

Reconstruction and Analysis for the DUNE 35-ton Liquid Argon Prototype



Michael Wallbank

Supervisor: Prof. Lee Thompson

Department of Physics & Astronomy
University of Sheffield

This dissertation is submitted for the degree of
Doctor of Philosophy

February 2018

Declaration

The material presented in this document has not been published or submitted for examination at another institute. It represents original work by the author unless stated otherwise, complete with references.

The work on the Deep Underground Neutrino Experiment (DUNE), described in Chapter 3, and within the Liquid Argon Software Framework (LArSoft), described in Chapter 6, was made possible due to a large number of collaborators both within and external to the DUNE and LArSoft experiments. Citations have been used where possible to reference work not done by the author. When this is not possible, the people responsible have been explicitly named.

The author contributed directly to the success of the 35-ton Phase II prototype, described in Chapter 4, alongside a number of collaborators. This consisted of involvement in all aspects of the experiment including testing, construction, installation, commissioning, data taking and data analysis, and serving as a deputy run coordinator during the commissioning and running period. The author also took sole responsibility for the Online Monitoring framework, described in Chapter 5, and in the leading of some of the main analyses from the data, discussed in Chapter 7.

The shower reconstruction development comprising the second half of Chapter 6 is original work by the author, along with its implementation in the DUNE ν_e CC selection, the subject of Chapter 8. The author notes the work of Dr Tingjun Yang (Fermi National Accelerator Laboratory) and Tyler Alion (University of Sussex) for the initial development of the selection, and Dr Dominic Brailsford (Lancaster University) for the relevant contributions from a parallel analysis.

Michael Wallbank
February 2018

Acknowledgements

First and foremost, I would like to thank my supervisor Prof. Lee Thompson for his support over the last few years and for giving me the opportunity to pursue a career in particle physics. Always happy to chat, even when on different continents, and always affording me the space to work freely and independently whilst reigning me in whenever my ideas got a little too mad, his supervision has really enabled me to work to my potential over the past four years.

I was privileged to spend a year and a half of my Ph.D. out at Fermilab and found it an incredible place to live and work. I am very grateful to Dr. Michelle Stancari, Dr. Tom Junk and Dr. Tingjun Yang for giving me all the opportunities and for being some of the most fun and productive people to work with. I also hugely enjoyed being a part of the 35-ton operations team and would like to extend my thanks to Dr. Alan Hahn for his endless knowledge and conversation.

I have loved being part of the Sheffield group and would like to thank everyone for making it such a great place to work. In particular, to those I have ever shared an office with and to those who congregate at the University Arms every Friday evening, I have many fond, crazy memories. Especial gratitude must be extended to Dr. Jon Perkin for his endless support during my first couple of years and for so memorably introducing me to karaoke in Tokai!

I would like to express my gratitude to the Science & Technology Facilities Council for funding my Ph.D. and giving me the opportunity to travel and learn as much as I have. I am also hugely grateful to the Neutrino Division at Fermilab for paying my rent for almost a year, allowing me to stay at the lab for as long as I needed. Additionally, I would like to thank Dr. Steve Boyd for introducing me to experimental neutrino physics during my Masters project and inspiring me to pursue research in this field during my Ph.D.

It's an endless source of relief we don't do this alone, and there are so many others I'm grateful to for helping make the past four years some of the most fun. In no particular order, to Calum, Dave, Manu, Andy, Matt, Amy and everyone from Sheffield, including all those I lived with in Rivelin; to Rob, Oksana, Tristan, Luke, David, Alex, Rob, all my other neighbours on Fermilab village, those from Social Soccer and everyone else I shared American life with; you guys are awesome and I'm grateful to you all for the friendships.

Karl deserves his own sentence because he treated me to one and, given he shared an office, house and car with me for the best part of three years and put up with ‘the laugh’ for as long as he did, few can argue it is not merited.

Next, to my friends from Southport and Warwick and to all others who have provided such strong friendships and support over many years, thank you. Most of this involves frequenting pubs and I very much appreciate the roles of The Zetland in Southport, The Benjamin Satchwell in Leamington, The Broomhill Tavern in Sheffield and Quigley’s in Naperville in providing all the fun and community (and beer).

Special thanks must be extended to my parents, and my brother, for their love and support throughout my entire life. I would never have got to the stage where I’m writing this mammoth document without the opportunities you have afforded me, so I hope this thesis goes some way towards showing this gratitude!

Finally, to my wife, Jen. How you have put up with me, and all the physics, over the last eight years is beyond me but I’m constantly amazed by the level of support and patience, particularly throughout the years we’ve been apart, you have shown. Thank you for everything you’ve done for me, even when I haven’t been appreciative, and for all the fun times and companionship – this achievement is as much yours as mine!

Abstract

Neutrino physics is approaching the precision era, with current and future experiments aiming to perform highly accurate measurements of the parameters which govern the phenomenon of neutrino oscillations. The ultimate ambition with these results is to search for evidence of CP-violation in the lepton sector, currently hinted at in the world-leading analyses from present experiments, which may explain the dominance of matter over antimatter in the Universe.

The Deep Underground Neutrino Experiment (DUNE) is a future long-baseline experiment based at Fermi National Accelerator Laboratory (FNAL), with a far detector at the Sanford Underground Research Facility (SURF) and a baseline of 1300 km. In order to make the required precision measurements, the far detector will consist of 40 kton liquid argon and an embedded time projection chamber. This promising technology is still in development and, since each detector module is around a factor 15 larger than any previous experiment employing this design, prototyping the detector and design choices is critical to the success of the experiment. The 35-ton experiment was constructed for this purpose and will be described in detail in this thesis. The outcomes of the 35-ton prototype are already influencing DUNE and, following the successes and lessons learned from the experiment, confidence can be taken forward to the next stage of the DUNE programme.

The main oscillation signal at DUNE will be electron neutrino appearance from the muon neutrino beam. High-precision studies of these ν_e interactions requires advanced processing and event reconstruction techniques, particularly in the handling of showering particles such as electrons and photons. Novel methods developed for the purposes of shower reconstruction in liquid argon are presented with an aim to successfully develop a selection to use in a ν_e charged-current analysis, and a first-generation selection using the new techniques is presented.

Table of contents

List of figures	xiii
List of tables	xxiii
List of acronyms	xxv
1 Introduction	1
2 Neutrino Physics	3
2.1 Historical Context	3
2.1.1 Prediction of the Neutrino	3
2.1.2 Discovery of the Neutrino	4
2.1.3 The Solar Neutrino Problem	8
2.1.4 The Atmospheric Neutrino Anomaly	9
2.2 Neutrino Oscillations	10
2.2.1 The Evidence for Neutrino Oscillations	11
2.2.2 Vacuum Oscillations	15
2.2.3 Matter Effects	17
2.2.4 CP Violation	18
2.3 Status of Neutrino Physics	19
2.3.1 Current and Future Experiments	19
2.3.2 Oscillation Parameters	20
2.3.3 Neutrino Mass	21
3 The Deep Underground Neutrino Experiment	25
3.1 The LAr TPC Concept	25
3.1.1 A Brief History of Time (Projection Chambers)	26
3.1.2 LAr TPC Operation	28
3.1.3 LArTPC Challenges	31

3.2	Overview of DUNE	32
3.3	Experimental Details	35
3.3.1	Beam	36
3.3.2	Far Detector	37
3.3.3	Near Detector	43
3.4	The Physics of DUNE	47
3.4.1	Mass Hierarchy and CP Violation	48
3.4.2	Oscillation Parameters	50
3.4.3	Other Physics	50
3.5	The Road to DUNE	54
3.5.1	The 35-ton Prototype	54
3.5.2	ProtoDUNE	54
4	The DUNE 35-ton Prototype	57
4.1	The Materials Test Stand and Liquid Argon Purity Demonstrator	58
4.1.1	The Materials Test Stand	58
4.1.2	The Liquid Argon Purity Demonstrator	62
4.1.3	LongBo	65
4.2	35-ton Experiment: Phase I	69
4.2.1	The 35-ton Cryostat	69
4.2.2	Filling the 35-ton	71
4.2.3	Outcomes of Phase I	72
4.3	35-ton Experiment: Phase II	73
4.3.1	The 35-ton Detector	74
4.3.2	Data Acquisition	79
4.3.3	The Sheffield Camera System	86
4.3.4	Phase II Run	91
4.3.5	Outcomes of Phase II	93
4.4	Summary	101
5	Online Monitoring and Event Displays for the 35-ton Experiment	103
5.1	The Online Monitoring Framework	103
5.1.1	Design of the Monitoring Framework	104
5.1.2	Interface with the DAQ Framework	105
5.1.3	Writing the Monitoring Data	106
5.2	Data Quality Monitoring	107
5.2.1	TPC Monitoring	107

5.2.2	Photon Detector Monitoring	107
5.2.3	External Counter Monitoring	108
5.2.4	General Monitoring	108
5.2.5	DQM Plots	108
5.3	Online Event Display	109
5.3.1	Selecting the Data	110
5.3.2	Representing the Data	110
5.4	Monitoring Web Interface	111
5.4.1	Automated Data Transfer	111
5.4.2	Web Page	113
5.5	Online Monitoring Summary	113
6	Reconstruction in a Liquid Argon TPC	115
6.1	The LArSoft Framework	115
6.2	The Reconstruction Chain	117
6.2.1	Raw Charge Calibration	117
6.2.2	Pattern Recognition	120
6.2.3	Calorimetry	122
6.3	Shower Reconstruction in LArTPCs	123
6.3.1	Showers Overview	124
6.3.2	BlurredCluster Algorithm	125
6.3.3	EMShower Algorithm	129
6.3.4	Track/Shower Separation	133
6.3.5	Performance of the Reconstruction	138
7	Analysis of 35-ton Data	145
7.1	Preparing 35-ton Data for Analysis	145
7.1.1	Selecting the Data	145
7.1.2	Improving Data Quality	147
7.1.3	Reconstructing Muon Tracks	147
7.1.4	Preparing Simulated Data	149
7.2	APA Gap-Crossing Muons	151
7.2.1	APA-Gap Offset Determination	153
7.2.2	Charge Deposited by APA Gap-Crossing Muons	164
7.3	APA-Crossing Muons	164
7.3.1	T0 Determination from APA Crossing Tracks	166
7.3.2	Charge Deposited by APA Crossing Tracks	178

7.3.3	Comparing Drift Regions with APA-Crossing Tracks	183
7.4	Shower Reconstruction in 35-ton Data	183
7.4.1	Data Specific Reconstruction	183
7.4.2	Calorimetry Reconstruction	186
7.4.3	Shower Reconstruction	188
7.4.4	π^0 Reconstruction	191
7.5	35-ton Data Analysis Summary	192
8	The ν_e Oscillation Signal at the DUNE Far Detector	193
8.1	Far Detector Samples	194
8.2	Cut-Based Tuning	194
8.2.1	Selection	194
8.2.2	Fiducial Volume Tuning	196
8.3	MVA-Based Selection	198
8.3.1	MVA Input Variables	198
8.3.2	Analysis Performance	200
8.4	Summary and Outlook for Future Selections	203
9	Conclusions	205
	References	207
	Appendix A APA Gap-Crossing Muons: Gap Measurements	223
	Appendix B DUNE Far Detector ν_eCC MVA Input Variables	229

List of figures

2.1	Energy spectrum of the electron produced in beta decay.	4
2.2	Measurements of the hadron production cross-section around the Z resonance.	7
2.3	Solar neutrino energy spectra as predicted by the Standard Solar Model.	8
2.4	Demonstrating the atmospheric neutrino anomaly using the double ratio R of muon to electron neutrino events showing data divided by expectation.	10
2.5	Zenith angle distributions of μ -like and e -like events for sub-GeV and multi-GeV SuperKamiokande data sets.	12
2.6	The ratio of the number of SuperKamiokande data events to expected events in the absence of oscillations as a function of reconstructed L/E_ν	13
2.7	Comparison of the predictions of the neutrino fluxes from the Standard Solar Model with measured rates from a variety of solar neutrino experiments.	14
2.8	General scattering mechanics which occur as neutrinos pass through matter.	18
2.9	Recent global 3-neutrino oscillation analysis.	22
2.10	Demonstration of the current uncertainties in the neutrino mass.	23
3.1	Cartoon showing the configuration of the Deep Underground Neutrino Experiment.	26
3.2	The original concept of the time projection chamber particle detector, drawn by David Nygren in 1974.	27
3.3	The LArTPC detector proposed by Carlo Rubbia in 1977.	28
3.4	Effect of electric field on luminosity of ionisation electrons and scintillation light in a LArTPC.	30
3.5	Schematic demonstrating the basic operation principles of a LArTPC.	31
3.6	Demonstration of how having access to multiple oscillation maxima facilitates measurements of both the neutrino mass hierarchy and leptonic CP violation using the same experiment.	34
3.7	Longitudinal section of the LBNF beamline facility at Fermilab.	36

3.8	The fluxes of the different neutrino flavour components of the DUNE beam in neutrino and antineutrino running mode.	37
3.9	The layout of the four cryostats underground at SURF comprising the DUNE far detector.	38
3.10	The basic design of a single-phase DUNE far detector module.	39
3.11	Design of a DUNE far detector Anode Plane Assembly (APA).	40
3.12	The design of the photon detectors for the DUNE far detector.	42
3.13	The DUNE dual-phase detector (partially open).	43
3.14	The extraction, amplification and readout of the ionisation electrons through the gaseous argon phase in the DUNE dual-phase LArTPC design.	44
3.15	Schematic of the DUNE near neutrino detector fine-grained tracker design.	45
3.16	The appearance probability at a baseline of 1300 km, as a function of neutrino energy, for $\delta_{CP} = -\pi/2, 0$ and $\pi/2$, neutrinos and antineutrinos, assuming normal hierarchy.	48
3.17	Sensitivity of the DUNE experiment to the neutrino mass hierarchy.	49
3.18	Sensitivity of the DUNE experiment to leptonic CP violation.	51
3.19	The sensitivity of DUNE to the octant and value of θ_{23}	52
3.20	The sensitivity of DUNE to the oscillation parameters describing ν_e and $\bar{\nu}_e$ appearance.	53
3.21	The 35-ton cryostat and detector designed to prototype the DUNE far detector design.	55
3.22	Schematic showing the layout of the ProtoDUNEs at the CERN neutrino platform.	56
4.1	The Materials Test Stand at FNAL.	59
4.2	Results from the Materials Test Stand showing the water contamination in LAr and the corresponding electron lifetime.	60
4.3	Schematic design of the purity monitors utilised at the FNAL LAr test stands.	61
4.4	The Liquid Argon Purity Demonstrator cryostat and purification system.	63
4.5	The piston purge technique in the Liquid Argon Purity Demonstrator to remove atmospheric impurities before filling.	64
4.6	The concentration of electronegative impurities during the gas circulation stage in the Liquid Argon Purity Demonstrator following the piston purge.	65
4.7	Schematic showing the recirculation of the LAr during commissioning and operations of the Liquid Argon Purity Demonstrator.	66
4.8	The electron lifetime achieved in the Liquid Argon Purity Demonstrator during a six week run.	67

4.9	The LongBo TPC detector shown within the Liquid Argon Purity Demonstrator Cryostat.	68
4.10	The LAr purity within the Liquid Argon Purity Demonstrator cryostat with the LongBo TPC present, measured using both data from the detector and information from the purity monitors.	68
4.11	The 35-ton cryostat.	70
4.12	The network linking the 35-ton cryostat, the Liquid Argon Purity Demonstrator and the purification system at PC4, Fermilab.	71
4.13	Filling the 35-ton cryostat in four stages: piston purge, gas recirculation, cooldown and liquid filling.	72
4.14	The electron lifetime in the 35-ton cryostat measured by two purity monitors over the course of the two month Phase I run.	73
4.15	The 35-ton detector operated during Phase II of the 35-ton programme. . .	75
4.16	Photograph of the partially installed 35-ton detector.	76
4.17	The 35-ton TPC geometry.	77
4.18	Photon detector units as implemented within the 35-ton APAs.	78
4.19	The location of the external counters positioned around the outer walls and in the ceiling above the 35-ton cryostat.	79
4.20	Block diagram showing the triggering system for the 35-ton Phase II. . . .	80
4.21	Demonstration of the format used in 35-ton raw data.	82
4.22	Demonstration of how TPC data from a triggered event in a LArTPC is saved when employing a DAQ with continuous readout.	84
4.23	Overview of the <i>lbne-artdaq</i> framework used for data acquisition by the DUNE 35-ton experiment.	85
4.24	The calibration images for the 8 cameras in the Sheffield Camera System installed in the 35-ton cryostat.	87
4.25	An example camera module developed for the 35-ton Sheffield camera system.	87
4.26	Two camera modules mounted on cryo piping in the 35-ton cryostat.	89
4.27	Full system block diagram for the camera modules in the DUNE 35-ton prototype.	89
4.28	The variation in picture quality degradation in the 35-ton camera system is illustrated by the changes in Camera 1 and Camera 4 over time.	90
4.29	The broken pipe, originally part of the framework introducing gaseous argon from LAPD into the 35-ton to maintain LAr levels, which resulted in the poisoning of the whole LAr volume by allowing the introduction of air into the system.	92

4.30	The data taking period of the 35-ton Phase II experiment.	93
4.31	Event display depicting the charge deposited by an electromagnetic shower during the 35-ton Phase II run.	94
4.32	Difference between optical hit peak times and muon counter trigger times for photon detector 3 in the 35-ton photon detection system.	96
4.33	Average Optical Hit Amplitude per Event vs. Counter Pair Positions for the 35-ton photon detection system.	96
4.34	Comparison between example muons tracks observed in 35-ton data and simulation, along with example waveforms for randomly selected channels.	98
4.35	FFT of ADC values for RCE00 for two different noise states.	99
4.36	The correlation between waveforms recorded on different channel combinations for all 2048 35-ton channels.	100
5.1	The software framework designed and built for online monitoring during the 35-ton Phase II run.	104
5.2	Selection of figures made by the Data Quality Monitoring framework during 35-ton Phase II running.	109
5.3	Example online event display made by the Online Monitoring framework.	112
5.4	Schematic showing the interface between the online monitoring system and the web.	113
5.5	Demonstration of the web page developed to display information produced by the online monitoring and event display.	114
6.1	The LArSoft architecture, highlighting support for both common and experiment-specific algorithms and methods and the interfacing with other packages.	116
6.2	The LArSoft reconstruction workflow to produce 3D reconstructed objects from the raw charge read out by the anode wires.	118
6.3	Detector response for charge readout in a LArTPC.	119
6.4	The process of deconvolution and hit finding to determine the correct charge from the measured pulses on the readout wires.	119
6.5	Demonstration of 2D reconstruction in LArSoft on a simulated ν_μ CC event.	120
6.6	Demonstration of 3D reconstruction in LArSoft on a simulated ν_μ CC event.	121
6.7	Correlation between the true deposited energy in an electromagnetic shower and the total charge in the detector for electron showers in the DUNE far detector.	123
6.8	An example particle gun π^0 event in the 35-ton geometry.	125

6.9	An example 2D Gaussian kernel alongside the contributions from each dimension.	126
6.10	Simplified demonstration, in 1D, of the blurring process used when smearing charge from a hit map.	128
6.11	The output of the blurring stage of the BlurredCluster algorithm on the hit maps from two planes of the 35-ton π^0 event illustrated in Figure 6.8. . . .	128
6.12	The 2D clusters made using BlurredCluster when applied to the 35-ton π^0 event shown in Figure 6.8.	129
6.13	Demonstration of the ‘global wire’ concept in single-phase LArTPC reconstruction.	130
6.14	Comparison of the 2D shower cluster reconstruction and the 3D tracking reconstruction for a 35-ton π^0 event.	131
6.15	Demonstration of the method utilised to ensure a correct shower orientation of the shower in a given view for the purpose of shower start finding in the EMShower algorithm.	133
6.16	The output from the full 3D shower reconstruction provided by BlurredCluster/EMShower when applied to the 35-ton π^0 event shown in Figure 6.8. . .	134
6.17	An example simulated ν_e CC interaction in the DUNE far detector.	135
6.18	A simplified event topology demonstrated typical track and shower topologies. 136	
6.19	The result of applying the track/shower separation, BlurredCluster and EMShower algorithms to the example DUNE far detector ν_e CC event demonstrated in Figure 6.17.	137
6.20	The fraction of total shower particles for which a shower object is created when using BlurredCluster/EMShower reconstruction.	139
6.21	The difference between the true and reconstructed shower conversion points. 139	
6.22	The vector dot product between the true and reconstructed initial shower direction.	139
6.23	The completeness of the reconstruction shower energy when compared with the true deposited energy from simulation.	140
6.24	The dE/dx information from the start of the reconstructed shower object. . .	140
6.25	The 35-ton fully reconstructed π^0 invariant mass peak from simulation. . .	141
6.26	The 35-ton π^0 mass peak considered using a combination of reconstructed and truth information.	142
6.27	The performance of the track/shower separation and shower reconstruction when applied to DUNE ν_e CC far detector interactions.	144
7.1	Comparison between noise levels for ‘good’ and ‘bad’ 35-ton runs.	146

7.2	Simulated demonstration of the method used to correct for stuck codes in the 35-ton data.	147
7.3	The effect of applying stuck bit mitigation to a waveform as seen in raw 35-ton data.	148
7.4	The effect of removing coherent noise from all channels on a voltage regulator in the 35-ton data.	148
7.5	Demonstration of the successive stages applied to hits on collection wires in the 35-ton data in order to select hits from the through-going track associated with the particle which caused the trigger.	150
7.6	Demonstration of crossing muons tracks in the 35-ton TPC.	151
7.7	Event display showing tracks passing across APA gaps and also through the APAs.	152
7.8	Schematic showing an example APA gap-crossing track as viewed looking down from the top of the detector.	153
7.9	Schematic showing an example track crossing two drift regions offset by an unknown quantity Δz	154
7.10	Illustration of the eight gaps between the four APA frames. There are four spaces between the APAs giving rise to eight gaps, measured in the long drift (odd-numbered) and short drift (even-numbered) regions. In the figure, the even-numbered gaps represent the distance between the APAs on the back-facing side and the odd-numbered gaps the distance on the side facing. Naively, one would expect the gaps to be identical but this will not necessarily be the case when assembled in the cryostat and, since they read out data from distinct drift regions, they may be considered separately.	154
7.11	The z -offset for the DV5/DV7 gap measured in the 35-ton data.	156
7.12	Demonstration of how an x -offset in the positions of the APAs can explain the degeneracy evident in the z -offset measured using the 35-ton data.	157
7.13	Studies of the effects of offsets in the positions of the APAs in simulation.	158
7.14	The χ^2 distribution for all APA-gap traversing tracks in Δz - Δx space.	159
7.15	Demonstration of the effects of offsets in both the x - and z -directions in the determination of Δz between DV5 and DV7.	160
7.16	Extraction of the true value of Δz from the full distribution of measured z -offsets.	161
7.17	Measurement of the x -offset between DV5 and DV7 after applying the z -gap correction.	161

7.18	of the z -offset between DV5 and DV7 after applying the x -offset previously determined.	162
7.19	The number of hits, and the average reconstructed hit charge, as a function of distance of the collection point from the nearest APA gap.	165
7.20	Event display of an APA gap-crossing track, focussed on the gap region. . .	165
7.21	Method to align track segments on either side of the APAs involving minimising residuals from linear least square fit.	168
7.22	Method to align track segments on either side of the APAs involving minimising the distance between the projected intersection of each with the centre of the APAs.	168
7.23	Difference between the T_0 calculated from TPC data and the T_0 provided by the counters representing the trigger time of the through-going muon, for simulation and data.	169
7.24	Correcting for T_0 using T_0^{counter} and T_0^{TPC}	170
7.25	Demonstration of the effect observed in the 35-ton data concerning tracks crossing the APAs.	170
7.26	Attempting to correct the track segment misalignment by assuming a misunderstanding of the spacing between the collection planes.	171
7.27	Attempting to correct the track segment misalignment by assuming a misunderstanding of the positioning of the collection wires inside the detector. . .	172
7.28	Attempting to correct the track segment misalignment by assuming an incorrect drift velocity.	172
7.29	Measuring the drift velocity of the ionisation electrons by taking tracks passing through opposite counter pairs and comparing the corresponding drift distance to the drift time.	173
7.30	The T_0 -corrected drift time for hits on APA-crossing tracks.	174
7.31	Difference between the interaction time measured by the TPC data and that provided by photon detector information.	174
7.32	The difference between the timestamps recorded by the PTB and the RCEs upon receiving a trigger.	175
7.33	The distribution of the drift times of all hits on APA-crossing tracks after correcting for the APA offsets along the direction parallel to the drift direction. 177	177
7.34	Accounting for the extra discrepancy in track alignment after fixing for all the measured offsets by assuming a misunderstanding in the collection plane spacing and the z -positions of the collection wires.	178
7.35	The T_0 -corrected drift time for all hits on an APA-crossing track in simulation. 179	179

7.36	Demonstration of the electron ionisation and hit collection for APA-crossing tracks.	180
7.37	Comparison between the T0-corrected hit time distributions on APAs with and without the grounded mesh.	181
7.38	Comparison between the distribution of T0-corrected hit times for hits on wires in front of the APA frame and away from the APA frame to validate the functionality of the mesh.	182
7.39	Average lifetime-corrected charge per hit for hits on an APA-crossing track separated according to whether or not the hit was collected around the interaction time.	182
7.40	Event display of an APA-crossing track with the charge deposited as it passes through the APAs evident.	184
7.41	Using APA-crossing tracks to confirm the drift velocity is consistent between the two drift regions.	185
7.42	Comparison of noise levels between the two drift regions using APA-crossing tracks.	185
7.43	The bias in the hit selection due to a high noise level in the 35-ton data.	187
7.44	The dE/dx distribution for mips passing through the 35-ton TPC.	188
7.45	Relationship between deposited charge and energy for 35-ton data, calculated using through-going mips.	189
7.46	Result of applying the shower reconstruction on a shower observed in the 35-ton data.	190
7.47	A candidate π^0 event observed in the online event display during the run.	191
8.1	The output of a multi-variate approach to particle identification when attempting to identify electrons.	195
8.2	The process of tuning the electron cut in the simple cut-based selection by maximising the effect of CP-violation on the oscillation probabilities.	196
8.3	The efficiency and purity of the ν_e CC Pandora cut-based selection as a function of a number of kinematic variables, after applying the selection.	197
8.4	Tuning the DUNE far detector fiducial volume, y-coordinate.	198
8.5	The MVA response when training ν_e (signal) against ν_μ and ν_τ (background).	200
8.6	The efficiency and purity of the ν_e CC MVA-based selection as a function of a number of kinematic variables, before applying the selection.	201
8.7	The efficiency and purity of the ν_e CC MVA-based selection as a function of a number of kinematic variables, after applying the selection.	202

A.1	Demonstration of the measurements of the x - and z -offsets in the 35-ton DV1/DV3 gap.	224
A.2	Demonstration of the measurements of the x - and z -offsets in the 35-ton DV1/DV5 gap.	225
A.3	Demonstration of the measurements of the x - and z -offsets in the 35-ton DV3/DV7 gap.	226
A.4	Demonstration of the measurements of the x - and z -offsets in the 35-ton DV5/DV7 gap.	227
B.1	MVA input variables related to event-level information for the DUNE far detector ν_e CC analysis.	230
B.2	MVA input variables related to information about the longest reconstructed track in the event for the DUNE far detector ν_e CC analysis.	231
B.3	MVA input variables related to information about the highest energy shower in the event for the DUNE far detector ν_e CC analysis.	233

List of tables

2.1	The current best-fit values for the neutrino oscillation parameters for normal and inverted hierarchies.	21
3.1	Properties of noble liquids relevant when considering a TPC medium for a neutrino experiment.	27
3.2	Exposures anticipated for the DUNE experiment for the first few years of operation.	47
4.1	Details and dimensions of the 35-ton cryostat.	70
6.1	The performance of track/shower separation and shower reconstruction on 50 000 DUNE ν_e events.	143
7.1	Measurements of all the APA offsets determined from the 35-ton TPC data.	163
7.2	The corrected gaps between the APAs, in x and z , based on the offsets measured in the data.	163
8.1	The dimensions and tuned fiducial volume of a single DUNE far detector module.	197
8.2	The input variable used in the MVA designed to separate ν_e events from ν_μ and ν_τ events.	199

List of acronyms

ADC	Analog-to-Digital Converter
APA	Anode Plane Assembly
ArgoNeuT	Argon Neutrino Test Stand
ASIC	Application-Specific Integrated Circuit
BDT	Boosted Decision Tree
BLM	Beamline Measurement System
BNL	Brookhaven National Laboratory
CC	Charged Current
CDF	Collider Detector at Fermilab
CERN	European Organisation for Nuclear Research
CKM	Cabibbo-Kobayashi-Maskawa Matrix
CMOS	Complementary Metal-Oxide-Semiconductor
CP	Charge-Parity Symmetry
CPA	Cathode Plane Assembly
CRC	Cosmic Ray Counter
CRP	Charge Readout Plane
CRY	Cosmic Ray Generator
CVN	Convolutional Visual Network
DAQ	Data Acquisition
DESY	German Electron Synchrotron
DONUT	Direct Observation of NuTau
DP	Dual-Phase
DQM	Data Quality Monitoring
DUNE	Deep Underground Neutrino Experiment
DV	Drift Volume
ECAL	Electromagnetic Calorimeter
ES	Elastic Scattering
eV	Electron Volt

EW	East-West
FD	Far Detector
FE	Front-End
FEMB	Front-End Mother Board
FFT	Fast Fourier Transform
FGT	Fine Grained Tracker
FNAL	Fermi National Accelerator Laboratory
FV	Fiducial Volume
GALLEX	Gallium Experiment
GAr	Gaseous Argon
GEANT4	Geometry And Tracking
GENIE	Generates Events for Neutrino Interaction Experiments
HTML	Hypertext Markup Language
HV	High Voltage
ICARUS	Imaging Cosmic And Rare Underground Signals
LAGUNA-LBNO	Large Apparatus for Grand Unification, Neutrino Astrophysics and Long Baseline Neutrino Oscillations
LAPD	Liquid Argon Purity Demonstrator
LAr	Liquid Argon
LArIAT	Liquid Argon In A Test Beam
LArSoft	Liquid Argon Software Framework
LArTPC	Liquid Argon Time Projection Chamber
LBNE	Long Baseline Neutrino Experiment
LBNF	Long Baseline Neutrino Facility
LEM	Large Electron Multiplier
LEP	Large Electron-Positron Collider
LHC	Large Hadron Collider
LNG	Liquified Natural Gas
LSB	Least Significant Bit
LZ	LUX-Zeplin
MC	Monte Carlo
MicroBooNE	Micro Booster Neutrino Experiment
MINOS	Main Injector Neutrino Oscillation Search
MSB	Most Significant Bit
MSW	Mikheyev-Smirnov-Wolfenstein Effect
MTS	Materials Test Stand

MuID	Muon Identifier
MVA	Multi-Variate Analysis
MW	Mega-Watt
NC	Neutral Current
NND	Near Neutrino Detector
NOvA	NuMI Off-Axis ν_e Appearance
NS	North-South
NuMI	Neutrinos from Main Injector
PCB	Printed Circuit Board
PDS	Photon Detection System
PID	Particle Identification
PIP	Proton Improvement Plan
PMA	Projection Matching Algorithm
PMNS	Pontecorvo-Maki-Nakagawa-Sakata Matrix
PMT	Photomultiplier Tube
POT	Protons On Target
PrM	Purity Monitor
PTB	Penn Trigger Board
RCE	Reconfigurable Computing Elements
RMS	Root Mean Square
ROC	Remote Operation Control
RPC	Resistive Plate Chamber
SAGE	Soviet-American Gallium Experiment
SLAC	Stanford Linear Accelerator Centre
SNBD	Short Baseline Neutrino Detector
SiPM	Silicon Photomultiplier
SNO	Sudbury Neutrino Observatory
SNU	Solar Neutrino Unit
SP	Single-Phase
SSM	Standard Solar Model
SSP	SiPM Signal Processor
STT	Straw-Tube Tracker
SURF	Sanford Underground Research Facility
T2K	Tokai-to-Kamioka
TPC	Time Projection Chamber
WLS	Wavelength Shifter

Chapter 1

Introduction

The theory of elementary particles, the Standard Model of Particle Physics, is an remarkably successful theory which has stood up to every experimental test since it was first formulated in the 1970s [1, 2]. The recent discovery of the Higgs boson at CERN [3, 4] was the final missing piece and establishes the Standard Model as *the* theory of physical phenomena at the electroweak scale (up to a few hundred GeV) [5, 6].

There are however many shortcomings to the theory and further theoretical and experimental work is necessary to advance our understanding of fundamental physics [7]. For example, it ignores gravity and requires a quantised theory of gravity to reconcile it with General Relativity. The observation of ‘Dark Matter’ and ‘Dark Energy’ in Astrophysics and Cosmology cannot be explained using the known particles in the Standard Model and needs an extension of the theory. It also offers no convincing explanation of the observed domination of matter over antimatter evident in the Universe, given they were created equally in the Big Bang. Additionally, there are many unresolved theoretical problems within the Standard Model, evidence of a more fundamental underlying theory which may replace it. It is for this reason that experiments are hoping to find phenomena which may only be understood as ‘Beyond the Standard Model’.

Neutrinos offer the most promising possibilities of new physics and are currently the subject of a great amount of research [6]. The observation of neutrino oscillations [8, 9], along with the associated implication of neutrino mass, represents physics which was not included in, or predicted by, the Standard Model. In recent years the field of neutrino physics has advanced rapidly and there is currently good understanding of most experimental results. Open questions remain, such as the origin and nature of neutrino mass, the characteristics of neutrino interactions and the exact features of neutrino mixing, and will define the future of the field for many years to come. This will be discussed in more detail in Chapter 2. There is also the possibility that neutrinos may explain the aforementioned matter-antimatter

asymmetry through CP-violation in the lepton sector and may even provide a potential dark matter candidate in the possible sterile neutrino.

Future understanding and discoveries in neutrino physics requires precise measurements from highly sensitive experiments. The future Deep Underground Neutrino Experiment (DUNE) is such an experiment and will be able to contribute towards many of the unanswered questions in the field. The DUNE experiment, along with its sensitivities to unexplained phenomena, is the subject of Chapter 3. It will use large quantities of liquid argon in order to make the necessary precision measurements and will be the largest experiment using this technology ever built by an order of magnitude. In order to ensure the experiment is successful and reaches its physics potential, prototyping the technology and detector design is essential. The experiences of operating such a prototype, the 35-ton experiment, are discussed in Chapter 4, and additionally in Chapter 5.

A major challenge in the design choice of DUNE is the successful and detailed reconstruction of particle interactions necessary to make the required measurements. This is discussed in depth in Chapter 6, with emphasis placed on the difficult task of reconstructing showering particles. The performance of the reconstruction in the selection of the main signal events for DUNE, and an analysis, at this early stage, of the current status of the DUNE software at meeting its required physics goals, is presented in Chapter 8. Conclusions are contained in Chapter 9.

Chapter 2

Neutrino Physics

This chapter contains an introduction to the field of neutrino physics to provide context for the main work presented in this thesis. The history of neutrino physics is an interesting story in its own right and provides the foundation for the present and future of the field. This will be briefly retold in Section 2.1 and will motivate a discussion of neutrino oscillations in Section 2.2. An overview of the current status of the field and its future is contained in Section 2.3.

2.1 Historical Context

2.1.1 Prediction of the Neutrino

The neutrino was first postulated in 1930 by Wolfgang Pauli [10] in order to account for an inconsistency in the theory of β -decay. In the apparent two-body decay

$$A \rightarrow B + e^-, \quad (2.1)$$

kinematically the electron must be emitted with an energy given by

$$E = \left(\frac{m_A^2 - m_B^2 + m_e^2}{2m_A} \right) c^2, \quad (2.2)$$

where m_α is the mass of particle α . This energy is fixed given the masses of the particles; it was observed however that the electron energy followed a distribution, shown in Figure 2.1, with Equation 2.2 giving the maximum permitted energy. The neutrino was postulated as a third final state particle in order to account for this result and retain energy conservation laws. Pauli initially called the particle a *neutron* (preempting the name Chadwick was to give his




Fig. 2.1 Energy spectrum of the electron produced in beta decay [14].

discovered particle in 1932) but his idea was met with much scepticism. It was Fermi who named the new particle *neutrino* ('little neutral one') when incorporating Pauli's hypothesis into his theory of beta decay [11–13]. With the great success and acceptance of this theory, the field of neutrino physics was born.

Further indications of the existence of the neutrino were provided by the studies of pion and muon decay by Cecil Powell's group at Bristol in 1947 [15, 16]. Topological investigations of the newly discovered π meson and its apparent decay into a lighter meson (now known to actually be the muon lepton) appear to hint at the presence of an additional, unknown, daughter particle [15]. Furthermore, subsequent studies of the decay of the muons implied a three-body decay involving two unknown final state particles, analogous to the implication of the neutrino in β -decay by considering the electron energy distribution [17]. It seemed a model involving neutrinos could explain these observations and provided more suggestions for the existence of such a particle.

2.1.2 Discovery of the Neutrino


The elegance of Fermi's theory convinced many physicists of the existence of the neutrino but until discovered experimentally it remained a hypothetical 'bookkeeping' device. Given the elusive nature of neutrinos this was not for many years, leading to Pauli famously declaring "I have done a terrible thing. I have postulated a particle that cannot be detected". However, a series of experiments conducted between 1953 and 1956 by Clyde Cowan and Frederick Reines confirmed the hypothesis and were later rewarded with the Nobel Prize in Physics

in 1953. Using the new technology of liquid scintillator detectors [18], they designed an experiment [19] to study the (anti)neutrinos from the Hanford nuclear reactor in Washington, U.S.A, via inverse beta decay;




Their signal comprised of an initial release of scintillation light when the positron annihilates with an electron, followed a characteristic time later by a gamma ray corresponding to the neutron capture. The initial results from 1953 [20] hinted at an excess over predicted background, but the background proved to be much larger than anticipated, mainly due to an underestimation of the effects of cosmic rays. A second experiment was conducted in 1956, this time 12 m underground and 11 m from the Savannah River reactor in South Carolina. A neutrino detection rate of 2.9 ± 0.2 per hour, greater than 20 times the accidental background rate was reported, confirming the previous indications [21]. The experimental discovery of the neutrino was confirmed.

Ray Davis was also using nuclear reactors to study the interaction rates of neutrinos. Using a detector comprised of 3000 gallons of carbon tetrachloride (CCl_4) also close to the Savannah River reactor, Davis and Harmer searched for the interactions



Since it was known from Reines and Cowan that inverse beta decay



occurs, this facilitated a comparison between the neutrino and the antineutrino. They found the interaction shown in Equation 2.4 occurred at a rate less than 20 times that represented in Equation 2.5, implying for the first time a difference between neutrinos and antineutrinos [22]. This gave rise to the notion of ‘lepton number’ and its conservation in physical interactions.

It was few years before the next chapter in the history of neutrinos, the discovery of the muon neutrino in 1962 at Brookhaven [23]. It was noted the apparently permitted decay



is never observed, inciting the possibility of two distinct neutrinos. In order to test this, Lederman, Schwarz and Steinberger used a muon neutrino beam to look for two separate

interactions:

$$\bar{\nu}_\mu + p^+ \rightarrow \mu^+ + n, \quad (2.7)$$

$$\bar{\nu}_\mu + p^+ \rightarrow e^+ + n. \quad (2.8)$$

With only one type of neutrino, each interaction would be expected to occur at around the same rate. The beam was produced by accelerating protons up to 15 GeV and using a Beryllium target to create secondary mesons, decaying to produce neutrinos with energies up to 1 GeV. A total of 34 muon tracks were detected (with an estimated background from cosmic muons of 5) and no events consistent with electrons were observed. This remarkable result can only be rivalled by the technological advancements required; it was the first experiment to construct and use an artificial neutrino beam (common to all contemporary long-baseline experiments) and used 13.5 m thick steel from a dismantled battleship in order to ensure only neutrinos arrived at the spark chamber detector. This discovery was rewarded with the Nobel Prize in 1988.

A third generation of lepton, the τ , was discovered in 1975 by Martin Perl and his team at SLAC [24], completing the set of three charged leptons. They reported 64 events of the form

$$e^+ + e^- \rightarrow e^\pm + \mu^\mp + \geq 2 \text{ undetected particles}, \quad (2.9)$$

using the energy and angle distributions to predict at least two additional particles. They claimed ‘no conventional explanation’ could account for these events and proposed the existence of a heavier charged lepton as an intermediate stage:

$$e^+ + e^- \rightarrow \tau^+ + \tau^- \rightarrow e^\pm + \mu^\mp + 4\nu. \quad (2.10)$$

The τ lepton was subsequently characterised by further experiments by the Mark I detector at SLAC [25] and by the PLUTO collaboration at DESY [26]. This result heavily implied the existence of an associated neutrino to complete the symmetry observed in the first two lepton couplets.

Further evidence for a third neutrino was provided by four experiments, using the Large Electron-Positron Collider (LEP) at CERN in 1989, which were studying the production of the newly discovered Z^0 boson [27–30]. The width Γ_Z of the Z^0 resonance is dependent on the partial widths relating to final state charged leptons, hadrons and neutrinos;

$$\Gamma_Z = N_\nu \Gamma_\nu + \Gamma_{ee} + \Gamma_{\mu\mu} + \Gamma_{\tau\tau} + \Gamma_{\text{hadron}}, \quad (2.11)$$




Fig. 2.2 Measurements of the hadron production cross-section around the Z resonance. The curves indicate the predicted cross-section for two, three and four neutrino species with SM couplings and negligible mass. Taken from [31].

where N_ν is the number of light ($m_\nu \leq \frac{m_Z}{2}$) active neutrinos. Figure 2.2 shows this resonance for a range of N_ν hypotheses; fitting to the data yields a value of 2.984 ± 0.008 neutrino flavours [31].

The extremely precise measurement reported by the LEP experiments was enough for many physicists to claim indisputable evidence for the existence of the tau neutrino; it was partly for this reason that its experimental discovery was not until 25 years after the addition of the τ lepton to the Standard Model. However, in 2000 the DONUT (Direct Observation of NuTau) experiment at Fermilab, IL, U.S.A. finally reported direct detection of the tau neutrino [32]. As its name suggests, DONUT was designed specifically for the purpose of finding the third neutrino. It did this by identifying the τ as the only lepton at the interaction vertex from a ν_τ beam created by firing 800 GeV protons from the Tevatron at a tungsten beam dump. The mean energy of the ν_τ s detected at the emulsion target 36 m downstream




Fig. 2.3 Solar neutrino energy spectra as predicted by the Standard Solar Model [35]. The solid lines represent neutrinos produced during the p-p chain and dashed line represent neutrinos from the CNO cycle. Each spectrum illustrates a particular reaction during the process of a given chain.

was 111 GeV, produced by the decay of a D_S meson to a τ lepton and a $\bar{\nu}_\tau$ neutrino followed by the decay of the τ to a ν_τ . Four events were found, above a predicted background of 0.34 ± 0.05 , consistent with the Standard Model description of the tau neutrino.

2.1.3 The Solar Neutrino Problem

It has been known since the 1930s, when Hans Bethe started developing the ideas of stellar nucleosynthesis [33], that electron neutrinos are created abundantly as byproducts of the nuclear processes powering the Sun. The Standard Solar Model (SSM), established by John Bahcall in 1968 [34], explains the nuclear fusion processes responsible for powering stars. For stars the size of the Sun, this is dominated by the *proton-proton chain*; heavier stars follow the *CNO cycle*. Figure 2.3 shows the energy spectra of neutrinos released during reactions occurring during both chains.

Ray Davis, in collaboration with Bahcall, conducted the first experiment to detect these solar neutrinos in 1968. Using a similar detection technique to his previous experiments, Davis used a 380 m³ tank of tetrachloroethene (C₂Cl₄) to detect neutrinos via the inverse beta decay reaction detailed in Equation 2.5. Given the threshold for this reaction is 0.814 MeV, the

main sources of neutrinos probed by this experiment were Be^7 and B^8 . In order to eliminate backgrounds from cosmic rays, Davis constructed his experiment 4850 ft underground at the Homestake mine near Lead, SD, U.S.A. It is worth noting, in a pleasing neutrino-full-circle, this is exactly where the far detector for the DUNE experiment will be housed. The Davis Homestake experiment ran for 25 years but the results obtained [36] disagreed quite strongly with the SSM [37], consistently measuring solar electron neutrinos at a rate around a third of that predicted by the model. This became known as the ‘solar neutrino problem’, and Davis was awarded the Nobel Prize for his work on this famous experiment in 2002.

The subsequent radiochemical experiments SAGE (from 1990) and GALLEX (from 1991) were sensitive to the large flux of pp neutrinos by utilising a Ga^{71} target and the lower threshold (233 keV) reaction



These experiments also reported ‘missing’ neutrinos, disagreeing with the SSM prediction of 130 SNU [38], with capture rates of $66.6_{-7.1}^{+6.8+3.8}$ SNU (SAGE) [39] and $77.5 \pm 6.2_{-4.7}^{+4.3}$ SNU (GALLEX) [40] determined. (1 ‘Solar Neutrino Unit’ (SNU) is equivalent to a neutrino flux producing 10^{-36} captures per target atom per second.) There appeared to be a problem – either the SSM was incomplete and incorrectly over-predicted the amount of electron neutrinos or hints of new physics were beginning to appear in the experimental data.

2.1.4 The Atmospheric Neutrino Anomaly

Another abundant source of natural neutrinos is the result of cosmic rays interacting with the upper atmosphere and producing ‘atmospheric neutrinos’, typically via the interactions [41]

$$\pi^{+} \rightarrow \mu^{+} + \nu_{\mu}, \quad \mu^{-} \rightarrow e^{-} + \bar{\nu}_e + \nu_{\mu} \quad (2.13)$$

$$\pi^{-} \rightarrow \mu^{-} + \bar{\nu}_{\mu}, \quad \mu^{+} \rightarrow e^{+} + \nu_e + \bar{\nu}_{\mu}. \quad (2.14)$$

These were discovered in the early 1960s in deep underground mines in South Africa [42] and India [43] with the detection of horizontally traversing muons, only explained as the result of interactions of neutrinos from all directions since it is not possible for cosmic muons to reach such depths of rock. Since the decay lengths and kinematics are well known, the predicted ratio of muon to electron neutrinos can be calculated to a good accuracy. This ratio can be compared to an experimentally determined ratio and analysed as a measure of the efficacy of the model.

It was first noticed as early as the 1980s, by experiments designed to search for nucleon decay predicted by the then-popular Grand Unified Theories, that the measured flux did not




Fig. 2.4 Demonstrating the atmospheric neutrino anomaly using the double ratio R of muon to electron neutrino events showing data divided by expectation [46]. Various underground atmospheric neutrino detectors are shown.

correspond to that predicted by the theory. The IMB [44] and Kamioka [45] experiments, whilst measuring the atmospheric neutrino flux as an important background for nucleon decay, both noticed deficiencies in the ratio between muon and electron neutrinos compared to that predicted by the models. These experiments utilised large tanks of pure water surrounded by Photo-Multiplier Tubes (PMTs) to detect neutrinos via the Cherenkov radiation created by their charged leptonic daughter particles. Using ring-imaging techniques, it is possible to distinguish between electron-like and muon-like events and therefore identify the flavour of the incoming neutrino. The problem implied by these measurements is known as the ‘atmospheric neutrino anomaly’.

Various other experiments over the following twenty years also reported similar measurements, suggesting an excess of electron neutrinos over prediction, a deficit in the number of muon neutrinos, or both. Results from numerous experiments are shown in Figure 2.4. Most experiments report a discrepancy, with its size seemingly dependent on the energy region being studied. The experiments reporting a ratio consistent with one were much smaller than the others, and with more statistics they also started observing similar effects. This anomaly, along with the issue of the solar neutrino problem, strongly hinted at an inconsistency in our understanding of neutrino physics. This is to be discussed in detail in the Section 2.2.

2.2 Neutrino Oscillations

The concept of neutrino oscillations involves the changing of the flavour of a neutrino as it propagates through time and space; a neutrino created in a certain flavour has a non-zero

probability of being later detected in a different flavour state. It was first postulated as an explanation of the solar neutrino problem by Pontecorvo in 1968 [47, 48], having initially proposed the phenomenon in 1957 as an analogy to $K^0 \rightarrow \bar{K}^0$ transition in the quark sector [49]. It offers an elegant solution to both the solar neutrino problem and the atmospheric neutrino anomaly by explaining where the ‘missing’ neutrinos had gone; it is possible they had simply ‘oscillated’ to a different flavour and therefore would not be detected as expected.

2.2.1 The Evidence for Neutrino Oscillations

Whilst there was speculation that neutrino oscillations may be the explanation behind the issues observed in the data much sooner [50, 51], definite proof was not provided until the late 1990s. In many ways, the story of neutrino physics can be considered a triumph of the scientific method; from the initial observations of the solar neutrino problem and atmospheric neutrino anomaly through the speculation and theoretical developments to the eventual proof.

The Kamiokande and Super-Kamiokande experiments in Japan (upgrades from the Kamioka experiment noted previously) and the SNO experiment in Sudbury, Canada produced the results which showed indisputable evidence for neutrino oscillations and provided explanations for all previous discrepancies observed. This result was monumental and the work of both collaborations was rewarded in 2015 when the Nobel Prize was awarded to T. Kajita and A. McDonald, from Super-Kamiokande and SNO respectively.

2.2.1.1 Super-Kamiokande and the Atmospheric Sector

In 1994, the Kamiokande experiment produced results which hinted at an angular dependence for the R -ratio deficit, implying a dependence on neutrino travel distance [52]. This result can be explained by invoking neutrino oscillations since the probability of oscillation is influenced by the propagation distance; its significantly larger successor, Super-Kamiokande, was constructed in order to make precise measurements of this phenomenon. Super-Kamiokande is located 1000 m underground and contains an inner detector consisting of 22.5 kton fiducial volume of pure water contained within a large stainless steel cylinder (37 m high, 34 m diameter) and surrounded by 13142 20-inch photo-multipliers. With 40% coverage, the photocathodes extend over nearly an acre and provide ten times more pixels than any other experiment at the time. Its results in 1998 confirmed the earlier angular dependence findings of Kamiokande, Figure 2.5, and also considered the data as a function of neutrino energy and propagation distance, as shown in Figure 2.6. The observed effects disagreed with a view of non-oscillating atmospheric neutrinos but were entirely consistent with a two-flavour

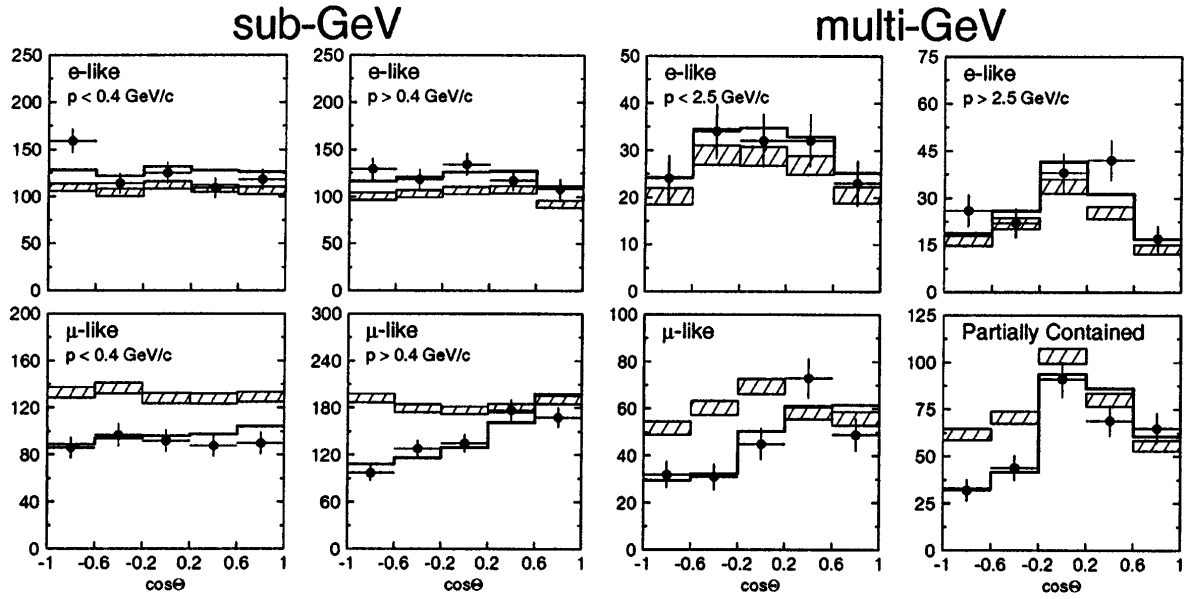


Fig. 2.5 Zenith angle distributions of μ -like and e -like events for sub-GeV and multi-GeV SuperKamiokande data sets. Upward-going particles have $\cos \Theta < 0$ and downward-going particles have $\cos \Theta > 0$. The hatched region shows the Monte Carlo expectation for no oscillations and the bold line is the best-fit expectation for $\nu_\mu \rightarrow \nu_\tau$ oscillations. Taken from [8].

oscillation model, $\nu_\mu \rightarrow \nu_\tau$. This resulted in the famous published claim for the experimental discovery of neutrino oscillations [8].

2.2.1.2 SNO and the Solar Sector

After the Super-Kamiokande results, it was clear that neutrino oscillations would probably also explain the deficit of electron neutrinos observed by solar neutrino experiments. It took a few years until the Sudbury Neutrino Observatory (SNO) in Canada provided some quite brilliant evidence of this in 2002 [9].

SNO was a water Cherenkov detector, like Super-Kamiokande, but used heavy water (D_2O) as a detector medium. The water is contained in a 12 m acrylic spherical shell and surrounded by 9456 photomultipliers at a depth of 6010 m water equivalent. The use of heavy water facilitated sensitivity to other neutrino interaction channels not accessible by Super-Kamiokande via the charge current (CC), neutral current (NC) and elastic scattering




Fig. 2.6 The ratio of the number of SuperKamiokande data events to expected events in the absence of oscillations as a function of reconstructed L/E_ν . The dashed lines show the expected shape for $\nu_\mu \rightarrow \nu_\tau$ oscillations. Taken from [8].

(ES) interactions;

$$\nu_e + d \rightarrow p + p + e^- \quad (\text{CC}) \quad (2.15)$$

$$\nu_x + d \rightarrow p + n + \nu_x \quad (\text{NC}) \quad (2.16)$$

$$\nu_x + e^- \rightarrow \nu_x + e^- \quad (\text{ES}). \quad (2.17)$$

The CC channel is sensitive exclusively to electron neutrinos, whilst the other two are accessible by neutrinos of any flavour. This allowed for the first time a simultaneous measurement of the total neutrino interaction rate as well as the electron neutrino interaction rate. The observations of SNO were the ‘smoking gun’ for neutrino oscillations; the total measured flux for all neutrinos, $\phi_{\text{NC}}^{\text{SNO}} = 6.42 \pm 1.57$ (stat.) $^{+0.55}_{-0.58}$ (sys.) $\text{cm}^{-2}\text{s}^{-1}$, agreed excellently with the electron neutrino flux predicted by the SSM, $\phi^{\text{SSM}} = 5.05^{+1.01}_{-0.81}$ $\text{cm}^{-2}\text{s}^{-1}$. However, the measured electron neutrino flux was around a third lower, $\phi_e^{\text{SNO}} = 1.76 \pm 0.05$ (stat.) ± 0.09 (sys.) $\text{cm}^{-2}\text{s}^{-1}$, consistent with previous measurements from the radiochemical experiments. The evidence seems conclusive: the solar models are correct and the neutrinos are not disappearing; they are simply changing their flavour state. A summary plot showing all solar neutrino experiments up until this point is depicted in Figure 2.7 [53].

Total Rates: Standard Model vs. Experiment
Bahcall–Serenelli 2005 [BS05(OP)]




Fig. 2.7 Comparison of the predictions of the neutrino fluxes from the Standard Solar Model with measured rates from a variety of solar neutrino experiments. The results of SNO (D_2O target, right two comparisons) show that the expected flux is observed, but not necessarily as electron neutrinos. This shows conclusively the oscillatory nature of neutrinos.

