

Calculating the Fluxes of Atmospheric Neutrinos

Laurence Cook

Lincoln College

University of Oxford

*A thesis submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy*

Trinity Term 2022

Abstract

The study of neutrinos produced in the interaction of cosmic rays with atmospheric nuclei was instrumental to the discovery of neutrino oscillations. As neutrino physics now enters the precision measurement era, precise measurement of atmospheric neutrino oscillations remains of interest due to its ability to answer questions such as the extent of CP violation in the neutrino sector and the ordering of the neutrino masses. An analysis of neutrino oscillations involves a comparison between the neutrino flux measured by a detector such as the Super-Kamiokande detector in Japan and the neutrino flux expected in the absence of oscillations. This thesis presents an estimate for the un-oscillated atmospheric neutrino flux using the Bartol atmospheric Monte Carlo simulation. The previous iteration of the Bartol flux estimate released in 2004 found that the dominant sources of systematic uncertainty on the neutrino flux were the uncertainty of the flux of cosmic rays and uncertainties in hadron production during the interaction of those cosmic rays with the atmosphere. Within this thesis, measurements of hadron production by fixed target experiments are used to develop a dataset-driven model which is applied to the simulation of pion and kaon production in atmospheric showers. The significant fraction of the relevant hadron production phase space covered by the datasets released since 2004 allows, for the first time, a calculation of the neutrino flux with hadronic interactions predominantly based on dataset measurements rather than Monte Carlo prediction. Additionally, an updated parameterisation of the cosmic ray fluxes is included within the neutrino flux calculation with parameters obtained in a global fit to recent cosmic ray flux measurements. The uncertainty of the atmospheric neutrino flux has been evaluated when using the updated hadronic interaction and cosmic ray flux models. The ν_μ flux uncertainty has been found to decrease from $\sim 15\%$ estimated on the 2004 fluxes to $\sim 5\text{--}10\%$ in this updated flux estimate.

Calculating the Fluxes of Atmospheric Neutrinos



Laurence Cook

Lincoln College

University of Oxford

A thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy

Trinity Term 2022

*The heavens declare the glory of God,
and the sky above proclaims his
handiwork.*

Psalm 19:1

Acknowledgements

First of all, I would like to thank my supervisors Giles Barr and Mark Hartz for their support, direction, advice and encouragement over the past few years without which this work would not have been possible. I particularly thank Giles for passing on his enthusiasm for atmospheric neutrinos and Mark for sharing his extensive knowledge of flux modelling, neutrino oscillations and detectors during our discussions. I have greatly valued our conversations over the past few years which have made this DPhil so intellectually stimulating and enjoyable.

I would also like to thank the group involved in the original Bartol Monte Carlo development, in particular Tom Gaisser for his helpful input at various stages in this work. I have greatly valued the input from the wider atmospheric neutrino and flux modelling community at meetings and workshops and I wish to thank the ATMPD group at Super-K for their input and the neutrino group at Nagoya University for their discussions, support and organisation of the semi-annual workshops on neutrino flux modelling which have been very valuable.

I have been fortunate to spend time at two physics institutes outside Oxford during my DPhil. I would like to thank the neutrino group at TRIUMF for hosting me over a very enjoyable and stimulating summer in 2019 and particular thanks to Mark for arranging the funding for this. I would also like to thank IPMU at the University of Tokyo for their funding of me through a joint fellowship with the University of Oxford and for hosting me from 2018 to 2021 when much of the work in this thesis was completed. Thank you to the administration at Oxford, IPMU and TRIUMF for all their assistance over the past few years.

My graduate life has been made immeasurably richer by the many people who have walked alongside me these past few years. I am especially grateful for the wonderful friendships with those around the world whom I would not have met were it not for this DPhil. Thank you to Alex for his companionship and encouragement throughout our graduate life, and Aniq for his wonderful friendship and for putting up with me as a flatmate in our first year in Oxford. Thank you to my whole community of friends at St. Aldates and the CiA group in Oxford for making my time in Oxford so memorable and challenging me in our discussions on the wider purpose of academic life. During my time in Vancouver, I was very thankful for the community at Tenth Church in particular for the hospitality of Ed during my visit

across the border to Washington and for introducing me to oysters. The memories of my 2-3 years in Tokyo will stay with me forever and I am very grateful for the support and friendship from everyone at Cornerstone Church in Tokyo and for making Tokyo a home. Particular thanks to Derick, Takumi, Amos, Elizabeth, Mike and Patti for all their support and love over the past few years. My final year and return to Oxford would not have been the same without the friendship of so many at Lincoln College in Oxford. I will remember with particular fondness the long breakfasts before a day of thesis writing!

I wish to thank my family, Mum, Emma, Amanda, Roland and my Grandparents for their support love and encouragement throughout this process. Finally to Hitomi, thank you for everything.

Abstract

The study of neutrinos produced in the interaction of cosmic rays with atmospheric nuclei was instrumental to the discovery of neutrino oscillations. As neutrino physics now enters the precision measurement era, precise measurement of atmospheric neutrino oscillations remains of interest due to its ability to answer questions such as the extent of CP violation in the neutrino sector and the ordering of the neutrino masses. An analysis of neutrino oscillations involves a comparison between the neutrino flux measured by a detector such as the Super-Kamiokande detector in Japan and the neutrino flux expected in the absence of oscillations. This thesis presents an estimate for the un-oscillated atmospheric neutrino flux using the Bartol atmospheric Monte Carlo simulation. The previous iteration of the Bartol flux estimate released in 2004 found that the dominant sources of systematic uncertainty on the neutrino flux were the uncertainty of the flux of cosmic rays and uncertainties in hadron production during the interaction of those cosmic rays with the atmosphere. Within this thesis, measurements of hadron production by fixed target experiments are used to develop a dataset-driven model which is applied to the simulation of pion and kaon production in atmospheric showers. The significant fraction of the relevant hadron production phase space covered by the datasets released since 2004 allows, for the first time, a calculation of the neutrino flux with hadronic interactions predominantly based on dataset measurements rather than Monte Carlo prediction. Additionally, an updated parameterisation of the cosmic ray fluxes is included within the neutrino flux calculation with parameters obtained in a global fit to recent cosmic ray flux measurements. The uncertainty of the atmospheric neutrino flux has been evaluated when using the updated hadronic interaction and cosmic ray flux models. The ν_μ flux uncertainty has been found to decrease from $\sim 15\%$ estimated on the 2004 fluxes to $\sim 5\text{-}10\%$ in this updated flux estimate.

Preface

This thesis and the work shown in it are my own. Figures and results included from other publications have been clearly attributed. This thesis has not been submitted either in whole or part for another qualification at this or any other university.

The focus of this thesis is the calculation of the atmospheric neutrino flux. My work updates and builds upon the 2004 Bartol atmospheric neutrino Monte Carlo by tuning hadron production and the primary cosmic rays within the simulation. This Monte Carlo was developed by Tom Gaisser, Todor Stanev and Giles Barr among others. The software packages which take the output of the Monte Carlo, calculate a tuned neutrino flux prediction and estimate the systematic uncertainties were all written by myself. In addition to the work described within this thesis, I made substantial changes to the Monte Carlo such that its output could be read using modern software packages.

Chapter 1 describes the motivation for the calculation of the flux of atmospheric neutrinos. In it I summarise the current theoretical and experimental understanding of the neutrino with relevant papers clearly referenced.

Chapter 2 discusses some of the features that exist within the atmospheric neutrino flux and their implication for this analysis. These features have already been discussed in previous publications but the plots in Chapter 2 which highlight these features have been produced by myself using the Bartol Monte Carlo and my analysis software.

Chapter 3 discusses the methodology of the Bartol Monte Carlo from 2004. In addition, I discuss a new method I developed within this work for simulating the neutrino flux by adapting the existing Monte Carlo.

In Chapter 4 I consider the relevant hadronic interactions which need to be understood to calculate the neutrino flux. To do this I use the Bartol Monte Carlo within my analysis framework to breakdown the contribution of different phase spaces of interactions to the flux. I then discuss the existing fixed-target experiments whose published data I will use to model these interactions.

In Chapters 5 and 6 I analyse the datasets published by these fixed-target experiments to develop a model for pion and kaon production in atmospheric showers. I adapt an existing parameterisation which has previously been used by beam neutrino flux calculations and fit it to the datasets. Issues of consistency between

the datasets are discussed. The software which analyses and fits parameterisations to the datasets was written by myself.

In Chapter 7 I use these pion and kaon production models to tune the neutrino flux prediction using the software framework developed by myself discussed earlier. In Chapter 8 I consider the prediction for hadron production outside the dataset phase space within Geant4 and Fluka hadron production generators, which were developed by their respective groups.

Chapter 9 discusses the uncertainties on the models developed in previous chapters and propagates them through to the neutrino flux estimate.

Chapter 10 introduces a previously published global fit to cosmic ray data by Evans et al. which I apply to the neutrino flux calculation. I propagate the uncertainties of the cosmic ray flux provided by Evans et al. through to the neutrino flux.

Finally, Chapter 11 provides the results for the tuning of the flux and uncertainty estimate along with my interpretation. In Chapter 12 I conclude by considering the outlook for this work and potential improvements to the precision of the neutrino flux which future hadron production experiments may provide.

Contents

List of Figures	xiii
1 Neutrinos and Neutrino Oscillation	1
1.1 Neutrino Physics	1
1.2 The Super-Kamiokande Experiment	9
1.3 Current Atmospheric Neutrino Research Questions	11
2 The Flux of Atmospheric Neutrinos	17
2.1 Production of Atmospheric Neutrinos	17
2.2 The Influence of Magnetic Fields	21
2.3 Symmetries and Asymmetries in the Un-Oscillated Neutrino Flux .	23
3 The Bartol Atmospheric Neutrino Monte Carlo	31
3.1 Overview of Bartol Monte Carlo	31
3.2 1D and 3D Calculations and the 3D Geometrical Effect	32
3.3 Injection of Primary Cosmic Rays	34
3.4 Hadronic Interactions with TARGET	35
3.5 Extended Detector and Decay Rotation Method	38
3.6 Obtaining the Neutrino Flux from Monte Carlo Neutrino Yields . .	45
4 Hadron Production Datasets for Atmospheric Neutrino Flux Pre-	47
dictions	
4.1 Hadron Interaction Phenomenology	47
4.2 Hadronic Interactions in Atmospheric Showers	49
4.3 Fixed Target Hadron Production Experiments	55
4.4 Summary	71
5 Modelling Pion Production in Air Showers	73
5.1 Using Fixed Target Data to Model Atmospheric Showers	73
5.2 Target Nucleus Scaling	75
5.3 Dataset Parameterisations	79
5.4 Comparing and Combining Datasets	86
5.5 Lower Energy Datasets	94

5.6	Data Phase Space Coverage of the Atmospheric Flux	107
5.7	Summary	111
6	Modelling Kaon Production in Air Showers	113
6.1	Kaon Component of Atmospheric Showers	113
6.2	Experimental Datasets	115
6.3	Neutral Kaon Cross Sections	116
6.4	Applying Kaon Data to Atmospheric Showers	118
6.5	Summary	120
7	Tuning the Atmospheric Neutrino Flux to Datasets	123
7.1	Applying the Hadronic Models to the Atmospheric Monte Carlo . .	123
8	Tuning Hadron Production Outside the Dataset Phase Space	133
8.1	Meson Production Outside Dataset Phase Space Coverage	134
8.2	Hadron Interaction Generators	134
8.3	Pion Production Comparisons Between Generators and Data	136
8.4	Kaon Production Comparisons Between Generators and Data . . .	141
8.5	Conclusions and the Weighting of the Interactions	142
9	Atmospheric Neutrino Flux Uncertainties due to Hadron Production	149
9.1	Uncertainties in Hadron Production Datasets	150
9.2	Interactions Occurring outside Dataset Phase Space	167
9.3	Energy Scaling Uncertainty	169
9.4	Additional Hadron Production Uncertainties	172
9.5	Summary	187
10	Primary Cosmic Rays	189
10.1	The Cosmic Ray Spectrum	189
10.2	Global Fits to Cosmic Ray Data	194
11	Results	201
11.1	Integrated Neutrino Fluxes	201
11.2	Neutrino Flux Flavour Ratios	204
11.3	Neutrino Flux Zenith and Azimuthal Directional Dependence	206
11.4	Different Detector Locations	210
12	Summary, Conclusions and Future Prospects	213
	References	217

List of Figures

1.1	Observation of neutrino oscillations with the Super-Kamiokande Experiment	6
1.2	Atmospheric neutrino events recorded by Super-Kamiokande shown as a function of zenith angle.	7
1.3	Elastic neutrino scattering in matter	8
1.4	Total modelled charged current neutrino interaction cross sections shown for ν_μ (left) and $\bar{\nu}_\mu$ (right).	10
1.5	Probability of atmospheric ν_μ to ν_e oscillation and atmospheric $\bar{\nu}_\mu$ to $\bar{\nu}_e$ oscillation.	13
1.6	Changes in ν_μ to ν_e oscillation probability caused by changing the CP violating phase, δ	14
2.1	The Bartol un-oscillated neutrino flux prediction at Super-Kamiokande Modelled at a solar minimum	18
2.2	Fraction of neutrino flux coming from pion and kaon decay modes.	20
2.3	Symmetry between the up-going and down-going atmospheric neutrino flux.	24
2.4	ν_μ fluxes as a function of cosine zenith angle for 3 neutrino energy bins	25
2.5	Bending of the primary cosmic rays with respect to the detector's local east and west as the origin of the east-west asymmetry in the atmospheric neutrino flux.	27
2.6	ν_μ fluxes as a function of azimuthal angle for 2 neutrino energy bins.	28
2.7	Ratio of the flux of different flavours of neutrinos.	29
3.1	The geometry of an isotropic source of neutrinos on a spherical shell as viewed by a detector.	33
3.2	The uncertainties on pion and kaon production as used by the Bartol calculation in 2006.	37
3.3	Absolute muon neutrino flux uncertainty estimate from the Bartol calculation in 2006.	38

3.4	A primary cosmic ray is generated at a small angle to the detector direction. The resultant down-going neutrino must be produced at a much larger angle to the primary cosmic ray in order to be able to reach the detector due to being produced much closer to the detector than the original primary cosmic ray.	42
3.5	A comparison of the ν_μ flux obtained via the decay rotation method and the extended detector method.	44
3.6	The graphs show the distribution of the neutrino weights due to the decay rotation probability at 1 GeV neutrino energy as a function of neutrino cosine zenith angle. They are made using a sample run of the decay rotation Monte Carlo with cosmic rays generated at a fixed energy of 100 GeV.	45
4.1	Breakdown of hadronic interaction component to the atmospheric ν_μ and ν_e flux.	50
4.2	For three different ν_μ energies (from left to right 0.2 GeV, 1.1 GeV, 7.1 GeV) contribution to the atmospheric neutrino flux by the primary energy of the hadron interaction.	52
4.3	Contribution of π^+ creation reactions of the form $pA \rightarrow \pi^+ + X$ to atmospheric muon neutrino flux.	53
4.4	Contribution of hadronic interactions to the atmospheric ν_μ flux at different neutrino zenith angles.	54
4.5	Figure showing the different beam momenta of the datasets used in this analysis.	56
4.6	Diagram showing the NA61/SHINE spectrometer alongside the detector setup in the beamline. This setup was used to take the data at 31 GeV beam momentum used to tune the atmospheric neutrino flux.	60
4.7	Diagram showing the NA49 spectrometer and the detector setup in the beamline. This setup was used to take the data at 158 GeV beam momentum.	62
4.8	Diagram showing the HARP spectrometer used to measure production cross sections from 3- 12 GeV on a variety of targets.	64
4.9	The E910 Spectrometer	66
4.10	The E802 Spectrometer	67
4.11	Single arm magnetic spectrometer used for production spectra measurements by Allaby et al. and Eichten et al.	68
4.12	The NA56 Experimental Setup	70
5.1	A Scaling Parameterisation with the ratio between pion production cross sections on aluminium and beryllium.	77

5.2	Deviation of A Scaling Parameterisation for carbon to nitrogen from simple $A^{\frac{2}{3}}$ scaling.	78
5.3	Lines of constant polar angle θ and outgoing pion momentum p in x_F - P_T phase space at 8,31 and 158 GeV beam momenta.	81
5.4	Extended BMPT parametrised fit to NA61 π^+ production data. The invariant differential cross section (defined in equation 5.4) is plotted against pion momentum in pion production angular bins	87
5.5	Extended BMPT parametrised fit to NA61 π^- production data. The invariant differential cross section (defined in equation 5.4) is plotted against pion momentum in pion production angular bins	88
5.6	Extended BMPT parametrised fit to NA49 π^+ production data. The invariant differential cross section (defined in equation 5.4) is plotted against x_F in transverse momentum (P_T) bins	89
5.7	Extended BMPT parametrised fit to NA49 π^- production data. The invariant differential cross section (defined in equation 5.4) is plotted against x_F in transverse momentum (P_T) bins	90
5.8	Energy dependence on $\pi^\pm$ inclusive production cross section.	92
5.9	Ratio of the NA61 pion production cross section at 31 GeV to the NA49 158 GeV pion production cross section.	93
5.10	Ratio of the NA49 pion production cross section at 158 GeV to the NA56 450 GeV pion production cross section.	95
5.11	NA49 and NA56 data where the bin centres overlap at $P_T = 0.15$ GeV and 0.5 GeV alongside fits to the NA49 data.	96
5.12	Datasets and joint fit in θ bins used to model interactions at $\sqrt{s} \approx 3.54$ GeV	100
5.13	Datasets and joint fit in θ bins used to model interactions at $\sqrt{s} \approx 5.17$ GeV	101
5.14	Datasets and joint fit in θ bins used to model interactions at $\sqrt{s} \approx 6.38$ GeV	101
5.15	Ratio between the joint BMPT fit to each of the datasets in the same energy region and the individual BMPT fit to each dataset for π^+ production.	102
5.16	Ratio between the joint BMPT fit to each of the datasets in the same energy region and the individual BMPT fit to each dataset for π^- production.	103
5.17	S error scaling factor as defined in the main text.	106
5.18	Fractional contribution of pion production interactions within each of the 7 energy regions defined in 5.6 to the ν_μ (left) and ν_e (right) atmospheric flux.	107

5.19	Contribution of different phase space components of pion creation reactions $pN \rightarrow \pi^+ + X$ to the atmospheric muon neutrino flux at 0.18 GeV	109
5.20	Contribution of different phase space components of pion creation reactions $pN \rightarrow \pi^+ + X$ to the atmospheric muon neutrino flux at 1.13 GeV	110
5.21	Contribution of different phase space components of pion creation reactions $pN \rightarrow \pi^+ + X$ to the atmospheric muon neutrino flux at 7.12 GeV	111
5.22	Extent to which the phase space of pion production interactions in atmospheric showers are covered by the experimental datasets used in this analysis.	112
6.1	Contribution to each flavour of atmospheric neutrino flux by decays of charged and neutral kaons.	114
6.2	Ratio of kaon production at 31 to 158 GeV using fits to NA61 and NA49 data.	119
6.3	Breakdown of the phase space coverage of the kaon datasets to the atmospheric neutrino flux.	120
7.1	Cartoon of an interaction of a cosmic ray on an atmospheric nuclei resulting in neutrino production.	125
7.2	Weights for π^+ production as the ratio between multiplicities from the dataset models to the multiplicities used by the TARGET2.4 generator.	128
7.3	Weights for π^- production as the ratio between multiplicities from the dataset models to the multiplicities used by the TARGET2.4 generator.	129
7.4	Weights for K^+ , K^- and K_L^0 production in interactions by incident protons on air as the ratio between multiplicities from the dataset models to the multiplicities used by the TARGET2.4 generator. . .	130
7.5	Results of weighting the neutrino fluxes using weights of Section 7.1.3 weighting only inside the region of dataset coverage	131
8.1	Underlying models used by different hadron production generators at different energies to simulate pion and kaon production in nucleon nucleus interactions.	135
8.2	Comparisons between Geant4 generators (solid), TARGET2.4 (dashed) and data points for pion production at low angles.	137

8.3	Comparisons between Geant4 generators (solid), TARGET2.4 (dashed) and data points for the lowest P_T Bin of 158 GeV NA49 pion production data.	139
8.4	Comparisons between Geant4 generators (solid), TARGET2.4 (dashed) and data points for pion production between 200 and 550 mrad at 5 GeV beam momentum of the HARP Data.	141
8.5	Comparisons between Geant4 generators (solid), TARGET2.4 (dashed) and data points for pion production between 200 and 550 mrad at 12 GeV beam momentum of the HARP Data.	142
8.6	Comparisons between Geant4 generators (solid), TARGET2.4 (dashed) and data points for pion production at large angles.	143
8.7	Comparisons between Geant4 generators (solid), TARGET2.4 and the NA61 data for kaon production at 31 GeV.	144
8.8	Weights for π^+ production outside the phase space coverage of the datasets.	145
8.9	Weights for π^- production outside the phase space coverage of the datasets.	146
8.10	Weights for K^+ , K^- and K_L^0 production by incident protons outside phase space coverage of the datasets.	147
8.11	The neutrino fluxes with pion and kaon production interactions weighted both inside and outside dataset coverage.	148
9.1	Example of perturbations to the NA49 π^+ production data for different values of the correlation parameter in equation 9.1	152
9.2	Uncertainty on the ν_μ flux due to the systematic uncertainty on the NA49 π^+ production data assuming different correlation models. . .	153
9.3	Uncertainty on the $\nu_e/\bar{\nu}_e$ flux ratio propagated from the bin-by-bin systematic uncertainties on the NA49 dataset assuming different correlation between π^+ and π^- production uncertainties.	154
9.4	Uncertainty on the ν_μ and $\nu_e/\bar{\nu}_e$ flux ratio due to the uncertainty on the NA61 and NA49 datasets.	155
9.5	Uncertainty on the ν_μ neutrino flux due to statistical uncertainties on the NA61 and NA49 datasets.	156
9.6	Uncertainty on the ν_μ flux due to the uncertainty on the <30 GeV datasets.	158
9.7	HARP OH and HARP CDP pion production cross section at 5 GeV beam momentum	159
9.8	Ratio of HARP OH to HARP CDP pion production cross section fits in x_F - P_T plane	160
9.9	Neutrino flux weighted to the HARP OH and HARP CDP analyses.	161

9.10	Histogram showing the fractional discrepancy between the HARP OH and HARP CDP analysis in the x_F - P_T plane weighted by the contribution of that region of x_F - P_T to the neutrino flux.	163
9.11	Uncertainty on the neutrino flux due to HARP high angle dataset .	164
9.12	Tuning of the neutrino fluxes with and without the target nucleus interpolation used to scale the NA61 and NA49 pion multiplicities from carbon to air.	165
9.13	Uncertainty on the ν_μ (left) and ν_e (right) flux due to the uncertainty on kaon production within experimental datasets	166
9.14	Contribution to the ν_μ flux from pion and kaon production interactions occurring outside the phase space of the experimental datasets.	168
9.15	Uncertainties on the ν_μ flux due to pion and kaon production outside the x_F - P_T dataset phase space.	169
9.16	Inclusive production cross section as a function of interaction beam momentum at fixed points in the x_F - P_T plane	171
9.17	Uncertainties on the ν_μ flux due to interpolation between the 31 GeV and 158 GeV datasets and the Feynman Scaling process from 158 GeV to 450 GeV	172
9.18	Example of a high energy cosmic ray undergoing multiple re-interactions to produce a lower energy neutrino	173
9.19	The mean number of hadronic interactions as a function of neutrino energy (left) and the energy of a primary cosmic ray prior to the neutrino producing interaction as a fraction of the initial cosmic ray energy (right).	174
9.20	Fractional contribution of each of the baryon to baryon hadronic interactions to the ν_μ neutrino flux.	175
9.21	Comparison of the multiplicity of $pN \rightarrow p + X$ interactions with TARGET2.4 and the HARP experimental dataset at a beam momenta of 12 GeV.	176
9.22	Comparison of the multiplicity of $pN \rightarrow p + X$ interactions with TARGET2.4 and the NA61 experimental dataset at a beam momenta of 31 GeV.	177
9.23	Comparison of the multiplicity of $pN \rightarrow p + X$ interactions with TARGET2.4 and the NA49 experimental dataset at a beam momenta of 158 GeV.	178
9.24	Effect of the tuning of the leading baryons on the neutrino flux. . .	179
9.25	The ν_μ (left) and ν_e (right) atmospheric neutrino flux after altering the multiplicity of fragmentation baryons.	180
9.26	The Λ^0 production component to the ν_μ neutrino flux broken down by separate contributions from each of the decay products.	181

9.27	The fractional component of the ν_μ neutrino flux created in atmospheric showers involving meson re-interaction separated into components coming from pion and kaon interaction with the pion component further separated by the daughter particle created in the interaction.	183
9.28	Contribution of $\pi^+ N \rightarrow \pi^+ + X$ interactions to the ν_μ neutrino flux shown as a function of incoming pion momentum.	184
9.29	Comparison between the multiplicity of $\pi^+ N \rightarrow \pi^+ + X$ interactions at incident momentum 60 GeV within TARGET2.4 and NA61/SHINE experimental data for interactions on a carbon target.	185
9.30	Ratio of the ν_μ neutrino flux after the weighting of $\pi^\pm N \rightarrow \pi^\pm + X$ interactions to the flux without pion re-interaction weighting.	186
10.1	Fluxes of primary cosmic ray nuclei	191
10.2	Fractional contribution of each cosmic ray nuclear group to the total nucleon flux and the atmospheric neutrino flux	193
10.3	Fractional deviation of the data points to the global fit shown for primary proton and helium fluxes	195
10.4	Ratio of the global fits of cosmic ray data by Evans et al. to the ICRC01 primary flux previously used to calculate the neutrino flux.	196
10.5	Ratio of the neutrino flux calculated using the global fit by Evans et al. to the flux calculated using the ICRC01 primary flux	196
10.6	Uncertainty of primary proton and helium fluxes as given by the global fit from Evans et al.	197
10.7	Uncertainty on the neutrino flux due to primary cosmic rays using the flux model by Evans et al.	198
10.8	Uncertainty on the neutrino flux due to primary cosmic rays estimated using the rescaling technique by Evans et al.	200
11.1	The atmospheric neutrino fluxes averaged over all incoming solid angles after the tuning of the pion and kaon production and application of the new primary cosmic ray model described within this thesis.	202
11.2	The uncertainty on the neutrino flux shown for each neutrino flavour.	203
11.3	The neutrino flux flavour ratios calculated using the model described within this thesis (labelled tuned flux) and the original Bartol 2004 model (labelled original flux)	204
11.4	The neutrino flux flavour ratio uncertainties.	205
11.5	The ν_μ neutrino fluxes shown as a function of neutrino cosine zenith angle at the detector binned in four neutrino energy bins.	207

11.6	Left: The total ν_μ neutrino flux uncertainty as a function of neutrino zenith angle for each of the four neutrino energy bins shown in Figure 11.5. The total has been calculated as a quadrature sum of each of the components listed in Figure 11.2. Right: A breakdown of the pion dataset uncertainty contribution to the first ν_μ energy bin (0.1-0.32 GeV).	208
11.7	The ν_μ neutrino fluxes shown as a function of neutrino azimuthal angle at the detector measured in degrees and binned in four neutrino energy bins.	209
11.8	The total uncertainty prediction on the ν_μ flux shown in Figure 11.7 as a function of neutrino zenith angle	209
11.9	The ν_μ and ν_e solid angle averaged atmospheric neutrino fluxes calculated using the hadron production modelling described within this thesis shown for three different detector sites: Kamioka, DUNE and IceCube.	211

1

Neutrinos and Neutrino Oscillation

1.1 Neutrino Physics

In 1931 Pauli proposed the existence of an as yet unobserved light neutral particle to explain the apparent lack of energy and angular momentum conservation in beta decays of atomic nuclei [1]. This particle was incorporated into Fermi's theory of beta decay (the first theory of weak interactions) as a neutral, massless, spin- $\frac{1}{2}$ fermion which he named a neutrino and which was ejected as a daughter product in beta decay such that the total energy, momentum and angular momentum were conserved [2]. The neutrino's interaction cross section, estimated by Bethe and Peierls, was so small ($\sim 10^{-44} \text{ cm}^2$) that they concluded there was no practical method to detect it [3]. Pauli predicted that he had proposed a particle which couldn't be observed experimentally and even bet a case of champagne that the neutrino wouldn't be detected. Indeed it wasn't until 1953 when Reines and Cowan were able to take advantage of the large neutrino flux from nuclear reactors and report the first evidence for the detection of the neutrino via inverse beta decay, 22 years after it was first proposed [4]. Although the neutrino is the second most abundant particle in the universe (after the photon) its small interaction cross

section with matter means it remains one of the least understood. Following the first detection of the neutrino in 1953 (now understood to be an electron antineutrino) the muon neutrino was discovered in 1962 as a neutrino which could interact to produce muons but not electrons [5]. A third flavour of neutrino was predicted following the discovery of the tau lepton [6] and subsequently the tau neutrino was discovered by the DONUT experiment in 2000 [7].

1.1.1 Neutrino Oscillations

In the 1960s the Homestake experiment [8] measured the flux of electron neutrinos coming from ${}^8\text{B}$ decay in the sun. Surprisingly the flux of solar neutrinos was measured to be around a third of the flux predicted by solar models [9]. At the time it was assumed that there must be an error in either the experimental measurement or the theoretical models which became known as the solar neutrino problem.

Measurements of the flux of neutrinos originating in the interaction of high-energy cosmic rays in the atmosphere encountered similar issues with a deficit in the flux of muon neutrinos measured by the Kamiokande [10] and IMB [11] experiments despite an agreement with the flux of electron neutrinos when compared to predictions. The prediction and modelling of the atmospheric neutrino flux is the focus of this thesis and will be described in detail within subsequent chapters.

The cause of both of these deficits is now understood to be neutrino oscillations, changes in the flavour of neutrinos as they propagate through space. In 1957 Pontecorvo had already suggested the possibility of oscillations in the lepton sector analogous to the neutral kaon oscillations in the hadron sector [12]. For neutrinos to oscillate they must be altered from Fermi's massless leptons and instead assigned a small non-zero mass.

Within the weak interaction, neutrinos interact as the flavour eigenstates ν_e , ν_μ and ν_τ . When they propagate through space they propagate as eigenstates of the free particle Hamiltonian, referred to as the mass eigenstates ν_1 , ν_2 and ν_3 . Under the hypothesis of neutrino oscillations the flavour and mass eigenstate bases are not

the same with the rotation between them described by the unitary matrix U [13]:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (1.1)$$

Thus the wave function of a neutrino of flavour α created at time $t = 0$ can be expressed as:

$$|\Psi(0)\rangle \equiv |\nu_\alpha\rangle = \sum_j U_{\alpha j}^* |\nu_j\rangle, \quad (1.2)$$

where $j = (1, 2, 3)$ indicates the mass eigenstates. The mass eigenstates each evolve according to the Schrödinger equation with the free particle Hamiltonian giving the evolution of the wave function as:

$$|\Psi(t)\rangle = \sum_j e^{-i\phi_j} U_{\alpha j}^* |\nu_j\rangle, \quad (1.3)$$

with:

$$\phi_j = E_j t - \mathbf{p}_j \cdot \mathbf{x} \quad (1.4)$$

with E as the energy of the mass eigenstate and $\mathbf{p}$ its momentum. This can be re-expressed in terms of the flavour states as:

$$|\Psi(t)\rangle = \sum_j \sum_k e^{-i\phi_j} U_{\alpha j}^* U_{kj} |\nu_k\rangle, \quad (1.5)$$

where $k = (e, \mu, \tau)$ indicates the flavour eigenstates. Hence after time t the probability of detecting a neutrino that was produced in flavour state α as flavour state β is given by:

$$P(\alpha \rightarrow \beta) = |\langle \nu_\beta | \Psi(t) \rangle|^2 = \left| \sum_j e^{-i\phi_j} U_{\alpha j}^* U_{\beta j} \right|^2 \quad (1.6)$$

To simplify this, the difference between the phases of the two mass eigenstates can be approximated as as:

$$\phi_j - \phi_i = (E_j - E_i)t - (\mathbf{p}_j - \mathbf{p}_i) \cdot \mathbf{x} \approx \frac{m_j^2 - m_i^2}{2E} t \approx \frac{\Delta m_{ij}^2 L}{2E}, \quad (1.7)$$

where Δm_{ij}^2 represents the difference between the square of the masses. Since neutrinos move at approximately the speed of light, time t has been replaced by distance L , the distance travelled by the neutrino in the detection frame.

The flavour probability can then be expressed as [14]:

$$P(\alpha \rightarrow \beta) = \delta_{ab} - 4 \sum_{k < j}^3 \text{Re}[U_{\alpha k} U_{\beta k}^* U_{\alpha j} U_{\beta j}^*] \sin^2\left(\frac{\Delta m_{kj}^2 L}{4E}\right) + 2 \sum_{k < j}^3 \text{Im}[U_{\alpha k} U_{\beta k}^* U_{\alpha j} U_{\beta j}^*] \sin\left(2 \frac{\Delta m_{kj}^2 L}{4E}\right) \quad (1.8)$$

Here we can see the requirement for neutrinos to be massive in order to exhibit oscillatory behaviour. In the limit of all the masses being zero, Δm_{kj}^2 is 0 and (using the unitarity of the matrix U) the oscillation probability reduces to the delta function and there is no possibility of flavour change. For oscillation of antineutrinos, each matrix element U_{ij} is replaced with its complex conjugate U_{ij}^* . Under this transformation the first two terms of 1.8 do not change but the final term changes sign. This final term describes differences between neutrino and antineutrino oscillation, this is known as CP violation and will be discussed later.

The mixing matrix U is known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [15, 16] and is often parameterised in terms of 3 separate two-dimensional rotations [14]:

$$U \equiv \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_{23} & \sin \theta_{23} \\ 0 & -\sin \theta_{23} & \cos \theta_{23} \end{pmatrix} \begin{pmatrix} \cos \theta_{13} & 0 & \sin \theta_{13} e^{-i\delta} \\ 0 & 1 & 0 \\ -\sin \theta_{13} e^{i\delta} & 0 & \cos \theta_{13} \end{pmatrix} \begin{pmatrix} \cos \theta_{12} & \sin \theta_{12} & 0 \\ -\sin \theta_{12} & \cos \theta_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (1.9)$$

In this parameterisation, the complex phase δ describes CP violation and much of modern experimental neutrino physics is concerned with the measurement of these oscillation parameters. This will be discussed further in Section 1.3.

1.1.2 Two Flavour Oscillations

The typical values of L/E for a specific neutrino source will determine which oscillations will dominate. If, in equation 1.8, $L/E \ll 1/\Delta m_{kj}^2$ the oscillation probability will be small and if $L/E \gg 1/\Delta m_{kj}^2$ the oscillations will be very fast over the energy resolution of the detector and $\sin^2(\frac{\Delta m_{kj}^2 L}{4E})$ averages to $\frac{1}{2}$. Often the more complex 3 flavour oscillation can be simplified to consider only the dominant oscillation between two neutrino flavours. Under this simplification the PMNS matrix can be re-written in terms of one parameter:

$$U = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}, \quad (1.10)$$

And the neutrino oscillation probability is simplified from equation 1.8 as [17]:

$$P(\alpha \rightarrow \beta) = \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L}{4E}\right) \quad (1.11)$$

Note that the oscillation probability is only dependent on the difference between the masses squared not the absolute mass meaning that analysis of neutrino oscillations cannot determine the scale of the neutrino mass only the difference between them. In addition to this, the two flavour oscillation probability is not dependent on the sign of Δm^2 meaning it cannot determine which mass eigenstate is heavier.

Detection of atmospheric neutrinos in 1998 with the Super-Kamiokande experiment (the successor to the Kamiokande experiment and discussed further in section 1.2) provided the first strong evidence of neutrino oscillations [18, 19]. Since atmospheric neutrinos are created at all points in the atmosphere, their flight distance L varies from ~ 10 km for neutrinos created directly over the detector to $\sim 10^4$ km for neutrinos created on the other side of the Earth enabling the oscillation hypothesis to be tested.

Figure 1.1, published in 1998, shows the ratio of neutrino events observed by Super-Kamiokande [19] to the predicted number in the absence of oscillations. A clear deficit can be seen for the ν_μ that have travelled the longest distance and is consistent with $\nu_\mu \leftrightarrow \nu_\tau$ two flavour oscillation (dashed line). Figure 1.1 was




Figure 1.1: Observation of neutrino oscillations with the Super-Kamiokande Experiment. The ratio of detected neutrinos by Super-Kamiokande against the prediction without oscillations is shown. The dashed line shows the best fit for $\nu_\mu \leftrightarrow \nu_\tau$ oscillations. Figure from [19].

published when Super-Kamiokande first showed evidence for neutrino oscillations [19]. The atmospheric neutrino data taken for that paper had a detector exposure of 535 days or 33.0 kt-years [20]. Figure 1.2 shows atmospheric neutrino data published in 2016 taken with a detector exposure of 4906 days and more than 300 kt-years [20]. The figure shows neutrino events as a function of neutrino cosine zenith angle measured by Super-Kamiokande. The neutrino events have been grouped into 15 labelled subgroups depending on energy and final state of the neutrino interaction within the detector, electron and muon neutrino like events are labelled. Similar to Figure 1.1 a deficit is seen for upward going muon neutrinos (corresponding to those which have travelled $\sim 10^4$ km) observed when compared to the un-oscillated Monte Carlo. The best fit line including oscillations can be seen to describe the data well.

In 2002, by measuring the flux of all three flavours of neutrinos using their




Figure 1.2: Atmospheric neutrino events recorded by Super-Kamiokande shown as a function of zenith angle in 15 separate subgroups. The blue line shows the prediction in the absence of oscillations and the red line shows the best fit oscillation prediction. Taken from [20]

neutral current interactions, the SNO experiment was able to confirm the deficit in solar electron neutrinos observed by the Homestake experiment was also caused by flavour-changing neutrino oscillations [21].

1.1.3 Neutrino Oscillations in Matter

Until this point, neutrino oscillations have been shown as a consequence of the mixing of the matter and flavour eigenstates where neutrino propagation is described by the free particle Hamiltonian. This is only true for neutrinos which propagate in a vacuum, within the presence of matter (such as for neutrinos traversing the Earth) neutrinos can scatter via the exchange of weak bosons. This scattering introduces an effective potential within the Hamiltonian for neutrinos traversing matter. Any flavour of neutrino can scatter via the exchange of the weak neutral Z




Figure 1.3: Elastic neutrino scattering in matter. Neutral current scattering is shown in the left-hand figure which affects all neutrino flavours equally while charged current scattering of the electron neutrino is shown in the right hand figure.

boson, as shown in Figure 1.3. Since this effect is equal for all neutrino flavours it will only introduce a common phase difference to the Hamiltonian and not affect oscillation probabilities. Electron neutrinos can interact with the electrons in matter via charged current interactions, see Figure 1.3. Since only electron neutrinos can interact with electrons this way it introduces an effective potential into the Hamiltonian which only affects electron neutrinos. This potential is given by: $V_e = \pm\sqrt{2}G_F N_e$ where G_F is the Fermi constant, and N_e the electron density number [17]. The potential is positive for ν_e and negative for $\bar{\nu}_e$. It is easiest to understand this effect in the two-flavour framework of oscillations between electron neutrinos and one other flavour. It is helpful to consider the Hamiltonian for the propagation of a free neutrino in the neutrino flavour basis. By transforming the diagonal free-particle Hamiltonian in the basis of mass eigenstates into the basis of flavour eigenstates via mixing matrix 1.10, the Hamiltonian for propagation in matter can be shown to be [22]:

$$H = \frac{\Delta m^2}{4p} \begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix} + 2\sqrt{2}G_F N_e \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad (1.12)$$

where the second term is the additional potential for electron neutrinos and the time-independent Schrödinger equation is given as:

$$H \begin{pmatrix} \nu_e \\ \nu_x \end{pmatrix} = E \begin{pmatrix} \nu_e \\ \nu_x \end{pmatrix} \quad (1.13)$$

This Hamiltonian can be rewritten as:

$$H = \frac{\Delta m_M^2}{4p} \begin{pmatrix} -\cos 2\theta_M & \sin 2\theta_M \\ \sin 2\theta_M & \cos 2\theta_M \end{pmatrix} \quad (1.14)$$

where:

$$\sin^2 2\theta_M \equiv \frac{\sin^2 2\theta}{\sin^2 2\theta + (\cos 2\theta - \Gamma)^2} \quad (1.15)$$

$$\Delta m_M^2 \equiv \Delta m^2 \sqrt{\sin^2(2\theta) + (\cos 2\theta - \Gamma)^2} \quad (1.16)$$

$$\Gamma \equiv \pm \frac{2\sqrt{2}G_F N_e E}{\Delta m^2} \quad (1.17)$$

Hence the original oscillation probabilities can be used but with a modification to the mixing angle and squared mass difference. It is to be noted that $\sin^2 2\theta_M$ is maximal (and thus the oscillation probability, see equation 1.11 is maximal) for $\Gamma \sim \cos 2\theta$. Crucially this depends on the sign of Δm^2 where if Δm^2 is positive this maxima will occur for neutrinos whereas if Δm^2 is negative it will occur for antineutrinos.

