

UNIVERSITY OF BIRMINGHAM

A novel method for track fitting in TPCs, and interactive computing for nuclear laboratories

By

Angus Hollands

A thesis submitted to
the University of Birmingham
for the degree of
Doctor of Philosophy

Nuclear Physics Group
School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham
Aug 24, 2023

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

ABSTRACT

It has long been postulated that light nuclei may exhibit molecular and clustered structures. Clustering in non-alpha conjugate nuclei is of particular interest in demonstrating the role of valence nucleons in stabilising the nucleus. One such example is the stability of the Borromean nucleus ^{9}Be over its alpha-conjugate sibling ^{8}Be [1]. Given the existing works investigating these structures in ^{14}C, it is the object of this thesis to explore the symmetries of the mirror nucleus ^{14}O and investigate the existence of low-lying clustered states, and their respective properties. Novel approaches to track reconstruction in TPCs have been developed, and used to make a preliminary investigation into the spectrum of the ^{14}O nucleus.

Shortly following the onset of the COVID-19 pandemic in early 2020, a concerted effort in the School of Physics and Astronomy at the University of Birmingham arose to establish novel solutions for distance learning provision. In particular, the challenge of providing robust, flexible teaching of undergraduate laboratory skills was identified as a key limitation in responding to the exigent challenges posed by public health interventions during 2020 and 2021. To meet this need, a virtual nuclear physics laboratory was developed and deployed to University undergraduate students, and their engagement with this new resource recorded. An exploration of these findings, and future directions of the project, was subsequently made.

CONTENTS

I Front Matter **19**

1 Acknowledgements **21**

2 Introduction **23**

II Experimental Nuclear Physics **25**

3 Nuclear Models **27**

3.1 Macroscopic Models . . . 27

3.2 Microscopic Models . . . 32

3.3 Conclusion . . . 52

4 Carbon-14 and Oxygen-14 Mirror Nuclei **53**

4.1 Mirror Nuclei . . . 53

4.2 Predicted Structures in Carbon-14 and Oxygen-14 . . . 55

4.3 Conclusion . . . 61

5 Thick Target Experiments **63**

5.1 Resonant Scattering . . . 63

5.2 Conservation Rules . . . 65

5.3 Time-Projection Chambers . . . 66

5.4 Thick-Target Inverse Kinematics . . . 69

5.5 Conclusion . . . 71

6 The TexAT Detector for Rare Isotope Beam Experiments **73**

6.1 Experiment Design . . . 73

6.2 Detector Overview . . . 73
6.3 Readout Electronics . . . 78
6.4 Conclusion . . . 79

7 Track Finding and Fitting **81**
7.1 Conventional Techniques . . . 81
7.2 Facility Location Problem . . . 87
7.3 Kinematic Fitting . . . 101
7.4 Conclusion . . . 104

8 Resonant Elastic Scattering of α-particles with ^{10}C **107**
8.1 Partitioning . . . 109
8.2 Noise and Baseline Removal . . . 109
8.3 Response Estimation . . . 113
8.4 Gain Matching . . . 119
8.5 Signal Fitting . . . 121
8.6 Signal Ringing . . . 123
8.7 Silicon Calibration . . . 127
8.8 Gas Simulation . . . 134
8.9 Beam Origin . . . 136
8.10 Cluster Reconstruction . . . 139
8.11 Track Fitting . . . 142
8.12 Particle Identification . . . 146
8.13 Kinematics . . . 150
8.14 Conclusion . . . 159

III O-Lab Online Nuclear Laboratory **161**

9 O-Lab Online Nuclear Laboratory **163**
9.1 Motivations . . . 163
9.2 Virtual Laboratories . . . 165
9.3 Software and Infrastructure . . . 168
9.4 Content . . . 177
9.5 Analytics . . . 187
9.6 Future Work . . . 194

IV Discussion 197

10 Conclusions 199

11 Future Work 201

V End Matter 203

12 Appendix 205

13 Bibliography 209

LIST OF FIGURES

3.1 Illustration of the physical interpretation of the terms that comprise the liquid-drop nature of the SEMF. Figure by Daniel FR, distributed under a CC BY-SA 3.0 license. 28

3.2 Schematic illustration of the measured quadrupole moment Q for an axially symmetric rotor. (a) shows the configuration in which the rotor symmetry axis 3 is maximally aligned with the space-fixed axis z. (b) depicts a state in which the 3 axis is approximately perpendicular to the z axis; in this case, the 3-projected momentum is zero. Under the rotor model, both the oblate doubly-even and prolate odd-mass nuclei measurements are explained by an intrinsic prolate deformation. Adapted from Ref. [4]. 29

3.3 Relationship between the total angular momentum **I**, the intrinsic angular momentum **J**, the rotor angular momentum **R**, and the components M of **I** along the fixed axis z; and K the symmetry axis. Adapted from Ref. [4]. 31

3.4 Labelling of particle and hole states relative to a sharp Fermi surface, that follows from the restricted Slater determinant form of the HF ground state $|\text{HF}\rangle$. Figure taken from Ref. [5]. 35

3.5 Schematic diagram of the convergence of different HF trial wave-functions to local minima with distinct quadrupole moments Q corresponding to a deformed isomer and subsequent fissioning of the nucleus. Figure taken from Ref. [15]. 36

3.6 Single-particle energies for (a) a simple harmonic oscillator (SHO); (b) a SHO with an l^2 term; (c) a shell-model potential with l^2 and $\mathbf{l} \cdot \mathbf{s}$ terms. Figure by Morten Hjorth-Jensen, distributed under a CC BY-NC 4.0 license. 39

3.7 Illustration of two nucleon orbits around a deformed core. K_1 depicts an orbit that lies closer to the bulk of the nucleus than K_2, and therefore lies at lower energy. Adapted from Ref. [18]. 40

3.8 Schematic representation of the eigenstates resulting from anti-symmetrisation of Gaussian wave-packets in the limit of $\frac{d}{\sqrt{a}} \to 0$, where $d = 0.75\sqrt{a}$ is the displacement between each wave-packet, and a the real width parameter. Figure taken from Ref. [23]. 43

3.9 Density distributions of the AMD wavefunctions for ^{9}Li obtained by (a) variation *before* parity and angular momentum projection, (b) *after* parity and *before* angular momentum projection, and (c) *after* parity and angular momentum projection. It can be seen that only (c) exhibits beyond mean-field structure. The proton (ρ_p), neutron (ρ_n), and total (ρ_m) densities of the AMD wave-function are illustrated. Figure adapted from Ref. [21]. 45

3.10 Density distributions of a singular AMD wave-function for ^{9}Li obtained by variation *after* parity and *before* angular momentum projection for a deformation constraint of $\beta = 0.35$. The proton (ρ_p), neutron (ρ_n), and total (ρ_m) densities of the AMD wave-function are illustrated. See Fig. 3.9 for comparison to other variational solutions. Figure adapted from Ref. [21]. 46

3.11 A myriad of clustering phenomena are seen in light nuclei, varying from small clusters outside of a closed shell to alpha condensates. Figure taken from Ref. [34]. 46

3.12 One-dimensional molecular orbits for cluster states in ^{9}Be; showing (a) σ bonding, (b) σ anti-bonding, (c) π bonding, (d) π anti-bonding. In this 2D projection, the deformation axis lies in parallel to the plane normal. Figure adapted from Ref. [35]. 47

3.13 One-dimensional molecular orbits from atomic physics; showing (a) π bonding, (b) π anti-bonding, (c) σ bonding, (d) σ anti-bonding. Figure adapted from Ref. [36]. 48

3.14 Illustration of the spatial overlap of two p-shell orbits, corresponding to (a) π and (b) σ molecular orbitals. Figure adapted from Ref. [1]. 48

3.15 Neutron level scheme of the two-centre shell-model as a function of the deformation eccentricity $z_0 \in [0, R]$, where R is the nuclear radius. The y-axis intercepts correspond to spherical shell-model states, whilst at large separations the spherical shell-model states of the fragments are indicated [38]. 49

3.16 Single-particle level scheme of the Nilsson deformed harmonic oscillator potential, for $N < 50$. Solid lines indicate positive parity solutions, whilst dashed lines represent negative parity states. Each state is labelled by $[Nn_z\Lambda\Omega]$, where N is the major oscillator number, n_z is the oscillator quanta along the 3-axis, Λ the z projection of the orbital angular momentum, and Ω the z projection of the total angular momentum J. Figure taken from Ref. [42]. 50

3.17 Correlation diagram of molecular orbitals for a two-center α-α system with the $p_{3/2}$ shell-model orbit of the valence neutron in ^{5}He indicated. The location of the potential minimum is labelled R_{min}. Each orbital state is labelled by $[K, \pi, u/g]$, where K is the projection of the total angular momentum J onto the cluster axis, π the orbital parity, and u/g the ungerade/gerade orbital family. From Fig. 3.16, the $p_{3/2}$ shell-model orbit for separate centres can be identified with Nilsson orbits in the deformed model. Figure taken from Ref. [1]. 51

4.1 Level schemes of the ^{11}C and ^{11}B nuclei. The order and spacing between the levels is well preserved under exchange of the protons and neutrons, with only the $\frac{7}{2}^-$ and $\frac{1}{2}^+$ states changing in order. Generated from data taken from the ENSDF database as of March 24th, 2023. Version available at http://www.nndc.bnl.gov/ensarchivals/. 54

4.2 Low lying energy levels of $A = 14$ isobars, with adjusted energies relative to ^{14}N. Levels in ^{14}C and ^{14}O have $T = 1$, whilst those in ^{14}N have $T = 0$ apart from those levels in common with the other members of the isobar. Figure taken from Ref. [46]. 55

4.3 Energy diagram of shell-model degeneracies for integer deformation ratios in an axially deformed harmonic oscillator. The open circles indicate the shells occupied by valence neutrons, whilst the shaded dots denote the shells occupied by the isospin-degenerate $J_z = 0$ nucleon pairs that form an α-particle. Figure taken from Ref. [1]. 56

4.4 One-dimensional molecular orbits for cluster states in ^{13}C; showing π. Note that the number of nodes $[n_1 n_2 n_3]$ increases with increasing oscillator quanta. In this 2D projection, the deformation axis lies in parallel to the plane normal. Figure adapted from Ref. [35]. 56

4.5 Energy diagrams of molecular chain states in carbon isotopes, adapted from Ref. [41]. The black boxes indicate the first two-center non-rotational isomeric states. The predicted configurations for ^{14}C are shown in the blue box. 57

4.6 Schematic illustration of the deformation modes of linear-chain states. (a) shows the *breathing* mode in which the centre-separation, d, varies, whilst (b) depicts the *bending* mode that depends upon the bending angle θ. Figure taken from Ref. [49]. 58

4.7 Energy levels in ^{14}C for positive (a) and negative (b) parity states. Energies are measured relative to the ^{10}Be$+\alpha$ threshold, and the experimental states are plotted against the predictions for comparison. Figure taken from Ref. [51]. 59

4.8 Schematic illustration of the cluster structures anticipated in ^{14}C for various valence neutron configurations. (a) depicts the symmetric, positive parity $\alpha+n+\alpha+n+\alpha$ configuration with the valence neutrons distributed outside of the symmetry axis; (b) shows the reflection-asymmetric inversion doublets, based upon ^{10}Be $+ \alpha$ structure with the valence neutrons in the same covalent π-bond; and (c) illustrates the two oblate triangular-cluster states with σ-bonds between two α-particles, whose valence particles wavefunctions have (i) a node at the center of the nucleus with $J^\pi = 3^-$ and (ii) a maximum at the center of the nucleus with $J^\pi = 0^+$. The dashed lines indicate three sites between which the two valence neutrons are distributed. Figure adapted from Ref. [52]. 60

4.9 (a) 2^+ and (b) 4^+ states in ^{14}C and ^{14}O determined by AMD calculation. Close correspondence between the π-bond chain states is observed, whilst the σ-bond levels are expected to differ significantly in excitation energy and decay width. Figure taken from Ref. [57]. 61

5.1 Semi-classical illustration of the limiting extremes to nuclear reactions: formation of a transitory compound nucleus, and direct reactions. Figure taken from Ref. [60]. 63
5.2 Comparison between the (left) Breit-Wigner distribution and the (right) elastic scattering cross section. (right) includes interference between the shape-elastic and resonance-elastic interactions. Subfigures taken from Ref. [60]. 65
5.3 Operational principle of a gridded ion-chamber, in which charged particles ionise the gas to liberate electron-ion charge carriers that move under an applied field. The movement of charge carriers induces a charge upon the anode and cathode. The grid serves to remove the positional dependence of the signal amplitude when operated in electron-sensitive mode; only the electron component contributes to the measured voltage across the resistor R once the electrons enter into the grid-anode region. 67
5.4 The different operation regions of gas-filled detectors, according to the applied electric field. Only a limited range of field strengths corresponds to a proportional regime in the detector response. Figure taken from Ref. [68]. 69
5.5 Schematic illustration of a non-relativistic elastic scattering reaction (lab frame) in (a) forward kinematics, in which the lighter nucleus a is used as the projectile; and (b) inverse kinematics, whereby the heavy nucleus A impinges upon the light target. In each panel, the upper region indicates a potential post-scattering configuration. In the inverse regime, due to conservation of momentum and energy, the reaction is *forward projected*; both nuclei continue travelling forward in the lab frame. Meanwhile, in forward kinematics it is possible for the beam to achieve a (smaller) negative velocity in the lab frame. 70
5.6 Stopping power of ^{10}C and ^{4}He ions in a gas-mixture of 96% ^{4}He, 4% CO_2 at 405 torr, simulated using SRIM [72]. 71

6.1 A 3D model of the "final" innermost TexAT detector components, comprised of silicon (purple), thallium doped cesium iodide (turquoise), and MicroMeGaS (yellow) detectors. The ion counter is not pictured here. 74
6.2 A schematic diagram of the TexAT detector configuration, indicating the location of the various detectors and auxiliary apparatus. A black arrow indicates the direction of the beam, which first enters the detector through the ion-counter (IC, indicated). The segmented anode is visible in blue and purple at the top of the detector, whilst the field-cage (FC) and the array of CsI(Tl) + Si detectors are visible around the outside of the field cage. These detectors occupy the same gas cell as the MicroMeGaS (MM), which sits at the top of the detector. The ion drift plane is given by the normal of the MicroMeGaS. The beam direction, given by a vector between the window and silicon wall, is indicated by a white arrow. Figure taken from Ref. [76]. 74

6.3 Charge signal formation from charged particle ionisation. A small amplification region, with an electric field of the order ~kV/cm, can be seen between the micromesh and the anode strips. The larger drift/conversion region, with a field strength of ~10^2V/cm, is bounded by the cathode and the micromesh. Gas multiplication predominantly occurs in the amplification region; the drift/conversion region serves only to separate the ion pairs liberated by ionising particle tracks, such as those of scattered alpha particles. Figure adapted from Ref. [77]. 75

6.4 The segmentation plan of the MicroMeGaS anode. There are three distinct regions: a central column of 6×128 pads, and two side regions with 64 interleaved rows (strips) and columns (chains). The strips are formed from single electrodes, whilst the chains are formed from staggered readout pads that are electrically connected together. Note that the measurements in this figure were subsequently revised for fabrication. Figure adapted from Ref. [73]. 76

6.5 A 3D rendering of the solid-state detector telescopes that surround the TPC: downstream array (top left); upstream array (bottom left); side array (top right); bottom array (bottom right). Figure taken from Ref. [73]. 77

6.6 An illustration of the relationship between the experimental detector, readout electronics, and data acquisition and storage equipment in an experiment that uses the GET system [78]. 78

6.7 A schematic diagram of the shaper bypass facility of AGET chips. External shapers are used for the ion counter signals (see above). Figure adapted from Ref. [73]. 79

7.1 Flowchart detailing the classic discrete Hough transform. Once the algorithm has terminated, maxima in the voting space correspond to models (parameterisations) with high support (large numbers of votes). These maxima can be treated as candidates that well describe the data. 82

7.2 Flowchart outlining the RANSAC algorithm. 83

7.3 3D scatter plot of a simulated elastic scattering reaction between a ^{10}C beam and a ^{4}He gas target. The simulation parameters match the experimental parameters described in Section 8. The simulation was performed using a heavily modified variant of TexATSim, the software developed by the group at Texaa A&M for TexAT reaction studies. The beam direction is given by the $\hat{y}$ axis. Three track labels given by sequential RANSAC are indicated by the colour map, with a fourth outlier label corresponding to -1. It can clearly be seen that the scattered beam (blue) labelling is over-supported, as it extends into the incident beam (yellow) track clusters. 84

7.4 3D visualisation of the tracks reconstructed from Fig. 7.3 using sequential RANSAC. Only the track direction should be considered; these fits correspond to unbounded lines. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. It can be seen that whilst the scattered beam track lies close to the simulated track, a common-vertex fit will pull the vertex away from its true value. 85

7.5 3D scatter plot of a second simulated elastic scattering reaction between a ^{10}C beam and a ^{4}He gas target. The simulation parameters match the experimental parameters described in Section 8. The beam direction is given by the $\hat{y}$ axis. Three track labels are indicated by the colour map, with a fourth outlier label corresponding to -1. Like Fig. 7.3, the RANSAC fit is greedy, but in this figure it can be seen that the indicent beam (purple) is over-supported and extends into the scattered beam (yellow) track clusters. 85

7.6 3D visualisation of the tracks reconstructed from Fig. 7.5 using sequential RANSAC. Only the track direction should be considered; these fits correspond to unbounded lines. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. Like Fig. 7.4, it can be seen that whilst the scattered beam track lies close to the simulated track, a common-vertex fit will pull the vertex away from its true value. 86

7.7 3D scatter plot of a third simulated elastic scattering reaction between a ^{10}C beam and a ^{4}He gas target. The simulation parameters match the experimental parameters described in Section 8. The beam direction is given by the $\hat{y}$ axis. Three track labels are indicated by the colour map, with a fourth outlier label corresponding to -1. Whilst the correct number of active tracks (3) have been identified, it can be seen that the light-product track (blue) extends far into the incident beam clusters (green), and the incident beam clusters extend into the scattered beam region (yellow). The beam track is thus discontinuous. . . 86

7.8 3D visualisation of the tracks reconstructed from Fig. 7.7 using sequential RANSAC. Only the track direction should be considered; these fits correspond to unbounded lines. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. It can be seen that the poor labelling shown in Fig. 7.7 produces a poor fit, with the reconstructed vertex (given by the intersection of the blue scattered and heavy beam tracks) displaced outside the measured clusters. 87

7.9 An illustration of a 2D track described by a proposed track model. 88

7.10 Illustration of a 2D smoothness prior between two labellings. The highest costs are observed between unalike labels at small distances. 90

7.11 Simple two-vertex graph. 91

7.12 Simple two-vertex graph with terminals s, t. The maximum flow between the terminals is 10. 91

7.13 Simple two-vertex graph with terminals s, t. A valid cut is indicated by dotted lines. 91

7.14 Graph representing a two-node system whose optimal labelling is given in Table 7.2. By choosing appropriate edge weights, the cost of the minimum cut is equal to the optimal labelling cost given in Eqn. 7.10. Cut links are indicated with dotted lines. Smoothness costs that penalise localised discontinuities in the labelling are introduced through the addition *auxillary nodes* at label boundaries, i.e. the $\star$ node between Ω and A. The edges between these auxiliary nodes and their neighbours encode the smoothness penalty. 93

7.15 Variation of Fig. 7.14 depicting a two-node system whose labelling is already optimal. As such, the weight between A and the terminal ω is set to infinity, so that the ω label does not grow i.e. the α label does not shrink. 94

7.16 Process diagram of the PeARL algorithm 95

7.17 3D volume plot of the likelihood function for a point described by a line interval between $\begin{pmatrix} 4 & 4 & 4 \end{pmatrix}$ and $\begin{pmatrix} 16 & 16 & 16 \end{pmatrix}$ with $\sigma = 1.0$. 97

7.18 3D scatter plot of the simulated elastic scattering reaction shown in Fig. 7.3. Three track labels given by the PeARL algorithm are indicated by the colour map. In comparison to Fig. 7.3, it can be seen that the PeARL fit (with line intervals) preserves a common vertex through distinct track separation. 98

7.19 3D visualisation of the tracks reconstructed from Fig. 7.18 using PeARL. Both the track direction and end-point should be considered; these fits correspond to bounded line intervals. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. For the same reaction as depicted in Fig. 7.4, it can be seen that the PeARL fit (with line intervals) better preserves a common vertex through distinct track separation. The line-interval parameterisation makes it possible to clearly identify the estimated vertex from the end-point of the incident beam track. 98

7.20 3D scatter plot of the simulated elastic scattering reaction shown in Fig. 7.5. Three track labels given by the PeARL algorithm are indicated by the colour map. In comparison to Fig. 7.5, it can be seen that the PeARL fit (with line intervals) avoids overfitting of the incident beam (blue) track, although a reduced degree of overfitting is observed in the scattered beam (purple) region. 99

7.21 3D visualisation of the tracks reconstructed from Fig. 7.20 using PeARL. Both the track direction and end-point should be considered; these fits correspond to bounded line intervals. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. For the same reaction as depicted in Fig. 7.6, it can be seen that the PeARL fit (with line intervals) better reconstructs the angle of each track in agreement with the simulation. A common vertex is much more easily identified by the intersection of the reconstructed tracks. 100

7.22 3D scatter plot of the simulated elastic scattering reaction shown in Fig. 7.7. Three track labels given by the PeARL algorithm are indicated by the colour map. In comparison to Fig. 7.7, it can be seen that the PeARL fit (with line intervals) avoids the discontinuity in the reconstructed incident beam (blue) track. As such, the labelling of the scattered beam (yellow) track has more support. The degree of overfitting in the light product (purple) track is also much diminished, and all clusters are considered inliers. 100

7.23 3D visualisation of the tracks reconstructed from Fig. 7.22 using PeARL. Both the track direction and end-point should be considered; these fits correspond to bounded line intervals. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. For the same reaction as depicted in Fig. 7.8, it can be seen that the PeARL fit (with line intervals) produces a far more reasonable reconstruction of the track angles, with the closest-point-of-approach between any two tracks producing a good starting-point for vertex estimation. As evidenced by its length, the incident beam track is slightly overfit. A bigger discrepancy lies in the light-product track, which is visually distinct from the simulated track. Here it can be presumed that secondary scattering in the simulation leads to a bending in the track, which lies below the resolution of the MicroMeGaS. As such, only a single track is observed here. 101

7.24 Diagram illustrating a constrained optimisation problem, in which the constraint (red curve) touches an equipotential contour (blue curve) of the function being optimised. By definition, the constrained solution occur where the two curves are touching, in order for the solution to be a minimum. 102

8.1 Process diagram of the TexAT analysis. 108

8.2 FPN channel waveforms digitised by the GET acquisition system, alongside their binned mean, for a random AGET chip. Each FPN waveform has a different vertical offset, which can be modelled to first-order as contributing a constant offset in the computed average basliebaseline. In practice, this offset is unimportant; even when using the nearest-FPN channel for noise estimation, there is typically a non-zero offset that necessitates a baseline removal step. The x-axis is given in terms of AGET circular-buffer cells (see Section 6.3.1). 110

8.3 Alternative estimates of the electronic noise for a singular MicroMeGaS GET channel. (a) The raw GET waveform. (b) The subtraction of a single FPN waveform. (c) The removal of the FPN waveform average. The leading samples of each estimate is fit with a linear baseline. 111

8.4 Estimate of the signal baseline used to normalise MicroMeGaS samples. The locus of the minimum variance indicates the “flattest” portion of the waveform. 112

8.5 Estimate of the signal baseline for the negative polarity silicon samples. The leading ten samples (shaded region) are used to compute an average constant baseline. 113

8.6 A plot of the waveforms that are employed by the fitting process to derive an estimated response function. 114

8.7 GET response function for τ given in Eqn. 8.1. . . . 115
8.8 A waterfall plot of the response function estimates computed over many sequential iterations. It can be seen that the estimated response function becomes broader with each iteration, ultimately converging upon a stable shape. . . . 115
8.9 A plot of the convolution fit solution cost as a function of the free parameter τ. . . . 116
8.10 Comparison of the response function for different values of τ. . . . 116
8.11 A plot of the estimated response function over several iterations, with an initial estimate given by optimising Eqn. 8.1. It can be seen that the estimated response function becomes broader with each iteration. 117
8.12 A comparison between the response functions computed for different starting estimates of $F^{(1)}$. . . . 117
8.13 A plot of the waveforms that are averaged to compute the silicon response function. . . . 118
8.14 A visual comparison of the response function estimates produced from iterative estimation and signal averaging. (a) A comparison of the raw response functions. (b) A comparison of the Gaussian broadened silicon average response function with the iterative estimate from the MicroMeGaS detector samples. . . 119
8.15 Illustration of the significant zones within the MicroMeGaS anode segmentation. The strip-chain region (a), held at the same potential as the high-gain pads (c), is shaded in green. The high-gain pads region (c) is shaded in orange, whilst the low-gain pads (b) are coloured in blue. The beam direction is indicated by a black vertical arrow. . . . 120
8.16 Ratio of the mean collected charge between the low gain and high gain region boundary in the Micromegas pads. . . . 120
8.17 A plot of the peaks identified within a MicroMeGaS waveform sample using boosted GOLD deconvolution. It can clearly be seen that the signal is formed of two separate peaks, whose convolution cannot trivially be separated into distinct peaks. Experimentally, this corresponds to overlapping charge clusters (in the drift plane). . . . 121
8.18 A plot of the saturated MicroMeGaS waveform sample, and its GOLD deconvolution. . . . 122
8.19 A plot of the saturated MicroMeGaS waveform sample, and the mask used for fitting. . . . 122
8.20 A plot of silicon waveform sample, and its linear response fit. . . . 123
8.21 A waveform with the ringing phenomenon visible at $t \sim 345$. A distinct non-ringing component is also visible at $t \sim 254$. The shaded region indicates the span covering ringing in this channel. . . . 124
8.22 The distribution of peak widths for ringing and non-ringing events. . . . 124
8.23 A waveform from the same event as Fig. 8.21 in which the signal is saturated at $t \sim 345$. The shaded region indicates the span covering ringing signals in other channels. . . . 125

8.24 A 2D hitmap of the MicroMeGaS anode for a random event. (a) A raw event heatmap with ringing signals. Note that the majority of the elements in the left side region have been activated. The linear track formed in the central pads region does not suggest such a pattern will be observed in the side region. (b) A modified hitmap with ringing signals and invalid fits removed. The resulting hitmap more closely agrees with the predicted track yielded by the central pads region. 126
8.25 A waveform with the ringing phenomenon visible at $t \sim 345$, and the deconvolution of the baseline adjusted signal overlaid. A distinct non-ringing component is also visible at $t \sim 254$. The shaded region indicates the span covering ringing in this channel. It can be seen that the green peaks in the shaded region have much higher relative amplitudes, and much narrower widths. 127
8.26 An illustration of the band structure in a semiconductor. The dashed line indicates the Fermi level that separates the valence and conduction bands. Figure by Tim Starling & Mliu92, distributed under a CC BY-SA 4.0 license. 127
8.27 An illustrative diagram of the minority-majority charge carriers in an (a) p-type semiconductor and (b) n-type semiconductor. A surplus of mobile charge carriers can be seen in the conduction band. The charge carriers in p-type semiconductors (a) are termed holes, whilst those in n-type materials (b) are electrons. 128
8.28 An illustrative diagram of the charge-carrier distribution across a P-N junction. Additional graphs show the charge, field, and electric potential, across the junction. Figure by TheNoise, distributed under a CC BY-SA 3.0 license. 129
8.29 A schematic diagram of the (a) forward and (b) side silicon detector arrays indicating quadrants with incomplete calibrations in pink. 130
8.30 A 2D histogram of the charge distribution measured by each silicon detector for a fixed calibration source containing 241Am, 148Gd, 239Pu, and 244Cm. 131
8.31 The charge distribution for a random silicon detector, superimposed with a plot of the cumulative distribution. The dashed lines indicate particular fractions of the cumulative distribution at which the peaks are localised. 132
8.32 A combined plot of the calibrated charge distribution histograms. The characteristic decay peak energies are shown in the staggered X axis. 133
8.33 The x position distribution of drifted electrons for y=5.90 cm fit with a normal distribution. The standard deviation region is shaded. 135
8.34 The standard deviation of the final electron transverse position as a function of the drift distance. Against the scatter plot, a linear regression indicates the slope and intercept of the data. 136
8.35 Distribution of the charge-cluster time coordinate in the early pads region of the MicroMeGaS detector. The origin of this coordinate corresponds to the modal value, which lies around 350 cells. 137

8.36 Fit of the histogram in Fig. 8.35 using a constant function and Gaussian distribution. Above the x-axis, the original histogram and minimiser result are superimposed. Below, the residual distribution of the fit against the measurement is shown. 138
8.37 3D scatter plot of the cluster reconstruction taken from experiment. The charge of each cluster is indicated by its colour. Three distinct tracks can be identified; a heavier incoming and scatter track, and a light recoiling track. A discontinuity in the light recoil track immediately after the interaction vertex is clearly visible. This follows from the low stopping power of the corresponding ion, whereby the energy deposited in the MicroMeGaS is insufficient to trigger the low-gain channels. Once the ion enters the high-gain region it forms a coherent track. 139
8.38 A 3D diagram of the pad reconstruction procedure. The electrical contacts that comprise the MicroMeGaS anode are represented with a plane of opaque quadrilaterals. The shaded volume above these contacts indicates the measured *drift time* of the charge cluster incident upon the element. The volume above the contact upon which the current signal is induced is shaded opaque. 140
8.39 A 3D diagram of the strip-chain pairing procedure. The electrical contacts that comprise the MicroMeGaS anode are represented with a plane of opaque quadrilaterals. The shaded volume above these contacts indicates the measured *drift time* of the charge cluster incident upon the element (strip/chain). The volume above the contact upon which the current signal is induced is shaded opaque, whilst the remaining (electrically connected) contacts are rendered partially transparent. It can be seen that the pairs with the closest drift times correspond to the elements over which the track passes. 141
8.40 A 3D diagram of the strip-chain pairing procedure for planar tracks. See Fig. 8.39 for more details. It can be seen that the pairs with the closest drift times no longer correspond to the elements over which the track passes. 141
8.41 A 3D plot of the track reconstruction for a single event using the PEARL method described in Section 7.2.4. Each distinct track is labelled with a distinct colour, and the line of best fit superimposed. Due to the uncertainty of the cluster positions in the pads region, the beam track is found to be longer than it is physically expected to be. This can be corrected in a subsequent constrained fit. The upper track corresponds to a beam-like pile-up track, whilst the lower track triplet belong to a single scattering reaction. 142
8.42 Histogram of the track multiplicity found during fitting of the usable experimental runs. A majority of events have a single track; these correspond to events in which the light product is not observed. 143
8.43 2D hitmaps for the forward silicon detector array, formed by the intersection of reconstructed tracks, whose endpoints lie in the final sector of the MicroMeGaS, with the silicon plane. The hitmap for all events is shown in (a). The subset of hits outside the central region is plotted in (b). 145

8.44 2D relative hitmaps for the forward silicon detector array, formed by the intersection of reconstructed tracks, whose endpoints lie in the final sector of the MicroMeGaS, with the silicon plane. The position of each reconstructed hit is plotted relative to the detector centroid. 145
8.45 Experimental distribution of the stopping power of ^{10}C ions in a gas-mixture of 96% ^{4}He, 4% CO_2 held at 405 torr. SRIM prediction for an after-window beam energy of 30.5 MeV indicated by the dashed line. 147
8.46 Experimental distribution of the stopping power of ^{10}C ions in a gas-mixture of 96% ^{4}He, 4% CO_2 held at 405 torr. MSTAR prediction for an after-window beam energy of 29.7 MeV indicated by the dashed line. 147
8.47 Stopping power curves predicted by SRIM for ^{1}H, ^{4}He, and ^{10}C ions in a gas-mixture of 96% ^{4}He, 4% CO_2 with density ρ. Each curve overlaps with the others for a given ion energy. 148
8.48 Illustration of a particle telescope. A thin ΔE detector measures the z-dependent energy loss of the incident radiation within a thin silicon target. The ion punches through this material, and is fully stopped within the "thick" silicon detector. Conventionally, the energy lost in both detectors is taken to be the total energy of the particle. However, to first approximation the contribution lost in the thin detector can be ignored. 149
8.49 2D histogram of the energy deposited in an individual silicon detector against the energy lost in the final N elements of the MicroMeGaS detector. Two loci are visible; the proton region at (5000, 250), and an upper alpha curve. 150
8.50 Schematic diagram of a static-target scattering reaction in a beam-centric coordinate system. The $\hat{1}$ vector is given by the beam direction, the $\hat{3}$ vector given by the reaction plane normal, and the $\hat{2}$ vector taken to be perpendicular to $\hat{1}$ and $\hat{3}$. Each particle $\mathbf{p}_i$ can be resolved along these canonical axis vectors. Each vector $\mathbf{p}_i$ indicates the momentum of the reaction participant i, whose components $\mathbf{p}_i^{(j)}$ have been resolved along the beam-centric coordinate system vectors $\hat{j}$. 151
8.51 A plot of the vertex position in non zero-degree scattering reactions as a function of energy measured in the silicon detectors. A clear locus of bins corresponding to elastic reactions can be observed, with a maximum near (-100, 7500). 152
8.52 A plot of the expected vertex position in zero-degree scattering as a function of energy measured in the central silicon detectors. The two curves show the interaction vertex for (solid line) elastic and (dashed line) inelastic scatters that produces the corresponding silicon hit for a known beam energy given by Section 8.13.4. 154

8.53 A histogram of the excitation energy of the ^{14}O nucleus integrated over all non-zero-degree scattering angles. States published in the ENSDF database are plotted in dashed lines, whilst the various AMD predictions discussed in this work for ^{14}O are, where plausible, indicated by solid lines. Clear peak-like structures are visible, against a wider background distribution. Despite this, further enhancement of these data through secondary fitting is required to provide more conclusive agreement with predicted states. . 155
8.54 A histogram of the excitation energy of the ^{14}O nucleus computed from the near-zero-degree scattering channel. States published in the ENSDF database are plotted in dashed lines, whilst the various AMD predictions discussed in this work for ^{14}O are, where plausible, indicated by solid lines. 156
8.55 A histogram of the silicon energy distribution, with inset axes indicating a candidate for the maximum value. 157
8.56 A histogram of the difference in projectile energy computed directly (through SRIM calculations of energy loss in the gas), and via indirect kinematic reconstruction (using the known track angles and measured energy deposit in the silicon detectors). 157

9.1 Screenshot of the `Lorenz` demonstration notebook within the JupyterLab next-generation IDE. Rich computational narratives can easily be constructed by interweaving code and markup (Markdown) cells. Rich outputs, such as interactive widgets and rendered plots, can be embedded to support exploratory analyses. 170
9.2 Screenshot of the `Lorenz` demonstration notebook presented using Voila. Voila templates can be customised and configured: from re-arranging content or changing the formatting, to building entire applications. 171
9.3 Architecture diagram of a Kubernetes deployment. Several black-box containers are run within Kubernetes “pods”, and communicate with one another over a shared network. Through horizontal and vertical autoscaling, Kubernetes can ensure that available resources are optimally distributed across the running pods. Healthy deployments are ensured through monitoring and subsequent redeployment, providing a mechanism for failure tolerance. Figure by Khtan66, distributed under a CC BY-SA 4.0 license. 174
9.4 Screenshot of the JupyterHub administrator panel, showing a list of previously seen users, and their server status. Individual user servers can be stopped and started by the administrator; enabling remote troubleshooting in the event of a problem. 177

9.5 (a) Schematic rendering of the page layout of a typical O-Lab experiment. Content is divided into several regions; the sidebar, which contains the global navigation links; the navbar, which contains the local navigation links; and the main content, in which the experiment details are presented. Each white rectangle represents a "card", which can contain rich content such as rich text, interactive widgets, or images. (b) Rendering of the "Getting Started" resource presented to new users of O-Lab. This page gives an overview of the purpose of O-Lab, its core features and user-guide, and where to access additional support. 178
9.6 Screenshot of the `60-detector-comparison.ipynb` notebook, described in Table 9.2. Each experiment followed a similar structure; the *Introduction* page was composed from several distinct cards: a list of tasks, a schematic diagram of the apparatus, additional details, safety notes, and any further reading. 179
9.7 Screenshot of the `10-nai-detector.ipynb` notebook introduction page, bearing close resemblance to Fig. 9.6. 180
9.8 Screenshot of the `10-nai-detector.ipynb` notebook details page, discussing the finer details of how the experimental equipment is operated, and points of consideration during operation. 182
9.9 Screenshot of the `10-nai-detector.ipynb` notebook properties page, discussing the various properties of a detector, and how they are measured. 183
9.10 Screenshot of the `10-nai-detector.ipynb` notebook experiment page, providing interactive widgets to configure the apparatus and produce a spectrum. 185
9.11 Screenshot of the `70-gamma-attenuation.ipynb` notebook experiment page, providing interactive widgets to explore the attenuation of gamma rays through a range of absorber materials and thicknesses. 186
9.12 The unique number of OLAB users aggregated by day. Anonymous user IDs were generated by hashing the unique Canvas student IDs. 188
9.13 The number of times each user started a server. Anonymous user IDs were generated by hashing the unique Canvas student IDs. 188
9.14 The start times for each user session, grouped according to whether the user started more than one session during the trial. 189
9.15 The total session duration of each user. 189
9.16 The duration of each started user server. 190
9.17 The number of concurrent users (users whose `start` and `end` times overlap). 191
9.18 Experiment notebook access patterns during the trial. 192
9.19 The number of notebook sessions that users opened during the trial. 193
9.20 The number of unique-notebook sessions that users opened during the trial. 193
9.21 The number of users that opened each notebook at least once during the trial. Anonymous user IDs were generated by hashing the unique Canvas student IDs. 194

9.22 Screenshot of the Voila preview panel running in JupyterLab courtesy of the `jupyterlab-preview` extension. With this extension, customised dashboards can be opened by the user, whilst conventional cellular notebooks can also be edited and run. 195
9.23 Screenshot of the Jitsi video chat panel running in JupyterLab courtesy of the `jupyterlab-videochat` extension. Jitsi provides the conventional video-conferencing features that students are already familiar with, conveniently integrated with the JupyterLab application. . . 196

LIST OF TABLES

7.1 Terms corresponding to the FLP interpretation of track labelling. 88
7.2 Lowest cost (correct) labelling for a two-node system 92

8.1 Isotopes present in the four-α calibration source, and their α decay radiation. From ENSDF database as of March 24th, 2023. Version available at http://www.nndc.bnl.gov/ensarchivals/. 130
8.2 The gas and field parameters that govern charge transport within the active chamber 134

9.1 High-level list of software packages used to build O-Lab. 169
9.2 List of preliminary O-Lab experiments 179

Part I

Front Matter

CHAPTER

ONE

ACKNOWLEDGEMENTS

PhDs are often referred to as journeys. Though it is tempting, therefore, to begin these acknowledgements with a quote, or reference to an epic voyage, to do so would feel foolish; I have not achieved a super-human feat, and there were no supernatural beings or powers at play. Yet, to ensure that the screenwriters of "2023: A Space-less Odyssey" are not made aware of this work and left without sufficient material, I will indulge such an interpretation (and commit to the full mustard-gogo-boots-level imagery).

With reflection upon the great epics such as Homer's Odyssey, Beowulf, or the IMAX screening of Interstellar, it suffices to start with a guiding force. To my supervisors, Tzany Wheldon and Carl Wheldon: thank you for your direction and counsel. Tzany, my painting from our stateside art class will serve as constant, slightly wonky testament to the role that creativity and perspective must play in one's life, and emphasise the enduring importance of getting paint on the canvas. Carl, I have reviewed my Grammarly statistics and the data shows that I've never before been so focused on my punctuation and grammar. Whilst I joke about this, and the pluarality of data, you have instilled in me the importance of the little details.

In every epic, one cannot fail to identify those characters whose contributions to the narrative before some pivotal chapter have hitherto gone unrecognised. It is my hope that Jack Bishop should land one of these leading roles, and that his salary be excessive and obscene (Jack, I've done all I can, I hope you can give a good audition). I sincerely wish to express my gratitude to Jack, Marina, and the wider team at the cyclotron institute for their knowledge and patience.

No expedition ever starts or finishes without friends or family. To my *companions*, you know who you are (but I don't, so *please*, get in touch). To embrace the melodramatic, I wouldn't be where I am now without you. Which is to say that I'd be somewhere else, and I suspect that where *that* world is, it does not have wind turbines or ABBA.

To my parents, I forgive you for not having any Physics qualifications that might have been of use to me. What you have imparted was far more important, and sown long before I started this journey. You have given me a work ethic, a belief in myself, and most importantly, safe and happy homes that I have always been able to return to. To Meg *MD*, I hereby note in formal academic literature, that you were the first twin to become a doctor. Though, should the call go out, in the spirit

of equality we shall both assist any mid-flight emergency. Thank you for taking the pressure off by ensuring that mum and dad would have at least one successful offspring. To my brother Sam, thank you for the welcome interludes helping you recover your email passwords provided, for teaching me to hate speed-chess, and for your excellent taste in memes. To the extensive Hollands and Fitzgerald clans: support comes in many different guises. Your never ending encouragement has meant more than you will ever know. Finally, to my additional family (the Rowes), thank you for always offering me both the first and the last potato.

Like all good models of reality, the metaphorical lens that I have used to describe the last four years is imperfect; there is one person still to be acknowledged, yet there are no illustrations, cursive or monospaced, sufficient to recognise her contributions. Ellie, this was an expedition in which you stood beside me each step of the way. Thank you for the unswerving kindness, interminable patience, and immeasurable time and energy that you have demonstrated unfalteringly over the years. Also, thank you for the thesaurus. Like most epics, this main character has been exceedingly lucky.

Along this journey, I count myself lucky not to have made any mortal enemies. At least, should they exist, I am unaware of them. Then again, isn't the greatest enemy of all: time?[1] Thus, as I come to the end of this first Epic journey, I am filled with nothing but gratitude. I am thankful for those I met along the way and hope you will be enduring characters in my next chapters, and I in yours. For those already in my life, thank you for your patience. I have no excuse not to do the dishes any more.

> The atmosphere was dominated by a feeling of mutual respect … There existed in those days a true camaraderie in our profession, whatever the small differences in our approach.
>
> ---The Remains of the Day, Kazuo Ishiguro.

---

[1] This, too, I am attestedly unaware of.

CHAPTER

TWO

INTRODUCTION

Experimental nuclear physics is fundamentally a discipline that is concerned with the strong interaction between quarks. At the low-energy scale, it is not the strong interaction itself, but the *residual* strong interaction that acts between nucleons, that is of interest. Investigations into the nature of this interaction, and the physics that emerges from its presence in the many-body system of the nucleus, lie at the heart of the field. An introduction to the contemporary understanding of the nucleus, and the role of the strong interaction in its behaviour, is given in Section 3.

Experiments that probe the nature of the residual strong interaction must ultimately proceed through scattering experiments, whereby systems of nucleons are brought together in order to glimpse the properties of the ensuing interaction through measurement of the products. These scattering reactions prove to be experimentally challenging; for example, one or more of the interacting systems may have microscopic lifetimes, or a range of interaction energies must be explored. A cornerstone technique in the exploration of these novel systems has been the use of time-projection chambers (TPC), such as that used in ALICE [2], in which the trajectories of particles moving through the chamber can be reconstructed. A discussion of TPCs, and the physics which governs their use, is conducted in Section 5.

Associated with the use of TPCs is the task of identifying the tracks left by particles moving within the detector chamber, and subsequently determining their kinematic properties. These two processes of track *identification* and track *fitting* are introduced in Section 7. Existing techniques include sequential RANdom SAmple Consensus (RANSAC) and the sequential Hough transform, which suffer from greedy fitting. A reframing of the track-fitting problem as an uncapacitated facility location problem (UFLP) is given in Section 7.2. There, the Propose Expand and Re-estimate Labels (PEaRL) algorithm is introduced as an approach to solving the UFLP for geometric track fitting, with a discussion of its performance when applied to several simulated track observations.

A scattering experiment was performed by the Texas A&M group in 2018, before the work in this thesis began. The TexAT TPC was used in conjunction with a radioactive beamline to perform (elastic) scattering of ^{10}C upon ^{4}He, which would facilitate an investigation into the presumed molecular structures in ^{14}O. An overview of the existing literature upon ^{14}O and its mirror nucleus ^{14}C is given in Section 4. A description of the TexAT detector, and the various detector

and experimental properties that were recorded during the experiment is given in Section 6. The workflow for an analysis of these (and similar) data is described in Section 8, with a focus upon the analysis challenges imposed by the early data from the TexAT TPC. These challenges include the handling of ringing artefacts, which must be distinguished from true MicroMeGaS measurements, and fitting of the GET-shaped waveforms in the absence of a measured response function.

Further developments to this analysis were impeded by the onset of the COVID-19 pandemic, which detracted from the collaborative nature of research in every way imaginable. During this time, a need for teaching laboratory skills in the face of legal and ethical barriers to in-person demonstration was identified. This led to the development of the Online nuclear LABoratory (O-Lab), whose development, features, and usage statistics are outlined in Section 9.

Underpinning the work in this thesis is a considerable body of software that was developed in order to perform the analysis work described in Section 8, and develop the fitting method described in Section 7. These software will be made publicly available under an open-source license.

Part II

Experimental Nuclear Physics

CHAPTER

THREE

NUCLEAR MODELS

When describing the nucleus, it may appear that the most natural degrees of freedom are those of the constituent nucleons: their spins, isospins, and positions. From these parameters arises a set of microscopic models, which vary according to the assumptions and symmetries used to render an otherwise intractable computation feasible. These models are sensitive to the nucleon-nucleon interaction, for which there is no analytical theory. It is also possible to describe the nucleus in terms of the *collective* degrees of freedom, such as the centre-of-mass, or quadrupole moment. The challenge of aligning these two descriptions of the nucleus is a significant, on-going effort in the research community.

3.1 Macroscopic Models

3.1.1 Liquid Drop Model

First described in 1935, the Semi Empirical Mass Formula (SEMF) is an approximation for the ground-state binding energy of a nucleus according to its proton and neutron numbers [3]. Inspired by the liquid drop model, which considers the nucleus as a drop of incompressible fluid, it comprises of several theoretically-motivated terms (see Fig. 3.1):

$$E_B = \underbrace{a_V A}_{\text{volume}} - \underbrace{a_S A^{\frac{2}{3}}}_{\text{surface}} - \underbrace{a_C \frac{Z(Z-1)}{A^{\frac{1}{3}}}}_{\text{coulomb}} - \underbrace{a_A \frac{(N-Z)^2}{A}}_{\text{asymmetry}} + \underbrace{\delta(N,Z)}_{\text{pairing}} \ . \tag{3.1}$$

Volume energy

The attractive-repulsive potential of the strong force. The short effective range of the strong interaction means that it is effectively constrained to a constant number of neighbours.

Surface energy

A corrective term to account for nucleons at the surface of the nucleus, which have fewer neighbours than those at the centre of the drop.

Coulomb energy

Coulombic repulsion between positively charged protons reduces the binding energy.

Asymmetry energy

Given the Paul exclusion principle, nuclei of $N = Z$ should have higher binding energies than those for which $N \neq Z$.

Pairing energy

Identical nucleons with anti-aligned spins in the same spatial state have maximal wavefunction overlap, and correspondingly an enhanced binding energy.

It will be useful to discuss some of the features of this model in subsequent sections. For now, it suffices to say that the liquid-drop model (or semi-empirical mass formula) produces a reasonable prediction of the binding energy of a wide range of nuclei, particularly for heavier nuclei.

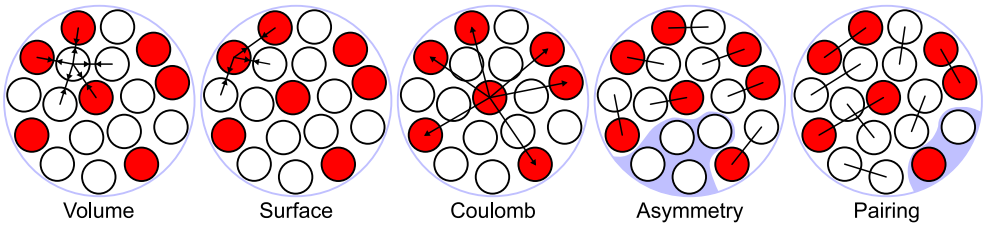


Fig. 3.1: Illustration of the physical interpretation of the terms that comprise the liquid-drop nature of the SEMF. Figure by Daniel FR, distributed under a CC BY-SA 3.0 license.

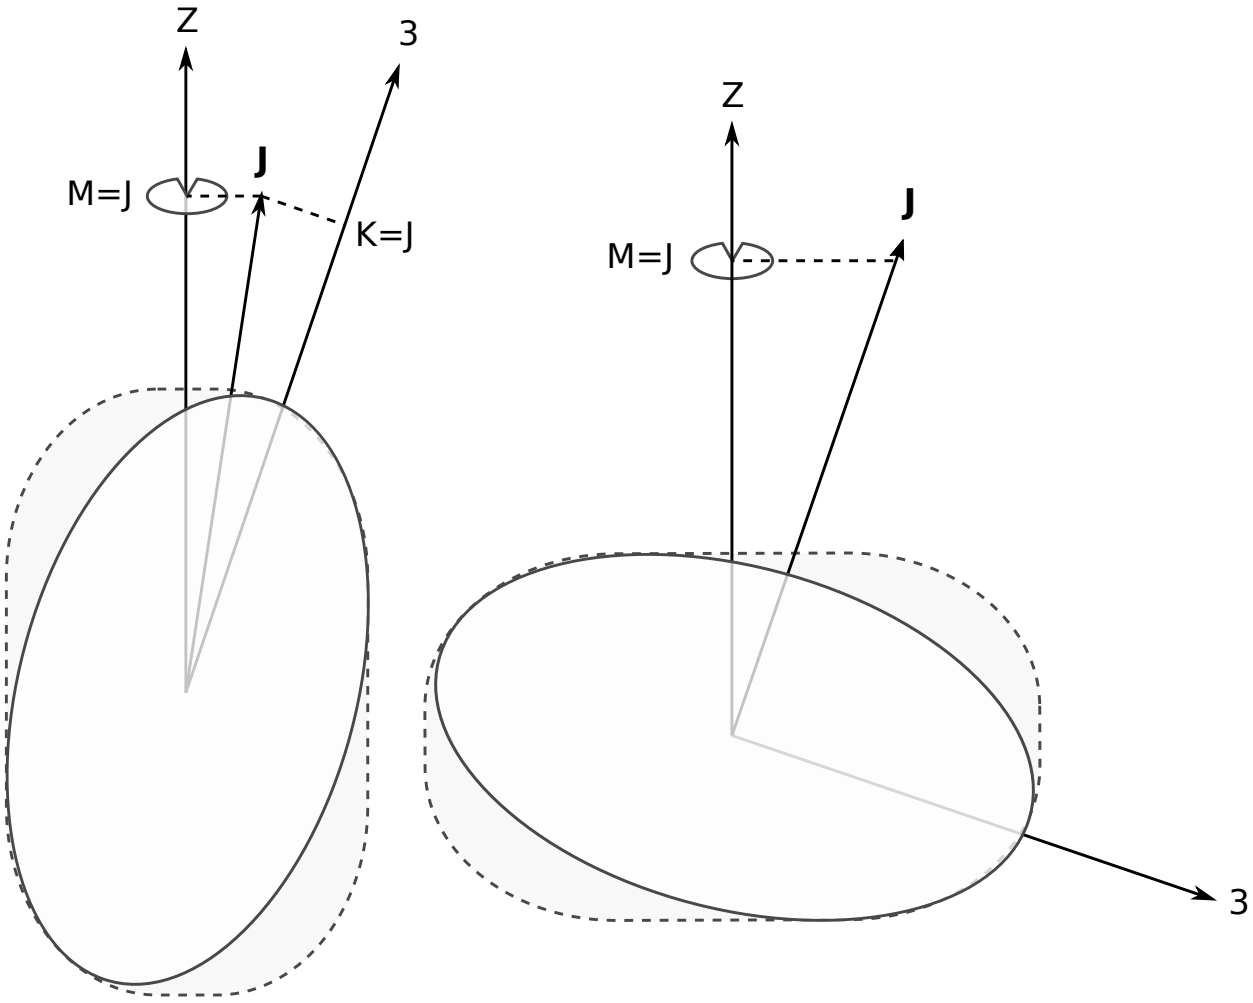
3.1.2 Rotor Model

Under the symmetry of the spherical shell model, eigenstates $|jm\rangle$ of the angular-momentum projection operator are degenerate with respect to the Hamiltonian; it is only the total momentum j that, under the spin-orbit interaction, yields distinct eigenvalues. Within a *deformed* system, however, it becomes possible to define an orientation, and the degeneracy of $|jm\rangle$ is lifted. This deformation may be invariant with respect to particular rotational subgroup(s), e.g. axially symmetric deformations, such that the orientation of the system is only partially defined and restrictions are imposed on the rotational degrees of freedom.

The lowest-order measure of stable deformation of a state is given by the (electric) quadrupole moment $Q = \langle JM = J|\hat{Q}_0|JM = J\rangle$ of the nucleus, which is measured such that the space-fixed $\hat{z}$ component is maximal. When measuring these values for the 2+ states even-even nuclei, a trend towards large, negative values is observed in the mid-shell region, whilst near to shell closures the quadrupole moment is significant diminished. Odd-mass near singly-closed nuclei demonstrate small negative Q values, corresponding to valence nucleons polarising a spherical core, whereas large

prolate deformations are seen in the mid-shell region [4].

When a charged object with a quadrupole moment undergoes rotation, it emits quadrupole (E2) radiation. It follows that the E2 transition rates serve as an indicator of deformation; when the reduced transition rate (B(E2)) values are expressed in terms of single-particle (Weisskopf) strengths (the typical strengths observed for single-nucleon transitions), it is clear that some of these transitions must involve *collective* excitations of many correlated nucleons. It transpires that both the large prolate and large oblate Q values observed in doubly-open mid-shell even-even and odd-mass nuclei can be explained through the *rotor model* (see Fig. 3.2).



(a) Prolate nucleus

(b) Oblate nucleus

Fig. 3.2: Schematic illustration of the measured quadrupole moment Q for an axially symmetric rotor. (a) shows the configuration in which the rotor symmetry axis 3 is maximally aligned with the space-fixed axis z. (b) depicts a state in which the 3 axis is approximately perpendicular to the z axis; in this case, the 3-projected momentum is zero. Under the rotor model, both the oblate doubly-even and prolate odd-mass nuclei measurements are explained by an intrinsic prolate deformation. Adapted from Ref. [4].

A quantum mechanical rotor has the Hamiltonian

$$\hat{H}_{\text{rot}} = \frac{\hbar^2}{2}\left[\frac{\hat{R}_1^2}{\mathcal{I}_1} + \frac{\hat{R}_2^2}{\mathcal{I}_2} + \frac{\hat{R}_3^2}{\mathcal{I}_3}\right], \tag{3.2}$$

where $\mathcal{I}_j$ is the moment of inertia with respect to the intrinsic axis j, and $\hbar$ the reduced Planck's constant. With the inclusion of intrinsic degrees of freedom, this Hamiltonian would also possess an intrinsic term that couples with the extrinsic components; if the nucleus has intrinsic angular momentum **J**, then the total angular momentum is given by the vector sum

$$\hat{\mathbf{I}} = \hat{\mathbf{R}} + \hat{\mathbf{J}}\,. \tag{3.3}$$

It can be shown that the Hamiltonian factors into intrinsic, rotational, and Coriolis terms

$$\underbrace{\hat{H}'_{\text{rot}}}_{\text{rotational}} + \underbrace{\hat{H}_{\text{intr}}}_{\text{intrinsic}} - \underbrace{\sum_i \frac{\hbar^2}{\mathcal{I}_i}\hat{I}_i\hat{J}_i}_{\text{Coriolis interaction}}\,, \tag{3.4}$$

where

$$\hat{H'}_{\text{rot}} = \sum_i \frac{\hbar^2}{2\mathcal{I}_i}\hat{I}_i^2 \tag{3.5}$$

(see [4]).

An *additional* coupling term arises if the moments of inertia depends upon the intrinsic variables. However, one can ignore both of these coupling interactions to define the *adiabatic approximation* which omits the coupling terms. For an *axially symmetric* rotor, $\hat{H}'_{\text{rot}}$ admits analytical solutions for the ground-state band, with energy eigenvalues

$$E'_{KI} = \frac{\hbar^2}{2}\left[\frac{I(I+1)}{\mathcal{I}} + \left(\frac{1}{\mathcal{I}_3} - \frac{1}{\mathcal{I}}\right)K^2\right], \tag{3.6}$$

and corresponding eigenvectors satisfying

$$\begin{aligned}\hat{I}^2\,|KIM\rangle &= I(I+1)\,|KIM\rangle\\ \hat{I}_z\,|KIM\rangle &= M\,|KIM\rangle\\ \hat{I}_3\,|KIM\rangle &= K\,|KIM\rangle\end{aligned} \tag{3.7}$$

where K is the eigenvalue of the $\hat{I}_3$ operator (see Fig. 3.3). In this model, M and K are good quantum numbers, and K defines a band of states for increasing I. For excited states, the K and $\mathcal{I}$ parameters may differ from the ground state,

generalising Eqn. 3.6 to

$$E_{\alpha KI} = \frac{\hbar^2}{2\mathcal{J}_\alpha} I(I+1) + E_{\alpha K}\,, \tag{3.8}$$

where α indicates the excited state.

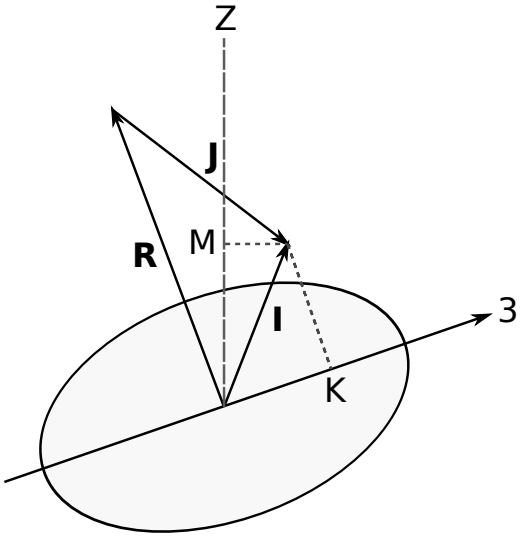


Fig. 3.3: Relationship between the total angular momentum **I**, the intrinsic angular momentum **J**, the rotor angular momentum **R**, and the components M of **I** along the fixed axis z; and K the symmetry axis. Adapted from Ref. [4].

Under the adiabatic limit, it is reasonable to consider the intrinsic structure of the nucleus to be independent of extrinsic rotational angular momentum [5]. We may partition the quantum state of the system into intrinsic (K) and extrinsic (I, M). Yet, the same cannot be said for the *symmetries* of each frame. For an axially symmetric rotor with reflection symmetry (in the plane whose normal is co-linear with the symmetry axis), it is not possible to distinguish rotations about the symmetry axis of the intrinsic frame. It follows that the total angular momentum component I_3 is given by the projection J_3 of the *intrinsic* angular momentum onto the symmetry axis, i.e. it is exclusively a function of the intrinsic state. This eliminates the degree of freedom associated with rotation around $\hat{\mathbf{3}}$.

Parity Doublets

Simultaneously, if the intrinsic Hamiltonian is invariant with respect to a rotation $\mathcal{R}_2(\pi)$ about an axis perpendicular to the symmetry axis, a further reduction in the rotational (extrinsic) degrees of freedom is possible. This intrinsic symmetry operator $\mathcal{R}_i = \mathcal{R}_2(\pi)$ has corresponding eigenvalues $r = \pm 1$ as $\mathcal{R}_i^2 = \mathcal{R}_2(2\pi) = +1$. The $K = 0$ state is an eigenfunction of this operator. Meanwhile, the same symmetry in the extrinsic frame $\mathcal{R}_e = \mathcal{R}_i$ yields eigenvalues of -1^I

for this $K = 0$ state, which follows from the behaviour of the spherical harmonics [6], such that we have

$$r = -1^I \tag{3.9}$$

Eqn. 3.9 admits two families of solutions; $I = 0, 2, 4, ...$ for the positive $r = +1$ eigenvalue of $\mathcal{R}_i$, and $I = 1, 3, 5, ...$ for the negative $r = -1$ counterpart. For reflection-symmetric axially-asymmetric structures, this gives rise to a phenomenon known as "parity inversion doublets", in which both solutions are seen with an associated energy splitting [7].

For the $K \neq 0$ states, it does not hold that they are eigenfunctions of $\mathcal{R}_i$. Upon expanding the intrinsic state in components of the total angular momentum, it can be shown that in order for $\mathcal{R}_i = \mathcal{R}_e$, a superposition of $K, \overline{K}$ solutions (where $\overline{K} = -K$) is required such that from two intrinsic states $K, \overline{K}$ a single rotational state is defined.

3.2 Microscopic Models

3.2.1 Many Body Systems

For a system of N indistinguishable particles, the canonical many-particle basis is given by the direct-product of the single-particle states:

$$\begin{aligned} |\lambda_1 \lambda_2 ... \lambda_N\rangle_{\text{C}} &= |\lambda_1\rangle \otimes |\lambda_2\rangle \otimes \cdots \otimes |\lambda_N\rangle \\ &= |\lambda_1\rangle\, |\lambda_2\rangle\, ...\, |\lambda_N\rangle \ . \end{aligned} \tag{3.10}$$

However, an arbitrary state in this basis is not guaranteed to satisfy the anti-symmetry requirements of a fermionic system such as the nucleus (under isospin symmetry). These anti-symmetric states are given by *permutations* of the canonical basis vectors (via the anti-symmetrisation operator $\hat{\mathcal{A}}$)

$$\begin{aligned} |\lambda_1 \lambda_2 ... \lambda_N\rangle_{\text{AS}} &= \hat{\mathcal{A}}\{|\lambda_1\rangle \otimes |\lambda_2\rangle \otimes \cdots \otimes |\lambda_N\rangle\} \\ &= \frac{1}{N!} \sum_{P \in S_N} (-1)^\pi \hat{P}(|\lambda_1\rangle \otimes |\lambda_2\rangle \otimes \cdots \otimes |\lambda_N\rangle)\,. \end{aligned} \tag{3.11}$$

It is usually more convenient to explore nuclear theory in terms of the language of second quantisation, which intrinsically embodies the anti-symmetrisation requirement through the creation and annihilation operators [8]. Subsequently, states are most naturally described using occupation number representation

$$|n_1 n_2 ...\rangle_{\text{OR}}\ , \tag{3.12}$$

where n_1 indicates the number of particles in the single-particle state 1.

The Hamiltonian of a quantum system is the operator $\hat{H}$, conventionally given as the sum of kinetic and potential energy operators

$$\hat{H} = \hat{T} + \hat{V}\,, \tag{3.13}$$

whose eigenvalues correspond to the set of measurable energies. Considering up-to two-body interactions, we can write Eqn. 3.13 in second quantization form as

$$\hat{H} = \sum_{\alpha\beta} T_{\alpha\beta} \underbrace{\hat{a}^\dagger_\alpha \hat{a}_\beta}_{\text{one-body}} + \frac{1}{2}\sum_{\alpha\beta\gamma\delta} V_{\alpha\beta\gamma\delta} \underbrace{\hat{a}^\dagger_\alpha \hat{a}^\dagger_\beta \hat{a}_\delta \hat{a}_\gamma}_{\text{two-body}}\,, \tag{3.14}$$

where $\hat{T}$ is the kinetic energy operator and $\hat{V}$ is the two-body inter-nucleon interaction.

3.2.2 Mean Field Theory

Eqn. 3.14 could be recast in terms of an artificial single-particle potential $\hat{U}$

$$\begin{aligned}\hat{H} &= \underbrace{\sum_{\alpha\beta} (T_{\alpha\beta} + U_{\alpha\beta})\hat{a}^\dagger_\alpha \hat{a}_\beta}_{\text{single particle Hamiltonian}} + \underbrace{\frac{1}{2}\sum_{\alpha\beta\gamma\delta} V_{\alpha\beta\gamma\delta}\hat{a}^\dagger_\alpha \hat{a}^\dagger_\beta \hat{a}_\delta \hat{a}_\gamma - \sum_{\alpha\beta} U_{\alpha\beta}\hat{a}^\dagger_\alpha \hat{a}_\beta}_{\text{residual interaction}}\,, \\ &= \hat{H}^{[1]} + \hat{V}'\,,\end{aligned} \tag{3.15}$$

where $\hat{H}^{[1]}$ is the single-particle Hamiltonian, and $\hat{V}'$ is a *residual* interaction between nucleons. In the event that this residual interaction is small, it can be treated as a perturbation of the single-particle Hamiltonian. Under this simplification, the problem of solving the many-body Schrodinger equation is reduced to that of solving the single-particle Schrödinger equation for an external potential

$$\hat{H}^{[1]} \left|\phi_k\right\rangle = E_k \left|\phi_k\right\rangle\,. \tag{3.16}$$

Hence, the introduction of a particular single-particle potential transforms the problem from a strongly interacting system of nucleons to a weakly interacting system of particles [11]. The residual interaction $\hat{V}'$, whose matrix elements are much reduced in amplitude with respect to the bare two-nucleon potential, is responsible for configuration mixing between single-particle states, i.e. correlating otherwise independent nucleons [4].

In the context of the nucleus, the existence of a single-particle potential corresponds to a *mean-field*, whose origin is the motion of the nucleons themselves. The physical justification for introducing a mean-field potential $\hat{U}$ depends upon

whether one-body behaviour or two-body collisions dominate within the nucleus, i.e. whether the nucleon mean-free-path is much greater than, or lesser than the nuclear size [12]. Despite the strong short-range repulsive nature of the strong nuclear force, through Pauli blocking of states close to the Fermi surface it can be shown that nucleons behave to first approximation as independent particles i.e. comprise an Independent Particle Model (IPM) [13].

In occupation representation, a general state $|\phi\rangle$ is written as

$$|\phi\rangle = \sum_{i_1 i_2 \dots i_N} c_{i_1 i_2 \dots i_N} \hat{a}^\dagger_{i_1} \hat{a}^\dagger_{i_2} \dots \hat{a}^\dagger_{i_N} |0\rangle \ , \tag{3.17}$$

where $|0\rangle$ is the vacuum state, and the number of combinations of $i_1 i_2 \dots i_N$ is given by $\binom{N_c}{N}$. Clearly, despite the introduction of a single-particle approximation (imposed by the supposition of a mean-field), Eqn. 3.16 remains an intractable problem unless a simplification is used to reduce the size of the model space. There are two predominant mechanisms for proceeding:

- Restrict the number of active particles and states (constrain N_c and N)
- Choose a restricted many-particle form (constrain $|\phi\rangle$) [4].

The former approach underpins the shell model, whilst the latter describes the high-level procedure of the Hartree-Fock Approximation.

3.2.3 Hartree Fock

The question that follows Eqn. 3.15 is that of how to determine $\hat{U}$ such that $\hat{V}'$ is minimised. A natural candidate would be one of the phenomenological potentials, such as the Woods-Saxon model, or an analytical potential such as the Harmonic Oscillator. However, these are ultimately simplifications of the nuclear potential. A proper treatment of the mean field must produce a *self-consistent* solution in which the field generated by the motion of the nucleons agrees with that used to derive the wavefunction of the system [10]. One such treatment is the Hartree-Fock (HF) approximation method.

Though a detailed discussion of HF will not be performed here, it suffices to establish the assumptions and conclusions of the method. In order to restrict the model space of the many-particle problem to a tractable subset, the HF method assumes that the ground state of the system is described by a single Slater-determinant of single-particle eigenfunctions. Therefore, instead of Eqn. 3.17, we have

$$|\mathrm{HF}\rangle = \hat{a}^\dagger_{i_1} \hat{a}^\dagger_{i_2} \dots \hat{a}^\dagger_{i_N} |0\rangle \ . \tag{3.18}$$

This assumption is most appropriate for doubly-closed nuclei, where there is a large energy gap between the first unfilled single-particle level and the ground state $|\mathrm{HF}\rangle$ such that the gound state is reasonably stable [14] [5].

Variational Principle

The mean field is ultimately derived through the variational principle, which seeks to determine $|\phi\rangle$ such that $\hat{H}^{[1]}$ is minimised with respect to small admixtures of Slater determinants:

$$\delta\Big(\langle\phi|\hat{H}|\phi\rangle - E\langle\phi|\phi\rangle\Big) = \langle\delta\psi|\hat{H}|\psi\rangle - E\langle\delta\psi|\psi\rangle = 0\,, \tag{3.19}$$

where $|\delta\psi\rangle = \varepsilon\hat{a}_k^\dagger\hat{a}_j|\psi\rangle$ for some arbitrary ε, and $k > N, j \leq N$ [4] [15]. If the N-lowest energy single-particle eigenfunctions are taken as the HF ground-state, then these admixtures correspond to single "particle-hole" (1p1h) excitations of the closed "core" that lies below the sharp Fermi surface (see Fig. 3.4).

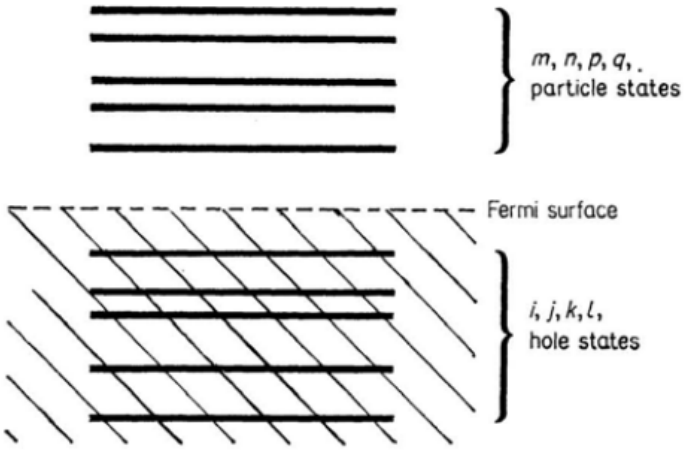


Fig. 3.4: Labelling of particle and hole states relative to a sharp Fermi surface, that follows from the restricted Slater determinant form of the HF ground state $|\text{HF}\rangle$. Figure taken from Ref. [5].

The process of determining the eigenstates of $\hat{H}^{[1]}$ is given by the following recurrence relation [5]:

$$\Big\{\Big|\text{HF}_{(0)}\Big\rangle\Big\} \to \hat{U}^{(0)} \to \Big\{\Big|\text{HF}_{(1)}\Big\rangle\Big\} \to \hat{U}^{(1)} \to ... \tag{3.20}$$

It can ultimately be shown that the HF Hamiltonian is

$$\hat{H} = \sum_{\alpha\beta}\left(T_{\alpha\beta} + \sum_\gamma \overline{V}_{\alpha\gamma\beta\gamma}\right)\hat{a}_\alpha^\dagger\hat{a}_\beta\,, \tag{3.21}$$

where $\overline{\overline{V}}$ is the anti-symmetrised interaction potential. This implies that the mean field $\hat{U}$ from Eqn. 3.15 is

$$\hat{U} = \sum_{\alpha\gamma\beta}\overline{V}_{\alpha\gamma\beta\gamma}\hat{a}_\alpha^\dagger\hat{a}_\beta\,, \tag{3.22}$$

i.e. the averaged two-body contributions from the $N-1$ remaining nucleons [4] [15]. Primitive excited states can then

be constructed from one-particle/one-hole (1p1h) and two-particle/two-hole (2p2h) excitations of this ground state [5].

Whether the minimum found by the HF procedure is a true global minimum depends upon the initial configuration. If the initial configuration determined by the trial wave-functions lies close to a shape-isomeric state, it is quite possible for the numerical procedure to settle upon a configuration that is deformed from the ground state [15]. The variational principle outlined in Eqn. 3.19 can be extended to include a deformation constraint that restricts e.g. the quadrupole moment Q of the nucleus (see Fig. 3.5).

