

SOLAR KK AXION SEARCH
WITH NEWS-G

by

FRANCISCO ANDRES VAZQUEZ DE SOLA FERNANDEZ

A thesis submitted to the
Department of Physics, Engineering Physics and Astronomy
in conformity with the requirements for
the degree of Doctor of Philosophy

Queen's University

Kingston, Ontario, Canada

September 2020

Copyright © Francisco Andres Vazquez de Sola Fernandez, 2020

Abstract

In theories with extra dimensions, the standard QCD axion has excited states with higher mass. These Kaluza-Klein (KK) axions would have a significantly shorter decay time and, when produced by the Sun, would remain gravitationally trapped in our Solar System, boosting their local decay rate. A low density detector would be able to distinguish such decays from background, by identifying the separate location of the capture of the two resulting photons. The NEWS-G collaboration uses low-pressure Spherical Proportional Counters, gas-filled metallic spheres with a high voltage electrode in their centre. This work aims to set exclusion limits on the solar KK axion model based on data from NEWS-G detectors.

Acknowledgments

What a time these past five years have been, in so many ways! As I prepare myself to submit the end product of uncountably many hours of work, I would like to mention the many people that helped write this document, in one way or another.

Gilles Gerbier, for bringing me to Canada, demonstrating leadership by example, and teaching me the ropes of what a research scientist does, yet always keeping the human component of his team in mind. And for his magical nose that can smell the results of any analysis before it's even been performed. Alexis Brossard, for surviving a PhD with me, from our humble beginnings suffering through Jackson's Classical Electrodynamics, to our hair-pulling thesis-writing struggles. Quentin Arnaud, for his technical expertise in all things analysis, and for not letting our laziness detract him from organizing social gatherings. Dan Durnford, for just being all around excellent, as a scientist, as a person, as a facial hair.

All our postdocs and research scientists, from Sabine Roth's intimidating yet kind competence, to Philippe Gros's uncanny capacity to keep the lab running despite the insanity of this past year, passing by Bei Cai, XiaoHe Zhang, Alvine Kamaha, Philippe Camus, Pierre Gorel, Guillaume Giroux, Jean François Caron, with a special mention to Ryan Martin for setting up the QUADIS library, the github repository, and our server. Julie "motormouth" McDonald and Anouchka Ronceray, who both

dealt with piles of paperwork and logistic issues so the rest of us wouldn't have to. And to the rest of the members of the NEWS-G collaboration, spread throughout the globe, especially Ali Dastgheibi-Fard, Michel Gros, and Ioannis Katsioulas, for helping me with their areas of expertise. And of course, to the incomparable Maurice Chapellier, may we all have half his "joie de vivre" when we reach his age. Finally, to the students I leave behind that are still working for their degree, Marie Vidal, Jean-Marie Coquillat, Georgios Savvidis; it appears there is a light at the end of the tunnel!

There are also multiple people that helped me stay sane throughout my years in Kingston, and helped me feel at home. People I met in the department, like Alex, Ananthan, Ryan, Ben, Philippe or Wolfgang, among many others, and the admin staff that kept my papers straight and my keys in order, most notably Loanne. But also outside of it, like the boardgame crew at Minotaur. Randy, Paul, Darius, John and Robbie, I expect you to have a spot for me when I next swing by. And the people in the Queen's Fencing Team, who were my family away from my family. Sean, who introduced me to the team (and to D&D), but who is just generally one of the most positive people I know. Luke, who is wise beyond his years, and about twice as patient, if he managed to put up with all of us for as long as he did. Brandon, for being Brandon. Hugh, who I didn't get to meet as much I would have liked. Rod, who stepped up when the team needed him most. Alice, who has been with me since before she could understand my accent, and stayed despite the accelerated hair loss. I couldn't hope for a kinder girlfriend. And her family, who took me in when COVID-19 set the world upside down.

Y, por supuesto, mi familia. Mi madre, una fuerza de la naturaleza que mantiene

todo en orden a pesar de nuestros esfuerzos en el sentido contrario. Mi padre, porque un padre es para cien hijos pero cien hijos no son para un padre. Lola, que uno de estos días se muerde la lengua y se envenena, pero que en el fondo es buena. Y Mar, que es más buena que el pan, y que esperamos todos que no la mate su trabajo. Esto hubiera sido *mucho* más difícil sin vuestro apoyo incondicional.

There is more I would like to say, and people I would like to mention, but this section is getting long already. Teachers, colleagues, friends I've made over the years... if you're reading this and thinking that you should have been included, rest assured that I thought of you too.

To all of you, thank you.

Contents

| | |
|--|-------|
| Abstract | i |
| Acknowledgments | ii |
| Contents | v |
| List of Tables | x |
| List of Figures | xi |
| List of Acronyms and Abbreviations | xxxix |
| Chapter 1: Introduction | 1 |
| Chapter 2: Solar Kaluza-Klein Axions | 4 |
| 2.1 Solar KK axions | 4 |
| 2.1.1 The PQ axion | 5 |
| 2.1.2 Extension in Kaluza-Klein theories | 6 |
| 2.1.3 Production in the Sun | 9 |
| 2.2 Evidence for solar KK axion | 15 |
| 2.2.1 Solar Corona hint | 15 |
| 2.2.2 Other | 19 |
| 2.3 Derived constraints | 21 |
| 2.3.1 Cosmological constraints | 22 |
| 2.3.2 Astrophysical constraints | 24 |
| 2.3.3 Laboratory searches | 25 |
| 2.3.4 Compiled constraints | 28 |
| Chapter 3: NEWS-G detectors | 31 |
| 3.1 Introduction to Gaseous Detectors | 31 |
| 3.2 Spherical proportional counter | 33 |
| 3.2.1 Description | 33 |

| | | |
|--|--|------------|
| 3.2.2 | Particle interaction within the SPC | 37 |
| 3.2.3 | Pulse formation | 38 |
| 3.2.4 | Event discrimination | 41 |
| 3.2.5 | Advantages and drawbacks | 44 |
| 3.3 | SEDINE | 45 |
| 3.3.1 | Setup | 45 |
| 3.3.2 | Backgrounds | 48 |
| 3.4 | NEWS-G at SNOLAB | 54 |
| Chapter 4: Data processing and Analysis | | 59 |
| 4.1 | Amplitude and Risetime estimation | 60 |
| 4.1.1 | SAMBA | 61 |
| 4.1.2 | Deconvolution | 62 |
| 4.1.3 | Modified decaytime deconvolution | 65 |
| 4.1.4 | Double Deconvolution method | 68 |
| 4.1.5 | Performance | 73 |
| 4.1.6 | Event discrimination | 78 |
| 4.2 | Analysis of Multi-Pulse events | 81 |
| 4.2.1 | Description | 81 |
| 4.2.2 | Sanity cuts | 86 |
| 4.2.3 | Performance | 88 |
| Chapter 5: Simulations | | 94 |
| 5.1 | Simple simulation | 95 |
| 5.2 | Electron drift simulation | 97 |
| 5.3 | Magboltz | 100 |
| 5.4 | COMSOL | 105 |
| 5.4.1 | Simple rod | 106 |
| 5.4.2 | Rod with umbrella | 108 |
| 5.4.3 | Note on the effect of space charge | 111 |
| 5.5 | Integration of results from Geant4 simulations | 114 |
| 5.6 | Axion simulations | 114 |
| 5.6.1 | Method | 115 |
| 5.6.2 | Photon absorption length | 118 |
| Chapter 6: Calibrating the simulations | | 122 |
| 6.1 | Electron drift time | 123 |
| 6.1.1 | Laser data | 123 |
| 6.1.2 | Track events | 124 |
| 6.1.3 | Muon veto | 127 |
| 6.1.4 | Limits of approach | 128 |

| | | |
|---|--|------------|
| 6.2 | Electron drift diffusion | 129 |
| 6.2.1 | Surface events | 130 |
| 6.3 | Avalanche statistics | 131 |
| 6.3.1 | Copper fluorescence | 133 |
| 6.3.2 | ^{37}Ar | 134 |
| 6.3.3 | Laser calibration | 135 |
| 6.4 | Ionization energy | 137 |
| 6.5 | Attachment | 138 |
| 6.5.1 | ^{37}Ar | 138 |
| 6.6 | Ion mobility | 141 |
| 6.6.1 | Pointlike events | 141 |
| Chapter 7: MPA applied to ^{55}Fe-induced Argon fluorescence | | 147 |
| 7.1 | Setup | 148 |
| 7.2 | Accounting for the high rate | 150 |
| 7.3 | Space charge effect | 153 |
| 7.4 | Expected results using a toy model | 155 |
| 7.5 | Escape peak | 159 |
| 7.5.1 | 200 mbar data | 159 |
| 7.5.2 | 110 mbar data | 161 |
| 7.6 | Double-pulse events | 162 |
| 7.6.1 | 200 mbar data | 162 |
| 7.6.2 | 110 mbar data | 168 |
| 7.7 | Performance conclusions | 172 |
| Chapter 8: Results | | 175 |
| 8.1 | SEDINE data | 175 |
| 8.1.1 | Basic axion-like cuts | 177 |
| 8.1.2 | Optimized cuts | 187 |
| 8.2 | Exclusion limit | 202 |
| 8.2.1 | Theory | 202 |
| 8.2.2 | KK Axion limit from SEDINE data | 204 |
| 8.3 | Projections for SNOGLOBE | 209 |
| 8.3.1 | Optimal running conditions | 209 |
| 8.3.2 | Axion-like backgrounds | 212 |
| 8.3.3 | Projections | 215 |
| Chapter 9: Summary and Conclusions | | 218 |
| Bibliography | | 221 |

| | |
|--|------------|
| Appendix A: The QCD Axion | 231 |
| A.1 Strong CP problem | 231 |
| A.2 Peccei-Quin mechanism and Axion solution | 232 |
| A.3 Model dependency | 234 |
| A.3.1 KSVZ vs DFSZ | 235 |
| A.4 Axion coupling to photons (and rest of SM?) | 236 |
| Appendix B: Electronics calibrations | 237 |
| B.1 Electronic components | 237 |
| B.1.1 Proportional Charge Counter | 237 |
| B.1.2 Digitizer | 238 |
| B.2 Calibration setup | 241 |
| B.2.1 Oscilloscope measurements | 243 |
| Appendix C: Theory of signal formation in an SPC | 246 |
| C.0.1 Ideal electric field | 246 |
| C.0.2 Signal induced by secondary ions | 247 |
| C.0.3 Signal induced by other charges | 251 |
| C.0.4 Ion space charge | 254 |
| Appendix D: Compute sensor charge from ion formation rate | 258 |
| Appendix E: SPC tests at Queen's University | 261 |
| E.0.1 Description | 261 |
| E.0.2 Available calibration sources | 262 |
| Appendix F: Trigger algorithm | 270 |
| F.1 Trapezoidal filter | 270 |
| F.2 Performance | 273 |
| F.2.1 Signal-to-Noise Ratio | 273 |
| F.2.2 Efficiency and False positives | 274 |
| Appendix G: Proper normalization of transforms of discrete data | 278 |
| G.1 Integrals and derivatives in discrete time | 278 |
| G.2 Fourier transforms | 281 |
| G.2.1 Convolutions and deconvolutions | 281 |
| G.2.2 Noise analysis | 284 |
| G.2.3 Discrete Fourier Transform | 285 |
| Appendix H: Parameter estimator performance | 289 |
| H.1 Amplitude resolution | 289 |

| | |
|---|------------|
| H.2 Risetime resolution | 294 |
| Appendix I: Parameter optimization for MPA | 297 |
| Appendix J: COMSOL simulations | 301 |
| J.1 Grid sensor | 301 |
| J.2 Achinos | 305 |
| Appendix K: Systematics of SEDINE axion searches | 313 |
| K.1 Control simulation | 314 |
| K.2 Drift time | 314 |
| K.3 Diffusion time | 318 |
| K.4 X-ray attenuation range | 321 |
| K.5 Energy calibration | 321 |
| K.6 Mean ionization energy | 322 |
| K.7 Attachment | 324 |
| K.8 Ion mobility | 324 |
| K.9 Combined systematics | 328 |

List of Tables

| | | |
|-----|---|-----|
| 2.1 | Coefficients for A in equation 2.10 [8]. | 12 |
| 3.1 | Summary of the main volume events background of SEDINE. | 52 |
| 3.2 | Summary of the main background of NEWS-G at SNOLAB. | 57 |
| 7.1 | Comparison of the ratio between escape peaks or fluorescence events to the total number of ^{55}Fe events, between simulations and data. . . | 173 |

List of Figures

| | | |
|-----|--|----|
| 2.1 | Effective coupling of axions to photons. Left: Decay into two photons (conversely, coalescence of two photons into an axion). Right: Primakoff conversion to a photon in the presence of a magnetic field. | 9 |
| 2.2 | First few revolutions of trapped solar KK axions. The coordinates are given in solar radii, with the shadowed region in the center outlining the solar disk [6]. | 11 |
| 2.3 | Velocity distribution for KK axions produced in the Sun via photon coalescence. a) All axions (normalised to unity); b) gravitationally trapped (normalized to f_{trap}). Note the different x scales for both plots [6]. | 13 |
| 2.4 | KK axions on Earth for $n_a = 1.0 \cdot 10^{14} \text{ m}^{-3}$ and $g_{a\gamma\gamma} = 1.0 \cdot 10^{-13} \text{ GeV}^{-1}$. Left: KK axion density. Right: KK axion decay rate [9]. | 14 |
| 2.5 | Temperature and density of the Sun surface and atmosphere [11]. The jump to $\sim 5 \cdot 10^5 \text{ keV}$ above 2000 km is in apparent defiance of the second law of thermodynamics | 15 |
| 2.6 | Energy flows to/from the corona (left), and two representative heating mechanisms [13]. | 16 |

| | |
|---|----|
| 2.7 Left: Altitude at which the Earth’s atmosphere attenuates incident radiation from the Sun by a factor e , depending on wavelength. Right: Temperature and density of Earth’s atmosphere [11]. The similarities between the temperature of the solar and Earth atmosphere suggest the behaviour of the Solar Corona could be explained by an external irradiation source [6]. | 18 |
| 2.8 X-ray photon image of the Moon as measured by ROSAT. The sunlit portion of the Moon is clearly visible, and conversely the shadowed region is blocking the diffuse X-Ray Background Radiation [18]. . . . | 20 |
| 2.9 Exclusion limits on $g_{a\gamma\gamma}$ and m_a for non-KK axion-like particles. Plot and references can be found in [19]. Red exclusion limits are derived from “experiments”, green from astronomical sources, and blue are cosmological. | 22 |
| 2.10 Exclusion limits on $g_{a\gamma\gamma}$ and n_a for the solar KK axion model. The exclusion limit from CAST [10] is the only one completely independent from any additional parameters in this model. The Solar Corona hint [6] and the exclusion limit from XMASS [10] depend on the mass distribution of trapped KK axions close to the Sun and to Earth, respectively. The constraint from the solar neutrino flux depends on the nature of the extra dimensions, but not on the fraction or distribution of trapped KK axions [6, 23]. The relationship between $g_{a\gamma\gamma}$ and n_a described in the solar KK axion model is the only one that depends on all these factors. | 29 |

| | | |
|------|---|----|
| 3.1 | Working principle (left side), with electric field lines (right side) [31]. 1: Ionization by incident particle. 2: Primary electrons drift to central electrode. 3: Charge multiplication in the avalanche region. 4: Secondary ions drift away from central electrode, inducing a current. This current is then integrated by a charge counter, forming our final signal. | 34 |
| 3.2 | Electric field lines for an SPC with a simple electrode. | 35 |
| 3.3 | Two different versions of SPC sensor with secondary umbrella electrode. | 36 |
| 3.4 | Theoretical output voltage from the preamplifier for an avalanche that creates 1000 electron-ion pairs at time $48\ \mu\text{s}$, assuming SEDINE-like conditions. | 39 |
| 3.5 | Gain distribution, i.e., number of secondary electrons created per primary electron. This is a parametrization of the actual, unknown distribution. | 41 |
| 3.6 | Risetime vs Amplitude of the signal for a 200 mbar Ar + CH ₄ (2%) + ³ He (0.4%) gas mixture. The horizontal line at $27\ \mu\text{s}$ corresponds to surface events. [30] | 42 |
| 3.7 | SEDINE. The S-shaped tube connects the sphere with the outside of its shielding [30]. | 46 |
| 3.8 | Shielding of SEDINE at LSM. The sphere is placed within the copper and lead castle [30]. | 47 |
| 3.9 | Muon flux in different laboratories, depending on their depth. LSM is at Modane, SNOLAB is at Sudbury. [30] | 48 |
| 3.10 | Decay chains from ²³⁸ U and ²³² Th. Energies listed are in MeV [43]. . | 49 |
| 3.11 | Decays of ²¹⁰ Pb and ²¹⁰ Bi [45]. | 53 |

| | |
|--|----|
| 3.12 Left: Schematic view of SNOGLOBE with shielding. Right: Installation of SNOGLOBE inside the lead shield at LSM. | 55 |
| 4.1 SAMBA processing on a pulse. In red, computed average baseline before the beginning of the pulse. The green vertical line points the position of the maximum of the pulse. The amplitude is the difference between the value at that point and the baseline; here, 12860 ADU. The two dashed black lines point the time at which the pulse reaches 10% and when it reaches 90% of the amplitude. The risetime is the difference between the two; here, $38\ \mu\text{s}$ | 61 |
| 4.2 Output of a theoretical preamplifier with infinite decaytime (dashed lines) and another with a decaytime of $50\ \mu\text{s}$ (solid lines), for events with same energy but different risetimes. For the real preamplifier, the amplitude of the pulse becomes lower as the risetime increases. | 63 |
| 4.3 Example of convolution, here with discrete time. For each of the peaks in the first plot, we add a transfer function starting at that time, multiplied by the height of the peak. | 64 |
| 4.4 MDec method applied to a high energy pulse ($\sim 10\text{ keV}$) from SEDINE's physics run. Left: Raw pulse. Middle: Deconvolved pulse. Right: Integrated deconvolved pulse. We note that the integration of the deconvolved pulse does not quite go back to a flat baseline, but instead makes a small dip before slowly increasing again. Taking the height before the dip offers a good estimator of the amplitude of the pulse. | 66 |

| | |
|---|----|
| 4.9 Energy resolution, depending on the reconstructed energy for each method. The “Ideal Energy” resolution is the normalized dispersion between the (normalized) amplitude computed by a method and the number of secondary charges in the simulated event, divided by the latter. The corresponding plots for the systematic uncertainties (Poisson from primary ionization, Polya from avalanche) are also shown for comparison. | 76 |
| 4.10 Electron arrival RMS resolution, depending on the normalized risetime for each method. The point at which the resolution shoots up for each method corresponds to the point where the risetime loses its discrimination power due to approaching its minimum value. | 78 |
| 4.11 Risetime vs Width. (a) and (b) are taken from a physics run, with (b) having a cut of 4s since the last event. (c) shows the same plot for a short neutron calibration run, justifying the lower cut on width ($80\mu\text{s}$) and risetime ($10\mu\text{s}$). | 79 |
| 4.12 Risetime vs Amplitude. Events with a Risetime $\sim 40\mu\text{s}$ are surface events. Below that are bulk events. Above, we have either tracks, multiple scatterings, or unphysical events. Below $\sim 12\mu\text{s}$ we have electronic events. | 80 |
| 4.13 Left: Simulated pulse. Right: Electron signal before convolution by the detector response. The two pulses can be difficult to distinguish when looking at the final event. | 81 |

| | |
|--|----|
| 4.14 An example of a double event found by the MPA method. Left: Raw pulse. Right: Deconvolved pulse after smoothing, with horizontal lines showing the threshold level. The event dipping under the threshold in the middle splits it into two different pulses | 82 |
| 4.15 Example of SearchHighRes finding peaks on a source spectrum[54]. The black curve is the raw data distribution, the red one is the data after TSpectrum treatment. The red markers denote the position and height of the found pulses. | 84 |
| 4.16 Multiple gaussian fit on real deconvolved pulses, based on the results from TSpectrum::SearchHighRes(). In blue, the deconvolved pulse; in red, the result from the fit. The method finds isolated peaks well, but struggles with close peaks, or when their height difference is too large. | 84 |
| 4.17 Left: MPA efficiency for KK axion events. Right: Total axion rate, and axion rate observed in our detector after applying our efficiency. | 88 |
| 4.18 Capacity of the MPA to separate two pulses (z-axis, showing the proportion of axion events that were properly reconstructed as double), based on the average width of both and the time difference between them. At smaller pulse widths, the effect of the smoothing from the method dominates. Otherwise, the MPA separates both pulses as long as the time difference is larger than their widths, as expected. The red line represents $y = x$ (no smoothing ideal scenario), the green line $y = \sqrt{x^2 + 20^2}$ (smoothing of 20 samples). | 91 |

| | |
|--|-----|
| 4.19 Capacity of the MPA to separate two pulses (z-axis, showing the proportion of axion events that were properly reconstructed as double), based on the average distance to the centre of the detector of both photon interactions, and the difference in distance between them. We require both photons to interact inside the detector. Running conditions of the detector are those from SEDINE’s physics run. From 15 cm and above, both photons have to have a radial distance difference of at least 2 cm to be resolved distinctly. | 92 |
| 5.1 The different steps of the simulation, after n_p and σ have been determined. | 97 |
| 5.2 Magboltz data for 3.1 bar of Neon with 0.7% of CH ₄ , and an Oxygen contamination of 16 ppb. | 101 |
| 5.3 Same Magboltz data, adapted to the scenario of a 30 cm radius sphere with a 3.15 mm sensor at 2520 V. The bottom-right plot shows the Townsend coefficient in green, shooting up below 1 cm, to indicate the point below which attachment can be ignored. | 103 |
| 5.4 Drift and diffusion time as a function of radius for an ideal electric field, with the same running conditions. The modelled functions are in black. | 104 |
| 5.5 Magboltz data for 1.0, 1.5, 2.00 bar of Neon with 2.0% of CH ₄ | 105 |
| 5.6 SPC setup with a simple sensor. | 107 |
| 5.7 Field lines for a detector with a simple rod and electrode configuration. Field lines away from the rod stay fairly isotropic, but become increasingly curved as we get closer to it. | 108 |

| | | |
|------|--|-----|
| 5.8 | SPC setup with a sensor with umbrella. | 109 |
| 5.9 | Comparison between COMSOL simulations of a detector with a simple rod, and one with a rod equipped with a secondary electrode close to the centre, called the “umbrella”. | 110 |
| 5.10 | Comparison of drift and diffusion times with and without space charge, for a test detector with a 15 cm radius shell and a 0.1 cm radius electrode with an applied voltage of 1200 V, with 0.2 <i>bar</i> of Argon ($\mu_0 = 1.5 \text{ cm}^2/\text{V} \cdot \text{s} \cdot \text{bar}$), and an avalanche gain of 5000. The ion formation rate is given by 500 Hz of 10 keV events. | 113 |
| 5.11 | Geant4 simulation of a 1 MeV electron inside the detector. Left: Visualization of its energy deposition inside the detector, with the star marking the starting position; while the track of the electron was relatively straight, the projection of its path into cylindrical coordinates introduces an artificial “bend” in its trajectory. Right: Pulse generated from this event. | 115 |
| 5.12 | Left: decay spectrum of solar KK axions on Earth, for an assumed axion local density of $1.0 \cdot 10^{14} \text{ m}^3$ [9]. Right: number density of solar KK axions depending on the distance from the Sun; the red fitted r^{-4} curve allows extrapolation to the distance from Earth to the Sun, $215 R_\odot$, for a local density of $4.075 \cdot 10^{13} \text{ m}^3$ [10]. | 116 |

| | |
|--|-----|
| 5.13 Left: X-ray mass attenuation coefficient μ/ρ and mass absorption coefficient μ_{en}/ρ in Neon [63]. For photons under 10 keV, the difference between μ/ρ and μ_{en}/ρ remains under 5%; the difference jumps to $\sim 25\%$ at 20 keV. Right: Absorption length of photons under SEDINE running conditions; the radius of the detector is shown with the horizontal red line. | 119 |
| 5.14 Left: In red, the probability of both photons being absorbed within the detector for a given attenuation length; in green, $0.54(1-e^{-1/x^2})$, which fits the containment probability increasingly well after $\lambda/R_{shell} > 1$. Right: separability of both photons, making the simplifying assumption that two photons are separable if their radial position differs by more than a fixed δ ; plots for various values of δ are shown. | 120 |
| 6.1 Juxtaposed SPC (blue) and laser trigger (red) channels. The delay between the 50% point of the rise in the SPC channel signal with respect to the beginning of the laser trigger signal gives the drift time of electrons from the surface to the central electrode. | 124 |
| 6.2 Distribution of widths for high width events in the background run, after cutting out alpha events. The amplitude cuts on the right plot are, from top to bottom, > 150 kADUs, > 200 kADUs, > 250 kADUs, > 300 kADUs. We can see that the maximum width depends on energy, but seems to be reached somewhere in the 200 – 250 kADUs range and above. | 125 |

| | | |
|-----|---|-----|
| 6.3 | Distribution of widths for high width events in the simulated track dataset. The amplitude cuts on the right plot are, from top to bottom, > 150 kADUs, > 200 kADUs, > 250 kADUs, > 300 kADUs. We can see that the width at the maximum of the distribution depends on energy, but stabilizes somewhere in the $200 - 250$ kADUs range and above. | 126 |
| 6.4 | Distribution of correction factors between data and simulations when the values of the measures for the maximum width of either is left to vary uniformly along their whole range. | 128 |
| 6.5 | $t = t_{max}(\frac{r}{r_{max}})^\alpha$, for different values of α . For $t_{max} = 470 \mu\text{s}$, $r_{max} = 30 \text{ cm}$, $\alpha = 2.5$, this reproduces the drift time vs radius curve obtained through Magboltz simulations. | 129 |
| 6.6 | Risetime of events in the $2 - 10 \text{ keV}$ energy range for the SEDINE physics data (blue), compared with simulations of electrons in same energy range generated at the surface (green). Simulations normalized to have a similar height as the data. | 130 |
| 6.7 | 8.05 keV peak from copper fluorescence during the SEDINE physics run. The fit results give a peak at 16820 ADU, with an uncertainty of 2.9%. | 134 |
| 6.8 | Amplitude of Ar^{37} events in a calibration run with the Queen's S30 detector. Both the 2.82 keV and 270 eV peaks are clearly apparent. . | 135 |
| 6.9 | Amplitude vs Risetime for the Ar^{37} calibration run. A linear fit is performed on the average amplitude in each slice of risetime. The value of the slope is $3.3 \pm 0.1 \cdot 10^2$ | 139 |

| | |
|--|-----|
| 6.10 Top: Simulated ^{37}Ar runs with different oxygen contaminations. Bottom: Fitted slopes from simulations [35]; the simulations were not calibrated for amplitude, so their units are arbitrary (au). The value of the slope from the calibration run is $3.3 \pm 0.1 \cdot 10^2$, corresponding to an oxygen contamination of 16 ppb. | 140 |
| 6.11 Pulse deconvolved by the response function of the detector, then integrated, for different values of the ion mobility in the processing. Low mobilities tend to “overdeconvolve” the pulse, producing a dip after it. Conversely, high mobilities “underdeconvolve” the pulse, leading to an increasing slope right after it. | 142 |
| 6.12 Distribution of computed ion mobilities for the two studied runs, with results from the fit. The final ion mobilities computed with this method are $6.18 \text{ cm}^2/\text{V/s} \pm 0.99\%$ and $7.45 \text{ cm}^2/\text{V/s} \pm 1.95\%$, respectively. . . | 144 |
| 7.1 Top left: S130 detector used for the calibration. Bottom left: ^{55}Fe source (without aluminium foil). Right: Close-up on the calibration source window, with and without the ^{55}Fe source installed in it. . . | 149 |
| 7.2 Event example from S130 with 200 mbar. Left: Raw pulse. Right: Processed. The laser-induced event is the smaller one at $\sim 4300 \mu\text{s}$. The pulse at $\sim 1800 \mu\text{s}$ is likely an ^{55}Fe event. The other three are likely muons. | 150 |

| | | |
|-----|--|-----|
| 7.3 | Maximum drift time calibration with laser events, with gaussian fit. The time of ionization is obtained with a fiber splitter that sends parts of the laser signal to a photodetector. The difference between that time and the average arrival time of the primary electrons gives the (maximum) drift time, around $240\ \mu\text{s}$ at 200 mbar, 900 V for this detector. | 152 |
| 7.4 | Black: Total energy distribution of all events after basic cuts. Red: Contribution from ^{55}Fe events. Green: Contribution from Muon events; they drop down to 0 around 180 keV. | 153 |
| 7.5 | Left: Electron drift simulation, total drift time depending on initial radial position. Simulations are run both with no space charge, and with the space charge induced by an ion current of 19.7 pA. Right: Fitted distribution of drift times in the data (blue/red, as seen in Fig. 7.3) and normalized distribution of drift times in simulations with adjusted ion current (green) | 154 |
| 7.6 | Attenuation length of photons in 200 mbar of Argon (NIST datapoints in green, log-log linear interpolations in black). In red, the radius of the detector. The drop at 3.2 keV corresponds to the K-shell of Argon. One of its consequences is that a 2.93 keV and a 5.9 keV photon have comparable attenuation lengths, despite the factor 2 difference in energy. | 156 |

| | | |
|------|---|-----|
| 7.7 | Top left: Escape peak event probability depending on gas pressure. Top right: double pulse event probability depending on gas pressure, for different minimum separability distances. Bottom: relative rate of escape peak events and double pulse events; the rate (in arbitrary units) of muons that fall in the energy range of ^{55}Fe events, based on the recorded data, is also added. | 157 |
| 7.8 | Ratio of contained Argon fluorescence events that can be separated with the processing, at 200 mbar with 1150 V. The separability dis- tance is taken to be ~ 8.5 cm. | 158 |
| 7.9 | Top: Risetime vs Amplitude after basic cuts, keeping only single- pulse events and removing laser events; ^{55}Fe are at 1300 ADU, with escape peaks around 600 ADU; the horizontal accumulation around $35\ \mu\text{s}$ across amplitudes comes from Muon events. Bottom: Amplitude distribution for events with risetime between $27\ \mu\text{s}$ and $40\ \mu\text{s}$; the sum of a linear function (shown separately in blue) and two gaussians is fitted to the distribution. | 160 |
| 7.10 | Amplitude distribution for events with risetime between $21\ \mu\text{s}$ and $35\ \mu\text{s}$; the sum of a linear function (shown separately in blue) and two gaussians is fitted to the distribution. | 162 |
| 7.11 | Effect of number of pulses in the event distribution (200 mbar data). . | 163 |

| | |
|--|-----|
| 7.12 3-parameter fit of equal-pulses event distribution. Top: whole range of the fit. Bottom: Zoom in the region under 20 keV. Red points: Data (200 mbar). Blue curve: Fit result with statistical uncertainty. Dotted lines: Contribution from coincident events (orange), Argon fluorescence (red), and False Positives (pink; here set to zero by fit). The overpopulation in the 2 – 4.5 keV range are improperly reconstructed events. | 165 |
| 7.13 Deconvolved events from the 200 mbar run. Left: Argon fluorescence event, with pulse threshold in red. Right: 3 keV event improperly reconstructed as being an equal-pulses event, due to it “jumping” above the threshold multiple times. At energies under 5 keV, the ratio of the second kind of event increases considerably. | 167 |
| 7.14 Effect of number of pulses in the event distribution (110 mbar data). . | 169 |
| 7.15 3-parameter fit of equal-pulses event distribution. Top: whole range of the fit. Bottom: Zoom in the region under 20 keV. Red points: Data (110 mbar). Blue curve: Fit result with statistical uncertainty. Dotted lines: Contribution from coincident events (orange), Argon fluorescence events (red), and False Positives (pink). The overpopulation in the 1 – 3.5 keV range are improperly reconstructed events. | 170 |
| 7.16 110 mbar data. Black: Total energy distribution of all events after basic cuts. Red: Presumed contribution from ^{55}Fe events. Green: Presumed contribution from Muon events. From their distribution, it is clear that the events naively attributed to ^{55}Fe are overestimated | 172 |

| | | |
|-----|--|-----|
| 8.1 | Risetime vs Width of raw pulses, excluding only very low amplitude events. Comparing this plot to results from calibrations, we choose to cut to only keep the events with a risetime above $10\ \mu\text{s}$ and width above $75\ \mu\text{s}$ | 176 |
| 8.2 | Distribution of time since previous event. Above 2 s, the exponential decrease with time is the behaviour expected from events happening randomly with respect with each other. The fast increase under 2 s reveals that there are periods with much higher event rates (or correlated events), that we should remove from the analysis. | 177 |
| 8.3 | Examples of events rejected by basic cuts. Top left: “electronic event” (zoomed-in) with a sharp rise characteristic of absence of ion-current structure. Top right: high-rate period after a large event. Bottom left: “pulse-like” baseline noise. Bottom right: transient noise. | 178 |
| 8.4 | Processed risetime vs amplitude of SEDINE physics data. | 179 |
| 8.5 | Distribution of amplitude of events for different cuts. Roughly, dark blue are physical events with two pulses; green are events with two pointlike pulses; light blue are events with two “simultaneous” pointlike pulses; red are events with two such pulses that are relatively close in energy. For comparison, we also show the results for a harsher cut on the allowed difference between both pulses (black). | 181 |

| | | |
|------|--|-----|
| 8.6 | Distribution of amplitude of simulated axion events for different cuts. Roughly, dark blue are physical events with two pulses; green are events with two pointlike pulses; light blue are events with two “simultaneous” pointlike pulses; red are events with two such pulses that are relatively close in energy. For comparison, we also show the results for a harsher cut on the allowed difference between both pulses (black). | 183 |
| 8.7 | Comparison between observed events that pass all axion-like cuts (blue), and normalized expectations from simulations of solar KK axions (red). Top: Energy distribution. Bottom: Time difference between first and second pulse. | 185 |
| 8.8 | Comparison between the first and the second pulse for events that pass all axion-like cuts. Top: Amplitudes. Bottom: Risetimes | 186 |
| 8.9 | Contribution of different radioactive contaminations to total axion-like backgrounds, using basic axion-like cuts. | 188 |
| 8.10 | Decay chain of ^{210}Pb , with ^{210}Bi [45]. | 188 |
| 8.11 | Example of ^{210}Pb event simulated by Geant4 that passes basic axion-like cuts. ^{210}Pb decays into an excited state of ^{210}Bi , which de-excites through emission of electrons and a photon. The electrons are either absorbed in the surface or escape out of the detector, while the photon is captured in the bulk of the detector, leading to two separate, simultaneous energy depositions. | 189 |
| 8.12 | Example of ^{210}Bi event simulated by Geant4 that passes basic axion-like cuts. A Bremsstrahlung photon interacting twice in the detector. | 191 |

| | |
|---|-----|
| 8.13 Example of ^{210}Bi event simulated by Geant4 that passes basic axion-like cuts. The high-energy electron from the decay of ^{210}Bi interacts in the gas for a short distance before escaping, while a Bremsstrahlung photon interacts deeper in the detector. | 192 |
| 8.14 Risetime of first pulse vs. Risetime of second pulse, after basic axion-like cuts. In both plots, the colour distribution is from axion simulations. The black line shows the requirement that the risetime of the first pulse be smaller than that of the second. Left: In red, distribution from simulations of SEDINE background. Right: In red, distribution from the SEDINE physics data. There is a larger proportion of background events where the second pulse has a risetime consistent with surface events. | 194 |
| 8.15 Distribution of Risetime of the second pulse for axion events (blue) and background events (red), based on simulations. | 195 |
| 8.16 Asymmetry between pulse amplitudes vs. Reconstructed event energy, after basic axion-like cuts. In both plots, the colour distribution is from axion simulations. The black line shows the asymmetry cut from the basic axion-like cuts. Left: In red, distribution from simulations of SEDINE background. Right: In red, distribution from the SEDINE physics data. | 195 |

| | |
|--|-----|
| 8.17 Asymmetry between pulse amplitudes vs. Reconstructed event energy of simulated axion events, after basic axion-like cuts, for high (blue), average (green) and low (red) values of the attachment. The solid black line shows the asymmetry cut from the basic axion-like cuts, the dashed line shows the improved asymmetry cut. | 197 |
| 8.18 Contribution of different radioactive contaminations to total axion-like backgrounds, using axion-like cuts with improved asymmetry and risetime cuts. The contributions have been smoothed out for better visualization. | 198 |
| 8.19 Top: Event rate vs. energy for axion and background events (black). Bottom: limit-like parameter when keeping all energies such that their SNR is above a given threshold. The green curves are for the expected value of electron attachment, the blue curves for its maximum allowed value, and the green for its minimum. | 199 |
| 8.20 Comparison of axion-like events in data (blue), and expectations from simulations of radioactive background (black), after improved axion-like cuts. Only 1 event in the data passes the improved cuts. | 201 |

| | | |
|------|---|-----|
| 8.21 | Exclusion limit for solar KK axions derived from this work (solid red line), with ranges due to systematics (red shaded areas). For comparison, we show the ideal exclusion limit in the absence of background (dashed red line), the previous limits on solar KK axions from CAST (orange line [26]), and from the XMASS collaboration (green line [10]). The preferred parameter space for the solar KK axion model is shown as the intersection between the solid black line (Solar KK axion model) and the dashed black line (Solar Corona hint) [6]. | 206 |
| 8.22 | Detector efficiency and detected events in the SEDINE detector, for a total exposure of $4.3 \text{ day} \cdot \text{m}^3$. Energies are corrected for attachment. The resolution of the energy reconstruction is $\sim 20\%$ at these energies. | 207 |
| 8.23 | Attenuation and absorption coefficients for different gases at 1 bar. The shaded area approximately covers the energies and distances involved in solar KK axion searches with SNOGLOBE. Helium is too transparent, and methane must be diluted to be used in SNOLAB, so the only available gases of interest are Neon and Argon. | 210 |
| 8.24 | Predicted solar KK axion “maximum” efficiency for different voltage and pressure configurations. Top: Neon. Bottom: Argon. Electron drift parameters stay roughly constant when E/P is constant, but photon attenuation length is inversely proportional to P , hence the diagonal feature in both plots. | 211 |
| 8.25 | Axion (red) and radioactive background (black) rate after cuts for SNOGLOBE with 600 mbar of neon and an applied voltage of 2000 V. | 214 |

| | |
|---|-----|
| 8.26 Projected limits for SNOGLOBE with 600 mbar of neon and an applied voltage of 2000 V. Both a plausible limit (dashed dark purple line) with a 30 day run and the expected background-to-axion rate of 20%, and an ideal limit (dashed light purple line) with a 180 day run and a background-to-axion rate of < 1%, are shown. | 215 |
| B.1 Functional diagram of a proportional counter, as provided by the Canberra 2006 documentation [49]. | 238 |
| B.2 Electronic diagram for the Calibox input. Top: Detail of the full diagram for one $+/ -$ pair of inputs. Bottom: Simplified diagram for a single input (although missing the 49.9Ω input resistance connected to the ground). Please note the left/right inversion between both diagrams. | 239 |
| B.3 Electronic setup to calibrate $K_{e^-/ADU}$ | 241 |
| B.4 Response function of the CR-Z-110 under low resistance load (50Ω). Red: old model (Rev2). Blue: new model (Rev2.1). The response of the new model is a decaying exponential with a time constant $\tau = 141\mu s$, while the old model had a distinct undershoot behaviour. . . . | 245 |
| C.1 Theoretical current induced on the electrode by an avalanche that creates 1000 electron-ion pairs at time $48\mu s$ | 248 |
| C.2 Theoretical output voltage from the preamplifier for an avalanche that creates 1000 electron-ion pairs at time $48\mu s$ | 250 |

| | | |
|-----|---|-----|
| C.3 | Composition of an SPC signal, compared to an ideal pure-ion-current signal, for an event with a single primary electron. SEDINE conditions are assumed. The signal from secondary electrons is visible, but the shape of the total induced pulse is effectively indistinguishable from the ideal one. The signal from primary electrons is too small to be observable, even for a relatively low value of the avalanche gain of 1000. | 252 |
| C.4 | Composition of an SPC signal, compared to an ideal instant-charge signal, for an event with a single primary electron, and no avalanche. SEDINE conditions are assumed. | 253 |
| E.1 | Some of the SPC detectors available at Queen's University. Top left: 30 cm-diameter SPC. Top right: 15 cm-diameter SPC. Bottom left: 130 cm-diameter SPC at Queen's, originally a radio-frequency cavity from the Large Electron-Positron collider at CERN [84]. Bottom right: 30 cm-diameter glass SPC, for outreach. | 263 |
| E.2 | Left: Experimental setup of a laser calibration at Queen's; the photodetector is used both as an external trigger and to monitor variations in the beam pulse intensity. Right: Example of energy spectrum from one such calibration with a relatively high number of primary electrons, together with the fit results (red: total, orange/green/black: contribution from different number of extracted primary electrons). Both of these diagrams are taken from [38]. | 267 |
| E.3 | Top: In the foreground, Backing detectors at TUNL; in the background, 15 cm SPC used for quenching measurements. Bottom: Schematic drawing of setup. | 268 |

| | | |
|-----|--|-----|
| F.1 | In blue, raw pulse (after removing average baseline, for comparison). In red, pulse after trapezoidal filter, scaled up by a factor of 4 to match the noise level of the raw pulse. The combination of a running average and a derivative allows the reduction of high and low frequency noise respectively, boosting the Signal-to-Noise Ratio. | 272 |
| F.2 | Signal-to-Noise ratio for the raw pulse, smoothed pulse, and the trape- zoidal filter of the pulse. The X axis is the gain of the primary electron, with respect to the average gain. The running conditions and noise traces were taken from the physics run. | 274 |
| F.3 | False positive rate depending on event detection efficiency, for thresh- olds applied to the raw pulse, smoothed pulse, and the trapezoidal filter of the pulse. Above the black horizontal line, the false positive rate was extrapolated from the proportion of empty traces that triggered. The small bump at low efficiencies for the trapezoidal filter is likely caused by very low energy events sneaking into the noise traces used to compute the false positive rate. | 276 |
| H.1 | Amplitude resolution of all three methods, depending on the secondary charges of the event (in units of the energy most likely to produce that number of secondary charges). For comparison, the relative resolution of the number of primary electrons produced by any given energy, and the relative resolution of the number of secondary charges produced by any number of primary electrons (in units of the energy most likely to produce that number of primary electrons). | 291 |

| | | |
|-----|--|-----|
| H.2 | Ideal secondary charge resolution, depending on the reconstructed energy from the amplitude for each method. The corresponding plots for the systematic uncertainties produced by primary electron (Poisson) and secondary charge (Polya) creation are shown again for comparison. | 293 |
| H.3 | Normalized MDec risetime vs the electron arrival time RMS. Note how the minimum risetime is $6\ \mu s$, despite the RMS going all the way down to 0. | 295 |
| H.4 | Electron arrival RMS resolution, depending on the normalized risetime for each method. The point at which the resolution shoots up for each method corresponds to the point where the risetime loses its discrimination power due to approaching its minimum value. | 296 |
| I.1 | Proportion of false positives out of uniformly distributed pointlike events, for different values of the processing parameters. Some mild smoothing was applied, to increase the readability of the plot. | 298 |
| I.2 | Left: Axion rate, depending on processing parameters. Right: Proportion of pointlike events that produce a False Positive, depending on processing parameters. | 298 |
| I.3 | Optimization metric, depending on processing parameters. The optimal parameters are taken to be a smoothing strength of 30 samples, and a threshold of 1.3. | 299 |
| I.4 | Left: Comparison of Axion rate, for default and optimized parameters. Right: Comparison of False Positive proportion of pointlike events, for default and optimized parameters. | 300 |
| J.1 | | 302 |

| | |
|--|-----|
| J.2 Top: field lines in the detector. The field lines are red when ending on the central electrode, dark blue when ending on the grid. Bottom: field strength vs radius. The red function is the ideal field created by a sensor with the radius of the grid and the same applied voltage. The blue lines are (starting from the bottom-most) the simulated field strengths at 0, 22.5, 45 and 90 degrees below the horizontal; they differ by less than 20%, at the edge of the detector. | 304 |
| J.3 | 306 |
| J.4 Top: 2D projections of isopotentials and field lines. Bottom: 3D field lines, blue for lines ending on the rod-side electrodes, red for lines ending on the other ones. The plots on the left have 1000 V applied on all electrodes, the ones on the right have 1080 V on rod-side electrodes and 1000 V on others. Despite the large observed difference in the avalanche for either configuration, the field lines are barely affected . | 307 |
| J.5 Left: field lines close to the electrode furthest away from the rod. Right: in red, ideal field for a single electrode inside the detector; in black, gray, cyan, blue and green, the field strength away from the electrode, at +90, +45, +0, -45 and -90 degrees respectively from the line away from the centre of the sensor. The +90 degree line points towards another electrode, the -90 degree line points in between two sensors. | 309 |

| | |
|---|-----|
| J.6 Left: field lines in the drift region. Right: in red, ideal field for a single electrode with radius that of the achinos sensor, and the same ideal field multiplied by 0.17; in cyan, blue, gray and black, the field strength away from the sensor, at 0, -22.5, -45 and -90 degrees respectively from the equator (all pointing away from the rod) | 311 |
| K.1 Dispersion on the detector efficiency from 200 simulations of 10k events each, with all parameters fixed, and assuming a local density of KK axions of $4.07 \cdot 10^{13} \text{ m}^{-3}$. For reference, the red line is the mean for this simulation. The relative dispersion of the efficiency is 2% | 315 |
| K.2 Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the drift time is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 9.5% | 316 |
| K.3 Top: Toy models for the drift time vs radius relationship, with fixed maximum drift; $\alpha = 2.5$ corresponds to the fit of the results from simulations. Bottom: Simulated distribution of separable events in the detector for different values of α . The relative rates (in order of increasing α) are 76%, 89%, 100%, 98%, 83% | 317 |
| K.4 Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the diffusion time is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 12% | 319 |

| | |
|--|-----|
| K.5 Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the photon attenuation range is left to vary according to NIST data, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 3%. | 320 |
| K.6 Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the mean avalanche gain is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 3%. | 322 |
| K.7 Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the mean ionization energy is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 2%. | 323 |
| K.8 Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the attachment parameter is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 5%. | 325 |

| | |
|---|-----|
| K.9 Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the ion mobility is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 2%. | 326 |
| K.10 Variation in efficiency when $2 \cdot 10^6$ events are simulated with an ion mobility drawn from a gaussian with a relative resolution of 18.9%. The y-axis is zoomed into the range of possible values to better appreciate the effect. | 327 |
| K.11 Dispersion on the detector efficiency to KK axions, from 1000 simulations of 10k events each, where all simulation parameters are left to vary according to our calibration uncertainties. The purple (resp. blue) does (resp. does not) take into account the effect from our lack of knowledge on the drift time in the bulk of the detector (cf. Sec. K.2). For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency (including the drift model uncertainty) is 21%. | 329 |

List of Acronyms and Abbreviations

| | |
|---------------|--|
| ADU | Analog-to-Digital Unit (DAQ unit) |
| AU | Arbitrary Unit |
| CMB | Cosmic Microwave Background |
| DAQ | Data AcQuisition system |
| DD | Double Deconvolution (pulse processing algorithm) |
| DM | Dark Matter |
| KK | Kaluza-Klein (a theory of quantum gravity) |
| MDec | Modified Decaytime (pulse processing algorithm) |
| MPA | Multi-Pulse Analysis (pulse processing algorithm) |
| NEWS-G | Novel Experiments With Spheres - Gas detectors (collaboration) |
| PQ | Peccei-Quinn (possible solution to strong CP problem) |
| preamp | Preamplifier (charge integrator) |
| QCD | Quantum Chromo-Dynamics |

SPC Spherical Proportional Counter

WIMP Weakly Interacting Massive Particle (Dark Matter candidate)

Chapter 1

Introduction

The axion is a theoretical boson first postulated in 1977 to solve the strong CP problem [1]. This elusive particle has evaded detection until the present day, but various cosmological and astrophysical constraints have set limits to its mass between $\sim 1 \mu\text{eV}$ and $\sim 1 \text{eV}$ [2], with preference towards the lower values. However, in quantum gravity theories with additional compactified dimensions, the axion gains a tower of excitations of much larger mass, up to $\sim 10 \text{keV}$ [3–5]. One of the consequences of the existence of these “KK modes” of the axion is their accumulation in the Solar System, due to their creation at lower speeds in the Sun. With the right number and size of the additional dimensions, and the right value of the axion-photon coupling, the solar KK axions could explain an old puzzle, the solar corona problem: the atmosphere of the Sun is hotter than its surface, defying thermodynamics. Decays from the accumulated solar KK axions could provide an external radiation for the Sun, solving this apparent contradiction [6].

This thesis aims to set limits on the solar KK axion model by looking for axion decays on Earth with a Spherical Proportional Counter (SPC) [7]. SPCs are a novel kind of detector developed by the NEWS-G collaboration. They consist of a grounded

metallic sphere holding some target gas, with a central electrode kept at a high voltage that collects the ionization signal produced by any incident particle interacting with the gas. Their strong advantage for testing the solar KK axion model is that axions decay into two photons travelling in opposite directions, which interact in distinct locations in SPCs due to low density of gas targets, compared to that of most other rare event detectors (liquids, crystals). This produces two signals shortly after each other, with the signal from the photon interacting closer to the electrode arriving before the one farther away. By looking exclusively for coincident events in the SPC, a very strong background discrimination can be achieved, allowing much stronger sensitivities to be reached than might be expected.

In Chapter 1, the theory behind the solar KK axion model will be covered, together with its consequence of a density of heavy axions in the Solar System and the potential hints for such a model. Already existing constraints on axions and axion-like particles will be revisited in the context of this model. Chapter 2 will present the functioning principle, advantages and drawbacks of SPCs. The setup for two existing low-background detectors, SEDINE and SNOGLOBE, will also be described, together with their radioactive backgrounds. Chapters 3, 4 and 5 heavily inform each other, and represent the bulk of the work performed for this thesis. Chapter 3 will get into the detail and performance of the algorithms and pulse processing tools that were developed to treat the data from SPCs. In particular, the method to distinguish between double-pulse events (characteristic of axion decays) from other events, and extract information on both pulses, will be related to its physical implications for axion detection. Chapter 4 will cover the full suite of tools used to simulate the signal from an SPC: energy deposition, electric field, electron drift, and pulse formation.

Chapter 5 provides the calibrations that were available for SEDINE. The processing tools described in Chapter 3 were applied simultaneously on this data and on the simulations described in Chapter 4 to refine the simulation parameters until there was agreement between the two. In Chapter 6, the search for double-pulse events was tested on a different type of source than axions: ^{55}Fe -induced argon fluorescence. This generates two photons of 2.9 keV each, for a very close approximation of the decay of a 5.8 keV axion. Finally, Chapter 7 gets into the physics data taken with SEDINE, the optimization of selection cuts to reject background, and the derivation of constraints on the solar KK axion model based on this data. We close the chapter with a discussion on the expected performance of the future SNOGLOBE detector, based on possible running conditions and predicted radioactive background.

Chapter 2

Solar Kaluza-Klein Axions

In models with additional compactified dimensions, the axion gains a tower of excitations, called Kaluza-Klein (KK) modes, regularly spaced in mass; depending on the nature of these additional dimensions, the mass of the excitations could be much higher than that of the axion, of the order of $\sim 1 - 10\text{keV}$, greatly boosting the rate of their decay into two photons. Furthermore, given their high mass, a proportion of heavy axions produced by the Sun would remain trapped in its gravity well, and accumulate until reaching densities high enough to observe their decays.

This chapter will first cover the theory behind KK axions, then move onto phenomena that hint at their existence, and finally existing constraints on their nature.

2.1 Solar KK axions

In this section, I will briefly cover the theoretical background behind the axion, and an extension to its “standard” description into Kaluza-Klein theories, based on [8]; a very concise description of the solar KK axion model can also be found in [9].

2.1.1 The PQ axion

The Peccei-Quinn (or Quantum Chromo-Dynamics) axion was first proposed to solve the Strong CP problem (see App. A for some extra detail). Before introducing its behaviour in Kaluza-Klein theories, we will quickly go over some of the properties of the “standard” axion. First, looking only at the coupling of axions to photons, the part of interest of the effective Lagrangian is:

$$\mathcal{L}_a \supset \frac{1}{2}(\partial_\mu a)^2 - \frac{1}{2}m_{PQ}^2 a^2 + \frac{g_{a\gamma\gamma}}{4}a F \tilde{F} \quad (2.1)$$

where a is the PQ axion field, F and \tilde{F} are the electromagnetic field-strength tensor and its dual, and

$$g_{a\gamma\gamma} = \frac{\xi \alpha_{em}}{\pi} \frac{1}{\nu_{PQ}} \quad (2.2)$$

is the effective axion-photon coupling, where ν_{PQ} is the breaking scale of the $U(1)_{PQ}$ symmetry (also noted f_{PQ} in the literature), and $\xi = O(1)$ a multiplicative factor that depends on the specific axion model. The mass of the axion, m_{PQ} , is also related to the breaking scale:

$$m_{PQ} \sim \frac{m_\pi^2}{\nu_{PQ}} \quad (2.3)$$

where $m_\pi \simeq 135$ MeV is the pion mass. Astrophysical and cosmological constraints (see Sec. 2.3) set bounds for the axion at 10^{-5} eV $\lesssim m_{PQ} \lesssim 10^{-2}$ eV. These values

can be used to compute the lifetime of the QCD axion:

$$\tau(a \rightarrow \gamma\gamma) = \frac{64\pi}{g_{a\gamma\gamma}^2 m_{PQ}^3} \sim 10^{48} \text{ days} \cdot \left(\frac{10^{-15} \text{ GeV}^{-1}}{g_{a\gamma\gamma}}\right)^2 \left(\frac{10^{-5} \text{ eV}}{m_{PQ}}\right)^3 \quad (2.4)$$

where $g_{a\gamma\gamma} = 10^{-15} \text{ GeV}^{-1}$, corresponding to $m_{PQ} = 10^{-5} \text{ eV}$, were used as the reference points, leading to a decay time much larger than the age of the universe $T_U \sim 5 \cdot 10^{12} \text{ days}$. This remains the case even for larger axion masses; even at the already excluded mass value of $m_{PQ} = 10^{-1} \text{ eV}$, the axion lifetime is still $\tau_{a \rightarrow \gamma\gamma} \sim 10^{27} \text{ days}$. Axion decay would remain undetectable, short of some prodigious source of the particle¹.

2.1.2 Extension in Kaluza-Klein theories

The properties of the PQ axion change in higher-dimensional theories of low-scale quantum gravity. The gauge hierarchy problem between the gravitational scale and the Planck scale could be explained if n extra compact dimensions through which gravity, but not the particles from the Standard Model, can propagate. In that case, the Planck scale M_P is just an effective coupling, related to the scale of ($4 = n$) dimensional gravity by:

$$M_P^2 = 4\pi R^n M_F^{2+n}$$

where R is the compactification radius of the extra dimensions. Singlet fields, such as the axion, could also propagate through some or all of those additional dimensions. We can introduce one such axion field $a(x^\mu, y)$, which feels the presence of $\delta \leq n$

¹For comparison, if all Dark Matter was composed of axions, and even in the very favourable case of $m_{PQ} = 10^{-1} \text{ eV}$, a detector of 1 km^3 would take $\sim 2e5 \text{ years}$ to see a single decay. Current experiments looking for axions as DM use other channels to detect them.

dimensions, denoted by $y = (y_1, y_2, \dots, y_\delta)$. This field is then compactified on a Z_2 orbifold with action $y \rightarrow -y$, giving rise to the KK decomposition:

$$a(x^\mu, y) = \sum_{n=0}^{\infty} a_n(x^\mu) \cos\left(\frac{xy}{R}\right) \quad (2.5)$$

where $n = (n_1, n_2, \dots, n_\delta)$ is a δ -dimensional vector that labels the individual KK excitations, and $\sum_{n=0}^{\infty} = \sum_{n_1=0}^{\infty} \sum_{n_2=0}^{\infty} \dots \sum_{n_\delta=0}^{\infty}$. Under this decomposition, the effective Lagrangian becomes:

$$\mathcal{L}_{eff} = \frac{1}{2} \sum_{n=0}^{\infty} (\partial_\mu a_n)^2 - \frac{1}{2} m_{PQ}^2 a_0^2 - \frac{1}{2} \sum_{n=0}^{\infty} \frac{n^2}{R^2} a_n^2 + \frac{\xi \alpha_{em}}{\pi} \sum_{n=0}^{\infty} \frac{r_n a_n}{\nu_{PQ}} F \tilde{F} \quad (2.6)$$

where ν_{PQ} denotes the original higher-dimensional PQ-breaking scale, $\nu_{PQ} = (\frac{M_F}{M_P})^{\delta/\mu} \nu_{PQ}$ ², $r_0 = 1$, and $r_{n \neq 0} = \sqrt{2}$. From this effective Lagrangian, we read the effective coupling of the KK axions to photons:

$$g_{a_n \gamma \gamma} = \frac{r_n \xi \alpha_{em}}{\pi} \frac{1}{\nu_{PQ}} \sim g_{a \gamma \gamma} \quad (2.7)$$

Since the correction between the PQ-photon coupling and the KK excitation-photon coupling is constant and of order unity, we will just take them to be equal for the sake of simplicity. The mass of the excitations, however, are considerably different:

$$m_{a_n} \sim \frac{n}{R} \gg m_{PQ} \quad (2.8)$$

²This correction could explain the largeness of the observed symmetry-breaking scale [3], motivating the study of the non-trivial phenomenology of KK axions in contexts other than the one covered in this work [4, 5].

with $n = |n| = \sqrt{n_1^2 + \dots + n_\delta^2} > 0$. A tower of excitations, all evenly spaced in mass by a factor of $1/R$, appear. More importantly, the masses of the excitations are much larger than the mass of the basic axion: while $m_{PQ} \lesssim 0.01$ eV, for $\delta = 2$ and $M_F \sim 100$ TeV, one obtains $1/R \sim 1$ eV³. An even higher value of $1/R \sim 10$ eV can be found for $\delta = 3$ and $M_F \sim 1$ TeV⁴.

Then, going back to the original question of axion decays into two photons, we get the following correction to the half-life of the KK axions:

$$\tau_{a_n \rightarrow \gamma\gamma} = \left(\frac{m_{PQ}}{m_{a_n}}\right)^3 \tau_{a_0 \rightarrow \gamma\gamma} \quad (2.9)$$

A single KK-axion with $m_{a_n} = 10$ keV, and for a coupling to photons of $g_{a\gamma\gamma} = 10^{-11}$ GeV⁻¹ (corresponding to an axion mass of $m_{PQ} = 10^{-1}$ eV) has a lifetime of $\tau \sim 10^{12}$ days, 15 orders of magnitude smaller than a PQ axion, and just under the age of the Universe. This opens the possibility of looking for the signature of axions through their decay into two photons⁵.

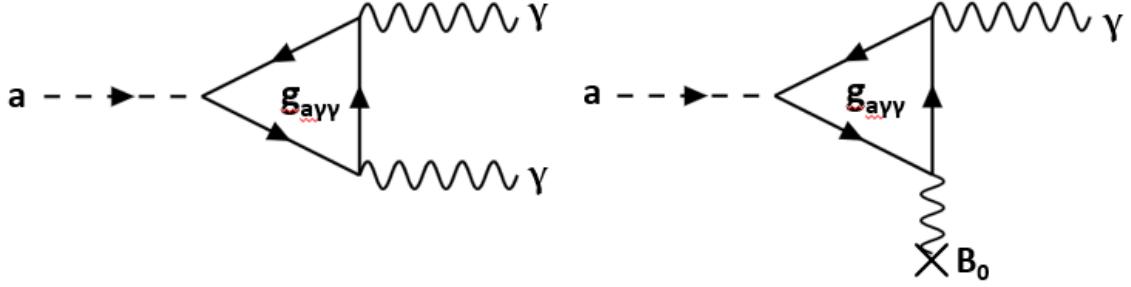


Figure 2.1: Effective coupling of axions to photons. Left: Decay into two photons (conversely, coalescence of two photons into an axion). Right: Primakoff conversion to a photon in the presence of a magnetic field.

2.1.3 Production in the Sun

One potential source for these massive KK axions, of the order of ~ 10 keV, is the Sun. Axions can be created through a number of processes. The first one, and the source of most PQ axions, is the Primakoff effect, $\gamma + Ze \rightarrow Ze + a$, where a thermal photon converts into an axion in the magnetic field of nucleons and electrons of the solar plasma. The second one is the coalescence of photons into an axion, $\gamma + \gamma \rightarrow a$, the reverse of an axion decay. The third one, which dominates in axion models where they couple to electrons at tree level, are the so-called ABC reactions

³In fact, since these higher mass modes are independent of the mass of the PQ axion, this model works for any axion-like particle that propagates in the extra dimensions with the same effective coupling to two photons and a mass much smaller than the inverse compactification radius. However, since this distinction does not bring any actionable difference for this work, compared to just considering the base state of the KK excitations as the PQ axion, we will continue to refer to it as such to stay consistent with the phrasing in [6]

⁴The compactification radius of the extra dimensions depends both on the number of extra dimensions n and the fundamental quantum-gravity scale M_F . Assuming it is common to all the extra dimensions, the relationship follows $R \sim 10^{32/n-12} (\frac{1 \text{ TeV}}{M_F})^{1+2/n} \text{ eV}^{-1}$ [5], with corrections of order unity depending on the type of compactification.

⁵To compare with the previous axionic Dark Matter footnote, if DM was dominated by $m_{a_n} = 10$ keV KK axions (but keeping the same base mass $m_{PQ} = 10^{-1}$ eV for the standard axion as before), a detector with a more reasonable size of 10^3 m^3 would see on the order of one event per year. Locally higher densities of KK axions would boost this even further.

(Atomic recombination and deexcitation, Bremsstrahlung, and Compton); to avoid introducing new coupling parameters into the solar KK axion scenario, we restrict ourselves to hadronic axion models where there is no tree-level coupling with electrons (e.g. the KSVZ model for the PQ axion). Diagrams for the first two processes can be found in Fig. 2.1.

Reference [6] used a standard model of the Sun to predict its production of KK axions through both processes. For Primakoff axions, they integrated over the black-body photon flux inside the Sun, all target species, and all scattering angles, considering solar photons as massless; due to target nuclei being essentially at rest, the conversion of photons into axions through what is essentially a “scattering” interaction suppressed slow-moving axions. For axions produced through coalescence, an approximate Maxwell-Boltzmann photon occupation number was used; while the absolute number of axions generated this way is smaller than through Primakoff, it becomes more important for heavier KK axion mode creation. The axion luminosity of the Sun was found to be:

$$L_a = A \cdot L_{\odot} \cdot \left(\frac{g_{a\gamma\gamma}}{10^{-10} \text{ GeV}^{-1}} \right)^2 \left(\frac{R}{1 \text{ keV}^{-1}} \right)^{\delta} \quad (2.10)$$

where $L_{\odot} = 3.85 \cdot 10^{33} \text{ erg/s}$ is the standard solar luminosity, and A is a multiplicative factor that depends on the axion formation process and the number of dimensions δ ; some values for $R \sim 1 \text{ eV}^{-1}$ are shown on Tab. 2.1. The R^{δ} term comes from the density of axion modes, since they are separated by an energy of $1/R$ in δ dimensions. Of note, additional exotic energy losses in the Sun would result in increased consumption of nuclear fuel in its core, with a corresponding increase in core temperature; we will get back to that point in Sec. 2.3.2.

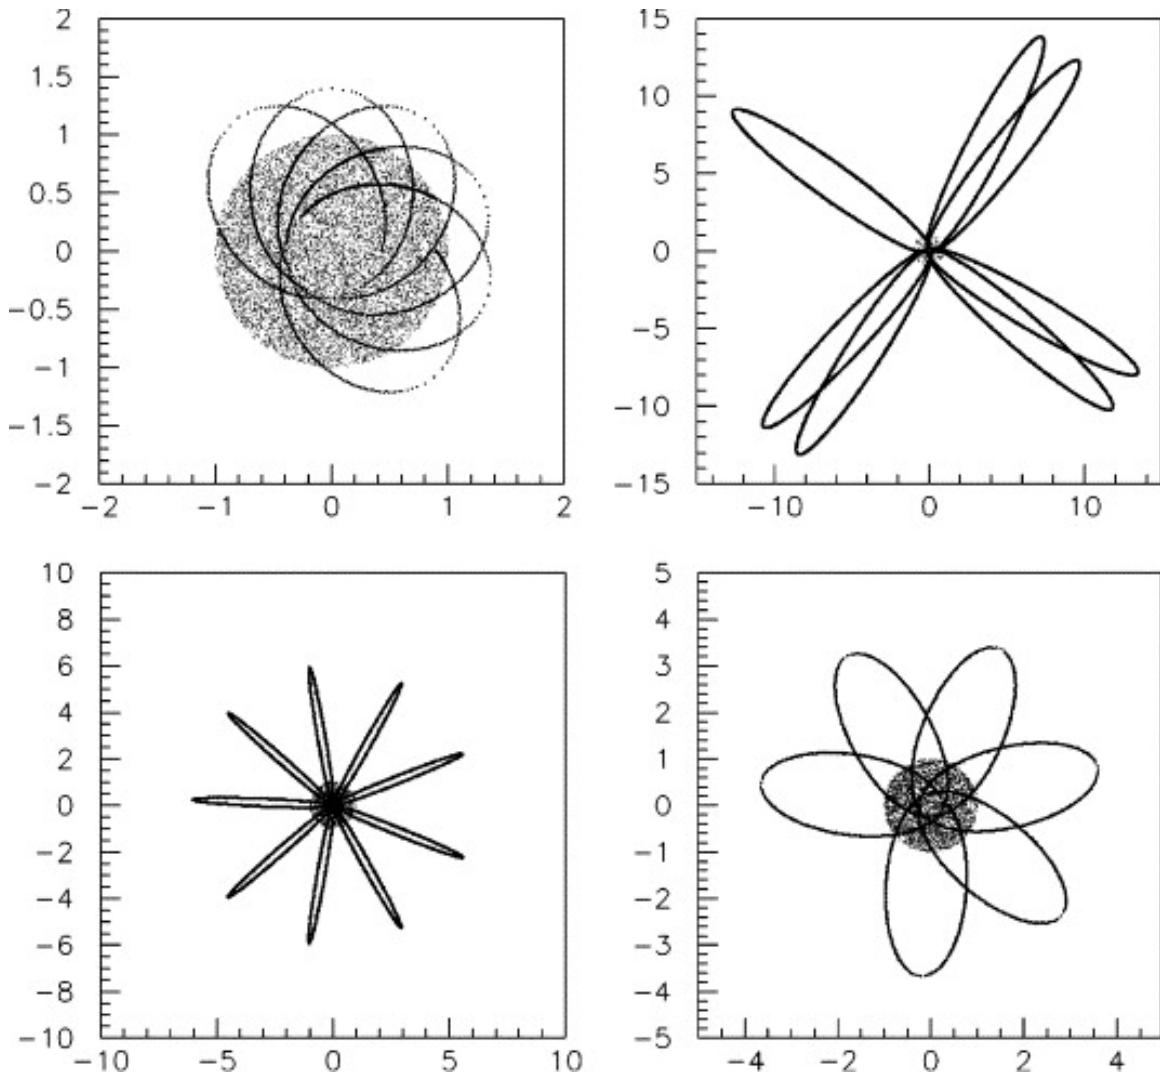


Figure 2.2: First few revolutions of trapped solar KK axions. The coordinates are given in solar radii, with the shadowed region in the center outlining the solar disk [6].

| | Primakoff | Coalescence | Sum |
|--------------|-----------|-------------|-------|
| $\delta = 1$ | 0.015 | 0.0033 | 0.018 |
| $\delta = 2$ | 0.12 | 0.067 | 0.19 |
| $\delta = 3$ | 0.99 | 1.06 | 2.05 |

Table 2.1: Coefficients for A in equation 2.10 [8].

A crucial consequence of this model is the accumulation of long-lived, massive axions in the vicinity of the Sun over its lifetime, up until reaching equilibrium with their decays. Indeed, some proportion of heavy KK axions created in the Sun will leave its surface with speeds under the escape velocity (see Fig. 2.3), and remain trapped in closed orbits around the Sun (see examples in Fig. 2.2).

The simulation in [6] found that the proportion of trapped KK axions for $R = 1 \text{ eV}^{-1}$ and $\delta = 2$ was $f_{trap} = 5 \cdot 10^{-11}$ for Primakoff axions, and $f_{trap} = 9 \cdot 10^{-8}$ for coalescence axions; the large difference between the proportion of trapped axions between both processes means that the resulting KK axion population is dominated by axions produced through coalescence of two photons in the Sun. The orbits of trapped KK axions were tracked in the simulations to determine the accumulation of axions at different distances from the Sun. A dependency on distance from the Sun of $1/r^4$ was obtained [10].

Assuming a steady-state Sun throughout its lifetime, $R = 1 \text{ eV}^{-1}$ and $\delta = 2$, the density of trapped solar KK axions on Earth today, and by extension their decay rates, are shown on Fig. 2.4j. For an axion density on Earth of $n_a = 1.0 \cdot 10^{14} \text{ m}^{-3}$, and an axion-photon coupling of $g_{a\gamma\gamma} = 1.0 \cdot 10^{-13} \text{ GeV}^{-1}$ (see next section for justification of these magnitudes), we expect a rate of axion decays of $\sim 0.2 \text{ evt/m}^3/\text{day}$, mostly in the $5 - 15 \text{ keV}$ range.

Note that, due to the solar axion luminosity varying as $g_{a\gamma\gamma}^2$ as shown in Eq. 2.10,

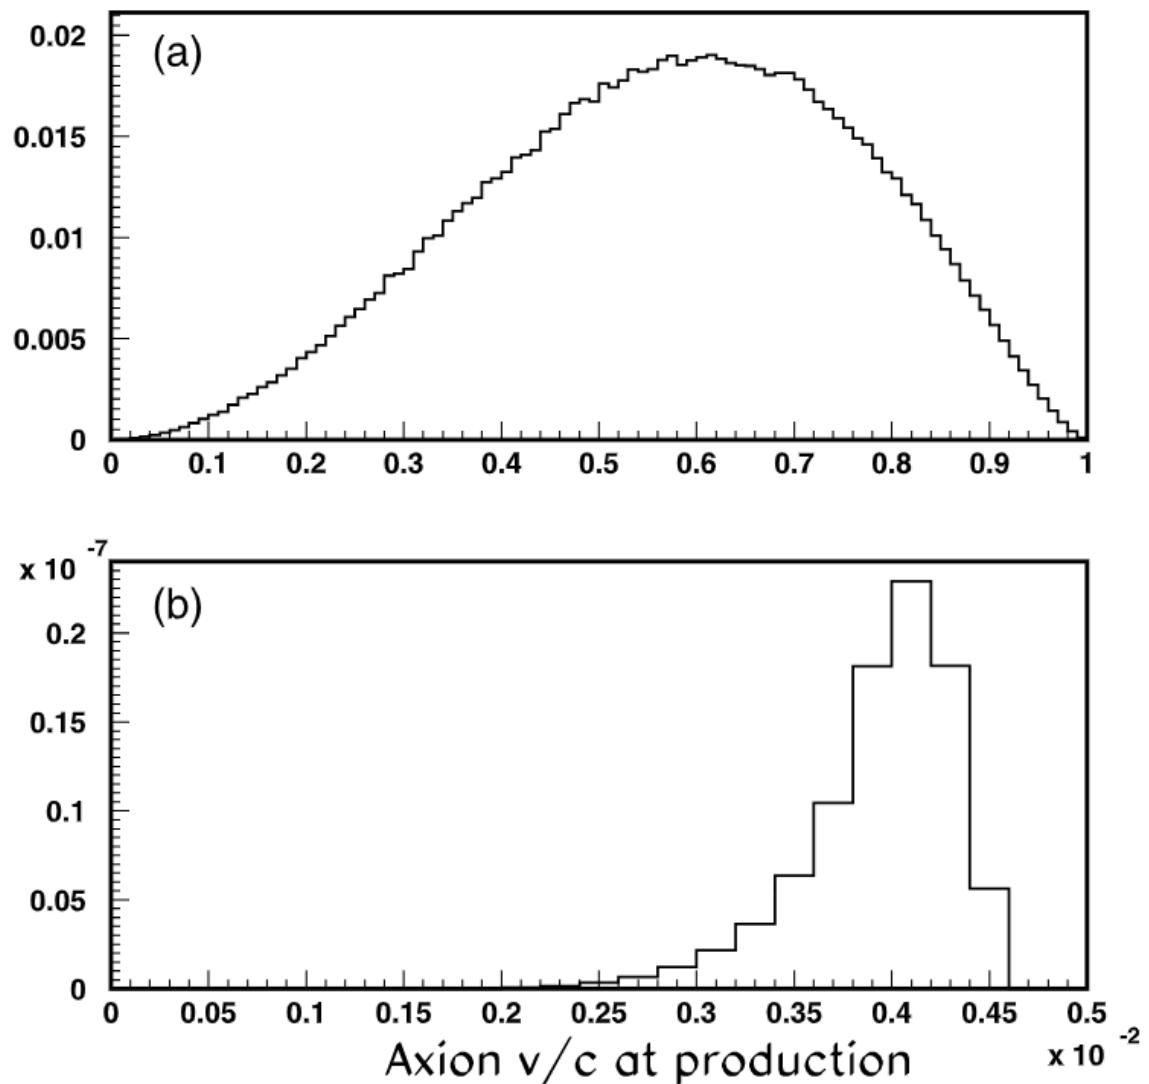


Figure 2.3: Velocity distribution for KK axions produced in the Sun via photon coalescence. a) All axions (normalised to unity); b) gravitationally trapped (normalized to f_{trap}). Note the different x scales for both plots [6].

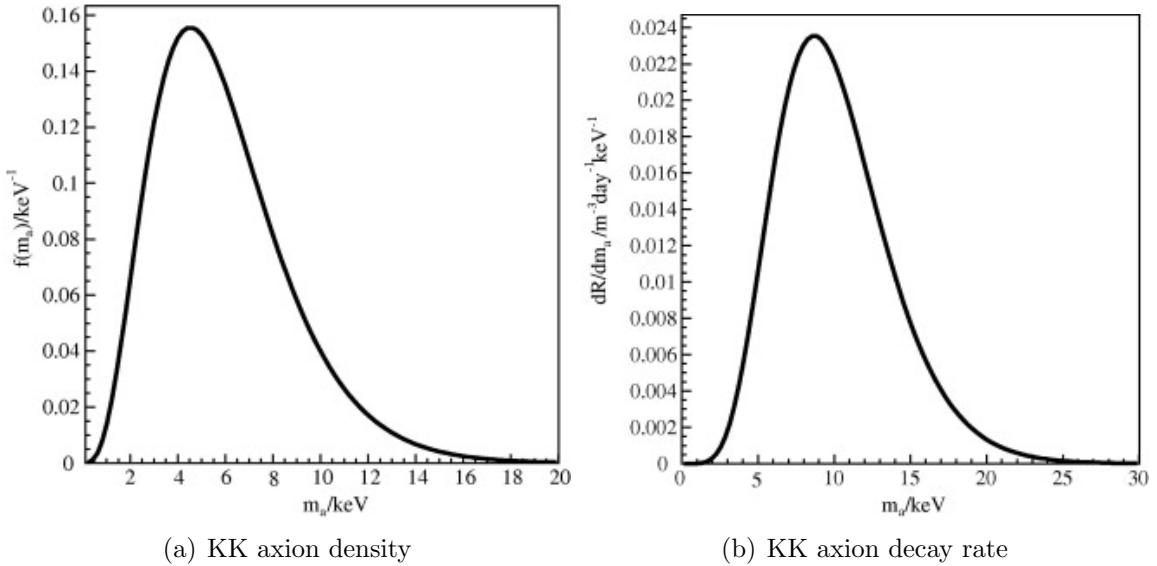


Figure 2.4: KK axions on Earth for $n_a = 1.0 \cdot 10^{14} \text{ m}^{-3}$ and $g_{a\gamma\gamma} = 1.0 \cdot 10^{-13} \text{ GeV}^{-1}$.
 Left: KK axion density. Right: KK axion decay rate [9].

the density of axions is expected to vary in the same way. And since the lifetime of axions also depends on $g_{a\gamma\gamma}^2$, this means that, everything else being equal, the expected decay rate of trapped axions on Earth should vary as $g_{a\gamma\gamma}^4$.

It should be emphasized that the model used assumed that electric charges in the Sun are isolated, and that initial state photons are massless. These are incorrect because the effective photon mass is given by plasma energy, which is ~ 300 eV in the Sun. A non-zero photon mass likely has an effect on simulation results; especially for trapped axions, since they are produced with low velocities. In fact, [6] explicitly states that their results may only be qualitatively correct, and insist on the importance of more precise simulations. Even for a given R and δ , the proportion of trapped KK axions, their mass distribution, and radial density of axions could all be affected by the simplifying assumptions made.

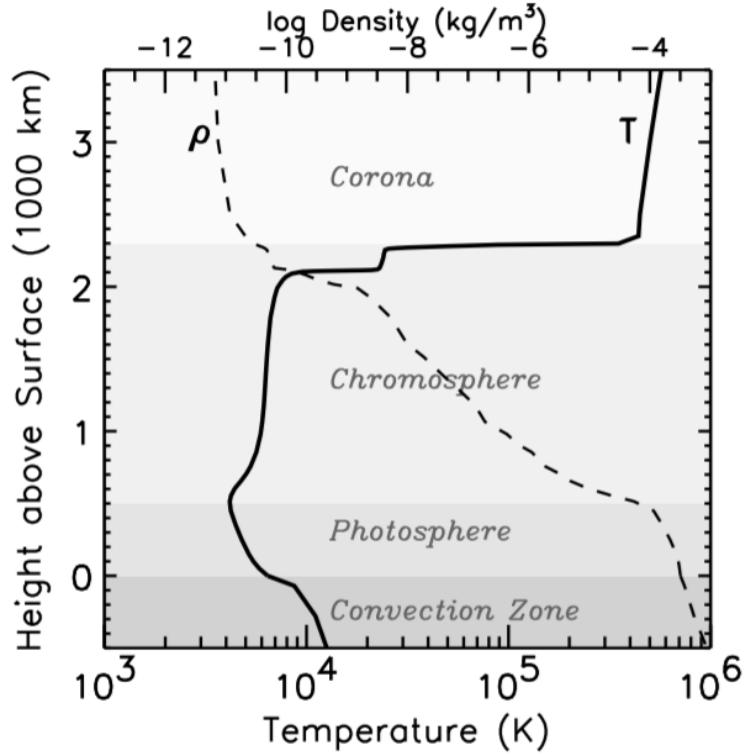


Figure 2.5: Temperature and density of the Sun surface and atmosphere [11]. The jump to $\sim 5 \cdot 10^5$ keV above 2000 km is in apparent defiance of the second law of thermodynamics .

2.2 Evidence for solar KK axion

We can now briefly describe some of the astrophysical observations that could be explained by an accumulation of KK axions trapped around the Sun, based on [6].

2.2.1 Solar Corona hint

Corona heating problem

Solar X-rays reveal an unexpectedly high temperature for the solar corona, as shown in Fig. 2.5. It is not well understood how it can be in equilibrium with the solar surface underneath, almost 300 times cooler; this would violate thermodynamics, since the

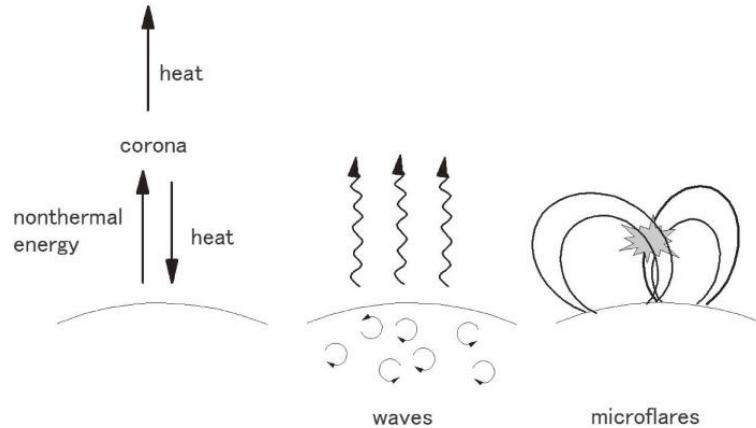


Figure 2.6: Energy flows to/from the corona (left), and two representative heating mechanisms [13].

source of energy of the Sun are the nuclear reactions in its centre, so the temperature should decrease with distance from its core. For the high temperature of the corona to be maintained, some other form of energy must be dissipated in the upper atmosphere. Since the amount required to compensate for thermal, radiative and convective losses in the Corona accounts for around 0.01% of the total solar output, the Sun makes for a likely source of that energy⁶. But while there are multiple processes that can transfer energy away from below the surface and into the atmosphere, the mechanism by which it is then dissipated thermally in the Corona remains an open question.

The leading categories rely on the shuffling of magnetic field lines in the Photosphere interacting with the Corona, and are usually split in two, as shown in Fig. 2.6: those where the Photosphere driver has time scales shorter than the Alfvén transit time in the Corona (“AC”), and those for which they are longer (“DC”) [12–17].

- AC motions generate both magnetohydrodynamic (MHD) and acoustic waves. Alfvén waves in particular do not reflect or refract at the Transition Region

⁶More precisely, likely from the magneto-convective motions at and under the Photosphere[12].

between the Corona and the Chromosphere, and their nominally weak damping in the Corona could be boosted through resonant absorption or phase mixing.

- The DC mechanisms involve magnetic tensions building gradually between highly localised magnetic current sheets, releasing their energy explosively through magnetic reconnection between field lines in opposite directions, in what are usually called micro or nanoflares. A sufficiently high rate of these events could explain the coronal heating.

Ultimately, the Solar Corona heating problem involves physics at multiple levels and scales. It is not enough to know the energy source and its conversion mechanism, but also how the solar plasma responds to the additional heating, and how that affects the emitted radiation from the Sun and hence the actual observables [12]. While the different hypotheses put forward could all be part of the solution, or even *be* the solution to the problem, due to the complexity of the system there is currently no confirmation of any of their contributions. Current observational capabilities cannot detect any heating mechanism at work, and computational resources cannot complete simulations covering all scales of their effects⁷ [12, 14, 15].

KK axion solution

An alternative explanation appears by analogy with the behaviour of Earth's atmosphere, as shown in Fig. 2.7. While the Earth's atmosphere is relatively transparent to visible and near-IR radiation above 310 nm, allowing $\sim 99\%$ of the Sun's energy to penetrate below 15 km, the remaining 1% of UV, EUV and X-ray radiation control

⁷It is also unclear whether they can explain the narrow width of the Transition Region between the Chromosphere and the Corona. I could not find an answer one way or another in the literature.

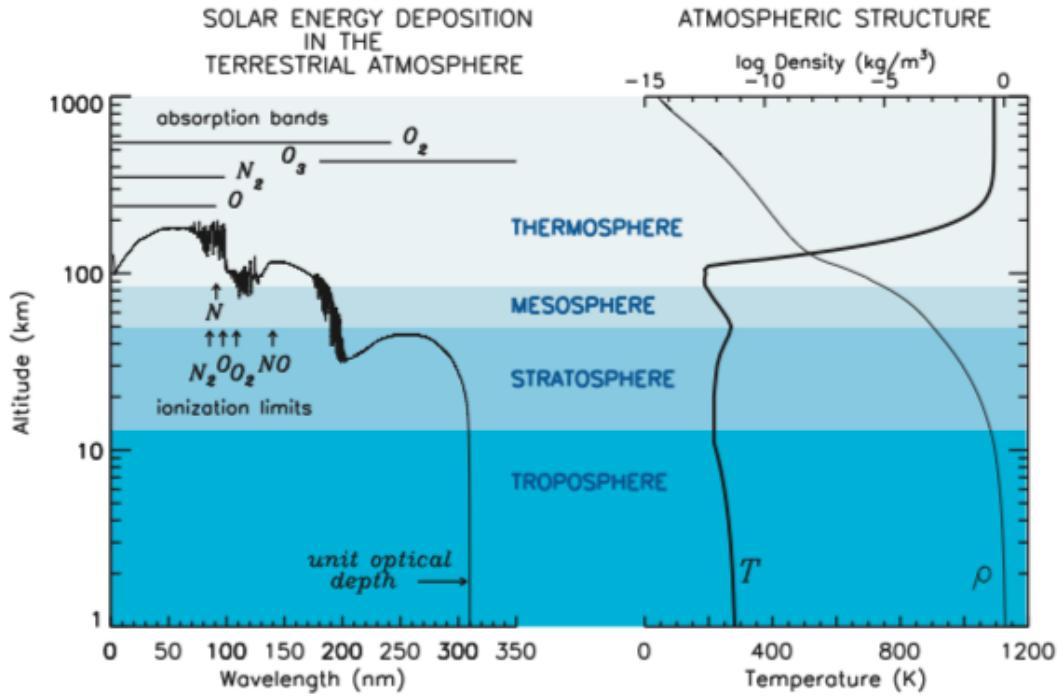


Figure 2.7: Left: Altitude at which the Earth’s atmosphere attenuates incident radiation from the Sun by a factor e , depending on wavelength. Right: Temperature and density of Earth’s atmosphere [11]. The similarities between the temperature of the solar and Earth atmosphere suggest the behaviour of the Solar Corona could be explained by an external irradiation source [6].

the overall thermal profile of the atmosphere above that height. In particular, all primary atmospheric gases (N_2 , O_2 and O) present in the upper atmosphere absorb radiation at wavelengths shorter than 100 nm, providing a way for solar radiation to dissipate in the thermosphere, heating it to more than 1000 K without impacting the lower atmosphere [11].

Given the similar temperature and pressure profile for the Sun’s atmosphere, an external irradiation source appears as a possible natural explanation for the Solar Corona heating problem. The source of this irradiation could be the Sun itself,

through the decay of the gravitationally trapped KK axions mentioned in the previous section. Radiation in the \sim keV range and below would be absorbed by the solar corona and the transition region, explaining the abnormally high temperatures there. Radiation with higher energies could in turn explain the slow increase in temperature with altitude of the chromosphere.

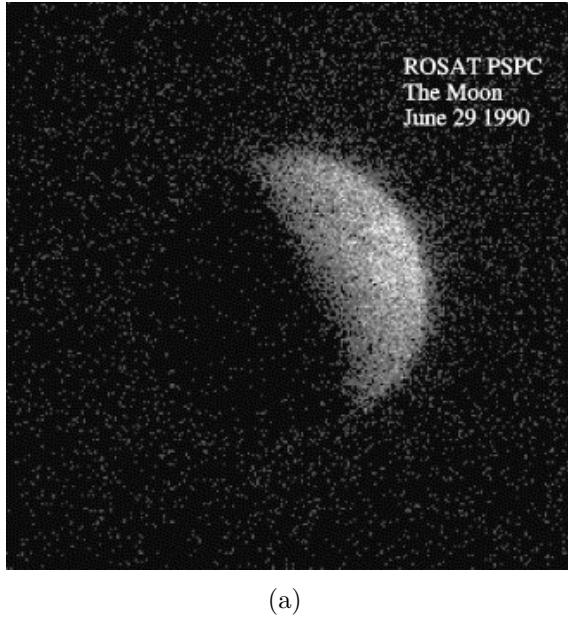
The value of $g_{a\gamma\gamma}$ is chosen so that the combined X-ray luminosity from decays of trapped KK axions matches the experimentally reconstructed one of $L_x^{2-8\text{keV}} \sim 10^{23}$ erg/s. This procedure gives $g_{a\gamma\gamma} = 9.2 \cdot 20^{-14} \text{ GeV}^{-1}$, and a density on Earth of $n_a = 4.0 \cdot 10^{13} \text{ m}^{-3}$ KK axions. This in turn gives a cross section via Primakoff under $\sim 10^{-54} \text{ cm}^2$, for which the mean free path is much larger than the total flight path of the axion even for the age of the Universe, so the fact that it spends part of its orbit in the Sun does not affect the final density of trapped axions.

While the total irradiation can be reproduced that way, the reconstructed solar x-ray spectrum contains a strong component below ~ 1 keV, which is not reproduced by the spectrum of trapped KK axions, mainly between 5 and 15 keV. This could potentially be explained by the processing of part of the ~ 10 keV photons into photons of lower energies in the outer parts of the Corona.

2.2.2 Other

[6] mentions other phenomena that could be explained by an accumulation of massive axions, although not all are necessarily consistent with each other. We describe here one more such hint for KK axions.

Another hint of solar KK axions comes from the ROSAT X-ray telescope. Aimed at the Moon, it observed X-rays coming from its dark side (see Fig. 2.8), primarily under



(a)

Figure 2.8: X-ray photon image of the Moon as measured by ROSAT. The sunlit portion of the Moon is clearly visible, and conversely the shadowed region is blocking the diffuse X-Ray Background Radiation [18].

2 keV. Their intensity was around 1% of those from the sunlit side of the Moon, and 30% compared to the X-Ray Background Radiation [18]. This is roughly 10 times higher than expected from either interaction of the solar wind with the Moon, or lunar X-ray fluorescence.

To produce this observation, trapped axions would have to decay at a rate of $\sim 200 \text{ evt/day/m}^3$ under 2 keV, and $\sim 2 \text{ evt/day/m}^3$ above. This is apparently incompatible with the Solar Corona hint, since that predicts a rate of only $\sim 0.08 \text{ evt/day/m}^3$, primarily in the 2–5 keV. With this number of decays on Earth, the total irradiation of the Sun from trapped KK axions would be orders of magnitude above the apparent one. Conversely, for the axion density that would match the solar irradiation, the density of axions on Earth would not be large enough to explain the X-rays from the dark side of the Moon above 2 keV, and even less for lower energies. Both could

potentially be explained at once if the radial density of axions decreased slower than $1/r^4$ with distance from the Sun.

This additional, if incompatible, hint for trapped solar KK axions is only mentioned to remark on the existence of multiple such phenomena. Even if the region of the parameter space suggested by the solar corona hint was ruled out, empirically or theoretically, other regions might replace it. Reference [6] suggests multiple other hints. In the rest of this work, only the axion density and axion-photon coupling suggested by the Solar Corona problem will be considered.

2.3 Derived constraints

The search for the axion and axion-like particles (ALPs) is a very active field, with many different approaches brought to bear at the problem, as can be seen in Fig. 2.9. However, not all constraints on ALPs are directly applicable to the KK axion model discussed in this work.

Since many of these methods are looking for a single mode of the axion at low energies (< 1 eV), they are insensitive to the excited modes of the axion postulated by this model. Any constraint on $g_{a\gamma\gamma}$ they set would only be applicable if it constrains *all* possible values for m_{PQ} mode, since the accumulated density of solar KK axions is composed only of KK modes, largely independent of its value.

A complete review of all ALP constraints as they apply to the solar KK axion model is beyond the scope of this work. A qualitative “translation” to our model will be offered for those most relevant to this study, based on the most recent Particle Data Group (PDG) review [2]. A more in-depth review from someone with the required expertise, especially for cosmological and astrophysical constraints, is sorely needed.

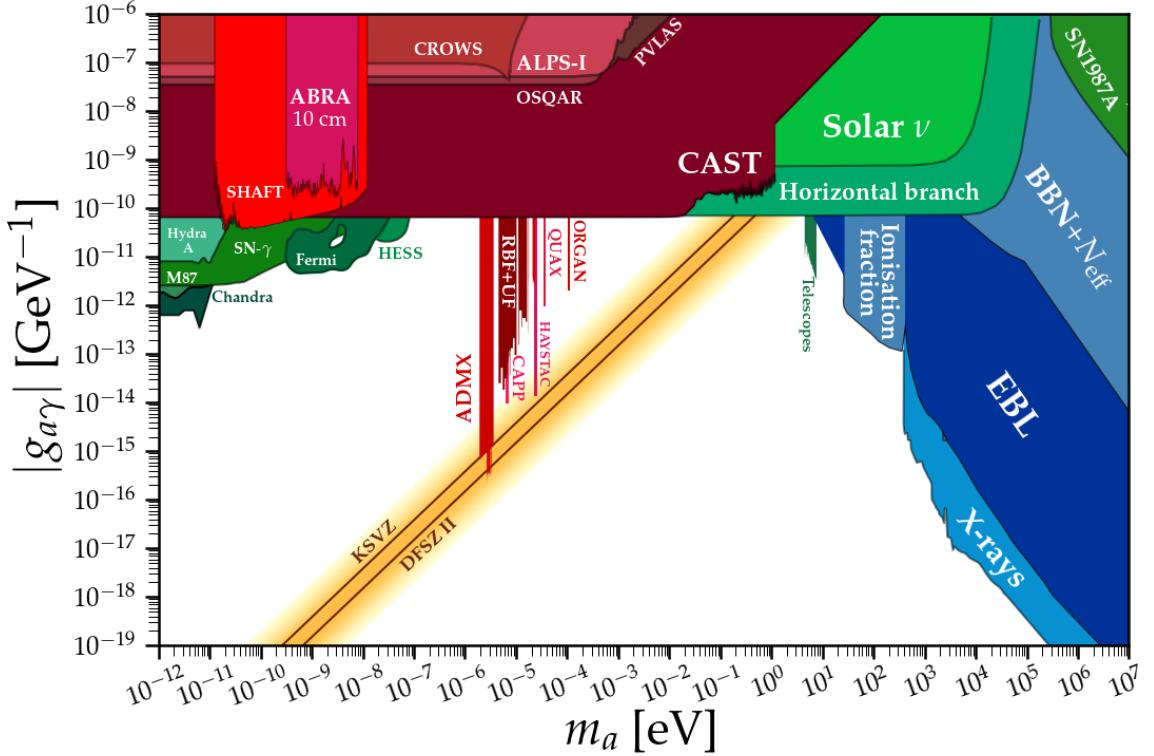


Figure 2.9: Exclusion limits on $g_{a\gamma\gamma}$ and m_a for non-KK axion-like particles. Plot and references can be found in [19]. Red exclusion limits are derived from “experiments”, green from astronomical sources, and blue are cosmological.

2.3.1 Cosmological constraints

Cosmological constraints are due to the consequences of axion production (either through thermal or non-thermal processes) in the early Universe.

