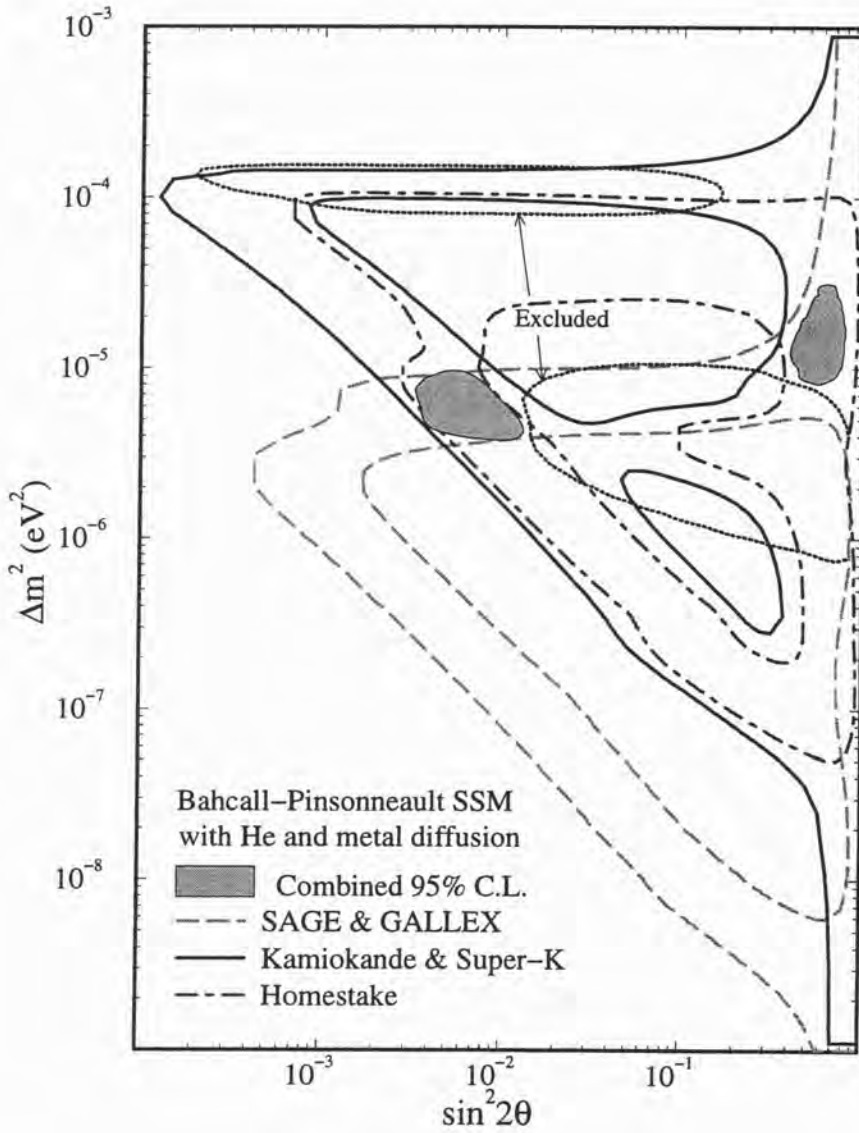


Figure 1.5: Allowed regions in the $(\sin^2 2\theta, \Delta m^2)$ plane for $\nu_e \rightarrow \nu_l$ solar neutrino oscillations by the MSW effect, assuming the SSM of Bahcall and Pinsonneault (1995). Allowed regions are shown for the two gallium experiments (dashed), the water Čerenkov experiments (solid line), and the Homestake experiment (dot-dashed); the regions excluded by Kamiokande spectrum and day-night data are dotted. The two shaded regions are allowed by the combined data at 95% confidence. (From Hata and Langacker, 1997, figure 13.)

1.4.3 Neutrino oscillation experiments

Apart from this important special case, experimental searches for neutrino oscillations may be understood with reference to the expression for the two-neutrino oscillation probability given above,

$$P(\nu_a \rightarrow \nu_b; L) = \sin^2 2\theta_{ab} \sin^2 \left(\frac{1.27 (\Delta m_{ij}^2 / \text{eV}^2) (L/\text{km})}{E/\text{GeV}} \right) \quad (1.40)$$

We suppose that a neutrino detector sits at a distance L from a source of neutrinos ν_a of energy E . In the case where the detector is sensitive to interactions of the second species ν_b , the experiment is known as an *appearance search*: a characteristic exclusion region in the parameter space $(\sin^2 2\theta_{ab}, \Delta m_{ij}^2)$, assuming no evidence for oscillation is found, is shown by the curve (a) in figure 1.6. Supposing that $\Delta m_{ij}^2 L/E \gg 1$ in the units shown (the limit of “large” Δm_{ij}^2) the second term in equation 1.40 is extremely sensitive to the exact propagation length and energy of the neutrinos, which are in general imperfectly known. As far as the sensitivity of the experiment is concerned, then, this term is replaced by its average value over an oscillation “cycle” and the oscillation probability is just

$$P(\nu_a \rightarrow \nu_b; L) = \frac{1}{2} \sin^2 2\theta_{ab} \quad (1.41)$$

for large Δm_{ij}^2 . A limit on the oscillation probability $P < P_{lim}$ thus translates straightforwardly into a limit on the mixing angle

$$\sin^2 2\theta_{ab} < 2P_{lim} \quad (1.42)$$

with no discrimination between nearby Δm_{ij}^2 values. The “reach” of an experiment in $\sin^2 2\theta_{ab}$ is therefore directly related to the size of the sample of neutrino interactions, and the suppression of backgrounds.

At large mixing angles $\sin^2 2\theta_{ab}$, on the other hand, the probability reduces to the second term of (1.40), with the limit for maximal mixing $\sin^2 2\theta_{ab} \rightarrow 1$ being given by

$$\Delta m_{ij}^2 < \frac{\sqrt{P_{lim}}}{1.27(L/\text{km})/(E/\text{GeV})}. \quad (1.43)$$

In this case, while a large dataset and well-controlled backgrounds will improve the limit (as \sqrt{P}), the experimental reach in Δm_{ij}^2 is set in practice by the characteristic L/E involved. An imprecise separation of experiments into *short baseline* and *long baseline* neutrino oscillation searches is made based on




Figure 1.6: Exclusion regions in the two-neutrino oscillation parameter space ($\sin^2 2\theta$, Δm^2) for appearance and disappearance experiments. See the text. (From Groom, 1998)

L/E , with L/E of order 1 km/GeV or less as a short baseline, and $L/E \gg 1$ in the same units as a long baseline. We should note that the important quantity is not the source-detector distance L , but rather L/E , so that an accelerator-based search with $L \approx 1$ km and $E \approx 25$ GeV (such as the NOMAD and CHORUS experiments) has a short baseline ($L/E \ll 1$ km/GeV) while a reactor-based experiment with the same physical baseline but much lower energy $E = 1 \sim 10$ MeV has a “long” baseline ($L/E \gg 1$ km/GeV).

Experiments of this second type are more typically *disappearance searches*, so-called because they are insensitive to the second neutrino species ν_b . Experiments using a nuclear reactor as a source of $\bar{\nu}_e$ and searching for oscillations $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ or $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$ are of this kind, as the respective interactions have too low a CMS energy to produce a μ or τ lepton. While the limits in the cases

of “large” Δm_{ij}^2 and large mixing are given by equations (1.42) and (1.43) as for the previous case, disappearance searches are qualitatively different in that low oscillation probabilities P cannot in practice be ruled out: the signal for oscillations is an observed number of interactions $N_{obs} < N_0$, where N_0 is the expected value; but the expected statistical fluctuations on N are of order $\sqrt{N_0}$ so that the limiting P_{lim} scales only as $1/\sqrt{N_0}$.

Where the expected neutrino flux is known, the expressions quoted above apply, and the typical excluded region for an experiment is given by the curve (b) in figure 1.6: the limit in the mixing angle is poor, as just discussed, but the experiment is sensitive to small mass-squared-differences due to the long baseline. The most sensitive limit of this kind at the time of writing is quoted by Apollonio et al. (1998), searching for disappearance of $\bar{\nu}_e$ from the nuclear reactors at Chooz in the Ardennes, in France. They exclude Δm_{ij}^2 above $9 \times 10^{-4} \text{ eV}^2$ at 90% confidence.

A forthcoming paper by the same group improves this limit to $7 \times 10^{-4} \text{ eV}^2$ and conducts an analysis of a second type, which does not rely on the knowledge of the expected flux but compares the event rate at different baselines L , in this case due to the two reactors at different distances from the detector ($\Delta L = 117 \text{ m}$; Apollonio et al., 1999). The typical exclusion region for an experiment of this type is shown as the curve (c) in figure 1.6: there is a loss of sensitivity at small Δm_{ij}^2 , as the effective baseline is reduced ($\Delta L < L$), and at large Δm_{ij}^2 as the analysis is insensitive to oscillations with $L_{osc} \ll L_{min}$, where L_{min} is the smallest effective baseline in the experiment (*cf.* Apollonio et al., 1999, figure 10).

In the categorisation of experiments defined in this section, the vacuum oscillation analysis of the solar neutrino experiments is a disappearance search with a “known” expected signal (curve (b)) and exceptionally long baseline ($L/E = 10^{12} \sim 10^{13} \text{ km/GeV}$). If, as in the case of solar neutrinos, evidence for oscillation is seen, the exclusion regions described in this section must be replaced by allowed regions for the parameters $\sin^2 2\theta_{ab}$ and Δm_{ij}^2 . In the simplest case such a region is given by the area between two (displaced) curves of the kind shown in figure 1.6: where the energy spectrum of events and other details are also available, both the analysis and the allowed region are more complex, and depend furthermore on the power of the statistical technique used, as discussed (for example) by Feldman and Cousins (1998, section V).

1.4.4 The motivation for a short baseline $\nu_\mu \rightarrow \nu_\tau$ appearance search

The present work describes a search for $\nu_\mu(\nu_e) \rightarrow \nu_\tau$ oscillations in the NOMAD experiment, a short-baseline ν_τ appearance search ($E \approx 24$ GeV, $L = 0.415 \sim 0.830$ km so that $L/E = 0.01 \sim 0.03$ km/GeV; Altegoer et al. (1998a)). Proposed in 1991, the experiment collected data from 1995 through 1998, concurrently with the CHORUS experiment (Eskut et al., 1997), proposed the year earlier and operating with the same neutrino beam and (therefore) similar L and E . The motivation for these experiments was fourfold (Astier et al., 1991a; Armenise et al., 1990):

1. the potential confirmation of the existence and behaviour of the ν_τ by observation of its charged-current interaction

$$\nu_\tau + N \rightarrow \tau^- + X;$$

2. testing for neutrino mass and mixing, of general interest due to the place of neutrinos in the Standard Model (section 1.3), and increasingly plausible given consistent MSW fits to the combined solar neutrino data (section 1.4.2), in an experimentally accessible part of the parameter space;
3. in particular, probing the $\Delta m^2 > 1 \text{ eV}^2$ region, where the implied neutrino mass(es) $m_i > 1 \text{ eV}$ are cosmologically interesting; and
4. addressing the controversy then-current surrounding the claim of a signal for a heavy (17 keV) neutrino contributing to certain β decays (Simpson, 1985; Simpson and Hime, 1989; Hime and Simpson, 1989), with several papers supporting this claim appearing in the same year as the experimental proposal (Hime and Jelley, 1991; Norman et al., 1991; Sur et al., 1991; Žlimen et al., 1991).

The first of these motivations remains as an outstanding task for experimental physics, at the time of writing; the last is no longer relevant, as the 17 keV neutrino claims were discredited prior to the start of NOMAD data-taking.

The parameter space for two-neutrino oscillations $\nu_\mu \rightarrow \nu_\tau$ excluded by previous experiments is shown in figure 1.7. At large mass-squared difference ($\Delta m^2 > 2 \times 10^2 \text{ eV}^2$) the best limit was set by the Fermilab nuclear emulsion experiment E531 ($\sin^2 2\theta < 4 \times 10^{-3}$; Ushida et al. (1986)); at large mixing the best limit was set by the ν_μ disappearance analysis of CDHS data for the low-energy ($E \approx 1 \text{ GeV}$) neutrino beam at the CERN Proton Synchrotron




Figure 1.7: Exclusion regions in the $(\sin^2 2\theta, \Delta m^2)$ plane for $\nu_\mu \rightarrow \nu_\tau$ two-neutrino oscillations from previous experiments: note the best previous limit due to E531 (Ushida et al., 1986), and the similar limits due to CHARM II (Gruwé et al., 1993), CCFR (McFarland et al., 1995), and the neutral- to charged-current ratio analysis of the CHARM/CDHS data by Loverre (1988, 1996). The ν_μ disappearance analyses of CCFR (Stockdale et al., 1984), CDHS (Dydak et al., 1984) and CHARM data (Bergsma et al., 1988) contribute to the limit only in the low- Δm^2 region; see the text. The projected sensitivity of the CHORUS experiment is shown: the sensitivity ultimately attained by the NOMAD experiment for its full dataset (1995-1998) should be similar. (From Eskut et al., 1997, figure 1.)

($\Delta m^2 < 0.23 \text{ eV}^2$). This second limit still represents the best state of our knowledge for large mixing angles, and it is not foreseen that either NOMAD or CHORUS will improve on this limit, while both experiments have already published analyses improving on the E531 limit, based on a subset of their data (Astier et al., 1999a; Eskut et al., 1998b,a). From this, and the projected sensitivity shown in figure 1.7, it is clear that the contribution of these experiments (and any discovery potential) is in the large- Δm^2 region.

Cosmological motivation

Apart from the general interest in the possibility of neutrino oscillations $\nu_\mu \rightarrow \nu_\tau$, mass-squared differences above 1 eV^2 would imply that neutrinos are cosmologically important: this is the key motivation for an oscillation search in this region of the parameter space.

In all of the foregoing discussion we have ignored, for obvious reasons, the gravitational interactions of neutrinos. On a cosmological or even a galactic scale, however, these interactions are potentially significant, because of the large number density of “relic” neutrinos expected from the Big Bang: about 100 cm^{-3} per species, with a 1.7K thermal distribution. If neutrinos are massless, then the corresponding energy density is insignificant as a gravitational source term; but even a small neutrino mass leads to a large gravitational effect, with neutrinos contributing a fraction

$$\Omega_\nu = \frac{\sum_a m_{\nu_a}}{92h_0^2 \text{ eV}} \quad (1.44)$$

of the “critical” energy density required to close the universe. (Here $\sum_a m_{\nu_a}$ is the sum of neutrino masses for $a = e, \mu, \tau$; and h_0 is the Hubble parameter in units of $100 \text{ km.s}^{-1}.\text{Mpc}^{-1}$.) By comparison, the critical density fraction for baryons is limited by Big Bang Nucleosynthesis arguments and the observed light element abundances to $\Omega_B < 0.024/h_0^2$ so that a total neutrino mass of only 2.2 eV is sufficient for the gravitational effect of neutrinos to exceed that of baryons on a large scale (Caso et al., 1998, sections 15, 16 and 18).

It has been known for some time that so-called *dark matter*—matter without significant emission or absorption of radiation—is important for understanding the dynamics of galaxies, clusters, etc.. Studies of the *rotation curves* of a large number of spiral galaxies (typical stellar speeds as a function of radius from the galactic centre) show their mass distribution, especially for smaller galaxies, to be dominated by a non-luminous component, extending well beyond the optical radius of the galaxy (Persic et al., 1996); the implied dark matter content is of order ten times the visible mass in a typical case. The corresponding dark matter critical density fraction is $\Omega_{dm} \approx 0.05$: model-dependent studies of galaxy and gas motions in clusters (White and Fabian, 1995), and galaxy motions on very large scales (Dekel, 1994), suggest even higher density fractions, $\Omega_{dm} \approx 0.2$ and $\Omega_{dm} > 0.3$ respectively.

Neutrinos contributing all or most of the dark matter would have masses in the 10 eV region; in the case of a mass hierarchy this implies $\Delta m_{32}^2 \approx 10^2 \text{ eV}^2$, in the range where the NOMAD and CHORUS experiments are most sensitive. Both the NOMAD and CHORUS experimental proposals are

cast in these terms (Astier et al., 1991a; Armenise et al., 1990), reflecting the fact that neutrinos are the only non-baryonic dark matter candidates which are actually known to exist.

A more sophisticated treatment of dark matter requires that the process of galaxy formation from density perturbations *etc.* be modelled under different assumptions. For this purpose, contributors to the dark matter are divided into two types: *hot dark matter*, which was relativistic during the period of galaxy formation, and *cold dark matter*, which was not; massive neutrinos, which would remain in the relativistic regime until temperatures dropped to the electron-volt level, are the prototypical hot dark matter. This area of cosmological study is in rapid flux, due to theoretical and modelling developments, and in particular new observations by astrophysical satellites. In recent years the fashion has shifted from models where all of the dark matter is cold; to *cold-hot dark matter* models (Primack et al., 1995), where an admixture of hot dark matter (with, say, $\sum_a m_{\nu_a} \approx 5 \text{ eV}$) smooths out the inhomogeneities caused by the cold dark matter, which would otherwise be too large at early times; to, most recently, models with a significant cosmological constant, cold dark matter, and a possible hot dark matter admixture (Primack and Gross, 1998).

The expected role of massive neutrinos in cosmology is not yet clear, and while the very high masses originally considered ($m_\nu \gtrsim 10 \text{ eV}$) are no longer favoured, so that neutrinos seem unlikely to *dominate* the dark matter, neutrinos with significant masses ($m_\nu \gtrsim 1 \text{ eV}$) are under active consideration. The information available from neutrino oscillation searches in the high-mass region is therefore of continuing cosmological interest.

Sensitivity to $\nu_e \rightarrow \nu_\tau$ and three-neutrino oscillations

While the CHORUS experiment, with its dense nuclear emulsion target, is unsuitable for tracking electrons, the NOMAD experiment was designed with electron tracking in mind, in order to be sensitive to the $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ decay (see the discussion in sections 2.3.4 and 2.4 of the next chapter). As a result the detector is sensitive to charged current interactions $\nu_e + N \rightarrow e^- + X$ of the ν_e in the neutrino beam: it may also probe oscillations $\nu_\mu \rightarrow \nu_e$ through the same interaction, although to do this a thorough understanding of the expected neutrino spectra, especially for the ν_e , is required. The relevant studies are still underway within the NOMAD collaboration.

Clearly, if oscillations $\nu_e \rightarrow \nu_\tau$ take place, the experiment is also sensitive to these, reinterpreting τ^- appearance as a signal for $\nu_e \rightarrow \nu_\tau$ rather than $\nu_\mu \rightarrow \nu_\tau$ oscillation. At the time of the NOMAD proposal, $\nu_e \rightarrow \nu_\tau$ oscillations were of particular interest due to the 17 keV neutrino claim, which was

most naturally interpreted as a sign that the weak eigenstate ν_e possessed some admixture of the heaviest mass eigenstate ν_3 , which presumably dominated the ν_τ . Accordingly the test of this claim, which we listed at (4) above, was included as a motivation for the experiment in the NOMAD proposal (Astier et al., 1991a). While this specific interest in $\nu_e \rightarrow \nu_\tau$ oscillations has since evaporated, the present experimental limits are weak: the CCFR collaboration find $\sin^2 2\theta_{e\tau} < 0.21$ at large Δm^2 in a ν_e disappearance search (Naples et al., 1998), while the $\bar{\nu}_e$ disappearance search at the Bugey reactor sets a limit on $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$ oscillations, which may reasonably be transferred to the $\nu_e \rightarrow \nu_\tau$ case, $\sin^2 2\theta < 0.15$ at large Δm^2 (Achkar et al., 1995, all limits at 90% confidence). The analysis of the NOMAD data for this hypothesis, when complete, is expected to be more than twice as sensitive in $\sin^2 2\theta_{e\tau}$.

At the simplest level, an observation or limit on τ^- appearance may be reinterpreted as an $\nu_e \rightarrow \nu_\tau$ oscillation signal by considering the ν_e rather than the ν_μ flux as the source of the ν_τ . More realistically, if $\nu_e \rightarrow \nu_\tau$ oscillations are being considered we should perform an analysis of the general case where mixing between all three neutrino species is considered. This is a much more difficult task than a two-neutrino mixing analysis as the number of important model parameters is five, rather than two; a precise understanding of the beam, in order to properly treat the (small) ν_e component, is furthermore required.

We noted in section 1.4.2 above that where all three mixing angles are small, it turns out that the formulae for two-neutrino oscillations apply to a good approximation, and this principle may be invoked to justify the use of the two-neutrino model up to this point. Another case where this model may be used as a guide is a hierarchy of neutrino masses $m_3 \gg m_2 \gg m_1$ so that two of the mass-squared differences are very close in value while the third is much smaller, $\Delta m_{32}^2 \approx \Delta m_{31}^2 \gg \Delta m_{12}^2$. In this case, the equation for two-neutrino oscillation probability, (1.36) above, is replaced by

$$P(\nu_a \rightarrow \nu_b; L) \approx |2U_{a3}U_{b3}|^2 \sin^2 \left(\frac{\Delta m_{32}^2}{4E} \right) \quad (1.45)$$

when an experiment is sensitive to the larger mass splitting $\Delta m_{32}^2 \approx \Delta m_{31}^2$ but not the smaller splitting Δm_{12}^2 . For τ^- appearance due principally to $\nu_\mu \rightarrow \nu_\tau$ oscillations, the reinterpretation of results due to (1.45) would be modest, with limits being set on the product of matrix elements $U_{\mu 3}U_{\tau 3}$, rather than on $\sin 2\theta_{\mu\tau}$. Oscillations $\nu_e \rightarrow \nu_\tau$ and $\nu_\mu \rightarrow \nu_e$ however would be governed by the mass splitting Δm_{32}^2 , rather than Δm_{31}^2 and Δm_{12}^2 as would be expected from the two-neutrino oscillation formula. These additional channels must therefore be treated with some care.

1.4.5 Recent evidence for neutrino oscillation

Since the beginning of NOMAD data-taking two experiments have interpreted their results as evidence of neutrino oscillations. The first of these was the somewhat unexpected positive result from the LSND experiment, using as its neutrino source an 800 MeV proton beam incident on a water target followed by a beam stop. The principal source of neutrinos is the decay at rest of positive pions

$$\pi^+ \longrightarrow \nu_\mu + \mu^+ \quad (1.46)$$

followed by muon decay

$$\mu^+ \longrightarrow \bar{\nu}_\mu + \nu_e + e^+ \quad (1.47)$$

in which $\bar{\nu}_e$ are not produced: the experiment is in the first instance a $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ appearance search. Despite the notorious initial confusion concerning candidate events (Athanasopoulos et al., 1995; Hill, 1995), the LSND collaboration have seen a persistent $\bar{\nu}_e$ signal above the expected background, corresponding to a $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation probability of $(0.31 \pm 0.12 \pm 0.05)\%$ (Athanasopoulos et al., 1996). Significantly, they have also recently published a $\nu_\mu \rightarrow \nu_e$ analysis, with the ν_μ produced by the minority of the produced π^+ which decay in flight, which finds a consistent oscillation probability, despite the rather different spectrum and background processes in that channel (Athanasopoulos et al., 1998). The confidence intervals in the $(\sin^2 2\theta, \Delta m^2)$ plane for both analyses are shown in figure 1.8. A search for $\nu_\mu \rightarrow \nu_e$ oscillations in the NOMAD experiment, sensitive to some (but not all) of the parameter space shown, is still in progress.

Less controversial was the recent claim by the Super-Kamiokande collaboration to have confirmed the neutrino oscillation interpretation of the longstanding “atmospheric neutrino problem” (Fukuda et al., 1998a). Cosmic ray interactions in the upper atmosphere initiate hadronic showers in which the pion decay sequence (1.46), (1.47) shown above, together with its charge-conjugate, produces neutrinos with a ratio of $(\nu_\mu + \bar{\nu}_\mu)$ to $(\nu_e + \bar{\nu}_e)$ fluxes of approximately two. Some experiments sensitive to the interactions of the different species have consistently seen a lower ratio: results are typically quoted in terms of the double ratio

$$R \stackrel{\text{def}}{=} \frac{(\nu_\mu : \nu_e)_{\text{data}}}{(\nu_\mu : \nu_e)_{\text{MC}}} \quad (1.48)$$

taking into account detector acceptances and Monte Carlo simulations of the cosmic ray airshower process, with values $R = 0.72 \pm 0.19_{-0.07}^{+0.05}$ seen at

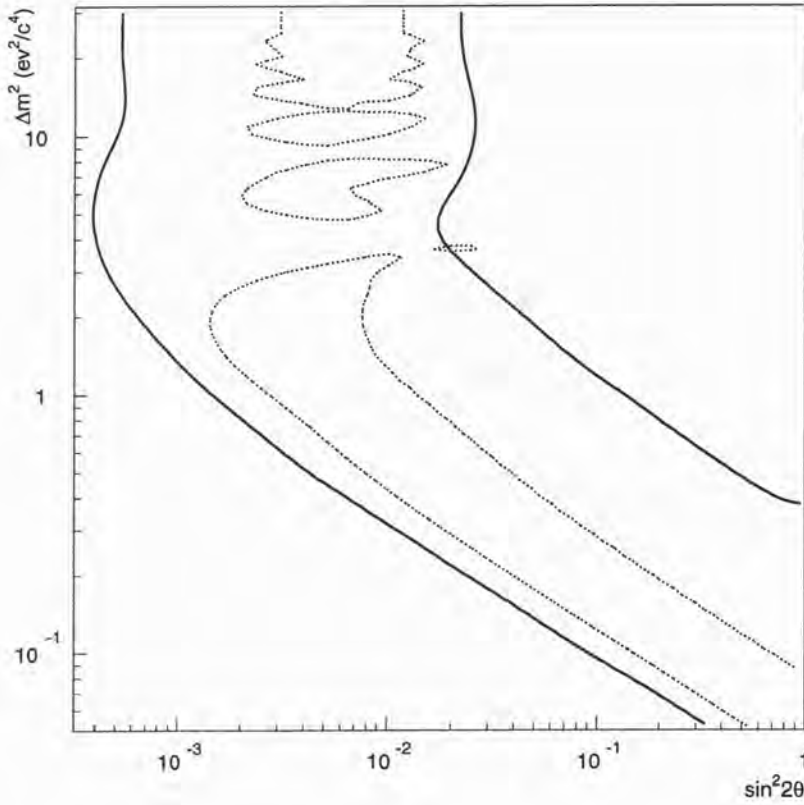


Figure 1.8: Confidence intervals in the $(\sin^2 2\theta, \Delta m^2)$ plane for the LSND experiment $\nu_\mu \rightarrow \nu_e$ decay-in-flight analysis (solid line, 95% confidence), and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ decay-at-rest analysis (dotted line, 99% confidence). (From Athanassopoulos et al., 1998, figure 33.)

Soudan 2 (Allison et al., 1997), and $R = 0.60_{-0.05}^{+0.06} \pm 0.05$ for low-energy and $R = 0.57_{-0.07}^{+0.08} \pm 0.07$ for high-energy events at Kamiokande (Fukuda et al., 1994); a deficit is also seen in the IMB 3 detector (Becker-Szendy, 1992), and found by Beier et al. (1992) to correspond to $R = 0.61 \pm 0.11 \pm 0.11$ and $R = 0.47 \pm 0.14 \pm 0.08$ for the same energy-ranges used by the Kamiokande group. Until recently, however, the statistical uncertainties on these ratios were much larger, and a number of earlier experiments reported values consistent with unity.

In the neutrino oscillation interpretation of this deficit, oscillations $\nu_\mu \rightarrow \nu_x$ deplete the muon neutrinos *en route* to the detectors. As a further test of this hypothesis, the experimental collaborations involved have considered the rate of muons passing through their detectors, originating *beneath* the




Figure 1.9: Confidence intervals in the $(\sin^2 2\theta, \Delta m^2)$ plane for the $\nu_\mu \rightarrow \nu_\tau$ two-neutrino oscillation interpretation of data from the Super-Kamiokande experiment. The 90% confidence interval from the Kamiokande experiment is also shown. (From Fukuda et al., 1998a, figure 2.)

detectors and hence presumably due to interactions $\nu_\mu + N \rightarrow \mu^- + X$ of atmospheric neutrinos. These neutrinos have travelled distances $\approx 10^4$ km from their production point in the upper atmosphere on the other side of the earth, so that a significant depletion is expected for $L_{osc} \lesssim 10^4$ km. The Super-Kamiokande collaboration have used the unprecedented size and large counting-rate of their detector to combine these approaches, forming both the double ratio (1.48) and a zenith-angle distribution for ν_e and ν_μ interactions in the detector. They find

$$R = 0.63 \pm 0.03 \pm 0.05 \quad (1.49)$$

for the low-energy and

$$R = 0.65 \pm 0.05 \pm 0.08 \quad (1.50)$$

for the high-energy events, a strong zenith-angle asymmetry for the $\nu_\mu + N \rightarrow \mu^- + X$ interactions, but no significant asymmetry for the $\nu_e + N \rightarrow e^- + X$ interactions.

The corresponding confidence intervals in $(\sin^2 2\theta, \Delta m^2)$, interpreting these results as a signal for $\nu_\mu \rightarrow \nu_\tau$ oscillations, are shown in figure 1.9. The alternative $\nu_\mu \rightarrow \nu_e$ interpretation is not favoured, as the expected zenith-angle asymmetry in the $\nu_e + N \rightarrow e^- + X$ events is not seen; the relevant oscillation parameters are in any case excluded by the results from the CHOOZ experiment (Apollonio et al., 1998, 1999). The most straightforward interpretation is therefore that the mixing between ν_μ and ν_τ is near-maximal, with a mass-squared-difference of order 10^{-3} eV^2 , corresponding to an oscillation length $L_{osc} \approx 2500 \text{ km}$ for $E_\nu \approx 1 \text{ GeV}$. Comparing this value with the baseline of the West Area Neutrino Facility beamline ($L < 1 \text{ km}$ for $E_\nu \approx 30 \text{ GeV}$) or considering the NOMAD and CHORUS sensitivity in $(\sin^2 2\theta, \Delta m^2)$ (figure 1.7), it is clear that the short baseline appearance searches would be completely insensitive to oscillations of this kind.

1.4.6 The current experimental situation

We should however note that it appears impossible to reconcile all of the present neutrino oscillation evidence with the mixing of only three neutrino species. This may be seen by comparing the mass-squared-differences of the different kinds of oscillation:

- solar neutrino oscillations, the longest established, require $\Delta m^2 = (5 - 8) \times 10^{-11} \text{ eV}^2$ (vacuum oscillations) or $\Delta m^2 \approx 10^{-5} \text{ eV}^2$ (MSW effect) for ν_e disappearance;
- the LSND collaboration find $\Delta m^2 > 10^{-1} \text{ eV}^2$ for oscillations $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ and $\nu_\mu \rightarrow \nu_e$;
- the Super-Kamiokande results require $\Delta m^2 \approx 10^{-3} \text{ eV}^2$ for ν_μ disappearance.

These Δm^2 values are each of a different order of magnitude, so that for three neutrinos, where the mass splittings obey the relation

$$\Delta m_{12}^2 + \Delta m_{23}^2 + \Delta m_{31}^2 = 0$$

we cannot accommodate all three results. Some alternative explanations of this situation are that one or more of the apparent results is in error; or that a physical effect other than neutrino oscillations is responsible (say) for the $\bar{\nu}_e$ excess in LSND; or that neutrino oscillations occur, but involve more than three species.

In the light of the result that only three light neutrinos couple to the Z^0 , this suggests the existence of at least one species of *sterile neutrino* and therefore the possibility of the active-to-sterile oscillation interpretation of the Super-Kamiokande result, which is currently open. We should note that the result of Fukuda et al. (1998a) may not be directly transferred to $\nu_\mu \rightarrow \nu_s$ oscillations since there is a different matter effect in this case: both ν_μ and ν_τ interact with matter as they pass through the earth while a sterile species ν_s does not. However an analysis of earlier Super-Kamiokande data taking this into account finds a good fit in the large-mixing region (Foot, Volkas, and Yasuda, 1998). Active-to-sterile oscillations also provide a possible solution to the solar neutrino problem, both in the vacuum oscillations case, with the same parameters as for the $\nu_e \rightarrow \nu_\mu$ interpretation; and for the MSW effect, with allowed regions similar to those shown in figure 1.5, but with only the small-angle solution allowed by all experiments (Hata and Langacker, 1997, figure 14 and discussion).

There is an additional constraint on oscillations involving sterile neutrinos due to the theory of Big Bang Nucleosynthesis, and the observed abundances of D, ^3He , ^4He and ^7Li in the universe (*e.g.* Caso et al., 1998, section 16). The effective number of neutrino species is restricted to be $N_\nu < 4.0$, and rather stringent constraints on the properties of sterile species, and also on the mixing angle and mass-squared difference for active-sterile oscillations, follow from this bound. It has however been shown that these bounds do not apply if there is a significant ($> 10^{-5}$) asymmetry between neutrinos and antineutrinos in the early universe: and, remarkably, that active-sterile oscillations can generate just such an anomaly (Foot and Volkas, 1995, 1997; Foot, Thomson, and Volkas, 1995).

The present situation is thus somewhat open. Once the possibility of sterile neutrinos is admitted, it becomes unclear whether the solar or the atmospheric neutrino problems are to be interpreted as due to active to sterile oscillations—or whether *both* are due to such oscillations—without further experimental tests. Accordingly, a number of new experiments designed to discriminate between various neutrino oscillation scenarios are in construction and development. As regards the existing short-baseline searches, $\nu_\mu \rightarrow \nu_\tau$ oscillations within the sensitivity of NOMAD and CHORUS remain an active possibility, so that the results of these experiments will help to constrain any neutrino-mixing solution to the world's neutrino data.

Chapter 2

The NOMAD experiment

2.1 Introduction

The CERN experiment WA-96, the “NOMAD” experiment, is a short-baseline ν_τ appearance search. One of two such experiments which operated in the CERN West Area from 1994 to 1998, its aim is to detect the appearance of tau neutrinos produced by oscillations

$$\nu_\mu \rightarrow \nu_\tau$$

or

$$\nu_e \rightarrow \nu_\tau,$$

through the charged current interaction

$$\nu_\tau + N \rightarrow \tau^- + X$$

in the material of an active target.

The neutrino beam to the West Area is produced by the West Area Neutrino Facility (WANF), which is described in section 2.2. The beam is predominantly ν_μ , with a contamination of $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$, but negligible ν_τ , so that its principal purpose is to provide sensitivity to oscillations $\nu_\mu \rightarrow \nu_\tau$; the median neutrino energy is well above the tau production threshold, so that for tau neutrinos produced by oscillation, the charged current process yielding the tau lepton is allowed.

The CHORUS and NOMAD experiments which share this beam rely on complementary methods to detect the short-lived tau: the observation of a short track followed by a decay kink in CHORUS (Eskut et al., 1997), and the identification of tau-decay events based on their kinematic properties in NOMAD. This kinematic method is presented in outline in section 2.3, where

the various decay modes of the tau and the main experimental backgrounds are also briefly discussed. The detector, which was designed with sensitivity to the electronic decay

$$\tau^- \longrightarrow \nu_\tau \bar{\nu}_e e^-$$

in mind, is then described in some detail in section 2.4. Experimental triggering, data acquisition and the event reconstruction code are presented in outline in section 2.5.

This present work is concerned with the search for the so-called three-prong decays

$$\tau^- \longrightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n \geq 0),$$

which provide a rich set of kinematic variables, and quantities sensitive to the structure of the decay product, as potential tools for analysis. These will be discussed in detail in chapters 3 and 4 respectively.

2.2 The West Area Neutrino Facility

Neutrinos for the CHORUS and NOMAD experiments are produced in CERN's West Area Neutrino Facility (WANF), which first operated in 1976 and has undergone occasional development, including substantial refurbishment in 1992-93 to suit the requirements of the present experiments (Acquistapace et al., 1995). In its current form, a 450 GeV proton beam from the Super Proton Synchrotron (SPS) interacts with a beryllium target, producing secondary particles which are focussed by a system of magnetic lenses, allowed to decay—producing a tertiary beam of neutrinos—and then removed from the beamline by extensive earth and iron shielding.

The elements of the beamline and its operation are discussed in more detail in section 2.2.1, while the proton beam from the SPS which drives the facility is briefly described in section 2.2.2. The resulting neutrino spectrum is then discussed in section 2.2.3.

2.2.1 The elements of the WANF beamline

The beamline of the West Area Neutrino Facility is shown in schematic form in figure 2.1. Protons of momentum 450 GeV are extracted from the SPS onto the “T9” target station, consisting of 11 beryllium rods, each 3 mm diameter and 10 cm in length, aligned with the proton beam axis and separated by 9 cm air gaps. Secondary particles produced in the p-Be interactions then pass through a pair of coaxial toroidal magnetic lenses: the *horn*, so-called because of the cone-shaped surface of its inner conductor, and the *reflector*.

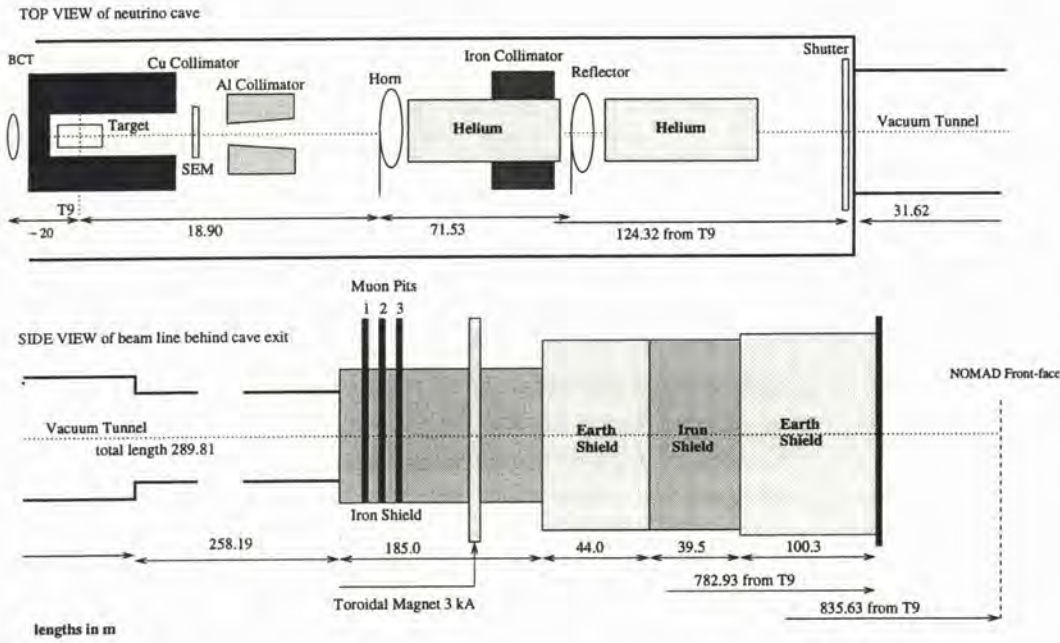


Figure 2.1: Schematic layout of the WANF beam line, with the main elements shown (see the discussion in the text). Not to scale. (Taken from Altegoer et al., 1998a, figure 3.)

The purpose of these lenses is to focus the positive secondary particles into a parallel beam, while defocussing the negative secondaries: since the decays of the positive (negative) particles yield neutrinos (anti-neutrinos), this arrangement increases the flux of neutrinos in the experimental area, and at the same time reduces the $\bar{\nu}/\nu$ ratio in the beam.

The hollow conical magnetic volume of the horn was designed so that positive particles of a given momentum—in this case $50 \text{ GeV}/c$ —emerging at a range of angles from the target are focussed into a parallel beam. The magnetic volume of the reflector was then designed to extend the focussing to a broader range of momenta. Particles travelling parallel to and close to the beam axis pass through the reflector undeflected; stiff high-momentum tracks are bent further towards the beam axis, while low-momentum tracks, which are over-focussed by the horn and would otherwise escape from the beamline, are bent back towards the beam axis. The result is an approximately parallel beam of positive particles across a range of momenta, as shown schematically in figure 2.2.

A system of collimators suppresses particles which are outside the acceptance of the horn, as well as those which have been bent outside the




Figure 2.2: The principle of the horn and reflector magnetic lens system. Positive charged secondaries of momentum $50 \text{ GeV}/c$ are focussed into a parallel beam by the horn, while the reflector improves the focussing of higher momentum tracks and corrects for the over-focussing of low momentum tracks. Positive tracks only are shown; not to scale. (Adapted from Acquistapace et al. (1995), figure [c9] by Boyd (1998).)

acceptance of the reflector. This ensures that neutrino-producing decays are dominated by the focussed particles, further reducing the $\bar{\nu}/\nu$ ratio in the beam. The focussed particles pass through two helium tubes along the beam axis, of length 63 m (between the horn and the reflector) and 18 m (downstream of the reflector) to reduce particle absorption by the ambient gas.

The particles then pass into a 290 m vacuum tunnel, where they undergo decays

$$\pi^+ \rightarrow \nu_\mu + \mu^+ \quad (2.1)$$

$$K^+ \rightarrow \begin{cases} \nu_\mu + \mu^+ & 63.5\% \\ \nu_\mu + \mu^+ + \pi^0 & 3.2\% \\ \nu_e + e^+ + \pi^0 & 4.8\% \end{cases} \quad (2.2)$$

$$K_L^0 \rightarrow \begin{cases} \nu_\mu + \mu^+ + \pi^- & 27.2\% \\ \nu_e + e^+ + \pi^- & 38.8\% \end{cases} \quad (2.3)$$

and the appropriate charge-conjugates, producing the neutrino beam.

Undecayed particles are allowed to range out in the extended iron and earth shielding (totals depths of 217 m and 144 m respectively) the first stage of which also serves as a beam dump for any beam protons which have not interacted with the target (see the figure). A toroidal magnet is placed inside the first stage of the shielding to focus any positive particles then-surviving into the core of the subsequent iron and earth shield, and to bend surviving negative particles away from the detectors.

Monitoring of the WANF beamline is provided by two beam current trans-

formers (BCT) upstream of the target station, to measure the flux of incident protons; secondary emission monitors (SEM) downstream of the target station, to measure the rate of production of secondary particles; and a system of silicon diode detectors (Heijne, 1983) to measure the accompanying muon flux, and its spatial distribution, at three different depths in the iron shield (the detectors are placed in “pits” at depths of 10.4 m, 30.9 m and 51.1 m of iron respectively). These detectors are cross-calibrated using a system of moving reference detectors, and an absolute calibration is periodically carried out using sheets of nuclear emulsion. The beam current transformers and silicon detectors then provide independent measurements of the flux along the beamline.

2.2.2 The SPS proton beam on the T9 target

During the operation of the WANF for CHORUS and NOMAD, the Super Proton Synchrotron was run on a 14.4 s cycle, in which the majority of the protons in the machine were devoted to the Neutrino Facility. The early part of the cycle—the one relevant to the WANF—is shown schematically in figure 2.3.

The SPS was filled with protons in two injections from the Proton Synchrotron, and the circulating protons then accelerated to the operating energy of 450 GeV (figures 2.3(a) and (b) respectively). Protons were then extracted onto the neutrino target T9 in two spills of ≈ 6 ms length separated by 2.6 seconds, the so-called *flat-top*, during which a proton energy of 450 GeV is maintained by the accelerator. During the flat-top, protons were directed to other targets continuously by slow extraction. Some of these targets feed test-beams in the area upstream of the NOMAD, and muons from these beams traversed the detector throughout this period.

Current to the horn and reflector was supplied in pulses co-ordinated with the SPS extraction onto the T9 target (figure 2.3(c)). Also co-ordinated with the SPS cycle were the main data-taking *gates* of the NOMAD experiment, shown in figure 2.3(d). Two *neutrino gates* of 10 ms each, centred on the two neutrino spills, defined the sensitive period of the NOMAD triggering system for neutrino interactions. The 2 second *muon gate* during the flat-top provided for triggering on throughgoing muons for subdetector calibration purposes *etc.*, as discussed in section 2.4 below where the various subdetectors are described. Other data-taking gates were also defined as needed during the remainder of the SPS cycle, which does not otherwise concern us here. (The NOMAD trigger system is more fully described in section 2.5.1 below.)

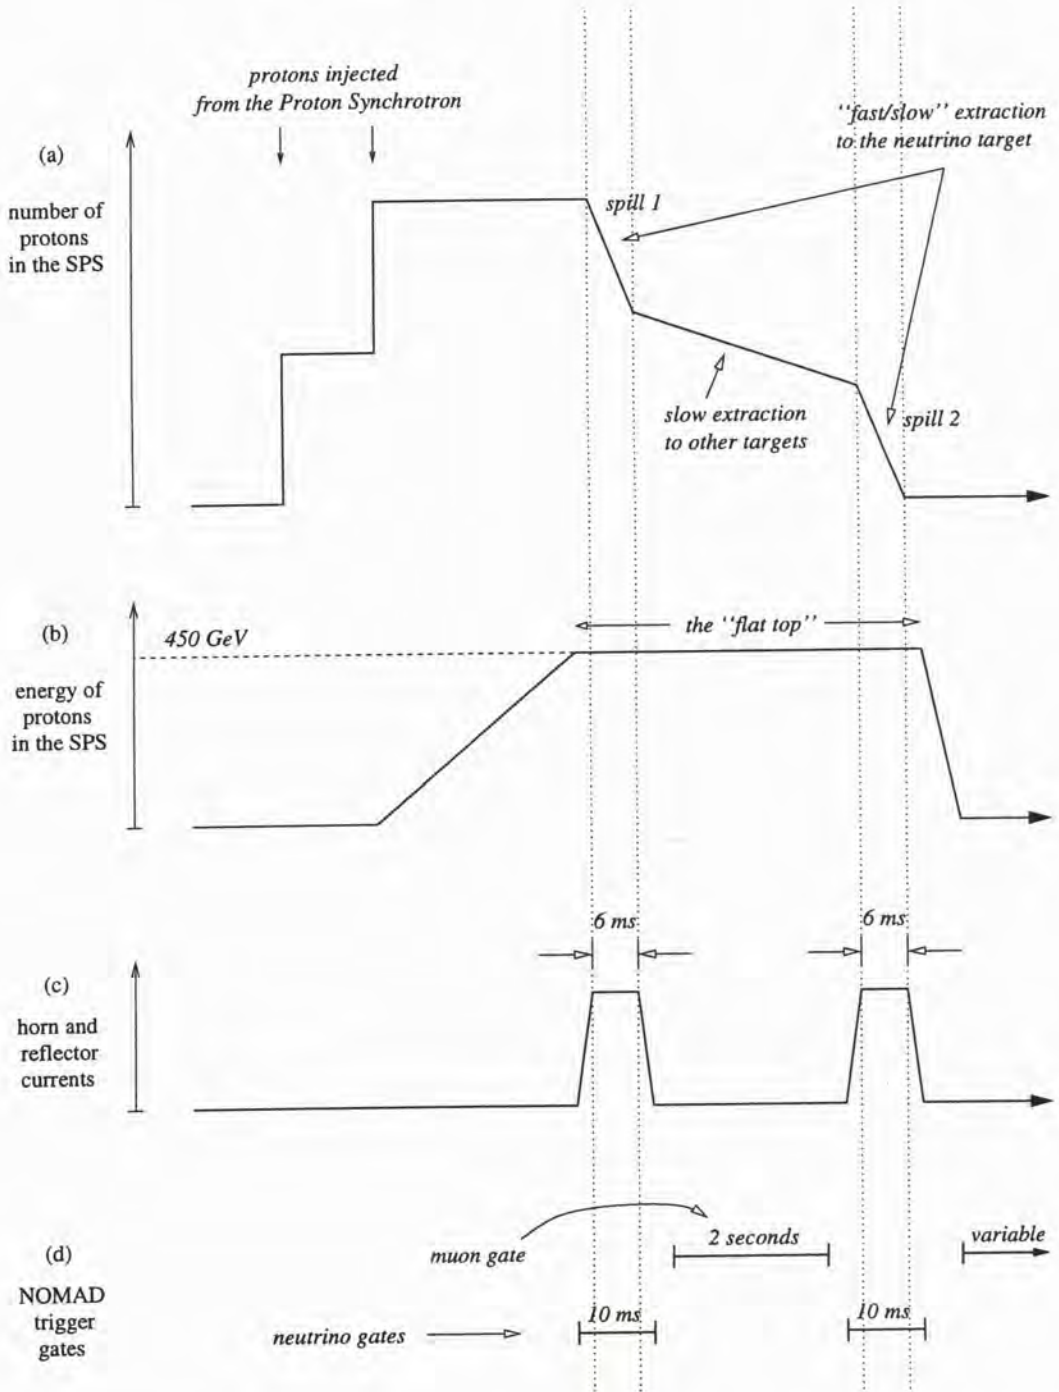


Figure 2.3: Schematic diagrams of the SPS cycle for the West Area Neutrino Facility (not to scale). See the text.

2.2.3 The expected neutrino spectrum at the NOMAD

A detailed GEANT-based simulation of the WANF beamline, prepared by and for the NOMAD collaboration, has been used to calculate the expected spectra of the various neutrino species at the detector (Ravndal et al., 1993); these spectra are used as the first stage of the Monte Carlo event generation process for the experiment. In addition to comparisons with the NOMAD data we should note that measurements from the silicon detectors in the pits in the iron shield provide an additional control on the simulation. The beamline simulation is undergoing continuous refinement, including incorporation of the results of the dedicated CERN experiment NA-56 (SPY) measuring the production rates of pions and kaons in proton collisions on a beryllium target (Ambrosini et al., 1998a,b, 1999).

Secondary particle production in interactions in the target (and the surrounding material) is simulated and the secondaries are then transported through a detailed model of the beamline, allowed to decay as appropriate, and the decay products themselves are then traced. If we consider muon

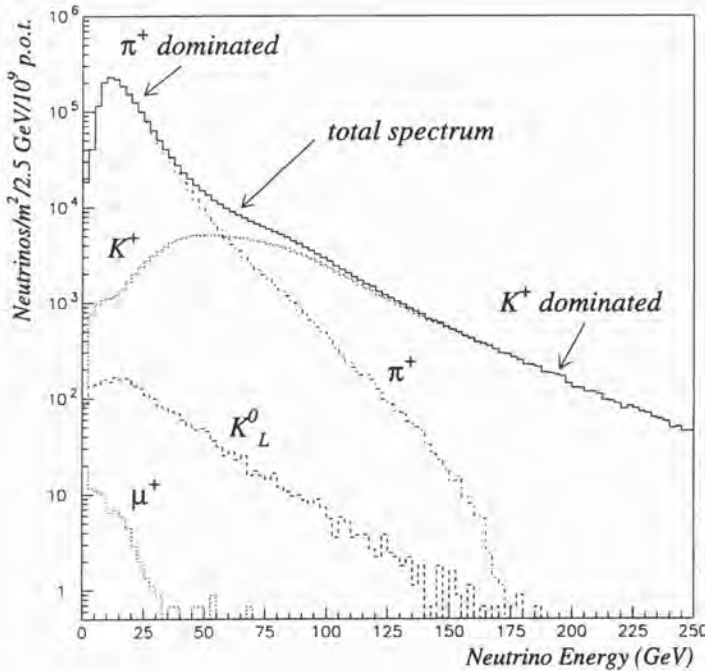


Figure 2.4: Parent particles for ν_μ passing through the NOMAD fiducial volume, as a function of energy, from the NOMAD beamline simulation (Godley, 1999).

neutrinos passing through a fiducial area of $2.6 \times 2.6 \text{ m}^2$ at the NOMAD position we find spectra of ν_μ due to the decay of π^+ , K^+ , K_L^0 and μ^+ as shown in figure 2.4. As expected from the list of decays presented above ((2.1) through (2.3)) ν_μ are principally due to π^+ and K^+ decay, with π^+ dominating the flux but K^+ dominating in the high-neutrino-energy tail due to the higher available energy in the kaon decay.

If we consider instead electron neutrinos, which are not produced in π^+ decays, the flux is much lower (only 1% of ν_μ), but the average energy is higher as the ν_e are produced in K^+ and K_L^0 decays. The predicted relative fluxes of all four species $\nu_\mu, \bar{\nu}_\mu, \nu_e, \bar{\nu}_e$ at the NOMAD are shown in table 2.1, together with the average neutrino energies; the expected neutrino energy spectra are shown in figure 2.5, where (*e.g.*) the high energy tail of the ν_e spectrum, due to K^+ decays, and the absence of a significant low- E_ν component, are apparent.

Species	$\langle E_\nu(\text{GeV}) \rangle$	Relative flux	Number of $\nu/\bar{\nu}$ CC interactions
ν_μ	23.6	1.00	1.15×10^6
$\bar{\nu}_\mu$	22.7	0.07	0.39×10^5
ν_e	37.0	0.01	0.17×10^5
$\bar{\nu}_e$	33.2	3×10^{-3}	0.22×10^4

Table 2.1: Predicted relative abundances and average neutrino energies at the NOMAD for a $2.6 \times 2.6 \text{ m}^2$ fiducial area. The number of charged-current (CC) interactions is normalised to 2.4×10^{19} protons on target.

Also shown in figure 2.5 are the distributions of neutrino positions in the plane transverse to the neutrino beam, where the effect of the focussing of positive secondaries (section 2.2.1) may be clearly seen. The ν_μ , principally due to decays of well-focussed π^+ and K^+ , are strongly concentrated in the core of the beam, while the $\bar{\nu}_\mu$, produced in the decay of defocussed π^- and K^- , have a relatively flat radial distribution. Electron neutrinos form an intermediate case as they are produced in the decays of focussed K^+ and un-focussed neutral K_L^0 .

More relevant for the NOMAD experiment than the neutrino flux is the expected rate of neutrino *interactions* in the detector: see the column of table 2.1 where the numbers of charged current interactions for the different species are shown (for a nominal two-year running period). The effect of the linear increase of the $\nu - N$ cross-section with energy is apparent (compare

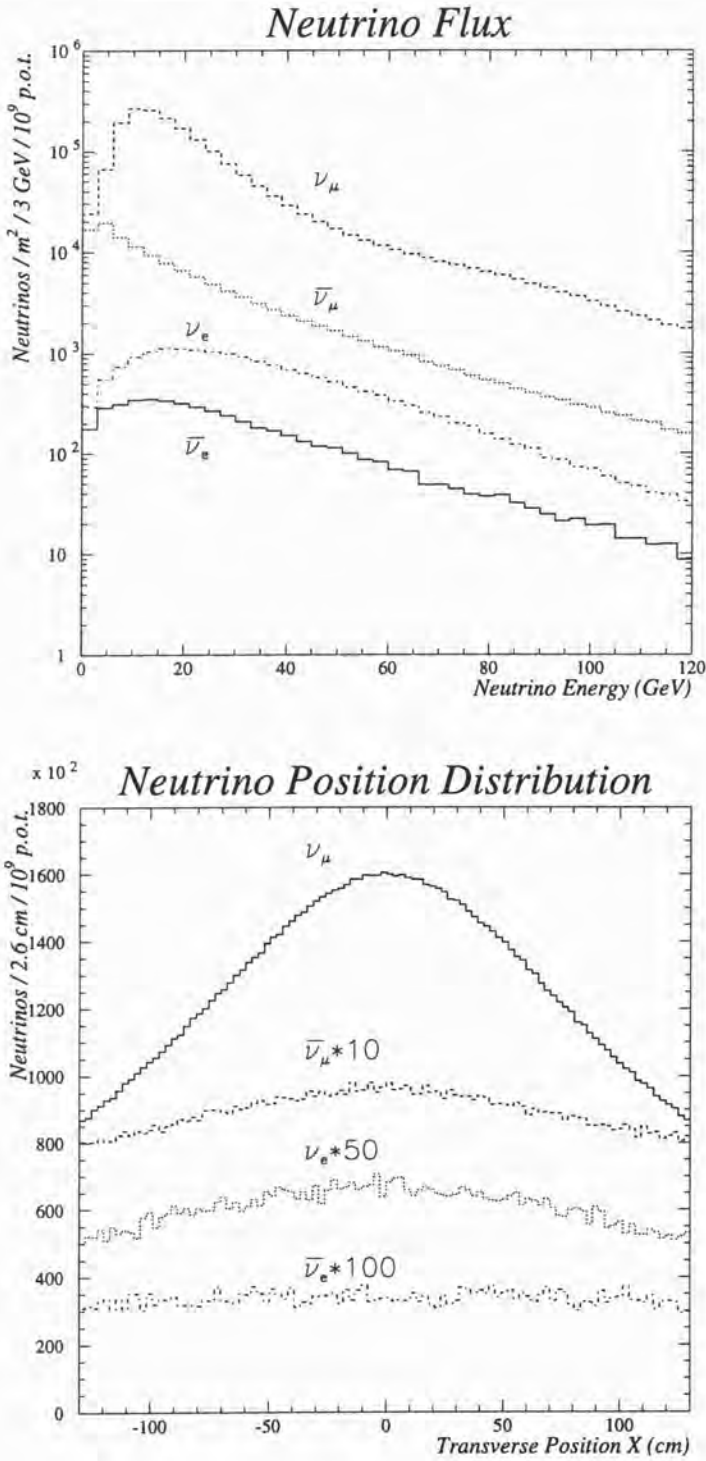



Figure 2.5: The predicted energy spectra and transverse position distributions of the four neutrino species at the NOMAD. (Taken from Altegoer et al., 1998a, figures 4 and 5.)

the relative flux to the relative number of interactions for the energetic ν_e) as is the lower interaction cross-section on nucleons for the antineutrinos. For hypothetical ν_τ produced in oscillations $\nu_\mu(\nu_e) \rightarrow \nu_\tau$, there is an additional effect due to the tau mass, which introduces an energy-dependent suppression of the interaction cross-section. The average suppression factor has been estimated at

$$\frac{\sigma_\tau}{\sigma_\mu} = 0.48 \quad (2.4)$$

for analyses sensitive to deep-inelastic interactions $\nu_\tau + N \rightarrow \tau^- + X$ in the NOMAD, after weighting for the neutrino energy spectrum (Astier et al., 1999a). We should note that the choice of the 450 GeV energy for SPS protons extracted onto the neutrino target—the maximum energy available—was dictated by the need to produce muon neutrinos of as high an energy as possible, in order to avoid too small a factor σ_τ/σ_μ . Several changes to the WANF beamline, such as adjustment of the horn and reflector positions *etc.*, were motivated by this same need to harden the neutrino spectrum.

So far we have neglected the production of ν_τ and $\bar{\nu}_\tau$ along the beamline, which provides an irreducible background to the neutrino oscillation search. The only significant mechanism for the production of the tau neutrinos, other than the hypothetical $\nu_\mu(\nu_e) \rightarrow \nu_\tau$ oscillations, is the production of charmed strange mesons in proton-nucleon interactions



followed by the decays

$$D_s^- \rightarrow \bar{\nu}_\tau + \tau^- \quad (2.5)$$

and

$$\tau^- \rightarrow \nu_\tau + X' \quad (2.6)$$

with the charge conjugate decay sequence beginning with the D_s^+ occurring at a similar rate. Interactions of beam protons with both the beryllium target and the iron beam dump must be considered. The rate of ν_τ charged current (CC) interactions been estimated by Van de Vyver and Zucchelli (1997) for the CHORUS location and fiducial volume at

$$\frac{\nu_\tau \text{ CC}}{\nu_\mu \text{ CC}} \approx 3.3 \times 10^{-6}$$

using an empirical parametrisation of the D_s production rate; Gonzalez-Garcia and Gomez-Cadenas (1997) find an equivalent rate using calculations based on the Quark Gluon String Model, and estimate the contamination at the NOMAD to be

$$\frac{\nu_\tau \text{ CC}}{\nu_\mu \text{ CC}} \cong 4.5 - 4.9 \times 10^{-6}.$$

This higher rate is due to the larger fiducial volume of the NOMAD, and the fact that the prompt ν_τ , produced in the decays of unfocussed mesons, are less concentrated in the core of the beam than the ν_μ .

Such a rate corresponds to an expectation of less than one $\nu_\tau/\bar{\nu}_\tau$ charged-current interaction over the whole NOMAD data-taking run. As the efficiencies of the various NOMAD tau decay analyses for tau decays are at the percent level, the resulting background to the oscillation search is negligible.

2.3 The strategy of the τ search in the NOMAD experiment

We now turn from the neutrino facility as a whole to the NOMAD experiment, and the search for ν_τ appearance through the $\nu_\tau + N \rightarrow \tau^- + X$ reaction. The technique of identifying tau decays via their kinematic signature, the approach chosen for the τ search in the NOMAD experiment, is outlined in section 2.3.1. The various decay modes of the tau are then reviewed in section 2.3.2. After considering the experimental backgrounds (section 2.3.3), the choice of decay modes for analysis is briefly discussed in section 2.3.4.

2.3.1 The kinematic method for detection of τ decay

The tau lepton is short-lived, with a proper decay length of only $87 \mu\text{m}$, so that even energetic taus produced in interactions of neutrinos from the WANF would decay over lengths of order 1 mm. The reconstruction of so short a track requires a detector system with extremely fine spatial resolution, such as the nuclear emulsion active target of CHORUS, or at least a high resolution tracking system (silicon vertex detectors *etc.*) to recover the decay vertex or the appropriate impact parameter of the decay product.

For the NOMAD experiment a complementary approach was chosen, relying instead on the distinctive kinematic properties of tau decay events. Tau neutrino interactions

$$\nu_\tau + N \rightarrow \tau^- + X$$

followed by tau decay

$$\tau^- \longrightarrow \nu_\tau + X'$$

exhibit the following features:

- net event momentum in the plane transverse to the neutrino beam, due to an energetic ν_τ in the final state;
- a visible decay product X' which reflects the kinematics of the tau itself, *i.e.* carrying a significant fraction of the event energy, and having opposite azimuth to the hadronic system X ;
- constraints on the visible decay product, and the missing transverse momentum (reflecting the final state neutrino(s)), due to the tau mass.

A sufficiently accurate reconstruction of the momenta of charged tracks in the final state, together with electromagnetic calorimetry to recover the neutral component of the event, should allow each of these features to be observed. The details of the analysis of course depend on the specific tau decay in question, and the properties and rate of the other neutrino interactions which act as a background. We discuss these in the following sections in turn.

2.3.2 Tau decay modes

Due to its large mass the tau lepton has a rich set of decay modes, and in an experiment with a sufficiently large τ sample (such as the experiments at the Large Electron Positron collider *etc.*) each of these modes would need to be considered. In our case, since even at the previous experimental limit for $\nu_\mu \rightarrow \nu_\tau$ oscillations the τ sample would be small (section 1.4.4 above), we may restrict our attention to decay modes of high branching fraction (B.F.) and/or those which offer the prospect of very efficient identification along with suppression of backgrounds.

Leptonic decays

Most important are the leptonic decay modes

$$\tau \longrightarrow \begin{cases} \nu_\tau \bar{\nu}_e e^- & \text{B.F. 17.8\%} \\ \nu_\tau \bar{\nu}_\mu \mu^- & \text{B.F. 17.4\%} \end{cases} \quad (2.7)$$

which in addition to having large branching fractions, provide a tagged decay product (e/μ) and an enhanced decay signature in the plane transverse to the neutrino beam due to the additional final state neutrino. In these decays

the neutrinos $\nu_\tau \bar{\nu}_{e/\mu}$ carry the majority of the energy of the tau lepton, so that the “missing” momentum in the transverse plane closely reflects the transverse momentum of the τ itself. The charged lepton furthermore will in typical cases emerge at some angle to the direction of motion of the τ so that a separation of tau decay events from the charged current backgrounds ($\nu_e + N \rightarrow e^- + X$ and $\nu_\mu + N \rightarrow \mu^- + X$ respectively) may take place based on the measurement of the e/μ direction, and not on its energy alone (see the following sections).

One-prong decays

The remaining branching fraction is dominated by the “one-prong” decays

$$\tau \longrightarrow \begin{cases} \nu_\tau \pi^- & \text{B.F. 11.1\%} \\ \nu_\tau \pi^- \pi^0 & \text{B.F. 25.3\%} \\ \nu_\tau \pi^- \pi^0 \pi^0 & \text{B.F. 9.2\%} \\ \nu_\tau \pi^- \pi^0 \pi^0 \pi^0 & \text{B.F. 1.1\%} \end{cases} \quad (2.8)$$

so-called because there is a single track in the decay product. The neutral-bearing modes together make up the majority of the branching fraction for decays of this type, but require the reconstruction of the π^0 (s) from calorimeter clusters and/or photon conversions for the decay-product momentum to be properly recovered. The practice in the NOMAD experiment has been to perform a dedicated analysis for $\tau^- \rightarrow \nu_\tau \pi^- \pi^0$ which contributes a full quarter of all tau decays, and to carry out in addition a consolidated one-prong analysis, optimised for the $\tau^- \rightarrow \nu_\tau \pi^-$ decay but having some efficiency for the other modes. When minor modes involving kaons *etc.* are included the one-prong decays contribute a total branching fraction of 49.5%.

The $\tau^- \rightarrow \nu_\tau \pi^-$ decay has some of the large-missing-momentum characteristics of the leptonic decays due to a τ spin/polarisation effect: as the τ^- and ν_τ are both spin- $\frac{1}{2}$ while the pion is spin-0, there is a preference for the (negative helicity) neutrino to be emitted in the direction opposite to the tau spin. For energetic taus from $\nu_\tau + N \rightarrow \tau^- + X$, the spin is itself partially aligned opposite to the laboratory momentum \vec{p}_τ (see section 4.5 below) so that the ν_τ is emitted in preference along the τ direction. The partitioning of lab-frame energy between the ν_τ and the pion therefore favours the neutrino.

Three-prong decays

By contrast with the preceding channels, the visible decay product tends to take a large fraction of the tau energy in the “three-prong” decays

$$\tau \longrightarrow \begin{cases} \nu_\tau \pi^- \pi^+ \pi^- & \text{B.F. 9.6\%} \\ \nu_\tau \pi^- \pi^+ \pi^- \pi^0 & \text{B.F. 4.4\%} \\ \nu_\tau \pi^- \pi^+ \pi^- \pi^0 \pi^0 & \text{B.F. 0.5\%} \end{cases} \quad (2.9)$$

where the multiparticle final state moreover offers a rich set of variables sensitive to the decay structure. The most important of these is the invariant mass of the three charged pions $M_{inv}^{3\pi}$, which has a special meaning in the case of the 3π decay, being dominated by the decay to the pseudovector meson a_1^- :

$$\tau \rightarrow \nu_\tau a_1^- \quad (2.10)$$

followed by

$$a_1^- \rightarrow \rho^0 \pi^- \quad (2.11)$$

and

$$\rho^0 \rightarrow \pi^+ \pi^-. \quad (2.12)$$

In this case the 3π mass recovers the large a_1^- mass ($m_{a_1} = 1.23 \pm 0.04 \text{ GeV}/c^2$) while the invariant masses of two-pion combinations are sensitive to the ρ substructure of the a_1 decay.

As with the one-prong decays, the practice to date has been to perform a consolidated analysis on all of the three-prong modes, which (including minor modes involving kaons *etc.*) contribute a total branching fraction of 15.2%. The use of the 3π mass in the analysis is not inconsistent with this practice as the width of the a_1^- is large (between 0.2 and 0.6 GeV/c^2), so that the mass distributions of 3π systems from $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- \pi^0$ and $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- \pi^0 \pi^0$ decays overlap the distribution due to the a_1^- .

Special modes

The analysis effort within the NOMAD experiment has been concentrated on the decay modes described above, *i.e.* the modes (or groupings-of-modes such as the one-prong decays) with large branching fraction. We should however note the existence of decay modes with branching fractions at the one percent level, and very distinctive decay products:

1. The decays involving a K_S^0 meson

$$\tau^- \longrightarrow \begin{cases} \nu_\tau K_S^0 \pi^- + n\pi^0 & (n \geq 0) \\ \nu_\tau K_S^0 K^- + n\pi^0 & (n \geq 0) \end{cases}$$

contribute a total branching fraction of 0.8%, with a strong tag on one decay product (the K_S^0) and, in the case of the $\nu_\tau K_S^0 \pi^-$ mode, a decay through the moderate-width resonance $K^*(892)^-$.

2. The decays

$$\tau^- \longrightarrow \begin{cases} \nu_\tau \omega \pi^- \\ \nu_\tau \omega K^- \end{cases}$$

followed by

$$\omega \rightarrow \pi^+ \pi^- \pi^0$$

have a total branching fraction of 1.7% (rising to 2.1% if decays with additional π^0 are considered) and while lacking a tag, proceed through an extremely narrow resonance ($\Gamma_\omega = 0.0084$ GeV).

In both cases the branching fraction is small compared to the decay modes considered above, but the decay product offers the prospect of efficient identification in the case of tau decay coupled with a powerful rejection of background processes. Whether the rejection of the background in either case is powerful enough to offset the modest branching fraction, requires a dedicated study which has not been carried out to date.

2.3.3 Experimental backgrounds

The decisive factor in choosing tau decay modes for analysis is the potential background from other neutrino interactions. We will briefly consider the principal backgrounds to each category of tau decays in turn.

Leptonic decays

For each of the leptonic decay modes

$$\tau \rightarrow \nu_\tau \bar{\nu}_l l^-$$

the principal background is the charged current interaction of the corresponding neutrino

$$\nu_l + N \rightarrow l^- + X$$

from the WANF neutrino beam. In this case the lepton l^- is interpreted as the tau decay product and the hadronic system X is interpreted straightforwardly. Since charged-current events are balanced in momentum perpendicular to the neutrino beam, up to Fermi motion effects, there must be a significant mis-reconstruction of the event—due to missed particles, momentum resolution effects, or bremsstrahlung in the case of an electron—to mimic the signature of the final state neutrinos $\nu_\tau \bar{\nu}_l$ from tau decay.

Consulting table 2.1 above we see immediately that this will be a much more serious background for the muonic decay of the tau, since the rate of ν_μ charged-current events exceeds that of the ν_e by two orders of magnitude. There is however an additional background to the electronic decay from ν_μ neutral current interactions

$$\nu_\mu + N \rightarrow \nu_\mu + X$$

not only from electrons in the final state X (such as from Dalitz decay of energetic π^0) but more seriously from hadrons π^- , K^- if they are mis-identified as electrons. In this case the final-state neutrino ν_μ mimics the expected $\nu_\tau \bar{\nu}_e$ signature from the tau decay; but to form a tau-decay-like event, the mis-identified hadron must also be well-isolated from the rest of the final state, and together with the missing transverse momentum due to the neutrino *etc.*) it must be consistent with belonging to a tau decay product $\nu_\tau \bar{\nu}_e e^-$.

One-prong decays

Neutral current events of this type also provide the principal background to the one-prong decays

$$\tau^- \longrightarrow \nu_\tau \pi^- + n\pi^0 (n \geq 0)$$

in the same manner, with the added disadvantage that no particle mis-identification is required for the hadron to be interpreted as the decay-product π^- .

There is a further background to the one-prong analysis from the charged current interactions

$$\nu_l + N \rightarrow l^- + X$$

in cases where the lepton is mis-identified as a π^- . As with the ν_μ CC background to the electronic decay, however, there must be significant mis-measurement of the event in order to mimic the final-state neutrino from the tau decay.

Three-prong decays

The backgrounds for the three-prong decays are similar in origin to those for the one-prong. The severity of these backgrounds is however more difficult to assess because of a number of competing effects. In the case of the neutral current background,

1. as the decay product contains three visible tracks $\pi^-\pi^+\pi^-$, a single isolated energetic π^- in the hadronic system cannot mimic a tau decay—instead the hadronic system has to possess two energetic groups relatively isolated from each other; but
2. there is a “combinatoric” background in high-multiplicity events, where of the many possible 3π groupings, one may have the right properties to pass kinematic (energetic and isolated 3π) and tau structure (*e.g.* three-pion mass) cuts.

Similarly for the charged-current backgrounds, a mis-identified lepton does not by itself constitute a three-prong decay candidate. However for the numerous $\nu_\mu + N \rightarrow \mu^- + X$ interactions there is an additional problem in that muon identification relies in general on penetrating power, so that low-momentum muons are difficult to identify. The sharing of the tau momentum among several decay products in $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n \geq 0)$ means that single decay tracks have a lower momentum on average than the pion in the one-prong decay (say), so that ν_μ CC interactions with intermediate-momentum muons form a more dangerous background to this channel.

2.3.4 Choice of decay modes for analysis

Even from the brief survey of the preceding section it is clear that the decay mode of choice for the kinematic method is the electronic decay

$$\tau^- \longrightarrow \nu_\tau \bar{\nu}_e e^-,$$

provided that electron/hadron discrimination in the detector can be made sufficiently powerful to suppress the neutral current background. In this case the remaining background is due to ν_e charged current events, which are less numerous than the corresponding ν_μ events by two orders of magnitude. This disproportion between the ν_μ and ν_e samples offers the added advantage that the charged current events

$$\nu_\mu + N \rightarrow \mu^- + X$$

may be used as a control sample, much larger than the $\nu_e + N \rightarrow e^- + X$ sample to be understood, to check the Monte Carlo simulation of the hadronic system X , detector resolution effects *etc.*.

The decision was accordingly taken to design the detector based on the needs of the $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ analysis. This design, and the detector performance, are discussed in some detail in the remainder of this chapter. It should however be clear that accurate reconstruction of both tracks and neutrals, sufficient to be sensitive to the momentum lost to final state neutrinos, is a requirement for analysis of *all* of the tau decay modes, not just the electronic decay. Furthermore some of the features of a detector optimised for the electronic decay are a direct benefit to the analysis of the hadronic decay modes (one prong and three prong): for example, accurate calorimetry allows the recovery of π^0 for an analysis of the $\tau^- \rightarrow \nu_\tau \pi^- \pi^0$ decay; and electron/hadron discrimination allows the suppression of the ν_e charged-current background to all of the hadronic decay modes (a question to which we will return in chapters 5 and 6).

Analysis of the other decay modes has therefore been foreseen since the earliest stages of the experiment (Astier et al., 1991b), and results have in fact been presented for the one- and three-prong modes discussed in the previous sections (Altegoer et al., 1998c; Astier et al., 1999a). In particular, the search for the three-prong decays

$$\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n \geq 0)$$

is the subject of this work, and will be treated in detail in chapters 3 and following.


2.4 Detector design and performance

The Neutrino Oscillation MAgnetic Detector (NOMAD) was designed to be sensitive to the electronic decay of the tau, $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$, using kinematic criteria only to distinguish this decay from charged current interactions $\nu_e + N \rightarrow e^- + X$. The principal requirements for the detector were therefore


- accurate reconstruction of the full kinematics of the event, to be sensitive to the missing transverse momentum \vec{p}_T^{miss} due to final-state neutrinos;
- efficient identification of electrons together with powerful electron/hadron discrimination, to efficiently select electronic decay candidates above the more numerous $\nu_\mu + N \rightarrow \mu^- + X$ and $\nu_\mu + N \rightarrow \nu_\mu + X$ backgrounds;
- accurate electromagnetic calorimetry, to aid in electron identification and to recover the electron energy in the case of bremsstrahlung (also improving the resolution of \vec{p}_T^{miss} by reconstructing photons);
- tagging of muons to both identify and reconstruct the kinematics of the $\nu_\mu + N \rightarrow \mu^- + X$ control sample;
- a significant target mass in order to provide a large ($\approx 10^6$) sample of $\nu_\mu + N \rightarrow \mu^- + X$ interactions, and sensitivity to low mixing angles for the oscillation $\nu_\mu \rightarrow \nu_\tau$.

These requirements and the basic design of the detector were laid down by Astier et al. (1991a,b), with refinements to the proposed analysis later in the same year (Astier et al., 1991c), and to the detector design during construction and testing (NOMAD Collaboration, 1992, 1993). A description of the detector and its performance has been published by Altegoer et al. (1998a); an overview of the detector, as it appeared during the late-1995 and 1996 data-taking runs, is shown below as figure 2.6.

The various components of the NOMAD are described in the following sections (2.4.1 through 2.4.10), noting the connection between the design and the requirements of the analysis. A B_4C -silicon target, added to the detector in 1997 (Barichello et al., 1998a), is briefly discussed in section 2.4.11. The triggering, data acquisition and reconstruction procedures are summarised in section 2.5.



(a) Side view



(b) Top view

Figure 2.6: The Neutrino Oscillation Magnetic Detector (NOMAD). (Adapted from Altegoer et al., 1998a, figures 1 and 2, together with Andrew Godley.)

2.4.1 The NOMAD magnet and subdetector layout

In order to provide momentum reconstruction for charged tracks most of the subdetectors of the NOMAD are located in a large-volume ($3.5 \times 3.5 \times 7.5 \text{ m}^3$) electromagnet, as shown in figure 2.6. The magnet itself is recycled from the UA1 experiment (Barranco Luque et al., 1980; Corden et al., 1985), and produces a dipole field directed horizontally, as shown (figure 2.6(b)). The operating field strength of 0.4 T was chosen as a compromise between acceptance and momentum measurement requirements; the operation of various photomultipliers within the magnet volume (sections 2.4.5 and 2.4.8) also favours a relatively modest field.

The modular iron-plate “C” pieces forming the flux return of the magnet are uninstrumented (*c.f.* Corden et al., 1985) while the former UA1 end-cap pieces (“I” modules) have been rearranged: two modules at the front and six at the back of the detector act as support pillars for a stainless steel frame inside the magnet coil, supporting the various subdetectors against gravity and holding them in position when the magnet is opened for detector maintenance. Both front and back supports have been partially instrumented to form iron-scintillator calorimeters, the front calorimeter (FCAL, section 2.4.3) serving as a target for a high-statistics sample of neutrino interactions, principally for multi-muon studies; the instrumented back support or hadronic calorimeter (HCAL, section 2.4.9) providing additional calorimetry for hadrons and providing the first stage of the muon filter.

Proceeding along the neutrino beam axis the detector consists of

- several planes of scintillators providing a veto signal on throughgoing particles (section 2.4.2);
- the front calorimeter (section 2.4.3);

and, inside the magnet volume,

- drift chambers, whose self-supporting walls also provide the neutrino interaction target (section 2.4.4);
- two planes of scintillation counters providing the main event trigger (section 2.4.5);
- a large transition radiation detector, providing electron/hadron discrimination (section 2.4.6);
- a preshower detector and lead-glass electromagnetic calorimeter (sections 2.4.7 and 2.4.8);

and finally, outside the magnet volume,

- an hadronic calorimeter (section 2.4.9); and
- a muon detection system placed behind the back support, comprising several scintillation counters and ten large-area drift chambers, with an additional iron wall providing further suppression of non-muon tracks for a subset of the chambers (section 2.4.10).

These subdetectors are described in turn in the following sections.

2.4.2 The veto counters

The neutrino beam at the NOMAD site is accompanied by a substantial flux of muons (section 2.2.3): to suppress experimental triggers due to these and other throughgoing particles, such as cosmic rays and products of interactions in CHORUS and other material upstream of the detector, a large-area system of scintillation counters is placed at the front of the detector to act as a veto.

An area of $5.4 \times 5.0 \text{ m}^2$ is covered by the counters, whose arrangement is shown in figure 2.7; the active area of the drift chambers (*q.v.*) is also shown. The central bank of ten counters (shaded in the figure) called “ V_8 ” is mounted on the front support pillar and acts as a veto for triggers selecting neutrino interactions in the front calorimeter (*q.v.*). The remaining eight banks of counters are arranged to cover the front-face of the magnet and flux return, including the region behind the support pillar (note the longitudinal placement of veto counters in figure 2.6). A small central region, occupied by the girder connecting the front pillar and the steel frame of the central detector, cannot be instrumented: interactions occurring in this area, and in the front-face of the magnet coil itself, are rejected off-line based on the reconstruction of the event vertex.

The counters, made from NE-110 scintillator, are 21×2 cm in cross-section and either 210 cm or 300 cm in length. Most (56) are viewed at both ends by photomultipliers; the remainder (3 short counters abutting the central support girder) are read out at one end. For the double-ended counters timing jitter is reduced by the use of mean timer modules (Cavestro et al., 1991), each handling up to eight scintillators, which also provide a logical OR of their outputs. A further OR of eight such mean-timer modules, encompassing the output of fifty double-ended and the three single-ended counters, is used to construct a 100 ns veto signal¹ for the main experimental trigger. The veto fires at a rate (790 ± 140) per 10^{13} protons on target (p.o.t.) depending

¹The width of veto signal is chosen to accommodate the timing jitter of the single-ended trigger scintillators (*q.v.*) and the difference in relative timing of trigger and veto




Figure 2.7: Front view of the NOMAD veto counters. The shaded area shows the central bank of veto counters, “ V_8 ” (see the text), and the dotted line shows the sensitive volume of the drift chambers ($q.v.$). (Taken from Altegoer et al., 1999, figure 2.)

on neutrino beam conditions, corresponding to muon rates of $(29 \pm 5) \text{ m}^{-2}$ per 10^{13} p.o.t. at the NOMAD front-face; this introduces a dead-time of approximately 4% for the main experimental trigger.

A ninth mean-timer module handles the three counters at the upper-left, and the three at the upper-right of the flux return, which lie outside the acceptance of the trigger planes ($q.v.$) and are excluded from the veto signal to reduce deadtime. The OR of the ten central counters (V_8) forms a

signals for beam-accompanying and cosmic ray particles. Signals from the banks of veto counters below the trigger planes, sensitive to backward-going cosmic rays for which the trigger scintillators fire first, are arranged to arrive 25 ns earlier than the signals from other banks.

separate veto signal for neutrino interaction triggers in the front calorimeter, as already noted.


The efficiency of the veto is monitored using flat-top muons, with the aid of the spectrometer of the CHORUS experiment (Eskut et al., 1997), and was stable within each year's run at 95.2% (1995), 96.5% (1996), and 97.8% (1997-98). Photomultiplier high voltages and currents, and discriminator thresholds are set and monitored by a slow control program written by the author; summary histograms are produced by a separate monitoring program. The veto and trigger system is discussed in further detail by Altegoer et al. (1999).

2.4.3 The front calorimeter


The front support pillar of the NOMAD comprises two of the UA1 "I" modules side-by-side, each consisting of 23 iron plates stacked longitudinally, 4.9 cm thick and separated by 1.8 cm air gaps. In 1995 the central region of this pillar was instrumented to form an iron-scintillator calorimeter—the "front calorimeter" or FCAL—as a large-mass active neutrino target for multi-muon studies and searches for neutral heavy particles.

The first 20 air gaps were instrumented with NE102A scintillation counters of dimension $175 \times 18.5 \times 0.6 \text{ cm}^3$, recycled from the CDHS experiment (Holder et al., 1978); the last three iron plates carry the main support girder for the NOMAD central detector and some related structures. Each gap contains ten scintillators stacked vertically, in two groups of five placed above and below structural rods at the vertical centre of the pillar, producing a detector of active area $175 \times 190 \text{ cm}^2$ at the centre of the neutrino beam, five nuclear interaction lengths deep with an instrumented mass of 17.7 tons (see figure 2.8). By comparison the NOMAD drift chambers have an active area of $260 \times 260 \text{ cm}^2$, so that muons from interactions in the FCAL are efficiently tracked (while muons from a significantly larger target would suffer from limited acceptance due to bending in the magnetic field); and a fiducial mass of 2.7 tons, so that the sample of interactions in the FCAL is relatively large.

The scintillators are ganged together in groups of five longitudinally using antlered light-guides, and read out at both ends by three inch photomultipliers. The FCAL is thus divided longitudinally into four segments, each consisting of a vertical stack of ten modules. Position resolution is limited by the size of these modules in the vertical and longitudinal directions ($\sigma_y = 5.8, \sigma_z = 9.7 \text{ cm}$), while the measured attenuation length of light in the scintillators allows an estimate of the x -coordinate of an interaction by comparing the response at each end ($\sigma_x = 3.8 \text{ cm}$; as discussed by Boyd, 1998, chapter 3).



(a) Top view.



(b) Front view.

Figure 2.8: The NOMAD front calorimeter (FCAL), with photomultiplier tubes and lightguides shown schematically. (a) Longitudinal groups of five scintillation counters are ganged together and viewed by a single photomultiplier; the FCAL is thus divided into four longitudinal segments. (b) The central 2 metres only of the front magnet support are instrumented. The active area of the drift chambers (*q.v.*) is also shown, dashed. (Adapted from Boyd, 1998, figure 3.4.)

Since the FCAL was built *in situ*, the energy calibration could not be performed in the usual manner with a test beam. Instead the relative calibration between modules was performed using throughgoing muons in the muon gate, while the absolute energy scale for the detector was set by a comparison of distributions of energy deposited in charged current events ($\nu_\mu + N \rightarrow \mu^- + X$) in both data and Monte Carlo under various cuts, and neutral current events ($\nu_\mu + N \rightarrow \nu_\mu + X$) in the Monte Carlo; both this procedure, and the performance of the FCAL, are discussed in detail by Boyd (1998). The energy resolution for interactions in the first three longitudinal segments or “stacks” is estimated to be $\sigma_E/E = 100\%/\sqrt{E(\text{GeV})}$, apart from the central gap where resolution is poorer.

The two data-taking triggers for interactions in the FCAL, incorporating signals formed by the FCAL response discriminated at different levels, are discussed by Altegoer et al. (1999). The first physics results based on neutrino interactions in the FCAL, a study of opposite-sign dimuon events, are currently being prepared for publication (Astier et al., 1999b).


2.4.4 The drift chambers

The principal active target for neutrino interactions in the NOMAD is provided by a set of 44 drift chambers. They were built to a novel design to effect a compromise between the requirements of


1. a small depth in radiation lengths for the detector as a whole, to minimise bremsstrahlung so that electrons can be tracked and their momentum reconstructed (and to restrict photon conversions to manageable levels);
2. a small depth in radiation lengths between successive track measurements, to minimise the effect of multiple scattering on tracks, and both secondary particle interactions and photon conversions near the event vertex; and
3. a significant target mass to maximise the total number of neutrino interactions.

Requirements (1) and (3) taken together dictate a target of material with low atomic number (Z), to minimise the radiation length for a given target mass (as radiation length scales approximately as Z^{-2}); requirement (2) suggests a low-density target with many measurement planes.

The solution chosen was to build large-area ($3 \times 3 \text{ m}^2$) self-supporting drift chambers whose walls act as the neutrino target. The walls are built from



(a) The arrangement of walls and gas gaps in a single drift chamber.



(b) Cross-section of a single drift cell.

Figure 2.9: An overview of a NOMAD drift chamber. Both sketches represent a cut by a $y - z$ plane, in the NOMAD coordinates. (Taken from Altegoer et al., 1998a, figure 7.)

a honeycomb of aramid fibre, sandwiched between skins of kevlar and epoxy resin to provide rigidity and flatness. A chamber consists of four wall panels separated by three instrumented gas gaps as shown in figure 2.9(a). Any two successive wire planes are separated by less than 1% of a radiation length ($0.01X_0$) in this arrangement. The total depth of a chamber is $0.02X_0$, for a fiducial mass of 2.7 tons for the 44 chambers (within an area $2.6 \times 2.6 \text{ m}^2$), reflecting their low average atomic mass: including internal supports, spacers, *etc.*, the chambers are 64% C, 22% O, 6% N, and 5% H by weight.

Each of the 8 mm gas gaps is filled with a 40% argon, 60% ethane mixture at atmospheric pressure, circulated in a closed circuit with a purifier to remove both oxygen and water vapour contamination. Field- and sense-wires are alternated at 3.2 cm intervals to form drift cells, as shown in figure 2.9(b). Additional field shaping is provided by aluminium strips, printed onto mylar sheets precision-glued (to better than $500 \mu\text{m}$) to the chamber walls; strip potentials are set to different values on the two sides of the gap to compensate for the effect of the magnetic field on drifting electrons. The result is a highly uniform drift field of 1 kV/cm throughout most of the cell, and an electron drift speed of about $50 \text{ mm}/\mu\text{s}$.

Within each chamber the successive wire planes are inclined at -5° , 0 and $+5^\circ$ respectively to the horizontal: in the ideal case where a track leaves a hit in each drift plane, this allows the resolution of the up-down ambiguity in each plane and the recovery of an x -coordinate estimate for the combined measurement using the stereo angle. Any residual up-down uncertainty is removed by staggering successive drift chambers by $\pm 1.6 \text{ cm}$ in y (the vertical). A Kalman filter technique is then used to fit the parameters of a helix to particle tracks, taking multiple scattering in the chamber walls into account. The momentum resolution thus achieved is

$$\frac{\sigma_p}{p} \approx \frac{0.05}{\sqrt{L}} \oplus \frac{0.008p}{\sqrt{L^5}}, \quad (2.13)$$

for p in GeV/ c and track length L in metres; the two terms added in quadrature correspond to multiple scattering and single hit resolution effects respectively. Taking the residuals of drift chamber hits with respect to these fitted tracks yields a measured resolution of $\sigma_y = 150 \mu\text{m}$ for tracks at normal incidence, with $\sigma_x \approx 1.5 \text{ mm}$.

An additional five drift chambers are interspersed with the modules of the transition radiation detector (*q.v.*), to extend the reconstruction of particle tracks into that region; this can be effectively carried out as each module of the TRD contributes only 1% of a radiation length to the detector.

The drift chambers, together with various tests of their performance, are further discussed by Altegoer et al. (1998a).

2.4.5 The trigger counters

The main event trigger for the NOMAD is provided by two planes of scintillators placed inside the central detection volume: the first at the end of the active target, the second immediately following the transition radiation detector (*q.v.*; see figure 2.6 for the relative disposition of the subdetectors). Each plane consists of 32 counters, made from BC-408 scintillator and $19.9 \times 0.5 \text{ cm}^2$ in cross-section: 28 counters of length 124 cm are arranged horizontally while 4 counters 130 cm in length are placed vertically across the lightguides of the horizontal counters, to increase the fiducial area to $280 \times 286 \text{ cm}^2$, matching the TRD. This arrangement is shown in figure 2.10.

Each scintillator is read out at one end by a proximity mesh photomultiplier (Hamamatsu R2490-05); when oriented parallel to the magnetic field,




Figure 2.10: Layout of a NOMAD trigger plane. Note the four vertical scintillation counters (shaded), providing coverage in the area occupied by the lightguides of the horizontal counters. The lightguides of the vertical counters are bent to align the corresponding photomultiplier tubes with the magnetic field axis (see the text). (Taken from Altegoer et al., 1999, figure 4.)

this type of photomultiplier suffers only a 30% loss of response in the 0.4 T field used in the NOMAD. The lightguides of the vertical counters are bent through 90° to allow the corresponding photomultipliers to be aligned with the magnetic field, as shown in the figure. Discriminated photomultiplier outputs are sent to a logic unit which forms the OR of the outputs of the 32 scintillators in each plane (T_1 and T_2 respectively) and the logical AND of the two planes ($T_1 \times T_2$). Together with the signal from the veto planes in anti-coincidence, this signal forms the main NOMAD event trigger ($\bar{V} \times T_1 \times T_2$).

The average efficiency for minimum ionising particles across all counters has been measured to be $(97.1 \pm 0.1)\%$ in the NOMAD data; for neutrino interactions with two or more tracks in the central detector, the efficiency for the coincidence $T_1 \times T_2$ is better than 99.6%. The design, operation and testing of the trigger counters are discussed in detail by Altegoer et al. (1999).

2.4.6 The transition radiation detector

Purpose and configuration:

A large transition radiation detector (TRD) is placed between the two planes of trigger counters. The purpose of the TRD is to identify electrons with high efficiency while imposing a large rejection factor on pions and other particles, suppressing the non-electronic background to the $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ search—in particular, the background due to isolated pions in $\nu_\mu + N \rightarrow \nu_\mu + X$ events (see section 2.3.3). Its intended performance is a rejection factor of 10^3 or better against pions in a broad momentum range for an electron identification efficiency of 90%: this must be combined with the rejection of pions provided by the preshower and electromagnetic calorimeter ($qq.v.$), in order to adequately suppress the background.

The distribution of subdetectors inside the NOMAD magnet volume is driven by this consideration. Placing the TRD immediately before the calorimetry makes electron tagging by the full TRD available for almost all particles whose tracks are within the calorimeter acceptance; by the same token tracks outside this acceptance, for which adequate pion rejection cannot be available, usually lack TRD information as well. The TRD occupies a smaller volume in this configuration than in one where it is distributed throughout the detector, for equivalent electron rejection: this allows an increase in the target mass of the experiment, granted that the drift chamber walls themselves serve as the neutrino interaction target, compacted into the first four metres of the magnet volume (as discussed in section 2.4.4 above; see also figure 2.6). As previously noted, an additional five drift chambers are included in the TRD volume, to extend the reconstruction of particle tracks

through this volume and up to the second trigger plane and preshower (Astier et al., 1991b).

Design and calibration:




Figure 2.11: Top view of a pair of TRD modules, and their associated drift chamber (left); structure of a TRD module (right). (From Bassompierre et al., 1998a, figure 3)

The TRD is built from nine modules, each consisting of a radiator of 315 polypropylene foils ($15\ \mu\text{m}$ thick, spaced by $250\ \mu\text{m}$) followed by a detection plane of 176 straw tubes of aluminised terphane ($16\ \text{mm}$ diameter, $3\ \text{m}$ long; tube wall thickness $28\ \mu\text{m}$) each strung with a $50\ \mu\text{m}$ gold-plated tungsten wire, and fed by an 80% xenon, 20% methane gas mixture. Its active area ($2.85 \times 2.85\ \text{m}^2$) is matched to that of the drift chambers and the electromagnetic calorimeter (*qq.v*). The modules are deployed as four doublets², each followed by a single drift chamber, with a single module followed by a single drift chamber at the end of the TRD volume. This arrangement and the main features of the TRD modules are shown schematically in figure 2.11. A radiator depth of 1% of a radiation length ($0.01 X_0$) was chosen so that only $0.02 X_0$ of material separates successive drift chambers in the TRD region, limiting bremsstrahlung by electrons and the effect of multiple scattering on all tracks (*cf.* the $0.02 X_0$ depth of a single drift chamber). Other features

²Four of the nine modules, found experimentally to have a lower photon yield than the remaining chambers, were placed in the downstream position within the module doublets: some fraction of photons not absorbed by the upstream module are then detected by the downstream module, compensating in part for the lower yield.

(foil thickness and spacing, the choice of foils over foam as radiator material, choice and parameters of the straw tube planes *etc.*) were chosen on the basis of simulation and test beam studies.

Electron/pion discrimination in the TRD is based on the difference in the straw tube signal due to pions, depositing energy only by ionisation of the gas; and electrons, for which soft X-ray transition radiation, emitted as the electron crosses foil/gas interfaces in the radiator, is absorbed by the xenon giving an additional signal. The response of a TRD module to a test-beam of pions and electrons, both at 10 GeV/ c , is shown in figure 2.12: the difference in both mean response and distribution shape for the two species is clear. Also shown is the response to 5.89 keV X-rays emitted by an ^{55}Fe source, in the form of a tape impregnated with ^{55}Fe -sulphate solution, stretched across the middle of each straw tube plane. These sources yield about 20 counts/s per straw, which are histogrammed during the off-spill period of each data taking run, providing run-by-run calibration of the gain for each straw.

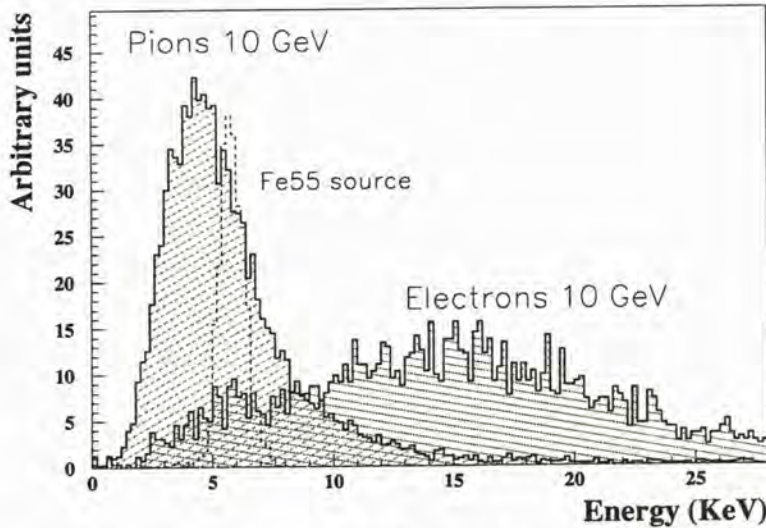


Figure 2.12: Response of a TRD module to 10 GeV/ c pions and electrons in a test beam, and to the ^{55}Fe source. (From Bassompierre et al., 1998a, figure 18)

Straw tube gain is a function of voltage, the temperature and pressure of the gas, percentage of CH_4 in the gas mix, and N_2 , O_2 and water-vapour contamination. These parameters, together with the gas-flow rate, are constantly monitored by a slow control system; gas losses are compensated automatically by feedback-controlled Xe and CH_4 input. O_2 and H_2O contaminations are limited by surrounding each module with an aluminised mylar skin, and maintaining a flow of N_2 between the skin and the straw tubes;

residual contamination is removed by including a palladium catalysor (for O_2) and a molecular sieve (for H_2O) in the Xe/CH_4 gas circulation line. The uniformity of response along the length of each straw tube, a function of gas-flow rate and temperature profile, is monitored by histogramming the response to throughgoing muons in the “flat top” between the two neutrino pulses, about 50μ per accelerator cycle (see section 2.2). The gain of the amplification chain for each straw-tube wire is also independently monitored.

Both the gain and uniformity of response of the detector have proven to be remarkably stable throughout the life of the experiment.

The design, construction, testing and calibration of the TRD are more fully described by Bassompierre et al. (1998a); the e/π discrimination algorithm and various tests of TRD performance are described in detail by Bassompierre et al. (1998b). The discrimination algorithm and its performance, together with various adaptations and extensions carried out by the author, are also discussed in chapter 5 of this work: the e/π discrimination algorithm and its use in the task of electron rejection in the $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n \geq 0)$ analysis are treated in section 5.2; the general case of discrimination between pairs of particles using information from the TRD is set out in section 5.3.