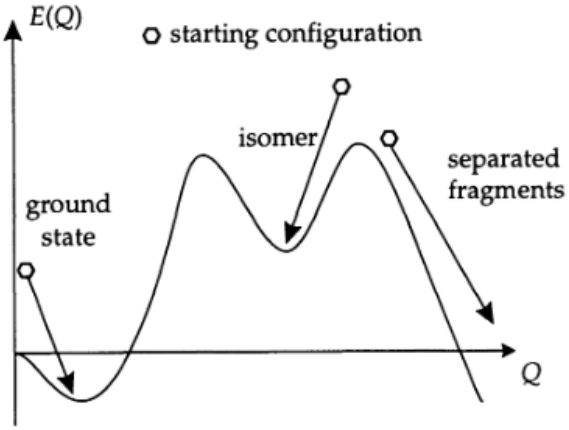


Fig. 3.5: Schematic diagram of the convergence of different HF trial wave-functions to local minima with distinct quadrupole moments Q corresponding to a deformed isomer and subsequent fissioning of the nucleus. Figure taken from Ref. [15].

Residual Interaction

Whilst the HF approximation generates a set of self-consistent single-particle states, it does so by explicitly ignoring the residual interaction. The variational principle requires that the single-particle Hamiltonian matrix element satisfies $\langle \delta\psi|\hat{H}|\psi\rangle = 0$, where $|\delta\psi\rangle = \sum_{mi} \epsilon_{mi}\hat{a}_m^\dagger \hat{a}_i \,|\psi\rangle$. This is equivalent to requiring that the Hamiltonian does not couple one-particle-hole excitations to the ground state (i.e., no residual interaction). A significant component of this interaction is the short-range pairing correlation, discussed in Section 3.1.1, that couples like-pairs of nucleons to states of zero angular momentum [4]. This is most evident in nuclei with pairs of protons or neutrons outside of doubly-closed shells [4].[1]

Extensions to the Hartree-Fock method have been developed to consider the pairing interaction:

- the Bardeen-Cooper-Schrieffer (BCS) model introduces a generalised pair creation operator for even-even nuclei of the form $A^+ = \sum_{k>0} \frac{v_k}{u_k} a_k^+ a_{\bar{k}}^+$ to build the BCS ansatz: $|\text{BCS}\rangle \propto \exp\left(A^+\right)|-\rangle = \sum_{r=0}^{\infty} \frac{1}{n!}\left(A^+\right)^r |-\rangle$, where

[1] Experimentally, it is found that the ground state of spherical nuclei with partially filled j-shells is not degenerate with respect to angular momentum as would be expected for a purely single-particle model. It follows that the residual interaction lowers the zero angular momentum state with respect to the other configurations [15].

$|-\rangle$ is the vacuum state.[16] The resulting Hamiltonian is constrained to conserve particle number N. Unlike the base HF method, the BCS model is not self-consistent; the pairing field does not change the single-particle wavefunctions [10], which are generated using traditional Hartree-Fock [17].

- the Hartree-Fock-Bogoliubov (HFB) model generalises HF theory to incorporate the quasi-particle concepts of the BCS model. Importantly, it replaces the single particle-pair operators with generalised Bogoliubov quasi-particle operators of the form $\sum_\mu \hat{a}^\dagger_\mu u_{\mu\nu} + \hat{a}_\mu v_{\mu\nu}$ [4].

Excited States

Meanwhile, *excited* states are primarily modelled in the HF approximation as 1p1h excitations. Experimentally, however, it is clear that for many nuclei these single 1p1h excitations are themselves insufficient to describe low-lying excited states. An example is ^{16}O, for which the lowest energy excitations predicted by the harmonic oscillator approximation should be seen at around 11.5 MeV (the difference between the 1p and 2s-1d shells). The spectrum of ^{16}O, however, has a $J^\pi = 3^-$ state at 6.13 MeV, indicating additional *coherent* excitation mechanisms [14]. The observation of these excitation modes directly underpins the use of collective models to describe the nucleus. These collective models usually ignore the single-particle degrees of freedom in favour of collective vibrational or rotational modes.

In addition to collective models, there are microscopic treatments of collective excitations. The most fundamental of these is the Tamm-Dancoff Approximation (TDA), which treats excited states of the nucleus as a superposition of 1p1h excitations of the ground state. These excitations are introduced by a creation operator [15]:

$$\hat{Q}^\dagger_\nu = \sum_{mi} c^\nu_{mi} \hat{a}^\dagger_m \hat{a}_i \,, \tag{3.23}$$

and a correlational potential that acts in the excited particle-hole space. It can be shown that such an ansatz admits collective behaviour through a *coherent* excitation of single particles [15]. Typically these solutions are obtained in a restricted model space comprising of several levels above and below the Fermi level [14].

A limitation of the TDA, besides treating only 1p1h excitations, is the use of the unmodified HF ground state; the TDA assumption is that only excited states may have particle-hole excitations above the ground state [13]. The Random-Phase Approximation (RPA) extends the particle-hole creation operator to

$$\hat{Q}^\dagger_\nu = \sum_{mi} x^\nu_{mi} \hat{a}^\dagger_m \hat{a}_i + y^\nu_{mi} \hat{a}^\dagger_i \hat{a}_m \,, \tag{3.24}$$

where $\{y_{mi}\}$ account for the account for the correlations in the *ground* state.

3.2.4 Shell Model

Having established the fundamental principles of microscopic models with Hartree-Fock theory, the nature of the shell model becomes fairly trivial to describe. Yet Hartree-Fock theory, though formally solvable, is computationally expensive and requires that the mean field be computed for each nucleus of interest. A simplified approach, the Shell Model, is found by determining the self-consistent field *empirically*, rather than from the two-body nucleon-nucleon interaction. The shell model potential should reconstruct the well-established shell-structure (magic numbers) of nuclear structure. The simplest model of the nucleus, that of the harmonic oscillator, does not on its own reconstruct these magic numbers. Instead, a *central potential* that lies between the HO potential and a square well, is combined with a spin-orbit interaction to reproduce the experimental values (see Fig. 3.6 and Eqn. 3.25). The Wood-Saxon potential is one such central-potential. The l^2 contribution shown in Fig. 3.6 accounts for the screening effect of the "finite" range of the strong nuclear force, and flattens the radial shape of the potential. The spin-orbit term, meanwhile, is primarily a surface term which equates to a dependence of the nuclear potential upon the total-spin of the nucleon [18]. The properties of the single-particle states upon which the shell model is based are determined experimentally through pick-up and stripping reactions, which determine the binding energy and angular momentum of particles near the Fermi surface [15].

$$u(\mathbf{r}) = V(\mathbf{r}) + f(\mathbf{r})\mathbf{l}\cdot\mathbf{s}\,. \tag{3.25}$$

For closed-shell nuclei, the shell model is simply an approximation to Hartree-Fock theory. An IPM Hamiltonian is chosen such that the low-energy eigenstates have the observed properties of the associated physical states, such as the nuclear density distributions, or angular momenta. The zeroeth order approximation yields the IPM itself, which describes closed-shell nuclei [5]. Thereafter, an improved approximation for open-shell nuclei is found by diagonalising a total Hamiltonian (such as Eqn. 3.15) in the restricted space spanned by the lowest-energy solutions of the IPM. These solutions correspond to placing valence nucleons in the unoccupied levels of the closed-shell core [4]. Further *second approximation*s include the independent-pair approximation, which accounts for the pairing contributions to the residual interaction.

In general, the shell model simplifies the many-body problem into a few-body problem by eliminating the core; considering only the exclusion principle and the potential well that the nucleons within the core generate. This approach is limited by the size of the model space required to perform a shell model calculation. In general, successful shell model calculations must select an appropriate subspace that is tailored to modeling the dynamics of interest to the calculation [4]. Rowe and Wood [4] describe the most significant role of the shell model as providing the

> language and concepts for the interpretation of observed phenomena, to gain understanding of the circumstances in which particular phenomena arise, and to extrapolate into other mass regions to see if the inter-

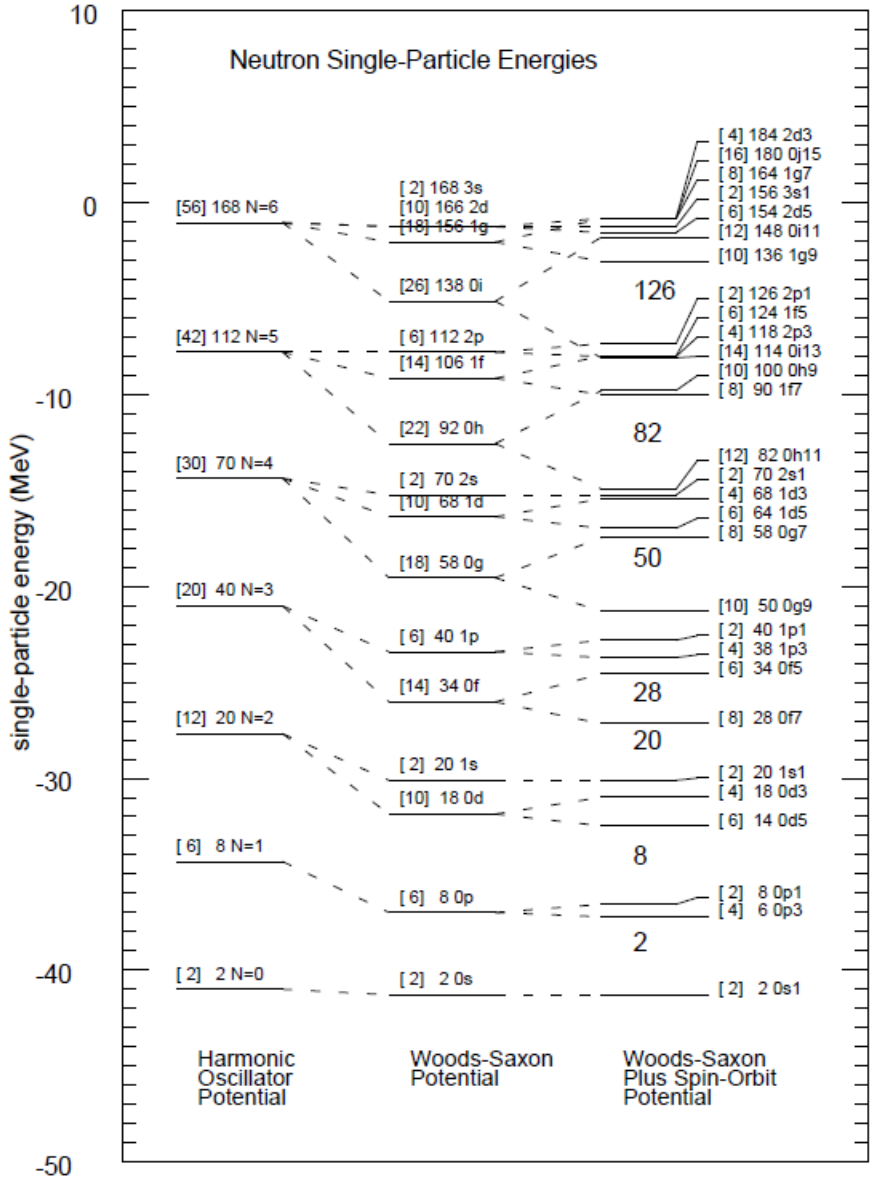


Fig. 3.6: Single-particle energies for (a) a simple harmonic oscillator (SHO); (b) a SHO with an l^2 term; (c) a shell-model potential with l^2 and $\mathbf{l} \cdot \mathbf{s}$ terms. Figure by Morten Hjorth-Jensen, distributed under a CC BY-NC 4.0 license.

pretations are consistent.

Deformed Shell Model

As discussed in Section 3.1.2, nuclear states away from closed shells show evidence of strong quadrupole deformation through their electric quadrupole moments. Given the success of the shell model in describing (near) closed-shell nuclei, it is reasonable to ask whether it might be extended to this deformed regime. One such model is the Nilsson model, considered one of the most successful nuclear models ever developed [4]. As with the rotor model, the adiabatic approximation is used to decouple the (deformed) core from the motion of the core; this model diminishes in effectiveness at high angular momenta.

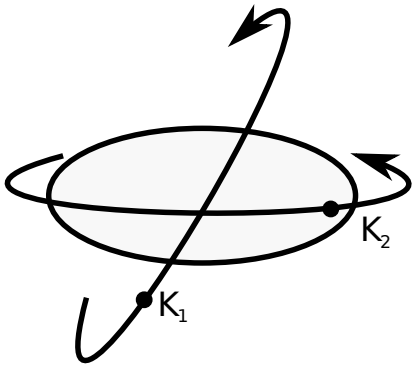


Fig. 3.7: Illustration of two nucleon orbits around a deformed core. K_1 depicts an orbit that lies closer to the bulk of the nucleus than K_2, and therefore lies at lower energy. Adapted from Ref. [18].

The distinguishing feature of the Nilsson model with respect to the spherical shell model is an orientation of nucleon states with respect to the core. This is illustrated in Fig. 3.7, which depicts two nucleon orbits K_1 and K_2. As illustrated for the rotor model in Fig. 3.3, the value K corresponds to the projection of the the intrinsic angular momentum (in this case, the single-particle momentum) onto the symmetry axis of the core.

Given the success of the SHO in the description of the single-particle basis of the spherical shell model, the spheroidal anisotropic (deformed) harmonic oscillator is well-motivated to establish these states in the deformed system. Unlike the *spherical* oscillator, in which there is a singular frequency ω that determines the per-quanta energy in each axis, the anisotropic harmonic oscillator distinguishes the frequency by axis. In the event of an axially-symmetric rotor, there are two distinct frequencies. Given that solutions to deformed oscillator should map onto the spherical oscillator in the deformation-free limit, the anisotropic frequencies must preserve the nuclear volume: $\omega_3\omega_\perp^2 = \omega_0^3$, where ω_0 is the spherical oscillator frequency, ω_3 is the axial oscillator frequency, and $\omega_\perp$ is the perpendicular oscillator frequency. This

yields the following relations

$$\begin{aligned} \omega_{\perp} &= \omega_0 \exp\left(\frac{\epsilon}{3}\right) \\ \omega_3 &= \omega_0 \exp\left(\frac{-2\epsilon}{3}\right). \end{aligned} \tag{3.26}$$

where $\hat{3}$ is the symmetry axis, and ϵ is the deformation parameter, with $\epsilon > 0$ for prolate deformation [4]. Under this deformed potential, the total angular momentum j is no longer a good quantum number; only the symmetry-axis projection remains invariant whilst the orthogonal components are partitioned between the nucleons and the core [4]. It follows that states of a given projection are two-fold degenerate as the Hamiltonian depends only upon the magnitude of the $\hat{3}$ projection.

The principle behind the shell model is that eigenstates of the single-particle Hamiltonian are well-separated into distinct bands. In the spherical harmonic oscillator, these are directly influenced by increments in the oscillator quanta as well as the spin-orbit interaction splitting. The deformed harmonic oscillator also produces shell-structure, with maximal degeneracy (corresponding to strongly-defined shell structure) appearing at integer ratios of the two oscillator frequencies [6]. Through an application of group theory, it can be shown that the shell closures at integer deformation ratios correspond to the decomposition of the deformed oscillator into multiple spherical potentials [19]. These findings have implications for understanding alpha-conjugate nuclei, nuclei which may be decomposed into a set of alpha particles, and the phenomenon of alpha clustering.

No-core Shell Model

The no-core shell model (NCSM) description of the nucleus is an *ab-initio* approach to determine the behaviour of the nucleus from two-body and three-body realistic interaction potentials. No further detail beyond the high-level principles of the model will be described here, but an overview is given to contextualize the landscape of microscopic models. The term *realistic* follows from the choice of nucleon-nucleon potential that fit experimental phase shifts up to a particular energy, and nucleon-nucleon-nucleon interactions that include terms corresponding to two-pion exchanges [20].

In the domain of nuclear physics, “ab-initio” assumes a large number of distinct meanings. In the context of the NCSM, the term is a description of a system in which all A nucleons are active members of the model space, and their collective behaviour is determined from the introduction of effective interactions. Clearly, a model that operates at the level of protons and neutrons is not concerned with the subnucleonic degrees of freedom that the term “ab-initio” might invoke.

The NCSM is somewhat paradoxically named; as aforementioned, the shell model is usually built upon a mean-field potential, which is notably absent in the NCSM. Rather, the harmonic-oscillator basis, truncated to a maximum total system energy, is used to define the single-particle states through *natural orbitals* [9]. The resulting system, using second

quantisation representation, is then amenable to treatment with techniques developed for shell model calculations, which justifies the *shell model* description.

3.2.5 Anti-symmetrised Molecular Dynamics

Anti-symmetrised Molecular Dynamics (AMD) attempts to model the coexistence of cluster and mean-fields phenomena within the nucleus. Like the HF methods, AMD employs a basis of Slater determinants. Instead of single-particle eigenfunctions, however, AMD uses a basis of Gaussian wave packets:

$$|\phi_{\text{AMD}}(Z)\rangle = \hat{\mathcal{A}}\{|\psi_1\rangle \otimes |\psi_1\rangle \otimes \cdots \otimes |\psi_N\rangle\}\,, \tag{3.27}$$

where $|\psi_i\rangle$ is the direct product state of the spatial (ϕ), intrinsic spin (χ), and isospin (τ) states of nucleus i:

$$\begin{aligned} |\psi_i\rangle &= \left|\phi_{\mathbf{Z}_i}\right\rangle \otimes |\chi_i\rangle \otimes |\tau_i\rangle \\ \left\langle\phi_{\mathbf{Z}_i}\middle|\phi_{\mathbf{Z}_i}\right\rangle &\propto \exp\left(-\nu\left(\mathbf{r}_j - \frac{\mathbf{Z}_i}{\sqrt{\nu}}\right)^2\right) \\ |\chi_i\rangle &= \left(\frac{1}{2} + \xi_i\right)|\chi_\uparrow\rangle + \left(\frac{1}{2} - \xi_i\right)|\chi_\downarrow\rangle\ . \end{aligned} \tag{3.28}$$

It follows that an AMD many-body basis state can be represented by a set of variational parameters, $\mathcal{Z} = \{\nu, \mathbf{Z}_1, \mathbf{Z}_2, \ldots, \mathbf{Z}_N, \xi_1, \xi_2, \ldots \xi_N\}$, where the width parameter ν is optimised for a given system, and $\mathbf{z}_i = \begin{pmatrix}\mathbf{x} & \mathbf{y} & \mathbf{z} & \sigma\end{pmatrix}$ describes the centre, momenta, and width of each localised Gaussian [21]. In the limit that the Gaussian centres converge upon a particular location, then the AMD wave-function becomes equivalent to the harmonic oscillator shell-model wave-function (see Fig. 3.8) [22].

Energy minimisation

To determine the optimum values for $\mathcal{Z}$, the time-dependent variational principle (using the principle of least action) [24]

$$\delta \int_{t_1}^{t_2} \langle\psi(t)|\hat{H} - \mathrm{i}\hbar\frac{\partial}{\partial t}|\psi\rangle = 0 \tag{3.29}$$

is used to derive the equations of motion of the system [23]:

$$i\hbar\frac{\mathrm{d}Z_k}{\mathrm{d}t} = (\lambda + i\mu)\frac{\partial}{\partial Z_k^*}\frac{\langle\phi^\pm|\hat{H}|\phi^\pm\rangle}{\langle\phi^\pm|\phi^\pm\rangle}\,, \tag{3.30}$$

where the term $(\lambda + i\mu)$ is an *artificial* viscous component designed to dissipate the energy of the system towards a minimum [25] [26].

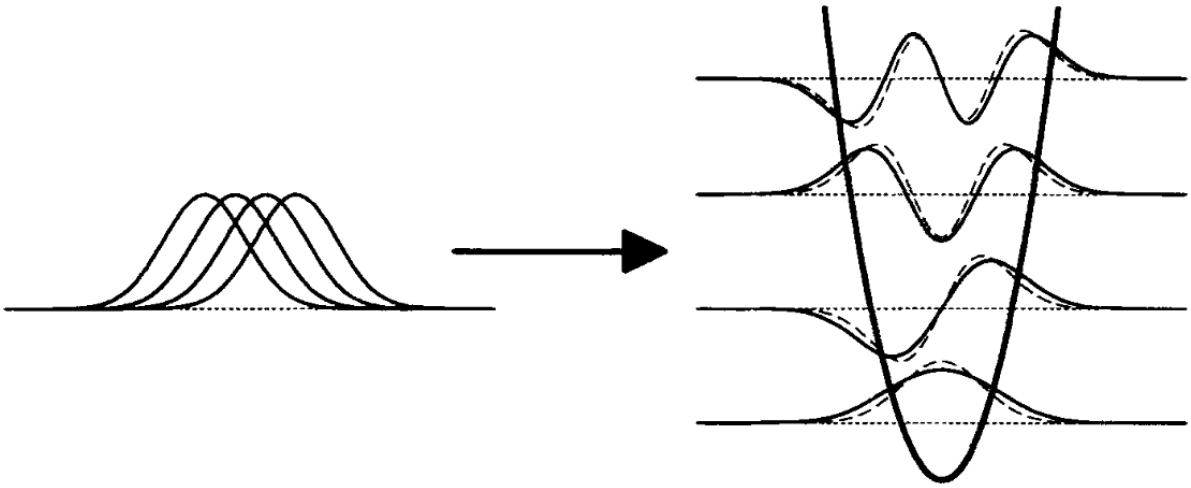


Fig. 3.8: Schematic representation of the eigenstates resulting from anti-symmetrisation of Gaussian wave-packets in the limit of $\frac{d}{\sqrt{a}} \to 0$, where $d = 0.75\sqrt{a}$ is the displacement between each wave-packet, and a the real width parameter. Figure taken from Ref. [23].

In the language of second quantisation, the AMD Hamiltonian

$$\hat{H} = \sum_{\alpha\beta} T_{\alpha\beta}\hat{a}^{\dagger}_{\alpha}\hat{a}_{\beta} + \frac{1}{2}\sum_{\alpha\beta\gamma\delta}\left(V^{\mathcal{N}}_{\alpha\beta\gamma\delta} + V^{\mathcal{C}}_{\alpha\beta\gamma\delta}\right)\hat{a}^{\dagger}_{\alpha}\hat{a}^{\dagger}_{\beta}\hat{a}_{\delta}\hat{a}_{\gamma} - \hat{T}_{\text{CoM}} \tag{3.31}$$

is identical to Eqn. 3.14 except for an additional two-body Coulomb interaction (usually ignored in the HFB procedure) [27], and the removal of the centre-of-mass energy via the operator $\hat{T}_{\text{CoM}}$. A range of effective nuclear potentials have been used for $\hat{V}_{\text{nuc}}$ including the Volkov, Gogny, or Skyme forces; whilst the Coulomb force is often approximated by a sum of (seven) Gaussians [22]. The expectation of this Hamiltonian is then minimised with respect to the wave-function parametrisation $\mathcal{Z}$

$$\min_{\mathcal{Z}} \frac{\langle\phi_{\text{AMD}}|\widehat{H} - \widehat{T}_{\text{cm}}|\phi_{\text{AMD}}\rangle}{\langle\phi_{\text{AMD}}|\phi_{\text{AMD}}\rangle} \tag{3.32}$$

to determine the lowest energy state $|\phi_{\text{AMD}}\rangle$.

Variation after Projection, Projection after Variation

States obtained by the variation procedure generally no longer demonstrate the symmetries of the Hamiltonian Eqn. 3.31 with respect to parity and rotational invariance [28]. These symmetries can be restored through projection onto angular momentum and parity eigenstates:

$$\begin{aligned} \left|\phi^{\pm}_{\text{AMD}}\right\rangle &\equiv P^{\pm}\left|\phi_{\text{AMD}}(\mathcal{Z})\right\rangle = \frac{1 \pm \hat{P}_r}{2}\left|\phi_{\text{AMD}}(\mathcal{Z})\right\rangle \\ \left|\phi^{J}_{MK}\right\rangle &\equiv P^{J}_{MK}\left|\phi_{\text{AMD}}(\mathcal{Z})\right\rangle = \int d\Omega D^{J*}_{MK}(\Omega)\hat{R}(\Omega)\left|\phi_{\text{AMD}}(\mathcal{Z})\right\rangle \ , \end{aligned} \tag{3.33}$$

where $D^{J}_{MK}(\Omega)$ is the Wigner's D function, and $\hat{R}(\Omega)$ is a rotation operator for some Euler angle Ω [22] [29].

Early AMD calculations were performed first by a method known as variation *before* projection (VBP, or properly VaP$^{\pi}$bP^{J}) [26]. By deferring angular momentum projection until after variation, VBP is numerically the least inexpensive means of obtaining states of good angular momentum. Nevertheless, it is possible to minimise the energy of the *projected state*, i.e.

$$\min_{\mathcal{Z}} \frac{\langle\phi_{\text{AMD}}|{P^{J\pm}_{MK}}^{\dagger}\left(\widehat{H} - \widehat{T}_{\text{cm}}\right)P^{J\pm}_{MK}|\phi_{\text{AMD}}\rangle}{\langle\phi_{\text{AMD}}|{P^{J\pm}_{MK}}^{\dagger}P^{J\pm}_{MK}|\phi_{\text{AMD}}\rangle} \tag{3.34}$$

for some state $\left|\phi^{J}_{MK}\right\rangle$, a method known as variation *after* projection (VAP, or properly VaP$^{\pi}$aP^{J}) [30]. Fundamentally, VAP approach necessitates a minimisation procedure for each eigenstate onto which the AMD state is projected. Although much more costly than VBP, the VAP procedure is necessary to describe states with intrinsic structures that differ from the mean-field state (see Fig. 3.9) [31].

The single Slater-determinant AMD basis representation ensures that the solution to the many-body problem is restricted to a subset of the anti-symmetric Hilbert space, i.e. that given by the mean-field approximation of independent nucleons in single-particle eigenstates moving in an external potential. Complex correlational phenomena, corresponding to configuration mixing, requires a linear superposition of Slater determinants. The admixtures of eigenstates yielded by the projection operators leads to the consideration of beyond mean-field phenomena during the variational procedure.

Excited States

In the HF approximation, a set of low-lying excited states can be constructed from superpositions of 1p1h excitations of the ground state $|HF\rangle$. This approximation is built upon the assumption that these excited states do not significantly deform the mean field, and that their correlations can be ignored [15]. In the extreme case of the structures that emerge in excited states near the cluster-decay threshold, these assumptions are unfounded and the single-particle states must be recalculated in a self-consistent manner.

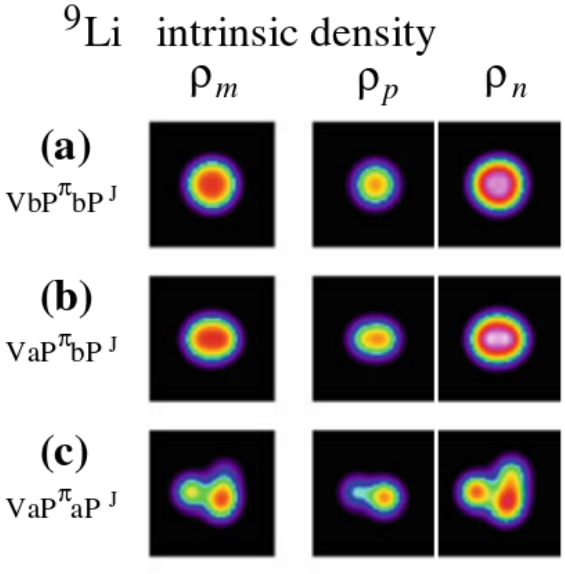


Fig. 3.9: Density distributions of the AMD wavefunctions for ^{9}Li obtained by (a) variation *before* parity and angular momentum projection, (b) *after* parity and *before* angular momentum projection, and (c) *after* parity and angular momentum projection. It can be seen that only (c) exhibits beyond mean-field structure. The proton (ρ_p), neutron (ρ_n), and total (ρ_m) densities of the AMD wave-function are illustrated. Figure adapted from Ref. [21].

Unlike HF, the AMD method does not readily admit a set of basis states from which to compose excited states. There are two favoured methods for determining such a basis; orthogonal VAP, and constrained variation. In orthogonal VAP, a set of basis functions is established through orthogonalisation to already-obtained lower states $\phi_k^{J\pm}(\{\mathcal{Z}_k\})$ for $k = 1, \ldots, n-1$, thus building a set of intrinsic AMD states:

$$\phi_n^{J\pm}(\mathcal{Z}) = P_{MK}^{J\pm}\phi_{\text{AMD}}(\mathcal{Z}) - \sum_{k=1}^{n-1} \frac{\left\langle \phi_k^{J\pm}\left(\mathcal{Z}_k^{J\pm}\right) \middle| P_{MK}^{J\pm}\phi_{\text{AMD}}(\mathcal{Z}) \right\rangle}{\left\langle \phi_k^{J\pm}\left(\mathcal{Z}_k^{J\pm}\right) \middle| \phi_k^{J\pm}\left(\mathcal{Z}_k^{J\pm}\right) \right\rangle} \phi_k^{J\pm}\left(\mathcal{Z}_k^{J\pm}\right) . \tag{3.35}$$

The eigenstates of the AMD Hamiltonian are constructed from these *intrinsic* states $\{|\phi_{\text{AMD}}(\mathcal{Z}_i)\rangle\}$ by simultaneously diagonalising the Hamiltonian $\left\langle P_{MK'}^{J\pm}\phi_{\text{AMD}}^i \middle| H \middle| P_{MK''}^{J\pm}\phi_{\text{AMD}}^j \right\rangle$ and normal $\left\langle P_{MK'}^{J\pm}\phi_{\text{AMD}}^i \middle| P_{MK''}^{J\pm}\phi_{\text{AMD}}^j \right\rangle$ matrices formed from these states [30]. It follows that these eigenstates are *superpositions* of intrinsic states.

Meanwhile, constrained variation closely resembles the Generator Coordinate Method (GCM) used to include collective phenomena in many-body wave-functions [32]. First, a set of basis functions is established through the imposition of a deformation constraint on the VBP variation equation. This establishes a set of states with a given quadrupole deformation for which the energy is locally minimal. These states are then projected onto states of good angular momentum, and used as a basis for the many-body state. Like orthogonal VAP, the eigenstates of the AMD Hamiltonian are determined through simultaneous diagonalisation of the norm and Hamiltonian matrices [21].

Constrained solutions are particularly useful in obtaining basis states for loosely bound systems, e.g. those of cluster or halo nature; the solution given by VAP usually corresponds to a compact configuration [31].

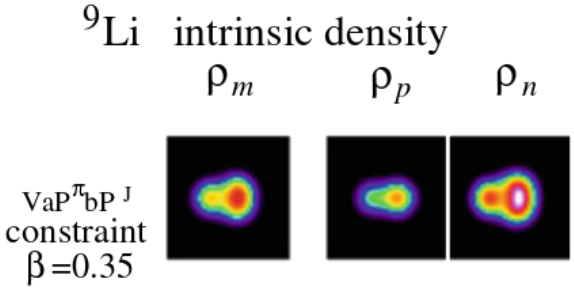


Fig. 3.10: Density distributions of a singular AMD wave-function for ^{9}Li obtained by variation *after* parity and *before* angular momentum projection for a deformation constraint of $\beta = 0.35$. The proton (ρ_p), neutron (ρ_n), and total (ρ_m) densities of the AMD wave-function are illustrated. See Fig. 3.9 for comparison to other variational solutions. Figure adapted from Ref. [21].

3.2.6 Nuclear Molecules

The pairing force, which anti-aligns the angular momenta of like-nucleons within the same orbital, is responsible for some of the most exciting phenomena in nuclear physics, e.g. the even-even energy gap between the ground state and lowest single-particle excitation, or the zero angular momentum of even-even ground state nuclei. It is most strikingly seen within the α-particle system, which is comprised to two correlated pairs of protons and neutrons. In combination with the small symmetry energy of the system, the pairing interaction is responsible for the significant binding energy (28 MeV), and first excited state (20 MeV). These two properties suggest that where an α-particle forms in the nucleus, it may remain stable for a considerable period of time [33]. Meanwhile, the emission of α-particles from the nucleus implies that such particles might be preformed and exist prior to the decay. Although it is known that such an assumption is overly reductive, the existence of many alpha-clustered structures is postulated by the Ikeda threshold rule, which predicts the emergence of such structures near the alpha breakup threshold [34].

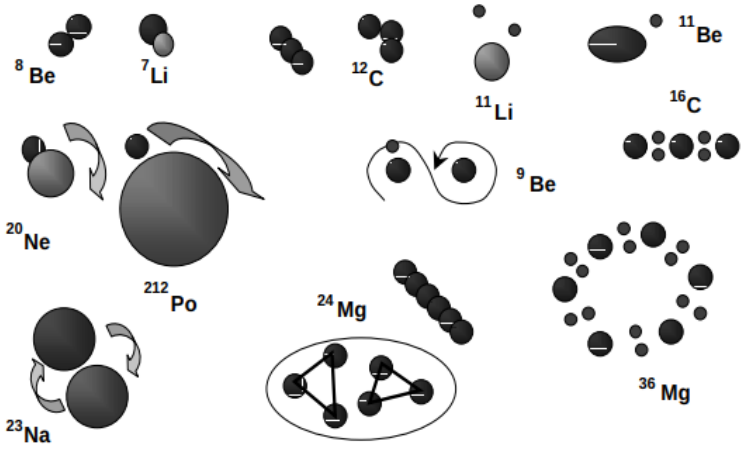


Fig. 3.11: A myriad of clustering phenomena are seen in light nuclei, varying from small clusters outside of a closed shell to alpha condensates. Figure taken from Ref. [34].

A number of examples of "clustering" phenomena, including alpha condensates (e.g. ^{8}Be), can be seen in Fig. 3.11. In recent times, a focus has been placed on neutron-rich nuclei, where the existence of valence neutrons has been shown to stabilise the cluster structure and give rise to a class of nuclear molecular states [1]. In the case of ^{8}Be; the alpha condensate has a half-life of $\sim 10^{-16}$s, whereas its sibling ^{9}Be is considered stable. This "Borromean"-like behaviour follows from the behaviour of the valence neutron, which may be thought of as occupying a delocalised orbit shared between separate alpha-particle clusters [35]. The term "nuclear molecule" is well-chosen; in the case of the beryllium isotopes, the effective α-α potential has three modes: weakly attractive at intermediate distances; strongly repulsive at smaller distances, due to the exchange interaction; and moderately repulsive at larger distances, due to Coulomb repulsion [1]. The nature of this interaction closely resembles the Van Der Waals molecular potential, giving rise to the term "nuclear molecule". In general, molecular structures are strongly predisposed to form chain-like states in order to minimise the Coulomb repulsion between the constituents

π bonds & σ bonds

The Hückel method, taken from molecular orbital (MO) theory, can be used to describe the molecular orbit occupied by a valence neutron as a superposition of the single-centre shell model states:

$$|\psi_{\pm}\rangle = \frac{1}{\sqrt{2}}(|\psi_1\rangle \pm |\psi_2\rangle)\,. \tag{3.36}$$

This yields two molecular orbitals, conventionally labelled *g* ($|\psi_+\rangle$) and *u* ($|\psi_-\rangle$). Each has both a total spin J and an orbital angular momentum l. The projection m_l of l admits two further families of orbitals; σ orbits with $m_l = 0$, and π orbits with $m_l = 1$. This labelling of the m_l projections follows from Molecular Orbital theory. Of course, in a non-spherical system the total angular momentum J is no longer a good quantum number, so the symmetry-axis projection K is taken instead. The parity of these states is given by $\pi = (-1)^l p$ where $p = 1$ for *g* orbits, and $p = -1$ for *u* orbits [7].

The Hückel wave-function amplitudes for the various orbital configurations are shown in Fig. 3.12.

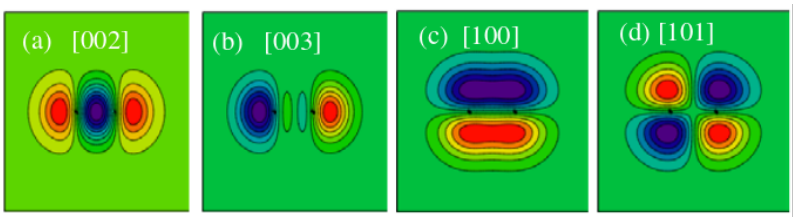


Fig. 3.12: One-dimensional molecular orbits for cluster states in ^{9}Be; showing (a) σ bonding, (b) σ anti-bonding, (c) π bonding, (d) π anti-bonding. In this 2D projection, the deformation axis lies in parallel to the plane normal. Figure adapted from Ref. [35].

It can be seen that the π and σ bonds in Fig. 3.12 bare close resemblence to their namesakes in atomic molecular orbitals, shown in Fig. 3.13.

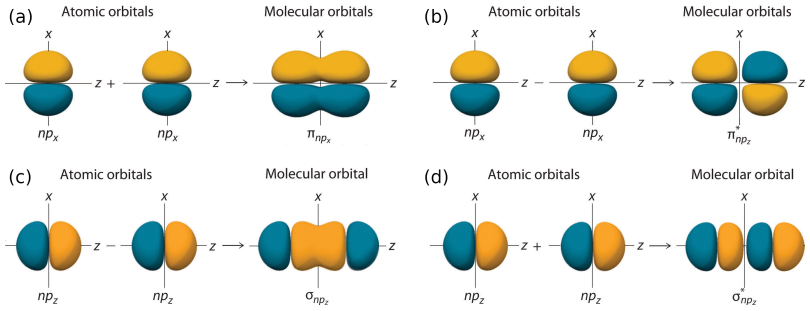


Fig. 3.13: One-dimensional molecular orbits from atomic physics; showing (a) π bonding, (b) π anti-bonding, (c) σ bonding, (d) σ anti-bonding. Figure adapted from Ref. [36].

The π orbit is also known as the *ring orbit*, because it spreads perpendicular to the α-α axis, whilst the σ orbit is called the *chain orbit* it spreads parallel to the α-α axis (see Fig. 3.14) [37].

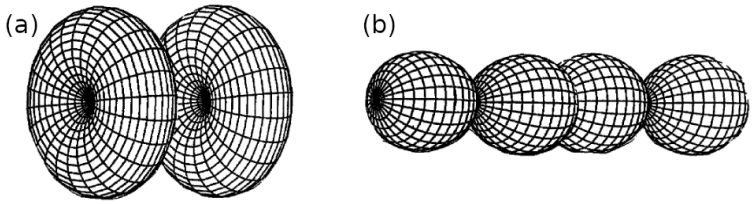


Fig. 3.14: Illustration of the spatial overlap of two p-shell orbits, corresponding to (a) π and (b) σ molecular orbitals. Figure adapted from Ref. [1].

In ^{9}Be, molecular states comprise two clusters $\alpha + \alpha + n \rightarrow {}^{5}\text{He} + \alpha$, whereby molecular orbitals are constructed from the $p_{3/2}$ shell-model orbit of the valence neutron in ^{5}He. A two-centre model of the nucleus has been explored in [38], and a neutron level scheme computed as a function of the centre separation (see Fig. 3.15). Though a formal model can be used to build this scheme [39], significant insight can be derived from mapping the separate-centre shell-model states onto states in the Nilsson deformed harmonic oscillator. With known spin and parity, the Nilsson orbits can be used to label the two-centre solutions and define a set of molecular orbitals [40]. By correlating the states in Fig. 3.15 with those in Fig. 3.16, one can identify the four molecular orbitals available to the $1p3/2$ neutron in ^{5}He: the $3/2^{-}$, $1/2^{+}$, and $1/2^{-}$ states in ^{9}Be correspond to the $\left(\pi_{3/2^{-}}, \mathrm{g}\right)$, $\left(\sigma_{1/2^{+}}, u\right)$, and $\left(\pi_{1/2^{-}}, g\right)$ orbitals [41] (see Fig. 3.17). The complete set of low-lying states in ^{9}Be are consistent with rotational bands built upon these basis [41].

The set of basis states established from this two-centre approach can be extended in the case of a three-center molecule, such as the carbon isotopes. This additional center primarily splits the basis states, and lowers the lowest state with respect to the three-body particle threshold [43]. In addition, the formation of asymmetric structures may lead to the phenomenon of parity doublets (see Section 3.1.2).

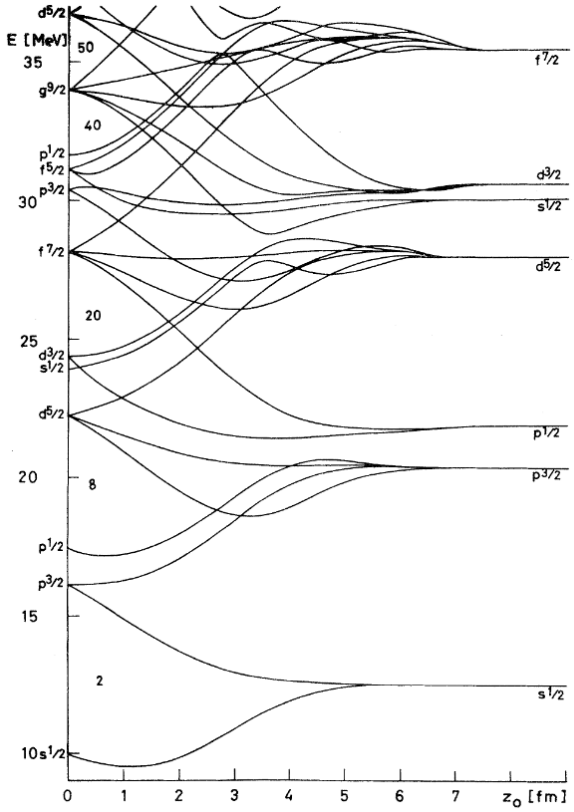


Fig. 3.15: Neutron level scheme of the two-centre shell-model as a function of the deformation eccentricity $z_0 \in [0, R]$, where R is the nuclear radius. The y-axis intercepts correspond to spherical shell-model states, whilst at large separations the spherical shell-model states of the fragments are indicated [38].

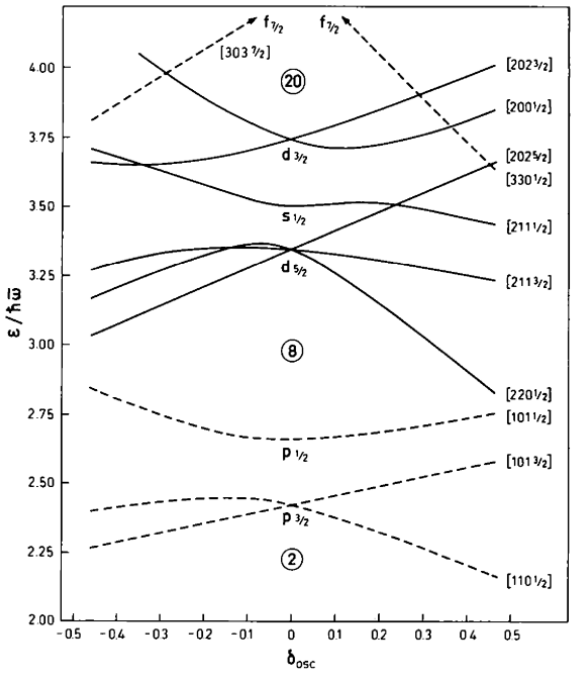


Fig. 3.16: Single-particle level scheme of the Nilsson deformed harmonic oscillator potential, for $N < 50$. Solid lines indicate positive parity solutions, whilst dashed lines represent negative parity states. Each state is labelled by $[Nn_z\Lambda\Omega]$, where N is the major oscillator number, n_z is the oscillator quanta along the 3-axis, Λ the z projection of the orbital angular momentum, and Ω the z projection of the total angular momentum J. Figure taken from Ref. [42].

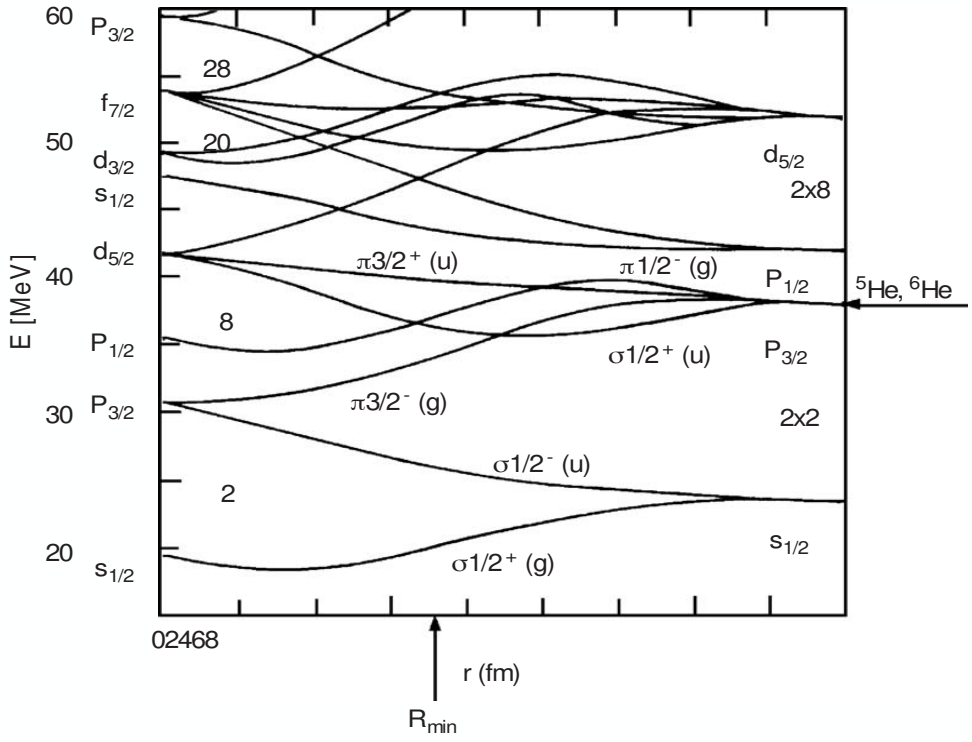


Fig. 3.17: Correlation diagram of molecular orbitals for a two-center α-α system with the $p_{3/2}$ shell-model orbit of the valence neutron in ^{5}He indicated. The location of the potential minimum is labelled R_{min}. Each orbital state is labelled by $[K, \pi, u/g]$, where K is the projection of the total angular momentum J onto the cluster axis, π the orbital parity, and u/g the ungerade/gerade orbital family. From Fig. 3.16, the $p_{3/2}$ shell-model orbit for separate centres can be identified with Nilsson orbits in the deformed model. Figure taken from Ref. [1].

3.3 Conclusion

From the discussion in this chapter, it is evident that the landscape of nuclear models is both rich and varied. It suffices to say that there is yet no single model that can describe the full set of physical phenomena that nuclei have been observed to demonstrate, despite new approaches such as the no-core shell model showing promise in the emergent description of collective behaviour alongside the established single-particle description of the nucleus.

Collective models provide a basis for describing the macroscopic phenomena that is easily visible in level schemes, such as rotational bands, or collective vibrational modes (though the latter is omitted in this chapter). Alongside single-particle models, collective descriptions of the nucleus provide a lens to interpret the behavior of the nucleus.

The Hartree-Fock description of the nucleus both serves to elucidate the single-particle behavior of many nuclear states, and provides a language for describing the further developments in nuclear models such as AMD, just as the shell-model is used to describe phenomena that escape single-particle descriptions. The Hartree-Fock description in this chapter is sufficient to understand its function, but leaves room for a more comprehensive, mathematically rigorous treatment.

The AMD model of the nucleus proves particular interesting to the discipline of cluster physics, as it provides computationally tractable predictions of the structure of these unusual nuclei. Later discussions in this thesis make particular reference to the AMD predictions of the structures in ^{14}O and ^{14}C, which provides some insight into the mechanisms that govern their observed level schemes. A focus was placed upon the variations of AMD that are in common practice; recent developments have led to VaP approaches becoming tractable and commonplace, providing richer predictions for far-from-mean-field structures.

Finally, a discussion of nuclear molecules provides a basis for describing the non alpha-conjugate cluster-like structures predicted in ^{14}O and ^{14}C. Understanding the level-schemes of these nuclei is an on-going field of research in nuclear structure, and molecular orbitals provide an approachable model for interpreting the sometimes unintuitive phenomena that arise, such as the binding behavior of valence nucleons.

CHAPTER

FOUR

CARBON-14 AND OXYGEN-14 MIRROR NUCLEI

4.1 Mirror Nuclei

In Section 3.1.1, the semi-empirical mass formula was introduced. Notably, it is only the Coulomb term of Eqn. 3.1 that explicitly depends upon a nucleon flavour (the proton). This hints at a symmetry of the strong interaction - that the Hamiltonian is degenerate with respect to nucleon charge. It is not an *exact* symmetry; the mass of the neutron is slightly greater than that of the proton, and thus the eigenvalues of the strong Hamiltonian are distinct. This difference can be treated as a *perturbation* of an exact symmetry, called Isospin, I, for its close resemblance to the spin-$\frac{1}{2}$ representation [44]. In this model, protons and neutrons form eigenstates of the spin-doublet system with 3 projection, $I_3 = \pm\frac{1}{2}$.

There are two interpretations of isospin symmetry:

Charge-symmetry

Charge-*symmetry* holds that the *pp* interaction is indistinguishable from that of the *nn* interaction. In terms of the isospin

$$\left[\hat{H}, \exp\left(\frac{i\pi\hat{I}_i}{\hbar}\right)\right] = 0\,, \tag{4.1}$$

where $\hat{I}_i$ is the operator that measures the $i \neq 3$ projection of the isospin I, and $\hat{H}$ is the strong Hamiltonian. The $\exp\left(\frac{i\pi\hat{I}_i}{\hbar}\right)$ rotation operator negates the $\hat{I}_3$ expectation value.

Charge-independence

Charge-*independence* holds that the *pp*, *nn*, and *pn* interactions are indistinguishable, i.e.

$$\left[\hat{H}, \exp\left(\frac{i\theta\hat{I}_i}{\hbar}\right)\right] = 0\,, \tag{4.2}$$

where θ is an arbitrary rotation in isospin space. It is therefore a stronger statement of charge-*symmetry*. The

eigenstates correspond to a set of *isobaric* nuclei; nuclei with the same total number of nucleons.

As implied from the construction of Eqn. 3.1, there is good evidence from measurements of nucleon-nucleon scattering lengths that the strong interaction is approximately charge-symmetric [45]. These same data indicate a stronger degree of charge-independence breaking, suggested by a larger (absolute) value for the *pn* scattering length. It follows that *mirror* nuclei, nuclei with exchanged proton and neutron numbers such as ^{11}C and ^{11}B, or ^{14}O and ^{14}C, should exhibit similar structures subject to modulation by the different Coulomb interactions.

The level schemes for a pair of mirror nuclei are shown in Fig. 4.1. It is clear that the two nuclei possess very similar structures; the ordering between levels is closely matched, with a general decrease in excitation energy associated with the additional long-range spin-independent Coulomb potential seen in ^{11}C. For this reason, mirror nuclei can serve as useful tools for probing the structure of hard-to-measure nuclei, and elucidating the role of valence nucleons in stabilising deformed structures.

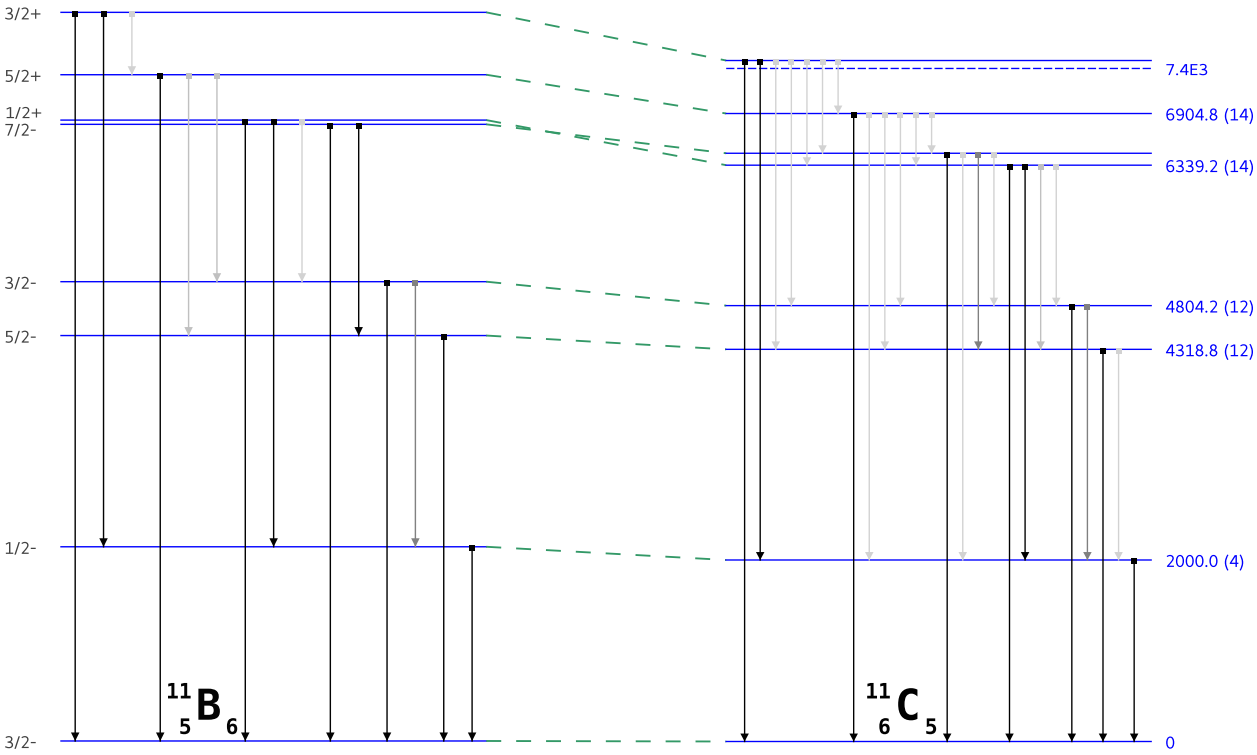


Fig. 4.1: Level schemes of the ^{11}C and ^{11}B nuclei. The order and spacing between the levels is well preserved under exchange of the protons and neutrons, with only the $\frac{7}{2}^-$ and $\frac{1}{2}^+$ states changing in order. Generated from data taken from the ENSDF database as of March 24th, 2023. Version available at http://www.nndc.bnl.gov/ensarchivals/.

4.2 Predicted Structures in Carbon-14 and Oxygen-14

Later in this chapter, an experimental measurement of the structure of ^{14}O is discussed. As the mirror of ^{14}O, ^{14}C is well studied. It therefore serves as a good foil against which to evaluate the performance of the measurement, in addition to the effects of charge-symmetry breaking upon the observed structure of the nucleus. As such, it is of interest to evaluate the landscape of structure predictions concerning the carbon analogue for which there is a greater body of research.

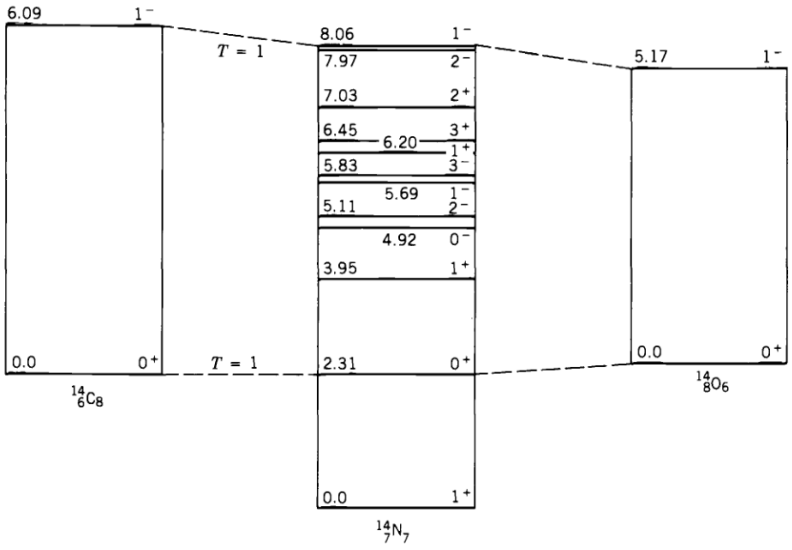


Fig. 4.2: Low lying energy levels of $A = 14$ isobars, with adjusted energies relative to ^{14}N. Levels in ^{14}C and ^{14}O have $T = 1$, whilst those in ^{14}N have $T = 0$ apart from those levels in common with the other members of the isobar. Figure taken from Ref. [46].

An exploration of the relationship between clustering and the deformed harmonic oscillator (DHO) model has been performed by Freer, Betts, and Wuosmaa in Ref. [47], demonstrating that spherical harmonic oscillator (SHO) shell-model structure appears within the DHO for integer deformation ratios ($\omega_x : \omega_y : \omega_z, \omega_i \in \mathcal{Z}$). The symmetries observed in these integral deformation ratios suggest that the deformed potential may instead be expressed in term of overlapping spherical oscillators. This provides a lens through which clustered structures may be understood as multi-centre spherical potentials. Subsequent work has applied this concept to propose the existence of a linear chain structure at a $3:1$ deformation. In this configuration, three quartets of paired protons and neutrons occupy the leading, degenerate SHO *s* shells; whilst the two valence neutrons lie on perpendicular orbitals given by the shared $p_{\frac{3}{2}}$ state (see Fig. 4.3). This chain is predicted to form a π-orbital, which the DHO model associated with the lowest-energy linear chain system by minimising the axial deformation [1]. See Fig. 4.4 for the Hückel wave-function amplitudes that correspond to the valence nucleon orbits in ^{13}C.

One of the earliest conjectures as to the existence of molecular structures in ^{14}C was that of a near-threshold ^{9}Be + ^{5}He isomeric state at ~17.24 MeV given by Ref. [48] (see (x) in Fig. 4.5). This prediction followed from the postulate

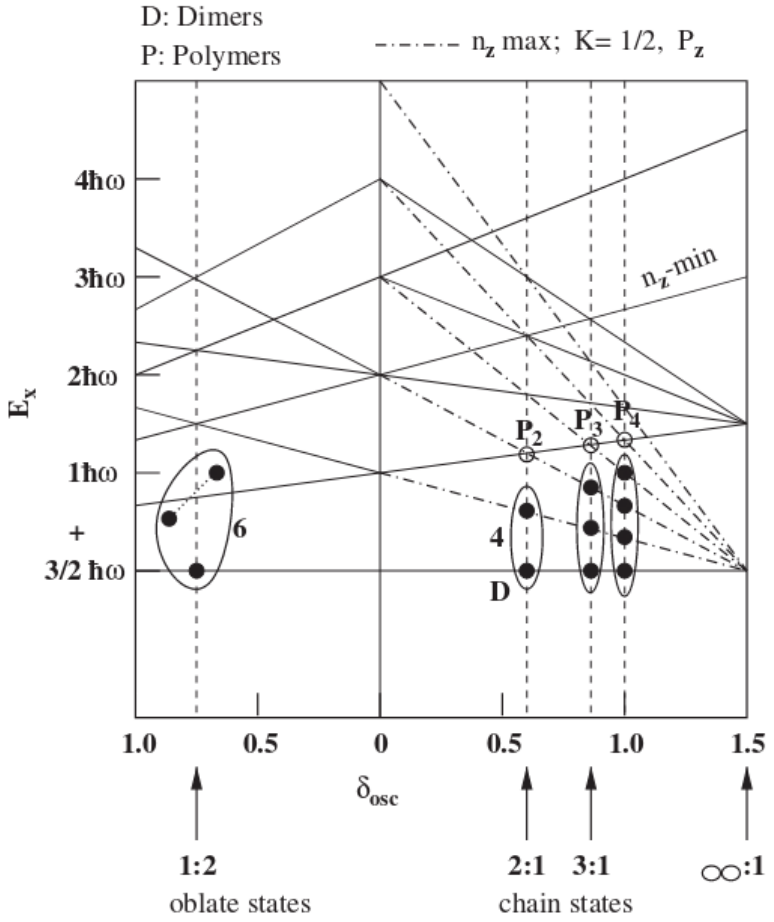


Fig. 4.3: Energy diagram of shell-model degeneracies for integer deformation ratios in an axially deformed harmonic oscillator. The open circles indicate the shells occupied by valence neutrons, whilst the shaded dots denote the shells occupied by the isospin-degenerate $J_z = 0$ nucleon pairs that form an α-particle. Figure taken from Ref. [1].

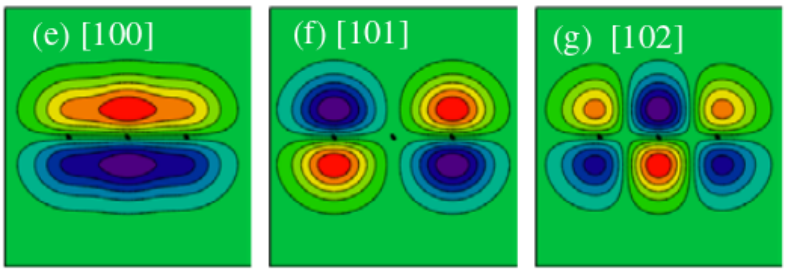


Fig. 4.4: One-dimensional molecular orbits for cluster states in ^{13}C; showing π. Note that the number of nodes $[n_1 n_2 n_3]$ increases with increasing oscillator quanta. In this 2D projection, the deformation axis lies in parallel to the plane normal. Figure adapted from Ref. [35].

that bound *multimers* could be formed by adding helium isotopes to the known beryllium dimer structures, whereby the additional valence neutrons would form covalent bonds between the next two α-particles.

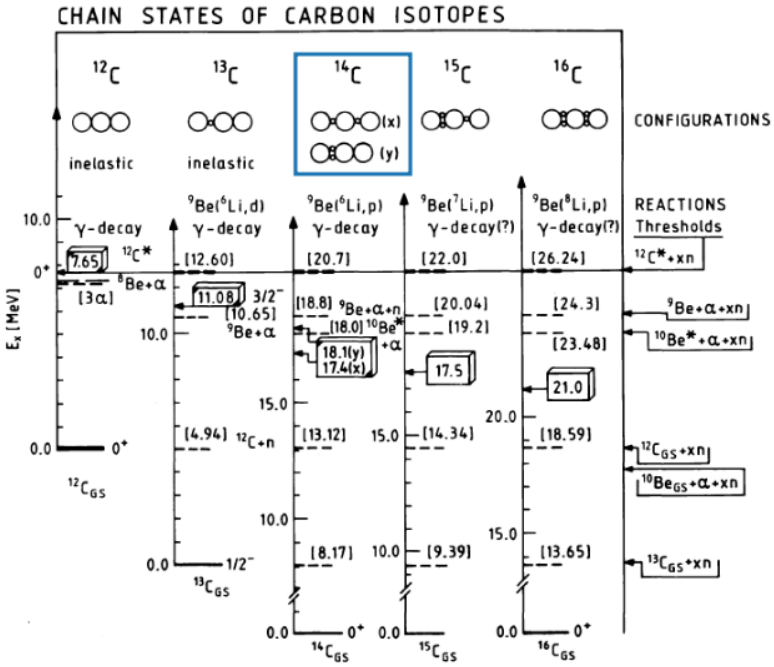


Fig. 4.5: Energy diagrams of molecular chain states in carbon isotopes, adapted from Ref. [41]. The black boxes indicate the first two-center non-rotational isomeric states. The predicted configurations for ^{14}C are shown in the blue box.

Extending the concept of near-threshold states, von Oertzen further anticipated resonances in the $^{10}\text{Be}^{*} + \alpha$ system at 18 MeV [41]. These resonances lie, below the ~18.8 MeV $^{8}\text{Be} + 2n + \alpha$ threshold, due to the covalent binding of the valence neutrons [1]. It was predicted that a series of four states would arise from the group of states at 6 MeV in $^{10}\text{Be}^{*}$ with spins 0^{+}, 2^{+}, 1^{-}, and 2^{-}. The reflection-symmetric σ-orbit configuration, with valence nucleons distributed equally among the alpha particles (see (x) in Fig. 4.5), is expected to be found at lower energies than the asymmetric π-orbit configuration due to the stronger binding effect of each valence neutron. It is anticipated to form a rotational band comprising of 0^{+}, 2^{+}, and 4^{+} [41].

The prolate linear-chain structures described above have two deformation modes, shown in Fig. 4.6. In the alpha-conjugate nucleus ^{12}C, there is an absence of experimental evidence for pure linear chain structures, indicating that it is unstable with respect to these bending and breathing modes. Itagaki, Okabe, Ikeda, and Tanihata have applied the Molecular Orbital model, discussed briefly in Section 3.2.6, to predict the behaviour of these conjected structures [49]. The orbitals that are obtained by this method are classified in terms of their projections onto states of good momentum projection and parity, K^{π}.

Calculations of ^{14}C for the given configurations identify minima in the binding energy with respect to molecular centre separation; the $^{14}\text{C}(3/2^{-}_{\pi})^{2}$ configuration appears to possess a minimum at an excitation energy of ~18 MeV [49]. Whilst the $^{14}\text{C}(1/2^{-}_{\sigma})^{2}$ configuration is predicted to exhibit a minimum at ~32 MeV, it is not sufficiently deep enough to be

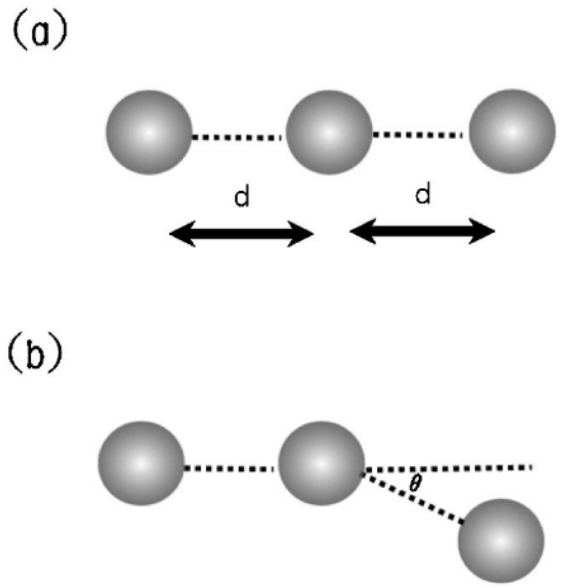
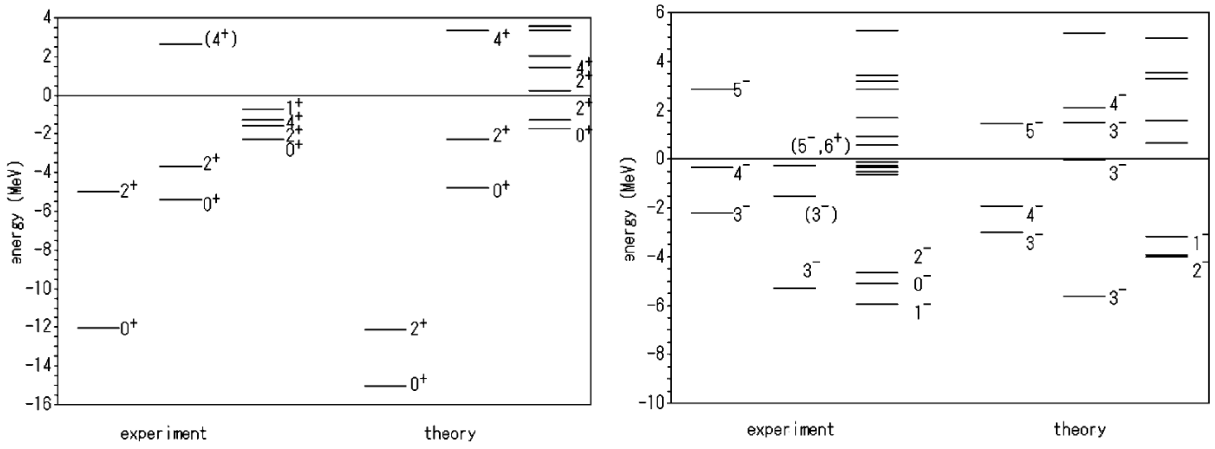


Fig. 4.6: Schematic illustration of the deformation modes of linear-chain states. (a) shows the *breathing* mode in which the centre-separation, d, varies, whilst (b) depicts the *bending* mode that depends upon the bending angle θ. Figure taken from Ref. [49].

associated with stability against breathing-mode deformation. Against bending-mode deformation, [49] do not predict stability for any of the molecular configurations in ^{14}C. Such stability is only predicted for $^{16}\mathrm{C}\left((3/2^{-}_{\pi})^{2}(1/2^{-}_{\sigma})^{2}\right)$, due to the overlap between the two neutrons in the π orbit and the two neutrons in the σ; an increase in mutual overlap between the valence orbitals leads to a decrease in the total-wavefunction overlap component due to Pauli blocking [49].

Milin and von Oertzen briefly predict the existence of three-centre states in ^{14}C that build upon the molecular orbital structures identified for ^{13}C [43]; $(\pi)^2$, $(\sigma)^2$, and $(\pi \otimes \sigma)$. The dominant structures are conjectured to be based upon the ground state $K = 0(\pi)^2$, the excited state $K = 0(\sigma)^2$, and the $K = 1^{-}(\sigma \times \pi)$ state of ^{10}Be.

A long conjectured state in three-alpha systems is that of an equilateral triangle configuration, which would exhibit a series of rotational bands starting at 3^{-} and 0^{+} [50] [51]. Though such a state lies above the three-alpha threshold in ^{12}C, it is expected that the valence neutrons of ^{14}C should stabilise the triangular configuration. Using the Generator Coordinate Method (GCM, see Section 3.2.5), Itagaki, Otsuka, Ikeda, and Okabe [51] make a series of predictions for the energy levels in ^{14}C, identifying the two aforementioned rotational bands (see Fig. 4.7). It is observed experimentally that the second 3^{-} band is easily excited by α transfer, which indicates an alpha-cluster structure. Simultaneously, the band based upon the (fourth) 0^{+} state is observed to exhibit strong B(E2) transition rates, supporting the indication of a rotational band. The authors emphasise the difference between the Ikeda near-threshold clustering, with localised excess neutrons, and the crystalline alpha-condensate structures that can form well below the threshold energy due to the stabilisation of the delocalised valence neutrons. This observation represents a distinct mechanism for cluster states beyond the localised molecular orbits seen in the linear-chain structures.



(a) Positive parity states.

(b) Negative parity states.

Fig. 4.7: Energy levels in ^{14}C for positive (a) and negative (b) parity states. Energies are measured relative to the ^{10}Be$+\alpha$ threshold, and the experimental states are plotted against the predictions for comparison. Figure taken from Ref. [51].

In a large work that undertakes a systematic review of experimental data, von Oertzen, Bohlen, Milin, Kokalova, Thummerer, Tumino, Kalpakchieva, Massey, Eisermann, Graw, Faestermann, Hertenberger, and Wirth identify single-particle states with oblate shapes, and prolate states with nontrivial structure associated with the low-lying oblate structure of ^{12}C [52]. These latter states are proposed to have strong α-clustering and form rotational bands with signature parity inversion doublets and a high moment of inertia. The complete set of anticipated structures is schematically depicted in Fig. 4.8. This work finds itself in agreement with Ref. [51] as to the existence of an oblate triangular configuration with $K^\pi = 0^+, 3^-$ band heads. von Oertzen, Bohlen, Milin, Kokalova, Thummerer, Tumino, Kalpakchieva, Massey, Eisermann, Graw, Faestermann, Hertenberger, and Wirth ascribe the prolate structures built upon the 0^+_3, 1^-_2states to chain structures [52], though do not go so far as to identify these as either the symmetric or anti-symmetric structures described above.