2.2.2 Vacuum Oscillations

The theory of neutrino oscillations is basically the quantum mechanics of mixed states and was developed on top of Pontecorvo's work by Ziro Maki, Masami Nakagawa and Shoichi Sakata [54]. If the neutrino flavour states can spontaneously convert from one to another, none can be considered as eigenfunctions of the Hamiltonian. The true stationary states are the *mass eigenstates* (ν_1, ν_2, ν_3), of which the flavour states (ν_e, ν_μ, ν_τ) can be considered linear superpositions:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U_{\text{PMNS}}^* \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}, \quad (2.18)$$

where U_{PMNS} is the PMNS mixing matrix which describes the flavour composition of each of the mass eigenstates, and vice versa. If the PMNS matrix were diagonal, each flavour state would correspond to a single mass state and oscillations would not occur.

Just as the flavour states are a superposition of mass states

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle, \quad (2.19)$$

the mass states can also be considered a superposition of flavour states

$$|\nu_i\rangle = \sum_\alpha U_{\alpha i} |\nu_\alpha\rangle. \quad (2.20)$$

For the three neutrino case, the PMNS matrix, decomposed into its three axial rotations, can be expressed as

$$U_{\alpha i} \equiv \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}_{\text{'Atmospheric' term}} \underbrace{\begin{pmatrix} c_{13} & 0 & e^{-i\delta} s_{13} \\ 0 & 1 & 0 \\ -e^{-i\delta} s_{13} & 0 & c_{13} \end{pmatrix}}_{\text{'Accelerator' or 'Reactor' term}} \underbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{'Solar' term}}, \quad (2.21)$$

where $c_{ij} \equiv \cos \theta_{ij}$, $s_{ij} \equiv \sin \theta_{ij}$ and δ is a CP-violating phase factor. Each axial component is often referred to by the means with which they are studied, as shown in Equation 2.21.

The weak interaction couples to the flavour eigenstates so neutrinos are always created and detected as flavour states. However, they propagate as mass states since it is these which are eigenstates of the Hamiltonian. Due to the misalignment of the flavour and mass states, oscillations can be shown to occur. A neutrino created with flavour α is a superposition of all the mass states (Equation 2.19). These states propagate as a plane wave, evolving in time

and space such that

$$|v_i(x,t)\rangle = |v_i(0)\rangle e^{-i\mathbf{x}\cdot\mathbf{p}}, \quad (2.22)$$

where \mathbf{x} and \mathbf{p} are the 4-position and 4-momentum of the neutrino respectively. From Equations 2.19 and 2.22, the evolution of the flavour states over space and time is therefore

$$\begin{aligned} |v_\alpha(x,t)\rangle &= \sum_i U_{\alpha i}^* |v_i(x,t)\rangle \\ &= \sum_i U_{\alpha i}^* e^{-i\mathbf{x}\cdot\mathbf{p}} |v_i(0)\rangle. \end{aligned} \quad (2.23)$$

In the ultra-relativistic limit, the mass of the neutrino is negligible compared to its momentum ($m_i \ll \vec{p}$) and $\vec{x} \approx ct$;

$$E_i = \sqrt{|\vec{p}|^2 + m_i^2} = \vec{p} \sqrt{1 + \frac{m_i^2}{|\vec{p}|^2}} \approx \vec{p} + \frac{m_i^2}{2\vec{p}} \quad (2.24)$$

$$\mathbf{x} \cdot \mathbf{p} = E_i t - \vec{x} \cdot \vec{p} = \vec{p} \cdot t + \frac{m_i^2}{2\vec{p}} t - \vec{x} \cdot \vec{p} \approx \frac{m_i^2}{2\vec{p}} \vec{x} = \frac{m_i^2}{2p} x, \quad (2.25)$$

assuming the neutrino displacement is in the direction of its momentum and using natural units ($c \equiv \hbar \equiv 1$). Thus, using Equations 2.23, 2.25 and 2.20,

$$\begin{aligned} |v_\alpha(x,t)\rangle &= \sum_i U_{\alpha i}^* e^{-i\frac{m_i^2}{2p}x} |v_i(0)\rangle \\ &= \sum_i \sum_\beta U_{\alpha i}^* e^{-i\frac{m_i^2}{2p}x} U_{\beta i} |v_\beta\rangle. \end{aligned} \quad (2.26)$$

The probability of a neutrino created in flavour state α being observed in flavour β can be determined from Equation 2.26

$$P(\alpha \rightarrow \beta) = |\langle v_\alpha | v_\beta(x,t) \rangle|^2 \quad (2.27)$$

$$= \left[\sum_i U_{\alpha i} e^{i\frac{m_i^2}{2p}x} U_{\beta i}^* \right] \left[\sum_j U_{\alpha j}^* e^{-i\frac{m_j^2}{2p}x} U_{\beta j} \right] \quad (2.28)$$

$$= \sum_{i,j} U_{\alpha i} U_{\alpha j}^* U_{\beta j} U_{\beta i}^* e^{i\frac{m_i^2 - m_j^2}{2p}x}, \quad (2.29)$$

and is observed to be dependent on the neutrino momentum, the difference between the squared masses of the flavour states, the propagation distance and the relative mixing of the two flavour states encoded in the matrix elements U .

An accelerator based neutrino experiment, such as DUNE, will typically use a ν_μ -dominated beam and look for *muon neutrino disappearance* ($P(\nu_\mu \rightarrow \nu_\mu)$) and *electron neutrino appearance* ($P(\nu_\mu \rightarrow \nu_e)$). In this case, also in the relativistic limit, the relevant appearance and disappearance probabilities can be approximated, respectively, as

$$P(\nu_\mu \rightarrow \nu_e) \approx \sin^2 2\theta_{13} \sin^2 \theta_{23} \sin^2 \left(1.27 \frac{\Delta m_{13}^2 L}{E} \right) \quad (2.30)$$

$$P(\nu_\mu \rightarrow \nu_\mu) \approx 1 - \cos^4 \theta_{13} \sin^2 2\theta_{23} \sin^2 \left(1.27 \frac{\Delta m_{23}^2 L}{E} \right), \quad (2.31)$$

where $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$ is the *mass squared splitting* in eV, L is the distance propagated in km and E is the neutrino energy in GeV.

From these equations, it can be seen the important controllable parameters relevant for observing oscillations are the neutrino energy and the distance they travel. An experiment will typically choose a ratio L/E which will attempt to maximise the effect of oscillations in order to make precision measurements.

2.2.3 Matter Effects

The oscillations considered thus far are *vacuum oscillations* which occur due to the mixing of the neutrino mass and flavour states. Whilst directly confirming the oscillation of solar neutrinos, the SNO experiment (along with every other solar neutrino experiment) reported more oscillations than can be explained using just the vacuum oscillation model discussed in Section 2.2.2 [55, 56]. When neutrinos propagate through matter, an additional potential can be shown to also produce an asymmetry between the different neutrino flavours. Since solar neutrinos travel through dense matter before exiting the Sun, it is possible these matter effects could explain this discrepancy.

Coherent scattering (scattering in which the neutrino state is unchanged) due to interactions with the medium cause neutrinos travelling through matter to feel a potential. As normal matter is composed of electrons, rather than their heavier counterparts muons and taus, electron neutrinos are affected more by this potential. The mechanism for this is demonstrated in Figure 2.8. This gives rise to an additional effective mass splitting between the electron neutrino and the other flavours and therefore results in the possibility of oscillations [57]. Due to the density of the Sun and the neutrino energies, the neutrinos actually feel a resonance which causes their oscillation probability to become dramatically higher than the vacuum oscillation probabilities. This is known as the Mikheev-Smirnov-Wolfenstein (MSW) effect [58, 59]. It should be noted that, since normal matter is composed of electrons




Fig. 2.8 General scattering mechanics which occur as neutrinos pass through matter. Neutral current scattering (Figure 2.8a) occurs for all neutrino flavour combinations whereas charged current scattering (Figure 2.8b) only occurs when the incoming leptons have the same flavour.

and not positrons, this effect is also different for ν_e s and $\bar{\nu}_e$ s; the importance of this becomes apparent when considering the additional effects of CP violation.

2.2.3.1 KamLAND and Reactor Neutrinos

In order to investigate the possible MSW effects in the Sun, measurements of electron neutrino disappearance from terrestrial neutrinos, which were not subject to these matter effects, were considered. The first experiment to publish results was KamLAND (Kamioka Liquid Scintillator Anti-Neutrino Detector) in 2003 [60, 61]. KamLAND occupied the site previously used by the Kamiokande experiment in Japan under 2700 m.w.e of rock and utilised 1 kton of ultra-pure liquid scintillator contained in a 13 m spherical balloon. It was surrounded by 53 Japanese nuclear power reactors with baselines ranging from 80 km to 800 km and detected the $\bar{\nu}_e$ s via the inverse beta decay reaction $\bar{\nu}_e + p \rightarrow e^+ + n$. Scintillation light from the delayed coincidence of a positron with the neutron capture was detected using 1879 PMTs and constituted a signal with very low background. The results from KamLAND confirmed the apparently large matter effect in solar neutrinos and completely solved for the first time the long-standing solar neutrino problem [62–64].

2.2.4 CP Violation

The δ terms in Equation 2.21 are CP-violating phase factors. They could be included in any of the diagonalised components but are generally added to the accelerator part since this is how current and future experiments will look for evidence of CP violation. As long as all the mixing angles are non-zero, there is the possibility of CP violation in the lepton sector.

This is an exciting prospect and one of the reasons for the current intense interest in neutrino physics. It is known CP-violating processes must have occurred in the early Universe

since matter has come to dominate massively over antimatter after they were created equally in the Big Bang. This has been observed in the quark sector but current experimental evidence can only account for a small amount of the necessary CP violation. It is also expected but has never been observed in strong interactions [65]. Leptonic CP violation could potentially account for the matter-antimatter asymmetry in the Universe and ultimately explain how it evolved to include our very existence [66, 67].

In neutrino experiments, effects of CP violation would be apparent as a difference in behaviour between neutrinos and antineutrinos. For example, since the sign of δ is different for neutrinos and antineutrinos, an asymmetry

$$\mathcal{A} = \frac{P(\nu_\mu \rightarrow \nu_e) - P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)}{P(\nu_\mu \rightarrow \nu_e) + P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)} \quad (2.32)$$

can be observed and measured.

2.3 Status of Neutrino Physics

The field of neutrino physics has advanced rapidly over the past twenty to thirty years (discussed in Sections 2.1 and 2.2) and there is currently a good understanding of most experimental results in the context of 3-flavour neutrino oscillations. Presently, the focus has shifted to making precise measurements of the oscillation parameters and trying to understand the nature of neutrino mass. The current understanding of each of these areas will be presented in Sections 2.3.2 and 2.3.3 respectively following a brief overview of current and future experiments in Section 2.3.1.

2.3.1 Current and Future Experiments

In recent years, neutrino experiments which utilise a custom-built artificial neutrino beam have been offering complimentary and world-leading results. These ‘accelerator experiments’ are used in order to have more control over the energy spectrum and composition of the neutrino beam and often use a long baseline, sampling the beam at different points to determine the effects of oscillation as the neutrinos propagate in between.

MINOS (Main Injector Neutrino Oscillation Search) was based at Fermilab, U.S., and detected neutrinos from the NuMI (Neutrinos at the Main Injector) beam at a ‘near detector’ and then again in Northern Minnesota, a baseline of 735 km. T2K (Tokai-to-Kamioka) follows a similar design and uses the Super-Kamiokande detector as the far detector, utilising a beam from J-PARC, Japan and a baseline of 295 km. T2K and NOvA (NuMI Off-Axis

ν_e Appearance), another current long-baseline experiment, were designed specifically to measure the mixing parameter θ_{13} by looking for ν_e appearance in a ν_μ beam. NOvA, like its predecessor MINOS, uses the NuMI beam and has a far detector at the same site. However, along with T2K, it is ‘off-axis’ by around 2° ; this produces a more monochromatic neutrino energy spectrum to maximise the effect of oscillations and make more accurate measurements in a specific region. T2K and NOvA still have many years left of their respective programmes and are currently making precision measurements of the mixing parameters along with constraining CP violation by combining neutrino and antineutrino analyses. Especially through joint analyses, these experiments will be able to offer increasingly statistically significant measurements before the next generation of experiments.

Future long-baseline experiments include DUNE [68], which will be discussed in depth in Chapter 3, and Hyper-Kamiokande [69], an upgrade of the current T2K experiment. Hyper-Kamiokande will also use water Cherenkov technology but will boast a far detector with a fiducial volume 25 times larger than that of Super-Kamiokande. The timescale of these projects is on the order of at least ten years from now and both pose incredible engineering challenges in their own right.

2.3.2 Oscillation Parameters

Through combinations of the choice of baseline and the neutrino energy spectrum (both peak energy and distribution), an oscillation experiment may be configured to make measurements of the desired mixing parameters. As is evident from Equations 2.30 and 2.31, the probability of oscillation varies sinusoidally and is dependent on these experimental design considerations. For example, both T2K and NOvA have an L/E chosen to coincide with the maximal probability of $\nu_\mu \rightarrow \nu_e$ oscillation, resulting in the possibility of precision measurements of the parameters governing ν_μ -disappearance and ν_e -appearance. By additionally searching for $\bar{\nu}_\mu$ -disappearance and $\bar{\nu}_e$ -appearance in a $\bar{\nu}_\mu$ dominated beam, constraints on CP-violation may be possible via the δ_{CP} parameter in the 3-neutrino mixing framework. Experiments studying natural sources of neutrinos, such as solar and atmospheric neutrinos, have less control over these experimental parameters but may still extract information about oscillations through shape fits over a range of energies or from different baselines (possible due to considerations in the neutrino direction and therefore trajectory).

The current status of the mixing angles and the mass-squared differences is depicted in Figure 2.9. The world-best measurements for θ_{12} and Δm_{12}^2 are provided by the solar neutrino experiments (Homestake [36], GALLEX [70], SAGE [71] and SNO [72]) and KamLAND [73]. The leading measurements in the atmospheric neutrino sector, θ_{23} and $|\Delta m_{32}^2|$, are from

Table 2.1 The current best-fit values for the neutrino oscillation parameters for normal (inverted) hierarchy. Taken from [81].

Parameter	Best fit ($\pm 1\sigma$)
$\sin^2 \theta_{12}$	0.306 ± 0.012
$\sin^2 \theta_{23}$	$0.441^{+0.027}_{-0.021}$ ($0.587^{+0.020}_{-0.024}$)
$\sin^2 \theta_{13}$	0.02166 ± 0.00075 (0.02179 ± 0.00076)
Δm_{12}^2 [10^{-5} eV 2]	$+7.50^{+0.19}_{-0.17}$
$ \Delta m_{3\nu}^2 $ [10^{-3} eV 2]	$2.524^{+0.039}_{-0.040}$ ($-2.514^{+0.038}_{-0.041}$)
δ_{CP} [$^\circ$]	261^{+51}_{-59} (277^{+40}_{-46})

Super-Kamiokande [74], IceCube [75] and the accelerator experiments MINOS [76, 77], T2K [78] and NOvA [79].

The value of θ_{13} was known, from limits determined from global fits to world data, to be much smaller than the others and was even consistent with zero. In addition to the accelerator experiments, reactor neutrino experiments are sensitive to θ_{13} via $\bar{\nu}_e$ disappearance and it was these experiments which produced the decisive results first. Daya Bay [82] in China and RENO [83] in South Korea found evidence of a non-zero value in 2012. There is good agreement between these reactor experiments and more recent measurements from T2K [84] and NOvA [85].

A summary of the best known values for all these oscillation parameters is shown in Table 2.1. The CP-violating phase δ_{CP} is the least known at present and provides a priority for current and future neutrino experiments. T2K have excluded the CP conservation regions with 90% confidence and currently favours a maximal CP violation value of $\delta_{\text{CP}} = -\pi/2$ [86], with recent results from NOvA broadly consistent [87]; this holds much promise for future experiments. The octant of θ_{23} , the location of the parameter in either the $> 45^\circ$ or $< 45^\circ$ octant, is also undetermined and requires high precision measurements; it is possible that the mixing in this sector is ‘maximal’ ($\theta_{23} = 45^\circ$).

2.3.3 Neutrino Mass

Neutrinos in the Standard Model are assumed massless as they are only known to exist in a left-handed chiral state and therefore cannot be used to construct the usual Dirac mass term. However, the observation of neutrino oscillations implies the existence of neutrino mass (the oscillation probabilities, such as Equations 2.30 and 2.31, would be zero if there was no mass splitting). Three active neutrino flavours gives rise to two independent mass splittings, Δm_{12}^2 and Δm_{32}^2 , as apparent in the oscillation probabilities. Unfortunately, fitting to the oscillation




Fig. 2.9 Global 3-neutrino oscillation analysis taken from [80, 81]. Each panel shows the two-dimensional projection of the allowed six-dimensional region after marginalisation with respect to the undisplayed parameters. The different contours correspond to 1 σ , 90%, 2 σ , 99%, 3 σ CL (2 dof).

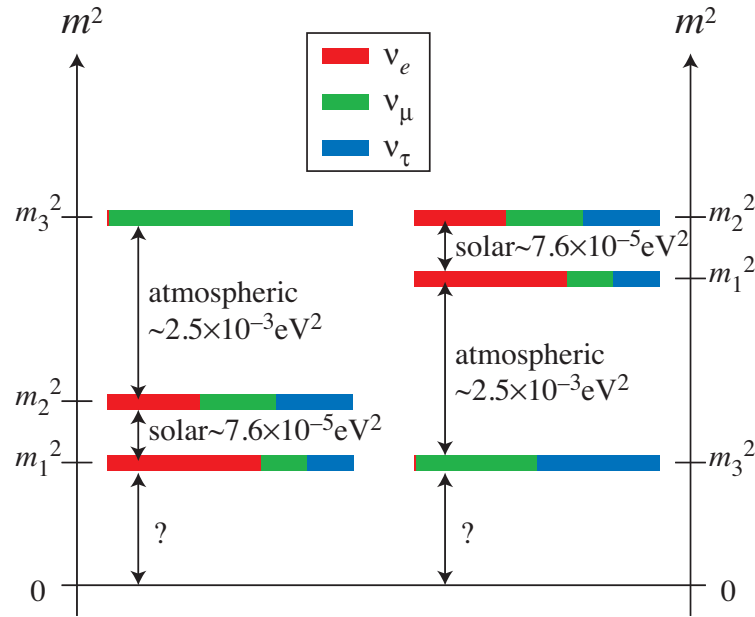


Fig. 2.10 Demonstration of the current uncertainties in the neutrino mass. The undetermined sign in the mass splitting between the 2 and 3 states leaves two possible ‘hierarchies’ open: normal (left figure) and inverted (right figure). The absolute scale of the masses is also currently unknown. The flavour composition of each of the mass states, given by the mixing angles, is denoted by the coloured bars.

results provides access to only the value of these splittings and not the signs, resulting in an ambiguity in the ordering of the mass states. This can be resolved in the solar sector by utilising the effect of the MSW resonance encountered by neutrinos in the Sun, allowing the sign of Δm_{12}^2 to be known (it must be positive as otherwise fewer oscillations, not more, will have been observed by SNO and the other solar neutrino experiments). This leaves two possible ‘hierarchies’, normal and inverted, which are possible given the experimental data. These mass splittings also do not offer any indication of an absolute mass scale for the neutrino mass states; this must be constrained using other methods. These uncertainties are illustrated in Figure 2.10. The best current constraints from NOvA disfavour the hypothesis of inverted mass hierarchy with θ_{23} in the lower octant with greater than 93% confidence for all values of δ_{CP} [87], with T2K in agreement [88].

DUNE will use the MSW effect present as neutrinos propagate through the Earth’s crust in order to resolve the hierarchy problem. It is essential that the hierarchy is resolved since the associated asymmetries between neutrinos and antineutrinos can mimic true CP violation, which therefore cannot be measured accurately until the mass splittings are completely understood. Due to the large matter effects associated with its long baseline, the NOvA

experiment is sensitive to the mass hierarchy and may be able to produce results before DUNE and Hyper-Kamiokande.

The absolute neutrino mass cannot be measured from oscillation experiments so other techniques have been developed. It is possible to use information from β -decay to measure the mass scale; the $\bar{\nu}_e$ mass alters the spectrum of electrons near the end point so precision measurements can study this effect. The current best limits on the mass are from H^3 β -decay experiments and yield $m_{\bar{\nu}_e} < 2.05$ eV at 95% C.L. [89, 90]. Cosmological analysis can also constrain the absolute neutrino mass by looking at the distribution of matter in the Universe and information such as galaxy clustering. The Planck collaboration reported the upper limit on the sum of all neutrino masses as $\sum_i m_{\nu_i} < 0.23$ eV at 95% C.L. in 2013 [91], indicating a significantly lower mass scale than is attainable using current experiments.

Chapter 3

The Deep Underground Neutrino Experiment

The Deep Underground Neutrino Experiment (DUNE) experiment [68, 92–94] is a future long-baseline neutrino experiment with a diverse physics program hosted by Fermilab, IL, U.S.. The far detector will be at the Sanford Underground Research Facility (SURF) near Lead, South Dakota, providing a baseline of 1300 km. A cartoon of the experiment is shown in Figure 3.1.

The DUNE experiment will be discussed in this present chapter. As the experiment utilises liquid argon TPCs, a brief history and description of this detector technology is provided as a basis in Section 3.1. An overview of the experiment, including its motivation, will be presented in Section 3.2 before the experimental details are discussed in Section 3.3. The sensitivities of the experiment and its potential discoveries are the subject of Section 3.4. Finally, the schedule and strategy implemented by the collaboration to ensure commencement of data taking in around ten years' time is outlined in Section 3.5.

3.1 The LAr TPC Concept

The use of a liquid argon (LAr) Time Projection Chamber (TPC), or LArTPC, as a high-precision fine-grained detector medium holds much promise for the successful resolution of the open questions in neutrino physics. A great amount of R&D work has taken place to advance the maturity of the technology and pioneering experiments, such as ICARUS [96], have further increased the understanding of the neutrino community of the detector techniques. Past and currently running experiments at Fermilab, such as ArgoNeuT [97],




Fig. 3.1 Cartoon showing the configuration of the Deep Underground Neutrino Experiment [95]. The experiment will be based at Fermilab, shown at the right of the figure, and will send neutrinos towards SURF, at the left hand side. The distance travelled, through the Earth's crust, will be 1300 km.

LArIAT [98] and MicroBooNE [99], are successfully using LArTPCs to take data for analysis and it seems certain to be the future of neutrino physics in the U.S. [100].

This section will provide a brief history of LArTPC technology and motivate its potential when used in a large experiment such as DUNE. The basic operation of such a detector will also be described to provide background for discussion of the DUNE and 35-ton experiments, and of reconstruction in LArTPCs, in future chapters.

3.1.1 A Brief History of Time (Projection Chambers)

The use of a time projection chamber as a potential particle detector was put forward by David Nygren in 1974 [101]. He envisioned bubble-chamber quality data but with the possibility of digital readout of the data, facilitating extremely fine spatial resolution, good timing resolution and fast recovery after triggering. The basic concept is a drift chamber containing a noble gas placed within a field to drift ionisation electrons created by a propagating particle towards a multielectron array. This setup allows full three-dimensional reconstruction by combining information from the two-dimensional readout plane with the drift time. Nygren also included a magnetic field to assist particle identification in his design, shown in Figure 3.2.

The extension of this concept to a liquid argon TPC and its potential as a high-precision fine-grained detector medium in neutrino physics was proposed by Carlo Rubbia in 1977 [102]. The use of a noble liquid rather than gas is necessary in neutrino experiments to provide a high enough target mass for increased probability of neutrino interactions. Noble liquids have high electron mobility and low diffusion, favourable properties as the detection of particles is from the ionisation and scintillation light created by the particles. Given the necessity of a high electric field in order to drift these electrons to the readout planes,




Fig. 3.2 The original concept of the time projection chamber particle detector, drawn by David Nygren in 1974 [101]. The sections are labelled as followed: methane-filled region (A), screen to establish electron field (B), end-cap detectors (C), superconducting solenoid (3.33 T) (D), iron return yoke for magnetic field (E), beam vacuum pipe (F).

excellent dielectric properties are also required; noble liquids possess such qualities. The properties of liquid argon which make it almost perfect for this use are demonstrated in Table 3.1.

An additional advantage of this technology is the low threshold for detection; this is set by the ionisation threshold of liquid argon and is only 23.6 ± 0.5 eV [104]. Rubbia realised that a LArTPC could be the digital replacement for the high quality particle detection methods used in bubble chambers, very common in neutrino physics in the 1970s. He proposed the first LArTPC detector design, shown in Figure 3.3, which bears a striking resemblance to the LArTPCs in use today.

Table 3.1 Properties of noble liquids relevant when considering a TPC medium for a neutrino experiment [103].

	Water	He	Ne	Ar	Kr	Xe
Boiling point [K] @ 1 atm	373	4.2	27.1	87.3	120.0	165.0
Density [g/cm^3]	1	0.125	1.2	1.4	2.4	3.0
Radiation length [cm]	36.1	755.2	24.0	14.0	4.9	2.8
Scintillation [γ/MeV]	-	19 000	30 000	40 000	25 000	42 000
dE/dx [MeV/cm]	1.9	0.24	1.4	2.1	3.0	3.8
Scintillation λ [nm]	-	80	78	128	150	175
Abundance (Earth atm) [ppm]	5×10^4	5.2	18.2	9340.0	1.10	0.09
Electron mobility [cm^2/Vs]	low	low	low	400	1200	2200




Fig. 3.3 The LArTPC detector proposed by Carlo Rubbia in 1977 [102].

Constructing and operating such a detector was beyond the technology of the time, and is still being understood today. The operation of a LArTPC detector is described in Section 3.1.2 and the associated challenges are the subject of Section 3.1.3.

3.1.2 LAr TPC Operation


A LArTPC typically consists of one or more anodes and cathodes at either end of an active drift region. An ionising particle passing through a LArTPC causes electrons to become free from argon atoms and, in the presence of a field, drift towards an anode where they are read out.

There are two differing designs of LArTPC: single-phase and dual-phase. In a single-phase design, only argon in a liquid state is utilised and the entire detector, including target medium, propagation, readout and signal processing, resides within the liquid. In contrast, a dual-phase design would employ a layer of gaseous argon above the liquid, through which the drift electrons are extracted before readout. This has the primary advantage of amplification of signal due to electron avalanches in an electric field in the gas and therefore a lower detection threshold but also results in a larger fiducial volume with the absence of dead regions in the LAr volume.

The readout consists of multiple wire planes with different orientations to facilitate the reconstruction. The wires are either ‘induction’ wires, which allow the electrons to deposit charge but continue past, or ‘collection’ wires, on which the electric field lines end and all

the charge on the electron is collected. Each wire plane is therefore held at a different ‘bias voltage’ to prevent any field lines ending on the induction wire, thus creating local electric fields which promote the continuing forward motion of the electrons. The signal seen is therefore dependent on the type of wire plane; a bipolar pulse on an induction plane wire and unipolar on a collection plane wire. It is also common, though not essential, to make use of a ‘grid plane’ upstream of the signal planes in order to shield them from the electron charge until the drift electrons are close. Without such a plane, the bipolar pulse would be highly asymmetric, though would still have zero integral. It also makes changing the drift voltage (controlling the electric field) slightly easier as the signal planes are somewhat shielded from its effects. MicroBooNE does not operate with a grid plane and, although the 35-ton and the DUNE reference design make use of a grid plane, it is uncertain whether the benefit outweighs the cost for a large LArTPC detector such as the DUNE far detector. There are alternative readout possibilities to this typical design which have been suggested but, given the scale of future LArTPCs, it is highly unlikely a viable solution which delivers superior readout at a comparable cost will be found.

Upon ionisation, an electron has a certain probability (around 60% at a field of 500 V/cm) of recombining before the field can separate it from its ion. Whilst this recombination compromises the signal observed, it is accompanied by a flash of scintillation light which may be detected and used to assign an ‘event time’ to the interaction, known as T_0 . Without this information, it would be impossible to place an absolute time scale on the event and result in an unresolved coordinate along the drift direction. Scintillation light in LAr is produced by two mechanisms:



where in Equation 3.1 the argon is excited to a higher energy state and in Equation 3.2 the argon is ionised. In the second stage, excited and ionised molecules are formed in the liquid before de-excitation leads to the production of scintillation light. The magnitude of the applied electric field must be chosen to balance between ionisation and scintillation; a larger field would result in less recombination and therefore compromise the scintillation light while a smaller field would have consequences on the signal received at the wire planes. Figure 3.4 demonstrates this and justifies the field value of 500 V/cm which is often chosen in current LAr neutrino experiments.

The basic operational principles of a LArTPC are demonstrated in Figure 3.5. The details of how the ionisation charge and scintillation light are collected and processed are




Fig. 3.4 Demonstration of the competing effect the electric field has on the luminosity of the ionisation electrons and scintillation light arriving at the detector readout [105–107]. The solid lines represent the charge recombination whilst the dashed lines demonstrate light recombination. The numeric labels are the specific energy loss (dE/dx) in units of mip. Since information from both is essential in reconstructing the complete interactions, a balance must be found. Image made by Craig Thorn, LBNE Document Database 4482.




Fig. 3.5 Schematic demonstrating the basic operation principles of a LArTPC [108]. The cathode is shown on the left-hand side and the anode readout wires, with three complementary planes (two induction, U and V, and one collection, Y), are demonstrated on the right-hand side. A through-going ionising particle will liberate electrons from the LAr which drift in the electric field; charge is induced on the first two wire planes and collected on the final one. Due to the differing orientations of the wires between planes, three views of the interaction are provided (two are demonstrated in the figure) and combining the information from multiple views with the drift time information facilitates a complete understanding of the original particle tracks.

experiment-specific and will be discussed in the context of DUNE in this chapter and the 35-ton experiment in Chapter 4. This information is all that is required to fully understand and analyse the interactions occurring in the detector; methods used to reconstruct particles and interactions in LAr will be the subject of Chapter 6.

3.1.3 LArTPC Challenges

There is no doubt of the promise of LArTPCs for the future of neutrino physics but with such expectation comes many challenges. This will be elaborated upon in more detail when discussing the 35-ton run in Section 4.3 but will be briefly mentioned here for completeness.

Given the drift fields required, and the necessary distances, the associated high voltage on the cathode must be on the order of ~ 100 kV. This presents engineering challenges related to the feedthrough and cryostat design but can also lead to dielectric breakdown of the liquid

nearby such large voltages. The properties of LAr and the design implications must be very well understood to ensure this does not endanger the quality of the detector medium.

The presence of electro-negative impurities in the argon can capture drift electrons as they travel towards the anode planes and hinder the observed signal. The probability of this recombination can be parametrised using an ‘electron lifetime’ τ , defined in

$$Q(t) = Q_0 e^{-t/\tau}, \quad (3.3)$$

where Q is the collected charge after drift time t and Q_0 is the initial charge deposited. The electron lifetime is directly affected by the maintained purity of the argon; DUNE expects a contamination no greater than 100 ppt O₂-equivalent and 20 ppm N₂ (to also ensure sufficient light yield) [94]. This necessitates a purification system to remove impurities and requires the constant recirculation of the liquid through it. A liquifier is also necessary to recondense any boiled-off gases at the surface.

Along with the possibility of lost signal through finite electron lifetimes, the electrons may also undergo interactions and drift off course either transversely or longitudinally. This ‘diffusion’ affects the location and size of the observed signal so must also be well understood.

With so much resting on the success of the DUNE experiment, and considering all these effects which must be understood, prototyping is essential. The 35-ton prototype was constructed as an attempt to better understand LArTPCs and is the subject of Chapter 4.

3.2 Overview of DUNE

The outstanding questions in neutrino physics discussed in Section 2.3, namely the resolution of the mass hierarchy, the determination of the CP-violating phase δ_{CP} , the measurement of the octant of θ_{23} and precision calculations of all the mixing angles, motivate the need for next generation experiments. The DUNE experiment will make decisive contributions to each of these areas; it will also search for nucleon decay with the ability to set world-leading proton lifetime limits and make detailed, unique measurements of the ν_e flux from a core-collapse supernova within our galaxy should one occur during the experiment. Along with this, DUNE will be used to look for Beyond Standard Model physics (such as non-standard interactions and sterile neutrinos), signatures of dark matter and, utilising the capable near detector, measurements of a range of neutrino cross-sections and nuclear effects including final state interactions.

The chosen technology for the DUNE far detector, in order to maximise sensitivity to all these factors, is a LArTPC, introduced and described in Section 3.1. The detector will contain

four modules, each comprised of 10 kt fiducial LAr and separate data acquisition and readout systems. The beam will be provided by Fermilab as part of its PIP-II program [109] and will be wide band, enabling the study of a range of neutrino energies. This facilitates a study of multiple oscillation peaks, essentially due to differing L/E ratios, and is relevant when considering the effects of an unknown CP-violating phase and unresolved mass hierarchy. Since the impact of both of these uncertainties is apparent as an asymmetry between neutrinos and antineutrinos (Equation 2.32), there is an implicit degeneracy which must be resolved to ensure both phenomena are correctly determined. Having access to multiple oscillation peaks means this may be dealt with in a single experiment, as demonstrated in Figure 3.6. In the plot, θ_{13} is held constant and the rates are determined by the number of neutrino and antineutrino events respectively. Assuming a baseline of 1300 km, as for DUNE, the first oscillation maximum is at $E_\nu = 2.4$ GeV and the second is at $E_\nu = 0.84$ GeV. The banana-shaped distributions are obtained as the value of δ_{CP} is varied from $-\pi$ to π . There is good separation between the distributions associated with each hierarchy at the first maximum whereas at the second maximum this is degenerate and the rates are similar for a given value of δ_{CP} regardless of the hierarchy. It can be seen how complimentary measurements at each maxima can be used to make unambiguous measurements of both the mass hierarchy and of CP violation with the same experiment [110]. It should be noted the rates are relatively low at the second oscillation maximum and designing an experiment with access to these events must take into consideration the ratio of the potential benefit to experimental factors such as being able to accurately detect, reconstruct and study these interactions; DUNE has decided this design choice has significant physics justifications and so has in its design a wide-band beam.

DUNE was officially formed in early 2015 following the dissolution and merging of two leading next generation long-baseline experiments: the Long Baseline Neutrino Experiment (LBNE) in the U.S. [111–113] and the Large Apparatus for Grand Unification, Neutrino Astrophysics, and Long Baseline Neutrino Oscillations (LAGUNA-LBNO) in Europe [114]. Given the scale of these projects, it was decided in 2014 that efforts should be focussed on one flagship experiment utilising as many experts in neutrino physics and LArTPC technology as possible [115]. The benchmark DUNE design is very similar to that of the former LBNE experiment, which also made use of an upgraded Fermilab neutrino beam and a large LArTPC at SURF, and gained the understanding of dual-phase LArTPC detectors developed by LAGUNA-LBNO. It is likely that at least one of the four DUNE detector modules will be a dual-phase LArTPC.

The experiment will be facilitated by the Long Baseline Neutrino Facility (LBNF), which will oversee the technical side of the project and ensure the DUNE experiment can function




Fig. 3.6 Demonstration of how having access to multiple oscillation maxima facilitates measurements of both the neutrino mass hierarchy and leptonic CP violation using the same experiment. Taken from [110].

as desired. The relationship between the LBNF and the DUNE projects is based on the model used at CERN to manage the Large Hadron Collider (LHC) and each of the experiments which use it. LBNF has its own management structure and operates separately from DUNE, though the two projects work closely together. It is supported mainly via the Department of Energy in the U.S. whereas DUNE is internationally funded. The DUNE collaboration is responsible for defining the scientific goals of the experiment and the corresponding technical requirements. Using these, LBNF will design and construct all technical facilities, such as the beam upgrade, the facilities for the near detectors at Fermilab and the excavation and outfitting of the large caverns for the far detectors underground at SURF along the required infrastructure to support the construction of the cryostats and the associated cryogenic systems. DUNE will provide the four massive LArTPCs and the near detector systems, to be constructed at the sites supplied by LBNF. These will be discussed further in Section 3.3. During the lifetime of the experiment, LBNF is responsible for the maintaining and operation of all the facilities whilst DUNE will commission and operate the detectors. The scientific research program conducted with the collected data is the duty of the DUNE collaboration and will be explored in Section 3.4.

Given the scale of the projects, work is already underway. Construction at the far detector began with groundbreaking on 21st July 2017, with installation of the first detector module due to commence in 2021. The start of the DUNE experiment will then correspond to the completion of this module, scheduled in 2024. The PIP-II upgraded 1.2 MW beam will be ready in 2025 and will signify the commencement of beam data taking. Subsequent detector modules will be added as soon as is feasible thereafter, increasing the fiducial volume up to the target mass of 40 kt. Further beam upgrades, up to 2.4 MW (PIP-III) are envisaged beyond this to bring the experiment up to full power and maximise the physics capability of the project. The timescales of both LBNF and DUNE, along with all the essential research which must be conducted as the plans progress, is the subject of Section 3.5.

3.3 Experimental Details

The design of the DUNE experiment is driven by the physics ambition of the project and all specific experimental details are motivated by the science DUNE wishes to study. This section will briefly report on the current understanding of how DUNE will be planned. It should be noted that whilst the important defining features of the design will endure it is likely, given the timescales of the project, particular details will change for the final proposal.

The plans for the beam will be overviewed in Section 3.3.1 before the present considerations for the far and near detectors are presented in Sections 3.3.2 and 3.3.3.

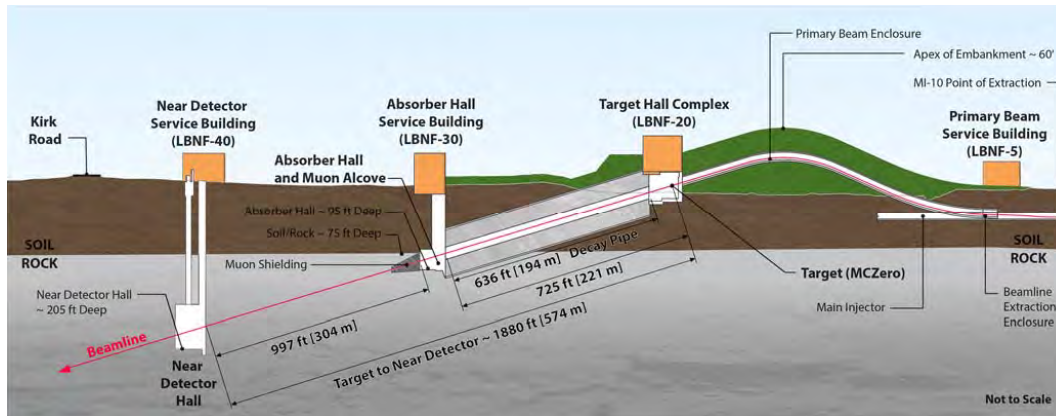


Fig. 3.7 Longitudinal section of the LBNF beamline facility at Fermilab [93].

3.3.1 Beam

The beam is required to be wide band to facilitate access to the first two oscillation maxima (at 1300 km, 2.4 GeV and 0.8 GeV). It is a conventional, horn-focussed, sign selected neutrino beam and is detailed in Figure 3.7.

A proton beam (60-120 GeV) is extracted from the Fermilab Main Injector, transported through a man-made embankment and bent downwards to establish the final trajectory towards the far detector. The beam is incident on a target to produce secondary mesons, which are subsequently focussed into a decay pipe by magnetic horns where they decay into muons and neutrinos. The polarity of the horns determines whether or not the beam is neutrino or antineutrino dominated. At the end of the decay pipe, the muons and any residual hadrons are stopped to produce a beam of neutrinos, tuned with energies between 0.5 GeV and 5 GeV. The fluxes of different neutrino species present in the beam in both neutrino- and antineutrino-running modes are shown in Figure 3.8. The target is based on the design used in the Fermilab NuMI (Neutrinos from Main Injector) beam, in operation since 2005, and consists of a graphite core with dual titanium water lines to dissipate heat, encased in a titanium containment tube.

The current beam design has been evolving since 2012 when it was first conceived for the LBNE experiment. Many potential optimisations have been identified and some are now part of the reference design; others are still in the process of evaluation. In the sensitivities presented in Section 3.4, two beam configurations are referred to: reference and optimised. The optimised design pertains to such ongoing considerations envisioned to improve performance, for example in flux at the first, and especially the second, oscillation maximum and in the reduction of wrong-sign neutrino backgrounds. The improvements are

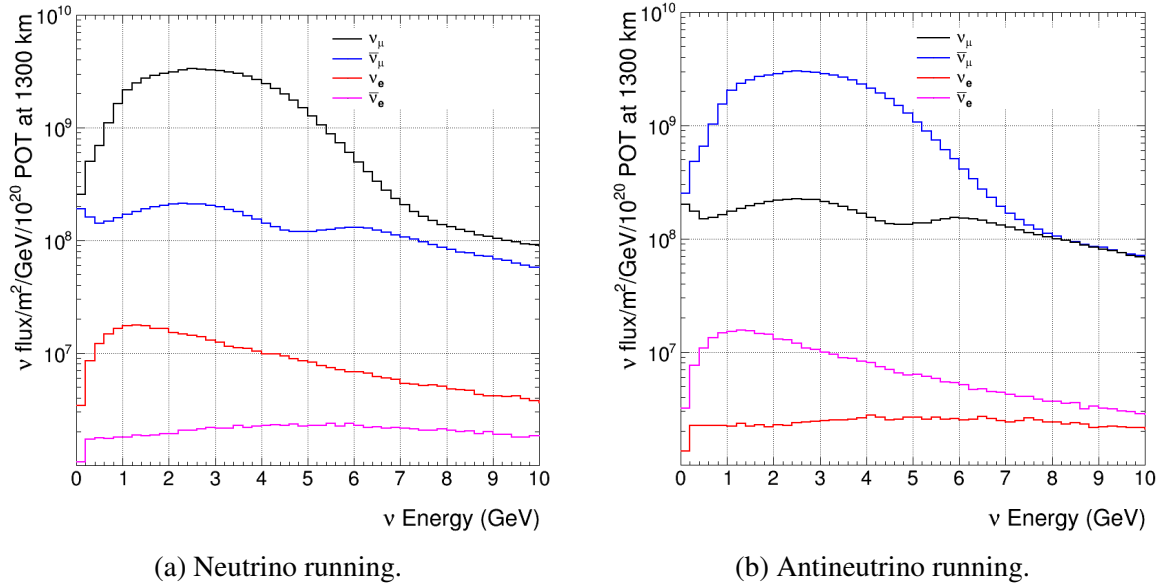


Fig. 3.8 The fluxes of the different neutrino flavour components of the DUNE beam in neutrino and antineutrino running mode [93].

foreseen to come from a thorough assessment of the target-horn system and will influence the reference design as the beam plans become more mature.

3.3.2 Far Detector

There are two potential designs for a far detector LArTPC, utilising either single-phase or dual-phase LArTPCs. Both are under consideration in the upcoming ProtoDUNE prototypes (described in Section 3.5) and final decisions will be taken upon their completion and on reflection on the status of the technology. The first module will certainly employ a single-phase design with future instalments potentially using either. The design of the cryostats, demonstrated in Figure 3.9, is such that only minor modifications would be required when switching between technologies. Each cryostat will hold a fiducial/active/total LAr mass of 10.0/13.3/17.1 kt.

The basic features of each detector design will be overviewed in the following sections. As the first detector will be single-phase, and the work detailed in this thesis pertains solely to this detector design, a much greater emphasis will be placed on this design choice in the proceeding discussions.




Fig. 3.9 The layout of the four cryostats underground at SURF comprising the DUNE far detector [94]. The detectors occupying the cryostats here are of the single-phase design, however they can also be used to accommodate detectors utilising dual-phase technology.

3.3.2.1 Single-Phase

The layout of the single-phase DUNE detector design is shown in Figure 3.10. The modular form of the detector ensures it is flexible, scalable and facilitates the transportation and assembly of component parts. The detector units are named Anode Plane Assemblies (APAs) and Cathode Plane Assemblies (CPA) and, as their names suggest, carry the required voltages necessary to establish the electric field. The active volume is 12 m high, 14.5 m wide and 58 m long in the beam direction, with the APAs and CPAs placed such that their planes are parallel to the beam. Each APA is 2.3 m wide and 6 m high, requiring them to be stacked two high and 25 long to instrument the required volume. The CPAs have the same width but half the height so must be stacked four high and, given the drift distance of 3.6 m, carry a voltage of -180 kV in order to provide the required field of 500 V/cm. Each module therefore contains 150 APAs and 200 CPAs, surrounded on its open sides by a field cage. The entire TPC is suspended from the ceiling of the cryostat on rails.

The basic design of an APA is demonstrated in Figure 3.11. Each comprises four sets of wire planes; from the outside in: a grid plane (G), two induction views (U and V) and a collection plane (Z). The APAs are designed to maximise active detector region and are therefore stacked as closely together as possible and read out charge from multiple drift regions simultaneously. This requires all readout be positioned at the top (or bottom of lower APAs) of the frame and make use of a wrapped wire approach for the induction wires, since these must be fixed at an angle. This wrapping also carries the added benefit of reducing the number of channels required since a wire from the U plane on one side of an APA is




Fig. 3.10 The basic design of a single-phase DUNE far detector module [94]. The beam direction is into the page and the electric field, provided by the Anode Plane Assemblies (APAs) and Cathode Plane Assemblies (CPAs), is perpendicular to this across the page.

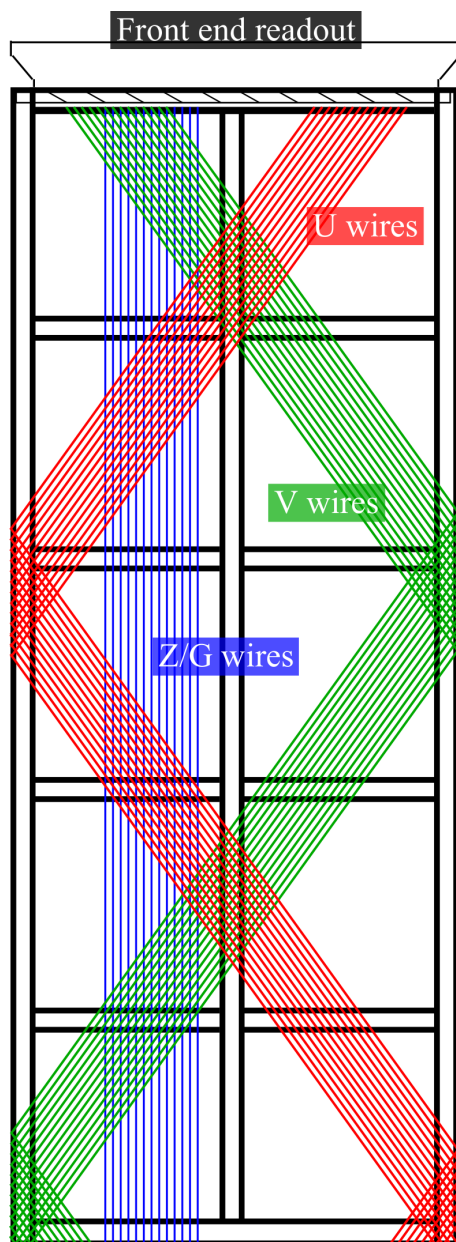


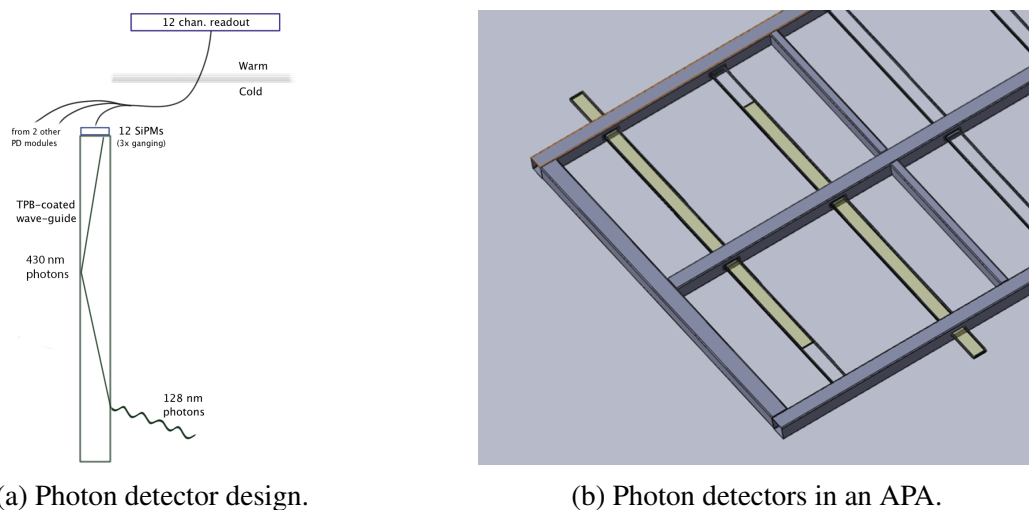
Fig. 3.11 Design of a DUNE far detector Anode Plane Assembly (APA). The three instrumented planes, collection (Z) and two inductions (U and V), are shown in blue, red and green respectively. The grid plane is not shown but is parallel to the collection wires. The induction wires are wrapped around the frame and make angles of $\pm 35.7^\circ$ to the vertical, chosen so that each wire segment crosses each collection channel no more than once. This ensures there is no ambiguity in the side of the APA at which the charge arrives.

part of the V plane on the opposite side. The collection and grid wires are both vertical, necessitating a set for each face of an APA. A single wire reading out multiple drift regions concurrently potentially leads to an ambiguity in the particular wire segment at which the charge arrived. This is broken in one of two ways; either the wires in the U and V planes are placed at slightly different angles, reducing the number of triple crossing points with a given collection wire, or the angle of the induction plane wires is chosen such that each wire segment intersects no more than one collection wire. The former is prototyped in the 35-ton experiment whilst the current plan for the DUNE far detector is to opt for the latter. An angle of $\pm 35.7^\circ$ ensures this is the case, demonstrated in Figure 3.11. A grounded mesh, with good optical transparency, is fixed to the face of the APA frame behind each collection plane to ensure a uniform electric field in the region inside the APA, between the two sets of biased collection wires and the grounded frame.

The front-end electronics are mounted on the APAs and function within the cryostat; it is for this reason they are referred to as ‘cold electronics’. They are implemented as two successive ASIC chips, the first providing amplification and signal shaping and the second the digitisation of the signal. Early versions of these chips were utilised in the 35-ton experiment for front-end readout.

The external triggering in the DUNE far detector relies on detecting the scintillation light from immediate recombination of ionisation electrons with the argon ions. This requires dedicated photon detectors which may detect this light with suitable efficacy. At a field of 500 V/cm, the photon yield is around 20,000 per MeV at wavelengths of 128 nm. The current reference design, depicted in Figure 3.12, involves placing light-guides, containing wavelength shifter, at the centre of each APA frame and using SiPMs for readout. There will be 10 photon detectors per APA, to be inserted after wire wrapping, with dimensions 2.2 m long, 83 mm wide and 6 mm thick.

Overall, optimisation of this detector design involves considerations of the physics reach as well as the associated cost, schedule and risks. Relevant decisions concern the spacing between the readout wires (‘pitch’), which directly affects the detector resolution; the spacing between the wire planes, affecting the shape of the field; the wire angle, necessary for reconstruction but with possible ambiguities; the wire length, affecting the noise levels observed; and the maximum drift length, which influences the size of signal observed. Each has been optimised to ensure the physics goals are reached while reducing as much as possible the cost of the detectors.



(a) Photon detector design.

(b) Photon detectors in an APA.

Fig. 3.12 The design of the photon detectors for the DUNE far detector [94]. The basic design of each bar is shown in Figure 3.12a and their positioning within an APA is shown in Figure 3.12b.

3.3.2.2 Dual-Phase

The DUNE dual-phase LArTPC design differs from its single-phase counterpart in that the cathode is placed at the bottom of the cryostat, with the anode readout at the top (above the gas layer), demonstrated in Figure 3.13. This sets up a field in the vertical direction and results in upwards electron drift. The detector module is a single volume, 60 m long, 12 m wide and 12 m high. Due to the gain provided by the gas phase, the required argon purity is equivalent to the single-phase detector despite the electron drift being more than three times greater. Each module additionally contains 180 PMTs (1 per 4 m²), located underneath the cathode.

The amplification uses Large Electron Multipliers (LEMs): printed circuit boards with electrodes at the top and bottom surfaces across which a potential difference is applied. The data collection is performed by modules named Charge Readout Planes (CRPs), each containing many LEM/anode sandwiches and two independent, orthogonal readout views. The basic operation, including the electric fields at each point in the electron extraction, amplification and readout, is demonstrated in Figure 3.14. As with the single-phase design, the nominal drift field in the liquid is 0.5 kV/cm. The electrons are extracted from the liquid with 100% efficiency using a 2 kV/cm field and are amplified in a field on the order of 30 kV/cm, providing a gain between 20-100. The readout utilises two perpendicular collection planes at the same level which allow full 3D reconstruction when combined with the location in the drift direction from photon detector information. Each CRP has




Fig. 3.13 The DUNE dual-phase detector (partially open) [94]. The cathode is at the bottom of the cryostat, with the PMTs responsible for detecting the scintillation light beneath. The anode is at the top, resulting in upwards electron drift, above a layer of gaseous argon present to introduce a gain in the collected signal.

dimensions $3 \times 3 \text{ m}^2$ (containing 36 $0.5 \times 0.5 \text{ m}^2$ LEM/anode sandwiches) and each module requires 80 such instruments to read out the active volume.

3.3.3 Near Detector

The primary role of the near detector is to characterise the energy spectrum and the composition of the neutrino beam before oscillations have occurred, and to make measurements of neutrino interaction cross-sections. This is necessary to control systematic uncertainties and requires good understanding of the muon- and electron-flavoured neutrino and antineutrino components of the beam. The unprecedented number of neutrino interactions collated by the near detector will also facilitate a broad program of ancillary physics and is an important constituent of the DUNE experiment.

The near detector will comprise of a near neutrino detector (NND) to perform these studies and a Beamline Measurement System (BLM) located downstream of the beam absorber designed to make measurements to constrain the neutrino flux at the near and far detectors. The NND and BLM systems will be the subject of Sections 3.3.3.1 and 3.3.3.2 respectively.




Fig. 3.14 The extraction, amplification and readout of the ionisation electrons through the gaseous argon phase in the DUNE dual-phase LArTPC design [94]. The voltages at each stage are shown on the left and the associated fields set up are described on the right.

3.3.3.1 Near Neutrino Detector

The NND is required to make precise measurements of neutrino interactions and distinguish between all four particle species present in the beam in order to characterise their spectra. The exact form of the detector is still under discussion with multiple designs being considered. The current consensus sees a composite detector consisting of a small LArTPC (to facilitate greater control over systematic uncertainties) followed by a low-mass argon-based magnetised spectrometer surrounded by an Electromagnetic Calorimeter (ECAL) and muon detector. There are two possibilities for both the magnet design and the tracker region leading to four configurations from which the final design will be selected: the magnet could either be a dipole or utilise the hardware from the former KLOE experiment [116] and the tracker may either consist of a Straw-Tube Tracker (STT) or a high-pressure gaseous argon TPC (which possess excellent tracking capabilities and momentum and energy resolution). The design which has been developed furthest and was documented in the CDR [94] includes a dipole magnet and STT (though without the additional upstream LArTPC) and so will be briefly presented here to provide some means of context.

This reference design utilises a magnetised Fine-Grained Tracker (FGT) detector containing a STT and ECAL, shown in Figure 3.15. The STT and ECALs are surrounded by a 0.4 T dipole magnet and also employ muon identifiers (MuIDs) in the magnet steel and upstream/downstream of the STT. The detector is yet to be finalised and possibly a different design will be chosen instead, or additional detectors added such as a small-scale LArTPC or high-pressure argon TPC to further reduce systematic uncertainties.

The STT consists of over 107,000 tubes made from carbon and aluminium and with outer diameter 1 cm to provide fine-grained tracking (thickness $< 0.1X_0$) with excellent angular and spatial resolution. The tubes will be filled with either 70% Ar, 30% CO₂ or 70% Xe, 30% CO₂, depending on the positioning of the tubes in the detector, and are read out at




Fig. 3.15 Schematic of the DUNE near neutrino detector fine-grained tracker design [94].

both ends. Nuclear targets will be placed upstream with multiple materials used to make complimentary measurements. Pressurised argon gas or calcium will allow nuclear effects of neutrino interactions on argon and modelling of signals and backgrounds in the DUNE far detector. Carbon (graphite) may also be used to study interactions on free protons and H₂O and D₂O targets will allow analysis of interactions on quasi-free neutrons via statistical subtractions.

The ECAL provides high segmentation in both the transverse and longitudinal directions and is composed of three separate systems: Forward ECAL, Barrel ECAL and Backward ECAL. It is designed to reconstruct photons from π^0 decay and electrons and positrons from bremsstrahlung radiation and consists of layers of either 1.75 mm thick (Forward ECAL) or 3.5 mm thick (Barrel and Backward ECALs) lead sheets positioned within 2.5 cm wide, 10 mm thick plastic scintillator bars.