2.4.7 The preshower detector

The second trigger plane is immediately followed by two lead-antimony sheets, 4% antimony by weight and 9 mm in total thickness, on either side of a 2 mm aluminium plate. This arrangement has a total depth of $1.6X_0$ and acts as a converter, initiating electromagnetic showers for gamma rays originating in the active target and TRD region. The converter is followed by two planes of proportional tubes, of active cross-section $9 \times 9 \text{ mm}^2$ with 1 mm thick walls, oriented horizontally (286 tubes) and vertically (288 tubes). Each tube is strung with a gold-plated tungsten anode at 1500 V and fed with a continuously circulating 80% argon, 20% CO_2 mixture. The resulting “preshower” detector, shown in figure 2.13, immediately precedes the electromagnetic calorimeter (ECAL) and allows converting photons to be localised to greater precision than would be possible with the ECAL alone. Since electromagnetic showers are also initiated by electrons, while pions deposit energy principally by ionisation of the proportional tube gas, the preshower also contributes to electron/pion discrimination.

The gains of the individual preshower tubes are monitored run-by-run using throughgoing muons from the SPS flat-top. The response of the combined preshower-ECAL system to electrons and pions has been measured in a test-beam, as discussed in the following section. The resolution of the impact point of photons converting in the lead has been estimated to be $\sigma = 1 \text{ cm}$,




Figure 2.13: An exploded view of the preshower detector. (Taken from Altegoer et al., 1998a, figure 14.)

based on measurement of charge exchange interactions of π^- in the converter in test beam runs. Taken together with the ECAL energy resolution, this implies a π^0 mass resolution of $11 \text{ MeV}/c^2$, in agreement with that found in the NOMAD data. This measurement, and the design and operation of the preshower, are discussed in somewhat more detail by Altegoer et al. (1998a).

2.4.8 The electromagnetic calorimeter

Accurate electromagnetic calorimetry is required in the NOMAD in order to

1. recover the initial momentum of electrons from the energy deposited by the electron track, and the energy of bremsstrahlung photons emitted along the track (the tracking detector has a total depth of $1.1X_0$);
2. discriminate between electrons, depositing their entire energy through an electromagnetic shower, and hadrons, depositing a fraction of their

energy chiefly through ionisation, including the early stages of an hadronic shower; and

3. reconstruct the π^0 component of the hadronic system.

Requirement (1) reflects the need to reconstruct the tau decay candidate for the $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ analysis; imposing (2) allows, together with the e/π separation already provided by the TRD, the suppression of the isolated π^\pm background to the $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ analysis; (3) reflects the need for accurate reconstruction of the total transverse momentum in all of the tau-decay and many other analyses. The dynamic range required is large (from 100 MeV and less to 100 GeV) and a good energy resolution and uniformity are desirable. An electromagnetic calorimeter (ECAL) built from lead-glass Čerenkov counters was chosen to meet these requirements.

The ECAL consists of a 25×35 grid of counters, each of cross-section $112 \times 79 \text{ mm}^2$ and 19 radiation lengths deep. The asymmetric shape is a function of a 45° cut at the back face, to which a 77 mm proximity mesh phototetrode (Hamamatsu R2186-01) is coupled. This arrangement, shown in figure 2.14, allows the efficient operation of the tetrode in the NOMAD magnetic field (there is a signal reduction of less than 20% with the tetrode




Figure 2.14: The geometry of an ECAL lead-glass counter, with the neutrino beam and magnetic field directions, and phototetrode shown. Taken respectively from Schmidt (1997), figure 2.16 and Boyd (1998), figure 2.11

axis at 45° to the field) together with good collection of Čerenkov light from the counter. The counter depth is a compromise between efficiency of light collection, and longitudinal containment of electromagnetic showers. Each tetrode is read out by a low-noise electronic chain providing a dynamic range better than 4×10^3 ; in addition to the energy measurement, a fast signal is also provided to allow off-line rejection of energy depositions out of time with the event trigger, and (if required) to allow triggering of NOMAD data-taking based on energy deposited in the ECAL (Altegoer et al., 1999).

The 875 counters were individually calibrated prior to assembly with a 10 GeV/ c electron test-beam. Some non-linearity of response to electromagnetic showers is expected, since the lead-glass is not perfectly transparent, and the shower maximum moves deeper into the counter (and hence closer to the tetrode) as energy increases. To test this effect a matrix of 10×5 counters was exposed to electron test beams with energies from 1.5 to 80 GeV. The deviation from linearity is well-described by a single logarithmic correction

$$E' = \frac{E}{1 + \epsilon \log E/E_0} \quad (2.14)$$

where $\epsilon = 0.038$, using the test-beam calibrated energy $E_0 = 10$ GeV/ c as a reference; this corresponds to a correction $(E' - E)$ of less than 10% for energies from 0.5 to 100 GeV. A separate correction

$$E'' = \frac{E}{\sqrt{\cos \theta}} \quad (2.15)$$

was found to describe the dependence of observed energy (E) on incidence angle θ , based on both test-beam and Monte Carlo studies. The resulting energy resolution is

$$\frac{\sigma_E}{E} = (1.04 \pm 0.01)\% + \frac{(3.22 \pm 0.07)\%}{\sqrt{E(\text{GeV})}} \quad (2.16)$$

A further correction, based on the energy deposited in the preshower, must be applied for both electrons and photons whose electromagnetic showers begin in the lead converter. Since there is no such effect for pions, a combined cut on the energies deposited in the preshower and ECAL provides powerful discrimination between pions and electrons, where their momentum is known. For particles of energy 2 to 10 GeV, a rejection factor of 10^3 against pions may be achieved together with 90% efficiency for electrons (see figure 2.15). Electron/pion discrimination is also possible based on the fit of the energy deposition pattern in an ECAL cluster to the well-parametrized electromagnetic shower shape (Autiero et al., 1999).




Figure 2.15: Scatter plots of the preshower vs. ECAL response to a 5 GeV/c test-beam of (a) electrons and (b) pions. A cut along the contour shown has a 90% efficiency for electrons. (Taken from Altegoer et al., 1998a, figure 32.)

Response is uniform across the face of the ECAL within $\pm 0.5\%$ and the ECAL is hermetic to better than 99%. The average position resolution for isolated electromagnetic showers is 4 mm (Autiero et al., 1999): the more typical case, where electromagnetic showers overlap the broader energy deposition by charged hadrons, may be resolved with the aid of the 1 cm resolution of the preshower for showers beginning in the lead converter. The ECAL is approximately 1.6 interaction lengths ($1.6\lambda_I$) deep, so the majority of hadrons begin the hadronic showering process within the lead-glass, depositing a variable amount of energy over an irregular region which (albeit with low probability) may extend many counters away from the hadron impact point (see for example Autiero et al., 1998, figure 2 and discussion). These overlapping energy deposits must be resolved using a clusterization algorithm, with the aid of the drift chamber tracking and the preshower.

The response of each of the lead-glass counters is continuously monitored off-spill using two blue light emitting diodes mounted on the same face as the tetrode; each electronic chain is also monitored directly. An independent check is provided by the response of counters to throughgoing muons from the SPS flat-top: this has been found to be stable within $\pm 1\%$. The design of the ECAL and its monitoring system, and the performance of the ECAL prototype in a test-beam, are discussed in much further detail by Autiero et al. (1996a,b, 1997).

2.4.9 The hadronic calorimeter

The back support pillar of the NOMAD comprises six of the UA1 I-modules side-by-side, each consisting of 23 iron plates 4.9 cm thick and separated by 1.8 cm air gaps, as for the modules of the front support: the total depth of the pillar is thus 6.7 interaction lengths ($6.7\lambda_I$). It roughly matches the width and height of the flux return (see figure 2.6) and greatly exceeds the cross-section of the central detector, and thus functions as an effective muon filter for the products of neutrino interactions in the central detector. The first eleven of the twenty-two air-gaps have been instrumented with scintillator, so that the support also acts as a crude hadronic calorimeter (HCAL), $3.1\lambda_I$ deep.

Scintillation counters of cross-section $18.3 \times 1 \text{ cm}^2$, 3.6 m long, have been fitted with long, tapered light pipes at each end and threaded into the gaps between the iron plates. A vertical arrangement of four groups of counters, with modification to the height and shape of counters in the central region, is required to accommodate the structural bolts holding together the back pillar, as shown in figure 2.16. Comparison with the same view of the FCAL (figure 2.8(b)) shows that a much larger region has been instrumented, since the HCAL acts as a calorimeter for particles produced in the drift chamber target (rather than a target for neutrino interactions, whose products are then tracked by the drift chambers). The instrumented region is $3.6 \times 3.5 \text{ m}^2$, to be compared with the $2.6 \times 2.6 \text{ m}^2$ fiducial area of the drift chambers.

One module of the HCAL consists of all eleven counters at a given height: at each end, the light pipes of the counters are ganged together and read out by a single five-inch photomultiplier. The HCAL thus has no longitudinal segmentation. The energy deposited in a module is recovered from the geometric mean of the phototube signals, while the horizontal position is determined from the ratio of signals and the measured attenuation length of the counters. Where more than one hadronic shower deposits energy in a counter, individual energies and positions are determined using the extrapolated positions and momenta of the particles (and the deposition in the




Figure 2.16: A front view of the hadronic calorimeter (HCAL). Note the three vertical gaps between groups of counters, and the notches in the eighth and eleventh counters, due to the bolts holding together the iron plates. (Taken from Altegoer et al., 1998a, figure 18.)

ECAL, if any) as a guide.

A signal based on energy deposition in the HCAL is also provided for use as an input to the NOMAD data-taking trigger, typically to exclude events with a penetrating component when filtering for muon-decay events (Altegoer et al., 1999).

As for the FCAL, the HCAL was built *in situ* and so could not be calibrated with a test beam. Its response to deposited energy has been estimated using a large sample of Monte Carlo events, with the relative scale between data and Monte Carlo being set by comparing the response to identified, isolated muons; the calibration is thus able to track any changes in response over time. An estimate of the *energy* of a hadron striking the HCAL (as opposed to the energy deposited) is relatively difficult to obtain: the HCAL is shallow ($3.1\lambda_I$) and is preceded by a relatively large depth of material in the electromagnetic calorimeter ($1.6\lambda_I$). Since a significant depth of uninstrumented magnet coil ($0.7\lambda_I$) lies between the ECAL and HCAL, and the HCAL is not longitudinally segmented, hadrons beginning to shower in the HCAL cannot be separated from hadrons beginning to shower in the material upstream. The energy deposited by hadrons must therefore be estimated in

all cases using a combination of the ECAL response (where present) and the response of the HCAL.

The procedure used begins with a linear combination of the ECAL and HCAL responses, optimised for Monte Carlo hadrons. Based on a comparison with particle momentum for well-isolated hadrons in the NOMAD data, a (nonlinear) correction is then applied to centre the energy estimate on the actual particle energy. The resulting energy resolution in the data is estimated to be

$$\frac{\sigma_E}{E} = \frac{0.965}{\sqrt{p(\text{GeV}/c)}} + 0.140; \quad (2.17)$$

the derivation of this value, and the procedure of HCAL calibration, are described in detail in an (unpublished) note of Hurst (1997b). The energy estimate thus obtained is an adjunct to the well-measured particle momentum for charged hadrons; for neutral hadrons, it is the only available energy measure. The use of such measurements in analysis is discussed in chapter 6.

2.4.10 The muon detection system

Identification of muons is performed using a system of 10 large-area drift chambers recycled from the UA1 experiment (Eggert et al., 1980). The chambers, each with an active area of $3.75 \times 5.55 \text{ m}^2$, are arranged as shown in figure 2.17 to cover the region behind the back support pillar (see also figure 2.6 for the relative arrangement of subdetectors). Each chamber consists of two planes of horizontal drift tubes (staggered by half a tube-width to resolve up-down ambiguity), followed by two similarly-staggered planes of vertical drift tubes. The drift tubes are $14.9 \times 4.4 \text{ cm}^2$ in cross-section, with a maximum drift distance of 7 cm, and are fed with the same 40% argon, 60% ethane gas mixture used for the central detector drift chambers. The chambers are arranged in pairs (or *modules*) as shown, with the two chambers in each module separated by 75.5 cm for track segment reconstruction.

The muon detection system is divided into two *stations*: the first, consisting of three modules³ immediately behind the back support, is reached by muons with momentum $p \gtrsim 2.3 \text{ GeV}/c$; the second, consisting of two modules, is separated from the first station by an additional iron wall 80 cm deep and is reached by muons with momentum $p \gtrsim 3.5 \text{ GeV}/c$, the threshold increasing in both cases for muons at non-normal incidence. Between the 1995

³As shown in the figure, the third module of the first station (“module 5”) is oriented at 90° to the others (*i.e.* with its long axis horizontal rather than vertical), and placed above them in order to provide sensitivity to muons emerging from the central detector at large angles; its constituent chambers separated by only 58.9 cm.




Figure 2.17: The NOMAD muon detection system. In (b) the position of module 5 is shown dotted.

and 1996 data-taking runs, the first station was augmented by four large scintillators ($80 \times 0.5 \text{ cm}^2$ in cross-section, 200 cm long) recycled from the CDHS experiment (Holder et al., 1978), placed as shown in figure 2.17 and read out at one end by five-inch photomultipliers. This addition prevented muons passing through the small gap between the modules of the first station from escaping detection.

Track segments are reconstructed separately in each of the two stations, from 3 or 4 hits per projection in the typical case, with an efficiency of 97%; the reconstruction procedure is discussed in greater detail in section 6.6.1. Residuals with respect to these track segments imply a single hit resolution from as low as 350 to 600 μm , depending on the gas quality. The chamber performance is monitored continuously using high energy muons both in the neutrino spills and the SPS flat-top.

2.4.11 The silicon target

For completeness we should note that for the 1997 and 1998 data-taking runs the first three drift chambers were removed and replaced with a prototype silicon strip detector. Discussions of the design for a “next generation” experiment to search for $\nu_\mu(\nu_e) \rightarrow \nu_\tau$ neutrino oscillations—one improving on the sensitivity of NOMAD and CHORUS by an order of magnitude—have centred on a sandwich of target plates with high resolution tracking detectors, *viz.* silicon microstrip detectors, interleaved with conventional tracking detectors inside a dipole electromagnet (Gomez-Cadenas et al., 1996; Gomez-Cadenas and Hernando, 1996; Ereditato et al., 1996; Ayan, 1997). Such a design would combine the reconstruction of event kinematics of the NOMAD experiment with the selection of tau decay candidates by event topology that is the basis of the sensitivity of CHORUS. Its feasibility turns on the performance of silicon detector surfaces comparable in size to the core of a neutrino beam, say $1.5 \times 1.5 \text{ m}^2$ for the WANF beam, yet constructed at modest cost. A small module of such a detector was therefore incorporated into the NOMAD as a test of the technique.

The design of the NOMAD silicon target (STAR) is shown in figure 2.18. The neutrino interaction target is provided by four plates of boron carbide (B_4C), each $72 \times 31.5 \times 2 \text{ cm}^3$, with a total mass of 45 kg; each plate is followed




Figure 2.18: Schematic side view of the silicon target. (Taken from Barichello et al., 1998a, figure 2.)

by a plane of silicon detectors, with a fifth detector plane added to ensure that interactions in the last plate are followed by two measurements. B_4C was chosen for its high density, a low- Z material being required to minimise multiple scattering.

Each silicon detector plane is made up of 10 overlapping modules, shown end-on in the figure. Each module is a “ladder” of twelve silicon microstrip detectors, each with a strip and readout pitch of 25 and 50 μm respectively, and 640 strips read out (figure 2.19). The result is a detector with an active surface of 1.14 m^2 and a resolution of $\approx 10 \mu\text{m}$ in the vertical (y) coordinate, based on charge-sharing between the strips (smaller ladders of nine detectors were found by Barichello et al. (1998b) to have a resolution of 6 μm). The low event rate and detector occupancy intrinsic to a neutrino detector make possible the use of long shaping-time, low-noise readout chips for the detectors, which in turn makes practical the long detector ladders (the longest implemented at the time of writing).

The silicon target is installed at the beginning of the inner volume of the NOMAD magnet, at its vertical and horizontal centre, coinciding with the core of the neutrino beam. Small scintillator planes immediately precede and follow the target to provide veto and trigger signals respectively for readout of the STAR. The readout electronics, performance, trigger conditions and data acquisition for STAR have been described elsewhere, together with a much more detailed account of the detector design (Barichello et al., 1998a,b).




Figure 2.19: Layout of an individual silicon layer. (Taken from Barichello et al., 1998a, figure 3.)

Results from the 1997 and 1998 data-taking runs will subsequently be published by the same authors, but will not be incorporated into this present work. The removal of the first three drift chambers in 1997 and 1998 to accommodate the STAR of course required a change to the fiducial volume for neutrino interactions in the NOMAD proper: this is discussed in chapter 6.

2.5 Data recording and handling

2.5.1 Triggering

The trigger logic for the NOMAD, and various bookkeeping tasks, are implemented using a purpose-built VME-based module, MOTRINO (Schmidt, 1997). Two “master” modules, handling triggers in the neutrino gates and the muon gate respectively,

- generate timing signals synchronised with the SPS cycle;
- form up to six triggers from eight subdetector signals using computer-programmable logic, introducing artificial deadtime for trigger-rate moderation where necessary;
- calculate trigger livetimes from coincidences of the trigger inputs and the detector busy-signal; and
- count rates for the various trigger inputs.

Five slave modules, one for each of the VME modules controlling subdetector readout (see the next section), transmit the global triggers, manage trigger information generated by the subdetectors, and generate local triggers outside the physics gates for subdetector calibration purposes.

The main NOMAD event trigger in the neutrino gates is formed from the coincidence of the OR signals from the two trigger planes, with the veto counter signal in anticoincidence ($\bar{V} \times T_1 \times T_2$, as noted in sections 2.4.2 and 2.4.5 above). The data used for the analyses presented in this thesis were taken with this trigger condition. Its rate is $\approx 5.0/10^{13}$ protons on target (p.o.t.): about 10% of these triggers are candidate neutrino interactions in the drift chamber target, the remainder being due to cosmic rays outside the veto acceptance (at a rate equivalent to $\approx 1/10^{13}$ p.o.t., depending on the beam conditions), throughgoing muons for which the veto plane fails to fire ($\approx 1.5/10^{13}$ p.o.t.), and neutrino interactions in the magnet coil and flux return ($\approx 2/10^{13}$ p.o.t.). Events of these latter types are rejected offline using various quality cuts, principally on the formation of a primary event vertex,

consistent with a neutrino interaction, within the fiducial volume of the drift chambers (see chapter 6).

The effective dead-time for this trigger is 15% on average, with fluctuations of RMS size 3% due to variations in the intensity of the proton beam on the T9 target, the width of the neutrino spills *etc.*. The dead-time, which is constantly monitored, is estimated using a signal from instruments in the WANF beamline which reflects the structure of the neutrino spill, so that both high- and relatively low-neutrino-flux periods are taken into account; it is dominated by the readout of the subdetectors to the front-end electronics.

Throughgoing muons from the SPS flat-top, which are extensively used for monitoring and calibration of the subdetectors (*qq.v.*), are selected in the muon gate by the veto and trigger planes in coincidence ($V \times T_1 \times T_2$), moderated by an artificial deadtime to a maximum of (typically) 50 triggers per gate; and a second, unmoderated trigger formed by the central veto bank (V_8 , see section 2.4.2) in coincidence with the trigger planes and a low-threshold trigger signal from the front calorimeter ($V_8 \times T_1 \times T_2 \times FCAL'$). The second of these triggers is used to better populate the central region of the detector—the majority of muons from the SPS flat-top originate from test-beams and other experiments upstream of the NOMAD and cross the detector at an angle.

A range of other triggers are implemented in both the neutrino and muon gates to study events in the FCAL (Boyd, 1998), neutrino interactions in the ECAL (Weiße, 1997), rare-process searches using the calorimeters (Steele, 1996; Altegoer et al., 1998b), events in the silicon target and simultaneous interactions in the NOMAD and CHORUS detectors. Dedicated triggers for measurement of subdetector efficiency, subdetector occupancy, and the selection of muon-decay and delta ray events are also implemented. All of these triggers, the subdetector trigger inputs and the MOTRINO module are discussed in detail by Altegoer et al. (1999).

2.5.2 Data acquisition

Hardware and data acquisition software

The NOMAD data acquisition electronics and network configuration are shown schematically in figure 2.20. Data from the various subdetectors are processed by three types of FASTBUS modules: CERN-designed 12-bit charge integrating ADCs (“cADC” in the figure) and 12-bit peak sensing ADCs (“pADC”); and commercial 16-bit TDCs (LeCroy 1876 Model 100). Up to

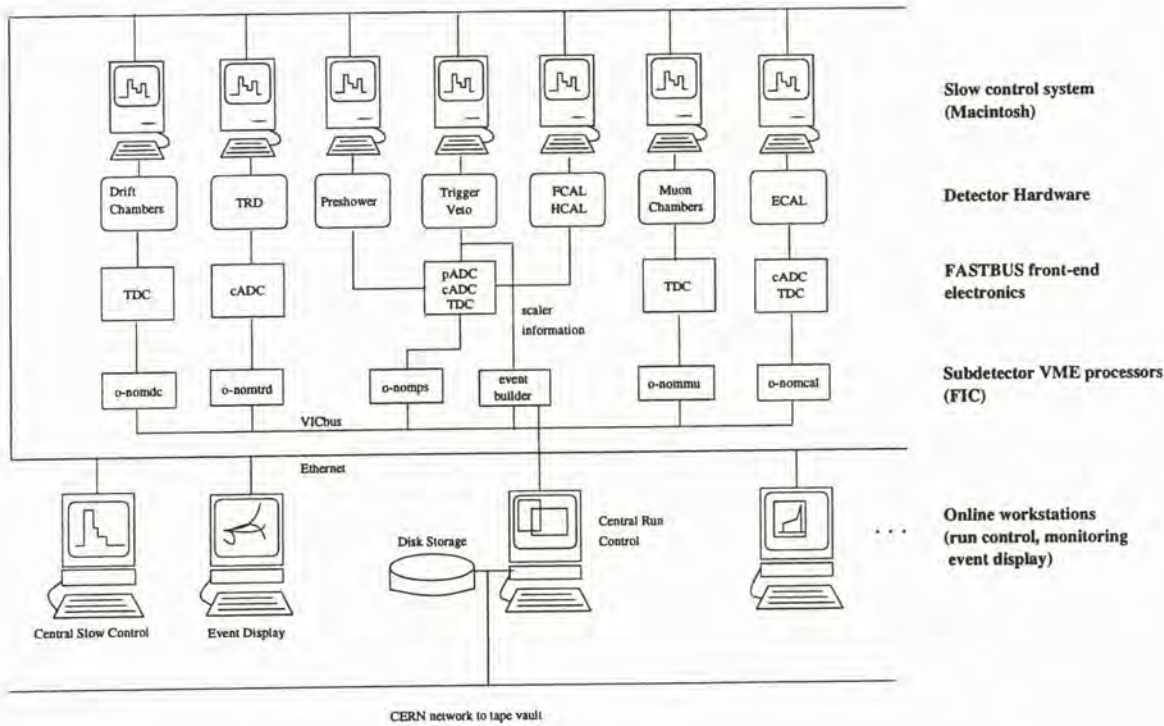


Figure 2.20: The NOMAD data acquisition system. (Taken from Boyd, 1998, figure 2.14)

11648 channels⁴ are read out per event. The digitized information from all events occurring within a given gate are buffered internally on these front-end modules.

The next level of processing is carried out by six VME-based boards (FIC 8234): one “master” board, then one each for the drift chamber, TRD, ECAL and muon chamber electronics, with a sixth board handling the preshower and all the scintillation counter information. Each FIC runs a data-acquisition process, the so-called *stage*, a data-driven program developed using the CERN-designed CASCADE package (Perrin et al., 1993). At the end of each gate the data from the front end modules are read into local buffers by their respective stages, assembled into sub-events, then passed to the stage running on the master FIC, the “event builder”, through a VME interconnect (VIC) bus. Beam, scaler, monitoring and slow control data are also passed by their respective systems to the event builder. The event

⁴Excluding the silicon target, which contributes an additional 32000 channels: data acquisition for the STAR is discussed by Barichello et al. (1998a).

builder stage assembles the various data into consistent events, which are then written by a data-recording process to one of two nine-gigabyte disks for temporary storage. A final process carries out a twice-daily transfer of data from disk to the CERN tape vault for permanent storage.

Monitoring and control of the VME processors, the run control program, and the central slow control program are executed by SUN workstations, connected to the other electronics by ethernet. Individual subdetector slow control programs are executed by dedicated Macintosh computers communicating directly with the subdetector hardware.

Subdetector monitoring and slow control

Monitoring programs running on the SUN workstations connect to the stages by ethernet, forming histograms of subdetector response in the neutrino and muon gates, and also the response to calibration “events” off-spill (*e.g.* to the LED light pulses and to directly injected charge for each of the counters of the ECAL, *q.v.*). Trigger, scaler and beam data are also accessed by these programs, allowing both histogramming of subdetector response by trigger type *etc.* and monitoring of the beam and elements of the trigger system. Histograms are variously recorded on a database as a first stage of subdetector calibration (such as the response of the TRD tubes to the ^{55}Fe source), printed out at eight-hourly intervals as a reference, and continuously updated to allow shift crews to monitor detector performance. Throughgoing muons from the SPS flat-top are processed by a special implementation of the event reconstruction code to allow calibration and monitoring of individual tubes, counters *etc.* to deposited energy, with reasonable statistics.

The Macintosh computers connected to the subdetector electronics, and various dedicated transducers, run “slow control” programs written in the LabVIEW graphical language to monitor high and low voltages, currents, gas flows, temperatures *etc.*. Subdetector voltage and other settings are typically controlled by the same programs. Summary data, and alarms in the case of readings outside the set tolerances, are regularly sent to a central slow control program running on a SUN workstation. This program alerts the shift crew in case of problems, and passes periodic data summaries, and all alarm conditions, to the event builder stage for inclusion in the permanent data storage.

Run control

High level control of the data acquisition system is carried out by a finite state machine, written in C++ and interacting with the shift crew via an X-

windows interface, so that the state of every stage and monitoring program is visible “at a glance”. Definite rules are set down for transitions from one state to another and the system operates semi-automatically, ending one data-taking run at a specified point and beginning a new run immediately all bookkeeping tasks have been performed; only occasional intervention by the shift crew is required. Transitions between runs, and down-time due to data acquisition and monitoring software problems cost less than 3% of the potential data-taking time.

2.5.3 The event reconstruction code in outline

Reconstruction of events in the NOMAD proceeds in two broad steps or “phases”. In phase one processing the raw data for each of the subdetectors—tetrode output in the ECAL, timed “hits” in the drift chambers and muon chambers, and so on—are converted to measurements in physical units using calibration information stored in a database. (In many cases such calibration is able to track small changes in subdetector response on a run-by-run basis.) Measurements by a given subdetector are also assembled into higher-level reconstruction “objects” during phase one: a Kalman filter process fits tracks to the hits in the drift chambers, and fits these tracks to vertices; simpler pattern-finding code fits straight-line tracks in the $x - z$ view to hits in the TRD straw tube planes; track segments in $x - z$ and $y - z$ projection are formed from hits in the muon chambers, and then into tracks in space where possible; a preliminary resolution of preshower and ECAL energy deposits into clusters is carried out, *etc.*.

During phase two of the reconstruction, objects from different subdetectors are “matched” to form a picture of the event in the detector as a whole. In the simplest case, track segments from the muon chambers are matched to extrapolations of drift chamber tracks, identifying these tracks as muons, when they are consistent in both position and direction. Re-fits of the subdetector data are triggered where appropriate: for example, track segments from the TRD are refined as the phase-one segments are matched to the longer three-dimensional fitted tracks reconstructed from the drift chamber data. In other cases, complex subdetector information is only resolved at phase two, such as for the ECAL, where overlapping clusters are disentangled using preshower information to localise converting photons and showering electrons, and extrapolated drift chamber tracks to identify the impact point of charged particles. The vertical “strip” of ECAL deposition due to bremsstrahlung by electron tracks is also reconstructed, where possible, at this stage.

The majority of phase two code is executed inside a loop, which is iterated

until a consistent picture of the event emerges. This procedure was chosen because the mode of reconstruction in one detector may depend on information from another—as for pion- or electron-track hypotheses for the Kalman filter, or ECAL clusterization—and preliminary identifications may change. In most cases sufficient information is stored to vary the default particle identification choices at the analysis level, for example the use of TRD information to separate pions and electrons (see chapter 5, especially section 5.2.4). Information from the HCAL is especially coarse (see section 2.4.9) and in most cases disentanglement of overlapping showers is strongly dependent on both track extrapolation and identification, and so HCAL reconstruction is implemented as a special case outside the loop, when other subdetector information has been resolved.

Physics analysis is typically carried out on a reduced data format, the so-called DST (“Data Summary Tape” format), reproducing the principal “physics” information from the phase two reconstruction, such as particle momenta, energies and positions, without some of the lower-level data. The $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n \geq 0)$ analysis presented in this thesis was carried out in this way. The full phase two output remains available, however, and was used for parts of the particle identification studies reported in chapter 5, chiefly when the DST format had not yet evolved to include the necessary information.

Chapter 3

Variables to distinguish τ decay from neutral current events

3.1 Introduction

A search for three-prong tau decays

$$\tau^- \longrightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n \geq 0)$$

must in the first place distinguish such events from those caused by the neutral current process $\nu_\mu + N \rightarrow \nu_\mu + X$. In both cases there is in general no muon or electron at the primary interaction vertex to “tag” the event, and the visible event momentum has a component perpendicular to the neutrino beam, due to the “missing” momentum of the final-state neutrino. We therefore seek a set of quantities sensitive to the differences which *do* exist between these two classes of event. A 3π candidate produced in the decay of a tau lepton, itself a product of a deep inelastic scattering, should in general have rather different properties to three pions chosen “at random” in a neutral current event: our task is to quantify those differences.

The choice of variables to distinguish three-prong decays from the neutral current background is discussed in this chapter and the chapter following. General matters concerning the construction of analysis variables from event measurements, and their classification, are treated in section 3.2; an historical discussion of some variables used in three-prong analysis follows in section 3.3. It turns out that the choice of variables to describe the overall properties of an event is a surprisingly difficult problem, and we devote the remainder of this chapter (section 3.4) to that task. The many variables sensitive to the internal structure of the $\tau \rightarrow \nu 3\pi$ decay are then discussed in chapter 4. In particular the selection of three pions from an event to form

the tau decay candidate, a non-trivial task, will be discussed in sections 4.2 and 4.6. We will take this selection for granted in the present chapter in order to focus on the overall properties of three-prong decay and neutral current events.

3.2 The construction and classification of analysis variables

At the simplest level a reconstructed event in the NOMAD can be considered as a set of tracks and neutral objects, with the three-momentum and charge listed for each. Once particle identification has been carried out, hypothetical particle types, and hence masses and energies, can also be assigned to each object. Moving to a further level of sophistication, the reconstruction code returns an error matrix along with the momentum for each track, as well as the beginning- and end-points of the track, number of “hits” in the drift chambers and so on—the count has already reached sixteen quantities per track—while neutral objects in the calorimeter can be assessed for their degree of overlap with other neutrals, or with clusters due to charged tracks. This additional information, in various combinations, can be used to carry out “quality cuts” on the data, as discussed in chapter 6, or to form uncertainties on higher-order measured quantities such as invariant masses.

3.2.1 The kinematic variables

Obviously the three-prong decay analysis must be carried out using a greatly reduced set of quantities: even the list of three-momenta for the tracks and neutrals contains much “irrelevant” detail, and does not present the most important information in an accessible form. At the simplest level we can treat an event as consisting of two three-vectors only: the momentum of the candidate 3π system, and the momentum of the remaining tracks and neutrals, which we denote “the hadronic system”. (Only in $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ events where we have correctly chosen the pions forming the 3π system, will this actually correspond to the hadronic system properly-so-called; in neutral current events all of the tracks and neutrals belong to the true hadronic system. In all subsequent discussions, except where otherwise stated, we will take the choice of the pions belonging to the “ 3π system” for granted, and speak of the remainder as “the hadronic system”.) This approximation is shown in figure 3.1.

For convenience we also define the “missing transverse momentum” as the additional momentum required to achieve momentum balance in the

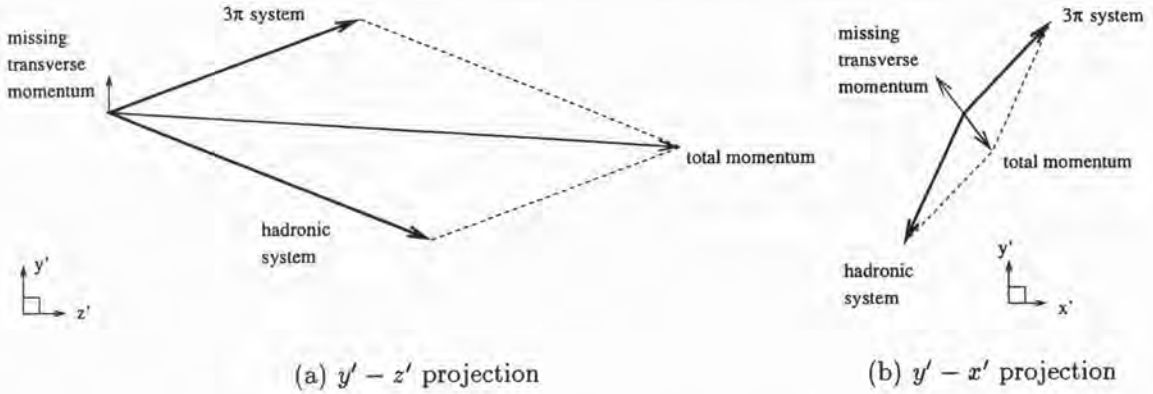


Figure 3.1: A three-prong candidate event, where the “internal structure” of both the 3π and the hadronic system has been ignored.

transverse plane:

$$\begin{aligned}\vec{p}_T^{miss} &\stackrel{\text{def}}{=} -\vec{p}_T^{\text{tot}} \\ &= -(\vec{p}_T^{3\pi} + \vec{p}_T^{\text{had}}),\end{aligned}\tag{3.1}$$

where the subscript “T” denotes quantities, especially 2-vectors, in the plane whose normal is aligned with the incoming neutrino beam. The vector \vec{p}_T^{miss} , which by construction is a function of the two three-vectors $\vec{p}^{3\pi}$ and \vec{p}^{had} which form the event in the simple approximation being treated here, will be used repeatedly in all following discussions.

At this level there are only five degrees of freedom in the event: six numbers are required to describe the pair of three-vectors, but we assume that the azimuth of $\vec{p}^{\text{tot}} \stackrel{\text{def}}{=} (\vec{p}^{3\pi} + \vec{p}^{\text{had}})$ about the neutrino beam is irrelevant for physics, whereupon we are free to make a rotation reducing the six numbers to five. Any quantity which is a function only of these degrees of freedom, we will call a *kinematic variable*. These variables, summarising the overall properties of the event, will be treated in sections 3.3 and 3.4 below.

3.2.2 Definition of coordinate axes

We have already referred to a coordinate system for variables, which we will now pause to define. The y -axis is given by the vertical, in the usual way. The z -axis is given by the longitudinal axis of the NOMAD (figure 2.6, with z increasing from left to right), with the x -axis then being determined by the usual right-handed convention (*i.e.* “into the page” on figure 2.6). Where

absolute quantities are necessary the centre of the front face of the magnet coil (see section 2.4.1) will be taken as the origin of coordinates.

This is the obvious physical coordinate system for the detector, but the neutrino beam suggests a more useful set of coordinates for physics analysis. The neutrino beam provided by the West Area Neutrino Facility is inclined upwards (but not sideways, by construction of the detector) from the z -axis, and we use its nominal angle α to define new coordinates

$$\begin{aligned}x' &= x, \\y' &= y \cos(\alpha) - z \sin(\alpha), \\z' &= z \cos(\alpha) + y \sin(\alpha), \\ \alpha &= 42.5278 \text{ mrad}.\end{aligned}$$

Thus the z' -axis is the axis of the neutrino beam, and the $y' - x'$ plane is the plane whose normal is given by the the neutrino beam, henceforth the *transverse plane*.

Conservation of momentum dictates that if all outgoing particles in a neutrino interaction were detected, and their momenta ideally reconstructed, *and* that the struck nucleon had no transverse component of Fermi momentum, *then* the total momentum of the event in the transverse plane should be identically zero. Thus the “missing” transverse momentum defined at (3.1) above is the sum of the transverse momenta of all undetected final-state particles, with additional terms due to mis-reconstruction of detected particles, and Fermi motion. To guide the analysis we will think of \vec{p}_T^{miss} as the transverse momentum of a final-state neutrino, from a neutral current event or $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ event as the case may be. Effects of mis-reconstruction and Fermi momentum will be discussed in chapter 6.

We will use the subscript “T” throughout to label vectors, magnitudes of vectors, or other quantities formed solely from vector components in the $y' - x'$ or transverse plane.¹

3.2.3 Internal structure variables

The approximation depicted in figure 3.1 effectively treats the 3π and hadronic systems as massless particles, so that their three-momenta are the only relevant quantities to be measured. The simplest extension of this approach is to also consider the invariant masses of the two systems: $M_{inv}^{3\pi}$ for the 3π system, which we would expect to reflect the a_1^- mass in the case of $\tau \rightarrow \nu 3\pi$

¹There is one, very important exception in the analysis variable called q_T , described in section 3.3.3 below: the name is retained here for historical reasons, *e.g.* consistency with published NOMAD papers.

decay; and W for the hadronic system. Both quantities reflect the relative momenta of the tracks (and neutrals) belonging to their respective systems, and are sensitive to particle mass assignments, through the particle energy terms.

More generally, we will call any quantity which depends on the relative momenta of individual particles, or their mass assignments, or on selections or maximisations *etc.* over sets of tracks or neutrals, an *internal structure variable*. Variables of the 3π system are the prototype of such quantities: opening angles of the 3π , variables sensitive to the τ^- polarisation, and so on; these will be discussed in chapter 4. But it is also necessary to construct quantities which do not rely on the division of the event into 3π and hadronic systems: for example, it is only by such a process that a 3π system can be chosen in the first place, since events where at least one 3π system can be formed (*i.e.* events with $n(\pi^-) \geq 2, n(\pi^+) \geq 1$) typically admit of many 3π combinations. This problem will be treated in section 4.6 of the next chapter. Other internal structure variables, reflecting the shape of the event quite apart from the model of 3π - and hadronic-systems, will be introduced in chapter 6 in the course of suppressing the charged-current background.

3.2.4 Quality cuts

Most of the quantities mentioned at the beginning of section 3.2—error matrices, positions, numbers of drift chamber hits and so on—are relegated to the background of the analysis. We are assuming throughout the basic reliability of the reconstructed momenta *etc.*, relying implicitly on the previous experience of analysis, and tests of detector- and event reconstruction-performance, some of which have been published by the collaboration (Altegoer et al., 1998a, section 3). Where more basic information does come to the fore is in decisions as to whether the event as a whole should be analysed or discarded, or whether some reconstructed track or neutral should be included in or excluded from the analysis. We will discuss *quality cuts* of this kind in chapter 6.

3.2.5 Three-pion and three-prong events

In section 2.3.2 above we noted the existence of several “three prong” tau decays: the three-pion decay proper, $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ through the a_1^- meson; the modes including one or more neutral pion; and the relatively rare modes including kaons. In common with other studies to date, we will develop a consolidated three-prong analysis which is guided by the dominant three-pion decay, but applied to a search for all of the three-prong decays:

for example, except when a comparison is being explicitly made between the different decay modes, Monte Carlo samples of the “signal” will be mixtures of all of the three-prong modes in the appropriate ratio.

Clearly some loss of efficiency, or at least of precision, is implied by this approach. Where neutral pions are lost entirely from the analysis, they will contribute to the missing transverse momentum which our analysis ascribes to the final-state neutrino: the resulting distortion of the analysis will be relatively subtle. Tests which rely explicitly on the decay through the a_1^- meson are not strictly applicable to the other decays, but as we have already noted the a_1^- is a broad resonance centred at a large fraction of the τ^- mass, so that the 3π system in the other three-prong decays will often approximate $a_1^- \rightarrow \pi^- \pi^+ \pi^-$ behaviour. More serious is the case where one or both gammas from the neutral pion decay(s) are reconstructed as part of the hadronic system. Here both momentum and energy are being assigned to the “wrong half” of the event in the 3π - and hadronic-system picture.

In the absence of dedicated analyses for the neutral-bearing modes, we take the pragmatic approach that even a relatively low efficiency to detect these modes will increase the overall sensitivity of the search for tau decays. While undoubtedly there are gains to be made by a dedicated analysis of these cases—in particular the four-pion decay $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- \pi^0$ through the narrow ω resonance—there are also additional difficulties, notably the relative inefficiency of π^0 reconstruction, and the additional combinatoric problems due to the many neutral energy clusters in the calorimeter, which arise before the combination of $\pi^- \pi^+ \pi^-$ and π^0 is even considered. These problems place such an analysis beyond the scope of the present work.

In all that follows we will use the terms “ 3π ” or “three-pion” when it is appropriate to emphasise the system in hand, which is being subjected to analysis, and the term “three-prong” when the final aim of the analysis, *i.e.* sensitivity to tau decays, is in view. The misallocation of tau-decay π^0 to the hadronic system, discussed above, will be ignored.

3.3 Some variables previously used in three-prong decay analysis

As we turn to the choice of kinematic variables for the analysis, it is useful to review the most important quantities which have previously been used. Section 3.3.1 presents the transverse momentum and azimuthal angle variables introduced when the kinematic method was first proposed, and their adaption to the NOMAD case. Related quantities introduced within the

NOMAD collaboration, to mitigate some of the resolution problems inherent in the angular variables, are then discussed in section 3.3.2. Finally the important isolation variable q_T , which has become a key feature of most tau decay analyses within this experiment, is presented in section 3.3.3.

3.3.1 Angles in the transverse plane: $\phi_{3\pi-h}$ and ϕ_{m-h}

The original discussion of the kinematic-variable method envisaged a search for $\tau^- \rightarrow \nu_\tau \bar{\nu}_\mu \mu^-$ decays above a background of $\nu_\mu + N \rightarrow \mu^- + X$ interactions: in this case the ν_μ and ν_τ were to be produced by the decays of (respectively) D and D_s mesons in a proton beam dump (Albright and Shrock, 1979). In this situation the muon is the tau-decay candidate; assuming the ability of a detector to measure the direction as well as the energy of the hadronic system, one constructs the missing transverse momentum as discussed in section 3.2.1 above and makes cuts requiring

- transverse momentum imbalance greater than the resolution of the detector, say $|\vec{p}_T^{\text{miss}}| > 1 \text{ GeV}/c$;
- the muon at a large angle to the hadronic system in the transverse plane, say $\phi_{\mu-h} > 90^\circ$; and
- the missing momentum opposite to the hadronic system in the transverse plane, say $\phi_{m-h} > 120^\circ$.

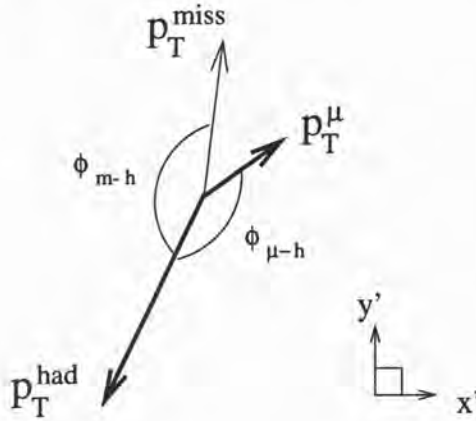


Figure 3.2: The transverse plane for $\tau^- \rightarrow \nu_\tau \bar{\nu}_\mu \mu^-$ analysis

These angles are shown, together with the type of configuration foreseen for a $\tau^- \rightarrow \nu_\tau \bar{\nu}_\mu \mu^-$ decay, in figure 3.2. The pair of neutrinos in the final state

is the source of both the large missing transverse momentum, and the strong anti-correlation between the directions of \vec{p}_T^{miss} and \vec{p}_T^{had} .

In the NOMAD proposal, a similar analysis was foreseen for the electronic decay $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$, above the relatively low $\nu_e + N \rightarrow e^- + X$ background, with the additional refinement that a contour on the plane (ϕ_{m-h}, ϕ_{e-h}) was used to separate signal- and background-regions (Astier et al., 1991a, figure 11, with the obvious $\phi_{\mu-h} \mapsto \phi_{e-h}$ substitution).

The extension of the analysis to three-prong decays involved a number of modifications, due to the dominance of the three-pion system over the ν_τ in such decays—the opposite of the situation for the leptonic decays, where the neutrinos $\nu_\tau \bar{\nu}_{\mu(e)}$ dominate the lepton $\mu^- (e^-)$. Thus the roles of the two azimuthal angles are reversed: one retains the structure of the analysis but performs the substitutions

$$\begin{aligned}\phi_{\mu-h} &\mapsto \phi_{m-h} \\ \phi_{m-h} &\mapsto \phi_{3\pi-h},\end{aligned}$$

so that the tighter cut is placed on $\phi_{3\pi-h}$; the various angles are shown in figure 3.3.




Figure 3.3: The transverse plane for $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$

In addition, the cut requiring large missing transverse momentum is replaced with empirically determined cuts requiring

- low missing transverse momentum, $p_T^{miss} < 1.6 \text{ GeV}/c$; and
- large transverse momenta for the 3π and hadronic systems, $p_T^{3\pi}, p_T^{had} > 1.6 \text{ GeV}/c$

(Astier et al., 1991b). These changes also reflect the need to suppress a *neutral current* background, where the typical p_T^{miss} is large, rather than a *charged current* background where p_T^{miss} is small, being an artefact of the event reconstruction.

3.3.2 Ratios of transverse momenta: $\rho^{3\pi}$ and ρ^{had}

During the course of analysis within the NOMAD collaboration, it became clear that angular variables such as ϕ_{m-h} suffer from a serious flaw. In the limit where $|\vec{p}_T^{miss}| \rightarrow 0$, ϕ_{m-h} depends on the direction of the infinitesimal \vec{p}_T^{miss} , so that events in the small region of phase-space with $\vec{p}_T^{miss} \approx \vec{0}$ are placed at arbitrary ϕ_{m-h} ; similar effects occur when $\vec{p}_T^{3\pi}$ or \vec{p}_T^{had} is small. There are thus dramatic variations in event topology for small movements on the plane $(\phi_{3\pi-h}, \phi_{m-h})$, with the resolution in $(\phi_{3\pi-h}, \phi_{m-h})$ becoming singular at the boundary of the physical region (Cousins, 1997a).

This problem is avoided if one uses ratios of transverse momenta, for example

$$\begin{aligned}\rho^{3\pi} &= \frac{p_T^{3\pi}}{p_T^{3\pi} + p_T^{had} + p_T^{miss}} \\ \rho^{had} &= \frac{p_T^{had}}{p_T^{3\pi} + p_T^{had} + p_T^{miss}}.\end{aligned}\quad (3.2)$$

These quantities are sensitive to both the relative size and orientation of \vec{p}_T^{miss} , while events with similar topologies are always close together in $(\rho^{3\pi}, \rho^{had})$; the only singularity occurs when all of $p_T^{3\pi}, p_T^{had}, p_T^{miss} \rightarrow 0$ (Cousins, 1997a).

At the same time, it became apparent that the empirically determined two-dimensional cuts in $(\phi_{3\pi-h}, \phi_{m-h})$ etc. (Astier et al., 1991b, figure 12 being the three-prong example), whose exact position and shape would always contain an arbitrary element, could be represented by simple cuts in ratios of transverse momenta (Geiser et al., 1996). As a result, description of the transverse plane in terms of the ratios (3.2) or similar quantities (see for example Altegoer et al., 1998c, section 9.1) has become favoured.

Transformation between $(\phi_{3\pi-h}, \phi_{m-h})$ and $(\rho^{3\pi}, \rho^{had})$

The information in the angles $(\phi_{3\pi-h}, \phi_{m-h})$ is reproduced in the ratios $(\rho^{3\pi}, \rho^{had})$: the transformations between them may be derived from the transverse momentum triangle relation

$$\vec{p}_T^{3\pi} + \vec{p}_T^{had} + \vec{p}_T^{miss} = \vec{0} \quad (3.3)$$

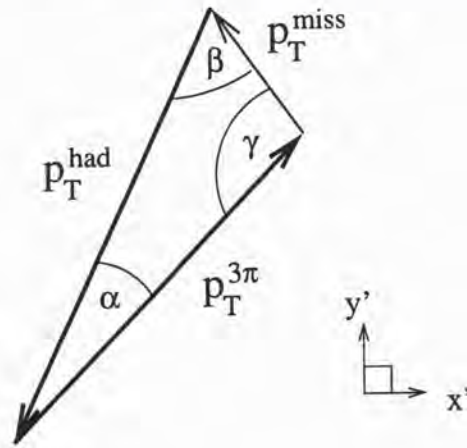


Figure 3.4: The transverse momentum vectors of figure 3.3 rearranged to illustrate the momentum triangle relation, with angles α, β, γ shown (their relation to the angles of the previous figure is given at (3.4)).

(*c.f.* (3.1)), defining interior angles

$$\begin{aligned}\alpha &= \pi - \phi_{3\pi-h} \\ \beta &= \pi - \phi_{m-h} \\ \gamma &= \pi - \alpha - \beta\end{aligned}\tag{3.4}$$

as shown in figure 3.4. The sine rule then immediately gives

$$\begin{aligned}\rho^{3\pi} &= \frac{\sin \beta}{\sin \alpha + \sin \beta + \sin \gamma} \\ \rho^{\text{had}} &= \frac{\sin \gamma}{\sin \alpha + \sin \beta + \sin \gamma},\end{aligned}$$

while setting

$$\begin{aligned}\rho^{\text{miss}} &= \frac{p_T^{\text{miss}}}{p_T^{3\pi} + p_T^{\text{had}} + p_T^{\text{miss}}} \\ &= 1 - \rho^{3\pi} - \rho^{\text{had}}\end{aligned}$$

we may write the inverse transformation using the cosine rule:

$$\begin{aligned}\cos \alpha &= \frac{(\rho^{\text{had}})^2 + (\rho^{3\pi})^2 - (\rho^{\text{miss}})^2}{2\rho^{3\pi}\rho^{\text{had}}} \\ \cos \beta &= \frac{(\rho^{\text{had}})^2 - (\rho^{3\pi})^2 + (\rho^{\text{miss}})^2}{2\rho^{\text{had}}\rho^{\text{miss}}}.\end{aligned}$$

(Following Cousins, 1997a, where the ρ variables were introduced).

3.3.3 The isolation variable q_T

An independent line of development in the analysis, specific to hadronic decays of the tau, has sought to quantify the *isolation* of the 3π system from the rest of the event. A very early expression of the idea can be found in Sciulli (1978): in the NOMAD experiment a critical study of the early proposals for analysis has led to the use of the variable we call q_T as a measure of the isolation of the 3π system.

In the early NOMAD study of the three-prong decay already mentioned, the final cuts on transverse momenta $p_T^{3\pi}, p_T^{had} > 1.6 \text{ GeV}/c$ yielded a rejection factor of over 3000 in the Monte Carlo then used (Astier et al., 1991b, table 2; the initial sample contained 3.7×10^5 events). In seeking the physical origin of this rejection it was noted that in the neutral current background, all the hadrons in the event are produced in the fragmentation of the struck quark and nucleon remnant (Astier et al., 1991c). The distribution of such hadrons in momentum perpendicular to the total hadronic system, $(1/p_\perp)(dn/dp_\perp)$, exhibits an approximately exponential falloff (*i.e.* equivalent to a Gaussian distribution in \vec{p}_\perp), as shown in figure 3.5 for the NOMAD data and full Monte Carlo.

This result is well-known from previous experiments (*e.g.* Arneodo et al., 1984) and a strongly falling distribution of p_\perp is expected from our understanding of the fragmentation process. In the Lund string model of fragmentation (Andersson et al., 1983), momentum of hadrons perpendicular to the total hadronic system is due to

- Gaussian p_\perp of $q\bar{q}$ pairs formed from the energy stored in the colour “flux tube” (the “string” of the model: the functional form follows from a quantum-mechanical “tunneling” argument);
- decays of unstable resonances into pions *etc.*;
- “primordial” transverse momentum k_\perp of the struck quark within the target nucleon;
- “soft” gluon emission; and
- first order QCD processes,

the last of which is not expected to be significant in the kinematic regime probed by NOMAD (see *e.g.* Jones et al., 1984). The first four sources contribute approximately Gaussian terms of $\langle p_\perp^2 \rangle \approx 0.2 \text{ GeV}^2/c^2$ or less. In NOMAD, where the most important constituents of the neutrino beam target




Figure 3.5: The distribution $\frac{1}{p_{\perp}} \frac{dn}{dp_{\perp}}$ for all tracks except the muon, measured relative to the hadronic system in identified $\nu_{\mu} + N \rightarrow \mu^{-} + X$ events in (a) the NOMAD data, and (b) the $\nu_{\mu} + N \rightarrow \mu^{-} + X$ Monte Carlo. (Compare the LEPTO simulation of Astier et al., 1991c, figure 5).

are carbon and oxygen nuclei, there is an additional contribution due to Fermi motion of the struck *nucleon*², which will be discussed further in chapter 6.

As a result, three-pion systems chosen from such events also exhibit an exponentially falling distribution of momentum transverse to the true hadronic system, which we can approximate by the total observed momentum in a reconstructed event (Astier et al., 1991c, figure 6; see also figure 3.6(b) below).

²The NOMAD Monte Carlo event generator is based on the LEPTO 6.1 program simulating the deep inelastic neutrino-nucleon scattering process (Ingelman, 1992), incorporating JETSET 7.4 to simulate the fragmentation of the hadronic system (Sjöstrand, 1994a,b); NOMAD-specific code treats the profile of the neutrino beam from the WANF (section 2.2 above) and Fermi motion effects in the target. First-order QCD processes are switched off in the code, while soft gluon effects are only partly taken into account, by broadening the p_{\perp} distribution used by JETSET. Some other parameters have been tuned to fit observed distributions.


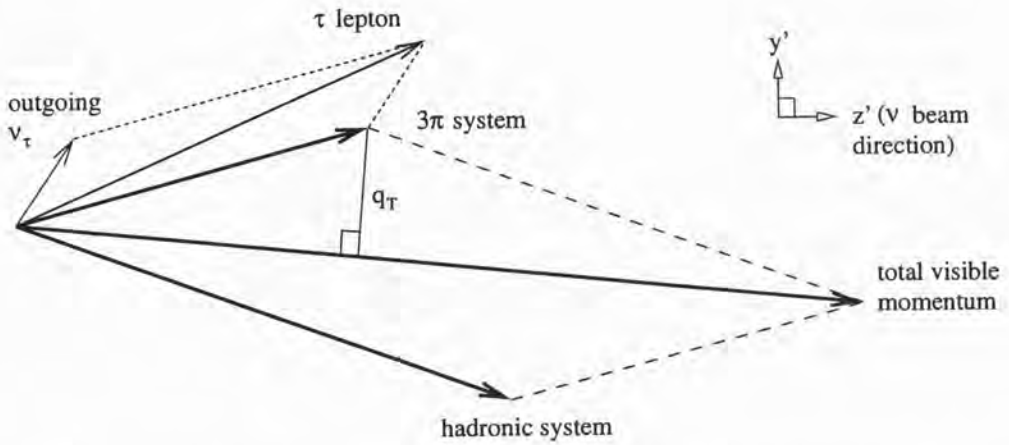


Figure 3.6: Distributions of q_T for Monte Carlo events from (a) the three-prong tau decay, and (b) the neutral current background sample.


For events which are “back-to-back” in the transverse plane (figure 3.3), such as those selected by a tight cut in $\phi_{3\pi-h}$ (e.g. Astier et al., 1991b, figure 12), momentum perpendicular to the total hadronic system (which we have been calling p_\perp in this section) is close to momentum transverse to the beam direction, p_T . The origin of the large rejection factor from transverse momentum cuts $p_T^{3\pi}, p_T^{had} > 1.6 \text{ GeV}/c$, referred to above, thus becomes clear (Astier et al., 1991c).

In three-prong decay events, on the other hand, the weak final-state neutrino ensures that

1. the total observed momentum is close in angle to the neutrino beam direction, and
2. the three-pion system carries a large fraction of the momentum, and also the transverse momentum, of the tau lepton,



(a) Three-prong tau decay event. The q_T of the 3π system reflects the momentum of the τ^- perpendicular to the neutrino beam, due to the deep inelastic scattering process.



(b) Neutral current event. The q_T of the "3 π system" reflects momentum of the individual pions perpendicular to the true hadronic system, resulting from the fragmentation process.

Figure 3.7: The physical origin of q_T is depicted for (a) three-prong tau decay and (b) neutral current events. The same outgoing lepton momentum (for the τ^- and ν_μ respectively), and the same true hadronic system momentum, have been used in each case. For simplicity, events with the main vectors lying entirely in the $y' - z'$ plane have been depicted. (This is a conservative choice: rotation about their own axes of the tau decay in (a) or the true hadronic system in (b) would *reduce* the resemblance between the events.)

as depicted in figure 3.7. As a result, the spectrum of momentum perpendicular to the total observed momentum for the three-pion system is hard (figure 3.6(a)).

The dramatic difference in both the fundamental origin and the distribution of this quantity has led to its use as a kinematic variable in its own right—as the most important single variable of the analysis—replacing the $p_T^{3\pi}$, p_T^{had} cuts. In keeping with the now-established NOMAD convention (Altegoer et al., 1998c) we will call this variable q_T , pausing to stress that the use of the subscript “T” is inconsistent with the usage of the rest of this work, where it denotes quantities of the plane perpendicular to the neutrino beam direction. The variable q_T does *not* lie in the transverse plane (except for the special case $\vec{p}_T^{miss} = \vec{0}$) and its relationship with quantities in that plane is complicated. This will emerge as an important point in the following section.

3.4 The choice of five kinematic variables

We wish to find a set of kinematic variables which is both complete, representing all of the information present in the approximation of $3\pi^-$ and hadronic-system “massless vectors” (see section 3.2.1 and figure 3.1), and effective at separating three-prong decay and neutral current events. Following the historical development just outlined, we begin with the transverse plane.

3.4.1 The transverse plane

Considering only quantities in the transverse plane, the event is represented at this level by three two-vectors: the transverse momenta of the three-pion and hadronic systems, and the “missing” transverse momentum, as shown in figure 3.8.




Figure 3.8: The transverse plane for $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$

Only two of these vectors are independent, as the missing transverse momentum is defined by imposing transverse momentum closure for the event (definition (3.1)), *i.e.*

$$\vec{p}_T^{miss} = -(\vec{p}_T^{3\pi} + \vec{p}_T^{had}).$$

So four numbers are sufficient to define the transverse momentum vectors. We are assuming that a rotation of the event about the neutrino beam axis has no effect on the physics, so subtracting one quantity (by rotating the event until one of the vectors is aligned with a predetermined axis, for example) it is apparent that the event has three degrees of freedom in the transverse plane.

As has already been noted, the ratios of transverse momenta introduced in section 3.3.2, describing the relative size of the vectors, are also sufficient to describe the angles between them, because of the momentum triangle relation (3.3) above. It remains only to specify the *size* of the vectors in the transverse plane, and so a suitable complete set is the pair of ratios $(\rho^{3\pi}, \rho^{had})$ together with the transverse momentum scale, *i.e.* we use

$$\begin{aligned}\rho^{3\pi} &= \frac{p_T^{3\pi}}{p_T^{3\pi} + p_T^{had} + p_T^{miss}} \\ \rho^{had} &= \frac{p_T^{had}}{p_T^{3\pi} + p_T^{had} + p_T^{miss}} \\ M &= p_T^{3\pi} + p_T^{had} + p_T^{miss}\end{aligned}\quad (3.5)$$

where the name “ M ” reflects the interpretation that massless particles radiating from a point with momentum vectors $\vec{p}_T^{3\pi}, \vec{p}_T^{had}, \vec{p}_T^{miss}$ would have an invariant mass equal to this quantity. We should stress that $\rho^{3\pi}$ and ρ^{had} are scalar quantities, constructed from the magnitudes of the transverse momentum vectors: the momentum triangle relation between those vectors (3.3) moreover places bounds on the ρ variables:

$$\begin{aligned}0 < \rho^{3\pi} < 0.5 \\ 0 < \rho^{had} < 0.5,\end{aligned}\quad (3.6)$$

and since the quantity $\rho^{miss} = p_T^{miss}/M = (1 - \rho^{3\pi} - \rho^{had})$ is similarly bounded, the sum of our two variables has a lower bound

$$(\rho^{3\pi} + \rho^{had}) > 0.5 \quad (3.7)$$

The physical region in $(\rho^{3\pi}, \rho^{had})$ is thus a triangle with corners at $(0, 0.5)$, $(0.5, 0)$ and $(0.5, 0.5)$.

3.4.2 The plane of momenta

Turning now to the isolation variable q_T , it is natural to consider the event not in the transverse plane but in the *plane of momenta*, the plane defined by the momentum vectors $\vec{p}^{3\pi}$ and \vec{p}^{had} , shown in figure 3.9:

If we consider relations within this plane, and ignore for the moment its orientation, then the event is a parallelogram with the 3π - and hadronic-system momenta along adjacent sides, and the total momentum along the diagonal emerging from their intersection; by construction, q_T is the momentum of the 3π system (or, equivalently, the hadronic system) perpendicular to the




Figure 3.9: The plane of momenta

total. This parallelogram has three degrees of freedom, of which q_T is one: it remains to make a suitable choice for the remaining two quantities.

From figure 3.7(a), it is clear that for tau decay events there will be a correlation between the value of q_T and the momentum of the 3π system, $p^{3\pi}$, since large values of both quantities are associated with energetic τ^- and 3π decay product. For the neutral current background we do not expect as strong a correlation, as some of the sources of hadron momentum perpendicular to the true hadronic system—specifically, p_\perp of $q\bar{q}$ pairs, and decays of resonances—are not correlated to the longitudinal momentum of the hadrons p_\parallel (section 3.3.3; see the discussion in Andersson et al., 1983, sections 3.2, 4.6.2).

Distributions of $p^{3\pi}$ in the Monte Carlo in bins of q_T , shown in figures 3.10 and 3.11, support this line of reasoning: the distribution of $p^{3\pi}$ for the signal is harder and broader than the distribution for the background, and its evolution with q_T is more pronounced, especially in the high- q_T region (figure 3.11). So $p^{3\pi}$ is a suitable choice for the second variable.

The remaining degree of freedom may be described by p^{had} , or, considering the scale of the parallelogram to have been fixed by cuts in the other two variables, by some ratio such as $E^{3\pi}/(E^{3\pi} + E^{had})$. Geometrically this quantity is just as important as the other two, but kinematically it is less significant: in three-prong-decay and neutral current events alike, we are dealing with systems of limited invariant mass boosted forward, so that cases where either the 3π or the hadronic system are completely dominant in the laboratory frame will be unlikely. Indeed, by proposing variables such as (unsigned) p^{had} or $E^{3\pi}/(E^{3\pi} + E^{had})$ we have left a whole class of geometries—where one of $\vec{p}^{3\pi}$, \vec{p}^{had} emerges at greater-than-90° to their sum—out of consideration.

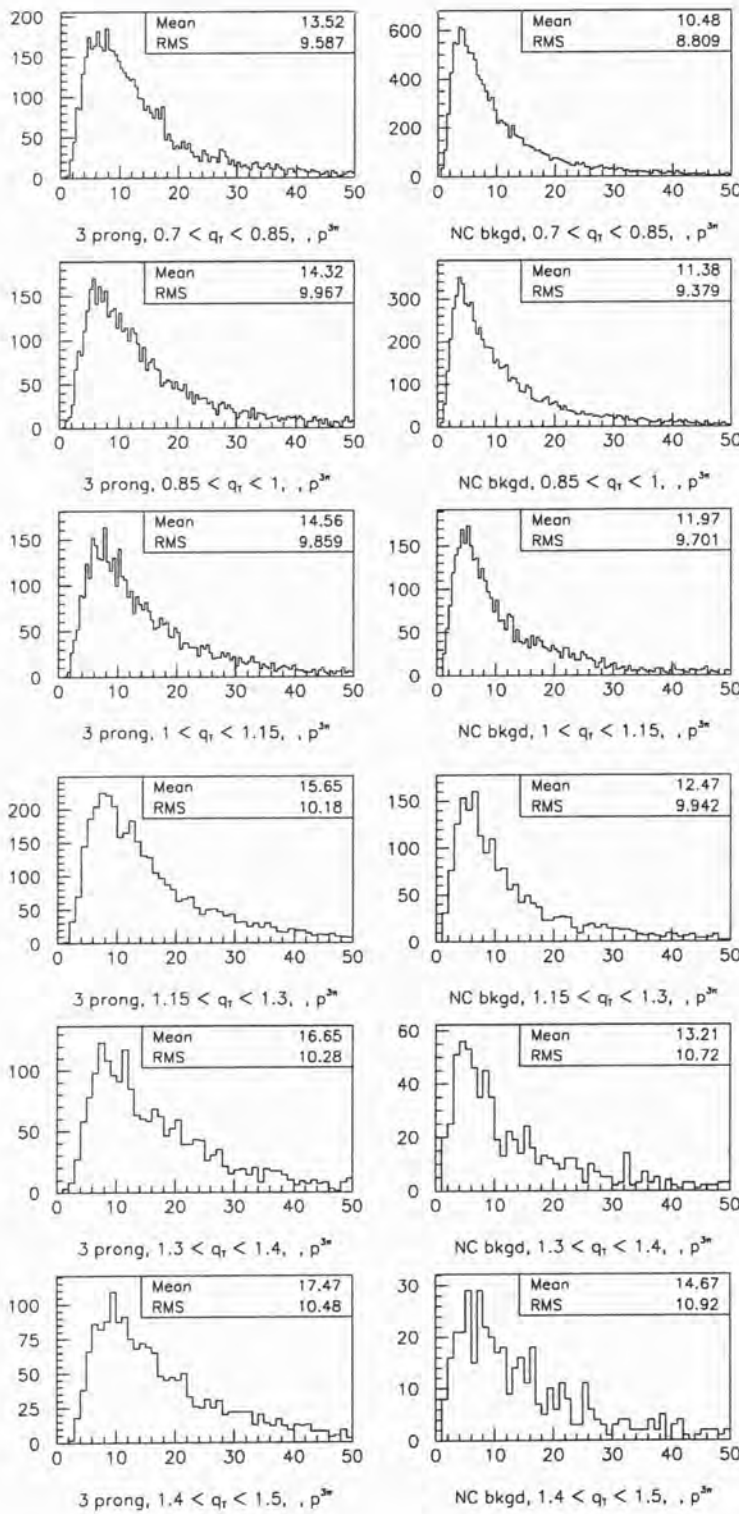


Figure 3.10: Distributions of 3π -system momentum for signal and background events in the Monte Carlo, plotted in bins of q_T .

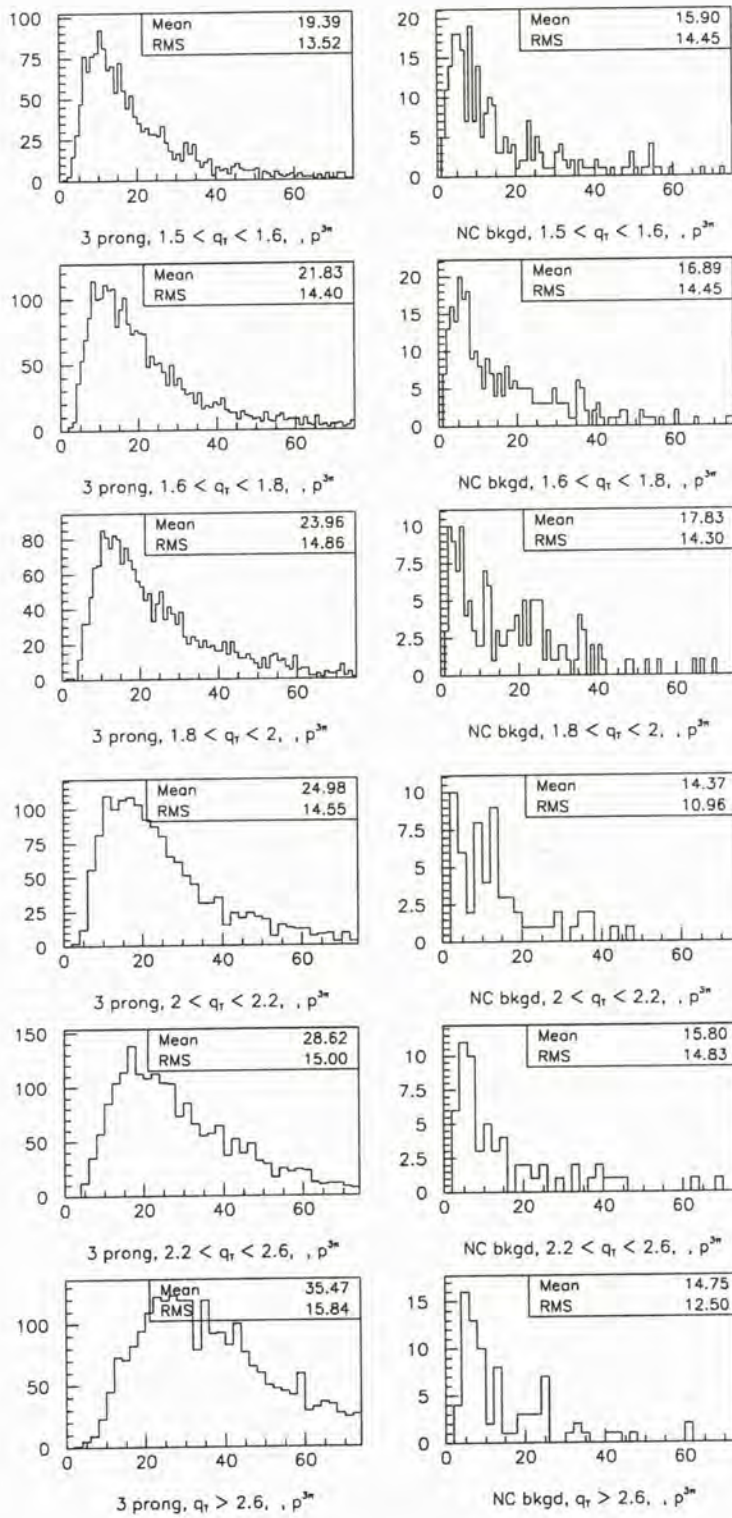


Figure 3.11: Distributions of 3π -system momentum for signal and background events in the Monte Carlo, plotted in bins of q_T . Continued from figure 3.10 with a change of horizontal scale.

3.4.3 The correlation between q_T and $(\rho^{3\pi}, \rho^{had})$

Reviewing these two descriptions of the event—by three quantities in the transverse plane, and another three quantities in the plane of momenta—it is clear that some redundancy exists, as there are only five degrees of freedom in the event at this level (section 3.2.1). More precisely, since the two sets of three variables each cover their own systems, there must be some redundancy *between* the set $(\rho^{3\pi}, \rho^{had}, M)$ and the set $(q_T, p^{3\pi}, \frac{E^{3\pi}}{E^{3\pi}+E^{had}})$.

We begin by considering the correlation between q_T , the most important variable of the plane-of-momenta, and the $(\rho^{3\pi}, \rho^{had})$ variables of the transverse plane. Monte Carlo distributions of $(\rho^{3\pi}, \rho^{had})$ in bins of q_T are shown in figures 3.12 and 3.13.

The $(\rho^{3\pi}, \rho^{had})$ distributions for three-prong decays and neutral current events at similar q_T differ from each other, but there are no clearly-defined signal- or background-regions. Certain trends are evident:

- at low q_T the background is concentrated at the diagonal ($\rho^{3\pi} + \rho^{had} = 0.5$) boundary;
- as q_T rises there is an increasing concentration of the signal at high $\rho^{3\pi}$ and $\rho^{had} \approx 0.5$, leaving a signal-free region along the diagonal boundary.

This effect can be understood if we consider figure 3.3 and the definition of the ρ variables (3.2): events with $\rho^{3\pi} + \rho^{had} \rightarrow 0.5$ correspond to the back-to-back configuration in which \vec{p}_T^{miss} dominates the transverse plane and is at $\approx 180^\circ$ to the (collinear) $\vec{p}_T^{3\pi}$ and \vec{p}_T^{had} . For three-prong decays, large q_T events correspond to cases where the 3π system is energetic and has large transverse momentum, and so dominant ν_τ (and hence \vec{p}_T^{miss}) are disfavoured.

The distinction between signal and background is much more clear if we suppress the dependence on q_T , superimposing the distributions from figures 3.12 and 3.13 to form summary plots for signal and background, as shown in figure 3.14. Three-prong decay events are concentrated at the upper boundary $\rho^{had} \approx 0.5$ (“back-to-back” events with approximately collinear $\vec{p}_T^{3\pi}$ and \vec{p}_T^{miss}), while the neutral current events are concentrated along the diagonal boundary (p_T^{miss} -dominated events) and in the upper-left corner ($\rho^{3\pi}, \rho^{had} \approx (0, 0.5)$) (events where $p_T^{3\pi}$ is very small compared to the other transverse momenta).

From the plots in bins of q_T , it is clear that these distinct concentrations of signal and background are a function of the harder q_T -distribution of the signal: the background concentration at the diagonal boundary is largely a low- q_T phenomenon, while only in the highest bins of q_T are the $(\rho^{3\pi}, \rho^{had})$ -distributions of signal and background clearly separated. On the other hand,

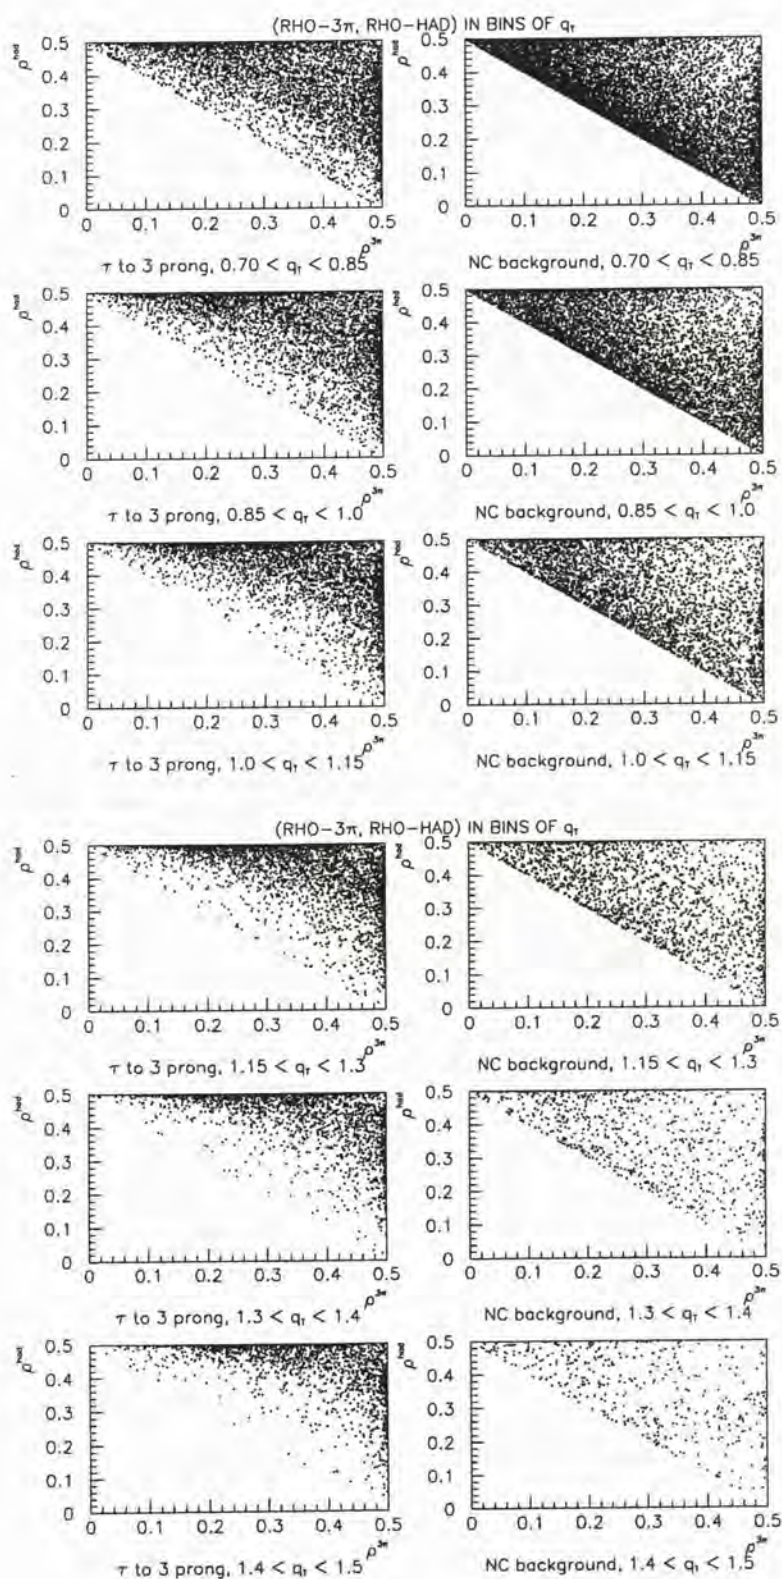


Figure 3.12: The $(\rho^{3\pi}, \rho^{had})$ plane for signal and background events in the Monte Carlo, plotted in bins of q_T .

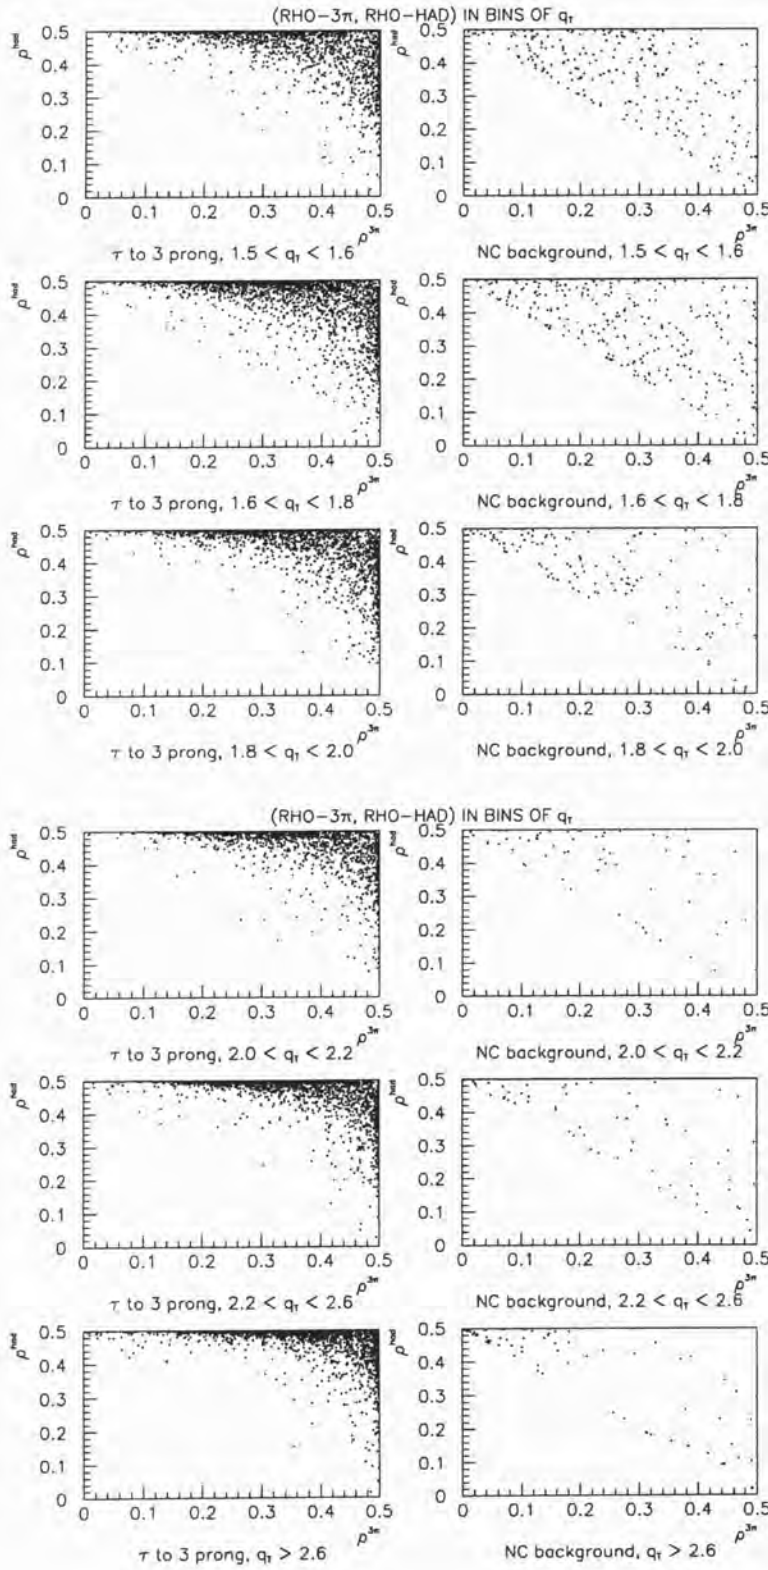


Figure 3.13: The $(\rho^{3\pi}, \rho^{had})$ plane for signal and background events in the Monte Carlo, plotted in bins of q_T , in the high- q_T region. Continued from figure 3.12.




Figure 3.14: Distribution in $(\rho^{3\pi}, \rho^{had})$ plane for signal and background events in the Monte Carlo, $q_T > 0.7 \text{ GeV}/c$, *i.e.* summing all of the distributions from figures 3.12 and 3.13.

the differences between signal and background in $(\rho^{3\pi}, \rho^{had})$ visible in figure 3.14 are not merely functions of q_T , as the distributions of $(\rho^{3\pi}, \rho^{had})$ in bins of q_T also differ. While a connection between the isolation variable and the ratios $(\rho^{3\pi}, \rho^{had})$ is therefore apparent, the way forward in the construction of the analysis is not yet clear.

3.4.4 The distribution of events in $(\rho^{3\pi}, \rho^{had}, M)$

One approach to this problem is to (temporarily) set the q_T variable to one side and consider the distribution of events in the space describing all configurations in the transverse plane: $(\rho^{3\pi}, \rho^{had}, M)$. A convenient way to do this is to plot the distribution of three-prong and neutral current events in $(\rho^{3\pi}, \rho^{had})$ in bins of the transverse momentum scale M , effectively taking a series of slices in the space $(\rho^{3\pi}, \rho^{had}, M)$. This succession of slices is presented in figures 3.15 and 3.16 below.

There is a dramatic separation of three-prong decay and neutral current events in this space, which can be described by the contours shown on the plots:

1. the diagonal contours separating high- and low-background regions; and
2. the horizontal contours separating higher- from lower-signal regions.




Figure 3.15: The $(\rho^{3\pi}, \rho^{had})$ plane for signal and background events in the Monte Carlo, plotted in bins of the p_T -scale M . Contours describing kinematic cuts are also shown: these are discussed in the text.

Each of these “moving” contours—tracing out a surface in the space $(\rho^{3\pi}, \rho^{had}, M)$ —may be described by a cut in a single variable, as described below for the two contours in turn.

The low- and high-background regions

The most obvious feature of figures 3.15 and 3.16 is the low-background region in the upper-right corner, *i.e.* at large $(\rho^{3\pi} + \rho^{had})$, which expands as the p_T -scale M increases. The division into low- and high-background regions is succinctly described by the diagonal contours shown on the plots, corresponding to a fixed cut in the variable

$$s_T \stackrel{def}{=} p_T^{3\pi} + p_T^{had} - p_T^{miss} \quad (3.8)$$

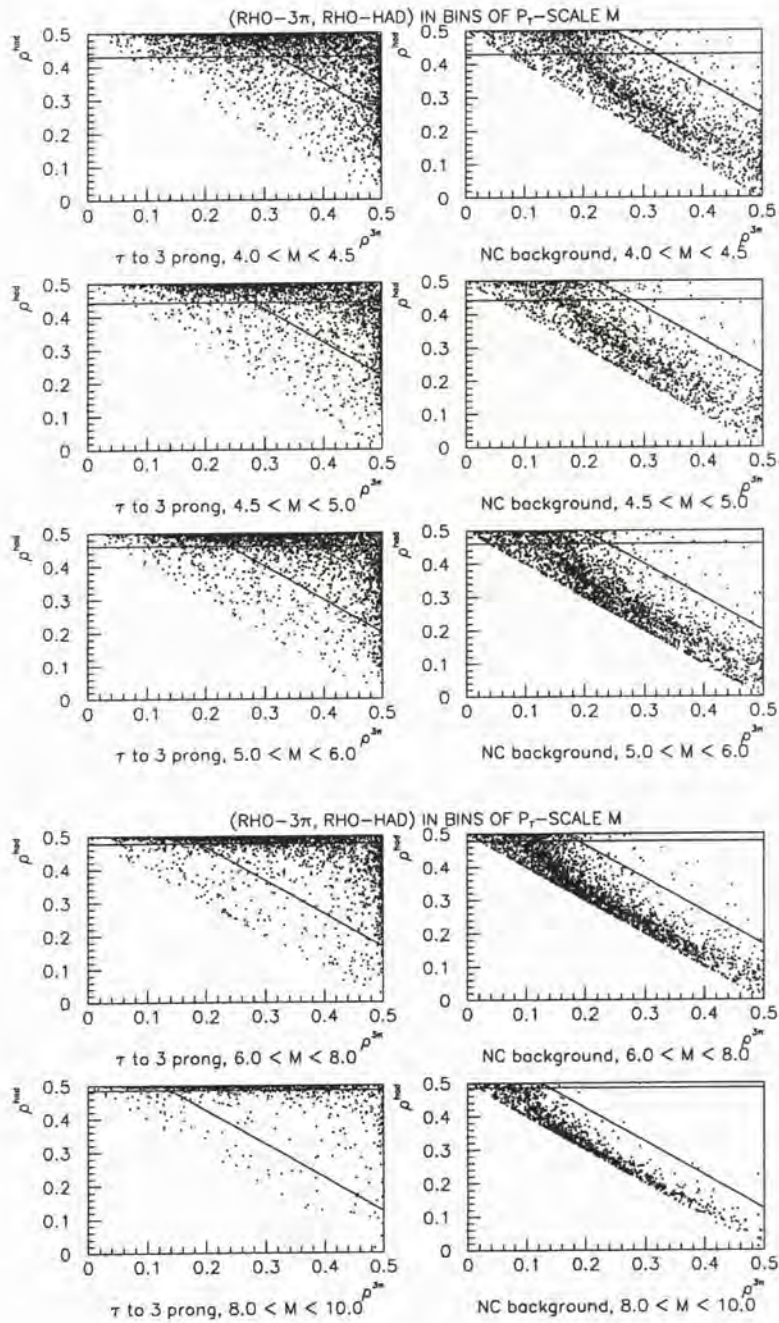


Figure 3.16: The $(\rho^{3\pi}, \rho^{had})$ plane for signal and background events in the Monte Carlo, plotted in bins of M . Continued from figure 3.15.

This quantity reflects the transverse component of the deep inelastic scattering process in ideal signal and background events, as depicted in figure 3.17:




Figure 3.17: Transverse momenta for ideal signal and NC events.

For ideal 3π decays, as shown in figure 3.17(a), the transverse momentum $\vec{p}_T^{3\pi}$ is close to that of the τ^- , so the sum $(p_T^{3\pi} + p_T^{had})$ approaches twice the transverse momentum of the deep inelastic scattering process ($p_T^\tau = p_T^{had}$); this quantity is then reduced by the relatively small p_T^{miss} to yield s_T . For ideal neutral current reactions, as shown in figure 3.17(b), the missing momentum dominates the transverse plane, and $(p_T^{3\pi} + p_T^{had}) \approx p_T^{miss}$, both quantities estimating the p_T of the final-state neutrino, so that $s_T \rightarrow 0$. (Compare with the respective events in figure 3.7.)

Of course, neutral current events resembling figure 3.17(a) are possible—given both a large energy-transfer y and a large hadronic system mass W —albeit highly suppressed. The empirically-chosen cut shown in figures 3.15 and 3.16, at $s_T = 2.0 \text{ GeV}/c$, neatly separates the preferred from the suppressed cases. The corresponding cut on $(\rho^{3\pi}, \rho^{had})$ is increasingly loose as

M rises, since from (3.8)

$$\begin{aligned} s_T &= p_T^{3\pi} + p_T^{had} - p_T^{miss} \\ &= M (\rho^{3\pi} + \rho^{had} - [1 - (\rho^{3\pi} + \rho^{had})]) \\ &= M (2 [\rho^{3\pi} + \rho^{had}] - 1) \end{aligned}$$

so that the cut in s_T corresponds to a contour

$$[\rho^{3\pi} + \rho^{had}] = \frac{1}{2} \left(\frac{2.0}{M} + 1 \right), \quad (3.9)$$

slowly approaching the physical boundary as $M \rightarrow \infty$.

The lower- and higher-signal regions

Also apparent in figures 3.15 and 3.16 is a concentration of the signal at the upper boundary (large ρ^{had}), tending strongly to the limiting value $\rho^{had} \rightarrow 0.5$ as M becomes large. A division into lower- and higher-signal regions in the high- M plots may be effected by the horizontal contours shown, corresponding to a fixed cut in the variable

$$M_{\perp} \stackrel{\text{def}}{=} \sqrt{(p_T^{miss} + p_T^{3\pi})^2 - (\vec{p}_T^{miss} + \vec{p}_T^{3\pi})^2}. \quad (3.10)$$

This is a ‘‘transverse mass’’ variable for the 3π -missing momentum system: it constructs a pseudo-invariant-mass for the 3π and the missing momentum, here considered as the momentum of an escaping body (*i.e.* the final-state neutrino), using transverse momenta only.

For $\tau \rightarrow \nu 3\pi$ events, where $\vec{p}_T^{miss} \approx \vec{p}_T^{\nu\tau}$ up to effects of reconstruction errors (which may of course be substantial), this quantity is bounded above by the tau mass (a result proven in appendix B.4). For neutral current events the tendency to large $\vec{p}_T^{miss} \approx \vec{p}_T^{\nu}$, and opposite alignment of \vec{p}_T^{miss} and $\vec{p}_T^{3\pi}$ (*cf.* figure 3.17(b)), leads to large values of M_{\perp} , the natural limit being imposed by the CMS energy of the event, $\sqrt{s} \gg m_{\tau}$.

The relationship between M_{\perp} and $(\rho^{3\pi}, \rho^{had})$ is non-linear:

$$\begin{aligned} M_{\perp} &= \sqrt{(p_T^{miss} + p_T^{3\pi})^2 - (\vec{p}_T^{miss} + \vec{p}_T^{3\pi})^2} \\ &= \sqrt{M^2 ([1 - \rho^{3\pi} - \rho^{had}] + \rho^{3\pi})^2 - M^2 (\rho^{had})^2} \\ &= M \sqrt{(1 - \rho^{had})^2 - (\rho^{had})^2} \\ &= M \sqrt{1 - 2\rho^{had}} \end{aligned}$$

so that the cut in M_{\perp} , chosen at $M_{\perp} = 1.7 \text{ GeV}/c$ (reflecting the bound due to the tau mass), corresponds to a contour

$$\rho^{had} = \frac{1}{2} \left(1 - \left(\frac{1.7}{M} \right)^2 \right) \quad (3.11)$$

which approaches the physical boundary at $\rho^{had} = 0.5$ very rapidly as M becomes large. This reflects the distribution of the signal in $(\rho^{3\pi}, \rho^{had})$: covering the plot almost uniformly when M is low, but concentrated into a vanishing region as $M \rightarrow \infty$.

3.4.5 Correlations between q_T and cuts in the transverse plane

We now return to the question of correlation between the plane of momenta—specifically q_T —and the transverse plane. From the results of the previous section we take the distribution of events in the space $(\rho^{3\pi}, \rho^{had}, M)$ to be usefully described by the variables s_T and M_{\perp} , so we consider the correlation between these quantities and q_T .

Distributions of the transverse mass M_{\perp} are shown for three-prong decay and neutral current events in figure 3.18, using broad bins of q_T . As noted in the previous section, the two samples have different scales in M_{\perp} : the three-prong events exhibit a fast-falling tail above the tau mass, the neutral currents a broad distribution more than twice as large in the mean. This behaviour is consistent across all bins and completely dominates over the slow hardening of the distributions with q_T , which occurs in both signal and background cases. So the distributions of events in M_{\perp} are largely independent of the isolation variable.

A very different situation is apparent in figures 3.19 and 3.20, where s_T is plotted in narrow bins of q_T . There is a marked evolution of the distributions with q_T , particularly for the three-prong decays: the bulk of events form a peak in the region $q_T < s_T < 2q_T$.

Neutral current events have a continuous distribution from a peak at $s_T \approx 0$ through to a limit at $s_T \lesssim 2q_T$, so that the distribution also evolves with q_T , but is markedly different from signal events in the same q_T bin. Most of this difference disappears, however, if a cut in the transverse mass is also imposed. In figures 3.19 and 3.20 the shaded area shows the distributions after the cut $M_{\perp} < 1.7 \text{ GeV}/c$ used in the previous section—the distributions for three-prong decays and neutral current events in a given q_T bin are now very similar.

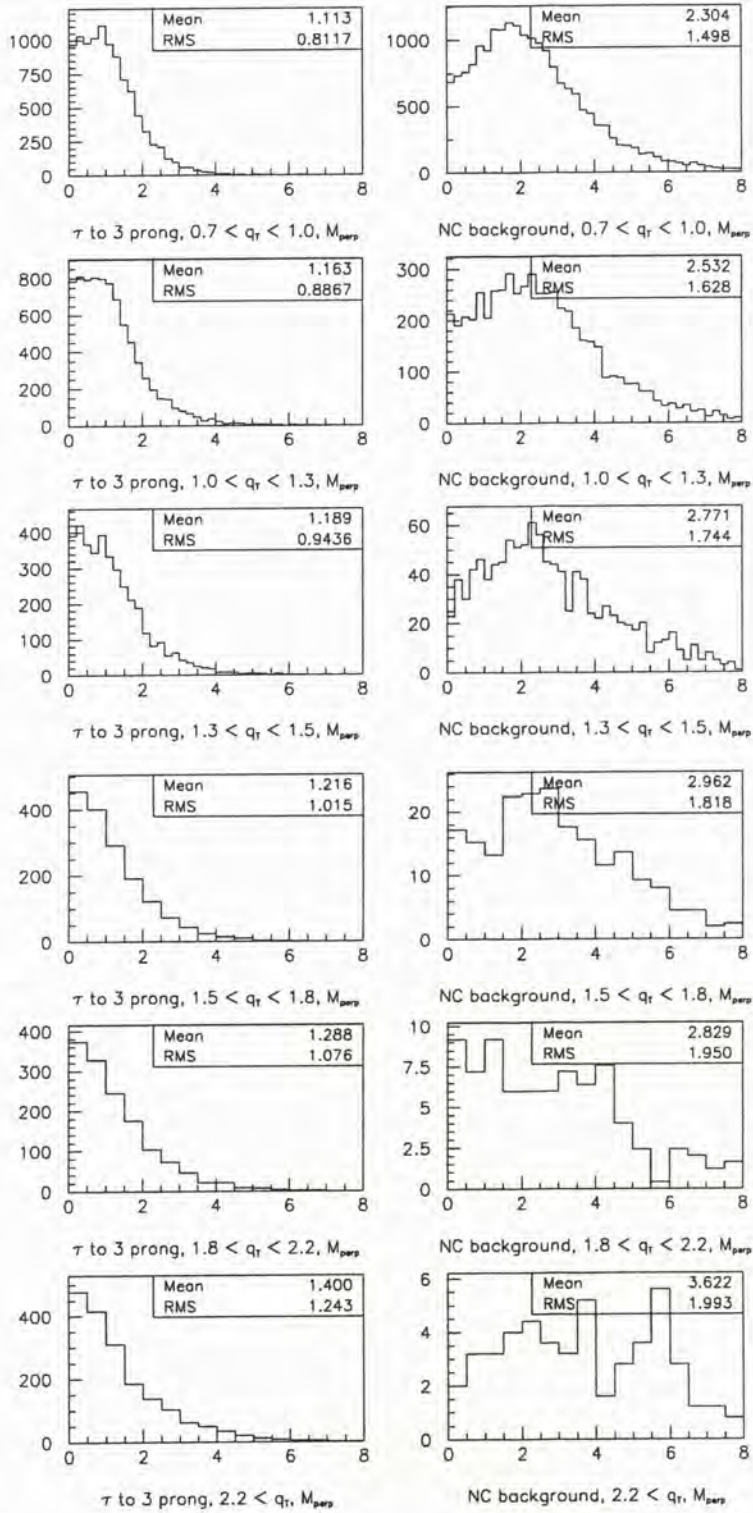


Figure 3.18: The transverse mass variable M_{\perp} for signal and background events in the Monte Carlo, plotted in broad bins of q_T .

