NNT : 2016SACLS582

Thèse de doctorat de l'Université Paris-Saclay préparée à l'Université Paris-Sud

Ecole doctorale N°576 PHENIICS
Spécialité de doctorat : Physique des particules

par

M. Valérian SIBILLE

Mesure de l'angle de mélange θ_{13} avec les deux détecteurs de Double Chooz

Thèse présentée et soutenue le 16 novembre 2016, à Saclay, devant le jury composé de :

M. Pierre DESESQUELLES	Professeur (CSNSM)	Président du jury
M. Fabrice PIQUEMAL	Directeur de recherches (CENBG)	Rapporteur
M. Michael WURM	Professeur (JGU)	Rapporteur
M. Thierry LASSERRE	Ingénieur-Chercheur (CEA)	Examinateur
M^{me} Lindley WINSLOW	Professeur (MIT)	Examinateur
M. David LHUILLIER	Ingénieur-Chercheur (CEA)	Directeur de thèse

An expert is a man who has made all the mistakes which can be made in a very narrow field.

Niels Bohr

Acknowledgements

L'épaisseur du manuscrit confère une certaine légitimité à l'auteur lorsque celui-ci entreprend l'écriture de remerciements tout aussi détaillés[1] que le contenu scientifique. Aussi vais-je en profiter et m'efforcer de ne laisser personne sur le bas-côté au cours de cette parade dithyrambique, tout en sachant qu'un *blitzkrieg* assure un déchargement rapide du pavé dans les rayons de la bibliothèque universitaire, ultime consécration. Je m'excuse donc par avance auprès de tous les participants dont le lever de pouce fut trop discret pour que je fisse des embardées à leurs signaux non moins valables.

La sphère

Dans un premier temps, je remercie ceux sans lesquels rien de toute cette aventure n'aurait été possible, aussi limitées fussent nos interactions (on ne s'intéresse pas aux neutrinos pour rien) : Philippe CHOMAZ et son successeur à la tête de l'IRFU, Anne-isabelle ETIENVRE. Bien que soumises à une érosion perpétuelle, les conditions de travail en haut du plateau ne sont pas sans rappeler l'avantage qu'octroierait le dénivelé avec la vallée lors d'une bataille rangée. Dès lors, je souhaite remercier les chefs successifs du SPhN, Héloïse GOUTTE et Franck SABATIE, ainsi que le chef adjoint, Jacques BALL. Je tiens à m'attarder tout particulièrement sur le soutien de ces derniers lors des joutes finales qui nous ont opposés à la matouchka doctorale.

Le service du SPhN ne serait pas ce qu'il est sans Isabelle RICHARD, Danielle CORET et Valérie POYETON, que je remercie pour leur gentillesse, leur attention et leur efficacité ! Qu'il est agréable de pouvoir se défausser de certaines tâches sans être trop inquiet quant à l'issue ! Je remercie également Patrick CHAMPION d'avoir réparé, à plusieurs reprises, nos stores blancs transparents, dont le pouvoir d'arrêt contre les rayons lumineux, quoique faible, est décisif lors de la rédaction d'un manuscrit.

Le labeur

Au cœur de mon travail se trouve mon directeur de thèse, David LHUILLIER, auquel je témoigne toute ma reconnaissance pour la pertinence et la justesse de son encadrement. Ce physicien hors pair a su être derrière moi lorsque j'en avais besoin (alors que je ne m'en rendais pas toujours compte), notamment pour donner des limites et un cadre à l'exploration

[1] La profusion de détails n'est pas un gage de qualité.

que constitue mon travail de thèse. Son approche directe et sa capacité à retranscrire des phénomènes physiques complexes, aussi aisément que pédagogiquement, ont opportunément perturbé mon équilibre chancelant. Bien que son temps libre fût victime d'un appauvrissement comparable à la décroissance d'un échantillon de ^{9}Li, il s'est toujours jeté – feutre en main – sur les problèmes physiques que je pouvais lui poser, mettant parfois fin à plusieurs jours de questionnement en quelques minutes. En sus de ses qualités techniques indéniables, David est un directeur accommodant et compréhensif, soit tout bonnement humain.

Had it not been for these bloated acknowledgements, this document would have been entirely reviewed by Fabrice PIQUEMAL and Michael WURM; I thank them dearly for appreciating my peculiar writing style and fondness for rigorous descriptions! I thank them equally for coming from Bordeaux and Mainz, respectively, to attend my viva voce. I also owe Lindley WINSLOW for being available remotely at 4am (Boston time) whilst retaining enthusiasm, Thierry LASSERRE for travelling from Munich, and Pierre DESESQUELLES for presiding over the committee.

On revient toujours à l'un des pères fondateurs de Double Chooz, Thierry. Je lui suis reconnaissant de m'avoir poussé à travailler sur les mesures de pesée, sujet, comme il le souligne justement, fort simple sur le principe, mais autrement plus épineux lorsqu'il s'agit d'afficher des systématiques maîtrisées. Je lui sais gré de sa confiance, sa gratitude et son soutien pour l'obtention d'un post-doc, il n'est pas pour rien dans ce qui m'attend à Boston. Ces analyses de mesures de pesée ont été accompagnées d'un travail quasi-archéologique sur un tissu de données suranné ; cette fouille au centre principal n'aurait jamais porté ses fruits sans la disponibilité de Jean-Christophe BARRIERE. Mentionner le centre conduit naturellement à remercier Guillaume MENTION, auquel je dois la relecture assidue de la note technique sur les simulations de spectres et des encouragements, qui ont assis la validité de mon entreprise. De la même manière, je ne saurais oublier cette semaine de calibration intense à Chooz, avec Matthieu VIVIER, agrémentée d'Ardennaises livrées par notre porte-parole, Hervé de KERRET, et consommées dans l'empressement. La parenthèse musicale dans la voiture, ainsi que les discussions sur le monde de la recherche, ont été autant de données pour mon étalonnage interne. Non sans poids dans le succès de cette thèse, je dois mon premier *shift* et mes premières incursions en territoire belge houblonné à Vincent FISCHER, qui figure inévitablement dans les présents balbutiements.

Double Chooz ne se limite pas au CEA ; je tiens ainsi à mettre l'accent sur Romain RONCIN pour son entrain et son ouverture d'esprit, et sur Emmanuel CHAUVEAU pour sa disponibilité sur site et à distance, et ce à n'importe quelle heure ! Merci également pour ce franc-parler si cocasse, ce goût pour les bonnes choses, et merci à l'humain qui se cache derrière le champ de compétences ! En ce qui concerne les permanents de plus longue date, je ne saurais omettre l'IPHC et Cécile JOLLET-MEREGAGLIA, pour tous les échanges que nous avons eus sous son œil bienveillant. Je me dois également de remercier les membres de l'APC et Anatael CABRERA expressément pour sa confiance et son regard sur mon travail.

Not only did I have the chance to travel Europe extensively during this three-year undertaking, but I even got the opportunity to spend one month in one of the queerest countries

I have set foot in, i.e. Japan. For this invaluable and unforgettable experience, I would like to express my deepest gratitude to Masaki ISHITSUKA, without whose steady support, kindliness, and open-mindedness, this exchange would have had a different taste. Many thanks also go to Michiru KANEDA for his help with the Japanese computing grid and for introducing Victoria and me to okonomiyaki! I owe Ralitsa SHARANKOVA much for fetching me at Haneda, and assisting me in the Tokyo underground and the O-Okayama Tokyu Store! Eventually, I am grateful to Masahiro KUZE for making my stay at Tokyo Tech possible.

I am deeply indebted to Susanne MERTENS for having shared her eye-opening experiences with me in Jyväskylä, and for having introduced me to Joseph FORMAGGIO in London, whom I thank sincerely for hiring me as a post-doctoral associate, thereby allowing me to tackle the question of neutrino masses.

Même si cet investissement ne figure pas, à proprement parler, dans ce document, je souhaite manifester toute ma gratitude à Sébastien DE ROSSI, qui m'a donné bien plus qu'un coup de pouce dans mes activités d'enseignement à travers les salles de SupOptique et du campus d'Orsay.

The mötley crüe

En dépit de leur proximité géographique du lundi au vendredi (hors jours fériés et ponts CEA), j'ai rarement travaillé avec les personnes citées ci-dessous ; pourtant, elles ont eu leur influence, sinon sur le manuscrit, au moins sur l'auteur.

J'ai effectivement eu la chance de rompre le jeûne de midi régulièrement avec ces théoriciens aux personnalités si dépaysantes que sont Jaume CARBONELL et Pierre GUICHON ; plus souvent au CERN que Pierre n'est en Australie, j'ai tout de même largement bénéficié de la bonne humeur de Nicole D'HOSE. De la même manière, le rez-de-chaussée neutrino s'est vu renforcé fréquemment par Alain LETOURNEAU et Thomas MATERNA, dont les caractères si différents se conjuguent en une alchimie théâtrale, montrant ainsi un autre visage de la recherche.

Chaque découpage a ses limitations, et si Antoine COLLIN a su me mettre patiemment sur les rails de la recherche et partager – encore aujourd'hui – son expérience emplie de sagacité, je retiendrai également sa culture débordante, ses opinions originales et son entrain merveilleux par-delà les grillages du CEA. Une partie non négligeable des résultats de cette thèse a été obtenue alors que je partageais le bureau 30A avec Maxime PEQUIGNOT ; sa bonne humeur, son rire facile et sa tolérance ont grandement facilité le partage des parcelles fertiles du lot 30A ! Ce lopin de terre fécond n'a pas laissé indifférent Thibaut METIVET, dont les incursions répétées ont non seulement permis de redessiner le monde, mais m'ont également laissé apercevoir le vaste océan de possibilités offertes par le langage de Bjarne STROUSTRUP. Sans sa sensibilisation à l'esthétique de la pieuvre, sa disponibilité, ou sa pédagogie, je serais – sans l'ombre d'un doute – toujours en train « d'écrire du C++ comme du Fortran » ; il m'a appris à naviguer avec mes propres voiles, et même si je retrouve

régulièrement de l'eau sur le pont, le navire n'a pas encore chaviré ! Je finis par une pensée pour Adrien BLANCHET, arrivé dans le 30A en tant que stagiaire au moment le moins opportun : préparation de Neutrino 2016 et spectre du manuscrit omniprésent ; je ne crois pas m'égarer en affirmant que n'importe quelle cohabitation lui sera désormais charmante !

Presque occupant du 30A lui aussi, véritable onde plane du SPhN, c'est à Raphaël BRISELET qu'incombe la transition. Je ne boude pas mon plaisir face à son ouverture d'esprit, sa curiosité fantastique, son sourire indéfectible et le piment qu'apportent au SPhN ses goûts singuliers. Occupant le bureau de Raphaël avec une parcimonie garante de résultats élaborés, je n'oublie pourtant pas Maud AIRIAU et sa conversation facile. Nous talonnant d'une année, je remercie Loïc THULLIEZ pour sa motivation lors de nos excursions forestières et sa démarche initiatique ayant abouti à l'éclosion des « précaires du SPhN ». Amongst the great runners of our department, I would like to thank Michał RĄPAŁA for his easy-going yet discreet personality (sought-after skills on the office-sharing market). I truly enjoyed hearing his experience and discussing French, Polish and Brazilian ways. Dziękuję bardzo! Indubitablement plus expansif que son homologue de l'est, la tournée des athlètes se poursuit avec Antoine VIDON, que je remercie vivement pour le torrent d'histoires poilantes qu'il apporte au moulin des discussions ! Grand chef, sportive de haut niveau, jardinier extraordinaire à ses heures perdues, je suis profondément reconnaissant envers Aurélie BONHOMME pour sa bienveillance bonhomme, son questionnement de l'ordre établi, ses talents de dessinatrice et le meilleur houmous de la région ! Pour avoir délicieusement partagé son vécu et son expérience de travail aux quatre coins du monde, pour le NFT de Boston et sa relecture stimulante des trois premiers chapitres de ce document, je remercie généreusement Nancy PAUL. Bande hétéroclite s'il en est, pour le moins charmante, je ne manque pas de mentionner les « théoriciens du fond », Pierre ARTHUIS, Benjamin BAILLY et Mehdi DRISSI. J'exprime toute ma sympathie envers le Barberousse du SPhN, Alessandro MINOTTI, pour ses avis tranchés comme par un sabre d'abordage, ses encouragements néanmoins affables et sincères, et ses propositions de sorties en pleine mer parisienne.

Le soutien indéfectible

Cette odyssée a été ponctuée d'escales en terres accueillantes, qu'elles soient musicalement fertiles, internationales, largement culturelles, ou qu'elles arborent les fruits bigarrés d'une longue amitié ; ces ports d'attache riment surtout avec : Sylvain, Théo, Alexis, Antoine, Geoffrey, Nathan. En particulier, je désire souligner la capacité d'écoute d'Antoine, sa ténacité et ses accueils répétés à Londres. Je voudrais également exprimer tout mon soutien à Alexis pour sa thèse, tout laisse à croire qu'il me rejoindra sans difficultés sous le dôme emblématique du MIT ; Boston s'ajouterait de bonne grâce à la liste des villes que nous avons conjointement foulées ! Naturligtvis har jag inte glömt ni, Birgit och Ben, det var jätteroligt att hänga ihop i Sverige och att träffas igen i Tyskland, vi hörs!

Le soutien de l'abbaye Notre-Dame-de-Saint-Rémy n'est pas en reste, et c'est principalement pour ce numéro 10 si onctueux et si parfumé – qu'il m'a été donné de déguster

régulièrement, grâce à la proximité providentielle de la congrégation avec le site de Chooz – qu'elle figure ici.

As a worthy substitute for coffee, I am acutely indebted to Ziggy Stardust, Tony Iommi, Phil Lynott, Steve Harris, Geddy Lee, Mikael Åkerfeldt, Robert Fripp, Steven Wilson, Gavin Harrisson, Maynard James Keenan, Mike Vennart, Mariusz Duda, Sean Reinert, Don Anderson, Eric Jernigan, Chris Hrasky, Mike Patton, Sully Erna, Kevin K.R. Starrs, Black Francis, Robert Smith, Jack White, Dave Brubeck, Bill Evans.

Pour ses conseils et l'œil distant mais sûr qu'il a gardé sur moi, pour s'être montré disponible au cours de tous ces déménagements, je suis reconnaissant envers mon père. Je remercie ma petite sœur d'être venue me faire une surprise le 16 novembre, et ma mère, pour ses appels téléphoniques répétés !

Хоть и последняя, но не менее важная: моя любимая Вика, я никогда не смогу отблагодарить тебя за всё то, что ты для меня сделала, пока я писал диссертацию. Без сомнения, диссертация была бы совсем другой, если бы ты не отдавала мне своё свободное время и не была бы такой терпеливой. Теперь ты даже почти шеф-повар! Ты во всём мне помогала, не щадила себя и заплатила дорогую цену. Но благодаря аспирантуре, мы посетили так много стран и многое пережили! Большое тебе мерси, дорогая!

Contents

Introduction

In the bestiary of elementary particles, neutrinos are quaint creatures. These neutral leptons are so elusive that they might be qualified as bystanders of particle physics, yet nothing could be further from the truth. In fact, their etherealness is rivalled but by their abundance and nearly lightning-fast speed, i.e. properties which make them ideal candidates for studying locations humanly inaccessible, be they stars, the core of the Earth, or man-made sources of energy such as nuclear power plants.

By all manner of means, neutrinos themselves come with their own secrets, and the more we learn about these messengers of new physics, the less we doubt their peculiarity. Although theories can accommodate for them by minimal extensions, it seems as though something is awry. For indeed, their individual masses, albeit unmeasured, are unquestionably small, suspiciously small. By the same token, neutrinos mingle amongst one another, in a fickle manner, presenting the largest mixing between particle species ever observed. If that was not enough, antineutrinos and neutrinos seemingly mix in different ways, which heralds consequences all the greater that these neutral fermions exist in copious quantities throughout the universe.

For all these prospects, we must first ensure that the gateway to further knowledge, i.e. the value of the smallest mixing parameter θ_{13}, is not biased. Quite apposite in the year 2016, let us emphasise that, like when attributing Nobel prizes, great care must be taken that values – or discoveries – have been cross-checked and identically observed by other experiments. Double Chooz gave the first direct indication of the non-zeroness of θ_{13} with reactor antineutrinos, and like other experiments, it still strives to refine the significance of its measurement.

Accurately measuring θ_{13} is no leisurely stroll, and the difficulty of such an endeavour is underscored by the length of this document. The latter is divided into four parts, of varying lengths.

The first part reviews the main properties of neutrinos by way of two chapters. Chapter 1 focuses on the discovery and first properties of neutrinos, while Chapter 2 introduces the neutrino oscillation phenomenon and eventually derives the antineutrino oscillation probability relevant to this thesis.

Part II is set on presenting a detailed picture of the Double Chooz experiment, from the production of electron antineutrinos (Chapter 3), to the design of the two detectors with which they may interact (Chapter 4). Additionally, in Chapter 5, the antineutrino production

model to which the recorded events are compared, along with the reconstruction algorithms, are reviewed. After the analysis cuts have separated the wheat from the chaff, the actual measurement of $\sin^2(2\theta_{13})$ – by means of a multi-detector configuration – is performed in Chapter 6.

Part III is dedicated to the background dominating the uncertainty on the $\sin^2(2\theta_{13})$ measurement from Chapter 6, caused by the decays of cosmogenic isotopes within the detector itself. These decays, that hamper the reliability of the antineutrino spectra, are simulated in Chapter 7. The building of spectra from these simulated events, complemented by a thorough error treatment, is presented in Chapter 8. The corresponding decays are partially selected within the data samples by dint of a cosmogenic veto, the performance of which is discussed in-depth in Chapter 9, for both detectors. The assessment of the rate of cosmogenic background remaining after the veto has been applied is the main topic of Chapter 10.

Chapter 11 is a loner in its Part IV, albeit no less stirring, for it addresses the absolute and relative normalisations of the antineutrino rates observed in both detectors.

Part I

Neutrino physics

Chapter 1

Admitting neutrinos in the field

Nowadays, neutrinos are perfectly valid contenders of the particle physics playground, and one may dare say, amongst the top-rated and most exciting players of the beginning of the twenty-first century. Now is also the time to remember that it has not always been like that. Not only did we not know that they came in several species – the number of which still being subject to debate – but a mere century ago, neutrinos we regarded as a pure construct of the mind, even for eminent physicists such Niels Bohr. At best, neutrinos were bookkeeping devices to rescue the conservation laws. And yet, if they would safeguard the energy conservation law, they would also emphasise parity violation, and even be detected with a different flavour than that with which they had been produced. But before studying neutrino oscillations, which we save as the main topic for Chapter 2, let us go back in time to the theoretical birth of this exceedingly abundant particle the neutrino is.

1.1 Discovery

1.1.1 Conundrum and postulate

Less than twenty years after the discovery of natural radioactivity by Henri Becquerel [1], the continuous nature of β-spectra was exhibited by James Chadwick in 1914 [2]. The kinematics of two and many-body decays will be extensively reviewed in part III, but it is not too much of a forecast to point out that two-body decays are characterised by a spectrum showing two distinct kinetic energy peaks. The disagreement between the experimental evidence and the two-body decay assumption

$$X \to Y + e^{-}, \tag{1.1}$$

where X stands for the decaying nucleus, Y its daughter, e^{-} the ejected electron, could not have been greater.

Undoubtedly, a third particle had to be involved in β-decays. This particle had to be neutral to conserve charge and remain hardly detectable. It must also be extremely

lightweight, for occasionally, there is no missing energy so that the maximum kinetic energy ${T_e}^{max}$ available for the electron is

$$ {T_e}^{max} = m_X - m_Y - m_e, \tag{1.2} $$

which still holds up to the accuracy of current measurements, with m indicating the nuclear masses. Besides, it must carry a $1/2$ spin to ensure the latter is conserved (thus, coupling with the electron would result in a 0 or 1 spin). Such a candidate was originally labelled "neutron" by Wolfgang Pauli, in 1930, in the originally-derided "Liebe Radioaktive Damen und Herren" open letter [3]. However, the "neutron" as we picture it today, would be discovered in 1932 by James Chadwiwick [4], and appear as far too heavy a fermion to meet Pauli's requirements. Enrico Fermi made the most of it all and offered, at the turn of the following year [5][1], a theory of β-decay so successful that Pauli's suggestion eventually had to be taken seriously. Thereafter, the lightweight neutral fermion, enforcing energy and spin conservation in β-decays, was labelled "neutrino" (symbol ν), from the Italian equivalent of "small neutron". Thus, the theoretical postulate – that "ν" embodies – corrects (1.1) into

$$ X \to Y + e^- + \nu. \tag{1.3} $$

1.1.2 Compelling evidence

Thanks to Fermi's theory, which can truly be regarded as a cornerstone in the building of the Standard Model, the theoretical foundations for the role of neutrinos had been laid. In 1947, the missing particles on the photographic emulsions of Cecil Frank Powell, that lead to the discovery of pions (π) [6], was more compelling evidence for the existence of neutrinos, as highlighted in the decay of the former into muons (μ)

$$ \pi \to \mu + \nu. \tag{1.4} $$

Less than three years later, the spectrum of the subsequent muon decay was confirmed to be continuous, with a mean energy of 34 MeV and an endpoint of 55 MeV [7][2], thereby leaving room for two neutrinos

$$ \mu \to e + \nu + \nu. \tag{1.5} $$

1.1.3 Detecting a poltergeist

Despite the theoretical motivations, there remained to directly detect this ghost particle, for it left no tracks, and did not decay. If truth be told, no one had every seen a neutrino do

[1] The first iterations appear in Italian and German journals because the famous Nature journal had deemed Fermi's theory "too remote from reality".

[2] Which is in stunningly good agreement with what is reconstructed from such decays in the Double Chooz detectors.

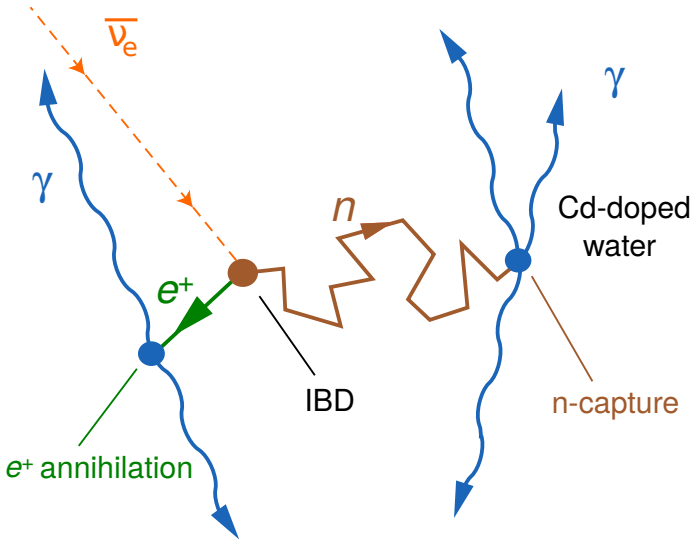

Figure 1.1 – Reactor antineutrino detection method used by F. Reines and C. L. Cowan at the Hanford and Savannah River experiments. Antineutrinos interact with water protons through inverse β-decay. The emitted e^+ annihilates with an e^- of the medium, and a few microseconds after, the n is captured by a Cd nucleus, which in turn, emits several γ's. The γ's can be detected thanks to horizontal liquid scintillator tanks not drawn on the schematic.

anything. And yet, as Fermi's interaction predicts, it should interact with a target full of water (and therefore of protons p) through the so-called "inverse β-decay"

$$p + \overline{\nu} \rightarrow n + e^+, \tag{1.6}$$

thereupon emitting a neutron (symbol n) and Dirac's positron (symbol e^+). In equation (1.6), we took a step forward by enforcing the antimatter character of the neutrino involved in this decay. This distinction first appears experimentally in 1955, when Raymond Davis reported his failed attempt to detect antineutrinos ($\overline{\nu}$) – allegedly emitted by the Brookhaven nuclear reactor – using the neutrino-sensitive reaction [8]

$$^{37}\mathrm{Cl} + \nu \rightarrow {}^{37}\mathrm{Ar} + e^-, \tag{1.7}$$

that was first advocated by Bruno Pontecorvo [9][3]. In order to unambiguously detect the interaction of the antineutrino with a proton from a water tank, both the double coincidence between a prompt and a delayed event (see Figure 1.1) and a significant shielding to cosmic rays, were crucial.

Indeed, if the requirement of a time coincidence between the annihilation of the positron and the capture of the neutron – a technique still put to good use in Double Chooz (with Gd

[3]Before he left Canada for USSR in 1950, B. Pontecorvo had initiated the building of a neutrino detector using the Chlorine-Argon technique at the Chalk River laboratory.

instead of Cd, cf. Chapter 4) – removes a large amount of background, in 1953, at the ^{239}Pu-producing Hanford reactor, the overwhelming rate of cosmic muons prevented Clyde Lorrain Cowan and Frederick Reines from detecting the antineutrino with enough significance [10]. On the other hand, when C. L. Cowan and F. Reines moved to the better shielded Savannah River experimental site, with a detector design aimed at rejecting backgrounds, they could detect a reactor-power dependent signal in good agreement with the predicted cross-section [11]. The antineutrino had been discovered.

1.2 First properties

1.2.1 Several neutrino flavours

Along with the experimental progress to detect (anti)neutrinos, theoretical works, initiated by E. J. Konopinski and H. M. Mahmoud [12], had introduced a quantity L, which we now call "lepton number". This supposedly conserved number evaluates to $L = 1$ for neutrinos and negatively charged leptons, and $L = -1$ for antineutrinos and the positively charged leptons. From that, it is clear that the Chlorine-Argon technique from (1.7) is insensitive to antineutrinos. Nevertheless, the theoretically allowed process

$$\mu \to e + \gamma \tag{1.8}$$

had never been observed [13], and the experiment was in accordance with process (1.5), hence suggesting the existence of different lepton numbers for different "kinds", or rather "flavours", of particles : L_e and L_μ at that time. Such an assumption implied that there were not only neutrinos and antineutrinos out there, but in fact, several flavours of them, as discussed in detail by B. Pontecorvo [14]. This hypothesis was confirmed by the 1962 Brookhaven experiment [15], in which muon (anti)neutrinos – produced by pion and kaon decays – successfully produced muons through

$$n + \nu_\mu \to p + \mu^- \tag{1.9}$$
$$p + \overline{\nu_\mu} \to n + \mu, \tag{1.10}$$

but during which the forbidden processes

$$n + \nu_\mu \to p + e^- \tag{1.11}$$
$$p + \overline{\nu_\mu} \to n + e^+, \tag{1.12}$$

were not observed in meaningful amounts. Everything comes in threes, at least that much can be said of the neutrinos that interact with the weak interaction[4], and a few decades later, the ν_τ was found by the DONUT collaboration [16].

[4]This thesis shall eagerly refrain itself from opening the "sterile neutrino" can of worms before the hurly-burly's done.

1.2.2 Parity and handedness

As emphasised in 1956 by T. D. Lee and C. N. Yang, the so-called "$\tau - \theta$ puzzle" was but an incentive to study parity violation in the weak interaction [17]; no experiment determining whether this interaction differentiated the right from the left, had ever been performed. Mathematically speaking, the parity operator changes the sign of all the spatial coordinates, and hereby the direction of motion of particles. The parity operator leaves the spin direction unchanged, thus, aligning the spin of particles or nuclei in the direction opposite to their previous ones, while retaining the observation along the same spatial direction, is effectively a parity transformation. Observing an asymmetric behaviour in the β-decays of nuclei, when reversing the direction of the magnetic field polarising them, is consequently a proof of parity violation.