The 0.4 T magnet has inner dimensions 4.5 m wide, 4.5 m high, 8.0 m long and assists with measurements of particle momentum and charge. It is also instrumented with the MuIDs designed to distinguish muons from hadrons, utilising the ability of muons to penetrate the iron without showering or interacting. The MuID is comprised of 432 resistive plate chamber (RPC) modules distributed within two 10 cm thick steel plates of the magnet and will also perform basic reconstruction on the muon track segments which can be matched with tracks from the STT tracker for global object reconstruction.

3.3.3.2 Beamline Measurement System

The beamline measurement system will operate for the life of the experiment and will monitor the beam profile on a spill-by-spill basis. Along with constraining the neutrino flux, it will be used to monitor the pulse-to-pulse variations for beam diagnostics and make measurements for validating the simulation of the thick target, horn material, decay tunnel and absorber.

The main concern of beamline measurements is to ensure a good understanding of the neutrino flux, critical for all the physics DUNE wishes to do, along with making measurements of the components of the beam and its stability over time. The flux will be measured by independently studying the muons which penetrate the absorber, a result of the meson decays which also created the neutrinos, in order to determine their rate to an accuracy of 5%. By combining this information with details of the energy spectrum observed by the muon monitors, the neutrino flux in the beam will be constrained at around the 4-5% level. Additionally, since a major background for neutrino oscillations at the far detector is the intrinsic ν_e component of the beam resulting from muon decays, detecting these muons with a low threshold will provide further measurements of the beam constituents. Due to the long distances travelled by the neutrinos, the relative fluxes at the near and far detectors may be

Table 3.2 Exposures anticipated for the DUNE experiment for the first few years of operation. Due to the staged approach in construction, it will take some time to reach full design capabilities. The first exposure column represents the exposure expected in that year and the next column the cumulative total. Adapted from [68].

Year	Exposure (kt·MW·year)	Total (kt·MW·year)	Detector stage
Year 1	10.7	10.7	10 kt far detector, no near detector, 1.07 MW 80 GeV proton beam (1.47×10^{21} pot per year)
Year 2	21.4	32.1	Addition of second 10 kt far detector module
Year 3	32.1	64.2	Addition of third 10 kt far detector module and initial constraints from near detector
Year 4	42.8	107.0	Addition of fourth 10 kt far detector module
Year 5	42.8	149.8	Inclusion of constraints from full near detector data analysis
Year 7	85.6	278.2	Upgrade beam power to 2.14 MW for 80 GeV proton beam

significantly altered by changes in direction; the stability of the beam must thus be assured to within 0.2 mrad, corresponding to a measurements of the position of the muons with an accuracy of 5 cm at the near site monitors.

3.4 The Physics of DUNE

The staged approach to the DUNE experiment will allow early preliminary results but will require more time for facilities from later phases to be constructed and commissioned. For this and other reasons, the accumulated data is often referred to as an ‘exposure’, a function of detector size, beam power and time with units kt·MW·year. The current assumptions on exposures for the first few years of operation are shown in Table 3.2. This staging will be assumed in all sensitivities presented in this section.

The appearance probability expected at the DUNE far detector is demonstrated in Figure 3.16 for various values of δ_{CP} . It can be seen why a broadband beam with the ability to operate in neutrino and antineutrino mode mode is critical; the value of δ_{CP} affects both the




Fig. 3.16 The appearance probability at a baseline of 1300 km, as a function of neutrino energy, for $\delta_{CP} = -\pi/2$ (blue), 0 (red) and $\pi/2$ (green) for neutrinos (left) and antineutrinos (right), for normal hierarchy. The black lines indicates the oscillation probability if θ_{13} were equal to zero. Taken from [92].


frequency and amplitude of the oscillations, with differing effects at the distinct oscillation nodes and between neutrinos and antineutrinos.

3.4.1 Mass Hierarchy and CP Violation

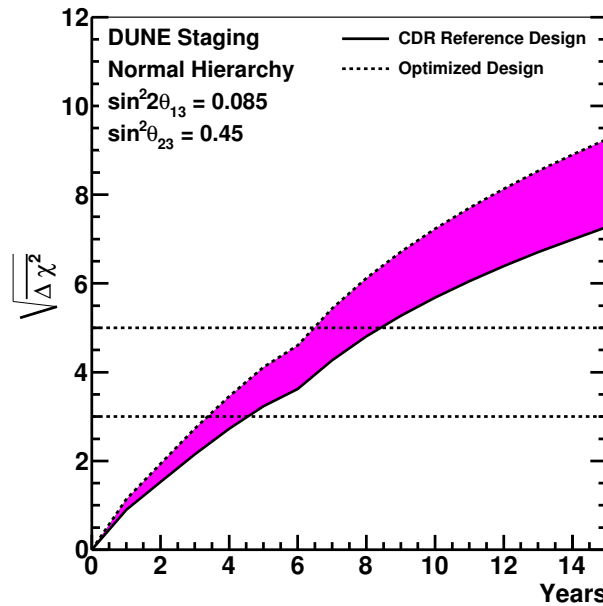
The measurements of the mass hierarchy and the degree of CP violation are determined by simultaneously fitting the $\nu_\mu \rightarrow \nu_\mu$, $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$, $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillated spectra, assuming 50% neutrino, 50% antineutrino exposure.

The sensitivity of DUNE to the neutrino mass hierarchy is shown in Figure 3.17. It is evident DUNE will discover the ordering of the mass states at 5σ significance within the first few years of running, regardless of the value of δ_{CP} . This is achievable because of the large matter effects present with a 1300 km baseline; the asymmetry between neutrinos and antineutrinos due to this is approximately $\pm 40\%$ in the region of peak flux.

The equivalent sensitivities for CP violation is displayed in Figure 3.18. It is impossible to cover 100% of δ_{CP} values because, in the case of CP conservation, the violation effects (i.e. disparities between neutrino and antineutrino oscillations) disappear. DUNE will be able to measure 75% of possible δ_{CP} values with 3σ significance after 1320 kt·MW·year and, in the case of near-maximal CP violation currently hinted at by T2K [86], 50% of these large CP-violating δ_{CP} values will be measured at 5σ confidence with an exposure of



(a) The significance with which the mass hierarchy can be determined as a function of the value of δ_{CP} for an exposure of 300 kt-MW-year assuming normal hierarchy (left) and inverted hierarchy (right). Taken from [92].



(b) Assuming normal hierarchy, the minimum significance (the lowest point on the curve in Figure 3.17a) with which the mass hierarchy can be determined for all values of δ_{CP} as a function of years of running under the assumptions in Table 3.2. Taken from [68].

Fig. 3.17 Sensitivity of the DUNE experiment to the neutrino mass hierarchy.

810 kt·MW·year, around 14 years of running following the staging described in Table 3.2. A full scope experiment operating with multi-megawatt beam power will eventually achieve around a 5% precision of the CP-violating phase, comparable to the equivalent precision in the quark sector described in the CKM matrix.

3.4.2 Oscillation Parameters

DUNE will make precision measurements of all parameters describing neutrino oscillations and improve our understanding of oscillation phenomenology. The least known mixing angle, θ_{23} , will be measured with a precision of at least 1° , even near 45° . This is possible by performing a combined analysis of the $\nu_\mu \rightarrow \nu_\mu$ and $\nu_\mu \rightarrow \nu_e$ channels, dependent on $\sin^2 2\theta_{23}$ and $\sin^2 \theta_{23}$ respectively (from Equations 2.31 and 2.30). The sensitivity of DUNE to the octant of θ_{23} , and the resolution of the value itself, is presented in Figure 3.19. The current world-leading measurements of θ_{13} from reactor experiments will also be achieved after sufficient exposure in the DUNE experiment, using analyses of ν_e and $\bar{\nu}_e$ appearance. The anticipated resolution of this parameter along with the associated mass splitting, Δm_{31}^2 , is shown in Figure 3.20.

3.4.3 Other Physics

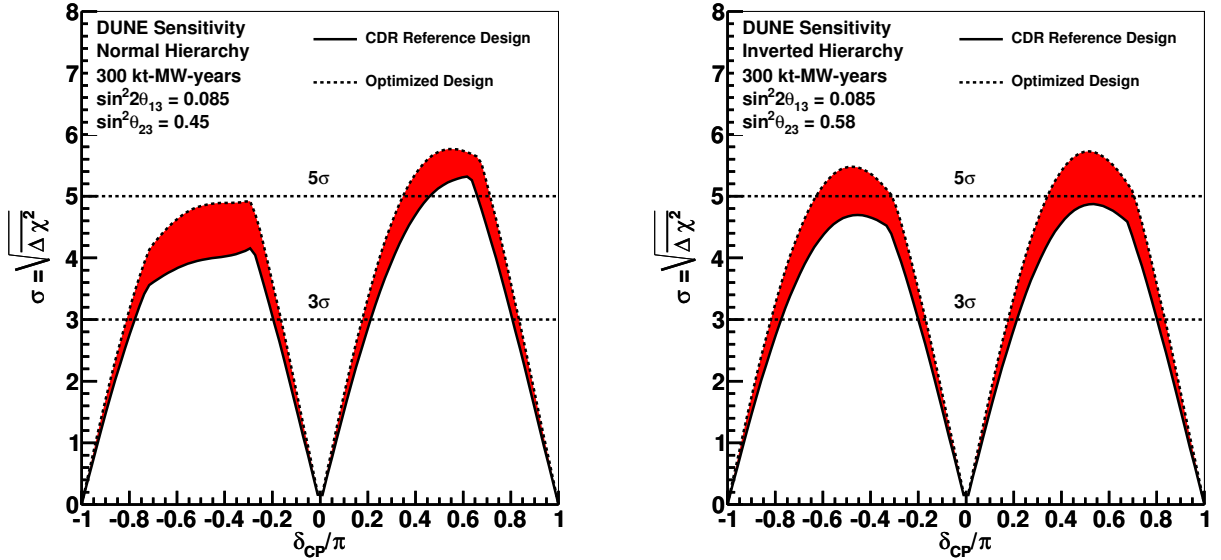
The LBNF/DUNE physics program is diverse and contains priorities unrelated to beam neutrino physics. These are not wholly relevant to this thesis and shall not be discussed in detail, but it would be inappropriate to ignore the additional physics potential of the experiment.

The capabilities to search for nucleon decay, mainly via the $p \rightarrow K^+ \bar{\nu}$ channel [117], will improve lifetime limits by an order of magnitude after 20 years' running. It will significantly extend sensitivities compared to water Cherenkov detectors as a consequence of a higher detection efficiency and low background rates. Many models in which this mode is dominant (e.g. supersymmetric GUT models) also favour other models with final-state kaons, enabling a wide range of nucleon decay physics to be studied at the DUNE far detector.


The unique sensitivity to supernova ν_e neutrinos, via the interaction

$$\nu_e + {}^{40}\text{Ar} \rightarrow e^- + {}^{40}\text{K}^*, \quad (3.4)$$

inaccessible to water or liquid scintillator detectors which detect the $\bar{\nu}_e$ component via inverse beta decay on free protons, will, along with $\bar{\nu}_e$ interactions, allow precision measurements of the neutrino spectrum and flavour composition of supernovae should one occur during the



(a) The significance with which CP violation can be determined as a function of the value of δ_{CP} for an exposure of 300 kt-MW-year assuming normal hierarchy (left) and inverted hierarchy (right). Taken from [92].



(b) Assuming normal hierarchy, the significance with which CP violation can be determined for 75% of δ_{CP} values (left) and the expected 1σ resolution (right) as a function of years of running under the assumptions in Table 3.2. Taken from [68].

Fig. 3.18 Sensitivity of the DUNE experiment to leptonic CP violation.




Fig. 3.19 The sensitivity of DUNE to the octant and value of θ_{23} . Figure 3.19a shows the significance with which DUNE can resolve the θ_{23} octant as a function of the true value of θ_{23} . The green shaded band around the curve represents the range in sensitivity due to potential variations in the beam design and in the true value of δ_{CP} . The yellow shaded regions indicate the current 1σ and 3σ bounds on the value of θ_{23} from a global fit. An exposure of 1320 kt-MW-year, which will provide a 3σ measurement of CP violation for 75% of the values of δ_{CP} is assumed. The shaded region represents the range in sensitivity due to potential variations in the beam design. Figure 3.19b shows the resolution of a measurement of $\sin^2\theta_{23}$ as a function of exposure assuming normal mass hierarchy and $\sin^2\theta_{23} = 0.45$ from the current global fit. Taken from [92].




Fig. 3.20 The sensitivity of DUNE to the oscillation parameters describing ν_e and $\bar{\nu}_e$ appearance as a function of exposure assuming normal hierarchy. Figure 3.20a shows the resolution of a measurement of $\sin^2 2\theta_{13}$ and Figure 3.20b shows the resolution of Δm^2_{31} assuming, from current global fits, $\sin^2 2\theta_{13} = 0.085$ and $\Delta m^2_{31} = 2.457 \times 10^{-3} \text{ eV}^2$. The shaded regions represent the range in sensitivity due to potential variations in the beam design. Taken from [92].

running of the experiment. These events may be identified via the coincidence of the emitted electron and the accompanying photon cascade from the $^{40}\text{K}^*$ de-excitation. The neutrinos from a core-collapse supernova are emitted in a burst of a few tens of seconds, with around half the luminosity in the first second. They are roughly divided equally between the three known neutrino flavours and have energies between 5-50 MeV. A large neutrino signal from a supernova will allow unique studies of astrophysical phenomena related to the end of the lifetime of a star.

3.5 The Road to DUNE

The staging presented in Table 3.2 is hugely ambitious and requires a great deal of preparation to ensure the project is a success. The preparations have commenced with a program of prototyping and construction to ensure assembly of the first 10 kt module can begin as scheduled in 2021. Work at the far detector site is underway and construction at the near detector will start next year to ensure the beam and near detector facilities are ready for 2026.

Each 10 kt module is larger than any LArTPC ever operated by well over an order of magnitude and it is crucial the detector technology is understood and the experiment operates as expected. To ensure this is the case, a comprehensive prototyping strategy is planned. This will be briefly overviewed in this section.

3.5.1 The 35-ton Prototype

The first prototype to test many of the design features of the DUNE far detector was the 35-ton, shown in Figure 3.21. This is the subject of the majority of this thesis and will be fully discussed in Chapter 4. Lessons learned from the 35-ton experience are already influencing the design choices and operating schedules of the next prototypes and the far detector, highlighting the benefits of a thorough research and development programme.

The 35-ton experiment demonstrated the cryostat and detector technologies over the course of two runs, in early 2014 and early 2016 respectively, by collecting and analysing data from cosmic-induced tracks and showers. Analysis of data from the second run is the subject of Chapter 7.

3.5.2 ProtoDUNE

The next prototypes, ProtoDUNE, will take data during the second half of 2018 and will be a full scale engineering prototype of the DUNE far detector. ProtoDUNE is hosted at the Neutrino Platform at CERN and consists of two demonstrators utilising single-phase




Fig. 3.21 The 35-ton cryostat and detector designed to prototype the DUNE far detector design [118].

(ProtoDUNE-SP) [119] and dual-phase (ProtoDUNE-DP) [120] designs. The layout is shown in Figure 3.22. The cryostats contain more than 700 t LAr and represent an intermediary stage between the 35-ton and the DUNE far detector modules.

The detectors are being constructed in a CERN test beam and will be subjected to various particles, including e^\pm , μ^\pm , K^\pm , p , \bar{p} . Along with characterising the detector performance and identifying potential improvements, ProtoDUNE will provide calibrations, such as electromagnetic shower energy resolution, electron/photon separation and particle identification, for the far detector and facilitate an assessment of the detector systematics. It will also provide opportunity, in a detector functionally identical to the far detector modules, to validate the simulation and reconstruction techniques developed in Monte Carlo. Upon completion of the ProtoDUNE program, DUNE will be prepared to begin final preparations and construction of the first, single-phase, far detector module.




Fig. 3.22 Schematic showing the layout of the ProtoDUNEs at the CERN neutrino platform [121]. The cryostat at the left, on a slight angle, is the dual-phase LArTPC and the single-phase detector is at the right hand side. The beamline, running from left to right through both detectors, may also be faintly discerned.

Chapter 4

The DUNE 35-ton Prototype

The 35-ton is the first experimental prototype of the DUNE far detector design and was briefly introduced in Section 3.5. It was originally constructed to demonstrate the unique design features of the LBNE far detector and was the only planned prototype for this experiment. Following the dissolution of LBNE and the subsequent formation of the DUNE collaboration, the 35-ton has become an integral part of the design and execution of the DUNE far detector design.

As discussed in Section 3.1, the use of LArTPCs in future long-baseline experiments shows great promise. To facilitate development of the detector technology, Fermilab has an extensive program of LArTPC experiments culminating in the flagship DUNE project. Prototyping is essential to the success of DUNE as understanding of how to operate progressively larger detectors evolves. The strategy is staged, with each subsequent phase building on previous success.

The most pertinent issues facing large-scale LArTPCs concern:

- the ability to achieve and maintain the necessary LAr purity for successful data taking;
- the design and construction of large underground cryostats.

The research and development performed thus far have demonstrated viable solutions to these obstacles and reinforced confidence in the design of the experiment and in the upcoming ProtoDUNE projects.

The outcomes of each of these projects at Fermilab are the subject of this present chapter. The first of the above issues, regarding LAr purity, is discussed in Section 4.1 with reference to the Materials Test Stand and the Liquid Argon Purity Demonstrator. The second complication, concerning the construction of large underground cryostats, was the main motivation for the 35-ton Phase I experiment and is the subject of Section 4.2. The author had no direct

involvement in these earlier prototyping efforts. The culmination of all these developments involved operating a small scale LArTPC alongside these improvements and was achieved in the 35-ton Phase II run, discussed in Section 4.3. Since this experiment forms the basis for later chapters and was actively worked on by the author, it will be reviewed in much greater detail. A summary of all this R&D is presented in Section 4.4.

4.1 The Materials Test Stand and Liquid Argon Purity Demonstrator

Work developing LArTPCs for future neutrino experiments began at FNAL in 2007 with a view to eventually facilitating a multi-kton LAr experiment. Even utilising a modular design, as with the DUNE far detector (Section 3.3.2), drift distances on the order of a few metres are realistically required, necessitating a low concentration of electronegative impurities. Attaining and holding the requisite LAr purity in a huge underground cryostat over many years of running is a considerable challenge addressed by the test stands reviewed in this section.

4.1.1 The Materials Test Stand

The Materials Test Stand (MTS) [122–125] was constructed at FNAL to develop LAr purification techniques and to characterise the effect of various materials on the electron lifetime when submerged in the liquid. It consists of a small cryostat and two filters containing activated-copper-coated granules and an adsorbent molecular sieve respectively; a schematic of the MTS setup is shown in Figure 4.1. The filters are designed to remove oxygen and water contaminants with functionality similar to that successfully demonstrated by the ICARUS collaboration [126]. Oxygen is removed by the copper beads using the chemical reaction



and water molecules are physically trapped in the microporous structure of the sieve. The filters additionally contain the ability to be regenerated in situ, a necessity when planning a long-running experiment, multi-kton experiment; those used previously were primarily proprietary [127, 128].

The MTS successfully demonstrated good argon purity (< 3 ppb H_2O) and showed the primary opposition to electron lifetime is water contamination, demonstrated in Figure 4.2. It was found that exposure to warm surfaces in the cryostat, such as in the ullage (the volume

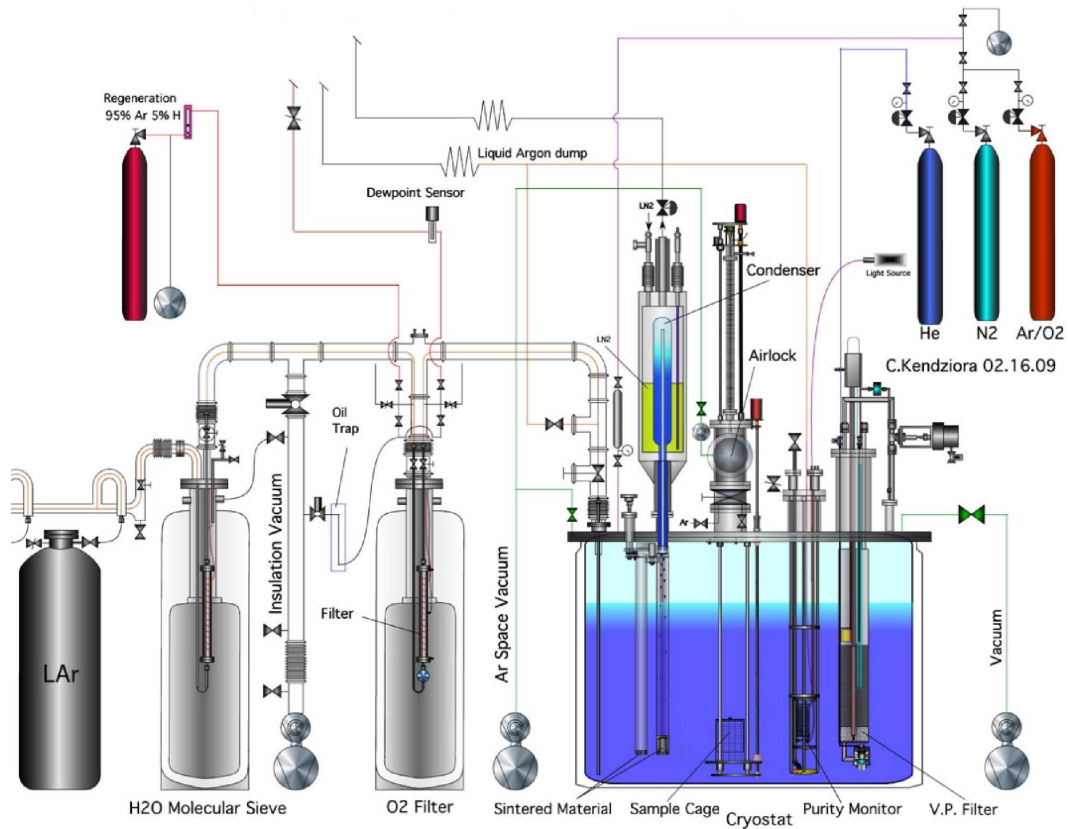


Fig. 4.1 The Materials Test Stand at FNAL [124]. Liquid argon used to fill the cryostat flows from left to right in the schematic, through two filters designed to reduce the H_2O and O_2 contamination respectively. A second filter system (the 'vapour pump', V.P.), using the same materials, is installed within the cryostat to remove impurities introduced by the materials being examined.




Fig. 4.2 Results from the Materials Test Stand showing the water contamination in LAr and the corresponding electron lifetime [123]. There is an obvious inverse correlation between the density of electronegative (H_2O) impurities and the resulting lifetime.

above the liquid level), facilitated contamination from water impurities as they remain on surfaces even in a vacuum. The condenser used in the MTS to recondense gaseous argon returned it directly to the liquid in the cryostat (as ‘raining’ condensation) and was found to dramatically reduce the LAr purity when in use. This is due to contaminants introduced into the gas by exposure to the warm cryostat walls, which could be negated by returning the liquid via a different path which maximised subjection to cold surfaces. Notably, the electron lifetime was found to be unaffected on the introduction of test materials, although as suspected the temperature of the materials did have an impact. This is a hugely promising result for the future of LArTPC design and construction.

4.1.1.1 Filter Regeneration

Over time, the filters become less effective as electronegative impurities accumulate. A significant success of the MTS was demonstrating the process of regenerating the filters in situ. This is achieved by heating the vessels to 250°C and, in the case of the molecular sieve, simply using a vacuum pump to remove the water vapour or, in the case of the activated copper, by pumping through a 95:5 mixture of Ar: H_2 gas to capture the oxygen through the reduction reaction



During the running of the test stand, the filters were regenerated after the passage of around 1000 litres of liquid argon. The process takes on the order of a week to heat the filter

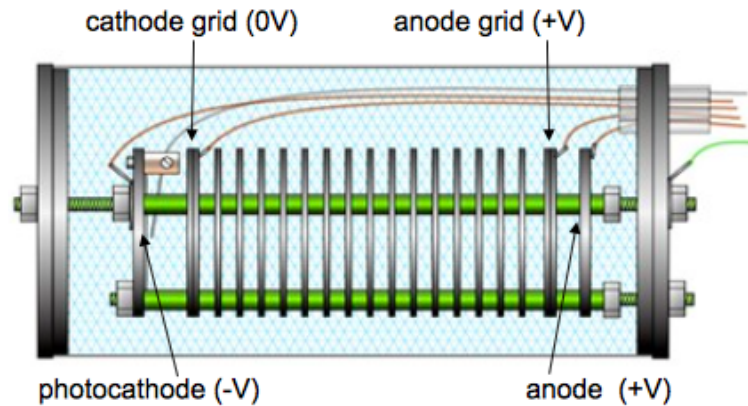


Fig. 4.3 Schematic design of the purity monitors utilised at the FNAL LAr test stands [130]. Purity monitors using this design were pioneered by ICARUS [129] and used in the MTS along with the subsequent Liquid Argon Purity Demonstrator (Section 4.1.2) and 35-ton Phases I (Section 4.2) and II (Section 4.3).

sufficiently and allow the impurities to be removed; DUNE will utilise replacement modules during this time to ensure continued recirculation.

4.1.1.2 Purity Monitoring

The ability to constantly evaluate the LAr purity during an experimental run is critically important to ensure high quality data. The impurity concentrations are typically beyond the capabilities of many conventional gas analysers and so a custom device, known as a ‘purity monitor’ (PrM), is utilised. The design is based on the purity monitors developed by ICARUS [129] and is shown in Figure 4.3.

The PrM consists of a cylindrical volume containing LAr from its surrounding environment and an anode and photocathode separated by a short drift region. When taking purity measurements, light from a Xenon flash lamp is incident on the cathode, liberating photoelectrons which traverse towards the anode. Electronegative impurities in the LAr will decrease the electron lifetime, and therefore the number of electrons reaching a certain point along the drift volume. A measurement of the ratio of the charge arriving at the anode to that at the cathode is hence a measurement of the inherent purity of the liquid.

The MTS cryostat contains a purity monitor and they were subsequently used in the Liquid Argon Purity Demonstrator and the 35-ton. When developed for the Liquid Argon Purity Demonstrator and 35-ton cryostats, two sizes were used; long (47 cm) and short (16 cm).

4.1.2 The Liquid Argon Purity Demonstrator

The Liquid Argon Purity Demonstrator (LAPD) [125, 131, 132] was designed to demonstrate the required purity of LArTPC experiments is possible without the use of large scale vacuum pumps. Previous and current LArTPC experiments, such as ICARUS, ArgoNeuT, LArIAT and MicroBooNE, have been constructed as flat plane vessels and have used an evacuation method as the first step in removing atmospheric impurities to facilitate the required LAr purity. The necessary mechanical capability of the cryostat to withstand this process, along with the associated equipment, results in unfeasible engineering challenges and costs as detectors increase to multi-kton scales.

In order to circumvent these issues, a design utilising multiple smaller-scale cryostats was proposed. This however leads to greater complexity relating to both the engineering requirements of the piping infrastructure and the reconstruction capabilities of interactions spanning multiple active volumes. LAPD successfully pioneered an alternative approach, using a ‘piston purge’ as a first purification step to remove atmospheric impurities. This is a very important result and has significantly influenced the design of future LArTPC experiments, including the 35-ton. Additionally, although designed to be evacuated with vacuum pumps, MicroBooNE was filled using the piston purge technique following the success of LAPD.

4.1.2.1 LAPD Experimental Setup

The LAPD cryostat is shown in Figure 4.4. It consists of a cylindrical tank, diameter 10 feet and height 10 feet, with a domed head and a capacity of 32.6 ton LAr. It is physically next to the MTS and uses the purification system prototyped by this previous effort. Insulation for the tank is provided by fibreglass sheets covering the outer volume which, along with the tank, is refrigerated by liquid nitrogen (LN_2) from an external supply. As with the MTS, a condenser is utilised above the cryostat to recondense argon gas using coils also cooled with LN_2 . This liquid is subsequently sent through the filtration system before being returned to the main volume, a consequence of the previous R&D with the MTS. After filling, the system is closed and a good LAr purity is maintained by constant circulation of the cryostat content through the filters.

The system is instrumented with PrMs, gas analysers and temperature sensors. Four PrMs are contained within the cryostat to measure the purity gradient, with an additional monitor just after the filters to sample to liquid before it is returned to the main volume. Along with purity, the temperature gradient is measured in order to study the effect of this on





Fig. 4.4 The Liquid Argon Purity Demonstrator cryostat and purification system [132]. The two cylinders at the bottom left are the filters described in Section 4.1.1. The piping facilitates the transport of LAr into and out of the cryostat so continuous purification within a closed system may be achieved.


electron drift velocity. The contaminants in the LAr are quantified using nitrogen, oxygen and water analysers outside of the main volume.

4.1.2.2 Filling LAPD

The piston purge technique involves injecting warm argon gas at high pressure at the bottom of the cryostat with the top open for venting, demonstrated in Figure 4.5a. The heavier than air argon gas acts as a piston, forcing the ambient air out of the top of the cryostat. Figure 4.5b demonstrates how this successfully reduces the impurity concentration in the cryostat, shown as a function of complete volume changes. After completion of the piston purging, the O₂ contamination had decreased from 21% to 6 ppm, N₂ from 78% to 18 ppm and H₂O from 200 ppm to 1.2 ppm.



(a) Schematic of the LAPD piston purge.



(b) LAPD impurity concentration during the piston purge.

Fig. 4.5 The piston purge technique in the Liquid Argon Purity Demonstrator to remove atmospheric impurities before filling [132]. The results from two LAPD runs are shown, the first with the cryostat only half filled to prototype the technique. Discontinuities between the impurity concentrations are caused by switches between gas analysers.




Fig. 4.6 The concentration of electronegative impurities during the gas circulation stage in the Liquid Argon Purity Demonstrator following the piston purge [132]. The stabilisation of the oxygen contamination signified a leak, which was fixed during the break in readings.

Following the filling of the cryostat with gaseous argon, the contents are then continuously circulated through the filters to further reduce the impurities present. The improved electronegative concentrations are shown, again with reference to the number of complete volume changes, in Figure 4.6. This lasted, as can also be observed in the figure, for a number of days and resulted in a much improved O_2 contamination of around 20 ppb and an H_2O level which balanced the outgassing rate from the warm cryostat surfaces.

The filling can thus proceed by transporting LAr into the cryostat, through the filter system to ensure a high purity is maintained. The impurity concentrations were inspected before filling and after filtration and, in total, a volume of 29.7 ton LAr was supplied to the LAPD cryostat. Once filled, and during the course of operations, the liquid argon volume was constantly recirculated through the filtration system to preserve the LAr purity. This is shown schematically in Figure 4.7.

4.1.2.3 LAPD Outcomes

LAPD successfully demonstrated achieving and maintaining the required LAr purity for a large neutrino detector is possible without the costly and challenging use of evacuation techniques, reaching purities of better than 60 ppt O_2 equivalent. The measured electron lifetimes over the course of a six week run is shown in Figure 4.8. Lifetimes of up to 4 ms were recorded, greater than the DUNE requirement of 3 ms although utilising a much smaller-scale cryostat. Nonetheless, the success of LAPD has great significance for future LArTPCs, including the 35-ton, and was an important stage in the FNAL LAr test program.

4.1.3 LongBo

Following the successful LAPD runs, a further phase involved the introduction of a small-scale TPC detector into the liquid argon [133]. The detector is named LongBo (an upgrade from the smaller Bo test detector) and is cylindrical with 25 cm diameter and 2 m length.

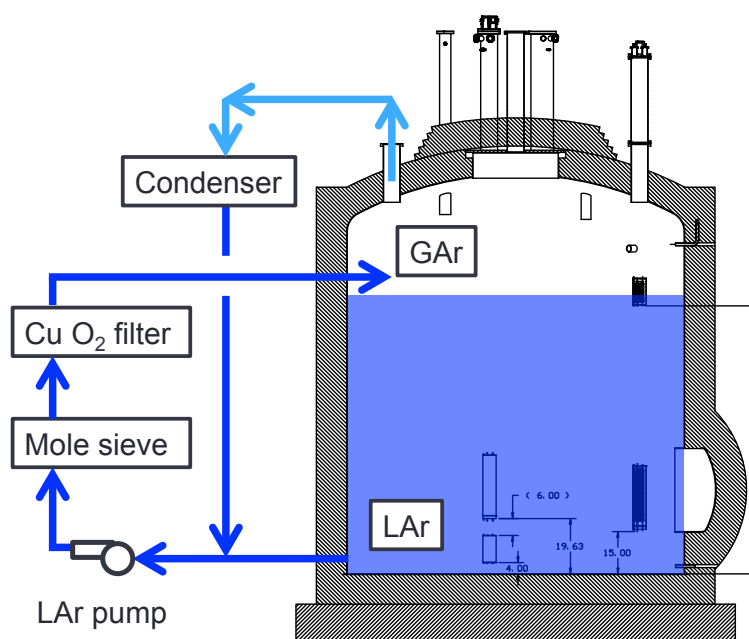


Fig. 4.7 Schematic showing the recirculation of the LAr during commissioning and operations of the Liquid Argon Purity Demonstrator [131]. Liquid is extracted from the bottom of the cryostat and pumped through the filters to remove any impurities which may have established in the medium. Following the experience of previous R&D with the MTS [123], the recondensed liquid is passed through the purification system before being reintroduced to the main volume inside the cryostat.




Fig. 4.8 The electron lifetime achieved in the Liquid Argon Purity Demonstrator during a six week run. Adapted from [131].

It was positioned vertically in the LAPD cryostat, demonstrated in Figure 4.9, and was equipped with a high voltage on the cathode to produce the drift field and three wire planes at the top of the detector for readout. External scintillator counters were placed around the outer wall of the cryostat to provide triggers on through-going cosmic muons which may deposit charge in the detector.

LongBo was the first LArTPC experiment to utilise ‘cold readout’ electronics to amplify and shape the signal at the front end. An early version of the ASICs being developed for MicroBooNE were used to read out 16 of the 144 channels, with the remaining using preamplifiers made with discrete circuitry. At the drift field of 350 V/cm, the signal/noise ratio, a useful number in quantifying the electronics, was around 30, with the channels read out by the ASICs reporting values up to 1.4 times larger.

The LAPD/LongBo experiment successfully maintained similar LAr purities to those without the presence of the detector, as predicted by the results of the MTS. By using TPC data, it was also possible to make measurements of the purity from through-going muons. Equation 3.3 may be used to extract a value for the electron lifetime from a plot of deposited charge determined as a function of drift time, where an exponential decay would be expected due to the attenuation from impurities in the LAr. A comparison between the measured values from the purity monitors and the TPC data may be found in Figure 4.10. A reasonable agreement is observed between these complimentary measurements, with values between 6 ms and 14 ms reported with 95% confidence. These promising results confirmed designing and operating a LArTPC within a non-evacuatable cryostat is viable and contributed to the development of the LAr programme towards the DUNE far detector, with the 35-ton experiment the next stage.




Fig. 4.9 The LongBo TPC detector shown within the Liquid Argon Purity Demonstrator Cryostat [133]. The black tube represents the high voltage feedthrough to the cathode at the bottom of the TPC.




Fig. 4.10 The LAr purity within the Liquid Argon Purity Demonstrator cryostat with the LongBo TPC present, measured using both data from the detector and information from the purity monitors [133]. The ratio Q/Q_0 is defined as in Equation 3.3.

4.2 35-ton Experiment: Phase I

The scale of the cryostats required for the DUNE experiment are such that constructing them as flat plane vessels 1.5 km underground would be unfeasibly expensive and pose great engineering challenges. Following the success of LAPD (discussed in Section 4.1.2), which eliminates the requisite to evacuate the cryostat prior to filling, the LBNE collaboration decided to utilise membrane cryostat technology well established in the liquefied natural gas (LNG) industry. The 35-ton [130, 134, 135] was therefore employed to demonstrate the application of a membrane cryostat to a LAr experiment. The DUNE project has maintained this design choice and the 35-ton has since become a recognised and integral part of the collaboration, providing the first test of the technologies envisioned for the eventual far detector.

The 35-ton cryostat was constructed in 2012 at PC4, a former proton facility in a decommissioned beamline at Fermilab. It has operated in two phases: Phase I (December 2013 – February 2014) was proposed to demonstrate the membrane cryostat technology with just the cryostat and purification systems; Phase II (February 2016 – April 2016) contained a small-scale DUNE-style detector to validate the integrated system and affirm the detector design elements. The Phase I run is the subject of Section 4.2 whilst Phase II is considered in detail in Section 4.3.

The 35-ton is the first membrane cryostat used for scientific purposes and the first overall constructed in the United States. It is also the first designed to contain LAr, which is around three times denser than LNG. The initial aims of the project (Phase I) are to demonstrate the feasibility of the cryostat technology for LAr, including thermal performance and leak tightness, and to show the required LAr purity may be achieved without evacuation and maintained through the use of the filtration system developed and validated by the MTS and LAPD. This first phase will be discussed in this section; the 35-ton cryostat and filling procedures will be described in Sections 4.2.1 and 4.2.2 respectively before outcomes of the experiment are presented in Section 4.2.3.

4.2.1 The 35-ton Cryostat

An overview of the 35-ton cryostat is shown in Figure 4.11. It contains a concrete shell within which the membrane cryostat is constructed from 2 mm thick stainless steel panels. An insulated region between these two segments reduces heat leaking. The roof consists of two plates; Plate A is flat with insulation and membrane beneath and Plate B contains all penetrations and services. Relevant properties of the 35-ton cryostat are listed in Table 4.1.




Fig. 4.11 The 35-ton cryostat [135].

Table 4.1 Details and dimensions of the 35-ton cryostat [135].

Parameter	Value
Cryostat volume	29.16 m ³
LAr total mass	38.6 metric tons
Depth of LAr	2.565 m (11% total ullage)
Inner dimensions	4.0 m (length) × 2.7 m (width) × 2.7 m (height)
Insulation	0.4 m polyurethane foam
Primary membrane	2.0 mm thick corrugated stainless steel
Secondary barrier system	0.1 mm thick fibreglass
Vapour barrier	1.2 mm thick carbon steel
Steel reinforced concrete	0.3 m thick layer
LAr temperature	89 ± 1 K
Operating gas pressure	70 mBar
Design pressure	207 mBar
Heat leak	< 13 W/m ²
Leak tightness	1 × 10 ⁻⁶ mBar·litre/s




Fig. 4.12 The network linking the 35-ton cryostat, the Liquid Argon Purity Demonstrator and the purification system at PC4, Fermilab [130].

The 35-ton was constructed geographically nearby the Liquid Argon Purity Demonstrator in order to utilise existing infrastructure. It is connected to the LAPD tank, which may be used to store LAr before transfer to the 35-ton, and uses the filtration setup designed and validated by the MTS and LAPD. This network is shown schematically in Figure 4.12. Unlike in LAPD, the pumps used in the 35-ton to circulate the LAr through the purification system are within the liquid but the framework operates in a similar way. An identical condenser is also employed above the cryostat to cool boiled off gaseous argon which is returned to the bottom of the cryostat, near to the pumps which subsequently extract the liquid for purification. The LAr was circulated at a rate of three volumes per day by the pumps, with the filters designed to remove 2 ppm O_2 from 35 tons of LAr before saturating. The argon used to fill the cryostat had an initial purity of around 1 ppm O_2 and so the filters were able to operate throughout the planned lifetime of the experiment without the need for regeneration.

The cryogenic environment is monitored and controlled using standard detectors, including temperature sensors, pressure transducers, flow meters and level sensors, along with a suite of commercial gas analysers. The height of the volume is instrumented with four PrMs, two large and two small, with an additional long monitor positioned after the filters, as with LAPD. Also as previously, the vertical temperature profile in the cryostat is monitored at 23 cm intervals with temperature detectors suspended on a chain.

4.2.2 Filling the 35-ton

The 35-ton cryostat is filled in a similar way to the Liquid Argon Purity Demonstrator, described in Section 4.1.2.2. Initially, a piston purge with warm gaseous argon is performed




Fig. 4.13 Filling the 35-ton cryostat in four stages: piston purge, gas recirculation, cooldown, liquid filling [135]. The gas filling is shown in Figure 4.13a and involves using a piston purge to fill the tank with warm gaseous argon before circulating this gas through the filtration system. Cooldown and liquid filling is demonstrated in Figure 4.13b, which shows the falling temperature of the cryostat, as a result of the injection of liquid argon through the cooldown sprayers, and the rising LAr level as the cryostat is filled from LAPD. The red lines represent temperature measured by various sensors, with the relevant scale on the left, and the liquid depth is shown by the dashed blue line and quantified by the right axis.

to remove atmospheric impurities before closing off the vents and redirecting argon at the top of the cryostat through the filters for purification. The impurity concentrations for this stage of filling are shown in Figure 4.13a. Before filling with liquid, the cryostat is cooled in an attempt to reduce outgassing and to create an appropriate environment in which to introduce LAr. This is achieved by injecting LAr through a spray at the top of the cryostat which generates a turbulent mixing of cold gas within the cryostat and gradually cools the walls of the vessel. Following this, LAr is transferred from LAPD into the 35-ton; since the 35-ton is slightly larger than LAPD, this is conducted in two stages. The cooldown and LAr filling stages are shown in Figure 4.13b.

4.2.3 Outcomes of Phase I

The 35-ton Phase I successfully demonstrated the feasibility of membrane cryostats for use with LAr and additionally showed the required LAr purity for future multi-kton LArTPC experiments may be achieved and held in such a vessel. The lifetime over the course of the ~ 2 month run, along with external changes to the system, is comprehensively summarised in Figure 4.14.

The observed lifetime reached the DUNE requirement and was maintained for a significant period of time; this is a major achievement in the context of the future of LArTPC experiments. Dips in the purity were observed when topping up the cryostat after initially




Fig. 4.14 The electron lifetime in the 35-ton cryostat measured by two purity monitors over the course of the two month Phase I run [130]. The measurements correspond to different positions in the cryostat, with the red points showing purity measurements at the bottom and blue points near the top. Major external factors affecting the observed LAr purity are shown at the top of the figure. The DUNE requirement of 3 ms lifetime, equivalent to 100 ppt O_2 , is illustrated by the dashed line.

filling one LAPD volume, and when switching between the two pumps installed to extract the liquid for purification. In both cases, good purity is recovered after a few volume exchanges.

The same variations of lifetime on temperature were observed as previously noticed in the MTS and LAPD, suggesting a genuine effect dependent on the ambient conditions. Additionally, during gas circulation, a leak was found and fixed in a seal and, during cold operations, a leak developed in the argon cryo-piping as the dielectric breaks necessary to electrically isolate the cryostat from the building were not leak tight at cryogenic temperatures. All associated 35-ton experience is useful as progress continues to larger and more complicated LAr cryostats.

The success of the 35-ton was exploited by utilising the existing setup for a second run, involving a small-scale DUNE-style detector. This would be the first time a membrane cryostat would facilitate a detector and is the next stage along in prototyping the DUNE far detector.

4.3 35-ton Experiment: Phase II

The first (and to date, only) particle detector housed within a membrane cryostat was the 35-ton Phase II. Following the positive outcomes of the 35-ton Phase I (discussed in Section 4.2),

it is natural to extend operations to include a prototype DUNE detector. The initial aims of the 35-ton Phase II experiment were to develop, build and install a working TPC within the existing cryostat and infrastructure and make measurements of particle interactions induced by cosmic muons whilst demonstrating the required LAr purity is still maintained within an integrated system. The far detector design was heavily constrained by construction, transport, assembly, time and cost requirements, and prototyping is essential to demonstrate the required spatial, time and energy resolution, signal-to-noise performance, detection efficiency and uptime may be achieved.

The operation of the second 35-ton phase will be discussed in detail in this section. An overview of the detector is provided in Section 4.3.1 before the data acquisition from the detector elements is discussed in Section 4.3.2. The custom camera system developed at Sheffield for detecting dielectric breakdown of the LAr is the subject of Section 4.3.3. Finally, the period of data taking is outlined in Section 4.3.4 before outcomes of the project are presented in Section 4.3.5.

4.3.1 The 35-ton Detector

A cutaway view of the 35-ton cryostat showing the detector installed is shown in Figure 4.15. The detector elements are designed to prototype as many features of the DUNE far detector (shown in Figure 3.10) as possible. The readout is performed by four APAs with wrapped induction wires and cold front end electronics (amplifiers and digitisers), which read out multiple drift regions simultaneously. Embedded within the APAs are photon detectors, representing three different design choices, to trigger on scintillation light. The drift field is enabled by cathodes at either end of the TPC. A flange placed on Plate A (referring back to Figure 4.11) facilitates a warm/cold interface through which all electrical signals and the high voltage (HV) feedthrough pass. Surrounding the walls of the cryostat are over 100 scintillation paddles (Cosmic Ray Counters, CRCs) to provide additional triggers from through-going cosmic muons.

The three main detector components, the TPC, photon detection system (PDS) and CRCs, are discussed in the following sections. A photograph of the partially installed detector is shown in Figure 4.16 highlighting most of the detector during construction.

4.3.1.1 TPC

The 35-ton TPC is very similar to the DUNE single phase design introduced in Section 3.3.2.1. It has a modular form, with multiple APAs reading out separate drift volumes, and two drift regions: the ‘long drift region’ of length 2.26 m and the ‘short drift region’, around 0.30 m




Fig. 4.15 The 35-ton detector operated during Phase II of the 35-ton program [118].

long. These were chosen to ensure the longest possible drift region in order to closely resemble the far detector drift distances, whilst ensuring the double-sided readout of the APAs may be tested. Four APAs are used with a very similar design to that depicted in Figure 3.11; each contains two wrapped induction views with a grid and collection plane on each face. The main difference between the APAs tested in the 35-ton and the current DUNE far detector design is the dimensions of the frames and the angle of the induction wires. There are three sizes of 35-ton APA; two tall (204 (height) \times 52 (width) cm) either side of two shorter structures stacked vertically (upper APA dimensions 112 (height) \times 52 (width) cm and lower APA dimensions 92 (height) \times 52 (width) cm). The induction wires are wrapped at an angle of around 45° , as opposed to 37° , with slight differences between the planes to ensure the degeneracy is broken (angles of 45.7° and 44.3° are used). The angle of 45° was initially chosen to optimise the physics reach by providing a high degree of spatial resolution for reconstruction of deposited charge but, following studies of the pattern-recognition performance, and experience with the 35-ton, the angles in the current design were chosen to facilitate a more straight forward disambiguation.

With four APAs and two separate drift regions, there are eight independent drift volumes (DVs), often also referred to as TPCs. These are demonstrated as part of the geometry in Figure 4.17. The coordinate system is defined in this figure; the drift direction is described by the x -coordinate, and the dimension across an APA face, along which the collection planes are spaced, uses the z -coordinate (explaining the denotation of this plane as the Z plane). The y -coordinate is parallel to the orientation of the vertical wires. The origin is at the edge of




Fig. 4.16 Photograph of the partially installed 35-ton detector [136]. The four APAs, with the embedded photon detectors, are visible and the field cage is under construction. Cameras and cold cabling from the Sheffield Camera System, the subject of Section 4.3.3, may be observed in a box, prior to installation, at the bottom of the photo.

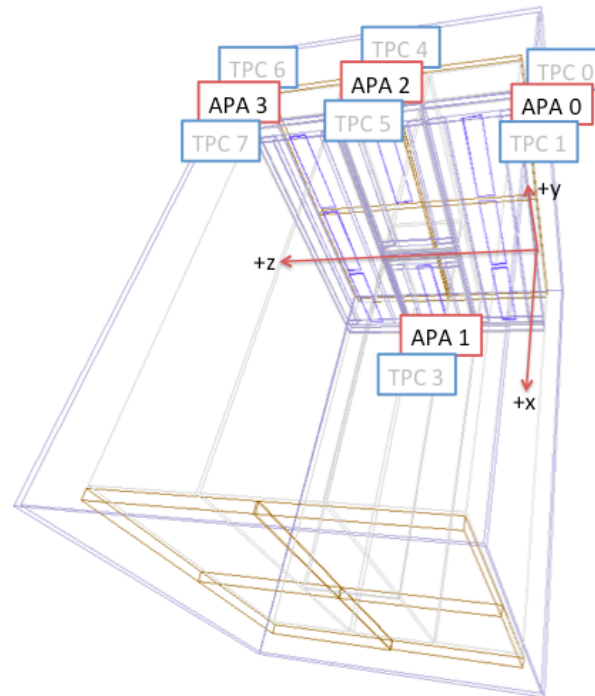


Fig. 4.17 The 35-ton TPC geometry and coordinate system [137]. The blue frames represent the APAs and the orange the CPAs. The eight separate drift volumes resulting from the modular TPC form are labelled TPC0–7.

one of the long APAs and is such that $x = 0$ is at the centre of the APA frames with positive x pointing into the long drift region, $y = 0$ is half way between the two short centre APAs and $z = 0$ is at the right hand side of the APAs when looking from the long drift region with positive z directed across the faces of the APAs.

The cathode and HV feedthroughs are designed to facilitate a voltage of 120 kV, providing the nominal field of 500 V/cm. A field cage constructed using FR4 printed circuit board surrounds the open sides of the TPC to establish the necessary electric field. This was the old LBNE design and has since evolved in the current DUNE outlook; it still enabled a study of the required field within a LArTPC however.

The TPC readout is similar to the DUNE design, with cold preamplifiers, signal shaping and digitisation implemented in ASICs mounted on front end boards at the ends of the APAs. This is the first time a fully cold signal readout has been implemented in a LArTPC experiment and will be discussed in more detail in Section 4.3.2.1.




Fig. 4.18 Photon detector units as implemented within the 35-ton APAs [138].

4.3.1.2 Photon Detectors

Three designs of photon detector were utilised in the 35-ton, none of which are current far detector considerations. These were implemented within APAs in between the wire planes as eight separate units, demonstrated in Figure 4.18 [138]. In the figure, the green detectors are the most similar to the current DUNE design and consist of a plastic bar with wavelength shifter (WLS); the blue and red detectors utilise designs of bundled fibres and plates embedded with WLS fibres respectively.

All detectors were read out by SiPMs which send analogue signals outside the cryostat, via optical cables, for processing. It was following experiences from the 35-ton that the DUNE far detector design evolved (shown in Figure 3.12). In the current plan, the detectors are orthogonal to the 35-ton versions and are inserted after the wire wrapping.

4.3.1.3 External Counters

In order to provide an additional external trigger system, the 35-ton detector is instrumented with CRCs repurposed from the CDF muon upgrade detectors [139]. Most are located on the outer walls of the cryostat, around all four sides and on top of Plate B on the roof. There are additional counters in the ceiling of the building directly above the 35-ton cryostat. The positioning of all the scintillator paddles is shown in Figure 4.19. There are two separate triggers provided by the counters: the ‘telescope trigger’ caused by coincident hits recorded by the counters in the ceiling and those on the cryostat roof, and the ‘horizontal trigger’




Fig. 4.19 The location of the external counters positioned around the outer walls and in the ceiling above the 35-ton cryostat [140].

caused by coincident counter hits on opposite walls of the cryostat (further subcategorised into ‘EW’ and ‘NS’ triggers). The trigger rate for telescope muons is on the order of 60 Hz whilst horizontal muons trigger at a rate of around 2-3 Hz.

4.3.2 Data Acquisition

The process of reading out the data from charge deposits on the anode planes through to the resulting data file on disk which may be utilised for subsequent analysis is the subject of this section. The hardware components, including all readout electronics and processing units, will be briefly described in Section 4.3.2.1. The data formats produced by the readout components are the subject of Section 4.3.2.2 before the software composing the data acquisition (DAQ) system is overviewed in Section 4.3.2.3.

4.3.2.1 Electronics and Readout

The Front End Mother Boards (FEMBs) mounted on the end of the APAs contain two ASICs [141, 142]; the ‘front end ASIC’, which provides signal time-shaping at either $0.5 \mu\text{s}$, $1.0 \mu\text{s}$, $2.0 \mu\text{s}$ or $3.0 \mu\text{s}$ and amplification at gain settings of either 4.7 mV/fC , 7.8 mV/fC , 14 mV/fC or 25 mV/fC ; and the ‘ADC ASIC’ (Analogue-Digital-Conversion) to perform

12-bit digitisation. A configuration of $3 \mu\text{s}$, 14 mV/fC was selected for normal data taking in order to maximise the signal/noise ratio in the collected data. The digitised signals are extracted by Reconfigurable Computing Elements (RCEs) [143], developed at SLAC, which trigger, buffer and format the data and send it downstream to the DAQ framework. The digitising rate is 2 MHz, with the unit of time corresponding to an ADC sample described as a ‘tick’ and equal to 500 ns.

The photon detector information is digitised outside the cryostat by custom built units named ‘SiPM Signal Processors’ (SSPs) [144], built at Argonne National Lab. Each SSP reads out 12 channels and contains a fully differential voltage amplifier and a 14-bit ADC for signal digitisation. Along with the RCE output, the SSPs transmit the processed signals to separate computers running the DAQ software.

The triggers are handled by the Penn Trigger Board (PTB), developed at the University of Pennsylvania, which also manages the CRC readout [145]. A simplified block diagram of the triggering system is shown in Figure 4.20. The PTB is designed to receive sub-system triggers (e.g. from the PDS) and generate global triggers, including internal triggers, and timestamps for the whole detector. It is additionally the front-end for the counter system and handles the reading out of all channels, forming counter ‘hits’ (when a counter has turned on or off) and constructing triggers based on coincidences of these counter hits. It has a backend which is designed to be compatible with the DAQ system, and sends all information downstream to the acquisition software after the on-board logic has formed the relevant data products.

4.3.2.2 35-ton Data Formats

The raw data format employs the concept of a ‘millislice’ as a unified data structure common across all detector subcomponents. An event is a collection of millislices, with one for each of the components being utilised (RCEs, SSPs, PTB). Each contains substructure unique to the detector element; a simplified overview is provided in Figure 4.21.

The raw format for the TPC data is complicated and has many levels of structure. The 2048 TPC channels are read out by 16 FEMBs, each processed by an RCE and represented by a separate millislice. For the TPC data, a millislice contains all the information for 128 channels. This data also has further substructure; a millislice is composed of N ‘microslices’, with each microslice containing M ‘nanoslices’. A nanoslice contains 128 ADC values, representing one tick worth of data for 1/16th of the detector. A microslice thus contains this information for a ‘drift window’ (N ticks) and a millislice a collection of M drift windows. For the normal data running, N was set to 20 and M was 1000.

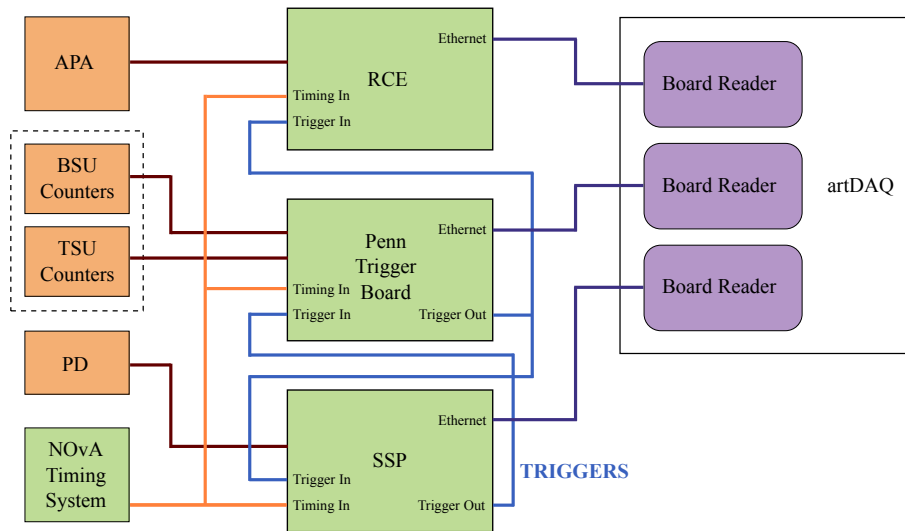


Fig. 4.20 Block diagram showing the triggering system for the 35-ton Phase II. The Penn Trigger Board handles triggers from the CRCs (labelled as TSU and BSU counters) and also from other detector components, such as the PDS. Adapted from [145].