Results of performing AMD simulations of the negative parity and positive parity states are given in Ref. [53]. The findings agree with the existing experimental level assignments such as those described in [52]. In particular, both shell-model and oblate/triaxial structures are identified at low energies below the ^{10}Be $+ \alpha$ threshold. The triaxial structures are observed in $K^\pi = 0^+$ and $K^\pi = 2^+$ bands, with the excess neutrons occupying the *sd*-like orbitals [53]. A prolate $K^\pi = 0^+$ band is also identified thereafter, with valence neutrons localised about two of the three alpha clusters. Such a configuration suggests a ^{10}Be correlation, in agreement with prior experiment. An investigation of the above-threshold decay in ^{4}He$(^{10}$Be, $\alpha)^{10}$Be in Ref. [54] provides some agreement with this structure determination, but notes that lower excitation energies would be useful to observe the angular distributions required to resolve the structural identity with more detail [54]. Further AMD work by Baba and Kimura in Ref. [55] adds to the theoretical support for linear-chain structures in ^{14}C, and remarks that the observation of decays to both the ground and 2^+ states in ^{10}Be would serve as additional evidence for linear chain formation. With these results, there is still little evidence for linear chain structures

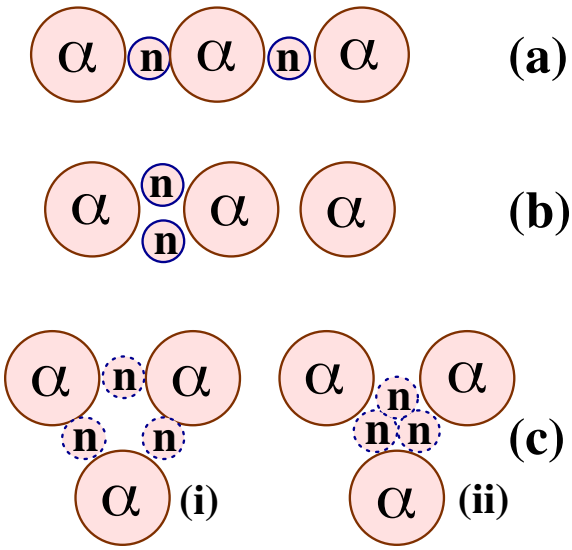
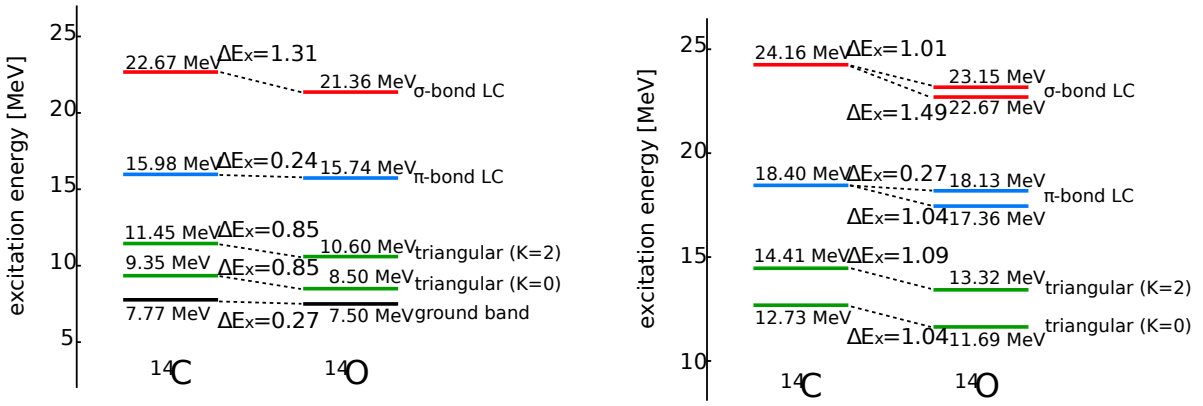


Fig. 4.8: Schematic illustration of the cluster structures anticipated in ^{14}C for various valence neutron configurations. (a) depicts the symmetric, positive parity $\alpha+n+\alpha+n+\alpha$ configuration with the valence neutrons distributed outside of the symmetry axis; (b) shows the reflection-asymmetric inversion doublets, based upon ^{10}Be $+ \alpha$ structure with the valence neutrons in the same covalent π-bond; and (c) illustrates the two oblate triangular-cluster states with σ-bonds between two α-particles, whose valence particles wavefunctions have (i) a node at the center of the nucleus with $J^\pi = 3^-$ and (ii) a maximum at the center of the nucleus with $J^\pi = 0^+$. The dashed lines indicate three sites between which the two valence neutrons are distributed. Figure adapted from Ref. [52].

with negative parity [55]. A recent experimental work [56] has been performed in a similar manner to this work, using resonant scattering in inverse kinematics. The findings of this work indicate the existence of several alpha-cluster-like states in the 14 MeV–19 MeV excitation energy region. The observed spin-parities, and level separation, is given as a strong indication of the linear-chain cluster state (LCCS) structure.

Of particular relevance to this work is that of Ref. [57], which considers the Coulomb shift of the mirror pair ^{14}C and ^{14}O. The AMD calculations performed in this work suggest that the excitation energy of the π-bond linear chains seen in ^{14}C should lie close to those in ^{14}O, with similar decay widths. The findings for the σ-bond chain are somewhat different; it is anticipated that the levels in ^{14}C have both larger decay widths and distinctly higher excitation energies, which are ascribed to a large Thomas-Ehrman shift (see Fig. 4.9). The Thomas-Ehrman shift describes a pronounced shift in the energy of the first $1/2+$ state; s-states are able to avoid significant Coulomb repulsion due to the lack of a centrifugal ($l > 0$) barrier which confines the wavefunction inside the Coulomb barrier [58].



(a) (a)

(b) (b)

Fig. 4.9: (a) 2^+ and (b) 4^+ states in ^{14}C and ^{14}O determined by AMD calculation. Close correspondence between the π-bond chain states is observed, whilst the σ-bond levels are expected to differ significantly in excitation energy and decay width. Figure taken from Ref. [57].

4.3 Conclusion

The symmetries of the residual strong interaction provide a mechanism by which to elucidate the structure of a poorly understood nucleus using that of its mirror. ^{14}O is one such nucleus, whose level scheme is less comprehensive than that of ^{14}C. In this chapter, the justification for using one nucleus as a means of understanding the structure of its mirror has been established. A discussion was made of the degree to which this model is appropriate, with particular emphasis on the shift in the potential that arises from the additional Coulomb contribution, and the discrepancy in the form of a Thomas-Ehrman shift.

Using ^{14}C as a reference, the literature was summarised to give an indication as to the predicted structures in ^{14}O, and their properties. From this discussion, the notion that a linear chain LCCS state may be observable in the 14 MeV—19 MeV region was established. Additional molecular structures were also conjectured, with two triangular configurations, and a second linear state. Many of the referenced works used AMD simulations to establish their predictions, which is discussed in Section 3.2.5.

CHAPTER

FIVE

THICK TARGET EXPERIMENTS

5.1 Resonant Scattering

Nuclear reactions may be classified with respect to two extremes modes; *compound* reactions in which the reactants fuse into a highly excited compound nucleus, and *direct* interactions that involve typically only the surface nucleons of the target (see Fig. 5.1). As these latter reactions transition directly between initial and final states, the entrance and exit channels are strongly correlated. Meanwhile, the decay channel of the liquid-drop formed by coalescence of the target and projectile is considered to be reasonably independent of the entrance channel [59]. Consequently, the reaction cross-section is typically modelled as two separable terms [60].

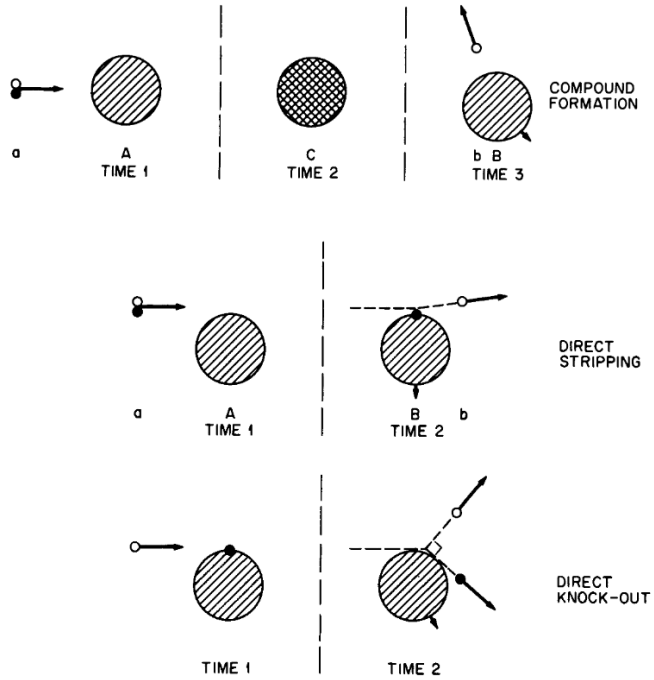


Fig. 5.1: Semi-classical illustration of the limiting extremes to nuclear reactions: formation of a transitory compound nucleus, and direct reactions. Figure taken from Ref. [60].

Between two charged nuclei, there exists a barrier to fusion that follows from the intrinsic positive charge of the nucleus. The magnitude of this barrier can be approximated through a classical treatment of the system, whereby the two nuclei are taken to be point charges separated by the sum of the nuclear radii. It follows that an interaction between two colliding nuclei may form a compound nucleus only if there is sufficient energy to overcome this barrier.

At low incident energies, the compound nucleus may exhibit discrete states, despite being unstable against particle emission. Although there is sufficient energy for ejection of one or more particles, this energy is typically partitioned among several nucleons. Decay of this compound state occurs only once the excitation energy localises itself sufficiently for particles to escape. It follows that the compound nucleus may demonstrate a lifetime much greater than the orbital period of a single nucleon ($\sim 10^{-22}$ s) i.e. the duration of a direct interaction [60]. In the low-energy region at which resonances are found, there are usually only two decay modes available to the liquid drop; rejection of the incident particle (scattering), or γ-emission [46]. Given that a quasi-bound state is unstable, it possesses an imprecise energy; from the Mandelstam-Tamm uncertainty principle [61]

$$\Delta H \Delta T \geq \frac{\hbar}{2}\,, \tag{5.1}$$

one might assume that the long lifetime $\tau = \Delta T$ of the quasi-bound compound state implies that the width of this distribution $\Gamma = \Delta H$ is smaller than 1 MeV. Where the energy of the entrance channel is sufficiently matched with that of the quasi-bound state, there will be a marked enhancement in the cross section of the reaction, a phenomenon known as a *resonance*. Typically this has the form of the Breit-Wigner distribution

$$\sigma_{\text{res}}(E) = \sigma_0 \frac{\frac{1}{4}\Gamma^2}{\left(E - E_{\text{r}}\right)^2 + \frac{1}{4}\Gamma^2}\,, \tag{5.2}$$

where Γ is the full-width at half-maximum (FWHM) of the resonance (see Fig. 5.2), E_r the centroid of the distribution, and σ_0 the amplitude of the resonance [60]. These resonant states are only visible at sufficiently low energies; as the excitation energy increases, the separation between levels decreases until the density is such that individual states can no-longer be resolved and form a continuum of unbound states. It transpires that in practical terms the energy relation in Eqn. 5.1 is inapplicable to the Breit-Wigner curve, yet it can be shown that the lifetime of the state τ is related to its width by Eqn. 5.3 [62] [63].

$$\Gamma = \frac{\hbar}{\tau}\,. \tag{5.3}$$

Among the exit channels of the compound nucleus is an elastic channel, i.e. that which corresponds to the $A(a,a)A$ reaction. When the elastic scattering cross section is measured, however, it typically does not exhibit the predicted Breit-Wigner shape. In addition to forming a resonant quasi-bound state in the compound nucleus, the incident particle may

simply scatter directly off the target. This leads to an additional shape-elastic potential contribution to the elastic scattering cross-section (see Fig. 5.2).

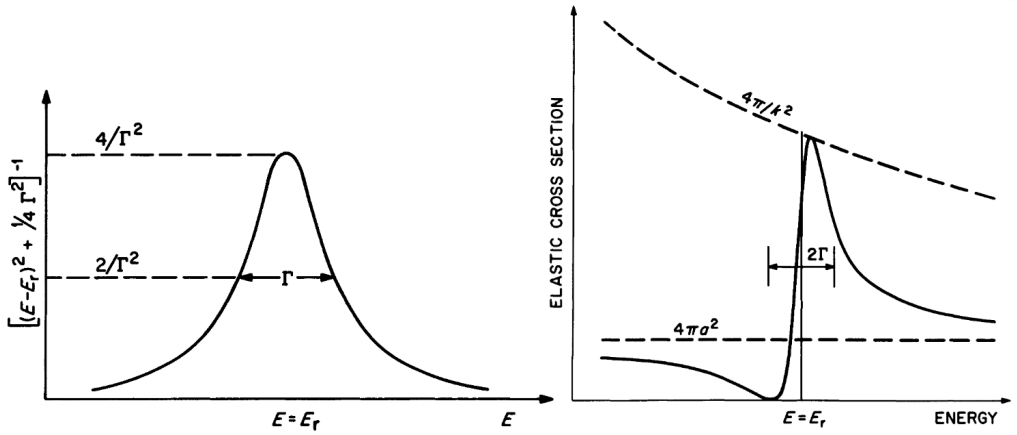


Fig. 5.2: Comparison between the (left) Breit-Wigner distribution and the (right) elastic scattering cross section. (right) includes interference between the shape-elastic and resonance-elastic interactions. Subfigures taken from Ref. [60].

5.2 Conservation Rules

As a closed physical system, the normal conservation rules still apply to a scattering reaction between two nuclei. The total energy of the system must therefore be conserved during the interaction. Assuming a resonance reaction of the form $A(a,b)B$,

$$T_A + m_A c^2 + T_a + m_a c^2 = T_B + m_B c^2 + T_b + m_b c^2 \tag{5.4}$$

where m_X is the *rest mass* of X, which may be excited, and T_X is the kinetic energy of X in the current frame. With the following definition of the reaction Q-value with respect to the ground-state masses (e.g. m_X^{gs})

$$Q = \left(m_A^{\text{gs}} + m_a^{\text{gs}} - m_B^{\text{gs}} - m_b^{\text{gs}}\right), \tag{5.5}$$

Eqn. 5.4 simplifies to

$$T_A + T_a + Q = T_B + T_b + E_{\text{ex}}, \tag{5.6}$$

where E_{ex} is the total excitation energy shared between B and b. For resonant reactions involving a compound nucleus X, one must consider the intermediate state. If both the reactants and the products of the total reaction remain in their

ground states, then

$$T_A + T_a + Q = T_X + E_{\text{ex}} = T_B + T_b \,, \tag{5.7}$$

i.e. measurement of the reactant and product kinetic energies is sufficient to determine the excitation energy for a known Q-value.

In addition to conservation of energy, the conservation of angular momentum must also be considered. For $J = 0$ reactants, the angular momentum of the compound nucleus is entirely determined by the orbital angular momentum of the beam

$$\begin{aligned} \hat{\mathbf{J}}_{\mathbf{X}} &= \hat{\mathbf{J}}_{\mathbf{A}} + \hat{\mathbf{J}}_{\mathbf{a}} \\ \hat{\mathbf{J}}_{\mathbf{X}} &= \hat{\mathbf{L}}_{\mathbf{A}} + \hat{\mathbf{I}}_{\mathbf{A}} + \hat{\mathbf{L}}_{\mathbf{a}} + \hat{\mathbf{I}}_{\mathbf{a}} \\ \hat{\mathbf{J}}_{\mathbf{X}} &= \hat{\mathbf{L}}_{\mathbf{a}} \,, \end{aligned} \tag{5.8}$$

where $\hat{\mathbf{J}}_{\mathbf{X}}$ is the total angular momentum of X, the compound nucleus; $\hat{\mathbf{I}}_{\mathbf{X}}$ the intrinsic spin angular momentum of X; A, the target; and a, the beam. The vector $\hat{\mathbf{J}}_{\mathbf{X}}$ is *by definition* perpendicular to the momentum vector of the beam, in the event of a two-body scatter. In combination with conservation of parity, the conservation of angular momentum implies that one can only populate either natural-parity ($\Pi = (-1)^l$) or unnatural parity ($\Pi = (-1)^{l+1}$) states for spin-0 reactants. The same holds for the subsequent decay of the resonant states; if the products are both spin-0, then the orbital angular momentum of the products is defined entirely by the parity of the populated state.

5.3 Time-Projection Chambers

A time-projection chamber (TPC) is a particle detector which employs the use of a time-sensitive anode to perform three-dimensional reconstruction of particle trajectories. The original design for the TPC was developed at Lawrence Berkeley Laboratory [64], and has since seen multiple variations that improve the spatial resolution and gain. At the heart of the TPC is the time-sensitive anode, upon which a current is induced due to the drift of charges liberated by ionising radiation incident upon the active volume (see Fig. 5.3). Conventionally, time-projection chambers use gas mixtures such as P10 (90% Argon, 10% Methane) as the interaction medium; gas targets may easily be pressurised to tune the density according to the energies and families of particles that are likely to be observed. In order to produce signals which can be discriminated against the electronic noise of the acquisition system, gas multiplication is typically employed.

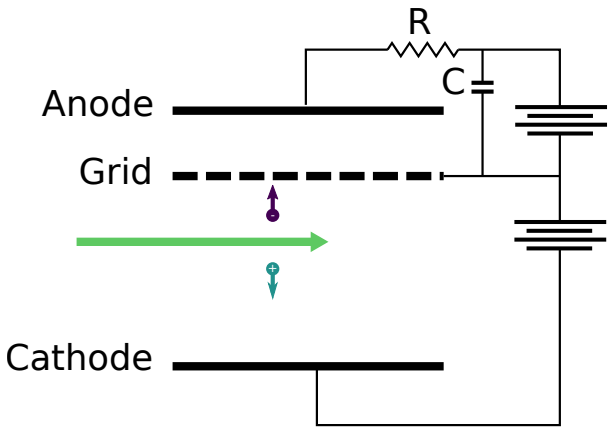


Fig. 5.3: Operational principle of a gridded ion-chamber, in which charged particles ionise the gas to liberate electron-ion charge carriers that move under an applied field. The movement of charge carriers induces a charge upon the anode and cathode. The grid serves to remove the positional dependence of the signal amplitude when operated in electron-sensitive mode; only the electron component contributes to the measured voltage across the resistor R once the electrons enter into the grid-anode region.

5.3.1 Gas Ionisiation

A charged particle moving through a gas medium loses energy by both excitation and ionisation of the target, to a similar extent. Each ionising collision is a stochastic process, obeying Poissonian statistics:

$$P(n) = \left(\frac{(s/\lambda)^n}{n!} \right) \exp(-s/\lambda) \,, \tag{5.9}$$

where s is the length of a given segment, $\lambda = \frac{1}{N_e \sigma_i}$ is the mean separation between clusters, N_e is the electron number density of the gas, and σ_i the ionisation cross-section of an electron [65]. The probability that zero clusters are formed in an interval is $\exp(\frac{-s}{\lambda})$, which tends to zero as $\lambda \to 0$. Thus, clusters tend to form quite close together along a given track. Whilst most of these clusters consist of single ions liberated by the charged particle, a small proportion of collisions induce secondary ionisation whereby the liberated electron ionises additional molecules. In some cases, these electrons travel sufficient distance from the primary ionisation that they leave detectable ionisation trails, and are referred to as δ-electrons. It is observed that clusters with greater than a single electron tend to contain a significant number, contributing strongly to the mean number of electrons per unit path length. Between the Poisson counting statistics and distribution of the cluster size, the spatial resolution of a drift-chamber is limited by the properties of the gas.

The mean energy-loss of an ion moving through a material is given by the Bethe-Bloch formula [66], which relates several material properties

$$-\left\langle \frac{\mathrm{d}E}{\mathrm{d}x} \right\rangle = \frac{4\pi}{m_e c^2} \cdot \frac{nz^2}{\beta^2} \cdot \left(\frac{e^2}{4\pi\varepsilon_0} \right)^2 \cdot \left[\ln\left(\frac{2m_e c^2 \beta^2}{I \cdot (1-\beta^2)} \right) - \beta^2 \right] , \tag{5.10}$$

where c is the speed of light, ε_0 the vacuum permittivity, I the mean excitation potential of the material, z the proton number of the ion, E the energy of the ion, $\beta = \frac{v}{c}, m_e$ the electronic rest mass, e the elementary charge, and n the electron number density of the material. In the non-relativistic limit, this simplifies to

$$-\left\langle \frac{\mathrm{d}E}{\mathrm{d}x} \right\rangle = \frac{4\pi n z^2}{m_e v^2} \cdot \left(\frac{e^2}{4\pi\varepsilon_0} \right)^2 \cdot \left[\ln\left(\frac{2m_e v^2}{I} \right) \right]. \tag{5.11}$$

5.3.2 Charge Multiplication

The applied field between the anode and cathode serves to accelerate the liberated charge carriers within the detector. Due to collisions with gas atoms, the drift velocity of each charge carrier approaches a constant proportional to $\frac{E}{P}$, where E is the applied field; and P the gas pressure [67]. The utility of a gas chamber for energy measurement lies in a proportional response between the measured current/voltage and the energy deposited within the detector. Below a threshold field strength, recombination of the liberated ion pairs contributes to non-linearity and defines the *ion-saturation* region of the pulse-amplitude applied-field curve. At higher field strengths, exceeding a threshold field value E_{av}, free electrons drifting under E gain sufficient kinetic energy between collisions with the gas molecules to cause secondary ionisations. This demarcates the start of a proportional response mode. In typical gases at atmospheric pressure, the secondary ionisation potential tends to exceed the order of 10^6 V/m. The first Townsend coefficient, α, of the gas determines the fractional increase in the number of electrons per unit path length

$$\frac{\mathrm{d}n}{\mathrm{d}x} = \alpha \, \mathrm{d}x \ . \tag{5.12}$$

It follows that for spatially invariant fields, the number density of electrons is given by $n(x) = n(0)\exp(\alpha x)$. Thus, increasing the voltage above E_{av} leads to an increased gain that facilitates measurements of clusters that are seeded by only a few primary electrons. Increasing the applied field further leads to the onset of non-proportionality, in which the build-up of slow-moving positive ions is sufficient to deform the electric field within the detector. Quickly thereafter begins the Geiger-Mueller region of operation, in which the space-charge effect is sufficiently strong that the multiplication process is self-limiting, and the pulse amplitude from the detector is independent of the properties of the incident radiation (see Fig. 5.4).

Frequently, noble gases such as argon or helium are used as the primary gas, as they have small electronegativities that prevent electrons liberated by ionisation being absorbed by the fill gas. However, alone these gases are susceptible to low breakdown-voltages in which the proportional region is lost, and typically have poor gain characteristics. Addition of a second component with a smaller ionisation energy than the fill gas leads to the Penning effect, in which the number of ion pairs produced per unit energy is enhanced [69]. Furthermore, additive gasses can be used to suppress the breakdown

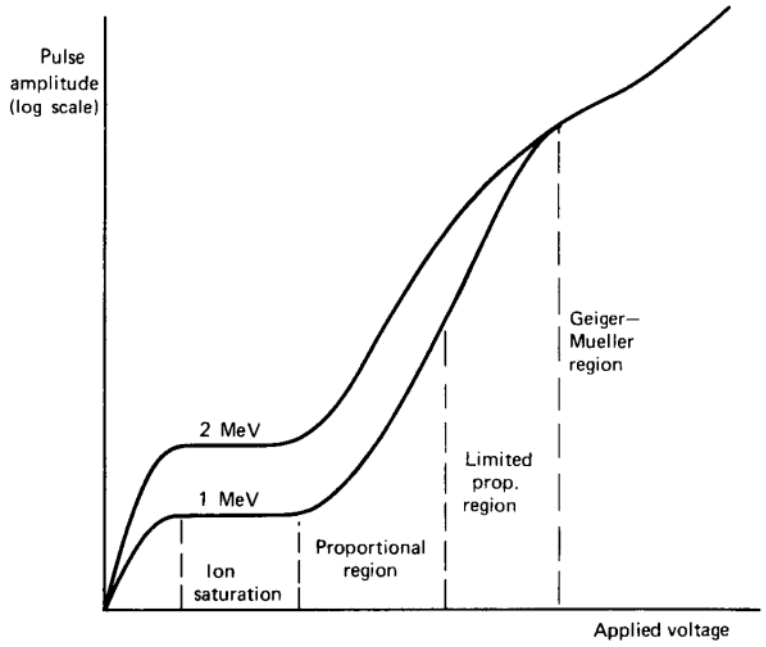


Fig. 5.4: The different operation regions of gas-filled detectors, according to the applied electric field. Only a limited range of field strengths corresponds to a proportional regime in the detector response. Figure taken from Ref. [68].

phenomenon; primary gas atoms excited by secondary electrons may de-excite through the emission of UV photons. These photons can seed additional avalanches within the active volume, leading to saturation. By adding a *quench* gas with additional degrees of freedom and a large UV photo-absorption cross section, e.g. diatomic hydrocarbons with rotational and vibrational modes, the UV photons emitted by secondary excitation can be absorbed without subsequent re-emission [70]. Finally, with a large enough field, the dielectric properties of a gas detector can break down. Admixtures that contain quench gases can also increase the value of the field required to induce this phenomenon [71].

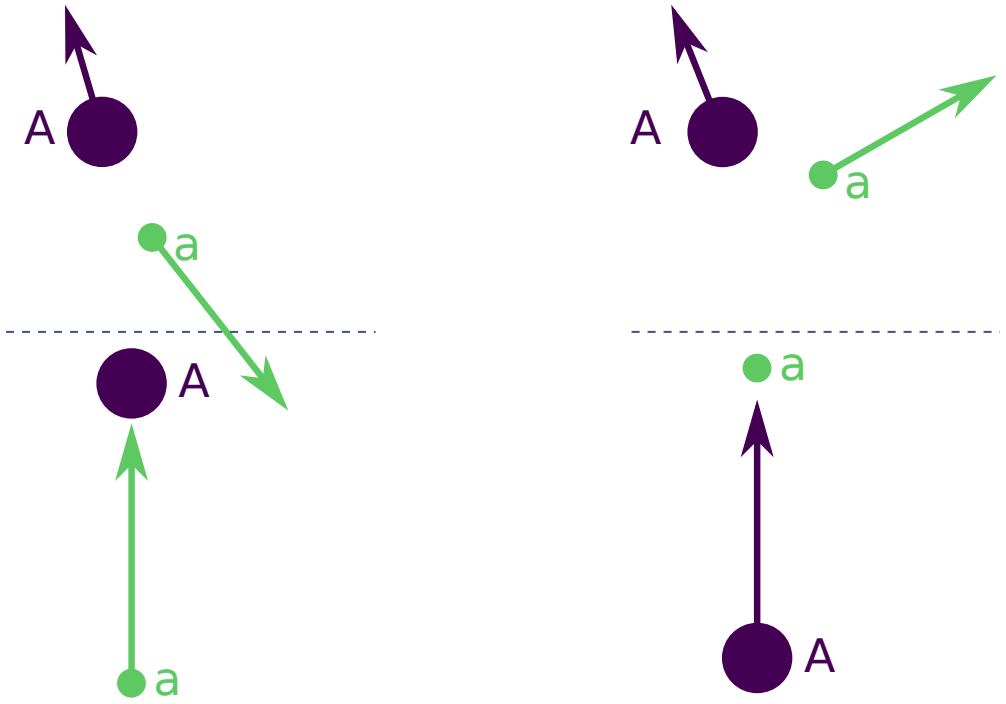
5.4 Thick-Target Inverse Kinematics

Nuclear reactions are typically performed in one of two configurations. In thin-target experiments, a beam of known energy is trained upon a target whose geometry is sufficiently thin that the majority of interactions occur at a similar beam energy. Thick-target experiments dispense with this constraint; the thickness of the target is employed as a degree of freedom. Whilst thin-target experiments pose a much reduced analysis challenge, to perform a measurement that spans a range of energies requires varying the energy of the beam, a slow process that is both costly in time and limited in the number of "stops" (energies) that can be observed. Meanwhile, thick-target experiments circumvent this restriction by using the target itself as both the reaction domain and energy loss medium. In this configuration, the energy of each interaction is given as a function of the interaction location, i.e. the distance travelled within the target by the projectile.

In a system consisting of two different-mass nuclei, thick-target resonant scattering can be performed either in forward

kinematics, or in *inverse* kinematics. In forward kinematics, the beam consists of the lighter nucleus, whilst the heavier nucleus serves as the target. In *inverse kinematics*, the reverse is true: the heavier ion serves as the beam, which is used to bombard a target consisting of the lighter isotope. This approach provides several advantages over forward kinematics:

- From the conservation of momentum, two-body reactions with negative (or zero) Q values are forward focussed (see Fig. 5.5), which improves the geometrical efficiency of the forward-angle detector for a particular solid-angle coverage.
- For short-lived isotopes, it is possible to use a particle accelerator to produce ions on-line to circumvent challenges of radioactive target production. This is particular feasible for alpha-scattering experiments; ^{4}He is one of the lightest nuclei in existence, and therefore most nuclei of interest have larger masses.



(a) Forward kinematics (b) Inverse kinematics

Fig. 5.5: Schematic illustration of a non-relativistic elastic scattering reaction (lab frame) in (a) forward kinematics, in which the lighter nucleus *a* is used as the projectile; and (b) inverse kinematics, whereby the heavy nucleus *A* impinges upon the light target. In each panel, the upper region indicates a potential post-scattering configuration. In the inverse regime, due to conservation of momentum and energy, the reaction is *forward projected*; both nuclei continue travelling forward in the lab frame. Meanwhile, in forward kinematics it is possible for the beam to achieve a (smaller) negative velocity in the lab frame.

Given the clear z dependence of Eqn. 5.11 (see Fig. 5.6), it is evident that thick-target inverse-kinematics configurations are particularly useful in isolating the heavy beam ions from the light target ions; the target dimensions can be configured such that the heavy ions are quickly arrested within the detector volume. From the $\frac{\ln(\frac{E}{I})}{E}$ terms, the stopping power acquires a maximum as a function of energy just prior to the particle coming to rest. When plotted as a function of the path length, the observed peak in the stopping power distribution is often called the Bragg-peak. The thickness of the

target additionally enhances the luminosity; unlike thin-target reactions in which a significant proportion of events will not interact with the target, thick-target reactions span a much wider portion of the reaction cross-section energy domain. In the event that the beam fully stops within the target, the efficiency of the experiment is limited only by secondary interactions, and the efficiency of the detectors.

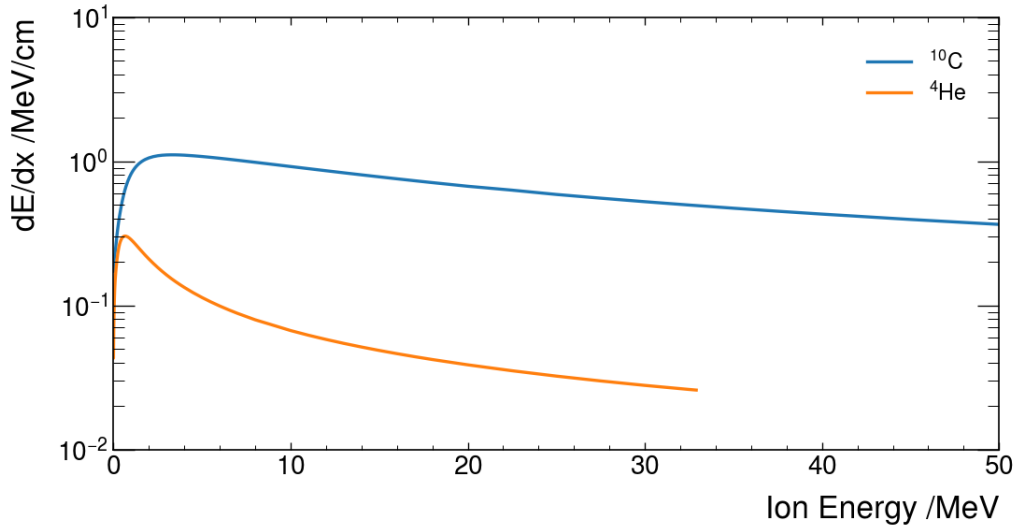


Fig. 5.6: Stopping power of ^{10}C and ^{4}He ions in a gas-mixture of 96% ^{4}He, 4% CO_2 at 405 torr, simulated using SRIM [72].

Unlike thin-target kinematics, reactions that take place within a gas chamber are not confined to a localised region. Therefore, the experimental advantages of operating in a thick-target inverse-kinematic regime can only realised if there are suitable detectors (see Section 6) and analysis techniques (see Section 7) available to recover the reaction kinematics.

5.5 Conclusion

Experimental nuclear physics is predicated in-part upon the ability to inspect the interactions that govern nuclear structure through scattering interactions. Within the domain of nuclear scattering, there are innumerable techniques that are used to scatter and measure two or more systems of nucleons, although they essentially fall into several extremes, indicated in Fig. 5.1. In this thesis, the phenomenon of resonant nuclear elastic scattering is of particular interest, as is discussed in subsequent chapters. In this regime, a compound nucleus is formed, whose structure can be elucidated from the properties of the outgoing products. As such, in this chapter, the conservation rules which govern nuclear interactions, and facilitate their kinematic reconstruction, were discussed.

In order to measure the products of a scattering reaction, an experimentalist has a choice of measurement techniques. A powerful approach that is outlined in this thesis is the time-projection chamber, used in inverse kinematics. The concept of

forward and inverse kinematics was introduced in this chapter, alongside the advantages of operating in inverse kinematics when working with short-lived radioactive reactants. In this chapter, the operating principles of TPCs were established, and the fundamental concepts of gas ionisation and charge multiplication discussed. Subsequent chapters build upon the concepts outlined here, in the description of the form and function of the TexAT TPC given in Section 6.

CHAPTER

SIX

THE TEXAT DETECTOR FOR RARE ISOTOPE BEAM EXPERIMENTS

TexAT (Texas Active Target) is a general-purpose detector developed by the Cyclotron Institute at Texas A&M University for low-energy nuclear physics experiments [73]. As an active-target detector, the fill gas serves as both the detection medium and the target, such that TexAT is well suited to conducting scattering and implantation experiments with rare isotope beams.

TexAT has been upgraded over several experiments. In this section, the initial design of the detector will be outlined, followed by an exploration of the modifications relevant to the experiment discussed later in the chapter.

6.1 Experiment Design

A primary $^{10}B^{3+}$ beam with 7 MeV/u was produced by the K500 cyclotron at the Cyclotron Institute, Texas A & M University [74]. These ions were directed towards a liquid-nitrogen (LN) cooled 9.2 cm gas cell, with 4 μm thick and 19 mm diameter Havar entrance and exit windows [75]. At a pressure of 870 torr, the H_2 gas within the cell was used to facilitate an $^{10}B(p,n)^{10}C$ exchange reaction. The resulting ^{10}C beam had an energy of 32.9 MeV, and an intensity of 16500 particles per second (pps).

6.2 Detector Overview

TexAT is comprised of a gas-filled TPC and MicroMeGaS (Micro-Mesh Gaseous Structure) particle detector, surrounded by an array of telescopes for particle identification and total energy measurement of ions that escape the active volume. A 3D model of these detectors is shown in Fig. 6.1.

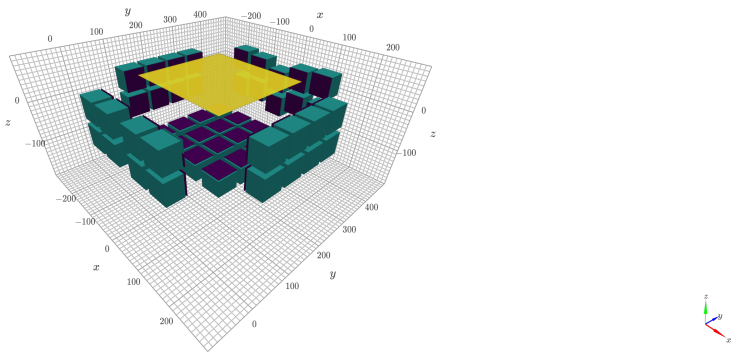


Fig. 6.1: A 3D model of the "final" innermost TexAT detector components, comprised of silicon (purple), thallium doped cesium iodide (turquoise), and MicroMeGaS (yellow) detectors. The ion counter is not pictured here.

To establish a uniform electric field within the ionisation region, a wire field-cage surrounds the active volume. An isographic 3D rendering of the combined detector with the field-cage and MicroMeGaS visible is shown in Fig. 6.2.

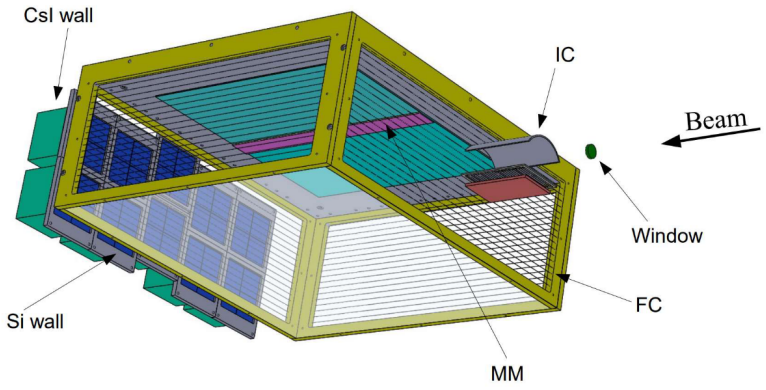


Fig. 6.2: A schematic diagram of the TexAT detector configuration, indicating the location of the various detectors and auxiliary apparatus. A black arrow indicates the direction of the beam, which first enters the detector through the ion-counter (IC, indicated). The segmented anode is visible in blue and purple at the top of the detector, whilst the field-cage (FC) and the array of CsI(Tl) + Si detectors are visible around the outside of the field cage. These detectors occupy the same gas cell as the MicroMeGaS (MM), which sits at the top of the detector. The ion drift plane is given by the normal of the MicroMeGaS. The beam direction, given by a vector between the window and silicon wall, is indicated by a white arrow. Figure taken from Ref. [76].

6.2.1 MicroMeGaS

The MicroMeGaS is a parallel-plate gas particle detector. It comprises of a micro-pattern anode, with a thin metallic mesh suspended at a height of 128 μm above the surface (see Fig. 6.3). This *micro*-mesh partitions the gas into two juxtaposed volumes of differing potentials, forming a large drift region above the mesh, and a small amplification region between the mesh and the anode. In this manner, the detector behaves like an ideal parallel plate detector; the signals formed on the anode readout are predominantly determined by the electron mobility of the gas rather than that of the ions. The irregular pixelation of the anode facilitates high position resolution measurement of these signals. The bias for

the mesh is constant across the detector geometry; individual region biasing is performed by applying different potentials to the corresponding elements of the anode.

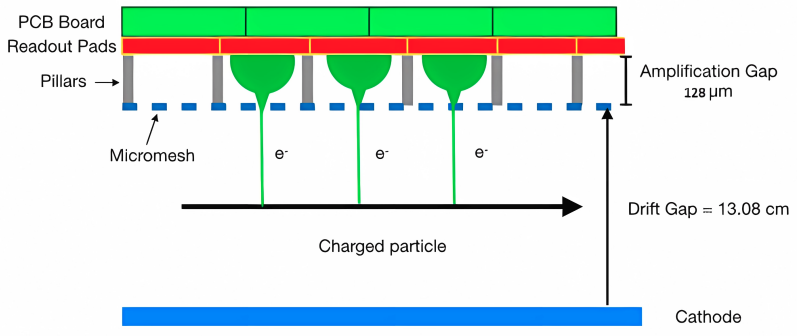


Fig. 6.3: Charge signal formation from charged particle ionisation. A small amplification region, with an electric field of the order ~kV/cm, can be seen between the micromesh and the anode strips. The larger drift/conversion region, with a field strength of ~10^2V/cm, is bounded by the cathode and the micromesh. Gas multiplication predominantly occurs in the amplification region; the drift/conversion region serves only to separate the ion pairs liberated by ionising particle tracks, such as those of scattered alpha particles. Figure adapted from Ref. [77].

The MicroMeGaS anode has an active area of $245 \times 224\,\mathrm{mm}^2$, which is composed of three distinct conductive elements:

- single electrode rows ($112 \times 1.67\,\mathrm{mm}^2$)
- chained electrode columns ($1.67 \times 224\,\mathrm{mm}^2$)
- single electrode pads ($3.42 \times 1.67\,\mathrm{mm}^2$)

These elements are arranged into a three-region configuration: the two side regions are formed by interleaving strips and chains, and the central columnar region is composed from a regular grid of pads (see Fig. 6.4). To ensure that adjacent elements remain electrically isolated from one another, the separation between disjoint strips and chains is 1.75 mm, whilst separate pads are separated by chains is 1.75 mm and 3.5 mm in the vertical and horizontal axes respectively.

The multiplexing scheme present in the side regions is used to reduce the number of channels from 64×128 to $64 + 64$, at a cost of lower resolution, sensitivity, and reconstruction ambiguity. The readout is mapped such that individual zones within the MicroMeGaS can be biased in order to create areas with different gas gains.

In the context of the wider TexAT TPC, the MicroMeGaS provides a readout of the active volume of the TPC chamber. Ionising radiation liberates charge carriers within the conversion region of the MicroMeGaS (i.e. the active volume of the TPC), which separate under the applied electric field (see the field cage in Fig. 6.3). Within the amplification region of the MicroMeGaS, the high-mobility electrons drifting from the conversion region under amplification, forming a proportional distribution of electrons which induce a charge upon the anode (see Fig. 6.4) [67]. This charge is subsequently recorded by the GET electronics to perform reconstruction of the incident particle track.

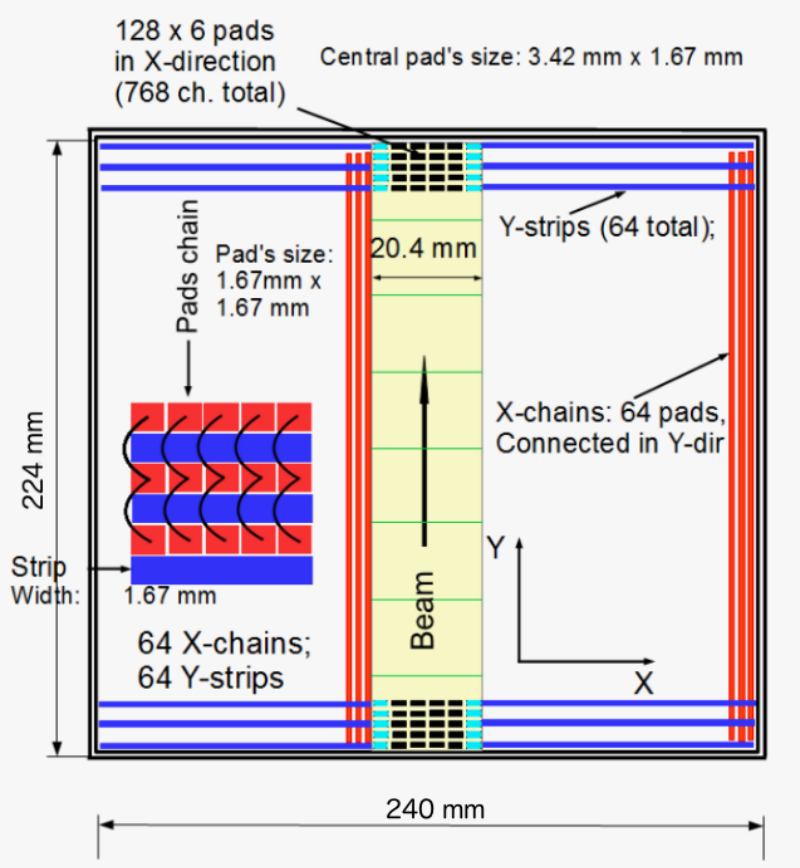


Fig. 6.4: The segmentation plan of the MicroMeGaS anode. There are three distinct regions: a central column of 6×128 pads, and two side regions with 64 interleaved rows (strips) and columns (chains). The strips are formed from single electrodes, whilst the chains are formed from staggered readout pads that are electrically connected together. Note that the measurements in this figure were subsequently revised for fabrication. Figure adapted from Ref. [73].

6.2.2 Silicon

In the TexAT configuration used in the experiment described in this chapter, only the silicon quadrant detectors were present. Without the thallium-doped caesium iodide detectors, only partial energy can be measured for escaping ions that punch through the silicon detectors. In these cases, particle identification must be performed by alternate means.

Two families of silicon detector are used in the TexAT detector:

- MSQ25-1000 (Micron Semiconductor)
- KDP-1K (JSC, "Institute in Physical-Technical Problems", Dubna, Russia)

Both of these detectors are configured with four ($25 \times 25\,\mathrm{mm}^2$) quadrants, as shown in Fig. 6.5. Although the MSQ25-1000 includes an additional rear quadrant ($50 \times 50\,\mathrm{mm}^2$), it was not read-out for this experiment. The thickness of these detectors ranges from 700-1000 μm [75]. Under alpha particles irradiation from an ^{241}Am source, the MSQ25-1000 detectors have been found to exhibit an energy resolution of 50-70 keV whilst the KDP-1K detectors have been observed to possess an improved resolution of 30-50 keV [73].

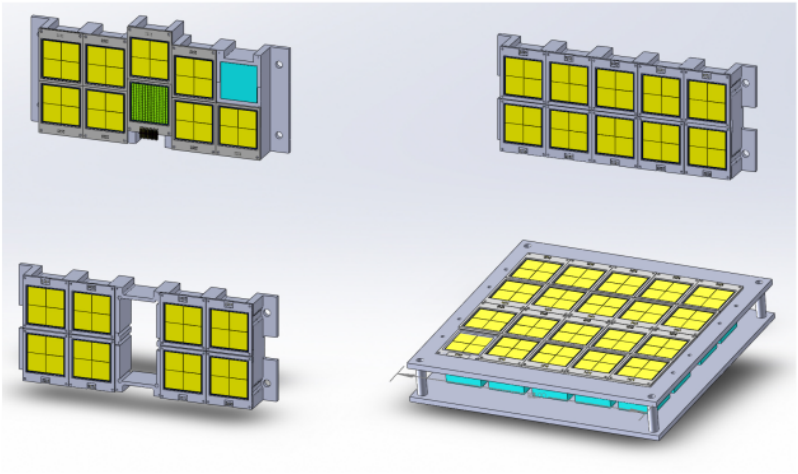


Fig. 6.5: A 3D rendering of the solid-state detector telescopes that surround the TPC: downstream array (top left); upstream array (bottom left); side array (top right); bottom array (bottom right). Figure taken from Ref. [73].

6.2.3 Ion Counter

For this experiment, a planar windowless ionisation chamber (IC) was placed immediately after the entrance window to the gas volume to detect the beam current (see Fig. 6.2). Through charge-time discrimination, the signals of this detector can be used to identify the beam ion, help to eliminate pile-up within the TPC, and trigger the acquisition system in coincidence with the other detectors.

6.3 Readout Electronics

6.3.1 GET Architecture

TexAT uses the General Electronics for TPCs (GET) system for readout of the detectors [78]. GET comprises of a hierarchy of electronic components, starting at the lowest level with the AGET (ASIC, Application-Specific Integrated Circuit, for GET) chip to amplify, shape, and store the signals from each detector (see Fig. 6.6).

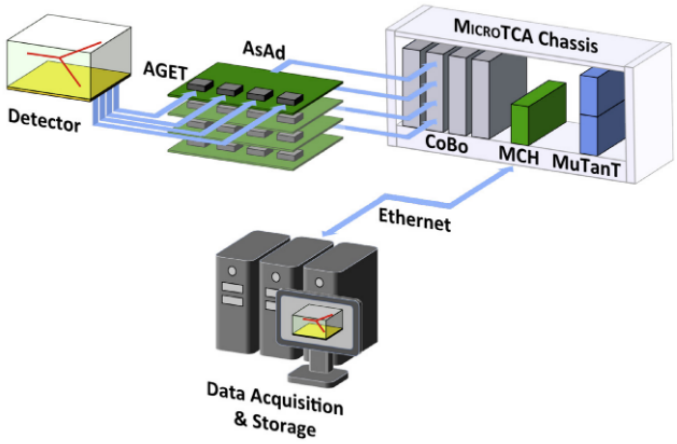


Fig. 6.6: An illustration of the relationship between the experimental detector, readout electronics, and data acquisition and storage equipment in an experiment that uses the GET system [78].

Each channel possesses an independent trigger and continuously samples the filtered analogue signal from the shaper, with a sampling frequency between 1 and 100 MHz, using a 512-cell switch-capacitor circular buffer [78]. For this experiment, the readout sampling frequency was set to 25 MHz. A single AGET chip has 64 independent channels in addition to 4 floating channels that provide a measure of the electronic noise, known as fixed-pattern noise (FPN) channels. These FPN channels are distributed across the AGET chip in order to place one FPN channel at the edge of each chip [78].

To read out all of the channels in a TexAT detector requires 24 AGET chips. There are 4 AGET chips on each AsAd (ASIC and Analog to Digital converter) board, which digitises the signals from the SCA memory of the AGET chips using a 12-bit ADC [78]. The data from these AsAd boards are collected by a series of CoBo's (Concentration Boards), which timestamp the information and transmit it to the storage infrastructure.

In order to synchronize the CoBos and generate a global trigger, an additional board called the MuTAnT (MUltiplicity Trigger ANd Time) is used. In this context, a *global* trigger describes a single trigger event derived from multiple CoBo boards. The MuTAnT supports multiple kinds (levels) of trigger:

- Level 0: an external trigger, e.g. from an ancillary detector
- Level 1: a "multiplicity" trigger which reduces the triggers from each CoBo board into a single trigger.

- Level 2: a pattern-based trigger that matches the 64-bit hit-pattern from each AGET

In this experiment, a Level 1 silicon-only trigger was used.

6.3.2 Signal Shaping

The internal pre-amplification stage (CSA) in each AGET chip can be bypassed in cases where external pre-amplifiers need to be read-out by the GET electronics [78]. In order to optimise the signal processing of the detector (bypassing the GET filtering stage) the IC is connected via a bypass circuit (see Fig. 6.7) to an external MESYTEC shaper (MSCF-16). The signals from this bypass circuit are fed back into the Gain-2 stage of AGET chip, as discussed in [73].

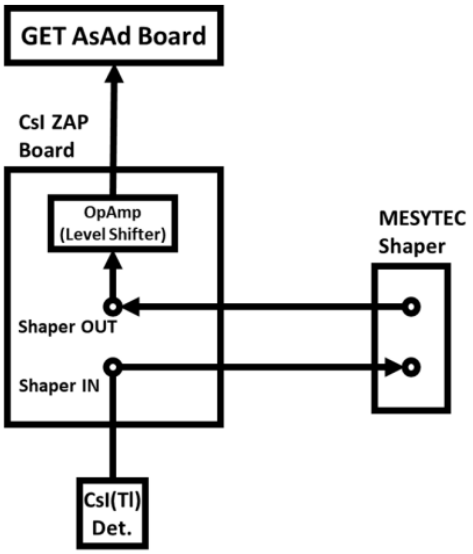


Fig. 6.7: A schematic diagram of the shaper bypass facility of AGET chips. External shapers are used for the ion counter signals (see above). Figure adapted from Ref. [73].

6.4 Conclusion

The TexAT detector is a novel TPC detector developed by the Cyclotron Institute at Texas A&M University. In this chapter, an overview of the function of this detector has been given, with a focus upon the high-level view of the detector. Previous chapters, such as Section 5, discuss some of the mechanisms by which the TPC functions. Following this discussion, the distinguishing characteristics of this TPC, namely the GET electronic readout and pixelated MicroMeGaS anode, have been established. In particular, the multiplexing scheme used to reduce the multiplicity of the MicroMeGaS readout is referenced, which is of particular importance in subsequent chapters concerning strip-chain reconstruction ambiguity (see Section 8.10).

The properties of the component detectors, such as the silicon detector array, have been outlined. This includes the thickness and energy resolution of these detectors. In addition to the detector characteristics, the experimental parameters were introduced; notably the use of a radioactive beamline of ^{10}C, with a gas target of $^{4}\text{He} + \text{CO}_2$. The information outlined here is of import to the later discussion of reconstructing and analysing the data from this experiment, given in Section 8. These parameters govern the phase space of the elastic scattering reaction that this experiment was designed to perform, and generally must be well chosen to ensure the physics of interest is observed. The conclusion in Section 8 refers to this constraint, whereby it is suggested that the experiment parameters were suboptimal for zero-degree scattering reconstruction.

CHAPTER

SEVEN

TRACK FINDING AND FITTING

Fundamental to the design of many particle physics experiments is the ability to detect reaction ejectiles and measure their physical properties. The considerable challenge that this poses, in even the most conventional forward-kinematics fixed-target experiments, is exacerbated in the thick-target regime described in Section 5. In the absence of a single interaction locus, it becomes of vital importance to be able to observe the *trajectories* of reaction products such that the underlying kinematic variables can be reconstructed. As discussed in Section 5.3, the TPC detectors in which these experiments are often conducted typically incorporate a segmented readout, whereby particle trajectories are described as a set of discretised samples. The resolution of these data is usually governed by a combination of the ejectile, gas, and readout properties.

The problems of identifying and fitting the particle trajectories recorded in such data are nominally distinct concepts, although many approaches necessarily combine the two. Track *finding* describes the task of identifying the existence of a particle track, whilst track *fitting* refers to the process of ascribing model variables to identified tracks subject to the parametrisation. Central to both subproblems is the importance of outlier rejection, which is often used as a metric by which to evaluate algorithm performance.

7.1 Conventional Techniques

7.1.1 Hough Transform

The *classical* Hough transform is a technique for estimating the parameters and multiplicity of a model family within a dataset. In this context, a *model* is a description of a set of data, such as a line or a plane. The Hough transform employs a scheme whereby *features* in the dataset, i.e. points along a line, vote for models with which they are compatible (see Fig. 7.1). For any continuous parametrisation, a single observation is a member of an infinite set of models. In order to make this tractable, the Hough transform is conventionally performed by discretising the n-dimensional parameter space

of the model. Thereafter, for each observation, the set of compatible models can be determined from the permutations of the discrete $n-1$-dimensional free parameter space.

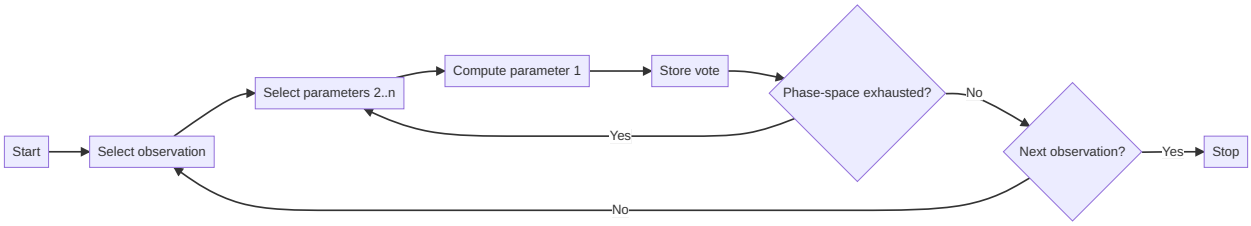


Fig. 7.1: Flowchart detailing the classic discrete Hough transform. Once the algorithm has terminated, maxima in the voting space correspond to models (parameterisations) with high support (large numbers of votes). These maxima can be treated as candidates that well describe the data.

This technique introduces an explicit relationship between the precision of the identified model parameters and the robustness to noise; fine discretisations are able to accurately determine the underlying model parameters, but the introduction of noise quickly leads to the identification of multiple models. Meanwhile, coarse discretisations are far less vulnerable to over-fitting, but at a cost of poor parameter resolution. This trade-off requires the user to balance the anticipated level of noise (signal-to-noise ratio) and the similarity of observations (how close any-two separate models are in parameter space. Furthermore, the search complexity increases exponentially with the number of model parameters.

Certain procedures can be used to mitigate this compromise:

- Bin smoothing: convolving the parameter space with a smoothing kernel to reduce the significance of spurious peaks
- Secondary optimisation: by asserting that identified models are close to their "true" representations, a secondary optimiser (e.g. least squares) can be used to improve the model parameters.

7.1.2 RANSAC

Random Sample Consensus (RANSAC) is a paradigm for fitting a singular model to experimental data. It recognises the model finding/fitting problem as a connected one; the challenge of identifying a unique set of models within a dataset is highly dependent upon estimating their free parameters given the observed data [79].

RANSAC draws a distinction between two kinds of error:

- *measurement* error: deviations from the "true" values observed in measurements
- *classification* error: gross errors that arise from misclassification of the dataset

The RANSAC procedure accounts for these two kinds of error by starting with a small, feasible, initial dataset, before enlarging this set with *consistent* data (see Fig. 7.2). Concretely, when fitting e.g a line within a set of N points, the

RANSAC approach would be to select two points, compute origin and direction vectors, and then determine the points which lie within an acceptance window (orthogonal distance) to the model. These new *inliers* can then be used to re-estimate the model; smoothing the measurement errors.

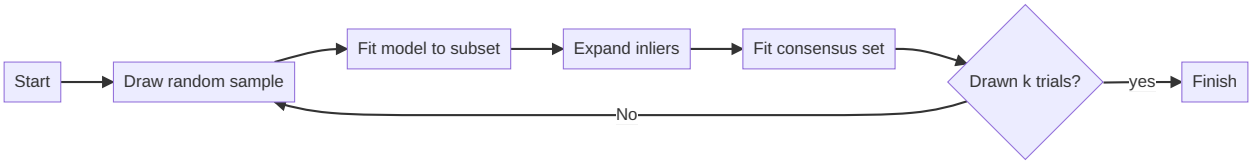


Fig. 7.2: Flowchart outlining the RANSAC algorithm.

A number of parameters determine the termination of the algorithm, namely:

- the acceptance window that separates inliers from outliers
- the minimum support required to instantiate a model (the number of inliers required to consider the model *valid*)
- the number of (random) subsets to draw in the search for model candidates.

The upper bound on the number of trials required to select a "good" subset from the dataset can be determined mathematically. If w is the probability that a selected data point is within the error tolerance of a given model, then the expectation value $E(K)$ of the number of trials k required to select a subset of n points is given by

$$\begin{aligned} E(k) &= b + 2(1-b)b + 3(1-b)^2 b + \cdots + i(1-b)^{i-1} b \ldots \,, \\ &= b[1 + 2a + 3a^2 + \cdots + ia^{i-1} + \ldots] \,, \end{aligned} \tag{7.1}$$

where $b = w^n$ and $a = (1-b)$.

It can be shown that the above simplifies to

$$E(k) = \frac{1}{b} = w^{-n} \tag{7.2}$$

(see [79]).

Although there is no analytic method for determining the minimum support required to instantiate a model, there exist methods by which to determine sensible values of this parameter for a given problem.

Sequential RANSAC

The RANSAC algorithm is designed to find the best-fit model given a dataset. In scenarios where there may be multiple models, RANSAC must be generalised to identify a family of consensus sets. A common approach is to use *Sequential* RANSAC, which iteratively applies the RANSAC algorithm to the dataset until no further models can be identified with sufficient support.

The RANSAC algorithm establishes a hard distinction between inliers and outliers. An important consequence of this behaviour is that the method is hierarchical: the first "winning" model has the greatest number of data points to choose from, whilst the final model has the least. As such, for models with overlapping observations, RANSAC can perform poorly. Two instances of ideal simulated track scattering reactions are shown in Fig. 7.3 and Fig. 7.5, whilst a reaction for which sequential RANSAC is ill-suited is shown in Fig. 7.7.

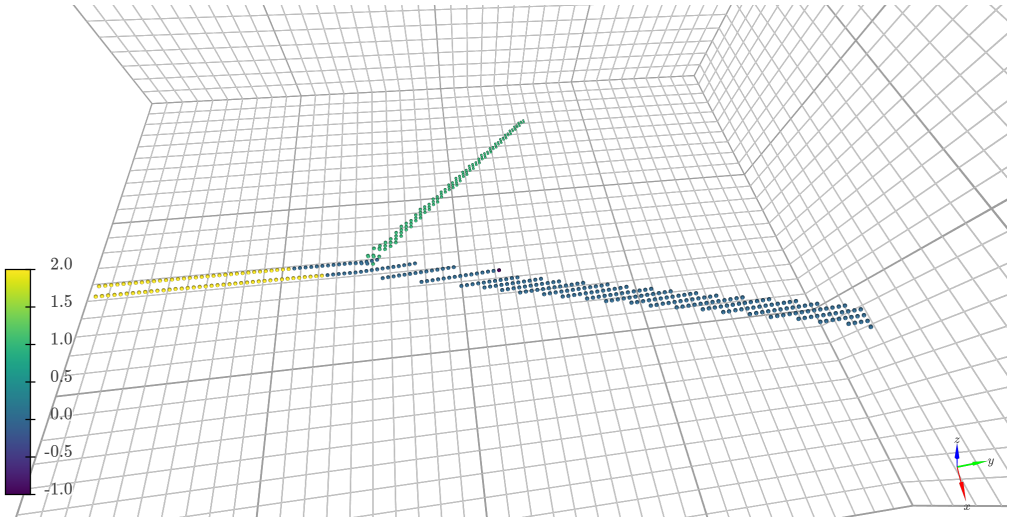


Fig. 7.3: 3D scatter plot of a simulated elastic scattering reaction between a ^{10}C beam and a ^{4}He gas target. The simulation parameters match the experimental parameters described in Section 8. The simulation was performed using a heavily modified variant of TexATSim, the software developed by the group at Texaa A&M for TexAT reaction studies. The beam direction is given by the $\hat{y}$ axis. Three track labels given by sequential RANSAC are indicated by the colour map, with a fourth outlier label corresponding to -1. It can clearly be seen that the scattered beam (blue) labelling is over-supported, as it extends into the incident beam (yellow) track clusters.

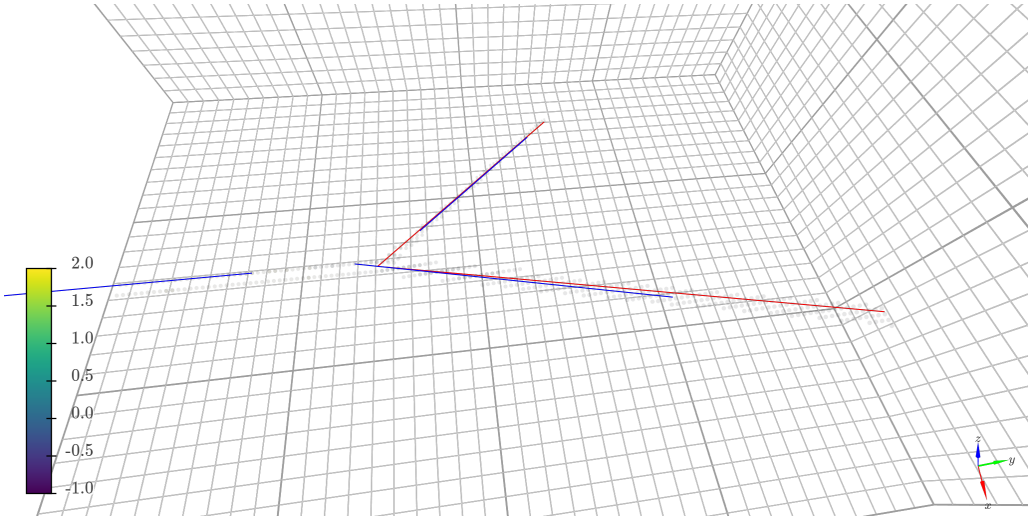


Fig. 7.4: 3D visualisation of the tracks reconstructed from Fig. 7.3 using sequential RANSAC. Only the track direction should be considered; these fits correspond to unbounded lines. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. It can be seen that whilst the scattered beam track lies close to the simulated track, a common-vertex fit will pull the vertex away from its true value.

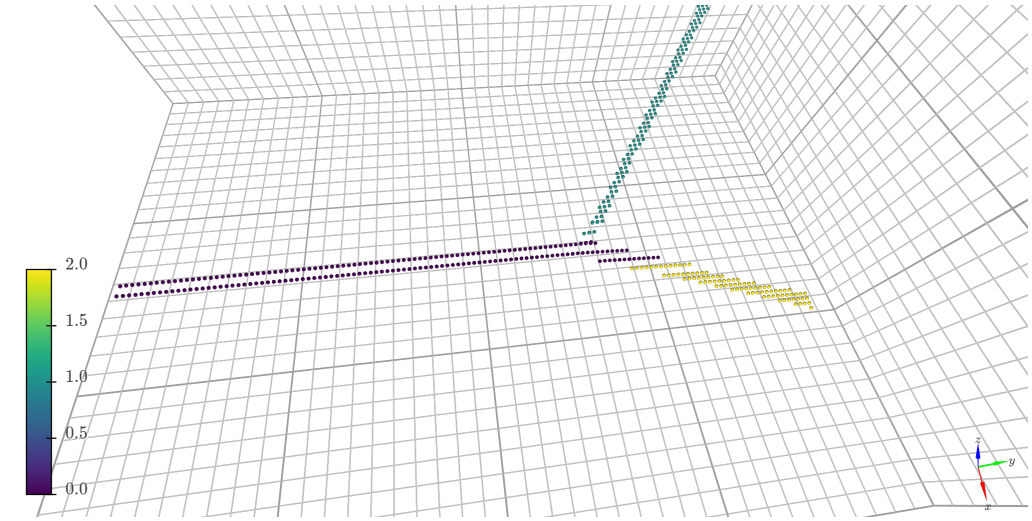


Fig. 7.5: 3D scatter plot of a second simulated elastic scattering reaction between a ^{10}C beam and a ^{4}He gas target. The simulation parameters match the experimental parameters described in Section 8. The beam direction is given by the $\hat{y}$ axis. Three track labels are indicated by the colour map, with a fourth outlier label corresponding to -1. Like Fig. 7.3, the RANSAC fit is greedy, but in this figure it can be seen that the indicent beam (purple) is over-supported and extends into the scattered beam (yellow) track clusters.

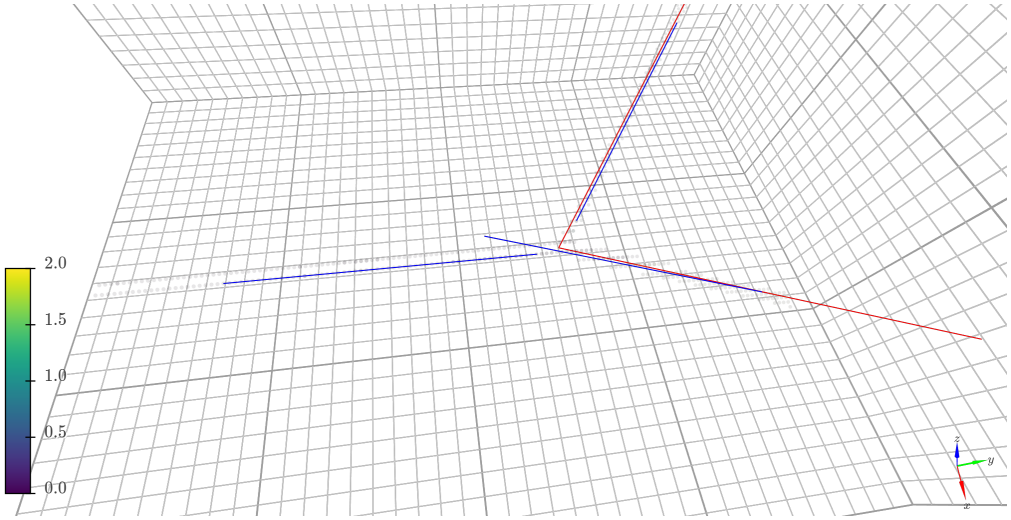


Fig. 7.6: 3D visualisation of the tracks reconstructed from Fig. 7.5 using sequential RANSAC. Only the track direction should be considered; these fits correspond to unbounded lines. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. Like Fig. 7.4, it can be seen that whilst the scattered beam track lies close to the simulated track, a common-vertex fit will pull the vertex away from its true value.

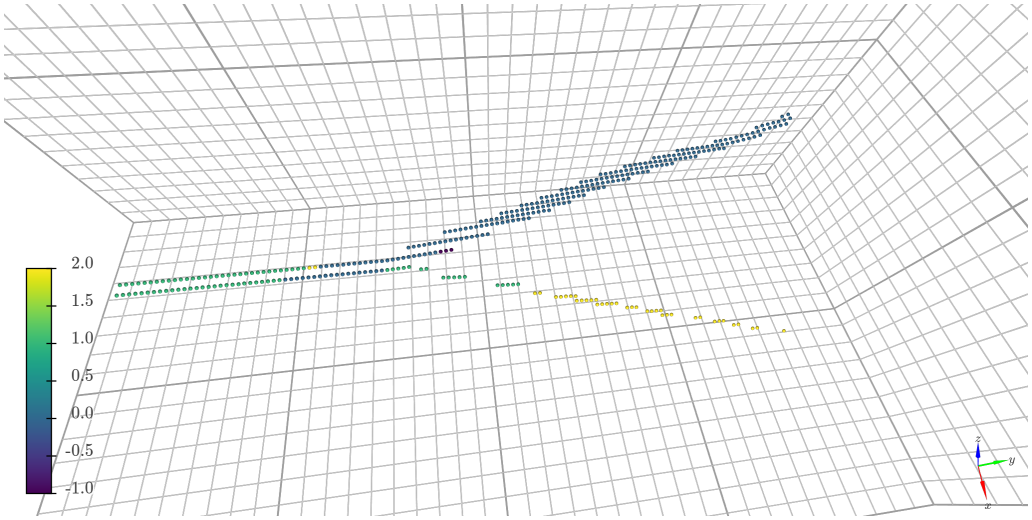


Fig. 7.7: 3D scatter plot of a third simulated elastic scattering reaction between a ^{10}C beam and a ^{4}He gas target. The simulation parameters match the experimental parameters described in Section 8. The beam direction is given by the $\hat{y}$ axis. Three track labels are indicated by the colour map, with a fourth outlier label corresponding to -1. Whilst the correct number of active tracks (3) have been identified, it can be seen that the light-product track (blue) extends far into the incident beam clusters (green), and the incident beam clusters extend into the scattered beam region (yellow). The beam track is thus discontinuous.

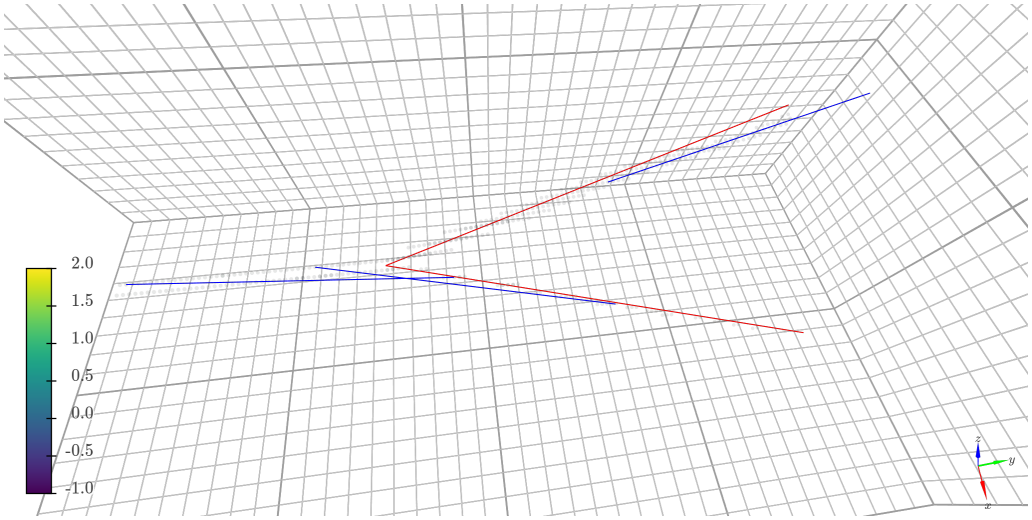


Fig. 7.8: 3D visualisation of the tracks reconstructed from Fig. 7.7 using sequential RANSAC. Only the track direction should be considered; these fits correspond to unbounded lines. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. It can be seen that the poor labelling shown in Fig. 7.7 produces a poor fit, with the reconstructed vertex (given by the intersection of the blue scattered and heavy beam tracks) displaced outside the measured clusters.