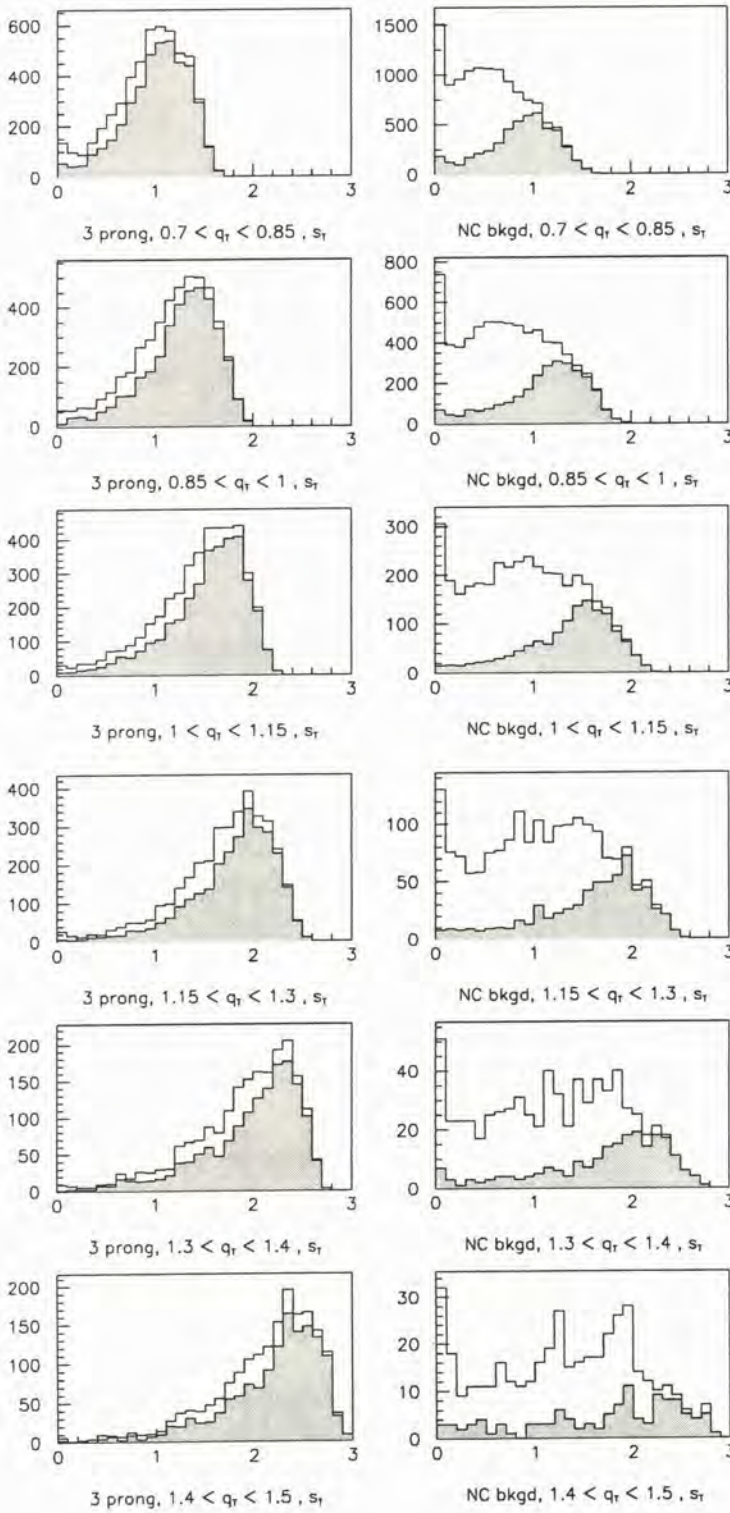


Figure 3.19: The transverse scattering variable s_T for signal and background events in the Monte Carlo, plotted in bins of q_T . The shaded region shows the effect of a cut $M_{\perp} < 1.7 \text{ GeV}/c$.

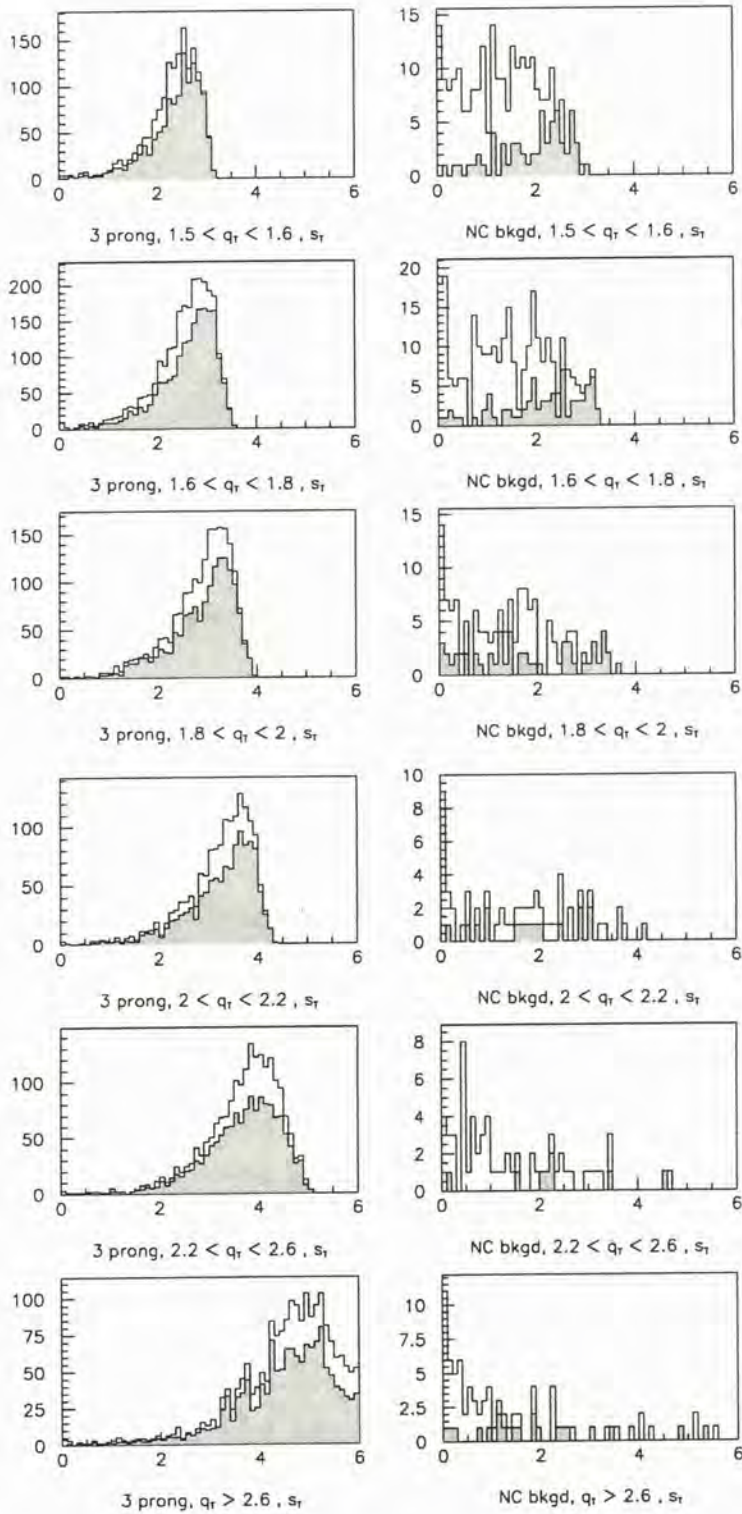


Figure 3.20: The transverse scattering variable s_T for signal and background events in the Monte Carlo, plotted in bins of q_T , for the high- q_T region. The shaded region shows the effect of a cut $M_{\perp} < 1.7 \text{ GeV}/c$. Note the change in horizontal scale from figure 3.19.

In retrospect the reason for the correlation between s_T and q_T is clear. Where the conditions that

1. $p_T^{3\pi}$ is large, and
2. $M_{\perp} < 1.7 \text{ GeV}/c$, or some similar restrictive value,

are imposed, then from the definition (3.10) the missing transverse momentum \vec{p}_T^{miss} is constrained to be small, or to be approximately aligned with the $\vec{p}_T^{3\pi}$, or both. We have effectively imposed the first of these conditions (where “large” means comparable to the tau mass, or rather to the cut in M_{\perp}) by considering events with sizeable q_T .

In such a case we are pre-selecting events with configurations such as the “ideal $\tau \rightarrow \nu 3\pi$ ” case shown in figure 3.17(a). Once q_T is specified the scale of the $p_T^{3\pi}$, p_T^{had} is set and s_T is largely determined through the momentum-triangle relation (3.3). Once a transverse mass cut has been applied, a cut in s_T will be equivalent to requiring large q_T in the event—a requirement which we are taking as given. This then is the key redundancy between the plane-of-momenta ($q_T, p^{3\pi}, \frac{E^{3\pi}}{E^{3\pi} + E^{\text{had}}}$) and the transverse plane ($\rho^{3\pi}, \rho^{\text{had}}, M$)³

3.4.6 A unified description of the kinematic variables

Thus the transverse mass M_{\perp} is revealed as a natural variable of the analysis, while the marked division of signal and background in $(\rho^{3\pi}, \rho^{\text{had}}, M)$ summarized by the s_T variable (section 3.4.4) turns out to be an effect of the q_T and M_{\perp} distributions acting in tandem. The clarity of the division by s_T in $(\rho^{3\pi}, \rho^{\text{had}}, M)$ (figures 3.19 and 3.20), and the connection between s_T and q_T just discussed, raise the question as to why the correlation between q_T and $(\rho^{3\pi}, \rho^{\text{had}})$, shown in figures 3.12 and 3.13, is relatively obscure.

Rotation of the plane of momenta

To clarify this it is necessary to set up a single framework within which the five degrees of freedom of the event (in the “massless vectors” approximation) can be described. An appropriate framework is provided by the plane-of-momenta representation of the event, shown previously in figure 3.9: we

³In passing we note that the analysis presented in this section is a variant on the argument of (Astier et al., 1991c), noted in section 3.3.3 above, explaining the high rejection factors against neutral current events achieved in the initial study: the transverse mass cut takes the place of the cut in $(\phi_{3\pi-h}, \phi_{m-h})$, an s_T cut replaces the $p_T^{3\pi}, p_T^{\text{had}}$ cuts, with the exponentially falling distributions of q_T , much harder for the signal than the background, taken for granted in both cases.

treat the total visible momentum vector \vec{p}^{tot} as if it were initially aligned with the neutrino beam axis (the z' -axis of section 3.2.2) and then perform two rotations,

1. a rotation about the y' -axis, *tilting* \vec{p}^{tot} away from the z' -axis; followed by
2. a rotation about \vec{p}^{tot} , *twisting* the plane-of-momenta about its own axis,

to give the actual event. These successive rotations are shown in figure 3.21.

The angles θ and ϕ correspond to the standard *Euler angles* describing an arbitrary rotation in space with the origin fixed, as discussed in appendix A. (The third Euler angle, ψ , may be set to zero as we are ignoring the overall azimuth of the event.)

The angle θ , or perhaps more intuitively the missing transverse momentum $p_T^{miss} = |\vec{p}^{tot}| \sin \theta$, describing the tilt of the plane of momenta from the neutrino beam axis, and the angle ϕ describing its twist about its own axis, are respectively the fourth and fifth degrees of freedom of the event—the first three being the variables $(q_T, p_T^{3\pi}, \frac{E^{3\pi}}{E^{3\pi} + E^{had}})$ of the plane of momenta. Clearly the transverse mass


$$M_{\perp} = \sqrt{(p_T^{miss} + p_T^{3\pi})^2 - (\vec{p}_T^{miss} + \vec{p}_T^{3\pi})^2}$$

is a strong function of both p_T^{miss} and ϕ ; the remaining dependence, on the magnitude $p_T^{3\pi}$ leads us to an interesting physical effect which will be discussed in the next chapter (section 4.3). For the present we note that the transverse mass provides a suitable variable to describe differences between three-prong and neutral current events in these two degrees of freedom.


The relationship between M and q_T

Equipped with this picture of the event we return to the question of the correlation between the q_T and $(\rho^{3\pi}, \rho^{had})$ variables. From section 3.4.4 it is clear that the distribution of events on the $(\rho^{3\pi}, \rho^{had})$ plane is a very strong function of the transverse momentum scale M , and so we may recast the question of q_T - $(\rho^{3\pi}, \rho^{had})$ correlation as a question about the relationship between q_T and M . For the $p_T^{miss} = 0$ case shown in figure 3.21(a), the relationship between M and q_T is especially simple:


$$\begin{aligned} p_T^{3\pi} &= q_T \\ p_T^{had} &= q_T \\ M &= p_T^{3\pi} + p_T^{had} + p_T^{miss} \\ &= 2q_T. \end{aligned} \tag{3.12}$$



(a) The plane-of-momenta, for $p_T^{miss} = 0$



(b) After a rotation of θ about the y' -axis



(c) After a final rotation of ϕ about \vec{p}^{tot}

Figure 3.21: A generic event, presented as a rotation of its plane-of-momenta representation: the event is shown in $x' - z'$ projection on the left, and $y' - x'$ projection (*i.e.* the transverse plane) on the right. Beginning with the plane-of-momenta in a standard position (a), we first perform a rotation of θ about the y' -axis (b), and then a rotation of ϕ about the *rotated* z' -axis, *i.e.* the total momentum \vec{p}^{tot} (c). The overall azimuth is irrelevant so we set the third Euler angle ψ to zero, equivalent to determining in advance that \vec{p}^{tot} will be aligned with the x' axis, as shown.

For the “tilted” plane-of-momenta shown in figure 3.21(b), corresponding to $p_T^{miss} \neq 0, \phi = 0$, and assuming $p_T^{miss} \lesssim 2q_T$ (see appendix A.2 for the exact condition, and other details),

$$\begin{aligned} p_T^{3\pi} &= q_T \cos \theta + k p_T^{miss} \\ p_T^{had} &= q_T \cos \theta - (1 - k) p_T^{miss} \\ M &\approx 2q_T + 2k p_T^{miss}, \end{aligned} \quad (3.13)$$

where the factor $0 \leq k \leq 1$ represents the sharing of the momentum component parallel to \vec{p}^{tot} between the 3π - and hadronic systems: for the $p^{3\pi} = p^{had}$ case, $k = \frac{1}{2}$ and $M \approx 2q_T + p_T^{miss}$; relaxing this condition, a correction as large as $\pm \frac{1}{3} p_T^{miss}$ can easily arise. For $\phi = \pi$, the roles of $p_T^{3\pi}$ and p_T^{had} are reversed.

Even in this much-simplified version of the event it is clear that the position of an event in $(\rho^{3\pi}, \rho^{had})$ will be a strong function of the relative sizes of p_T^{miss} and q_T . Furthermore if the condition $p_T^{miss} \lesssim 2q_T$ is not met, then we have instead

$$\begin{aligned} p_T^{3\pi} &= q_T \cos \theta + k p_T^{miss} \\ p_T^{had} &= -q_T \cos \theta + (1 - k) p_T^{miss} \\ M &= 2p_T^{miss}, \end{aligned} \quad (3.14)$$

and q_T has no influence on the value of M .

Considering now a general event where $p_T^{miss} \neq 0, \phi \neq 0$, corresponding to figure 3.21(c), the situation becomes more complicated, but we can make the simplifying assumption $p^{3\pi} = p^{had}$ and consider the extreme case of $\phi = \frac{\pi}{2}$, where

$$\begin{aligned} p_T^{3\pi} &= \sqrt{q_T^2 + \left(\frac{p_T^{miss}}{2}\right)^2} \\ p_T^{had} &= \sqrt{q_T^2 + \left(\frac{p_T^{miss}}{2}\right)^2} \end{aligned}$$

and under the further assumption $p_T^{miss} \lesssim q_T$,

$$M = 2q_T + p_T^{miss} \left(1 + \frac{p_T^{miss}}{4q_T}\right) \quad (3.15)$$

A correction, nonlinear in p_T^{miss} , is thus introduced, and relaxing the assumption $p^{3\pi} = p^{had}$, the two corrections will interact in a complicated way. The size and exact form of the correction are functions of the Euler angle ϕ .

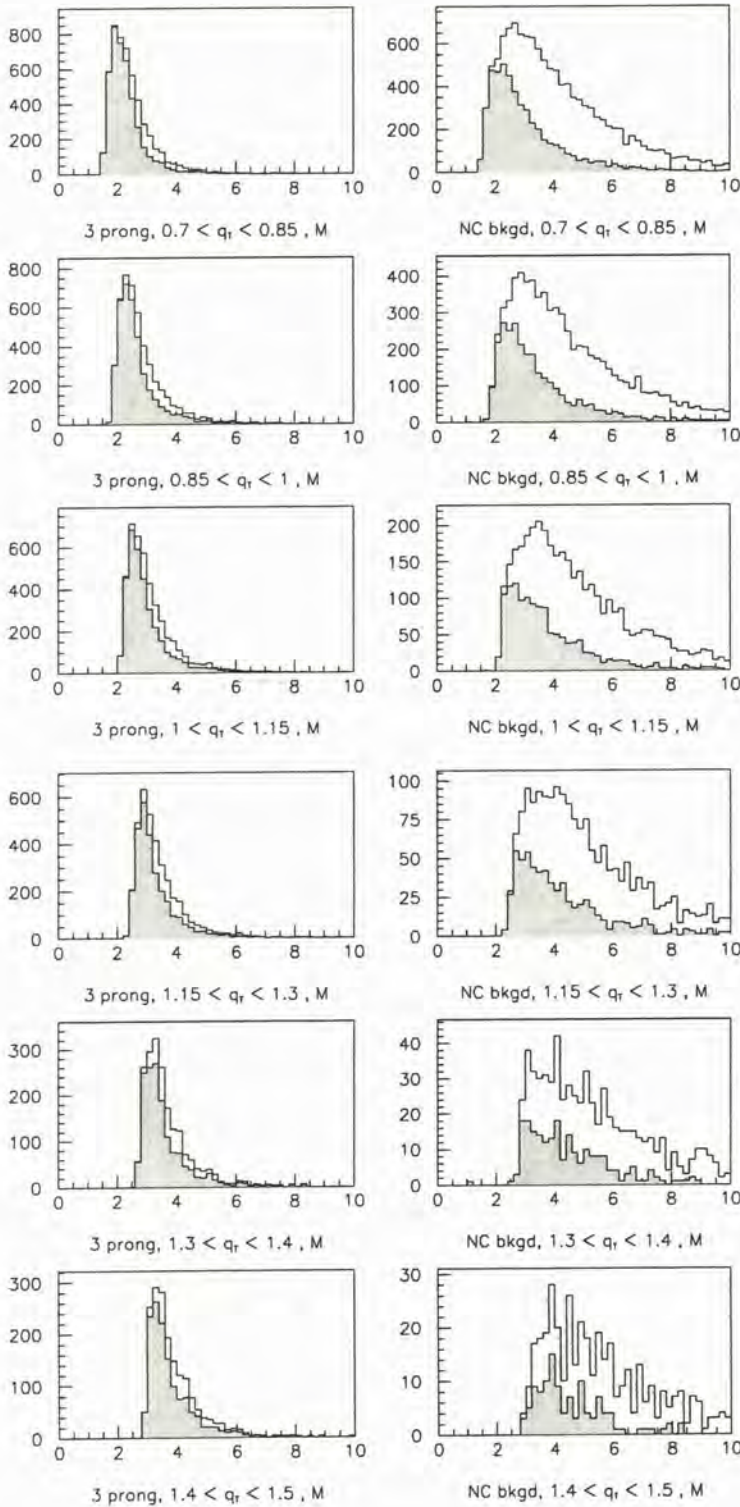


Figure 3.22: The transverse momentum scale M for signal and background events in the Monte Carlo, plotted in bins of q_T . The shaded region shows the effect of a cut $M_{\perp} < 1.7 \text{ GeV}/c$.

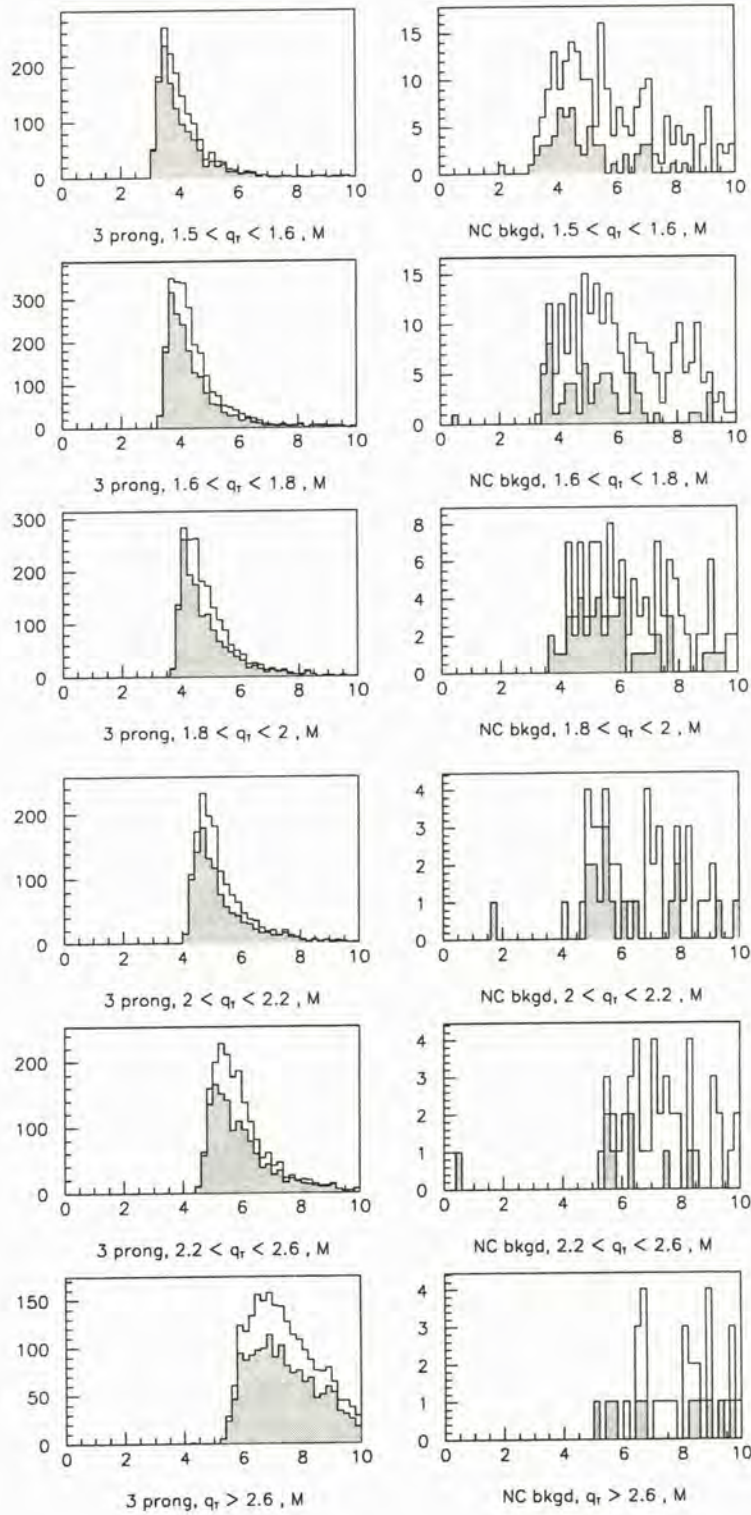


Figure 3.23: The transverse momentum scale M for signal and background events in the Monte Carlo, plotted in bins of q_T , for the high- q_T region. The shaded region shows the effect of a cut $M_{\perp} < 1.7 \text{ GeV}/c$.

The p_T -scale M is thus seen to be a strong function not only of q_T , but also of $p_T^{miss} = p^{tot} \sin \theta$ and the angle ϕ ; there is also a dependence, represented above by the factor k and the $\cos \theta$ terms, on the other degrees of freedom $p^{3\pi}$ and $\frac{E^{3\pi}}{E^{3\pi} + E^{had}}$. This is reflected in the broad distribution of M when plotted in bins of q_T , as shown in figures 3.22 and 3.23: especially in the neutral current, and even after a cut in the transverse mass, a tail to large values is apparent.

Summary

The dependence of the event distribution in $(\rho^{3\pi}, \rho^{had})$ on q_T , then, is more obscure than the dependence on M because

- in addition to q_T , the transverse mass M_{\perp} is an important and “natural” variable of the analysis, with a consistent, distinct and largely-independent-of- q_T difference of distribution between signal and background;
- the “footprint” of an M_{\perp} cut on the $(\rho^{3\pi}, \rho^{had})$ plane moves according to the *square* of M (section 3.4.4);
- M is a strong function of p_T^{miss} and the Euler angle ϕ as well as q_T , with an additional dependence on the remaining degrees of freedom,

so that the distinct signal- and background-regions in $(\rho^{3\pi}, \rho^{had}, M)$, shown in figures 3.15 and 3.16, are washed out when considering $(\rho^{3\pi}, \rho^{had})$ in bins of q_T alone (figures 3.12 and 3.13). Rather than attempting to recover this information through a complicated correlation of $(\rho^{3\pi}, \rho^{had})$ and q_T with other variables, it is expedient to use the transverse mass M_{\perp} as the principal variable to represent the fourth and fifth degrees of freedom, respectively p_T^{miss} (or θ) and ϕ .

3.4.7 Choice of a fifth kinematic variable

Finally, one more variable is required to describe the (p_T^{miss}, ϕ) degrees of freedom. In seeking such a variable we take our cue from the residual differences in the plots of the transverse scattering variable s_T (figures 3.19 and 3.20), and the p_T -scale M (figures 3.22 and 3.23), even after the M_{\perp} cut has been applied.

While for neutral current events the distribution of $s_T = p_T^{3\pi} + p_T^{had} - p_T^{miss}$ is slightly softer, the distribution of $M = p_T^{3\pi} + p_T^{had} + p_T^{miss}$ exhibits a tail to large values absent in the signal: comparison of the signs in the

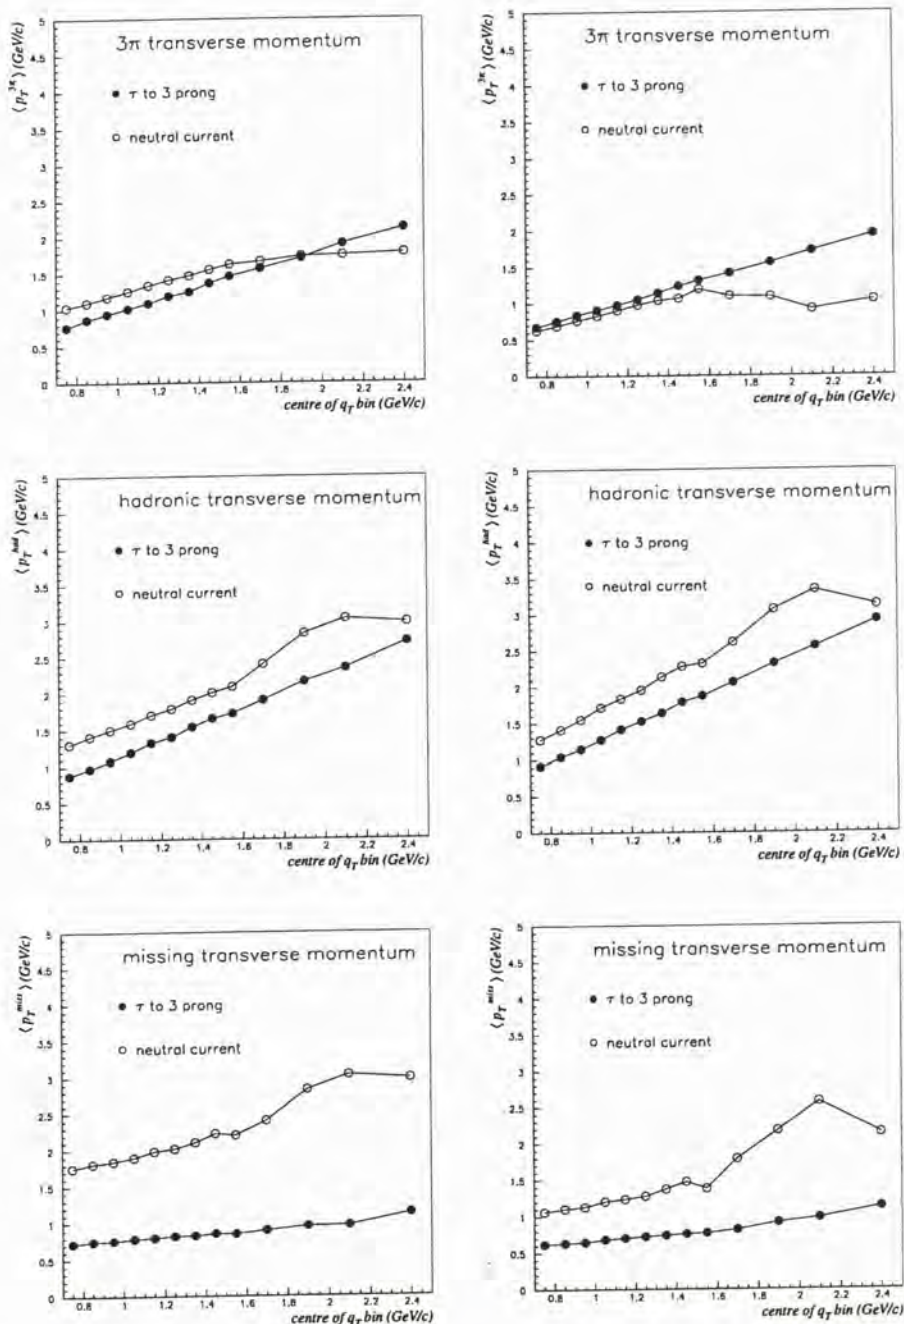


Figure 3.24: Mean values of the p_T variables in the Monte Carlo in bins of q_T , before (left plots), and after (right plots) a cut $M_{\perp} < 1.7 \text{ GeV}/c$ is imposed. Note that the surviving neutral current events become *softer* in $p_T^{3\pi}$ than the tau decays after the M_{\perp} cut is imposed. Missing transverse momentum in neutral current events is reduced by the cut, but remains larger than p_T^{miss} for the tau decays in the mean.

formulae for s_T and M suggest immediately a residual difference in p_T^{miss} between signal and background. If we consider the evolution of the individual p_T variables with q_T , with the transverse mass cut in place, we see that as q_T rises, $p_T^{3\pi}$ becomes slightly softer for neutral currents than for three-prong decays; both p_T^{miss} and p_T^{had} , on the other hand, remain systematically larger for neutral current events (figure 3.24). The latter differences add to produce the tail in M , while their near-cancellation limits the effect on s_T .

This pattern of differences suggests a combination of the transverse momenta such as

$$n_T = p_T^{had} + p_T^{miss} - p_T^{3\pi} \quad (3.16)$$

as a discriminator: this measure however has a very broad distribution in both signal and background (figure 3.25) and evolves with q_T . Dividing through by either $p_T^{3\pi}$ or M to obtain a scale-free variable yields a more satisfactory behaviour but in both cases the resulting measure maps one-to-one to $\rho^{3\pi}$:

$$\begin{aligned} \frac{p_T^{had} + p_T^{miss} - p_T^{3\pi}}{p_T^{3\pi} + p_T^{had} + p_T^{miss}} &= \rho^{had} + [1 - \rho^{3\pi} - \rho^{had}] - \rho^{3\pi} \\ &= 1 - 2\rho^{3\pi} \end{aligned}$$

while

$$\begin{aligned} \frac{p_T^{had} + p_T^{miss} - p_T^{3\pi}}{p_T^{3\pi}} &= \frac{\rho^{had} + [1 - \rho^{3\pi} - \rho^{had}]}{\rho^{3\pi}} - \frac{p_T^{3\pi}}{p_T^{3\pi}} \\ &= \frac{1}{\rho^{3\pi}} - 2 \end{aligned}$$

so for simplicity we choose $\rho^{3\pi}$ itself as the final summary kinematic variable.

Distributions of $\rho^{3\pi}$ are shown for signal and background in broad bins of q_T in figure 3.26: after the M_\perp cut the $\rho^{3\pi}$ distribution for three-prong decays depends only weakly on q_T , while for the neutral current there is a (desirable) q_T -dependent effect, with the $\rho^{3\pi}$ distribution becoming very soft for the small sample of large- q_T events. There are clear differences between the tau decay and neutral current distributions, but no signal- or background-free regions, so that the additional discrimination power obtained is weak, except for the very high- q_T region.

A method for exploiting the difference in distribution of $\rho^{3\pi}$ other than a crude cut, such as a likelihood analysis combining $\rho^{3\pi}$ with other variables, would be useful at this point. In fact it turns out that a simple cut in a related variable is sufficient for our purposes, but this requires that the $\tau \rightarrow \nu 3\pi$ decay structure be taken into account: we will take up this question in the following chapter (section 4.3).

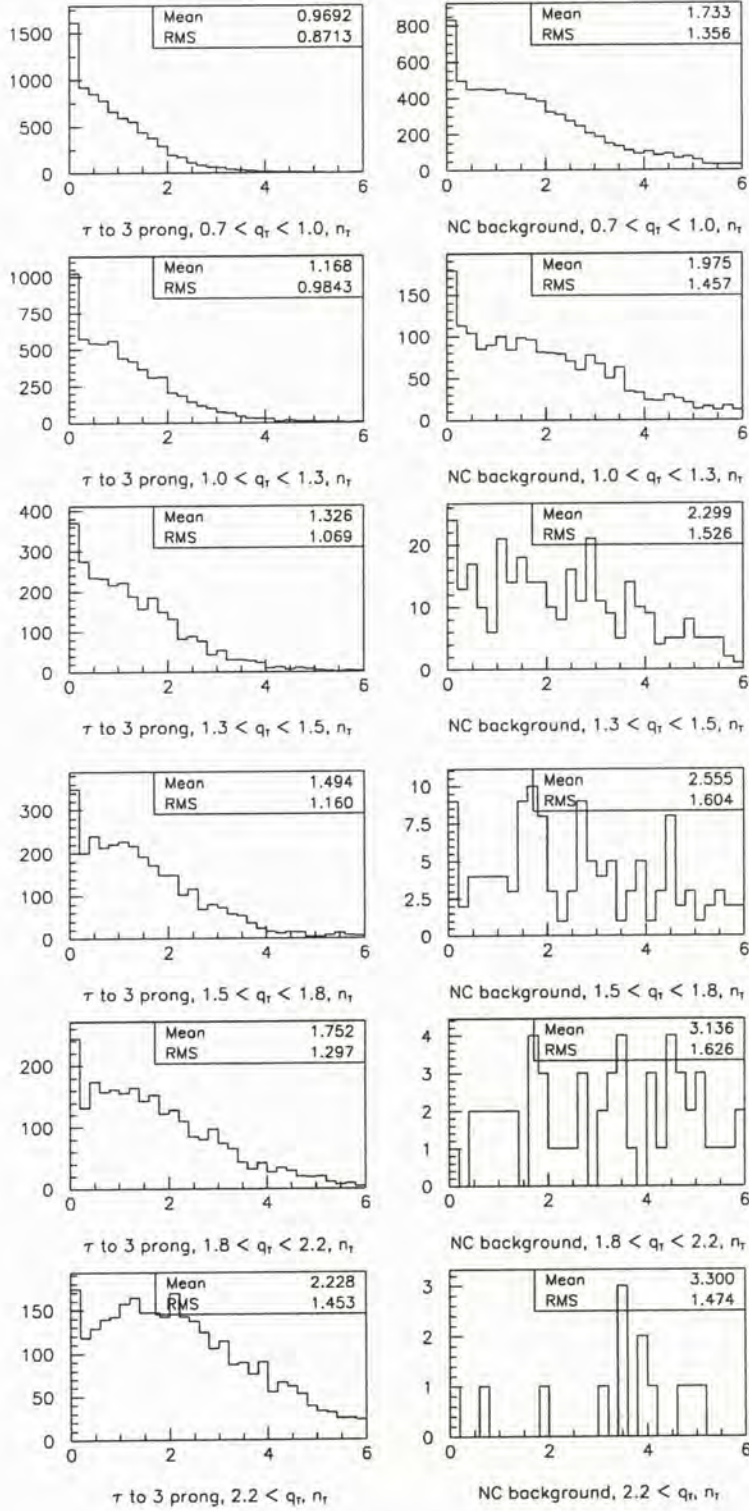


Figure 3.25: Distribution of the quantity $(p_T^{had} + p_T^{miss} - p_T^{3\pi})$ in broad bins of q_T , with a cut $M_{\perp} < 1.7 \text{ GeV}/c$ imposed.

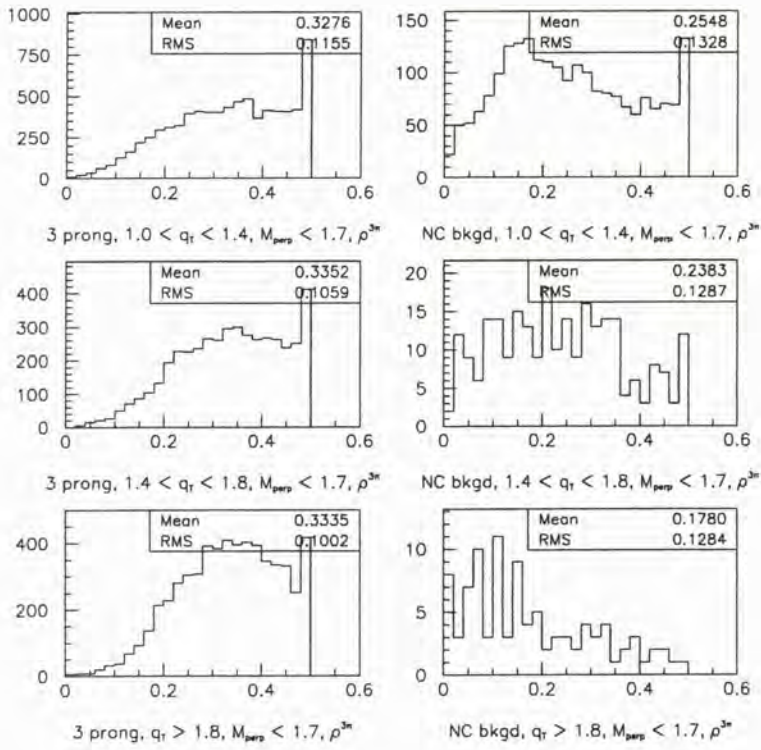


Figure 3.26: The $\rho^{3\pi}$ variable for signal and background in very broad bins of q_T , after a cut $M_{\perp} < 1.7 \text{ GeV}/c$.

3.5 Summary: The five kinematic variables

In the approximation where the internal structure of the three-pion and hadronic systems can be ignored, we describe the event by five summary kinematic variables, *viz.*

- variables describing the plane-of-momenta (see figure 3.9):
 - q_T , the isolation variable discussed in section 3.3.3,
 - $p^{3\pi}$, the three-pion system momentum, and
 - $E^{3\pi}/(E^{3\pi} + E^{had})$, the three-pion energy fraction;
- variables sensitive to its orientation (see figure 3.21):
 - M_{\perp} , the transverse mass of the 3π -missing- p_T system, and
 - $\rho^{3\pi}$, the three-pion transverse momentum fraction.

The three variables q_T , $p^{3\pi}$ and M_{\perp} are the most important for distinguishing three-prong tau decay events from the neutral current background.

The kinematic method for identifying tau decays was originally conceived using “transverse plane” variables: the magnitudes of significant transverse momenta, the angles between them, and other variables related to these quantities. We have found however that the transverse plane variables have a strong and rather complicated correlation with the the isolation variable q_T . Since q_T powerfully discriminates between tau decay and neutral current events, and furthermore has a clear and fundamental physical meaning, it is desirable to use this variable as the basis of our analysis. In this case we need to find a consistent set of variables to complete the task of separating the three-prong decays from the background.

The use of the q_T variable leads us to consider the plane of momenta—the structure describing the magnitudes and relative orientation of the 3π and hadronic-system momentum vectors, while ignoring their overall orientation. This accounts for three of the five degrees-of-freedom of the event in the “kinematic” approximation, and suggests the 3π momentum $p^{3\pi}$ as an important variable. By then considering the orientation of the plane of momenta in space, we can account for the remaining two degrees of freedom in a natural way, leading to a set of five variables whose internal correlations are minimised. As well as improving the clarity of the analysis, we find that the most important quantities—the isolation q_T , the three-pion momentum $p^{3\pi}$, and the transverse mass M_{\perp} —all have a clear physical motivation.

Chapter 4

Variables sensitive to the $a_1^- \rightarrow \pi^- \pi^+ \pi^-$ structure

4.1 Introduction

In the previous chapter we considered variables describing a simple approximation of a $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ candidate event: a pair of three-vectors, representing the momenta of the 3π and hadronic systems respectively. While there are several striking differences between true tau decays and the neutral current background in this approximation, it is equivalent to treating each of the 3π and hadronic systems as featureless, massless particles. In a true $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ decay, however, the 3π system has a well-defined structure: the tau lepton decays to the massive pseudovector meson a_1^- ,

$$\tau^- \longrightarrow \nu_\tau a_1^-,$$

which then itself decays to the observed three pions,

$$a_1^- \rightarrow \pi^- \pi^+ \pi^-.$$

We now turn our attention to variables which are sensitive to this structure, in order to improve the discrimination between tau decay events, and the neutral current and other backgrounds.

The most obvious implication of the decay through the a_1^- meson is that the three-pion system has a significant invariant mass—the mass distribution of the a_1^- is centred at $1.23 \text{ GeV}/c^2$ —and that this mass is a large fraction of the mass of the tau. Most of the quantities considered in this chapter rely, in one way or another, on this fact. In section 4.2 we consider the 3π invariant mass as a discriminator, along with the masses of pairs of pions from the 3π system, reflecting the so-called “substructure” of the a_1^- decay.

In section 4.3 it is shown that the *transverse mass* variable M_\perp (introduced in section 3.4.4 above) can be modified to take the 3π mass into account, resulting in an improved rejection of the neutral current background. The 3π mass also proves to be an important influence on the *opening angle* of the decay pions in the laboratory, which is discussed in section 4.4.

A more subtle effect concerns the polarisation of the tau lepton. The effect of any tau polarisation on the $\tau^- \rightarrow \nu_\tau a_1^-$ decay is complicated by the unit spin of the a_1^- meson: in contrast to the $\tau^- \rightarrow \nu_\tau \pi^-$ decay, where tau polarisation has a straightforward effect on the decay pion momentum, the many polarisation-sensitive variables in $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ are complicated functions of the momenta of the three pions. The feasibility of using these variables to discriminate between tau decay and background events is discussed in section 4.5.

Finally, in section 4.6, we consider the choice of the 3π decay candidate in a given event where (as is often the case) there is more than one $\pi^- \pi^+ \pi^-$ combination. A summary of the structure-sensitive quantities chosen for this analysis is then presented in section 4.7.

4.2 Invariant masses of the 3π system

Unlike the other principal decay channels of the tau lepton, the 3π decay through the a_1^- meson, and the three-prong decays in general, have decay products with a significant invariant mass, forming a large fraction of the mass of the tau itself. In sections 4.3 and 4.4 below, treating the transverse mass and opening angles of the 3π system respectively, some implications of this large mass are drawn out. However as we now consider the 3π invariant mass itself as a variable (section 4.2.1) it turns out to be more important that it has a natural *upper bound* at the tau mass. Three-pion systems selected from neutral current and other background events do not have this bound, making the invariant mass a powerful tool for discriminating against background.

The $a_1^- \rightarrow \pi^- \pi^+ \pi^-$ decay also has a more complicated structure than that due to the a_1^- mass alone. The most straightforward way to approach this structure is by considering the invariant masses of pairs of pions within the 3π system: we turn to this question in section 4.2.2.

4.2.1 The three-pion mass $M_{inv}^{3\pi}$

Forming the invariant mass of the three-pion system

$$M_{inv}^{3\pi} \stackrel{\text{def}}{=} \sqrt{(E_1 + E_2 + E_3)^2 - (\vec{p}_1 + \vec{p}_2 + \vec{p}_3)^2} \quad (4.1)$$

where \vec{p}_i, E_i are the momenta and energies of the three tracks, we expect to recover the a_1^- mass in the case of tau decays $\tau^- \rightarrow \nu_\tau a_1^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$. In the Monte Carlo we indeed see a peak near the a_1^- mass, with a small tail out to high values of $M_{inv}^{3\pi}$ due to reconstruction errors (figure 4.1). For three-prong decays with one or more accompanying π^0 , on the other hand, we expect in general lower 3π masses, because the tau energy is being partitioned among a larger group of particles. The distribution for $\tau \rightarrow \nu 3\pi + n\pi^0, n > 0$ decays in the Monte Carlo is also shown in figure 4.1, where in addition to the lower mean mass, we note a broader mass peak rising sharply from the lower bound $M_{inv}^{3\pi} \geq 3m_\pi$, resembling the distribution we would expect from the phase space of the decay alone, in the absence of a single object from which the three charged pions, and only the three charged pions, are produced.

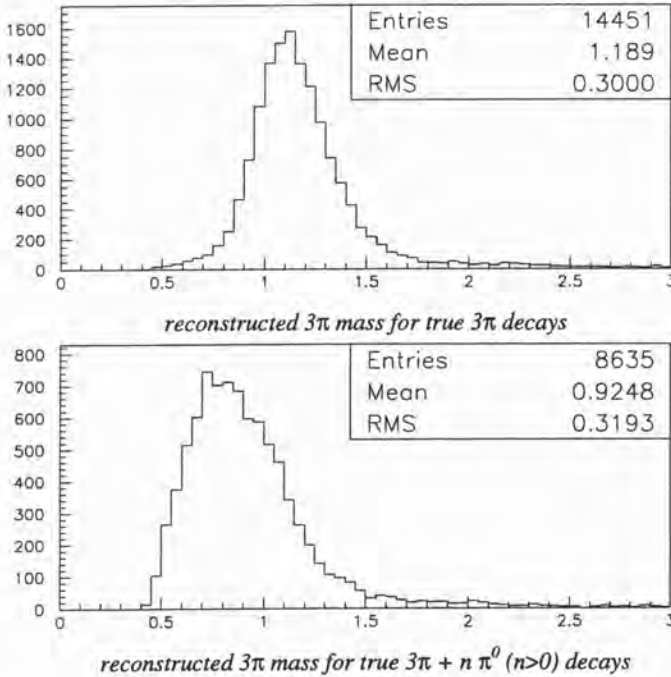


Figure 4.1: Reconstructed three-pion mass $M_{inv}^{3\pi}$ for $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ (upper plot) and $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n > 0)$ decays (lower plot) in the Monte Carlo.

The distributions shown may not however be recovered in the data, because an event where a 3π combination may be formed—with at least two negative and one positive track at the primary interaction vertex—in general admits of many 3π combinations. We must therefore select a 3π decay candidate from among these possibilities. A range of selection schemes has been used within the NOMAD collaboration, from a simple selection of the most energetic 3π combination (Astier et al., 1991b) to the selection of the 3π combination which maximises a likelihood function constructed from a basket of 3π -structure-sensitive variables (Altegoer et al., 1998c; Astier et al., 1999a). We will defer the discussion of this second approach to section 4.6, after the full range of structure-sensitive variables has been considered. For the present we will choose the 3π combination with the largest value of the dimensionless quantity q_{lep}/p^{had} , constructed as shown in figure 4.2. This quantity was found to select the true 3π combination more efficiently than the 3π momentum, in some early studies on event-level simulations of $\tau \rightarrow \nu 3\pi$ decays (*i.e.* prior to the full detector simulation being carried out).

With this simple choice of 3π combination we obtain the $M_{inv}^{3\pi}$ distribu-




Figure 4.2: Construction of the quantity q_{lep}/p^{had} for ranking 3π combinations. First we form the momentum of the 3π candidate perpendicular to the remaining event momentum, called q_{lep} as it is used as a measure of the isolation of the lepton in the $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ decay search (Altegoer et al., 1998c; Astier et al., 1999a). We then divide by the remaining, “hadronic system” momentum p^{had} to provide a scale for the momenta involved.




Figure 4.3: Reconstructed three-pion mass $M_{inv}^{3\pi}$ for the first-ranked 3π candidate in three-prong tau decay (upper plot) and neutral current events (lower plot).

tions shown in figure 4.3 for three-prong decays and neutral current events in the Monte Carlo. The mass rises to a noticeable peak in the three-prong sample, over a low spectrum extending to high values. In the neutral current sample, on the other hand, we see a much broader structure consistent with the origin of the 3π system as a combination of largely unrelated tracks.

Under cuts in the isolation variable q_T and the transverse mass the distribution for three-prong decays retains the peak near the a_1^- mass, which also narrows; the neutral current distribution, on the other hand, broadens and increases in mean value (figure 4.4). With the addition of an energy cut, there is a small effect on the three-prong distribution, while the additional broadening of the neutral current distribution is marked (figure 4.5). The difference in behaviour reflects the fact that passing kinematic cuts and having a mass near m_{a_1} are consistent in three-prong decay events, where a true 3π decay product exists, but not in neutral current events, where three particles from the true hadronic system are being artificially grouped.

We may take advantage of this fact to suppress the neutral current background while losing very little efficiency for true three-prong decays. If we modify the 3π selection procedure so that we first require a mass close to

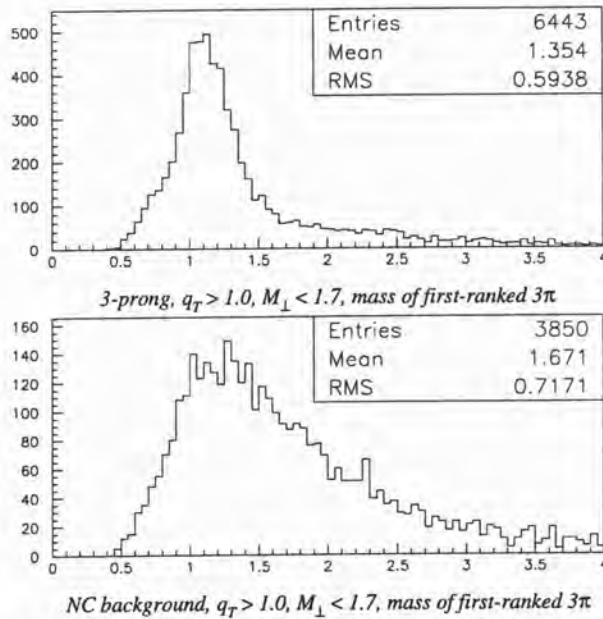


Figure 4.4: Reconstructed three-pion mass $M_{inv}^{3\pi}$ for the first-ranked 3π candidate under kinematic cuts $q_T > 1.0 \text{ GeV}/c$ and $M_\perp < 1.7 \text{ GeV}/c$ in three-prong tau decay (upper plot) and neutral current events (lower plot).

the peak in figure 4.4, $0.80 < M_{inv}^{3\pi} < 1.50 \text{ GeV}/c^2$, and then choose the surviving 3π combination with the highest value of q_{lep}/p^{had} (hereafter, the highest-ranked combination), we lose only 7% of the three-prong events surviving the kinematic cuts described, but reject more than 50% of the background (figure 4.6). Most of the three-prong decay events in figure 4.5 where the selected 3π combination has $M_{inv}^{3\pi} > 1.50 \text{ GeV}/c^2$ also contain other 3π combinations, high-ranked in q_{lep}/p^{had} and passing the kinematic cuts, but with masses in the window (0.80, 1.50): including, rather often, the true 3π decay product. There is no such effect for the neutral current background so that we retain most of the rejection power of a post-selection cut in the invariant mass. This rejection of the background, as well as the efficiency of the selection for three-prong decays, improve as the kinematic cuts are made more strict.

For the remainder of this chapter we will use the selection requiring $0.80 < M_{inv}^{3\pi} < 1.50 \text{ GeV}/c^2$, and then choosing the highest-ranked 3π combination, as our default, before returning to discuss the problem of 3π selection in section 4.6. We will also use the earlier selection, of the first-ranked




Figure 4.5: Reconstructed three-pion mass $M_{inv}^{3\pi}$ for the first-ranked 3π candidate under kinematic cuts, with a cut in the three-pion energy $E^{3\pi} > 10$ GeV/c added.




Figure 4.6: Reconstructed three-pion mass $M_{inv}^{3\pi}$ where a window cut in the mass, $0.80 < M_{inv}^{3\pi} < 1.50$ GeV/c², has been imposed as a condition prior to selecting the first-ranked 3π candidate. The distributions have been formed from the same event sample that were used for figure 4.5.

combination without restriction on the mass, from time to time when it is convenient to avoid imposing a mass bias.

4.2.2 Two-pion masses: 3π “substructure”

We noted in section 2.3.2 above that the a_1^- meson itself decays through an intermediate state $\rho^0 \pi^-$, rather than producing the final pions $\pi^- \pi^+ \pi^-$ directly.¹ By forming an invariant mass from a $\pi^+ \pi^-$ combination within the 3π candidate we therefore expect to be able to recover the rho mass distribution as a peak of moderate width ($m_\rho = 0.770, \Gamma_\rho = 0.151 \text{ GeV}/c^2$). It is not immediately clear which mass to choose, as two $\pi^+ \pi^-$ combinations are available: a simple approach is to form both masses and choose the one closest to the rho mass. This quantity, which we will call $M_\rho^{2\pi}$, is plotted for three-prong decays and neutral current events in the Monte Carlo in figure 4.7.

A peak in the $M_\rho^{2\pi}$ distribution close to the rho mass is apparent for the three-prong decays, but is also present for the neutral current background, albeit with an additional shoulder to low values. We are in fact producing this peak artificially by considering the two possible $\pi^+ \pi^-$ invariant masses and choosing the one closest to m_ρ . The situation becomes even worse when we apply restrictions on the 3π invariant mass in selecting our 3π candidate (figure 4.8) and the tail to high $M_\rho^{2\pi}$ values is removed in the neutral current.

There is in any case a physical objection to the picture on which this approach is based, *i.e.* making a definite choice as to which negative pion is the rho daughter. The two π^- in the final state are indistinguishable, so that both possible negative pion assignments in $a_1^- \rightarrow \pi^- \rho^0, \rho^0 \rightarrow \pi^+ \pi^-$ must be considered to occur (with some amplitude) and interfere.

We will therefore follow a conventional approach in the literature which is to form both $\pi^+ \pi^-$ invariant masses and order them in a consistent way, which does not explicitly involve the rho mass. Writing \vec{p}_1 and \vec{p}_2 for the momenta of the two π^- , and $\vec{p}_3 = \vec{p}^{\pi^+}$, we form the squares of the two-pion invariant masses

$$\begin{aligned} s_1 &= (E_2 + E_3)^2 - (\vec{p}_2 + \vec{p}_3)^2, \\ s_2 &= (E_1 + E_3)^2 - (\vec{p}_1 + \vec{p}_3)^2. \end{aligned} \quad (4.2)$$

Where \vec{p}_1^* and \vec{p}_2^* are the momenta of the negative pions in the rest-frame of the 3π , we choose the pion indices so that

$$|\vec{p}_2^*| > |\vec{p}_1^*| \quad (4.3)$$

¹The alternative decay $a_1^- \rightarrow \rho^- \pi^0$ leads to the one-prong decay product $\nu_\tau \pi^- \pi^0 \pi^0$, of comparable branching fraction to the $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ decay; see section 2.3.2.

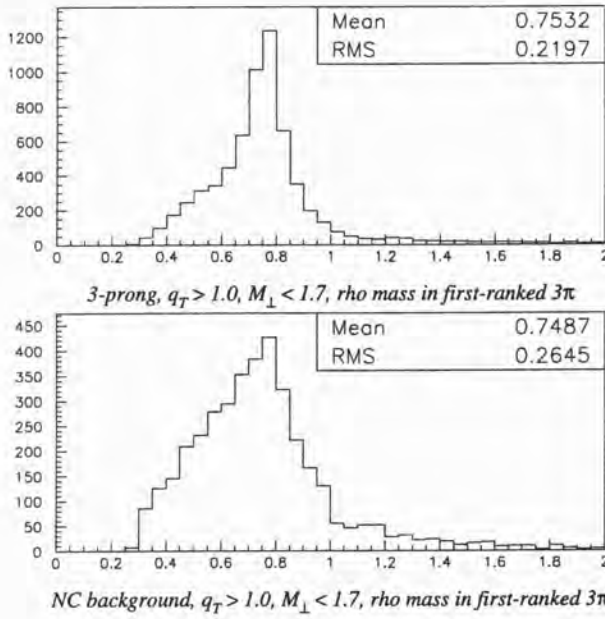


Figure 4.7: Reconstructed two-pion (rho) mass $M_\rho^{2\pi}$ for the first-ranked 3π candidate under kinematic cuts $q_T > 1.0 \text{ GeV}/c$ and $M_\perp < 1.7 \text{ GeV}/c$ in three-prong tau decay (upper plot) and neutral current events (lower plot).

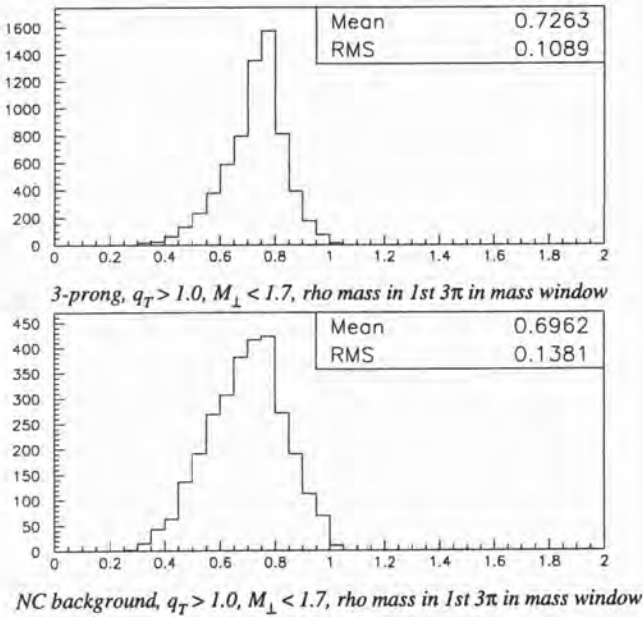


Figure 4.8: The two-pion (rho) mass $M_\rho^{2\pi}$ with a cut $0.80 < M_{inv}^{3\pi} < 1.50 \text{ GeV}/c^2$ imposed prior to candidate selection. Both signal (upper plot) and background (lower plot) spectra are more concentrated towards the central rho mass.

in which case,

$$s_1 > s_2 \quad (4.4)$$

always. The larger of these variables yields a mass $\sqrt{s_1}$ which is similar in distribution to the quantity $M_\rho^{2\pi}$ defined above, but with some differences (figure 4.9).

Distributions of the variables s_1 and s_2 are shown in figures 4.10(a) and (b) respectively, for both three-prong decay and neutral current events under kinematic cuts. A cut on the 3π mass has been imposed in the selection so these distributions show differences between the samples *in addition* to those described in the preceding section. Modest differences are apparent in both variables: the peak in s_1 near m_ρ^2 is less pronounced in the background, and the distribution is more broad; by contrast s_2 has a broad distribution in the three-prong decays but is concave to the origin in the neutral current,

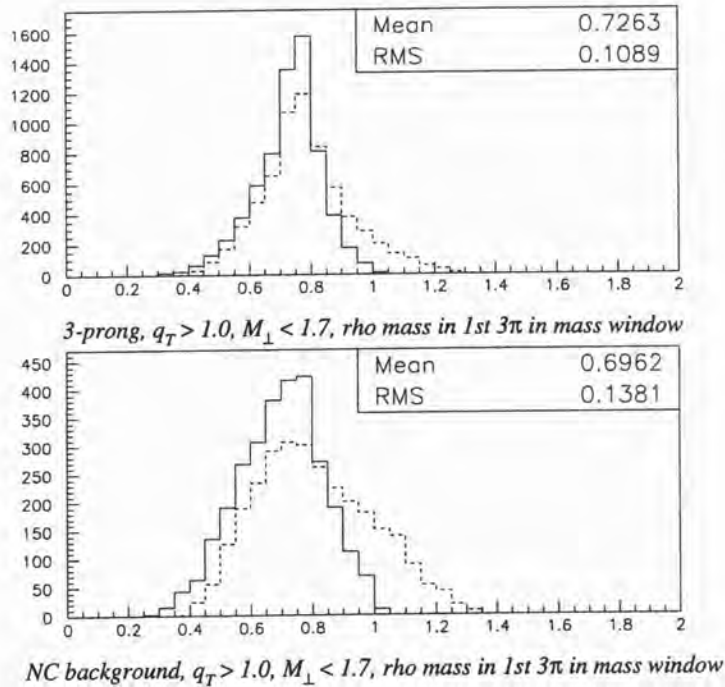


Figure 4.9: A comparison between the two-pion (rho) mass (solid line), and the larger of the possible masses, $\sqrt{s_1}$ (dotted line). A three-pion mass $0.80 < M_{inv}^{3\pi} < 1.50 \text{ GeV}/c^2$ has been required prior to candidate selection, and kinematic cuts $q_T > 1.0 \text{ GeV}/c$ and $M_\perp < 1.7 \text{ GeV}/c$ have been imposed.


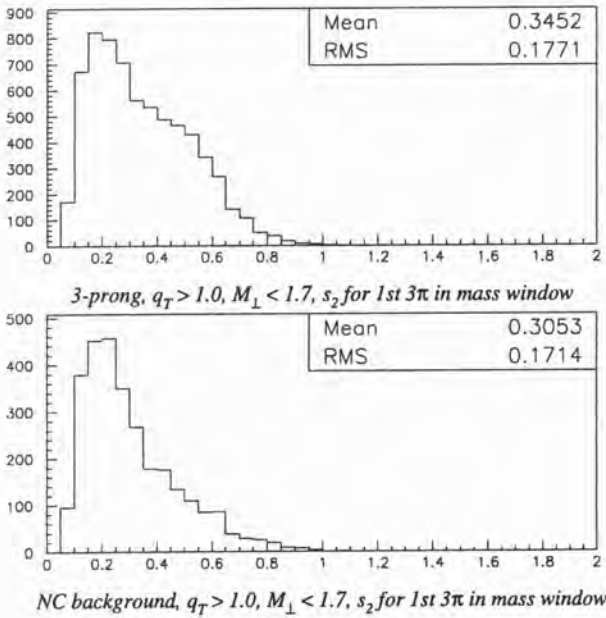
(a) The larger mass-squared, s_1 .(b) The smaller mass-squared, s_2 .

Figure 4.10: Distributions of the two-pion mass-squared variables for signal (upper plots) and background (lower plots). The $0.80 < M_{inv}^{3\pi} < 1.50 \text{ GeV}/c^2$ candidate selection, and cuts $q_T > 1.0 \text{ GeV}/c$ and $M_{\perp} < 1.7 \text{ GeV}/c$ have been imposed.

above the rise due to the threshold at $(2m_\pi)^2$. Both variables however cover the same range of values and are therefore unsuitable for use in simple cuts, unless large losses in efficiency for three-prong decays are countenanced. An appropriate way to exploit the differences in distribution would be to form a likelihood variable combining the two quantities, and perhaps $M_{inv}^{3\pi}$ itself, since there are clear differences between the two samples in its distribution, even after a cut in $M_{inv}^{3\pi}$ has been applied (figure 4.6). We will return to this question in section 4.7.

4.3 The choice of transverse mass variable

The three-pion mass is not the first mass-like quantity to have been introduced into this analysis: we have already selected the *transverse mass* M_\perp , as one of the one of the five kinematic variables describing the overall momenta of the 3π and hadronic systems. We now return to the transverse mass, considering it from the perspective of the $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ structure, and more specifically the structure of the 3π system.

4.3.1 The transverse mass M_\perp

As noted in the previous chapter, the direction of the neutrino beam is known, so the choice of this direction for the longitudinal or z' -coordinate axis implies that the total event momentum in the (x', y') plane—the *transverse momentum* \vec{p}_T^{tot} —should be zero up to Fermi motion effects. The so-called *missing transverse momentum*

$$\vec{p}_T^{miss} \stackrel{\text{def}}{=} -\vec{p}_T^{tot}$$

thus provides an estimate of the transverse momentum of any final-state neutrino \vec{p}_T^ν , with lost or mismeasured final-state particles contributing an error, often substantial, to this estimate. As the energy spectrum of the neutrino beam is broad, there is no corresponding estimate of the longitudinal momentum of the final-state neutrino p_L^ν .

It therefore makes sense to consider the “mass” of the combined three-pion/missing-momentum system, constructed from transverse plane quantities only, since we expect its value to reflect the tau mass in true $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ decays. In section 3.4.4 above we introduced the transverse mass

$$M_\perp = \sqrt{(p_T^{miss} + p_T^{3\pi})^2 - (\vec{p}_T^{miss} + \vec{p}_T^{3\pi})^2} \quad (4.5)$$

which indeed proved to be a powerful discriminator between three-prong decay events (with M_\perp values concentrated below the tau mass) and the

neutral current background (with a distribution of M_\perp peaking above the tau mass and extending to large values). While not a Lorentz invariant quantity, M_\perp is a longitudinal invariant—its value is unchanged under any boost along the neutrino beam direction—and its physical interpretation as a measure of the mass of the three-pion/missing-momentum system is clear.

4.3.2 The “required mass” M_R

A more careful approach to this question reveals that M_\perp is not in fact the most natural choice of transverse mass variable for the three-pion/missing-momentum system. By analogy with the construction of a true invariant mass, we consider quantities \mathcal{M} where

$$\mathcal{M}^2 = (E_T^{miss} + E_T^{3\pi})^2 - (\vec{p}_T^{miss} + \vec{p}_T^{3\pi})^2; \quad (4.6)$$

the relevant momentum vectors are shown in figure 4.11. That is, we identify the missing transverse momentum \vec{p}_T^{miss} with the transverse momentum of a final-state neutrino ν_τ , and form a mass for the presumed decaying object $\tau \rightarrow \nu 3\pi$ using the ν_τ and 3π transverse momenta only. It is natural to set $E_T^{miss} = p_T^{miss}$, since we are assuming a final state ν_τ with negligible mass. The internal structure of the 3π system, on the other hand, presents us with a number of choices for the three-pion transverse energy:

1. $E_T^{3\pi} = p_T^{3\pi}$, treating the 3π system as a massless vector, yields $\mathcal{M} = M_\perp$, the variable previously defined (3.10);
2. $E_T^{3\pi} = \sum_{\pi \in 3\pi} \sqrt{m_\pi^2 + (p_T^\pi)^2}$, summing the transverse energies of the constituent pions, yields another transverse mass we will call $\mathcal{M} = M_T$;
3. $E_T^{3\pi} = \sqrt{(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2}$, which yields a third variable $\mathcal{M} = M_R$ which we will call the *required mass*, for reasons which should become apparent.

In passing we should note that all three definitions are equivalent for the $\tau \rightarrow \nu \bar{\nu} e$ decay, due to the negligible electron mass $m_e \ll m_\tau$; similarly for the $\tau \rightarrow \nu \bar{\nu} \mu$ and $\tau \rightarrow \nu \pi$ decays, the distinction between M_\perp and M_T is small, while the variables M_T and M_R are equivalent. The variables M_T and M_R differ only for decays to multi-pion final states, $\tau \rightarrow \nu \rho$, $\tau \rightarrow \nu 3\pi$, and so on.

The choice leading to the variable M_T ,

$$E_T^{3\pi} = \sum_{\pi \in 3\pi} \sqrt{m_\pi^2 + (p_T^\pi)^2},$$


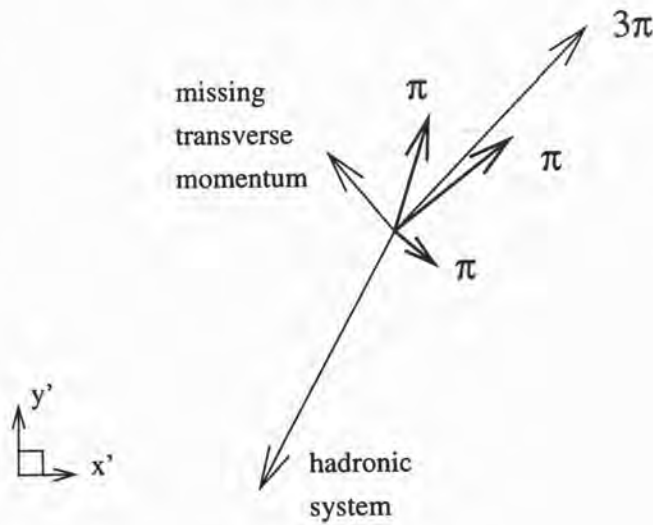
(a) $y' - z'$ projection(b) $x' - y'$ projection

Figure 4.11: Momentum vectors for calculation of the transverse mass. The hadronic system momentum is used indirectly, through the definition of the missing transverse momentum $\vec{p}_T^{miss} = -(\vec{p}_T^{3\pi} + \vec{p}_T^{had})$.

recommends itself as taking some account of the internal structure of the 3π system, which is ignored by the identification $E_T^{3\pi} = p_T^{3\pi}$. The choice leading to M_R ,

$$E_T^{3\pi} = \sqrt{(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2}, \quad (4.7)$$

introduces longitudinal momenta through $M_{inv}^{3\pi}$, but we note that since $M_{inv}^{3\pi}$ is a Lorentz invariant it is *a fortiori* a longitudinal invariant: put another way, $M_{inv}^{3\pi}$ is sensitive to the relative p_L of the three pions, but not the overall $p_L^{3\pi}$. This choice provides the most straightforward way to take the three-pion structure into account:

- it effectively treats the 3π system as a single particle of mass $M_{inv}^{3\pi}$ (and is thus a generalisation from, say, the $\tau^- \rightarrow \nu_\tau \pi^-$ case, where $E_T^\pi = \sqrt{m_\pi^2 + (p_T^\pi)^2}$, to $\tau^- \rightarrow \nu_\tau a_1^-$);
- the transverse energy² then satisfies the relation

$$(E_T^{3\pi})^2 = (E^{3\pi})^2 - (p_L^{3\pi})^2. \quad (4.8)$$

In an alternative approach, we may consider the identification of the missing transverse momentum \vec{p}_T^{miss} with the transverse momentum of an escaping neutrino \vec{p}_T^ν . On the hypothesis that the 3π system and this neutrino are the products of the decay of a particle of mass M , one can easily derive the longitudinal momentum of the neutrino p_L^ν . In practice, a solution for p_L^ν does not necessarily exist for fixed $M = m_\tau$: if we ask for the *minimum* possible mass M for which a solution exists—so that the identification $\vec{p}_T^{miss} = \vec{p}_T^\nu$ makes sense—we find that it is given by $M = M_R$. This rather fundamental derivation of M_R , which suggests the name “required mass”, is straightforward but lengthy, and is presented in appendix B.

For any event there is a strict hierarchy in the values of the different transverse masses,

$$M_\perp < M_T \leq M_R, \quad (4.9)$$

²In relativistic kinematics, discussions of parton fragmentation and so on, the quantity $\sqrt{m^2 + p_T^2}$ is called “transverse mass”, whereas in collider experiments *etc.* this name is reserved for pseudo-invariant masses constructed using transverse momenta from both visible and invisible particles—we of course follow this second convention. It is therefore necessary to have another name for $\sqrt{m^2 + p_T^2} = \sqrt{E^2 - p_L^2}$, and “transverse energy” seems the most natural. We allow the same name for the quantities $p_T^{3\pi}$ and $\sum_{\pi \in 3\pi} \sqrt{m_\pi^2 + (p_T^\pi)^2}$ above, as approximations to $\sqrt{(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2}$.

and on the assumption of perfect momentum reconstruction, including $\vec{p}_T^{miss} = \vec{p}_T^{\nu}$, the largest of these obeys the bound

$$M_R \leq m_\tau. \quad (4.10)$$

The comment in section 3.4.4 above that M_\perp is bounded by the tau mass follows from the combination of (4.9) and (4.10). Both results are proven in appendix B.

4.3.3 Relative performance of M_R and M_\perp

The required mass M_R has signal/background discrimination power superior to M_\perp in the Monte Carlo. It is instructive to show this using distributions of the transverse momentum fraction $\rho^{3\pi} = p_T^{3\pi} / (p_T^{3\pi} + p_T^{had} + p_T^{miss})$: plots of this variable after a cut $M_R < 2.2 \text{ GeV}/c$ are included as figure 4.12 below, with the effect of the transverse mass cut $M_\perp < 1.7 \text{ GeV}/c$ also shown. (The value $2.2 \text{ GeV}/c$ was chosen to approximately reproduce the signal-efficiency of the $M_\perp < 1.7 \text{ GeV}/c$ cut. Note that a mass cut $0.80 < M_{inv}^{3\pi} < 1.50 \text{ GeV}/c^2$ has already been imposed in the selection of the 3π combinations for these plots, so the effect of the differing mass spectra in three-prong and neutral current events is removed to first order. In particular, neutral current events are *not* failing the M_R cut because their invariant masses $M_{inv}^{3\pi}$ are too large.)

While the three-prong decay events are largely unaffected by the change from an M_\perp to an M_R cut, many of the low- and medium- $\rho^{3\pi}$ neutral current events are cut away. While this is an “internal structure” effect—there is a correlation between $M_{inv}^{3\pi}$, $p_T^{3\pi}$ and p_T^{miss} values for three-pion decays which is absent for 3π combinations in neutral current events—we can gain some understanding of the phenomenon by treating the 3π invariant mass as fixed, say

$$\begin{aligned} M_{inv}^{3\pi} &= m_a, \text{ constant,} \\ &= 1.2 \text{ GeV}/c^2. \end{aligned}$$

In this case M_R becomes a function of $\vec{p}_T^{3\pi}$ and \vec{p}_T^{had} only, and thus a kinematic variable describing the transverse plane. As was done for the transverse plane variables s_T and M_\perp in section 3.4.4, we may express a cut in M_R as a surface dividing the $(\rho^{3\pi}, \rho^{had}, M)$ space—describing all configurations of the transverse momenta—into two regions, one signal- and one background-dominated. For a given value of the transverse mass scale $M = p_T^{3\pi} + p_T^{had} + p_T^{miss}$, this surface describes a contour on the $(\rho^{3\pi}, \rho^{had})$ plane: the expression for the contours is derived in appendix B.5, while the contours for a cut $M_R < 2.2 \text{ GeV}/c$ are shown in figure 4.13 for the large- M region. The contours for

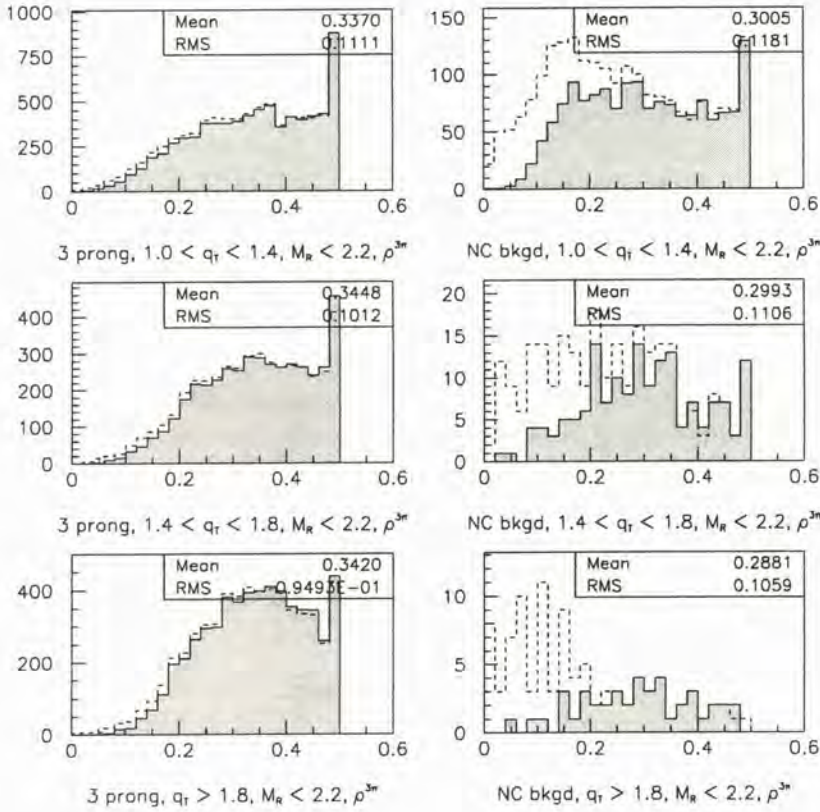


Figure 4.12: The $\rho^{3\pi}$ variable for signal and background in very broad bins of q_T , after a cut $M_R < 2.2$ GeV/c (shaded area), and a cut $M_\perp < 1.7$ GeV/c (dashed curve). The statistics shown apply to the M_R -cut sample (compare with figure 3.26).

an $M_\perp < 1.7$ GeV/c cut are also shown for comparison (effectively these plots are a detail of figure 3.16, with the M_R contours added).

The M_R contour effectively encompasses the high-signal region of the plot in each bin of the p_T -scale M , which was the reason for introducing M_\perp in section 3.4.4 above, while excluding background otherwise accepted by the M_\perp cut. We had noted much earlier (section 3.4.3) a concentration of the background in the top-left corner ($\rho^{3\pi}, \rho^{had}$) \approx (0, 0.5), not present in the signal: for genuine $\tau \rightarrow \nu 3\pi$ events, the massive 3π system must have significant $p_T^{3\pi}$ —or else p_T^{miss} must be relatively small—otherwise the difference in size between E_T and p_T terms in

$$M_R = \sqrt{(p_T^{miss} + E_T^{3\pi})^2 - (\vec{p}_T^{miss} + \vec{p}_T^{3\pi})^2} \quad (4.11)$$

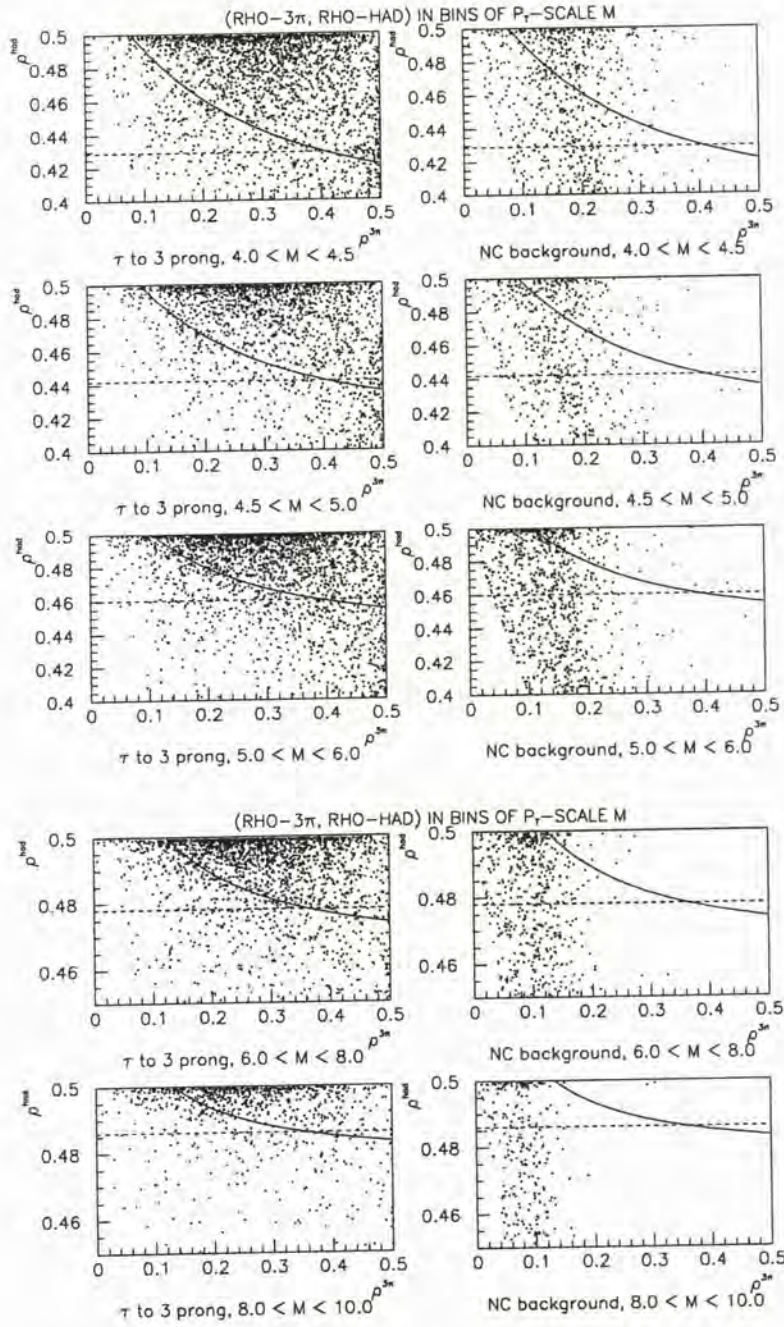


Figure 4.13: A comparison of M_R and M_\perp cuts on the transverse plane. Three-prong decay (left plots) and neutral current events (right plots) in the Monte Carlo are plotted on the $(\rho^{3\pi}, \rho^{had})$ plane in bins of the p_T -scale M for the higher- M region. The effect of a cut $M_R < 2.2 \text{ GeV}/c$ (assuming a fixed $M_{inv}^{3\pi} = 1.2 \text{ GeV}/c^2$; see the text) is shown with the solid contour; that of a cut $M_\perp < 1.7 \text{ GeV}/c$ by the dashed line. Compare with figure 3.16: the vertical scale is expanded by five for the upper three plots, by ten for the lower two plots; the s_T cut is not shown.

takes M_R to large values.

The required mass M_R is thus a more powerful discriminator between three-prong decay and neutral current events than the transverse mass M_{\perp} .

4.3.4 A kinematic effect of the three-pion mass

The physical basis of this extra restriction to which M_R is sensitive may be understood by considering the tau decays shown in figure 4.14, where high- and low- $p_T^{3\pi}$ decays of a τ with a given momentum are shown. In order to achieve the low- $p_T^{3\pi}$ configuration shown in figure 4.14(a), the neutrino must be “thrown backwards” in the rest frame of the tau: naïvely one would suppose that if the 3π system were thrown backwards with the same momentum, we would achieve the low- $p_T^{3\pi}$ configuration shown in figure 4.14(b). However, writing (E^*, \vec{p}^*) for the 3π four-momentum in the tau frame, (E, \vec{p}) for its four-momentum in the laboratory frame, and β, γ for the tau velocity and Lorentz factor respectively, we note that the Lorentz transformation between the frames is

$$\begin{aligned} p_{\parallel} &= \gamma(p_{\parallel}^* + \beta E^*) \\ &= \gamma(-q + \beta E^*) \end{aligned} \tag{4.12}$$

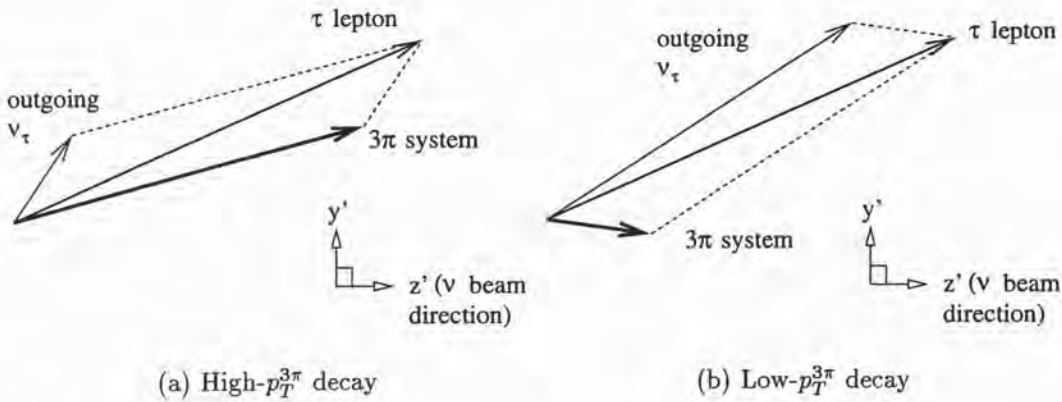


Figure 4.14: Tau decays with (a) high $p_T^{3\pi}$ and (b) low $p_T^{3\pi}$. The τ decay from figure 3.7(a) has been used in (a); for (b) the τ momentum has been retained, but the ν_{τ} and 3π momenta have been exchanged and then rotated by 180° about the τ momentum. Up to polarization effects this configuration would seem equally likely to (a), but in fact it is suppressed due to the mass of the 3π system, as discussed in the text.

where we have written $p_{\parallel}^* = -q$, ($q > 0$) for the momentum component of the 3π system along the τ direction. For a given 3π momentum \vec{p}^* in the tau frame, E^* and hence p_{\parallel} increases with the 3π mass,

A (massive) 3π system, thrown backwards by \vec{p}^* in the tau frame, is therefore more energetic in the lab than a (negligible mass) ν_{τ} would be if it were thrown backwards with an identical momentum \vec{p}^* : the straightforward exchange of laboratory momenta which we assumed for the sketch in figure 4.14(b) is in fact impossible. The required mass M_R , taking the 3π mass into account in the proper way, is sensitive to this physical effect; the transverse mass M_{\perp} , treating the 3π system as massless, is not.

4.3.5 Summary: M_R as a transverse mass variable

We introduced the transverse mass M_{\perp} in section 3.4.4 as a means of summarising the division between three-prong decay and neutral current events in the space $(\rho^{3\pi}, \rho^{had}, M)$, *i.e.* the set of configurations of the transverse momenta $\vec{p}_T^{3\pi}$ and \vec{p}_T^{had} . While M_{\perp} is well-justified on the pragmatic ground that it discriminates well between three-prong decay events and the background, and is moreover compatible with the use of the isolation variable q_T , we also noted that M_{\perp} has a clear physical interpretation as a measure of the mass of the three-pion/missing-momentum system.

Proceeding instead from this physical interpretation, we find ourselves led to the required mass M_R , which reduces to M_{\perp} in the limit where the 3π system is massless, but otherwise takes the 3π mass into account. This quantity, which occupies a middle ground between the purely “kinematic” variables of the previous chapter (describing the overall momenta of the 3π and hadronic systems) and the “structure” variables of the present chapter (sensitive to the $a_1^- \rightarrow \pi^- \pi^+ \pi^-$ decay) provides a natural way of handling a striking kinematic effect of the 3π structure: a bias towards cases where the 3π system dominates in the transverse plane over the missing transverse momentum, due to the 3π mass.

4.4 Opening angles of the 3π system

In the laboratory frame one of the striking differences between true three-prong tau decay and neutral current events is that the tracks of the pions of the 3π system are more collimated in space. A natural way to quantify this is to use an *opening angle* between the pions, and distributions of such angles for tau decay and neutral current background are presented in section 4.4.1, showing the clear differences between the two types of event. Critical examination however reveals that the differences in opening angle are due to a kinematic effect (section 4.4.2) and a 3π structure effect (section 4.4.3) which have already been taken into account in the analysis. While there are some residual differences between the two cases, one of which is considered in section 4.4.4, the apparently important angle variables prove to be of only marginal use in improving tau-decay/neutral-current discrimination.

4.4.1 The opening angles and their distributions

A number of definitions of opening angle for a three-pion system are possible: two of the more obvious are illustrated in figure 4.15. We can choose the largest angle in space between any two of the pions, which we will call $\theta_{\pi\pi}$, as a measure of the angular size of the system in the laboratory frame; alternatively, we can choose the largest of the angles between a pion momentum \vec{p}_i and the total momentum of the 3π system $\vec{p}^{3\pi}$, which we will call θ_{\max} , as giving the largest angular deviation between the direction of the hypothetical decaying particle a_1^- and any of its decay products.

Distributions of both variables for three-prong decay and neutral current events in the Monte Carlo are shown in figures 4.16 ($\theta_{\pi\pi}$) and 4.17 (θ_{\max}).




Figure 4.15: An example of a 3π system, with the opening angles $\theta_{\pi\pi}$ and θ_{\max} shown; see the text.

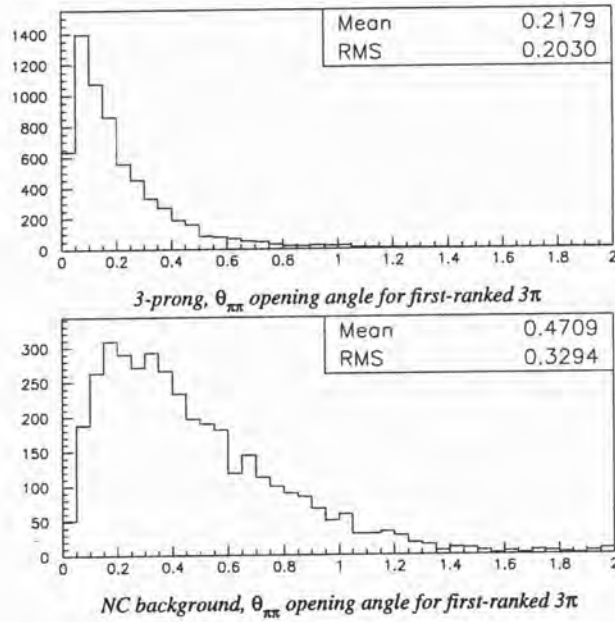


Figure 4.16: The opening angle $\theta_{\pi\pi}$ of the first-ranked 3π candidate in three-prong tau decay (upper plot) and neutral current events (lower plot).

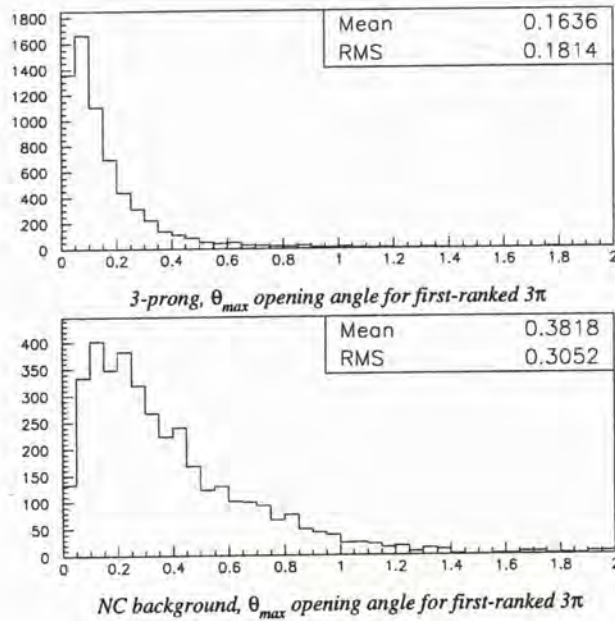


Figure 4.17: The opening angle θ_{\max} of the first-ranked 3π candidate in three-prong tau decay (upper plot) and neutral current events (lower plot).

In both cases there are very clear differences between the two samples. The opening angles for three-prong decays are less than half as large in the mean, and more concave to the origin than the neutral currents, and if a small tail to large values is neglected, the three-prong distributions have a smaller effective range: cuts at (say) $\theta_{\pi\pi} < 0.5$ or $\theta_{\max} < 0.4$ would reject half of the background while keeping most of the signal, and we would anticipate a greater discrimination power than this if a likelihood approach were used.

4.4.2 The $1/\gamma^{3\pi}$ scale of the opening angle

Thus it appears that the 3π opening angle offers strong discrimination between 3π systems produced in τ decay and those selected from neutral current events. At least part of the difference in opening angle distribution, however, is due to kinematic effects which have already been described. In particular, it is easily shown that for individual pions produced in a 3π decay, the opening angle of the pion with respect to the 3π system, θ , is a simple function of the 3π Lorentz factor:

Suppose we have a particle A , of mass M , travelling in some direction in the laboratory frame with speed β and Lorentz factor $\gamma = 1/\sqrt{1-\beta^2}$; it then undergoes a decay $A \rightarrow B + C (+D + \dots)$, and we consider one of the decay products B of mass m .

In the rest frame of the decaying particle A , the particle B has momentum \vec{p}^* , with components parallel to the motion of A ³ and perpendicular to that motion, p_L^* and p_T^* respectively; there is no reference direction perpendicular to the motion of A , so we will ignore the azimuth of \vec{p}_T^* throughout. If we transform to the laboratory frame, B has momentum \vec{p} with components

$$\begin{aligned} p_T &= p_T^* \\ p_L &= \gamma(p_L^* + \beta E^*), \end{aligned} \quad (4.13)$$

E^* being the energy of B in the rest frame of A . The situation is shown in figure 4.18 where $A = 3\pi$ and B is one of the decay pions. When the "opening angle"—the polar angle θ of B in the laboratory, with respect to the direction of motion of A —is small, we may write

$$\begin{aligned} \theta &\simeq \tan \theta \\ &= \frac{p_T}{p_L} \\ &= \frac{p_T^*}{\gamma(p_L^* + \beta E^*)}. \end{aligned} \quad (4.14)$$

³That is, opposite to the motion of the laboratory, as seen from the frame of A .




Figure 4.18: The opening angle θ for a pion belonging to the 3π system. The momentum of the pion in the laboratory frame \vec{p} , and the direction of $\vec{p}^{3\pi}$ in the lab are shown: θ is the angle in space between these vectors. The components of \vec{p} are also shown, in terms of the pion momentum in the 3π -rest frame, \vec{p}^* (see the text).

Since the variation of $\beta = \sqrt{1 - 1/\gamma^2}$ with γ is small for $\gamma \gg 1$, we see immediately that the distribution of opening angle θ has a natural scale of $1/\gamma$ when the decaying particle is energetic: the other quantities in equation 4.14 depend only on the dynamics of the decay in the rest frame, and will not vary with the boost γ of A . That is, for $\gamma \gg 1$,

$$\theta = \frac{1}{\gamma} \times \{\text{terms independent of } \gamma\}. \quad (4.15)$$

As the opening angles of the individual pions scale with $\gamma^{3\pi}$ in this way, so will the largest of these angles θ_{\max} , and, to a good approximation, the largest angle between any two of the pions $\theta_{\pi\pi}$.

4.4.3 The scale-free part of the opening angle

For true $\tau \rightarrow \nu 3\pi$ decays, then, the distribution of any of the opening angles will be the product of a scale-free distribution dictated by the dynamics of the $\tau \rightarrow \nu 3\pi$ decay, and a factor

$$\frac{1}{\gamma^{3\pi}} = \frac{M_{\text{inv}}^{3\pi}}{E^{3\pi}}$$

so that the mean of the opening angle distribution will be inversely proportional to the 3π energy.

For 3π systems selected from neutral current events, the situation is less straightforward: for example, there is in general no true decaying system. Indeed, if we multiply the largest of the individual opening angles, θ_{\max} , by

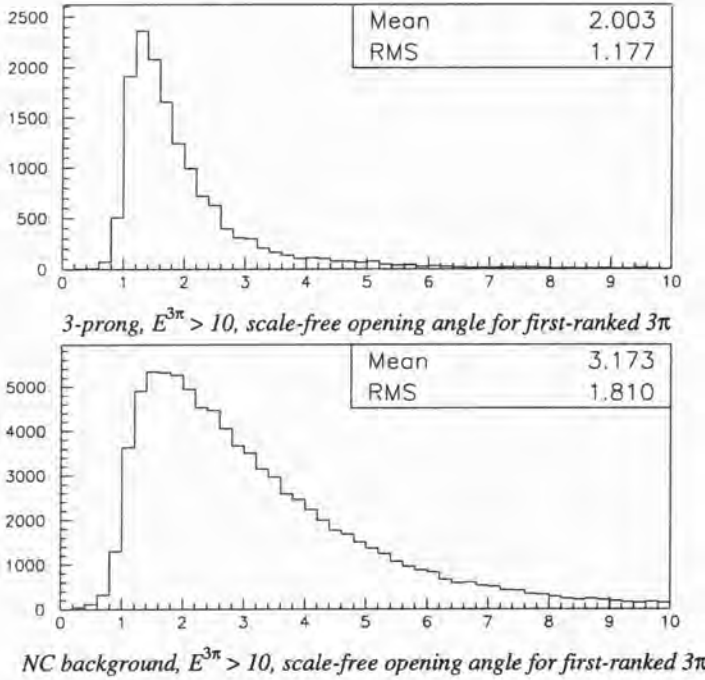


Figure 4.19: The scale-free opening angle $\gamma^{3\pi}\theta_{\max}$ for the first-ranked 3π candidate in three-prong tau decay (upper plot) and neutral current events (lower plot). A weak cut $E^{3\pi} > 10$ GeV in the 3π energy has been imposed.

the 3π boost $\gamma^{3\pi}$ in order to recover a scale-free quantity $\gamma^{3\pi}\theta_{\max}$, significant differences between the tau decay and neutral current events are still apparent, as shown in figure 4.19.

Given the invariant mass $M_{inv}^{3\pi}$ of such a 3π system, however, we may transform to the rest-frame of the three pions, where the momentum vectors of the pions will sum to zero, as for a true $a_1^- \rightarrow \pi^- \pi^+ \pi^-$ decay. This simple manoeuvre turns out to be sufficient to derive strong conditions on the opening angle, regardless of the origin of the three pions. We find (appendix C) that requiring an m_{a_1} -like mass for $M_{inv}^{3\pi}$ imposes the following form on the distribution of $\gamma^{3\pi}\theta_{\max}$:

- an approximate lower bound, $\gamma^{3\pi}\theta_{\max} \gtrsim 0.8$, which holds apart from a small tail;
- an upper bound, $\gamma^{3\pi}\theta_{\max} \lesssim 5$;
- a falling distribution up to the upper bound.

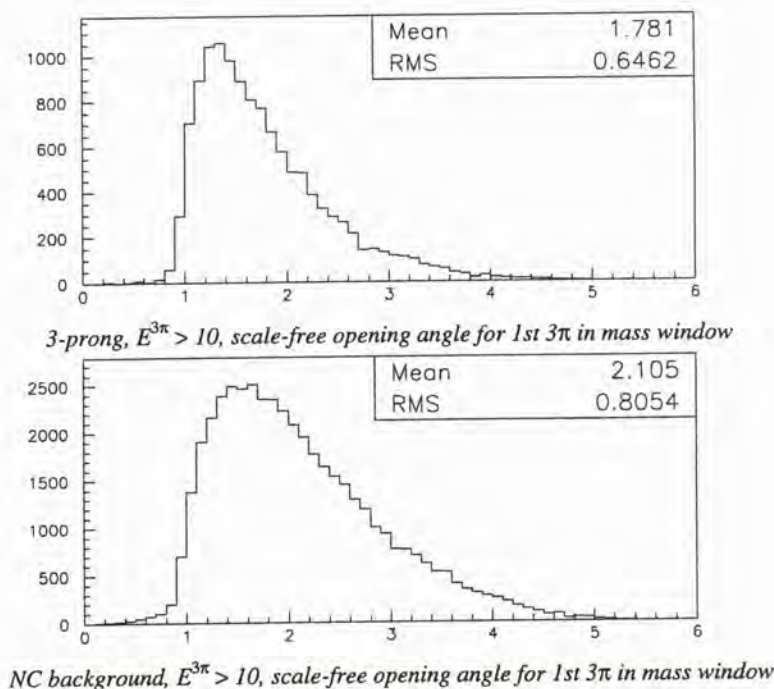


Figure 4.20: The scale-free opening angle $\gamma^{3\pi}\theta_{\max}$ when a 3π mass $0.80 < M_{inv}^{3\pi} < 1.50 \text{ GeV}/c^2$ is required prior to candidate selection: most of the differences between the three-prong tau decay (upper plot) and neutral current event distributions (lower plot) are thus removed (*cf.* figure 4.19 above). A weak cut $E^{3\pi} > 10 \text{ GeV}$ in the 3π energy has been imposed.