1.2 The Super-Kamiokande Experiment

Super-Kamiokande is a 50-kiloton water Čerenkov neutrino detector situated 1 km under Mt. Ikeno within the Kamioka mine in Japan [23]. It consists of a cylindrical inner detector containing 11,000 photomultiplier tubes (PMTs) facing inwards to detect light produced within the inner detector and a 2 m thick outer detector surrounding the inner detector. A neutrino that interacts via a charged current interaction inside the detector will produce a charged lepton which if travelling faster than the speed of light in water will emit Čerenkov radiation. This radiation is detected by the PMTs and observed as a ring on the wall of the detector, the original neutrino flavour, energy and direction can be reconstructed by analysis of




Figure 1.4: Total modelled charged current neutrino interaction cross sections shown for ν_μ (left) and $\bar{\nu}_\mu$ (right). The total cross section shown in black is split into components with different final states alongside the charged lepton: Charged current quasi-elastic interactions (red), charged current single pion production (blue) and deep inelastic scattering (purple). Data points from experimental measurement are shown. Figure taken from [24].

these rings. Crucially for the atmospheric neutrino analysis, the SK detector can determine the direction of the incoming neutrino and therefore the approximate distance it will have travelled from creation to detection. The cross section for charged current neutrino and antineutrino interactions is shown in Figure 1.4. Notice that the cross section decreases at energies <1 GeV, while it will be shown in Chapter 2 that the atmospheric neutrino flux decreases with increasing energy, the decrease in cross sections at lower neutrino energies means the event rate for atmospheric neutrinos is typically highest around ~ 1 GeV [18].

Electron and muon neutrinos can be distinguished within the detector by the features of the Čerenkov ring. Electrons and positrons scatter and produce electromagnetic showers, resulting in Čerenkov rings which appear blurred in comparison to the sharper muon rings [25]. Detection of tau neutrinos via charged current interactions requires the production of the tau lepton which is significantly more massive than the electron and muon. The tau neutrino charged current interaction is therefore suppressed in comparison to that of the electron and muon neutrino. The kinematic limit imposed by the large tau lepton mass also means it is not possible to detect ν_τ below 3.5 GeV and the rapid decay of the tau lepton into muons and electrons make identifying tau neutrinos above this limit challenging [26].

It is difficult to distinguish between neutrinos and antineutrinos with the Super-Kamiokande detector which does not have a magnetic field and therefore cannot easily distinguish the charge of the lepton produced in the charged current interaction. Some separation is possible by considering the differences between the interaction of different daughter particles created in neutrino and antineutrino charged current interactions within the detector [27].

The Super-Kamiokande detector is also used as the far detector for the T2K experiment, which measures the oscillation of a beam of neutrinos created 295 km from the Super-Kamiokande detector [28]. To form a neutrino beam, a beam of protons is aimed at a carbon target resulting in the production of pions and kaons which later decay into neutrinos. Due to the similarities between neutrinos produced via this method and the production of atmospheric neutrinos, some of the techniques, methods and datasets used by the T2K experiment to calculate the flux of neutrinos will be used within the atmospheric flux calculation presented in this thesis.

1.3 Current Atmospheric Neutrino Research Questions

Historically Super-Kamiokande studied the dominant $\nu_\mu \leftrightarrow \nu_\tau$ atmospheric neutrino oscillations within the two flavour oscillation framework (as in Figure 1.1). Due to the large number of neutrino events detected since Super-Kamiokande started data taking in 1996 and the decrease in uncertainties, Super-Kamiokande is now sensitive to the full three flavour oscillation [20]. In particular atmospheric neutrino oscillations are sensitive to the following sub-dominant oscillation parameters:

- While the magnitude of Δm_{31}^2 is known from oscillations of ν_μ to ν_τ its sign is as yet unknown. As shown in equation 1.11 the vacuum oscillation probability is independent of the sign of Δm_{31}^2 and can only be determined through additional effects of propagation through matter. It has already been established that $m_2 > m_1$ from the measurement of the matter effects in solar oscillations however it is unknown whether the third matter eigenstate is

heavier or lighter than the other two. $m_3 > m_2 > m_1$ is referred to as normal ordering and $m_2 > m_1 > m_3$ as inverted ordering. In addition to reducing the degeneracy of other oscillation parameters, understanding the ordering of the neutrino mass eigenstates helps point toward different fundamental models explaining the neutrino masses [29].

- The phase δ in equation 1.9 is a measure of the CP violation in neutrino oscillations. The CP transformation changes a particle into its corresponding antiparticle and CP violation indicates a breaking of the symmetry between matter and antimatter. CP violation is one of the conditions proposed by Sakharov needed to explain the imbalance between matter and antimatter observed in the universe [30]. CP violation has already been observed in the weak interaction of quarks through the CKM matrix but not at a scale that could fully explain the matter-antimatter asymmetry observed in the universe [14]. A value of δ other than 0 or π would imply CP violation in neutrino oscillations. The value of the CP violating phase is not yet known precisely. Recent results from the T2K experiment have constrained the value of δ to be between $[-3.41, -0.03]$ if the normal mass ordering is normal and $[-2.54, -0.32]$ if the mass ordering is inverted [31].

Precise measurement of the oscillation of atmospheric neutrinos within the following channels is therefore of interest to future analyses. The leading terms of the expansion of equation 1.8 for the probability of oscillation in a vacuum are given [27, 32]:

$$P(\nu_e \rightarrow \nu_e) \approx 1 - \sin^2 2\theta_{13} \sin^2 \left(\frac{1.27 \Delta m_{31}^2 L}{E} \right) \quad (1.18)$$

$$P(\nu_\mu \rightarrow \nu_\mu) \approx 1 - 4 \cos^2 \theta_{13} \sin^2 \theta_{23} (1 - \cos^2 \theta_{13} \sin^2 \theta_{23}) \times \sin^2 \left(\frac{1.27 \Delta m_{31}^2 L}{E} \right) \quad (1.19)$$

$$P(\nu_\mu \leftrightarrow \nu_e) \approx \sin^2 \theta_{23} \sin^2 2\theta_{13} \sin^2 \left(\frac{1.27 \Delta m_{31}^2 L}{E} \right). \quad (1.20)$$




Figure 1.5: Probability of ν_μ to ν_e oscillation (left) and $\bar{\nu}_\mu$ to $\bar{\nu}_e$ oscillation (right) with atmospheric neutrinos as a function of neutrino energy and neutrino zenith angle. This figure assumes normal ordering. There is a resonance in oscillation probability between 2 and 10 GeV for neutrino oscillation. If the mass ordering is inverted the resonance instead occurs for antineutrino oscillations. The CP violating phase δ is taken as 0 for these plots. Neutrinos with a negative zenith correspond to those which travel through the Earth to the detector. Figure taken from [27]

The factor of 1.27 arises from changing from natural units to distance L in km, energy E in GeV and Δm_{31}^2 in eV^2 . For neutrino propagation through matter the mixing angle and squared mass differences are changed (as in equation 1.15) as follows [27]:

$$\Delta m_{31,M}^2 = \Delta m_{31}^2 \sqrt{\sin^2 2\theta_{13} + (\Gamma - \cos 2\theta_{13})^2} \quad (1.21)$$

$$\sin^2 2\theta_{13}^M = \frac{\sin^2 2\theta_{13}}{\sin^2 2\theta_{13} + (\Gamma - \cos 2\theta_{13})^2}, \quad (1.22)$$

where Γ follows the same definition as equation 1.15.

Figure 1.5 shows the probability of atmospheric ν_μ to ν_e oscillation and $\bar{\nu}_\mu$ to $\bar{\nu}_e$ oscillation as a function of cosine zenith angle of the neutrino at the detector (related to its path length) and energy. Under the hypothesis of normal mass ordering, there is a resonance in oscillation probability between 2-10 GeV for neutrinos travelling through the Earth to the detector which does not occur for antineutrino oscillation. This occurs where Γ is approximately $\cos 2\theta_{13}$ in equation 1.21. If the mass ordering were inverted, the sign of Γ is flipped and the resonance would be seen in the




Figure 1.6: Changes in atmospheric ν_μ to ν_e oscillation probability caused by changing the CP violating phase, δ . The difference between the oscillation probability when $\delta = 20^\circ$ and $\delta = 220^\circ$ is shown for $P(\nu_\mu \rightarrow \nu_\mu)$ (left) and $P(\nu_\mu \rightarrow \nu_e)$ (right). The oscillation probability is shown as a function of neutrino energy and cosine zenith angle. Figure taken from [33]

antineutrino oscillation and not neutrino oscillation.

At leading order the probabilities $P(\nu_\mu \rightarrow \nu_e)$ and $P(\nu_e \rightarrow \nu_\mu)$ are the same and given by equation 1.20. This is an example of time reversal or T symmetry. The CPT operator is the combination of the CP and T operator, asymmetry under the CPT operation (CPT violation) implies violation of Lorentz invariance for which there has been no experimental observation [34]. Under the assumption of CPT symmetry, if T is symmetric so also is CP. The next terms in $P(\nu_\mu \rightarrow \nu_e)$ expansion are given by [35]:

$$\begin{aligned}
 P(\nu_\mu \rightarrow \nu_e) \approx & \sin^2 \theta_{23} \sin^2 2\theta_{13} \sin^2 \left(\frac{1.27 \Delta m_{31}^2 L}{E} \right) \\
 & + \sin 2\theta_{23} \sin 2\theta_{13} \cos \theta_{13} \sin 2\theta_{12} \sin \left(\frac{1.27 \Delta m_{31}^2 L}{E} \right) \\
 & \times \sin \left(\frac{1.27 \Delta m_{21}^2 L}{E} \right) \cos \left(\frac{1.27 \Delta m_{32}^2 L}{E} + \delta \right), \quad (1.23)
 \end{aligned}$$

which includes the CP violating phase δ and therefore at this order in the expansion T symmetry does not hold. Instead $P(\nu_e \rightarrow \nu_\mu)$ is obtained by the action of the T operator on equation 1.23, equivalent to the CP operator and therefore $\delta \rightarrow -\delta$. Figure 1.6 shows the sensitivity of the oscillation probability of atmospheric ν_μ

to ν_μ and ν_μ to ν_e to the CP violating phase δ . Accounting for detector energy resolution and event rate, the sensitivity mainly comes from neutrinos <1 GeV [33].

The oscillated flux of atmospheric neutrinos as measured by a neutrino detector alone is not sufficient to determine the probability of oscillation and therefore to measure oscillation parameters. Alongside the detection of atmospheric neutrinos, a model for their production is needed, providing an estimate of the neutrino flux at the detector in the absence of oscillation effects. The modelling and estimate of the un-oscillated atmospheric neutrino flux is the focus of this thesis. As a consequence of figures 1.5 and 1.6, this thesis will focus on modelling the atmospheric neutrino flux between 0.1 and 10 GeV, the energy region most relevant to both CP violation and the neutrino mass ordering.

2

The Flux of Atmospheric Neutrinos

This chapter introduces neutrinos created in atmospheric showers, the modelling of which is the main focus of this thesis. This will start by considering their production and the effects which contribute to the features in their flux. To illustrate features in the atmospheric neutrino flux, fluxes modelled with the Bartol atmospheric Monte Carlo are presented. This Monte Carlo and the methodology behind calculating the neutrino flux are described in greater detail in Chapter 3.

2.1 Production of Atmospheric Neutrinos

Atmospheric neutrinos are produced in the decay of particles created due to the collision of cosmic rays with atmospheric nuclei. As such, the spectra of primary cosmic rays, their propagation through geomagnetic and heliomagnetic fields, their interaction producing daughter hadrons and the decay of those daughter hadrons and muons to neutrinos all affect the flux of atmospheric neutrinos. This chapter considers the underlying features and symmetries of the atmospheric neutrino flux caused by the flux of cosmic rays incident on the top of the atmosphere and neutrino creation in the decays of pions and kaons within atmospheric showers. These physical phenomena which impact the shape of the neutrino flux will be




Figure 2.1: The Bartol un-oscillated neutrino flux prediction at Super-Kamiokande Modelled at a solar minimum. Antineutrino fluxes are shown with a dashed line.

illustrated using the fluxes predicted by the Bartol atmospheric Monte Carlo. A detailed discussion of this Monte Carlo and its methodology including how it models some of the physical effects discussed in this chapter is reserved for Chapter 3. Subsequent chapters will also consider the hadronic interactions of cosmic rays within the atmosphere and the precise form of the primary cosmic ray flux.

The un-oscillated atmospheric neutrino fluxes are shown between 0.1-10 GeV neutrino energy in Figure 2.1. Notice the rapid decrease of the flux with increasing energy. This feature is a consequence of the shape of the primary cosmic ray flux which can be approximately described by a simple power law in energy decreasing as $\phi \sim E^{-2.7}$ [36] (greater detail on the flux parameterisation is given in Chapter 10).

These primary cosmic rays interact to produce mesons with neutrinos created in their decays. Neutrinos within the 0.1-10 GeV energy range are created in pion and kaon decay chains.

Atmospheric neutrinos are created from pions via the following modes:

$$\begin{aligned}\pi^\pm &\rightarrow \mu^\pm \nu_\mu && (99.9\%) \\ \mu^\pm &\rightarrow e^\pm \nu_e \nu_\mu && (\sim 100\%),\end{aligned}\tag{2.1}$$

and from kaons (with daughter muon and charged pion decays into neutrinos as above) via the modes:

$$\begin{aligned}K^\pm &\rightarrow \mu^\pm \nu_\mu && (64\%) \\ K^\pm &\rightarrow \pi^0 e^\pm \nu_e && (5\%) \\ K^\pm &\rightarrow \pi^0 \mu^\pm \nu_\mu && (3\%) \\ K^\pm &\rightarrow \pi^\pm \pi^0 && (21\%) \\ K^\pm &\rightarrow \pi^\pm \pi^\pm \pi^\mp && (6\%) \\ K^\pm &\rightarrow \pi^\pm \pi^0 \pi^0 && (2\%) \\ K_L^0 &\rightarrow \pi^\pm e^\mp \nu_e && (40\%) \\ K_L^0 &\rightarrow \pi^\pm \mu^\mp \nu_\mu && (27\%) \\ K_L^0 &\rightarrow \pi^+ + \pi^- + \pi^0 && (13\%) \\ K_S^0 &\rightarrow \pi^+ + \pi^- && (69\%) \end{aligned}\tag{2.2}$$

The branching ratio of the decay channels is shown (from Ref. [14]). Whether the decay produces a neutrino or antineutrino will depend on the charge of the outgoing lepton via lepton number conservation. Neutral kaon decay modes that only include neutral pions are not shown as neutral pions decay electromagnetically and will not produce neutrinos.

Figure 2.2 shows the fraction of the neutrino flux as a function of energy arising from these decay modes. The majority of muon neutrinos within this energy range originate in decays of pions and muons from pion decay while most electron neutrinos originate in decays of muons from pion decay. There is a smaller component arising from the decay of kaons and kaon daughter products. Since electron neutrinos are typically created further down the chain of decays they will typically be created at lower altitudes in the atmosphere than muon neutrinos. It can immediately be seen that the creation of muon neutrinos is highly coincident with the creation of




Figure 2.2: Fraction of neutrino flux coming from pion and kaon decay modes. Shown separately for each neutrino flavour. The fraction originating from pion decay and decay of pion decay products is shown with a solid line and the fraction originating from kaon decay and kaon decay product decay is shown with a dashed line. Kaon decay includes decays of both charged and neutral kaons.

electron neutrinos. These pion and kaon decay channels will be considered further in Section 2.3.3 when considering the ratios of fluxes of different flavours of neutrino.

Beyond the 10 GeV neutrino energy considered here, neutrinos coming from kaon decays become increasingly important, dominating above 100 GeV neutrino energy [37]. At higher energies still, the decay modes of the D meson (which contain charm quarks) become important, dominating the creation of neutrinos beyond 10^6 GeV [38]. At the neutrino energies considered within this thesis, the D meson (which can only be created when a heavy charm anti-charm pair is created) has a

negligible contribution ¹ to the neutrino flux and will not be considered.

This section has only considered the production of electron and muon neutrinos. The creation of the tau neutrino, which is created alongside the heavy tau lepton is not kinematically allowed in pion and kaon decays. It can be produced in the decays of the D meson (predominantly the D_s meson [39]); however, as already discussed, the production of the D meson is negligible within the energy ranges relevant here. The un-oscillated tau neutrino flux is therefore negligible within the energy ranges considered here with tau neutrinos coming from the oscillation of muon and electron neutrinos.

2.2 The Influence of Magnetic Fields

The primary cosmic rays which instigate the atmospheric showers and thus give rise to the neutrino flux are of galactic origin [36]. While they are assumed to be isotropic and homogeneous in deep space, their propagation to the top of the Earth's atmosphere involves these charged particles encountering both the magnetic fields within the solar system and the magnetic field of the Earth. Both of these magnetic fields will affect the flux of primary cosmic rays and thus the flux of neutrinos. When considering the propagation of charged particles through magnetic fields it is often more helpful to consider their rigidity, defined as the ratio of their momentum to their charge, than their energy or momentum directly. This is because rigidity is proportional to their bending radius in a magnetic field.

2.2.1 Solar Magnetic Fields and Cosmic Ray Flux Modulation

The Sun's magnetic field has been observed to fluctuate between a minimum and maximum in solar activity with an 11-year cycle, during which the magnetic field flips polarity. During periods of higher solar activity, there is an increase in the emission of plasma from the surface of the sun into the interplanetary medium

¹From Ref. [38] the D meson component of the 30 GeV neutrino flux is estimated to be 0.01 %. This component will be smaller at the lower neutrino energies considered in this thesis.

(solar wind) [40]. This solar wind and its associated magnetic field scatter the primary cosmic rays reducing their flux at the top of the atmosphere [41]. Therefore, during periods of higher solar activity, there is an associated decrease in cosmic ray (and therefore neutrino) flux. Since the magnetic fields are turbulent this effect is isotropic and homogeneous at the top of the Earth's atmosphere. It does not affect the directionality of the primary flux only its overall normalisation at a single rigidity. The solar wind has the greatest impact on the lowest energy cosmic rays (below ~ 50 GeV). This effect is known as the solar modulation of the cosmic ray flux.

Historically solar activity was tracked by observing the appearance of sunspots (cooler regions on the surface of the sun associated with large magnetic fields) [40]. These sunspots occur more frequently at periods of solar maximum. Modern measurements use neutron monitor data to track solar activity, these are measurements of the daughter particles produced in atmospheric interactions which reach the surface of the Earth [42]. These measurements will be directly correlated to cosmic ray intensity at the top of the atmosphere and thus can be used to estimate the effects of solar modulation on the primary flux.

For the remainder of this thesis, neutrino fluxes will be shown at a solar minimum. To calculate the neutrino flux at a different point in the solar cycle, neutron monitor data is first used to determine the modulated cosmic ray flux. The neutrino flux can then be calculated using this modulated cosmic ray flux. This has previously been done by considering the modulation of the primary cosmic ray flux using neutron monitor data taken during the specific time period an experiment has detected atmospheric neutrinos [43].

2.2.2 The Earth's Magnetic Field and Geomagnetic Cut-off

The Earth's magnetic field will also bend cosmic rays and at lower energies, cosmic rays can be completely deflected before they are able to interact with the atmosphere [37]. Unlike the solar magnetic field, the Earth's magnetic field is far more stable and will not vary significantly with time. This shielding of the lower energy cosmic rays will vary significantly as a function of position on the Earth. For example,

a cosmic ray travelling with respect to the local vertical at the north pole will experience a far lower magnetic force where its direction is approximately parallel to the magnetic field lines than a cosmic ray at the equator where it is perpendicular to those field lines. This effect creates, for each trajectory and at every point in the Earth's atmosphere, a minimum cosmic ray rigidity below which it will not be possible for a cosmic ray to penetrate the atmosphere. This effect is known as the geomagnetic cut-off of cosmic rays [37]. Its effect on the neutrino flux will be further discussed in Section 2.3 and the method by which it is modelled within the simulation is discussed in Chapter 3.

2.3 Symmetries and Asymmetries in the Un-Oscillated Neutrino Flux

Within experimental physics, it is often found that while the systematic uncertainty on an absolute measurement may be high, these systematic uncertainties can be highly correlated with similar measurements made within the same experiment. Therefore it is often desirable to consider the ratio between two measurements where there can be a significant cancellation of uncertainties enabling features to be seen more clearly. This is true for neutrino fluxes where many of the uncertainties on both the measurement of the neutrino fluxes in a detector and the flux modelling will cancel when considering ratios between the fluxes. This enables the effects of neutrino oscillation to be seen more clearly. Chapter 1 briefly introduced some of these concepts within the context of the initial experimental measurement of neutrino oscillation by the Super-Kamiokande experiment. This section will consider in more detail some of the symmetries and asymmetries within the un-oscillated fluxes and discuss how they can be sensitive to neutrino oscillations. Chapter 9 will later demonstrate the cancellation of uncertainties associated with flux ratios when compared to absolute flux measurements within the uncertainty model developed in this thesis.




Figure 2.3: Symmetry between the up-going and down-going atmospheric neutrino flux. Diagram shows an up-going neutrino produced in the atmosphere on the far side of the Earth traversing the Earth and entering the detector at zenith angle θ_Z^{up} . A down-going neutrino produced in the atmosphere above the detector is shown at zenith angle θ_Z^{down} such that $\theta_Z^{\text{up}} = \pi - \theta_Z^{\text{down}}$. Both neutrinos are produced in the atmosphere, for clarity their paths have been extended beyond the atmosphere in the diagram. Figure 2.3a demonstrates the symmetry between the up-going and down-going neutrino fluxes in the absence of the geomagnetic field while the figure 2.3b demonstrates the breaking of this symmetry by the geomagnetic field. Figure 2.3a has also labelled the distances travelled by each neutrino from production to the detector as d_{up} and d_{down} respectively.

2.3.1 Up-Down Ratio

Figure 2.3 shows two neutrinos entering the detector, the up-going neutrino is produced in the atmosphere on the far side of the Earth and traverses through the Earth to the detector while the down-going neutrino is produced in the atmosphere directly above the detector. The zenith angle of each neutrino at the detector is shown, by convention zenith angle is defined such that $\cos \theta_Z < 0$ ($\cos \theta_Z > 0$) correspond to the up-going (down-going) neutrinos. As can be seen in Figure 2.3a the up-going neutrino at θ_Z^{up} is produced symmetrically within the atmosphere to the down-going neutrino at $\theta_Z^{\text{down}} = \pi - \theta_Z^{\text{up}}$ with both being produced at angle α to the local vertical. The area of atmosphere subtended by $\theta_Z^{\text{up}} + d\theta$ is proportional to the distance to the detector squared (d_{up}^2 from Figure 2.3a) and hence more




Figure 2.4: ν_μ fluxes as a function of cosine zenith angle for 3 neutrino energy bins, 0.1-0.32 GeV (top left), 1-3.16 GeV (top right) and 3.16-10 GeV (bottom). In each graph, the un-oscillated flux is calculated using the Bartol Monte Carlo described in Chapter 3.

neutrinos are produced in that region of atmosphere for up-going neutrinos than the down-going neutrinos. The probability for neutrinos to reach the detector a distance d away will follow an inverse square law with the fraction of up-going neutrinos reaching the detector proportional to d_{up}^{-2} cancelling with the d_{up}^2 factor. Therefore it can be concluded that in the absence of magnetic fields where the cosmic ray flux is the same at all points on the Earth, the un-oscillated flux for both the down-going and up-going neutrinos is identical [44].

With the inclusion of the geomagnetic field (as in Figure 2.3b) it can immediately be seen that this symmetry is broken with the differences in the magnetic field at the up-going and down-going neutrino production site introducing differences in the cosmic ray geomagnetic cut-off and therefore changes in neutrino production. This effect will primarily be seen in the lower energy neutrino fluxes which are most sensitive to changes in the flux of the lower energy primary cosmic rays. Figure 2.4 shows the un-oscillated flux at Super-Kamiokande for three different neutrino energies. It can be seen that the up-going and down-going flux is less

symmetric at lower energies.

The up-going and down-going flux in Figure 2.4 are binned in neutrino energy and cosine zenith angle, effectively averaging over all azimuthal angle directions for the arriving neutrino. For the up-going flux, this will average neutrino production across a large range of geomagnetic fields whereas for the down-going flux the averaging over azimuthal angle is only over the region of the atmosphere directly above the detector. The down-going flux is therefore highly dependent on the local geomagnetic field above the detector and will vary significantly for detectors in different locations. This effect will be seen in Chapter 11 when the neutrino fluxes calculated using the models presented within this thesis are shown at different detector sites.

While not exact, this symmetry in the un-oscillated flux is important for the analysis of neutrino oscillations, as was discussed in Chapter 1 and can be seen in Figure 2.3 there is an approximate relationship between the zenith angle of detection and the distance travelled by a neutrino. The effects of neutrino oscillations, which depend on the ratio of neutrino energy to distance travelled can then be studied by comparing the ratio of the up-going to down-going flux.

2.3.2 East-West Asymmetry

Having considered symmetries in neutrino zenith angle, neutrino azimuthal angle is now considered. The east-west asymmetry is an asymmetry between the un-oscillated flux arriving from the east and west direction at the detector [45]. By convention, the azimuthal angle is defined as being 0° for neutrinos arriving from the north, 90° for neutrinos from the east and so forth. Figure 2.5 demonstrates the source of this asymmetry for a detector viewed in the plane of the equator where the magnetic field is coming out of the page and the detector east is to the left and west to the right. This figure shows the trajectory that two low-energy primary cosmic rays would need to take in order to approach the detector at the local east and west. The trajectory for the cosmic ray arriving from the west is an allowed trajectory, originating from outside the Earth's magnetic field. By contrast the trajectory of an equivalent cosmic ray arriving from the east circles back towards




Figure 2.5: Bending of the primary cosmic rays with respect to the detector's local east and west as the origin of the east-west asymmetry in the atmospheric neutrino flux. For illustrative purposes the detector is positioned at the equator with the magnetic field coming out of the page, the north pole is marked alongside the local east and west at the detector. This diagram shows a cosmic ray coming from the west on a physical trajectory it would be possible for a cosmic ray originating in deep space to take. A similar cosmic ray coming from the east would need to originate on the Earth, such a trajectory is unphysical for cosmic rays from deep space.

the Earth meaning it is on a disallowed trajectory and cosmic rays cannot arrive from that direction. This in turn leads to a reduction in neutrinos coming from the east direction with respect to the west direction. The azimuthal distribution of the un-oscillated neutrino flux is shown in Figure 2.6 the deficit in the flux from the east (90°) can be seen with respect to the west (270°). Similar to the up-down fluxes it can be seen that the strength of the asymmetry decreases at higher neutrino energies.

2.3.3 Neutrino Flavour Ratios

As was discussed in Chapter 1 the relationship between muon and electron neutrinos was instrumental in establishing experimental evidence for neutrino oscillations.




Figure 2.6: ν_μ fluxes as a function of azimuthal angle for 3 neutrino energy bins, 0.1-0.32 GeV (left) and 1-3.16 GeV (right). In each graph, the un-oscillated flux is calculated using the Bartol Monte Carlo described in Chapter 3.

Considering the decay modes of pions in equation 2.1 it can be seen that for every pion and subsequent muon decay, two muon neutrinos (ν_μ and $\bar{\nu}_\mu$) and one electron neutrino (either ν_e or $\bar{\nu}_e$) are produced. It does not immediately follow that the un-oscillated muon neutrino flux should be twice that of the electron neutrino as is often stated. However, as a consequence of the pion and muon decay kinematics, the three neutrinos have roughly the same energies and therefore the flavour ratio of muon to electron neutrinos as 2 : 1 is obtained [36]. The muon to electron neutrino flux ratio is shown in Figure 2.7. It can be seen that while the 2 : 1 ratio holds well at low neutrino energies it breaks down as the neutrino energy increases. This is a consequence of the muons which decay to higher energy neutrinos having a greater Lorentz boost and therefore a longer lifetime as measured in the frame of the Earth. This increases the probability that they will reach the ground before decaying to a higher energy neutrino and therefore the ratio of muon to electron neutrinos increases.

The ratio of both ν_e to $\bar{\nu}_e$ and ν_μ to $\bar{\nu}_\mu$ are important due to their sensitivity to neutrino oscillations within the full three flavour framework discussed in Section 1.3. It can be seen from Figure 1.5 that the $\nu_e/\bar{\nu}_e$ flux flavour ratio observed by the detector will be sensitive to the resonance in oscillation probability and therefore to the Δm_{31} mass ordering. The $\nu_\mu/\bar{\nu}_\mu$ flavour ratio will likewise be sensitive to this effect via the deficit of muon neutrinos caused by the resonance in their oscillation to electron neutrinos.




Figure 2.7: Ratio of the flux of different flavours of neutrinos as a function of neutrino energy. $\nu_e/\bar{\nu}_e$ (top left), $\nu_\mu/\bar{\nu}_\mu$ (top right) and $(\nu_\mu+\bar{\nu}_\mu)/(\nu_e+\bar{\nu}_e)$ (bottom). In each graph, the un-oscillated flux is calculated using the Bartol Monte Carlo described in Chapter 3.

The un-oscillated $\nu_e/\bar{\nu}_e$ and $\nu_\mu/\bar{\nu}_\mu$ flavour ratios are also shown in Figure 2.7. Considering the decay modes in 2.1, it can be seen that the un-oscillated $\nu_e/\bar{\nu}_e$ ratio will depend predominantly on the ratio between π^+ and π^- production in the atmosphere. Uncertainties affecting the total $\pi^\pm$ production, such as the total flux of cosmic rays, will largely cancel reducing the uncertainty in the flavour ratio. The greater production of π^+ than π^- in atmospheric showers, caused by the greater number of protons than neutrons in the cosmic rays, causes the ν_e flux to be greater than the $\bar{\nu}_e$ flux. At lower neutrino energies where the probability of muon decay is larger, each pion produced will result in both a ν_μ and $\bar{\nu}_\mu$ being created regardless of its charge and the flavour ratio is approximately one. At higher neutrino energies the ν_μ and $\bar{\nu}_\mu$ depends on the $\pi^\pm$ production ratio and therefore there is a greater flux of ν_μ .

Due to their importance to the neutrino oscillation analysis, the new models developed within this thesis will be used to calculate the flavour ratios of the fluxes described here. The neutrino flux ratios and uncertainties obtained by this

work are presented in Chapter 11.

3

The Bartol Atmospheric Neutrino Monte Carlo

This chapter provides a summary of the Bartol atmospheric neutrino flux Monte Carlo simulation. The analysis in this thesis primarily focuses on an update to the main sources of uncertainty within this calculation: the primary cosmic ray flux and the hadronic interactions of these primary cosmic rays. The details of the Bartol calculation are described in detail in Ref. [44] and [46], this chapter will summarise the methodology used by the Monte Carlo simulation to estimate the neutrino flux. The uncertainty estimate on the last atmospheric neutrino flux release from 2004 is summarised. In addition, Section 3.5 describes a new method implemented into the Monte Carlo within this work to calculate the neutrino flux based on the probability that an individual neutrino intersects the detector. This method is referred to as the decay rotation method.

3.1 Overview of Bartol Monte Carlo

The Bartol¹ atmospheric neutrino simulation calculates the un-oscillated atmospheric muon and electron neutrino (and antineutrino) flux using Monte Carlo methods. Detailed descriptions of the simulation were published after the last major updates to the flux prediction in 2004 and 2006 when the one-dimensional calculation was

¹Named after the Bartol Institute in Delaware where it was originally developed

extended to a three-dimensional calculation [46, 47]. This thesis presents a further update to this prediction and its uncertainties by incorporating recent hadron production and cosmic ray data. A brief overview of the techniques used within the Monte Carlo simulation is given in this chapter. In addition to this, Section 3.5 describes an alternative method investigated within this work to calculate the neutrino flux by modifying the Monte Carlo. One of the major challenges within any atmospheric neutrino flux simulation is the comparatively small fraction of the Earth's surface covered by a neutrino detector. When three-dimensional atmospheric showers are simulated most of the neutrinos produced will miss the detector and not contribute to the flux calculation. As a consequence, the Monte Carlo has been designed to be as efficient as possible in other areas to maximise the neutrino statistics and minimize Monte Carlo statistical errors.

3.2 1D and 3D Calculations and the 3D Geometrical Effect

To remove the inefficiencies related to most neutrinos missing the detector, early neutrino flux calculations approximated atmospheric showers as one-dimensional [48, 49]. Within this limit, and ignoring the effects of the bending of secondary particles by the geomagnetic field, any neutrinos produced will be parallel to the original cosmic ray generated. Therefore only primary cosmic rays pointing directly at the detector need to be considered. While the one-dimensional approximation is accurate for higher energy neutrinos ($\sim > 4$ GeV) which are produced in atmospheric showers with large Lorentz boost factors, it breaks down for lower energy neutrinos where the correlation between the neutrino direction and the cosmic ray direction reduces. Naively, it might be assumed that the difference between the one and three-dimensional calculation is minimal since the neutrino flux is calculated by averaging over the effects of many atmospheric showers from different directions. In fact, this assumption is false due to a geometrical effect first demonstrated by Lipari [50] which enhances the neutrino flux arriving at a neutrino detector with respect to the




Figure 3.1: The geometry of an isotropic source of neutrinos on a spherical shell as viewed by a detector. The area of atmosphere subtended the vertical solid angle (around $\cos \theta_{\text{zenith}} = 1$) is given by dA_1 and horizontal ($\cos \theta_{\text{zenith}} = 0$) by dA_2 . The main text compares the one-dimensional limit to a second limit where the flux can be described as an isotropic source of neutrinos on a spherical shell at altitude L . dA'_2 is the equivalent perpendicular area subtended by the solid angle at the horizontal.

local horizontal. This geometrical effect is in addition to the increased atmospheric depth for horizontal moving particles which separately increases the neutrino flux due to a greater probability of meson and muon decay before reaching the ground.

To demonstrate this geometrical effect, consider this illustration based on the arguments by Lipari [50]. First consider a simplified atmospheric neutrino production model where all incident primary cosmic rays create a single neutrino in an “interaction” at a fixed altitude above the surface of the Earth. Figure 3.1 shows an example of the detector and atmosphere in this setup with neutrinos created at altitude L . To simplify this illustration, magnetic fields and the cut-off of primary cosmic rays are temporarily ignored and therefore the flux of cosmic rays reaching altitude L is homogeneous and isotropic. Consider two limits:

- Firstly consider the one-dimensional limit where the neutrino is produced parallel to the cosmic ray. Since the cosmic ray flux is isotropic at all points in space, the neutrino flux at the detector must also be isotropic and an equal number of neutrinos will reach the detector from all directions.
- Secondly consider the limit where neutrinos are produced at a random direction with respect to the cosmic ray direction. Within this limit, there is effectively a source of neutrinos being produced isotropically in all directions at each point on the shell at altitude L (see Figure 3.1). To first order there is

a cancellation with the area of the atmosphere (and thus rate of neutrino production) increasing with r^2 , distance from the detector squared, and the probability of a neutrino reaching the detector decreasing with increasing distance as r^{-2} (similar to the up-down symmetry discussed in Section 2.3.1). This cancellation would be exact between the vertical and horizontal neutrinos if the area of atmosphere subtended by the horizontal were given by dA'_2 , the area perpendicular to the horizontal direction in Figure 3.1. In fact, the area is given by dA_2 ($dA_2 > dA'_2$) due to the additional geometrical effect caused by the detector not being at the centre of the atmospheric shell and therefore the atmospheric area not being perpendicular to the horizontal direction. This causes an enhancement of the neutrino flux coming from the local horizontal due to this greater area of the atmosphere subtended.

This toy example illustrates the origin of the 3D effect which is visible in Figure 2.4 where it can be seen that the neutrino flux is greatest at the horizontal ($\cos\theta_Z \sim 0$). This effect will be important when considering hadronic interactions since the flux arriving at the horizontal has greater dependence on daughter particle production at greater angles with respect to the original primary particle direction. Since the one-dimensional approximation cannot accurately predict the neutrino flux in the 0.1-10 GeV energy range, the full three-dimensional Monte Carlo will be used for the studies in this thesis.

3.3 Injection of Primary Cosmic Rays

The Bartol Monte Carlo begins with the creation of primary cosmic rays. As discussed in Section 2.2.2, while the flux of cosmic rays is isotropic in deep space the effects of the Earth's magnetic field alter the flux of cosmic rays at the top of the Earth's atmosphere. The brute force solution to simulating this would generate cosmic rays isotropically and homogeneously a large distance away from the Earth and propagate them forward through the Earth's magnetic field. Since the vast majority of cosmic rays generated via this method would miss the Earth it is too

computationally inefficient. Instead, a method described by Favier et al. in Ref. [51] is used. This method invokes Liouville’s phase space theorem to state that if a cosmic ray trajectory at the top of the Earth’s atmosphere can be traced backwards in time to a large distance away from the Earth where the flux can be assumed isotropic then the cosmic ray flux around that trajectory at the top of the atmosphere will be isotropic also. Due to the effects of the Earth’s magnetic field, not all cosmic ray trajectories can be traced back to deep space (as was shown in Section 2.3.2) and any trajectories which circle back towards the Earth are unphysical and are excluded from the flux. This method of generating the flux isotropically at the top of the atmosphere, backtracing them to deep space and excluding those that intersect with the Earth was demonstrated by Favier et al. to provide equivalent results to the brute force method [51]. The Bartol Monte Carlo injects primary cosmic rays with an isotropic distribution at an altitude of 80 km [46]. To improve the efficiency of the calculation only those primary cosmic rays which eventually produce the neutrinos in the detector undergo the backtracing calculation. Only primary cosmic rays which can be traced back to a distance further than thirty Earth radii away are then included as part of the primary cosmic ray flux [46]. For the 2004 Bartol neutrino flux calculation, the primary fluxes are modelled using the ICRC01 parameterisation from 2001 [52]. Cosmic ray parameterisations will be discussed further in Chapter 10 where a new cosmic ray model will be incorporated.

3.4 Hadronic Interactions with TARGET

Once the primary cosmic rays are generated (via the methods described in 3.3) they are propagated forwards in time through the Earth’s geomagnetic field and variable density atmosphere. Their interaction within the atmosphere is modelled using the TARGET hadron interaction generator [44]. To model the hadronic interactions of the helium and heavier nuclei component of the cosmic rays the superposition approximation will be used [46]. Under this approximation, the interactions of a nucleus of mass A and energy E is treated as equivalent to A individual nucleons

each of energy $\frac{E}{A}$. Since the vast majority of the flux comprises protons and helium nuclei this approximation is reasonable. Further details on the primary cosmic ray composition are discussed in Chapter 10. The Bartol fluxes released in 2004 were calculated using the TARGET2.1 generator [46]. An updated version of this generator, developed prior to the work in this thesis, called TARGET2.4 [44] will be used within the calculation presented within this thesis. Prior to this work, the neutrino flux calculated using TARGET2.4 had not been published. A brief description of the general principles of TARGET is given here alongside some of the features added between the TARGET2.1 and TARGET2.4 releases, further details can be found in reference [44].

TARGET simulates hadronic interactions by generating daughter hadrons and distributing the interaction energy amongst them, ensuring the conservation of energy and other relevant quantum numbers within each interaction. To determine the spectra of daughter particles, TARGET uses a combination of parameterisations to datasets that existed when it was developed and phenomenological models. TARGET has been specifically designed to be computationally fast for atmospheric neutrino production. At lower interaction energies (below 4 GeV) TARGET incorporates separate modules to handle resonance production. TARGET2.4 was a direct upgrade to TARGET2.1 which modified TARGET2.1 in order to include fragmentation nucleons (nucleons emitted from the target nucleus as a result of nuclear fragmentation) and anti-baryon pair production.

The main focus of this thesis will be the tuning of the neutrino flux by comparing recent hadron production datasets to TARGET. A significant improvement in the flux prediction is possible due to the large number of datasets published since TARGET2.4 was developed in 2005. Of the datasets described in Chapter 4 which will be used to tune the flux, only the datasets by Eichten et al. and Allaby et al. were used in the development of TARGET2.4.