The data structure output from the SSPs and the PTB is a millislice consisting of a series of triggers filled with relevant information when required. In the case of the photon detectors, an ‘SSP trigger’ simply describes the waveform for a given channel as a list of ADC values, one for each tick. A ‘PTB trigger’ is either a counter hit, a trigger or a timestamp and contains a ‘trigger word’ with the type, the timestamp and a variable payload describing relevant further information, such as the channel number or trigger type. Triggers are created and saved in the SSP and PTB millislices either regularly, when self-triggering, or on the occurrence of an external triggering event, identified by the PDS or the CRCs and handled by the PTB as demonstrated in Figure 4.20.

As data are collected, the RCEs continuously create and save microslices to send to the DAQ to form a millislice. These microslices are empty (contain no nanoslices) until a trigger is received, at which point nanoslices are made and saved within each microslice. Additionally, a buffer is in place to save a certain number of full microslices (containing nanoslices) before the microslices containing the trigger. A certain number of full microslices preceding the trigger are also recorded by the RCEs. During normal running, a ‘5 + 1 + 9’ format was employed; five microslices containing nanoslices before the trigger was received, the microslices containing the trigger, and the nine following microslices. It should be further noted that, since the DAQ was designed for continuous data readout, these microslices need not necessarily be within the same millislice: it is possible for the trigger to occur in, for example, Microslice 18 of a certain millislice, resulting in the 15 filled microslices straddling successive millislices. This is demonstrated in Figure 4.22, and results in real ‘physics events’

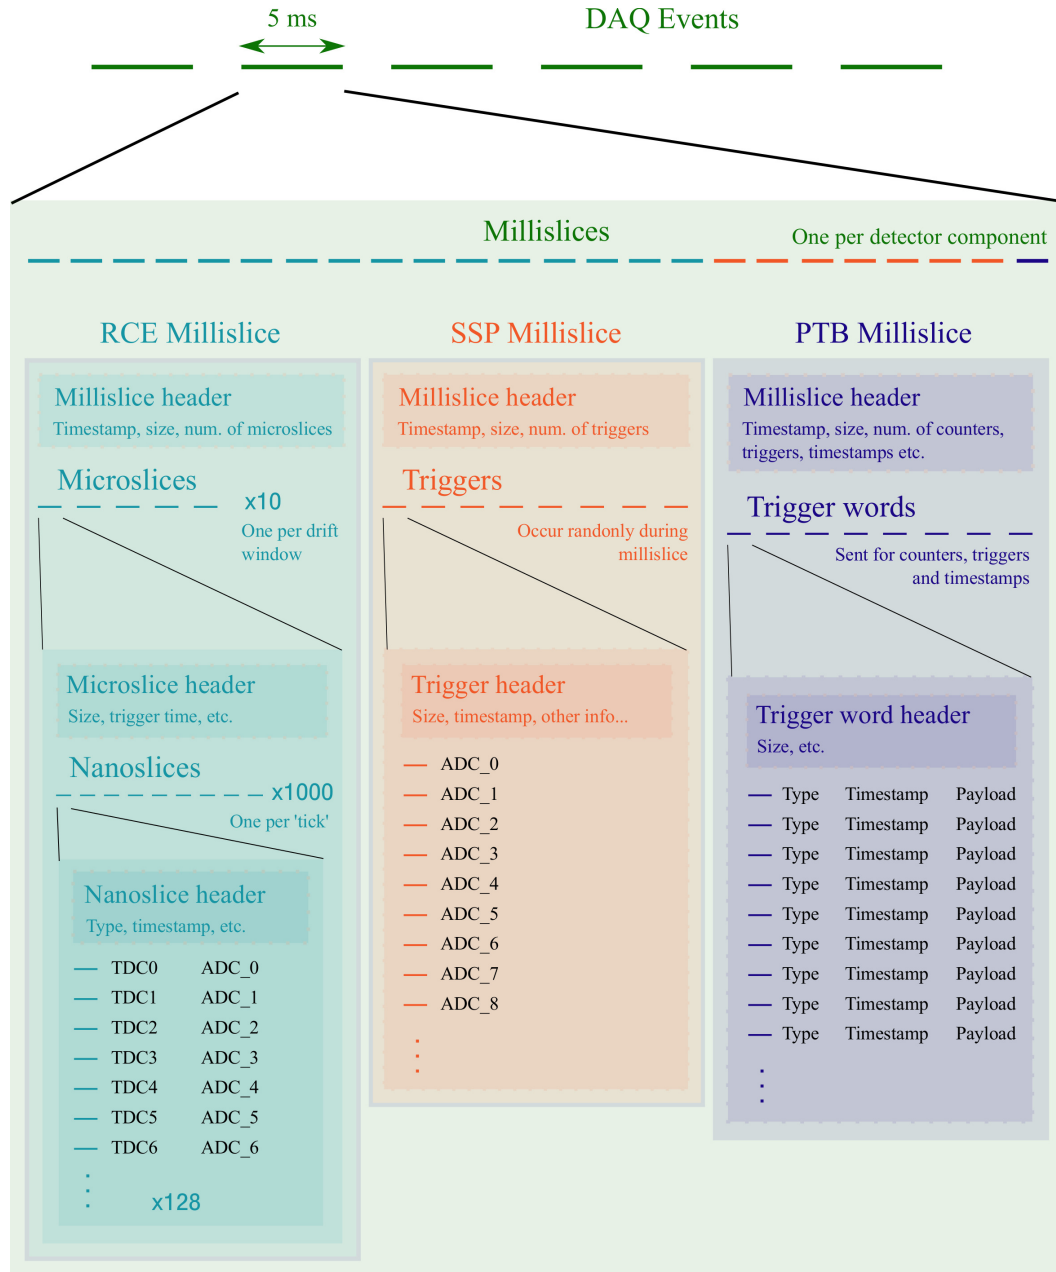


Fig. 4.21 Demonstration of the format used in 35-ton raw data. A ‘DAQ event’ is composed of a single millislice from each component, each containing further substructure unique to the readout elements.

being saved in separate ‘DAQ events’. To account for this, a splitter/stitcher module has been designed to extract the actual triggered events from the raw data and repackage them into a useful event structure. This is the first stage before all offline analysis with the 35-ton data.

4.3.2.3 35-ton DAQ

Experiments at FNAL are migrating to *artdaq* [146], a centrally-maintained data acquisition system built on the *art* [147, 148] framework utilised by all offline software written for experiments hosted at the lab. The DUNE 35-ton experiment was one of the first to use this new software (only LArIAT had previously used it for data taking), and utilised an experiment specific system named *lbne-artdaq*.

A general overview of *lbne-artdaq* is shown in Figure 4.23. Data flow from left to right and pass through components common to most DAQ systems. Closest to the detector components (i.e. the RCEs, SSPs and PTB) are the board readers which take the output from the firmware and send it downstream to the event builders. There exists a board reader for each of the detector components (totalling 24), with each unaware of the existence of the others. It is the job of the event builders to assemble a full ‘event’ from these individual ‘fragments’ passed on from each of the detector elements. An event is complete once composed of a full set of fragments and the event builders will wait to receive them all before sending the data onwards to the aggregators.

There are two aggregators which take the full events but process them in very different ways. All the data passes through the first aggregator, whose function it is to write the output to disk and thus end processing by the DAQ. The second aggregator receives no events but instead has access to the shared memory occupied by the data as it passes through the first aggregator; it is thus designed specifically for the purpose of monitoring and in no way affects the data or the output from the first aggregator. This will be discussed further in Chapter 5.

The 35-ton DAQ was designed to be ‘triggerless’, with the ability to perform continuous readout with a design event rate of 200 Hz. This is an important feature of the far detector DAQ which is required to ensure data may be recorded safely for non-neutrino beam events, such as nucleon decay or a supernova burst. This requires high levels of suppression and buffering to ensure rapid movement of data through the system. In particular, zero suppression for the TPC has been designed such that only ADC values around a window of interest will be kept, vastly reducing the amount of data for the framework to handle. The DAQ was additionally designed to run in various ‘modes’, such as ‘scope mode’, which focusses on a single channel during running, and ‘burst mode’, designed to collect a sample of data from all channels for a given time upon receiving a trigger.

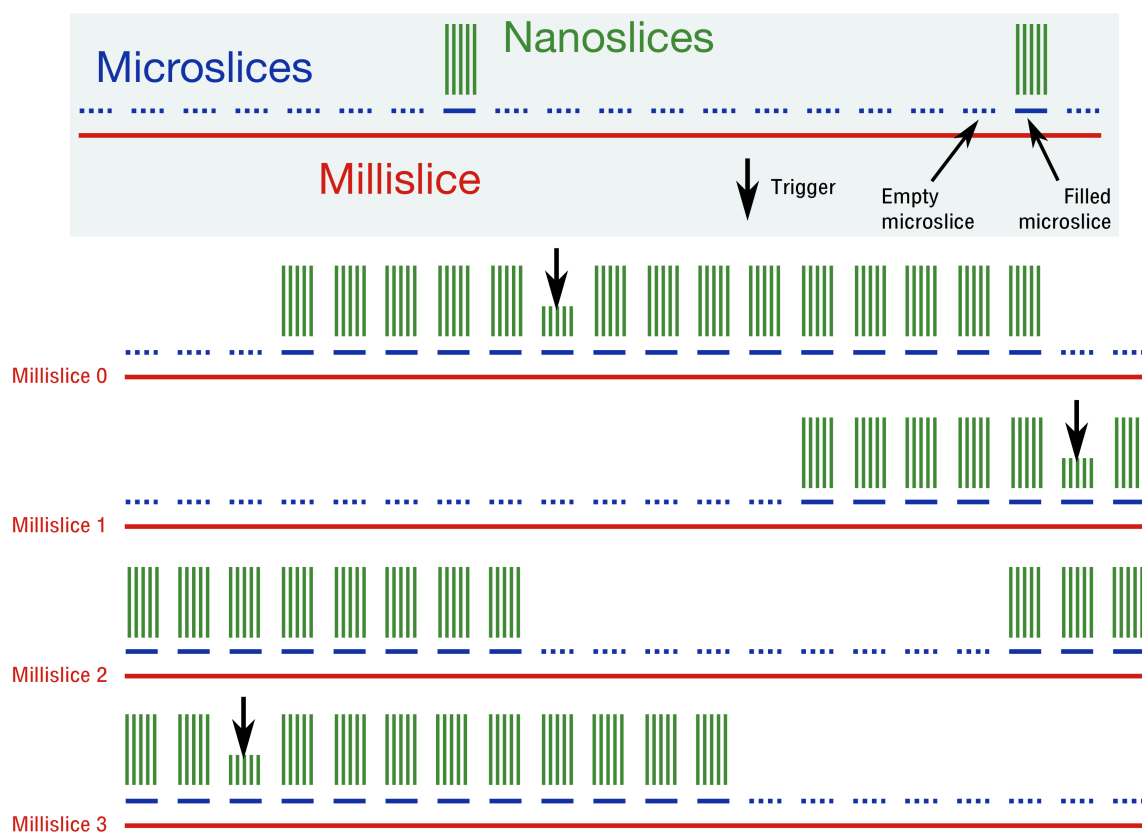


Fig. 4.22 Demonstration of how TPC data are saved when using a DAQ designed for continuous readout. The black arrows represent hypothetical triggers occurring within the duration of a particular millislice. In each case, the five preceding microslices and the nine preceding microslices are filled with nanoslices and saved; all other microslices are saved with no nanoslices since they contain no useful data. An example of such an event is shown occurring in Millislice 0 in the figure. As described in the text, a trigger can cause the useful microslices to straddle consecutive millislices; this is represented in the following millislices in the figure.

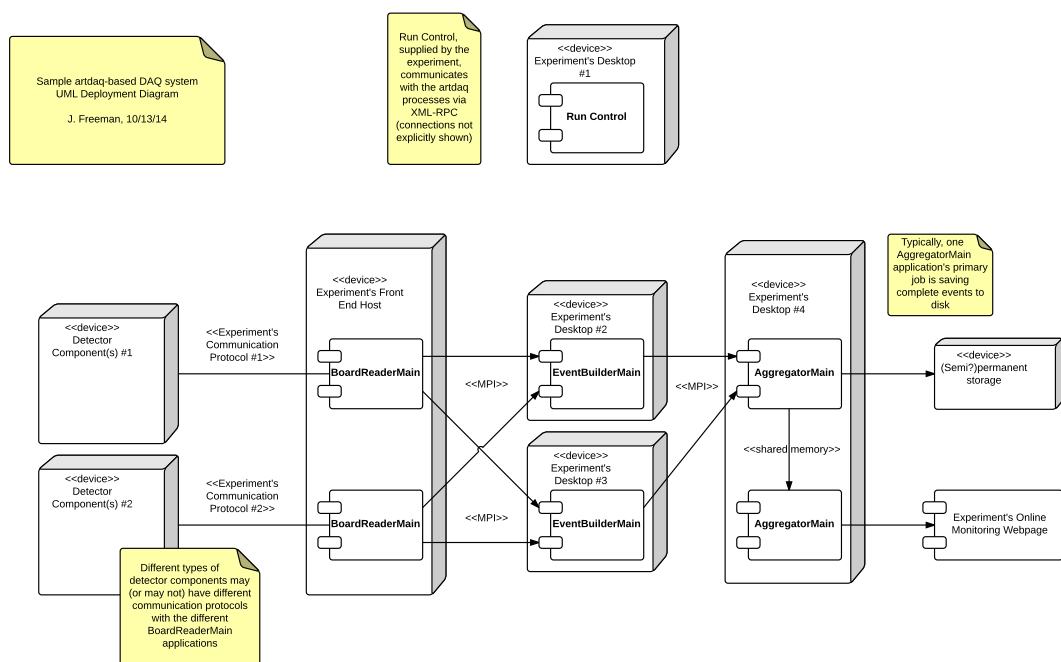


Fig. 4.23 Overview of the *lbne-artdaq* framework used for data acquisition by the DUNE 35-ton experiment [149]. The detector components are shown on the left before the board readers, event builders and aggregators.

4.3.3 The Sheffield Camera System

There are many motivations for developing a camera system which operates at cryogenic temperatures as interest in experiments utilising LAr and LXe (as many dark matter experiments, such as Lux-Zeplin [150], are considering) progresses. These include visual monitoring of the cryostat after sealing, including observing the cooldown and filling with cryogenic liquids, and to monitor HV discharge problems. This latter issue has become cause for concern as LArTPC experiments with very large voltages are being developed; DUNE, for example, will require a cathode HV of -190 kV. Understanding the dielectric properties of LAr is therefore of paramount importance, with recent research suggesting breakdowns occurring at only 40 kV/cm [151]. An additional aim of the 35-ton Phase II experiment was to study the effects of HV and to search for evidence of HV breakdown of the LAr, which may be used to influence the design of future LArTPC experiments in order to mitigate against these effects. This is the primary motivation of the camera system deployed in the 35-ton cryostat [152], designed at the University of Sheffield and described in this section.

The 35-ton was instrumented with eight cameras; six to monitor high-field locations within the cryostat and for detecting visual sparks from HV breakdowns, and two for diagnosis of different cryogenic systems including the cooldown sprayer and the phase separator. The fields of view of each of the cameras are demonstrated in the calibration images shown in Figure 4.24.

4.3.3.1 The Camera System

Previous cameras designed to study cryogenic liquids have either been placed outside the volume or have been maintained in a heated vessel for protection from the cold surroundings. A system which operates directly in cryogenic temperatures is desirable when applying the technology to larger-scale cryostats and for possible use in the detection of secondary scintillation light. Achieving this without an actively heated region in the cryostat is also advantageous to avoid boiling and disturbing the LAr in close proximity. The camera system developed utilised Complementary Metal-Oxide Semiconductor (CMOS) cameras contained within a module alongside a temperature sensor and small resistive heater. This is demonstrated in Figure 4.25.

The cameras are commercially sold as car-reversing cameras and are rated by the manufacturer down to -40°C (233 K). A wide range of cameras were tested and those which consistently performed well in tests whilst at cryogenic temperatures (submerged in liquid nitrogen) were selected. Around half of these were found to reliably endure power cycling when cold (the inconsistency arising from operating the cameras outside of the recommenda-


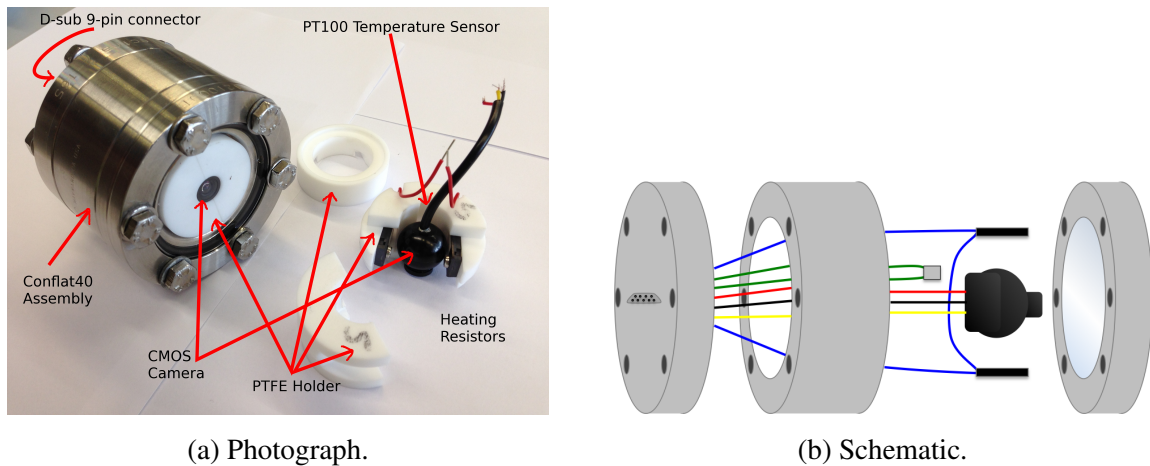


Fig. 4.24 The calibration images for the 8 cameras in the Sheffield Camera System installed in the 35-ton cryostat. Upper (left to right): phase separator, ullage, cathode top right, cathode bottom right. Lower (left to right) cooldown sprayers, cathode top left, cathode bottom left and high voltage feedthrough. The upper images were taken with a halogen light illuminating the cryostat, prior to it being sealed up. The lower images were taken with the LED ring light on, with the cryostat sealed up. All images are left-right inverted due to software. Taken from [152].



(a) Photograph.

(b) Schematic.

Fig. 4.25 An example camera module developed for the 35-ton Sheffield camera system, taken from [152]. Figure 4.25a shows a sealed camera module and the components of such a module. Figure 4.25b demonstrates schematically the composition of a camera module: from left to right a CF40 flange with 9-pin D-sub feedthrough, double sided CF40 flange, PT100 sensor (green wires), camera (red, black and yellow wires), two heating resistors (blue wires) on either side of the camera connected in series, optical viewport on CF40 flange.

tions) and it was these which were included in the modules used in the 35-ton. The heating elements were included as a failsafe mechanism in case the cameras developed a requirement of warmer local temperatures to turn on after sustained periods in the cold.

Each camera contains 712×486 pixels and has a roller shutter rate of 50 frames per second. Their resolution at 10 mm was found to be (2.0 ± 0.5) mm at room temperature and (1.5 ± 0.5) mm at 77 K, with the improvement at lower temperatures due to a higher refractive index of LN_2 resulting in the light becoming less diffuse. The minimum measurable light pulse width the cameras could trigger on, in both the warm and the cold, was observed to be 20 ns. One notable change when operating the cameras at cryogenic temperatures was the chrominance output of the video signal. The usual colour signal is observed as monochromatic when in the cold, possibly due to partial failures on the on-board encoding electronics.

Before installation, the response of the cameras to sparks was characterised by applying a HV across a printed circuit board (PCB) in LAr until breakdown was observed. The discharge was between 40 and 60 ms and the cameras showed localised sparks persisting over multiple frames of exposure. The trigger system, which relies on a percentage change in the number of different pixels between successive frames, was also able to successfully detect and automatically record on occurrence of the sparks.

4.3.3.2 Operation and Outcomes of 35-ton Camera System

The camera modules were mounted on the existing piping from the cryogenic system within the 35-ton. An example is shown in the photograph in Figure 4.26. Data acquisition, operation and control was performed using a rack-based system containing a power supply, a temperature sensor reader, DAQ and computer control system. Full details of the entire arrangement and all the interconnects are available in Figure 4.27.

The cameras were characterised in room temperature following installation and the software trigger tested on the Xe flash light from the purity monitors (described in Section 4.1.1.2). The system ran continuously throughout the 10 weeks of the 35-ton Phase II cooldown. It was heavily utilised during cooldown and filling to monitor the inside of the cryostat and observe the rising liquid level (an excellent video of the LAr when level with one camera module is available at Reference [153]). The entire system was power cycled successfully three times during TPC debugging and following the FNAL site wide power outage on 4th March 2016. The downtime ranged from 30 minutes to 9 days, with the cameras turning on without assistance each time.

The picture quality was observed to degrade noticeably over time, demonstrated in Figure 4.28, with significant variation between different camera modules. When in darkness,



Fig. 4.26 Two camera modules mounted on cryo piping in the 35-ton cryostat. Taken from [152].

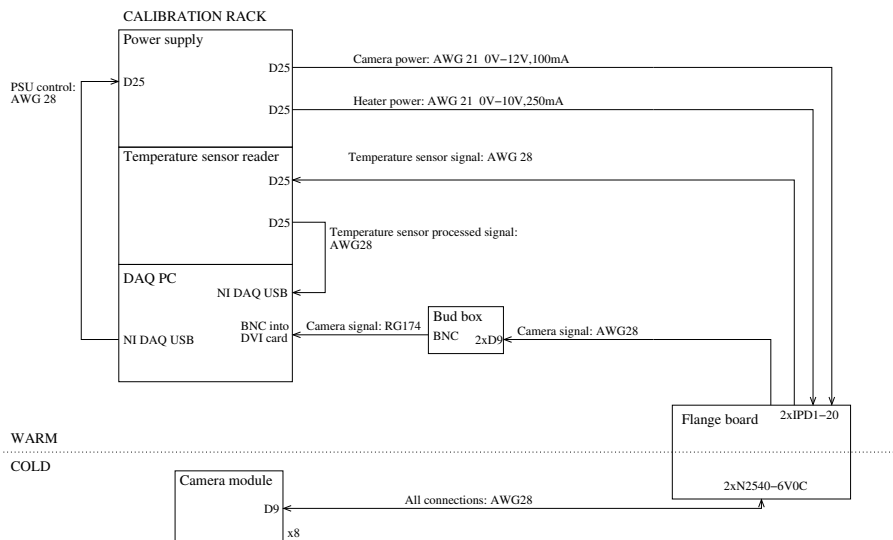


Fig. 4.27 Full system block diagram for the camera modules in the DUNE 35-ton prototype. Taken from [152].




Fig. 4.28 The variation in picture quality degradation is illustrated by the changes in Camera 1 (upper) and Camera 4 (lower) over time. Left: prior to cooldown, centre: immediately post-cooldown, right: after 10 weeks submerged in LAr. The field of view changes due to the change in refractive index. Note that these are full colour images with no post-processing. Taken from [152].

a greater number of saturated or noisy pixels is observed across the cameras and when illuminated by the LED ring, the noise increase is noticeable with a decreased colour depth. This is likely due to signal transmission length, power cycling and prolonged exposure to the cold.

Two suspected HV breakdowns occurred during normal operations at 60 kV but the system was unoperational as a result of the power outage during both. Following the end of running, when testing the HV at 135 kV, four breakdowns occurred with three detected and triggered on by the camera system. However, the location of the spark could not be determined clearly from the recorded video. This could be due to either the spark occurring outside the cryostat or the field of view of the cameras, an insufficient intensity or duration of the flash or the degradation in picture quality being such that the efficiency and sensitivity of the triggering system were compromised.

The camera system was shown to be successful and a useful aid in 35-ton operations. Despite not showing HV breakdowns clearly, the modules remained operational during the 35-ton Phase II run and were valuable for monitoring purposes. They were shown to trigger successfully on a test bench so it seems reasonable to conclude their inability to do so within the LArTPC was solely due to the degradation in picture quality, which must be improved if such a system were to be used in future LAr experiments.

4.3.4 Phase II Run

Following a long period of testing the detector components at FNAL, installation of the TPC and field cage was carried out in October 2015. This was followed by the final parts of the system, such as the long drift region cathode, the purity monitors, HV feedthrough and cameras, in November 2015. Following the Fermilab readiness clearance, operations began in December 2015. This involved piston purging both LAPD and the 35-ton, filling LAPD with LAr delivered from the suppliers, cooling down the 35-ton cryostat and finally transferring the liquid argon from LAPD into the 35-ton. This was completed by the end of January 2016.

The 35-ton Phase II run officially started on 11th February 2016 upon the final liquid transfer into the cryostat and the starting of the pumps and recirculation of the LAr through the filtration system. A week later, the HV on the cathode was ramped up to half nominal value: 60 kV, providing a drift field of 250 V/cm. The intention was to ensure a sufficient amount of collected data was on disk before proceeding with increasing the HV up to the design voltage of 120 kV (500 V/cm) and even up to the maximum of 135 kV.

The start of the run was dedicated to many noise tests; it was immediately clear the noise on the TPC channels was much larger than anticipated even after the testing from the previous summer. These tests involved studying each of the FEMBs separately and considering effects from other non-TPC detector elements by removing power from all systems in the cryostat before reintroducing components iteratively. An additional ‘high noise state’ was also identified, corresponding to a very high oscillatory noise level instantaneously appearing on all channels simultaneously and remaining for up to hours at a time. The noise problems in the 35-ton Phase II will be discussed in more detail in Section 4.3.5.

This commissioning time was also important as the stability of the DAQ was improved. Near the beginning of data taking, it was uncommon for the DAQ to run for more than a few minutes with even a small subsection of components (RCEs, SSPs, PTB), with issues such as data throughput, disk writing speed and hardware interface issues contributing to a very unstable system. In the months of installation and commissioning, the DAQ was the subject of much attention and progress on improving the framework progressed in parallel with the final installation, LAr filling and noise hunting.

Following the completion of the designated noise runs and the stabilising of the DAQ, the focus was on collecting as much data as possible before raising the HV, with the plan to run for at least a week at 90 kV and 120 kV respectively. However, the run was unfortunately cut short in the early hours of the morning of 19th March 2016 when a tube, part of the system which was introducing GAr from LAPD to the 35-ton purification network in order to maintain the LAr level, sheared and facilitated the introduction of air directly into the



Fig. 4.29 The broken pipe, originally part of the framework introducing gaseous argon from LAPD into the 35-ton to maintain LAr levels, which resulted in the poisoning of the whole LAr volume by allowing the introduction of air into the system.

filters. Within a few minutes, faster than it would have been possible to respond even if this incident had not occurred at 3 a.m., the filters were saturated and the entire volume of LAr in the 35-ton was poisoned. The offending pipe break is shown in Figure 4.29. This incident effectively concluded the data collection prematurely and meant the design HV could not be tested in good quality LAr and no data could be taken at nominal drift field.

The run is summarised in Figure 4.30, showing the LAr purity as a function of time and notable incidents. The bulk of collected data was either side of a site-wide FNAL power outage on 4th March 2016, after which it took a few days to recover the LAr purity. After recuperating from this incident, an issue with the LN₂ values resulted in a cooling failure and the boiling off of a large portion of the LAr in the cryostat. The pipe break occurred shortly after rectifying this issue. These issues plagued the final few weeks of data taking, evident from the corresponding LAr purity in the detector.

Most of the data taken were triggered using the horizontal muon trigger. In the last week of running, the telescope trigger was deployed, with a large prescaling due to the high rate of cosmic muons, and the PDS was also used to trigger data taking. Both systems appeared to work as intended but thorough testing proved impossible due to the temporal proximity to the unforeseen termination of run. Throughout data taking, the DAQ recorded data to disk at a rate of 1 Hz. Because of the electronics noise and small signal/noise ratio, tests of data taking using zero suppression were unable to be performed.




Fig. 4.30 The data taking period of the 35-ton Phase II experiment. The electron lifetime measured by the two long PrMs in the cryostat is shown as a function of time, with the horizontal axis covering the period 11th February – 19th March 2016. The numbers within the green boxes represent the amount of data taken, in days, with the drift field of 250 V/cm present. The major incidents which affected the LAr purity are shown on the figure.

Overall, the run provided 22 days of high quality (good LAr purity, high stable voltage, stable DAQ) data, albeit with much higher noise than anticipated. An example electromagnetic shower observed in the data with strong signals in all planes is depicted in Figure 4.31. The noise problems have resulted in limitations to the analyses possible with the 35-ton data and focus has shifted to studies utilising datasets unique to the 35-ton. Some such analyses are the subject of Chapter 7.

4.3.5 Outcomes of Phase II

The 35-ton Phase II collaboration successfully designed, constructed, installed and ran a small-scale DUNE-style LArTPC and collected data whilst maintaining a good LAr purity, with electron lifetimes consistently reported above the DUNE requirement of 3 ms. This is the first time a detector has been operated within a membrane cryostat and the integrated system has been strongly validated. The complete process has been instructive and a great many lessons have been learned alongside the successes of the project [154].

This section will review all these outcomes and discuss how the experience will influence the DUNE programme as it progresses towards the first far detector module. In general the majority of subsystems achieved or superseded expectations and, following the 35-ton




Fig. 4.31 Event display depicting the charge deposited by an electromagnetic shower during the 35-ton Phase II run. The three views are, from the top down, the collection plane and the V and U induction planes. Each shows the wire number on the horizontal axis and time, measured in units of tick ($\equiv 500$ ns) on the vertical axis. Charge is represented by the colour scale on the z -axis. The shower is clearly visible in all three planes and demonstrates the functionality of the 35-ton detector.

Phase II experience, there is no reason for reservation over ProtoDUNE as rapid development continues to be made.

4.3.5.1 Cryostat and TPC

The cryostat and most TPC components behaved as expected and resulted in no unexpected functionality. When filled with GAr, before the introduction of LAr, the cryostat was leak tested. When this was performed in Phase I a few issues were identified and had to be addressed; there were no complications during Phase II commissioning however. The pumps were not tested between phases and required a large current to begin their operation with the cryostat already filled with LAr; this demonstrates how vital it is to assess all detector components before commissioning. Other than the failing in the cooling system, all cryogenics performed excellently. Since this incident occurred not long after the power outage, the alarm system had not been correctly brought back online, resulting in an avoidably large loss of LAr. These are two of many examples of lessons learned from the 35-ton.

The HV and drift field presented no issues during the course of the run. No confirmed breakdowns were observed at 60 kV but testing in clean LAr at 120 kV was not possible. Although a voltage of 135 kV was attained and held for multiple days in contaminated argon, the impurities are presumed to alter the dielectric properties of the material and therefore complete validation remains unproven.

Results from the purity monitors and temperature sensors suggest a stratification along the height of the LAr volume within the cryostat, similar to observations made during the Phase I run. The cause of this is likely due to returning LAr from the purification system being cooled below the ‘bulk temperature’ by the phase separator and reentering near the bottom of the cryostat, resulting in reduced convection and poor mixing. Resolutions, such as returning warmer LAr to the main volume, are being considered for future LArTPCs in an attempt to mitigate these effects and ensure a good, isotropic purity.

The TPC electronics were the largest source of shortfalls in the experiment and have significantly compromised the utility of the data. During warm tests over summer 2015, it was evident the intrinsic noise levels in the ASIC electronics were higher than anticipated and an additional issue with the ADC ASIC was observed. The digitisers are affected by bit-level corruption whereby the six least-significant bits (LSB) or most-significant bits (MSB) are erroneously reported as either 0x0 or 0x3F at a rate, between 20% and 80%, which is strongly dependent on the proximity of the true value to these ‘sticky’ codes, and also on the temperature, the input current and the channel. Along with this ‘stuck code’ problem are further issues with ‘stuck bits’, where a particular bit is never set or cleared. These issues may be somewhat mitigated in software but work is ongoing to rectify concerns before use of




Fig. 4.32 Difference between optical hit peak times and muon counter trigger times for photon detector 3 in the 35-ton photon detection system. The binning reflects the digitisation time of the photon detector electronics. Taken from [138].




Fig. 4.33 Average Optical Hit Amplitude per Event vs. Counter Pair Positions for the 35-ton photon detection system. Error bars are statistical errors on mean hit amplitudes per bin. Taken from [138].

the ADC ASIC in ProtoDUNE. The multiple problems with coherent and incoherent noise which characterise the 35-ton dataset are discussed further in Section 4.3.5.4.

4.3.5.2 Triggering Systems: Photon Detectors and Muon Counters

The photon detectors and muon counters also achieved expectations. Although the CRCs are unnecessary for the far detector, they proved critical to the success of the 35-ton. The vast majority of data was recorded whilst triggering on through-going muons and, as will be discussed further in Chapter 7, all worthwhile analyses rely heavily on counter information.

The PDS was shown to successfully record data in both externally triggered (when using the CRCs or an internal trigger from the PTB) and self-triggered modes, where the PDS sends a trigger to the PTB upon receiving a sufficient level of scintillation light. The timing resolution of the detectors was shown to be better than 100 ns with respect to the counter timing, as shown in Figure 4.32, with signals as low as a single p.e. detected. A characteristic length of light in LAr may be determined by considering the signal size of scintillation flashes, using counter trigger information to determine how far from the detectors the interaction occurred. This is demonstrated in Figure 4.33 and yields a measurement of 155 ± 28 cm.

Given the noise problems in the TPC data, it was not possible to do joint analyses using the photon detectors as planned. The system performed well however and validated the concept of using WLS bars with SiPM readout as opposed to PMTs for the DUNE far detector design.

4.3.5.3 DAQ and Computing

The DAQ was remarkably consistent throughout data taking following the stabilisation period. All components could be operated simultaneously with data written to disk at a steady rate, successfully demonstrating continuous readout of the detector systems. In total, $\sim 500\text{k}$ cosmics were recorded during the 35-ton Phase II data taking, with an impressive capacity on disk of ~ 30 TB.

It proved imperative to monitor the data during running as detector issues spontaneously arose on a regular basis. The large volume of data was an additional issue and finding an optimum output file size, balancing number of data files on disk with size of each file and potential for data loss upon a DAQ crash, occupied a sizeable amount of commissioning time. Additionally, a potentially disastrous failure in the alarm system for one of the computing racks resulted in serious overheating and the loss of all the machines which were running most of the online processes.

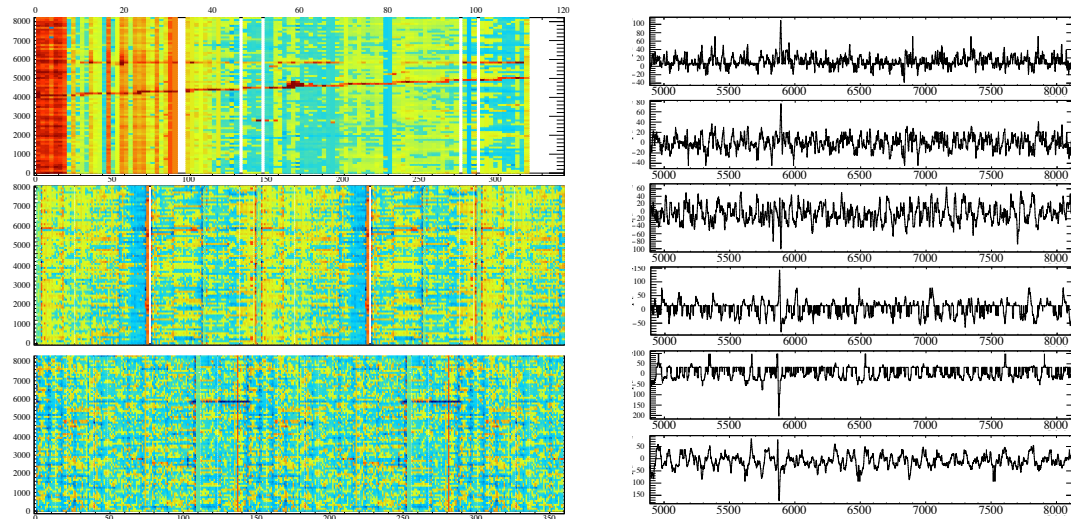
Data from the cold electronics were shown to be processed by the RCEs at a rate of 1 Gb/s but a bottle neck in the framework restricted disk writing to 60 MB/s, resulting in an enforced reduced data flow through the system. An event rate of 1 Hz was utilised during the run, much smaller than the design rate of 200 Hz. This could have been improved by employing zero suppression in the TPC data but this was unable to be tested as planned in the 35-ton. The event rate requires improvement before the far detector DAQ but work is underway and the experience with the 35-ton will be taken forward with most of the existing framework under development for use in ProtoDUNE.

4.3.5.4 Noise Issues

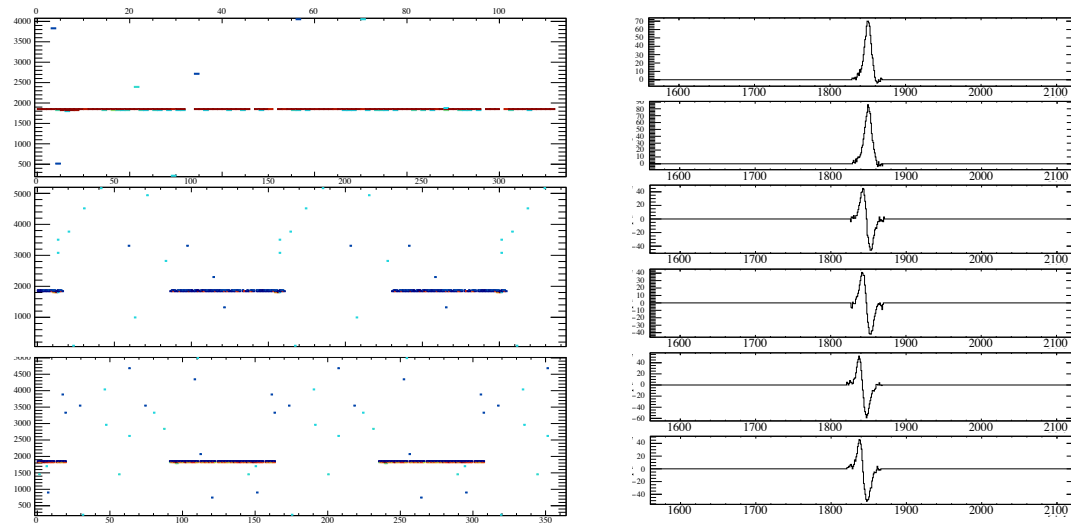
An example muon track observed in the 35-ton data, along with typical waveforms recorded on the anode wires, is shown alongside an analogous muon track and detector response from simulation in Figure 4.34. The difference between the simulated and observed data is immediately apparent, with the stuck code and noise problems evident in each of the example waveforms recorded on the various detector channels. Understanding this noise so it may be alleviated in future experiments is the subject of this section.

There were multiple sources of noise in the 35-ton detector with distinct ‘modes’: the ‘normal noise state’ (which still contains numerous issues) and the ‘high noise state’ [155]. The frequency bands of noise in each state is demonstrated in Figure 4.35.

The normal noise state is characterised by 11 kHz and 100 kHz bands. The phase of the 11 kHz noise appears to alter every 64 channels, corresponding to the blocks of channels read out by ASICs sharing a common voltage regulator (four 16 channel ASICs). The correlation



(a) 35-ton data.



(b) 35-ton simulation.

Fig. 4.34 Comparison between example muons tracks observed in 35-ton data and simulation, along with example waveforms for randomly selected channels. The event displays each show the time against wire for the Z, V, U planes from the top downwards. Example waveforms, illustrated as ADC values as a function of tick, observed on two channels from each plane are shown alongside the two dimensional view

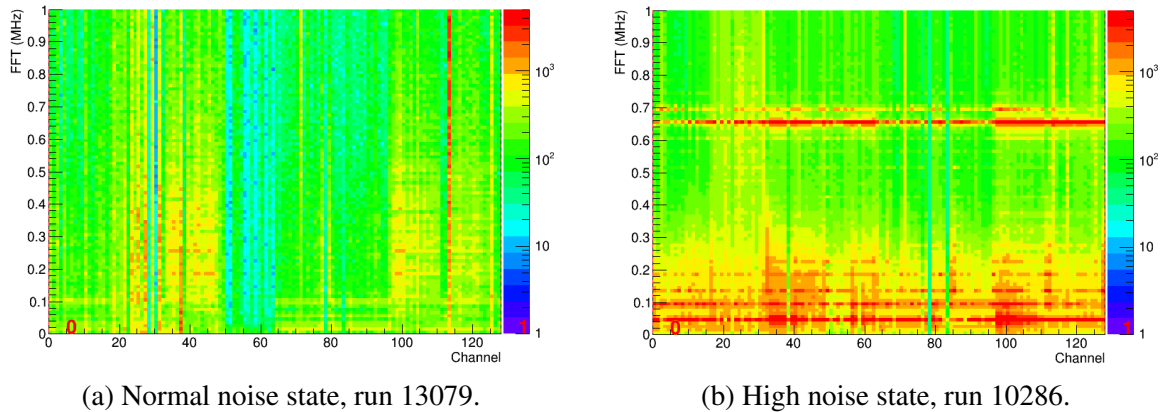


Fig. 4.35 FFT of ADC values for RCE00 for two different noise states. During the normal noise state, the noise band at 11 kHz (faintly visible at 0.011 MHz in Figure 4.35a) is present across all channels in the detector and a lot of channels also see 100 kHz frequency noise. The high noise state manifests across all channels in the detector as multiple frequency bands and render any collected data useless when present.

between the waveforms observed on the channels maintained by the same regulator is evident in the plot shown in Figure 4.36. This was shown to be removed following the run by the addition of a 1 Ω resistor in series, effectively forming a low pass filter, and can be removed crudely in software using a coherent noise subtractor. A similar phase shift in the 100 kHz noise is observed at the boundaries between FEMBs, which are each maintained separately by the low voltage power supply. Again following the completion of the run, close inspection of the cabling found a short between the supply return line for the FE ASICs and the chassis ground for the supply. Correcting this removed all noise sources and, along with the correlated component from the voltage regulators, explained all prominent noise frequencies in the normal mode.

The high noise state was entirely unanticipated but was characterised by several features: a very high noise level is observed without saturating the ASICs; multiple frequency bands, most under 300 kHz, are observed simultaneously across all channels in the detector; these frequency bands are consistent for the duration of the high noise state but change each time the state is entered; the frequencies are also observed on a spectrum analyser connected to an APA grid plane; the current draw of the ASICs is observed to drop when in the high noise state. Furthermore, the high noise state was not observed when the cryostat was at room temperature and so could not be investigated subsequent to the end of the run. It has been understood as a collective oscillation of all detector components which is spontaneously entered, roughly every few hours, during running. Often, after a time period on the order of an hour, the system may egress from the state; it was also noted that power cycling the front end




Fig. 4.36 The correlation between waveforms recorded on different channel combinations for all 2048 35-ton channels.

ASICs may also return the detector to the normal noise state. The noise investigations after data taking were unable to definitively identify the conditions of the abnormality but have offered suggestions as to the likely causes. The frequency of the oscillations, and the inability to induce the state in the warm system, argues strongly against external influences. The source cannot be the anode wires as this would saturate the front end electronics and, given the necessary power required to sustain the oscillations on the grid plane, the only candidate is the low voltage power supply. The main difference between the 35-ton and MicroBooNE, which uses the same supplies and has not observed similar problems, is the length of cabling used in the 35-ton being around 10 times greater. This may turn the negative feedback in a remote sense system into a positive feedback loop, causing the circuit to search for the correct voltage settings by overshooting and subsequently undershooting (i.e. oscillating) due to the round trip cable delay being longer than the circuit response time. The strong frequency bands at 650 kHz, which are always present across all channels whenever the high noise state is entered, unlike the other frequencies, is likely due to the oscillating cable acting as a cable resonator. During the run, it was observed that APA1 (the short, bottom centre, APA) was most prone to these issues and was actually left unpowered during much of the data taking. This is explained by considering the most likely coupling of the oscillating power cable to the detector is via the FE electronics for this APA (the only one where these are at the bottom). These oscillations may then be transferred to the grid plane and subsequently

to the cathode on the short drift side, which couples to the other APAs in the detector. The decreased capacitance of the cable in air compared to that when submerged in LAr explains why this state could not be induced following the end of operations.

Finally, it is observed that the minimum noise in the detector is higher than in Micro-BooNE. Although the induction wires are much longer, there is still an increase greater than could be accounted for by the larger capacitance of the wires. The noise experts suggest there may be a common mode noise on the supply line which may intensify the overall noise levels without inducing the high noise state; this would enter via the cathode, then the grid planes and followed by the induction wires and would explain why these planes see more noise than the collection view.

The noise issues encountered in the 35-ton, though unexpected, have been critical to understanding the issues which may be present in large scale LArTPCs and would be seriously detrimental to the DUNE project if encountered in the far detector. Every effort has been made to understand the issues with the 35-ton and ensure the eventual success of the experiment.

4.4 Summary

As work progresses towards the final, full-scale, DUNE experiment, the research and development performed by the test stands and prototypes discussed in this chapter represent crucial understanding which will influence design, schedule and planning of the far detector. Along with the important comprehension of critical or unexpected phenomena, the experience gained through the construction and operation of prototype experiments contributes to a better understanding of the technology and matured expertise which may be taken forward towards the full-scale DUNE project.

The Materials Test Stand and the Liquid Argon Purity Demonstrator pioneered a new, necessary method for achieving extreme LAr purity in a non-evacuatable cryostat and the 35-ton vindicated the design choices of the DUNE cryostats which greatly simplifies the associated engineering requirements. Both of these results have had a profound impact on the development of LArTPCs, with a particular significance for the eventual success of DUNE.

The 35-ton Phase II experience, while unable to deliver the high quality data anticipated for the purpose of physics analyses, was invaluable to the DUNE strategy. A large number of ‘lessons learnt’ are already influencing the ProtoDUNE experiments and even the far detector considerations. It was a significantly important phase of the overall plan and, as a prototype experiment, can be considered a notable success. The 35-ton Phase II experiment

will be discussed again in detail with reference to the Online Monitoring and Event Display framework in Chapter 5 and for the purposes of data analysis in Chapter 7.

Chapter 5

Online Monitoring and Event Displays for the 35-ton Experiment

Monitoring of the data collected during the running of an experiment is imperative to ensure a high quality is maintained. Such monitoring is often provided in real-time (‘online monitoring’), summarising the data from the current run, or in near real-time (‘nearline monitoring’), summarising data over runs from typically the previous day, week or month to represent the longer term fluctuations in the data quality. An event display, designed to illustrate physics events as they occur in the detector, is another desirable feature that is particularly useful during data collection. The system developed to provide online feedback, including a basic event display, for the 35-ton Phase II data taking period, described in detail in Chapter 4, is the subject of this present chapter.

The framework was designed to be flexible and provide prompt feedback for those operating the experiment; it was thus included as part of the DAQ system, discussed in Section 4.3.2.3. The monitoring framework itself is the subject of Section 5.1, with its two functions, data quality monitoring and producing online event displays, presented in Sections 5.2 and 5.3 respectively. Finally, the web interface developed to allow synchronisation of this monitoring data to a dedicated web page for ease of access is briefly described in Section 5.4.

5.1 The Online Monitoring Framework

The framework developed for the monitoring system had the following design goals:

- to be able to analyse the data read out of memory in its raw ‘DAQ format’;

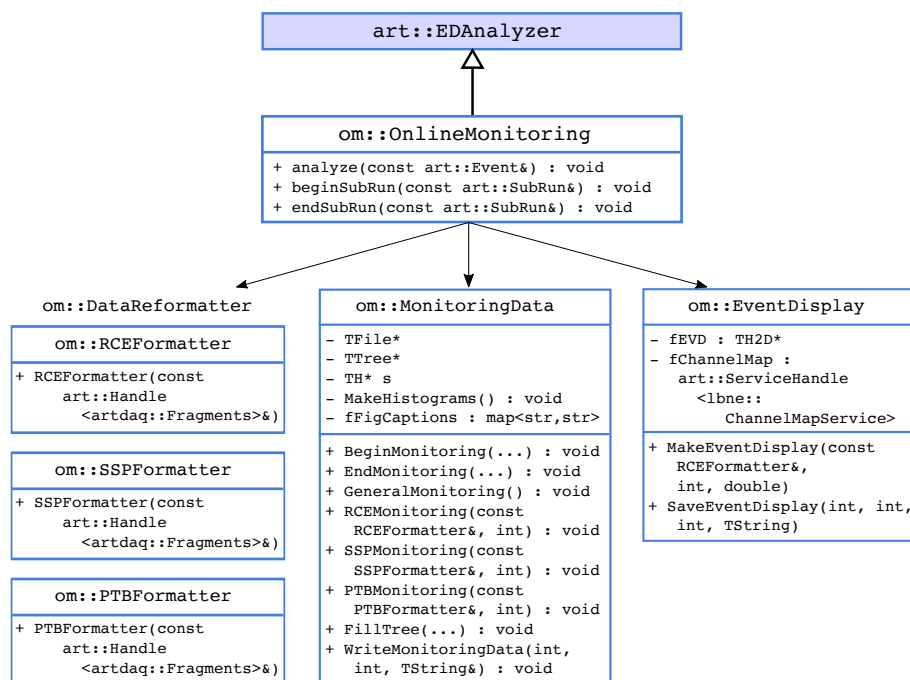


Fig. 5.1 The software framework designed and built for online monitoring during the 35-ton Phase II run.

- to be computationally efficient to allow for processing at the event rate (data taking rate);
- to provide the flexibility for further monitoring plots to be added with ease;
- to allow for use of an online event display to provide comprehensible images of the raw data.

In general, the developed system succeeded in all these goals and provided invaluable information, becoming an integral tool in the commissioning and the data taking during the 35-ton Phase II run.

5.1.1 Design of the Monitoring Framework

The setup consists of a central ‘module’, `OnlineMonitoring_module.cc`, which is configured within the *art* framework through its base class. An illustration of the full monitoring framework is shown in Figure 5.1. The `OnlineMonitoring` class manages the running of the system and owns instances of further classes each designed for a specific purpose, controlling the data flow by calling the relevant methods when required. Once an event has been obtained,

the data for each component is processed and repackaged into RCEFormatter, SSPFormatter and PTBFormatter objects. The purposes of this method are

- to provide an interface between the raw data and the methods which analyse the data. This is important as it provides a single point of maintenance for when formats change and allows for various ‘DAQ modes’ to use the same analysis code;
- to separate interaction with the DAQ from the handling of output data objects;
- to facilitate random access of the data for more detailed analysis which would not be possible if just processing linearly.

The main drawback to performing this step is it requires all the data to be held in memory until the end of the event and represents basically the same information as initially present. However, it was decided the advantages were worth the required compromises in memory usage and no problems were apparent during the course of the run except when operating at the very limits of DAQ capabilities.

These reformatted data objects are then passed to the methods in the MonitoringData class for analysis. This class owns all of the data products which are output from the monitoring (e.g. histograms, graphs, trees and files) and deals with their filling and writing out when required. This is discussed further in Section 5.2.

The event display is handled by its own dedicated class, EventDisplay, which has methods for making the displays and saving them as an image in the correct place when required. It is designed to accept the reformatted RCE object and presents the data in as meaningful a way as possible; this is detailed fully in Section 5.3.

5.1.2 Interface with the DAQ Framework

The 35-ton DAQ, previously discussed in Section 4.3.2.3, was based on the *lbne-artdaq* framework illustrated in Figure 4.23. This system has support for running online monitoring embedded into its design philosophy, with the Aggregator2 process allowed to access data from shared memory as it is managed by Aggregator1. The controlling monitoring module, discussed in Section 5.1.1, may be configured to run within this second aggregator process and thus receive events in real-time as they pass through the data acquisition system.

The events are passed to the Aggregator2 process by the framework when resources are available; if this is not possible then the event is simply skipped. This behaviour does not affect the processing of the data through Aggregator1 and any events missed by the monitoring will still be saved to disk. The issues arise primarily when the monitoring runs slower than the data taking rate (i.e. when producing monitoring information for an event

takes longer than the length of the event itself) and were largely inevitable due to the number of required plots and the computational resources necessary for tasks such as FFTs and event displays. As most monitoring plots, such as TPC noise, require only a few events, they are mostly unaffected; however, there are implications when calculating rates and similar quantities. During normal running, as many as half of the events may be missed by the monitoring, depending on the detail of the plots being produced. Using multiple threads detached from the main processes was considered, particularly when making events displays, as way to increase the event exposure but, due to the potential computing issues which may arise, it was decided not to implement this for the purposes of a short prototype run.

Each of the DAQ processes run on a machine on the private DAQ network and are configured as normal within *art*, using *fhicl*. Two nodes on the main FNAL network (lbne-gateway01/02) provide access to these private machines, of which there are 7 (lbnedaq1-7), and contain all scripts necessary to setup, configure and run the DAQ via a command line interface.

5.1.3 Writing the Monitoring Data

The data objects are newly created for each subrun and are written out at three points during data taking:

- an initial write out N seconds after the start of the subrun;
- at frequent intervals during the subrun, every M seconds;
- at the end of the subrun.

The parameters N and M are configurable and were set to 30 and 500 respectively for normal data taking. The data products are only cleared at the end of a subrun, so any intermediate writing out of data simply refreshes the current plots.

The event displays are computationally expensive to make and so were only created once per subrun during normal running. However, since a subrun was automatically stopped, and a new one started, by the DAQ once the output file had reached 5 GB in size, and (since zero suppression was not utilised at any point during the run) this occurred on average every four minutes, a new event display was made relatively frequently.

All the output data were saved on a shared disk on the gateway DAQ machines for further use. This is discussed in Section 5.4 below.

5.2 Data Quality Monitoring

The overarching aim of the online monitoring system was to provide direct feedback to the experimental operators with information about the status of the data taking and the quality of the data. This is vital for various different aspects of data taking, for example

- ensuring all detector components being used in the current run are receiving and processing data;
- noting the TPC readout has entered the ‘high noise state’ and acting accordingly;
- checking the trigger rates from the external cosmic muon counters are feasible.

The monitoring was diagonalised in a similar way to the DAQ readout with data from the TPC, photon detector and external counters processed separately.

5.2.1 TPC Monitoring

Monitoring of the TPC data involved mainly considering various distributions of the ADC values provided by the front-end boards, separated by channel, board and APA. The mean and RMS of the ADC values for a given channel provides information such as the measured pedestal and the level of noise being read out. The uncorrelated component of the noise can be monitored using the concept of ‘DNoise’; this considers the difference in ADC value between two neighbouring channels at a given readout time and represents the level of noise which would be impossible to remove by the use of coherent noise filters only. Unfortunately, for the 35-ton, this uncorrelated component made up most of the noise across all channels (see Figure 5.2a). FFTs of the signal waveforms, performed separately for each RCE, were also useful in monitoring bands of noise in frequency space.

Monitoring of various other problems, such as the digitiser stuck code issue, synchronisation concerns resulting in a different number of microslices being saved in corresponding RCE millislices, and the asymmetry of bipolar pulses, were added as these issues became apparent during the commissioning.

5.2.2 Photon Detector Monitoring

Analogously to the TPC data, monitoring of the photon detectors mainly involved considering various ADC distributions separated by optical channel and by photon detector. The peak height, pedestal and integral of each waveform were also considered as a function of channel to ensure each were operating consistently.

The triggers sent on by the SSPs were also studied; unfortunately, due to the design of the monitoring framework (with it not guaranteed to receive each event), trigger rates were challenging to compute. It was decided to leave them in the monitoring but only consider the relative rates; the monitoring code may be utilised offline, processing closed files on disk, to determine accurate rates by ensuring all events are considered. Along with the trigger rate, the number of triggers, the fraction of events containing a trigger and the number of readout ticks within each trigger were also considered.

During installation, one photon detector was erroneously left unconnected to its SSP and so was unavailable during the run. This was discovered using the online monitoring framework but unfortunately only following the completion of the installation and the sealing of the cryostat.

5.2.3 External Counter Monitoring

Since monitoring the external counters primarily involves considering trigger rates, a similar issue to the photon detector monitoring was encountered; as with the SSP triggers, the rates were only considered relative to different counters. For each counter, the hit rate and the average activation time were monitored to ensure counters in similar positions were recording comparable cosmic muon data. The number and type of payloads sent on from the PTB were also detailed so the amount of data, along with information about what the data are comprised of, could be monitored.

5.2.4 General Monitoring

A variety of useful quantities not pertaining to any specific subcomponent were also monitored to assure smooth data taking. These include the size of output files and the average event size from recent runs, information about which detector subcomponents are taking data and the number of events seen by each, and also synchronisation information between various detector components.