That is, the distribution is largely determined. If we once again plot the distribution of the scale-free angle $\gamma^{3\pi}\theta_{\max}$ with an energy cut in place, but include a cut in the invariant mass $M_{inv}^{3\pi}$ in the 3π selection, all of the features listed are apparent, and most of the differences between the three-prong and neutral current samples in the previous plots have disappeared (figure 4.20). There is a residual difference in that the neutral current distribution is more broad than the three-prong distribution, while respecting the same bounds. Even here, however, we are to some extent considering an effect which has already been noted elsewhere: the distribution of 3π invariant mass within the window $0.80 < M_{inv}^{3\pi} < 1.50 \text{ GeV}/c^2$ is biased to large values for the neutral current sample (*cf.* figure 4.6), and the probability of cases $2 \lesssim \gamma^{3\pi}\theta_{\max} \lesssim 4$ rises with the maximum pion momentum \vec{p}^* in the 3π rest frame, which rises with the mass $M_{inv}^{3\pi}$ (appendix C.2, especially table C.1).

It is worth noting how weak a condition we have imposed in order to

produce the constraints listed. The $a_1 \rightarrow \rho\pi$ substructure of the $a_1^- \rightarrow \pi^-\pi^+\pi^-$ decay is not required, nor is any special momentum distribution for the three pions in the 3π rest-frame. All we have required is an a_1 -like range of invariant masses for the 3π system: for any $\tau \rightarrow \nu 3\pi$ analysis this is an indispensable 3π -structure cut, and so the bounds exhibited above should apply regardless of the scheme used to select the three pions, and of other cuts.

4.4.4 The isolation angle

There is however an additional effect on the angles between tracks in the event, due not to the 3π structure but to the structure of the *hadronic system* in neutral current events $\nu_\mu + N \rightarrow \nu_\mu + X$. In such events all of the measurable particles arise in the fragmentation of the system X which is produced in the interaction of the target nucleon and the exchanged Z^0 boson. Apart from first order QCD processes leading to hard gluon emission *etc.*, which we expect to be a small effect at NOMAD experiment energies, there is no reason for such systems to produce “gaps”—significant angular regions between subsystems in the final state where no particles are present.

For tau decay events, on the other hand, the 3π and hadronic systems are produced from quite different objects, so such gaps are foreseen. A line of reasoning of this kind suggested the variable q_T to quantify the isolation of the 3π system as a whole (section 3.3.3) and we now seek a way of quantifying the isolation of the 3π tracks from individual tracks in the hadronic system. One suitable quantity is the isolation angle θ_{isol} defined in figure 4.21. Significantly positive values of θ_{isol} indicate a gap in space between the tracks of the 3π and hadronic systems, while negative values indicate an overlap between the systems.

Considering 3π combinations from events without other cuts, the differences between three-prong decay and neutral current events are in fact minor, as shown in figure 4.22. After kinematic cuts are imposed, the difference between the two classes of event becomes more marked—with three-prong decays having higher values of θ_{isol} than neutral current events—but recalling the close relationship between the opening angle and the energy of the 3π system (section 4.4.2) we should be careful to consider events of similar $E^{3\pi}$. If we compare the distribution of events under cuts, in narrow bins of $E^{3\pi}$, the difference between tau decay and neutral current events remains modest (an example is shown in figure 4.23.)

We may understand this result by considering two effects which work against the observability of an angular “gap” in tau decay events:

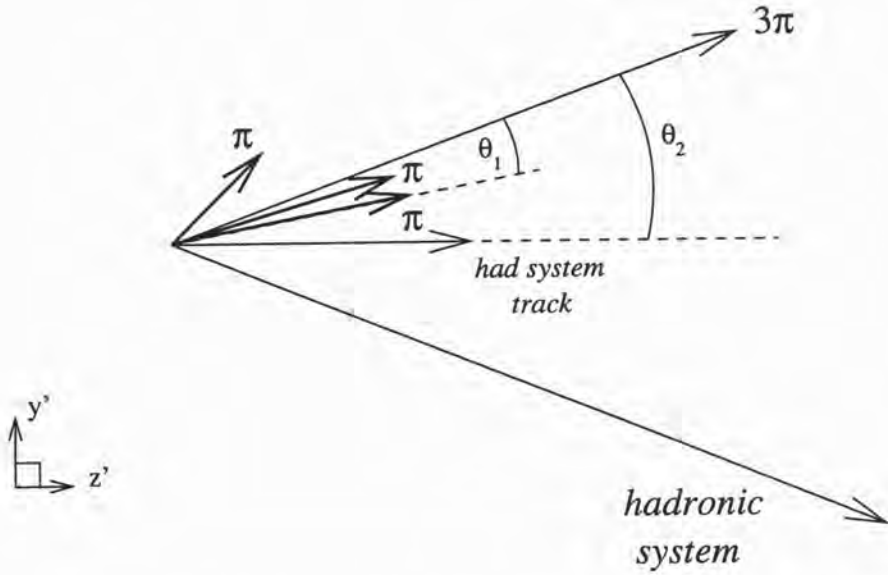
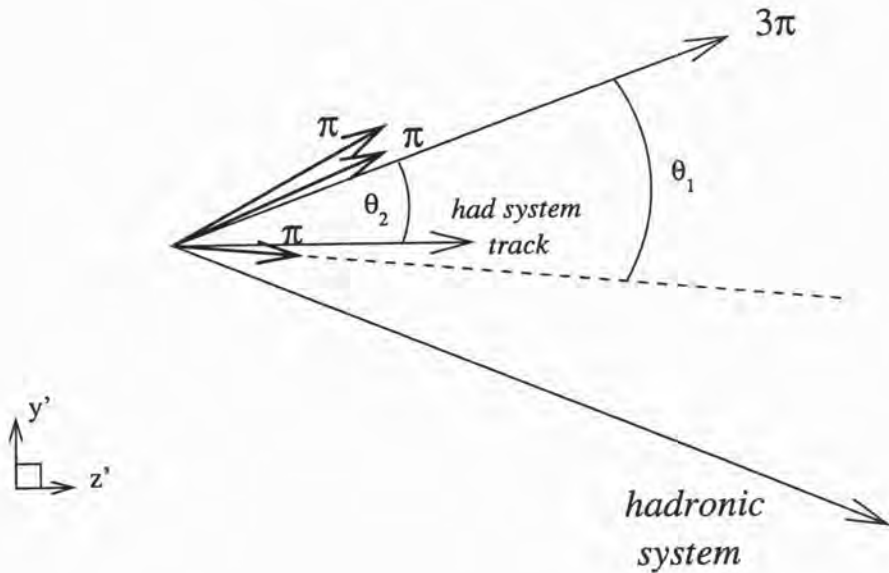
(a) Well-separated 3π and hadronic systems.(b) Overlapping 3π and hadronic systems.

Figure 4.21: Definition of the isolation angle θ_{isol} : the difference of the angle between the 3π momentum and the closest track from the hadronic system (θ_2), and the angle between the 3π momentum and the $\pi \in 3\pi$ track closest to the hadronic system (θ_1). For (a) well-separated 3π and hadronic systems, $\theta_{\text{isol}} = \theta_2 - \theta_1$ is positive; where the systems overlap, as in (b), θ_{isol} is negative.

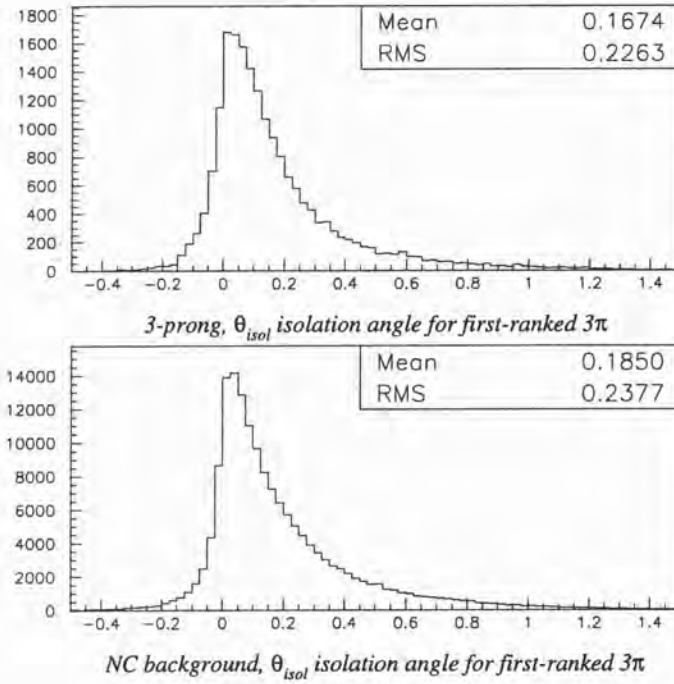


Figure 4.22: Distributions of the isolation angle θ_{isol} for the first-ranked 3π candidate in three-prong tau decay (upper plot) and neutral current events (lower plot).

- in both tau decay and neutral current interactions, the whole event is boosted forward, which “bunches” most of the tracks together into a relatively narrow angular region;
- the mass of the tau is a significant fraction of the CMS energy of a $\nu_\tau + N \rightarrow \tau^- + X$ event as a whole, so that the angular spread of tracks in three-prong decay is comparable to the separation of the 3π and hadronic systems—hence the occurrence of negative values of θ_{isol} for three-prong decays, even under cuts in the 3π isolation q_T (figure 4.23).

In any case, due to the limited discrimination power which seems to be available, we will omit the angle θ_{isol} from the analysis of the event.




Figure 4.23: Distributions of the isolation angle θ_{isol} when a 3π mass $0.80 < M_{inv}^{3\pi} < 1.50 \text{ GeV}/c^2$ is required prior to candidate selection, and kinematic cuts $q_T > 1.0 \text{ GeV}/c$ and $M_{\perp} < 1.7 \text{ GeV}/c$ are imposed. In this example events have been chosen in an energy bin $10 < E^{3\pi} < 15 \text{ GeV}$.

4.5 Sensitivity to the τ polarisation?

We now consider the sensitivity of the NOMAD experiment to the polarisation state of the τ lepton, which we have ignored in the discussion of tau structure variables up to this point. The polarisation of taus in leptoproduction is discussed in section 4.5.1, while the quantities used to measure tau polarisation in electron-positron colliders are introduced in section 4.5.2; those quantities which are measurable in the rather different experimental situation at the NOMAD are then constructed for the Monte Carlo data in section 4.5.3. The consequences for the present analysis, and potential analyses at the NOMAD and “next generation” detectors, are discussed in section 4.5.4.

4.5.1 Polarisation of the τ lepton in $\nu_\tau + N \rightarrow \tau^- + X$

Polarisation-sensitive variables are attractive as τ -decay/neutral-current discriminators in the NOMAD experiment because in leptoproduction, $\nu_\tau + N \rightarrow \tau^- + X$, the tau lepton is fully polarised in its rest frame. This striking effect, and the direction of polarisation, may both be deduced from a simple argument.

Tau polarisation in the quark-lepton frame

We consider the quarks and leptons taking part in the fundamental reaction $d + \nu_\tau \rightarrow u + \tau^-$, in their centre-of-momentum frame as depicted in figure 4.24, using the direction of motion of the ν_τ as the reference (z) direction for spin; we write \hat{s} for the squared-total-energy of the quark-lepton system ($\hat{s} = xs$, where x is the Bjorken scaling variable and s the CMS energy-squared of the full nucleon-lepton system). The spin of the left-handed ν_τ , with $m_{\nu_\tau} \approx 0$, will be aligned opposite to its motion, so $J_z^{\nu_\tau} = -\frac{1}{2}$; the incoming d quark, with an energy in this frame much larger than its mass ($\frac{1}{2}\sqrt{\hat{s}} \gg m_d$) may also be treated in the ultrarelativistic approximation where positive (negative) chirality implies positive (negative) helicity. The weak interaction couples only the left-handed states of the particles and so negative chirality and helicity are chosen: since we have chosen $p_z^d = -|\vec{p}^d|$, $J_z^d = +\frac{1}{2}$. The total spin of the initial state is therefore zero.

In the final state the outgoing u quark is polarised by the same argument: setting the z' axis such that $p_z^u = -|\vec{p}^u|$, we have $J_z^u = +\frac{1}{2}$. But conservation of the total (zero) angular momentum requires the total $J_{z'} = 0$, so the tau is constrained to have $J_{z'}^\tau = -\frac{1}{2}$, *i.e.* to be totally left-hand polarised. Any deviation from this behaviour is due to the (small) masses of the quarks




Figure 4.24: Quark and lepton spin directions in the quark-lepton centre-of-momentum system.

(terms of $O(2m_{u,d}/\sqrt{\hat{s}}) \ll 1$), rather than the large tau mass. Furthermore, since the transformation from the quark-lepton centre-of-momentum frame to the tau rest frame is a boost along the direction of motion of the tau, *i.e.* parallel to the spin axis of the tau in the quark-lepton frame, the spin axis is unaffected by this transformation: the spin vector of the tau in its rest frame, \vec{s}_τ , is thus given by the negative of the tau direction of motion *in the quark-lepton frame*:

$$\vec{s}_\tau = -\frac{\vec{p}_{q\ell}^\tau}{|\vec{p}_{q\ell}^\tau|} \quad (4.16)$$

A caveat

In fact the preceding argument is incomplete, as it assumes that the quark-neutrino system has overall angular momentum $J = 0$: the most general combination of two $J = \frac{1}{2}$ systems to give $J_z = 0$ as shown would be a superposition of $J = 0$ and $J = 1$ states. A $J = 1$ state would clearly allow deviation from the behaviour described since $J_{z'} = 0$ would no longer be required. Noting however that neutrino-quark scattering is isotropic we may deduce that only the $J = 0$ component of the initial state contributes to the scattered amplitude, in which case the argument presented above would still

hold.

A rigorous treatment of the interaction $\nu_\tau + N \rightarrow \tau^- + X$ in fact shows that the $J = 0$ component alone participates in the scattering, and that the τ^- is indeed fully polarised in its rest frame, left-handed with respect to its direction of motion in the quark-lepton frame (Iacopini, 1997).

4.5.2 Measurement of τ polarisation at e^-e^+ colliders

We now turn to the appropriate technique for measuring this polarisation. The literature, in the main, assumes that measurements are taking place at an electron-positron collider—especially the Large Electron Positron Collider (LEP) at CERN, where taus are produced in the reaction $e^- + e^+ \rightarrow Z^0 \rightarrow \tau^- + \tau^+$. The differences in the coupling of the Z^0 to positive- and negative-helicity fermions introduces a small polarisation of the tau leptons, which may be accurately reconstructed from the distribution of the tau decays, in an appropriate set of variables, over many events. We will follow the treatment of Kühn and Mirkes (1992a,b) where the $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ decay is described by⁴

- the three-pion mass $M_{inv}^{3\pi}$;
- the two-pion masses $\sqrt{s_1}$ and $\sqrt{s_2}$ discussed in section 4.2;
- the polar angle θ_a of the 3π system in the τ rest frame, taking the direction of motion of the τ in the laboratory as the reference, which may be recovered using the expression

$$\cos \theta_a = \frac{(2xm_\tau^2 - m_\tau^2 - (M_{inv}^{3\pi})^2)}{(m_\tau^2 - (M_{inv}^{3\pi})^2)\sqrt{1 - m_\tau^2/(E^\tau)^2}} \quad (4.17)$$

where E^τ is the energy of the τ in the lab and $x = E^{3\pi}/E^\tau$;

- the polar angle β_a of the normal to the 3π decay plane in the 3π rest frame, taking the direction of motion of the laboratory ($\vec{n}_{lab} = -\vec{p}^{3\pi}/|\vec{p}^{3\pi}|$) as the reference; and
- the azimuthal angle γ_a of the odd-sign pion (the π^+ , in our case) in the 3π decay plane, relative to the plane defined by the normal to the 3π plane \vec{n}_\perp , and the direction of motion of the laboratory \vec{n}_{lab} .

⁴Among the wide range of choices in the literature, we use the notation of Kühn and Mirkes, with an “a” subscript added to the angles (referring to the a_1 or 3π system) to distinguish them from the quantities with the same name used elsewhere in this work.

A final angle α_a , giving the azimuth of the normal to the 3π decay plane, cannot be measured unless the direction of motion of the τ is known (see the discussion and figures 1 and 2 of Kühn and Mirkes, 1992b). This is in contrast to the polar angle θ_a , which is defined using the τ direction, but may be obtained from (4.17) given the tau energy E^τ , available in a collider as $\frac{1}{2}\sqrt{s}$.

The tau polarisation is then recovered by fitting a linear combination of positive and negative helicities to moments of distributions of the quantities listed above: $\cos \theta_a$; and then each of $\langle \frac{1}{2}(3 \cos^2 \beta_a - 1) \rangle$, $\langle \cos 2\gamma_a \rangle$, $\langle \sin 2\gamma_a \rangle$, and $\langle \cos \beta_a \rangle$ as functions of $\cos \theta_a$. This method is described in detail by Privitera (1993), and used by the DELPHI collaboration to obtain the tau polarisation presented by Abreu et al. (1995).⁵ Among these distributions, $\cos^2 \beta_a$ and $\cos 2\gamma_a$ carry the bulk of the information.

4.5.3 Polarisation measurement in the NOMAD experiment

Adaption of the method of moments

In the NOMAD experiment neither the tau direction nor its energy are reconstructed, and so the angle θ_a is unavailable. The angles β_a and γ_a may, however, be recovered from the laboratory momenta of the three pions, using

$$\cos \beta_a = \frac{\vec{p}_3 \cdot (\vec{p}_1 \times \vec{p}_2)}{|\vec{p}^{3\pi}| T}, \quad (4.19)$$

$$\cos \gamma_a = \frac{A_3}{|\vec{p}^{3\pi}| \sqrt{B_3} \sin \beta_a}, \quad (4.20)$$

$$\sin \gamma_a = -\frac{\cos \gamma_a}{T} \left(B_3 \frac{A_1}{A_3} - \frac{B_2 - B_1 - B_3}{2} \right), \quad (4.21)$$

⁵The last two of these moments are sensitive to the ordering of the negative pions and the full expressions read $\langle S_{12} \sin 2\gamma_a \rangle$, and $\langle S_{12} \cos \beta_a \rangle$, where

$$S_{12} \stackrel{\text{def}}{=} \frac{s_1 - s_2}{|s_1 - s_2|}; \quad (4.18)$$

with the convention chosen at (4.3), $S_{12} = 1$ always so that the ordering is automatically satisfied.

An alternative procedure for measuring polarisation, based on the same analysis of the underlying physics, constructs a single polarisation-sensitive variable ω from the quantities defined here; this method is described and advocated by Davier et al. (1993) and used by the ALEPH and L3 collaborations (Buskulic et al., 1993; Acciarri et al., 1994). Since it is less transparent this method is less suitable for the adaption being carried out in this section.

where

$$A_i = \frac{E^{3\pi} \vec{p}^{3\pi} \cdot \vec{p}_i - E^i (\vec{p}^{3\pi})^2}{M_{inv}^{3\pi}}, \quad (4.22)$$

$$B_i = \frac{(E^i E^{3\pi} - \vec{p}_i \cdot \vec{p}^{3\pi})^2 - (M_{inv}^{3\pi})^2 m_\pi^2}{(M_{inv}^{3\pi})^2}, \quad (4.23)$$

$$T = \frac{1}{2} \sqrt{-\lambda(B_1, B_2, B_3)}, \quad (4.24)$$

$$\lambda(X, Y, Z) = X^2 + Y^2 + Z^2 - 2XY - 2XZ - 2YZ \quad (4.25)$$

following Privitera (1993).

Since we cannot reconstruct θ_a we are unable to use the method of the previous section, which relies in each case on distributions with respect to $\cos \theta_a$. We may however take advantage of the kinematic cuts defined in section 3.4, specifically the cut on the 3π momentum $p^{3\pi}$: the selection of energetic 3π systems introduces a bias in the efficiency of the cuts towards




Figure 4.25: The polar angle of the tau decay θ_a for $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ decays in the Monte Carlo (upper plot), and after a strong energy cut on the selected 3π combination $E^{3\pi} > 40$ GeV (lower plot).

$\tau \rightarrow \nu 3\pi$ decays in which the 3π system is forward in the tau frame. So we see in figure 4.25 that while $\cos \theta_a$ has an essentially uniform distribution over all $\tau^- \rightarrow \nu \pi^- \pi^+ \pi^-$ decays, a strong cut in the decay product energy $E^{3\pi}$ yields a distribution with a marked bias to large values of $\cos \theta_a$.

Much of the power of the moments analysis should then be recovered as

1. the distribution of $\cos \beta_a$ changes only slowly with $\cos \theta_a$, and clearly distinguishes between tau polarisations $P_\tau = -1$ and $P_\tau = +1$ cases, for decays in the range, say, $\cos \theta_a > -0.2$ (Rougé, 1990, figure 1); and
2. the distributions of $\langle \cos 2\gamma_a \rangle$ and $\langle \sin 2\gamma_a \rangle$ show the same evolution with $\cos \theta_a$ as does $\langle \frac{1}{2}(3 \cos^2 \beta_a - 1) \rangle$ (Privitera, 1993, equations 28 through 30)

Plots of $\cos^2 \beta_a$, $\cos 2\gamma_a$ and $\sin 2\gamma_a$ may therefore be expected to discriminate between positive- and negative-helicity τ decays. In the method of moments these three quantities are responsible for 93% of the sensitivity of the measurement (Privitera, 1993).

Potential problems with this adaption

This adaption of the moments analysis may not in fact be appropriate for our purposes, for two reasons:

1. We require discrimination between (the 100% polarised, $P_\tau = -1$) $\nu_\tau + N \rightarrow \tau^- + X$ interaction, and *neutral current* events containing isolated 3π combinations—and it is unclear what the apparent “polarisation” of these combinations of (unrelated) pions will have.
2. The treatment of Kühn and Mirkes (1992b) assumes that the spin vector of the tau \vec{s}_τ lies in the plane defined by \vec{n}_{lab} and \vec{n}_τ , the direction of motion of the tau as seen from the 3π rest frame. This follows from the assumption that the tau is polarised along its direction of motion in the laboratory, an assumption which holds at an e^-e^+ collider—while in leptoproduction, as noted in section 4.5.1, the tau is polarised along its direction of motion in the center of momentum system of the interacting quark and neutrino. In general this is by no means the same as the direction of motion in the laboratory.

The derivation of the $(\cos \theta_a, \cos \beta_a)$ distribution by Rougé (1990) depends explicitly on this assumption.⁶ In the more complete, and complex treatment of Kühn and Mirkes the role played by this assumption

⁶Rougé calls these variables $(\cos \theta, \cos \psi)$.

is unclear. A more general treatment, not relying on the alignment of the tau polarisation and laboratory line-of-flight axes, does not seem to be available in the literature.

We may gain some feel for the magnitude of this second problem by constructing the angle δ between the direction of motion of the tau lepton in the quark-lepton CMS, and in the laboratory. While there is a bias towards cases where $\cos \delta > 0$, there is no strong relationship between the two directions for three-pion decays in general; a strong cut in the decay product energy $E^{3\pi}$, however, introduces a very noticeable bias towards $\cos \delta > 0.4$ (figure 4.26). So while the laboratory flight direction and the polarisation axis are not the same, there is for energetic 3π a relationship between the two.

For the purpose of this analysis we take the empirical approach of constructing the distributions of $\cos^2 \beta_a$, $\cos 2\gamma_a$ and $\sin 2\gamma_a$ for $\tau \rightarrow \nu 3\pi$ and neutral current events, after the appropriate kinematic cuts, and testing these variables for discrimination power. We rely at this point on the accuracy of




Figure 4.26: The cosine of the angle δ between the tau flight direction in the quark-lepton CMS and the laboratory frame, for $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ events (upper plot), and after a strong energy cut on the selected 3π combination $E^{3\pi} > 40$ GeV (lower plot).

event generation in the NOMAD Monte Carlo, which has been shown to precisely reproduce the results of the argument of section 4.5.1, and the rigorous derivation of Iacopini (1997), in the $\tau^- \rightarrow \nu_\tau \pi^-$ case (Cousins, 1997b).

Distributions in the NOMAD Monte Carlo

Using a strong energy cut $E^{3\pi} > 40 \text{ GeV}/c$, distributions of $\cos^2 \beta_a$, $\cos^2 \beta_a$ and $\cos 2\gamma_a$ are shown for three-prong decays and neutral current events in figures 4.27 through 4.29. In each case we see a similar effect: the distributions differ significantly in the mean, but cover the same range with no excluded regions; and the most noticeable differences are not the presence or absence of concentrations at low or high values, but the relative sizes of these concentrations. In other words, these variables are unsuitable for the application of cuts to separate the two samples, and even combining the three together by a weighting or likelihood method we would anticipate no significant discrimination power at the event-by-event level. Over a sample




Figure 4.27: The polarisation-sensitive variable $\cos^2 \beta_a$ for three-prong tau decay (upper plot) and neutral current events (lower plot) after a strong energy cut $E^{3\pi} > 40 \text{ GeV}$ has been imposed. The $0.80 < M_{inv}^{3\pi} < 1.50 \text{ GeV}/c^2$ candidate selection has been used.




Figure 4.28: The polarisation-sensitive variable $\cos 2\gamma_a$ for three-prong tau decay (upper plot) and neutral current events (lower plot) after a strong energy cut $E^{3\pi} > 40$ GeV has been imposed. The $0.80 < M_{inv}^{3\pi} < 1.50$ GeV/ c^2 candidate selection has been used.




Figure 4.29: The variable $\sin 2\gamma_a$ under the same cuts.

of many events, on the other hand, some discrimination would clearly be possible, but this is not the situation in the NOMAD experiment.

4.5.4 Prospects for the use of polarisation-sensitive variables

In the light of the distributions shown in the previous section, polarisation-sensitive variables are not included in the present analysis. There is however some potential for the use of these variables in an experimental design such as those originally proposed for “next-generation” neutrino oscillation searches in the CERN West Area (Gomez-Cadenas et al., 1996; Gomez-Cadenas and Hernando, 1996; Ereditato et al., 1996; Ayan, 1997), or acting as a near detector for long-baseline oscillation searches in the proposed CERN neutrino beam to the Gran Sasso laboratory (Acquistapace et al., 1998).⁷

In such a design there is either high-precision tracking near the interaction vertex by silicon detectors, or an active target of nuclear emulsion, or both: in the case of a three-prong tau decay this means that the decay point, and hence the direction of tau flight, are recovered. Together with the momentum of the decay product $\vec{p}^{3\pi}$, and the assumption of the tau mass m_τ for the decaying object, this supplies a sufficient number of constraints to solve for the momentum of the tau lepton in the laboratory, \vec{p}^τ . Therefore in addition to the variables available at the NOMAD, it would be possible to recover not only the tau decay angle θ_a but also the azimuthal angle α_a which is not even available at an e^+e^- collider (section 4.5.2). This would be an unprecedented experimental situation, and we should recall in addition that unlike the collider environment, tau leptons produced in one of these experiments would be completely polarised (section 4.5.1).

There is however still the difficulty that the tau polarisation axis is its direction of flight in the quark-lepton CMS, and there is no obvious way to recover this frame from the experimental data. As already noted there does not appear to be a discussion of the recovery of tau polarisation information in such a situation: in view of the otherwise complete information which would be available in the type of experiment described, a sustained theoretical treatment of this problem is called for.

⁷We should note that it appears that a near detection station is no longer foreseen for this facility.

4.6 Selection of the 3π -decay candidate

We now return to the problem of choosing the three-prong decay candidate—that is, a $\pi^-\pi^+\pi^-$ combination—in any given event. Where necessary in this chapter we have chosen the 3π according to one of the two schemes

- (a) selection of the 3π combination with the largest value of q_{lep}/p^{had} ;
- (b) imposition of a mass precut $0.80 < M_{inv}^{3\pi} < 1.50 \text{ GeV}/c^2$, selecting the 3π combination with the largest value of q_{lep}/p^{had} among those which remain;

as discussed in section 4.2.1. The second of these schemes was used unacknowledged throughout chapter 3. As noted in section 4.2.1, the addition of the precut in the 3π mass improves the rejection of the neutral current background, but the efficiency of the three-prong decay selection is still rather poor: in the Monte Carlo, the selection (b) chooses the correct three tau decay tracks in only 61% of cases (figure 4.30(b)). There is some improvement




Figure 4.30: Comparison of the efficiency of (a) the simple 3π ranking scheme, and (b) the same scheme with a prior cut in the 3π mass, in selecting the true 3π combination from three-prong tau decays in the Monte Carlo.




Figure 4.31: Improvement in 3π selection efficiency when weak kinematic cuts, $q_T > 1.0 \text{ GeV}/c$ and $M_{\perp} < 1.7 \text{ GeV}/c$, are imposed. Compare with figure 4.30 above.

when kinematic cuts are imposed, so that where the true three-prong decay product is reconstructed, and has $q_T > 1.0 \text{ GeV}/c$ and $M_{\perp} < 1.7 \text{ GeV}/c$, the selection (b) chooses this combination in 73% of cases (figure 4.31(b)).

In interpreting this performance we should note that nothing in the analysis actually relies on the selection of the correct 3π combination.⁸ A number

⁸In vertexing studies it has been noted that the quality of the fit at the primary vertex can be poor in the case of tau decay Monte Carlo events, reflecting the origin of several tracks at the tau decay point, a millimetre or more from the true primary vertex in energetic events. In other words, the NOMAD experiment *does* appear to have some modest sensitivity to the tau decay length. Exploitation of this in the three-prong decay analysis would require (say) the removal of the decay candidate tracks from the primary vertex, together with a trial fit to vertex of their own, and a study of the χ^2 of the fit and various other properties such as the distribution of implied flight distances, leading to a figure of merit which could be used to discriminate between true tau decays and backgrounds (Geiser, 1997).

In such a case it would clearly be of great importance to identify the correct decay pions in an efficient way in the analysis. However no such study has been carried out to date, and none is presently foreseen.




(a) With cuts $q_T > 1.0 \text{ GeV}/c$, $M_{\perp} < 1.7 \text{ GeV}/c$.(b) With cuts $q_T > 1.4 \text{ GeV}/c$, $M_R < 2.2 \text{ GeV}/c$.(c) With $q_T > 1.4 \text{ GeV}/c$, $M_R < 2.2 \text{ GeV}/c$, $E^{3\pi} > 20 \text{ GeV}$.

Figure 4.32: The improvement in the efficiency of the simple 3π selection procedure described in the text, as kinematic cuts are tightened. If the pre-cut in the invariant mass is omitted, the corresponding efficiencies fall to 77%, 69% and 79% respectively.

of possible *indirect* effects may be identified:

1. “false” 3π combinations will not necessarily have the same distributions of 3π structure quantities as the true decay products;
2. since at least one track is not produced from the tau decay, false 3π combinations may not be as energetic as the true 3π , or may otherwise fail the kinematic cuts;
3. in the case of the selection (b), and many other selection schemes, when the true 3π combination is not selected, it may be that *no* combination passes the precut, so that the event in question is lost from the analysis.

The first of these effects we may suppose to be minor, as we have found throughout this chapter that 3π structure effects are dominated by the 3π invariant mass: any false 3π combination selected (at least by scheme (b)) will have $M_{inv}^{3\pi}$ in the same range as a true decay product; we are furthermore relying only on $M_{inv}^{3\pi}$, and M_R which is a function of $M_{inv}^{3\pi}$, among the structure-sensitive variables. Effects (2) and (3) have been anticipated in figure 4.31, where for *true* three-prong decay products passing cuts, we show for the two selections respectively the cases where no 3π candidate is selected, where a “false” candidate is selected but itself fails the cuts, where a false candidate is selected and *passes* the same cuts as the true 3π , and where the true 3π combination is identified.

For our purposes the last two categories are indistinguishable. We therefore find that under the weak cuts shown, 87% of the potential 3π combinations are effectively recovered; as cuts are made more stringent, the case where the true 3π is recovered becomes dominant, and the total efficiency improves slightly to 91% (figure 4.32). By comparison the selection (a) performs relatively poorly under cuts, with a significant fraction of the selected candidates failing cuts which are passed by the true 3π . It may be seen that the inclusion of 3π structure effects, at least in the form of the 3π mass, is necessary to gain an efficient selection which is robust under cuts; and furthermore that a very simple treatment of the 3π structure can lead to an effective 3π selection.

In the light of its high efficiency when kinematic cuts are imposed, we take the selection (b) to be sufficient for our present purposes, rather than seeking out some more elaborate selection scheme. We note that a full optimisation of 3π selection would also take into account the rate at which “tau-decay-like” neutral current events are selected: effectively, monitoring how the final signal-to-background ratio of the analysis changes as the 3π selection procedure is varied between a number of options. Such a study is beyond the scope of this present work.

4.7 Summary: Variables sensitive to the 3π structure

There are clear advantages to the analysis of taking the 3π structure into account. The use of the invariant mass $M_{inv}^{3\pi}$ both improves the efficiency with which the true decay product is selected in τ decay events, and provides additional rejection of the neutral current background (sections 4.2.1 and 4.6); the inclusion of 3π structure information in the transverse mass calculation, yielding the required mass variable M_R , also affords improved rejection of the neutral current background (section 4.3).

Beyond this, however, the results are somewhat disappointing. A careful consideration of the kinematics of opening angle variables shows that they provide negligible discrimination power when the 3π mass and energy are already taken into account in the analysis (section 4.4); and the polarisation-sensitive variables, despite their apparent usefulness for future experiments, provide no discrimination for individual events under NOMAD experimental conditions (section 4.5). Accordingly, the use of 3π structure variables in this analysis is limited to the employment of the invariant mass in 3π candidate selection, and the use of the required mass M_R in place of the traditional variable M_{\perp} .

We should note, however, that we are not yet using the invariant mass information in an optimal way. Even after the cut in $M_{inv}^{3\pi}$ is imposed during the 3π selection, there is a clear difference in the distribution of $M_{inv}^{3\pi}$ between the tau decay and neutral current events (figure 4.6). Furthermore the two-pion mass variables s_1 and s_2 , which we have not used in the analysis as they are not amenable to simple cuts, also exhibit signal/background differences at the distribution level (figure 4.10). There is clear potential for the use of a likelihood function to combine the $M_{inv}^{3\pi}$, s_1 and s_2 information, providing an additional tool for neutral current rejection. Such a function would be most easily employed if cuts in the kinematic variables were also implemented through a likelihood technique, allowing a combination of the kinematic and structure information, as the discrimination power of the mass variable distributions alone is rather modest.

This is precisely the approach which has been taken in the established analysis of the three-prong decays within the NOMAD experiment (Astier et al., 1999a, section 8). In that case a likelihood technique has been employed from the outset, whereas we have adopted the complementary approach of a critical study of the analysis variables, and have not yet proceeded to implement likelihood functions. The combination of the two approaches in the future should lead to an improvement in the sensitivity of the analysis: at

the very least, an improvement over the sensitivity of the analysis presented in this work.

Chapter 5

Particle identification using TRD information

5.1 Introduction

We now set the three-prong decay analysis to one side in order to discuss an important problem in particle identification.

In section 2.3 we noted that the electronic tau decay $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ is the most important signal of ν_τ appearance in the NOMAD experiment, since the small ν_e contamination of the neutrino beam makes the $\nu_e + N \rightarrow e^- + X$ background relatively easy to suppress. The background due to the numerous neutral current events $\nu_\mu + N \rightarrow \nu_\mu + X$, on the other hand, threatens to overwhelm the $\tau \rightarrow \nu \bar{\nu} e$ search if pions in these events can be misidentified as electrons. It is therefore necessary to identify electrons with high efficiency while achieving very high rejection factors against pions and other particles.

Accordingly the NOMAD includes a large transition radiation detector (TRD) to distinguish electrons from other particles using the soft X-ray component of the transition radiation. The TRD itself has been described in section 2.4.6 above; the algorithm used to discriminate between electrons and pions, based on TRD information, is described in some detail in this chapter (section 5.2).

The three-prong and other hadronic decay searches are presented with the opposite problem: the rejection of electrons from ν_e and $\bar{\nu}_e$ charged-current interactions, while retaining a high efficiency for pions and other particles. While this task of “electron veto” is not one of the design-aims of the NOMAD, the electron-pion discrimination algorithm can easily be adapted to this purpose, as described in section 5.2.4 below. This technique is useful to any analysis requiring a sample of hadrons where the contamination

of electrons must be controlled.

A further adaption of the TRD algorithm allows discrimination between any two particle species, based on the variation of *ionisation* energy loss (dE/dx) with particle mass at a given momentum. The TRD provides some sensitivity to this variation as the straw tubes sample the dE/dx of a particle track at up to nine points (see section 2.4.6 above). Work carried out to provide and test this discrimination is described in section 5.3. Some tests of TRD-based particle identification are then discussed in sections 5.3.6 through 5.3.8, while its application to a search for ϕ decays is described in section 5.3.9.

5.2 Electron/pion discrimination

5.2.1 Transition radiation in the NOMAD

Transition radiation (TR) was first predicted by Ginzburg and Franck (1946), and the theory of the phenomenon is well-known (Garibian, 1958, 1960; Durand, 1975; Artru et al., 1975; Garibian et al., 1975). A charged particle crossing an interface between two media of different dielectric constants emits radiation at a typical angle of $1/\gamma$, where γ is the Lorentz factor of the particle. Classically the emission follows from the requirement that the electromagnetic field due to the particle be continuous at the boundary of the media: the emission of photons can be described as part of the field being “shaken loose” as it makes the transition from one medium to the other.

Both the total energy emitted and the typical energy of the emitted photons scale with γ , so the effect may be used to discriminate between electrons and other particles on the basis of the small electron mass (*e.g.* $m_\pi/m_e \approx 273$). For $\gamma \gtrsim 1000$, corresponding to electrons of momenta $p \gtrsim 0.5 \text{ GeV}/c$, the radiation is in the soft X-ray (keV) range and thus comparable with the energy deposited by a particle in a practical gas-filled detector due to ionisation. Choosing a detector with xenon as the principal gas allows efficient TR detection due to the strong X-ray absorption by xenon. However the quantum yield at a single radiating boundary is very low (of order of the fine structure constant α) so a practical detector requires a “radiator” made from foam or a stack of many foils, to provide many transitions at which radiation can occur. In this case both absorption of X-rays in the foils, and interference between the radiation from the various interfaces must be taken into account.

In both the design of the TRD (Bassompierre et al., 1998a) and the analysis to be described in the following sections, a dedicated Monte Carlo




Figure 5.1: Comparison of experimental (points) and simulated (histogram) energy deposition by 10 GeV/c muons in the TRD straw tubes. The same data are shown in linear and logarithmic scales: note the excellent agreement at both low and high deposited energies. (From Bassompierre et al., 1998b, figure 10.)




Figure 5.2: Comparison of experimental (points with error bars) and simulated (histogram) energy deposition by delta rays in the TRD straw tubes. (From Bassompierre et al., 1998b, figure 14.)

simulation program (Fazio et al., 1994) was used, simulating ionization losses in the straw tubes using the model of Ermilova et al. (1977) and Allison and Cobb (1980), and the emission and absorption of transition radiation in the radiators using the formulae of Garibian et al. (1975). A comparison of the energy deposition predicted with this program with experimental results has been published by Bassompierre et al. (1998b): results for muons, a test of the dE/dx simulation, are reproduced as figure 5.1; results for a sample of delta-rays, a test of the dE/dx and TR simulation for electrons, as figure 5.2. The agreement in both cases is excellent over several orders of magnitude.




Figure 5.3: The number of TR photons emitted by a NOMAD TRD 315-foil radiator as a function of the Lorentz factor γ , as predicted by the simulation described in the text (Fazio et al., 1994). (From Bassompierre et al., 1998b, figure 2.)

For the TRD radiators installed in the NOMAD (section 2.4.6) the mean number of (X-ray) TR photons emitted, according to the simulation just described, is shown in figure 5.3. It is clear that a TRD module effectively functions as a threshold detector for particles at a boost $\gamma \approx 2000$, although the small number of TR photons involved means that the actual response is subject to large fluctuations.

5.2.2 A likelihood ratio for e/π discrimination

For all charged particles, the Landau distribution of ionisation energy deposition (see section 5.3.1), and for electrons the probabilistic nature of the transition radiation process, mean that the energy deposited in the TRD straw tubes follows a broad distribution. As an example, the probability distribution for the deposition of energy E in a straw tube is shown in figure 5.4 for both pions and electrons at $10\text{ GeV}/c$. The absorption of transition radiation causes the distribution for electrons to be both larger in the mean and broader in shape, although there is still a substantial overlap with the distribution for pions.




Figure 5.4: Normalised spectra of energy deposited in the TRD straw tubes by $10\text{ GeV}/c$ pions and electrons at normal incidence, in the Monte Carlo simulation. (From Bassompierre et al., 1998b, figure 3.)

For a given track in the detector there will be as many as nine associated “hits” in the TRD straw tubes, *i.e.* one for each TRD module. The optimum way of using the information from these hits in separating electrons from

pions is to form a *likelihood ratio*¹

$$\mathcal{L} = \sum_{i=1}^{N_h} \log \frac{P(E_i | e; p)}{P(E_i | \pi; p)}, \quad (5.1)$$

and use its value as the basis for electron/pion discrimination. In this expression $N_h \leq 9$ is the number of TRD planes recording an energy deposit by the particle, and E_i are the individual energies. As shown in figure 5.4, $P(E_i | e; p)$ is the probability for an electron to deposit the energy E_i , and $P(E_i | \pi; p)$ the probability for a pion to deposit E_i ; the P are normalised spectra of energy deposited by e, π , and depend on the momentum of the particle p . In principle the energy deposition spectra depend on a great many other variables, but the continuous calibration procedure described in section 2.4.6 effectively removes the dependence on gas composition, temperature, *etc.* (see also Bassompierre et al., 1998a); studies indicating that the dependence on other variables, such as impact point on the radiator, may be ignored, are discussed by Bassompierre et al. (1998b).

5.2.3 Construction of acceptance functions

By inspection the likelihood ratio defined at (5.1) is sensitive to differences both of scale and of shape between the distributions for electrons and pions. Consulting also figure 5.4 it is clear that for the majority of TRD hits due to electron tracks, the probability $P(E_i | e; p)$ will be larger than $P(E_i | \pi; p)$, so $\log \frac{P(E_i | e; p)}{P(E_i | \pi; p)}$ will contribute a positive term to the sum; for the majority of hits due to pion tracks, $P(E_i | \pi; p)$ will be larger, contributing a negative term. The result is a strong separation between the distributions of \mathcal{L} for pions and electrons, as shown in figure 5.5.

To achieve a given rejection factor k against pions it is necessary to choose a cut in \mathcal{L} which accepts only $\frac{1}{k}$ of the distribution for pions. (The broadly cross-hatched sample of “good electrons” in figure 5.5 corresponds to $k = 1000$, accepting only 10^{-3} of the pion distribution.) So given likelihood distributions $\frac{dN_e}{dl}$ and $\frac{dN_\pi}{dl}$ we define *acceptance functions*

$$A_j(\mathcal{L}; p, N_h) = \frac{1}{N_j} \int_{\mathcal{L}}^{\infty} dl \left. \frac{dN_j}{dl} \right|_{p, N_h}, \quad j = e, \pi \quad (5.2)$$

and apply a cut $A_\pi(\mathcal{L}; p, N_h) < \frac{1}{k}$ to reject pions.

¹In keeping with normal practice the logarithm of the likelihood ratio is the actual quantity used: the likelihood ratio $\prod_{i=1}^{N_h} \frac{P(E_i | e; p)}{P(E_i | \pi; p)}$ extends over many orders of magnitude and would be intractable.




Figure 5.5: Distributions of likelihood ratio \mathcal{L} for pions (left) and electrons (right) for $p = 10 \text{ GeV}/c$ and $N_h = 9$ in the Monte Carlo. The cross-hatched “good π ” and “good e ” obey cuts which reject all but 10^{-3} of the electrons and pions respectively (see the text).

In practice, the spectra $P(E_i | j; p)$, $j = e, \pi$ are stored for six standard momenta ($p = 0.5, 1.0, 2.0, 3.5, 10.0$ and $50.0 \text{ GeV}/c$) and intermediate values of p are treated by logarithmic interpolation at run-time, to produce the likelihood ratio \mathcal{L} in (5.1). Tables of the acceptance functions $A_j(\mathcal{L}; p, N_h)$, $j = e, \pi$ are stored for the same six standard momenta, and for $4 \leq N_h \leq 9$, with two-dimensional linear interpolation used in \mathcal{L} and p to give the final result. The choice of the limiting momenta 0.5 and $50.0 \text{ GeV}/c$ reflects the “threshold” behaviour in γ noted above: for electrons in this momentum range, $10^3 < \gamma < 10^5$, *i.e.* near or above the threshold, while for pions, $3.6 < \gamma < 360$, below the threshold.

The pion rejection which can be achieved by this technique is a function of

- the number of hits N_h , the rejection improving with the additional information from each hit;
- the momentum p , the rejection improving with momentum until elec-




Figure 5.6: Expected pion acceptance ϵ_π for fixed values of electron identification efficiency ϵ_e , as a function of the number of hits N_h , for 1 GeV/c (left) and 10 GeV/c (right). The value of ϵ_e for each line is indicated where it meets the right-hand boundary of the plot. (From Bassompierre et al., 1998b, figure 7.)

trons reach the saturation level of $N_{TR} \approx 3$ around $p \approx 50$ GeV/c; the probability for pions to emit TR photons becomes appreciable in this same regime (see figure 5.3);

- the electron identification efficiency, with the pion rejection increasing as progressively lower electron ID efficiencies are tolerated.

This dependence is shown graphically in figure 5.6. The standard electron identification cut within NOMAD analyses requires a pion efficiency

$$\begin{aligned} \epsilon_\pi &= A_\pi(\mathcal{L}; p, N_h) \\ &< 10^{-3} \end{aligned} \quad (5.3)$$

which corresponds to an electron identification efficiency of about 90% over a broad sample of tracks. From figure 5.6 we can see that for nine-hit tracks at high momentum, the electron ID efficiency improves to 95%.

5.2.4 Interpretation of acceptance functions in analysis

Acceptance functions for electron rejection

The choice of $\epsilon_\pi < 10^{-3}$ for electron identification reflects the need for a pure sample of electrons in the $\tau \rightarrow \nu\bar{e}$ search. For many other analyses, such as QCD or fragmentation studies or searches for resonances, the requirement is instead to form a sample of hadrons with low contamination by electrons. In a $\tau^- \rightarrow \nu_\tau + \text{hadrons}$ search, it is also necessary to veto events where the basic process is in fact $\nu_e + N \rightarrow e^- + X$, typically with a hard electron from the deep inelastic scattering which could form a false tau-decay candidate ($\tau \rightarrow \nu\pi$) or part of such a candidate ($\tau \rightarrow \nu 3\pi$).

It follows from the approximately 90% efficiency for electrons of the $\epsilon_\pi < 10^{-3}$ cut that about 10% of electrons will fail such a requirement and therefore contaminate the sample; that is, the cut imposes a rejection factor of about 10 on electrons. For some purposes this will be too loose a cut, and furthermore the rejection of electrons will be a function of both p and N_h . One would prefer a cut which gave a consistent rejection of electrons, and could be adjusted according to the needs of each analysis.

Such a cut is easily performed using the likelihood ratio \mathcal{L} , by applying the logic of the previous section directly to this new case. If we construct alternative acceptance functions

$$\tilde{A}_j(\mathcal{L}; p, N_h) = \frac{1}{N_j} \int_{-\infty}^{\mathcal{L}} dl \left. \frac{dN_j}{dl} \right|_{p, N_h}, \quad j = e, \pi \quad (5.4)$$

then a cut $\tilde{A}_j(\mathcal{L}; p, N_h) < \frac{1}{k}$ requires that particles have a likelihood ratio \mathcal{L} sufficiently *low*, that only a fraction $\frac{1}{k}$ of particle type j will have a likelihood ratio $l \leq \mathcal{L}$. (A cut in the functions $A_j(\mathcal{L}; p, N_h)$ requires that \mathcal{L} be sufficiently *high* ... *etc.*) This choice is illustrated by the “good π ” sample marked by the close cross-hatching in figure 5.5, corresponding to a cut

$$\begin{aligned} \tilde{\epsilon}_e &= \tilde{A}_e(\mathcal{L}; p, N_h) \\ &< 10^{-3}. \end{aligned} \quad (5.5)$$

The symmetry of this approach is apparent. It should also be clear from equations (5.2) and (5.4) that

$$\tilde{A}_j(\mathcal{L}; p, N_h) \equiv 1 - A_j(\mathcal{L}; p, N_h) \quad (5.6)$$

One should note that an electron acceptance of 10^{-3} can also be achieved using a cut $\epsilon_e < 10^{-3}$, and many other arbitrary choices of region in \mathcal{L} . By construction of the likelihood ratio \mathcal{L} the choice of cut variable $\tilde{\epsilon}_e = \tilde{A}_e(\mathcal{L}; p, N_h)$, which integrates over the cases where the pion-to-electron ratio is highest, provided maximal acceptance for pions among all of these choices.

“Flatness” and “cutoffs” of acceptance functions

When considering a sample of pions it is also useful to consider the acceptance function $\tilde{\epsilon}_\pi = \tilde{A}_\pi(\mathcal{L}; p, N_h)$. This has the important property that a pure sample of pions should have a flat distribution of $\tilde{\epsilon}_\pi$ values from 0 to 1 (see the definition (5.4)). Considering the position of the electron distribution in figure 5.5, and the choice of integration bounds in equation (5.4), it is clear that the tail of the electron distribution will appear as an excess of particles with high values of $\tilde{\epsilon}_\pi$, while the bulk of the electrons will form a spike at $\tilde{\epsilon}_\pi = 1$: an example is shown in figure 5.7.




Figure 5.7: The distribution of the pion acceptance variable $\tilde{\epsilon}_\pi$ for negative primary-vertex tracks in the NOMAD data. Tracks passing an electron rejection cut $\tilde{\epsilon}_e < 0.10$ are shown cross-hatched: the accumulation of electrons at $\tilde{\epsilon}_\pi \approx 1$ is apparent.

The effect of a cut in $\tilde{\epsilon}_e$ will also be apparent in such a distribution: the “ambiguous” region excluded by the cut contains the pions with the highest values of $\tilde{\epsilon}_\pi$, so a $\tilde{\epsilon}_\pi$ distribution will run from 0 to some value $f < 1$ which will approximate the efficiency of the cut for pions. In practice this cutoff will be blurred as the likelihood ratio distributions “move together” or apart, as p and N_h vary—a fixed rejection factor for electrons does not yield a fixed efficiency for pions, and vice versa. An example is shown in figure 5.8, where progressively stronger cuts in $\tilde{\epsilon}_e$ remove the accumulation of electrons at $\tilde{\epsilon}_\pi \approx 1$ and then introduce a sloping cutoff in the distribution at high $\tilde{\epsilon}_\pi$.




Figure 5.8: The distribution of the pion acceptance variable $\tilde{\epsilon}_\pi$ for negative primary tracks, under progressively tighter electron rejection cuts: (solid line) $\tilde{\epsilon}_e < 10^{-1}$, (dashed line) $\tilde{\epsilon}_e < 10^{-2}$, (hatched) $\tilde{\epsilon}_e < 10^{-3}$, (cross-hatched) $\tilde{\epsilon}_e < 10^{-4}$. Note the effect on the acceptance for pions for cuts harder than $\tilde{\epsilon}_e < 10^{-2}$.

Electron contamination and electron rejection

To date the discussion has focussed on nominal electron rejection levels k , achievable by a cut $\tilde{\epsilon}_e < \frac{1}{k}$. In passing we should note that the *contamination* of electrons in a sample of tracks is the more relevant quantity for analysis—hence the usefulness of the distribution of pion acceptance $\tilde{\epsilon}_\pi$, just discussed. Depending on the fraction of electrons in the initial sample, weaker or stronger cuts in $\tilde{\epsilon}_e$ will have to be chosen, in order to achieve a pion sample with an acceptably low contamination of electrons.

It is convenient to use primary-vertex tracks from $\nu_e + N \rightarrow e^- + X$ Monte Carlo events to illustrate this point (figure 5.9). The positive track sample has a small initial contamination of positrons due to photon conversions close to the primary vertex, Dalitz decays of π^0 , *etc.*, while the negative track sample contains in addition the hard electrons from the deep inelastic scattering process, visible as a striking peak at $\tilde{\epsilon}_\pi = 1$ (figure 5.9(a)). A cut $\tilde{\epsilon}_e < 0.10$ is sufficient to remove most of the peak in the positive sample, while a significant contamination remains in the negatives, about 8% in this case (figure 5.9(b)).

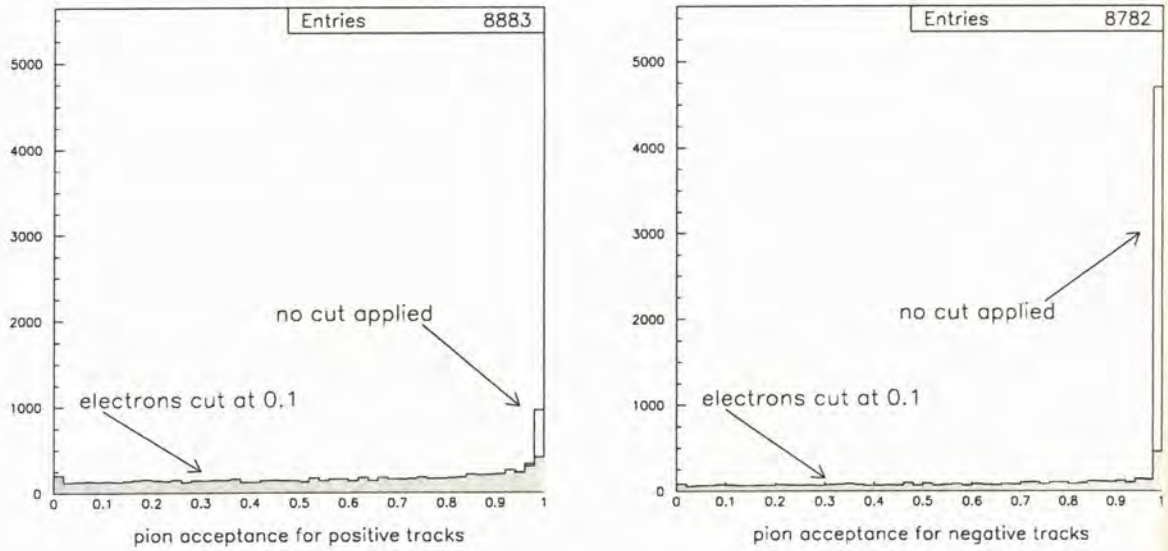
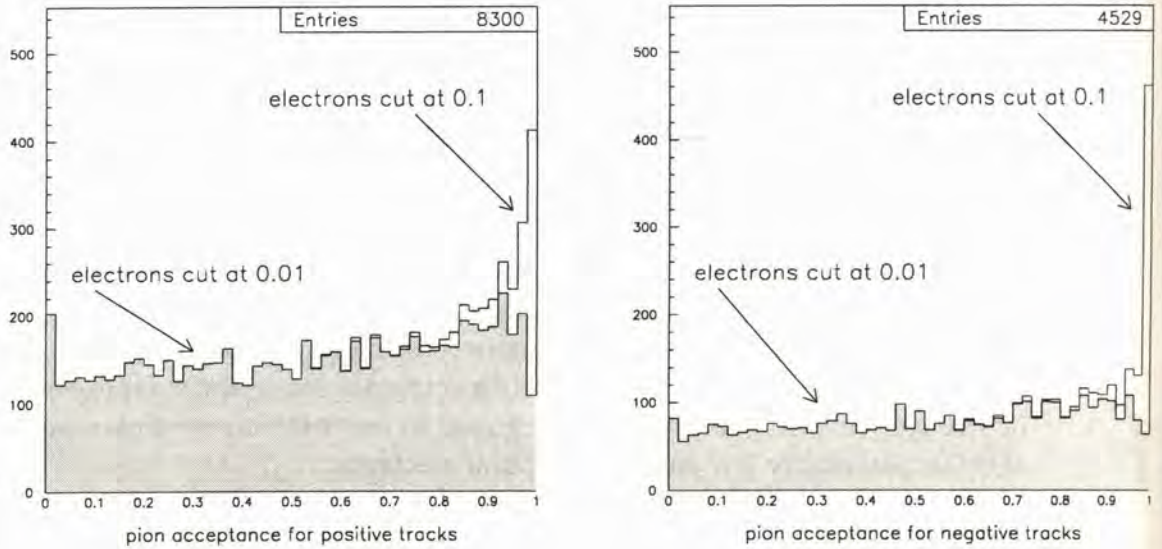
(a) Tracks in the $\nu_e + N \rightarrow e^- + X$ Monte Carlo(b) Tracks after an $\tilde{\epsilon}_e < 0.1$ cut

Figure 5.9: The distribution of the pion acceptance variable $\tilde{\epsilon}_\pi$ for primary tracks in ν_e charged-current Monte Carlo events, separated into (left plots) positive tracks and (right plots) negative tracks. Electrons form a significant peak at $\tilde{\epsilon}_\pi = 1$ even after the $\tilde{\epsilon}_e < 0.1$ cut in the negative sample.

5.2.5 Electron/proton discrimination

In figure 5.9(b) an additional population of tracks is visible in the positives sample: concentrated at $\tilde{\epsilon}_\pi \approx 1$, as are electrons, but not efficiently removed by the cut $\tilde{\epsilon}_e < 0.1$, unlike electrons. Dividing the Monte Carlo sample by track type (figure 5.10) it is apparent that these tracks are protons, with an ionisation energy deposition larger than pions at low momentum: they form an intermediate case between pions and electrons.

(If we suppress protons in the sample of positive tracks and reproduce the comparison between positive and negative tracks—as shown in figure 5.11—then the effect of the larger initial proportion of electrons in the negative sample is much more clear *cf.* figure 5.9.)

This difference in the ionisation energy deposition behaviour of different species provides an opportunity for charged particle identification in NO-




Figure 5.10: The distribution of $\tilde{\epsilon}_\pi$ for positive primary tracks in ν_e charged-current Monte Carlo events, after a cut $\tilde{\epsilon}_e < 0.1$ against electrons, with particles separated by species: pions are shown cross-hatched, kaons hatched; the remainder of the distribution is made up of protons, excepting a very small admixture of positrons, visible in the highest bins (the gap between the dashed and solid histogram).




Figure 5.11: The distribution of the pion acceptance variable $\tilde{\epsilon}_\pi$ for primary tracks in ν_e charged-current Monte Carlo events, after a cut $\tilde{\epsilon}_e < 0.1$ against electrons. In the positive track sample (left) protons have been excluded from the distribution; compare figure 5.9(b).

MAD, which will be discussed in section 5.3 below. From the point of view of the $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ search and other electron-oriented analyses, the effect is a nuisance, since it leads to positron samples (useful as control samples for electron analyses, for example) contaminated by protons. This is prevented by constructing an additional likelihood ratio

$$\mathcal{L}_{e/p} = \sum_{i=1}^{N_h} \log \frac{P(E_i | e; p)}{P(E_i | p; p)} \quad (5.7)$$

for separation of protons and positrons, and the corresponding acceptance functions

$$A_j^{e/p}(\mathcal{L}_{e/p}; p, N_h) = \frac{1}{N_j} \int_{\mathcal{L}_{e/p}}^{\infty} dl \left. \frac{dN_j}{dl} \right|_{p, N_h}, \quad j = e, p. \quad (5.8)$$

Where necessary, protons are suppressed by performing a suitable cut $A_p^{e/p}(\mathcal{L}_{e/p}; p, N_h) < \frac{1}{k}$ on positive tracks in the low-momentum ($p < 1.5 \text{ GeV}/c$) region.

5.2.6 The effect of track angle

If we consider plots of the pion acceptance $\tilde{\epsilon}_\pi$ such as figure 5.11, a residual right-to-left downward slope is apparent even when protons have been suppressed. This corresponds to a (sub-) sample with energy deposition in the TRD straw tubes systematically higher than pions of the same momentum.

The effect is due to particles crossing the TRD at a non-zero angle θ to the $x-z$ (horizontal) plane of NOMAD: such particles have ionisation tracks longer than those of normally incident particles, by a factor $1/\cos\theta$; to first order the mean energy deposited by such tracks is increased by the same factor. Low momentum particles are especially subject to this effect due to the bending of tracks in the vertical plane by the 0.4 T magnetic field. (The distribution of $\tilde{\epsilon}_\pi$ for the soft π^- from Λ^0 decays in the data is shown in figure 5.12 as an example.)

For electrons the situation is more complicated. In addition to the energy deposition due to ionisation there is of course the transition radiation: at non-normal incidence both the production and absorption of TR in the radiators is enhanced. The various effects were found to approximately cancel each




Figure 5.12: Pion acceptance $\tilde{\epsilon}_\pi$ for π^- from identified $\Lambda^0 \rightarrow p\pi^-$ decays in the data.

other by Bassompierre et al. (1998b) and so dependence on track angle was left out of the construction of the probability distributions for the likelihood ratio defined at (5.1).

5.2.7 Treatment of overlapping tracks

Up to this point it has been tacitly assumed that all TRD straw tube “hits” are associated to only one particle track. In fact it is relatively common for particle tracks to be at least partially overlapping in the $x - z$ projection in the TRD region (the straw tubes are aligned vertically; see section 2.4.6). The measured energy deposition in a straw tube crossed by two tracks is, to first order, equal to the sum of the depositions by the two tracks in isolation, and so clearly some special treatment is required for such cases.

In practice, tracks with so-called shared hits are divided into two groups, *viz.*

tracks with 1, 2 or 3 shared hits: These tracks are treated by constructing the likelihood ratio (5.1) using the *isolated* hits only, and ignoring the shared hits. Distributions of acceptance functions in this chapter have in fact included tracks treated in this way.

tracks with 4 or more shared hits: In this case the overlapping tracks are treated together. Noting that for a given pair of momenta (p_1, p_2) , the distributions of deposited energy for each of the four particle identity hypotheses $\pi(p_1)\pi(p_2)$, $e(p_1)\pi(p_2)$, $\pi(p_1)e(p_2)$ and $e(p_1)e(p_2)$ may typically be distinguished (figure 5.13), four likelihoods

$$\mathcal{L}_{kl} = \sum_{i=1}^{N_s} \log P(E_i | k(p_1)l(p_2)) + \sum_{j=1}^{(N_h - N_s)} \log P((E_j^1 + E_j^2) | k(p_1)l(p_2)) \quad (5.9)$$

are constructed, where

k, l run over e, π ;

N_s is the number of shared hits;

N_h is the total number of TRD modules with hits; and

P are probability densities for the energy E_i to be deposited, under the hypotheses shown.

Where modules contain non-shared hits, these are added to form an effective shared hit $(E_j^1 + E_j^2)$ to allow a consistent treatment of all hit modules. The hypothesis corresponding to the largest of the likelihoods is chosen and the particles identified as electrons or pions accordingly.




Figure 5.13: Simulated response in the TRD straw tubes to isolated tracks (upper plot) and overlapping tracks (lower plot), with the momenta in GeV/c shown in parentheses. (From Bassompierre et al., 1998b, figure 9.)

Note that because likelihood ratios, acceptance distributions, *etc.* are not defined in this second case, the identification efficiency for one species or the rejection factor for the other species cannot be set at will. The maximum-likelihood method was chosen in order to prevent overlapped pion tracks being identified as electrons, *i.e.* to remove a potentially dangerous contamination of the electron sample, in the $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ search. It is not necessarily suited to the suppression of electrons in the $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n \geq 0)$ search, a problem which will be further discussed in chapter 6 (section 6.7.5).

Where three or more tracks overlap, but there are only 1–3 shared hits, likelihood ratios are calculated using the isolated hits only, as in the two-track case. Where there are 4 or more shared hits no special treatment is attempted.

5.3 Discrimination between two particle species: the general case

The likelihood ratio method described in sections 5.2.2 and 5.2.3 is not specific to the electron/pion (TR/non-TR) case. Provided the spectra of energy deposition in the TRD straw tubes for two types of particle are different, they can in principle be distinguished by the same technique. Accordingly the likelihood ratio method has been adapted for several common pairs of particle species, as discussed in sections 5.3.1 through 5.3.4; its anticipated performance is discussed in section 5.3.5; tests of the method on various samples of tracks in the data are presented in sections 5.3.6 through 5.3.8.

With full particle identification for events in the NOMAD it would be possible to routinely construct Lorentz invariant quantities (rather than being restricted to the laboratory frame), study exclusive decays and search for various rare processes, and recover some information on lost neutral particles, the Fermi momentum of the struck nucleon, and the incoming neutrino energy. While the achievable discrimination between particle species is not sufficient for all this, some new avenues for analysis have been opened, one of which—a search for decays of the ϕ meson—is discussed in some detail below (section 5.3.9). A summary of the particle discrimination technique is given in section 5.3.10.

5.3.1 Energy loss due to ionisation

In the general case we rely on differences in deposited energy due to ionisation to distinguish particle species. In the momentum range of interest for the NOMAD experiment, the mean rate of energy loss due to ionization is given by the Bethe-Bloch formula,

$$-\frac{dE}{dx} = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \log \frac{2m_e c^2 \beta^2 \gamma^2 T_{max}}{I^2} - \beta^2 - \frac{\delta}{2} \right], \quad (5.10)$$

where z, β, γ are the particle charge, velocity and Lorentz factor respectively, and Z and A are the atomic number and mass number of the medium. The definition of T_{max} and other details are given in (Caso et al., 1998, section 23). The formula is plotted for singly-charged particles in various media in figure 5.14.

For a given material the mean dE/dx is a function of velocity β ; if the momentum p is also given, the only remaining dependence is on the particle mass M . Because of this, and the functional form of (5.10), it is convenient




Figure 5.14: Mean rate of energy loss due to ionization in various materials, taken from (Caso et al., 1998, section 23).

to consider the mean dE/dx as a function of the generalised velocity $\beta\gamma = p/Mc$. From figure 5.14 it is clear that there are three regimes:

1. the low-velocity regime, extending from sub-relativistic velocities to $\beta\gamma \lesssim 3$, where the energy lost is a rapidly falling function of velocity (as $\beta^{-5/3}$);
2. “minimum ionisation”, near $\beta\gamma \approx 3$; and
3. the “relativistic rise” for $\beta\gamma \gg 3$, where ionisation is a weakly rising function of $\beta\gamma$.

The formula (5.10) gives only the mean value of dE/dx : at any given velocity the actual energy loss in a thin layer of a medium has a broad distribution with a long (“Landau”) tail to high values. Thus even in cases where the mean dE/dx are well-separated, distributions for individual “hits” will




Figure 5.15: The distribution of energy deposited in the TRD straw tubes by pions and protons at normal incidence, for various momenta: (a) $p = 0.5$ GeV/c ($\beta\gamma = 3.6$ for π and 0.53 for p); (b) $p = 1.0$ GeV/c ($\beta\gamma = 7.2, 1.07$ respectively); (c) $p = 2.0$ GeV/c ($\beta\gamma = 14.3, 2.2$); and (d) $p = 10.0$ GeV/c ($\beta\gamma = 72, 10.7$). Pions are on the “relativistic rise” throughout, while protons move from the sub-relativistic region (a), through minimum ionisation (near $p = 2.0$ GeV/c), to the relativistic rise (d). All distributions were generated using the code of Fazio et al. (1994) described in section 5.2.1.

overlap, so that several measurements will be required for effective discrimination between particle mass hypotheses. As an example, distributions for pions and protons are compared at four momenta in figure 5.15 (*cf.* figure 5.4 above for pions and electrons).

Pions and protons are well-separated in mass ($m_p/m_\pi = 6.7$) and hence in $\beta\gamma$ at any given momentum. In this case it will be clear from figures 5.14 and 5.15 that there are three momentum regimes:

1. $p \ll 3m_p c$, where protons are at low velocities, depositing much more energy than pions so that discrimination will be very good (figure 5.15(a)), albeit a strong function of momentum (*cf.* figure 5.15(b), where discrimination is much weaker);
2. $p \approx 2 \text{ GeV}/c$, where protons ($\beta\gamma \approx 2.2$) and pions ($\beta\gamma \approx 14.3$) lie on each side of the minimum ionisation, with similar energy deposition, so that discrimination will not be possible (figure 5.15(c)); and
3. $p > 3m_p c$, where both particles are on the relativistic rise, with pions depositing more energy on average (figure 5.15(d)), and relatively weak discrimination will be possible.

(In other cases of two species q and r with $m_q > m_r$, the mass difference will be smaller, so that regime (1) will be given instead by $p < 3m_r c$ (the lighter particle is at or below minimum ionisation), and regime (2) where there is no discrimination will be given to a good approximation by $p \approx 3\sqrt{m_q m_r} c$, while discrimination in regime (3) may be too weak for practical use.)

In the first of these regimes, even a very crude method would yield an effective separation of protons and pions, and indeed many other pairs of particle species. In the third regime efficient use must be made of all the information available, and so we adapt the likelihood ratio method described in sections 5.2.2 and 5.2.3 above to this new case. In principle a simplified method could be used with little loss of efficiency, since the width and other properties of the dE/dx distribution approximately scale with the mean deposition² (Allison and Cobb, 1980), but the approach of re-implementing the existing e/π discrimination procedure for the general case was felt to be more straightforward.

²*cf.* the e/π case, where the shape of the energy deposition spectrum for electrons, as the sum of dE/dx and TR components, evolves with γ .

5.3.2 Adaption of the likelihood ratio method

Likelihood ratio and acceptance functions

To discriminate between two species q and r , then, we construct a likelihood ratio as in equation 5.1,

$$\mathcal{L}_{q/r} = \sum_{i=1}^{N_h} \log \frac{P(E_i | r; p)}{P(E_i | q; p)}, \quad (5.11)$$

where $N_h \leq 9$ is the number of TRD planes recording an energy deposit by the particle, E_i are the individual energies, p is the particle momentum, and the $P(E_i | j; p)$ $j = q, r$, are the probabilities for deposition of the energy E_i under the hypothesis shown—equivalent to normalised energy deposition spectra such as those shown in figure 5.15. We then define acceptance functions

$$\tilde{A}_j^{q/r}(\mathcal{L}_{q/r}; p, N_h) = \frac{1}{N_j} \int_{-\infty}^{\mathcal{L}_{q/r}} dl \left. \frac{dN_j}{dl} \right|_{p, N_h}, \quad j = q, r \quad (5.12)$$

for a given $\mathcal{L}_{q/r}$. By convention we will assign the labels q and r such that $m_q > m_r$. From the form of (5.11) it is apparent that the distribution of $\mathcal{L}_{q/r}$ for the heavier particle q will be dominated by low values $\mathcal{L}_{q/r} < 0$ while the distribution for the lighter particle r will be dominated by high values $\mathcal{L}_{q/r} > 0$. This is shown for the most spectacular case, protons and pions at low momentum, in figure 5.16 (*cf.* equation (5.1) and figure 5.5).




Figure 5.16: Distribution of the likelihood ratio $\mathcal{L}_{p/\pi}$ for protons (left) and pions (right) at $p = 0.5 \text{ GeV}/c$ and $N_h = 9$.

To impose a rejection factor of k on the lighter particle r , we make a cut $\tilde{\epsilon}_r = \tilde{A}_r^{q/r}(\mathcal{L}_{q/r}; p, N_h) < \frac{1}{k}$. We would then expect a pure sample of the heavier particle q to have a flat distribution of the acceptance function $\tilde{\epsilon}_q = \tilde{A}_q^{q/r}(\mathcal{L}_{q/r}; p, N_h)$ up to some cutoff, characteristic of its efficiency (*cf.* section 5.2.4 above).

In the less-likely case where a clean sample of the lighter particle r is required, we may impose a rejection factor of k on the particle q by performing a cut at $\frac{1}{k}$ in the acceptance function

$$\begin{aligned}\epsilon_q &= A_q^{q/r}(\mathcal{L}_{q/r}; p, N_h) \\ &\stackrel{\text{def}}{=} 1 - \tilde{A}_q^{q/r}(\mathcal{L}_{q/r}; p, N_h)\end{aligned}$$

Implementation of the method

In addition to the case of proton/pion discrimination, which has already been used in examples above (figures 5.15 and 5.16) the likelihood ratio method has been implemented for proton/kaon, kaon/pion, kaon/muon and pion/muon discrimination.

In each case the construction of the likelihood ratios $\mathcal{L}_{q/r}$ is in principle straightforward, but the evolution of particle discrimination power with momentum, which is extremely rapid in some cases (such as p/π at low momentum) creates difficulties whenever approximation is used. As for the electron/pion case (section 5.2.3) spectra of the energy deposited in a TRD straw by each particle type (p, K, π, μ) were generated at a limited set of reference momenta:

- the six standard momenta used for the electron/pion case (0.5, 1.0, 2.0, 3.5, 10.0, and 50.0 GeV/ c);
- an additional momentum bin at 20.0 GeV/ c to accommodate the change in kaon/pion discrimination there; and
- *thirteen* additional momenta from 0.25 to 1.25 GeV/ c , to account for the extremely rapid change in the power of kaon/pion and proton/pion discrimination over this range.

For each particle type and momentum value, one million instances were simulated. In all subsequent applications, in particular the calculation of $\mathcal{L}_{q/r}$ based on the hits of a particular track, intermediate momenta were handled using logarithmic interpolation at run-time, as for the e/π case.

The tables of acceptance functions $\tilde{A}_j^{q/r}(\mathcal{L}_{q/r}; p, N_h)$, each generated using one million simulated tracks of the relevant particle types (q and r),

also required some special handling. The electron/pion discrimination power evolves only slowly with momentum (figure 5.6) and in that case it was sufficient to tabulate acceptance functions within the bounds in \mathcal{L} shown in figure 5.5. In the general case this approach was impractical as the effective range of distributions of $\mathcal{L}_{p/\pi}$, to take the most extreme example, changes by two orders of magnitude as a function of momentum. The expedient of setting bounds on the acceptance function tables by hand for each case, after inspecting histograms for several trial bounds, was adopted. Two dimensional linear interpolation in $\mathcal{L}_{q/r}$ and p was then used at run time to obtain actual acceptance function values, as for the e/π case. The spacing of the twenty reference momenta noted above was chosen so that the interpolation in momentum was practical in all cases.

5.3.3 A correction for track angle

In the discussion of electron/pion discrimination it was noted that tracks at non-normal incidence to the TRD have a longer ionisation trail in the straw tubes, and proportionately larger deposited energy, than normally incident tracks (section 5.2.6). In the general case tracks corresponding to both particle ID hypotheses deposit energy only by ionisation (dE/dx) and so a straightforward correction for this effect has been added: all TRD energy depositions are multiplied by a factor

$$\cos(\theta) = \frac{\sqrt{p_x^2 + p_z^2}}{p}$$

before comparison with the reference distributions.

As an example of the effect of the correction, distributions of the proton/pion discrimination variable $\tilde{\epsilon}_\pi = \tilde{A}_\pi^{p/\pi}(\mathcal{L}_{p/\pi}; p, N_h)$ are shown in figure 5.17 for a sample of π^- tracks from identified $\Lambda^0 \rightarrow p\pi^-$ decays, rich in large-angle tracks. The variable $\tilde{\epsilon}_\pi$ should have a flat distribution from 0 to 1 for pions: with no correction applied (top plot) there is a clear enhancement at low values of $\tilde{\epsilon}_\pi$, reflecting the increased energy deposition by these tracks; with the $\cos \theta$ correction applied (bottom plot) the distortion is greatly reduced. (Compare with figure 5.12 above, showing the distribution of the electron/pion discrimination variable on the same track sample.)

When processing the full output of the NOMAD reconstruction code, as in the figure shown, the momentum $\vec{p} = (p_x, p_y, p_z)$ is estimated by taking the extrapolated track momentum at each of the individual TRD straw tubes. Most of the distributions shown in this chapter have been prepared by applying the particle identification algorithm described in section 5.3.2 to




Figure 5.17: Pion acceptance variable $\tilde{\epsilon}_\pi^{p/\pi}$ for the negative arm of identified $\Lambda^0 \rightarrow p\pi^-$ decays, with the track angle correction disabled (upper plot), and enabled (lower plot).

the reduced NOMAD data format (DST) where this level of detail in track extrapolation is not available: in this case the track extrapolation to the preshower detector, immediately following the TRD (see figure 2.6) is used. The approximation should be sufficient for all but the lowest-momentum tracks.

5.3.4 Treatment of overlapping tracks

As noted in section 5.2.7 above, it is relatively common for two tracks to partially overlap in the $x-z$ projection in the TRD region. In this case one or more of the straw tube “hits” associated to a given track in fact correspond to the combined energy deposition due to the two tracks, which is equal to the sum of the depositions of the two tracks in isolation, to first order.

To date no special treatment of such shared hits has been implemented for the general case of discrimination between two particle ID hypotheses. In the plots shown in the remainder of this chapter, tracks for which particle discrimination variables are calculated are required to have no shared hits,

to avoid distortion of the distributions, or false identification of a track as corresponding to the particle with larger energy deposition.

At the next reprocessing of the NOMAD data, a change to the reduced data format (DST) is planned which will flag individual shared hits in the TRD: at that time the algorithm of section 5.3.2 will be adapted to exclude shared hits from the calculation of likelihood ratios $\mathcal{L}_{q/r}$ as is done for the electron/pion likelihood ratio \mathcal{L} .

5.3.5 Five cases of discrimination between species

As noted above, the particle identification algorithm has been implemented for five cases: p/π , p/K , K/π , K/μ and π/μ discrimination. We will consider these briefly in turn.

p/π discrimination

The mass difference between protons and pions is the largest among the five pairs of particles ($m_p/m_\pi = 6.7$) and the discrimination in the p/π case is accordingly the most powerful. Distributions of the proton/pion likelihood ratio $\mathcal{L}_{p/\pi}$ are shown in figure 5.18 for three momenta: $p = 0.5$ and 1.0 GeV/ c from the low momentum region and $p = 10$ GeV/ c from the high momentum region. At low momentum the distributions are well-separated so that individual protons and pions may be readily distinguished: at $p = 0.5$ GeV/ c the distributions overlap only in the extreme tails, so that pure samples of one or other particle may be obtained without any significant loss of efficiency; at $p = 1.0$ GeV/ c particle discrimination is deteriorating, with the proton and pion distributions overlapping so that there is a trade-off between efficiency and purity in obtaining a sample of either particle. At high momentum there is in all cases a significant overlap between the distributions.

It is convenient to quantify the discrimination power available by setting various desired proton identification efficiencies (99%, 95%, 90%, 80% *etc.*) and plotting the corresponding acceptance for pions: this is done in figure 5.19 for the same three momenta, $p = 0.5, 1.0, 10.0$ GeV/ c . Thus, taking the left-most figure ($p = 0.5$ GeV/ c), if we consider 8- and 9-hit tracks we may achieve efficiencies for protons greater than 99% with a cut in the pion acceptance $\tilde{\epsilon}_\pi = \tilde{A}_\pi^{p/\pi}(\mathcal{L}_{p/\pi}; p, N_h) < 10^{-4}$, *i.e.* imposing a rejection factor of 10^4 against pions. Note that in all cases the discrimination power is a strong function of the number of hit planes N_h , so that for $N_h = 4$, the cut against pions must be loosened to $\tilde{\epsilon}_\pi < 0.02$ in order to gain 99% efficiency for protons. Nonetheless, even for $N_h = 3$, the proton identification efficiency exceeds 90% for a cut as tight as $\tilde{\epsilon}_\pi < 10^{-2}$, imposing a rejection factor of




Figure 5.18: Distributions of the proton/pion likelihood ratio $\mathcal{L}_{p/\pi}$ for protons (solid line) and pions (dashed line) at three momenta: (upper plot) $p = 0.5 \text{ GeV}/c$, where discrimination is excellent; (middle plot) $p = 1.0 \text{ GeV}/c$, as the discrimination power is decreasing; (lower plot) $p = 10 \text{ GeV}/c$, where modest discrimination has returned with rising momentum. Number of hit planes $N_h = 9$.

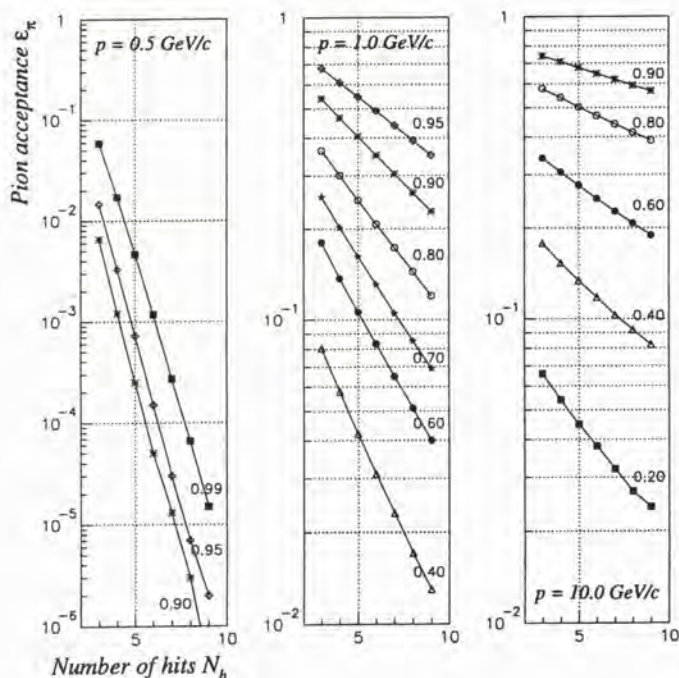


Figure 5.19: Expected pion acceptance versus number of TRD planes for different proton identification efficiencies for particles at $p = 0.5 \text{ GeV}/c$ (left), $1.0 \text{ GeV}/c$ (centre), and $10.0 \text{ GeV}/c$ (right). Compare figure 5.6 above.

100 on pions: at this momentum, we have ideal discrimination between the two particles.

This discrimination power is rapidly lost as momentum rises, so that at $p = 1.0 \text{ GeV}/c$ even a modest cut against pions, say $\tilde{\epsilon}_\pi < 8 \times 10^{-2}$ (a rejection factor of 12.5, chosen for ease of reading from the figure) yields efficiencies for protons ranging from 40% (for $N_h = 3$) to better than 70% (for $N_h = 9$). Above this momentum the discrimination continues to fail until about $p = 2.0 \text{ GeV}/c$, where $\mathcal{L}_{p/\pi}$ distributions for the two particles effectively coincide (figure 5.20) so that no discrimination is possible: this is the second momentum regime discussed above, with the proton ($\beta\gamma = 2.1$) and pion ($\beta\gamma = 14.3$) “on either side” of the minimum ionisation ($\beta\gamma \approx 3$). As momentum increases, with both particles on the relativistic rise in ionisation, some discrimination is recovered (the $10 \text{ GeV}/c$ case in figure 5.19): but in this case a cut at $\tilde{\epsilon}_\pi < 8 \times 10^{-2}$ only reaches an efficiency for protons of 40% in the most favourable case, $N_h = 9$.




Figure 5.20: Distribution of the likelihood ratio $\mathcal{L}_{p/\pi}$ for protons (left) and pions (right) at $p = 2.0 \text{ GeV}/c$ and $N_h = 9$.

When distinguishing protons and pions, then, we have excellent discrimination at low momentum even for a small number of hits in the TRD region, albeit that this discrimination is a strong function of momentum; close to $p = 2.0 \text{ GeV}/c$, discrimination fails completely; and modest discrimination is recovered as momentum rises above this value.

p/K discrimination

Turning to the proton/kaon case, where the two particles are much closer in mass ($m_p/m_K = 1.9$), we find that discrimination is much less power-

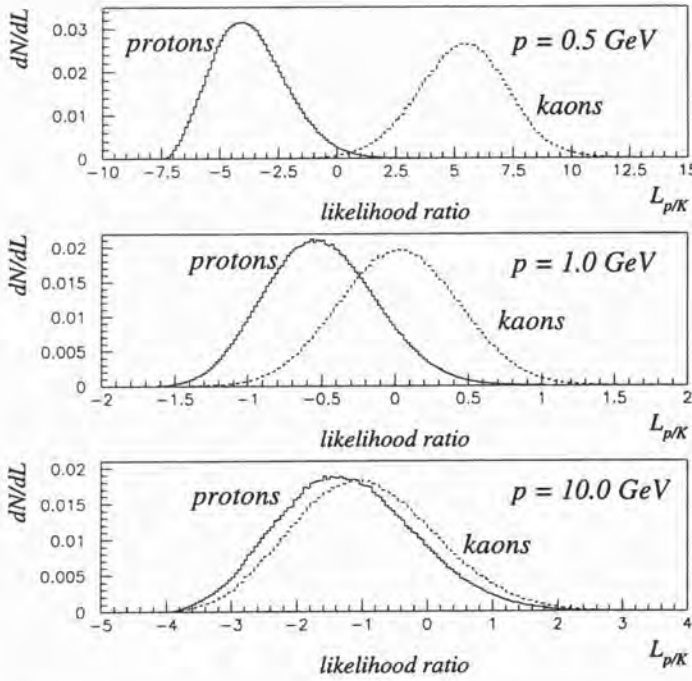


Figure 5.21: Distributions of the proton/kaon likelihood ratio $\mathcal{L}_{p/K}$ for protons (solid line) and kaons (dashed line) at three momenta: (upper plot) $p = 0.5 \text{ GeV}/c$, (middle plot) $p = 1.0 \text{ GeV}/c$, (lower plot) $p = 10.0 \text{ GeV}/c$. Compare with figure 5.18 above.




Figure 5.22: Expected kaon acceptance versus number of TRD planes for different proton identification efficiencies for particles at $p = 0.5 \text{ GeV}/c$ (left), $p = 1.0 \text{ GeV}/c$ (centre), and $p = 10.0 \text{ GeV}/c$ (right). Compare with figure 5.19 above.

ful. Distributions of the proton/kaon likelihood ratio $\mathcal{L}_{p/K}$ are shown in figure 5.21, and kaon acceptance $\tilde{\epsilon}_K = \tilde{A}_K^{p/K}(\mathcal{L}_{p/K}; p, N_h)$ for various proton identification efficiencies in figure 5.22, for the same three momenta $p = 0.5, 1.0, 10.0$ GeV/c.

At low momenta the proton and kaon $\mathcal{L}_{p/K}$ are well-separated, so that discrimination is very good, with a cut $\tilde{\epsilon}_K < 10^{-2}$ imposing a rejection factor of 100 against kaons yielding a 90% efficiency for protons at $N_h = 5$, and better than 99% for a nine-hit track. While falling well short of the spectacular proton/pion discrimination in this region (the rejection factor against pions is 10^{-4} for similar efficiency, see above) this is adequate for any but the most demanding analysis.

Again, the discrimination falls away with rising momentum, so that at 1.0 GeV/c even the most modest kaon rejection requires a loss of efficiency for protons. In the high-momentum region, distributions of $\mathcal{L}_{p/K}$ almost entirely overlap, so that discrimination on a particle-by-particle basis will in no case be practical.

K/ π (and K/ μ) discrimination

In kaon/pion discrimination, the spacing in mass is intermediate between the previous two cases ($m_K/m_\pi = 3.5$). Considering the distribution of the kaon/pion likelihood ratio (figure 5.23), and the pion acceptance $\tilde{\epsilon}_\pi = \tilde{A}_\pi^{K/\pi}(\mathcal{L}_{K/\pi}; p, N_h)$ for various kaon identification efficiencies (figure 5.24) we see a similar pattern as a function of momentum, with good discrimination at $p = 0.5$ GeV/c, falling away as momentum rises ($p = 1.0$ GeV/c), with modest discrimination recovered at higher momentum ($p = 10.0$ GeV/c).

In fact, the use of the same reference momenta ($p = 0.5, 1.0, 10.0$ GeV/c) in this case is somewhat misleading, as the lower masses of the particles concerned shift the three discrimination regimes (“low” momentum, minimum ionisation, “high” momentum) downward in momentum. At the fixed momenta $p = 0.5, 1.0$ GeV/c, K/ π discrimination is *less* powerful than discrimination between protons and kaons; but noting the difference in the geometric mean of masses in the two cases ($\sqrt{m_p m_K} = 0.681, \sqrt{m_K m_\pi} = 0.263$ GeV/c²) a more appropriate comparison may be made with kaon/pion distributions at 0.25 GeV/c (the lowest momentum for which reference spectra have been generated) and 0.35 GeV/c (*cf.* $\frac{0.263}{0.681} \times 1.0$ GeV/c ≈ 0.38 GeV/c). These distributions are shown in figure 5.25. The discrimination power is greater than for the proton/kaon distributions to which they approximately correspond, $p = 0.5, 1.0$ GeV/c respectively (*cf.* figure 5.21).

A consequence of this better intrinsic discrimination is that kaons and pions may be distinguished at high momentum, albeit with a loss of efficiency




Figure 5.23: Distributions of the kaon/pion likelihood ratio $\mathcal{L}_{K/\pi}$ for kaons (solid line) and pions (dashed line) at three momenta: (upper plot) $p = 0.5 \text{ GeV}/c$, (middle plot) $p = 1.0 \text{ GeV}/c$, and (lower plot) $p = 10 \text{ GeV}/c$. Compare with figures 5.18 and 5.21 above.




Figure 5.24: Expected pion rejection versus number of TRD planes for different kaon identification efficiencies for particles at $p = 0.5 \text{ GeV}/c$ (left), $p = 1.0 \text{ GeV}/c$ (centre) and $10.0 \text{ GeV}/c$ (right). Compare with figures 5.19 and 5.22 above.




Figure 5.25: Distribution of the likelihood ratio $\mathcal{L}_{K/\pi}$ for kaons (left) and pions (right) at $p = 0.25 \text{ GeV}/c$ (upper plot) and $p = 0.35 \text{ GeV}/c$ (lower plot) and $N_h = 9$. For comparison with proton/kaon likelihood distributions at $p = 0.5, 1.0 \text{ GeV}/c$ respectively, figure 5.21.

for kaons. So considering the $p = 10.0 \text{ GeV}/c$ case (figure 5.24, right), for a cut against pions of $\tilde{\epsilon}_\pi = \tilde{A}_\pi^{K/\pi}(\mathcal{L}_{K/\pi}; p, N_h) < 10^{-1} (10^{-2})$, the nominal kaon acceptance is 36% (8%) for a 9-hit track, falling to 23% (4%) for a 3-hit track. The proportion of kaons in a track sample could thus be enriched by a factor $\approx 3-8$ depending on the efficiency loss which could be sustained—not sufficient to pick “good kaons” from a random selection of tracks, but useful if dealing with a sample whose kaon content were already enhanced, or in a case where the discrimination may be applied more than once (as in the $\phi \rightarrow K^+K^-$ search described in section 5.3.9 below).

Results for kaon/muon discrimination are similar, with slightly better performance due to the larger mass difference ($m_K/m_\mu = 4.7$).

π/μ discrimination

Finally, considering discrimination between pions and muons, the two particles are so closely spaced in mass ($m_\pi/m_\mu = 1.32$) that no particle-by-particle discrimination is possible.

At low momenta ($p = 0.5, 1.0 \text{ GeV}/c$) the pion/muon likelihood ratio $\mathcal{L}_{\pi/\mu}$ is distributed almost identically for the two species (figure 5.26). Recalling




Figure 5.26: Distributions of the pion/muon likelihood ratio $\mathcal{L}_{\pi/\mu}$ for pions (solid line) and muons (dashed line) at (upper plot) $p = 0.5 \text{ GeV}/c$, and (lower plot) $p = 1.0 \text{ GeV}/c$, with a greatly expanded horizontal scale to show the small difference in the distributions. Compare with figures 5.18, 5.21 and 5.23 above.




Figure 5.27: Distributions of the pion/muon likelihood ratio $\mathcal{L}_{\pi/\mu}$ for pions (solid line) and muons (dashed line) at (upper plot) $p = 2.0 \text{ GeV}/c$, and (lower plot) $p = 10.0 \text{ GeV}/c$, with an expanded horizontal scale to show the difference in the distributions. Compare with the previous figure.

however that for protons and kaons, also very close in mass ($m_p/m_K = 1.9$), effective discrimination is available at low momentum, it is clear that these distributions are somewhat misleading. Momenta of $p \approx 0.5 \text{ GeV}/c$ in fact correspond to the minimum ionisation region for these light species, while the “low momentum” region of good discrimination is inaccessible in our experiment ($p < 0.25 \text{ GeV}/c$). Due to the same effect, however, momenta as low as $2.0 \text{ GeV}/c$ are on the relativistic rise, so that there is some distinction between $\mathcal{L}_{\pi/\mu}$ distributions there, as at higher momenta (figure 5.27). This will become important in section 5.3.8 below.

Summary

Based on the Monte Carlo simulation, then, we expect particle discrimination based on TRD information as follows:

- at low momenta ($p < 1 \text{ GeV}/c$), good discrimination is available in the p/π , p/K , K/π and K/μ cases (p/π being the most effective), with the discrimination power falling as a function of momentum in this region;
- where both particle-hypotheses are close to the minimum ionisation ($\beta\gamma \approx 3$), no discrimination is possible;
- in the high momentum region ($p \gtrsim 2 \text{ GeV}/c$) modest discrimination is available in the p/π , K/π and K/μ cases;
- pion/muon discrimination is too weak to be applied on a particle-by-particle basis at momenta accessible in the NOMAD experiment.

5.3.6 A test of p/π discrimination using Λ^0 decays

We now turn to tests of the particle discrimination on the NOMAD data, beginning with the p/π mode. Clearly it is necessary to obtain a tagged sample of the appropriate particle: in this case, protons. The most straightforward sample of this kind is that provided by Λ^0 decays $\Lambda^0 \rightarrow p\pi^-$. Along with the decays $K_S^0 \rightarrow \pi^+\pi^-$ and $\bar{\Lambda}^0 \rightarrow \bar{p}\pi^+$, and photon conversions $\gamma \rightarrow e^+e^-$, Λ^0 decays appear as an isolated vertex “downstream” of the primary vertex of the event, with two tracks of opposite sign emerging, and a total momentum vector aligned with the direction between the primary and the secondary vertex. These so-called V^0 vertices are treated as a special case by the event reconstruction code and may trivially be separated from other vertices for analysis purposes.

The Armenteros plot

It is convenient to discuss such a V^0 sample using the variables introduced by Armenteros et al. (1951):

$$\begin{aligned}\alpha &= \frac{p_{\parallel}^+ - p_{\parallel}^-}{p_{\parallel}^+ + p_{\parallel}^-} \\ &= \frac{p^+ \cos \theta^+ - p^- \cos \theta^-}{p^+ \cos \theta^+ + p^- \cos \theta^-} \\ p_{\perp} &= p^+ \sin \theta^+ \\ &= p^- \sin \theta^-\end{aligned}\tag{5.13}$$

where p^+ , p^- are the momenta of the positive and negative arms of the V^0 and θ^+ , θ^- the angles in space between those momenta and the total momentum of the V^0 , $\vec{p}^{V^0} = (\vec{p}^+ + \vec{p}^-)$.

The relationship between the variables (α, p_{\perp}) and quantities in the rest-frame of the V^0 is especially straightforward:

$$\begin{aligned}\alpha &= \frac{E^{*,+} - E^{*,-}}{M} + \frac{2p^* \cos \theta^*}{\beta M} \\ p_{\perp} &= p^* \sin \theta^*,\end{aligned}\tag{5.14}$$

where M and β are the V^0 mass and velocity, $E^{*,\pm}$ are the energies of the decay products in the V^0 rest frame, and p^* and θ^* the momentum and decay angle of the positive particle; note that energy and momentum conservation fix the values of $E^{*,+}$, $E^{*,-}$ and p^* for a given V^0 species, while momentum conservation prescribes $p^{*,+} = p^{*,-} = p^*$ and $\sin \theta^{*,+} = -\sin \theta^{*,-} = \sin \theta^*$.




Figure 5.28: The Armenteros plot (α, p_{\perp}) for V^0 in the 1995-1997 NOMAD data. Cuts $p_{\perp} > 0.030$ GeV/c (to suppress photon conversions) and $p_{\perp}^{V^0} < 0.100$ GeV/c (to suppress spurious V^0) have been applied.

That is, for a given V^0 species, and under the approximation $\beta \rightarrow 1$, events describe a half-ellipse in (α, p_{\perp}) , about a centre

$$\begin{aligned}\alpha_0 &= \frac{E^{*,+} - E^{*,-}}{M} \\ &= \frac{m_+^2 - m_-^2}{M^2} \\ p_{\perp,0} &= 0.\end{aligned}\tag{5.15}$$

For the K_S^0 , $m_+ = m_- = m_{\pi}$ and so $\alpha_0 = 0$: events are distributed in a broad ($2p^*/M = 0.83$) ellipse centred at the origin. The Λ^0 are offset by $\alpha_0 = (m_p^2 - m_{\pi}^2)/m_{\Lambda^0}^2 = 0.69$ while the ellipse is smaller ($2p^*/M = 0.18$). Both

distributions are clearly visible in the sample of V^0 vertices taken from the 1995-1997 NOMAD data, shown in figure 5.28. Also visible is the small population of $\bar{\Lambda}^0$ centred at $\alpha_0 = -0.69$. The extremely numerous photon conversions $\gamma \rightarrow e^+e^-$ have been suppressed by a cut requiring $p_{\perp} > 0.030 \text{ GeV}/c$ (although the tail of the photon conversion distribution is visible at the base of the plot $p_{\perp} \approx 0.030 \text{ GeV}/c$), and spurious V^0 have been suppressed by a cut requiring the component of the total V^0 momentum perpendicular to its presumed line-of-flight from the primary vertex, $p_{\perp}^{V^0} < 0.100 \text{ GeV}/c$.




Figure 5.29: Detail of the Armenteros plot (α, p_{\perp}) for V^0 in the 1995-1997 NOMAD data, showing the Λ^0 sample and the overlap region with the K_S^0 sample. Cf. figure 5.28.