Examining the parity conservation in β-decays requires an allowed[5] transition, with a spin change $\Delta J = 1$, so that the electron and the antineutrino always align the spin they carry away with that of the daughter nucleus. Moreover, orienting the spins of nuclei demands a H/T ratio (with H the magnetic field, and T the temperature) so high to overcome the tiny value of the Bohr magnetron μ_B, that nuclei with a large coupling between the nuclear spin and the electronic moment had to be used. On top of having a relatively manipulable half-life of 5.27 y, the $J^\pi = 5^+$ (with π the parity of the state) ground state of ^{60}Co, which mainly β-decays to the 4^+ state of ^{60}Ni in

$$^{60}\mathrm{Co} \rightarrow {}^{60}\mathrm{Ni} + e^- + \overline{\nu_e}, \tag{1.13}$$

meets all the aforementioned requirements. An asymmetry between the number of detected electrons when polarizing the ^{60}Co upwards or downwards was observed by Chien-Shiung Wu in 1957 [18]. It follows that the antineutrino is always[6] emitted in the half-space into which the ^{60}Ni spin points, the direction of which being identical to its polarized ^{60}Co mother. Recalling that the spin of the antineutrino is aligned with the ^{60}Ni spin, it means that its helicity – defined as the sign of the projection of the spin on the direction of motion – is always positive (see Figure 1.2).

A few months later, M. Goldhaber confirmed that the neutrino emitted in the electron capture reaction

$$^{152}\mathrm{Eu} + e^- \rightarrow {}^{152}\mathrm{Sm} + \nu_e, \tag{1.14}$$

had a negative helicity [19]. In a nutshell, neutrinos are left-handed, and antineutrinos are right-handed. It was natural to assume that neutrinos were massless, for in addition of the apparent mass conservation in β-decays (see 1.1.1), massive neutrinos could have been overtaken by a Lorentz boost.

[5] In allowed decays, the electron and the antineutrino do not carry any orbital angular momentum, which unambiguously determines the direction of their spins.

[6] Due to the difficulty of aligning all the spins of all the nuclei, C. S. Wu could only observe a significant asymmetry.

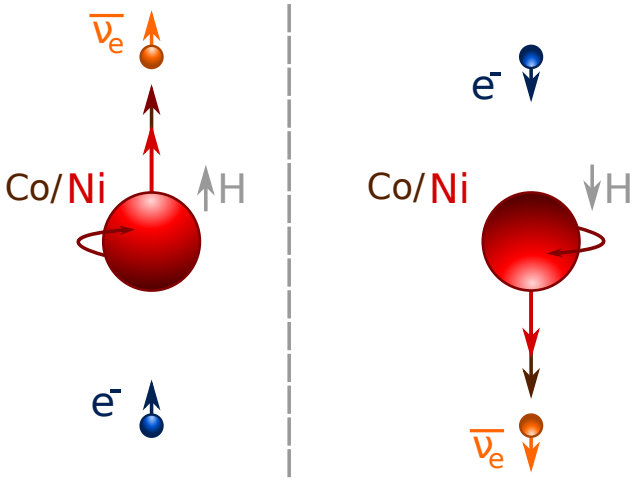

Figure 1.2 – In the β-decay of ^{60}Co into ^{60}Ni, the electrons are emitted in the direction opposite the nuclear spins – which have been polarized by a magnetic field H – and the antineutrinos, thus, have a positive helicity. Spin conservation is emphasised by the colourful arrows.

Undoubtedly, the weak interaction – that is to say, the only interaction through which neutrinos interact – does not conserve parity. It is worthwhile noting that although the mirror image of the left-handed neutrino is the yet to be detected right-handed neutrino, combining the charge conjugation C – which turns matter into anti-matter – with the parity operator P, does output a particle we are well acquainted with : the right-handed antineutrino. Obviously, it took us but a few more years to understand that the CP symmetry was not a symmetry good enough for this world, at least that much can be said for mesons [20]; the measurement to which this thesis contributes, is closely related to the global picture.

Chapter 2

Oscillating neutrinos

If it were not for neutrinos, the Standard Model, as formalised in the sixties for its electroweak part, and in the seventies for its strong interaction component, would hold in triumph, and the 2012 discovery of the Higgs particle would but vouch for it. However, over the last two decades, irrefutable evidence for neutrino oscillations has been exhibited by several collaborations of physicists spread across the terrestrial globe, using either natural or man-made neutrino sources. Neutrino oscillations were world-acknowledged by the 2015 Nobel prize in Physics, attributed to Takaaki Kajita and Arthur B. McDonald, pillars of the Super-Kamiokande and SNO experiments, respectively. The oscillation phenomenon is not highlighted so much for it is the main topic of this thesis, but rather because it unequivocally proves that neutrinos have a mass, and in doing so, unveils a new area of physics beyond the Standard Model. Therefore, building up an increasingly accurate knowledge of the parameters characterising neutrino oscillations, amongst which θ_{13} is a peculiar contender, paves the way for a better understanding of the physical world.

2.1 Experimental signs

2.1.1 Disturbing experimental data

Parity violation favoured the development of chiral symmetries in which neutrinos were massless [21, 22]. And yet, disturbing data started to accumulate at the end of the sixties, shaping what would be referred to as the "solar neutrino problem".

Indeed, the very same Raymond Davis who had confirmed the impossibility to detect reactor antineutrinos with the Chlorine-Argon technique (see 1.1.3), turned his attention [23] to the number of neutrinos emitted by nuclear fusion inside the sun. To do so, R. Davis and his colleagues placed their 380 m^3 tetrachloroethylene Brookhaven detector at the Homestake Gold Mine, 1478 m below the surface, and for more than a decade, issued frequent reports [24, 25] acknowledging a two-third deficit in the rate of detected neutrinos, when compared to the predictions of John Bahcall [26, 27]. The detected rate was exceedingly low and

about one ^{37}Ar atom was produced every two days, which had to be collected regularly by helium purging, to examine its subsequent decay through electron capture. Of course, such an indirect radiochemical method, along with the expected low rate, arose doubts amongst experimentalists. Similarly, the solar model was called into question [28], and the rather high energy threshold of the Chlorine detection method limited the comparison to the production of ^{8}B in the sun[1], which is far from dominant.

2.1.2 Supporters

More than twenty years after it had recorded its first neutrino, the Homestake experiment eventually saw its results backed by data from the Kamiokande and Baksan detectors.

In 1990, through the use of elastic scattering on electrons

$$\nu + e^- \rightarrow \nu + e^-, \tag{2.1}$$

the Japanese water Cherenkov[2] detector measured the boron-related neutrino flux, and found it to be $0.46 \pm 0.05(\text{sys.}) \pm 0.06(\text{stat.})$ of the value predicted by the solar model [29]. At this point, it is apposite to stress that if elastic scattering on electrons can proceed via the weak interaction mediator Z^0 for all neutrino flavours, there is an extra Feynman diagram mediated by the W boson for ν_e. Therefore, if a part of the ν_e flux had converted – or oscillated – to other flavours, the total neutrino flux could not be assessed, unless one were to rely on the charged current results from Homestake to isolate the ν_e contribution [30], which was not a persuasive disentangling procedure to all.

On the other hand, the Soviet-American Gallium Experiment (SAGE), sensitive also to low-energy neutrinos released by deuteron production[3] in the sun – which account for more than 90% of the solar neutrino production, and are less dependant on solar models than boron-induced neutrinos – reported a neutrino-capture rate from

$$^{71}\text{Ga} + \nu_e \rightarrow {}^{71}\text{Ge} + e^-, \tag{2.2}$$

40% lower than its prediction with a 90% confidence level [31]. At Gran Sasso, the GALLEX experiment would soon swell the ranks of the Gallium supporters [32].

2.1.3 Crowning achievements

2.1.3.1 Solar neutrinos

If the aforementioned experiments had put the lid on the coffin of the solar neutrino problem, the Canadian Sudbury Neutrino Observatory (SNO) genuinely nailed the pine box. When

[1] The signal comes from the β^+-decay of ^{8}B into ^{8}Be, which is accompanied by the emission of a ν_e.

[2] Pet peeve: great diligence is often taken to substitute "Ch" for "Č", thereby heralding a stylistic Czech or Slovak transcription, but let us emphasise that the Nobel laureate is actually Russian.

[3] Two protons fusion into a deuteron, thereby emitting a ν_e.

conceding that neutrinos change flavour according to the oscillation mechanism, one must needs design an experiment able to detect all species with the same efficiency. In addition, the original sun-produced ν_e's must be counted over the same period of time, to compare both fluxes unambiguously. To do so, SNO made the most of the relatively low break-up threshold of deuteron's (d):

$$d + \nu \to n + p + \nu, \tag{2.3}$$

a process equally sensitive to all flavours (only the Z^0 diagram is available to ν_e's), and for which neutron emission is a characteristic signature. Moreover, since the incoming neutrinos do not have a kinetic energy large enough to produce heavy leptons, the charged current channel

$$d + \nu_e \to p + p + e^-, \tag{2.4}$$

is – as made explicit – relevant to ν_e's only. The cherry on the cake is neutrino detection via elastic scattering on electrons (2.1), as in the Kamiokande (and their upgraded Super-Kamiokande version) detectors, which allows to cross-check the purely neutral and charged current channels. Taking into account the available Feynman diagrams, the expected fluxes from equations (2.3), (2.4) and (2.1), respectively read

$$\phi_{neutral} = \phi_{\nu_e} + \phi_{\nu_\mu} + \phi_{\nu_\tau} \tag{2.5}$$

$$\phi_{charged} = \phi_{\nu_e} \tag{2.6}$$

$$\phi_{elastic} = \phi_{\nu_e} + 0.15\left(\phi_{\nu_\mu} + \phi_{\nu_\tau}\right) \tag{2.7}$$

A last, in 2002, by means of a 1 kt spherical heavy water detector [33], located 2039 m below the surface, SNO published the following flux measurements

$$\phi_{neutral} = 5.09^{+0.44}_{-0.43}(\text{stat.})^{+0.46}_{-0.43}(\text{syst.}) \tag{2.8}$$

$$\phi_{charged} = 1.76^{+0.06}_{-0.05}(\text{stat.})^{+0.09}_{-0.09}(\text{syst.}) \tag{2.9}$$

$$\phi_{elastic} = 2.39^{+0.24}_{-0.23}(\text{stat.})^{+0.12}_{-0.12}(\text{syst.})\ . \tag{2.10}$$

Without reserve, all the previous experiments had seen a deficit because all but a third of the ν_e's produced by fusion reactions in the sun, had oscillated to the muon or tau flavour. As if that was not enough, the total neutrino flux was in conformity with the solar standard model [34].

2.1.3.2 Atmospheric neutrinos

Assuredly, the sun is not the only supplier of neutrinos and these neutral fermions are also produced in copious quantities in the upper atmosphere, particularly owing to high energy protons hitting air molecules, whereby pions are created (see Figure 2.1). Setting the right flavours on the pion and subsequent muon decays, we find

$$\begin{array}{l} \pi^+ \to \mu^+ + \nu_\mu \\ \pi^- \to \mu^- + \overline{\nu_\mu} \end{array} \quad \Longrightarrow \quad \begin{array}{l} \mu^+ \to e^+ + \nu_e + \overline{\nu_\mu} \\ \mu^- \to e^- + \overline{\nu_e} + \nu_\mu\ . \end{array} \tag{2.11}$$

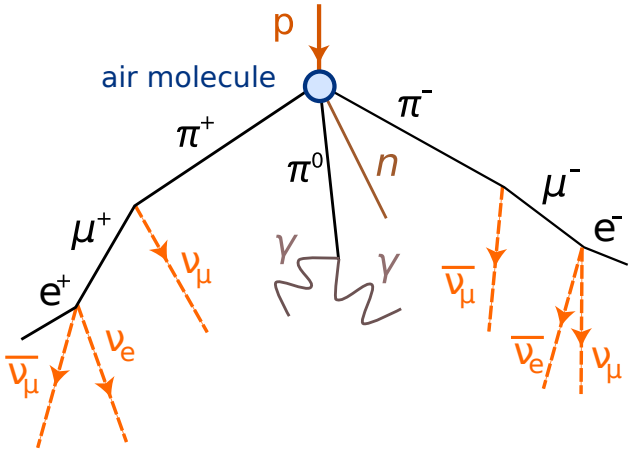

Figure 2.1 – Production of electron and muon neutrinos and antineutrinos in the upper atmosphere by bombardment of high energy protons.

Discarding neutrino oscillations for a just an moment, one would expect the ratio of the number of muon neutrinos (and antineutrinos alike) over that of electron neutrinos, to be close to 2, as is clear according to (2.11). The role of atmospheric kaons and the highly suppressed pion decays with an electron flavour are hereby overlooked, but this does not change the argument. In 1998, by way of charged currents – allowing flavour identification via the emitted charged lepton – the 50 kt water Cherenkov detector of Super-Kamiokande reported with great accuracy the value of the R ratio [35]

$$R = 0.63 \pm 0.03(\text{sys.}) \pm 0.05(\text{stat.}) \; . \tag{2.12}$$

The R quantity is defined as the data to Monte-Carlo ratio

$$R = \frac{N_\mu}{N_e} \frac{N_e^{MC}}{N_\mu^{MC}} \simeq \frac{N_\mu}{N_\mu^{MC}}, \tag{2.13}$$

with N standing for the sum of the number of neutrinos and antineutrinos. To put it differently, nearly half of the muon (anti)neutrinos were missing with respect to the prediction, and their oscillation to the tau flavour was a competing explanation.

Furthermore, the Super-Kamiokande collaboration was able to detect the direction of the incoming neutrinos. Neutrinos coming downwards onto the detector travel a distance that is of the order of $L = 10\,\text{km}$, on the other hand, neutrinos coming upwards have travelled a distance of the order of $L = 10\,000\,\text{km}$. Thus, defining a quantity that is a function of the zenith angle θ (cf. Figure 2.2), directly challenges the alleged distance-dependence of neutrino oscillations. The asymmetry A embodies such a quantity, it is defined as

$$A = \frac{N_U - N_D}{N_U + N_D}, \tag{2.14}$$

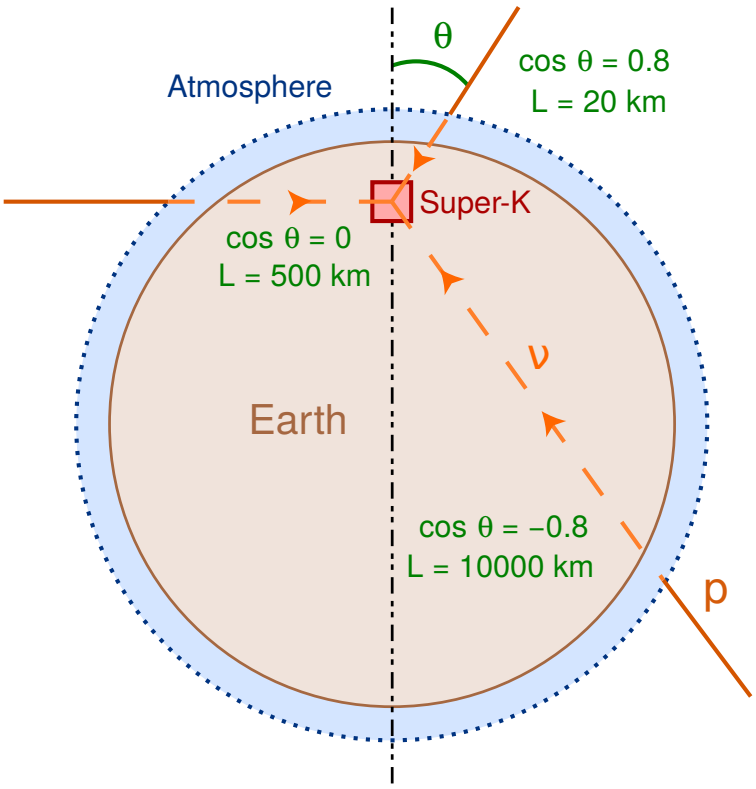

Figure 2.2 – Distance travelled by incoming atmospheric neutrinos as function of the cosine of the zenith angle θ, whose value is defined with respect to the axis of the Super - Kamiokande cylindrical detector. Thus, downward neutrinos have a positive cosine, and neutrinos coming upwards have a negative one.

where N_U and N_D are the number of upward and downward[4] events, respectively. If the asymmetry for electron-like events is consistent with zero (see [35]), and therefore no oscillations of electron neutrinos presumably happen, muon-like events with an energy larger than 1330 MeV exhibit a staggering asymmetry

$$A = -0.296 \pm 0.048(\text{sys.}) \pm 0.01(\text{stat.}) \,, \tag{2.15}$$

which deviates from zero by more than 6 standard deviations. Although Super-Kamiokande did not quite solve the solar neutrino problem, it did bear the most conclusive testimony to atmospheric neutrino oscillations. Not only did it show data compatible with $\nu_\mu \to \nu_\tau$ oscillation[5], but it also highlighted the dependence of this phenomenon on the neutrino energy and the distance it had travelled.

[4]To be precise, downward events are defined by $\cos\theta \in [0.2, 1]$ and upward events by $\cos\theta \in [-1, -0.2]$.

[5]Tau neutrinos were the only ones that Super-Kamiokande could not detect, hence the deficit. The OPERA collaboration would later endeavour to detect the $\nu_\mu \to \nu_\tau$ oscillation.

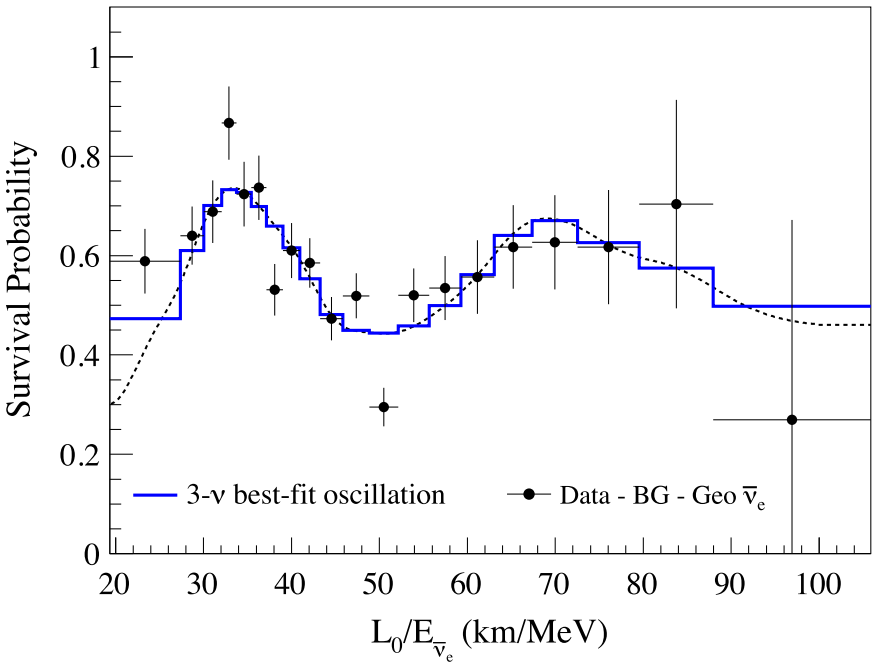

Figure 2.3 – Ratio of the background-subtracted number of $\overline{\nu_e}$'s to the expectation for no-oscillation, as a function of $L_0/E_{\overline{\nu_e}}$, with $L_0 = 180\,\text{km}$ the effective baseline and $E_{\overline{\nu_e}}$, the $\overline{\nu_e}$ energy. Extracted from KamLAND's data [38].

2.1.3.3 Anthropogenic neutrinos

Man-made sources do not fall short of expectations, they allowed us to discover the neutrino, and it is to them that we shall turn – in the framework of this thesis – to understand one of the last properties of neutrino oscillations.

Shortly after SNO's marvellous results, at the former Kamiokande site, the Kamioka Liquid Scintillator Anti-Neutrino Detector (KamLAND) was bent on observing oscillating $\overline{\nu_e}$'s, which were produced by the sizeable number of Japanese nuclear reactors. To characterise the $\overline{\nu_e}$'s, that had journeyed along a typical flux-weighted 180 km baseline, the detection method of choice was the inverse β-decay reaction (1.6), with which F. Reines and C. L. Cowan had raised their profiles (see 1.1.3). Unsurprisingly, in 2003, a deficit with respect to the standard expectation was observed, with a 4σ significance [36]. The unimpeachable energy-dependence of this deficit would latter be refined [37], thereby upholding it as the trademark of neutrino oscillations (at fixed distances). An up-to-date plot of this signature is to be found in Figure 2.3.

At much shorter baselines – essentially, a few hundreds of meters – the two thousand and tens have cast the limelight onto the Daya Bay, RENO (Reactor Experiment for Neutrino Oscillation), and Double Chooz reactor experiments, whose achievements we shall review later, for they all aim at precisely measuring the smallest (anti)neutrino deficit ever observed.

However discreet neutrinos can be, accelerator neutrinos are by no means bystanders.

In fact, in June 2011, the Tokai to Kamioka (T2K) collaboration provided the first positive indication of the value of the smallest deficit, with a 2.5σ significance [39]. Besides, with their energies of the order the GeV, and their baselines of a few hundreds of kilometres, experiments such as MINOS [40] tackle the same type of oscillation that affects atmospheric neutrinos. In addition, they do it with a different sensitivity, as will become obvious after we have laid the theoretical groundwork for measuring all the physical quantities that drive neutrino oscillations, in section 2.2.

2.2 Theory of neutrino oscillations

2.2.1 First draft

All the deficits and patterns observed by the previously introduced experiments can be explained in the framework of neutrino oscillations.

Although B. Pontecorvo, inspired by the kaon transitions $K^0 \to \overline{K^0}$, had speculated on the oscillation of neutral[6] particles in 1957 [41], and thus on that of $\nu \to \overline{\nu}$ [42]– each of which possibly being a quantum superposition of other particles – it is not quite the phenomenon that has been observed over the last fifty years[7]. In 1962, Ziro Maki, Masami Nakagawa and Shoichi Sakata, in an attempt to unify all particules into a unitary scheme, put forth the quaint idea that baryons could be compound systems of leptons and a new sort of "B-matter" [43]. Stirred by the ν_μ discovery (cf. 1.2.1), they also contemplated, for the first time, that the "weak" neutrinos ν_e and ν_μ, could be a mixture of "true" neutrinos ν_1 and ν_2, and therefore that $\nu_e \leftrightarrows \nu_\mu$ conversions, or rather, oscillations, were possible. They postulated that both representations were related by an orthogonal transformation, and that the weak neutrinos – which should separately conserve the leptonic numbers (see 1.2.1) – were but rotated versions of the true ones[8],

$$\begin{cases} |\nu_e\rangle &= \cos\theta\,|\nu_1\rangle + \sin\theta\,|\nu_2\rangle\,, \\ |\nu_\mu\rangle &= -\sin|\nu_1\rangle + \cos\theta\,|\nu_2\rangle\,, \end{cases} \tag{2.16}$$

with $\theta \in [0, \pi/2]$, the angle between the states $|\nu_e\rangle$ and $|\nu_1\rangle$.

[6] B. Pontecorvo paid attention to neutral particles because he focused on particle to anti-particle transitions, which are otherwise intricate...

[7] It does relate to the search for neutrinoless double-beta decays, though.

[8] In [43], the angle is actually opposite ours, and the true neutrinos are the ones rotated by θ with respect to the weak neutrinos, but we here follow the modern convention, that will simplify later comparisons and ease the reading.

2.2.2 Mixing matrix

2.2.2.1 Two-dimensional formalism

It would, however, take more than ten years, and the vital progress made by Cabibbo in the quark sector, to make plain that the weak neutrinos do not have a definite mass [44], and that there exists mixing with what are the true mass eigenstates [45]. Accordingly, we shall henceforth call the former "flavour eigenstates", and the latter, "mass eigenstates".

Mathematically, the system (2.16) defines the coordinates of the flavour neutrinos in the mass basis, from which is derived the change of basis matrix

$$P_m^{fl} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}. \tag{2.17}$$

The P_m^{fl} matrix embodies the rotation of the mass eigenstates by an angle of θ and allows to convert back into the mass basis, the coordinates of neutrinos which are expressed in terms of the flavour eigenstates. Explicitly writing the basis as a subscript turns the counter-intuitive observation into a rule of thumb (superscripts and subscripts cancel)

$$|\nu_m\rangle = P_m^{\cancel{fl}} \left|\nu_{\cancel{fl}}\right\rangle. \tag{2.18}$$

From there, the aptly-named change of basis matrix U_{MNS} reads

$$U_{MNS} = P_{fl}^{m} = \left(P_m^{fl}\right)^{-1} = {}^{\mathrm{t}}P_m^{fl} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}, \tag{2.19}$$

with the transpose being indicated as a pre-superscript. Consequently, the U_{MNS} matrix corresponds to a rotation by an angle $(-\theta)$ and its columns are simply the coordinates of the mass eigenstates in the flavour basis, hence partly justifying the abusive writing

$$U_{MNS} = \begin{pmatrix} U_{e1} & U_{e2} \\ U_{\mu 1} & U_{\mu 2} \end{pmatrix}. \tag{2.20}$$

The U_{MNS} matrix is the rotation we wish to apply to neutrinos in the "true" mass basis in order to get their coordinates in the flavour basis

$$|\nu_{fl}\rangle = U_{MNS} |\nu_m\rangle. \tag{2.21}$$

Any state from the mass basis reads $|\nu_m\rangle = x_1 |\nu_1\rangle + x_2 |\nu_1\rangle$, with x_1 and x_2 arbitrary complex numbers. In the mass basis, $|\nu_m\rangle$ simply reads

$$|\nu_m\rangle = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}_m. \tag{2.22}$$

Applying U_{MNS}, we obtain (x_e, x_μ), the coordinates of this state in the flavour basis

$$|\nu_{fl}\rangle = \begin{pmatrix} x_e \\ x_\mu \end{pmatrix}_{fl} = \begin{pmatrix} \cos\theta x_1 + \sin\theta x_2 \\ -\sin\theta x_1 + \cos\theta x_2 \end{pmatrix}_{fl}. \tag{2.23}$$

Taking $x_1 = 1$ and $x_2 = 0$ confirms that

$$|\nu_1\rangle = \cos\theta\,|\nu_e\rangle\ - \sin\theta\,|\nu_\mu\rangle\,. \tag{2.24}$$

The diligence of this paragraph may seem overly pedantic to the reader, but (2.24) is there to stress that U_{MNS} allows to expand the mass eigenstates into the flavour basis, and not the other way around, which is, more often than not, stated in other documents, based on an erroneous interpretation of (2.21). Such statements, in the generalisation that is to follow, usually lead to magical air-drops of stars in a crooked attempt to get the complex conjugates on the coefficients of U_{MNS}[9].