Relativistic axions produced thermally increase the amount of radiation in the Universe, modifying the Cosmic Microwave Background’s (CMB) temperature angular power spectrum. Once non-relativistic, they contribute to Hot Dark Matter (along with massive neutrinos), suppressing low-scale structure formation and leaving an imprint on CMB anisotropies. Combining measurements of CMB anisotropies,

Halo Power Spectrum, and Hubble constant from various sources, an upper bound is set for the PQ axion at $m_{PQ} \lesssim 0.5 \text{ eV}$ [20]. It should be noted that the axion contribution is degenerate with that from neutrinos to Hot Dark Matter, so the exact value of the limit depends on the model used.

Conversely, axion production through non-thermal processes produces a lower bound on m_{PQ} instead: low-mass PQ axions (and hence with weak couplings to photons) would bring about an over-abundance of axions as Cold Dark Matter. While their exact relic density depends on whether the PQ symmetry was broken before or after inflation, and is subject to large uncertainties⁸, a minimum mass of $\sim 6 \mu\text{eV}$ is offered in the PDG review.

Both thermal and non-thermal production of PQ axions set limits on the possible values of m_{PQ} . However, since the KK axion model is relatively insensitive to m_{PQ} , they do not constrain it. The only way in which m_{PQ} is directly involved in the model is in the assumption that $m_{PQ} \ll 1/R$; an upper bound of $m_{PQ} \lesssim 0.5 \text{ eV}$ is reassuring in that sense.

On the other hand, the cosmological bounds become more complex when considering general ALPs, instead of only the PQ axion. In particular, masses above 154 eV for the ALP would generate a thermal relic in overabundance compared to Dark Matter, while ALPs with a decay time smaller than 10^{20} days would generate a photon background incompatible with the Extragalactic Background Light (EBL), among many others [21]. The KK axion model produces examples of both. It is entirely unclear how the introduction of a tower of excitations evenly spaced in mass, rather

⁸For PQ symmetry breaking before inflation, the relic density is proportional to the arbitrary misalignment angle in the Vacuum Realignment mechanism; for after inflation, the decay of cosmic strings and domain walls introduces further populations of DM axions, themselves suffering from significant uncertainties too.

than a single particle with a discrete mass, changes the limits from these sources.

2.3.2 Astrophysical constraints

Low mass particles are produced in the hot plasma of stars, and can contribute to their total energy loss. The solar KK axion model introduces new such particles, boosting the energy loss of stars by the total axion luminosity. This would have to be matched by additional consumption of nuclear fuel in the core of the star, increasing its temperature and shortening its lifetime.

In the case of the Sun, any exotic energy losses are constrained by helioseismology [22] and measurement of the core temperature through solar neutrino flux [23]: $L_{exotic} < 0.2L_{\odot}$ at 2σ . Translated to the solar KK axion model, this sets an upper bound of $g_{a\gamma\gamma} < 10^{-13} \text{ GeV}^{-1}$ for $\delta = n = 2$, $M_F = 100 \text{ TeV}$, and $R = 1 \text{ eV}^{-1}$, based on E. 2.10; the much stronger constraint than for the standard PQ axion ($g_{a\gamma\gamma} < 10^{-9} \text{ GeV}^{-1}$) is due to the increased energy loss from all the additional KK modes.

This constraint is in tension with the preferred axion-photon coupling of $g_{a\gamma\gamma} \sim 9.2^{-14} \text{ GeV}^{-1}$. For the most part, the tension cannot be relaxed by changing the size of the additional dimensions nor their number. While that could decrease the total KK axion luminosity, it would also decrease the density of trapped axions by the same factor. So to solve the Solar Corona problem, a higher axion-photon coupling would be necessary, hence restoring this tension. On the other hand, if the assumptions noted in Sec. 2.1.3 are such that the proportion of trapped KK axions was underestimated, then a lower axion-photon coupling would be enough to solve the Solar Corona problem, weakening this tension. In the absence of a more precise model for solar KK axions

that would allow to resolve this tension one way or the other, direct detection of these particles remain the only way to conclude.

Constraints also exist based on stars other than the Sun. In particular, the lifetimes of stars in the Horizontal Branch (HB) of Galactic Globular Cluster (and in turn their ratio over stars in the Red Giant Branch (RGB)) would decrease due to the additional energy losses from axions as $\sim L/(L + L_a)$, where L is their visible luminosity, and L_a their axion luminosity. For ALPs, this provides the strongest limit on axion-photon coupling for a wide range of masses, at $g_{a\gamma\gamma} < 6.6 \cdot 10^{-11} \text{ GeV}^{-1}$ [24]. It seems reasonable that it would also provide stronger limits for the KK axion model than those derived from the Sun. Unfortunately, their KK axion luminosity is not trivial to compute, since it depends on their internal structure. This work has not been performed, so there are currently no constraints based on HB star lifetimes, beyond the overly conservative limit from PQ axion production alone.

2.3.3 Laboratory searches

Light-shining-through-wall experiments

Light-Shining-Through-Wall (also known as “beam dump”) experiments, use a laser in a transverse magnetic field to induce photon-axion oscillations through a photon barrier. A first dipole magnet induces conversion of photons into axions before reaching the photon barrier, while a second one placed after the photon barrier induces photon regeneration from the axion flux. By examining the photon flux after the barrier, or lack thereof, limits on sub-eV ALPs can be set.

At time of writing, leading constraints come from OSQAR, with $g_{a\gamma\gamma} \lesssim 3 \cdot 10^{-8} \text{ GeV}^{-1}$ for $m_a < 0.3 \cdot 10^{-3} \text{ eV}$ [25]. This approach has the advantage to be

independent from the chosen axion model, and in particular independent from our knowledge of astrophysical sources of axions, since they are generated in the laboratory itself. Unfortunately, it does not set strong enough constraints to be of much interest for solar KK axion searches.

Helioscopes

Helioscopes aim at detecting the flux of PQ axions created by the Sun. A dipole magnet is aimed at the Sun, using a dipole magnet to convert axions to X-rays by Primakoff effect. CAST has set the most stringent limits with this method, with $g_{a\gamma\gamma} < 0.66 \cdot 10^{-10} \text{ GeV}^{-1}$ for $m_a < 0.02 \text{ eV}$, and $g_{a\gamma\gamma} \lesssim 2 \cdot 10^{-10} \text{ GeV}^{-1}$ for $m_a < 0.5 \text{ eV}$ [26].

Since the solar KK axion model does *not* have any effect on the flux of PQ axions generated by the Sun, only *adding* onto it a flux of heavier modes, limits on $g_{a\gamma\gamma}$ based on it remain valid. Notably, since the flux of PQ axions does not depend on the number of extra dimensions δ or their compactification radius R (unlike constraints based on star energy losses), this limit on $g_{a\gamma\gamma}$ remains constant for all KK axion models. Since we are assuming $m_{PQ} \ll 1 \text{ eV}$, we will use the limit for $m_a < 0.02 \text{ eV}$.

Annual modulation (XMASS)

XMASS is a large-volume liquid xenon scintillation detector located 1000 m (2700 m water equivalent) underground at the Kamioka Observatory in Japan. With a sensitive volume of 0.288 m^3 of liquid xenon, it accumulated data for over a year.

XMASS is currently the only pre-existing experiment that explicitly targeted solar KK axions. While it cannot distinguish axion decays from other electronic interactions directly, they exploited the seasonal component of the rate of KK axion decays. Indeed, since the density of trapped axions n_a varies as $1/r^4$ with distance from the Sun, their decay rate should experience an annual modulation as the Earth moves between its perihelion (where $n_a = 4.36 \cdot 10^{13} \text{ m}^{-3}$) and aphelion (where $n_a = 3.81 \cdot 10^{13} \text{ m}^{-3}$). The absence of modulation in the event rate in the detector allowed them to set an upper bound on the axion-photon coupling of $g_{a\gamma\gamma} < 4.8 \cdot 10^{-12} \text{ GeV}^{-1}$ for an average density on Earth of $n_a = 4.07 \cdot 10^{13} \text{ m}^{-3}$ [10].

A note on Haloscopes

One of the arguments in favour of the existence of the axion is that for some values of the symmetry breaking scale ν_{PQ} , the axion is a candidate for the Dark Matter of the universe. Haloscopes work on the assumption that the local density of Dark Matter, $\rho_{DM} \sim 0.3 \text{ GeV/cm}^3$, is entirely composed of axions. A microwave resonant cavity inside a magnetic field is used to convert axion into photons through the Primakoff effect⁹. The power generated by the resulting microwave photons is extremely low, of the order of $\sim 10^{-23} \text{ W}$, requiring cryogenic systems and ultra low noise microwave receivers to detect them. ADMX is the leading haloscope in terms of sensitivity, setting limits of $g_{a\gamma\gamma} < 4 \cdot 10^{-16} \text{ GeV}^{-1}$ for $m_{PQ} \sim 3 \mu\text{eV}$ [27], lower than any other detector at time of writing.

However, these constraints are not relevant for the solar KK axion model, due to its very restricted mass ranges: KK axions are essentially independent of the mass of

⁹The resonant frequency of the cavity has to be tuned to convert axions of a specific mass, meaning haloscopes are only sensitive to relatively short mass ranges.

their base state, to which existing haloscopes might not be sensitive. Furthermore, even if $m_{PQ} \sim 3 \mu\text{eV}$ was indeed the case, it might still not be the case that it composes the majority of Dark Matter. The axion relic density could be much lower, in which the constraints from haloscopes would be artificially stronger than they should be.

2.3.4 Compiled constraints

The applicable constraints described in the previous sections are compiled in Fig. 2.10, in terms of axion-photon coupling strength $g_{a\gamma\gamma}$ and KK axion density on Earth n_a . The prediction from the solar KK axion model, and the Solar Corona hint, are also shown; the preferred region is the intersection between both.

The strongest constraint is the one from the limit on exotic energy losses in the Sun derived from the Sun core temperature from SNO [23], with strong tension on the KK axion model as predicted by the Solar Corona hint. However, as noted in Sec. 2.1.3, the assumptions on the solar production of KK axions mean that this model remains only qualitative. In particular, if the fraction of trapped axions produced by the Sun is larger than expected, this would allow a similar density of axions to be generated with a smaller solar axion luminosity, weakening this tension.

Since the Solar Corona hint is based on the rate and energy of decays coming from an accumulation of massive axions around the Sun (and, consequently, Earth), adapted experiments should be looking for a signal coming from this same source. For some constraints derived from other types of observations, a more precise model is required to be certain of their translation into the $g_{a\gamma\gamma} - n_a$ parameter space shown in Fig. 2.10.

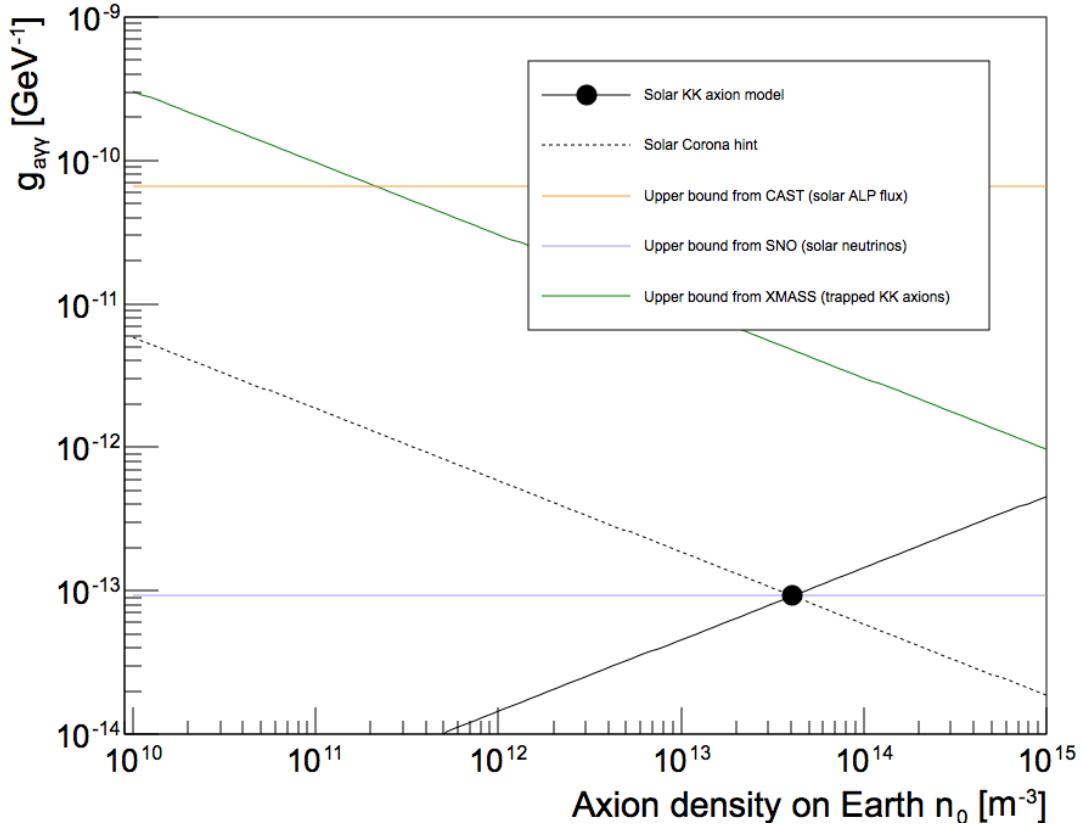


Figure 2.10: Exclusion limits on $g_{a\gamma\gamma}$ and n_a for the solar KK axion model. The exclusion limit from CAST [10] is the only one completely independent from any additional parameters in this model. The Solar Corona hint [6] and the exclusion limit from XMASS [10] depend on the mass distribution of trapped KK axions close to the Sun and to Earth, respectively. The constraint from the solar neutrino flux depends on the nature of the extra dimensions, but not on the fraction or distribution of trapped KK axions [6, 23]. The relationship between $g_{a\gamma\gamma}$ and n_a described in the solar KK axion model is the only one that depends on all these factors.

Summary

In this section, we have introduced KK axions, excited modes of the standard axion reaching masses of up to ~ 10 keV. In this model, the Sun produces heavy axions with speeds too low to escape its gravity well, which accumulate throughout the history of the Solar System. The decay of trapped KK axions would provide an external irradiation for the Sun. This would solve the Solar Corona problem (i.e. why the solar atmosphere is hotter than the Sun’s surface) if the axion-photon coupling is $g_{a\gamma\gamma} = 9.2 \cdot 20^{-14} \text{ GeV}^{-1}$; for such a coupling, the predicted axion density on Earth is of $n_a = 4.0 \cdot 10^{13} \text{ m}^{-3}$ KK axions, for a total decay rate on Earth of $\sim 0.08 \text{ evt/m}^3/\text{day}$, mainly in the $5 - 15$ keV range.

From the reviewed constraints on standard axions and axion-like particles, the largest tension with this model comes from the constraint on exotic energy losses of the Sun, at $L_a < 0.1 \cdot L_\odot$. An update to the model would likely be able to weaken this tension. Instead, it is useful to attempt to detect the accumulated axions themselves. Their decay into two photons of same energy should provide a distinctive signal, specifically in low-density detectors, as we will see in the next chapter.

Chapter 3

NEWS-G detectors

In this chapter, we will be covering the technology behind NEWS-G detectors, Spherical Proportional Counters (SPCs). We will start with a description of the working principle of particle detection with SPCs, then move on to the theory of signal formation within the detector (which will be important when discussing pulse processing in the next chapter), and finally the specifics of the most relevant SPCs in the context of this thesis.

3.1 Introduction to Gaseous Detectors

In the context of particle physics, a gas detector is a volume of gas contained between conductors, where usually the anode consists of one (or more) thin wire(s) kept at high voltage, and the signal recorded depends on the charge reaching the anode. While many types of gas detectors have been developed since the invention of Multi-Wire Proportional Chambers by G. Charpak in 1968 (for which he was awarded the Nobel prize in 1992), the principle remains fairly constant among all them.

First, an incident particle interacts with the gas in the detector, exciting some

atoms. Depending on the particle and the gas, the excitation energy will dissipate as radiation, heat, or will ionize the gas, releasing electrons. For any incident particle, its quenching factor Q will be the ratio between the mean ionization energy for an electronic recoil and that for the incident particle considered¹. With this definition, an incident electron will have a quenching factor close to one, while incident nucleons will have a lower quenching factor, with the precise value depending on the exact incident particle, its energy, and the type of gas. However, for a given type of interaction, generally the ionization energy will be proportional to the total energy deposited, so the former can be a good estimation of the latter.

Second, due to the electric field in the detector, the electrons released in the first step (referred to as “primary electrons”) will drift towards the anode². Depending on the configuration of the gas detector, it might be possible to estimate the time the primary electrons take to reach the electrode, which in turn can be a way to determine the position of the event.

Finally, when the primary electrons reach the region close to the anode, the intense field surrounding it will create an avalanche: the kinetic energy gained by an electron between two collisions is larger than the ionization energy of the gas atoms, so, one electron releases another, then they release two more, then four, etc., in such a way that a single primary electron can release thousands or more of secondary charges. Depending on the strength of the electric field, this charge multiplication process can be kept in proportional mode, so that the collected charge remains proportional to

¹For various reasons, different experiments choose different definitions for the numerical value of the quenching factor, but the central concept remains the same.

²The primary ions will in turn drift towards the cathode, but they produce no discernible signal.

the initial energy ³. Because of this multiplication process, we can observe events that would be too faint to observe otherwise.

For an in-depth overview of the technology, [28] cannot be recommended enough. A shorter and less recent (but still fairly comprehensive) set of lectures can also be found online [29].

3.2 Spherical proportional counter

A new development in the world of gas detectors is the Spherical Proportional Counter (SPC, initially proposed by I. Giomataris [7]), which will allow the exploration of lower energy ranges probed by particle detectors. The following sections on the SPC are based on the thesis by A. Dastgheibi Fard [30].

While the coming sections are meant to be generally applicable to all of our detectors, sometimes numerical values are needed. The SEDINE SPC, on which the results of this thesis were based, will be used as the reference point in those cases; that specific detector is introduced properly in Sec. 3.3.

3.2.1 Description

The SPC is extremely simple in principle, as seen in Fig. 3.1. It consists of a grounded large metallic sphere (from 0.3 m to 1.3 m in diameter) and a small ball or sensor (from 1 mm to 16 mm in diameter) kept at a high voltage (around 1000 – 2500 V) located at the centre of the vessel, forming a proportional counter. The sensor is maintained at the centre of the sphere by a grounded metallic rod, carrying inside a wire to feed

³This is not true of all gas detectors. For example, the Geiger-Muller tube is a gas detector where the high electric field saturates the avalanche. Its objective is to measure the rate of background radiation, rather than its energy.

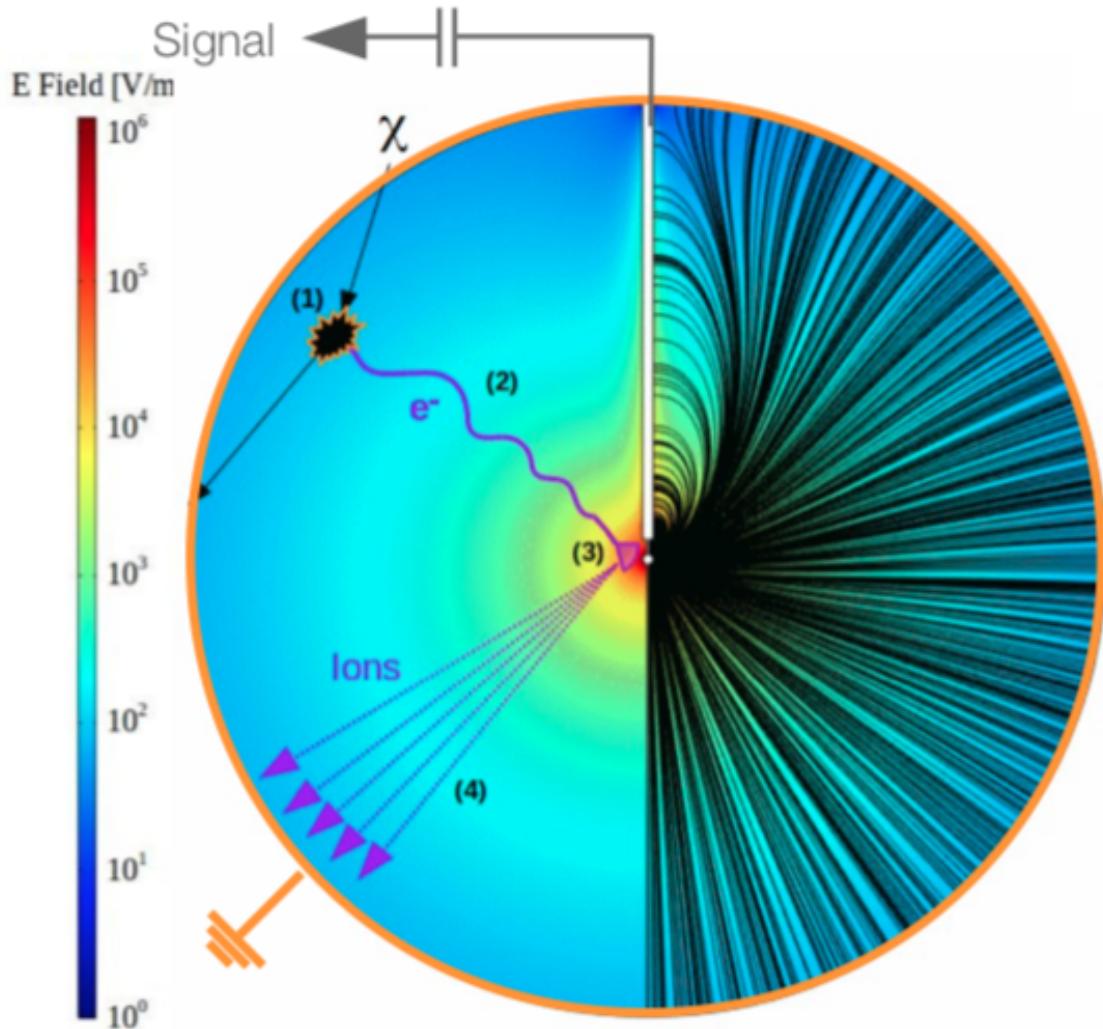


Figure 3.1: Working principle (left side), with electric field lines (right side) [31].
 1: Ionization by incident particle. 2: Primary electrons drift to central electrode. 3: Charge multiplication in the avalanche region. 4: Secondary ions drift away from central electrode, inducing a current. This current is then integrated by a charge counter, forming our final signal.

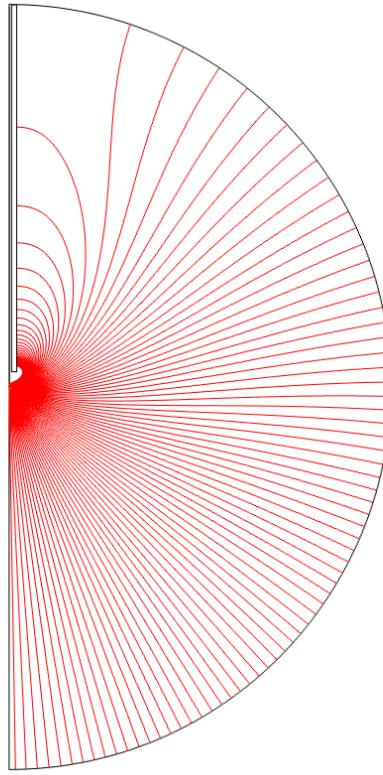


Figure 3.2: Electric field lines for an SPC with a simple electrode.

the high voltage to the sensor. Setting aside the effect of the rod for the time being⁴, the electric field varies as $1/r^2$; this allows electrons to drift to the central sensor in the low field regions that constitute most of the volume, while still triggering an avalanche within a few mm around the sensor. The charges created in the avalanche region will generate an electric current as they drift, a process described in detail in App. 3.2.3.

Unfortunately, the presence of the rod alters the electric field inside the detector, as shown on Fig. 3.2. To help solve this problem, some of the rods include a second electrode close to the sensor. The shape and voltage of this second electrode, which

⁴The effect of the rod on the electric field inside the SPC will be covered in more detail in Sec. 5.4.

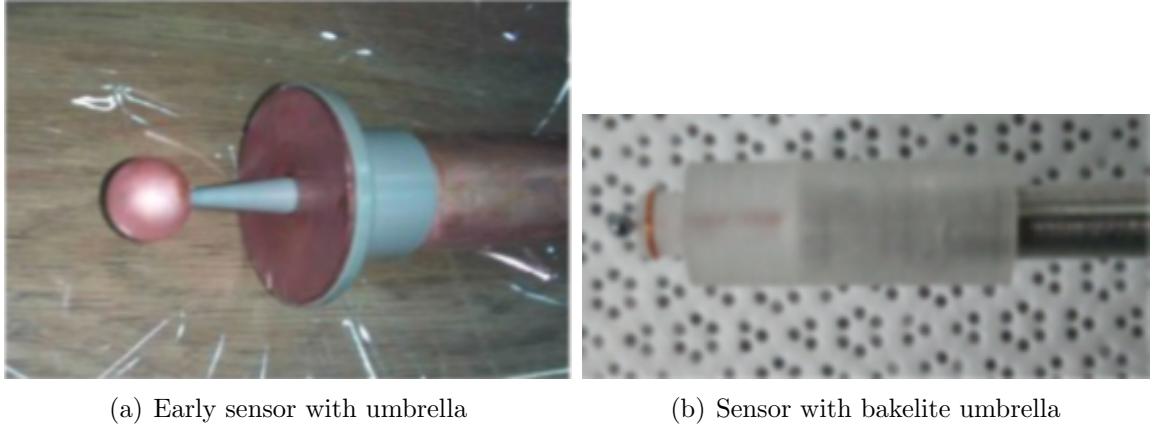


Figure 3.3: Two different versions of SPC sensor with secondary umbrella electrode.

we call the “umbrella” (cf. Fig. 3.3), depend on the goal of the data-taking. With the right setup, we can either get the electric field to keep its spherical symmetry in most of the volume of the detector, or instead redirect all the primary electrons to the SPC hemisphere farther away from the rod, lowering the dependency of the avalanche gain on the direction the primary electrons come from⁵.

For the proper functioning of the detector, we also require:

- A pumping and gas handling system.
- A high voltage feed.
- Electronics (amplifiers, filters, etc.) to record the signal coming from the sensor.
- An acquisition and/or analysis system (digitizer, computer).

Aside from the electronics (cf. App. B for how they were calibrated), I will not

⁵The development of SPC sensors is a topic of active research and development within the collaboration. For an overview of the most recent work done on sensors, see [32–34]; for a more in-depth look at the advantages of sensors with umbrella, see [35]. In this thesis, sensors will only be covered from the point of view of electric field simulations; see Sec. 5.4 for simple and umbrella sensors, and App. J for grid and achinos sensors.

cover the details of these in this document. They are still essential components to keep the detector working as intended.

3.2.2 Particle interaction within the SPC

For all possible incident particles, the different ways in which they can interact with the SPC are the following:

- Charged particles, mainly electrons and alphas. They can either:
 - create ion-electron pairs through interaction with the electronic cloud of an atom.
 - excite an atom, which releases a photon. While the photon will generally be less energetic than the ionization energy of the atom, if there are a mixture of gases in the detector, the photon might still lead to ionization if it interacts with an atom with lower ionization energy. This is called the Penning effect, and in our detectors will depend on the percentage of CH₄ in the gas mixture.
 - release Bremsstrahlung radiation through interaction with the electric field. In particular, α or β radiation can produce a γ background by interacting with the copper or lead that shields the detector.
 - scatter against a particle in the target gas. In the case of heavier charged particles, such as ions, the resulting ionization will be quenched.
- Photons. They can interact either through:
 - Photoelectric effect. The photon is entirely absorbed by an atom. The excitation energy is then dissipated either by emission of an electron or radiation. For photons with less than 10 keV in Neon, this is the main interaction.

- Compton effect. The photon interacts with a free (or quasi-free) electron without being fully absorbed, transforming some of its energy into kinetic energy. Becomes predominant in Neon for photons with more than 20 keV.
- Electron-Positron creation. The photon annihilates, producing the pair of particles. This interaction requires at least $\sim 1\text{MeV}$, and will remain pretty rare for the energy levels we are interested in.
- Neutral Particles. In the case of neutrons, they can scatter against a particle in the target gas through the Strong interaction, mediated by a pion exchange. Recoils produced by more exotic neutral particles, such as WIMPs, would likely involve either the Weak interaction, or entirely new forces. Since these events leave their energy primarily in the nucleus, they are always quenched.

3.2.3 Pulse formation

During the avalanche, a large number of electron-ion pairs are created very close to the central electrode of an SPC. As these move away from their starting point, the charge they induce on the electrode changes (per the Shockley-Ramo theorem [36]), which gets integrated by the preamp, and a pulse forms on the digitizer. The bulk of that pulse is generated by the ions, as they drift through the large potential difference between the central electrode and the shell of the detector.

For an ideal spherically symmetric detector (i.e., with no rod), with detector radius r_1 , electrode radius r_2 , applied voltage V_0 , gas pressure P , and ion mobility in the gas mu_0 , the current induced by the drifting ions on the electrode is (cf. App. C):

$$I_{ind}(t) = -q_{ions}\alpha\rho(r_2^3 + 3\alpha t)^{\frac{-4}{3}} \quad (3.1)$$

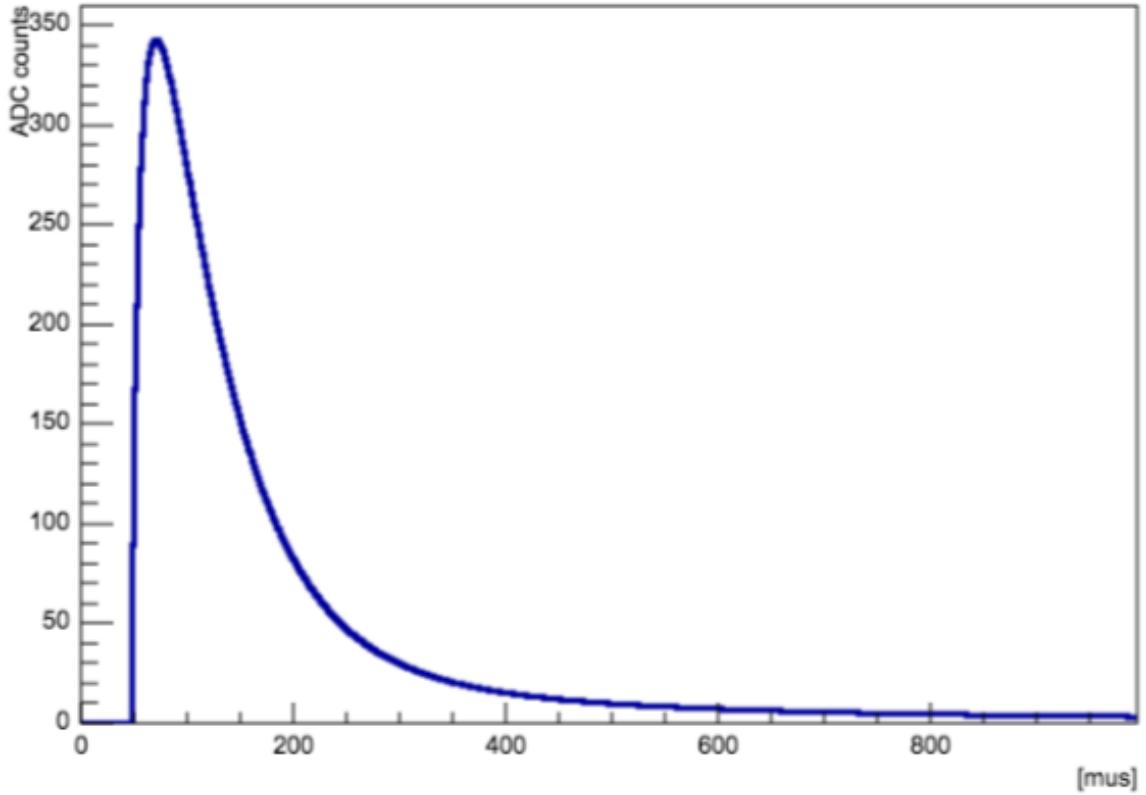


Figure 3.4: Theoretical output voltage from the preamplifier for an avalanche that creates 1000 electron-ion pairs at time $48 \mu s$, assuming SEDINE-like conditions.

where $\alpha = \mu_0 \frac{V_0}{P} \rho$, ρ is the density, r_2 is the radius of the detector, and q_{ions} is the total charge from the secondary ions. The ions drift for a total time of $t_{max} = \frac{r_1^3 - r_2^3}{3\alpha}$, which is usually in the range 1 – 50 s, increasing with detector size. This current is then convolved with the response function of the preamp:

$$f(t) = G_{preamp} e^{-t/\tau} \quad (3.2)$$

where G_{preamp} is the gain of the preamp (in V/C) and τ is its decaytime. The combined pulse shape is shown in Fig. 3.4.

The reason why understanding the signal formation is important is because, if we know what a single ion looks like in our detector, we can extract the actual physical parameters we are interested in from the pulses we record, as will be discussed in chapter 4.

Furthermore, because we are usually interested in relatively low energy events, a few extra statistical considerations have to be taken into account to understand the signal formation inside the SPC. First, the number of primary electrons released is not deterministic, but rather follows a Poisson distribution with mean $n_p = E_I/W_I$. E_I is the proportion of the recoil energy that appears as ionization energy in our detector, $E_I = Q \cdot E_R$. For a nuclear recoil, we typically have $Q \sim 0.2 - 0.7$, depending on the energy of the incoming particle. W_I is the mean ionization energy of the target gas, i.e., the average energy required to extract an electron from an atom. Second, the avalanche isn't deterministic either. The exact process of the avalanche is complex, and complicated by the non-uniform electric field, but assuming that photoelectric and charge accumulation effects are negligible, the gain distribution can be parametrized by a Polya distribution [37], as plotted in Fig. 3.3,

$$P(\nu) = \frac{(1 + \theta)^{1+\theta}}{\Gamma(1 + \theta)} (\nu)^\theta \exp(-(1 + \theta)\nu) \quad (3.3)$$

where $\nu = \frac{n}{\bar{n}}$ is the ratio between the number of secondary charges created and the average avalanche gain, and θ is a form factor for the distribution. Calibrations have led to values of the order of a few thousands for \bar{n} (highly variable depending on running conditions), and $\theta < 0.1$ for our spherical detector [38].

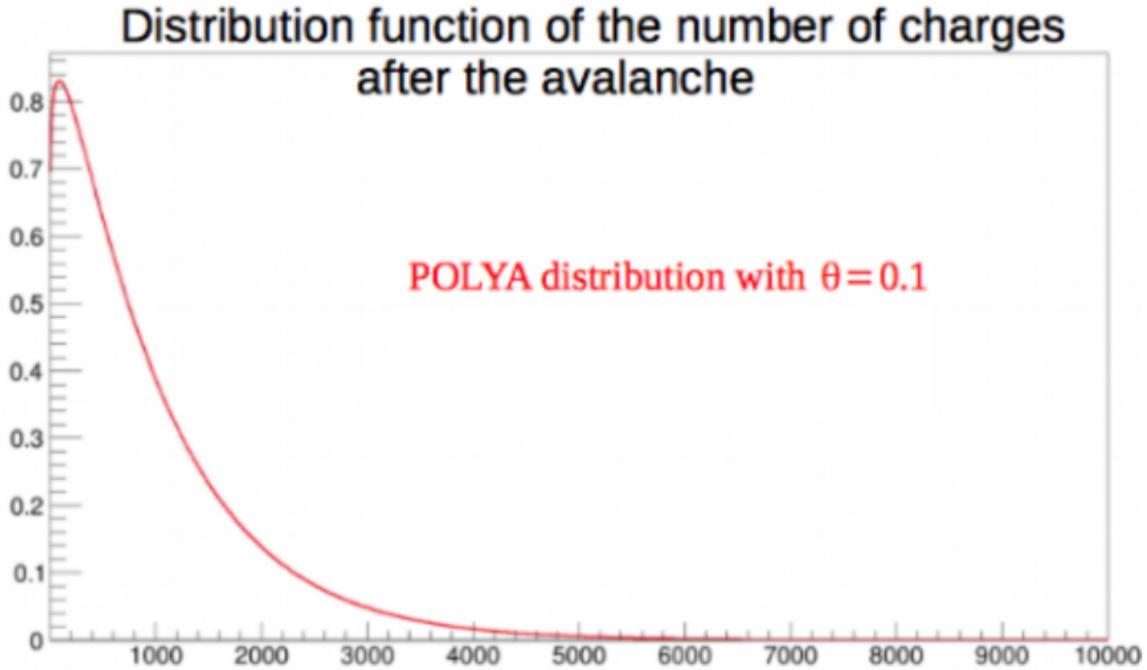


Figure 3.5: Gain distribution, i.e., number of secondary electrons created per primary electron. This is a parametrization of the actual, unknown distribution.

3.2.4 Event discrimination

An important specificity of our detector comes from the analysis of the shape of our pulses; more concretely, their risetime, which we define as the time between the point where the pulse reaches 10% and 90% of its amplitude⁶. Through the analysis of the risetime, we can determine the position of the event in our detector, and separate pointlike events from tracks.

For pointlike events, all the electron-ion pairs are created in the same spot of the detector. The electrons will drift towards the electrode, but as they do, they will diffuse, and so they will not reach it at the same time. While we are not directly

⁶Generally, since we expect a roughly gaussian-like distribution in the arrival time of primary electrons, a wider range (e.g. 5% to 95%) will tend to improve the discrimination power of the risetime metric, while a narrower one (e.g. 25% to 75%) will improve its robustness against baseline noise. 10% to 90% was chosen as a reasonable compromise between both effects.

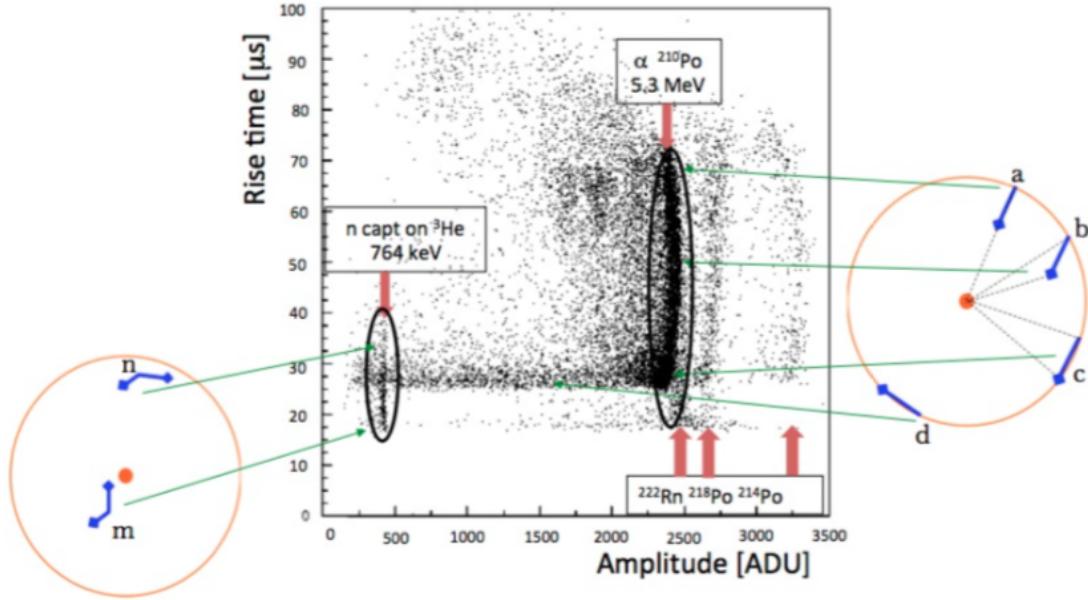


Figure 3.6: Risetime vs Amplitude of the signal for a 200 mbar Ar + CH₄ (2%) + ³He (0.4%) gas mixture. The horizontal line at $27 \mu\text{s}$ corresponds to surface events. [30]

sensitive to the drift time itself (since we don't know when exactly the event happens), the larger the “spread” in arrival time, the wider the pulse will be, and the longer the risetime we record. And since the longer the drift time, the more the electrons will be subject to diffusion, we conclude that the risetime is related to the drift distance. As such, we can use our measure of the risetime to discriminate events happening closer to the electrode than to the outer surface of the detector.

For track events, the behaviour is different. Charged particles do not leave energy in a single spot, but instead interact with the target mass all throughout their near-instant trajectory through the detector. In effect, the risetime is no longer dominated by the diffusion time at any single point, but instead by the difference in drift time between the parts of the track that are respectively closer to and farther from the

electrode. So track events will have a longer risetime than pointlike events.

Fig. 3.6 is useful to understand how different type of events relate to our measure of the risetime. This data was taken with a 200 mbar Ar + CH₄ (2%) + ³He (0.4%) gas mixture. The two main backgrounds we see are α particles released from the surface of our detector (amplitude of ~ 2400 ADU⁷), and neutron capture by ³He (amplitude of ~ 450 ADU). The different cases are:

- a) The α particle is emitted radially. The difference in drift time between the start and end points of the track is longer, and we record a longer risetime.
- b) The α particle is emitted obliquely. The difference in drift time between the start and end points of the track is shorter than before, even though the itself track is just as long.
- c) The α particle is emitted nearly tangentially. The difference in drift time between the start and end points is negligible, and the diffusion time dominates again. The risetime is roughly the same than for surface pointlike events.
- d) The α particle is emitted tangentially, and leaves the detector before depositing all its energy. The risetime remains the same as in the previous case, but the amplitude is lower.
- m) The neutron capture happens close to the electrode. The released electrons don't have much time to diffuse before reaching the electrode, and we record a very short risetime.
- n) The neutron capture happens far away from the electrode. The electrons have more time to diffuse, and we record a longer risetime.

⁷Analog-to-Digital Unit, the unit of the acquisition system.

3.2.5 Advantages and drawbacks

Now that we have properly described the functioning of the SPC, we note that the advantages of this type of detector over others used for rare event searches are:

- Simplicity. The physics of the detector are easy to understand, so adapting simulation and analysis to new conditions is practical.
- Sensitivity to low energy events, down to single electron (~ 10 eV). This is in part due to a very low electronic noise due to the low capacitance of the detector, $C < 1\text{pF}$, which is mostly independent of its size.
- Surface and track-like event discrimination via risetime cutoff.
- Flexibility in the target. By changing the gas type and pressure, we can change the sensitivity of the detector to different incident particles. For KK axion searches, gas and pressure can be chosen to optimize photon capture and separability.

The last point should be emphasized. Due to the target mass inside the detector being a gas, the photons from the decay of a KK axion will interact in two separate locations. This allows for excellent background rejection, by keeping only events with two pulses. A detector with a liquid or solid target mass would absorb both photons at the location of the decay, negating this advantage.

Its main disadvantages are the relatively small target mass of gas detectors, and the lack of a second channel that would allow discrimination of electronic from nuclear recoils. The first disadvantage is not a problem for KK axion searches, since the detection rate depends only on the volume, not the target mass; however, it is

a problem for WIMP direct detection searches. The second disadvantage remains valid: an additional readout channel could potentially help rejecting some forms of background, such as neutron captures.

3.3 SEDINE

SEDINE is the long-running SPC installed at the Laboratoire Souterrain de Modane (LSM) in France [39]. The data taken by this SPC detector was the basis for all my analysis. I will now provide a quick overview of the more concrete practical aspects of SEDINE, but more details can be found in [30, 35, 40, 41].

3.3.1 Setup

The sphere, rod and electrode are made of electropure copper (98% of the total composition of the detector). The sphere has an inner radius of 30 cm. It is connected to the outside of the shielding via an S-shaped copper tube which serves both to connect the electrode to the high voltage source, and the inner volume to the gas handling system. See Fig. 3.7.

The sphere is filled with a gas mixture, usually 98 – 99.3 % noble gas (Argon, Neon or Helium) with a purity < 1 ppm, and 0.7 – 2 % of CH₄ with a purity < 5 ppm. The role of the CH₄ is to serve as a quencher gas, i.e., a gas that can absorb radiation and release its energy through other channels (like vibration or molecule breakdown), slowing down the avalanche process so as to stay in the proportional regime. To maximize exposure for rare event searches ⁸, the detector was kept at as high a

⁸The main goal of the NEWS-G collaboration being to detect low-mass WIMPS (see [42] for first results), we need as much target mass as possible to increase the interaction rate, hence the high pressures. This is of course irrelevant for solar KK axion searches, since we are looking for particle decays, rather than interactions.



Figure 3.7: SEDINE. The S-shaped tube connects the sphere with the outside of its shielding [30].

pressure as possible, up to $3 +$ bar. In turn, to keep a signal amplification in the order of a few thousands, the central electrode had a high voltage in the ~ 2500 V applied on it.

Another crucial part of the detector is its protection against natural radioactivity (see Sec. 3.3.2 for details on its sources). Three layers of shielding protect SEDINE, as shown in Fig. 3.8:

- The first, outer layer is composed of polyethylene $(C_2H_4)_n$ bricks, 30 cm wide. It thermalizes the neutrons coming from the rock inside the LSM. Neutrons at LSM come mostly from cosmic muons interacting with the walls of the lab, or from spontaneous fission or (α, n) reaction by ^{238}U and ^{232}Th in the rock. The polyethylene also blocks part of the ambient γ radiation.
- The second layer is composed of lead bricks, 10 cm wide. Its purpose is to stop the ambient γ radiation coming from the Uranium, Thorium and Potassium in the rock walls and the concrete covering them. Together with the polyethylene

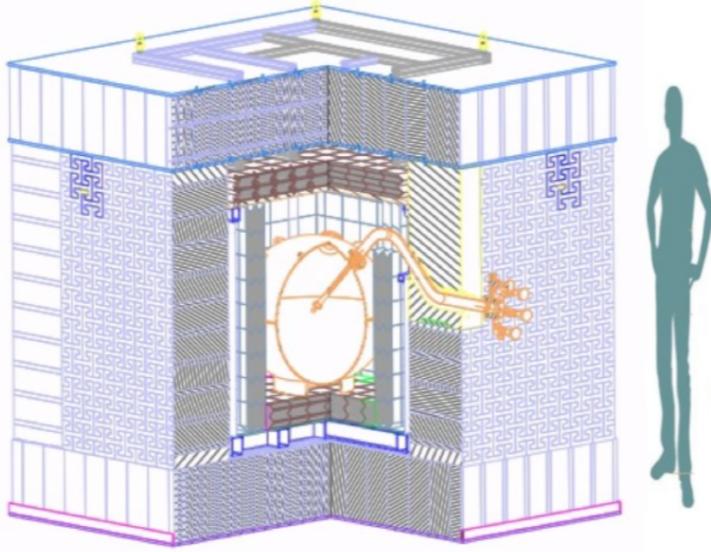


Figure 3.8: Shielding of SEDINE at LSM. The sphere is placed within the copper and lead castle [30].

shield, it stops all γ rays, aside from those coming from ^{208}Tl and ^{40}K .

- The third, inner layer is composed of copper sheets, 5 cm wide. It was added to stop the radiation coming from the lead shielding, namely γ rays from ^{210}Pb , and electrons generated via Bremsstrahlung by ^{210}Bi .

Finally, the detector was installed in the previously mentioned Laboratoire Souterrain de Modane, in the border between France and Italy, under the Frejus mountain. The large rock overhead protects against cosmic rays: high energy particles coming from space that create a shower of secondary particles when they interact with Earth's atmosphere. At ground level, this mostly manifests as a high rate of so-called "cosmic muons". Their energy is too high to be stopped by regular shields, so the only protection against them is to place our detector underground. Fig. 3.9 shows a comparison of the muon rate at different underground laboratories.

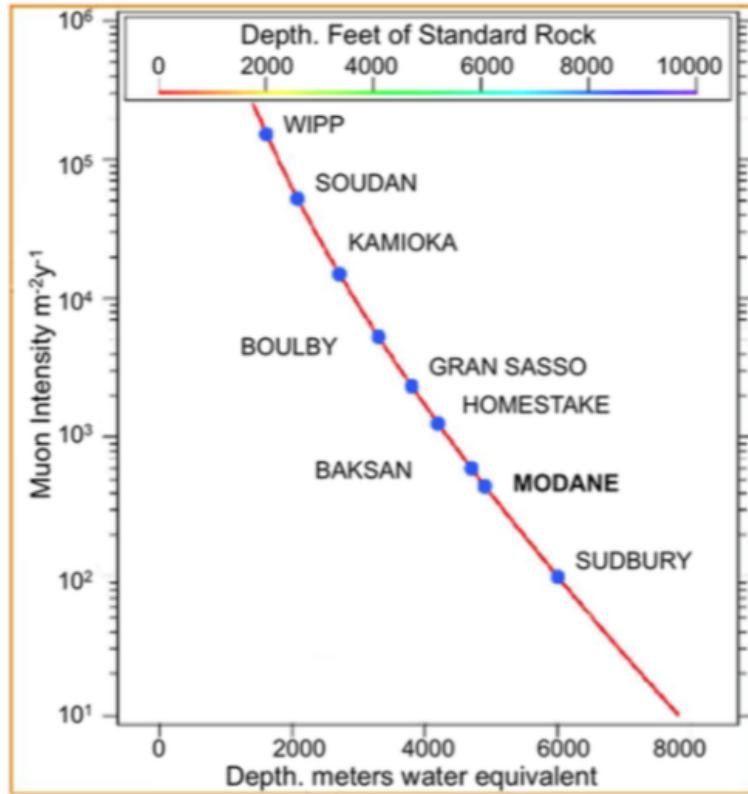


Figure 3.9: Muon flux in different laboratories, depending on their depth. LSM is at Modane, SNOLAB is at Sudbury. [30]

3.3.2 Backgrounds

Sources

Background radiation is a limiting factor for all rare event detection experiments. Because we are looking for a very small excess of events over our background, the statistics of any background that we cannot discriminate against could hide our signal. As such, it is of utmost importance to understand the origin of all backgrounds we are subject to, the event rate they might result in, and the ways in which we can limit them. We will give a very brief overview of the background sources we typically have to contend with.

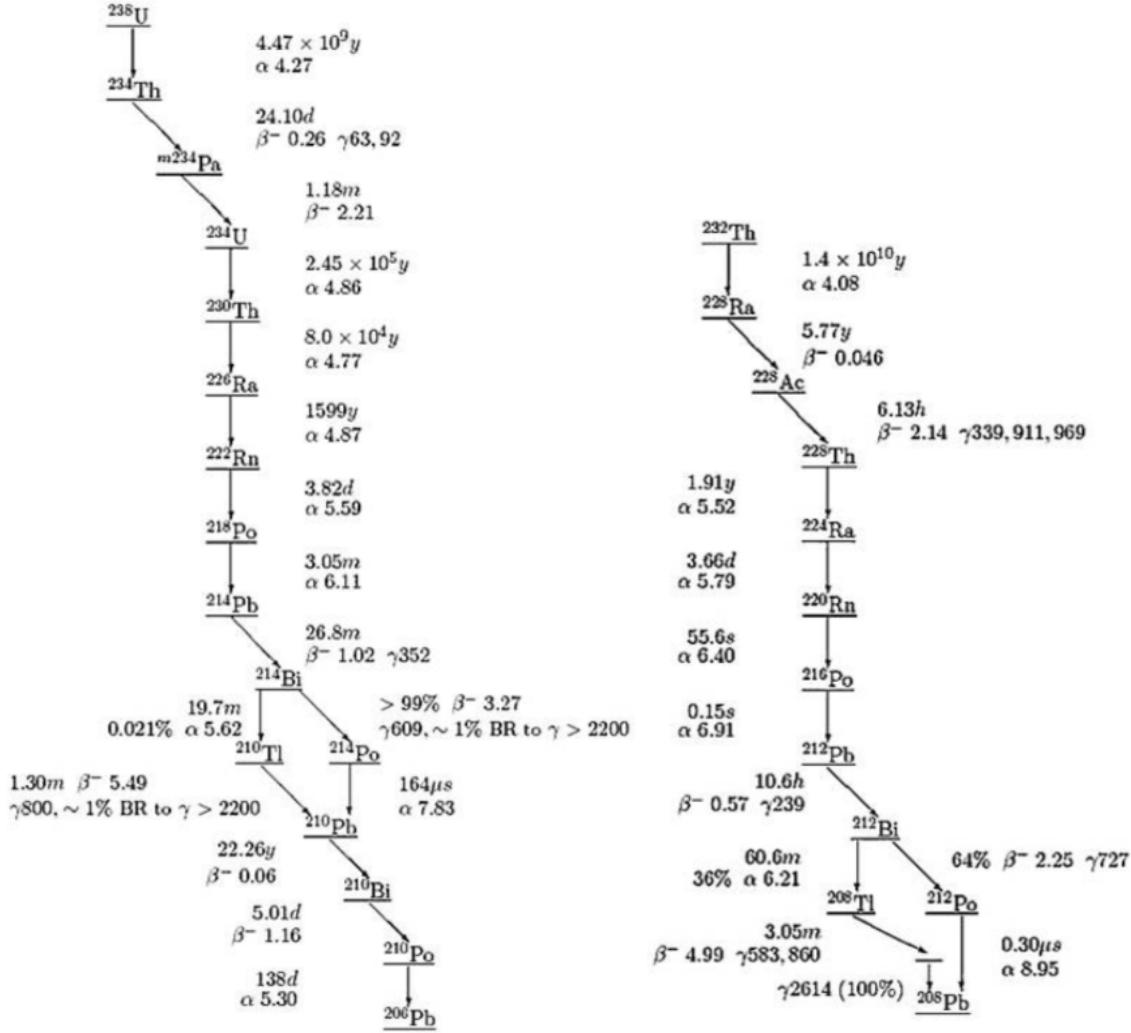


Figure 3.10: Decay chains from ^{238}U and ^{232}Th . Energies listed are in MeV [43].

One of the main components of cosmic radiation is the flux of high energy particles (up to 10^{20} eV), mainly protons (87%) and alpha particles (12%), that flow through the interstellar vacuum. These particles, called primary, interact with our atmosphere, creating electromagnetic and hadronic jets. While the majority of the secondary particles from these jets do not reach the Earth surface, muons ($\sim 75\%$) do, for a

flux of about $5.106 \mu/\text{m}^2/\text{day}$ [44]. The most common way to shield any experiment from these muons is to place it underground, where their flux is much lower, as seen on Fig. 3.9.

Furthermore, cosmic radiation can activate radioactive isotopes in our copper, defeating the purpose of using extremely pure copper for our detector. The main one is ^{60}Co , which can be activated from the copper in our detector while it is transported from its production plant to the underground laboratory⁹. ^{60}Co has a half life of around five years, so once created, it will contaminate the whole lifetime of our experiment. It decays into ^{60}N in an excited state, which in turn emits two γ rays at 1.17 and 1.33 MeV. As such, it is very important to limit the time any part of our detector spends above ground, but especially the sphere itself, because of its larger mass.

A side effect from setting the experiment underground is the increased radioactivity from the surrounding rock. The main background sources are ^{40}K and ^{208}Tl present in the rock walls and the concrete covering them. Both emit high energy γ rays, at 1.46 and 2.61 MeV respectively. This is however a much smaller problem than cosmic radiation itself.

The other main source of background radiation are the Uranium and Thorium decay chains. ^{238}U and ^{232}Th can be found in pretty much all materials in varying concentrations, in particular in the copper of our detector (minimized by using only extremely pure copper) or its shielding. All together, they emit α , β and γ radiation. See Fig. 3.10 for the complete decay chains.

Furthermore, there is one particularity with ^{210}Pb from the Uranium chain. ^{222}Rn

⁹Cosmic-induced neutrons produce radioactive cobalt from copper through the following reaction:
 $^{63}\text{Cu} + n \rightarrow ^{60}\text{Co} + \alpha$.

is a gas, which means it can be present anywhere, and in particular, in the air. It decays into ^{218}Po , ^{214}Pb , ^{214}Bi (all with very short lifetimes), and finally into ^{210}Pb , which has a half life of 22 years. This means that ^{210}Pb can attach itself to all surfaces, and once there, stay indefinitely. ^{210}Pb decays into ^{210}Bi in an excited state, which in turn emits $X -$ rays and Auger electrons, producing low energy events . As such, it is important to reach a very high vacuum before filling the detector with gas, and regularly clean the inside to remove any remaining traces of ^{210}Pb . Finally, the lead shielding is itself radioactive, so, in the absence of low-radioactivity archaeological lead, some extra shielding has to be placed in between it and the detector proper.

Expectations for SEDINE

The expected rate of background events was determined with a combination of calibrations and Geant4 simulations by A. Brossard, with results shown in Tab. 3.1.

The table shows all combined backgrounds, but only the rate of events reconstructed as pointlike and in the bulk of the gas. However, for solar KK axion searches, we are only interested in events that deposit energy in distinct locations within around a microsecond. For our detector, this generally involves photons, although not always exclusively. Photons of a few keV and higher can travel far enough from the decay site to be clearly distinct from any other particles generated there; unlike electrons, they do not leave a track of ionized atoms in their wake. Neutrons could potentially be another source of such interactions, but their rate is much lower than that of gammas.

The production of energy depositions at two distinct locations comes from either a single gamma doing two Compton interactions in the gas, or two or more different particles interacting in the detector at different locations at once; in the latter case,

| | source | contamination / flux | Unit | Event rate [0.5;1] keV (events/keV/kg/day) | Event rate [1;5] keV (events/keV/kg/day) | Total rate (mHz) |
|------------------|-------------------------------------|----------------------|-------------------------------|--|--|------------------|
| Copper sphere | ^{210}Pb | 26 | mBq/kg | 1.03 | 2.0 | 0.13 |
| | ^{210}Bi | 26 | mBq/kg | 2.6 | 2.1 | 0.11 |
| | ^{238}U | < 16 | $\mu\text{Bq}/\text{kg}$ | 0.036 | 0.042 | 0.0037 |
| | ^{232}Th | < 12 | $\mu\text{Bq}/\text{kg}$ | 0.053 | 0.045 | 0.0038 |
| | ^{40}K | < 110 | $\mu\text{Bq}/\text{kg}$ | 0.023 | 0.019 | 0.0013 |
| | ^{60}Co | 41 | $\mu\text{Bq}/\text{kg}$ | 0.067 | 0.052 | 0.0051 |
| Copper Shielding | ^{57}Co | 23 | $\mu\text{Bq}/\text{kg}$ | 0.044 | 0.046 | 0.0026 |
| | ^{210}Pb chain | 26 | mBq/kg | 0.36 | 0.33 | 0.033 |
| | ^{238}U | < 16 | $\mu\text{Bq}/\text{kg}$ | 0.11 | 0.089 | 0.0096 |
| | ^{232}Th | < 12 | $\mu\text{Bq}/\text{kg}$ | 0.082 | 0.076 | 0.0088 |
| | ^{40}K | < 110 | $\mu\text{Bq}/\text{kg}$ | 0.038 | 0.035 | 0.0039 |
| | ^{60}Co | 25 | $\mu\text{Bq}/\text{kg}$ | 0.16 | 0.12 | 0.016 |
| Lead Shielding | ^{210}Pb | 37.4 | Bq/kg | 6.9 | 5.9 | 0.7 |
| | ^{238}U | 79 | $\mu\text{Bq}/\text{kg}$ | 0.034 | 0.023 | 0.0028 |
| | ^{232}Th | 9 | $\mu\text{Bq}/\text{kg}$ | 0.0047 | 0.0035 | 0.0008 |
| | ^{40}K | < 1.46 | mBq/kg | 0.0059 | 0.0048 | 0.0006 |
| | Gamma ray | 1460 keV | $\gamma/\text{cm}^2/\text{s}$ | 0.53 | 0.47 | 0.059 |
| | ^{2614}keV | 0.042 | $\gamma/\text{cm}^2/\text{s}$ | 1.22 | 0.97 | 0.10 |
| Inner Surface | ^{210}Pb | 1.8 | mBq | 3.42 | 0.77 | 0.039 |
| | ^{210}Bi | 1.8 | mBq | 0.44 | 0.11 | 0.0040 |
| | ^{210}Po | 1.8 | mBq | 0.22 | 0.014 | 0.0011 |
| | ^{214}Pb | 0.2 | mBq | 0.074 | 0.003 | 0.0044 |
| | $^{214}\text{Bi} + ^{214}\text{Po}$ | 0.2 | mBq | 0.05 | 0.02 | 0.0008 |
| | Total | | | 17.5 | 13.2 | 1.2 |

Table 3.1: Summary of the main volume events background of SEDINE based on Geant4 simulations [35].

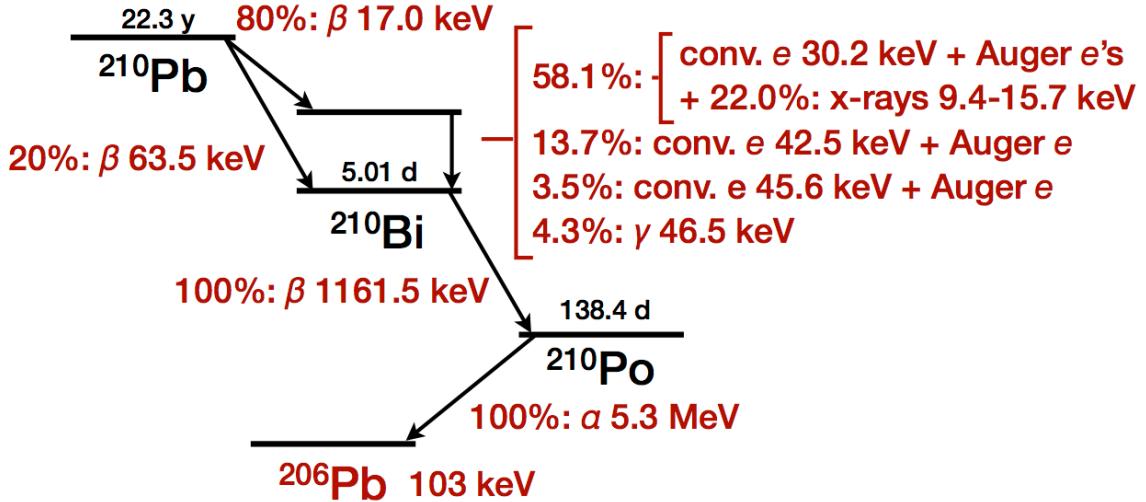


Figure 3.11: Decays of ^{210}Pb and ^{210}Bi [45].

usually at least one of the particles is a gamma. As such, for the search for KK axions with SEDINE, the main backgrounds were:

- ^{210}Pb on the inner surface of the detector. The excited state of ^{210}Bi releases both photons and electrons simultaneously in the energy range of interest for this search (see Fig. 3.11).
- ^{210}Bi in the bulk of the copper of the detector. Through Bremsstrahlung, the high energy electron released in the decay generates multiple gammas of high enough energy to perform Compton interactions, and may even enter the detector internal volume.
- ^{210}Bi in the lead shield. While they start farther away from the target gas than decays in the bulk of the detector shell, the larger mass of lead shield and its much worse radiopurity make up for the longer distance the Bremsstrahlung radiation has to cross.

More details on the nature of these backgrounds, and their effect on SEDINE data and its analysis, can be found in Sec. 8.1.2. We expect a rate around $1 - 10 \text{ evt/day}$ from these sources, after some unoptimized selection cuts.

Finally, it could be expected that random coincidences of two different events could also be a background for this search, but a quick computation shows this to not be the case. After some basic cuts, the physics data of SEDINE had a total event rate of around 20 mHz . Taking a relatively wide coincidence window of 1 ms , this gives an upper bound on the rate of random coincidences during the run of $\sim 0.035 \text{ evt/day}$. As such, this source of background can be safely ignored in this work.

3.4 NEWS-G at SNOLAB

The next step for the NEWS-G collaboration is the installation of a new detector at SNOLAB, nicknamed the SNOGLOBE (pending an official name). SNOLAB is one of the deepest low-background laboratories in the World (as seen on Fig. 3.9), which will lead to a reduced muon flux over that at LSM. It will also be larger than SEDINE, with a high purity copper (C10100) sphere 140 cm wide and 10 cm thick; a new kind of sensor electrode, achinos [34], was developed to accommodate for the larger size. An engineering drawing and picture of the real detector can be seen in Fig. 3.12. The internal $500 \mu\text{m}$ of the detector shell have been electroplated with pure copper to eliminate background from the ^{210}Pb contamination in the internal surface of the detector. The sphere will be enclosed in a 25 cm thick lead sphere, to protect against γ radiation; the internal 3 cm of this shield are made of archaeological lead, with a much lower presence of ^{210}Pb . Finally, 40 cm thick polyethylene walls protect against neutrons coming from the cavern walls. The detector and most of the

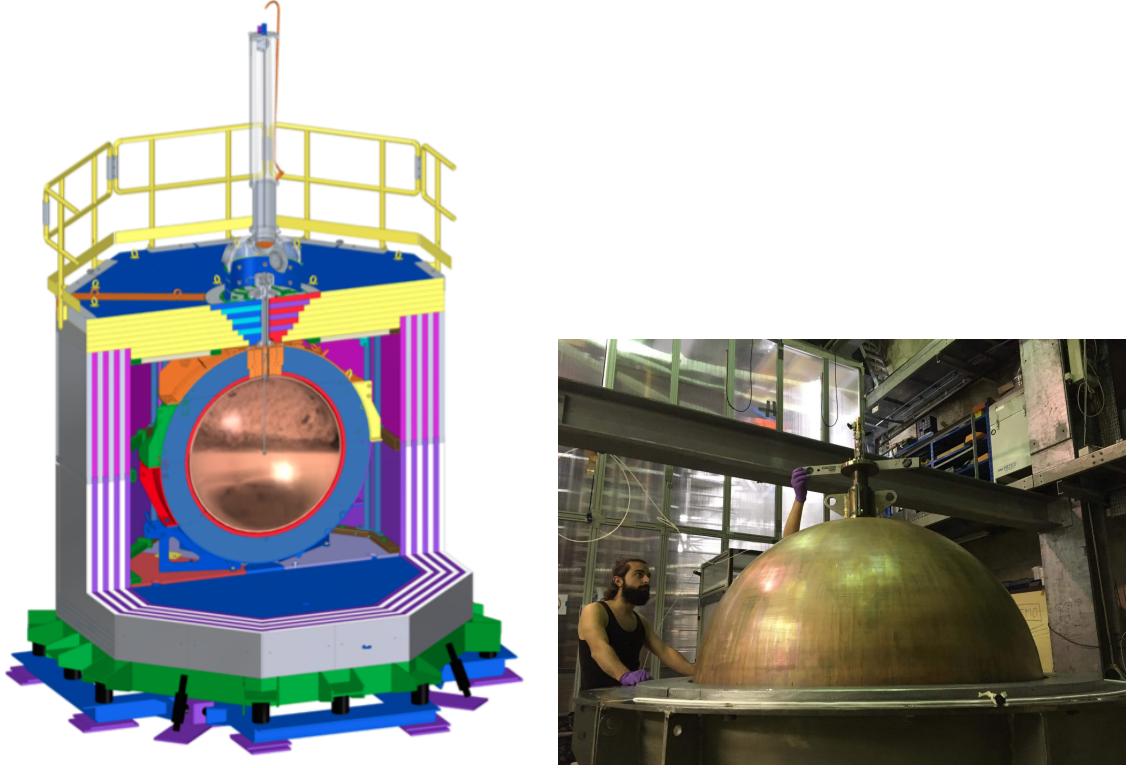


Figure 3.12: Left: Schematic view of SNOGLOBE with shielding. Right: Installation of SNOGLOBE inside the lead shield at LSM.

components of its shield is already waiting inside SNOLAB at the time of writing, awaiting relaxation of COVID-19 safety measures for installation, hopefully before the end of 2020. Papers on both the details of the detector setup and the electroplating procedure are currently in preparation.

Tab. 3.2 shows an early estimation of the nature and distribution of the backgrounds expected from the new version of the experiment, based on Geant4 simulations performed by A. Brossard. This table includes all energy depositions, not only the ones expected to be reconstructed as pointlike. The higher radiopurity of the lead shield explains the decrease in background from that source. Furthermore, due to the electroplating procedure, we expect almost no background from ^{210}Pb plated on the

internal surface of the copper shell. As such, radioactive backgrounds at low energies will be dominated by bremsstrahlung from the β decays of ^{210}Bi , and ^{60}Co from the cosmogenic activation of the copper, both from the copper shell of the detector. The latter should become sub-dominant after a few months to a year.

Before being shipped to SNOLAB, SNOGLOBE was temporarily installed in the LSM, without the polyethylene shield. Its time there was used to familiarize for Canadian collaborators to familiarize themselves with its setup before installation in the harder-to-access SNOLAB, and for running various tests and calibrations. The nature of the data taken with the new achinos sensor is still in the process of being understood, so no result from that campaign was used for the work in this thesis. Still, a discussion on the projected performance of this detector for the search of solar KK axions can be found in Sec. 8.3.

Summary

Spherical Proportional Counters are gas detectors which use a central electrode at a high voltage to collect the ionization signal left by particles interacting with the gas. The use of a low density target mass means an axion decaying in the volume produces two axions that are absorbed at distinct locations, unlike for rare event detectors that use liquids or crystals. Given the increase in the electron drift time with its radial position, axion decays appear as two pulses arriving shortly after each other. The specificity of such a signal allows for rejection of all background events that leave energy either at a single point or as a track inside the detector.

SEDINE is a low-background, 60 cm wide SPC, running at the underground lab of LSM. The data taken with this detector was the one used to search for solar KK

| | Source | Contamination / flux | Unit | Event rate <1 keV [dru] | Event rate in [1:5] keV [dru] | Total rate [mHz] |
|--|-------------------|-----------------------|----------------------------------|-------------------------|-------------------------------|------------------------|
| Gas mixture | ^{3}H | 13 | $\mu\text{Bq}/\text{kg}$ | 0.05 | 0.06 | 0.005 |
| | ^{222}Rn | 111 | $\mu\text{Bq}/\text{kg}$ | 0.05 | 0.04 | 0.2 |
| Copper sphere 500 μm electrolyte | ^{210}Pb | 28.5 | mBq/kg | 1.04 | 1.01 | 0.86 |
| | ^{238}U | 3 | $\mu\text{Bq}/\text{kg}$ | 0.0117 | 0.115 | 0.028 |
| | ^{232}Th | 13 | $\mu\text{Bq}/\text{kg}$ | 0.0754 | 0.0692 | 0.163 |
| Roman lead | ^{40}K | 0.1 | mBq/kg | 0.0157 | 0.0186 | 0.0622 |
| | ^{210}Pb | <25 | mBq/kg | <0.14 | <0.12 | 0.057 |
| | ^{238}U | 44.5 | $\mu\text{Bq}/\text{kg}$ | 0.142 | 0.094 | 0.277 |
| Low activity lead | ^{232}Th | 9.1 | $\mu\text{Bq}/\text{kg}$ | 0.0256 | 0.0161 | 0.0577 |
| | ^{40}K | <1.3 | mBq/kg | <0.28 | 0.23 | 0.65 |
| | ^{210}Pb | 4.6 | Bq/kg | 0.053 | 0.055 | 0.17 |
| | ^{238}U | 79 | $\mu\text{Bq}/\text{kg}$ | 0.17 | 0.132 | 0.5 |
| | ^{232}Th | 9 | $\mu\text{Bq}/\text{kg}$ | 0.0251 | 0.0201 | 0.075 |
| | ^{40}K | <1.46 | mBq/kg | <0.35 | 0.26 | 0.67 |
| Cavern | Gamma | 4.87×10^{-8} | $\gamma/\text{cm}^2/\text{s}$ | 0.0084 | 0.0095 | 0.00464 |
| | Neutron | 4000 | neutron/ m^2/day | 0.0044 | 0.0004 | 3.54×10^{-11} |
| | Muon | 0.27 | muon/ m^2/day | 0.00062 | 0.00044 | 5.04×10^{-8} |
| Total | | | | 1.67 | 1.54 | 2.4 |
| Total + cosmogenic activation of the copper sphere | | | | 5.20 | 5.20 | 5.4 |
| Total + cosmogenic activation of the copper sphere and 6 months of cooling | | | | 2.8 | 2.5 | 3.4 |
| Total + cosmogenic activation of the copper sphere and 1 years of cooling | | | | 2.1 | 1.9 | 3.0 |
| Total + cosmogenic activation of the copper sphere and 2 years of cooling | | | | 1.9 | 1.7 | 2.9 |

Table 3.2: Summary of the main backgrounds of NEWS-G at SNOLAB, without rise time selection, based on Geant4 simulations. The upper limits of activities in the lead are not taken into account in the total [35].

axions, as will be covered in upcoming chapters. The main sources of radioactive background are the ^{210}Pb contamination in the copper shell of the detector and in the lead shield, producing around $\sim 5 \text{ evt/keV/kg/day}$ and $\sim 6 \text{ evt/keV/kg/day}$ pointlike events in the $1 - 5 \text{ keV}$ energy range, respectively. The upcoming detector at SNOLAB, 140 cm wide, made with more radiopure materials and electroplated, should only see $\sim 1 \text{ evt/keV/kg/day}$ and $\sim 0.05 \text{ evt/keV/kg/day}$ events in that range, from those same sources.

In the next chapter on pulse processing, we will be using our knowledge of the detector's working principle to extract all relevant information from the events recorded by an SPC.

Chapter 4

Data processing and Analysis

The objective of automatic pulse processing is twofold:

- Have a standard, reproducible, quantifiable way to identify the characteristics (energy, risetime...) of any event;
- Be able to process large amounts of data quickly.

The latter is important due to the size of the datasets we need to manipulate. The former allows us to do the actual physics we are interested in, and potentially compare the performance of different algorithms to each other, and to minimum benchmarks.

Simulations will be mentioned multiple times in this chapter. Systematic uncertainties in our calibration data can make it difficult to characterize the behaviour of the different algorithms, and hence predict their performance for situations in which no calibration data is available. By accounting for those uncertainties in the simulations, the performance of the different algorithms can be assessed with the necessary precision. The process to create these simulations is described in detail in chapter 5.

4.1 Amplitude and Risetime estimation

After recording an event, we need to be able to extract useful information from it. The main parameter we are interested in is the energy of the event. While we do not generally have direct access to that information, the signal coming out of the preamplifier is directly proportional to the total charge created during the avalanche by an energy deposition. Since the average charge produced is proportional to the energy deposited, the “height”, or amplitude, of an event is a good estimator for its energy. Knowing the energy distribution of our data will then allow us a comparison with known radioactive backgrounds, or with any potential particle signal.

The second parameter is what we normally call the “risetime”, or the time it takes for the pulse to go from the “baseline” to its “maximum”¹. Its a measure of the spread in arrival times of primary electrons, which in turn depends on the position of the energy deposition: the farther away from the central electrode, the larger their dispersion, and hence the longer the risetime. This is useful to distinguish events coming from the bulk of the gas detector, from the surface of the outer shell, or leaving energy depositions in multiple spots at once.

This section will cover three different algorithms for computing these estimators, how to use those estimators to discriminate between different type of events, and their relative performance.

¹Those terms are not well defined. In practice, the risetime is the time between the pulse reaching some low percentage (e.g. 10%) and some high percentage (e.g. 90%) of its maximum amplitude. The wider the range, the stronger the discrimination power of the risetime, but the more sensitive to noise it becomes, so the choice will depend on the type of processing performed on the pulse beforehand.

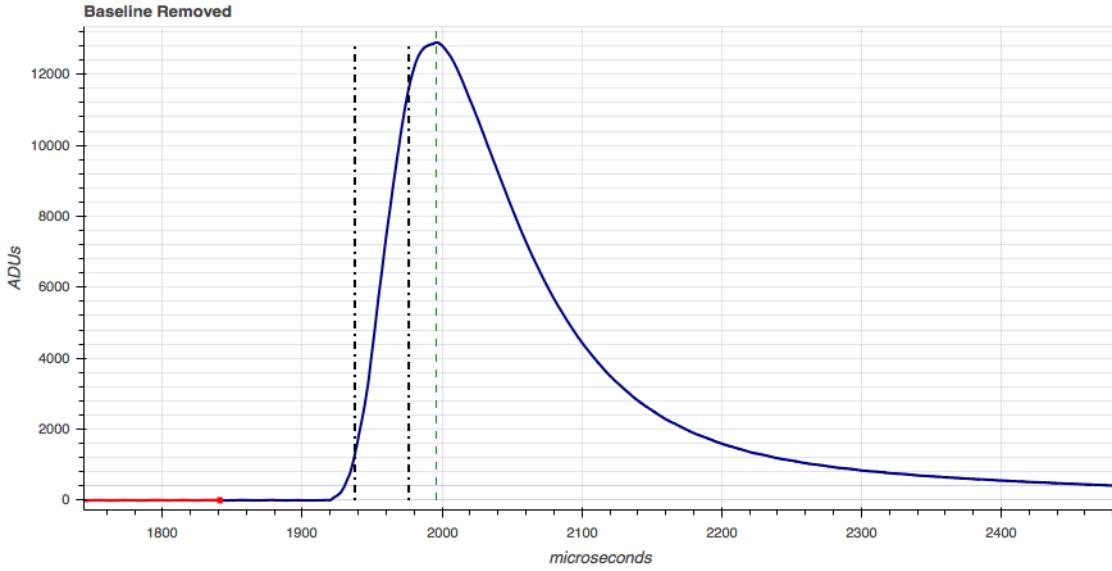


Figure 4.1: SAMBA processing on a pulse. In red, computed average baseline before the beginning of the pulse. The green vertical line points the position of the maximum of the pulse. The amplitude is the difference between the value at that point and the baseline; here, 12860 ADU. The two dashed black lines point the time at which the pulse reaches 10% and when it reaches 90% of the amplitude. The risetime is the difference between the two; here, $38\ \mu\text{s}$

4.1.1 SAMBA

The first method used to compute the amplitude and risetime of our events is the one implemented by our data acquisition software, SAMBA, originally designed for EDELWEISS [46]. The amplitude is calculated as the difference between the maximum of the pulse, and the mean of the baseline before the beginning of the pulse. The risetime is calculated as the difference between the time when the pulse reaches 10% and the time when it reaches 90% of the amplitude. An example is shown in Fig. 4.1.

The advantage of this algorithm is its simplicity, allowing for simultaneous computation while recording the data. This can be exploited to, for example, set extra

constraints for the trigger, rejecting events with risetimes that are too short or too long. Even more importantly, it gives instant feedback to anyone operating the detector.

However, the main disadvantage of this method is that its amplitude estimation is biased. When we use a preamplifier (cf. App. B) with a decaytime that is too short with respect to the event length, then the signal produced will start decaying noticeably before all the charge has been collected. This means that, for two given events with the same energy, the one with the longer risetime will appear with a lower amplitude than the one with the shorter risetime, as pictured in Fig 4.2. We see that in the case with infinite decaytime, all signals have very similar amplitudes, but in the case with a short decaytime, the difference in amplitude is of more than 20%. The relative difference in amplitude between the scenario with an infinite decaytime and the real scenario is called the *ballistic deficit*. Note that the ballistic deficit is not a problem *per se*, but its value changing depending on the risetime is.

For this reason, alternate processing algorithms were developed to correct this effect. The two methods described below solve this problem. Since they both involve a deconvolution step, we will describe that first.

4.1.2 Deconvolution

Let us introduce a few definitions that will be useful in what will follow.

The transfer function $R(t)$ of a system is the output of that system when the input is a Dirac delta-function (henceforth, “dirac”). Since we can decompose any function $f(t)$ into a (continuous) “sum” of time-displaced diracs, then the output of our system for any arbitrary input is just a (continuous) “sum” of shifted transfer

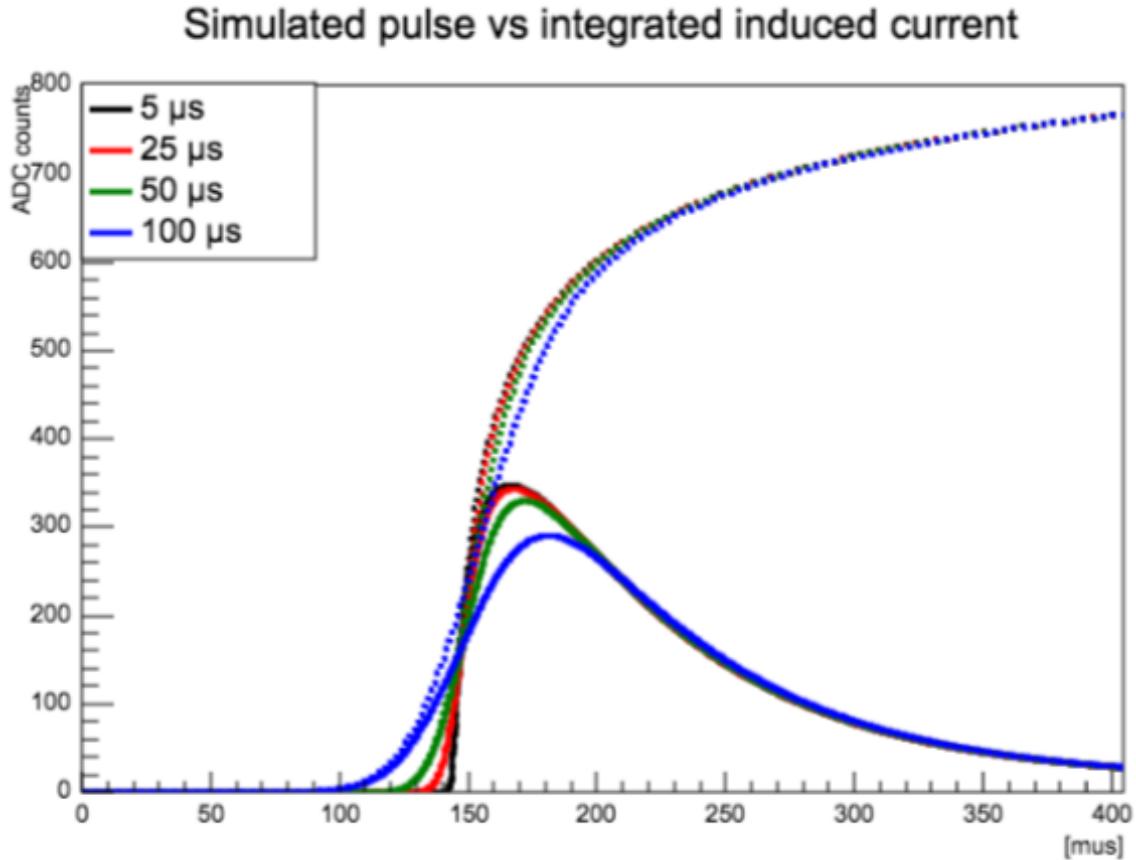


Figure 4.2: Output of a theoretical preamplifier with infinite decaytime (dashed lines) and another with a decaytime of $50\ \mu s$ (solid lines), for events with same energy but different risetimes. For the real preamplifier, the amplitude of the pulse becomes lower as the risetime increases.

functions, scaled by the amplitude of each dirac. This is called a convolution:

$$f * R(t) = \int f(\tau)R(t - \tau)d\tau \quad (4.1)$$

where $f * R$, the convolution of f (input to the system) and R (system response function), is the output of our system. Since we are working with discrete time

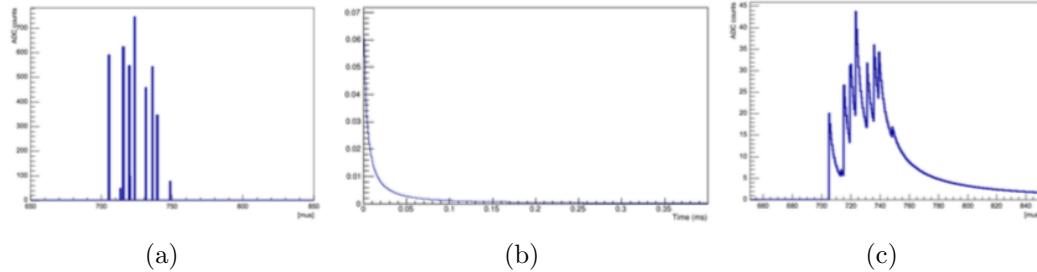


Figure 4.3: Example of convolution, here with discrete time. For each of the peaks in the first plot, we add a transfer function starting at that time, multiplied by the height of the peak.

windows, this becomes

$$f * R(t_i) = \sum_{j=0}^i f(t_j)R(t_i - t_j) \quad (4.2)$$

where f and R are defined from t_0 to some t_n , and $t_i = i \Delta t$, where Δt is our sampling period (we will be using $\Delta t = 0.48\mu s$ as an example). You can see an example of convolution in Fig. 4.3, with a function f composed of discrete diracs, a transfer function R (based on the induced current from Sec. C.0.2, and the output from the convolution of the two.

An important consequence of the definition is that convolution is linear, i.e., $(f + g) * R = f * R + g * R$. The amplitude of a pulse is thus proportional to the amplitude of the original signal, and the combination of two signals can be separated without worry.

While convolutions are easy enough to do numerically, and can sometimes be done analytically, deconvolutions are trickier. As the name implies, a deconvolution is the process of obtaining f , when you know R and $f * R$. While there is no general analytical method to solve the problem, one way to solve it numerically is by using

the convolution theorem, which states that convolution transforms into product in Fourier space, i.e.,

$$\widetilde{f * g} = \tilde{f} \cdot \tilde{g} \quad (4.3)$$

where \tilde{f} denotes the Fourier transform of f , and f and g are two arbitrary functions. This gives us an easy way to deconvolve a transfer function, just by doing a division in Fourier space:

$$\tilde{f} = \widetilde{f * R} / \tilde{R} \quad (4.4)$$

and then we just have to take the inverse Fourier transform of \tilde{f} . Of course, this is only the most basic approach on the question of deconvolution. The topic of deconvolving transfer functions is a well-established and extremely active research area, especially in image processing. An approachable review of techniques and pitfalls can be found in chapter 5 of [47].

Coming back to the processing of our data, if our model for signal formation is correct, we can deconvolve the resulting response function from our recorded pulses. After integration, the resulting signal has an amplitude proportional to the total charge, but is now unaffected by the ballistic deficit.

4.1.3 Modified decaytime deconvolution

Before work started on this thesis, G. Gerbier had already been providing a secondary analysis of our data. As previously mentioned, the original motivation was to remove the bias due to the ballistic deficit from our amplitude estimation. Fig. 4.2 shows what the pulse would look like if we had a preamplifier with an infinite decaytime.

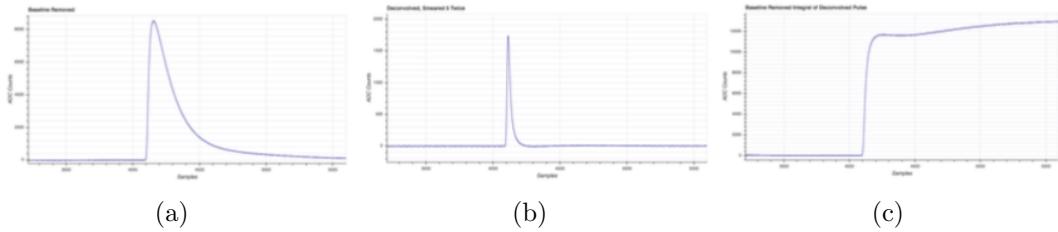


Figure 4.4: MDec method applied to a high energy pulse (~ 10 keV) from SEDINE’s physics run. Left: Raw pulse. Middle: Deconvolved pulse. Right: Integrated deconvolved pulse. We note that the integration of the deconvolved pulse does not quite go back to a flat baseline, but instead makes a small dip before slowly increasing again. Taking the height before the dip offers a good estimator of the amplitude of the pulse.

This can be reproduced digitally by deconvolving the response of the preamplifier (a decaying exponential), then integrating the resulting signal. The final pulse now has an amplitude that is independent of the spread in arrival time of primary electrons.

However, the amplitude keeps rising for some time after the arrival of the last electron, which means that the estimation of the amplitude is now dependent on the time at which the height of the pulse is calculated (see Fig. 4.4c.). This downside was solved by G. Gerbier by using a modified, longer decaytime for the deconvolution than the real time constant of the preamplifier. Tuning this ad hoc decaytime stabilizes the pulse and removes the long tail characteristic of the ion-induced current. This allows for a consistent computation of the amplitude, much less sensitive to the exact position at which it is calculated. The method will be referred to as the Modified Decaytime method (MDec) from now on.

An example of the result is shown on Fig. 4.4. For context, this was performed on a relatively high energy (~ 10 keV) event from SEDINE’s physics run, where the decaytime of the preamplifier was $47\ \mu\text{s}$; the modified decaytime used in the deconvolution was $83\ \mu\text{s}$ instead. The result is a much flatter baseline after the pulse, compared to

what a “proper” deconvolution of the preamplifier response would produce.

The risetime, in turn, is computed as the time difference between the point at which the pulse reaches 10% and the 75% (rather than the previous 10% and the 90%) of the amplitude. This change is due to the “roundedness” this process generates at the end of the pulse. The less sharp a feature is, the more sensitive its timing information is to noise oscillations. The tapering off of the top of the pulse, compared to the raw one, would decrease the resolution of our risetime estimation if we were to still use the 90% point for its computation. Note that this change doesn’t come without a drawback. The risetime of a pulse is only a stand-in for the actual dispersion in arrival time of the primary electrons, and the quality of this estimator tends to improve (in the absence of noise) with the number of primary electrons “englobed” by the end points of the risetime.

While the MDec method is fundamentally an ad hoc modification for the purpose of improving on the processing of the raw pulse, theoretical justification can be found in the parametrization of the signal mentioned in Eq. C.13 in Sec. C.0.2, $S(t) = -q_{ions}k(e^{-at} - e^{-bt})$. In fact, it does produce very satisfactory results, especially for amplitude resolution. The disadvantages of the method are the need to tune τ over a potentially wide range whenever we change the operating conditions of our experiment, a relatively poor risetime discrimination for short events, and the tapering off and non-flat baseline at the end and after the pulse, respectively.

One last advantage of this method is that it does not require going through Fourier space, avoiding problems such as edge effects. In discrete time with sampling period Δt , the recursive form of the convolution $a[i]$ of a signal $b[i]$ with a decaying exponential with time constant T follows the formula:

$$a[i] = a[i-1] e^{-1/T_{sp}} + (N \Delta t) b[i]$$

Where N is a normalization factor given by the height of the response to an impulse of integral 1, and $T_{sp} = T/\Delta t$ is the time constant of the decay, in units of the sampling period. If we want the convolution to preserve the integral of the signal, then $N = (1 - e^{-1/T_{sp}})/\Delta t \sim 1/T$, for $T_{sp} \gg 1$. More importantly, this relation is easy to invert, giving:

$$b[i] = \frac{1}{N \Delta t} (a[i] - a[i-1] e^{-1/T_{sp}})$$

Not only do we not need to go through Fourier space, but the existence of a linear formula involving only two points of the final signal allows for very fast deconvolution. This is even potentially useable at the digitizer level, even though it has not been implemented there as of writing.²

4.1.4 Double Deconvolution method

The second deconvolution method was developed in the context of this thesis. Based on the signal formation mechanism discussed in chapter 3, we can deconvolve the full response function of the detector. Ideally, this would result in a current signal consisting of short impulses, corresponding to the arrival of primary electrons to the avalanche region, with the height of each impulse being proportional to the number of secondary charges created by each primary electron. After integration, this gives a series of step functions, whose total height is proportional to the total number of

²Note that, for $T_{sp} \gg 1$, this formula can be rewritten as $b[i] = \frac{1}{N} \left(\frac{a[i] - a[i-1]}{\Delta t} \right) + a[i-1]/T$. This is reminiscent to the equivalent deconvolution formula in continuous time, $b(t) = \frac{d}{dt} a(t) + a(t)/T$, which can be found through judicious use of the Laplace transform, as shown in [48].

secondary charges produced in the avalanche region during the event.

In practice, differences between the ideal response function and the real one create artificial current tails or negative current spikes (we refer to the latter as undershoots). Furthermore, to moderate the amplification of noise due to the deconvolution process, the data has to be smoothed. This widens the recovered current impulses produced by primary electrons, and in general makes them impossible to distinguish from one another. However, the amplitude of the total integral is not affected³.

This method was called the Double Deconvolution (DD) method due to its implementation: we deconvolve the response function of the preamplifier, a decaying exponential, with the numerical algorithm described in the previous subsection; then, since there is no equivalent for the deconvolution of the ion-induced current formula, we do that deconvolution by division in Fourier space.

Substantial work was done to integrate this method into the analysis suite developed for the collaboration, and several collaborators now use this algorithm and its associated visualization tools.

Method validation

We apply this procedure to pulses recorded during SEDINE's physics run previously mentioned. The run was taken under the following conditions: applied voltage of $V_0 = 2520$ V, with a Neon and CH₄ (0.7%) gas mixture at $P = 3.1$ bar, using the Canberra preamplifier [49]. Together with the dimensions of the sphere ($r_2 = 30$ cm) and electrode ($r_1 = 0.315$ cm), this means that the only parameter we haven't fixed to generate the response function of the detector is the ion mobility of Neon in our gas

³The smoothing algorithm used is a simple running average applied twice, but this would remain true for any smoothing algorithm that can be described as a convolution with a transfer function that is zero-valued outside a finite window.

mixture. We can start with $\mu_0 \sim 6 \text{ cm}^2\text{V}^{-1}\text{s}^{-1}$, the value for Neon ions in pure Neon at normal temperature and pressure [50]. A first validation of the method comes from deconvolving the transfer function of the preamplifier to so-called “electronic” events, non-physical signals that are created in the electronics of the detector, rather than in the detector itself, and as such have a different shape than “physical” events. As seen in Fig. 4.5, the recovery of a peaked event⁴ from this process, that goes back to a flat baseline nearly instantly and without undershoot, proves that:

- The deconvolution method is adapted to our problem.
- We have the correct transfer function for our preamplifier.
- This event is indeed purely electronic in nature, since we don’t observe the characteristic shape of the ion induced current.

A second validation comes from looking at events formed by an energy deposition on our target mass. These are what we call “physical” events, and we select them based on preliminary parameters extracted from the shape of the raw pulse (see Sec. 4.1.6). We had to tune the value of the ion mobility manually to get optimal results, from the expected $\mu_0 \sim 6 \text{ cm}^2\text{V}^{-1}\text{s}^{-1}$, to $\mu'_0 = 7.5 \text{ cm}^2\text{V}^{-1}\text{s}^{-1}$. The value of μ_0 was selected so that the majority of the physical pulses would exhibit neither a long tail (characteristic of the current induced by the drifting ions) nor an undershoot. While the optimal value was not exactly the same for all pulses, the range of possibilities was within 5% of the previously given value⁵.

⁴The pulse would be even more peaked if we didn’t smooth it out, but then noise might obscure the second peak.