It is apparent from the figure that the right half of the Λ^0 distribution and the lower-right corner of the K_S^0 distribution overlap. This is not fortuitous: it follows from the fact that in both cases the V^0 rest-frame momentum of

the decay products p^* is at least comparable to the lighter decay product mass (m_π). In such a case, for both $\beta \rightarrow 1$ and $\cos \theta^* \rightarrow 1$ we have

$$\begin{aligned} \alpha &\rightarrow \frac{E^{*,+} - E^{*,-} + 2p^*}{M} \\ &\approx \frac{E^{*,+} + E^{*,-}}{M} \\ &= 1, \end{aligned}$$

with the actual limiting values being 0.83 and 0.87 for K_S^0 and Λ^0 respectively. This behaviour is visible on the detail of the Armenteros plot shown as figure 5.29: note that V^0 with $\beta \lesssim 1$ extend to values $\alpha > \alpha_{lim}$ (see expression (5.14)).

The Λ^0 sample used for this study

It is desirable to remove the K_S^0 contamination from a Λ^0 sample in order to obtain a pure sample of protons. Removal of all of the overlap region results in a relatively small sample ($\approx \frac{1}{3}$ of the original) where low-momentum protons are enhanced (since energetic protons have a bias to high α), and so this approach is not generally suitable. Ideally one would perform fits to the hypotheses that the primary and V^0 are joined by a Λ^0 and K_S^0 respectively, with these particles constrained to have their true masses and with all momentum and position errors taken into account, and make a selection based on the respective χ^2 of the fits.

Unfortunately an error in the NOMAD data processing during 1998 has rendered the particle identification algorithm described in section 5.3.2 unusable on the standard data samples³ (the e/π variables described in section 5.2 may still be used). We are therefore restricted to a sample of Λ^0 decays filtered from the 1995 data, for which the full TRD information may be readily recovered. In addition to the cuts previously described,

- $p_\perp > 0.030$ GeV/ c against photon conversions, and
- $p_\perp^{V^0} < 0.100$ GeV/ c against spurious V^0 ,

the following cuts were applied (Schahmaneche, 1997):

- $p_\perp < 0.120$ GeV/ c against K_S^0 ;

³ A re-processing of the NOMAD data, with the bug responsible for this problem removed, is foreseen, but has not been carried out at the time of writing.

- the proper decay length $|\vec{x}_{prim} - \vec{x}_{V^0}| \times \frac{M^{V^0}}{p^{V^0}} < 47$ cm, corresponding to six times the mean proper decay length $c\tau$ of the Λ^0 , *i.e.* rejecting candidates which are implausibly long-lived;
- $|M^{V^0} - m_{\Lambda^0}| < 0.015 \text{ GeV}/c^2$.

The V^0 mass M^{V^0} is evaluated on a Λ^0 hypothesis, *i.e.* assigning proton and pion masses to the positive and negative arms respectively. The resulting sample is small (1091 events) and has not been subjected to any further cuts. Its distribution on the Armenteros plot is shown in figure 5.30.




Figure 5.30: The Armenteros plot (α, p_{\perp}) for V^0 in the 1995 data surviving the Λ^0 selection cuts described in the text.

Distributions for the positive arm of the V^0

Considering first the positive arm of the selected V^0 s, we expect a flat distribution in the proton acceptance variable $\tilde{\epsilon}_p = \tilde{A}_p^{p/\pi}(\mathcal{L}_{p/\pi}; p, N_h)$, and a concentration at low values of the pion acceptance variable $\tilde{\epsilon}_\pi = \tilde{A}_\pi^{p/\pi}(\mathcal{L}_{p/\pi}; p, N_h)$. Distributions are shown in figure 5.31 below, excluding any tracks with shared hits as in previous sections. The $\tilde{\epsilon}_p$ distribution is not flat but is well-populated in all bins and shows no systematic bias to high or low values; the $\tilde{\epsilon}_\pi$ distribution shows a peak at zero and falls away consistently towards unity. Following the earlier discussion, for a pure proton sample we expect the $\tilde{\epsilon}_\pi$ distribution to be a strong function of momentum: strongly peaked at zero for low momenta, showing little discrimination at $p \approx 2$ GeV/c, and showing a concentration at low values for high momenta. Plots exhibiting this behaviour are shown in figure 5.32 below for momentum bins $p < 1.5$, $1.5 < p < 3$ and $p > 3$ GeV/c; the events in the lower-momentum bin are mostly due to tracks with $p > 1$ GeV/c due to the relatively hard spectrum of the proton, which carries off most of the Λ^0 momentum.

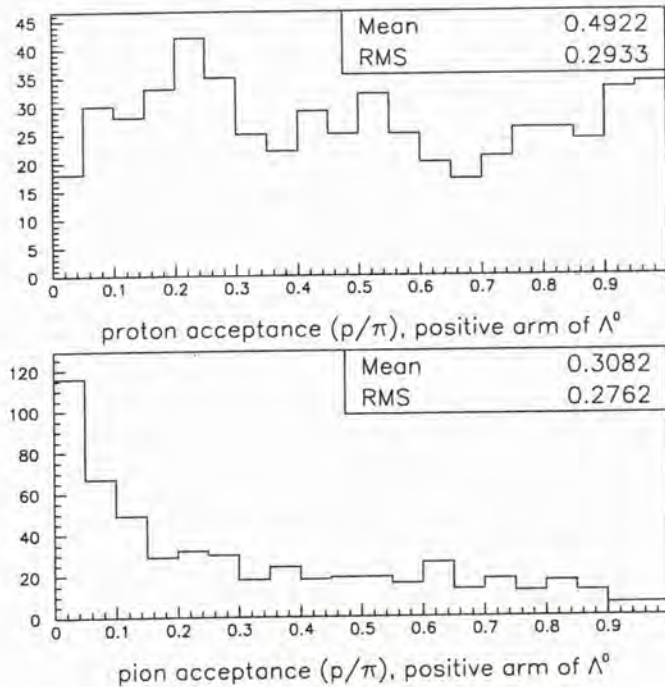


Figure 5.31: Distributions of proton/pion discrimination variables $\tilde{\epsilon}_p$ (upper plot) and $\tilde{\epsilon}_\pi$ (lower plot) for the positive arm of V^0 s passing Λ^0 cuts. Tracks with detected shared hits have been excluded.

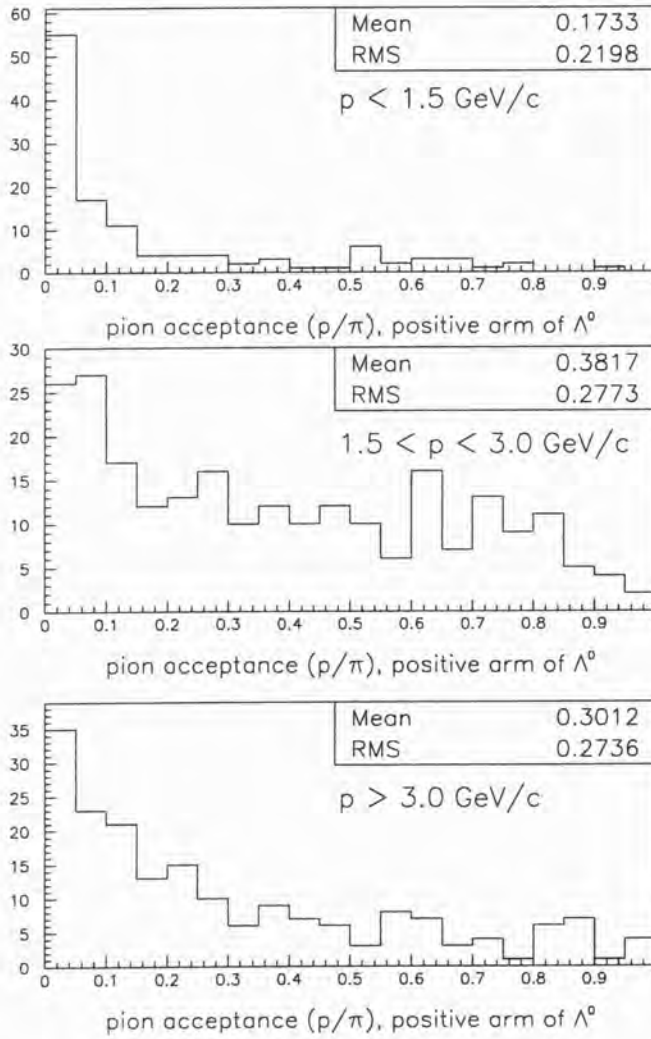


Figure 5.32: Distributions of the pion acceptance variable $\tilde{\epsilon}_\pi$ for the positive arm of the Λ^0 sample in three momentum bins: (upper plot) $p < 1.5 \text{ GeV}/c$, (middle plot) $1.5 < p < 3 \text{ GeV}/c$, (lower plot) $p > 3 \text{ GeV}/c$.

As noted above, the Λ^0 sample is known to contain some contamination due to K_S^0 . Choosing the momentum range $p > 3 \text{ GeV}/c$, where we expect some discrimination and have modest statistics, this contamination should be visible. In figure 5.33 below, Λ^0 s with $p^+ > 3 \text{ GeV}/c$ are shown on the Armenteros plot, with cases where the positive track passes cuts $\tilde{\epsilon}_\pi > 0.5$ and $\tilde{\epsilon}_\pi > 0.7$ singled out. The concentration at the top and along the right of the plot—the region of overlap with the K_S^0 distribution—is clear.




Figure 5.33: Armenteros plot of V^0 s passing Λ^0 cuts with $p^+ > 3 \text{ GeV}/c$, and pion rejection cuts on the positive track. V^0 s where the positive track has shared hits in the TRD have been suppressed. Cases where the positive track has $\tilde{\epsilon}_\pi > 0.5$ are shown as circles, with $\tilde{\epsilon}_\pi > 0.7$ shown as filled circles. These are clearly concentrated in the region of overlap with the $K_S^0 \rightarrow \pi^+\pi^-$ distribution (cf. figure 5.29).

Distributions for the negative arm of the V^0

Turning to the negative arm of the V^0 s, we expect very good discrimination due to the low momentum of the tracks, typically $p < 1 \text{ GeV}/c$. The distribution of the proton acceptance variable $\tilde{\epsilon}_p$ is shown in figure 5.34 below for the negative arm, with the distribution for the positive arm for the few cases in the same momentum regime ($p^+ < 1.5 \text{ GeV}/c$) included for comparison. Tracks are concentrated in an abrupt peak at $\tilde{\epsilon}_p > 0.99$ for the negatives (π s), *cf.* the distribution across all bins for the positives (protons). In fact, the low straggle of negative tracks with $\tilde{\epsilon}_p < 0.90$ is almost entirely due to the tracks with $p > 1.0 \text{ GeV}/c$, the momentum region where discrimination is deteriorating. Applying a cut to remove these tracks (figure 5.35), we find the pions in a peak at $\tilde{\epsilon}_p > 0.9999$. In a track sample with a similar momentum spectrum, a very pure pion sample could be obtained by a cut in $\epsilon_p = 1 - \tilde{\epsilon}_p$ of $\epsilon_p < 10^{-3}$ —imposing a rejection factor of 1000 on protons—with

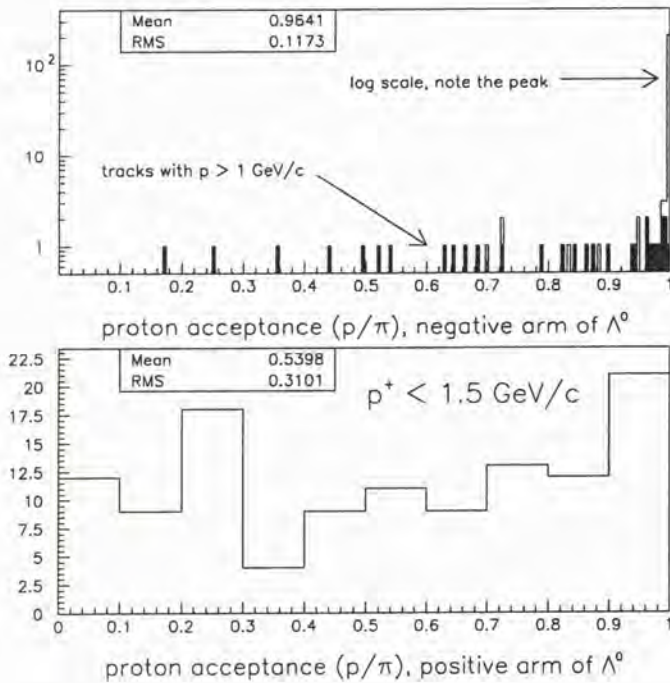


Figure 5.34: Proton acceptance variable $\tilde{\epsilon}_p$ for the negative arm of the Λ^0 s (upper plot, note the log scale), compared to the positive arm (protons) for $p^+ < 1.5 \text{ GeV}/c$ (lower plot, linear scale). Tracks above $1 \text{ GeV}/c$, where we expect poorer discrimination, have been emphasized in the upper plot.

95% efficiency for pions.




Figure 5.35: Proton acceptance $\tilde{\epsilon}_p$ for the negative arm of the Λ^0 s for $p^- < 1.0 \text{ GeV}/c$, on an expanded scale. Note the small number of tracks falling below $\tilde{\epsilon}_p = 0.99$ (the underflow bin).

The more likely physics interest in this momentum range will be to select protons. Finally, then, the distribution of $\tilde{\epsilon}_\pi$ is shown in figure 5.36 below for the negative arm, with the distribution for the positive arm in cases $p^+ < 1.0 \text{ GeV}/c$ only included for comparison. For the negative arm, the population of all bins is clear, as is the near-flat distribution (see the discussion in section 5.3.3). By comparison, the $\tilde{\epsilon}_\pi$ distribution for the positive arm is strongly concentrated in a spike at values below 0.005. For a representative sample of protons below $1.0 \text{ GeV}/c$, the acceptance for any given cut (say, $\tilde{\epsilon}_\pi < 0.005$) will greatly exceed that implied by this graph, as the hard momentum spectrum of protons from Λ^0 (see figure 5.37) has constrained us to use many tracks in the region $0.8 < p < 1.0 \text{ GeV}/c$ where discrimination is deteriorating.




Figure 5.36: Pion acceptance $\tilde{\epsilon}_\pi$ for the negative arm of the Λ^0 s (upper plot), compared to the positive arm for $p^+ < 1.0 \text{ GeV}/c$ (lower plot; note the expanded scale and the overflow bin.)




Figure 5.37: Momentum spectrum for both negative and positive arms of the Λ^0 below $1.5 \text{ GeV}/c$ used in the previous plots. Proton/pion discrimination falls off rapidly above $0.8 \text{ GeV}/c$ (see the text); the tracks used for the most important comparisons, with $p < 1.0 \text{ GeV}/c$, are shown hatched.

5.3.7 A test of K/π and K/μ discrimination using the calorimeters

A sample of isolated tracks

A tagged sample of kaons, to test the behaviour of the K/π (K/μ) discrimination variables on the data, is not readily available⁴. Furthermore the particle identification algorithm described in section 5.3.2 may not be used on the standard data samples presently available (see section 5.3.6 and note 3 on page 222). As an expedient we work with a sample of tracks isolated in the calorimeters, filtered from the 1995 data, for which the particle identification algorithm may be used (Hurst, 1997a). Only events with twenty tracks or less are considered, and the following cuts are applied to each track:

- the track extrapolation to the front face of the ECAL (section 2.4.8) lies within the restricted fiducial volume $|x| < 135$ cm, $|y| < 135$ cm;
- the track extrapolation to the front face of the HCAL lies within $|x| < 170$ cm, and intersects one of the 18 HCAL counters (*i.e.* avoids the gaps between counters and the gap in the central region; see section 2.4.9);
- of all the other tracks in the event (other than electrons⁵), the closest track-extrapolation to the HCAL must be more than three counters away.

These tracks have isolated, well-defined energy deposition in the calorimeters which may be used to facilitate the separation of a kaon-rich sample.

⁴The most straightforward source of kaon-enriched tracks is the sample of “leakage tracks” from neutrino interactions in the front calorimeter (section 2.4.3) where two muons $\mu^- \mu^+$ are produced, the so-called opposite-sign dimuon events. The leading muon (typically μ^-) is produced in the deep inelastic scattering $\nu_\mu + N \rightarrow \mu^- + X$ while many of the secondary muons (μ^+) are due to the decay of charm quarks $c \rightarrow s \bar{\nu}_\mu \mu^+$ produced in the scattering. Energetic negative hadrons produced in such events will be strongly enriched in K^- due to the strange quark produced in the charm decay.

Based on the results of a preliminary study of the 1995 FCAL data, processing of the full 1995-1998 FCAL data would be required to secure a suitable sample (Boyd, 1997). The first round of processing for this data is still underway at the time of writing and so the corresponding analysis has not yet been carried out.

⁵The reason for the restriction is that this sample was originally produced to study energy deposition by isolated hadrons in the HCAL (Hurst, 1997a). Electrons deposit all their energy in the ECAL with insignificant leakage of the electromagnetic shower into the HCAL (section 2.4.8), whence the special treatment at this point.

The distribution of acceptance variables

To consider the distribution of K/π discrimination variables we first prepare a kaon- and pion-dominated sub-sample by suppressing muons and electrons. To suppress muons we require that tracks, when extrapolated through the calorimeters, must have a probability to both reach the muon chambers (*i.e.* not range out in the lead-glass, iron, support structures *etc.*) and be within their geometric acceptance, of at least 99%; and *not* be matched to a muon-chamber track (see section 2.4.10). Furthermore, another track in the same event must be matched to a muon-chamber track. To suppress electrons we perform a cut $\tilde{\epsilon}_e = \tilde{A}_e(\mathcal{L}; p, N_h) < 0.01$. After these cuts, the kaon acceptance $\tilde{\epsilon}_K = \tilde{A}_K^{K/\pi}(\mathcal{L}_{K/\pi}; p, N_h)$ and the pion acceptance $\tilde{\epsilon}_\pi = \tilde{A}_\pi^{K/\pi}(\mathcal{L}_{K/\pi}; p, N_h)$ are distributed as shown in figure 5.38. Both distributions are consistent with the expected behaviour if most of the tracks are pions (concentrated at large $\tilde{\epsilon}_K$, and flat in $\tilde{\epsilon}_\pi$), and there is a small admixture of kaons (flat in $\tilde{\epsilon}_K$, and concentrated at small $\tilde{\epsilon}_\pi$).

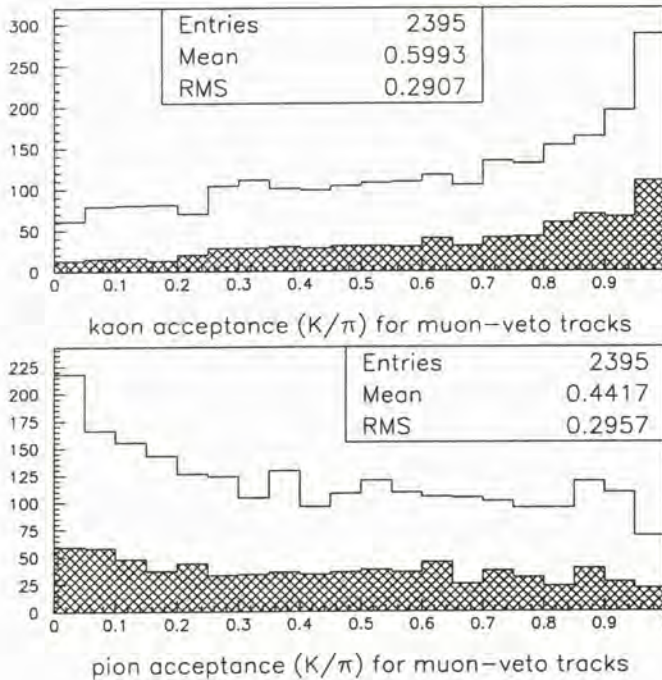


Figure 5.38: Kaon/pion discrimination variables $\tilde{\epsilon}_K$ (upper plot) and $\tilde{\epsilon}_\pi$ (lower plot) for tracks isolated in the HCAL, with both muons and electrons suppressed (see the text). In both plots the subset of tracks with $p > 7 \text{ GeV}/c$ is shown cross-hatched.

Regarding K/μ discrimination variables it is trivial to obtain a sub-sample of tracks dominated by muons. After requiring that tracks be matched to a muon-chamber track (section 2.4.10), the kaon acceptance $\tilde{\epsilon}_K^{K/\mu} = \tilde{A}_K^{K/\mu}(\mathcal{L}_{K/\mu}; p, N_h)$ and the muon acceptance $\tilde{\epsilon}_\mu^{K/\mu} = \tilde{A}_\mu^{K/\mu}(\mathcal{L}_{K/\mu}; p, N_h)$ are distributed as shown in figure 5.39. As expected for a sample of muons, the distribution of $\tilde{\epsilon}_K^{K/\mu}$ is concentrated at high values (since the muon is the lighter-mass particle; see the discussion of section 5.3.2) while the distribution of $\tilde{\epsilon}_\mu^{K/\mu}$ is flat.




Figure 5.39: Kaon/muon discrimination variables $\tilde{\epsilon}_K^{K/\mu}$ (upper plot) and $\tilde{\epsilon}_\mu^{K/\mu}$ (lower plot) for tracks isolated in the HCAL and matched as muons.

A test of kaon/muon discrimination

The muon-chamber system allows us to “tag” particles as muons and hence to evaluate how well they are identified by other means: in this case, the TRD-based K/μ discrimination variables. In order to make use of the muon-chamber information we take the sample of isolated tracks and consider the sub-sample in the small momentum bin $5 < p < 7 \text{ GeV}/c$ (to suppress any momentum-dependent effects) with a probability to both reach and hit the

chambers of the first muon station > 0.999 : we may then use the presence or absence of a match to a muon-chamber track to classify the particle as a muon or K/π respectively. In order to suppress the otherwise-dominant pions in this second group, we use information from the electromagnetic calorimeter.

A cut requiring the energy of the ECAL cluster associated to the track to be less than 0.8 GeV selects tracks with an energy deposition consistent with a single “minimum ionising” track in the calorimeter. This cut efficiently selects muons, as the mean deposition by a muon is 0.566 GeV with a limited tail to higher values (see Altegoer et al., 1998a, figure 17); it also admits those pions and kaons which do not undergo a nuclear interaction in the lead-glass. This last group will be enriched in kaons, as their interaction length is longer than that of the π^\pm due to their strange-quark content, as shown in figure 5.40. (See also the interaction cross sections recorded by Caso et al., 1998, as figures 38.20 through 38.22)


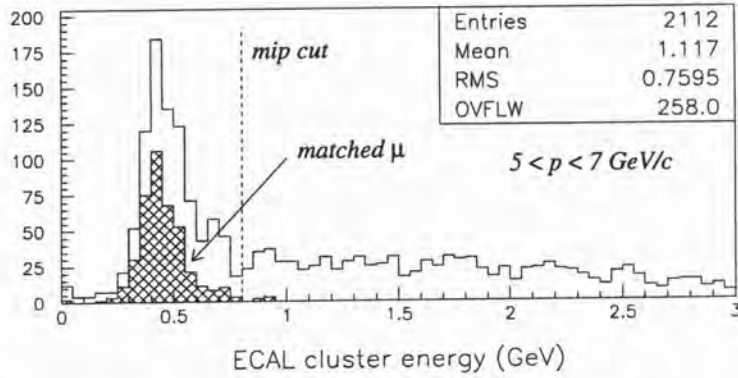



Figure 5.40: Interaction length for various hadrons as a function of momentum. Note the significantly longer interaction lengths for kaons at low momentum. (Taken from Vilain et al., 1999, figure 10).

A cut in the muon acceptance variable $\tilde{\epsilon}_\mu^{K/\mu}$ then successfully discriminates between the two groups (see figure 5.41): tracks passing a $\tilde{\epsilon}_\mu^{K/\mu} < 0.2$ cut are dominated 2:1 by non-muons, while tracks with $\tilde{\epsilon}_\mu^{K/\mu} > 0.2$ are evenly split between muons and non-muons. (The remaining contamination of pions works against the visibility of the effect, since the pion is much closer to the muon than to the kaon mass—the $\tilde{\epsilon}_\mu^{K/\mu} < 0.2$ cut will suppress pions in the non-muon sample, reducing its numbers.) This is consistent with the variable distinguishing between kaons and muons as expected.



(a) A minimum ionising particle (mip) cut $E_{\text{ECAL}} < 0.8 \text{ GeV}$ selecting muons and non-interacting hadrons (see the text).



(b) Tracks passing the mip cut, divided into those *not matched* to muon chamber tracks (mostly hadrons), and those *matched* (μ).

Figure 5.41: Discrimination between kaons and muons in the data using a cut in the muon acceptance $\tilde{\epsilon}_{\mu}^{K/\mu}$.

5.3.8 A test of the π/μ discrimination variables

The distribution of acceptance variables

The track sample introduced in the previous section may also be used to test the π/μ discrimination variables. As discussed in section 5.3.5, the anticipated discrimination in this mode is too weak to separate out a high purity sample of either pions or muons based on the TRD information alone: however, using the muon chambers to select both muon-dominated and muon-free subsamples, we may check the distributions of the π/μ discrimination variables for their expected behaviour. As previously, we will suppress electrons by applying a cut $\tilde{\epsilon}_e < 10^{-2}$, and we will exclude any tracks having shared hits in the TRD region.

Figure 5.42 shows distributions for the muon acceptance variable $\tilde{\epsilon}_\mu^{\pi/\mu} = \tilde{A}_\mu^{\pi/\mu}(\mathcal{L}_{p/\pi}; p, N_h)$ for muon and non-muon track samples. For muons the distribution is flat, as expected (upper plot). For tracks passing the muon veto




Figure 5.42: Distribution of the pion/muon discrimination variable $\tilde{\epsilon}_\mu^{\pi/\mu}$ for (upper plot) matched muons, and (lower plot) negative tracks passing a muon veto. The track sample and cuts are the same as those used in the previous section (provided by Hurst, 1997a).

conditions—another track in the same event has been identified as a muon; this track has a probability to both reach and hit the muon chambers greater than 0.999, but is not matched to a muon-chamber track—we expect a falling distribution (since $m_\pi > m_\mu$; see section 5.3.2), as shown in the lower plot. (Negative tracks only have been shown, to suppress any contamination due to protons.) Note that the distribution is well-populated from zero to one, and shows only a slight accumulation at low values, so any cut in $\tilde{\epsilon}_\mu^{\pi/\mu}$ to suppress muons would be almost as expensive in efficiency for pions. This is in line with the assessment of section 5.3.5 that particle-by-particle discrimination in this case will be impractical.

The distribution in this last case is complicated by kaon tracks (and by protons, if positive tracks are plotted), which unlike electrons cannot be easily removed. A more straightforward comparison is provided by the inverse case, where we plot the distribution of pion acceptance for muons, where a pure sample may readily be obtained. Figure 5.43, then, shows distributions of the pion acceptance variable $\tilde{\epsilon}_\pi^{\pi/\mu} = \tilde{A}_\pi^{\pi/\mu}(\mathcal{L}_{p/\pi}; p, N_h)$ for both muon and non-muon track samples. The muon sample should show a bias to large values of $\tilde{\epsilon}_\pi^{\pi/\mu}$: in the upper plot this is unmistakable. The same distribution is also shown for the restricted momentum bin $6 < p < 10 \text{ GeV}/c$ (inset), to suppress the effect of the broad momentum spectrum for muons: the slope is reduced but still clear. The sample passing the muon-veto conditions, dominated by pions, yields a near-flat histogram as expected, with some distortion towards lower values consistent with a contribution due to heavier particles (K^-). Note that this slope is opposite in sense to that for muons, so that the distinction between the two cases is clear.

An application to muon-veto studies

While particle-by-particle discrimination will not be possible, we may exploit the the difference in the *distribution* of $\tilde{\epsilon}_\pi^{\pi/\mu}$ between a sample of pions and a sample of muons, to test for the presence of muons in a set of tracks. The most obvious application is a test of “muon veto” algorithms in the $p \approx 2 \text{ GeV}/c$ region: muons below about $2.5 \text{ GeV}/c$ do not penetrate to the muon chambers, so for tracks in this momentum range it is not possible to carry out muon identification by conventional means (section 2.4.10); events $\nu_\mu + N \rightarrow \mu^- + X$ with unidentified muons at this momentum nonetheless form an important background to tau hadronic decay searches, in particular the three-prong decay search (section 2.3.3; see also the further discussion in chapter 6). Since these muons must be suppressed by kinematic methods, an independent test of the muon contamination of the resulting samples is desirable.




Figure 5.43: Distribution of the pion/muon discrimination variable $\tilde{\epsilon}_\pi^{\pi/\mu}$ for (upper plot) identified muons, with the subset of muons with $6 < p < 10$ GeV/c inset; and (lower plot) for negative tracks with muons and electrons suppressed as discussed in the text. The parameters of a straight-line fit $\delta N(\tilde{\epsilon}_\pi^{\pi/\mu}) = A0 + A1 \times \tilde{\epsilon}_\pi^{\pi/\mu}$ are also shown: there is a marginal slope with the opposite sense, possibly due to kaons (see the text).

An example demonstrating that this is possible in-principle is shown in figure 5.44: a histogram of 10,000 9-hit tracks at 2.0 GeV/c, 20% of which are muons and the remainder pions, generated using an extension of the Monte-Carlo code of Fazio et al. (1994). The slope of the distribution in this case is readily detected. The number of tracks N required to detect a fractional contamination f goes approximately as $N \propto \frac{1}{f^2}$, so the technique is not capable of detecting subtle leakages of a muon-veto algorithm, only fairly gross failures. Table 5.1 shows the required N for a number of contaminations f at various levels of confidence; the rough argument and formula used to derive the numbers follows below.

The unavailability of the particle identification algorithm on the principal NOMAD data samples, due to the bug noted in section 5.3.6 above, means that a test of this procedure on the muon-veto algorithm presented in this thesis (section 6.7) may not at present be carried out.




Figure 5.44: Simulation of a sample of 10,000 tracks at 2.0 GeV/c, 80% π , 20% μ ; the variable $\tilde{\epsilon}_\pi^{\pi/\mu}$ is plotted, together with the straight-line fit $\delta N(\tilde{\epsilon}_\pi^{\pi/\mu}) = A0 + A1 \times \tilde{\epsilon}_\pi^{\pi/\mu}$. The slope due to the presence of muons is apparent (*cf.* the fit to a flat histogram yields a χ^2/ndf of 46.1/19).

fractional contamination of muons f	number of tracks N required for contamination to be “visible”			
	at 2σ	at 3σ	at 4σ	at 5σ
0.10	16900	38000	68000	105600
0.20	4200	9500	16800	26300
0.50	670	1510	2700	4200
0.80	260	590	1040	1630
0.90	210	260	820	1280

Table 5.1: For a sample of tracks of the size N shown, a fractional contamination of muons f should be visible from the slope of the $\tilde{\epsilon}_\pi^{\pi/\mu}$ distribution, at the “level” ($2\sigma, 3\sigma$, etc.) shown.

Estimating the visibility of muon contamination

We can model the task of finding a slope in the distribution as finding a difference in the populations of two wide bins, $0 \leq \tilde{\epsilon}_\pi^{\pi/\mu} < 0.5$ and $0.5 \leq \tilde{\epsilon}_\pi^{\pi/\mu} \leq 1.0$, which we will call N_1 and N_2 . The difference between the two numbers will be significant when

$$\frac{N_2 - N_1}{\sigma(N_2 - N_1)} < k \tag{5.16}$$

where $k = 3$, for example. We take $N = N_1 + N_2$ to be fixed and the probability of a track falling in the first bin to be p ; following the binomial distribution, N_1 will have mean value pN and variance $p(1 - p)N$. The difference $(N_2 - N_1) = N - 2N_1$ will then have mean value $(1 - 2p)N$ and variance $= (-2)^2 \text{Var}(N_1) = 4p(1 - p)N$. Thus

$$\begin{aligned} \frac{N_2 - N_1}{\sigma(N_2 - N_1)} &= \frac{(1 - 2p)N}{\sqrt{(1 - (1 - 2p)^2)N}} \\ &= \sqrt{\frac{N}{\left(\frac{1}{1-2p}\right)^2 - 1}}, \end{aligned} \tag{5.17}$$

and substitution into (5.16) then gives

$$N > k^2 \left(\frac{1}{(1 - 2p)^2} - 1 \right) \tag{5.18}$$

for the difference $(N_2 - N_1)$ to be visible at $k\sigma$. Now we suppose that the sample of N tracks consists of a fractional contamination f of muons, each

with probability q to fall in the first bin; pions will fall in the first bin with probability 0.5. Then

$$\begin{aligned} p &= fq + (1 - f) \times 0.5 \\ &= 0.5 + f(q - 0.5) \end{aligned} \tag{5.19}$$

and so finally

$$N > k^2 \left(\frac{1}{f^2(1 - 2q)^2} - 1 \right). \tag{5.20}$$

The values of table 5.1 were generated using this formula, and the value $q = 0.4231$, taken from the same program used to generate figure 5.44.

5.3.9 A search for $\phi \rightarrow K^+K^-$ decays

Finally we turn to the first significant application of the particle discrimination techniques developed in this chapter: a search for decays of resonances. The most straightforward example is the decay of the vector meson $\phi(1020) \rightarrow K^+K^-$, which in the absence of particle identification is completely swamped by the combinatoric background; since both of the decay products are kaons we may apply the K/π discrimination variables *twice*, suppressing the background sufficiently to make the signal visible. The analysis of the data is at a preliminary stage, but sufficient to demonstrate the usefulness of the technique.

Data selection

The data samples used in this section were prepared in 1997 from a special processing of the 1995 and 1996 data. The TRD particle identification information was available at that time, but is not available on the present standard samples (*cf.* section 5.3.6 and note 3 on page 222), so it has not yet become possible to extend the analysis to the 1997 and 1998 data.

After performing standard quality cuts (see chapter 6) and requiring the presence of a muon at the primary vertex (restricting ourselves to the majority $\nu_\mu + N \rightarrow \mu^- + X$ events) we form all combinations of one positive with one negative track from the primary vertex, excluding

- the identified muon, and any track flagged as having a possible (albeit lower probability) match to the muon-chamber track;
- electrons e^\pm , imposing a rejection factor of 10 by a weak cut $\tilde{\epsilon}_e < 0.10$;
- any track without TRD information;
- any track having shared hits in the TRD.

We then divide the data into two subsamples according to momentum:

1. $p^{tot} = |\vec{p}^+ + \vec{p}^-| > 5.0 \text{ GeV}/c$, where we expect modest K/π discrimination, and an initial enrichment in this as kaons from $\phi \rightarrow K^+K^-$ have a harder momentum spectrum in the Monte Carlo than do random tracks;
2. $p^{tot} < 2.0 \text{ GeV}/c$, where the Monte Carlo leads us to expect negligible $\phi \rightarrow K^+K^-$, but we have the advantage of powerful K/π discrimination.

The mass spectrum $M(h^+h^-)$, assigning the kaon mass to each track, is shown for these subsamples in figure 5.45. In both cases the spectrum is peaked at low values; the tail to high values, apparent for the high-momentum sample, is suppressed in the low-momentum sample due to the restricted phase space. Note that $\phi \rightarrow K^+K^-$ will appear as a narrow peak at $M \approx 1.019 \text{ GeV}/c^2$, and in both cases the combinatoric background in this region is large.

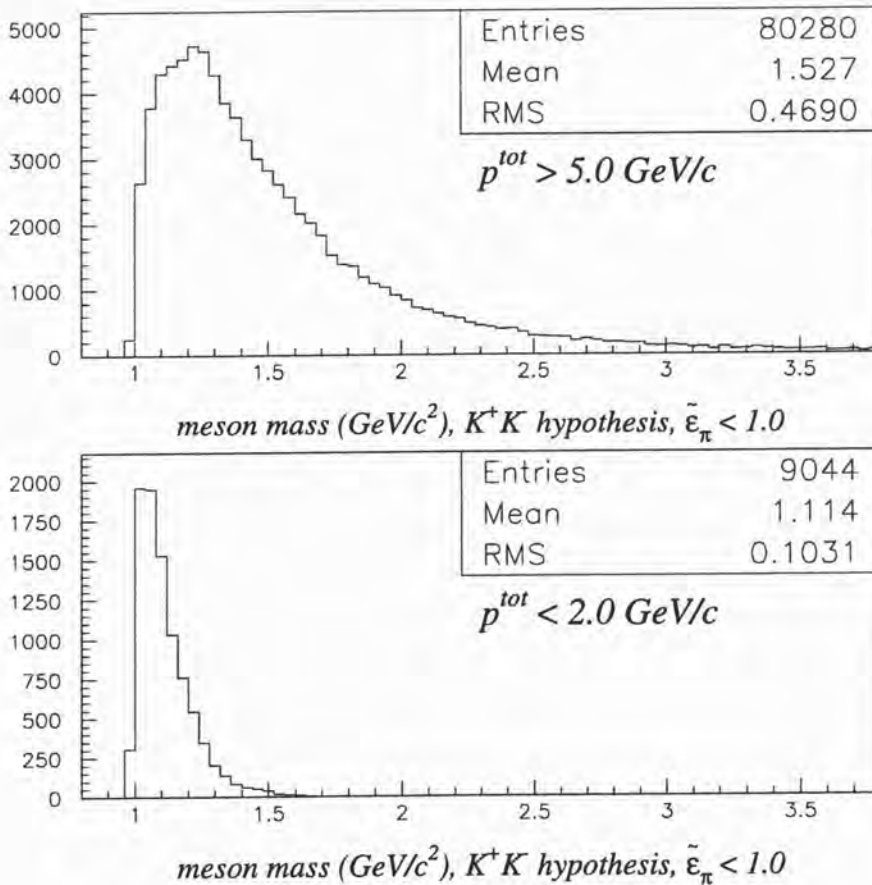


Figure 5.45: The two-particle invariant mass $M(h^+h^-)$, using reconstructed momenta and ascribing the charged kaon mass to each particle, for the high-momentum (upper plot) and low-momentum (lower plot) samples described in the text.

We will now apply K/π discrimination cuts to each of these samples in turn.

The high-momentum sample

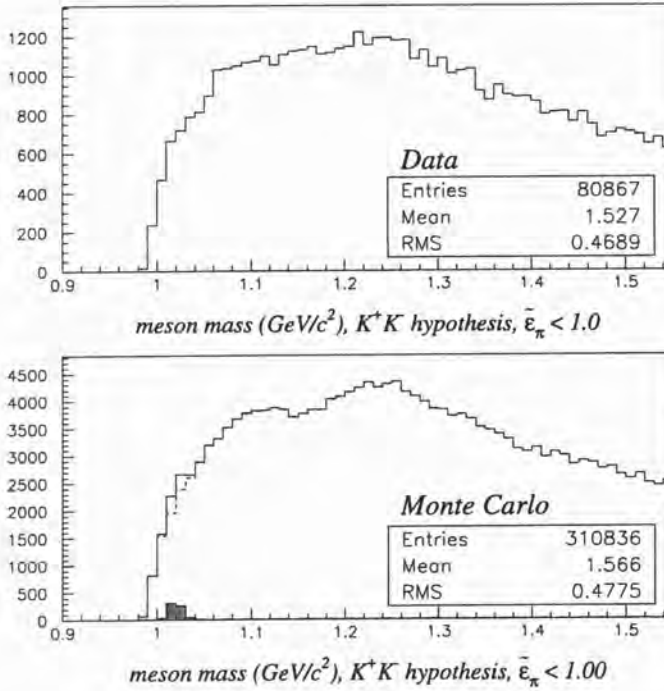


Figure 5.46: The low-mass region for the high-momentum ($p^{tot} > 5.0 \text{ GeV}/c$) sample in the data (upper plot) and Monte Carlo (lower plot), before K/ π discrimination cuts. The statistics apply to the entire distribution. In the Monte Carlo plot, candidates which are K^+K^- pairs from the decay of a single meson are shown both as the gap between the solid and dashed histogram, and as a shaded contribution at the base of the plot: note the small ϕ sample on the rising edge of the combinatoric background.

For the high-momentum sample, shown for both data and Monte Carlo in figure 5.46, the K/ π discrimination is modest and we adopt the simple procedure of applying cuts on the pion acceptance $\tilde{\epsilon}_\pi = \tilde{A}_\pi^{K/\pi}(\mathcal{L}_{K/\pi}; p, N_h)$ for each track. Figure 5.47 shows the effect on the data of increasingly strict cuts $\tilde{\epsilon}_\pi < 0.50, 0.20, 0.15$ and 0.10 ; these correspond to rejection factors against $\pi^+\pi^-$ combinations, assuming a flat distribution of $\tilde{\epsilon}_\pi$ for pions, of 4, 25, 44 and 100 respectively. With each cut the total population of the distribution is reduced by factors 3.6, 17, 27 and 52 respectively, *i.e.* with an increasing gap between the nominal rejection factor and the actual reduction in population as the $\tilde{\epsilon}_\pi$ becomes more tight, consistent with the expected behaviour given

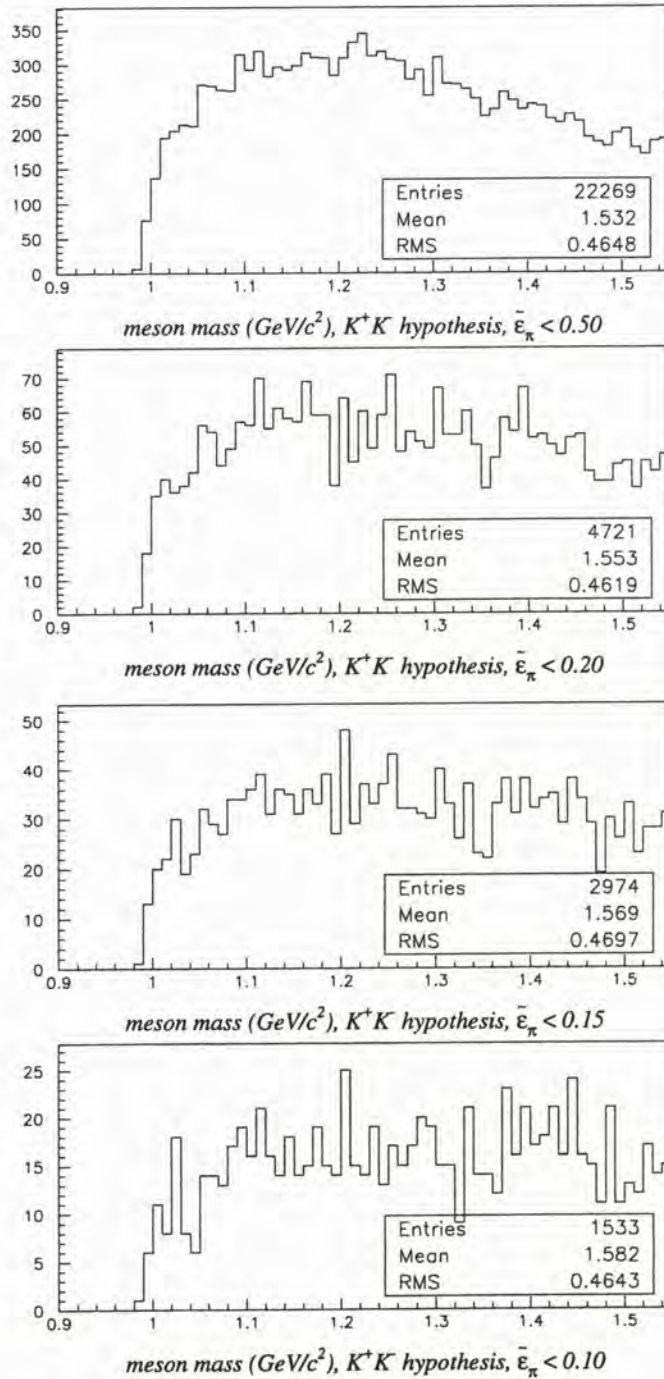


Figure 5.47: The low-mass region for the high-momentum ($p^{\text{tot}} > 5.0 \text{ GeV}/c$) sample as progressively tighter cuts in the K/π discrimination variable $\bar{\epsilon}_\pi$ are applied to both positive and negative tracks: from top to bottom, $\bar{\epsilon}_\pi < 0.50, 0.20, 0.15, 0.10$ respectively. A small peak at the $\phi(1020)$ mass emerges from the background as tighter cuts are applied.

an admixture of kaons in the track sample. The shape of the distribution, as reflected in its mean and first moment and also the visual impression, is largely unchanged by the cuts, except for some flattening of the profile in the region above threshold ($M(h^+h^-) = 1.1 \approx 1.3 \text{ GeV}/c^2$).

A peak, consistent with a contribution due to $\phi \rightarrow K^+K^-$ decay, emerges at the ϕ mass $1.019 \text{ GeV}/c^2$ (immediately above threshold) as the cuts in $\tilde{\epsilon}_\pi$ tighten. The identification is marginal with these statistics and depends on our knowledge of the ϕ mass: the Poisson noise on the distribution is substantial in the latter plots where the total population is low, as we have been forced to choose narrow bins by the small natural width of the ϕ ($\Gamma = 0.0044 \text{ GeV}/c^2$); any unconstrained search for peaks would be impractical.

Turning to the Monte Carlo, we begin with a population larger than the data by a factor of 3.8 after the cuts to define the sample are applied, but before any K/π discrimination cuts have been imposed (figure 5.46). The contribution from inclusive meson decays producing K^+K^- is also shown: there is a small contribution at the ϕ mass. As cuts in the pion acceptance $\tilde{\epsilon}_\pi$ are imposed (figure 5.48) the total distribution is reduced by factors 3.2, 14, 23, 43 respectively, consistent with a slightly larger proportion of K^\pm in the Monte Carlo than the data in this region. The contribution due to ϕ s remains, but is too small for ready identification even under cuts.

(The enhancement of ϕ s by cuts in the Monte Carlo is only modest, whereas the peak in the data becomes progressively more distinct as cuts tighten in the data; the median momentum of the ϕ s is $\approx 11 \text{ GeV}/c$ in the Monte Carlo while the meson candidates producing the peak in the tightest-cut plots for the data (figure 5.47) appear to be all at this momentum or greater. The efficiency of a fixed $\tilde{\epsilon}_\pi$ cut for kaons increases with momentum in this region.)

In summary, then, the behaviour of the high-momentum data sample under K/π discrimination cuts is consistent with the presence of ϕ s at a larger level than that implied in the Monte Carlo. With the next processing of the NOMAD data it will be possible to confirm this interpretation with greatly increased statistics: the addition of the 1997 and 1998 runs will increase the number of events by a factor of 2.5, while the full treatment of shared hits (section 5.3.4) will further increase the available number of combinations h^+h^- by a factor of two. Any estimate of the *rate* of ϕ production for $p(\phi) > 5 \text{ GeV}/c$ will require a detailed study on the Monte Carlo (clearly using a larger sample than that utilised here) to determine efficiencies *etc.*, noting that these will be a function of momentum.

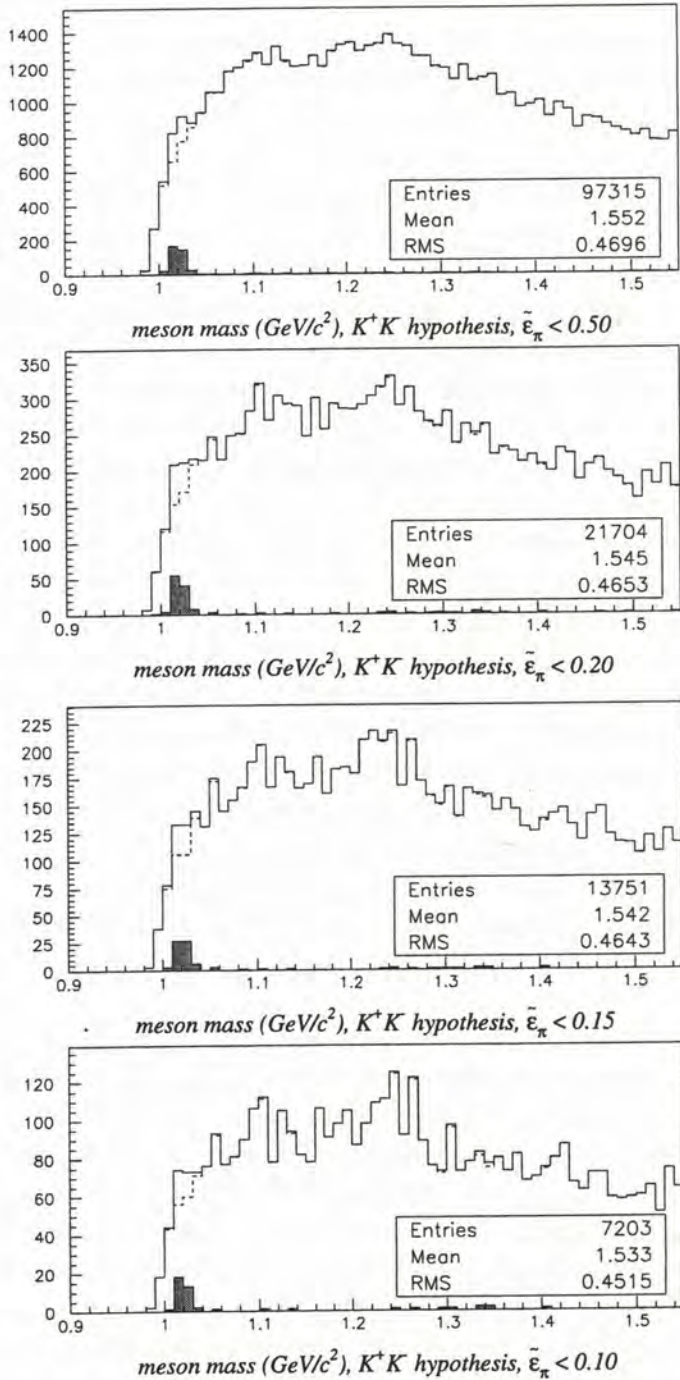


Figure 5.48: The mass $M(h^+h^-)$ for the high-momentum ($p^{tot} > 5.0$ GeV/c) sample in the Monte Carlo, under progressively tighter cuts in the K/π discrimination variable $\tilde{\epsilon}_\pi$. K^+K^- decays are shown both as the gap between the solid and dashed histogram, and as the shaded contribution, as previously: the small ϕ sample remains obscure beneath the fluctuation in bin contents; note also that there is no apparent phase-space effect producing a “false peak” at the ϕ mass. Cf. the same distributions for data, figure 5.47.

The low-momentum sample

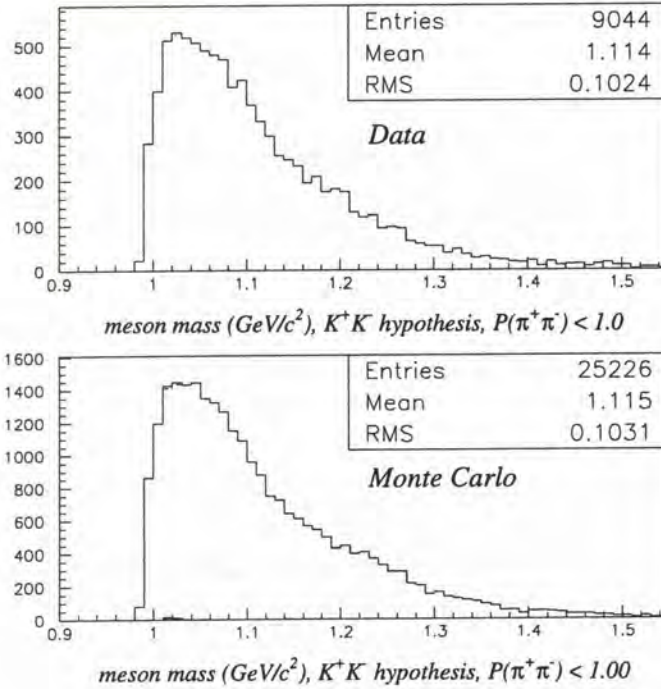


Figure 5.49: The low-mass region for the low-momentum ($p^{tot} < 2.0 \text{ GeV}/c$) sample in the data (upper plot) and Monte Carlo (lower plot), before K/ π discrimination cuts. In the Monte Carlo plot, candidates which are K^+K^- pairs from the decay of a single meson are shown shaded at the base of the plot: note that these cases are a negligible fraction of the total.

For the low momentum ($p^{tot} < 2.0 \text{ GeV}/c$) sample, shown in figure 5.49 for data and Monte Carlo, some care must be taken in combining the pion acceptance values $\tilde{\epsilon}_\pi(h^+)$ and $\tilde{\epsilon}_\pi(h^-)$ for the positive and negative track respectively. For the high-momentum sample we used a “box cut”, $\tilde{\epsilon}_\pi(h^+) < k$, $\tilde{\epsilon}_\pi(h^-) < k$; in the low-momentum region the discrimination between kaons and pions is strong (figure 5.50) and we may take better advantage of the information present in the two values by combining them using the method of Fisher (1932), equivalent in this case to

$$\mathcal{P} = Q(1 - \ln(Q)), \text{ where} \quad (5.21)$$

$$Q = \tilde{\epsilon}_\pi(h^+)\tilde{\epsilon}_\pi(h^-)$$




Figure 5.50: The kaon/pion discrimination variable $\tilde{\epsilon}_\pi$ for kaons at $p = 0.70 \text{ GeV}/c$ in the Monte Carlo. Pions have a flat distribution in the same variable.

which should be generally applicable where the distribution of the $\tilde{\epsilon}_\pi$ for kaons is convex to the origin⁶ as in figure 5.50.

Figure 5.51 shows the effect on the data of increasingly strict cuts in the $\pi^+\pi^-$ probability \mathcal{P} , corresponding to rejection factors of 4, 25, 100, 200 respectively against $\pi^+\pi^-$ combinations. With each cut the total population is reduced by factors 1.6, 4.3, 10, 16 respectively, *i.e.* significantly less than the nominal amounts, and a prominent peak appears at the ϕ mass, together with further structure in the region above threshold ($M(h^+h^-) = 1.05 \approx 1.15 \text{ GeV}/c^2$). This behaviour is consistent with a significant population of both ϕ s and heavy particles (protons and K^\pm , the latter possibly from $K^* \rightarrow K\pi$ decays) which are not suppressed as are pions by the \mathcal{P} cuts.

⁶This method gives in all cases the correct value of the probability for the *product of probabilities* $\tilde{\epsilon}_\pi(h^+)\tilde{\epsilon}_\pi(h^-)$ to have the value Q or less, if the h^\pm are pions so that $\tilde{\epsilon}_\pi$ is flat. Whether it provides optimum discrimination between the $\pi^+\pi^-$ (null) and K^+K^- (alternative) hypotheses depends upon the distribution of $\tilde{\epsilon}_\pi$ for kaons. Where this distribution is convex to the origin, as is typical for a quantity derived from a comparison of *displaced* distributions such as those of $\mathcal{L}_{K/\pi}$ for kaons and pions in figure 5.23, the discrimination should be close to ideal (Wallis, 1942, section III); where it takes the form $(m+1)\epsilon^m$, $-1 < m \leq 0$ it will be the most powerful method possible (Pearson, 1938). Distributions of $\tilde{\epsilon}_\pi$ for kaons are in fact poorly fitted by functions $(m+1)\epsilon^m$ in both momentum ranges.

Empirically the Fisher technique yields a *larger* ϕ peak in the low-momentum sample, where $\tilde{\epsilon}_\pi$ is convex for kaons, than the $\tilde{\epsilon}_\pi(h^\pm) < k$ box cut; although the ϕ peak is still clearly visible under such a cut. The box cut yields a larger peak on the high-momentum sample. A realistic study of the efficiency of various cut strategies for ϕ s at a range of momenta, above a background, will require special preparation due to the low rate of ϕ production in the Monte Carlo with default settings; this has not been carried out at the time of writing.

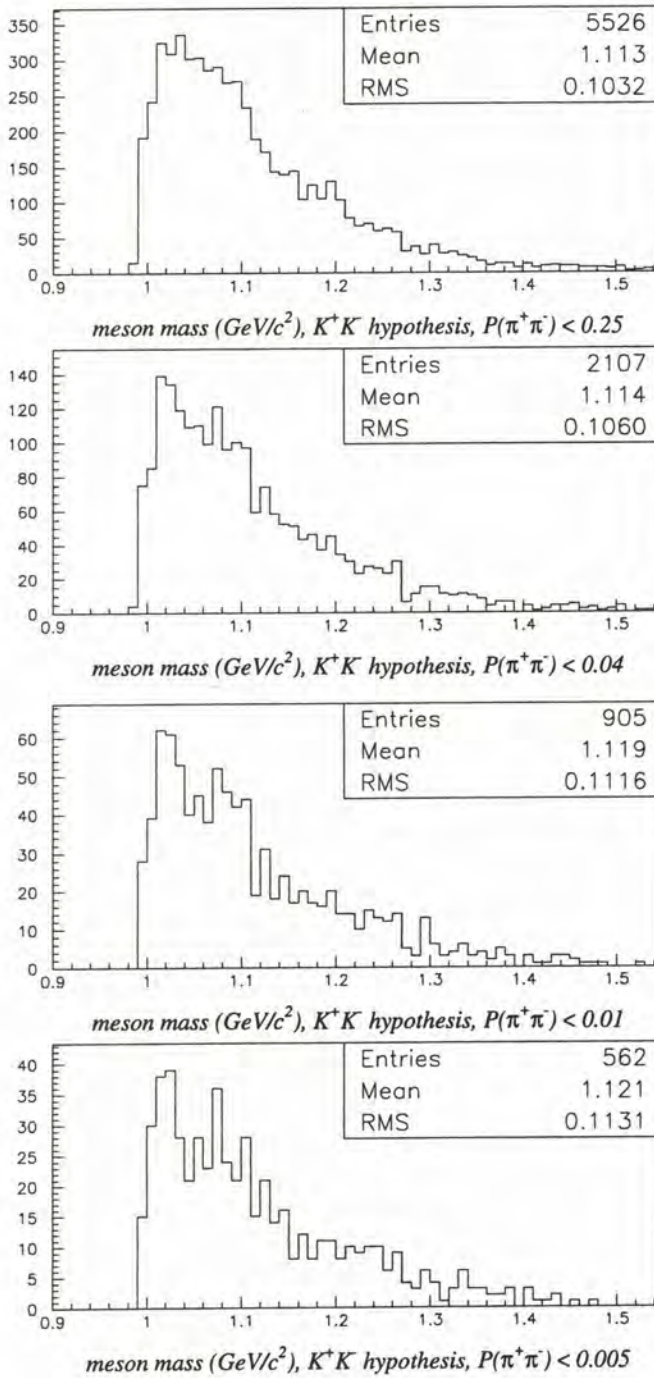


Figure 5.51: The mass $M(h^+h^-)$ for the low-momentum ($p^{tot} < 2.0$ GeV/c) sample as progressively tighter cuts in the probability of the $\pi^+\pi^-$ hypothesis are applied: from top to bottom, $P(\pi^+\pi^-) < 0.250, 0.040, 0.010, 0.005$ respectively. Note the well-developed peak at the $\phi(1020)$ mass.

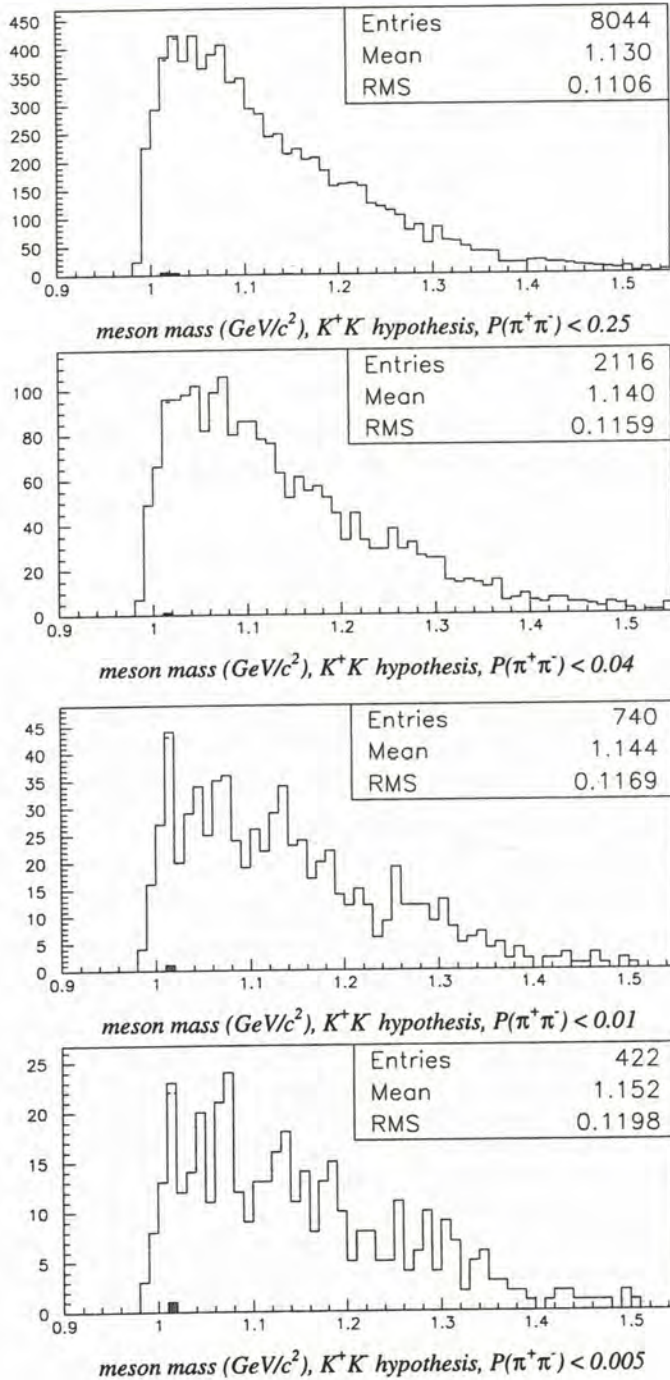


Figure 5.52: The mass $M(h^+h^-)$ for the low-momentum ($p^{tot} < 2.0 \text{ GeV}/c$) sample in the Monte Carlo, under progressively tighter cuts in the probability of the $\pi^+\pi^-$ hypothesis. No peak develops at the $\phi(1020)$ mass and the total population of the plot falls with the value of the $P(\pi^+\pi^-)$ cut: cf. the same distributions for data, figure 5.51; see also the discussion in the text.

The same cuts applied to the Monte Carlo (figure 5.52) yield very different results: the population of the plots are reduced by factors 3.1, 11.9, 34, 60 respectively, much closer to the nominal amounts, and no peak appears at the ϕ mass. Indeed, the initial population of low-momentum ϕ s in the Monte Carlo is insignificant (figure 5.49), but by the same token no peak appears at the ϕ mass as an artefact of the cuts.

The difference between the results for data and Monte Carlo is striking. As a precaution against artefacts a number of variations on the selection procedure have been tried:

1. tightening the cut against electrons to $\tilde{\epsilon}_e < 0.05$;
2. adding a cut $\tilde{\epsilon}_\pi > 0.01$ on positive tracks to suppress protons⁷;
3. selecting K^+K^- candidates using the box cut $\tilde{\epsilon}_\pi(h^\pm) < k$ applied to the high-momentum sample, or a “diagonal” cut $(\tilde{\epsilon}_\pi(h^+) + \tilde{\epsilon}_\pi(h^-)) < k'$.

The plot population is reduced in the first two cases, and the relative size of the peak at m_ϕ is reduced in the last case, but the peak is still apparent in all cases.

The presence of protons, the lack of a ϕ -rich Monte Carlo sample to use as a basis of comparison, and the significant kinematic distortion introduced by the cut against electrons (an effective lower momentum limit at 0.50 GeV/c, the lowest momentum for which the standard e/π discrimination algorithm of section 5.2 may be applied⁸) make it impossible to adopt the simplest interpretation of these results—that there is a significant population of low-momentum ϕ mesons not described by the NOMAD Monte Carlo—without

⁷Removal of protons from the positive track sample in this case is in fact an ideal application of the p/K discrimination variables, as the discrimination is so strong in this momentum region (see figures 5.21 and 5.22) that protons can be separated out without a large effect on the acceptance for kaons. Unfortunately this was not realised at the time of the special processing used in this analysis and the information has been lost (until the NOMAD data is re-processed). The cut actually used, $\tilde{\epsilon}_\pi > 0.01$, is based on a spike at $\tilde{\epsilon}_\pi \approx 0.01$ for positive tracks in the low-momentum sample, absent for the negative tracks which otherwise have a similar $\tilde{\epsilon}_\pi$ distribution.

⁸This is not an intrinsic limitation but an implementation choice which was made from the perspective of the $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ analysis where very powerful e/π discrimination is required. Rejection factors of 10, 20 or 50 against either electrons or pions are achievable in the $0.25 < p < 0.50$ GeV/c momentum region and the algorithm is straightforward to implement for this purpose; this will be carried out when the full TRD information is once again available, after the re-processing of the NOMAD data.

It is not practical to attempt the current analysis without specific cuts against electrons, since electrons pass K/π discrimination cuts against pions without difficulty, swamping the distributions where we would expect a ϕ signal.

further study. Both the proton contamination and the effect of the electron cuts can be addressed after the next round of NOMAD data processing, where a much larger data sample will also become available: given the additional generation of a suitable Monte Carlo sample, with low-momentum ϕ s artificially introduced if necessary, such a study will become possible.

5.3.10 Discrimination between two particle species using TRD information: a summary

The likelihood-ratio based method of discrimination between two particle ID hypotheses, initially developed in the NOMAD experiment for electron/pion discrimination using the TRD, has been successfully implemented for p/π , p/K , K/π , K/μ and even π/μ discrimination:

- at low momenta ($p < 1 \text{ GeV}/c$) powerful discrimination is available in all cases except π/μ ;
- in common with all particle identification methods based on dE/dx , there is a momentum region (both particles close to the minimum ionisation $\beta\gamma \approx 3$) where discrimination is not possible;
- in the high momentum region ($p \gtrsim 2 \text{ GeV}/c$) modest discrimination is available in the p/π , K/π and K/μ cases;
- *populations* of muons and pions may be distinguished based on the distribution of the π/μ discrimination variables.

Tests on track samples from $\Lambda^0 \rightarrow p\pi^-$ decay, kaon-enriched tracks selected by cuts in the electromagnetic calorimeter, and both muon-dominated and muon-suppressed track samples all confirm the behaviour expected on the basis of the Monte Carlo.

A correction for the effect of tracks at non-normal incidence to the TRD has been successfully introduced; a further correction to allow the equal treatment of tracks with “shared hits” in the TRD region is foreseen. With this improvement, and the increase in statistics which will follow the next reprocessing of the NOMAD data, it should be possible to complete the search for $\phi \rightarrow K^+K^-$ decays discussed in section 5.3.9. The results already presented are encouraging and are consistent with a measurable sample of ϕ s in both the low- and high-momentum regions. This would be the first measurement of ϕ production in neutrino interactions, and serves as an example of the usefulness of the technique which has been developed here.

Chapter 6

Event-quality and charged current rejection

6.1 Introduction

In chapters 3 and 4 we considered in turn kinematic variables, and variables sensitive to the 3π structure, for separating three-prong tau decay events from neutral current interactions. In the process we took the reconstruction of the event—in particular, the total momentum vector \vec{p}^{tot} —for granted. We also ignored the problem of backgrounds to the analysis due to charged current interactions of the dominant muon neutrinos

$$\nu_\mu + N \longrightarrow \mu^- + X$$

as well as the contaminating (anti)neutrinos

$$\begin{aligned}\bar{\nu}_\mu + N &\longrightarrow \mu^+ + X \\ \nu_e + N &\longrightarrow e^- + X \\ \bar{\nu}_e + N &\longrightarrow e^+ + X\end{aligned}$$

where the lepton μ^\mp/e^\mp is not identified. We now turn our attention to these backgrounds, and to quality-of-reconstruction issues, taking the structure of the analysis developed in chapters 3 and 4 as fixed.

The definition of a fiducial volume for neutrino interactions, and other cuts on the properties of the primary interaction vertex, are presented in section 6.2. The reconstruction of the total event momentum \vec{p}^{tot} , including both the choice of candidate particles for inclusion in \vec{p}^{tot} and the method of assigning individual momenta \vec{p}^i , is then described in section 6.3. Detector malfunctions, and event simulation and reconstruction problems are

discussed in section 6.4. The danger posed by mis-identified charged-current (CC) events is then briefly reviewed in section 6.5.

Electron/hadron separation using information from the TRD, the most important tool for the rejection of $\nu_e/\bar{\nu}_e$ CC backgrounds, has already been discussed at length in the previous chapter. In section 6.6 we survey the presently available tools for muon identification, and more generally the rejection of the $\nu_\mu/\bar{\nu}_\mu$ CC backgrounds. The choice of cuts to reject both muonic and electronic backgrounds is then presented in section 6.7.

With all of these cuts in place, we will finally turn to the conduct of the tau decay search in chapter 7.

6.2 Cuts on the primary vertex

The key feature returned by the NOMAD event reconstruction code is the primary interaction vertex. Drift chamber tracks emerging from the same location are fitted to common vertices by a Kalman filter process with an appropriate χ^2 threshold, taking the variances of the reconstructed positions and momenta into account (see also section 2.5.3). The most “upstream” of these, *i.e.* the one closest to the beam origin, save for trivial “start-vertices” due to backward-going tracks *etc.*, forms the *primary vertex* and in well-reconstructed events it closely approximates the position of the interaction

$$\nu_\mu + N \rightarrow \mu^- + X$$

(or similar for $\bar{\nu}_\mu$, neutral current interactions *etc.*). The primary vertex is used as the basis of subsequent reconstruction and analysis, and is especially important for the reconstruction of the neutral component of the event: for example the plentiful gamma rays from $\pi^0 \rightarrow \gamma\gamma$ are assigned momenta by combining the energy E_{ECAL} of the resulting cluster in the electromagnetic calorimeter, and the vector joining the cluster position \vec{r}^{ECAL} to the primary vertex position \vec{r}^{prim} :

$$\vec{p}^\gamma = E_{\text{ECAL}} \frac{\vec{v}^\gamma}{|\vec{v}^\gamma|}$$

where

$$\vec{v}^\gamma = \vec{r}^{\text{ECAL}} - \vec{r}^{\text{prim}}$$

The distribution of primary vertex positions \vec{r}^{prim} in fact reproduces the distribution of material in the detector, as it should. The (x, y) positions of




Figure 6.1: Distribution of primary interaction vertices in (x, y) in the NOMAD data. The concentration of interactions where the neutrino beam is most intense (close to $(0, 0)$ in these coordinates) and where the material of the detector is most dense (the nine points where melamine spacers maintain the gas widths of the individual drift planes) are easily visible.

vertices are shown in figure 6.1, and longitudinal (z) positions in figure 6.2: in both cases the correspondence between the interaction rate and the construction of the detector is clearly visible. Straightforward cuts in \vec{r}^{prim} , corresponding to the edges of the fully-sensitive region of the drift chambers, are used to define a *fiducial volume* for neutrino interactions:

$$\begin{aligned}
 -130 < x^{\text{prim}} < 130 \text{ cm} \\
 -125 < y^{\text{prim}} < 135 \text{ cm} \\
 5 < z^{\text{prim}} < 405 \text{ cm},
 \end{aligned}
 \tag{6.1}$$

The cuts in x and y correspond to the region where information from at least two out of the three sets of drift planes (those with wires at 0° , -5° and $+5^\circ$ to the x -axis: see section 2.4.4) is available. The cuts in z correspond to the beginning and the end of the main drift chamber target: the cut

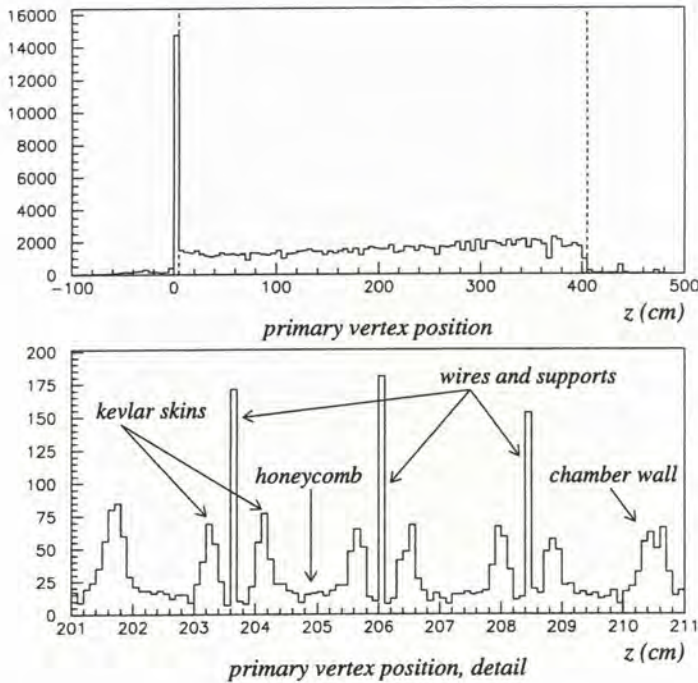


Figure 6.2: Distribution of primary interaction vertices in z , the longitudinal coordinate of the detector. In the upper plot the even distribution of interactions along the drift chamber volume, expected for a neutrino beam (apart from a small effect due to variation of trigger efficiency with z) is clear. The spike at zero is due to non-vetoed throughgoing muons and interactions in the magnet coil *etc.*: the cuts in z defining the fiducial volume for neutrino interactions are shown dotted.

In the lower plot a detail of an individual drift chamber is shown. The longitudinal distribution of material in the chamber is clearly seen (*cf.* the drift chamber description in section 2.4.4).

$z^{\text{prim}} > 5$ cm is especially important for rejecting interactions taking place in material upstream of the drift chamber target (where resolution will be degraded by the intervening matter) and false “primary” vertices due to throughgoing muons (see figure 6.2).

These cuts in the z coordinate correspond to the 1996 data-taking run when the full complement of 11 modules (44 drift chambers in total) was installed in the target region. During the early part of the 1995 data-taking run the detector was operated with only part of the target installed, so fiducial

cuts of

$$\begin{aligned} 260 < z^{\text{prim}} < 405 \text{ cm, and} \\ 115 < z^{\text{prim}} < 405 \text{ cm} \end{aligned} \quad (6.2)$$

are imposed for the so-called 4-module and 8-module running periods respectively.

Using the presence of an identified muon to tag $\nu_\mu + N \rightarrow \mu^- + X$ events we may assess the resolution of the detector for neutrino interactions by considering the total event momentum in the transverse plane, p_T^{tot} . Up to Fermi motion effects this quantity should be zero for interactions where all particles in the final state are measured, so its mean value is a measure of loss of particles, momentum resolution effects *etc.*. Figure 6.3 shows the distribution of $\langle p_T^{\text{tot}} \rangle$ as a function of the x - and y -coordinates for events with




Figure 6.3: Mean total transverse momentum p_T^{tot} of events with an identified muon, as a function of the x (upper plot) and y (lower plot) coordinate of the primary vertex. The degradation in resolution at the edges of the drift chamber volume, due to missed tracks, information from some but not all drift planes *etc.* is clearly visible. The cuts defining the fiducial volume are shown dotted.

an identified muon, within the fiducial region in z , but without other cuts applied. The degradation in resolution towards the edges of the detector, and particularly outside the fiducial region defined above, is clear.

Loss of charged tracks may be assessed by noting that neutrino interactions on nucleons should have a total charge of either 0 or +1 in the final state, corresponding to interactions on neutrons and protons respectively (ignoring nuclear reinteractions *etc.* producing additional particles). The total charge of tracks at the primary vertex is shown for events with an identified muon in figure 6.4: the distribution is indeed dominated by cases with $\sum q = 0, 1$, with a bias towards $\sum q = 0$ as expected due to the presence of two rather than one negatively-charged valence quarks in the neutron. Cases $\sum q = -1, +2$, principally due to the loss of short tracks or the failure to attach such tracks to the primary vertex, are also numerous. A quality cut requiring

$$-1 \leq \sum q \leq 2 \quad (6.3)$$

is imposed: events outside this range are assumed to have at least two lost or non-attached tracks and therefore an unreliable estimate of the total momentum \vec{p}^{tot} .




Figure 6.4: The total charge of tracks attached to the primary interaction vertex, for muon-producing interactions within the fiducial volume of the detector.

Events where one of the tracks has a very poorly reconstructed momentum are also discarded. The Kalman filter track fitting code returns a covariance matrix for the fitted parameters of each track, from which variances on track properties, in particular its momentum p , may be derived. The maximum root-mean-square momentum variance $\max\{\sigma(p)\}$ among all primary tracks is shown for events in the data in figure 6.5. A small tail out to very large




Figure 6.5: The maximum standard deviation on track momentum $\sigma(p)$ returned by the track-fitting code for all tracks at the primary vertex. Note the logarithmic scale.

values, due to tracks which are short and/or have very high momentum, is apparent (compare expression (2.13) for the resolution of the drift chambers, and the discussion in section 2.4.4). A quality cut is placed at

$$\max\{\sigma(p)\} < 5 \text{ GeV}/c \quad (6.4)$$

6.3 Choice of objects for estimation of \vec{p}^{tot}

After the primary vertex the next most important feature of the event reconstruction is the total momentum \vec{p}^{tot} of particles assumed to be emerging from this vertex. Most straightforwardly this includes the reconstructed momenta of drift chamber tracks fitted to the primary, but neutrals in the ECAL (presumed to be gammas, as discussed above), V^0 vertices due to K_S^0 , Λ^0 and converting gammas $\gamma \rightarrow e^+e^-$, and other objects must also be taken into account. This presents some questions as to both the *choice* of objects for inclusion in the total event momentum, and the method of calculating their momentum contributions. We turn to these questions in the following sections.

The appropriate quality cuts on secondary vertices (including V^0) and the so-called *hanging* tracks are presented in sections 6.3.1 and 6.3.3. The procedure for assigning a momentum vector to neutral interactions is then discussed in section 6.3.4. The hadronic and electromagnetic calorimeters, taken together, offer some sensitivity to heavy neutral particles, and the effect of including reconstructed neutrals of this kind is assessed in section 6.3.5; the need for quality cuts on short drift chamber tracks is then assessed in

section 6.3.6. The total momentum reconstruction algorithm is summarised in section 6.3.7.

6.3.1 Treatment of secondary vertices

The procedure currently used in the NOMAD reconstruction code for fitting tracks to vertices is imperfect, and vertices are on occasion split into two or more reconstructed vertices close together in space. This is especially true of the primary vertex in events which are “crowded” by many tracks in the space immediately beyond the interaction point. Where the primary and the secondary vertex are joined by a drift chamber track—and the Kalman filter procedure has successfully fit this track to both vertices—then the spacing of the vertices presents no difficulties. Secondary vertices without “incoming” tracks which are close to the primary, however, are predominantly fragments of the true primary vertex in Monte Carlo events. In common with other NOMAD analyses we therefore merge “neutral” (no-incoming-track) secondary vertices within

$$\begin{aligned} |\Delta x| &< 5.0 \text{ cm,} \\ |\Delta y| &< 5.0 \text{ cm,} \end{aligned}$$

and

$$|\Delta z| < 15.0 \text{ cm} \tag{6.5}$$

of the primary into an expanded primary vertex, the so-called *box vertex*. This is used as the basis of subsequent reconstruction and analysis: in particular, all tracks emerging from this box vertex are considered to come directly from the interaction point and may potentially be included in a tau decay candidate $\tau \rightarrow \nu 3\pi$ etc..

6.3.2 Cuts on V^0 vertices

An exception is made for the so-called V^0 vertices—vertices with two tracks of opposite charge emerging from a neutral secondary vertex, and moreover consistent in invariant mass with one of the long-lived strange particle decays

$$\begin{aligned} K_S^0 &\longrightarrow \pi^+ \pi^- \\ \Lambda^0 &\longrightarrow p \pi^- \\ \bar{\Lambda}^0 &\longrightarrow \bar{p} \pi^+ \end{aligned}$$




Figure 6.6: A comparison of the total V^0 momentum \vec{p}^{V^0} and the presumed line-of-flight of the neutral particle from the primary vertex: (upper plot) closest approach of the linear extrapolation of \vec{p}^{V^0} to the primary; (lower plot) the component perpendicular to the line-of-flight, $p_{\perp}^{V^0}$, used as a quality cut (see the text).

or a photon conversion

$$\gamma \longrightarrow e^+e^-$$

in the electromagnetic field of one of the nuclei in the detector material. As noted in the discussion of Λ^0 decays in section 5.3.6 above, these vertices form a special case as they have already been tested on a well-defined physical hypothesis. Accordingly a V^0 vertex is only merged into the box vertex if it lies within

$$\begin{aligned} |\Delta x| &< 0.5 \text{ cm,} \\ |\Delta y| &< 0.5 \text{ cm,} \end{aligned}$$

and

$$|\Delta z| < 1.5 \text{ cm} \tag{6.6}$$

of the reconstructed primary.¹

The total momentum $\vec{p}^{V^0} = \vec{p}^+ + \vec{p}^-$ allows the projection of the presumed decaying neutral back to the primary vertex. In most cases this extrapolation passes close to the primary vertex position (figure 6.6, upper plot). In practice the V^0 momentum component perpendicular to the line $\vec{v}^{\text{join}} = (\vec{r}^{V^0} - \vec{r}^{\text{prim}})$ connecting the primary and V^0 vertex is used as a quality cut. Cases satisfying

$$p_{\perp}^{V^0} < 0.100 \text{ GeV}/c \quad (6.7)$$

are assumed to come from the primary vertex, and a neutral object with momentum \vec{p}^{V^0} is accordingly added to the particles from the primary vertex; cases failing the cut are assumed to be mis-reconstructions, or conversions of bremsstrahlung photons from electron tracks, *etc.*, and ignored for the purpose of this analysis.

6.3.3 Treatment of hanging tracks

Neutral secondary vertices with only one emerging track also form a special case. These so-called *hanging* tracks can be due to

1. interactions of neutrals produced at the primary vertex, where only one outgoing particle is reconstructed, such as neutron-proton scattering;
2. the event reconstruction algorithm breaking the true primary vertex into fragments, as discussed above;
3. the reconstruction breaking a single *track* into two pieces, the second of which is the hanging track;
4. so-called *asymmetric conversions* $\gamma \rightarrow e^+e^-$, where one of the tracks has very low momentum, or emerges at a large angle *etc.* and is not reconstructed; or
5. mis-reconstruction effects of other kinds, fragments of low-momentum “looping” tracks not properly recovered by the tracking algorithm, *etc.*

¹Note that this special treatment for V^0 vertices is motivated by V^0 studies such as those presented in section 5.3.6, in particular the need to be sensitive to low-momentum Λ^0 decaying close to the primary vertex ($c\tau_{\Lambda} = 7.9 \text{ cm}$). It is not necessarily well-suited to an oscillation analysis where efficient recovery of primary tracks—especially those due to potential tau decays—is more important than identification of the strange component of the hadronic system. This point is discussed in section 7.6.

Accordingly a rather complex treatment of hanging tracks is required. If the first drift-chamber “hit” of the track lies within the cut defined at (6.6) above, the track is merged into the box vertex at the primary (corresponding to case 2). If the first hit lies within the cylinder

$$\begin{aligned} \Delta r &< 4 \text{ cm,} \\ -5 &< \Delta z < 20 \text{ cm} \end{aligned} \quad (6.8)$$

about the end of another track (where $\Delta r = \sqrt{(\Delta x)^2 + (\Delta y)^2}$) the hanging track is assumed to be a continuation or scattering of this track (case 3 above) and is excluded from the momentum reconstruction to avoid double-counting.




Figure 6.7: The closest approach by the momentum vector of a “hanging” track, in a linear extrapolation, to the primary interaction vertex. See the text.

For the remainder, the hypothesis of an asymmetric conversion (case 4 above) is tested using a linear extrapolation of the supposed γ backward along the path defined by the hanging track momentum \vec{p}^{hanger} . The distance of closest approach of this extrapolation to the primary vertex is shown in figure 6.7: *cf.* the same distribution for true V^0 vertices, figure 6.6 above. Where the extrapolation passes within

$$\begin{aligned} |\Delta x| &< 3 \text{ cm, and} \\ \Delta r &< 4 \text{ cm} \end{aligned} \quad (6.9)$$

of the primary, the hanging track is assumed to be due to an asymmetric conversion of a primary photon, and a neutral object of momentum \vec{p}^{hanger} is added to the primary vertex. (The differing treatment of the position mismatch in x and y reflects the differing resolution of the drift chambers for the two coordinates: see section 2.4.4 above.) The same cut and momentum assignment is assumed to be satisfactory for np scattering and the like (case 1).

The rather numerous hanging track vertices failing this cut are assumed to be due to misreconstruction effects (case 5 above) and are not treated further in this analysis.²

6.3.4 Treatment of neutral secondary “interactions”

Neutral secondary vertices with more than one emerging track which are *not* V^0 present a problem for the total momentum reconstruction. The most straightforward interpretation of such vertices is that a neutral particle such as a neutron or K_L^0 from the primary vertex has interacted with the detector material to produce the observed tracks. The total momentum $\sum_i \vec{p}_{int}^i$ of tracks at the (secondary) interaction vertex does not however appear to be a good estimate of the momentum of the interacting particle: its component perpendicular to the assumed line-of-flight from the primary, p_{\perp}^{int} , is large in the mean and has a significant tail to very high values (figure 6.8). Moreover high values of p_{\perp}^{int} are not an artefact of misreconstruction of en-




Figure 6.8: The component of the total momentum at a “neutral interaction” vertex perpendicular to the presumed line-of-flight of the neutral particle from the primary. Unlike the V^0 case (figure 6.6) the distribution is dominated by a tail to large values.

ergetic secondary tracks, *etc.*: if we consider the angle in space between the total momentum at the secondary vertex $\sum_i \vec{p}_{int}^i$ and the line-of-flight from the primary, we find that interactions with low p_{\perp}^{int} approximately “point” to

²Asymmetric conversion of bremsstrahlung photons from electrons in the event are an additional source of hanging tracks. No allowance is made at present for the bremsstrahlung process in this analysis, although events with identified primary electrons are subject to a veto: see section 6.7.4 below.

the primary vertex, while the numerous cases with large p_{\perp}^{int} point at large angles to the primary (figure 6.9).




Figure 6.9: The angle between the total momentum at a “neutral interaction” vertex and the presumed line-of-flight of the neutral particle from the primary. A cut in the perpendicular momentum component p_{\perp}^{int} separates the interactions into (a) pointing, and (b) non-pointing cases.

Assuming the interaction of a neutral from the primary, we interpret this phenomenon as being due to an unreconstructed neutral component in the final state, in addition to the reconstructed tracks, introducing lost momentum p_{\perp}^{int} at the interaction vertex. On the hypothesis that this unreconstructed neutral component is a single neutron, and that the interacting (incoming) neutral particle is also a neutron, we may solve for the momentum of the incoming particle. This has been attempted and in the great majority of cases no solution exists. A variant on the same idea, hypothesising an interacting K_L^0 and a lost π^0 in the final state, has also been attempted. This is less restrictive but still fewer than 10% of cases admit of a solution.

It appears that hypotheses of this kind are in general too restrictive, and

indeed we would expect many interactions to yield a final state with more than one neutral particle (several π^0 , or a π^0 in addition to a final-state neutron *etc.*). Unfortunately there are insufficient constraints in these cases to allow for a solution for the momentum of the incoming particle.

Accordingly we apply the simple expedient of rotating the neutral interaction momentum onto the line-of-flight from the primary. That is, we take the modulus of the total momentum at the secondary interaction vertex

$$\left| \sum_i \vec{p}_{int}^i \right|$$

as the best available estimate of the magnitude of the incoming neutral momentum, and the line-of-flight from the primary

$$\vec{v}^{join} = (\vec{r}^{int} - \vec{r}^{prim})$$

as the best estimate of its direction, and form a momentum vector

$$\vec{p}^{int} = \left| \sum_i \vec{p}_{int}^i \right| \frac{\vec{v}^{join}}{|\vec{v}^{join}|} \quad (6.10)$$

A neutral object with this momentum is then added to the primary vertex.

Clearly in cases where there is a lost neutral component in the final state, $|\sum_i \vec{p}_{int}^i|$ will under-estimate the energy of the interacting neutral. There is not however an obvious way to correct for this without introducing spurious additional momentum in cases where the mismatch p_{\perp}^{int} is due to finite momentum resolution *etc.* (presumably the “pointing” cases in figure 6.9 above).

Comparison of the “rotated momentum” treatment with other reconstruction options is encouraging. The total event transverse momentum p_T^{tot} is shown in figure 6.10 for events with an identified muon in the data: events with neutral interaction vertices have larger p_T^{tot} than those without, in part because these events are more energetic on average and in part due to mis-reconstruction of the neutral interaction momentum. Assigning the rotated interaction momentum to \vec{p}^{int} (c) leads to lower p_T^{tot} in the mean than either (b) the $\sum_i \vec{p}_{int}^i$ reconstruction, or (d) omitting the neutral interactions entirely. We see similar behaviour in the ν_{μ} CC Monte Carlo.

The equivalent test cannot be carried out in the data for neutral current events but in the Monte Carlo we may construct the vector difference between the “missing” transverse momentum of the event, and the true transverse momentum of the final-state neutrino, $\Delta\vec{p}_T = (\vec{p}_T^{miss} - \vec{p}_T^{\nu})$, as a measure of the transverse momentum mis-reconstruction. Distributions of $|\Delta\vec{p}_T|$ are shown

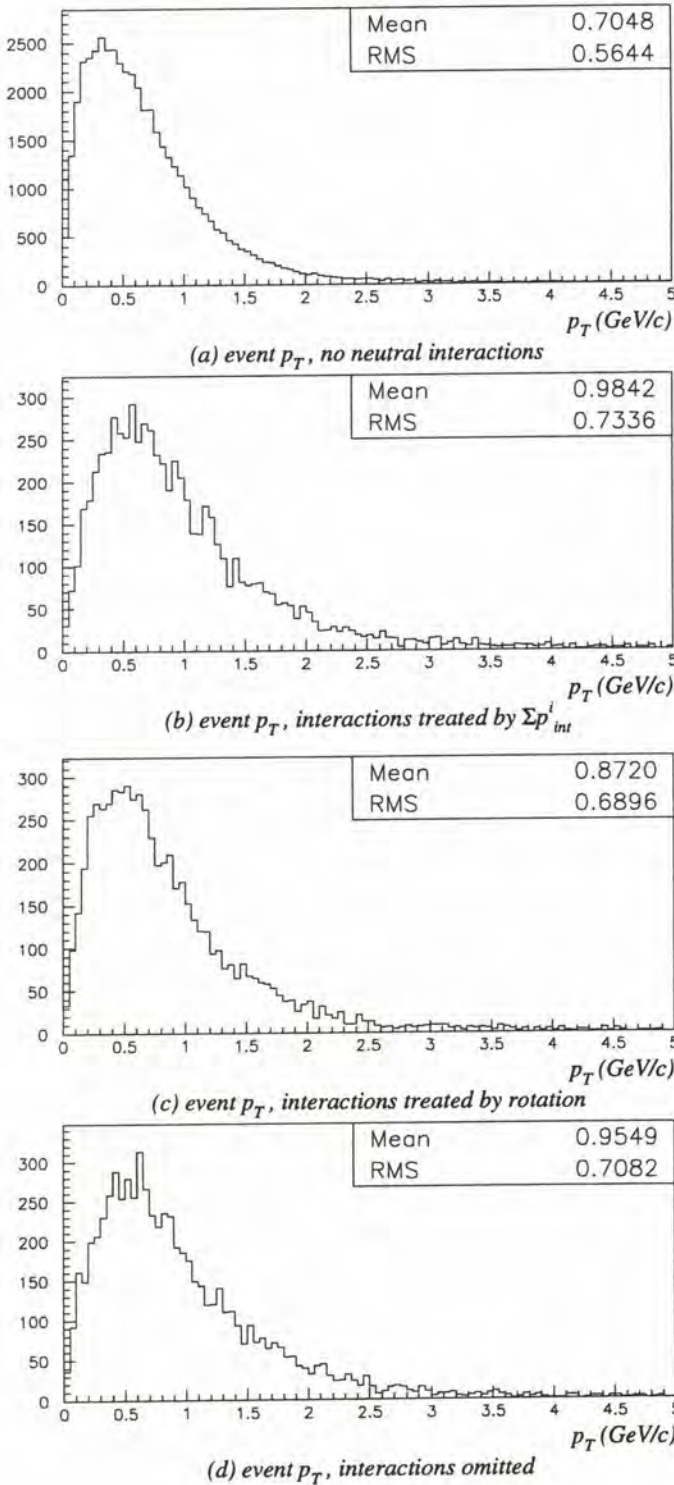


Figure 6.10: Transverse component of the total momentum p_T^{tot} for events with an identified muon in the data: (a) where there are no neutral interactions in the event; (b) where there is ≥ 1 neutral interaction, with momentum estimated by $\sum_i \vec{p}_{int}^i$; (c) where momentum is rotated onto the line-of-flight of the neutral; and (d) where the interaction is omitted from the momentum reconstruction.

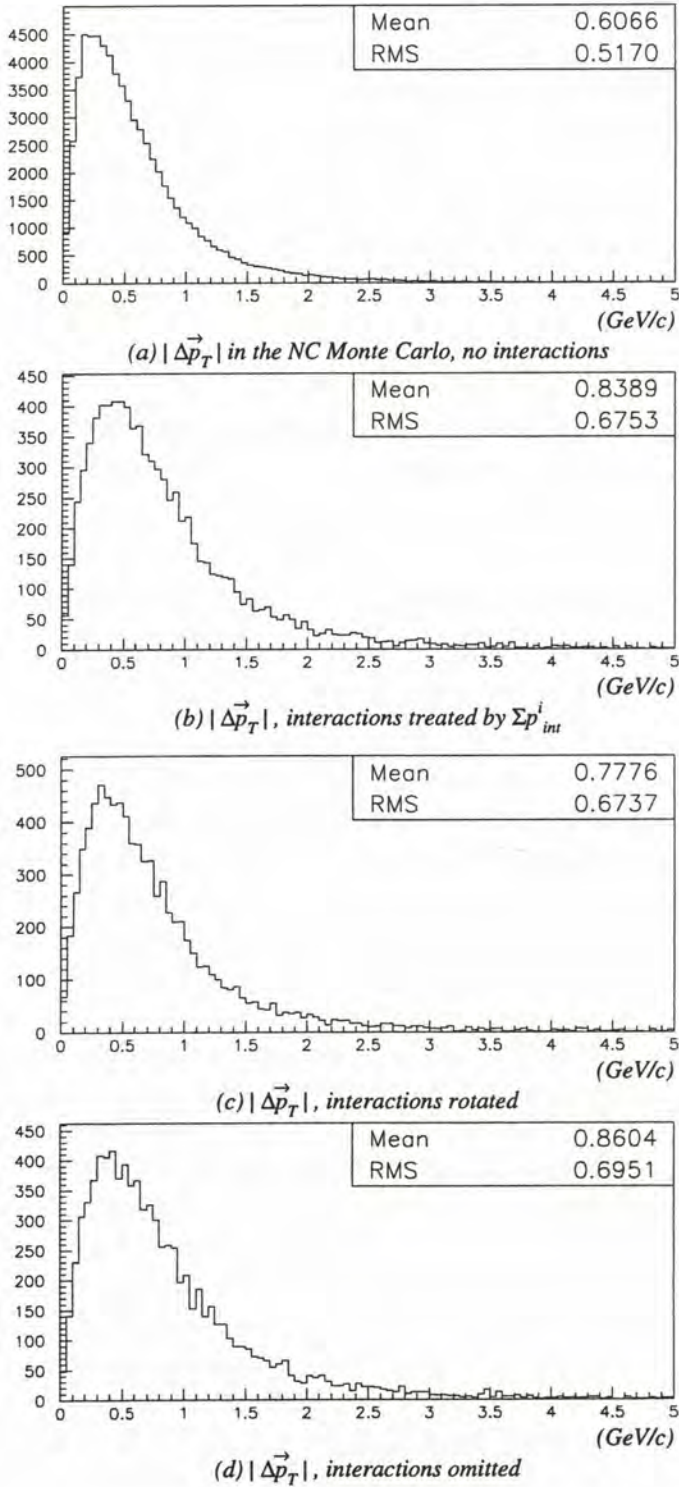


Figure 6.11: Transverse momentum mis-reconstruction $\Delta\vec{p}_T = (\vec{p}_T^{miss} - \vec{p}_T^\nu)$ in neutral current Monte Carlo events: (a) where there are no neutral interactions in the event; (b) where there is ≥ 1 neutral interaction, with momentum estimated by $\sum_i \vec{p}_{int}^i$; (c) where momentum is rotated onto the line-of-flight of the neutral; and (d) where the interaction is omitted from the momentum reconstruction.

in figure 6.11: as in the previous case, the rotated interaction momentum yields superior reconstruction in the mean. This behaviour is preserved under cuts in the kinematic variables q_T , $E^{3\pi}$ and the required mass M_R to select high-probability tau decay candidates.

6.3.5 Use of hadronic calorimeter information

Up to this point the momentum reconstruction is sensitive to heavy neutrals through the decay of K_S^0 , and the nuclear interaction of n , K_L^0 *etc.* in the drift chamber volume: neutrons and K_L^0 which do not interact are lost, leading to a degradation in the momentum resolution of the detector. The hadronic calorimeter (HCAL; see section 2.4.9) however offers some sensitivity to the energy of hadrons within its acceptance. The HCAL was originally conceived as a “veto” for neutrino oscillation analyses—allowing the rejection of events where a neutral hadronic component plays a significant role in the total momentum—rather than as a contributor to the momentum reconstruction on a par with the other subdetectors. The poor energy resolution of the device (section 2.4.9) is consistent with this. The feasibility of using HCAL information in the momentum reconstruction should however be assessed.

We noted in section 2.4.9 that an energy estimate for charged hadrons is obtained by forming a linear combination of the measured energy of the ECAL cluster *and* the HCAL cluster matched to the track, and then making a correction for the nonlinearities of the system. This procedure is well-tested for charged hadrons and we expect on physical grounds that the same procedure will yield a reasonable estimate of the particle energy in the case of neutral hadrons. There is however the complication that, in general, several tracks deposit energy in the poorly-segmented HCAL so we are dealing with the energy remaining after a subtraction algorithm has been applied.