2.2.2.2 Three-dimensional case

Change of basis

When adding the tau neutrinos, the U_{MNS} matrix sometimes gets an additional subscript letter, hereby turning into U_{PMNS}

$$U_{PMNS} = P_{fl}^{m} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix}, \tag{2.25}$$

which is a unitary matrix, and thus satisfies an equation similar to (2.19), that is

$$P_{m}^{fl} = (U_{PMNS})^{-1} = (U_{PMNS})^{\dagger}\,, \tag{2.26}$$

where the † denotes the Hermitian transpose. Accordingly

$$P_{m}^{fl} = \begin{pmatrix} {U_{e1}}^{*} & {U_{\mu 1}}^{*} & {U_{\tau 1}}^{*} \\ {U_{e2}}^{*} & {U_{\mu 2}}^{*} & {U_{\tau 2}}^{*} \\ {U_{e3}}^{*} & {U_{\mu 3}}^{*} & {U_{\tau 3}}^{*} \end{pmatrix}, \tag{2.27}$$

and any flavour eigenstate $|\nu_\alpha\rangle$ where $\alpha \in \{e, \mu, \tau\}$, represented by a 1 at the α-th line of a column vector in the flavour basis, has the following coordinates in the mass basis

$$|\nu_\alpha\rangle = \begin{pmatrix} {U_{\alpha 1}}^{*} \\ {U_{\alpha 2}}^{*} \\ {U_{\alpha 3}}^{*} \end{pmatrix}_{m}. \tag{2.28}$$

In terms of states, we have for all $\alpha \in \{e, \mu, \tau\}$,

$$|\nu_\alpha\rangle = \sum_{k=1}^{3} {U_{\alpha k}}^{*}\,|\nu_k\rangle = \sum_{k=1}^{3} U_{k\alpha}^{-1}\,|\nu_k\rangle = \sum_{k=1}^{3} \left(P_{m}^{fl}\right)_{k\alpha}|\nu_k\rangle\,, \tag{2.29}$$

hereby corroborating the first-class role of P_{m}^{fl}, which is often improperly peddled as U_{PMNS}.

[9]When tackling this issue the proper way, that is, starting from quantum fields, one gets complex conjugates from the creation operators for particles, as opposed to antiparticles, the former bearing a "dagger".

Parametrisation

Let us express the U_{PMS} matrix in the case of N neutrino flavours and N mass eigenstates, with $N > 1$. Each eigenvalue λ, of a unitary matrix U, has modulus $|\lambda| = 1$, thus, there exists θ such that $\lambda = \exp(i\theta)$. By diagonalising the matrix, it is easy to see that it can be written as the exponential of a Hermitian matrix H so that

$$U = \exp(iH). \tag{2.30}$$

Since $H = H^\dagger$, Hermitian matrices have N diagonal real terms, and $N(N-1)/2$ off-diagonal independent complex terms, thereby amounting to a total of N^2 independent real coefficients. Likewise, orthogonal matrices O can be written

$$O = \exp(A), \tag{2.31}$$

where A is antisymmetric and verifies ${}^{\mathrm{t}}A = -A$, leaving room for

$$N_\theta = \frac{N(N-1)}{2} \tag{2.32}$$

real coefficients, or angles. Consequently, of the N^2 real coefficients that parametrise a unitary matrix, there remains

$$N_\delta = \frac{N(N+1)}{2} \tag{2.33}$$

phases, which cannot be expressed as angles. Nevertheless, $2N-1$ of these phases are already free parameters of the lepton fields[10], which leaves

$$N_\delta^{free} = \frac{N(N+1)}{2} - (2N-1) = \frac{(N-1)(N-2)}{2} \tag{2.34}$$

free phases in the U_{PMS} matrix.

Plugging $N = 3$ in (2.32) and (2.34), provides us with three mixing angles, which we baptise θ_{12}, θ_{13}, θ_{23}, and one phase δ. Abiding by the usual parametrisation (see [46]), for Dirac neutrinos, we can write

$$U_{PMNS} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{2.35}$$

with $c_{ij} = \cos(\theta_{ij})$, $s_{ij} = \sin(\theta_{ij})$. When multiplying all matrices out, we find

$$U_{PMNS} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}c_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}. \tag{2.36}$$

Solar experiments deal with parameters related to $|U_{e2}/U_{e1}| = \tan(\theta_{12})$ and atmospheric experiments, with $|U_{\mu 3}/U_{\tau 3}| = \tan(\theta_{13})$. In the meantime, reactor experiments, such as Double Chooz, are particularly interested in $|U_{e3}| = \sin(\theta_{13})$, as the title of this thesis hints at.

[10] For Dirac neutrinos, global U(1) gauge transformations are indeed allowed for e, μ, τ and ν_e, ν_μ, ν_τ.

2.2.3 Quantum-mechanical approach

2.2.3.1 Plane wave derivation

If truth be told, all this does not tell us how to experimentally extract the aforementioned parameters. In fact, all the experiments we have reviewed in section 2.1 involve neutrino production of a certain flavour, may it be in the sun as ν_e's, in the atmosphere as ν_e's, $\overline{\nu_\mu}$'s $\overline{\nu_e}$'s and ν_μ's, or in nuclear reactors, as $\overline{\nu_e}$'s. They also imply propagation from the production area to the detector, and eventually, detection, by charged or neutral currents. As 2.2.2 recommends, all these operations imply thoughtful change of coordinates, which is precisely how the U_{PMNS} coefficients – that we strive to measure – end up in the equations.

We hereby present the simpler quantum-mechanical derivation of the oscillation probabilities in vacuum, an approach based on quantum field theory can be studied in [47]. To do so, we start with a flavour eigenstate $|\nu_\alpha\rangle$ with $\alpha \in \{e, \mu, \tau\}$, produced at the space-time origin

$$|\nu_\alpha(0)\rangle = |\nu_\alpha\rangle \,. \tag{2.37}$$

In order to easily propagate this state to another location, let us first expand it with respect to the mass eigenstates $(\nu_k)_{k\in[\![1,3]\!]}$. As underscored by (2.29), we obtain

$$|\nu_\alpha(0)\rangle = \sum_{k=1}^{3} U_{k\alpha}^{-1} |\nu_k\rangle = \sum_{k=1}^{3} {U_{\alpha k}}^* |\nu_k\rangle \,. \tag{2.38}$$

Assuming that the mass eigenstates follow the time-dependent Schrödinger equation with no potentials – as is the case in vacuum – at a different point x in space-time, for all $k \in [\![1, 3]\!]$, we have

$$|\nu_k(x)\rangle = e^{-ip_k\cdot x} |\nu_k\rangle \,, \tag{2.39}$$

where p_k is the four-momentum of the k-th state. Thus, the state $|\nu_\alpha(x)\rangle$ evolves as

$$|\nu_\alpha(x)\rangle = \sum_{k=1}^{3} {U_{\alpha k}}^* e^{-ip_k\cdot x} |\nu_k\rangle \,. \tag{2.40}$$

Detecting $|\nu_\alpha(x)\rangle$ amounts to projecting it on a flavour eigenstate $|\nu_\beta\rangle$ with $\beta \in \{e, \mu, \tau\}$. For this reason, it is suitable to expand the $|\nu_k\rangle$'s back to the flavour basis, which, for once, does rely on U_{PMNS}, from which we can directly read the column vectors introduced in (2.25). When doing so, we find

$$|\nu_\alpha(x)\rangle = \sum_{k=1}^{3} \sum_{\gamma=e,\tau,\mu} {U_{\alpha k}}^* U_{\gamma k} e^{-ip_k\cdot x} |\nu_\gamma\rangle \,. \tag{2.41}$$

Since the flavour eigenstates form an orthonormal set, the probability for a neutrino of a

given flavour α, produced at origin, to be detected with a β flavour at x is

$$\begin{aligned} P_{\nu_\alpha \to \nu_\beta} =& |\langle \nu_\beta | \nu_\alpha(x) \rangle|^2 \\ =& \left| \sum_{k=1}^{3} {U_{\alpha k}}^* U_{\beta k} \; e^{-ip_k \cdot x} \right|^2 \\ =& \sum_{k=1}^{3} \sum_{j=1}^{3} U_{\alpha k}{}^* U_{\beta k} U_{\alpha j} {U_{\beta j}}^* e^{-i(p_k - p_j)\cdot x} \; . \end{aligned} \tag{2.42}$$

It is worthwhile noting that for antineutrinos, the expansion of the flavour eigenstates into the mass basis proceeds through U_{PMNS} directly, which would exchange the conjugate terms in (2.42).

2.2.3.2 Ultra-relativistic approximation

A few dodgy assumptions are now needed to simplify further (2.42) and derive its usual form for neutrino experiments. The masses of the neutrinos are very small compared to the energies at which they are detected, in a nutshell, neutrinos are ultra-relativistic particles[11]. In natural units, this implies that $t \simeq \| \overrightarrow{x} \| = L$, where L is the distance between the neutrino source and the detector (on whose direction the three-momenta $\overrightarrow{p_k}$'s are taken to be aligned). Consequently, the phases can be approximated by

$$(p_k - p_j) \cdot x \simeq \left[(E_k - \| \overrightarrow{p_k} \|) - (E_j - \| \overrightarrow{p_j} \|) \right] L \; . \tag{2.43}$$

Carrying on with the ultra-relativistic assumption, we assume that the energy E, given by the kinematics of the production process neglecting neutrino masses, verifies, for all $k \in [\![1, 3]\!]$

$$E_k - \| \overrightarrow{p_k} \| = \frac{{m_k}^2}{E_k + \| \overrightarrow{p_k} \|} \simeq \frac{{m_k}^2}{2E}, \tag{2.44}$$

so that we can write

$$(p_k - p_j) \cdot x \simeq \frac{\Delta {m_{kj}}^2}{2E} L, \tag{2.45}$$

where $\Delta {m_{kj}}^2 = {m_k}^2 - {m_j}^2$. Inserting (2.45) into (2.42) eventually leads to the probability to detect a neutrino oscillation at a distance L from a source which produces them with a (kinetic) energy E

$$P_{\nu_\alpha \to \nu_\beta}(L, E) = \sum_{k=1}^{3} \sum_{j=1}^{3} U_{\alpha k}{}^* U_{\beta k} U_{\alpha j} {U_{\beta j}}^* \; e^{-i \frac{\Delta m_{kj}^2}{2E} L}. \tag{2.46}$$

[11] We are treating relativistic particles as plane waves, we could hardly do with less sense regarding that plane waves are everywhere but nowhere, nevertheless, imposing coherent contributions of wave packets yields identical results, see [48] for discussions.

2.2.3.3 Oscillation probabilities with trigonometric functions

We can expand (2.46) into a more usable form that has only trigonometric functions. For each space-time position x, regardless of any relativistic approximation, (2.42) and (2.46) bring out a Hermitian matrix in $(k,j) \in [\![1,3]\!]^2$, whose coefficients read

$$h_{kj} = U_{\alpha k}{}^* U_{\beta k} U_{\alpha j} U_{\beta j}{}^* e^{-i(p_k - p_j)\cdot x}. \tag{2.47}$$

Evidently, these satisfy $h_{kj} = h_{jk}{}^*$. It is apposite to take advantage of the Hermitian symmetry and split the integral into diagonal and off-diagonal terms

$$\begin{aligned}\sum_{k=1}^{3}\sum_{j=1}^{3} h_{kj} &= \sum_{k=1}^{3} h_{kk} + \sum_{k=1}^{3}\sum_{j=1}^{k-1} h_{kj} + \sum_{k=1}^{3}\sum_{j=k+1}^{3} h_{kj} \\ &= \sum_{k=1}^{3} h_{kk} + \sum_{k=2}^{3}\sum_{j=1}^{k-1} h_{kj} + \sum_{k=1}^{2}\sum_{j=k+1}^{3} h_{kj}.\end{aligned} \tag{2.48}$$

Up to a conjugate, the last term is identical to the second one, as a result

$$\sum_{k=1}^{3}\sum_{j=1}^{3} h_{kj} = \sum_{k=1}^{3} h_{kk} + 2\sum_{k=2}^{3}\sum_{j=1}^{k-1} \mathfrak{Re}(h_{kj}), \tag{2.49}$$

where $\mathfrak{Re}$ denotes the real part. Along these lines,

$$\begin{aligned}P_{\nu_\alpha \to \nu_\beta}(L,E) &= \sum_{k=1}^{3} |\ U_{\alpha k}|^2 |U_{\beta k}|^2 + 2\sum_{k=2}^{3}\sum_{j=1}^{k-1} \mathfrak{Re}\left(U_{\alpha k}{}^* U_{\beta k} U_{\alpha j} U_{\beta j}{}^*\ e^{-i\frac{\Delta m^2_{kj} L}{2E}}\right) \\ &= \sum_{k=1}^{3} |U_{\alpha k}|^2 |U_{\beta k}|^2 + 2\sum_{k=2}^{3}\sum_{j=1}^{k-1} \cos\left(\frac{\Delta m^2_{kj} L}{2E}\right) \mathfrak{Re}\left(U_{\alpha k}{}^* U_{\beta k} U_{\alpha j} U_{\beta j}{}^*\right) \\ &\quad + 2\sum_{k=2}^{3}\sum_{j=1}^{k-1} \sin\left(\frac{\Delta m^2_{kj} L}{2E}\right) \mathfrak{Im}\left(U_{\alpha k}{}^* U_{\beta k} U_{\alpha j} U_{\beta j}{}^*\right),\end{aligned} \tag{2.50}$$

since $\exp\left(-i\frac{\Delta m^2_{kj} L}{2E}\right) = \cos\left(\frac{\Delta m^2_{kj} L}{2E}\right) - i\sin\left(\frac{\Delta m^2_{kj} L}{2E}\right)$. The unitarity of U_{PMNS} also implies that

$$\left(\sum_{k=1}^{3} U_{\alpha k} U^\dagger_{k\beta}\right)^2 = \delta_{\alpha\beta}. \tag{2.51}$$

Using the same property that yielded (2.49), we obtain

$$\sum_{k=1}^{3} |U_{\alpha k}|^2 |U_{\beta k}|^2 = \delta_{\alpha\beta} - 2\sum_{k=2}^{3}\sum_{j=1}^{k-1} \mathfrak{Re}\left(U_{\alpha k}{}^* U_{\beta k} U_{\alpha j} U_{\beta j}{}^*\right). \tag{2.52}$$

Substituting (2.52) into (2.50), we conclude that

$$\begin{aligned}P_{\nu_\alpha \to \nu_\beta}(L,E) =& \delta_{\alpha\beta} - 4\sum_{k=2}^{3}\sum_{j=1}^{k-1} \sin^2\left(\frac{\Delta m^2_{kj} L}{4E}\right) \mathfrak{Re}\left(U_{\alpha k}{}^* U_{\beta k} U_{\alpha j} U_{\beta j}{}^*\right) \\ &+ 2\sum_{k=2}^{3}\sum_{j=1}^{k-1} \sin\left(\frac{\Delta m^2_{kj} L}{2E}\right) \mathfrak{Im}\left(U_{\alpha k}{}^* U_{\beta k} U_{\alpha j} U_{\beta j}{}^*\right).\end{aligned} \tag{2.53}$$

2.2.3.4 Survival probability

Nuclear power plants, in which we are most interested, are actually generators of antineutrinos. One could go back to 2.2.3.1 and read the comments there that concern antineutrinos to derive the antineutrino version of (2.53), but we can also use some ingenuity and recall that CP transforms left-handed neutrinos into right-handed antineutrinos (cf. 1.2.2), the time operator T, on the other hand, essentially changes the direction of the arrow in the probability $P_{\nu_\alpha \to \nu_\beta}$. To put it differently,

$$P_{\overline{\nu_\alpha} \to \overline{\nu_\beta}} = P_{\nu_\beta \to \nu_\alpha} \tag{2.54}$$

as long as CPT holds, and it does, so far. Recalling that the real part of a complex number is unaffected by conjugation, and that the imaginary part takes a minus sign, we can explicitly write

$$\begin{aligned} P_{\overline{\nu_\alpha} \to \overline{\nu_\beta}}(L,E) =& \delta_{\alpha\beta} - 4 \sum_{k=2}^{3} \sum_{j=1}^{k-1} \sin^2 \left(\frac{\Delta m^2_{kj} L}{4E} \right) \mathfrak{Re} \left({U_{\alpha k}}^* U_{\beta k} U_{\alpha j} {U_{\beta j}}^* \right) \\ & - 2 \sum_{k=2}^{3} \sum_{j=1}^{k-1} \sin \left(\frac{\Delta m^2_{kj} L}{2E} \right) \mathfrak{Im} \left({U_{\alpha k}}^* U_{\beta k} U_{\alpha j} {U_{\beta j}}^* \right) . \end{aligned} \tag{2.55}$$

The survival probability, which determines the chances to project the incoming neutrino back to the flavour state $|\nu_\alpha\rangle$ with which one started, is even simpler than (2.55)

$$P_{\overline{\nu_\alpha} \to \overline{\nu_\alpha}}(L,E) = P_{\nu_\alpha \to \nu_\alpha}(L,E) = 1 - 4 \sum_{k=2}^{3} \sum_{j=1}^{k-1} |U_{\alpha k}|^2 |U_{\alpha j}|^2 \sin^2 \left(\frac{\Delta m^2_{kj} L}{4E} \right) . \tag{2.56}$$

2.3 Parameters of the model

2.3.1 Neutrino masses

2.3.1.1 Individual masses

On top of a non-diagonal mixing matrix, it should now be clearer – with the help of section 2.2 – why neutrino oscillations demand the neutral leptons to have mass. All the fuss about physics beyond the Standard Model comes from the neutrino masses, whose simplest experimental signature is neutrino oscillations. To be fair, at least two of the mass eigenstates $(\nu_k)_{k \in [\![1,3]\!]}$ should have a non-vanishing mass, since the oscillation probabilities boil down to the squared mass differences introduced in (2.45). The individual masses are indeed difficult to access experimentally, and direct measurements of the electron neutrino mass – which is but a U_{PMNS}-weighted average of the mass eigenvalues $(m_k)_{k \in [\![1,3]\!]}$ – relying on the endpoint of the β-decay of tritium, have only been able to set an upper-limit [49]

$$m_{\overline{\nu_e}} = \langle \overline{\nu_e} | M | \overline{\nu_e} \rangle = \sum_{k=1}^{3} |U_{ek}|^2 {m_k}^2 \ < \ 2.05\,\text{eV} \quad (95\%\text{CL}), \tag{2.57}$$

where M designs the mass operator, which is diagonal in the mass basis. Other model-dependent limits, coming from cosmology, have set more stringent limits on the sum of the mass eigenvalues themselves [50]

$$\sum_{k=1}^{3} m_k \ < \ 0.136\,\text{eV} \quad (95\%\text{CL}). \tag{2.58}$$

2.3.1.2 Squared mass differences

As is apparent from (2.53), the frequencies that drive the probability to detect one flavour or another, as a function of L/E, are explicitly related to the squared mass differences, and are responsible for the pattern from Figure 2.3. There are various regimes, determined by the energy of the neutrino source, the location of the detector with respect to the source, and of course, by the distribution of the squared mass differences. Fortunately, the latter are well-separated and allow experiments to focus on one particular set of values. Echoing back to our comments from 2.1.3.3, that is precisely because they can adjust relatively well the L/E ratio – hence overshadowing irrelevant terms in the oscillation probabilities – that man-made neutrinos particularly shine when it comes to estimating certain parameters.

From solar neutrinos[12], we have learnt that [51]

$$\Delta m^2_{21} = (7.53 \pm 0.18) \times 10^{-5}\,\text{eV}^2. \tag{2.59}$$

On the other hand, atmospheric neutrinos[13] have us wavering between the two conflicting values that follow [51]

$$\Delta m^2_{32} = (2.44 \pm 0.06) \times 10^{-3}\,\text{eV}^2 \tag{2.60}$$

$$\text{OR} \quad \Delta m^2_{32} = -\,(2.49 \pm 0.06) \times 10^{-3}\,\text{eV}^2. \tag{2.61}$$

How come we seem to only have knowledge about the absolute value $|\Delta m^2_{32}|$, which cannot be distinguished from $|\Delta m^2_{31}|$? Why are there two different numbers for Δm^2_{32} ? As it happens, we do not know yet if the two close mass eigenstates, ν_1 and ν_2, are much lighter than ν_3, or if there actually are two "heavy" neutrinos out there. The first situation, in which $\Delta m^2_{32} > 0$, is referred to as the "normal mass hierarchy", the second one corresponds to $\Delta m^2_{32} < 0$ and goes by the name of "inverted mass hierarchy". To get a better grasp of these ambiguities, we must look for bare sines in (2.53) and (2.55), which are the sole terms that can tell us about the signs of the squared mass differences. From (2.56), it is clear that no such information can effortlessly come from the study of neutrino survival, for one would have to endeavour to see sub-leading order differences between Δm^2_{32} and Δm^2_{31}[14]. Unfortunately, in apparition

[12] KamLAND is effectively looking at the solar regime although the lightweight leptons it observes come from nuclear reactors.

[13] Accelerator neutrinos, observed by MINOS and K2K, shed light on the same parameters as atmospheric neutrinos.

[14] By and large, the JUNO experiment will actually try to understand if $|\Delta m^2_{32}| > |\Delta m^2_{31}|$, which would correspond to a normal ordering.

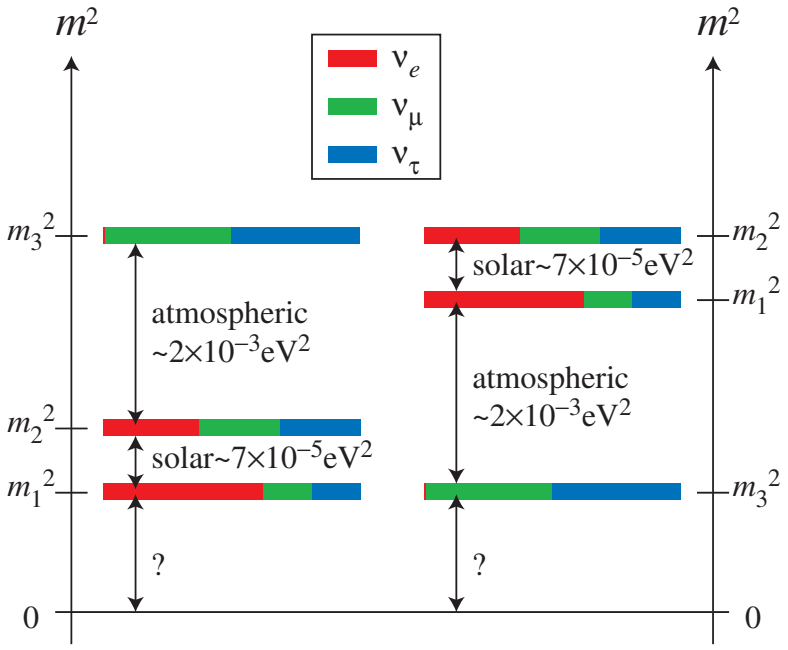

Figure 2.4 – Visual summary of the knowledge about neutrino oscillations. On the left-hand side, the normal mass hierarchy is assumed, on the right, the inverted ordering is presented. The colourful bars indicate the components of the mass eigenstates in the flavour basis. For instance, the state of mass m_1 has its largest component along $|\nu_e\rangle$ whereas $|\nu_3\rangle$ has barely any weight along it, on account of the tininess of θ_{13}. Figure taken from [53].

experiments, which answer for $P_{\nu_\alpha \to \nu_\beta}$ with $\alpha \neq \beta$, the set of model parameters is such that the bare sine terms are also second order corrections. The different behaviour of neutrinos and antineutrinos in matter can, however, help to disentangle the two mass hierarchies [52]. While on the subject, that it is precisely due to the considerable matter effect in the sun that the ordering $m_2 > m_1$ has been resolved.

A visual representation of the knowledge about neutrino masses can be found in Figure 2.4, it presents the two possible mass hierarchies and serves as an appetiser for the next sub-section, which tackles the coefficients of U_{PMNS}.

2.3.2 Matrix coefficients

2.3.2.1 Large mixing angles

The squared mass differences appear in the sines of the oscillation formula (2.53), and accordingly, determine the frequency with which the values of the probabilities are repeated when moving along a L/E axis. On the other hand, the amplitude of the oscillations is fixed by the U_{PMNS} matrix coefficients. As reviewed in 2.2.2.2, the mixing matrix consists of three angles of rotation $\theta_{12}, \theta_{13}, \theta_{23}$ and one phase δ. Nonetheless, the coefficients of the

matrix remain to be predicted by any theory, and we must rely on experiments to set their values[15]. For the solar sector, these experiments have observed [51]

$$\sin^2(\theta_{12}) = 0.304 \pm 0.014, \tag{2.62}$$

and for the atmospheric angle, we have [51]

$$\sin^2(\theta_{23}) = 0.514 \pm 0.056 \tag{2.63}$$

$$\text{OR} \quad \sin^2(\theta_{23}) = 0.511 \pm 0.055, \tag{2.64}$$

where (2.64) is the estimation in the case of an inverted mass hierarchy.

2.3.2.2 Minute mixing angle

With regards to the remaining mixing angle θ_{13}, it has proven to be the most challenging to quantify and it remained unmeasured until just a few years ago. In the pursuit of its measurement, two types of experiments have been undertaken: accelerator experiments and reactor experiments.

The former, such as T2K and MINOS, are relying on ν_e appearance in a ν_μ beam. To put it differently, these experiments are measuring $P_{\nu_\mu \to \nu_e}$, which involves the real and imaginary parts of the terms ${U_{\mu k}}^* U_{ek} U_{\mu j} {U_{ej}}^*$ for $2 \leqslant k \leqslant 3$ and $1 \leqslant j \leqslant k-1$. Thus, these terms are dependent on the unknown δ phase, that will be presented more in-depth in 2.3.2.3. Besides, as underlined in 2.3.1.2, there are second order corrections related to the mass hierarchy in the apparition formula[16].

On the contrary, insofar as the survival formula (2.56) bears only moduli, disappearance experiments are utterly agnostic as to whether the U_{PMNS} matrix is complex or purely real, i.e. they are independent of whether $\delta \neq 0$ or $\delta = 0$. Moreover, inasmuch as there are only squared sines in (2.56), they are, for the most part, free from assumptions on the mass hierarchy provided that $\Delta m^2_{32} \simeq \Delta m^2_{31}$. Therefore, it is scarcely surprising that reactor experiments such as Daya Bay, RENO, and Double Chooz, have set the most stringent bounds on the value of $\sin^2(\theta_{13})$.

Electron antineutrino survival in the vicinity of a nuclear power plant

Let us explicitly write the survival probability for the $\overline{\nu_e}$'s produced in nuclear reactors, and detected a few hundreds of meters farther, via inverse β-decay (1.6). Expanding (2.56), we

[15] In that regard, the neutrino masses are scarcely different since they require fine-tuned Yukawa couplings.