Figure 3.2: The uncertainties on pion and kaon production as estimated by the Bartol calculation in 2006. The pion and kaon production phase space is broken into phase space regions based on incident particle energy E_i and outgoing particle x_{LAB} , the ratio of longitudinal momentum of the outgoing meson to incident particle momentum. An additional uncertainty on production at the highest energies is included and described in more detail in [47]. Figure taken from Ref. [47]

3.4.1 2006 Bartol Neutrino Flux Uncertainties

An uncertainty estimate of the neutrino flux calculated within the original Bartol calculation with TARGET2.1 and ICRC01 primary fluxes was presented in 2006 [47]. The uncertainties in the atmospheric neutrino flux estimate are dominated by the uncertainties in primary cosmic ray and particle production in hadronic interactions with all other effects estimated to have an uncertainty of less than 2% [44]. The uncertainty on pion and kaon production was obtained by dividing the phase space for particle production into regions labelled A-I for pion production and W-Z for kaon production, the regions and uncertainties are shown in Figure 3.2. The regions are integrated over pion and kaon transverse momentum and uncertainties are assumed correlated within the same region but uncorrelated between regions [47]. In addition, uncertainties are assumed to be correlated between positive and negative meson production with an additional uncertainty of 5% assigned for an overall uncertainty on the ratio of positive to negative meson production. Within




Figure 3.3: Absolute muon neutrino flux uncertainty estimate from the Bartol calculation in 2006. Upper case labels are for hadron production uncertainties and correspond to the regions in Figure 3.2. Lower case labels are due to the uncertainty of the cosmic ray fluxes within the ICRC01 parameterisation which is discussed further in Chapter 10. Figure taken from [47]

this model, the absolute neutrino flux uncertainty is estimated to be around 15% (see Figure 3.3). In Chapter 9 and Chapter 11 the uncertainty estimate will be updated after the incorporation of the hadron production datasets with hadron production phase space being considered in three dimensions for the first time (i.e. without the integration over transverse momentum).

3.5 Extended Detector and Decay Rotation Method

The daughter particles created in hadronic interactions continue to be propagated through the atmosphere and the geomagnetic field where they may interact further or decay. Particles which reach the surface of the Earth will stop in the Earth before decaying and only create very low-energy daughter particles which are not considered within the calculation. The simulation stores all neutrinos which are produced on a trajectory which intersects with the detector.

To improve the efficiency of the three-dimensional calculation the Bartol simulation extends the detector size across the Earth's surface to a disc centred on the detector location of radius 500 km [46]. This improves the efficiency of simulation but has two downsides:

- Firstly, the calculated neutrino flux will be the averaged neutrino flux over the area of the disc, with any differences due to variation in the geomagnetic field across the disc averaged over. Although this average flux is not the same as the true flux at the detector the difference was estimated to be smaller than 0.5% [46] and not a large effect.
- Secondly, while the detector is extended horizontally across the surface of the Earth it cannot be extended vertically via the same method where artificially reducing or increasing the atmospheric depth would introduce significant unphysical effects on the flux. Consequently, neutrinos travelling vertically towards the detector experience a much larger effective detector area than those travelling horizontally, where the effective detector becomes zero for a neutrino moving exactly at the horizontal. The Monte Carlo simulation, therefore, records far fewer horizontal neutrinos which each have larger contributions towards the flux and there is a corresponding increase in statistical uncertainty.

To resolve these issues an alternative method was implemented into the simulation for the first time within this work. It is described in the subsection below.

3.5.1 Decay Rotation Method

At the point of meson or muon decay into a neutrino, the original Bartol Monte Carlo chooses a direction in which to produce the neutrino relative to the parent direction. If that neutrino is created on a trajectory which intersects with the extended detector it is recorded and contributes to the neutrino flux. There is an underlying probability distribution function for the angular distribution of neutrino creation. The alternative method presented here calculates the probability that any neutrino produced will intersect the detector. The neutrino flux can then be

calculated as before with each neutrino assigned a weight equal to this probability. This method of calculating the flux will be referred to as the *decay rotation* method.

The probability that a neutrino is produced at solid angle Ω in the laboratory frame with respect to the decaying parent trajectory is then calculated to be:

$$f(\Omega) = \frac{f'(\Omega')}{\gamma^2(1 - \beta \cos \theta)^2}, \quad (3.1)$$

with $f'(\Omega')$ the decay probability in the centre of mass frame, θ as the angle between the neutrino and meson trajectory and γ and β taking their usual definitions within special relativity. For the decay of the pion and kaon (which have no spin), the neutrino will be produced isotropically within the rest frame of the meson and $f'(\Omega') = (4\pi)^{-1}$. Muons are produced polarised, and the effects of muon polarisation have previously been found to impact the neutrino flux [48]. As such neutrinos produced within muon decays are not produced isotropically in the muon rest frame but instead depend on the polarisation angle of the muon via distributions given in [53] from which the value of $f'(\Omega')$ is calculated. The probability of a neutrino intersecting the detector is given by:

$$P_{\text{det}} = f(\Omega)\Delta\Omega_{\text{det}}, \quad (3.2)$$

where $\Delta\Omega_{\text{det}}$ is the solid angle subtended by a small finite detector at the point of neutrino production. The detector can then be modelled as having a unitary cross-sectional area from all angles of approach (effectively a sphere of unit cross-sectional area) which removes the directional dependence of the detector size within the extended detector method and the averaging of the flux over the large disc. The factor $\Delta\Omega_{\text{det}}$ will then scale as r^{-2} where r is the distance between neutrino production and the detector.

While the neutrino flux can be calculated via this method a number of changes needed to be made to the Monte Carlo in order to produce acceptable statistical errors. Without any additional changes, this new version of the simulation records a large number of neutrinos with very low weights and a small number of neutrinos

with very large weights which dominate the statistical uncertainties. The following alterations were made to the Monte Carlo to increase the generation of neutrinos with higher weights and have a greater contribution to the flux:

- Firstly, a weight cut-off was introduced. This removes the large number of low-weight neutrinos where their overall contribution to the flux is dwarfed by smaller numbers of high-weight neutrinos. Any neutrino weight lower than a prescribed weight threshold is either removed from the calculation or upscaled to the weight threshold with the probability of upscaling given by the ratio of the neutrino weight before upscaling to the weight threshold. For example, if the weight threshold is set at 0.1, a neutrino of weight 0.01 is removed with a probability of 90% and upscaled to 0.1 with a probability of 10%. The weight threshold was chosen to be dependent on the neutrino energy. At 10 GeV the neutrino flux is a factor of $\sim 10^3$ smaller than at 0.1 GeV and the weight threshold was chosen to reflect that. A weight threshold of 10^{-17} was chosen at a neutrino energy of 0.1 GeV which decreased to 10^{-20} at 10 GeV. The weight threshold in between these two energies was determined using a linear interpolation between the logarithm of the neutrino energies. Note that this threshold is only for the weight due to the probability of decay into a unit area, overall normalisation factors such as the number of primary particles generated within within a run and other factors discussed in Section 3.6 are not included.
- Secondly, the location at which the primary cosmic rays are injected into the simulation needed to be tweaked. Generating primary cosmic rays the “natural” way (uniformly on the sphere of injection at 80 km above the surface of the Earth) within the decay rotation method, led to a very small number of cosmic rays produced above the detector and therefore very few down-going neutrinos being stored. Within the extended detector method, more neutrinos are produced in the up-going direction also but they have a much smaller chance of intersecting the detector, leading to a roughly equal number of




Figure 3.4: A primary cosmic ray is generated at a small angle to the detector direction. The resultant down-going neutrino must be produced at a much larger angle to the primary cosmic ray in order to be able to reach the detector due to being produced much closer to the detector than the original primary cosmic ray.

up-going and down-going neutrino events being recorded. To improve the down-going neutrino statistics the Monte Carlo was altered to inject primary cosmic rays uniformly as a function of their position measured as cosine zenith angle at the neutrino detector. Each neutrino produced by a primary cosmic ray is then assigned an additional weight:

$$W_{\nu} = \frac{P_{\text{true}}(\mathbf{r})}{P_{\text{gen}}(\mathbf{r})}, \quad (3.3)$$

as the ratio between the probability of generating a primary cosmic ray at position $\mathbf{r}$ with the physical (true) homogeneous distribution to the distribution which was actually used to generate primary cosmic rays.

- After application of the first two changes it was found that there were still very few neutrinos being recorded at higher energies (>1 GeV). Higher energy neutrinos are typically produced at small angles to the decaying parent particle due to an increase in their Lorentz boost factor. The rotation method forces the produced neutrino to the detector direction and therefore if the simulation produces too few pions directed towards the detector there will be too few high-energy neutrinos. To resolve this, the primary cosmic ray generation angle was altered to preferentially produce primary cosmic rays directed towards the detector with a Gaussian distribution. The resultant neutrino

is then assigned a weight via equation 3.3. It is important to note that the meson or muon which decays into the neutrino will be produced at some angle to the primary cosmic ray, with larger angles at smaller energies. This means that a primary cosmic ray generated at a direction pointing away from the detector can still produce the neutrino parent directed towards the detector and thus a high-energy neutrino with a larger weight. Producing primary cosmic rays too tightly distributed towards the detector will mean too few neutrinos produced via the aforementioned method will be created, leading to large statistical uncertainties. The width of the Gaussian distribution which describes the angular generation of the primary cosmic rays therefore needed to be tuned at each primary energy with a larger width at smaller energies. This effect is different for up-going and down-going neutrinos. As an example consider a primary cosmic ray generated at an altitude of 80 km at a small angle to the detector in a toy simulation where the primary cosmic ray creates a neutrino at 20 km. A down-going neutrino will be produced much closer to the detector than the cosmic ray and therefore will need to be produced at a much larger angle to the primary cosmic ray direction and consequently be lower in energy. This is shown in Figure 3.4. An up-going neutrino is produced at approximately the same distance from the detector as the primary cosmic ray and therefore will not need to be produced at such a large angle from the cosmic ray direction in order to reach the detector. Within the decay rotation method, primary cosmic rays above the detector need to be generated more tightly towards the detector in order to produce mesons and muons pointed towards the detector. These particles are then more likely to decay to a high-energy neutrino that reaches the detector increasing the number of these events.

Even after altering the distribution of primary cosmic rays, there were still individual neutrino events occurring with low probability where the neutrino was produced very close to the detector and therefore had a very large weight (see equation 3.2), which resulted in large statistical uncertainties. Figure 3.6 illustrates




Figure 3.5: A comparison of the ν_μ flux obtained via the decay rotation method (blue) and the extended detector method (orange). The flux is shown binned in neutrino cosine zenith angle within 4 neutrino energy bins. The statistical uncertainties are based on the weights of neutrinos contributing to each bin.

this problem alongside the need for the weight cut-off discussed in the first bullet point above. This figure is produced for a shorter Monte Carlo sample run with the cosmic rays produced at a fixed energy of 100 GeV. This figure shows the distribution of the weights which contribute to the neutrino flux at 1 GeV as a function of neutrino cosine zenith angle. The left hand plot shows the number of neutrinos recorded by this sample run and the right hand plot shows the fractional contribution to the neutrino flux by neutrinos in each weight bin. A large fractional contribution in a weight bin with a small number of neutrinos (such as the $\sim 20\%$ contribution by the single neutrino with weight $\sim 10^{-17}$ at $\cos \theta_Z \sim 0.15$) will lead to large statistical uncertainties. The weight threshold described above is shown with the red line.

A long Monte Carlo run was completed to try to generate as many of these events as possible and reduce statistical uncertainties. The ν_μ neutrino flux calculated using the decay rotation method is shown compared to the extended detector method




Figure 3.6: The graphs show the distribution of the neutrino weights due to the decay rotation probability at 1 GeV neutrino energy as a function of neutrino cosine zenith angle. They are made using a sample run of the decay rotation Monte Carlo with cosmic rays generated at a fixed energy of 100 GeV. The left hand plot shows the number of neutrinos within this Monte Carlo run at each rotational weight. The right hand plot shows the fractional contribution of those neutrinos to the total neutrino flux. The red line shows the threshold weight described in the main text. A total of 15,900 neutrinos were created at 1 GeV during this run.

calculation in Figure 3.5. The predictions are typically consistent within Monte Carlo statistical uncertainties which confirms that the original approximations made by the extended detector method are acceptable. Concerns about the small possibility of neutrino production very close to the detector creating very large statistical events need to be considered further before this method can be used within the final flux prediction. If the decay rotation method samples such events too infrequently they will cause large statistical errors where these events occur and where they don't occur the overall flux will be slightly lower than the true value. The results of this thesis will use the extended detector method with the inclusion of the decay rotation method left to future flux releases.

3.6 Obtaining the Neutrino Flux from Monte Carlo Neutrino Yields

The calculation of the neutrino flux is performed by executing multiple Monte Carlo runs at separate fixed cosmic ray energies. The cosmic ray energies relevant to

the production of 0.1-10 GeV neutrinos have previously been determined to span from 1 GeV to 10 TeV [46]. The Monte Carlo is run with 10 bins per decade in primary cosmic ray energy within this range. The yield of neutrinos at each primary cosmic ray energy is obtained from the Monte Carlo simulation. The neutrino flux is calculated by multiplying this yield by the primary cosmic ray flux at each energy bin and integrating over all primary cosmic ray energy bins and species. The effects of the solar wind can also be accounted for by modulating the primary cosmic ray flux at this stage; however, all the neutrino fluxes presented within this thesis will be shown assuming a solar minimum. This method has the advantage of being able to alter the cosmic ray flux without re-running the Monte Carlo simulation and reducing Monte Carlo statistical effects from high-energy cosmic rays which occur infrequently but with large neutrino yields due to their high energies. The atmospheric neutrino flux is given as $dN_\nu/d\log E$ in units of $m^{-2}s^{-1}sr^{-1}$. Where the neutrino flux is not explicitly binned in cosine zenith or azimuthal angle it is presented as averaged over incoming directions and hence the sr^{-1} in the units.

4

Hadron Production Datasets for Atmospheric Neutrino Flux Predictions

The hadron interaction model in air showers is the dominant source of uncertainty on the modelled atmospheric neutrino flux. This chapter focuses on the measurement of hadron production by fixed target experiments in the context of the relevant hadron production for modelling atmospheric showers. This begins by considering the relative contribution to the atmospheric neutrino flux from different regions of pion and kaon production phase space. The general principles behind hadron production measurements are then discussed alongside a summary of each of the experiments used in this analysis.

4.1 Hadron Interaction Phenomenology

Atmospheric neutrinos are created in air showers when a primary cosmic ray such as a proton or light atomic nucleus interacts with atomic nuclei in atmospheric gasses. The hadronic interaction creates unstable mesons which decay via the weak force with charged leptons and neutrinos being produced as decay products, as was discussed in Chapter 2. While the decay of mesons and leptons to neutrinos is well described by the theory of electroweak interactions, their creation occurs via the strong force which is theoretically described by QCD and characterised

by the α_s running coupling constant [13]. At the interaction energies relevant to this atmospheric neutrino study this coupling constant is large and perturbative QCD can not be used to directly calculate their production cross sections [54]. The uncertainty on these strong interactions forms the dominant uncertainty on the spectra of particles in the air showers and hence on the flux of neutrinos [47]. In order to understand the neutrino flux, the underlying hadronic interactions must be first to be understood. This is done through fixed target hadron production experiments which measure the spectra of mesons coming from interactions of a proton beam with a fixed nuclear target.

The basic form of each of the datasets used in this analysis are similar and consist of an inclusive double differential cross section for the production of a specific species of meson. Inclusive cross sections integrate over all possible final states alongside the meson produced [55]. This differs from exclusive cross sections where all of the daughter particles are specified. The inclusive cross sections work well in modelling the hadronic interactions within the neutrino flux as the coincidence of two particles from the same hadronic interaction is not of interest. This is because the probability of detecting two neutrinos from the same hadronic interaction is vanishingly small. Inclusive cross sections are given in the form:

$$\frac{d^2\sigma(pC \rightarrow \pi^+ + X)}{dx dy}, \quad (4.1)$$

where x and y are kinematic quantities of the outgoing pion (e.g. momentum p and production angle θ) [56]. Note that when using unpolarised proton beams, the particle interaction is cylindrically symmetric and the outgoing final state of the pion can be fully described by two independent kinematic variables at each proton beam energy. The X indicates all possible final states which result in the production of at least one positive pion making this an inclusive cross section.

Within this chapter, the terms primary particle or parent particle will refer to the projectile in either fixed target experiments or atmospheric interactions (the proton in equation 4.1). The term daughter particle refers to the outgoing hadron of interest in that interaction (the pion in equation 4.1).

4.2 Hadronic Interactions in Atmospheric Showers

As explained in Section 1.3, the focus of this analysis is primarily the 0.1-10 GeV neutrino flux region due to its sensitivity to the neutrino mass hierarchy and δ_{cp} in neutrino oscillation measurements. Initial studies were conducted to determine the relative contribution of different regions of hadron production phase space to the atmospheric neutrino flux at these energies. For each neutrino that contributes to the atmospheric flux, there is a chain of hadronic interactions starting with the original primary cosmic ray and ending in a pion or kaon which decays to either a neutrino directly or to a muon which subsequently decays to a neutrino.

Initially, the Bartol Monte Carlo was run with the TARGET2.4 generator with the hadronic shower for each neutrino recorded. This will be used to estimate the contribution of different hadronic interactions to the neutrino flux. The results for the contributions to the flux in this section are all presented for fluxes modelled at Kamioka for the Super-Kamiokande detector.

4.2.1 Relevant Hadron Production Species

To begin, we consider Figure 4.1 which shows a breakdown of each type of hadronic interaction component in the chain of hadronic interactions as a function of neutrino energy. From this we see the most important interactions to consider are:

- Pion production interactions
- Kaon production interactions
- Baryon-to-baryon interactions (e.g. interactions of the form $pN \rightarrow p + X$ and $pN \rightarrow n + X$)

It can be seen that of the meson production interactions, pion production reactions of the form $pN \rightarrow \pi^\pm + X$ are the dominant meson creation reactions across the full energy range with kaon production becoming increasingly relevant at higher energies. This is to be expected as kaons are more likely to be produced




Figure 4.1: Breakdown of hadronic interaction component to the atmospheric ν_μ (left) and ν_e (right) flux. Baryon to baryon refers to interactions with either a proton or neutron incident on a target nucleus where the outgoing particle is also a proton or neutron. Baryon to pion refers to any pion creation interaction with a proton or neutron primary particle. Similarly, baryon to kaon refers to kaon production interactions. Meson re-interaction is an interaction where a pion or kaon re-interacts within the atmosphere. Anti-baryon is any interaction featuring an anti-proton or antineutron and is found to be negligible. In each case, the contribution of each type of interaction is shown as a percentage of all hadronic interactions which contribute to the neutrino flux as a function of neutrino energy.

at higher energies and decay into higher energy neutrinos. The remainder of this chapter and Chapter 5 will focus in detail on the modelling of the pion production interactions. The tuning of the kaon interactions will follow a similar method and is discussed in Chapter 6.

Baryon-to-baryon interactions are governed by baryon number conservation. Such interactions only contribute to neutrino production when the daughter baryon later re-interacts to form a pion or kaon which decays. The only way to create additional baryons is via baryon anti-baryon pair production which is estimated to not contribute significantly to the neutrino flux (see Figure 4.1). For most of these interactions, the outgoing baryon is designated by the TARGET2.4 generator as the leading baryon (i.e the same baryon which initiated the interaction). Unlike the meson production interactions which can create multiple mesons, the number of

baryons before and after the interaction will be conserved.¹ As such the neutrino flux is expected to be less sensitive to the tuning of these interactions than pion and kaon production. Such interactions are discussed further in Section 9.4.1.

Finally there is a smaller component from meson re-interaction (e.g. interactions of the form $\pi^\pm N \rightarrow \pi^\pm + X$). Since the typical distance within which a meson decays in the atmosphere will be shorter than its interaction length, most mesons will decay before they have an opportunity to re-interact. Therefore meson re-interaction is only expected to have a small effect on the neutrino flux. Mesons with greater kinetic energy which decay to higher energy neutrinos have a greater decay length due to their Lorentz-dilated lifetime and are therefore more likely to interact before decaying than lower energy mesons. This is reflected in the increase in the meson re-interaction contribution with neutrino energy in Figure 4.1. At the neutrino energies we are interested in, these are estimated to only have a small contribution and their tuning to data is beyond the scope of this work. A systematic uncertainty estimate is discussed in Section 9.4.3.

4.2.2 Incident Particle Momentum in Meson Production Interactions

Having established the dominant hadronic interactions which contribute to neutrino production, this section considers the momentum of the incident particle in the interaction, this momentum is analogous to the beam momentum in fixed target experiments. Figure 4.2 shows the contribution of different parent momenta in pion and kaon production interactions to the atmospheric muon neutrino flux at three neutrino energies. It can be seen that the peak in the response curves is approximately where the parent momentum is an order of magnitude more than the neutrino momentum (this can serve as a helpful approximation). It should also be noted that the width of these peaks spans several decades in momentum and thus

¹A small component of these interactions will be due to nuclear fragmentation. These are interactions where the outgoing baryon is an ejected nucleon from the target nucleus. In this case, more than one baryon can be "produced" in an interaction. This is estimated to only affect the low-energy neutrino flux and to only have a small effect. It will be discussed further in Section 9.4.1.




Figure 4.2: For three different muon neutrino energies (from left to right 0.2 GeV, 1.1 GeV, 7.1 GeV) contribution to the atmospheric neutrino flux by the primary momentum of the hadronic interaction. Shown separately for pion and kaon production interactions. The y-axis shows the contribution to the neutrino flux in arbitrary units. The contributions for kaon production have been multiplied by the factors shown in the legends of each plot for greater visibility.

to calculate the flux at a single neutrino energy, interactions across a wide range of parent momentum must be considered. This is in direct contrast to beam neutrino experiments where a single proton energy is used and hadronic interactions only need to be understood in a much narrower range of energies (see for example Ref. [57, 58]).

4.2.3 Outgoing Pion Phase Spaces

Finally, it is necessary to consider the neutrino flux dependence on the kinematic phase space of the outgoing pion. The neutrino flux dependence on the momentum of the primary particle allows us to determine the relevant beam momenta to consider for fixed-target experiments. The contribution to the atmospheric flux of daughter mesons produced at specific a momentum and angle (relative to the incoming primary particle direction) is important when considering the acceptances of those fixed target experiments. A detector with a wider acceptance that more closely overlaps the relevant outgoing daughter particle phase space is able to better constrain the atmospheric flux. This is also useful when considering the design of future hadron production experiments to ensure that detectors are able to measure the cross section in regions most relevant to the atmospheric flux and further lower its uncertainty.




Figure 4.3: Contribution of hadron interactions from π^+ creation reactions of the form $pA \rightarrow \pi^+ + X$ to atmospheric ν_μ flux. Shown in the phase space of outgoing pion momentum p , production angle relative to incoming primary direction θ for ν_μ at four different neutrino energies in the 0.1 - 10 GeV region. The ν_μ contribution arises directly from π^+ decay. The relative contribution at each neutrino energy is in arbitrary units.

Considering this phase space is a multi-dimensional problem. The flux at each neutrino energy has a dependence on the momentum of the primary particle (Figure 4.2) and for each interaction at this momentum, there is a two-dimensional dependence on the daughter meson kinematics. Figure 4.3 shows the contribution to the ν_μ flux by π^+ momentum and production angle θ , integrated over all primary particle momenta. For lower neutrino energies there is a greater contribution from lower momentum pions as could be expected. There is a cut-off in pion momentum at each neutrino energy below which pions cannot decay to a neutrino of that energy due to kinematic limits. These lower energy neutrinos also have a greater dependence on pions produced at higher angles. These pions have lower energies due to receiving less energy from the Lorentz boost and are consequently more likely to decay to




Figure 4.4: Contribution of hadronic interactions to the atmospheric ν_μ flux at different neutrino zenith angles. Interaction phase space is shown for $pA \rightarrow \pi^+ + X$ reactions broken down by pion production angle θ relative to incoming proton direction (shown on the y-axis in mrad). This is shown at three different neutrino energies (left to right top to bottom 0.2 GeV, 1.1 GeV, 7.1 GeV). The contribution is in arbitrary units and normalised for each neutrino zenith and energy bin.

lower energy neutrinos. Consequently, experiments measuring daughter particles at higher angles need to be considered to model neutrino fluxes at lower energies.

Considering Neutrino Zenith Angles

As an aside, the neutrino energy bin can be further broken down by the zenith angle of the neutrino at the detector. Until now we have only considered the hadronic interaction phase space relevance to a single neutrino flux bin; however, the flux expressed in terms of neutrino zenith angle is important for neutrino oscillation analyses as discussed in Chapter 2. Figure 4.4 shows the contribution to the atmospheric flux by pions produced at angle θ relative to primary particle direction. This is shown for three different neutrino energy bins broken down into separate neutrino cosine zenith bins. This contribution is integrated over all incoming primary energies.

As with Figure 4.3, pions produced over a much broader range of angles contribute towards the lower energy neutrino flux. At higher neutrino energies, pions produced at high angles become less important and the 3D atmospheric showers tend towards the 1D approximation. It is seen that this effect is not uniform across the neutrino zenith angle. For a specific neutrino energy (and particularly noticeable at lower energies), the horizontal neutrinos ($\cos \theta_Z \sim 0$) come from pions produced over a far broader range of angles. Moving away from the horizon the more vertical incoming neutrinos are less dependent on pions produced at these larger angles. This is a consequence of the 3D geometric effect described in Section 3.2 where a pion produced at a greater angle with respect to the incoming primary is more likely to intersect the detector at the local horizontal. This dependence on a wider pion phase space will be important to consider when propagating uncertainties to the neutrino flux. It should be expected that the uncertainties on horizontal-going neutrinos to be greater due to increased dependence on larger angle regions of the interaction phase space which fewer experiments measure. Equally, while ratios between different neutrino fluxes often reduce systematic errors, any ratio comparing horizontal to vertical neutrinos will depend on different regions of hadron interaction phase space and as such should not be expected to cancel as significantly.

4.3 Fixed Target Hadron Production Experiments

Having explored the relevant hadron production phase spaces for estimating the neutrino flux between 0.1-10 GeV it is now necessary to consider the available datasets from fixed target experiments which cover this phase space. The experiments are chosen to maximise the overlap with the relevant phase space (both primary momentum and daughter particle phase space) to the neutrino flux. A summary of the different beam momenta of the datasets used in this analysis is shown in Figure 4.5.

Since a number of different experimental datasets will be used in this thesis to tune the atmospheric neutrino flux, it is helpful to begin with an overview of how




Figure 4.5: Figure showing the different beam momenta of the datasets used in this analysis. A summary of each of the datasets is included in Section 4.3. References for each dataset are included in Section 4.3 and Table 4.1.

hadron production is measured experimentally. This section will start by considering concepts common to all measurements before looking at the specific experimental setup and analysis methods used by each of the experiments. Understanding how the different measurements are made will aid our understanding of how to tune the atmospheric flux to these datasets and propagate uncertainties.

A fixed target hadron production experiment aims a particle beam at a fixed nuclear target. While a variety of different beams can be used to investigate hadronic physics, for the datasets used to tune the neutrino flux this particle beam will be a proton beam. Interactions between the proton and the target produce daughter hadrons. The number of hadrons produced downstream of the target in two-dimensional kinematic phase space bins is measured. Measuring this number of hadrons alongside the number of incident beam protons it is possible to measure the inclusive production cross section.

Detectors upstream of the target establish the species of particle in the incident proton beam (often secondary beam lines are used with additional particles alongside the protons) and the position and direction of the particles in the beam (see the

inset in Figure 4.6 for an example). These come together to define a beam trigger which records when a proton with desired kinematics is incident on the target. Downstream of the target is a spectrometer which identifies the downstream particle species (either by energy loss in a tracking detector, specific PID sub-detectors or a combination of both) and its kinematics (typically the angle of the particle relative to the beam and the momentum of the particle). The fixed target experiments considered here use thin nuclear targets (typically around 5% of the interaction length) to minimise the probability of multiple interactions within the target and improve the estimate of the cross section on a single nucleus.

4.3.1 Measuring the Double Differential Cross Section

This subsection will adopt the formalism of the NA61/SHINE experiment[56]; however, it is broadly applicable to all hadron production measurements. Where there are significant deviations from this method they will be expanded upon below when discussing individual experiments. Unfortunately, the available hadron production datasets are not presented in an entirely consistent manner within the literature with different binning schemes used and different analysis corrections made. In addition, not all datasets provide complete uncertainty estimates (discussed further in Section 5.4.1). Each experiment will be considered in turn which will provide a basis for the comparisons between datasets in Chapter 5.

Interactions in the target can be considered elastic (where the beam particle coherently scatters off the whole target nuclei) or inelastic (where the target nucleus disintegrates). The NA61 formalism separates inelastic collisions further into quasi-elastic interactions and production interactions [56]. A quasi-elastic interaction results in the ejection of additional nucleons from the nucleus whereas a production interaction results in the creation of new hadrons. Thus within this formalism:

$$\sigma_{\text{inelastic}} = \sigma_{\text{total}} - \sigma_{\text{elastic}} \quad (4.2)$$

$$\sigma_{\text{production}} = \sigma_{\text{inelastic}} - \sigma_{\text{quasi-elastic}} \quad (4.3)$$

It is these production interactions that are of specific interest to tuning pion and kaon production within the atmospheric neutrino flux, as by definition all interactions involving the production of a meson must be a production interaction.

As an initial estimate the double differential cross section (equation 4.1) can be experimentally measured in two-dimensional outgoing hadron kinematic bins as [56]:

$$\frac{d^2\sigma_h}{dpd\theta} = \frac{A}{N_A l \rho} \frac{1}{\Delta p \Delta \theta} \frac{\Delta n_h(p, \theta)}{N_{\text{beam}}} \quad (4.4)$$

where

- σ_h is the production cross section of hadron h
- A is the atomic mass number of the target material
- N_A is Avogadro's number, l the target thickness
- ρ the target density
- Δp and $\Delta \theta$ the width of the momentum and angular phase space bin
- Δn_h the number of daughter particles (e.g. pions, kaons) in that bin
- N_{beam} the number of beam protons incident on the target.

Corrections are then applied to Δn_h the raw number of hadrons measured to account for detector and cut inefficiencies (typically done via correction factors from Monte Carlo simulations). Some of the experiments also make corrections to the target length l to account for beam attenuation across the target. Corrections are also made to subtract the empty target background by performing the same experiment without the target in place to estimate the number of mesons being produced in interactions of the proton beam outside the target. Since out-of-target interactions can occur both upstream and downstream of the target, the correction is not as simple as just subtracting one number from the other (see Ref. [56] for the full approach to this).

It can be seen from equation 4.4 that the uncertainty on the normalisation of the double differential cross section measurements are related to understanding the

number of beam protons N_{beam} , target thickness l and density ρ . This normalisation uncertainty is common to the production of all mesons regardless of species. The bin-to-bin uncertainties are related to understanding Δn_h , the statistical uncertainty relating to the number of mesons detected and systematic uncertainty to the corrections and reconstruction processes that have been applied.

At higher interaction energies the production of short-lived neutral strange particles such as the neutral lambda baryon (Λ^0) and short lived neutral kaon (K_s^0) meson have significant production cross sections. These mesons can decay to pions inside the spectrometer. The NA61/SHINE [56] and NA49 [59] datasets correct their pion production measurements to remove pions that come from the weak decays of such particles from being included in the pion production cross section, this correction is referred to as a feed-down correction of neutral particle decays. For the atmospheric neutrino flux prediction, it does not matter whether a pion originates in the initial hadronic interaction or from the feed-down decay of a secondary strange particle; only total pion production is of interest. However, it is necessary to understand this effect to compare the equivalent cross section between data and Monte Carlo. I.e. if the cross section for pion production in Monte Carlo includes those produced in feed-down decays but the cross section in the datasets does not, then the re-weighting methods discussed in Chapter 7 will be invalid. The Λ^0 contribution to the neutrino flux is discussed in Section 9.4.2.

The following sub-sections summarise each of the experiments which make the hadron production measurements used in this thesis to estimate the atmospheric neutrino flux. In each case, the information is summarised from the relevant cited publications.

4.3.2 NA61/SHINE

The NA61/SHINE [56] experimental setup is shown in Figure 4.6. NA61/SHINE has a wide physics program, taking data with a combination of different beams and targets. The datasets most relevant to the atmospheric analysis are the production of pions and kaons in proton carbon interactions ($pC \rightarrow \pi^\pm/K^\pm + X$). Specific




Figure 4.6: Diagram showing the NA61/SHINE spectrometer alongside the detector setup in the beamline. This setup was used to take the data at 31 GeV beam momentum used to tune the atmospheric neutrino flux. A summary of subsystems is provided in the main text. Image is taken from [56]

focus will be given to the experimental setup used to take this data. Upstream of the target the beam particles are identified and counted by two Čerenkov detectors, the differential Čerenkov detector (CEDAR) and threshold Čerenkov counter (THC) which identify the proton component of the beam. Scintillation counters (S0 and S1) and veto counters (V0, V1 and $V1^p$) provide the beam definition and constrain the beam position and angles that are incident on the target. These scintillation and veto counters together with the THC and CEDAR form the beam trigger. The three beam position detectors (BPD) are proportional chambers which measure the incoming beam trajectories.

The target is graphite with a length of 4% of the nuclear interaction length and data is taken using 31 GeV momentum incident protons. Downstream of the target the S4 scintillation counter is used to determine an interaction trigger with the absence of a signal in S4 coincident with a beam trigger indicating that an interaction has taken place. Essentially S4 is included in order to remove the large number of non-interacting through-going protons from the analysis and ensure production interactions are mainly read out. Corrections are applied for those elastic and quasi-elastic interactions which miss the S4 counter and thus create

an interaction trigger. There will also be some production interactions where the charged daughter particle hits the S4 counter and there is no interaction trigger [60]. This will limit the measurement of interaction daughter particles at small angles.

Immediately downstream of the target are two vertex time projection chambers (VTPC1 and VTPC2) placed within a dipole magnetic field of bending power 1.14 Tm followed by two main TPCs (MTPC-L and MTPC-R) operated with a Ar : CO₂ gas mixture. The magnetic field the VTPCs sit in serves to reduce track density and provide a measurement of particle momentum. Further downstream is a time of flight (TOF) wall which consists of scintillator bars read out by photomultipliers. It provides a measurement of the length of time it takes a daughter particle to reach this point in the spectrometer.

Particle identification is performed by both time-of-flight measurements and $\frac{dE}{dx}$ energy loss of the particle in the TPCs. At lower momentum only $\frac{dE}{dx}$ was used due to a large number of the lower momentum mesons decaying before reaching the TOF wall. Above 1 GeV daughter momentum, a combined TOF $\frac{dE}{dx}$ approach was taken to reduce uncertainty where $\frac{dE}{dx}$ bands overlap for different hadron species.

The normalisation uncertainty is estimated to be 1% and dominated by the uncertainty on the trigger cross section driven by the uncertainty on the number of protons incident on the target. Individual bins have separately quoted systematic and statistical uncertainties with the systematic uncertainties due to the particle identification and track selection processes alongside the corrections applied to the bins which account for the detector acceptance. At lower daughter particle momentum systematic uncertainties typically dominate the individual bin uncertainties whereas for higher momentum statistical uncertainties dominate. The differential cross sections is given in terms of meson momentum p and production angle θ as $\frac{d^2\sigma}{d\theta dp}$ and binned in θ and p .

4.3.3 NA49

A schematic view of the NA49 experiment is shown in Figure 4.7. NA49 [59] uses a 158 GeV momentum beam incident on a graphite target of thickness 1.5% of the




Figure 4.7: Diagram showing the NA49 spectrometer and the detector setup in the beamline. This setup was used to take the data at 158 GeV beam momentum. As can be seen, it is highly similar to the NA61/SHINE apparatus in Figure 4.6. Image is taken from [61].

interaction length. NA49 was an earlier experiment than NA61/SHINE and much of the detector setup from NA49 was reused by the NA61/SHINE experiment which also added additional sub-detectors such as the TOF system. The collaboration working on the measurements was mostly different and there has been no suggestion by either collaboration that their uncertainties should be treated as correlated. For the atmospheric analysis, these measurements will be treated as independent. The experimental setup of the beamline, vertex TPCs and main TPCs are the same in NA61/SHINE and NA49 and they will not be described again here. The magnetic bending power of the magnets is larger than NA61/SHINE at 9 Tm to account for higher momentum daughter particles. Unlike NA61/SHINE there are no TOF

detectors and PID is performed exclusively by $\frac{dE}{dx}$.

Uncertainties on the normalisation procedure are estimated to be at 2.5% and systematic uncertainties on the individual bins are estimated at 3.8% for each bin. Bin-by-bin statistical errors are also given and dominate in the higher momentum bins. Bin-by-bin systematic uncertainties are dominated by the feed-down corrections (correcting for neutral strange particles which later decay to pions and kaons in the detector) and bias in the interaction trigger where inelastic interactions trigger the S4 counter and are rejected by the interaction trigger. Cross sections are presented as $E \frac{d^3\sigma}{d\mathbf{p}^3}$ and binned in x_F (Feynman-x) and P_T . x_F is defined as the ratio between the longitudinal momentum of the daughter particle in the centre of mass frame to the maximum possible momentum of that particle kinematically allowed [62]. The cross section $E \frac{d^3\sigma}{d\mathbf{p}^3}$ is known as the Lorentz invariant cross section as is invariant under the Lorentz boost. In Chapter 5 these variables will be shown to be particularly useful to describe hadron production within atmospheric showers. The use of these variables is discussed in more detail in Section 5.3.1, where Feynman scaling of hadron production cross sections will be explored.

4.3.4 HARP

An overview of the HARP spectrometer is shown in Figure 4.8. HARP [63–65] took data on a variety of different nuclear targets with beam momenta from 3 - 12 GeV. The beryllium, carbon and aluminium target data will be used for the atmospheric analysis. Each target had a thickness of 5% their interaction length. HARP, unique to all the experiments considered in this section, has been designed with acceptance to cover the full 4π phase space of the daughter particles and achieves this by placing the target within the inner field cage of a TPC which itself is positioned inside a solenoidal magnet of strength 0.7 T. In addition to this, there is also a forward spectrometer consisting of five drift chambers to track the daughter particles with one upstream and one immediately downstream of a dipole magnet used to measure momentum. A TOF wall, threshold Čerenkov detector and electromagnetic calorimeter are used for particle identification. The forward




Figure 4.8: Diagram showing the HARP spectrometer used to measure production cross sections from 3- 12 GeV on a variety of targets. The target is placed inside the TPC. For the low angle data, only the forward spectrometer (this excludes the TPC) was used in the analysis whereas for the high angle data only the TPC is used without the forward spectrometer. Details of the subsystems are discussed further in the main text. Figure taken from [63].

spectrometer has acceptance up to polar angle 250 mrad. A resistive plate chamber (RPC) outside the TPC is used to make TOF measurements to aid with particle identification in the high-angle region not covered by the forward spectrometer. Upstream of the target, scintillation triggers ensure only beam particles incident on the target register as a beam trigger and Halo counters (scintillators with a hole at the beam axis) act similar to the NA61/SHINE veto to reject particles too far from the beam axis. Additionally, four multiwire proportional chambers measure the position and direction of the incoming beam.

Due to hardware effects with the TPC discovered after data taking, a number of unexpected effects needed to be accounted for in the analysis. Unfortunately, the

experiment was not given the opportunity to repeat data collection. Inhomogeneity of the TPC's electric field together with a number of dead TPC readout pads lead to difficulty in reconstructing particle trajectories in the TPC. As such, the measurement of daughter particles at small production angles (<250 mrad), which used the forward spectrometer, did not take input from the TPC. The high production angle data, which used only the TPC and RPC, was analysed separately from the small angle data. In both analyses, the spectra are normalised to the total number of incident particles on the target recorded by the beam trigger system described above. The normalisation uncertainty is estimated to be 2%. For both the low and high angle data the results are given as $\frac{d\sigma}{dpd\Omega}$ binned in θ and p .

Low Angle HARP Data <250 mrad

Momentum is calculated from the deflection in the track measured in the drift chambers due to the dipole magnet. For particle identification, the TOF wall is used below 3 GeV and the threshold Čerenkov above 3 GeV, the electromagnetic calorimeter is also used to separate electrons and pions. Runs are taken without the target and the background from out-of-target interactions is subtracted as in other experiments. No corrections are made for pion production from neutral particle decay (the feed-down correction); however, this is estimated to be small at the interaction energies measurements were made for. Systematic and statistical uncertainties are given for each bin.