Given these “neutral” HCAL clusters remaining after energy subtraction, momenta were assigned in the following way:

- Where the HCAL neutral has been “matched” to a neutral energy cluster in the ECAL by the NOMAD reconstruction code, the corrected energy estimate described in section 2.4.9 is used, and the presumed neutral hadron is assigned a position at the ECAL cluster centroid in (x, y) , and one nuclear interaction length into the ECAL in z . A vector \vec{v}^{join} connecting the primary vertex to this point is formed and the neutral is assigned a momentum

$$\vec{p}^{neut} = E(E_{ECAL}, E_{HCAL}) \frac{\vec{v}^{join}}{|\vec{v}^{join}|}.$$

- Where the HCAL neutral has not been matched in this way, the measured energy in the HCAL alone is corrected for nonlinearities as described in section 2.4.9, and the presumed neutral hadron is assigned a position at the *HCAL* cluster centroid in (x, y) , and one nuclear interaction length into the HCAL in z . A momentum vector is then formed by analogy with the preceding case.

When examining the performance of this momentum assignment, we may ignore events which have no neutral energy deposited in the HCAL.³ In the remainder of this section we consider events with such a deposit, and an identified muon, in the ν_μ charged-current Monte Carlo. In cases where there is a neutron or K_L^0 within the calorimeter acceptance in the Monte Carlo (figure 6.12) there can be a small reduction in the mean event p_T when the HCAL neutrals are included in the reconstruction, depending on the other cuts *etc.* which are applied. For the inclusive sample shown in the figure, the performance is in fact slightly worse in the mean.

However, in cases where there is no neutron or K_L^0 within the calorimeter acceptance (figure 6.13), momentum resolution is badly degraded: $\langle p_T^{tot} \rangle$ rises by over 0.150 GeV/ c in the sample shown. In events of this kind, energetic tracks deposit more energy than expected, and/or deposit energy far from the projection of the track into the HCAL volume, and this deposit is misinterpreted as the footprint of a neutral hadron. Spurious momentum is thus added to the event. We should also note that these cases are more numerous than neutral energy clusters where a true heavy neutral exists (*cf.* figure 6.12).

We therefore do *not* include neutral HCAL energy deposits in the total momentum reconstruction for this analysis. The question of whether to use a momentum reconstruction with HCAL neutrals included as a cross-check and veto for the standard reconstruction, is more difficult to answer, requiring an analysis of Monte Carlo background events surviving cuts, tails of

³For completeness we should note that the HCAL clustering algorithm depends upon clustering in the ECAL: the ECAL energy matched to a charged track is taken into account when calculating the energy deposit in the HCAL “expected” for the track; this can determine whether deposited energy in one or more HCAL slabs is deemed to be a neutral energy deposit or just a fluctuation of the deposit by the charged track(s).

The ECAL clustering algorithm used in this connection is different to the one used in the standard reconstruction of the event, for historical reasons. To maintain a consistent analysis, this alternative ECAL clustering was therefore used to reconstruct neutral ECAL energy into γ s from the primary vertex, when HCAL neutrals were included in the event momentum. As a result there *are* small differences in the reconstructed momentum of the event, even when no neutral energy deposit is found in the HCAL. These differences are however found to be negligible in the mean.




Figure 6.12: Transverse component of the total momentum p_T^{tot} for events with an identified muon in the ν_μ CC Monte Carlo: (upper plot) using the momentum reconstruction previously described, and (lower plot) with HCAL neutrals included in the momentum reconstruction. Cuts requiring that ≥ 1 heavy neutral particle is produced at the primary vertex and is within the calorimeter acceptance, and that ≥ 1 neutral hadron is reconstructed, have been imposed.

distributions *etc.*. These studies are in a preliminary stage and no such veto has been included in the analysis presented here.

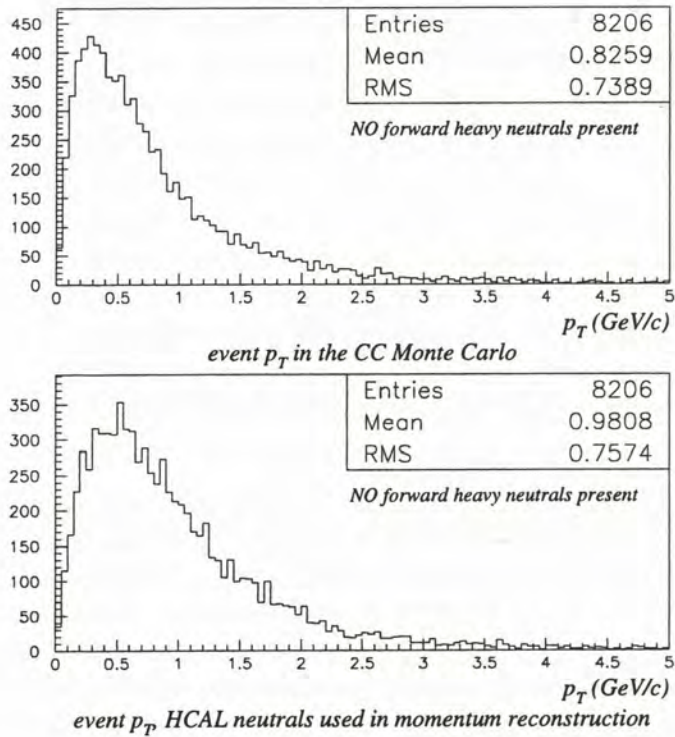


Figure 6.13: Transverse component of the total momentum p_T^{tot} for events with an identified muon in the ν_μ CC Monte Carlo. Events with ≥ 1 neutral hadron reconstructed, but no true heavy neutral within the calorimeter acceptance, are shown.

6.3.6 A note on the use of charged tracks

We now return briefly to the inclusion of charged tracks in the primary vertex. Studies by groups conducting other tau decay searches within the NOMAD data (such as for $\tau^- \rightarrow \nu_\tau \pi^-$) have found a significant extra background in the Monte Carlo, specifically in the form of a large- q_T tail in neutral current events, due to short tracks being included in the momentum reconstruction. Some effect of this kind is expected, as the resolution on measured momentum becomes poor as the number of drift chamber “hits” on the track becomes small (see for example the role of track length L in the parametrisation of the momentum resolution (2.13) in section 2.4.4). Accordingly cuts on the number of drift chamber hits such as $N_{hit} \geq 7$ and $N_{hit} \geq 13$, for a track to be included in the total event momentum, have been advocated.

A range of cuts on N_{hit} has been tried within the framework of this analysis: we have found that

- the “missing” transverse momentum p_T^{miss} in reconstructed $\nu_\mu + N \rightarrow \mu^- + X$ Monte Carlo events increases in the mean as a cut in N_{hit} increases;
- after constructing the difference between the missing transverse momentum and the true final-state neutrino momentum, $\Delta\vec{p}_T = (\vec{p}_T^{miss} - \vec{p}_T^\nu)$, its magnitude $|\Delta\vec{p}_T|$ in reconstructed $\nu_\mu + N \rightarrow \nu_\mu + X$ Monte Carlo events falls to a minimum for $N_{hit} \geq 7 \sim 10$, but the effect is slight (at the fraction of a percent level in the mean);
- under q_T and M_R cuts, $|\Delta\vec{p}_T|$ is minimised when an N_{hit} cut is absent.

A cut on N_{hit} cannot therefore be justified for this analysis by appeal to the (transverse) momentum reconstruction alone. Turning to the rate at which background events in the Monte Carlo survive cuts in the important kinematic and structure variables, we find that the rate of surviving events for $M_R < 2.2 \text{ GeV}/c$, and high q_T , $E^{3\pi}$ cuts shows no significant dependence on an N_{hit} cut for the total momentum \vec{p}^{tot} reconstruction.

There is however a small effect in the *selection* of tracks to form the 3π candidate. If we require $N_{hit} \geq 7$ for tracks contributing to the 3π candidate, we impose an efficiency loss of 1% on the three-prong decay sample for events passing strong q_T , M_R and $E^{3\pi}$ cuts; this rises to 9% for a cut $N_{hit} \geq 13$. The equivalent losses in the neutral current background are 7% and 30%. Such a cut is not compelling on event-rejection grounds, but on the *a priori* ground that momentum reconstruction is less reliable for short tracks we impose a cut $N_{hit} \geq 7$ when performing the 3π selection.

6.3.7 Summary: reconstruction of event momentum

$$\vec{p}^{tot}$$

In summary the total event momentum \vec{p}^{tot} is constructed in the following way:

- All tracks fitted to the primary vertex are included, using the momentum returned by the Kalman filter track fit.
- Tracks emerging from “charged secondary vertices” (those with a fitted incoming track) are ignored.
- Tracks emerging from neutral secondary vertices (other than V^0) within the box

$$\begin{aligned} |\Delta x| &< 5.0 \text{ cm,} \\ |\Delta y| &< 5.0 \text{ cm,} \\ |\Delta z| &< 15.0 \text{ cm} \end{aligned}$$

around the primary are merged into a box vertex, and treated as primary tracks in all respects.

- Tracks from V^0 vertices within

$$\begin{aligned} |\Delta x| &< 0.5 \text{ cm,} \\ |\Delta y| &< 0.5 \text{ cm,} \\ |\Delta z| &< 1.5 \text{ cm} \end{aligned}$$

of the primary are merged into the box vertex.

- No cut on the number of drift-chamber hits of tracks is applied.
- Neutral clusters in the ECAL are assumed to be gamma rays from the primary interaction vertex, and are assigned a momentum

$$\vec{p}^{neut} = E_{ECAL} \frac{\vec{v}^{join}}{|\vec{v}^{join}|}$$

where E_{ECAL} is the cluster energy and \vec{v}^{join} the vector connecting the primary vertex and the cluster centroid at the ECAL front-face.

- Neutral clusters in the HCAL are ignored.

- V^0 vertices with a momentum component perpendicular to the line from the primary vertex satisfying

$$p_{\perp}^{V^0} < 0.100 \text{ GeV}/c$$

are assumed to be due to neutrals from the primary: the total V^0 momentum \vec{p}^{V^0} is used.

- Hanging tracks not merged to the primary are ignored if their first hit lies within

$$\begin{aligned} \Delta r &< 4 \text{ cm}, \\ -5 &< \Delta z < 20 \text{ cm} \end{aligned}$$

of the end of another track. Otherwise a neutral object of momentum \vec{p}^{hanger} is included if the hanger extrapolates within

$$\begin{aligned} |\Delta x| &< 3 \text{ cm, and} \\ \Delta r &< 4 \text{ cm} \end{aligned}$$

of the primary vertex.

- Neutral secondary interactions with ≥ 1 dependent track, which are not V^0 , are treated by summing the momenta \vec{p}_i of the dependent tracks and including a neutral object with momentum

$$\vec{p}^{int} = \left| \sum_i \vec{p}_{int}^i \right| \frac{\vec{v}^{join}}{|\vec{v}^{join}|}$$

where \vec{v}^{join} is defined as above.

All tracks fitted or merged to the primary vertex, also satisfying the condition $N_{hit} \geq 7$ on the number of drift-chamber hits, are considered for the formation of a 3π decay candidate. If the remaining “hadronic system” does not contain at least one track fitted or merged to the primary vertex, the event is rejected.

6.4 Treatment of detector malfunctions and software problems

In addition to requiring a well-recovered primary interaction vertex, and performing as accurate a reconstruction of the event momentum \vec{p}^{tot} as possible, we must take care to remove events which are affected by gross failures of the detector, or of the simulation or reconstruction software. In the following sections we briefly deal with some cuts of this kind: cuts against known problems in the NOMAD subdetectors and software are treated in sections 6.4.1 and 6.4.2 respectively, and a special cut on the generated Fermi momentum in the Monte Carlo is then discussed in section 6.4.3.

6.4.1 Events with known subdetector problems

In the assembly of data from the various subdetectors to form the record of an event, discussed in section 2.5.2 above, certain tests on the integrity of the data are possible. Most straightforwardly, information from a subdetector may be missing; or it may arrive out of proper sequence, be incomplete, or not conform to the expected format due to some malfunction *etc.*. Failures of this kind are flagged in the reconstruction code, and any event with one of these flags set is rejected from the analysis.

Known subdetector malfunctions affecting many events in sequence, noted by the experimental shift crew due to slow control alarms, subdetector monitoring histograms, or other means, were also routinely recorded in the experimental log. Any anomalous state of a subdetector due to an expert intervention *etc.* was recorded in a similar way. A list of data-taking runs affected by malfunctions or interventions of this kind is available, and events from these runs are rejected from the analysis, with two exceptions:

1. if a run is flagged as “bad” only because the information from the WANF beamline monitoring (section 2.2.1) failed to arrive or could not be properly recorded, then events from this run are still analysed, as beam data are not used directly in a $\nu_\mu \rightarrow \nu_\tau$ search;
2. if the only subdetector with a recorded malfunction is the front calorimeter (FCAL, section 2.4.3), events from from this run are processed since the FCAL is not used in this analysis.

In addition to these tests, a special treatment of electromagnetic calorimeter (ECAL) data is required, due to the relatively long shaping time of the readout electronics (1.4 ms). This makes the ECAL data subject to a background due to muons traversing the detector close to, but out of time with,

the event trigger. Interactions which fail to initiate a data-taking trigger may similarly deposit energy in the lead-glass which is read out with a subsequent or preceding event. A fast timing signal is provided by the ECAL readout chain in order to allow the rejection of deposits of this kind (section 2.4.8; see also Autiero et al. (1996a)); matching of ECAL clusters to out-of-time hits in the trigger scintillators, or out-of-time track segments in the muon chambers during “phase two” of the event reconstruction (section 2.5.3) also allows some out-of-time deposits to be flagged. Prior to undertaking a study of this problem, however, and noting in particular that the fast timing signal is efficient only for energetic deposits, so that out-of-time energy may be only partially flagged, we adopt for now the conservative approach of rejecting any event with a known out-of-time deposit.

6.4.2 Events with reconstruction or simulation problems

During the event reconstruction process a number of tests are performed to detect events with too many drift chamber hits to allow reconstruction in a reasonable time, an unusually large fraction of drift chamber or TRD hits at the edge of the detector (suggesting an interaction in the magnet coil or flux return), and the like. Problems with the process of “matching” information between different subdetectors during phase two of the reconstruction (section 2.5.3) are also flagged. Any event with problems of this kind is rejected from the analysis.

Information regarding “unused” drift chamber hits, *i.e.* hits which the tracking code was unable to incorporate into fitted tracks, is also available, both at the whole-event level and in the vicinity of reconstructed tracks and vertices. No cut on such information is carried out in the present analysis.

Finally, the Monte Carlo code used to produce simulated events contains certain buffers of limited size, which can overflow during the simulation of very complex events and under some other circumstances. The “raw data” from which such simulated events are reconstructed is incomplete and therefore potentially inconsistent: simulated events of this kind are rejected from the analysis.

6.4.3 A special cut on the Fermi momentum in the Monte Carlo

We now consider a cut required not to control the quality of the data, but to prevent artificial backgrounds emerging in the Monte Carlo. When simu-

lating the fundamental neutrino-nucleon interaction

$$\nu_\mu + N \rightarrow \mu^- + X$$

etc. in the target, the NOMAD Monte Carlo uses the Fermi momentum distribution of Bodek and Ritchie (1981) for the nucleon $N = n, p$ if it belongs to a larger target nucleus. This distribution, shown in figure 6.14, possesses a small tail out to very large values. Studies on quantities and processes sensitive to the Fermi momentum have found that these events with $p_{\text{Fermi}} > 1 \text{ GeV}/c$ lead to results in the simulation which are not observed in the NOMAD data. Accordingly, it has become standard practice to truncate this distribution at $1 \text{ GeV}/c$.

In the $\tau^- \rightarrow \nu_\tau \pi^-$ analysis the tail $p_{\text{Fermi}} > 1 \text{ GeV}/c$ introduces a serious additional background of high- q_T events in the ν_μ neutral- and charged-

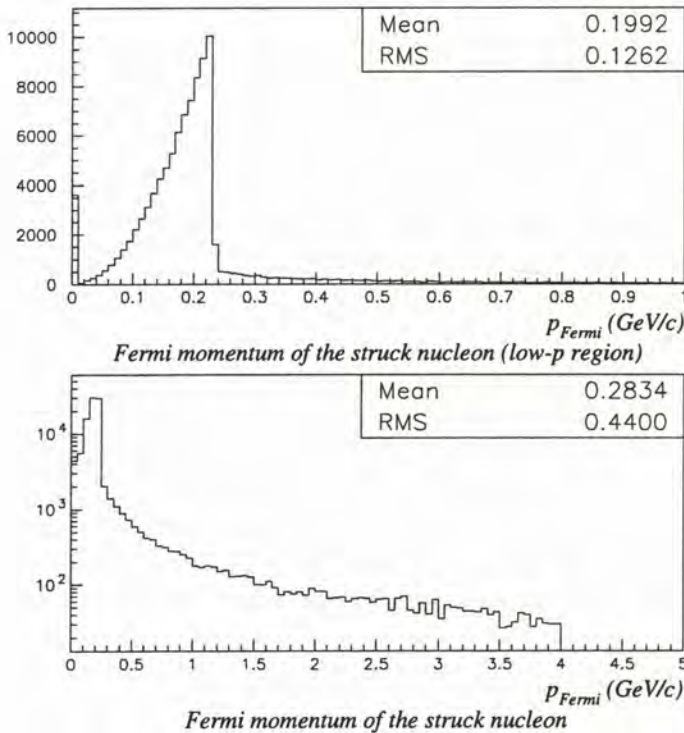


Figure 6.14: Fermi momentum of the struck nucleon for ν_μ CC events in the NOMAD Monte Carlo: (upper plot) the low momentum region, with the cutoff at $\approx 0.230 \text{ GeV}/c$ and the peak at zero (due to interactions on hydrogen nuclei) clearly visible; (lower plot) the full range of the distribution, logarithmic scale.

current Monte Carlo. In the present analysis the Fermi tail also makes a disproportionate contribution to events with high q_T , as shown in figure 6.15. This effect is however largely offset by the other key variables of the analysis: events in the Fermi tail produce 3π candidates which have larger required mass M_R , and are less energetic, than events in the bulk of the distribution. Moreover, these effects are preserved and even enhanced when high- q_T events are considered (figures 6.16 and 6.17).

Noting that the extended tail of the distribution appears to be unphysical, however, we follow the practice of excluding Monte Carlo events which fail the cut

$$p_{\text{Fermi}} < 1 \text{ GeV}/c \quad (6.11)$$

from the analysis.

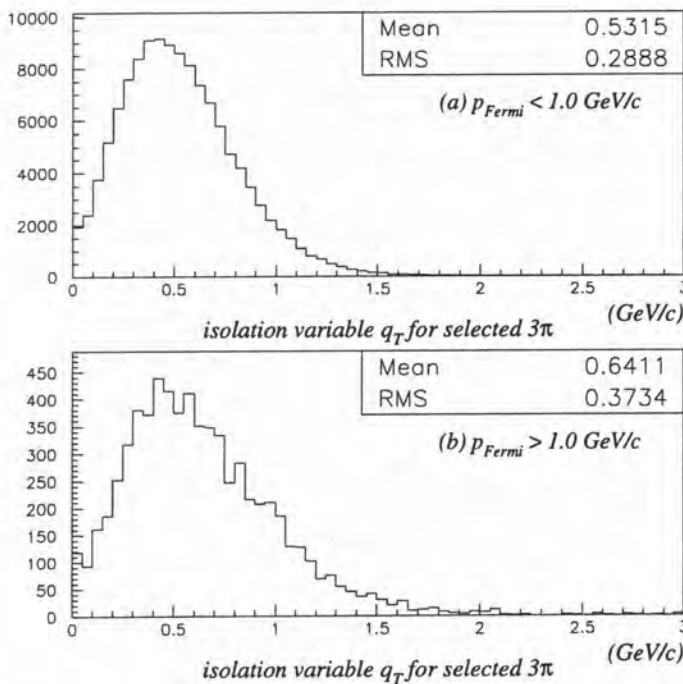


Figure 6.15: The isolation variable q_T for selected 3π candidates in the ν_μ neutral current Monte Carlo, in (a) the bulk of the Fermi momentum distribution, and (b) the extreme tail.

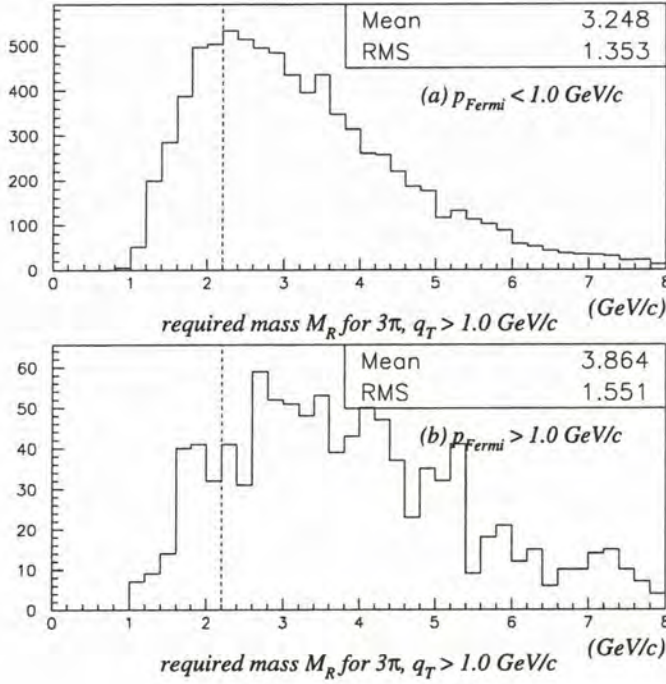


Figure 6.16: The required mass M_R for high- q_T 3π candidates in the ν_μ neutral current Monte Carlo, in (a) the bulk of the Fermi momentum distribution, and (b) the extreme tail. The standard cut at $M_R < 2.2 \text{ GeV}/c$ is also shown.




Figure 6.17: Energy $E^{3\pi}$ for high- q_T 3π candidates in the ν_μ NC Monte Carlo, in (a) the bulk of the Fermi momentum distribution, and (b) in the extreme tail.

6.5 The charged-current backgrounds

Even granted a properly recovered event momentum \vec{p}^{tot} , the kinematic and 3π structure variables discussed in chapters 3 and 4 are not sufficient to remove the background to the three-prong decays from other neutrino interactions. The quantities introduced to date rely on the differences between tau decay and *neutral current* interactions, whereas the majority of the neutrino interactions in the NOMAD are due to the charged current process

$$\nu_\mu + N \rightarrow \mu^- + X.$$

Where the muon is identified, such an event is “tagged” and may be rejected immediately. Where the muon is *not* identified, the event may potentially provide a three-prong decay candidate, as shown in figure 6.18.

As noted in section 2.4, the NOMAD includes a system of muon chambers, providing for the identification of muons which penetrate the back detector support. Clearly any event with a well-reconstructed muon should be excluded from the analysis, since in general a three-prong tau decay event will contain no muon. This will not however be sufficient to suppress the $\nu_\mu + N \rightarrow \mu^- + X$ background: the sample of events is overwhelming ($> 10^6$) and the muon detection system is not perfectly efficient. In addition to positive muon *identification* we therefore also require criteria to *veto* events because certain tracks might be muons (even if the identification is uncertain); or because a significant track, if it were a muon, would be unidentifiable due to low momentum *etc.*.

The construction of such a muon veto procedure depends upon the detail of the muon chamber construction, the reconstruction algorithms used, *etc.* These, together with some other resources for the rejection of $\nu_\mu + N \rightarrow \mu^- + X$ events, are rehearsed in section 6.6 below. The particular muon identification and veto algorithm chosen for this analysis, based on these tools, is then presented in section 6.7.

The closely-related $\bar{\nu}_\mu + N \rightarrow \mu^+ + X$ background may be treated using the same technique. However the electronic backgrounds

$$\begin{aligned}\nu_e + N &\longrightarrow e^- + X \\ \bar{\nu}_e + N &\longrightarrow e^+ + X,\end{aligned}$$

while kinematically similar, require a special treatment due to the differing behaviour of muons and electrons in the detector. For example, the effectiveness of electron identification in the NOMAD is not a strong function of momentum, so that unlike the muonic case, there is no serious background due to low momentum electrons. On the other hand there *is* a background

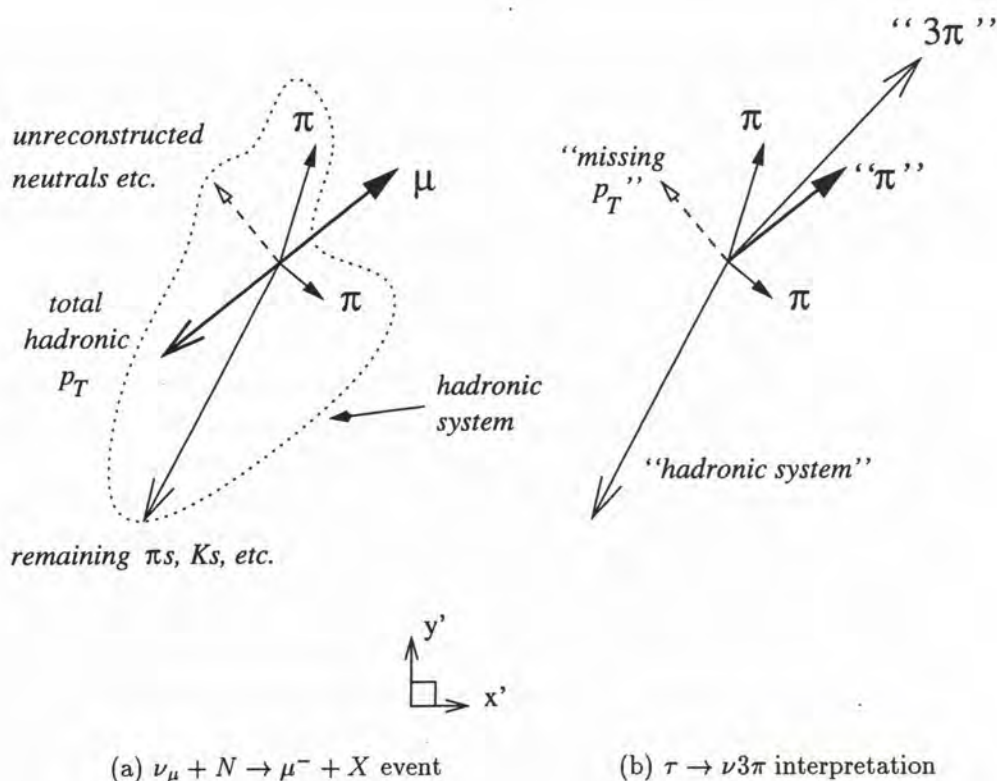


Figure 6.18: A ν_μ charged current event interpreted as a three-prong tau decay after misidentification of the muon; the transverse plane of the event is shown. An energetic ν_μ undergoes a highly inelastic charged-current scatter producing a high-multiplicity, "open" hadronic system, shown in (a). If the muon is insufficiently energetic to penetrate to the muon chambers, or is outside their acceptance, then it is misidentified as a pion, and the event readily forms a tau decay candidate (b).

due to energetic electrons, and despite the relatively low flux of ν_e and $\bar{\nu}_e$ at the NOMAD (section 2.2.3) this background must be treated with some care. The principal tool is the electron/hadron discrimination provided by the TRD, which has been discussed in detail in the preceding chapter. The choice of cuts against electrons for this analysis, using the TRD and some other variables, is presented in sections 6.7.4 and following.

6.6 Elements of muon identification and veto algorithms

The most important, but not the only criterion for both muon identification and veto in the NOMAD, is penetration of the back detector support by a track, which then leaves a signal in the muon detection system (section 2.4.10). The reconstruction of such signals is discussed in sections 6.6.1 through 6.6.3 below. Additional tools for use on individual tracks are discussed in sections 6.6.4 and 6.6.5, while a measure of the rate at which muons escape identification is presented in section 6.6.6. Finally, an independent approach, allowing charged-current *events* to be rejected based on a cut on an event structure variable, is discussed in section 6.6.7.

6.6.1 Reconstructed track segments in the muon system

The most important element of event reconstruction in the muon detection system is the *module*, a pair of muon chambers separated in z to allow the efficient reconstruction of track segments (see section 2.4.10 and figure 2.17). In each module, track segments are first reconstructed in *projection* ($x-z$ and $y-z$), and then (where possible) in *space* (x, y, z).

Tracks-in-projection

A cut through a single module is shown in figure 6.19, together with the response to a single muon. In each projection a muon may leave up to four “hits”, each corresponding to the measurement of the x (or y) coordinate at that $x-z$ ($y-z$) plane; due to the inherent left-right ambiguity of drift tubes there is also a “mirror hit” on the opposite side of the drift wire in each case. Ideally these four hits are sufficient to reconstruct a *track-in-projection*, and resolve the left-right ambiguities: this case is shown for the rightmost track-in-projection in the figure. Reconstruction of three-hit segments is also common, as the single hit efficiency of the system is 92.5%, due principally to dead regions near the metal walls of the tubes: this case is shown for the leftmost track-in-projection in the figure. In the absence of other possible tracks-in-projection, a segment formed from two hits and two dead regions will also be reconstructed.

In other cases (not shown) there may be more than one track-in-projection reconstructed in a given module: two choices for resolving left-right ambiguities may be close in χ^2 , in which case both are retained, or there may be

additional hits in a chamber due to noise or (more rarely) the passage of more than one penetrating particle, *etc.*. Where many hits are present the number of potential tracks-in-projection becomes large, in the usual combinatorial manner. Cuts in χ^2 are applied to reject low-probability cases, especially




Figure 6.19: A cut through a muon detection *module* illustrating the reconstruction of *tracks-in-projection*. The longitudinal (z) axis of the experiment lies in the plane of the page, running from bottom to top.

The path of a muon traversing the module is shown as a solid line, and the “hits” registered in the drift tubes as crosses. Note the pair of hits in each tube: one close to the true path of the muon and one “mirror hit” on the other side of the sense wire. The dotted lines show track segments in the $x-z$ and $y-z$ projections reconstructed from the hits (see the text). (Taken from Altegoer et al., 1998a, figure 21.)

where a good (low χ^2) 3- or 4-hit track also uses one or more of the hits concerned. The tracks-in-projection surviving these cuts, some of which may be mutually exclusive, are then returned as a list.

Tracks-in-space

Where a module contains at least one track-in-projection in both $x-z$ and $y-z$ projections, all possible combinations yielding *tracks-in-space* are formed. In each case the hit position estimates are refined by correcting drift times for signal propagation time along the sense wires, on the hypothesis that this combination is the true track. No attempt is made to resolve ambiguous cases at this level, and all possible tracks-in-space are returned as a list.

6.6.2 Association of drift-chamber and muon tracks

The reconstruction process for muon chamber track segments described so far takes place during “phase one” of the event reconstruction (see section 2.5.3) and uses information from the muon chambers only, together with the global event trigger. Subsequent treatment during “phase two” of the reconstruction involves the association of these track segments to reconstructed drift chamber tracks, which are extrapolated into the muon chamber region with appropriate uncertainties on their position and direction, taking magnetic fields, energy loss, multiple scattering *etc.* into account. It will be convenient to define three levels in this process: *association*, where a track-in-projection is flagged as potentially being due to the same particle as left the drift chamber track; *good association*, where the quality of the association meets some defined higher standard; and *matching*, where the association is taken to be proven, and the drift chamber track is identified as a muon.

Association of tracks-in-projection

All processing at phase two takes place on tracks-in-projection, rather than tracks-in-space. A track-in-projection is associated to an extrapolated drift chamber track if they lie within 100 cm of each other in x (y) at a reference plane within the module; differ in slope $\frac{dx}{dz}$ ($\frac{dy}{dz}$) by less than 0.5; and if the hypothesis that the tracks are identical, taking extrapolation errors in x and $\frac{dx}{dz}$ (y and $\frac{dy}{dz}$) into account, has a χ^2 below 100. For a given drift chamber track, if more than one track-in-projection meets these criteria, only the one with the best χ^2 is associated.

As a result an extrapolated track can have at most 2 tracks-in- x and 2 tracks-in- y associated to it (one for each of the muon detection stations),

except for the case of the region at the horizontal centre of the detector where modules overlap (see figure 2.17(b)), where up to 4 track segments in each projection may be associated.

For each of the modules of the muon detection system, if there are several associated tracks-in-projection, then the “best” of these (the closest in x (y) to their respective drift chamber tracks) are flagged, and any association of the same track-in-projection to another drift chamber track is ruled out.⁴

Matching criteria

For the first muon detection station, a tighter cut of 40 cm in x (y) is used to define a good association of track-in-projection and extrapolated drift chamber track; in the second muon station, the cut is 50 cm. If an extrapolated drift chamber track has at least one good association in each projection, x and y , regardless of station, then these tracks-in-projection are flagged as *matched* to the drift chamber track, which is then identified as a muon.

Association of muon chamber hits

In parallel to the process just described, an independent procedure is used to associate individual hits (detected points-in-projection) to the extrapolation of a drift chamber track. Hits within a “road” of 40 cm around the track are associated to it, as being potentially caused by the track, regardless of their association to other tracks at the track-in-projection or hit level. There is no criterion for positive matching of hits.

6.6.3 The muon veto scintillators

For each of the four scintillation counters in the muon detection system (section 2.4.10) a discriminator threshold defining a signal above the noise level, and a time window around the global event trigger, have both been empirically set: the NOMAD data structure records the presence or absence of a signal meeting these criteria. No formal association of scintillator hits to extrapolated drift chamber tracks *etc.* is provided, although during analysis the probability for a track to leave a signal in the scintillators, were it a muon, may be estimated as described in the following section.

⁴Any association involving an “ambiguous” track-in-projection—one constructed from several of the same hits, or the corresponding mirror hits, as the flagged track-in-projection—would also be ruled out at this point.

6.6.4 Acceptance of the muon detection system

As a complement to the elements of the reconstruction code described in the previous sections—which allow the detection of muons passing through the muon chambers and scintillators—the probability for any given track to intersect the muon detection system is also routinely calculated. This allows the vetoing of drift chamber tracks as possible muons, in the case where, were a muon to have that position and momentum, it would not pass through the muon detection system.

For convenience the acceptance of the muon detection system is considered as having two independent components: its *geometric acceptance* and its *momentum acceptance*, each of which is discussed below. In both cases, calculations for a given track are based on an extrapolation of the track, assuming it to be a muon, through the back detector support (and also the iron wall in the case of the second muon station; see section 2.4.10), yielding a position (x_0, y_0) and momentum \vec{p}_0 at any given z plane, together with an appropriate covariance matrix for these quantities.

Geometric acceptance

At a given z plane, the probability density for a track to intersect the point (x, y) is approximated by the Gaussian function

$$P(x, y|x_0, y_0, \sigma_x, \sigma_y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp\left(-\frac{(x-x_0)^2}{2\sigma_x^2} - \frac{(y-y_0)^2}{2\sigma_y^2}\right) \quad (6.12)$$

where (x_0, y_0) is the extrapolated position of the track and σ_x^2, σ_y^2 the variances on the extrapolation in x and y respectively. That is, $x-y$ correlations and non-Gaussian tails due to multiple scattering *etc.* are ignored. The geometric acceptance of various elements of the muon detection system for the given track is then estimated by integrating⁵ the function P over the active area of the detector. This is done for each of the muon detection modules, keeping track of the regions where the modules overlap; the acceptance for each muon detection station is then calculated. The acceptance of the muon veto scintillators is calculated separately using a similar procedure.

⁵Since the muon chambers and scintillators are rectangular in cross-section this may be done easily using calls to the error function. It was realised during the preparation of these notes that an unnecessary approximation in the calculation causes the acceptance to be overestimated in the case where σ_x and/or σ_y is very large ($\gtrsim 1$ m). Since the acceptance is used only in the muon veto procedure, the effect of this error is conservative, rejecting more events than necessary.

Momentum acceptance

The probability for a track to reach a certain z -plane without ranging out, assuming it to be a muon, is calculated in a similar manner: given the central estimate $p_0 = |\vec{p}_0|$ for the track momentum at z , based on extrapolation through the detector support *etc.*, and its variance σ_p^2 , the probability for $p = |\vec{p}|$ to exceed zero is calculated assuming a Gaussian distribution for the probability density

$$P(p|p_0, \sigma_p) = \frac{1}{\sqrt{2\pi}\sigma_p} \exp\left(-\frac{(p - p_0)^2}{2\sigma_p^2}\right) \quad (6.13)$$

Clearly, when the central value of the momentum of the track, extrapolated, reaches $p_0 = 0$, there is still a significant probability for the track to reach the muon detection system (50%) although continued extrapolation in this case is problematic. The procedure adopted is to freeze the variances on the track position at this point, and progressively reduce p_0 to fictitious negative values as it travels through matter, using the energy loss rate for a minimum ionising particle. Together with equation (6.13) this yields sensible values for the probability for $p > 0$.

For both $p_0 > 0$ and $p_0 < 0$ cases, the standard deviation of the central momentum estimate is approximated by the simple formula

$$\sigma_p = \sigma_{\text{DC}} + \left(\frac{p_{\text{DC}} - p_0}{\Delta p}\right) \sigma_{\Delta p} \quad (6.14)$$

where p_{DC} and σ_{DC} are respectively the central momentum estimate and its standard deviation, as returned by the drift chamber track fit; Δp and $\sigma_{\Delta p}$ are the mean energy loss and its standard deviation for a particle at normal incidence reaching the muon system (determined empirically at $\Delta p = 2.3 \text{ GeV}/c$ for the first muon station and $\Delta p = 3.5 \text{ GeV}/c$ for the second). This is found to be more reliable than an estimate of σ_p built up from many increments as tracks are extrapolated through the iron.

6.6.5 Energy deposition in the calorimeters

When a drift chamber track is associated to energy deposited in either of the calorimeters (ECAL and HCAL, sections 2.4.8 and 2.4.9 above) a confidence level is constructed for the hypothesis that the energy deposited is due to a muon; where levels are constructed for both calorimeters, they are combined to yield a single confidence level. This quantity provides an additional tool for muon identification and veto.

In each case the spectrum of deposited energy for well-identified muons exhibits a single peak, with the distribution falling away monotonically for both smaller and larger energies. An appropriate confidence level P for the hypothesis that the energy E was deposited by a muon may in this case be constructed by setting

$$P(E) = \frac{\int_0^E de \frac{dN}{de}}{\int_0^{E_{\text{peak}}} de \frac{dN}{de}} \text{ for } E \leq E_{\text{peak}}, \text{ and}$$

$$= \frac{\int_E^\infty de \frac{dN}{de}}{\int_{E_{\text{peak}}}^\infty de \frac{dN}{de}} \text{ for } E > E_{\text{peak}}$$

By construction, $P(E)$ is uniformly distributed from 0 to 1 for particles following the distribution $\frac{dN}{dE}$, *i.e.* muons. (The evolution with momentum of energy deposited by muons is very slight, so a distribution $\frac{dN}{dE}$ independent of momentum is used as an approximation.) It furthermore has the meaning that among all muons with deposited energy E' on the same side of the distribution to E (*i.e.* both $E, E' < E_{\text{peak}}$ or $E, E' > E_{\text{peak}}$) a fraction $P(E)$ are further away from the peak value, and have lower likelihood, than E .

Where only one of the ECAL or HCAL deposited energies is reliable (due to serious overlap of energy deposition, or a track traversing an uninstrumented region of the other calorimeter (see section 2.4.9 and figure 2.16)) the respective quantity $P(E)$ is returned as the confidence level. Where both energies are reliable, the levels P_{ECAL} and P_{HCAL} are combined using the method of Fisher (1932) to yield a single confidence level \mathcal{P} .⁶ A cut requiring $\mathcal{P} < k$ should then retain only a fraction k of muons while accepting hadrons with energy depositions which are not muon-like.⁷

⁶During the preparation of these notes it was discovered by the author that the technique in fact used in the NOMAD reconstruction code contains an error at this point.

The confidence level \mathcal{P} may be calculated by transforming the levels $P_i \in (0, 1]$ into variables $X_i \in [0, \infty)$ using the χ^2 probability distribution for 2 degrees of freedom, and then calculating the probability of $X = \sum_i X_i = X_{\text{ECAL}} + X_{\text{HCAL}}$ following the χ^2 distribution for $\sum_i 2 = 4$ degrees of freedom. (The simple formula at (5.21) is equivalent to this procedure in the case of two levels P_i .) In the code, degrees-of-freedom of 1 and 2 respectively are used instead of 2 and 4 for the χ^2 transformations.

The result is a quantity \mathcal{P}' which will still have a uniform distribution from 0 to 1 for muons, but will not order the cases ($E_{\text{ECAL}}, E_{\text{HCAL}}$) in likelihood as expected for a confidence level. This error should therefore only affect the efficiency of a cut in \mathcal{P}' for hadrons, not its rejection power for muons. At the time of writing a study of this problem has not been undertaken, and none is foreseen (Hurst, 1999).

⁷For the distinction between this and an optimum muon/hadron discriminator see note 6 to page 248.

6.6.6 Muon decay, large angle scattering, etc.

Muons traversing the detector during the SPS flat-top provide a large sample on which to test the effectiveness of reconstruction of muons in the muon chambers and scintillators, and cuts on the calorimeter response. At the time of writing a sample of 2 million muons, within the acceptance of the muon detection system (section 6.6.4) and meeting certain quality cuts, collected during 1996 using the $V \times T_1 \times T_2$ trigger condition (section 2.5.1), have been analysed. Muon identification criteria have been applied to the single throughgoing track and the rate at which this track (taken to be a muon) escapes identification has been measured.

The muon is considered to be identified if muon chamber track segments are *matched* to the drift chamber track, if tracks-in-projection are *associated* in both projections, or if hits are associated in both projections (section 6.6.2); if there is *any* track-in-space found in modules 1 through 4 of the muon system (see figure 2.17); or if the energy deposited in both of the calorimeters is low, consistent with the deposit of a minimum ionising particle (Bueno et al., 1999). The ratio of the surviving cases to the total sample, as a function of momentum, is a measure of the rate of failure of muon identification, and a fit to the simple function

$$R(p) = \frac{a}{p + b} \quad (6.15)$$

is performed, motivated by the $1/p$ dependence expected if the effect is due to muon decay.

In the range $3 < p < 41 \text{ GeV}/c$, where good statistics are available, the data are well-described by the fit

$$R_{\text{data}}(p) = \frac{0.74 \times 10^{-3}}{p/(\text{GeV}/c) - 2.13}; \quad (6.16)$$

the χ^2 per degrees-of-freedom is 0.98. If we suppose that these events are due entirely to muon decays within some fixed longitudinal distance d_{decay} then we expect a rate

$$\begin{aligned} R &= 1 - \exp\left(\frac{-d_{\text{decay}}}{\gamma c \tau_{\mu}}\right) \\ &\approx \frac{d_{\text{decay}}}{\gamma c \tau_{\mu}} \end{aligned}$$

in the small-rate ($R \ll 1$) approximation;

$$= \frac{m_{\mu} d_{\text{decay}} / c \tau_{\mu}}{E_{\mu} - \Delta E_{\mu}} \quad (6.17)$$

where τ_μ is the proper lifetime of the μ , γ its Lorentz factor in the decay region, and E_μ and ΔE_μ respectively the muon energy in the drift chambers and the typical energy loss between the drift chambers and the decay region. Identifying equations (6.17) and (6.15) in the approximation $p_\mu = E_\mu$ yields

$$\begin{aligned} d_{\text{decay}} &= \frac{ac\tau_\mu}{m_\mu} & (6.18) \\ &= 4.6 \text{ m} \end{aligned}$$

and

$$\Delta E_\mu = 2.13 \text{ GeV} \quad (6.19)$$

given the values found in (6.16).

These values are to be tested against a picture of a “muon decay region” from (say) the end of the HCAL volume to the front face of the iron wall which ends the first muon station (see figure 2.6), since a muon decaying in this region will fail to form a track segment in the muon chambers: this picture would imply $d_{\text{decay}} = 2.5 \text{ m}$ and $\Delta E_\mu \approx 2.2 \text{ GeV}$ (taking an average across the region).

That is, the result is consistent with muon decay being a contributor to the non-detection of muons, but suggests an additional cause(s) as d_{decay} (equivalently, the rate factor a) is implausibly large. In fact, the results at (6.16) include a cut on the energy deposited by tracks in the calorimeters, not taken into account by the model used in this section; without this cut, the rate of survivors is twice as large, and the estimate of d_{decay} increases in the same manner. We must therefore seek an additional source of inefficiency in the reconstruction. One possibility is a failure to reconstruct muons in the tail of the multiple scattering distribution, which have been deflected by large angles: the scale angle for multiple scattering has a $1/p$ dependence, so such an effect would be consistent with the good fit of $R(p)$ to the form (6.15). At the time of writing a study of this problem has not been undertaken.

In any case, the function at (6.16) provides an estimate for the rate at which muons of a given momentum will pass a veto algorithm constructed from the elements described up to this point, and is therefore suitable for inclusion in the analysis to predict expected backgrounds.

6.6.7 Event structure: The transverse size ratio R_T

Finally, it is possible to suppress charged-current events relative to neutral-current and tau-decay events by applying cuts to event structure variables.

Specifically, given the total momentum \vec{p}^{tot} of the event, and n tracks at the primary vertex each with momentum $\vec{p}_j, j = 1 \dots n$, we consider the squared momentum component of a track j perpendicular to the total momentum, $p_{\perp}^2(\vec{p}_j; \vec{p}^{tot})$ and take its average value

$$\frac{1}{n} \sum_{j=1}^n p_{\perp}^2(\vec{p}_j; \vec{p}^{tot})$$

to be a measure of the “spread” of the tracks about the direction of the total momentum. Each track i is then excluded in turn from the total momentum—thus treating $(\vec{p}^{tot} - \vec{p}_i)$ as the momentum of the event—and the spread of the remaining tracks about this new direction is calculated.

The resulting *transverse size ratio*

$$R_T^{(i)} \stackrel{\text{def}}{=} \frac{\frac{1}{n-1} \sum_{j \neq i} p_{\perp}^2(\vec{p}_j; \vec{p}^{tot} - \vec{p}_i)}{\frac{1}{n} \sum_{j=1}^n p_{\perp}^2(\vec{p}_j; \vec{p}^{tot})} \quad (6.20)$$

is a measure of the isolation of track i from the rest of the event (Gangler, 1998). Where $R_T^{(i)}$ is close to 1, the spread of tracks around the total momentum is largely unaffected by the presence or absence of track i . Where $R_T^{(i)}$ is significantly smaller than 1, the spread of tracks in momentum about \vec{p}^{tot} is dominated by the i^{th} track: that is, the remaining tracks $j \neq i$ are closely associated in momentum while track i is at an angle to the group. In $\nu_{\mu} + N \rightarrow \mu^{-} + X$ and $\nu_e + N \rightarrow e^{-} + X$ events this may occur when track i is the lepton from the charged current scattering.

6.7 Choice of lepton rejection cuts

With these tools in hand we turn to the lepton-veto needs of the present analysis. The most straightforward cuts, against identified muon tracks, are presented in section 6.7.1. The seriousness of the surviving $\nu_{\mu} + N \rightarrow \mu^{-} + X$ background, and some of its kinematic properties, are then discussed in section 6.7.2. The principal muon veto procedure—a set of cuts to reject tracks which are possible muons, or unidentifiable as muons—is presented in section 6.7.3, together with the choice of tracks to which this veto should be applied. We then address the electronic backgrounds. A simple cut against identified electrons is presented in section 6.7.4; in section 6.7.5 we then describe cuts in the TRD particle discrimination variables of the previous chapter, designed to veto possible electrons, together with a cut on a preshower variable in cases where the TRD information is unavailable. Finally, cuts

in transverse size ratio variables, providing further suppression of both the muonic and electronic backgrounds, are described in section 6.7.6, where the potential further development of such cuts is also discussed. The various lepton veto procedures are then summarised in section 6.7.7.

6.7.1 Cuts against identified muons

Criteria for associating muon chamber track segments to a drift chamber track were presented in section 6.6.2: in particular, the strong association or *matching* of a track to ≥ 2 segments was discussed. An event with such a matched track is assumed to be a $\nu_\mu/\bar{\nu}_\mu$ charged current interaction, and is rejected from the analysis.

In addition to these cases we reject events with a muon which appears to come from the central detector, regardless of association to drift chamber tracks. As noted in section 6.6.1, in the case where a muon module contains track segment(s) in both projections, these are combined to form potential *tracks-in-space* for which both position (x, y, z) and direction $(p_x/p_z, p_y/p_z)$ estimates are available. The muon chambers cover a large area extending well beyond the rest of the detector, intercepting the “halo” of muons accompanying the neutrino beam (section 2.2) so it is inappropriate to veto every event with a potential track-in-space. Instead we allow such events but only after applying rather strict tests.

A muon chamber track-in-space is considered to be a halo muon if its position at the muon chambers is well outside the area of the detector,

$$|x| > 250 \text{ cm},$$

or

$$|y| > 250 \text{ cm}; \tag{6.21}$$

when extrapolated back towards the central detector, it must also remain well outside the detector, so that either

$$|x_{\text{extrap}}| > 200 \text{ cm for all } z > 0$$

or

$$|y_{\text{extrap}}| > 200 \text{ cm for all } z > 0 \tag{6.22}$$

($z = 0$ is the front-face of the detector). All other cases cause the event to be rejected.

6.7.2 The kinematic properties of surviving events

These cuts reject ν_μ CC events rather efficiently in the Monte Carlo data, with only 4% of the total sample surviving to contribute 3π candidates. To assess the effect of this remaining background we define a set of cuts which reduces the *neutral current* background to very low levels in the Monte Carlo, so that only a handful of events are expected in the data, and consider the ν_μ charged-current events which survive. A suitable set of cuts is

$$\begin{aligned} q_T &> 1.4 \text{ GeV}/c \\ M_R &< 2.2 \text{ GeV}/c \\ E^{3\pi} &> 20.0 \text{ GeV}/c, \end{aligned} \quad (6.23)$$

under which only 5×10^{-5} of neutral current events contribute a candidate. We will also use a weaker cut in the isolation variable,

$$q_T > 1.0 \text{ GeV}/c \quad (6.24)$$

together with the cuts above in M_R and $E^{3\pi}$ when it is convenient to consider a larger population of events as a cross-check.

The q_T distribution for ν_μ CC events under the M_R and $E^{3\pi}$ cuts is shown in figure 6.20. It is clear that the charged-current events provide the dominant background when q_T is large:




Figure 6.20: Distribution of the isolation variable q_T for unidentified ν_μ CC events surviving kinematic cuts. The ordinate shows the events in each bin as a fraction of all ν_μ CC interactions. The same distribution for neutral current events is shown dashed: note its much softer spectrum. The distribution for three-prong tau decays is shown dotted, with arbitrary normalisation, for comparison.

- approximately 10^{-4} have $q_T > 1.4 \text{ GeV}/c$, which is twice the rate in the neutral current;
- the q_T distribution is hard, so that more stringent cuts will not be effective without a significant loss of 3π efficiency;
- there are in any case three times as many charged- as neutral-current events.

The cut in the required mass M_R is even less effective. Figure 6.21 shows the M_R distribution for unidentified ν_μ CC events without cuts applied: it shows a clear peak at $M_R \approx 1.6 \text{ GeV}/c$ and is dominated by events below the cut at 2.2, reflecting the tau-decay-like, “back-to-back” 3π candidates produced by charged-current events (*cf.* figure 6.18 above). This distribution is largely unchanged by cuts in the isolation or 3π energy.




Figure 6.21: Distribution of the required mass M_R for unidentified ν_μ CC events, without cuts. The ordinate shows the events in each bin as a fraction of all ν_μ CC interactions. Dashed line: the same distribution for neutral current events, with an arbitrary normalisation; dotted line: for three-prong tau decays, arbitrary normalisation. Note the similarity of the ν_μ CC and tau decay distributions.

Discrimination is however available if we turn to the 3π candidate energy. Unidentified ν_μ CC events produce 3π decay candidates which are even less energetic than those from neutral current events: the distribution is strongly peaked at low energies $E^{3\pi} \approx 5 \text{ GeV}$, and the tail to large $E^{3\pi}$ is distributed as for the neutral current; this is to be compared with the very hard spectrum for three-prong tau decays, which is three times as energetic in the mean (figure 6.22: the weaker q_T cut has been chosen to yield well-populated




Figure 6.22: Distribution of the candidate energy $E^{3\pi}$ for unidentified ν_μ CC events surviving kinematic cuts. The ordinate shows the events in each bin as a fraction of all ν_μ CC interactions. There is a pronounced peak at very low energies, and a tail almost identical in distribution to the neutral current background (shown dashed). The distribution for three-prong tau decays is shown dotted, with arbitrary normalisation.

distributions). This behaviour is to be expected if we consider the argument of section 6.5 above: while misidentification of a muon as a hadron naturally produces events with a back-to-back configuration, yielding high- q_T and low required mass, there is no constraint requiring the three “pions” to be energetic. The distribution thus reflects the spread of particle energies from the hadronic system, and resembles that of neutral current events. Indeed there is a bias towards *soft* 3π candidates in that the muon, not being identified by a muon-chamber track, will typically have ranged out in the back detector support due to low momentum $p \lesssim 2.5 \text{ GeV}/c$.

The kinematic cut in $E^{3\pi}$ then, in addition to rejection of the neutral current background, is an indispensable tool in the suppression of misidentified charged current events. If it is removed completely, for example, the rate at which these events survive cuts rises to 2×10^{-3} of the total charged current sample: an increase by a factor of twenty.

The background from $\bar{\nu}_\mu$ charged currents is qualitatively similar, although much reduced due to the relatively low $\bar{\nu}_\mu$ flux, the diminished cross-section for antineutrino interactions on nucleons, and the larger average momentum of muons produced in $\bar{\nu}_\mu N$ interactions, which makes muon identification (section 6.7.1) more efficient than for the $\nu_\mu N$ case⁸. The process of 3π

⁸The reduced flux and larger average muon momentum for $\bar{\nu}_\mu N$ have a common cause,

candidate formation is somewhat different for this background—the positive charge of the muon makes it less likely to be part of the supposed 3π system than in the ν_μ CC case—so that the $E^{3\pi}$ cut is less effective, reducing the background by a factor of only $4 \sim 5$. After this cut however the background is so low, with events surviving at a rate $\approx 6 \times 10^{-5}$ and thus contributing at the one-event level for several years of NOMAD running, that further cuts could be omitted. This is clearly not the case for the ν_μ charged current, where the final background is a hundred times as large. We therefore adopt additional tools to reject the surviving events.

6.7.3 The muon veto for high- p_T tracks

The most straightforward way to improve the rejection of ν_μ charged-current events is to loosen the criteria for muon identification. The matching criteria defined in section 6.6.2 above are suitable for the positive identification of a track as being a muon, but are too restrictive for the task of rejecting possible ν_μ CC events from a sample which should contain neutral current (and possibly tau decay) events alone. Consequently we will treat certain tracks with relatively weak association to muon chamber hits or track segments as possible muons, and reject an event containing such a track. There are a number of issues to consider: what kind of association to allow; how to treat tracks which are unlikely to leave signals in the muon chambers; which tracks should be included in this “muon veto” procedure; and how to treat special-case tracks. We now treat each of these matters in turn.

in the effect of the fixed helicity of (anti-)neutrinos on their scattering properties. Considering only the (dominant) interactions on up- and down-quarks in the nucleon, the cross section for neutrino scattering, which proceeds through a $J = 0$ state, varies as

$$\frac{d\sigma}{dy} (\nu_\mu d \rightarrow \mu^- u) = \frac{G_F x s}{\pi},$$

while that for antineutrino scattering has an extra term

$$\frac{d\sigma}{dy} (\bar{\nu}_\mu u \rightarrow \mu^+ d) = \frac{G_F x s}{\pi} (1 - y)^2,$$

as the reaction proceeds through a $J = 1$ state of fixed helicity; G_F is the Fermi constant, and x and y are Bjorken’s variables. As a result the total cross-section for $\bar{\nu}_\mu$ scattering is reduced, and the mean energy fraction carried by the scattered muon, $\langle 1 - y \rangle$, is increased.

Weak association of hits or track segments

In section 6.6.2 we described the standard conditions for associating a muon chamber track segment (track-in-projection), or a single “hit”, to a drift chamber track. Due to noise hits, tracks left by halo muons which overlap an extrapolated drift chamber track in one projection, *etc.*, it has been found that association to a single track segment or hit is too loose a condition for muon veto: the solution adopted is to require an association in both x - and y -projections. We note that the association of hits and track segments is to some extent complementary: several hits are required to define a track segment, but the cut on the distance of the segment from the extrapolated drift chamber track, 100 cm, is less restrictive than the 40 cm required for the association of a hit. Accordingly we consider a track with

an associated hit in x OR an associated segment in x

AND

an associated hit in y OR an associated segment in y (6.25)

as a potential muon.

Detectability

Clearly such a test is not relevant for a track which, when extrapolated, does not intersect the muon detection system; nor do we expect a track to have associated hits or segments if it has insufficient momentum to penetrate the back detector support, even on the hypothesis that it is a muon. An additional test is therefore required to determine whether a track is within the acceptance of the muon chambers. Using the acceptance functions described in section 6.6.4 above, we define the *detectability* of a track to be

$$d \stackrel{\text{def}}{=} a_1 r_1 + (1 - a_1) a_2 r_2 \quad (6.26)$$

where

- a_i = the geometric acceptance of the i^{th} muon station ($i = 1, 2$), and
- r_i = the probability to reach the i^{th} muon station without ranging out.

The $(1 - a_1)$ term takes into account the fact that the geometric acceptance of the two muon stations is different, so that a track may be outside the acceptance of the first station, but inside the acceptance of the second, in

which case the track is detectable; conversely in the limit where $a_1 \rightarrow 1$ the second station offers redundant information only, since the probability to reach the second station $r_2 < r_1$ always. We consider a track with

$$d < 0.95 \quad (6.27)$$

to be unacceptably difficult to detect.

The gap in the first muon detection station, described in section 2.4.10, is treated as a special case (see also figures 2.6(a) and 2.17(a)). Tracks with some probability to pass between muon detection module 5, and the overlapping modules 1 & 2, clearly have a limited probability to intersect the first muon station ($a_1 < 1$), but unlike other such tracks, are in the “forward” region at relatively low angles to the neutrino beam. Muons of this kind have been found to be a dangerous source of background for hadronic decay analyses and accordingly a more stringent criterion is applied to these cases. Constructing a_{gap} as the “geometric acceptance” of the gap between the modules, using the technique of section 6.6.4, we consider a track with

$$a_{gap} > 10^{-4} \quad (6.28)$$

to have an unacceptable risk of passing undetected through the gap. The low threshold chosen reflects the use of a Gaussian probability distribution in (6.12), whereas the distribution of multiple scattering angles and displacements has larger tails.

From 1996 onwards, the gap between the modules was instrumented with scintillators, as discussed in section 2.4.10. We discard an event with a hit in one of these scintillators as a potential $\nu_\mu + N \rightarrow \mu^- + X$ event.⁹

Choice of tracks for application of the veto

The muon in a $\nu_\mu + N \rightarrow \mu^- + X$ interaction, whether or not it is identified, in general has the largest transverse momentum of all of the tracks in the event. The largest- p_T track is therefore our first choice as a muon veto track: if the largest- p_T track in an event has associated muon chamber activity as defined at (6.25), or is difficult to detect as a muon as defined at (6.27) or (6.28), we reject the event. For the standard cuts defined in the previous section, we find that the rate of surviving events is reduced by a factor of four, from 1.2×10^{-4} to 2.7×10^{-5} . An additional reduction is desirable as we

⁹Due to an error of the author, the information required to construct a_{gap} was corrupted, and the information on muon scintillator hits was unavailable, in the data files used to conduct this analysis. A partial re-processing of these files has been conducted to rectify this error. The effect on the analysis is discussed in section 7.3.2.

expect 6.7×10^5 ν_μ CC events in the 1995-96 data and more than twice this number when the analysis is extended to the full NOMAD data-taking run (1995-1998). Therefore we also apply the veto procedure to the track with second-highest p_T : the rate of surviving events then falls to 1.2×10^{-5} , so that an overall reduction of a factor of ten has been achieved. By comparison, the rate at which three-prong events survive the cuts is reduced by 14% by the leading- p_T veto, and a total of 30% for the leading- and 2nd-leading- p_T veto.




Figure 6.23: The fraction of ν_μ CC Monte Carlo events surviving standard cuts in M_R and $E^{3\pi}$ as a function of the isolation cut $q_T > q_T^{\text{cut}}$ applied. A significant reduction is obtained when the muon veto procedure is applied to the track with highest p_T (dashed line); a further reduction, for $q_T^{\text{cut}} < 2$, is obtained if, in addition, the veto is applied to the track with second-highest p_T (dotted line).

For events surviving the weaker cut defined above, $q_T > 1.0 \text{ GeV}/c$, the effect is similar. The rate of events surviving the cuts, for different values of the q_T cut, is shown in summary form in figure 6.23. We note that at very large values of the isolation q_T , the improvement due to the veto on the first-ranked track falls to a factor ≈ 2 , and the effect of the additional cut on the second-ranked track becomes negligible. A study of the Monte Carlo shows that in a small but significant fraction of $\nu_\mu + N \rightarrow \mu^- + X$ events, the muon has a relatively low ranking in transverse momentum (3rd-ranked or below) so that cuts of the general type described here can have only an indirect effect on this type of background.

Since we apply the veto procedure to tracks regardless of charge, the

$\bar{\nu}_\mu + N \rightarrow \mu^+ + X$ background is also reduced, albeit by a smaller factor (≈ 3 overall) due to the different properties under 3π candidate selection *etc.* of these events. The expected background for $\bar{\nu}_\mu$ CC is then 0.3 events for the 1995-96 data.

Special treatment of interacting tracks

For completeness we should note that tracks which are fitted at their end to a vertex, rather than leaving the drift-chamber volume, are treated as a special case, since the concept of extrapolating such a track makes little sense. Wholesale rejection of events where the high- p_T tracks have this behaviour is undesirable, as over 20% of three-prong events surviving the cuts defined above have a leading- or 2nd-ranked-in- p_T track with an end vertex.

The case where more than one track emerges from the end vertex is straightforward. These correspond chiefly to nuclear interactions in the material of the drift chamber target and as such occur for hadrons but not for muons. We therefore immediately classify tracks of this kind as hadrons. In tests on the Monte Carlo, no additional survivor is introduced by this classification in a run over 2.6 million events.

Where only one track emerges from the end vertex, the vertex can be due to a scattering or merely a break in a single track introduced during the track reconstruction. As muons are subject to both processes some test should be applied on *a priori* grounds to these cases: we therefore apply the same muon-association and detectability cuts defined above to the track *emerging* from the vertex. This incurs some loss of efficiency on three-prong decay events, as tracks emerging from such "scattering vertices" are in general low in momentum, with a limited probability to reach the muon chambers. The condition is however justified, as additional surviving charged-current events are introduced at the 10^{-5} level if this cut is not applied.

6.7.4 Cuts against identified electrons

We now turn to the choice of cuts to suppress the remaining charged current backgrounds $\nu_e + N \rightarrow e^- + X$ (ν_e CC) and $\bar{\nu}_e + N \rightarrow e^+ + X$ ($\bar{\nu}_e$ CC). As with the muonic backgrounds we begin by rejecting events with identified electrons. However in this case we must make a distinction between electrons fitted or merged to the primary interaction vertex, and other electrons due to the conversion of photons in the detector. Three-prong tau decays are subject to this process in the same way as other classes of event, so that we will impose an unacceptable (and arbitrary) efficiency loss by rejecting all

events where electrons are present. We therefore restrict this cut to identified electrons which are attached to the primary vertex.

The standard reconstruction algorithm of the NOMAD experiment identifies a track as an electron if

- it has a TRD pion acceptance $\epsilon_\pi < 10^{-3}$ (section 5.2.3);
- it is flagged as an electron by the maximum-likelihood algorithm, described in section 5.2.7, for handling overlapping tracks;
- it is flagged as an electron by the preshower reconstruction, in cases where the TRD information is unavailable, by a condition with 90% efficiency for recognising electrons, *and* the probability of thereby misidentifying a pion as an electron is less than 10%,

together with some special cases applying to V^0 tracks which do not concern us here. We extend this scheme slightly by “propagating” an electron identification backwards along any chain of tracks and vertices, so that a track with an identified electron emerging from its end vertex is also considered to be an identified electron for the purpose of this analysis; and so on back along the chain, if there is more than one intermediate vertex. If any such track is fitted or merged to the primary vertex, then the event is rejected.

Both of the electronic backgrounds pass the cuts at (6.23) at a significant rate, even after these identified electrons have been rejected, and the effect of the muon veto cuts defined in the previous section are minor (figure 6.24). Although the ν_e and $\bar{\nu}_e$ fluxes are low, the survival rate for these events is seventy times that for the corresponding ν_μ and $\bar{\nu}_\mu$ -induced events (*cf.* figure 6.23) so that the electronic backgrounds would be equal contributors to the final background passing all cuts, if no additional electron-rejection measures were employed.

6.7.5 The electron veto for high- p_T tracks

The electron-identification criteria described in the previous section were drawn up with the $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ analysis in mind, and are therefore intended to provide a relatively pure electron sample. We have the opposite requirements and seek a way of identifying electrons with high efficiency so that the corresponding events may be discarded. The technique for using the TRD information in this way was described in the previous chapter (section 5.2.4); we will use this information to define a suitably loose criterion to identify a track as a possible electron. Proceeding by analogy with the muonic case, we define the level of the cut, the selection of tracks to which it will be applied, and the handling of special cases as follows.

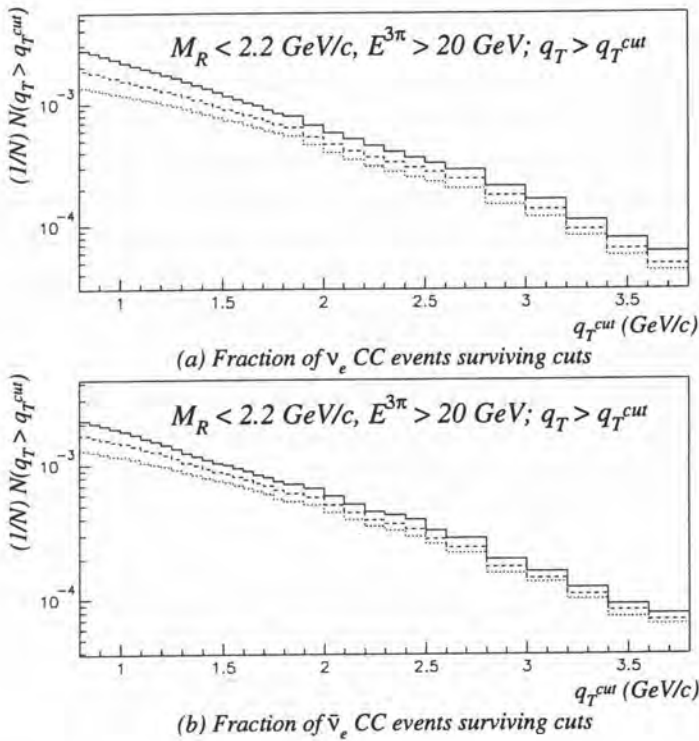


Figure 6.24: The fraction of (a) ν_e CC and (b) $\bar{\nu}_e$ CC Monte Carlo events surviving standard cuts in M_R and $E^{3\pi}$ as a function of the isolation cut $q_T > q_T^{\text{cut}}$ applied. The effect of the muon veto procedure applied to the leading- p_T track (dashed line), and in addition the track with second-highest p_T (dotted line), are also shown.

Choice of cut in the TRD electron acceptance

We noted in section 5.2.4 that the criterion $\epsilon_\pi < 10^{-3}$ identifies electrons with an efficiency of $\approx 90\%$. The remaining cases exhibit a distribution of the electron acceptance variable $\tilde{\epsilon}_e$ from zero up to a cutoff of order 0.1. (In any given case this depends on the exact efficiency with which electrons of that momentum p and number of TRD hits N_h fail to be identified: this varies, in both directions, by a factor of a few about the 0.1 value.) We expect pions, by contrast, to be concentrated at the lowest values of $\tilde{\epsilon}_e$: examining the acceptance function tables we find that in the ideal case of tracks with nine associated TRD hits, more than 99% of pions have $\tilde{\epsilon}_e$ values of 0.01 or lower. A cut at this level would therefore have a marginal effect on the efficiency with which three-prong decay events pass the cuts, while imposing an additional rejection factor of order 10 ($= \frac{0.010}{0.10}$) on electrons leading in

transverse momentum.

A more careful examination of the Monte Carlo events confirms this line of argument. The leading- p_T track in three-prong decay events surviving other cuts exhibits electron acceptance values strongly peaked at $\tilde{\epsilon}_e < 0.01$, with a small tail to ≈ 0.05 (figure 6.25(a)); those in ν_e and $\bar{\nu}_e$ CC events show a much diminished peak, due to cases where a pion is first-ranked in p_T , together with a larger sample of cases, due to electrons first-ranked in p_T , where $\tilde{\epsilon}_e$ is distributed out to 0.10 and larger values (figure 6.25(b)). These cases are efficiently removed by a cut $\tilde{\epsilon}_e < 0.010$: the rejection does not undergo significant further improvement until values $\tilde{\epsilon}_e \lesssim 0.002$ are approached: there is then a significant loss of efficiency in the three-prong decay sample.

We therefore set the level of the TRD cut to $\tilde{\epsilon}_e < 0.010$ in all that follows.




Figure 6.25: The TRD electron/pion discrimination variable $\tilde{\epsilon}_e$ for the leading p_T track in (a) three-prong decay and (b) ν_e CC Monte Carlo events, after muon veto, rejection of identified electrons, and kinematic cuts $q_T > 1.4 \text{ GeV}/c$, $M_R < 2.2 \text{ GeV}/c$, $E^{3\pi} > 20 \text{ GeV}$. Events shown as a fraction of the sample for which TRD information is available: see the text.

Choice of tracks for application of the veto

As for the muonic background, we apply this $\tilde{\epsilon}_e < 0.010$ cut in the first instance to the track with the largest p_T in the event. If the track is outside the geometric acceptance of the TRD, so that $\tilde{\epsilon}_e$ may not be calculated, we also reject the event. The effect on the number of events expected to survive the other cuts, normalised to the 1995-1996 data (see the discussion in the next chapter, section 7.2.1) is shown in figure 6.26. The background is reduced by a factor ≈ 3.4 , for a corresponding reduction in the efficiency for three-prong decays of 11%. Extending the veto to the second-ranked track in p_T improves the background rejection to ≈ 4.6 , while the three-prong efficiency is reduced by 18% in total. The decision between these two cases depends on the overall background level and we defer the treatment of this question to following sections, when additional cuts are in place.

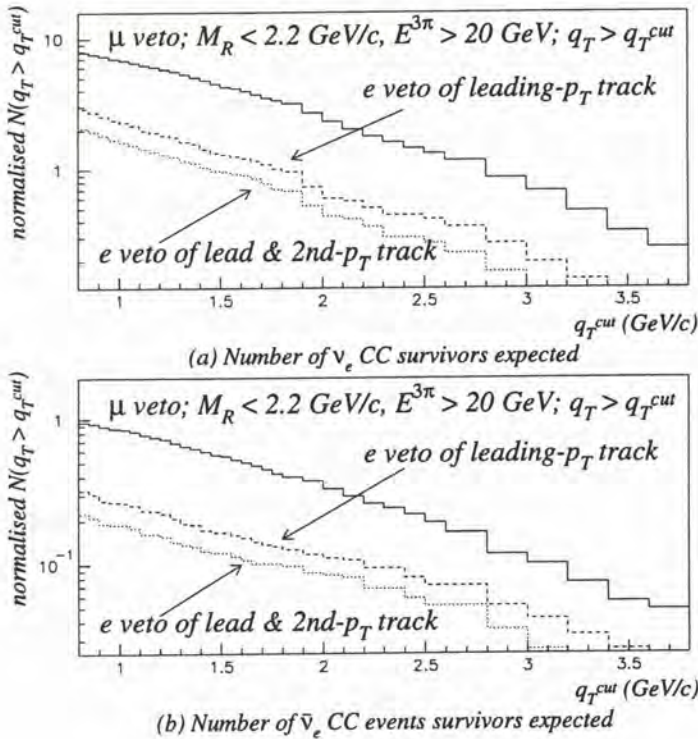


Figure 6.26: The number of expected (a) ν_e CC and (b) $\bar{\nu}_e$ CC surviving events in the 1995-1996 data, as a function of the isolation cut $q_T > q_T^{cut}$ applied. The effect of the TRD electron acceptance cut applied to the leading- p_T track (dashed line), and in addition to the track with second-highest p_T (dotted line), are also shown.

Treatment of special cases

We noted in the previous chapter that electron/pion discrimination for tracks with four or more shared hits in the TRD region is not carried out by the likelihood ratio method of section 5.2: instead, a maximum likelihood method is used to choose between four hypotheses for the overlapping tracks (two electrons, e & π , π & e or two pions; see section 5.2.7). This procedure is unsuitable for use in electron veto algorithms as the probability for mis-identifying an electron as a pion is at the 10% level, is a function of the momentum of both the examined track and the overlapping track, and (by the nature of the method) may not be tuned (Bassompierre et al., 1998b).

Some treatment of such tracks is however required as 15–20% of signal events surviving cuts have four or more shared hits on the leading track (a consequence of the kinematic cuts, which select energetic events with tracks bunched into the forward region by the Lorentz boost of the centre of mass); a blanket rejection of these events is undesirable. Acceptance of all such cases is also unsuitable as this would increase the rate of surviving events in the ν_e and $\bar{\nu}_e$ charged-current samples by $\approx 50\%$.

As an expedient, we use the deposition of energy in the preshower to discriminate electrons from pions in these cases. Based on the exposure of the prototype of the preshower to electron and pion test beams, three momentum-dependent thresholds on the combined response of the horizontal and vertical tubes have been defined, identifying electrons above 1 GeV/ c with efficiencies of 85%, 90% and 95% respectively (Schmid, 1996). These efficiencies are well-reproduced in both the NOMAD Monte Carlo, and samples of electrons selected from the data by strict cuts in the TRD and the electromagnetic calorimeter. The threshold with 95% efficiency is suitable for our purpose: by requiring that the tracks in question *not* be identified as electrons by this criterion, we impose a rejection factor of 20 on electrons, independent of momentum. The rate of rejection of pions is in general a strong function of momentum, but is about 10% on average for high-transverse-momentum tracks in three-prong decay events surviving other cuts, an acceptable loss of efficiency since we are already dealing with a subset of the signal.

As with the definition of muon-veto cuts for tracks in section 6.24 above, we must also determine how to proceed when a track ends on a fitted vertex in the drift chamber volume. By analogy with that case, if more than one track emerges from the vertex we classify the incoming track immediately as a hadron; if only one track emerges from the vertex, we apply the electron veto cuts to that track.

6.7.6 Choice of cuts on the transverse size ratio

While these cuts are implemented to reject the electronic backgrounds, they also have some effect on the ν_μ CC background, largely as a result of the condition requiring the leading- p_T track *etc.* to be within the acceptance of the TRD, a more restrictive condition than requiring that the track be within the geometric acceptance at the muon chambers. The survival rate is shown in figure 6.27: for some values of the isolation cut q_T^{cut} the improvement is as large as a factor of two. We should note however that for a sample of order 10^6 in size, the survival rate is still significant, and some of the remaining events have such large values of the isolation variable, $q_T > 2 \text{ GeV}/c$, that further kinematic cuts against these events would impose an unacceptable loss of efficiency for the three-prong decays. Furthermore, as noted above, extending the muon veto procedure to additional tracks is ineffective against these cases.

For the final reduction in the background we rely instead on the event structure variable introduced in section 6.6.7. For a given track i the *transverse size ratio* R_T^i defined at (6.20) provides a measure of the extent to which the transverse size or “width” of the event is dominated by the inclusion of the track i , with smaller values of R_T^i indicating a track set apart




Figure 6.27: The effect of the electron track veto on the rate at which ν_μ CC events pass the cuts, as a function of q_T^{cut} . The kinematic cuts shown, and the muon track veto for the leading- and 2nd-ranked track in p_T have been applied: the addition of the electron veto on the leading track yields the dashed line; if also applied to the 2nd-ranked track, the dotted line.

in momentum space from the remainder: possibly the lepton produced in a deep inelastic scatter. In order to implement a cut on this quantity we must first determine how to choose the potential lepton from the various tracks in the event, and then choose the level of the cut in R_T^i . The simple options considered to date for this analysis are the following:

Finding the minimum of R_T^i over negative tracks

In this case, noting that the ν_μ CC background is the most serious, followed by the ν_e CC background, we choose as our candidate lepton the track i with the smallest value R_T^i among the *negative* tracks at the primary vertex. The resulting transverse size ratio

$$R_T^- \stackrel{\text{def}}{=} \min\{R_T^i | q(i) = -1\} \quad (6.29)$$

has a broad distribution with a peak at ≈ 0.9 for three-prong decays, and an almost identical distribution for the neutral current background (figure 6.28(a)). That is, the transverse size of the event is largely unchanged if any single negative track is excluded. For these cases, in general there is no single track which dominates the transverse size of the event: for the neutral current interactions, all tracks are produced in the fragmentation of a single hadronic system; for the three-prong decays, while the three-pion *combination* from the tau decay is usually isolated from the remaining tracks, an *individual* decay pion will not be isolated in the same way, not least because its companion decay pions are included in the remainder of the event. This remains the case when lepton vetos and kinematic cuts are applied: remarkably, even under very strong cuts in the 3π isolation q_T , the R_T^- distribution is little altered, and remains neutral-current-like.

In the ν_μ and ν_e charged-current backgrounds, on the other hand, the most likely value for the transverse size ratio is $R_T^- \lesssim 0.2$, with the remainder of the distribution falling off steadily so that only a handful of cases have R_T^- near unity (figures 6.29(a) and (b)). That is, the transverse size of the event is dominated by the grouping of the lepton from the charged current scattering with the hadronic system: when this track is excluded, the measured transverse size falls dramatically, being due only to the width of the hadronic system. No such effect is seen for the $\bar{\nu}_\mu$ and $\bar{\nu}_e$ backgrounds, where the produced lepton is positive: the distribution of R_T^- in these cases remains neutral-current-like (figures 6.29(c) and (d)).

As the 3π candidate isolation q_T rises, in the ν_μ (ν_e) CC background, the “narrowing” of the event with the removal of the lepton becomes increasingly pronounced: the R_T^- distribution becomes more concave to the origin and its




Figure 6.28: The transverse size ratio R_T^- for three-prong decays. (a) R_T^- for all events with selected 3π , without kinematic cuts or vetoes; the very similar distribution for neutral current events is also shown, dashed. (b) R_T^- for three-prong decays under standard kinematic cuts $q_T > 1.4 \text{ GeV}/c$, $M_R < 2.2 \text{ GeV}/c$, $E^{3\pi} > 20 \text{ GeV}$, with muon- and electron-vetos also applied. The same distribution for a much stronger isolation cut $q_T > 2.2 \text{ GeV}/c$ is also shown, dotted.

mean falls dramatically, with events at very high isolation $q_T > 2.2 \text{ GeV}/c$ having values $R_T^- < 0.2$ in most cases (figure 6.30(d)). This is to be contrasted with the behaviour of the three-prong decay events, where R_T^- is largely unchanged under q_T cuts, as noted above (figure 6.28(b)). It is clear that to take full advantage of this behaviour, a q_T -dependent cut on R_T^- is indicated, determined either empirically or (ideally) set by a two-dimensional likelihood function. For the present we choose a simple cut in R_T^- , and aim to retain some rejection power against ν_μ CC at intermediate values of q_T , where the ν_μ CC background is highest. A cut requiring

$$R_T^- > 0.5 \quad (6.30)$$

is suitable for this purpose.

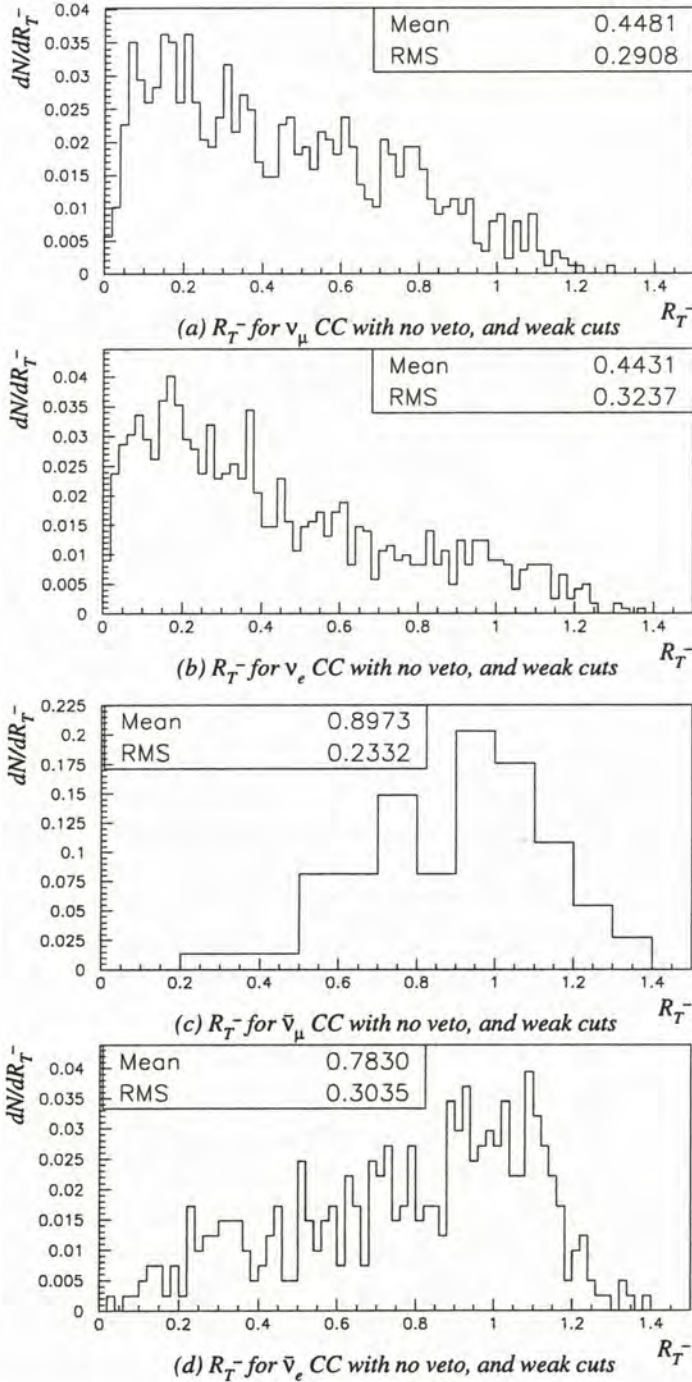


Figure 6.29: The transverse size ratio R_T^- for the charged-current backgrounds without muon- and electron-track vetos; kinematic cuts $q_T > 1.0$ GeV/c, $M_R < 2.2$ GeV/c, $E^{3\pi} > 20$ GeV have been applied. For the (a) ν_μ CC and (b) ν_e CC backgrounds, R_T^- has a falling distribution, concentrated below 0.5, while for the (c) $\bar{\nu}_\mu$ CC and (d) $\bar{\nu}_e$ CC backgrounds, where the produced lepton is positively charged, the R_T^- distribution is neutral-current-like.

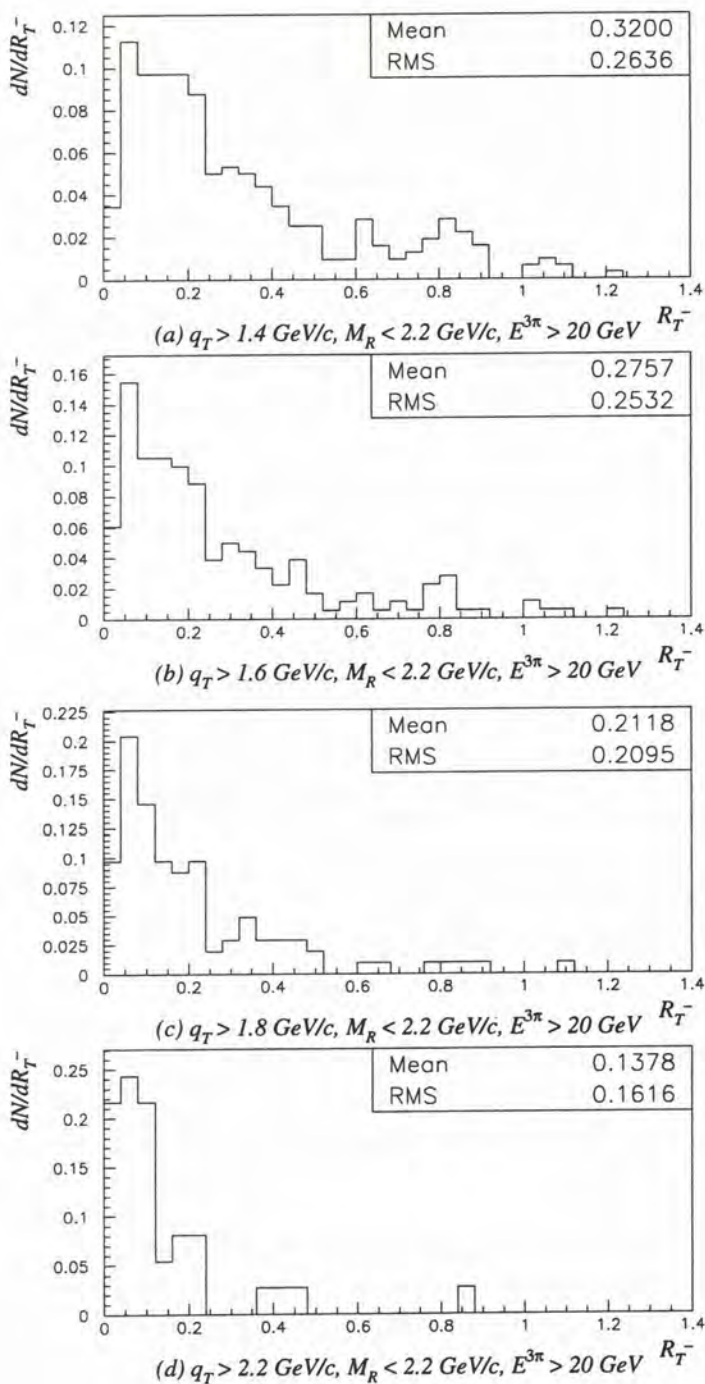


Figure 6.30: The transverse size ratio R_T^- for the ν_μ charged-current background as the cut in the isolation variable q_T is made more strict. No muon veto has been applied; cuts $M_R < 2.2 \text{ GeV}/c$, $E^{3\pi} > 20 \text{ GeV}$ are imposed throughout. At very high q_T , R_T^- is strongly concentrated at low values $R_T^- < 0.2$.

Minimising R_T^i over *all* tracks

As an alternative we may take the more even-handed approach of choosing the candidate lepton from *all* of the tracks at the primary vertex: distributions of the resulting quantity

$$R_T^{\text{all}} \stackrel{\text{def}}{=} \min\{R_T^i\} \quad (6.31)$$

for three-prong tau decays and neutral current interactions are shown in figure 6.31(a). Clearly the distribution must be shifted towards zero, as by construction $R_T^{\text{all}} \leq R_T^-$, but for the three-prong decays there is in fact a

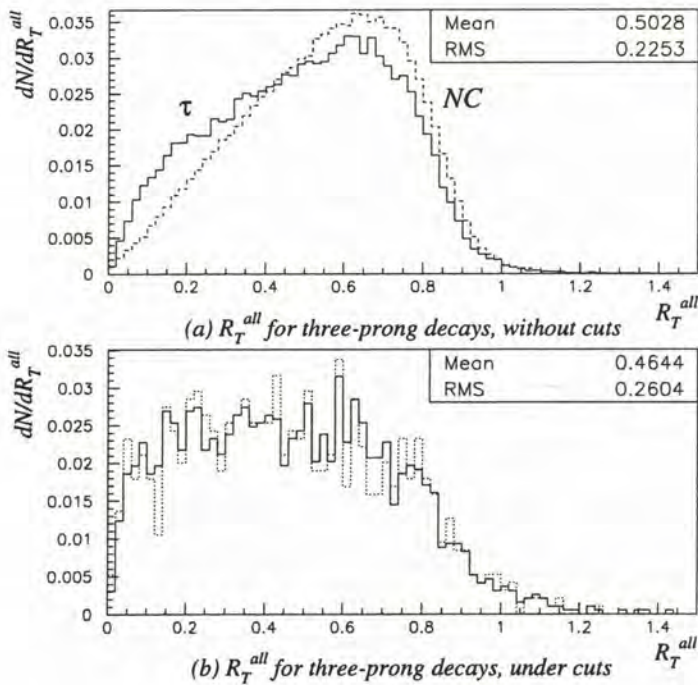


Figure 6.31: The transverse size ratio R_T^{all} for three-prong decays. (a) R_T^{all} for all events with selected 3π , without kinematic cuts or vetoes; the similar distribution for neutral current events is also shown, dashed. (b) R_T^{all} for three-prong decays under standard kinematic cuts $q_T > 1.4 \text{ GeV}/c$, $M_R < 2.2 \text{ GeV}/c$, $E^{3\pi} > 20 \text{ GeV}$, with muon- and electron-vetos also applied. Note the change in the distribution at low R_T^{all} forming a plateau. This is an effect of removing soft and/or low- q_T 3π systems: for a much more stringent isolation cut, $q_T > 2.2 \text{ GeV}/c$ (shown dotted) there is no further change to the distribution.

significant change in the distribution, largely due to cases where one positive track dominates the hadronic system: such events are superficially $\bar{\nu}_\mu$ -CC-like, and yield low values of R_T^{all} . As low-energy and low- q_T events are removed by kinematic cuts, this effect becomes more marked, with the R_T^{all} distribution becoming plateau-like, although it does not evolve further as the isolation cut q_T^{cut} is raised to higher values (figure 6.31(b)). In order to retain efficiency for three-prong decays, any cut in R_T^{all} will have to be set at a low value.

Turning to the charged-current backgrounds, we see a falling distribution of R_T^{all} in each case, including the $\bar{\nu}_\mu$ - and $\bar{\nu}_e$ -induced events (figure 6.32). As the isolation cut q_T^{cut} increases, each of these distributions evolve as for the case previously shown (figure 6.30), concentrating at values below 0.2. Bearing in mind also the problem of three-prong decay efficiency, then, a suitable cut in this variable is

$$R_T^{\text{all}} > 0.2 \quad (6.32)$$

Choice of cut between these options

Based on the Monte Carlo samples, the choice between these cuts depends on the relative importance of the different backgrounds. Both have a dramatic (and near-identical) effect on the $\nu_\mu + N \rightarrow \mu^- + X$ background at high q_T , reducing the level of these otherwise robust surviving events by a factor as large as five (figure 6.33(a)). For intermediate cuts $1.0 \lesssim q_T^{\text{cut}} \lesssim 1.6$, where the typical R_T^i values for the lepton are less concentrated towards zero, the more stringent cut $R_T^- > 0.5$ achieves a higher background rejection; by contrast the cut in R_T^{all} has little effect on the $\bar{\nu}_\mu + N \rightarrow \mu^+ + X$ background (figure 6.33(b)) as expected from the previous discussion. The level of the $\bar{\nu}_\mu$ CC background is already sufficiently low (0.2 events with $q_T > 1.4 \text{ GeV}/c$ expected in the 1995-96 data under the cuts already defined) that this is not a significant obstacle to the use of the R_T^- cut, however it would be a drawback with larger statistics (such as the full 1995-1998 dataset) if the total background were reduced to the $\lesssim 1$ event level by other cuts.

The effect on the electronic backgrounds is more modest, with the effect on the number of events surviving cuts never exceeding 30% (figure 6.34). By contrast with the muonic backgrounds, where the track veto and R_T cuts are to some extent complementary (achieving their best rejection of events at intermediate and high- q_T respectively) the track veto for electrons and the R_T cuts are largely redundant, as shown in figure 6.35(a). Taking the level of survivors after the muon track veto alone as a reference, we see that the reduction due to an $R_T^- > 0.5$ cut, and that due to a TRD cut on the

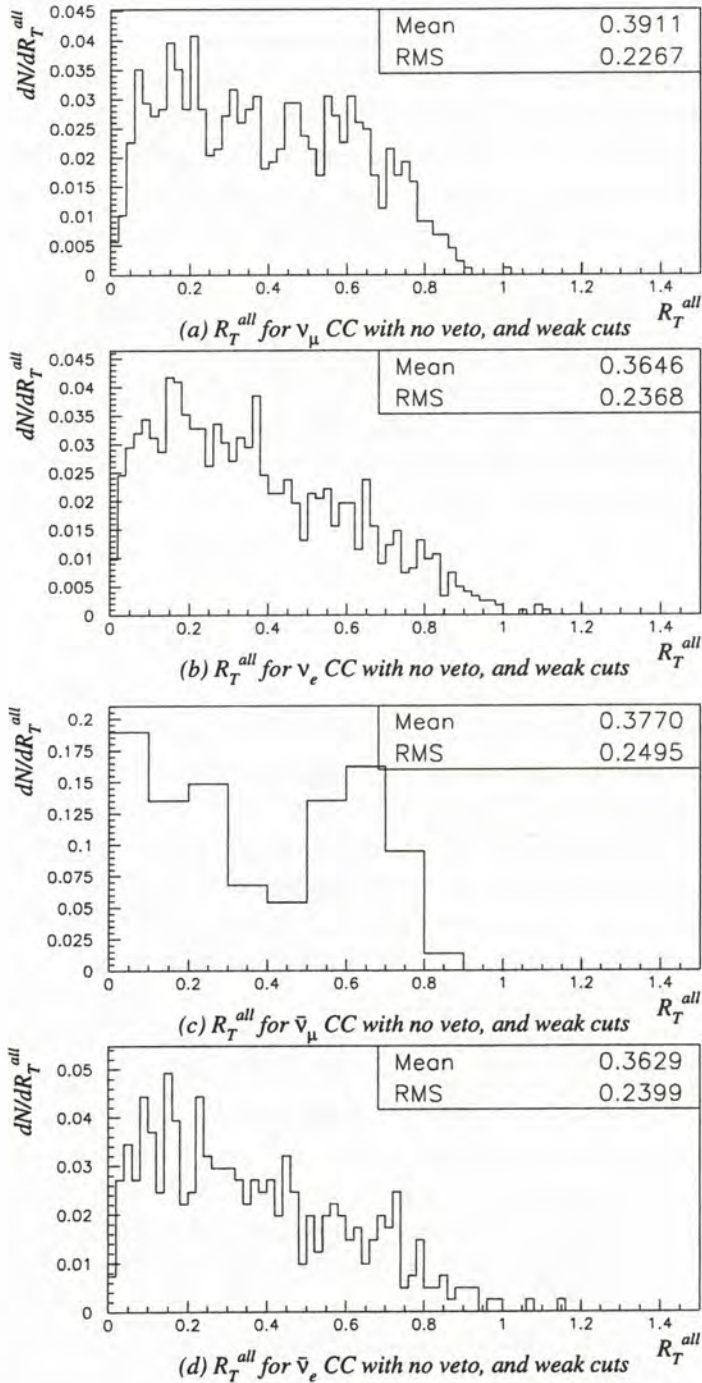


Figure 6.32: The transverse size ratio R_T^{all} for the charged-current backgrounds without muon- and electron-track vetos; kinematic cuts $q_T > 1.0 \text{ GeV}/c$, $M_R < 2.2 \text{ GeV}/c$, $E^{3\pi} > 20 \text{ GeV}$ have been applied. In this case each of the backgrounds exhibits a falling distribution.

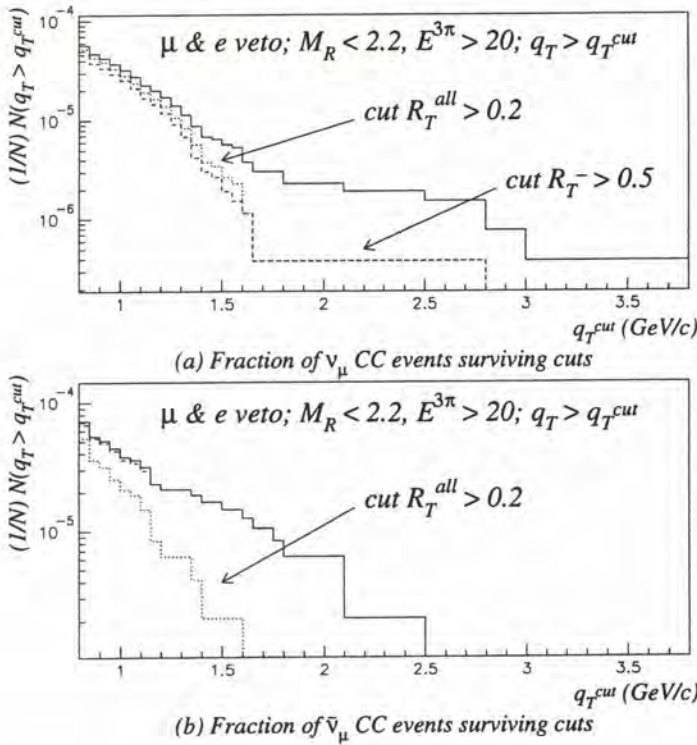


Figure 6.33: The effect of transverse size ratio cuts on the (a) ν_μ CC and (b) $\bar{\nu}_\mu$ CC backgrounds: the rate of events surviving the muon and electron track vetos is shown with a solid line; the rate after an additional cut $R_T^- > 0.5$ is shown dashed, the rate after $R_T^{\text{all}} > 0.2$ dotted. Note that the R_T^- cut has a negligible effect on the $\bar{\nu}_\mu$ background.

leading- p_T track, are very similar. A more detailed comparison confirms this effect, with the tail in electron acceptance $\tilde{\epsilon}_e$ shown in figure 6.25 for the ν_e CC Monte Carlo largely disappearing under the $R_T^- > 0.5$ cut. In the less important $\bar{\nu}_e$ background, the effect of the TRD cut is significantly stronger (figure 6.35(b)), a situation duplicated even for the R_T^{all} cut until the very highest values of q_T are approached.