5.2.5 DQM Plots

The DQM component of the online monitoring produced around 60 figures for each subrun, illustrating the data discussed in Sections 5.2.1, to 5.2.4. A sample subset of these figures is shown in Figure 5.2. Figure 5.2a shows the TPC noise; the total noise (RMS of the ADC values) is shown in blue and the uncorrelated component of this noise in green (the gaps indicating dead wires). The FFT of a waveform read out by the first RCE (channels 1–128)

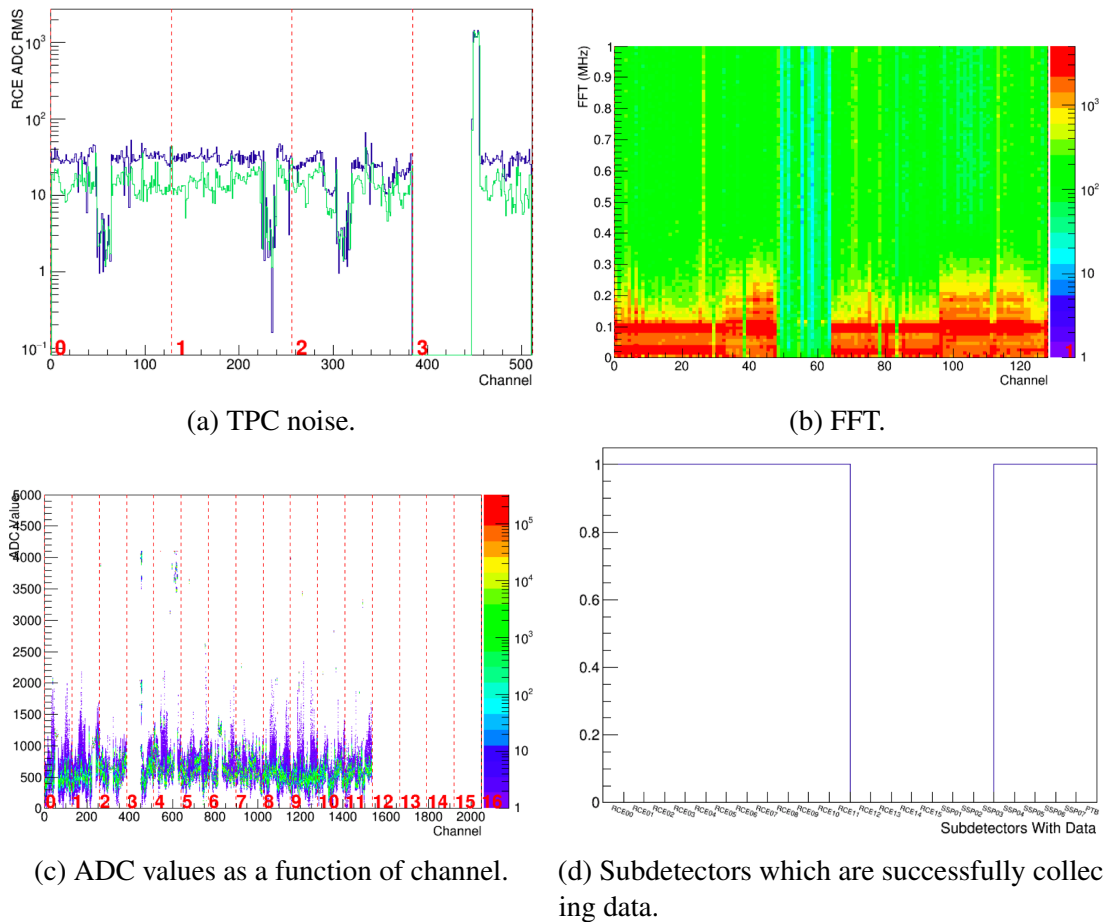


Fig. 5.2 Selection of figures made by the Data Quality Monitoring framework during 35-ton Phase II running.

is shown in Figure 5.2b and a 2D plot showing the ADC values for each channel, hugely useful as it demonstrates both the mean and RMS for all channels together, is depicted in Figure 5.2c. Figure 5.2d shows the subdetectors which are successfully collecting data and may be used to note one quarter of the TPC readout, along with three photon detector readouts, were turned off in this subrun.

5.3 Online Event Display

One of the highlights of the data taking was being able to watch events refresh on the online event display. In addition to the interesting visual display of interactions in the detector, it was an additional tool used in data monitoring; high noise states, poor LAr purity and drift field problems were all immediately evident from the display.

Given the structure of the data when the detector was read out, it proved to be challenging finding a comprehensible way to represent events. The construction of such a display is the subject of this section.

5.3.1 Selecting the Data

The raw data formats for the various 35-ton data streams were discussed in detail in Section 4.3.2.2. Each DAQ event comprises a collection of millislices, one for each of the detector subsystems (RCEs, SSPs, PTB), with further structure specific to each system and comprehensively illustrated in Figure 4.21. An example triggered event in the 35-ton data is demonstrated in Figure 4.22.

Since the event display runs online, a suitable selection must be applied to ensure the full physics event occurs within the current DAQ event; proceeding and preceding events are inaccessible during running. This is achieved by noting whether or not a trigger occurred (i.e. microslices contain nanoslices), and in which microslice it occurred, when reformatting the RCE data in `DataReformatter`. For the event display, an event is only useful if the trigger occurred within a certain range (e.g. Microslice 5 to Microslice 10), ensuring all the filled microslices are present within the current millislice. The event display is then filled for a given range of microslices around the trigger to capture all the physics data.

5.3.2 Representing the Data

The wrapped nature of the induction wires, and the inability to perform disambiguation without full reconstruction, results in only data from the collection plane being useful for an online event display. Use of a second dimension is possible if the detector is viewed from above and by using the drift time as a coordinate. This necessitates the two centre APAs to be shown together as one combined readout structure and a global two-dimensional coordinate system established for the entire detector. The wire coordinate is defined simply by counting wires from the collection planes across all APAs and incorporating fake wires between the frames, and the time coordinate may be used to take multiple drift regions into account by correcting all charge deposited in the short drift region to negative ticks.

The event displays are filled with the raw ADC values provided by the FE readout without the use of reconstruction. An approximate pedestal subtraction is possible by working with the system used to record these pedestal values. During data taking, the shifter would perform a run at least once per shift to produce a file containing all the calculated pedestals on each channel for subsequent uploading to a database for offline use. By ensuring a copy of the most recent file is always available to the monitoring framework, the pedestals may be

corrected for and the charge represented as accurately as possible. To limit noisy channels and to correct for accidental negative charge, the pedestal-subtracted ADC values are only included if within the range 0 – 250. Finally, given the relatively low signal-to-noise ratio, it was decided a grey-scale image showed the best resolution for observing tracks in the cryostat.

An example event display is shown in Figure 5.3.

5.4 Monitoring Web Interface

The output of the monitoring is vital in assuring the experiment continues to take high quality, analysable data. To facilitate this process, a web interface was developed to enable all useful information to be displayed and accessed in a convenient, universal location. This interface, along with the complementary web page, was relatively basic but was functional and performed all that was required for the purposes of a short prototype run. The method of automating the transfer of the monitoring data from where it was saved by the DAQ process to somewhere accessible by the web server is briefly described in Section 5.4.1 and the web page itself is overviewed in Section 5.4.2.

5.4.1 Automated Data Transfer

Ensuring the monitoring output was available in the correct place when needed was the most complicated part of the web interface. This was achieved using a combination of disk mounting and automated scripts, demonstrated in Figure 5.4.

The DAQ aggregator processes run on the lbne35t-gateway01 and lbne35t-gateway02 nodes, requiring any saved output be placed in a location accessible to these machines. Mounting a disk belonging to a gateway node onto these private machines and saving the output directly onto this ensured the data may be available outside of the private network. The constraints placed on the configuration by the DAQ group, which preferred nothing other than DAQ processes to run on lbne35t-gateway01, required a second gateway node, lbne35t-gateway02, be utilised. The web transfer framework was completed by mounting the Fermilab web area onto this machine and utilising an automated job to copy the monitoring data from the disk to the relevant part of the web server. The frequency of this job, 30 seconds, defined the maximum latency one could expect between data being written out and images appearing online.




Fig. 5.3 Example online event display made as part of the online monitoring framework for run 14306 (2nd March, 2016). The view is from the top of the detector looking down; the red lines represent the spaces between the APAs and the blue line the location of the APA frames, separating the long and short drift regions.

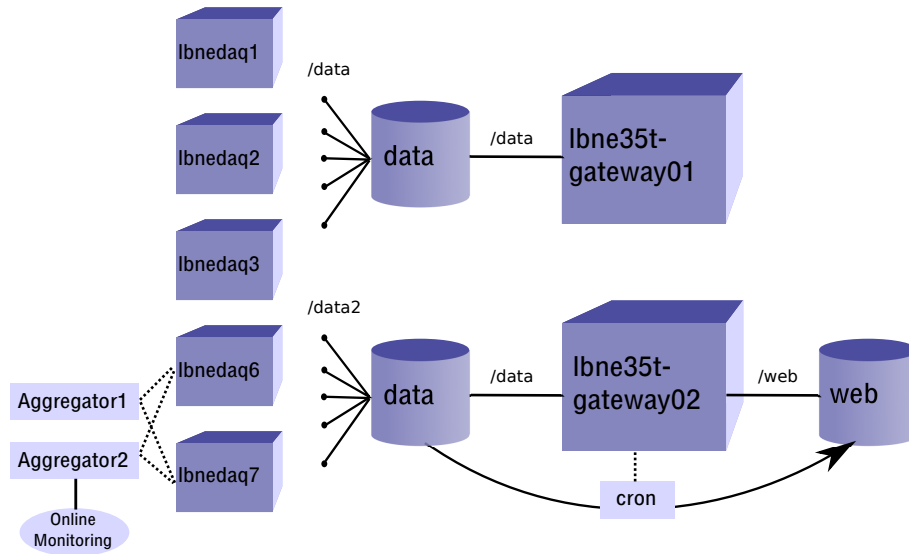


Fig. 5.4 Schematic showing the interface between the online monitoring system and the web. The DAQ machines are shown as rectangles with their disks represented as cylinders. Connections between a node and a disk are shown as straight lines, with dotted lines representing processes running on the machine.

5.4.2 Web Page

The web page was hosted at FNAL and located at lbne-dqm.fnal.gov. When the monitoring framework initiates a write out of all data products, the HTML necessary to correctly display these images is also written and saved as part of the output. This is copied, along with all the images and data files, to the web area as discussed in Section 5.4.1. The web page was basic but fulfilled all fundamental requirements for 35-ton monitoring; it had dedicated pages for all the data quality monitoring information and the online event display (the nearline monitoring was also hosted at this website but is not described here). See Figure 5.5 for a demonstration of web page and example navigation.

5.5 Online Monitoring Summary

The monitoring, with web support, was imperative for the success of the 35-ton. During the ongoing vertical slice tests in summer 2015, the majority of the setup was in place and enabled progress in testing and commissioning the APAs to be completed significantly faster than it otherwise would have been. During this time, and also during commissioning, the framework was the only way of analysing the data without reading it into LArSoft and writing specific software. Overall, the framework provided essential feedback and contributed

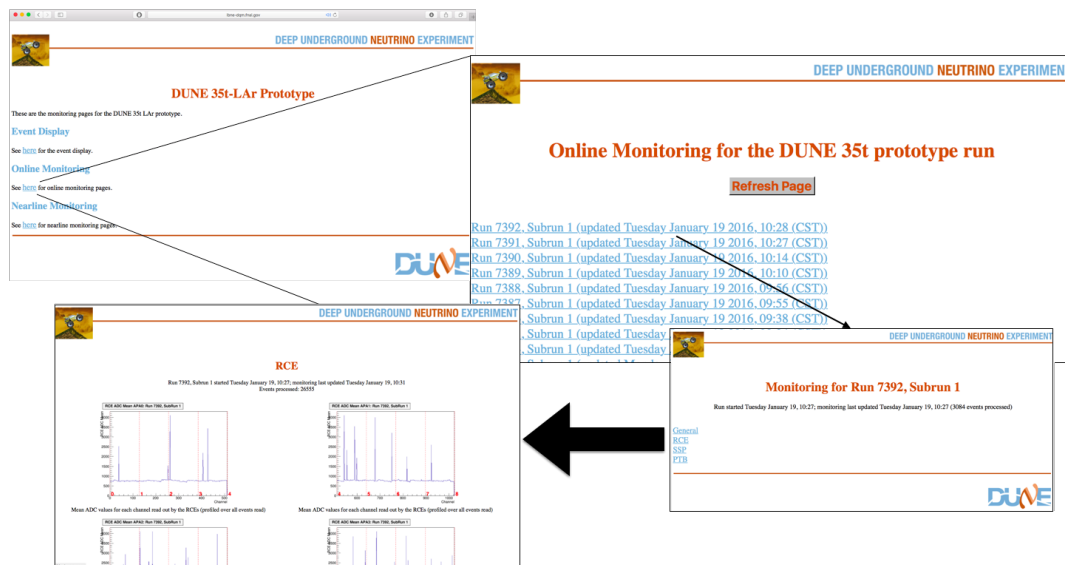


Fig. 5.5 Demonstration of the web page developed to display information produced by the online monitoring and event display. The pages are written in HTML and allowed prompt and convenient feedback directly from the DAQ be accessed anywhere and assist in remote monitoring of the experiment. All previous runs are also kept on the website for reference.

positively towards DAQ uptime during the data taking period. It is currently in the process of being adapted for future use in DUNE, specifically as part of the ProtoDUNE DAQ for the run in 2018.

Chapter 6

Reconstruction in a Liquid Argon TPC

The use of LArTPCs in future high-precision projects, such as long-baseline neutrino experiments, is very well motivated by the unprecedented spatial and energy resolution available to detectors utilising the technology. In order to take advantage of all this accessible information, accurate reconstruction methods must be developed to perform pattern recognition and energy determination for use by the proceeding analyses. The techniques and status of reconstruction in LArTPCs is the subject of this chapter, with particular focus on novel methods developed for the reconstruction of electromagnetic showers.

The implementation of the reconstruction algorithms discussed in this chapter utilises the Liquid Argon Software framework (LArSoft), developed at FNAL and shared between all experiments in the LAr program. LArSoft will be overviewed in Section 6.1 before the reconstruction chain is described in Section 6.2. The development of new techniques in the reconstruction of showers will compose the main discussion in this chapter and is contained in Section 6.3.

6.1 The LArSoft Framework

The Liquid Argon Software (LArSoft) [156–159] collaboration supports the development, use, sharing and distribution of code utilised by all LAr experiments at FNAL. The LArSoft framework is written in C++ and built on *art* [147, 148], the event-processing system established at FNAL and used by offline code developed for most experiments hosted at the lab. As data from most LArTPC experiments share a similar basic format, LArSoft is envisioned to be agnostic to the detector specifics and to provide a common interface, infrastructure and algorithms for simulation, reconstruction and analysis. Along with vastly reducing duplicated effort, this also allows access to the most advanced software developments for smaller collaborations who otherwise may not have the required resources.

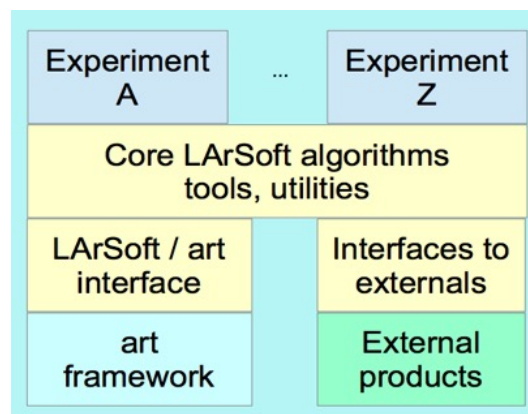


Fig. 6.1 The LArSoft architecture, highlighting support for both common and experiment-specific algorithms and methods and the interfacing with other packages. Taken from [158].

LArSoft provides well defined interfaces to various external packages, such as GENIE [160] and GEANT4 [161], and access by particular experiments utilises configurable descriptions of the detector geometry, the electronics, detector response and other unique features. This structure is demonstrated in Figure 6.1. This architecture ensures a flexible structure which facilitates the addition and evolution of algorithms and with contingency for future developments to be introduced with ease.

The *art* framework provides an interface to the information stored for each event by user-written ‘modules’, and handles execution by processing each entry and making the data available to these plug-in modules. Additionally, ‘services’, which exist outside of the event structure, are provided and may be used to obtain general information such as detector geometry or LAr properties whenever needed. The two main module types, named Analyzers and Producers, can access the data products stored in a particular event and, in the case of Producers, have the ability to place data into the event for use in future processes. All modules, regardless of their type, have read-only access to the existing information in the event. Configuration of an *art* job utilises the custom Fermilab Hierarchical Configuration Language (*fhicl*, pronounced ‘fickle’) which may be used to define the modules (including their order) and services to be run and to provide run-time parameters for use by these products.

The end-to-end configuration for a given experiment involves the following standard stages: generation (provided by a generator, such as GENIE), propagation (executed with GEANT4), detector response simulation, and reconstruction. The results from the first three steps aim to reproduce as closely as possible the expectations from real data, with the same reconstruction applied to both the simulated output from the detector and data. The processes

are configured, using Producer modules, in *art* using *fhicl*, generally separated into four jobs representing each of the stages.

LArSoft was initially developed for use in ArgoNeuT in around 2011 and has since progressed into the large collaboration it is today, with more than 100 code authors spanning multiple experiments. The constant progression of algorithms has resulted in very well-developed, advanced simulation and reconstruction tools with corresponding shared expertise. The recent use of LArSoft reconstruction on real data, in MicroBooNE [162] and the 35-ton (Chapter 7), is providing an excellent test of the efficacy of the simulation along with the validation of reconstruction applied to data. The current reconstruction chain is discussed in Section 6.2.

Until recently, convincing electromagnetic shower reconstruction had not existed within LArSoft. Despite this being a major advantage of LArTPCs, it is particularly challenging and requires significant investment of resources to fully understand. This motivated the development of new algorithms within the LArSoft framework, BlurredCluster and EMShower, which will be discussed in detail in Section 6.3.

6.2 The Reconstruction Chain

Reconstruction in LArSoft is the process of forming particle objects, with enough information to be able to perform identification, from the raw charge read out by the anode planes. The process may be considered as consisting of three main components: calibrating the raw charge to remove detector effects, pattern recognition, and calorimetry. These will be discussed in Sections 6.2.1, 6.2.2 and 6.2.3 respectively. The general workflow is shown schematically in Figure 6.2.

6.2.1 Raw Charge Calibration

The charge induced and collected on the readout wires is modified by detector effects which must be well understood in order to be properly accounted for. The measured waveform, $p(t)$, obtained from the signal $s(t)$, measured as a function of time t , can be represented in pseudocode as

$$p(t) = (s(t) \otimes e(t) \otimes f(t)) + n(t), \quad (6.1)$$

where e is the electronics response, f is the field response and n is the noise in the detector. Together, $e(t) \otimes f(t)$ are referred to as the detector response.

The first step in the reconstruction, referred to as the ‘deconvolution stage’, involves removing these detector responses. This proceeds by subtracting the noise profile from the




Fig. 6.2 The LArSoft reconstruction workflow to produce 3D reconstructed objects from the raw charge read out by the anode wires [158]. The status of the reconstruction is shown on the right, with the various algorithms and their outputs represented on the left.

measured waveform before Fourier transforming into frequency space and dividing out the field and electronics components. The detector responses must be accurate and the models used have been developed at test stands to ensure they best represent the detector effects. This process is applied in reverse during the detector simulation stage of the simulation to reproduce the expectations from the data as much as possible.

The detector response is demonstrated in Figure 6.3. The principle of current induction on the anode wires is described by the Shockley-Ramo [163, 164] theorem; the instantaneous induced current i is given by

$$i = q\vec{E}_w \cdot \vec{v}_q, \quad (6.2)$$

where q is the charge of an element of ionisation, \vec{E}_w is the field vector at the location of the charge and \vec{v}_q is the velocity of the charge packet [165]. The electric potential for the full wire planes is based on a 2D Garfield simulation [166]; simulated TPC signals for each of the planes are demonstrated in Figure 6.3a. The induced current is received, amplified and shaped by the preamplifier in the front end electronics. Typical shaping times and gains are demonstrated in Figure 6.3b.

Following the deconvolution stage, the waveforms all have the form of a unipolar pulse. The reconstruction proceeds with ‘hit finding’, with the purpose to accurately determine the properties of the collected charge. In particular, the peak time, width and total charge of the

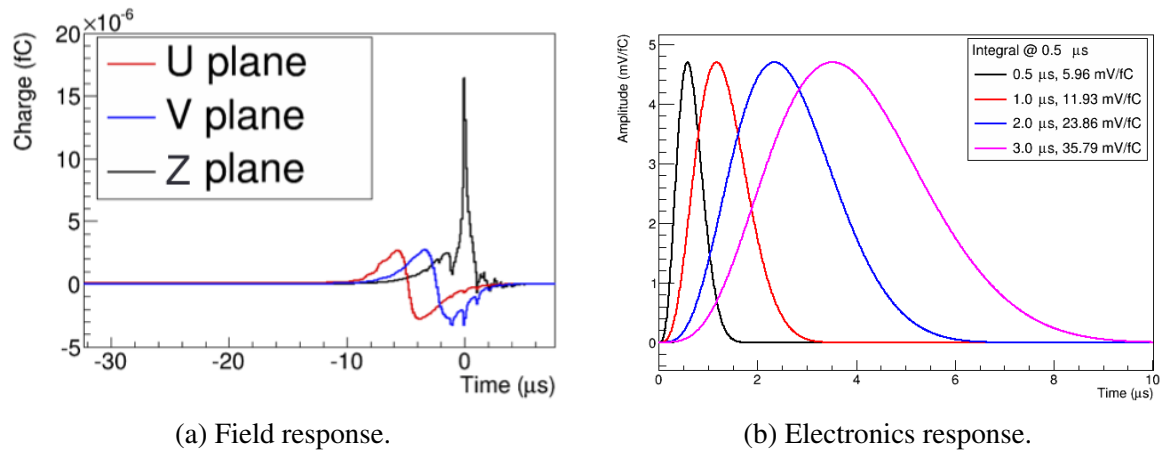


Fig. 6.3 Detector response for charge readout in a LArTPC. The field response is shown in Figure 6.3a and the electronics response, defined by configurable front end electronics settings, is demonstrated in Figure 6.3b. Based on descriptions in [165].




Fig. 6.4 The process of deconvolution and hit finding to determine the correct charge from the measured pulses on the readout wires. Each 2D plane is represented as charge (ADC) as a function of time (tick, 500 ns). In each case, the black waveform represents the raw measurement, the blue shows the outcome of deconvolution and the pink peak indicates the reconstructed hit. Note the normalisation of the deconvoluted signal has been fixed to provide a factor four amplification for illustrative purposes. The shift in time and shape is a result of the deconvolution from the electronics response (demonstrated in Figure 6.3b).

‘hits’ are pertinent for future reconstruction algorithms. This is typically achieved by fitting a Gaussian to the pulse and using this to aid the determination of the hit properties. The result of the deconvolution and hit finding stages are represented for simulated hits on three separate planes in Figure 6.4. The hit finder used here is the Gaussian Hit Finder in LArSoft.

Due to the wrapped wires, multiple hits will be reconstructed on each channel, one for each possible wire segment. The final stage of hit processing is to perform disambiguation to select the correct hit for use in subsequent algorithms. As previously discussed (Section 3.3.2.1 and Section 4.3.1.1), this is trivial in the far detector design as the wire wrapping angles are chosen such that no induction wire segment crosses each collection channel more than once. In the 35-ton, the slight difference in angle between the two induction planes results in ‘triple points’, where there are hits on each plane almost instantaneously, which facilitates a deduction of the correct induction channel hits.

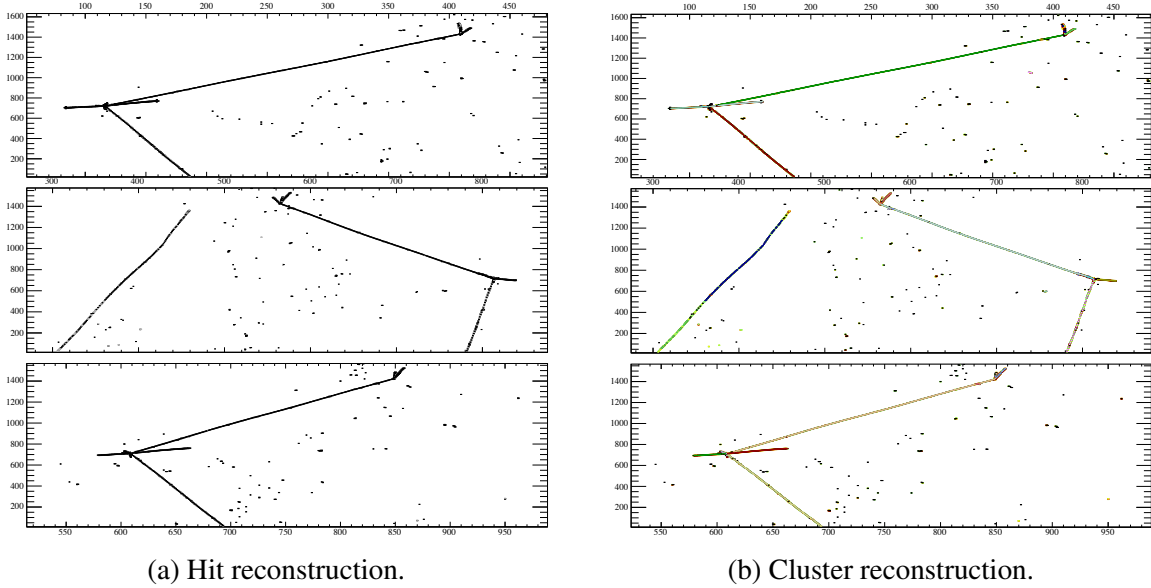


Fig. 6.5 Demonstration of 2D reconstruction in LArSoft on a simulated ν_{μ} CC event. The result of hit finding is shown in Figure 6.5a, where each black rectangle represents a separate hit, and the grouping of these hits into clusters is illustrated in Figure 6.5b, where blocks of hits sharing a common colour are associated to the same cluster object.

6.2.2 Pattern Recognition

The next typical stage of the reconstruction following hit finding is the process of forming 2D objects, or ‘clustering’. In LArSoft, all 2D objects, regardless of their topology, are named ‘clusters’ and are simply a collection of hits identified as being associated with a common ionising particle. These 2D structures exist only on a given plane and do not have contributions from multiple views; the extension to 3D reconstruction involves combining numerous clusters between the views. Multiple cluster algorithms exist in LArSoft, each designed for different specialisations.

An example 2D view of a ν_{μ} CC event, simulated in a reduced far detector volume, is shown in Figure 6.5. This particular interaction is a deep inelastic scatter producing, along with the muon, multiple final state protons and a very high energy π^+ meson. The outcome of hit reconstruction is demonstrated in Figure 6.5a and the result of applying clustering to these hits is displayed in Figure 6.5b, with each colour representing a separate cluster. The 2D cluster finder used here is Cluster Crawler [167].

Combining 2D information from multiple views in a LArTPC enables the formation of 3D objects. In LArSoft, several 3D products exist: ‘space points’, ‘vertices’, ‘tracks’ and ‘showers’. Space points are a form of 3D hit, defined simply by a point in the detector Cartesian geometry, and vertices are a specific example of these locations created to represent

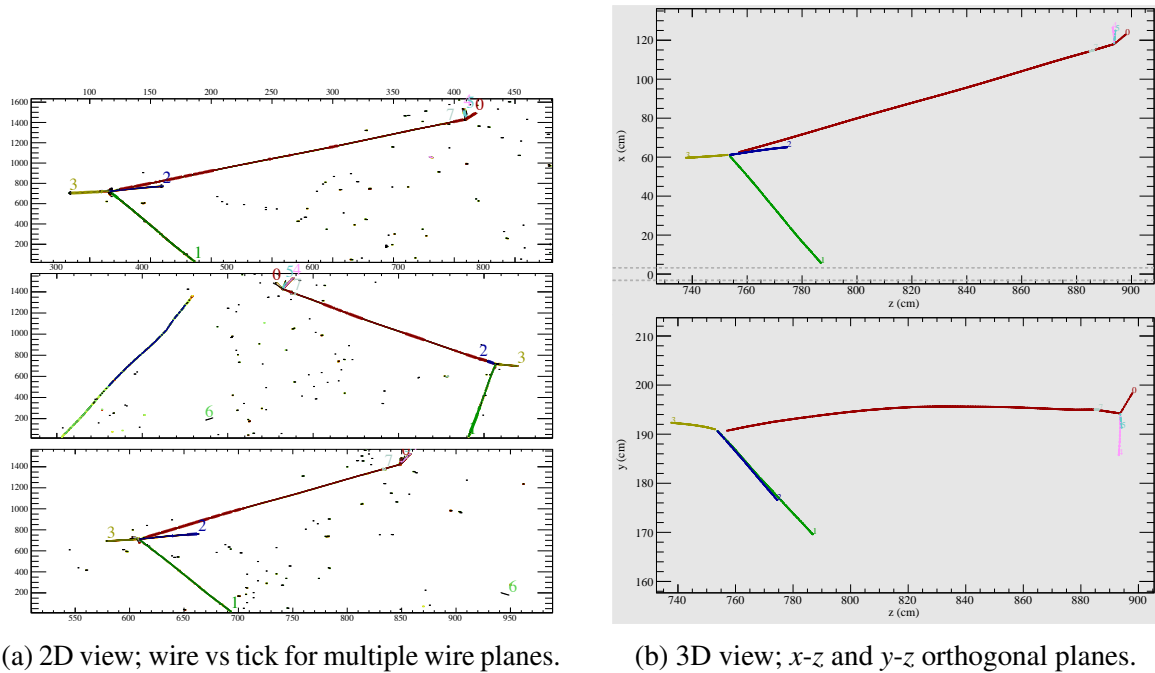


Fig. 6.6 Demonstration of 3D reconstruction in LArSoft on a simulated ν_{μ} CC event. Each colour represents a unique 3D track object, with information shared between each of the anode planes; this is evident in Figure 6.6a. An additional view is demonstrated in Figure 6.6b where two orthogonal planes of the detector geometry are shown.

an important part of the event (e.g. the neutrino interaction point or even just a track start). Tracks and showers are more complex objects which represent individual particles and contain associated properties relevant to the topology of each. Tracking is relatively advanced in LArSoft, with multiple track fitters shown to produce well-formed track objects efficiently; shower reconstruction is less so and will be discussed in detail in Section 6.3.

The result of applying track reconstruction to the ν_{μ} CC interaction discussed in Figure 6.5 is shown in Figure 6.6. As with clusters, each colour represents a unique track object and it is clear from Figure 6.6a that hits across multiple planes contribute to a given track. A complementary view is also shown in Figure 6.6b, where the objects are represented in two orthogonal views in the detector coordinate system. The track finder used here is Projection Matching Algorithm (PMA) [168].

It is worth noting the approach to forming the eventual 3D objects, from which particle identification and analysis may be executed, is not unique. The method outlined above is most relevant to the shower reconstruction discussed in Section 6.3 and is the chain used in the majority of LArSoft processing; however, progress is being made on an additional technique which focusses directly on 3D reconstruction. This is called ‘Wire-Cell’ [169] and utilises a tomographic method, following the LArTPC principle that each plane observes the

same amount of ionisation electrons, to build up a 3D image of the event. This may then be characterised to find the properties of the particles and make the track and shower products. All reconstruction methods utilise the same information provided by the detector but not necessarily in the same order.

6.2.3 Calorimetry

The possibility of extremely precise energy measurements is a significantly attractive feature of LArTPCs and is achieved by careful calibration of the charge received from the drift electrons. There are a number of distinct calibrations which must be conducted to ensure reliable calorimetric measurements and the relevant procedures will be outlined in this section. Setting the ADC to energy translation is the subject of Section 6.2.3.1 and determining the electromagnetic shower and neutrino energy conversion from the total deposited charge will be discussed in Sections 6.2.3.2 and 6.2.3.3 respectively.

6.2.3.1 Energy Determination

Following deconvolution and hit finding, the charge carried by each packet of electrons is understood as the Gaussian integral of the hit and is normalised, by default, in units of $200 e^-$ (often just ‘ADC’ in certain following plots). The ‘calorimetry constant’, responsible for converting between digitised readout charge and the number of electrons which deposited the charge, is therefore 5×10^{-3} in simulation ($1/200$), but is not necessarily equivalent in data due to limitations of the simulation. This will be discussed further in Section 7.4.2 in Chapter 7.

The precise calorimetry constant may be tuned by considering minimally ionising particles (mips), which are known to deposit 2.1 MeV/cm energy when propagating through LAr. The number of electrons contributing to a hit is determined using the calorimetry constant and corrected for lifetime and recombination to yield the total number of ionised electrons initially forming the ionisation packet. Utilising the ionisation energy of LAr (23.6 eV/ion), the total energy deposited by the particle may be inferred. Combining this information for hits on successive wires, accounting for the pitch of the particle track, enables a measurement of dE/dx to be made. By plotting the distribution for multiple tracks, the calorimetry constant may be tuned until the dE/dx is as expected.

6.2.3.2 Shower Energy Reconstruction

For large groups of hits from a common source, such as for electromagnetic showers, an alternative approach may be employed. A linear relationship is observed between the lifetime-




Fig. 6.7 Correlation between the true deposited energy in an electromagnetic shower and the total charge in the detector for electron showers in the DUNE far detector. The parameters from the linear fit are extracted for use in the reconstruction.

corrected charge deposited and the true deposited energy, demonstrated in Figure 6.7. This was made by simulating 1000 electrons in the DUNE far detector geometry at each of the energies 0.5, 1.0, ..., 5.0 GeV and taking a Gaussian fit of the total charge deposited in the detector for each energy. An almost-perfect linear fit is observed and the associated parameters are used when determining the shower energy in reconstruction.

6.2.3.3 Neutrino Energy Reconstruction

In order to reconstruct the neutrino energy from beam interactions, the total charge deposited in an event is considered under the assumption that it all originated from the incoming neutrino. Leptonic energy (i.e. muon tracks and electron showers) is first determined from either the electron shower energy or, in the case of a muon, by using either containment or multiple Coulomb scattering (MCS) considerations. The remaining charge is then treated as hadronic energy and a similar conversion to that demonstrated in Section 6.2.3.2 is employed to perform the conversion [170]. This approach ensures the correct calibration for the hadronic energy component, which often includes unseen contributions from neutrons.

6.3 Shower Reconstruction in LArTPCs

Reconstructing showers in a LArTPC is challenging and until recently was an unsolved problem in LArSoft. This section describes the development of novel techniques for the

purposes of complete shower reconstruction within a LArTPC. The primary motivations for the undertaking were initially the desire to perform a π^0 analysis in the 35-ton dataset and the lack of available tools. The reconstruction was thus developed with π^0 s in mind, specifically focussing on the separation of closely occurring showers from the two photon daughters, but has utility in the reconstruction of all showering particles.

Excellent shower reconstruction is essential for the success of the DUNE experiment. In order to distinguish between the main oscillation signal (discussed further in Chapter 8)

$$\nu_e + n \rightarrow e^- + p^+ \quad (6.3)$$

and the potentially tricky neutral current background

$$\nu_\mu + n \rightarrow \nu_\mu + \pi^0 + \text{hadrons}, \quad (6.4)$$

the electron and the π^0 decay photons must be well discriminated. This motivates the requirement of high quality reconstruction of these electromagnetically showering particles.

The shower reconstruction problem is first overviewed in Section 6.3.1 before the specific algorithms are discussed in Sections 6.3.2 and 6.3.3 respectively. A further problem in LArTPC reconstruction, track/shower separation, is the subject of Section 6.3.4, before finally the performance of the reconstruction is evaluated in Section 6.3.5.

6.3.1 Showers Overview

High energy electrons and photons produce electromagnetic showers in a LArTPC via pair production and bremsstrahlung processes. They may be distinguished using the excellent calorimetric properties of the detector by careful analysis of the initial part of the object, before showering occurs. A photon, since it has no electric charge, does not ionise and so is only observable following pair production. Thus, one would expect the initial part of a photon shower to be twice as ionising (two electrons) than the start of an electron shower (a single electron); the dE/dx at the beginning of an electron and a photon shower is 2.1 MeV/cm and 4.2 MeV/cm respectively.

The required properties provided by the reconstruction of electromagnetic showers include an accurate start position and direction, a dE/dx measurement for the first 3 cm of the shower and the total deposited energy left by the shower. It is thus vital the object is as complete as possible and is orientated correctly. Additionally, good separation between π^0 decay photon showers is required for calibration purposes and to reduce the neutral current




Fig. 6.8 An example particle gun π^0 event in the 35-ton geometry. Hit reconstruction (including disambiguation in the induction planes) has been applied with individual hits shown as black rectangles. The two decay photons are visible by eye but shower very close to each other.

π^0 backgrounds previously discussed. An example 35-ton particle gun π^0 event is shown in Figure 6.8 to demonstrate the task presented to the reconstruction.

The algorithms presented in the following section successfully meet these criteria and produce high quality reconstructed showers with high efficiency. Two methods are outlined, BlurredCluster and EMShower, which are designed to perform 2D and 3D reconstruction respectively and together produce fully reconstructed shower objects.

6.3.2 BlurredCluster Algorithm

The BlurredCluster algorithm is a 2D clustering algorithm [171] implemented with LArSoft [172, 173]. It is optimised for shower reconstruction in general and, more specifically, for the reconstruction of π^0 s. The method convolves the hit map with a Gaussian kernel in order to ‘smear’ out the deposited charge and introduce ‘fake’ hits to facilitate a more complete and accurate reconstruction.




Fig. 6.9 An example 2D Gaussian kernel, made using Equation 6.5, alongside the contributions from each dimension. In the context of convolving with a hit map for reconstruction purposes, consider the charge at point (0, 0) to be a hit. The neighbouring bins, within a given blurring radius r (in this case, 4 in each direction), are assigned charge according to the value of the kernel at each point. If no charge were originally deposited at that location, a ‘fake’ hit is introduced. In this example, the Gaussian in the horizontal direction is much wider ($\sigma_{\text{wire}} = 2.6$) than in the vertical direction ($\sigma_{\text{tick}} = 1$).

6.3.2.1 Algorithm Description

The energy deposited by showering particles is contained on several discrete readout wires and is not necessarily representative of the initial deposition of charge, which is more diffuse. This motivates the concept of ‘smearing’ the charge between adjacent wires, reproducing the more realistic distribution. In addition, since ionisation packets naturally spread in time (longitudinal diffusion), it is natural to take this into account when redistributing the deposited energy.

The smearing is performed using a 2D Gaussian kernel, given by

$$\frac{1}{\sqrt{2\pi}\sigma_{\text{wire}}} \frac{1}{\sqrt{2\pi}\sigma_{\text{tick}}} e^{-\frac{r_{\text{wire}}^2}{2\sigma_{\text{wire}}^2}} e^{-\frac{r_{\text{tick}}^2}{2\sigma_{\text{tick}}^2}} \quad (6.5)$$

and defined by two parameters in each dimension (wire and tick); the standard deviation σ of the Gaussian function and the ‘blurring radius’ r , the distance the blurring is extended in each direction. An example of such a kernel is demonstrated in Figure 6.9.

The Gaussian widths and blurring radii in each direction are determined dynamically for each individual hit map and even for individual hits. This is achieved by initially guessing a direction of each shower (using least squares and Principal Component Analysis (PCA))

considerations) and assigning the parameter values relevant to the assumed directionality. For example, if it appears a particle shower is travelling primarily in the wire direction, an appropriate 2D kernel would be wider in the wire direction and would extend further in this dimension (i.e. $r_{\text{wire}} \gg r_{\text{tick}}$). This example is demonstrated in Figure 6.9 where it is clear the charge is noticeably smeared more in the wire (horizontal) than the tick (vertical) direction. Additionally, the width of a particular hit in time is taken into account when selecting an appropriate kernel with which to smear its charge. A hit with a greater width justifies a much wider kernel in the tick direction to take into account this distribution of charge.

A demonstration, in 1D, of the convolution process is the subject of Figure 6.10. In this example, the width of the kernel is fixed and the blurring radius is set to 2. The final charge distribution may be compared to the input hit map in the bottom right figure, in which a more realistic smearing of the deposited energy in the detector may be observed. The process in 2D is identical to this simplified example but uses contributions from bins in two dimensions rather than just one.

The result of the Gaussian blurring process on the 35-ton π^0 depicted in Figure 6.8 is shown in Figure 6.11. The efficacy with which the blurring encompasses the showers, whilst maintaining separation, is evident. This is the aim for the smearing process and the resulting shower quality is directly dependent on how well the charge is redistributed before clustering.

Following completion of the hit map blurring, clustering proceeds utilising a simple nearest neighbour algorithm. This is sufficient to group all hits of a shower together once the Gaussian smearing has smoothed out the charge distribution and introduced fake hits to populate gaps in the collected charge. The final stage of the process involves identifying the original hits and producing the relevant output clusters for use in further algorithms. The output clusters produced by the BlurredCluster algorithm for the 35-ton π^0 event are demonstrated in Figure 6.12.

6.3.2.2 Configuring the Clustering

The BlurredCluster algorithm takes reconstructed hits as input, previously determined by a LArSoft hit finder (such as gaushit). It places the output clusters, along with associations to the hits which comprise them, back into the event record. There are multiple user-defined parameters which must be tuned to provide the best reconstruction. This has been done for both the DUNE 35-ton and FD geometries and is dependent on detector properties such as the induction wire angle and anode wire pitch.

There are two ways in which the reconstruction can be performed in a given detector; either for each cryostat, drift volume (DV) and plane separately or alternatively across all DVs for each plane in a given cryostat. This distinction is necessary due to the detector




Fig. 6.10 Simplified demonstration, in 1D, of the blurring process used when smearing charge from a hit map. The process follows from left to right, top to bottom. The Gaussian width is fixed for each hit and the blurring radius is set to 2. For each bin, any real hits within the blurring radius (represented by red lines) contribute to the charge introduced as a fake hit. The original hit map is shown underneath the blurred version, which is completed as the bins are looped through. The final blurred hit map shows a smearing of the input charge according to the original distribution of charge and the particular Gaussian kernel used in the convolution.




Fig. 6.11 The output of the blurring stage of the BlurredCluster algorithm on the hit maps from two planes of the 35-ton π^0 event illustrated in Figure 6.8. The greyscale represents charge deposited in the wire/tick space.




Fig. 6.12 The 2D clusters made using BlurredCluster when applied to the 35-ton π^0 event shown in Figure 6.8. The different colours represent separate cluster objects.

readout, which arbitrarily breaks up the active region into distinct DVs depending on the APA which reads the charge out. When performing reconstruction, it is unnatural to consider each DV separately and so instead a ‘global’ volume is defined. Such a coordinate assigns a wire number for all wires on a given plane in a single cryostat, agnostic to the boundaries defined by the detector hardware. This scheme is demonstrated in Figure 6.13 and is used by default by BlurredCluster.

6.3.3 EMShower Algorithm

The EMShower algorithm is a 3D shower reconstruction algorithm implemented within LArSoft [174, 175]. It was conceived as an extension of the 2D BlurredCluster reconstruction method and takes these clusters as input. The philosophy behind such a design choice is to perform as much of the reconstruction as possible in 2D, where most of the information exists, before simply applying a 3D matching algorithm to the output to produce complete objects. Since BlurredCluster has been shown to produce very well formed, complete clusters, EMShower essentially just combines these clusters together to form 3D shower objects.

The shower reconstruction is intended to be very high-level and depend heavily on previous reconstruction to take advantage of as much available information as possible. It is




Fig. 6.13 Demonstration of the ‘global wire’ concept in single-phase LArTPC reconstruction. A single hypothetical induction plane is shown across the 4 APAs in the 35-ton geometry. The blue boxes represent the APAs and the black lines and numbering at the top show the wires in one plane. In the left-hand figure, the APAs are considered separately and the conventional wire numbering is shown, with each wire represented by four numbers: the cryostat, DV, plane and wire number. Reconstruction may be performed in each DV independently using the wire number as the spatial coordinate. In the right-hand figure, the entire volume is considered as a whole and the wires are renumbered to reflect their position in the global geometry. An imaginary line connecting wires which would overlap others on different APAs (and thus different DVs) if extended is used to give each of the wires a common index. In this scheme, each is described using just three numbers: the cryostat, plane and global wire.

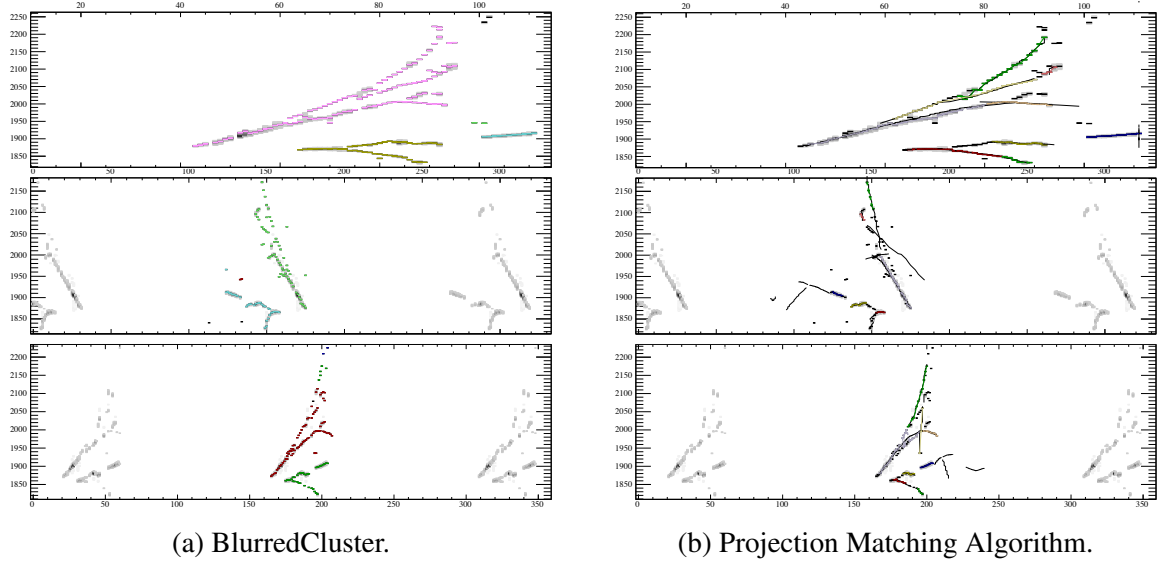


Fig. 6.14 Comparison of the 2D shower cluster reconstruction and the 3D tracking reconstruction for a 35-ton π^0 event. In both cases, the different colours represent individually reconstructed objects. It should be noted that, due to the 3D nature of the tracks, the same colours are evident across multiple views in Figure 6.14b, representing the hits used to construct the shower in each plane.

written very generically, making no assumptions on the specifics of the detector geometry such as the number of planes (LArIAT, which only possesses two readout planes, successfully uses EMShower to reconstruct showers). The shower-forming proceeds in two general steps: pattern recognition provides the geometrical shower shapes by combining 2D clusters from multiple views before analysis of the hits across the views is performed to determine properties of these shower objects. These stages will be discussed in the following Sections 6.3.3.1 and 6.3.3.2 respectively.

6.3.3.1 Shower Pattern Forming

Taking advantage of the well developed tools already within LArSoft, EMShower performs cluster matching by utilising previously conducted 3D track reconstruction. This is demonstrated in Figure 6.14, which shows the output of BlurredCluster adjacent to the 3D tracks found by the Projection Matching Algorithm for the 35-ton π^0 event shown in Figure 6.8. It is clear how using the respective hit associations facilitates connections between 2D clusters and 3D tracks, essentially matching the clusters between the views and forming 3D objects. The output from this stage is a collection of hits which are assumed to be part of the same shower object; the shower properties are determined by analysing these hits.

The algorithm performs checks to ensure two cluster merge candidates are indeed related to the same particle. The number of track hits associating the clusters and the length of the track are taken into account before matching two objects across views. Additionally, with three planes, it is possible to analyse the 2D reconstruction to determine, for example, whether or not one plane contains one large cluster whilst the other two planes contain two smaller clusters, which may be indicative of poor reconstruction in the first view.

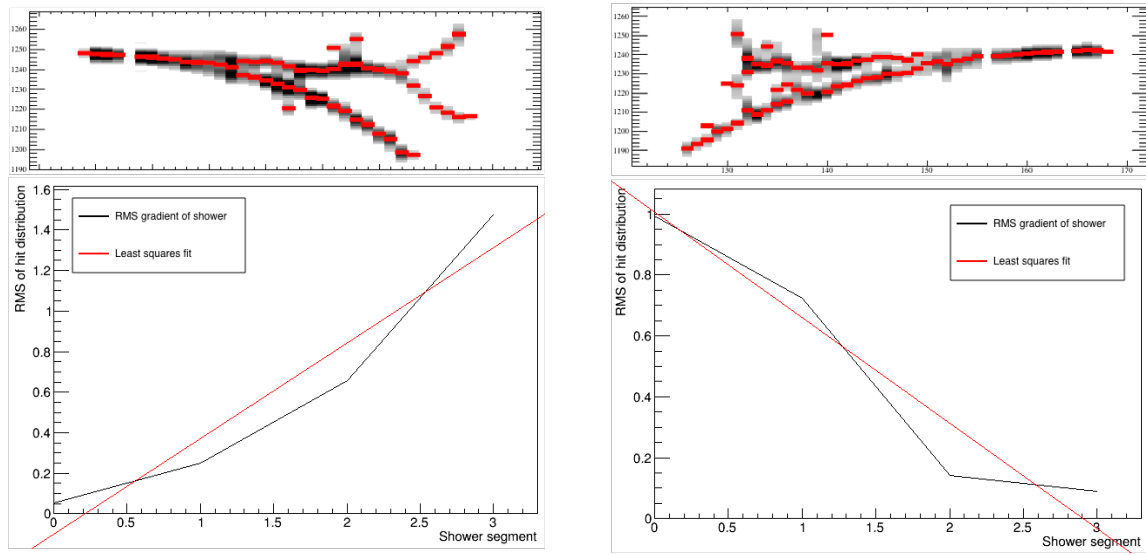
6.3.3.2 Shower Properties Reconstruction

Following the formation of the showers and the identification of all the associated hits in each view, the properties of the object may be determined. The shower energy may be determined for each plane using the method discussed in Section 6.2.3.2. For all other relevant properties, it is imperative the start of the shower is successfully found and efficiently reconstructed. From this initial track before the particle begins to shower, the start point (conversion point), initial direction and dE/dx may be determined.

In order to orient the shower, the hits are analysed in each plane separately whilst using information from other planes to validate. Initially, the hits in each plane are assembled into a rough ‘order’, corresponding to how far along the shower each is. This is achieved by projecting each hit onto the axis determined using a least squares fit on all the hits in a given view. In most cases, this order is accurate but not necessarily oriented correctly. Occasionally, if the shower is not very well formed (such as if it travels very parallel to the wires) in one plane, the hit order is not defined and the view is discarded for the purposes of shower start finding. This is accomplished by comparing the RMS of the perpendicular distance of all hits from the central axis and removing views in which this differs significantly.

Once there is sufficient confidence that the hits are correctly ordered in the remaining planes, they are then oriented correctly. This is done by evaluating the ‘RMS gradient’, the RMS of the distribution of the hits in discrete segments along the length of the shower. Figure 6.15 demonstrates how this may be utilised to correctly orient the shower. As is evident from the figure, a positive gradient, implying the distribution of the hits is becoming more diffuse along its axis, indicates the assumed orientation of the shower is correct and a negative gradient is indicative of a reverse-facing shower. This represents an overly simplified case but this method has been found to work surprisingly well across a wide range of differing shower and event topologies, at least as an initial guess. Additional checks are in place to ensure the orientations agree between views and to re-evaluate the shower ordering if necessary.

After correctly orientating the hits along the axis of the shower in each view separately, it is straightforward to determine the initial track-like parts. For each plane, the hits are



(a) Positive gradient.

(b) Negative gradient.

Fig. 6.15 Demonstration of the method utilised to ensure a correct shower orientation of the shower in a given view for the purpose of shower start finding in the EMShower algorithm.

considered in order until it becomes obvious the shower is diverging (e.g. multiplicity of hits on the wires). The initial shower hits from each of the planes are combined in order to construct a 3D track object, which is used to provide a start point and direction. These hits are additionally analysed to provide a dE/dx for the start of the shower.

The result of applying the full EMShower algorithm to the 35-ton π^0 event from Figure 6.8 is presented in Figure 6.16. The complete 3D shower objects for each photon are evident with well reconstructed properties, demonstrating the successful shower reconstruction provided by BlurredCluster and EMShower together. The reconstruction is characterised and validated further in Section 6.3.5.

6.3.4 Track/Shower Separation

The shower reconstruction described in Sections 6.3.2 and 6.3.3 is optimised for shower-like hits and makes no attempt to ignore hits which originate from tracks. This issue of track/shower separation is a complex problem in LArTPCs and remains largely unsolved. Significant progress has been made in the last year however, mainly via the use of machine learning. The vastly differing topologies, wide band energies in the DUNE beam and highly detailed event information all contribute to a challenging reconstruction issue unlikely to be solved without the use of deep learning algorithms.

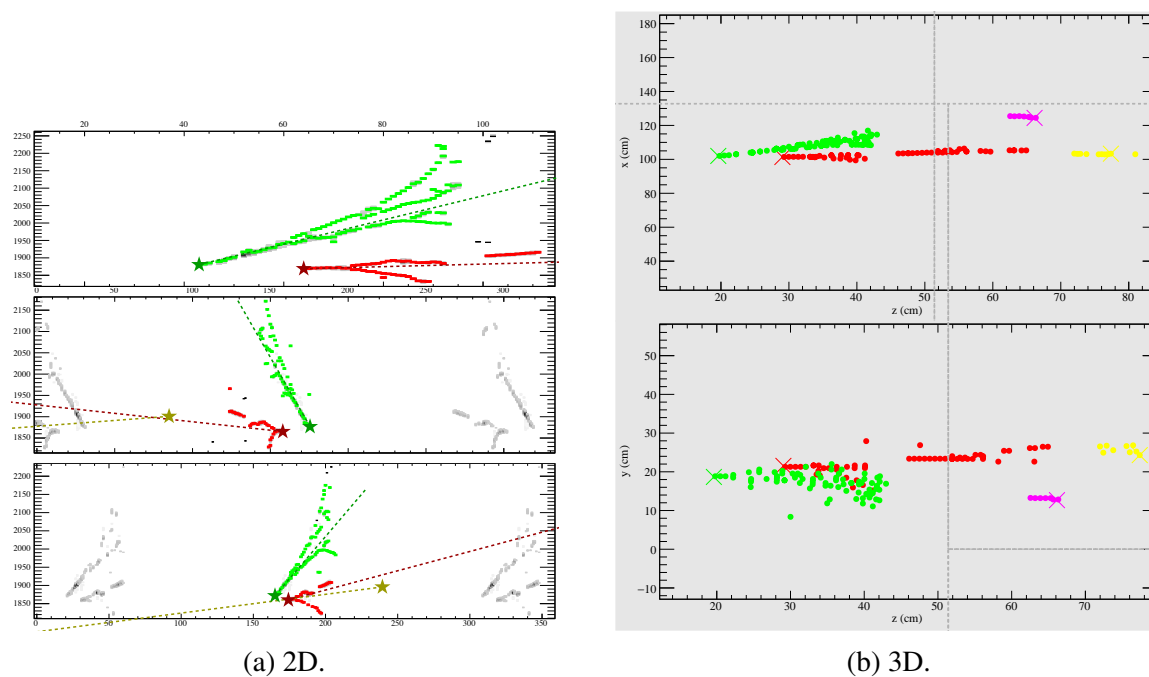


Fig. 6.16 The output from the full 3D shower reconstruction provided by BlurredCluster/EMShower when applied to the 35-ton π^0 event shown in Figure 6.8. The showers are shown in both the 2D wire/tick view in Figure 6.16a and in the orthogonal 3D representation in Figure 6.16b. The stars/crosses represent the reconstructed start point and the dotted lines the reconstructed direction.




Fig. 6.17 An example simulated ν_e CC interaction in the DUNE far detector. The Gaussian hit finder has been used to reconstruct hits, shown as black rectangles. The electron shower is evident, along with many other smaller hadron tracks. The ν_e energy in this event is 3.7 GeV and the electron has an energy of 1.9 GeV.

This section briefly details the development of an alternative approach to the problem, which was found to be inferior to the more recent developments. It is nonetheless instructive to overview the subject and assess the attempted method to evaluate how the experience may inform other potential strategies in addressing the issue.