[16] Let us emphasise again that T2K and MINOS, on top of looking for ν_e appearance, are set on measuring θ_{23} and Δm^2_{32} by ν_μ disappearance and think anyhow bigger than reactor experiments.

get

$$\begin{aligned} P_{\overline{\nu_e}\to\overline{\nu_e}}(L,E) = 1 &- 4|U_{e2}|^2|U_{e1}|^2\sin^2\left(\frac{\Delta m^2_{21}L}{4E}\right) \\ &- 4|U_{e3}|^2|U_{e1}|^2\sin^2\left(\frac{\Delta m^2_{31}L}{4E}\right) \\ &- 4|U_{e3}|^2|U_{e2}|^2\sin^2\left(\frac{\Delta m^2_{32}L}{4E}\right). \end{aligned} \tag{2.65}$$

Recalling that

$$|U_{e1}|^2 = \cos^2(\theta_{12})\cos^2(\theta_{13}) \simeq \cos^2(\theta_{12}) = 1-\sin^2(\theta_{12}) \tag{2.66}$$

$$|U_{e2}|^2 = \sin^2(\theta_{12})\cos^2(\theta_{13}) \ \simeq \sin^2(\theta_{12}) \tag{2.67}$$

$$|U_{e3}|^2 = \sin^2(\theta_{13})\ , \tag{2.68}$$

and in light of the value from (2.62), it may seem as though the last two terms in (2.65) are negligible compared to the first. That would be too hasty a judgement, for we must first discuss the value of the phase. As such, rewriting back the c^2 next to each mass, with c being the speed of light, the phases are in MeV · fm or any equivalent unit. To cut the matter short, the phases are missing $\hbar c$, with $\hbar$ the reduced Planck constant, and read in conventional units

$$\frac{\Delta m^2 c^4 L}{4\hbar c E} \simeq 1.27\times 10^6\,\mathrm{eV}^{-1}\mathrm{m}^{-1}\frac{\Delta m^2 c^4 L}{E}. \tag{2.69}$$

In reactor experiments looking for θ_{13}, typically $E = 3\,\mathrm{MeV}$ and $L = 10^3\,\mathrm{m}$, so we may benefit from writing

$$\frac{\Delta m^2 c^4 L}{4\hbar c E} \simeq 4.2\times 10^2\,\mathrm{eV}^{-2}\Delta m^2 c^4. \tag{2.70}$$

Considering the measurements from 2.3.1.2, it is plain to see that the first sine falls in the approximation $\sin^2 x \sim x^2$ with $x \simeq 3\times 10^{-2}$ and is consequently three orders of magnitude smaller than the last two sines, whose argument is around 1. All things considered, in our case, we have

$$P_{\overline{\nu_e}\to\overline{\nu_e}}(L,E) \simeq 1-\sin^2(2\theta_{13})\left[\cos^2(\theta_{12})\sin^2\left(\frac{\Delta m^2_{31}L}{4E}\right)+\sin^2(\theta_{12})\sin^2\left(\frac{\Delta m^2_{32}L}{4E}\right)\right], \tag{2.71}$$

since $4\sin^2(\theta_{13})\cos^2(\theta_{13}) = \sin^2(2\theta_{13})$. To shrink (2.71) further down, we consider that $\Delta m^2_{32} \simeq \Delta m^2_{31}$. It follows that the oscillation probability relevant for moderately short baseline reactor experiments is well-approximated by

$$P_{\overline{\nu_e}\to\overline{\nu_e}}(L,E) \simeq 1-\sin^2(2\theta_{13})\sin^2\left(\frac{\Delta m^2_{31}L}{4E}\right). \tag{2.72}$$

In accordance, a detector located 1 km away from a nuclear power plant must show signs of an energy-dependent deficit in the neutrino spectrum it observes, as epitomised by Figure 2.5.

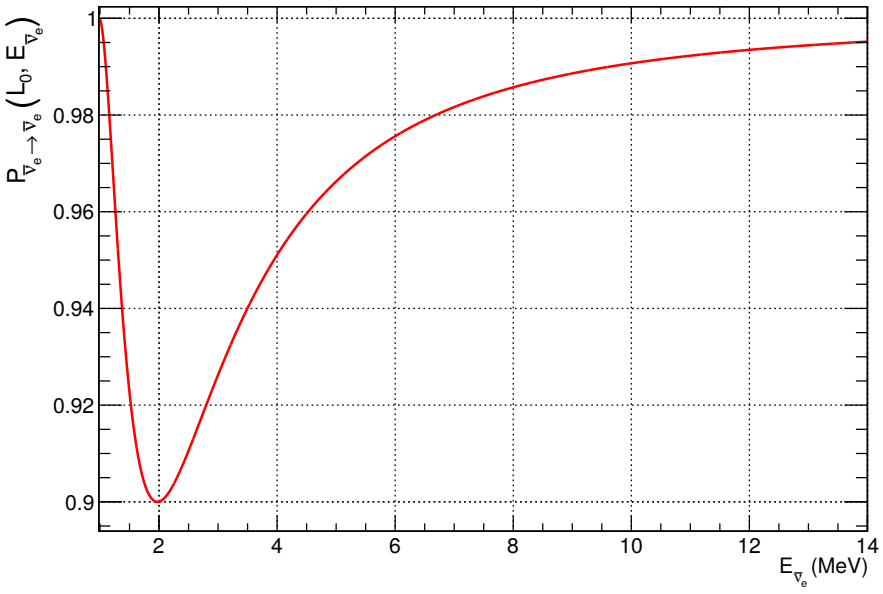

Figure 2.5 – Electron antineutrino survival probability $E_{\overline{\nu_e}} \longrightarrow P_{\overline{\nu_e}\to\overline{\nu_e}}(L_0, E_{\overline{\nu_e}})$ at a distance $L_0 = 1\,\text{km}$ from the source. The kinetic energy of the antineutrino is designated by $E_{\overline{\nu_e}}$. The textbook case $\sin^2(2\theta_{13}) = 0.1$ is assumed.

Experimental values

In spite of the robustness of (2.72), the CHOOZ experiment – located on the current "far site" of Double Chooz (more details will follow in Chapter 3) – could but set a lower limit on the smallness of θ_{13}. Indeed, mostly plagued by liquid scintillator degradation and accidental background from several origins, at the turn of the twenty-first century, CHOOZ reported $\sin^2(2\theta_{13}) < 0.17$ (90%CL) [54].

Considerably improving the detector design used by CHOOZ, the Double Chooz collaboration showed the first indication of $\overline{\nu_e}$ disappearance in January 2012 [55], reporting $\sin^2(2\theta_{13}) = 0.086 \pm 0.041(\text{sys.}) \pm 0.030(\text{stat.})$. Notwithstanding the excellence of the Double Chooz analysis, the irrefutable evidence came from the China-based Daya Bay experiment, a few months later. In April 2012, by means of two experimental halls near Hong Kong, at distinct flux-weighted distances from the nuclear cores $L_n \simeq 500\,\text{m}$ and $L_f \simeq 1600\,\text{m}$, the Daya Bay collaboration indeed issued [56]

$$\sin^2(2\theta_{13}) = 0.092 \pm 0.016(\text{sys.}) \pm 0.05(\text{stat.}). \tag{2.73}$$

Comparing the spectra from the detectors of the "near hall" (at which the survival probability $P_{\overline{\nu_e}\to\overline{\nu_e}}(L_n, E) \simeq 1$ for all E) to that of the "far hall" (located where the survival probability $P_{\overline{\nu_e}\to\overline{\nu_e}}(L_f, 3\,\text{MeV}) < 1$), they obtained Figure 2.6, which bears gratifying resemblance[17] to

[17] Daya Bay's far hall actually corresponds to $L_f \simeq 1.6\,\text{km} \neq 1\,\text{km}$ so the oscillation maximum is expected at $E_{\overline{\nu_e}} \simeq 3.2\,\text{MeV}$. And yet, Figure 2.6 uses the prompt energy, which is about 0.8 MeV lower than $E_{\overline{\nu_e}}$, thus, it is hardly surprising that both curves look alike.

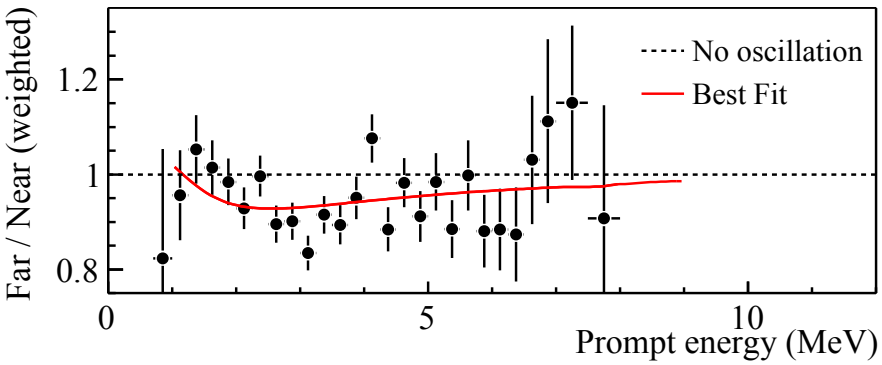

Figure 2.6 – Ratio of the far and near hall Daya Bay spectra. To all intents, the near hall spectrum is a no-oscillation prediction. The solid curve is the best-fit solution with $\sin^2(2\theta_{13}) = 0.092$, obtained from a so-called rate-only analysis, which integrates the deficit over all energies. Figure extracted from [56].

Figure 2.5. Such a differential measurement, using a "near" and a "far site", is actually the key to great accuracy. This method was first advocated by Double Chooz, at the end of the year 2002, to cancel the systematic errors originating from the prediction of the reactor spectra and to dwarf the detector-related uncertainties [57].

It is worthwhile noticing that the size of the deficit along the energy axis is driven by the phase $\Delta m^2_{31} L/4E$. In 2015, taking advantage of more statistics, Daya Bay thus provided an interesting measurement of m^2_{31} [18] [58].

The Korean RENO experiment confirmed Daya Bay's θ_{13} value shortly after, in May 2012, with a comparable 4.9σ significance [59]. Although Double Chooz latest's paper [60], uses only the "far" experimental site, and exhibits a lower 3σ significance than its competitors, with $\sin^2(2\theta_{13}) = 0.090^{+0.032}_{-0.029}$, the collaboration has not kept idle hands. Not only did it build its near detector, which is put to good use in this thesis, but it also developed an unprecedented understanding of its detectors, as we shall explain in Chapter 6.

A best-fit from global analyses published up to 2015 yields [51]

$$\sin^2(2\theta_{13}) = 0.085 \pm 0.050 \tag{2.74}$$
$$\sin^2(\theta_{13}) = (2.19 \pm 0.12) \times 10^{-2}. \tag{2.75}$$

2.3.2.3 CP violation phase

To this day, the most enigmatic parameter is unmistakably the CP violation phase δ. It is called that way because if it were zero, the imaginary parts would be but nought in the

[18] Actually, access to $\Delta m^2_{ee} \simeq m^2_{31} \simeq m^2_{32}$ – which can effortlessly be defined from the square bracket in (2.71) – is offered by electron antineutrino survival.

oscillation probabilities, and the difference between (2.53) and (2.55)

$$\left(P_{\nu_\alpha\to\nu_\beta} - P_{\overline{\nu_\alpha}\to\overline{\nu_\beta}}\right)(L,E) = 4\sum_{k=2}^{3}\sum_{j=1}^{k-1}\sin\left(\frac{\Delta m^2_{kj}L}{2E}\right)\mathfrak{Im}\left(U_{\alpha k}{}^{*}U_{\beta k}U_{\alpha j}U_{\beta j}{}^{*}\right). \tag{2.76}$$

would vanish identically. Put another way, should $\delta = 0$, the behaviour of neutrinos and antineutrinos would be identical, and there would be no CP violation indications coming from oscillation experiments. Understanding CP violation in the neutrino sector would fill a piece in the leptogenesis puzzle, and enlighten us as to why our universe is mostly made of electrons, and not of positrons, which is a bit akin to understanding why there are so few left-handed persons in our societies[19].

As is obvious in (2.36), θ_{13} has no ordinary position in the U_{PMNS} matrix. In fact, the mixing matrix has been parametrised in such a way that the smallest mixing angle is in front of the CP violation phase [46]. The smaller θ_{13}, the more difficult δ is to measure. Nevertheless, the values from (2.75) are actually large, and certainly not much smaller than CHOOZ's limit. Consequently, the very same accelerator experiments that are sensitive to θ_{13}, but hindered by corrections related to the unknown mass hierarchy and the δ value, may use the input from reactor experiments to better assess and constrain theses corrections. With encouraging results already published at the beginning of the current year, NOνA swells the list of experiments looking at ν_e appearance in a ν_μ beam [61]. Using a classical Likelihood event selector, they disfavour $0.1\pi < \delta < 0.5\pi$ for the inverted mass hierarchy at 90% CL whereas a new-fangled Library Event Matching classifier provides much bolder results, disfavouring all δ values in the case of an inverted ordering, and in this manner, the inverted mass hierarchy altogether (see Figure 2.7).

Of course, all these results are dependent on the value of θ_{13}, which is precisely the reason why the inputs handed over to these new undertakings must be cross-checked by several experiments. Beyond the shadow of a doubt, both the provided central values and the errors are of paramount importance.

[19] Although shooting footballs with the left foot, as well as wielding a battle-axe in the left hand, should boost the survival probability of individuals with the quaint handedness.

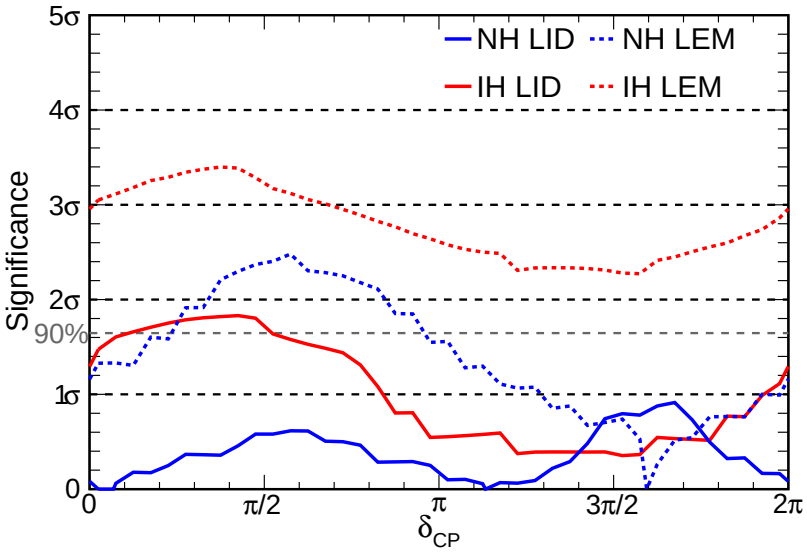

Figure 2.7 – Significance of the difference between the selected and predicted number of events as a function of δ and the mass hierarchy (designated by NH or IH). The disagreement with the observed data is shown in solid lines for the primary Likelihood Identifier (LID), and in dotted lines for the secondary Library Event Matching (LEM) classifier. Figure taken from [61].

Part II

Neutrinos in Double Chooz

Chapter 3

$\overline{\nu_e}$ production at Chooz

As long as particle physicists trade in patience with the safety authorities, nuclear power plants are inexpensive (anti)neutrino sources. All things considered, they are remarkably easier to get by than thermonuclear weapons. Besides, commercial nuclear reactors provide a much steadier flux than the latter or potentially malfunctioning accelerators. For indeed, the money-makers know how willing the folks are to light up their households and blast music at night, and accordingly, try hard to keep the machine roaring. In addition, fuel reloading does shrink the emitted antineutrino rate in a very predictable way. All that being said, the main strength of reactor experiments undeniably lies in the copiousness and purity of this electron antineutrino rate, which the reactors pour isotropically at the near and far detectors.

3.1 Production site

The 757-inhabitant Chooz village – located in northern France, slyly protruding into Belgium, by the Meuse river (cf. Figure 3.1)– has a long history in pioneering nuclear engineering. Indeed, following the American guidelines, Chooz was the target of choice to build the first Pressurised Water Reactor (PWR) in France, in 1967, with a 320 MW generating capacity. Later, the picturesque village welcomed on its banks the first powerful 1450 MW PWR's, B1 and B2, in 1996 and 1997, respectively.

On top of being a first-class nuclear power site, over the last two decades, Chooz has had tight links with research in particle physics. Before the new generation reactors B1 and B2 were built, the Chooz experiment took advantage of the vast network of tunnels, at the Chooz A site – where the former 320 MW nuclear reactor entered the decommissioning phase in 2001 – to set up its 5 t liquid scintillator antineutrino detector, 100 m underground.

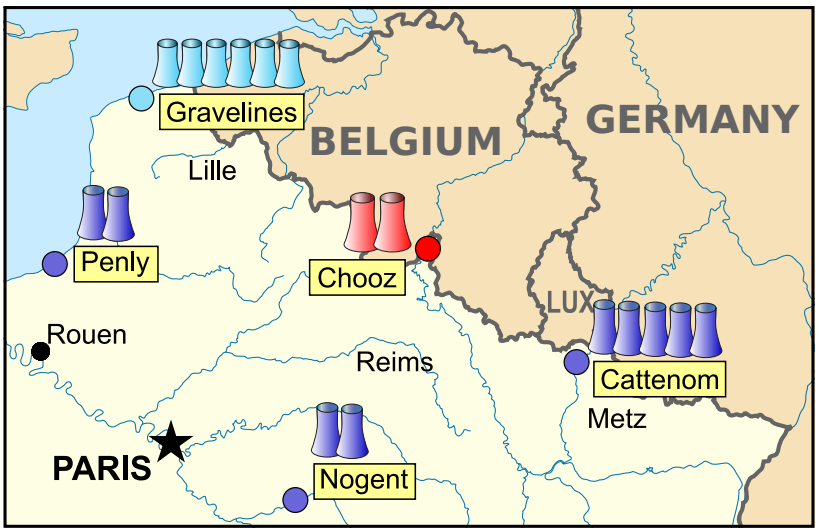

Figure 3.1 – Location of the $\overline{\nu_e}$ factory for the Double Chooz experiment. In red, the two cooling towers of the 2×1450 MW Chooz nuclear power station. For information purposes only, older nuclear reactors are also shown in light blue (900 MW cores) and in purple (1300 MW cores).

3.2 Nuclear fission

3.2.1 Overcoming the Coulomb barrier

Predicting the antineutrino production of a nuclear reactor is no leisurely stroll, for fission itself, is no simple matter.

Nuclear fission is a process in which the nucleus of an atom breaks into lighter nuclei, referred to as "fission fragments". The process can be initiated by a nuclear reaction, such as neutron capture; it may also occur spontaneously, as a usual decay. Fission primarily results from the competition between the nuclear binding force, which increases roughly in proportion to A, the mass number, and the Columb repulsion of protons, growing faster as Z^2, with Z the atomic number. In other words, the higher the ratio Z^2/A, the easier it is for the nucleus to split apart. Naively, if the unstable nucleus ${}^{A}_{Z}\mathrm{U}$ were to split into two equal-mass fragments ${}^{A/2}_{Z/2}\mathrm{V}$, the characteristic ratio would be divided in two for both fragments, hence the increase in stability, at least with respect to fission. To this end, the fission fragments must first overcome the Coulomb barrier, which inhibits spontaneous fission (see Figure 3.2) in way analogous to α-decay.

The absorption of a relatively small amount of energy, however, forms an intermediate state, which is above the Columb barrier, so that fission occurs readily. The absorbed energy may not be too large though, for the cross section decreases with energy: the slower the incoming neutron, the higher the probability to interact with the nucleus. Isotopes with an odd mass but an even atomic number willingly welcome another neutral nucleon, by

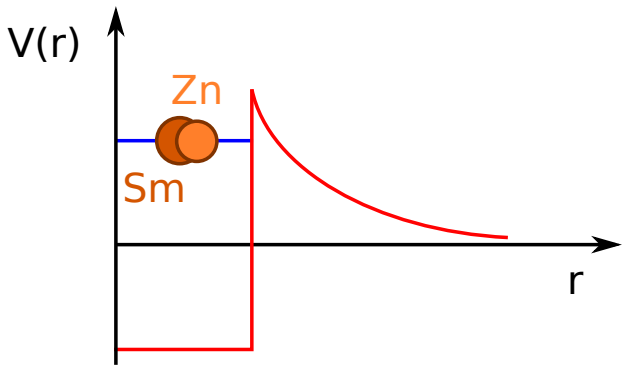

Figure 3.2 – Unless it receives an additional energy from fast neutrons, ^{238}U does not come apart. The two fission fragment candidates, ^{159}Sm and ^{79}Zn, cannot overcome the Coulomb barrier, and are thus trapped by the nuclear potential V, at small r, with r the separation distance between the fission fragments.

means of the pairing force, to which the large fission cross sections of ^{235}U and ^{239}Pu owe much. An example of a fission process, induced by neutron capture, is given in (3.1)

$$^{235}\mathrm{U} + n \rightarrow {}^{236}\mathrm{U}^{*} \rightarrow {}^{141}\mathrm{Cs} + {}^{93}\mathrm{Rb} + 2n\ . \tag{3.1}$$

3.2.2 Chain reaction in nuclear reactors

Nuclear reactors indulge in the latter type of mechanism, i.e. neutron capture, and there is many a reaction similar to (3.1) happening in the reactor cores. In accordance, the fission products are not always ^{141}Cs and ^{93}Rb and there actually is a whole distribution of them, which shows two distinct peaks (see Figure 3.3).

Taking into account Figure 3.3, one can show that when uranium is bombarded by a neutron flux in a PWR, its fission is accompanied by the emission of an average of 2.4 neutrons [63]. Out of these 2.4 neutrons, all but one of them is to be lost, either by being absorbed by another nucleus which cannot undergo fission, or by leaving the reactor core. If each fission is accompanied by the effective release of exactly one neutron, the process is self-sustaining, considering that uranium fission is induced by one neutron, as exemplified by (3.1). If too few neutrons are lost, the fission is explosive. If too many are absorbed or spill out, the reaction is just a wet firecracker.

Inasmuch as the fission cross-sections decrease with energy, maintaining a self-sustaining chain reaction implies lowering the energy of the released neutrons. To this end, the nuclear fuel is placed in a neutron moderator. Ideally, out of all choices, the moderator should have a mass closest to $A = 1$ (so that each collision adequately slows the neutrons) while meagrely capturing neutrons. In PWR's, a pressure of 155 bar [63] ensures that ordinary water, the moderator, remains liquid well above its atmospheric boiling point, thereby retaining its capacity to slow neutrons down and transfer the heat from the reactor cores to electrical

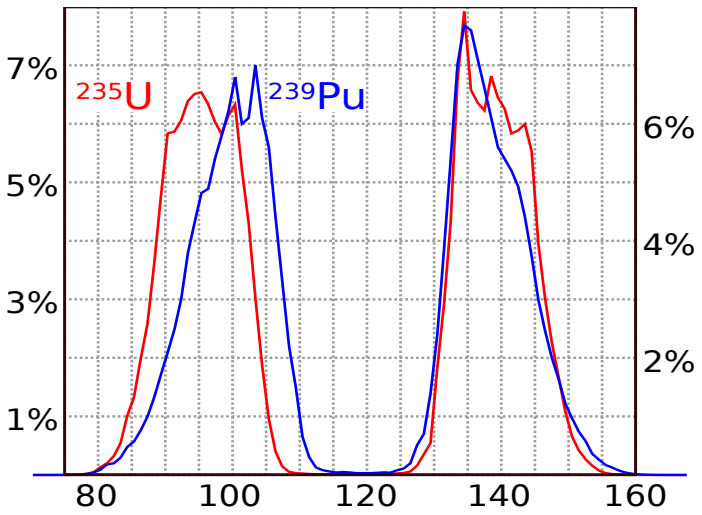

Figure 3.3 – Mass distribution of the fission fragments when fissioning ^{235}U (red) or ^{239}Pu (blue) under a thermal neutron flux. The mass on the x-axis is expressed in terms of the mass number A. Curves based on [62].

generators, with a larger heat capacity than steam would. Unfortunately, light water (H_20), as opposed to heavy water (D_20), has a high neutron-capture cross section[1], which produces deuterium. Making up for these captures entails preparing a special fuel.

3.2.3 Reactor fuel composition

By and large, natural uranium is made up of 99.27% of ^{238}U, 0.72% of ^{235}U and traces of ^{234}U, which are not relevant to the prediction of the antineutrino rate.

Echoing back to our comments concerning the pairing force, the heavier ^{238}U is not easily split. In fact, it is branded as a "fissionable" isotope, whose fission cross-section only rises past 1 MeV (see Figure 3.4), but even with a heavy moderator, the emitted neutrons are not that fast[2]. This leaves us with ^{235}U, the fissile material, whose fission cross-section at thermal energies (typically 25 meV) is more than five orders of magnitude larger than that of ^{238}U.

As alluded in 3.2.2, the natural abundance of ^{235}U will not serve and PWR's demand to burn fuel enriched in a fissile isotope. Enrichment usually involves turning the extracted ore into uranium hexafluoride UF_6 so as to pull the lighter ^{235}U gas molecules into the centre of a centrifuge, thereby leaving the heavier ^{238}U closer to the edges. Out of fashion methods include gaseous diffusion: the lighter isotope diffuses more rapidly, but thousands of passages through the diffuser are necessary to deliver fuel usable in thermal nuclear reactors.

[1] Heavy water has a lower capture cross-section, but when it does trap a neutron, it produces tritium (^{3}H), a radioactive isotope with which the Canadian facilities must cope.

[2] The so-called fast neutrons reactors, rather than trying to fission ^{238}U, usually breed fissile material via neutron capture on the heavy isotope.

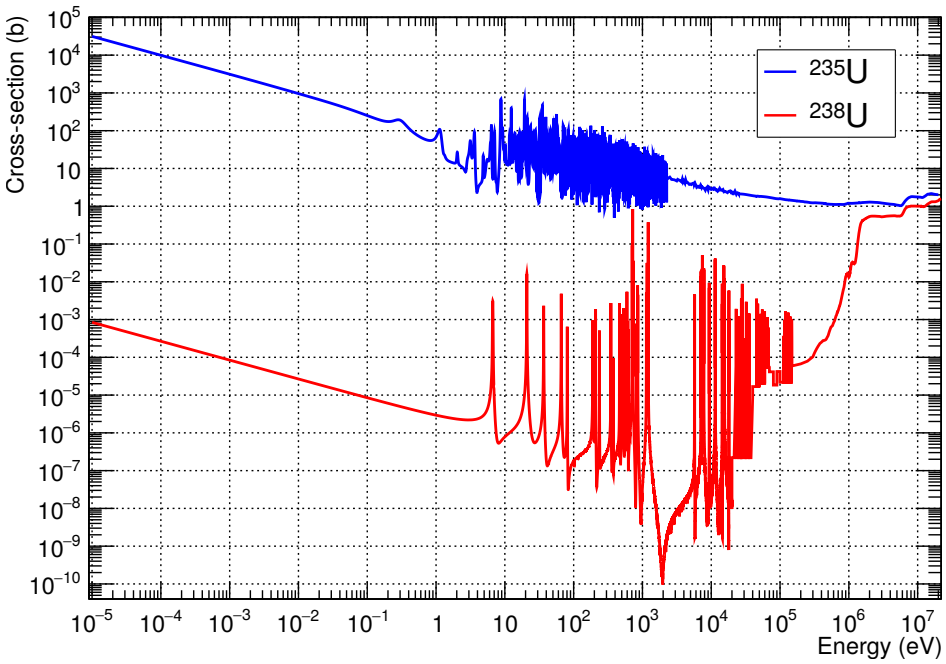

Figure 3.4 – Fission cross-sections of the fissile ^{235}U (blue) and the fissionable ^{238}U (red). In conventional nuclear reactors, the tiny cross-section of ^{238}U below 1 MeV prevents it from maintaining a chain reaction. Evaluated nuclear data extracted from [64].

The fuel feeding the Chooz reactors comes in the form of pellets of uranium dioxide (UO_2), which have been enriched in ^{235}U to between 1.8% and 4% [63], though most assemblies are 3.94% and 4%. PWR's are designed to run on 12 to 18-month cycles and a third of the core is changed at each refuelling outage, which undoubtedly has an impact on the neutrino rate and spectrum. Indeed, the main fuel component, ^{238}U, aside from being fissionable, is known above all as a fertile material. In a nutshell, after a certain period of neutron irradiation in the reactor core, ^{238}U can be converted into fissile material. Breeding proceeds as follows

$$^{238}\mathrm{U} + n \rightarrow {}^{239}\mathrm{U} \xrightarrow[23\,\mathrm{min}]{\beta} {}^{239}\mathrm{Np} \xrightarrow[2.3\,\mathrm{d}]{\beta} {}^{239}\mathrm{Pu}, \tag{3.2}$$

hence involving one neutron capture and two β-decays. Thereafter, ^{239}Pu may capture another neutron

$$^{239}\mathrm{Pu} + n \rightarrow {}^{240}\mathrm{Pu}, \tag{3.3}$$

which in turn, forms up the fissile ^{241}Pu isotope, following

$$^{240}\mathrm{Pu} + n \rightarrow {}^{241}\mathrm{Pu}. \tag{3.4}$$

Out of all the plutonium isotopes that can be produced in a nuclear reactor, ^{239}Pu has the longest half-life, which makes it easy to manipulate[3]. In addition, ^{239}Pu scarcely fissions

[3]The half-life of ^{241}Pu is about 14.3 y whereas ^{239}Pu's exceeds 24×10^3 y.

spontaneously, unlike the heavier ^{240}Pu, whose natural activity could set off reactions prematurely. Last but not least, ^{239}Pu emits more neutrons per fission than ^{235}U, thus reducing the critical mass to detonate thermonuclear weapons. To conclude, ^{239}Pu is the silver bullet of nuclear warfare, and accordingly, a material whose production must be adamantly overseen. In another life, we would have had the time to present our endeavours to correlate the evolution of the reactor fuel, i.e. the accumulation of ^{239}Pu in the Chooz reactors, to the detected $\overline{\nu_e}$'s.