In summary, then, if we include a transverse size ratio cut, as is desirable for reducing the ν_μ CC background at large q_T , there is a case for weakening the electron track veto previously defined (section 6.7.5). It would be possible to omit this track veto entirely, but this leads to the unsatisfactory situation where $\bar{\nu}_e$ charged-current events contribute to the final background at the event level, despite the very low (and rather poorly understood) $\bar{\nu}_e$ flux at

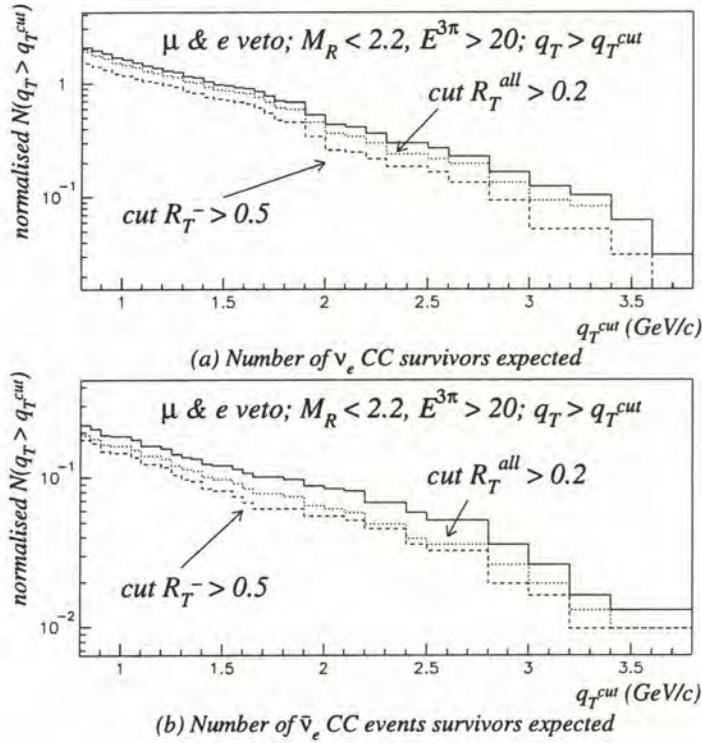


Figure 6.34: The effect of transverse size ratio cuts on the (a) ν_e CC and (b) $\bar{\nu}_e$ CC backgrounds: the number of events surviving the muon and electron track vetos is shown with a solid line; the number after an additional cut $R_T^- > 0.5$ is shown dashed, the number after $R_T^{\text{all}} > 0.2$ dotted. In each case there is a small additional improvement. (Normalised to the 1995-96 data.)

the NOMAD. There is also an additional reduction of $\gtrsim 30\%$ in the ν_e CC background if the electron track veto is retained, and (surprisingly) a similar reduction for the ν_μ CC due to the acceptance effect previously discussed; the corresponding loss of three-prong decay efficiency is 7%.

If we retain the electron track veto, then, suitable alternatives are presented by the following two schemes:

1. the TRD cut applied to the leading- p_T track only (leading e veto) and a cut $R_T^- > 0.5$;
2. the TRD cut applied to both the leading- and 2nd-leading- p_T track (strong e veto) and a cut $R_T^{\text{all}} > 0.2$.

For the electronic backgrounds, the effect on the number of surviving events

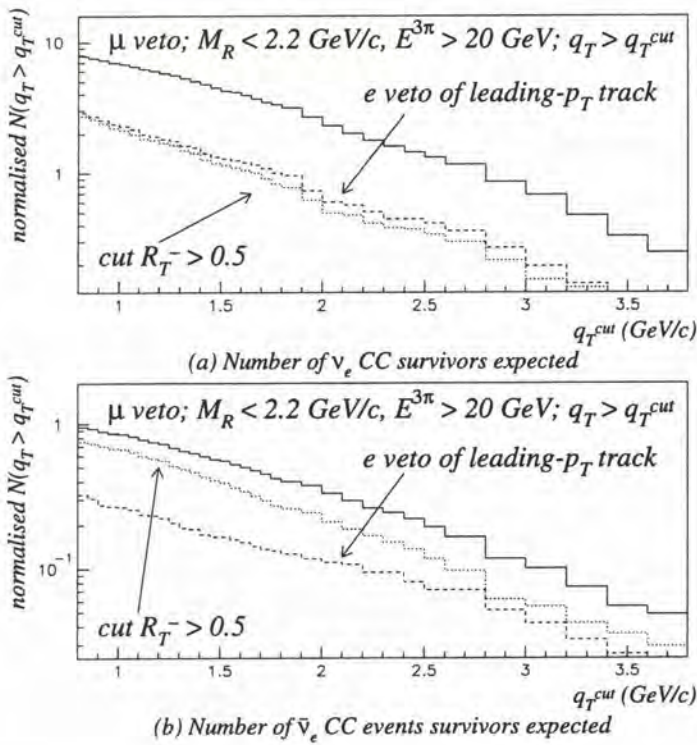


Figure 6.35: Comparison of the effect of electron veto and R_T^- cuts on the (a) ν_e CC and (b) $\bar{\nu}_e$ CC backgrounds: the number of events surviving the muon veto only is shown with a solid line; the number after an electron veto on the leading- p_T track is shown dashed, the number after a cut $R_T^- > 0.5$, with no electron veto, is shown dotted. The behaviour is comparable in (a) but the R_T^- cut has only a minor effect on the $\bar{\nu}_e$ (b). (Normalised to the 1995-96 data.)

is shown in figure 6.36: it is very similar for the two schemes. The first relies on the R_T^- cut to suppress ν_μ/ν_e backgrounds at intermediate and high- q_T , and the electron track veto to suppress the $\bar{\nu}_\mu$ background; the second uses the electron track veto to suppress the electronic backgrounds, and the R_T^{all} cut to suppress all of the charged-current backgrounds, $\bar{\nu}_\mu$ included, at high q_T .

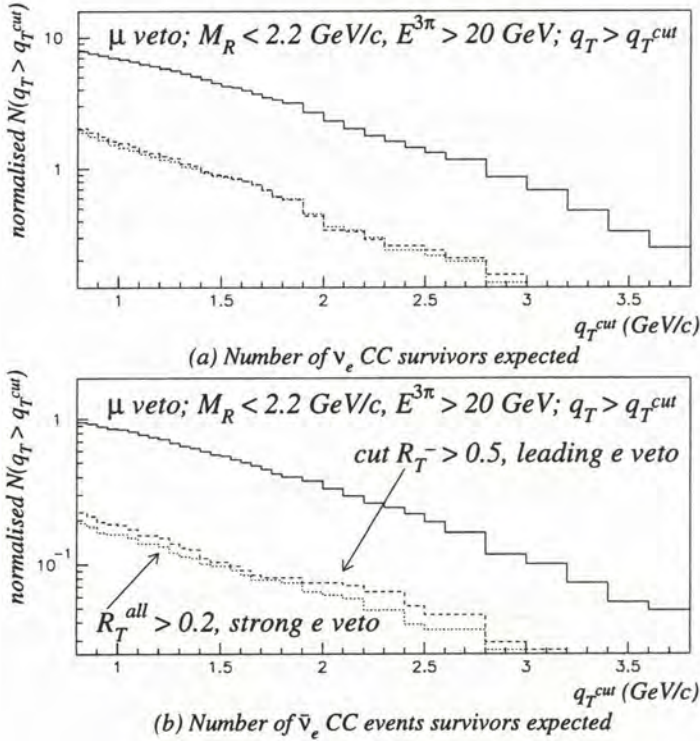


Figure 6.36: Comparison of two lepton-veto schemes on the (a) ν_e CC and (b) $\bar{\nu}_e$ CC backgrounds: the number of events surviving the muon veto only is shown with a solid line; the number after an $R_T^- > 0.5$ cut and an electron veto on the leading- p_T track is shown dashed; the number after a cut $R_T^{\text{all}} > 0.2$, and a “strong” electron veto, *i.e.* applying to the leading- and second-leading p_T track, is shown dotted. The behaviour is comparable in both cases. (Normalised to the 1995-96 data.)

6.7.7 Summary: The lepton-veto cuts

In summary, we reject as a $\nu_\mu/\bar{\nu}_\mu$ charged-current event any event with

- a drift-chamber track matched to muon-chamber track segments in x - and y -projections;
- a muon chamber track-in-space which points close to the central detector volume;
- hits in the muon-detection scintillator counters; or
- the leading or second-leading- p_T track either
 1. associated to a muon chamber hit or track segment in each projection;
 2. with a probability $d < 0.95$ of intersecting the muon detection system, if it were a muon; or
 3. having a probability $> 10^{-4}$ to cross the gap in the first muon detection station,

with special cases treated as described in section 6.7.3.

Events with an identified electron at the primary vertex are rejected as $\nu_e/\bar{\nu}_e$ background. We define additional cuts in the TRD electron acceptance $\tilde{\epsilon}_e < 0.010$ to reject electrons, and in transverse size ratio variables, to suppress both muonic and electronic backgrounds, in the combination

1. the TRD cut applied to the leading- p_T track only (leading e veto) and a cut $R_T^- > 0.5$; or alternatively
2. the TRD cut applied to both the leading- and 2nd-leading- p_T track (strong e veto) and a cut $R_T^{\text{all}} > 0.2$.

Special cases are treated as described in section 6.7.5. The choice between these two schemes will be made based on a study of the behaviour of the cuts in the data, to which we now turn.

Chapter 7

The search for three-prong tau decays

7.1 Introduction

With techniques to reject both neutral- and charged-current backgrounds in hand, we now turn to a first application of the analysis to the NOMAD data.

A number of preliminary matters are treated in section 7.2: the simulated event samples corresponding to the neutral current and each of the charged-current background processes are normalised to correspond to the 1995-96 NOMAD data; we then define a set of cuts which reduce these backgrounds to acceptable levels. In section 7.3 we compare distributions of variables in the Monte Carlo with the data, and compare the number of events surviving the cuts, in a search for the charge-conjugated process $\tau^+ \rightarrow \bar{\nu}_\tau \pi^+ \pi^- \pi^+ + n\pi^0 (n \geq 0)$. Since we do not anticipate that any τ^+ are present in the data, the number of events passing the cuts provides an estimate of the combined background due to neutral- and charged-current processes. A comparison of the results of this “positives search” on the data and Monte Carlo thus serves to test how well the backgrounds are currently understood.

In section 7.4 we make the same comparisons for the true three-prong decay search, outside the region of the parameter space where tau decays would be clearly visible above the background, as an additional test. Up to this point we have not looked at the data in this region, which may be considered as a sealed box containing any potential signal, to prevent bias being introduced into the choice of variables, cuts *etc.* by a knowledge of their effect in the data. We finally “open the box”, and determine the number of events in the data surviving all cuts, in section 7.4.3. These results are then used to provide an estimate of the rate of $\nu_\mu \rightarrow \nu_\tau$ oscillation in a simple

two-flavour oscillation model in section 7.5. A discussion of the results, and the potential for future development of this analysis, follows in section 7.6.

7.2 Definition of cuts on the normalised Monte Carlo events

In the preceding chapters we have dealt with the neutral current and the various charged current backgrounds in turn, defining variables to discriminate against neutral current events, and veto conditions to reject charged current events. We now seek a set of cuts which will reduce the *total* background to an acceptable level, while retaining some efficiency for three-prong tau decays. To do this it is first necessary to normalise the various Monte Carlo samples to the data, so that both the total rate of events and the relative proportion of neutral current, ν_μ CC *etc.* events correspond to those expected from our understanding of the WANF beamline and NOMAD acceptance. This normalisation is described in section 7.2.1. We then proceed to choose the variables on which cuts are to be applied, and to choose the values of the cuts. After discussing a possible source of systematic error associated with this procedure (section 7.2.2), we present the cuts, together with the distributions of the relevant variables in the Monte Carlo, in section 7.2.3.

7.2.1 Normalisation of the simulated events

The determination of the rate of neutral current, ν_e charged current, and other events in the data, compared to the rate for the dominant ν_μ charged current process, is a study in its own right, requiring a dedicated analysis. For the purposes of the present work we rely instead on the relative proportion of interactions determined by the NOMAD simulation of the WANF beamline, described in section 2.2.3 above. The resulting normalisation of the simulated event samples to the data proceeds in several stages, as follows.

Normalisation of the $\nu_\mu + N \rightarrow \mu^- + X$ Monte Carlo to the data

In setting the total number of each class of event expected in a given period, we use the $\nu_\mu + N \rightarrow \mu^- + X$ sample as the reference since the majority of these events have an identified μ^- providing a well-defined tag. We take the number of events

- with a fitted primary vertex in the fiducial volume,
- with a matched negative muon fitted to the primary vertex,

Event sample	N_{vert} fitted vertices in the fiducial volume	$N_{\mu-ID}$ events with a μ^- identified and fitted
“4-module” data	70481	22915
“8-module” data	80811	34173
“11-module” data	79574	38938
total 1995 data	230866	96026
total 1996 data	440949	219117
TOTAL 1995-96 DATA	671815	315143
ν_μ CC simulation	2618899	2106080
normalisation factor $r_{\nu_\mu CC}$		0.1496

Table 7.1: Normalisation of the ν_μ CC sample to the 1995-96 data.

- with at least one other track fitted to the primary vertex, and
- passing the $p_{Fermi} < 1.0 \text{ GeV}/c$ cut (in the case of the Monte Carlo; see section 6.4.3)

as a measure of the sample size for both the data and the ν_μ CC Monte Carlo, and adjust the Monte Carlo accordingly. The resulting values are shown in table 7.1. Note that at this stage we have made no correction for varying acceptance *etc.* in the 4- and 8-module data, nor for events passing the cuts in the data which are due not to ν_μ CC but to meson decay and other processes. In both cases the effect on the final estimate should be small.

Normalisation of the minor charged-current samples

The remaining charged-current samples were then normalised by using the number of events

- with a fitted primary vertex in the fiducial volume, and
- passing the $p_{Fermi} < 1.0 \text{ GeV}/c$ cut

as a measure of the sample size: its ratio with the corresponding number for the ν_μ charged current was corrected to match the ratio of charged-current interactions predicted by the simulation of the neutrino beam (section 2.2.3). The ν_μ CC normalisation was then used to provide an absolute normalisation to the data. The resulting normalisation factors are shown in table 7.2. The samples used correspond to 35 ($\bar{\nu}_\mu$), 97 (ν_e) and 308 ($\bar{\nu}_e$) times the

Sample	N_{vert}	$r_{vert} = \frac{N_{vert}}{N_{\nu\mu CC}}$	$r_{pred}^{CC} = \frac{N_{pred}^{CC}}{N_{pred}^{\nu\mu CC}}$	$r_{norm} = \frac{r_{pred}}{r_{vert}} \times r_{\nu\mu CC}$
ν_μ CC	2618899	1.000	1.000	0.1496
test ν_μ CC	694352	0.265	1.000	0.564
$\bar{\nu}_\mu$ CC	475526	0.1816	0.0339	0.0279
test $\bar{\nu}_\mu$ CC	172314	0.0658	0.0339	0.0771
ν_e CC	559918	0.2138	0.0148	0.01036
$\bar{\nu}_e$ CC	230953	0.0882	0.00191	3.24×10^{-3}

Table 7.2: Relative normalisation of the charged-current Monte Carlo samples, and the final normalisation to the data. The columns show for each sample respectively the number of vertices in the fiducial volume, the ratio of this number to the ν_μ CC, the expected ratio based on the beamline simulation (section 2.2.3), and the overall normalisation factor applied. The role of the “test” sub-samples is described in the following section.

expected number of events in the data, so that the contribution of Monte Carlo statistics to the uncertainty on the predicted background from these processes will be small.

Normalisation of the neutral current sample

Finally, to normalise the neutral current simulation to the data, it was first necessary to set the expected ratio of neutral- to charged-current interactions. Using the relations of Llewellyn-Smith (1983)

$$R_\nu = \rho^2 \left(\frac{1}{2} - \sin^2 \theta_W + \frac{5}{9} \sin^4 \theta_W (1 + r) \right) \quad (7.1)$$

$$R_{\bar{\nu}} = \rho^2 \left(\frac{1}{2} - \sin^2 \theta_W + \frac{5}{9} \sin^4 \theta_W \left(1 + \frac{1}{r} \right) \right) \quad (7.2)$$

we set the ratio of neutral- to charged-current interactions strengths to $\rho = 1$, use the world-average for the Weinberg angle (in the on-shell convention) $\sin^2 \theta_W = 0.226$ (Caso et al., 1998, section 10), and take the fraction of $\bar{\nu}$ to ν CC interactions on nucleons to be $r = \frac{0.34}{0.67} = 0.507$ (Hikasa et al., 1992); this yields

$$R_\nu = \frac{\sigma(\nu_l + N \rightarrow \nu_l + X)}{\sigma(\nu_l + N \rightarrow l^- + X)} = 0.317 \quad (7.3)$$

Neutrino species	r_{pred}^{CC}	$R_\nu(R_{\bar{\nu}})$	$r_{pred}^{NC} = r_{pred}^{CC} \times \frac{R}{R_\nu}$
ν_μ	1.000	0.317	1.000
$\bar{\nu}_\mu$	0.0339	0.358	0.0383
ν_e	0.0148	0.317	0.0148
$\bar{\nu}_e$	0.00191	0.358	0.0022
correction r_{corr}^{NC}	1.055		
effective R_ν	0.334		

Table 7.3: Correction of the ν_μ neutral current normalisation in order to take $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$ neutral current processes into account.

Quantity	test ν_μ NC	ν_μ NC
N_{vert}	320162	1060321
$r_{vert} = \frac{N_{vert}}{N_{\nu_\mu}^{CC}}$	0.1223	0.405
R_ν	0.317	0.317
r_{corr}^{NC}	1.055	1.055
$r_{\nu_\mu}^{CC}$	0.1496	0.1496
$r_{norm} = \frac{R_\nu}{r_{vert}} \times r_{corr}^{NC} \times r_{\nu_\mu}^{CC}$	0.409	0.1236

Table 7.4: Relative normalisation of the ν_μ neutral-current Monte Carlo sample, and the final normalisation to the data. The figures shown are respectively the number of vertices in the fiducial volume, the ratio of this number to the ν_μ CC sample, the expected ratio, the correction due to other neutrino species, the normalisation of the ν_μ CC, and the overall normalisation factor for neutral currents. The role of the “test” sub-sample is described in section 7.2.2.

for the neutrino species, and

$$R_{\bar{\nu}} = \frac{\sigma(\bar{\nu}_l + N \rightarrow \bar{\nu}_l + X)}{\sigma(\bar{\nu}_l + N \rightarrow l^- + X)} = 0.358 \quad (7.4)$$

for the antineutrinos.

Separate runs on simulated $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$ neutral current event samples have not yet been carried out, so these processes were taken into account by incrementing the ν_μ NC sample by the appropriate number of events. Clearly this does not give the correct radial distribution of vertices for these processes, and in the case of the ν_e and $\bar{\nu}_e$ samples the interactions will be insufficiently energetic in the mean (section 2.2.3). As shown in table 7.3,

the overall correction to the expected neutral current interaction rate is only 5.5% in this approximation so any additional corrections due to the differing radial- and energy-spectra should be small.

The normalisation of the neutral current sample to the data is then shown in table 7.4. The Monte Carlo statistics used correspond to eight times the expected number of interactions.

7.2.2 A problem in the estimation of efficiency of cuts

Both the choice of cuts against background, and estimates of the efficiency with which background and three-prong tau decay events pass the cuts, have been made based on the normalised samples just described. In this situation there is a risk of incorrectly estimating the efficiency with which a class of events passes the cuts. Where a variable is continuous, the decision on where to place the cut is based on the rejection it imposes on the background sample B , and the efficiency with which events from the signal sample S pass the cut: an estimate based on the effect of the cut on the samples B and S is in each case biased by the preceding decision.

For example, supposing that events from the background B pass some cut requiring $x > x_0$ at an (unknown) small rate b , corresponding to only 3 events in the Monte Carlo sample being considered. However due to a statistical fluctuation only 1 event in fact passes the cut. This influences the decision to place the cut at this value rather than some more strict value $x > x_1 = x_0 + \delta$, and the cut is considered to be more effective at reducing the background than it is in fact. Upward fluctuations of course also occur, so that in some situations estimates of this kind are too pessimistic. However their effect is to deter placing the cut at the value x_0 , so that it is placed at some more strict value x_1 where the fluctuation is reduced or removed. A bias towards over-estimating the rejection power of cuts against the background is thus introduced.

In a similar way, an upward fluctuation in the number of events from the signal sample passing a cut will influence the decision to place the cut at that value, introducing a bias towards over-estimating the rate at which signal events pass the cut. Where a “cut” is discrete (such as the decision on how many tracks should be subjected to muon- or electron-veto conditions) there is likewise a bias due to fluctuations in the effect of the cut on signal and background events.

This problem has been addressed as follows: for the three-prong decay sample, and the ν_μ neutral-current and ν_μ charged-current backgrounds, decisions have been based on approximately one-third of the final Monte Carlo sample in each case: these are the “test” sub-samples included in the nor-

malisation tables 7.2 and 7.4 above for completeness. The estimate of the rate at which events pass the cuts has then been made using the full sample. The effect of statistical fluctuations on the decision-making process is in fact increased by this process (the relative size of fluctuations in the test samples will be greater by a factor $\sqrt{3} \approx 1.7$ in the mean) but the estimate of the effect of the actual cuts chosen is more reliable.

For the ν_e and $\bar{\nu}_e$ charged current backgrounds, the Monte Carlo samples used were in both cases a very large multiple of the expected sample in the data, and the number of surviving events in the full samples were of order 100, so that statistical fluctuations are expected at the 10% level: the full samples were therefore used throughout. For the most part the $\bar{\nu}_\mu$ charged-current sample was treated as a variation of the main ν_μ CC, with no additional decisions needing to be taken: the exception being studies of the transverse size ratio cuts described in section 6.7.6 above, where the relative size of the $\bar{\nu}_\mu$ background was important to establish: a test sample comprising $\approx 40\%$ of the total was used.

For the remainder of this chapter, all efficiencies quoted and distributions shown will be for the full Monte Carlo samples, to avoid the unnecessary duplication of quoting test- and full-sample efficiencies. It should therefore be stressed that all *decisions*, including those of the previous chapter, were taken based on the test samples.

7.2.3 The choice of variables and cuts against the total background

In the previous chapter (section 6.7) we defined nominal cuts against the neutral current events, *requiring*

1. $q_T > 1.4 \text{ GeV}/c$ for the isolation,
2. $M_R < 2.2 \text{ GeV}/c$ for the required mass, and
3. $E^{3\pi} > 20.0 \text{ GeV}/c$ for the energy of the 3π candidate.

Given these cuts, we set additional conditions to suppress the $\nu_\mu/\bar{\nu}_\mu$ charged-current backgrounds which remained. An event is *rejected* if

4. any drift chamber track is *matched* to track segments in the muon chambers (section 6.6.2);
5. a muon chamber *track-in-space* points close to the central detector volume (section 6.7.1);

6. the highest- or second-highest-ranked track in p_T satisfies any of the conditions
 - (a) in both x - and y -projections, a muon chamber hit or track segment is associated to the track (loose muon identification), or
 - (b) the detectability criterion d defined at (6.26) is less than 0.95 (track not detectable as a muon), or
 - (c) the probability for the track to pass through the gap in the first muon station, assessed using the method of section 6.6.4, is greater than 10^{-4} ;

tracks with end-vertices being treated according to the prescription of section 6.7.3.

To suppress the electronic backgrounds, the event is *rejected* if

7. any track fitted to the primary vertex, or its daughter, is identified as an electron (section 6.7.4); or
8. either the highest- or second-highest-ranked track in p_T has a TRD electron acceptance $\tilde{\epsilon}_e > 0.01$, with special cases treated as defined in section 6.7.5: note that this implies that the event is rejected if the tested track is outside the TRD acceptance.

To suppress all of the charged-current backgrounds, particularly in the high- q_T region, a final cut on the event structure is made, *requiring*

9. $R_T^{\text{all}} > 0.2$, where R_T^{all} is the minimum value of the transverse size ratio R_T^i over all tracks i at the primary vertex (sections 6.6.7 and 6.7.6).

In (8) and (9) we have chosen for definiteness the second of the two options canvassed in section 6.7.6.

These cuts were determined by treating the neutral- and charged-current backgrounds in isolation. Applying them to the combined background, as normalised in section 7.2.1, we find that 5.6 events are expected to survive. By comparison, 2.4% of three-prong tau decays survive these cuts: for oscillations occurring at the previous experimental limit, $\sin^2 2\theta = 5 \times 10^{-3}$ at large Δm^2 (Ushida et al., 1986), 1.7 three-prong decay events would be expected; for the strongest limit currently published, set by the combined NOMAD analysis of several tau decay modes, $\sin^2 2\theta = 1.2 \times 10^{-3}$ (Astier et al., 1999a), 0.40 events would be expected.¹ Any signal would not therefore be visible above the background. A more stringent q_T cut yields better

¹The method for calculating the expected number of tau events in the case of oscillations is discussed in section 7.5.1.

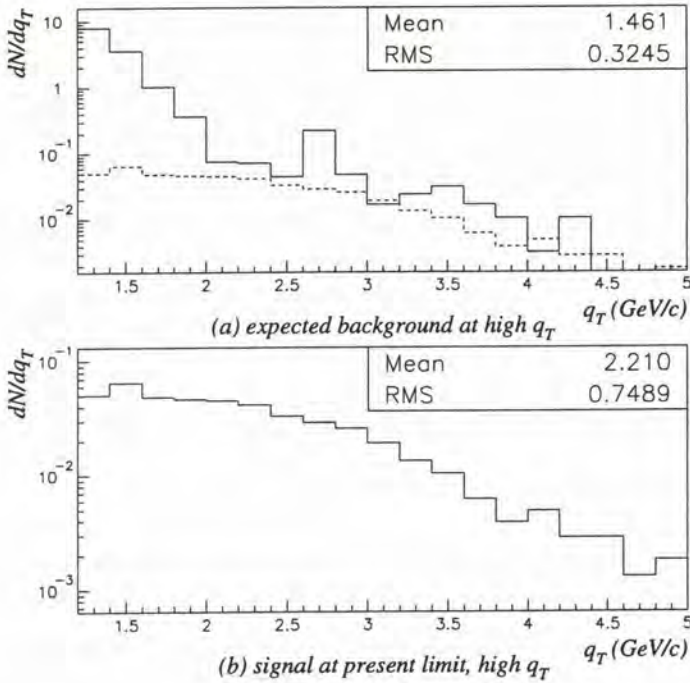


Figure 7.1: The expected total background in the high- q_T region for 1995-96 data statistics (a); the expected tau signal if oscillations occurred at the present published limit is superimposed, dotted, and shown separately at (b).

performance, with an expected number of observed three-prong decays 1.2 and 0.3 respectively, above a background expectation of ≈ 1 .

The situation is not of course as simple as just presented:

- the final analysis would take place using the full 1995–1998 dataset, which is more than twice as large as the 1995-96 dataset used here;
- in practice this analysis would be combined with analyses of other decay modes $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$, $\tau^- \rightarrow \nu_\tau \pi^-$, $\tau^- \rightarrow \nu_\tau \pi^- \pi^0$, with the oscillation level or limit being set by the four results taken together;
- to properly exploit the differing q_T distribution of the expected background and any three-prong decay signal (figure 7.1) we would carry out the analysis in several bins in q_T , the number of surviving events with (say) $2.0 < q_T < 3.5$ GeV/ c being significantly more sensitive to the presence of a τ signal than events with $1.4 < q_T < 2.0$ GeV/ c or events in the extreme tail $q_T > 3.5$ GeV/ c .

Notwithstanding these points, the sensitivity of the analysis to oscillations should clearly be improved if possible. Apart from the more fundamental improvements canvassed below (section 7.6) the values of the cuts (1) through (3), (6b) and (6c), (8) and (9) are in many cases rather arbitrary, and should in principle be tuned. Prior to testing the prediction of the background in the data, however, we will treat the cuts (2) through (9) as fixed, and allow the value of the isolation cut (1) to vary if needed. More subtle changes to the cuts may be made, if warranted, at a later stage.

7.3 A search for three-prong τ^+ decays using an equivalent analysis

We may gauge the effect of these cuts in the data, without prejudicing the final result, by conducting a search for the charge-conjugated decay $\tau^+ \rightarrow \bar{\nu}_\tau \pi^+ \pi^- \pi^+ + n\pi^0 (n \geq 0)$. The flux of $\bar{\nu}_\mu$ is sufficiently low, with 1.3×10^4 charged-current events expected in the 1995-96 data as normalised above, that even if $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ oscillations occurred at the previous experimental limit, we would expect only two three-prong τ^+ decays to occur in the data. Given the efficiency of the analysis we would not therefore expect to see any signal, even in this most optimistic case; and the present NOMAD limit is in any case more stringent by a factor of three.

It turns out that a τ^- signal, if it exists, may also contribute to the apparent $\tau^+ \rightarrow \bar{\nu}_\tau \pi^+ \pi^- \pi^+ + n\pi^0 (n \geq 0)$ rate. The dominant process is for a $\pi^- \pi^+$ pair from three-prong τ^- decay to be grouped together with a π^+ from the hadronic system, yielding a $\pi^+ \pi^- \pi^+$ combination which in some cases approximates the kinematics of the τ^- and has acceptable invariant mass. (This may be compared with $\pi^- \pi^+ \pi^-$ combinations with ≥ 1 “wrong pion” being selected, and passing cuts, in τ^- decay events; section 4.6.) Recent simulations have shown that such an “induced” τ^+ signal would be reduced by a factor of 10 with respect to the true τ^- signal (DiLella, 1999): $\nu_\mu \rightarrow \nu_\tau$ oscillations at the level of the previous experimental limit would therefore induce a negligible $\pi^+ \pi^- \pi^+$ signal in the 1995-96 data.

We may therefore use the number of surviving events in a “positives search” as a measure of the background, to gauge how well the backgrounds have been modelled in the Monte Carlo, and to test whether the analysis is robust against the inevitable differences between simulated and actual events. As a first test (section 7.3.1) we compare distributions of significant variables in the Monte Carlo and the data, with the veto conditions (4) through (9) in place, but without applying the cuts (1), (2) and (3). In section 7.3.2 we then

compare the rate at which events survive the full set of cuts, as the isolation cut (1) is varied, for the simulation and the data. Certain differences are found, and a correction to the simulation is defined in section 7.3.3 as a means of taking these differences into account. A further means of correcting for differences between simulation and data is briefly discussed in section 7.3.4.

7.3.1 Distributions for Monte Carlo and data positives

Beginning with the most important analysis variable, the isolation q_T , we are immediately confronted with significant differences between the simulation and the data. Based on figure 7.2, we may identify two key discrepancies, as follows:

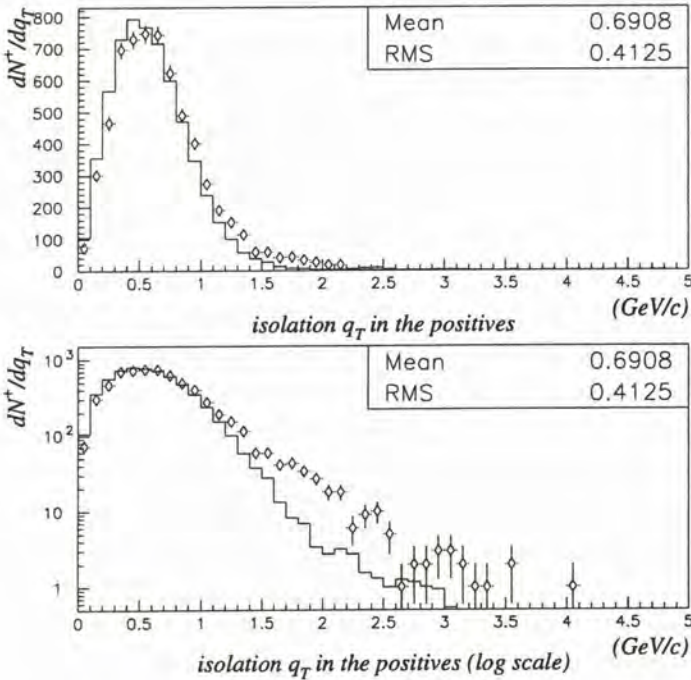


Figure 7.2: The isolation variable q_T in the normalised Monte Carlo (histogram) and the data positives (points with Poissonian error bars). Statistics refer to the data. Events with identified muons and electrons have been rejected, and the vetos against muons and electrons for the leading- and 2nd-leading p_T tracks have also been applied; no kinematic cuts have been imposed. The isolation of the background is underestimated in the Monte Carlo: there is also a significant additional population at the highest q_T values.

1. The isolation is in general underestimated in the simulation, an effect visible as a horizontal displacement between the otherwise-similar data and Monte Carlo distributions (linear scale, upper plot). This effect is seen in all NOMAD analyses and is imperfectly understood, although poor simulation of the details of fragmentation at low Q^2 , nuclear reinteraction effects, and the omission of some subtleties of drift-chamber response from the Monte Carlo have all been suggested. In any case, a technique to at least partially correct this problem exists, and is discussed below in section 7.3.4.
2. There is a significant additional background at large values $q_T > 1.3 \text{ GeV}/c$ (say) for which mere underestimation of q_T appears an insufficient explanation: for some values the number of survivors in the data is in excess by factors of four or five. We note in addition that the q_T -underestimation effect is in the first instance a problem of the neutral current background, and that the expected level of this background at very large $q_T \gtrsim 2.0 \text{ GeV}/c$ is insignificant, while the observed data excess is prominent in this region. This suggests that the excess is due to charged-current events which have not been properly controlled.

A consistent picture emerges if we survey the other important variables of the analysis (figures 7.3 through 7.7). In general the agreement between the simulation and the data is good. Among the four remaining summary kinematic variables defined in chapter 3 the candidate energy $E^{3\pi}$ is overestimated in the Monte Carlo (figure 7.3(b)), which may have the same underlying cause as the underestimate of q_T : hadronic systems are more “open”, with particles having more momentum perpendicular to the system axis, in the data than in the simulation. On the other hand, the small but significant excess of events with $0.25 < \rho^{3\pi} < 0.50$ and $M_R \approx 2.3 \text{ GeV}/c$, visible above distributions which are otherwise remarkably well-reproduced, is consistent with an additional charged-current background, these events typically having required mass values in the signal region (*cf.* figure 6.21) and large transverse momentum fraction $\rho^{3\pi}$.

Turning to the 3π structure variables, the invariant mass distribution is well-predicted for the higher values $1.1 < M_{inv}^{3\pi} < 1.5 \text{ GeV}/c^2$, while there is a small but consistent excess at lower masses $0.8 < M_{inv}^{3\pi} < 1.1 \text{ GeV}/c^2$, not presently understood (figure 7.4). Of the remaining quantities, the two-pion mass-squared variables s_1 and s_2 are very well predicted by the simulation (figure 7.5(a) and (b)). Most of the distribution of the scale-free opening angle $\gamma^{3\pi}\theta_{\max}$ is similarly in close agreement, although there is a significant excess at the peak, representing the lowest values readily attainable,

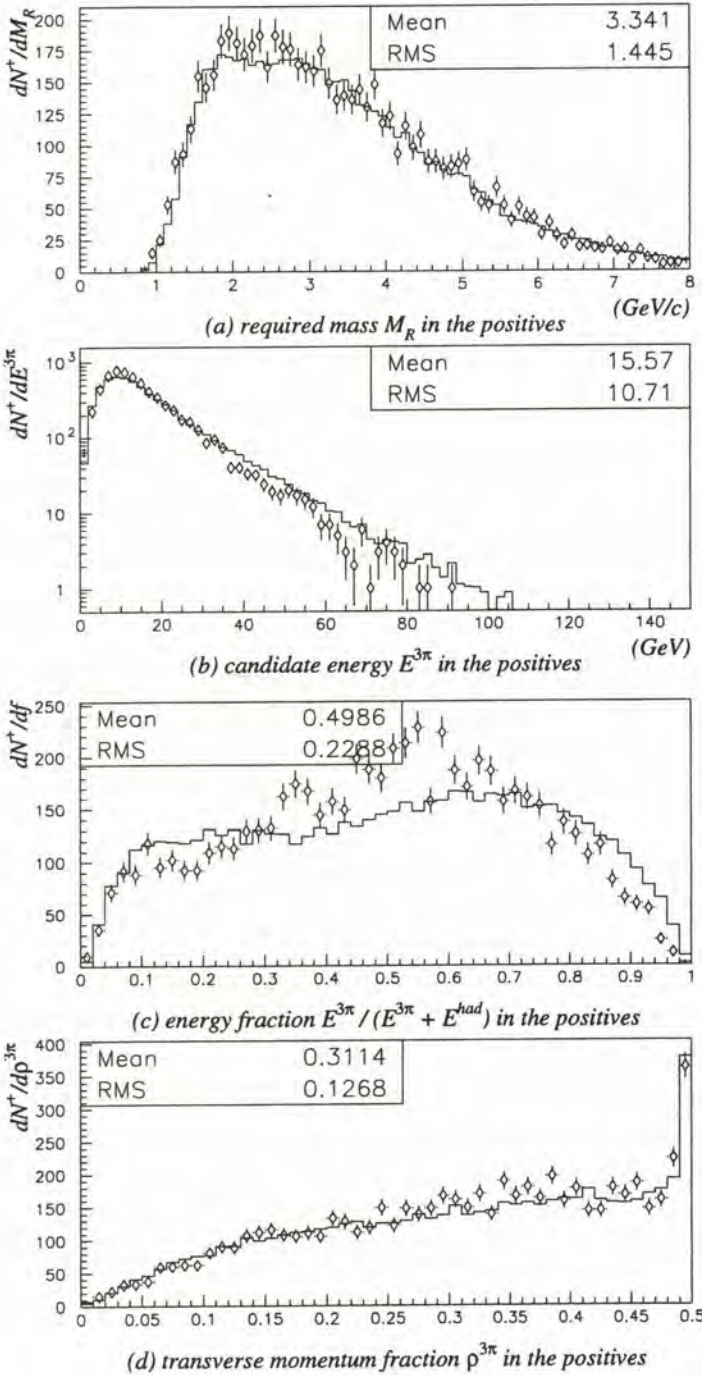


Figure 7.3: The four remaining summary kinematic variables in the Monte Carlo (histogram) and data positives (points with Poissonian error bars). The agreement across the distribution is good for (a) the required mass and (d) the 3π transverse momentum fraction. The candidate energy (b) is overestimated in the Monte Carlo, reflected also in the 3π energy fraction (c) where the agreement in the shape of the distribution is poor. The small excess in the required mass distribution about the peak $M_R \approx 2.3$ GeV/c is consistent with an additional charged-current background; see the text.




Figure 7.4: The 3π invariant mass $M_{inv}^{3\pi}$ in the Monte Carlo and data positives. There is a small excess of candidates at low masses in the data.

$\gamma^{3\pi}\theta_{\max} \gtrsim 1$ (figure 7.5(c)). An examination of the processes giving rise to this distribution (appendix C.2) shows that for low 3π masses $M_{inv}^{3\pi}$, leading to limited pion momenta \vec{p}^* in the 3π rest frame, $\gamma^{3\pi}\theta_{\max}$ is restricted to these low values. The excess is thus consistent with that seen in the invariant mass.

Finally, the isolation angle θ_{isol} is seen to be significantly overestimated in the Monte Carlo, with a clear excess of cases in the data where the 3π and hadronic systems overlap ($\theta_{\text{isol}} < 0$), and a deficit of cases where all of the 3π tracks are well-isolated from the hadronic system ($\theta_{\text{isol}} > 0.3$; figure 7.6). This may reflect the same differences in the hadronic system between the simulation and the data as have been seen in the q_T and $E^{3\pi}$ distributions.

The most significant evidence supporting the interpretation of an additional charged-current background ((2) above) is seen in the distributions of the transverse size ratio variables (figure 7.7). For both R_T^- and R_T^{all} , the distribution for $R_T > 0.3$ is very well predicted, while an excess of events is clearly visible in the $R_T < 0.3$ region where the charged-current backgrounds are concentrated.

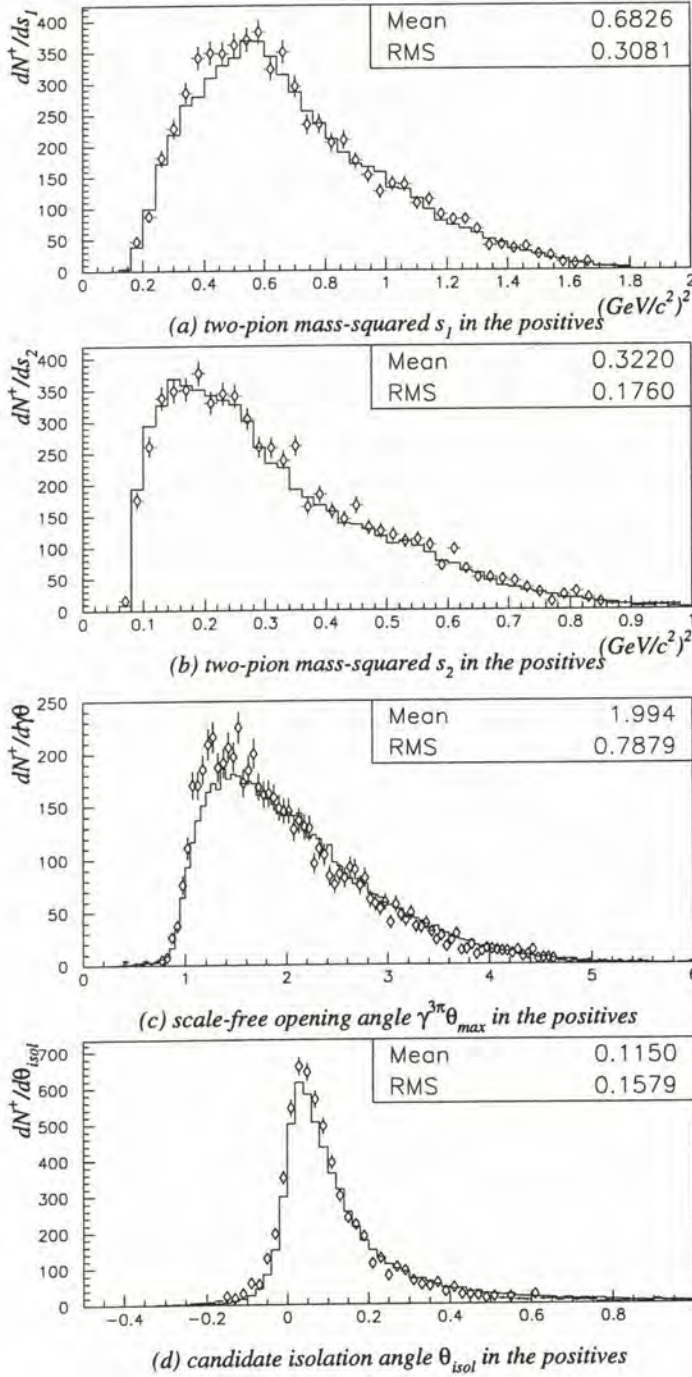


Figure 7.5: The remaining 3π structure variables in the Monte Carlo (histogram) and data positives (points with Poissonian error bars). The agreement is good for the two-pion mass-squared variables (a) s_1 and (b) s_2 . The scale-free opening angle (c) is overestimated in the Monte Carlo, consistent with the underestimate of the invariant mass $M_{inv}^{3\pi}$; see the text. The isolation angle (d) is overestimated in the Monte Carlo: see also figure 7.6 and the discussion in the text.




Figure 7.6: The candidate isolation angle θ_{isol} in the Monte Carlo and data positives, shown with a logarithmic scale. The Monte Carlo predicts 3π systems whose tracks are systematically more isolated from the rest of the event than those observed in the data.

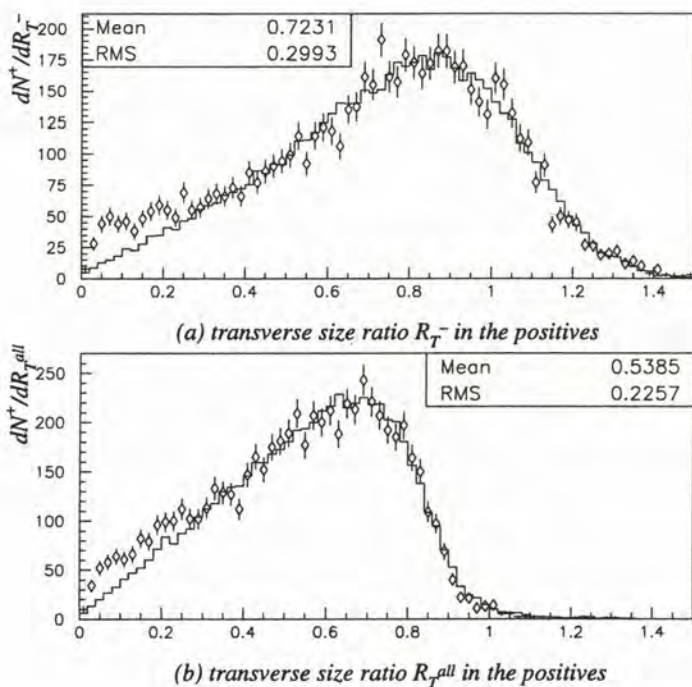


Figure 7.7: The transverse size ratio, minimised over (a) negative tracks and (b) all primary vertex tracks, in the Monte Carlo and data positives. In both cases agreement across the distribution is good, excepting a significant excess of events at low values $R_T < 0.3$.

7.3.2 Events surviving cuts in the data positives

Turning to the events surviving kinematic cuts and the $R_T^{all} > 0.2$ structure cut, we find that the excess of events noted in the previous section is still present. There is a substantial additional population at large values of the isolation q_T , leading to a dramatic excess of events surviving any given cut $q_T > q_T^{cut}$ in the data positives (figure 7.8(a) and (b) respectively).

Considering the 28 events passing the cuts defined above (*i.e.* with the isolation cut set to $q_T > 1.4 \text{ GeV}/c$) we find that the remaining kinematic variables are consistent with a charged-current background, so far as the limited statistics allow a comparison to be made (figure 7.9). The required mass distribution is peaked about $M_R \approx 1.6 \text{ GeV}/c$, as is typical for 3π candidates in charged-current events, rather than having a strictly rising distribution

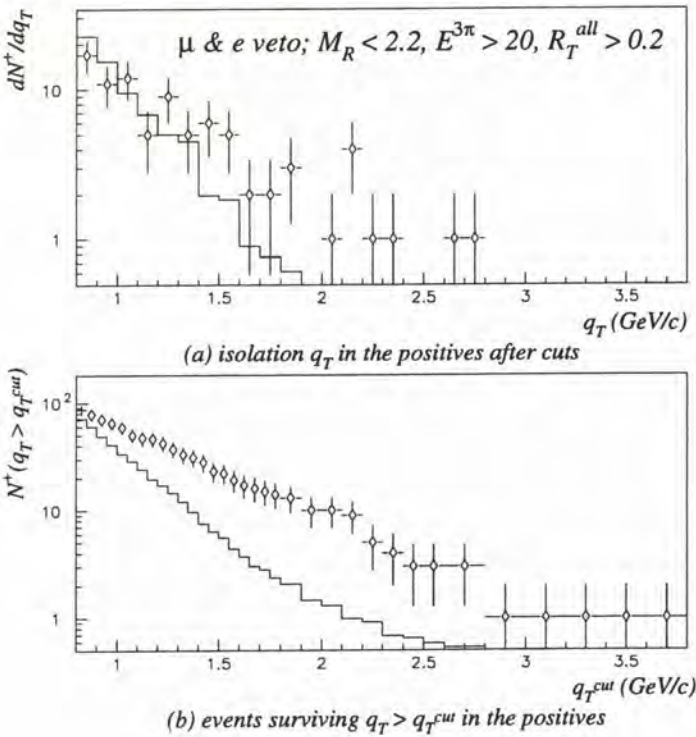


Figure 7.8: (a) The isolation variable q_T in the normalised Monte Carlo (histogram) and the data positives (points with Poissonian error bars), after vetos and all other cuts have been applied. (b) The number of events surviving all cuts as a function of the isolation cut q_T^{cut} .

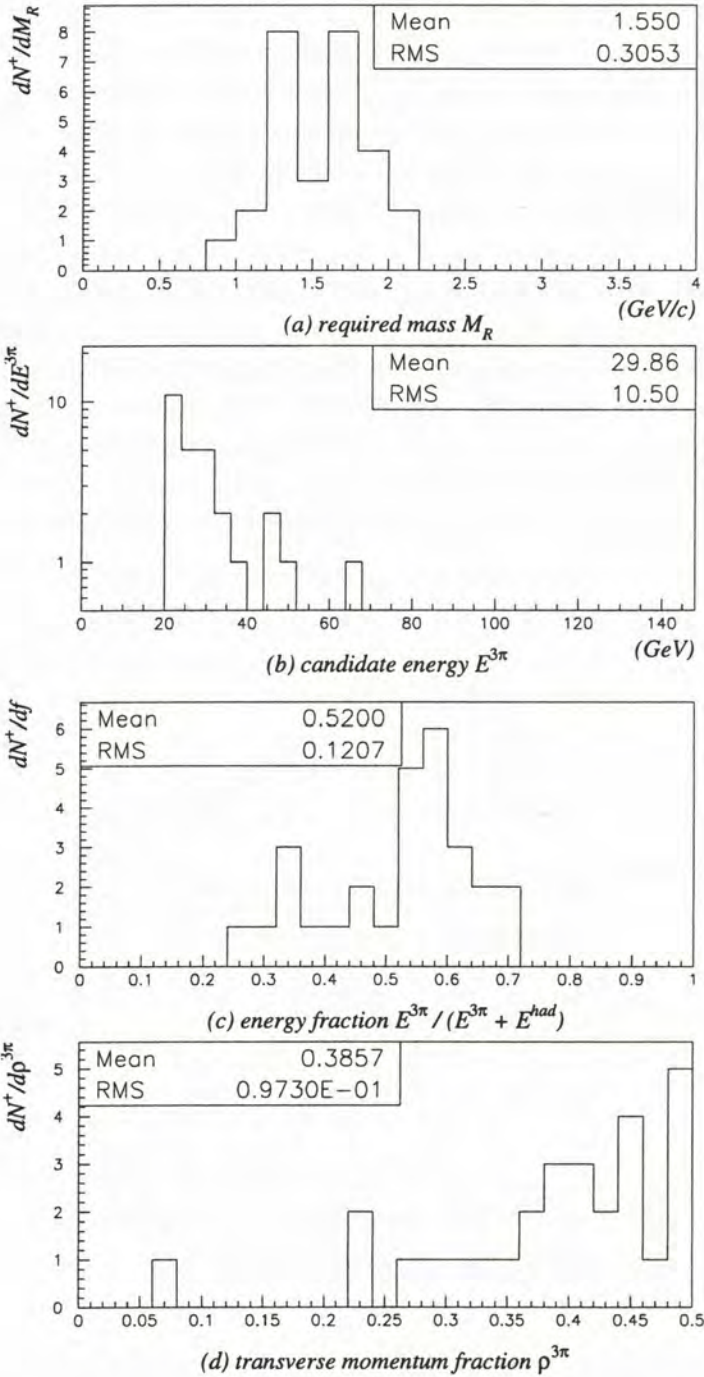


Figure 7.9: The remaining summary kinematic variables for the events surviving all cuts (including $q_T > 1.4$ GeV/c) in the data positives.

as would 3π from neutral-current events (7.9(a), *cf.* figure 6.21). Another variable offering some discrimination between the two classes of background is the transverse momentum fraction $\rho^{3\pi}$, with a rather flat distribution at large q_T values in neutral current events (see *e.g.* figure 4.12) which is retained under additional kinematic cuts, while the distribution for charged current events is strongly concentrated at high $\rho^{3\pi}$, with only a small tail to values $\rho^{3\pi} < 0.3$. The surviving events exhibit a $\rho^{3\pi}$ distribution of this second kind (figure 7.9(d)).




Figure 7.10: The 3π invariant mass $M_{inv}^{3\pi}$ for the events surviving all cuts (including $q_T > 1.4$ GeV/c) in the data positives.

For completeness we also show the distribution of invariant masses for the surviving events (figure 7.10) although this offers no discrimination between background hypotheses. By contrast the transverse size ratio variables, bearing in mind that a cut $R_T^{\text{all}} > 0.2$ has already been applied, have a distribution consistent with that expected for charged current events (figure 7.11, *cf.* figure 6.30(a) of the previous chapter) rather than the concentrations about $R_T^- \approx 0.9$ and $R_T^{\text{all}} \approx 0.7$ characteristic of neutral current events (figures 6.28(a) and 6.31(a)); these distributions are robust under kinematic cuts).

Some tests have therefore been made to check for gross failures of the muon- and electron-rejection procedures defined in the previous chapter. As noted in section 6.7.3, the data files on which this analysis is based have corrupt information concerning the probability of tracks to pass through the gap in the muon detection system, and any “hits” in the muon veto scintillators are furthermore unavailable. It is not possible at present to include this information in the analysis in a consistent manner, since this would require the re-production of the analysis data format for the large




Figure 7.11: The transverse size ratio, minimised over (a) negative tracks and (b) all primary vertex tracks, for the events surviving all cuts (including $q_T > 1.4 \text{ GeV}/c$) in the data positives.

Monte Carlo samples listed in section 7.2.1; however, carrying out this task on the 1995 and 1996 data we find that when the relevant information is restored, the number of surviving events falls from 28 to 26. The higher-than-expected rate of survivors is therefore *not* due to this oversight, at least in the first instance. Studies on the nature of the background are continuing.

It was partly in anticipation of possible problems of this kind that this analysis was temporarily restricted to the 1995-96 data. The higher-than-expected background suggests that the control of event quality and charged-current processes is not sufficiently mature to justify finalising the momentum reconstruction, quality and signal-defining cuts. However for completeness we will apply the cuts so far defined to the 1995-96 data, and assess the results as a test of the method developed up to this point. The 1997 and 1998 data will be kept in reserve, so that any influence which our knowledge of the data might have on subsequent refinements to the analysis is minimised.

7.3.3 Adjustment of cuts and construction of a correction factor

We noted in the previous chapter (section 6.7.6) that transverse size ratio cuts are most effective against charged current events at large q_T values. Since we appear to have a problem with a poorly-controlled background of this type, further R_T cuts suggest themselves as a way of reducing the number of survivors in the data. For example, if we apply a cut in the variable R_T^- in addition to the cuts already defined, we find that the number of surviving events falls significantly; the disagreement with the Monte Carlo prediction is also reduced (figure 7.12). Clearly we incur a loss of efficiency for three-prong decays by merely adding a cut in this way (17%), but short of an extended study of the problem there is little present alternative. In addition to the cut (9), then, we require that

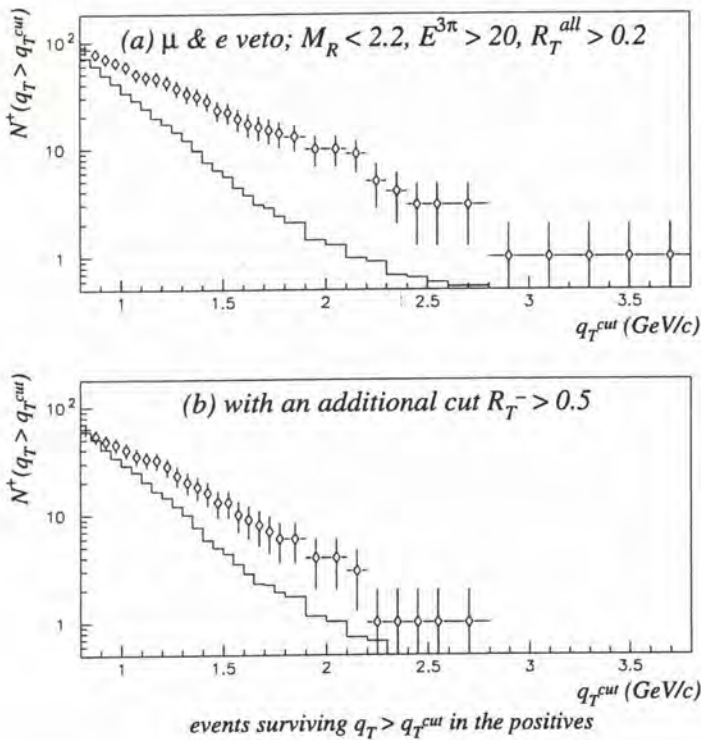


Figure 7.12: Events surviving all cuts in the Monte Carlo (histogram) and data positives (points with Poissonian error bars), (a) for the vetoes and cuts previously defined, and (b) when an additional cut $R_T^- > 0.5$ is imposed.

10. $R_T^- > 0.5$, where R_T^- is the minimum value of the transverse size ratio R_T^i over all *negatively-charged* tracks i at the primary vertex.

Even with this cut in place, we see from figure 7.12(b) that the observed background to the $\tau^+ \rightarrow \bar{\nu}_\tau \pi^+ \pi^- \pi^+ + n\pi^0 (n \geq 0)$ search exceeds that predicted by the Monte Carlo, and we would anticipate a similar phenomenon in the $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n \geq 0)$ search. Accordingly we define an *ad hoc* correction factor

$$C(q_T^{cut}) \stackrel{\text{def}}{=} \frac{N^+(q_T > q_T^{cut})_{data}}{N^+(q_T > q_T^{cut})_{MC}} \quad (7.5)$$

by which we will multiply the Monte Carlo background prediction for the τ^- search at any given q_T^{cut} . Plots of this quantity for the cuts defined above,

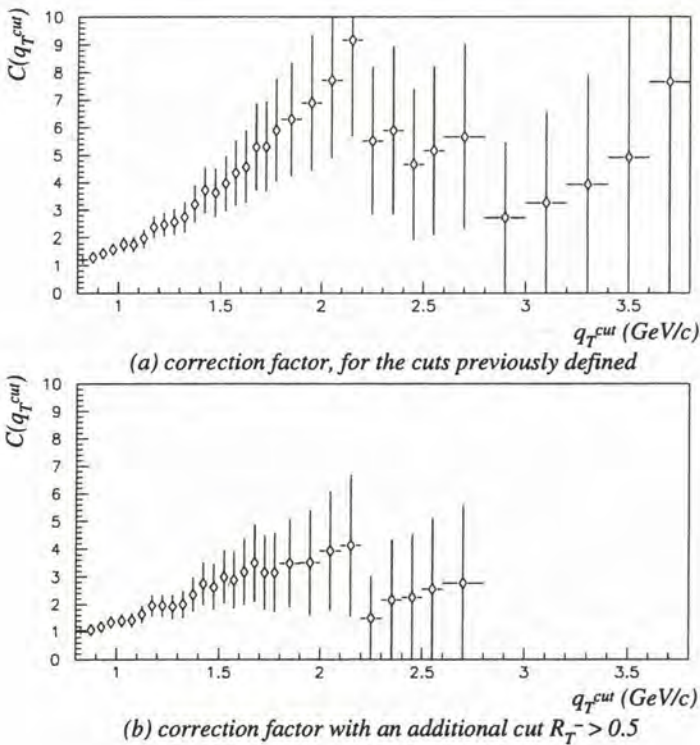


Figure 7.13: A multiplicative factor $C(q_T^{cut})$, based on the results of the positives search in the data, to correct the number of survivors predicted by the Monte Carlo, (a) for the vetos and cuts previously defined, and (b) when the additional $R_T^- > 0.5$ cut is imposed. The 1σ uncertainty due to the data and Monte Carlo statistics used is shown using error bars: clearly the uncertainty for neighbouring q_T^{cut} values is strongly correlated.

and after the additional cut (10) is imposed, are shown in figure 7.13: the improvement with the R_T^- cut is clear.

The limited data statistics, and the low number of surviving events at high q_T in the Monte Carlo, lead to a significant uncertainty on the value of the correction factor, also shown in the figure. This is unsatisfactory as uncertainty in the background prediction greatly complicates the calculation and interpretation of physical parameters from experimental results. In the following sections we will therefore restrict ourselves to values of the isolation cut where the correction factor C , and its uncertainty, are relatively modest.

7.3.4 “Data simulator” corrections for the neutral current background

For completeness we should also note an established technique for correcting the neutral current Monte Carlo for the apparent mis-simulation of the hadronic system, noted above (section 7.3.1). Beginning with identified ν_μ charged-current events in the data, the muon track and any associated calorimeter clusters, muon chamber hits *etc.* are removed from the data structure, and the measured muon momentum \vec{p}^μ is assigned to a Monte Carlo lepton l which is added to the event. In the case $l = \nu$ no addition is required, and this *data simulator* technique yields “fake” neutral current events, with hadronic system properties taken directly from those in the data, although there is a charge-bias effect due to the fact that a charged lepton rather than a neutrino was actually produced in the deep inelastic scattering. For the case $l = \tau$, granted a proper simulation of the τ decay and detector acceptance *etc.* for the decay products (the latter being a non-trivial task), fake tau decay events are obtained.

The efficiency with which neutral current backgrounds pass cuts in the established NOMAD analyses is then calculated as

$$\epsilon = \epsilon_{MC} \times \frac{\epsilon_{DS}}{\epsilon_{MCS}} \quad (7.6)$$

where ϵ_{MC} and ϵ_{DS} are the efficiencies for Monte Carlo and “fake” neutral current events respectively; ϵ_{MCS} is the efficiency for Monte Carlo ν_μ charged-current events where the muon has been removed in an analogous way (the “Monte Carlo simulator”). The ratio $\epsilon_{DS}/\epsilon_{MCS}$ is a suitable correction factor for the Monte Carlo efficiency, describing the differences between hadronic systems in the Monte Carlo and the data, on the assumption that any distortions due to the muon removal, or charge bias, cancel in the ratio. (Astier et al., 1999a, section 4.1)

A data simulator for the present analysis has not yet been implemented. Since the correction factor C defined in the previous section is applied to the combined background we are already taking Monte Carlo underestimates of the background into account, albeit in a much more crude manner than is done with a data simulator. We further note that

- the correction $\epsilon_{DS}/\epsilon_{MCS}$ for the neutral current background to the three-prong analysis reported by Astier et al. (1999a) was found to be 1.1, modest compared to the correction already employed for this analysis; and
- based on the distributions shown above we suppose, although we cannot prove, that this analysis is in any case dominated by a poorly-controlled *charged-current* background.

It is also conventional to apply a correction as at (7.6) to the efficiency with which simulated tau decay events pass analysis cuts. However, no such data simulator study has yet been carried out for a full three-prong decay analysis in the NOMAD experiment. This matter is discussed further in section 7.5.2 below.

7.4 Results of the analysis for the 1995-1996 data

As discussed in section 7.3.1, the presence of a background at high- q_T not predicted in the Monte Carlo suggests that the data-quality controls in this analysis—especially for charged current events—require further development. Nevertheless we may make a preliminary application of the analysis to the data using the correction factor defined by the positives search (section 7.3.3) to model the differences between the simulation and the data.

With this correction in hand, we choose a final q_T cut in section 7.4.1, in such a way as to maximise the sensitivity of the analysis to oscillations. We also define a blind region or sealed box with a more loose cut in q_T : in section 7.4.2 we compare distributions of events in the data with the Monte Carlo, *excluding* any events in the blind region where a potential signal might be apparent, as a last test on the data; the rate of surviving events in the Monte Carlo, multiplied by the correction factor defined above, is also compared to the data outside the blind region. Finally, in section 7.4.3 we open the box to discover the number of events in the data surviving all cuts.

7.4.1 Choice of final q_T cut and blind region

With all elements of the analysis now fixed, it remains to choose the final cut in the isolation variable q_T which defines the signal. Clearly we wish to do this in such a way as to maximise the sensitivity of the analysis to any three-prong tau decays which may be present, and alternatively to set the most stringent possible upper limit in the case that no signal is observed.

Once the value of q_T^{cut} is chosen, we proceed as follows. Given an observed number of events N_{obs} with $q_T > q_T^{cut}$, and passing all other cuts, and the expected number of background events with these same properties N_{exp} , we determine either an upper limit or a confidence belt for a neutrino oscillation signal using the unified approach of Feldman and Cousins (1998). In this way the choice between setting a limit and setting a confidence belt is made consistent with *coverage* at the desired confidence level (see the discussion in the work cited). In order to measure the intrinsic *sensitivity* of an analysis to the underlying physics—rather than the result actually obtained, which may be a strong function of background fluctuations, mis-estimates *etc.*—these authors recommend the use of the following quantity: the average upper limit on the number of signal events which would be found in an ensemble of experiments with expected background N_{exp} , but no true signal. This quantity, which we will label N_{sens} , provides a suitable way of choosing the value for the cut in q_T .

We use table XII of the work cited to find N_{sens} as a function of the expected background N_{exp} , choosing a 90% confidence level for the final limit or confidence belt. In order to choose the best sensitivity to the oscillation probability P , we minimise the ratio

$$\frac{N_{sens}}{\epsilon_\tau} \quad (7.7)$$

where ϵ_τ is the efficiency with which three-prong tau decays pass the cuts. Clearly ϵ_τ is a function of q_T^{cut} ; N_{sens} is also a function of q_T^{cut} due to its dependence on N_{exp} . The ratio (7.7), which will be more fully motivated in section 7.5.1, is plotted for a range of q_T^{cut} values in figure 7.14.

The resulting curve is relatively flat for a broad range of cuts $q_T > q_T^{cut}$, rising significantly only for q_T^{cut} values above 2.4 and below 1.2, and so some additional criteria are appropriate in choosing a value.² We choose to

²We cannot at present make a consistent plot of the sensitivity for q_T^{cut} values of 1.3 or below as the table of N_{sens} values in Feldman and Cousins (1998) stops at $N_{exp} = 15$. Accordingly the sensitivity (7.7) with N_{sens} calculated using a 90% confidence Poisson upper limit, given $N_{obs} = N_{exp}$, is shown dotted in the figure. This quantity provides an under-estimate for the sensitivity of Feldman and Cousins, as can be seen from the region




Figure 7.14: The sensitivity to the neutrino oscillation probability, N_{sens}/ϵ_{τ} , as a function of the q_T cut. The solid line shows the sensitivity as calculated using the N_{sens} values of Feldman and Cousins, while the dotted line shows the sensitivity where N_{sens} is set using the standard Poissonian upper limit (both at 90% confidence; see the text). A detail of the lower- q_T^{cut} region is shown in (b).

1. improve the efficiency with which any true three-prong tau decays would be detected; and
2. minimise the correction defined in section 7.3.3 which must be applied to the expected background.

Both criteria select a value at the low end of the minimum-region, as ϵ_{τ} is a monotonically falling function of q_T^{cut} , and the correction is a monotonically increasing function of q_T^{cut} in this region. The value used as a nominal cut up to this point, $q_T > 1.4 \text{ GeV}/c$, is therefore suitable as the final value for the cut.

where the curves overlap, and is included as a guide to the evolution of the sensitivity as q_T^{cut} is lowered.

We will use the lower value $q_T = 1.2 \text{ GeV}/c$ to define a blind region or sealed “box” which we will exclude from consistency plots, *etc.*, when the data is first examined.

7.4.2 Comparison of distributions for Monte Carlo and data events

As a first test of the analysis we compare the number of events surviving all cuts with the corrected background prediction, for values of the isolation cut outside the blind region: the results are shown in figure 7.15(a). A

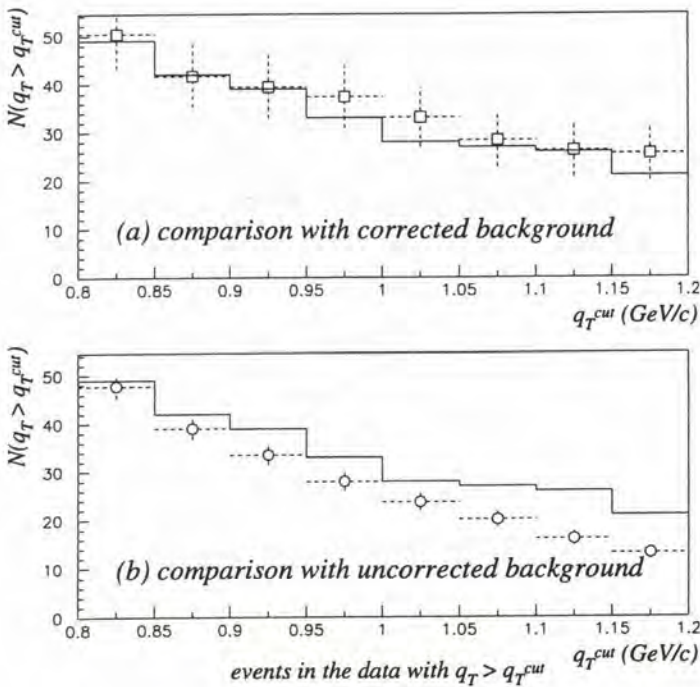


Figure 7.15: The number of events surviving all cuts in the data, as a function of the isolation cut q_T^{cut} , outside the blind region. Note that in this case the *data* is shown as the histogram: points with error bars show the expected background and its uncertainty, (a) after correction with the factor $C(q_T^{cut})$ described above, and (b) with no correction applied. The agreement with the data is greatly improved by the correction, although the uncertainty on the background also rises substantially. (Note also that the uncertainty for neighbouring q_T^{cut} values is strongly correlated.)

satisfactory agreement is seen, although we should note that a substantial uncertainty has been introduced into the background prediction by the limited number of events used to define the correction, as discussed above. (The statistical uncertainty is shown; we do not here attempt to assign a systematic uncertainty to this procedure.) By contrast, the uncorrected background prediction has a relatively low uncertainty, governed by the Monte Carlo statistics, but the agreement with the data is very poor, even for isolation cuts as low as $q_T > 0.9 \text{ GeV}/c$ where we would not expect any oscillation signal to be clearly manifested (figure 7.15(b); *cf.* figure 7.1 and the discussion of present limits in section 7.2.3 above).

Considering the distributions of analysis variables in the data, with the lepton vetos in place but without other cuts applied (figures 7.16 through 7.20), we see in general a similar behaviour to that of the positives sample (*cf.* figures 7.2 through 7.7). The distribution of q_T outside the signal region shows the well-known offset from the Monte Carlo prediction discussed above (figure 7.16). The remaining kinematic variables are in better agreement with the Monte Carlo prediction than was the case in the positives search: the excess at the peak in the required mass M_R is less significant, and the candidate energy $E^{3\pi}$ reproduces the Monte Carlo prediction in the tail, whereas a deficit was seen in the positives (figure 7.17, *cf.* figure 7.3).




Figure 7.16: The isolation variable q_T in the normalised Monte Carlo (histogram) and the data (points with Poissonian error bars). Statistics refer to the data. Events with identified muons and electrons have been rejected, and the vetos against muons and electrons for the leading- and 2nd-leading p_T tracks have also been applied; no kinematic cuts have been imposed. The signal or “blind” region $q_T > 1.2 \text{ GeV}/c$ has been excluded.

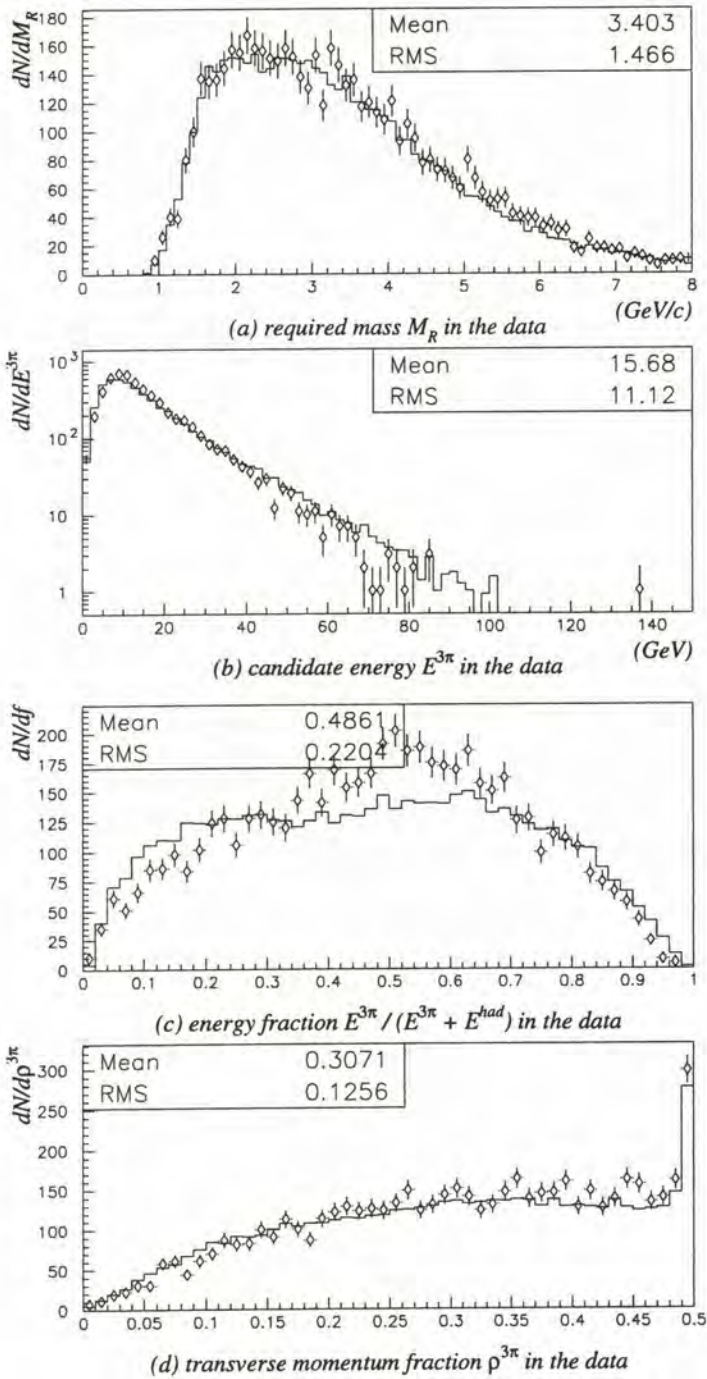


Figure 7.17: The four remaining summary kinematic variables in the Monte Carlo (histogram) and data (points with Poissonian error bars). The agreement across the distribution is good for all variables except the energy fraction (c).

By contrast in the 3π structure variables there is greater disagreement with the Monte Carlo: the difference in the shape of the invariant mass distribution is noticeable (figure 7.18) and there is a corresponding, and rather pronounced distortion in the opening angle distribution (figure 7.19(c)). For the transverse size ratio variables, the excess at $R_T < 0.3$, discussed above, is seen (figure 7.20).




Figure 7.18: The 3π invariant mass $M_{inv}^{3\pi}$ in the Monte Carlo and data. The excess of candidates at low masses is more pronounced than in the positives search (*cf.* figure 7.4).

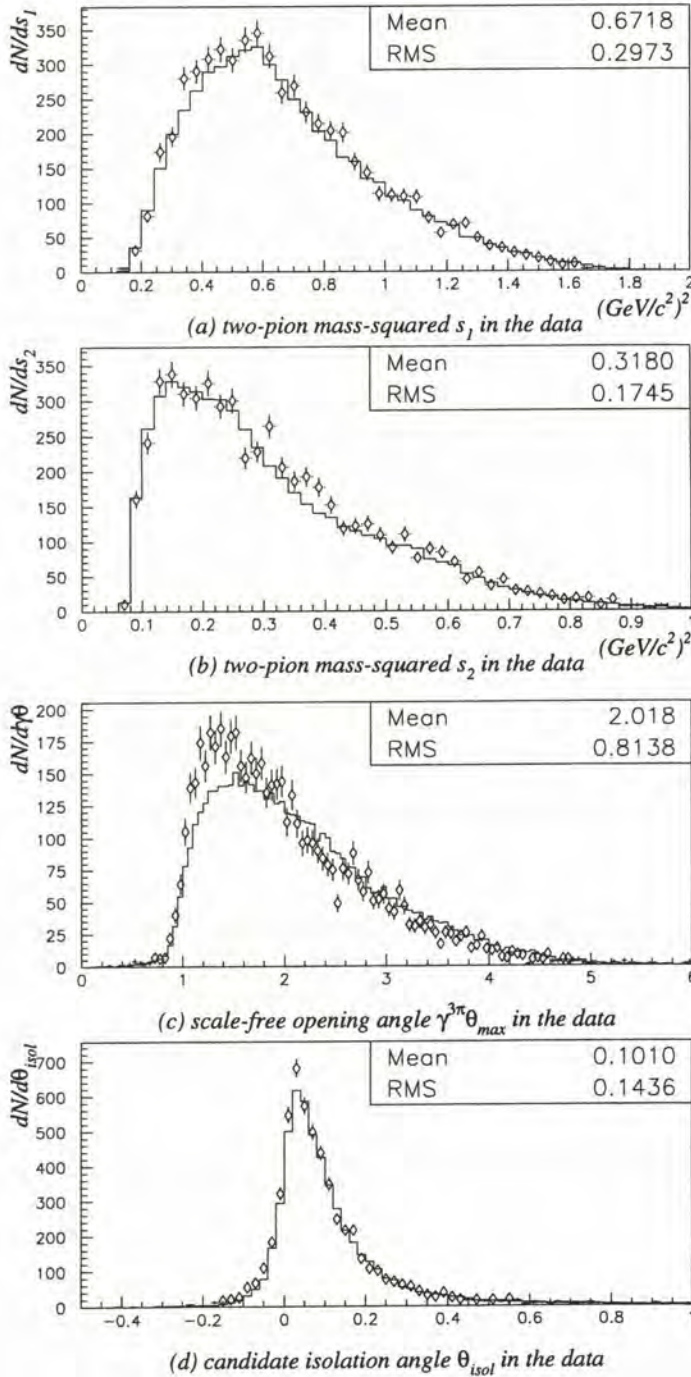


Figure 7.19: The remaining 3π structure variables in the Monte Carlo (histogram) and data (points with Poissonian error bars). The agreement is reasonable for the two-pion mass-squared variables (a) s_1 and (b) s_2 . The additional candidates at the peak in the scale-free opening angle (c) are much more pronounced than in the positives search (*cf.* figure 7.5), consistent with the behaviour of the invariant mass $M_{\text{inv}}^{3\pi}$.

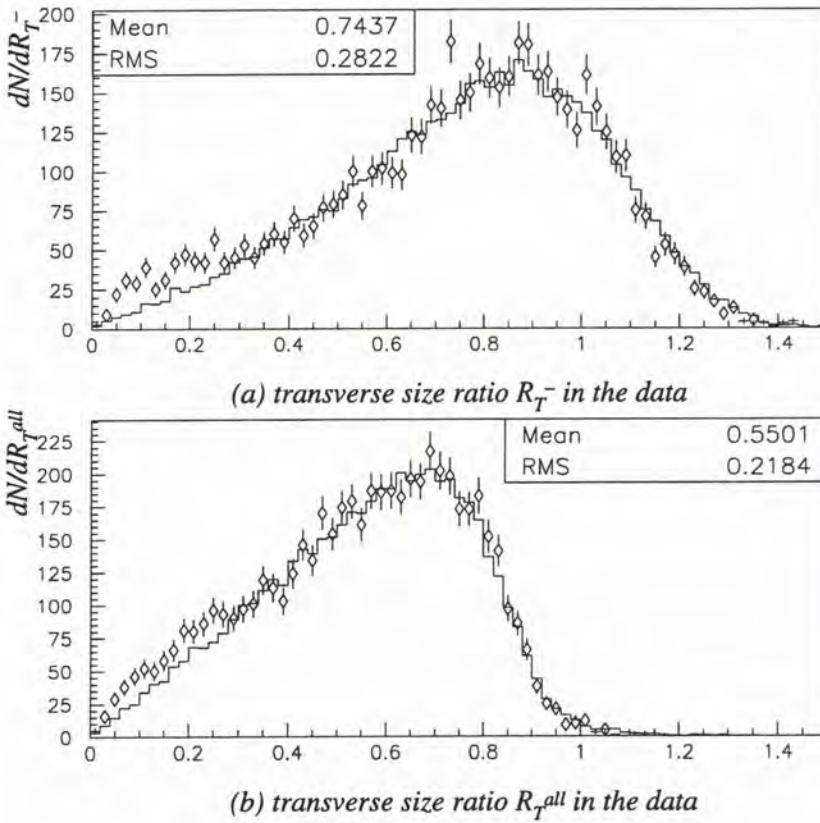


Figure 7.20: The transverse size ratio, minimised over (a) negative tracks and (b) all primary vertex tracks, in the Monte Carlo and data. Note the excess of events at low values $R_T < 0.3$.

7.4.3 Results in the signal region

Turning to the results within the blind region we find for the isolation cut previously defined, $q_T > 1.4 \text{ GeV}/c$, and all other cuts and vetos ((2) to (10)) in place, 10 surviving events in the data; the corrected background expectation is 12 events. Examining the number of surviving events across a large range of q_T^{cut} values (figure 7.21) we find that the corrected background estimate agrees well with the data within its uncertainties: the behaviour seen is thus rather consistent between the positives search and the true $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n \geq 0)$ search shown here.

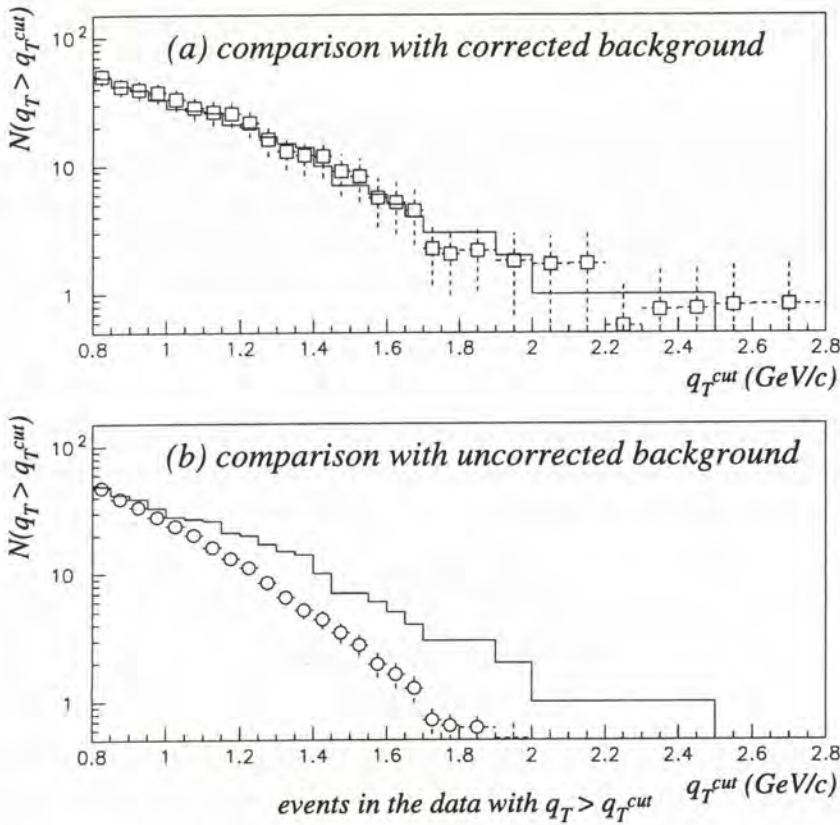


Figure 7.21: The number of events surviving all cuts in the data, as a function of the isolation cut q_T^{cut} . Good agreement is seen with the corrected background (a), within its substantial uncertainty, for all values of the isolation cut. By contrast the data is in significant excess of the uncorrected background (b).

7.5 Estimation of neutrino oscillation parameters

Intuitively it is clear that this result shows no evidence for neutrino oscillations within the present sensitivity of this analysis. In order to quantify this statement we now proceed to calculate the limit on the oscillation probability based on the observed number of events passing the cuts. The basic formulae and quantities are set out in section 7.5.1; limits on the $\nu_\mu \rightarrow \nu_\tau$ oscillation probability and mixing parameters, as well as the sensitivity of the analysis, are then determined in section 7.5.2.

7.5.1 Quantities determining the oscillation probability

In addition to the numbers of events seen in the data N_{obs} , and the number expected from the background N_{exp} , we require the mean number of three-prong decays which would survive all cuts if $\nu_\mu \rightarrow \nu_\tau$ oscillations occurred with unit probability. This quantity, which we label N_{max} , is estimated using the formula

$$N_{max} = N_\mu \times \frac{\sigma_\tau}{\sigma_\mu} \times B(3\text{-prong}) \times \epsilon_\tau \quad (7.8)$$

where the suppression of the ν_τ interaction cross-section (section 2.2.3), the combined three-prong branching fraction, and the efficiency of the analysis for three-prong decays are taken to be

$$\frac{\sigma_\tau}{\sigma_\mu} = 0.48, \quad (7.9)$$

$$B(3\text{-prong}) = 0.152 \quad \text{and} \quad (7.10)$$

$$\epsilon_\tau = 0.0196 \quad (7.11)$$

respectively. We should note that the NOMAD Monte Carlo simulates the three-prong decays using only the three modes listed at (2.9), but as these have a total branching fraction of 0.145 and the remaining modes (involving kaons) have a similar ratio of neutral-bearing to charged-only types, we assume that the simulated modes are representative of the full set, and therefore use the larger branching fraction at (7.10).

The total number of ν_μ charged-current events in the fiducial volume, N_μ , is itself an estimate, since we require the number prior to the effect of muon identification efficiencies *etc.*. Using the quantities defined in section 7.2.1 where the Monte Carlo was normalised to the data (see table 7.1) we use the efficiency of muon identification in the ν_μ CC Monte Carlo to correct the number of $\nu_\mu + N \rightarrow \mu^- + X$ events identified in the data, setting

$$N_\mu = N_{\mu-ID}^{data} \times \frac{N_{vert}^{MC}}{N_{\mu-ID}^{MC}} \quad (7.12)$$

$$= 391880 \quad (7.13)$$

Where a limit N_{lim} is set on the number of three-prong decay events seen, as is the relevant case here, the corresponding limit on the oscillation probability is then

$$P_{osc}(\nu_\mu \rightarrow \nu_\tau) < \frac{N_{lim}}{N_{max}}. \quad (7.14)$$

For alternative case of a confidence interval $[N_{osc}^1, N_{osc}^2]$ on the number of events seen, the interval for the oscillation probability is determined in an analogous manner.

We should also note that to determine the sensitivity of the analysis, in the sense defined by Feldman and Cousins, we simply substitute the quantity N_{sens} for N_{lim} in expression (7.14). Consulting (7.8) we therefore find that the sensitivity to the oscillation probability depends on the level of the q_T cut, and indeed on any of the analysis cuts, only through the quantities N_{sens} and ϵ_τ . The ratio

$$\frac{N_{sens}}{\epsilon_\tau},$$

which was used without explanation in section 7.4.1 to set the level of the isolation cut, is thus seen to be a suitable measure of the effect of the cuts on the sensitivity of the analysis.

7.5.2 Limit and exclusion region for $\nu_\mu \rightarrow \nu_\tau$ oscillation

Given the values listed above, we find from (7.8)

$$N_{max} = 560. \quad (7.15)$$

The values for the cross-section suppression, branching fraction, and number of ν_μ CC interactions used ((7.9), (7.10) and (7.13) above) are the best presently available, but we should note that the estimate for the efficiency with which three-prong decay events pass the cuts at (7.11) is rather poorly controlled, being determined from the Monte Carlo simulation alone. Bearing in mind the known mis-simulation of the hadronic system in the NOMAD Monte Carlo, we should make some correction to this efficiency. As discussed in section 7.3.4 above, “data simulator” studies for three-prong decay efficiency have not yet been carried out: noting that a reduction as large as 18% has been found for the tau efficiency in previous analyses of other decay channels, we use this figure as a correction for want of any other estimate. Replacing the efficiency at (7.11) by

$$\epsilon_\tau^* = 0.0161, \quad (7.16)$$

then, we find the corrected estimate of the signal at maximal mixing to be

$$N_{max}^* = 460 \quad (7.17)$$

Consulting table V of Feldman and Cousins (1998) we find that for 10 events observed, given an expected background of 12 events, the 90% confidence interval for the signal mean is $[0.00, 4.71]$ so that

$$N_{lim} = 4.71 \quad (7.18)$$

is the upper limit on the signal mean, yielding a limit

$$P_{osc}(\nu_\mu \rightarrow \nu_\tau) < 0.010 \quad (7.19)$$

on the oscillation probability.

As discussed in section 7.4.1 above, it is also appropriate to estimate the *sensitivity* of the analysis to oscillations, apart from any fluctuations which might affect the result actually observed. Consulting table XII of the work cited we find that an average upper limit

$$N_{sens} = 7.28 \quad (7.20)$$

is found for experiments with an expected background of 12 events, if no true signal is present; this corresponds to an oscillation probability

$$P_{osc}^{sens}(\nu_\mu \rightarrow \nu_\tau) = 0.016 \quad (7.21)$$

We translate these estimates into an exclusion region (or sensitivity) in the neutrino oscillation parameters as follows. Applying the simple two-neutrino (or “two-flavour”) oscillation model described in section 1.4, the oscillation probability is given by the function

$$P(\nu_\mu \rightarrow \nu_\tau; L) = \sin^2 2\theta \sin^2 \left(\frac{1.27 (\Delta m^2 / \text{eV}^2) (L / \text{km})}{E_\nu / \text{GeV}} \right), \quad (7.22)$$

where the parameters $(\sin^2 2\theta, \Delta m^2)$ describe the neutrino mixing, and the physical baseline L and neutrino energy E_ν are characteristic of the experiment. For sufficiently large Δm^2 (*i.e.* $\Delta m^2 L / E_\nu \gg 1$) the effect of the second term is “washed out” even by small variations of L and E_ν about their mean values, so that we may replace it by its mean value; the oscillation probability is then given by

$$P(\nu_\mu \rightarrow \nu_\tau) = \frac{1}{2} \sin^2 2\theta \quad (7.23)$$

as previously discussed. This justifies our use of a single oscillation probability above to characterise all the experimental results, regardless of the features of individual events.

In this case the limit on the mixing parameter is just twice the limit on the oscillation probability, so that

$$\sin^2 2\theta < 0.020 \quad (7.24)$$

in the large- Δm^2 region, with a corresponding sensitivity

$$(\sin^2 2\theta)_{sens} = 0.032 \quad (7.25)$$

Where the condition $\Delta m^2 L/E_\nu \gg 1$ fails, the determination of mixing parameters is rather more complicated. For a given set of surviving events, for example, the total energies of individual events contain information on the Δm^2 value (in the hypothesis of neutrino oscillations) through the second term in (7.22). Since we make no such test in this analysis, the properties of the surviving events do not contribute to the calculation: however the properties of the *simulated signal* do contribute, since the oscillation probability (7.22) is a function of E_ν for each individual event.




Figure 7.22: The range of mixing parameters ($\sin^2 2\theta, \Delta m^2$) excluded by this analysis, at 90% confidence, based on the 1995-96 NOMAD data. The published NOMAD limits at large Δm^2 and large mixing, based on an analysis over several decay channels for the 1995-96 and 1997 data, are also shown (Astier et al., 1999a)

We therefore determine the limit in $\sin^2 2\theta$ as a function of Δm^2 by assigning a weight

$$w_i = \frac{\int_{L_{min}}^{L_{max}} dL \sin^2(1.27\Delta m^2 L/E_\nu)}{0.5(L_{min} - L_{max})} \quad (7.26)$$

to each event i ($i = 1 \dots N_\tau^{MC}$) surviving cuts in the three-prong tau decay Monte Carlo, for a range of values of Δm^2 (Feldman and Geiser, 1999). In this expression the oscillation lengths $L_{min} = 0.415$ km and $L_{max} = 0.830$ km correspond to neutrinos produced in decays at the end of the vacuum decay tunnel, and the region immediately following the T9 target respectively (see section 2.2.1); E_ν is the true energy of the incident ν_τ taken from the simulation. The value

$$\tilde{N}_{max}(\Delta m^2) = N_{max} \frac{1}{N_\tau^{MC}} \sum_{i=1}^{N_\tau^{MC}} w_i \quad (7.27)$$

is then substituted for N_{max} in equation (7.14) and the limit on $\sin^2 2\theta$ recalculated.

The resulting exclusion region is shown in figure 7.22.

7.6 A discussion: The present state and prospects of this analysis

It is clear, from the results of the “positives search” in section 7.3, and our resort to a correction factor for the expected background, that this analysis is limited by its poor control of background processes. We should also note that the excess at large values of the isolation q_T has not yet been subjected to a sustained study: it is therefore *not* clear whether the background currently seen is characteristic of the analysis, or whether some straightforward measure would bring the background into agreement with the Monte Carlo prediction.

For this reason we treat the limit on $\nu_\mu \rightarrow \nu_\tau$ oscillations determined in the previous section as an illustration of the present state of the analysis, rather than a serious estimate of the underlying physics. Supposing in the worst case that the background seen is characteristic of the analysis in its present form, the *expected level* of the background would need to be accurately determined before a limit could be calculated using the simple technique of section 7.5; alternatively, the effect of the substantial uncertainty on the background would have to be taken into account by the implementation of a toy Monte Carlo, or some other means. Supposing on the other hand that the background is reducible, a more reliable (and sensitive) limit would be drawn from a modified analysis.

With these points in mind we may define the following agenda for the development of the analysis presented in this work.

1. Most straightforwardly, the two known “bugs” in the analysis should be removed:
 - (a) The corrupt information describing the muon scintillators and the “gap” in the muon system has already been corrected for the data, and needs to be extended to the Monte Carlo simulation. Its effect appears to be small based on studies to date (sections 6.7.3 and 7.3.2).
 - (b) We noted in section 6.3.2 that the handling of V^0 vertices is a holdover from dedicated V^0 studies: in particular, reconstructed V^0 very close to the primary vertex are treated as such, rather than being merged to the primary as is done for other classes of vertex. This creates potential problems for the momentum reconstruction, for example the formation of fake V^0 by primary tracks, followed by the discarding of the tracks due to the V^0 failing pointing cuts. As with the previous case, this is straightforward to test and/or

rectify but requires the reprocessing of the analysis data format over several million events.

2. The events passing cuts in the positives should be studied, in order to establish if possible the nature of the surviving background. The quality cuts defined in the previous chapter (sections 6.2 through 6.4) should also be re-examined in case misreconstruction effects are contributing to the background excess.
3. The charged-current veto for this analysis is relatively unsophisticated, and could be extended using a number of techniques already developed elsewhere. For example,
 - the NOMAD reference analysis for this decay mode uses a transverse momentum fraction for the candidate lepton ρ_t in conjunction with the transverse size ratio variable (Astier et al., 1999a, section 8);
 - electron identification and momentum reconstruction in the $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ and $\tau^- \rightarrow \nu_\tau \pi^-$ analyses is augmented by reconstruction of the “bremsstrahlung strip” in the electromagnetic calorimeter for candidate electrons (*op. cit.*, sections 5 and 6).

Various other quantities, such as variables measuring the calorimeter response to tracks (section 6.6.5) are yet to be exploited in this analysis.

4. While retaining the description of neutrino interactions and 3π candidates developed in chapters 3 and 4, the analysis could be extended by the use of likelihood functions to properly exploit
 - (a) the differing correlations between q_T , $E^{3\pi}$ and the R_T cuts for three-prong decays and the backgrounds; and
 - (b) the differing distributions of the 3π structure variables $M_{inv}^{3\pi}$, s_1 and s_2 for three-prong decays and the neutral current background, as discussed in section 4.7.

An improvement in sensitivity is to be expected from each of these measures. Some loss of τ efficiency may be incurred by tightening the charged-current veto, so that it is not possible to have assurance about an improved final result: the effect may merely be to exchange the correction factor of section 7.3.3 for a higher background estimate. Extending the analysis presented here to include a likelihood function, however, may be expected to yield a significant reduction in the ratio of background to signal efficiency,

and hence an improvement in the final sensitivity of the analysis. The differing correlations between q_T , $E^{3\pi}$ and R_T are clear, as are the differing distributions of the 3π structure variables. Having identified these distinctions between signal and background, we are not yet exploiting them in an optimal way.

A comparison between the present work and the NOMAD reference analysis (Astier et al., 1999a, section 8) is to be made in the light of this last point. The reference analysis is undeniably more mature, as may be determined from the additional features implemented, and its more accurate prediction of the background seen in the positives. It is also clearly more sensitive, with a background expectation half that of the present analysis, and a tau efficiency nearly twice that found here. We have however shown that the key kinematic features of three-prong decay candidate events may be described by a remarkably economic set of variables (chapter 3); and that much of the confusing array of “ 3π structure” variables may be set aside when kinematic and 3π mass effects are properly taken into account (chapter 4). The present work thus contributes to the transparency of the analysis: it is to be hoped that once the likelihood method, and certain other features of the reference analysis are incorporated into this $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n \geq 0)$ search, a significant improvement in the sensitivity to neutrino oscillations will be achieved.

Chapter 8

Conclusion

The NOMAD experiment, while optimised for the search for $\tau^- \rightarrow \nu_\tau \bar{\nu}_e e^-$ decays, is also sensitive to hadronic decays of the the tau, in particular the three-prong decays $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^- + n\pi^0 (n \geq 0)$. Both these and other analyses rely in part on efficient particle identification to suppress various backgrounds.

In the course of the particle identification studies reported here, we have shown the suitability and effectiveness of the standard NOMAD TRD variables for the task of electron rejection, tuned to a level required by a given analysis. The technique for discriminating between protons, kaons, pions and muons, based on ionisation energy deposited in the TRD straw tubes, has also been shown to provide useful results in certain kinematic ranges. In particular the application to a search for $\phi \rightarrow K^+ K^-$ decays, at present limited by the NOMAD data processing, offers the potential for a first measurement of ϕ production in neutrino interactions.

The search for three-prong tau decays, developed in this work, relies in the first instance on the separation of candidate three-prong decay events from neutral current interactions, using summary kinematic variables. A simple argument shows that five such variables should be sufficient to describe an event, in the approximation where particle and multiparticle-system masses are ignored.

Both our understanding of the fragmentation process, and detailed Monte Carlo simulation, suggest that the isolation variable q_T is the single most important kinematic quantity for separating tau decay from neutral current events. Once this quantity is in hand, an analysis of the remaining degrees of freedom shows that the use of the traditional “transverse plane” variables is no longer strictly appropriate. An analysis of events based on the “plane of momenta” leads instead to a consistent set of five summary variables, of which the most important are the isolation q_T , the decay candidate energy

$E^{3\pi}$, and the transverse mass M_{\perp} .