⁵This value was originally found by visual inspection of a few dozen values, with noticeable tails or undershoots appearing when changing the value of μ_0 by more than $\sim 0.6 \text{ cm}^2\text{V}^{-1}\text{s}^{-1}$. An automated way to find the “optimal” value is described in Sec. 6.6.

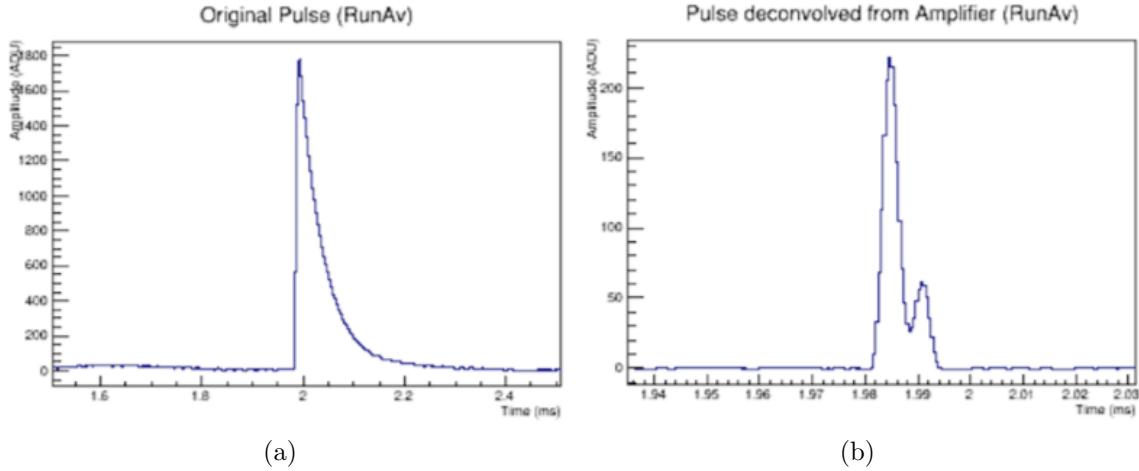


Figure 4.5: Example of electronic event. Left figure: raw pulse as recorded by our detector. Right figure: same pulse, zoomed in, after deconvolution of the transfer function of the preamplifier. Both are smoothed out to remove some of the noise. Note the different time scales: the width of the left pulse is ~ 0.2 ms, while the right one is only ~ 0.015 ms.

The disparity in values between events is not well understood. A possible explanation is that this method is absorbing variations of the electric field into the ion mobility: events happening close to the rod will experience a lower field, but since we assume the field to be isotropic (for the sake of the response function used in the deconvolution), this appears as a lower ion mobility instead. As for the difference in the average value for the ion mobility with respect to the literature, this can be a combination of the previous effect, together with the presence of CH_4 in our gas mixture, since ion mobility is very sensitive to gas composition. Still, it is fairly remarkable that the optimal value of μ_0 found this way is so close to the values in the literature, despite using a highly variable electric field instead of a constant one, and doing the measure via the shape of the induced current instead of the ion drift time. The optimal value is also relatively constant from run to run. A more rigorous

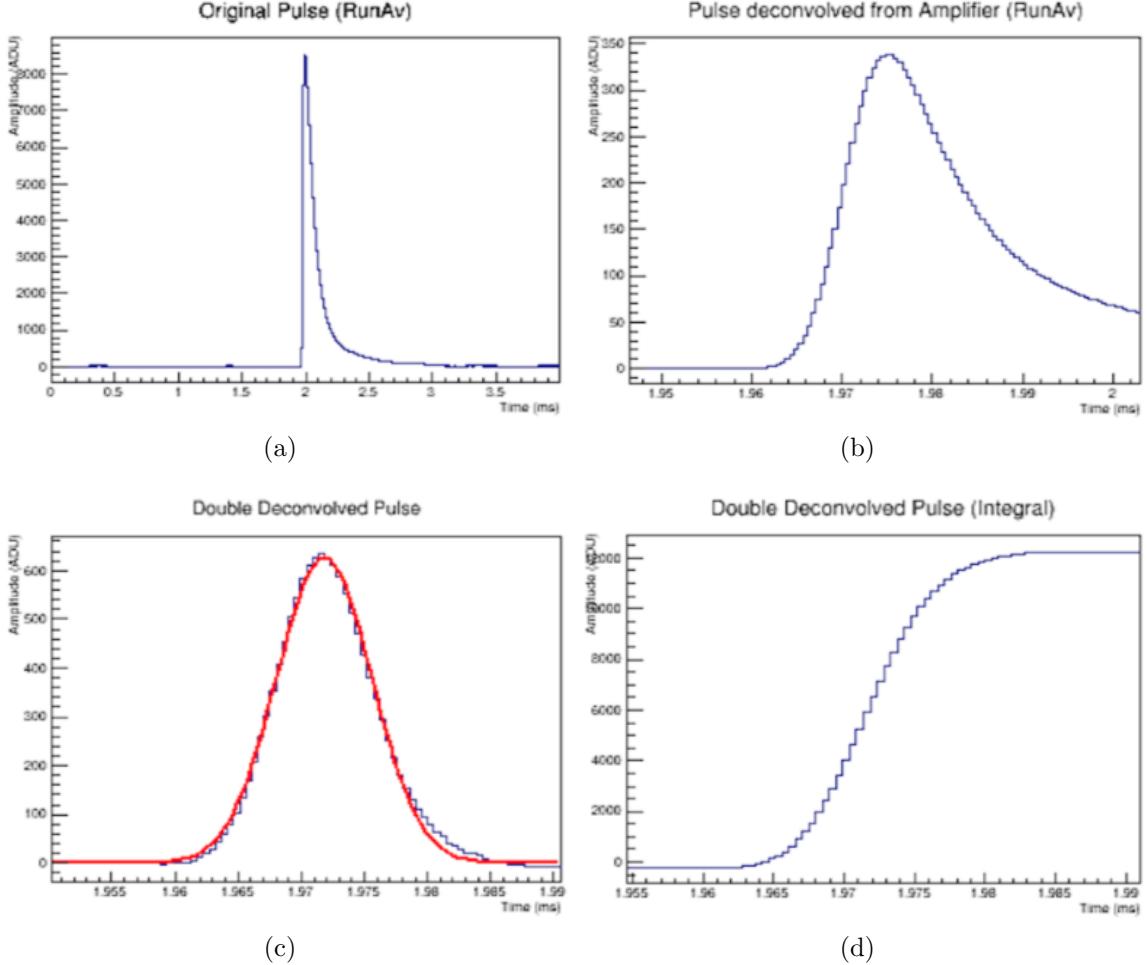


Figure 4.6: Same example of a high energy (~ 10 keV) physical event as in Fig. 4.4. Top left: raw pulse as recorded by our detector. Top right: same pulse, zoomed in, after deconvolution of the transfer function of the preamplifier. Bottom left: same pulse, after also deconvolving the current induced by a drifting ion, with gaussian fit superimposed. Bottom right: integral of the previous pulse. For pointlike events with a sufficiently high number of primary electrons, we do expect them to reach the electrode forming a gaussian distribution in arrival time.

calibration of the ion mobility is described in Sec. 6.6.

As shown in Fig. 4.6, we can recover a signal that is very well localized in time ($\sim 20 \mu s$, compared to $\sim 500 \mu s$ for the raw pulse), without affecting the flat baseline before or after the pulse. This is fairly strong evidence not only that the method works and that we can use it to extract the parameters of any given physical pulse, but also that our understanding of the physics of the detector from which we derived the ion induced current is sound. This works even for lower energies and longer events, as shown in Fig. 4.7.

4.1.5 Performance

Let's compare the performance of the three previously mentioned methods: SAMBA, MDec, and DD estimators.

Ballistic deficit

As mentioned earlier, SAMBA estimators are affected by varying degrees of ballistic deficit, biasing its amplitude estimator. This effect is easy to see when looking at plots of the risetime vs energy distribution of events for a given run. For example, we should see vertical lines for ^{210}Po alphas at 5.3 MeV. However, as seen in Fig. 4.8, these lines are instead slanted towards the left for the SAMBA estimators, because of the ballistic deficit. We can correct for that effect with either of the previously exposed methods. Then we obtain a ^{210}Po line that is now mostly vertical, with a slight slant towards the right instead. As we will cover in Sec. 6.5, we assume that this residual dependency is likely due to electron attachment, although whether it is

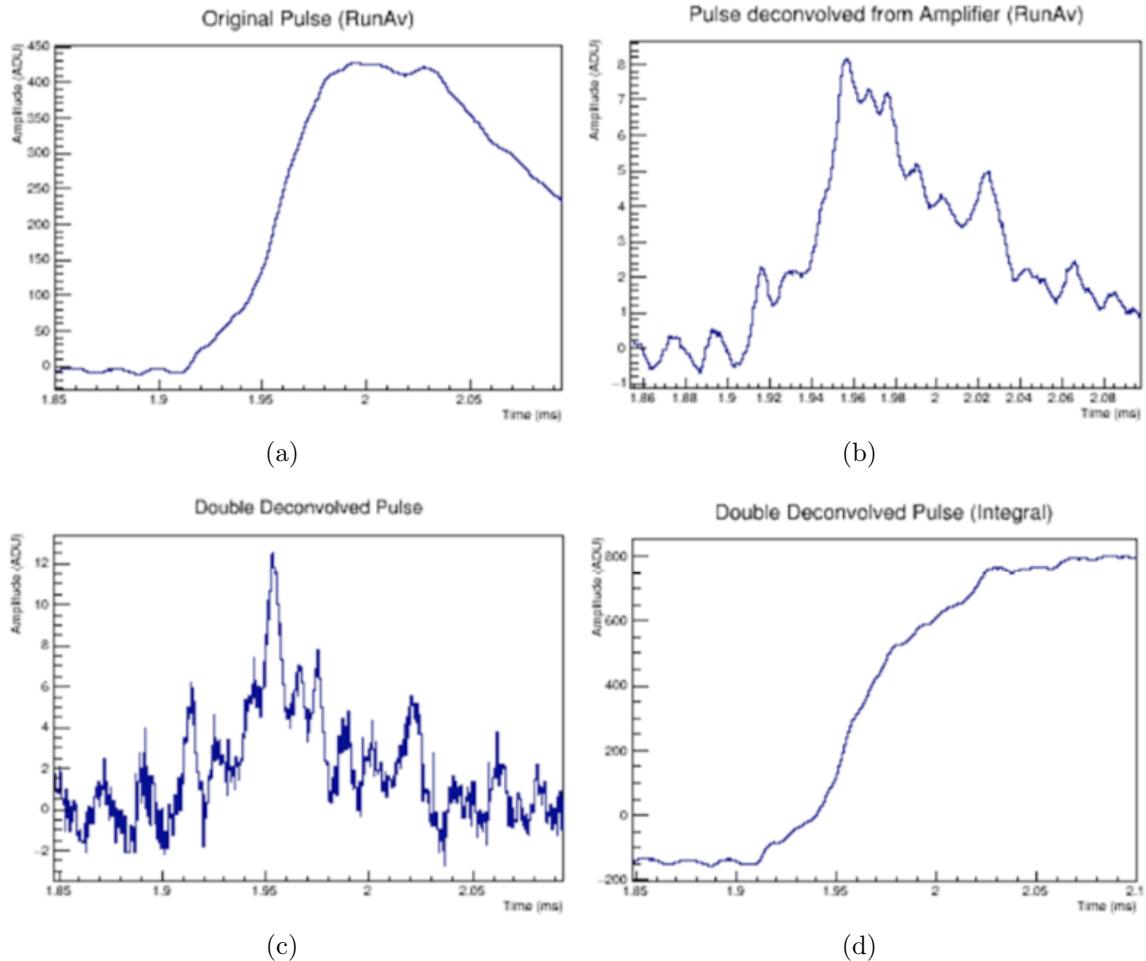


Figure 4.7: Example of a lower energy (~ 600 eV) physical event. The noise that survives the running average is similar to the shape of an avalanche, and makes it hard to tell where the pulse starts or ends by looking at the deconvolved pulse.

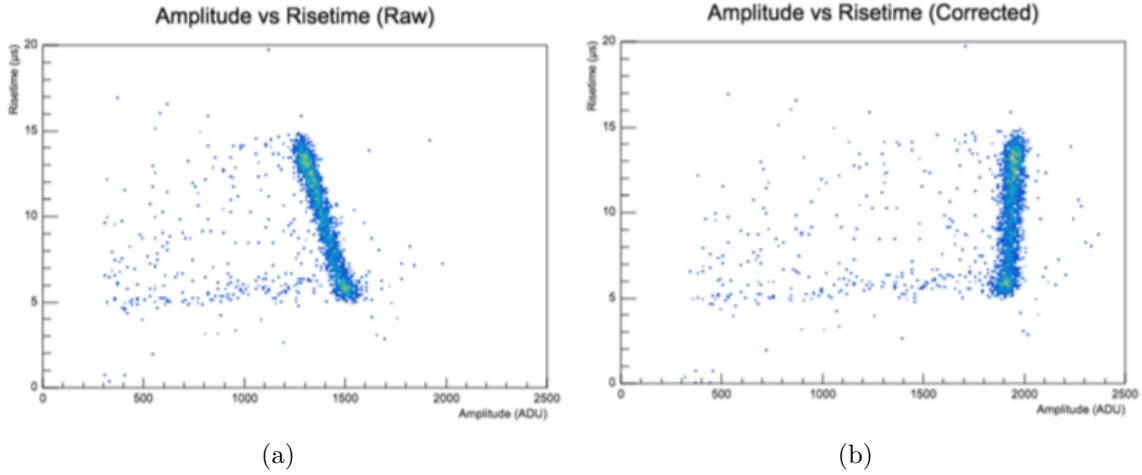


Figure 4.8: Risetime vs Energy for an alpha run in a small detector (15 cm wide). Left: raw data. Right: processed data. We clearly see that the offline analysis corrects the dependency of the amplitude on the risetime.

the only factor or not remains to be proven.

Note that the behaviour of high energy alphas is different from other type of lower energy events. These all start from the surface of the sphere, then move into the volume leaving a track of electron-ion pairs behind. The length of this track is mainly determined by the energy of the alpha, with a higher energy leading to a longer penetration depth. And in turn, the length of the track dominates our risetime estimator, since the spread in arrival time of the electrons will be mainly determined by the difference in drift time from the different points of the track. Furthermore, an alpha track pointing towards the electrode, rather than being tangential, will be less subject to attachment, since part of its track will start very close to the detector.

Estimator resolution

Both the MDec and DD methods correct the amplitude bias due to the ballistic deficit, so we need to compare them on the basis of the precision of their estimators.

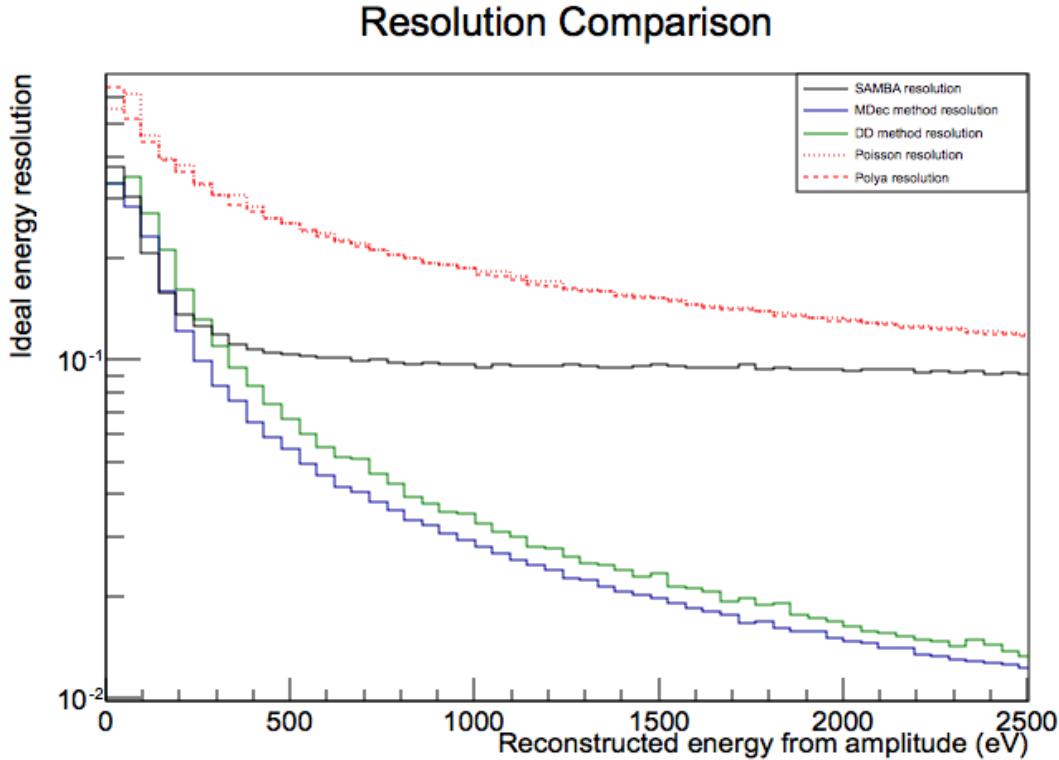


Figure 4.9: Energy resolution, depending on the reconstructed energy for each method. The “Ideal Energy” resolution is the normalized dispersion between the (normalized) amplitude computed by a method and the number of secondary charges in the simulated event, divided by the latter. The corresponding plots for the systematic uncertainties (Poisson from primary ionization, Polya from avalanche) are also shown for comparison.

The first, and arguably more important, is the amplitude estimator, since it gives a measure of the energy of the event. We want a metric for the reconstruction power of each estimator. This is a more nuanced question than it might at first appear, as expanded on in App. H. We will just cover the results of that discussion in this section, based on simulations of pointlike events distributed uniformly in energy and position inside the detector. The result for the amplitude (or reconstructed energy) resolution is shown on Fig. 4.9.

We see that the MDec method has the best amplitude resolution across all ranges of energy. However, given that in practice the effective resolution also includes the effect due to systematic uncertainties (due to primary ionization and avalanche fluctuations) of the detector, the marginal improvement of MDec over DD is not significant. In turn, SAMBA has similar performance to those two methods at very low energies (< 200 eV), but compared to them, degrades significantly beyond that point. The bias introduced by the ballistic deficit (which does not affect DD or MDec) limits its resolution to 10%.

The second estimator studied is the risetime. The main difference with the amplitude study is that the risetime of the pulse has a minimum strictly positive value that depends on the processing used. This makes it insensitive to the physical parameter being estimated (the RMS of the arrival time of the primary electrons) under a certain time. After normalizing the risetime to convert it into a proper estimator of the electron arrival RMS, we can plot the relative resolution of the estimator, as shown on Fig. 4.10.

While both MDec and DD offer the same performance for larger values of the electron arrival RMS, the minimum dispersion discernible by MDec is quite a bit worse than that for DD. This makes the DD method much more adapted for separating events from each other, as required for the search of KK axion events. In a similar vein, it will also produce better results when trying to discriminate events with a single electron (which always have an arrival RMS of 0) from other events, since multiple electrons have to arrive closer together before they can't be distinguished; this is useful in searches for very low energy events, since they always produce a single electron.

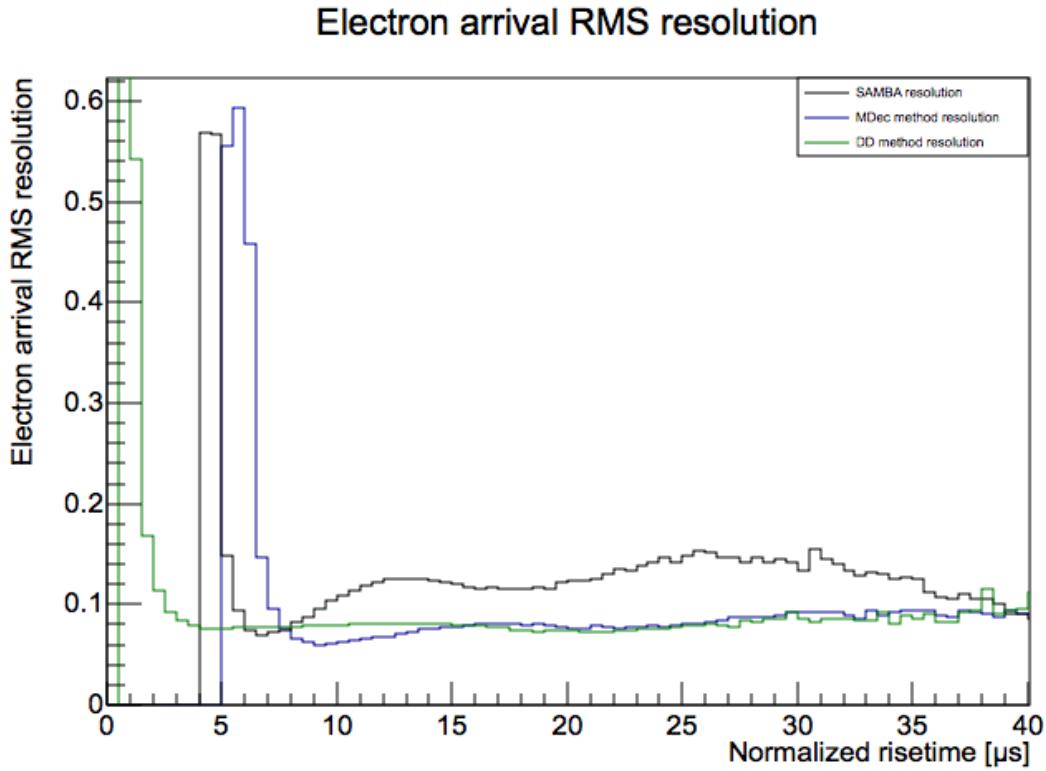


Figure 4.10: Electron arrival RMS resolution, depending on the normalized risetime for each method. The point at which the resolution shoots up for each method corresponds to the point where the risetime loses its discrimination power due to approaching its minimum value.

4.1.6 Event discrimination

We can use the raw data to make a first cut on our total database. The main parameters we use for discrimination are⁶

- Amplitude. The maximum of the pulse minus the value of the baseline. A basic estimate of the energy of the event.
- Risetime. The time difference between 10% and 90% of the amplitude. A basic

⁶The cuts described in this section are only for illustrative purposes. We will have to wait until Sec. 8.1.1 to go into the details of cut choice and its effect on our data.

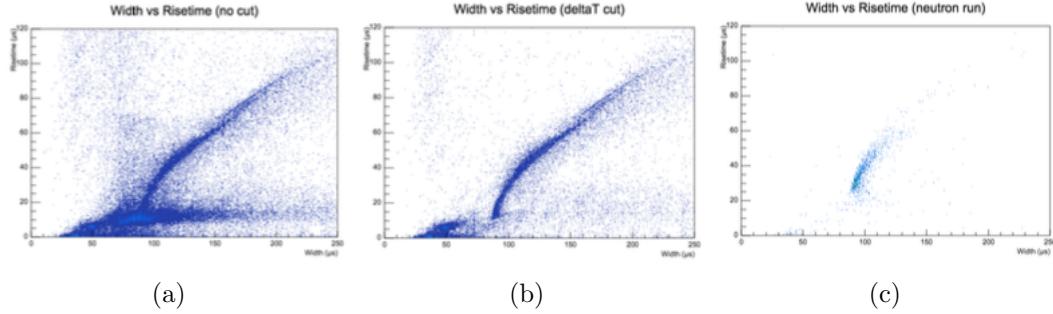


Figure 4.11: Riserise vs Width. (a) and (b) are taken from a physics run, with (b) having a cut of 4s since the last event. (c) shows the same plot for a short neutron calibration run, justifying the lower cut on width ($80\mu\text{s}$) and risetime ($10\mu\text{s}$).

estimate of the diffusion time of the event.

- Width. The Full-Width Half-Maximum (FWHM) of the pulse. Together with the Riserise, allows for preliminary shape discrimination.
- Time since last event. We ignore events too close in time to the previous one (a few seconds). We have noticed sporadic accumulations of fast, low-energy events, whose origin is uncertain. This cut removes them.

Making plots of these parameters against each other makes distinct populations appear. After examining these populations, we can determine which ones to keep, as seen on Fig. 4.11 and 4.12. Cuts on risetime and width eliminate most of the signals with “unphysical” shapes, that are absent from our calibration runs. Further cuts on the risetime allow to discriminate tracks from surface events. We are then left with the amplitude to select the range of energies we are interested in. Other parameters, such as the value of the baseline before the event, are also stored, but are currently not used for further discrimination.

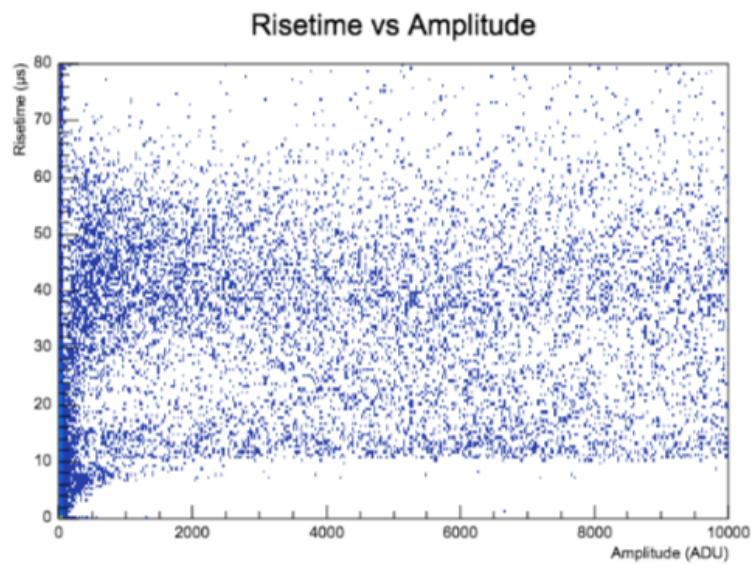


Figure 4.12: Risetime vs Amplitude. Events with a Risetime $\sim 40 \mu\text{s}$ are surface events. Below that are bulk events. Above, we have either tracks, multiple scatterings, or unphysical events. Below $\sim 12 \mu\text{s}$ we have electronic events.

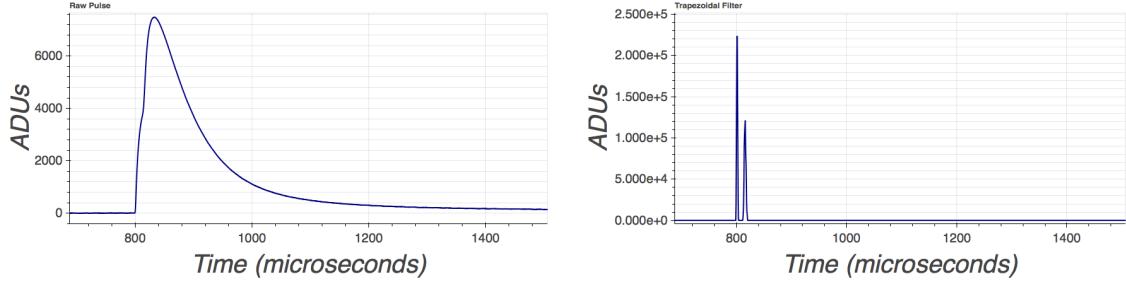


Figure 4.13: Left: Simulated pulse. Right: Electron signal before convolution by the detector response. The two pulses can be difficult to distinguish when looking at the final event.

4.2 Analysis of Multi-Pulse events

4.2.1 Description

Moving on to the analysis of axion-like data, we are now interested in events where two pulses arrive in close succession. For that, we need to be able to separate pulses within a given event window, and estimate the amplitude and risetime of each independently.

In Fig. 4.13, we see that multiple pulses arriving shortly after each other can be hard to distinguish, when looking at the raw signal. The distinction is much easier with the deconvolved pulse, although that comes with its own set of problems, notably the amplification of the noise.

To try and separate the pulses from the deconvolved signal, two algorithms were compared: a division of the window into sub-windows based on a threshold check, and a peak finder algorithm coupled with a multi-gaussian fit.

Window subdivision algorithm

After deconvolving the detector response from the signal, we do a running average over the result, twice. The running average is done over 20 bins ($\sim 10 \mu s$), as in

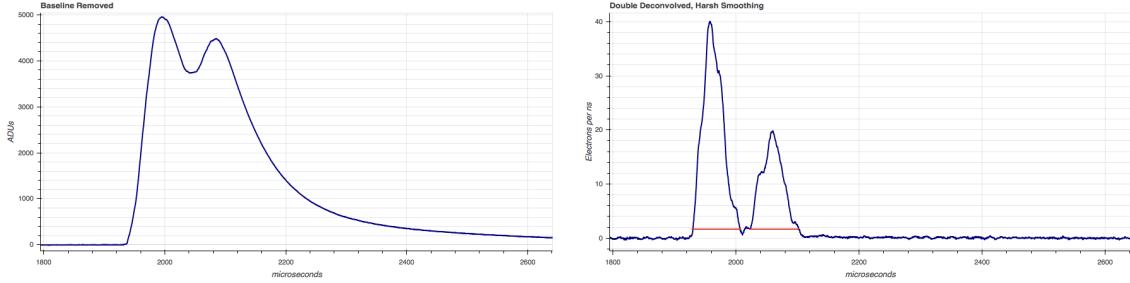


Figure 4.14: An example of a double event found by the MPA method. Left: Raw pulse. Right: Deconvolved pulse after smoothing, with horizontal lines showing the threshold level. The event dipping under the threshold in the middle splits it into two different pulses

the DD method step to find the end of the pulse. After this, we divide the signal window into smaller subwindows where the signal stays above a certain threshold. An example of such a pulse can be seen on Fig. 4.14. Intuitively, we can tell that the lower the threshold, the more “certain” we are that the different subwindows do split the event into different pulses, rather than just sections of the same one. Similarly, a stronger smoothing reduces the noise, but also spreads out pulses, making it harder to split pulses apart.

Since noise fluctuations may produce signals that dip above and below the threshold repeatedly, an extra step is performed. If any two subwindows are closer than 5 bins, both subwindows are merged together, including the section between them. Then, if any subwindow is shorter than 5 bins, it is dropped. Finally, safety margins of 50 bins are applied to each side of each subwindow, or half the distance to the closest subwindow in either direction (whichever is larger).

After the signal is divided in subwindows, we do the same process for the pulse in each one as in the regular DD method. The deconvolved signal, smoothed twice over 5 bins ($\sim 2.5\mu s$), is integrated. The amplitude of each pulse is then the difference

between the signal at the beginning and the end of its subwindow, and its risetime is the time difference between the point where it reaches 10% and 90% of its amplitude.

Since this algorithm was the one chose to produce results for the search of KK axions, it was name the Multi-Pulse Analysis (MPA) method.

Peak finder

The first step of the peak finder method is the same as the window subdivision method: I deconvolve the detector response from the raw pulse and do some running average smoothing on the result. But then, instead of subdividing the result, I use ROOT's TSpectrum *SearchHighRes()* method to find any peaks in the resulting event [51–53]. This method repeatedly deconvolves a gaussian of a given width from its input, then looks for peak positions by scanning the first and second derivative of the signal for sign changes and (negative) local minima, respectively.

While this does give the position of peaks and their height, it cannot find their width, which has a large effect on the final width and amplitude of the event. To remedy this, a fit of a function made of N gaussians is performed on the deconvolved pulse, where N is the number of peaks found by *SearchHighRes()*, and each gaussian is centred on the peak position; their width is set to some arbitrary value with the right order of magnitude, generally one to a few microseconds. For low electron count events, we would expect the fit result for σ to be the same for all peaks, and correspond to the timing resolution of the double deconvolution method; the theoretical electron signal should be shorter than a microsecond, so the observed width is just the smoothing from the method. Any signal with a larger width is then likely due to multiple electrons arriving too close to each other to be separated.

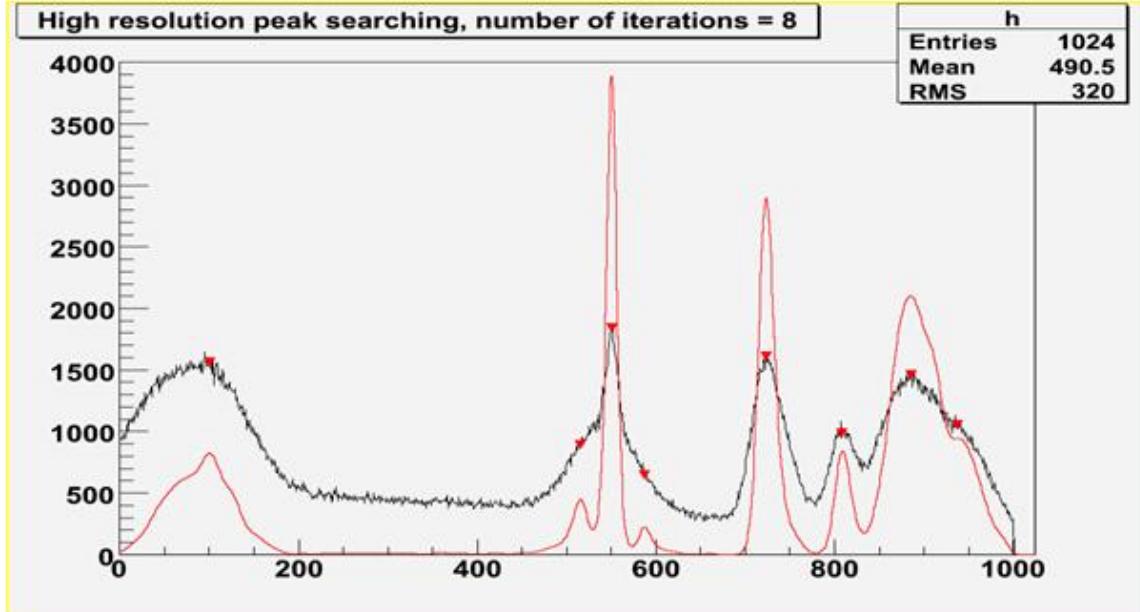


Figure 4.15: Example of SearchHighRes finding peaks on a source spectrum[54]. The black curve is the raw data distribution, the red one is the data after TSpectrum treatment. The red markers denote the position and height of the found pulses.

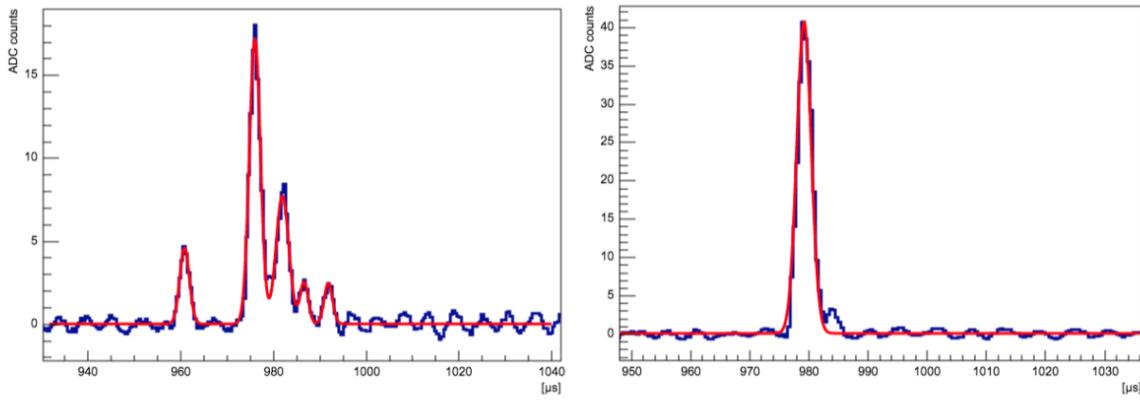


Figure 4.16: Multiple gaussian fit on real deconvolved pulses, based on the results from TSpectrum::SearchHighRes(). In blue, the deconvolved pulse; in red, the result from the fit. The method finds isolated peaks well, but struggles with close peaks, or when their height difference is too large.

The original motivation behind this approach was to obtain the number of primary electrons in an event directly, rather than trying to “guess” it from the number of secondary electrons from the event (which is essentially what we are measuring with the other methods). This would have a multitude of applications (see App. E.0.2 for what can already be done with statistical identification of primary electron counting), not the least of which would be a dramatically more precise measure of the energy of the event, since we wouldn’t be affected by the systematic uncertainty due to the avalanche.

Unfortunately, after running the simplest test simulation with 2 electron bulk events, the results were that the PF method could only find both electrons 62% of the time. This goes up to 74% when looking only at events where both electrons are separated by at least 6 bins ($3\ \mu\text{s}$), but we have no way to identify those events during real data taking. For comparison, with the same data set, applying a cut on the DD risetime above $4.5\ \mu\text{s}$ correctly identifies 65% of all double electron events, and more than 78% when looking only at events where both electrons are separated by at least 6 bins. Since being able to distinguish events with 1 electron from events with multiple electrons is both the easiest and the most important task we could give to an electron counting algorithm, the performance of this method is considered insufficient.

Finally, the major drawback of this method is how slow it is, almost exclusively due to the fitting. Even when forcing it to fit a single gaussian (instead of one per peak found by *SearchHighRes()*) we are looking at a few milliseconds per event, and roughly three times slower than the Double Deconvolution method for no observable gain. However, if left to fit a function with one gaussian per peak found, an event

with a dozen peaks can take upwards of a second to be processed. For comparison, a simulated run with one million events takes 2.5h to be processed with the DD method, 6.5h to be processed with the PF method fixed at 1 peak maximum, and more than 10 days with the full PF method. While not a dealbreaker by itself, the lack of prospects for the method just made it not worth it to try to optimize processing times into something more usable.

4.2.2 Sanity cuts

Once an event is reconstructed as having multiple pulses, further cuts are applied to guarantee that the event is physical, and consistent with what we expect from axion-like events.

- Cut on the number of pulses found. There has to be at least two of them.
- Cut on the amplitude of the two largest pulses. They have to be above a minimum value to prevent “triggering” on fake pulses.
- Cut on the minimum of the deconvolved pulse. If the deconvolved pulse goes too low into the negatives, the pulse was over-deconvolved, likely splitting a single pulse into two.
- Cut on the total amplitude of the two largest pulses. Their sum has to be within the energy range of interest (e.g., 2 to 20 keV).
- Cut on the relative amplitude of both pulses. Ideally, they should both be the same, within the uncertainties due to our systematics, since both photons have the same energy. However the presence of attachment requires this cut to be a bit more lax than that.

- Cut on risetime and width of both pulses. Their values should be consistent with those found for physical pulses.
- Cut on relative risetime of both pulses. The second pulse should be wider than the first, since the electrons come from further away and experience more diffusion.
- Cut on the time separation of both pulses. The pulses should come within a time difference consistent with the maximum drift time of electrons in the sphere.

An important consideration is that the cuts depend not only on our expectations of what an axion event looks like, but also on our certainty about it. When reporting limits, we need to also report what effect our uncertainties have on those limits. Ideally, when optimizing a cut, we should consider both its rejection efficiency, but also its sensitivity to calibration errors. For example, a stringent cut on the relative amplitude of both pulses would be much better at rejecting non-axion events. However, if we underestimate the effect of electron attachment, then it would also reject more axion-like events, lowering our sensitivity to axions. This could be translated into a large systematic uncertainty on our final results, if our calibration of the attachment is not very precise.

In practice, optimizing the cuts based on our calibration uncertainties would also require good knowledge of our background. Absent that, the cuts that had the largest effect on our expected sensitivity to axions were relaxed until their associated uncertainty fell down to tolerable levels. This was the case mainly for the relative amplitude of both pulses, and their maximum width and risetime.

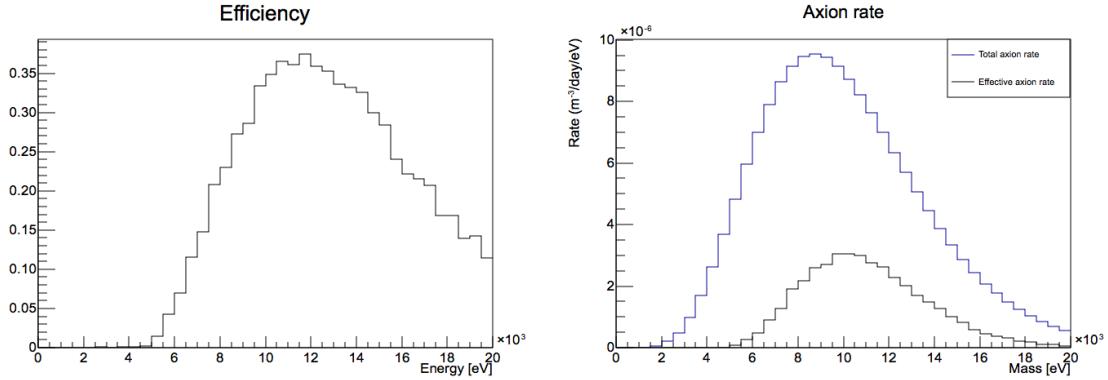


Figure 4.17: Left: MPA efficiency for KK axion events. Right: Total axion rate, and axion rate observed in our detector after applying our efficiency.

4.2.3 Performance

Efficiency and False Positive rate

To be able to set constraints on the solar KK axion model, we must be able to determine the expected rate of events on our detector. A necessary step of this process is determining the efficiency of our detector. To test the different methods and thresholds, a simulation with a uniform energy and position distribution of axion decays was performed, to get the efficiency depending on the total energy of the event. We can then multiply that by the energy distribution of the axion decay rate to get the total expected rate in the detector. Both are shown in Fig. 4.17.

The best axion detection efficiency we can achieve is about 0.35 for a 12 keV axion, although this depends on the processing parameters chosen (cf. App. I). Considering that this includes both the loss of efficiency due to one of the photons escaping the detector and the loss of efficiency due to the MPA method missing events where the pulses are too close, this is a satisfying result (cf. Sec. 5.6.2 for the loss of efficiency due to containment probability alone). But just as important as the maximum efficiency,

is where that is reached. Our detector and algorithm combined are most sensitive to KK axions in the 10 to 14 keV range. This has to be contrasted with the expected energy distribution of solar KK axions, which are at a maximum in the 7 to 11 keV range. The fact that both don't quite align properly means that our sensitivity to axions will be lower than the optimistic 0.35 previously announced. Indeed, when integrating over the whole range of energies after applying our efficiency, we see that we only keep about 25 % of the total rate. While this could definitely be improved, it is suitable to produce first results.

The second, and arguably even more important metric to be measured, is the false positive rate of the MPA method. By this we mean the number of events created by single energy depositions that are reconstructed as being composed of multiple pulses. While the expected number of axion events in the detector for the physics data available is less than one (barring background), the total number of events in that data is more than a million. Even applying some preliminary cuts to the data leave at least 10000 events that could potentially be falsely reconstructed as KK axion events. As such, we need an extremely high rejection rate of single energy deposition events, of the order of at least 10^{-4} , or else we will be dominated by those events that manage to pass the cuts. This can be achieved with proper processing parameter choice (cf. App. I). The question of the rate of background events that generate axion-like events, such as a photon doing two Compton interactions, is more complex, and will be covered in Chap. 8.

Separity

While the performance of the algorithm is very encouraging, especially considering that the running conditions were not chosen for this kind of study, we should dive in further into them to understand which double pulses are correctly resolved as “double”. This will allow us to tune the setup of the detector (gas, pressure, voltage...) to achieve higher efficiencies while keeping a low false positive rate.

The most direct way to recognize which events will be reconstructed as “double” is to study the time separation we need. This is unfortunately not a fixed number. Since the algorithm used relies on both pulses going back to baseline in between them, the “spread” in electron arrival time for both pulses will affect the necessary time difference between both. This is shown in Fig. 4.18. The time difference is defined from the centre of one pulse to the other, and the width is defined as the time between the first and last electron for each pulse.

As expected, the necessary condition to separate both pulses is that the average width of both is smaller than the time difference between them⁷. The only time when this is not true is for sufficiently short pulses, where the smoothing of the algorithm dominates over their natural width. This will normally only happen for energy depositions very close to the sensor, which are comparatively rare, so unlikely to make much of a difference. Note that these conclusions are independent of the running conditions of the detector. They just formalize the requirements to have a double-pulse event be separable.

It is more fruitful to think in terms of the position of the photon interactions, instead of the time difference between pulses. This, however, strongly depends on the

⁷This condition is equivalent to having the first pulse end before the second one begins.

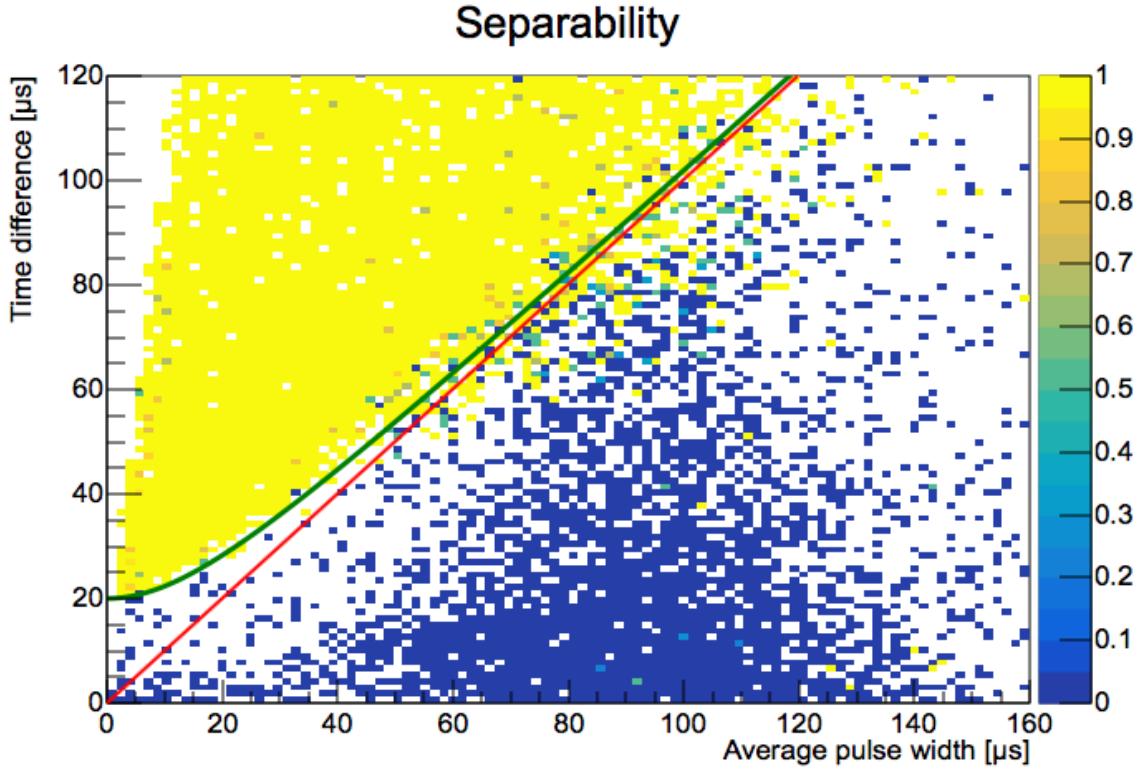


Figure 4.18: Capacity of the MPA to separate two pulses (z-axis, showing the proportion of axion events that were properly reconstructed as double), based on the average width of both and the time difference between them. At smaller pulse widths, the effect of the smoothing from the method dominates. Otherwise, the MPA separates both pulses as long as the time difference is larger than their widths, as expected. The red line represents $y = x$ (no smoothing ideal scenario), the green line $y = \sqrt{x^2 + 20^2}$ (smoothing of 20 samples).

running conditions of the detector: the drift and diffusion time of electrons varies with the gas, the pressure, the voltage, among others. We will just assume the conditions in SEDINE's physics run.

Fig. 4.19 shows that in this scenario, the interaction point of both photons have to have a radial distance difference of at least 2 cm for them to be separable. While the difference has to be longer than that for interactions happening within 15 cm of the

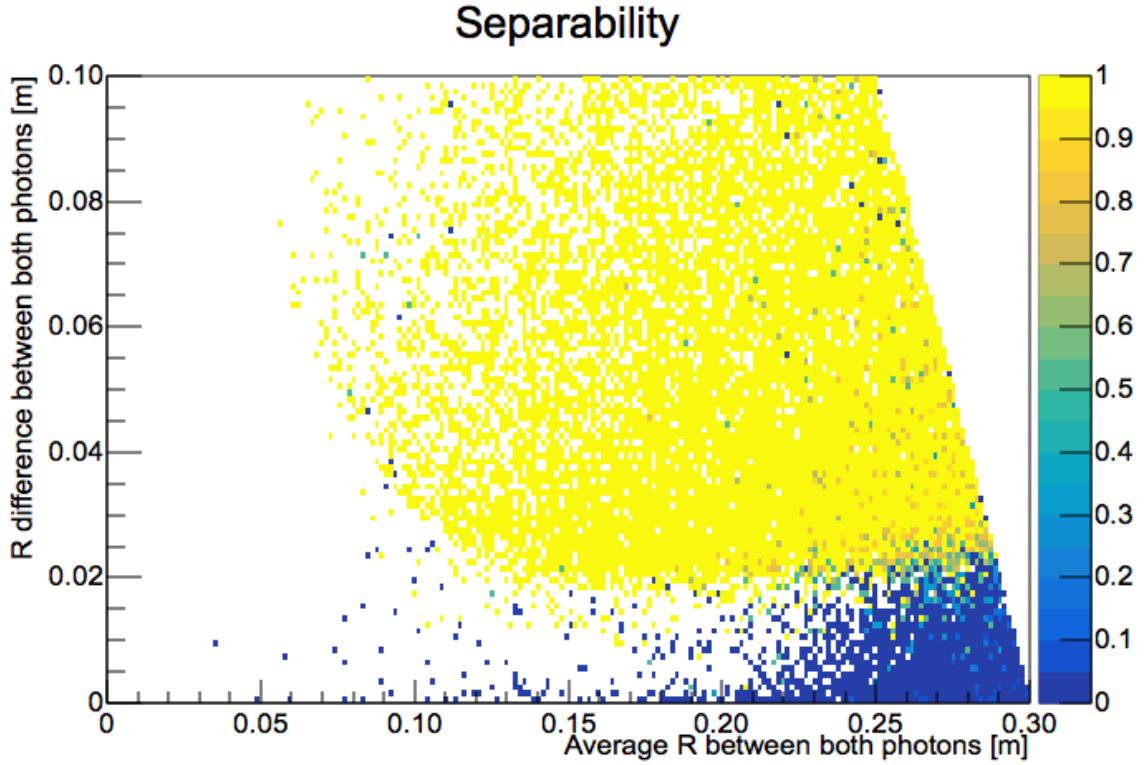


Figure 4.19: Capacity of the MPA to separate two pulses (z-axis, showing the proportion of axion events that were properly reconstructed as double), based on the average distance to the centre of the detector of both photon interactions, and the difference in distance between them. We require both photons to interact inside the detector. Running conditions of the detector are those from SEDINE’s physics run. From 15 cm and above, both photons have to have a radial distance difference of at least 2 cm to be resolved distinctly.

centre, the spherical symmetry of the setup means that volume represents at most $1/8$ of the whole detector, so the generalization is appropriate. Note that this difference is not the actual distance between photons, but the difference in radial positions. Since the photons are unlikely to be aligned with the radius of the sphere, the real distance between the photon interactions will often need to be longer than that before they become separable.

If we wanted to decrease this minimum distance, we would have to change the running conditions to either decrease the diffusion time of electrons, or increase the dependency of the drift time on the radius. The former would make for narrower pulses, and the latter would spread both pulses further apart from each other. Unfortunately, both are strongly correlated, and change roughly at the same rate when changing the voltage or the pressure of any given gas (cf. Sec. 5.3). So this distance will mostly be determined by the choice of gas used.

Summary

In this chapter, we have explored algorithms to identify the type and energy of recorded events, and if multiple pulses were recorded at once (and if so, what were the characteristics of each). The main processing applied to the recorded pulses was their deconvolution by the response function of the SPC, based on our understanding of the detector. This procedure corrects for the effect of the ballistic deficit (i.e. wider pulses having a smaller amplitude than narrower pulses of same energy), and separating pulses as close as $20\mu\text{s}$ from each other. For SEDINE conditions, this corresponds to energy depositions as close as 2 cm radially from each other, and an axion event detection efficiency of up to 35% at 12 keV.

The performance of the described algorithms was tested against simulations, which will be covered in the next chapter.

Chapter 5

Simulations

Whenever trying to interpret results from an experiment, a crucial step is to compare it with the results we were expecting before doing the experiment. In some situations, this can be as straightforward as comparing two different experiments in different conditions. For example, if we could take two datasets with the KK axion signal respectively “switched off” and “switched on”, we would then just need to compare the number of events in the region of interest for both to either set constraints or prove their existence. Unfortunately, it’s impossible to “switch off” the KK axion signal, so we need to turn to simulations to produce that dataset artificially.

In this chapter, I will mainly cover the different steps we take to simulate events. The focus will be on the technical aspects of the simulation, but also on developing an understanding of the response of our detector to varying conditions, trying to give quantitative rules-of-thumb wherever possible. This chapter will primarily cover the simulation of the drift of primary electrons, with the last part going over considerations specific to KK axions. The signal formation process is covered in App. C.0.2 instead.

Most applications of the simulations are described elsewhere. For the comparison

of simulations with data for calibration purposes, see Chap. 6. For the characterization of different pulse processing algorithms, see the relevant sections in Chap. 4. Finally, for the application of the simulations to extract KK axion limits based on our data, see Chap. 8.

5.1 Simple simulation

Since, as seen in Sec. 4, we understand how the pulse formation of our detector works, we can use this understanding to simulate pulses to estimate the efficiency of our analysis and cuts. The steps for our original, simple simulation, as seen on Fig. 5.1, are:

- 1 Selection of energy E (electron-equivalent) and position r of a pointlike event in the detector. At this level, we may use an uniform distribution of energy, and positions either uniformly distributed in the volume or very close to the inner surface of the shell. For pointlike deposition in the bulk, we take $r = R_{shell} u_1^{1/3}$, and $\theta = \text{acos}(2u_2 - 1)$, where u_1, u_2 are the results from two uniform distributions between 0 and 1 (θ is not relevant in spherically symmetric situations, but will be for the cylindrically symmetric ones).
- 2 Get the number of primary electrons from the energy. We take this from a Poisson distribution with mean $n_p = E/W_I$, where W_I is the mean ionization energy for the gas. If deemed necessary, apply an adhoc correction based on r , to account for attachment.
- 3 Get the standard deviation σ of the drift time of the primary electrons from the position r . Some preliminary Magboltz simulations, contrasted with the data,

suggest that $\sigma = \sigma_{max}(r/r_{max})^\alpha$, with $\alpha \sim 2.75$ (see Sec. 5.3).

- 4 Draw the arrival time of each primary electron from a Gaussian distribution with standard deviation σ .
- 5 For each primary electron avalanche, draw a gain G from a Polya distribution¹.
- 6 Convolve the resulting signal with the current induced by a drifting ion.
- 7 Convolve the resulting signal with the transfer function of our preamplifier.
- 8 To simulate the noise of our detector, add noise taken from the pre-trigger region of a randomly selected event from a run with same physical conditions as our simulation. To match the simulated pulse to the noise, we first need to multiply the pulse by a constant, to match the equivalence between ADUs and eVs given by our calibrations.
- 9 Store the energy, number of primary electrons, total number of secondary electrons, diffusion time, RMS of our signal (pre-convolutions), and any other variables we may need for our posterior analysis.
- 11 Repeat N times to generate a simulated run.

This simple simulation has the advantage of being extremely fast, generating thousands of events in a second. It was also the first code that could generate simulated runs for testing purposes. It's still usable for any tests that are mostly concerned

¹The nature of the charge multiplication in the avalanche region might lead one to assume the gain distribution should follow a decaying exponential distribution. However, experimental measurements with Micromegas [55, 56] and Gas Electron Multipliers [57] show that at high ratios of electric field to pressure, the gain from each primary electron is better described by a Polya distribution.

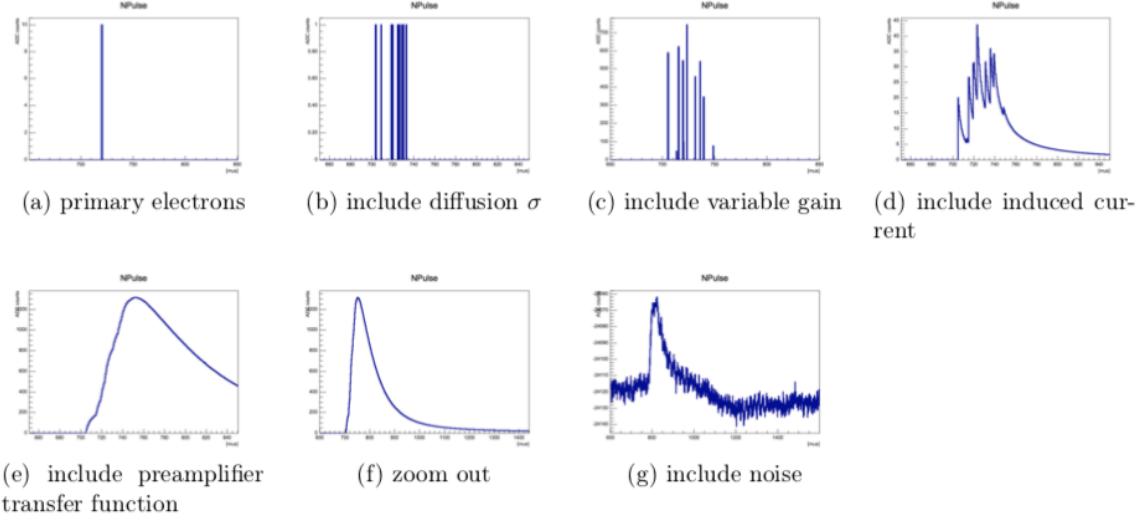


Figure 5.1: The different steps of the simulation, after n_p and σ have been determined.

about the general shape of events, rather than their distribution. The main limitation of this approach is that it does not take into account the real electron drift diffusion inside the detector. For that, we need to go deeper into steps 3 and 4.

5.2 Electron drift simulation

For the full simulation, we require a proper way to model the transport of electrons in the drift region, and their spread in arrival times, which will in turn affect the risetime of the final pulse. To do this, we need information on the speed at which electrons drift inside the gas, and how they diffuse as they drift. This information can be obtained by using Magboltz [58] to generate electron drift velocity and diffusion inside a given gas composition and for a given electric field (through Monte Carlo simulations of electron collisions), and using COMSOL [59] to calculate the electric field strength and direction at each point inside the detector (through finite-element computations). More details on Magboltz and COMSOL will be given in the next

sections.

Once we have this information, simulating the drift of an electron is relatively straightforward. We assume a cylindrically symmetric detector. For an electron in a given position (ρ, z) , we obtain the electric field from COMSOL, E_ρ, E_z . We can get the field strength $E = \sqrt{E_\rho^2 + E_z^2}$. From Magboltz, we then obtain the drift velocity of the electron, $v(E)$, and the longitudinal and transverse coefficients $DL(E)$, $DT(E)$. Since the strength of the electric field depends on the position of the electron, these values have to be updated after small “steps” of the electron, dr , usually taken to be half a centimetre or less. Thus, before taking the diffusion of electrons into account, the new position of the electron is:

$$\begin{aligned}\rho &\rightarrow \rho + dr \frac{E_{rho}}{E} \\ z &\rightarrow z + dr \frac{E_z}{E}\end{aligned}$$

Note that so far the process is deterministic. Since we are fixing the length of the step to dr , this gives us a time for the step of $dt = dr/v(E)$. This is important to compute the diffusion of the electron, since it depends on the amount of time the electron spends in the drift region. In particular, the variance in the position of the electron in the longitudinal direction is $DL(E) dt$, and similarly for both transverse directions. If we name dr_{DL} , dr_{DT1} and dr_{DT2} the difference in the final position of the electron after one step due to diffusion in the longitudinal and both transverse directions, we have:

$$dr_{DL} = \mathcal{N}(0, \sqrt{DL(E) dt})$$

$$dr_{DT1} = \mathcal{N}(0, \sqrt{DT(E) dt})$$

$$dr_{DT2} = \mathcal{N}(0, \sqrt{DT(E) dt})$$

where $\mathcal{N}(0, \sigma)$ is the result from a normal distribution centred at 0 with a standard deviation of σ . As such, these values are random, not deterministic. Separating the effect of each direction on the final position of the electron after one step, we get the following corrections in cylindrical coordinates:

- The longitudinal direction

$$\rho \rightarrow \rho + dr_{DL} \frac{E_\rho}{E}$$

$$z \rightarrow z + dr_{DL} \frac{E_z}{E}$$

- The transverse direction in the $\rho - z$ plane

$$\rho \rightarrow \rho - dr_{DT1} \frac{E_z}{E}$$

$$z \rightarrow z + dr_{DT1} \frac{E_\rho}{E}$$

- The transverse direction orthogonal to the $\rho - z$ plane

$$\rho \rightarrow \sqrt{\rho^2 + dr_{DT2}^2}$$

$$z \rightarrow z$$

Technically, the correction for the diffusion orthogonal to the $\rho - z$ plane should be applied last, but in practice it's a negligible contribution to the final position of the electron. The other two corrections can be applied in either order.

This process is then repeated, recomputing the new values for the electric field, the drift velocity and the diffusion coefficients after every step, until the electron either reaches the sensor, or escapes the detector. The arrival time of all electrons that reach the sensor is recorded, and a pulse is generated for each of them starting at their arrival time (including the effect from the avalanche), then added together. The final pulse is then shifted so that it starts at the centre of the event window.

Magboltz also has information on the probability of an electron being captured during their drift. For a step of length dr , the odds of survival for the electron are:

$$p = e^{-\text{ATT}(E) dr}$$

If we want to simulate attachment at the same time as drift times, the additional step is to draw a random number between 0 and 1, and only keep the electron if it's smaller than p .

5.3 Magboltz

Magboltz is a numerical solver for the Boltzmann transport equations for electrons in gas, under the influence of an electric and magnetic field[58]. We use it to find the drift speed and diffusion of electrons in our detectors depending on the run conditions. In our case, we are only interested in the effect from a pure electric field, which is the reason why the speed of the electrons is purely in the direction of the field (longitudinal direction), and both of the transverse diffusion coefficients are identical. In Fig. 5.2, you can see the computed drift velocity, longitudinal diffusion coefficient, attachment coefficient, and the Townsend coefficient for an electronic avalanche.

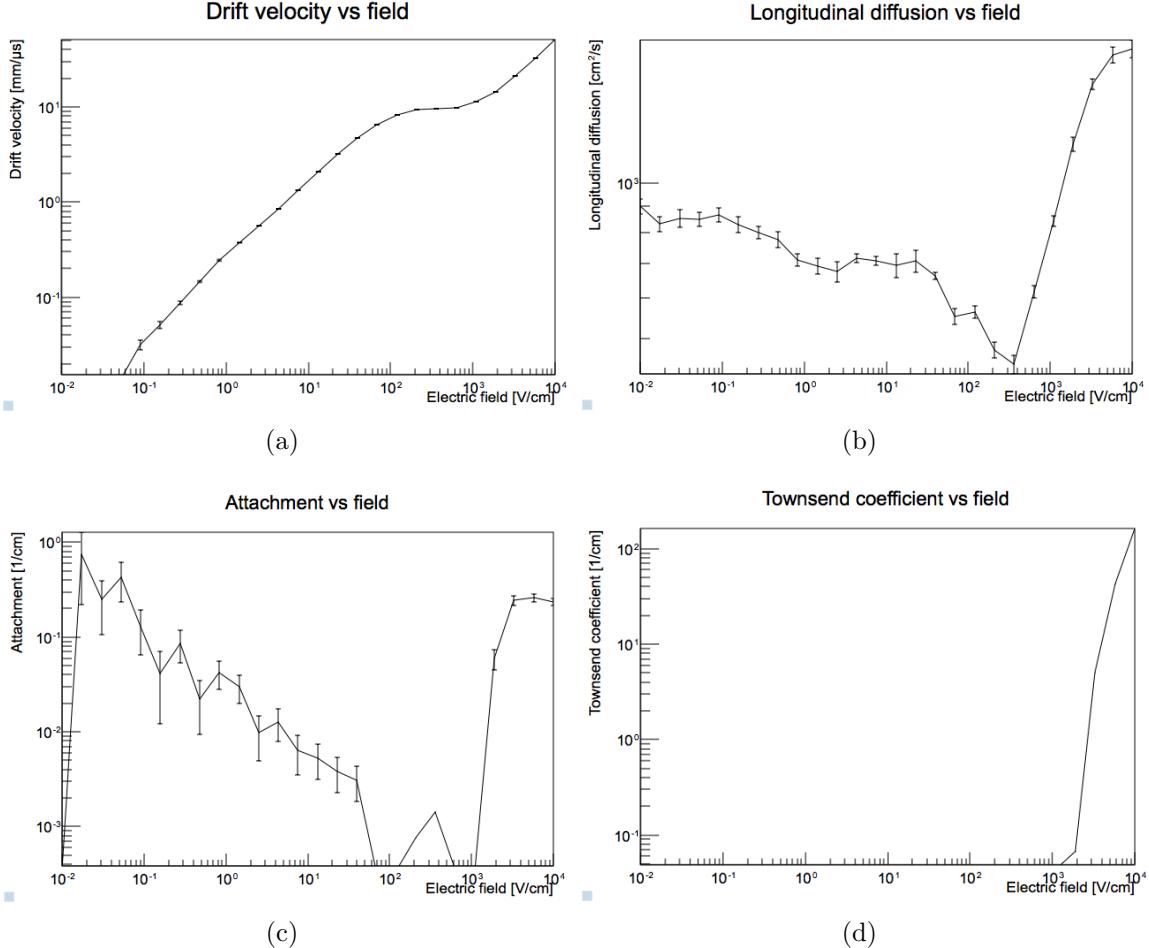


Figure 5.2: Magboltz data for 3.1 bar of Neon with 0.7% of CH₄, and an Oxygen contamination of 16 ppb.

For the purpose of visualizing the behaviour of electrons inside the detector, drawing these values as a function of the electric field is not useful. Instead, we can assume the ideal electric field:

$$E(r) = V_0 \frac{r_{\text{sensor}} R_{\text{shell}}}{r_{\text{sensor}} + R_{\text{shell}}} / r^2$$

to show these same values in terms of the radius from the centre of the detector,

as shown on Fig. 5.3. This clearly shows that the electron drift speed becomes slower farther away from the sensor, while the attachment probability becomes larger. A potentially misleading conclusion from looking at the longitudinal diffusion coefficient is that, since it stays roughly constant throughout most of the drift region of the detector, that would mean that the diffusion time depends roughly linearly on the radius from which the electron started drifting. This is incorrect, however, since this does not take into account the fact that the drift velocity is not constant in that same region.

To provide a more accurate plot of the contribution to the total diffusion time, I introduce a new parameter, the “temporal diffusion parameter”, defined as $\text{DIFF} = DL/v^3$, where I’ve assumed that either transverse diffusion will have little effect on the final arrival time of the electron (as is the case for our data). This new term, with units of time squared by distance, can be integrated over the length of the electron path to obtain the variance in arrival times of multiple electrons drifting through the same path.

A simple demonstration of the definition of this term can be done by assuming steps in a single dimension, and that the effect of the electron diffusion is small compared to the distance covered by the electron drift. After some time dt , the electron will move a distance $dx' = v dt + \mathcal{N}(0, \sqrt{DL dt})$. But, if instead of looking for the distance travelled after a time dt , we want to know the time it took to move a distance $dx = v dt$, we then have $dt' = dt \frac{dx}{dx'} = dt \frac{1}{1 + \mathcal{N}(0, \sqrt{DL dt})/(v dt)}$. Then, since the diffusion effect is small compared to the step size, $dt' \simeq dt (1 - \mathcal{N}(0, \sqrt{DL dt})/(v dt)) = dt - \mathcal{N}(0, \sqrt{DL dt/v^2})$, where I have also used that $\mathcal{N}(0, A\sigma) = A\mathcal{N}(0, \sigma)$. After simplification, we end up with:

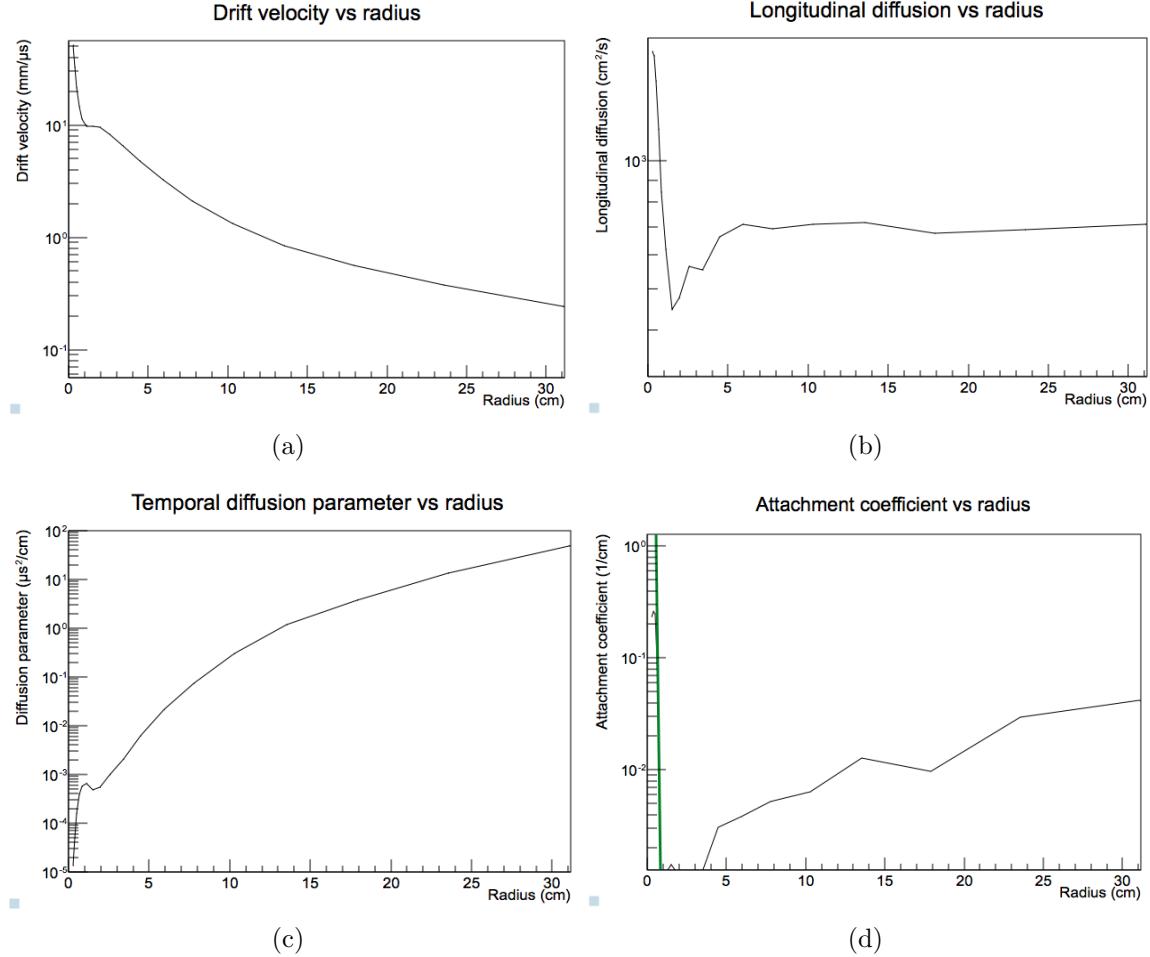


Figure 5.3: Same Magboltz data, adapted to the scenario of a 30 cm radius sphere with a 3.15 mm sensor at 2520 V. The bottom-right plot shows the Townsend coefficient in green, shooting up below 1 cm, to indicate the point below which attachment can be ignored.

$$dt' = dt - \mathcal{N}(0, \sqrt{(DL/v^3) dx})$$

or, equivalently, $DIFF = DL/v^3$. We can obtain the diffusion time from $DIFF$ with the following relationship:

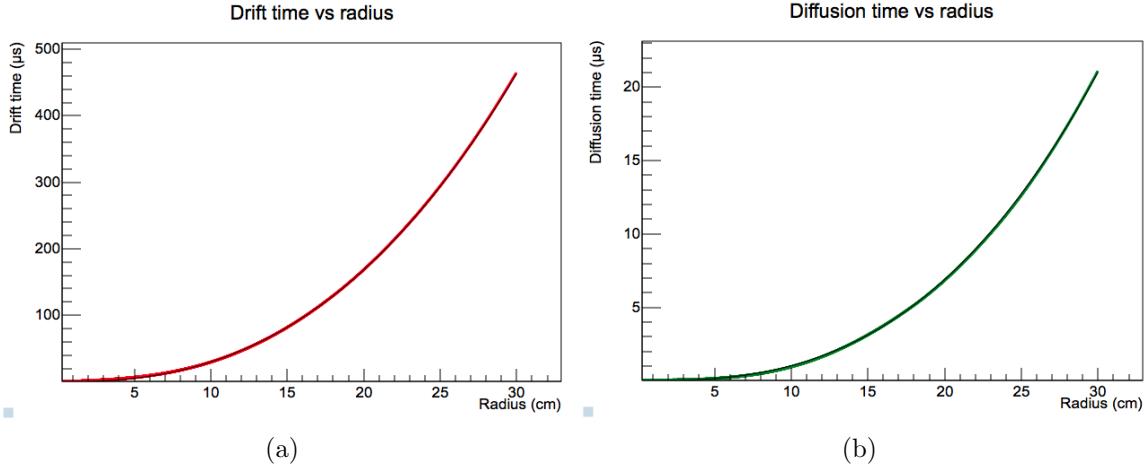


Figure 5.4: Drift and diffusion time as a function of radius for an ideal electric field, with the same running conditions. The modelled functions are in black.

$$\sigma_{diff} = \sqrt{\int_{r_{sensor}}^{r_{event}} \text{DIFF}(E(x)) dx}$$

With both $v(r)$ and $\text{DIFF}(r)$ known in the ideal field scenario, it becomes easy to plot the dependency between drift time or diffusion time and the radius, as shown on Fig. 5.4. Some quick fits show that, for the case of a 30 cm radius sphere with a sensor radius of 3.15 mm, an applied voltage of 2520 V, and 3.1 bar of Neon with 0.7% CH₄, the drift time is very well modeled by $drift(r) = 463 (r/30)^{2.5}$ above 7 cm, while the diffusion time is modeled by $\sigma_{diff}(r) = 21.0 (r/30)^{2.75}$ above 13 cm

Finally, while the results shown so far allow us to quickly understand how the behaviour of the electrons drift will change when we modify V , and hence the electric field, we haven't yet mentioned the effect of modifying the pressure. Fig. 5.5 shows the drift speed and longitudinal diffusion for Neon with 2% CH₄ for varying pressures. Noting that the usual electric fields in the drift region of our detectors is in the $10^{-1} - 10^1$ V/cm range, we clearly see an approximate dependency of $1/P$ for both

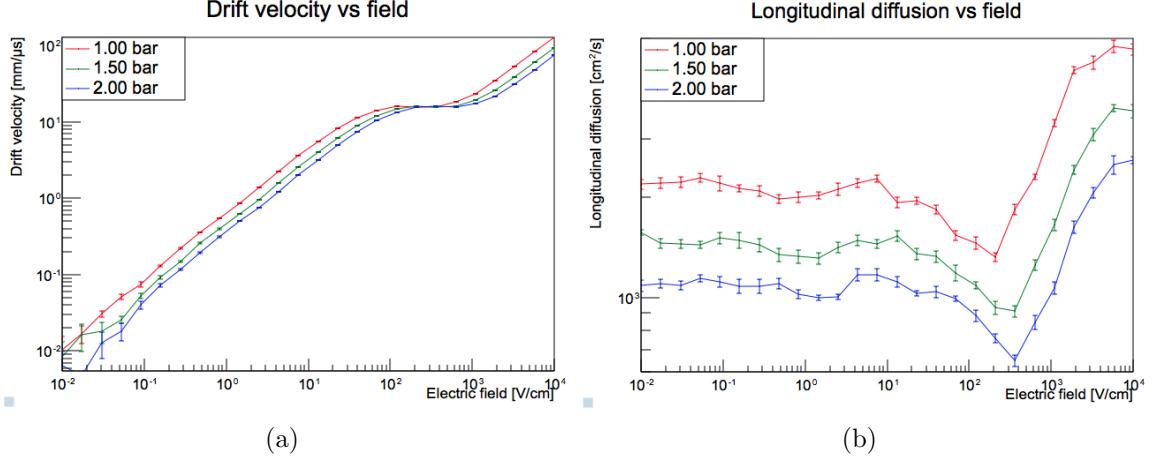


Figure 5.5: Magboltz data for 1.0, 1.5, 2.00 bar of Neon with 2.0% of CH_4 .

the drift speed and the longitudinal diffusion. Furthermore, the drift speed can be approximated by $v(E) = \alpha E^{0.91}$ for electric fields typical for the drift region, while the longitudinal diffusion is close to constant. Putting all this together, we get, in the case of Neon with traces of CH_4 :

$$v \propto V^{0.91}/P$$

$$\sigma_{diff} \propto P/V^{1.33}$$

5.4 COMSOL

COMSOL Multiphysics [59] is a finite-element analysis tool and solver. We use it to solve Maxwell's equations within the detector to get its electric field configuration. The input we provide to COMSOL is the geometry of the detector, the physical properties of the materials involved (mainly conductors, insulators, and gas), and the voltage applied on the boundaries of the detector. COMSOL then generates a

meshing for the detector, and computes the electric field and potential for all the edges of the meshing. The results are then interpolated and stored for late use². This allows our simulations to take the anisotropies of the field (mainly due to the rod) into account.

For the purpose of the simulations described in this chapter, the detector configurations studied were all kept cylindrically symmetric. I will be giving some examples of configurations where that is not the case, but they were not used for generating any simulated runs, and are only shown for illustrative purposes. The electron drift code is written assuming a cylindrically symmetric geometry, but could be adapted by an enterprising new grad student³ to accept 3D geometries.

The physics run I used to set limits on KK axions used a simple rod configuration, and hence was the one whose behaviour interested me the most. The field configuration with a secondary “umbrella” electrode is also shown here for comparison. For the work I performed to simulate more recent models of central electrodes, see App. J; they also serve as a good demonstration of the limits of what we can do with COMSOL at time of writing. More quantitative comparisons between simulations and calibration data will be discussed on Chapter 6.

5.4.1 Simple rod

The simple scenario has a grounded metallic rod holding the high voltage electrode in the centre, and shielding the wire carrying the high voltage to the electrode. A picture of the setup in COMSOL and a close-up of an example of such a sensor can

²We save the interpolated data in a ROOT histogram, because looking up values in a histogram is much faster than interpolating values from the mesh. Doing the latter would severely slow down the simulation of the electrons drifting in the detector.

³Just not the one writing this document.

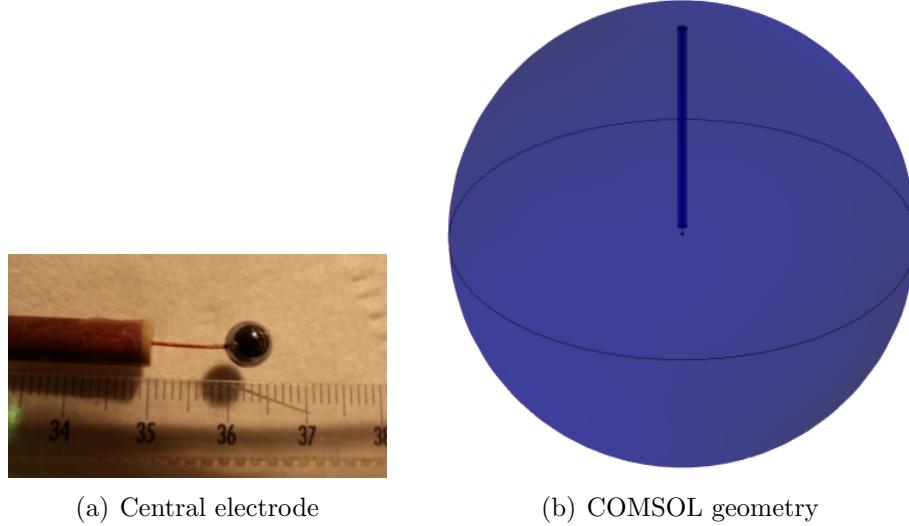


Figure 5.6: SPC setup with a simple sensor.

be seen in Fig. 5.6.

The presence of the grounded rod disturbs the local electric field, and moves us away from the ideal scenario of an electrode floating in the centre of the detector. The field becomes more anisotropic closer to the rod in two ways: the field lines curve around towards the rod, leading to longer drift paths for electrons towards the electrode; and the lower field close to the rod means a smaller avalanche will be produced for primary electrons reaching the electrode in that area.

These changes to the field configuration help explain the larger dispersions in diffusion and drift times that we observe in our data, compared to what we would see for the ideal geometry. We can do this quantitatively with our electron drift code, obtaining a tail of long drift times corresponding to events happening close to the rod. And, while we can also use this simulation to explain the presence of low amplitude events in calibration data, even after correcting for attachment, this explanation remains qualitative, since we do not currently have any avalanche simulations reliable enough

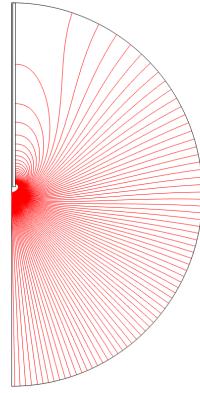


Figure 5.7: Field lines for a detector with a simple rod and electrode configuration. Field lines away from the rod stay fairly isotropic, but become increasingly curved as we get closer to it.

to consistently match the data. Any corrections to the distribution in amplitudes for an event of a given energy have to stay ad hoc for now, based on those calibrations; one such correction was applied for the first WIMP results with NEWS-G [42].

5.4.2 Rod with umbrella

One of the first attempts at correcting the field for better isotropic behaviour was to add a second electrode, called the “umbrella”. Multiple designs of umbrellas have been developed and tested, always with the objective to improve the amplitude resolution of the detector. While we do not have a reliable way to translate the strength of the electric field into a value for the avalanche, we can still gain qualitative insights into the homogeneity of the gain from COMSOL simulations.

Fig. 5.8 shows one such sensor with an umbrella. Usage of this sensor to take calibration data at Queen’s revealed that applying a negative voltage on the secondary electrode improved the energy resolution of the detector. This was contrary to the common wisdom of the collaboration, which would have expected the resolution to

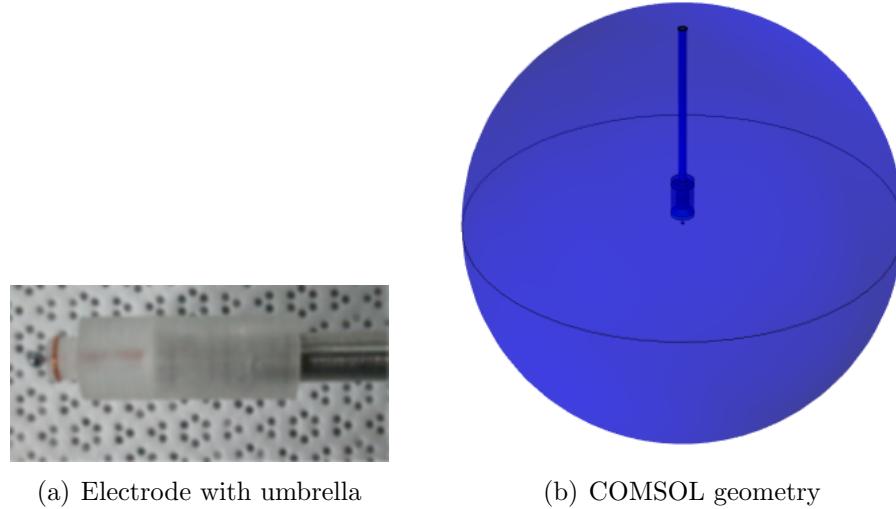


Figure 5.8: SPC setup with a sensor with umbrella.

improve if the umbrella matched the potential it would experience in an ideal geometry, since it would help reduce the anisotropy of the field due to the rod. In fact, if we show the electric field close to the electrode for both the simple rod configuration and the rod with umbrella, we do see a slight degradation of the isotropy of the field for the negatively charged umbrella, as shown on Fig. 5.9. This is even more apparent when looking at the field lines in the drift region of the detector, which are significantly modified by the presence of the umbrella.

While this may appear puzzling at first, COMSOL reveals a straightforward explanation. The field lines for the simple rod drop off in density at an angle of 2.14 rad above the vertical line. However, the presence of a negatively charged umbrella makes field lines twist away from it, and they are pushed back to the other side of the central electrode: as such, the field line density drops off at an angle of 1.36 rad . So, even though the electric field close to the central electrode is *less* isotropic with the addition of a negatively charged umbrella, the field lines are concentrated in the section

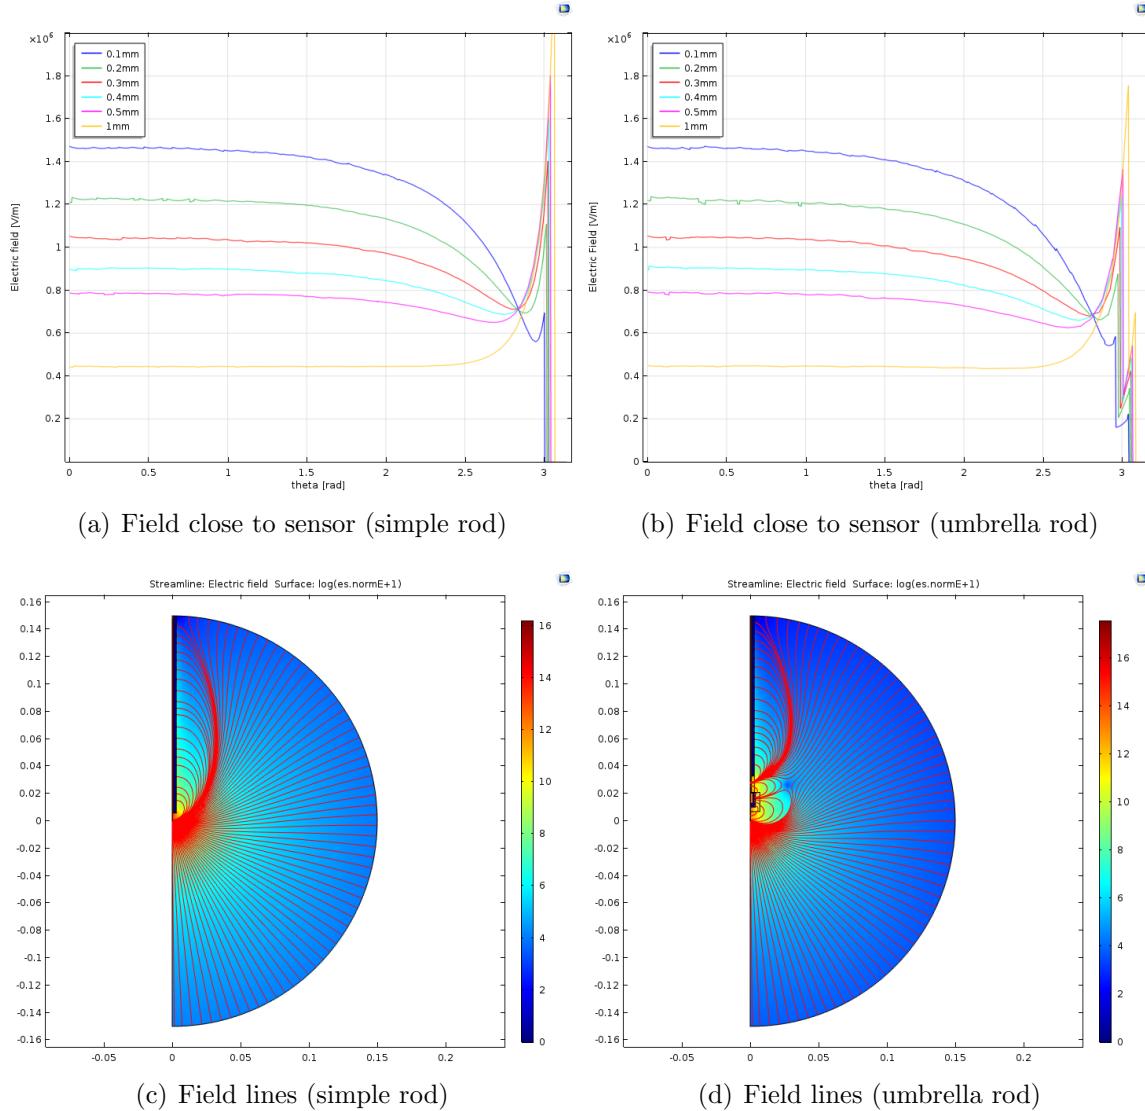


Figure 5.9: Comparison between COMSOL simulations of a detector with a simple rod, and one with a rod equipped with a secondary electrode close to the centre, called the “umbrella”.

of the electrode with a homogeneous field, paradoxically leading to an *improvement* of the energy resolution. This does come at the cost of an increase in “dead volume”, where released primary electrons are not collected by the central electrode, but this is almost always a reasonable trade-off.

5.4.3 Note on the effect of space charge

All the simulations done with COMSOL assume static conditions, and ignore the space charge effect from ions slowly drifting through the detector. While this is unlikely to be a problem during physics runs, when the event rate is low, we may expect calibration runs to exhibit different behaviour. Recall the equation C.18 for the static electric field under continuous ion formation rate of I_A in an ideal geometry:

$$E(r) = \frac{1}{r^2} \sqrt{\left(\frac{Q}{4\pi\epsilon_0}\right)^2 + \frac{I_A}{6\pi\epsilon_0\mu} (r^3 - r_{sensor}^3)} \quad (5.1)$$

This gives us an idea of when space charge effects start dominating over the field created by the central electrode. If we want the relative difference between the field with and without space charge to be less than a certain factor α at the edge of the detector (the point where space charge has maximum effect), we need

$$I_A < ((1 + \alpha) - 1) 6\pi\epsilon_0 \mu V_0^2 \frac{r_0^2}{R_{shell}^3 - r_{sensor}^3}$$

where $r_0 = R_{shell} r_{sensor} / (R_{shell} - r_{sensor})$. Taking $\alpha = 0.1$, and making an approximation for $R_{shell} \gg r_{sensor}$, we get:

$$I_A < 2.2 \cdot 10^6 \mu V_0^2 \frac{r_{sensor}^2}{R_{shell}^3}$$

where $[I_A] = \#ions/s$, $[\mu] = \text{cm}^2/\text{V/s}$, and $[R_{shell}] = [r_{sensor}] = \text{cm}$. For illustrative purposes, this corresponds to:

- For SEDINE, with a 30 cm radius shell and a 0.315 cm radius electrode with an applied voltage of 2520 V, with 3 *bar* of Neon ($\mu_0 = 7.5 \text{ cm}^2/\text{V/s} \cdot \text{bar}$), and assuming an avalanche gain of 7000 secondary charges per primary charge, we start seeing space charge effects at 66 Hz of 10 keV events.
- For SNOGLOBE, with a 67.5 cm radius shell and a 1.4 cm radius achinos with an applied voltage of 1600 V, with 1 *bar* of Neon ($\mu_0 = 7.5 \text{ cm}^2/\text{V/s} \cdot \text{bar}$), and assuming an avalanche gain of 5000 secondary charges per primary charge, we start seeing space charge effects at 5.7 Hz of 10 keV events.
- For a test detector, with a 15 cm radius shell and a 0.1 cm radius electrode with an applied voltage of 1200 V, with 0.2 *bar* of Argon ($\mu_0 = 1.5 \text{ cm}^2/\text{V/s} \cdot \text{bar}$), and assuming an avalanche gain of 5000 secondary charges per primary charge, we start seeing space charge effects at 51 Hz of 10 keV events.

Of course, the choice of 10 keV events is arbitrary, for the sake of getting numerical values for the rate. The necessary rate will be correspondingly higher (resp. lower) if the events are lower (resp. higher) in energy. It's the total number of secondary ions they indirectly create that matter, not their individual energies. Please note that the values given in the previous list are upper bounds on the ion formation rate so that we can neglect the effect of the space charge everywhere in the detector, and so we can trust COMSOL simulations. Due to the square root factor and the r^3 term in the electric field equation, we need a ~ 50 times higher ion formation rate for the space charge field to dominate in the outer half of the detector.

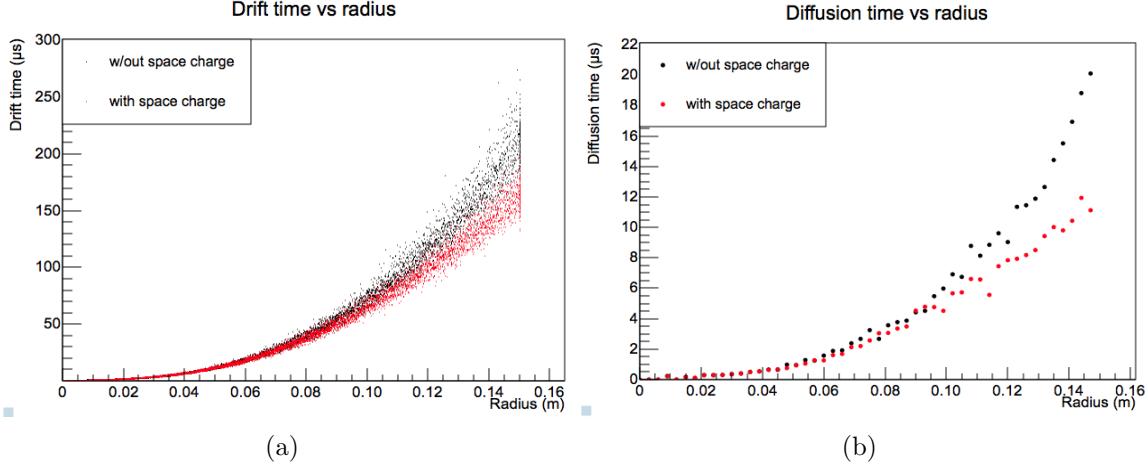


Figure 5.10: Comparison of drift and diffusion times with and without space charge, for a test detector with a 15 cm radius shell and a 0.1 cm radius electrode with an applied voltage of 1200 V, with 0.2 bar of Argon ($\mu_0 = 1.5 \text{ cm}^2/\text{V}\cdot\text{s} \cdot \text{bar}$), and an avalanche gain of 5000. The ion formation rate is given by 500 Hz of 10 keV events.