3.3 $\overline{\nu_e}$ release

The lighter nuclei produced by fission are neutron-rich, far from the nuclear valley of stability (cf. Figure 3.5), and thus strive to shed neutrons by undergoing negative β-decays. While the fission products β-decay to reach the stable isobars, as in (3.5) and (3.6)

$$^{141}\mathrm{Cs} \to {}^{141}\mathrm{Ba} \to {}^{141}\mathrm{La} \to {}^{141}\mathrm{Ce} \to {}^{141}\mathrm{Pr}, \tag{3.5}$$

$$^{93}\mathrm{Rb} \to {}^{93}\mathrm{Sr} \to {}^{93}\mathrm{Y} \to {}^{93}\mathrm{Zr}, \tag{3.6}$$

additional energy is released in the reactor cores, in the form of e^-, γ and $\overline{\nu_e}$'s.

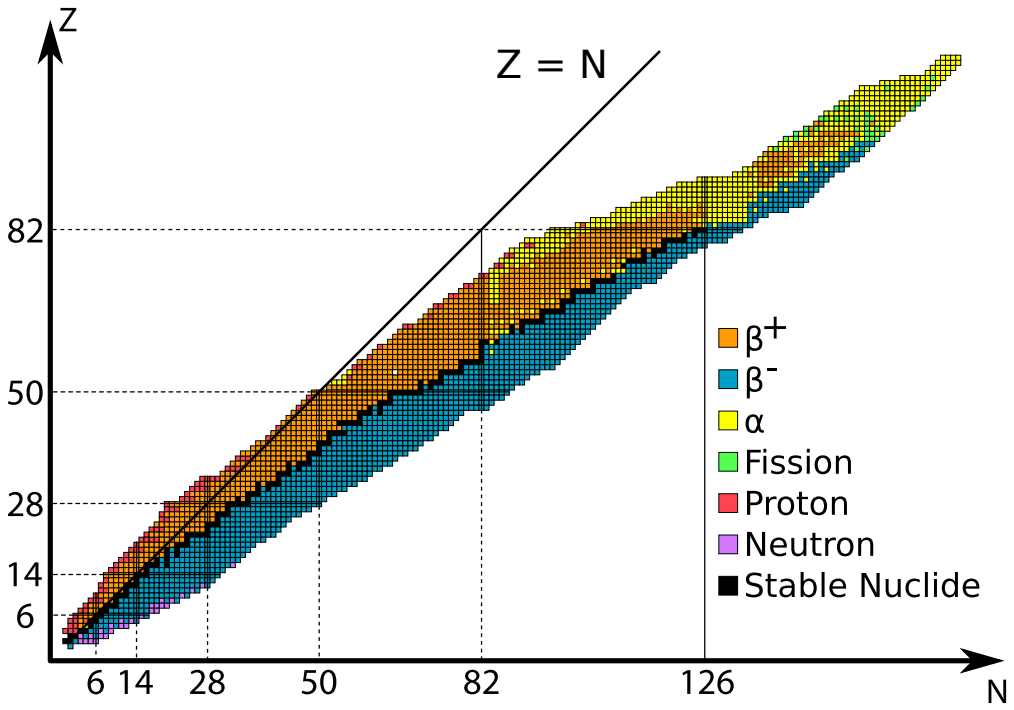

Figure 3.5 – Chart of nuclides. The stability valley is indicated by the black squares. The legend indicates the type of decay that the nuclides are likely to undergo. Inasmuch as they have $(N, Z) \simeq (85, 55)$ and $(N, Z) \simeq (60, 35)$, fission products in a PWR belong to the blue region, which musters the emitters of β^- and $\overline{\nu_e}$'s. Figure produced by means of the Python chart drawer [65].

Obviously, the $\overline{\nu_e}$'s barely interact in the power plant and steadfastly head into all directions, some of them ending up in the liquid scintillator of the Double Chooz detectors. More than 99.7% of the thermal power in the Chooz reactors is related to the fission of the natural uranium isotopes (^{235}U and ^{238}U) and the plutonium bred by neutron capture (^{239}Pu and ^{241}Pu) [63]. The average energies $\langle E_{\overline{\nu_e}} \rangle$ that the $\overline{\nu_e}$'s steal away for each fission of these four nuclei, along with the energies retrieved in the reactor E_f, are displayed in Table 3.1.

Isotope	$\langle E_{\overline{\nu_e}} \rangle$ (MeV)	E_f (MeV)
^{235}U	9.06 ± 0.13	202.36 ± 0.26
^{238}U	10.85 ± 0.39	205.99 ± 0.52
^{239}Pu	7.41 ± 0.18	211.12 ± 0.34
^{241}Pu	8.42 ± 0.12	214.26 ± 0.33

Table 3.1 – Average energy carried away by $\overline{\nu_e}$'s and energy released in the reactor cores for each fission of the considered isotope. Values from [66].

In consequence, fission events always result in electron antineutrinos, even though these particles are not produced directly by the fission event itself. As implied by Figure 3.3, the stunning variety of the fission products makes predicting a most accurate antineutrino rate a great endeavour, let alone spectra, a task left for Chapter 5. Be that as it may, on average, each fission product undergoes three β-decays, as shown on Figure 3.6.

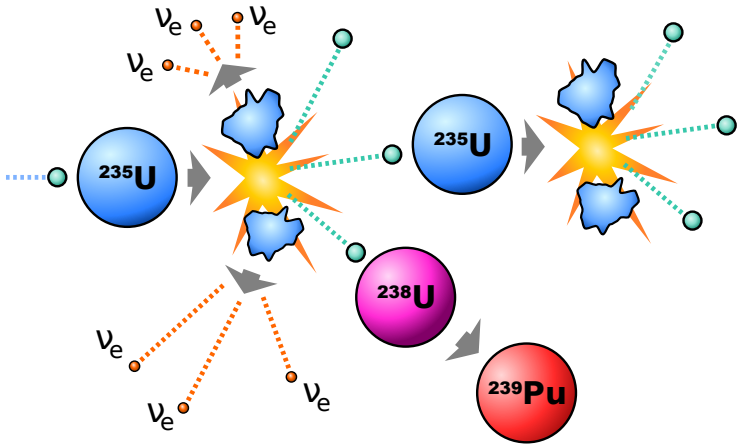

Figure 3.6 – The neutron-induced fission chain reaction in a nuclear reactor is responsible for the unflinching $\overline{\nu_e}$ flux. On average, each fission product undergoes three negative β-decays to reach a sufficiently stable state. Building up of ^{239}Pu in the cores – by dint of neutron-capture on ^{238}U – is also alluded. The grey arrows are short-cuts to otherwise non-trivial processes.

Accordingly, each fission releases a mean number of six electron antineutrinos, which we

denote $\langle n_{\overline{\nu_e}} \rangle$. The two cores – close to each other – produce a maximum thermal power of $P_{th} = 2 \times 4.25\,\mathrm{GW_{th}}$, considering an ever so slightly underestimated energy per fission of $E_f = 200\,\mathrm{MeV}$, we obtain the order of magnitude of the emitted $\overline{\nu_e}$ rate

$$\langle r_{\overline{\nu_e}} \rangle = \frac{\langle n_{\overline{\nu_e}} \rangle\, P_{th}}{E_f} \simeq 1.6 \times 10^{21}\,\mathrm{s}^{-1}. \tag{3.7}$$

Obviously, these $\overline{\nu_e}$'s spread into a 4π solid angle and they will simply be remembered as a $10^{10}\,\mathrm{s}^{-1}\mathrm{cm}^{-2}$ luminosity at 1 km, that is to say a few orders of magnitude higher than the boron-induced solar neutrino flux, thus suitably vouching for the whole reactor experiment trend.

Chapter 4

$\overline{\nu_e}$ detection

When it comes to stopping the $\overline{\nu_e}$'s pouring out of nuclear power stations, the main detection method is no imaginative process, inverse beta decay has been writing the history of antineutrinos for decades, and it must still fulfil its purpose. On the other hand, as hinted by the limitations of the CHOOZ experiment at the turn of the millennium, room remained for improvement in detector design.

The $\overline{\nu_e}$'s effortlessly go through the stainless steel vessels of the reactors and the heavy concrete shielding of the containment walls. Considering the etherealness of the neutral leptons, neutrino experiments are bound to be low-background undertakings, and if building a detector with low-radioactivity materials and considerable shielding is necessary, a distinguishable reaction is a must-have.

4.1 Reaction

4.1.1 Signature

The inverse β-decay (IBD) reaction, already set forth in 1.1.3 as the charged current interaction helping C. L. Cowan and F. Reines to detect the poltergeist, meets the discernibility requirements with flying colours. For the sake of simplicity, it is repeated below with the right flavour

$$\overline{\nu_e} + p \rightarrow n + e^+. \tag{4.1}$$

Albeit more details are to come in 4.3, it is apposite to state that the Double Chooz detectors are filled with organic liquid scintillator, and are therefore targets very rich in hydrogen atoms, whose protons allow (4.1).

Being a charged particle, e^+ quickly looses energy in the liquid scintillator and annihilates with an e^- of the medium. The annihilation γ-rays, along with the kinetic energy deposited by the e^+ as it ionises the scintillator, result in a trigger which is known as the "prompt event". The "delayed event" refers to the neutron capture. In the Double Chooz liquid scintillator, before annihilating, the e^- and e^+ bind together into a short-lived exotic atom,

called positronium. The lifetime of the positronium state observable in Double Chooz – a triplet state known as "ortho-positronium", decaying into three γ's[1] – is $\tau_{\text{o-Ps}} = (3.68 \pm 0.23)$ ns [67]. Moreover, the time necessary for a typical 2 MeV e^+ to deposit all of its kinetic energy is well-below the nanosecond. In light of the characteristic neutron-capture time, which is of the order of several tens of micro-seconds, it is evident that the prompt event is point-like both in space and time for most Double Chooz analyses, all the more so when the ortho-positronium state is not formed, which is the case for more than half of the events.

The abundance of hydrogen atoms in the liquid scintillator does not only offer a low-threshold interaction, as shall be developed in 4.1.2, but it also thermalises neutrons efficiently and captures them, exactly as light water would in nuclear reactors. Thereafter, the capturing nucleus may rearrange its internal structure and release one or several γ-rays. Capture on hydrogen leads to the emission of one 2.2 MeV photon, which unluckily competes with natural radioactivity, the highest natural peak being that of ^{208}Tl, at 2.6 MeV. As can be observed in Table 4.1, the gadolinium isotopes however emit γ-rays with a total energy well above natural radioactivity. In addition, the latter also have a capture cross-section several orders larger than that of hydrogen. In short, Gd-doping of the liquid scintillator provides the analysers with a fast high-energy delayed signal.

Isotope	E_{tot} (MeV)	σ (b)	Number of γ's
^{1}H	2.22	0.332	1
^{155}Gd	8.54	60.9×10^3	~ 5
^{157}Gd	7.94	25.4×10^3	~ 5

Table 4.1 – Total energy E_{tot} released after neutron-capture on the isotopes most relevant to the Double Chooz experiment. The neutron-capture cross-sections (symbol σ), as well as the number of γ's accounting for E_{tot}, are presented. Other gadolinium isotopes, with cross-sections similar to that of hydrogen, leading to ~ 6 MeV γ-rays, have been omitted due to their lower abundances[2]. More data can be found in [54].

The double coincidence of a prompt and a delayed event, within the right energy, time and distance windows is the signature to look for in the lavish amount of data the Double Chooz detectors are recording. A cartoon of this double-coincidence is shown in Figure 4.1.

[1] Although momentum conservation could be achieved by two photons, a triplet has spin $S = 1$, whose conservation entails coupling with a third photon.

[2] Gd-doping is minute, and hydrogen somehow makes up for its lower cross-section by its overwhelming numbers, but the other gadolinium isotopes with similarly low cross-section cannot compete with the abundance of the former.

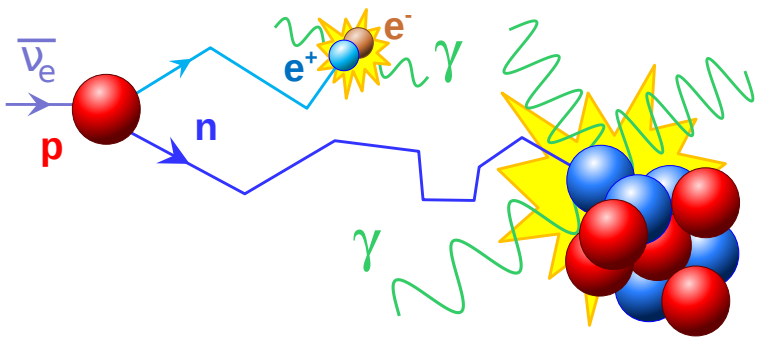

Figure 4.1 – After the inverse β-decay reaction, the e^+ quickly deposits its kinetic energy in the liquid scintillator (resulting γ's not displayed) and annihilates with an e^-, hence the prompt event. The neutron, after it has thermalised, scatters following a random walk until it is captured on a gadolinium nucleus (or possibly hydrogen). The latter then emits several γ-rays, which are detected as a delayed event.

4.1.2 Threshold

The maximum energy that a single $\overline{\nu_e}$ can carry away in the β-decays of the fission products is roughly 8 MeV. As we shall demonstrate, the threshold of (4.1) is thankfully below this endpoint and allows to retain a good proportion of the emitted $\overline{\nu_e}$'s. Nevertheless, the threshold $E_{\overline{\nu_e}}^{thres}$ is not negligible and if we are yet to observe proton decay, we are well acquainted with β-decay and the value of the mass difference $\Delta_{np} = m_n - m_p \simeq 1.293\,\text{MeV}$. In order to get a rough estimate of $E_{\overline{\nu_e}}^{thres}$, the mass of e^+ must be added to Δ_{np}, but let us derive it more rigorously.

We consider a proton p, at rest[3] in an hydrogen atom, and introduce the Mandelstam variable

$$s = (P_{\overline{\nu_e}} + P_p)^2 = (P_{e^+} + P_n)^2\,, \tag{4.2}$$

where P denotes the four-momentum of the particle indicated in subscript. From the evaluations of the inner product $P_{\overline{\nu_e}} \cdot P_p$ in the frame of the proton, which happens to be the lab frame, it follows that

$$s = {m_{\overline{\nu_e}}}^2 + 2E_{\overline{\nu_e}} m_p + {m_p}^2, \tag{4.3}$$

with $E_{\overline{\nu_e}}$ the energy of $\overline{\nu_e}$ in this frame. At threshold, the three-momenta of e^+ and n are $\overrightarrow{0}$, so that we also have

$$s = (m_e + m_n)^2\,. \tag{4.4}$$

Thus, the threshold energy $E_{\overline{\nu_e}}^{thres}$ for $\overline{\nu_e}$ in the lab frame is

$$E_{\overline{\nu_e}}^{thres} = \frac{(m_e + m_n)^2 - {m_{\overline{\nu_e}}}^2 - {m_p}^2}{2m_p} \simeq 1.806\,\text{MeV}, \tag{4.5}$$

[3]The thermal agitation is of course negligible compared to the other energies at play.

where the expression has been evaluated neglecting the mass of $\overline{\nu_e}$, conforming to 2.3.1.1. Writing $m_n = m_p + \Delta_{np}$ and keeping only leading order[4] terms in $\Delta_{np} > m_e > m_{\overline{\nu_e}}$, we find

$$E_{\overline{\nu_e}}^{thres} \simeq \Delta_{np} + m_e, \tag{4.6}$$

which echoes back to our naive picture. For comparison, the threshold of

$$\overline{\nu_e} + {}^{12}\mathrm{C} \rightarrow {}^{12}\mathrm{B} + e^+ \tag{4.7}$$

is 14.39 MeV [68][5] and reactor neutrinos somehow enforce an hydrogen-rich target.

4.1.3 Energy relations

4.1.3.1 $\overline{\nu_e}$ energy

The prompt signal is triggered by the total energy E_{e^+} the positron deposits in the detector, but the antineutrino deficit induced by θ_{13} depends on $E_{\overline{\nu_e}}$. Thus, it is most appropriate to try and relate both energies. However small, we would like to avoid dealing with the energy of the neutron[6], this is done by isolating its four-momentum P_n and squaring the conservation relation

$$(P_{\overline{\nu_e}} - P_{e^+} + P_p)^2 = {P_n}^2. \tag{4.8}$$

As in 4.1.2, we are to evaluate the inner products $(P_{\overline{\nu_e}} - P_{e^+}) \cdot P_p$ and $P_{\overline{\nu_e}} \cdot P_{e^+}$ in the lab frame, where the target proton is at rest. Expanding (4.8), we obtain

$${m_{\overline{\nu_e}}}^2 + {m_e}^2 + 2(\overrightarrow{p_{e^+}} \cdot \overrightarrow{p_{\overline{\nu_e}}} - E_{\overline{\nu_e}} E_{e^+}) + 2(E_{\overline{\nu_e}} - E_{e^+})m_p + {m_p}^2 = {m_n}^2. \tag{4.9}$$

Under the ultra-relativistic approximation $\|\overrightarrow{p_{\overline{\nu_e}}}\| \simeq E_{\overline{\nu_e}}$, we have

$${m_e}^2 + 2E_{\overline{\nu_e}}(\|\overrightarrow{p_{e^+}}\| \cos(\theta) - E_{e^+}) + 2(E_{\overline{\nu_e}} - E_{e^+})m_p + {m_p}^2 = {m_n}^2, \tag{4.10}$$

with θ the lab angle between $\overrightarrow{p_{e^+}}$ and $\overrightarrow{p_{\overline{\nu_e}}}$. We infer that the energy at play in the antineutrino deficit is related to the energy of the positron following

$$E_{\overline{\nu_e}} = \frac{{m_n}^2 - {m_p}^2 - {m_e}^2 + 2E_{e^+}m_p}{2\left(\sqrt{{E_{e^+}}^2 - {m_e}^2}\cos(\theta) - E_{e^+} + m_p\right)}. \tag{4.11}$$

[4] This ordering should only be understood when all terms are divided by m_p so that indeed $\Delta_{np}/m_p \ll 1$.

[5] Using the analogue of (4.5) yields a threshold of 13.89 MeV, which is not far from the [68] value including nuclear corrections.

[6] The kinetic energy distribution of neutrons is peaked below 5 keV and barely extends above 70 keV [69].

4.1.3.2 Visible energy

If we neglect the tiny energy deposited by the neutron, the energy visible in an IBD interaction corresponds to the sum of the kinetic energy of the positron T_{e^+} and that of the annihilation γ's

$$E^{vis} = T_{e^+} + 2m_e = E_{e^+} + m_e. \tag{4.12}$$

Rightfully neglecting some terms in (4.11), we can inverse the relationship between E_{e^+} and $E_{\overline{\nu_e}}$. Indeed, writing $\sqrt{{E_{e^+}}^2 - {m_e}^2} = \|\overrightarrow{p_{e^+}}\| = E_{e^+}\beta_{e^+}$, with β_{e^+} being the ratio of the speed of e^+ over c, and exploiting that for reactor experiments [70]

$$\langle \cos(\theta) \rangle \simeq 0.034\,\beta_{e^+} \ll 1, \tag{4.13}$$

we may reduce the denominator of (4.11) to $m_p - E_{e^+}$. Making the ordering explicit leads to

$$E_{\overline{\nu_e}} \simeq \frac{\Delta_{np} + \frac{{\Delta_{np}}^2}{2m_p} - \frac{{m_e}^2}{2m_p} + E_{e^+}}{1 - \frac{E_{e^+}}{m_p}}. \tag{4.14}$$

Rearranging and Taylor expanding gives

$$E_{\overline{\nu_e}} \simeq \left(E_{e^+} + \Delta_{np} + \frac{{\Delta_{np}}^2 - {m_e}^2}{2m_p} \right) \left(1 + \frac{E_{e^+}}{m_p} \right), \tag{4.15}$$

which is a second-order equation in E_{e^+}. Straightforward solving yields the only physical solution

$$E_{e^+} = -\frac{m_p + \epsilon}{2} + \frac{1}{2}\sqrt{\left(m_p + \epsilon\right)^2 + 4\left(E_{\overline{\nu_e}} - \epsilon\right) m_p}, \tag{4.16}$$

where ϵ is an energy defined as $\epsilon = \Delta_{np} + \frac{{\Delta_{np}}^2 - {m_e}^2}{2m_p}$, evidently satisfying $\epsilon < E_{\overline{\nu_e}}^{thres} \leqslant E_{\overline{\nu_e}}$. From there, substituting into (4.12) is effortless.

4.1.3.3 Visible energy at zeroth-order

Starting from either (4.16) or a zeroth-order form of (4.14), as is usually found in the literature, it is easy to deduce the leading-order expression

$$E_{e^+} \simeq E_{\overline{\nu_e}} - \Delta_{np}, \tag{4.17}$$

which is but the threshold equation (4.6) with m_e swapped for E_{e^+}. At zeroth-order, the visible energy plainly reads

$$E^{vis} \simeq E_{\overline{\nu_e}} - 0.78\,\text{MeV}, \tag{4.18}$$

and the visible energy threshold is indeed $2m_e = 1.022\,\text{MeV}$, in accordance with (4.5). Please note that the difference between (4.16) and the usual approximation (4.17) is larger than meets the eye (cf. Figure 4.2), and the bias may exceed the energy of the neutron.

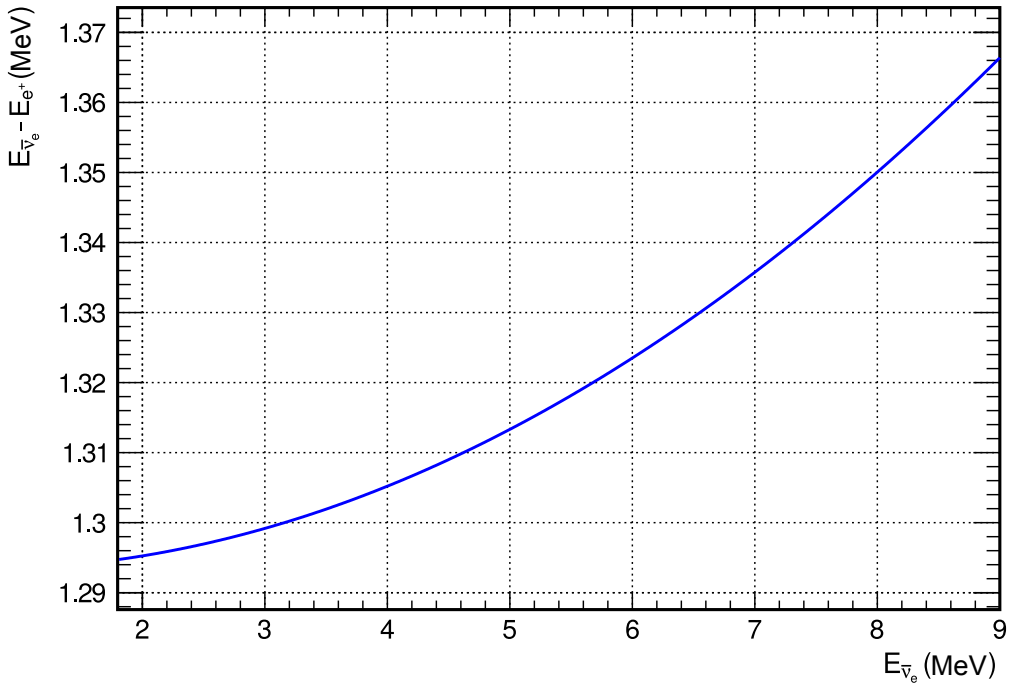

Figure 4.2 – Energy difference $E_{\overline{\nu_e}} - E_{e^+}$ as a function of the incoming antineutrino energy $E_{\overline{\nu_e}}$. If the zeroth order approximation predicts a $\Delta_{np} \simeq 1.293\,\text{MeV}$ constant shift across all reactor energies, as suggested by (4.17), the difference with Δ_{np} exceeds 50 keV at 8 MeV.

4.1.4 Cross-section

The cross-section for the IBD reaction may be rearranged in terms of the measured lifetime of the neutron $\tau_n = (880.3 \pm 1.1)\,\text{s}$ [51] and it increases with the energy of the incoming antineutrino. This time, the derivation of the final expression is beyond the scope of this document, and we will solely quote [71]. For the energy-range that reactor experiments study, the cross-section is well-approximated by

$$\sigma_p\left(E_{e^+}\right) \simeq \frac{2\pi^2}{m_e^5 f \tau_n}\, E_{e^+}\sqrt{{E_{e^+}}^2 - {m_e}^2}\,. \tag{4.19}$$

with $f = 1.7152$, a phase-space factor also accounting for Coulomb, weak magnetism, recoil, and outer radiative corrections. The cross-section may be expressed in terms of the energy of the antineutrino, substituting (4.16) into (4.20). For the reactor energies, which satisfy $E_{\overline{\nu_e}} \geqslant E_{\overline{\nu_e}}^{thres}$, the cross-section σ_p roughly exhibits a quadratic growth.

Natural units may be economical, but for experimentalists, it is always suitable to be able to plug-in numbers. Adding back the missing c's next to m_e, σ_p is currently homogeneous to $\text{MeV}^{-3}\text{s}^{-1}$, yet τ_n is only short of $\hbar^{-1}$ to effectively look like the inverse of an energy, so that we are left with MeV^{-2}. To land on our feet, we need but multiply the pre-factor by

$\hbar^2 c^2$ whose units are $\mathrm{MeV}^2 \cdot \mathrm{fm}^2$, thus

$$\sigma_p\left(E_{e^+}\right) \simeq \frac{2\pi^2\hbar^3 c^2}{m_e^5\, c^{10} f \tau_n}\; E_{e^+}\sqrt{{E_{e^+}}^2 - {m_e}^2 c^4}\,. \tag{4.20}$$

The constant pre-factor evaluates to

$$\frac{2\pi^2\hbar^3 c^2}{m_e^5\, c^{10} f \tau_n} \simeq (9.617 \pm 0.012) \times 10^{-44}\,\mathrm{cm}^2 \cdot \mathrm{MeV}^{-2}, \tag{4.21}$$

where the dominant error, i.e. that on τ_n, has been propagated. With the comprehensive formula (4.20) in mind, the cross-section values plotted in Figure 4.3 are easily understood. As an appetiser for 5.1, Figure 4.3 also depicts a typical $\overline{\nu_e}$ reactor flux – which roughly decreases as the exponential of a polynomial in $E_{\overline{\nu_e}}$ – and the expected detected spectrum. Aside from its dependency on the size of the target and its detection efficiency, the detected spectrum is essentially the product of the cross-section from (4.20) and the reactor flux, hence the presence of a maximum in Figure 4.3.