7.2 Facility Location Problem

The previously explored methods of the Hough transform and RANSAC are both *greedy* methods; the order in which solutions are generated favours those models which are found first. The consequence of greedy optimisation is that the global solution across several models may not be optimal.

In order to obtain a *global* solution to track labelling, the task must be reframed as a global optimisation problem. The facility location problem (FLP) is an optimisation problem whose solution determines the optimal location for a set of facilities such that the transportation costs for each facility are minimised. There are several formulations of the problem with auxillary constraints, such as the *capacity* of the facilities, or the addition of an *opening cost* for each facility.

The significance of "facility" and "transportation cost" depends upon the framing of the problem. Here, the *uncapacitated facility location problem* (UFLP) will be explored in the context of particle track labelling.

7.2.1 Uncapacitated Facility Location Problem

The uncapacitated facility location problem is a variant of the FLP in which each facility has an unlimited capacity and an associated opening cost. In prose, the problem may be defined by three requirements:

- A manufacturing company wants to minimise costs.
- Delivery is expensive, so facilities need to be close to consumers.

- Facilities are also expensive, so as few as possible should be opened.

The solution to this problem is given by the minimisation of the following cost function:

$$E(F) = \sum_{f \in F} \left(O_f + \sum_{c \in C} T_{fc} D_{fc} \right) , \tag{7.3}$$

where

- F is the set of open facilities.
- O_f is the cost of opening facility f.
- T_{fc} is the cost of transporting goods between facility f and consumer c.
- D_{fc} is the demand of consumer c satisfied by facility f.

In the domain of track labelling, the terms “facility”, “distance”, and “consumer”, acquire new meaning:

Term	Interpretation
Facility	A hypothetical track model
Consumer	An observation that belongs to an unknown track model
Transport cost	The goodness-of-fit for an observation given a particular track model

Table 7.1: Terms corresponding to the FLP interpretation of track labelling.

Meanwhile, the consumer demand D_{fc} is simply the delta function of the labelling $\delta_{ff(c)}$, i.e. a single observation c is served by exactly one facility $f(c)$. See Fig. 7.9 for an illustration of these components.

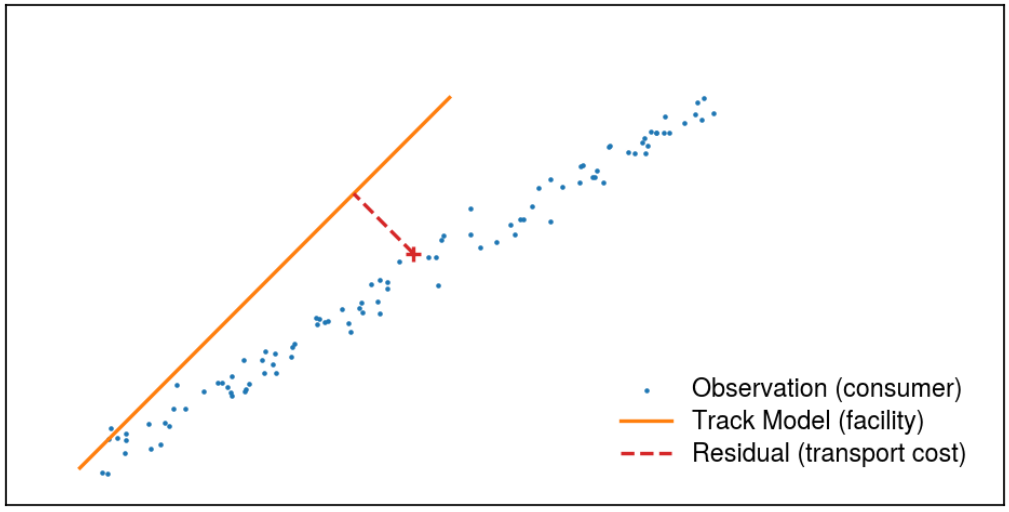


Fig. 7.9: An illustration of a 2D track described by a proposed track model.

The probability of observing set C of observations for facility f is the product of the probabilities of observing each c, i.e.

$$P(f \mid C) = \prod_{c} P(f \mid c)\,. \tag{7.4}$$

If the per-consumer transport cost is given as a log-likelihood, then it follows that the sum over the transport costs for each facility is equivalent to the log-likelihood of the facility given C:

$$\begin{aligned} \sum_{c} \ln\left(P(f \mid c)\right) &= \ln\left(\prod_{c} P(f \mid c)\right) \\ &= \ln\left(P(f \mid C)\right)\,. \end{aligned} \tag{7.5}$$

7.2.2 Metric Labelling Problem

The trivial form of Eqn. 7.3, in which the facility cost O_f is zero

$$E(F) = \sum_{f \in F,\, c \in C} T_{fc} D_{fc}\,, \tag{7.6}$$

is ill-posed. A *well-posed* problem is guaranteed to have a *unique* solution that depends continuously on the input data [80].

In order to solve Eqn. 7.6, the set of admissible solutions must be restricted by the introduction of *a priori* knowledge, i.e. the problem must be *regularised*. The label cost term given in Eqn. 7.3 is one such regulariser which imposes a preference for solutions with fewer labels.

A common feature of the labelling problem (of which the track labelling problem is a derivative) is that coherent groups of observations are typically known *a priori* to be positively correlated.

This behaviour can be accounted for by an additional regulariser in the energy function

$$\sum_{cc' \in C} V_{f(c),f(c')}(c, c')\,, \tag{7.7}$$

where $V_{f(c),f(c')}(c, c')$ is the pairwise prior that penalises $f(c) \neq f(c')$, such that

$$E(F) = \underbrace{\sum_{f \in F} O_f}_{\text{label cost}} + \underbrace{\sum_{f \in F,\, c \in C} T_{fc} D_{fc}}_{\text{data cost}} + \underbrace{\sum_{cc' \in C} V_{f(c),f(c')}(c, c')}_{\text{smooth cost}}\,. \tag{7.8}$$

A two dimensional model potential is shown in Fig. 7.10.

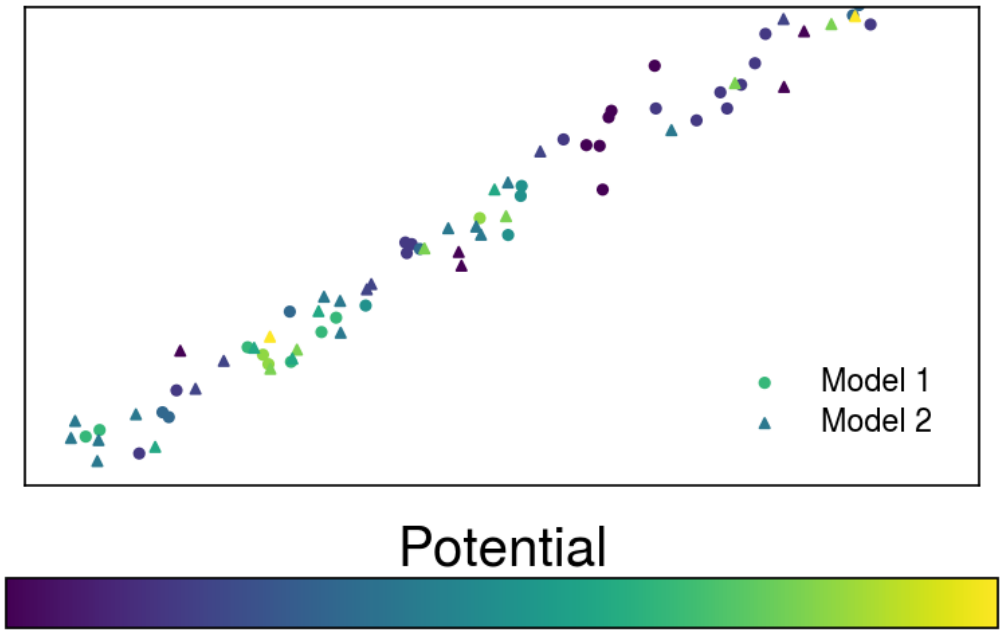


Fig. 7.10: Illustration of a 2D smoothness prior between two labellings. The highest costs are observed between unalike labels at small distances.

Clearly, the *label cost* leads to minimisation of the model count, whilst the *smooth cost* leads to the preference of spatially coherent labellings. As before, the *data cost* prefers models which are well-described by their inliers.

7.2.3 Graph Cut Optimisation

In the absence of smooth costs, the UFLP can be solved by integer programming techniques, but remains NP-hard [81]. There exists an heuristic method which solves Eqn. 7.3 in $O(|F|^2|C|)$ time [82], where $O(n)$ represents an asymptotic time complexity that is linear in n. Meanwhile, the task of minimising Eqn. 7.8 is also NP-hard for $|F| \geq 3$ [83], but there exist approximate solutions using *graph cuts* that are guaranteed to find the local minima within a fixed bound of the global optimum [84].

Graph Cuts

One family of approximations to the solution of Eqn. 7.8 are based upon the method of *graph cuts*. Here, the fundamentals of the α-expansion method will be discussed, with the treatment of label costs and $\alpha\beta$-swap moves omitted for brevity.

Without label costs, Eqn. 7.8 becomes

$$E(F) = \underbrace{\sum_{f \in F,\, c \in C} T_{fc} D_{fc}}_{\text{data cost}} + \underbrace{\sum_{cc' \in C} V_{f(c),f(c')}(c, c')}_{\text{smooth cost}} . \tag{7.9}$$

A *graph* $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is fundamentally a set of vertices $\mathcal{V}$ and a set of pairwise edges $\mathcal{E}$ between them. Each edge may have an associated weight, and may further be directed such that $A \rightarrow B$ is distinct from $A \leftarrow B$.

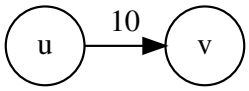


Fig. 7.11: Simple two-vertex graph.

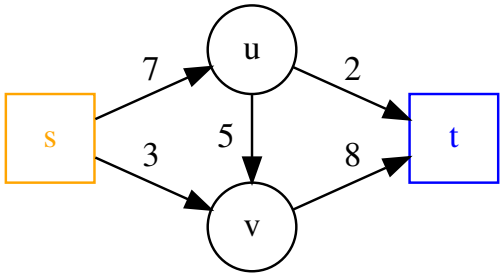


Fig. 7.12: Simple two-vertex graph with terminals s, t. The maximum flow between the terminals is 10.

A *cut* $\mathcal{C} \subset \mathcal{E}$ is a set of edges which partitions the graph into two disjoint subsets, separating the terminals. Importantly, the *cost* $|\mathcal{C}|$ equals the sum of the edge weights. Furthermore, no proper subset of $\mathcal{C}$ may also be a cut.

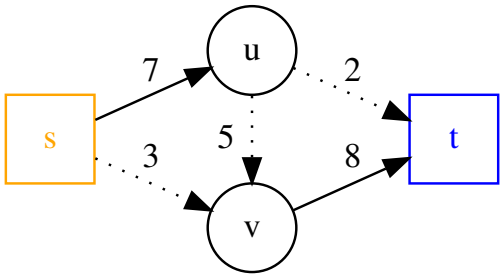


Fig. 7.13: Simple two-vertex graph with terminals s, t. A valid cut is indicated by dotted lines.

Under the *min-cut max-flow* theorem , the *minimum*-cut $\mathcal{C}^\star$ (cut for which $|\mathcal{C}|$ is minimal) is equivalent to the maximum-*flow*, which can be computed in an efficient manner [85] [86].

Alpha Expansion

The α-expansion algorithm is a method of finding the labelling F such that the potential Eqn. 7.9 is minimised. It has been extended to solve Eqn. 7.8 through the treatment of label costs. The method is based upon the observation that, if Eqn. 7.9 is encoded in terms of a particular graph construction G_α, the minimum cut corresponds to the optimal (local) labelling. The choice of this graph, and the mechanism by which the optimal labelling is determined by the minimum cut, will be outlined briefly below.

Consider an example comprised of two nodes (observations) Ω and A. One can introduce two labels: α and ω. The true labelling is given in Table 7.2: the Ω is labelled by ω, and A is labelled by α.

Node	Label
A	α
Ω	ω

Table 7.2: Lowest cost (correct) labelling for a two-node system

The optimal labelling given in (Table 7.2) should have a cost that evaluates to

$$E(F) = D(A \cdot \alpha) + D(\Omega \cdot \omega) + V(\Omega \cdot \omega, A \cdot \alpha)\,, \tag{7.10}$$

i.e. the sum of

- the data cost for A under the α model
- the data cost for Ω under the ω model
- the smoothness cost between Ω under the ω model, and A under the α model.

The graph for this two-node system whose minimum cut yields Eqn. 7.10 is shown in Fig. 7.14.

Clearly, the cut illustrated in Fig. 7.14 yields the optimal cost. However, how does this cut invoke a (new) labelling? One implication of the property

> no proper subset of $\mathcal{C}$ may also be a cut

is that exactly *one* edge between the terminals ($\{\,\omega, \alpha\,\}$) and any node can be cut [83]. This naturally defines a correspondence between a cut and a labelling: the cutting of a terminal edge assigns the corresponding label to the node. In this case, there is a cut through $\alpha - A$ and $\omega - \Omega$, which assigns label α to A and Ω to ω respectively.

Although it is already known that this labelling is the correct (optimal) one, in general one might not have this prior information. The *expansion-move* algorithm can be used to find an improved labelling F'_α through the expansion of a

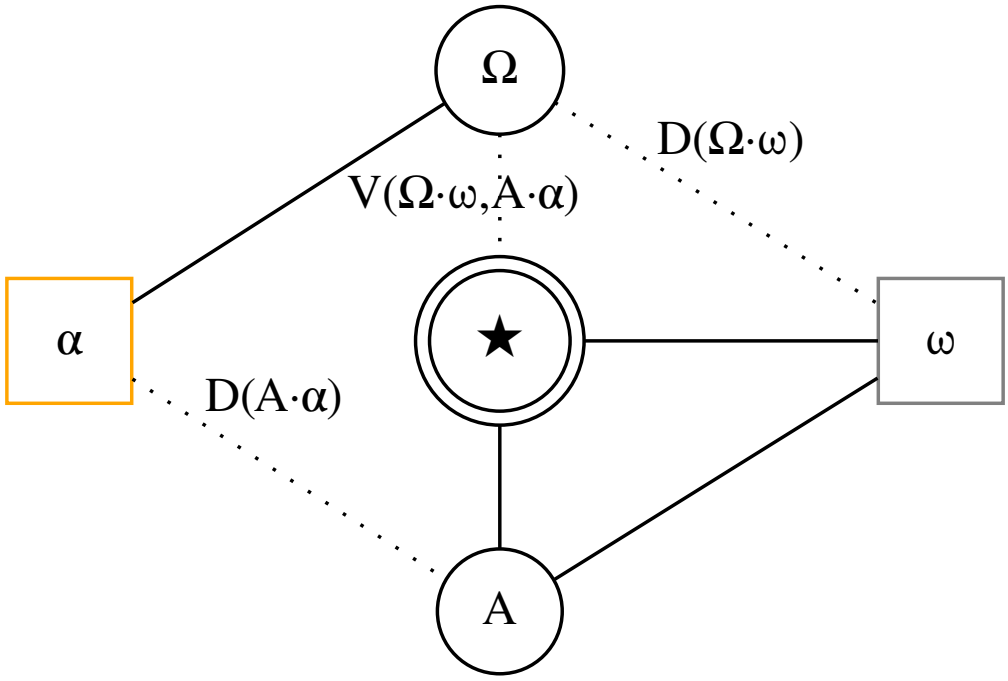


Fig. 7.14: Graph representing a two-node system whose optimal labelling is given in Table 7.2. By choosing appropriate edge weights, the cost of the minimum cut is equal to the optimal labelling cost given in Eqn. 7.10. Cut links are indicated with dotted lines. Smoothness costs that penalise localised discontinuities in the labelling are introduced through the addition *auxillary nodes* at label boundaries, i.e. the $\star$ node between Ω and A. The edges between these auxiliary nodes and their neighbours encode the smoothness penalty.

particular label α. If the expansion does not modify the labelling, then it is optimal i.e. the current labelling F is a (local) minimum. The *expansion* of a label α refers to an equivalence class of cuts that preserve the edges between the alternate label and each node. For the two-node system shown in Fig. 7.14, the *expansion* property of the expansion-move algorithm is introduced by setting the cost of the edges between the terminal ω and the nodes A, Ω to infinity. Edges with infinite costs will belong to any minimum cut, ensuring that the ω label does not grow (see Fig. 7.15).

This simple model only considers two labels ($|F| = 2$), whilst in practice $|F| \gg 2$. To generalise the expansion of label α to $|F| > 2$, ω can be replaced with a *meta* label $\overline{\alpha}$ that represents not-α. Thereafter, the procedure remains identical. In order to determine the optimum expansion-move, this procedure must be repeated for each label.

The α-expansion algorithm is therefore

1. Build a graph $\mathcal{G}_\alpha(\mathcal{V}_\alpha, \mathcal{E}_\alpha)$ composed of a set of vertices $\mathcal{V}_\alpha$ and a set of pairwise edges $\mathcal{E}_\alpha$.
2. Add a pair of additional *terminal* vertices α and $\overline{\alpha}$ that represent the α label and the set of complement labels $\overline{\alpha} = F - \{\,\alpha\,\}$.
3. Establish a *weight* for each edge in $\mathcal{E}_\alpha$.

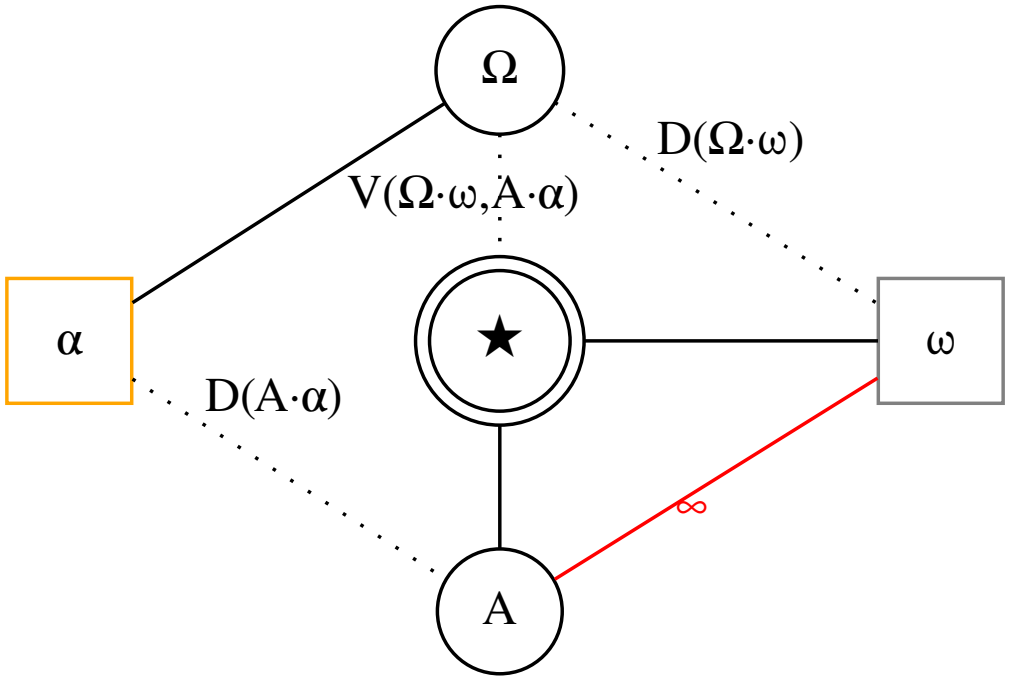


Fig. 7.15: Variation of Fig. 7.14 depicting a two-node system whose labelling is already optimal. As such, the weight between A and the terminal ω is set to infinity, so that the ω label does not grow i.e. the α label does not shrink.

4. Determine the minimum cut of $\mathcal{G}(\mathcal{V}, \mathcal{E})$ to propose a new labelling.

5. Accept the new labelling if the cost of the cut, which is equivalent to Eqn. 7.9, is lower than the existing solution.

6. Go to 1. for the next label.

The choice of weights is the fundamental mechanism by which the expansion-move algorithm operates. The rules for finding these values are given in [84]. In addition, it can be shown that this algorithm can be generalised to account for label costs in order to solve Eqn. 7.8 [84].

A condition of the α-expansion method is that the interaction potential between labels be *metric*, i.e. it must satisfy [83]

$$\begin{aligned} V_{\alpha,\beta} &\leq V_{\alpha,\gamma} + V_{\gamma,\beta} \\ V_{\alpha,\beta} &= V_{\beta,\alpha} \geq 0\,. \end{aligned} \tag{7.11}$$

The choice of a $V_{\alpha,\gamma} = \delta_{\alpha\gamma}$ interaction potential (the Potts model) combined with a geometric distance meets this requirement [87].

7.2.4 PeARL

So far, the application of graph-cuts to solving the metric labelling problem has been explored. It is not immediately obvious how this approach maps onto the problem of track fitting. At first glance, the model f seen in Eqn. 7.8 resembles the label described in Section 7.2.3. However, our track models are members of $\mathcal{L} = \mathbb{R}^n$, and therefore a one-to-one mapping between F and $\mathcal{L}$ would be infinitely large. In order to explore the continuum of model parameters whilst working within a feasible subset of the parameter space, the PeARL algorithm may be used [83].



Fig. 7.16: Process diagram of the PeARL algorithm

A process diagram for the PeARL algorithm is shown in Fig. 7.16, and is comprised of the following steps:

1. *Propose* an initial set of models $\mathcal{L}_0$ by random sampling from a set of observations
2. Compute the optimal labelling F_t from $\mathcal{L}_t$ by *expanding* each label $f \in F$
3. *Re-estimate* the model parameters of $\mathcal{L}_t$ using F_t to find $\mathcal{L}_{t+1}$; go to (2.)

The proposal step (1) of the PeARL algorithm closely resembles that of RANSAC. Unlike RANSAC, however, PeARL evaluates these proposals against a single optimisation objective, rather than hierarchically.

7.2.5 Track Fitting

Unbounded Line

Conventional track fitting approaches often model a track as an unbounded line in three dimensional space. Such a model has an non-normalisable probability density; naturally over an infinite interval, the probability of observing any particular value approaches 0. As such, it is not appropriate (or indeed possible) to use the log-likelihood in Eqn. 7.5 to describe the data cost required by Eqn. 7.8. Instead, a simple orthogonal displacement cost function, analogous to the total least-squares cost function, can be used

$$T_{fc} = k \frac{\left\| (\vec{r}(c) - \vec{o}(f)) \wedge \hat{\vec{n}}(f) \right\|^2}{\sigma(f)^2}, \tag{7.12}$$

where $\vec{r}(c)$ is the position of observation c, $\vec{o}(f)$ a point belonging to model f, $\hat{\vec{n}}(f)$ the direction of model f, $\sigma(f)$ the width of the transverse distribution, and k a scaling factor. This corresponds to the residual shown in Fig. 7.9.

The unbounded line may not be the most suitable model to describe linear data. Crucially, it has no concept of regularity; there is no parameterisation by which to describe the distribution along the linear axis.

Line Segment

Another model applicable to track fitting is to treat a set of observations as a line *interval*. With a finite length, the model has a normalisable probability density, and as such the log likelihood can be used to determine the data cost.

A curve can be defined as an infinite sum of point distributions, where the probability $p(\vec{\mathbf{r}})$ of observing $\vec{\mathbf{r}}$ is

$$p(\vec{\mathbf{r}}) = \int_L f(\vec{\mathbf{r}}; \vec{\mu}(t), \vec{\Sigma})\, \mathrm{d}t \ . \tag{7.13}$$

The PDF for a non-degenerate multivariate normal distribution is

$$f(\vec{\mathbf{r}}) = \frac{\exp\left(-\frac{1}{2}(\vec{\mathbf{r}} - \vec{\mu})^{\mathrm{T}} \vec{\Sigma}^{-1} (\vec{\mathbf{r}} - \vec{\mu})\right)}{\sqrt{(2\pi)^k \left|\vec{\Sigma}\right|}} \ , \tag{7.14}$$

for $\vec{\mathbf{r}} \in \mathbb{R}^k$, where $\vec{\Sigma}$ is a symmetric matrix. For a line interval of normally distributed points, bounded between $t \in [0, 1]$, Eqn. 7.13 expands to

$$p(\vec{\mathbf{r}}) = \int_0^1 \frac{\exp\left(-\frac{1}{2}(\vec{\mathbf{r}} - \vec{\mu}(t))^{\mathrm{T}} \vec{\Sigma}^{-1} (\vec{\mathbf{r}} - \vec{\mu}(t))\right)}{\sqrt{(2\pi)^k \left|\vec{\Sigma}\right|}} \, \mathrm{d}t \, . \tag{7.15}$$

With considerations for symmetry, it can be shown that the log of the integral in Eqn. 7.15 reduces to

$$p(\vec{\mathbf{r}}) = \underbrace{\log\left(\frac{1}{4\pi L \sigma^2}\right)}_{\text{length penalty}} + \underbrace{\log\left(\operatorname{erfc}\left[\frac{1}{\sqrt{2}\sigma}\left(\left|\Delta\vec{\mathbf{r}} \cdot \hat{\mathbf{n}} - \frac{L}{2}\right| - \frac{L}{2}\right)\right]\right)}_{\text{end-point penalty}} + \underbrace{\frac{1}{2\sigma^2}\left[\left(\Delta\vec{\mathbf{r}} \cdot \hat{\mathbf{n}}\right)^2 - \Delta\vec{\mathbf{r}} \cdot \Delta\vec{\mathbf{r}}\right]}_{\text{orthogonal distance penalty}} , \tag{7.16}$$

where L is the length of the line, $\hat{\mathbf{n}}$ is the unit direction of the line, and $\Delta\vec{\mathbf{r}}$ the relative position of $\vec{\mathbf{r}}$ with respect to the line origin.

A visualisation of this likelihood function can be see in Fig. 7.17.

From the same dataset used to produce Fig. 7.3 (and associated figures), a series of track fits using the PeARL algorithm (using the line interval model given by Eqn. 7.16) were produced. Whilst Fig. 7.18 and Fig. 7.20 offer moderate improvements in track labelling over their respective RANSAC results in Fig. 7.3 and Fig. 7.5, these represent ideal behaviour of the RANSAC algorithm. Under narrower track angles, the sequential RANSAC approach given in Section 7.1.2 suffers



Fig. 7.17: 3D volume plot of the likelihood function for a point described by a line interval between $\begin{pmatrix} 4 & 4 & 4 \end{pmatrix}$ and $\begin{pmatrix} 16 & 16 & 16 \end{pmatrix}$ with $\sigma = 1.0$.

from poor fits due to earlier track fits intersecting with separate tracks. Such a pathological worst-case is depicted in Fig. 7.7, with the improved PeARL result shown in Fig. 7.22. In this case it is both the choice of bounded line intervals *and* a global fitting approach that regularises the result.

Of fundamental importance to any track-fitting approach is the treatment of each track on an equal basis, and establishing the concept of a global (total) fit. The PeARL method provides a mechanism for this, and the relative advantages can be seen in particular in Fig. 7.22. A particular challenge of using PeARL with line *intervals* (as defined above) is the increase in dimensionality of the space from whic PeARL samples candidate models. Unlike unbounded lines, line intervals do not well describe points that lie beyond the interval bounds (indeed, unbounded lines have infinite bounds). As such, the problem is more sensitive to the set of available models, as prior samples are poorly incentivised to expand their inlier set. Further work could be performed here to investigate the prior distribution from which candidate models are sampled, such as the use of scattering-vertex seeding. The use of unbounded lines with PeARL is also feasible, and retains many of the advantages of the PeARL algorithm over sequential RANSAC. However, the problem of small-angle scattering is less tractable through unbounded lines, due to the strong incentive to grow the inlier set without impedence due to a track length penalty.

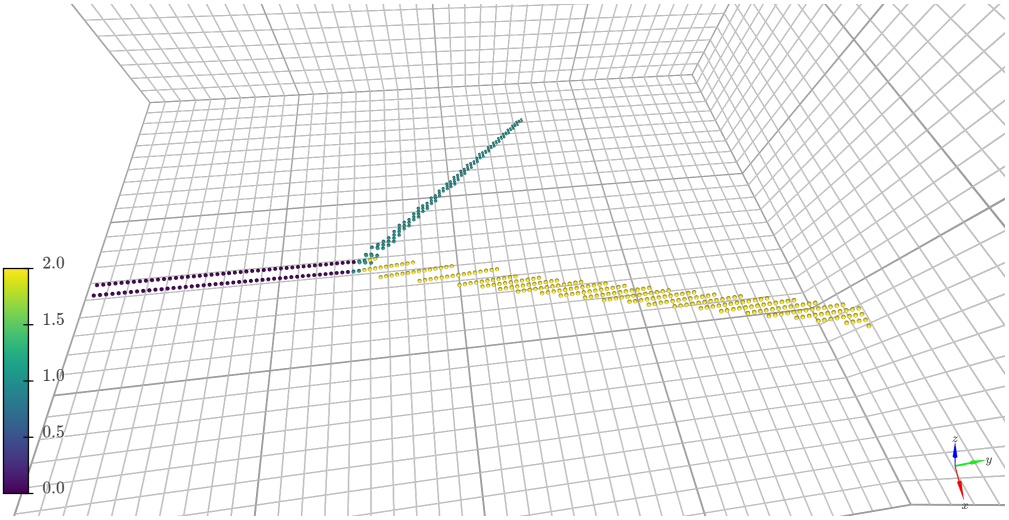


Fig. 7.18: 3D scatter plot of the simulated elastic scattering reaction shown in Fig. 7.3. Three track labels given by the PeARL algorithm are indicated by the colour map. In comparison to Fig. 7.3, it can be seen that the PeARL fit (with line intervals) preserves a common vertex through distinct track separation.

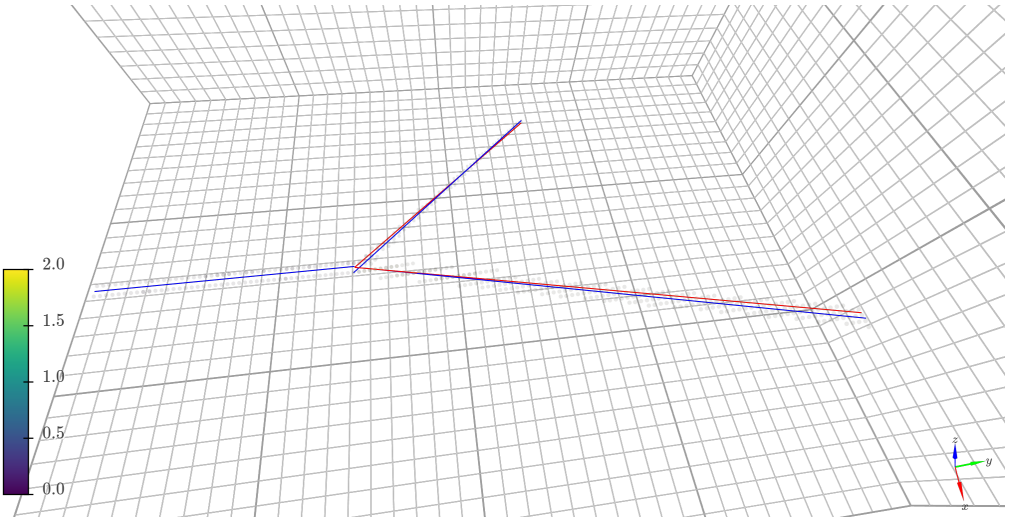


Fig. 7.19: 3D visualisation of the tracks reconstructed from Fig. 7.18 using PeARL. Both the track direction and end-point should be considered; these fits correspond to bounded line intervals. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. For the same reaction as depicted in Fig. 7.4, it can be seen that the PeARL fit (with line intervals) better preserves a common vertex through distinct track separation. The line-interval parameterisation makes it possible to clearly identify the estimated vertex from the end-point of the incident beam track.

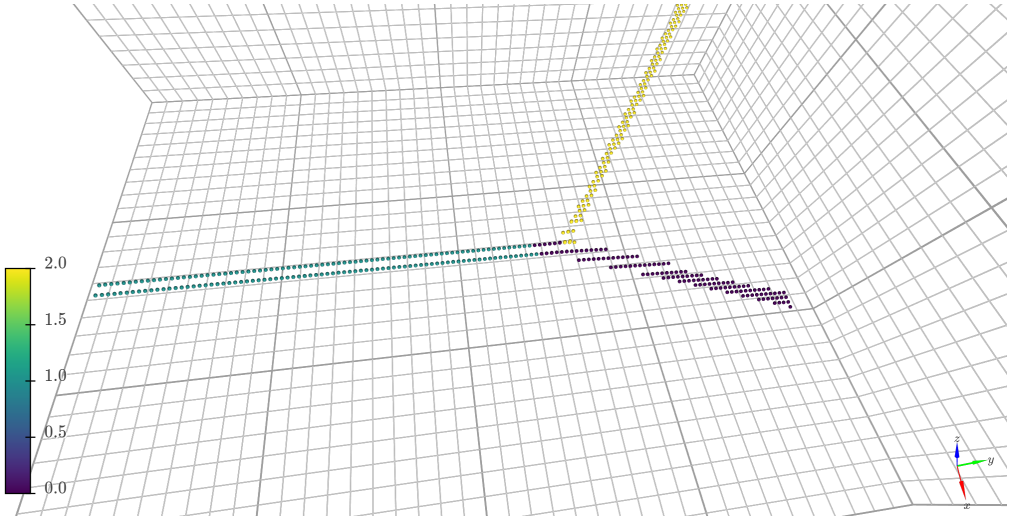


Fig. 7.20: 3D scatter plot of the simulated elastic scattering reaction shown in Fig. 7.5. Three track labels given by the PeARL algorithm are indicated by the colour map. In comparison to Fig. 7.5, it can be seen that the PeARL fit (with line intervals) avoids overfitting of the incident beam (blue) track, although a reduced degree of overfitting is observed in the scattered beam (purple) region.

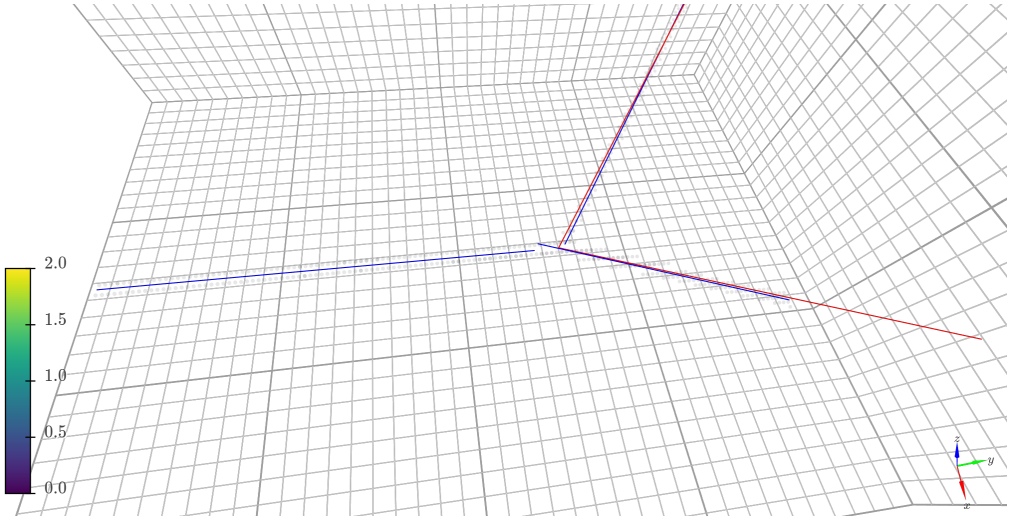


Fig. 7.21: 3D visualisation of the tracks reconstructed from Fig. 7.20 using PeARL. Both the track direction and end-point should be considered; these fits correspond to bounded line intervals. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. For the same reaction as depicted in Fig. 7.6, it can be seen that the PeARL fit (with line intervals) better reconstructs the angle of each track in agreement with the simulation. A common vertex is much more easily identified by the intersection of the reconstructed tracks.

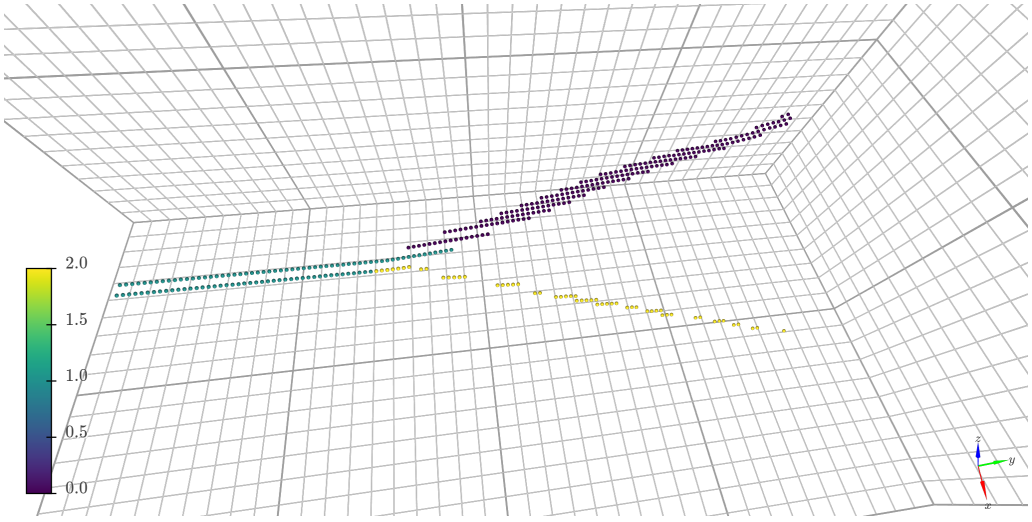


Fig. 7.22: 3D scatter plot of the simulated elastic scattering reaction shown in Fig. 7.7. Three track labels given by the PeARL algorithm are indicated by the colour map. In comparison to Fig. 7.7, it can be seen that the PeARL fit (with line intervals) avoids the discontinuity in the reconstructed incident beam (blue) track. As such, the labelling of the scattered beam (yellow) track has more support. The degree of overfitting in the light product (purple) track is also much diminished, and all clusters are considered inliers.

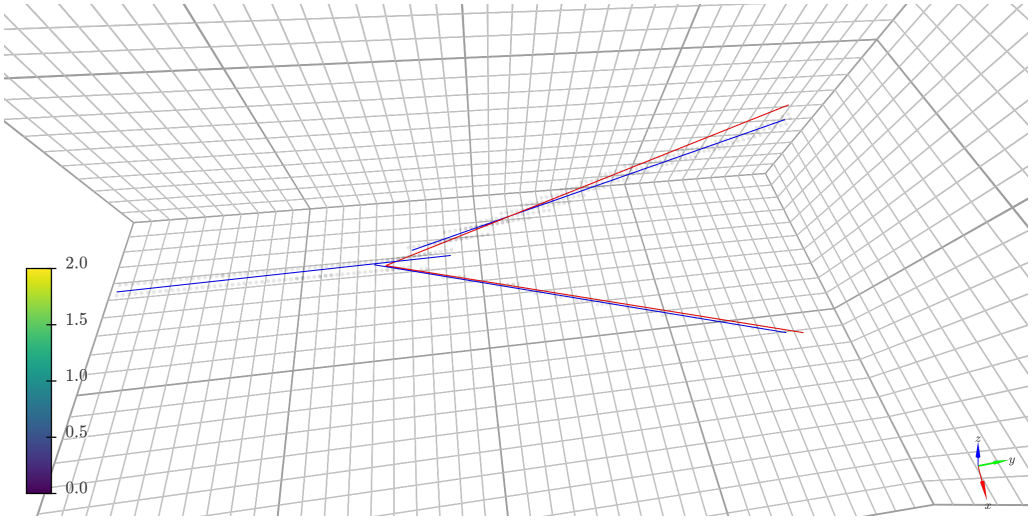


Fig. 7.23: 3D visualisation of the tracks reconstructed from Fig. 7.22 using PeARL. Both the track direction and end-point should be considered; these fits correspond to bounded line intervals. The simulated (red) tracks and reconstructed (blue) tracks are superimposed over a semi-transparent point-cloud of the measured track clusters. For the same reaction as depicted in Fig. 7.8, it can be seen that the PeARL fit (with line intervals) produces a far more reasonable reconstruction of the track angles, with the closest-point-of-approach between any two tracks producing a good starting-point for vertex estimation. As evidenced by its length, the incident beam track is slightly overfit. A bigger discrepancy lies in the light-product track, which is visually distinct from the simulated track. Here it can be presumed that secondary scattering in the simulation leads to a bending in the track, which lies below the resolution of the MicroMeGaS. As such, only a single track is observed here.

7.3 Kinematic Fitting

Many classes of problem, including those seen in track fitting, are amenable through least squares optimisation. In the context of track fitting, it is often the case that the solution is required to satisfy a set of constraints. These constraints ensure that determined solution is physically admissible.

The general case of least squares estimation with constraints is solved by linearising the constraint equations, and solving the problem iteratively. Let $\vec{\eta} = \{\eta_1, \eta_2, \dots, \eta_n\}$ be a vector of N observables, whose measured values are $\vec{y}$, with errors given by the covariance matrix $V(\vec{y})$. From the Least-Squares Principle [88], the objective function χ^2 is given by

$$\chi^2 = (\vec{y} - \vec{\eta})^T V(\vec{y})^{-1} (\vec{y} - \vec{\eta}) \,. \tag{7.17}$$

It follows that the best estimates of $\vec{\eta}$ satisfy

$$\nabla_{\vec{\eta}} \chi^2(\vec{\eta}) = \vec{0} \,. \tag{7.18}$$

Frequently one is interested in a solution to Eqn. 7.18, subject to constraints. Consider a single constraint g, for which $S = \{\vec{r} : g(\vec{r}) = 0\}$. The minimum for f subject to g is the point $\vec{r}_1 \in S$ at which any small displacement along g yields

no change in f (to first order). If this were not the case, there would exist another point $\vec{r}_2 \in S$ for which $f(\vec{r}_2) < f(\vec{r}_1)$. It follows that level curve $f(\vec{r}) = k$ must be parallel to g at $\vec{r}_1$, which is equivalently stated as

$$\nabla f = \lambda \nabla g \,, \tag{7.19}$$

where λ is a constant (see Fig. 7.24).

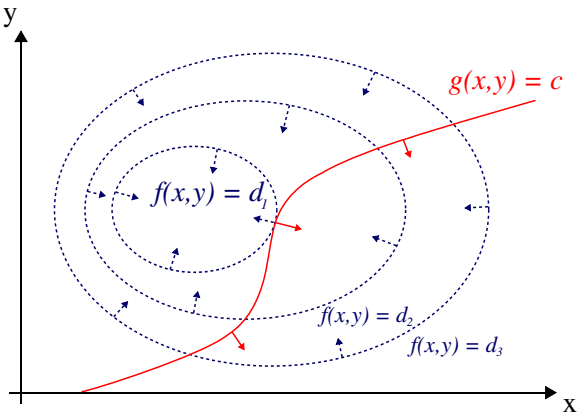


Fig. 7.24: Diagram illustrating a constrained optimisation problem, in which the constraint (red curve) touches an equipotential contour (blue curve) of the function being optimised. By definition, the constrained solution occur where the two curves are touching, in order for the solution to be a minimum.

Instead of a single constraint $g(\vec{r})$, one might have K constraint equations,

$$g_k\left(\vec{\eta}, \vec{\xi}\right) = 0 \,, \qquad\qquad k = 1, 2, \ldots, K \,. \tag{7.20}$$

where $\vec{\xi}$ is a set of unmeasured model variables $\vec{\xi} = \{\xi_1, \xi_2, \ldots, \xi_J\}$.The solution to Eqn. 7.18 is then given by reformulating the problem in terms of the Lagrangian

$$\mathcal{L} = \chi^2 + \sum_k \lambda_k g_k \,, \tag{7.21}$$

and solving $\nabla_{\lambda,\vec{r}} \mathcal{L} = 0$ for $\vec{r}$.

Our original problem in Eqn. 7.17 may be rewritten using Lagrange multipliers, such that one minimises

$$\chi^2\left(\vec{\eta}, \vec{\xi}, \vec{\lambda}\right) = \left(\vec{y} - \vec{\eta}\right)^T V(\vec{y})^{-1} \left(\vec{y} - \vec{\eta}\right) + 2\vec{\lambda}^T \vec{g}\left(\vec{\eta}, \vec{\xi}\right) . \tag{7.22}$$

Setting the gradient of χ^2 to zero gives

$$\begin{aligned}\nabla_{\vec{\eta}}\chi^2 &= 2V^{-1}(\vec{y}-\vec{\eta}) + 2G_{\vec{\eta}}^T\vec{\lambda} = \vec{0}\\ \nabla_{\vec{\xi}}\chi^2 &= 2G_{\vec{\xi}}^T\vec{\lambda} = \vec{0}\\ \nabla_{\vec{\lambda}}\chi^2 &= 2\vec{g}\left(\vec{\eta},\vec{\xi}\right) = \vec{0}\,,\end{aligned} \tag{7.23}$$

where the matrices $G_{\vec{\mu}}$ are defined as

$$\left(G_{\vec{\mu}}\right)_{ki} = \frac{\partial g_k}{\partial \mu_i}\,. \tag{7.24}$$

It follows that to minimise the objective function, one must solve

$$V^{-1}(\vec{\eta}-\vec{y}) + G_{\vec{\eta}}^T\vec{\lambda} = \vec{0}\,, \tag{7.25}$$

$$G_{\vec{\xi}}^T\vec{\lambda} = \vec{0}\,, \tag{7.26}$$

and

$$\vec{g}\left(\vec{\eta},\vec{\xi}\right) = \vec{0}\,. \tag{7.27}$$

Taylor expanding $g(\vec{\eta},\vec{\xi})$ to first order about the point $\vec{\eta}^v,\vec{\xi}^v$ gives

$$\vec{g}\left(\vec{\eta},\vec{\xi}\right) = \vec{g}\left(\vec{\eta}^v,\vec{\xi}^v\right) + G_{\vec{\eta}}^v(\vec{\eta}-\vec{\eta}^v) + G_{\vec{\xi}}^v\left(\vec{\xi}-\vec{\xi}^v\right). \tag{7.28}$$

Solving Eqn. 7.25 for $\vec{\eta}$ and setting Eqn. 7.28 to zero, it follows that

$$\vec{g}^v + G_{\vec{\eta}}^v\left(\vec{y} - V\left(G_{\vec{\eta}}^v\right)^T\vec{\lambda} - \vec{\eta}^v\right) + G_{\vec{\xi}}^v\left(\vec{\xi}-\vec{\xi}^v\right) = \vec{0}\,. \tag{7.29}$$

By the introduction of

$$\begin{aligned}\vec{r} &= \vec{g}^v + G_{\vec{\eta}}^v(\vec{y}-\vec{\eta}^v)\\ S &= G_{\vec{\eta}}^v V\left(G_{\vec{\eta}}^v\right)^T,\end{aligned} \tag{7.30}$$

Eqn. 7.29 becomes

$$\vec{r} - S\vec{\lambda} + G_{\vec{\xi}}^v\left(\vec{\xi}-\vec{\xi}^v\right) = \vec{0}\,. \tag{7.31}$$

Left multiplying by S^{-1} one obtains

$$\vec{\lambda} = S^{-1}\left(\vec{r} + G^v_{\vec{\xi}}\left(\vec{\xi} - \vec{\xi}^v\right)\right). \tag{7.32}$$

Eqn. 7.32 may then be substituted into Eqn. 7.26 to give

$$\left(G^v_{\vec{\xi}}\right)^T S^{-1}\left(\vec{r} + G^v_{\vec{\xi}}\left(\vec{\xi} - \vec{\xi}^v\right)\right) = \vec{0}\,, \tag{7.33}$$

which may be solved for $\vec{\xi}$:

$$\vec{\xi} = \vec{\xi}^v - \left(\left(G^v_{\vec{\xi}}\right)^T S^{-1} G^v_{\vec{\xi}}\right)^{-1}\left(G^v_{\vec{\xi}}\right)^T S^{-1}\vec{r}\,. \tag{7.34}$$

Finally, Eqn. 7.25 can be rearranged to give

$$\vec{\eta} = \vec{y} - V\left(G^v_{\vec{\eta}}\right)^T\vec{\lambda} \tag{7.35}$$

The solutions to the system are obtained by repeated application of the update equations as follows:

- Determine $\vec{\xi}^{i+1}$ from $\vec{r}^i$ and S^i using Eqn. 7.34.
- Find $\vec{\lambda}^{i+1}$ from $\vec{\xi}^{i+1}$ using Eqn. 7.32.
- Use $\vec{\lambda}^{i+1}$ to find $\vec{\eta}^{i+1}$ using Eqn. 7.35.

In this manner, the unknown variables $\vec{\xi}$ are initially computed, then the Lagrangian multipliers found, and finally the estimates $\vec{\eta}$ of the measured variables $\vec{y}$ are updated. This process repeats until the χ^2 reaches a minimum, e.g. it becomes stable within successive iterations to some tolerance,

$$\left|\frac{\left(\chi^2\right)^{i+1} - \left(\chi^2\right)^{i+1}}{\left(\chi^2\right)^i}\right| < \epsilon\,. \tag{7.36}$$

7.4 Conclusion

The fundamental challenge involved with reconstructing the information captured within a TPC is the ability to identify and fit the kinematics of the tracks left behind by the movement of ionising particles. This is a space of on-going research, with classical statistical techniques such as the Hough transform forming the basis of many analyses. In recent times, the RANSAC method has found popularity in its ability to overcome the limitations of the Hough transform, which suffers from poor performance in the phase space of poor track resolution and similar track properties.

In this chapter, these two methods have been introduced and discussed, with reference to their limitations as track-finding approaches. The RANSAC method was explored in greater detail, with an application to scattering reactions simulated by the TexATSim simulation package. This package was developed by the Texas A&M group, and further enhanced during the work on this thesis, in order to improve the user-experience and the behavior of the cluster resolution fit model. From these simulations, it was shown that the sequential RANSAC technique suffers from poor performance in a number of scenarios, which represent the limits of the greedy fitting approach.

Subsequently, the reframing of the track fitting & finding task as an Uncapacitated Facility Location Problem was made, in order to explore the set of similar problems outside of nuclear physics that have already been solved. Through an introduction to the method of Graph Cuts, a novel track fitting approach that leverages the PEaRL algorithm was introduced. In support of this novel approach, the cost function of a line-interval model was derived to provide the ability to fit finite-length tracks. This new track fitting approach using line intervals was later applied to the simulation reactions shown in Section 7.1.2, and the improved reconstruction performance emphasised as a function of the global, interval based fitting technique.

Finally, a discussion of the theory behind kinematic fitting (constrained optimisation) was made, in order to support the next steps in the analysis workflow described in Section 8.

CHAPTER

EIGHT

RESONANT ELASTIC SCATTERING OF α-PARTICLES WITH ^{10}C

This chapter outlines an analysis workflow that has been developed to reconstruct experimental data taken from an experiment conducted at Texas A&M University in 2018 using the TeXAT detector at the Cyclotron Institute. For much of this analysis work, there were no public references for best-practice or existing analysis techniques beyond high-level discussions in a preliminary paper published by the physics group at Texas A&M University. As such, all methods described here are original work, although the group at Texas A&M were helpful in resolving particular lines of questioning concerning detector function and other analysis techniques being developed. Reconstruction of the GET response function, and its use in signal fitting; track fitting using the PeARL method; and ringing elimination may be considered to be particularly novel contributions. Application of kinematic fitting to the reconstructed kinematics was also a novel contribution to this kind of analysis, although a summary of the results from this work is not included following a refocussing on the challenges associated with signal ringing in the detector. The development of this analysis using modern high-energy physics packages, and the Dask parallel computing library was also novel, although discussion of this work was omitted here for brevity.

The analysis process described in this chapter was performed in a series of discrete stages (see Fig. 8.1).

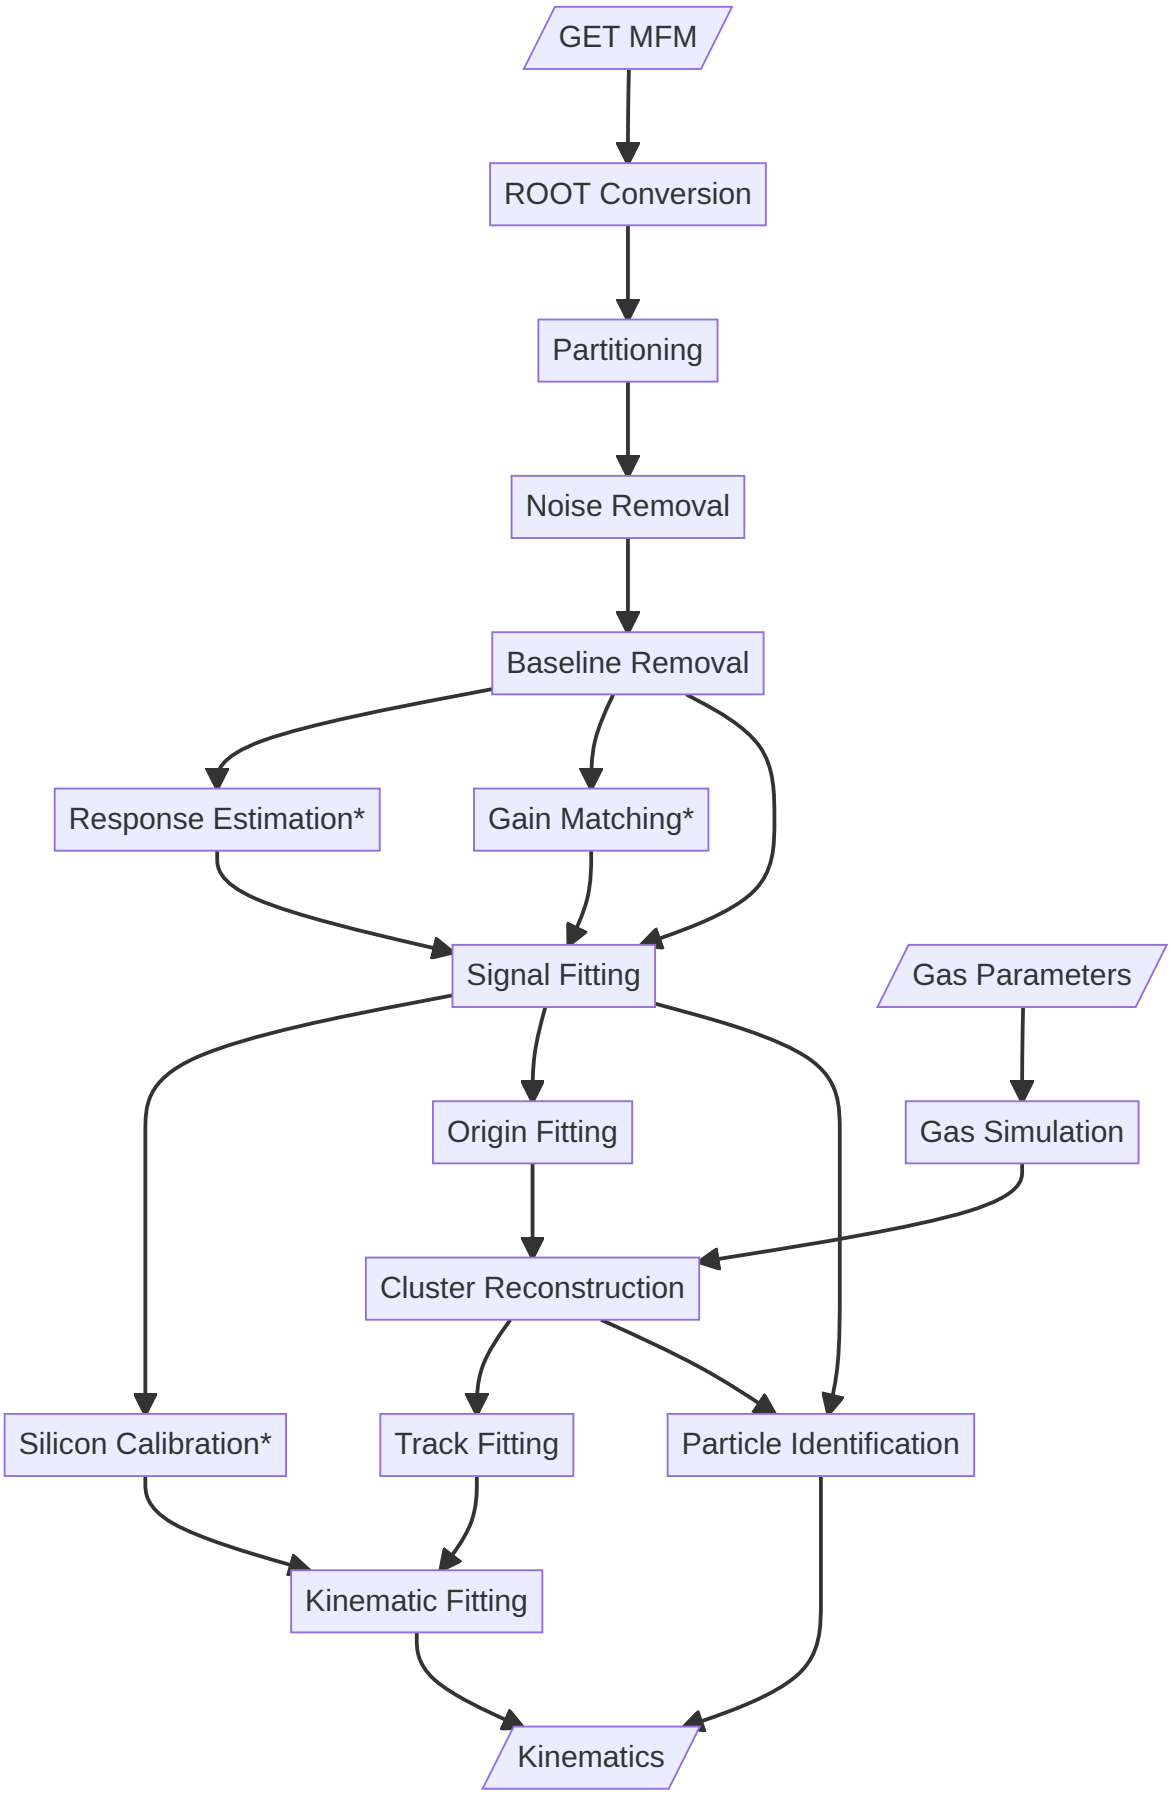


Fig. 8.1: Process diagram of the TexAT analysis.

8.1 Partitioning

The GET acquisition system serialises data using the MultiFrame Metaformat (MFM), which is described as: [89]

> a binary format for data acquisition and serialization that are self-contained, layered, adapted to network transfers[,] and evolving.

This file format is more suitable for disk storage than in-memory representation of data, as it is a hierarchical format in which each *frame* must be read, regardless of whether it is needed. The use of separate header and data sections means that the cost of skipping a frame is limited to parsing the header, but MFM is still a struct-based format. In many cases, columnar processing is both more performant and more expressive than an equivalent event-loop based approach [90].

At the earliest stages of the reconstruction pipeline, performance is primarily bounded by available memory, and therefore it is sensible to partition the data by the event multiplicity: the number of waveforms stored for each event. This process constitutes the first two stages of this analysis, seen in Fig. 8.1. In the ROOT dataset, the `mmMul` branch contains the per-event multiplicity. To regularise the dataset according to this multiplicity, it was partitioned into approximate chunks of a given aggregate multiplicity. Given that the memory of an MFM event is primarily determined by the number of waveforms that it contains, this effectively balances the memory footprint of each partition.

In subsequent stages, the memory bound gives way to compute costs, and so the partition size is refined to balance naive parallelisation against parallelisation overhead. Where data is written to disk, the Hierarchical Data Format version 5 (HDF5) format [91] is used as a consequence of its good language support, high compression ratios (with GZip [92]), and columnar representation. The Parquet format [93] is another strong format for columnar representation but, due to some library constraints, HDF5 was preferred.

8.2 Noise and Baseline Removal

8.2.1 Electronic Noise Removal

As discussed in Section 6.3.1, the AGET chip provides four additional channels for measuring the electronic noise (see Fig. 8.2). These channels can be used to eliminate some of the electronic noise inherent within the system.

As seen in Fig. 8.2, the fixed-pattern noise (FPN) waveforms typically include a non-linear baseline component. This component, when subtracted from non-FPN channels, restores the constant floor of the measured signals (see Fig. 8.3).

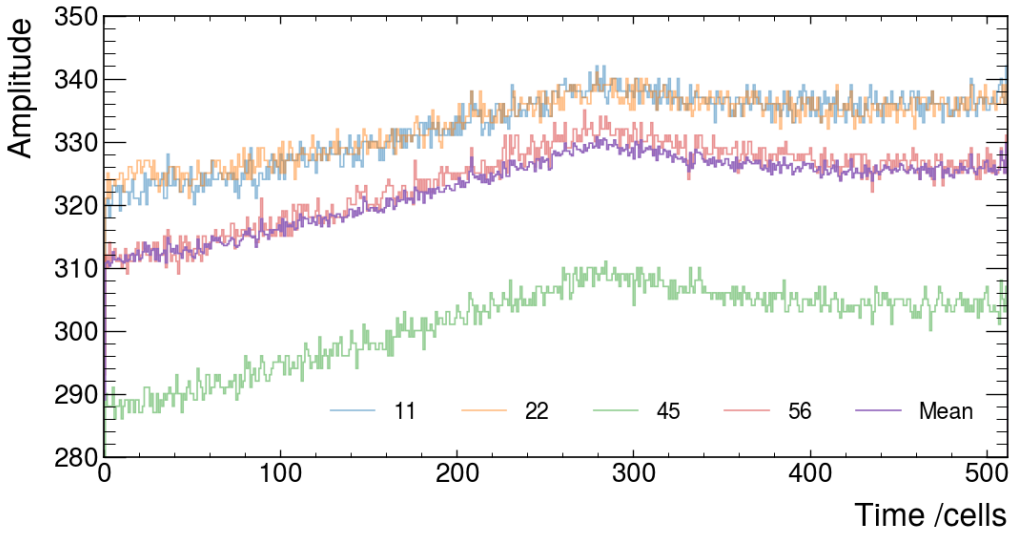


Fig. 8.2: FPN channel waveforms digitised by the GET acquisition system, alongside their binned mean, for a random AGET chip. Each FPN waveform has a different vertical offset, which can be modelled to first-order as contributing a constant offset in the computed average basliebaseline. In practice, this offset is unimportant; even when using the nearest-FPN channel for noise estimation, there is typically a non-zero offset that necessitates a baseline removal step. The x-axis is given in terms of AGET circular-buffer cells (see Section 6.3.1).

A number of noise removal strategies are discussed in [94], where it is established that an average over the set of FPN channels per-chip is optimal.

In addition to a low frequency component of the electronic noise, a separate baseline offset was observed in the MicroMeGAS and silicon detector channels. This component was removed using two separate methods:

8.2.2 MicroMeGaS Baseline Estimation

Given that time in the MicroMeGaS is primarly determined by the drift time of the liberated electrons in the gas volume, the baseline observed in the MicroMeGaS detector was not seen to occur at specific times. Therefore, in order to identify the region of baseline that would be linearly subtracted from the recorded signal, a local variance was computed for the channel samples, with a window of 96 cells. The value of the baseline was given by the locus of the minimum computed variance (see Fig. 8.4).

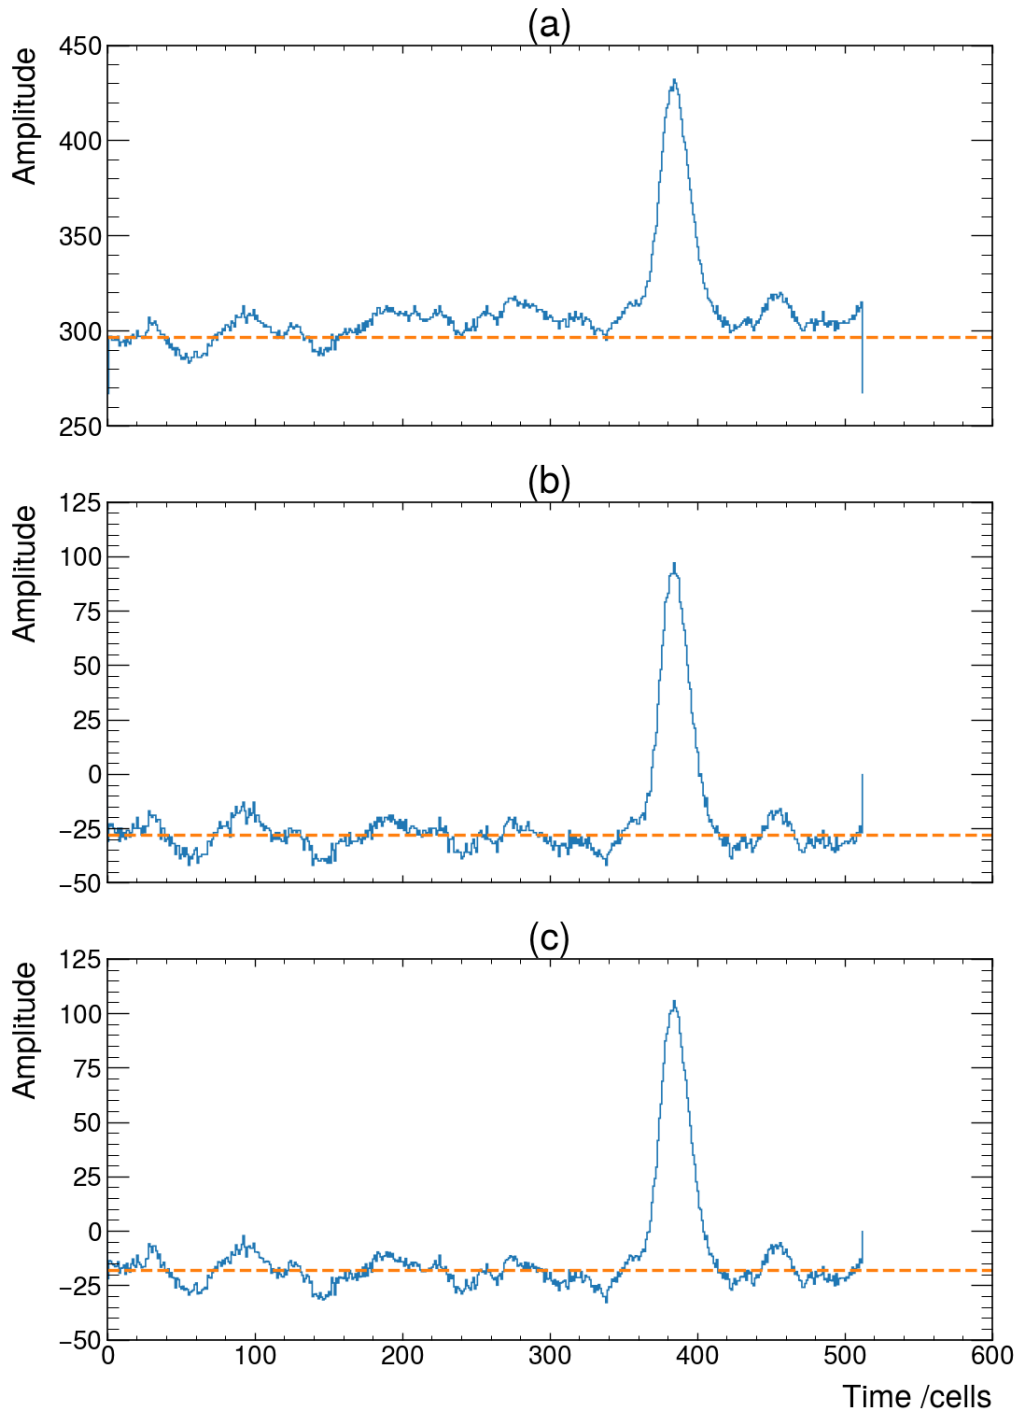


Fig. 8.3: Alternative estimates of the electronic noise for a singular MicroMeGaS GET channel. (a) The raw GET waveform. (b) The subtraction of a single FPN waveform. (c) The removal of the FPN waveform average. The leading samples of each estimate is fit with a linear baseline.

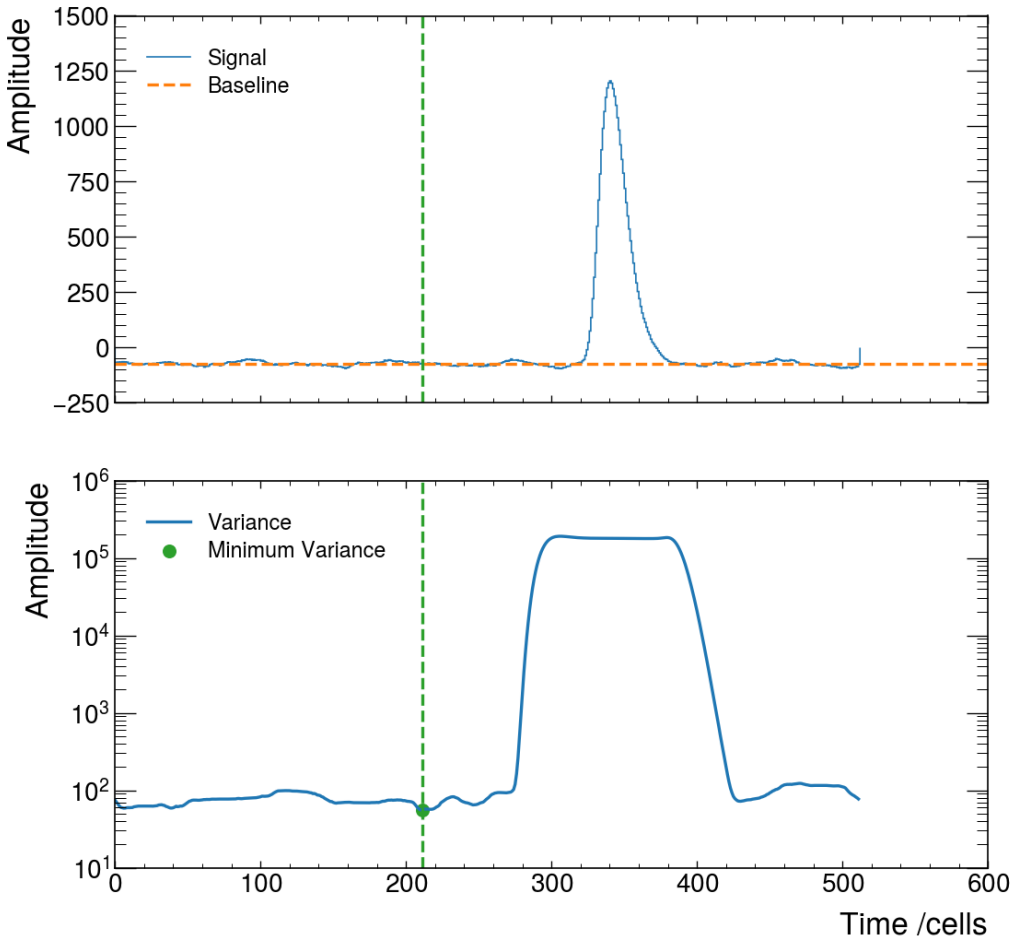


Fig. 8.4: Estimate of the signal baseline used to normalise MicroMeGaS samples. The locus of the minimum variance indicates the “flattest” portion of the waveform.

8.2.3 Silicon Baseline Estimation

Unlike the MicroMeGaS channels, the silicon quadrant detectors have a well defined boundary between the baseline and the interaction signals. This boundary occurs at a similar time for each detector, as it is the silicon detectors upon which the GET system was configured to trigger. To determine the baseline value, the mean of the sample values for the preceding ten cells was used (see Fig. 8.5).