With a high enough ion formation rate, the electric field in the drift region of the detector will be dominated by the contribution from the space charge. In those situations, the results from COMSOL simulations are not appropriate, unless we are just interested in upper value constraints on the drift and diffusion time. Fig 5.10 shows the decrease in drift and diffusion times when applying a sizeable ion formation rate. The analytical formula for the electric field with space charge can be used for the electron drift simulation, but then any anisotropic effect from the presence of the rod or the umbrella won't be taken into consideration. As such, calibration runs with high event rates remain difficult to compare to simulations.

5.5 Integration of results from Geant4 simulations

For the simulations of the background rate for our detector, it was also important to be able to use the results produced by A. Brossard with Geant4 [60–62] in the pulse simulation code. Geant4 is a Monte Carlo based toolkit to simulate the passage of particles through matter. Particles are propagated through the various levels of shielding of the detector, then through the gas inside, and their interactions are recorded with their position, energy deposited, mother particle, and so on. For more information on the nature of these simulations as applied to NEWS-G detectors, please see A. Brossard’s thesis [35].

The files where this information is recorded are then used as the basis for events in our run simulation code. The position and energy deposition of all interactions within a short time window (~ 1 ms) are gathered, then the corresponding primary electrons generated by those interactions are drifted, and their arrival times used to generate a pulse. An example can be seen on Fig. 5.11.

5.6 Axion simulations

The only difference between axion simulations and the rest is the nature of the energy deposition in the detector, and their energy distribution. The latter is rather straightforward, but the former requires further explanations as to the behaviour of the created photons. Just as for the previous sections, the focus here will be on the technical aspects of the simulation, but some time will also be spent on the conclusions drawn for choosing optimal running conditions, expanded upon in Sec. 8.3.

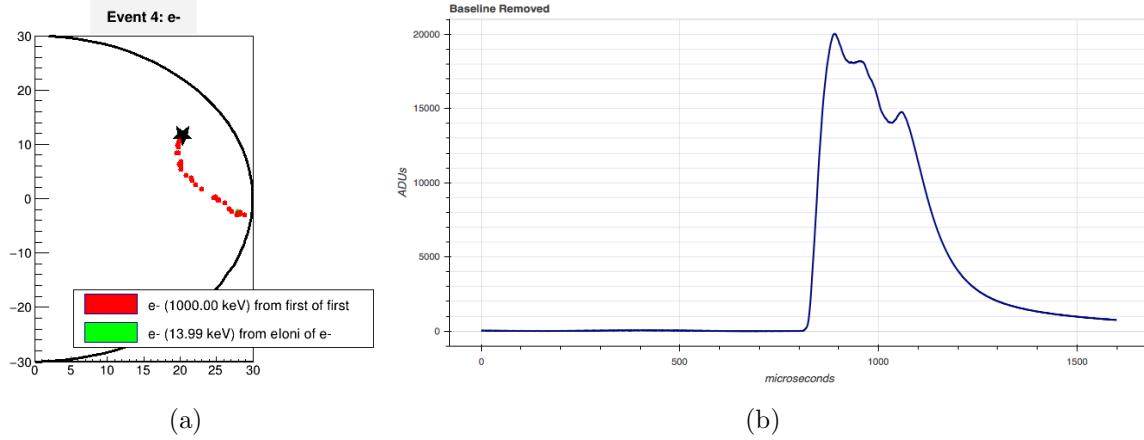


Figure 5.11: Geant4 simulation of a 1 MeV electron inside the detector. Left: Visualization of its energy deposition inside the detector, with the star marking the starting position; while the track of the electron was relatively straight, the projection of its path into cylindrical coordinates introduces an artificial “bend” in its trajectory. Right: Pulse generated from this event.

5.6.1 Method

Since solar KK axion are decays, rather than interactions, we expect the events to be uniformly distributed in the volume of the detector. An axion decay will generate two photons in opposite directions, isotropically. The mass/energy spectrum of KK axions is taken from [9] and normalized by the expected solar KK axion density on Earth, as can be see on Fig. 5.12. For an axion event simulation, a random mass will be taken from this distribution, and the created photons will be given half the mass of the axion as their energy.

The position of the decay is chosen randomly the same way as for bulk pointlike simulations:

$$r = R_{shell} u_1^{1/3}$$

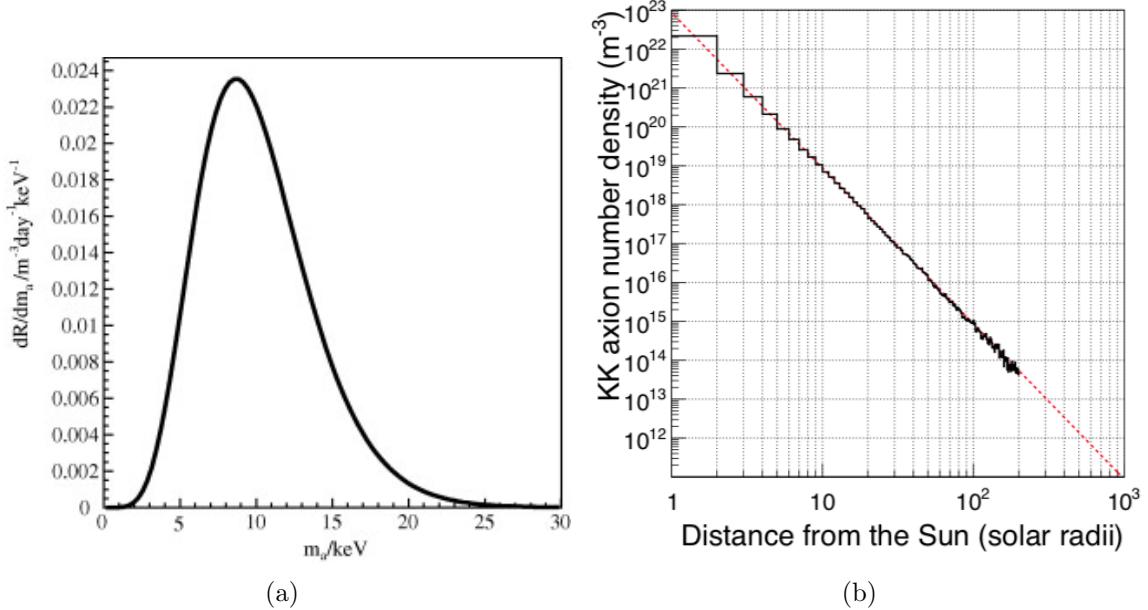


Figure 5.12: Left: decay spectrum of solar KK axions on Earth, for an assumed axion local density of $1.0 \cdot 10^{14} \text{ m}^3$ [9]. Right: number density of solar KK axions depending on the distance from the Sun; the red fitted r^{-4} curve allows extrapolation to the distance from Earth to the Sun, $215 R_\odot$, for a local density of $4.075 \cdot 10^{13} \text{ m}^3$ [10].

$$\theta = \arccos(2u_2 - 1)$$

where u_1 and u_2 are two independent random draws from a uniform distribution between 0 and 1. This is then transformed into Cartesian coordinates:

$$x = r \sin(\theta)$$

$$y = 0$$

$$z = r \cos(\theta)$$

where $y = 0$ can be chosen without loss of generality due to the cylindrical symmetry of the setup. The direction of one of the produced photons is generated by drawing three times from a normal distribution to produce an isotropic result:

$$\overline{x_{ph}} = n_1/N$$

$$\overline{y_{ph}} = n_2/N$$

$$\overline{z_{ph}} = n_3/N$$

where $N = \sqrt{n_1^2 + n_2^2 + n_3^2}$ is the normalization constant. The direction for the second photon is just the opposite one. Then, knowing the mean absorption length λ of the photon for a given energy (half the mass of the decaying axion for each), we generate the actual absorption lengths l_1 and l_2 from an exponential distribution with mean λ . This gives a final position for both interactions of:

$$x_1 = r \sin(\theta) + l_1 \overline{x_{ph}} \quad x_2 = r \sin(\theta) - l_2 \overline{x_{ph}} \quad (5.2)$$

$$y_1 = +l_1 \overline{y_{ph}} \quad y_2 = -l_2 \overline{y_{ph}} \quad (5.3)$$

$$z_1 = r \cos(\theta) + l_1 \overline{z_{ph}} \quad z_2 = r \cos(\theta) - l_2 \overline{z_{ph}} \quad (5.4)$$

which we convert back into cylindrical coordinates for use with the electron drift simulation code, assuming either position is still within the bounds of the detector. Note that while an axion could decay inside the detector but still have both photons exit it before interacting, the fact that both photons travel in opposite directions means it's not possible for an axion to decay outside the detector yet have both created photons interact within. As such, we only need to simulate axion decays

inside the detector.

5.6.2 Photon absorption length

The photon absorption length is computed from the NIST X-ray absorption coefficient database [63]. The absorption length of a photon of a given energy in a pure gas is $\lambda = 1/(\mu_{en}/\rho)/\rho$, where ρ is the mass density of the gas. For gas mixtures, the effective μ_{en}/ρ is built from the NIST data with the relationship:

$$\mu_{en}/\rho = \sum_i (\mu_{en}/\rho)_i \frac{PP_i \rho_i}{\rho}$$

where $(\mu_{en}/\rho)_i$ is the absorption coefficient of gas element i , PP_i is the partial pressure of gas element i , ρ_i is its density at standard temperature and pressure, and ρ is the total density of the gas in the experimental conditions. The absorption length of the photon will then depend on the energy of the photon, as do the absorption coefficients shown on Fig. 5.13. The use of the lower absorption coefficients instead of the attenuation coefficients leads to a conservative underestimation of the expected number of double events in the detector⁴. For the run conditions of SEDINE, using Neon with 0.7% CH₄ at 3.1 bar, the absorption coefficients are dominated by the contribution from the Neon. The absorption lengths for the range of energies considered are shown on Fig. 5.13 too. The lack of data under 1 keV effectively limits our search for axions to masses above 2 keV, but the impact on our exposure is negligible.

⁴To see this, consider three categories of events: one in which a photon is completely absorbed, one in which a photon escapes the detector, and another in which the photon interacts without leaving all its energy in the gas (or interacts multiple times before doing so). By taking only the lower absorption coefficient, I am effectively setting all events in the latter category into the second one, ignoring them completely. As such, the total number of expected axion events will be lower than the real one, giving conservative results.

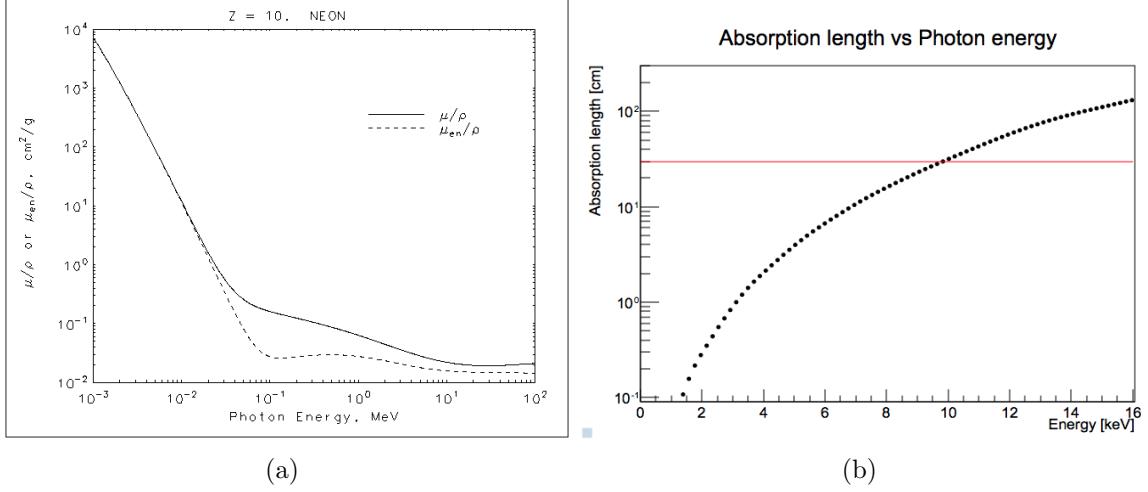


Figure 5.13: Left: X-ray mass attenuation coefficient μ/ρ and mass absorption coefficient μ_{en}/ρ in Neon [63]. For photons under 10 keV, the difference between μ/ρ and μ_{en}/ρ remains under 5%; the difference jumps to $\sim 25\%$ at 20 keV. Right: Absorption length of photons under SEDINE running conditions; the radius of the detector is shown with the horizontal red line.

For a given absorption length and detector radius, the chances of both photons being captured is a purely geometric consideration that depends only on the ratio λ/R_{shell} . While an analytical formula does not exist, a simulation with the ideal geometry is pretty straightforward; the results are shown on Fig. 5.14. This gives a containment probability of 0.5 at $\lambda/R_{shell} = 0.39$, and 0.1 at $\lambda/R_{shell} = 1.55$, quickly decreasing as $\sim 1/\lambda^2$ past that point. Since the peak of the solar KK axion decay rate is at ~ 8 keV, we want the attenuation length of 4 keV photons to be shorter than $\sim 0.4 R_{shell}$. The attenuation length is inversely proportional to the gas pressure, so this can be achieved easily by increasing the pressure.

However, increasing the pressure will in turn decrease the attenuation length at low energies, making both photon interactions hard to separate. This effect is much harder to relate to a single variable: the minimum radial distance between both

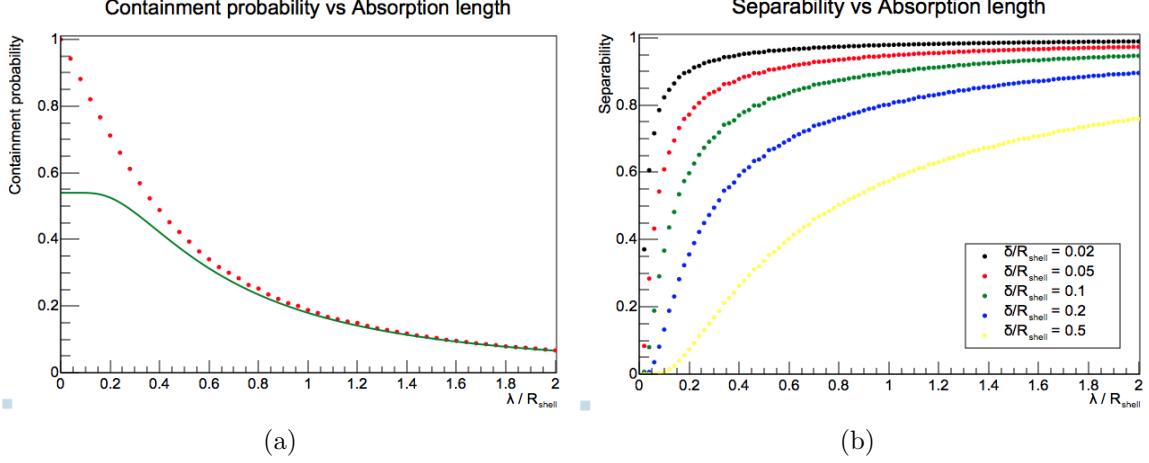


Figure 5.14: Left: In red, the probability of both photons being absorbed within the detector for a given attenuation length; in green, $0.54(1 - e^{-1/x^2})$, which fits the containment probability increasingly well after $\lambda/R_{shell} > 1$. Right: separability of both photons, making the simplifying assumption that two photons are separable if their radial position differs by more than a fixed δ ; plots for various values of δ are shown.

interactions needed to separate them depends on the position of the interactions, since the drift time is not proportional to the radial position. Moreover, the drift time - radial position dependency changes in a non-trivial manner with gas pressure. As such, any optimization of experimental setup has to be done through comparing full simulations in varying conditions. Still, some qualitative conclusions can be made if we make the simplifying assumption that two photons are separable if and only if their radial position differs by more than a fixed distance δ . The results from simulations for various values of δ are shown on Fig. 5.14. They suggest that a separability of $\sim 50\%$ is reached when $\lambda > 1.5\delta$. So, assuming for example that $\delta = 0.1 R_{shell}$, and to keep the peak of the solar KK axion decay rate above this cut-off point, we want the attenuation length of 4 keV photons to be larger than $\sim 0.15 R_{shell}$.

Looking back at Fig. 5.13, we see that in the running conditions of SEDINE, a

4 keV photon had an attenuation length of 2 cm, or $\sim 0.07 R_{shell}$. This is a shorter attenuation length than what would be optimal, unless we are overestimating δ . By decreasing the pressure in the detector by a factor of up to 2, we would improve the separability of axions with energies around the maximum of the distribution, and hence increase our sensitivity. Still, for a physics run taken in a completely different context (search for WIMPs), the conditions are more than sufficient to set first constraints on solar KK axions. For a more quantitative optimization of the running conditions for SNOGLOBE, full simulations have to be run and compared to each other.

Summary

In this chapter, we have described the multiple steps involved in the creation of a signal in our detector, together with the tools we have used to simulate them. They can be categorized roughly into four big areas:

- Energy deposition within the detector (either GEANT4 or *ad hoc* code);
- Electric field configuration (COMSOL);
- Electron drift within that field (both Magboltz and *ad hoc* code);
- Pulse formation (*ad hoc* code).

Altogether, a run can be fully simulated in any configuration of (cylindrically symmetric) SPC, and then treated in the same way as real data. But to use these simulations to estimate the efficiency of our detector to axions, they must first be calibrated, as will be shown in the next chapter.

Chapter 6

Calibrating the simulations

In an ideal world, we would be able to understand everything about our detector just from an engineering drawing, working our way up from first principles, and maybe running some simulations. Unfortunately, that's not the world we live in, so we need to calibrate the detector to verify that our observations match our expectations. In this section, I will describe

- the different unknown physical parameters that affect SPCs;
- what theory and/or simulations predict;
- which calibrations can verify those predictions;
- whether (and how) we can alter our models to match the calibrations, if relevant.

I will mainly cover calibrations that were used for the SEDINE physics run. Other calibrations used within the collaboration that were unavailable for one reason or another will only be touched upon for context. For more information on the available calibration methods themselves, refer to Sec. E.0.2. The effect on axion search results of uncertainties in SEDINE calibrations will be covered in Sec. K.

6.1 Electron drift time

We don't have direct access to any measure that would inform us on the speed of electrons drifting in the SEDINE gas detector. Normally, this sort of measures are done through knowledge of both the beginning (t_0) and the end of the drift process of an electron, with their associated position and time. For example, a TPC uses the time difference between the scintillation and ionization signals to find the drift time, and the diffusion in the ionization signal to determine the drift distance. With an SPC, we only have the ionization signal, so we cannot use this technique. Instead, we must find other ways to find the starting time and position of the drift.

6.1.1 Laser data

The first and most reliable method is to use a pulsed laser. By shining a $\lambda = 213\text{ nm}$ pulsed laser on the inner surface of the detector's shell, photoelectrons can be extracted on demand. Better yet, they are at a known position (the inner surface of the shell) and at a known time (laser sends a trigger signal before each pulse). This provides one calibration point: the drift time of electrons from the surface to the sensor. An example of one such calibration event is shown in Fig. 6.1.

Unfortunately, while this technique has been used multiple times to test the match between calibrations and simulations in multiple detectors, SEDINE did not benefit from a laser setup. As such, other techniques had to be developed to calibrate the drift speed of electrons.

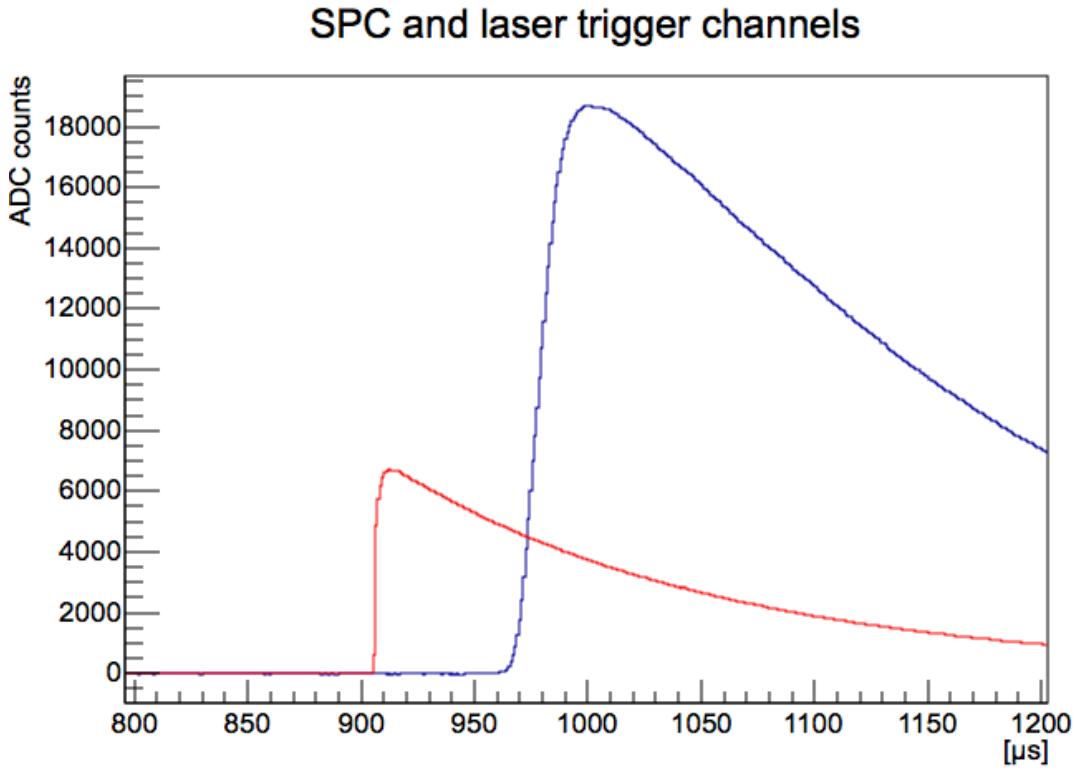


Figure 6.1: Juxtaposed SPC (blue) and laser trigger (red) channels. The delay between the 50% point of the rise in the SPC channel signal with respect to the beginning of the laser trigger signal gives the drift time of electrons from the surface to the central electrode.

6.1.2 Track events

In the absence of a pulsed laser, obtaining a calibration point for the drift time is still possible, if trickier. The measure relies on track events that start at the inner surface of the detector and pass next to the central electrode. For such an energy deposition, the electrons released closest to the sensor will induce a signal almost instantly, providing a t_0 of sorts for the drift time of the electrons released close to the surface. In other words, the maximum time length of a track will be given by the drift time of surface electrons, plus a term that depends on the spread in arrival times

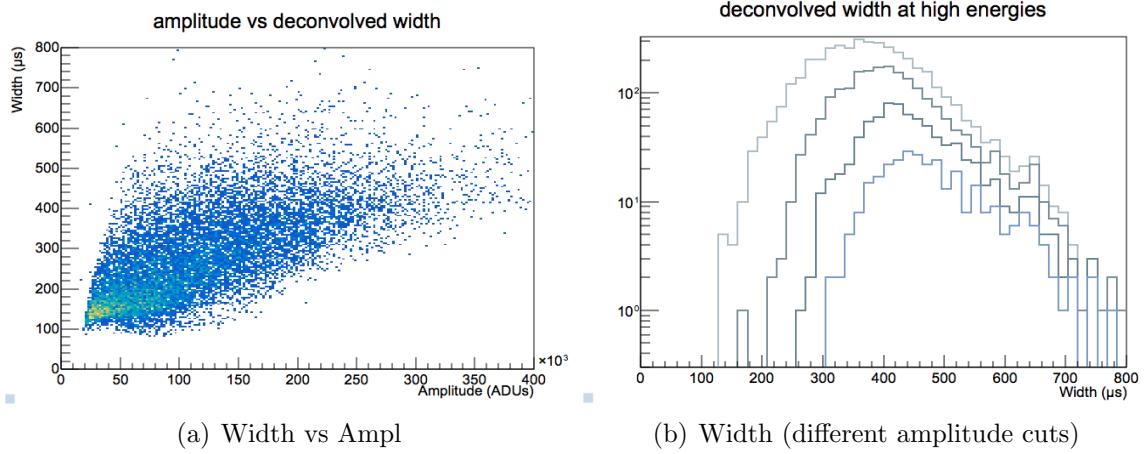


Figure 6.2: Distribution of widths for high width events in the background run, after cutting out alpha events. The amplitude cuts on the right plot are, from top to bottom, > 150 kADUs, > 200 kADUs, > 250 kADUs, > 300 kADUs. We can see that the maximum width depends on energy, but seems to be reached somewhere in the $200 - 250$ kADUs range and above.

of those electrons. Since the “end” of a distribution is not usually a feature that is easy to pinpoint, we can look instead at the point where the pulse width distribution starts decaying in number. The distribution itself can be seen in Fig. 6.2.

We can see that the maximum width depends on the energy. While this is not completely unexpected, since the time difference between the first and last electron from a given energy deposition will tend to increase with the number of primary electrons generated due to their diffusion, this dependency is higher than what would be expected from this effect alone. The simple explanation is just that at lower amplitudes, these events are not energetic enough to actually cover the whole range of distances from the sensor. The widths seem to stabilize somewhere above the $200 - 250$ kADUs range and above. To be conservative about the uncertainty of this measure, we will take the “maximum” width of an event in this configuration to be between $398.0 \mu s$ and $446.8 \mu s$, which are the results from fitting a gaussian on the

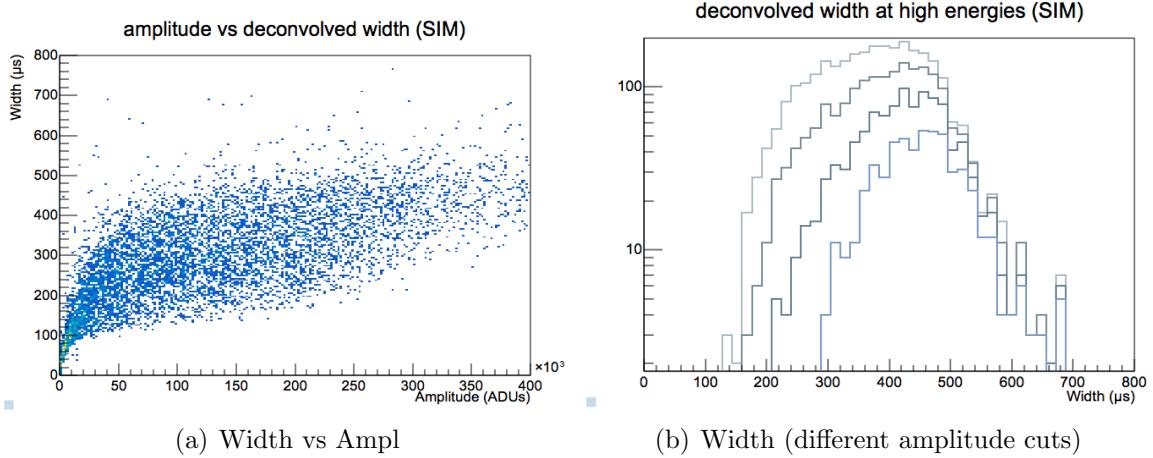


Figure 6.3: Distribution of widths for high width events in the simulated track dataset. The amplitude cuts on the right plot are, from top to bottom, > 150 kADUs, > 200 kADUs, > 250 kADUs, > 300 kADUs. We can see that the width at the maximum of the distribution depends on energy, but stabilizes somewhere in the $200 - 250$ kADUs range and above.

width distribution above 200 kADUs and 300 kADUs, respectively.

Due to the complexity of this measure, a straight interpretation of the width of these events as the drift time is not appropriate. Instead, a simulation of tracks produced by 1000 keV electrons generated in the bulk of the detector was done with Geant4, then fed into the pulse generator simulations. The same processing as the one described for the background run was performed on this simulated dataset, giving the results shown on Fig. 6.3.

These show a much sharper drop-off in event count past a certain point around $450 - 500 \mu s$, close to the expected surface drift time of $460 \mu s$. We also see a dependency of the drop-off point based on the amplitude of the event. Fitting the different width distributions based on amplitude cuts gives us a maximum in the width distributions between $395.7 \mu s$ and $449.9 \mu s$, very close to the measure from the actual

data.

However, since we are not doing a direct measurement of the drift time, but rather of a derived value (the pulse width), and comparing the results from data and simulations, the uncertainty of this process will be a combination of the uncertainty of both measures. Since we don't have a real measure either so much as a range of values, we must make the assumption that any value in that range is equally likely, leading to a higher uncertainty than if we were assuming the likelihood of each value to follow a gaussian distribution with similar width ranges. This gives a data measurement of $422.4 \pm 24.4 \mu s$, and a simulation measurement of $422.8 \pm 27.1 \mu s$. While the averages are remarkably similar, the correction factor due to the ranges in the measures still has a variance of 5%, as shown in Fig. 6.4.

6.1.3 Muon veto

A potential method to get more calibration points for the drift time has been put forward, but not tested at the time of writing. Detectors on the surface are constantly crossed by cosmic muons. By placing a muon veto with a good enough position discrimination (e.g. a Micromegas) above and below the detector, both the time of crossing and the path of the muon through the detector can be determined. In turn, the point of closest approach to the sensor can be computed. This means that we have both a t_0 and a position for the beginning of the drift. While this approach would likely suffer from some similar precision issues than the track method discussed above, the ability to have measurement points for potentially *any* position inside the detector, rather than just the surface, would make it extremely useful.

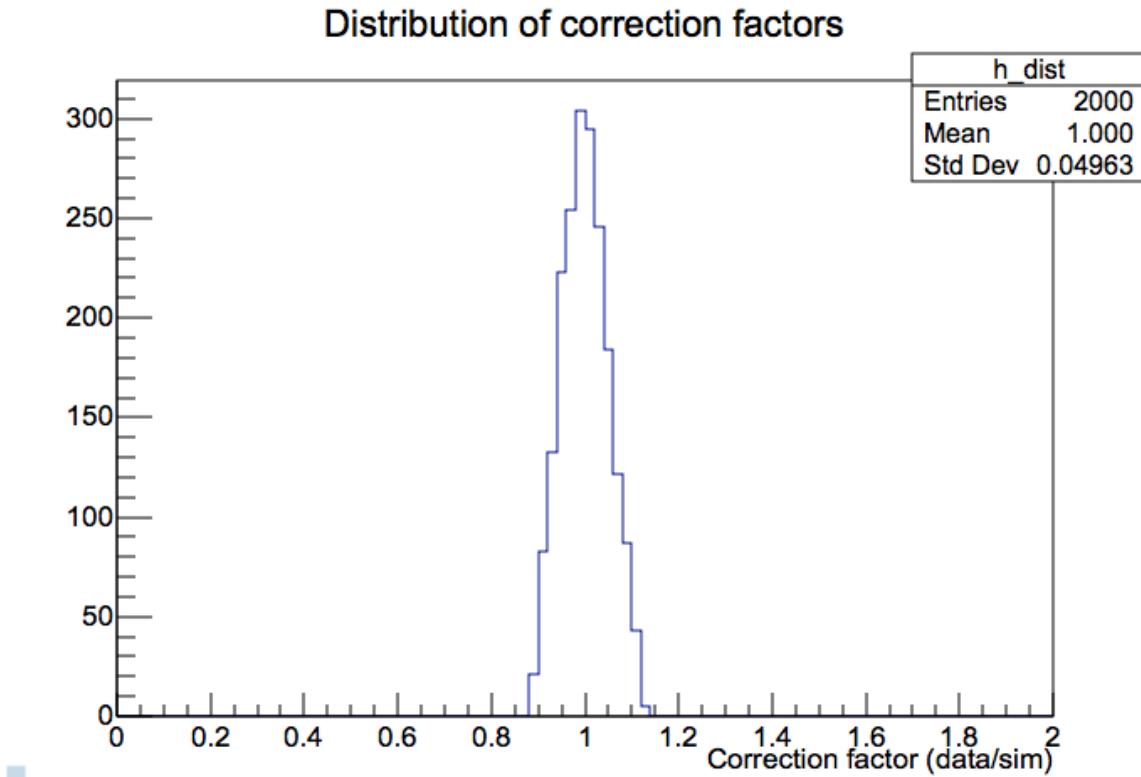


Figure 6.4: Distribution of correction factors between data and simulations when the values of the measures for the maximum width of either is left to vary uniformly along their whole range.

6.1.4 Limits of approach

The main issue with the calibrations we have access to is that we have a single measurement point (the surface), while we are interested in knowing the drift time from any point inside the detector. This makes us reliant on simulations to find the values for all intermediate points. This could fail if either the electron drift data (from Magboltz) or the electric field (from COMSOL) are wrong.

An illustration of the problem is shown on Fig. 6.5. Based on SEDINE simulations, the drift time can be approximated by the function $t = t_{max} \left(\frac{r}{r_{max}} \right)^\alpha$, with $t_{max} = 470 \mu\text{s}$, $r_{max} = 30 \text{ cm}$, $\alpha = 2.5$. Let us assume that, even if the simulations are wrong,

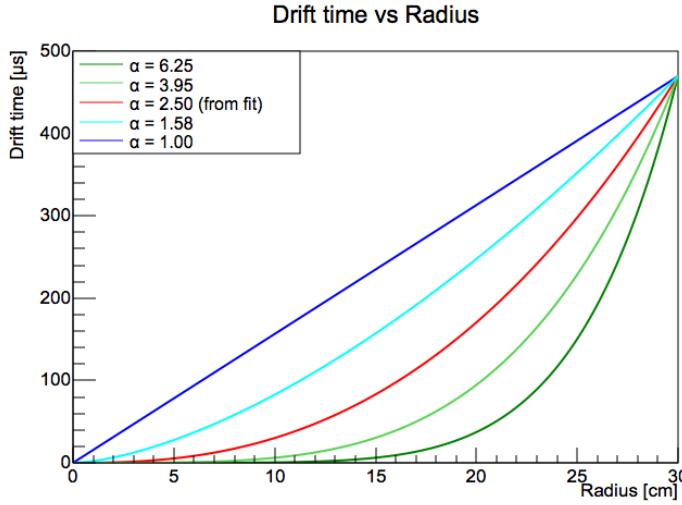


Figure 6.5: $t = t_{max} \left(\frac{r}{r_{max}} \right)^\alpha$, for different values of α . For $t_{max} = 470 \mu\text{s}$, $r_{max} = 30 \text{ cm}$, $\alpha = 2.5$, this reproduces the drift time vs radius curve obtained through Magboltz simulations.

the drift time has to follow this type of relationship. Then r_{max} is fixed by the detector size, and t_{max} is given by the calibration of the surface drift time, but we do not have a good way to constrain α . An attempt to characterize this effect in the context of axion searches is described in App. K.2.

6.2 Electron drift diffusion

Many of the points presented for the electron drift speed are just as valid for the calibration of their spread, and so will not be repeated. The main difference between the two is the slightly increased complexity of the relationship between simulation data and calibration data.

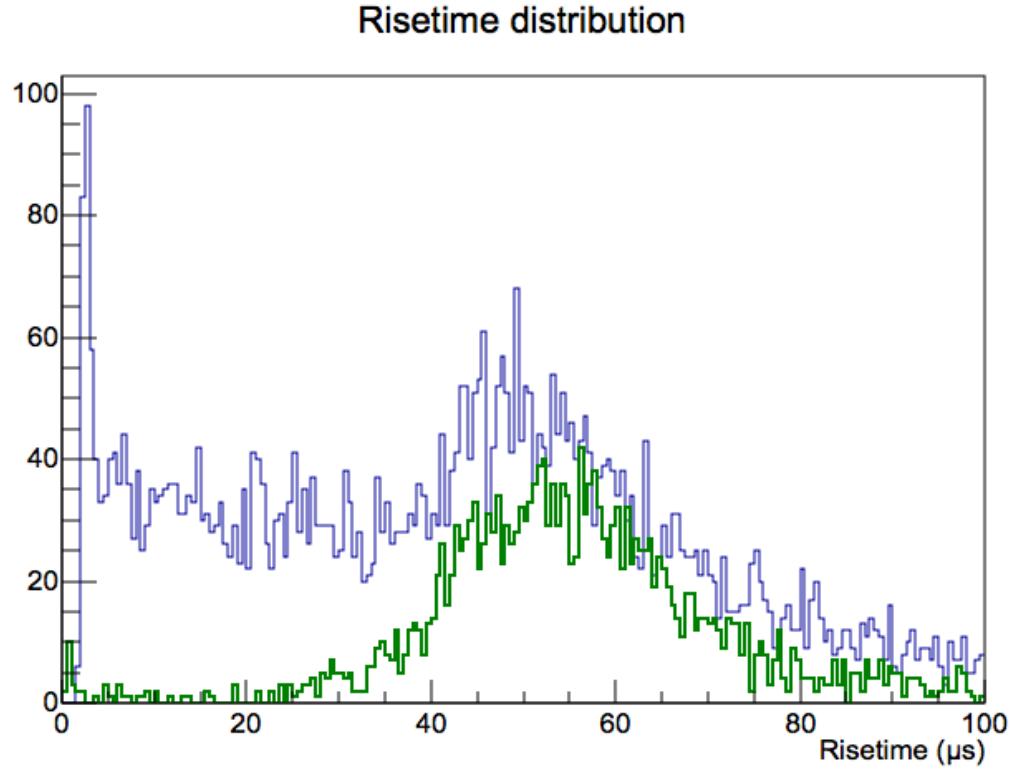


Figure 6.6: Risetime of events in the $2 - 10$ keV energy range for the SEDINE physics data (blue), compared with simulations of electrons in same energy range generated at the surface (green). Simulations normalized to have a similar height as the data.

6.2.1 Surface events

Due to the presence of impurities on the shell of the detector, a larger proportion of events are generated there than in the bulk. In the context of rare event searches, we would use the estimation of pulse risetimes to reject this background; however, they also offer a natural calibration of the spread in arrival time of primary electrons coming from the surface. Alternatively, a laser can be used to shine on the inner surface of the shell, reproducing this population while discriminating against non-surface events, but the principle remains the same.

The results from simulations give a surface diffusion time slightly higher than the data: $53.0 \pm 1.1 \mu\text{s}$ instead of $50.0 \pm 0.6 \mu\text{s}$, both numbers obtained by fitting a gaussian over the peak. A corrective factor has to be applied to the longitudinal diffusion generated by Magboltz, to obtain the correct diffusion times in the simulations. Since the uncertainty from the fits is smaller than the correction between simulations and data, we take the conservative approach of using the magnitude of the correction as the uncertainty of our simulations (relative error of 6% on diffusion times).

Based on the relation between diffusion time and Magboltz parameters described in Sec. 5.3:

$$\sigma_{diff} \propto \sqrt{\frac{DL}{v^3}} \quad (6.1)$$

where DL is the longitudinal diffusion, and v is the electron speed obtained from Magboltz. Then, if we want to modify the diffusion times from the simulations by a factor of α_σ , while also modifying the electron speed by a factor of α_v (to simultaneously match drift time calibrations), the longitudinal diffusion from Magboltz needs to be modified by a factor of:

$$\alpha_{DL} = \alpha_\sigma^2 \alpha_v^3 \quad (6.2)$$

6.3 Avalanche statistics

The number of secondary charges produced in the avalanche region is directly linked to the size of pulses we record with our detector. What we extract from an event is its amplitude in ADUs (Analog-to-Digital Units), which is not a physical quantity. We can, however, convert that into a number of charges:

$$1[e-] = G_{\text{preamp}}[V] = G_{\text{preamp}} G_{\text{dig}}[\text{ADU}]$$

where G_{preamp} is the voltage outputted by the preamplifier when one elementary particle goes through it, and G_{dig} is the number of ADUs recorded by the digitizer for a 1 V signal. For more details about these values, please see the App. B. Assuming the processing used to compute the amplitude of the event is taking care of the ballistic deficit (cf. Sec. ??), this means we can get the number of secondary charges created:

$$Q_{\text{secondary}} = A_{\text{event}} / (G_{\text{preamp}} G_{\text{dig}})$$

where A_{event} is the amplitude in ADUs of the event. In turn, for a given energy deposition, the number of secondary charges created is:

$$Q_{\text{secondary}} = E_{\text{event}} / W * G_{\text{av}}$$

where W is the mean ionization energy of the gas mixture, and G_{av} is the charge gain due to the avalanche. I will ignore for the time being any effect from electron attachment in the drift region (or rather, absorb it into G_{preamp}). Putting both relationships together, we get that

$$E_{\text{event}} = A_{\text{event}} \frac{W}{G_{\text{preamp}} G_{\text{dig}} G_{\text{av}}} \quad (6.3)$$

Since we can calibrate our electronics to find G_{preamp} and G_{dig} , and assuming for now that W is known (more on that later in this section), this means that we only have one unknown: G_{av} , the mean number of secondary ion-electron pairs one primary electron creates in the avalanche region. We can then measure this value if

we have a source producing events at a fixed, known energy.

Given that the avalanche process strongly depends on the electric field around the sensor, the gas composition, and the pressure, G_{av} may change wildly from one run (i.e., data taken in any given set of operating conditions) to the next. Emptying and refilling the vessel, in particular, is known to change this value. This means that each new run usually requires a new energy calibration. We now move on to some of those calibrations.

6.3.1 Copper fluorescence

During our background run, no calibration source was introduced in the detector. However, the fluorescence of the copper that constitutes the detector shell provided a natural calibration source. We expect a monoenergetic source of fully absorbed X-rays at 8.05 keV from copper atoms excited by higher energy γ s. The results are shown on Fig. 6.7¹. While the fact that it wasn't a proper calibration source with a high rate does mean that it is harder to separate from other backgrounds, after appropriate cuts in energy and risetime, the peak rises clearly above the background.

Since the peak appears at 16820 ADU, this gives us a conversion factor between ADU and eV of $C = 2.077 \text{ ADU/eV}$. Note that this is for the normalizations done with the Double Deconvolution method (cf. Sec. 4.1.4); other analysis methods will likely give other conversion factors.

Using the gain of the Canberra preamplifier [49] used for this run, $G_{preamp} = 235 \cdot 10^{-9} \text{ V/e}^-$, the effective gain of the Calibox digitizer (cf. App. B), $G_{dig} = 48120 \text{ ADU/V}$, and the mean ionization energy of Neon, 36 eV [29], this gives us a

¹In other runs that did not use the shielding, exposing the detector to a source of γ s above that energy did in fact increase the rate of this population, confirming its nature as copper fluorescence.

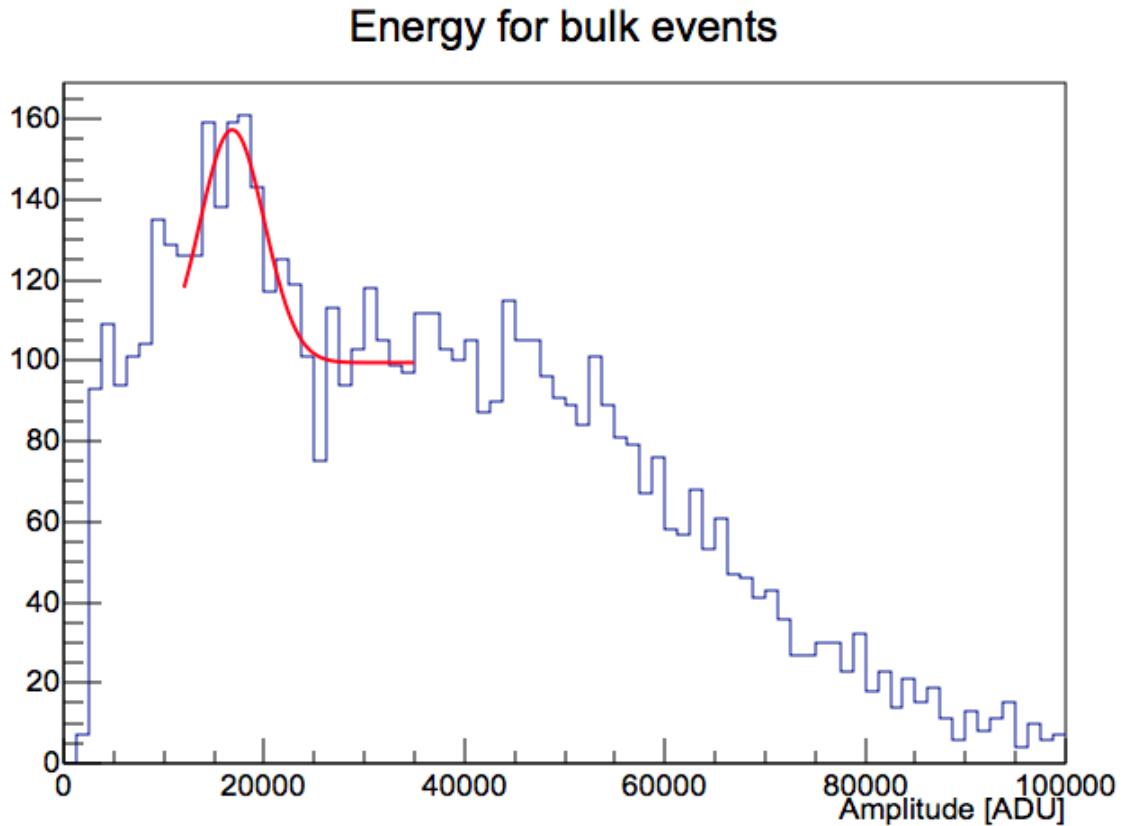


Figure 6.7: 8.05 keV peak from copper fluorescence during the SEDINE physics run. The fit results give a peak at 16820 ADU, with an uncertainty of 2.9%.

mean avalanche gain of 6610 secondary ion-electron pairs per primary electron.

6.3.2 ^{37}Ar

Similarly to the copper fluorescence, ^{37}Ar provides monoenergetic lines, one at 2.82 keV and another at 270 eV. ^{37}Ar was pumped into SEDINE right after the SEDINE physics run, without otherwise changing the gas or voltage. Unfortunately, the gain found with this calibration turned out to be only 2/3 of that found from the copper fluorescence during the physics run. This suggests that the act of filling in the detector with the calibration source also introduced contaminating elements, producing a

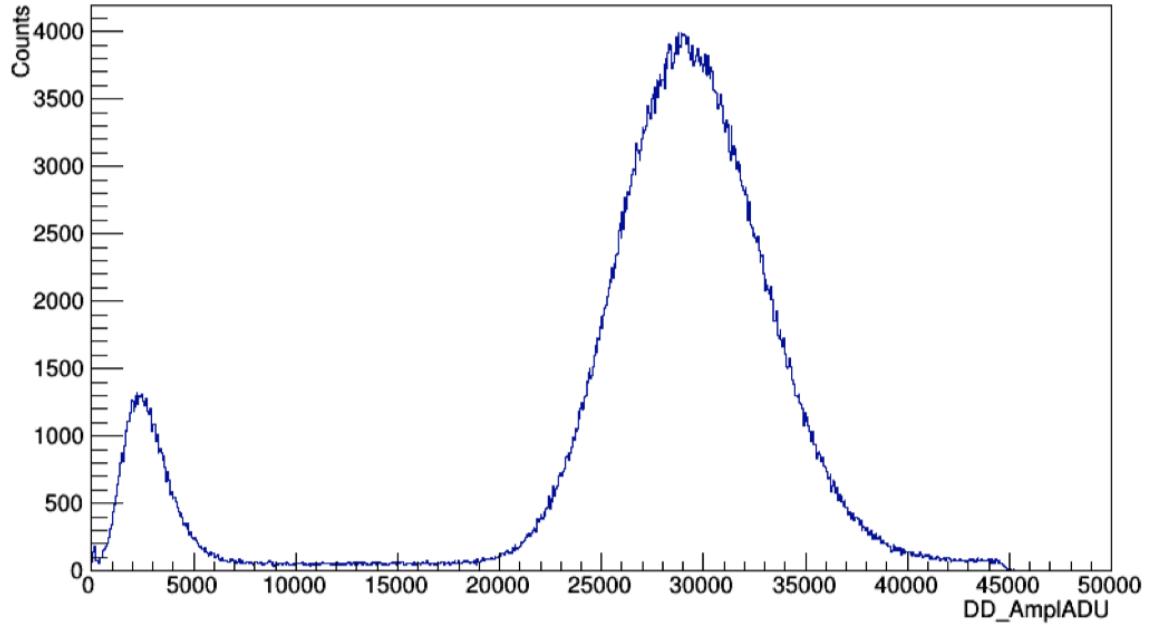


Figure 6.8: Amplitude of Ar^{37} events in a calibration run with the Queen’s S30 detector. Both the 2.82 keV and 270 eV peaks are clearly apparent.

drop in gain.

While we could not use that run for gain calibration, this method has been used to great success with other SPC detectors, with excellent resolution, as shown on Fig. 6.8. Together with the 0 eV / 0 ADU point, the double monoenergetic lines allow confirmation of the linearity of the energy response of the detector at low energies.

6.3.3 Laser calibration

While no laser was available for SEDINE, relevant data for the physics run could still be gathered from the work done in other detectors; namely, information on the behaviour of the avalanche process (and on the mean ionization energy, see next section). Knowing the mean value of the avalanche only gives us half the information about the process. The shape of the distribution of the number of secondary charges

created also depends on parameter θ , from the Polya distribution:

$$P(\nu) = \frac{(1+\theta)^{1+\theta}}{\Gamma(1+\theta)} (\nu)^\theta \exp(-(1+\theta)\nu)$$

where $\nu = \frac{n}{\bar{n}}$ is the ratio between the number of created secondary charges and \bar{n} , the average of the distribution. The difficulty with finding a measure for θ is that the dispersion in the number of secondary charges created is large. Two events with a different but similar number of primary electrons may have end up with same amplitude. And since the distribution of amplitudes for an event depends not only on θ , but also on the number of primary electrons, it is almost impossible to get a proper measure using standard radioactive measures.

One way around the issue is to generate events with a single primary electron. If a process to generate these events is found, then the distribution of amplitudes from those events will follow the relevant Polya distribution (ignoring any spread due to electronic noise and so on). A fit can then be performed to find the values of both θ and $G_{avalanche}$.

The way we did this measure in one of our Queen's prototype sensors was to use a laser, as described in [38]. By shining a $\lambda = 213$ nm pulsed laser on the inner surface of the detector's shell, photoelectrons could be extracted on demand. A variable attenuator in front of the laser allows us to reduce the intensity of the laser until mostly only 0 or 1 photoelectrons were created per pulse. A trigger signal coming from the laser allows us to identify these events. While this setup still occasionally creates events with more than one electron, the odds of that can be made very low.

Then, while events cannot be individually tagged as having 0, 1, 2 or more electrons, the total distribution can be fitted with a formula whose free parameters are

G_{av} the mean gain of the avalanche, θ the shape factor of the distribution, σ the RMS of the baseline, and μ the average number of photoelectrons extracted per pulse that reach the avalanche region, assuming they follow a Poisson distribution.

It is important to note that this fit is completely independent from the effects of electron attachment, or the mean ionization energy of the gas mixture. Only the number of primary electrons reaching the avalanche region matters. Furthermore, the laser can be turned on and off as necessary, allowing calibrations to be performed without changing anything about the operational setup. As such, this is an extremely powerful calibration technique. Its only downside is that it requires an appropriate laser and a detector with a fibre feedthrough, which is comparatively more complicated and expensive than some of the calibrations with radioactive sources. In particular, the SEDINE detector did not have access to laser calibrations.

6.4 Ionization energy

As referred to on the previous Sec. 6.3, a difference in the mean ionization energy W can be “absorbed” by the avalanche gain G_{av} without changing any physical observation. Most physical observations only depend on the ratio between the two. Technically, changing the average number of primary electrons μ while keeping the average number of secondary changes (μG) constant would lead to slightly better relative resolution for the event with more primary electrons. In practice, the effect is too small to be observed in most circumstances.

This limitation can be avoided if we have access to both a laser calibration and a fixed energy source calibration. Indeed, as previously mentioned, the laser calibration can be used to determine G_{av} independently of W . Once that is known, a known

energy source can be used to get W by rearranging the formula 6.3:

$$W = \frac{A_{event}}{E_{event}} \frac{1}{G_{preamp} G_{dig} G_{av}}$$

again, assuming that there is no attachment for the data from the known energy source, and that the ballistic deficit is properly accounted for.

More details on the procedure to find this value with NEWS-G detectors can be found in [38], which showed a mean ionization energy of 27.6 eV in neon with 2% CH₄, as opposed to the literature value in pure neon of 36 eV [29].

6.5 Attachment

The last parameter related to the drift of electrons is the chance that one of them will be captured by a particle inside the gas before reaching the avalanche region. If all primary electrons were equally likely to be captured, then this would not be much of an issue, since it would be almost equivalent to having a lower avalanche gain. However, since attachment chance increases with time, electrons that drift from farther away will be more likely to be captured. This leads to a risetime dependency of the amplitude, beyond what we would expect from the ballistic deficit alone, degrading our energy resolution.

To measure the strength effect, we need a source of fixed energy that produces events everywhere inside the detector.

6.5.1 ^{37}Ar

Adding some ^{37}Ar to the detector provides a monoenergetic calibration at 2.8 keV to study the dependence of amplitude with changing risetimes. Such a calibration was

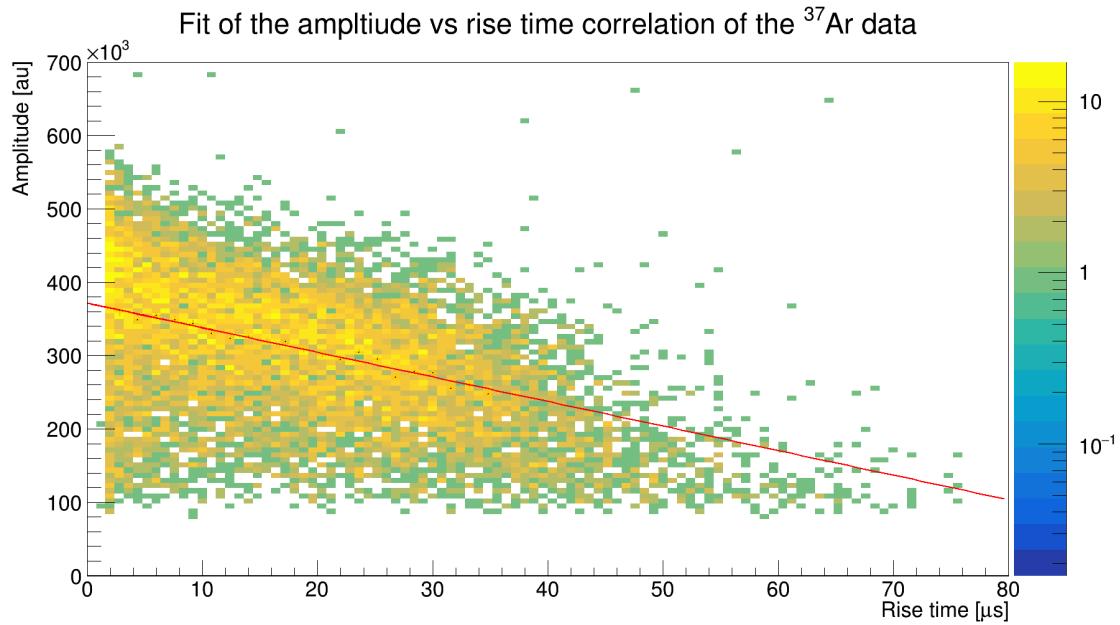


Figure 6.9: Amplitude vs Risetime for the Ar^{37} calibration run. A linear fit is performed on the average amplitude in each slice of risetime. The value of the slope is $3.3 \pm 0.1 \cdot 10^2$.

performed for SEDINE right after the physics run, with results shown on Fig. 6.9. We see that the events at low risetime have amplitudes up to three times as high as events with high risetime. This effect can be reproduced qualitatively in the simulations by introducing some oxygen into the gas mixture, with higher concentrations leading to stronger dependence between amplitudes and risetimes.

A. Brossard ran multiple simulations at different concentrations of oxygen, then fitted the relationship between amplitude and risetime for all, as shown on Fig. 6.10 [35]. The comparison with the result from the calibration run showed the best match at an oxygen contamination of 16 ppb.

It should be stressed that this number is not be taken as a measure of the actual oxygen contamination. Experience with other SPC detectors showed contamination

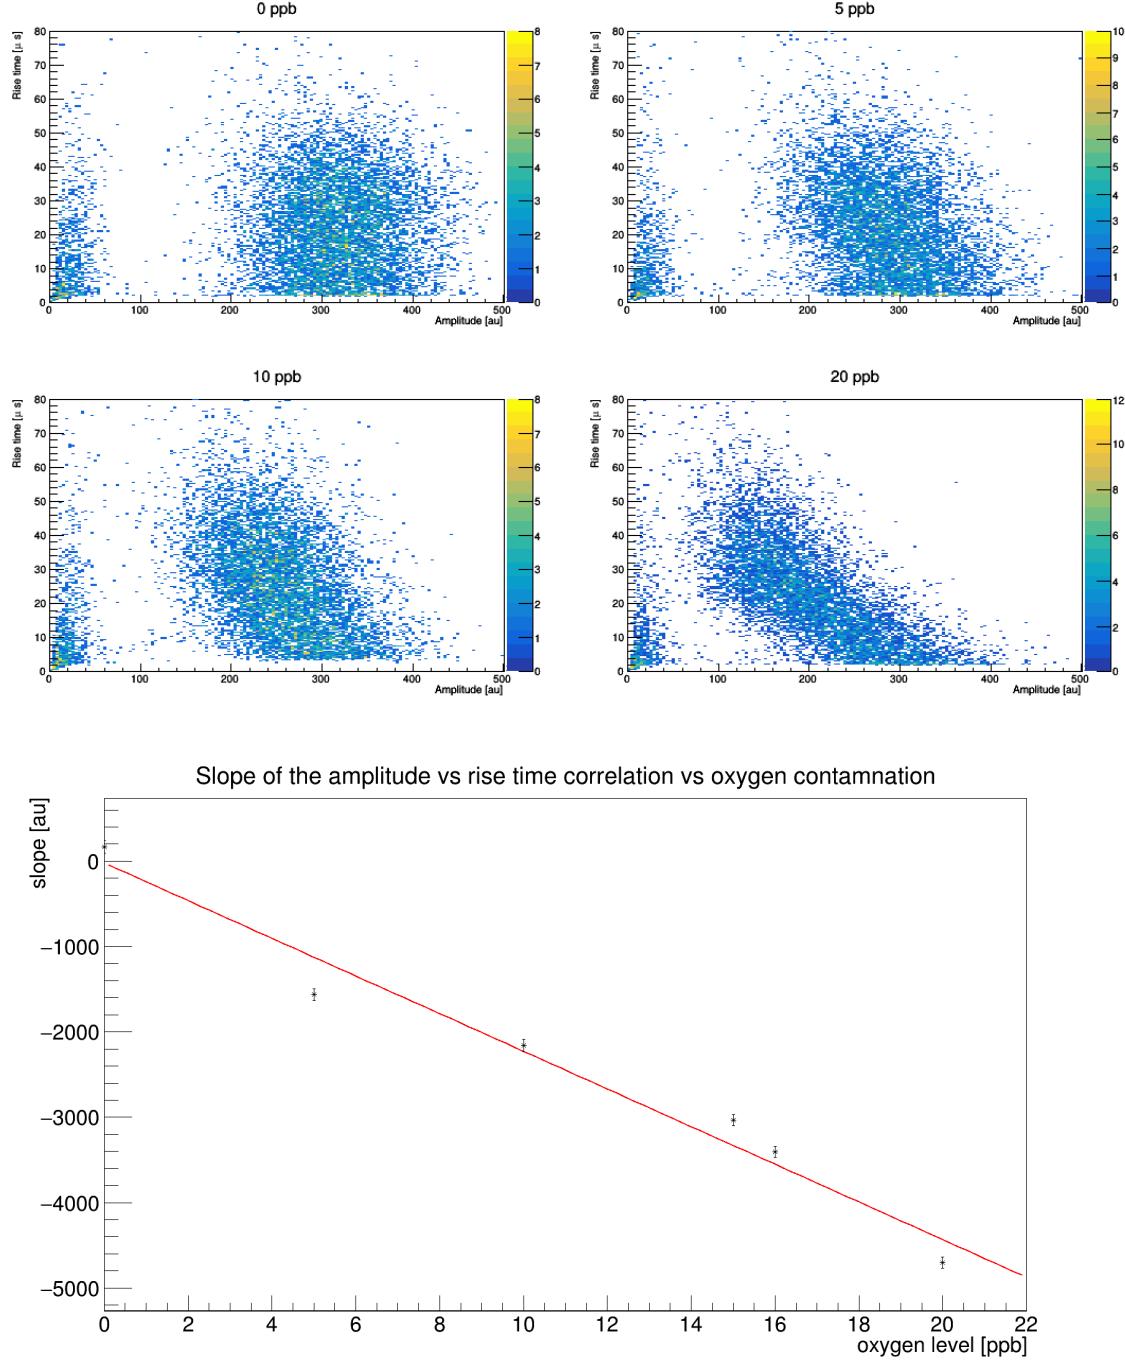


Figure 6.10: Top: Simulated ^{37}Ar runs with different oxygen contaminations. Bottom: Fitted slopes from simulations [35]; the simulations were not calibrated for amplitude, so their units are arbitrary (au). The value of the slope from the calibration run is $3.3 \pm 0.1 \cdot 10^2$, corresponding to an oxygen contamination of 16 ppb.

values around the tens of part per million, primarily from oxygen trapping in tubing and flanges. While it is in principle possible to achieve such a small oxygen contamination level, given the precision of Magboltz at low pressures and fields, it is safer to only take it as a simulation parameter that quantifies attachment.

Note that the value of the oxygen contamination is not used anywhere after generating the Magboltz files. When modifying the strength of the attachment (to study the effect systematic uncertainties, see App. K), instead of generating new Magboltz files, the attachment coefficients are all multiplied by a constant factor.

6.6 Ion mobility

The mobility of ions in the gas affects the response function of the detector. Using the wrong value in our pulse processing may lead to the improper deconvolution from the raw pulses, producing either undershoots or long upwards tails in the final pulse. This leads to an increased uncertainty on our estimation of the amplitude of the pulse, and considerably shorter or longer risetimes. An illustration of this phenomenon is shown in Fig. 6.11.

While this is relatively easy to notice when looking at individual pulses, a good calibration should be able to put some level of confidence on the ion mobility observed. To do so, the task of finding the ion mobility that leads to the best deconvolution has to be automated, then applied to all the events that pass some basic cuts.

6.6.1 Pointlike events

To determine if an event is properly deconvolved, the method I used was to check for tails and undershoots for pointlike events. The best method appeared to be to

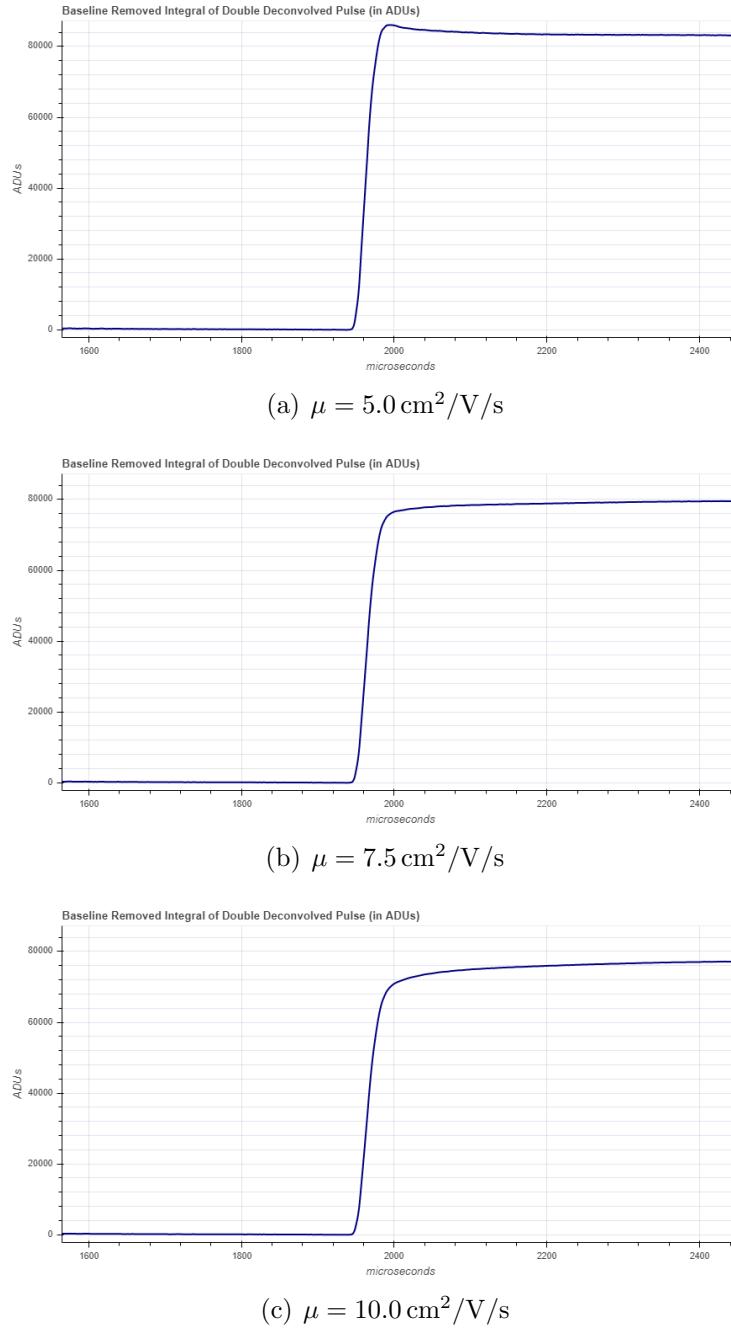


Figure 6.11: Pulse deconvolved by the response function of the detector, then integrated, for different values of the ion mobility in the processing. Low mobilities tend to “overdeconvolve” the pulse, producing a dip after it. Conversely, high mobilities “underdeconvolve” the pulse, leading to an increasing slope right after it.

check for the maximum of the deconvolved integrated pulse, and to compare it to the value of the new baseline after the pulse. Some simulations revealed that for events deconvolved properly, the maximum should be within 3% of the value of the baseline after the pulse, with the reference point being the value of the baseline *before* the pulse. While this does not remove all overdeconvolved events, it did remove the vast majority.

Conversely, to check if the event is underdeconvolved, the heavily smoothed deconvolved (but not integrated) pulse was fitted with a gaussian. If the pulse has a tail, the normalized χ^2 would be increasingly higher with the size of the tail. An attempt was made to see if events with undershoots could be rejected this way too, but the fit turned out to not be sensitive to them, likely due to the heavy smoothing.

The method devised to produce an ion mobility calibration was to select all events that passed some basic pointlike-shape cuts, then deconvolve each of them multiple times with different ion mobilities. For each event, the ion mobility selected was the one producing the lowest χ^2 , as long as the undershoot cut was passed and the χ^2 was under a certain value (of the order of 2 – 3). The resulting mobilities were then put into a histogram for further analysis.

This procedure was carried out for two runs: the ^{37}Ar calibration run, and the background run. Out of 35900 and 26360 physical events, respectively, only 242 and 123 events were deconvolved well enough. The low statistics are likely due to the χ^2 cut being equivalent to a combination of low noise, high amplitude, and small risetime requirement for the pulses. The results are shown on Fig. 6.12.

The first thing to note is that the mobilities obtained for the two runs do not agree with each other: $6.18 \pm 0.99\% \text{ cm}^2/\text{V/s}$ and $7.45 \pm 1.95\% \text{ cm}^2/\text{V/s}$, respectively. The

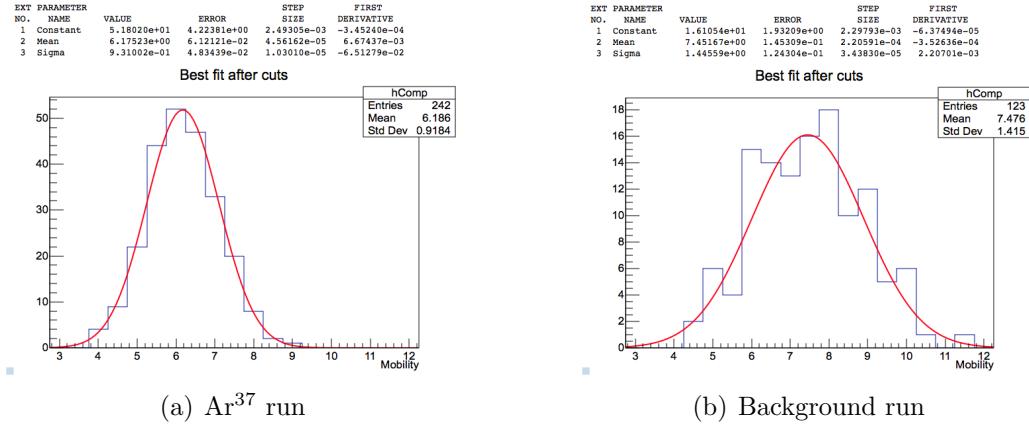


Figure 6.12: Distribution of computed ion mobilities for the two studied runs, with results from the fit. The final ion mobilities computed with this method are $6.18 \text{ cm}^2/\text{V/s} \pm 0.99\%$ and $7.45 \text{ cm}^2/\text{V/s} \pm 1.95\%$, respectively.

difference cannot be explained away with just low statistics, since the fits are actually rather precise. Two possible explanations come to mind for this discrepancy. The first one is the change in gas when introducing ^{37}Ar in the detector between both runs, as mentioned earlier. While ideally the introduced ^{37}Ar was in such low quantity as to not affect the effective gas composition (neon with 0.7% methane) at 3.1 bar, any gas filling procedure may introduce impurities into the gas mix. The effective mobility of ions in a gas is highly dependent on the gas composition, partially due to the change in the drift gas composition, but also due to the change in the ions that end up drifting (e.g., Ar^+ instead of CH_4^+ or Ne^+ ions). At higher pressures, the nature of collisions between ions and surrounding particles also change [64]. All these effects combine into turning even a relatively small modification to the gas composition into an appreciable ion mobility difference.

The second explanation is due to the different rate of events between the calibration run and the background run. An increase in rate leads to an increase in

the number of secondary ions drifting through the detector volume, in turn changing the internal electric field. The calibration run may have had a large enough space charge due to ions to decrease the field close to the sensor compared to that of the background run. Since the response function used in the deconvolution assumes an ideal electric field, a decrease in the real field would translate into a decrease in the ion mobility with this method.

At any rate, these unknowns mean that this measure of the ion mobility should be taken as an *ad hoc* value to match the response function of our detector, rather than a true measure of the ion mobility. Since that is what we use it for, the limited physical interest of the values found is not a detriment. However, it is still encouraging that the values are relatively close to that found on the literature for Ne^+ ions in pure Neon, $\sim 4 \text{ cm}^2/\text{V/s}$ [50, 65]. The difference between the two is most likely explained by a larger-than-expected field close to the electrode being “absorbed” into the computed ion mobility². If we were interested in a real measurement of the ion mobility, a setup with parallel plates (and uniform electric field) would be much preferable.

Summary

We have covered all our calibrations that were available for SEDINE. Aside the electron drift and diffusion time in the bulk of the detector (for which no calibration method was available), the simulation parameters obtained were (in parenthesis, the calibration method used):

- surface electron drift time: $422.4 \mu\text{s} \pm 5\%$ (by studying maximum track length);

²A difference larger than an order of magnitude between the computed ion mobility and the one found in the literature would be harder to justify that way, and would suggest a lack of understanding of our detector.

- surface electron diffusion time: risetime of $50.0 \mu\text{s} \pm 6\%$ (by studying events from surface contamination);
- energy (or avalanche gain): $2.077 \text{ ADU/eV} \pm 2.9\%$ (from copper fluorescence);
- ionization energy: $27.6 - 36 \text{ eV}$, (with a pulsed laser and ^{37}Ar in a test SPC), although this uncertainty is effectively absorbed by the energy calibration;
- electron attachment: “effective” O_2 Magboltz contamination of 16 ppb, with an electron attachment uncertainty of $\sim 100\%$ (by introducing ^{37}Ar in SEDINE);
- ion mobility: $7.45 \text{ cm}^2/\text{V/s} \pm 1.95\%$ (by studying SEDINE’s response function).

A different sort of calibration, which reproduces axion-like events, is studied in the next chapter.

Chapter 7

MPA applied to ^{55}Fe -induced Argon fluorescence

To test the multi-pulse analysis methods in preparation for the search of KK axions, a calibration was performed with ^{55}Fe -induced Argon fluorescence. ^{55}Fe decays into ^{55}Mn by electron capture, leaving an electron vacancy in the K-shell, which is then filled by an electron from a higher shell. The difference in energy is then released by either an Auger electron of 5.2 keV, a K- α X-ray of 5.9 keV, or a K- β X-ray of 6.5 keV, with α 's being 8 times as likely as β 's [66, 67]. In turn, an Argon atom that absorbs a photon will fluoresce, with a probability of 12 % (known as the fluorescence yield¹), emitting a photon of 2.9 keV, leaving behind a bit under 3.0 keV of energy, to be dissipated through ionization.

At the right pressures, this will lead to two simultaneous energy depositions at different positions of the detector, both of the same energy: the leftover ionization energy from the original ^{55}Fe photon, and the energy from the absorbed fluorescence photon. With some divergence in the volume distribution of these events, this is

¹Early measurements of the fluorescence yield of Argon pointed towards 8 – 9 % [68, 69], but measurements of 12 – 14 % started cropping up in the early sixties [70–72]. Later reviews converged on a value of 12.0 %, combining experimental data of Argon and fits of the yield over different elements [73–75].

the same signal we would expect from 6 keV axions decaying within the detector, providing an excellent calibration for multiple-pulse analysis of our signal of interest.

7.1 Setup

To optimize our double-pulse event frequency, a low gas pressure is preferred to increase the absorption length of 3 keV photons, leading to better separation between both energy depositions. This in turn means we require a larger detector, to limit the number of photons escaping. The final setup used the 130 cm wide SPC at Queen's (S130 detector), with a 2 mm diameter electrode, with either 200 mbar or 110 mbar, and a high voltage of 1150 V or 900 V, respectively. The pressures were selected based on Monte Carlo simulations of the detector to maximize the rate of recognizable fluorescence events, and the voltages were tuned to obtain similar gains in both setups.

The 37 MBq ^{55}Fe source used was collimated with an aperture of 1 mm, approximately 5 mm away from the source. The aperture was covered by two sheets of aluminium foil to block β radiation, roughly 20 μm thick each, and placed at the end of a 4 cm-long window into the detector. Pictures can be seen on Fig. 7.1. A pulsed laser with a frequency of 10 Hz was used concurrently, to calibrate for the drift and diffusion time of primary electrons coming from the surface of the detector.

Unfortunately, using a large detector on the surface comes with an added difficulty, in the form of cosmic radiation. We expect a muon flux at sea level of $\Phi(\theta) \simeq 70 \text{ m}^{-2}\text{s}^{-1}\text{sr}^{-1}\cos^2(\theta)$ (depending on atmospheric conditions[76]), which corresponds to an event rate of $\simeq 390 \text{ Hz}$ for a spherical detector of radius 60 cm, from muons alone. This is a problem, since we expect an event rate from the ^{55}Fe source in the

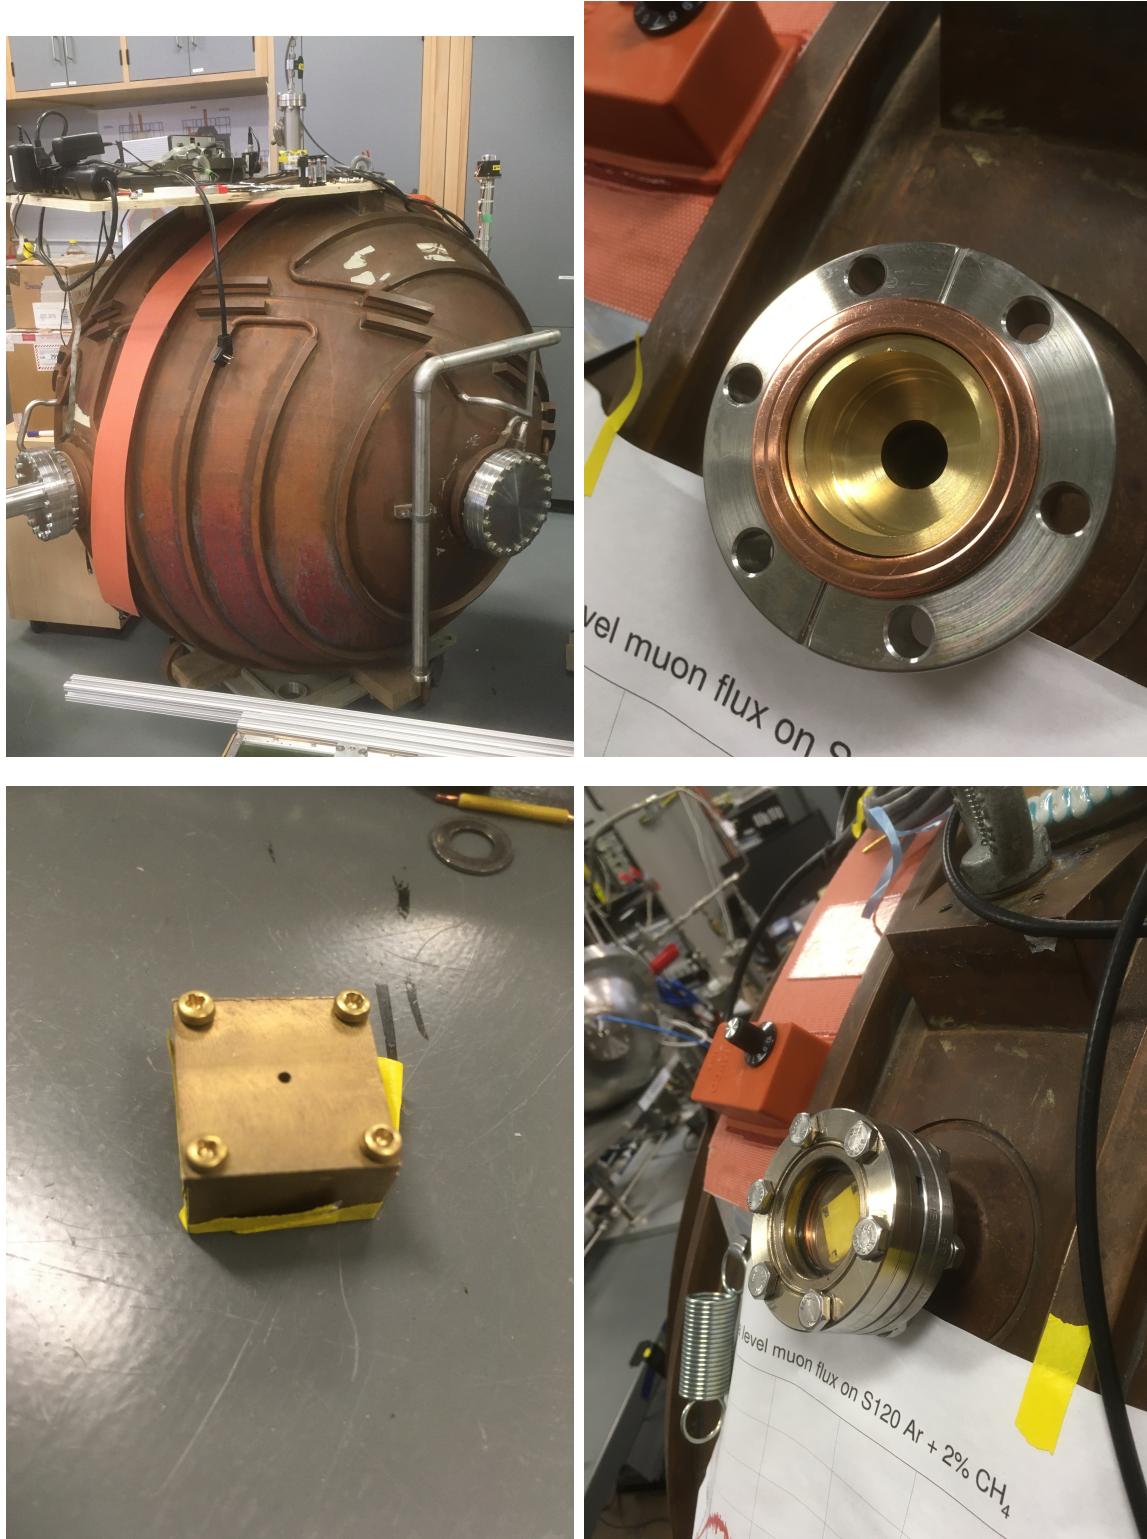


Figure 7.1: Top left: S130 detector used for the calibration. Bottom left: ^{55}Fe source (without aluminium foil). Right: Close-up on the calibration source window, with and without the ^{55}Fe source installed in it.

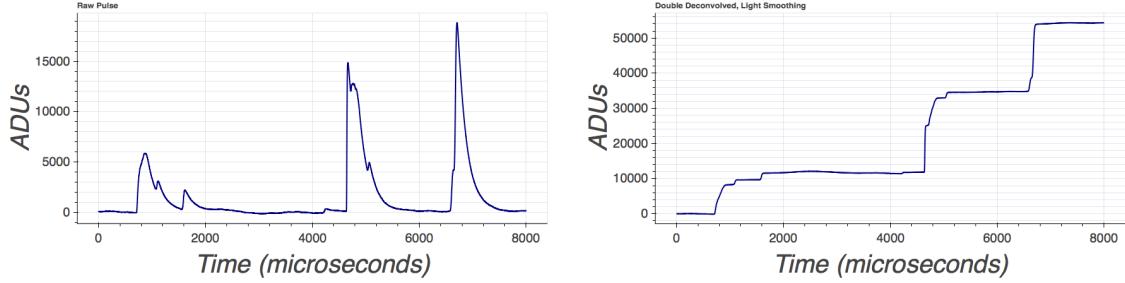


Figure 7.2: Event example from S130 with 200 mbar. Left: Raw pulse. Right: Processed. The laser-induced event is the smaller one at $\sim 4300 \mu\text{s}$. The pulse at $\sim 1800 \mu\text{s}$ is likely an ^{55}Fe event. The other three are likely muons.

order of tens of Hertz. The first way in which it is problematic is the large amount of pileup in the data; for recorded event windows of 8 ms, we expect 3.2 events to fall within *in addition to the one that triggered*, from muons alone. A representative example is shown in Fig. 7.2. Care has to be taken to separate the true simultaneous events we are interested in from random coincidences. The second way in which this is a problem is that for such a large detector, with such an important rate of high-amplitude events, the electric field far from the sensor is dominated by that from secondary ions drifting away from the avalanche region. As such, we cannot use COMSOL files to get the field strength everywhere in the detector, but have to resort to the analytical approximation described in 5.4.3. Fortunately, the high rate means that the continuous space current assumption should hold and give appropriate results.

7.2 Accounting for the high rate

To simplify the corrections that had to be done, the runs studied used a laser-dependent trigger, running at 10 Hz. Due to the relatively wide event window recorded,

at 8 ms per trigger, and the very high event rate, enough ^{55}Fe statistics could be obtained from this approach, despite the effectively random nature of the trigger. The advantage was that instead of having to compute how much dead time there was in between triggers, *all* such time was effectively dead. The basic run length was thus $N_{triggers} \cdot 8\text{ ms}$. For the run at 200 mbar, this meant a total of 650 s; $\sim 2810\text{ s}$ for the considerably longer run at 110 mbar, after removing unstable stretches of time. Aside from the duration, both runs provided similar results, so the following will only be referring to the 200 mbar run, unless otherwise specified.

The Multi-Pulse Analysis was used on these recorded windows, splitting each one into one or more separate pulses. Since we are looking for double-pulse events, multiple pulses close to each other were joined into a single event. The laser events provided a calibration for the maximum drift time of electrons in the detector, as shown in Fig. 7.3. This maximum drift time is necessarily the maximum time separation between two pulses generated from simultaneous energy depositions in the detector. Adding a safety margin based on the maximum diffusion time from laser events, pulses are iteratively joined into a single event as long as they are within $260\text{ }\mu\text{s}$ ($260\text{ }\mu\text{s}$ at 110 mbar) of each other.

This leaves another source of dead time, in the form of the width of the joined pulses themselves, plus an extra “tail” of $260\text{ }\mu\text{s}$ for each of them. While this would normally be ignorable in setups with low event frequency, it was a sizeable contribution in this scenario. It reduced the effective run time from $\sim 650\text{ s}$ to $\sim 505\text{ s}$ (from $\sim 2810\text{ s}$ to $\sim 2330\text{ s}$ at 110 mbar). After dead time corrections, we get a total event rate of 372 Hz, discounting the contribution from laser events.

Once this has been accounted for, cuts are selected to remove unwanted events.

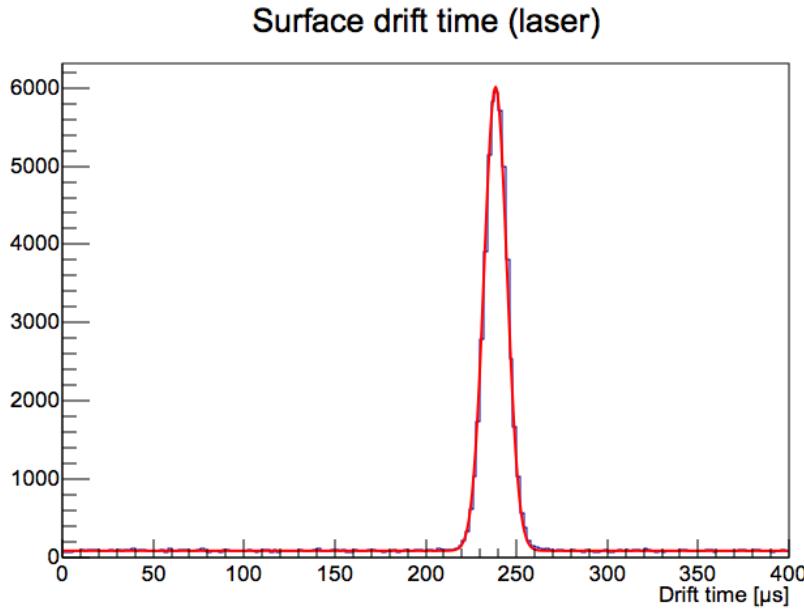


Figure 7.3: Maximum drift time calibration with laser events, with gaussian fit. The time of ionization is obtained with a fiber splitter that sends parts of the laser signal to a photodetector. The difference between that time and the average arrival time of the primary electrons gives the (maximum) drift time, around $240\ \mu\text{s}$ at 200 mbar, 900 V for this detector.

The first round of basic cuts removes laser events, saturated events, events at the edges of the window, and events that are over-deconvolved (usually a sign that they are not signals induced by ions drifting). The resulting distribution is shown in Fig. 7.4, and corresponds to a total rate of 325 Hz. The ^{55}Fe is very well defined, allowing for an energy calibration at 5.9 keV. To be able to get estimations of the proportion of Argon fluorescence events, the contribution of ^{55}Fe events has to be separated from all other events, which we will assume to be mostly Muon events. To do this, the distribution under 20 keV is fit with a linear function plus a Gaussian. The ratio between the Gaussian and the linear part of the fit is used to split the total distribution in two, into a ^{55}Fe contribution and a Muon distribution, again shown in

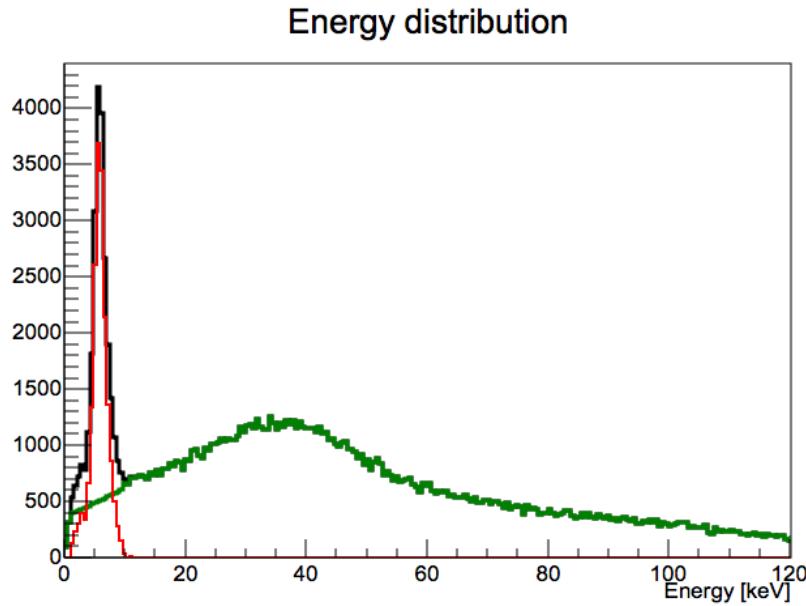


Figure 7.4: Black: Total energy distribution of all events after basic cuts. Red: Contribution from ^{55}Fe events. Green: Contribution from Muon events; they drop down to 0 around 180 keV.

Fig. 7.4. The rate of ^{55}Fe is found to be 29 Hz, leaving around 300 Hz of Muon events that pass the basic cuts.

7.3 Space charge effect

Given the high rate of events, the space charge (or ion) current for a given run can be taken as the total integrated charge recorded, divided by the effective length of the run. The integrated charge from an event is obtained by dividing the amplitude of the processed pulse by the amplification factor of the digitizer (in terms of ADC units per Volt) and the gain of the preamplifier (in terms of Volt per integrated charge). We obtain 85 pA. This value can be used in simulations to try to reproduce the drift time. These show that, in the absence of space charge, the drift time of electrons

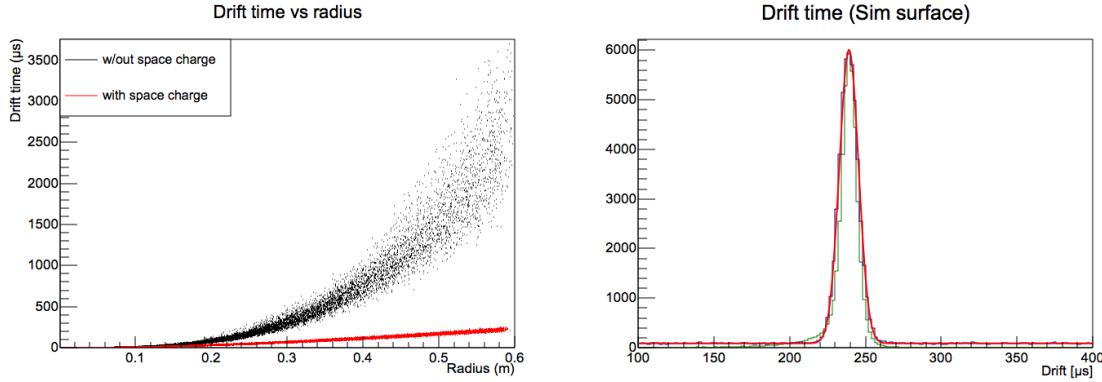


Figure 7.5: Left: Electron drift simulation, total drift time depending on initial radial position. Simulations are run both with no space charge, and with the space charge induced by an ion current of 19.7 pA. Right: Fitted distribution of drift times in the data (blue/red, as seen in Fig. 7.3) and normalized distribution of drift times in simulations with adjusted ion current

generated at the edge of the detector would take up to 2.5 ms to reach the central electrode. With the measured ion current, the drift time drops all the way down to 360 μ s, much closer to the real one, yet still not quite the same. Some trial and error finds that an ion current of 19.7 pA, slightly over twice the measured one, reproduces a surface drift time of 240 μ s. The two extremes are represented on Fig. 7.5.

Multiple factors could explain this discrepancy:

- Pile up effects or erroneous calculations of the dead time could bias this relatively simple measurement of the ion current.
- Total charge from high energy events is not properly reconstructed, because they saturate. This likely leads to an underestimation of their contribution to total space charge. Despite this, they still account for $\sim 10\%$ of the total.
- We are folding the uncertainty over the ion mobility (of roughly $\sim 10\%$) into

the ion current. For a given number of ions created, a smaller ion mobility will lead to a proportionally larger amount of ions present in the volume of the detector.

- Differences between the ideal field configuration assumed in the space charge field formula, and the real detector configuration. These could come in many forms, mainly affecting the path ions take, and so the time they spend inside the detector.

All in all, a difference of a factor 2 between the measured ion current and the one required to match drift times is in good qualitative agreement.

7.4 Expected results using a toy model

A simple simulation of ^{55}Fe -induced Argon fluorescence events can be performed, considering only the geometry of the detector, and the attenuation length of photons of different energies (see Fig. 7.6). An Argon atom that fluoresces after absorbing a 5.9 keV photon from the ^{55}Fe source can produce three different kinds of events:

- Escape peak event: the 2.93 keV fluorescence photon leaves the detector without interacting. This leaves 2.9 – 3.0 keV energy in the detector. The longer the attenuation length of 2.93 keV photons, the more frequent this is.
- Double pulse event: the 2.93 keV fluorescence photon is absorbed in the detector after travelling some non-negligible distance. This will appear as two pulses of ~ 2.9 keV amplitude arriving shortly after each other. The maximum time separation between both is the drift time of electrons coming from the surface:

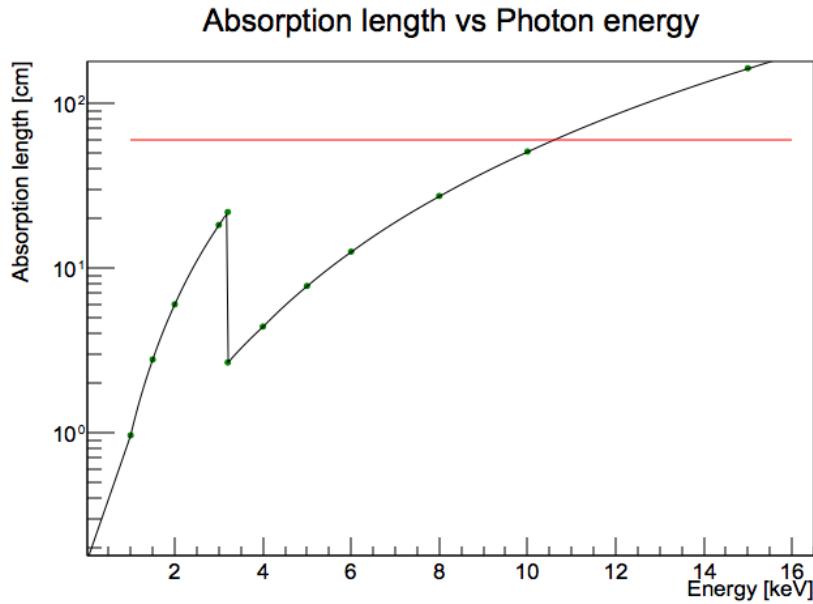


Figure 7.6: Attenuation length of photons in 200 mbar of Argon (NIST datapoints in green, log-log linear interpolations in black). In red, the radius of the detector. The drop at 3.2 keV corresponds to the K-shell of Argon. One of its consequences is that a 2.93 keV and a 5.9 keV photon have comparable attenuation lengths, despite the factor 2 difference in energy.

this happens if the ^{55}Fe photon was absorbed close to the surface, while the fluorescence photon was absorbed close to the central electrode, or viceversa.