The algorithm was developed primarily for the purposes of the ν_e CC oscillation appearance selection, the subject of Chapter 8. An example MC ν_e CC event is demonstrated in Figure 6.17. The topologies of the different particle objects are evident by eye and the purpose of the track/shower separation is to remove the track-like hits from the events, leaving just shower-like hits over which the shower reconstruction algorithms may be run.

Following the success of EMShower, which utilised as much previously executed reconstruction as possible to access all available information about an event, the track/shower separation was also conceived to benefit in this way. It is designed to be performed following track finding and, in essence, simply analyses each of the reconstructed track objects to identify which are the result of track-like particles. The associated hits are then removed from the sample passed on to BlurredCluster for shower reconstruction.




Fig. 6.18 A simplified event topology demonstrated typical track and shower topologies. The three distinct features are highlighted; the blue label represents a track, the red a ‘shower track’ and the green a ‘shower cone’.

A typical, though simplified, event, containing a track and shower, is demonstrated in Figure 6.18. The three topological objects are noted as a *track*, a *shower track* (the initial part of a shower before the electromagnetic cascade occurs) and a *shower cone* (the rest of the shower downstream of the shower track). The track/shower separation treats each of these objects independently to ensure a complete evaluation of the entire topology. The algorithm proceeds in the following general steps:

- track-like objects (including *tracks* and *shower tracks*) are identified;
- *shower tracks* are identified;
- *shower cone* activity is identified.

The *tracks* and *shower tracks* are deduced by considering additional activity along the length of each previously reconstructed track object, in both 2D and 3D. Tracks which do not belong to the shower cone have a characteristically lower amount of external hits in a surrounding cylinder or rectangle, and are identified with good efficiency. Following this, the number of previously determined tracks and hits in a projection from either end of these track-like objects may be used to distinguish *shower tracks*. The *shower cones* are determined by examining the event following the determination of any *shower tracks*. Additional checks are in place, such as reevaluating a track-like object upon discovering a large amount of activity downstream of its end and reversing the orientation of objects based on other identified activity in the event, to ensure the method is as robust as possible.

The result of applying the algorithm to the DUNE ν_e CC event shown in Figure 6.17 is demonstrated in Figure 6.19. The electron shower is observed to be well separated from the other tracks from the vertex and reconstructed nicely. This is a very successful reconstruction




Fig. 6.19 The result of applying the track/shower separation, BlurredCluster and EMShower algorithms to the example DUNE far detector ν_e CC event demonstrated in Figure 6.17. Both the 2D views and 3D projections are shown. The colours represent individual shower and track objects (note the same colour may be used for both a track and a shower without implying the objects are related) and the crosses demonstrate the reconstructed start position of the showers.

of the original event and may be utilised in the ν_e CC selection discussed in Chapter 8. A more general analysis of the performance of the reconstruction is considered in Section 6.3.5.

The reasoning behind performing this separation following the completion of track finding, as mentioned previously, is to utilise as much information as possible about the interaction. There are however many drawbacks which were identified whilst developing the algorithm. For example, any mistakes in the track reconstruction has implications for the success of the overall analysis of the event and may be a considerable problem. Additionally, track reconstruction is not optimised in the important region around the vertex and often fails to distinguish separate particles or too aggressively approaches a complex hit distribution. If the *shower track* is not well reconstructed, or even just some hits misidentified, the method detailed above will often fail. It became clear during the building of the framework that the track reconstruction would also benefit from prior separation of hits, and that this problem must be solved at the hit, or even charge, level. Such fine scrutiny is exceptionally challenging and it was only when utilising machine learning tools that the issue appeared resolvable. As noted earlier, this is likely how the eventual solution will be found.

6.3.5 Performance of the Reconstruction

The performance of the reconstruction will be discussed in this section, with shower properties compared with true values (from the simulation) and with expectations. The validation sample includes the following particles (generated using a particle gun with no other event activity):

- Electrons, 0.1–5.0 GeV, 10 000 events, DUNE reduced far detector (10kt) geometry;
- Photons, 0.1–2.0 GeV, 10 000 events, DUNE reduced far detector (10kt) geometry;
- π^0 s, 0.4–1.0 GeV, 10 000 events, DUNE 35-ton geometry.

The same, standard reconstruction was applied to each sample and the resulting shower objects analysed to quantify its functionality.

6.3.5.1 Shower Properties

The efficacy in creating complete shower objects, including a reconstructed start point, direction, dE/dx , energy and associated hits, is demonstrated in Figure 6.20. It is clear the reconstruction performs excellently at all energies above the very lowest, where the least amount of information is available to the reconstruction. This is a very important feature of the reconstruction and great care was taken to ensure the reconstruction operates as effectively as possible. It should be noted that this is only for single particle events, and less reliable reconstruction should be expected when the detector contains multiple particles, but it is nonetheless important in demonstrating basic capabilities.

The quality of the reconstructed shower start point and direction is demonstrated in Figures 6.21 and 6.22 respectively. In both cases, it is clear in general the reconstruction performs very well at identifying the shower conversion point and extracting the relevant parameters. The small (arguably insignificant) spike around zero in the direction plots results from challenging shower topologies in which a start point is successfully found but there is insufficient information from multiple planes to make an accurate direction measurement.

The calorimetric information provided by the reconstruction, namely the shower energy and the dE/dx from the initial shower track, is evaluated in Figure 6.23 and 6.24 respectively. As with the previous validation plots, the reconstruction is shown to produce high quality shower objects. The energy completeness peaks at around 80% for both particle species demonstrating a good reconstruction efficiency and good energy conversion. The reasons for the slightly low peak are due to small issues with fragmentation in the pattern recognition stage, especially towards the lower energy end of a shower, with more sparsely distributed hits, and misunderstandings in the shower energy conversion (described in Section 6.2.3.2).




Fig. 6.20 The fraction of total shower particles for which a shower object is created when using BlurredCluster/EMShower reconstruction. A complete reconstructed shower object is required to have a reconstructed start point, direction, dE/dx , energy and associated hits.




Fig. 6.21 The difference between the true and reconstructed shower conversion points. The resolution for both electron and photon showers is around 0.1%.

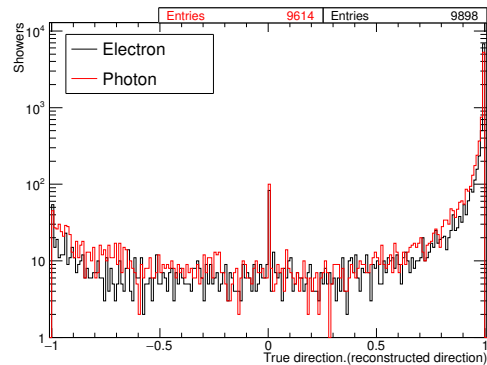


Fig. 6.22 The vector dot product between the true and reconstructed initial shower direction. The angular resolution for electrons is 19% and for photons is 23%.




Fig. 6.23 The completeness of the reconstruction shower energy when compared with the true deposited energy from simulation. The energy resolution for electrons is 25% and for photons is 28%.




Fig. 6.24 The dE/dx information from the start of the reconstructed shower object. The resolution of the dE/dx is 2% for electrons and 10% for photons.

The dE/dx ionisation information shows very nice separation between electrons and photons, as expected, with the peaks occurring at the anticipated values.

6.3.5.2 π^0 Reconstruction

As the primary motivations for developing the shower reconstruction were to demonstrate successful π^0 reconstruction, it is instructive to consider the invariant π^0 mass peak. This is determined from the two decay photons using

$$m_{\pi^0} = \sqrt{4E_{\gamma 1}E_{\gamma 2} \sin^2\left(\frac{\Delta\theta_{\gamma 1,2}}{2}\right)}, \quad (6.6)$$

where $E_{\gamma n}$ is the energy of each photon and $\Delta\theta_{\gamma 1,2}$ is the angle between the two photon trajectories.

The invariant mass is determined for all π^0 s where both decay photons are considered well reconstructed (i.e. populate the plot shown in Figure 6.20) and the distribution is shown in Figure 6.25. The reconstructed π^0 mass appears very reasonable when compared to the true value of 135 MeV. This is a particular challenging reconstruction problem and, when developed via BlurredCluster/EMShower, became the first time it had been solved in LArSoft.

The fully reconstructed π^0 mass peak differs from expectation in two principal ways: it peaks lower than the π^0 mass and is relatively broad. Further analysis is possible by substituting reconstructed quantities for the equivalent true values in the determination of




Fig. 6.25 The 35-ton fully reconstructed π^0 invariant mass peak from simulation. The mean of the distribution is 0.120 GeV with a width of 0.066 GeV.

the mass. This is demonstrated in Figure 6.26, where the photon energy and direction are successively replaced by their associated true properties. It may be observed the energy reconstruction causes the peak at a lower mass and the direction reconstruction is responsible for the wide distribution. This is entirely consistent with observations from the shower reconstruction discussed in Section 6.3.5.1.

6.3.5.3 Track/Shower Separation Performance

The simple track/shower separation algorithm discussed in Section 6.3.4 is assessed in this present section. In order to quantify the performance of the reconstruction, the following metrics are defined:

- *basic track/shower separation*: the electron shower must be reconstructed as a shower, the longest hadron track originating from the vertex must be identified as a track;
- *full track/shower separation*: the electron shower must be reconstructed as a shower, all hadron tracks from the vertex must be identified as tracks;
- *good shower*: the reconstructed shower start point must be within 10 cm of the true start, the reconstructed direction within 10° and the shower must be at least 50% complete;




Fig. 6.26 The 35-ton π^0 mass peak considered using a combination of reconstructed and truth information.

- *good reconstruction*: the event must have basic track/shower separation and the electron shower must be classified as ‘good’;
- *very good reconstruction*: the event must have full track/shower separation and the electron shower must be classified as ‘good’.

Without strictly requiring perfect reconstruction, these metrics have been shown to represent the status of the software reasonably well. The validation utilised 50 000 DUNE ν_e interactions in a reduced far detector volume, intended to represent as closely as possible the oscillation appearance signal.

The efficacy of separating the electron shower from the main hadronic tracks from the vertex, and the reconstruction performance of the electron shower, is demonstrated in Table 6.1. It is clear the challenge is being able to produce well separated particle objects whilst maintaining the shower qualities. In general the reconstruction performs acceptably but is far from the final solution to this problem. The fraction of events which were well reconstructed, as a function of true neutrino energy, is demonstrated in Figure 6.27. The reconstruction performs consistently except at the very lowest energies, when shower completeness becomes the main issue, but only at around 50% efficiency. This performance was comparable to, and even slightly superseded, other solutions within LArSoft when developed, but still requires significant improvements. Much effort has been made recently

Table 6.1 The performance of track/shower separation and shower reconstruction on 50 000 DUNE ν_e events, classified using the metrics discussed in the text. The number of events categorised into each sample is shown, along with the proportion of the total number of ν_e CC events represented. In the case of poor shower reconstruction, the particular features which failed the requirements are also detailed.

	Number of events	Proportion of events
ν_e CC events in FV	43 942	100%
Good shower	23 023	52%
<i>Poor shower – start point</i>	17 891	-
<i>Poor shower – direction</i>	9 831	-
<i>Poor shower – completeness</i>	6 367	-
Basic track/shower separation	36 200	82%
Full track/shower separation	32 785	75%
Good reconstruction	20 399	46%
Very good reconstruction	18 971	43%

within the LArSoft collaboration to address this issue directly and good progress has been made. The experience here has been instructive in how to approach the topic but has demonstrated the successful solution will be attained using a different approach.




Fig. 6.27 The performance of the track/shower separation and shower reconstruction when applied to DUNE ν_e CC far detector interactions. The fraction of events which are classified with ‘good reconstruction’, as defined in the text, is shown as a function of true neutrino energy.

Chapter 7

Analysis of 35-ton Data

The 35-ton run, discussed in Section 4.3, provided 22 days of good quality (high purity, stable field (250 V/cm), stable DAQ), analysable data. Due to the issues encountered, high quality physics analyses proved very challenging and instead studies, particularly those presented here, focused on trying to understand the detector and characterise previously untested responses. In this respect, the 35-ton proves to be a vital experiment in informing the next generation of prototypes and even the final DUNE far detector design. It also boasts unique datasets which no other planned experiment will possess before the full DUNE modules; it is therefore essential as much information as possible is extracted from the 35-ton analyses.

Before analyses are presented, techniques developed to enhance the quality of the data, and the data selection, will be discussed in Section 7.1. The main studies, concerning tracks passing across APA gaps and through the APA frames, are presented in Section 7.2 and Section 7.3 respectively before a brief investigation into the performance of basic shower and calorimetric reconstruction on the 35-ton data is discussed in Section 7.4. A summary is provided in Section 7.5.

7.1 Preparing 35-ton Data for Analysis

To ensure analyses are as accurate as possible, careful preselection and preprocessing of the data is performed. Methods for producing the analysable sample are discussed in this section.

7.1.1 Selecting the Data

The level of noise present in the TPC data varied significantly between runs and is immediately evident when analysing the RMS of the charge read out on a particular channel. Figure 7.1 shows a comparison of this metric for ‘good’ and ‘bad’ runs. In this case, the




Fig. 7.1 Comparison between noise levels for ‘good’ and ‘bad’ 35-ton runs. The channels shown are on APA2 (online convention, APA0 offline) and are read out by RCEs 8 through 11 (labelled). The increase in read out charge RMS is evident in the case of the noisy run. These plots are from runs 15797 (Fig 7.1a) and 15790 (Figure 7.1b) and were taken only 50 minutes apart.

figures show a ‘normal noise’ run and one taken with the detector in the ‘high noise’ state. The noise worsened as the experiment progressed but, in general, runs taken with normal noise and most of the TPC operating during the period of stable data-taking were used to assemble a ‘good run’ list by the collaboration.

In all there were 1269 runs used for analysis, containing some data taken before the FNAL site wide power outage (3rd March 2016) with most the week after stabilising the experiment again (9th March – 17th March 2016), and representing around 80% of all runs from the period. A selection of bad channels, classified as either ‘dead’ (electrically) or ‘bad’ (exhibit sufficiently more than average noise), constitute 13–25% of the total number of channels [176].

Due to the continuous nature of data taking, there is a non-trivial correlation between a ‘DAQ event’, a collection of fragments read out by the DAQ, and a ‘physics event’, an event in which particle interactions occurred. The external triggers used in the 35-ton, namely the external muon scintillators and the photon detectors, are used to define the event time. Given the trigger rate at which most data was taken (~ 1 Hz), a typical run comprising a few thousand events will only contain $\mathcal{O}(10)$ triggered events. Furthermore, as described in Section 4.3.2.2, these events often straddle multiple DAQ events (previously demonstrated in Figure 4.22), requiring the use of a splitter/stitcher module to search for triggers within runs and construct physics events containing the information useful for analysis.




Fig. 7.2 Simulated demonstration of the method used to correct for stuck codes in the 35-ton data [177]. On a given channel, ADCs exhibiting the consequences of this problem are corrected by interpolating charge at neighbouring time units. This is tested by simulating a waveform and applying the observed stuck code effect, demonstrated by the black and red waveforms respectively. The efficacy of the developed algorithm at correcting the afflicted bits may then be evaluated, shown by the blue trace.


7.1.2 Improving Data Quality

Two issues present in the raw data, namely the presence of correlated noise and the stuck codes in the digitiser (both described in Section 4.3.5.1), are dealt with as an initial step of the reconstruction. Initially, an algorithm attempts to correct for the stuck bits analyses waveforms on a wire and identifies problematic ADCs; interpolating between charges read out at neighbouring times is successful at reconstructing the initial waveform in most cases. Figure 7.2 demonstrates this interpolation method on simulated data. The effect of applying this algorithm on a full waveform, to correct for all the stuck bits, is apparent in Figure 7.3.

Following this process, a coherent noise removal stage is applied. This simply looks at the average noise across channels sharing a front-end voltage regulator and removes this component from the readout ADC for each channel. The effect of this correction is seen in Figure 7.4.

7.1.3 Reconstructing Muon Tracks


All analyses discussed below only make use of information recorded on the collection planes. Since the induction wires are longer (a necessity for wrapping), a larger capacitance



(a) Raw waveform before correcting for stuck bits.

(b) After applying stuck bit mitigation.

Fig. 7.3 The effect of applying stuck bit mitigation to a waveform as seen in raw 35-ton data. This particular waveform is from run 15660, channel 722 (induction channel).



(a) Waveform before removing coherent noise.

(b) After removing coherent noise.

Fig. 7.4 The effect of removing coherent noise from all channels on a voltage regulator in the 35-ton data. This waveform is from run 15660, channel 2010 (collection channel). As a result of this process the noise (RMS of ADC values) decreases from 37.1 to 18.8 leaving a signal which is noticeably larger, considerably improving reconstruction performance.

results in higher noise levels, complicating the reconstruction. In general, after applying the refinements outlined in Section 7.1.2, the signals on the collection channels are prominent enough for competent analyses. The methods used to select tracks are described in this section and applied during the subsequent studies.

Using only the collection plane presents challenges, the most obvious being the impossibility of full 3D reconstruction. A hit on a collection wire at a given time gives well-defined x and z coordinates but cannot give any information in the y -direction. ‘Quasi-3D’ reconstruction is achieved by making use of the external counters. Through-going muons are triggered by the coincidence of hits in two opposite counters; this information can be used to give a crude handle on the y position of hits.


Figure 7.5 outlines the stages of selecting hits originating from the particle track which caused the trigger. Figure 7.5a shows all hits from an example event containing a through-going muon. The first stage of track selection involves taking those hits which lie in the ‘counter shadow’, the narrow section of the collection plane area physically in between the opposing counters through which the triggering particle passed. The hits which remain are shown in Figure 7.5b. The track hits are visible along with further, unrelated hits. These are removed by requiring that only hits on wires with single occupancy be kept, and then applying a linear fit and removing all hits with residual > 2 cm. The final output after these stages is shown in Figure 7.5c.

The result of this track selection, as evident from Figure 7.5c, is a well-formed, high quality track with which it is possible to perform analyses. These will be the focus of the remainder of this chapter.


7.1.4 Preparing Simulated Data

Comparisons with simulated data are often essential in understanding various phenomena in the data. Throughout the analyses presented in this chapter, simulations were used to aid investigations and therefore it is important to ensure the Monte Carlo is as similar to the real data as possible.


The standard LArSoft simulation tools were used as described in Section 6.1, employing the CRY cosmic ray generator [178]. The data passing through the detector were filtered on counter coincidences, exactly as the raw data is triggered. The simulated data was then processed in the same way as the real data and reconstructed using the methods described in Section 7.1.3.



(a) All hits before any track selection. The red lines represent the boundary defined by the edges of the two counters causing the trigger.



(b) Hits in the counter shadow.



(c) Hits on single wire occupancy and with residual < 2 cm.

Fig. 7.5 Demonstration of the successive stages applied to hits on collection wires in the 35-ton data in order to select hits from the through-going track associated with the particle which caused the trigger. The hits left after all stages are taken forward into the analyses.

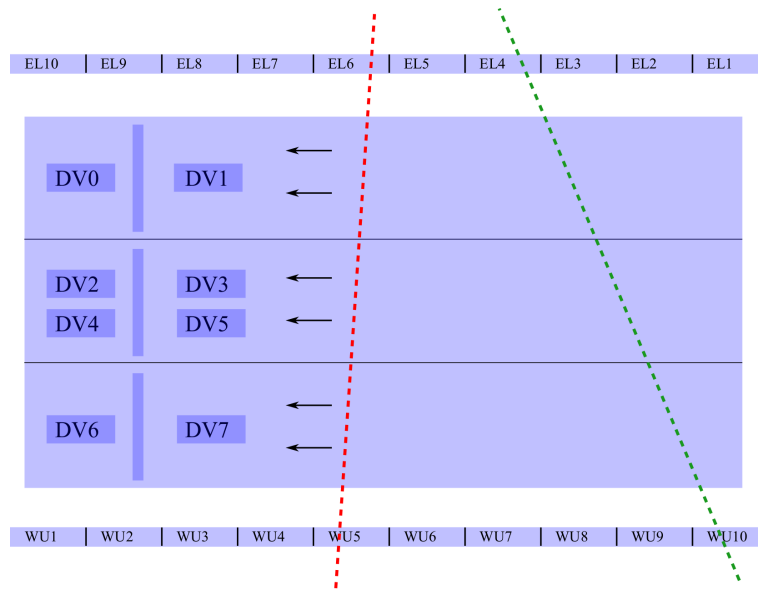


Fig. 7.6 Demonstration of crossing muon tracks in the 35-ton TPC. The view is from the top of the cryostat looking down, with EW counters shown on either side of the TPC. The APAs, defining the eight DVs, are represented by the darker rectangles towards the left of the detector. The ‘counter gradient’ is defined as the angle a track makes with respect to the face of the APAs. Two example tracks are shown; the red track has a counter gradient of zero whilst the green track has a gradient of three.

7.2 APA Gap-Crossing Muons

One of the primary motivations for the design of the 35-ton TPC was to test its modular form, where a single drift region is read out by multiple anode assemblies. Particles passing through the detector will leave deposits in multiple DVs and will pass uninstrumented regions of the detector, such as gaps in between neighbouring APAs; some example tracks are demonstrated in Figure 7.6. Many APA gap-crossing tracks are evident from the event display in Figure 7.7 and an example of such a track is depicted schematically in Figure 7.8. It is essential the implications of this design choice are understood before constructing the far detector modules, each of which will contain 150 APAs.

The 35-ton dataset consisting of muons which pass across the face of APAs and therefore deposit charge in consecutive DVs is discussed in this present section. An analysis of these tracks to calculate the size of the gaps is presented in Section 7.2.1 and a study of the charge deposited by such tracks is the subject of Section 7.2.2.




Fig. 7.7 Event display showing tracks passing across APA gaps and also through the APAs. A study of the tracks which pass across gaps between the APAs (the red lines) is the subject of Section 7.2. There is a visible offset apparent as the track crosses through the APAs (the blue line); correcting for T0 would eliminate this and yield a single accurately connected track. This is discussed further in Section 7.3.




Fig. 7.8 Schematic showing an example APA gap-crossing track as viewed looking down from the top of the detector. The vertical direction represents the drift direction (x); the horizontal direction represents the z -direction. In general, these tracks make an angle with respect to the face of the APAs, as shown in the figure. As the gap in between the APAs is uninstrumented, no charge is deposited in this region.

7.2.1 APA-Gap Offset Determination

It is possible to use these gap-crossing tracks to make measurements of the gaps between each of the APAs. This involves aligning the track segments from neighbouring DVs, demonstrated in Figure 7.9. The value of the z -offset, Δz , is determined by considering a range of offset hypotheses, performing a linear fit and finding the offset which minimises the residual least squares

$$L = \sum_i^{nhits} (o_i - e_i)^2, \quad (7.1)$$

where $o_i - e_i$ is the distance from hit i to the best fit line.

There are eight gaps which can be measured from the data, demonstrated in Figure 7.10. Due to very low statistics, it was found measurements of the gaps on the short drift volume side of the APAs were not possible using the 35-ton data. Analysis of the gaps using tracks passing through the long drift volume, hereafter named DV1/DV3, DV1/DV5, DV3/DV7 and DV5/DV7, was therefore the focus of this study.

A number of cuts were applied to ensure only high quality tracks were included for analysis:

- Only hits greater than 1 cm and less than 15 cm away from the gap were included in the track segments. The purpose of this cut is to limit the effect of multiple scatterings and the poorly understood region closest to the gap, where charge deposited in the uninstrumented region may later be collected.




Fig. 7.9 Schematic showing an example track crossing two drift regions offset by an unknown quantity Δz . The effect of this is evident from the track deposits (Figure 7.9a) and can be corrected by ensuring the segments are aligned between the DVs (Figure 7.9b).




Fig. 7.10 Illustration of the eight gaps between the four APA frames. There are four spaces between the APAs giving rise to eight gaps, measured in the long drift (odd-numbered) and short drift (even-numbered) regions. In the figure, the even-numbered gaps represent the distance between the APAs on the back-facing side and the odd-numbered gaps the distance on the side facing. Naively, one would expect the gaps to be identical but this will not necessarily be the case when assembled in the cryostat and, since they read out data from distinct drift regions, they may be considered separately.

- Each track segment must contain at least ten hits to allow an accurate measure of the gradient.
- The angle between the track segments either side of the gap must be less than 2° to remove any poorly reconstructed tracks, or segments originating from different particle tracks.
- The angle the track makes with respect to the APA face must be large enough that the gap offset effect can be measured to an acceptable accuracy. It is common in the 35-ton to refer to a ‘counter gradient’, the offset between the two counters forming the through-going particle trigger in the drift direction, in units of counter length (demonstrated in Figure 7.6). The tracks must have a counter gradient of at least three.


7.2.1.1 Measuring the APA Gaps

The gap which possesses the largest number of crossers is DV5/DV7 and so the procedure will be demonstrated using data from this channel. The z -offset determined using the method and cuts described above is shown in Figure 7.11. An unexpected feature is evident from this distribution; there is not a single peak but two, seemingly related to the angle which the through-going particle makes with respect to the APAs.


One explanation for this observed double-peak effect involves considering the possibility of additional offsets from the assumed positions of the APAs. This is demonstrated in Figure 7.12. It appears an offset in the x -position of the APAs could result in the problems encountered in the data. In order to test this, these offsets were artificially introduced into the simulation; the findings are presented in Figure 7.13. It appears the distribution of Δz measured from the data is consistent with APAs with offsets from expectation in both x and z . Moreover, it may be possible to measure both offsets from the same data set.

The offsets may be determined using a joint χ^2 fit, in Δz - Δx space, demonstrated in Figure 7.14. However, since hits are absolute and do not possess inherent positional uncertainties, it is particularly challenging estimating errors on the χ^2 values, resulting in meaningless experimental uncertainties for the measured offsets. Instead, since it is evident the offset values are uncorrelated, the process may be separated into two stages to ensure the uncertainties are correctly estimated. This method will be the subject of the remainder of this section; the obtained values, in general, agree excellently with those extracted from simultaneous 2D fits.

It is clear from Figure 7.13 that the z -offset may be determined as the minimum between the angular-separated distributions. This can be justified by geometrical considerations,



(a) Full distribution.



(b) Separated by the angle the track makes to the APAs.

Fig. 7.11 The z -offset for the DV5/DV7 gap measured in the 35-ton data. A very noticeable double-peak structure is evident in Figure 7.11a; this bias appears to be related to the sign of the angle the particle track makes to the APA planes, demonstrated in Figure 7.11b.

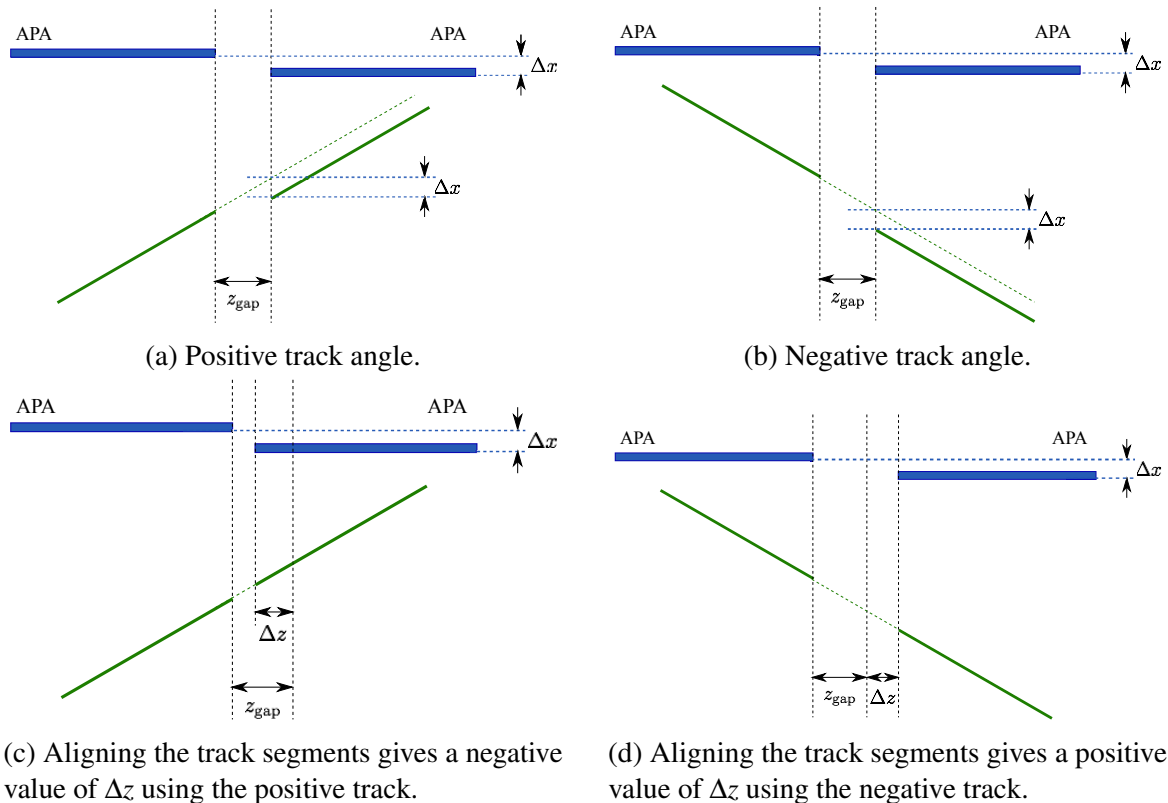


Fig. 7.12 Demonstration of how an x -offset in the positions of the APAs can explain the degeneracy evident in the z -offset measured using the 35-ton data (Figure 7.11). In the left-hand plots, Figures 7.12a and 7.12c, the through-going particle makes a positive angle to the face of the APAs and in the right-hand plots, Figures 7.12b and 7.12d, the particle is travelling with a negative gradient. In both cases, the offset of the APAs in the x -direction is the same. It is clear from Figures 7.12c and 7.12d how the sign of the measured Δz is dependent on the angle of the track.




(a) z -offset = 2 cm, x -offset = 0 cm.(b) z -offset = 2 cm, x -offset = 0.5 cm.(c) z -offset = 2 cm, x -offset = 0.5 cm.

Fig. 7.13 Studies of the effects of offsets in the positions of the APAs in simulation. Artificial z - and x - offsets are introduced and their impact observed in the measurements of Δz . Figure 7.13a shows the effect of an offset in the z -direction; as expected, there is a single peak measuring the input value. Figures 7.13b and 7.13c show the consequence of offsets in both the x - and z -directions. This appears to show exactly what is seen in the 35-ton data (Figure 7.11).




Fig. 7.14 The χ^2 distribution for all APA-gap traversing tracks in Δz - Δx space.

explained in Figure 7.15. In this case, this may be achieved by fitting a function of the form

$$f(x) = a(x - b)^2 + c \quad (7.2)$$

and extracting parameter b as the true value of Δz . This is shown in Figure 7.16.

Using this measured value of Δz , the offsets can be analysed again, this time measuring the x -offset by correcting for the z -offset. The measured x -offset distribution is shown in Figure 7.17. With this value of Δx , the z -offset can be evaluated once more to ensure the distribution contains a single peak, as initially expected. This is confirmed in Figure 7.18.

7.2.1.2 Measurements of the APA Offsets

The offsets apparent from the data for all of the gaps accessible using TPC tracks in the long drift volume were determined as described in Section 7.2.1.1. Appendix A contains all figures relevant for each gap measurement. Table 7.1 details all the measurements, and the new gaps are presented in Table 7.2 after taking all the offsets into account. The determined errors are statistical only; the effects of systematic uncertainties were not considered and assumed to be negligible in comparison.

Since there are two longer APAs either side of two short APAs, by combining the gap offset information from the gaps either side of the centre frames it is possible to make two implicit measurements of the offsets between the outer ones. This is shown in Table 7.2 for both the x - and z -offsets, with the number in brackets representing the DV across which the

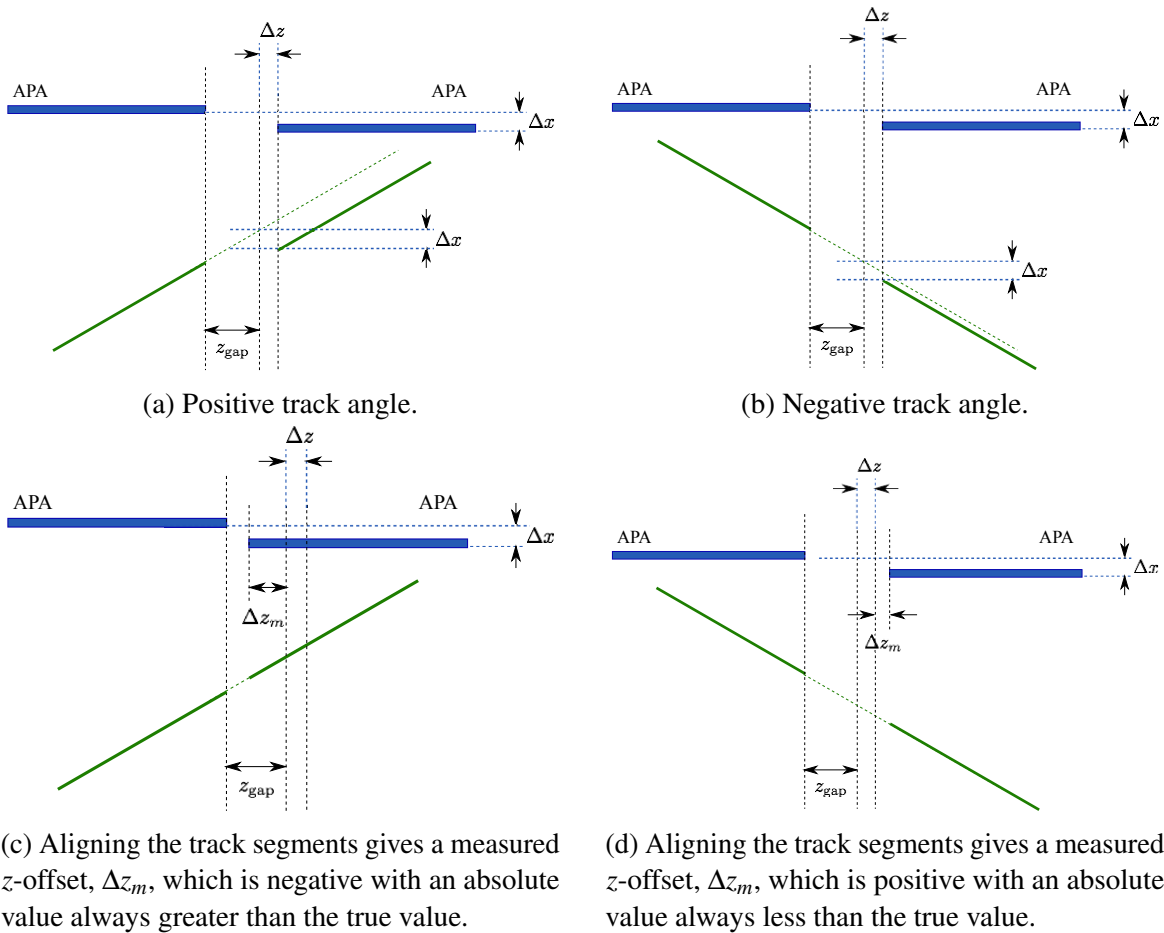


Fig. 7.15 Demonstration of the effects of offsets in both the x - and z -directions in the determination of Δz between DV5 and DV7. With an x -offset present, it is impossible for the true value of Δz to be measured – this is evident from Figure 7.13. It is clear from these geometrical considerations how the measured offset Δz_m will populate distributions either side of the true value; the true value Δz is given by the minimum between the two distributions.




Fig. 7.16 Extraction of the true value of Δz from the full distribution of measured z -offsets. A measured value of 0.117 ± 0.007 cm is found.




Fig. 7.17 Measurement of the x -offset between DV5 and DV7 after applying the z -gap correction determined using the method described in the text and Figure 7.16. A measurement of -0.286 ± 0.002 cm is determined.




Fig. 7.18 Measurement of the z -offset between DV5 and DV7 after applying the x -offset determined from Figure 7.17. As initially anticipated, there is a single peak distributed around the true value of the offset. This validates the method used and confirms the initial presence of an x -offset between the neighbouring APAs. The final measurement of Δz is 0.103 ± 0.004 cm which agrees reasonably with the value measured previously (0.117 ± 0.007 cm from Figure 7.16).

Table 7.1 Measurements of all the APA offsets determined from the 35-ton TPC data. The method followed is described in Section 7.2.1.1. The first row represents the initial measurements of the z -offset from the two-peak distribution, with the following two lines detailing the measured offsets that follow from these results.

	DV1/DV3	DV1/DV5	DV3/DV7	DV5/DV7
Initial z -offset (cm)	-0.64 ± 0.04	0.15 ± 0.01	0.58 ± 0.06	0.117 ± 0.007
x -offset (cm)	-0.377 ± 0.006	-0.252 ± 0.002	-0.16 ± 0.01	-0.286 ± 0.002
z -offset (cm)	-0.63 ± 0.02	0.131 ± 0.007	0.55 ± 0.03	0.103 ± 0.004

Table 7.2 The corrected gaps between the APAs, in x and z , based on the offsets measured in the data (Table 7.1).

	Assumed (cm)	Offset (cm)	Corrected (cm)
DV1/DV3 x -gap	0	-0.377 ± 0.006	-0.377 ± 0.006
DV1/DV5 x -gap	0	-0.252 ± 0.002	-0.252 ± 0.002
DV3/DV7 x -gap	0	-0.16 ± 0.01	-0.16 ± 0.01
DV5/DV7 x -gap	0	-0.286 ± 0.002	-0.286 ± 0.002
DV1/(3)/DV7 x -gap	0	-0.538 ± 0.003	-0.538 ± 0.003
DV1/(5)/DV7 x -gap	0	-0.537 ± 0.010	-0.537 ± 0.010
DV1/DV3 z -gap	2.08	-0.18 ± 0.02	1.90 ± 0.02
DV1/DV5 z -gap	2.08	0.131 ± 0.007	2.211 ± 0.007
DV3/DV7 z -gap	2.08	0.10 ± 0.03	2.18 ± 0.03
DV5/DV7 z -gap	2.08	0.103 ± 0.004	2.183 ± 0.004
DV1/(3)/DV7 z -gap	4.16	-0.08 ± 0.04	4.08 ± 0.04
DV1/(5)/DV7 z -gap	4.16	0.23 ± 0.01	4.39 ± 0.01

total gap offset was determined. There appears to be some consistency in the measurements of the x -offsets between DV1 and DV7, found by considering the successive offsets between DV1/DV3 and DV3/DV7, and DV1/DV5 and DV5/DV7. An exceptional agreement is seen between the two values, within 0.1 mm. There also seems to be slight evidence of a rotation between DV1 and DV7 when considering the associated z -offsets; the offset at the top of the APA (when measured via DV5) is greater than at the bottom (when measured via DV3).

The method demonstrated here will have direct implications for similar studies using the full DUNE far detector. All the gaps between the APAs, both in the drift and z directions, will need to be understood for accurate reconstruction and are essential in order for DUNE to make precise physics measurements. For example, accurate calorimetric reconstruction is imperative in order to perform particle identification and shower energy determination (discussed in Section 6.2.3) and is directly related to the drift time of the ionisation electrons; any offsets in APA positions will lead to systematic uncertainties in this information.

7.2.2 Charge Deposited by APA Gap-Crossing Muons

The charge deposited by gap-crossing particles cannot be collected in the dead regions between the APA frames. It is interesting to consider where the charge is read out in order to further understand the implications of a modular TPC design.

Figure 7.19 demonstrates the properties of hits as a function of distance from the nearest DV edge. It appears more hits are found as charge is collected near a gap but the charge of these hits do not differ significantly. This may be interpreted as hits arriving at a slightly later time near the APA gaps after drifting towards the nearest wire to the gap from a more centre-gap position. One may expect to observe this in the data as a smearing in the tick direction where charge is deposited over more time, leading to a small gradient change. Although not as noticeable as anticipated, this effect is observable in the event display shown in Figure 7.20.

7.3 APA-Crossing Muons

The 35-ton is the only proposed experiment before the full DUNE far detector modules that has fully implemented anode planes within the cryostat reading out data from multiple drift regions simultaneously (ProtoDUNE will have wrapped wire APAs but will only read out one drift region each and SBND has the CPAs in the centre of the cryostat with the APAs at the edges). Referring to Figure 3.10, this is a design consideration that features prominently in the eventual detector so any implications in the data must be well understood. Analysis of

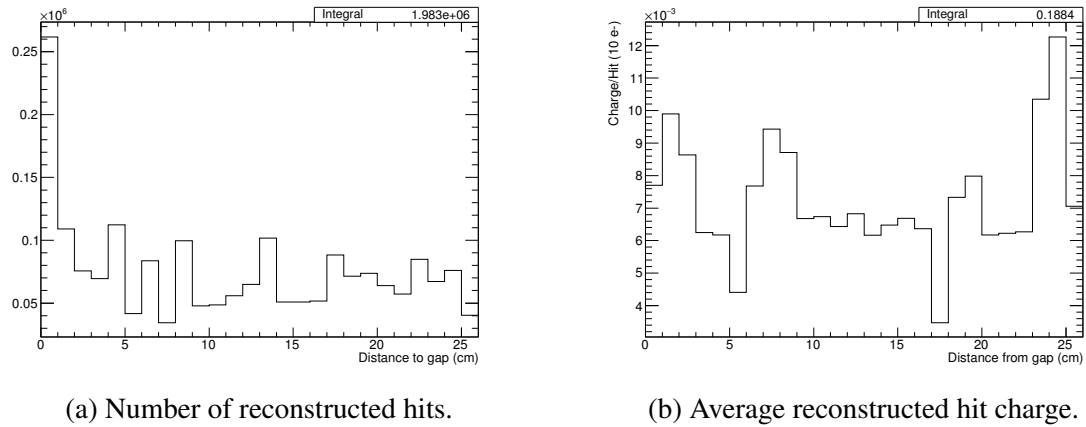


Fig. 7.19 The number of hits, and the average reconstructed hit charge, as a function of distance of the collection point from the nearest APA gap.




Fig. 7.20 Event display of an APA gap-crossing track, focussed on the gap region. Charge arriving at the centre of a gap deflects toward the nearest wire and is collected at a slightly later time. This results in more charge being deposited on wires nearest the gap, with a larger spread in time. This is subtly observable in the charge distributions shown here.

tracks which pass through the APAs and deposit charge in both drift regions is the subject of this section.

In Section 7.3.1, a method to determine the absolute event time, T_0 , from APA-crossing tracks is presented and in Section 7.3.2 the charge deposited by these tracks, particularly when crossing through the planes, is studied. Comparisons between the two drift regions, made possible by comparing tracks left by the same particle, are contained in Section 7.3.3.

7.3.1 T_0 Determination from APA Crossing Tracks

Given the nature of a TPC detector, an ‘event time’ (T_0) must be known in order to set an absolute timescale, and therefore absolute position, on all interactions within the detector. An accurate T_0 is essential for calorimetric reconstruction: in order to understand how much charge a hit had when it was created, a lifetime correction dependent on the total drift time must be applied. An incorrect T_0 would lead to a systematic under- or over-estimation of the reconstructed energy and have implications in particle identification and shower energy determination.

In a LArTPC, an event time is usually given by an external triggering system. The DUNE far detector will rely on the instantaneous detection of photons produced from the immediate recombination of the ionisation electrons with positive Ar ions. In the 35-ton, an additional external system was provided by the scintillation counters. Since the sample of APA-crossing muons used in this analysis were all selected and reconstructed using counter information, an interaction time is immediately known.

Without correctly accounting for T_0 , the tracks on each side of the APAs appear offset from the planes. This is evident from the event display shown in Figure 7.7. By aligning the track segments on either side of the APAs, a measurement of T_0 can be made directly from the TPC data.

7.3.1.1 Aligning APA Crossing Tracks

Two complementary methods were used to accurately align the track segments across the APA. Both involved initially correcting for the counter T_0 , T_0^{counter} , before considering a range of alternative T_0 hypotheses and minimising a relevant metric to determine the most likely value. In the first method, demonstrated in Figure 7.21, a least square linear fit is applied to the track and the residual minimised (the ‘residual method’). The second method, demonstrated in Figure 7.22, involves fitting a line to each segment in turn and minimising the projected distance between the intersections of the lines with the centre of the APAs ($x = 0$)

(the ‘separation point method’). As will be shown, and can be seen from Figures 7.21b and 7.22b, the two methods agree very well with each other.


Naively, one would expect the T0 determined using these methods, T_0^{TPC} , to agree with T_0^{counter} . This is confirmed by applying the analysis to simulated data and demonstrated in Figure 7.23a. However, there appears to be a systematic offset between the T0 given by the counters and measured from the TPC data in the 35-ton. The distribution of this discrepancy is shown in Figures 7.23b and 7.23c for each of the two methods described; it peaks around 61 ticks ($30.5 \mu\text{s}$) and is importantly incompatible with zero. This suggests an inconsistency somewhere in the data taking and attempts to understand this track misalignment will be the subject of the remainder of this section. Figure 7.24 shows an example track before and after this disparity is corrected for. As is evident from Figure 7.23, the separation point method provides more consistent results so this will be used exclusively for alignment measurements in the rest of this section.

7.3.1.2 Understanding the Misalignment of APA-Crossing Tracks

The underlying issue described above is essentially a misalignment of the same particle track between the two drift regions, demonstrated plainly in Figure 7.25. This obviously implies a misunderstanding somewhere in the data processing and stems from an issue with the detector or data readout. The most obvious cause is a miscalibration of the DAQ timing systems for the separate detector components, as previously assumed. There are however other possible solutions to the problem, and it is likely the effect arises from a combination of different factors.

Geometry Apart from timing, a misunderstanding of the geometry could explain this perceived misalignment. The spacing between the collection planes is one such example, as demonstrated in Figure 7.26a; the spacing necessary to explain this effect, determined by aligning the tracks using the methods discussed above over a range of collection plane spacing hypotheses, is demonstrated in Figure 7.26b. As is evident from the figure, the collection planes must be repositioned in such a way that they would be reversed; the track alignment complications cannot be explained solely by this.


A further problem is related to the wire positioning on the APAs in the z -direction; it is understood there may be a discrepancy between the two sides of the APA resulting in hits from the long and short drift regions at the same z -position reconstructed with a systematic offset. Figure 7.27a shows how this could be utilised to explain the apparent track misalignment with Figure 7.27b showing the distribution of corrected z positions necessary to resolve the issue. Offsets of ~ 2 cm, as suggested by these results, are highly unlikely



(a) Demonstration of the calculation of residuals from a linear fit through all hits. The red points are hits and the green line represents a linear fit through all points on both sides of the APA.

(b) The residuals to the linear fit of the track over a range of T0 candidates. The value of T0 which minimises this distribution (62 ticks in this case) is considered the most likely interaction time.


Fig. 7.21 Method to align track segments on either side of the APAs involving minimising residuals from a linear least square fit. A fit is applied to all hits and the resulting residual, a representation of the 'goodness of fit', is minimised over a range of T0 candidates to find the most likely interaction time for the particle leaving the track.




(a) Demonstration of the determination of the distance between the track segments at the centre of the APAs. The red and green lines represent linear fits to the hits (applied separately on each side of the APA) for different values of T0.

(b) The separation distance over a range of T0 candidates. The value of T0 which minimises this distribution (63 ticks in this case) is considered the most likely interaction time.


Fig. 7.22 Method to align track segments on either side of the APAs involving minimising the distance between the projected intersection of each with the centre of the APAs. A fit is applied to each track segment separately and the distance between the intersection of these lines with the centre of the APA is minimised over a range of T0 candidates to find the most likely interaction time for the particle leaving the track.



(a) 35-ton simulation. The difference in the two measurements of T0 is distributed around zero, as expected, and validates the method. The peak is actually at 1 tick, indicating a slight systematic offset.



(b) 35-ton data using the separation point method.



(c) 35-ton data using the residual method.

Fig. 7.23 Difference between the T0 calculated from TPC data and the T0 provided by the counters representing the trigger time of the through-going muon, for simulation (Figure 7.23a) and data (Figures 7.23b and 7.23c). If the two measurements of T0 agree the distribution would peak around zero, confirmed in simulation; the fact this is not the case for data is indicative of a systematic offset somewhere in the data taking.

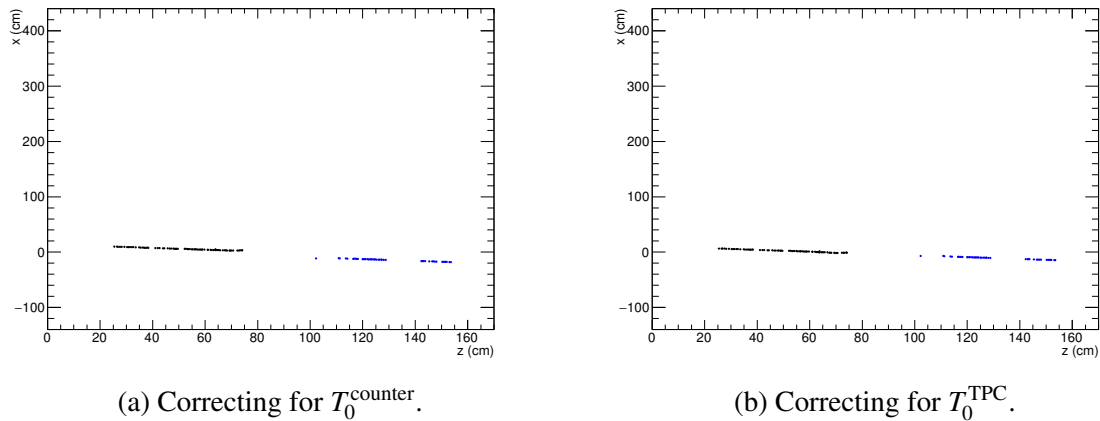


Fig. 7.24 Correcting for T0 using T_0^{counter} (Figure 7.24a) and T_0^{TPC} (Figure 7.24b). The difference is subtle but noticeable; the method for determining T0 directly from the TPC data can be validated by eye. The minimisation of the metrics to determine T_0^{TPC} in this case are demonstrated in Figs. 7.21b and 7.22b.




Fig. 7.25 Demonstration of the effect observed in the 35-ton data concerning tracks crossing the APAs. Even after correcting for the T0 provided by the counters, there is still a misalignment of the track segments across the APA frames.

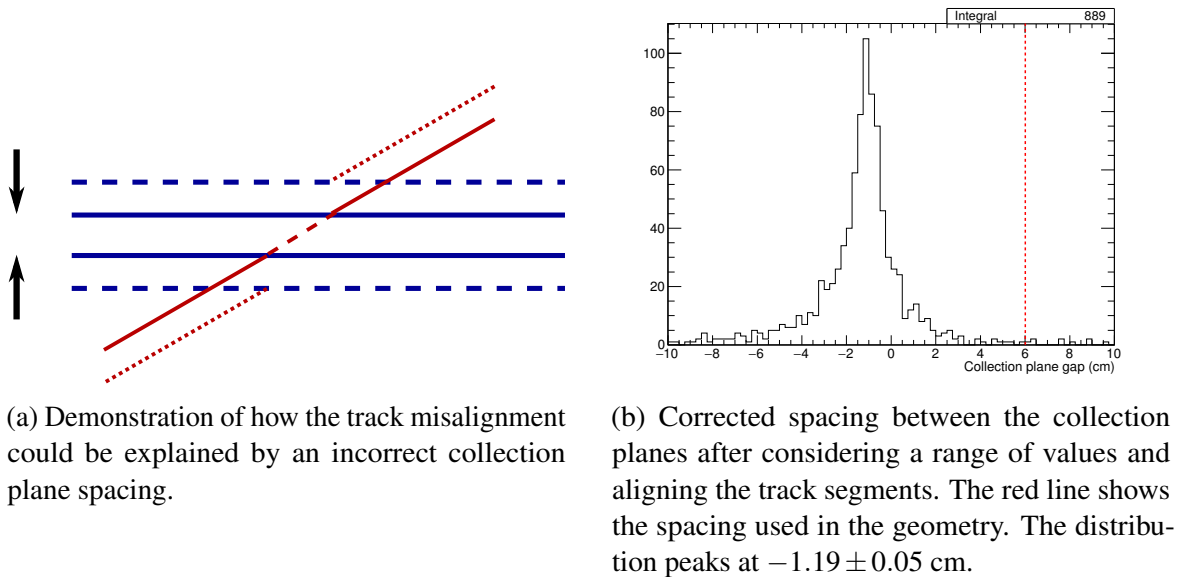



Fig. 7.26 Attempting to correct the track segment misalignment by assuming a misunderstanding of the spacing between the collection planes. It appears the resulting spacing necessary to correct for this issue would involve physically reversing the order of the planes.

given the scale of offsets identified in Section 7.2.1.2, indicating again the track alignment problem cannot be resolved in this way.

Drift velocity The drift velocity affects the angle of the tracks in wire/time space; a high velocity would result in a refraction-like effect towards the APA planes. As demonstrated in Figure 7.28a, this could explain the track segment misalignment if the effect was large enough. Figure 7.28b shows the necessary drift velocity required to account for the disparity observed in data; compared to a nominal value of 109 cm/ms, the scale of the change required to explain the oddity is unreasonably large, around a factor of five.

This can be tested by measuring the drift velocity directly from the data. Taking tracks which pass through opposite counter pairs (zero counter gradient) and comparing this drift distance with drift time is a trivial exercise, demonstrated in Figure. 7.29. The measured value of 110.2 ± 0.4 cm/ms agrees very well with the aforementioned value, determined theoretically, of 109 cm/ms. It may therefore be assumed the drift velocity is as expected and does not contribute at all to the track alignment anomaly.


Timing The timing offset of $32 \mu\text{s}$, calculated in Section 7.3.1.1, is so large it was assumed another explanation for the track segment misalignment was likely. However, after reviewing all possibilities it appears there must be a significant timing offset present somewhere in the



(a) Demonstration of how the track misalignment could be explained by an offset in the wire z -position on either side of the APA.

(b) Corrected z -positions of the APA wires after considering a range of values and aligning the track segments.


Fig. 7.27 Attempting to correct the track segment misalignment by assuming a misunderstanding of the positioning of the collection wires inside the detector. The wire offset would have to be around 2 cm to fix this issue.



(a) Demonstration of how the track misalignment could be explained by an incorrect drift velocity.

(b) Corrected drift velocity required to align the track across the APAs. The red line shows the assumed value of 109 cm/ms.

Fig. 7.28 Attempting to correct the track segment misalignment by assuming an incorrect drift velocity. In order to account for the effect noted in the data the drift velocity would have to be around five times larger than that initially calculated from models.



(a) Distribution of hit drift times for eight sets of counter pairs, assuming all tracks pass through the centres of the counters.


(b) The eight points found from taking the Gaussian mean of the time distributions for each rough drift distance.

Fig. 7.29 Measuring the drift velocity of the ionisation electrons by taking tracks passing through opposite counter pairs and comparing the corresponding drift distance to the drift time. Assuming all tracks pass through the geometric centres of the counters, a poor assumption, a distribution of hit time for this drift distance can be found; this is shown in 7.29a. Taking each counter pair separately and fitting a Gaussian to the distribution of drift times nullifies the assumptions necessary due to a lack of exact knowledge, on a track by track basis, of the exact x -position. This is shown in the graph in Figure 7.29b.

data. Further evidence for this hypothesis is presented in Figure 7.30 which displays the T_0 -corrected time distribution for all hits on the APA-crossing track. The minimum drift time these hits may have, since they pass directly through the planes, is the interaction time, T_0 . As is evident from the distribution in Figure 7.30b, this is around 60 ticks ($30 \mu\text{s}$) and is again notably inconsistent with zero. The curious spike at the interaction time motivates the work presented in Section 7.3.2 and will be discussed there. Additionally, it is possible to compare the T_0 provided by the counters with information from the photon detectors. This is shown in Figure 7.31 and provides further confirmation for a timing miscalibration in the TPC readout.

This interesting result provoked further investigation into the notion of a timing offset between detector components, specifically the TPC and counter readout (RCEs and PTB respectively). Confirmation of this miscalibration is displayed in Figure 7.32 which shows the difference between the timestamps recorded by each of the subcomponents upon receiving the trigger.

There are now three measurements of the timing offset with a slight disagreement between each. This will be discussed further in Section 7.3.1.3.



(a) Over the full range of drift times. The sharp dip around 500 ticks corresponds to the maximum drift time for hits in the short drift region; beyond this only hits in the long drift region contributes to the distribution.

(b) Zoomed in on the interaction time. The red line is drawn at 59.5 ticks ($29.8 \mu\text{s}$) and represents, by eye, the centre of the rising edge of the distribution.