High Angle HARP Data >250 mrad

The high-angle data uses only the TPC and RPC. Momentum was reconstructed from tracks in the TPC (within the solenoid magnetic field) and particle identification was performed using $\frac{dE}{dx}$ energy loss in the TPC which were calibrated by the TOF measurements made by the RPC. Build-up of ion-charge density during a beam spill caused distortions to reconstructed track trajectories and so a cut was placed to only consider events during the beginning of each spill before the charge density had built up. Errors are quoted as systematic and statistical for each bin with the dominant




Figure 4.9: The E910 Spectrometer. Subsystems are described in the main text. Figure taken from [67]

error source coming from the absolute scale of the momentum measurement. There are two published analyses of the HARP high angle data and they differ significantly from each other [65, 66]. The discrepancies between these datasets form the greatest uncertainty in this region and its effect on the flux is addressed in Section 9.1.3.

4.3.5 E910

The E910 experiment [67, 68] measured pion production double differential cross section spectra with an incident proton beam of momenta 6.4, 12.3 and 17.5 GeV on a beryllium target. The beryllium target has a thickness of 4.5% of its interaction length. The E910 spectrometer is shown in Figure 4.9. Upstream of the target (not shown in Figure 4.9) are three Čerenkov detectors to identify protons in the beam. Two trigger counters and veto counters are used to define the beam trigger and limit the trajectories of protons on the target. Pairs of scintillator slats downstream of the target (known as the "bullseye" counter) function to detect non-interacting protons in much the same way as the S4 counter in NA61/SHINE and NA49. A beam trigger coincident with the absence of a signal in the bullseye counter defines the event trigger. Upstream of the target, a proportional chamber determines the trajectory of the incoming protons.

The daughter particles are detected in the spectrometer. Particle identification is completed using $\frac{dE}{dx}$ energy loss in the TPC, threshold Čerenkov counters (CKOV counter) and a TOF wall. For pions from 0.4 - 1.2 GeV, $\frac{dE}{dx}$ is the primary PID,




Figure 4.10: The E802 Spectrometer. The cross-hatched area refers to the region where the magnetic field is present. The spectrometer is capable of rotating to cover a wider acceptance and in this diagram is at 5° from the beam axis. Subsystems are described in the main text. Figure taken from [69]

above 1.2 GeV TOF becomes the primary PID with pions also producing Čerenkov light above 2.8 GeV that further aids PID. The TPC and drift chambers DC1-DC3 are inside a 0.5T magnetic field and used for track reconstruction and momentum determination. Uncertainties are quoted as a single value for each bin which combines bin statistics and detector systematic effects. Systematic uncertainties arise from the particle identification process and the efficiency of the bullseye trigger where inelastic interactions are counted as non-interacting protons. Additionally, there is a 5% normalisation uncertainty from the beam trigger and uncertainties on the target thickness. Cross sections are given in $\frac{d^2\sigma}{dpd\Omega}$ and binned in p and θ .

4.3.6 E802

E802 [69] measured production cross sections at 14.6 GeV beam energy on Be, Al, Cu and Au targets. The beryllium and aluminium data will be used in the atmospheric analysis since they are closest in mass number to oxygen and nitrogen nuclei in the atmosphere. The experiment was designed to measure these spectra at higher daughter particle angles. Figure 4.10 shows the E802 spectrometer layout. To achieve this, the spectrometer can be rotated from 5° - 58° relative to the beamline. Four drift chambers are used for tracking, two upstream and two




Figure 4.11: Single arm magnetic spectrometer used for production spectra measurements by Allaby et al. and Eichten et al. Details of the measurements in the main text. Figure taken from [70]

downstream of a dipole magnet (the “Henry Higgins Magnet” in Figure 4.10) of bending power 1.5 Tm. Deflection in the magnet is used to determine a daughter particle’s momentum. Particle identification is completed by a TOF wall and Čerenkov counter with the TOF used for lower momenta particles and Čerenkov counter for higher momenta. The production measurement is given as $\frac{d^2\sigma}{dm_T dy}$ (double differential cross section with respect to transverse mass, m_T and rapidity, y) and is binned in rapidity and transverse kinetic energy. Transverse kinetic energy is defined as the rest mass subtracted from the transverse mass, $m_T - m_0$, where the transverse mass is defined as $m_T = \sqrt{m_0^2 + P_T^2}$ in natural units. Bin-to-bin systematic uncertainties are not quoted, only statistical uncertainties given. An overall normalisation uncertainty of 15% is given due to the procedure of normalising the particle yield to the number of protons on target.

4.3.7 Eichten et al. and Allaby et al.

The measurements by Eichten et al. (1972) [71] and Allaby et al. (1970) [70] use a single-arm magnetic spectrometer to determine the momentum and angle of the daughter particles as opposed to the tracking chambers that have been used by the experiments considered so far. The main difference is that the differential production

cross section is measured at discrete points in daughter particle phase space chosen by the spectrometer settings rather than continuously and binned afterwards. Both experiments use the same spectrometer and the experimental apparatus and analysis is highly similar. The experimental setup is shown in Figure 4.11.

Magnets SM1 and SM2 can be moved such that the spectrometer accepts a specific daughter particle production angle. Particle identification is performed by Čerenkov counters C_1 , C_2 and C_3 along with a differential Čerenkov counter (DISC). Dipole magnets M1-M3 and quadrupole magnets Q_i provide a dispersive focus of the target at F_1 . Scintillator plates (A_i) are moved through this dispersive focus with each position corresponding to a reconstructed momentum. Thus the spectrometer can be set to measure the number of particles at a specific momentum and specific angle. Momentum calibration is performed using elastic proton-proton scattering (using a hydrogen target). The incoming proton flux is determined by a secondary emission chamber upstream of the target.

Corrections to the final spectra are applied. These include subtracting the empty target background, corrections due to the decay of pions/kaons inside the spectrometer and the absorption of pions and kaons inside the detector. Allaby provides the spectra at specific angles and momenta of daughter particles as $\frac{d^2\sigma}{dpd\Omega}$, Eichten provides a Lorentz invariant multiplicity defined as $\frac{2E}{p^2\sigma_a} \frac{d^2\sigma}{dpd\Omega}$ where the double differential cross section is divided by the total inelastic cross section (σ_a) measured by Bellettini et al. [72]. Both are provided as a function of radial angle θ and momentum.

The Allaby measurement provides bin-to-bin uncertainties as a combination of statistical and systematic uncertainties. An overall normalisation uncertainty of 15% is applied due to the calibration of the secondary emission chamber monitoring the beam flux. The reproducibility of the results when repeating the experiment after a long time interval was found to be 10%. For the atmospheric analysis in this thesis, the combined reproducibility and bin-to-bin uncertainty was used as a conservative estimate of the total bin-to-bin uncertainty.




Figure 4.12: The NA56 experimental setup, showing the target and downstream detectors described in the main text. Figure taken from [73]

The Eichten measurement gives a bin-by-bin uncertainty based on a combination of the systematic uncertainties during the application of corrections (5% for each bin) and the irreproducibility of the spectra upon repeating the measurement (also 5% for each bin). The statistical uncertainty is stated to be negligible when compared to the systematic uncertainty and not given. For the atmospheric analysis in this thesis, we take the combination of these two uncertainties as an upper limit for the bin-to-bin uncertainties. Additionally, there is also a 15% normalisation uncertainty from the same source as the Allaby measurement.

4.3.8 NA56

The NA56 experiment [73, 74] measures pion and kaon yield for interactions of a 450 GeV proton beam incident on a beryllium target. The experimental setup is shown in Figure 4.12. Secondary emission monitors around the target (not shown in Figure 4.12) measure the proton beam intensity. Three sets of motorised collimators downstream of the target define the angular and momentum acceptance of the daughter particles. Particle identification is performed using a combination of the TOF detectors, threshold Čerenkov, CEDAR detector (another type of Čerenkov detector) and hadronic calorimeter. The daughter particles can be tracked by the five wire chambers (WiT in Figure 4.12). The data is provided in the form of pion and kaon yields for different beryllium target thicknesses without correction for beam attenuation across the target or feed-down decays of short-lived neutral

Experiment	p/GeV	$\sqrt{s}/\text{GeV}$	Target	Inclusive Interaction	Unit
HARP [64–66]	3, 8, 5, 12	2.77, 3.36, 4.11, 4.93	Be,C,Al	$pN \rightarrow \pi^\pm + X$	$\frac{d\sigma}{dpd\Omega}$
E910 [67]	6.4, 12.3, 17.5	3.72, 4.99, 5.86	Be	$pN \rightarrow \pi^\pm + X$	$\frac{d^2\sigma}{dpd\Omega}$
E802 [69]	14.6	5.40	Be Al	$pN \rightarrow \pi^\pm + X$ $pN \rightarrow K^\pm + X$	$\frac{d^2\sigma}{dm_T dy}$
Eichten et al. [71]	24	6.84	Be Al	$pN \rightarrow \pi^\pm + X$	$\frac{2E}{p^2\sigma_a} \frac{d^2\sigma}{dpd\Omega}$
Allaby et al. [70]	19.2	6.15	Be Al	$pN \rightarrow \pi^\pm + X$ $pN \rightarrow K^\pm + X$	$\frac{d^2\sigma}{dpd\Omega}$
NA61/SHINE [56]	31	7.74	C	$pN \rightarrow \pi^\pm + X$ $pN \rightarrow K^\pm + X$	$\frac{d^2\sigma}{d\theta dp}$
NA49 [59, 76]	158	17.27	C	$pN \rightarrow \pi^\pm + X$ $pN \rightarrow K^\pm + X$	$E \frac{d^3\sigma}{d\mathbf{p}^3}$
NA56 [74]	450	29.01	Be	$pN \rightarrow \pi^\pm + X$	$E \frac{d^3\sigma}{d\mathbf{p}^3}$

Table 4.1: Summary of the hadron production datasets used within this analysis to tune the atmospheric neutrino flux. Shown for each experiment are the beam momentum and $\sqrt{s}$ at which the data are taken, the target materials, the inclusive production cross section measurements which are used and the form that double differential cross section is given as.

particles. A separate study using this data [75] derived the inclusive production differential cross sections with attenuation and feed-down corrections, this data will be used for this analysis. Statistical and Systematic uncertainties are provided separately for each bin. Additionally, there is a 10% normalisation uncertainty.

4.4 Summary

The datasets used within the pion and kaon production modelling in this thesis are summarised in Table 4.1.

5

Modelling Pion Production in Air Showers

The goals of this chapter are to develop a model for the inclusive pion production cross sections relevant to the atmospheric flux based on the datasets from Chapter 4. The methods developed in this chapter will then also be applied to kaon production in Chapter 6. An empirical target nucleus scaling is developed and used to interpolate between the differential production cross section measurements taken on different target nuclei. The use of parameterisations to the datasets is explored which, alongside the target nucleus scaling, is used to compare different dataset measurements of the production cross section. The extent to which the production cross section scales with interaction energy is explored. Lower energy datasets at similar interaction energies are found to be in disagreement with each other and a method for estimating the uncertainties in this region is proposed. Finally, the extent to which the datasets cover the phase space of pion production for the neutrino flux calculation is shown.

5.1 Using Fixed Target Data to Model Atmospheric Showers

Having reviewed the available datasets alongside the hadronic interactions in the atmospheric showers, there are several key differences that need to be addressed before the datasets can be applied to the model.

- Firstly the fixed target datasets exclusively use proton beams. While the

cosmic ray flux is mostly comprised of free protons, approximately 5% are light nuclei [36] (see Chapter 10). It is therefore necessary to model particle production for interactions with incident neutrons from these nuclei.¹ Despite the lack of hadron production data from neutron interactions, differential cross sections for pion production with incident neutrons can be obtained from the incident proton data using isospin symmetry [77]. This symmetry implies that for inclusive differential cross sections:

$$\sigma(pN \rightarrow \pi^\pm + X) = \sigma(nN \rightarrow \pi^\mp + X), \quad (5.1)$$

if N is an isoscalar nucleus, such as carbon, oxygen and nitrogen. This allows results from fixed target experiments using proton beams to be used to model neutron interactions in the atmosphere.

- In the fixed target experiments the beam is incident on solid targets such as carbon, aluminium and beryllium whereas in atmospheric showers the primary particles are incident on oxygen and nitrogen in the atmospheric gases. In order to apply the data to model atmospheric showers, a scaling between cross sections on different targets needs to be developed. A full discussion on how differential cross sections scale between interactions on different target masses is included in Section 5.2.
- In Section 4.2.2 it was seen that to model the atmospheric neutrino flux at a specific energy, hadronic interactions over a wide centre of mass energy range need to be understood. Therefore data taken at different beam momenta needs to be compared. As the datasets each use a slightly measure for double differential cross section (for example $\frac{d^2\sigma}{dpd\theta}, \frac{d^2\sigma}{dx_F dP_T}$) with different histogram binning schemes, a direct comparison is typically not possible. To allow for comparison of these datasets and to model the hadronic interactions across the

¹It will be assumed that the multiplicity of meson production in an interaction of a bound proton or neutron on air is the same as that of their free counterparts. This follows the superposition approximation, described in section 3.4, which has previously been used to model the light nuclei interactions within air showers [46].

Experiment	Target Nuclei	Beam Momentum /GeV
HARP	Be C Al	3, 5, 8, 12
E802	Be Al	14.6
Allaby	Be Al	19.2
Eichten	Be Al	24

Table 5.1: Summary of the pion production datasets used in the fitting of the target nucleus scaling parameters. In each case, the ratio between the cross section on the target nucleus is taken to the beryllium dataset and used for the fit.

whole phase space a parameterisation of the double differential cross section is used in Section 5.3.

5.2 Target Nucleus Scaling

A simple approach to modelling the scaling of the production cross section between different nuclei could be to model the nucleus as a sphere comprised of fixed volume nucleons. If the interaction cross section is proportional to the physical cross section of this spherical nucleus, we obtain that the interaction cross section is proportional to $A^{\frac{2}{3}}$ where A is the number of nucleons in the nucleus [75]. This is the dominant effect of changing the target nucleus size in a fixed target interaction, the larger the nucleus the greater the probability an interaction occurs. This simple proportionality alone does not account for the subdominant effects of differences in the outgoing pion kinematics with different target nuclei.

To model the effect of nuclei of different masses on the differential pion production cross section and therefore allow us to interpolate cross sections on beryllium, carbon and aluminium targets to oxygen and nitrogen, an empirical parameterisation in Feynman x (x_F) and transverse momentum (P_T) of the outgoing pion is used as an extension of the $A^{\frac{2}{3}}$ scaling. This is able to account for variations in the outgoing pion kinematic distribution with different target nuclei. This parameterisation has previously been used to scale between hadron production on different target materials [78]. The use of the x_F and P_T variables to describe hadron production

Parameter	π^+	π^-
p_0	0.687	0.804
p_1	-0.381	-0.556
p_2	0.251	-0.135
p_3	-0.027	0.033
p_4	-0.127	0.149
p_5	0.115	0.320
$\chi^2/\text{d.o.f}$	463.0/491	404.9/436

Table 5.2: Results for the fitting of the target nucleus scaling parameterisation to the datasets described in the text.

is discussed further in Section 5.3.1. This target nucleus (sometimes referred to as “A scaling”) parameterisation is as follows:

$$\frac{d^2\sigma(pN_1 \rightarrow \pi^+ + X)/dx_F dP_T}{d^2\sigma(pN_2 \rightarrow \pi^+ + X)/dx_F dP_T} = \left(\frac{A_1}{A_2}\right)^{\alpha(x_F, P_T)} \quad (5.2)$$

$$\alpha(x_F, P_T) = p_0 + p_1 x_F + p_2 P_T + p_3 x_F^2 + p_4 P_T^2 + p_5 x_F P_T.$$

In this parameterization, A_i is the atomic number of the nucleus N_i and p_j are parameters to be fit. This parameterisation was used to fit ratios between cross sections on aluminium and carbon from the datasets listed in Table 5.1. It is assumed that the fit parameters will not vary significantly with beam energy across the energies listed here. A reduced χ^2 fit was performed taking the uncertainty on the ratio in each bin as the quadrature sum of the uncertainty on the beryllium and heavier nuclei pion production double differential cross section. The values of the parameters obtained are shown in Table 5.2 separately for both π^- and π^+ daughters. Figure 5.1 shows a comparison of the ratio of the pion production cross section on aluminium to beryllium with the Allaby measurement to this parameterisation. As can be seen by the data this ratio is about $1.9 \approx \left(\frac{A_{\text{Al}}}{A_{\text{Be}}}\right)^{\frac{2}{3}}$ with differences across the x_F - P_T space of the produced pion described well by the parameterisation.

Figure 5.2 shows the deviation of this parametrised fit from the simple $A^{\frac{2}{3}}$ scaling across x_F - P_T phase space for scaling between carbon and nitrogen cross sections. An enhancement of the cross section in the central region around $x_F \approx 0$




Figure 5.1: The result of the A Scaling Parameterisation fitted to the datasets in Table 5.1 shown alongside the ratio between pion production cross sections on aluminium and beryllium measured by Allaby et al. [70]. Shown for the production of π^+ (top) and π^- (bottom).




Figure 5.2: Deviation of A Scaling Parameterisation for carbon to nitrogen from simple $A^{2/3}$ scaling. Plotted is the A Scaling parameterisation in equation 5.2 for nitrogen to carbon divided by the simple $A^{2/3}$ scaling.

(the region corresponding to pion production at higher angles) is seen alongside a decrease in the cross section in the low angle forward region (high x_F , low P_T). The interaction at a nucleon level is the same for pion production on both carbon and nitrogen nuclei, the proton either interacts with a proton or a neutron inside the nucleus. Once the pion is created it must exit the nucleus, this secondary interaction between the pion and the remaining nucleons has the effect of scattering the pion to higher angles. Since the nitrogen nucleus has a greater number of nucleons than carbon the scattering effect is greater as seen in Figure 5.2.

This parameterisation has demonstrated a method for interpolating cross sections on solid nuclear targets to nitrogen and oxygen for modelling hadronic interactions in atmospheric showers. This parameterisation is shown to be effective over a wider range of atomic masses (beryllium to aluminium) than is typically needed to interpolate to nitrogen and oxygen for the atmospheric flux. Parameters in Table 5.2 are obtained for fits to datasets between 3 and 24 GeV. In addition to these energies, these parameters will also be used for scaling the 31 GeV NA61 data.

The scaling of the cross section with target material at higher energies has previously been studied with a similar empirical parameterisation in x_F and P_T . Within this parameterisation $\alpha(x_F, P_T)$ in equation 5.2 is given as [75, 79]:

$$\alpha(x_F, p_T) = (0.74 - 0.55x_F + 0.26x_F^2)(0.98 + 0.21p_T^2). \quad (5.3)$$

This parameterisation will be used in place of the lower energy parameterisation for the NA49 (158 GeV) and NA56 (450 GeV) datasets. The effect of the target nucleus scaling on the neutrino flux uncertainty is discussed in Section 9.3.

5.3 Dataset Parameterisations

5.3.1 Kinematic Variables and Feynman Scaling

As seen in Table 4.1, there are many potential variables with which the differential cross section for meson production can be expressed. A convenient choice for this analysis is the Lorentz invariant differential cross section, which was first seen in the discussion of the NA49 data in Section 4.3.3. It is defined as [80]:

$$E \frac{d^3\sigma}{d^3\mathbf{p}}, \quad (5.4)$$

where E is the energy and $\mathbf{p}$ is the momentum of the outgoing meson. It presents a convenient choice as it describes the hadron production invariantly without the need for consideration of a specific frame. In a collision with an unpolarised beam particle incident along the z-axis with cylindrical symmetry, it can be shown that:

$$E \frac{d^3\sigma}{d^3\mathbf{p}} \equiv \frac{E}{p^2} \frac{d^2\sigma}{dp d\Omega} \equiv \frac{E}{2\pi p^2 \sin\theta} \frac{d^2\sigma}{dp d\theta}, \quad (5.5)$$

where p is the absolute value for the momentum of the outgoing meson and θ the angle between the outgoing meson and the z-axis. This allows us to obtain the invariant cross section from the datasets in Table 4.1.

Hadron production experiments express the cross section as a function of the kinematics of the outgoing meson. There are a number of different choices for which kinematic variables to use. Of particular relevance to this analysis are the Feynman scaling variable x_F and scaling variable x_R . Although x_F was previously defined in Section 4.3.3 and used in Section 5.2, its definition will be repeated here in anticipation of the discussion of Feynman scaling below. x_F and x_R are defined as [62, 81]:

$$x_F = \frac{P_L^*}{P_{\text{Max}}^*} \quad (5.6)$$

$$x_R = \frac{E^*}{E_{\text{Max}}^*} \quad (5.7)$$

where E^* and P_L^* are the energy and longitudinal momentum of the outgoing meson in the centre of mass frame respectively. E_{Max}^* and P_{Max}^* are the maximum energy and momentum for that particle, also in the centre of mass frame.

Feynman suggested, from theoretical motivation, that the inclusive differential cross section could be approximately described by x_F and P_T without an additional variable to take the centre of mass energy of the interaction into account [62]. In other words, the scaling of the double differential cross section with the centre of mass energy is fully encapsulated in the Feynman scaling variable x_F . Following Feynman's scaling hypothesis the use of the x_R variable as an alternative to x_F was suggested by Yen [82] and Taylor [81] and demonstrated to extend the scaling to a wider energy range. This scaling is particularly useful in this analysis as it provides a method to compare datasets across different beam momenta. While it will be shown that this approximation does not hold well at low interaction energies (see Section 5.4.2) an approximate energy scaling is nonetheless useful.

x_F - P_T Phase Space

Throughout the remainder of this section, the differential production cross section will be shown in x_F - P_T phase space of the outgoing meson due to this approximate energy scaling. It is sometimes more intuitive to consider the phase space in terms of the momentum and polar angle of the outgoing meson in the production interaction. Figure 5.3 demonstrates the relationship between x_F - P_T and p - θ phase space. It is seen that the low θ region refers to positive x_F at low values of P_T while the negative x_F region (where pions are travelling backwards in the centre of mass frame) corresponds to high angles. This will be particularly useful when considering interactions occurring outside experimental datasets in Chapter 8.




Figure 5.3: Lines of constant polar angle θ and outgoing pion momentum p in x_F - P_T phase space at 8,31 and 158 GeV primary (beam) momenta.

5.3.2 BMPT Parameterisation

The BMPT parameterisation is a suggested parameterisation of the invariant differential cross section as a function of P_T and x_R [75]. It takes the following form:

$$E \frac{d^3\sigma}{dp^3} = A(1 - x_R)^\alpha (1 + Bx_R)x_R^{-\beta} \times \left[1 + \frac{a}{x_R^\gamma} p_T + \frac{a^2}{2x_R^\delta} p_T^2 \right] e^{-\frac{ap_T}{x_R^\gamma}}. \quad (5.8)$$

The parameterisation is motivated by a combination of theoretical expectations within the quark-parton model and observed experimental behaviour of soft hadronic interactions. Parameters A , B , a , α , β , γ and δ are determined by fits to the datasets. An additional factor of $r_0(1 + x_R)^{r_1}$ included for π^- and $r_0(1 - x_R)^{r_1}$ included for K^- production, with r_0 and r_1 parameters determined in the fit. These factors were originally intended for parametrising the ratio between positive and negative

meson production however its inclusion as an additional factor to equation 5.8 was found to improve the quality of the fits to negative meson production data when used by the T2K experiment and therefore, are used here also [78].

This parameterisation was originally developed for use in modelling hadronic interactions and has been shown to work well at fitting a smooth curve to hadron production data points. As this is a parameterisation of the invariant double differential cross section as a function of the energy scaling variables x_F , x_R and P_T , it is also suitable to test the validity of the Feynman scaling discussed in Section 5.3.1. During the development of this parameterisation, it was used to compare proton-beryllium interactions at 24 and 450 GeV and demonstrated good agreement with the Feynman scaling hypothesis. These reasons make it a suitable starting point to be used as a parameterisation to compare the datasets that were discussed in Chapter 4.

For each of the datasets, the value of BMPT is calculated by converting the binning from the experiment into units of x_R and P_T . The BMPT function is then integrated across each two-dimensional bin and the parameters are evaluated in a reduced χ^2 fit. At this stage, each dataset is fit independently for each meson species and incoming primary particle momentum. The parameterisation was originally intended to jointly fit meson and anti-meson spectra (e.g. the π^+ and π^- spectra having the same parameters) with only the additional component for the negative meson as in equation 5.8; however, this was found to result in poor modelling of the data with more the precise modern datasets.

The exact format of uncertainties varies from dataset to dataset and is summarised in Section 5.4.1. For the purposes of the fits in this section the following equation is used to calculate the χ^2 :

$$\chi^2 = \sum_i \left(\frac{\left(\frac{d^2\sigma}{dx dy} \right)_i - f_{\text{BMPT}}}{\sigma_i} \right)^2 \quad (5.9)$$

where

- $\left(\frac{d^2\sigma}{dx dy} \right)_i$ is the experimentally measured value of the cross section in the binning scheme (x, y) of the dataset

- f_{BMPT} the value of the BMPT function at that bin which has been averaged over the area of the bin where applicable
- σ_i is the uncertainty on that bin given in the datasets. If the dataset gives statistical and systematic uncertainties separately σ_i is taken as the sum of these uncertainties in quadrature.

Note that the overall normalisation error has been ignored at this stage as it will not affect the χ^2 minimisation.

The fitting was performed by writing fitting software which used the MINUIT [83] package for minimisation of the χ^2 . MINUIT was configured to use the MIGRAD [83] algorithm for the minimisation. Fitting required choosing initial values for the BMPT parameters prior to minimisation from which the software finds the minimum. The initial values of the BMPT parameters were taken as those previously determined in Ref. [75]. To ensure the robustness of the minimum in χ^2 that was found, perturbations to the initial values were applied and the fitter was re-run to ensure the same minimum was found. This ensured that the minimum in χ^2 found by the software did not depend on the initial parameter values that had been selected.

Modifications

Initial fits using the BMPT parameterisation revealed fits to several datasets with poor χ^2 per degree of freedom (see Table 5.3). This is particularly true for datasets with wide phase space coverage measuring production at large angles (low and negative x_F) such as E802 and NA49. This is likely a consequence of the BMPT parameterisation describing the cross section as a function of x_R and P_T only. As x_R is a measure of the energy in the centre of mass frame it is symmetric for both forward and backward moving particles in the centre of momentum frame (CMF) and therefore the BMPT formula has this inherent symmetry. For interactions of protons with heavy nuclei, this symmetry is manifestly broken. For neutrino beam experiments (for which this parameterisation was originally developed), the daughters moving backwards in the CMF are of less

Experiment	Beam Momen- tum/ GeV	Target	π^+ χ^2 per d.o.f.		π^- χ^2 per d.o.f.	
			BMPT Only	Modified BMPT	BMPT Only	Modified BMPT
HARP low angle	3	C	2.38/4	1.5/-1	1.44/1	1.2/-4
	5	C	30.9/18	21.2/13	11.0/14	10.0/9
	8	C	71.3/27	34.0/22	25.9/23	23.1/18
	12	C	38.2/32	30.1/27	29.1/28	25.8/13
HARP high angle	3	C	129.6/74	71.6/69	128.9/72	97.0/67
	5	C	129.9/74	77.0/69	220.1/72	134.2/67
	8	C	149.7/74	83.3/69	151.1/72	88.0/67
	12	C	168.2/74	111.1/69	149.5/72	82.4/67
E910	6.4	Be	10.5/23	8.9/18	43.5/22	27.4/17
	12.3	Be	58.9/64	45.2/59	118.3/57	68.6/52
	17.5	Be	65.1/67	49.9/62	118.8/65	60.5/60
E802	14.6	Al	396.6/194	299.5/189	852.5/178	285.8/173
Allaby	19.2	Al	34.1/35	17.9/30	17.1/25	12.9/20
Eichten	24	Al	22.7/57	17.2/52	11.5/42	6.0/37
NA61	31	C	594.4/407	524.2/402	1133.8/438	638.2/433
NA49	158	C	1081.2/263	487.0/258	1166.2/261	353.6/256

Table 5.3: χ^2 per d.o.f. when fitting the BMPT parameterisation to individual datasets for $\pi^\pm$ production. Compared here is the result of using the BMPT parameterisation directly (equation 5.8) and also with the modification described in Section 5.3.2.

importance. In contrast, for the modelling of atmospheric neutrinos, these high production angle particles form a significant contribution and need to be modelled correctly, as was shown in Section 4.2.3.

As the breakdown in BMPT validity is related to the asymmetry caused by massive nuclei, there is a strong similarity to the differences in physics that the A Scaling parameterisation models. The following is then suggested as a modification to the BMPT formula:

$$E \frac{d^3\sigma}{d^3\mathbf{p}} = \left(\frac{A}{A_{\text{Be}}}\right)^{\alpha_1} \left(\frac{A_{\text{Be}}}{2}\right)^{\alpha_2} \left[E \frac{d^3\sigma}{d^3\mathbf{p}} \right]_{\text{BMPT}} \quad (5.10)$$

$$\alpha_2(x_F, P_T) = \frac{2}{3} + a_0 x_F + a_1 P_T + a_2 x_F^2 + a_3 P_T^2 + a_4 x_F P_T \quad (5.11)$$

where a_i are parameters to be determined simultaneously to the BMPT parameters, α_1 is the previously fitted scaling from beryllium to heavier nuclei in Section 5.2 and A is the atomic mass of the target species. The use of x_F dependent on the direction of the daughter particle in the CMF in α_2 breaks the inherent forward-backward asymmetry of the BMPT parameterisation. In this new formalism, BMPT models the interaction of a proton on an isoscalar deuterium nucleus with the new parameterisation modelling the additional asymmetry that appears when scaling to heavier nuclei. The A-Scaling term α_1 then models the scaling of the cross section from beryllium to heavier species already discussed in Section 5.2. This builds on the flux modelling by the T2K collaboration [84] where x_F dependent terms were included to break the forward-backward symmetry. In contrast to the approach by T2K, additional free parameters are also included here to further improve the quality of the fits.

Fitting with these additional parameters gives significantly improved results. The new values for χ^2 are summarised in Table 5.3 under “Modified BMPT”. For datasets with a smaller number of data points (such as the 3 and 5 GeV HARP low angle data), there are too many additional degrees of freedom within this new fitting framework resulting in under-constrained BMPT parameters. When these energy regions are modelled in Section 5.5.2, the 5 GeV dataset will be grouped together with other datasets and this will no longer be a problem. The 3 GeV dataset will be parameterised by the original BMPT parameterisation without the additional parameters.

Fits to the NA61 and NA49 data are shown in figures 5.4-5.7, fits to the lower energy datasets are shown later when considering joint fits to these lower energy regions (see Section 5.5.2).

Despite the improvement in χ^2 per degree of freedom with the additional parameters, the NA61 and NA49 fits both have a χ^2 per degree of freedom greater than 1. It is worth noting at this stage that the extent of correlation of systematic

uncertainties between bins is unknown and the χ^2 has not taken into account any correlation between these bin-to-bin systematic uncertainties. The χ^2 per degree of freedom should not be used as the sole indicator of the goodness of the fit. Considering figures 5.4-5.7, we see the parameterisation typically describes the data well with most deviations appearing more reminiscent of bin-to-bin noise than an overall trend. The NA61 π^- production data does show some deviation at the higher momentum higher angle bins however these don't show an overall trend between theta bins. The negative x_F region of the NA49 data also shows deviations from the fit likely due to issues with the strength of the forward-backward asymmetry allowed with the extended BMPT parameterisation. It is seen later in figures 5.21 that pions produced in this region do not have a large contribution to the neutrino flux and therefore we will continue using the parameterisation to model hadronic interactions.

5.4 Comparing and Combining Datasets

Now that a parameterisation has been developed and applied to the datasets alongside a method of scaling between cross sections on different target materials it is possible to compare cross sections between different datasets and binning schemes. The goal of this section is to consider how the cross section varies with beam energy and explore how the datasets can be combined to model pion production cross sections across the range of energies relevant to the atmospheric flux. The uncertainty due to pion production will be shown in Chapter 9 to be the most significant source of uncertainty on the neutrino flux, so the methods of modelling pion production in the atmospheric shower will be described in detail within this section and Section 5.5.

This section will initially focus on energy scaling between the higher energy datasets, comparing the 31 GeV NA61 data to the 158 GeV NA49 data and then the NA49 data to the 450 GeV NA56 data. The results from this will form the basis for modelling pion production cross sections at interaction momenta greater than 31 GeV. The NA61 and NA49 datasets will be shown to cover a far wider




Figure 5.4: Extended BMPT parametrised fit to NA61 π^+ production data. The invariant differential cross section (defined in equation 5.4) is plotted against pion momentum in pion production angular bins

range of the relevant pion production phase space in the x_F - P_T plane than the lower energy datasets. It will also be seen that there the variation between cross sections is greatest at lower energies. As such we will consider the lower energy datasets separately in Section 5.5.

5.4.1 Individual Dataset Uncertainty

To compare datasets it is not only necessary to be able to interpolate the production cross section across x_F - P_T phase space but also to understand the uncertainty on




Figure 5.5: Extended BMPT parametrised fit to NA61 π^- production data. The invariant differential cross section (defined in equation 5.4) is plotted against pion momentum in pion production angular bins

that cross section at each point in the phase space. For each of the experimental datasets, there is a normalisation error (correlated across all bins and all production species) and bin-to-bin errors. These bin-to-bin errors comprise statistical and systematic effects with some datasets providing both statistical and systematic uncertainties separately and others only providing one of the statistical or systematic errors. See Table 5.4 for a summary of the errors provided for each of the datasets. For this section, the uncertainty of the parameterisation is calculated by refitting the data with each data point perturbed by its bin-to-bin error (with statistical




Figure 5.6: Extended BMPT parametrised fit to NA49 π^+ production data. The invariant differential cross section (defined in equation 5.4) is plotted against x_F in transverse momentum (P_T) bins




Figure 5.7: Extended BMPT parametrised fit to NA49 π^- production data. The invariant differential cross section (defined in equation 5.4) is plotted against x_F in transverse momentum (P_T) bins

Experiment	Normalisation Uncertainty	Bin-to-bin uncertainty
HARP	5%	Separate Statistical and Systematic Uncertainties
E910	5%	Combined statistical and systematic uncertainty
Eichten	15%	Systematic Uncertainties Only
Allaby	15%	Combined statistical and systematic uncertainty
E802	15%	Statistical Uncertainty Only
NA61/SHINE	1%	Separate Statistical and Systematic Uncertainties
NA49	2.5%	Separate Statistical and Systematic Uncertainties
NA56	10%	Separate Statistical and Systematic Uncertainties

Table 5.4: Summary of uncertainties given for each experimental dataset.

and systematic components added in quadrature). This gives an error band within x_F - P_T space with which to compare dataset parameterisations. While it is implied that these bin-to-bin errors will have a significant degree of correlation, the datasets do not provide correlation matrices. Further consideration of the effects of the hadron production uncertainty across x_F - P_T phase space on the neutrino flux will be explored in Chapter 9.




Figure 5.8: Energy dependence on $\pi^\pm$ inclusive production cross section in proton-proton interactions. Figure taken from [85]. This is used as a guide for how we may expect the energy dependence of the pion production cross section in proton nuclear interactions.

5.4.2 Scaling with Energy

We start by exploring how pion production cross sections scale with energy. A study by Rossi et al. [85] from 1974 looking at the energy dependence in proton-proton inclusive pion production cross sections found that at high energies the Feynman scaling discussed in Section 5.3.1 held, and the cross sections expressed as a function of x_F and P_T were approximately independent of energy i.e.:

$$\sigma(\sqrt{s}, x_F, P_T) \rightarrow \sigma(x_F, P_T), \quad (5.12)$$

but this did not hold at lower energies. Figure 5.8 from that study demonstrates this result for $x_F = 0$ at a series of transverse momenta points. We now have datasets across large regions of x_F - P_T phase space with many more data points and can investigate this scaling using the higher energy datasets for proton interactions with a target nucleus.




Figure 5.9: Ratio of the NA61 pion production cross section at 31 GeV to the NA49 158 GeV pion production cross section. Shown for π^+ (top) and π^- (bottom). The absolute ratio between the parameterised fits is shown on the left-hand side. The deviation of this ratio from 1 in terms of number of standard uncertainty intervals σ is shown on the right-hand side. These standard uncertainty intervals are propagated to the fit as described in Section 5.4.1 and are the propagation of the bin-to-bin uncertainties only. In addition to this, there is a combined normalisation uncertainty however it is clear from these plots that the shape uncertainty dominates the difference between the two datasets. In both cases, the phase space of the measurement for each experiment is shown and the ratio is only shown within these regions.

Production Cross Section Scaling Between 31 and 158 GeV

The ratio of the production cross section between 31 and 158 GeV is shown in Figure 5.9. Considering again Figure 4.2, it is seen that this energy region covers the majority of the interaction phase space for neutrinos between 2 and 10 GeV. The ratio and significance in Figure 5.9 show clear deviation from 1 indicating that across this region Feynman scaling does not hold. Since there are currently no other datasets within this large interaction energy region there will be uncertainty due to the interpolation between these datasets. A linear interpolation of the production cross section will be used to interpolate production cross sections at energies between these datasets. The systematic uncertainty associated with the interpolation of the production cross section within these regions is discussed in 9.3.

Production Cross Section Scaling Between 158 and 450 GeV

We now turn to the highest energy datasets and compare the cross section between 158 and 450 GeV beam energy. The NA56 data covers a smaller phase space than the NA49 data and the ratio between their parametrised fits (shown in Figure 5.10) is shown for the region of overlap between the two datasets. The deviation between these datasets is smaller than the lower energy comparison and typically within 20%. This is also consistent with the findings from the proton-proton cross sections that at higher interaction energies Feynman scaling holds. In addition to the fits, Figure 5.11 shows NA49 data points in two P_T bins where the NA49 bin centres overlap with NA56 bins. For these reasons, the 158-450 GeV region will be modelled under the assumption of Feynman scaling using the NA49 data. Additional uncertainties arising from the extent of the agreement of these datasets with the Feynman scaling hypothesis are discussed in Section 9.3.2.

5.5 Lower Energy Datasets

Having established a method to model pion production at higher energies attention will now be turned to the lower (<31 GeV) energy datasets. Individually, these




Figure 5.10: Ratio of the NA49 pion production cross section at 158 GeV to the NA56 450 GeV pion production cross section. The ratio and significance are calculated as in Figure 5.9 with the significance due to the bin-to-bin uncertainty only. The ratio has been shifted by a normalisation factor of 8% to account for differences in normalisation between the datasets. This is within the overall normalisation uncertainty of 11% which is obtained by summing the NA49 and NA56 normalisation in quadrature.




Figure 5.11: NA49 and NA56 data where the bin centres overlap at $P_T = 0.15$ GeV and 0.5 GeV alongside fits to the NA49 data. Plots show the invariant production cross section for π^+ (top) and π^- (bottom) as a function of x_F at specific values of P_T .

datasets cover a smaller fraction of the relevant x_F - P_T phase space and will be shown to be in poor agreement in the phase space regions where they overlap. This is in addition to the expectation that the cross section will vary more significantly with the interaction energy as shown in Figure 5.8. Understanding pion production and its uncertainty in this lower energy region is a major challenge within this analysis and appropriate methods of inflating uncertainties to cover differences between the existing datasets will be discussed.

This section will start by attempting a combined fit to the HARP high and low angle datasets before grouping datasets by interaction energy and attempting a joint fit to datasets within each group. The differences between the datasets will be explored and a method of error estimation on pion production at these lower

		3 GeV	5 GeV	8 GeV	12 GeV
π^+	χ^2 per dof	80.7/80	120/94	196/103	187/108
	norm. diff	0.5 %	10%	10%	20%
π^-	χ^2 per dof	107/77	154/89	148/98	200/103
	norm. diff	2%	4%	3%	5%

Table 5.5: χ^2 per degree of freedom for the joint fits to the HARP High and Low Angle datasets. The normalisation uncertainty for each dataset is 2%.

interaction energies will be developed. This will be built upon when propagating this uncertainty through to the neutrino flux in Chapter 9.

5.5.1 HARP High and Low Angle Datasets

As an initial attempt of fitting multiple datasets, we try a combined fit of the high and low-angle data published by the HARP experiment. As discussed in 4.3.4, HARP was uniquely designed to measure high-angle production cross sections however due to hardware issues with the TPC the high-angle data was analysed and published separately.

Attempted fits of the BMPT parameterisation between the high and low angle data at each energy did not work well with poor χ^2 per degree of freedom and large normalisation differences between the datasets chosen by the fitter (see Table 5.5). Given this, alongside the previous concerns with the differences between two different analyses of the HARP high angle data, we will treat this region separately and not use the joint fit or interpolation between the HARP high and low angle data. The uncertainties in the high-angle data are discussed in Section 9.1.3.