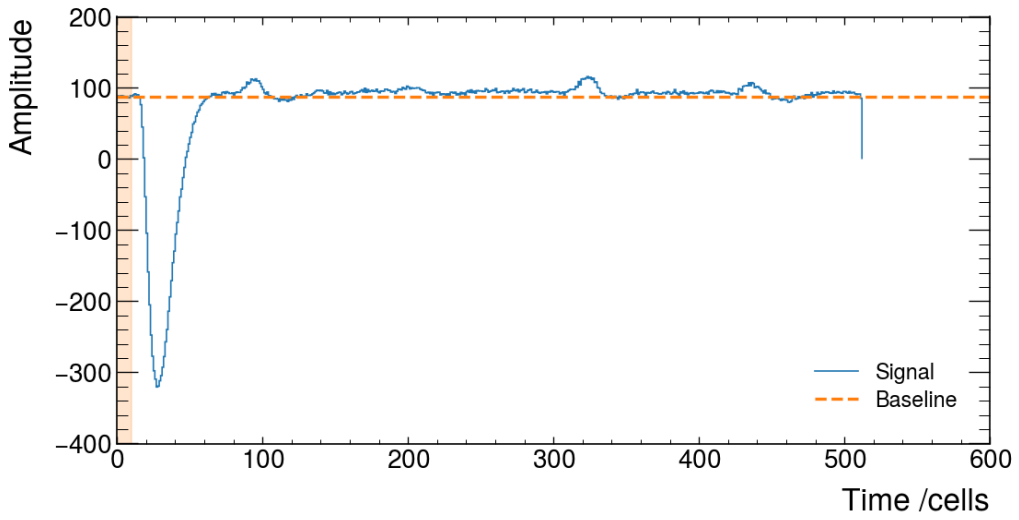


Fig. 8.5: Estimate of the signal baseline for the negative polarity silicon samples. The leading ten samples (shaded region) are used to compute an average constant baseline.

8.3 Response Estimation

Channels which employ the AGET shaper require deconvolution in order to recover the underlying signal. This process requires a good knowledge of the underlying response function of the system. An approximate analytical solution is given in [94]:

$$h(t) = A \cdot \exp\left(-3\frac{t}{\tau}\right) \cdot \left(\frac{t}{\tau}\right)^3 \cdot \sin\left(\frac{t}{\tau}\right) , \tag{8.1}$$

for $\tau = \frac{\tau_s}{\delta t}$, where τ is the shaping time and δt the reciprocal sampling rate. The authors of Ref. [94] note, however, that it is insufficiently accurate to be used with deconvolutional methods. Whilst the electronics can be configured such that the response function can be directly measured, this would have required further access to the detector which resides at the Cyclotron Institute, Texas A&M University [94]. Without the ability to experimentally derive a response function, an estimate was determined through two separate methods.

8.3.1 Iterative Estimation

In order to determine a better approximation for the response function, an iterative deconvolve-fit process can be applied to a set of MicroMeGaS waveforms. These waveforms should be taken from a random spanning sample of the recorded MicroMeGaS waveforms (see Fig. 8.6).

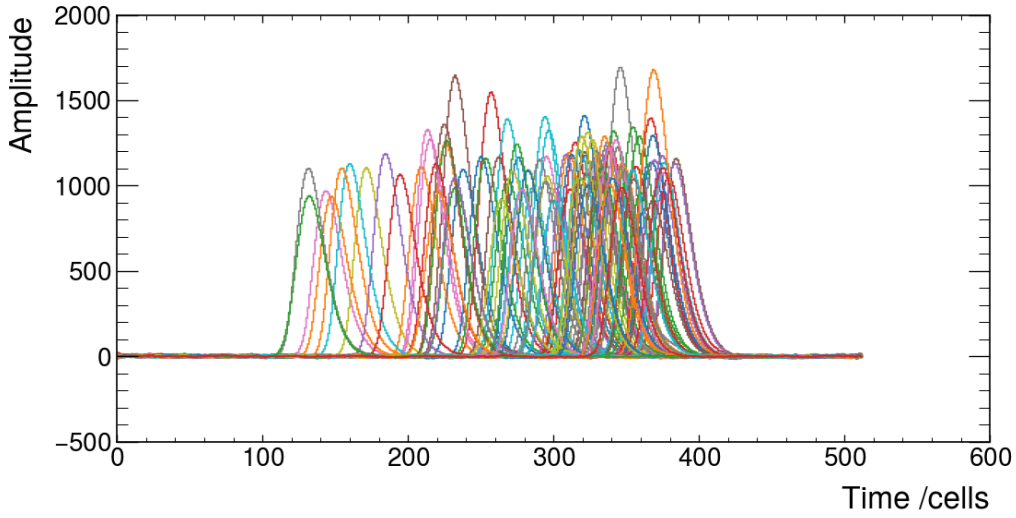


Fig. 8.6: A plot of the waveforms that are employed by the fitting process to derive an estimated response function.

Given such a sample, an initial estimate of the response function, $F^{(1)}$, can be used to fit each waveform with the convolution of a Gaussian function:

$$Y^{(1)} = F^{(1)} * \mathcal{N}^{(1)}\,, \tag{8.2}$$

where the initial function $F^{(1)}$ can be determined from Eqn. 8.1 with $\tau = 502\,\text{ns}$ (see Fig. 8.7) [94].

The resulting fit $\mathcal{N}^{(1)}$ can then be *deconvolved* from the original sample Y to yield $F^{(2)}$. This process can be repeated until the solution becomes stable, and the solutions for a set of random waveforms may be averaged to produce a singular

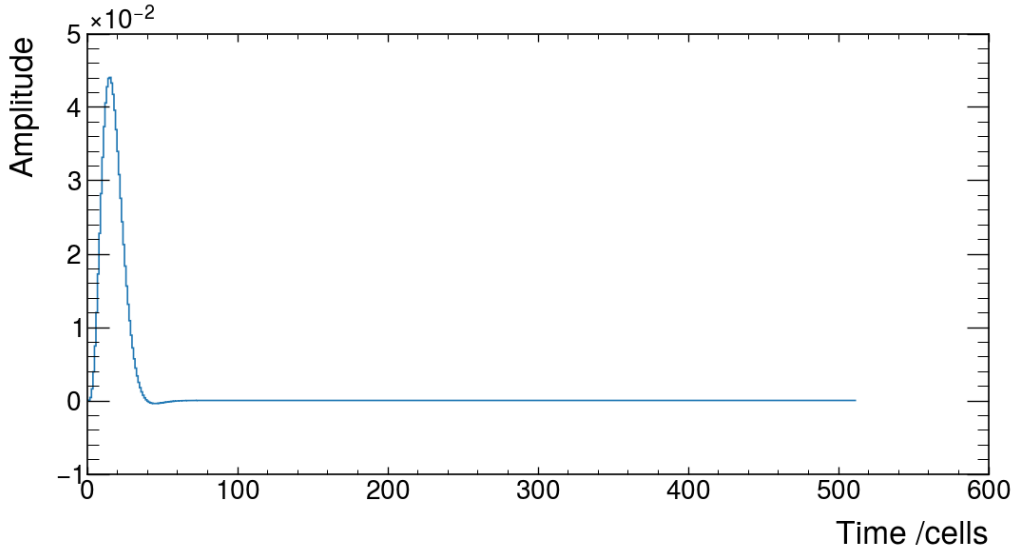


Fig. 8.7: GET response function for τ given in Eqn. 8.1.

response function F (see Fig. 8.8).

$$
\begin{aligned}
Y^{(1)} &= F^{(1)} * \mathcal{N}^{(1)}\,, \\
F^{(2)} &= \text{Deconv}(Y^{(1)}, \mathcal{N}^{(1)})\,, \\
&\vdots \\
Y^{(n-1)} &= F^{(n-1)} * \mathcal{N}^{(n-1)}\,, \\
F^{(n)} &= \text{Deconv}(Y^{(n-1)}, \mathcal{N}^{(n-1)})\,,
\end{aligned}
\tag{8.3}
$$

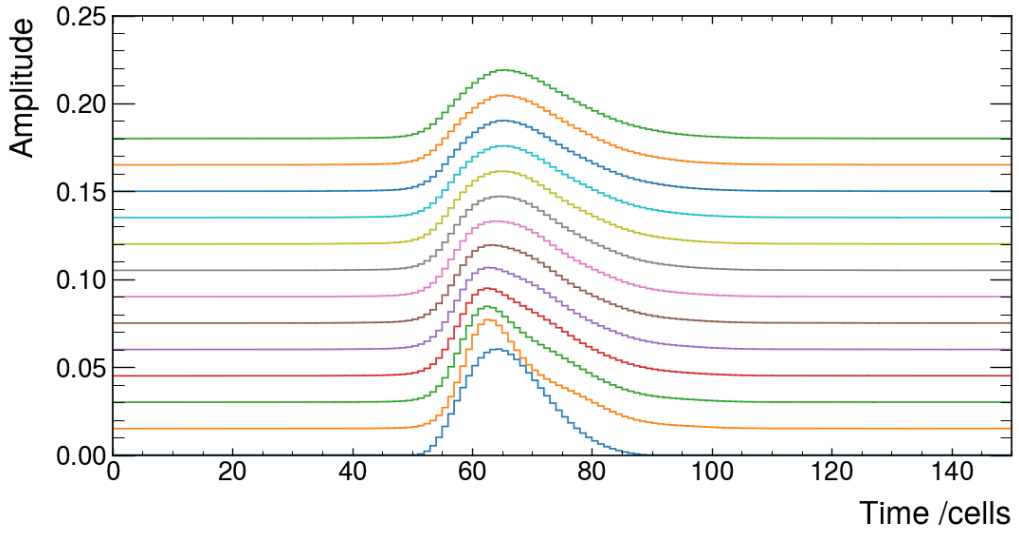


Fig. 8.8: A waterfall plot of the response function estimates computed over many sequential iterations. It can be seen that the estimated response function becomes broader with each iteration, ultimately converging upon a stable shape.

The initial trial function shown in Fig. 8.8 is a poor fit with respect to the final response function. The convergence of the fit can be improved by first optimising the parameter τ using a convolution fit (see Fig. 8.9).

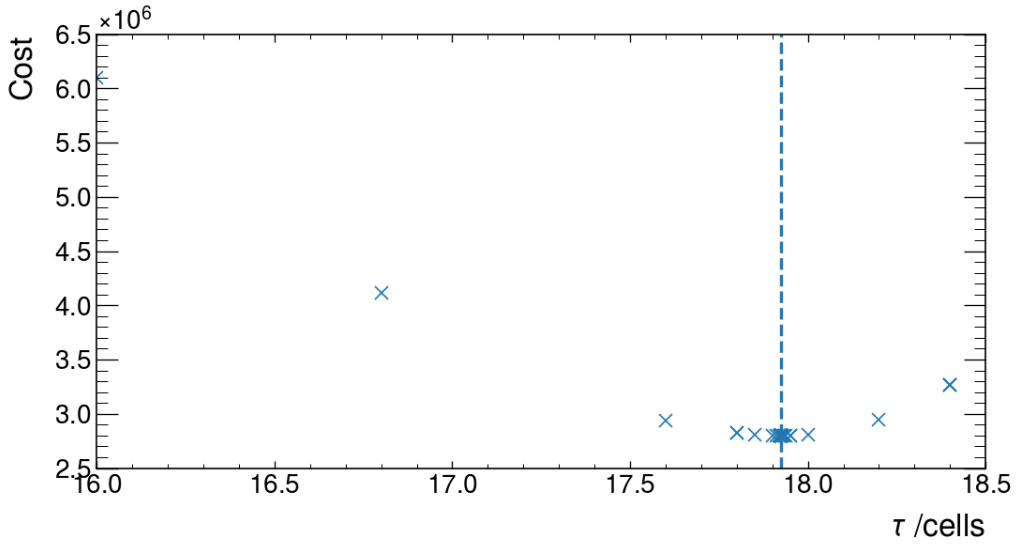


Fig. 8.9: A plot of the convolution fit solution cost as a function of the free parameter τ.

It can be seen that the fit value of τ is larger than the $\tau = 502\,\text{ns}$ that was taken from the literature (see Fig. 8.10).

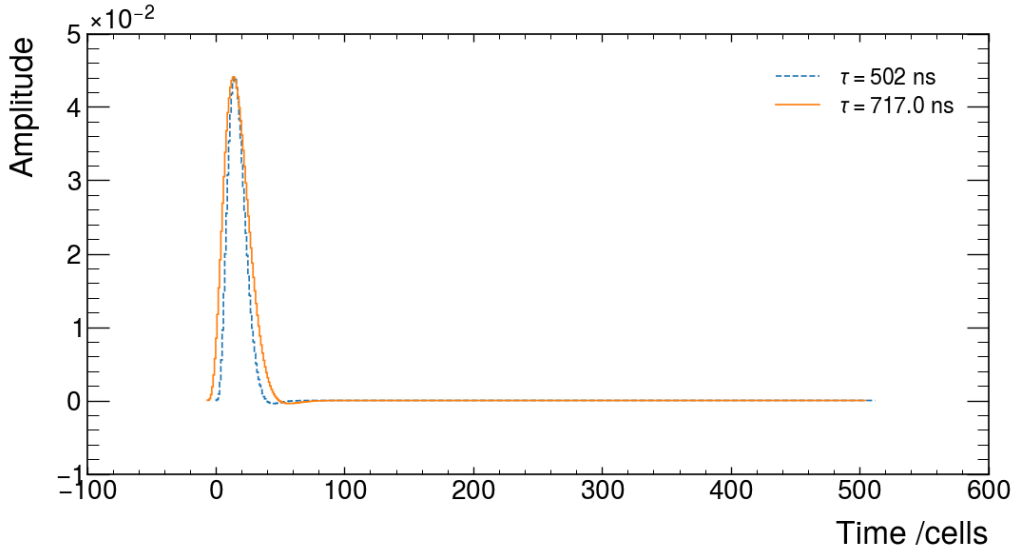


Fig. 8.10: Comparison of the response function for different values of τ.

Performing another estimate of F using Eqn. 8.3 converges to a solution much more quickly (see Fig. 8.11), with little change in macroscopic shape.

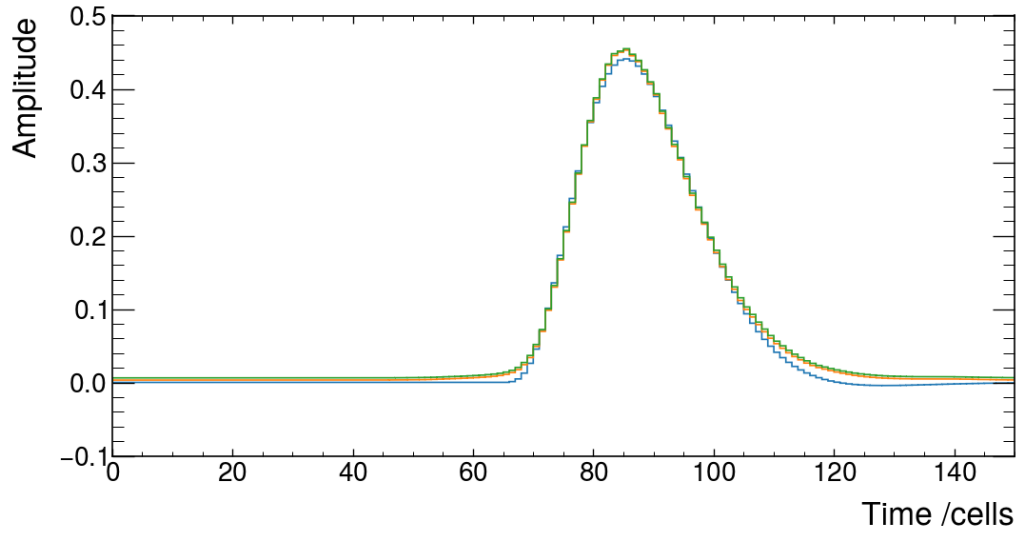


Fig. 8.11: A plot of the estimated response function over several iterations, with an initial estimate given by optimising Eqn. 8.1. It can be seen that the estimated response function becomes broader with each iteration.

The final response function estimate also agrees with that computed from the published τ. A comparison of the two solutions is shown in Fig. 8.12.

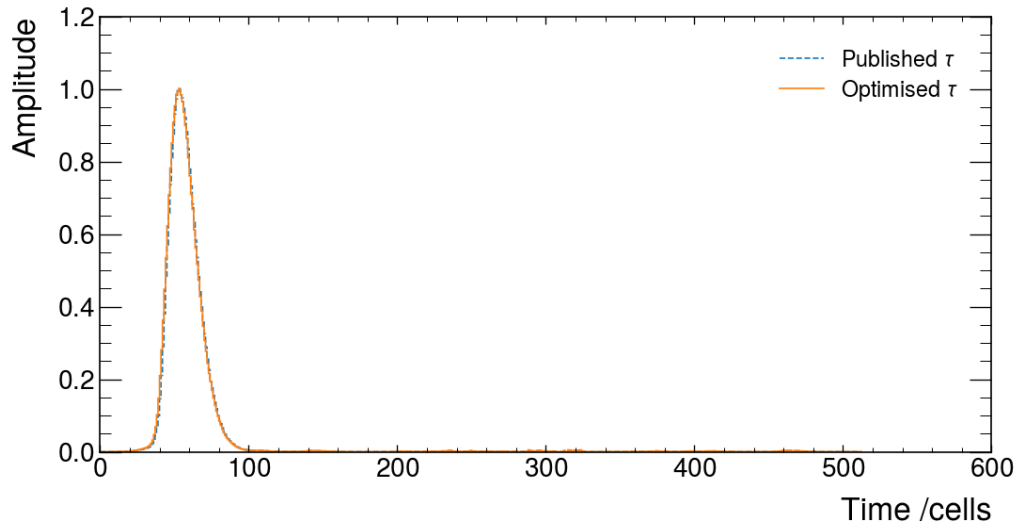


Fig. 8.12: A comparison between the response functions computed for different starting estimates of $F^{(1)}$.

That the published value of τ does not agree with the value determined in Fig. 8.9 is not necessarily surprising. Given that the convolution of any two Gaussians is also Gaussian, the solution from Eqn. 8.2 is not guaranteed to be unique, i.e. the resulting response function may include some Gaussian component common to the sampled waveforms.

8.3.2 Sample Averaging

As discussed in Section 6.2.1, the width of the deconvolved GET waveforms in the MicroMeGaS is partially determined by the shape of the charge clusters formed in the detector, i.e. the elevation of the path travelled by ionising radiation. This additional degree of freedom is accounted for above through the Gaussian convolution fit. It follows that the silicon waveforms may provide a more robust estimate of the GET response function, as the shape of the current signal produced by interactions with the silicon detector is predominantly determined by the shaper.

To first approximation, the measured signal can be assumed to be the convolution of the delta function at some time t and the intrinsic response function. To build a response function estimate, a random sample of silicon waveforms spanning a range of amplitudes is taken. The bin-wise mean of this set of waveforms is taken to produce an average waveform (see Fig. 8.13). This waveform is an approximate response function, which can then be time-shifted such that the peak coincides with $t = 0$.

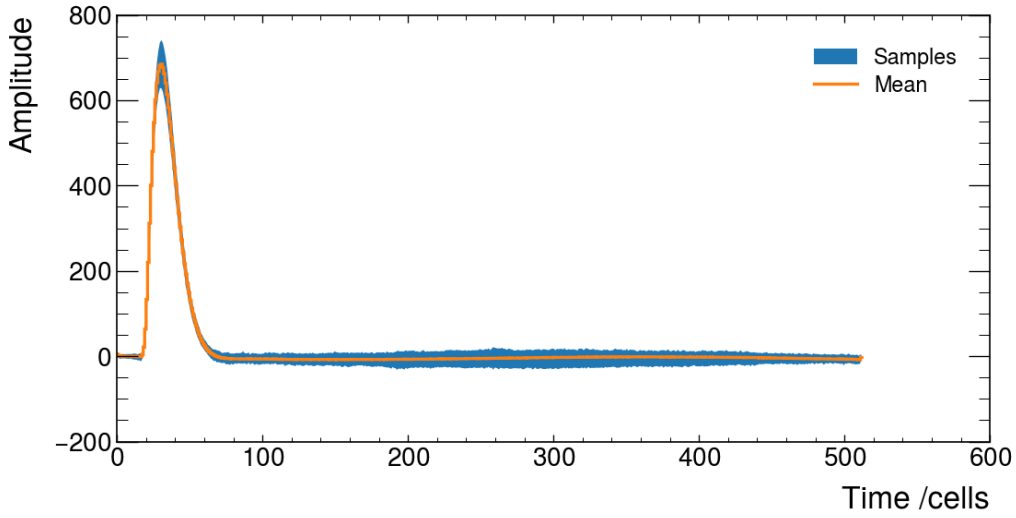


Fig. 8.13: A plot of the waveforms that are averaged to compute the silicon response function.

These two methods of estimating the response function of the GET electronics can be compared. In Fig. 8.14 the two response function estimates are visually compared. It can be seen that the MicroMeGaS iterative-estimate response function closely resembles the silicon average response function after Gaussian broadening. As the silicon average response function is computed directly, and is narrower than that of the MicroMeGaS response, it is better suited as the response function of choice.

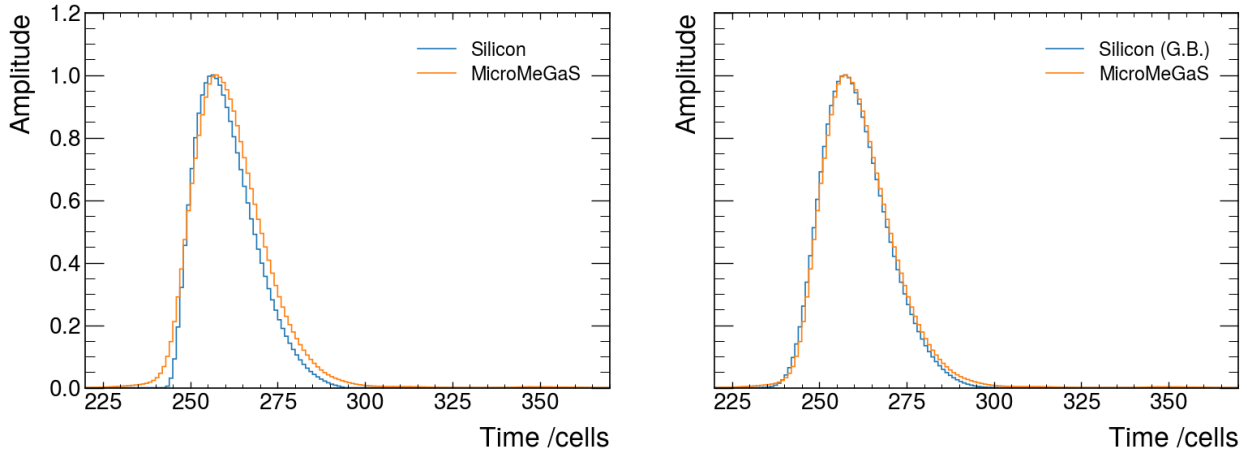


Fig. 8.14: A visual comparison of the response function estimates produced from iterative estimation and signal averaging. (a) A comparison of the raw response functions. (b) A comparison of the Gaussian broadened silicon average response function with the iterative estimate from the MicroMeGaS detector samples.

8.4 Gain Matching

In Section 6.2.1, it was noted that the MicroMeGaS anode is subdivided into a set of distinct zones that can be held at different potentials in order to spatially vary the gain. In this experiment, both the strip-chain (side) regions and final block of pads in the central region were held at a high gain in order to better resolve light particle tracks with low stopping powers (see Fig. 8.15).

Following position reconstruction, these gains must be accounted for if the collected charge is to be used for particle identification. Given that the strips, chains, and final central pads were held at the same potential, the relative gain of these regions can be determined solely by looking at the relative gain of the final central pad region. This was experimentally measured by observing the change in charge collected by the rows of pads either side of the high-low gain boundary (see regions (b) and (c) of Fig. 8.15).

After fitting this distribution with a binned likelihood estimator, the relative gain between the low and high gain regions was found to be 0.105 with a standard deviation of 0.018.

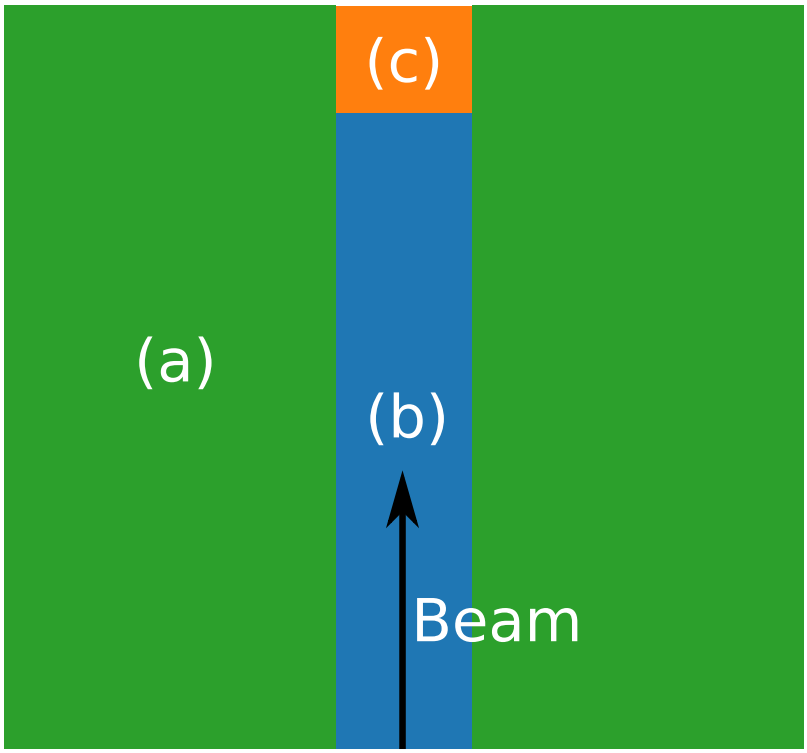


Fig. 8.15: Illustration of the significant zones within the MicroMeGaS anode segmentation. The strip-chain region (a), held at the same potential as the high-gain pads (c), is shaded in green. The high-gain pads region (c) is shaded in orange, whilst the low-gain pads (b) are coloured in blue. The beam direction is indicated by a black vertical arrow.

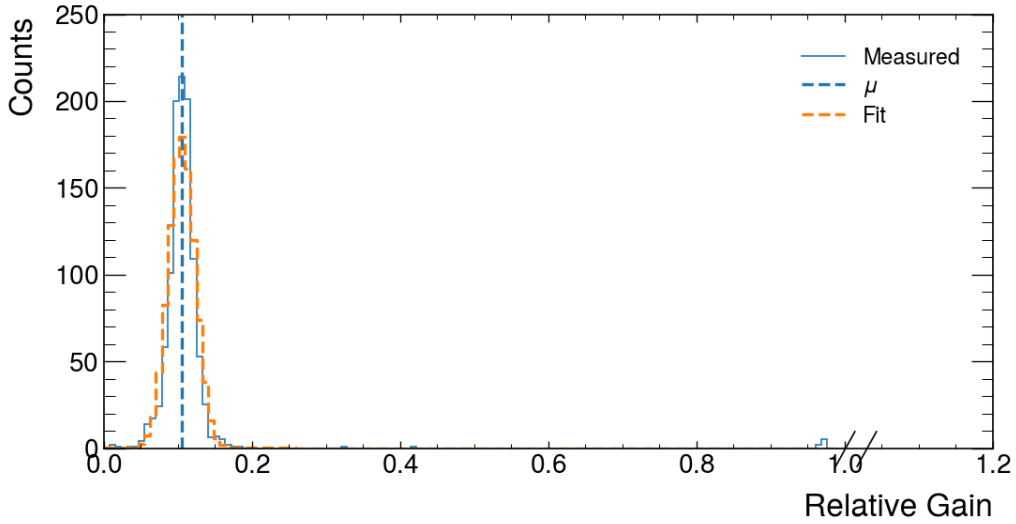


Fig. 8.16: Ratio of the mean collected charge between the low gain and high gain region boundary in the Micromegas pads.

8.5 Signal Fitting

In the above sections, it was established that the sampled GET waveforms are effectively convolutions of an original source signal and the intrinsic GET shaper response function. Having determined an approximation for this response function, the source signals can then be recovered by deconvolutional methods.

8.5.1 MicroMeGaS Signals

For each preprocessed MicroMeGaS waveform, a convolutional Gaussian fit is performed using a nonlinear least squares optimiser that solves $y = F * \mathcal{N}(\phi)$ for the set of model parameters ϕ. This is preferred over computing the Gaussian fit of the deconvolved signal; although repeated convolution is much more computationally expensive, it is less vulnerable to artefacts produced by the deconvolution technique. For samples which do not exceed the dynamic range of the acquisition system, the peak multiplicity is estimated using the GOLD deconvolution algorithm with boosting [95], which collapses broad peaks into smooth, sharp peaks. These peaks are then identified using a simple turning-point localiser, which is used to seed the initial fit parameters (see Fig. 8.17).

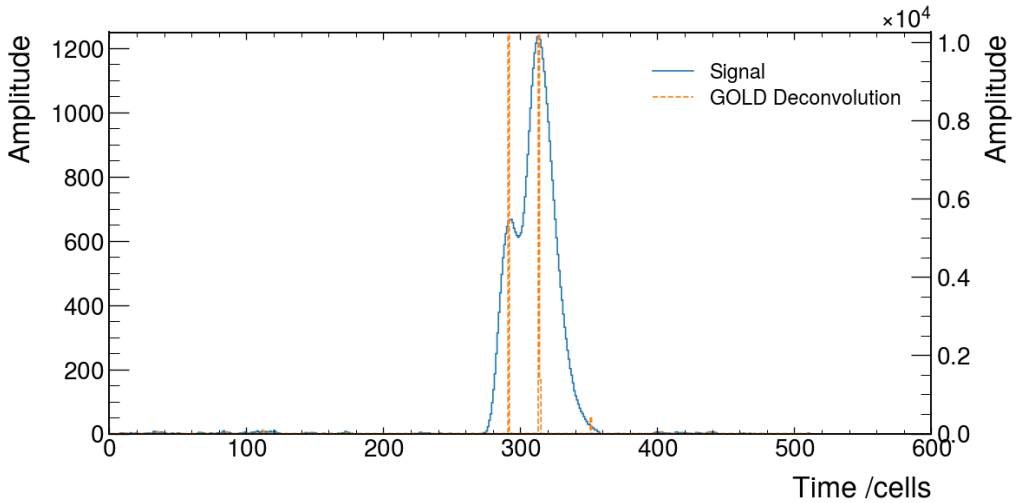


Fig. 8.17: A plot of the peaks identified within a MicroMeGaS waveform sample using boosted GOLD deconvolution. It can clearly be seen that the signal is formed of two separate peaks, whose convolution cannot trivially be separated into distinct peaks. Experimentally, this corresponds to overlapping charge clusters (in the drift plane).

In some cases, the measured sample exceeds the dynamic range of waveform amplitudes that can be measured. This saturation introduces a discontinuity in the derivative of the sample, which renders it un-amenable to deconvolution (see Fig. 8.18).

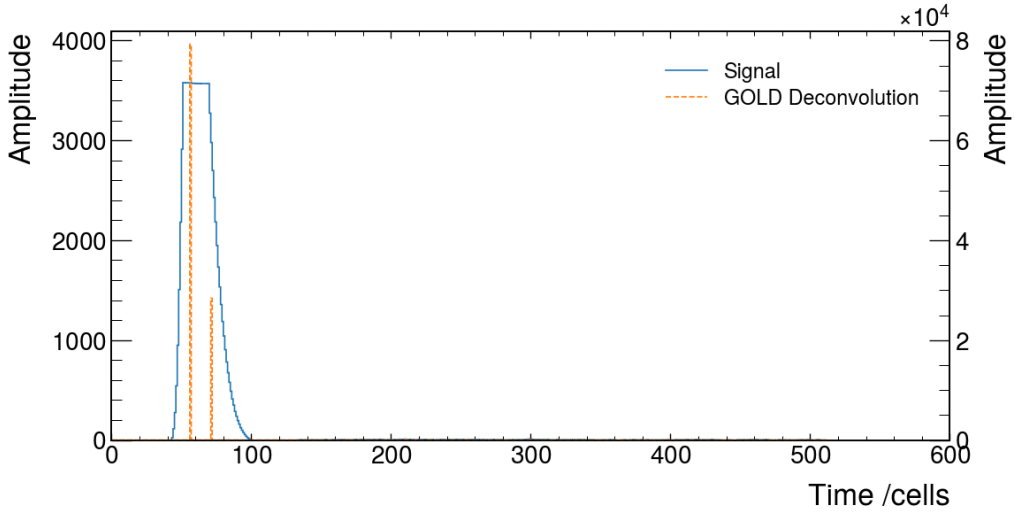


Fig. 8.18: A plot of the saturated MicroMeGaS waveform sample, and its GOLD deconvolution.

In such instances, the peak multiplicity of the sample cannot easily be determined. Given that MicroMeGaS waveforms provide both charge and temporal (spatial) information, it is important that these signals are not discarded prematurely. To recover these signals, it is therefore assumed that the measured sample contains only one peak, which is fit as described above using an additional mask. This mask that forces the fit algorithm to ignore the saturated regions. Consequently, these fits are more prone to poor estimation of the signal amplitude where a significant portion of the signal is lost to saturation effects.

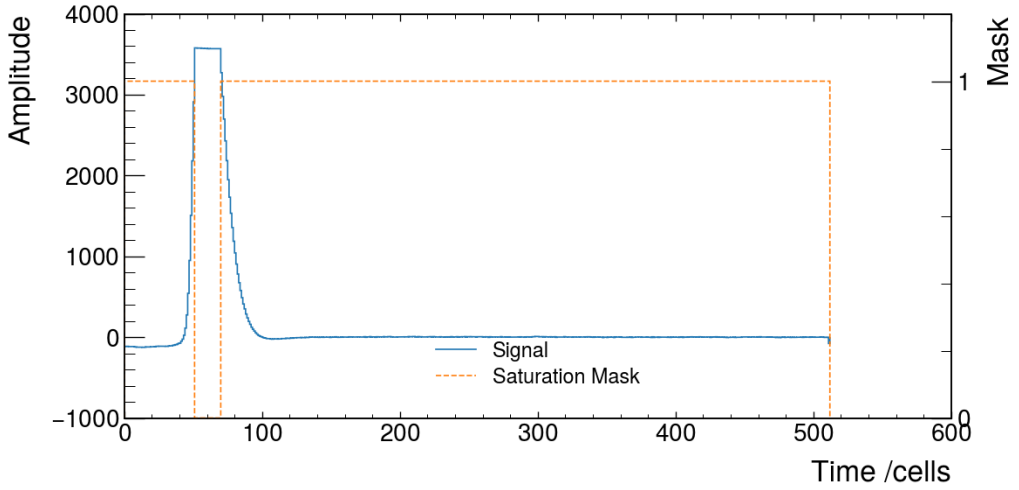


Fig. 8.19: A plot of the saturated MicroMeGaS waveform sample, and the mask used for fitting.

8.5.2 Silicon Signals

In the silicon detector, the preliminary calibration run that is analysed in Section 8.7 was used to ensure that the dynamic range of the silicon detectors was sufficiently large during the acquisition phase of the experiment. As such, signals measured by the silicon detectors do not suffer from the saturation phenomenon described above.

Meanwhile, the response function derived in Section 8.3 is, by virtue of its construction, a poor choice for (de)convolution fits of the silicon signals as the necessary source waveforms approach widths of ~1 time cell. Therefore, a simple linear transformation of the response function was used to fit silicon waveforms, which correctly recovers the peak locus and amplitude (see Fig. 8.20).

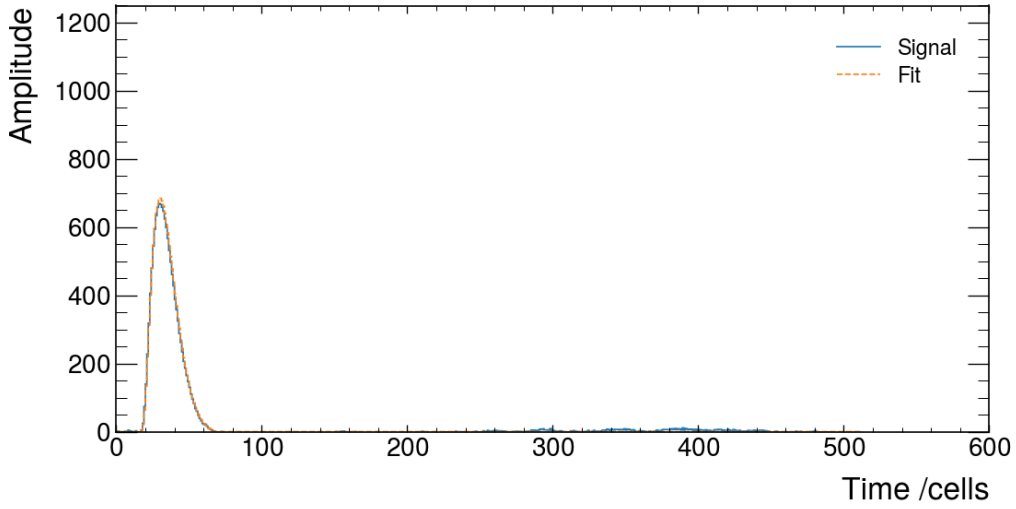


Fig. 8.20: A plot of silicon waveform sample, and its linear response fit.

8.6 Signal Ringing

In the presence of shaped signals whose amplitudes exceed the dynamic range of the ADC, spurious waveforms are observed in the other channels of the acquisition system. This phenomenon, known as “ringing”, presents a challenge to the reconstruction of particle trajectories within the TPC. An example waveform is shown in Fig. 8.21.

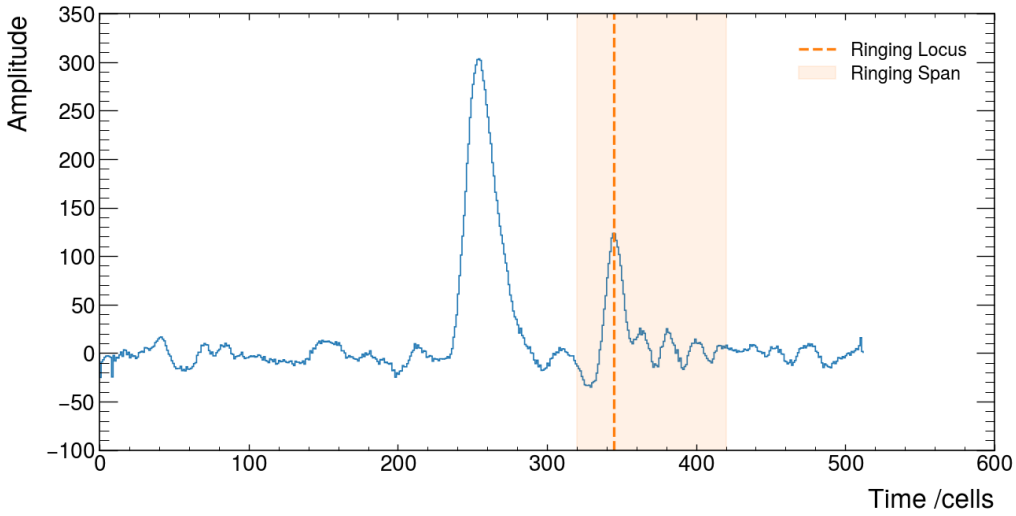


Fig. 8.21: A waveform with the ringing phenomenon visible at $t \sim 345$. A distinct non-ringing component is also visible at $t \sim 254$. The shaded region indicates the span covering ringing in this channel.

8.6.1 Pulse Width Identification

In order to identify and eliminate these spurious signals, one approach might be to consider the narrow width of the primary peaks. On average, it can be seen that ringing waveforms posses sharper peaks than true waveforms (see Fig. 8.22)

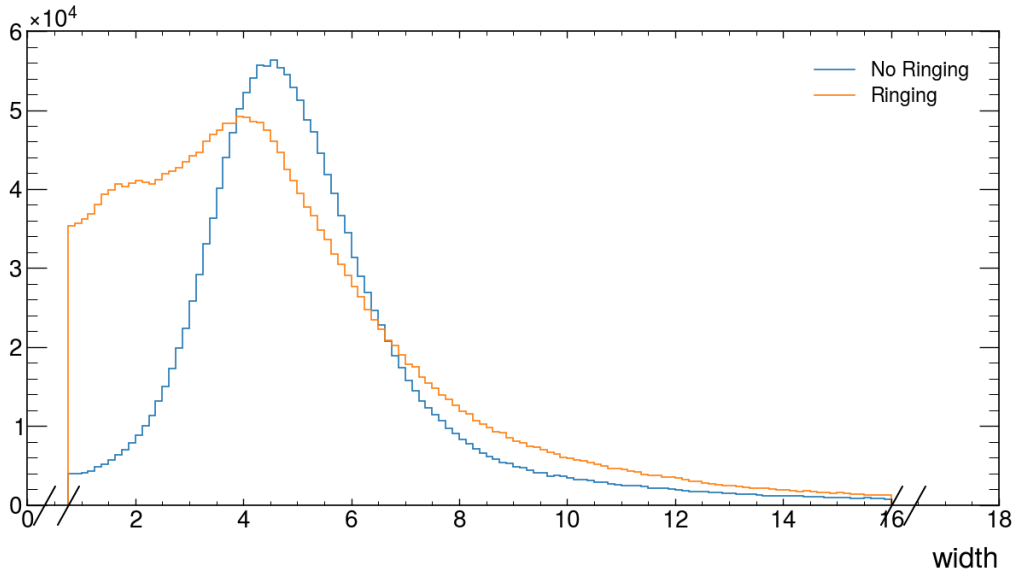


Fig. 8.22: The distribution of peak widths for ringing and non-ringing events.

However, given that the shape of these ringing waveforms is not that of a typical GET waveform (see Fig. 8.21), it follows that we cannot rely on the fit (which assumes a regular response function) as a robust description of the ringing peak. In

order to identify and remove these waveforms, a different approach must be taken.

8.6.2 Pulse Height Identification

As outlined above, the appearance of ringing waveforms in the sampled GET waveforms follows from the presence of saturated channels. Predictably, these ringing waveforms also occur at a similar time to that of the saturated signals (see Fig. 8.22)

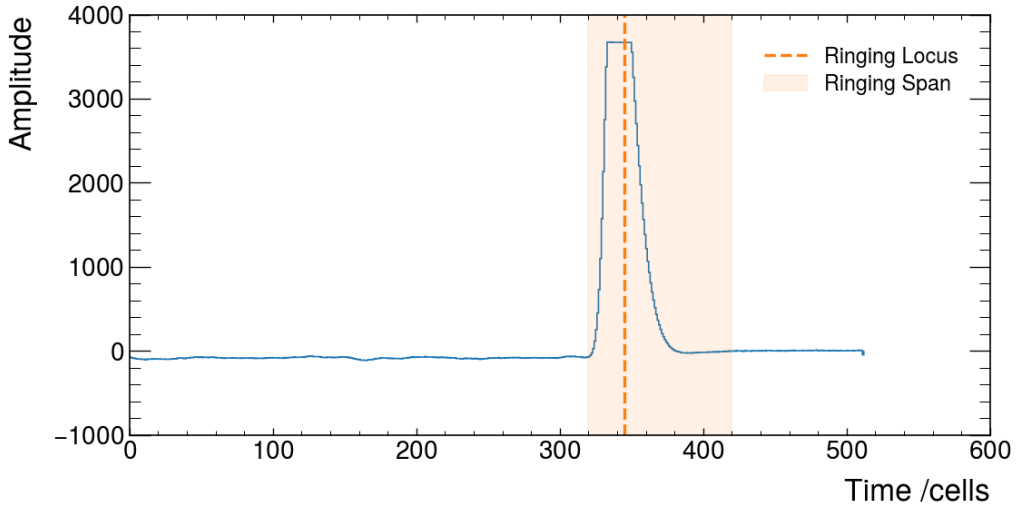


Fig. 8.23: A waveform from the same event as Fig. 8.21 in which the signal is saturated at $t \sim 345$. The shaded region indicates the span covering ringing signals in other channels.

Consequently, the loci of peaks fit to saturated waveforms can be used to define a set of time intervals in which ringing signals may be anticipated. Given that one cannot use the GET fit to perform a robust classification between a measurement and a ringing signal, the amplitude of the signal may be used instead. It is observed that the majority of ringing signals have amplitudes well below that of the saturated signal, on the order of $y < 750$ units. By taking the logical AND of these two conditions, a better classification of ringing signals can be performed.

8.6.3 Periodicity

The presence of saturated signals in the MicroMeGaS detector gives a locus (in time) for ringing signals to occur. It is not sufficient to discriminate those elements which are trigged by legitimate signals at this locus from those which are experiencing the ringing phenomenon. The damped periodicity of the ringing signals can be used to further discriminate one from the other. Fig. 8.25 shows the result of computing a deconvolution of the baseline-adjusted waveform. It

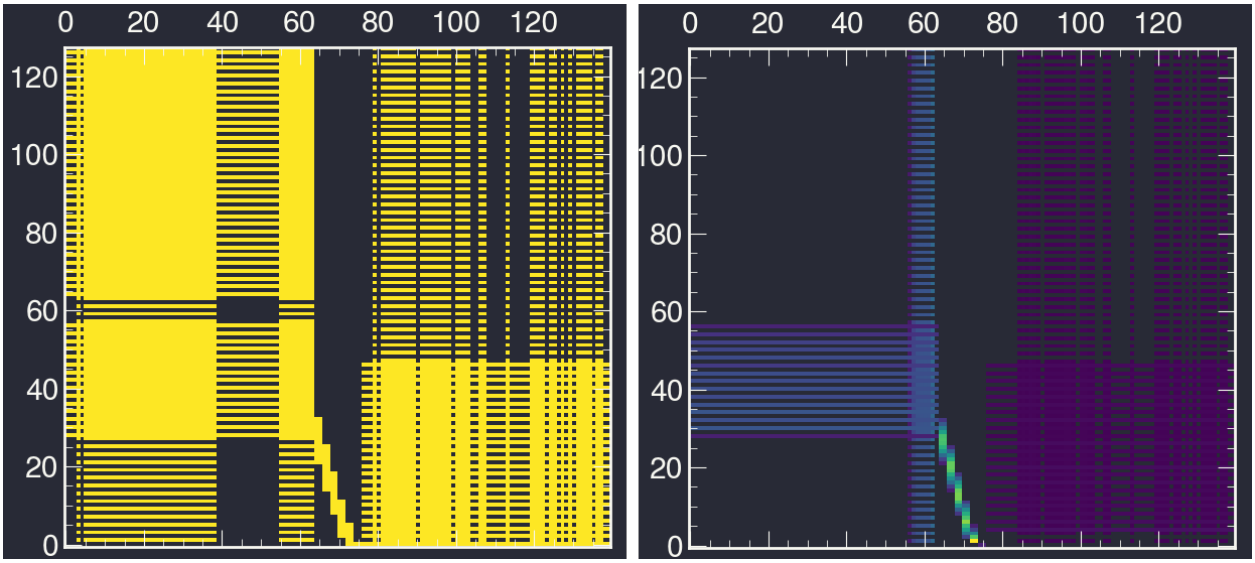


Fig. 8.24: A 2D hitmap of the MicroMeGaS anode for a random event. (a) A raw event heatmap with ringing signals. Note that the majority of the elements in the left side region have been activated. The linear track formed in the central pads region does not suggest such a pattern will be observed in the side region. (b) A modified hitmap with ringing signals and invalid fits removed. The resulting hitmap more closely agrees with the predicted track yielded by the central pads region.

can clearly be seen that the peaks in the shaded region are of a similar amplitude, and narrow widths, whilst the those associated with the locus at $t = 250$ are dominated by a single, broad peak.

The baseline is adjusted such that the waveform is positive definite. The GOLD algorithm [95] used to deconvolve ringing waveforms requires such a condition. Furthermore, the baseline estimation procedure outlined in Section 8.2.2 does not produce reasonable solutions in the event that the baseline is dominated by a strong noise component. This is a significant problem for signals with ringing artefacts, which therefore require baseline adjustment in order to properly recover the peaks.

It follows that a veto window that start immediately before the saturated waveform locus, and extends several peak-widths beyond it, can be defined. Waveform peaks that are identified within this veto window are only retained iff. they demonstrate a reasonable width (see Section 8.6.1), and they have no secondary ringing peaks (identified as above) with a primary-amplitude ratio that exceeds a threshold value.

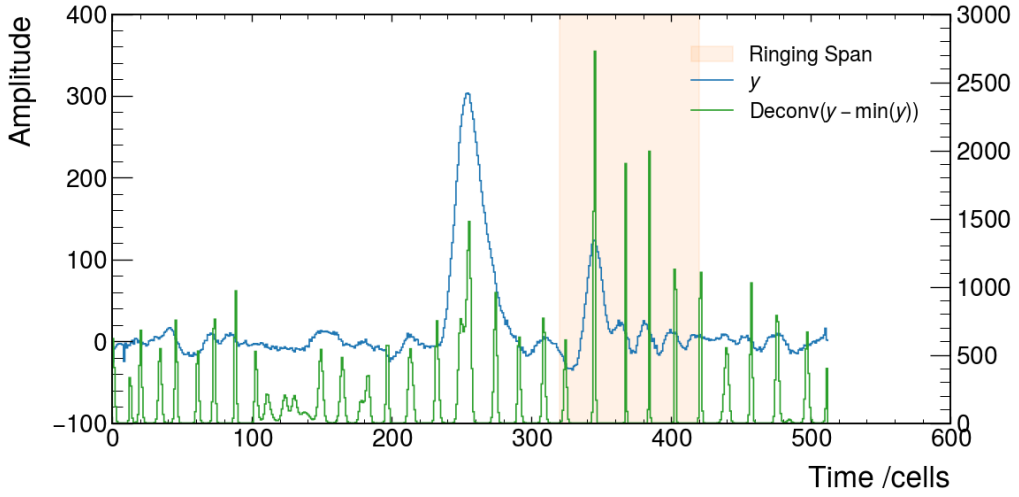


Fig. 8.25: A waveform with the ringing phenomenon visible at $t \sim 345$, and the deconvolution of the baseline adjusted signal overlaid. A distinct non-ringing component is also visible at $t \sim 254$. The shaded region indicates the span covering ringing in this channel. It can be seen that the green peaks in the shaded region have much higher relative amplitudes, and much narrower widths.

8.7 Silicon Calibration

8.7.1 Semiconductor Detectors

The silicon detectors that surround the TexAT TPC are used to measure the energy of light products that escape the active volume. Within each detector, the periodic structure of the semiconductor's crystalline lattice establishes a set of electron energy bands separated by ranges of forbidden energies. At the Fermi level, the lower energy (*valence*) band comprises a continuum of fully occupied states that are bound to particular lattice sites, whilst the higher energy (*conduction*) band contains entirely unfilled states in which the electrons are free to move throughout the lattice (see Fig. 8.26).

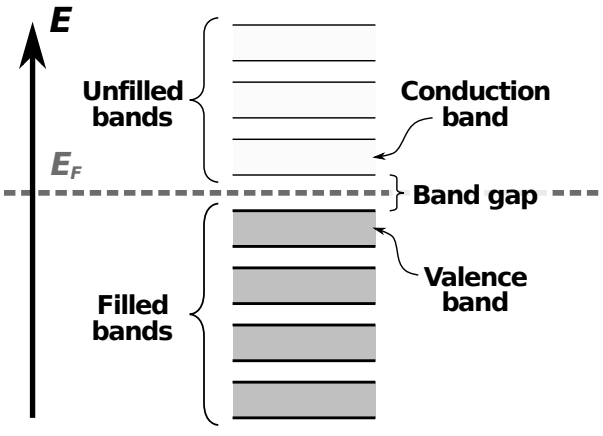


Fig. 8.26: An illustration of the band structure in a semiconductor. The dashed line indicates the Fermi level that separates the valence and conduction bands. Figure by Tim Starling & Mliu92, distributed under a CC BY-SA 4.0 license.

The excitation of an electron across the gap between the valence and conduction bands establishes an electron-hole pair, which is free to migrate throughout the lattice. It is the motion of these charge-pairs under an electric field that gives rise to a current. Above 0 K there will exist some steady-state concentration (given by the Boltzmann distribution) of thermal electrons in the conduction band such that the material will be conductive.

In order to observe the charge-pairs produced by ionisation, there must exist an electric field such that the charge-carriers do not recombine, but rather induce a current through their motion. As any externally developed potential will induce an opposing potential in the material, it follows that this field must be established through intrinsic diffusion of the charge carriers. To this end, one can *dope* the semiconductor lattice with substitutional impurities that introduce a surplus of a particular charge carrier. These dopant atoms do not conserve the electron-hole equilibrium, i.e. they introduce a majority of the singular charge carrier that exclusively determines the material's electric conductivity. Semiconductors doped with donor (electron) impurities are denoted N-type semiconductors whilst those with acceptor impurities are termed P-type materials (see Fig. 8.27).



(a) (a) (b) (b)

Fig. 8.27: An illustrative diagram of the minority-majority charge carriers in an (a) p-type semiconductor and (b) n-type semiconductor. A surplus of mobile charge carriers can be seen in the conduction band. The charge carriers in p-type semiconductors (a) are termed holes, whilst those in n-type materials (b) are electrons.

Upon interfacing two alternately doped semiconductors, a P-N junction is formed. An initial gradient in the concentration of charge carriers leads to the establishment of a "depletion zone" in which electrons and holes combine, filling the valence band. This displacement of the majority charge carriers in turn establishes an electric field that inhibits further diffusion (see Fig. 8.28). These two phenomena render the depletion zone an ideal region for radiation detection; the electron-hole pairs formed by ionisation will quickly move under the electric field to establish a measurable current.

The *contact potential* established by the space charge at either end of the junction is of the order $\sim 1\,V$, which yields a poor electron-hole mobility. An electric potential can be applied to either side of the junction, widening the depletion region and reducing the likelihood of charge-pair recombination. Known as reverse biasing, this significantly improves the detector performance.

It is the time-integral of the current induced within the semiconductor that is proportional to the energy deposited within the material. In order to determine the constant of proportionality between the measured signal and the deposited energy, the detector must first be calibrated with a known source.

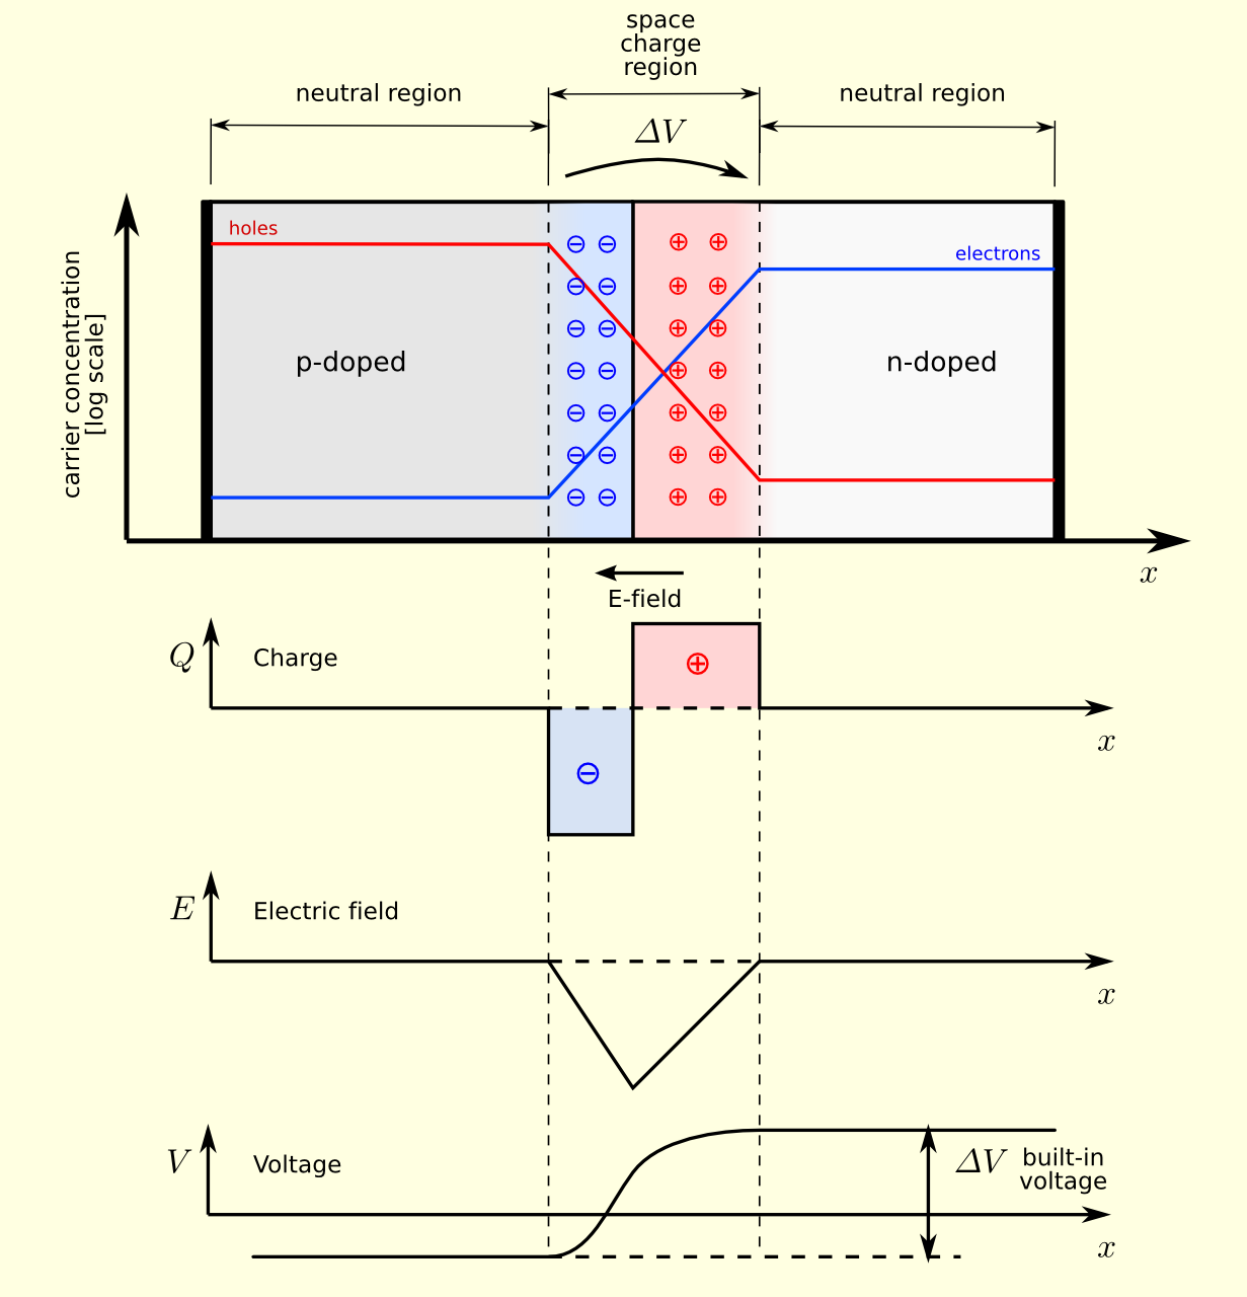


Fig. 8.28: An illustrative diagram of the charge-carrier distribution across a P-N junction. Additional graphs show the charge, field, and electric potential, across the junction. Figure by TheNoise, distributed under a CC BY-SA 3.0 license.

For this experiment, a four-peak alpha source was suspended within the TexAT TPC chamber, which was held under vacuum. A list of constituent isotopes and their decay radiation is given in Table 8.1.

Isotope	Energy /keV	Intensity (%)
241Am	5442.8	13.1
	5485.6	84.8
244Cm	5762.6	23.1
	5804.8	76.9
239Pu	5105.5	11.9
	5144.3	17.1
	5156.6	70.8
148Gd	3182.7	100

Table 8.1: Isotopes present in the four-α calibration source, and their α decay radiation. From ENSDF database as of March 24th, 2023. Version available at http://www.nndc.bnl.gov/ensarchivals/.

The waveforms recorded by the GET acquisition system were fit using the methods described in Section 8.5, and the corresponding charge accumulated in a histogram. It can be seen in Fig. 8.30 that the majority of the silicon quadrants measured similar signal amplitudes for the given sources. A total of 11 quadrants were found to have an insufficient number of hits to be able to build a calibration, primarily those in the forward silicon detector array. See Fig. 8.29 for a visual representation of these missing quadrants.

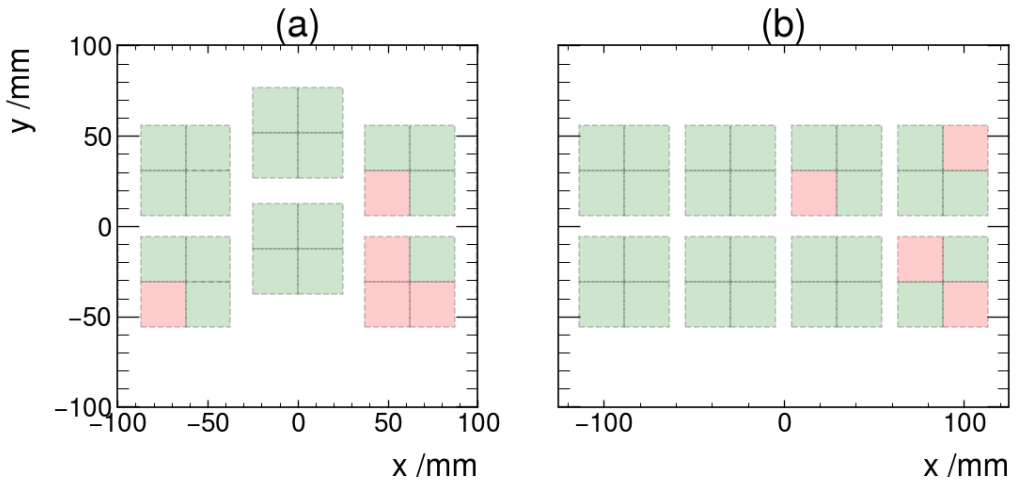


Fig. 8.29: A schematic diagram of the (a) forward and (b) side silicon detector arrays indicating quadrants with incomplete calibrations in pink.

Where charge-pair formation is not the primary mode of energy loss within the semiconductor, the number of created

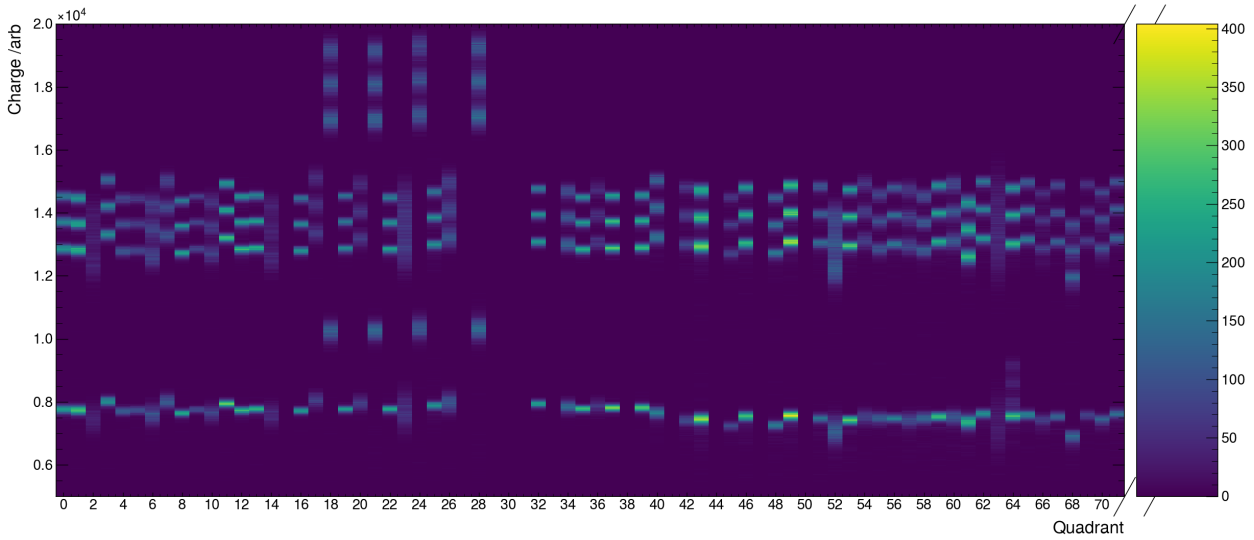


Fig. 8.30: A 2D histogram of the charge distribution measured by each silicon detector for a fixed calibration source containing 241Am, 148Gd, 239Pu, and 244Cm.

pairs is governed by Poisson counting statistics $\mathrm{Var}(N) = N$. Meanwhile, in the limit that charge-pair formation is the singular energy-loss mechanism, the number of charge pairs is exactly determined by the deposited energy (ignoring the distribution on the ionisation energy), i.e. $\mathrm{Var}(N) = 0$. The ratio between the observed variance and its Poissonian expectation value is the Fano factor: [96]

$$F = \frac{\langle \delta N^2 \rangle}{\langle N \rangle}\,, \tag{8.4}$$

where $\langle \delta N^2 \rangle$ is the observed variance and $\langle N \rangle$ the observed mean. In silicon detectors, which exhibit a high charge-pair formation efficiency, F is experimentally found to be of the order 10^{-1} [97][98].

Note: Simplified "final-state" models can predict the experimental values for F with a microscopic assuming a single-particle-distribution function. One such model, the "bathtub model", assumes no inter-charge-pair correlations and draws from this single-particle-distribution to partition the deposited energy. It has been shown that correlational effects are required to fully reproduce experimental measurements of F, through models such as the random parking problem (RPP) [96].

For each detector, the charge distribution was fit as a sum of four Gaussian distributed random variables.

Note: The motivation for a Gaussian fit function follows from the application of the central-limit-theorem to the final-

state model of charge-pair formation [96], yielding

$$P(E_N) = C_N \exp\left[-\frac{(E_N - N\epsilon)^2}{2N\langle\delta\epsilon^2\rangle}\right] . \tag{8.5}$$

In order to determine the set of initial fit parameters, a cumulative mass lookup was used to estimate the peak loci (see Fig. 8.31), and sensible values chosen for the peak width and scaling factor.

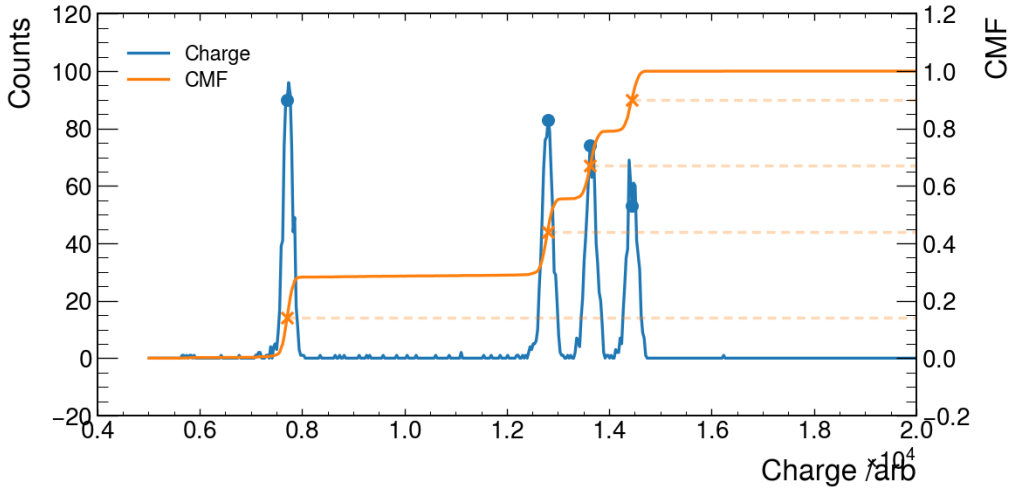


Fig. 8.31: The charge distribution for a random silicon detector, superimposed with a plot of the cumulative distribution. The dashed lines indicate particular fractions of the cumulative distribution at which the peaks are localised.

This method of localising the decay peaks from the calibration source is robust in the face of varying detector noise.

8.7.2 Calibration

Following the peak fitting process, a linear regression was then performed against the known energies given in Table 8.1 to obtain a set of calibration coefficients for each detector. A superposition of the calibrated charge distributions can be seen in Fig. 8.32.

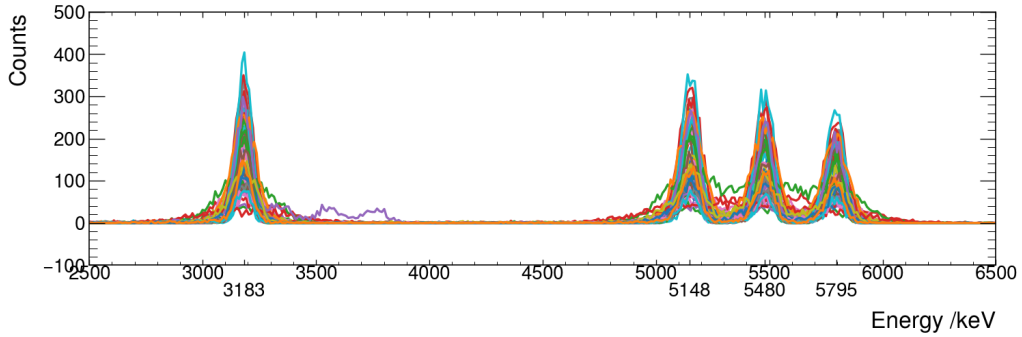


Fig. 8.32: A combined plot of the calibrated charge distribution histograms. The characteristic decay peak energies are shown in the staggered X axis.

8.7.3 Resolution

The energy resolution of a detector is a measure of its performance with respect to the resolution of fine detail in the deposited energy.

It is conventionally given as the dimensionless quantity

$$R = \frac{\text{FWHM}}{E}\,, \tag{8.6}$$

where E is the mean, and FWHM the full-width at half-maximum (FWHM) of the energy distribution. Given Eqn. 8.4, this expands to

$$R = 2\sqrt{\frac{2\ln 2Fk}{E}}\,, \tag{8.7}$$

where F the Fano factor, and k is the mean number of charge carriers liberated per unit energy deposited in the detector. Experimental measurements have found k to be approximately 3.6 eV for alpha particles in silicon [97][99][98].

Whilst the majority of the calibrated detectors were found to have a resolution of 70-120 keV at 3182.7 keV, it can be seen that some detectors exhibited clear signs of degraded resolution. The waveforms measured by these detectors were subsequently discarded.

8.8 Gas Simulation

8.8.1 Drift Velocity

The transformation from energy-time space of the GET waveforms to the energy-position space of the reconstructed charge clusters requires knowledge of the intrinsic drift velocity of the MicroMeGaS detector. The drift velocity is a function of the gas parameters (pressure, temperature, etc.) and the applied electric field within the drift chamber. The parameters for the experimental run described in this chapter are recorded in Table 8.2

Name	Value
Pressure	405 Torr
Temperature	290 K
Composition	96% ^{4}He, 4% CO_2
Drift Field	53.85 V/cm

Table 8.2: The gas and field parameters that govern charge transport within the active chamber

The phenomenon of electron diffusion within a gas is described by the Boltzmann transport equation (BTE), which models the phase-space evolution of a system in the absence of significant ionisation. Given the distribution function $f(\mathbf{r}, \mathbf{v}, t)$ describing the number density of electrons at the locus $(\mathbf{r}, \mathbf{v}, t)$ in phase space, the BTE may be expressed as

$$\frac{\partial f}{\partial t} + \frac{\partial f}{\partial \mathbf{r}} \cdot \dot{\mathbf{r}} + \frac{\partial f}{\partial \mathbf{v}} \cdot \dot{\mathbf{v}} = \left(\frac{\partial f}{\partial t}\right)_{\text{coll}}, \tag{8.8}$$

where $\left(\frac{\partial f}{\partial t}\right)_{\text{coll}}$ represents the rate of change due to collision processes.