- Non-separable event: the 2.93 keV fluorescence photon is absorbed very close to the point where it was generated. The pulses will then arrive to close to each other to be separated, and will just look like a single pulse. The shorter the attenuation length of 2.93 keV photons, the more frequent this is.

Notably, the rate of escape peak events depends only on the size of the detector and the attenuation length of photons. For a given gas mixture and photon energy, the attenuation length is given solely by the pressure of the gas. In turn, for double pulse events, we also need to introduce a “separability distance”, i.e., the minimum

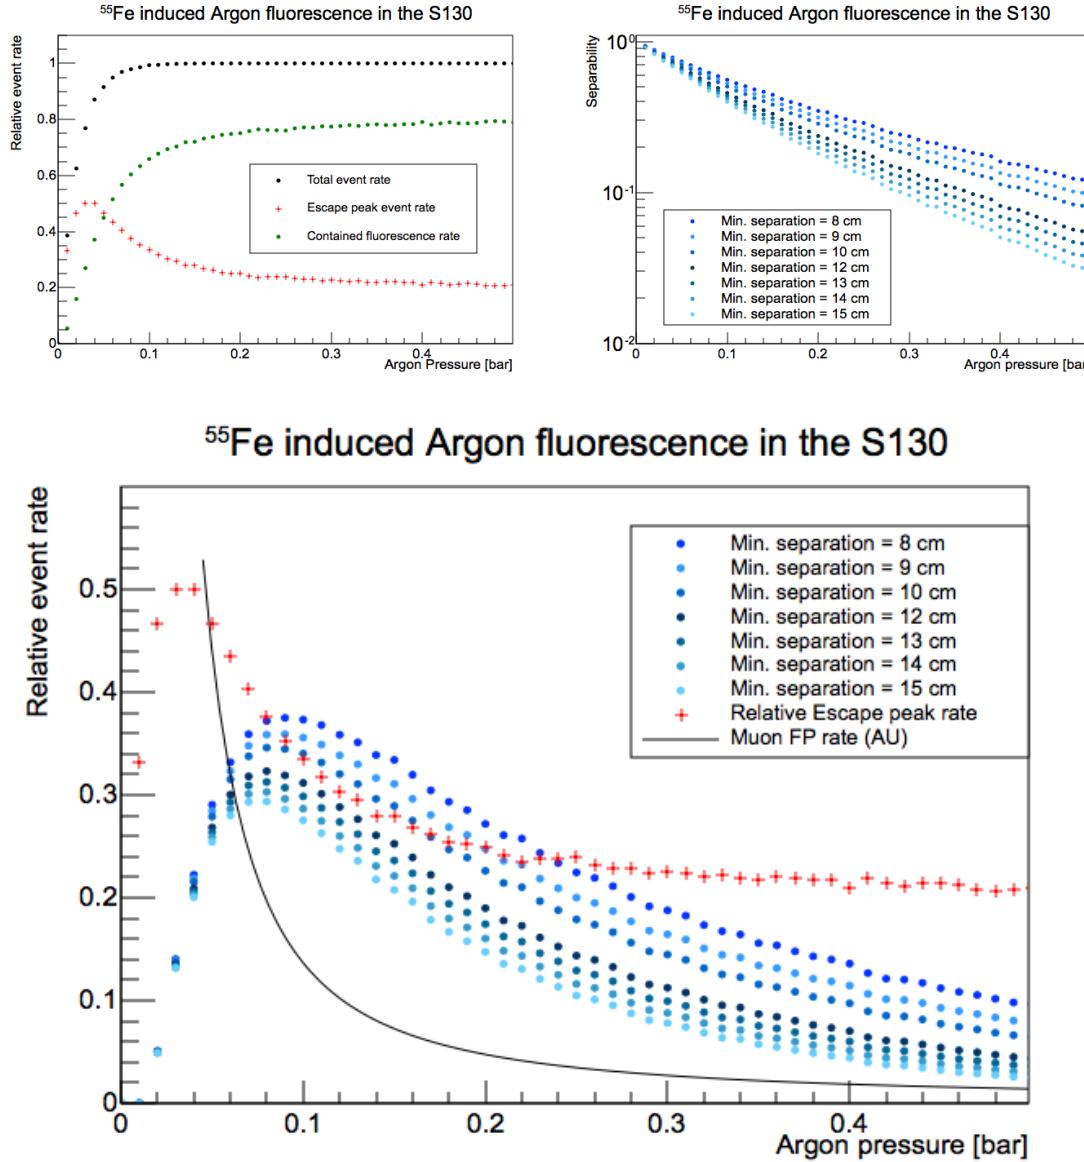


Figure 7.7: Top left: Escape peak event probability depending on gas pressure. Top right: double pulse event probability depending on gas pressure, for different minimum separability distances. Bottom: relative rate of escape peak events and double pulse events; the rate (in arbitrary units) of muons that fall in the energy range of ^{55}Fe events, based on the recorded data, is also added.

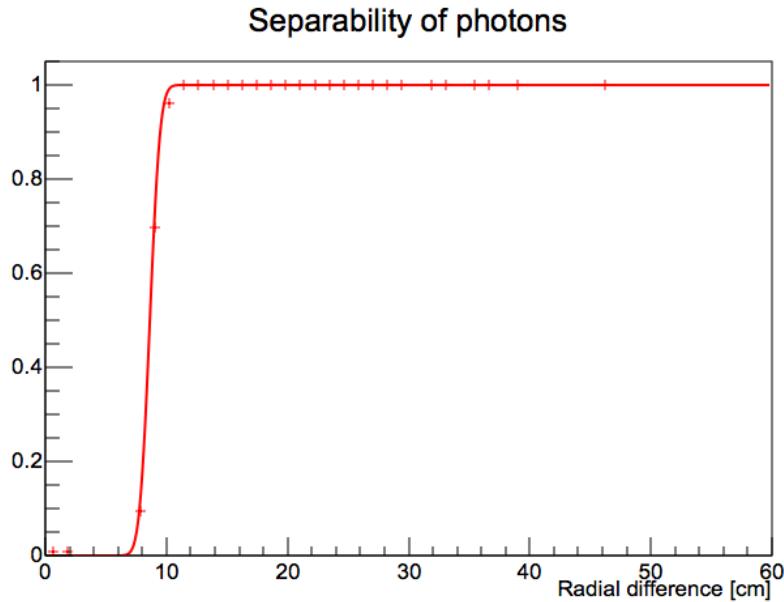


Figure 7.8: Ratio of contained Argon fluorescence events that can be separated with the processing, at 200 mbar with 1150 V. The separability distance is taken to be ~ 8.5 cm.

distance between both interactions before they can be told apart. Taking all this into account allows us to make a simple toy model to find the relative rate of all these events. The results are shown on Fig. 7.7. We see that the rate of escape peak events increases as pressure decreases, until we reach 40 mbar; at that point, the original ^{55}Fe photons stop being absorbed by the detector, and the total event rate plummets.

The separability of both photon interactions for contained fluorescence events also tends to improve with decreasing pressure, for a given minimum separability distance. However, the separability distance itself will depend on the gas pressure and voltage, so the relationship between pressure and separability is not as straightforward. To find the separability distance for both run conditions, full simulations of double events where performed and processed, taking into account the space charge effect as described in the previous subsection. By plotting the ratio of events properly separated

depending on the radial distance between them, we get the separability distance for each run. This is shown in Fig. 7.8, giving a separability distance of ~ 8.5 cm (~ 14.0 cm at 110 mbar). This is used together with the plot in Fig. 7.7 to find the rate of double-pulse events we expect from ^{55}Fe -induced Argon fluorescence.

Note that the previous results are given as a ratio of all Argon fluorescence events. In reality, we do not have access to that number, but to the observed number of ^{55}Fe events instead (potentially missing escape-peak or double-pulse events). Including the effect of the fluorescence yield, and taking the error as dominated by an uncertainty in the pressure measurement of ± 5 mbar (for the escape peak events) and in the separability distance of ± 0.5 cm (for double-pulse events), the expected ratios from this toy model are:

- 3.11 – 3.25% proportion of escape peak events, 2.80 – 3.22% proportion of double-pulse events, for 200 mbar and 1150 V;
- 3.97 – 4.38% proportion of escape peak events, 3.18 – 3.47% proportion of double-pulse events, for 110 mbar and 900 V.

7.5 Escape peak

7.5.1 200 mbar data

Before looking at double events, we look at Argon fluorescence escape peaks, i.e., Argon fluorescence events where the fluorescence photon escapes the detector without interacting. This leads to a 2.9 keV event in the detector. To boost the ratio of ^{55}Fe events over Muon background, we keep only events with a single pulse, within a narrow range of risetimes between 27 μs and 40 μs .

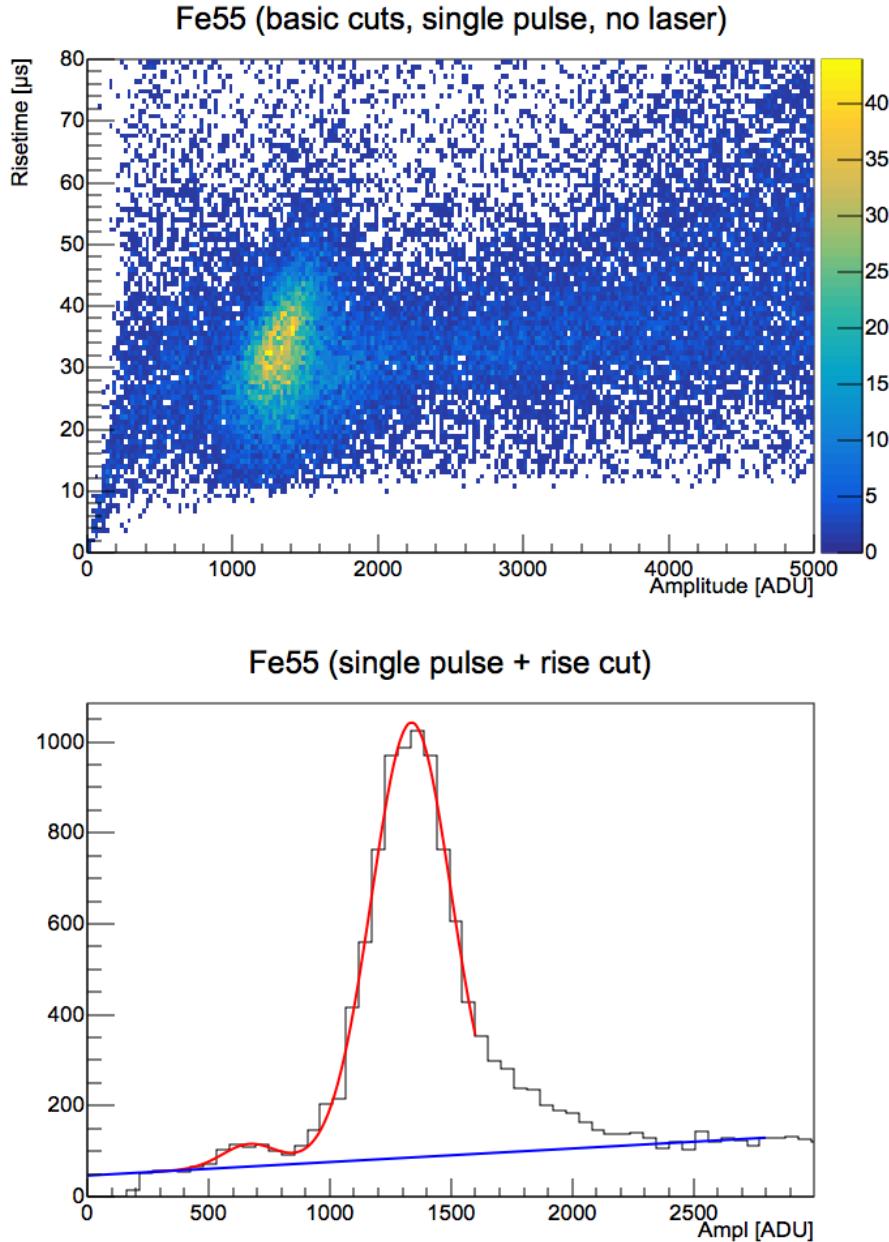


Figure 7.9: Top: Risetime vs Amplitude after basic cuts, keeping only single-pulse events and removing laser events; ^{55}Fe are at 1300 ADU, with escape peaks around 600 ADU; the horizontal accumulation around $35\ \mu\text{s}$ across amplitudes comes from Muon events. Bottom: Amplitude distribution for events with risetime between $27\ \mu\text{s}$ and $40\ \mu\text{s}$; the sum of a linear function (shown separately in blue) and two gaussians is fitted to the distribution.

The resulting distribution is fitted with the sum of a linear function and two gaussians, as shown on Fig. 7.9. The ratio between the mean of both gaussians is found to be 0.500 ± 0.015 , in great agreement with expectations. The integral of the gaussian corresponding to the escape peak is of 258 counts, with an error derived from the fit results of ± 77 . The large error is to be expected, due to the low statistics of the escape peak, and it being on top of a sizeable background; it comes primarily from the uncertainty on the width of the peak. The integral of the gaussian corresponding to the main ^{55}Fe has much better statistics, with an integral of 7349 counts and an error of 130. This leads to a measured proportion of escape peaks to other single-pulse ^{55}Fe events of $3.5 \pm 1.0\%$.

Another issue is the tail above 1600 ADU, that is not consistent with a Gaussian. It is unclear whether that tail comes from actual ^{55}Fe events, or some other form of background. This introduces a significant source of uncertainty in the measured ratio: assuming all events in the tail come from ^{55}Fe , and that the statistical uncertainty is dominated by the number of escape peak events, the final ratio becomes $3.0 \pm 0.9\%$. Ultimately, both sources of uncertainty (statistical and “tail”) originate from the difficulty to extract the distribution of ^{55}Fe events from the large Muon background. In either scenario, they are consistent with the 3.2 found with the toy model.

7.5.2 110 mbar data

The integral of the gaussian corresponding to the escape peak is of 2330 counts, with an error derived from the fit results of ± 417 . Once again, the large error is to be expected, this time primarily due to resting on top of a proportionally larger background, despite the statistics being larger than before; it still comes primarily from

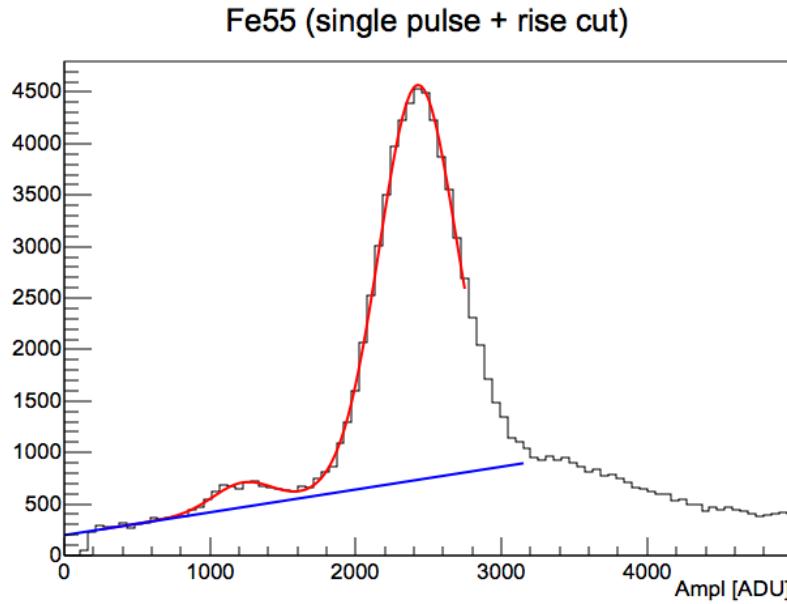


Figure 7.10: Amplitude distribution for events with risetime between $21\ \mu\text{s}$ and $35\ \mu\text{s}$; the sum of a linear function (shown separately in blue) and two gaussians is fitted to the distribution.

the uncertainty on the width of the peak. The integral of the gaussian corresponding to the main ^{55}Fe has much better statistics, with an integral of 47163 counts and an error of 1244. This leads to a measured proportion of escape peaks to other single-pulse ^{55}Fe events of $4.9 \pm 0.9\%$. The tail effects mentioned for 200 mbar appear less important, only shifting this result to $4.7 \pm 0.9\%$. Both are relatively high, but still consistent with the expected results of $3.97 - 4.38\%$ obtained with the toy model.

7.6 Double-pulse events

7.6.1 200 mbar data

Unlike for escape peak events, we now want to look at events where two pulses happen closer to each other than the maximum drift time, and we also cut events where there

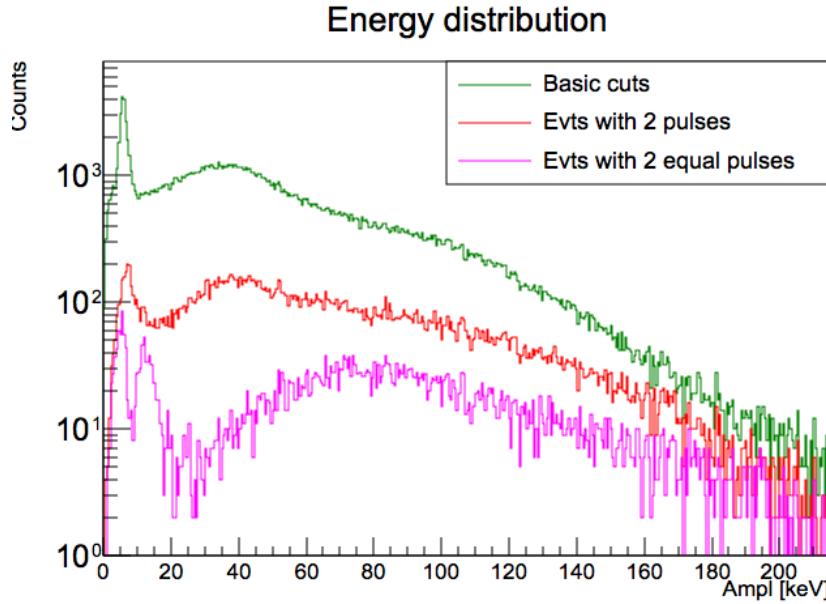


Figure 7.11: Effect of number of pulses in the event distribution (200 mbar data).

are more than just two pulses within that same window. This turns out to still leave a large number of events, so an extra condition is added that both pulses are “equal”, defined as having their amplitudes within 30% of their average; the choice of this fairly wide tolerance in amplitude difference will be justified in a moment. The results are shown on Fig. 7.11.

We see a very clear population of events with two equal amplitude around 6 keV, corresponding to the Argon fluorescence events that the processing can separate. Perhaps surprisingly, we see a second population suddenly appear around 12 keV, and a large diffuse population centred around 80 keV. These are random coincidences of uncorrelated events. That is, two events of same energy happening shortly after each other by pure chance, due to the high event rate in this run.

To prove this, the final distribution of events with two equal pulses is fitted with a normalized sum of three distributions:

- the total distribution of ^{55}Fe events. These represent separable Argon fluorescence events (after normalization);
- the total distribution of Muon events. These represent False Positives (after normalization), i.e., single-pulse events improperly reconstructed as being multiple;
- the distribution of random coincidences with equal energy. This is built by taking two random events from the distribution of events with a single pulse, and keeping the sum of their amplitudes as an entry if they are within 30% of their average. This process is repeated $N^2 \frac{\Delta t}{T}$ times, where N is the total number of single-pulse events, $\Delta t = 260 \mu\text{s}$ is the maximum time separation we keep, and $T = 505 \text{s}$ is the effective length of the run.

The results of the fit are shown on Fig. 7.12.

The most remarkable part of the fit is how closely the distribution of random coincident events matches both the peak at 12 keV, and the wide population of events 80 keV. This is despite having no parameter driving its shape, with the distribution being built straight from the data, and adding just a normalization factor afterwards. However, the normalization factor given by the fit is 1.70. That is, we find 70% *more* events than what the $N^2 \frac{\Delta t}{T}$ formula predicts. This could be for three reasons. The first is an underestimation of N : if events other than single-pulse ones that pass the basic cuts could somehow combine into events that pass the final level of cuts (e.g., one large pulse followed by two smaller ones too close to each other to be separated); this seems unlikely, given that it would probably affect the final shape of the distribution. The second is an overestimation of T , most likely due to missing sources of dead time;

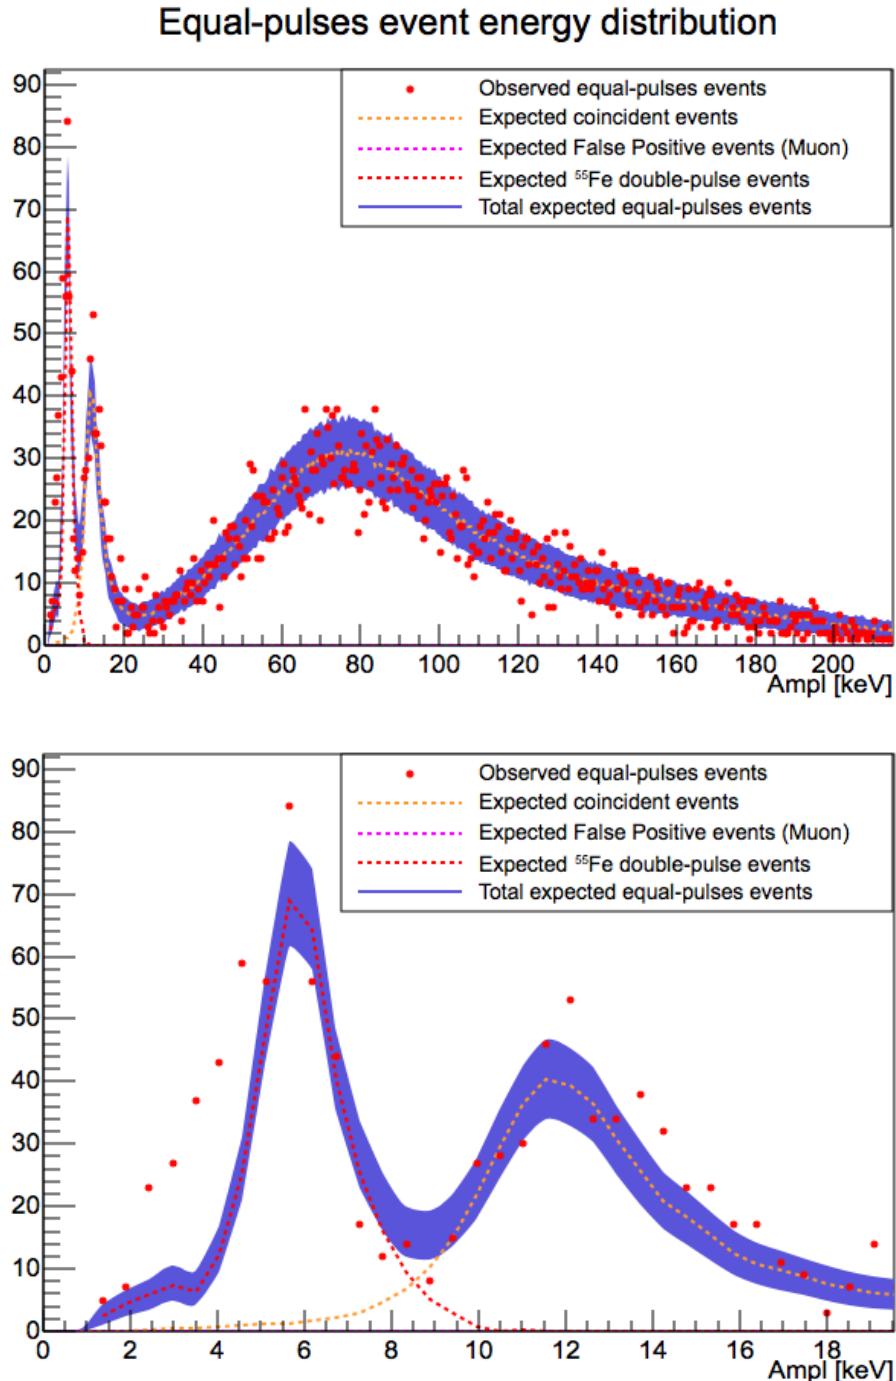


Figure 7.12: 3-parameter fit of equal-pulses event distribution. Top: whole range of the fit. Bottom: Zoom in the region under 20 keV. Red points: Data (200 mbar). Blue curve: Fit result with statistical uncertainty. Dotted lines: Contribution from coincident events (orange), Argon fluorescence (red), and False Positives (pink; here set to zero by fit). The overpopulation in the 2 – 4.5 keV range are improperly reconstructed events.

this one is relatively likely, although the extra dead time would have to account for $\sim 40\%$ of the effective run length. The last one is a mistake hidden somewhere in the calculations, given the delicate nature of computing the rate in these conditions.

The False Positive (FP) rate for Muons given by the fit is $\sim 10^{-9}$. Given that there are only ~ 150000 Muon events in total, this just means the fit is setting the FP rate to zero. However, given the statistical uncertainty, believing this number would be overly optimistic. For a more reasonable upper limit on the FP rate, we can consider the behaviour of both distributions around 20 keV, where the fit reaches a local minimum. On the corresponding bin, the Muon distribution has 850 events, while the final fit has a statistical uncertainty of 2.7 events. Since FP events have no discernible effect at that point, that sets a very conservative upper limit for the rate of FP at 0.3%. The real rate is likely much lower, but it cannot be constrained further under such high rate conditions.

Finally, the last term in the fit is the ratio of equal-pulse events in the main ^{55}Fe peak. Before any corrections, the fitted proportion is $2.38 \pm 0.15\%$, if we take only the fitted gaussian to determine the total number of events in the peak, or $1.87 \pm 0.12\%$, if we take everything above the “fitted” Muon distribution as being ^{55}Fe events, with the truth being likely somewhere in between. For the sake of reporting a single number, we get $2.13 \pm 0.20\%$, by taking any value in between the two assumptions on the ^{55}Fe peak to be equally likely, and adding the fit errors in quadrature ².

The observed overpopulation of events in the 2 – 4.5 keV range are not real equal-pulse events, but events with a low maximum after deconvolution, leading to pulses

²If we multiply the results from a gaussian of mean m and standard deviation σ with that of a uniform distribution of bounds A and B , we can approximate the result to a distribution with a mean of $m\frac{A+B}{2}$ and standard deviation $m\frac{A+B}{2}\sqrt{(\frac{\sigma}{m})^2 + \frac{1}{3}(\frac{B-A}{B+A})^2}$.

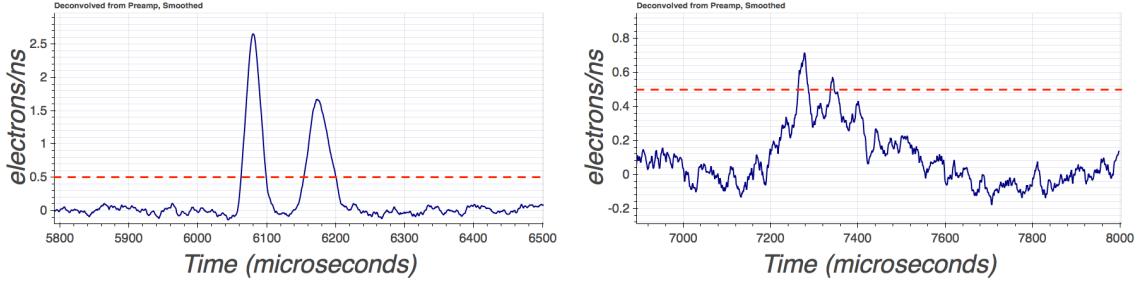


Figure 7.13: Deconvolved events from the 200 mbar run. Left: Argon fluorescence event, with pulse threshold in red. Right: 3 keV event improperly reconstructed as being an equal-pulses event, due to it “jumping” above the threshold multiple times. At energies under 5 keV, the ratio of the second kind of event increases considerably.

oscillating around the pulse threshold. The difference between the two kind of events is demonstrated in Fig. 7.13, with the latter kind comprising the majority of events in the mentioned energy range. This problem would normally be solved by increasing the strength of the smoothing and reducing the value of the threshold. However, in these conditions, the value of the baseline shifts too much due to the high rate of high energy events, so this would not improve the results. Instead, we make the reasonable assumption that the distribution between 5 and 8 keV is driven by Argon fluorescence events, which means that the results from the fit are not biased by this population of “fake” equal-pulses events. This appears justified by inspection of the events in that range.

An extra correction needs to be accounted for, since the cuts to select for equal-pulses events also included a cut to have *exactly* two pulses. This means that events in which a third pulse happened concurrently with the main 2 of the Argon fluorescence event were rejected. Given a measured average width of equal-pulse events of $684 \mu\text{s}$, and a total event rate of 498 Hz, this gives the odds of some other pulse coinciding

with any given fluorescence event of 28.9%. Correcting for this, we get a proportion of equal-pulse events of $2.99 \pm 0.28\%$.

This is to be contrasted with the expected results from the toy model of 2.80 – 3.22%. This number, however, does not take into account the effect of only keeping events when the amplitudes of both pulses are close enough to each other. To estimate the contribution of this cut, we use the results from the fit of the escape peak: for each event in our toy model, we draw two pulses from a gaussian distribution with the mean and width of that fit, and only keep the event if they are within the predefined amplitude tolerance. Given the large uncertainty of the fit on the width of that peak, an excessively restrictive tolerance would in turn lead to a large uncertainty on the effect of this cut. For example, a tolerance of 15% leads to an efficiency of the cut in the range of 67.8 – 91.9%, for values of the peak width within one standard deviation of the mean. Taking 30% instead, the efficiency range is 94.4 – 99.9%, for a much narrower spread. To avoid compounding uncertainties, the larger tolerance was chosen, even though it leads to less rejection of background. We finally get an expected ratio of equal-pulse events in the main ^{55}Fe peak of 2.64 – 3.22%. This result is also consistent with the data, although the ranges of values remain relatively wide.

7.6.2 110 mbar data

The final fit results for the 110 mbar data are shown on Fig. 7.15, with comparison over the complete distribution on Fig. 7.14. There are three main qualitative differences observed with the 200 mbar data. The first is the much higher statistics, with around ~ 10 times more events. The second is that the distribution of Muon events is shifted down in energy by roughly a factor 2; this is expected, since the energy deposited

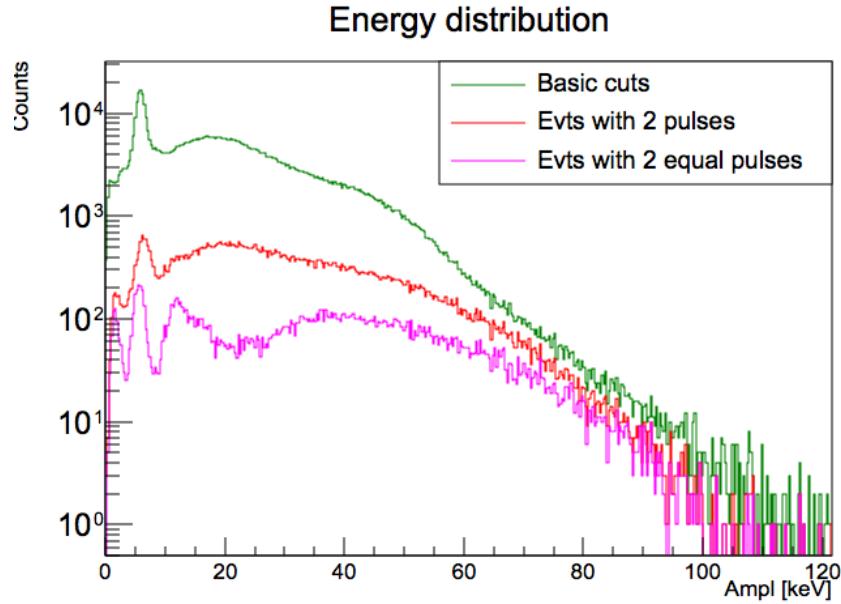


Figure 7.14: Effect of number of pulses in the event distribution (110 mbar data).

by high-energy charged particles crossing a given length of material is proportional to its density, and we roughly halved it. The third is that the population of “fake” equal-pulses events is now clearly distinct from the Argon fluorescence peak; this is due to a higher gain for this run, which has a charge multiplication factor of ~ 700 , instead of ~ 390 for the run at 200 mbar, so the events hovering around the pulse threshold are lower in energy. As such, they can clearly be discriminated against.

For this dataset, we get a normalization factor for random coincident events of 1.41. While we still underestimate the rate of coincidences, the difference is less pronounced. This might be either because of the longer run leading to a more accurate estimation of the absolute event rate; or the time cuts that select for stable detector operation improving our understanding of the effect of same-energy cuts.

The Muon FP rate is actually visible in the fit this time, at $0.26 \pm 0.4\%$. We note that the contribution of this FP rate to the fit is driven by the shape difference

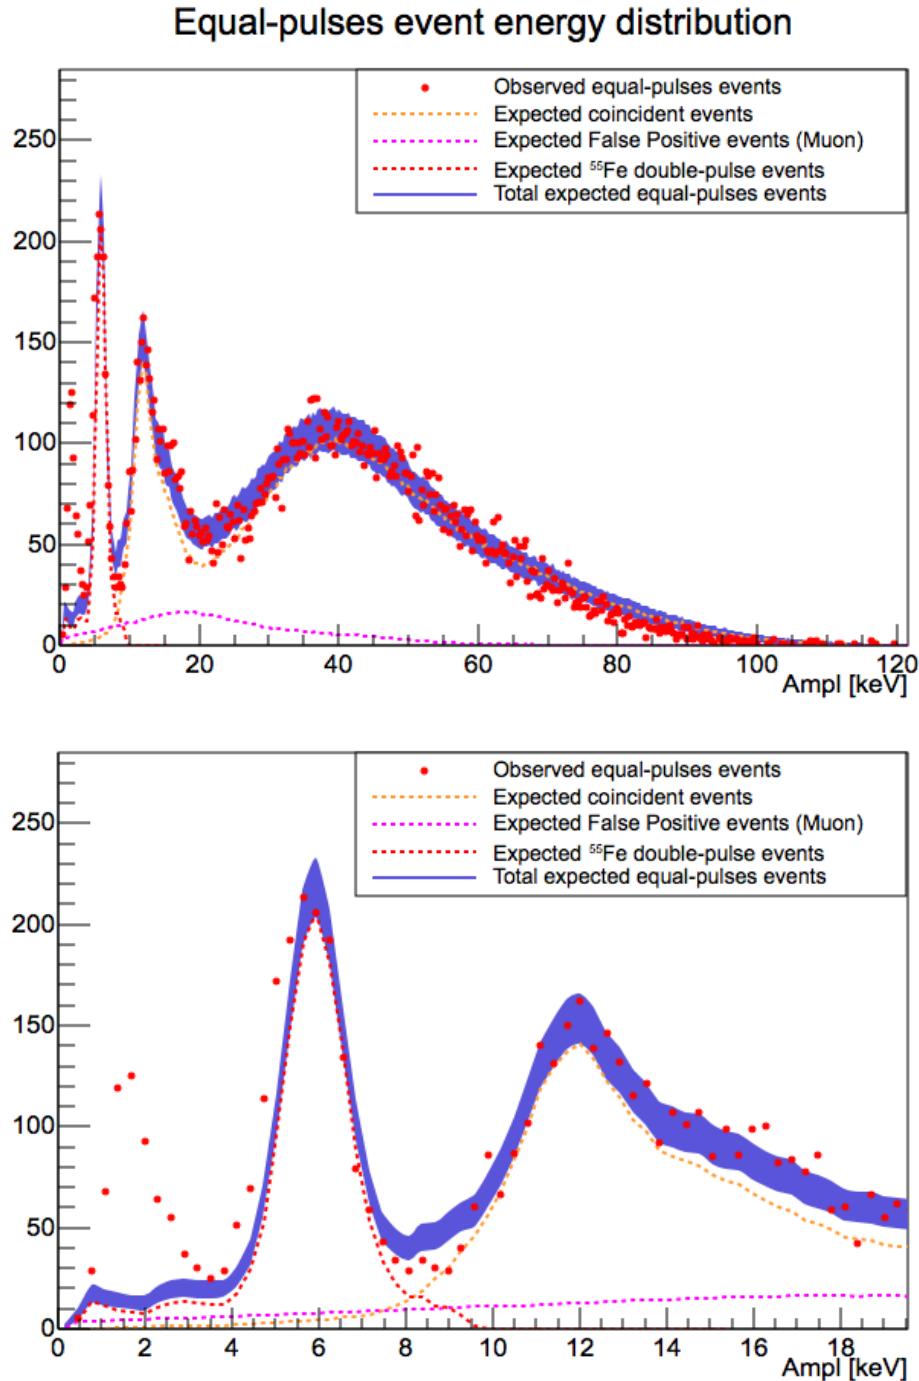


Figure 7.15: 3-parameter fit of equal-pulses event distribution. Top: whole range of the fit. Bottom: Zoom in the region under 20 keV. Red points: Data (110 mbar). Blue curve: Fit result with statistical uncertainty. Dotted lines: Contribution from coincident events (orange), Argon fluorescence events (red), and False Positives (pink). The overpopulation in the 1 – 3.5 keV range are improperly reconstructed events.

between the distribution of expected random coincidence events and the actual events observed around 20 keV. This is the same region we previously used to set an upper limit on the FP rate, which happens to match the actual value found by the fit with this dataset. Due to the fact that no energy range is dominated by Muons falsely reconstructed as equal-pulse events, it is hard to study the apparent increase in their rate under these conditions. Possible culprits may be the higher gain or shorter drift times involved, but how exactly remains unclear.

Finally, the proportion of equal-pulse events from the ^{55}Fe peak found with the fit is $1.96 \pm 0.06\%$. For this run, the higher contribution of Muon events under 10 keV made it harder to disentangle ^{55}Fe events before equal-pulse events, as shown on Fig. 7.16. As such, we only took the number of events from the fitted gaussian as the reference, leading to an artificially lower reported error. Then, taking into account the same correction as before due to the possibility of random coincidences, the measured proportion becomes $2.74 \pm 0.09\%$.

This is to be compared with the expectations from the toy model, after accounting for the effect of the equal-pulse cut, of $3.09 - 3.45\%$. The results are somewhat close, but not quite as consistent as for the 200 mbart data, in part due to the considerably smaller uncertainty on the measured rate. We observe *fewer* distinguishable fluorescence events than expected. While this could be due to an actual mismatch between simulation and reality, the observed increase in Muon events (or other background) in the region of interest, both before and after equal-pulse cuts, negatively affects the confidence with which we take these numbers.

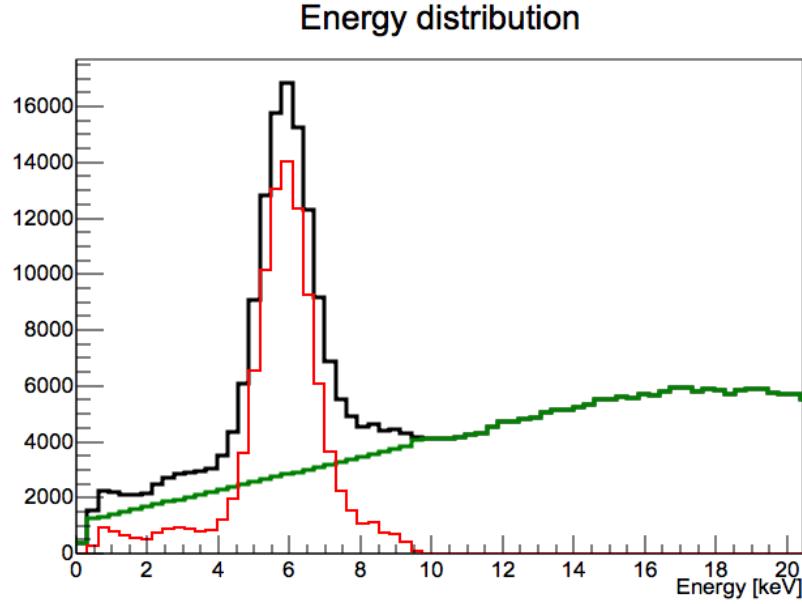


Figure 7.16: 110 mbar data. Black: Total energy distribution of all events after basic cuts. Red: Presumed contribution from ^{55}Fe events. Green: Presumed contribution from Muon events. From their distribution, it is clear that the events naively attributed to ^{55}Fe are overestimated

7.7 Performance conclusions

The final numerical results from both datasets are summarized in Table 7.1. As previously mentioned, the large rate of Muon events presents itself in these results as very large uncertainties for the escape peak under both run conditions. Conversely, since there was no practical way to account for this effect for the calculation of the ratio of distinguishable fluorescence events at 110 mbar, the uncertainty reported in that case is artificially low for the actual data. Nevertheless, a great agreement exists between data and simulations for the 200 mbar data, within the reported uncertainties, even if that agreement is weaker at the lower pressure.

This is further remarkable given the less than ideal conditions in which the data were taken. Due to the large cross-section of the detector and its situation at the

| | Escape peak | | Fluorescence | |
|----------|-----------------|-----------------|-------------------|-----------------|
| | Data | Sim | Data | Sim |
| 200 mbar | $3.2 \pm 1.0\%$ | $3.11 - 3.25\%$ | $2.99 \pm 0.28\%$ | $2.64 - 3.22\%$ |
| 110 mbar | $4.8 \pm 0.9\%$ | $3.97 - 4.38\%$ | $2.74 \pm 0.09\%$ | $3.09 - 3.45\%$ |

Table 7.1: Comparison of the ratio between escape peaks or fluorescence events to the total number of ^{55}Fe events, between simulations and data.

surface, a flood of Muon events made the analysis considerably more complex. Three additional problems had to be fixed: identifying and separating multiple events within the same window, accounting for the effect of space charge in the simulations, and extracting the distribution of ^{55}Fe events from the background for rate normalization. Repeating these same calibrations underground would remove all these difficulties, and drastically improve the final uncertainties. Yet even in these circumstances, a distribution of only ~ 300 events was successfully extracted from a run with $\sim 3 \cdot 10^5$ events, thanks to a False Positive rate for Muons (*i.e.*, track events) under 0.3%, and likely much lower.

Summary

^{55}Fe -induced argon fluorescence provides an appropriate source of axion-like events at 5.9 keV for 12% of ^{55}Fe decays. A calibration was performed with a large 130 cm wide SPC at Queen’s University. The size of the detector and the lack of cosmic radiation protection brought about a rate of ~ 400 Hz of cosmic muons, compared to only 30 Hz of ^{55}Fe . The high muon rate caused 10% uncertainties on the determination of the total ^{55}Fe rate, induced a random coincidence rate of ~ 40 Hz, and changed the electric field inside the detector. Despite these difficulties, the data agreed with the results from a toy model of the detector, with a ^{55}Fe -induced argon fluorescence

detection efficiency of 25% at 200 mbar, and a rate of random coincidences of equal energy within a factor 2 of predictions. Additionally, 99.7% of muon tracks were rejected based on shape alone.

To conclude, this experiment serves as a proof-of-concept for this approach for solar KK axions detection:

- Efficient identification of double events (both fluorescence and random coincidences);
- Strong rejection of non-double events;
- Agreement between simulations and data;

with the first two points proving the adequacy of the method to the problem, and the last being a requirement for the extraction of results from real physics data, which we will tackle in the next chapter.

Chapter 8

Results

After discussing the working principle of the detector (cf. Ch. 3), the analysis tools developed to process its data (cf. Ch. 4), the simulations to estimate the expected signal (cf. Ch. 5), and calibrations to match them with data (cf. Ch. 6), we can finally move onto the dataset used to set constraints on solar KK axions. This chapter covers a description of SEDINE’s physics run, the procedure used to build the aforementioned limits (with related uncertainties), and conclude with the future of solar KK axion searches in the NEWS-G collaboration with the upcoming SNOGLOBE detector.

8.1 SEDINE data

The SEDINE detector was run for 42.7 days, filled with 3.1 bar of Neon and 0.7% of CH_4 , with a central electrode at 2520 V. In that time, 1639360 events were recorded. A preliminary analysis of pulse shapes, as shown in Fig. 8.1, allows discrimination of non-physical events (like noise transients, events without ion signal, voltage supply spikes, etc.).

By looking at the distribution of the time since the previous event (Fig 8.2), we

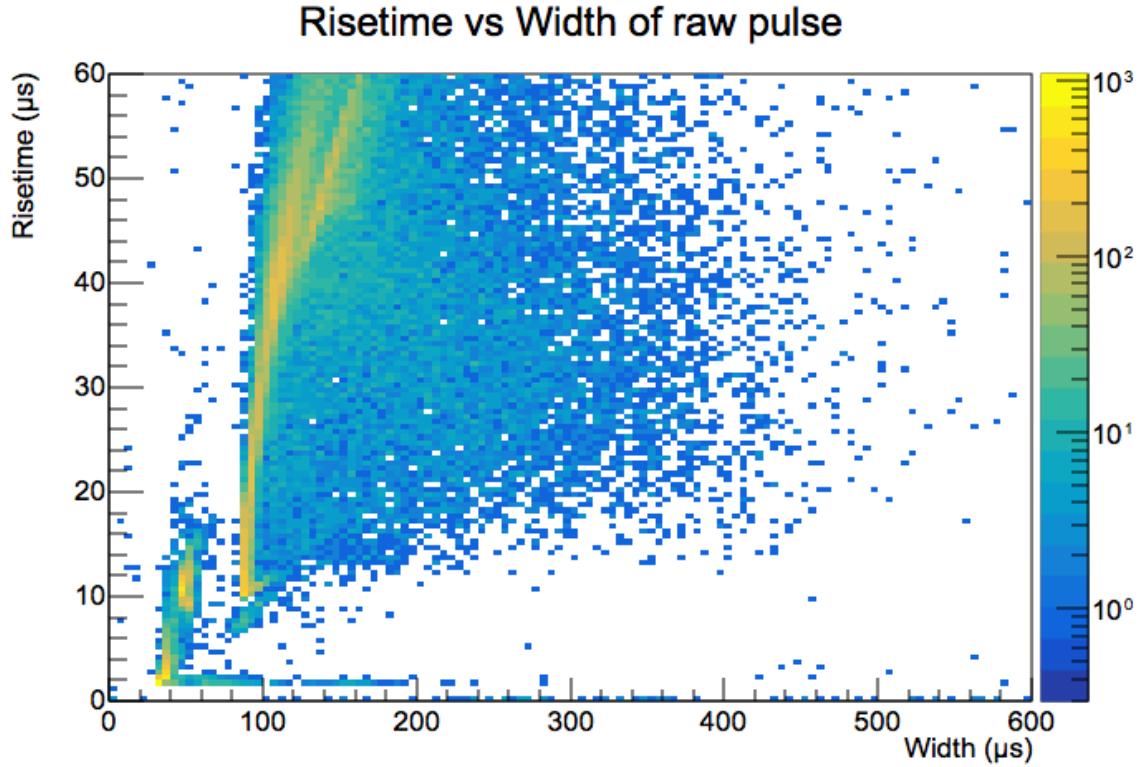


Figure 8.1: Risetime vs Width of raw pulses, excluding only very low amplitude events. Comparing this plot to results from calibrations, we choose to cut to only keep the events with a risetime above $10\ \mu\text{s}$ and width above $75\ \mu\text{s}$.

see that events are not purely independent from each other, with periods of time where the rate is much higher. We are interested in pulses happening close in time to each other, so we need to remove these periods of time where the rate is much higher, since they could induce false positives. Removing all events that happened less than 2 s after the previous one just leads to a loss of effective run time from 42.7 to 38.0 days.

Together with a cut on the risetime and width of raw pulses, we reject events such as the ones shown in Fig. 8.3. After processing the data according to the double

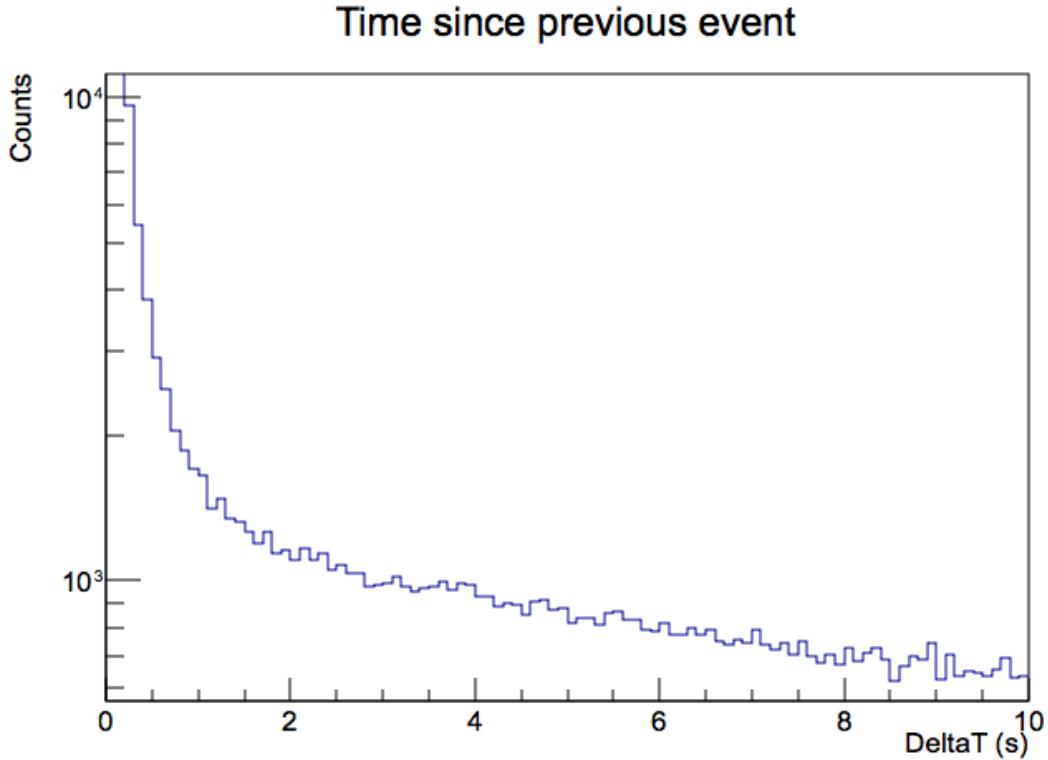


Figure 8.2: Distribution of time since previous event. Above 2 s, the exponential decrease with time is the behaviour expected from events happening randomly with respect with each other. The fast increase under 2 s reveals that there are periods with much higher event rates (or correlated events), that we should remove from the analysis.

deconvolution method (cf. Sec. 4.1.4), we obtain the plot in Fig. 8.4. Some features are visible, such as the accumulation of events at a risetime of $\sim 50 \mu\text{s}$, corresponding to surface events, and a more subtle one around 5 – 8 keV, corresponding to copper fluorescence events (see Ch. 6 for details on how these were used for calibrations).

8.1.1 Basic axion-like cuts

The cuts chosen to select axion-like events are different than the ones chosen to study the general population of physical events. Basic cuts were chosen based on general

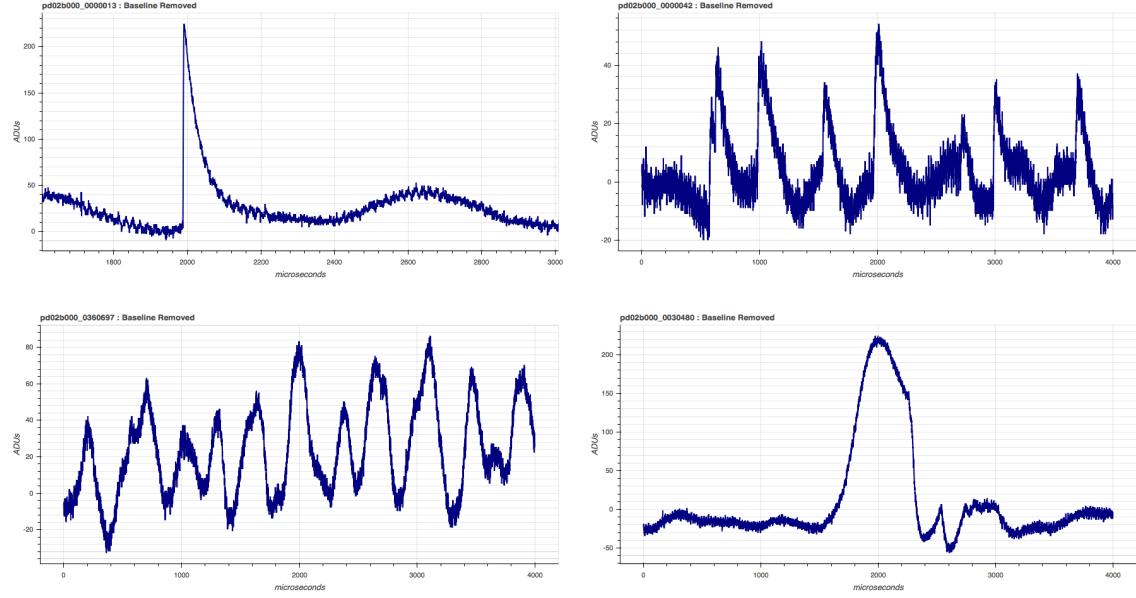


Figure 8.3: Examples of events rejected by basic cuts. Top left: “electronic event” (zoomed-in) with a sharp rise characteristic of absence of ion-current structure. Top right: high-rate period after a large event. Bottom left: “pulse-like” baseline noise. Bottom right: transient noise.

principles of what the expected signal should look like (two pulses of same energy close in time to each other), while trying to limit any loss of sensitivity due to those cuts. The cuts, along with their justification, are:

- $\text{RawRise} > 10 \mu\text{s}$ & $\text{RawAmpl} < 55000 \text{ ADU}$ & $\Delta\text{t} > 2 \text{ s}$

These basic cuts are based on the cuts for regular events. The cut on the width of the raw pulse is dropped, since the pile-up from two pulses close to each other (what we are looking for) distorts the computation of this parameter.

- $(\text{DecMinimum} > -6 \parallel \text{DecMinimum}/\text{DecMaximum} > -0.4)$
- $\text{NPulses} > 1$ & $\text{Ampl}_{P1} > 0.5 \text{ keV}$ & $\text{Ampl}_{P2} > 0.5 \text{ keV}$

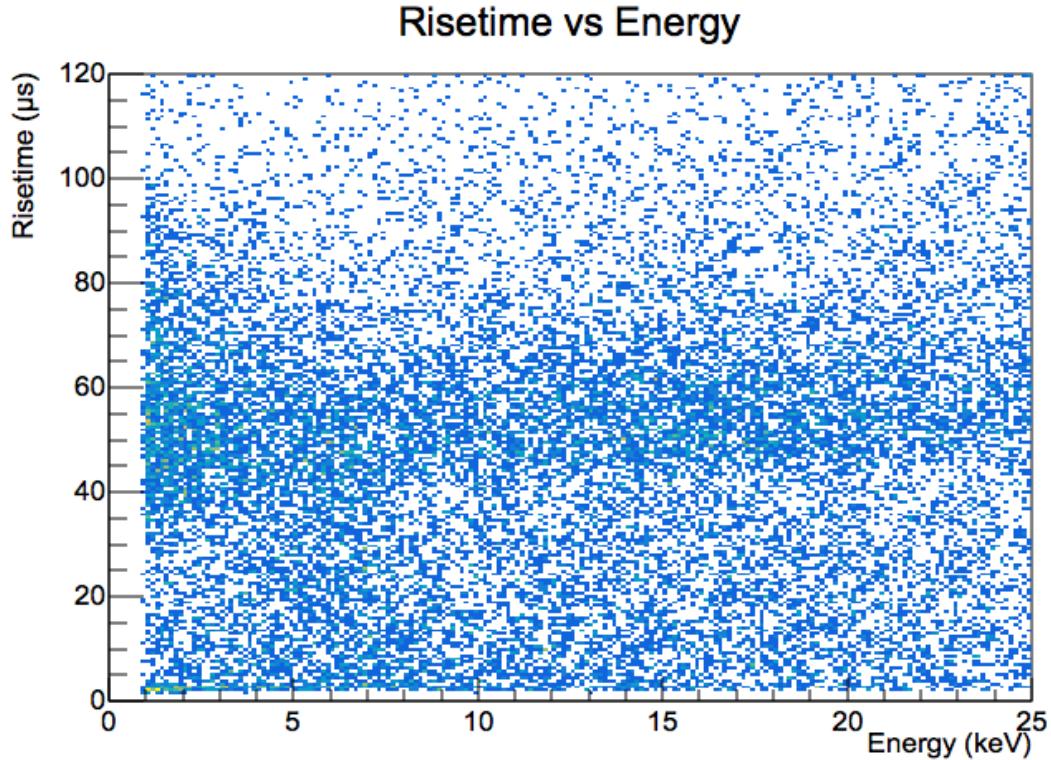


Figure 8.4: Processed risetime vs amplitude of SEDINE physics data.

We then add the most basic cuts for selecting events with two pulses. The restriction on the minimum amplitude of the found pulses, and the cut on negative values of the deconvolved pulse are there to prevent artificial 2-pulse events generated by undershoots or small noise oscillations around the pulse-finding threshold.

- $\text{Ampl}_{P1} + \text{Ampl}_{P2} > 2 \text{ keV} \& \text{Ampl}_{P1} + \text{Ampl}_{P2} < 23 \text{ keV}$ (or replace by energy range)

We restrict the energy range to that where axion decays are expected. $2 - 23 \text{ keV}$ is the maximum range, likely suboptimal; the search for a better range is described in the next section.

- $\text{Rise}_{P1} < 70 \text{ } \mu\text{s} \& \text{Rise}_{P2} < 70 \text{ } \mu\text{s}$

- $\text{Width}_{P1} < 150 \mu\text{s}$ & $\text{Width}_{P2} < 150 \mu\text{s}$

The cuts on the maximum risetime and width remove tracks and other events improperly reconstructed as multiple pulses. The value for the maximum is chosen to be large enough that the vast majority of pointlike events will pass, even if our diffusion time calibrations were too low for some reason.

- $\text{Rise}_{first} < \text{Rise}_{second}$
- $|\text{Centre}_{P1} - \text{Centre}_{P2}| < 500 \mu\text{s}$

Simultaneous event depositions will produce multiple pulse so that the one observed earlier will be narrower (its primary electrons spent less time drifting, and so diffused less with respect to each other). Similarly, the time difference between both pulses is constrained by the maximum drift time in the detector. The maximum value is once again taken on the larger side as a safety measure against miscalibrations.

- $\text{Ampl}_{P1} < 3 \cdot \text{Ampl}_{P2} + 5 \text{ keV}$

Finally, since a decaying axion produces two photons of same energy, the two pulses should be of similar energy. However, due to the resolution of the detector, and the presence of relatively strong attachment, this cut is kept very wide: we just require the reconstructed energy of the largest pulse to be no larger than three times the smallest pulse. An extra safety margin of 5 keV is added to avoid unwanted behaviour due to noise at the lowest amplitudes.

The effect of applying all these cuts is shown on Fig. 8.5. After applying basic cuts to remove unphysical events, the additional cuts with the strongest effect are:

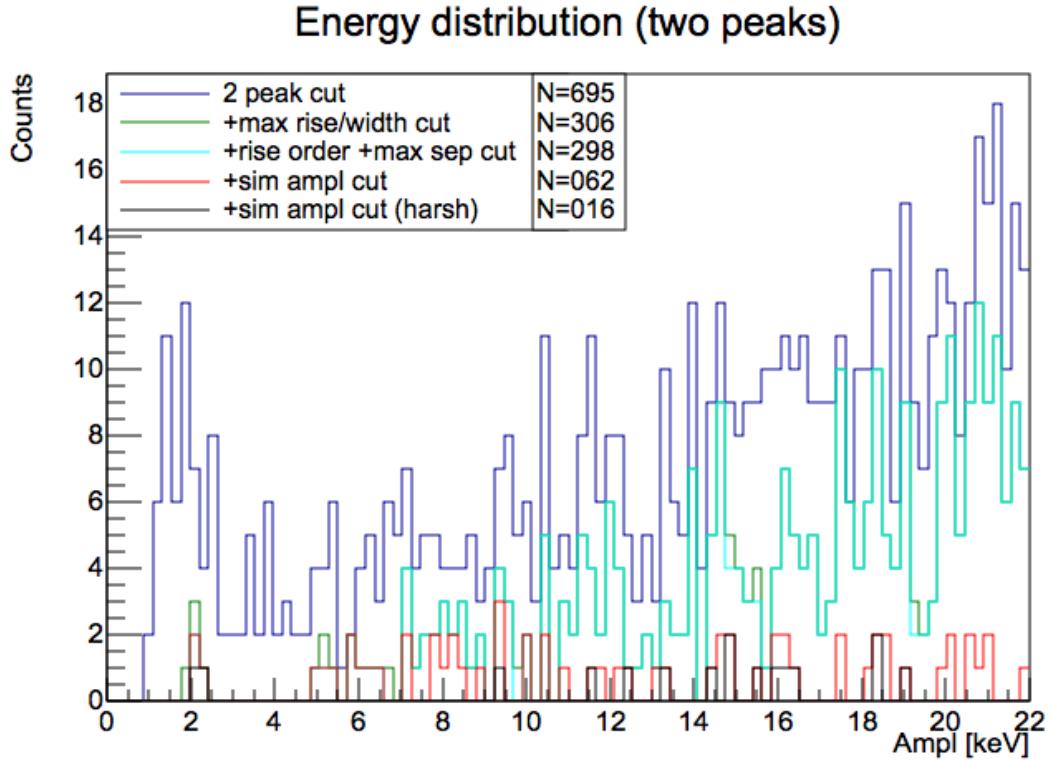


Figure 8.5: Distribution of amplitude of events for different cuts. Roughly, dark blue are physical events with two pulses; green are events with two pointlike pulses; light blue are events with two “simultaneous” pointlike pulses; red are events with two such pulses that are relatively close in energy. For comparison, we also show the results for a harsher cut on the allowed difference between both pulses (black).

- The cut that selects events with two pulses from all “physical” events (not pictured; goes from ~ 21200 to 695 events);
- The cut that removes events with risetimes or widths that are too large (goes from ~ 695 events to 306);
- The cut that keeps only events where both pulses are close in energy (goes from 298 to 62 event, or even down to 16 for a harsher version).

Some conclusions from the effectiveness of the cuts can be drawn. Of note, the combination of the first two cuts just mentioned reduce the number of observed events by 98.5%. This is the great advantage of the search for solar KK axions in a gas detector: by looking only for events with multiple pulses, an extreme rejection power is achieved over solid or liquid detectors, for which both photons from the axion decay would be absorbed at the same location. By itself, the cut that removes events with risetimes or widths that are too large removes events with undershoots, and the high rate periods after a large event, that survived the basic cuts mentioned earlier.

It is interesting to note that the cut that selects for events in which the first pulse is shorter than the second barely has an effect on the number of events. This suggests that random coincidences of pointlike events are extremely rare, if there are any at all; otherwise, we would expect the number of events in which the first pulse is larger to be close to those in which it is not ¹. The few events cut is likely to only be due to the imperfect resolution of the risetime estimation, or events spread out by more than the “maximum” drift time due to electric field anisotropies very close to the rod.

The other very efficient cut is the selection for pulses of similar amplitude. While we do expect some backgrounds that generate simultaneous event depositions (like ^{210}Pb on the inner surface of the shell), no background generates simultaneous pulses of same energy. Unfortunately, due to resolution effects on our amplitude estimation, and the presence of relatively high attachment, the reconstructed amplitude of two events with same energy might be quite different. A lax cut that takes this into account still reduces the event count by 79%. But if we had a clean was with minimum

¹This is consistent with expectations: given the rate of physical events was around 20 mHz, the number of random coincidences expected was ~ 0.7 , even before taking into account any cuts

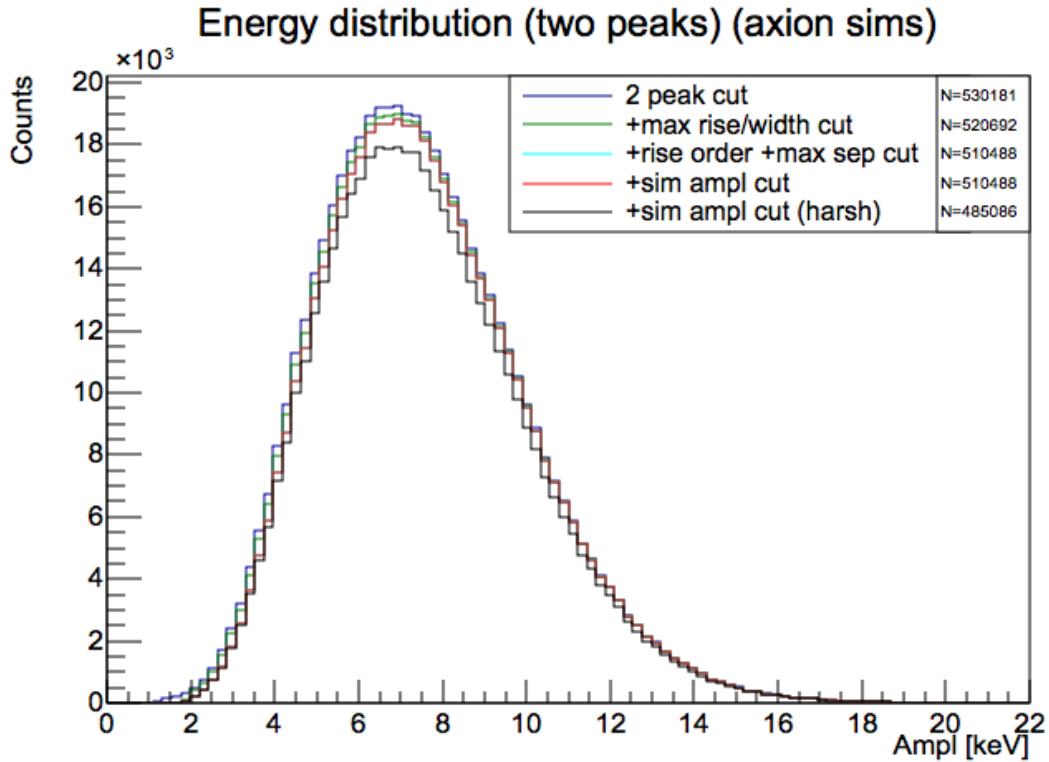


Figure 8.6: Distribution of amplitude of simulated axion events for different cuts. Roughly, dark blue are physical events with two pulses; green are events with two pointlike pulses; light blue are events with two “simultaneous” pointlike pulses; red are events with two such pulses that are relatively close in energy. For comparison, we also show the results for a harsher cut on the allowed difference between both pulses (black).

or no attachment², a harsher cut would have reduced it by 95% instead (under the assumption that the events seen by SEDINE are mostly background where the two pulses have “unrelated” energies).

The effect of these cuts on the data should be contrasted with their effect on our expected axion signal. $2 \cdot 10^6$ decays were simulated, of which $1.96 \cdot 10^6$ left energy inside the detector. Applying the same cuts as for the data leaves us with the results

²SNOGLOBE will have an oxygen getter, allowing higher levels of purity in the target gas than SEDINE.

shown in Fig. 8.6. Keeping only events with two separate pulses has the biggest impact on our efficiency, keeping only 26.5% of all axion decay events (for reference, only 76.5% of all axion decays even had both photons captured within the detector). However, past that point, all other cuts have marginal effects on the total efficiency. Applying all cuts, including the harsher version of the same-amplitude cut, would only drop the efficiency by 8.5% of all double-pulse axion events, compared to up to 97.7% for observed events in our data. This justifies the use of these cuts, since they reject a much larger proportion of background than of signal.

We can contrast data and expected signal in other ways too. After applying all cuts, we end up with 62 candidate events across the whole energy range (which we will be restricting in a moment). Their distribution in energy and time difference, compared to the expected one (normalized) from solar KK axions, are shown on Fig. 8.7. We see a definite difference between the two of them, with the data being more uniform in energy than the simulations, and the time difference peaking around $\sim 170\ \mu\text{s}$ or higher, instead of $\sim 100\ \mu\text{s}$. An “unfortunate” side effect of the low event count is that the statistics do not really allow for a confirmation on the true source of these events, or applying more elaborate background subtraction techniques.

For a final check on the data, we look at the comparison between the first and second pulse for each of these events, shown on Fig. 8.8. Not much can be said from the comparison of risetimes, due to the low statistics. The distribution of amplitudes points against these events coming from two simultaneous events of same energy, since then the first pulse would usually be the largest, due to attachment; we observe the opposite. Then again, we were not expecting to be sensitive enough with SEDINE to observe solar KK axion events, so this is as expected. The main use of analysing

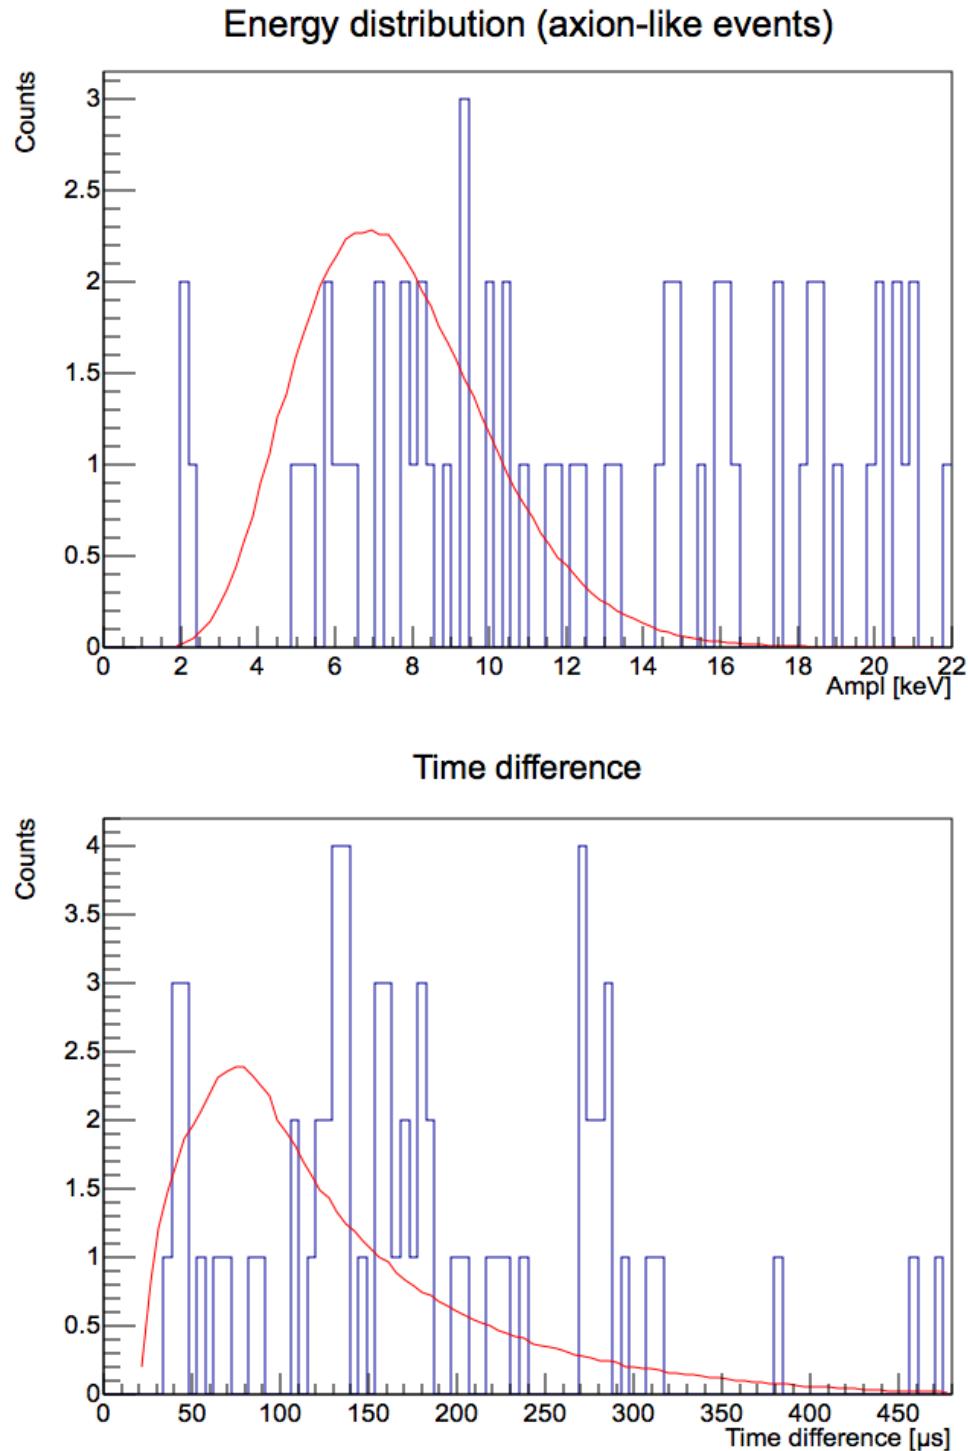


Figure 8.7: Comparison between observed events that pass all axion-like cuts (blue), and normalized expectations from simulations of solar KK axions (red). Top: Energy distribution. Bottom: Time difference between first and second pulse.

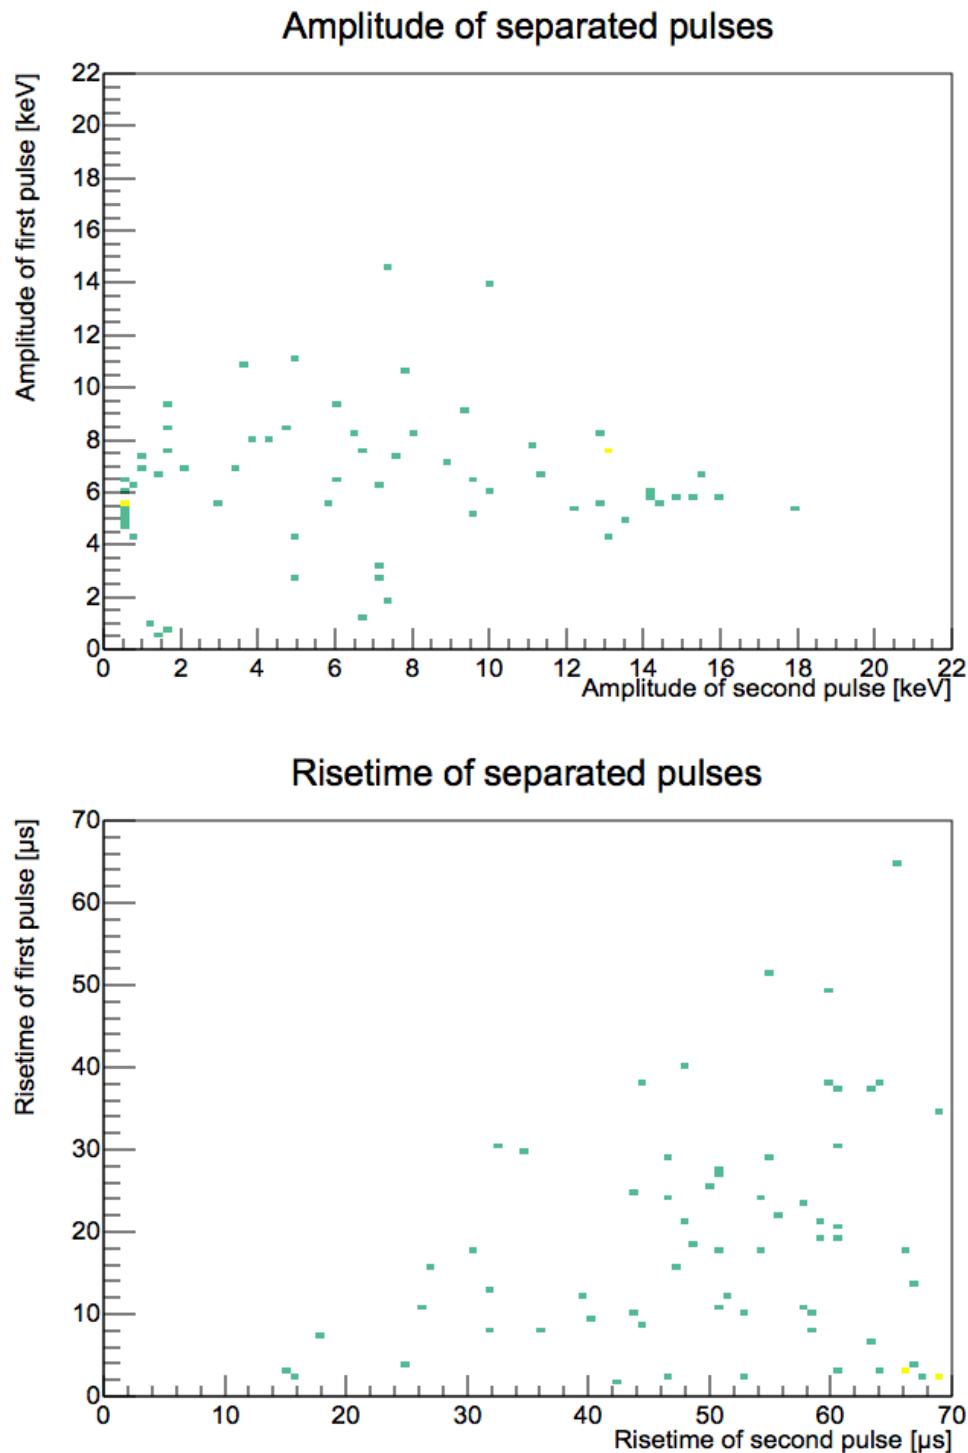


Figure 8.8: Comparison between the first and the second pulse for events that pass all axion-like cuts. Top: Amplitudes. Bottom: Risetimes

these events is to prepare for SNOGLOBE data and its respective backgrounds.

8.1.2 Optimized cuts

Axion-like backgrounds

As mentioned in the previous section, there is no reason to expect the basic axion-like cuts to be the optimal ones for this analysis. In fact, it is reasonable to think that by restricting the range of the search, we can improve the Signal-over-Noise Ratio (SNR, here the number of expected axion events over background events), at the cost of some detector sensitivity to axions³.

As such, to optimize our final results, we require knowledge of all background that might potentially appear as solar KK axion events. The possible background sources of SEDINE were mentioned in Sec. 3.3.2, and quantifying them was one of the main objectives of A. Brossard’s thesis [35]. With the help of his Geant4 simulations, the contributions from the different radioactive contaminations were computed, as shown on Fig. 8.9⁴. These simulations reveal that the primary background is the presence of ^{210}Pb deposited on the inner surface of the detector shell from the ^{222}Rn chain. The second most important contribution is the presence of ^{210}Bi in the bulk of the shell.

We should take a moment to understand how these radioactive contaminations

³How to produce an exclusion limit from number of observed events is explained in the next section. For the rest of this section, it is enough to know that we want as high an efficiency and as few background events as possible. We do not consider the possibility of discovery at this time: the solar KK axion model predicts ~ 0.1 events for SEDINE’s physics data, at least a couple orders of magnitude below the necessary for a discovery claim.

⁴The visible statistical uncertainty for all but surface ^{210}Pb is due to how *rare* these events are, and how long it takes to simulate them. 10^5 (resp. 10^8) events were simulated for ^{210}Pb on the inner surface (resp. in the bulk of the shell) of the detector; 2258 events (resp. 118) passed all cuts. 10^5 (resp. 10^8) events were simulated for ^{210}Bi on the inner surface (resp. in the bulk of the shell) of the detector; 46 events (resp. 331) passed all cuts.

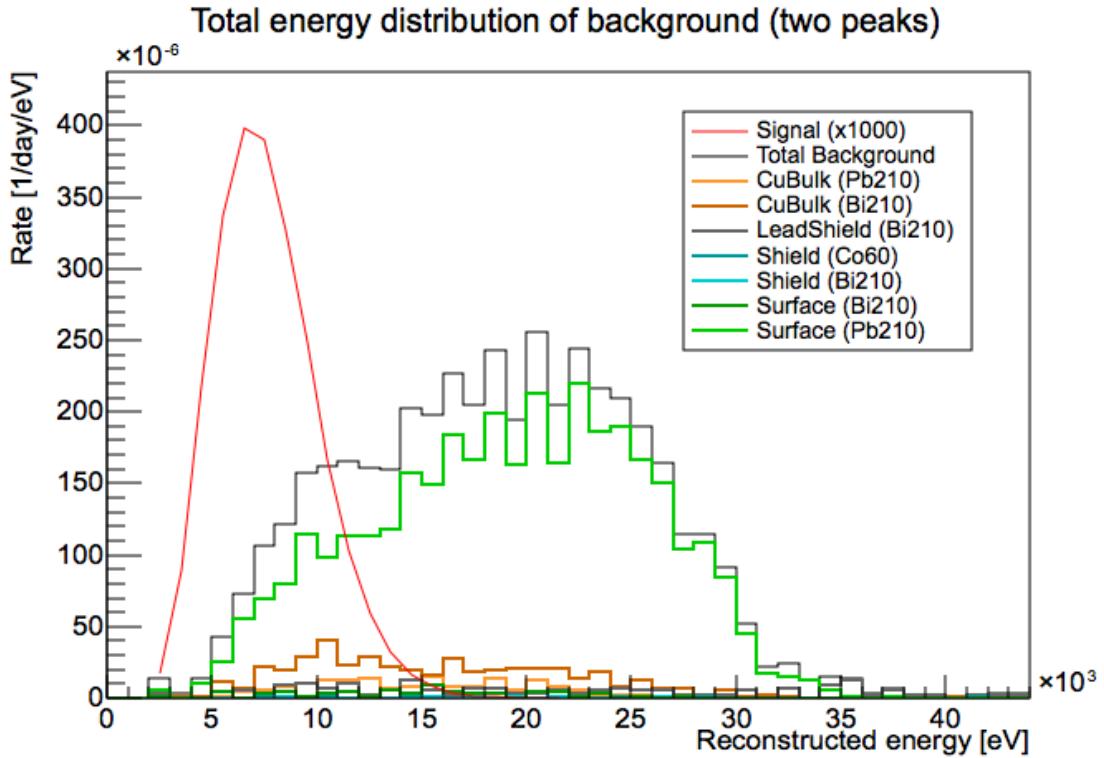


Figure 8.9: Contribution of different radioactive contaminations to total axion-like backgrounds, using basic axion-like cuts.

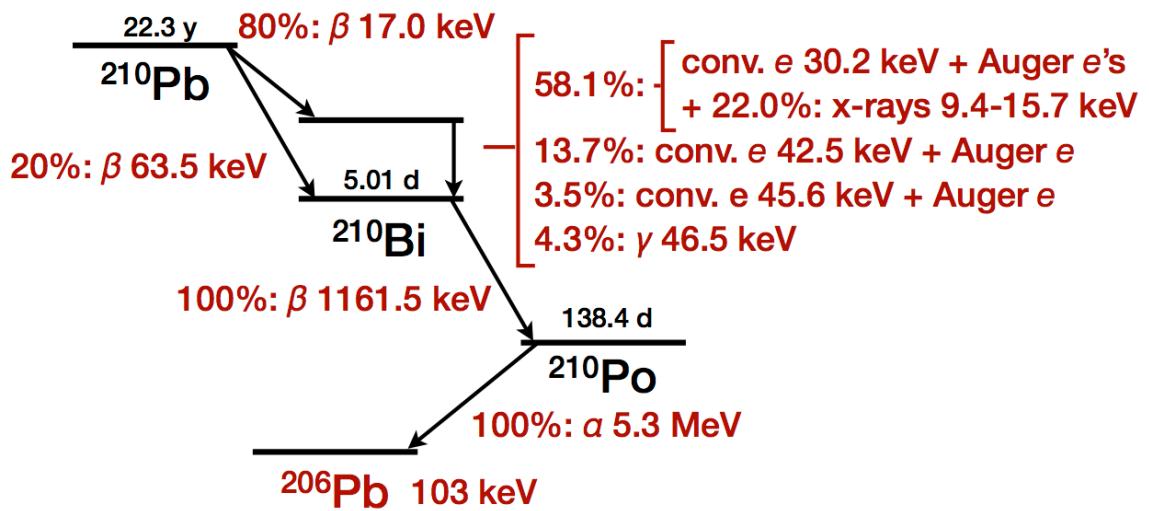


Figure 8.10: Decay chain of ^{210}Pb , with ^{210}Bi [45].

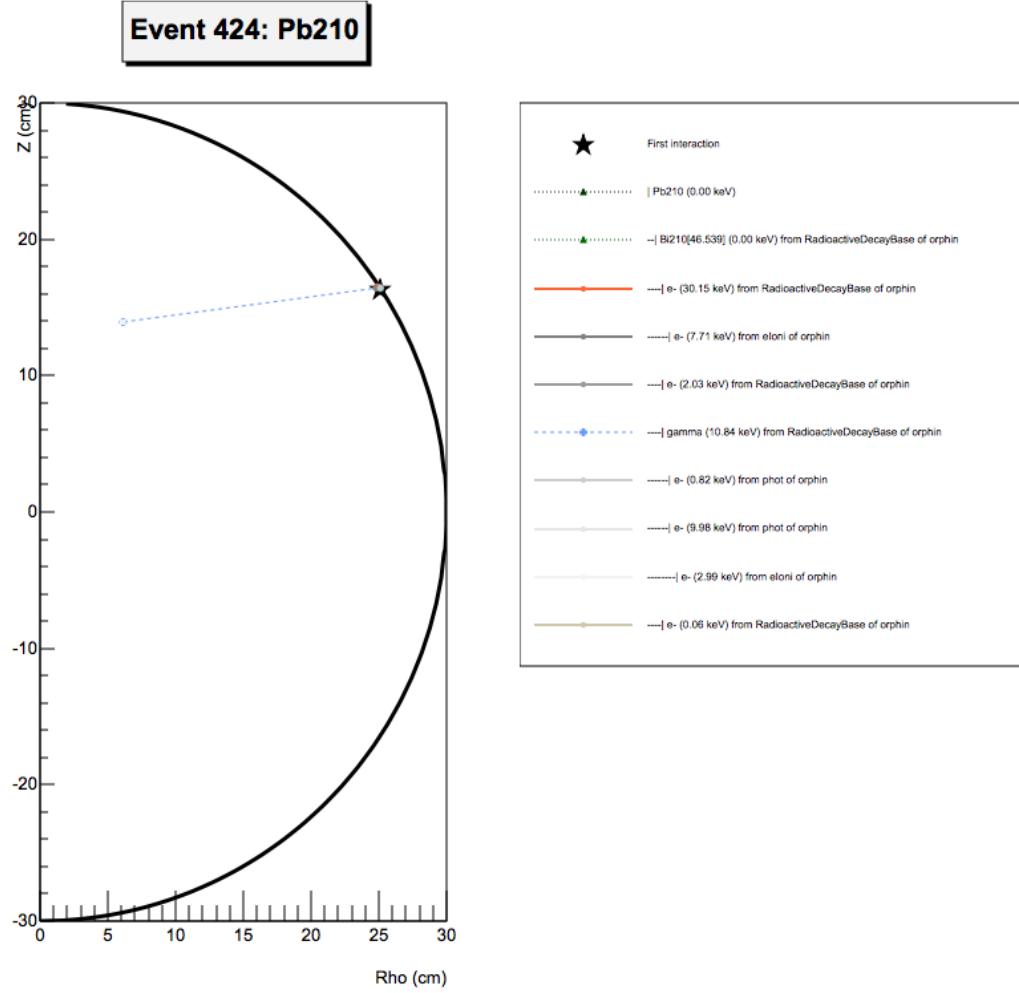


Figure 8.11: Example of ^{210}Pb event simulated by Geant4 that passes basic axion-like cuts. ^{210}Pb decays into an excited state of ^{210}Bi , which de-excites through emission of electrons and a photon. The electrons are either absorbed in the surface or escape out of the detector, while the photon is captured in the bulk of the detector, leading to two separate, simultaneous energy depositions.

produce axion-like events. Fig. 8.10 shows the decay chains of ^{210}Pb . Other radioactive backgrounds either emit alphas (which, being heavy and charged, cannot produce separate energy depositions), or their base rate is too low to be relevant.

For ^{210}Pb , the de-excitation of the Bismuth daughter happens most often by a combination of emission of electrons and photons. A photon in the $\sim 10\text{ keV}$ energy range travels some distance in the gas before being captured, while the electrons interact almost instantly. Since these events all involve relatively low energies, they are dominated by the ^{210}Pb contamination at the inner surface of the detector shell, where the electrons are not stopped by the bulk of the copper. From the decays that leave energy inside the detector, up to 2.3% (for surface decays) produce axion-like signals. We end up with a rate of 3.5 evt/day from surface contamination, and an extra 0.16 evt/day from ^{210}Pb in the rest of the copper bulk.

For ^{210}Bi decays, come in two varieties: either we see a Bremsstrahlung photon (from the β particle crossing the copper bulk) interacting twice in the detector (more likely for decays farther from the detector, cf. Fig. 8.12), or we see both the electron crossing the detector (generally staying close to the surface) and a Bremsstrahlung photon interacting once deep inside (more likely for decays closer to the inner surface of the copper shell, cf. Fig. 8.13). In either case, since the released particles are so high in energy, the location of the decay in the detector shell matters less than for ^{210}Pb , so the extra contribution from the bulk dominates: 0.46 evt/day from the bulk, compared to 0.07 evt/day from the surface. We note also that the proportion of axion-like signals from all decays that leave energy in the detector is lower than for ^{210}Pb , with only 0.07% of surface ^{210}Bi decays producing an axion-like event, and only 0.0003% of ^{210}Bi from the copper bulk.

Another “background” to take into account, while not directly related to any radioactive source, is the fluorescence of copper, at 8.1 keV. While ^{210}Bi decays in particular generate particles at high energies, and as such are not in theory more

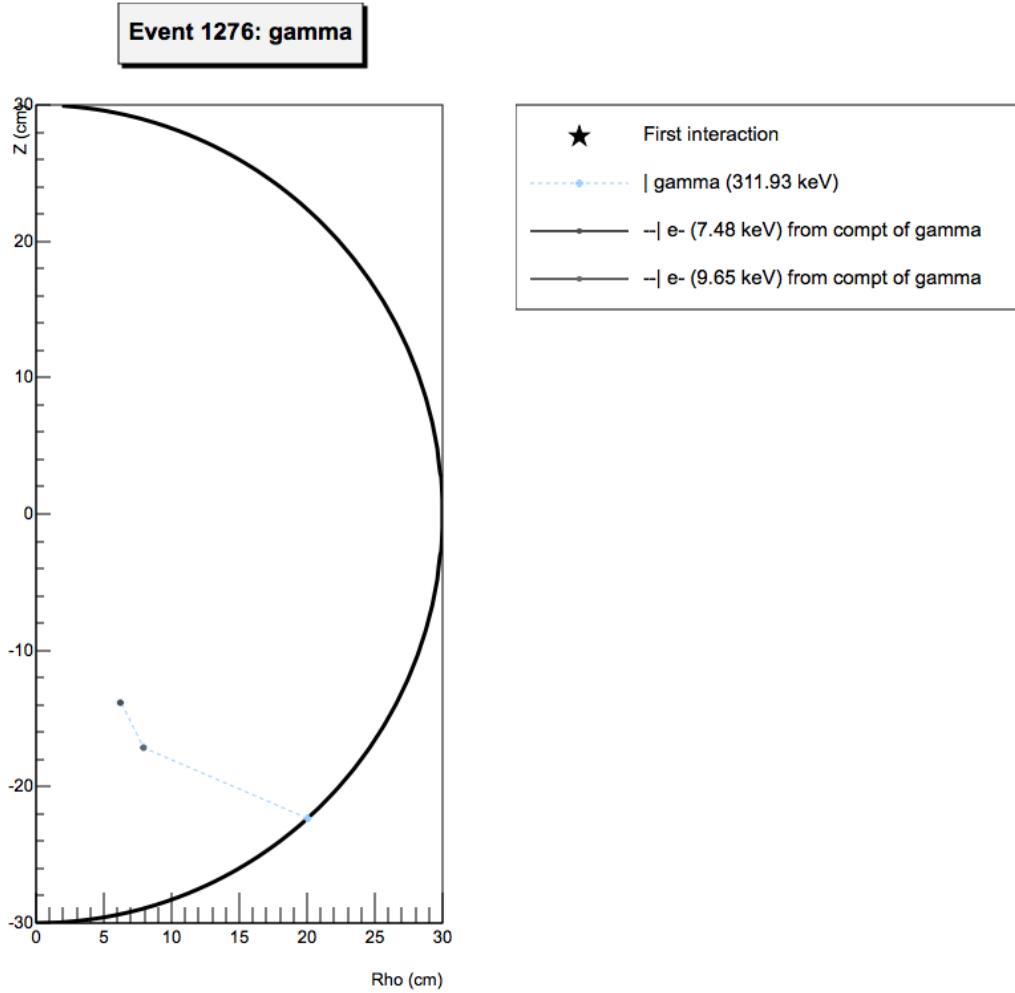


Figure 8.12: Example of ^{210}Bi event simulated by Geant4 that passes basic axion-like cuts. A Bremsstrahlung photon interacting twice in the detector.

likely to generate events at any specific energy, they may produce photons at a fixed energy through fluorescence in the copper bulk. As such, we expect a higher rate of axion-like background events around 16 keV than at other energies, though this is likely drowned out by resolution effects, and the pairing of one fluorescence photon with a random energy deposition.