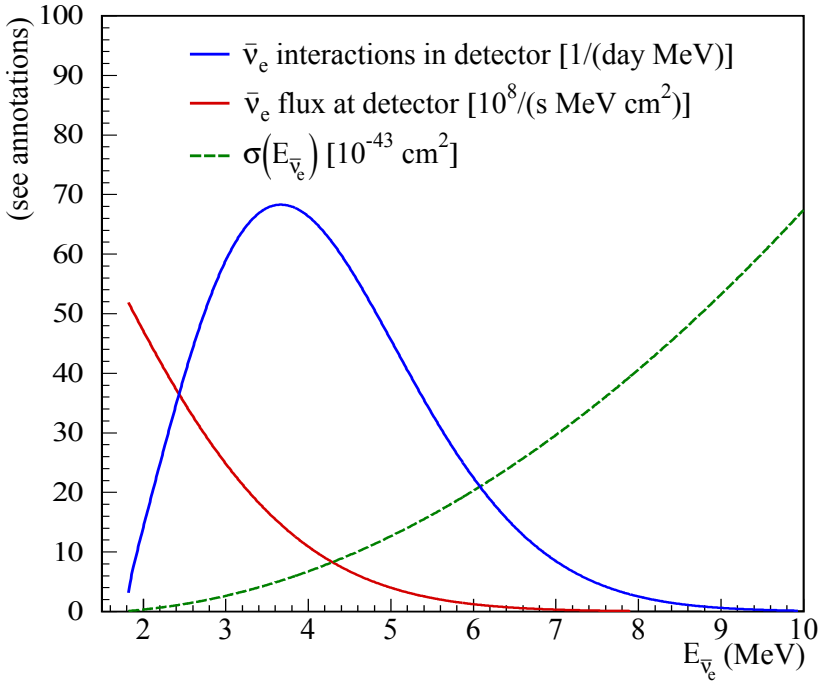

Figure 4.3 – IBD cross-section as a function of the energy $E_{\overline{\nu_e}}$ of the antineutrino (green dashed line). The energy-dependent antineutrino flux from a $12\,\mathrm{GW_{th}}$ nuclear reactor at 800 m (red solid line) is also plotted. With the help of the IBD cross-section, the antineutrino spectrum observed by a detector with a 12 t fiducial mass can be derived (blue solid line). Although nuclear reactors emit $\overline{\nu_e}$'s neutrinos below the IBD energy threshold, note that $E_{\overline{\nu_e}}$ starts at the threshold value $E_{\overline{\nu_e}}^{thres} \simeq 1.806\,\mathrm{MeV}$. Curves from [71].

4.2 Detector positioning

4.2.1 Striking a balance

Insofar as the probability to detect back an $\overline{\nu_e}$ depends on the distance from the $\overline{\nu_e}$ factory at which one tries to do so, it comes as no surprise that detector positioning is of paramount importance.

Assuredly, there are other factors at play, for as we she shall extensively stress in III, the overburden determines the amount of cosmogenic background – the worst of all, by far – the experiment has to suffer. The French Ardennes, which harbour the Double Chooz detectors, are quite hilly, so there comes another parameter to be picky about. Besides, since nuclear power plants have several cores – regardless of whether you have at your disposal a reactor flux monitor, i.e. a near detector – the flux that goes through each detector must be thoroughly understood. Indeed, since the $\overline{\nu_e}$ flux varies as the fuel burns up, as introduced in 3.3, and that no company would be foolish enough to refuel with the same fuel and start at the same time, all of its cores – which would deprive the electricity network of all its power during the vast outage – the instantaneous composition of the fuels, and the reactor powers, must be precisely monitored.

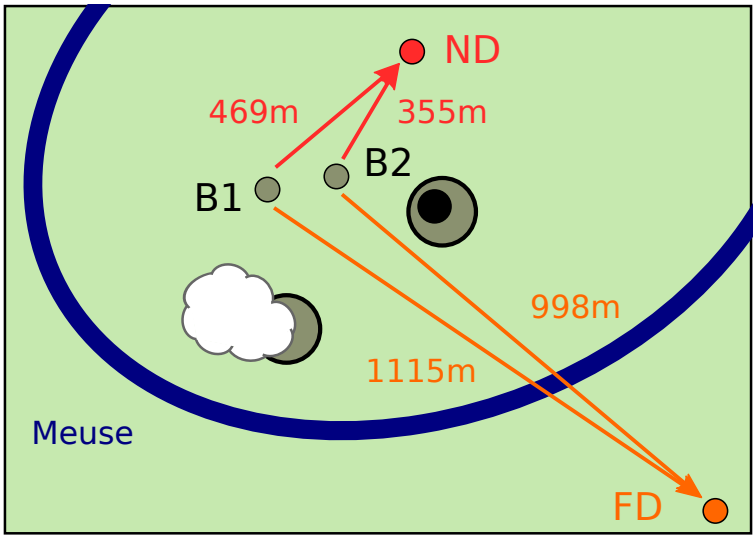

Figure 4.4 – Layout of the Double Chooz experiment. The reactor cores B1 and B2 are indicated in black. From the cooling towers, which make use of the water from the Meuse river, we can see that reactor B2 is off in this schematic, thereby temporarily turning the near detector (ND, red) into a perfect flux monitor of the far detector (FD, orange), at the cost of lower statistics. The distances between the reactor cores and the near and far detectors are next to the arrows representing the direction of the relevant part of the $\overline{\nu_e}$ flux, when both reactors are on.

There is a way around relying on overly-complicated flux predictions, though. Thereafter

focusing on the uncomplicated Chooz site, if the near detector were to be observe the same flux share (between reactor $B1$ and $B2$) as the far detector, the reactor fuel compositions and powers could be overlooked, this is referred to as "iso-fluxness".

To summarise, detector positioning is rooted in a compromise between the distance from the source, the overburden, iso-fluxness, the famed willingness of the French authorities and administrations, and of course, the essence of all undertakings: money. Weighing all these considerations lead the collaboration to the layout drawn in Figure 4.4. The far detector has been taking data since April 2011 whereas the near detector recorded its first antineutrinos in December 2014.

4.2.2 Averaged oscillation

As underscored in 4.2.1, unless one reactor is off[7], the near detector cannot be thought of as a perfect monitor of the far detector, and the fluxes must be weighted by means of the squared distances from Figure 4.4 and the state at which the reactors are operating. The layout of the Double Chooz experiment, although unrivalled in the θ_{13} field when it comes to geometrical suppression of the flux uncertainty, does not flawlessly fulfil iso-fluxness. Still, somewhat overlooking the time-dependence of the difference between both reactor fluxes, plotting the antineutrino survival probability for $\overline{\nu_e}$'s with a fixed energy $E^0_{\overline{\nu_e}}$, as in Figure 4.5, provides a compelling example of the effect of θ_{13}.

4.2.3 Overburden values

Although not indicated in Figure 4.4, the overburden values are of interest. The near detector stands under about 40 m of rocks whereas the far detector is not as shallow, lying under about 100 m of rocks. These overburdens are best used in their water equivalent ("mwe" for meter of water equivalent), which simplifies comparisons to other experiments, particularly so when it comes to cosmogenic backgrounds. The near detector has a 140 mwe overburden whereas the far detector has a 300 mwe one. For comparison, Daya Bay has 250 mwe and 860 mwe coverings for the near and far sites [56], respectively. Similarly, RENO has 120 mwe and 450 mwe overburdens [59]. In other words, the shallowness of the Double Chooz detectors brings about unique results concerning the cosmogenic studies, and they come with sizeable statistics.

4.3 Detector design

Apart from the overburdens and the external shieldings, the near and far Double Chooz detectors are virtually identical; a unique description shall cover both. With respect to

[7] It must be pointed out that even if the single reactor case provides a simple geometrical suppression of the reactor flux uncertainty, having several identical reactors also brings down the overall flux uncertainty.

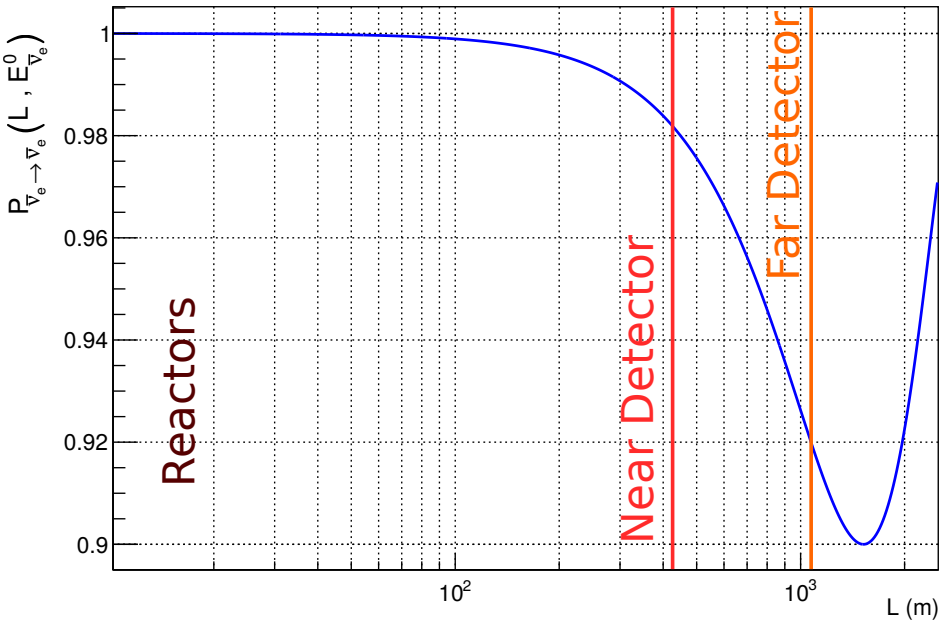

Figure 4.5 – Electron antineutrino survival probability $L \longrightarrow P_{\overline{\nu_e}\to\overline{\nu_e}}(L, E^0_{\overline{\nu_e}})$ for $\overline{\nu_e}$'s satisfying $E^0_{\overline{\nu_e}} = 3\,\text{MeV}$. The position of the near (red) and far (orange) detectors, via their average distances from the reactors, is stressed by the vertical lines crossing the oscillation probability. The textbook case $\sin^2(2\theta_{13}) = 0.1$ is assumed.

the CHOOZ experiment, the main improvements in detector design are the addition of a non-scintillating buffer, the tighter control over the natural radioactivity of all the detector parts, and the use of liquid scintillator demonstrating extreme stability over several-year time periods.

A cutaway view of the Double Chooz far detector, whose legend will be referred to in the following, can be found in Figure 4.6. The main components of the Inner Detector (ID) are indicated in blue (numbers 1 to 4), the additional passive and active shieldings in red (numbers 5 to 7), the target calibration systems in black (numbers 8 to 10).

4.3.1 Inner Detector

4.3.1.1 Neutrino Target (1)

The neutrino target (NT) is the part in which the neutrino interactions via inverse β-decay are expected. This innermost vessel of the ID is a 8 mm-thick cylindrical shape (number 1 in Figure 4.6), made of acrylic transparent to scintillation light above 300 nm, with a radius of 1.150 m and a height of 2.458 m, hence accounting for a 10.3 m^3 fiducial volume.

The Double Chooz detectors are filled with a newfangled liquid scintillator originating from a single batch to ensure identical compositions. In the aftermath of the degradation of the CHOOZ scintillator, the Double Chooz liquids were designed with emphasis on stability,

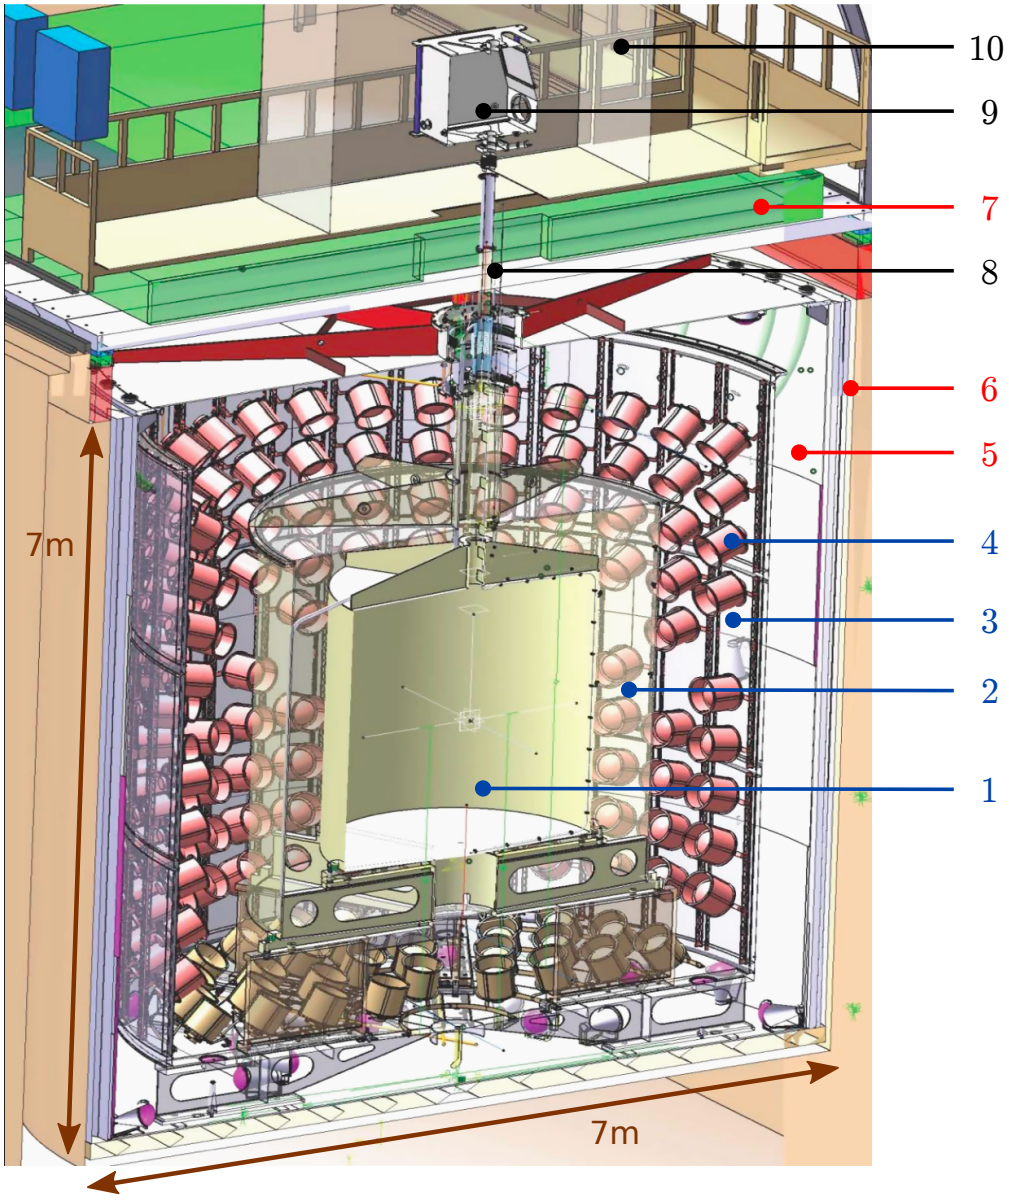

Figure 4.6 – CAD cutaway view of the Double Chooz far detector. Cutaway from [72].

transparency, and radiopurity [73]. As mentioned in 4.1.1, the target is full of protons, mostly provided by the $\sim 80\%_{\text{vol}}$ of dodecane ($C_{12}H_{26}$) and $\sim 20\%_{\text{vol}}$ of PXE ($C_{16}H_{18}$) that compose the liquid. PXE is a colourless aromatic liquid fluorescing in the UV range [74], it is meant to signal energy deposits in the detector. Two fluorescent materials – abbreviated "fluors" – are however added to bring the scintillation wavelengths from the UV range up to the visible range. The first one, short-named "PPO" ($C_{15}H_{11}NO$), allows to extract the energy of the PXE molecules in a non-radiative way, so that only the former emits photons. These photons are then wavelength-shifted by means of the added bis-MSB fluor ($C_{24}H_{22}$). Owing to these two fluors – taking into account absorption corrections – the main scintillation light lies between 430 nm and 500 nm, where the Photo-Multiplier Tubes (PMT) are most sensitive. To decrease capture times and increase the delayed energy (cf. 4.1.1), Gd is added to the target mixture. Since the rare earth Gd scarcely dissolves in organic solvents, it was first incorporated into a metal-organic complex, thereby allowing to reach a $0.123\%_{\text{w.}}$ concentration by weight in Gd, which corresponds to about 1 g/L.

4.3.1.2 Gamma Catcher (2)

The primary purpose of the gamma catcher (GC) is to ensure the conversion of γ-rays emitted near the edge of the NT into visible energy deposits. Indeed, recalling that Gd de-excites by emitting several γ-rays toting up to about 8 MeV (cf. Table 4.1), delayed-neutron captures near the edge of the target volume would otherwise have a lower detection efficiency[8]. In a nutshell, the GC is meant to allow the Gd-doped region, i.e. the NT, to fully define the fiducial volume of the detector. To this end, the 12 mm-thick acrylic of the GC is filled with 22.6 m^3 of liquid scintillator, hence forming a 55 cm-thick shell around the NT (see number 2 in Figure 4.6). The internal radius and height of the GC are 1.708 m and 3.572 m, respectively.

Inasmuch as the GC is only designed to catch escaping γ-rays, but not to capture neutrons – although it may also be used as such owing to its hydrogen fraction, at the cost of an increased accidental background – no Gd is present in the GC liquid. Gd-loading has non negligible consequences on the transparency of a liquid scintillator [73], which entails tweaking the composition of the GC liquid to match the properties of the target one. Indeed – aside from the stability and safety considerations – the design of the GC liquid is fundamentally aimed at matching the light yield in the NT, along with its volumetric mass density of $(804 \pm 1)\,\text{kg}\,\text{m}^{-3}$ at 14 °C. The latter requirement is rooted in the need to guarantee the mechanical stability of the acrylic vessels; accordingly, all the densities of the liquids must agree within less than 1%. To fit the bill, dodecane and PXE were mixed with a third solvent: medicinal white mineral oil. Last but not least, to shift the PXE light into the visible energy range, PPO and bis-MSB fluors were also added to the GC liquid.

[8] Unless one were to substantially lower the delayed energy threshold.

4.3.1.3 Buffer (3)

The buffer is the outermost volume of the Inner Detector, and the largest. Above all, it is the innermost passive shielding[9], aimed at protecting the scintillating regions it surrounds (see number 3 in Figure 4.6). However careful the production of materials may be, a small radioactive ^{40}K-contamination in the PMT glasses is indeed unavoidable. Furthermore, radioactive isotopes such as ^{235}U, ^{238}U, and ^{232}Th, are naturally present in all materials, including the surrounding rocks and the 3 mm-thick stainless steel vessel that encloses the 114.2 m^3 of non-scintillating buffer liquid. This 105 mm-thick layer, with a 2.760 m radius and a 5.680 m height, is therefore a noteworthy upgrade of the CHOOZ detector. On its wall, 392 Hamamatsu R7081 10-inch PMT's effectively cover about 13.5% of the steel' surface (number 4 in Figure 4.6) to collect the wavelength-shifted PXE scintillation light, effortlessly[10] travelling through the transparent liquids and acrylics.

Eventually, it is worth highlighting that all the materials that needed to be in contact with the liquids were screened for compatibility. A table summarising the composition of the liquids in the inner detector vessels, along with their respective volumes, can be found in Table 4.2. A visual illustration of the resemblance between the near and far inner detectors is given in Figure 4.7.

Vessel	**V** (m^3)	**Solvent** (%$_{\mathbf{vol}}$)			**Solute** (g/L)		
		dodecane	PXE	oil	PPO	bis-MSB	Gd
NT	10.3	80	20	0	7	0.02	1
GC	22.6	30	4	66	2	0.02	0
Buffer	114.2	43	0	57	0	0	0

Table 4.2 – Volume and chemical composition of the liquids in the inner detectors.

[9]The buffer is referred to as a passive shielding in the sense that there is no trigger to veto events penetrating it.

[10]The attenuation length is higher than 5 m in both detectors.

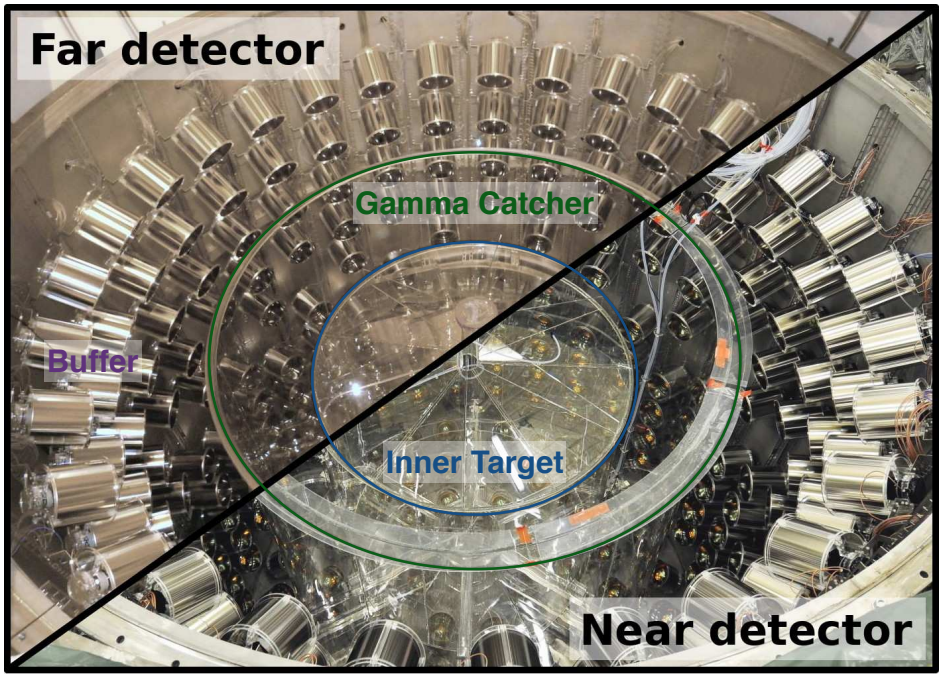

Figure 4.7 – Photos of the near and far Double Chooz inner detectors before the filling phases. The Target is indicated in blue, the Gamma Catcher in green, and the Buffer in purple. On the latter, 392 10-inch PMT's are attached [75].

4.3.2 Additional shieldings

4.3.2.1 Inner Veto (5)

Surrounding the Inner Detector are additional shieldings, the first of which being an active one, the Inner Veto (IV, number 5 in Figure 4.6). The latter contains 78 Hamamatsu R1408 8-inch PMT's recycled [76] from the Irvine-Michigan-Brookhaven (IMB) and Super-Kamiokande nucleon decay experiments. The layout of these PMT's was optimised via Monte-Carlo simulations [72] in regards to maximising the detection of muons and the correlated background they produce. To have a sufficient amount of light – in terms of photoelectrons per MeV – allowing the detection of clipping muons with short track lengths in the IV, the PMT's were divided into three parts: 24 PMT's on the top, 12 PMT's in a ring on the centre, and 42 PMT's at the bottom, thereby taking into consideration the support structures. Adjacent PMTs were set up to point in alternate directions, as Figure 4.8 shows. It is worth making note of the fact that adding more PMT's would have had an overshadowing effect, i.e. it would have reduced the field of view of the PMT's already present in the IV.

Despite impairing track reconstruction capabilities, the outwards facing buffer wall (grey vessel in 4.8) was covered with reflective foils, as was the IV wall of the near detector (the IV at the far detector was only painted white), thus increasing the amount of collected light,

and consequently the vetoing efficiency. Indeed, when the IV charge threshold is passed, the ID triggers following the IV trigger may be vetoed in subsequent (so-called "offline") analyses.

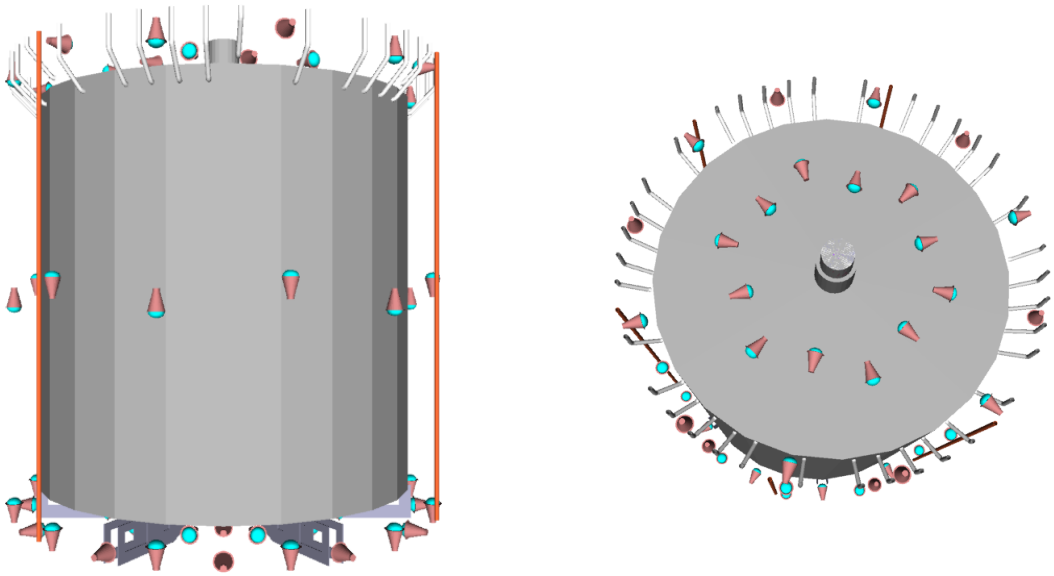

Figure 4.8 – Distribution of the 78 Hamamatsu 8-inch IV PMT's (easily identifiable by their light blue glasses) on the buffer exterior, which is represented by the grey cylindrical vessel. The inner veto vessel, on which the PMT's are installed, is not shown. The left cutaway is a side-view, the right one, a top-view [72].

While the inner detector liquids were developed by the Max-Planck-Institut für Kernphysik in Heidelberg, the 90 m^3 of IV liquid scintillator – enclosed in a 10 mm-thick 6.830 m-high cylindrical stainless steel vessel with a 3.250 m radius – were the responsibility of the Technische Universität München. The two solvents in this 50 cm-thick region are dodecane and the LAB aromatic fluor[11], in equal proportions. The solutes are the PPO and bis-MSB wavelength shifters, in concentrations identical to that of the GC liquid.

4.3.2.2 Passive shielding (6)

At the far detector laboratory, 250 t of demagnetised steel – in the shape a 15 cm-thick 7.150 m-high cylinder with a 3.300 m radius – surround the outside of the IV. This layer (number 6 in Figure 4.6) shields the scintillating volumes against naturally occurring γ-radiation in the rock and cavern surrounding the detector. The steel must be demagnetised to suppress the interactions with the electric circuits of the PMT's. And yet, low radioactivity sand or water could have served the shielding purpose. In fact, the use of such a steel

[11] If dodecane offers many target protons for $\overline{\nu_e}$ interactions, it does not scintillate; there is little sense in using dodecane where light collection matters most, i.e. in the IV.

shell originates from the space constraints that arouse when recycling the former CHOOZ laboratory, all the more so with a target $\sim 60\%$ larger. At the near detector, a 1 m-thick water shell shields the detector instead; additionally, water is a peerless neutron moderator, as advertised in 3.2.2.

4.3.2.3 Outer Veto (7)

The Outer Veto (OV, number 7 in Figure 4.6) is the outermost shielding of a Double Chooz detector. It consists of an array of solid scintillator modules topping the detector parts previously described. The OV is an improvement with respect to CHOOZ as it allows to tag muons passing near or slightly clipping the liquid scintillator volumes, hence allowing the analysers to correlate spallation neutron captures to OV triggers. When covering the chimney allowing the introduction of radioactive calibration sources (see 4.3.3), it also serves to identify muons decaying near the chimney, which are hardly visible to the IV.