Eqn. 8.8 is nontrivial to solve. Under the assumption of a constant uniform electric field, the distribution function may be reparametrised in terms of the electron energy ϵ and angle to the field θ such that f may be written as an expansion in the Legendre polynomials [100]. This simplification yields a set of coupled differential equations, which can be solved with numerical methods. An experimental code that numerically solves the BTE, Magboltz, is described in [100]. Recent versions of the Magboltz program solve the problem of electron transport with Monte Carlo techniques, which yields solutions whose accuracy depends solely upon computation time [101].

A Magboltz simulation was performed to compute the mean drift velocity within the TPC for the parameters given in Table 8.2. The drift velocity was found to be 5.302 μm/ns.

8.8.2 Cluster Resolution

Using a microscopic Monte Carlo model, the spread of the drift electrons can be determined as a function of displacement. The Garfield++ detector simulation toolkit includes a microscopic tracking component that consumes the per-scattering-process collision rates $\tau_i^{-1}(\epsilon)$ computed by Magboltz in order to evolve single electrons through time. An aggregate of these single-particle walks can be used to determine the distribution of the drift time and widths in the transverse and drift axes. Fig. 8.33 shows the transverse position resolution of electrons produced at a drift height of 5.90 cm.

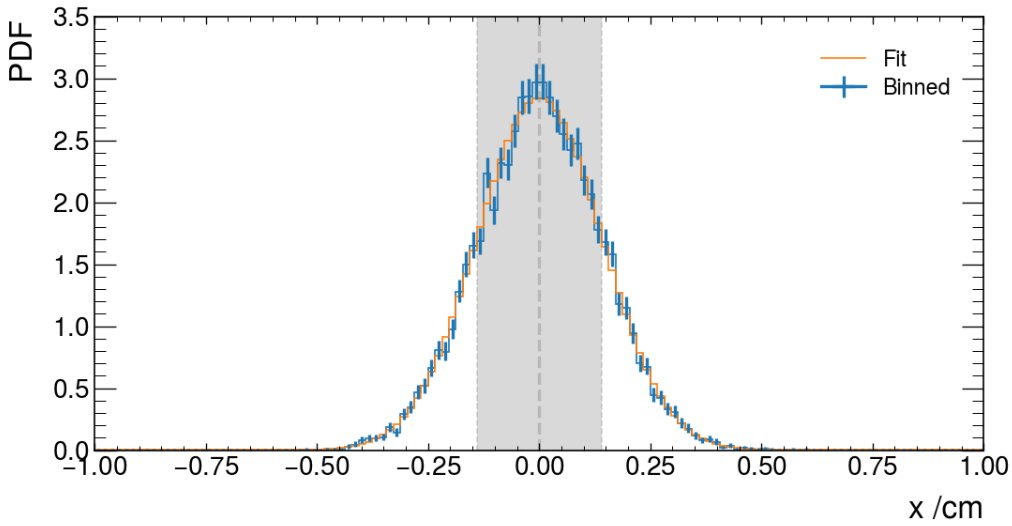


Fig. 8.33: The x position distribution of drifted electrons for y=5.90 cm fit with a normal distribution. The standard deviation region is shaded.

The observed width of a charge cluster resolved perpendicular to the drift plane is determined by two components; the intrinsic resolution (e.g. of the readout system), and the spread determined by transverse diffusion of the charge cluster:

$$\sigma^2_{\text{track}} = \sigma_0^2 + C_D^2 z\,, \tag{8.9}$$

where C_D is the transverse diffusion coefficient, z the drift distance, and σ_0 the intrinsic resolution [102].

A regression of the transverse variance against the drift height yields C_D of 0.016365 mm (see Fig. 8.34), with the median standard deviation of the charge cluster distribution within the conversion region of the TPC found to be on approximately 1 mm. This value is significantly smaller than the width of a MicroMeGaS pad (3.5 mm), which (subject to the initial cluster size) implies that tracks within the central region of the detector should predominantly comprise of single pads.

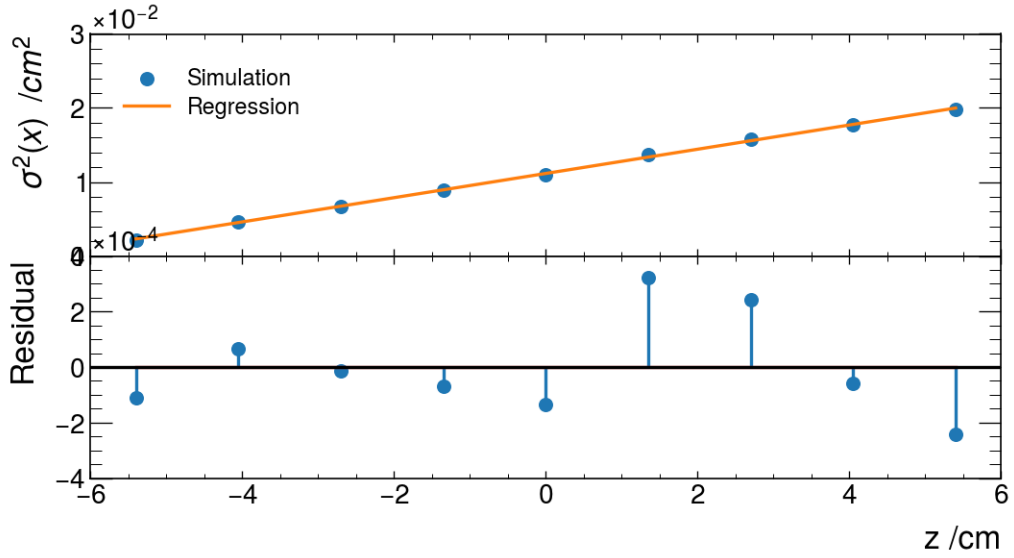


Fig. 8.34: The standard deviation of the final electron transverse position as a function of the drift distance. Against the scatter plot, a linear regression indicates the slope and intercept of the data.

8.9 Beam Origin

In the MicroMeGaS detector, the zero-point of the $\hat{z}$ coordinate is a computed property of both the detector geometry, and the readout electronics. The incoming beam is expected to be distributed according to some uncertainty about the $\hat{x} - \hat{y}$ plane, therefore a histogram of the time distribution in the early pads region of the MicroMeGaS should yield a symmetric distribution. The mode of this distribution represents the origin in the time domain (see Fig. 8.35).

The distribution observed in Fig. 8.35 was fit as the sum of a constant function, and a Gaussian distribution, using a binned negative log-likelihood optimiser. The linear contribution of the constant function corresponds to interactions that occur in the region prior to the MicroMeGaS. These interactions occur over a wide range of stops and over a range of angles, such that the charge clusters liberated by these tracks appear evenly distributed when viewed over a narrow slice. Despite some periodicity in the residual plot, indicating a mismatch between model and data, the fit shown in Fig. 8.36 demonstrates good agreement with the data and yields an origin of 350.8 cells.

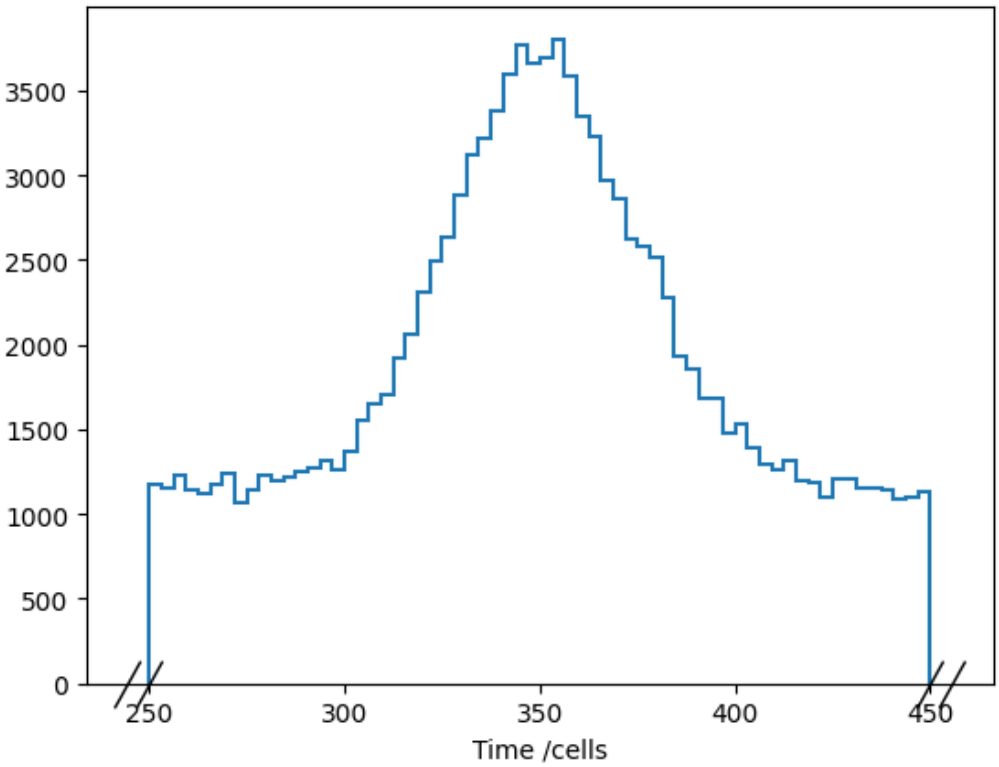


Fig. 8.35: Distribution of the charge-cluster time coordinate in the early pads region of the MicroMeGaS detector. The origin of this coordinate corresponds to the modal value, which lies around 350 cells.

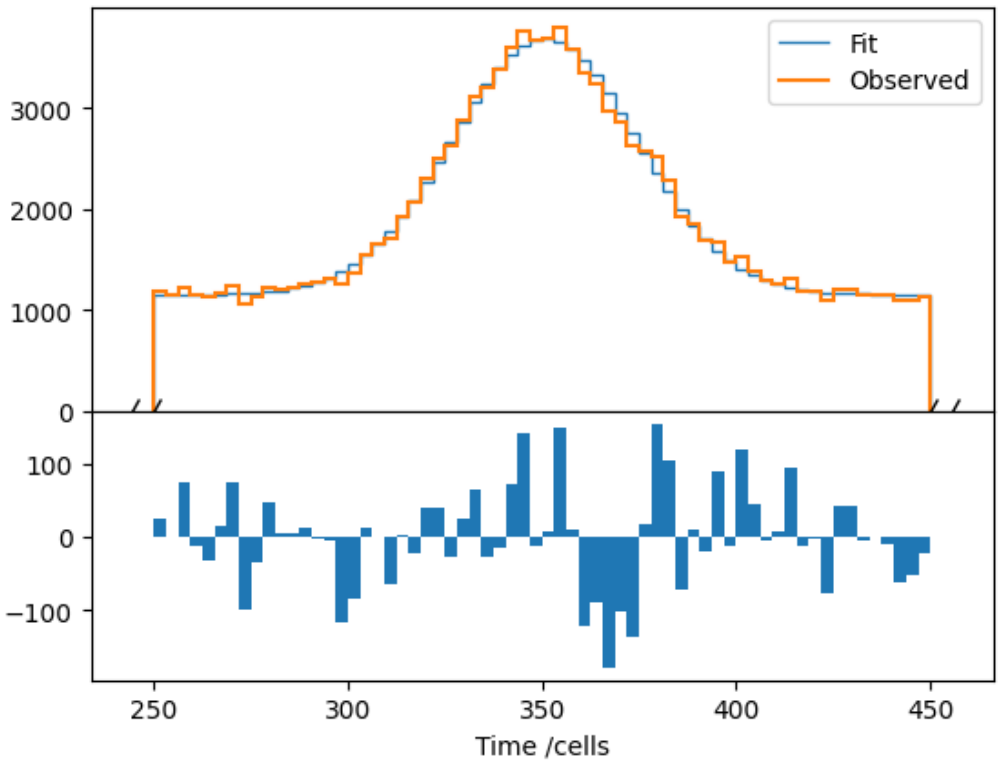


Fig. 8.36: Fit of the histogram in Fig. 8.35 using a constant function and Gaussian distribution. Above the x-axis, the original histogram and minimiser result are superimposed. Below, the residual distribution of the fit against the measurement is shown.

8.10 Cluster Reconstruction

As outlined in Section 6.2.1, the MicroMeGaS anode is highly segmented to give good position resolution measurements of the charge clusters produced by ionisation. In order to reconstruct the track trajectories in 3D, it is necessary to determine the height of the charge cluster. Given that the MicroMeGaS waveforms are energy-time signals, the height coordinate can be derived from the drift time, given by the locus of the peak in the measured waveform, and the drift velocity (see Section 8.8). The *origin* of this coordinate was discussed in Section 8.9.

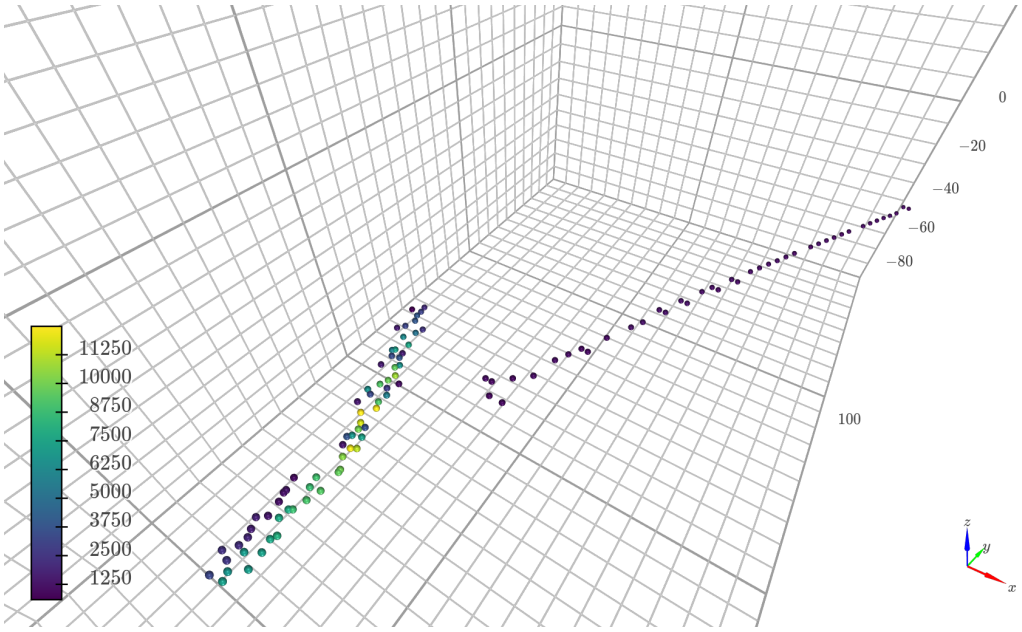


Fig. 8.37: 3D scatter plot of the cluster reconstruction taken from experiment. The charge of each cluster is indicated by its colour. Three distinct tracks can be identified; a heavier incoming and scatter track, and a light recoiling track. A discontinuity in the light recoil track immediately after the interaction vertex is clearly visible. This follows from the low stopping power of the corresponding ion, whereby the energy deposited in the MicroMeGaS is insufficient to trigger the low-gain channels. Once the ion enters the high-gain region it forms a coherent track.

8.10.1 Central Region

Within the central "pads" region, the anode position of a waveform is uniquely determined by the channel from which it was measured (see Fig. 8.38). Yet, the resolution in this region is diminished in the horizontal (x) axis; each pad has a width of 3.5 mm, and a length of 1.75 mm. As such, to avoid biasing subsequent fitting routines with aliasing artefacts yielded by taking the exact pad location, a uniform random offset was added to each (x, y) coordinate .

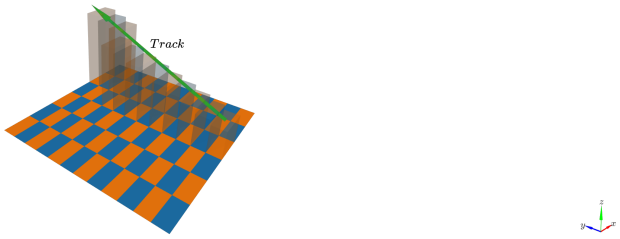


Fig. 8.38: A 3D diagram of the pad reconstruction procedure. The electrical contacts that comprise the MicroMeGaS anode are represented with a plane of opaque quadrilaterals. The shaded volume above these contacts indicates the measured *drift time* of the charge cluster incident upon the element. The volume above the contact upon which the current signal is induced is shaded opaque.

8.10.2 Side Region

Meanwhile, for the multiplexing side regions, only *one* coordinate is uniquely defined: the column for the *chains*, and the row for the *strips*. In order to localise a waveform in these regions, *another* waveform must be recorded in the complementary family (strip-chain or chain-strip) of channels such that *both* coordinates are uniquely determined.

The set of valid positions P defined for the side regions is naively given by the Cartesian product of the strips S and chains C, i.e.

$$P = S \times C \,. \tag{8.10}$$

Among these ordered pairs, only a small subset P_{valid} provides a reasonable interpretation of the observed ion path in the detector. In order to identify this subset, an additional discriminator is required.

For tracks with a significant inclination to the drift plane, we can use the drift *time* to order strip-chain pairings by their vertical separation. In the extreme case of a track that lies along the drift axis, the drift time is an exact proxy for the position of the charge distribution along the track (see Fig. 8.39). It follows that for some deviation from this ideal case, taking the strip-chain pairs with the smallest time difference will yield a reasonable set of charge clusters.

Meanwhile, the charge distributions formed along a track have a non-zero width and relative position to the average track direction; the electrons liberated by scattering of the incident ion arrest over a finite distance within the gas. There is therefore a *minimum* time-separation below which a “reasonable set” of pair can not be discriminated from those which arise from random correlations. In the extreme case, a track that lies exactly within the drift plane will yield strip-chain pairs that are entirely separated in time by the uncertainty on the cluster position, and this method cannot be used (see

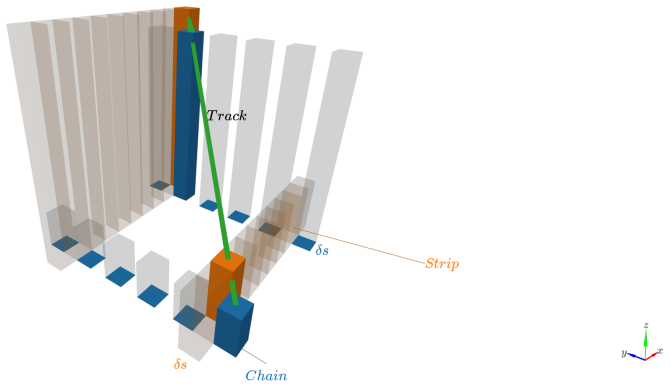


Fig. 8.39: A 3D diagram of the strip-chain pairing procedure. The electrical contacts that comprise the MicroMeGaS anode are represented with a plane of opaque quadrilaterals. The shaded volume above these contacts indicates the measured *drift time* of the charge cluster incident upon the element (strip/chain). The volume above the contact upon which the current signal is induced is shaded opaque, whilst the remaining (electrically connected) contacts are rendered partially transparent. It can be seen that the pairs with the closest drift times correspond to the elements over which the track passes.

Fig. 8.40).

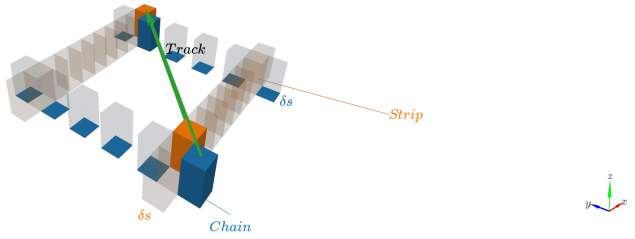


Fig. 8.40: A 3D diagram of the strip-chain pairing procedure for planar tracks. See Fig. 8.39 for more details. It can be seen that the pairs with the closest drift times no longer correspond to the elements over which the track passes.

To discern the range over which time-separation holds as a valuable reconstruction method, consider a track with the gradient $\frac{\mathrm{d}z}{\mathrm{d}s}$ at 45° to the element stacking axis, where *element* refers to a strip or chain. Any two elements (strip-strip or chain-chain) are separated by a fixed spacing δs. The separation between an adjacent hit on a strip-chain pair is therefore $\sqrt{2\left(\delta s\right)^2}$. The change in height between the members of the pair consequently given by $\delta z = \frac{\mathrm{d}z}{\mathrm{d}s} \cdot \sqrt{2}\delta s$. For some intrinsic width σ of the charge distribution about the track, it follows that the resolving limit is

$$\frac{\mathrm{d}z}{\mathrm{d}s} \geq \frac{\sigma}{\sqrt{2}\delta s} \tag{8.11}$$

In-plane Events

A fundamental limitation to the strip-chain time-matching approach detailed above is observed for events in which all strip-chain times lie beneath the resolving limit. For these events, the set of possible matches grows to encompass the Cartesian product of the elements. It follows that an alternative approach must be taken to recover the track information for these events. Under the assumption of a single, linear track (which holds reasonably well in the side regions), a pair of vertices can be drawn from the innermost and outermost pairs of strips and chains. These vertices correspond to the conjectured start and end positions of a forward-focussed track that originates from the central pads region. One can then interpolate along this idealised track to recover the charge clusters.

8.11 Track Fitting

Tracks within the TPC volume were fit with line-interval models (see Section 7.2.5) using the methods described in Section 7.2.4. From these fits, a set of labels was generated, which permits ascribing each cluster to a given track model. Fig. 8.41 shows an event that was fit according to these techniques.

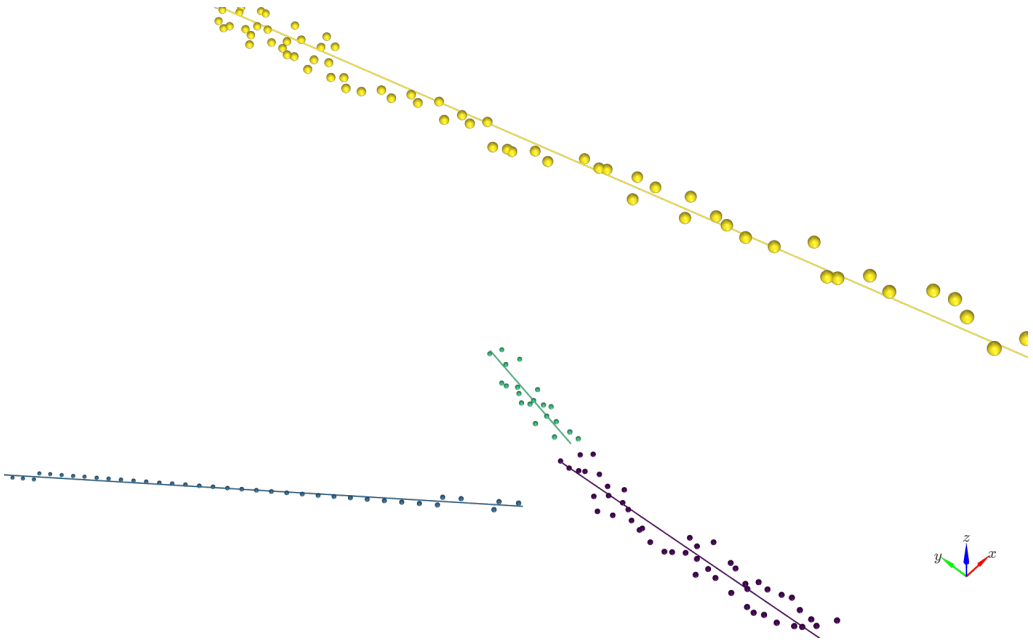


Fig. 8.41: A 3D plot of the track reconstruction for a single event using the PEARL method described in Section 7.2.4. Each distinct track is labelled with a distinct colour, and the line of best fit superimposed. Due to the uncertainty of the cluster positions in the pads region, the beam track is found to be longer than it is physically expected to be. This can be corrected in a subsequent constrained fit. The upper track corresponds to a beam-like pile-up track, whilst the lower track triplet belong to a single scattering reaction.

8.11.1 Track Multiplicity

In Section 7.2, the PeARL method for multi-model discovery and fitting was discussed. For this experiment, a random sample of 1024 models was drawn from pairwise combinations of the clusters computed in Section 8.10. Each model comprised of a pair of clusters, with corresponding positions and charges, and assigned a characteristic width drawn from a Gaussian distribution whose width was determined by inspection. Fig. 8.42 shows the track multiplicity of these fits, i.e. the number of tracks per event. Perfect scattering reactions correspond to three-track events, in which the beam, light product recoil, and heavy reactant scatter are observed. However, these events are in the minority; interactions which occur near (or before) the start of the MicroMeGaS are only partially captured by the detector; the beam track falls outside of the active region of the TPC. Similarly, the heavy product track is often too short to be identified from the resolution of the detector. As such, the majority of events are fewer-than-three-track events. Beyond-three-track events typically correspond to coherent noise, which is reconstructed as ghost tracks, or pileup, in which there are demonstrably greater than three tracks present in the event.

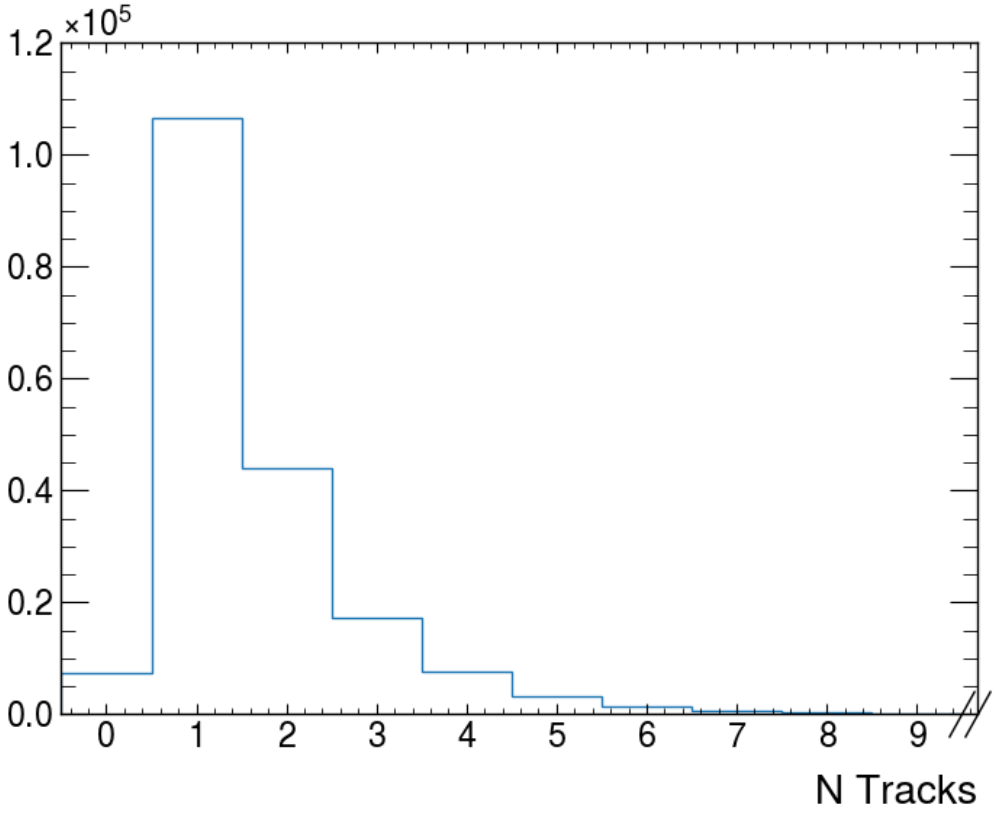


Fig. 8.42: Histogram of the track multiplicity found during fitting of the usable experimental runs. A majority of events have a single track; these correspond to events in which the light product is not observed.

These track fits are unconstrained; in the preliminary fitting procedure there is no imposition of reaction kinematics upon

the result. Though it might seem prudent to pursue only those solutions which satisfy kinematic constraints, at this early juncture in the analysis there is insufficient information to introduce these constraints; in order to apply the rules of energy and momentum conservation, each track must be identified in terms of the particle kind, and the reaction to which it belongs. This requires that the fitting process be diphasic; a primary track finding routine, followed by a kinematic fitting procedure (see Section 7.3).

8.11.2 Silicon Projection

For each event, with the trigger placed upon the silicon detector channels, it follows that there should be a track associated with a low-z ion incident upon the detector. Therefore, one can establish a 2D histogram corresponding to the projection of these tracks upon the triggered silicon detector. Without any gating logic, it should follow that the these intersections agree closely to the known geometrical dimensions of each quadrant. Fig. 8.43 shows a pair of hitmaps for the forward silicon detector array. Due to the high incidence of proton tracks in the central region, it is necessary to ignore the central detectors in order to best visualise the quality of the hitmap in other regions. These protons originate from the production reaction of the ^{10}C secondary beam, whereby one of the available exit channels is the elastic scattering of the ^{10}B primary beam with the H_2 gas-cell ions.

A plot of the *relative* hitmap formed by taking the difference of the reconstructed hit position with the known silicon centroid is shown in Fig. 8.44 to have good agreement with the quadrant dimensions. It is expected that the agreement will not be exact; small tracks incur a larger angular error, which manifests as significant displacements in the reconstructed hit position. In both figures, the hitmap is formed by the union of all track intersections in order to remove bias from gating the track hit position. A position gate was applied to the track endpoint; tracks whose endpoints lie within the final sector of the MicroMeGaS are likely to be the light particles that reach the silicon; the beam and scattered beam tracks are typically found earlier in the detector.

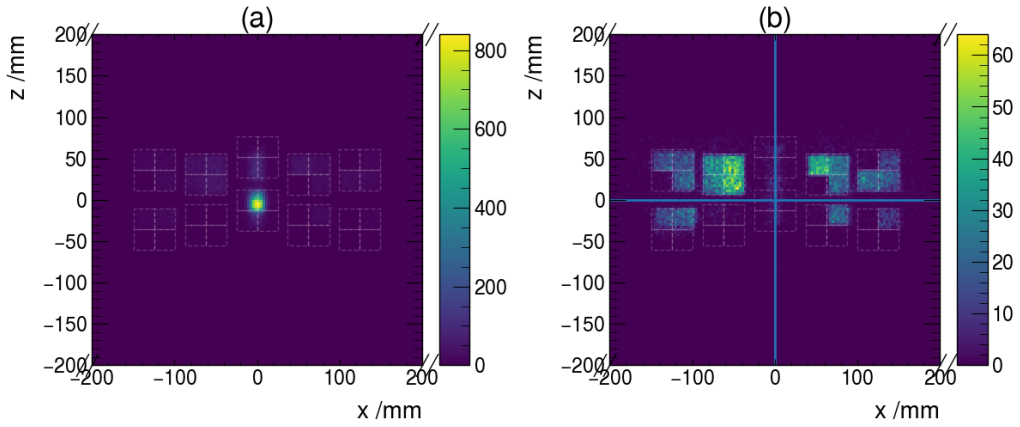


Fig. 8.43: 2D hitmaps for the forward silicon detector array, formed by the intersection of reconstructed tracks, whose endpoints lie in the final sector of the MicroMeGaS, with the silicon plane. The hitmap for all events is shown in (a). The subset of hits outside the central region is plotted in (b).

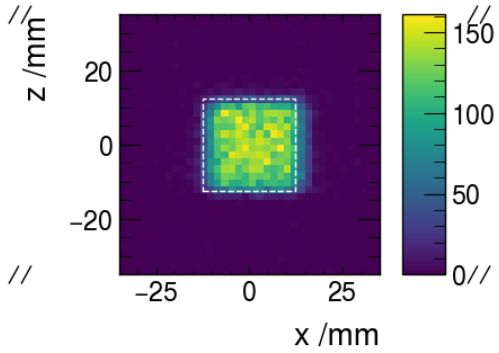


Fig. 8.44: 2D relative hitmaps for the forward silicon detector array, formed by the intersection of reconstructed tracks, whose endpoints lie in the final sector of the MicroMeGaS, with the silicon plane. The position of each reconstructed hit is plotted relative to the detector centroid.

8.12 Particle Identification

8.12.1 Beam Tracks

The majority of signals formed in the MicroMeGaS belong to charge clusters liberated by the passage of ^{10}C beam ions within the TPC. Given the linear response between the signal formed on the anode and the energy deposited by the beam, it is possible to correlate the energy loss curve of the central MicroMeGaS pads region with a known stopping power curve of ^{10}C within the TPC gas mixture.

A 2D histogram of the measured charge against the pad row was computed across all events (see Fig. 8.45). Superimposed over this distribution is the stopping power for the beam ions in the active target, indicated by the dashed line, as simulated by SRIM [72] for a beam energy given by a fit parameter (30.5 MeV). There are two loci clearly visible within the histogram, one of which lies at much lower charges than the other. The primary locus, that of the beam, coincides with the predicted stopping power curve, which has been displaced along the y axis in order to treat the initial beam energy as an unknown. The Bragg peak, corresponding to the maximum stopping power of the ^{10}C ions, is indicated by the dotted line at row ~60. It follows that the beam is fully stopped in the active target, and any hits measured in the central silicon detectors originate from lighter particles which have sufficient energy to escape the chamber. The secondary locus, at near constant charge of 1500 units, corresponds to the scattered light alpha-particles and ^{10}C ions produced by reactions before the sensitive region of the MicroMeGaS anode. The far pads region, has been omitted due to the high-gain pads region located in the last 16 rows of pads. These pads are sensitive to both 180°-scattered alpha particles and low stopping-power protons. Meanwhile, the beam is fully stopped well before this region.

By approximately fitting the stopping power curve to the arbitrary charge units seen in Fig. 8.45, a rudimentary calibration of the MicroMeGaS gain can be achieved yielding $\sim 1.219 \times 10^{-1}$ keVcm^{-1} in the low-gain region.

The same distribution can be fit using stopping powers predicted by MSTAR [103] (see Fig. 8.46). The predicted after-window beam energy is 29.7 MeV.

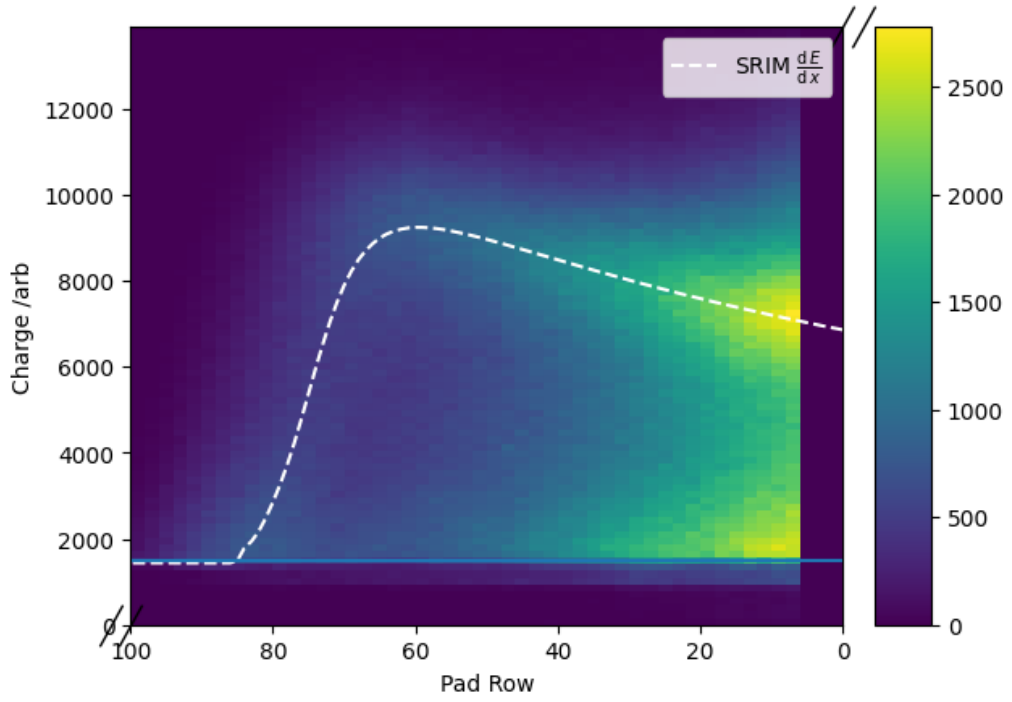


Fig. 8.45: Experimental distribution of the stopping power of ^{10}C ions in a gas-mixture of 96% ^{4}He, 4% CO_2 held at 405 torr. SRIM prediction for an after-window beam energy of 30.5 MeV indicated by the dashed line.

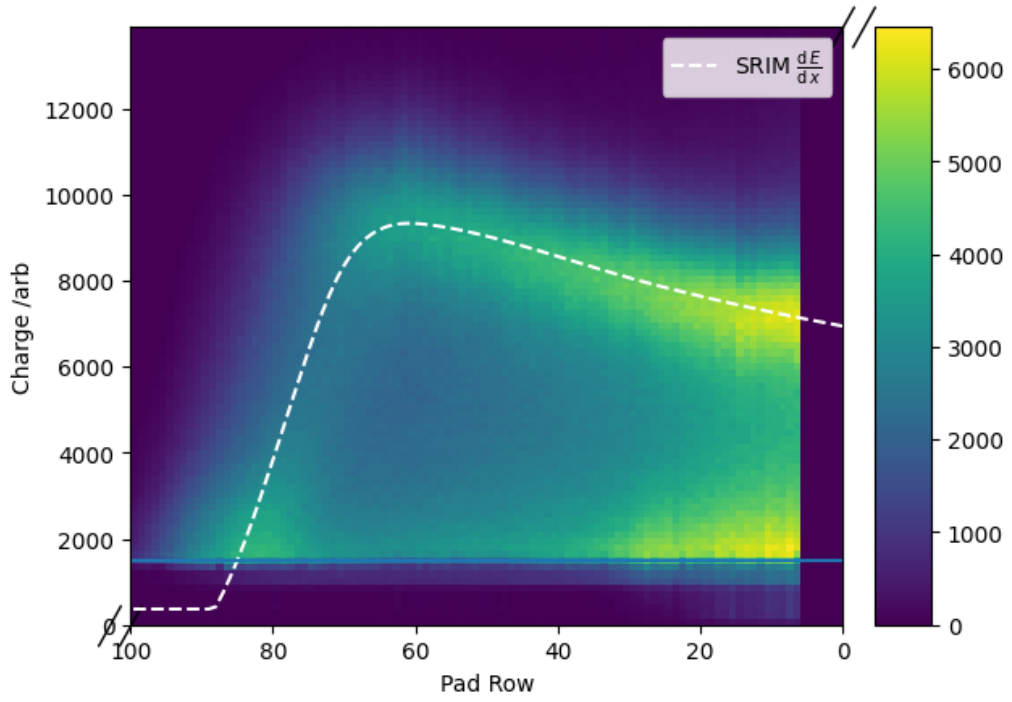


Fig. 8.46: Experimental distribution of the stopping power of ^{10}C ions in a gas-mixture of 96% ^{4}He, 4% CO_2 held at 405 torr. MSTAR prediction for an after-window beam energy of 29.7 MeV indicated by the dashed line.

8.12.2 Separating Light-Product Tracks

The primary set of particles observed in the TPC are ^{10}C, ^{1}H, and ^{4}He. Each has a characteristic stopping power curve, as determined by SRIM (see Fig. 8.47). Notably, the protons are easily distinguished from the beam ions, which typically vary by 1–2 orders of magnitude over the active region of the MicroMeGaS. Meanwhile, the stopping power of the ^{4}He ions lies nearly equidistant from the other two ions (in log space), and if the ion energies are not known, experimental uncertainty can render it difficult to distinguish between the two ions simply from the energy deposition within the TPC.

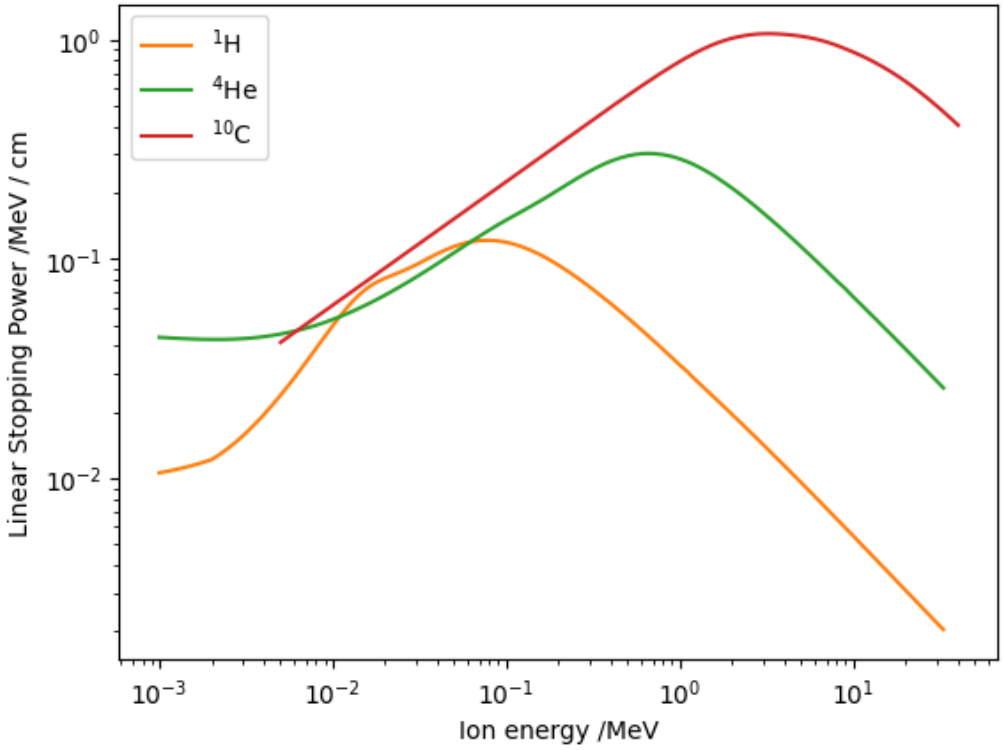


Fig. 8.47: Stopping power curves predicted by SRIM for ^{1}H, ^{4}He, and ^{10}C ions in a gas-mixture of 96% ^{4}He, 4% CO_2 with density ρ. Each curve overlaps with the others for a given ion energy.

In order to identify the ions that comprise the tracks found in Section 8.11, the MicroMeGas was used in tandem with the silicon detector array to establish a particle telescope (see Fig. 8.48). Conventionally, a particle telescope is established with two silicon detectors placed in close succession. In the TexAT detector, only a single "thick" silicon array was available (and chosen to prevent punch-through of the alpha particles). Therefore, in order to establish the ΔE component of the telescope, the MicroMeGas itself was used. Following from Eqn. 5.11, the energy deposited by an ion within the active region of the MicroMeGaS is approximately proportional to $\frac{1}{E}$, such that a plot of the measured energy E against the charge deposited within the final region of the MicroMeGaS should yield a locus unique to the z of the ion. A series of histograms plotting these two quantities for each silicon detector (and the associated region of the MicroMeGaS) were

produced, and clearly demonstrate such a phenomenon. See Fig. 8.49 for an example of such a histogram.

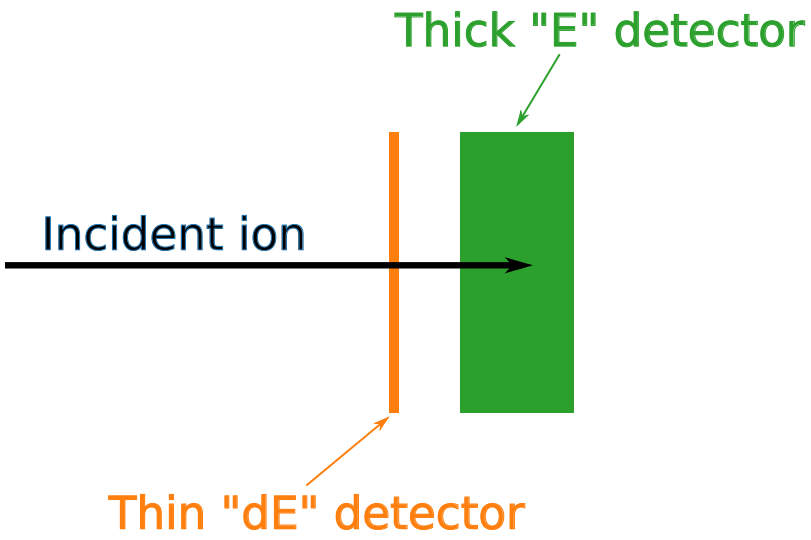


Fig. 8.48: Illustration of a particle telescope. A thin ΔE detector measures the z-dependent energy loss of the incident radiation within a thin silicon target. The ion punches through this material, and is fully stopped within the "thick" silicon detector. Conventionally, the energy lost in both detectors is taken to be the total energy of the particle. However, to first approximation the contribution lost in the thin detector can be ignored.

The $\frac{\mathrm{d}E}{\mathrm{d}x}$—$E$ plot is useful in discriminating between alpha particles and the lighter protons, but is only applicable in the case that a given track interacts with the silicon detector. For any other track, it is not possible to use this approach to identify the particle. Furthermore, the presence of pile-up events, which superimpose in the drift plane, can frustrate efforts to use the MicroMeGaS as a "thin" detector (see Fig. 8.48) due to the difficulty in disambiguating the charge clusters of separate tracks. This phenomenon is confined predominantly to the central-angle detectors; tracks in the high-gain side regions are typically more well reconstructed.

For non-backscatter (lab) events, such as those observed in inverse kinematics, track identities can be assumed by the process of elimination; identifying a vertex using the light-product track and the ideal beam axis makes it possible to partition the anode into forward and backward regions. Near co-linear tracks in the backward region (towards the ion counter) are likely to be the beam (or pile-up). Remaining tracks are highly likely to be the scattered beam ion.

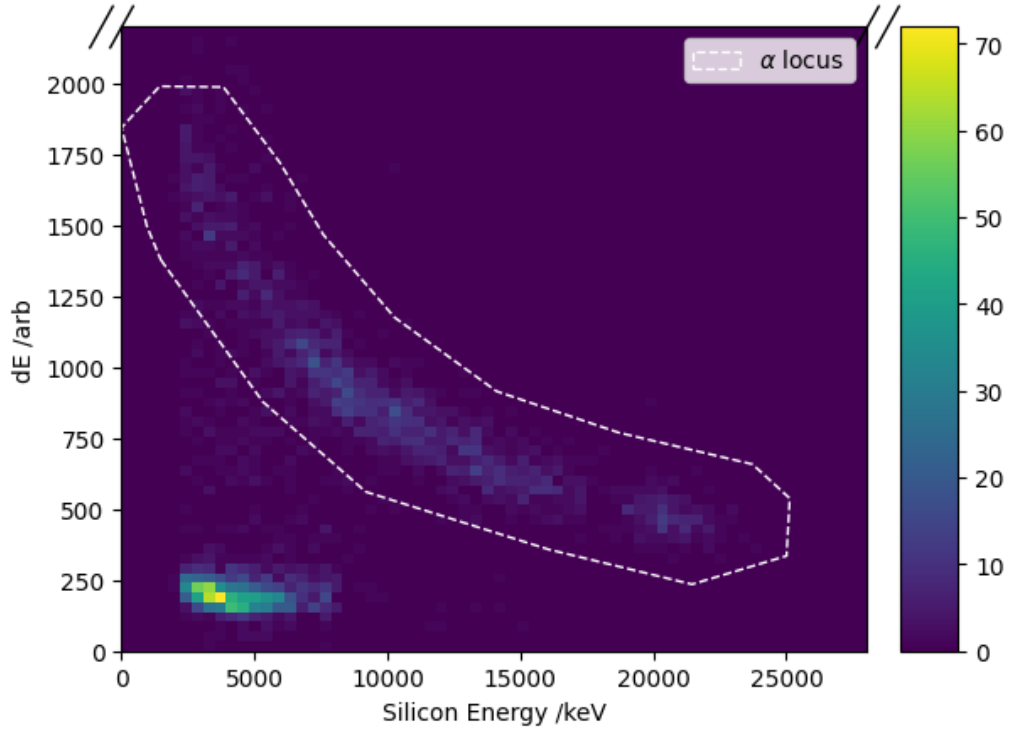


Fig. 8.49: 2D histogram of the energy deposited in an individual silicon detector against the energy lost in the final N elements of the MicroMeGaS detector. Two loci are visible; the proton region at (5000, 250), and an upper alpha curve.

8.13 Kinematics

8.13.1 Angular Scattering

Fundamentally, in the sub-relativistic regime, the following conservation laws typically form the basis of any kinematic reconstruction of an elastic interaction

$$\begin{aligned} \sum_j \mathbf{p_i^j} &= \sum_j \mathbf{p_f^j} \\ \sum_j E_i^j &= \sum_j E_f^j \,. \end{aligned} \tag{8.12}$$

In the context of a scattering reaction between an incident heavy (larger A, Z) ion and a stationary light (smaller A, Z) target (gas), the reaction will be forward projected and the range of the scattered light ion will be significantly greater than that of the recoiling heavy ion. As such, the tracks left by light ions moving through the TPC are much more easily identified, and may be fit with greater precision than fits performed using the heavy recoil.

Given the energy and angular straggling of the beam following interactions with the window, it is also evident that indirect measurements of the beam energy will be less accurate than those of the light ion after it interacts with the silicon detector

array. As such, the most accurate measurements of the reaction kinematics are given by the track scattering angles and known energy of the light product. In such a configuration, the beam and recoil momentum vectors are initially unknown. We can define a beam-centric coordinate system (see Fig. 8.50) with the $\hat{1}$ vector given by the beam direction, the $\hat{3}$ vector given by the normalised cross product of the beam and scatter tracks, and the $\hat{2}$ vector taken to be perpendicular to $\hat{1}$ and $\hat{3}$.

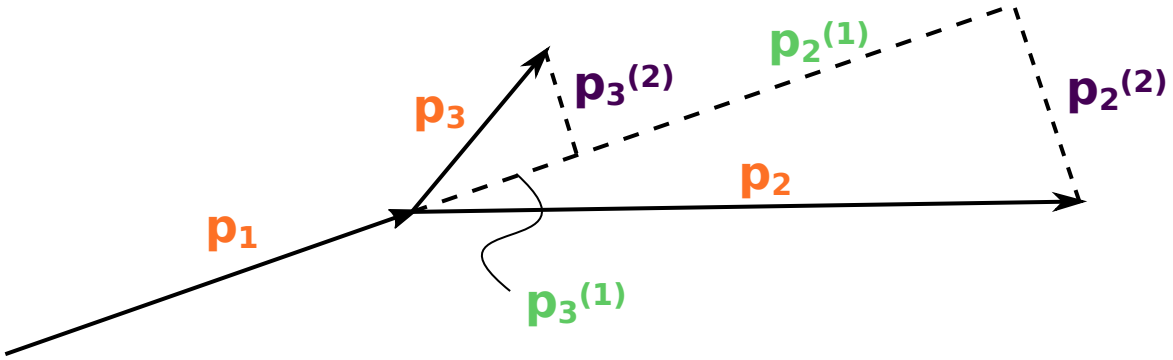


Fig. 8.50: Schematic diagram of a static-target scattering reaction in a beam-centric coordinate system. The $\hat{1}$ vector is given by the beam direction, the $\hat{3}$ vector given by the reaction plane normal, and the $\hat{2}$ vector taken to be perpendicular to $\hat{1}$ and $\hat{3}$. Each particle $\mathbf{p}_i$ can be resolved along these canonical axis vectors. Each vector $\mathbf{p}_i$ indicates the momentum of the reaction participant i, whose components $\mathbf{p}_i^{(j)}$ have been resolved along the beam-centric coordinate system vectors $\hat{j}$.

From Eqn. 8.12, conservation of momentum can be applied separately to each axis such that we have

$$\begin{aligned} p_1 &= p_2^{(1)} + p_3^{(1)} \\ 0 &= p_2^{(2)} + p_3^{(2)} \,. \end{aligned} \tag{8.13}$$

An application of energy conservation yields

$$\frac{p_1^2}{2m_1} = \frac{p_2^2}{2m_2} + \frac{p_3^2}{2m_3} + E_{\text{ex}} \,, \tag{8.14}$$

where E_{ex} is the total excitation energy of particles 2 and 3. For elastic scattering, $E_{\text{ex}} = 0$. It can be shown that if $m_1 = m_3$, the recoil momentum is uniquely determined by

$$\begin{aligned} p_3^{(1)} &= \frac{\left(\left(p_2^{(2)}\right)^2 + \left(p_2^{(1)}\right)^2\right)\frac{m_1}{m_2} + \left(p_2^{(2)}\right)^2 - \left(p_2^{(1)}\right)^2}{2p_2^{(1)}} \\ &= \frac{\left(p_2^{(2)}\right)^2(m_1 + m_2) + \left(p_2^{(1)}\right)^2(m_1 - m_2)}{2p_2^{(1)}m_2} \,, \end{aligned} \tag{8.15}$$

ultimately yielding an expression for the beam energy in terms of the reaction products.

Fig. 8.51 shows a histogram of the observed silicon energy as a function of the reconstructed event vertex position. The

vertex is computed from the tracks fit in Section 8.11, using the closest point of approach between the beam and the light recoil tracks. In the event that one does not observe the beam track directly, an idealised track can be assumed in which the beam is aligned along the forward ($+\hat{y}$) vector. For these tracks, the vertex of the interaction is given by taking the closest point of approach between the ideal beam track and the track of the observed light-recoil.

Two clear loci are visible, corresponding to the elastic (upper) and inelastic (lower) scattering channels; with a portion of the available centre-of-mass energy partitioned into formation of an excited state in the (heavy) scattered beam nucleus, there is less kinetic energy available to the products of inelastic scattering reactions.

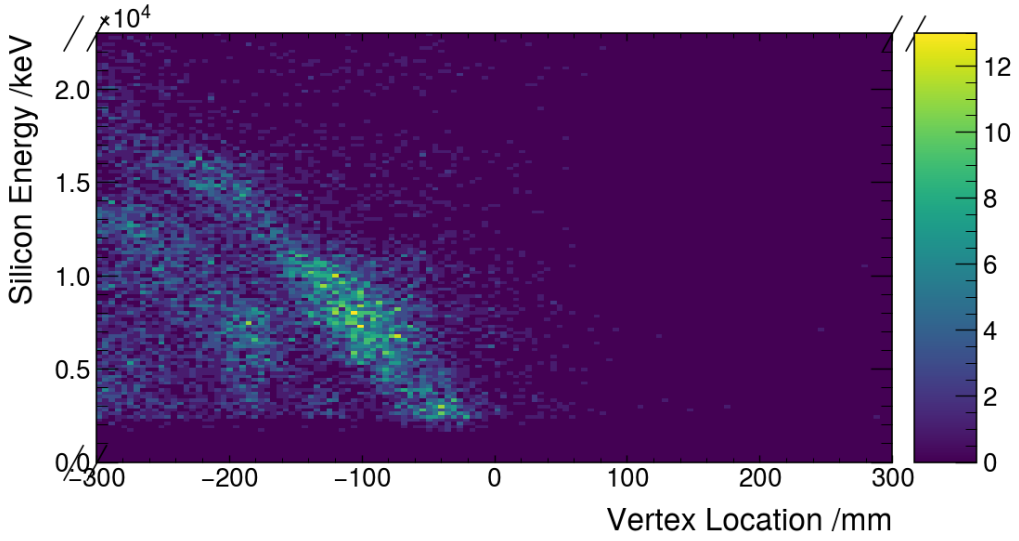


Fig. 8.51: A plot of the vertex position in non zero-degree scattering reactions as a function of energy measured in the silicon detectors. A clear locus of bins corresponding to elastic reactions can be observed, with a maximum near (-100, 7500).

This approach necessarily requires the knowledge of the interaction vertex (to identify the light particle energy) and the relative angle of the beam to the scattered ion track. If both of these cannot be determined, then there is insufficient knowledge of the system to find a solution. In the central detectors, in which the light product track is close to zero degrees, it is possible to leverage the assumption of a linear reaction in order to once-again determine the beam energy (see Section 8.13.2).

8.13.2 Zero-degree Scattering

Unlike the method described in Section 8.13.1, reactions in the zero-degree regime are unsuitable for reconstruction using track fit information; the resolution of the TexAT MicroMeGaS in the drift plane is too poor to appropriately identify the reaction vertex in the majority of these interactions. Reactions confined to the drift axis are more amenable to the track-fit method; the temporal resolution of the MicroMeGaS is superior to the anode spatial resolution, but only a subset of events are suitable for this kind of reconstruction.

An alternative reconstruction approach leverages the approximation that angular information can be neglected in the extreme zero-degree (lab frame) scattering regime. Under this assumption, Eqn. 8.15 simplifies to

$$p_3 = \frac{p_2(m_1 - m_2)}{2m_2}\,. \tag{8.16}$$

This in turn admits a relationship between the beam and scatter energies

$$\frac{E_2}{E_1} = \frac{4m_1m_2}{(m_1 + m_2)^2}\,. \tag{8.17}$$

In this regime, the reaction vertex (and hence kinematics) is given by the point at which a separate model of the beam energy, such as that described in Section 8.13.4, agrees with that given by Eqn. 8.17. Yet, it was observed that the beam energy as predicted by direct and indirect methods did not show strong agreement and thus choice of a particular method introduces significant error in the reconstructed kinematics.

The infeasibility of using a directly-reconstructed beam-energy model to identify the reaction vertex derives from from the sensitivity of energy loss calculations upon the path-length and stopping-power model of beam ion. For high-Z ions and large path lengths, this uncertainty dominates all other measures of uncertainty (such as the resolution of the silicon detectors and appropriateness of the zero-degree assumption).

An alternative reconstruction approach is to identify the vertex of the reaction using the position of the Bragg peak of the scattering projectile. Knowledge of the Bragg peak location provides a measure of the heavy-product energy at that position. The interaction vertex is then defined by the position at which the heavy-product energy given by Eqn. 8.17 agrees with that which yields the Bragg peak at the observed location. Such an approach requires precise tabulation of the low-energy stopping powers for the projectile, unlike the methodology outlined in Section 8.13.1 which requires only the stopping powers of the light-product. Furthermore, and prohibiting the use of such a method in this analysis, it must be possible to identify the Bragg peak within the active volume. For these data, the kinematics given by simulated stopping powers (using both SRIM and MSTAR) do not admit the possibility of observing the Bragg peak of the scattered beam; for an elastic scattering reaction to have sufficient energy to be measurable by the silicon detectors, it is found that the

heavy-product must stop before the start of the MicroMeGaS.

With Eqn. 8.17, the beam energy is uniquely determined for elastic scattering by the energy of the light ion at the reaction vertex. If the beam energy is approximately known, it is possible to identify the vertex at the intersection between the curves that describe the beam energy as reconstructed from the initial beam energy and the energy deposited in the silicon detectors. See Fig. 8.52 for a plot of the vertex position as a function of silicon energy as reconstructed by this method.

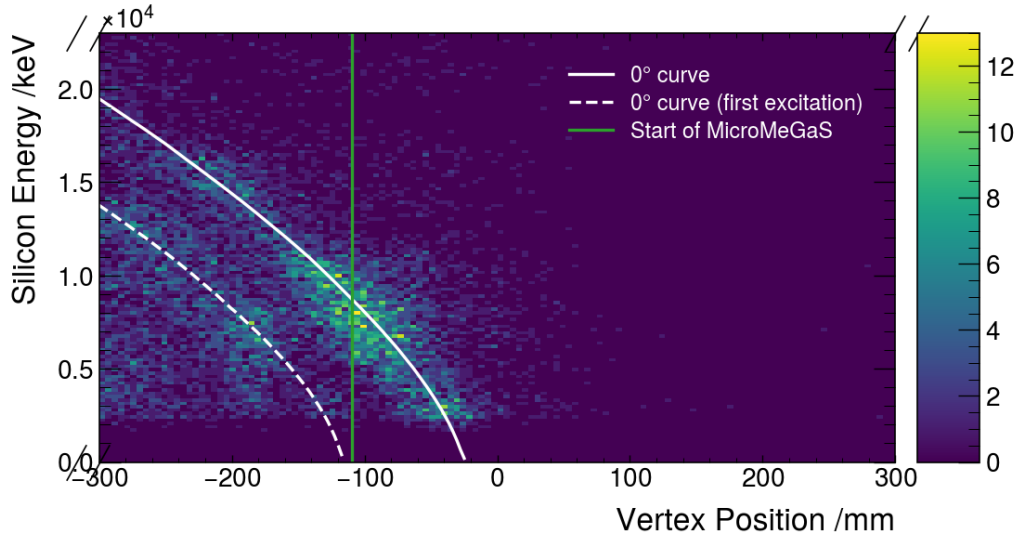


Fig. 8.52: A plot of the expected vertex position in zero-degree scattering as a function of energy measured in the central silicon detectors. The two curves show the interaction vertex for (solid line) elastic and (dashed line) inelastic scatters that produces the corresponding silicon hit for a known beam energy given by Section 8.13.4.

For *inelastic* scatters, it is anticipated that the measured silicon energy be reduced with respect to the elastic curve for a given vertex position. This describes the loci below the elastic curve shown in Fig. 8.52. By treating the excitation of the heavy product, it can be shown that Eqn. 8.17 is the ground-state simplification of the following relation:

$$\frac{E_2}{E_1} = \frac{4m_1m_2}{\left(m_1\left(1+\frac{E^*}{E_2}\right)+m_2\right)^2}\,, \tag{8.18}$$

where E^* is the excitation energy of the scattered heavy-ion.

Without additional information such as the Bragg peak to disambiguate between elastic and inelastic loci, it is not possible with these data to identify the subset which corresponds to elastic interactions; for any event with an observed energy deposit in the central silicon detector, either an elastic scattering reaction at vertex v_1 took place, or an inelastic reaction at vertex v_2, where v_1 lies closer in the beam axis towards the silicon detector than v_2.

8.13.3 Excitation Functions

The calculation of a reaction Q-value is described in Section 5.2. For elastic scattering of ^{10}C ions upon a ^{4}He target, the Q-value is 10117 keV. This value is used to build the elastic non-zero-degree excitation function shown in Fig. 8.53, which is given by $E_{\text{ex}} = E_1 - E_{\text{CoM}} + Q = E_1 \frac{m_2}{m_1+m_2} + Q$. Without the ability to isolate the elastic and inelastic contributions to the excitation function, it is not possible to combine the data in Fig. 8.53 with those in Fig. 8.54 which plots the zero-degree counterpart. As such, Fig. 8.54 plots the excitation function derived under the assumption that these data were obtained via the elastic, and inelastic channels.

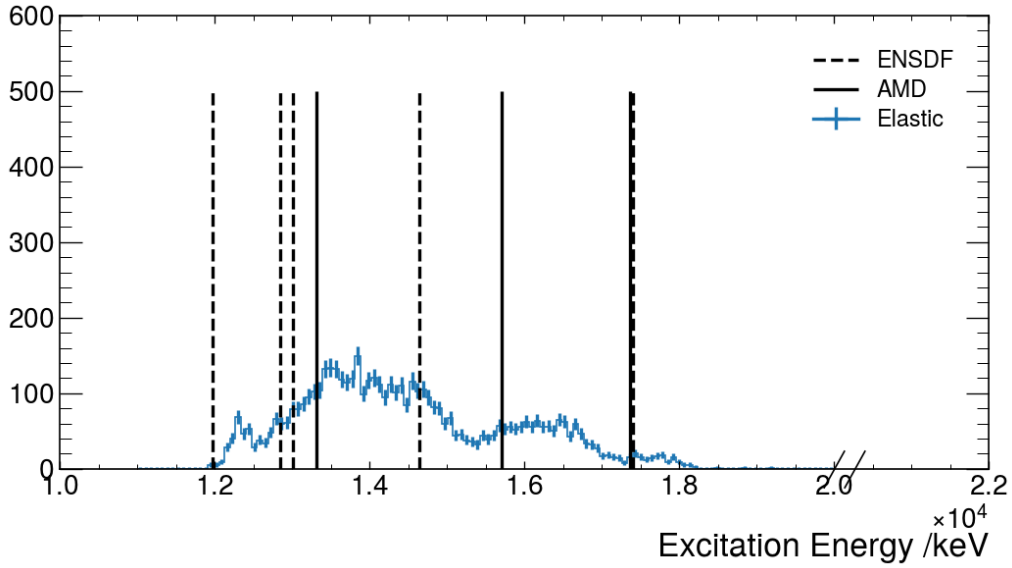


Fig. 8.53: A histogram of the excitation energy of the ^{14}O nucleus integrated over all non-zero-degree scattering angles. States published in the ENSDF database are plotted in dashed lines, whilst the various AMD predictions discussed in this work for ^{14}O are, where plausible, indicated by solid lines. Clear peak-like structures are visible, against a wider background distribution. Despite this, further enhancement of these data through secondary fitting is required to provide more conclusive agreement with predicted states.

Although the analysis work described in this chapter does not yield sufficient results to elucidate the structure of the ^{14}O nucleus, it can be seen that a number of peaks in the excitation function coincide with known or predicted energies. Yet, further analysis is still required, including kinematic fitting of the angular measurements, and elastic-inelastic subtraction, to make any deeper observations. The completion of this work is planned by the nuclear physics group at the University of Birmingham, which will identify the cross sections associated with the elastic scattering channel, and use them to infer the states present in ^{14}O.

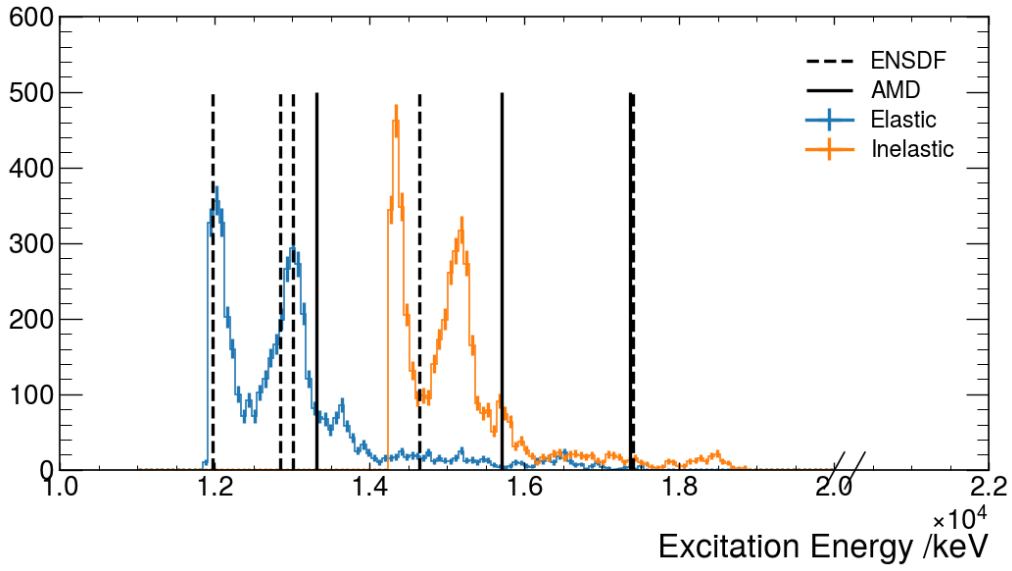


Fig. 8.54: A histogram of the excitation energy of the ^{14}O nucleus computed from the near-zero-degree scattering channel. States published in the ENSDF database are plotted in dashed lines, whilst the various AMD predictions discussed in this work for ^{14}O are, where plausible, indicated by solid lines.

8.13.4 Beam Energy Estimation

The beam energy after the Havar window is predicted by SRIM calculations to be ~25 MeV. A preliminary evaluation of this estimate can be performed using the maximum energy deposited in the central silicon detectors. This should correspond to a near zero-degree scatter from the farthest point in the active volume, i.e. the window (see Fig. 8.55). Given a maximum energy of 23 MeV, the after-window beam energy computed by the alpha stopping power is 30.576 MeV. This value is much larger than the initial prediction, but lies in close agreement with the fit parameter used to produce the theoretical stopping power curve in Fig. 8.45.

A more robust approach is to reconstruct the reaction kinematics in the non zero-degree regime, and plot the difference of the projectile energies computed by this reconstruction and directly by energy loss through the gas. Fig. 8.56 clearly shows a non-central distribution, indicating that the direct prediction lies below the reconstructed value by ~1 MeV, i.e. the true beam energy after the window is .

It is evident that these calculations are highly sensitive to the stopping power tabulation; the disagreement in Fig. 8.56 implies that the stopping powers used in this analysis over-predict the energy loss of the beam in the energy region of interest. There is preliminary experimental evidence to suggest that this is the case for both MSTAR and SRIM [104]. Future work should measure these stopping powers, both of ^{10}C in ^{4}He and ^{4}He in ^{4}He in order to reduce the discrepancy

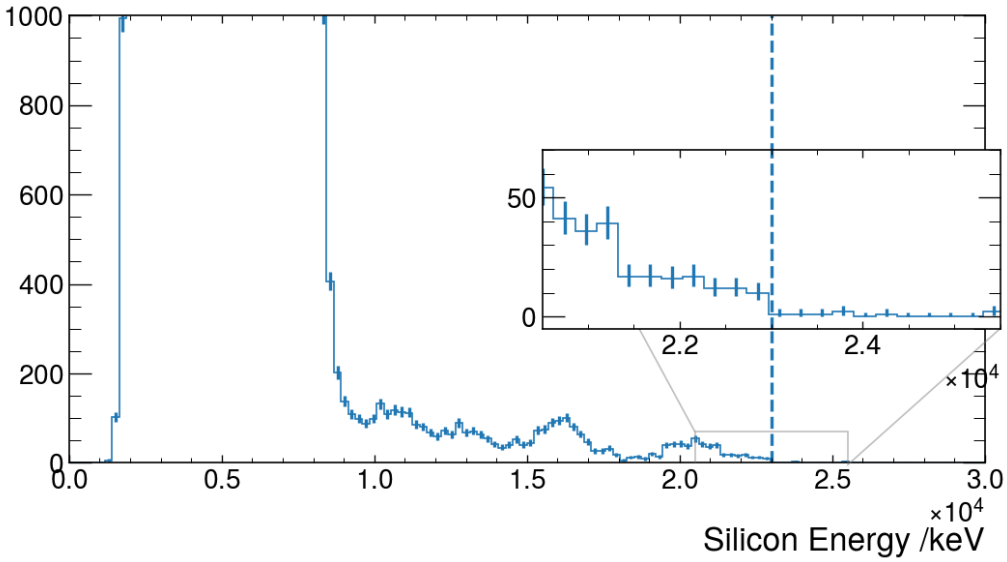


Fig. 8.55: A histogram of the silicon energy distribution, with inset axes indicating a candidate for the maximum value.

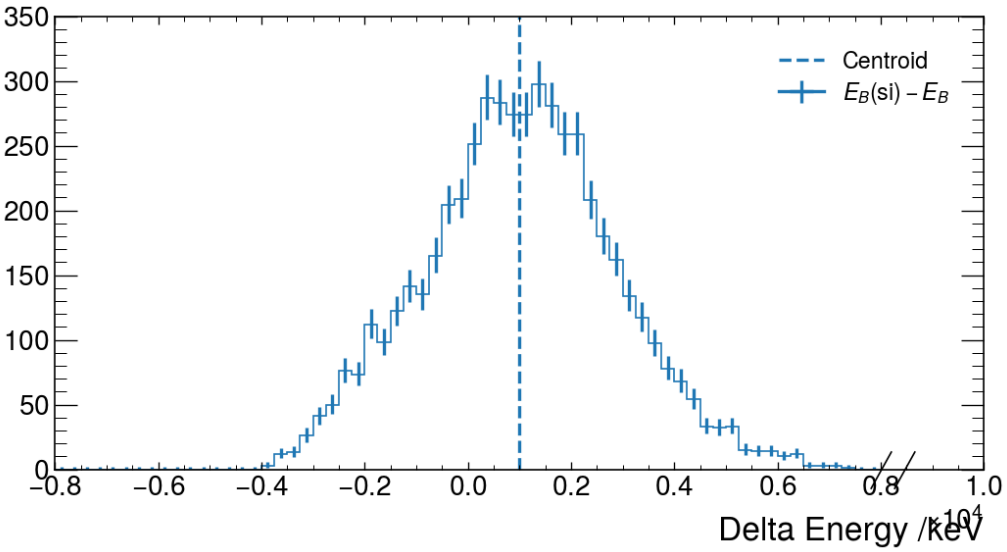


Fig. 8.56: A histogram of the difference in projectile energy computed directly (through SRIM calculations of energy loss in the gas), and via indirect kinematic reconstruction (using the known track angles and measured energy deposit in the silicon detectors).

in the agreement of these two beam-energy estimates.