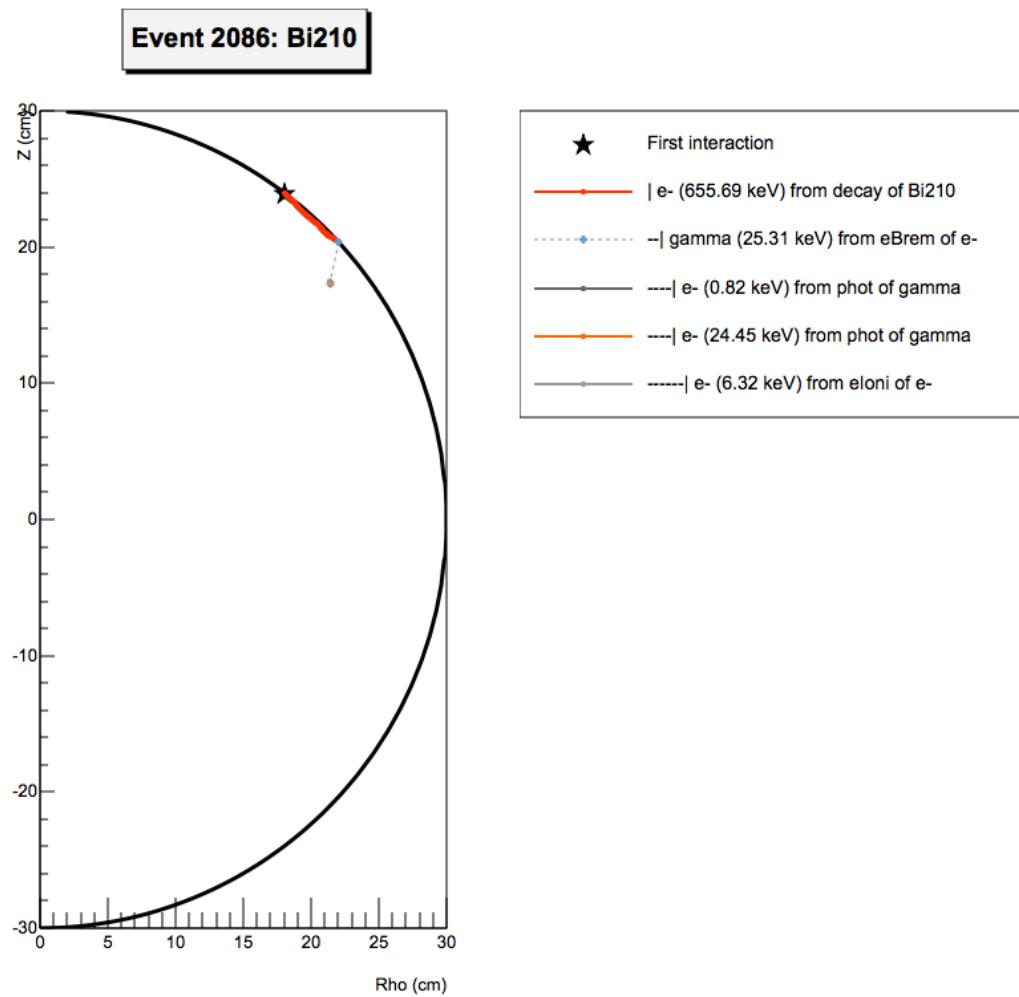


Figure 8.13: Example of ^{210}Bi event simulated by Geant4 that passes basic axion-like cuts. The high-energy electron from the decay of ^{210}Bi interacts in the gas for a short distance before escaping, while a Bremsstrahlung photon interacts deeper in the detector.

To improve our background rejection, we can consider each event to be characterized by 5 parameters:

- Its reconstructed energy;
- The asymmetry between the amplitude of the first and the second pulse, $(A_1 - A_2)/(A_1 + A_2)$;
- The risetime of the first and the risetime of the second pulse;
- The time separation between both pulses.

Ideally, cuts would be selected taking into account all these parameters in once, through methods like Boosted-Decision-Trees. In practice, even just optimizing for some of these parameters one by one provides significantly improved results.

Improved risetime cut

Fig. 8.14 shows a clear difference in the distribution of events in the Risetime vs. Risetime plot for axion events and background events. While the risetime of the first pulse is pretty evenly distributed among all possible values for both axion events and background events, the risetime of the second pulse is concentrated at the high values for background events. The simulations agree with the data on this point: although in the data we see both events in which the risetime of the second pulse is low and in which it is high, there are clearly more of the second category.

This is in accord with the source of the background as we just described. The dominant source of background is ^{210}Pb on the inner surface of the copper shell, releasing low energy electrons and a photon at the same time. The low energy electrons

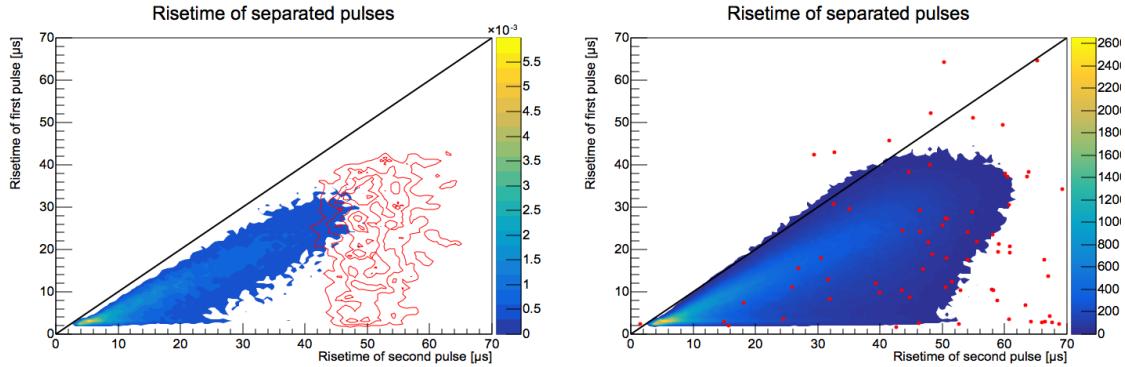


Figure 8.14: Risetime of first pulse vs. Risetime of second pulse, after basic axion-like cuts. In both plots, the colour distribution is from axion simulations. The black line shows the requirement that the risetime of the first pulse be smaller than that of the second. Left: In red, distribution from simulations of SEDINE background. Right: In red, distribution from the SEDINE physics data. There is a larger proportion of background events where the second pulse has a risetime consistent with surface events.

always release their energy at the decay site, on the surface, generating a pulse at high risetimes. Conversely, axions can decay anywhere in the detector, so both the pulses they generate can have any possible risetime.

By looking at the distribution of risetimes of the second pulse for both axion and background simulated events, as shown on Fig. 8.15, we can set a more stringent cut on the risetime to improve our Signal-over-Noise Ratio. In this case, removing all events with a risetime above $40\ \mu\text{s}$ should remove the vast majority of surface background, while keeping most axion events.

Improved asymmetry cut

Fig. 8.16 shows that for axion events, the value of the asymmetry is concentrated around ~ 0.1 . While both the pulses from an axion decay have the same energy, the

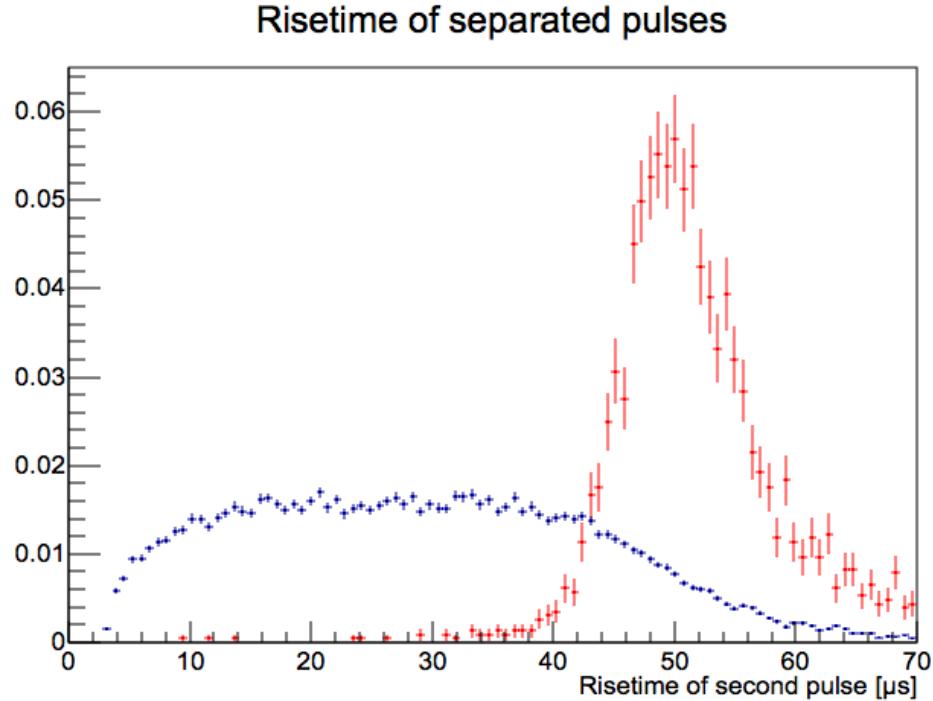


Figure 8.15: Distribution of Risetime of the second pulse for axion events (blue) and background events (red), based on simulations.

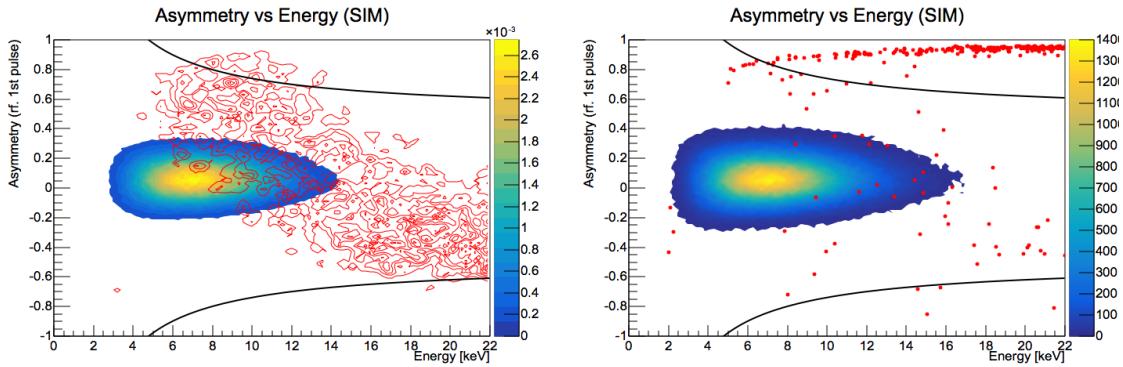


Figure 8.16: Asymmetry between pulse amplitudes vs. Reconstructed event energy, after basic axion-like cuts. In both plots, the colour distribution is from axion simulations. The black line shows the asymmetry cut from the basic axion-like cuts. Left: In red, distribution from simulations of SEDINE background. Right: In red, distribution from the SEDINE physics data.

effect of electron attachment means that the photon absorbed closer to the central electrode will have more of its charge going through the avalanche process than the photon absorbed farther away. The asymmetry for background events is spread out over all possible values, although with a clear dependency on total energy. This overall trend is likely due to the combination of a fluorescence photon (of fixed energy) with a surface electron (leaving a more “randomized” amount of energy in the detector). In the data, the additional population with very high asymmetries is due to a small population of events that were deconvolved improperly: an artificial oscillation behind the pulse is formed in some cases, which is incorrectly reconstructed as a small pulse.

As shown in Chap. 6, our calibration data for attachment does not allow for a precise measure of this phenomenon to be reproduced in our simulations. As such, the basic cuts on the asymmetry were left deliberately lax. In the simulations, the odds of one electron not being captured during a step of length dt were set to $e^{(\alpha(E)-\beta(E))*dt}$, where α is the Townsend coefficient and β is the attachment coefficient, both taken from Magboltz and depending on the energy of the electron⁵. To include our uncertainty on the measure of the attachment, β was multiplied by a random number chosen uniformly between 0 and 2, giving a very conservative sweep of the possible values.

Due to the large uncertainties on attachment, a different method was used to select the improved asymmetry cut than for the risetime cut. In Fig. 8.17, the distribution of simulated axion events is shown for various values of the attachment, including the most extreme values. The lower the attachment, the lower the asymmetry between both pulses, and the higher the reconstructed event energy, and vice-versa for higher

⁵This was set to 1 if larger than 1, the Townsend coefficient only being there to offset the large chance of attachment once inside the avalanche region.

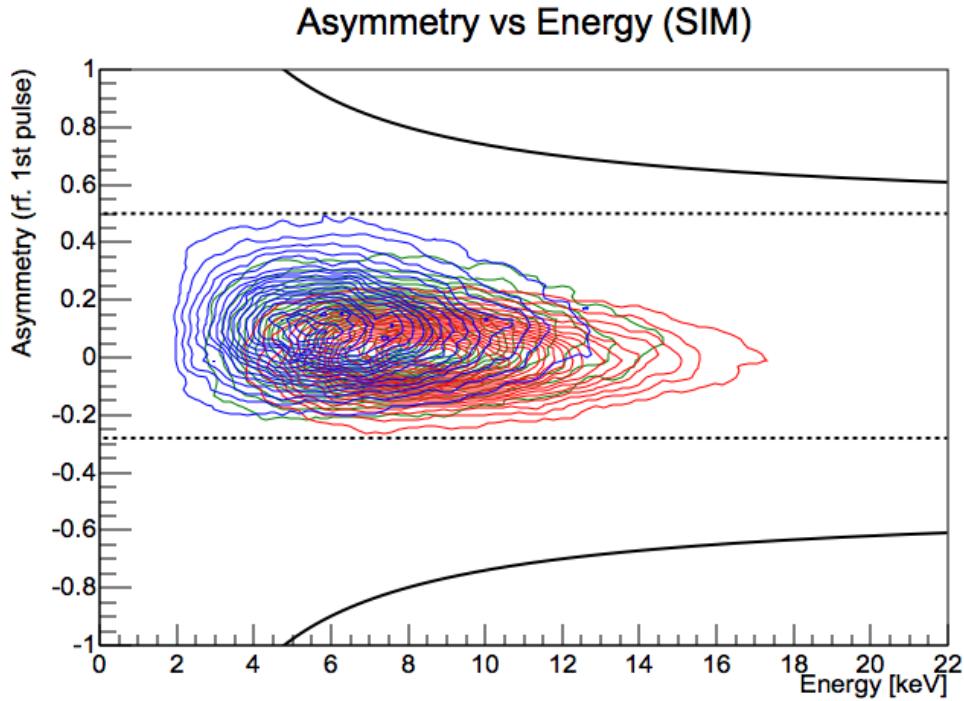


Figure 8.17: Asymmetry between pulse amplitudes vs. Reconstructed event energy of simulated axion events, after basic axion-like cuts, for high (blue), average (green) and low (red) values of the attachment. The solid black line shows the asymmetry cut from the basic axion-like cuts, the dashed line shows the improved asymmetry cut.

attachments. To minimize the effect on systematics (cf. App. K) due to attachment uncertainties, the new asymmetry cut was selected as to include both extremes: a maximum asymmetry of 0.50, and a minimum of -0.28 .

Improved energy range

Using the improved risetime and asymmetry cuts leads to a reduction of over a factor 10 in background in the energy range in which we expect axion events, as shown on Fig. 8.18. The reduction is especially due to rejection of surface ^{210}Pb events, leaving the remaining background evenly distributed between surface ^{210}Pb , ^{210}Bi in

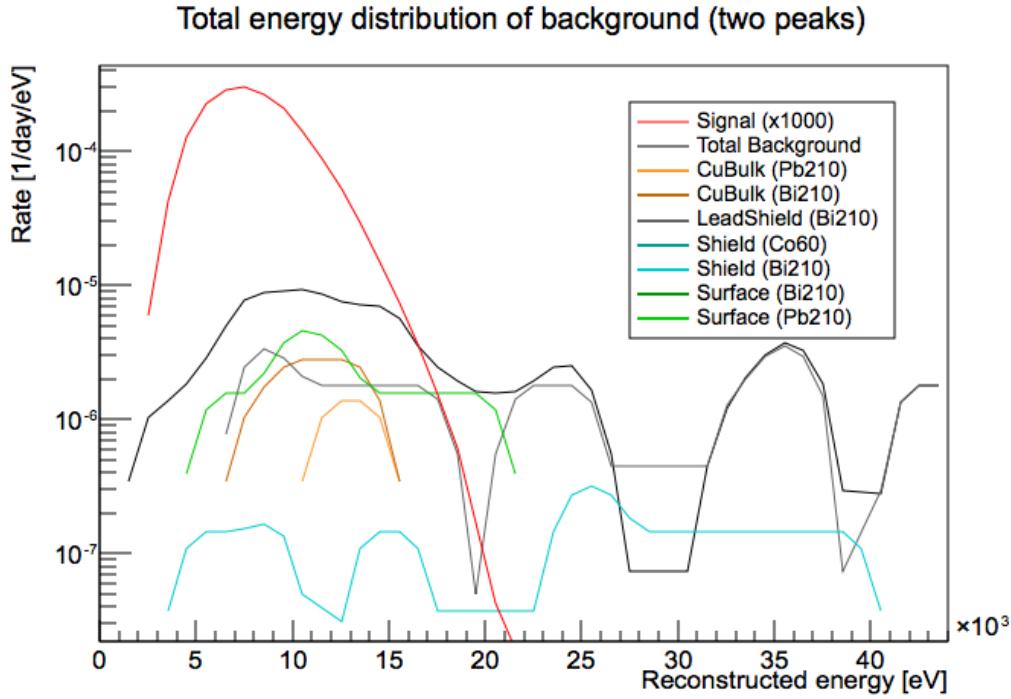


Figure 8.18: Contribution of different radioactive contaminations to total axion-like backgrounds, using axion-like cuts with improved asymmetry and rise-time cuts. The contributions have been smoothed out for better visualization.

the copper bulk, and ^{210}Bi in the lead shield. But improvements can still be obtained by selecting a restricted energy range.

One approach to do so is by comparing the expected axion signal to the expected background at different energies. By selecting only the energy ranges with a SNR above a given threshold, integrating the expected number of axion and background events separately, and comparing them to each other, we can draw a function with a “limit-like” parameter depending on the SNR threshold (cf. Fig. 8.19). It is then just a matter of selecting the value that optimizes the limit, and converting that back into an energy range, to be used with the real data.

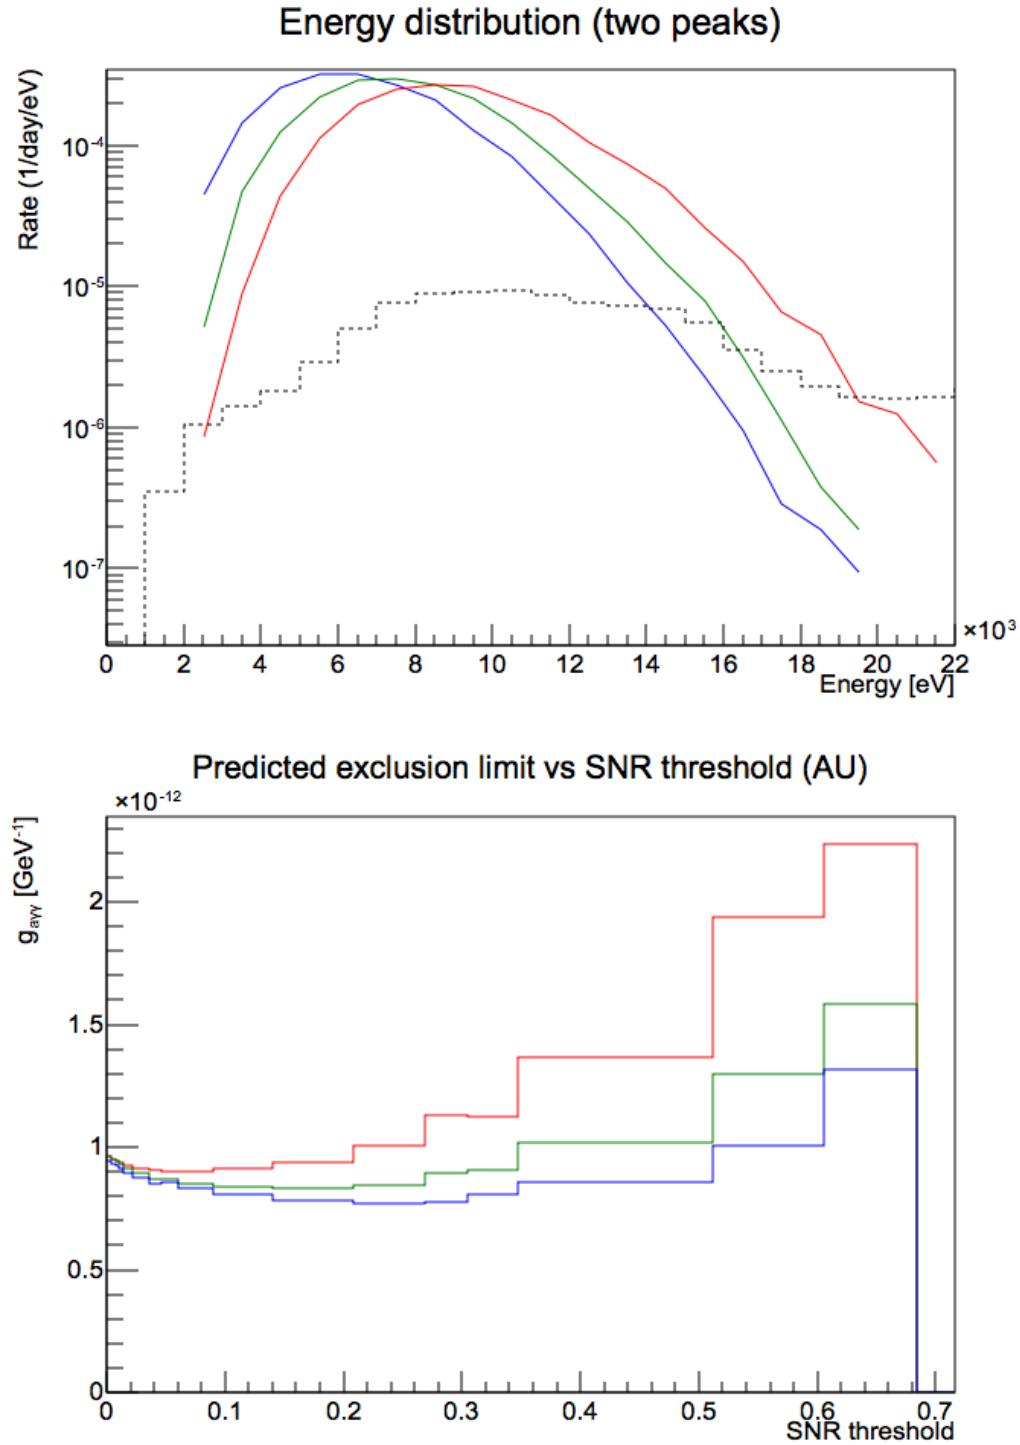


Figure 8.19: Top: Event rate vs. energy for axion and background events (black). Bottom: limit-like parameter when keeping all energies such that their SNR is above a given threshold. The green curves are for the expected value of electron attachment, the blue curves for its maximum allowed value, and the green for its minimum.

There is one additional difficulty due to the previously mentioned uncertainty on electron attachment: the reconstructed energy of an event will tend to be lower with higher electron attachment. This is clearly visible in the top figure in Fig. 8.19, in which the axion events are reconstructed at different values for different attachment coefficients. If we were to choose the SNR threshold that optimizes the case with the expected attachment ($SNR_{min} = 0.17$), we would end up with an energy range of $3 - 10$ keV. This would lead to a drastic decrease in detector efficiency if the low-attachment scenario.

Instead, we choose the SNR threshold so that it minimizes the expected exclusion limit in the *worst* case scenario, which here ends up being the low-attachment case. This results in a decrease in the threshold ($SNR_{min} = 0.08$), and an increase in the energy range, to $3 - 12$ keV. While this effectively means less background will be rejected (due to the weaker energy constraint), it severely reduces the impact of attachment uncertainties on the detector efficiency: the low-attachment case only loses a relative 15.8% with a $3 - 12$ keV energy cut (compared to no energy cut), but 36.5% with a $3 - 10$ keV energy cut.

Final cuts

The improved asymmetry, risetime and energy cuts have a combined effect of reducing the efficiency of the detector by 6.1% of the efficiency with the basic axion-like cuts, for a total efficiency of 16.9% and axion event rate of 0.015 evt/day. On the other hand, the expected rate for background events is reduced from 2.9 evt/day to 0.062 evt/day,

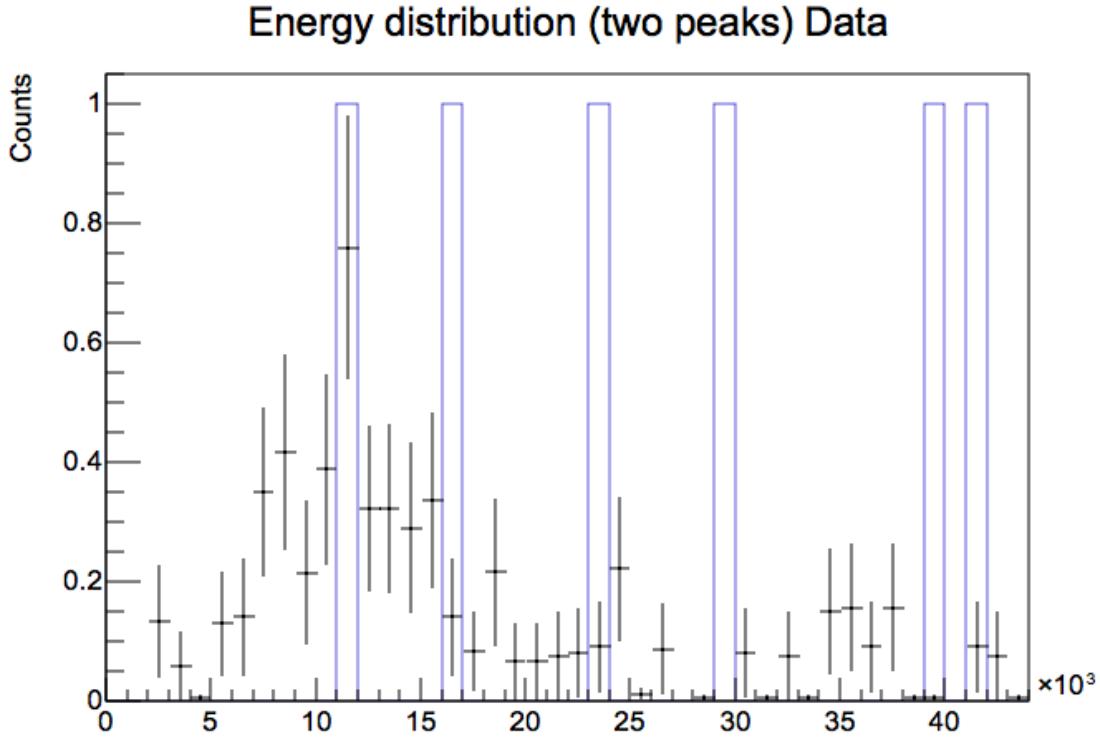


Figure 8.20: Comparison of axion-like events in data (blue), and expectations from simulations of radioactive background (black), after improved axion-like cuts. Only 1 event in the data passes the improved cuts.

for a reduction of almost a factor 50⁶. The comparison with the real data is shown in Fig. 8.20: only 1 event passes the improved cuts, compared to the 62 with the basic axion-like cuts. We see a qualitative agreement in the rate and energy distribution of events in an extended energy range up to 44 keV, although the low statistics prevents drawing conclusive results.

It is very important to note that, while we did look at the data in the previous section, the optimized cuts are based *entirely* on simulations. Selecting the energy

⁶While the rate of background events is still four times higher than that of axion events even after cuts, this does not have much of an effect on the results drawn from the SEDINE data, due to the relatively low exposure of $4.3 \text{ day} \cdot \text{m}^3$. However, it does severely limit the usefulness of increasing the exposure with this detector in an attempt to obtain improved constraints on the KK axion.

range (or risetime, asymmetry, etc.) for analysis based on the data itself would obviously grossly bias the final result (barring special cases like Yellin’s optimum interval method [77]). This method of selecting the region of interest produces conservative conclusions: even if all the improved cuts had turned out to be “wrong” due to an extreme misunderstanding of our background, our final limit would not be “false”; it just would not be as strict as it could have been by selecting a better range.

8.2 Exclusion limit

8.2.1 Theory

The number of expected events for a given axion-photon coupling coefficient $g_{a\gamma}$ is

$$N_{exp} = \text{exposure} * \text{rate} * \text{efficiency} \quad (8.1)$$

The differential rate of solar KK axion decays is:

$$\frac{dR}{dm_a} = \frac{g_{a\gamma\gamma}^2}{64\pi} n_0 m_a^3 f(m_a)$$

where $g_{a\gamma\gamma}$ is the coupling between axions and photons, n_0 is the local density of trapped solar KK axions on Earth, m_a is the mass of the axion, and $f(m_a)$ is the function describing their mass distribution.

The exposure of a given physics run is simply the volume of the detector $V = \frac{4\pi}{3} r_{det}^3$ multiplied by the time length of the run T . The efficiency is the ratio of solar KK axion events that are properly identified as such. Nominally, it depends on the mass of the decaying axion, $\epsilon(m_a)$, but since their mass distribution is fixed in this model, we can integrate that dependency away:

$$\int m_a^3 f(m_a) \epsilon(m_a) dm_a = F \epsilon$$

where we define $F = \int m_a^3 f(m_a) dm_a$, and ϵ -the value that satisfies this equation- as the total efficiency.

Putting all this together, we get:

$$N_{exp} = \frac{4\pi}{3} r_{det}^3 T * \frac{g_{a\gamma\gamma}^2}{64\pi} n_0 F * \epsilon \quad (8.2)$$

This equation can be turned into an exclusion limit based on data. For a given excluded number of events N_{excl} from any given run, and denoting g_{ZDL} and n_{ZDL} as the predicted values of $g_{a\gamma\gamma}$ and n_0 in [6], then the excluded coupling between axions and photons, depending on the density of axions on Earth is:

$$g_{a\gamma\gamma_{excl}}(n_0) = \sqrt{\frac{48 N_{excl}}{r_{det}^3 T n_0 F \epsilon}} = g_{ZDL} \sqrt{\frac{N_{excl}}{N_{exp}} \frac{n_{ZDL}}{n_0}} \quad (8.3)$$

Note that, given a fixed solar KK axion model, and given some excluded number of events N_{excl} , the only variable left to produce a limit on $g_{a\gamma\gamma}$ is ϵ , the efficiency of the detector. This value is obtained through simulations, and is the source of the uncertainty in the final limit. This effect will be explored more in detail in the next section.

Finally, to obtain the number of *excluded* events based on the number of *observed* events, we need to use the properties of the Poisson distribution. Since axion events are random in time, and independent from each other, the number of observed events in a given time follows a Poisson distribution, with only parameter the average number of events in that time (equal to the event frequency, times the duration of the observation). For a given experiment, the lowest average number of events such

that the observed number of events or higher is at most 10% likely (i.e., the upper confidence level at 90%) follows the following formula:

$$N_{excl} = F_{\chi^2}^{-1}(0.9, 2(N_{obs} + 1))/2 \quad (8.4)$$

where $F_{\chi^2}^{-1}(p, N)$ is the quantile function of the χ^2 distribution with N degrees of freedom [78].

One consequence of this formula is that, to be able to reject the existence of solar KK axions as currently modelled at 90% confidence level, we need to run an experiment with enough exposure to expect *at least* 2.3⁷ signal events, assuming no other background. Below that, even observing 0 events is still consistent with existence. The converse, rejecting non-existence at 99.9% confidence level, depends on the distribution of background events: the number of observed events has to be at least three standard deviations above the expected from background alone.

8.2.2 KK Axion limit from SEDINE data

We now have enough information to generate the exclusion limit on solar KK axions derived from our data. Referring to Fig. 8.20, we see that the number of candidate events in SEDINE’s physics data after optimizing our region of interest was 1, leading to an excluded average number of events at 90% of 3.89. This is compared to the expected number of axion events for the run: at a total rate of $0.08371 \text{ m}^{-3}\text{day}^{-1}$, an exposure of $4.303 \text{ m}^3\text{day}$, and an efficiency of 16.34% (after accounting for systematics, cf. App. K), the expected number of events is 0.0589. While the value is almost two orders of magnitude too low for a discovery claim, exclusion limits can still be set based

⁷ $F_{\chi^2}^{-1}(0.9, 2)/2 \simeq 2.3$.

on the Eq. 8.3. At $n_0 = n_{ZDL} = 4.07 \cdot 10^{13} \text{ m}^{-3}$, and given $g_{ZDL} = 9.2 \cdot 10^{-14} \text{ GeV}^{-1}$, we get an exclusion limit of $g_{a\gamma\gamma} = 7.76 \cdot 10^{-13} \text{ GeV}^{-18}$.

The effect of the different systematics on the detector efficiency is covered in detail in App. K. The parameters considered were the uncertainties on the calibration of the electron drift and diffusion time, electron attachment, photon attenuation length, energy and mean ionization energy calibrations, and ion mobility in the gas. For the chosen optimized cuts, the main contributions to the final uncertainty on the exclusion limit were the diffusion time (due to the strengthened risetime cut to reject surface events), at 12% of the total efficiency, and the drift time (due to being the main factor affecting the separability of axion pulses), at $\gtrsim 9.5\%$ of the total efficiency. The electron attachment uncertainty could have had an effect of up to 35% on the total efficiency if we had tried to optimize the cuts ignoring our poor calibration of this phenomenon. Instead, the judicious choice of cuts in Sec. 8.1.2 reduced its effect to only 5%. All other calibrations induced lesser uncertainties on the efficiency.

The final exclusion limit plot is shown in Fig. 8.21. Compared to the only other existing exclusion limit at n_{ZDL} , set by the XMASS collaboration at $g_{a\gamma\gamma} = 4.8 \cdot 10^{-12} \text{ GeV}^{-1}$ [10], NEWS-G sets a limit 6.2 times lower (or just 4.7 times lower if taking the higher $+2\sigma$ limit).

For a model independent plot representing the results of this search, see Fig. 8.22. As an example of interpretation, a signal in the $9 - 14 \text{ keV}$ range, consisting of non-relativistic particles decaying into two photons of same energy, is excluded at 90% C.L. for rates higher than $\sim 2 \text{ evt/day/m}^3$.

⁸For comparison, if we had taken the basic axion-like cuts, we would have had an efficiency of 25.52%, but a total of 66 events in the region of interest, for an exclusion limit of $g_{a\gamma\gamma} = 2.48 \cdot 10^{-12} \text{ GeV}^{-1}$. Despite a relative loss of efficiency of $\sim 33\%$, the additional cuts improved the limit by almost one third.

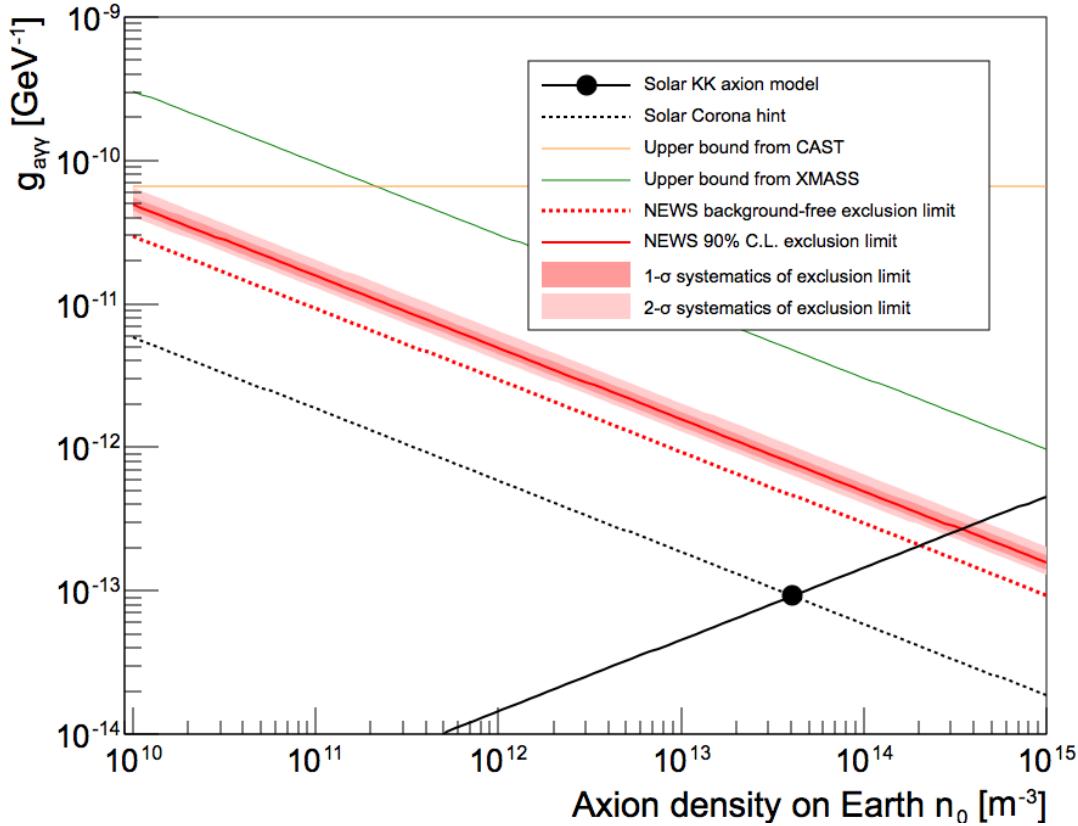


Figure 8.21: Exclusion limit for solar KK axions derived from this work (solid red line), with ranges due to systematics (red shaded areas). For comparison, we show the ideal exclusion limit in the absence of background (dashed red line), the previous limits on solar KK axions from CAST (orange line [26]), and from the XMASS collaboration (green line [10]). The preferred parameter space for the solar KK axion model is shown as the intersection between the solid black line (Solar KK axion model) and the dashed black line (Solar Corona hint) [6].

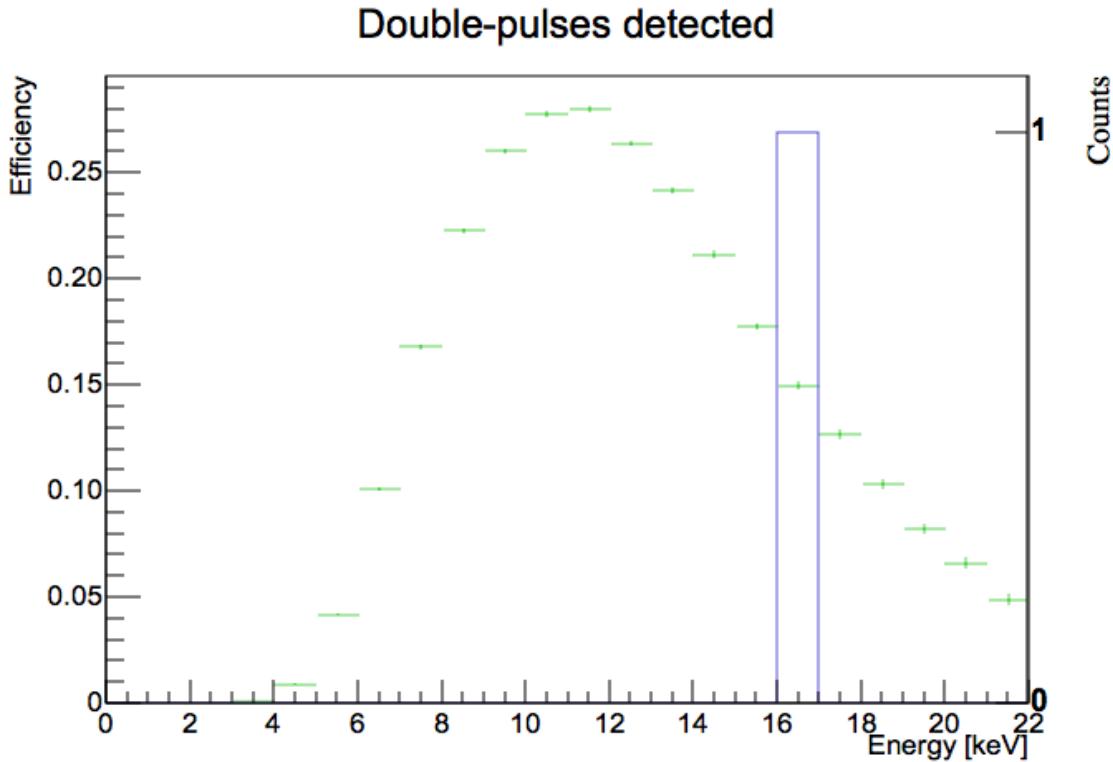


Figure 8.22: Detector efficiency and detected events in the SEDINE detector, for a total exposure of $4.3 \text{ day} \cdot \text{m}^3$. Energies are corrected for attachment. The resolution of the energy reconstruction is $\sim 20\%$ at these energies.

The main conclusions from this study are:

- The large advantage of a gaseous detector in the search for KK axions, in the form of background rejection. Selecting only equal-energy multi-pulse events allows rejection of up to 99.99% of all background in the $2 - 22 \text{ keV}$ energy range (if also rejecting events close to the detector shell). Liquid and solid target detectors cannot make use of this effect.
- World-leading limits are set with a relatively small detector, and less than one and a half months of data. Increasing exposure for improved results is still very possible, both through detector size and run length.

- The main issue with this dataset was radioactive background, mainly in the form of ^{210}Pb deposited in the inner surface, and ^{210}Bi in the detector shell itself. A more radio-pure detector is necessary to probe the preferred parameter space of the solar KK axion model. Even with optimized cuts, the background rate was 4 times the axion rate.
- The lack of proper calibrations for electron drift and diffusion times was the main source of systematic error. Similarly, the large uncertainty on electron attachment prevented the use of more stringent energy or asymmetry cuts, which would improve background rejection even further.

8.3 Projections for SNOGLOBE

On that note, we segue into the future of the collaboration with SNOGLOBE⁹. This detector offers improved radiopurity, increased size, gas purification, and continuous calibrations via laser. These should drastically enhance its sensitivity to solar KK axions. This section will describe the expected performance of this new detector.

8.3.1 Optimal running conditions

The first step is to figure out the optimal running conditions, in terms of detector sensitivity to axions. To limit the scope of the question down to a workable level, the electrode used will be kept as the one tested while SNOGLOBE was at the LSM, during the first leg of its journey to SNOLAB: an achinos with 11 arms, of radius 1.4 cm, with 1 mm electrodes. While our simulation software for 3-dimensional electrodes is not yet mature or tested enough to perform this study, the field can be approximated with a 2-dimensional, cylindrically-symmetric electrode, as shown in App. J.2.

Once that is fixed, the question remains about which gas to use, which pressure, and which voltage. The gas choice is limited to noble gases, to allow regular functioning of the SPC. The photon attenuation properties of the relevant gases are shown in Fig. 8.23, together with methane (our choice of quencher). Since we want photons in the few keVs to have an attenuation length in the 10s of cm, the preferred gas is Neon if going to higher pressures, and Argon otherwise. Helium is too transparent to photons to be of use at those energies, and keeping a high-pressure vessel full of methane underground is impractical.

⁹At time of writing, the detector is waiting at SNOLAB, locked away while we await the passing of the COVID-19 pandemic.

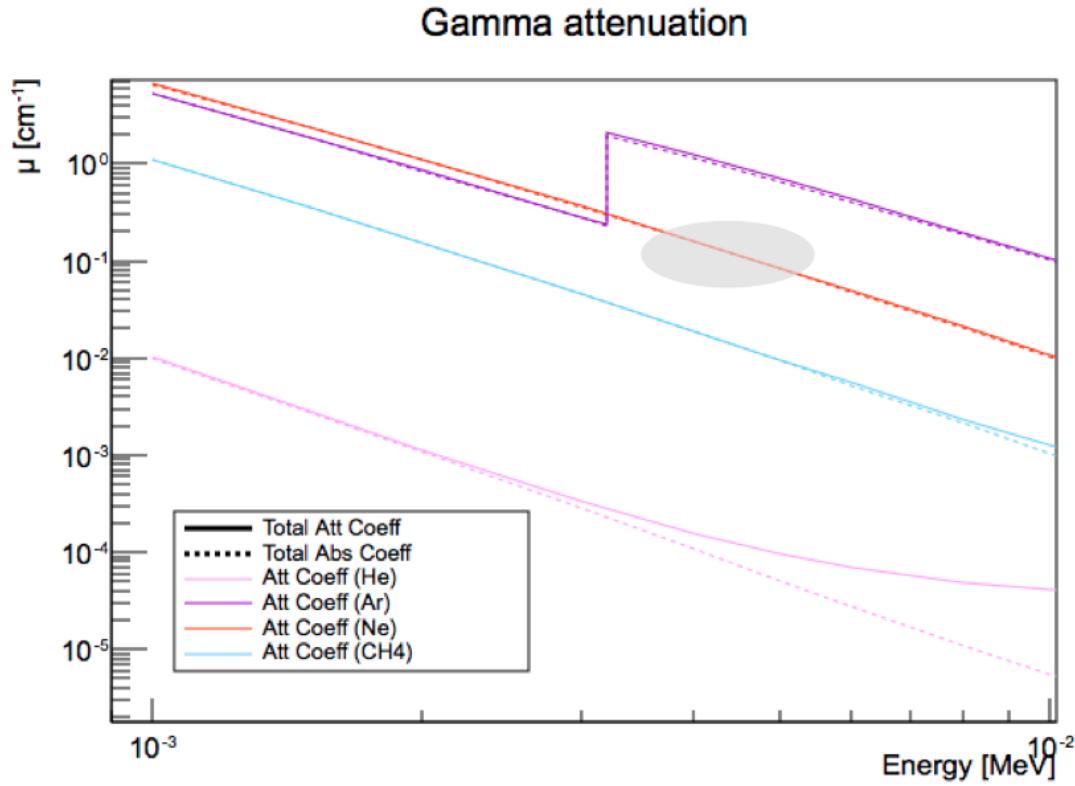


Figure 8.23: Attenuation and absorption coefficients for different gases at 1 bar. The shaded area approximately covers the energies and distances involved in solar KK axion searches with SNOGLOBE. Helium is too transparent, and methane must be diluted to be used in SNOLAB, so the only available gases of interest are Neon and Argon.

This leaves the question of voltage and pressure. To find the optimal conditions, simulations of 10^5 events were performed and processed, for both Neon and Argon, with varying pressures and voltages. A “maximum” efficiency was computed for each condition by finding the proportion of axion events in which both photons are contained inside the gas volume, and the last primary electron from the first photon arrives before the first primary electron from the second photon (i.e., no “pile-up” between both pulses). The results are shown in Fig. 8.24. The optimal efficiency is achieved with 1.0 bar of Neon, and 4000 V, at 29.1%. In practice, operating conditions

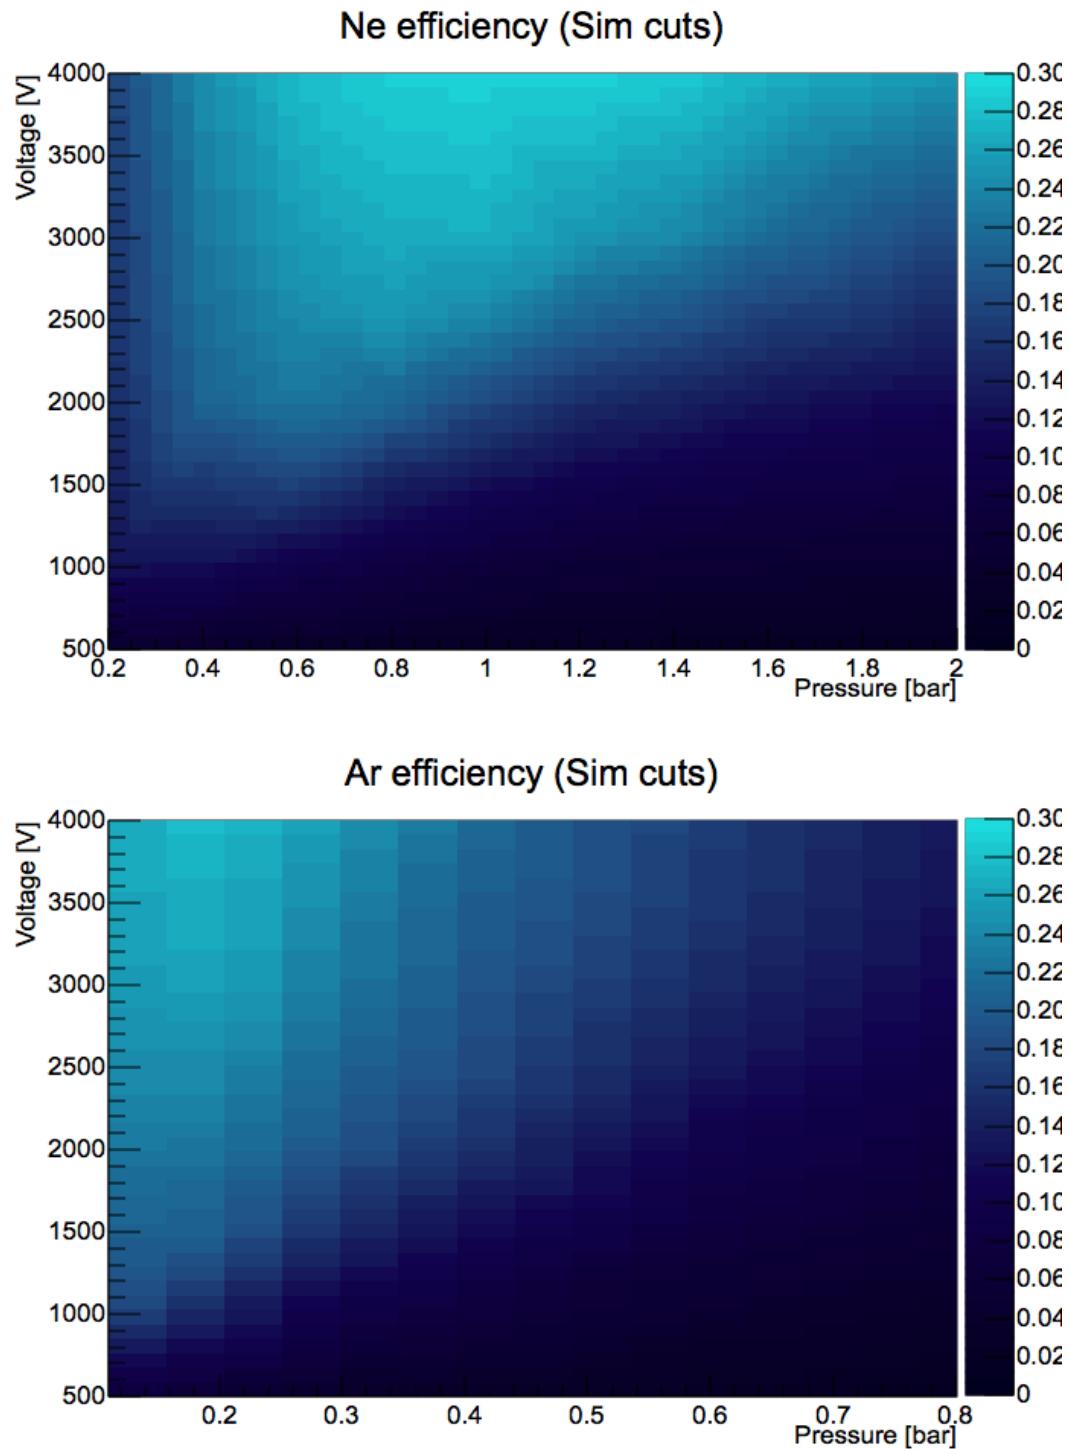


Figure 8.24: Predicted solar KK axion “maximum” efficiency for different voltage and pressure configurations. Top: Neon. Bottom: Argon. Electron drift parameters stay roughly constant when E/P is constant, but photon attenuation length is inversely proportional to P , hence the diagonal feature in both plots.

are limited by sparks forming spontaneously around the electrode at high voltages. If the maximum voltage we can set is only 2000 V, then either Neon at 0.6 bar or Argon at 0.15 bar both reach an efficiency of $\sim 22\%$. At even lower voltages, the performance of Neon keeps going down, while that of Argon remains relatively constant if the pressure is decreased at the same rate as the voltage.

Taking an efficiency of 22% as the reference, together with SNOGLOBE’s internal radius of 67.5 cm, and a combined solar KK axion decay rate of $0.08371 \text{ m}^{-3}\text{day}^{-1}$, we get an event rate of 0.0237 evt/day. It would take 100 days to have enough exposure to exclude the solar KK axion model, assuming it is indeed wrong, and no background at all. A discovery would require higher exposure still, depending on our expected background.

8.3.2 Axion-like backgrounds

For SEDINE, the three main sources of axion-like backgrounds were ^{210}Pb contamination on the inner surface of the copper shell, ^{210}Bi contamination in the bulk of the copper shell, and ^{210}Bi in the lead shield. In terms of improved radiopurity, the largest improvement for SNOGLOBE is the 0.5 mm of copper electroplated on the inner surface of the detector shell [79]. This effectively removes all axion-like events from surface ^{210}Pb . The lead used in the lead shield is also more pure, with a ^{210}Pb contamination of 4.6 Bq/kg (compared to 37.4 Bq/kg for SEDINE). Since the ^{210}Bi contamination in the bulk of the copper shell remains relatively unchanged at 28.5 Bq/kg (compared to 26 Bq/kg for SEDINE), we can assume that this will be the

dominating background for SNOGLOBE¹⁰.

To get a qualitative estimate of the rate of axion-like background events to actual axion events, we use the field approximation for the SNOGLOBE detector with 600 mbar of Neon and an applied voltage of 2000 V, and simulate both axions and $^{210}\text{Pb}/^{210}\text{Bi}$ decays in the copper shell of the detector, taking into account the 0.5 mm of electroplated copper. Due to the gas purification system used in SNOGLOBE, we also assume no gas impurities, and so no attachment.

A Geant4 simulation with $3 \cdot 10^9$ decays for both ^{210}Pb and ^{210}Bi in the copper shell was performed; for a total mass of 521.4 kg of copper (removing the electroplated region) with a contamination of 28.5 Bq/kg of ^{210}Pb , this corresponds to 2336 days of exposure. From all the simulated events, 81982 left energy in the detector, and 36 passed axion-like cuts adapted to the new geometry and gas composition. Of those 36:

- 18 were events leaving less than 6 keV at the detector surface; the low number of primary electrons and long diffusion times produce sharp peaks in the pulse, which are wrongly reconstructed as different pulses. The few events due to ^{210}Pb decays are all in this category.
- 12 were Bremsstrahlung photons from the decay of ^{210}Bi interacting twice in the detector. Their energies ranged from 9.6 to 33.8 keV.
- 6 were two photons interacting in the detector at the same time. Their energies were contained in a narrow range, 14.8 – 18.3 keV. Events with more than one

¹⁰For details on radioactive backgrounds, see the relevant sections in Chap. 3, or A. Brossard’s thesis [35]. In particular, cosmogenic activation of ^{60}Co in the copper shell might actually be the dominating source of background, although its contribution should become subdominant after half a year to one year of “cooling”.

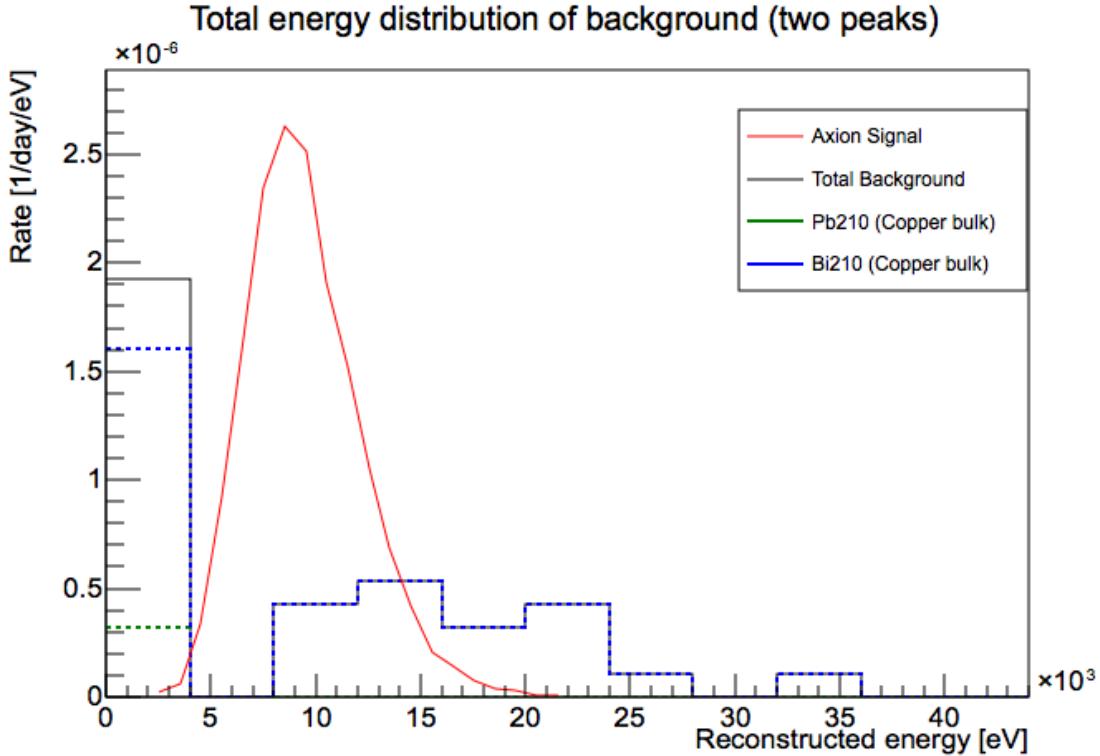


Figure 8.25: Axion (red) and radioactive background (black) rate after cuts for SNOGLOBE with 600 mbar of neon and an applied voltage of 2000 V.

photon generally involve at least one copper fluorescence photon (8.05 keV), hence the narrow range of energies for these events.

Their distribution, compared to the expected rate of axion events after cuts, is shown in Fig. 8.25. Keeping only events in the 5 – 15 keV range, the axion event rate is $16.5 \cdot 10^{-3}$ evt/day. Note that this value is around 33% lower than the one we obtained with the “maximum” efficiency, due to the use of actual processing-like cuts, rather than idealized ones. On the other hand, even making the conservative estimate that there is a uniform background of $3.2 \cdot 10^{-7}$ evt/day/eV in the same energy range, the total background rate is of $3.2 \cdot 10^{-3}$ evt/day, roughly five times lower than the

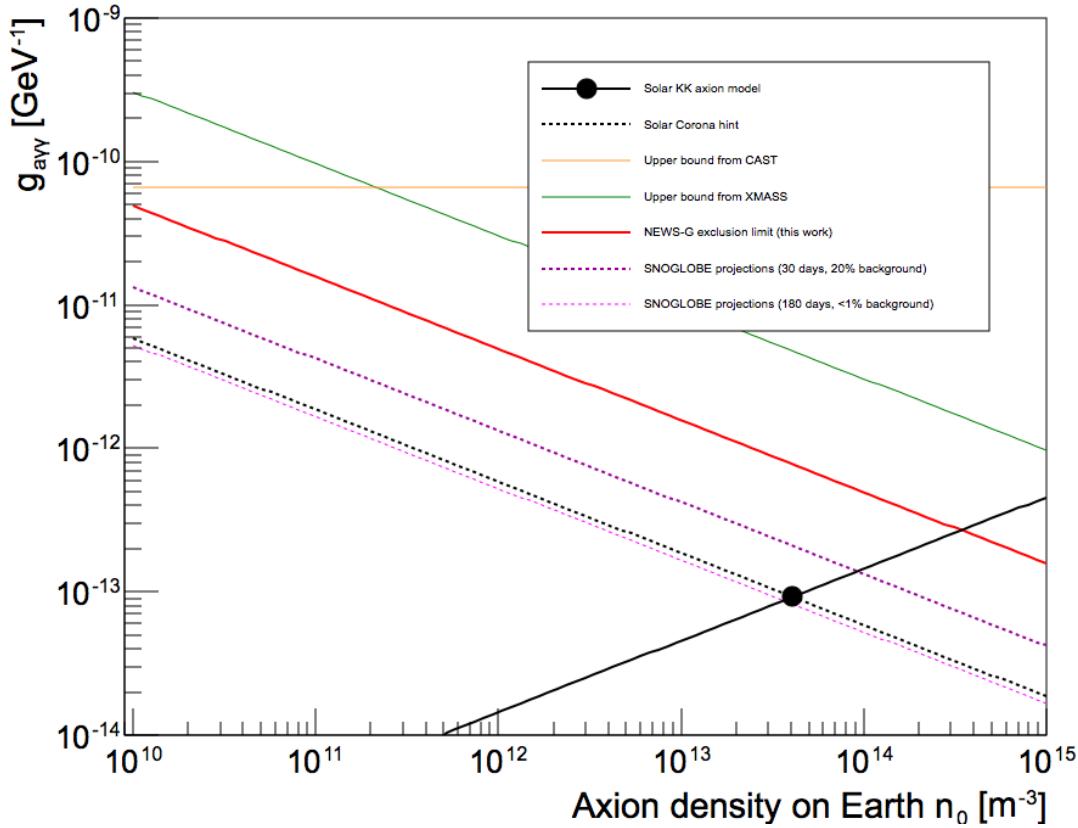


Figure 8.26: Projected limits for SNOGLOBE with 600 mbar of neon and an applied voltage of 2000 V. Both a plausible limit (dashed dark purple line) with a 30 day run and the expected background-to-axion rate of 20%, and an ideal limit (dashed light purple line) with a 180 day run and a background-to-axion rate of < 1%, are shown.

axion rate. This is a considerable improvement over SEDINE, for which the expected background rate after cuts was four times *higher* than the axion rate.

8.3.3 Projections

Based on the expected axion and background event rate, we can draw projections for the SNOGLOBE detector. They are shown in Fig. 8.26. The weaker limit is

set assuming the 20% background-to-axion ratio we obtained for SNOGLOBE, and a moderate run length of 30 days. With a local density of KK axions of $n_{ZDL} = 4.07 \cdot 10^{13} \text{ m}^{-3}$, the projected excluded axion-photon coupling strength is $g_{a\gamma\gamma} < 2.09 \cdot 10^{-13} \text{ GeV}^{-1}$, 3.7 times stronger than the limit set with SEDINE ¹¹.

The stronger limit, that would be able to reject the solar KK axion model at 90% C.L., requires a rather optimistic < 1% background-to-axion ratio, and 180 days of exposure. Given that neither are particularly likely to be achievable for SNOGLOBE in the near future, it is only presented as a benchmark. For comparison, taking the actual expected background rate of 20%, the necessary exposure would be ~ 370 days (for an expected 1.2 background events and 6 axion events).

We can also consider the necessary exposure for a potential discovery. For a discovery at 5σ , we would require a signal that is more than 5 standard deviations away from what could be produced from background alone. For SNOGLOBE, given the relatively small exposures and backgrounds, we can take the uncertainty on the background to be dominated by its statistics. Given a Poissonian background, and an expected number of background events of \bar{n}_{bk} , the standard deviation is $\sigma_{bk} = \sqrt{\bar{n}_{bk}}$. As such, we want enough exposure so that $\bar{n}_a > \bar{n}_{bk} + 5\sigma_{bk}$. Given that we expect a background-to-axion ratio of 20%, this means we want enough exposure so that $\bar{n}_a = 8.0$ and $\bar{n}_{bk} = 1.6$. For an axion rate of $16.5 \cdot 10^{-3} \text{ evt/day}$, this means an exposure of 485 days.

Given the limited time available for SNOGLOBE at SNOLAB, and that the primary mission of the NEWS-G collaboration is the search for WIMPs, it is unlikely

¹¹Remember that for a given local density of axions, their decay rates depend on $g_{a\gamma\gamma}^2$, so the limit on the coupling strength only decreases with the square root of the exposure. This explains why, despite SNOGLOBE having 11 times the volume of SEDINE, the constraint on $g_{a\gamma\gamma}$ did not decrease as much.

that enough detector time will be made available to either reject or discover solar KK axions. However, barring the practical exposure limits of this detector generation, the background study demonstrates that this technology is fully capable of probing the preferred parameter space of the model.

Summary

Based on simulations for our detector response to both axion events and other radioactive backgrounds, optimized cuts were chosen to keep the former while rejecting the latter. Their application to $4.3 \text{ day} \cdot \text{m}^3$ of physics data taken with the SEDINE detector left only one candidate event. Compared to the expected 0.0589 detected axion decays, this allowed us to set a world-leading exclusion limit on solar KK axions: $g_{a\gamma\gamma} < 7.76 \cdot 10^{-13} \text{ GeV}^{-1}$ for a KK axion density on Earth of $n_{ZDL} = 4.07 \cdot 10^{13} \text{ m}^{-3}$, six times stronger than the previous limit set by XMASS. The limited radiopurity of SEDINE, mainly ^{210}Pb on the inner surface of the detector and ^{210}Bi in the copper bulk and lead shield, prevents setting much stronger constraints, even assuming considerably longer exposures.

The upcoming SNOGLOBE detector benefits from significant improvements on its radiopurity compared to SEDINE, leading to a background-to-axion rate of only 20%. This makes it fully capable of probing the parameter space preferred by the Solar Corona problem, $g_{ZDL} = 9.2 \cdot 10^{-14} \text{ GeV}^{-1}$. However, given limited available time for axion searches, we may be limited to setting stronger constraints on the axion-photon coupling. A 30 day run (exposure of $38.6 \text{ day} \cdot \text{m}^3$) is projected to set limits $g_{a\gamma\gamma} < 2.09 \cdot 10^{-13} \text{ GeV}^{-1}$, 3.7 times stronger than with SEDINE.

Chapter 9

Summary and Conclusions

This thesis documents the work done to set limits on the coupling between axion and photons, based on the solar KK axion model developed in [6]. In the parameter space in which KK axions solve the solar corona problem, the model predicts axions decaying into two photons at a rate on Earth of $\sim 0.08 \text{ evt/m}^3/\text{day}$, mainly in the $5 - 15 \text{ keV}$ range. A review of the literature on the QCD axion was performed to derive the experimental and observational constraints on the aforementioned model. XMASS was the only pre-existing limit that targeted it explicitly, while tension with constraints on exotic energy losses from the Sun urge for a more detailed treatment of the model than the one available.

Pulse processing methods were developed and tested, based on our understanding of the detector. In turn, software for simulating pulses was also developed. In SEDINE-like conditions, the MPA pulse processing method achieved an axion event detection efficiency of up to 35%, with a pointlike event rejection of $\sim 99.99\%$. In the same conditions, the simulations showed that two energy depositions in the detector can be resolved independently if their radial distance to the central electrode differs by at least 2 cm. To test the pulse processing on real axion-like data, a run

with ^{55}Fe -induced argon fluorescence was taken with a 130 cm wide SPC at Queen's University. While the ~ 400 Hz of muon tracks severely degraded the quality of the results, we found an agreement between simulations and data. In particular, for the data at 200 mbar, the rate of fluorescence events to total ^{55}Fe events were predicted to be $2.93 \pm 0.29\%$, and found to be $2.99 \pm 0.28\%$.

For the 42 day long neon data taken with the SEDINE detector, the main source of radioactive background was found to be ^{210}Pb deposited on the inner surface of the shell (^{210}Bi de-excites by simultaneous release of photons and electrons), followed by ^{210}Bi in the copper bulk and lead shield (double Compton interaction from a Bremsstrahlung photon). Strengthening the cuts on surface events and pulse asymmetry allowed reaching a background rejection of 99.99%, although the total efficiency to axions was lowered to 16.34%. With a total exposure of $4.3 \text{ day} \cdot \text{m}^3$, and only one candidate event left in the region of interest, NEWS-G can set a world-leading exclusion limit on solar KK axions: $g_{a\gamma\gamma} < 7.76 \cdot 10^{-13} \text{ GeV}^{-1}$ for a KK axion density on Earth of $n_{ZDL} = 4.07 \cdot 10^{13} \text{ m}^{-3}$, six times stronger than the previous limit set by XMASS.

Testing the preferred parameter space of the solar KK axion model is impractical for SEDINE due to its radioactive contamination. However, given the significant improvements in the radiopurity of the upcoming 140 cm wide SPC detector at SNO-LAB (mainly through electroplating of the inner shell and use of cleaner lead in the shield), we predict a ratio of only 20% between background and axion decay rates in the region of interest. This would allow either the discovery or rejection of the model, given a run of ~ 480 or ~ 370 days, respectively, under appropriate operating conditions. It might not be possible to achieve the necessary exposure during its time

at SNOLAB, but even a 30 day run (exposure of $38.6 \text{ day} \cdot \text{m}^3$) is projected to achieve an exclusion limit of $g_{a\gamma\gamma} < 2.09 \cdot 10^{-13} \text{ GeV}^{-1}$, 3.7 times stronger than with SEDINE.

Bibliography

- [1] R. D. R.D. Peccei and H.R. Quinn. CP conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, 38:1440–1443, 1977. 1
- [2] PARTICLE DATA GROUP. The review of particle physics. *Phys. Rev. D*, 98, 2019. 1, 21, 231
- [3] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali. Phenomenology, astrophysics, and cosmology of theories with submillimeter dimensions and tev scale quantum gravity. *Phys. Rev. D*, 59:086004, Mar 1999. 1, 7
- [4] K.R. Dienes, E. Dudas, and T. Gherghetta. Invisible axions and large-radius compactifications. *Phys. Rev. D*, 62:105023, Oct 2000. 7
- [5] S. Chang, S. Tazawa, and M. Yamaguchi. Axion model in extra dimensions with TeV scale gravity. *Phys. Rev. D*, 61:084005, 2000. 1, 7, 9
- [6] L. DiLella and K. Zioutas. Observational evidence for gravitationally trapped massive axion(-like) particles. *Astropart. Phys.*, 19(1):145–170, 2003. xi, xii, xxx, 1, 9, 10, 11, 12, 13, 14, 15, 18, 19, 21, 29, 203, 206, 218
- [7] I. Giomataris *et al.* A novel large-volume spherical detector with proportional amplification read-out. *JINST*, 3(2):P09007, 2008. 1, 33

- [8] L. Lella, A. Pilaftsis, G. Raffelt, and K. Zioutas. Search for solar kaluza-klein axions in theories of low-scale quantum gravity. *Phys. Rev. D*, 62, 2000. x, 4, 12
- [9] B. Morgan *et al.* Searches for solar Kaluza-Klein axions with gas TPCs. *Astropart. Phys.*, 23:28–302, 2005. xi, xix, 4, 14, 115, 116
- [10] XMASS Collaboration. Search for solar Kaluza-Klein axion by annual modulation with the XMASS-I detector. *PTEP*, 2017(10):103C01, 2017. xii, xix, xxx, 12, 27, 29, 116, 205, 206
- [11] J. Lean. The sun’s variable radiation and its relevance for earth. *Ann. Rev. Astron. Astrophys.*, 35:33–67, 1997. xi, xii, 15, 18
- [12] J. Klimchuk. On solving the coronal heating problem. *Solar Physics*, 234, 12 2005. 16, 17
- [13] T. SAKURAI. Heating mechanisms of the solar corona. *Proceedings of the Japan Academy, Series B*, 93:87–97, 02 2017. xi, 16
- [14] C. Parnell and I. De Moortel. A contemporary view of coronal heating. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 370:3217–40, 07 2012. 17
- [15] R. Erdélyi and I. Ballai. Heating of the solar and stellar coronae: A review. *Astronomische Nachrichten*, 328:726 – 733, 10 2007. 17
- [16] I. Moortel and P. Browning. Recent advances in coronal heating. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 373, 05 2015.

- [17] S. Cranmer, M. Asgari-Targhi, M.P. Miralles, J. Raymond, L. Strachan, H. Tian, and L. Woolsey. The role of turbulence in coronal heating and solar wind expansion. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 373, 12 2014. 16
- [18] J.H.M.M. Schmitt *et al.* A soft X-ray image of the Moon. *Nature*, 349:583–587, 1991. xii, 20
- [19] C. O’Hare. Axion limits, 2020. xii, 22
- [20] E. Di Valentino *et al.* Cosmological Axion and neutrino mass constraints from Planck 2015 temperature and polarization data. *Phys. Lett. B*, 752:182–185, 2016. 23
- [21] Davide D. Cadamuro and J. Redondo. Cosmological bounds on pseudo Nambu-Goldstone bosons. *JCAP*, 02:032, 2012. 23
- [22] H. Schlattl, A. Weiss, and G. Raffelt. Helioseismological constraint on solar axion emission. *Astropart. Phys.*, 10:353–359, 1999. 24
- [23] Q.R. Ahmad *et al.* Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002. xii, 24, 28, 29
- [24] A. Ayala, I. Domínguez, M. Giannotti, A. Mirizzi, and O. Straniero. Revisiting the bound on axion-photon coupling from Globular Clusters. *Phys. Rev. Lett.*, 113(19):191302, 2014. 25
- [25] R. Ballou et al. New exclusion limits on scalar and pseudoscalar axionlike particles from light shining through a wall. *Phys. Rev. D*, 92(9):092002, 2015. 25

- [26] V. Anastassopoulos *et al.* New CAST Limit on the Axion-Photon Interaction. *Nature Phys.*, 13:584–590, 2017. xxx, 26, 206
- [27] T. Braine *et al.* Extended Search for the Invisible Axion with the Axion Dark Matter Experiment. *Phys. Rev. Lett.*, 124(10):101303, 2020. 27
- [28] F. Sauli. *Gaseous Radiation Detectors: Fundamentals and Applications*. Cambridge University Press, 2014. 33
- [29] F. Sauli. Principles of operation of multiwire proportional and drift chambers. Technical report, CERN, 1977. 33, 133, 138, 322
- [30] A.D. Fard. *Etude d'un detecteur spherique gazeux pour la recherche d'évenements rares à bas seuil en énergie*. PhD thesis, Faculte des Sciences d'Orsay, 2014. xiii, 33, 42, 45, 46, 47, 48
- [31] Q. Arnaud. First results from NEWS-G at LSM and outlook, 2018. Presented at SuperCDMS Science Meeting. xiii, 34
- [32] I. Katsioulas. Light dark matter searches with the NEWS-G experiment, 2018. 36
- [33] G. Giroux *et al.* The search for light dark matter with the NEWS-G spherical proportional counter. *Journal of Physics: Conference Series*, 1312:012008, 2019.
- [34] A. Giganon *et al.* A multiball read-out for the spherical proportional counter. *JINST*, 12:P12031, 2017. 36, 54
- [35] A. Brossard. *Optimization of spherical proportional counter backgrounds and*

response for low mass dark matter search. PhD thesis, Queen's University, 2020. xxii, 36, 45, 52, 57, 114, 139, 140, 187, 213

[36] W. Shockley. Currents to conductors induced by a moving point charge. *J. Appl. Phys.*, 9:635–636, 1938. 38, 247

[37] G.D. Alkhazov. Statistics of electron avalanches and ultimate resolution of proportional counters. *Nucl. Instrum. Methods*, 89:155–165, 1970. 40

[38] Q. Arnaud *et al.* Precision laser-based measurements of the single electron response of spherical proportional counters for the NEWS-G light dark matter search experiment. *Phys. Rev. D*, 99(10), 2019. xxxii, 40, 136, 138, 266, 267

[39] F. Piquemal. Modane underground laboratory: Status and project. *Eur. Phys. J. Plus*, 127(1):110–114, 2012. 45

[40] G. Gerbier *et al.* News : a new spherical gas detector for very low mass wimp detection. 2014. 45

[41] A. Dastgheibi-Fard *et al.* Background optimization for a new spherical gas detector for very light wimp detection. *PoS*, TIPP2014:375, 2015. 45

[42] Q. Arnaud *et al.* First results from the NEWS-G direct dark matter search experiment at the LSM. *Astropart. Phys. J. C*, 97:54–62, 2018. 45, 108

[43] M. Walker. *Calibration Hardware Research and Development for SNO+*. PhD thesis, 2014. xiii, 49

[44] K.A. Olive *et al.* (Particle Data Group). Review of particle physics. *Chinese Physics C*, 38(9):090001, 1970. 50

- [45] R. Agnese *et al.* Demonstration of surface electron rejection with interleaved germanium detectors for dark matter searches. *Applied Physics Letters*, 103(16):164105, 2013. xiii, xxvii, 53, 188
- [46] E. Armengaud *et al.* Performance of the EDELWEISS-III experiment for direct dark matter searches. *JINST*, 12(08):P08010, 2017. 61
- [47] B.R. Frieden. *Picture processing and digital filtering*. Springer, Berlin, Heidelberg, 1975. 65
- [48] T.J. Kennet and W.V. Prestwich. On the deconvolution of exponential response functions. *Physics in Medicine and Biology*, 24(6):1107–1122, 1979. 68
- [49] Mirion Technologies (CANBERRA) SAS. Model 2006 proportional counter preamplifier. xxxi, 69, 133, 238, 251
- [50] H R Skullderud and P H Larsen. Mobility and diffusion of atomic Helium and Neon ions in their parent gases. *J. of Phys. B*, 23(6):1017–1041, 1990. 70, 145
- [51] M. Morhac, J. Kliman, V. Matousek, M. Veselsky, and I. Turzo. Background elimination methods for multidimensional coincidence gamma-ray spectra. *Nucl. Instrum. Meth. A*, 401:113–132, 1997. 83
- [52] M. Morhac, J. Kliman, V. Matousek, M. Veselsky, and I. Turzo. Efficient one- and two-dimensional gold deconvolution and its application to gamma-ray spectra decomposition. *Nucl. Instrum. Meth. A*, 401(2):385 – 408, 1997.
- [53] M. Morhac, J. Kliman, V. Matousek, M. Veselsky, and I. Turzo. Identification of peaks in multidimensional coincidence gamma-ray spectra. *Nucl. Instrum. Meth. A*, 443:108–125, 03 2000. 83

- [54] M. Morhac. Tspectrum class reference. xvii, 84
- [55] J. Derre, Y. Giomataris, P. Rebourseard, H. Zaccone, J.P. Perroud, and G. Charpak. Fast signals and single electron detection with a MICROMEGAS photodetector. *Nucl. Instrum. Meth. A*, 449:314–321, 2000. 96
- [56] T. Zerguerras, B. Genolini, V. Lepeltier, J. Peyre, J. Pouthas, and P. Rosier. Single electron response and energy resolution of a Micromegas detector. *Nucl. Instrum. Meth. A*, 608:397–402, 2009. 96
- [57] M. Kobayashi, T. Ogawa, T. Kawaguchi, K. Fujii, T. Fusayasu, K. Ikematsu, Y. Kato, S. Kawada, T. Matsuda, R.D. Settles, A. Sugiyama, T. Takahashi, J. Tian, T. Watanabe, and R. Yonamine. A novel technique for the measurement of the avalanche fluctuation of gaseous detectors. *Nucl. Instrum. Meth. A*, 845:236–240, 2017. 96
- [58] S.F. Biagi. Monte Carlo simulation of electron drift and diffusion in counting gases under the influence of electric and magnetic fields. *Nucl. Instr. and Meth. A*, 421:234–240, 1999. 97, 100
- [59] COMSOL, Inc. COMSOL Multiphysics reference manual (version 5.3). 97, 105
- [60] S. Agostinelli *et al.* Geant4—a simulation toolkit. *Nucl. Instr. and Meth.*, 506(3):250–303, 2003. 114
- [61] J. Allison *et al.* Geant4 developments and applications. *IEEE Transactions on Nuclear Science*, 53(1):270–278, 2006.
- [62] J. Allison *et al.* Recent developments in Geant4. *Nucl. Instr. and Meth.*, 835:186–225, 2016. 114

- [63] J.H. Hubbell and S.M. Seltzer. Tables of X-ray mass attenuation coefficients and mass energy-absorption coefficients (version 1.4), 2004. xx, 118, 119, 321
- [64] M.A. Biondi and L.M. Chanin. Blanc's law—ion mobilities in Helium-Neon mixtures. *Phys. Rev.*, 122:843–847, 1961. 144
- [65] J.A. Hornbeck. The drift velocities of molecular and atomic ions in Helium, Neon, and Argon. *Phys. Rev.*, 84:615–620, 1951. 145
- [66] M.-M. Bé, V. Chisté, C. Dulieu, E. Browne, C. Baglin, V. Chechev, N. Kuzmenko, R. Helmer, F. Kondev, D. MacMahon, and K.B. Lee. *Table of Radionuclides*, volume 3 of *Monographie BIPM-5*. Bureau International des Poids et Mesures, 2006. 147
- [67] S.Y.F. Chu, L.P. Ekstrom, and R.B. Firestone. WWW table of radioactive isotopes, 1999. 147
- [68] G.R. Harrison, R.C. Crawford, and J.I. Hopkins. K-series fluorescence yields of argon, copper-65, and indium-113. *Phys. Rev.*, 100(3):841–844, 1955. 147
- [69] C. Godeau. Determination of the K fluorescence yield of argon by proportional counter spectrometry. 91, 1961. 147
- [70] T. Watanabe, H.W. Schnopper, and F.N. Cirillo. K X-ray fluorescence yield of argon. *Phys. Rev.*, 127:2055–2057, 1962. 147
- [71] L.E. Bailey, K.G. Harrison, and J.B. Swedlund. X-ray fluorescence yields of Al, Cl, Ar, Sc, Ti, V, Mn, Fe, Co, Y, and Ag. *Phys. Rev.*, 158:6–13, 1967.

[72] H. Tawara, K.G. Harrison, and F.J. De Heer. X-ray emission cross sections and fluorescence yields for light atoms and molecules by electron impact. *Physica*, 63(2):351–367, 1973. 147

[73] W. Bambynek *et al.* X-ray fluorescence yields, Auger, and Coster-Kronig transition probabilities. *Rev. Mod. Phys.*, 44:716–813, 1972. 147

[74] M.O. Krause. Atomic radiative and radiationless yields for K and L shells. *J. Phys. & Chem. Ref. Data*, 8(2):307–327, 1979.

[75] J.H. Hubbell *et al.* A review, bibliography, and tabulation of K, L, and higher atomic shell X-ray fluorescence yields. *J. Phys. & Chem. Ref. Data*, 23(2):339–364, 1994. 147

[76] P. Shukla and S. Sankrith. Energy and angular distributions of atmospheric muons at the Earth. *International Journal of Modern Physics A*, 33(30), 2018. 148

[77] S. Yellin. Finding an upper limit in the presence of an unknown background. *Phys. Rev. D*, 66:032005, 2002. 202

[78] F. Garwood. Fiducial limits for the poisson distribution. *Biometrika*, 28(3-4):437–442, 1936. 204

[79] S. Alcantar Anguiano *et al.* Background suppression through copper electroplating for the NEWS-G Dark Matter search. 2020. 212

[80] G. Grilli di Cortona, E. Hardy, J. Pardo Vega, and G. Villadoro. The QCD axion, precisely. *JHEP*, 01:034, 2016. 231

- [81] L. Di Luzio, M. Giannotti, E. Nardi, and L. Visinelli. The landscape of QCD axion models. 3 2020. 231
- [82] Inc. Cremat. Cr-110-r2 charge sensitive preamplifier. 238, 251
- [83] J. Derre. Pulse shape in the SPC prototype. Technical report, 2007. 246
- [84] CERN. Lep radio frequency copper cavity. xxxii, 263
- [85] F. Kelly *et al.* The production of ar-37 using a thermal neutron reactor flux. *Journal of Radioanalytical and Nuclear Chemistry*, 318, 08 2018. 264
- [86] TUNL. Tunl facilities. 266
- [87] RMTL. Reactor materials testing laboratory. 266
- [88] V.T. Jordanov, G.F. Knoll, A.C. Huber, and J.A. Pantazis. Digital techniques for real-time pulse shaping in radiation measurements. *Nucl. Instr. and Meth. A*, 353(1):261 – 264, 1994. 270
- [89] M. Frigo and S.G. Johnson. FFTW. 285
- [90] M. Frigo and S.G. Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231, 2005. 285
- [91] P.D. Higgins, F.H. Attix, J.H. Hubbell, S.M. Seltzer, and M.J. Berger. Mass energy-transfer and mass energy-absorption coefficients, including in-flight positron annihilation for photon energies 1 kev to 100 MeV. Technical report, NIST (NML), 1991. 321

Appendix A

The QCD Axion

This section was based on [2], [80] and [81], in increasing order of theoretical detail. This is only meant as a primer to introduce some of the vocabulary, concepts and behaviour that are relevant for understanding the KK axion model.

A.1 Strong CP problem

The Standard Model QCD Lagrangian includes a CP violating term, the so-called topological term:

$$\mathcal{L}_\Theta = -\overline{\Theta}(\alpha_s/8\pi) G^{\mu\nu a} G_{\mu\nu}^a \quad (\text{A.1})$$

where $-\pi < \overline{\Theta} < \pi$ is the effective Θ parameter after diagonalizing quark masses, $G^{\mu\nu a}$ is the color field strength tensor, and $G_{\mu\nu}^a$ its dual.

The neutron electric dipole moment (nEDM) is an observable consequence of CP violation induced by $\overline{\Theta}$. We get the following relationship between both:

$$d_n \sim 1.2 \cdot 10^{-2} \overline{\theta} e \text{ GeV}^{-1}$$

note that for the topological term of the QCD Lagrangian to contribute to nEDM, this relationship requires the absence of any massless quark.

Given experimental constraints on nEDM of $d_n < 1.5 \cdot 10^{-12} e \text{ GeV}^{-1}$, we get extremely small values of $\bar{\theta} < 10^{-10}$, even though we were expecting $\theta = O(1)$. The smallness of θ is the so-called Strong CP problem.

A.2 Peccei-Quin mechanism and Axion solution

To solve this issue, we introduce a new boson, the axion, with Lagrangian:

$$\mathcal{L}_a = \frac{1}{2}(\partial_\mu a)^2 + \frac{a}{f_a} \frac{\alpha_s^2}{8\pi} G\tilde{G} + \frac{1}{4}a g_{a\gamma\gamma}^0 F\tilde{F} + \frac{\partial_\mu a}{2f_a} j_{a,0}^\mu \quad (\text{A.2})$$

where f_a is a parameter characterizing the energy scale of the axion, and axionic pseudo-shift symmetry was used to absorb the θ QCD topological term. If the potential of the axion is such that it set the vacuum expectation value of the axion to zero (i.e. $\langle a \rangle = 0$), then the CP-violating term $G\tilde{G}$ disappears from the Lagrangian, solving the Strong CP problem.

Applying a field-dependent axial transformation of the quark fields allows for the elimination of the linear coupling to QCD, at the cost of generating other axion-dependent terms. In particular, choosing the change of field variables on the up and down quarks:

$$q = \begin{pmatrix} u \\ d \end{pmatrix} \rightarrow e^{i\gamma_5 \frac{a}{2f_a} Q_a} \begin{pmatrix} u \\ d \end{pmatrix} \quad (\text{A.3})$$

produces the new form of the Lagrangian:

$$\mathcal{L}_a = \frac{1}{2}(\partial_\mu a)^2 + \frac{1}{4}a g_{a\gamma\gamma} F\tilde{F} + \frac{\partial_\mu a}{2f_a} j_a^\mu - \bar{q}_L M_a q_R + \text{h.c.} \quad (\text{A.4})$$

where

$$g_{a\gamma\gamma} = \frac{\alpha}{2\pi f_a} \left[\frac{E}{N} - 6\text{tr}(Q_a Q^2) \right] \quad (\text{A.5})$$

$$j_a^\mu = j_{a,0}^\mu - \bar{q} \gamma^\mu \gamma_5 Q_a q \quad (\text{A.6})$$

$$M_a = e^{i\frac{a}{2f_a} Q_a} M_q e^{-i\frac{a}{2f_a} Q_a} \quad (\text{A.7})$$

and

$$M_q = \begin{pmatrix} u & 0 \\ 0 & d \end{pmatrix} \quad (\text{A.8})$$

$$Q = \begin{pmatrix} 2/3 & 0 \\ 0 & -1/3 \end{pmatrix} \quad (\text{A.9})$$

Taking the leading order in the chiral expansion of the Lagrangian, all the the non-derivative terms for the axion are contained in the pion mass terms. Choosing Q_a proportional to the identity, and expanding for $a f_a \ll 1$, provides the following axion-pion potential:

$$V(a, \pi^0) = -m_\pi^2 f_\pi^2 \sqrt{1 - \frac{4m_u m_d}{(m_u + m_d)^2} \sin^2\left(\frac{a}{2f_a}\right)} \cos\left(\frac{\pi^0}{f_\pi} - \phi_a\right) \quad (\text{A.10})$$

which is minimized at $\langle a \rangle = 0$, solving the Strong CP problem. Furthermore, expanding to quadratic order gives the formula for the axion mass:

$$m_a^2 = \frac{m_u m_d}{(m_u + m_d)^2} \frac{m_\pi^2 f_\pi^2}{f_a^2} \quad (\text{A.11})$$

$$m_a \simeq 5.7 \left(\frac{10^9 \text{GeV}}{f_a} \right) \text{meV} \quad (\text{A.12})$$

A.3 Model dependency

The previous section remained general, covering a wide range of models that could solve the strong CP problem. The PQ symmetry needs to be precisely defined, giving rise to model-dependent interactions.

In particular, the PQ current is conserved up to anomalies¹:

$$\partial^\mu J_\mu^{PQ} = \frac{g_s^2 N}{16\pi^2} G \tilde{G} + \frac{e^2 E}{16\pi^2} F \tilde{F} \quad (\text{A.13})$$

where N and E are the model-dependent QCD and electromagnetic anomaly coefficients, respectively. The Goldstone theorem applied to the axion gives an effective Lagrangian containing the terms:

$$\mathcal{L}_a \supset \frac{a}{\nu_a} \frac{g_s^2 N}{16\pi^2} G \tilde{G} + \frac{a}{\nu_a} \frac{e^2 E}{16\pi^2} F \tilde{F} + \frac{\partial_\mu a}{\nu_a} J_\mu^{PQ} \quad (\text{A.14})$$

where ν_a is the order parameter for the symmetry breaking of the PQ symmetry, and J_μ^{PQ} depends on the "global charges of the fields transforming under $U(1)_{PQ}$ ".

¹An "Anomaly" is the breaking of a classical symmetry (or equivalently conservation law) in the quantum theory introduced by renormalization.

Normalizing this expression in terms of the decay constant $f_a = \frac{\nu_a}{2N}$ gives:

$$\mathcal{L}_a \supset \frac{a}{f_a} \frac{g_s^2}{32\pi^2} G\tilde{G} + \frac{1}{4} g_{a\gamma\gamma}^0 a F\tilde{F} + \frac{\partial_\mu a}{2f_a} \bar{f} c_f^0 \gamma_\mu \gamma_5 f \quad (\text{A.15})$$

where f is the field for a given fermion. The model-dependent coupling of the axion to photons and fermions are thus, respectively:

$$g_{a\gamma\gamma}^0 = \frac{\alpha}{2\pi f_a} \frac{E}{N} \quad (\text{A.16})$$

and

$$c_f^0 = \frac{\chi_{H_f}}{N} \quad (\text{A.17})$$

where χ_{H_f} is the PQ charge of the Higgs between the left and right-handed fermion.

A.3.1 KSVZ vs DFSZ

The main differences between models are the number of anomalies, and the PQ charges of the fields transforming under $U(1)_{PQ}$. While a wide range of models exist, the main two used as benchmarks are the Kim-Shifman-Vainshtein-Zakharov (KSVZ) and the Dine-Fischler-Srednicki-Zhitnitsky (DFSZ)axions.

The KSVZ or Hadronic axion extends the Standard Model with new heavy quark with a PQ charge, leaving all standard fermions untransformed under $U(1)_{PQ}$. If the new quarks have neutral electric charge, then $E/N = 0$, and so the axion does not couple to electrons (nor any other standard fermions) at tree-level.

The DFSZ or Grand Unified Models includes at least two Higgs doublets H_u and H_d , but unlike KSVZ, ordinary quarks and leptons do carry a PQ charge. In these

models, $E/N = 8/3$ (DFSZ-I) or $E/N = 2/3$ (DFSZ-II), with the difference between the sub-models being the type of coupling between the Higgs doublets and leptons.

A.4 Axion coupling to photons (and rest of SM?)

From previous sections, the coupling strength of axions to photons is:

$$g_{a\gamma\gamma} = \frac{\alpha}{2\pi f_a} \frac{E}{N} - \frac{\alpha}{2\pi f_a} \left(\frac{2}{3} \frac{4m_d + m_u}{m_u + m_d} \right) \quad (\text{A.18})$$

where the first term is model dependent, and the second term arises from the cancellation of the axion-gluon coupling in the effective axion Lagrangian, independently of the model. Numerically, we obtain:

$$g_{a\gamma\gamma} \simeq (0.203E/N - 0.39) \frac{m_a}{\text{GeV}^2}$$

From the effective Lagrangian of the axion, its two photon decay-width (or decay time) is:

$$\Gamma_{a \rightarrow \gamma\gamma} = \frac{g_{a\gamma\gamma}^2 m_a^3}{64\pi} \quad (\text{A.19})$$

with axions decaying faster than age of universe if $m_a \gtrsim 20 \text{ eV}$

Appendix B

Electronics calibrations

B.1 Electronic components

B.1.1 Proportional Charge Counter

Proportional Charge counters, also known as Charge Sensitive Preamplifiers (referred to as “preamp” from now on), convert an ionization charge into a step voltage signal, whose amplitude is proportional to the charge collected. A functional schematic of one such preamp is shown on Fig. B.1.

The “integrator” part of the diagram is responsible for the response function of the preamp, a decaying exponential with time constant $\tau = RC$, where R and C are the feedback resistance and capacitance of the preamp, respectively¹. This response function behaves like an integrator for signals that are short compared to its time constant, while still returning to the baseline after some time, limiting pileup.

The combination of resistors and capacitances at the HV and Detector inputs

¹The response of the preamp is not perfectly instant, with a risetime of the order of $\sim 10\text{ ns}$ (increasing with the input capacitance). Given our usual sampling frequencies are no shorter than 480 ns , we can effectively consider it to be instant, but this may be different for faster digitizers.