Fig. 7.30 The T0-corrected drift time for hits on APA-crossing tracks. The lower leading edge of this distribution is an indication of the interaction time, T0.




Fig. 7.31 Difference between the interaction time measured by the TPC data and that provided by photon detector information. Only events with a single reconstructed flash are considered, with each assumed to have been caused by the triggering particle. This results in very few events, but clear supporting evidence of a timing offset on the order of 60 ticks is found.




Fig. 7.32 The difference between the timestamps recorded by the PTB and the RCEs upon receiving a trigger. The absolute timing for the DAQ system is given, along with most experiments at FNAL, by ‘NOVA time’: a 64 MHz clock starting on 1st January, 2010 (with one NOVA tick therefore being 15.625 ns). The distribution peaks sharply at 1705 NOVA ticks, or 26.6 μ s.

7.3.1.3 Combined Offset Analysis

The discussion in Section 7.3.1.2 hints strongly at an intrinsic timing offset present in the data. However, as already shown in Section 7.2, it is understood there are geometrical offsets in the positions of the APAs in the x - and z -dimensions. Attempting to measure all these offsets simultaneously presents challenges since they all affect each other. It is possible the tension between the measurements of the timing offset may be resolved by combining the results from each of the offset calculations.

The timing offset will not influence the determination of the geometrical APA gaps (found in Section 7.2) unless the track segments used to measure the gaps cross through the APA frames; the timing is consistent for each drift region. A simple cut was used to exclude such events when making these measurements. However, the geometrical offsets will have an impact on the APA crosser analysis. For example, the drift times measured for each hit will be affected by the physical positions of the APAs. Figure 7.33 shows the distribution of the drift times for all track hits corrected for the offsets implied by the x -gap measurements. It can be seen this accounts for the disparity between the previous measurements. It does not appear to agree completely with the offsets found between the timestamps but serves to demonstrate differences from the assumed positions of the APAs have a very sizeable effect on distributions such as these.

Correcting for the time offset of 53 ticks implied by the sub-component timestamp mis-calibration, along with those in the x - and z -positions of the APAs, does not entirely account for the initial inconsistency observed in Figure 7.23. A similar evaluation to that undertaken in Section 7.3.1.2, namely considering the required disparities in various quantities to account for this, may be used to facilitate a complete understanding. After correcting for the three aforementioned offsets, Figure 7.34 demonstrates the necessary misunderstandings in the collection plane spacing and the z -positions of the collection wires to account for the remaining discrepancy. It seems highly likely that the offsets between the APA gaps left unresolved in the short drift region, incalculable in the 35-ton data, can account for the outstanding misalignment between the track segments. Nothing conclusive can be extracted from Figure 7.34b with regards to the values of these uncertainties since this considers differences between all short drift region TPCs and long drift region TPCs together, but implies further offsets at a similar scale to those measured in the long drift region may still be present. With corrected APA gaps in the short drift region, it is reasonable to argue the track segment misalignment between drift regions would be completely resolved.

This is the first time tracks crossing the readout planes have been used in a LArTPC experiment and have proven to be a valuable way of calibrating inter-detector components and finding other inconsistencies in the data. Without studying this data set, the timing





Fig. 7.33 The distribution of the drift times of all hits on APA-crossing tracks after correcting for the APA offsets along the direction parallel to the drift direction, found in Section 7.2. The red line represents a T_0 of 53 ticks, representing the difference observed between the trigger timestamps between the scintillation counter and TPC readout systems. The hit time distribution appears to agree with this value to a greater extent than previously (Figure 7.30b).



(a) Assuming a misunderstanding in the spacing between the collection planes, a value of 4.74 ± 0.04 cm is measured. This is a difference of 1.27 ± 0.04 cm from the assumed spacing, a discrepancy which is highly unlikely.

(b) Assuming a misunderstanding in the alignment of the collection planes in z between the two drift regions, an offset of -0.24 ± 0.03 cm is found. Given the scale of the corrections determined in Section 7.2.1.2, and the incapability to measure the gaps in the short drift regions, this is eminently credible.

Fig. 7.34 Accounting for the extra discrepancy in track alignment after fixing for all the measured offsets by assuming a misunderstanding in the collection plane spacing (Figure 7.34a) and the z -positions of the collection wires (Figure 7.34b).

offset between the TPC and the external counters would not have been discovered and all analyses would naively use the incorrect T_0 . The experience in characterising the offsets in the 35-ton, in time, x and z , will be crucial when understanding the eventual DUNE far detector. Based on experience here, it is imperative these misunderstandings are mitigated as much as possible at the far detectors, with each module containing 150 APAs and four drift regions.

7.3.2 Charge Deposited by APA Crossing Tracks

The intriguing distribution of the T_0 -corrected hit times observed in the data, shown in Figure 7.30a, hints at some aspect of the detector response that needs to be understood. In the DUNE far detector, a large number of events will contain particles which pass through the APA frames, so characterising resulting effects is critical. The equivalent plot for simulated data is shown in Figure 7.35. Comparing these distributions, there is a very obvious difference around the interaction time. It appears there is an effect present in the data, not currently being simulated, which manifests in around twice as many hits occurring at T_0 on the collection planes for APA-crossing tracks. This is described in Section 7.3.2.1 and the phenomenon is visible on event displays presented in Section 7.3.2.2.




Fig. 7.35 The T0-corrected drift time for all hits on an APA-crossing track in simulation. The equivalent plot for 35-ton data is shown in Figure 7.30a.

7.3.2.1 Interaction Time Hits

The excess of hits at the interaction time is due to the use of a grounded ‘mesh’ at the centre of the APAs. The purpose of such a design choice is to ensure a uniform electric field across the face of the APA; without it the field would be ill-defined given the presence of the grounded, rectangular APA frames with positively biased planes on either side. It is plausible therefore to consider a ‘backward-facing’ field being set up between the grounded mesh and the positively biased collection planes which would lead to hits drifting the ‘wrong’ way when produced in this region; APA-crossing tracks would hence leave significantly more hits on the collection plane as the other planes. This is demonstrated schematically in Figure 7.36.

A convenient way of confirming whether or not the mesh can explain this excess of hits at the interaction time is possible since one of the four APAs in the 35-ton was constructed without the mesh, precisely for this purpose. Unfortunately, this was the APA which was most plagued by noise issues, so very little good data is available from channels on this APA. It is however possible to make a crude comparison; this is shown in Figure 7.37. This appears to confirm the sharp peak of hits occurring at the interaction time comes from the APAs which use a mesh.

Using the 35-ton dataset, it is also possible to confirm that the mesh is functioning as expected. Without a mesh, one may expect a difference between the hits deposited on wires





Fig. 7.36 Demonstration of the ionisation and hit collection for APA-crossing tracks. The red line represents a track passing through the anode planes, shown in black. The grey region is the centre of the APA frame on which the grounded mesh is affixed. The red dots correspond to the ionisation of electrons which then drift, depositing charge (black dots) on the readout wires. The three planes shown are, from left to right, the collection plane and the two induction planes. The biasing of each of the planes and mesh sets up field lines which all terminate on the collection wires, resulting in charge collected from before the track passes through and after.



(a) Hit times for all hits on APAs 0, 2 and 3; these are the three APAs containing the grounded mesh at the centre.

(b) Hit times for all hits on APA 1, the APA without a grounded mesh at its centre.

Fig. 7.37 Comparison between the T0-corrected hit time distributions on APAs with and without the grounded mesh. Even given the very low statistics in Figure 7.37b, there is a noticeable difference in the distribution of hits around the interaction time.

towards the centre of an APA face and wires at the edges, in front of the grounded frame. The functionality of the grounded mesh ensures there is no difference between any wires on a given APA. Figure 7.38 confirms this is the case.

A natural question to pose at this point is to ask if these ‘extra’ hits deposited by APA-crossing tracks as a result of this ‘backwards’ field have similar properties to the ‘correct’ hits. The most important property to consider is the charge of the hits; Figure 7.39 shows the average charge per hit for hits occurring at the interaction time and all other hits. It is clear from this there is nothing different about these additional hits and they can be treated in the same way.

As alluded to earlier, the DUNE simulation software is simplistic and does not simulate any ionisation within the region of the APA planes; in the case of APA-crossing muons this results in no hits being created after the track passes through the first induction wires. Evidently, this is an important region and must be understood and well simulated in order to test reconstruction and analyses. When this is added to the software, the 35-ton data will be essential for validation purposes.

7.3.2.2 Event Displays of APA-Crossing Tracks

The effect investigated in Section 7.3.2.1 is directly observable in the raw data, as shown in Figure 7.40. The electrons ionised as the particle track passes between the collection





Fig. 7.38 Comparison between the distribution of T0-corrected hit times for hits on wires in front of the APA frame and away from the APA frame to validate the functionality of the mesh. Both distributions are normalised by the number of entries. There is no evidence of any differences between the two distributions so this suggests the mesh is working as intended.



(a) Hits occurring around the interaction time; $50 < \text{tick} < 70$. A fitted Gaussian of the peak yields a mean of 149 and width of 49.

(b) Hits occurring away from the interaction time; $\text{tick} < 50, \text{tick} > 70$. A fitted Gaussian of the peak yields a mean of 152 and a width of 28.

Fig. 7.39 Average lifetime-corrected charge per hit for hits on an APA-crossing track separated according to whether or not the hit was collected around the interaction time. There is no evidence to suggest the ‘extra’ hits collected around the interaction time have significantly more or less average charge than ‘regular’ hits.

plane and the mesh are observable as hits which appear to have drifted in the negative time direction. The outcome is a little ‘hook’ shape in the data.

7.3.3 Comparing Drift Regions with APA-Crossing Tracks

APA-crossing tracks may be utilised to make unique, specific measurements of the detector, made possible since they originate from the same particle. For example, any drift velocity differences between the drift regions may be observed and the noise levels on the collection readouts on either side of the APA can be studied and compared.

The drift velocity is given by the angle of the track in wire/time space and any difference between this velocity in the two drift regions would be noticeable in a refraction-like effect. This is demonstrated in Figure 7.41a. A measure of the angle between the track segments in the different regions would therefore be a measure of the change in drift velocity; this is shown in Figure 7.41b.

The relative noise on the two collection planes can be evaluated by considering the number of hits present in the counter shadow, in each drift region, which were not reconstructed as part of the track associated with the triggering particle. The difference between each collection plane for a given event should peak at zero if similar levels of noise were observed in each drift region; this is confirmed in Figure 7.42.

7.4 Shower Reconstruction in 35-ton Data

The developments in the reconstruction for LArSoft, discussed in Chapter 6, were originally motivated by an interest in reconstructing and analysing π^0 mesons in the 35-ton data. Given the unfortunate eventual problems prevalent in the data, such analyses would be extremely challenging and likely impossible. Since it is still interesting and instructive to analyse how well the reconstruction performs on a sample of real data, this will be briefly explored in the present section.

Considerations relevant when applying the reconstruction developed on simulation to data are discussed in Section 7.4.1 before the necessary reanalysis of the calorimetry is presented in Section 7.4.2. The algorithms are applied to a shower and a π^0 candidate found in the data in Sections 7.4.3 and 7.4.4 respectively.

7.4.1 Data Specific Reconstruction

The BlurredCluster and EMShower algorithms, outlined in Sections 6.3.2 and 6.3.3 respectively, were applied to the data in an attempt to reconstruct particle objects. In general,

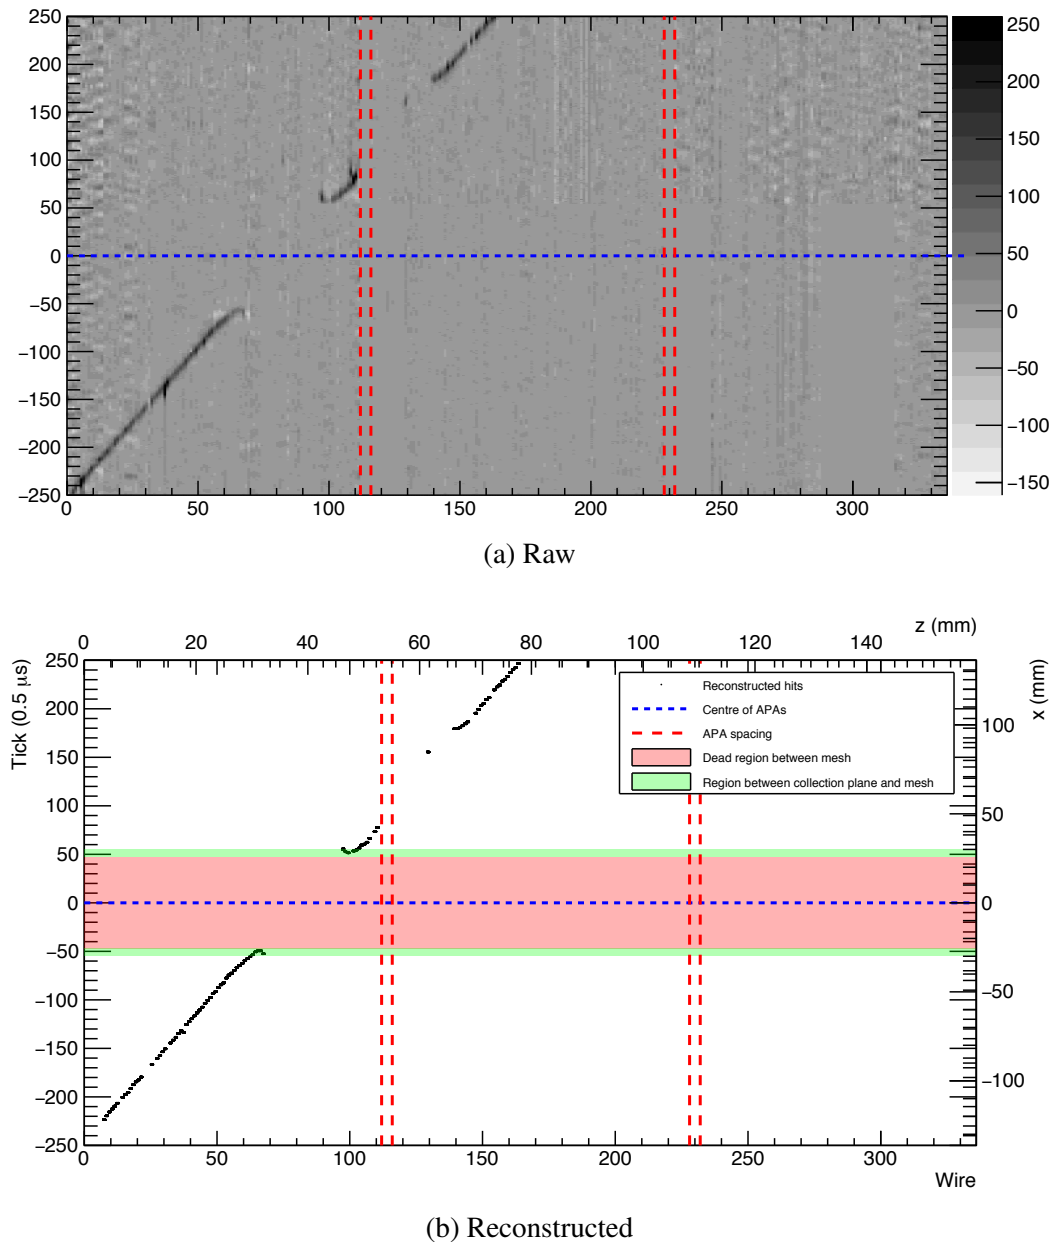



Fig. 7.40 Event display of an APA-crossing track with the charge deposited as it passes through the APAs evident. Figure 7.40a shows the raw charge and Figure 7.40b shows the reconstructed hits. The 'hook'-like effect is visible, with hits at apparently negative drift time. The cm scale on Figure 7.40b is provided as a guide and is not completely correct due to the differing fields.



(a) Demonstration of how differing drift velocities between the drift regions would manifest in the data. A refraction-like effect would result in an angle between the two track segments.

(b) The angle between the track segments on either side of the APAs. The distribution peaks around zero, implying, as expected, the drift velocity is constant in both regions.

Fig. 7.41 Using APA-crossing tracks to confirm the drift velocity is consistent between the two drift regions.




Fig. 7.42 Comparison of noise levels between the two drift regions using APA-crossing tracks. The number of noise hits in the counter shadow for each drift region was considered separately by neglecting all hits identified as track hits, all hits on noisy wires and all hits which have a large number of closely neighbouring hits (which could be symptomatic of unrelated particle tracks). The difference between the number of noise hits in each peaks around zero, implying similar levels of noise.

the algorithms worked out of the box and required no tuning. Since this requires real 3D reconstruction, as opposed to the subtle techniques developed to circumvent the issues with the induction planes (described in Section 7.1.3), the use of more than just the collection plane is necessitated. Reconstruction is therefore only possible for showers with large enough signals on induction planes, following the coherent noise removal and hit disambiguation.

As showering particles are likely to have a different interaction time, T_0 , to the triggering muon, an unassociated method for obtaining this information is required. In general, the photon detectors are designed for this purpose so the use of these seems natural.

Since it is highly unlikely the electronics models and detector responses used in the simulation are perfectly accurate, applying the calorimetric reconstruction to the data without modification would be inappropriate. The relevant calorimetric constants and functions must be determined from the data; this is essential for complete reconstruction and is discussed in Section 7.4.2.

7.4.2 Calorimetry Reconstruction

There are two relevant calorimetric conversions which are pertinent to shower reconstruction (both previously discussed in Section 6.2.3): the calorimetry constant and the shower energy conversion. The methods used to determine these for data will be discussed in this section.

It should be stressed that due to the large noise levels, accurate calorimetry will not be possible in the 35-ton data. This may be understood by considering the distribution of charge deposited by ionising particles; typically this is sampled from a Landau distribution with a most probable value dependent on the electron drift distance (due to lifetime effects). Since hit reconstruction tends to cut on the hit ‘threshold’, the height of the peak above pedestal, this compromises lower energy hits populating the full charge distribution and biases the reconstruction toward higher energies. This is demonstrated in Figure 7.43. As far as possible, steps to mitigate these effects have been applied in the preceding discussion. There will however be an inevitable bias so the following should not be treated as a full, rigorous assessment.

Calorimetric reconstruction is only attempted for the collection planes where the effects of noise are mitigated somewhat compared to the induction views. Since the data used were taken at a drift field of 250 V/cm (half the nominal voltage), the recombination factor used must take this into account. At 500 V/cm the value is 0.63 whilst at 250 V/cm a factor of 0.52 is used.

The procedure invoked to determine the calorimetry constant is largely identical to that used in simulation: the dE/dx of a through-going mip is calculated and the constant varied until the expected distribution is obtained. The through-going muons used in the analyses




Fig. 7.43 The bias in the hit selection due to a high noise level in the 35-ton data [179]. The charge distribution for through-going muons is shown for three different displacements along the drift direction, $20 \text{ cm} < x < 40 \text{ cm}$, $80 \text{ cm} < x < 100 \text{ cm}$ and $140 \text{ cm} < x < 160 \text{ cm}$ and takes the form of a Landau convoluted with a Gaussian. The red line represents a typical hit finding threshold. The most probable value of the distribution is close to this boundary in each of the cases, resulting in the lower charge hits being missed. This introduces a bias towards higher charge and has implications for the reliability of calorimetry in the 35-ton data sample.

described in Sections 7.3 and 7.2 were utilised to make these measurements. Additional necessary information, such as the interaction time (to correct the charge for lifetime) and the track angle (to correct the dE/dx for track pitch), is provided by the counters. In order to produce reliable results, an additional cut requiring at least 20 consecutive wires with a single hit on each was applied, with the dE/dx measurement obtained using just these hits. The eventual dE/dx distribution is demonstrated in Figure 7.44 and implies a calorimetry constant of 7.4×10^{-3} (for comparison, the value used for the collection plane in simulation is 5.4×10^{-3}).

In simulation, truth information was used to find a general charge to energy conversion used in, for example, the determination of total shower energy. This obviously is not possible in data so a similar technique to the calculation of the calorimetry constant described above was used. The lifetime-corrected charge and track pitch information can be utilised to find a value of dQ/dx (ADC/cm), which may then be converted into a measure of dE/dx (MeV/cm) using the calorimetry constant previously determined. This may in turn be used to find the total deposited energy by taking into account the distance travelled by the associated track in the collection view. As demonstrated in Figure 6.7, there exists a linear relationship between total deposited lifetime-corrected charge and the particle energy; this is also seen in data in Figure 7.45. This may then be used as a conversion in shower energy reconstruction.




Fig. 7.44 The dE/dx distribution for mips passing through the 35-ton TPC. The calorimetry constant is chosen to ensure the peak of the distribution, which ideally follows a Landau, is around 1.8-1.9 MeV/cm.

7.4.3 Shower Reconstruction

Using the modifications discussed in Sections 7.4.1 and 7.4.2, the performance of the showering reconstruction on real data can be assessed by applying it to an electromagnetic shower. The result of applying the algorithms to the famous 35-ton shower depicted in Figure 4.31 is shown in Figure 7.46. The calorimetric reconstruction yields a dE/dx of 1.1 MeV/cm and a total shower energy of 188 MeV. These results appear feasible and are consistent with an electron shower, for which one would expect a dE/dx peaked around 2.1 MeV/cm; 1.1 MeV/cm is not an unreasonable value in the tail of this distribution. Given its dE/dx and energy, the most likely origin for this particle is Compton scattering of a high-energy photon produced from bremsstrahlung radiation from the cosmic muon, or from a delta ray produced by it [180].

The T0 for this particle was determined to be 4740 ticks from reconstructing flash information collected by the photon detectors – this makes this shower the only fully automated reconstructed particle object in the 35-ton dataset.




Fig. 7.45 Relationship between deposited charge and energy for 35-ton data, calculated using through-going mips. The lower edge of the distribution follows a linear pattern and it is this which the conversion is chosen to represent. Deviations from this linear fit observed above it are related to the fundamental issues with the 35-ton data and arise from missed charge due to the problems illustrated in Figure 7.43. This results in hits reconstructed with a lower charge for the implied energy deposited by the mip. It should be noted there are no cases of extra charge deposited; this concurs with this interpretation and ensures confidence in the displayed line as the correct conversion may be assumed. A possible way to restore linearity would be to truncate the thresholding demonstrated in Figure 7.43.




Fig. 7.46 Result of applying the shower reconstruction on a shower observed in the 35-ton data. Each small rectangle represents a reconstructed hit and the colour associated with each corresponds to a reconstructed shower object. The stars and dotted lines represent the reconstructed start point and direction for each shower.




Fig. 7.47 A candidate π^0 event observed in the online event display during the run.

7.4.4 π^0 Reconstruction

An important calorimetric test of particle detectors involves demonstrating a reasonable reconstructed π^0 mass peak. It was for this reason that the shower reconstruction discussed in Chapter 6 was developed. An analysis of a π^0 candidate event is briefly considered here.

Without full reconstruction and selection, identifying candidate events is very difficult. Such an event was observed in the online event display however and is shown in Figure 7.47.

Unlike the shower discussed in Section 7.4.3, there is no associated photon detector information for this event; however, one of the candidate photons passes through the APA frames so techniques developed for the APA-crosser analysis (Section 7.3) may be employed to determine the relevant interaction time. Applying the calorimetry reconstruction, the dE/dx information associated with the high energy candidate photon (the one which crosses the APAs) gives a value of 4.75 MeV/cm, entirely consistent with the expectation for a photon of a distribution centred around 4.2 MeV/cm. The low energy candidate photon travels almost completely along the collection view direction resulting in unreliable dE/dx information. The total energy for each shower is determined to be 161.8 MeV and 500.5 MeV with an implied

invariant mass of

$$m_{\pi^0} = 156.6 \text{ MeV}, \quad (7.3)$$

comparable to the true π^0 mass of 135 MeV.

Without fully considering uncertainties and biases present, it is not possible to make a judgement as to the performance of the basic calorimetric reconstruction discussed in Section 7.4.2 or to confirm whether or not the event displayed in Figure 7.47 represents a π^0 decay. However, dE/dx values of 1.1 MeV/cm and 4.75 MeV/cm for different showers appear consistent with electron and photon particles respectively and, within the limits of the analysis presented here, it is conceivable the particle with invariant reconstructed mass of 156.6 MeV is indeed a π^0 .

7.5 35-ton Data Analysis Summary

Despite initial problems with the 35-ton, good progress has been made in analyses, specifically focussing on understanding the detector. Techniques developed will be directly applicable to the data collected with the eventual far detector and for the commissioning and operation of ProtoDUNE.

The studies using TPC-crossing muons are highly useful in developing a better understanding of the detector and will be vital for future stages in the DUNE programme. Indeed, both measurements of APA-crossing and APA-gap crossing muons yielded unexpected offsets, demonstrating the utility of such methods and again highlighting the need for prototyping. Despite being unable to perform a π^0 analysis with the 35-ton data, good progress was made applying the novel reconstruction methods described in Chapter 6 to real data. This is hugely important for the future of the DUNE experiment, along with the demonstration of automated reconstruction in the 35-ton data.

In general, the 35-ton was a successful first prototype for DUNE and lessons learned are already being carried forward to ProtoDUNE and are influencing the design of the far detector.

Chapter 8

The ν_e Oscillation Signal at the DUNE Far Detector

A primary aim of the DUNE experiment, as discussed in Chapter 3, is to make precision measurements of the PMNS matrix parameters describing neutrino oscillations by searching for electron neutrino appearance from the predominantly muon neutrino beam (i.e. $\nu_\mu \rightarrow \nu_e$ oscillation, described by Equation 2.30). This channel is of critical importance for all oscillation-related physics and thus requires very efficient reconstruction and selection. Methods developed to provide reconstruction of these events were discussed in detail in Chapter 6 and utilised in this present chapter in the selection of simulated charged-current (CC) ν_e events (ν_e CC) in the DUNE far detector.

The selection presented in the following sections represents the very first generation analysis for the DUNE experiment and serves primarily to demonstrate the principle of selecting these events in a large LArTPC.¹ Initially, the samples and simulation methods utilised in the selection are very briefly discussed in Section 8.1. Much work is needed to advance the analysis to comply with DUNE requirements and many improvements may be expected as further developments progress. It should also be noted the reconstruction discussed in Chapter 6 is not the only solution and various other techniques, primarily using the Pandora toolkit [181], have been assessed, notably in Section 8.2. The selection presented in Section 8.3 does utilise the novel reconstruction detailed in this thesis; however, due to significant recent progress, it is likely the selection will continue to explore all reconstructions in LArSoft and take advantage of the continuing developments. This outlook will be briefly discussed in Section 8.4.

¹This work was undertaken with Dominic Brailsford (Lancaster University), who was working on the complementary ν_μ CC selection.

8.1 Far Detector Samples

The selections discussed in this chapter aim to separate ν_e events from ν_μ and ν_τ events. Samples of 100,000 of each of these particle species were generated within LArSoft for use in the tuning and the running of the analysis.

The neutrinos in the samples were generated as a simulated particle beam, with mostly ν_μ s and a small accidental contamination from other species, before exchanging the flavours to simulate oscillations. This ensures the correct energy spectrum for all neutrinos arriving at the far detector. In the case of the ν_e sample, the muon neutrinos are swapped for electron neutrinos and the beam electron neutrinos for tau neutrinos; in the ν_τ sample, the muon neutrinos are swapped for tau neutrinos and the beam electron neutrino component for muon neutrinos. When used, these events are then weighted according to the relevant oscillation probabilities for the specific transitions simulated, determined by the neutrino energy and the current understanding of the mixing parameters (described in Section 2.3.2).

The events are also scaled to ensure the correct POT (protons-on-target) weighting. The charged-current events in each sample are de-weighted by the number of events and the neutral-current events are de-weighted by the number of events in all three samples. The scaling then weights all events to the same POT (here, 1×10^{21} is used).

8.2 Cut-Based Tuning

This section details the early developments of a selection which utilises only Pandora reconstruction and a cut-based approach. This procedure shows promise at performing as well as the more established multi-variate analysis (MVA) method, discussed in Section 8.3, but is used here mainly to tune the selection.

As will be discussed in Section 8.3, the current MVA implementation of the ν_e CC selection contains a mixture of particle-level and event-level variables and requires significant understanding before developments may progress. The motivation behind considering a simplified selection is to facilitate a careful evaluation of the events, and the analysis utilises a distinct particle identification (PID) system to perform particle level discrimination, separate from the event level classification. As with the majority of this present chapter, this work is very preliminary and represents the first studies into these areas for the DUNE experiment.

8.2.1 Selection

The elementary cut-based selection utilises Pandora reconstruction for tracks and showers and an MVA approach to PID for each of the reconstructed objects. This PID framework




Fig. 8.1 The output of a multi-variate approach to particle identification when attempting to identify electrons.

[182] utilises properties of the particles such as charge, dE/dx and hit distribution to form a hypothesis for each of the particle types muon, electron, proton, photon, pion. The selection simply selects events which contain candidate electron showers and places a cut on the value of the electron-MVA result to attempt to identify the electrons. The separation between electrons and photons for the electron-MVA value of the highest energy shower in each event is demonstrated in Figure 8.1. It is immediately evident that, given the strengths of LArTPC technology to separate these showering particle types, significant progress is required in the next decade.

To tune this cut, and for further tunings discussed in Section 8.2.2, a physics-based approach was taken. Arguably the most important consequence of this analysis, since the value of θ_{13} is reasonably well understood, is to search for CP-violation, and so the tuning was designed to maximise its capability for this. For each electron-MVA value, the ν_e -appearance energy spectrum was determined for the CP-conservation ($\delta_{CP} = 0$) and CP-violation ($\delta_{CP} = \pi/2$) hypotheses; maximising the χ^2 between the distributions ensures the greatest discriminating ability of the analysis to discern hints of CP-violation. This is demonstrated in Figure 8.2, and a cut value of 0.02 is found to provide the best differentiation. The events were POT-weighted and the oscillation probabilities were provided extemporaneously by the Prob3++ framework [183].

In order to assess the efficacy of the analysis, the commonly used efficiency and purity metrics may be employed. The efficiency represents how many of the total number of signal events are chosen by the analysis and the purity describes the fraction of the selected events




Fig. 8.2 The process of tuning the electron cut in the simple cut-based selection by maximising the effect of CP-violation on the oscillation probabilities.

which are indeed signal. The efficiency and purity of the selection with just this simple cut is demonstrated in Figure 8.3. It is found the selection already looks reasonable, with an overall efficiency of 76% and purity of 35%, and, as will be observed and discussed further in Section 8.3, is competitive with the MVA-based analysis.

8.2.2 Fiducial Volume Tuning

A similar approach to the tuning described in the previous section may be utilised to optimise the fiducial volume (FV) applied in the selection. This has not previously been performed in the DUNE far detector, with estimations and assumptions applied during the development of the analysis. The distance of the start point of the electron shower from the walls of the cryostat is considered over a range of energies and oscillation probabilities, as before, and the χ^2 between the distributions for maximum CP-conservation and CP-violation maximised to provide the optimal parameters.

Example plots showing the tuning of the y -coordinate are the subject of Figure 8.4. The tuned FD coordinates, along with the cryostat dimensions, are shown in Table 8.1. It should again be noted this is a preliminary study and this volume would be expected to change with more advanced simulation and reconstruction. In particular, a fully inclusive region in the upstream beam direction ($-z$), implied by these results, will include much activity from the rock and will be better tuned when using a more detailed simulation including deeper rock interactions.

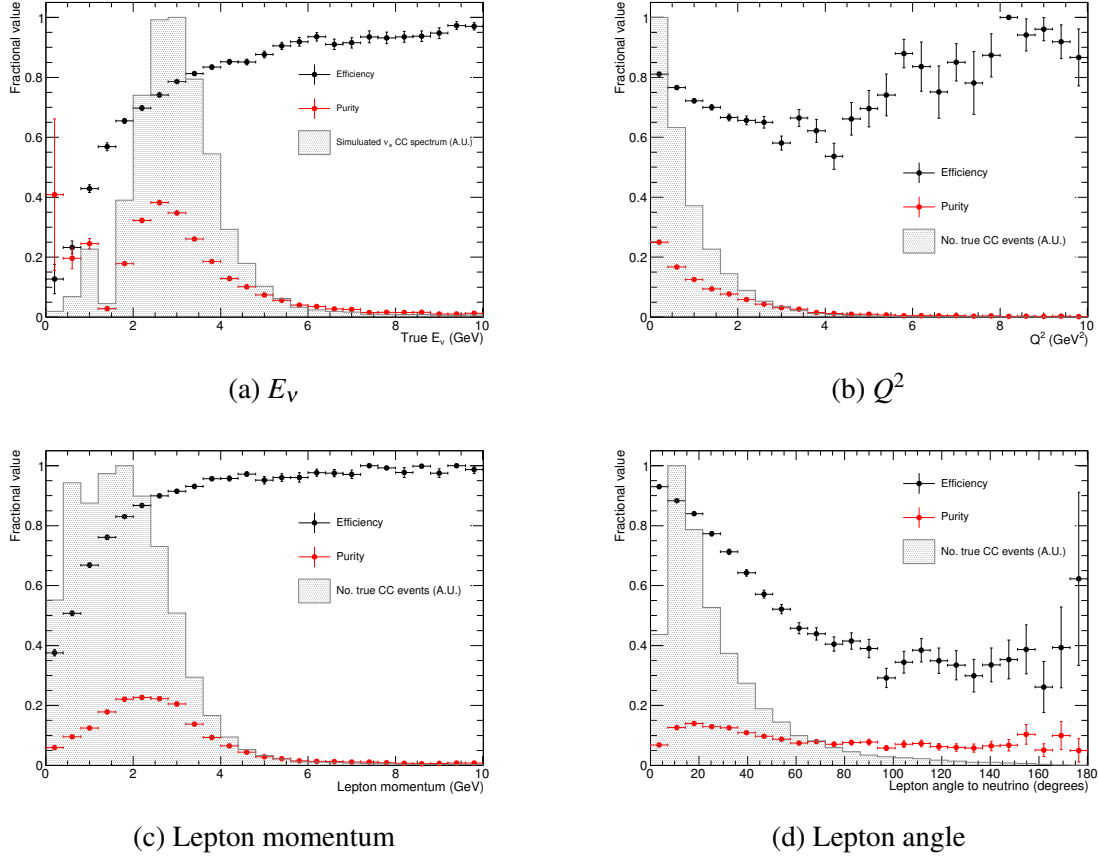


Fig. 8.3 The efficiency and purity of the ν_e CC Pandora cut-based selection as a function of a number of kinematic variables, after applying the selection. The variables are the true neutrino energy (Figure 8.3a), the true momentum transfer, Q^2 (Figure 8.3b), the electron momentum (Figure 8.3c) and the angle the electron makes to the neutrino beam (Figure 8.3d). The plots are filled for each event which passes the cut (highest energy shower PID MVA response > 0.02).

Table 8.1 The dimensions and tuned fiducial volume of a single DUNE far detector module.

	$-x$	$+x$	$-y$	$+y$	$-z$	$+z$
Dimensions (cm)	-360.0	360.0	-600.0	600.0	0.0	1394.0
Fiducial volume (cm)	-358.5	358.0	-597.5	581.5	0.0	1245.0




Fig. 8.4 Tuning the DUNE far detector fiducial volume, y -coordinate. The asymmetry is a result of the direction of the beam; it is not exactly parallel to the z -direction due to the curvature of the Earth and thus results in a different response from the top and bottom of the detector.

8.3 MVA-Based Selection

The current leading ν_e CC selection utilises a multi-variate approach, as briefly discussed previously, with a mixture of event-level and particle-level variables designed to separate ν_e events from ν_μ and ν_τ events.² The current performance of this analysis, when applied to events reconstructed using the methods developed and described in Chapter 6, is the subject of this section.

8.3.1 MVA Input Variables

In total there are 30 input variables to the MVA [184], summarised in Table 8.2, containing information about the event, the longest reconstructed track and the reconstructed shower with the highest energy in the event. The separation between signal (ν_e) and background (ν_μ and ν_τ) events for each of these variables are presented in Appendix B.

The use of all of these variables is still being understood and it is possible some are biasing the selection. This is the motivation behind studying the events using the cut-based approach introduced in Section 8.2 and is only just beginning to be explored by the collaboration. The selection presented in this section should be viewed as a ‘proof-of-principle’, demonstrating the current simulation and reconstruction tools may be used to begin development of analyses, rather than as a complete analysis which may be used to study data.

²This implementation of the selection was developed by Tingjun Yang (Fermi National Accelerator Laboratory) and Tyler Alion (University of Sussex) and is unchanged from their developments.

Table 8.2 The input variable used in the MVA designed to separate ν_e events from ν_μ and ν_τ events. The variables describe the event, the longest track and the highest energy shower in the event.

Variable	Description
Event charge	Total charge deposited in the event
Number of tracks	Number of reconstructed tracks in the event
Maximum track length	The length of the longest reconstructed track in the event
Average track length	The average reconstructed track length for all tracks in the event
Track dE/dx	The dE/dx of the longest track in the event across its length
Signal fluctuation	Ratio of lowest 50% to highest 50% of measured charge
Transverse track profile	Fraction of charge within 200 ticks of longest track
Fraction of track charge	Fraction of the event charge deposited by the longest track in the event
Track PIDA	The output of a particle identification algorithm which considers the ionising power of tracks as a function of their residual range [185]
Maximum fraction of charge in 5, 10, 50, 100 wires (4)	Maximum fraction of charge contained on neighbouring 5, 10, 50 and 100 wires
Track angle (x, y, z) (3)	The angle the longest track makes to each dimension (x, y, z)
Fractional transverse energy	The fractional transverse energy of the longest track
Number of showers	Number of reconstructed showers in the event
Shower dE/dx	The dE/dx of the highest energy shower
Shower energy	The energy of the highest energy shower
Fraction of shower charge	Fractional position along the highest energy shower to maximal charge
Number of hits per shower wire	Number of reconstructed hits from the highest energy shower per shower hitting wire
Shower length	The length of the shower in units of wire number
Shower maximum	The fractional distance along the length of the shower to maximal charge
Displacement of shower start (x, y, z) (3)	Distance of the highest energy shower start point from the event neutrino vertex in x, y, z
Shower angle (x, y, z) (3)	The angle the highest energy shower makes to each dimension (x, y, z)




Fig. 8.5 The MVA response when training ν_e (signal) against ν_μ and ν_τ (background) using the variables described in Section 8.3.1.

8.3.2 Analysis Performance

The MVA response to the input variables discussed in the previous section is shown in Figure 8.5. A good separation is observed, demonstrating successful reconstruction and well discriminating variables. The cut was tuned in an identical way to the method described earlier (in Section 8.2.1), by maximising the distinction between the oscillated spectra for the maximum CP-conserving and CP-violating hypotheses, and a value of 0.81 was found to provide the optimum selection.

The efficiency and purity distributions for the selection, over a number of observable kinematic variables, are shown in Figure 8.6 for all events before applying the selection and in Figure 8.7 following the implementation of the MVA cut.

Overall, the selection is shown to have an efficiency of 52% and a purity of 63%. The biases in the selection around the distribution of events are evident from the plots in Figure 8.7 and have not yet been fully understood. This potentially may not be a significant issue but does require further comprehension and is being addressed using the selection discussed in Section 8.2, where the effects observed in the assessment of the MVA-based analysis were not present and instead the performance was more directly correlated with the kinematic variables (shown in Figure 8.3). Given that the PID MVA utilised in the cut-based selection employs many of the same particle-level variables as in this present implementation, it appears a more natural approach to the problem. By incorporating the event-level information to the

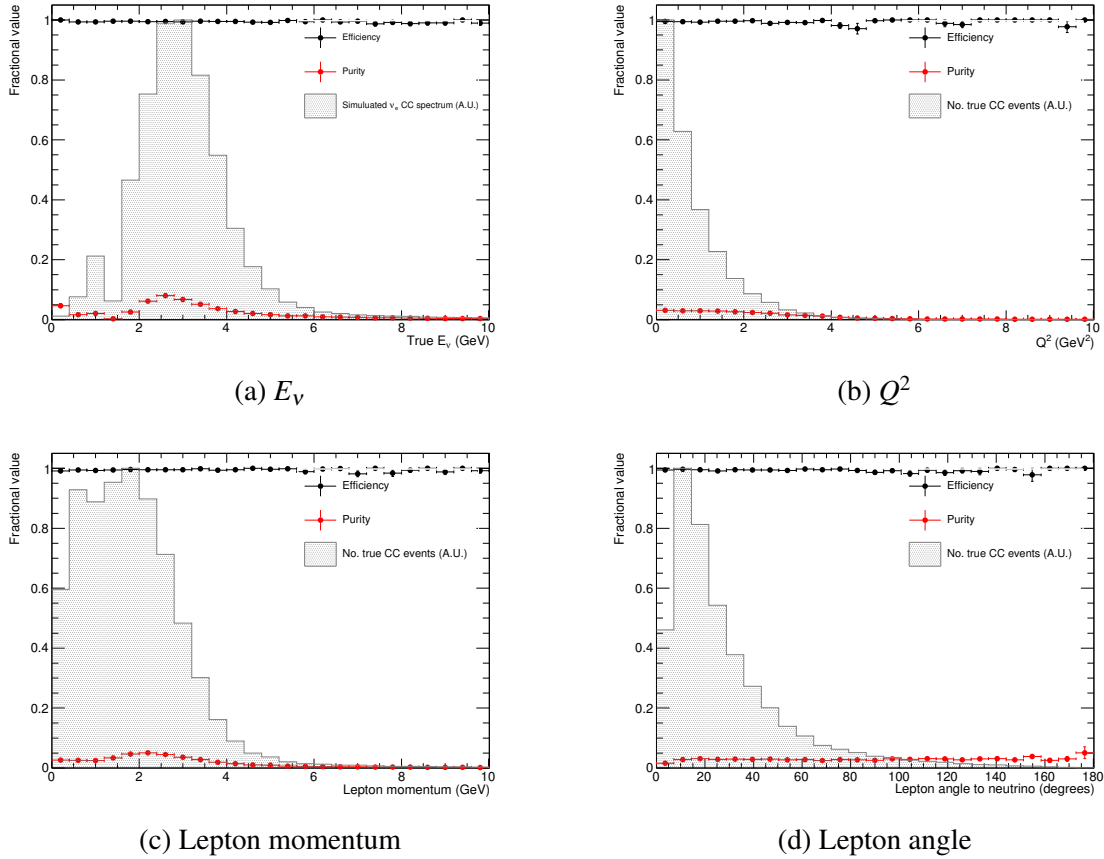


Fig. 8.6 The efficiency and purity of the ν_e CC MVA-based selection as a function of a number of kinematic variables, before applying the selection. The variables are the true neutrino energy (Figure 8.6a), the true momentum transfer, Q^2 (Figure 8.6b), the electron momentum (Figure 8.6c) and the angle the electron makes to the neutrino beam (Figure 8.6d). The plots are filled for each event with reconstruction within the fiducial volume. Given the relaxed restrictions on this region currently, described in Section 8.2.2, these distributions may be expected to change significantly when using a more realistic FV.

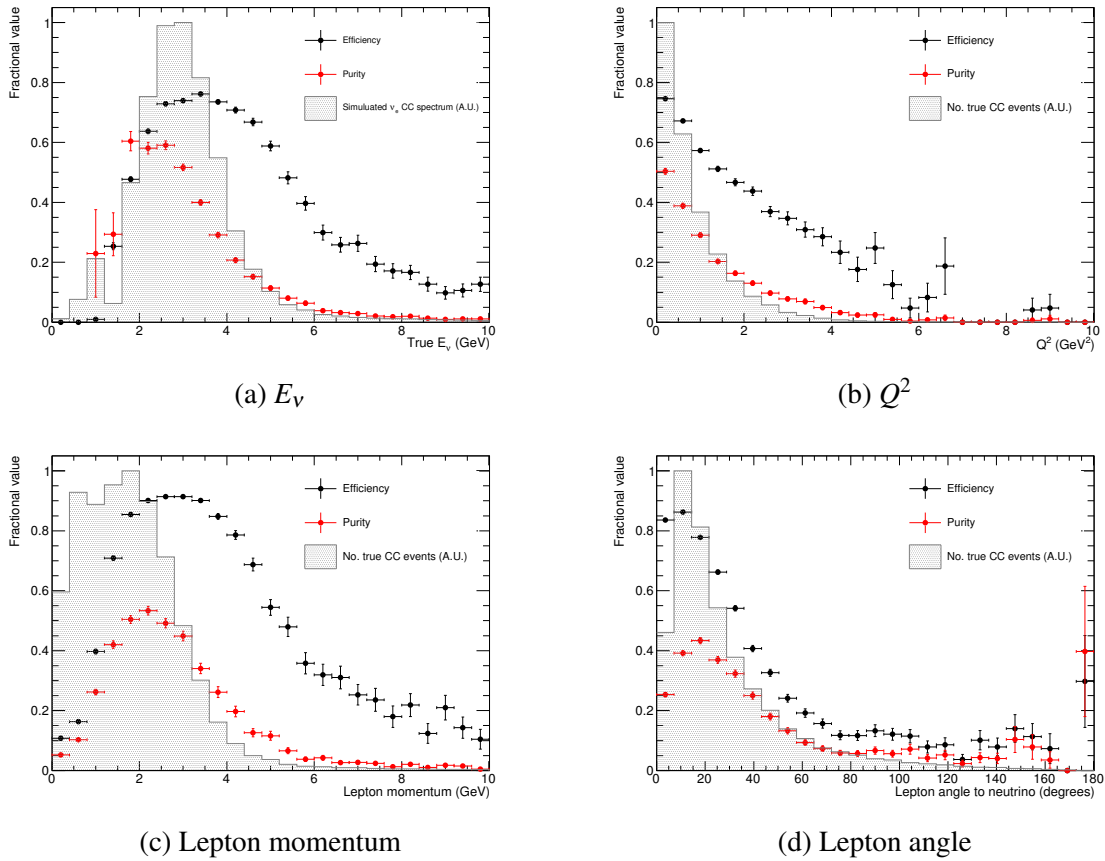


Fig. 8.7 The efficiency and purity of the ν_e CC MVA-based selection as a function of a number of kinematic variables, after applying the selection. The variables are the true neutrino energy (Figure 8.7a), the true momentum transfer, Q^2 (Figure 8.7b), the electron momentum (Figure 8.7c) and the angle the electron makes to the neutrino beam (Figure 8.7d). The plots are filled for each event which passes the cut (MVA response > 0.81).

selection on top of this basis, a more thorough evaluation of the interactions is possible. As already discussed, this work is ongoing.

8.4 Summary and Outlook for Future Selections

The analyses presented in this chapter represent a promising first step towards the eventual DUNE ν_e CC selection. The MVA-based selection, described in Section 8.3, utilises discriminating variables which have proved useful in preceding neutrino experiments and currently yields an event sample with an overall efficiency of 52% and purity of 63%. Despite these reasonable performance metrics, the selection still needs to be completely understood and has motivated the development of the stripped-back analysis discussed in Section 8.2, which may be used to introduce the quantities input to the MVA iteratively to facilitate a thorough evaluation of their relative merits. However, with an efficiency of 76% and a purity of 35% after just one simple cut, it is already competitive to the more established selection and should be pursued as a genuine alternative to this early implementation of the ν_e CC analysis. Further cuts would be expected to improve the purity by removing the remaining background events and, if chosen astutely, may not impact the efficiency as harshly as the MVA-based selection.

Future iterations will include many improvements in general but specifically in the event reconstruction, as discussed in Section 6.3.5, and in the current implementation of the selection. In addition to the MVA and Pandora cut-based studies, a significant effort in utilising the CVN machine learning techniques, which have been successfully demonstrated in the NOvA experiment [186], for this purpose is underway.

The current rate of development within the LArSoft and DUNE collaborations is significant and progress over the next decade will facilitate a high-precision study of the electron neutrino appearance channel in the DUNE far detector. This is of critical importance to the success of the DUNE project and, considering the efforts presented in this thesis in both the reconstruction (Chapter 6) and the selection (Chapter 8), there can be much confidence in the future of the analysis and the potential discovery power of the experiment.

Chapter 9

Conclusions

The developments presented in this thesis represent crucial steps on the experimental path towards the eventual Deep Underground Neutrino Experiment and the world-leading physics aims of the programme. The DUNE experiment is essential in advancing our current understanding of neutrino physics as we approach the precision-era, as discussed in Chapter 2, and has much exciting physics potential, described in detail in Chapter 3, notably the search for CP-violation in the lepton sector, which would represent the most significant progress in the field since the initial discovery of neutrino oscillations.

Chapter 4 described the engineering challenges facing the construction and operation of such an ambitious experiment and discussed the current test facilities and prototyping efforts designed to address these concerns. Of particular interest was the 35-ton experiment, the first prototype of the DUNE far detector, which represented the opportunity to facilitate a deeper understanding of the detector technology and the potential associated issues. Despite many problems, which were discussed in detail, the experiment was a success and represented significant progress as we advance towards the final DUNE far detector. Alongside the lessons learned, many positive outcomes ensure encouragement may be taken forward into future prototyping efforts, including ProtoDUNE next year at CERN, and, given the time and resources available before construction begins in 2021, the final full-scale detector for DUNE.

The 35-ton prototype was discussed in more detail in Chapter 5 and Chapter 7, which described the framework for the data quality monitoring during the running of the 35-ton Phase II, and multiple data analyses performed with data collected from the run, respectively. As the 35-ton is the first detector utilising all elements from the DUNE far detector design, much can be learned from the experiences of data taking and from the data itself. It is imperative as much as possible is extracted from the 35-ton prototype and the analyses presented successfully studied multiple effects and datasets for the first time. All developed

ideas and techniques will be vital for further detector development considerations and for the calibration of future LArTPC experiments.

Alongside the engineering obstacles associated with LAr experiments on such large scales, the reconstruction of physics objects within the highly detailed events provided by the detector presents many additional challenges to the success of current and future experiments. Although a lot of progress has been made, further developments are essential to ensure the relevant analyses may be achieved once data taking starts. The development of novel reconstruction techniques for showering particles was discussed in detail in Chapter 6, with specific considerations for ν_e CC events at the DUNE far detector. This channel is essential for the physics required of the experiment and it is critical they may be well reconstructed, selected and understood. The reconstruction was found to be effective only around 50% of the time, highlighting the necessity of further developments, but representing good progress on what is undoubtedly a very challenging problem. The application of this reconstruction was lastly demonstrated in Chapter 8 on simulated far detector neutrino events, with a very first-generation selection utilised to analyse the interactions. The status of the DUNE analysis capabilities was discussed and the route forwards was outlined. As with present reconstruction abilities, further improvements are required but the current developments represent significant progress. Overall, given the current timescales, there is much cause for optimism for the progress of the exciting experimental programme of the DUNE project and its ambitious physics goals.