8.14 Conclusion

In this chapter, the workflow for an analysis of the experimental data recorded by the TexAT TPC was introduced and discussed. In the process of developing this workflow, a range of problems associated with the early variant of the TexAT detector were identified, such as the poor performance of the silicon detector array, the considerable problem of signal ringing and its indistinguisability from valid signals, and the challenging nature of the inhomogenous pixellated MicroMeGaS. The raw GET data files that this analysis was built to ingest are typically of considerable size; even with a partial read-out mode, a single event nominally contains on the order of 300 waveforms of 512 samples. The workflow described in this chapter was designed to make such a problem tractable; breaking down the data into compute-friendly chunks, and leveraging modern distributed computation frameworks to facilitate its analysis. Each stage of the workflow given in Fig. 8.1 was original work, with the exception of the ROOT Conversion tool, which was improved but otherwise pre-existing; and the gas-simulation tool described in Section 8.8, which was strongly inspired by the Texas A&M approach.

The track-fitting technique established in this thesis was used to perform preliminary steps towards an investigation of the structure of ^{14}O. With the challenges posed by COVID-19 and the quirks with the aforementioned TexAT data, this analysis was not carried beyond an initial exploration of the problem space. During this process, limitations of the data were identified, such as the missing Bragg peak of the heavy scattered beam that would be required for zero-degree reconstruction. Before the COVID pandemic, a future iteration of the TexAT detector was in its development phase, and this author made some contributions to the discussions and planning around future analysis software. This new detector, TeBAT, will solve many of the problems encountered in this analysis, such as the strip-chain matching for planar events, and the poor resolution of the central columnar region. Future efforts to finalise the analysis of the data discussed in this chapter should look to resolving the question of the zero-degree elastic-inelastic discrimination, before proceeding to account for the silicon detector efficiencies in the cross-section reconstruction of the reconstructed data. The preliminary excitation function shown in Section 8.13.3 indicates that the prediction of a linear-chain state in the 14 MeV—19 MeV region may be feasibly identified in these data, should it exist.

Part III

O-Lab Online Nuclear Laboratory

CHAPTER

NINE

O-LAB ONLINE NUCLEAR LABORATORY

> And now for something completely different.
>
> ---Monty Python's Flying Circus.

A foundational theme running through this thesis is the use of interactive computing and modern software development practices to support research and communication. In this chapter, an application of these principles is outlined in response to the COVID-19 pandemic, and the unique, time-sensitive challenges that it posed for teaching provision within the School of Physics and Astronomy at the University of Birmingham.

It is my hope that this work helps to shape the learning of future student cohorts, from both a technical and cultural perspective, and serves as a testament to the versatility of interactive, virtual learning.

9.1 Motivations

The identification of the SARS-COV-2 virus in 2019 marked the onset of an unprecedented period of social and economic disruption across the global community. Its impact has been felt in the majority of countries around the world, both in terms of the real-terms consequences of wide-spread infection, and the subsequent public health interventions in pursuit of elimination or suppression. Within the domain of Higher Education, which has placed a particular importance upon in-person learning and interaction, the complex ongoing challenges associated with the COVID-19 pandemic have placed significant burden upon educators and their institutions [105].

At the University of Birmingham, a central aspect of the Nuclear Physics teaching provision is the availability of undergraduate laboratories. Over the course of the teaching semester, students enrolled in the appropriate modules are given the opportunity to learn the fundamentals of practical nuclear physics, including

- the principles which underpin the measurement of radiation,
- the process of building and testing radiation experiments,

- practices and procedures involved in working with radioactive sources,
- the analytical techniques required to extract physics from laboratory measurements.

The specialised nature of the laboratory sessions in which these skills are taught imposes rigid constraints on the kinds of accommodations that can be made in light of University public health measures e.g. occupancy limits, staffing availability, and ventilation requirements. Several experiments that undergraduates are expected to undertake include fundamentally immovable apparatus, such as a $8\,\mathrm{m}^3$ graphite stack for neutron moderation, or a 150 L water bath used for neutron activation experiments.

The detector equipment used to observe and measure incident radiation is typically fragile, bulky, and extremely expensive, such that undergraduates are not expected to move such equipment within the laboratory. The radioactive sources used within the various experiments are also controlled substances; requiring significant care during handling, preparatory safety training, and diligent monitoring to ensure that the sources are properly accounted for at all times. These factors conspire together such that conventional approaches to mitigating COVID-19 risks during the pandemic, such as the introduction of a “two meter rule”, or reducing teaching group sizes, represent impractical solutions in the context of the nuclear laboratory [106].

Unlike other areas of the taught programmes such as lectures, in which provisions for distance learning (such as Panopto) were already in-use within the undergraduate teaching programs before the pandemic, no such solutions could be found to facilitate the development of practical laboratory skills in the domain of Nuclear Physics [107]. To address this challenge for the new and existing students impacted by the COVID-19 pandemic, an Online nuclear LABoratory (OLAB) was developed by a team of postgraduate nuclear physics students at the University of Birmingham. The objectives of this project were to implement an online learning resource which

- minimises the impact of digital inequality
- scales to meet the demand of the existing and future student cohorts
- provides an extensible platform upon which to develop new course material
- integrates with the University E-Learning platform(s)
- leverages existing, open-source technologies and
- facilitates content-first development

9.2 Virtual Laboratories

9.2.1 Simplicity and Reproducibility

Despite being strongly motivated by the need to provide students with laboratory experience, the opportunities afforded by an online physics laboratory transcend the immediate problem of tackling the disruption of a public health emergency. Through a carefully defined environment, experiment designers can guide students towards successful application of learned theory, but also towards "interesting" or educational failure modes [108]. Crucially, the experiment designer can design experiments in which these "error conditions" are directly addressed as part of the narrative content, such as through user-prompts, or remote demonstrator questions. For example, the background radiation observed during a measurement can be modified by the designer to exhibit particular spectral peaks that the student might be tasked with explaining. Due to the curated nature of the virtual environment, these failure modes can be introduced in accordance with the objectives of the laboratory; educators can *teach* through failure in a controlled manner. Furthermore, by designing failure into an experiment, educators are given the ability to prepare in advance for adverse outcomes, and guide students through the experience. This is particularly relevant in the event that experiment failures are irrecoverable, e.g. loss of experimental data, or damage to the electronics of a high-voltage power supply. Through the careful selection of a basis of representative failure modes, students can be better prepared for failures in a real-world experiment, which are often multi-factorial and represent poor educational experiences.

Given the removal of the physical constraints associated with on-premise in-person teaching, online laboratory experiences are free to explore experiment designs that are infeasible or indeed impossible to realise in a conventional setting. Experiments can be simplified to eliminate confusing details such as cumbrous software, or simplify characteristics of the model such that the results are easier or more convenient to interpret e.g. increasing the rate of a decay process [108]. Furthermore, it is possible to treat physical invariants as model parameters that students can interact with e.g. the material properties of a radiation absorber. This freedom to modify "reality" can be used to build experiments in which the challenge to the student is not building the experiment, but rather interpreting the observed results. Where suitable, this may provide a mechanism to improve student understanding and retention.

9.2.2 Hybrid Learning

In addition to providing educators with greater freedoms over experimental outcomes and narrative towards deeper understanding of fundamental concepts, a significant benefit of virtual experiments is the improved efficiency associated with reduced setup time, and time-to-results. Nuclear physics experiments are typically constrained by the time required to configure the necessary detection equipment, and the subsequent interval over which the measurement must be gathered. In some cases, the time investment required to perform a measurement may serve as a lens through which to develop planning and reflection skills [108]. In the virtual laboratory, educators can choose whether to impose these restrictions. Where hybrid learning is possible, it may be favourable to distinguish between virtual and practical experiments through these real-world constraints in order to prioritise different learning outcomes. Equally, by preparing an experiment in the virtual space, students would have the opportunity to identify areas of concern ahead of time, and optimise their time in the physical laboratory e.g. determining ideal detector locations. It has been shown that hybrid approaches that combine in-person and virtual laboratories can improve learning outcomes with respect to either method in isolation [108].

For educators, the added flexibility of hybrid learning can also be beneficial; supervision of multiple groups of students is made more feasible by the elimination of spatial constraints associated with particular in-person demonstrations [109]. In the event that physical/spatial resources are limited, such as equipment confined to particular laboratories with a finite capacity, virtual experiments can be used to increase the number of concurrent students that may engage with the session. Furthermore, demonstrators of virtual laboratories would not be required to undergo the same kinds of safety training that in-person demonstration necessitates. This follows from the improved health and safety risk-profile of distance learning.

9.2.3 Accessibility and Student-Centred Learning

As noted in the introduction to this chapter, a significant motivation for the development of OLAB was provided by the need to support students engaging with distance learning as a consequence of University, domestic, or foreign policy. Outside of the pandemic control measures, the opportunity for students to revisit practical experiments in order to aid retention and understanding has been identified by students as a welcome consequence [110]. Remote labs can be used to extend the window of availability of an experiment outside the limits of contact time, which provides tangible benefits to part-time learners, or students who wish to revisit material at a later date in order to consolidate their learning. It is also noted that the flexibility to arrange their own learning process can give students a sense of ownership that correlates with increased motivation [109].

It should also be noted that hybrid virtual-practical laboratories offer the opportunity to increase accessibility and equity of learning provision, particularly in such cases where buildings or experiments cannot be readily made accessible. For students with learning difficulties such as dyslexia, the ability to decouple the process of learning laboratory skills from

the spatio-temporal requirements of timetabled learning may offer a more inclusive learning experience. Furthermore, this flexibility may also offer benefits to students who are unable to access on-campus teaching, but wish to continue their learning in some capacity, such as students on temporary withdrawal or off-sick.

Undergraduate laboratories typically imposes significant constraints on the teaching schedule due to the long sessions required to setup, measure, and restore the experimental equipment. This, in turn, places a higher cost on mistakes during these sessions, and prevents the students from repeating an experiment with the knowledge learned through failure. The ability to revisit experiments has been recognised by students as a positive aspect of virtual laboratories. Some students have also reported that the ability to work alone was a preferential mode of learning that was not possible in the physical laboratory due to equipment constraints [110]. Alongside this, research has shown that communication opportunities within virtual laboratories are highly correlated with student learning achievement; care must be taken by educators to ensure that these opportunities are established within the virtual space [111].

9.2.4 Powerful Learning Tools

The provision of a virtual laboratory offers the opportunity for educators to directly influence the kinds of software available to students on a per-course or even per-experiment basis. In particular, unlike practical laboratories in which the hardware often has a lifespan of many years (with the associated software lifespan), virtual laboratories may not be so far constrained in the software that students are required to use, and are better positioned to modify the tools that students are presented with on a needs basis. This is particularly relevant in the context of nuclear laboratories, in which the lifespan of the apparatus can span decades. Virtual laboratories are well positioned to eliminate these kinds of "lock-in" issues. Within virtual experiments, educators may choose to provide interfaces that make it easier for students to analyse their acquired data. The elimination of this boundary between measurement and analysis is a contributing factor to reducing the "time to result" that distracts from learning outcomes. Whilst there is a recognisable place for teaching students the skills required to operate real hardware, virtual laboratories enable educators to determine when these skills are taught.

The promise of virtual laboratories includes the ability to interleave hands-on experimental tasks with curated learning material. Outside of the virtual laboratory space, Massive Online Only Courses (MOOCs) are increasingly built around interactive notebooks, which employ the principles of computational narratives and interactive computing to build expressive and communicative resources that permit students to employ an explorative approach to their learning [112]. These tools show similar promise in the context of university education, where they are often already used to facilitate undergraduate programming classes. Developments in these tools offer the prospect for educators to interact with multiple remote students, providing a foundation for remote-first teaching. Students in virtual laboratories would arguably benefit from reusing the same tools that they are familiar with outside of the laboratory to reduce inertia. Furthermore, the challenge of preparing a software environment in order to perform an analysis is ever-growing; students wishing to

use the tools that are commonplace in modern high-energy physics (HEP) research must increasingly demonstrate an understanding of package management, code compilation, unix, and containerisation technologies before they can even start to begin looking at their data. By utilising managed, remote software, educators can reduce the burden of supporting students with installing these software packages on their own computers via a write-once, deploy many strategy. In some institutions, it has been possible for a small team of developers to support upwards of 800 students using such strategies [113].

9.3 Software and Infrastructure

The objectives outlined in Section 9.1 can be partitioned into development-focused priorities and product-focused outcomes. Though members of the development team had varying degrees of experience in writing Physics analyses, the landscape of modern software development spans a significant number of software packages, programming languages, and domain-specific knowledge. In order to deliver a usable, extensible platform which could be used to facilitate remote learning within the constrained time frame, a content-first, iterative approach to developing OLAB would be imperative. From these constraints, the following technological decisions were made:

Software / Tool	Motivation
Python	Python is the fastest growing programming language as measured by the TIOBE index, and is frequently used in data analytics and dashboarding contexts [114]. The development group already had varying experience with the language, which is renowned for its quick time-to-first-result.
Jupyter Notebook	Jupyter Notebooks are widely used in academic research. There are already a significant number of software packages that add interactive widgets, and the notebook format serves as a standard across the Jupyter ecosystem.
JupyterHub	JupyterHub provides a battle-tested multi-tenant platform around which to serve customised computing environments to individual users. It provides hooks to integrate with LTI authentication systems such as those used by the Canvas LMS.
ROOT	ROOT is a standard tool used in the particle physics domain for all stages of the analysis lifecycle.
ipywidgets	ipywidgets provides a well-tested suite of interactive widgets, which can easily be composed to build complex layouts.
Voilà	The Voilà server was used alongside jupyter-flex to present linear notebook documents as a customised user-friendly dashboard.
Kubernetes on GCP	To orchestrate the multi-tenant JupyterHub containers, Kubernetes was deployed on Google Cloud Platform.

Table 9.1: High-level list of software packages used to build O-Lab.

Further discussion of the merits and motivations behind this set of software packages is given later in this chapter.

9.3.1 Jupyter & Interactive Dashboarding

The Jupyter (a loose acronym meaning Julia-Python-R) project is a

> non-profit, open-source project, born out of the IPython Project in 2014 as it evolved to support interactive data science and scientific computing across all programming languages. Jupyter will always be 100% open-source software, free for all to use and released under the liberal terms of the modified BSD license [115].

It encompasses a range of open source protocols and repositories around which a collaborative, interoperable ecosystem has been established. Through these common protocols and standards, a project building a particular Jupyter project is able to integrate with a range of other Jupyter-aware applications and software libraries. The Jupyter kernel standard, and the Jupyter notebook format are of particular relevance to this project; around these two standards, a strong separation between code, evaluation, and presentation has been established. This separation facilitates specialised workflows in which

content authorship can use different tools to those used for the final presentation. Fig. 9.1 shows an example notebook of the Lorenz attractor open in the JupyterLab IDE. It can be seen that both narrative and computational components are present in the same editor workflow [116]. The same notebook can be served via the Voila server, which presents end-users with pre-rendered interactive representations of Jupyter notebooks. Fig. 9.2 shows a screenshot of Voila serving the Lorenz attractor notebook. It can be seen that the editable code cell has been hidden (removed) from the end user, which places greater emphasis on the narrative and interactive portions of the document.

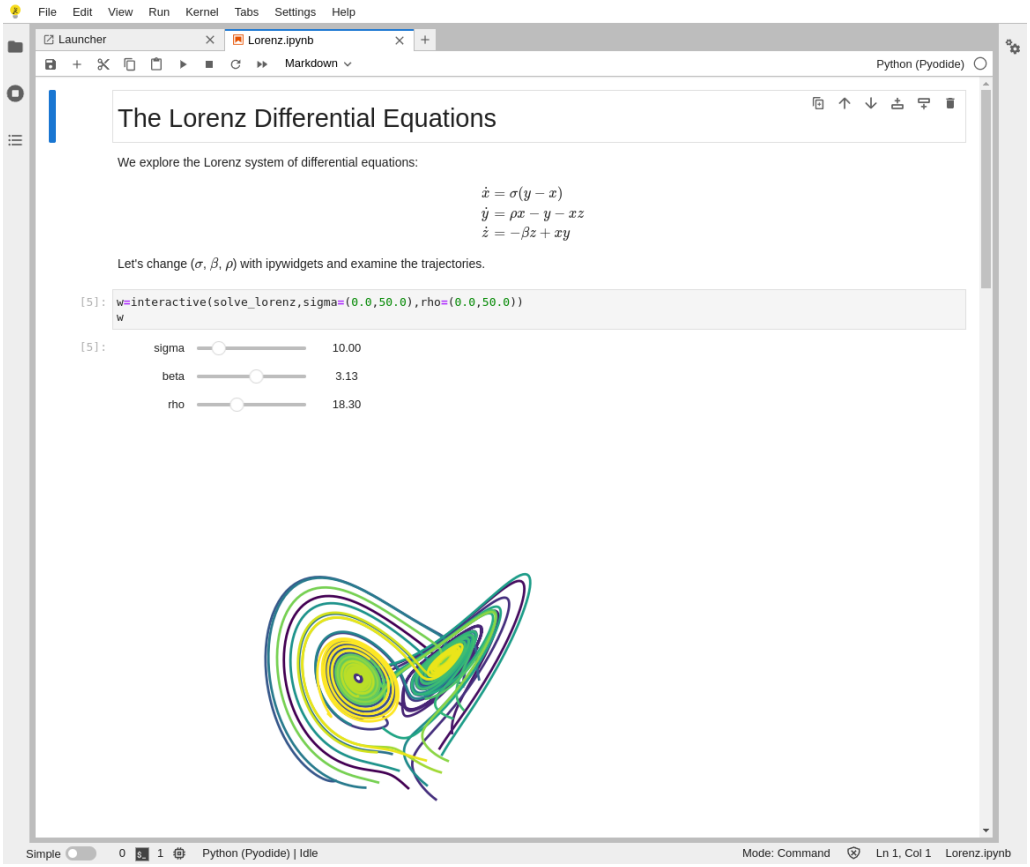


Fig. 9.1: Screenshot of the `Lorenz` demonstration notebook within the JupyterLab next-generation IDE. Rich computational narratives can easily be constructed by interweaving code and markup (Markdown) cells. Rich outputs, such as interactive widgets and rendered plots, can be embedded to support exploratory analyses.

Jupyter Notebook (or Lab) is already an established tool for aiding research within the academic sphere. It is used across a range of disciplines, and is particularly popular in the science and engineering fields. It follows that the OLab development team were already partially familiar with parts of the Jupyter ecosystem, and the choice to commit to the Jupyter software stack provided a low barrier to entry for authoring virtual laboratory content.

As discussed in Section 9.2.4, a prospective benefit of virtual laboratories is the ability to unify and maintain the software and content deployment process. A key component of this within the Jupyter project is JupyterHub.

JupyterHub brings the power of notebooks to groups of users. It gives users access to computational environ-

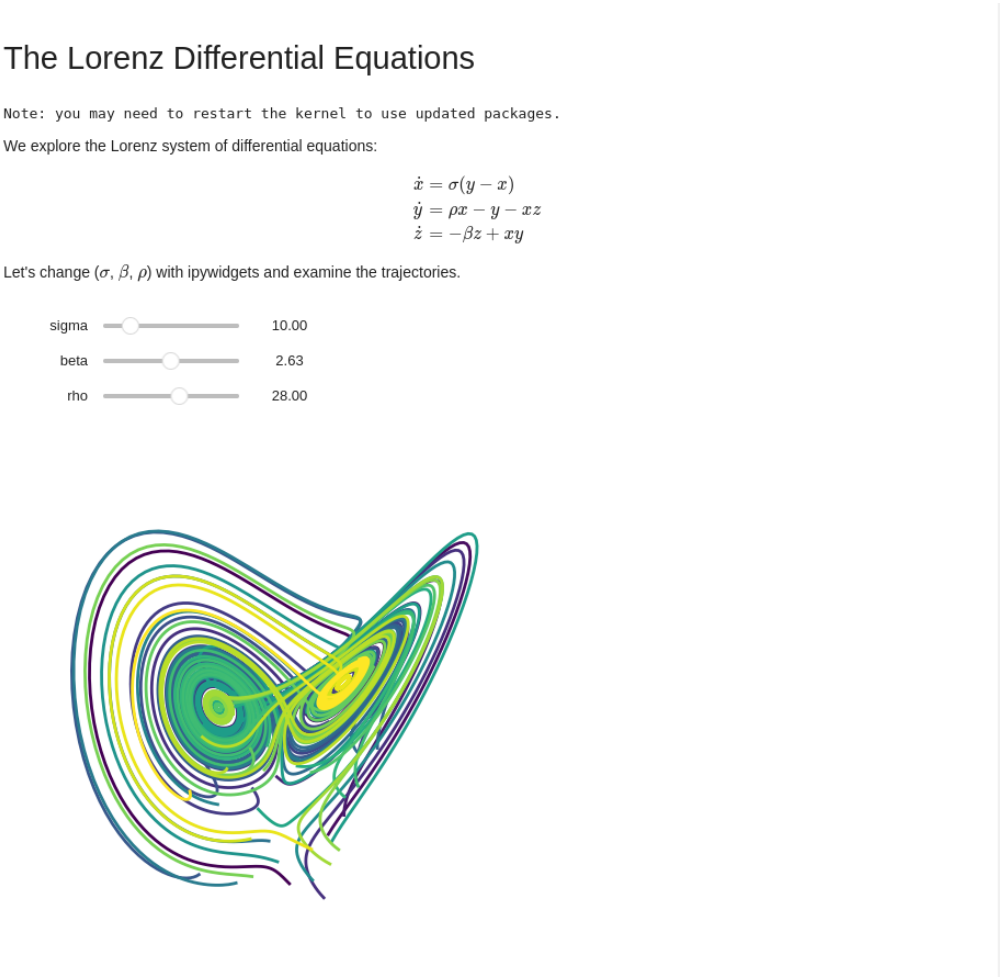


Fig. 9.2: Screenshot of the `Lorenz` demonstration notebook presented using Voila. Voila templates can be customised and configured: from re-arranging content or changing the formatting, to building entire applications.

ments and resources without burdening the users with installation and maintenance tasks. Users - including students, researchers, and data scientists - can get their work done in their own workspaces on shared resources which can be managed efficiently by system administrators [117].

Within institutions, JupyterHub is increasingly popular, with a number of universities building MOOCs upon the platform such as Berkley's Data8 course [118]. Though these institutions predominantly utilise JupyterHub to provision JupyterLab instances, i.e. to provide a partially homogeneous computing environment, there are an increasing number of groups using JupyterHub to serve third-party software such as RStudio. JupyterHub is built from a series of composable resources that facilitate integration of a multi-tenant compute platform with existing virtual learning systems and compute resources. Users can be authenticated using OAuth or LTI solutions. User environments (containers) can easily be configured by platform maintainers, and can be run on a range of cloud providers e.g. via Kubernetes, or using on-prem infrastructure via SLURM, Torque, or even Docker spawners. It was for these reasons that the Jupyter technologies listed in Table 9.1 were employed to build O-Lab.

9.3.2 Containerisation

Just as there exists a challenge in supporting students with installation and management of software packages, the same concerns are present when working with a team of developers to build software. Modern software development typically involves the usage of solutions that guarantee the reproducibility of an environment, such as `conda`, or Node JS's `npm` [119]. These package managers make use of a "lockfile" to solve a set of software dependencies ahead-of-time for a known set of architectures (in the later case, exclusively Node). In doing so, it is guaranteed that users of these lockfiles can reproduce the same environment, avoiding the often seen "this works on my machine" complaint. Of these two examples, Conda is the more general; it provides binary packages for multiple platforms, fulfilling the role of a general-purpose package manager.

Despite this, there are times when even an equivalent set of packages is insufficient to guarantee a working environment; existing software installed on the host operating system can sometimes "bleed" into the Conda environment, e.g. through environment variables, or additional constraints may be required, such as a particular file-system layout, or port availability. It follows that reproducibility of packages alone is not sufficient to fully reproduce a software environment. Containers are a tool which provide this abstraction by not only "locking" the application state (installed packages), but also the filesystem. They support mechanisms for remapping ports, binding external files, and can be deployed as immutable artefacts, avoiding problems of state corruption. For this reason, they can be a powerful tool in distributing the same development environment to a team of software developers. Indeed, tools such as Visual Studio Code and IntelliJ IDEA now provide remote development workflows that are powered by containers.

The container concept has undergone many iterations and realisations in the last decade, with recent popularity driven by

Docker, which has seen vast adoption since its introduction in 2013 [120]. Since that time, industry-wide efforts have focused upon standardisation, with the Open Container Initiative developing standards for container images (that contain the container filesystem) and container runtimes (that extract the container image and setup the relevant namespaces for execution). JupyterHub supports a number of spawners which work with user-defined container images, including `KubeSpawner` and `DockerSpawner`. The benefit of using containers in this context is predominantly to reduce development time; user containers can be tested locally on a local JupyterHub instance, and subsequently deployed with high confidence to remote / cloud-based Hubs. The ease of local development conserves developer time that would otherwise be spent on administering the remote deployment. At the time of O-Lab development, Docker was still a popular choice of image builder and container runtime, and was used to implement both the reproducible development environment, and the final production user image (see Listing 12.1).

9.3.3 Cloud Deployment & Provisioning

An impediment to deploying a JupyterHub within an academic institution is the lack of appropriate compute resources. Whilst universities often maintain on-prem compute facilities, these are usually designed to tackle conventional (homogeneous) HPC workloads, rather than the heterogeneous demands of general-purpose user containers. Whilst these systems can be leveraged to host a JupyterHub instance, e.g. using the aforementioned SLURM spawner, in practice the limitations of these systems are encountered fairly immediately. On an administrative level, HPC resources are typically earmarked for research purposes; deployment of teaching resources usually requires a separate treatment. Nonetheless, a trial deployment using the SLURM spawner was trialed to evaluate the viability of the HPC platform as a compute layer. There were several notable drawbacks identified during this trial period: users required VPN access in order to access the Hub from outside the University network, and the startup times were considerable; in some instances, on the order of tens of minutes. In general, this is a drawback of HPC-oriented schedulers; users are rewarded for short-running jobs. Meanwhile, JupyterHub jobs are unknown in duration, and therefore we must ask the schedule for long runtimes on the order of many hours, which leads to low job priorities.

It follows that a more versatile deployment solution was needed. Without dedicated on-prem compute to power O-Lab, a cloud solution was sought after. In this space, the dominant approach is to use Kubernetes as the lingua-franca of cloud compute providers. Kubernetes defines an abstract set of primitives that collectively facilitate large-scale container orchestration and deployment (see Fig. 9.3). It is supported by the majority of cloud providers, meaning that an investment in a particular cloud provider is not fraught with the perils of vendor lock-in, as is common with bespoke cloud platforms.

These building blocks are configured through declarative YAML templates, which can be parameterised by user values to easily customise the deployment through centralised configuration files. See Listing 9.1 for an example of such a template. Kubernetes is a “lingua franca” for cloud providers because it abstracts over the details of how these platforms expose

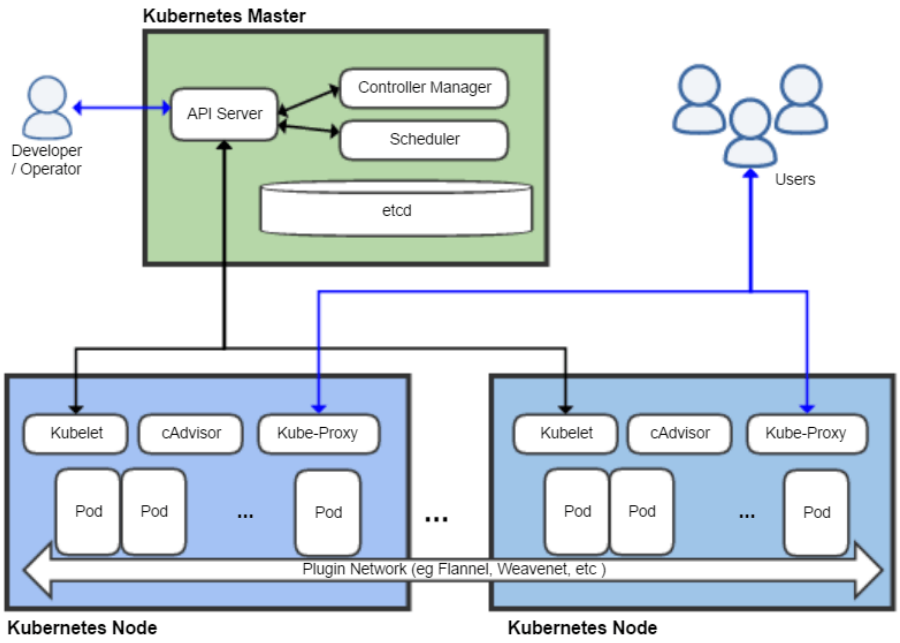


Fig. 9.3: Architecture diagram of a Kubernetes deployment. Several black-box containers are run within Kubernetes "pods", and communicate with one another over a shared network. Through horizontal and vertical autoscaling, Kubernetes can ensure that available resources are optimally distributed across the running pods. Healthy deployments are ensured through monitoring and subsequent redeployment, providing a mechanism for failure tolerance. Figure by Khtan66, distributed under a CC BY-SA 4.0 license.

e.g. routing, disk allocation, etc. Instead, Kubernetes provides a standard interface with which tools like JupyterHub's `KubeSpawner` can communicate, allowing a write-once use-everywhere design. This flexibility was employed during the O-Lab trials; an initial deployment was made onto Microsoft Azure, before subsequent migration to Google Kubernetes Engine (GKE).

Listing 9.1: YAML template for a Kubernetes resource configuring the HTTPS certificates of the ingress node.

```
apiVersion: networking.gke.io/v1
kind: ManagedCertificate
metadata:
  name: ingress-certificate
spec:
  domains:
    {{- range $.Values.domains }}
    - {{ . }}
    {{- end }}
```

Whilst Kubernetes is a powerful tool for managing deployments of scalable applications in the cloud, it does not address the challenge of setting up the infrastructure upon which it runs. For this, Terraform (after investigations with Ansible) was used. Like Kubernetes, a fundamental advantage of Terraform is its declarative configuration files; infrastructure state is expressed through intent (infrastructure as code), rather than imperative procedures. Although a unique Terraform configuration is required for each cloud provider, the building blocks provided by Terraform significantly reduce the maintenance burden. When building cloud-hosted distributed applications, the ability to automate the provisioning and deployment of infrastructure is invaluable.

Listing 9.2: Example Terraform Resource configuration file, provisioning a node pool of three compute nodes on Google Kubernetes Engine (GKE).

```
resource "google_container_node_pool" "user_nodes" {
  name       = "user"
  cluster    = google_container_cluster.primary.name
  node_count = 3

  node_config {
    preemptible  = false
    machine_type = "e2-standard-2"
```

(continues on next page)

(continued from previous page)

```
    # Google recommends custom service accounts that have
    # cloud-platform scope and permissions granted via IAM Roles.
    service_account = google_service_account.default.email

    oauth_scopes = [
      ...
    ]

    # To prevent non-user nodes being scheduled here
    labels = {
      "hub.jupyter.org/node-purpose" = "user"
    }

    taint = [
      {
        effect = "NO_SCHEDULE"
        key    = "hub.jupyter.org_dedicated"
        value  = "user"
      }
    ]
  }

  autoscaling {
    min_node_count = 0
    max_node_count = 3
  }
}
```

A significant drawback of using cloud infrastructure is the runtime costs. Whilst these services take responsibility for managing the infrastructure upon which they operate, this comes with an associated penalty with respect to usage tariffs; users are typically charged according to the number of nodes (and their respective configurations) required to run the cluster. Cloud operators increasingly offer managed Kubernetes deployments, such as Google Kubernetes Engine's Autopilot, in which the baseline runtime cost is governed by the number of *pods* running, rather than the nodes upon which they run. In this model, even more of the work associated with deploying an application on Kubernetes is managed by the cloud provider, simplifying the task of deploying an application. In general, the tradeoffs and benefits of using a managed

service over a self-managed deployment yield an inflection point in the baseline load below which it is more affordable to manage the deployment one's self. Furthermore, managed Kubernetes services often do not benefit from the same "free tier" benefits that providers such as Google Cloud make available. As such, this self-managed scheme was used in the O-Lab. For this reason, infrastructure-as-code was of particular importance; the ability to programmatically tear down and subsequently re-establish the cluster infrastructure enabled the implementation cost-saving shutdowns during quiet period (e.g. weekends).

JuypterHub has several distributions according to the needs of the user. For small deployments *The Littlest JupyterHub* (TLJH) runs on a single machine, without containerisation. For larger deployments in which scaling is important, and containers are required, there is *Zero to JupyterHub with Kubernetes*. It is this distribution that was chosen for O-Lab. In this distribution, JupyterHub employs two primary Kubernetes pods: a proxy pod, and a Hub pod, that together comprise "JupyterHub". The proxy is used to route network traffic between different endpoints according to a cookie that is associated with each user request. If a user already exists, the proxy pod routes their traffic to the appropriate user pod. Otherwise, the request is instead sent to he Hub pod, which creates the user container. The Hub pod, meanwhile, is responsible for administrative tasks, such as authentication (is the user logged in?), container spawning (provisioning an environment for new users), and an administrative panel for maintenance and deployment information (e.g. who is logged in) (see Fig. 9.4).

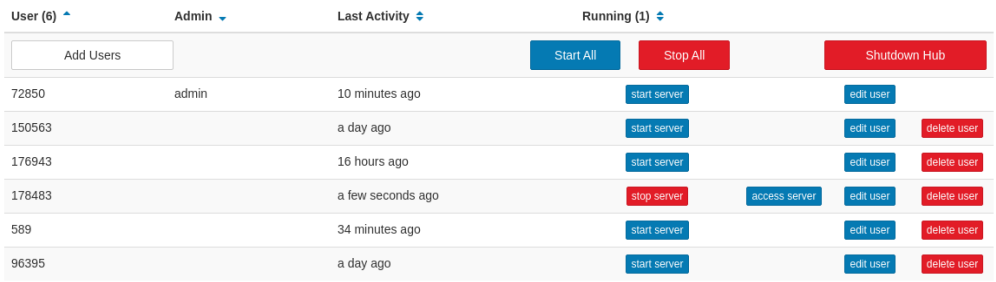


Fig. 9.4: Screenshot of the JupyterHub administrator panel, showing a list of previously seen users, and their server status. Individual user servers can be stopped and started by the administrator; enabling remote troubleshooting in the event of a problem.

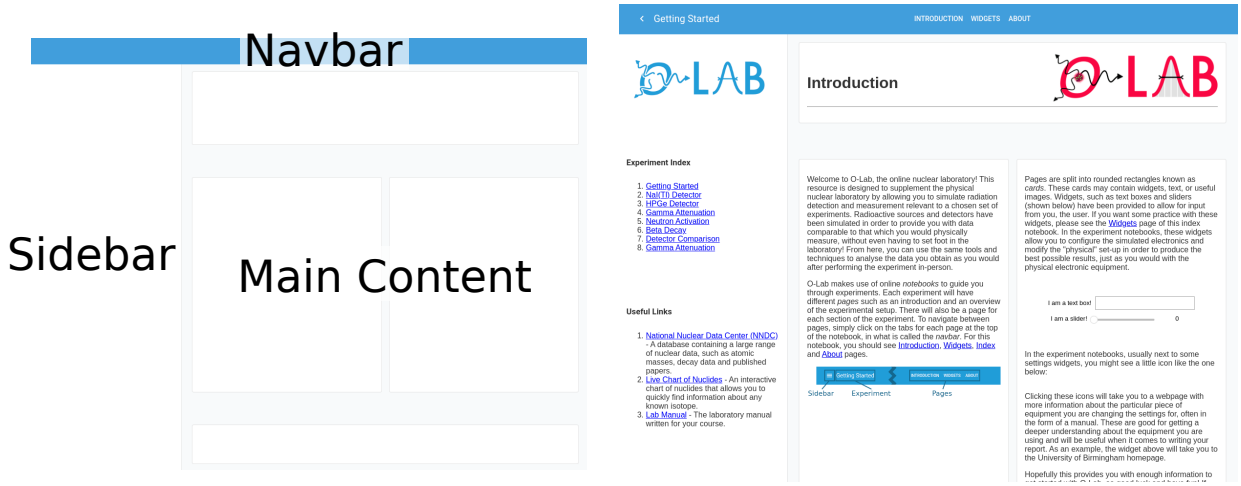
\

9.4 Content

9.4.1 Experiment Layout

Each page of the O-Lab deployment consisted of three main sections: the sidebar, the navbar, and the main-content (see Fig. 9.5). The sidebar, which was hidden by default but could be maximised, was used to present a set of global

navigation links between the various experiments and other resource links, such as the laboratory manual. These links could be customised according to each deployment; different teaching groups made use of different manuals, and so the links generated in the sidebar could be configured for each user. The navbar was used for *local* navigation within the current notebook. For the experiment notebooks, this was used to separate the different stages of an experiment. Within each page, the per-page content was shown in the main-content region, which comprised a grid system populated by discrete "cards". Each card contained rich content, e.g. rich-text; interactive widgets; or images; which was used to convey the relevant information about the experimental setup, analysis tasks, and equipment configuration to the user. An example notebook that employs this card system can be seen in Fig. 9.5.



(a) (a) (b) (b)

Fig. 9.5: (a) Schematic rendering of the page layout of a typical O-Lab experiment. Content is divided into several regions; the sidebar, which contains the global navigation links; the navbar, which contains the local navigation links; and the main content, in which the experiment details are presented. Each white rectangle represents a "card", which can contain rich content such as rich text, interactive widgets, or images. (b) Rendering of the "Getting Started" resource presented to new users of O-Lab. This page gives an overview of the purpose of O-Lab, its core features and user-guide, and where to access additional support.

9.4.2 Preliminary Experiments

As a first iteration of the online laboratory concept, a set of six preliminary experiments with close resemblance to their real-world counterparts was designed, listed in Table 9.2. In addition to reducing the scope of the laboratory design, replicating on-premise experiments also offers the opportunity for students to prepare for practical in-situ experimentation, with an understanding of the real-world constraints. The landing page of each notebook was designed to introduce the experiment, and discuss the various tasks that the user must complete (see Fig. 9.6. Rendered schematics of the experiment were generated by the the team, and included in the introduction to give a visualisation of what a practical experiment configuration would look like.

Notebook	Aim
`10-nai-detector.ipynb`	To investigate and study the properties of the NaI(Tl) detector and understand the mechanisms involved in the detection of gamma rays.
`20-hpge-detector.ipynb`	To investigate and study the properties of the HPGe detector and understand the mechanisms involved in the detection of gamma rays.
`30-gamma-attenuation.ipynb`	To measure the absorption cross sections of γ-rays for a wide range of materials, and to understand the mechanisms involved.
`40-neutron-activation.ipynb`	To demonstrate the principles of identifying elements in a sample by means of neutron activation analysis.
`50-beta-decay.ipynb`	To measure simultaneously the half-life of a beta decaying nucleus and the decay energy from the beta spectrum.
`60-detector.comparison.ipynb`	To compare detector performance of different size NaI(Tl) scintillation detectors.

Table 9.2: List of preliminary O-Lab experiments

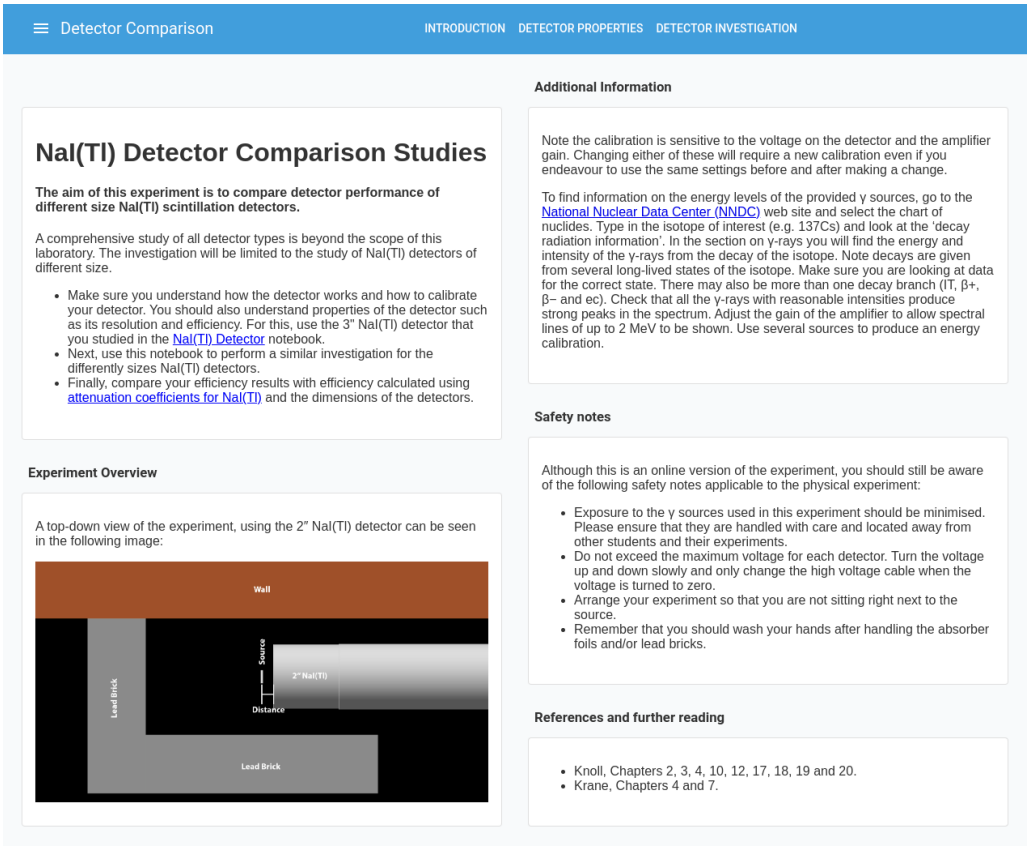


Fig. 9.6: Screenshot of the `60-detector-comparison.ipynb` notebook, described in Table 9.2. Each experiment followed a similar structure; the *Introduction* page was composed from several distinct cards: a list of tasks, a schematic diagram of the apparatus, additional details, safety notes, and any further reading.

From this initial set of experiment proposals, a single candidate was chosen for the proof-of-concept demonstration: *characterisation of a NaI(Tl) scintillation detector* (`10-nai-detector.ipynb`). In this notebook (see Fig. 9.7), users were charged with building a comprehensive understanding of the observable characteristics of the detector in relation to the observed gamma spectra. In the physical laboratory, this experiment involved the setup of lead shielding (to minimise background radiation, and to reduce exposure to ionising radiation), the selection of a range of gamma sources such as 137Cs, and the configuration of a NaI(Tl) scintillation detector with the necessary Multi Channel Analyser (MCA), high-voltage amplifier, and computer (for analysis). Fig. 9.7 shows the landing page for the experiment. Students were expected to read this page, and identify the key tasks and further reading in order to support self-assisted learning. As discussed in Fig. 9.6, the layout of this page was standardised with subsequent experiments to provide a familiar workflow. Once students were comfortable with the requirements of the experiment, they were encouraged to move to the next page (section) in order to learn about the details of the experiment that they would be performing.



Fig. 9.7: Screenshot of the `10-nai-detector.ipynb` notebook introduction page, bearing close resemblance to Fig. 9.6.

The *Experiment Details* page (see Fig. 9.8) was designed to provide auxiliary information about the experiment setup to the students. Students are expected to gain an understanding of the operating limits of the detectors (e.g. the safe voltage range), and learn how to connect the experimental modules together in order to perform a measurement. Some of these tasks have been replicated within O-Lab, whilst others (particularly the task of working with physical experimental modules) have been left for future work. As discussed in Section 9.2.2, it is an educational question as to whether to fully virtualise the physical experiment, or whether to build a distinct learning experience through the form of a virtual experiment. In the early stages of this pilot, it was decided that these experiments would attempt to prepare students for the real-world laboratory experiments. This would minimise the burden on educators to adapt assessment workflows, and account for the considerable uncertainty in the student timetabling as to whether students would return to the physical laboratories in the academic year. In future revisions to these experiments, this information will still be pertinent in order to contextualise the subsequent results that the students obtain.

The *Detector Properties* page (see Fig. 9.8) was designed to strike a balance between establishing O-Lab as a self-contained learning resource, and leaning upon existing literature to develop student understanding. Although virtual platforms are at an advantage in terms of integrating third-party content, such as by providing hyperlinks to additional websites, it was considered important that students were not required to move between multiple resources during an experiment. The cards explore the fundamental principles of efficiency, resolution, and calibration, in order to prepare students for the objectives of the experiment. The importance of mixed visual-text imagery was emphasised during the development of this page, with interactive widgets used to visually convey the principle of solid angle (see the *Geometric Efficiency* card in Fig. 9.8). These kinds of interactivity, as discussed in Section 9.2.4, provide a compact interrogatory mechanism for explaining visual phenomenon. The ability to mix rich-prose (Markdown) with mathematical expressions (LaTeX) and interactive widgets was found to significantly reduce the time to author this content. In particular, diagrams could be written in code and directly embedded into the document in favor of authoring the same content offline (see the *Intrinsic Efficiency* card in Fig. 9.8). This significantly reduced the time to produce these kinds of diagrams, and made the collaborative effort easier; content authors were able to synchronise changes to the document (using Git) and iterate upon the designs.

The *Detector Investigation* page (see Fig. 9.10) provides students with a series of interactive cards. Some of these cards were designed to represent the interface of experimental modules in the laboratory, such as the *ORTEC 671 Amplifier*. The *ORTEC 671 Amplifier* card depicts a spectroscopic amplifier manufactured by ORTEC; it provides the user with sliders to control the coarse and fine gain of the amplifier. This level of granularity (one card per module) was chosen in order to maintain a logical relationship between the experimental setup diagrams and the practical experiment. Each module embeds a link to the relevant manufacturer documentation, shown with a blue external link icon. This supports students in furthering their understanding of the apparatus, and encourages them to think about these cards as representative of a physical device. This is the first step towards designing a skeumorphic user interface.

In the *Spectrum Viewer* card, users are presented with a histogram display. This can be panned, zoomed, and supports

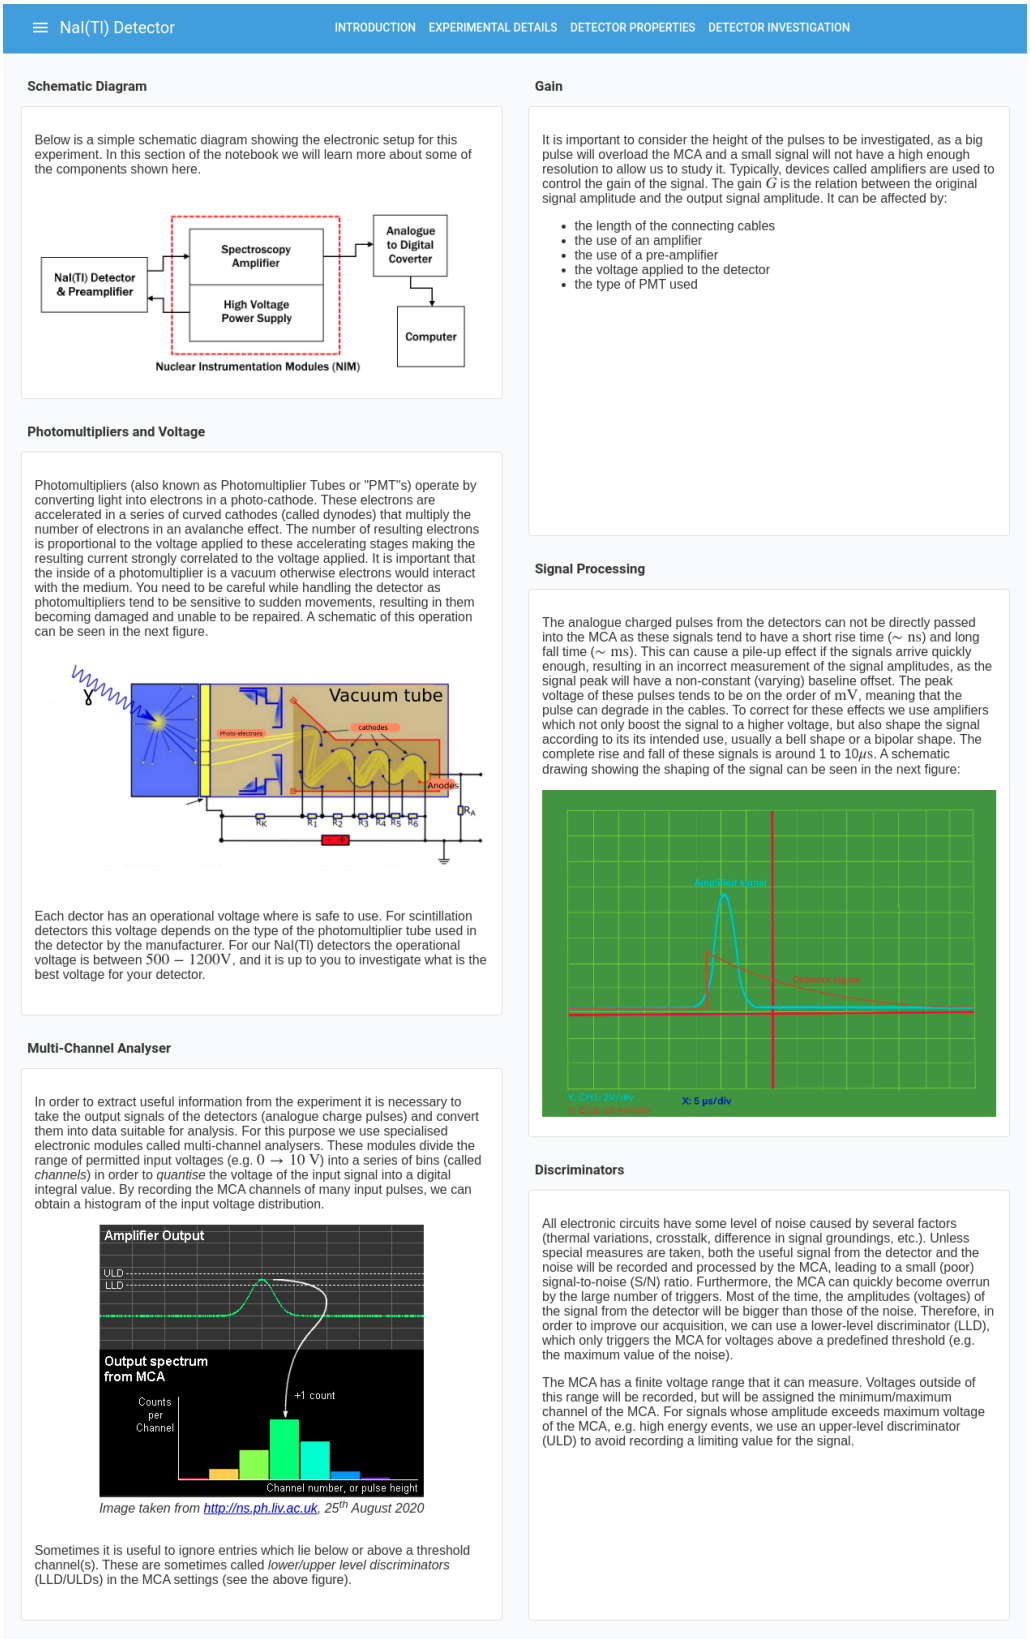


Fig. 9.8: Screenshot of the `10-nai-detector.ipynb` notebook details page, discussing the finer details of how the experimental equipment is operated, and points of consideration during operation.

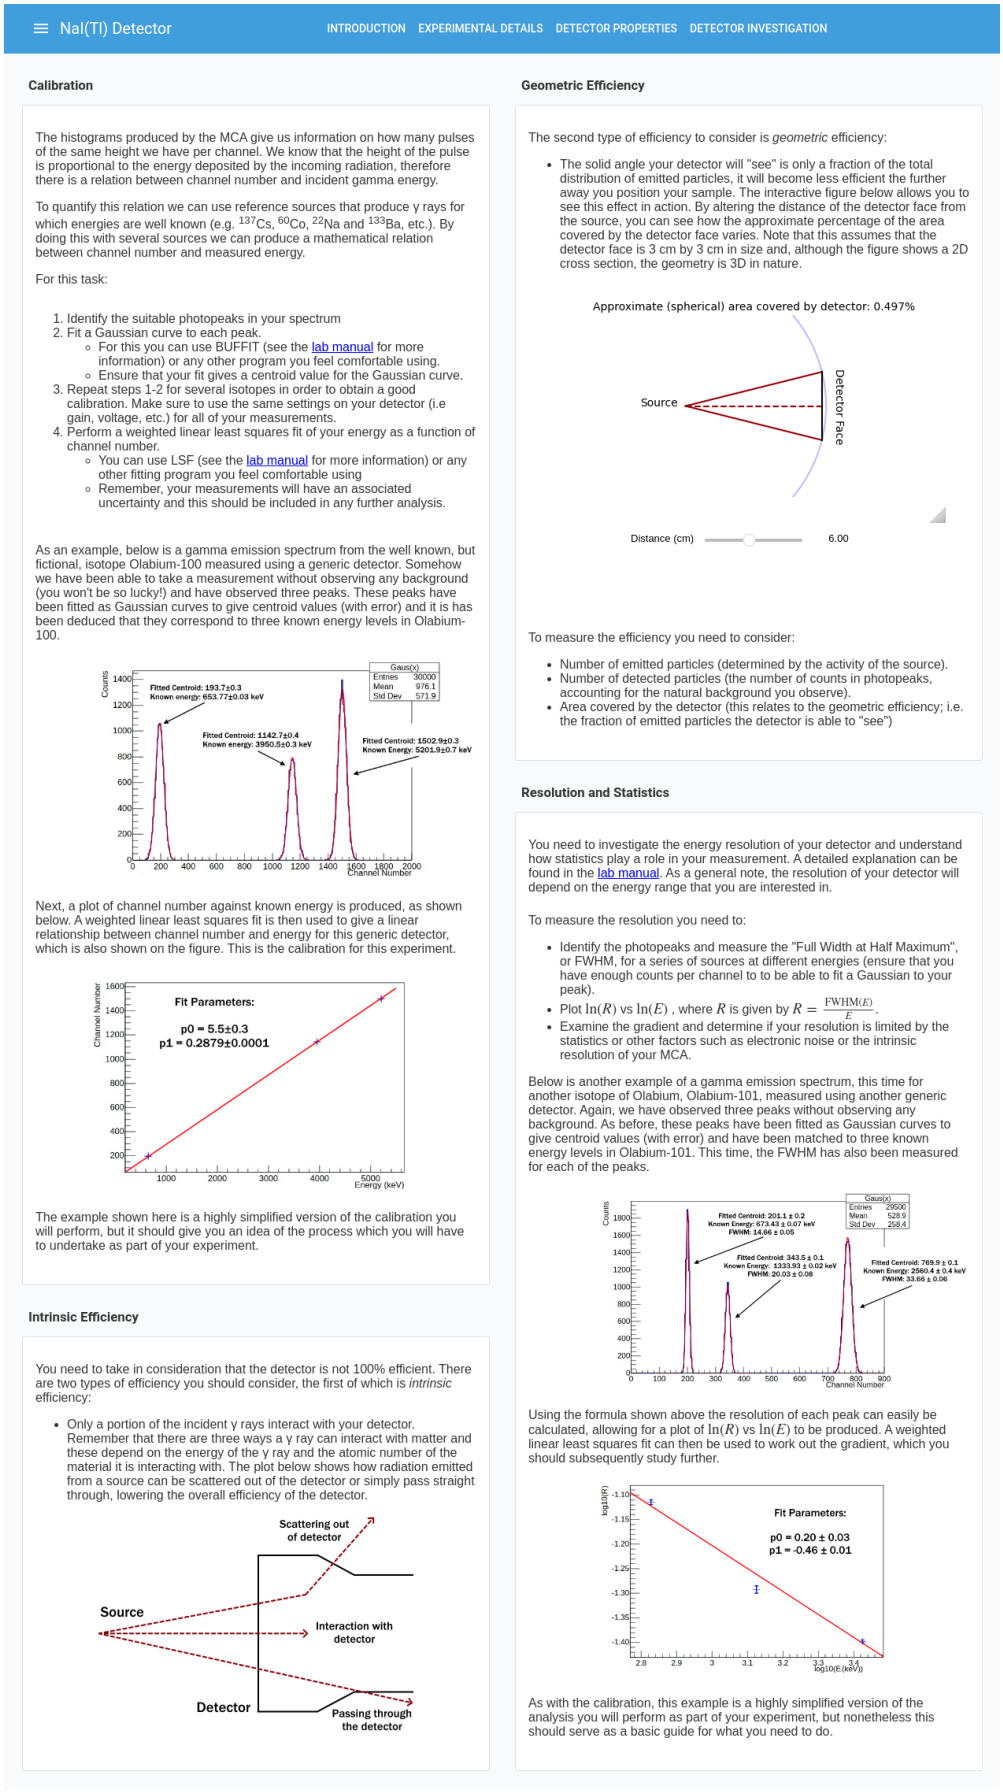


Fig. 9.9: Screenshot of the `10-nai-detector.ipynb` notebook properties page, discussing the various properties of a detector, and how they are measured.

switching to a logarithmic (Y) scale. A separate *Export* widget provides the ability to export the rendered histogram to a range of formats, such as ROOT, Spe, CSV, PNG, JPEG, and SVG. The result can be renamed according to the student preference. The exported plots can then be used by students in their subsequent write-up of the experiment.

The *NaI(Tl) Detector* card represents the sodium iodide scintillation detector that underpins the experiment. The user is presented with a voltage slider, controlling the preamplifier voltage, and shown a status checkmark (*Operational Voltage*) which describes whether the detector is running at a supported voltage. This was designed in order to prepare users for the practical considerations of operating a detector in laboratory conditions, a technique discussed in Section 9.2.1.

The *Environment* card is an entrypoint for the user-configured aspect of the experiment. Each experiment features a customised *Environment* card that implements the non-standard aspects of the experiment. In Fig. 9.10, the environment corresponds to the source type and distance to the detector. The user can also choose to hide the background, which would otherwise require a manual background subtraction in the laboratory.

Finally, the *Acquisition* card corresponds to the MCA discussed above. Users can perform multiple acquisitions, and configure the characteristic parameters of the corresponding device software such as the lower and upper channels, or the live time of the measurement. Repeated acquisitions are superposed via the *Acquire* button, whilst the user can clear the buffer using the *Clear* button.

An additional experiment was virtualised in order to assist with educational outreach provision. A simplified variant of the *Gamma Attenuation* experiment was built, in which the absorption cross sections of gamma-rays is determined for a range of materials (see Fig. 9.11). Using the rich-content mechanism, a spreadsheet widget was created, whose contents were mapped onto a scatter plot to visualise the Beer-Lambert law. A group of Year-9 students was tasked with exploring the relationship between absorber material, thickness, and the number of gamma rays observed. In a practical laboratory, this would have required careful instruction about how to use the apparatus, such as learning how to use the proprietary acquisition software, and handling of shielding materials. These present both a safety and supervision burden on staff, limiting the kinds of experiments that can be performed by visiting students. With O-Lab, it was possible to design an experiment that did not require such knowledge — exemplifying the flexibility of virtual laboratories that was explored in Section 9.2.2. In this notebook, some of the interactive elements were simplified to minimise the risk of confusion. The *Acquisition* card was modified to prevent superposition of measurements; students were not required to clear the buffer before each new measurement. The addition of a *Peak Area* card avoided the need for off-line analysis tools e.g. ROOT, and allowed students to continue to work in-browser. Finally, the plotting cards (consisting of a spreadsheet and scatter graph) were added to reduce the time-to-results for the students. These modifications ensured that remote demonstrators could effectively cover the experiment material in the required time frame.

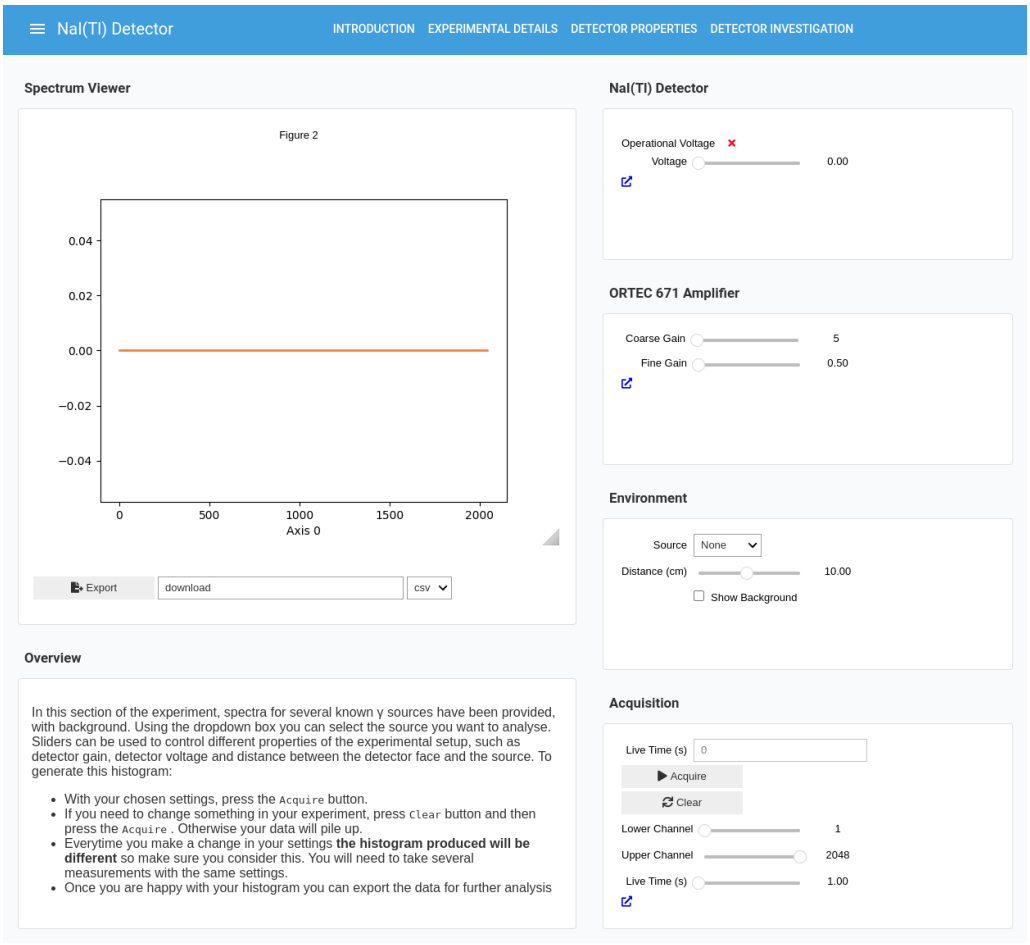


Fig. 9.10: Screenshot of the `10-nai-detector.ipynb` notebook experiment page, providing interactive widgets to configure the apparatus and produce a spectrum.

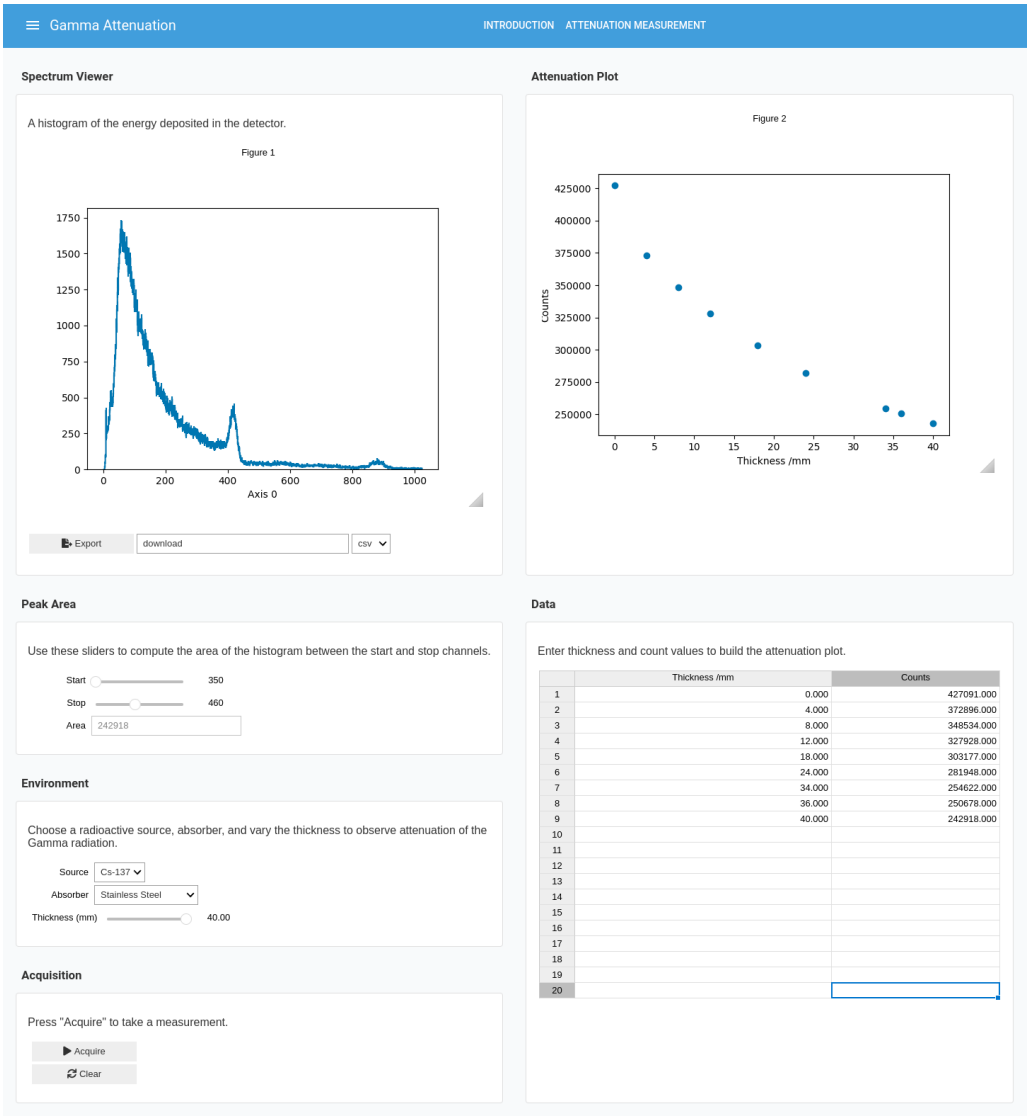


Fig. 9.11: Screenshot of the `70-gamma-attenuation.ipynb` notebook experiment page, providing interactive widgets to explore the attenuation of gamma rays through a range of absorber materials and thicknesses.

9.5 Analytics

9.5.1 User Analytics

In order to provision appropriate resources for the Kubernetes JupyterHub deployment such that costs could be minimised, a number of metrics were recorded. These data were then parsed from the Google Cloud Logging backend, and built into a dataset consisting of JSON records. Each record was built by correlating analytics events, e.g. the adjacent start and stop events for a particular user server. These records were then aggregated by notebook path (experiment name), time (concurrent users), and analysed to give indications towards usage trends.

The produced JSON records consisted of two families: a notebook record (see Listing 9.4), and a server record (see Listing 9.3). Both records contain a user ID field which uniquely identifies the user to whom the record pertains. The server record includes only the start and stop times of the server.

Listing 9.3: An example server-session record taken from recorded analytics.

```
{
    "user": "5ac27e656c88c5cddd3ebe956b57819586541a810ffac363eab2906e0539d6db",
    "start": 1605711173,
    "start_delay": 2.231,
    "stop": 1605717001,
    "stop_delay": 6.851
}
```

The `start` time of each user server was downsampled to a granularity of days, and the distinct number of unique users per-day computed. A total of 66 unique users (including demonstrators) were identified in this dataset. See Fig. 9.12 for a subset of these records. The OLAB cluster was closed over the weekend to minimise running costs, which is reflected in the sharp, constant drop in users during these two-day windows. The median number of concurrent users for the entire trial period was found to be 5 users per day, but this value represents the entire distribution; in totality, the trial was phased as different student cohorts were given access to the resource. Throughout the month of November, a steady increasing trend in the number of concurrent users can be seen. In the subsequent period, user engagement remained fairly closely distributed around 6 daily users. Strongly evident in these data is the end of term on the 21st of December 2020; in the week preceding, student engagement was much reduced.

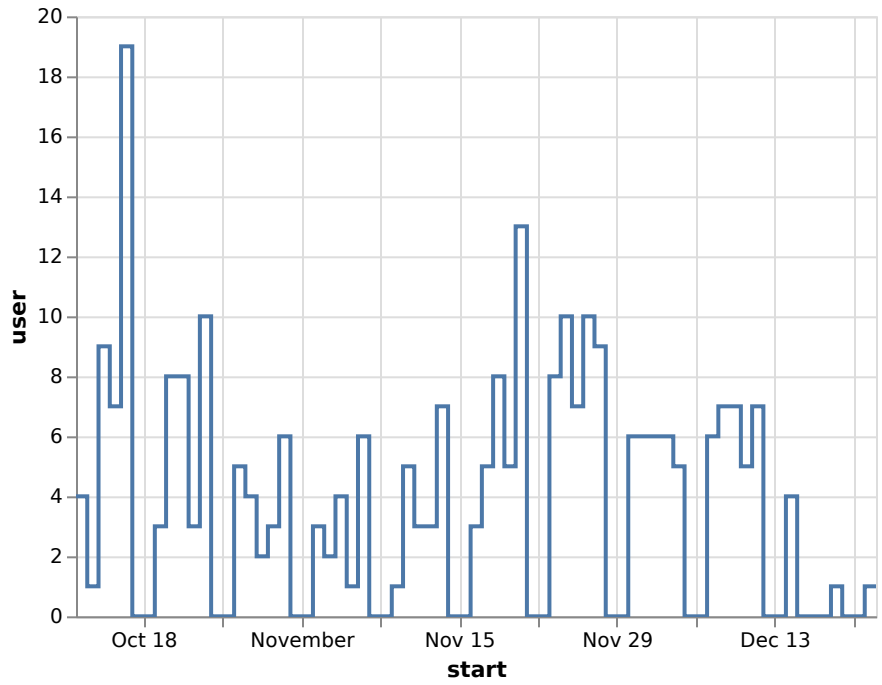


Fig. 9.12: The unique number of OLAB users aggregated by day. Anonymous user IDs were generated by hashing the unique Canvas student IDs.

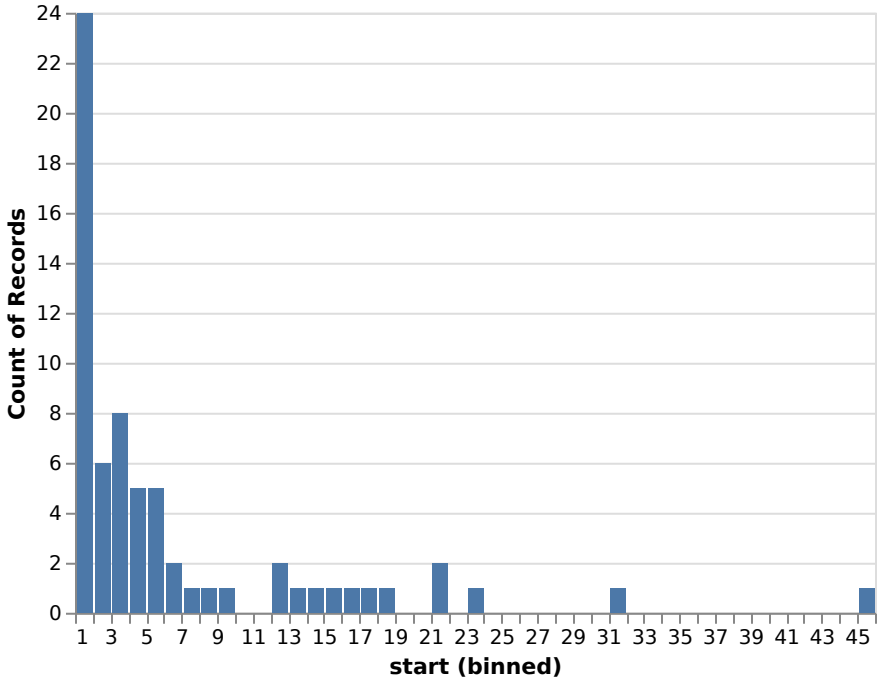


Fig. 9.13: The number of times each user started a server. Anonymous user IDs were generated by hashing the unique Canvas student IDs.