5.5.2 Jointly Fitting Low Angle Energy Datasets

The datasets in the lower energy region cover smaller pion production phase spaces than the NA61/SHINE and NA49 datasets. No single dataset covers enough of the

Region	Average $\sqrt{s}$ Energy/ GeV	Experimental Datasets	χ^2 per dof of joint fit π^+	χ^2 per dof of joint fit π^-
1	2.77	3 GeV HARP Low Angle	2.38/4	1.44/1
2	3.54	5 GeV HARP Low Angle 6.4 GeV E910	48.9/43	51.7/40
3	4.11	8 GeV HARP Low Angle	37.4/22	22.3/18
4	5.17	12 GeV HARP Low Angle 14.6 GeV E802 12.3 GeV E910	634.5/299	475/276
5	6.38	17.5 GeV E910 19.2 GeV Allaby 24 GeV Eichten	478.0/168	315.6/145
6	7.74	31 GeV NA61	524.2/402	638.2/433
7	17.27	158 GeV NA49	487.0/258	353.6/256

Table 5.6: Grouping of the datasets into energy regions and results from the joint fit to the cross sections. In addition to the 5 low (<31 GeV) energy regions we define regions 6 and 7 as being modelled by the NA61 and NA49 data respectively. The region numbers are used in the text and future plots. The average energy is shown in $\sqrt{s}$ (where s is the usual Mandelstam variable [77]), alongside the beam momentum for each of the datasets. The Mandelstam variable s is calculated for the interaction between the incident proton and a single nucleon within the target nucleus. When calculated this way for a fixed target experiment, s is the same for all interactions with the same beam energy regardless of target nucleus mass. As explained in the main text the HARP high angle data is not included within the joint fit but it is used to model the high angle region separately in regions 1-4. The χ^2 per degree of freedom for the joint fits 1-5 is calculated before the error scaling technique described in Section 5.5.2 takes place and as such will decrease by a factor of S from Table 5.7 after error inflation.

hadronic production phase space to sufficiently model pion production interactions alone. Therefore to maximise phase space coverage multiple datasets at similar interaction energies are combined (based on the invariant $\sqrt{s}$ Mandelstam variable [77]). The interaction energy for pion production is separated into five energy regions, labelled 1 to 5, between 3 GeV and 24 GeV. The datasets used to model pion production at each of these energies are shown in Table 5.6. In addition, we label the higher energy NA61 and NA49 datasets discussed in Section 5.4 as 6 and

7 respectively and include them in this table. The final goal of this subsection will be to perform a combined fit of each dataset at each of these 5 energy regions which will then be used to model pion production at lower energies in the atmospheric showers. In addition, separate fits to the HARP high angle datasets will be used for pion production at high angles.

A joint fit of the BMPT function with the additional parameters discussed earlier in Section 5.3.2 was attempted at each energy region modelled by multiple datasets (regions 2, 4 and 5). The BMPT parameters were determined by minimising the following χ^2 :

$$\chi^2_{\text{joint}} = \sum_{i,j} \left(\frac{\left(\frac{d^2\sigma}{dx dy} \right)_i - N_j \cdot f_{\text{BMPT}}}{\sigma_i} \right)^2 + \sum_j \left(\frac{1 - N_j}{\sigma_{\text{norm}}} \right)^2 \quad (5.13)$$

where N_j is the normalisation parameter for experimental dataset j and σ_{norm} is the normalisation uncertainty of that dataset. The other parameter definitions follow equation 5.9.

The values of the χ^2 per degree of freedom are shown in Table 5.6 and can be seen for several of the regions to be large and indicate a poor combined fit. The reason for this is found to be due to differences between measurements made by each experiment within the region, despite their highly similar interaction energies. As previously, discussed direct comparisons between data points can be difficult to visualise due to different two-dimensional binning schemes. Figures 5.12 - 5.14 show data points where the bin centres lie within 10 mrad of each other. To make these figures, the target nucleus scaling parameterisation has been used to scale the invariant cross sections to a nitrogen target. Similarly, the $p - \theta$ values have been scaled to the same centre of mass energy such that the values of $p - \theta$ correspond to the same values of $x_F - P_T$ across the slight difference in beam energies of the different experiments. Discrepancies between the data points can be observed with the greatest differences being between the 17.5 GeV E910 data and the Eichten and Allaby data in Figure 5.14. Even where there is minimal overlap between bin centres, differences between the datasets and the combined fit are different for




Figure 5.12: Datasets and joint fit in θ bins used to model interactions at $\sqrt{s} \approx 3.54$ GeV. For the HARP and E910 datasets the bin centres do not overlap within a 10 mrad wide bin to compare side by side. The top graphs show the invariant cross section for π^+ production while the bottom are for π^- production.

each experiment indicating incompatibility between measurements. For example in Figure 5.13, the E802 and E910 data points typically fall above the combined fit line while the HARP data points typically fall below the combined fit line.

Similarly, ratios of the individual fit to each dataset to the joint fit at each energy region are shown in figures 5.15 and 5.16 where large differences between the combined and individual fits can be observed.

These discrepancies have likely arisen from a combination of an incomplete understanding of the uncertainties on the data (as highlighted in Table 5.4), uncertainties not covering the differences due to the slight differences in interaction energy or uncertainties arising due to comparing cross sections on different nuclear targets. Regardless of the source of these uncertainties a method needs to be developed to estimate the uncertainty on the production cross section to cover these discrepancies that will be used for the modelling of atmospheric neutrino fluxes.

To deal with these discrepancies, a method is developed based on an error scaling technique used by the PDG group to scale uncertainties where the data is found to




Figure 5.13: Datasets and joint fit in θ bins used to model interactions at $\sqrt{s} \approx 5.17$ GeV. The top graphs show the invariant cross section for π^+ production while the bottom are for π^- production. Data points from different experiments are shown where the centre of the data points overlap within the same θ bin.




Figure 5.14: Datasets and joint fit in θ bins used to model interactions at $\sqrt{s} \approx 6.38$ GeV. The top graphs show the invariant cross section for π^+ production while the bottom are for π^- production. Data points from different experiments are shown where the centre of the data points overlap within the same θ bin.




Figure 5.15: Ratio between the joint BMPT fit to each of the datasets in the same energy region and the individual BMPT fit to each dataset for π^+ production. Shown for top: region 2 ($\sqrt{s} \approx 3.54$ GeV), middle: region 4 ($\sqrt{s} \approx 5.17$ GeV), bottom: region 5 ($\sqrt{s} \approx 6.38$ GeV)




Figure 5.16: Ratio between the joint BMPT fit to each of the datasets in the same energy region and the individual BMPT fit to each dataset for π^- production. Shown for top: region 2 ($\sqrt{s} \approx 3.54$ GeV), middle: region 4 ($\sqrt{s} \approx 5.17$ GeV), bottom: region 5 ($\sqrt{s} \approx 6.38$ GeV)

be in disagreement [80]. While this technique was developed for single measurements, it is extended here across two-dimensional x_F - P_T space. The parameter S is defined at each point in x_F - P_T phase space to be:

$$S(x_F, P_T) = \sqrt{\frac{\chi^2}{N(x_F, P_T)}} \quad \chi^2 = \sum_i^N \frac{(f_{\text{joint}} - f_{\text{single}})^2}{\sigma_{\text{single}}^2}. \quad (5.14)$$

Where $f_{\text{joint/single}}$ is the value of the joint/single fit respectively, N the number of experiments at that point in x_F - P_T phase space and σ_{single} the uncertainty on the fit as estimated in Section 5.4.1. In effect, S is the ratio of the variance of the data points around the joint fit to their uncertainties. A factor of $S \leq 1$ indicates datasets in good agreement with one another while $S > 1$ would indicate discrepancies (which side-by-side comparisons of the datasets have already revealed). Where $S > 1$ it is an indicator of the factor by which the errors should be inflated. We consider the plots of the S factors for the regions with multiple experiments in Figure 5.17.

Considering the small number of datasets (2 or 3) within each energy region, there are two caveats which need to be considered when judging the scale of disagreement between datasets using the S factor. Firstly, statistical fluctuations could cause an S factor greater than 1 in small regions of x_F - P_T phase space even if the datasets are in agreement with each other. Indeed, with a large number of data points across x_F - P_T phase space it would not be expected that every data point from two measurements would lie within one standard deviation of each other, even if the datasets are in agreement. Secondly, by the same argument, two discrepant datasets may have regions in x_F - P_T phase space where the results agree by chance. Only inflating the uncertainties in the regions of x_F - P_T space where $S > 1$ may have the unintended consequence of underestimating uncertainties in the regions where there are underlying differences but statistical fluctuations cause the measurements to be closer together. For each energy region in Figure 5.17 a value of the S parameter is chosen based on the largest significant value of S based on these principles. It is then used to inflate all experimental uncertainties

Energy Region	$\sqrt{s}$	Error scaling factor π^+	Error scaling factor π^-
2	3.54	1.5	2.0
4	5.17	2.5	2.5
5	6.38	2.5	2.5
1	2.77	2.0	2.0
3	4.11	2.0	2.0

Table 5.7: Summary of the S error inflation factors chosen for <31 GeV data regions for π^+ and π^- production cross sections. The method for obtaining the factor is discussed in Section 5.5.2 with regions 2,4,5 using Figure 5.17. An argument for the value of the error scaling factors for regions 1 and 3 where there are no overlapping data is presented in the main text.

across x_F - P_T phase space so that the uncertainties reflect the variance between the datasets. The S factors chosen are shown in Table 5.7.

Energy regions 1 and 3

Energy regions 1 and 3 are both modelled using fits to the HARP data at 3 GeV and 8 GeV respectively with no other dataset for comparison. Since the 5 GeV and 12 GeV HARP data when been compared to other datasets has resulted in subsequent inflation of uncertainties, it would be unwise to assume that the uncertainties on these datasets do not also require inflation. As such, an error inflation factor of 2 is assigned to these regions corresponding to an approximate average of the error inflation factors used for regions 2 and 4 where other HARP data is used. Without other datasets at similar energies for comparison this factor is only a rough estimate; however, it is expected to be a conservative one.

For each of the regions 1-5, the combined fit to the data is used to model the production cross section with the scaled uncertainties used to estimate the uncertainty on the cross section in this region. As the uncertainty scaling is applied to each dataset equally, it does not change the parameters obtained in the combined fit. It does change the uncertainty, and the effect that this has on the neutrino flux is discussed in Chapter 9.




Figure 5.17: S error scaling factor as defined in the main text equation 5.14. Shown for energy regions 2,4 and 5 (top to bottom). $\pi^+(\pi^-)$ production on the left(right) hand side.




Figure 5.18: Fractional contribution of pion production interactions within each of the 7 energy regions defined in 5.6 to the ν_μ (left) and ν_e (right) atmospheric flux.

5.6 Data Phase Space Coverage of the Atmospheric Flux

In the previous section, we demonstrated how experimental data can be used to model pion production interactions in atmospheric showers. The datasets were grouped into seven distinct regions (five low energy and the NA61 and NA49 regions) with combined BMPT fits used to calculate the cross section across x_F - P_T pion creation phase space at each interaction energy region. It is now helpful to consider how well the datasets cover the phase space of hadronic interactions which contribute to the neutrino flux. The plots in this section show the fractional contributions to hadronic interactions at different energy regions and daughter pion phase spaces where TARGET2.4 has been used to model the hadronic interactions. It is to be expected that there will be slight changes to these contributions once the simulation has been tuned to the datasets used in this study.

Figure 5.18 shows the fractional contribution from pion production interactions at each of the seven energy regions to the atmospheric neutrino flux. The total

contribution due to pion production interactions is shown in black and seen to decrease from almost 1 at neutrino energy 0.1 GeV to around 0.8 at 10 GeV due to the increase in atmospheric neutrinos created in kaon decays. We see that for the 2 - 10 GeV neutrino energy region, the NA61 and NA49 datasets (region 6 and 7) cover the majority of the interaction phase space. For the lower energy neutrinos, the datasets discussed in Section 5.5.2 are more relevant. At 0.1 - 0.3 GeV neutrino energies there is a significant component from interactions at energies below the lowest dataset available (<3 GeV), this region is discussed further in Section 9.4.4.

We then need to consider what fraction of pion production phase space in atmospheric showers is covered by experimental data. For each of three different neutrino energies, figures 5.19 - 5.21 break down the pion production contribution to the ν_μ flux into 2D x_F - P_T pion production phase space at each interaction energy. The phase space covered by each of the datasets is overlaid. Interactions lower than the tuning energy (<3 GeV and labelled ‘Energy Bin 0’) and above the tuning energy (>450 GeV labelled ‘Energy Bin 8’) are also shown. These figures show firstly the relative contribution of each interaction energy bin (similar to figure 5.18), and then how well each of the energy regions is covered by data. The HARP high angle dataset coverage is included also.

Figure 5.22 summarises this result, showing the fraction of total pion production interactions inside and outside the dataset coverage. The importance of the HARP high angle datasets is shown separately in this figure and is shown to be most significant for low-energy neutrino fluxes. This is also seen in Figure 5.19 for the 0.2 GeV neutrino flux with a significant proportion of the total x_F - P_T phase space falling into the region covered by this high angle data.

The experimental datasets cover around 80% of the pion production interactions relevant to higher energy neutrino fluxes (1-10 GeV). For interactions relevant to the lower energy neutrino fluxes (0.1-1 GeV) a greater proportion of the interactions occur outside the dataset coverage - this is a combination of interactions that occur below the energy of the datasets (<3 GeV) and interactions occurring within the energy where there is data but outside the x_F - P_T phase space coverage of




Figure 5.19: Contribution of different phase space components of pion creation reactions $pN \rightarrow \pi^+ + X$ to the atmospheric muon neutrino flux at 0.18 GeV. The contribution is broken down into the 2D phase space of different interaction energy bins chosen by experimental data given in Table 5.6. Interaction energy bin 0 is an underflow bin for interactions < 3 GeV while interaction bin 8 is an overflow bin for interactions > 450 GeV. The first plot shows the total contribution of each energy region to the atmospheric flux in arbitrary units. Each subsequent plot shows the contribution in x_F - P_T phase space to the flux overlaid with the coverage of the experimental data. The percentage indicates the percentage of the interaction phase space covered by the experimental data.




Figure 5.20: Contribution of different phase space components of pion creation reactions $pN \rightarrow \pi^+ + X$ to the atmospheric muon neutrino flux at 1.13 GeV. Equivalent to Figure 5.19 for the flux of muon neutrinos at 1.13 GeV. Note interaction energy bin 0 is too low in energy to produce neutrinos at this energy.

the datasets. This second component is greater for the lower energy neutrinos both due to their greater dependence on pion production at higher angles and that the lower energy datasets have smaller phase space coverages than the higher energy NA49 and NA61 datasets. The method for tuning and uncertainties for the interactions occurring outside the experimental datasets is discussed in Chapter 8. Pion production (<3 GeV) occurs in a region where the resonances in the interaction cross section dominate and as such has a greater variation with beam energy.




Figure 5.21: Contribution of different phase space components of pion creation reactions $pN \rightarrow \pi^+ + X$ to the atmospheric muon neutrino flux at 7.12 GeV. Equivalent to Figure 5.19 for the flux of muon neutrinos at 7.12 GeV. Note interaction energy bins 0-3 are too low in energy to produce neutrinos at this energy.

Systematic uncertainties associated with this region are discussed in Section 9.4.4.

5.7 Summary

This chapter has developed a model for pion production interactions ($pN \rightarrow \pi^\pm + X$) based on the parameterisation of the datasets described in Chapter 4. The variation of the production cross section with energy has been considered, with Feynman scaling used to scale between the production cross sections of the highest energy




Figure 5.22: Extent to which the phase space of pion production interactions in atmospheric showers are covered by the experimental datasets in Section 4.3. Inside refers to the percentage of pion creation phase space covered by the experimental datasets. Outside phase space refers to interactions that occur in a primary (beam) energy region where data has been taken but the pion is created outside the phase space coverage of the experiment. Outside energy refers to pions created in a primary energy region where data was not taken. This is further broken into interactions lower in energy than the datasets (< 3 GeV) and higher in energy than the available datasets (> 450 GeV). Finally, the inside region has been separately plotted excluding the HARP High Angle data to show the importance of this data at low energies and how it is less important at high energies.

datasets. An empirical scaling was developed and used to interpolate between the pion production cross section in proton interactions with different target nuclei. Discrepancies between the lower energy datasets led to the development of an uncertainty inflation technique to cover those differences. Finally, it was shown that within this model the datasets cover $\sim 80\%$ of the pion production phase space relevant to modelling the 2-10 GeV neutrino flux. This decreases to $\sim 40\%$ when modelling the 0.1 GeV neutrino flux due to a combination of both smaller outgoing pion phase space coverage by the datasets and increased dependence on interactions occurring below 3 GeV - the lowest energy dataset. Having developed this model it will be applied to the calculation of neutrino fluxes in Chapter 7.

6

Modelling Kaon Production in Air Showers

This chapter considers the modelling of kaon production interactions in atmospheric showers. This builds on the tools and methods used in Chapter 5 to model pion production interactions. The contribution of both charged and neutral kaons to the atmospheric neutrino flux is considered, with charged kaon production cross section datasets used to develop a model to apply to the flux calculation. A method of estimating neutral kaon production and kaon production during incident neutron interactions is introduced. Neutrinos coming from kaon decays have been shown in Chapter 2 to be subdominant to those from pion decays with around $\sim 20\%$ of neutrinos at 10 GeV originating in kaon decays. While important to the modelling of the neutrino fluxes kaon production is less important than pion production.

6.1 Kaon Component of Atmospheric Showers

Having considered the dominant pion contribution to the 0.1-10 GeV atmospheric neutrino flux, the smaller kaon component will now be considered. The analysis techniques used here are similar to the pion modelling in Chapter 5 and where they are the same will not be explained again in detail. There are a few differences that need to be considered when modelling kaons. One of the key differences is




Figure 6.1: Contribution to each flavour of atmospheric neutrino flux by decays of charged and neutral kaons. For each atmospheric neutrino flavour the fraction of the flux attributed to $K^\pm$ and $K_{L/S}^0$ decays are shown.

that while neutral pions decay into photons and therefore do not contribute to the neutrino flux, neutral kaons can decay into neutrinos (either directly or via the decays of their daughter products).

Figure 6.1 shows the relative contribution from both charged and neutral kaons to the atmospheric flux. It is to be noted that the total kaon contribution is higher for the neutrino (ν_μ and ν_e) as opposed to the antineutrino ($\bar{\nu}_\mu$ and $\bar{\nu}_e$) fluxes. This can be explained by considering firstly that K^+ are the most abundantly produced kaon in atmospheric showers. This is in part due to their being comprised of an up quark (alongside the strange quark component) which exists within the valance quarks of the interacting nucleons, unlike K^- and K^0 which consist of an anti-up and anti-down quark respectively. K^+ can also be produced alongside the neutral lambda baryon in an interaction which conserves both strangeness and baryon number without the need for additional strange hadrons. Since the most common decay modes of the K^+ favour neutrino over antineutrino final states, this explains the increased kaon component of the ν_μ and ν_e neutrino fluxes.

The second point to notice from Figure 6.1 is that the ν_e and $\bar{\nu}_e$ flux both have a significant contribution from the K_L^0 meson. This is due to its dominant decay mode into electron neutrinos (branching ratio 40%) via semi-leptonic decay:

$$K_L^0 \rightarrow \pi^\pm e^\mp \nu_e. \quad (6.1)$$

Looking back at the meson decays to neutrinos discussed in Section 2.1, it can be seen that alongside the charged kaon decay mode $K^+ \rightarrow \pi^0 e^+ \nu_e$ (with branching ratio 5%), these are the only decays of a meson directly to electron neutrinos with all other electron neutrinos coming from decays of muons. Neutral kaon production is therefore of importance to the atmospheric neutrino flux. The production of neutral kaons in hadronic interactions is difficult to measure directly in fixed target experiments due to their being electrically neutral and not leaving tracks in TPCs or drift chambers. A method for estimating their production from charged kaon production is discussed in Section 6.3.

6.2 Experimental Datasets

There are fewer experimental datasets for kaon production cross sections than pion production and the datasets that are available typically have larger uncertainties. This is both due to fewer kaon events than pion events (leading to larger statistical uncertainties) and difficulties in particle identification separating kaons from protons and pions arising from there being fewer kaons. Looking back at Figure 4.2 it is seen that the relevant kaon production interactions occur at energies similar to the relevant pion interactions with kaon production at higher energies being more relevant due to higher energy neutrinos increasingly being produced in kaon decays.

The following experiments provide charged kaon production cross sections at a specific proton beam energy and are chosen to model kaon production in atmospheric showers:

- 14.6 GeV E802 $pBe \rightarrow K^\pm + X$ [69]

- 19.2 GeV Allaby $pBe \rightarrow K^\pm + X$ [70]
- 31 GeV NA61/SHINE $pC \rightarrow K^\pm + X$ [56]
- 158 GeV NA49 $pC \rightarrow K^\pm + X$ [76]

While these experiments are described in detail in Section 4.3, it is worth considering a couple of additional points on how they are designed to accommodate the measurement of daughter kaons. Both E802 and Allaby use threshold Čerenkov detectors which allow them to separate pions and kaons. The NA49 detector uses $\frac{dE}{dx}$ energy loss in the TPC for particle identification. Due to significant overlap between pion, proton and kaon $\frac{dE}{dx}$ curves the kaon spectra can only be measured over a comparatively small kaon phase space [76]. NA61/SHINE experiment, which uses much of the same apparatus as NA49, also includes a TOF measurement which can be combined with the $\frac{dE}{dx}$ energy loss to allow measurement over a larger phase space of the outgoing kaon [56].

6.3 Neutral Kaon Cross Sections

Earlier it was demonstrated that neutral kaon production and decays provide an important component to the electron neutrino fluxes. The direct measurement of K_L^0 production with either TPC or single-arm magnetic spectrometer-based detectors is difficult. In principle K_S^0 decays to pions can be used to measure its production by accounting for the displaced vertex [86] but their uncertainties are typically large and the use of such measurements is beyond the scope of this thesis. An estimate for neutral kaon production using a weighted combination of charged kaon production cross sections based on their quark contents is presented here, adapted from a technique developed in Ref. [57]. A simplified quark-parton model is assumed, with the partons from both the sea and valence quarks. In this framework, it is assumed that:

$$u_s = \bar{u}_s = d_s = \bar{d}_s = s_s = \bar{s}_s \quad (6.2)$$

and

$$n = \frac{u_v}{d_v} \quad (6.3)$$

with $n = 2$ for proton-proton collisions and $n = 1$ for proton-neutron collisions. Within this framework the neutral kaon production cross sections are derived as [56, 57]:

$$\sigma(K_L^0) = \sigma(K_S^0) = \frac{1}{8}(3\sigma(K^+) + 5\sigma(K^-)) \quad (6.4)$$

where $\sigma(K)$ is the production cross section for a specific species of kaon for proton collisions with an isoscalar nucleus (such as the carbon target).

6.3.1 Kaon Production by Incident Neutrons

As already discussed, a small fraction of incident cosmic rays are light nuclei. The cross section of kaon production in interactions of an incident neutron on air needs to be understood to be able to model these cosmic ray interactions. These interactions are not expected to significantly contribute to the neutrino flux within the energy range considered since cosmic rays mainly comprise protons and most neutrinos come from pion decays rather than kaons. The following argument is used to estimate their interaction cross section. This starts by invoking isospin symmetry (as with modelling pion production by neutrons) to reach:

$$\sigma(nN \rightarrow K^+) = \sigma(pN \rightarrow K^0) \quad (6.5)$$

$$\sigma(nN \rightarrow K^-) = \sigma(pN \rightarrow \bar{K}^0) \quad (6.6)$$

Repeating the arguments above by weighting the incident proton cross sections by their sea and valence quark content:

$$\sigma(nN \rightarrow K^+) = \frac{1}{4}(3\sigma(pN \rightarrow K^+) + \sigma(pN \rightarrow K^-)) \quad (6.7)$$

$$\sigma(nN \rightarrow K^-) = \sigma(pN \rightarrow K^-) \quad (6.8)$$

Energy Region	Datasets	$K^+\chi^2$ per dof	$K^-\chi^2$ per dof
1	14.6 GeV E802 and 19.2 GeV Allaby	106.1/ 125	116.3/ 85
2	31 GeV NA61	66.3/ 82	101.6/ 75
3	158 GeV NA49	37.7/ 36	32.6/ 35

Table 6.1: Results for the modified BMPT parameterisation to the Kaon datasets

The production of K^- , being the same for incident protons or neutrons naively makes sense as K^- only contains quarks that are sea quarks in both protons and neutrons.

Similarly, the production of neutral kaons in incident neutron interactions can be estimated using isospin symmetry exclusively:

$$\sigma(nN \rightarrow K_{L/S}^0) = \frac{1}{2}(\sigma(pN \rightarrow K^+) + \sigma(pN \rightarrow K^-)) \quad (6.9)$$

6.4 Applying Kaon Data to Atmospheric Showers

The extended BMPT parameterisation discussed in Section 5.3.2 for pion production cross sections is fit to the kaon production data. The 14.6 GeV E802 and 19.2 GeV Allaby data were jointly fit to achieve wider phase space coverage of the lower energy region. Since these joint fits worked well, no additional increase of the uncertainties is necessary and the region of phase space between the E802 and Allaby datasets is modelled by this joint BMPT parameterisation. The results from these fits are shown in Table 6.1.

Initially, the NA49 and NA61 measurements were compared using these fits in order to determine the extent of variation in production cross sections between these two energies and whether Feynman scaling arguments could be used. This is shown in Figure 6.2. These demonstrate clear differences between the cross sections at these two energies and the need for using the NA49 data even though it has significantly lower phase space coverage than the NA61 data. The fact that kaon production is greater at 158 GeV beam momentum measured by the NA49 than the 31 GeV measured by NA61 could be expected from the higher mass of the strange quark.




Figure 6.2: Ratio of Kaon production at 31 to 158 GeV using fits to NA61 and NA49 data. Top are the production of K^+ with the 31 to 158 GeV ratio (left) and deviation of that ratio from 1 (right) using the uncertainty on the data added in quadrature. Bottom are repeated for the production of K^- . Each ratio is only plotted in the phase space region where each experiment has made measurements. Note the significantly lower NA49 phase space coverage, the reasons for this were discussed in Section 6.2. These plots demonstrate a clear difference between kaon production at these two energies.

To model kaon production in atmospheric showers, the three energy regions in Table 6.1 will be used. Scaling of the kaon production cross section from the carbon and beryllium target data is done using the parameters derived for pions in Section 5.2. The kaon production cross section can be interpolated between these three energy regions to model kaon production from 14-158 GeV. A summary of the proportion of the atmospheric flux inside the dataset phase space (including the use of both charged kaon production datasets and the estimated production of neutral kaons from Section 6.3) is shown in Figure 6.3. The datasets typically cover the majority of the relevant kaon production phase space, with only a small component of kaon production occurring below the dataset energies (<14 GeV). There is a significant component of kaon production from interactions within the dataset energy range but outside the outgoing kaon phase space coverage. This is




Figure 6.3: Breakdown of the kaon contribution to the atmospheric neutrino flux. Shown for all 4 neutrino flavours ν_μ (top left), $\bar{\nu}_\mu$ (top right), ν_e (bottom left) and $\bar{\nu}_e$ (bottom right). In each case, the fractional contribution is shown as a function of neutrino energy. The kaon contribution is broken down by interactions covered by the datasets in the 14-158 GeV Region, interactions occurring above and below this energy region and interactions occurring at this energy region but outside the x_F - P_T phase space where the experiments measured kaon production.

most likely due to the comparatively small NA49 phase space coverage. Additionally, there is a significant component due to interactions occurring above the 158 GeV energy of the NA49 dataset. The uncertainties on the production outside dataset phase space are discussed in Chapter 9.

6.5 Summary

Within this chapter kaon production within atmospheric showers has been considered. Four different datasets are used to describe kaon production for interactions between 14 and 158 GeV incident nucleon momentum. Similar to pion production, kaon

production is modelled using the extended BMPT parameterisation with parameters determined in a fit to the datasets. A method for estimating neutral kaon production and production of kaons by incident neutrons has also been described. This model will be applied to the calculation of neutrino fluxes in Chapter 7.

7

Tuning the Atmospheric Neutrino Flux to Datasets

This chapter explores a method of tuning the atmospheric neutrino flux with the pion and kaon production dataset models from the previous chapters. The differences between the production cross section and multiplicity are discussed in the context of interactions occurring within the atmosphere. A method of weighting the multiplicities in the atmospheric neutrino calculation is presented. The interactions which fall inside the phase space coverage of the datasets are weighted using this method and the results for tuned neutrino fluxes are shown.

7.1 Applying the Hadronic Models to the Atmospheric Monte Carlo

Having developed models for pion and kaon production in Chapters 5 and 6, the next stage is to apply this to the atmospheric neutrino Monte Carlo simulation. Until this point we have considered hadron production in terms of a differential cross section. This cross section combines two pieces of information: firstly the probability of an incoming particle interacting on a given target nucleus (the total production cross section) and secondly the probability distribution for hadron production with

specific kinematics given an interaction has occurred (the differential multiplicity). Mathematically this can be expressed as [56]:

$$\frac{d^2\sigma_h}{dx_F dP_T} = \sigma_{\text{prod}} \frac{d^2n_h}{dx_F dP_T}, \quad (7.1)$$

where σ_{prod} is the interaction cross section for all hadron production interactions and $\frac{d^2n_h}{dx_F dP_T}$ is the double differential multiplicity of production of a specific species of hadron.

The production cross section is inversely proportional to the interaction length of primary protons in air. Since the atmosphere has an approximate depth of at least 10 interaction lengths¹ virtually all protons will have interacted before reaching the surface of the Earth. Because of this, perturbations to the production cross section will not significantly change the number of proton interactions occurring in the atmosphere only alter the altitude at which these interactions occur. Therefore, changes in the production cross section will not significantly change the total production of mesons in the atmosphere. The production cross section will still affect the neutrino flux due to mesons produced at higher altitudes being more likely to decay before reaching the surface of the Earth resulting in an increase in the neutrino flux. Artificially altering the total interaction cross section has previously been shown to have minimal impact on the neutrino flux within the energy range considered and its uncertainty within TARGET2.4 is considered subdominant to the uncertainty on pion and kaon production [44]. As such it will not be considered any further within this work.

Conversely the multiplicity will have a significantly greater effect on the flux with changes in the multiplicity directly affecting the number and kinematics of the mesons created in each interaction which will directly affect the neutrino flux.




Figure 7.1: Cartoon of an interaction of a cosmic ray on an atmospheric nuclei resulting in neutrino production. The relevant hadron production interaction for neutrinos marked 1 is $pN \rightarrow \pi^+ + X$ and for neutrinos marked 2 $pN \rightarrow \pi^- + X$

7.1.1 Tuning the Hadron Production Multiplicities

To tune the atmospheric neutrino flux prediction a method of weighting the multiplicity of hadronic interactions is used. In this method the atmospheric neutrino Monte Carlo is run with the unaltered TARGET2.4 generator used to simulate hadronic interactions. For each neutrino which contributes to the flux the full chain of inclusive hadron production interactions preceding it is stored. For example, in Figure 7.1 for the neutrinos marked ‘1’ the hadronic interaction $pN \rightarrow \pi^+ + X$ is stored. For each of these interactions a weight can be calculated:

$$W = \frac{dn_h^{\text{data}}}{dx_F dP_T} / \frac{dn_h^{\text{TARGET2.4}}}{dx_F dP_T} \quad (7.2)$$

as the ratio between the multiplicity of the interaction from the models of chapters 5 and 6 based on the data, to the multiplicity of the TARGET2.4 generator. These weights are then applied to each of the neutrinos when they are combined to produce

¹assuming an average atmospheric depth of 1000 g/cm²[87] for down-going protons and an approximate proton interaction length of 100 g/cm² at the energies relevant to 0.1-10 GeV neutrino flux. This will be greater for protons travelling closer to the local horizontal.

the flux. Where a single neutrino is the product of multiple hadronic interactions the product of the weights of each interaction is used. This method of weighting multiplicities has previously been used by beam neutrino flux calculations [78].

For a toy example to demonstrate how this method works, let us say the TARGET2.4 generator is under-producing π^+ in a specific x_F - P_T bin at a specific interaction energy. This will have the effect of underestimating the neutrino flux for the neutrinos produced in the decay of these pions. By weighting up these neutrinos by a factor which corrects for this underproduction of pions by the generator, the final neutrino flux has been tuned to this new value of pion production from the data.

This tuning method has the advantage of not needing to re-run the Monte Carlo simulation each time a new hadron production model is used, saving significant computational time. The next section describes the process used to obtain the numerator in equation 7.2 from the models previously developed.

7.1.2 Multiplicity on Air from the Datasets

For the purposes of this section the air is assumed to be comprised of 79% nitrogen and 21% oxygen by number of atoms ². Since the Monte Carlo does not distinguish between interactions on nitrogen and oxygen the mean multiplicity for interactions on air is needed. To start, we consider the probability of an interaction occurring on either nitrogen or oxygen as the probability of a randomly selected nucleus being either nitrogen or oxygen multiplied by the probability of an interaction occurring on that species. The probability of an interaction occurring on a specific nucleus is proportional to the total production cross section. Therefore, the probabilities can be written:

$$P(O) \propto 0.21 \cdot \sigma_{\text{prod}}^O \quad (7.3) \quad P(N) \propto 0.79 \cdot \sigma_{\text{prod}}^N \quad (7.4)$$

In effect this is saying that although oxygen comprises of 21% of the nuclei in the atmosphere, it is responsible for marginally more than 21% of atmospheric interactions as it has more nucleons than nitrogen and a larger production cross section.

²Assuming within the ideal gas law that the volume is proportional to the number of molecules

The multiplicity for interactions on air can then be written as the sum of the multiplicities on oxygen and nitrogen weighted by their respective probabilities:

$$\frac{dn_{\text{air}}}{dx_F dP_T} = \frac{0.21\sigma_{\text{prod}}^O \frac{dn_O}{dx_F dP_T} + 0.79\sigma_{\text{prod}}^N \frac{dn_N}{dx_F dP_T}}{0.21\sigma_{\text{prod}}^O + 0.79\sigma_{\text{prod}}^N}, \quad (7.5)$$

then applying equation 7.1:

$$\frac{dn_{\text{air}}}{dx_F dP_T} = \frac{0.21 \frac{d\sigma_O}{dx_F dP_T} + 0.79 \frac{d\sigma_N}{dx_F dP_T}}{0.21\sigma_{\text{prod}}^O + 0.79\sigma_{\text{prod}}^N} \quad (7.6)$$

The multiplicity for interactions on air is obtained.

The double differential cross sections for nitrogen and oxygen can be obtained using the target nucleus scaling of Section 5.2. The NA61 and NA49 experiments which cover pion production energy regions 6 and 7 and kaon production energy regions 2 and 3 both provide a value for the total production cross section on carbon. The total production cross sections are assumed to scale via the $A^{\frac{2}{3}}$ scaling and this is used to calculate the total production cross sections on oxygen and nitrogen to calculate the multiplicities. The experiments used to model regions 1-5 for pions and Region 1 for kaons do not publish a value for this cross section and instead the prediction for the total production cross section on air from the TARGET2.4 generator is used.

7.1.3 Weights from TARGET2.4 to Data Ratios

Having obtained the estimated multiplicities based on the models from the datasets, it is now possible to produce the weights comparing these models to TARGET2.4. The results for weights to $pN \rightarrow \pi^+ + X$ and $pN \rightarrow \pi^- + X$ are shown in figures 7.2 and 7.3 with isospin symmetry used to weight incident neutron interactions. Within the phase space regions that contribute most significantly to the flux (recalling figures 5.19-5.21) in the positive x_F low P_T region, weights indicate the dataset model and generator are typically within $\sim 20\%$ of each other with larger weights appearing in regions where there is lower pion production which have smaller contributions to the flux. There are also large weights in the lowest interaction momentum (3 GeV) bin.




Figure 7.2: Weights for π^+ production as the ratio between multiplicities from the dataset models to the multiplicities used by that TARGET2.4 generator. Weights are shown in the produced pion x_F and P_T space for interactions with an incident proton at each of the 7 interaction energies used to model pion production from Chapter 5. Weights are only shown within the x_F - P_T phase space coverage of the datasets.




Figure 7.3: Weights for π^- production as the ratio between multiplicities from the dataset models to the multiplicities used by that TARGET2.4 generator. The labels are the same as Figure 7.2

Figure 7.4 shows the weights for the production of charged and neutral kaons by incident protons. The weights are also obtained for incident neutron production and found to be highly similar to the weights for incident proton production so are not shown again. This indicates that TARGET2.4 is using similar models to those discussed in Chapter 6 to relate the kaon production by incident neutrons and protons. It is of note that TARGET2.4 is significantly under-producing positive kaons in the higher P_T (higher angular region) with large weights between 1.5 and 2 assigned.




Figure 7.4: Weights for K^+ (top), K^- (middle) and K_L^0 (bottom) production by incident interactions by incident protons on air as the ratio between multiplicities from the dataset models to the multiplicities used by the TARGET2.4 generator. Weights are shown in the produced kaon x_F and P_T space for interactions with an incident proton at each of the 3 interaction energies used to model kaon production from Chapter 6. The weights for K_S^0 production are not shown as they are found to be the same as K_L^0 indicating that as with the model in Chapter 6, TARGET2.4 assumes their production is the same. Weights are only shown within the x_F - P_T phase space coverage of the datasets.




Figure 7.5: Results of weighting the neutrino fluxes using weights of Section 7.1.3 weighting only inside the region of dataset coverage. The ratio of the weighted to unweighted fluxes is shown for (left to right, top to bottom) ν_μ , $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$ fluxes. Outside the regions of phase space coverage the original TARGET2.4 generator is used. Shown as dashed lines the weighting to only the pion production and only kaon dataset models separately. Shown as a function of neutrino energy between 0.1-10 GeV

7.1.4 Weighting the Flux with Data Model Weights

The neutrino fluxes obtained when applying the weights from Section 7.1.3 are shown in Figure 7.5. The weight for an interaction occurring between two energy regions is linearly interpolated in $\sqrt{s}$ between the weights at both regions. Note that at this stage the weighting is only done for the interactions covered by dataset phase space. Since a smaller fraction of the lower energy neutrino fluxes is covered by the data ($\sim 50\%$ - see Figure 5.22) a smaller fraction of the fluxes have been weighted at this point. Also shown are the results when weighting separately to only the pion or kaon weights. The kaon weighting is seen to contribute to an increase in neutrino flux across the full energy range due to the increase in kaon production suggested by the datasets when compared with TARGET2.4 in the higher P_T region.

The increase in neutrino flux at the lower neutrino energies is due to an increase in pion production seen in the region covered by the high angle datasets (the negative x_F region) when compared with TARGET2.4 in figures 7.2 and 7.3.

An estimate for the production of pions and kaons occurring outside the phase space of the dataset measurements is discussed in Chapter 8 with the systematic uncertainties on the neutrino flux due to the datasets described in Chapter 9.

8

Tuning Hadron Production Outside the Dataset Phase Space

This chapter explores the methods for estimating the production cross section outside the x_F - P_T daughter phase space coverage of the experimental datasets. In Chapter 5 neutrinos from pions produced outside the dataset phase space were estimated to contribute to $\sim 35\%$ of the neutrino flux at lower (<1 GeV) energies dropping to $< 10\%$ at higher (>2 GeV) neutrino energies. Similarly, Chapter 6 demonstrated that neutrinos from kaons outside dataset phase space contributed $< 5\%$ to the total neutrino flux, mostly due to the smaller number of neutrinos coming from kaon decays. This already represents a substantial improvement of the phase space coverage that was available for the 2004 Bartol Monte Carlo [47]. Within this chapter, the prediction of pion and kaon production by different Monte Carlo hadron production generators is explored alongside the TARGET2.4 generator. Comparisons between these generators and the datasets highlight which generators may perform better at interpolating outside the dataset phase space. The average value of these generators is then used to tune pion and kaon production within these phase space regions. This weighting is shown to have a $\sim 2\%$ effect on the neutrino flux prediction. This chapter considers interactions outside the dataset x_F - P_T phase space but within the interaction energy range where pion and kaon interactions are tuned. Interactions occurring above and below experimental dataset energies are treated separately in Section 9.4.4

8.1 Meson Production Outside Dataset Phase Space Coverage

Initial attempts to estimate pion and kaon production outside the dataset phase space considered extrapolations of the fits to the extended BMPT parameterisation from Section 5.3. These were found to work poorly for two reasons. Firstly, while the fits at different energy regions tended to vary by only modest amounts within the phase space coverage of the data, the extrapolated regions had much larger variations (by several orders of magnitude). The greatest variations were in the lowest energy regions (beam momenta 3-10 GeV) where the dataset coverage is smallest. Secondly, as discussed in Section 5.5.1, joint fits to the HARP high and low angle data proved unsuccessful. Fits to the high and low angle data separately similarly had large variations in the region between the dataset coverage and it was unclear which (if either) parameterisation provides the most physical extrapolation into that region.