Consideration of the $\tau^{-} \rightarrow \nu_{\tau} \pi^{-} \pi^{+} \pi^{-}$ decay structure then suggests the modification of M_{\perp} to include the kinematic effects of the 3π mass, leading to the required mass variable M_R which provides superior rejection of the neutral current background. Among the many remaining structure-sensitive variables, only the three-pion invariant mass $M_{inv}^{3\pi}$, and the invariant masses of $\pi^{+} \pi^{-}$ combinations among the three pions (conventionally written $\sqrt{s_1}$ and $\sqrt{s_2}$) are found to provide significant discrimination under the conditions of the NOMAD experiment.

Once the separation of tau decay and neutral current events is established, strong additional measures must be taken to suppress the background due to charged-current processes. In addition to the rejection of events with identified muons and primary-vertex electrons, it is found necessary to veto events with high- p_T tracks which are identified as muons or electrons by rather weak criteria, or which are hard to identify in this way due to acceptance effects *etc.*. Cuts in the “transverse size ratio”, an event-structure quantity, are also found to be important in suppressing the charged-current backgrounds.

With all of these cuts in place, the application of the analysis to the 1995-96 data shows that a background is present in excess of the Monte Carlo prediction; a study of kinematic and structure variables suggests that this background is due to a charged-current process. After correcting for this background we find no evidence for $\nu_{\mu} \rightarrow \nu_{\tau}$ oscillations in the 1995-96 data, setting a limit

$$\sin^2 2\theta < 0.020$$

at 90% confidence in the large- Δm^2 region. The corresponding sensitivity of the analysis to the mixing parameter is found to be

$$(\sin^2 2\theta)_{sens} = 0.032.$$

The analysis presented in this work has potential for further development. In addition to a number of measures to improve the charged-current veto, which is at present its limiting feature, the analysis should be extended by the use of likelihood functions. Both the treatment of tau structure variables, and the correlation between the two kinematic variables q_T and $E^{3\pi}$ and the event structure variable R_T , are well-suited to this technique, which would lead to a more powerful discrimination between tau decays and the background processes. Once the analysis is extended in this way, a substantial improvement in the sensitivity to neutrino oscillations may be foreseen.

Appendices:

Appendix A

Formulae for the relation between planes of the event

In section 3.4.6 above the connection between the plane-of-momenta representation of an event (figure 3.9 and section 3.4.2) and the variables of the transverse plane (section 3.4.1) were discussed using Euler angles. For completeness the relevant formulae are presented here as an appendix.

A.1 Representation of the rotation using Euler angles

An arbitrary rotation $(x', y', z') \mapsto R(x', y', z')$, with the origin fixed, may be decomposed into three successive rotations:

1. a rotation by ψ about the z' -axis,
2. a rotation by θ about the *rotated* y' -axis, and
3. a rotation by ϕ about the *rotated* z' -axis,

where the order of the rotations is of course significant. The matrix of the rotation is then

$$R = \begin{bmatrix} \cos \phi \cos \theta \cos \psi - \sin \phi \sin \psi & -\sin \phi \cos \theta \cos \psi - \cos \phi \sin \psi & \sin \theta \cos \psi \\ \cos \phi \cos \theta \sin \psi + \sin \phi \cos \psi & -\sin \phi \cos \theta \sin \psi + \cos \phi \cos \psi & \sin \theta \sin \psi \\ -\cos \phi \sin \theta & \sin \phi \sin \theta & \cos \theta \end{bmatrix} \quad (\text{A.1})$$

Granted $0 \leq \psi, \theta, \phi < 2\pi$, and $\theta \neq 0, \pi$, the decomposition into the *Euler angles* ψ, θ, ϕ is unique (Zwillinger, 1996).




Figure A.1: The plane-of-momenta representation of the event, in a standard position in the $x' - z'$ plane, with coordinates shown.

In section 3.4.6 an event, in the “massless vectors” approximation, was described by placing its plane-of-momenta representation in a standard position and then performing two rotations. The standard position chosen was for the momentum vectors $\vec{p}^{3\pi}$, \vec{p}^{had} , \vec{p}^{tot} to begin at the origin of coordinates and lie in the $x' - z'$ plane, with the total momentum vector \vec{p}^{tot} aligned with the z' -axis. This arrangement, together with the (x', y', z') coordinates for the vectors which it entails, is shown in figure A.1. Since these vectors will be rotated it is convenient to have a subscript other than z or z' to describe the components of $\vec{p}^{3\pi}$ and \vec{p}^{had} parallel to the total momentum, equal to $\sqrt{(p^{3\pi})^2 - q_T^2}$ and $\sqrt{(p^{had})^2 - q_T^2}$ respectively: we choose ζ , which we will use to label the image of the z' -axis under rotation, where necessary.

As is easy to verify from (A.1), the azimuth of the ζ - or rotated- z' -axis is given by the angle ψ . Since we are assuming throughout the analysis that the overall azimuth of the event is irrelevant, we set $\psi = 0$ whereupon the matrix of the rotation takes the much-simplified form

$$R = \begin{bmatrix} \cos \phi \cos \theta & -\sin \phi \cos \theta & \sin \theta \\ \sin \phi & \cos \phi & 0 \\ -\cos \phi \sin \theta & \sin \phi \sin \theta & \cos \theta \end{bmatrix} \quad (\text{A.2})$$

The final or “rotated” momentum vectors are then given by

$$\vec{p}^{3\pi} = R \begin{bmatrix} +q_T \\ 0 \\ p_\zeta^{3\pi} \end{bmatrix} = \begin{bmatrix} q_T \cos \phi \cos \theta + p_\zeta^{3\pi} \sin \theta \\ q_T \sin \phi \\ -q_T \cos \phi \sin \theta + p_\zeta^{3\pi} \cos \theta \end{bmatrix} \quad (\text{A.3})$$

$$\vec{p}^{had} = R \begin{bmatrix} -q_T \\ 0 \\ p_\zeta^{had} \end{bmatrix} = \begin{bmatrix} -q_T \cos \phi \cos \theta + p_\zeta^{had} \sin \theta \\ -q_T \sin \phi \\ q_T \cos \phi \sin \theta + p_\zeta^{had} \cos \theta \end{bmatrix} \quad (\text{A.4})$$

while the “missing” transverse momentum is given by

$$\vec{p}_T^{miss} = \begin{bmatrix} -(p_\zeta^{3\pi} + p_\zeta^{had}) \sin \theta \\ 0 \end{bmatrix} \quad (\text{A.5})$$

These expressions were used to generate the event sketches in figure 3.21.

A.2 Some notes on expressions for $p^{3\pi}$, p_T^{had} and p_T^{miss}

The expressions for the transverse momenta used in the discussion of variables in section 3.4.6 were taken from the formulae (A.3) through (A.5), making the following observations and approximations:

- Since $p_T^{miss} = (p_\zeta^{3\pi} + p_\zeta^{had}) \sin \theta$, if we set

$$k = \frac{p_\zeta^{3\pi}}{p_\zeta^{3\pi} + p_\zeta^{had}}$$

we can write $p_\zeta^{3\pi} \sin \theta = k p_T^{miss}$ and $p_\zeta^{had} \sin \theta = (1 - k) p_T^{miss}$, making clear the dependence of $p_T^{3\pi}$ and p_T^{had} on p_T^{miss} .

- The comment that a correction of $\pm \frac{1}{3} p_T^{miss}$ can easily arise in the $M \approx 2q_T + 2k p_T^{miss} \approx 2q_T + p_T^{miss}$ expression at (3.13) follows from an assignment $p_\zeta^{3\pi} = 2p_\zeta^{had}$, whence $k = \frac{2}{3}$.
- From (A.4), p_T^{had} is given by

$$p_T^{had} = |-q_T \cos \theta + p_\zeta^{had} \sin \theta|$$

when $\phi = 0$, with the change between q_T - and p_ζ^{had} -term dominance occurring at $p_\zeta^{had} = q_T \cos \theta / \sin \theta$. The approximate condition for q_T to dominate quoted in the text, $p_T^{miss} \lesssim 2q_T$, holds for $\cos \theta \approx 1$ and $p_\zeta^{3\pi} = p_\zeta^{had}$ so that $p_\zeta^{had} \sin \theta = \frac{1}{2} p_T^{miss}$. If instead the longitudinal (*i.e.* ζ) component of either the 3π - or hadronic systems dominates the event, then the condition will no longer be true even as an approximation.

- The standard approximation

$$\sqrt{q_T^2 + \left(\frac{p_T^{miss}}{2}\right)^2} \approx q_T \left(1 + \frac{(p_T^{miss})^2}{8q_T^2}\right),$$

taking the first two terms in the Taylor expansion of $\sqrt{1+x}$, was used to clarify the expression for M at (3.15).

- M depends on *both* of the minor degrees of freedom of the plane-of-momenta, $p^{3\pi}$ and $\frac{E^{3\pi}}{E^{3\pi}+E^{had}}$:
 - given p_T^{miss} and ϕ , say, the *sum* $p^{tot} = p_\zeta^{3\pi} + p_\zeta^{had}$ is needed to determine the $\cos\theta$ terms ($p_T^{miss} = p^{tot} \sin\theta$);
 - the factor k depends by construction on the *relative size* of the p_ζ terms.

Appendix B

Some results concerning transverse mass variables

A number of results concerning the transverse mass variables M_\perp , M_T and M_R , in particular the latter, were quoted in chapter 4: the corresponding derivations and proofs are assembled in this appendix. The derivation of M_R is given in B.1. In B.2, M_R is shown to be a type of transverse mass variable, obeying the inequality $M_\perp < M_T \leq M_R$. In B.3, the interpretation of M_R as the minimum mass “required” for the τ if we identify $\vec{p}_T^{\nu\tau} = \vec{p}_T^{miss}$, quoted in section 4.3 above, is formally proven. In B.4 it is shown that given ideal measurement of momenta by the detector, M_R is bounded above by the tau mass. Finally, the form of the surface in the $(\rho^{3\pi}, \rho^{had}, M)$ space showing the effect of a cut in M_R on the variables of the transverse plane (as discussed in section 4.3.3 of the text) is derived in B.5.

B.1 Derivation of M_R

Suppose we interpret an event as due to the interaction $\nu_\tau + N \rightarrow \tau^- + X$ followed by the decay $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$, and further that we identify the missing transverse momentum \vec{p}_T^{miss} with the transverse momentum of the final-state tau neutrino $\vec{p}_T^{\nu\tau}$. We may then use the equation for the tau mass to solve for the longitudinal momentum of the final-state neutrino $p_L^{\nu\tau}$,

$$\begin{aligned} m_\tau^2 &= (E^{\nu\tau} + E^{3\pi})^2 - (\vec{p}^{\nu\tau} + \vec{p}^{3\pi})^2 \\ &= 2E^{3\pi} \sqrt{(p_T^{\nu\tau})^2 + (p_L^{\nu\tau})^2} + (M_{inv}^{3\pi})^2 - 2\vec{p}_T^{\nu\tau} \cdot \vec{p}_T^{3\pi} - 2p_L^{\nu\tau} p_L^{3\pi}, \end{aligned} \quad (\text{B.1})$$

where various terms have been cancelled, and the longitudinal and transverse momentum components separated, in the second line. Rearranging,

$$2E^{3\pi} \sqrt{(p_T^{\nu\tau})^2 + (p_L^{\nu\tau})^2} = m_\tau^2 - (M_{inv}^{3\pi})^2 + 2\vec{p}_T^{\nu\tau} \cdot \vec{p}_T^{3\pi} + 2p_L^{\nu\tau} p_L^{3\pi}; \quad (\text{B.2})$$

and squaring,

$$4(E^{3\pi})^2 [(p_T^{\nu\tau})^2 + (p_L^{\nu\tau})^2] = (D + 2p_L^{\nu\tau} p_L^{3\pi})^2, \quad (\text{B.3})$$

where we write

$$D = m_\tau^2 - (M_{inv}^{3\pi})^2 + 2\vec{p}_T^{\nu\tau} \cdot \vec{p}_T^{3\pi}. \quad (\text{B.4})$$

Expanding the square and then gathering terms, we arrive at a quadratic equation in $p_L^{\nu\tau}$:

$$[(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2] (p_L^{\nu\tau})^2 - [p_L^{3\pi} D] (p_L^{\nu\tau}) + \left[(p_T^{\nu\tau})^2 (E^{3\pi})^2 - \frac{D^2}{4} \right] = 0. \quad (\text{B.5})$$

Now frequently no solution can be found to this equation. In this case we replace the constant m_τ by a variable M , so that

$$D = M^2 - (M_{inv}^{3\pi})^2 + 2\vec{p}_T^{\nu\tau} \cdot \vec{p}_T^{3\pi}, \quad (\text{B.6})$$

and seek the smallest value of M which will allow a solution, *i.e.* that value for which the determinant is zero:

$$\begin{aligned} \Delta &= (p_L^{3\pi} D)^2 - 4 [(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2] \left[(p_T^{\nu\tau})^2 (E^{3\pi})^2 - \frac{D^2}{4} \right] \\ &= 0. \end{aligned} \quad (\text{B.7})$$

Collecting terms we have

$$D^2 = 4(p_T^{\nu\tau})^2 [(p_T^{3\pi})^2 + (M_{inv}^{3\pi})^2],$$

and solving with (B.6) for M , finally,

$$M^2 = (M_{inv}^{3\pi})^2 + 2p_T^{\nu\tau} \sqrt{(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2} - 2\vec{p}_T^{\nu\tau} \cdot \vec{p}_T^{3\pi}.$$

Instead of using the solution for $p_L^{\nu\tau}$ which this allows, we consider the properties of the solution for the tau mass M itself. We call

$$M_R = \sqrt{(M_{inv}^{3\pi})^2 + 2p_T^{miss} \sqrt{(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2} - 2\vec{p}_T^{miss} \cdot \vec{p}_T^{3\pi}} \quad (\text{B.8})$$

the *required mass*, as $M = M_R$ is the smallest nominal tau mass for which the 3π -missing momentum system can be interpreted as the product of the decay $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ (identifying $\vec{p}_T^{miss} = \vec{p}_T^{\nu\tau}$; see also section B.3 below). Intuitively, a cut on this variable will be more demanding than a cut on the transverse mass, since the derivation of M_R incorporates information on longitudinal momenta within the 3π system which is ignored by M_\perp and M_T .

B.2 Comparison of M_R , M_T and M_\perp

The required mass was originally developed as an analysis variable following the argument of the previous section. Comparing the expression (B.8) for M_R with the generic transverse mass expression (4.6), *viz.*

$$\mathcal{M}^2 = (E_T^{miss} + E_T^{3\pi})^2 - (\vec{p}_T^{miss} + \vec{p}_T^{3\pi})^2,$$

we see that M_R is in fact a transverse mass variable with

$$E_T^{3\pi} = \sqrt{(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2},$$

as discussed in section 4.3. This raises the question of the relationship between the value of M_R and that of the other transverse mass variables discussed in that section,

$$\begin{aligned} M_\perp &= \sqrt{(p_T^{miss} + p_T^{3\pi})^2 - (\vec{p}_T^{miss} + \vec{p}_T^{3\pi})^2} \\ &= \sqrt{2p_T^{miss} p_T^{3\pi} - 2\vec{p}_T^{miss} \cdot \vec{p}_T^{3\pi}} \end{aligned} \quad (\text{B.9})$$

and

$$\begin{aligned} M_T &= \sqrt{(p_T^{miss} + E_T^{3\pi})^2 - (\vec{p}_T^{miss} + \vec{p}_T^{3\pi})^2} \\ &= \sqrt{[(E_T^{3\pi})^2 - (p_T^{3\pi})^2] + 2p_T^{miss} E_T^{3\pi} - 2\vec{p}_T^{miss} \cdot \vec{p}_T^{3\pi}}, \end{aligned} \quad (\text{B.10})$$

where

$$E_T^{3\pi} = \sum_{\pi \in 3\pi} \sqrt{m_\pi^2 + (p_T^\pi)^2}. \quad (\text{B.11})$$

Immediately we have the result that $M_T > M_\perp$, since (B.9) and (B.10) are identical apart from the first term in (B.10), and a substitution $p_T^{3\pi} \mapsto E_T^{3\pi}$

in the second term, while

$$\begin{aligned} \sum_{\pi \in 3\pi} \sqrt{m_\pi^2 + (p_T^\pi)^2} &> \sum_{\pi \in 3\pi} \sqrt{(p_T^\pi)^2} \\ &= \sum_{\pi \in 3\pi} |\vec{p}_T^\pi| \\ &\geq \left| \sum_{\pi \in 3\pi} \vec{p}_T^\pi \right| \\ &= p_T^{3\pi}. \end{aligned}$$

It is also true that $M_R \geq M_T$ for all events, a proof of which follows:

Comparing equations (B.8) and (B.10), we see that the third terms are identical, while the first and second terms in the expression for M_R exceed or equal the first and second terms respectively of M_T when

$$(M_{inv}^{3\pi})^2 - ((E_T^{3\pi})^2 - (p_T^{3\pi})^2) \geq 0. \quad (\text{B.12})$$

Writing $\vec{p}^{3\pi} = \vec{p}^{\pi^+} + \vec{p}^{\pi_1^-} + \vec{p}^{\pi_2^-}$ and similarly for $E^{3\pi}$, $\vec{p}_T^{3\pi}$ and $E_T^{3\pi}$, and using equation (B.11) for $E_T^{3\pi}$, we evaluate the squares and then cancel terms, leaving

$$\begin{aligned} &2(E^{\pi^+} E^{\pi_1^-} - E_T^{\pi^+} E_T^{\pi_1^-} - p_L^{\pi^+} p_L^{\pi_1^-}) + \\ &2(E^{\pi^+} E^{\pi_2^-} - E_T^{\pi^+} E_T^{\pi_2^-} - p_L^{\pi^+} p_L^{\pi_2^-}) + \\ &2(E^{\pi_1^-} E^{\pi_2^-} - E_T^{\pi_1^-} E_T^{\pi_2^-} - p_L^{\pi_1^-} p_L^{\pi_2^-}) \geq 0. \end{aligned} \quad (\text{B.13})$$

It will be sufficient for the inequality (B.12) to hold if each of the terms in parentheses in (B.13) is non-negative definite; for simplicity of notation we write a and b for a pair of pions (say π^+ and π_1^-), and must show

$$E^a E^b - p_L^a p_L^b \geq E_T^a E_T^b. \quad (\text{B.14})$$

Both sides of this expression are positive, so squaring preserves the inequality,

$$\begin{aligned} [m_\pi^2 + (p_T^a)^2 + (p_L^a)^2] [m_\pi^2 + (p_T^b)^2 + (p_L^b)^2] + (p_L^a)^2 (p_L^b)^2 - 2E^a E^b p_L^a p_L^b \\ \geq [m_\pi^2 + (p_T^a)^2] [m_\pi^2 + (p_T^b)^2], \end{aligned}$$

where the E^2 terms have been expanded. Cancelling terms and regrouping, we have

$$(E^b)^2 (p_L^a)^2 + (E^a)^2 (p_L^b)^2 - 2E^a E^b p_L^a p_L^b \geq 0.$$

When $p_L^a p_L^b < 0$, this holds trivially. When $p_L^a p_L^b \geq 0$ we may change sides and square:

$$(p_L^a)^4 (E^b)^4 + 2(E^a E^b p_L^a p_L^b)^2 + (p_L^b)^4 (E^a)^4 \geq 4(E^a E^b p_L^a p_L^b)^2.$$

Regrouping, we have

$$[(p_L^a)^2 (E^b)^2 - (p_L^b)^2 (E^a)^2]^2 \geq 0$$

which always holds; thus the inequality (B.14) holds for each of the three 2π combinations in (B.13), so inequality (B.12) holds always. That is,

$$M_\perp < M_T \leq M_R. \tag{B.15}$$

B.3 The interpretation of M_R as the “required mass” of the tau

As noted in section B.1 above, the quantity M_R was originally introduced to allow a solution of equation (B.5) in $p_L^{\nu\tau}$, in cases where no solution for $M = m_\tau$ exists. Labelling the solutions of the quadratic as $p_L^{(1)}$ and $p_L^{(2)}$ (with $p_L^{(1)} > p_L^{(2)}$), then, a sensible choice is

$$p_L^{\nu\tau} = \begin{cases} p_L^{(1)}(M = m_\tau), & \text{when } \Delta(M = m_\tau) \geq 0, \\ p_L^{(1)}(M = M_R), & \text{otherwise,} \end{cases}$$

the second of these choices being unique as $\Delta(M = M_R) = 0$ by construction. Instead, since we are not using the $p_L^{\nu\tau}$ estimate in the analysis, we have interpreted the quantity M_R in all cases as the *minimum* mass M which can be ascribed to the τ in order to obtain a solution for $p_L^{\nu\tau}$, and thus for the 3π -missing momentum system to be interpreted as a tau decay. This interpretation requires that there are no cases where $M_R > M'$ but a solution exists for $M = M'$ (*i.e.* $\Delta(M = M') \geq 0$). In the Monte Carlo such cases, although rare, can be found for $\nu_\mu + N \rightarrow \nu_\mu + X$ events analysed as tau decays, where $M' = m_\tau$; such cases *cannot* be found when a cut at or below the tau mass is placed on the invariant mass of the 3π system. In fact any solution where $\Delta(M = M') \geq 0$ for $M' < M_R$ is due to an unphysical value of $M_{inv}^{3\pi}$, as will be shown in the remainder of this section.

Suppose that a solution to equation (B.5) exists for $M = M' < M_R$. That is,

$$\begin{aligned} (M_R)^2 &= (M_{inv}^{3\pi})^2 + 2p_T^{miss} \sqrt{(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2} - 2\vec{p}_T^{miss} \cdot \vec{p}_T^{3\pi} \\ &> (M')^2 \end{aligned} \tag{B.16}$$

(compare with equation (B.8)) and

$$\begin{aligned}
 \Delta &= (p_L^{3\pi} D)^2 - 4 [(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2] \left[(p_T^{miss})^2 (E^{3\pi})^2 - \frac{D^2}{4} \right] \\
 &= (E^{3\pi})^2 \{ D^2 - 4(p_T^{miss})^2 [(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2] \} \\
 &= (E^{3\pi})^2 \left\{ [(M')^2 - (M_{inv}^{3\pi})^2 + 2\vec{p}_T^{miss} \cdot \vec{p}_T^{3\pi}]^2 - 4(p_T^{miss})^2 [(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2] \right\} \\
 &> 0
 \end{aligned} \tag{B.17}$$

(compare with equation (B.7)). Writing

$$\begin{aligned}
 A &= (M')^2 \\
 B &= (M_{inv}^{3\pi})^2 - 2\vec{p}_T^{miss} \cdot \vec{p}_T^{3\pi} \\
 C &= 2p_T^{miss} \sqrt{(M_{inv}^{3\pi})^2 + (p_T^{3\pi})^2},
 \end{aligned}$$

the inequality (B.16) yields

$$A - [B + C] < 0, \tag{B.18}$$

and the inequality (B.17) yields

$$\begin{aligned}
 (A - B)^2 - C^2 &> 0, \\
 A^2 - 2AB + (B^2 - C^2) &> 0, \\
 (A - [B + C])(A - [B - C]) &> 0.
 \end{aligned} \tag{B.19}$$

Inequalities (B.18) and (B.19) together require

$$(A - [B - C]) < 0$$

and since $C \geq 0$, this also requires

$$(A - B) < 0,$$

that is,

$$(M')^2 - (M_{inv}^{3\pi})^2 + 2\vec{p}_T^{miss} \cdot \vec{p}_T^{3\pi} < 0. \tag{B.20}$$

Now M' must be interpretable as the invariant mass of the 3π -missing momentum system, in which case

$$\begin{aligned}
 (M')^2 &= (E^{miss} + E^{3\pi})^2 - (\vec{p}^{miss} + \vec{p}^{3\pi})^2 \\
 &= (M_{inv}^{3\pi})^2 + 2E^{miss} E^{3\pi} - 2\vec{p}_T^{miss} \cdot \vec{p}_T^{3\pi} - 2p_L^{miss} p_L^{3\pi}
 \end{aligned}$$

(compare with equation (B.1)), where p_L^{miss} is given by the corresponding solution of equation (B.5). Substitution into the inequality (B.20) yields

$$2(E^{miss} E^{3\pi} - p_L^{miss} p_L^{3\pi}) < 0$$

which is impossible. That is, for any 3π -missing momentum system, M_R is the smallest mass which can be ascribed to the τ if the system is to be interpreted as the product of $\tau^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$ decay, under the identification $\vec{p}_T^{\nu\tau} = \vec{p}_T^{miss}$.

Considering the inequality (B.20), it will be apparent that the cases mentioned above (where $M_R > m_\tau$ but $\Delta(M = m_\tau) > 0$) require that the mass of the 3π system $M_{inv}^{3\pi}$ be larger than the mass $M = M' = m_\tau$ ascribed to the tau: hence such cases are excluded by cuts on $M_{inv}^{3\pi}$. The unphysical solution to $\Delta(M) \geq 0$ arises because the two sides of an expression derived from the tau mass equation ((B.2) above) are squared (to give (B.3)) in the derivation of the equation for $p_L^{\nu\tau}$.

B.4 Bounds on the transverse mass given ideal measurement

We now prove the result quoted in section 3.4.4, that if $\vec{p}_T^{3\pi}$ etc. were perfectly measured, the transverse mass would be bounded above by the tau mass.

Assuming all measured quantities to correspond exactly to those of the particles, we may write for the tau mass

$$\begin{aligned} m_\tau^2 &= (E^{\nu\tau} + E^{3\pi})^2 - (\vec{p}^{\nu\tau} + \vec{p}^{3\pi})^2 \\ &= (E^{\nu\tau})^2 + 2E^{\nu\tau} E^{3\pi} + (E^{3\pi})^2 - (\vec{p}_T^{\nu\tau} + \vec{p}_T^{3\pi})^2 - (p_L^{\nu\tau} + p_L^{3\pi})^2, \end{aligned}$$

where we have partially expanded both squared terms; expanding the energies,

$$\begin{aligned} &= ((E_T^{\nu\tau})^2 + (p_L^{\nu\tau})^2) + 2E^{\nu\tau} E^{3\pi} + ((E_T^{3\pi})^2 + (p_L^{3\pi})^2) \\ &\quad - (\vec{p}_T^{\nu\tau} + \vec{p}_T^{3\pi})^2 - ((p_L^{\nu\tau})^2 + 2p_L^{\nu\tau} p_L^{3\pi} + (p_L^{3\pi})^2); \end{aligned} \tag{B.21}$$

regrouping,

$$\begin{aligned} &= \left[(E_T^{\nu\tau})^2 + (E_T^{3\pi})^2 - (\vec{p}_T^{\nu\tau} + \vec{p}_T^{3\pi})^2 \right] + 2E^{\nu\tau} E^{3\pi} - 2p_L^{\nu\tau} p_L^{3\pi} \\ &= [M_R^2 - 2E_T^{\nu\tau} E_T^{3\pi}] + 2E^{\nu\tau} E^{3\pi} - 2p_L^{\nu\tau} p_L^{3\pi}. \end{aligned} \tag{B.22}$$

So to show that $m_\tau \geq M_R$ under the assumptions listed, we must prove that

$$E^{\nu_\tau} E^{3\pi} - E_T^{\nu_\tau} E_T^{3\pi} \geq p_L^{\nu_\tau} p_L^{3\pi}. \quad (\text{B.23})$$

Since both sides of (B.23) are positive, it will be sufficient to show that

$$(E^{\nu_\tau} E^{3\pi} - E_T^{\nu_\tau} E_T^{3\pi})^2 \geq (p_L^{\nu_\tau} p_L^{3\pi})^2. \quad (\text{B.24})$$

Expanding the left-hand side, we have

$$\begin{aligned} & (E^{\nu_\tau})^2 (E^{3\pi})^2 - 2E^{\nu_\tau} E^{3\pi} E_T^{\nu_\tau} E_T^{3\pi} + (E_T^{\nu_\tau})^2 (E_T^{3\pi})^2 \\ &= [(E_T^{\nu_\tau})^2 + (p_L^{\nu_\tau})^2] [(E_T^{3\pi})^2 + (p_L^{3\pi})^2] - 2E^{\nu_\tau} E^{3\pi} E_T^{\nu_\tau} E_T^{3\pi} + (E_T^{\nu_\tau})^2 (E_T^{3\pi})^2; \end{aligned}$$

regrouping,

$$\begin{aligned} &= [(E_T^{\nu_\tau})^2 (E_T^{3\pi})^2 + (p_L^{\nu_\tau})^2 (E_T^{3\pi})^2] + [(E_T^{\nu_\tau})^2 (E_T^{3\pi})^2 + (E_T^{\nu_\tau})^2 (p_L^{3\pi})^2] \\ &\quad - 2E^{\nu_\tau} E^{3\pi} E_T^{\nu_\tau} E_T^{3\pi} + (p_L^{\nu_\tau} p_L^{3\pi})^2; \end{aligned}$$

and re-forming energies from $(E_T^2 + p_L^2)$ terms,

$$\begin{aligned} &= [E^{\nu_\tau} E_T^{3\pi}]^2 + [E_T^{\nu_\tau} E^{3\pi}]^2 - 2E^{\nu_\tau} E^{3\pi} E_T^{\nu_\tau} E_T^{3\pi} + (p_L^{\nu_\tau} p_L^{3\pi})^2 \\ &= (E^{\nu_\tau} E_T^{3\pi} - E_T^{\nu_\tau} E^{3\pi})^2 + (p_L^{\nu_\tau} p_L^{3\pi})^2 \\ &\geq (p_L^{\nu_\tau} p_L^{3\pi})^2 \end{aligned} \quad (\text{B.25})$$

as required; so (B.24) is proven.

That is, given ideal measurement,

$$M_R \leq m_\tau \quad (\text{B.26})$$

so that finally

$$M_\perp < M_T \leq M_R \leq m_\tau. \quad (\text{B.27})$$

B.5 A transverse-plane approximation to the required mass

In section 4.3.3 of the text it is noted that if the 3π invariant mass is treated as fixed,

$$M_{inv}^{3\pi} = m_a, \text{ constant,}$$

the required mass becomes a function of $\vec{p}_T^{3\pi}$ and \vec{p}_T^{had} only. That is, it becomes a purely *kinematic* variable in the terms of chapter 3; more specifically, it becomes a variable of the transverse plane. We may therefore express M_R in this approximation, as a function of the three variables $(\rho^{3\pi}, \rho^{had}, M)$ describing the transverse plane (sections 3.3.2 and 3.4.1):

$$\begin{aligned} M_R &= \sqrt{\left(\sqrt{m_a^2 + (p_T^{3\pi})^2} + p_T^{miss}\right)^2 - (\vec{p}_T^{3\pi} + \vec{p}_T^{miss})^2} \\ &= \sqrt{M^2 \left(\sqrt{\left(\frac{m_a}{M}\right)^2 + (\rho^{3\pi})^2} + [1 - \rho^{3\pi} - \rho^{had}]\right)^2 - M^2(\rho^{had})^2}. \end{aligned}$$

Writing now for compactness

$$\eta = \sqrt{\left(\frac{m_a}{M}\right)^2 + (\rho^{3\pi})^2} - \rho^{3\pi},$$

and dividing through by the transverse mass scale M , we have

$$\begin{aligned} \left(\frac{M_R}{M}\right)^2 &= ([1 + \eta] - \rho^{had})^2 - (\rho^{had})^2 \\ &= [1 + \eta]^2 - 2\rho^{had}[1 + \eta]. \end{aligned}$$

Given a value for M_R , then, we have an equation relating any one of the three transverse plane variables to the other two: expressing ρ^{had} as a function of $\rho^{3\pi}$ and M , say,

$$\rho^{had} = \frac{(1 + \eta)^2 - \left(\frac{M_R}{M}\right)^2}{2(1 + \eta)}$$

which describes a surface in $(\rho^{3\pi}, \rho^{had}, M)$ or equivalently, at fixed M , a contour in $(\rho^{3\pi}, \rho^{had})$. A cut $M_R < 2.2 \text{ GeV}/c$, in the fixed $M_{inv}^{3\pi} = m_a$ approximation, thus corresponds to contours

$$\rho^{had} = \frac{\left(1 + \sqrt{\left(\frac{m_a}{M}\right)^2 + (\rho^{3\pi})^2} - \rho^{3\pi}\right)^2 - \left(\frac{M_R}{M}\right)^2}{2\left(1 + \sqrt{\left(\frac{m_a}{M}\right)^2 + (\rho^{3\pi})^2} - \rho^{3\pi}\right)} \quad (\text{B.28})$$

on the $(\rho^{3\pi}, \rho^{had})$ plot as a function of the transverse mass scale M . These are the contours shown in figure 4.13, with M_R set to the value $2.2 \text{ GeV}/c$.

Note that if $m_a \rightarrow 0$, then $\eta \rightarrow 0$, and we recover the expression

$$\rho^{had} = \frac{1}{2} \left(1 - \left(\frac{M_R}{M} \right)^2 \right)$$

for the contour of a cut in M_{\perp} (equation 3.11), as we should, since $M_R \rightarrow M_{\perp}$ in the approximation of vanishing 3π mass.

Appendix C

Constraints on the opening angle of $\tau \rightarrow \nu 3\pi$ pions

In section 4.4.2 of the text it was demonstrated that individual opening angles θ of the 3π system scale as $1/\gamma^{3\pi}$ when the three-pion system is energetic ($\gamma^{3\pi} \gg 1$). The same behaviour holds for the largest of these angles, θ_{\max} . This immediately raises the question of the form of the distribution of the scale-free part of the opening angle, $\gamma^{3\pi}\theta_{\max}$. It can be shown that once the invariant mass of the 3π system is decided, the form of the $\gamma^{3\pi}\theta_{\max}$ distribution is largely set, regardless of whether the 3π system is the true product of a decay $a_1^- \rightarrow \pi^- \pi^+ \pi^-$, or a group of pions arbitrarily chosen.

The arithmetic involved is straightforward but lengthy, and is included here as an appendix. The basic result, relating the opening angle of a pion in the lab to its momentum in the rest-frame of the 3π system, is derived in C.1. This is then applied to pions in the backward hemisphere, $\theta^* > 90^\circ$, in the rest-frame of the 3π system (C.2); then to pions in the forward hemisphere (C.3). Finally, in C.4, the resulting constraints on the distribution of $\gamma^{3\pi}\theta_{\max}$ are deduced.

C.1 The constraint on an individual pion

C.1.1 Motivation

In section 4.4.2 it was shown that for a decay $A \rightarrow B + C (+D + \dots)$ of a particle with velocity β and Lorentz factor γ , a decay product B with




Figure C.1: The components of the momentum \vec{p}_T of a pion in the laboratory frame, in terms of the pion momentum in the 3π -rest frame, \vec{p}_T^* . The direction of $\vec{p}^{3\pi}$ in the laboratory is shown; θ is the angle between the pion \vec{p} and $\vec{p}^{3\pi}$.

momentum \vec{p}^* and energy E^* in the rest-frame of A emerges at an angle

$$\begin{aligned} \theta \simeq \tan \theta &= \frac{p_T}{p_L} \\ &= \frac{p_T^*}{\gamma(p_L^* + \beta E^*)} \end{aligned} \quad (\text{C.1})$$

to the direction of motion of A in the laboratory: this situation is shown in figure C.1 where $A = 3\pi$ and B is one of the decay pions. (We use the direction of motion of A in the lab as the reference or longitudinal direction, and ignore the azimuth about this direction throughout.) From this we see immediately that a slow pion with $p^* \ll m_\pi$ (where $E^* \approx m_\pi$) will emerge at a very small angle $\theta \ll 1/\gamma$ regardless of its polar angle θ^* in the rest-frame of A . In this case, however, at least one of the remaining pions will have a much larger momentum, taking a large fraction of the energy of the 3π system. For such a pion it turns out that there are non-trivial lower bounds on the angle θ in the laboratory for large regions in θ^* .

In the more general case, assuming $M_{inv}^{3\pi}$ close to the a_1 mass (which is large, $m_{a_1} = 1.230 \text{ GeV}/c^2 \gg m_\pi$) at least one of the three pions *must* have significant momentum $p^* > m_\pi$, and in many cases two or all three of the pions are relativistic in the rest-frame of A . By construction of the 3π rest-frame $\sum_{i \in 3\pi} \vec{p}_i^* \equiv \vec{0}$, and so

- all of the pion momenta \vec{p}_i^* lie in a plane, and
- the polar angle θ_i^* of a pion, as well as its momentum p^* , is determined once the momenta \vec{p}_i^* of the other two pions are set.

Because of this, the weak condition that the 3π system be close to the a_1 mass will be sufficient to derive an approximate lower bound on the *largest* of the opening angles in the laboratory, θ_{\max} . An upper bound and the general form of the distribution will also follow.

C.1.2 The condition to escape a cone of given opening angle

To determine these bounds we must first derive the condition for a particle B of momentum \vec{p}^* in the rest-frame of A to reach a certain angle θ with respect to A in the laboratory. Since we are interested in energetic 3π systems so that $\gamma \gg 1$, we see from (C.1) that the angle θ obeys an approximate scale $1/\gamma$; it is therefore convenient to measure angles in terms of this scale. We require that B escape the cone of opening angle a/γ around the direction of A , where a is some (positive) constant:

$$\begin{aligned} \frac{p_T^*}{\gamma(p_L^* + \beta E^*)} &> \frac{a}{\gamma} \\ p_T^* &> a(p_L^* + \beta E^*) \\ (p_T^* - ap_L^*) &> a\beta E^*. \end{aligned} \quad (\text{C.2})$$

The term on the right is positive definite, so we find that $p_T^* < ap_L^*$ will be sufficient (although not necessary) for B to be *contained* by the a/γ cone in the laboratory. Considering instead the case $p_T^* > ap_L^*$, we may square both sides of the inequality (C.2), obtaining

$$(p_T^*)^2 - 2ap_T^*p_L^* + a^2(p_L^*)^2 > a^2\beta^2 [m^2 + (p_T^*)^2 + (p_L^*)^2],$$

where the $(E^*)^2$ term has been expanded. Writing $\beta^2 = 1 - 1/\gamma^2$, then,

$$\left[\frac{a^2}{\gamma^2} - (a^2 - 1) \right] (p_T^*)^2 + \frac{a^2}{\gamma^2} (p_L^*)^2 - 2ap_T^*p_L^* > a^2 \left(1 - \frac{1}{\gamma^2}\right) m^2.$$

Collecting the a^2/γ^2 terms, and writing $p_L^* = p^* \cos \theta^*$ and $p_T^* = p^* \sin \theta^*$, we obtain

$$\frac{1}{\gamma^2} (E^*)^2 - \left(\frac{a^2 - 1}{a^2} \right) (p^*)^2 \sin^2 \theta^* - \left(\frac{2}{a} \right) (p^*)^2 \sin \theta^* \cos \theta^* > m^2.$$

Finally then,

$$\sin \theta^* \left[\left(\frac{2}{a} \right) \cos \theta^* + \left(\frac{a^2 - 1}{a^2} \right) \sin \theta^* \right] < -k, \quad (\text{C.3})$$

which will be the basic condition from which the remaining results will follow.

The left-hand side of (C.3) depends only on the decay angle θ^* ; the factor k is given by

$$k \stackrel{\text{def}}{=} \frac{m^2 - \frac{1}{\gamma^2}(E^*)^2}{(p^*)^2} = \frac{(1 - \frac{1}{\gamma^2})m^2}{(p^*)^2} - \frac{1}{\gamma^2} \quad (\text{C.4})$$

$$\rightarrow \frac{m^2}{(p^*)^2} \text{ as } \gamma \rightarrow \infty, \quad (\text{C.5})$$

the deviation from scaling with $1/\gamma$ being contained in the variation of this term with γ , which is small for $\gamma \gg 1$.

As we will be applying these results to $\tau \rightarrow \nu 3\pi$ candidates which are energetic, we will ignore the case where γ is of order unity and the factor k is potentially negative: we will assume $k > 0$ throughout. It will be convenient to work with the limiting value

$$k_\infty \stackrel{\text{def}}{=} \frac{m^2}{(p^*)^2} \quad (\text{C.6})$$

for most purposes.

C.1.3 Variation of the factor k with momentum

Clearly for $p^* \ll m$, the factor k becomes large and solutions of the inequality (C.3) will be possible only for small $a \ll 1$: that is, the decay product will emerge at a small angle $\theta \ll 1/\gamma$ in the lab. This result was already derived from the basic formula (C.1) in section C.1.1 above, where it was noted that in a 3π system of a_1 -like mass, such a slow pion will be accompanied by at least one fast pion ($p^* > m_\pi$).

It is these faster pions on which the constraints on θ_{max} will ultimately depend, and the two most useful reference cases are as follows:

The maximum pion momentum

The largest possible pion momentum is attained where one pion emerges back-to-back with two slower, collinear pions which share the balancing momentum. Now for a decay to three particles, the decay momentum p^* obeys

$$p^* = \frac{1}{2M} \sqrt{[M^2 - (m_{12} + m)^2][M^2 - (m_{12} - m)^2]} \quad (\text{C.7})$$

where m_{12} is the invariant mass of the system formed by the other two decay products (compare (Caso et al., 1998, equation 35.20b)). In this case $m_{12} = 2m_{\pi^\pm}$ so we have

$$\begin{aligned} p_{\max}^* &= \frac{1}{2M} \sqrt{(M^2 - 3m_\pi^2)(M^2 - m_\pi^2)} \\ &= 0.599 \text{ GeV}/c \text{ for } M = 1.230 \text{ GeV}/c^2, \end{aligned} \quad (\text{C.8})$$

for which

$$k_\infty = 0.054, \quad (\text{C.9})$$

where the central value for the a_1 mass distribution has been used. Setting the full width (which is uncertain) to $\Gamma_{a_1} = 0.400 \text{ GeV}/c^2$, the factor k_∞ changes from 0.127 through 0.030 as we move across the distribution from $M = (m_{a_1} - \Gamma_{a_1})$ to $(m_{a_1} + \Gamma_{a_1})$. The constraints derived below will vary only weakly with k for values of this type, where k is significantly less than unity.

The momentum of equally-shared-energy

Since slower pions will be confined to smaller angles θ in the lab, but we are interested in the *largest* of these angles θ_{\max} , it is instructive to consider the smallest possible value for the largest of the three rest-frame momenta, $\inf\{\max\{p_i^*, i \in 3\pi\}\}$. Since each of the three pions has the same mass, this is given by the momentum for which the energy of the system is shared equally among the three pions. Trivially

$$\begin{aligned} p_{\text{share}}^* &= \sqrt{\left(\frac{M}{3}\right)^2 - m^2} \\ &= 0.386 \text{ GeV}/c \text{ for } M = 1.230 \text{ GeV}/c^2, \end{aligned} \quad (\text{C.10})$$

for which

$$k_\infty = 0.131, \quad (\text{C.11})$$

this value changing from 0.342 through 0.071 as M moves across the a_1 mass distribution. It should be noted that in this case all three pions are relativistic, and are separated from each other by 120° in the plane of the decay.

C.2 Decay angles in the backward hemisphere

We now consider solutions to the fundamental inequality (C.3) for particles emerging in the backward hemisphere, i.e. $\pi/2 \leq \theta^* \leq \pi$. Setting $\psi = \pi - \theta^*$ we have $\sin \psi = \sin \theta^*$ and $\cos \psi = \cos \theta^*$; writing $s = \sin \psi$ and $\sqrt{1 - s^2} = \cos \psi$, then,

$$s \left[-\left(\frac{2}{a}\right) \sqrt{1 - s^2} + \left(\frac{a^2 - 1}{a^2}\right) s \right] < -k, \\ \left(\frac{2}{a}\right) s \sqrt{1 - s^2} > k + \left(\frac{a^2 - 1}{a^2}\right) s^2.$$

The left-hand side of this inequality is non-negative definite: when the right-hand side is negative (for some $a < 1$), the expression holds trivially; otherwise we may square both sides, yielding

$$\left(\frac{4}{a^2}\right) s^2(1 - s^2) > k^2 + 2k \left(\frac{a^2 - 1}{a^2}\right) s^2 + \left(\frac{a^2 - 1}{a^2}\right)^2 s^4$$

and so

$$\left(\frac{a^2 - 1}{a^2}\right)^2 s^4 + \left(\frac{2k(a^2 - 1) - 4}{a^2}\right) s^2 + k^2 < 0. \quad (\text{C.12})$$

The left-hand side is a quadratic expression in s^2 : given roots $s_{(1)}^2 > s_{(2)}^2$ (so that we have $(s^2 - s_{(1)}^2)(s^2 - s_{(2)}^2) < 0$), the inequality will be satisfied for $s_{(2)}^2 < s^2 < s_{(1)}^2$. The two roots are given by

$$s^2 = \frac{a^2}{(a^2 + 1)^2} \left(-k(a^2 - 1) + 2 \pm \sqrt{k^2(a^2 - 1)^2 - 4k(a^2 - 1) + 4 - k^2(a^2 + 1)^2} \right),$$

so that the $\theta > a/\gamma$ condition will be satisfied when $\sin \psi$ lies between the two values

$$\frac{a}{a^2 + 1} \sqrt{2 - k(a^2 - 1) \pm 2\sqrt{(1 + k)(1 - ka^2)}}, \quad (\text{C.13})$$

having simplified the expression under the inner square root.

C.2.1 Escape to large angles $\theta > 1/\gamma$

For the case $a = 1$ the inequality takes the simple form

$$\sqrt{\frac{1 - \sqrt{1 - k^2}}{2}} < \sin \psi < \sqrt{\frac{1 + \sqrt{1 - k^2}}{2}}; \quad (\text{C.14})$$

and as noted above, for $k \ll 1$, the variation of these bounds with k will be weak. For the case of equally shared energy ($p^* = p_{\text{share}}^*$ and $M = m_{a_1}$, so $k_\infty = 0.131$) the bounds are $0.066 < \sin \psi < 0.998$, corresponding to the range of polar angles $176^\circ > \theta^* > 94^\circ$; for the maximum pion momentum ($p^* = p_{\text{max}}^*$, $k_\infty = 0.054$) the range is slightly larger, $178^\circ > \theta^* > 92^\circ$.

That is, a fast pion entering almost any part of the backward hemisphere in the rest-frame of A will emerge with an opening angle $\theta > 1/\gamma$ in the laboratory.

As the factor a is increased, the region which a fast decay pion must enter to escape to large angles $\theta > a/\gamma$ becomes progressively smaller. Some sample values are given in table C.1.

a	shared momentum p_{share}^*			maximum momentum p_{max}^*		
	$\theta > a/\gamma$ region		solid angle $\Omega/4\pi$	$\theta > a/\gamma$ region		solid angle $\Omega/4\pi$
	θ_{min}^*	θ_{max}^*		θ_{min}^*	θ_{max}^*	
1.0	93°	177°	0.47	92°	178°	0.49
1.5	119°	174°	0.26	115°	178°	0.29
2.0	135°	171°	0.138	130°	177°	0.177
2.5	149°	167°	0.059	141°	176°	0.112
3.0	—	—	—	148°	175°	0.071
3.5	—	—	—	155°	173°	0.044
4.0	—	—	—	161°	171°	0.022

Table C.1: The range of decay angles θ^* for which a pion of momentum $p^* = p_{\text{share}}^*$ (left) and $p^* = p_{\text{max}}^*$ (right) will escape a cone of opening angle $\theta = a/\gamma$ about the direction of a decaying a_1 in the laboratory frame (for $\gamma^{3\pi} \gg 1$).

The solutions (C.13) exist only for $ka^2 < 1$, so it follows that

$$\begin{aligned}
 a_{\text{max}} &= \frac{1}{\sqrt{k}} & (C.15) \\
 &= 2.76 \text{ for } p^* = p_{\text{share}}^*, \text{ and} \\
 &= 4.23 \text{ for } p^* = p_{\text{max}}^*
 \end{aligned}$$

yields the largest opening angle a_{max}/γ to which any pion can escape. Only the fastest pions can escape to angles of a significant multiple of $1/\gamma$

C.2.2 Lower bounds on θ for the backward hemisphere

From the table we can see that there is a region at $\theta^* \approx 180^\circ$, which grows slowly with a , for which pions fail to escape the $\theta = a/\gamma$ cone due to their

small p_T^* ; and a region around the “rim” of the backward hemisphere, $\theta^* > 90^\circ$, which grows quickly with a , for which the pion p_T^* is large but the value of $|p_L^*|$ ($p_L^* < 0$) is sufficiently small that the particle is “thrown forward” by the Lorentz boost into a cone $\theta < a/\gamma$. Pions emerging at intermediate angles in the rest-frame of the decaying particle emerge at large angles $\theta > a/\gamma$.

For sufficiently small a values the small-angle region at $\theta^* \approx 90^\circ$ closes for pions of a given momentum p^* , *i.e.* the larger solution of (C.13) goes to one,

$$\frac{a^2}{(a^2 + 1)^2} \left(-k(a^2 - 1) + 2 + \sqrt{(1+k)(1-ka^2)} \right) = 1,$$

expanding to

$$(k+1)^2 a^8 + 2k(k+1)a^6 + (k^2 - 2k - 2)a^4 - 2ka^2 + 1 = 0$$

with unique solution

$$\begin{aligned} a_{\text{lim}}^{90} &= \frac{1}{\sqrt{1+k}} & (C.16) \\ &= 0.940 \text{ for } p^* = p_{\text{share}}^*, \text{ and} \\ &= 0.974 \text{ for } p^* = p_{\text{max}}^* \end{aligned}$$

The smaller solution of (C.13) never goes to zero (assuming $k > 0$ throughout), but the solid angle subtended by the $\theta^* \approx 180^\circ$ region is always small (see the table), so we calculate the fraction of (4π) solid angle subtended when $a = a_{\text{lim}}^{90}$ and find

$$\begin{aligned} \frac{\Omega_{\text{lim}}}{4\pi} &= \frac{1 + \frac{k}{2} - \sqrt{1+k}}{k+2} & (C.17) \\ &= 9.5 \times 10^{-4} \text{ for } p^* = p_{\text{share}}^*, \text{ and} \\ &= 1.7 \times 10^{-4} \text{ for } p^* = p_{\text{max}}^*; \end{aligned}$$

to obtain a more intuitive expression than (C.17) we expand the square root (for $k \ll 1$) and find

$$\frac{\Omega_{\text{lim}}}{4\pi} \approx \frac{k^2}{8(k+2)}. \quad (C.18)$$

So while a strict lower bound cannot be set, for a fast pion in the backward hemisphere the region yielding $\theta < a_{\text{lim}}^{90}/\gamma$ is tiny.

C.3 Decay angles in the forward hemisphere

Considering instead solutions to the inequality (C.3) for particles emerging in the *forward* hemisphere, i.e. $0 \leq \theta^* < \pi/2$: if we write $s = \sin \theta^*$ and $\sqrt{1-s^2} = \cos \theta^*$,

$$s \left[\left(\frac{2}{a} \right) \sqrt{1-s^2} + \left(\frac{a^2-1}{a^2} \right) s \right] < -k, \\ \left(\frac{2}{a} \right) s \sqrt{1-s^2} < -k - \left(\frac{a^2-1}{a^2} \right) s^2. \quad (\text{C.19})$$

Now the left-hand side of this expression is non-negative definite, so solutions will exist only when

$$\left(\frac{1-a^2}{a^2} \right) s^2 > k. \quad (\text{C.20})$$

Since this condition cannot be met for $a > 1$ we immediately have the result that *all pions, regardless of decay momentum p^* , which enter the forward hemisphere will emerge within the $\theta = 1/\gamma$ cone around the direction of the decaying particle*

This result depends only on the $k > 0$, i.e. $\gamma \gg 1$ condition noted above.

C.3.1 Lower bounds on θ near $\theta^* = 90^\circ$

For sufficiently low values $a < 1$, however, we can define a region near the rim of the forward hemisphere, $\theta^* = 90^\circ$, where pions will escape to angles $\theta > a/\gamma$ in the lab. Assuming the condition (C.20), we can square both sides of (C.19) yielding

$$\left(\frac{4}{a^2} \right) s^2 (1-s^2) < k^2 + 2k \left(\frac{a^2-1}{a^2} \right) s^2 + \left(\frac{a^2-1}{a^2} \right)^2 s^4$$

and so

$$\left(\frac{a^2-1}{a^2} \right)^2 s^4 + \left(\frac{2k(a^2-1)-4}{a^2} \right) s^2 + k^2 > 0. \quad (\text{C.21})$$

which is the same as the inequality (C.12), with the sense reversed. Given roots $s_{(1)}^2 > s_{(2)}^2$ (so that we have $(s^2 - s_{(1)}^2)(s^2 - s_{(2)}^2) > 0$), the inequality will be satisfied for $s^2 > s_{(1)}^2$ and $s^2 < s_{(2)}^2$, although the latter solution will in practice be ruled out by the additional condition (C.20) above. The $\theta > a/\gamma$

condition will thus be satisfied when $\sin \theta^*$ exceeds the larger of the values at (C.13), *i.e.*

$$\sin \theta^* > \frac{a}{a^2 + 1} \sqrt{2 - k(a^2 - 1) + 2\sqrt{(1+k)(1-ka^2)}} \quad (\text{C.22})$$

A pion of maximum momentum ($p^* = p_{\max}^*$, $k_\infty = 0.054$) at an angle $\theta^* > 79^\circ$ in the rest-frame of the 3π will thus escape to an angle $\theta > 0.8/\gamma$ in the lab.

For the following section it will be more practical to know the angle to which a particle of given momentum p^* will escape once its polar angle θ^* is determined. We may determine this by solving the earlier, fundamental expression (C.3) for a given θ^* and the factor k : after a calculation similar to those already shown, we have

$$a(\theta^*, k) = \frac{\sin \theta^*}{k + \sin^2 \theta^*} \left(\sqrt{1+k} - \cos \theta^* \right) \quad (\text{C.23})$$

C.4 Constraints on the distribution of $\gamma^{3\pi} \theta_{\max}$

Noting that the three pions are constrained to lie in a plane in the 3π -rest system, with $\sum_{i \in 3\pi} \vec{p}_i^* \equiv \vec{0}$, several constraints on $\gamma^{3\pi} \theta_{\max}$ follow straightforwardly from the various results derived above. We will list these in turn, and also relax some of the assumptions previously made, to determine how strictly these constraints will apply in non-ideal cases.

C.4.1 The upper bound on $\gamma^{3\pi} \theta_{\max}$

The simplest of the constraints is a straightforward upper bound: it was noted at (C.15) that $\theta < a_{\max}/\gamma$ regardless of decay angle θ^* , and for the fastest possible pion $a_{\max} = 4.23$; since this applies to each of the angles θ it applies to θ_{\max} , so $\gamma^{3\pi} \theta_{\max} < 4.23$.

We have up to now considered principally the case $M = m_{a_1}$: for smaller masses the bound is more strict, while for larger masses, taking the extreme value $M = (m_{a_1} + \Gamma_{a_1})$, the bound increases only to 5.8. From table C.1 it is clear that the available solid angle for cases approaching the upper bound is very small, even when p^* takes the largest possible value; thus values greater than (say) 5 will not in practice occur.

The bound will however not apply when the $\gamma^{3\pi} \gg 1$ condition is not met. In this case the $1/\gamma^2$ correction terms in equation (C.4) become important: scaling with $1/\gamma$ fails and a “tail” in $\gamma^{3\pi} \theta_{\max}$ grows as $\gamma^{3\pi}$ falls. So assuming the central mass $M = 1.23 \text{ GeV}/c^2$, for $\gamma^{3\pi} = 5$ and $p^* = p_{\max}^*$ the factor k

has fallen to 0.0121 (*c.f.* $k_{\infty} = 0.054$) and $a_{\max} = 9.1$. Because the rest of the analysis relies on the 3π system being energetic—indirectly through the q_T cut, and directly through cuts on $p^{3\pi}$ —such a case, where $p^{3\pi}$ is only 6 GeV/ c , will not affect the analysis in practice. For the more realistic $\gamma^{3\pi} = 10$, k falls to 0.044 and a_{\max} rises only to 4.8.

C.4.2 The shape of the $\gamma^{3\pi}\theta_{\max}$ distribution

In general the angle θ_{\max} will be due to a pion emerging in the backward hemisphere in the 3π -rest frame: consulting section C.2.1 and especially table C.1 we see that the available solid angle for the $\theta > a/\gamma$ region falls rapidly as a rises towards its limit a_{\max} ; furthermore, this limit falls rapidly with pion momentum. In the absence of conspiratorial effects on pion decay angles θ^* , therefore, the upper part of the $\gamma^{3\pi}\theta_{\max}$ distribution will be concave, approaching its limiting value as an asymptote.

On the other hand, since the three pions lie in a plane, at least one of the pions is constrained to be in the backward hemisphere (including the rim at $\theta^* = 90^\circ$). For a fast pion, as noted in section C.2.1, this implies $\theta > 1/\gamma$ for almost all decay angles θ^* . Even if a pion as slow as $p^* = \frac{1}{2}p_{\text{share}}^* = 0.193$ GeV/ c is considered, the $\theta > 1/\gamma$ region spans $0.345 \times 4\pi$ in solid angle, *i.e.* 69% of the backward hemisphere. (For this and slower pions, there are significant contributions to θ_{\max} from forward hemisphere pions, as will be discussed below.) The peak value of the $\gamma^{3\pi}\theta_{\max}$ distribution will therefore occur somewhere above one.

C.4.3 Approximate lower bounds for $\gamma^{3\pi}\theta_{\max}$

No true lower bound on $\gamma^{3\pi}\theta_{\max}$ may be derived, due to a handful of extreme cases: most importantly, where one of the $p^* = p_{\max}^*$ the pions are collinear, so that if the decay angle θ^* is very close to either 0° or 180° all pions will be confined to small angles $\theta \ll 1/\gamma$. By considering possible 3π configurations in turn, however, it will become clear that *apart from a small tail* the distribution of $\gamma^{3\pi}\theta_{\max}$ is bounded below.

A fast pion in the backward hemisphere

If the backward pion is fast then for almost all decay angles θ^* a strong limit $\theta > a_{\text{lim}}^{90}/\gamma$ applies, so that for fast pion momenta $p^* > p_{\text{share}}^*$, $\gamma^{3\pi}\theta_{\max}$ is bounded below by a number 0.940 or greater. For a small fraction of cases where the fast pion escapes at $\theta^* \approx 180^\circ$, it is thrown forward to small

angles by the Lorentz boost, so that if both companions lie in the forward hemisphere, $\gamma^{3\pi}\theta_{\max}$ will take on a value significantly less than 1.

The case of equally-shared-energy

In this case the small-angle region at $\theta^* \approx 180^\circ$ subtends less than $10^{-3} \times 4\pi$ in solid angle (see (C.17)), and supposing that one pion escapes through this region its companions are constrained to lie at $\theta^* \approx 60^\circ$ in the forward hemisphere. In this case we may apply the formula at (C.23), finding $a = 0.55$ for the forward pions. For all remaining decay angles at least one of the pions, lying in the backward hemisphere ($\theta^* \geq 90^\circ$), obeys $\theta > a_{\text{lim}}^{90}/\gamma$: $\gamma^{3\pi}\theta_{\max}$ is then bounded below by 0.940.

For neighbouring cases with a harder backward pion $p^* > p_{\text{share}}^*$, this bound increases in value. For $M > m_{a_1}$ the bound increases in value, while for $M = (m_{a_1} - \Gamma_{a_1})$ it falls only to 0.863. If the 3π system has only modest $\gamma^{3\pi} > 1$, then the factor $k < k_\infty$ and the bound increases in value, while the region at $\theta^* \approx 180^\circ$ shrinks.

A relatively soft backward pion

We now consider the remaining case, where the pion emerging in the backward hemisphere is relatively soft: $p^* = \frac{1}{2}p_{\text{share}}^* = 0.193 \text{ GeV}/c$ will be a suitable reference value. For the majority of decay angles $\theta^* > 90^\circ$ this pion will escape to angles $\theta > a_{\text{lim}}^{90}/\gamma$, that is, $> 0.81/\gamma$ in this case.

The region at $\theta^* \approx 180^\circ$ for which the pion is contained by a cone $\theta < a_{\text{lim}}^{90}/\gamma$ takes up $0.011 \times 4\pi$ in solid angle, i.e. 2% of the backward hemisphere, and this will grow as the momentum p^* falls below the value considered. A soft pion of this kind, however, takes only a small fraction of the 3π energy and its companions become energetic: if they share the remaining energy equally they reach $p^* = 0.476 \text{ GeV}/c$, i.e. 80% of p_{max}^* . If the soft pion emerges at $\theta^* \approx 180^\circ$ these companions emerge at $\theta^* \approx 78^\circ$, close to the rim of the forward hemisphere, and an angle $\theta = 0.79/\gamma$ in the lab (equation C.23). Even a slight unequal sharing of energy forces one of the companions to wider angles θ^* , and angles close to $1/\gamma$ in the lab; the same applies if the slow pion emerges at (say) $\theta^* = 170^\circ$, close to the boundary of the $\theta < a_{\text{lim}}^{90}/\gamma$ region.

So when the backward pion is as soft as $p^* = \frac{1}{2}p_{\text{share}}^*$ or softer, the limit on θ_{\max} due to the companion pions becomes comparable to or greater than that of the backward pion. Values significantly below $\gamma^{3\pi}\theta_{\max} \approx 0.8$ cannot therefore be reached in this configuration. For 3π masses $M > m_{a_1}$ the limit hardens slightly, while for masses as low as $M = (m_{a_1} - \Gamma_{a_1})$, the limit softens only to ≈ 0.75 . The backward pion is largely unaffected if the $\gamma^{3\pi} \gg 1$

condition fails, but the limits on the companions become *more* strict.

C.4.4 Summary

Bringing together these results,

- there is an approximate lower bound on $\gamma^{3\pi}\theta_{\max}$ at 0.8, with a tail to ≈ 0.5 (and perhaps smaller values, due to energetic pions escaping at $\theta^* = 180^\circ$, but with very low probability);
- the $\gamma^{3\pi}\theta_{\max}$ distribution then rises to a peak above 1;
- the upper part of the distribution falls asymptotically to a limit above 5.

The shape of the distribution is thus largely determined: in particular, apart from low-probability tails it is bounded. It is worth noting that masses close to m_{a_1} , and an energetic 3π system, have been the only strong assumptions made to derive these results, so the constraints deduced should hold for groups of (unrelated) pions selected by cuts, just as well as for true $a_1^- \rightarrow \pi^- \pi^+ \pi^-$ decay pions. It is true that an implicit appeal has been made throughout to a distribution of 3π systems in a phase space, but none of the results depend on a *uniform* distribution in such a space—only that the pathological cases do not occur with high probability. The constraints listed above should therefore be robust.

Bibliography

- J.N. Abdurashitov et al. Results from SAGE (the Russian-American gallium solar neutrino experiment). *Physics Letters B*, 328:234–248, 1994.
- J.N. Abdurashitov et al. Measurement of the response of a gallium metal solar neutrino experiment to neutrinos from a ^{51}Cr source. *Physical Review C*, 59:2246–2263, 1999.
- G.S. Abrams et al. Measurements of Z^0 -boson resonance parameters in e^+e^- annihilation. *Physical Review Letters*, 63:2173, 1989.
- P. Abreu et al. Measurements of the τ polarisation in Z^0 decays. *Zeitschrift für Physik C – Particles and Fields*, 67:183–201, 1995.
- M. Acciarri et al. A measurement of τ polarization at LEP. *Physics Letters B*, 341:245–256, 1994.
- B. Achkar et al. Search for neutrino oscillations at 15, 40 and 95 metres from a nuclear power reactor at Bugey. *Nuclear Physics B*, 434:503–534, 1995.
- G. Acquistapace et al. The CERN neutrino beam to Gran Sasso (NGS): Conceptual technical design. CERN 98-02, CERN and INFN, Geneva, 1998.
- G. Acquistapace, V. Falaleev, J.M. Maugain, G. Olesen, S. Rangod, and J. Zaslavsky. The West Area Neutrino Facility for CHORUS and NOMAD experiments (94–97 operation). Technical Report CERN-ECP/95-14, CERN, 1995.
- Carl H. Albright and Robert E. Shrock. Signals for tau neutrino interactions in a beam dump experiment. *Physics Letters B*, 84(1):123–127, 1979.
- ALEPH, 1993. From the web pages of the ALEPH collaboration, at <http://alephwww.cern.ch/ALEPHGENERAL/reports/figures/ew/index.html>.

- W.W.M. Allison and J.H. Cobb. Relativistic charged particle identification by energy loss. *Annu. Rev. Nucl. Part. Sci.*, 30:253–298, 1980.
- W.W.M. Allison et al. Measurement of the atmospheric neutrino flavour composition in Soudan 2. *Physics Letters B*, 391:491–500, 1997.
- J. Altegoer et al. The NOMAD experiment at the CERN SPS. *Nucl. Instr. and Meth. in Phys. Res. A*, 404:96–128, 1998a.
- J. Altegoer et al. Search for a new gauge boson in π^0 decays. *Physics Letters B*, 428:197–205, 1998b.
- J. Altegoer et al. A search for $\nu_\mu \rightarrow \nu_\tau$ oscillations using the NOMAD detector. *Physics Letters B*, 431:219–236, 1998c.
- J. Altegoer et al. The trigger system of the NOMAD experiment. *Nucl. Instr. and Meth. in Phys. Res. A*, 428:299–316, 1999.
- L.W. Alvarez. The capture of orbital electrons by nuclei. *Physical Review*, 54:486, 1938.
- G. Ambrosini et al. K/ π production ratios from 450 GeV/c protons on beryllium. *Physics Letters B*, 420:225–232, 1998a.
- G. Ambrosini et al. Pion yield from 450 GeV/c protons on beryllium. *Physics Letters B*, 425:208–214, 1998b.
- G. Ambrosini et al. Production of π^\pm , K^\pm , p and \bar{p} from 450 GeV/c protons on beryllium. Accepted for publication in The European Physical Journal C., 1999.
- B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand. Parton fragmentation and string dynamics. *Physics Reports*, 97:31–145, 1983.
- M. Apollonio et al. Initial results from the CHOOZ long baseline reactor neutrino oscillation experiment. *Physics Letters B*, 420:397–404, 1998.
- M. Apollonio et al. Limits on neutrino oscillations from the CHOOZ experiment. Preprint hep-ex/9907037, 1999.
- N. Armenise et al. A new search for $\nu_\mu - \nu_\tau$ oscillations. Technical Report CERN-SPSC/90-42, CERN, 1990.
- R. Armenteros, K.H. Barker, C.C. Butler, and A. Cachon. The properties of neutral V-particles. *Philosophical Magazine, Series 7*, 42(333):1113–1135, 1951.

- M. Arneodo et al. Transverse momentum and its compensation in current and target jets in deep inelastic muon-proton scattering. *Physics Letters B*, 149:415–420, 1984.
- X. Artru, G.B. Yodh, and G. Mennessier. Practical theory of the multilayered transition radiation detector. *Physical Review D*, 12(5):1289–1306, 1975.
- K. Assamagan et al. Upper limit of the muon-neutrino mass and charged-pion mass from momentum analysis of a surface muon beam. *Physical Review D*, 53(11):6065–6077, 1996.
- P. Astier et al. Proposal – Search for the oscillation $\nu_\mu \rightarrow \nu_\tau$. Technical Report CERN-SPSLC/91-21, CERN, 1991a.
- P. Astier et al. Addendum to Proposal P261 – Search for $\nu_\mu - \nu_\tau$ oscillation. Technical Report CERN-SPSLC/91-48, CERN, 1991b.
- P. Astier et al. Addendum to Proposal P261 – Search for $\nu_\mu - \nu_\tau$ oscillation. Technical Report CERN-SPSLC/91-53, CERN, 1991c.
- P. Astier et al. A more sensitive search for $\nu_\mu \rightarrow \nu_\tau$ oscillations in NOMAD. *Physics Letters B*, 453:169–186, 1999a.
- P. Astier et al. Neutrino production of opposite sign dimuons in the NOMAD experiment. To be submitted to *The European Physical Journal C*, 1999b.
- C. Athanassopoulos et al. Candidate events in a search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations. *Physical Review Letters*, 75(14):2650–2653, 1995.
- C. Athanassopoulos et al. Evidence for neutrino oscillation from muon decay at rest. *Physical Review C*, 54(5):2685–2708, 1996.
- C. Athanassopoulos et al. Results on $\nu_\mu \rightarrow \nu_e$ oscillations from pion decay in flight neutrinos. *Physical Review C*, 58(4):2489–2511, 1998.
- D. Autiero et al. The electromagnetic calorimeter of the NOMAD experiment. *Nucl. Instr. and Meth. in Phys. Res. A*, 373:358–373, 1996a.
- D. Autiero et al. A high stability light emitting diode system for monitoring lead glass electromagnetic calorimeters. *Nucl. Instr. and Meth. in Phys. Res. A*, 372:556–561, 1996b.
- D. Autiero et al. Test beam performance of the electromagnetic calorimeter of the NOMAD experiment. *Nucl. Instr. and Meth. in Phys. Res. A*, 387:352–364, 1997.

- D. Autiero et al. A study of the transverse fluctuations of hadronic showers in the NOMAD electromagnetic calorimeter. *Nucl. Instr. and Meth. in Phys. Res. A*, 411:285–303, 1998.
- D. Autiero et al. Parametrization of e and γ initiated showers in the NOMAD lead-glass calorimeter. *Nucl. Instr. and Meth. in Phys. Res. A*, 425:188–209, 1999.
- A.S. Ayan. Letter of intent: A high sensitivity short baseline experiment to search for $\nu_\mu \rightarrow \nu_\tau$ oscillation. Technical Report CERN-SPSC/97-5, CERN, 1997.
- J.N. Bahcall and M.H. Pinsonneault. Solar models with helium and heavy-element diffusion. *Reviews of Modern Physics*, 67:781–808, 1995.
- John N. Bahcall, M.H. Pinsonneault, Sarbani Basu, and J. Christensen-Dalsgaard. Are standard solar models reliable? *Physical Review Letters*, 78(2):171–174, 1997.
- R. Barate et al. An upper limit on the τ neutrino mass from three- and five-prong tau decays. *The European Physical Journal C*, 2:395–406, 1998.
- G. Barichello et al. A B_4C -silicon target for the detection of neutrino interactions. *Nucl. Instr. and Meth. in Phys. Res. A*, 419:1–15, 1998a.
- G. Barichello et al. Performance of long modules of silicon microstrip detectors. *Nucl. Instr. and Meth. in Phys. Res. A*, 413:17–30, 1998b.
- M. Barranco Luque et al. The construction of the central detector for an experiment at the CERN $\bar{p} - p$ collider. *Nucl. Instr. and Meth. in Phys. Res.*, 176:175–180, 1980.
- G. Bassompierre et al. A large area transition radiation detector for the NOMAD experiment. *Nucl. Instr. and Meth. in Phys. Res. A*, 403:363–382, 1998a.
- G. Bassompierre et al. Performance of the NOMAD transition radiation detector. *Nucl. Instr. and Meth. in Phys. Res. A*, 411:63–74, 1998b.
- L. Baudis et al. Limits on the majorana neutrino mass in the 0.1 eV range. *Physical Review Letters*, 83(1):41–44, 1999.
- R. Becker-Szendy. Electron- and muon-neutrino content of the atmospheric flux. *Physical Review D*, 46(9):3720–3724, 1992.

- E.W. Beier et al. Survey of atmospheric neutrino data and implications for neutrino mass and mixing. *Physics Letters B*, 283:446–453, 1992.
- A.I. Belesev. Results of the Troitsk experiment on the search for the electron antineutrino rest mass in tritium beta-decay. *Physics Letters B*, 350:263–272, 1995.
- A. Benvenuti et al. Observation of muonless neutrino-induced inelastic interactions. *Physical Review Letters*, 32:800–803, 1974.
- F. Bergsma et al. A search for neutrino oscillations. *Zeitschrift für Physik C – Particles and Fields*, 40:171–192, 1988.
- H. Bethe and R. Peierls. The “neutrino”. *Nature*, 133:532, 1934.
- R.M. Bionta et al. Observation of a neutrino burst in coincidence with supernova 1987a in the Large Magellanic Cloud. *Physical Review Letters*, 58:1494–1496, 1987.
- J.D. Bjorken and S.L. Glashow. Elementary particles and SU(4). *Physics Letters*, 11:255–257, 1964.
- A. Bodek and J.L. Ritchie. Fermi-motion effects in deep-inelastic lepton scattering from nuclear targets. *Physical Review D*, 23:1070–1091, 1981.
- S.B. Boyd, 1997. Private communication.
- S.B. Boyd. *Neutrino production of opposite sign dimuons in the NOMAD experiment*. PhD thesis, University of Sydney, 1998.
- A. Bueno, J. Rico, and A. Rubbia. Study of muon rejection power in $\tau \rightarrow 1$ prong analysis using flat-top muons. NOMAD Memo 98–030, 1999.
- D. Buskulic et al. Measurement of the tau polarisation at the Z resonance. *Zeitschrift für Physik C – Particles and Fields*, 59:369–386, 1993.
- Nicola Cabibbo. Unitary symmetry and leptonic decays. *Physical Review Letters*, 10(12):531–533, 1963.
- Robert N. Cahn and Gerson Goldhaber. *The experimental foundations of particle physics*. Cambridge University Press, Cambridge, 1989.
- C. Caso et al. Review of particle physics. *The European Physical Journal C*, 3:1–794, 1998.

- A. Cavestro et al. An analog mean-timer for long scintillation counters. *Nucl. Instr. and Meth. in Phys. Res. A*, 305:488–491, 1991.
- Bruce T. Cleveland et al. Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *The Astrophysical Journal*, 496: 505–526, 1998.
- Claude Cohen-Tannoudji, Bernard Diu, and Franck Laloë. *Quantum mechanics*, volume 1. Hermann, Paris, France, 2nd edition, 1977. Translated from the French by S. Reid Hemley, N. Ostrowsky, D. Ostrowsky.
- M.J. Corden et al. Central hadron calorimeter of UA1. *Nucl. Instr. and Meth. in Phys. Res. A*, 238:273–287, 1985.
- Bob Cousins. Ills of the $\phi - \phi$ plot and how to cure them. NOMAD Memo 97-035, 1997a.
- Bob Cousins. Tau polarisation. Summary of $\nu_\mu \rightarrow \nu_\tau$ studies, presented at the NOMAD Collaboration Meeting, 5th September, 1997b.
- C.L. Cowan et al. Detection of the free neutrino: A confirmation. *Science*, 124:103, 1956.
- G.T. Danby et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Physical Review Letters*, 9:36–44, 1962.
- Arnon Dar and Giora Shaviv. Standard solar neutrinos. *The Astrophysical Journal*, 468:933–946, 1996.
- M. Davier, L. Duflot, F. Le Diberder, and A. Rougé. The optimal method for the measurement of tau polarization. *Physics Letters B*, 306:411–417, 1993.
- D. Decamp et al. Improved measurements of electroweak parameters from Z decays into fermion pairs. *Zeitschrift für Physik C – Particles and Fields*, 53:1–20, 1992.
- Avishai Dekel. Dynamics of cosmic flows. *Annu. Rev. Astron. Astrophys.*, 32:371–418, 1994.
- Luigi DiLella, 1999. Private communication.
- Loyal Durand. Transition radiation from ultrarelativistic particles. *Physical Review D*, 11(1):89–105, 1975.

- F. Dydak et al. A search for ν_μ oscillations in the Δm^2 range $0.3 - 90 \text{ eV}^2$. *Physics Letters B*, 134:281–286, 1984.
- K. Eggert et al. Large area drift tube chambers for a muon detector at the $\bar{p}p$ -collider. *Nucl. Instr. and Meth. in Phys. Res.*, 176:217–222, 1980.
- A. Ereditato, P. Strolin, and G. Romano. Study of a new experiment for the search of $\nu_\mu - \nu_\tau$ oscillations. *Nuclear Physics B (Proc. Suppl.)*, 54: 139–150, 1996.
- V.C. Ermilova, L.P. Kotenko, and G.I. Merzon. Fluctuations and the most probable values of relativistic charged particle energy loss in thin gas layers. *Nucl. Instr. and Meth. in Phys. Res.*, 145:555–563, 1977.
- E. Eskut et al. The CHORUS experiment to search for $\nu_\mu \rightarrow \nu_\tau$ oscillation. *Nucl. Instr. and Meth. in Phys. Res. A*, 401:7–44, 1997.
- E. Eskut et al. Search for $\nu_\mu \rightarrow \nu_\tau$ oscillation using the τ decay modes into a single charged particle. *Physics Letters B*, 434:205–213, 1998a.
- E. Eskut et al. A search for $\nu_\mu \rightarrow \nu_\tau$ oscillations. *Physics Letters B*, 424: 202–212, 1998b.
- V.V. Ezhela et al. *Particle physics: One hundred years of discoveries: an annotated chronological bibliography*. AIP Press, Woodbury, NY, 1996.
- T. Fazio, P. Nédélec, and S. Valuev. TRDSM: TRD simulation package. NOMAD Internal Software Note, 1994.
- G. Feldman and A. Geiser. Calculation of $\nu_\mu \rightarrow \nu_\tau$ oscillation limits. NOMAD Memo 99–001, 1999.
- G.J. Feldman and R.D. Cousins. A unified approach to the classical statistical analysis of small signals. *Physical Review D*, 57:3873–3889, 1998.
- E. Fermi. Towards the theory of β -rays. *Il Nuovo Cimento*, 11:1, 1934.
- R.P. Feynman and M. Gell-Mann. Theory of the fermi interaction. *Physical Review*, 109:193–198, 1958.
- R.A. Fisher. *Statistical methods for research workers*. Oliver and Boyd, Edinburgh, 4th edition, 1932.
- R. Foot, M.J. Thomson, and R.R. Volkas. Large neutrino asymmetries from neutrino oscillations. *Physical Review D*, 53:5349–5353, 1995.

- R. Foot and R.R. Volkas. Reconciling sterile neutrinos with big bang nucleosynthesis. *Physical Review Letters*, 75:4350–4353, 1995.
- R. Foot and R.R. Volkas. Studies of neutrino asymmetries generated by ordinary-sterile neutrino oscillations in the early universe and implications for big bang nucleosynthesis bounds. *Physical Review D*, 55:5147–5176, 1997.
- R. Foot, R.R. Volkas, and O. Yasuda. Comparing and contrasting the $\nu_\mu \rightarrow \nu_\tau$ and $\nu_\mu \rightarrow \nu_s$ solutions to the atmospheric neutrino problem with SuperKamiokande data. *Physical Review D*, 58:013006, 1998.
- Y. Fukuda et al. Atmospheric ν_μ/ν_e ratio in the multi-GeV range. *Physics Letters B*, 335:237–245, 1994.
- Y. Fukuda et al. Solar neutrino data covering solar cycle 22. *Physical Review Letters*, 77(9):1683–1686, 1996.
- Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Physical Review Letters*, 81:1562–1667, 1998a.
- Y. Fukuda et al. Measurement of the solar neutrino flux from Super-Kamiokande's first 300 days. *Physical Review Letters*, 81:1158–1162, 1998b.
- E. Gangler. *Recherche de l'oscillation $\nu_\mu \rightarrow \nu_\tau$ dans le canal $\tau \rightarrow 3\pi$* . PhD thesis, Université de Paris VI et VII, 1998.
- G.M. Garibian. Contribution to the theory of transition radiation. *JETP*, 6(6):1079–1085, 1958.
- G.M. Garibian. Transition radiation effects in particle energy losses. *JETP*, 10(2):372–376, 1960.
- G.M. Garibian, L.A. Gevorgian, and C. Yang. The calculation of x-ray transition radiation generated in regular- and irregular-layered media. *Nucl. Instr. and Meth. in Phys. Res.*, 125:133–137, 1975.
- A. Geiser, D. Autiero, and J. Long. First results from 11 module ν_μ CC data. Presented at the NOMAD Collaboration Meeting, 28th March, 1996.
- Achim Geiser, 1997. Private communication.
- M. Gell-Mann. A schematic model of baryons and mesons. *Physics Letters*, 8:214–215, 1964.

- V.L. Ginzburg and I.M. Franck. *Zh. Eksp. Teor. Fiz.*, 16:15, 1946.
- S.L. Glashow. Partial-symmetries of weak interactions. *Nuclear Physics*, 22: 579, 1961.
- S.L. Glashow, J. Iliopoulos, and L. Maiani. Weak interactions with lepton-hadron symmetry. *Physical Review D*, 2(7):1285–1292, 1970.
- Andrew Godley, 1999. Private communication.
- M. Goldhaber, L. Grodzins, and A.W. Sunyar. Helicity of neutrinos. *Physical Review*, 109:1015–1017, 1958.
- J.J. Gomez-Cadenas and J.A. Hernando. Search for $\nu_\mu(\nu_e) \leftrightarrow \nu_\tau$ oscillations with a detector based on a emulsion-silicon target. *Nucl. Instr. and Meth. in Phys. Res. A*, 381:223–235, 1996.
- J.J. Gomez-Cadenas, J.A. Hernando, and A. Bueno. A neutrino apparatus with improved capabilities for a short baseline $\nu_\mu(\nu_e) \leftrightarrow \nu_\tau$ search. *Nucl. Instr. and Meth. in Phys. Res. A*, 378:196–220, 1996.
- M.C. Gonzalez-Garcia and J.J. Gomez-Cadenas. Prompt ν_τ fluxes in present and future τ neutrino experiments. *Physical Review D*, 55:1297–1306, 1997.
- D.E. Groom. Searches for massive neutrinos. *The European Physical Journal C*, 3:320–321, 1998. Included in the Review of Particle Physics, C. Caso *et. al.*, pages 1–794.
- M. Gruwé et al. Search for $\nu_\mu \rightarrow \nu_\tau$ oscillation. *Physics Letters B*, 309: 463–468, 1993.
- Francis Halzen and Alan D. Martin. *Quarks and leptons: An introductory course in modern particle physics*. John Wiley & Sons, New York, 1984.
- W. Hampel et al. GALLEX solar neutrino observations: Results for GALLEX III. *Physics Letters B*, 388:384–396, 1996.
- W. Hampel et al. Final results of the ^{51}Cr neutrino source experiments in GALLEX. *Physics Letters B*, 420:114–126, 1998.
- F.J. Hasert et al. Observation of neutrino-like interactions without muon or electron in the Gargamelle neutrino experiment. *Physics Letters B*, 46: 138–140, 1973a.
- F.J. Hasert et al. Search for elastic muon-neutrino electron scattering. *Physics Letters B*, 46:121–124, 1973b.

- Naoya Hata and Paul Langacker. Solutions to the solar neutrino anomaly. *Physical Review D*, 56:6107–6166, 1997.
- E.H.M. Heijne. Muon flux measurement with silicon detectors in the CERN neutrino beams. Yellow Report 83–06, CERN, 1983.
- Peter .W. Higgs. Spontaneous symmetry breakdown without massless bosons. *Physical Review*, 145(4):1156–1163, 1966.
- P.W. Higgs. Broken symmetries and the masses of gauge bosons. *Physical Review Letters*, 13(16):508–509, 1964a.
- P.W. Higgs. Broken symmetries, massless particles and gauge fields. *Physics Letters*, 12(2):132–133, 1964b.
- K. Hikasa et al. Review of particle properties. *Physical Review D*, 45(11):S1–S584, 1992.
- James E. Hill. An alternative analysis of the LSND neutrino oscillation search data on $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations. *Physical Review Letters*, 75(14):2654–2657, 1995.
- A. Hime and N.A. Jelley. New evidence for the 17 keV neutrino. *Physics Letters B*, 257:441–449, 1991.
- A. Hime and J.J. Simpson. Evidence of the 17-keV neutrino in the β spectrum of ${}^3\text{H}$. *Physical Review D*, 39(7):1837–1850, 1989.
- E.P. Hincks and B. Pontecorvo. Search for gamma-radiation in the 2.2-microsecond meson decay process. *Physical Review*, 73:257–258, 1948.
- E.P. Hincks and B. Pontecorvo. The penetration of μ -meson decay electrons and their bremsstrahlung radiation. *Physical Review*, 75:698–699, 1949.
- K. Hirata et al. Observation of a neutrino burst from the supernova SN1987A. *Physical Review Letters*, 58:1490–1493, 1987.
- M. Holder et al. A detector for high-energy neutrino interactions. *Nucl. Instr. and Meth. in Phys. Res.*, 148:235–249, 1978.
- Peter Hurst, 1997a. Private communication.
- Peter Hurst. The performance of the hadron calorimeter. NOMAD Memo 97–042, 1997b.
- Peter Hurst, 1999. Private communication.

- Enrico Iacopini. Presentations at the NOMAD $\nu_\mu \rightarrow \nu_\tau$ Analysis Meetings, 8th August, 20th August and 3rd September, 1997.
- G. Ingelman. LEPTO version 6.1 — the Lund Monte Carlo for deep inelastic lepton-nucleon scattering. In W. Buchmüller and G. Ingelman, editors, *Physics at HERA*, pages 1366–1394, Hamburg, 1992. DESY.
- G.T. Jones et al. Study of transverse momenta of charged hadrons produced in νp and $\bar{\nu} p$ charged current interactions. *Zeitschrift für Physik C – Particles and Fields*, 25:121–127, 1984.
- Chung Wook Kim and Aihud Pevsner. *Neutrinos in physics and astrophysics*, volume 8 of *Contemporary Concepts in Physics*. Harwood, Langhorne, PA, 1993.
- Makato Kobayashi and Toshihide Maskawa. CP-violation in the renormalizable theory of weak interactions. *Progress of Theoretical Physics*, 49(2): 652–657, 1973.
- E.J. Konopinski and H.M. Mahmoud. The universal Fermi interaction. *Physical Review*, 92:1045–1049, 1953.
- J.H. Kühn and E. Mirkes. Angular distributions in semileptonic τ decays. *Physics Letters B*, 286:381–386, 1992a.
- J.H. Kühn and E. Mirkes. Structure functions in τ decays. *Zeitschrift für Physik C – Particles and Fields*, 56:661–671, 1992b.
- L.D. Landau. On one possibility for polarization properties of the neutrino. *Zh. Eksp. Teor. Fiz.*, 32:407, 1957. (JETP 5:337 (1957)).
- Benjamin W. Lee and Robert E. Shrock. Natural suppression of symmetry violation in gauge theories: Muon- and electron-lepton-number nonconservation. *Physical Review D*, 16(5):1444–1473, 1977.
- T.D. Lee, M.N. Rosenbluth, and C.N. Yang. Interaction of mesons with nucleons and light particles. *Physical Review*, 75:905, 1949.
- T.D. Lee and C.N. Yang. Parity nonconservation and two-component theory of the neutrino. *Physical Review*, 105:1671–1675, 1957a.
- T.D. Lee and C.N. Yang. Possible nonlocal effects in μ decay. *Physical Review*, 108(6):1611–1614, 1957b.

- R.B. Leighton, C.D. Anderson, and A.J. Seriff. The energy spectrum of the decay particles and the mass and spin of the mesotron. *Physical Review*, 75:1432–1437, 1949.
- C.H. Llewellyn-Smith. On the determination of $\sin^2 \theta_w$ in semileptonic neutrino interactions. *Nuclear Physics B*, 228:205–215, 1983.
- Thomas J. Loredo and Don Q. Lamb. Neutrinos from SN 1987A: Implications for cooling of the nascent neutron star and the mass of the electron antineutrino. *Annals of the New York Academy of Sciences*, 571:601–630, 1989.
- P.F. Loverre. Limits on neutrino oscillations derived from the ratio of neutral- to charged-current cross sections of neutrinos measured at the CERN SPS NBB. *Physics Letters B*, 206(4):711–714, 1988.
- P.F. Loverre. Limits on ν_μ oscillations from the measurement of the ratio of 0μ to 1μ events at the CERN narrow band neutrino beam. *Physics Letters B*, 370:156–158, 1996.
- E. Majorana. Symmetrical theory of electrons and positrons. *Il Nuovo Cimento*, 14:171, 1937.
- K.S. McFarland et al. Limits on $\nu_\mu(\bar{\nu}_\mu) \rightarrow \nu_\tau(\bar{\nu}_\tau)$ and $\nu_\mu(\bar{\nu}_\mu) \rightarrow \nu_e(\bar{\nu}_e)$ oscillations from a precision measurement of neutrino-nucleon neutral current interactions. *Physical Review Letters*, 75(22):3993–3996, 1995.
- S.P. Mikheyev and A.Yu. Smirnov. Resonance enhancement of oscillations in matter and solar neutrino spectroscopy. *Soviet Journal of Nuclear Physics*, 42:913–917, 1985.
- S.P. Mikheyev and A.Yu. Smirnov. Resonant amplification of ν oscillations in matter and solar-neutrino spectroscopy. *Il Nuovo Cimento C*, 9:17–26, 1986.
- K. Nakamura. Solar neutrinos. *The European Physical Journal C*, 3:327–330, 1998. Included in the Review of Particle Properties, C. Caso et al., pages 1–794.
- A. Naples et al. High statistics search for $\nu_e(\bar{\nu}_e) \rightarrow \nu_\tau(\bar{\nu}_\tau)$ oscillations. *Physical Review D*, 59:031101, 1998.
- K. Nishijima. Vanishing of the neutrino rest mass. *Physical Review*, 108:907–908, 1957.

- The NOMAD Collaboration. NOMAD progress report. Technical Report CERN-SPSLC/92-51, CERN, 1992.
- The NOMAD Collaboration. NOMAD status report. Technical Report CERN-SPSLC/93-19, CERN, 1993.
- Eric B. Norman et al. Evidence for the emission of a massive neutrino in nuclear beta decay. *Journal of Physics G: Nuclear and Particle Physics*, 17:S291–S299, 1991.
- W. Pauli. In *Rapports du Septième Conseil de Physique Solvay, 1933*, page 324, Paris, 1934. Gauthier-Villars.
- E.S. Pearson. The probability integral transformation for testing goodness of fit and combining independent tests of significance. *Biometrika*, 30: 134–148, 1938.
- M.L. Perl et al. Evidence for anomalous lepton production in e^+e^- annihilation. *Physical Review Letters*, 35:1489–1492, 1975.
- M.L. Perl et al. Properties of anomalous $e\mu$ events produced in e^+e^- annihilation. *Physics Letters B*, 63:466–470, 1976.
- M.L. Perl et al. Properties of the proposed τ charged lepton. *Physics Letters B*, 70:487–490, 1977.
- Y. Perrin et al. CASCADE: A toolkit for the construction of distributed, real-time, data-acquisition systems. Technical Report CERN-ECP/93-13, CERN, 1993.
- Massimo Persic, Paolo Salucci, and Fulvio Stel. The universal rotation curve of spiral galaxies - I. The dark matter connection. *Mon. Not. R. Astron. Soc.*, 281:27–47, 1996.
- O. Piccioni. Search for photons from meson-capture. *Physical Review*, 74: 1754–1758, 1948.
- B. Pontecorvo. Nuclear capture of mesons and the meson decay. *Physical Review*, 72:246–247, 1947.
- B. Pontecorvo. Inverse β -processes and lepton charge nonconservation. *Zh. Eksp. Teor. Fiz.*, 34:247, 1958.
- B. Pontecorvo. Electron and muon neutrinos. *Zh. Eksp. Teor. Fiz.*, 37: 1751–1757, 1959. (JETP 10:1236, 1960.).

- Joel R. Primack and Michael A.K. Gross. Cold + hot dark matter after superkamiokande. Preprint astro-ph/9810204, to be published in the proceedings of the tenth Rencontres de Blois, *The Birth of Galaxies.*, 1998.
- Joel R. Primack, Jon Holtzman, Anatoly Klypin, and David O Caldwell. Cold + hot dark matter cosmology with $m(\nu_\mu) \approx m(\nu_\tau) \approx 2.4 \text{ eV}$. *Physical Review Letters*, 74(12):2160–2163, 1995.
- Paolo Privitera. Structure functions in $\tau \rightarrow 3\pi\nu_\tau$ and the τ polarization measurement. *Physics Letters B*, 308:163–173, 1993. •
- S. Ravndal et al. GEANT: Detector description and simulation tool. CERN Program Library Long Writeup W5013, CERN, 1993.
- F. Reines and C.L. Cowan. Detection of the free neutrino. *Physical Review*, 92:830–831, 1953.
- F. Reines and C.L. Cowan. The neutrino. *Nature*, 178:446, 1956.
- A. Rougé. Polarization observables in the $3\pi\nu$ decay mode of the τ . *Zeitschrift für Physik C - Particles and Fields*, 48:75–77, 1990.
- A. Salam. Weak and electromagnetic interactions. *Elementary Particle Theory*, 1968.
- R.D. Sard and E.J. Althaus. Test of the hypothesis that the sea-level cosmic-ray meson disintegrates into a photon and an electron. *Physical Review*, 73:1251, 1948.
- K. Schahmaneche. *Recherche du charme dans les canaux hadroniques*. PhD thesis, Université de Paris VI et VII, 1997.
- Thierry Schmid. Travail de diplôme: étude des performances d'un prototype de détecteur de pied de gerbe. Université de Lausanne, 1996.
- B. Schmidt. *Study of ν_μ -Charged-Current Interactions and Search for $\nu_\mu \leftrightarrow \nu_\tau$ Oscillations in the Channel $\tau^- \rightarrow \mu^- \bar{\nu}_\mu \nu_\tau$ with the NOMAD Detector at the CERN Wideband Neutrino Beam*. PhD thesis, Universität Dortmund, 1997.
- M. Schwartz. Feasibility of using high energy neutrinos to study the weak interactions. *Physical Review Letters*, 4:306–307, 1960.
- Julian Schwinger. A theory of the fundamental interactions. *Annals of Physics*, 2:407–434, 1957.

- Frank Sciulli. What can 1000 GeV protons do for neutrino physics? In Earle C. Fowler, editor, *Neutrinos-78*, pages 863–870. Purdue University, 1978.
- R.E. Shrock. Neutrinos. *Physical Review D*, 50:1385–1389, 1994. Included in the Review of Particle Properties, L. Montanet *et al.*, pages 1173–1826.
- R.E. Shrock. Neutrinos. *Physical Review D*, 54:275–280, 1996. Included in the Review of Particle Physics, R.M. Barnett *et al.*, pages 1–720.
- J.J. Simpson. Evidence of heavy-neutrino emission in beta decay. *Physical Review Letters*, 54(17):1891–1893, 1985.
- J.J. Simpson and A. Hime. Evidence of the 17-keV neutrino in the β spectrum of ^{35}S . *Physical Review D*, 39(7):1825–1836, 1989.
- Torbjörn Sjöstrand. High-energy-physics event generation with PYTHIA 5.7 and JETSET 7.4. *Computer Physics Communications*, 82:74–89, 1994a.
- Torbjörn Sjöstrand. PYTHIA 5.7 and JETSET 7.4: Physics and manual. Technical Report CERN-TH.7112/93, CERN, 1994b.
- D. Steele. *A Search for Neutral, Heavy Particles Decaying to a Neutrino and a Single Photon at the SPS Wide-Band Neutrino Beam*. PhD thesis, Johns Hopkins University, Baltimore, 1996.
- I.E. Stockdale *et al.* Limits on muon-neutrino oscillations in the mass range $30 < \Delta m^2 < 1000 \text{ eV}^2/c^4$. *Physical Review Letters*, 52:1384–1388, 1984.
- Wolfgang Stoeffl and Daniel J. Decman. Anomalous structure in the beta decay of gaseous molecular tritium. *Physical Review Letters*, 75(18):3237–3240, 1995.
- E.C.G. Sudarshan and R.E. Marshak. Chirality invariance and the universal Fermi interaction. *Physical Review*, 109:1860–1862, 1958.
- Bhaskar Sur *et al.* Evidence for the emission of a 17-keV neutrino in the β decay of ^{14}C . *Physical Review Letters*, 66(19):2444–2447, 1991.
- Yasutaka Tanikawa. Renormalizable theory for Fermi interactions. *Physical Review*, 108(6):1615–1619, 1957.
- N. Ushida *et al.* Limits to $\nu_\mu, \nu_e \rightarrow \nu_\tau$ oscillations and $\nu_\mu, \nu_e \rightarrow \tau^-$ direct coupling. *Physical Review Letters*, 57(23):2897–2900, 1986.

- B. Van de Vyver and P. Zucchelli. Prompt ν_τ background in wide band ν_μ beams. *Nucl. Instr. and Meth. in Phys. Res. A*, 385:91–99, 1997.
- P. Vilain et al. Leading-order QCD analysis of neutrino-induced dimuon events. Technical Report CERN-EP/98-128, CERN, 1999. Submitted for publication in *The European Physical Journal C*.
- I. Žlimen, A. Ljubičić, S. Kaučić, and B.A. Logan. Evidence for a 17-keV neutrino. *Physical Review Letters*, 67(5):560–563, 1991.
- W.Allen Wallis. Compounding probabilities from independent significance tests. *Econometrica*, 10:229–248, 1942.
- Steven Weinberg. A model of leptons. *Physical Review Letters*, 19:1264–1266, 1967.
- Steven Weinberg. Physical processes in a convergent theory of the weak and electromagnetic interactions. *Physical Review Letters*, 27(24):1688–1691, 1971.
- T. Weiße. *Comparison of Muonic and Electronic Final States in the NOMAD Experiment at the CERN Wide-band Neutrino Beam*. PhD thesis, Universität Dortmund, 1997.
- D.A. White and A.C. Fabian. *Einstein Observatory* evidence for the widespread baryon overdensity in clusters of galaxies. *Mon. Not. R. Astron. Soc.*, 273:72–84, 1995.
- L. Wolfenstein. Neutrino oscillations in matter. *Physical Review D*, 17:2369–2374, 1978.
- C.S. Wu et al. Experimental test of parity conservation in beta decay. *Physical Review*, 105:1413–1415, 1957.
- C.N. Yang and R.L. Mills. Conservation of isotopic spin and isotopic gauge invariance. *Physical Review*, 96(1):191–195, 1954.
- H. Yukawa. On the interaction of elementary particles. *Proc. Phys. Math, Soc. Japan*, 17:48, 1935.
- G. Zweig. An SU(3) model for strong interaction symmetry and its breaking I. Technical Report CERN-8182-TH-401, CERN, 1964a.
- G. Zweig. An SU(3) model for strong interaction symmetry and its breaking II. Technical Report CERN-8419-TH-412, CERN, 1964b.

Daniel Zwillinger, editor. *CRC standard mathematical tables and formulae*.
CRC Press, Boca Raton, Florida, 30th edition, 1996.