On a per-user basis, a long-tailed distribution was observed in the number of times that a server session was started (see Fig. 9.13). The first quartile of the number of starts was found to be 1. The cumulative session duration of each user shown in Fig. 9.15 indicates that the majority of these single-use servers were opened for very short periods of time. Meanwhile, Fig. 9.14 shows that the start times of each user session are skewed earlier in the single-use server dataset. This is likely due to the following phenomena:

- demonstrators walking through the login process with students; these demonstrators were not required to use OLAB during the sessions.
- servers that failed to launch for particular users.
- students from outside the teaching cohort visiting the deployment page.
- students failing to meaningfully engage with the laboratory.

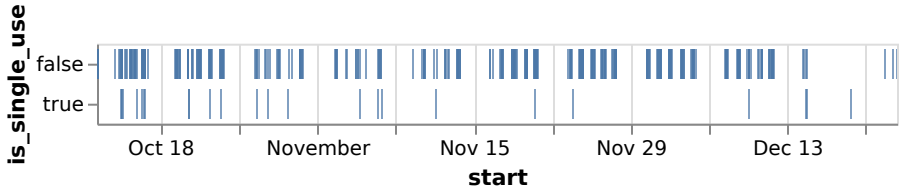


Fig. 9.14: The start times for each user session, grouped according to whether the user started more than one session during the trial.

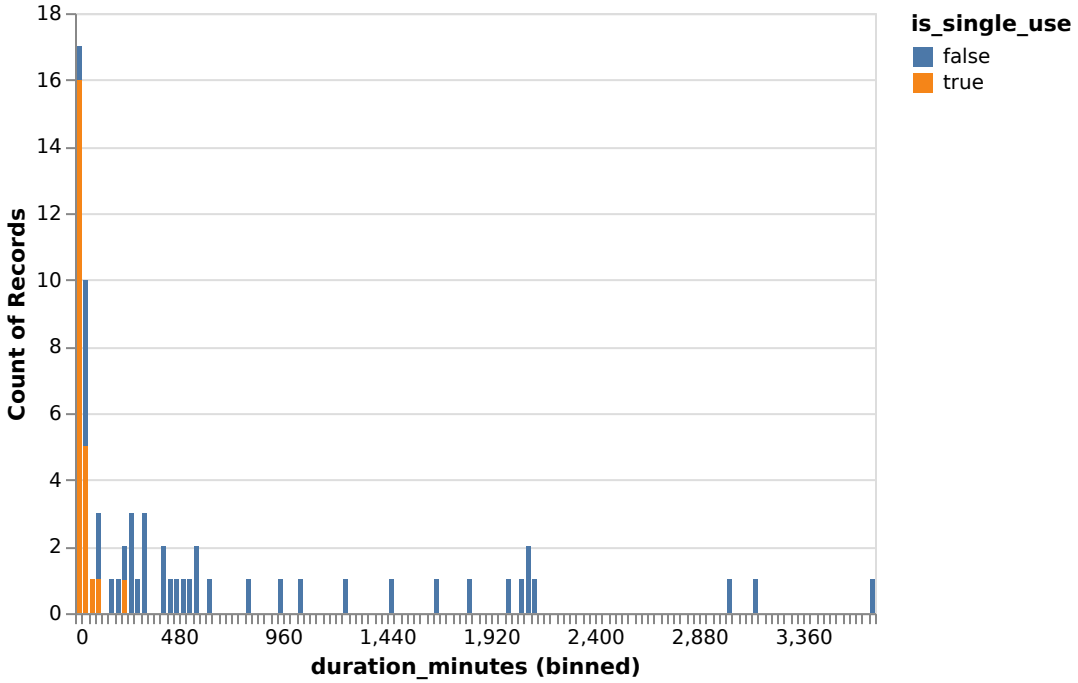


Fig. 9.15: The total session duration of each user.

The median session length for single-use servers was found to be 14.86 minutes. This includes an additional margin of one minute for the server to detect and stop servers after users close the browser tab. Due to the anonymisation of records, it is not possible to isolate instructors from students in these data. The median session length across all server sessions was 204.6 min, with a long tailed distribution; the mean session length was 567.1 min with a standard deviation of 858.7 min. One reason for this long-tail is the possibility of idle servers; servers with which users are no longer interacting, but remain open and consuming resources. A prompt was generated to remind users to close unused servers, but the component responsible for automatically closing the server (after a grace period elapsed) was not successful, following some technical problems with API authentication.

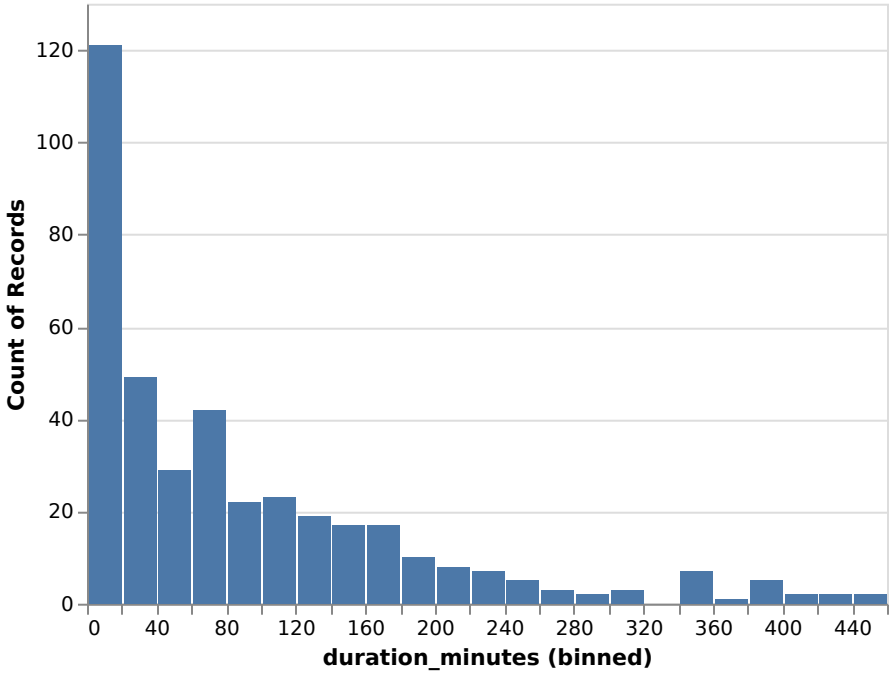


Fig. 9.16: The duration of each started user server.

To better understand resource provision, it is useful to look at the number of concurrent users over the course of the trial (see Fig. 9.17). From a peak of 9 users, the median number of concurrent users was 2. Again, partitioning the dataset around the 9th of November reveals two trends; the median number of concurrent sessions in the pre-November period was , whilst during the second period this value was . This finding was significantly far from early predictions; to ensure sufficient capacity exists for 90% of the trial period, provision for only 4 concurrent users was required. With the near-continuous roll-out of OLAB within the separate student cohorts, and without the necessary further study into usage trends and student engagement, only limited conclusions can be drawn from these data. Specifically, it is not within scope to use these data in order to predict the capacity required for future deployments.

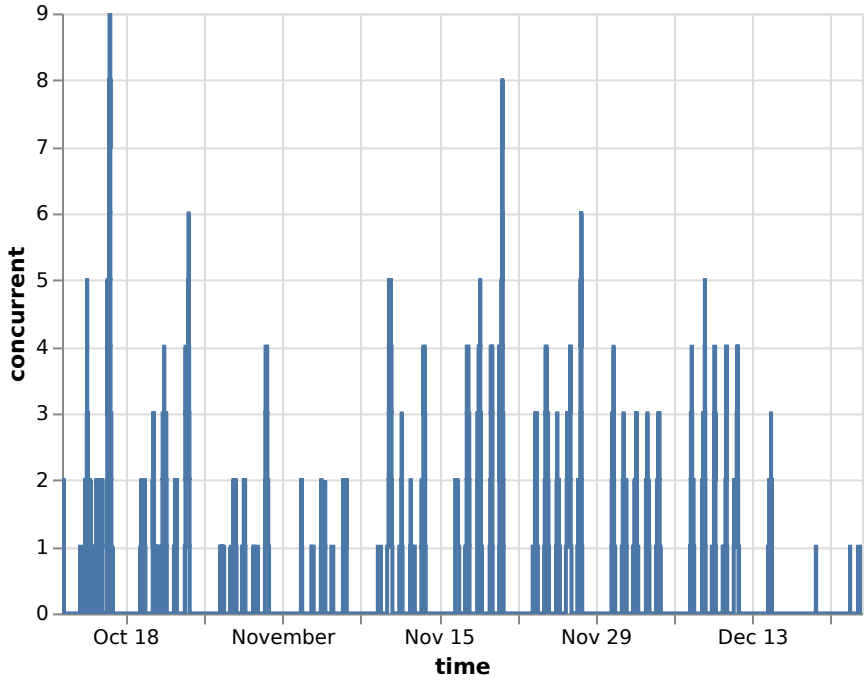


Fig. 9.17: The number of concurrent users (users whose `start` and `end` times overlap).

9.5.2 Notebook Analytics

The notebook record includes information about the time that a notebook session started, stopped, and to which notebook the record corresponds (see Listing 9.4).

Listing 9.4: An example notebook-session record taken from recorded analytics.

```
{
    "user": "62a39aa76ab34147ef996f9f8f052dd663d67b38bf13c55e8ce6b911404b2e8e",
    "kernel_id": "a5f751d0-1650-480a-b5b6-3384f6f72c96",
    "start": 1605878281,
    "notebook": "index.ipynb",
    "stop": 1605878396
}
```

Within the notebook analytics, clear trends emerge with respect to access patterns. In Fig. 9.18, a timeline plot of the per-notebook session intervals is shown. It can clearly be seen that the `30-gamma-attenuation.ipynb` resource was accessed predominantly after the 13th of November. Meanwhile, students most-frequently visited the

`10-nai-detector.ipynb` NaI(Tl) detector investigation notebook throughout the trial, with similar degree of regularity to that of the HPGe detector investigation notebook. It is likely that these notebooks were used as a reference on detector behaviour and principles (e.g. solid angle calculations) even after students had completed the experiment in the early (relative) weeks of the trial.

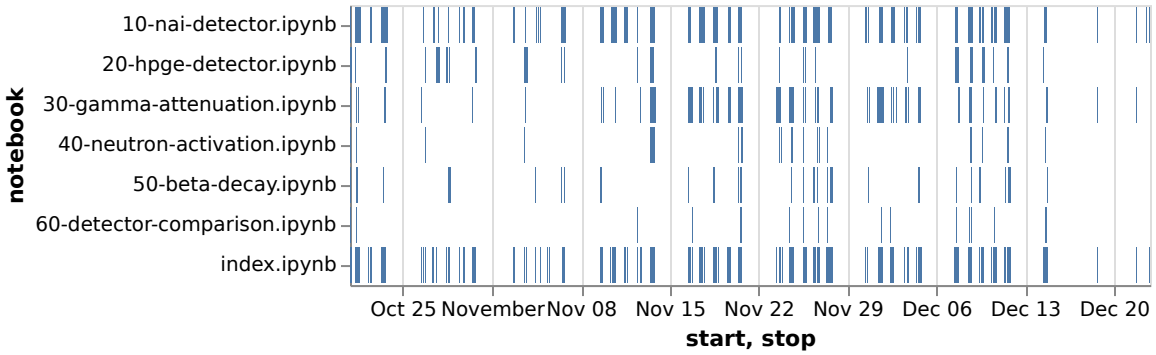


Fig. 9.18: Experiment notebook access patterns during the trial.

The number of times a notebook server was launched differed significantly across users. Fig. 9.19 shows the distribution of the number of started notebooks across the trial participants. A significant number of users opened under two notebook sessions, with the median number of notebook sessions equal to 8 sessions per user. Within these sessions, it is useful to consider the unique notebooks that were opened by each user. Fig. 9.20 plots the distribution of the number of unique notebooks opened by users. It is clear that the median student opened 2 notebooks. A breakdown of the number of users that visited each notebook (see Fig. 9.21) shows a distinct pattern; nearly all users visited the index page (those who didn't likely experienced connectivity problems), with the most popular experiment being the NaI(Tl) detector investigation notebook (`10-nai-detector.ipynb`). Besides this experiment, the majority of the virtual experiments had similar occupancy, with the detector comparison notebook (`60-detector-comparison.ipynb`) being the least popular of the available assignments.

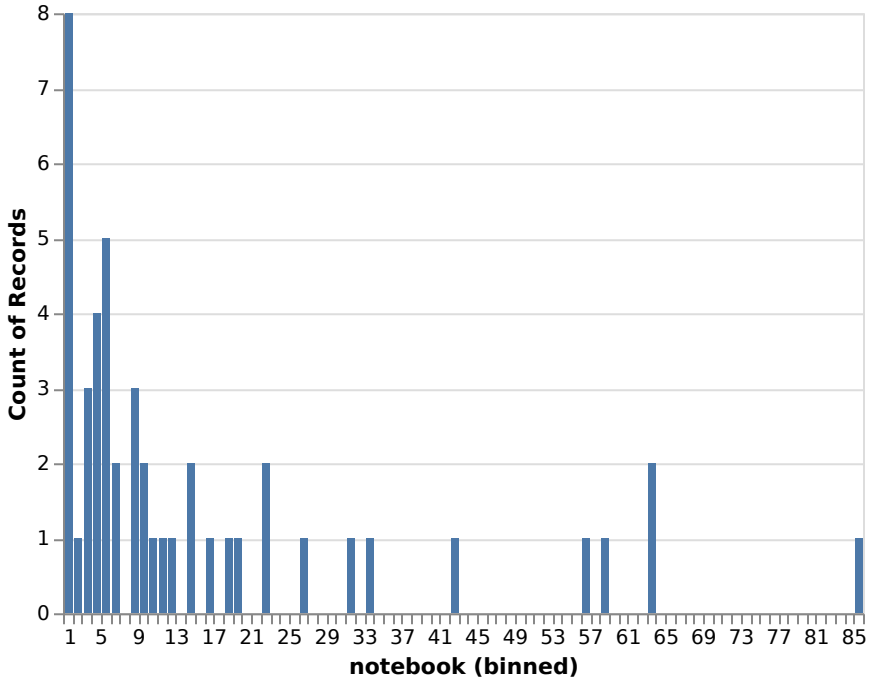


Fig. 9.19: The number of notebook sessions that users opened during the trial.

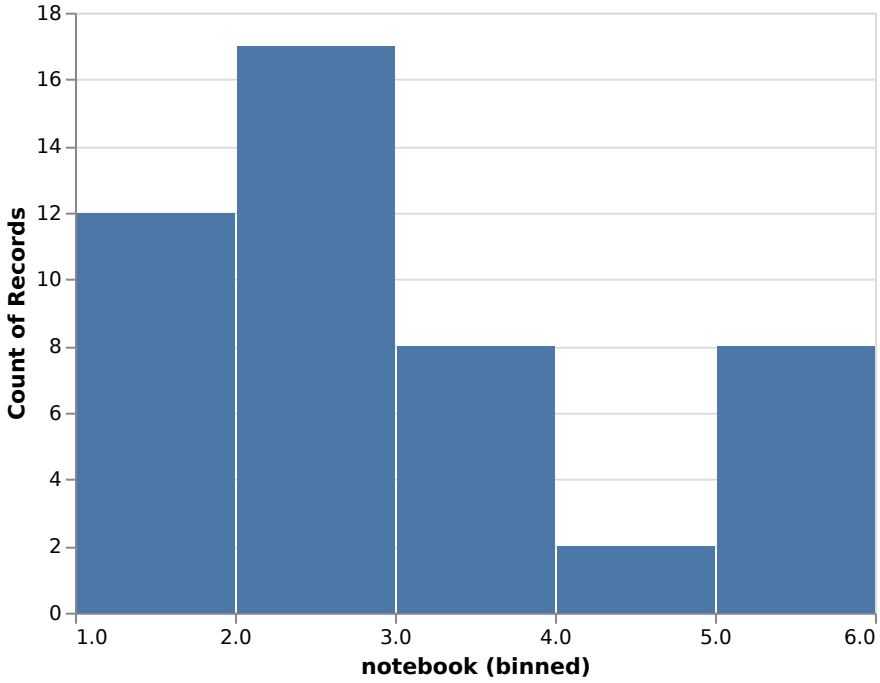


Fig. 9.20: The number of unique-notebook sessions that users opened during the trial.

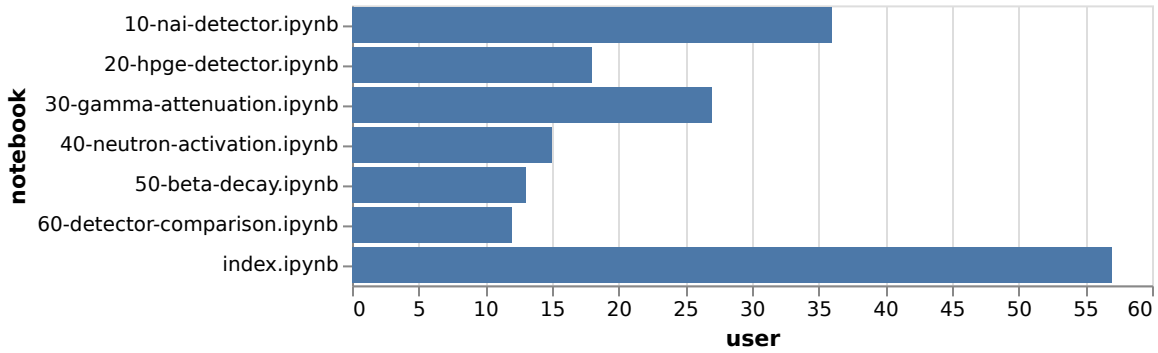


Fig. 9.21: The number of users that opened each notebook at least once during the trial. Anonymous user IDs were generated by hashing the unique Canvas student IDs.

9.6 Future Work

It was discussed in Section 9.4 that the provisional O-Lab deployment was designed with reduced scope in order to meet the requirements of the academic timetabling. these constraints. Yet, in Section 9.2, it was established that the wider landscape available to virtual laboratory design is much broader; the decision to replicate in-person experiments is both a practical decision *and* an educational one. Whilst the primary goal of this early O-Lab deployment was to facilitate the student laboratory provision during the semesters disrupted by public health measures, subsequent investigations could be designed to further explore the utility of a virtual nuclear laboratory in both hybrid and dedicated contexts. In order to evaluate the value of such a proposal, it would be necessary to predetermine the scope of a subsequent pilot study, and establish the criteria by which success can be measured. Both educators and students are stakeholders in laboratory work, and it would be of fundamental importance to establish a dialogue with these groups in the wake of this O-Lab deployment to determine those aspects which proved of value to students. At the same time, educators should be encouraged to provide feedback as to the impact that the virtual format of these laboratories had upon student engagement and understanding, particularly in relation to assessed material such as lab reports. During this preliminary trial, a feedback form was designed in order to assess student sentiment towards the trial, but was not released to students following technical errors. Such feedback would prove invaluable in directing future development of the O-Lab platform.

In addition to a broader remit, a subsequent iteration on the O-Lab platform could look to expand its scope for student interaction. O-Lab was designed upon a platform that could be extended to meet future need, such as student analyses. In the majority of experiments there is a need to analyse data, which is conventionally done with in-house software. Future development could provide students with a JupyterLab instance, and implement additional widgets to export data from the O-Lab frontend to these sessions in order to facilitate further investigation. It is of growing importance to be able to perform analysis work using programming languages such as Python; providing an ability for students to leverage these skills within a virtual laboratory would provide both a mechanism to reinforce student understanding and cross-

discipline knowledge. Additional JupyterLab plugins such as `jupyterlab-preview` enable integration of voila-backed dashboards within the JupyterLab environment (see Fig. 9.22).

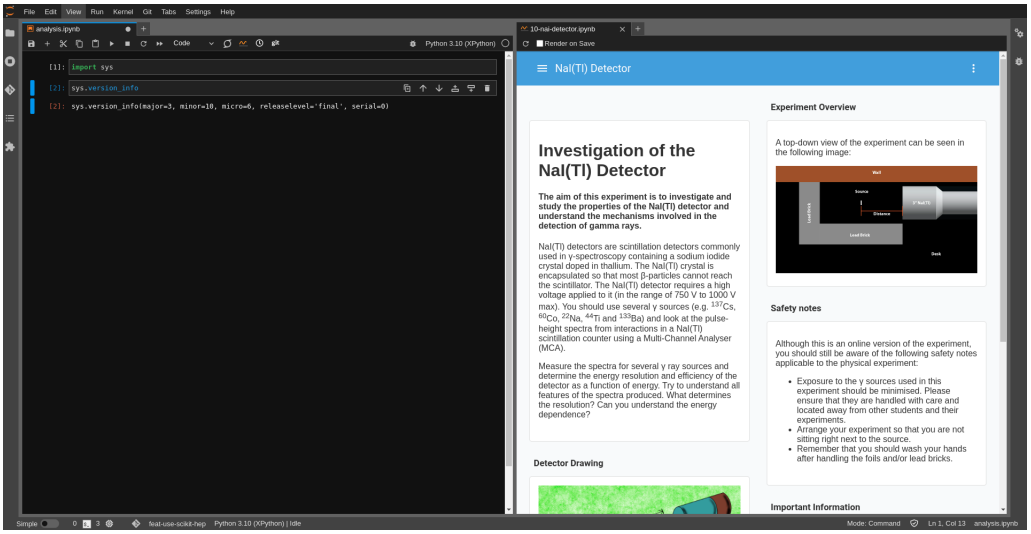


Fig. 9.22: Screenshot of the Voila preview panel running in JupyterLab courtesy of the `jupyterlab-preview` extension. With this extension, customised dashboards can be opened by the user, whilst conventional cellular notebooks can also be edited and run.

Students were able to engage with remote demonstrators during the O-Lab trial period using the Zoom video-conferencing software. JupyterHub plugins such as `jupyter-videochat` could be added to the JupyterHub deployment in order to reduce the barrier to engaging with demonstrators (see Fig. 9.23) [121]. It would also then be easier for students to form breakout rooms to work with peers on group projects. This would be increasingly useful as the scope of the O-Lab project grows, where the logistics concerning organising Zoom meetings could be circumvented in favour of drop-in sessions facilitated by the Hub plugin.

JupyterLab also implements support for collaborative editing. As part of a remote-first workflow, teaching assistants could utilise such a feature to work with students to resolve problems with their analysis. The standard working practices in the practical laboratory are designed around resource limitations; students are often tasked with working in pairs. The O-Lab pilot demonstrated that this was no longer a practical issue, and introduced the ability for students to work individually. Therefore, the collaborative features of JupyterLab could be employed to permit students collaborate on an analysis. Alongside a push towards building and publishing reports as interactive documents (notebooks), this kind of shared editing provides the opportunity for deeper collaboration than conventional pen-and-paper based workflows.

The Apple iOS operating system popularised the concept of skeumorphic design; making items resemble their real-world counterparts. Though there is continuing debate as to whether skeumorphism aids retention and learning, there is evidence to suggest that the *affordances* (ability to determine as-yet unseen behaviour from visual cues) it provides reduce cognitive overhead through familiar metaphors [122]. The real-world equipment that O-Lab imitates, such as the high-voltage amplifiers, could be extended to acknowledge this design trend. One approach would be to retain the existing data-

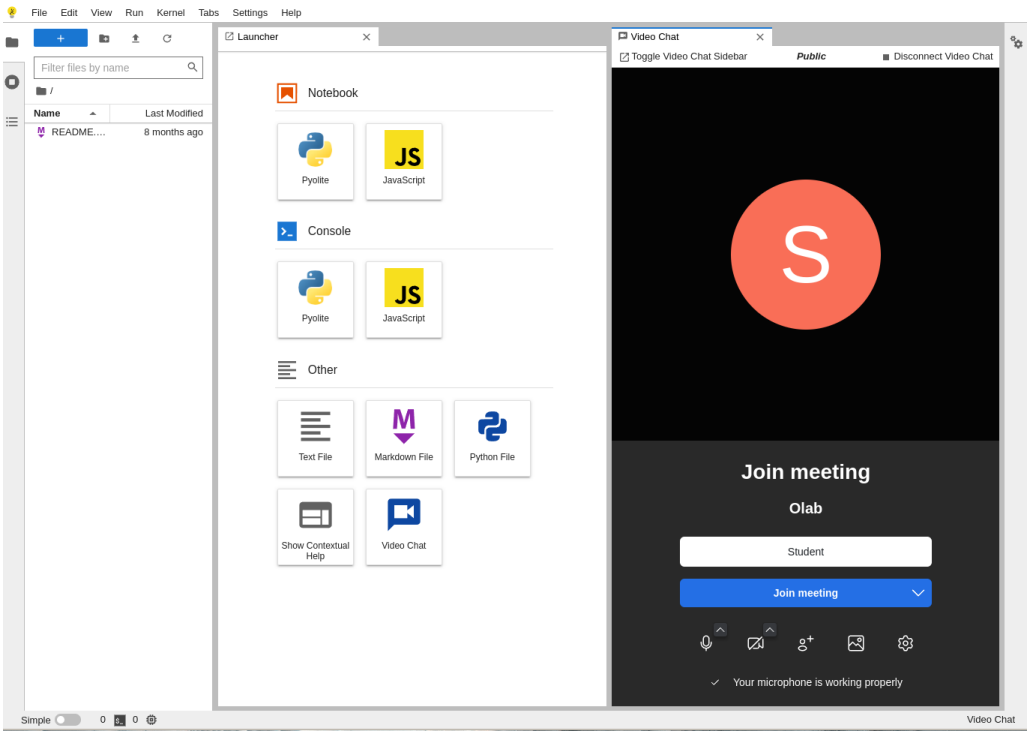


Fig. 9.23: Screenshot of the Jitsi video chat panel running in JupyterLab courtesy of the `jupyterlab-videochat` extension. Jitsi provides the conventional video-conferencing features that students are already familiar with, conveniently integrated with the JupyterLab application.

oriented user interface, and design a visual widget that represents the equipment state (e.g. the front display panel of the HV amplifier). Another, more involved approach would be to replace the form input fields with an entirely skeumorphic UI, e.g. replacing sliders with control knobs. As a component of these approaches, additional development could be spent in the direction of teaching students how to assemble an experiment in an interactive manner. As these are inherently visual changes, considerable effort would be required to implement and validate the approaches.

Part IV

Discussion

CHAPTER

TEN

CONCLUSIONS

A workflow for feature analysis of the raw GET waveforms was established, using modern columnar analysis tools and techniques to ensure high throughput and convenient re-analysis. Alongside this analysis, a suite of investigative tools was developed to aid the investigation of single events at different stages of the analysis workflow, such as 3D event viewers, 2D hit-map reconstruction, and 1D waveform analysis. Behind the discussion in this thesis lies a considerable amount of software designed to provide high-throughput analysis, with a strong consideration of the importance of iterative feedback and exploratory analysis. Despite its omission due to the length of this thesis, this programming work included the application of columnar programming techniques, JIT-compiled analysis routines, and research into random sampling techniques. An understanding of the limitations associated with the TexAT detector was developed, with particular focus upon the problem of "ringing" in the high-gain MicroMeGaS regions, and the challenge associated with strip-chain agreement for near-planar events. As such, strategies for identifying the GET waveforms with ringing artefacts, and extracting a usable signal were developed. This proved to take considerable amounts of work, which involved identifying commonalities amongst signals that suffered this phenomenon, and exploring different approaches, such as wavelet transforms, deconvolutions, and spatio-temporal pattern matching to search for a solution. Though amenable to recovery using symmetries of the detector, poor performance of the silicon detector array was observed to diminish the statistics of subsequent analysis work. In the absence of information concerning the detector configuration, such as the displacement of the silicon detector wall with respect to the center of the detector, hitmaps of reconstructed tracks against the silicon quadrants were made to infer an adjustment factor.

Simulations of the properties of the $^{4}\mathrm{He} + \mathrm{CO}_2$ gas mixture were made to support simulation and analysis of novel reconstruction approaches in the $^{4}\mathrm{He} + ^{10}\mathrm{C}$ scattering reaction. The drift velocity, drift resolution, and stopping powers of relevant ions were computed for the gas mixture; wrapper codes for both SRIM and MSTAR were improved / developed, and these codes investigated for the heavy ion (^{10}C) energy loss calculations. Investigations into the use of Hough transform, RANSAC, and graph-cut methods for track finding and fitting were performed, with a comparison made between state of the art sequential RANSAC and the novel PEARL method for a set of exemplar events. To generate these events, the reaction and electronic simulation packages (TexATSim and TexATResponse) were heavily modified to improve the

user experience, and correctness of the cluster diffusion resolution model. These tools were then used to generate test cases against which the PEaRL-based method for track fitting was developed. A line-interval model for finite-length tracks was derived, and used alongside the PEARL method to reconstruct particle tracks recorded by the TexAT TPC detector at the Texas A & M Cyclotron Institute, College Station, United States.

An exploration of the literature was conducted to determine the structures that might be populated in the elastic alpha-scattering entrance channel of the ^{14}O nucleus. An initial, highly preliminary exploration of the experimental results was made towards investigating the states populated in ^{14}O via the elastic scattering channel. A preliminary excitation function was derived for the different reconstruction methods, whose shortcomings were discussed. The inability to discriminate between the elastic and inelastic contributions to the zero-degree excitation function was identified as a fundamental limitation to further analysis of these data. In support of future analysis of these data, the theory behind constrained optimisation of track fitting was explored. Future work should investigate the use of R-Matrix fitting codes such as AZURE2 in the fitting of angular cross-section functions in order to assign spin-parities to measured states. Alongside work to compute the final reaction kinematics, an exploration of the beam energy reconstructed through a full kinematic analysis was made in comparison to that derived from the stopping power beam in the TPC fill gas. A focus was placed upon the auxiliary data required to reconstruct TPC reactions, with reference to the need for measured stopping powers of the light particle within the active medium. In the absence of these data, the aforementioned simulated values were used.

In the second half of this thesis, a virtual online laboratory for nuclear physics students was developed, with a strong focus on reducing the time-to-deployment, familiarity to existing laboratory users, and integration with existing Learning Tools Interoperability (LTI)-aware Virtual Learning Environments (VLE)s. A series of virtual laboratory experiments were designed around existing in-person experiences, with requisite support material to facilitate self-sufficient learning. A cloud-based JupyterHub deployment was developed using Kubernetes to provide distance-learning and self-isolating students with access to per-user compute environments outside of the University campus network. Provision was made for future extensions of the laboratory including cloud-based analysis to facilitate future extensions into hybrid-learning. The high-level usage data for an initial cohort of students were recorded and analysed to infer usage patterns, and task engagement.

Future work to further the O-Lab provision should look to extend the self-supervised analysis capabilities of the system, such as integration of a data analysis environment that permits students to develop and extend their programming and analysis skills. This could be done in-line with existing undergraduate taught programming modules, to ensure that these skills are applied in a practical context. Additional experiments may be added to provide a more comprehensive introduction to the nuclear laboratory, with a particular focus on beyond-practical experiments i.e. experiments which take advantage of the low-risk, digital nature of the laboratory.

CHAPTER

ELEVEN

FUTURE WORK

The data analysed in this thesis are amenable to further investigation. A particular challenge in this work has been the measurement of the ion stopping powers in the TPC gas, which are fundamental to the reconstruction of the reaction kinematics. Future efforts to measure these stopping powers in order to accurately model energy loss in the gas would likely bring the different beam-energy reconstruction approaches outlined in Section 8.13.4 into agreement, subject to the energy and angular straggling incurred by the beam as it moves through the TexAT TPC entrance window.

The problem of ringing in the MicroMeGaS signals may also be further addressed in future works. Whilst the methods outlined in this thesis address many of these waveforms, a more robust approach could be introduced to better classify non-significant waveform regions. Machine learning methods are of an increasing relevance in the space of signal classification, and future works may be well positioned to take advantage of the rapid pace of contemporary developments in this space.

The ability to separate the elastic and inelastic contributions to the zero-degree reconstruction outlined in Section 8.13.2 depends upon the ability to identify the interaction vertex. Without a well defined Bragg peak, one could instead identify a step change in the $\frac{\mathrm{d}E}{\mathrm{d}x}$ of the pads region which corresponds to the interaction vertex. Otherwise, it may also be possible to identify the position at which the energy lost in the MicroMeGaS no longer includes contributions from the heavy-scatter track. In such a scenario it would be kinematically possible to identify the interaction vertex, and thus discriminate elastic and inelastic scatters.

With these data, future work may compute differential cross sections to aid in the assignment of spins and parities to the observed levels, in order to support theoretical predictions of the structure of excited states.

It would be beneficial to build upon the success of the TexAT TPC detector to improve the position resolution, and reduce the artefacts originating from inhomogeneity in the MicroMeGaS anode. Replacing the thick silicon quadrant detectors with particle telescopes (e.g. thin silicon, thick scintillator) would remove a degree of ambiguity in the particle identification procedure that derives from the poor charge sensitivity of the MicroMeGaS anode. Furthermore, increasing the coverage of the particle detectors to 2π would improve the efficiency of the detector, enhancing the statistics of the dataset and its derived measures. By replacing the strip-chain layout of the MicroMeGaS with an homogenous pixelated

anode would eliminate the challenges of track reconstruction that pertain to the interaction between the central and side regions of the anode. This could be further enhanced with the introduction of a resistive strip anode, which would provide sub-pad resolution of charge clusters measured in the TPC [123].

Part V

End Matter

CHAPTER

TWELVE

APPENDIX

Listing 12.1: Sample from the O-Lab Dockerfile configuration. This Dockerfile was used to build both the production and development targets, using a multi-stage build.

```
# Base image ##############################################################################
 ↪######
FROM rootproject/root:6.22.00-ubuntu20.
 ↪04@sha256:f1013a28f03e91a886ce2e2a3954428258c4fe8901770645b286b159e5fcf20e as base

# Derived from https://github.com/thomasms/docker-g4root/blob/master/ubuntu/
 ↪Dockerfile.ubuntu18.04
MAINTAINER Angus Hollands [redacted]

# See http://label-schema.org
LABEL org.label-schema.name="O-Lab Runtime Container" \
      org.label-schema.description="Runtime environment for O-Lab project" \
      org.label-schema.url="https://gitlab.com/uob-nuclear-olab/olab" \
      org.label-schema.vcs-url="https://gitlab.com/uob-nuclear-olab/olab.git" \
      org.label-schema.vendor="" \
      org.label-schema.license="Apache-2.0" \
      org.label-schema.schema-version="1.0"



# Allow builder to specify UID
ARG NB_USER=olab
ARG NB_UID=1000
```

(continues on next page)

(continued from previous page)

```
ENV NB_USER ${NB_USER}
ENV NB_UID ${NB_UID}
ENV PYTHONDONTWRITEBYTECODE 1
ENV POETRY_VIRTUALENVS_CREATE 0
ENV POETRY_HOME=/opt/poetry
ENV POETRY_VERSION=1.1.0rc1
ENV POETRY_PREVIEW=1

# Load requirejs for ROOT
ADD https://requirejs.org/docs/release/2.3.6/minified/require.js ${ROOTSYS}/js/
 ↪components/requirejs/require.js

# Install packages
RUN apt-get update -qq \
 && apt-get -y install python3-pip python3-venv git ca-certificates \
 && rm -rf /var/lib/apt/lists/* \
 # Add a user
 && adduser --disabled-password \
     --gecos "Default user" \
     --uid ${NB_UID} \
     ${NB_USER} \
  && chown -R ${NB_UID} ${HOME} \
  && echo "${NB_USER}:${NB_USER}" | chpasswd \
 # Fix PyROOT v22.
 && sed -i "s/gStyle, style)/gStyle, obj)/" ${ROOTSYS}/js/scripts/JSRootPainter.v6.js␣
 ↪\
 # Change permissions of downloaded requirejs
 && chmod 775 ${ROOTSYS}/js/components/requirejs/require.js

 # Install poetry
RUN curl -sSL https://raw.githubusercontent.com/python-poetry/poetry/master/get-
 ↪poetry.py | python3
RUN chmod og+x -R ${POETRY_HOME}/bin
# Install poetry as user
ENV PATH="${POETRY_HOME}/bin:${PATH}"
```

(continues on next page)

(continued from previous page)

```
# Setup virtualenv
ENV VIRTUAL_ENV=/opt/olab-venv
RUN python3 -m venv "$VIRTUAL_ENV" \
  && chown -R "${NB_USER}:${NB_USER}" "${VIRTUAL_ENV}"
ENV PATH="$VIRTUAL_ENV/bin:$PATH"



# Production (base) image  ############################################################
 ↪###############
FROM base as production-base

WORKDIR /usr/src/olab

# Setup user
RUN chown -R "${NB_USER}:${NB_USER}" /usr/src/olab
USER ${NB_USER}

# Install dependencies
COPY --chown="${NB_USER}:${NB_USER}" pyproject.toml poetry.lock ./
COPY --chown="${NB_USER}:${NB_USER}" extern extern/
RUN poetry install --no-root --no-dev

# Copy Jupyter configuration directory
COPY --chown="${NB_USER}:${NB_USER}" etc/jupyter /usr/local/etc/jupyter/
COPY --chown="${NB_USER}:${NB_USER}" share/jupyter /usr/local/share/jupyter/

# Install olab
COPY --chown="${NB_USER}:${NB_USER}" olab olab/

# Install dependencies
RUN poetry install --no-dev

# Copy resources
COPY --chown="${NB_USER}:${NB_USER}" notebooks notebooks/
```

(continues on next page)

(continued from previous page)

```
COPY --chown="${NB_USER}:${NB_USER}" templates templates/



# Production image  ##################################################################
 ↪##############
FROM production-base as production
WORKDIR /usr/src/olab/notebooks
ENTRYPOINT ["jupyterhub-singleuser"]



# Test image ###########################################################################
 ↪##############
FROM production-base as test
COPY tests tests/
RUN poetry install
ENTRYPOINT ["/usr/bin/env", "bash"]



# Development image ####################################################################
 ↪##############
FROM base as development
USER ${NB_USER}
ENTRYPOINT ["/usr/bin/env", "bash"]
```

CHAPTER

THIRTEEN

BIBLIOGRAPHY

[1] W. von Oertzen, Martin Freer, and Yoshiko Kanada-En'yo. Nuclear clusters and nuclear molecules. *Physics Reports*, 432(2):43–113, September 2006. doi:10.1016/j.physrep.2006.07.001.

[2] J. Alme, Y. Andres, H. Appelshäuser, S. Bablok, N. Bialas, R. Bolgen, U. Bonnes, R. Bramm, P. Braun-Munzinger, R. Campagnolo, P. Christiansen, A. Dobrin, C. Engster, D. Fehlker, Y. Foka, U. Frankenfeld, J. J. Gaardhøje, C. Garabatos, P. Glässel, C. Gonzalez Gutierrez, P. Gros, H. -A. Gustafsson, H. Helstrup, M. Hoch, M. Ivanov, R. Janik, A. Junique, A. Kalweit, R. Keidel, S. Kniege, M. Kowalski, D. T. Larsen, Y. Lesenechal, P. Lenoir, N. Lindegaard, C. Lippmann, M. Mager, M. Mast, A. Matyja, M. Munkejord, L. Musa, B. S. Nielsen, V. Nikolic, H. Oeschler, E. K. Olsen, A. Oskarsson, L. Osterman, M. Pikna, A. Rehman, G. Renault, R. Renfordt, S. Rossegger, D. Röhrich, K. Røed, M. Richter, G. Rueshmann, A. Rybicki, H. Sann, H. -R. Schmidt, M. Siska, B. Sitár, C. Soegaard, H. -K. Soltveit, D. Soyk, J. Stachel, H. Stelzer, E. Stenlund, R. Stock, P. Strmeň, I. Szarka, K. Ullaland, D. Vranic, R. Veenhof, J. Westergaard, J. Wiechula, and B. Windelband. The ALICE TPC, a large 3-dimensional tracking device with fast readout for ultra-high multiplicity events. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 622(1):316–367, October 2010. doi:10.1016/j.nima.2010.04.042.

[3] C. F. v. Weizsäcker. Zur Theorie der Kernmassen. *Z. Physik*, 96(7):431–458, July 1935. doi:10.1007/BF01337700.

[4] David J. Rowe and John L. Wood. *Fundamentals of Nuclear Models*. World Scientific Publ, Singapore [u.a.], 2010. ISBN 981-256-956-1.

[5] D. J. Rowe. *Nuclear Collective Motion: Models and Theory*. World Scientific, 2010. ISBN 978-981-279-064-4.

[6] A. Bohr, B. R. Mottelson, Gregory Breit, and Gerald E. Brown. Nuclear Structure, Vol. 2: Nuclear Deformations. *Physics Today*, 30(3):59–62, March 1977. doi:10.1063/1.3037453.

[7] Christian Beck, editor. *Clusters in Nuclei*. Volume 3 of Lecture Notes in Physics. Springer International Publishing, 2014. ISBN 978-3-319-01076-2 978-3-319-01077-9. doi:10.1007/978-3-319-01077-9.

[8] J. M. Eisenberg and W. Greiner. *Microscopic Theory of the Nucleus*. North-Holland Publishing Company, 1976. ISBN 978-0-7204-0484-5. 76359606

76359606.

[9] Chrysovalantis Constantinou. Natural orbitals for the no-core configuration interaction approach. April 2017. Comment: 85 pages, 25 figures, thesis. arXiv:1704.02363, doi:10.48550/arXiv.1704.02363.

[10] Vladimir Zelevinsky and Alexander Volya. *Physics of Atomic Nuclei*. John Wiley & Sons, Incorporated, Somerset, GERMANY, 2017. ISBN 978-3-527-69360-3.

[11] Jouni Suhonen. *From Nucleons to Nucleus : Concepts of Microscopic Nuclear Theory*. Springer, Theoretical and Mathematical Physics, 1st ed. 2007. edition, 2007. ISBN 978-3-540-48859-0.

[12] J. W. Negele. The mean-field theory of nuclear structure and dynamics. *Rev. Mod. Phys.*, 54(4):913–1015, October 1982. doi:10.1103/RevModPhys.54.913.

[13] A. L. Fetter, J. D. Walecka, and Leo P. Kadanoff. Quantum Theory of Many Particle Systems. *Physics Today*, 25(11):54–55, November 1972. doi:10.1063/1.3071096.

[14] Peter Ring and Peter Schuck. *The Nuclear Many-Body Problem*. Springer Berlin, Heidelberg, 2004. ISBN 978-3-540-21206-5.

[15] Walter Greiner and Joachim A. Maruhn. *Nuclear Models*. Springer, Berlin, Heidelberg, 1996. ISBN 978-3-540-78046-5 978-3-642-60970-1. doi:10.1007/978-3-642-60970-1.

[16] J. M. Eisenberg and W. Greiner. *Nuclear Theory: Nuclear Models*. North-Holland Publishing Company, 1975. ISBN 978-0-7204-0326-8.

[17] M. Anguiano, A. M. Lallena, G. Co', and V. De Donno. A study of self-consistent Hartree–Fock plus Bardeen–Cooper–Schrieffer calculations with finite-range interactions. *J. Phys. G: Nucl. Part. Phys.*, 41(2):025102, January 2014. doi:10.1088/0954-3899/41/2/025102.

[18] R. F. Casten. *Nuclear Structure from a Simple Perspective*. Oxford University Press, Cary, January 1990. ISBN 978-0-19-504599-4.

[19] W. Nazarewicz and J. Dobaczewski. Dynamical symmetries, multiclustering, and octupole susceptibility in superdeformed and hyperdeformed nuclei. *Phys. Rev. Lett.*, 68(2):154–157, January 1992. doi:10.1103/PhysRevLett.68.154.

[20] Bruce R. Barrett, Petr Navrátil, and James P. Vary. Ab initio no core shell model. *Progress in Particle and Nuclear Physics*, 69:131–181, March 2013. doi:10.1016/j.ppnp.2012.10.003.

[21] Christian Beck, editor. *Clusters in Nuclei*. Volume 1 of Lecture Notes in Physics. Springer, Berlin, Heidelberg, 2010. ISBN 978-3-642-13898-0 978-3-642-13899-7. doi:10.1007/978-3-642-13899-7.

[22] Yoshiko Kanada-En'yo, Masaaki Kimura, and Akira Ono. Antisymmetrized molecular dynamics and its applications to cluster phenomena. *Progress of Theoretical and Experimental Physics*, 2012(1):2898, 2012. doi:10.1093/ptep/pts001.

[23] H. Feldmeier and J. Schnack. Fermionic molecular dynamics SCV252SCV133 V2. *Progress in Particle and Nuclear Physics*, 39:393–442, January 1997. doi:10.1016/S0146-6410(97)00047-1.

[24] Peter Kramer and Marcos Saraceno, editors. *The Time-Dependent Variational Principle (TDVP)*, pages 3–14. Lecture Notes in Physics. Springer, Berlin, Heidelberg, 1981. doi:10.1007/3-540-10579-4_20.

[25] L. Wilets, E. M. Henley, M. Kraft, and A. D. Mackellar. Classical many-body model for heavy-ion collisions incorporating the Pauli principle. *Nuclear Physics A*, 282(2):341–350, May 1977. doi:10.1016/0375-9474(77)90220-2.

[26] H. Horiuchi, Y. Kanada-En'yo, and A. Ono. Neutron-rich nuclei studied with AMD. In G. S. Anagnostatos and W. von Oertzen, editors, *Atomic and Nuclear Clusters*, 142–146. Berlin, Heidelberg, 1995. Springer. doi:10.1007/978-3-642-79696-8_32.

[27] M. Anguiano, J. L. Egido, and L. M. Robledo. Coulomb exchange and pairing contributions in nuclear Hartree–Fock–Bogoliubov calculations with the Gogny force. *Nuclear Physics A*, 683(1):227–254, February 2001. doi:10.1016/S0375-9474(00)00445-0.

[28] T. Neff and H. Feldmeier. Nuclear structure and reactions in the fermionic molecular dynamics approach. *J. Phys.: Conf. Ser.*, 111:012007, May 2008. doi:10.1088/1742-6596/111/1/012007.

[29] Gaotsiwe J. Rampho and Sofianos A. Sofianos. *Antisymmetrized Molecular Dynamics and Nuclear Structure*. IntechOpen, April 2012. ISBN 978-953-51-0443-8. doi:10.5772/34963.

[30] Y. Kanada-En'yo. Variation after Angular Momentum Projection for the Study of Excited States Based on Antisymmetrized Molecular Dynamics. *Phys. Rev. Lett.*, 81(24):5291–5293, December 1998. doi:10.1103/PhysRevLett.81.5291.

[31] Hans Feldmeier and Thomas Neff. Nuclear clustering in fermionic molecular dynamics. In *Nuclear Particle Correlations and Cluster Physics*, pages 71–94. WORLD SCIENTIFIC, February 2017. doi:10.1142/9789813209350_0003.

[32] Chun Wa Wong. Generator-coordinate methods in nuclear physics. *Physics Reports*, 15(5):283–357, January 1975. doi:10.1016/0370-1573(75)90036-8.

[33] Martin Freer. The clustered nucleus—cluster structures in stable and unstable nuclei. *Rep. Prog. Phys.*, 70(12):2149–2210, November 2007. doi:10.1088/0034-4885/70/12/R03.

[34] W. N. Catford. Clustering in Nuclei from N/Z=1 to N/Z=2. February 2013. Comment: Proc. Int. Conf. held at Faure, Cape Town, January 1999. arXiv:1302.3849.

[35] P. McEwan and M. Freer. Characterization of molecular structures in the deformed harmonic oscillator. *J. Phys. G: Nucl. Part. Phys.*, 30(4):447–463, February 2004. doi:10.1088/0954-3899/30/4/005.

[36] 3.5: Diatomic Molecules. https://chem.libretexts.org/Sandboxes/marzluff_at_grinnell.edu/Username%3A_marzluff@grinnell.edu/C January 2020.

[37] N. Itagaki and S. Okabe. Molecular orbital structures in 10Be. *Phys. Rev. C*, 61(4):044306, March 2000. doi:10.1103/PhysRevC.61.044306.

[38] D. Scharnweber, U. Mosel, and W. Greiner. Asymptotically Correct Shell Model for Nuclear Fission. *Phys. Rev. Lett.*, 24(11):601–603, March 1970. doi:10.1103/PhysRevLett.24.601.

[39] Michio Seya, Michio Kohno, and Sinobu Nagata. Nuclear Binding Mechanism and Structure of Neutron-Rich Be and B Isotopes by Molecular-Orbital Model. *Progress of Theoretical Physics*, 65(1):204–223, January 1981. doi:10.1143/PTP.65.204.

[40] Wolfram von Oertzen. Nuclear Clusters and Covalently Bound Nuclear Molecules. *Progress of Theoretical Physics Supplement*, 146:169–178, March 2002. doi:10.1143/PTPS.146.169.

[41] W. von Oertzen. Dimers based on the α + α potential and chain states of carbon isotopes. *Z Phys A - Particles and Fields*, 357(4):355–365, December 1997. doi:10.1007/s002180050255.

[42] Kris L. G. Heyde. *Basic Ideas and Concepts in Nuclear Physics*. Inst. of Physics Publ, Bristol [u.a.], 3. ed. edition, 2004. ISBN 978-0-7503-0980-6.

[43] M. Milin and W. von Oertzen. Search for molecular bands in 13C. *Eur Phys J A*, 14(3):295–307, July 2002. doi:10.1140/epja/i2001-10199-6.

[44] E. Wigner. On the Consequences of the Symmetry of the Nuclear Hamiltonian on the Spectroscopy of Nuclei. *Phys.Rev.*, 51(2):106–119, January 1937. doi:10.1103/PhysRev.51.106.

[45] R. Machleidt. High-precision, charge-dependent Bonn nucleon-nucleon potential. *Physical Review C*, January 2001. doi:10.1103/PhysRevC.63.024001.

[46] Kenneth S. Krane. *Introductory Nuclear Physics*. Wiley, New York u.a, 1987. ISBN 978-0-471-80553-3.

[47] M. Freer, R. R. Betts, and A. H. Wuosmaa. Relationship between the deformed harmonic oscillator and clustering in light nuclei. *Nuclear Physics A*, 587(1):36–54, May 1995. doi:10.1016/0375-9474(94)00820-D.

[48] W. von Oertzen. Two-center molecular states in 9B, 9Be, 10Be, and 10B. *Z. Phys. A — Hadrons and Nuclei*, 354(1):37–43, December 1996. doi:10.1007/s002180050010.

[49] N. Itagaki, S. Okabe, K. Ikeda, and I. Tanihata. Molecular-orbital structure in neutron-rich C isotopes. *Phys. Rev. C*, 64(1):014301, May 2001. doi:10.1103/PhysRevC.64.014301.

[50] R. Bijker and F. Iachello. The Algebraic Cluster Model: Three-Body Clusters. *Annals of Physics*, 298(2):334–360, June 2002. doi:10.1006/aphy.2002.6255.

[51] N. Itagaki, T. Otsuka, K. Ikeda, and S. Okabe. Equilateral-Triangular Shape in 14C. *Phys. Rev. Lett.*, 92(14):142501, April 2004. doi:10.1103/PhysRevLett.92.142501.

[52] W. von Oertzen, H. G. Bohlen, M. Milin, Tz Kokalova, S. Thummerer, A. Tumino, R. Kalpakchieva, T. N. Massey, Y. Eisermann, G. Graw, T. Faestermann, R. Hertenberger, and H.-F. Wirth. Search for cluster structure of excited states in 14C. *Eur. Phys. J. A*, 21(2):193–215, August 2004. doi:10.1140/epja/i2003-10188-9.

[53] Tadahiro Suhara and Yoshiko Kanada-En'yo. Cluster structures of excited states in 14C. *J. Phys.: Conf. Ser.*, 321:012047, September 2011. doi:10.1088/1742-6596/321/1/012047.

[54] M. Freer, J. D. Malcolm, N. L. Achouri, N. I. Ashwood, D. W. Bardayan, S. M. Brown, W. N. Catford, K. A. Chipps, J. Cizewski, N. Curtis, K. L. Jones, T. Munoz-Britton, S. D. Pain, N. Soić, C. Wheldon, G. L. Wilson, and V. A. Ziman. Resonances in 14C observed in the 4He(10Be,α)10Be reaction. *Phys. Rev. C*, 90(5):054324, November 2014. doi:10.1103/PhysRevC.90.054324.

[55] T. Baba and M. Kimura. Structure and decay pattern of the linear-chain state in 14C. *Phys. Rev. C*, 94(4):044303, October 2016. doi:10.1103/PhysRevC.94.044303.

[56] H. Yamaguchi, D. Kahl, S. Hayakawa, Y. Sakaguchi, K. Abe, T. Nakao, T. Suhara, N. Iwasa, A. Kim, D. H. Kim, S. M. Cha, M. S. Kwag, J. H. Lee, K. B. Lee, E. J. Lee, P. Lee, C. S. Lee, K. Y. Chae, Y. Wakabayashi, N. Imai, N. Kitamura, J. Y. Moon, C. Akers, H. S. Jung, N. N. Duy, and L. H. Khiem. Experimental investigation of a linear-chain structure in the nucleus 14C. *Physics Letters B*, 766(C):11–16, March 2017. doi:10.1016/j.physletb.2016.12.050.

[57] T. Baba and M. Kimura. Coulomb shift in the mirror pair of 14C and 14O as a signature of the linear-chain structure. *Phys. Rev. C*, 99(2):021303, February 2019. doi:10.1103/PhysRevC.99.021303.

[58] Makoto Ito. Cluster Thomas-Ehrman Shift in 10Be–10C. *EPJ Web of Conferences*, 117:06014, 2016. doi:10.1051/epjconf/201611706014.

[59] Niels Bohr. Neutron Capture and Nuclear Constitution. *Nature*, 137(3461):344–348, February 1936. doi:10.1038/137344a0.

[60] G. R. Satchler. *Introduction to Nuclear Reactions*. Palgrave Macmillan UK, London, 1990. ISBN 978-0-333-51484-9 978-1-349-20531-8. doi:10.1007/978-1-349-20531-8.

[61] L. Mandelstam and \relax Ig. Tamm. The Uncertainty Relation Between Energy and Time in Non-relativistic Quantum Mechanics. In I. E. Tamm, Boris M. Bolotovskii, Victor Ya. Frenkel, and Rudolf Peierls, editors, *Selected Papers*, pages 115–123. Springer, Berlin, Heidelberg, 1991. doi:10.1007/978-3-642-74626-0_8.

[62] B. H. Bransden. *Quantum Mechanics*. Pearson, Harlow, England ; New York, 2nd edition edition, January 2000. ISBN 978-0-582-35691-7.

[63] Jos Uffink. The rate of evolution of a quantum state. *American Journal of Physics*, 61(10):935–936, October 1993. doi:10.1119/1.17368.

[64] Jay N. Marx and David R. Nygren. The Time Projection Chamber. *Physics Today*, 31(10):46–53, October 1978. doi:10.1063/1.2994775.

[65] H. J. Hilke. Time projection chambers. *Rep. Prog. Phys.*, 73(11):116201, October 2010. doi:10.1088/0034-4885/73/11/116201.

[66] J. D. Cockcroft. Experimental Nuclear Physics. *Nature*, 175(4445):53–54, January 1955. doi:10.1038/175053a0.

[67] Kristen A. Recine, James B. R. Battat, and Shawn Henderson. Understanding Current Signals Induced by Drifting Electrons. *American Journal of Physics*, 82(4):322–330, April 2014. Comment: 21 pages, 12 figures. arXiv:1404.3802, doi:10.1119/1.4864642.

[68] Glenn F. Knoll. Radiation detection and measurement / Glenn F. Knoll. 1989. ID: 44BIR_ALEPH_DS000123890.

[69] Ö Şahin, İ Tapan, E. N. Özmutlu, and R. Veenhof. Penning transfer in argon-based gas mixtures. *J. Inst.*, 5(05):P05002, May 2010. doi:10.1088/1748-0221/5/05/P05002.

[70] Marco Cortesi, Wolfgang Mittig, Daniel Bazin, Yassid Ayyad Limonge, Saul Beceiro-Novo, Rim Soussi Tanani, Michael Wolff, John Yurkon, and Andreas Stolz. Recent advances with a hybrid micro-pattern gas detector operated in low pressure H2 and He, for AT-TPC applications. *EPJ Web Conf.*, 174:01007, 2018. doi:10.1051/epjconf/201817401007.

[71] D. Suzuki, M. Ford, D. Bazin, W. Mittig, W. G. Lynch, T. Ahn, S. Aune, E. Galyaev, A. Fritsch, J. Gilbert, F. Montes, A. Shore, J. Yurkon, J. J. Kolata, J. Browne, A. Howard, A. L. Roberts, and X. D. Tang. Prototype AT-TPC: Toward a new generation active target time projection chamber for radioactive beam experiments. *Nu-*

clear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 691:39–54, November 2012. doi:10.1016/j.nima.2012.06.050.

[72] James F. Ziegler, M. D. Ziegler, and J. P. Biersack. SRIM – The stopping and range of ions in matter (2010). *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 268(11):1818–1823, June 2010. doi:10.1016/j.nimb.2010.02.091.

[73] E. Koshchiy, G. V. Rogachev, E. Pollacco, S. Ahn, E. Uberseder, J. Hooker, J. Bishop, E. Aboud, M. Barbui, V. Z. Goldberg, C. Hunt, H. Jayatissa, C. Magana, R. O'Dwyer, B. T. Roeder, A. Saastamoinen, and S. Upadhyayula. Texas Active Target (TexAT) detector for experiments with rare isotope beams. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 957:163398, March 2020. doi:10.1016/j.nima.2020.163398.

[74] D. H. Youngblood. The Texas A & M K500 cyclotron facility. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 56–57:991–995, May 1991. doi:10.1016/0168-583X(91)95079-S.

[75] J. Hooker, G. V. Rogachev, E. Koshchiy, S. Ahn, M. Barbui, V. Z. Goldberg, C. Hunt, H. Jayatissa, E. C. Pollacco, B. T. Roeder, A. Saastamoinen, and S. Upadhyayula. Structure of 9C through proton resonance scattering with Texas Active Target detector. *Phys. Rev. C*, 100(5):054618, November 2019. Comment: 9 pages, 12 figures. arXiv:1903.01402, doi:10.1103/PhysRevC.100.054618.

[76] M. Barbui, A. Volya, E. Aboud, S. Ahn, J. Bishop, V. Z. Goldberg, J. Hooker, C. H. Hunt, H. Jayatissa, \relax Tz. Kokalova, E. Koshchiy, S. Pirrie, E. Pollacco, B. T. Roeder, A. Saastamoinen, S. Upadhyayula, C. Wheldon, and G. V. Rogachev. α-cluster structure of 18Ne. *Phys. Rev. C*, 106(5):054310, November 2022. doi:10.1103/PhysRevC.106.054310.

[77] TexAT TPC – Texas Active Target Time Projection Chamber.

[78] E.C. Pollacco, G.F. Grinyer, F. Abu-Nimeh, T. Ahn, S. Anvar, A. Arokiaraj, Y. Ayyad, H. Baba, M. Babo, P. Baron, D. Bazin, S. Beceiro-Novo, C. Belkhiria, M. Blaizot, B. Blank, J. Bradt, G. Cardella, L. Carpenter, S. Ceruti, E. De Filippo, E. Delagnes, S. De Luca, H. De Witte, F. Druillole, B. Duclos, F. Favela, A. Fritsch, J. Giovinazzo, C. Gueye, T. Isobe, P. Hellmuth, C. Huss, B. Lachacinski, A.T. Laffoley, G. Lebertre, L. Legeard, W.G. Lynch, T. Marchi, L. Martina, C. Maugeais, W. Mittig, L. Nalpas, E.V. Pagano, J. Pancin, O. Poleshchuk, J.L. Pedroza, J. Pibernat, S. Primault, R. Raabe, B. Raine, A. Rebii, M. Renaud, T. Roger, P. Roussel-Chomaz, P. Russotto, G. Saccà, F. Saillant, P. Sizun, D. Suzuki, J.A. Swartz, A. Tizon, A. Trifiró, N. Usher, G. Wittwer, and J.C. Yang. GET: A generic electronics system for TPCs and nuclear physics instrumentation. *Nuclear Instru-*

ments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 887:81–93, April 2018. doi:10.1016/j.nima.2018.01.020.

[79] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, June 1981. doi:10.1145/358669.358692.

[80] J. Hadamard. Sur les problèmes aux dérivés partielles et leur signification physique. *Princeton University Bulletin*, 13:49–52, 1902.

[81] J. Cheriyan and R. Ravi. Approximation algorithms for network problems. 1998.

[82] Alfred A. Kuehn and Michael J. Hamburger. A Heuristic Program for Locating Warehouses. *Management Science*, 9(4):643–666, 1963. arXiv:2627368.

[83] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, November 2001. doi:10.1109/34.969114.

[84] Andrew Delong, Anton Osokin, Hossam N. Isack, and Yuri Boykov. Fast approximate energy minimization with label costs. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2173–2180. June 2010. doi:10.1109/CVPR.2010.5539897.

[85] L. R. Ford and D. R. Fulkerson. Maximal Flow Through a Network. *Canadian Journal of Mathematics*, 8:399–404, 1956/ed. doi:10.4153/CJM-1956-045-5.

[86] V. Kolmogorov and R. Zabin. What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):147–159, February 2004. doi:10.1109/TPAMI.2004.1262177.

[87] Hossam Isack and Yuri Boykov. Energy-Based Geometric Multi-model Fitting. *Int J Comput Vis*, 97(2):123–147, April 2012. doi:10.1007/s11263-011-0474-7.

[88] A. G. Frodesen, O. Skjeggestad, and H. Tøfte. *Probability and Statistics in Particle Physics*. Universitetsforlaget, Bergen, 1979.

[89] S Anvar, R Fox, B Raine, F Saillant, and N Usher. The MultiFrame Metaformat Specification.

[90] Nick Smith. The Case for Columnar Analysis (a two-part series). 2019.

[91] The HDF5 Group. Hierarchical Data Format, version 5. 1997.

[92] P. Deutsch. GZIP file format specification version 4.3. Technical Report RFC1952, RFC Editor, May 1996. doi:10.17487/rfc1952.

[93] Apache Parquet. Apache Software Foundation.

[94] J. Giovinazzo, T. Goigoux, S. Anvar, P. Baron, B. Blank, E. Delagnes, G. F. Grinyer, J. Pancin, J. L. Pedroza, J. Pibernat, E. Pollacco, A. Rebii, T. Roger, and P. Sizun. GET electronics samples data analysis. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 840:15–27, December 2016. doi:10.1016/j.nima.2016.09.018.

[95] Miroslav Morha\textasciicaron c and Vladislav Matou\textasciicaron sek. MULTIDIMENSIONAL FFT BASED POSITIVE DEFINITE GOLD DECONVOLUTION ALGORITHM. *Journal of ELECTRICAL ENGINEERING*, pages 5, 2005.

[96] Arsen Subashiev and Serge Luryi. Correlation effects in sequential energy branching: an exact model of the Fano statistics. *Phys. Rev. E*, 81(2):021123, February 2010. Comment: 11 pages, 4 figures. Submitted to Physical Review E. arXiv:0911.1532, doi:10.1103/PhysRevE.81.021123.

[97] I. V. Kotov, H. Neal, and P. O'Connor. Pair creation energy and Fano factor of silicon measured at 185 K using 55Fe X-rays. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 901:126–132, September 2018. doi:10.1016/j.nima.2018.06.022.

[98] R. C. Alig, S. Bloom, and C. W. Struck. Scattering by ionization and phonon emission in semiconductors. *Phys. Rev. B*, 22(12):5565–5582, December 1980. doi:10.1103/PhysRevB.22.5565.

[99] F. Scholze, H. Rabus, and G. Ulm. Mean energy required to produce an electron-hole pair in silicon for photons of energies between 50 and 1500 eV. *Journal of Applied Physics*, 84(5):2926–2939, 1998. doi:10.1063/1.368398.

[100] V. Palladino and B. Sadoulet. Application of Classical Theory of Electrons in Gases to Drift Proportional Chambers. *Nucl.Instrum.Meth.*, 128:323, 1975. doi:10.1016/0029-554X(75)90682-5.

[101] S. F. Biagi. Monte Carlo simulation of electron drift and diffusion in counting gases under the influence of electric and magnetic fields. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 421(1):234–240, January 1999. doi:10.1016/S0168-9002(98)01233-9.

[102] R. K. Carnegie, M. S. Dixit, J. Dubeau, D. Karlen, J. -P. Martin, H. Mes, and K. Sachs. Resolution studies of cosmic-ray tracks in a TPC with GEM readout. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 538(1):372–383, February 2005. doi:10.1016/j.nima.2004.08.132.

[103] Helmut Paul and Andreas Schinner. Empirical stopping power tables for ions from 3Li to 18Ar and from 0.001 to 1000MeV/nucleon in solids and gases. *Atomic Data and Nuclear Data Tables*, 85(2):377–452, November 2003. doi:10.1016/j.adt.2003.08.003.

[104] D. Torresi, J. Bishop, D. Carbone, M. Cavallaro, A. Di Pietro, J. P. Fernández Garcia, P. Figuera, M. Fisichella, Tz. Kokalova, M. Lattuada, C. Parascandolo, C. Wheldon, and M. Zadro. Measurement of the stopping power

for 12 C in 4 He gas. This paper is taken from private communications; it is yet to be submitted to a journal., December 2017.

[105] Richard Watermeyer, Tom Crick, Cathryn Knight, and Janet Goodall. COVID-19 and digital disruption in UK universities: afflictions and affordances of emergency online migration. *High Educ*, 81(3):623–641, March 2021. doi:10.1007/s10734-020-00561-y.

[106] Nicholas R. Jones, Zeshan U. Qureshi, Robert J. Temple, Jessica P. J. Larwood, Trisha Greenhalgh, and Lydia Bourouiba. Two metres or one: what is the evidence for physical distancing in covid-19? *BMJ*, 370:m3223, August 2020. doi:10.1136/bmj.m3223.

[107] Petar Jandrić, David Hayes, Paul Levinson, Line Lisberg Christensen, Happiness Onesmo Lukoko, Jimmy Ezekiel Kihwele, James Benedict Brown, Charles Reitz, Peter Mozelius, Harry G. Nejad, Ana Fuentes Martinez, Janine Aldous Arantes, Liz Jackson, Ulrika Gustafsson, Sandra Abegglen, Tom Burns, Sandra Sinfield, Michael Hogan, Pallavi Kishore, Paul R. Carr, Ivana Batarelo Kokić, Paul Prinsloo, Dennis Grauslund, Anne Steketee, Charlotte Achieng-Evensen, Blessing Funmi Komolafe, Juha Suoranta, Nina Hood, Marek Tesar, Jennifer Rose, Niklas Humble, James D. Kirylo, Julia Mañero, Lilia D. Monzó, Mikkel Lodahl, Jimmy Jaldemark, Susan M. Bridges, Navreeti Sharma, Jacob Davidsen, Jānis (John) Ozoliņš, Peter Bryant, Carlos Escaño, Jones Irwin, Kulpreet Kaur, Sarah Pfohl, Kevin Stockbridge, Thomas Ryberg, Olli Pyyhtinen, Suzanne SooHoo, Moses Kayode Hazzan, Jake Wright, Stephanie Hollings, Sonja Arndt, Andrew Gibbons, Shreya Urvashi, Daniella J. Forster, Ian Truelove, Peter Mayo, Glenn Rikowski, Paul Alexander Stewart, Michael Jopling, Georgina Tuari Stewart, Rachel Buchanan, Nesta Devine, Richa Shukla, Rene Novak, Madhav Mallya, Eva Biličić, Sean Sturm, Sahar D. Sattarzadeh, Abey P. Philip, Bridgette Redder, E. Jayne White, Derek R. Ford, Quaylan Allen, Mousumi Mukherjee, and Sarah Hayes. Teaching in the Age of Covid-19—1 Year Later. *Postdigit Sci Educ*, 3(3):1073–1223, October 2021. doi:10.1007/s42438-021-00243-7.

[108] Ton de Jong, Marcia C. Linn, and Zacharias C. Zacharia. Physical and Virtual Laboratories in Science and Engineering Education. *Science*, 340(6130):305–308, April 2013. doi:10.1126/science.1230579.

[109] Lysanne S. Post, Pengyue Guo, Nadira Saab, and Wilfried Admiraal. Effects of remote labs on cognitive, behavioral, and affective learning outcomes in higher education. *Computers & Education*, 140:103596, October 2019. doi:10.1016/j.compedu.2019.103596.

[110] Gülsüm AÅŸıksoy and Didem Islek. The Impact of the Virtual Laboratory on Students' Attitude in a General Physics Laboratory. *International Journal of Online and Biomedical Engineering (iJOE)*, 13(04):20–28, April 2017. doi:10.3991/ijoe.v13i04.6811.

[111] P. Klein, L. Ivanjek, M. N. Dahlkemper, K. Jeličić, M.-A. Geyer, S. Küchemann, and A. Susac. Studying

physics during the COVID-19 pandemic: Student assessments of learning achievement, perceived effectiveness of online recitations, and online laboratories. *Phys. Rev. Phys. Educ. Res.*, 17(1):010117, March 2021. doi:10.1103/PhysRevPhysEducRes.17.010117.

[112] Brian E. Granger and Fernando Perez. Jupyter: Thinking and Storytelling With Code and Data. *Comput. Sci. Eng.*, 23(2):7–14, March 2021. doi:10.1109/MCSE.2021.3059263.

[113] Jeffrey M. Perkel. Why Jupyter is data scientists' computational notebook of choice. *Nature*, 563(7729):145–146, October 2018. doi:10.1038/d41586-018-07196-1.

[114] TIOBE Index. https://www.tiobe.com/tiobe-index/.

[115] Project Jupyter. https://jupyter.org.

[116] Ian Stewart. The Lorenz attractor exists. *Nature*, 406(6799):948–949, August 2000. doi:10.1038/35023206.

[117] Project Jupyter | JupyterHub. https://jupyter.org/hub.

[118] Ani Adhikari, John DeNero, and Michael I. Jordan. Issue 3.2, Spring 2021. *Harvard Data Science Review*, April 2021. doi:10.1162/99608f92.cb0fa8d2.

[119] Björn Grüning, John Chilton, Johannes Köster, Ryan Dale, Nicola Soranzo, Marius van den Beek, Jeremy Goecks, Rolf Backofen, Anton Nekrutenko, and James Taylor. Practical Computational Reproducibility in the Life Sciences. *Cell Systems*, 6(6):631–635, June 2018. doi:10.1016/j.cels.2018.03.014.

[120] Hong Zhu and Ian Bayley. If Docker is the Answer, What is the Question? In *2018 IEEE Symposium on Service-Oriented System Engineering (SOSE)*, 152–163. March 2018. doi:10.1109/SOSE.2018.00027.

[121] Jupyter-videochat — jupyter-videochat 0.6.0 documentation. https://jupyter-videochat.readthedocs.io/en/latest/.

[122] David Oswald and Steffen Kolb. Flat Design vs. Skeuomorphism – Effects on Learnability and Image Attributions in Digital Product Interfaces. *DS 78: Proceedings of the 16th International conference on Engineering and Product Design Education (E&PDE14), Design Education and Human Technology Relations, University of Twente, The Netherlands, 04-05.09.2014*, pages 402–407, 2014.

[123] Jonathan Bortfeld. Development of Micro - Pattern Gaseous Detectors – Micromegas.