References

- [1] S. L. Glashow; *Partial-symmetries of weak interactions*; Nuclear Physics **22**, 4, 579–588 (1961).
URL [http://dx.doi.org/10.1016/0029-5582\(61\)90469-2](http://dx.doi.org/10.1016/0029-5582(61)90469-2)
- [2] S. Weinberg; *A Model of Leptons*; Phys. Rev. Lett. **19**, 21, 1264–1266 (1967).
URL <http://dx.doi.org/10.1103/PhysRevLett.19.1264>
- [3] G. Aad, T. Abajyan, B. Abbott *et al.*; *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*; Physics Letters B **716**, 1, 1–29 (2012).
URL <http://dx.doi.org/10.1016/j.physletb.2012.08.020>
- [4] S. Chatrchyan, V. Khachatryan, A. M. Sirunyan *et al.*; *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*; Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics **716**, 1, 30–61 (2012).
URL <http://dx.doi.org/10.1016/j.physletb.2012.08.021>
- [5] T. Shears; *The Standard Model*; Phil. Trans. Roy. Soc. Lond. **A370**, 805–817 (2012).
URL <http://dx.doi.org/10.1098/rsta.2011.0314>
- [6] S. M. Bilenky; *Neutrino in standard model and beyond*; Physics of Particles and Nuclei **46**, 4, 475–496 (2015).
URL <http://dx.doi.org/10.1134/S1063779615040024>
- [7] J. Ellis; *Outstanding questions: Physics beyond the Standard Model*; Phil. Trans. Roy. Soc. Lond. **A370**, 818–830 (2012).
URL <http://dx.doi.org/10.1098/rsta.2011.0452>
- [8] Y. Fukuda, T. Hayakawa, E. Ichihara *et al.*; *Evidence for Oscillation of Atmospheric Neutrinos*; Phys. Rev. Lett. **81**, 8, 1562–1567 (1998).
URL <http://dx.doi.org/10.1103/PhysRevLett.81.1562>
- [9] Q. R. Ahmad, R. C. Allen, T. C. Andersen *et al.*; *Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory*; Phys. Rev. Lett. **89**, 1, 11301 (2002).
URL <http://dx.doi.org/10.1103/PhysRevLett.89.011301>
- [10] W. Pauli; *Open letter to the participants of the conference in Tübingen (1930)*.
URL <http://microboone-docdb.fnal.gov/cgi-bin/RetrieveFile?docid=953;filename=pauliletter1930.pdf>

- [11] E. Fermi; *Trends to a Theory of beta Radiation. (In Italian)*; Nuovo Cim. **11**, 1–19 (1934).
URL <http://dx.doi.org/10.1007/BF02959820>
- [12] E. Fermi; *Versuch einer Theorie der β -Strahlen. I*; Zeitschrift für Physik **88**, 3, 161–177 (1934).
URL <http://dx.doi.org/10.1007/BF01351864>
- [13] F. Wilson; *Fermi's Theory of Beta Decay*; American Journal of Physics **36**, 12, 1150–1160 (1968).
URL <http://dx.doi.org/10.1119/1.1974382>
- [14] G. M. Lewis; *Neutrinos*; Wykeham publications, London; Winchester (1970); ISBN 0851091407 9780851091402.
- [15] C. M. G. Lattes, H. Muirhead, G. P. S. Occhialini *et al.*; *Processes Involving Charged Mesons*; Nature **159**, 694–697 (1947).
URL <http://dx.doi.org/10.1038/159694a0>
- [16] C. M. G. Lattes, G. P. S. Occhialini, and C. F. Powell; *Observations on the Tracks of Slow Mesons in Photographic Emulsions. I*; Nature **160**, 453–456, 486–492 (1947).
URL <http://dx.doi.org/10.1038/160453a0>
- [17] R. Brown, U. Camerini, P. H. Fowler *et al.*; *Observations With Electron Sensitive Plates Exposed to Cosmic Radiation*; Nature **163**, 82 (1949).
URL <http://dx.doi.org/10.1038/163082a0>
- [18] C. L. Cowan, F. Reines, F. B. Harrison *et al.*; *Large Liquid Scintillation Detectors*; Phys. Rev. **90**, 3, 493–494 (1953).
URL <http://dx.doi.org/10.1103/PhysRev.90.493>
- [19] F. Reines and C. L. Cowan; *A proposed experiment to detect the free neutrino*; Physical Review **90**, 3, 492–493 (1953).
URL <http://dx.doi.org/10.1103/PhysRev.90.492.2>
- [20] F. Reines and C. L. Cowan; *Detection of the free neutrino*; Physical Review **92**, 3, 830–831 (1953).
URL <http://dx.doi.org/10.1103/PhysRev.92.830>
- [21] C. L. Cowan, F. Reines, F. B. Harris *et al.*; *Detection of the Free Neutrino: a Confirmation*; Science **124**, 3212, 103–104 (1956).
URL <http://dx.doi.org/10.1126/science.124.3212.103>
- [22] R. Davis Jr. and D. S. Harmer; *Attempt to observe the $Cl^{37}(\bar{\nu}e^-)Ar^{37}$ reaction induced by reactor antineutrinos*; Bull. Am. Phys. Soc. **4**, 217 (1959).
- [23] G. Danby, J.-M. Gaillard, K. Goulianos *et al.*; *Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos*; Phys. Rev. Lett. **9**, 1, 36–44 (1962).
URL <http://dx.doi.org/10.1103/PhysRevLett.9.36>

- [24] M. L. Perl, G. S. Abrams, A. M. Boyarski *et al.*; *Evidence for anomalous lepton production in e^+e^- annihilation*; Physical Review Letters **35**, 22, 1489–1492 (1975).
URL <http://dx.doi.org/10.1103/PhysRevLett.35.1489>
- [25] G. J. Feldman, F. Bulos, D. Lüke *et al.*; *Inclusive Anomalous Muon Production in e^+e^- Annihilation*; Phys. Rev. Lett. **38**, 3, 117–120 (1977).
URL <http://dx.doi.org/10.1103/PhysRevLett.38.117>
- [26] J. Burmester, L. Criegee, H. C. Dehne *et al.*; *Anomalous muon production in e^+e^- annihilations as evidence for heavy leptons*; Physics Letters B **68**, 3, 297–300 (1977).
URL [http://dx.doi.org/10.1016/0370-2693\(77\)90292-1](http://dx.doi.org/10.1016/0370-2693(77)90292-1)
- [27] D. DeCamp, B. Deschizeaux, J. P. Lees *et al.*; *Determination of the number of light neutrino species*; Physics Letters B **231**, 4, 519–529 (1989).
URL [http://dx.doi.org/10.1016/0370-2693\(89\)90704-1](http://dx.doi.org/10.1016/0370-2693(89)90704-1)
- [28] B. Adeva, O. Adriani, M. Aguilar-Benitez *et al.*; *A determination of the properties of the neutral intermediate vector boson Z^0* ; Physics Letters B **231**, 4, 509–518 (1989).
URL [http://dx.doi.org/10.1016/0370-2693\(89\)90703-X](http://dx.doi.org/10.1016/0370-2693(89)90703-X)
- [29] M. Z. Akrawy, G. Alexander, J. Allison *et al.*; *Measurement of the Z^0 mass and width with the opal detector at LEP*; Physics Letters B **231**, 4, 530–538 (1989).
URL [http://dx.doi.org/10.1016/0370-2693\(89\)90705-3](http://dx.doi.org/10.1016/0370-2693(89)90705-3)
- [30] P. Aarnio, P. Abreu, W. Adam *et al.*; *Measurement of the mass and width of the Z^0 -particle from multihadronic final states produced in e^+e^- annihilations*; Physics Letters B **231**, 4, 539–547 (1989).
URL [http://dx.doi.org/10.1016/0370-2693\(89\)90706-5](http://dx.doi.org/10.1016/0370-2693(89)90706-5)
- [31] S. Schael, R. Barate, R. Bruneliere *et al.*; *Precision electroweak measurements on the Z resonance*; Physics Reports **427**, 5-6, 257–454 (2006).
URL <http://dx.doi.org/10.1016/j.physrep.2005.12.006>
- [32] K. Kodama, N. Ushida, C. Andreopoulos *et al.*; *Observation of tau neutrino interactions*; Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics **504**, 3, 218–224 (2001).
URL [http://dx.doi.org/10.1016/S0370-2693\(01\)00307-0](http://dx.doi.org/10.1016/S0370-2693(01)00307-0)
- [33] H. A. Bethe; *Energy Production in Stars*; Phys. Rev. **55**, 5, 434–456 (1939).
URL <http://dx.doi.org/10.1103/PhysRev.55.434>
- [34] J. N. Bahcall, N. A. Bahcall, and G. Shaviv; *Present Status of the Theoretical Predictions for the ^{37}Cl Solar-Neutrino Experiment*; Phys. Rev. Lett. **20**, 21, 1209–1212 (1968).
URL <http://dx.doi.org/10.1103/PhysRevLett.20.1209>
- [35] J. N. Bahcall, A. M. Serenelli, and S. Basu; *New Solar Opacities, Abundances, Helioseismology, and Neutrino Fluxes*; The Astrophysical Journal **621**, 1, L85–L88 (2005).
URL <http://dx.doi.org/10.1086/428929>

- [36] B. T. Cleveland, T. Daily, R. Davis *et al.*; *Update on the measurement of the solar neutrino flux with the Homestake chlorine detector*; Nuclear Physics B (Proceedings Supplements) **38**, 1-3, 47–53 (1995).
URL [http://dx.doi.org/10.1016/0920-5632\(94\)00732-B](http://dx.doi.org/10.1016/0920-5632(94)00732-B)
- [37] J. N. Bahcall, M. H. Pinsonneault, and G. J. Wasserburg; *Solar models with helium and heavy-element diffusion*; Rev. Mod. Phys. **67**, 4, 781–808 (1995).
URL <http://dx.doi.org/10.1103/RevModPhys.67.781>
- [38] W. Hampel, G. Heusser, J. Kiko *et al.*; *GALLEX solar neutrino observations: Results for GALLEX IV*; Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics **447**, 127–133 (1999).
URL [http://dx.doi.org/10.1016/S0370-2693\(96\)01121-5](http://dx.doi.org/10.1016/S0370-2693(96)01121-5)
- [39] J. N. Abdurashitov, E. L. Faizov, V. N. Gavrin *et al.*; *Results from SAGE (The Russian-American gallium solar neutrino experiment)*; Physics Letters B **328**, 1-2, 234–248 (1994).
URL [http://dx.doi.org/10.1016/0370-2693\(94\)90454-5](http://dx.doi.org/10.1016/0370-2693(94)90454-5)
- [40] P. Anselmann, W. Hampel, G. Heusser *et al.*; *Solar neutrinos observed by GALLEX at Gran Sasso*; Physics Letters B **285**, 4, 376–389 (1992).
URL [http://dx.doi.org/10.1016/0370-2693\(92\)91521-A](http://dx.doi.org/10.1016/0370-2693(92)91521-A)
- [41] E. Gaisser, T. K.; Engel, R.; Resconi; *Cosmic Rays and Particle Physics*; Cambridge University Press (1990); ISBN 9780521016469.
URL <http://www.cambridge.org/uk/catalogue/catalogue.asp?isbn=0521326672>
- [42] F. Reines, M. F. Crouch, T. L. Jenkins *et al.*; *Evidence for High-Energy Cosmic-Ray Neutrino Interactions*; Phys. Rev. Lett. **15**, 9, 429–433 (1965).
URL <http://dx.doi.org/10.1103/PhysRevLett.15.429>
- [43] C. V. Achar, M. G. K. Menon, V. S. Narasimham *et al.*; *Detection of muons produced by cosmic ray neutrinos deep underground*; Physics Letters **18**, 2, 196–199 (1965).
URL [http://dx.doi.org/10.1016/0031-9163\(65\)90712-2](http://dx.doi.org/10.1016/0031-9163(65)90712-2)
- [44] T. J. Haines, R. M. Bionta, G. Blewitt *et al.*; *Calculation of Atmospheric Neutrino-Induced Backgrounds in a Nucleon-Decay Search*; Phys. Rev. Lett. **57**, 16, 1986–1989 (1986).
URL <http://dx.doi.org/10.1103/PhysRevLett.57.1986>
- [45] K. S. Hirata, T. Kajita, M. Koshiba *et al.*; *Experimental study of the atmospheric neutrino flux*; Physics Letters B **205**, 2, 416–420 (1988).
URL [http://dx.doi.org/10.1016/0370-2693\(88\)91690-5](http://dx.doi.org/10.1016/0370-2693(88)91690-5)
- [46] W. A. Mann; *Atmospheric neutrinos and the oscillations bonanza*; Int. J. Mod. Phys. **A15S1**, 229–256 (2000).
URL <http://dx.doi.org/10.1142/S0217751X00005176>
- [47] B. Pontecorvo; *Neutrino Experiments and the Problem of Conservation of Leptonic Charge*; Sov. Phys. JETP **26**, 984–988 (1968).

- [48] V. Gribov and B. Pontecorvo; *Neutrino astronomy and lepton charge*; Physics Letters B **28**, 7, 493–496 (1969).
URL [http://dx.doi.org/10.1016/0370-2693\(69\)90525-5](http://dx.doi.org/10.1016/0370-2693(69)90525-5)
- [49] B. Pontecorvo; *Mesonium and anti-mesonium*; Sov. Phys. JETP **6**, 429 (1957).
- [50] D. Casper, R. Becker-Szendy, C. B. Bratton *et al.*; *Measurement of atmospheric neutrino composition with the IMB-3 detector*; Phys. Rev. Lett. **66**, 20, 2561–2564 (1991).
URL <http://dx.doi.org/10.1103/PhysRevLett.66.2561>
- [51] R. Becker-Szendy, C. B. Bratton, D. Casper *et al.*; *Electron- and muon-neutrino content of the atmospheric flux*; Phys. Rev. D **46**, 9, 3720–3724 (1992).
URL <http://dx.doi.org/10.1103/PhysRevD.46.3720>
- [52] Y. Fukuda, T. Hayakawa, K. Inoue *et al.*; *Atmospheric ν_μ/ν_e ratio in the multi-GeV energy range*; Physics Letters B **335**, 2, 237–245 (1994).
URL [http://dx.doi.org/10.1016/0370-2693\(94\)91420-6](http://dx.doi.org/10.1016/0370-2693(94)91420-6)
- [53] J. N. Bahcall; *Solar Models and Solar Neutrinos*; Physica Scripta **T121**, 46–50 (2005).
URL <http://dx.doi.org/10.1088/0031-8949/2005/T121/006>
- [54] Z. Maki, M. Nakagawa, and S. Sakata; *Remarks on the Unified Model of Elementary Particles*; Progress of Theoretical Physics **28**, 5, 870 (1962).
URL <http://dx.doi.org/10.1143/PTP.28.870>
- [55] J. N. Bahcall, C. M. Gonzalez-Garcia, and C. Pena-Garay; *Before and After: How has the SNO NC measurement changed things?*; Journal of High Energy Physics **2002**, 07, 54 (2002).
URL <http://dx.doi.org/10.1088/1126-6708/2002/07/054>
- [56] A. Y. Smirnov; *The MSW effect and solar neutrinos*; in: *Neutrino telescopes. Proceedings, 10th International Workshop, Venice, Italy, March 11-14, 2003. Vol. 1+2*; 23–43 (2003).
- [57] L. Wolfenstein; *Neutrino oscillations in matter*; Phys. Rev. D **17**, 9, 2369–2374 (1978).
URL <http://dx.doi.org/10.1103/PhysRevD.17.2369>
- [58] S. P. Mikheev and A. r. Y. Smirnov; *Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos*; Sov. J. Nucl. Phys. **42**, 913–917 (1985).
- [59] S. P. Mikheev and A. r. Y. Smirnov; *Resonant amplification of neutrino oscillations in matter and solar neutrino spectroscopy*; Nuovo Cim. **C9**, 17–26 (1986).
URL <http://dx.doi.org/10.1007/BF02508049>
- [60] K. Eguchi, S. Enomoto, K. Furuno *et al.*; *First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance*; Phys. Rev. Lett. **90**, 2, 21802 (2003).
URL <http://dx.doi.org/10.1103/PhysRevLett.90.021802>

- [61] T. Araki, K. Eguchi, S. Enomoto *et al.*; *Measurement of Neutrino Oscillation with KamLAND: Evidence of Spectral Distortion*; Phys. Rev. Lett. **94**, 8, 81801 (2005).
URL <http://dx.doi.org/10.1103/PhysRevLett.94.081801>
- [62] A. Bandyopadhyay, S. Choubey, R. Gandhi *et al.*; *The Solar neutrino problem after the first results from KamLAND*; Phys. Lett. **B559**, 121–130 (2003).
URL [http://dx.doi.org/10.1016/S0370-2693\(03\)00340-X](http://dx.doi.org/10.1016/S0370-2693(03)00340-X)
- [63] P. C. de Holanda and A. r. Y. Smirnov; *LMA MSW solution of the solar neutrino problem and first KamLAND results*; JCAP **0302**, 1 (2003).
URL <http://dx.doi.org/10.1088/1475-7516/2003/02/001>
- [64] G. L. Fogli, E. Lisi, A. Marrone *et al.*; *Evidence for Mikheyev-Smirnov-Wolfenstein effects in solar neutrino flavor transitions*; Phys. Lett. **B583**, 149–156 (2004).
URL <http://dx.doi.org/10.1016/j.physletb.2003.12.045>
- [65] T. Mannel; *Theory and Phenomenology of CP Violation*; Nuclear Physics B - Proceedings Supplements **167**, 115–119 (2007).
URL <http://dx.doi.org/10.1016/j.nuclphysbps.2006.12.083>
- [66] T. Ohlsson, H. Zhang, and S. Zhou; *Radiative corrections to the leptonic Dirac CP-violating phase*; Phys. Rev. D **87**, 1, 13012 (2013).
URL <http://dx.doi.org/10.1103/PhysRevD.87.013012>
- [67] T. Ohlsson, H. Zhang, and S. Zhou; *Probing the leptonic Dirac CP-violating phase in neutrino oscillation experiments*; Physical Review D - Particles, Fields, Gravitation and Cosmology **87**, 5, 1–8 (2013).
URL <http://dx.doi.org/10.1103/PhysRevD.87.053006>
- [68] DUNE Collaboration; *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): The LBNF and DUNE Projects*; arXiv:1601.05471 [physics.ins-det] **1** (2016).
- [69] K. Abe, H. Aihara, C. Andreopoulos *et al.*; *Physics potential of a long-baseline neutrino oscillation experiment using a J-PARC neutrino beam and Hyper-Kamiokande*; Progress of Theoretical and Experimental Physics **2015**, 5, 053C02 (2015).
URL <http://dx.doi.org/10.1093/ptep/ptv061>
- [70] F. Kaether, W. Hampel, G. Heusser *et al.*; *Reanalysis of the Gallex solar neutrino flux and source experiments*; Physics Letters B **685**, 1, 47–54 (2010).
URL <http://dx.doi.org/10.1016/j.physletb.2010.01.030>
- [71] J. N. Abdurashitov, V. N. Gavrin, V. V. Gorbachev *et al.*; *Measurement of the solar neutrino capture rate with gallium metal. III. Results for the 2002–2007 data-taking period*; Phys. Rev. C **80**, 1, 15807 (2009).
URL <http://dx.doi.org/10.1103/PhysRevC.80.015807>
- [72] B. Aharmim, S. N. Ahmed, A. E. Anthony *et al.*; *Combined analysis of all three phases of solar neutrino data from the Sudbury Neutrino Observatory*; Phys. Rev. C **88**, 2, 25501 (2013).
URL <http://dx.doi.org/10.1103/PhysRevC.88.025501>

- [73] A. Gando, Y. Gando, H. Hanakago *et al.*; *Reactor on-off antineutrino measurement with KamLAND*; Phys. Rev. D **88**, 3, 33001 (2013).
URL <http://dx.doi.org/10.1103/PhysRevD.88.033001>
- [74] R. Wendell, C. Ishihara, K. Abe *et al.*; *Atmospheric neutrino oscillation analysis with subleading effects in Super-Kamiokande I, II, and III*; Phys. Rev. D **81**, 9, 92004 (2010).
URL <http://dx.doi.org/10.1103/PhysRevD.81.092004>
- [75] M. G. Aartsen, M. Ackermann, J. Adams *et al.*; *Determining neutrino oscillation parameters from atmospheric muon neutrino disappearance with three years of IceCube DeepCore data*; Phys. Rev. D **91**, 7, 72004 (2015).
URL <http://dx.doi.org/10.1103/PhysRevD.91.072004>
- [76] P. Adamson, I. Anghel, C. Backhouse *et al.*; *Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS*; Phys. Rev. Lett. **110**, 25, 251801 (2013).
URL <http://dx.doi.org/10.1103/PhysRevLett.110.251801>
- [77] P. Adamson, I. Anghel, C. Backhouse *et al.*; *Electron Neutrino and Antineutrino Appearance in the Full MINOS Data Sample*; Phys. Rev. Lett. **110**, 17, 171801 (2013).
URL <http://dx.doi.org/10.1103/PhysRevLett.110.171801>
- [78] K. Abe, J. Adam, H. Aihara *et al.*; *Precise Measurement of the Neutrino Mixing Parameter θ_{23} from Muon Neutrino Disappearance in an Off-Axis Beam*; Phys. Rev. Lett. **112**, 18, 181801 (2014).
URL <http://dx.doi.org/10.1103/PhysRevLett.112.181801>
- [79] P. Adamson, C. Ader, M. Andrews *et al.*; *First measurement of muon-neutrino disappearance in NOvA*; Phys. Rev. D **93**, 5, 51104 (2016).
URL <http://dx.doi.org/10.1103/PhysRevD.93.051104>
- [80] M. C. Gonzalez-Garcia, M. Maltoni, and T. Schwetz; *Updated fit to three neutrino mixing: status of leptonic CP violation*; Journal of High Energy Physics **2014**, 11, 52 (2014).
URL [http://dx.doi.org/10.1007/JHEP11\(2014\)052](http://dx.doi.org/10.1007/JHEP11(2014)052)
- [81] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni *et al.*; *Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity*; Journal of High Energy Physics **2017**, 1, 87 (2017).
URL [http://dx.doi.org/10.1007/JHEP01\(2017\)087](http://dx.doi.org/10.1007/JHEP01(2017)087)
- [82] F. P. An, J. Z. Bai, A. B. Balantekin *et al.*; *Observation of Electron-Antineutrino Disappearance at Daya Bay*; Phys. Rev. Lett. **108**, 17, 171803 (2012).
URL <http://dx.doi.org/10.1103/PhysRevLett.108.171803>
- [83] J. K. Ahn, S. Chebotaryov, J. H. Choi *et al.*; *Observation of Reactor Electron Antineutrinos Disappearance in the RENO Experiment*; Phys. Rev. Lett. **108**, 19, 191802 (2012).
URL <http://dx.doi.org/10.1103/PhysRevLett.108.191802>

- [84] K. Abe, J. Adam, H. Aihara *et al.*; *Observation of Electron Neutrino Appearance in a Muon Neutrino Beam*; Phys. Rev. Lett. **112**, 6, 61802 (2014).
URL <http://dx.doi.org/10.1103/PhysRevLett.112.061802>
- [85] P. Adamson, C. Ader, M. Andrews *et al.*; *First Measurement of Electron Neutrino Appearance in NOvA*; Phys. Rev. Lett. **116**, 15, 151806 (2016).
URL <http://dx.doi.org/10.1103/PhysRevLett.116.151806>
- [86] K. Abe, J. Amey, C. Andreopoulos *et al.*; *Combined Analysis of Neutrino and Antineutrino Oscillations at T2K*; Phys. Rev. Lett. **118**, 15, 151801 (2017).
URL <http://dx.doi.org/10.1103/PhysRevLett.118.151801>
- [87] P. Adamson, L. Aliaga, D. Ambrose *et al.*; *Constraints on Oscillation Parameters from ν_e Appearance and ν_μ Disappearance in NOvA*; Phys. Rev. Lett. **118**, 23, 231801 (2017).
URL <http://dx.doi.org/10.1103/PhysRevLett.118.231801>
- [88] K. Abe, J. Amey, C. Andreopoulos *et al.*; *Measurement of neutrino and antineutrino oscillations by the T2K experiment including a new additional sample of ν_e interactions at the far detector*; Phys. Rev. D **96**, 9, 92006 (2017).
URL <http://dx.doi.org/10.1103/PhysRevD.96.092006>
- [89] V. N. Aseev, A. I. Belesev, A. I. Berlev *et al.*; *Upper limit on the electron antineutrino mass from the Troitsk experiment*; Phys. Rev. D **84**, 11, 112003 (2011).
URL <http://dx.doi.org/10.1103/PhysRevD.84.112003>
- [90] C. Kraus, B. Bornschein, L. Bornschein *et al.*; *Final results from phase II of the Mainz neutrino mass search in tritium β -decay*; The European Physical Journal C - Particles and Fields **40**, 4, 447–468 (2005).
URL <http://dx.doi.org/10.1140/epjc/s2005-02139-7>
- [91] Planck Collaboration, P. A. R. Ade, N. Aghanim *et al.*; *Planck 2013 results. XVI. Cosmological parameters*; Astronomy & Astrophysics **571**, A16 (2014).
URL <http://dx.doi.org/10.1051/0004-6361/201321591>
- [92] DUNE Collaboration; *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): The Physics Program for DUNE at LBNF*; arXiv:1512.06148 [physics.ins-det] **2** (2015).
- [93] DUNE Collaboration; *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Long Baseline Neutrino Facility for DUNE*; arXiv:1601.05823 [physics.ins-det] **3** (2016).
- [94] DUNE Collaboration; *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): The DUNE Detectors at LBNF*; arXiv:1601.02984 [physics.ins-det] **4** (2016).
- [95] *Deep Underground Neutrino Experiment*.
URL dunescience.org

- [96] S. Amerio, S. Amoruso, M. Antonello *et al.*; *Design, construction and tests of the ICARUS T600 detector*; Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **527**, 3, 329–410 (2004).
URL <http://dx.doi.org/10.1016/j.nima.2004.02.044>
- [97] C. Anderson, M. Antonello, B. Baller *et al.*; *The ArgoNeuT detector in the NuMI low-energy beam line at Fermilab*; Journal of Instrumentation **7**, 10, P10019 (2012).
URL <http://dx.doi.org/10.1088/1748-0221/7/10/P10019>
- [98] F. Cavanna, M. Kordosky, J. Raaf *et al.*; *LArIAT: Liquid Argon In A Testbeam*; arXiv:1406.5560 [physics.ins-det] (2014).
- [99] R. Acciarri, C. Adams, R. An *et al.*; *Design and construction of the MicroBooNE detector*; Journal of Instrumentation **12**, 02, P02017 (2017).
URL <http://dx.doi.org/10.1088/1748-0221/12/02/P02017>
- [100] B. Baller, C. Bromberg, N. Buchanan *et al.*; *Liquid Argon Time Projection Chamber research and development in the United States*; Journal of Instrumentation **9**, 05, T05005 (2014).
URL <http://dx.doi.org/10.1088/1748-0221/9/05/T05005>
- [101] D. R. Nygren; *The Time Projection Chamber - A New 4π Detector for Charged Particles*; eConf **C740805**, PEP-0144, 58–78 (1974).
- [102] C. Rubbia; *The Liquid Argon Time Projection Chamber: A New Concept For Neutrino Detectors*; Technical Report CERN-EP-INT-77-08 (1977).
- [103] M. Soderberg; *The MicroBooNE Proposal*; MicroBooNE Public Note **70** (2008).
URL <http://microboone-docdb.fnal.gov/cgi-bin/ShowDocument?docid=70>
- [104] V. Chepel and H. Araújo; *Liquid noble gas detectors for low energy particle physics*; Journal of Instrumentation **8**, 04, R04001 (2013).
URL <http://dx.doi.org/10.1088/1748-0221/8/04/R04001>
- [105] J. Thomas and D. A. Imel; *Recombination of electron-ion pairs in liquid argon and liquid xenon*; Phys. Rev. A **36**, 2, 614–616 (1987).
URL <http://dx.doi.org/10.1103/PhysRevA.36.614>
- [106] S. Amoruso, M. Antonello, P. Aprili *et al.*; *Study of electron recombination in liquid argon with the ICARUS TPC*; Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **523**, 3, 275–286 (2004).
URL <http://dx.doi.org/10.1016/j.nima.2003.11.423>
- [107] S. Kubota, A. Nakamoto, T. Takahashi *et al.*; *Recombination luminescence in liquid argon and in liquid xenon*; Phys. Rev. B **17**, 6, 2762–2765 (1978).
URL <http://dx.doi.org/10.1103/PhysRevB.17.2762>
- [108] R. Acciarri, C. Adams, R. An *et al.*; *Noise Characterization and Filtering in the MicroBooNE Liquid Argon TPC*; Journal of Instrumentation **12**, 08, P08003 (2017).
URL <http://dx.doi.org/10.1088/1748-0221/12/08/P08003>

- [109] P. Derwent, S. Holmes, J. Hunt *et al.*; *Proton Improvement Plan-II (PIP-II)*; Technical Report ProjectX Document 1232 (2013).
URL <http://pip2.fnal.gov/index.html>
- [110] P. Huber and J. Kopp; *Two experiments for the price of one? The role of the second oscillation maximum in long baseline neutrino experiments*; *Journal of High Energy Physics* **2011**, 3, 13 (2011).
URL [http://dx.doi.org/10.1007/JHEP03\(2011\)013](http://dx.doi.org/10.1007/JHEP03(2011)013)
- [111] LBNE Collaboration; *Long-Baseline Neutrino Experiment (LBNE) Project: The LBNE Project 1* (2012).
- [112] LBNE Collaboration; *Long-Baseline Neutrino Experiment (LBNE) Project: Detectors At The Near Site 3* (2012).
- [113] LBNE Collaboration; *Long-Baseline Neutrino Experiment (LBNE) Project: Liquid Argon Detector At The Far Site*; LBNE Document Database **4**, 1 (2012).
URL http://lbne2-docdb.fnal.gov/0048/004892/012/CDR-LArFD-volume-final-2012Oct22_smaller.pdf
- [114] M. B. Avanzini; *The LAGUNA-LBNO Project*; *Physics Procedia* **61**, 524–533 (2015).
URL <http://dx.doi.org/10.1016/j.phpro.2014.12.117>
- [115] HEPAP Subcommittee; *Building for Discovery: Strategic Plan for U.S. Particle Physics in the Global Context* (2014).
URL https://science.energy.gov/~media/hep/hepap/pdf/May-2014/FINAL_P5_Report_053014.pdf
- [116] F. Bossi, E. De Lucia, J. Lee-Franzini *et al.*; *Precision Kaon and Hadron Physics with KLOE*; *Riv. Nuovo Cim.* **31**, 531–623 (2008).
URL <http://dx.doi.org/10.1393/ncr/i2008-10037-9>
- [117] S. Dimopoulos, S. Raby, and F. Wilczek; *Proton decay in supersymmetric models*; *Physics Letters B* **112**, 2, 133–136 (1982).
URL [http://dx.doi.org/10.1016/0370-2693\(82\)90313-6](http://dx.doi.org/10.1016/0370-2693(82)90313-6)
- [118] M. Convery and Z. Djurcic; *The 35-Ton Liquid Argon TPC Prototype for the Long Baseline Neutrino Experiment (Poster)*; in: *Neutrino 2014 Conference Proceedings* (2014).
- [119] T. Kutter; *Proposal for a Full-Scale Prototype Single-Phase Liquid Argon Time Projection Chamber and Detector Beam Test at CERN*; Technical Report CERN-SPSC-2015-020. SPSC-P-351; CERN; Geneva (2015).
URL <https://cds.cern.ch/record/2022751>
- [120] L. Agostino, B. Andrieu, R. Asfandiyarov *et al.*; *LBNO-DEMO: Large-scale neutrino detector demonstrators for phased performance assessment in view of a long-baseline oscillation experiment*; arXiv:1409.4405 [physics.ins-det] (2014).
- [121] CERN Courier; *Neutrinos take centre stage* (2016).
URL <http://cerncourier.com/cws/article/cern/65502>

- [122] D. Finley, W. Jaskierny, C. Kendziora *et al.*; *Work at FNAL to achieve long electron drift lifetime in liquid argon*, FERMILAB-TM-2385-E (2006).
URL <http://lss.fnal.gov/archive/test-tm/2000/fermilab-tm-2385-e.pdf>
- [123] R. Andrews, W. Jaskierny, H. Jostlein *et al.*; *A system to test the effects of materials on the electron drift lifetime in liquid argon and observations on the effect of water*; Nucl. Instrum. Meth. **A608**, 251–258 (2009).
URL <http://dx.doi.org/10.1016/j.nima.2009.07.024>
- [124] A. Curioni, B. T. Fleming, W. Jaskierny *et al.*; *A regenerable filter for liquid argon purification*; Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **605**, 3, 306–311 (2009).
URL <http://dx.doi.org/10.1016/j.nima.2009.04.020>
- [125] B. Rebel, M. Adamowski, W. Jaskierny *et al.*; *Results from the Fermilab materials test stand and status of the liquid argon purity demonstrator*; J. Phys. Conf. Ser. **308**, 12023 (2011).
URL <http://dx.doi.org/10.1088/1742-6596/308/1/012023>
- [126] P. Cennini, S. Cittolin, L. Dumps *et al.*; *Argon purification in the liquid phase*; Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **333**, 2, 567–570 (1993).
URL [http://dx.doi.org/10.1016/0168-9002\(93\)91209-6](http://dx.doi.org/10.1016/0168-9002(93)91209-6)
- [127] P. Benetti, A. Bettini, E. Calligarich *et al.*; *A three-ton liquid argon time projection chamber*; Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **332**, 3, 395–412 (1993).
URL [http://dx.doi.org/10.1016/0168-9002\(93\)90297-U](http://dx.doi.org/10.1016/0168-9002(93)90297-U)
- [128] F. Arneodo, P. Benetti, M. Bonesini *et al.*; *Performance of a liquid argon time projection chamber exposed to the CERN West Area Neutrino Facility neutrino beam*; Phys. Rev. D **74**, 11, 112001 (2006).
URL <http://dx.doi.org/10.1103/PhysRevD.74.112001>
- [129] G. Carugno, B. Dainese, F. Pietropaolo *et al.*; *Electron lifetime detector for liquid argon*; Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **292**, 3, 580–584 (1990).
URL [http://dx.doi.org/10.1016/0168-9002\(90\)90176-7](http://dx.doi.org/10.1016/0168-9002(90)90176-7)
- [130] A. Hahn, M. Adamowski, D. Montanari *et al.*; *The LBNE 35 Ton Prototype Cryostat*; in: *Proceedings, 21st Symposium on Room-Temperature Semiconductor X-ray and Gamma-ray Detectors (RTSD 2014): Seattle, WA, USA, November 8-15, 2014*; 7431158 (2014).
URL <http://dx.doi.org/10.1109/NSSMIC.2014.7431158>
- [131] T. Tope and Others; *Extreme argon purity in a large, non-evacuated cryostat*; AIP Conf. Proc. **1573**, 1169–1175 (2014).
URL <http://dx.doi.org/10.1063/1.4860838>

- [132] M. Adamowski, B. Carls, E. Dvorak *et al.*; *The Liquid Argon Purity Demonstrator*; *Journal of Instrumentation* **9**, 07, P07005 (2014).
URL <http://dx.doi.org/10.1088/1748-0221/9/07/P07005>
- [133] C. Bromberg, B. Carls, D. Edmunds *et al.*; *Design and operation of LongBo: a 2 m long drift liquid argon TPC*; *Journal of Instrumentation* **10**, 07, P07015 (2015).
URL <http://dx.doi.org/10.1088/1748-0221/10/07/P07015>
- [134] D. Montanari and Others; *First scientific application of the membrane cryostat technology*; *AIP Conf. Proc.* **1573**, 1664–1671 (2014).
URL <http://dx.doi.org/10.1063/1.4860907>
- [135] D. Montanari, M. Adamowski, A. Hahn *et al.*; *Performance and results of the LBNE 35 ton membrane cryostat prototype*; *Physics Procedia* **67**, 308–313 (2015).
URL <http://dx.doi.org/10.1016/j.phpro.2015.06.092>
- [136] *Fermi National Accelerator Laboratory Visual Media Services* (2015).
URL vms.fnal.gov
- [137] T. Alion; *35-ton Geometry* (2014).
URL https://cdcvs.fnal.gov/redmine/projects/lbnecode/wiki/LBNE_Geometries
- [138] 35-ton Collaboration; *Photon Detector System Performance in the DUNE 35-ton prototype LAr-TPC detector*; Publication in preparation .
- [139] A. Artikov, J. Budagov, I. Chirikov-Zorin *et al.*; *Design and construction of new central and forward muon counters for CDF II*; *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **538**, 1–3, 358–371 (2005).
URL <http://dx.doi.org/10.1016/j.nima.2004.09.009>
- [140] M. Stancari; *35-ton Counter Locations* (2015).
URL https://cdcvs.fnal.gov/redmine/projects/35ton/wiki/TSU_Counter_Locations
- [141] G. D. Geronimo, A. D’Andragora, S. Li *et al.*; *Front-End ASIC for a Liquid Argon TPC*; *IEEE Transactions on Nuclear Science* **58**, 3, 1376–1385 (2011).
URL <http://dx.doi.org/10.1109/TNS.2011.2127487>
- [142] C. Thorn, G. D. Geronimo, A. D’Andragora *et al.*; *Cold Electronics Development for the LBNE LAr TPC*; *Physics Procedia* **37**, 1295–1302 (2012).
URL <http://dx.doi.org/10.1016/j.phpro.2012.02.473>
- [143] R. Herbst, R. Claus, M. Freytag *et al.*; *Design of the SLAC RCE Platform: A general purpose ATCA based data acquisition system*; in: *2014 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*; 1–4 (2014).
URL <http://dx.doi.org/10.1109/NSSMIC.2014.7431254>
- [144] J. Anderson, P. De Lurgio, Z. Djurcic *et al.*; *SiPM Signal Processor User Manual*; Technical report (2016).
URL <https://indico.fnal.gov/getFile.py/access?contribId=1&resId=0&materialId=7&confId=12392>

- [145] N. Barros, S. Glavin, P. Keener *et al.*; *The Penn Trigger Board*; Technical report (2016).
URL <https://indico.fnal.gov/getFile.py/access?contribId=13&resId=0&materialId=slides&confId=12639>
- [146] K. Biery, C. Green, J. Kowalkowski *et al.*; *artdaq: An Event-Building, Filtering, and Processing Framework*; IEEE Transactions on Nuclear Science **60**, 5, 3764–3771 (2013).
URL <http://dx.doi.org/10.1109/TNS.2013.2251660>
- [147] *art*.
URL <http://art.fnal.gov>
- [148] C. Green, J. Kowalkowski, M. Paterno *et al.*; *The art framework*; Journal of Physics: Conference Series **396**, 2, 22020 (2012).
URL <http://dx.doi.org/10.1088/1742-6596/396/2/022020>
- [149] J. Freeman; *Courtesy of John Freeman, Fermilab* (2014).
- [150] The LZ Collaboration; *LUX-ZEPLIN (LZ) Conceptual Design Report*; arXiv:1509.02910 [physics.ins-det] (2015).
- [151] A. Blatter, A. Ereditato, C. C. Hsu *et al.*; *Experimental study of electric breakdowns in liquid argon at centimeter scale*; Journal of Instrumentation **9**, 04, P04006 (2014).
URL <http://dx.doi.org/10.1088/1748-0221/9/04/P04006>
- [152] N. McConkey, N. Spooner, M. Thiesse *et al.*; *Cryogenic CMOS cameras for high voltage monitoring in liquid argon*; Journal of Instrumentation **12**, 03, P03014 (2017).
URL <http://dx.doi.org/10.1088/1748-0221/12/03/P03014>
- [153] *35-ton Phase II LAr filling monitored by camera system* (2016).
URL <https://web.fnal.gov/collaboration/DUNE/ImagesforWeb/LAr-in-DUNE-prototype.gif>
- [154] M. Convery, T. Junk, and M. Stancari; *Lessons Learned from the DUNE 35-ton Prototype*; DUNE Document Database **1315** (2016).
URL <https://docs.dunescience.org/cgi-bin/private/RetrieveFile?docid=1315&filename=lessons-learned-dune%284%29.pdf&version=4>
- [155] L. Bagby, M. Johnson, B. Kirby *et al.*; *35-ton Noise*; DUNE Document Database **1834** (2016).
- [156] *LArSoft*.
URL larsoft.org
- [157] E. D. Church; *LArSoft: A Software Package for Liquid Argon Time Projection Drift Chambers*; arXiv:1311.6774 [physics.ins-det] (2013).
- [158] R. Pordes and E. Snider; *The Liquid Argon Software Toolkit (LArSoft): Goals, Status and Plan*; 38th International Conference on High Energy Physics (ICHEP 2016): Chicago, IL, USA, August 3-10, 2016 182 (2016).

- [159] E. L. Snider; *LArSoft: Toolkit for Simulation, Reconstruction and Analysis of Liquid Argon TPC Neutrino Detectors*; in: *22nd International Conference on Computing in High Energy and Nuclear Physics (CHEP 2016) San Francisco, CA, October 14-16, 2016* (2016).
URL <http://lss.fnal.gov/archive/2017/conf/fermilab-conf-17-052-cd.pdf>
- [160] C. Andreopoulos, A. Bell, D. Bhattacharya *et al.*; *The GENIE neutrino Monte Carlo generator*; Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **614**, 1, 87–104 (2010).
URL <http://dx.doi.org/10.1016/j.nima.2009.12.009>
- [161] J. Allison, K. Amako, J. Apostolakis *et al.*; *Geant4 developments and applications*; IEEE Transactions on Nuclear Science **53**, 1, 270–278 (2006).
URL <http://dx.doi.org/10.1109/TNS.2006.869826>
- [162] R. Acciarri and Others; *Michel Electron Reconstruction Using Cosmic-Ray Data from the MicroBooNE LArTPC*; arXiv:1704.02927 [physics.ins-det] (2017).
- [163] W. Shockley; *Currents to Conductors Induced by a Moving Point Charge*; Journal of Applied Physics **9**, 10, 635–636 (1938).
URL <http://dx.doi.org/10.1063/1.1710367>
- [164] S. Ramo; *Currents Induced by Electron Motion*; Proceedings of the IRE **27**, 9, 584–585 (1939).
URL <http://dx.doi.org/10.1109/JRPROC.1939.228757>
- [165] The MicroBooNE Collaboration; *A Method to Extract the Charge Distribution Arriving at the TPC Wire Planes in MicroBooNE*; MicroBooNE Public Note **1017** (2016).
URL <https://www-microboone.fnal.gov/publications/publicnotes/MICROBOONE-NOTE-1017-PUB.pdf>
- [166] *Garfield simulation*.
URL <http://garfield.web.cern.ch/garfield/>
- [167] B. Baller; *Cluster Crawler Manual* (2014).
URL <https://cdcvs.fnal.gov/redmine/documents/727>
- [168] M. Antonello and Others; *Precise 3D track reconstruction algorithm for the ICARUS T600 liquid argon time projection chamber detector*; Adv. High Energy Phys. **2013**, 260820 (2013).
URL <http://dx.doi.org/10.1155/2013/260820>
- [169] Brookhaven National Laboratory; *Wire-Cell LArTPC Reconstruction*.
URL <http://lar.bnl.gov/wire-cell/>
- [170] N. Grant and T. Yang; *Neutrino Energy Reconstruction in the DUNE Far Detector*; Technical report (2017).

- [171] A. C. Kaboth; *Detecting the Invisible Universe with Neutrinos and Dark Matter*; Ph.D. thesis; Massachusetts Institute of Technology (2012).
URL <http://hdl.handle.net/1721.1/76980>
- [172] M. Wallbank; *The BlurredCluster Reconstruction Technique*; DUNE Document Database **54** (2016).
URL <https://docs.dunescience.org/cgi-bin/private/RetrieveFile?docid=54&filename=blurredClustering.pdf&version=1>
- [173] M. Wallbank; *BlurredCluster Algorithm* (2016).
URL <http://larsoft.org/single-record/?pdb=110>
- [174] M. Wallbank; *The EMShower Reconstruction Technique*; DUNE Document Database **1369** (2016).
URL <https://docs.dunescience.org/cgi-bin/private/RetrieveFile?docid=1369&filename=emShower.pdf&version=1>
- [175] M. Wallbank; *EMShower Algorithm* (2016).
URL <http://larsoft.org/single-record/?pdb=113>
- [176] B. Kirby; *35-ton Cold Electronics Testing, Integration and Commissioning Experience*; Technical report (2016).
URL https://docs.dunescience.org/cgi-bin/private/RetrieveFile?docid=1704&filename=20160826_35t_Experience.pdf&version=1
- [177] J. Insler; *Module to Unstick ADC Codes*; Technical report (2016).
URL <https://indico.fnal.gov/getFile.py/access?contribId=2&resId=0&materialId=slides&confId=11627>
- [178] C. Haggmann, D. Lange, and D. Wright; *Cosmic-ray shower generator (CRY) for Monte Carlo transport codes*; in: *2007 IEEE Nuclear Science Symposium Conference Record*; volume 2; 1143–1146 (2007).
URL <http://dx.doi.org/10.1109/NSSMIC.2007.4437209>
- [179] D. Brailsford; *Longitudinal Electron Diffusion*; Technical report (2016).
URL <https://indico.fnal.gov/getFile.py/access?contribId=3&resId=0&materialId=slides&confId=13014>
- [180] The MicroBooNE Collaboration; *Cosmic Shielding Studies at MicroBooNE*; MicroBooNE Public Note **1005** (2016).
URL <https://www-microboone.fnal.gov/publications/publicnotes/MICROBOONE-NOTE-1005-PUB.pdf>
- [181] J. S. Marshall and M. A. Thomson; *The Pandora software development kit for pattern recognition*; *The European Physical Journal C* **75**, 9, 439 (2015).
URL <http://dx.doi.org/10.1140/epjc/s10052-015-3659-3>
- [182] G. Barker, N. Grant, and M. Haigh; *Particle ID and energy reconstruction in the DUNE far detector*; Technical report (2016).
URL <https://indico.fnal.gov/getFile.py/access?contribId=85&sessionId=20&resId=0&materialId=slides&confId=10613>

-
- [183] Duke University; *Prob3++* (2012).
URL <http://webhome.phy.duke.edu/~raw22/public/Prob3++/>
- [184] A. Hoecker, P. Speckmayer, J. Stelzer *et al.*; *TMVA - Toolkit for Multivariate Data Analysis*; arXiv:physics/0703039 [physics.data-an] (2007).
URL <https://arxiv.org/pdf/physics/0703039.pdf>
- [185] R. Acciarri, C. Adams, J. Asaadi *et al.*; *A study of electron recombination using highly ionizing particles in the ArgoNeuT Liquid Argon TPC*; *Journal of Instrumentation* **8**, 08, P08005 (2013).
URL <http://stacks.iop.org/1748-0221/8/i=08/a=P08005>
- [186] A. Aurisano, A. Radovic, D. Rocco *et al.*; *A convolutional neural network neutrino event classifier*; *Journal of Instrumentation* **11**, 09, P09001 (2016).
URL <http://stacks.iop.org/1748-0221/11/i=09/a=P09001>

Appendix A

APA Gap-Crossing Muons: Gap Measurements

This appendix contains all figures used to make measurement of the gaps between the APA frames in the 35-ton experiment, as discussed in Section 7.2 and presented in Table 7.1. The method used is described in detail in Section 7.2.1.1.

The following pages contain the information for each of the four gaps in the long drift region in the 35-ton detector: DV1/DV3 (Figure A.1), DV1/DV5 (Figure A.2), DV3/DV7 (Figure A.3) and DV5/DV7 (Figure A.4). The top figure shows the initial measurement of the z -offset, with the double-peak effect and the associated initial extraction of a z -offset; the centre figure shows the measurement of the x -offset after correcting for the initial z -offset estimate; and the bottom figure shows the measurement of the z -offset after correcting for the determined x -offset. In each case, the initial z -offset is found by fitting a parabola between the peaks except for the DV3/DV7 case. In this instance, the number of tracks in the sample is too low to make meaningful measurements of this effect and instead a coarse binning is chosen to mask the peaks and a Gaussian fit is used on the resulting distribution. As is evident from the following figures, this works reasonably well in the case of small sample size.




(a) Initial z -offset(b) x -offset(c) z -offset

Fig. A.1 DV1/DV3 gap.




(a) Initial z -offset(b) x -offset(c) z -offset

Fig. A.2 DV1/DV5 gap.




(a) Initial z -offset(b) x -offset(c) z -offset

Fig. A.3 DV3/DV7 gap.




(a) Initial z -offset(b) x -offset(c) z -offset

Fig. A.4 DV5/DV7 gap.

Appendix B

DUNE Far Detector ν_e CC MVA Input Variables

This appendix shows the separation between ν_e and ν_μ, ν_τ events provided by the input variables to the DUNE far detector MVA-based ν_e CC selection, discussed in Section 8.3. The events were reconstructed using the developments presented in Chapter 6 and the selection applied as discussed in Chapter 8.

The variables are listed and described in Table 8.2 in Section 8.3.1. They are separated into event-level variables, shown in Figure B.1, variables pertaining to the longest reconstructed track in the event, shown in Figure B.2, and variables describing the shower with the highest energy in the event, shown in Figure B.3.

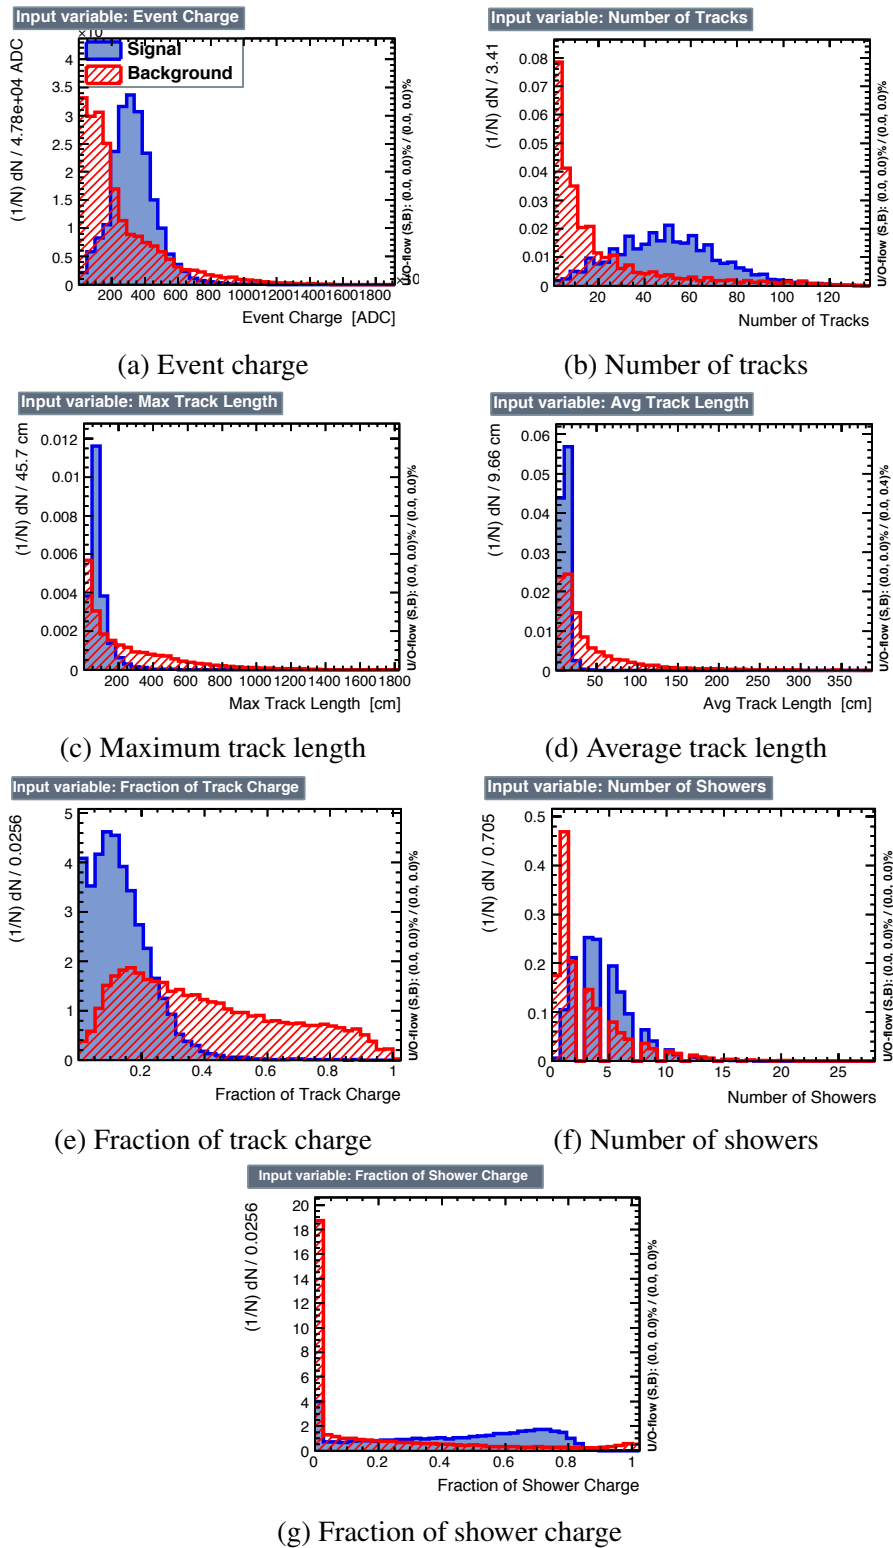


Fig. B.1 MVA input variables related to event-level information for the DUNE far detector ν_e CC analysis.

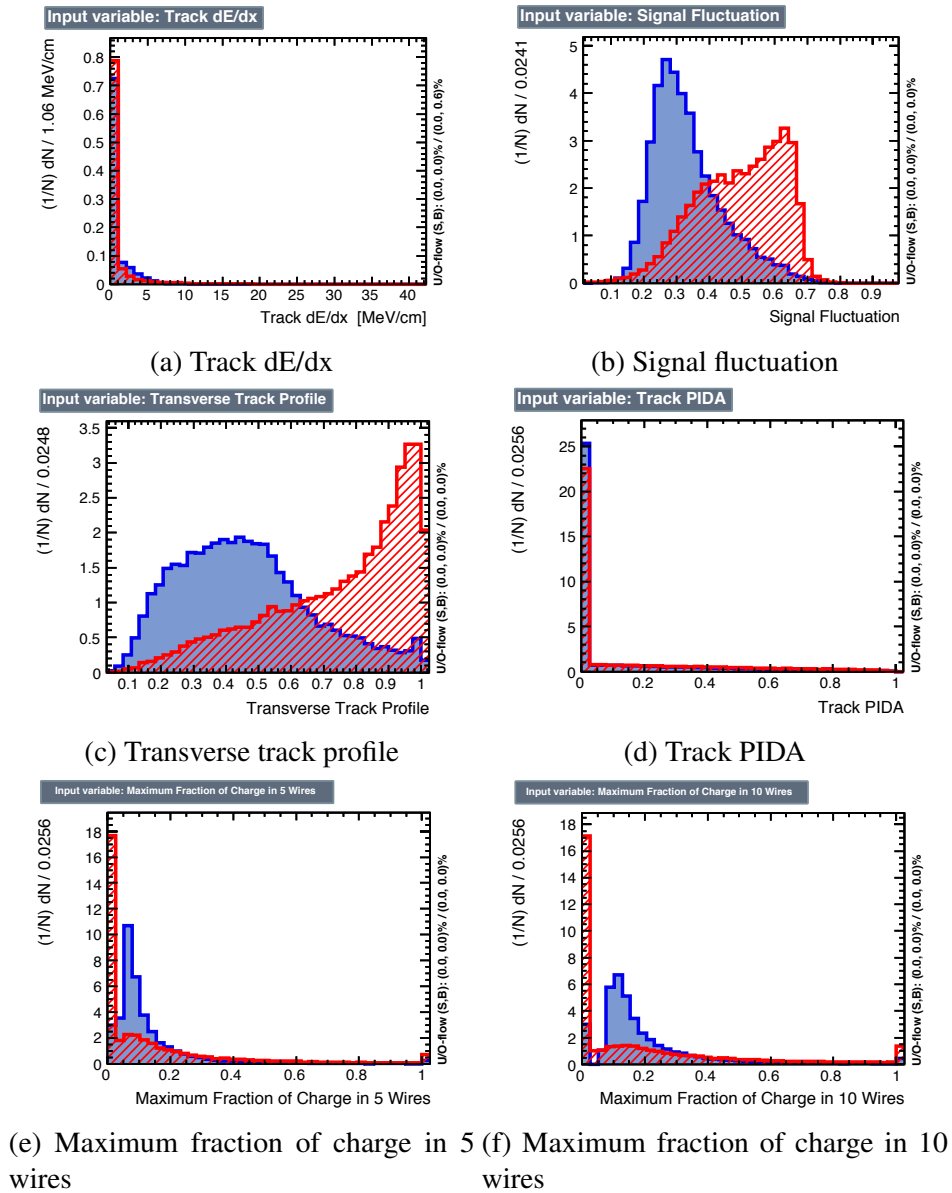
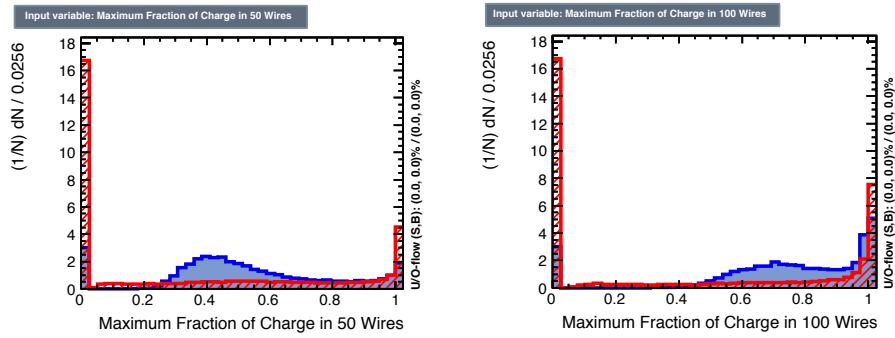



Fig. B.2 MVA input variables related to information about the longest reconstructed track in the event for the DUNE far detector ν_e CC analysis.




(g) Maximum fraction of charge in 50 wires (h) Maximum fraction of charge in 100 wires



(i) Track angle (x -direction)

(j) Track angle (y -direction)



(k) Track angle (z -direction)

(l) Fractional transverse energy

Fig. B.2 MVA input variables related to information about the longest reconstructed track in the event for the DUNE far detector ν_e CC analysis.

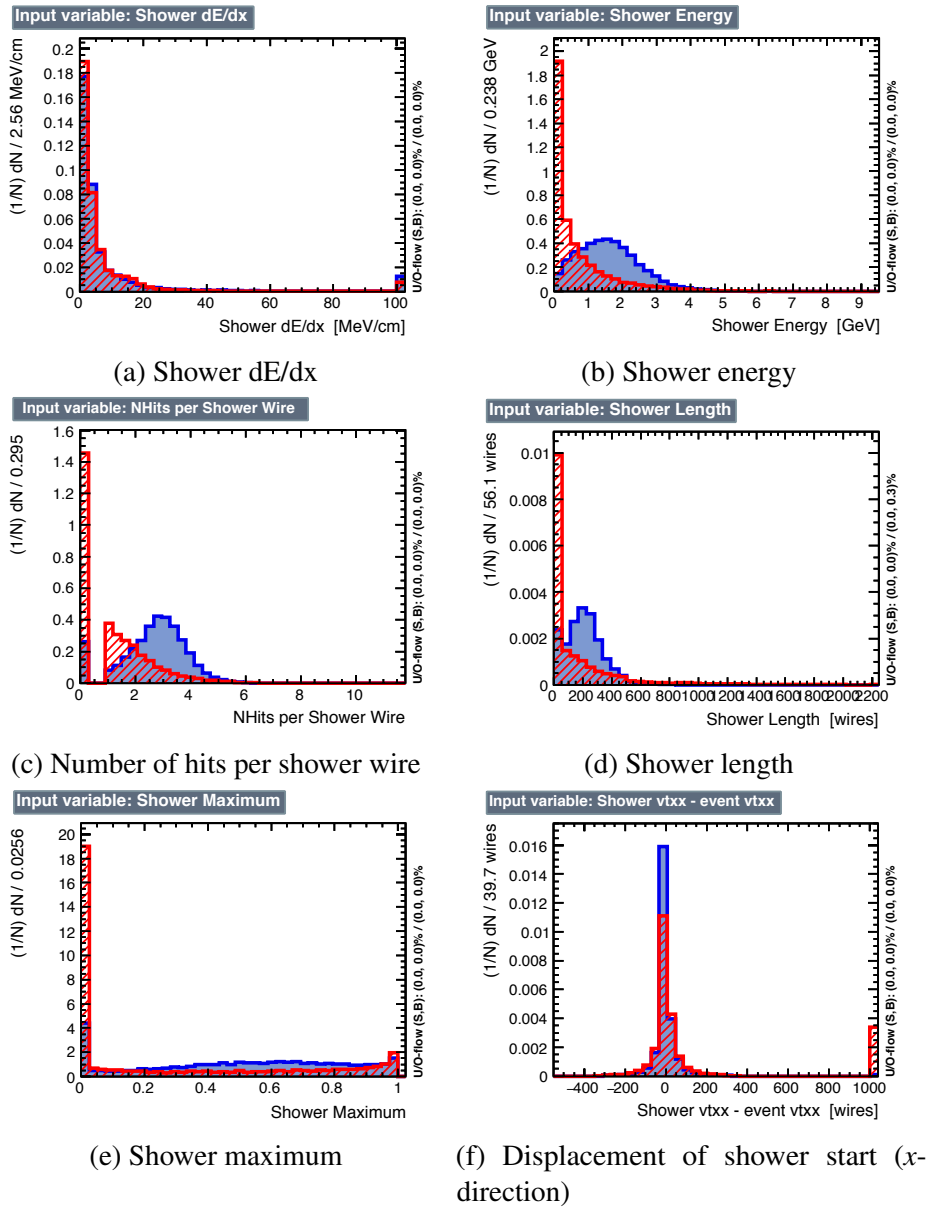


Fig. B.3 MVA input variables related to information about the highest energy shower in the event for the DUNE far detector ν_e CC analysis.






Fig. B.3 MVA input variables related to information about the highest energy shower in the event for the DUNE far detector ν_e CC analysis.