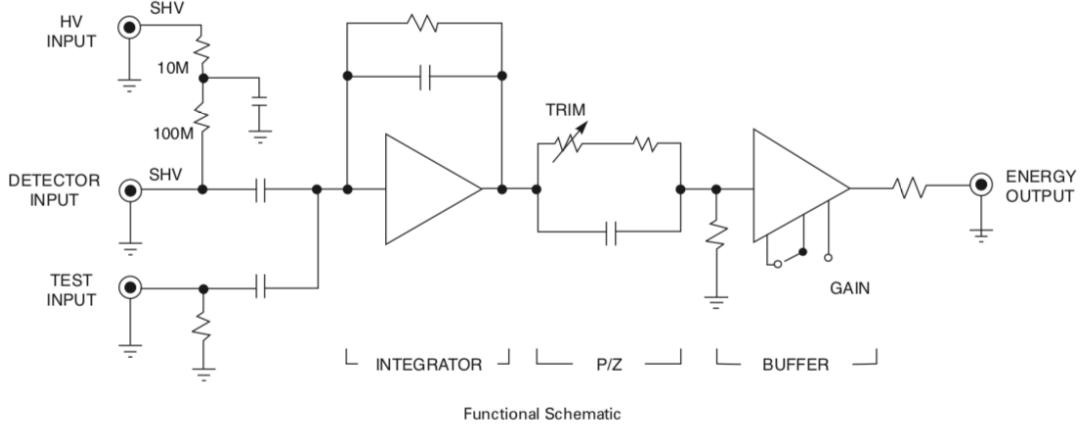


Figure B.1: Functional diagram of a proportional counter, as provided by the Canberra 2006 documentation [49].

serve a dual purpose. At the HV input, it behaves like a low-pass filter, eliminating electronic noise from the bias voltage. At the Detector input, the bias resistor presents a high impedance to the detector signal, which instead goes through the low-impedance capacitance (protecting the amplifier from the high voltage) into the integrator.

For our purposes, a preamp is characterized by two parameters: its decaytime τ (50 μ s for Canberra 2006 [49], 140 μ s for Cremat Z-110 [82]) and its amplification gain G_{preamp} (235 mV/Me⁻ for Canberra 2006, 1.4 V/pC for Cremat Z-110, both roughly equivalent). In practice, its *effective* gain will also be affected by their output resistance (93 Ω for Canberra 2006, 50 Ω for Cremat Z-110).

B.1.2 Digitizer

The role of the digitizer is to convert the analog voltage signal coming from the preamp into a digital (i.e. computer-readable) format. While multiple models exist,

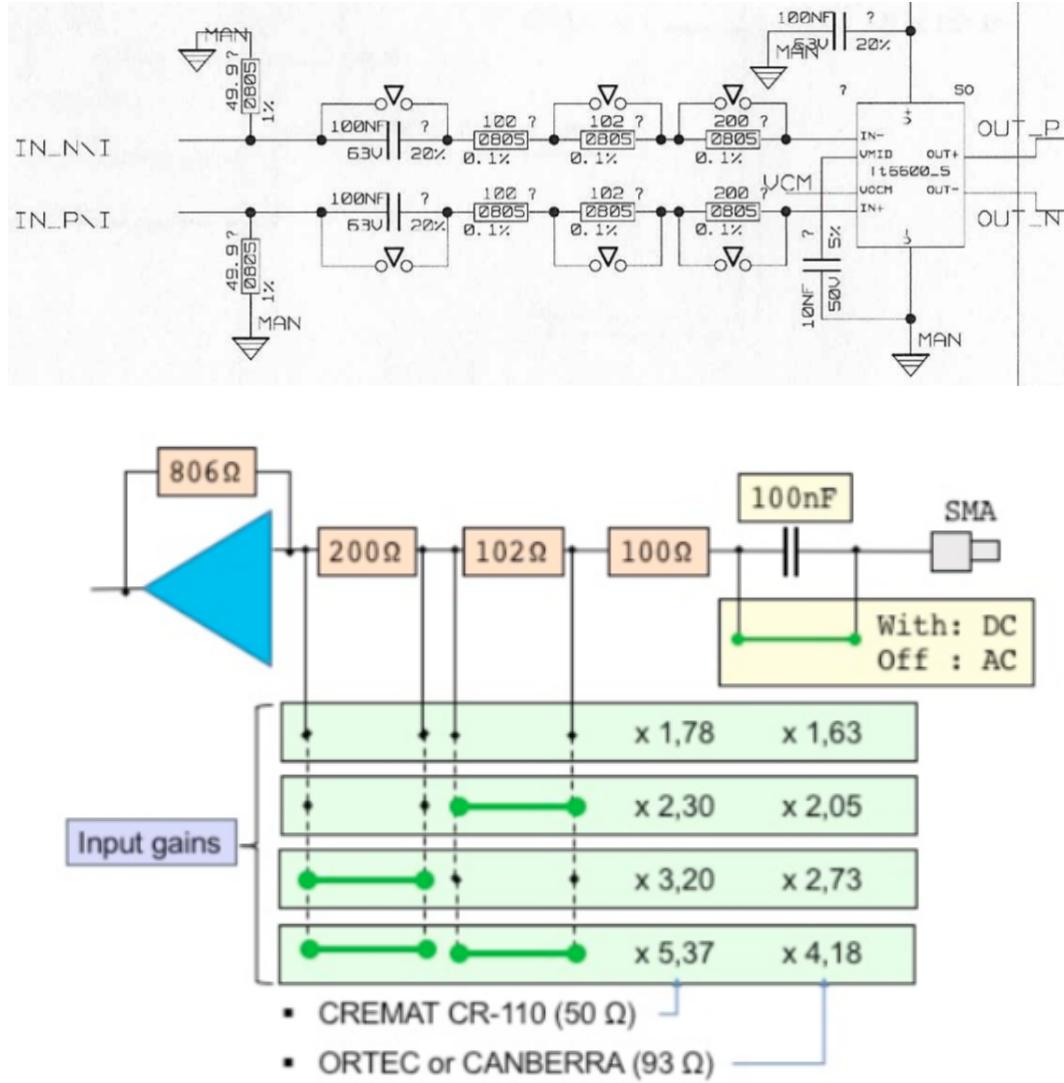


Figure B.2: Electronic diagram for the Calibox input. Top: Detail of the full diagram for one +/- pair of inputs. Bottom: Simplified diagram for a single input (although missing the $49.9\ \Omega$ input resistance connected to the ground). Please note the left/right inversion between both diagrams.

I will focus on the Calibox, the in-house acquisition system developed primarily by M. Gros. In fact, we will only be discussing the part that is relevant to the electronic calibrations we want to perform: the input electronics.

The Calibox records up to 4 sensors of 16 bits each (i.e., values between -32768 and 32767 ADU each). To lose a minimum of information to digitization, we want 1 ADU to correspond to as small a voltage difference as possible; however, to avoid losing information to pulse saturation (values larger than 32767 ADU or smaller than -32768 ADU), we want 1 ADU to represent as large a voltage difference as possible. Due to these two conflicting constraints, the digitizer gain ($G_{ADU/V}$) should be tuned for the range of voltages we are interested in.

The requirement to have a flexible digitizer gain lead to a system of “jumpers” in the input electronics, as shown in Fig. B.2. The effective gain of the digitizer is theoretically proportional to the ratio between the feedback resistance of the input amplifier (the LT66000-5) and the total input resistance. Using the simplified diagram, the “Gain 1” setup has no jumpers installed, so the gain is proportional to $\frac{806}{100+102+200}$; conversely, the “Gain 4” setup has jumpers on both available resistances, so the gain is proportional to $\frac{806}{100}$.

In reality, a number of factors change this calculation:

- The presence of an input resistance in the Calibox coupled directly to the ground (not shown in the simplified diagram, but visible in the full one);
- The output resistance of the preamp connected to the digitizer changes the total input resistance, so the digitizer gain depends on the preamp used (e.g. CR-Z-100 has $R_{O_{preamp}} = 50 \Omega$, Canberra 2006 has $R_{O_{preamp}} = 93 \Omega$).
- The input amplifier used in the Calibox is not the same for all boxes, with some having an LT6600-15 instead (with a feedback resistance of 536Ω , instead of 806Ω).

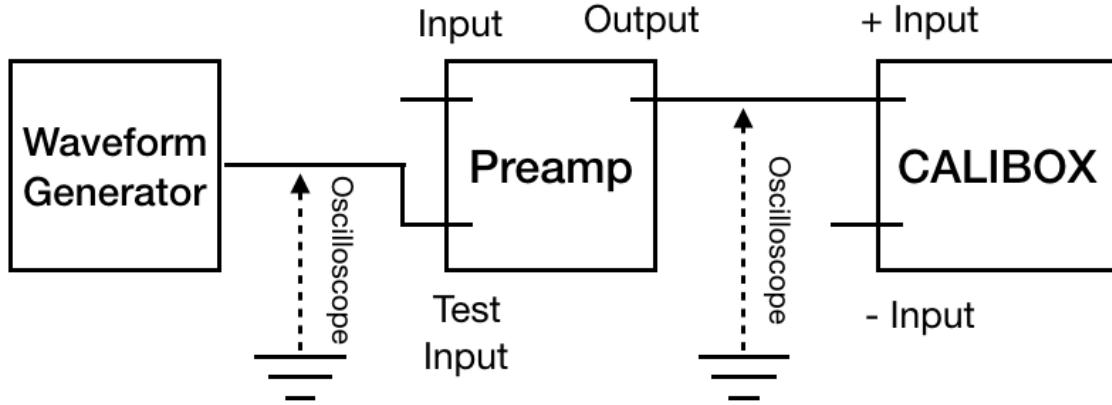


Figure B.3: Electronic setup to calibrate $K_{e^-/ADU}$.

- The input amplifier is a differential amplifier² (with a low-pass filter), so the behaviour of the + input depends on the setup of its coupled – input, and viceversa; the same jumpers should be set on both inputs at once for regular behaviour.

What these considerations mean in practice is that the gain of any electronic setup using a Calibox cannot be easily predicted. Any new setup should be calibrated, independently of any other previous calibrations.

B.2 Calibration setup

To find the average avalanche gain, or to estimate the “space charge current” in the detector, we need to find the conversion factor between the digitizer unit and the number of secondary charges integrated by the preamp, $K_{e^-/ADU}$.

²The use of a differential amplifier could in theory be exploited to increase the dynamic range of the Calibox by choosing the right vias voltage to apply on the coupled input, or even to decrease electronic noise. To the best of my knowledge, this has not been done for any data I have interacted with.

This calibration requires both the preamp and the digitizer, and a waveform generator (WFG), as shown on Fig. B.3. The WFG is connected into the test input of the preamp (leaving the regular input and voltage bias connectors unused), with the output of the preamp connected to whichever digitizer is to be tested.

The WFG is used to generate a square wave with, for example, an amplitude peak-to-peak of 500 mV (choosing the voltage so that the digitizer does not saturate) and a frequency of 100 Hz (choosing the frequency so that there is no pulse pileup). The reason for a square wave is that the voltage shift between “valleys” and “crests” in the signal will be converted into charge pulses by the capacitance in the test input of the preamp. The total charge of each such pulse is $Q = C_{Test}V_{Test}$, where C_{Test} is the test capacitance and V_{Test} is the amplitude peak-to-peak of the square wave as it enters the preamp. This value will likely be different from the amplitude peak-to-peak of the WFG signal, due to the output resistance of the WFG and the input resistance of the preamp behaving like a voltage divider.

The signal on the digitizer should be pulses with a decaying exponential, alternating between positive and negative. The amplitude of one such pulse in digitizer units provides the conversion factor through the following formula:

$$K_{e^-/ADU} = V_{WFG} \frac{R_{I_{preamp}}}{R_{I_{preamp}} + R_{O_{WFG}}} C_{Test}/A_{ADU} \quad (B.1)$$

where V_{WFG} is the amplitude of the square wave created by the waveform generator, C_{Test} is the test capacitance of the preamp, $R_{O_{WFG}}$ and $R_{I_{preamp}}$ are the output resistance of the WFG and input resistance of the preamp, and A_{ADU} is the amplitude of the pulse on the digitizer in ADU.

Notably, the conversion factor depends on both the preamp and the digitizer. The

values found for a few setups were:

- *Canberra + CALI 16 (Gain 1)*: $306 \text{ e}^-/\text{ADU}$
- *Canberra + CALI 16 (Gain 4)*: $98 \text{ e}^-/\text{ADU}$
- *CR-Z-110 + CALI 16 (Gain 1)*: $170 \text{ e}^-/\text{ADU}$
- *CR-Z-110 + CALI 16 (Gain 4)*: $54 \text{ e}^-/\text{ADU}$
- “*METALBOX*” + *CALI 16 (Gain 1)*: $194 \text{ e}^-/\text{ADU}$
- “*METALBOX*” + *CALI 16 (Gain 4)*: $60 \text{ e}^-/\text{ADU}$

One important caveat is that the tolerance of the capacitance in the test input of the preamps can be quite large. For the in-house “*METALBOX*” preamp, the test capacitance is $1 \text{ pF} \pm 10\%$. For the Cremat Z-110, it is $1 \text{ pF} \pm 25\%$. The tolerance of the test capacitance in the Canberra 2006 was not specified. Since the vendors tune the gain of their preamps based on the used test capacitance, this leaves an inescapable uncertainty on the true value of $K_{e^-/\text{ADU}}$, up to 25% for the CR-Z-110!

B.2.1 Oscilloscope measurements

The oscilloscope measurements are not actually necessary for this calibration, but they can be useful to verify its correctness. To minimize the effect of the oscilloscope on the calibration setup, a high input resistance should be used for the oscilloscope (typically, $1 - 10 \text{ M}\Omega$); some minor “ringing” might still be unavoidably introduced. The oscilloscope measurement right after the WFG measures a square wave of amplitude $V'_{WFG} = V_{WFG} \frac{R_{I_{\text{preamp}}}}{R_{I_{\text{preamp}}} + R_{O_{WFG}}}$, due to the voltage division between the output

resistance of the WFG and the input resistance of the preamp. If either of those resistances are unknown, this value can just be used directly in Eq. B.1.

The oscilloscope measurement before the Calibox would theoretically measure a decaying exponential of amplitude $V'_{preamp} = V_{preamp} \frac{R_{I_{CALI}}}{R_{I_{CALI}} + R_{O_{preamp}}}$. However, as mentioned in the previous section, the input electronics of the Calibox are complex, with no equivalent input resistance (however, there might be for other digitizers). Instead, the Calibox can be unplugged, then the oscilloscope measures V_{preamp} directly (the amplitude of the decaying exponential). The gain of the preamp should satisfy:

$$G_{preamp} = V_{preamp} / (C_{Test} V'_{WFG})$$

where G_{preamp} is provided by the preamp supplier.

Keeping the Calibox plugged in, the second oscilloscope measurement can also be used to verify the response function of the preamplifier. Setting the oscilloscope trigger properly to filter the negative pulses, a few thousand pulses can be averaged out together to get the response function with minimal noise. This pulse shape can then be fit with a decaying exponential to verify the decay constant provided by the vendor (e.g., the Canberra 2006 states it has a decaytime of $50\ \mu\text{s}$, but the unit tested had a decay time of only $47\ \mu\text{s}$). In the case of old Cr-Z-110 chips, this test showed an altered response function when under low resistance loads (cf. Fig. B.4). Exchanging with the vendor revealed it to be an issue with the internal power supply, which was solved in the new revised model.

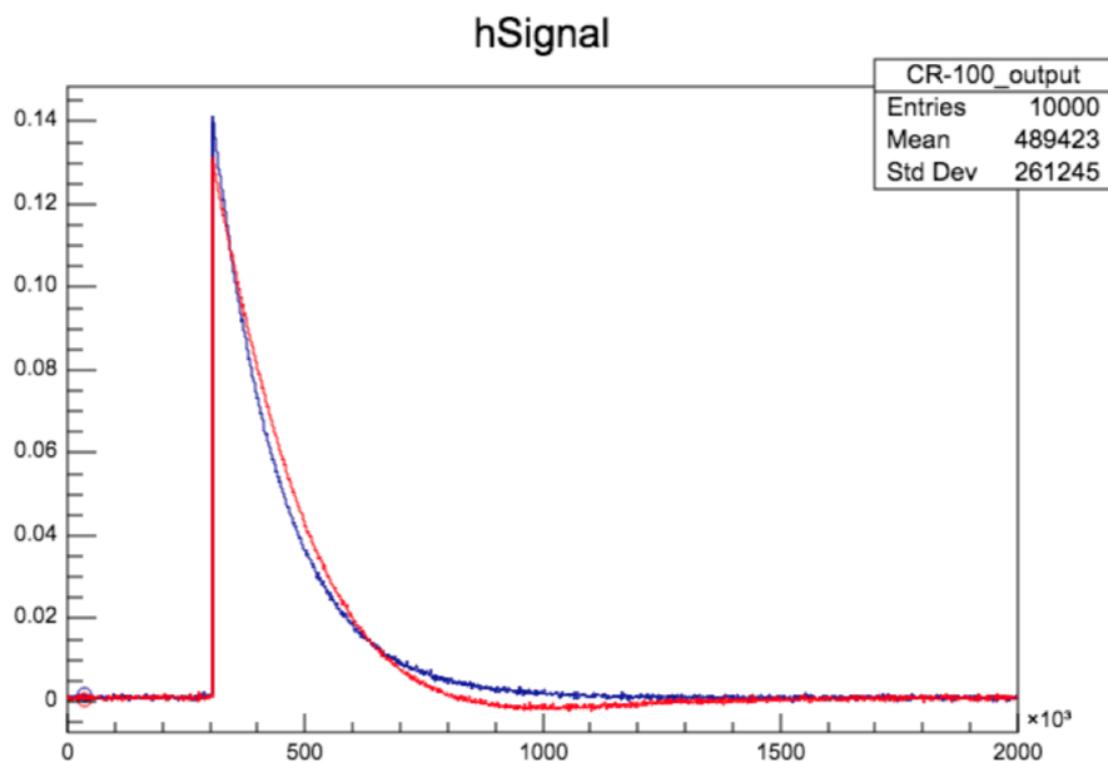


Figure B.4: Response function of the CR-Z-110 under low resistance load ($50\ \Omega$). Red: old model (Rev2). Blue: new model (Rev2.1). The response of the new model is a decaying exponential with a time constant $\tau = 141\ \mu\text{s}$, while the old model had a distinct undershoot behaviour.

Appendix C

Theory of signal formation in an SPC

During the avalanche, a large number of electron-ion pairs are created very close to the central electrode of an SPC. As these move away from their starting point, the charge they induce on the electrode changes, which gets integrated by the preamp, and a pulse forms on the digitizer. The bulk of that pulse is generated by the ions, as they drift through the large potential difference between the central electrode and the shell of the detector.

In the simple approximation of an ideal spherically symmetric detector, the electric field and signal induced by the drifting ions can be worked out analytically, as was shown in an unpublished note by J. Derre [83]. We will cover the demonstration in this section. We will also go over two aspects that were not described in that note: the signal induced by drifting electrons, and the effect of an ideal, continuous space current due to the drifting ions on the electric field.

C.0.1 Ideal electric field

We will consider the ideal case scenario of a big metallic sphere of radius r_1 , kept grounded, with a small electrode of radius r_2 and kept at a voltage V_0 . While not the

real scenario, due to the presence of the rod, we find that the results from such an approach are close enough to the real output we get from our detector. We can use Gauss law and the spherical symmetry of the system to determine:

$$E(r) = \frac{V_0}{r^2} \rho \quad (\text{C.1})$$

$$C = 4\pi\epsilon\rho \quad (\text{C.2})$$

where $E(r)$ is the magnitude of the electric field at a distance r from the centre of the sphere, C is the capacitance of the detector, and $1/\rho = 1/r_2 - 1/r_1$. While this doesn't hold exactly in the real detector due to the presence of the rod, the correction of the field introduced by the umbrella makes it a good enough approximation for most of our volume.

C.0.2 Signal induced by secondary ions

Now that we have the electric field, we will determine the current induced on the electrode by the ions drifting away from it. Using the Shockley-Ramo theorem [36], we know that the change of the charge induced on our electrode by a moving ion is

$$dQ_{ind} = -q_{ion} \frac{E(r)}{V_0} v_{ions} dt \quad (\text{C.3})$$

We can figure out the velocity of the drifting ions via the ion mobility in the target gas, μ . We take

$$\mu = \frac{\mu_0}{P} \quad (\text{C.4})$$

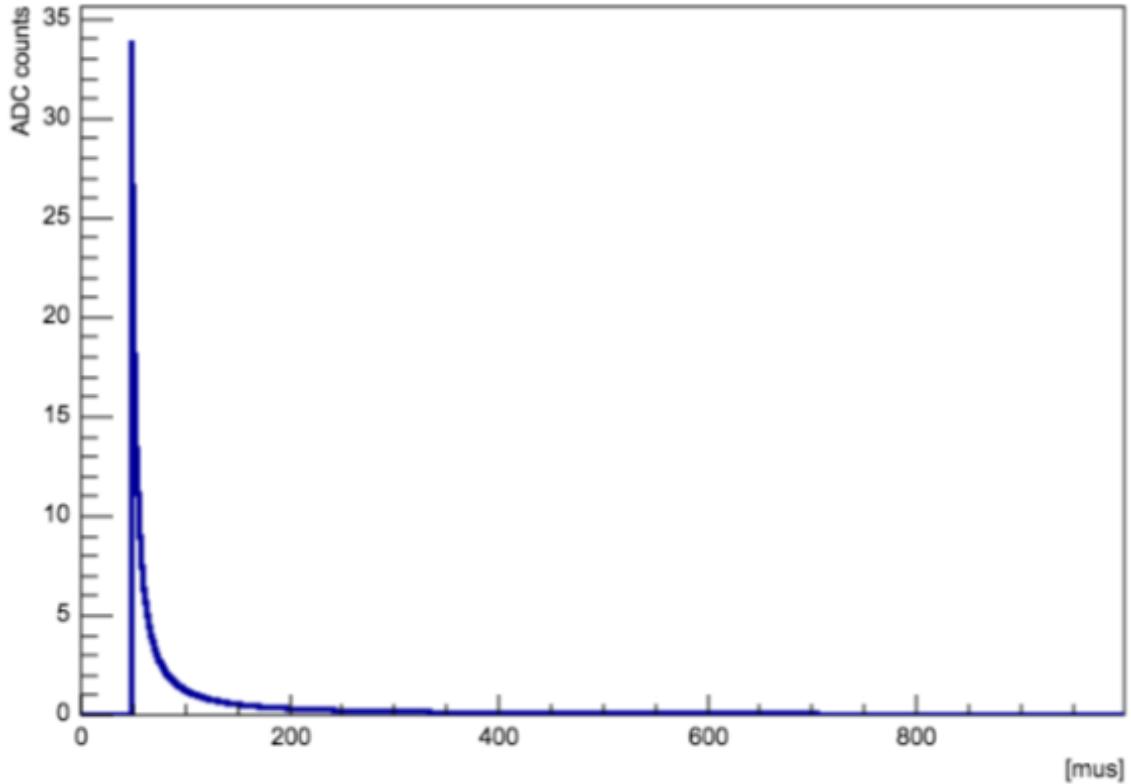


Figure C.1: Theoretical current induced on the electrode by an avalanche that creates 1000 electron-ion pairs at time $48 \mu\text{s}$.

where μ_0 is the ion mobility in the gas at normal temperature and pressure, and P is the pressure, and then, by definition,

$$\mu E(r) = v_{ion} \quad (\text{C.5})$$

Note that μ is also directly proportional to the absolute temperature, but we will ignore this effect, as temperature remains mostly constant in our experiment; μ also depends on the strength of the electric field in complicated ways depending on the exact mixture, but we will ignore this effect for these derivations. Some values of μ_0 are $\sim 2 \text{ cm}^2 \text{V}^{-1} \text{s}^{-1}$ for Argon and $\sim 6 \text{ cm}^2 \text{V}^{-1} \text{s}^{-1}$ for Neon.

We can then use Eq. C.1 in Eq. C.5 to get

$$r^2 dr = \alpha dt \quad (C.6)$$

where

$$\alpha = \mu_0 \frac{V_0}{P} \rho \quad (C.7)$$

Assuming that our ions drift from the surface of the electrode starting at time $t = 0$, we get the following integration:

$$\int_{r_2}^r u^2 du = \int_0^t \alpha du \quad (C.8)$$

$$\Rightarrow r = (r_2^3 + 3\alpha t)^{\frac{1}{3}} \quad (C.9)$$

The total ion drift time is then given by $r = r_1$, so $t_{\max} = \frac{r_1^3 - r_2^3}{3\alpha}$. Note that for our larger detectors, this gives $t_{\max} \sim 10s$! So by the time the next event happens, we will still have ions drifting from the previous one.

Putting Eq. C.6 and C.1 back into Eq. C.3, we get

$$dQ_{ind} = -q_{ions} \alpha \rho \frac{dt}{r^4} \quad (C.10)$$

and, with Eq. C.9

$$dQ_{ind} = -q_{ions} \alpha \rho (r_2^3 + 3\alpha t)^{-\frac{4}{3}} \quad (C.11)$$

Note that, if we integrate this between $t = 0$ and $t = t_{\max}$, we do get $Q_{ind} = -q_{ions}$,

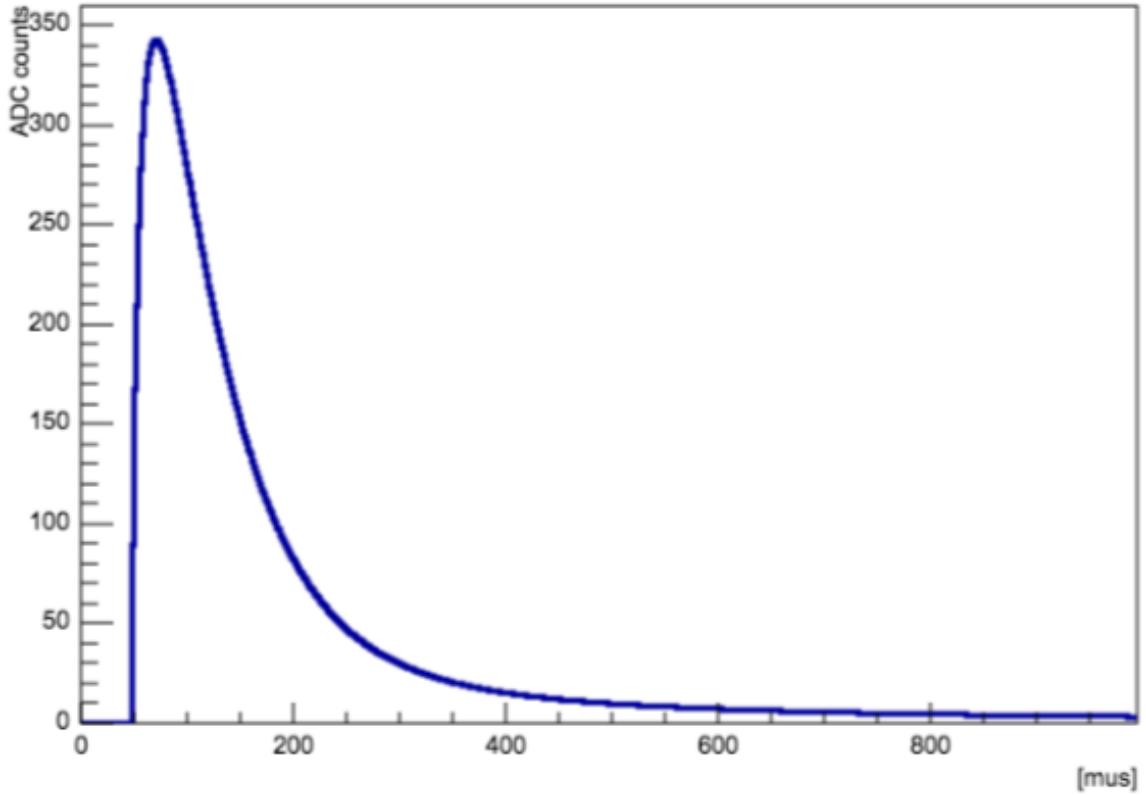


Figure C.2: Theoretical output voltage from the preamplifier for an avalanche that creates 1000 electron-ion pairs at time $48 \mu s$.

as expected. The shape of the induced current is shown in Fig. C.1.

The reason why we are interested in the induced current, $I_{ind} = \frac{dQ_{ind}}{dt}$, rather than the induced charge itself, is because of the following step in our signal formation, the “integration” of the signal by the charge preamplifier. The aim of the preamplifier, aside from further amplifying the signal, is to integrate the instant charge (in other words, a current) arriving on the electrode, and output a voltage that is proportional to the total collected charge. But to avoid our baseline voltage at the electrode shifting away from the one before the event, the preamplifier has to dissipate the charge fast enough as to not interfere with the following event.

This can be achieved with a preamplifier whose current response function is a decaying exponential

$$f(t) = e^{-t/\tau} \quad (\text{C.12})$$

where τ is the decay constant of the preamplifier. Note that there is no perfect preamplifier, but getting one that is close enough to the ideal case simplifies the following analysis. The main preamplifiers used both at LSM and with the Queen's prototypes are the Canberra 2006 [49] and Cremat CR-110 [82]. I will mostly discuss cases using the former, with $\tau = 47 \mu\text{s}$.

So, going back to our signal, we do not actually see the current induced by the ions drifting towards the outer sphere, but the convolution of that with the response function of the preamplifier, giving the shape seen in Fig. C.2. Unfortunately, there is no analytical way (to the best of my knowledge) to do this convolution, so any further steps would have to be done numerically. For example, the resulting function could be parametrized by:

$$S(t) = -q_{ions}k(e^{-at} - e^{-bt}) \quad (\text{C.13})$$

with a, b depending on τ and V_0 , $a < b$. While we don't use this parametrization directly in our analysis, it helps to understand the behaviour of our signal.

C.0.3 Signal induced by other charges

While the previous section covers the signal induced by secondary ions, it is reasonable to ask what signal is induced by secondary electrons, or even primary charges.

As the secondary electrons drift from the avalanche region to the central electrode,

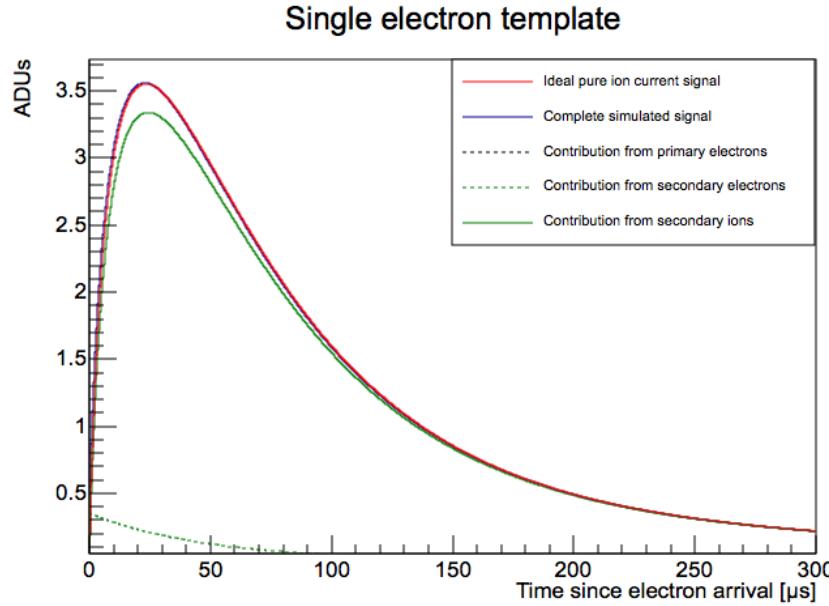


Figure C.3: Composition of an SPC signal, compared to an ideal pure-ion-current signal, for an event with a single primary electron. SEDINE conditions are assumed. The signal from secondary electrons is visible, but the shape of the total induced pulse is effectively indistinguishable from the ideal one. The signal from primary electrons is too small to be observable, even for a relatively low value of the avalanche gain of 1000.

they move in the field generated by said electrode, and so generate a current per the Shockley-Ramo theorem. Since they start so close to the electrode, they effectively induce all their charge instantly (less than 25 ns in the SEDINE SPC, while we are looking at effects in the tens of μ s). On the other hand, precisely because they start so close to the electrode, they don't induce all of their charge, but just a small portion of it. Ions, which move in the opposite direction, induce most of their charge, and so will tend to dominate the signal.

The exact proportion of electron to ion signal depends on the position at which they are created, which is in turn given by the avalanche process. Due to its exponential nature, the average secondary charge is created one mean-free path away

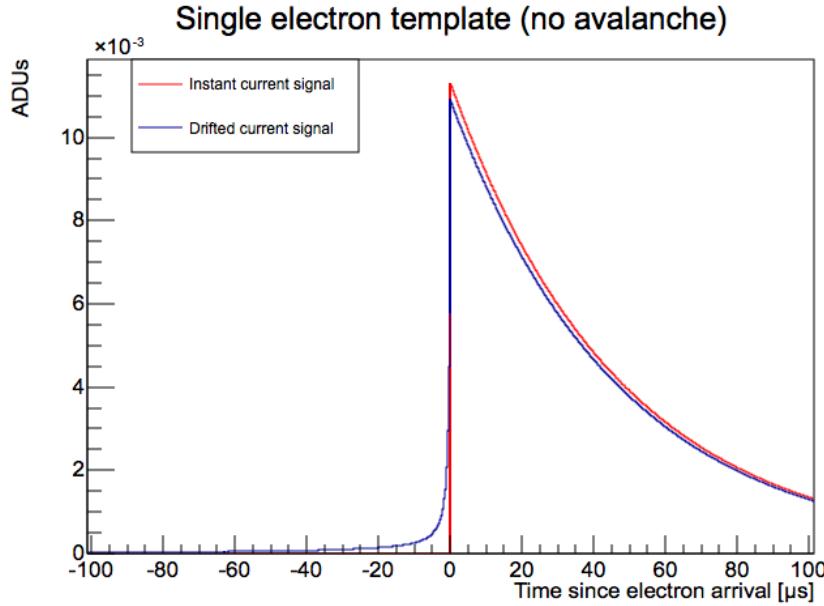


Figure C.4: Composition of an SPC signal, compared to an ideal instant-charge signal, for an event with a single primary electron, and no avalanche. SEDINE conditions are assumed.

from the electrode. In SEDINE's case, the mean-free path in the last few hundreds of μm of the avalanche region is under $100 \mu\text{m}$ ¹. Using the Shockley-Ramo theorem, assuming all secondary charges are created $100 \mu\text{m}$ away from the surface of the sensor, we obtain that secondary electrons account for 3.1% of the total induced signal. An illustration of the composition of the total signal is shown in Fig. C.3. We see that we can safely approximate the SPC signal as being produced by secondary ions generated at the surface of the central electrode.

Primary electrons can be similarly neglected. While they start farther away from the sensor, and so they induce most of their charge, their signal is negligible because they do not benefit from the multiplicative effect of the avalanche, which is generally in the $1000 - 10000$ range. However, they can sometimes be observed for events that

¹As derived from the Townsend coefficient obtained with Magboltz; see Sec. 5.3.

do not undergo an avalanche, if in large enough numbers. In that scenario, the signal they induce can be computed using the Shockley-Ramo theorem again, and the result is shown in Fig. C.4. Some structure is observed, with the electron inducing a signal before it reaches the electrode proper. This is not likely to be visible in practice: either it will be drowned out by baseline noise, or by the signal from other primary electrons arriving at slightly offset times. We can approximate it by primary electrons depositing all their charge when they reach the central sensor.

Finally, primary ions, which do not benefit from the avalanche, and start far from the electrode (inducing a small amount of charge spread out over a long period of time), are completely ignorable.

C.0.4 Ion space charge

As the ions slowly drift away from the avalanche region, their charge affects the electric field inside the detector, potentially changing the drift behaviour of charged particles and the strength of the avalanche. While we are not capable of directly identifying the space charge inside the detector, we can do some assumptions to find an analytical expression of this effect on the electric field, and get some qualitative conclusions.

We will assume again an ideal spherical detector, and a isotropic rate of ion creation at the sensor of R_A , corresponding to a charge flux of $I_A = R_A q_e$ in C/s . The electric field will still be radial due to the symmetry. We will also assume that the detector is in a steady state with the creation of ions from the avalanche happening at the surface of the sensor. Then, using conservation of charge on any sphere centred

on the sensor, we have:

$$I_A = 4\pi r^2 \cdot \rho(r)v(r) \quad (\text{C.14})$$

as the ions drift away from the sensor at the same rate as they are created in the avalanche region in the steady state. Then, using eq. C.5, we have:

$$I_A = 4\pi r^2 \rho(r) \mu E(r)$$

$$\Rightarrow \rho(r) = \frac{I_A}{4\pi r^2 \mu E(r)} \quad (\text{C.15})$$

In turn, the differential form of Gauss's law in spherical coordinates gives:

$$\vec{\nabla} \cdot \vec{E} = \frac{1}{r^2} \frac{d}{dr} (r^2 E(r)) = \frac{\rho(r)}{\epsilon_0} \quad (\text{C.16})$$

which, together with eq. C.15, gives a differential equation on $E(r)$:

$$\begin{aligned} \frac{d}{dr} (r^2 E(r)) &= \frac{I_A}{4\pi \epsilon_0 \mu E(r)} \\ \Rightarrow \frac{d(E^2(r))}{dr} + \frac{4}{r} E^2(r) &= \frac{I_A}{2\pi \epsilon_0 \mu r^2} \end{aligned}$$

The solutions to this differential equation, when the right-hand term is null, are of the form $E_0^2(r) = B/r^4$. They are strictly positive for $r > 0$, so eq. can be rewritten as:

$$\frac{d}{dr} (r^4 \cdot E^2(r)) = r^4 \cdot \frac{I_A}{\pi \epsilon_0 \mu r^2} = \frac{I_A}{2\pi \epsilon_0 \mu} r^2$$

$$\Rightarrow E^2(r) = \frac{1}{r^4} \left(B + \frac{I_A}{6\pi\epsilon_0\mu} r^3 \right) \quad (\text{C.17})$$

Finally, to “solve” for B , let’s call Q the electric charge on the sensor. If we apply Gauss’s law on the surface of the sensor, we need to satisfy:

$$\frac{Q}{\epsilon_0} = 4\pi r_{in}^2 E(r_{in}) = 4\pi r_{in}^2 \left(\frac{1}{r_{in}^2} \sqrt{\left(B + \frac{I_A}{6\pi\epsilon_0\mu} r_{in}^3 \right)} \right)$$

So we can rewrite eq. C.17 to get the final form:

$$E(r) = \frac{1}{r^2} \sqrt{\left(\frac{Q}{4\pi\epsilon_0} \right)^2 + \frac{I_A}{6\pi\epsilon_0\mu} (r^3 - r_{in}^3)} \quad (\text{C.18})$$

To solve for Q , we use $\int_{r_{in}}^{r_{out}} E(r) dr = V$, where V is the voltage applied on the sensor. Of note, if $I_A = 0$, then

$$Q = Q_0 = 4\pi\epsilon_0 r_0 V$$

and so

$$E(r) = V \frac{r_0}{r^2}$$

where $1/r_0 = 1/r_{in} - 1/r_{out}$, and we recover the expression found in the previous section.

Otherwise, $Q \lesssim Q_0$, but unfortunately has no analytical formula. It has to be solved for numerically so that the voltage difference between r_{in} and r_{out} matches the applied voltage on the sensor. A numerical integrator and a numerical solver are needed to find the value. An example is given in App. D.

Using eq. C.15 again, we can also find the analytical expression for the charge

density some distance away from the centre of the sphere:

$$\rho(r) = \frac{I_A}{4\pi\mu} / \sqrt{\left(\frac{Q}{4\pi\epsilon_0}\right)^2 + \frac{I_A}{6\pi\epsilon_0\mu} (r^3 - r_{in}^3)} \quad (\text{C.19})$$

Appendix D

Compute sensor charge from ion formation rate

Run the following code in Python to get the charge on the sensor for a given SPC setup and ion avalanche rate.

```

##First, we need to load the relevant modules and libraries:
import math
import scipy.integrate as integrate
from scipy.optimize import fsolve

##Define functions:
eps_0 = 8.854187e-12 #F/m = C/V/m

#Electric field at r
def E(r, Q, I_A, mu, r_in):
    return math.sqrt( (Q/(4*math.pi*eps_0))**2 +
                      I_A/(6*math.pi*mu*eps_0)*(r**3-r_in**3) ) / (r**2)

```

```
#Electric potential at r

def Delta_V(r, Q, I_A, mu, r_in):
    return integrate.quad( E, r_in, r, args=(Q, I_A, mu, r_in) )[0]

#Difference between electric potential at r and applied voltage

def V_equation(Q, I_A, mu, r_in, r_out, V):
    return Delta_V(r_out, Q, I_A, mu, r_in) - V

#Solve for Q to match potential difference with applied voltage

def Q_sensor(V, I_A, mu, r_in, r_out):
    if (I_A==0):
        rho = 1. / (1./r_in - 1./r_out)
        Q_sol = [4*math.pi*eps_0*V*rho, 0]
    else:
        Q_sol = fsolve(V_equation, (1.0e-9), args=(I_A, mu, r_in, r_out, V))
    return Q_sol[0]

##Run the solver for a given setup (change example values!)

V_test = 2520 #V (voltage at sensor)
r_in_test = 0.00315 #m (radius of sensor)
r_out_test = 0.30 #m (radius of sphere)
mu_ref = 7.5e-4 #m2/V/s * bar (ion mobility)
P = 3.1 #bar (pressure)
```

```
freq = 1000. #s-1 (event rate)
ener = 400000. #eV (energy per event)
gain = 3000. #(secondary electrons per primary)
W_e = 36. #eV (mean ionization energy)

mu_test = mu_ref/P #m2/V/s
q_ev = 1.6e-19/W_e #C/eV
I_A_test = freq*ener*gain*q_ev #C/s (avalanche ion current)

Q_test = Q_sensor(V_test, I_A_test, mu_test, r_in_test, r_out_test)

print "Computed sensor charge:", Q_test, "C"
```

Appendix E

SPC tests at Queen's University

E.0.1 Description

While the SEDINE detector is the only *low-background* SPC that is operational at time of writing, many more functioning SPC detectors exist. Queen's University possesses a range of them, which are used for a number of objectives:

- Sensor tests;
- Electronic tests;
- Gas handling system tests;
- Gas characterization;
- Other calibrations that cannot be performed with SEDINE.

They vary in size (ranging between 15 cm and 130 cm in diameter) and building material, although made mostly with stainless steel. This thesis focuses only on the results from SEDINE, and no further time will be spent discussing these other detectors (with one exception in Chap. 7). However, it is still important to point

out that the concepts discussed throughout this chapter have been tested in multiple configurations. Pictures of some of these test detectors are shown in Fig. E.1.

E.0.2 Available calibration sources

I will be briefly describing the main calibration sources used with Queen's SPCs, to give some context on the state of the art at time of writing. Not all of these were available for SEDINE, but some of the results extracted from the work performed at Queen's are still applicable; more detail on SEDINE calibrations will be provided in Chap. 6. Other laboratories that collaborate with NEWS-G

However, it is important to note that the exact gas composition (especially the presence of impurities), and the electrode sensor and voltage applied, have a large effect on electron drift characteristics, and especially on electron attachment and avalanche gain. In practice, this means that energy and drift and diffusion times must be recalibrated after any change in running conditions, even refilling the detector with the same nominal gas mixture and pressure.

^{37}Ar

^{37}Ar decays through electron capture into ^{37}Cl . Depending on whether the electron is absorbed from the K or L shell, the resulting atom will generally desexcite by emitting a characteristic X-ray of 2.82keV or a 270eV, respectively. Since argon is a gas, it spreads throughout the detector, so these events are reconstructed as uniformly distributed volume events. Due to the double monoenergetic lines, this source is useful for energy calibrations, and to test the linearity of detector response, together with the 0-energy point. It has the advantage that it can be put inside a

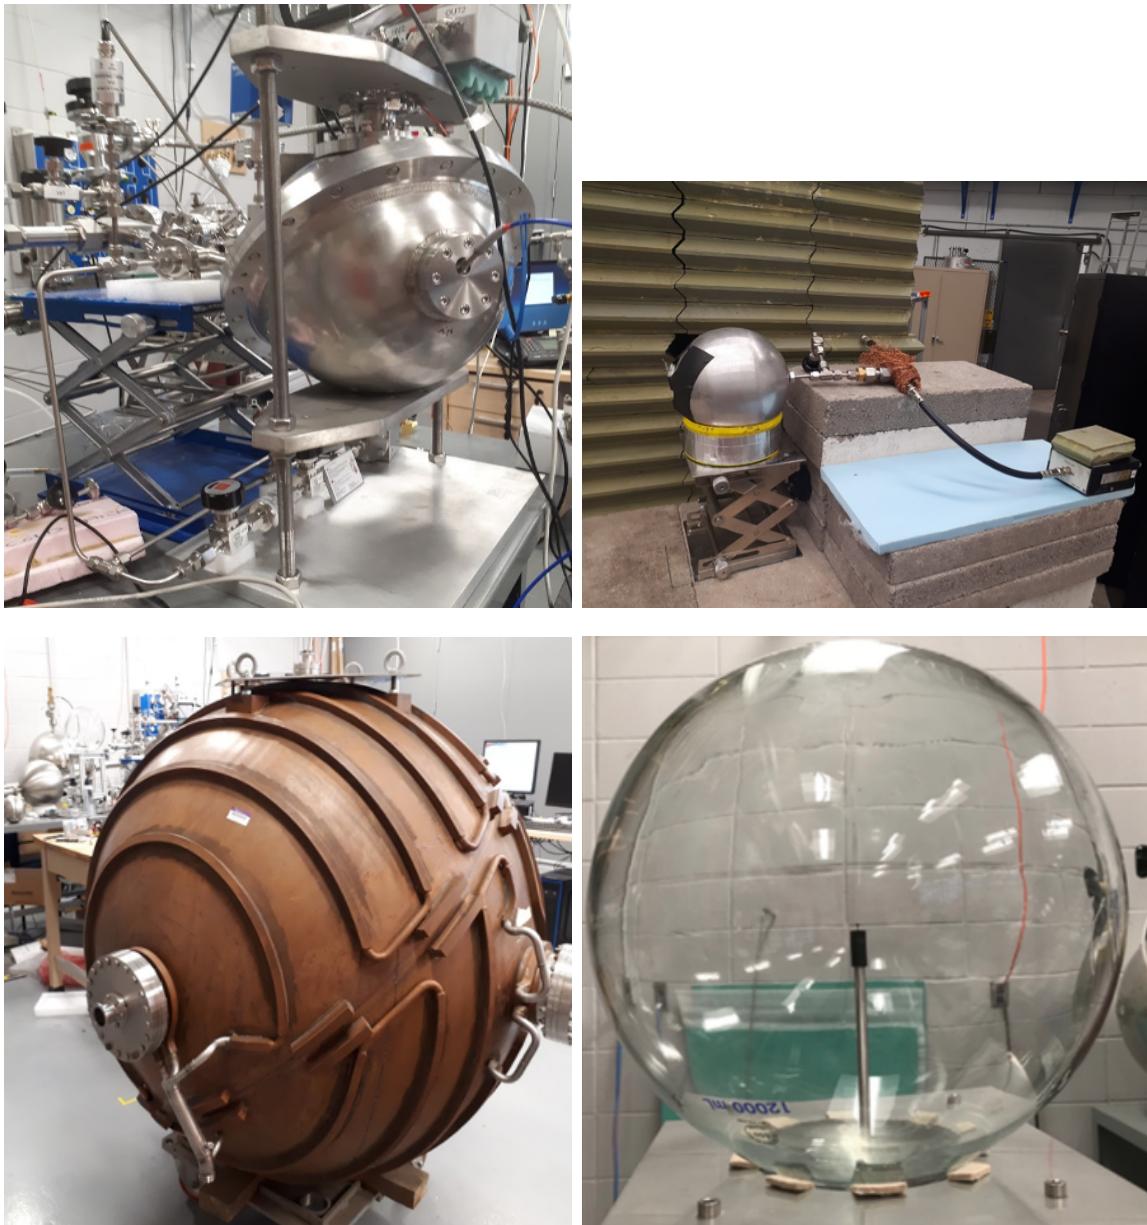


Figure E.1: Some of the SPC detectors available at Queen's University. Top left: 30 cm-diameter SPC. Top right: 15 cm-diameter SPC. Bottom left: 130 cm-diameter SPC at Queen's, originally a radio-frequency cavity from the Large Electron-Positron collider at CERN [84]. Bottom right: 30 cm-diameter glass SPC, for outreach.

detector without pumping, hence only minimally disturbing its running conditions, assuming no impurities are introduced during the procedure. However, it can only be removed by pumping the whole gas, restricting its use for calibrations of physics runs to only *after* they are complete.

Due to the relatively short lifetime of 35 days of ^{37}Ar , frequent calibrations require a renewable source of the isotope. For calibrations with Queen's test SPCs, we are provided with ^{37}Ar produced from neutron irradiation of calcium at the SLOWPOKE-2 reactor at Royal Military College of Canada [85], located in the same city as Queen's University.

^{55}Fe

^{55}Fe also decays through electron capture, generating 5.9keV photons. Unlike argon, iron is a solid, so it cannot be put inside the detector through the gas handling system. Since the attenuation length of photons of this energy in copper or steel is of the order of a few hundred micrometers, an external calibration with this source requires a window in the detector shell (typically plastic or aluminium foil covering a hole in the metal), in which case it can be freely put on or removed during any run. The attenuation length of photons of this energy in the gases we use is of the order of tens of centimetres, so it can be considered to generate uniformly distributed volume events for small to mid-sized SPCs. As for ^{37}Ar , ^{55}Fe is useful for energy calibrations, although it cannot be used to test the linearity of the detector response by itself.

In the absence of a window, the detector needs to be opened to place the source directly inside. This procedure is possible for easy-to-access test detectors that can either be opened or have large enough gas feedthrough to fit the source through one,

but impractical enough for shielded ones (such as SEDINE) that it is rarely, if ever, performed for them.

Copper fluorescence

Always present in copper detectors, but at a low rate, which can be enhanced by irradiating it with a strong gamma source. The energy of the emitted photon is 8.05 keV. While the source of the fluorescence is at the surface of the detector, the relatively high energy of these photons compared to the attenuation length of the gas compositions we use means that they are reconstructed as volume events; we assume them to be uniformly distributed in the detector, but this may depend on gas choice and pressure, and detector size. It is useful for energy calibrations when other sources are not available, or are otherwise impractical; it was the energy calibration used for the SEDINE physics data.

Pulsed Laser

By generating a beam of photons with energies in the ~ 10 eV range and hitting the internal surface of the detector shell, photoelectrons are released. Due to pulsed nature of the laser, an external trigger can be used to perfectly identify laser events from other background. The laser and a variable attenuator can be tuned to modify the average number of primary electrons extracted per pulse. This turned out to be a very versatile calibration: with judicious fitting, we can obtain average gain of avalanche, the theta of the Polya function, the Fano factor, and surface electron drift and diffusion time. Finally, combined with an energy calibration (such as with ^{37}Ar), it also serves as a calibration of the mean ionization energy of the gas. A paper

describing the full setup and analysis techniques to obtain this data in argon with a 30 cm Queen’s detector has already been published [38]. The same work is currently being performed for other gas mixtures in the upcoming detector at SNOLAB.

The two main difficulties are that this type of calibration requires a leak-tight laser feedthrough, which many of our detectors do not have, and that lasers suitable for photoelectron extraction are expensive, and difficult to relocate from one lab to another. In particular, SEDINE did not have such a feedthrough.

Neutron beams

Triangle University National Laboratory facility has a Tandem 10 MV accelerator, that can produce pulsed beams of quasi-monoenergetic neutrons [86]. The test SPC is put in the path of the neutron beam. Backing detectors are put behind the SPC, offset from the path of the beam, and are used as an external trigger; pulse-shape discrimination allows rejection of gamma background. The recoil energy of a neutron depends on its scattering angle, so by altering the position of the backing detectors, we can measure the electron-equivalent recoil energy for different nuclear recoil energies, even though the neutron beam is mono-energetic. A schematic drawing of the setup is shown in Fig. E.3. In parallel, the option to use the 1 – 8 MeV proton beam at the Reactor Materials Testing Laboratory at Queen’s University [87], impinged onto a nuclear target such as LiF, as an alternative neutron source, is being investigated by J.F. Caron.

A campaign with Neon as the target gas has been performed, with results currently being analysed; a paper by M. Vidal is in preparation. Neutron calibrations are used for quenching factor measurements. This is crucial for WIMPs searches, since they

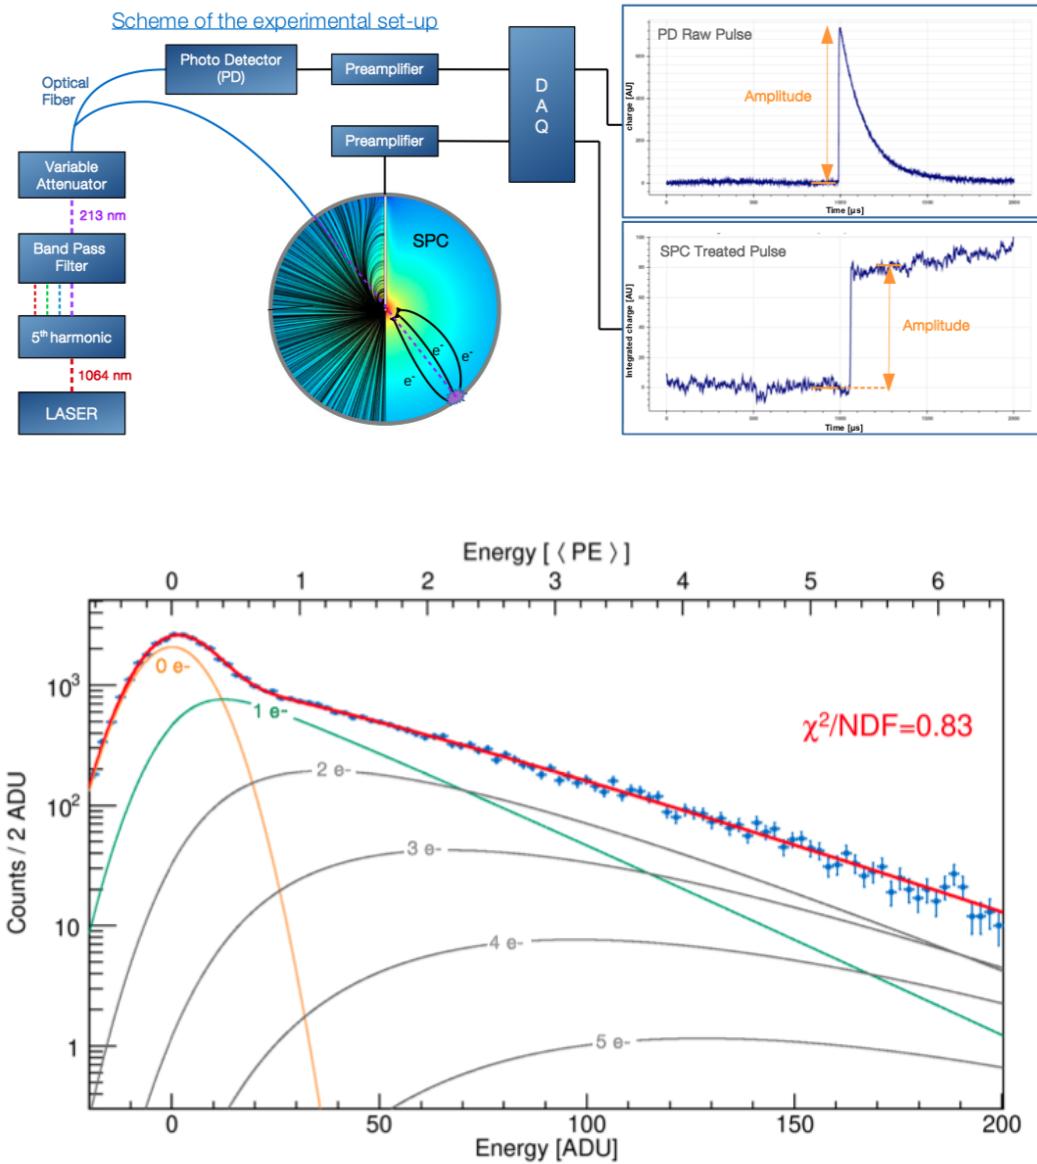


Figure E.2: Left: Experimental setup of a laser calibration at Queen's; the photodetector is used both as an external trigger and to monitor variations in the beam pulse intensity. Right: Example of energy spectrum from one such calibration with a relatively high number of primary electrons, together with the fit results (red: total, orange/green/black: contribution from different number of extracted primary electrons). Both of these diagrams are taken from [38].

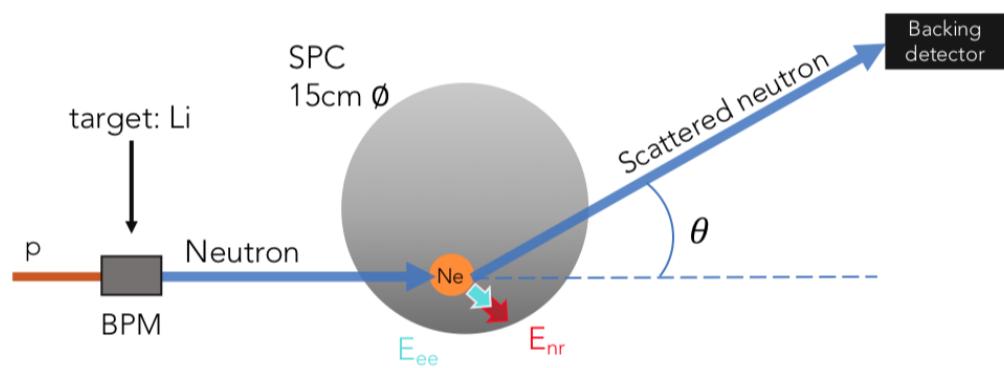
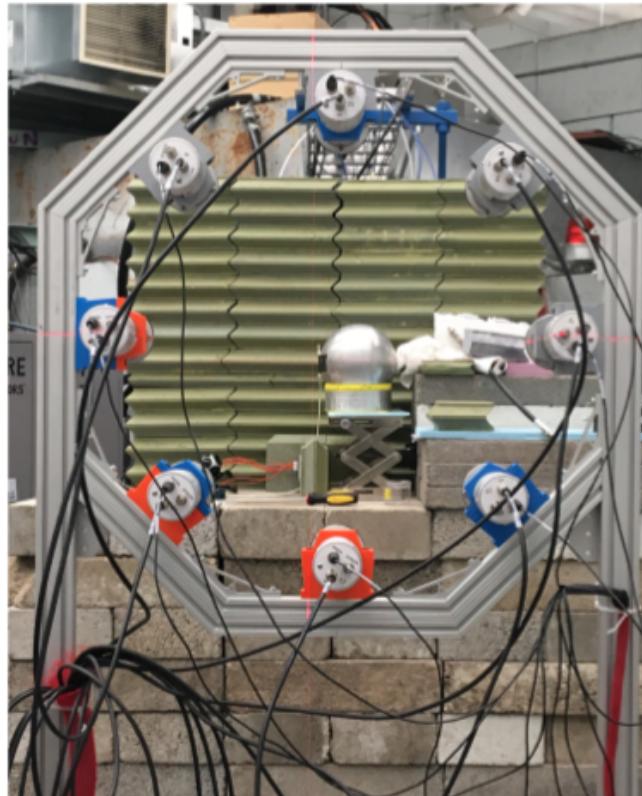


Figure E.3: Top: In the foreground, Backing detectors at TUNL; in the background, 15 cm SPC used for quenching measurements. Bottom: Schematic drawing of setup.

produce nuclear recoils. However, since KK axions decay into photons, which produce electronic recoils, this is not applicable for the work in this thesis (unless we were dominated by neutron backgrounds, but we did not explore that possibility).

Appendix F

Trigger algorithm

The trigger algorithm of our digitizer is the method that determines whether to save a slice of time from the voltage stream coming out of our preamplifier as an "event", or to just keep reading. The proper choice of this algorithm is important, because while we want to retain any potentially physically relevant event, we do not want to store noise (if for no other reason that we don't want to produce TeraBytes of "data" per day).

F.1 Trapezoidal filter

For almost all of the data we have taken, we have used the same algorithm: a trapezoidal filter combined with a threshold check. A trapezoidal filter is a generalized derivative with two parameters: an averaging window length L , and a gap size G [88]. The average over L samples is subtracted from the average, G samples earlier, over L samples:

$$f[i] = \frac{1}{L} \sum_{k=0}^{L-1} a[i + L + G + k] - \frac{1}{L} \sum_{k=0}^{L-1} a[i + k]$$

where $a[i]$ is the sampled values, and $f[i]$ are the filtered values. This is equivalent to taking the running average of the pulse with a smoothing window of L samples, then taking the derivative over $L+G$ samples of the resulting pulse. Note that the resulting pulse is shorter by $2L + G$ samples at the end. We use a derivative-like algorithm for our trigger algorithm because it is mostly unaffected by drifts in the baseline voltage coming out of the preamplifier; a simple threshold check would either trigger constantly or not at all if the baseline drifted too far up or too far down, respectively. Another advantage is that the simplicity of the algorithm makes it fast¹, so the digitizer can keep up with the data taking while filtering it. For an example of a trapezoidal filter applied to a low amplitude event, please see Fig. F.1. Note how the amplitude of the filtered pulse is higher than that of the raw pulse, for a comparable level of baseline noise to the baseline noise. A judicious choice of the gap (closer to the expected risetime of our pulse) and averaging window (smaller than the risetime, but close to the period of high-frequency noise) will tend to boost the signal over noise ratio of the result.

There is also a time shift between the results obtained with this formula and what one may expect from doing a running average, and then a derivative. This becomes particularly noticeable when we use the result from a trapezoidal filter to determine the beginning of the pulse for our advanced processing algorithms. If we want a filtered pulse whose features start at the same sample as the raw pulse, the filtered pulse must be shifted forward by $2L + G$ samples; if instead we want a filtered pulse whose features are centred around the same samples as the raw pulse, the final pulse must be shifted forward $\frac{3}{2}L + G$ samples. Note that while the former is often more

¹A recursive version of the formula can easily be found: $f[i+1] = f[i] + ((a[i+2L+G] - a[i+L+G]) - (a[i+L] - a[i]))/L$, requiring only four sums and a division for each new sample; this can be programmed into an FPGA.

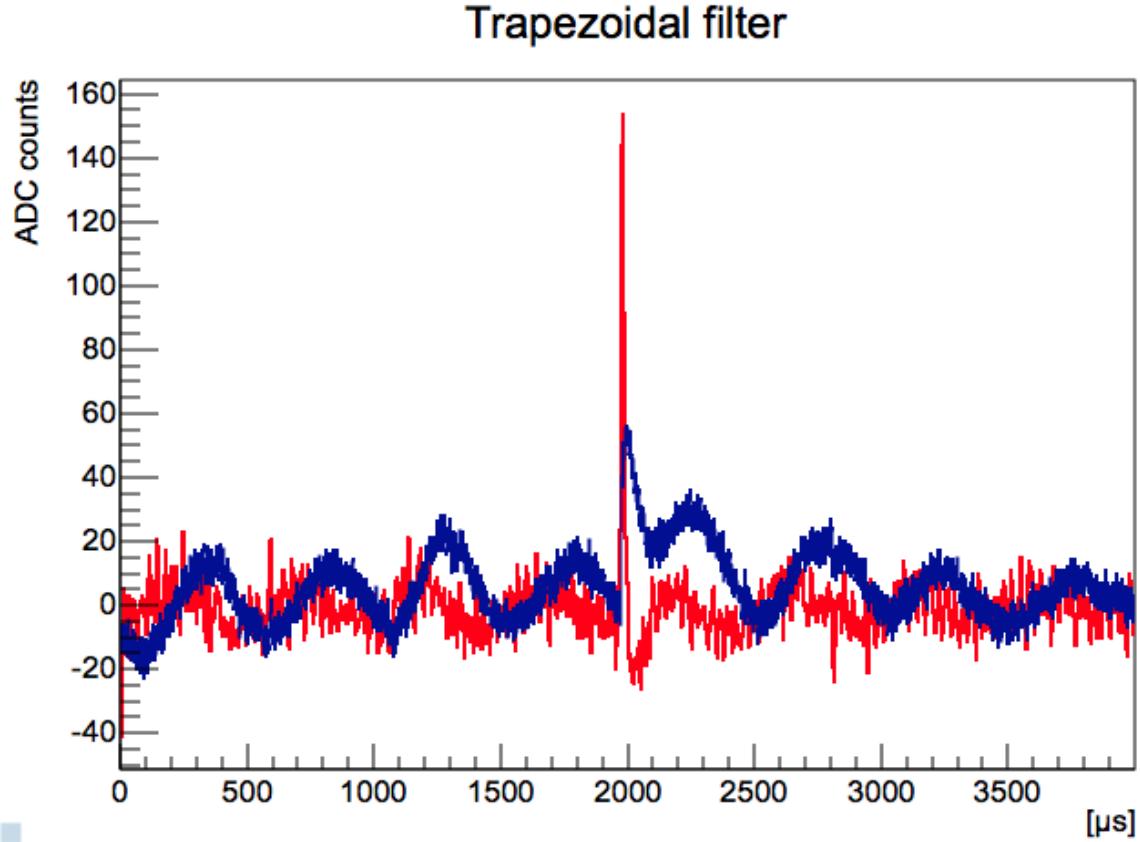


Figure F.1: In blue, raw pulse (after removing average baseline, for comparison). In red, pulse after trapezoidal filter, scaled up by a factor of 4 to match the noise level of the raw pulse. The combination of a running average and a derivative allows the reduction of high and low frequency noise respectively, boosting the Signal-to-Noise Ratio.

useful for trigger time considerations, the latter has the more “intuitive” behaviour. Obviously both are equivalent if the only use of the algorithm is to check whether the filtered pulse does or not reach a certain threshold.

F.2 Performance

F.2.1 Signal-to-Noise Ratio

Two different metrics will be presented for qualifying the performance of the algorithm. The first, and more naive one, is based on the concept of Signal-to-Noise Ratio (SNR). The “signal” is defined as the maximum of the event with respect to the baseline. For the “noise”, we can use the Root Mean Square (RMS)² of the baseline. This is what the digitizer computes to give an estimate of how “noisy” any given run is: the higher the RMS, the wider the baseline oscillations are, and hence the harder it will be to distinguish a pulse in between those variations.

We can compare the SNR for different algorithms. For the sake of illustrating the concept, I will be comparing the raw event, the running average of the event (both after subtracting the baseline), and the event after applying a trapezoidal filter. We would logically expect the processing with the highest ratio to be the optimal one. This was tested on simulated events containing a single primary electron. The results are shown on Fig. F.2. Unsurprisingly, the SNR is lowest for the raw, unprocessed event. However, the average is fairly close for both the smoothed event and the trapezoidal filter, with the latter having more consistent values. This could be explained by the fact that removing the baseline and doing a running average are essentially the same two steps that the trapezoidal filter does, only less optimized for the search of rising edges. This would suggest that both approaches should perform similarly well as triggering algorithms.

²In Physics, the term RMS is often used to refer to the standard deviation of the samples, rather than to its “real” definition: the square root of the mean of the squares of the samples.

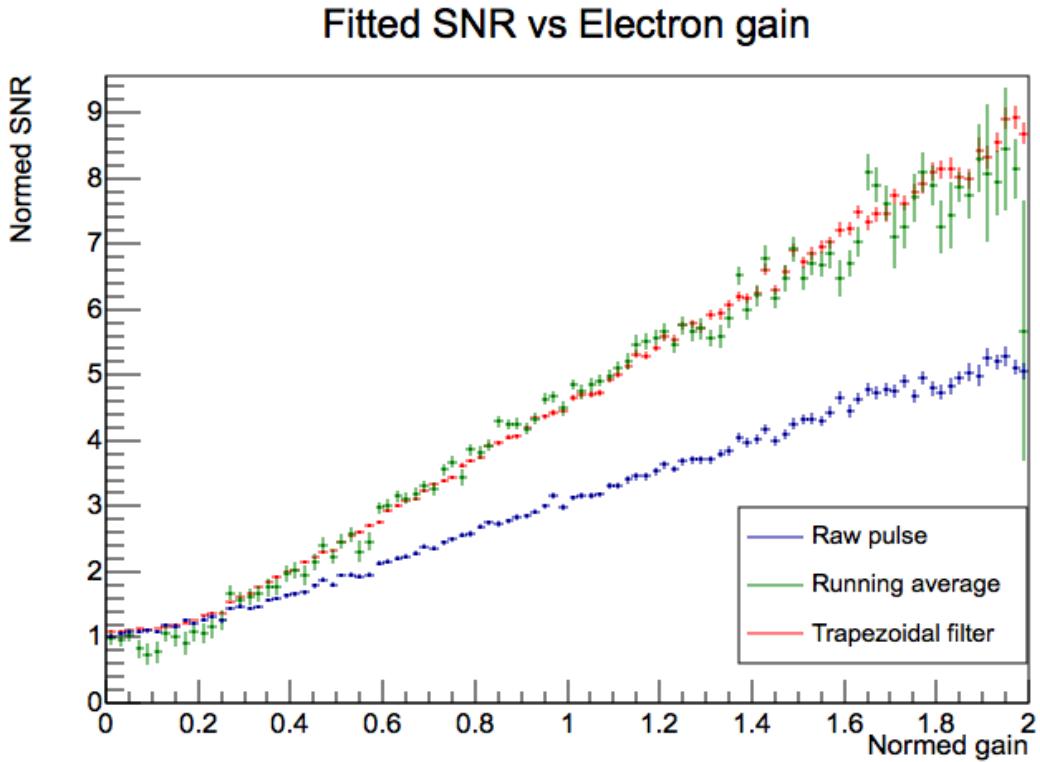


Figure F.2: Signal-to-Noise ratio for the raw pulse, smoothed pulse, and the trapezoidal filter of the pulse. The X axis is the gain of the primary electron, with respect to the average gain. The running conditions and noise traces were taken from the physics run.

F.2.2 Efficiency and False positives

The SNR turns out to be a misleading metric, as we will see now. Instead of looking at SNR as a proxy for the capacity of algorithms to distinguish the signal among random oscillations, we can look directly at the efficiency and false positive rate of each algorithm when detecting signals among noise. The efficiency is the proportion of electron signals that triggered; since this depends on the amplitude and time distribution of those signals, we once more take single electron events to simplify comparisons. The false positive rate is the frequency of triggers produced by pure noise

with no physical signal. This can be tested by running the algorithms on empty noise traces taken with our detector, giving a result in Hertz (number of triggers divided by total time length of all empty traces).

Notably, since all three algorithms end by comparing to a fixed threshold to determine whether the signal triggered, we can modify the value of the threshold to “sweep” the efficiency - false positive parameter space for each. The results are shown on Fig. F.3. The false positive rate is computed as $f = -\ln(1 - p)/\Delta t$, where p is the proportion of noise windows of length Δt that triggered; this gives $f \simeq p\Delta t$ when $p \ll 1$, the horizontal black line in the plot shows the point above which this approximation does not hold any more.

It is immediately clear that the relative performance of the different algorithms revealed by this plot is different from the one suggested by the SNR. We see that the raw pulse and the smoothed pulse have almost identical curves, with the trapezoidal filter well below them. In fact, for any desired detection efficiency of single electron events, the trapezoidal filter has one to two orders of magnitude less false positives than the other two.

The disagreement between both metrics as to the performance of applying a threshold to the smoothed pulse can only be explained if our estimation of the “noise”, the RMS of the baseline, is not actually well-chosen for this context. More likely, the main contributor to the lower performance of the raw and the smoothed pulses is the need to compute a fixed baseline, despite the low frequency components of the noise. The running average does improve the baseline RMS, and hence the SNR of the method. However, they both suffer from depending on the value of the baseline being computed properly, before it is subtracted for comparison with the threshold.

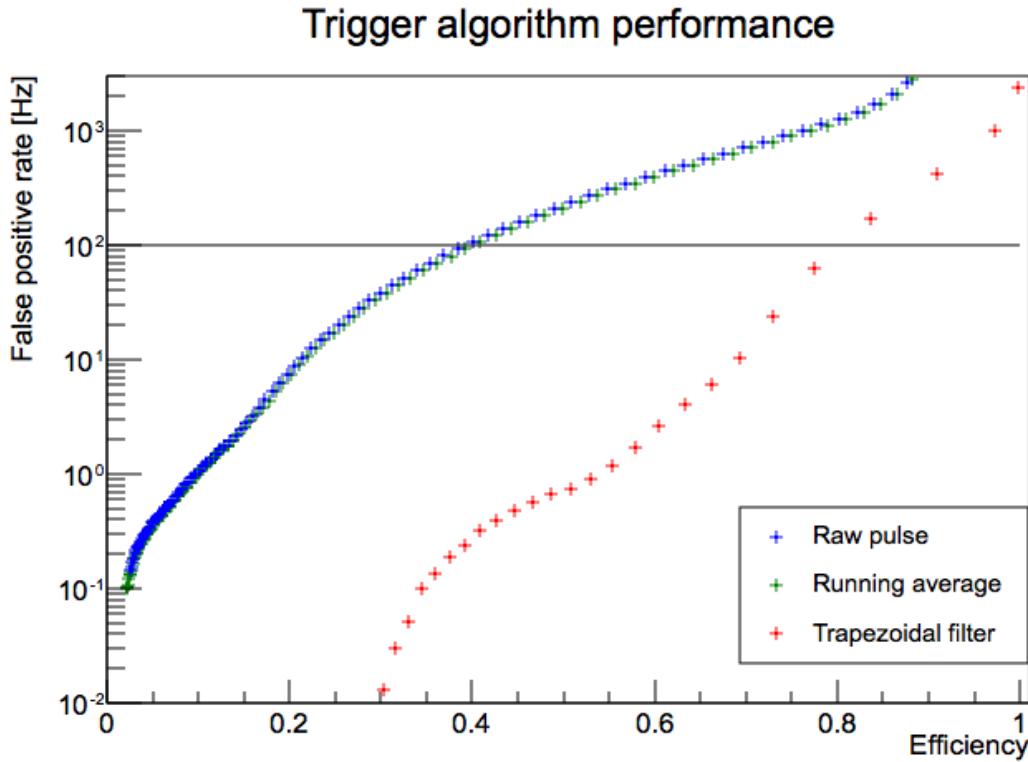


Figure F.3: False positive rate depending on event detection efficiency, for thresholds applied to the raw pulse, smoothed pulse, and the trapezoidal filter of the pulse. Above the black horizontal line, the false positive rate was extrapolated from the proportion of empty traces that triggered. The small bump at low efficiencies for the trapezoidal filter is likely caused by very low energy events sneaking into the noise traces used to compute the false positive rate.

As such, both the raw and the smoothed pulse would still perform poorly in terms of efficiency and false positives.

This study should be taken as a word of caution against using performance metrics that are not well understood. While there might be a definition of the noise that produces a useful SNR for the sake of comparing triggering algorithms, it is clear that the baseline RMS is not it. Caution is recommended when trying to interpret

changing values of the RMS ³. This conclusion can be generalized to other data processing algorithms: we must be careful not to draw conclusions from variables that we believe are good stand-ins for what we are really interested in, until we confirm they actually are.

³Unlike we often do in our collaboration, where the baseline RMS pre-defined in our digitizer is routinely used to compare noise levels between setups.

Appendix G

Proper normalization of transforms of discrete data

There are many subtleties when doing Fourier transforms, and when transferring mathematical concepts from continuous time to the discrete time of computers. I will try to cover the solutions I ended up going with when trying to solve normalization problems that arose during pulse processing. An effort was made to make analysis methods that are as independent as possible from the sampling frequency of our digitizer, since it is not a relevant physical quantity for any of our searches; at times, we had to change digitizer settings while keeping the running conditions of the detector the same. Note that I will assume that all functions involved in this section are sufficiently well-behaved for me to invert the order of integrals liberally.

G.1 Integrals and derivatives in discrete time

Without going into the details of the inner workings of our digitizers, the data we have access to is the value of the voltage coming out of the preamplifier at discrete times. From a coding perspective, this should be represented by a graph of (x,y)

points where x is the time and y is the voltage. However, due to practical concerns about the implemented methods for TGraph's and TH1's in ROOT (our framework of choice for data analysis), we often represent our signal as histograms instead. The height of each bin then represents the value of the voltage at each point in time, with the point in time being either the time value of the first edge of the bin, or the centre of the bin. Note that either definition of the time is equivalent, since we ever only care about the relative position of features in a signal, not their absolute position. We just need to stay consistent throughout.

The choice to represent signals as histograms¹, instead of graphs, tends to obscure the nature of differentiation and integration, leading to problems with normalization down the line. Noting $S(i)$ the value of the signal at bin i , we may be tempted to note the derivative of the signal as $s(i) = S(i + 1) - S(i)$. This leads to an obvious problem: the values of the derivative depends on the sampling period. Indeed, let's assume we have a steadily increasing signal, $S(t) = at$. Now, let's consider what happens for two different sampling periods, Δt and $\Delta t'$, then the digitized signals will respectively be $S(i) = a(i \Delta t)$ and $S'(i) = a(i \Delta t')$. If we use the naive definition for differentiation mentioned above, we get $s(i) = S(i + 1) - S(i) = a \Delta t$ and $s'(i) = S'(i + 1) - S'(i) = a \Delta t'$. Both are constant, but their values depend on the sampling period of the digitizer, which is not a physical value that should affect our measure! This is solved by defining the derivative as $s(i) = (S(i + 1) - S(i)) / \Delta t$. Care has to be taken that the derivative has the proper units: assuming S was in Volts and Δt was in μs , $s(i)$ is now in $V \mu s^{-1}$. While this may appear obvious in hindsight, it is an easy issue to overlook when the sampling period stays constant for

¹When using ROOT, it is common to use histograms to present all data, because it comes with a better axis system.

long periods of data taking. Normalizing derivatives properly allows analysis methods to keep working as expected when changing sampling periods.

Note that while I defined $s(i) = (S(i + 1) - S(i))/\Delta t$, in our case where we are often looking for the start of features in our signal, it makes more sense to use $s(i) = (S(i) - S(i - 1))/\Delta t$. Indeed, the previous definition will show a feature for the derivative appear on bin i if the raw signal has a feature appear at bin $i + 1$, inducing an off-by-one error. I did not attempt to use any other kind of more advanced derivative definition, since they tend to smooth out the signal, and I wanted the smoothing to be done independently. If someone were to use some other definition, I would advise to keep in mind this sort of off-by-one error that may appear between features in the raw signal and the differentiated signal.

Lastly, once we have defined $s(i) = (S(i) - S(i - 1))/\Delta t$, then the formula for integration has to perfectly reverse this process: $S(i) = \sum_{k=0}^{k=i} s(k) \Delta t$. Since $s(0)$ is not defined (because $S(-1)$ does not exist), to keep differentiation reversible we fix $s(0) = S(0)/\Delta t$. This is a meaningless value, but the upsides of having a reversible differentiation compensate for it; otherwise, you may set $s(0) = 0$ or $s(0) = s(1)$, depending on the context. Note again the change in units: we want Δt to be in the same time unit as $s(i)$ (e.g. μs and $V \mu s^{-1}$).

It was important to establish this before talking about the normalization of Fourier transforms and convolutions/deconvolutions, since their respective definitions also involve integrations.

G.2 Fourier transforms

Fourier transforms are notorious for not having generally agreed-upon conventions. The choice of normalization and windowing is often done depending on the specific use, as long as it maintains $\tilde{\mathcal{F}}(\mathcal{F}(f)) = f$. In our case, we have two uses for Fourier transforms: deconvolutions, and comparing noise levels between different runs.

G.2.1 Convolutions and deconvolutions

Let's start with convolutions. Remember that the definition of the convolution between f and g is $f * g(t) = \int f(\tau)g(t - \tau)d\tau$. We want the convolution theorem to be true without adding further normalization constants:

$$\widetilde{f * g} = \tilde{f} \cdot \tilde{g}$$

Let's go back to the proof of the theorem. Expanding the left hand side, and noting $h = f * g$ and k the Fourier normalization, constant, we get:

$$\tilde{h} = k \int \left(\int f(\tau)g(t - \tau)d\tau \right) e^{-2\pi i \nu t} dt = k \int \left(\int g(t - \tau)d\tau e^{-2\pi i \nu t} dt \right) f(\tau)d\tau$$

Substituting $T = t - \tau$, we get

$$\begin{aligned} \tilde{h} &= k \int \left(\int g(T)d\tau e^{-2\pi i \nu (T + \tau)} dT \right) f(\tau)d\tau \\ &= k \int f(\tau) e^{-2\pi i \nu \tau} d\tau \left(\int g(T)d\tau e^{-2\pi i \nu T} dT \right) \end{aligned}$$

$$= \frac{1}{k} \tilde{f} \cdot \tilde{g}$$

So, if we want $\widetilde{f * g} = \tilde{f} \cdot \tilde{g}$ to be true without extra normalization terms, then $k = 1$. So our definition of the direct and inverse Fourier transforms will be:

$$\begin{aligned}\mathcal{F}(f)(\nu) &= \int f(t) e^{-2\pi i \nu t} dt \\ \bar{\mathcal{F}}(\tilde{f})(t) &= \int \tilde{f}(\nu) e^{+2\pi i \nu t} d\nu\end{aligned}$$

Note that if we want to describe the Fourier transform in terms of the angular frequency ω instead of the “real” frequency ν , with $\omega = 2\pi\nu$, then the definitions become:

$$\begin{aligned}\mathcal{F}(f)(\omega) &= \int f(t) e^{-i\omega t} dt \\ \bar{\mathcal{F}}(\tilde{f})(t) &= \frac{1}{2\pi} \int \tilde{f}(\omega) e^{+i\omega t} d\omega\end{aligned}$$

where a normalization factor appears on the inverse Fourier transform, to compensate for the fact that ω and t are not in inverse units any more (*rad/s* and *s*, respectively). For that reason, I will avoid using angular frequencies in this document where Fourier transforms may be involved.

Going on a bit of a tangent, the physical interpretation of the convolution will depend on the dimension of the ‘response’ function. For example, assuming that f and g are such that $f * g$ exists, and that our ‘response’ function g is normalized to 1, i.e., $\int g(t) dt = 1$, then:

$$\int f * g(t) dt = \int \int f(\tau)g(t - \tau) d\tau dt = \int f(\tau) \left(\int g(t - \tau) dt \right) d\tau = \int f(\tau) d\tau$$

In other words, convolving f by a 'normalized' function does not change the numeric value of the integral of f . This has two noteworthy consequences. The first one is that if the response function of our detector is short compared to the signal, then the integral of the raw pulse will have the same amplitude as the integral of the deconvolved pulse, aside from a proportionality term; this is, however, rarely the case. The second one is a bit less obvious and has to do with the fact that a convolution of f with g has dimensions of $[f][g] \cdot s$, which in general are not the same dimensions as either f or g .

In our case, we have three signals that we convolve with each other: the instant current from primary electrons, the induced current from a secondary ion, and the response function of the preamplifier. If we want the convolution of all three together to be in Volts, while staying consistent with the definition of the convolution, then the proper dimension for each becomes:

- Primary charge instant current: charges/s
- Ion induced current: C/s
- Preamplifier response: V/C

G.2.2 Noise analysis

Noise analysis is commonly done in frequency space. The preferred way to study the noise is via a Power Spectral Density (PSD), defined as:

$$\text{PSD}_n(\nu) = |\tilde{n}(\nu)|^2/T_0$$

where $\tilde{n}(\nu)$ is the Fourier transform of the noise $n(t)$ in a window of duration T_0 . Note that the unit of the PSD is AU²/Hz, where AU (arbitrary unit) is the unit of $n(t)$. This formula has the advantage that for stationary noise (i.e., noise that does not change with time), the PSD does not depend on the size of the window. This is due to Parseval's theorem, which states that:

$$\int_{-\infty}^{+\infty} x(t)^2 dt = \int_{-\infty}^{+\infty} x(\nu)^2 d\nu$$

For a stationary noise with mean 0, $\text{Var}(x) = (\int_t^{t+T_0} x(t)^2 dt)/T_0$ is the variance of the noise computed on a window between t and $t + T_0$, which necessarily does not depend on t . Furthermore, the only dependency on T_0 is whether or not low frequency components of the noise are contained in the window. If all components have periods either much smaller or much larger than T_0 , then $\text{Var}(x)$ is mostly independent of the window size. Parseval's theorem then gives us the same property for the PSD.

The reason why we want to study the noise in a window-size independent way is that we often have to compute the Fourier transform of the noise on windows of different sizes. While this will necessarily have an effect on our estimation of the noise at low frequencies, no matter the approach we choose, the PSD at high frequencies remains more or less unaffected.

It is important to note that this property of the PSD does not hold true for transient signals. Unlike for stationary processes, for which the total energy increases linearly with the time of integration, a transient has a finite energy contained in a finite time support. As such, its Power Spectral Density will be inversely proportional to the length of the window over which it is computed, as long as the window is larger than the support. This means that we must be careful while doing any procedure that requires comparing noise and transient in frequency space, since the ratio between the two will roughly vary linearly with the window size.

G.2.3 Discrete Fourier Transform

We need to check that the concepts that work in continuous time can be translated properly to discrete time, since that is the nature of our data. To do Fourier transforms, I have access to the FFTW3 [89] library through the ROOT class `TVirtualFFT`. FFTW3 uses a planner to choose between different strategies for the Cooley-Tukey algorithm and other non-Cooley-Tukey algorithms to perform the Discrete Fourier Transform (DFT) of an array X of n complex numbers [90]. While the interpretations of the DFT depend on the scenario, it is useful to understand it as periodically extending the original signal beyond its original window (potentially leading to edge effects), and then computing the finite number of coefficients in Fourier space necessary to reproduce that periodic discrete signal. The algorithm itself is not relevant for this section, only the final values computed, which are:

$$\text{DFT}(X)[k] = \sum_{j=0}^{n-1} X[j] e^{-\frac{2\pi i}{n} jk}$$

$$\text{IDFT}(Y)[j] = \frac{1}{n} \sum_{k=0}^{n-1} Y[k] e^{+\frac{2\pi i}{n} jk}$$

where $0 \leq k < n$, though it is easy to see that $\text{DFT}(X)[k]$ and $\text{IDFT}(Y)[j]$ are periodic with period n ². A particularity of the FFTW3 library is that its computation of the inverse Discrete Fourier Transform does *not* normalize its output, either by the length of the series or otherwise, leaving that to the user. Other libraries may do that automatically.

Of note, the definition of the DFT and IDFT do not include any information on time. We can change the definition so that it includes that information. Assuming we have a function f that we sample n times during a time window $[0, T]$, then

$$\text{DFT}f(k \nu_0) = \sum_{j=0}^{n-1} f(j t_0) e^{-2\pi i j t_0 \cdot k \nu_0}$$

where $t_0 = T/n$ is the sampling period, and $\nu_0 = 1/T$ is the smallest non-zero frequency we can access with a window of length T . This definition has a similar problem than the naive derivative described above: if we halve the sampling period, then the number of terms in the sum doubles. By taking the simple case of a constant function, that would also double the values of the DFT of the sampled function. This is not a desirable behaviour: I want the DFT to be as agnostic as possible from the sampling frequency. This can easily be achieved by multiplying the DFT by the value of the sampling period:

²Do not confuse the Discrete Fourier Transform with the similarly-named Discrete-Time Fourier Transform. The latter is defined for all frequencies, instead of just a discrete, finite amount of them. All FFT algorithms compute the DFT of an array. I assume the unambiguous "Discrete-Time Discrete-Frequency Fourier Transform" was too much of a mouthful to catch on.

$$\text{nDFT}f(k \nu_0) = t_0 \sum_{j=0}^{n-1} f(j t_0) e^{-2\pi i j t_0 \cdot k \nu_0}$$

which now has same dimension as a continuous Fourier transform. In fact, this formula tends towards the continuous Fourier transform of f defined between $[0, T]$ as $n \leftarrow \infty$. This shouldn't come as a surprise, since this modification is equivalent to the one we did earlier to go from an integral in continuous time to one in discrete time. The inverse DFT then becomes:

$$\text{nIDFT}\tilde{f}(j t_0) = \frac{1}{N t_0} \sum_{j=0}^{n-1} \tilde{f}(k \nu_0) e^{+2\pi i j t_0 \cdot k \nu_0}$$

which also has the same dimensions as a continuous inverse Fourier transform, with the extra term $\frac{1}{N t_0} = \nu_0$.

Thanks to this non-standard normalization of the DFT, we can now reuse all the relationships for continuous Fourier transforms previously described in this section without having to add extra proportionality factors:

$$\text{nDFT}(f * g) = \text{nDFT}(f) \cdot \text{nDFT}(g)$$

$$\text{nDFT}(f *^{-1} g) = \text{nDFT}(f) / \text{nDFT}(g)$$

$$\sum_{j=0}^{n-1} [f(j t_0)]^2 t_0 = \sum_{k=0}^{n-1} [\text{nDFT}f(k \nu_0)]^2 \nu_0$$

Note that none of this are new results. They are just the DFT relations for the convolution theorem and Parseval's theorem, where my normalization convention happens to absorb the extra terms. I believe that the straightforward conversion from the continuous Fourier transform version of these relations minimizes the potential

for errors in many situations.

Appendix H

Parameter estimator performance

Choosing the right metric to compare algorithms can be a more subtle topic than it might seem at first view. In this section, we go into more detail in the comparison of the estimator resolutions for the algorithms described in Sec. 4.1: SAMBA, MDec, and DD estimators.

H.1 Amplitude resolution

Both the MDec and DD methods correct the amplitude bias due to the ballistic deficit, so we need to compare them on the basis of the precision of their estimators. The first, and arguably more important, is the amplitude estimator, since it gives a measure of the energy of the event. We want a metric for the reconstruction power of each estimator.

The first thing to note is that, while generally what we are interested in is the energy of our event, this is not what we have access to in our data. What we do measure is the total charge created in the avalanche region. This is proportional to the energy of the event *on average*, but this correlation is not exact event by event. This is because the number of primary electrons created is randomly drawn from a Poisson

distribution based on the energy of the event, and in turn the number of secondary charges created is drawn from a Polya distribution based on the number of primary electrons. As such, the if we tried to compare the resolution of the energy estimator from the different methods, we would be folding in the systematic uncertainty from those two physical processes, making for a poor comparison metric.

For this reason, the comparison between algorithms should be based on their capacity to reconstruct the total secondary charges created during the event. Since we do not have any calibration source capable of producing events with a known number of primary electrons, let alone a known number of secondary charges, this necessarily relegates the comparison to the application of the different methods to simulations, where we can keep track of all this information.

The “naive” way of comparing the methods is to look at the relative dispersion of the amplitude estimator for events with a fixed number of secondary charges. The result for all three methods is shown in Fig. H.1. The X-axis for the relative resolution of the methods is the number of secondary charges for each event, scaled to correspond with the event energy most likely to produce that number of charges. This scaling was performed so that the systematic uncertainty due to Poisson (primary ionization) and Polya (avalanche) may be shown in the same plot. As can be seen, those uncertainties are at least one order of magnitude higher than the resolution from the MDec and DD methods for energies above 1.5 keV, and are only of the same order under 100 eV. This puts the comparison of those two methods into perspective, since any improvement in resolution will be washed out by the systematics.

However, both MDec and DD do provide a substantial gain over the SAMBA estimator. The resolution of SAMBA is limited to a minimum of 8 %, due to the

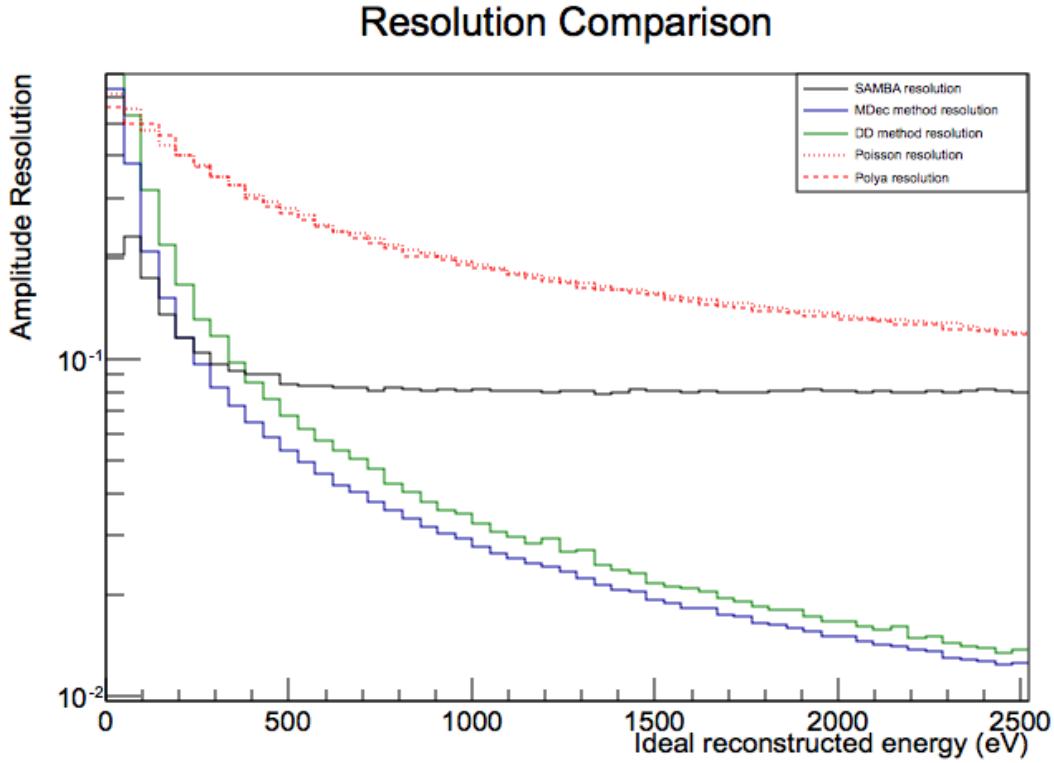


Figure H.1: Amplitude resolution of all three methods, depending on the secondary charges of the event (in units of the energy most likely to produce that number of secondary charges). For comparison, the relative resolution of the number of primary electrons produced by any given energy, and the relative resolution of the number of secondary charges produced by any number of primary electrons (in units of the energy most likely to produce that number of primary electrons).

bias introduced by the changing ballistic deficit. MDec and DD are not affected, so the resolution keeps improving with increasing energies, down to under 1% at 5 keV. However, at low energies, the odds of having a single primary electron (or a low number of electrons arriving at the same time) increases, so the SAMBA estimator is not affected by the ballistic deficit anymore, and so the resolution of all methods converge together.