It is possible that a combined fit of pion and kaon production data such that charge and energy are conserved in the extrapolated region as a constraint would provide a better result. This is beyond the scope of this work and would additionally require an alternative parameterisation to correctly describe both the high and low-angle data. As an alternative to extrapolations of dataset BMPT fits, the production of pions and kaons by hadron interaction generators are explored within the extrapolated region. Much like TARGET2.4, these hadron interaction generators simulate interactions using Monte Carlo methods such that energy, momentum and charge are conserved and prevent the unphysical extrapolations that were seen with the BMPT parameterisation.

8.2 Hadron Interaction Generators

This section considers the hadron interaction modelling within both the Geant4 and Fluka Monte Carlo simulations. Both of these software packages are widely used in a variety of contexts for simulating the propagation and interactions of particles. While these packages have various features, this analysis will focus on




Figure 8.1: Underlying models used by different hadron production generators at different energies to simulate pion and kaon production in nucleon nucleus interactions. The approximate overlap region between two generators is shown separately as the kinetic energy of the beam particle. A summary of the meanings of the interaction models is found in the main text. Information on the Geant4 physics list models is taken from [88–91] and the FLUKA simulation from [92]

their modelling of hadron production in proton-nucleus interaction. Within these simulations, different models are used to simulate hadronic interactions at different energies. This section briefly summarises these models.

The Geant4 package [93–95] has the option to choose from a variety of different models and processes used to generate nucleon-nucleus interactions. These choices are provided through different pre-defined physics lists which specify which interaction models are used at each energy. The physics lists used within this analysis are the FTFP_BERT, QGSP_BERT, FTF_BIC [90] and NuBeam [89]. A summary of the models used by each of the physics lists is shown in Figure 8.1. The low energy region is modelled by an intranuclear cascade with either the Binary (BIC) or Bertini (BERT) cascade model. Within these models, an interacting hadron enters the nucleus and interacts within the density of the nuclear material [96]. String models describe the high energy region with either the Fritof (FTFP) or

the quark-gluon string model (QGSP). These model the hadronic interactions in terms of their constituent quarks with the colour field between quarks forming QCD strings [96, 97]. In addition to these models, de-excitation models are used, which handle the excited nucleus after an interaction has occurred. The QGSP_BERT physics list also uses Geant4 low energy parameterisation model (LEP) which unlike the other models discussed here is not a theory-based model but based on a parameterisation of datasets.

Fluka version 2011.2x6 [98, 99] is used within this analysis. A summary of the models used by Fluka is also shown in Figure 8.1. The low energy region is modelled by the pre-equilibrium cascade model (PEANUT) and the higher energy region is modelled by Gribov-Glauber multiple scattering, which incorporates the dual parton model [92]. Similar to the FTFP and QGSP models in Geant4 the dual parton model is a string model, modelling partons as quarks held by the gluon-gluon interaction which forms a QCD string [100].

In the following section, comparisons will be drawn between these generators, the TARGET2.4 generator and the datasets. Since there is significant overlap between the physics lists used by Geant4 they will provide highly similar flux predictions in the energy regions where the underlying models overlap. For this reason, we chose different physics lists for comparison in each energy region such that the underlying physics models are as independent as possible. Table 8.1 shows the physics lists chosen for comparison to data at each energy region.

8.3 Pion Production Comparisons Between Generators and Data

The Geant4 and Fluka packages are used to simulate interactions on a thin fixed carbon target. The results of these simulations are then used to determine the multiplicities of pions and kaons across x_F - P_T space at the interaction energy regions in Table 5.6.




Figure 8.2: Comparisons between Geant4 generators (solid), TARGET2.4 (dashed) and data points for pion production at low angles. π^+ (left) and π^- (right) multiplicities for the lowest polar angle bin for each dataset (50-100 mrad for the HARP 5 and 12 GeV datasets and 0-10 mrad for the NA61 dataset) are compared to both the TARGET2.4 and Geant4 generators. Also shown are the BMPT fits (dotted) discussed in sections 5.4 and 5.5. For each graph, the pion momentum is plotted within the specific pion angular bin against the invariant multiplicity $E \frac{d^3n}{d^3p}$ on the y axis.

Energy Region	$\sqrt{s}$ (Beam Momentum)/ GeV	Generators for comparison
1	2.77 (3.0)	FTF_BIC, NuBeam, Fluka
2	3.54 (5.7)	FTF_BIC, NuBeam, Fluka
3	4.11 (8.0)	QGSP_BERT, NuBeam, Fluka
4	5.17 (13.3)	QGSP_BERT, NuBeam, Fluka
5	6.38 (20.7)	QGSP_BERT, NuBeam, Fluka
6	7.74 (31.0)	(QGSP_BERT), NuBeam, Fluka
7	17.27 (158.0)	FTFP_BERT, NuBeam, Fluka

Table 8.1: Summary of generators chosen to model each energy region of hadron production. For each of the pion production energy regions (defined before in 5.6), the generators listed are used to estimate pion production outside dataset phase space. The generators are initially chosen to minimize the overlap of their underlying models (see Figure 8.1). The average value of the generator predictions is calculated and used within the neutrino flux calculation to estimate pion production outside the dataset phase space. The QGSP_BERT in region 6 is found to be in poor agreement with the data (see Section 8.3) and as such only NuBeam and Fluka are used in this region to calculate the neutrino flux.

The aims of this section are to judge how well each generator may extrapolate into the phase space not covered by data. This is done by comparing their multiplicities in the region where there is data¹. This begins by reconsidering plots 5.19 - 5.22 from Chapter 5 which show the coverage of the relevant pion production x_F - P_T phase space overlaid with the coverage of datasets. Pion production not covered by datasets mainly arises at the lower interaction energy regions where, despite combining multiple datasets, there are still areas of phase space not covered by any experiment. The contribution of these interactions outside the dataset coverage

¹Direct comparisons of Fluka results with the datasets are not permitted by Fluka licence conditions so only Geant4 results will be shown here




Figure 8.3: Comparisons between Geant4 generators (solid), TARGET2.4 (dashed) and data points for the lowest P_T Bin of 158 GeV NA49 pion production data. π^+ (left) and π^- (right) multiplicities for the 0.05 GeV P_T bin of the NA49 158 GeV data are compared to Geant4 and TARGET2.4 Hadron production generators. Also shown is the BMPT fit (dotted) discussed in Section 5.4

has the most significant effect on neutrino fluxes below 2 GeV. We group these regions into 3 groups to consider separately.² These groups are pion production at small angles below dataset coverage, pion production in the central region between the HARP high and low angle datasets and pion production at high angles above dataset coverage.

8.3.1 Pion production at Small Angles

Figures 5.19 and 5.20 in Chapter 5 demonstrated that for both 0.1 and 1.1 GeV neutrino fluxes there are significant contributions from pion production at angles below dataset coverage (the small P_T region). In addition, Figure 5.21 shows a smaller but still significant contribution from pion production below the lowest P_T bin of the NA49 data. Comparisons between generators in Table 8.1, the TARGET2.4 generator, and the data in the lowest angular bins of HARP at 5 and 12 GeV and NA61/SHINE at 31 GeV for pion production are shown in Figure 8.2. The same is shown for the lowest P_T bin of NA49 data in Figure 8.3.

²As plots 5.19-5.21 are in x_F - P_T phase space Figure 5.3 is useful to relate this to the p - θ phase space bins of some of the experimental datasets.

In each of these cases, it is seen that as pion momentum approaches zero (or as x_F approaches 0) the multiplicity of TARGET2.4 sharply increases in contrast to both the data and the other generators. The other generators typically produce multiplicities closer to the datasets than TARGET2.4 with the exception of the QGSP_BERT Geant4 generator at 31 GeV which also diverges substantially from the data. This indicates that (with the exception of QGSP_BERT at 31 GeV) these generators will likely provide a more physical extrapolation into the regions not covered by data than TARGET2.4.

8.3.2 Pion Production in the Central Region

At low interaction energies, there is a contribution to the neutrino flux from pions produced in the phase space region between the HARP high and low angle datasets and also at a higher momentum than the high angle dataset. Figure 8.4 shows a comparison between generators and data on this boundary. The lowest angle of the 5 GeV HARP high angle dataset is shown alongside the highest angle of the low angle data. Figure 8.5 is the equivalent plot for the 12 GeV data. Once again, TARGET2.4 is seen to diverge from the datasets and other generators at low momenta - this is especially noticeable for the 12 GeV data. None of the generators reproduce the π^+ production at 450 mrad well however the Geant4 generators are marginally closer to the data points than TARGET2.4.

8.3.3 Pion Production at High Angles

Finally, we consider pion production at high angles. Figure 8.6 shows a comparison between the highest (390 mrad) 31 GeV NA61 angular bin and generators and also the (1250 mrad) angular bins of the 5 and 12 GeV HARP high angle data. Pion production at higher angles than the NA61 data are seen in Figure 5.19 to contribute to the neutrino flux at low energies. The comparisons with the HARP high angle data are included to consider how the generators behave as the pion production angle increases further; however, this phase space region has a smaller contribution to the neutrino flux. None of the generators reproduces the data




Figure 8.4: Comparisons between Geant4 generators (solid), TARGET2.4 (dashed) and data points for pion production between 200 and 550 mrad at 5 GeV energy of the HARP Data. π^+ (left) and π^- (right) multiplicities are shown for the 200-250 mrad bin of the HARP low angle datasets and the 350-550 mrad angle bin for the HARP high angle datasets. This is done to investigate how well the generators may extrapolate into the region between these two datasets. Also shown is the BMPT fit (dotted) discussed in Section 5.5

exactly however it is seen that as the angle increases the TARGET2.4 generator significantly underproduces pions when compared to the HARP data. The other generators perform better in this region.

8.4 Kaon Production Comparisons Between Generators and Data

Having compared the production of pions with generators to data, attention is now turned to the smaller kaon production component. Using the same methodology




Figure 8.5: Same as Figure 8.4 for 12 GeV beam momentum data.

as Section 8.3, kaon production within 3 angular bins of the 31 GeV NA61 data is compared to the generators. The results are shown in Figure 8.7. It can be seen that, similar to pion production at this energy in Figure 8.2, the Geant4 QGSP_BERT generator does not match the data well with large divergences at the low angle low momentum region. Similarly, TARGET2.4 also overproduces kaons at low angles and underproduces at high angles. In contrast, the Geant4 NuBeam generator reproduces kaon data much better.

8.5 Conclusions and the Weighting of the Interactions

Sections 8.3 and 8.4 have demonstrated that the Geant4 physics lists are typically in better agreement with the data than TARGET2.4. The Fluka simulation was




Figure 8.6: Comparisons between Geant4 generators (solid), TARGET2.4 (dashed) and data points for pion production at large angles. π^+ (left) and π^- (right) multiplicities are shown for the 1150-1350 mrad bin of the HARP high angle data (5 GeV top and 12 GeV middle). The highest angle bin 360-420 GeV of the 31 GeV NA61 data is also shown.




Figure 8.7: Comparisons between Geant4 generators (solid), TARGET2.4 and the NA61 data for kaon production at 31 GeV. K^+ (left) and K^- (right) multiplicities are shown for the 0-20 mrad bin (top), 100-140 mrad bin (middle) and 180-240 mrad bin (bottom) compared against data from NA61.




Figure 8.8: Weights for π^+ production outside the phase space coverage calculated by taking the ratio of generator average multiplicities to the TARGET2.4 multiplicity as described in Section 8.5

separately found to be in better agreement with the data also. TARGET2.4 has shown large divergent behaviour from the data. This is particularly noticeable in the low x_F and low P_T region where TARGET2.4 overproduces pions and kaons compared to data and in the regions of high angle pion production where TARGET2.4 underproduces pions. It is therefore reasonable to assume that these generators provide a more physical estimate of pion and kaon production outside the phase space coverage of the datasets than TARGET2.4. TARGET2.4 will accordingly be weighted using these generator predictions. The exception is the




Figure 8.9: Weights for π^- production outside the phase space coverage calculated by taking the ratio of generator average multiplicities to the TARGET2.4 multiplicity as described in Section 8.5

QGSP_BERT model which was shown to have divergent behaviour compared to the data in the 31 GeV region and will not be used.

For each interaction energy region, the average value of the multiplicities of the generators listed in Table 8.1 are calculated. These average multiplicities are scaled from carbon to air interactions as described in Section 7.1.2. A ratio of this mean multiplicity to the TARGET2.4 multiplicity is used to tune the interactions which occur outside dataset coverage in the neutrino flux simulation via the methods discussed in Chapter 7. These multiplicity ratios for π^+ and π^- production are




Figure 8.10: Weights for K^+ (top), K^- (middle) and K_L^0 (bottom) production by incident protons outside phase space coverage of the datasets. Only the region of positive rapidity is shown as target does not produce kaons in the backward going direction as described in Section 8.5

shown in Figures 8.8 and 8.9. Multiplicity ratios for kaon production are shown in Figure 8.10. It should be noted that TARGET2.4 does not produce kaons in the region of rapidity $y < 0$ (i.e. kaons produced backwards in the laboratory frame). The contribution to the neutrino flux from this region is expected to be negligible and so this region is not weighted.

These weights are applied to the calculation of the neutrino flux. The neutrino flux tuned to both the dataset models of the previous chapters and the model




Figure 8.11: The neutrino fluxes with pion and kaon production interactions are weighted both inside dataset coverage (where they are weighted to the dataset model) and outside dataset coverage (where they are weighted to the mean generator multiplicities described in this chapter). Shown separately are the fluxes weighted only outside the region of dataset coverage (dashed red), only inside the dataset coverage which is the same as Figure 7.5 (dashed green) and the combined weighting to the whole x_F - P_T phase space region (solid orange). In each case the ratio of the weighted to unweighted fluxes is shown for (going left to right then top to bottom) ν_μ , $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$ fluxes.

based on generators described within this chapter is shown in Figure 8.11. It can be seen that the tuning of interactions outside the phase space alone has only a small $\sim 2 - 3\%$ effect. This figure shows the neutrino fluxes with all pion production interactions with momentum between 3-450 GeV and kaon production interactions between 15-158 GeV now tuned. Overall, the weighting changes the flux by between 3 and 8% well within the original 15% uncertainty previously assigned to the original flux estimate (see Section 3.4.1).

9

Atmospheric Neutrino Flux Uncertainties due to Hadron Production

This chapter considers the uncertainty on the neutrino flux estimate due to the uncertainty in hadronic interactions in atmospheric showers. This will involve considering the uncertainties on the hadron production datasets and the estimate of the hadron production cross section outside the dataset phase space. In many cases, the uncertainty and correlation between the sources of uncertainties are not precisely known. In each case, this involves justifying the approach taken to the treatment of uncertainties with the aim to produce the best estimate of the uncertainty on the neutrino flux prediction with the available information.

Having tuned the multiplicity of pion and kaon production in Chapter 7 and 8 the next stage is to estimate the uncertainty of these multiplicities and propagate this through to the neutrino flux. In doing so, this chapter will demonstrate that there is a decrease in the uncertainty of the neutrino flux prediction compared to the original Bartol 2006 estimate now that the dataset tuning has been incorporated. This chapter considers the uncertainties on meson production datasets in Section 9.1 and then the uncertainty on the estimate of meson production outside the dataset phase space in Section 9.2. Finally, subdominant uncertainties such as the uncertainty due to the interpolation between meson production at different energies

and the target nucleus scaling process are considered in Section 9.4. Particular care is taken to consider the effects of correlation between sources of uncertainty and to choose the approach which gives the most conservative uncertainty estimate where the effect of correlations is unclear.

9.1 Uncertainties in Hadron Production Datasets

This section starts by considering the uncertainty on the NA61 and NA49 pion production datasets, before extending this approach to the lower energy pion production datasets where the uncertainty inflation of Section 5.5.2 is applied. The uncertainty of the HARP High Angle datasets is considered separately in Section 9.1.3. An initial summary of dataset uncertainties was given in Section 5.4.1. This section takes this further by applying these uncertainties to the neutrino flux. It will be assumed within this section that each experiment gives an independent measurement of the hadron production cross sections however the effects of correlation between measurements given by the same experiment will be discussed.

9.1.1 NA61 and NA49 Pion Production

The uncertainties in the measurements of hadron production (binned across the two-dimensional daughter hadron phase space) are given in three forms:

- Statistical uncertainties on each hadron production invariant cross section bin due to the number of hadrons observed in that bin.
- Systematic uncertainties on each cross section bin due to corrections, particle identification and the track reconstruction process.
- The uncertainty on the normalisation of the production cross section measurements. These uncertainties affect each cross section bin equally and arise from uncertainties on the number of interactions that have taken place in the target. This was seen in Chapter 4 to be principally due to the uncertainty

on the number of protons incident on the target and the uncertainty of the target thickness.

Understanding the correlation of the sources of these uncertainties across each of the bins within a single experiment will be important to understanding their effect on the neutrino flux. The normalisation uncertainty is by definition correlated across all bins and each meson production measurement (e.g. correlated between $\pi^\pm$ and $K^\pm$ production). The bin-by-bin statistical uncertainty is likewise not correlated between bins as it arises from uncorrelated statistical effects. It is less clear how to treat the correlation of the bin-by-bin systematic uncertainties across daughter particle phase space as while both NA61 and NA49 imply that these uncertainties are correlated, the exact value of this correlation in the form of a covariance matrix is not given. The next sections will discuss each of these uncertainties in turn and explore the effects of different correlation models on the neutrino flux uncertainty.

Bin-by-bin Systematic Uncertainties

To investigate the effect of the bin-by-bin systematic uncertainty on the flux, we start by considering the NA49 π^+ production dataset. A toy covariance matrix for each of the dataset bins is constructed with the correlation between two bins' systematic uncertainties given by:

$$\rho_{ij} = A^{d(i,j)}, \quad (9.1)$$

where ρ_{ij} is the correlation between bins i and j , $d(i,j)$ is the separation between bins i and j (measured in number of bins in the x_F - P_T plane) and A is a parameter $0 \leq A \leq 1$ which determines the strength of the correlation between neighbouring bins. In this scheme, A equal to 1 would result in the complete correlation of systematic uncertainties while A equal to 0 would result in no correlation between systematic uncertainties. The correlation matrix was used to draw toy samples of the datasets using the Cholesky decomposition method [101] for 4 values for the A parameter ($A = 0.0$, $A = 0.5$, $A = 0.9$, $A = 1$). An example of a perturbation




Figure 9.1: Example of perturbations to the NA49 π^+ production data for different values of the correlation parameter in equation 9.1. Shown is the ratio of a systematic perturbation to the original unperturbed cross section in the x_F - P_T plane. The value of A changes the correlation of the systematic uncertainty between bins and thus the perturbed cross section.

for each value of the A parameter is shown in Figure 9.1 where it can be seen how the larger the A parameter the greater the correlation between the bins in the perturbation across the x_F - P_T plane. 100 sample datasets are drawn from each correlation model, each of which is fit to the BMPT parameterisation and used to calculate a neutrino flux. The uncertainty on the flux is then calculated from the variance between each of these toy fluxes.

Figure 9.2 shows the uncertainty obtained via each of the correlation models. This demonstrates that maximal correlation between the bin-by-bin systematic uncertainties results in the greatest uncertainty on the neutrino flux and is therefore the most conservative approach.

While computing a large number of sample datasets and using it to re-weight the flux is the most accurate method to estimate the uncertainty, it is computationally intensive. Also shown in Figure 9.2 is a simpler method for estimating the maximally correlated bin-by-bin systematic where the cross section is shifted up and down by




Figure 9.2: Uncertainty on the ν_μ flux due to the systematic uncertainty on the NA49 π^+ production data assuming different correlation models. Also shown is an approximate method to estimate the uncertainty when maximally correlated described in the main text.

1σ (the systematic error) in each bin and both shifted differential cross sections are used to obtain a flux estimate. The two re-weighted fluxes are then used to obtain an uncertainty. As can be seen, the uncertainty obtained via this simpler method is consistent with the more computationally intensive method and will be used going forward.

In addition to the correlation of systematic uncertainties across bins in the x_F - P_T plane, the possibility of the bin-by-bin systematic uncertainties being correlated between measurements of different meson production must be considered (e.g. the extent to which bins for π^+ production in the NA49 are correlated with π^- production). It is less likely that there is a significant correlation between the production of oppositely charged meson species due to the use of different analysis techniques and different particle identification methods used for different meson species. Additionally, overestimating the correlation between the production of different meson species will have the unintended consequence of overestimating the correlation of uncertainties between different flavour neutrino fluxes. This will cause the uncertainty of the flux flavour ratios to be underestimated (discussed in Chapter 2). For example the ν_e to $\bar{\nu}_e$ ratio depends strongly on the π^+ to π^-




Figure 9.3: Uncertainty on the $\nu_e/\bar{\nu}_e$ flux ratio due to the bin-by-bin systematic uncertainties on the NA49 dataset assuming different correlation between π^+ and π^- production uncertainties. The black dashed line shows the uncertainty when $\pi^\pm$ production uncertainties are uncorrelated and the green line shows uncertainties when $\pi^\pm$ production uncertainties are maximally correlated.

production ratio. Overestimating the correlation between the π^+ and π^- production will result in an underestimate of the uncertainty on the $\pi^\pm$ production ratio and therefore underestimate the uncertainty on the ν_e to $\bar{\nu}_e$ ratio.

This last point is highlighted in Figure 9.3 where the uncertainty on the ν_e to $\bar{\nu}_e$ ratio is shown for maximally correlated π^+ and π^- production and uncorrelated π^+ and π^- production. While a negatively correlated π^+ and π^- production would yield a greater still uncertainty on the ν_e to $\bar{\nu}_e$ ratio there is no motivation from the analyses of the production measurements which suggests a significantly negative correlation of systematic uncertainties.

Based on these results, the bin-by-bin systematic uncertainties of the production of a single meson species will be taken to be maximally correlated across the x_F - P_T plane and uncorrelated between the production of different mesons. Within this framework, the fractional systematic uncertainties of the neutrino flux are shown in Figure 9.4.




Figure 9.4: Uncertainty on the ν_μ and $\nu_e/\bar{\nu}_e$ flux ratio due to the uncertainty on the NA61 and NA49 datasets.

Normalisation Uncertainties

This section considers the propagation of the normalisation uncertainty to the neutrino flux. The neutrino flux is weighted according to the multiplicity rather than the production cross section (see equation 7.1) however recalling the discussion from Section 4.3.1 it can be seen that the effects which contribute to the uncertainty on the total production cross section will correlate with the normalisation uncertainty on the differential cross section measurements. As such the normalisation uncertainty on the production cross section will be taken as an estimate for the upper limit of the normalisation uncertainty of the multiplicity.

The normalisation uncertainty is propagated through to the neutrino flux via the methods described above and shown in Figure 9.4.

Bin-by-bin Statistical Uncertainties

Finally, an alternative method is developed to propagate the uncorrelated bin-by-bin statistical uncertainties. This method constructs the weights for pion production within the differential cross section binning scheme used by each experiment. This is as an alternative to using the BMPT parameterisation to weight hadron production smoothly across the two-dimensional phase space. For each bin, the weight is the




Figure 9.5: Uncertainty on the ν_μ neutrino flux due to statistical uncertainties on the NA61 (left) and NA49 (right) datasets. In addition, the systematic uncertainty (solid line) is shown for comparison. It can be seen that the statistical uncertainty (dashed line) is negligible in comparison.

ratio of the multiplicity as measured by the experiment to the multiplicity of the TARGET2.4 generator. 100 toy datasets are generated from each of the NA61 and NA49 pion production datasets in each case perturbing the values of the differential cross section within each bin by their statistical uncertainty. Each of these toy datasets is used to produce weights via the method above which are used to calculate a neutrino flux, the uncertainty is calculated as the variance of these fluxes.

The propagated bin-to-bin statistical uncertainties are shown in Figure 9.5. As can be seen, these uncertainties are much smaller than the correlated systematic uncertainties. Similar to the considerations of the correlation between systematic errors, the effect of the uncorrelated fluctuations is greatly reduced when integrating over pion production within the x_F - P_T plane to produce a neutrino flux. When considering the total uncertainty due to the NA61 and NA49 datasets the statistical effect is negligible and will not be considered further.

9.1.2 Lower Energy Dataset Pion Production

In Section 5.5.2 it was demonstrated that the uncertainties assigned to the lower energy datasets were smaller than the disagreements between the datasets. A method of inflating the assigned uncertainties to cover the discrepancy between the

datasets was developed and the total bin-by-bin and normalisation uncertainties were increased by factors defined in Table 5.7. Due to the inflation process and the differences between the types of bin-by-bin uncertainties presented by each dataset, the total bin-by-bin uncertainty will be treated as a systematic uncertainty when propagating it through to the neutrino flux. Following the conclusions of the previous section (9.1.1) this bin-by-bin uncertainty will be treated as correlated across the x_F - P_T plane within a single experimental measurement but not correlated across the production of different meson species.

In addition, the HARP and E910 datasets present cross section measurements at a number of different beam momenta. Since the same detector and analysis method are used at each beam energy, it is reasonable to assume a significant correlation between these measurements. Treating each measurement as independent would underestimate the uncertainty on the neutrino flux. Therefore the bin-by-bin systematic uncertainty will be treated as correlated for the production of the same meson by the same experiment (e.g. the bin-by-bin uncertainties for π^+ production by the HARP experiment at 3,5,8 and 12 GeV beam momenta). The normalisation uncertainty is treated as correlated between the production of all mesons at a single beam energy but not between beam energies. This is because the most significant uncertainty is due to the uncertainty on the number of beam protons which hit the target and will not be correlated between runs at different momenta (see for example Ref. [67]). Similar to the approach taken with the NA61 and NA49 data, the normalisation uncertainty which is given for the cross section measurement is taken as an upper limit on the normalisation of the multiplicity.

In keeping with the approach taken in Section 9.1.1, the inflated uncertainty on the production cross section will be assumed to be uncorrelated between different experiments (i.e the 5 GeV HARP and 6.4 GeV E910 pion production measurements will be assumed independent and their uncertainties uncorrelated).

To propagate these systematics through to the neutrino flux, the effect of each systematic is considered in turn and used to perturb the data via the same method




Figure 9.6: Uncertainty on the ν_μ flux due to the uncertainty on the <30 GeV datasets. The left-hand plot is the bin-by-bin uncertainties where the π^+ and π^- uncertainties are added in quadrature for each experiment. The right-hand plot shows the correlated normalisation uncertainty.

as Section 9.1.1. In each case, the combined dataset fit is rerun and used to re-weight the neutrino flux. The variance on the neutrino flux for each systematic perturbation is taken as the uncertainty.

Figure 9.6 shows the uncertainties due to these lower energy datasets. In both cases the HARP and E910 uncertainties dominate. This is not surprising as it is a reflection that these measurements cover the widest range of beam momenta. It should not be interpreted that these datasets have the largest individual uncertainties. The total uncertainty shown in Figure 9.6 is calculated as the quadrature sum of the components shown in each plot, which (as explained above) are treated as independent.

9.1.3 HARP High Angle Data

Understanding the uncertainty on pion production at large angles covered by the HARP high angle data presents a unique challenge due to the different cross sections obtained in the two different analyses of the same data caused by disagreements within the experiment on how best to analyse the TPC data after hardware operating failures (see Section 4.3.4). In addition to this, the two results are binned in different




Figure 9.7: HARP CDP and HARP OH pion production invariant differential cross section comparison between analyses at 5 GeV beam momentum. Shown for $\pi^+(\pi^-)$ production left(right) at ~ 436 mrad(~ 610 mrad) top(bottom) demonstrating the differences between the analyses.

binning schemes making a direct comparison of these data points difficult. The two different analyses are the analysis by the official HARP collaboration (referred to as the OH analysis [65]) and the analysis by the CDP breakaway group (known as the CDP analysis [66]). Until this stage the OH analysis has been used to weight the flux, this section will consider the systematic effect due to the differences between the two analyses. Unlike the discrepancies between datasets discussed earlier, these different analyses are not independent datasets as they use the same underlying data.

The BMPT parameterisation described in Section 5.3.2 is separately fit to the HARP OH and HARP CDP analysis. Figure 9.7 shows both of these datasets where there is an approximate overlap of angular bins. Figure 9.8 shows the ratio of the BMPT parameterisation to each analysis in the x_F - P_T plane. Both of




Figure 9.8: Ratio of HARP OH to HARP CDP pion production cross section fits in x_F - P_T plane. Shown for each of the 4 beam momenta 3,5,8,12 GeV for π^+ (π^-) production top (bottom).




Figure 9.9: Neutrino flux weighted to the HARP OH and HARP CDP analyses. Shown as the ratio to the fluxes without weighting hadronic interactions. The absolute ν_μ flux left and the $\nu_e/\bar{\nu}_e$ flux flavour ratio right. The flux of other flavours are not shown but follows similar behaviour to the ν_μ flux.

these figures demonstrate the significant differences between the analyses which are clearly outside the dataset uncertainties. Recalling Figure 5.19 which shows contributions across the x_F - P_T phase space to the neutrino flux, it can be seen that pions produced in the smallest angular bins of the large angle datasets (i.e. the largest/least negative x_F) contribute most significantly to the neutrino flux. Within this region, the differences between the datasets are approximately 20-30%.

Each of the HARP high angle analyses is used to separately produce pion production weights which are used to weight the neutrino flux. Figure 9.9 shows the ν_μ flux and ν_e to $\bar{\nu}_e$ flux ratio obtained when the hadron production weighting is performed using each of these analyses separately. It can be seen that despite the significant differences between the two analyses the effect on the neutrino flux is comparatively small. This is due to a cancellation in the differences between the two analyses across the x_F - P_T plane when the flux calculation integrates over the production of pions to produce a weighted flux. Further, it can be seen that the greatest difference between the two flux predictions is in the ν_e to $\bar{\nu}_e$ ratio where a $\sim 1\%$ difference arises due to the differences in relative π^+/π^- production between these analyses.

The minimal differences between the fluxes produced when using each analysis provide confidence that the approach taken so far (of weighting to the HARP OH analysis alone) is sufficient. The approach to estimating the systematic uncertainty in this region is less clear. One approach may be to take the small difference between the fluxes caused by the differing analyses as the uncertainty. There is concern that this approach would underestimate the underlying hadron production uncertainty. When calculating the flux there is clearly a significant cancellation of differences between each analysis across the x_F - P_T plane; however, taking the difference in the fluxes as the uncertainty does not allow for the possibility that one analysis is closer to the ‘true’ multiplicity in one region of phase space and the other to be closer in another region.

As discussed above, it is typically the lowest angular bins within these datasets which contributes most significantly to the neutrino flux. To estimate the uncertainty, the discrepancy between the two datasets needs to be understood in the phase space region which contributes most significantly to the flux. To this end, Figure 9.10 shows a histogram of the fractional discrepancy in x_F - P_T phase space weighted by the fractional contribution of that region of phase space to the neutrino flux. From this, it is concluded that the fractional discrepancy between the two analyses is typically within 30% in the region which contributes to the flux. A constant 30% uncertainty is applied to the high angle pion production and propagated through to the neutrino flux via the methods discussed in Section 9.1.1. Similarly to Section 9.1.2 the uncertainty is treated as uncorrelated between $\pi^\pm$ production but correlated across the energy ranges. The effect of this uncertainty on the neutrino flux is shown in Figure 9.11.

Ultimately, there is insufficient information for a precise estimate of the uncertainty within this phase space region and more data will need to be taken by future experiments to resolve these issues. While the 30% uncertainty may prove to be too high, with the datasets currently available it is judged to be a reasonable conservative estimate.




Figure 9.10: Histogram showing the fractional discrepancy between the HARP OH and HARP CDP analysis in the x_F - P_T plane weighted by the contribution of that region of x_F - P_T phase space to the neutrino flux. The flux contribution on the y-axis is in arbitrary units. Shown for π^+ production top and π^- production bottom for each of the 4 beam momenta. The 30% uncertainty chosen as the uncertainty in this region is also shown.




Figure 9.11: Uncertainty on the neutrino flux due to HARP high angle dataset. Shown are the uncertainty on the absolute ν_μ flux (left) and $\nu_e/\bar{\nu}_e$ ratio (right). The total uncertainty is the addition in quadrature of the individual uncertainties.

9.1.4 Target Nucleus Interpolation Uncertainty

In Section 5.2 a method was developed to interpolate between production cross sections for particle production on different targets, this was then used in Chapter 7.1.2 to interpolate the production cross section to air. The parameterisation of the lower energy (<31 GeV) datasets involved comparisons between datasets using different targets with the uncertainties inflated to cover the spread in data points between the datasets, therefore any uncertainty on the target nucleus interpolation will have already been included in that inflation of those dataset systematic uncertainties. For pion production at higher energies the NA61/SHINE and NA49 datasets which measured the production cross section on a carbon target are used together with the target nucleus interpolation to estimate the multiplicity of pion production on air. To investigate the sensitivity of the interpolation to the neutrino flux calculation, the neutrino fluxes are tuned to the multiplicities of pion production on carbon directly without the interpolation to air. Figure 9.12 shows the neutrino fluxes where pion production is weighted directly to the multiplicity on carbon measured by NA61/SHINE and NA49 compared to the neutrino flux using the full interpolation from carbon to air. As can be seen,




Figure 9.12: Tuning of the neutrino fluxes with and without the target nucleus interpolation used to scale the NA61 and NA49 pion multiplicities from carbon to air. The ratio of the ν_μ neutrino flux tuned using pion multiplicity on a carbon target is shown as a ratio to the neutrino flux where those multiplicities have been interpolated to air.

the effect of this change on the flux is minimal. This is because although the target nucleus scaling has a large effect on the differential cross section most of its behaviour can be approximated by the $A^{2/3}$ scaling which will cancel with the scaling of the total production cross section when calculating the multiplicity. As such the uncertainty on the interpolation will be negligible compared to the other uncertainties discussed and will not be considered further.

9.1.5 Kaon Dataset Uncertainties

The E802 and Allaby datasets used to describe kaon production at ~ 17 GeV beam momentum provide a single bin-by-bin uncertainty which is treated as a systematic uncertainty. Although there are no direct comparisons to indicate discrepancies (as exist with the pion datasets), the kaon datasets already have larger uncertainties and the joint fit between the two datasets works well (see Chapter 6). As such the original uncertainties are used without additional inflation factors. Note that the normalisation uncertainty is treated as correlated with the pion production normalisation uncertainty, and included in the discussion above.




Figure 9.13: Uncertainty on the ν_μ (left) and ν_e (right) flux due to the uncertainty on kaon production within experimental datasets.

The approach to the propagation of uncertainties in the NA61 and NA49 kaon production datasets follows a similar approach to that of the pion datasets and will not be discussed in detail again here. When considering pion production dataset uncertainties, it was found that the uncorrelated bin-by-bin statistical uncertainties only had a negligible contribution when propagated to the neutrino flux compared to the correlated systematic uncertainties. It is, therefore, to be expected that the correlated systematic uncertainties on the kaon data will similarly have a much greater effect on the neutrino flux than the uncorrelated statistical uncertainties. In keeping with the approach taken with the pion datasets, the uncorrelated statistical uncertainties will not be considered for the kaon datasets. The uncertainty on neutral kaon production, which was estimated from the charged kaon production cross sections, will be derived from the uncertainty on the charged kaon production via the relationship established in Section 6.3. Since neutral kaons contribute to less than $\sim 5\%$ of the flux, additional uncertainty on the neutral kaon production will not significantly change the overall neutrino flux uncertainty estimate.

These uncertainties are shown for the ν_μ and ν_e flux in Figure 9.13. Note that the NA61 K^+ production uncertainty is the largest not due to NA61 having the largest individual errors but due to NA61 covering a larger region of the relevant phase space for pion and kaon production. By contrast, NA49 has a relatively small contribution to the uncertainty due to its much smaller phase space coverage.

Additional uncertainties are assigned for kaon production outside experimental phase space coverage in Section 9.2.

9.2 Interactions Occurring outside Dataset Phase Space

Chapter 8 explored estimates of pion and kaon production outside the x_F - P_T phase space coverage of the datasets by different hadron production generators. The pion and kaon multiplicity outside the dataset phase space coverage was weighted to the mean value of the generators that most accurately replicated the datasets. This section considers the uncertainty of these estimates and their propagation through to the neutrino flux. This section specifically considers the uncertainty on the estimate of pion and kaon production outside the x_F - P_T phase space coverage but within the interaction energy range of the datasets (i.e. 3-450 GeV incident particle momenta for pion production and 14-158 GeV for kaon production).

Recalling Figure 8.1 and the discussion in Section 8.1 the generators use different underlying models at different incident particle momenta. Different physics lists were selected for each of the 7 pion production regions in Table 8.1 such that the underlying models were independent. As the same models are used across several of these 7 regions (e.g. the binary cascade model is used at both regions 1 and 2 (incident momenta 3 and 5 GeV)), it would be incorrect to treat the generator hadron production estimate as independent between these 7 regions. To estimate the systematic uncertainty, three new broader pion and kaon production energy regions are defined for production outside experimental phase space. These regions are labelled A-C and very roughly correspond to the regions where the different hadron production models from Figure 8.1 are used. Region A covers interactions with incident momenta from 3 to ~ 6 GeV covering energy regions 1 and 2. B covers interactions from 6 to ~ 100 GeV covering energy regions 3-6 and region C covers interactions > 100 GeV covering region 7. Since kaon production is tuned over a smaller range of energies only regions B and C will be relevant. The uncertainty on




Figure 9.14: Contribution to the ν_μ flux from pion and kaon production interactions occurring outside the phase space of the experimental datasets. The fractional contribution for pion production (left), and kaon production interactions (right) is shown in the 3 energy regions described in the main text.

the hadron production outside experimental datasets modelled by the generators is taken as correlated within each of these energy regions but not between these energy regions. Figure 9.14 shows the contribution of each of the A-C regions to the weighted neutrino flux. Note that the fraction of the neutrino flux coming from pions produced outside the dataset phase space coverage is lower in Figure 9.14 than was previously shown in Figure 5.22 which showed the fraction prior to the tuning of hadronic interactions.

Initially, the variance in the multiplicity between each generator in Table 8.1 was explored as a possible measure of the uncertainty. Each of the individual generators was used to weight hadronic interactions within energy regions A-C and used to produce a weighted flux. The variance between the flux produced by different models was taken as the uncertainty. This method was found to work poorly as the variation between the generators outside the experimental phase space was found to be lower than the uncertainty from experimental measurements inside the phase space. This resulted in smaller uncertainties assigned to hadron production where no measurements were made than were measurements had been made.

As an alternative approach a constant 40% uncertainty is applied to each of the energy regions A-C. Without measurements of the multiplicity in this region, the only constraints are due to charge and energy conservation. A 40% uncertainty is chosen to reflect the maximum uncertainty that was applied to the original




Figure 9.15: Uncertainties on the ν_μ flux due to pion (left) and kaon (right) production outside the x_F - P_T dataset phase space. Broken down into energy regions which are described in the main text. In each case, the total is the quadrature sum of the individual parts.

uncertainty estimate on hadron production with the original TARGET2.1 flux prediction in the region without data available. In addition, the comparisons between generators and datasets in Chapter 8 showed that generators are typically well within a 40% agreement with the data points. This uncertainty is treated as correlated within each region A-C but not between these three regions. Additionally and consistent with the approach taken in Section 9.1, these uncertainties are assumed uncorrelated between the production of different meson species. Within this scheme, the uncertainties are shown in Figure 9.15. Since these uncertainties are treated as uncorrelated, the total uncertainty in Figure 9.15 is calculated by summing the parts in quadrature.

9.3 Energy Scaling Uncertainty

Until this point the uncertainties on hadron production have focused on pion/kaon production uncertainties in the x_F - P_T plane. These have been considered at each of 7 interaction energies (3 for kaon production) grouped according to projectile momentum. This section considers the uncertainty in the interpolation of the production cross section between these interaction energies.

The interpolation uncertainty will be dominated by the interpolation between the 31 GeV NA61 dataset and 158 GeV NA49 dataset and separately the Feynman scaling of the NA49 dataset to the NA56 450 GeV dataset. It will be assumed that the uncertainty in the interpolation between the lower (<31 GeV beam momenta) energy regions is not significant. This is due to the differences in beam energy between these lower energy regions being much smaller (recall Figure 4.5) and the experimental uncertainties being inflated in Section 5.5.2 to cover known differences between the datasets at slightly different beam momenta.