However, the OV is unquestionably where both detector designs differ most. If the far detector was blessed with 44 veto-modules – arranged in perpendicular layers to provide precise coordinate information in both the x and y directions – covering an area of about $100\,\mathrm{m}^2$ above the main antineutrino detector, the near only received 10 modules [77], despite the considerably larger experimental hall. There is no denying that savings got the better of the near OV. Nonetheless, it must also be pointed out that the three-year gap, which separated the installation of the far and near detectors, saw the development of multiple techniques, aimed at cutting down on the stopping muon and fast neutron backgrounds without the OV. Still, when it is willing to work[12], the OV provides an interesting cross-check of the data analysed with these novel rejection techniques (reviewed in 6.2).

4.3.3 Calibration systems

In light of the energy-dependence of the antineutrino oscillation, if one is to precisely evaluate the induced-spectral distortion, one must needs understand flawlessly its energy scale, and thence have dedicated calibration systems.

4.3.3.1 Light injection system

The least invasive calibration device is the LED-fibre system. Since it does not make use of any radioactive sources, it harmlessly remains in both the ID and IV. Although not directly yielding the energies of the positrons, which relate to that of the $\overline{\nu_e}$'s through (4.16), LED flashes help to synchronise the time offsets of all the PMT's, as long as the right time of flight corrections are applied.

Moreover, for positron energies below 3 MeV – which is where the corresponding $\overline{\nu_e}$ spectrum from Figure 4.3 is maximum – for each ID PMT, one electron at best is ejected from

[12] Shifting duties saw a few OV-blaming followers, including the author.

the cathode. Consequently, quantifying the charge collected by one PMT for a single photo-electron (PE) signal is a valuable information, which can be obtained via low-intensity LED flashes. In turn, this charge[13] may be used to inverse the charge-PE relationship and eventually offer a common quantity, whose value can be compared between calibrated-PMT's. Higher light intensities, which are set up remotely, provide the so-called "multi-PE" calibration points. Altogether, the LED's allow to extract the number of PE's actually seen by each PMT during an inverse β-decay reaction, which paves the way for a thorough energy reconstruction, described further in 5.2.3.

With regard to hardware, the LED-based light injection system of the ID is channelled to 46 injection points interspersed between the PMT's. It offers three different wavelengths, which allow to study the properties of the scintillator. The 385 nm wavelength is mostly absorbed in the GC, the 425 nm one is partially absorbed in the GC and NT, whereas the 470 nm light blue is not appreciably attenuated, thus echoing back to the transparency of the liquids introduced in 4.3.1. The IV system is similar, although it only has two wavelengths: 365 nm and 475 nm.

Having the light injection systems at hand when taking shifts allows to monitor the time stability of the PMT gains, which are the physical quantities relating the integrated charge of each PMT to the corresponding number of PE's. This is particularly relevant after power outages[14].

4.3.3.2 Target calibration system (z-axis)

For each trigger, after having summed all the photo-electrons of all the PMT's – which are quantities we can compute thanks to the Light Injection system previously reviewed – one can contemplate matching the total number of photo-electrons to the actual energy of the source. This is where the so-called z-axis system comes into play.

Deployment of radioactive calibration sources can proceed by way of a glove box (number 9 in Figure 4.6), flushed with nitrogen so as to preserve the liquids, and accessed via a clean tent (number 10). From there, a motorised pulley-and-weight system can be operated to position sources along the chimney (number 8), the location of which corresponding to the symmetry axis of the cylindrical NT and going by the name of "z-axis". The positions of the sources are known to within 1 mm and range from 1 cm above the bottom of the NT to the chimney's end.

4.3.3.3 Gamma Catcher calibration system (Guide Tube)

Even when the GC is not used to extend the target volume via neutron-captures on hydrogen, comprehending the response in the GC is essential. Indeed, as neutrons easily travel from

[13] Or rather, the mean of the charge distribution for a given PMT, obtained after several light flashes.

[14] However close to the B1 and B2 reactors the Double Chooz detectors are, it sometimes feel as though they are located north of the Korean Demilitarised Zone.

one volume to another, using a neutron source near the edge of the target yields estimable information about border effects. Additionally, introducing sources in the GC allows to study the properties of the liquid scintillator in that region, and consequently the so-called light-non-linearity, which impacts the energy reconstruction (see 5.2.3.4). Last but not least, considering the position of sources in the Guide Tube (GT) is known to within 1 cm, the GT data allow to tweak the performance of the position reconstruction algorithm (explained in 5.2.2). The author being an official GT calibration expert, and having set up the system described thereafter for the first calibration of the near detector, this part will be expanded on.

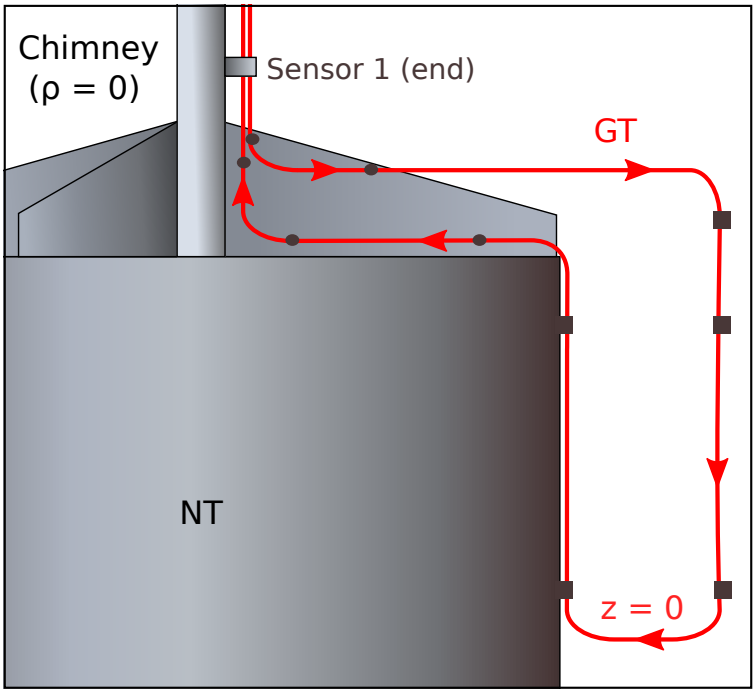

Figure 4.9 – The Guide Tube (GT, red) runs down along the chimney, the wall of the GC (not displayed) and eventually follows the NT before coming back next to the chimney. Radioactive sources are pushed into the GT in the direction indicated by the red arrows; they are stopped thanks to Sensor 1.

In short, the GT is a system of nested tubes guiding source capsules attached to a thin wire controlled by a pulley. The innermost tube is Teflon-made and it runs through a 5 mm stainless steel pipe hold at several points visible in Figure 4.9: NT lid, internal GC wall, exterior NT wall.

Capsules containing radioactive sources start their tour by passing through Sensor 2 (S2)[15], which must be set up in the clean tent for each deployment (see Figure 4.10). They

[15]The ordering is counter-intuitive and a little unfortunate, as is the American Gauge System for expressing the thickness of the tubes; we don't always get to choose.

reach the end of their lengthy $\sim 11.5\,\mathrm{m}$ journey at Sensor 1 (S1, only shown in Figure 4.9), below the clean tent. Both sensors are inductive and can be set to "low" or "high" sensibilities. The former setting allows sources to pass through whereas the latter blocks them.

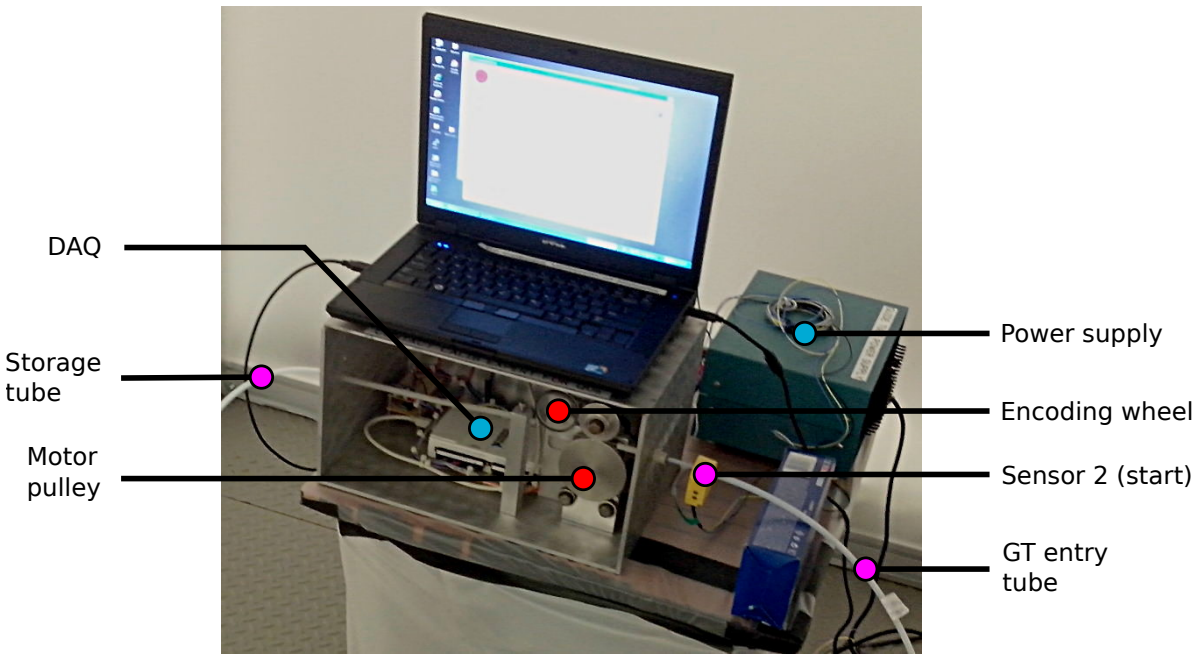

Figure 4.10 – Wire driver system set up in the clean tent of the near detector, above the chimney. The DAQ (light blue) is connected to the computer and allows to push sources past Sensor 2 (pink) – down into the GT itself – by way of a motor-driven pulley (red). The laptop accesses the distance travelled in the GT thanks to the encoding wheel (red). Sensor 2 is not fixed on any Teflon tube and it must be removed to attach sources to the wire from the storage tube (pink), which comes through the pulleys.

Each deployment implies zeroing the sensors. This is achieved by pushing a dummy capsule source from S2 to S1, measuring the travelled distance, lowering the sensibility of S1 to get $\sim 20\,\mathrm{cm}$ past its position so as to remove additional cable slack, slightly pulling back the capsule to zero S1, and eventually pulling the capsule all the way back to S2. The travelled distances S2S1 and S1S2 should agree within less than 1 cm. Regardless of the data acquisition point, radioactive sources are always sent to the end of the GT and pulled from there to minimise cable slack. Seeing that the capsule is still manually driven in the near detector, some skill is required to stop the ride on time, unless one wishes to start all over again.

While pushing the capsule, particularly past S1, the remaining wire cable that may still go through the driver system must be industriously monitored lest it fall off the pulley. Similarly, when installing the complete driver system, the storage tube must be unfolded with extreme diligence and taped as best as possible to the ground of the small clean tent, for if

the Teflon of the storage tube bends, the tenuous wire pushing the capsule may bend in turn and be wedged there. Attaching the capsule to the flange on the wire is not undemanding either; not to mention transporting the full wire driver system from one detector to the other, its non-removable fragile storage tube turning this task into a harrowing endeavour. Far be it from the author to denigrate a system that can hardly jeopardise the detector while providing helpful calibration data, but the wire driver system does require some habit and expertise.

4.3.3.4 Artificial sources

Up to four radioactive sources can be introduced in the ID. Of paramount importance is the ^{252}Cf source since with a 2.65 y half-life, it spontaneously fissions with a 3.09% branching ratio [78], the remaining 96.91% being α-decays, as expected from such heavy transuranium elements, in accordance with the chart from Figure 3.5. As exemplified by (3.1), on top of releasing fission products with asymetric masses, fission is accompanied by the emission of several neutrons, in the case of ^{252}Cf, about 3.7 prompt neutrons are released on average [79]. In view of the relatively short half-life, the current Double Chooz ^{252}Cf sources correspond to a neutron emission rate of $\sim 10\,\mathrm{s}^{-1}$, which implies long night runs to reliably study the neutron detection efficiency of the detectors, along with the energy scale corresponding to γ-rays emitted by H and Gd captures (cf. Table 4.1).

Lower energy sources are also put in service. With a burdensome[16] half-life of 30.08 y, ^{137}Cs offers the lowest energy out of all sources – natural and artificial included – available in Double Chooz. Although it is often peddled as a 662 keV γ-emitter, it is apposite to stress that ^{137}Cs mostly β-decays to the $11/2^+$ state of ^{137}Ba. In turn, the latter rapidly reaches its ground state by emitting a γ-ray with the aforementioned energy, as presented in Figure 4.11. Given that all sources are encapsulated, only the γ's make it to the detector, hence the hasty shortcut. This 662 keV energy is below the IBD threshold of $2m_e = 1.022$ MeV in terms of positron energy, but the analysis threshold of the experiment is advisedly below, as presented more in-depth by 6.1.1. Perhaps of higher interest is ^{60}Co. With a comfortable half-life of about 5.27 y, it offers γ-rays above the $\overline{\nu_e}$ detection threshold. As ^{137}Cs, ^{60}Co is by no means a γ-emitter but it mostly β-decays to the 4^+ state of ^{60}Ni, which rapidly de-excites by emitting two γ-rays of 1.17 MeV and 1.33 MeV (see Figure 4.11). These two γ's add up to an energy near the H-capture peak, a helpful input to better quantify the properties of the scintillators and the energy response near the oscillation maximum.

Eventually, a ^{68}Ge source may be introduced in the inner detector, though its 270.95 d half-life demands frequent preparations of new sources, a task all the more trying that both the z-axis and the GT require custom-shaped capsules. It has yet to be used in the near detector GT, but it provided a calibration point right at the detection threshold for the far detector. Indeed, compared to its stable ^{74}Ge counterpart, ^{68}Ge is willing to let go of a

[16] The ^{137}Cs isotope is produced by ^{235}U fissions and is particularly troublesome in the case of nuclear accidents.

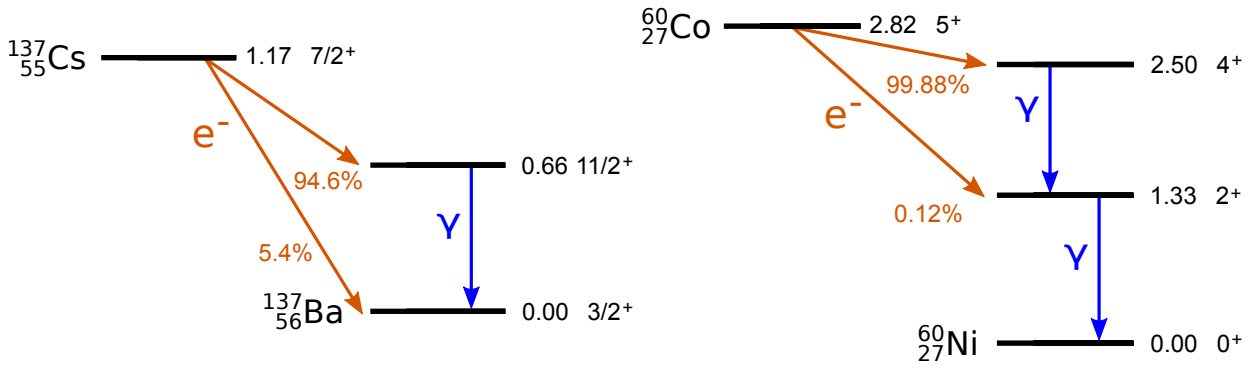

Figure 4.11 – Decay scheme of ^{137}Cs (left) and ^{60}Co (right) relevant for calibration purposes. The energies, spins, and parities of the levels are indicated next to their positions.

few protons, a result surely achieved via electron capture (or β^+ decay), leading to ^{68}Ga, which still falls short of one neutron and usually β^+ decay to the stable ^{68}Zn. As usual, charged particles remain in the capsule and two 551 keV annihilation γ's hit the scintillator, mimicking a positron signal at threshold.

4.4 Data acquisition

4.4.1 ν-DAQ

Detection of $\overline{\nu_e}$'s could not be complete without a readout and data acquisition system (referred to as ν-DAQ), whose schematic is given in Figure 4.12. Extensive descriptions of all the modules can be found in [80, 81, 82]; they are briefly summarised in the paragraphs that follow.

Each one of the 392 ID and 78 IV PMT's is connected via a single coaxial cable to a high voltage (HV) module. These modules consist of HV power supplies – each distributing a 1.3 kV voltage to its PMT – and custom HV splitters allowing to decouple the HV from the PMT outputs. The PMT signals, which have typical output voltages of 5 mV per PE after the splitters, are then passed to the front-end electronics (FEE) modules. The FEE's subsequently amplify and optimise these signals so as to feed the trigger boards and the flash analog-to-digital converters (FADC). The FADC's constantly write on their internal memory buffers 256 ns-long optimised PMT waveforms with 2 ns-time bins, thus providing accurate PMT hit times and helping to reject some backgrounds through their specific pulse shapes (see 6.2.2.1).

The trigger is composed of a master board, two ID boards, and one IV board. Each ID board receives signals from half of the ID PMT's, evenly distributed throughout the detector. The trigger boards do not receive single PMT outputs however, these are indeed gathered and summed so that one input to the trigger board corresponds to a variable-size PMT subgroup (3 to 6 PMT's for the IV and mostly 16 for the ID). For the ID and IV

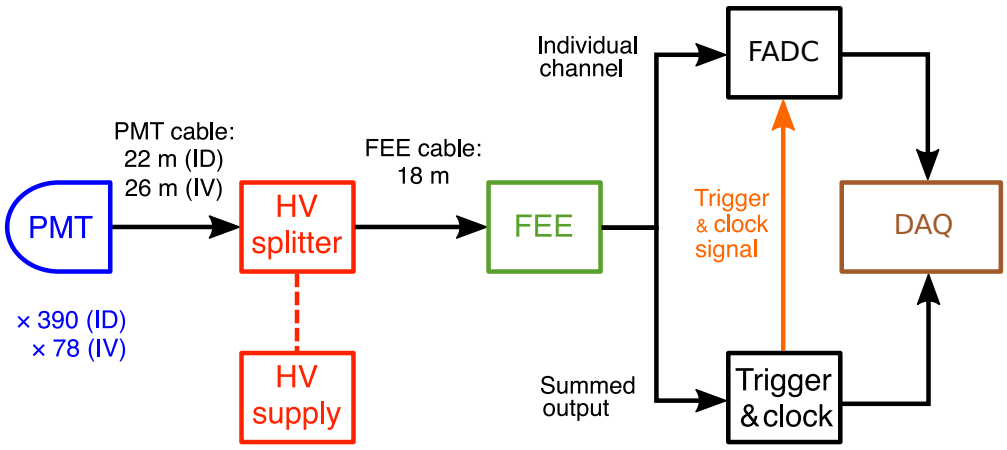

Figure 4.12 – Main readout and data acquisition system providing the waveforms of all the 392 ID and 78 IV PMT's. For each trigger, the waveforms are saved in 256 ns-long time windows with a 2 ns binning. Largely inspired by [77].

channels to be read out on the FADC buffers and written to disk (on the DAQ computers), any of the two ID trigger boards or the IV board may pass its custom threshold. The former is equivalent to about 350 keV while the IV threshold amounts to about 10 MeV, thereby corresponding to what an 8 cm minimum-ionizing muon track should deposit. On top of requiring one board to pass its threshold, the master board will trigger the recording of both the ID and IV waveforms only if the multiplicity of the PMT subgroups exceeds 2, i.e. if at least two PMT subgroups are hit.

4.4.2 OV-DAQ

With regard to the OV, the 64 wavelength-shifting fibres of each module are rooted to 8×8 multi-anode PMT's, each of which being attached to a board housing the Multi-Anode Read Out Chip (MAROC) and providing the PMT's HV. The read out chip seeks out hits above a certain threshold and sends this information to a FPGA for trigger decision.

When the trigger condition is achieved, signals are digitised by FADC's and stored in the FPGA memory. Groups of up to 10 OV modules are then read out together in a daisy-chain and conveyed via USB to the DAQ computers. The latter may also propagate commands to the PMT boards. To allow the OV-DAQ data stream to be later merged with that of the ν-DAQ, both are synchronised by means of periodic pulses. More details about the OV DAQ can be found in [83].

Chapter 5

Flux prediction and event reconstruction

Chapters 3 and 4 told us how to produce and detect antineutrinos, but they did not quite tell us how to actually extract the θ_{13} value from the recorded data. We do not even know how to reconstruct the energy needed in the oscillation formula (2.72). In addition, the rejection of backgrounds goes in hand in hand with a reliable knowledge of the positions of energy deposits in the scintillator. Likewise, the location of IBD candidates with respect to the tracks of the cosmic background producers, i.e. muons, helps to draw compelling correlations.

Before we even contemplate measuring θ_{13} with some set of selected interactions, based on the knowledge accumulated from Chapter 3, we must needs predict the $\overline{\nu_e}$ spectrum seen by each detector. Not only must the fuel composition of the reactor be related to the production of antineutrinos, but the modelling of interactions in the liquids ought to be simulated as well.

5.1 Prediction

5.1.1 $\overline{\nu_e}$ flux

Thereafter, we focus on the case in which only one reactor is monitored, as if the other reactor were off; fluxes and rates being additive quantities, they will be summed at the end of this section with no effort.

5.1.1.1 Fission rate

Back in (3.7), we approximated the rate of fissions happening in a nuclear reactor by the ratio of its thermal power P_{th} over an underestimated energy per fission of 200 MeV; although boasting about the magnitude of the reactor $\overline{\nu_e}$ flux, it will not serve to accurately quantify

the small $\sin^2(2\theta_{13})$ deficit. Overlooking isotopes other than the ones in

$$\mathcal{F} = \left\{^{235}\mathrm{U}, ^{238}\mathrm{U}, ^{239}\mathrm{Pu}, ^{241}\mathrm{Pu}\right\}, \tag{5.1}$$

the thermal power of a nuclear reactor is actually the sum of the fission rates r_f^k of these elements k in $\mathcal{F}$, weighted by the energies E_f^k released in the fission of each one of these isotopes (see 3.1 for values). Thus, the thermal power reads

$$P_{th} = \sum_{k\in\mathcal{F}} r_f^k E_f^k, \tag{5.2}$$

which is no different from a sum over the powers generated by each isotope. Introducing the total fission rate

$$r_f = \sum_{k\in\mathcal{F}} r_f^k, \tag{5.3}$$

and the fractional fission rates α_k for $k \in \mathcal{F}$,

$$\alpha_k = \frac{r_f^k}{r_f}, \tag{5.4}$$

we may rewrite the total rate of fissions happening in a nuclear reactor as

$$r_f = \frac{P_{th}}{\sum\limits_{k\in\mathcal{F}} \alpha_k E_f^k}. \tag{5.5}$$

5.1.1.2 Fractional fission rates

The fractional fission rates from (5.4) are actually corrected so that $\sum_{k\in\mathcal{F}} \alpha_k = 1$, which somehow makes up for the non-simulated isotopes accounting for about $\sim 0.3\%$ of the fissions happening in a reactor, and whose energies per fission and spectra should not be fundamentally different from the main four contributors [63]. Following the processes described in 3.2.3, the fission fractions evolve over time; they were simulated during the single-detector phase by way of the MCNP Utility for Reactor Evolution (MURE)[84, 85], using detailed fuel inventories and thermal power data. The results were successfully benchmarked against the DRAGON code [86]. For the multi-detector phase, less computing-intensive models able to reproduce the former MURE outputs have been developed [87].

An example of the evolution of the fission fractions for a typical reactor cycle, that of B2, is shown in Figure 5.1. In accordance with 3.2.3, while ^{235}U is burnt, ^{239}Pu builds up via neutron captures on ^{238}U, whose baffling amount provides a virtually constant contribution to the $\overline{\nu_e}$ rate. The presence of ^{239}Pu and ^{241}Pu at the beginning of the cycle is rooted in the use of fuel mixing new and spent assemblies, as mentioned in 3.2.3.

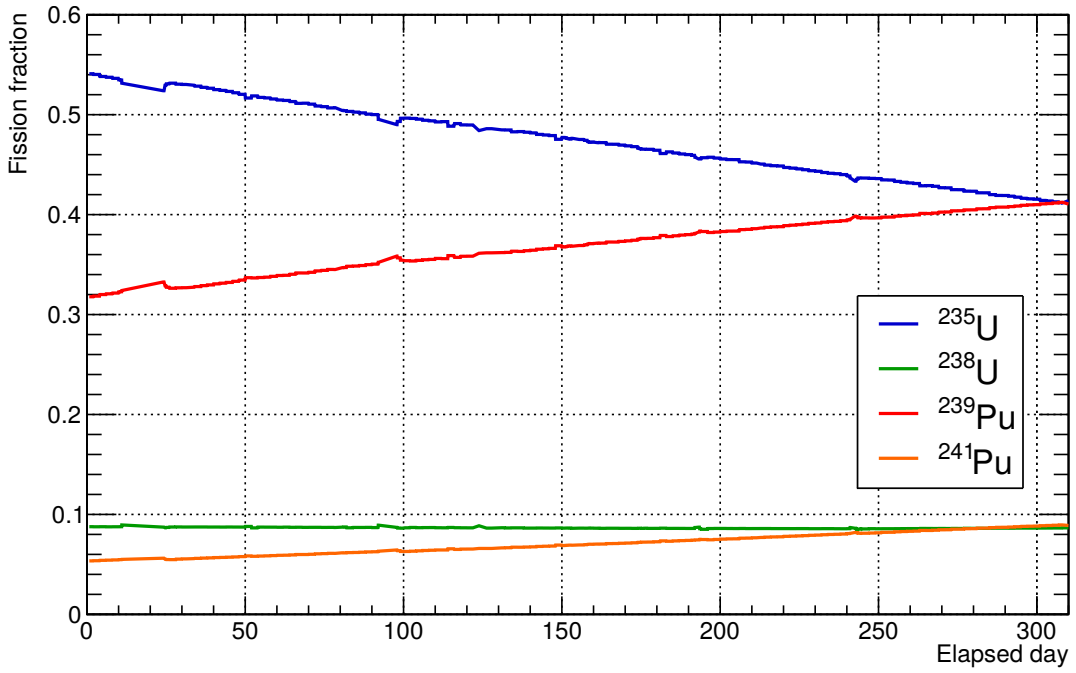

Figure 5.1 – Simulated evolution of the fractional fission rates for reactor B2 over a typical 300 d cycle starting with mixed fuel. The small glitches are due to corrupted or unavailable data in the power plant's database, from where the inputs of the simulations are retrieved.

5.1.1.3 Infinitesimal flux

At a given time t, the flux $\mathrm{d}\phi$ of antineutrinos with an energy between $E_{\overline{\nu_e}}$ and $E_{\overline{\nu_e}} + \mathrm{d}E_{\overline{\nu_e}}$, seen by a detector monitoring a reactor from a distance L, is the product of the total fission rate $r_f\,(t)$ in this reactor and the number of $\overline{\nu_e}$'s emitted for each fission[1] with the right energy $S_{\overline{\nu_e}}\,(E_{\overline{\nu_e}}, t)\,\mathrm{d}E_{\overline{\nu_e}}$, corrected for the isotropic spread. Thus, we have

$$\mathrm{d}\phi\,(E_{\overline{\nu_e}}, t) = \frac{r_f\,(t)}{4\pi L^2}\; S_{\overline{\nu_e}}\,(E_{\overline{\nu_e}}, t)\,\mathrm{d}E_{\overline{\nu_e}}\;, \tag{5.6}$$

and $\mathrm{d}\phi$ is expressed in $\mathrm{s}^{-1}\mathrm{cm}^{-2}$ or any equivalent combination of units.