The reason why the systematics from Poisson and Polya match is just a coincidence

due to a chosen θ for Polya of 0. In that scenario, the Polya distribution is just an exponential distribution, for which the standard deviation and the average are equal. Then, the relative resolution of from Poisson where the average is n is

$$r = \frac{\sigma_{Poisson}}{\mu_{Poisson}} = \frac{\sqrt{n}}{n} = 1/\sqrt{n}$$

And, in turn, the relative resolution from adding n results from an exponential distribution of average G is:

$$r = \frac{\sigma_{Exp^n}}{\mu_{Exp^n}} \simeq \frac{1}{\sqrt{n}} \frac{\sigma_{Exp}}{\mu_{Exp}} = 1/\sqrt{n}$$

Going back to the characterization of this metric as “naive”: the relative resolution of an estimator for a fixed value of the quantity being estimated is not often well-behaved. The obvious counterexample is that of an estimator that always predicts the same value: this would have a perfect resolution of 0, but would obviously be of no interest. This is because what we actually want from an estimator is a good resolution on the *real* value of the quantity, for a given *estimated* one, instead of looking at it the other way around, as previously described.

This comes with one added nuance: the results now depend on the energy distribution of our (simulated) data. This is conceptually similar to the concept of a prior in bayesian analysis: the most likely value of the quantity estimated depends on the a priori probability of each value, before the measurement. To be as agnostic as possible on the energy distribution, we can assume a uniform one, but ideally we would use the expected distribution from our data of interest. The results are shown on Fig. H.2. They turn out to be essentially the same as in the “naive” approach,

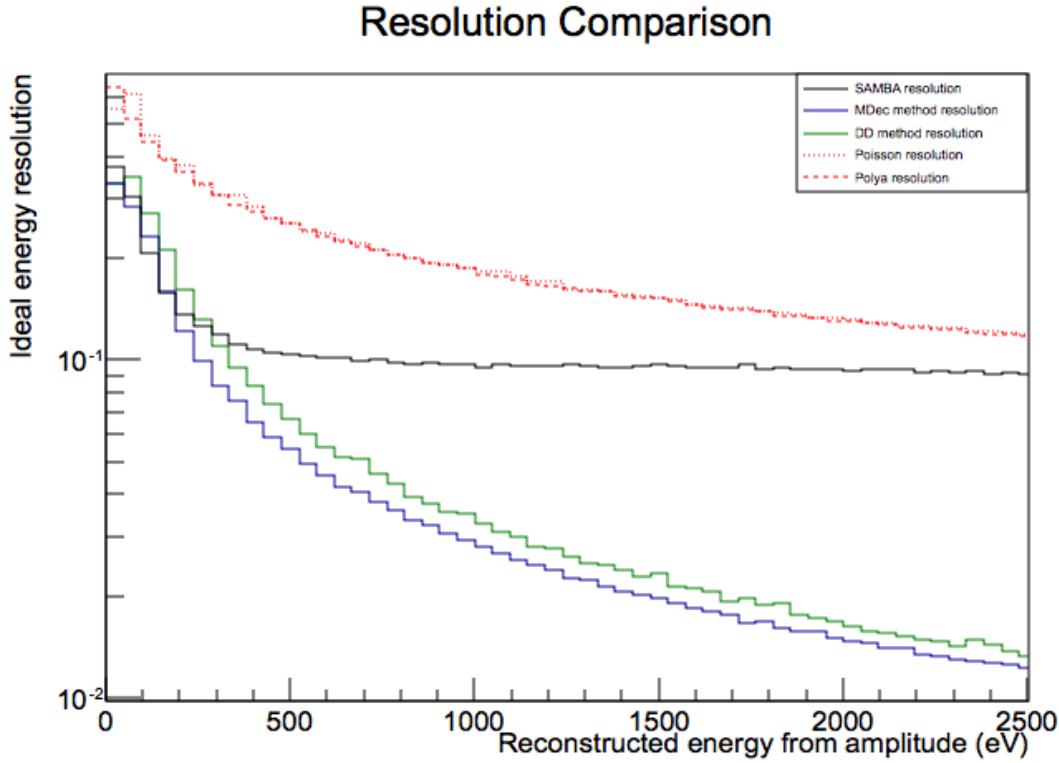


Figure H.2: Ideal secondary charge resolution, depending on the reconstructed energy from the amplitude for each method. The corresponding plots for the systematic uncertainties produced by primary electron (Poisson) and secondary charge (Polya) creation are shown again for comparison.

because of the linearity of our estimator with respect to the quantity estimated, and our choice of a uniform energy distribution. The only differences are a marginally worse resolution for SAMBA, and a marginally better resolution for all approaches at the lowest energies. The latter should be taken with a grain of salt, since the definition of *relative* resolution necessarily breaks down as the average of the value studied approaches 0.

In conclusion, we see that the MDec method has the best amplitude resolution across all ranges of energy, although the marginal improvement over DD disappears

once we take into account the systematic uncertainties of the detector. Unless interested in figuring out the avalanche gain process and we have some way to know the primary electron content of events (individually or statistically), the performance of both methods is essentially equivalent. In turn, SAMBA has similar performance to those two methods at very low energies (< 200 eV), where the ballistic deficit is not an issue anymore due to the low electron count. But beyond that point, its biases degrade its performance significantly.

H.2 Risetime resolution

The study of the risetime resolution is very similar to the previous one on the amplitude, with two differences. The first is that we have a preferred distribution of our events, which is a uniform distribution inside the detector. The second is that our estimators are no longer purely linear with respect to the quantity they are estimating. More specifically, the risetime of the pulse has a minimum value that depends on the processing used, so at low values it stops being sensitive to the real RMS of the arrival time of the primary electrons. This can be seen most clearly for the MDec method, as shown on Fig. H.3. This is primarily due to the response function of our detector to a single electron, which has a non-zero width, even after processing. Another effect contributing to this is any smoothing introduced during the processing, which also widens the shape of the pulse produced by a single electron.

With the reasonable assumption that the spread in electron arrival time roughly follow a gaussian distribution, the risetime from SAMBA and DD (time between 10 % and 90 % of the total amplitude), and the one from MDec (time between 10 % and 75 %) can be converted into an estimator of the electron arrival RMS by dividing by

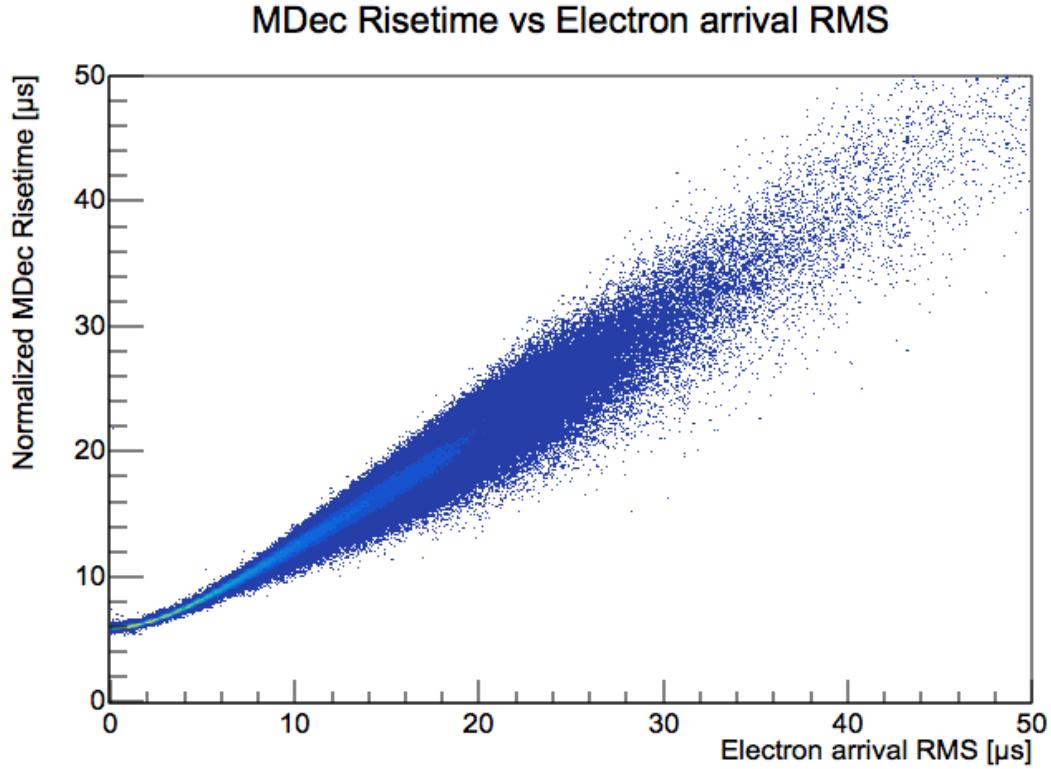


Figure H.3: Normalized MDec risetime vs the electron arrival time RMS. Note how the minimum risetime is $6 \mu\text{s}$, despite the RMS going all the way down to 0.

2.560 and 1.955, respectively. This allows us to plot the relative resolution of the real RMS, for a given estimated RMS. The results are shown on Fig. H.4.

The conclusions to be drawn from the plot are that while both MDec and DD offer the same performance for larger values of the electron arrival RMS, the minimum dispersion discernible by MDec is quite a bit larger than that for DD. This makes the DD method much more adapted for separating events from each other, as required for the search of KK axion events. In a similar vein, it will also produce better results when trying to discriminate events with a single electron (which always have an arrival RMS of 0) from other events, since multiple electrons have to arrive closer together

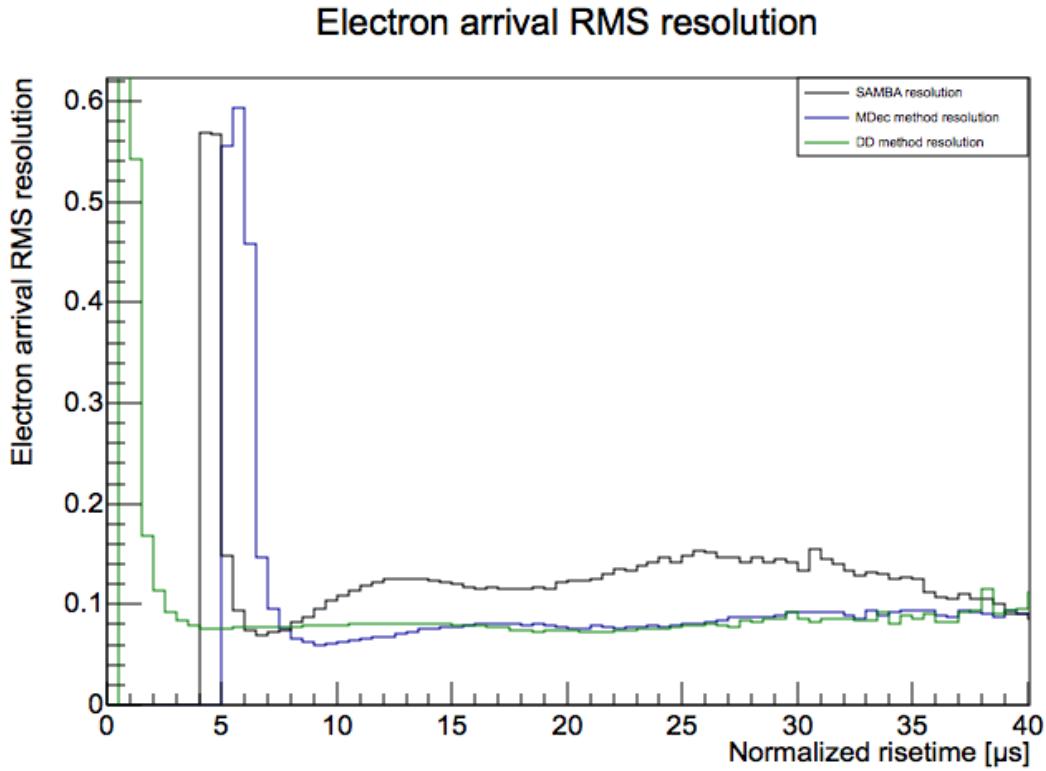


Figure H.4: Electron arrival RMS resolution, depending on the normalized risetime for each method. The point at which the resolution shoots up for each method corresponds to the point where the risetime loses its discrimination power due to approaching its minimum value.

before they can't be distinguished; this is useful in searches for very low energy events, since they always produce a single electron. As for the SAMBA estimator, while it technically has a moderately smaller minimum risetime than the MDec method, its resolution is considerably worse for almost all values of the risetime.

Appendix I

Parameter optimization for MPA

A relevant question for any algorithm is whether its internal parameters have been chosen for optimal results. In the case of the Multi-Pulse Analysis described in Sec. 4.2, we see in particular that the rate of false positives strongly depends on the chosen processing parameters, potentially by orders of magnitude, as shown on Fig. I.1.

To find the optimal parameters for MPA, different values for the smoothing strength and threshold were used. The sensitivity to axions and the false positive rate from pointlike events was computed for each smoothing/threshold couple. The results are shown on Fig. I.2. The first conclusion is that the higher the threshold chosen, the better the algorithm is capable of separating pulses, but the more false positives are also created.

To chose the optimal parameters based on both metrics, a simplified form of an exclusion limit was used. We assume 10000 pointlike events uniformly distributed in the 2 to 20 keV range, and 0 observed KK axion events, in the conditions of the physics run of the SEDINE detector. The expected rate of axion events (usually equivalent to less than 1 for our running conditions) is divided by the 90 % excluded average

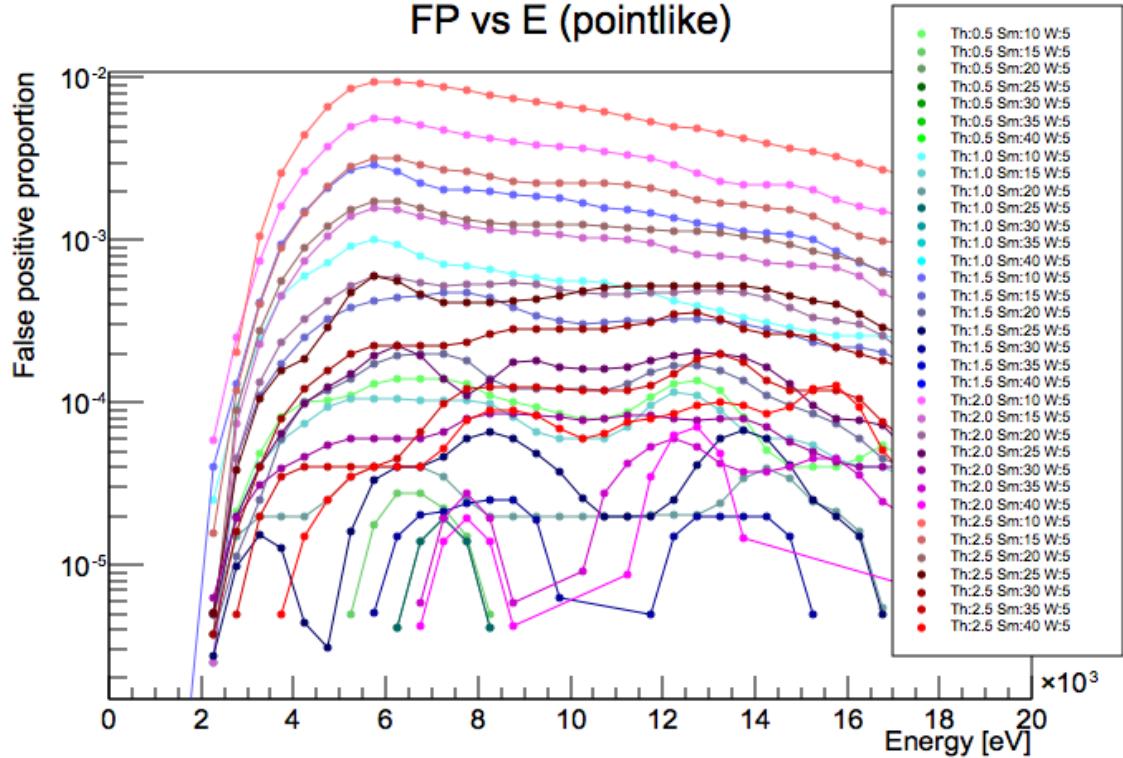


Figure I.1: Proportion of false positives out of uniformly distributed pointlike events, for different values of the processing parameters. Some mild smoothing was applied, to increase the readability of the plot.

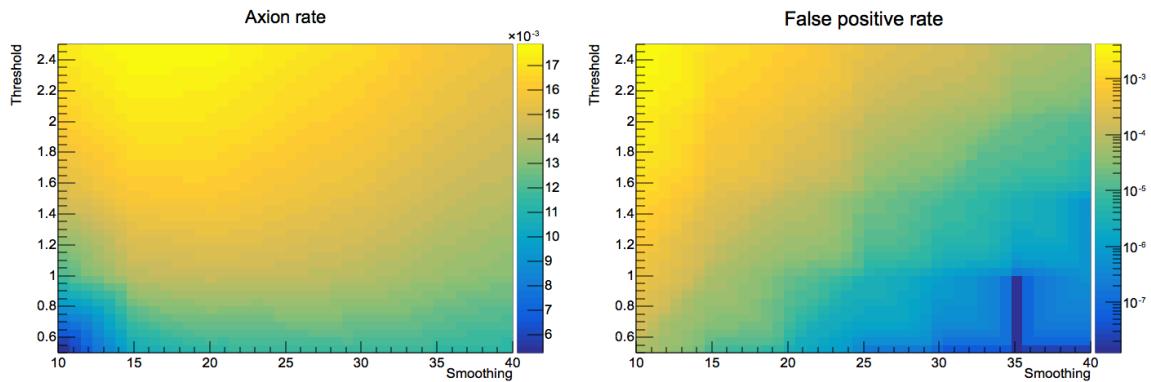


Figure I.2: Left: Axion rate, depending on processing parameters. Right: Proportion of pointlike events that produce a False Positive, depending on processing parameters.

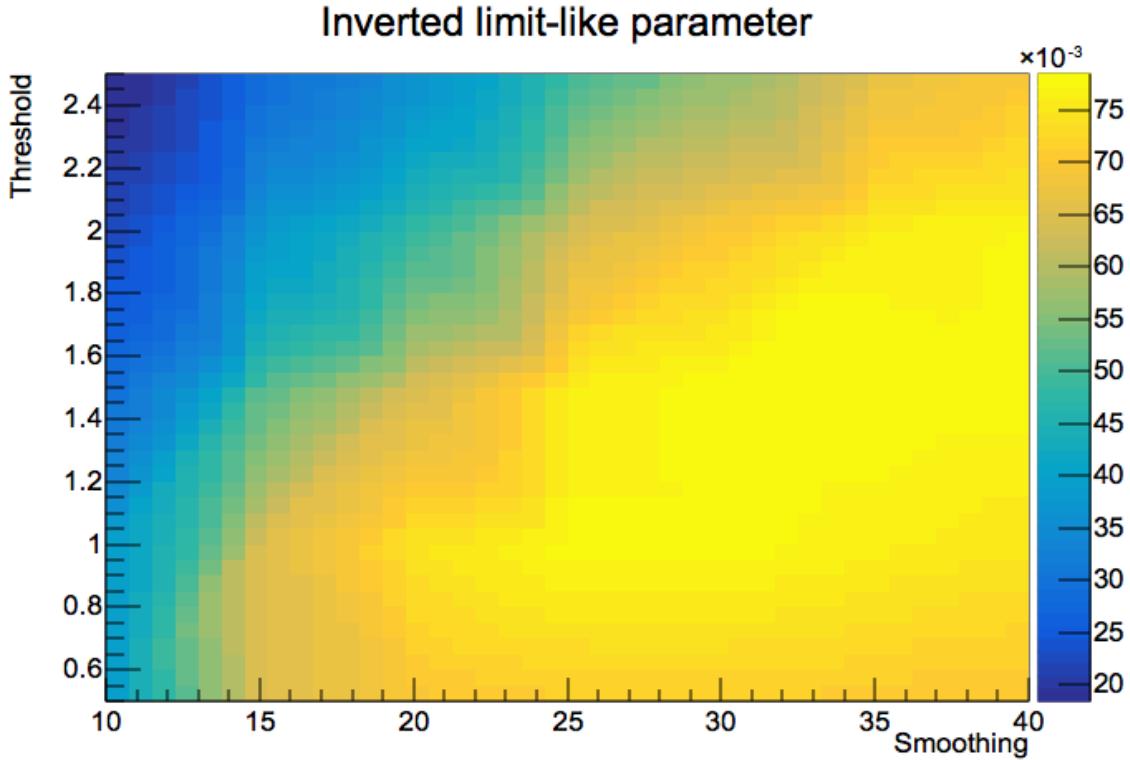


Figure I.3: Optimization metric, depending on processing parameters. The optimal parameters are taken to be a smoothing strength of 30 samples, and a threshold of 1.3.

number of events based on the number of events seen (all false-positive pointlikes). This parameter is then proportional to the inverse of the square of the exclusion limit (cf. Sec. 8.2.1), so it makes for an adapted metric to optimize the processing. The results are shown on Fig. I.3.

We can further compare the expected results from the default parameters, to the optimal ones. In fact, Fig. ?? shows that not only the optimal parameters represent a marginal increase in terms of sensitivity to axions ($\sim +5\%$), but also a sizeable decrease in the proportion of false positive events ($\sim -45\%$). About the latter, note that only 37 events were falsely reconstructed as double, out of $2 \cdot 10^6$ simulated events,

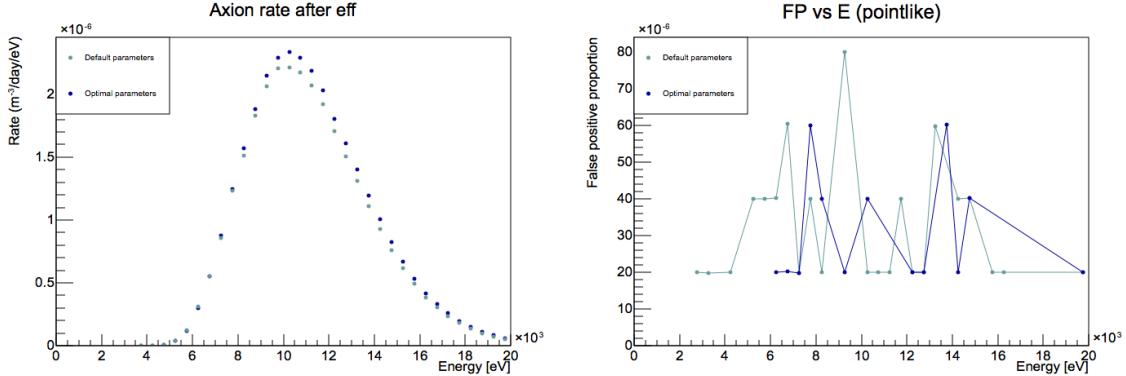


Figure I.4: Left: Comparison of Axion rate, for default and optimized parameters. Right: Comparison of False Positive proportion of pointlike events, for default and optimized parameters.

so it should be taken with caution due to the relatively low statistics. Nevertheless, putting both together, and assuming once again 10000 pointlike events, we expect an improvement of $\sim 10\%$ of our solar KK axion limit, compared to the default parameter choice. This is only true if our dominant background is pointlike events improperly reconstructed as being double; if instead it turns out to be some source of physical events that reproduce double interactions, then the limit may only improve by $\sim +2.5\%$.

Appendix J

COMSOL simulations

The choice of central electrode for an SPC has far-ranging effects, from its stability in time to the level of isotropy of the field it generates, but also in how it affects the drift of primary electrons within the bulk of the detector, or the maximum voltage that can be applied before continuous sparking develops. Multiple types of electrodes have been developed to try to fix some of the issues with earlier versions, like the simple ball electrode, or the ball and umbrella one. In this appendix, COMSOL simulations are performed for two newer models of the central electrode: the grid sensor, and the achinos.

J.1 Grid sensor

To try and solve the issue of the anisotropy of the gain and drift time inside the detector, another prototype of sensor included a “grid” of wires around a central electrode, as show in Fig. J.1. The rationale for this type of sensor is that, close to the central electrode, any anisotropies due to the rod will be shielded by the grid; conversely, in the drift region of the detector, the larger size of the grid compared to the central electrode will soften the effect of the grounded rod on the field.

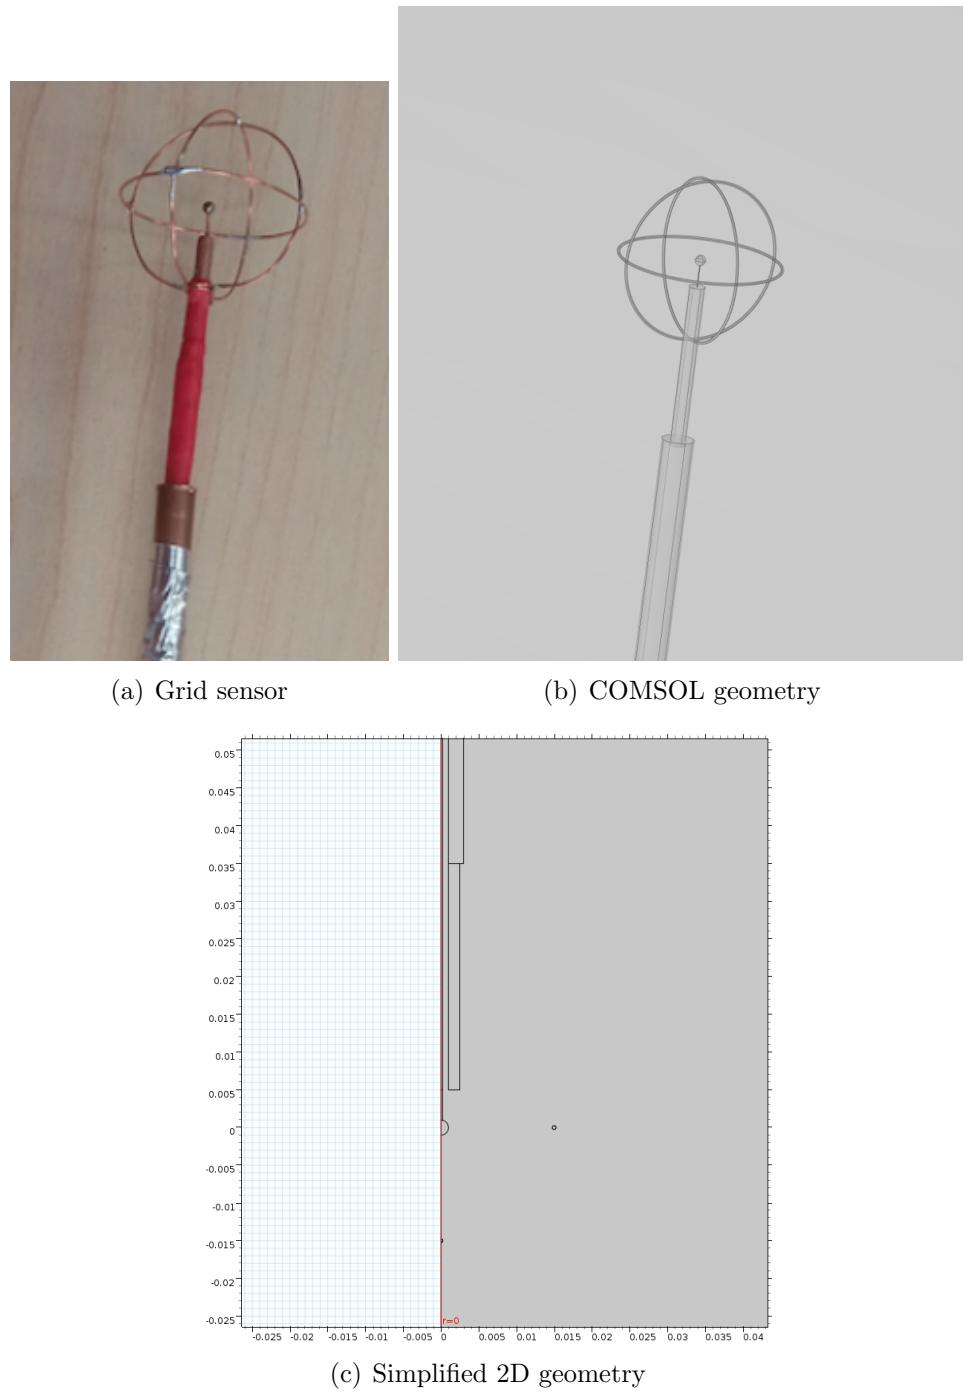


Figure J.1

As is immediately visible from the picture of the sensor, this is not cylindrically symmetric. While I did perform 3D simulations of the setup, simplified cylindrically symmetric configurations lent themselves much better to studying the effects of changing the applied electric potential on the grid. Specifically, the focus was on studying the effect of the grid potential on the dead region of the detector. Intuitively, we would expect a higher voltage on the grid to lead to an increase in the volume of the detector with field lines ending on the grid, and hence experiencing no avalanche.

The experimental setup used had the central sensor at 1200 V. The grid was kept at 70 V, since that is roughly the value of the potential that we would expect at the radius of the grid in the ideal scenario. The isotropy of the field was also tested by comparing the electric field as a function of radius for different angles. The results are shown on Fig. J.2. From the plots, it is clear that changing the potential on the grid has a large effect on the size of the dead region inside the detector, with less and less volume being visible as the potential increases. The second conclusion is that the electric field in the hemisphere away from the rod is pretty isotropic, no matter the voltage chosen (a difference of 20% maximum in field strength at the edge of the detector). On the other hand, the grid is not enough to shield the contribution from the central sensor: the field strength increases by roughly a factor of 2 between a grid potential of 30 V and 120 V, while the theoretical field produced by a sensor with the radius of the grid and the same potential increases by a factor of 4 instead.

The conclusions drawn from this are that the grid sensor may be good for improving the amplitude resolution of the detector, but is unlikely to be useful to control the behaviour of the field in the drift region. This is because the field changes slower than the potential applied, and small changes in the potential applied have large effects on

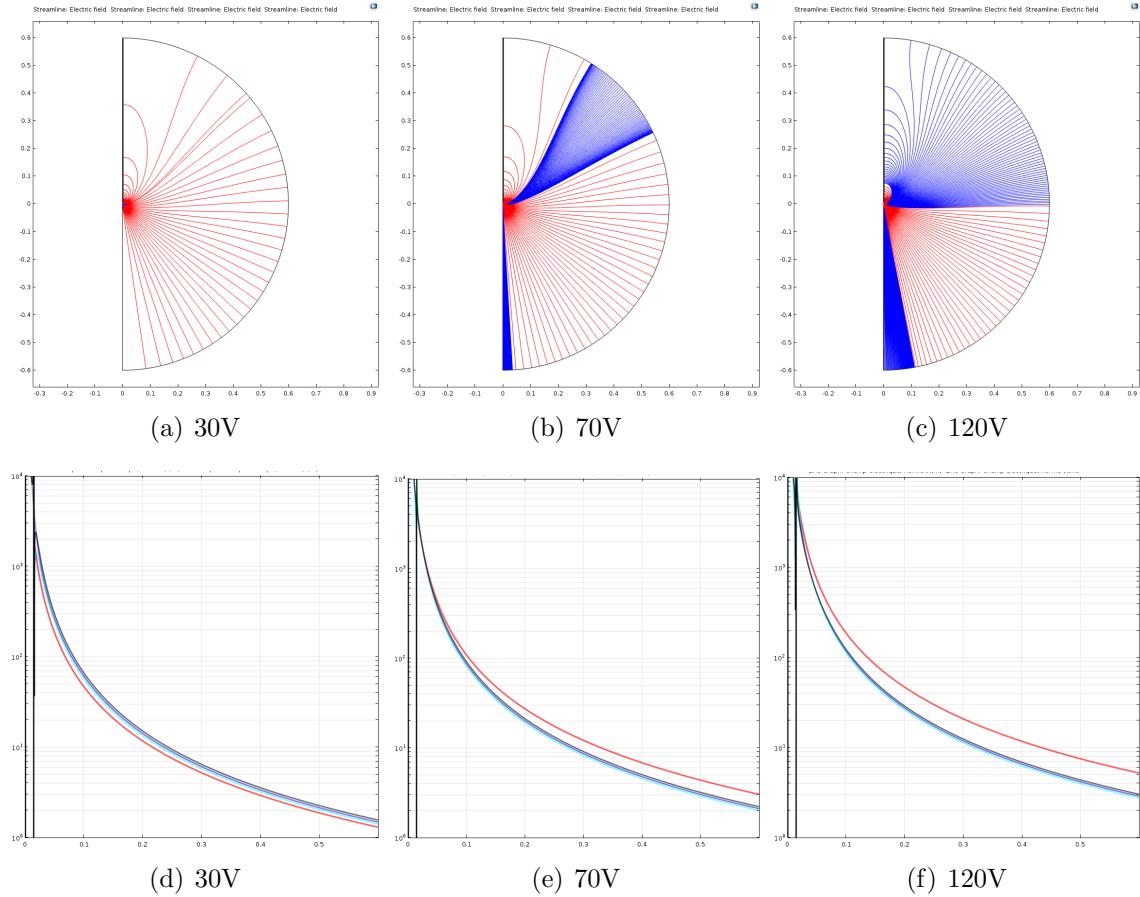


Figure J.2: Top: field lines in the detector. The field lines are red when ending on the central electrode, dark blue when ending on the grid. Bottom: field strength vs radius. The red function is the ideal field created by a sensor with the radius of the grid and the same applied voltage. The blue lines are (starting from the bottom-most) the simulated field strengths at 0, 22.5, 45 and 90 degrees below the horizontal; they differ by less than 20%, at the edge of the detector.

the strength of the avalanche, which we generally want to avoid. Furthermore, the introduction of relatively large dead regions in the detector, even larger than those shown in the plots (since this simulation only includes one of the three wires in the real sensor), will lead to a decrease in the effective event rate. Short of experimental data proving a vastly improved amplitude resolution for this kind of sensor, this does not seem like an upgrade from the simpler ones.

J.2 Achinos

The Achinos (greek for "sea urchin") sensor was initially developed to solve a limitation of the simpler ball sensors: the strength of the avalanche gain is coupled to the strength of the field in the drift region. In practice, if we want to increase the field in the drift region while keeping the maximum voltage constant (to avoid sparking), we have to use a larger sensor, since the electric field varies roughly as $E(r) \sim Vr_{sensor}/r^2$. However, this will decrease the electric field close to the sensor, since that roughly varies as $E(r_{sensor}) \sim V/r_{sensor}^2$, so the avalanche may become too weak to observe the events of interest.

The solution to the problem was to use a sensor geometry where the "size" of the sensor "depend" on the distance from which we look at it. By using small electrodes kept some distance away from the centre, an electron drifting towards the sensor will "see" the larger radius at which the electrodes are kept while drifting. But, once it reaches the sensor, it will produce an avalanche according to the smaller radius of these electrodes. By choosing the right radii for the electrodes and their position with respect to the centre of the detector, we can get a good enough amplification for the avalanche, while keeping a drift time that is fast enough for the primary electrons not

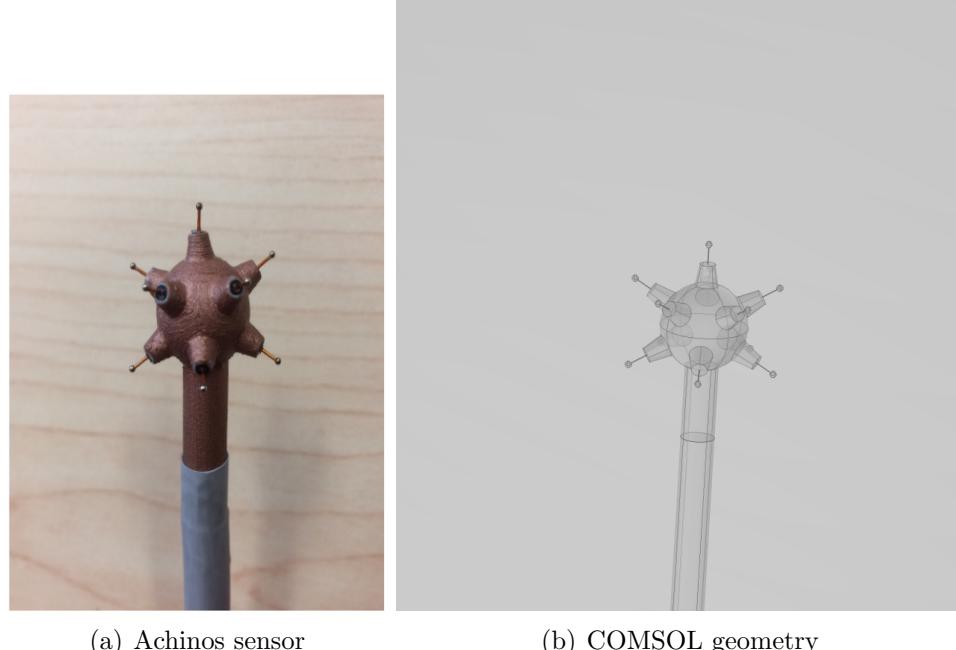


Figure J.3

to get captured before reaching the sensor, but not so fast that we lose discrimination power against surface events.

I performed a 3D simulation of the detector with COMSOL to get qualitative answers about any dead region in the detector. Another check was to see how many field lines end up in any of the rod-side electrodes, compared to the bottom electrodes. Somewhat unexpectedly, the hardest part of finding the answer to these questions was to find a good way to represent 3D results as a 2D plot.

In Fig. J.4, some plots showing the field configuration are shown. For context, experimentally we observed that applying the same voltage to all electrodes lead to only the electrodes farther away from the rod recording any signal. If the voltage on the electrodes closest to the rod was increased a moderate amount, signals could be recorded everywhere. If that voltage was increased even further, then it was only the

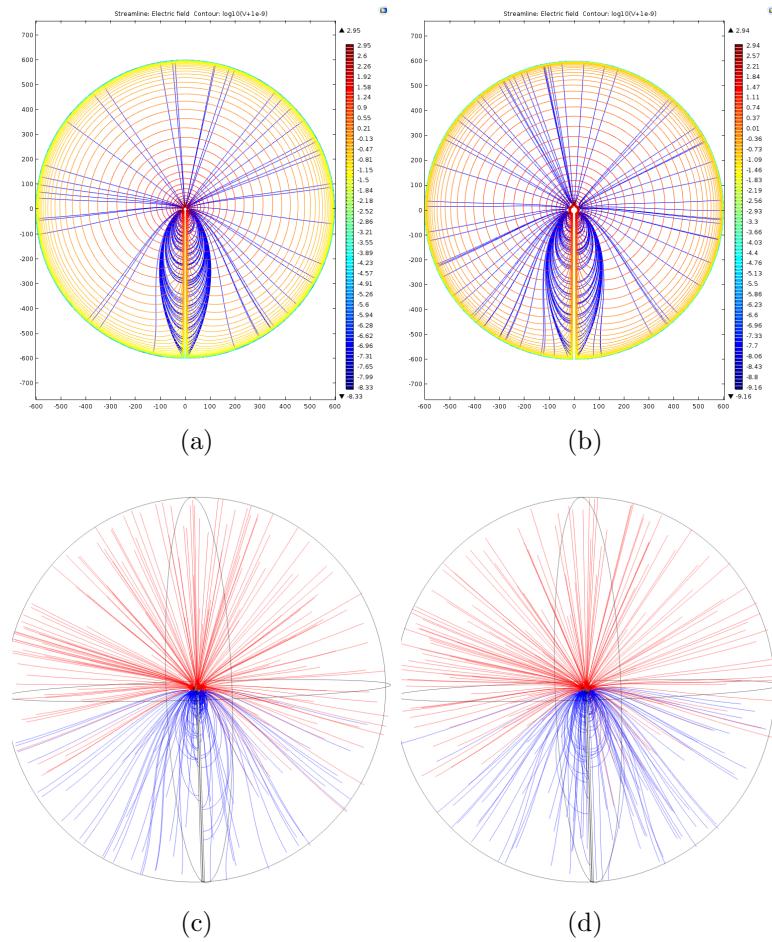


Figure J.4: Top: 2D projections of isopotentials and field lines. Bottom: 3D field lines, blue for lines ending on the rod-side electrodes, red for lines ending on the other ones. The plots on the left have 1000 V applied on all electrodes, the ones on the right have 1080 V on rod-side electrodes and 1000 V on others. Despite the large observed difference in the avalanche for either configuration, the field lines are barely affected

electrodes closest to the rod that recorded a signal. A relative increase of 8% on the voltage of the rod-side electrodes gave the most balanced results.

A possible explanation for this behaviour was that maybe most of the volume of the detector had field lines that ended on either the top or the bottom electrodes. This was quickly rejected with the COMSOL simulations: changing the relative voltage of the electrodes had barely any effect on the field configuration in the drift region of the detector. Furthermore, almost no area of the detector had field lines ending on a “dead” region of the sensor. Whatever the explanation for the discrepancy between the electrodes, it is unlikely to be due to the field configuration in the drift region. An alternative explanation is that the relative proximity of the sensors has a large effect on the strength of the avalanche, and so it takes a precise voltage setup to have all electrodes experience a large enough avalanche while in the same configuration.

The field close to each electrode was also observed to check whether we do expect the same behaviour for each one of the achinos electrodes than for a simple sensor with a single electrode. The results can be seen on Fig. J.5. For the fieldlines coming from the drift region that reach the electrode, the electric field away from the electrode is remarkably close to the ideal field (just $\sim 10\%$ smaller), up to $\sim 4\text{ mm}$ away. This corresponds to a voltage drop of $85 - 90\%$ from the maximum, which means the ion induced current will match the ideal one for that same percentage of the total amplitude, and only diverge for the final $10 - 15\%$. For the specific case of SNOGLOBE, with 1 bar of Neon with 5% CH_4 , and an applied voltage of 1630 V (the conditions tested at the time of writing this section), this means the first $\sim 100\mu\text{s}$ of signal match the ideal scenario; the difference in shape for the tail after that will lead to undershoots or overshoots after processing, depending on whether the

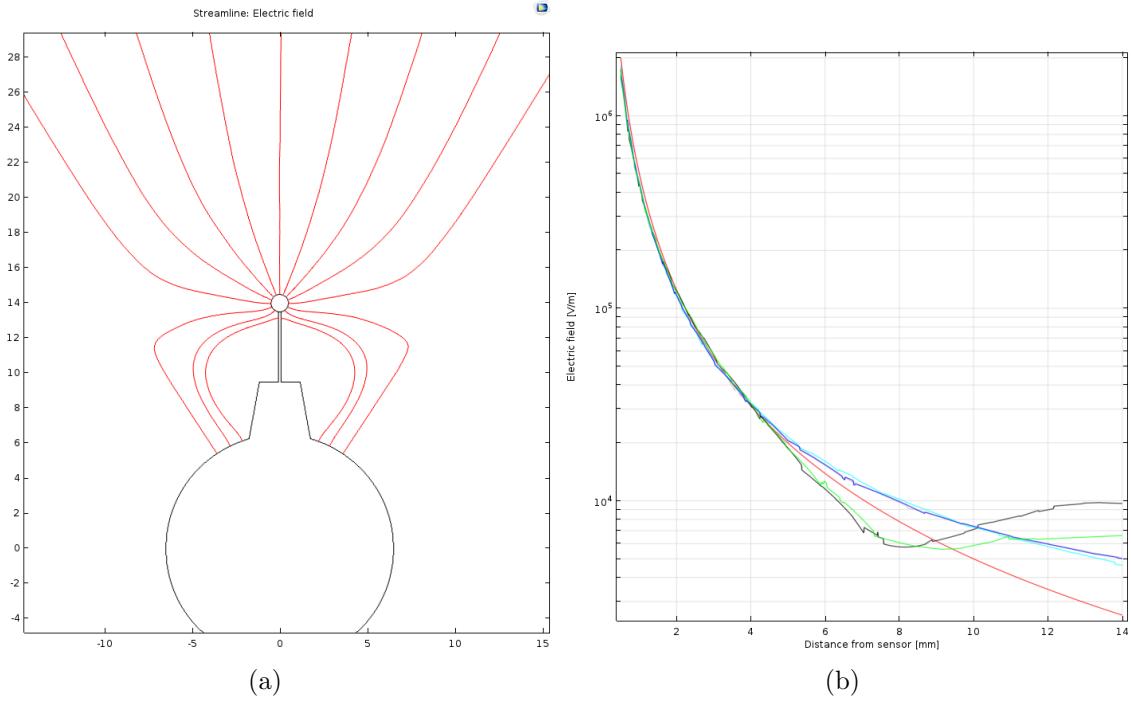


Figure J.5: Left: field lines close to the electrode furthest away from the rod. Right: in red, ideal field for a single electrode inside the detector; in black, gray, cyan, blue and green, the field strength away from the electrode, at +90, +45, +0, -45 and -90 degrees respectively from the line away from the centre of the sensor. The +90 degree line points towards another electrode, the -90 degree line points in between two sensors.

primary electrons reach the electrode straight on or obliquely, respectively. As such, the approximation of the field close to the electrode being that of an isolated electrode is good enough for processing data, or for producing qualitative simulations. But if we are interested in simulating data where the pulse tails are important, the ideal ion induced current won't do.

One consequence of the electric field around the electrodes matching the ideal one so closely is that it implies the total charge on each electrode is the same as the charge in the ideal scenario. Far away from the sensor, the field should behave like

the one from one central charge (the combined charge of all electrodes), due to the approximate spherical symmetry of the detector and Gauss's law. This is shown to only be partially true with COMSOL, as shown in Fig. J.6. While the electric field does roughly follow a $1/r^2$ behaviour, especially farther than 30 cm away from the centre, we need to add a corrective factor of 0.17 to the ideal electric field to match the actual values found with COMSOL. Remember that the charge of a spherical conductor of radius r held at a potential V inside a hollow conductor of much larger radius is $Q \approx 4\pi\epsilon Vr$. So the total charge from 11 electrodes of radius 0.5 mm should be $\sim 0.4 Q_{ideal}$, where Q_{ideal} is the charge from a single electrode of radius that of the Achinos, 14 mm, and the strength of the electric field in the drift region should follow the same ratio. If we get 0.17 times the electric field instead, that means that the total charge in the achinos sensor must be $0.17 Q_{ideal}$. But, since the field close to each of the achinos electrodes follows the ideal field, this means that the support of the achinos, which is kept grounded, must hold a charge of $Q_{support} = -0.23 Q_{ideal}$.

I do not know of a way to get an estimate of the charge on the support analytically, so other Achinos configurations would have to be run through COMSOL to find it. The importance of this computation is that it has a large effect on the drift field; indeed, in the configuration described above, the field is only $\sim 44\%$ of the field produced by a sensor with 11 floating electrodes, and $\sim 17\%$ of the naive approximation with a single electrode of radius that of the achinos. This is still a sizeable improvement over a single electrode of radius 0.5 mm, which would only produce 3.6% of the naive approximation field. However, the accumulation of charge on the support to compensate the field from the electrodes will limit how much the field close to the electrodes (which governs the avalanche process) can be decoupled from

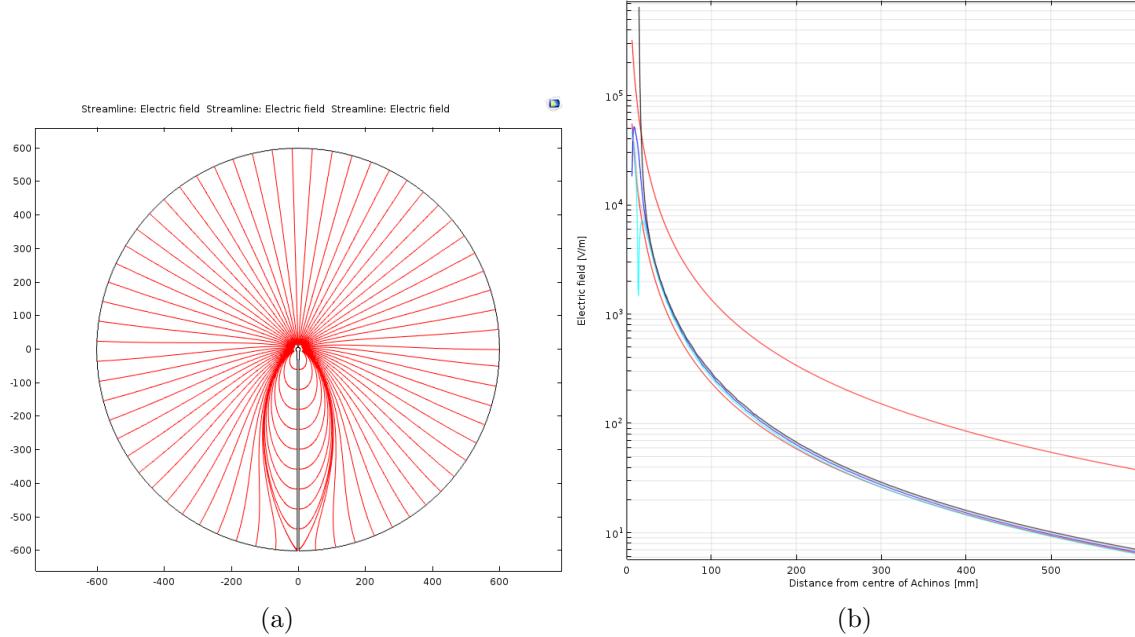


Figure J.6: Left: field lines in the drift region. Right: in red, ideal field for a single electrode with radius that of the achinos sensor, and the same ideal field multiplied by 0.17; in cyan, blue, gray and black, the field strength away from the sensor, at 0, -22.5, -45 and -90 degrees respectively from the equator (all pointing away from the rod).

the field in the drift region (which governs the drift and diffusion time, and indirectly the attachment).

Appendix K

Systematics of SEDINE axion searches

One concern with the calculation of exclusion limit is that it depends on two factors: the observed number of axion-like events in the data, and the expected sensitivity of the detector to KK axion events. Once the processing and cuts are fixed, independently of how those cuts are determined, the first of those two factors is fixed too. However, the expected sensitivity of the detector depends on the quality of our simulations. Indeed, if our simulations are wrong, then so is our computed sensitivity.

To include this effect in the final exclusion limit, we need study the systematics of the method. This means that, for every parameter that may be expected to have an impact on the final result, we determine how well that parameter is known, as already shown on the calibrations chapter of this thesis, Chap. ???. Then, by running more simulations where the parameter is changed by a factor that is consistent with its measurement uncertainty, we can establish its effect on the final exclusion limit.

This has to be run with all parameters left to vary independently, but for the sake of understanding what we are most sensitive to, we will first show the results when sweeping only one parameter at a time.

K.1 Control simulation

For the sake of verifying that the following studies have enough statistics, we first ran 200 simulations of 10k axion-like events each, where all parameters were fixed to their expected values (i.e., no parameter was swept nor left to vary randomly). These simulations were then processed, and the value of the total efficiency was extracted from each. Small deviations from the mean will lead to half as large variations on the final relative uncertainty on the exclusion limit, due to the square root term in the formula for the limit. In the specific case of the control simulation, this gives a measure of the statistical uncertainty of the approach, for this number of events per simulation. It also provides a point of comparison for the effects of changing each parameter separately. As per Fig. K.1, we see that the relative dispersion σ/μ of the detector efficiency to KK axion events is 2%, well within reason. While the statistics of this test could be increased to reduce this error, we will soon show that this is far from the dominant source of uncertainty on the final exclusion limit.

K.2 Drift time

The drift time of electrons in the detector have a strong effect on the efficiency. Indeed, the longer the drift times involved, the easier it is for the processing to separate different locations in the detector, as shown on Sec. 5.6.2. We leave the maximum drift time to vary with a standard deviation of 5%, based on our calibrations. The resulting spread in efficiency is shown in Fig. K.2. Despite the previously stated difficulties with calibrating drift times without the use of a laser, the relative spread in efficiency is of only 9.5%.

On the other hand, another effect adds to that. As explained in Sec. 6.1.4, we only

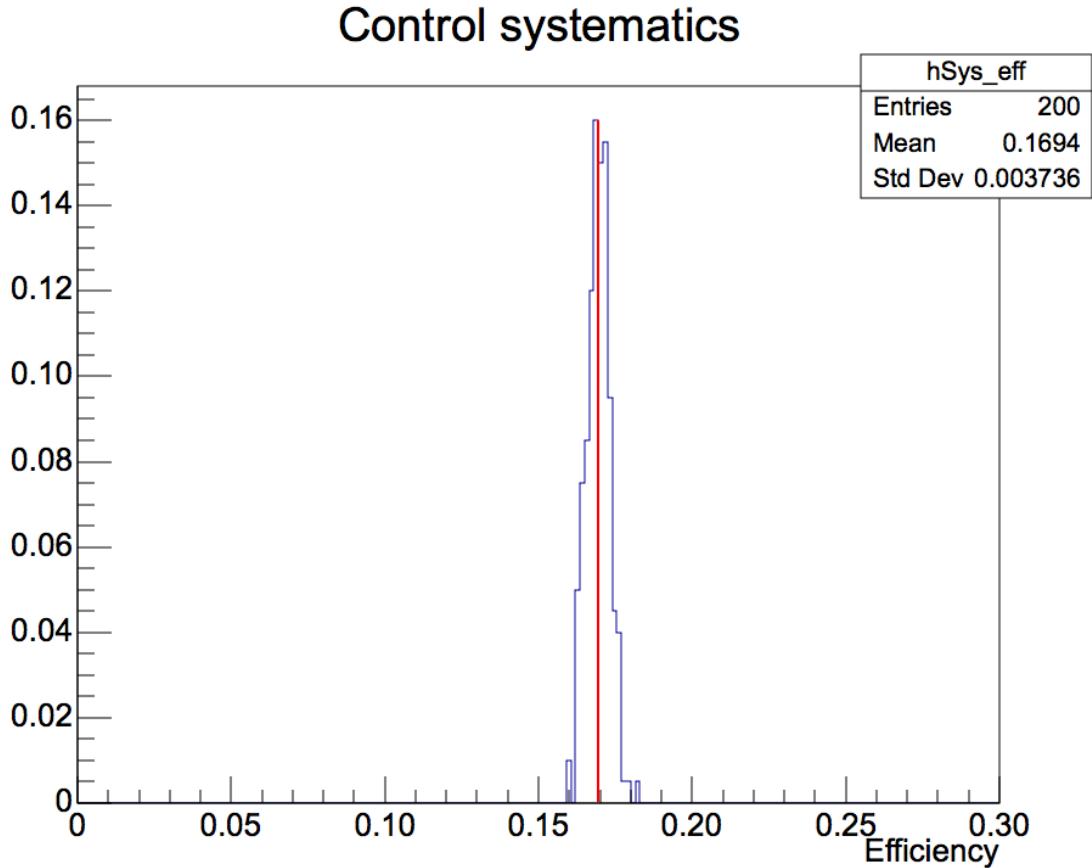


Figure K.1: Dispersion on the detector efficiency from 200 simulations of 10k events each, with all parameters fixed, and assuming a local density of KK axions of $4.07 \cdot 10^{13} \text{ m}^{-3}$. For reference, the red line is the mean for this simulation. The relative dispersion of the efficiency is 2%.

have a calibration for the drift time of surface events, leaving intermediate drifts to be determined solely through simulations. We can still make an estimate of the error introduced by this issue, using a toy model for the drift time. The relationship between drift time can be approximated by $T_{drift} = T_{drift_{max}} \left(\frac{r}{r_{max}} \right)^\alpha$, with $T_{drift_{max}} = 470 \mu\text{s}$ taken from calibrations, $r_{max} = 30 \text{ cm}$ the size of the detector, and $\alpha = 2.5$ the value obtained by fitting results from simulations. By letting α change, while keeping the conditions for separability the same as in Sec. 5.6.2, we can get an estimate of the

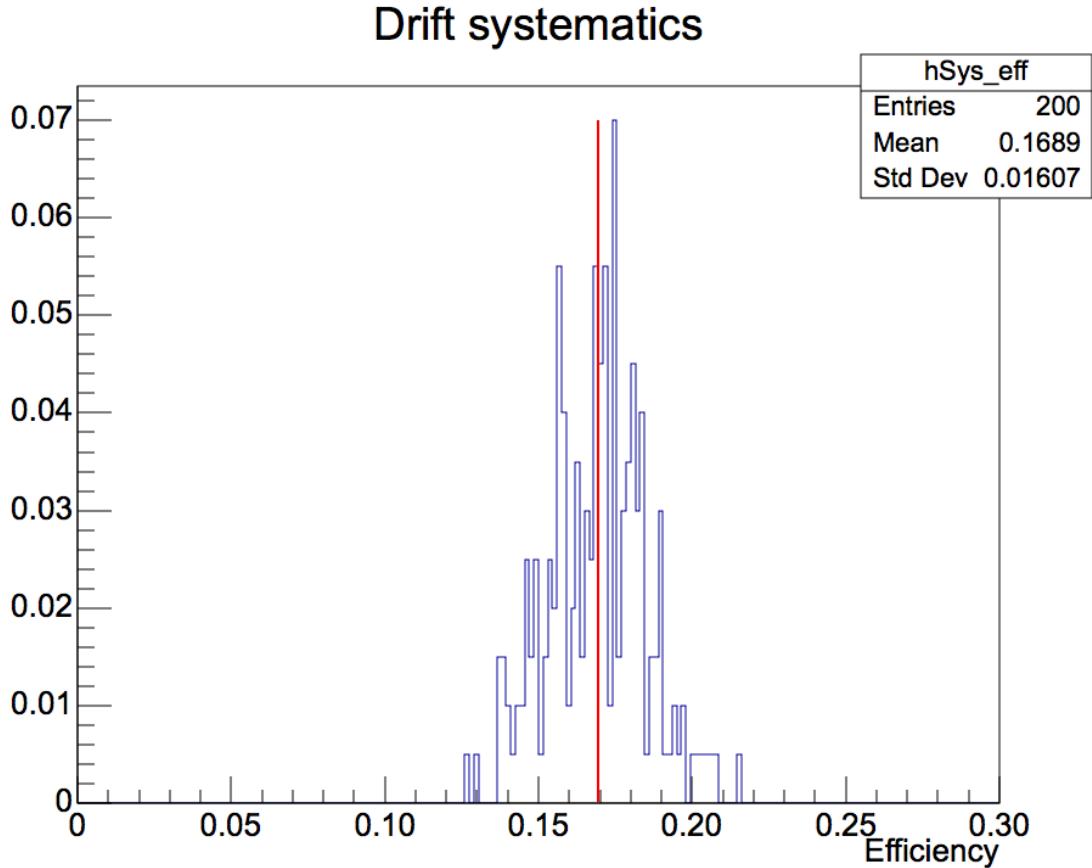


Figure K.2: Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the drift time is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 9.5%.

uncertainty on the expected axion event rate due to this effect. The results are shown in Fig. K.3.

The loss in separability when changing α in either direction can be explained by two conflicting effects. In the one hand, separability improves the more a small change in position induces a large change in drift time; the higher the α , the steeper the curve is in the high radius region (where most of the volume of the detector is). On the

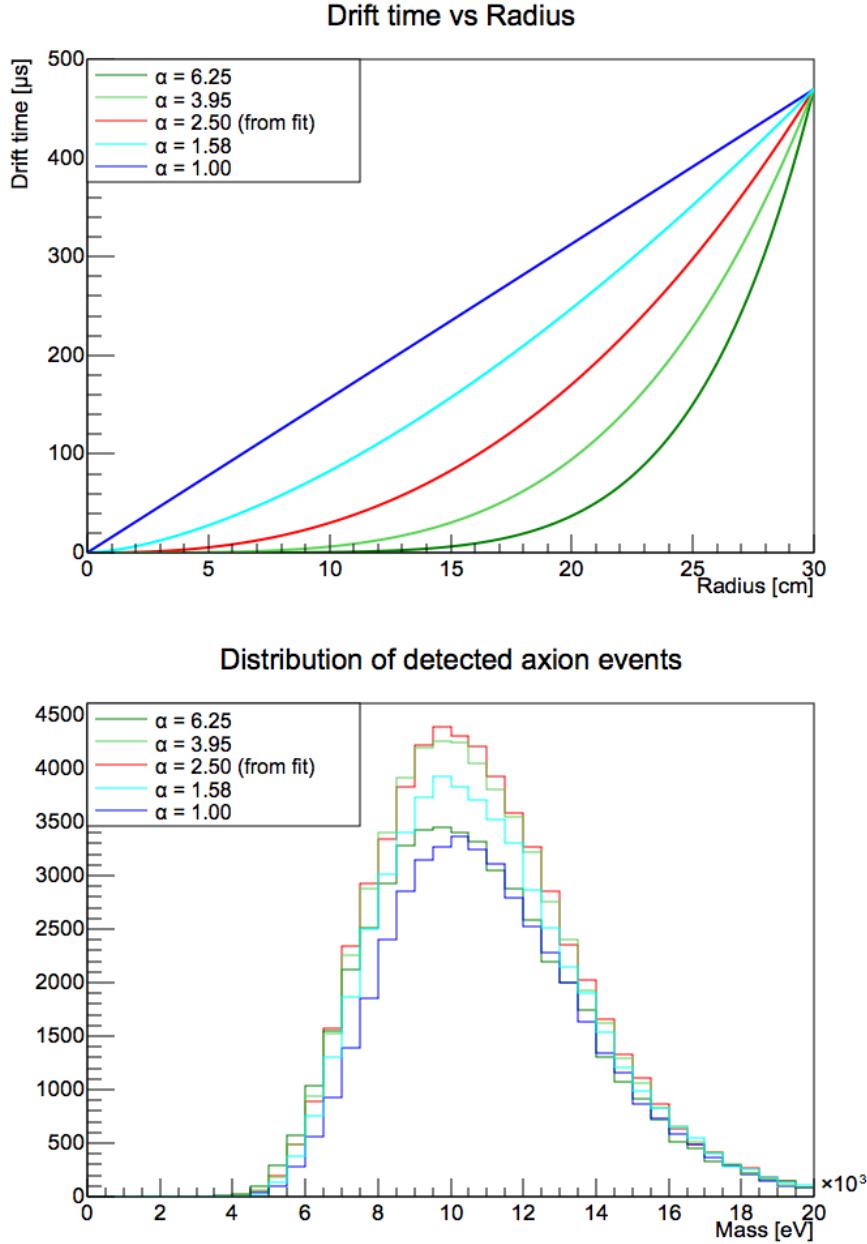


Figure K.3: Top: Toy models for the drift time vs radius relationship, with fixed maximum drift; $\alpha = 2.5$ corresponds to the fit of the results from simulations. Bottom: Simulated distribution of separable events in the detector for different values of α . The relative rates (in order of increasing α) are 76%, 89%, 100%, 98%, 83%.

other hand, the steeper the curve, the smaller the region of the detector where the curve is steep enough for events to be separable, since we are keeping the maximum drift time constant. According to this toy model, the ideal separability is likely for α somewhere between 2 and 4.

The values of α tested were chosen somewhat arbitrarily, between 1 and α_{fit}^2 . In practice, values of α below 2, or much above it, are unlikely, due to the observed distribution in the data of events at low and mid risetimes. An α below 2 would lead to a decreasing number of events with lower risetime, while an alpha above it would lead to the opposite effect. Yet what we see is a roughly uniform number of events at all risetimes, with a marked increase at very low risetimes (likely to minimum risetime effects from processing). While this only directly constrains diffusion times, they are closely related to drift speeds. Together with the results from simulations, we can constrain α between 1.5 and 4.0 with a reasonable amount of confidence.

Altogether, and to err on the side of caution, our lack of knowledge on the values of the drift time *in the bulk* of the detector means its sensitivity to KK axions is anywhere between 85% and 100% of the one computed without taking this effect into account. Its impact on the final exclusion limit is shown in Sec. 8.2.2, but can already be estimated to weaken the limit by $1/\sqrt{1 - 0.15/2} \simeq 1.04$ (i.e., increase the limit by $\lesssim 4\%$).

K.3 Diffusion time

The diffusion time of electrons also has an effect on the efficiency of the detector, if lesser than the drift time. Roughly, for two pulses to be separable, their difference in arrival time must be larger than their respective spreads in arrival times. However,

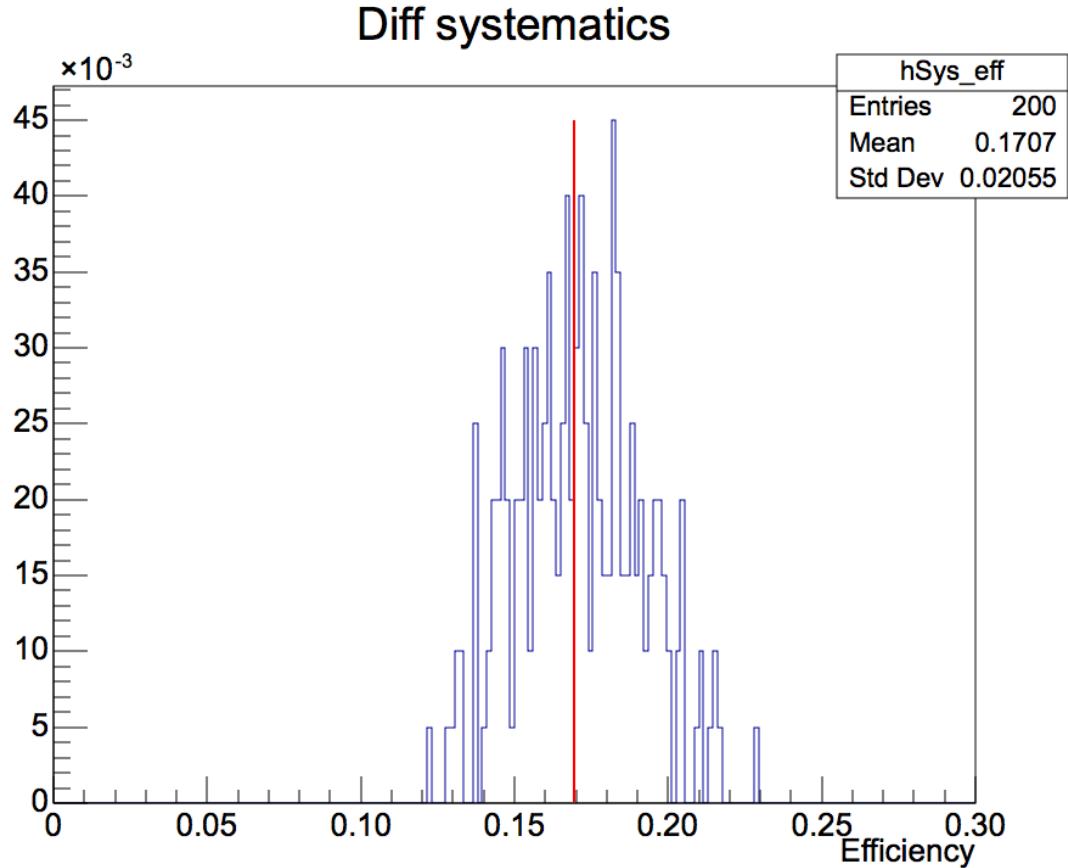


Figure K.4: Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the diffusion time is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 12%.

since diffusion times are roughly one order of magnitude smaller than drift times, the effect of uncertainties on the efficiency ought to be smaller. On the other hand, due to the improved risetime cut digging into some of the axion event rate, while the cut on time separation between pulses was left intentionally wide, a miscalibration of the diffusion time will affect the proportion of bulk events rejected. The two effects end up cancelling each other: the diffusion time calibration has a somewhat larger

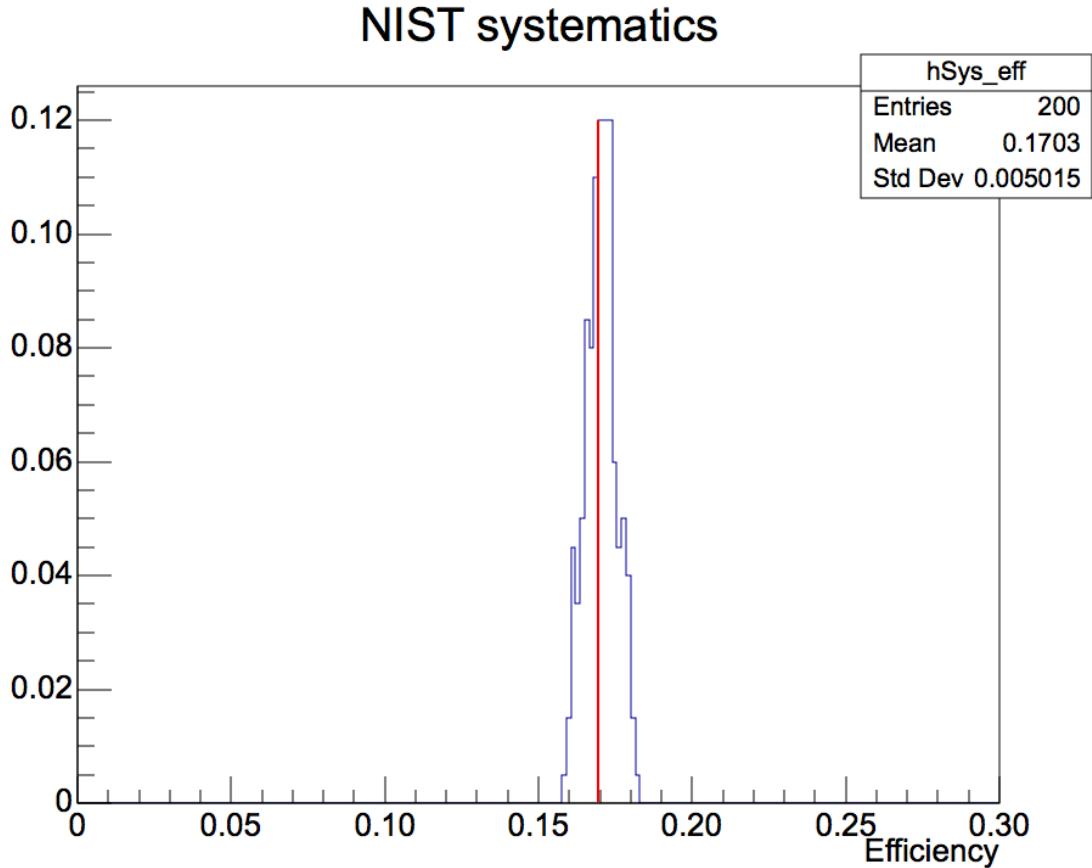


Figure K.5: Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the photon attenuation range is left to vary according to NIST data, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 3%.

uncertainty, of 6% (compared to 5% for the drift time), and an equally larger effect on the efficiency uncertainty, at 12% (compared to 9.5%). The distribution is shown in Fig. K.4.

K.4 X-ray attenuation range

The attenuation length of X-rays in gas was taken from the NIST database [63]. While the precision of their simulations reach $\sim 0.1\%$, we chose a more conservative uncertainty of 10 %, according to the report in [91]. This is unfortunately one piece of information we could not calibrate with our detectors.

We drew a normal distribution centred around the nominal attenuation length, with a relative resolution 10 %, and a minimum of 10 % of the nominal value, to avoid non-physical simulations. The results are shown in Fig. K.5. A relatively small effect of 3% is observed, only twice the size of the statistical uncertainty. This is likely due to the fact that the attenuation length of a photon depends strongly on its energy: any small differences due to uncertainties are going to be dwarfed by the variations due to energy changes.

K.5 Energy calibration

The energy (or, equivalently, mean avalanche gain) calibration is also one of our systematics. An error in this calibration will in turn lead to a cut in amplitude that won't match the cut in energy that we believe we are doing. This effect would be completely negligible for the default region of interest (2 – 22keV), since the expected number of events at the edges is negligible. However, for our optimized region of interest, this is not the case any more: a miscalibrated energy range will shift which events pass our cuts, potentially diminishing the sensitivity of our detector considerably. For the reported uncertainty of 5% on our energy calibration, we end up with an efficiency uncertainty of 3% (cf. Fig. K.6).

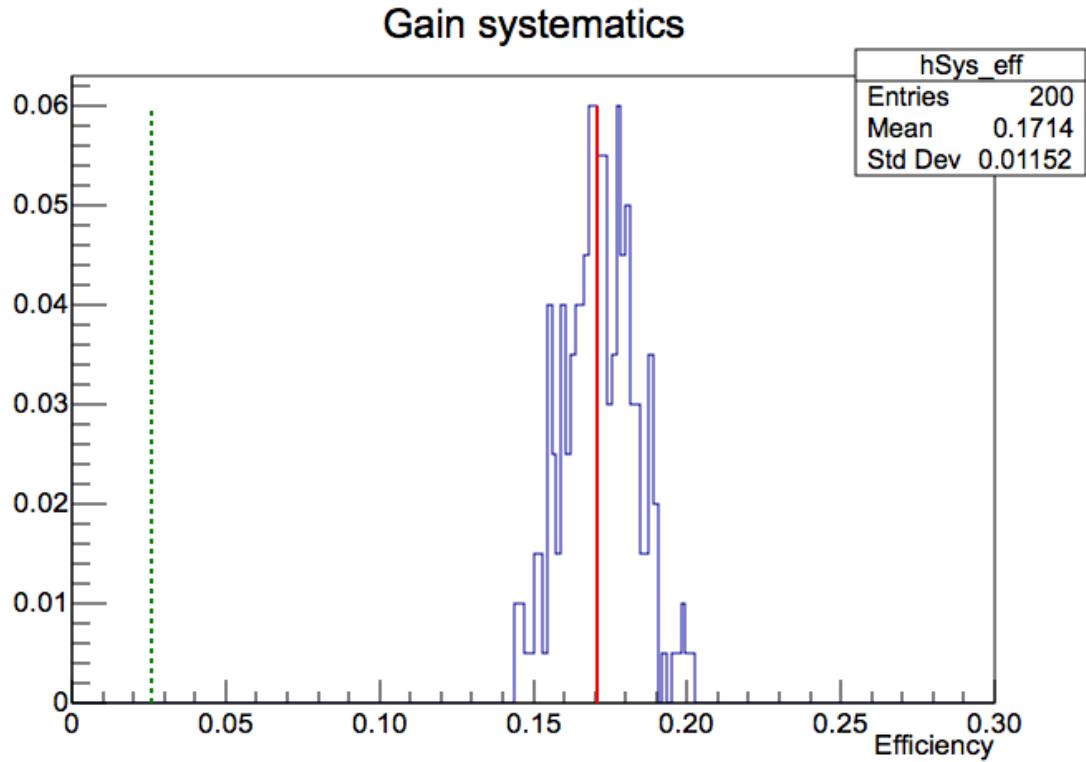


Figure K.6: Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the mean avalanche gain is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 3%.

K.6 Mean ionization energy

Naively, we would expect the mean ionization energy of Neon with 0.7% CH₄ to be the same as that of pure Neon, 36 eV [29]. However, recent calibrations of one of our Queen's prototypes revealed the mean ionization energy of Neon with 2.0% CH₄ to be 27.6 ± 0.2 eV, significantly closer to that of pure methane (27.7 eV) than that of pure Neon. Unfortunately, the dataset we used had no calibrations that we could use to determine the actual ionization energy in that gas mixture. Instead, we left

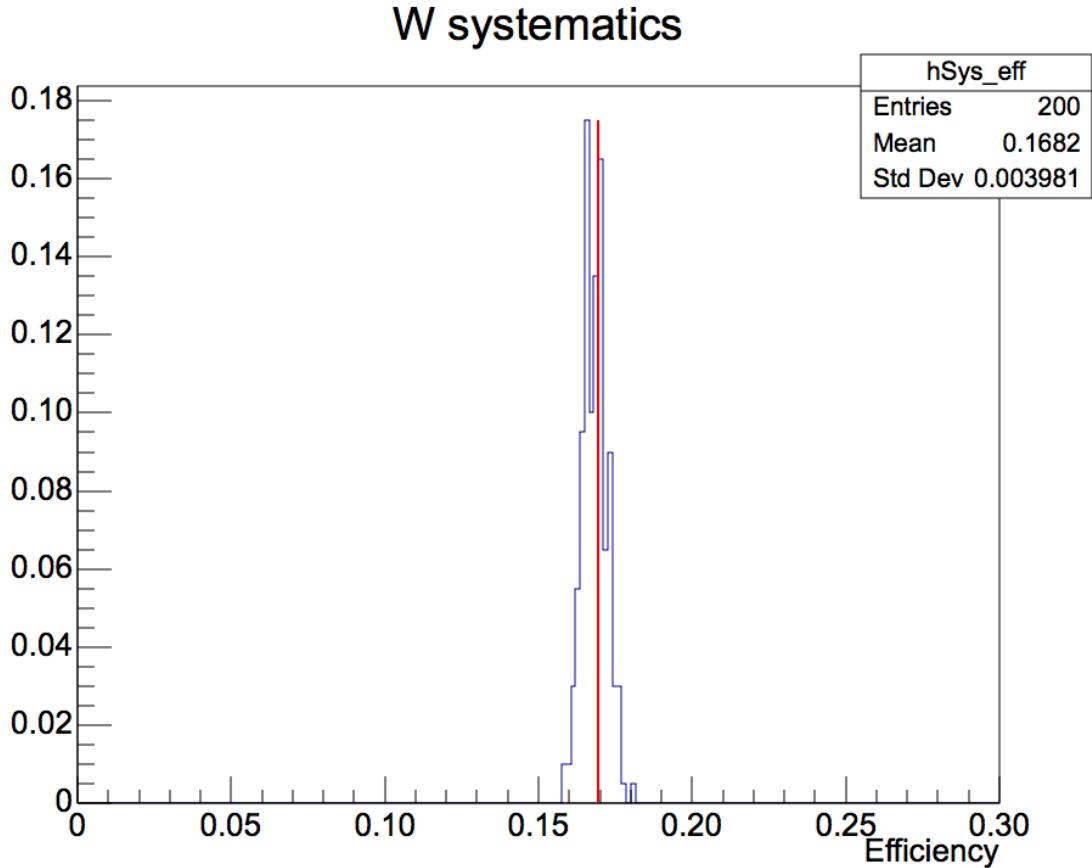


Figure K.7: Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the mean ionization energy is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 2%.

the value vary uniformly between the two extreme values (while keeping the energy calibration constant!), which should give a conservative estimate of its uncertainty on the exclusion limit. Fortunately, as shown on Fig. K.7, the effect is small enough to be dominated by the statistical uncertainty of this study. This is likely due to the fact that the lower energy values considered in this study, 1 keV, are still large compared to the mean ionization value. The discreteness of the number of electrons produced

is not a significant effect at this stage.

K.7 Attachment

As already mentioned in Sec. 8.1.2, the uncertainty on the electron attachment calibrations has a large effect on both the reconstructed energies of axion events, and their asymmetries. Choosing cuts without taking this effect into account would lead to very large systematics; for example, a naive energy range of $3 - 6.5\text{keV}$, optimal for basic axion-like cuts and the expected attachment, would induce a relative dispersion on the efficiency due to attachment uncertainties of 35%.

Fortunately, the energy and asymmetry cuts selected in Sec. 8.1.2 already took this uncertainty into account, opting to maximize the efficiency in the *worst* case scenario, rather than the *expected* one. This decreased the systematics that attachment uncertainties induce down to 5%, as shown on Fig. K.8.

K.8 Ion mobility

When processing the pulses, one of the main steps is to deconvolve the response of the detector from the signal recorded. The main unknown in the process is the mobility of ions in the gas, since all other parameters are either fixed experimentally or can be measured. Indeed, while the literature suggests a mobility of Neon ions in pure Neon at atmospheric pressure to be around $4\text{ cm}^2/\text{V/s bar}$, our calibration gives a value in the $6 - 8\text{ cm}^2/\text{V/s bar}$ range, depending on running conditions. This difference could be due to a combination of effects (increased methane concentration, anisotropies of electric field absorbed by mobility parameter).

An error in the estimated mobility is not an issue for the computation of the

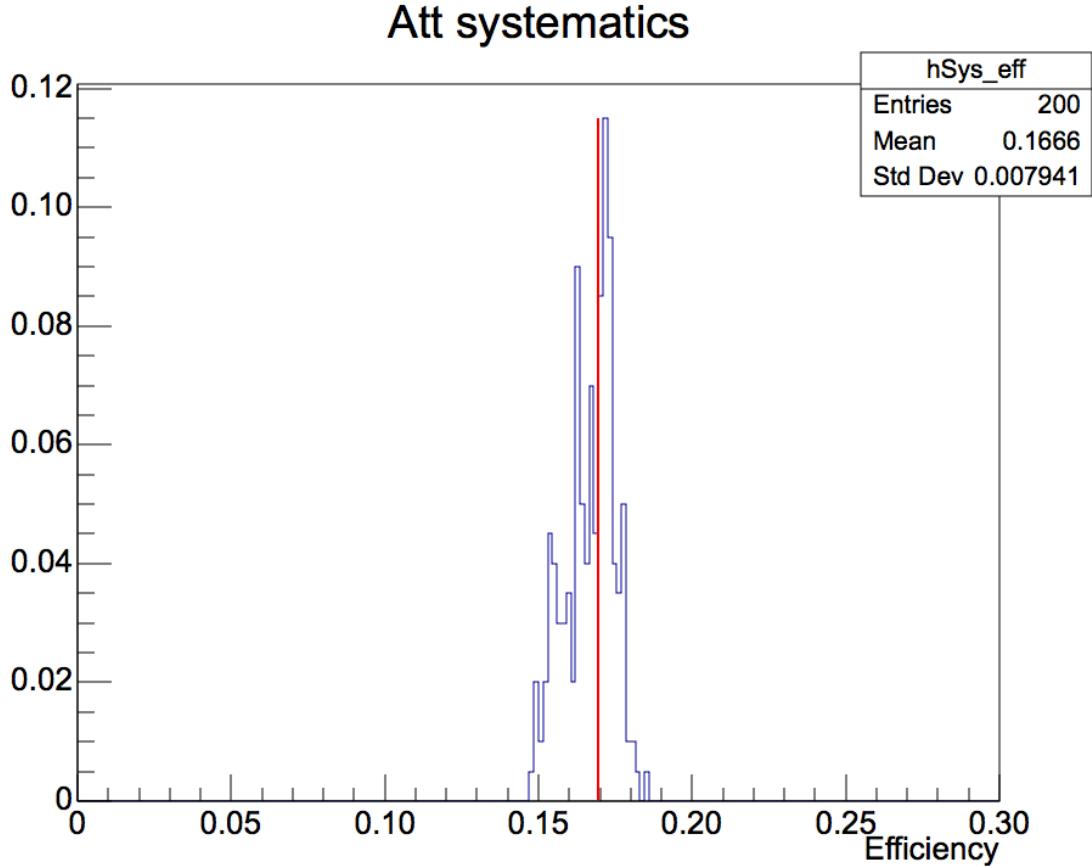


Figure K.8: Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the attachment parameter is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 5%.

amplitude of events, which is only weakly dependent on the shape of the response of the detector (for small variations). However, it can have a larger effect on the estimation of the width of events and the time until the return to baseline, since an improper deconvolution will either leave events with long tails (pushing the end of the event farther from the real end), or with undershoots (making the end of the event be computed sooner than it should). Either of those will lead to simultaneous

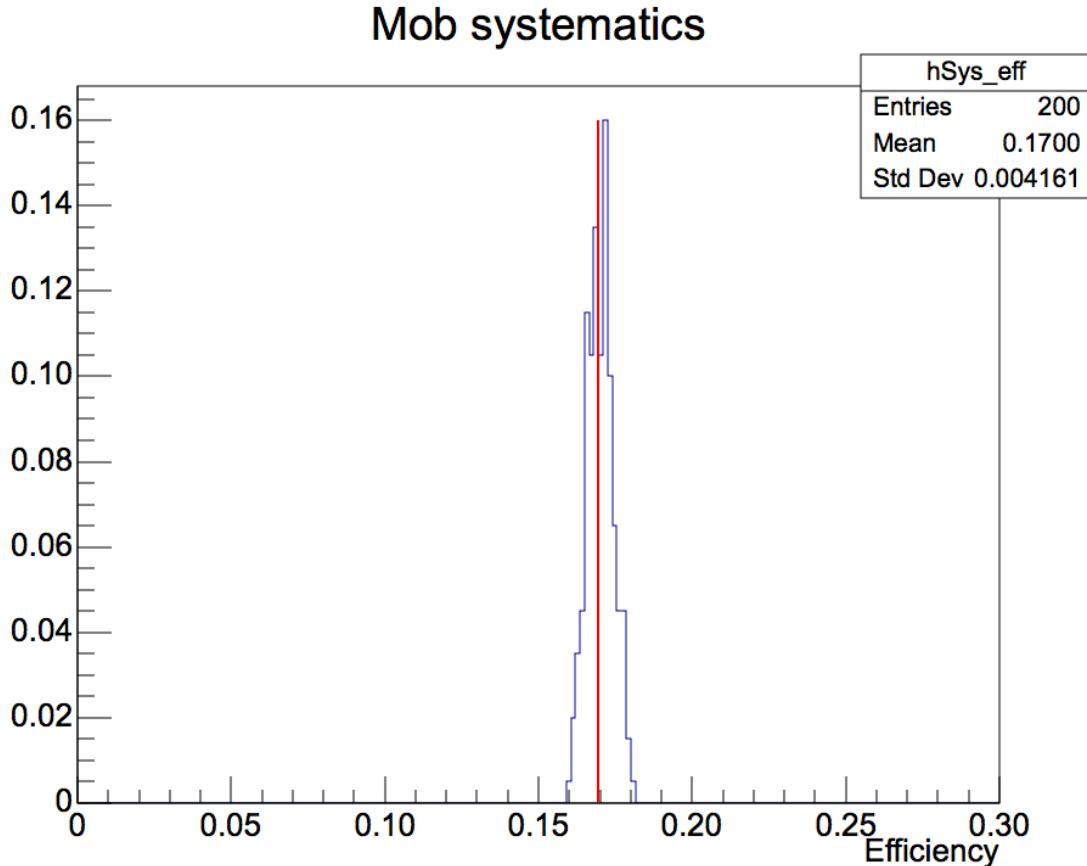


Figure K.9: Dispersion on the detector efficiency to KK axions, from 200 simulations of 10k events each, where the ion mobility is left to vary according to our calibrations, with all other parameters fixed. For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency is 2%.

events not being separated properly, or single events being wrongly reconstructed as two different events, respectively.

Fortunately, the disagreement between the literature and the calibrations are ultimately secondary to our actual use of the mobility: we use it as a simple pulse shape parameter. Understood that way, we actually have a good measure of the mobility, with a relative resolution of only 2%. The resulting variability in detector efficiency

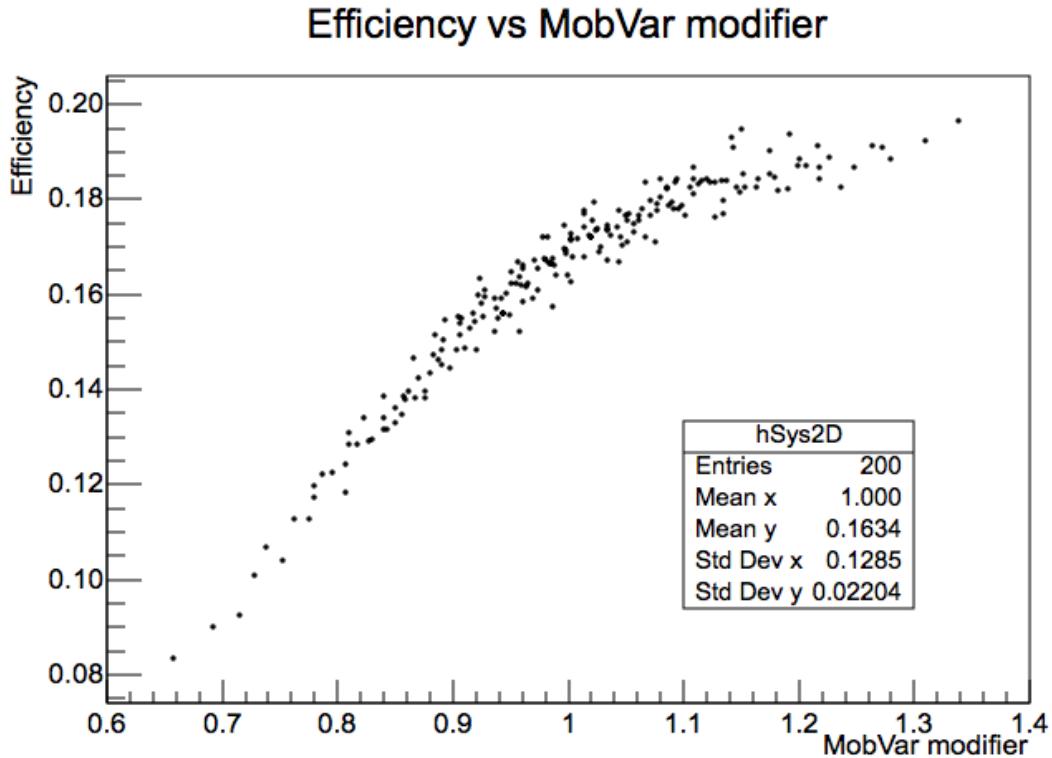


Figure K.10: Variation in efficiency when $2 \cdot 10^6$ events are simulated with an ion mobility drawn from a gaussian with a relative resolution of 18.9%. The y-axis is zoomed into the range of possible values to better appreciate the effect.

is shown in Fig. K.9; the effect is dominated by the statistic uncertainty.

This interpretation of the mobility parameter comes with an added complication: it is not a physical constant that depends only on the gas used, but instead varies pulse-to-pulse, since it “absorbs” the small differences in shape from one pulse to another. An upper bound on this effect is taken from our mobility calibration, by assuming the spread in measurements comes solely from actual shape variation, ignoring any contribution from the resolution on the calibration method. A new simulation was performed in which the ion mobility was left to vary with the measured dispersion

of 18.9%. The results are shown in Fig. K.10.

We observe that, for pulses in the lower tail end of mobility values, the processing only detects 9% of axion events, compared to 19.5% at high mobilities (events with low mobility are underdeconvolved, leaving a tail that makes it hard for the processing to split two pulses close to each other). However, the *average* efficiency of the detector is 16.34%, only marginally lower than the efficiency with a fixed mobility value, at 16.94%. This loss in efficiency results in an exclusion limit 1.8% higher, a negligible contribution to the final result.

K.9 Combined systematics

To generate the combined systematics, 1000 simulations with $10k$ events each, where the simulation parameters were left to vary freely according to the calibration uncertainties. The result is shown on Fig. K.11, including the effect due to our lack of knowledge on the drift time in the bulk of the detector (cf. Sec. K.2), whose main consequence is lowering the mean efficiency by $\sim 7.5\%$. The combined effect of all systematics leads to an uncertainty on the total efficiency of 21%.

Given that the exclusion limit depends as $1/\sqrt{\epsilon}$ on the total efficiency, we expect systematics to induce an uncertainty of order $\sim 10\%$ on the limit. This can be verified by computing the exclusion limit for the efficiency computed for each of these simulations. The median exclusion limit is $7.76 \cdot 10^{-13} \text{ GeV}^{-1}$, the exclusion limit larger than 84% of all limits ($\sim +1\sigma$) is $8.75 \cdot 10^{-13} \text{ GeV}^{-1}$, the exclusion limit larger than 97.5% of all limits ($\sim +2\sigma$) is $10.1 \cdot 10^{-13} \text{ GeV}^{-1}$.

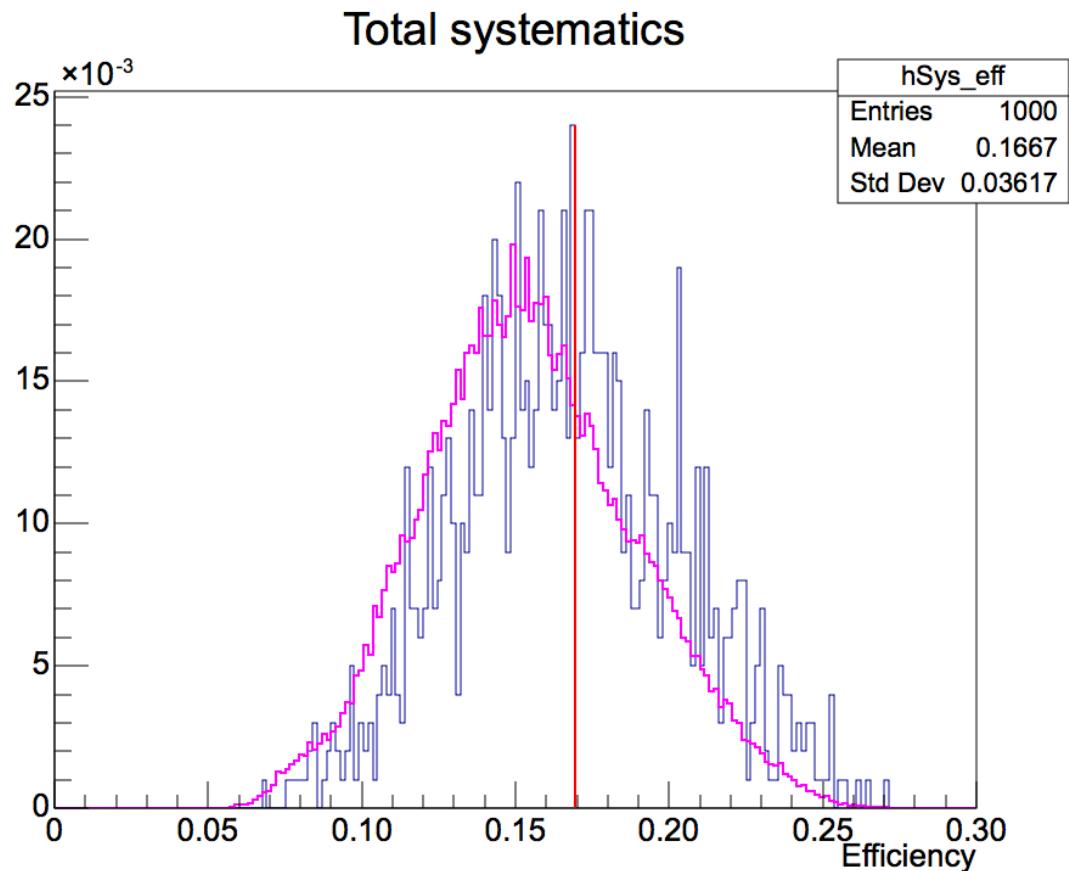


Figure K.11: Dispersion on the detector efficiency to KK axions, from 1000 simulations of 10k events each, where all simulation parameters are left to vary according to our calibration uncertainties. The purple (resp. blue) does (resp. does not) take into account the effect from our lack of knowledge on the drift time in the bulk of the detector (cf. Sec. K.2). For reference, the red line is the mean of the control simulations. The relative dispersion of the efficiency (including the drift model uncertainty) is 21%.