Figure 9.16 shows the pion production cross section as a function of interaction energy at fixed points in the x_F - P_T plane obtained from the fits from Chapter 5. This figure highlights both the gaps between the datasets at 31-158 GeV and 158-450 GeV where the interpolation uncertainties are considered and also demonstrates the Feynman scaling behaviour discussed in Section 5.4.2 where the dependence of the production cross section on interaction energy decreases with increasing energy.

9.3.1 NA61-NA49 Interpolation

To estimate the uncertainty of the linear interpolation between the NA61 and NA49 datasets, the limiting case is taken where the 158 GeV NA49 dataset weights are used to weight all interactions in the full 31-158 GeV energy range. A second limiting case is taken where the 31 GeV NA61 dataset weights are used to weight the full energy range. The neutrino flux in both cases is calculated with the difference between these limiting cases and the linear interpolation is taken as the uncertainty. This uncertainty is shown in Figure 9.17.

9.3.2 NA49-NA56 Energy Scaling

Section 5.4.2 demonstrated approximate Feynman scaling behaviour between the NA49 pion production dataset at 158 GeV and the 450 GeV NA56 dataset. As the NA56 dataset has smaller phase space coverage than the NA49 dataset, the NA49 dataset is used to weight interactions up to 450 GeV based on this scaling approximation. A 20% uncertainty is assigned due to this scaling based on the




Figure 9.16: Inclusive production cross section as a function of interaction beam momentum at fixed points in the x_F - P_T plane. Shown for π^+ (left) and π^- production (right) at $x_F = 0.1, 0.2, 0.4$ top, middle, bottom. Each point represents the invariant differential cross section of the combined dataset fit, with the uncertainty on the points being a quadrature sum of the uncertainties discussed in Section 9.1. These points are only shown where they are inside the dataset phase space. A straight line is drawn between each of the points.




Figure 9.17: Uncertainties on the ν_μ flux due to interpolation between the 31 GeV and 158 GeV datasets and the Feynman Scaling process from 158 GeV to 450 GeV for pion production

differences between the NA56 and NA49 datasets in Figure 5.10. This uncertainty is linearly interpolated between 0% at 158 GeV where there is no scaling and the full 20% at 450 GeV. This uncertainty is applied separately to π^+ and π^- interactions and is also shown in Figure 9.17.

9.4 Additional Hadron Production Uncertainties

When estimating the effect of hadron production uncertainties on the neutrino flux the 2006 Bartol flux uncertainty estimate only considered the uncertainty on pion and kaon production (see Section 3.4.1). The reduction in the uncertainty of pion and kaon production obtained when using the tuning procedure developed within this thesis necessitates the consideration of uncertainties which were previously subdominant. This section will estimate the uncertainty of meson re-interaction and baryon-to-baryon interactions for the first time. Additionally, the effect of neutral lambda baryon production will be considered. These interactions which have a smaller effect on the neutrino flux have not been tuned to data as pion and kaon production were. This section will now consider the uncertainties in these interactions where the original TARGET2.4 generator has been used.




Figure 9.18: Example of a high energy cosmic ray undergoing multiple re-interactions to produce a lower energy neutrino

9.4.1 Baryon-to-Baryon Interactions

As shown in Figure 4.1 in Chapter 4, a significant component of the inclusive hadronic interactions within atmospheric showers are of the form $p/nN \rightarrow p/n + X$, where both the projectile and the daughter particle are baryons. Since conservation of baryon number directly governs the multiplicity of the daughter baryons it was postulated that their effect on the flux would be small and the tuning instead focused on pion and kaon production. This section considers these interactions in more detail.

When a high-energy cosmic ray first interacts within the atmosphere the mesons produced will have higher energies which will subsequently decay to higher energy neutrinos. However, as the cosmic ray undergoes interactions and loses energy, it will create lower energy pions and kaons which decay to lower energy neutrinos. Figure 9.18 shows a cartoon of such a chain of interactions where the initial primary cosmic ray re-interacts several times before creating a lower energy pion which decays to a lower energy neutrino. It is worth noting that daughter mesons can similarly re-interact to produce lower energy neutrinos; however, as was discussed in Chapter 4, their short lifetime mean they are more likely to decay than re-interact.

Figure 9.19a shows the mean number of hadronic interactions in the atmospheric shower as a function of neutrino energy. Figure 9.19b shows the average momentum of the incident baryon in the interaction which creates the meson that decays into the neutrino as a fraction of the initial cosmic ray momentum (i.e. the ratio of the momentum of the proton after the first two interactions to the initial cosmic ray momentum in Figure 9.18). These show as expected that the primary cosmic ray will undergo several re-interactions with greater energy loss to produce lower energy




Figure 9.19: The mean number of hadronic interactions as a function of neutrino energy (left) and the energy of a primary cosmic ray prior to the neutrino producing interaction as a fraction of the initial cosmic ray energy (right).

neutrinos. Note that a cosmic ray must undergo at least one interaction to produce a neutrino as it must create a pion or kaon which decays to a neutrino. Notice also that the average energy loss undergone by the baryon within ν_e production is lower than for ν_μ . This is a consequence of the ν_e being mainly produced in muon decay and therefore produced at a lower average energy than ν_μ which can be produced directly during pion decay.

These interactions can be further broken down and are grouped into three classes:

- **Leading baryon continuation:** an interaction of the form $pN \rightarrow p + X$ or $nN \rightarrow n + X$, where the incident baryon exits the interaction having lost energy within the interaction.
- **Leading baryon charge exchange:** an interaction of the form $pN \rightarrow n + X$ or $nN \rightarrow p + X$, where the incident baryon undergoes a charge exchange and the initial proton exits as a neutron (or vice versa).
- **Target fragmentation:** an interaction which causes the breakup of the target atmospheric oxygen or nitrogen nucleus and results in the ejection of a lower energy nucleon.

Baryon number conservation ensures that within the first two classes of interaction the number of interacting baryons within the atmospheric shower before and




Figure 9.20: Fractional contribution of each of the baryon to baryon hadronic interactions to the ν_μ neutrino flux.

after the interaction are the same. On the other hand, the fragmentation interactions essentially introduce additional baryons which can later interact to form neutrinos.¹ Figure 9.20 shows the contribution of each of these three types of interaction to the neutrino flux. In addition, baryon anti-baryon pair production is shown; however, (as was also shown in Chapter 4) this is estimated to only have a very small contribution to the neutrino flux in the energy range considered here. Note that a single neutrino can be a product of multiple baryon to baryon interactions (as in Figure 9.18) and therefore it is possible for the sum of these fractions to exceed one.

To assign systematic uncertainties to these interactions the multiplicity for proton production within the TARGET2.4 generator is compared to three experimental datasets for $pN \rightarrow p + X$ interactions taken at beam momenta 12 GeV by HARP [102], 31 GeV by NA61/SHINE [56] and 158 GeV by NA49 [103]. Experimental measurement of daughter neutrons is not possible in the tracking detectors typically used by these experiments (see Chapter 4) and although the NA49 experiment did obtain an estimate for neutron production using a hadronic calorimeter this result had significantly greater uncertainties than the proton measurements and will not be considered here.

¹While baryon number conservation must also apply, these interactions are able to release nucleons bound in nuclei and thus the number of free baryons increases after the interaction




Figure 9.21: Comparison of the multiplicity of $pN \rightarrow p + X$ interactions with TARGET2.4 and the HARP experimental dataset at a beam momenta of 12 GeV. Shown the HARP outgoing proton angle bins 75 mrad (top left), 125 mrad (top right), 175 mrad (bottom left) and 225 mrad (bottom right). The TARGET2.4 cross section is broken down into components from target fragmentation and leading baryon continuation. The contribution from pair production is also shown to be consistent with subsequent plots showing multiplicities at higher energies. At the 12 GeV energy of this figure, this component is negligible.

Figures 9.21 - 9.23 show comparisons between the TARGET2.4 generator and the data points for the multiplicity of $pN \rightarrow p + X$ interactions. The Monte Carlo truth information is used to break down the component of these interactions coming from target fragmentation, the leading proton continuation and pair-production.

A rough minimisation is attempted between the data and Monte Carlo which will be used to estimate the systematic effect on the neutrino flux. The fragmentation and continuation components are allowed to vary by a constant factor across the whole range of interaction energies and daughter proton phase space covered by each experiment i.e.:

$$\begin{aligned}
 \langle pN \rightarrow p + X \rangle_{\text{total}} = & A \cdot \langle pN \rightarrow p + X \rangle_{\text{continuation}} + B \cdot \langle pN \rightarrow p + X \rangle_{\text{fragmentation}} \\
 & + \langle pN \rightarrow p + X \rangle_{\text{pair production}},
 \end{aligned}
 \tag{9.2}$$




Figure 9.22: Comparison of the multiplicity of $pN \rightarrow p + X$ interactions with TARGET2.4 and the NA61 experimental dataset at a beam momenta of 31 GeV. Shown for 4 of the NA61 outgoing proton angle bins 15 mrad (top left), 30 mrad (top right), 120 mrad (bottom left) and 270 mrad (bottom right). The TARGET2.4 cross section is broken down into components as in Figure 9.21

where $\langle pN \rightarrow p + X \rangle_{\text{total}}$ is the total multiplicity for proton production broken down by the classes of interaction.

The A and B parameters are allowed to vary such that the function:

$$\chi^2 = \sum_{\text{data}} \left(\frac{\langle pN \rightarrow p + X \rangle_{\text{MC}} - \langle pN \rightarrow p + X \rangle_{\text{data}}}{\sigma_{\text{data}}} \right)^2, \quad (9.3)$$

is minimised. While this is an approximation to a full weighting of these interactions which only produces a single average weight across all proton kinematic phase space (as opposed to the full weighting of pions and kaons where the weight was calculated as a function of kinematic variables) it will be sufficient for an estimate of systematic effects.




Figure 9.23: Comparison of the multiplicity of $pN \rightarrow p + X$ interactions with TARGET2.4 and the NA49 experimental dataset at a beam momenta of 158 GeV. Shown for 4 of the NA49 outgoing proton transverse momentum, p_T bins 0.2 GeV (top left), 0.4 GeV (top right), 0.6 GeV (bottom left) and 1.1 GeV (bottom right). The TARGET2.4 cross section is broken down into components as in Figure 9.21

The results of this minimisation are as follows:

$$A = 0.78 \quad (9.4)$$

$$B = 0.74$$

Tuning the Leading Baryon

The tuning of the A parameter is performed by weighting individual neutrinos in the flux in a similar fashion to the pion and kaon multiplicity tuning described in Chapter 7. Each leading baryon interaction ($pN \rightarrow p + X$) is weighted by a factor of A . To preserve baryon number, the $pN \rightarrow n + X$ charge exchange interaction is also weighted in such a way that the total multiplicity of leading baryons is




Figure 9.24: Effect of the tuning of the leading baryons on the neutrino flux. The effect on both the ν_μ flux is shown (left) and the ν_e to $\bar{\nu}_e$ flux ratio (right), in both cases the value of the flux and flux ratio are shown as a ratio to their values before the weighting has been applied.

unaffected by the tuning. In other words such that:

$$\langle pN \rightarrow p + X \rangle_{\text{continuation}} + \langle pN \rightarrow n + X \rangle_{\text{charge ex.}} \quad (9.5)$$

is conserved under the tuning. Similar to the pion production tuning, isospin symmetry is used to weight equivalent interactions with incident neutrons.

Figure 9.24 shows the result of tuning the neutrino flux to the value of A given in 9.4. This large $\sim 22\%$ tuning is seen to have a very small ($< 1\%$) effect on the absolute ν_μ neutrino flux prediction. Since the tuning will alter the ratio of protons to neutrons after a hadronic interaction and since protons will produce π^+ more abundantly in further hadronic interactions while neutrons produce more π^- , this tuning will directly impact the ratio of π^+ to π^- production in atmospheric showers. Since the ν_e to $\bar{\nu}_e$ ratio is highly dependant on that ratio this tuning will affect that ratio to a greater degree as seen in Figure 9.24 where the tuning has a larger but still comparatively small 2% effect on the electron neutrino to antineutrino flux ratio. The comparatively small tuning effect confirms the previous assumption that the baryon number conservation will limit the sensitivity of the neutrino flux to these interactions. The $\sim 2\%$ effect on the neutrino flux ratio




Figure 9.25: The ν_μ (left) and ν_e (right) atmospheric neutrino flux after altering the multiplicity of fragmentation baryons. Two methods are used to estimate the tuning of these interactions, the first alters the TARGET2.4 generator and re-runs the Monte Carlo (labelled Monte Carlo Re-Run), and the second weights neutrinos which come from fragmentation protons after the original Monte Carlo run. In both cases, the resultant flux is shown as a ratio to the flux without the weighting of fragmentation baryon production.

will be taken as a systematic uncertainty on the leading baryon component of the baryon to baryon interactions although it will be subdominant to the uncertainties already considered in the pion and kaon production model.

Tuning the Fragmentation Baryons

The tuning of the B parameter is more involved than the tuning of the A parameter. If the number of fragmentation nucleons decreases (as the data suggests it should be) then the kinetic energy that the simulation assigns to these nucleons will instead be distributed to other particles (such as by increasing the production of pions or kaons or reducing the energy lost by the primary baryon in the interaction). Note that although the data only considered proton daughter particles, it is assumed that the probability of the fragmentation nucleon being a neutron or proton is equal since atmospheric nuclei are isoscalar i.e.:

$$\langle pN \rightarrow p + X \rangle_{\text{fragmentation}} = \langle pN \rightarrow n + X \rangle_{\text{fragmentation}} \quad (9.6)$$

Therefore both proton and neutron fragmentation production will be weighted by the same factor. Two methods were used to estimate the effect on the neutrino flux. The first altered TARGET2.4 to reduce the number of fragmentation nucleons by




Figure 9.26: The Λ^0 production component to the ν_μ neutrino flux broken down by separate contributions from each of the decay products. In each case, the fraction of the neutrino flux that was produced via an interaction chain involving either of the Λ^0 decay products is shown.

30% in line with the approximate fitted value for B , this allowed the additional energy to be distributed amongst the other daughter nuclei but required a full re-run of the neutrino Monte Carlo simulation. A new set of pion and kaon production weights with this altered version of TARGET2.4 are produced and used to obtain a weighted flux. The second method simply weights target fragmentation interactions down in the final neutrino flux via the weighting technique already described. Both methods are shown in Figure 9.25. Despite the second method not actively re-distributing this energy to other daughter particles, it can be seen that it gives the same result as the first method to within the fluctuation of Monte Carlo statistical errors. As both the altered and original generator are tuned to the same pion and kaon production model the differences between the two generators after tuning will be the energy distribution of the outgoing leading baryon and pion and kaon production below the energy range where tuning is applied. The deviation between the original TARGET2.4 and this tuned model will be taken as the uncertainty which as expected can be seen to be greatest at lower neutrino energies.

9.4.2 Neutral Lambda Contribution to the Neutrino Flux

The strange neutral lambda baryon decays into a proton or neutron and pion. As the lambda lifetime is short it will have always decayed before it could re-interact within the atmosphere and therefore for the purpose of calculating the neutrino flux it does not matter whether the decay daughter particles were produced in the initial hadronic interaction or as a subsequent decay of the Λ^0 . It is only necessary to differentiate them because, as was discussed in Chapter 4, the NA61 and NA49 pion production spectra are explicitly corrected to remove pions produced in Λ^0 decays (the feed-down decay).

Figure 9.26 shows the contribution to the ν_μ neutrino flux from Λ^0 decays to both pions and protons. The contribution to other flux flavours is similar to that for ν_μ and not shown. As can be seen around 2% of the neutrino flux originates in atmospheric showers interactions involving Λ^0 production with the majority of the component coming from the subsequent re-interaction of the protons and neutrons produced in the Λ^0 decay with only a much smaller component coming from the pions produced in Λ^0 decay.

Since baryon number must be conserved in Λ^0 production and decay, changes to the Λ^0 baryon production will not result in an overall increase or decrease in proton and neutron production. Since Section 9.4.1 found a 20% change in the significantly larger baryon re-interaction component resulted in a $< 1\%$ change in the neutrino flux it is not expected that the uncertainty in the rate of Λ^0 production and subsequent decay to protons and neutrons will have a significant impact on the neutrino flux. While the component of the neutrino flux arising from pions produced in Λ^0 will be affected by a change in the rate of Λ^0 production, this component is $< 0.5\%$ and therefore any systematic uncertainty on the Λ^0 production will be negligible in comparison to the pion and kaon production uncertainties and will not be considered.

9.4.3 Secondary Meson Interactions

In Chapter 4, Figure 4.1 showed that only a small ($< 5\%$) fraction of all hadronic interactions which contribute to neutrino production are due to mesons produced in




Figure 9.27: The fractional component of the ν_μ neutrino flux created in atmospheric showers involving meson re-interaction separated into components coming from pion and kaon interaction with the pion component further separated by the daughter particle created in the interaction.

hadronic interactions subsequently re-interacting. In this subsection, a systematic uncertainty on the generation of these interactions with TARGET2.4 will be estimated and propagated through to the neutrino flux. To begin it is necessary to separate these interactions by meson parent and daughter particle. Figure 9.27 shows the fraction of the ν_μ neutrino flux produced within an atmospheric shower involving re-interactions of different species of meson. It can be seen that the dominant meson re-interaction is of the form $\pi^\pm N \rightarrow \pi^\pm + X$ where the outgoing daughter is a pion of the same charge as the parent. Note that isospin symmetry links the inclusive cross sections of interactions of positive and negative pions on isoscalar nuclei as follows:

$$\sigma(\pi^+ N \rightarrow \pi^+ + X) = \sigma(\pi^- N \rightarrow \pi^- + X), \quad (9.7)$$

hence the interactions are grouped together.

Figure 9.28 shows the fractional contribution of interactions of the form $\pi^\pm N \rightarrow \pi^\pm + X$ at different incident pion momentum to the neutrino flux. It can be




Figure 9.28: Contribution of $\pi^+ N \rightarrow \pi^+ + X$ interactions to the ν_μ neutrino flux shown as a function of incoming pion momentum. The red line at incoming pion momentum 60 GeV indicates the beam momentum of the NA61/SHINE dataset used as a comparison to TARGET2.4. The contribution at each neutrino energy and incoming pion momentum is shown in arbitrary units.

seen that the relevant incident pion momentum is between 20 and 100 GeV which follows the expectation that pions with higher momentum are more likely to re-interact due to their longer decay length. The NA61/SHINE experiment measured the multiplicity of daughter particles produced in interactions of a 60 GeV π^+ beam with a carbon target [86]. Comparisons between the measured multiplicity for $\pi^+ N \rightarrow \pi^+ + X$ interactions and TARGET2.4 will be used to estimate the uncertainty of the TARGET2.4 interaction generation.

Figure 9.29 shows the multiplicity of the NA61/SHINE data compared to the multiplicity of TARGET2.4 in four bins of daughter pion production angle. It can be seen that TARGET2.4 over-produces pions in the low production angle region but under-produces pions in the higher angle production bins.

To estimate the uncertainty a similar approach to the uncertainty of baryon to baryon interactions (Section 9.4.1) is taken, with a single parameter used to adjust the normalisation of the multiplicity of $\pi^+ N \rightarrow \pi^+ + X$ interactions within




Figure 9.29: Comparison between the multiplicity of $\pi^+N \rightarrow \pi^+ + X$ interactions at incident momentum 60 GeV within TARGET2.4 and NA61/SHINE experimental data for interactions on a carbon target. Shown for 4 bins of daughter pion production angle θ , 5 mrad (top left), 30 mrad (top right), 80 mrad (bottom left) and 210 mrad (bottom right). In each bin, the double differential multiplicity ($\frac{d^2n}{dpd\theta}$) is shown as a function of pion momentum. The data is from [86].

TARGET2.4. A minimisation of the χ^2 comparing the TARGET2.4 multiplicity to the data (similar to equation 9.3) is used to determine the value of the normalisation parameter where the TARGET2.4 Monte Carlo generation best matches the dataset. The normalisation factor is an approximation to a full weighting which does not tune the shape of the TARGET2.4 but only alters the average multiplicity across the outgoing pion phase space. As such, shape differences in the multiplicity between TARGET2.4 and the data are not accounted for. This approach will be sufficient as an estimate of systematic uncertainty as these interactions only have a small contribution to the flux.

The minimisation found the normalisation factor to be 0.5, driven by the over-production of pions at small production angles. These pion interactions are then




Figure 9.30: Ratio of the ν_μ neutrino flux after the weighting of $\pi^\pm N \rightarrow \pi^\pm + X$ interactions to the flux without pion re-interaction weighting. The difference between the original and the tuned fluxes will be taken as the uncertainty of the neutrino flux due to these interactions.

assigned a weight of 0.5 and the neutrino flux is calculated and shown in Figure 9.30 and the difference between the weighted flux and the original TARGET2.4 estimate is taken as the uncertainty, which is approximately 2% at the highest neutrino energies. As there is no expectation of significant correlation between the remaining meson re-interaction components in Figure 9.27 and as they are all seen to individually contribute to less than 1% of the neutrino flux, the effect of their uncertainty will be subdominant to the $\pi^\pm N \rightarrow \pi^\pm + X$ interactions and will not be considered.

9.4.4 Pion and Kaon Production Outside the Tuning Energy Range

Within this thesis pion and kaon production have been tuned for interactions with incident particle momenta in the range 3-450 GeV for pion production and 17-158 GeV for kaon production. Figure 5.22 in Chapter 5 and Figure 6.3 in Chapter 6 show the fraction of the total neutrino flux originating from interactions outside this range. The most significant contribution is from pion production at beam momenta < 3 GeV to the lowest energy neutrino flux. Resonant production of particles will be significant within this energy range. Resonances cause the production cross

section to vary significantly as a function of energy. As such, the approach taken so far of smoothly weighting pion and kaon production within interaction energy bins, will not work at these lower energies. TARGET2.4, which includes resonant production models, will be used without any tuning outside the energy range. The original uncertainties assigned to TARGET2.1 (discussed in Section 3.4) are taken as the uncertainties on pion and kaon production within this region and propagated through to the neutrino flux.

9.5 Summary

This chapter has considered the uncertainty of the neutrino flux due to pion and kaon production in atmospheric showers. Uncertainties associated with the pion and kaon production dataset models from Chapters 5 and 6 have been discussed in addition to the uncertainty for the estimates of pion and kaon production outside the dataset phase space. These uncertainties were then propagated through to the neutrino flux prediction. Additionally, uncertainties on baryon to baryon interactions and meson re-interaction as modelled by TARGET2.4 have been estimated for the first time. In Chapter 11 the uncertainties on the neutrino flux due to the hadronic interactions presented within this chapter will be combined with the uncertainty on the primary cosmic ray flux to obtain an estimate for the overall neutrino flux uncertainty.

A summary of the uncertainty on the atmospheric ν_μ flux prediction due to the components discussed in this chapter is shown in Table 9.1.

Section	Uncertainty Source	0.1 GeV	1 GeV	10 GeV
9.1.1	NA61 NA49 Dataset Bin-by-bin Uncertainty	0.2%	0.7%	1.3%
9.1.1	NA61 NA49 Dataset Normalisation Uncertainty	0.0%	0.4%	1.3%
9.1.2	Low Energy (<30 GeV) Dataset Bin-by-bin Uncertainty	2.1%	6.5%	0%
9.1.2	Low Energy (<30 GeV) Dataset Normalisation Uncertainty	0.8%	3.0%	0.2%
9.1.3	HARP High Angle Data Uncertainty	6.5%	0.0%	0.0%
9.1.5	Kaon Dataset Bin-by-bin Uncertainty	0.0%	0.2%	1.0%
9.1.5	Kaon Dataset Normalisation Uncertainty	0.0%	0.3%	0.6%
9.2	Outside Pion Dataset Phase space	4.2%	2.1%	0.7%
9.2	Outside Kaon Dataset Phase space	0.2%	0.4%	0.9%
9.3.1	NA61-NA49 Interpolation	0.5%	1.5%	0.6%
9.3.2	NA49-NA56 Feynman Scaling	0.0%	0.0%	2.1%
9.4.1	Target Nucleus Fragmentation	6.8%	0.5%	0.0%
9.4.1	Leading Baryon Flavour Change	0.0%	0.2%	0.5%
9.4.3	Pion Re-interaction	0.2%	0.7%	2.0%

Table 9.1: Summary of the hadron production uncertainties discussed in Chapter 9 on the atmospheric ν_μ flux prediction. In each case, the uncertainty source is named alongside the section of this chapter where it is discussed in detail. The percentage uncertainty due to each source on the ν_μ flux prediction is given at three neutrino energies, 0.1 GeV, 1 GeV and 10 GeV. Uncertainties greater than 1% have been highlighted for ease of identifying the dominant components at each neutrino energy.

10

Primary Cosmic Rays

This chapter considers the flux of cosmic rays and their effect on the atmospheric neutrino flux. The composition and features of the cosmic ray flux spectrum are discussed alongside the different methods for their detection. A parameterisation and uncertainty obtained by Evans et al. in a global fit to proton and helium cosmic ray data is applied to the neutrino flux calculation and used to estimate the uncertainties on the neutrino flux due to the spectra of primary cosmic rays.

10.1 The Cosmic Ray Spectrum

The atmospheric showers which bring about neutrino production are initiated by cosmic rays incident on the top of the Earth's atmosphere. As such, understanding the energy spectra and composition of these cosmic rays is highly important to the atmospheric neutrino flux estimate. The origin and acceleration mechanism of these cosmic rays is still a topic of active research with leading theories suggesting that they originate in galactic supernovas [104]. Since the energy spectra of cosmic rays span many orders of magnitude (with the highest energy cosmic ray recorded at 10^{20} eV [105]) they provide a unique opportunity to study the standard model at energies inaccessible to particle accelerators. In addition to this, the chemical composition

of cosmic rays and their spectra provide a method of investigating the composition of the universe and astrophysical phenomena. While understanding the origin of these cosmic rays is not as important for estimating the neutrino flux, the neutrino flux estimate is able to take advantage of recent high accuracy cosmic ray flux measurements. This chapter will focus on the cosmic ray flux between 1 GeV-10 TeV previously shown to be relevant to the neutrino flux between 0.1-10 GeV [48].

Chapter 2 has already discussed the modulation of the primary flux by the solar wind and the geomagnetic cut-off where the effect of the Earth's magnetic field on the propagation of the flux to different regions of the atmosphere is taken into account. This chapter will focus on the measurement of the isotropic cosmic ray flux. Geomagnetic cut-off corrections are applied to this isotropic flux within the Monte Carlo at the point of the primary cosmic ray generation using the backtracing technique as discussed in Chapter 3. Solar modulation corrections are performed during oscillation analysis using neutron monitor data to estimate the cosmic ray modulation during the period in which a neutrino detector is taking data [106] and will not be discussed here.

10.1.1 Composition and Features

The cosmic ray flux consists of protons (approximately 90%), helium nuclei ($\sim 9\%$), and heavier nuclei ($\sim 1\%$) [36]. There is also an electron, positron and photon component but these will not contribute toward neutrino production and will not be discussed further.

Figure 10.1 taken from Ref. [107] shows the fluxes of the most abundant cosmic rays by species of nuclei. It is seen that above about 10 GeV ¹ the fluxes can approximately be described by an inverse power law:

$$\Phi \sim E^{-\gamma} \quad (10.1)$$

where γ known as the spectral index is found to be ~ 2.7 from 10 GeV to 1 PeV [36].

¹Below ~ 10 GeV solar modulation causes deviation from the simple power law




Figure 10.1: Fluxes of primary cosmic ray nuclei. For species other than free protons (H) the flux has been multiplied by the factor labelled to aid with clarity. Figure taken from [107]

Under the superposition approximation discussed in Section 3.4, the flux of each species of nuclei is less important and instead the total flux of nucleons is needed to calculate the neutrino flux. Conventionally the nuclei are separated into 5 groups [107]:

- Free protons (H nuclei)
- Helium nuclei
- medium nuclei ($6 \leq Z \leq 9$) of which the C and O fluxes are the most significant

- heavy nuclei ($10 \leq Z \leq 20$) of which the Ne, Mg and Si fluxes are the most significant
- very heavy nuclei ($21 \leq Z \leq 30$) of which the Fe fluxes are the most significant

The final 3 groups are sometimes called the CNO, MgSi and Fe groups [108]. The contribution from $3 \leq Z \leq 5$ is negligible.

The average number of protons and neutrons in each group of bound nuclei is then used to calculate a flux of bound protons and bound neutrons. Within the Monte Carlo simulation, these differ only from their free counterparts in their interaction with the Earth's magnetic field. Bound protons and neutrons are treated as having a charge-to-mass ratio of half that of a free proton to account for being bound in light nuclei.

The 2004 Bartol neutrino flux estimate used the ICRC01 parameterisation to model the cosmic rays with parameters obtained in a global fit to the data, the fitted parameters and the parameterisation are described in more detail in [52]. The parameters obtained from this global fit were presented at the ICRC conference in 2001, after which the parameterisation is named. Within this chapter the same parameterisation will be used; however, the parameters will be updated with a global fit to more recent data (see Section 10.2). The ICRC01 parameterisation is as follows [52]:

$$\Phi(E_K) = a[E_K + b \exp\{(c\sqrt{E_K})\}]^{-d}, \quad (10.2)$$

with $\Phi(E_K)$ as the primary cosmic ray flux, E_K the kinetic energy per nucleon and a - d parameters to be obtained in the fit to datasets. The dominant feature of this parameterisation is the spectral index d , describing the power law discussed earlier. To avoid confusion, care will be taken to emphasise the difference between the ICRC01 parameters obtained in the 2001 fits and the newly obtained global fit parameters discussed in 10.2, both are parameters of the same function, the ICRC01 parameterisation shown above.




Figure 10.2: Fractional contribution of each cosmic ray nuclear group to the total nucleon flux (left) and the atmospheric neutrino flux(right) as estimated with the ICRC01 parameterisation with the 2001 ICRC01 parameters. The primary energy in the left-hand graph is energy per nucleon. The neutrino contribution is calculated after the hadron interaction weighting described in earlier chapters. The groups are defined within the main text.

The fraction of each group of nuclei to the total nucleon flux (as estimated by the 2001 ICRC01 parameters [47, 52] of cosmic rays) is shown in Figure 10.2. Also shown is the estimated contribution of each nuclei group to the neutrino flux using the Bartol Monte Carlo with the TARGET2.4 generator. As can be seen, the majority of the nucleons originate from the free proton and helium nuclei with a much smaller component due to heavier nuclei. Note that although the heavier nuclei only contribute $\sim 1\%$ to the total flux of cosmic rays as measured by the number of nuclei they have a greater contribution to the total flux of nucleons as a single heavy nucleus will contain multiple nucleons.

The remainder of this chapter will focus on the calculation of the neutrino flux with updated measurements of the primary proton and helium flux. The heavier nuclei will continue to use the original ICRC01 parameters and uncertainties (as in [47]).

10.1.2 Detection of Cosmic Rays

With the cosmic ray energy spectra spanning many orders of magnitude and the flux at the highest energies being many orders of magnitude lower than at lower energies, it is unsurprising that the experimental methods used will vary depending on the energy region being measured.

For measurements of the lower energy flux where the flux of particles is greatest a magnetic spectrometer-based detector provides the most accurate cosmic ray momentum measurements [36]. Such detectors can be installed on satellites (such as PAMELA [109]), the international space station (such as AMS-02 [110]) or balloon flights (such as BESS [111]) to measure cosmic ray fluxes at the top of the atmosphere before scattering on atmospheric gasses can occur. Space station and satellite measurements have the advantage of being both higher above the atmosphere and being able to take data for longer periods of time than balloon experiments. Such experiments are used to measure the flux from 1 GeV-1 TeV per nucleon.

High energy cosmic rays have too low a flux for the small acceptance of a magnetic spectrometer to record sufficient events [36]. To measure fluxes from 1-100 TeV per nucleon, calorimeter-based detectors are used without a magnetic field. The absence of the magnetic spectrometer reduces the accuracy of the momentum measurement but increases the detector acceptance allowing for more events to be recorded. The CREAM [112] experiment is a calorimeter-based detector that took data on balloon flights.

At the highest energies, the fluxes are so small that the detector acceptance needed is too large and the detector becomes physically too large to install on a satellite or balloon. Instead, ground-based detectors with large acceptances measure the products of atmospheric showers from which the original primary kinematics can be reconstructed. These cosmic ray energies are higher than what is necessary for the neutrino flux modelling in this thesis and will not be considered further.

10.2 Global Fits to Cosmic Ray Data

Evans et al. presented a combined fit of several separate primary cosmic ray measurements to the ICRC01 parameterisation [113] in 2017. A statistical technique was used to estimate the uncertainty of the combined fit in the absence of information about the correlation of uncertainties and to account for inconsistencies between different measurements. This was presented separately for both the free proton




Figure 10.3: Fractional deviation of the data points to the global fit shown for primary proton (left) and helium fluxes (right). This figure is taken from [113].

and helium cosmic ray fluxes. Figure 10.3 (taken from [113]) shows the deviation of individual data points from the global central fit value. For cosmic rays with energies less than 10 GeV/n only data taken during a period with low solar activity is used such that the flux parameterisation describes a minimum in solar activity.

Within this section, the global fit by Evans et al. will be used to tune the neutrino flux prediction and the uncertainty given on the fit parameters propagated through to the neutrino flux. As was described in Section 3.6, the neutrino fluxes are obtained by convolving the yield of neutrinos at each cosmic ray energy obtained by the Monte Carlo with the primary cosmic ray flux. The neutrino fluxes can therefore be calculated using a different cosmic ray model without re-running the Monte Carlo simulation. The parameters obtained with the Evans et al. global fit are used for proton and helium fluxes while the original ICRC01 parameters are used for heavier nuclei. Figure 10.4 shows the ratio of the primary fluxes with the Evans et al. parameters to the original proton and helium fluxes with the original ICRC01 parameters. The results for the neutrino flux using these new measurements are shown in Figure 10.5.




Figure 10.4: Ratio of the global fits of cosmic ray data by Evans et al. to the ICRC01 primary flux previously used to calculate the neutrino flux.




Figure 10.5: Ratio of the neutrino flux calculated using the global fit by Evans et al. to the flux calculated using the ICRC01 primary flux. Shown for ν_μ (left) and ν_e neutrino flux (right)

10.2.1 Uncertainty on the Neutrino Flux

To propagate the uncertainty to the neutrino flux, the covariance matrix for the fit parameters obtained in the global fit given by Evans et al. is used [113]. This covariance matrix is used to draw 1000 toy primary flux models for the primary fluxes via the Cholesky decomposition method [101]. The variance between these models is used to obtain the uncertainty on the primary fluxes and shown in Figure 10.6, these uncertainties recreate the 1σ band uncertainties given by Evans et al.




Figure 10.6: Uncertainty of primary proton (left) and helium (right) fluxes as given by the global fit from Evans et al. Plotted also are the uncertainties on the ICRC01 fluxes which were previously used within in the Bartol neutrino flux estimate. The uncertainties are shown as a function of cosmic ray kinetic energy per nucleon (the total nuclear kinetic energy divided by the number of nucleons.)

in Figure 10.3. Notice the jump in uncertainties at 200 GeV, which account for an increase in uncertainty on the spectral index at higher energies in the region primarily measured by the less accurate calorimeter-based experiments without magnetic spectrometers discussed in Section 10.1.2. Both the Evans global fit uncertainties and the original ICRC01 uncertainties increase the uncertainty on the spectral index by a factor between 2 and 3 above 200 GeV cosmic ray kinetic energy per nucleon.

Each of these 1000 toy models is used to calculate a neutrino flux and the standard deviation of these fluxes is used to estimate the uncertainty. The original systematics of the heavy nuclei ($Z > 2$) cosmic ray fluxes [47] are included also. The uncertainty on the neutrino flux due to primary cosmic ray fluxes following this method is shown in Figure 10.7. Although the heavy nuclei only contribute $< 10\%$ to the total neutrino flux, the new proton and helium uncertainties within the Evans et al. global fit are low enough that the original ICRC01 uncertainties on the heavy nuclei dominate. The uncertainty on heavy nuclei could be reduced further by incorporating more recent datasets to model their flux however this is beyond the scope of this work and the primary flux uncertainty of between 1-3% is already a significant improvement on the original model and is typically subdominant to the hadron production uncertainties seen in Chapter 9.




Figure 10.7: Uncertainty on the ν_μ (left) and ν_e (right) neutrino flux due to primary cosmic rays using the flux model by Evans et al. The original ICRC01 uncertainties are propagated for the heavy ($Z < 2$) nuclei. The total uncertainty is obtained by summing the proton, He and heavy nuclei fluxes in quadrature.

As a rough calculation to check that these uncertainties seem reasonable, consider the neutrino flux at 1 GeV where there is an uncertainty of around 0.6% coming from the proton flux uncertainty. Figure 10.2 shows that the proton cosmic rays contribute to around 60% of the total neutrino flux. The neutrino flux at 1 GeV will primarily originate from hadronic interactions of primary protons between 10-100 GeV where the uncertainty on the proton flux is between 0.6% and 1% this is consistent with the proton component of the flux uncertainty being $\sim 0.6\%$.

Alongside estimating the uncertainty of the cosmic ray flux, the paper by Evans et al. provided an estimate of the muon neutrino flux uncertainty when in the absence of a full atmospheric Monte Carlo simulation run, which has now been performed in this thesis. The estimate employed a rescaling technique whereby uncertainty on the ICRC01 parameters, which was propagated to the neutrino flux in the 2006 Bartol flux uncertainty paper [47], is rescaled to its new lower value in the Evans global fit. The rescaling technique obtained a significantly higher neutrino flux uncertainty than was obtained via the method above (shown in Figure 10.7) by directly propagating the uncertainties within the neutrino Monte Carlo. The uncertainties obtained by Evans et al. using the rescaling method are shown in Figure 10.8. After careful consideration of both methods, the discrepancy is believed to be caused by an error in the original cosmic ray flux uncertainties

within the 2006 Bartol flux uncertainty. The 2006 paper gives two values for the spectral index uncertainty with an increase in uncertainty at 200 GeV similar to the Evans' uncertainty in Figure 10.6. It appears that only the higher value for the uncertainty has been used across the whole cosmic ray energy range when the uncertainty is propagated to the neutrino flux. Attempts to reproduce the neutrino flux uncertainty from 2006 are only successful if the higher uncertainty is used across the whole energy range. As a consequence, the rescaling technique scales down an uncertainty which is higher than stated and the final neutrino flux uncertainties are therefore too high.

As a further confirmation that this is the cause of the difference between the uncertainty estimates consider Figure 10.8. It is noted that there is not a sharp increase in the neutrino flux uncertainty, which would be expected at the neutrino energies where the increased uncertainty in the spectral index (observed through the discontinuity in Figure 10.6) becomes increasingly important. By contrast, this behaviour is visible in Figure 10.7 with a sharp increase in uncertainties above 2 GeV.

Since the propagation of the uncertainties directly within the Monte Carlo in Figure 10.7 agrees with what would be expected with the primary flux uncertainties in Figure 10.6 via the rough calculation, and since it appears that the discrepancy is caused by an error in the propagation of the ICRC01 uncertainties to the neutrino fluxes, the uncertainties shown in Figure 10.7 will be assumed to be correct. In Chapter 11 the uncertainties obtained directly within the Monte Carlo (Figure 10.7) will be used as the uncertainty on the neutrino flux due to the uncertainty on the primary cosmic ray spectrum.




Figure 10.8: Uncertainty on the neutrino flux due to primary cosmic rays estimated using the rescaling technique by Evans et al. The neutrino energy is shown on the x-axis. This estimated uncertainty on the neutrino flux is found to disagree with uncertainties obtained by propagating the flux model by Evans et al. directly into the Monte Carlo which is discussed further in the main text. The lines labelled “this work” are the uncertainty predictions on the neutrino flux using the Evans et al. global fit, shown both broken down by each ICRC01 parameter and as a total (labelled “primary”). The lines labelled “Barr et al.” refer to the neutrino flux uncertainties due to the uncertainty of the ICRC01 parameters. Figure is taken from [113]

11

Results

Following the tuning of the pion and kaon production model within atmospheric showers and the inclusion of the new cosmic ray flux model, the updated flux and uncertainties are presented within this chapter. Section 11.1 presents the neutrino fluxes integrated over all incoming neutrino directions. The uncertainties due to hadron production described in Chapter 9 and the cosmic ray flux model in Chapter 10 are combined to estimate the total neutrino flux uncertainty. Section 11.2 presents the ratio between neutrino flavours together with an uncertainty estimate. Section 11.3 presents the neutrino fluxes as a function of neutrino cosine zenith and azimuthal angle. Until Section 11.4 all fluxes and ratios within this chapter will be shown for the Super-Kamiokande (Kamioka) detector site. Section 11.4 shows the neutrino fluxes modelled at the DUNE and IceCube detector sites and discusses differences between the flux prediction at these sites and at Kamioka.

11.1 Integrated Neutrino Fluxes

Figure 11.1 shows the un-oscillated neutrino flux prediction at Kamioka obtained using the hadron production and primary cosmic ray flux model tuning described within this thesis. In addition the ratio of the tuned flux to the flux without this tuning (the original Bartol 2004 flux) is shown. The tuned fluxes are within the 15% uncertainties originally assigned to the Bartol flux calculation in 2006 (see Section 3.4.1). In addition to the tuning, the Monte Carlo was run with




Figure 11.1: The atmospheric neutrino fluxes averaged over all incoming solid angles after the tuning of the pion and kaon production and application of the new primary cosmic ray model described within this thesis. Shown are the absolute neutrino fluxes (left) and the ratio of these fluxes, labelled tuned, to the original Bartol 2004 fluxes without the tuning, labelled original (right).

approximately ten times the number primary cosmic rays that were used in the 2004 three-dimensional flux estimate, taking advantage of the advances in computational power since the 2004/2006 flux releases.

Figure 11.2 shows the systematic uncertainty on this flux estimate combining the hadron production uncertainties described in Chapter 9 and primary cosmic ray uncertainties described in Chapter 10. For clarity, each of the individual experimental uncertainties are added in quadrature and shown as a single uncertainty (labelled pion/kaon production data). Unsurprisingly the pion production uncertainties typically dominate within the neutrino energy range considered here. The kaon production uncertainties become increasingly important at higher energies. Statistical effects can be seen in uncertainties below 10^{-2} which do not impact the total neutrino flux uncertainty. The total uncertainty (shown as a black line) is obtained by adding each of the independent uncertainties together in quadrature. The flux uncertainty in the 0.1-1 GeV neutrino energy region is between 8-10% falling to 4-5% between 1 and 10 GeV. This reduction is due to the lower uncertainties in pion production within the hadronic interaction energy region covered by the NA61/SHINE (31 GeV) and NA49 (158 GeV) experiments, partially due to the lower systematic uncertainties on the data points and partially due to the wider




Figure 11.2: The uncertainty on the neutrino flux shown for each neutrino flavour. Individual uncertainty components from chapters 9 and 10 are added together in quadrature to obtain the total neutrino flux uncertainty (shown in black). The pion dataset uncertainty is obtained as the quadrature sum of the uncertainty on each dataset (discussed in Section 9.1) (including the HARP high angle dataset, section (9.1.3)), the energy scaling uncertainty (Section 9.3) and the dataset normalisation uncertainty. The baryon to baryon uncertainty combines both the fragmentation and leading particle uncertainty and the cosmic ray uncertainty combined the proton, helium and heavy nuclei flux uncertainties. The original Bartol uncertainties applied outside the interaction energy range, where tuning of pion and kaon production was applied is shown separately for each. Shown for ν_μ (top right), $\bar{\nu}_\mu$ (top left), ν_e (bottom right) and $\bar{\nu}_e$ (bottom left).




Figure 11.3: The neutrino flux flavour ratios calculated using the model described within this thesis (labelled tuned flux) and the original Bartol 2004 model (labelled original flux). $\nu_e/\bar{\nu}_e$ (top left), $\nu_\mu/\bar{\nu}_\mu$ (top right) and $(\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e)$ (bottom).

pion production x_F - P_T phase space coverage of these experiments compared to the datasets at lower energies. Additionally the uncertainty on the fragmentation baryons (included within the baryon to baryon uncertainty in Figure 11.2) which was discussed in Section 9.4.1 is more important at lower neutrino energies.

11.2 Neutrino Flux Flavour Ratios

Figure 11.3 shows the tuned neutrino $\nu_e/\bar{\nu}_e$, $\nu_\mu/\bar{\nu}_\mu$ and $(\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e)$ flavour ratios together with the original flux estimate without the tuning. Note the reduction in the statistical fluctuations in the tuned flux due to increased Monte Carlo statistics. It can be seen that the deviation between the tuned and original flux ratios are much smaller than their absolute values (shown in Figure 11.1) due to correlation between the production of different flavours of neutrinos.




Figure 11.4: The neutrino flux flavour ratio uncertainties. Individual flux uncertainty components are shown and added in quadrature to obtain the total flux ratio uncertainty. The uncertainties are shown on $\nu_e/\bar{\nu}_e$ (top left), $\nu_\mu/\bar{\nu}_\mu$ (top right) and $(\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e)$ (bottom). The individual uncertainty components are described in Figure 11.2.

Figure 11.4 shows the uncertainty on each of these flavour ratios where it is seen that the uncertainty on the flux ratios are lower than the absolute flux prediction similarly due to the high degree of correlation within the hadron production and primary neutrino flux which give rise to different flavours of neutrino fluxes.

As explained in Section 2.3.3 the $\nu_e/\bar{\nu}_e$ ratio has strong dependence on the ratio between π^+ and π^- production. Care was taken within Chapter 9 not to overestimate the degree of correlation between $\pi^\pm$ production within experimental datasets. It is therefore unsurprising that the uncertainty on this ratio is dominated by pion production. The treatment of the dataset bin-by-bin uncertainties as uncorrelated between $\pi^\pm$ production increases the uncertainty on the $\pi^\pm$ production

ratio and hence the uncertainty on the $\nu_e/\bar{\nu}_e$ flux ratio. The $\pi^\pm$ treatment within this model represents a significant improvement from the 2006 Bartol flux uncertainty estimate where a constant 5% uncertainty was applied regardless of dataset coverage. Additionally, the baryon to baryon interaction affects this ratio to a smaller degree. The baryon to baryon interaction affects the rate of charge exchange of the leading baryon within an interaction and therefore the ratio of $\pi^\pm$ production by that baryon's subsequent re-interaction.

The high degree of correlation between ν_μ and $\bar{\nu}_\mu$ production which leads to the prediction of the ratio being approximately one at lower energies (discussed in Section 2.3.3) can equally be seen by the lower uncertainties on the ν_μ to $\bar{\nu}_\mu$ ratio. Lower energy neutrinos are produced in lower atmospheric showers where muons produced in pion decay have a greater probability of decaying before reaching the Earth's surface. As a consequence each pion will produce a ν_μ and $\bar{\nu}_\mu$ regardless of its charge (from direct decay and subsequent muon decay). Therefore at lower neutrino energies there is a greater degree of correlation between the ν_μ and $\bar{\nu}_\mu$ fluxes. At higher energies where the probability of the muon decaying before reaching the ground decreases this correlation is reduced. This is reflected within the uncertainty on the ν_μ and $\bar{\nu}_\mu$ ratio shown in Figure 11.4.

The uncertainty on the $(\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e)$ flux ratio is also shown and seen to depend dominantly on the pion production uncertainties.

11.3 Neutrino Flux Zenith and Azimuthal Directional Dependence

11.3.1 Neutrino Cosine Zenith Dependence

Until this point this chapter has considered the tuning of the neutrino fluxes averaged over all incoming directions. To consider the tuned neutrino flux prediction for neutrinos arriving at the detector from different directions the ν_μ neutrino flux is shown as a function of cosine zenith angle at the detector. The tuned neutrino flux is shown alongside the original (untuned) neutrino flux prediction in four neutrino




Figure 11.5: The ν_μ neutrino fluxes shown as a function of neutrino cosine zenith angle at the detector shown separately in four neutrino energy bins. The fluxes calculated using the tuning methods described in this thesis are shown (labelled tuned flux) alongside the original Bartol 2004 prediction (labelled original flux).

energy bins in Figure 11.5. The $\cos \theta_Z < 0$ ($\cos \theta_Z > 0$) correspond to the up-going (down-going) following the same convention as Section 2.3.1.

In the lowest neutrino energy bin it can be seen that the tuning causes an enhancement of the flux of horizontal neutrinos ($\cos \theta_Z \sim 0$). As discussed in Chapter 4, as a consequence of the three-dimensional effect, the neutrino fluxes arriving at the local horizontal have a greater dependence on pion production at large angles. The increased production of pions at high production angles within the data, compared to TARGET2.4, is therefore responsible for this similar increased horizontal neutrino flux.

The total uncertainty on the neutrino flux as a function of cosine zenith angle is shown in Figure 11.6a for each of these energy bins. The total uncertainty for each energy bin is calculated as the quadrature sum of the individual uncertainty




Figure 11.6: Left: The total ν_μ neutrino flux uncertainty as a function of neutrino zenith angle for each of the four neutrino energy bins shown in Figure 11.5. The total has been calculated as a quadrature sum of each of the components listed in Figure 11.2. **Right:** A breakdown of the pion dataset uncertainty contribution to the first ν_μ energy bin ($0.1\text{--}0.32$ GeV). The contribution to the total uncertainty from all the pion datasets is broken down into the component from neutrino production at high angles within the HARP high angle datasets (dashed) and the component from the remaining pion datasets (also dashed).

sources (the same as the total uncertainty in Figure 11.2). For brevity a breakdown of each component for each energy bin as a function of cosine zenith angle will not be shown. Figure 11.6b shows the total uncertainty due to pion datasets for the first ($0.1\text{--}0.32$ GeV) neutrino energy bin which is then separately broken down into the uncertainty on the HARP high angle datasets and the uncertainty on the remaining datasets used to model pion production. As a consequence of the three-dimensional effect discussed above, it can be seen that the neutrino flux at ($\cos \theta_Z \sim 0$) has an increased uncertainty due to its dependence on pion production at larger angles.

11.3.2 Neutrino Azimuthal Angle Dependence

The ν_μ neutrino flux is similarly shown as a function of azimuthal angle at the detector in Figure 11.7 with the uncertainties on each neutrino energy bin given in Figure 11.8.




Figure 11.7: The ν_μ neutrino fluxes shown as a function of neutrino azimuthal angle at the detector measured in degrees and binned in four neutrino energy bins. As in Figure 11.5 the fluxes obtained using the methods described within this thesis are shown alongside the 2004 Bartol flux prediction. In addition to the tuning, the longer Monte Carlo run completed within this thesis reduces the statistical fluctuations visible at higher neutrino energies within the original three dimensional flux prediction.




Figure 11.8: The total uncertainty prediction on the ν_μ flux shown in Figure 11.7 as a function of neutrino zenith angle. The total has been calculated as a quadrature sum of each of the components listed in Figure 11.2.

11.4 Different Detector Locations

Finally the Bartol Monte Carlo was re-run at two different detector locations to compare to the fluxes at Kamioka. The hadron production and primary cosmic ray tuning developed within this thesis to estimate the neutrino fluxes at Kamioka was applied to calculate the neutrino flux at these alternative locations. For brevity only the angle averaged fluxes at each detector site will be shown here to illustrate the key differences between the flux at each detector site. The models produced within this thesis can also be used to improve the uncertainty on the atmospheric neutrino modelling at each of these detector sites also. Figure 11.9 shows the neutrino flux estimate at the Homestake mine in North America - the location of the future DUNE detector [114] and the location of the IceCube experiment at the South Pole [115]. As can be seen, there are large differences between the neutrino fluxes at these detector sites compared to Kamioka. These differences are caused by the differences in geomagnetic field at each detector site, as briefly discussed in Chapter 2. Kamioka, which is near the geomagnetic equator, has a large geomagnetic cut-off, suppressing the lower energy cosmic ray flux locally to the detector. The geomagnetic cut-off local to the DUNE and IceCube detectors, which are both closer to geomagnetic poles, is smaller and hence the low energy cosmic ray flux is less suppressed. Consequently, the low energy neutrino fluxes are much larger at DUNE and IceCube than Kamioka. At higher neutrino energies where the cut-off of lower energy primary cosmic rays is less important, the fluxes at each detector site can be seen to converge. This effect has previously been demonstrated in comparisons of the Bartol flux at the Soudan and Kamioka detectors [46]. Of note here is the increased contribution of lower energy cosmic rays to the flux of lower energy neutrinos at DUNE and IceCube which will cause an increase in the importance of hadron production at lower energies to the modelling of these fluxes. The work within this thesis has found that hadron production at lower energies have higher uncertainties due to dataset disagreements and smaller dataset phase space coverage. The increased importance of hadron production at lower energies




Figure 11.9: The ν_μ and ν_e solid angle averaged atmospheric neutrino fluxes calculated using the hadron production modelling described within this thesis shown for three different detector sites: Kamioka, DUNE and IceCube.

at these alternative locations will cause an increase in the systematic uncertainties when compared to neutrino fluxes at Kamioka.

12

Summary, Conclusions and Future Prospects

Within this thesis, I have presented a prediction for the un-oscillated atmospheric neutrino flux which is tuned to hadron production datasets and incorporates a new model for primary cosmic ray fluxes. This work updates the 2004 atmospheric neutrino flux estimate published by the Bartol group. In addition, I have conducted a detailed study of the uncertainties of the newly developed hadron production model and propagated them through to the neutrino fluxes.

Chapter 4 explored the relevant hadron interaction phase space for neutrino production within atmospheric showers and the fixed target hadron production experimental datasets that cover that interaction phase space. Within Chapter 5 I explored the use of a parameterisation to model pion production, fitting the parameterisation to the datasets and exploring the extent to which the datasets cover the phase space of the hadronic interactions within atmospheric showers. Significant discrepancies between the pion production measurements at lower energies were found and I developed an estimate for the uncertainties in the phase space covered by these measurements which accounts for the differences between the data points. This

methodology was extended to model kaon production in Chapter 6. The modelling of hadron production outside the phase space regions covered by the datasets by different hadron interaction generators is explored in Chapter 8. The pion and kaon production interactions are weighted to this newly developed hadronic interaction model which is used together with the cosmic ray parameterisation from Evans et al. [113] to calculate the neutrino flux. A large number of high-quality hadron production measurements with large phase space coverage have been published since the 2004 neutrino flux calculation. This calculation provides, for the first time, an atmospheric neutrino flux prediction that is predominantly based on dataset measurements rather than hadron generator Monte Carlo predictions.

The uncertainties due to the modelling of pion and kaon production are described in Chapter 9 and the uncertainty due to the cosmic ray model is described in Chapter 10. As these uncertainties are smaller than before, additional previously subdominant uncertainties needed to be considered for the first time. Chapter 9 explores the uncertainty due to the interaction of secondary mesons within atmospheric showers and the re-interaction of baryons produced within atmospheric showers. Combining all the uncertainties in quadrature the uncertainty on the neutrino flux in the 0.1 and 1 GeV energy range was found to be $\sim 8 - 10\%$, decreasing to $\sim 5\%$ for the flux between 1-10 GeV. This represents a significant reduction from the previous $\sim 15\%$ uncertainty across the 0.1-10 GeV neutrino energy range within the 2004 Bartol flux prediction.

The higher uncertainty on the lower energy neutrino flux is caused by a combination of the discrepancy between existing measurements of pion production at lower interaction energies, uncertainty in the treatment of different analyses of the HARP high angle data and the smaller coverage of pion production phase space relevant to atmospheric neutrino production at lower interaction energies. Further hadron production measurements at low interaction energy (< 31 GeV) and covering wider phase spaces than the existing datasets could resolve discrepancies and further reduce the neutrino flux uncertainty. In addition to the work presented in this thesis I have worked alongside two fixed target hadron production experiments

planning to measure hadron production at low interaction energies: EMPHATIC [116] and NA61/SHINE [117], to clarify the phase space regions where newer measurements would improve neutrino flux uncertainties. EMPHATIC plans to run at beam momenta between 2 and 20 GeV alongside higher energy runs, and NA61/SHINE similarly plans measurements within the 1 to 20 GeV region. These future measurements promise to further lower the neutrino flux prediction uncertainty. In addition, new measurements of hadron production at 60 and 120 GeV taken by the NA61/SHINE collaboration which are currently in analysis and soon to be released [117] will reduce the uncertainty of the interpolation between hadron production at 31 and 158 GeV discussed in Chapter 9.

The tuned model and uncertainties developed within this thesis can now be used within neutrino oscillation analyses to better constrain the CP violating phase δ and the neutrino mass ordering. In future, the accuracy of the experimentally measured atmospheric neutrino flux will continue to improve as existing neutrino detectors have longer live times and with the development of the DUNE [114] and Hyper-Kamiokande neutrino detectors [118]. The combination of the reduction of the un-oscillated flux uncertainty demonstrated within this thesis, combined with improved accuracy in neutrino flux measurement will further lower the uncertainty of the oscillation parameters.

References

- [1] W. Pauli. “Dear radioactive ladies and gentlemen”. *Phys. Today* 31N9 (1978), p. 27.
- [2] E. Fermi. “An attempt of a theory of beta radiation. 1.” *Z. Phys.* 88 (1934), pp. 161–177.
- [3] H. Bethe and R. Peierls. “The ‘neutrino’”. *Nature* 133 (1934), p. 532.
- [4] F. Reines and C. L. Cowan. “Detection of the free neutrino”. *Phys. Rev.* 92 (1953), pp. 830–831.
- [5] G. Danby et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. *Phys. Rev. Lett.* 9 (1962), pp. 36–44.
- [6] Martin L. Perl et al. “Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation”. *Phys. Rev. Lett.* 35 (1975), pp. 1489–1492.
- [7] K. Kodama et al. “Observation of tau neutrino interactions”. *Phys. Lett. B* 504 (2001), pp. 218–224. arXiv: {hep-ex/0012035}.
- [8] Raymond Davis Jr., Don S. Harmer, and Kenneth C. Hoffman. “Search for neutrinos from the sun”. *Phys. Rev. Lett.* 20 (1968), pp. 1205–1209.
- [9] B. T. Cleveland et al. “Measurement of the solar electron neutrino flux with the Homestake chlorine detector”. *Astrophys. J.* 496 (1998), 505–526.
- [10] K. S. Hirata et al. “Experimental Study of the Atmospheric Neutrino Flux”. *Phys. Lett. B* 205 (1988). Ed. by J. Tran Thanh Van, p. 416.
- [11] D. Casper et al. “Measurement of atmospheric neutrino composition with IMB-3”. *Phys. Rev. Lett.* 66 (1991), 2561–2564.
- [12] B. Pontecorvo. “Inverse beta processes and nonconservation of lepton charge”. *Zh. Eksp. Teor. Fiz.* 34 (1957), p. 247.
- [13] Mark Thomson. *Modern particle physics*. New York: Cambridge University Press, 2013.
- [14] R. L. Workman et al. “Review of Particle Physics”. *PTEP* 2022 (2022), p. 083C01.
- [15] B. Pontecorvo. “Neutrino Experiments and the Problem of Conservation of Leptonic Charge”. *Zh. Eksp. Teor. Fiz.* 53 (1967), pp. 1717–1725.

- [16] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. “Remarks on the unified model of elementary particles”. *Prog. Theor. Phys.* 28 (1962), pp. 870–880.
- [17] Giles Barr et al. *Particle Physics in the LHC Era*. Oxford master series in particle physics, astrophysics and cosmology. Oxford: Oxford University Press, 2016.
- [18] Takaaki Kajita. “ATMOSPHERIC NEUTRINOS AND DISCOVERY OF NEUTRINO OSCILLATIONS”. *Proc. Japan Acad. B* 86 (2010), pp. 303–321.
- [19] Y. Fukuda et al. “Evidence for oscillation of atmospheric neutrinos”. *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. arXiv: hep-ex/9807003.
- [20] T. Kajita, E. Kearns, and M. Shiozawa. “Establishing atmospheric neutrino oscillations with Super-Kamiokande”. *Nucl. Phys. B* 908 (2016), pp. 14–29.
- [21] Q. R. Ahmad et al. “Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory”. *Phys. Rev. Lett.* 89 (2002), p. 011301. arXiv: nucl-ex/0204008.
- [22] Boris Kayser. “Neutrino physics”. *eConf C040802* (2004). Ed. by Joanne Hewett et al., p. L004. arXiv: hep-ph/0506165.
- [23] M. Nakahata. “Super-Kamiokande”. *10th Moriond Workshop: New and Exotic Phenomena*. 1990, pp. 605–608.
- [24] Yusuke Koshio. “Observation of Atmospheric Neutrinos”. *Universe* 6.6 (2020), p. 80.
- [25] M. Fechner et al. “Kinematic reconstruction of atmospheric neutrino events in a large water Cherenkov detector with proton identification”. *Phys. Rev. D* 79 (2009), p. 112010. arXiv: 0901.1645 [hep-ex].
- [26] K. Abe et al. “Evidence for the Appearance of Atmospheric Tau Neutrinos in Super-Kamiokande”. *Phys. Rev. Lett.* 110.18 (2013), p. 181802. arXiv: 1206.0328 [hep-ex].
- [27] K. Abe et al. “Atmospheric neutrino oscillation analysis with external constraints in Super-Kamiokande I-IV”. *Phys. Rev. D* 97.7 (2018), p. 072001. arXiv: 1710.09126 [hep-ex].
- [28] K. Abe et al. “The T2K Experiment”. *Nucl. Instrum. Meth. A* 659 (2011), pp. 106–135. arXiv: 1106.1238 [physics.ins-det].
- [29] X. Qian and P. Vogel. “Neutrino Mass Hierarchy”. *Prog. Part. Nucl. Phys.* 83 (2015), pp. 1–30. arXiv: 1505.01891 [hep-ex].
- [30] A. D. Sakharov. “Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe”. *Pisma Zh. Eksp. Teor. Fiz.* 5 (1967), pp. 32–35.

- [31] K. Abe et al. “Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations”. *Nature* 580.7803 (2020). [Erratum: *Nature* 583, E16 (2020)], pp. 339–344. arXiv: 1910.03887 [hep-ex].
- [32] M. Jiang et al. “Atmospheric Neutrino Oscillation Analysis with Improved Event Reconstruction in Super-Kamiokande IV”. *PTEP* 2019.5 (2019), 053F01. arXiv: 1901.03230 [hep-ex].
- [33] C. Bronner. “Super-Kamiokande latest results (Atmospheric neutrinos)”. <https://indico.ictp.it/event/8312/session/88/contribution/540/material/slides/0.pdf>, May 2018.
- [34] O. W. Greenberg. “CPT violation implies violation of Lorentz invariance”. *Phys. Rev. Lett.* 89 (2002), p. 231602. arXiv: hep-ph/0201258.
- [35] Hiroshi Nunokawa, Stephen J. Parke, and Jose W. F. Valle. “CP Violation and Neutrino Oscillations”. *Prog. Part. Nucl. Phys.* 60 (2008), pp. 338–402. arXiv: 0710.0554 [hep-ph].
- [36] Gaisser, Thomas K. and Engel, Ralph and Resconi, Elisa. *Cosmic Rays and Particle Physics. 2.* Cambridge University Press, 2016.
- [37] T. K. Gaisser and M. Honda. “Flux of atmospheric neutrinos”. *Ann. Rev. Nucl. Part. Sci.* 52 (2002), pp. 153–199. arXiv: hep-ph/0203272.
- [38] Anatoli Fedynitch et al. “Hadronic interaction model sibyll 2.3c and inclusive lepton fluxes”. *Phys. Rev. D* 100.10 (2019), p. 103018. arXiv: 1806.04140 [hep-ph].
- [39] L. Pasquali and M. H. Reno. “Tau-neutrino fluxes from atmospheric charm”. *Phys. Rev. D* 59 (1999), p. 093003. arXiv: hep-ph/9811268.
- [40] Malcolm S. Longair. *High Energy Astrophysics.* 3rd ed. Cambridge University Press, 2011.
- [41] Grace Adagba and Charles Wang. “The effects of solar wind on galactic cosmic ray flux at Earth”. *Astrophysics and Space Science* 361 (Dec. 2015).
- [42] John A. Simpson. “The Cosmic Ray Nucleonic Component: The Invention and Scientific Uses of the Neutron Monitor - (Keynote Lecture)”. *Space Science Reviews* 93 (July 2000), pp. 11–32.
- [43] E. Richard et al. “Measurements of the atmospheric neutrino flux by Super-Kamiokande: energy spectra, geomagnetic effects, and solar modulation”. *Phys. Rev. D* 94.5 (2016), p. 052001. arXiv: 1510.08127 [hep-ex].
- [44] S. Robbins. “Atmospheric neutrino predictions and the influence of hadron production”. PhD thesis. Oxford U., 2004.

- [45] Paolo Lipari. “The East west effect for atmospheric neutrinos”. *Astropart. Phys.* 14 (2000), pp. 171–188. arXiv: [hep-ph/0003013](#).
- [46] G. D. Barr et al. “A Three - dimensional calculation of atmospheric neutrinos”. *Phys. Rev. D* 70 (2004), p. 023006. arXiv: [astro-ph/0403630](#).
- [47] G. D. Barr et al. “Uncertainties in Atmospheric Neutrino Fluxes”. *Phys. Rev. D* 74 (2006), p. 094009. arXiv: [astro-ph/0611266](#).
- [48] G. Barr, T. K. Gaisser, and T. Stanev. “Flux of Atmospheric Neutrinos”. *Phys. Rev. D* 39 (1989), pp. 3532–3534.
- [49] T. K. Gaisser et al. “The Flux of Atmospheric Neutrinos”. *Phys. Rev. Lett.* 51 (1983), pp. 223–226.
- [50] Paolo Lipari. “The Geometry of atmospheric neutrino production”. *Astropart. Phys.* 14 (2000), pp. 153–170. arXiv: [hep-ph/0002282](#).
- [51] Jean Favier, R. Kossakowski, and J. P Vialle. “A 3-D simulation of the atmospheric neutrinos”. *Phys. Rev. D* 68 (2003), p. 093006. arXiv: [astro-ph/0305460](#).
- [52] T Gaisser et al. “Primary spectrum to 1 TeV and beyond” (Aug. 2001).
- [53] Stephen M. Barr et al. “Ratio of ν_e/ν_μ in Atmospheric Neutrinos”. *Phys. Lett. B* 214 (1988), pp. 147–150.
- [54] H. G. Fischer et al. “A comprehensive study of the inclusive production of negative pions in p+p collisions for interaction energies from 3 GeV to 13 TeV covering the non-perturbative sector of the Strong Interaction” (Feb. 2022). arXiv: [2202.09137 \[hep-ex\]](#).
- [55] R. Devenish and A. Cooper-Sarkar. *Deep inelastic scattering*. 2004.
- [56] N. Abgrall et al. “Measurements of $\pi^\pm$, $K^\pm$, K_S^0 , Λ and proton production in proton-carbon interactions at 31 GeV/c with the NA61/SHINE spectrometer at the CERN SPS”. *Eur. Phys. J. C* 76.2 (2016), p. 84. arXiv: [1510.02703 \[hep-ex\]](#).
- [57] K. Abe et al. “T2K neutrino flux prediction”. *Phys. Rev. D* 87.1 (2013). [Addendum: *Phys.Rev.D* 87, 019902 (2013)], p. 012001. arXiv: [1211.0469 \[hep-ex\]](#).
- [58] A. A. Aguilar-Arevalo et al. “The Neutrino Flux prediction at MiniBooNE”. *Phys. Rev. D* 79 (2009), p. 072002. arXiv: [0806.1449 \[hep-ex\]](#).
- [59] C. Alt et al. “Inclusive production of charged pions in p+C collisions at 158-GeV/c beam momentum”. *Eur. Phys. J. C* 49 (2007), pp. 897–917. arXiv: [hep-ex/0606028](#).

- [60] Magdalena Zofia Posiadala. “Measurements of Hadronic Cross Section for Precise Determination of Neutrino Beam Properties in T2K Oscillation Experiment”. 2011. URL: <http://cds.cern.ch/record/2240138>.
- [61] C. Alt et al. “Inclusive production of charged pions in p+p collisions at 158-GeV/c beam momentum”. *Eur. Phys. J. C* 45 (2006), pp. 343–381. arXiv: [hep-ex/0510009](#).
- [62] Richard P. Feynman. “Very high-energy collisions of hadrons”. *Phys. Rev. Lett.* 23 (1969). Ed. by L. M. Brown, pp. 1415–1417.
- [63] M. G. Catanesi et al. “The HARP detector at the CERN PS”. *Nucl. Instrum. Meth. A* 571 (2007), pp. 527–561.
- [64] M. Apollonio et al. “Forward production of charged pions with incident protons on nuclear targets at the CERN PS”. *Phys. Rev. C* 80 (2009), p. 035208. arXiv: [0907.3857 \[hep-ex\]](#).
- [65] M. G. Catanesi et al. “Large-angle production of charged pions by 3-GeV/c - 12-GeV/c protons on carbon, copper and tin targets”. *Eur. Phys. J. C* 53 (2008), pp. 177–204. arXiv: [0709.3464 \[hep-ex\]](#).
- [66] A. Bolshakova et al. “Cross-sections of large-angle hadron production in proton- and pion-nucleus interactions VII: tin nuclei and beam momenta from $\sqrt{s}$ 3 GeV/c to $\sqrt{s}$ 15 GeV/c”. *Eur. Phys. J. C* 71 (2011), p. 1719. arXiv: [1105.2239 \[hep-ex\]](#).
- [67] I. Chemakin et al. “Pion production by protons on a thin beryllium target at 6.4-GeV/c, 12.3-GeV/c, and 17.5-GeV/c incident proton momenta”. *Phys. Rev. C* 77 (2008). [Erratum: *Phys.Rev.C* 77, 049903 (2008)], p. 015209. arXiv: [0707.2375 \[nucl-ex\]](#).
- [68] Chemakin, I, Cianciolo, V, Cole, B A, Fernow, R C, Frawley, A D, Gilkes, M, Gushue, S, Hartouni, E P, Hiejima, H, Justice, M, Kang, J H, Kirk, H G, Link, J M, Maeda, N, McGrath, R L, Mioduszewski, S, Monroe, J, Morrison, D, Moulson, M, and Namboodiri, M N. “Erratum: Pion production by protons on a thin beryllium target at 6.4, 12.3, and 17.5 GeV/c incident proton momenta [*Phys. Rev. C* 77, 015209 (2008)]”. 77 (4 Apr. 2008).
- [69] T. Abbott et al. “Measurement of particle production in proton induced reactions at 14.6-GeV/c”. *Phys. Rev. D* 45 (1992), pp. 3906–3920.
- [70] J. V Allaby et al. “High-energy particle spectra from proton interactions at 19.2-GeV/c” (Apr. 1970).

- [71] T. Eichten et al. “Particle production in proton interactions in nuclei at 24 GeV/c”. *Nuclear Physics B* 44.2 (1972), pp. 333–343. URL: <https://www.sciencedirect.com/science/article/pii/0550321372901204>.
- [72] G. Bellettini et al. “Proton-nuclei cross sections at 20 GeV”. *Nucl. Phys.* 79 (1966), pp. 609–624.
- [73] G. Ambrosini et al. “Pion yield from 450-GeV/c protons on beryllium”. *Phys. Lett. B* 425 (1998), pp. 208–214.
- [74] G. Ambrosini et al. “Measurement of charged particle production from 450-GeV/c protons on beryllium”. *Eur. Phys. J. C* 10 (1999), pp. 605–627.
- [75] M. Bonesini et al. “On Particle production for high-energy neutrino beams”. *Eur. Phys. J. C* 20 (2001), pp. 13–27. arXiv: [hep-ph/0101163](https://arxiv.org/abs/hep-ph/0101163).
- [76] Gemma Maria Tinti. “Sterile neutrino oscillations in MINOS and hadron production in pC collisions”. PhD thesis. 2010.
- [77] F. Halzen and Alan D. Martin. *QUARKS AND LEPTONS: AN INTRODUCTORY COURSE IN MODERN PARTICLE PHYSICS*. 1984.
- [78] Tomislav Vladislavljevic. “Predicting the T2K neutrino flux and measuring oscillation parameters”. PhD thesis. Oxford University, Oxford U., 2020.
- [79] D. S. Barton et al. “Experimental Study of the a -Dependence of Inclusive Hadron Fragmentation”. *Phys. Rev. D* 27 (1983), p. 2580.
- [80] R. L. Workman et al. “Review of Particle Physics”. *PTEP* 2022 (2022), p. 083C01.
- [81] F. E. Taylor et al. “Analysis of radial scaling in single-particle inclusive reactions”. *Phys. Rev. D* 14 (5 Sept. 1976), pp. 1217–1242. URL: <https://link.aps.org/doi/10.1103/PhysRevD.14.1217>.
- [82] E. Yen. “New scaling variable and early scaling in single-particle inclusive distributions for hadron-hadron collisions”. *Phys. Rev. D* 10 (3 Aug. 1974), pp. 836–843. URL: <https://link.aps.org/doi/10.1103/PhysRevD.10.836>.
- [83] F. James and M. Roos. “Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations”. *Comput. Phys. Commun.* 10 (1975), pp. 343–367.
- [84] A. Fiorentini et al. *Flux Prediction and Uncertainty Updates with NA61 2009 Thin Target Data and Negative Focussing Mode Predictions*. Tech. rep. 2018.
- [85] A. M. Rossi et al. “Experimental Study of the Energy Dependence in Proton Proton Inclusive Reactions”. *Nucl. Phys. B* 84 (1975), pp. 269–305.

- [86] A. Aduszkiewicz et al. “Measurements of hadron production in $\pi^+ + \text{C}$ and $\pi^+ + \text{Be}$ interactions at 60 GeV/c”. *Phys. Rev. D* 100.11 (2019), p. 112004. arXiv: 1909.06294 [hep-ex].
- [87] Stefano Cecchini and Maurizio Spurio. “Atmospheric muons: experimental aspects” (Aug. 2012). arXiv: 1208.1171 [astro-ph.EP].
- [88] Vladimir Ivanchenko et al. “Hadronic Physics in Geant4: Improvements and Status for LHC Start”. *PoS* (Jan. 2008).
- [89] Julia Yarba. *Geant4 Physics Lists: Optimizing for Hadron Production on Light Targets*. Meeting on Neutrino Flux Predictions. Sept. 22, 2014. URL: <https://indico.fnal.gov/event/8863/contributions/110280/attachments/71808/86169/G4-HAD-physicslists-NBI-talk.pdf> (visited on 05/30/2022).
- [90] Geant4. *Geant4 Reference Physics Lists*. CERN. May 30, 2022. URL: <https://geant4.web.cern.ch/node/155> (visited on 05/30/2022).
- [91] Dennis Wright (SLAC). *Physics I: Physics Lists*. Geant4 Tutorial at Stanford. Mar. 4, 2014. URL: <https://www.slac.stanford.edu/xorg/geant4/SLACTutorial14/Physics1.pdf> (visited on 05/30/2022).
- [92] A Ferrari et al. *FLUKA: A multi-particle transport code (program version 2005)*. CERN Yellow Reports: Monographs. Geneva: CERN, 2005. URL: <https://cds.cern.ch/record/898301>.
- [93] John Allison et al. “Geant4 developments and applications”. *IEEE Trans. Nucl. Sci.* 53 (2006), p. 270.
- [94] J. Allison et al. “Recent developments in Geant4”. *Nucl. Instrum. Meth. A* 835 (2016), pp. 186–225.
- [95] S. Agostinelli et al. “GEANT4—a simulation toolkit”. *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303.
- [96] Dennis H. Wright et al. “Low and high energy modeling in GEANT4”. *AIP Conf. Proc.* 896.1 (2007). Ed. by M. Albrow and R. Raja, pp. 11–20.
- [97] V. V. Uzhinsky. “The Fritiof (FTF) Model in Geant4”. *International Conference on Calorimetry for the High Energy Frontier*. 2013, pp. 260–264.
- [98] Alfredo Ferrari et al. “FLUKA: A multi-particle transport code (Program version 2005)” (Oct. 2005).
- [99] T. T. Böhlen et al. “The FLUKA Code: Developments and Challenges for High Energy and Medical Applications”. *Nucl. Data Sheets* 120 (2014), pp. 211–214.

- [100] Paola R. Sala. *The physics of the FLUKA code: The physics of the FLUKA code: hadronic models*. Fluka. Sept. 6, 2006. URL: <https://indico.cern.ch/event/3734/contributions/883244/attachments/743579/1020017/fluka-models.pdf> (visited on 05/30/2022).
- [101] James Gentle. *Computational Statistics*. Jan. 2009.
- [102] M. Apollonio et al. “Measurements of Forward Proton Production with Incident Protons and Charged Pions on Nuclear Targets at the CERN Proton Synchrotron”. *Phys. Rev. C* 82 (2010), p. 045208. arXiv: 1006.1223 [hep-ex].
- [103] T. Anticic et al. “Inclusive production of protons, anti-protons and neutrons in p+p collisions at 158-GeV/c beam momentum”. *Eur. Phys. J. C* 65 (2010), pp. 9–63. arXiv: 0904.2708 [hep-ex].
- [104] Pasquale Blasi. “The Origin of Galactic Cosmic Rays”. *Astron. Astrophys. Rev.* 21 (2013), p. 70. arXiv: 1311.7346 [astro-ph.HE].
- [105] D. J. Bird et al. “Detection of a cosmic ray with measured energy well beyond the expected spectral cutoff due to cosmic microwave radiation”. *Astrophys. J.* 441 (1995), pp. 144–150. arXiv: astro-ph/9410067.
- [106] E. Richard et al. “Measurements of the atmospheric neutrino flux by Super-Kamiokande: energy spectra, geomagnetic effects, and solar modulation”. *Phys. Rev. D* 94.5 (2016), p. 052001. arXiv: 1510.08127 [hep-ex].
- [107] Benedetto D’Ettorre Piazzoli and Si-Ming Liu and Domenico della Volpe and Zhen Cao and Andrea Chiavassa and Benedetto D’Ettorre Piazzoli and Yi-Qing Guo and Leonid T. Ksenofontov and Olivier Martineau-Huynh and Diane Martraire and Ling-Ling Ma and Xin-Hua Ma and Yuri Stenkin and Qiang Yuan and Hou-Dun Zeng and Shou-Shan Zhang and Yi Zhang and Hui Zhu. “Chapter 4 Cosmic-Ray Physics”. *Chinese Physics C* 46.3 (Mar. 2022), p. 030004. URL: <https://doi.org/10.1088/1674-1137/ac3faa%7D>.
- [108] Thomas K. Gaisser. “Spectrum of cosmic-ray nucleons, kaon production, and the atmospheric muon charge ratio”. *Astropart. Phys.* 35 (2012), pp. 801–806. arXiv: 1111.6675 [astro-ph.HE].
- [109] O. Adriani et al. “PAMELA Measurements of Cosmic-ray Proton and Helium Spectra”. *Science* 332 (2011), pp. 69–72. arXiv: 1103.4055 [astro-ph.HE].
- [110] M. Aguilar et al. “Precision Measurement of the Proton Flux in Primary Cosmic Rays from Rigidity 1 GV to 1.8 TV with the Alpha Magnetic Spectrometer on the International Space Station”. *Phys. Rev. Lett.* 114 (2015), p. 171103.

- [111] T. Sanuki et al. “Precise measurement of cosmic ray proton and helium spectra with the BESS spectrometer”. *Astrophys. J.* 545 (2000), p. 1135. arXiv: astro-ph/0002481.
- [112] Y. S. Yoon et al. “Cosmic-Ray Proton and Helium Spectra from the First CREAM Flight”. *Astrophys. J.* 728 (2011), p. 122. arXiv: 1102.2575 [astro-ph.HE].
- [113] Justin Evans et al. “Uncertainties in Atmospheric Muon-Neutrino Fluxes Arising from Cosmic-Ray Primaries”. *Phys. Rev. D* 95.2 (2017), p. 023012. arXiv: 1612.03219 [astro-ph.HE].
- [114] R. Acciarri et al. “Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 2: The Physics Program for DUNE at LBNF” (Dec. 2015). arXiv: 1512.06148 [physics.ins-det].
- [115] M. G. Aartsen et al. “The IceCube Neutrino Observatory: Instrumentation and Online Systems”. *JINST* 12.03 (2017), P03012. arXiv: 1612.05093 [astro-ph.IM].
- [116] T. Akaishi et al. “EMPHATIC: A Proposed Experiment to Measure Hadron Scattering and Production Cross Sections for Improved Neutrino Flux Predictions” (Dec. 2019). arXiv: 1912.08841 [hep-ex].
- [117] Yoshikazu Nagai. “NA61/SHINE Experiment Status and Plans after Long Shutdown 2”. Jan. 2021. URL: %7Bhttps://indico.cern.ch/event/873509/contributions/3690202/attachments/2170145/3663837/NA61%5C_AtμNuWorkshop%5C_YNagai.pdf%7D.
- [118] K. Abe et al. “Physics potential of a long-baseline neutrino oscillation experiment using a J-PARC neutrino beam and Hyper-Kamiokande”. *PTEP* 2015 (2015), p. 053C02. arXiv: 1502.05199 [hep-ex].