5.1.1.4 Reference spectra

In usual terms, $S_{\overline{\nu_e}}$ is plainly the antineutrino spectrum per fission. The number of emitted $\overline{\nu_e}$'s not only depends on the energy that you consider, but it also depends on the time at which the data are taken. Indeed, as presented in 3.2.2, ^{235}U and ^{239}Pu have different fission product distributions, and thence, different decays chains and corresponding β-spectra. The same goes for ^{238}U and ^{241}Pu, although their contribution to the total fission rate is meagre,

[1] In the 3.3 appetiser, this quantity was approximated by 6 $\overline{\nu_e}$'s per fission, though without requirements on the energy window, thus completely overlooking the IBD threshold.

in accordance with Figure 5.1. More precisely, $S_{\overline{\nu_e}}$ is defined at each instant in time t by

$$S_{\overline{\nu_e}}\left(E_{\overline{\nu_e}}, t\right) = \sum_{k \in \mathcal{F}} \alpha_k\left(t\right) S_k\left(E_{\overline{\nu_e}}\right), \tag{5.7}$$

where S_k is the antineutrino spectrum corresponding to one fission of $k \in \mathcal{F}$.

The β-spectra that account for the $\overline{\nu_e}$ spectrum we observe when fissioning one nucleus $k \in \mathcal{F}$ are numerous, several thousands, if truth be told. Stating that calculations from first principles – the so-called "ab initio" method – are challenging is putting it mildly. Consequently, the Double Chooz experiment relies on the conversion of the cumulated β-spectra, measured in the eighties at the Institut Laue-Langevin (ILL), in Grenoble. At the ILL research reactor, thin foils containing each isotope were exposed to the thermal neutron flux 80 cm from the core, thereby inducing fission, as explained in 3.2.1. Thereafter, the electrons from the β-decays of the fission products were analysed with a magnetic spectrometer. The cumulated β-spectrum of ^{235}U was published in [88], while that of ^{239}Pu and ^{241}Pu were published later in [89]. Hardly undergoing fission under thermal neutron fluxes, as discussed in 3.2.3, the cumulated β-spectrum of ^{238}U remained unmeasured until late 2013, and ab initio predictions had to be relied upon for this isotope. The Forschungs-Neutronenquelle Heinz Maier-Leibnitz reactor (FRM), near Munich, covers a broad range of applications and is equipped, amongst other apparatuses, with a fast neutron source. The latter provided the ^{238}U $\overline{\nu_e}$ spectrum in 250 keV bins between 2.875 MeV and 7.625 MeV, with an energy-dependent error of 3.5% at 3 MeV [90]. As for the ILL spectra, the lower energy region is interpolated down to the IBD threshold defined in (4.5).

The conversion procedure is based on [91, 92]. In short, it exploits energy conservation during β-decays, rightfully neglecting the recoil energies of the daughter nuclei. However, the cumulated spectrum of an isotope includes many branches; accordingly, several virtual branches – which somehow average the nuclear data – are fitted to each cumulated electron spectrum as measured at ILL and FRM. Once a linear combination of virtual electron spectra convincingly matches the cumulated electron spectrum, the conversion into a cumulated $\overline{\nu_e}$ spectrum – via each of one the virtual branches – is straightforward by way of energy conservation.

For each fission of one isotope $k \in \mathcal{F}$, the spectrum S_k, according to which the energies of the few $\overline{\nu_e}$'s emitted above the IBD threshold are distributed, is shown in Figure 5.2. As can be noticed, the integral of the spectrum corresponding to the fission of ^{239}Pu is marginally lower than that of ^{235}U, i.e. as the reactor fuel burns up (see 5.1.1.2), less $\overline{\nu_e}$'s are expected, although this is somewhat mitigated by the resemblance the ^{241}Pu-induced spectrum bears to that of ^{235}U.

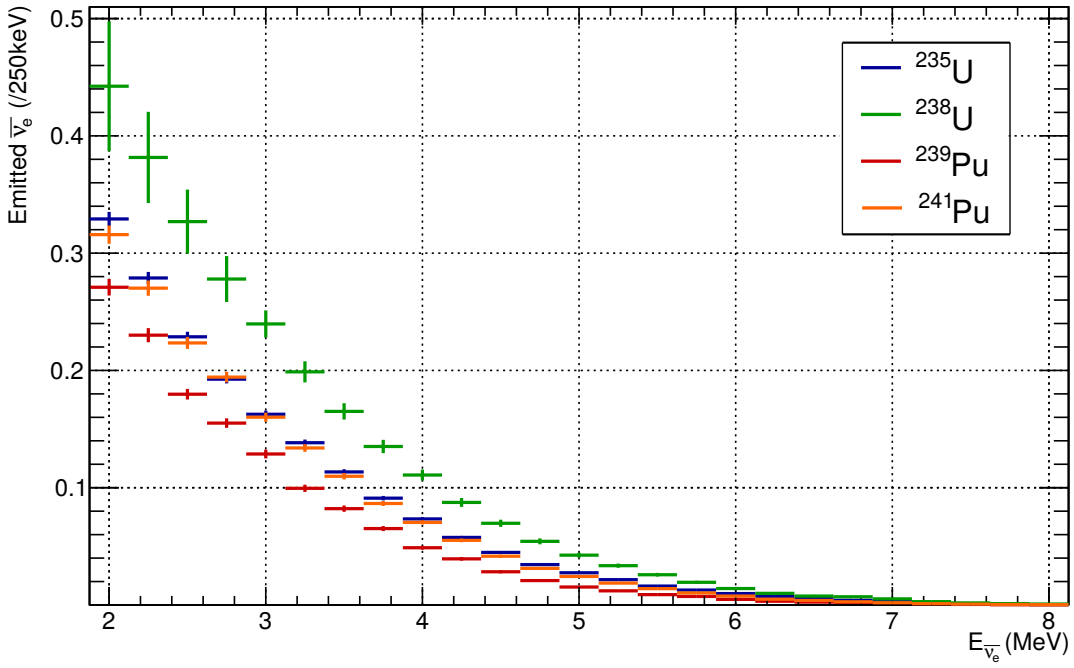

Figure 5.2 – Electron antineutrino spectra induced by the fission of the labelled isotopes. The spectra are normalised to one fission and the bin contents integrate over a 250 keV range.

5.1.2 Expected signal

5.1.2.1 Differential form

By and large, the electron antineutrino rate $r^{exp}_{\overline{\nu_e}}$ expected at a detector is the product of the effective flux ϕ_{eff} it sees – directly related to the total fission rate r_f in the reactor it monitors – and the cross-section σ_{NT} its target offers. However, since the cross-section on a single proton σ_p (cf. 4.1.4) depends on the energy of the incoming $\overline{\nu_e}$[2] and that

$$\sigma_{NT} = n_p \, \sigma_p \, , \tag{5.8}$$

with n_p the number of target atoms, i.e. the so-called "proton number", we must supply the differential form $\mathrm{d}\phi_{eff}$. Proceeding on this track, the rate of $\overline{\nu_e}$'s expected to have an energy between $E_{\overline{\nu_e}}$ and $E_{\overline{\nu_e}} + \mathrm{d}E_{\overline{\nu_e}}$ is

$$\mathrm{d}r^{exp}_{\overline{\nu_e}} \left(E_{\overline{\nu_e}}, t \right) = n_p \, \sigma_p \left(E_{\overline{\nu_e}} \right) \mathrm{d}\phi_{eff} \left(E_{\overline{\nu_e}}, t \right) . \tag{5.9}$$

We define the effective flux ϕ_{eff} as the flux seen by the detector, yet corrected for its detection efficiency ϵ. In Double Chooz, as presented in 4.4.1, the prompt energy threshold is low enough to be viewed as 100% efficient[3], and the detection inefficiency is essentially

[2] In (4.20), σ_p is expressed in terms of E_{e^+} but we painstakingly related E_{e^+} to $E_{\overline{\nu_e}}$ in (4.16), we hereby gracelessly retain the symbol σ_p although its parametrisation has changed.

[3] The analysis threshold is a little higher than the trigger threshold, but that does not change the argument.

driven by the ability to detect neutrons on Gd[4], plus minor corrections to be discussed in 6.3.2. In other words, ϵ is virtually independent from $E_{\overline{\nu_e}}$. Moreover, considering that the data must be compared to the non-oscillated expectation to assess how many electron antineutrinos θ_{13} discards from the flux, ϕ_{eff} need not include assumptions on the non-zero value of $\sin^2(2\theta_{13})$. It follows that

$$\mathrm{d}\phi_{eff}\left(E_{\overline{\nu_e}}, t\right) = \epsilon \; \mathrm{d}\phi\left(E_{\overline{\nu_e}}, t\right). \tag{5.10}$$

By substituting (5.6) into (5.10), we can obtain the short form of the rate

$$\mathrm{d}r^{exp}_{\overline{\nu_e}}\left(E_{\overline{\nu_e}}, t\right) = \frac{n_p\,\epsilon}{4\pi L^2}\, r_f\left(t\right) \sigma_p\left(E_{\overline{\nu_e}}\right) S_{\overline{\nu_e}}\left(E_{\overline{\nu_e}}, t\right) \mathrm{d}E_{\overline{\nu_e}}\ . \tag{5.11}$$

5.1.2.2 Binned expected antineutrino rate

Making explicit the time dependency in (5.11) – which arises from the evolution of the fractional fission rates and the thermal power – by applying (5.5) and (5.7), we find

$$\mathrm{d}r^{exp}_{\overline{\nu_e}}\left(E_{\overline{\nu_e}}, t\right) = \frac{n_p\,\epsilon}{4\pi L^2} \frac{P_{th}\left(t\right)}{\sum\limits_{m\in\mathcal{F}} \alpha_m\left(t\right) E_f^m} \sum_{k\in\mathcal{F}} \alpha_k\left(t\right) \sigma_p\left(E_{\overline{\nu_e}}\right) S_k\left(E_{\overline{\nu_e}}\right) \mathrm{d}E_{\overline{\nu_e}}\ . \tag{5.12}$$

However many events they may record each day, experiments are bound to bin their data. Unlike $\mathrm{d}E_{\overline{\nu_e}}$, the variable-length bin-width ΔE_i of the i-th energy bin (values defined in visible energy in 6.3.1.2), used in the analysis, is not an infinitely small. In a nutshell, (5.12) must be integrated over between the bin edges E_i and $E_i + \Delta E_i$ so that

$$R^{exp}_{\overline{\nu_e}\ i}\left(t\right) = \frac{n_p\,\epsilon}{4\pi L^2} \frac{P_{th}\left(t\right)}{\sum\limits_{m\in\mathcal{F}} \alpha_m\left(t\right) E_f^m} \sum_{k\in\mathcal{F}} \alpha_k\left(t\right) \int_{E_i}^{E_i+\Delta E_i} \sigma_p\left(E_{\overline{\nu_e}}\right) S_k\left(E_{\overline{\nu_e}}\right) \mathrm{d}E_{\overline{\nu_e}}\ , \tag{5.13}$$

where we introduced the binned rate $R^{exp}_{\overline{\nu_e}\ i}$, expressed in s^{-1}, and formally defined as the integral of the spectral rate $\mathrm{d}r^{exp}_{\overline{\nu_e}} \big/ \mathrm{d}E_{\overline{\nu_e}}$ over the i-th energy bin

$$R^{exp}_{\overline{\nu_e}\ i}\left(t\right) = \int_{E_i}^{E_i+\Delta E_i} \frac{\mathrm{d}r^{exp}_{\overline{\nu_e}}}{\mathrm{d}E_{\overline{\nu_e}}}\left(E_{\overline{\nu_e}}, t\right) \mathrm{d}E_{\overline{\nu_e}}\,. \tag{5.14}$$

5.1.2.3 Total binned antineutrino count

In the same way that the energy range is binned, the time resolution is not infinitely small; the number of $\overline{\nu_e}$'s detected between the acquisition times t and $t + \Delta T$, with an energy

[4] Analysing only delayed events around 8 MeV, i.e. Gd-captures, misses out about 14% of the $\overline{\nu_e}$'s.

between E_i and $E_i + \Delta E_i$, is computed as

$$N^{exp}_{\overline{\nu_e}\ i}(t, \Delta T) = \int_t^{t+\Delta T} R^{exp}_{\overline{\nu_e}\ i}(t')\,\mathrm{d}t' = \frac{n_p\,\epsilon}{4\pi L^2} \sum_{k\in\mathcal{F}} \left(\int_t^{t+\Delta T} \alpha_k(t') \frac{P_{th}(t')}{\sum_{m\in\mathcal{F}} \alpha_m(t')\,E_f^m}\,\mathrm{d}t' \right) \left(\int_{E_i}^{E_i+\Delta E_i} \sigma_p(E_{\overline{\nu_e}})\,S_k(E_{\overline{\nu_e}})\,\mathrm{d}E_{\overline{\nu_e}} \right). \tag{5.15}$$

Notwithstanding the diligence of (5.15), not only are the time steps at which P_{th} is recorded 6 s at best [63] but the simulation of the fission fractions are made such that there is only one value per run with a typical run length of $\Delta T = 3600\,\mathrm{s}$. In other words, the counts expected during a run starting at t and lasting ΔT artlessly read

$$N^{exp}_{\overline{\nu_e}\ i}(t, \Delta T) = \frac{n_p\,\epsilon}{4\pi L^2} \frac{P_{th}(t)\,\Delta T}{\sum_{m\in\mathcal{F}} \alpha_m(t)\,E_f^m} \sum_{k\in\mathcal{F}} \alpha_k(t) \int_{E_i}^{E_i+\Delta E_i} \sigma_p(E_{\overline{\nu_e}})\,S_k(E_{\overline{\nu_e}})\,\mathrm{d}E_{\overline{\nu_e}}\,. \tag{5.16}$$

Eventually summing over both reactors $R \in \mathcal{R}$ with $\mathcal{R} = \{B1, B2\}$, and writing explicitly the dependencies on R, we get

$$N^{tot}_{\overline{\nu_e}\ i}(t, \Delta T) = \sum_{R\in\mathcal{R}} \frac{n_p\,\epsilon}{4\pi {L_R}^2} \frac{P^R_{th}(t)\,\Delta T}{\sum_{m\in\mathcal{F}} \alpha^R_m(t)\,E_f^m} \sum_{k\in\mathcal{F}} \alpha^R_k(t) \int_{E_i}^{E_i+\Delta E_i} \sigma_p(E_{\overline{\nu_e}})\,S_k(E_{\overline{\nu_e}})\,\mathrm{d}E_{\overline{\nu_e}}\,, \tag{5.17}$$

i.e. the count to which the data reconstructed in the right energy bin must compared.

In fact, Double Chooz is somehow a three-detector set up and constraints correlated across the near and far detectors, retrieved and adapted from the results of the Bugey4 experiment [93], are applied to the energy integral in (5.17) [82]. To make a long story short, the rate measurement from Bugey4 is corrected for the different fuel composition between its reactor and the Double Chooz cores, which serves to constrain the spectra-weighted cross-section, thereby yielding a 1.7% uncertainty on the $\overline{\nu_e}$ rate prediction.

5.1.3 Monte-Carlo Signal

5.1.3.1 Random generation of events

Although in 4.1.3.2 we derived a relationship between the visible energy the positron should deposit in a $\overline{\nu_e}$ interaction E^{vis} and the incoming antineutrino energy $E_{\overline{\nu_e}}$, converting E^{vis} back into $E_{\overline{\nu_e}}$ through (4.14) and (4.12), would amount to overlooking the modelling of the scintillator and electronics, which are such that the energy actually detected is not the idealised E^{vis}. As a matter of fact, quite the opposite is done: the expected $\overline{\nu_e}$ spectrum for each run, as defined in (5.17), is not directly compared to the recorded data, rather, it paves the way for randomly generating $\overline{\nu_e}$ interactions – which mirror the state of the reactors during the acquisition period – in the detector.

Practically, for each simulated $\overline{\nu_e}$, a random interaction vertex is drawn in the target. Once a direction for the outgoing positron has been pulled by means of the relevant angular distribution (see [70] for instance), which depends on $E_{\overline{\nu_e}}$, the kinematic quantities of the positron and the neutron are fully determined. The latter can then proceed through the modelling of the detector.

5.1.3.2 Main simulation

The input momenta of the positron and neutron – generated for each simulation of an inverse β-decay in the detector – are processed through a custom Geant4-based [94, 95] detector simulation package, called DCGLG4sim [96]. This package reproduces the detector geometry and scintillator characteristics, with refinements – built on top of the 2009 9.2.p02 Geant4 version – particularly aimed at improving the models for thermal neutron diffusion and radiative capture. Furthermore, these additions have been benchmarked against a software developed at CEA, under the name "Tripoli-4", which is dedicated to modelling the behaviour of thermal neutrons. This comparison helped to assess the systematic uncertainties associated to neutron detection in the Double Chooz detectors [69]. The geometry – particularly when it comes to the position of the PMT's – is modelled in intricate detail, with sub-millimetre accuracy on account of photographic surveys [82].

5.1.3.3 Quenching model

Overview

With regard to the liquids, their parameters, such as scintillation light yields, Cherenkov light production amplitudes, attenuation lengths, refraction indices, and reflectivities, were thoroughly measured in the laboratory and subsequently tuned in the Monte-Carlo simulation [73, 97, 98].

Of particular importance to understand the results of the simulations in III is the so-called "quenching effect", which the loss of scintillation light (measured for instance in photoelectrons collected by a PMT), for particles with a large energy loss per distance travelled, embodies. The higher the energy loss $\mathrm{d}E/\mathrm{d}x$, the smaller the scintillation light L, and therefore, the detected energy. Birk's law reads

$$\frac{\mathrm{d}L}{\mathrm{d}x} = S\frac{\frac{\mathrm{d}E}{\mathrm{d}x}}{1 + b\frac{\mathrm{d}E}{\mathrm{d}x}}, \tag{5.18}$$

with S the scintillation efficiency (in any equivalent of MeV^{-1} when L is measured in photoelectrons), and b Birk's constant (in $\mathrm{cm\,MeV}^{-1}$).

Echoing to the Bragg peak, the energy loss $\mathrm{d}E/\mathrm{d}x$ is itself a function of the actual energy of the particle, and the former increases whilst the latter decreases. As is manifest from (5.18), the light output per distance travelled $\mathrm{d}L/\mathrm{d}x$ is consequently proportional to

the energy loss $\mathrm{d}E/\mathrm{d}x$ at relatively high energies, satisfying

$$b\frac{\mathrm{d}E}{\mathrm{d}x} \ll 1. \tag{5.19}$$

Electron quenching

This effect does not only impact on charged particles, it also indirectly affects γ-rays, via Compton electron quenching. In particular, regarding neutron-captures on Gd, the several γ-rays emitted deposit their energies in multiple Compton scatter events, thereby casting out numerous low-energy electrons, hence the need to fully model quenching for the charged leptons. Electrons, themselves, mostly suffer from a loss of scintillation light below 100 keV, as was measured in the laboratory utilising 662 keV γ's from ^{137}Cs decays (cf. 4.3.3.4) scattering off electrons. In a Compton scattering event

$$e + \gamma \to e' + \gamma', \tag{5.20}$$

the outgoing electron e' carries a kinetic energy $T_{e'} = E_{e'} - E_e = E_\gamma - E_{\gamma'}$, subject to quenching. In our case, for each decay of ^{137}Cs, $E_\gamma = 662\,\mathrm{keV}$ and $E_{\gamma'}$ can be evaluated by a germanium detector, which enables us to match a particular $T_{e'}$ value to the number of photo-electrons collected in coincidence by a PMT.

The light yield per distance travelled is difficult to measure and its integral along the path, or over the deposited energy, is actually provided by the PMT pulses. We may indeed change the integration variable; using the chain rule to compute the derivative of the composition $L \circ E$, we get

$$\frac{\mathrm{d}L}{\mathrm{d}x} = \frac{\mathrm{d}L}{\mathrm{d}E}\frac{\mathrm{d}E}{\mathrm{d}x}. \tag{5.21}$$

As a consequence, the light output L generated by a particle of kinetic energy T can be computed as

$$L(T) = L(T) - L(0) = \int_0^T \frac{\mathrm{d}L}{\mathrm{d}T'}\,\mathrm{d}T' = S\int_0^T \frac{1}{1 + b\frac{\mathrm{d}T'}{\mathrm{d}x}}\,\mathrm{d}T'\,, \tag{5.22}$$

where we exploited the fact that taking the derivative with respect to T or E makes no difference.

Plugging in the Berger-Seltzer equation [99] describing $\mathrm{d}E/\mathrm{d}x$ $(= \mathrm{d}T/\mathrm{d}x)$, (5.22) can be evaluated numerically for various sets of the parameters (S, b). Doing so, for each $T_{e'}$ measured in the lab, allows to find the set of parameters (S, b) which best reproduce the data, and thus can be saved[5] in DCGLG4sim, whose quenching model is Birk's. In the ideal case in which the loss $\mathrm{d}T/\mathrm{d}x$ is small compared to b^{-1}, across the energies of interest, (5.22) simplifies to

$$L_{lin}(T) = ST, \tag{5.23}$$

[5]This is a starting point for b; the simulation takes into account the production of secondary particles whereas numerical integration of Birk's law only tackles the energy loss of the primary particle.

which satisfies $L_{lin}(T) \geq L(T)$ for all T. Plotting the ratio L/L_{lin}, with S gauged using a relatively high energy reference, demonstrates the effect of quenching. Laboratory measurements epitomising electron quenching, along with the model built from the best-matching (S, b) set, are displayed in Figure 5.3. In the NT, Birk's constant for electrons is evaluated to $b_e^{NT} \simeq 0.016\,\mathrm{cm\,MeV^{-1}}$ [98].

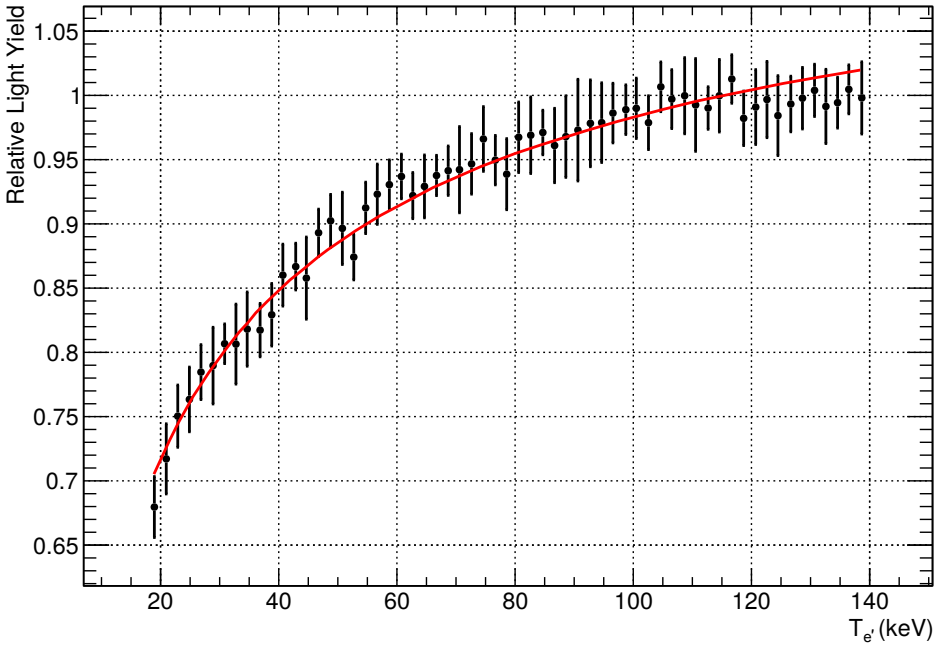

Figure 5.3 – Ratio of the light yield L of electrons in the GC over the linear expectation L_{lin} as a function of the kinetic energy of electrons $T_{e'}$. The linear behaviour takes 140 keV as a reference, i.e. $L_{lin}(T) = L(140\,\mathrm{keV})\,T/140\,\mathrm{keV}$. Values below 1 indicate a quenched light output to be reproduced by the detector simulation. Birk's constant is evaluated to $b_e^{GC} \simeq 0.028\,\mathrm{cm\,MeV^{-1}}$ in the GC liquid, and serves to produce the model (red). Data points from [98].

Alpha quenching

Naturally occurring radioactive isotopes include α-emitters, which – despite the radiopurity requirements and shieldings – cannot be completely left out. The decay of muon-induced cosmogenic isotopes (see Chapter 7) also release α particles. As a result, a proper modelling of the interaction of these heavy particles – which, as a side note, can also be used for calibration purposes – is essential for the success of the experiment. Their large charge and weight go hand in hand with a substantial energy loss along the track $\mathrm{d}E/\mathrm{d}x$, leading to light yield quenching at the MeV scale, also modelled by Birk's law.

Contrary to ^{137}Cs, whose decays release γ-rays which can easily penetrate a liquid cell and subsequently knock off electrons, α-emitters must be incorporated in the liquid itself. This was achieved by flushing onto a powder of PPO fluor (see 4.3.1) an illustrious product

of the ^{238}U chain: ^{222}Rn. The latter decays within days to ^{218}Po, which in turn produces ^{214}Po and ^{210}Po. All of them decay by emitting α's with fixed kinetic energies ranging from $\sim 5.3\,\mathrm{MeV}$ to $\sim 7.7\,\mathrm{MeV}$. In this case, the light output can be directly related to the energy of the source as the stopping power of a $T = 7.7\,\mathrm{MeV}$ α-particle is about $80\,\mu\mathrm{m}$ [97]. A similar procedure as that described for electrons can be carried out to deduce a value for the Birk's constant of α's: $b_\alpha \simeq 0.010\,\mathrm{cm\,MeV^{-1}}$ in the NT.

Even at 7.7 MeV, the light output is strongly reduced, thus, there is no α source exhibiting a linear behaviour. Instead, it is enlightening indeed to seek the electron kinetic energies T_e, which are necessary to produce the same amount of light as the aforementioned α sources L_α. Looking at different Compton scattering angles, and using other γ sources, enables to study the light yield of electrons L_e in the $500 - 700\,\mathrm{keV}$ region, where $L_e \sim L_{lin}$. In that energy range, electrons produce as much light as α's with more than ten times their kinetic energies. Therefore, this study provides a one-to-one correspondence $T_e \leftrightarrow L_e$, i.e. a function $T_e \rightarrow f(T_e) = L_e$ whose inverse can be used to associate a L_α measurement to a T_e value: $T_e = f^{-1}(L_\alpha)$. From there, the quenching factor Q of an α particle of kinetic energy T_α which generates a L_α strong light is defined as

$$Q(T_\alpha) = \frac{T_\alpha}{f^{-1}(L_\alpha)}, \tag{5.24}$$

where $f^{-1}(L_\alpha)$ is often shortened into $T_e(L_\alpha)$. The ratios of the kinetic energies of α's to the electron ones producing the same light outputs are summarised in Table 5.1 for the four ^{222}Rn-descending sources.

T_α	Q_{NT}	Q_{GC}
5.30 MeV	13.1 ± 0.2	17.4 ± 0.5
5.49 MeV	12.7 ± 0.3	16.2 ± 0.4
6.00 MeV	11.9 ± 0.3	15.3 ± 0.2
7.69 MeV	9.8 ± 0.4	12.6 ± 0.6

Table 5.1 – Quenching factors for α sources of energy T_α in the NT (Q_{NT} column) or GC (Q_{GC} column) liquids. Values from [97].

5.1.3.4 Readout and common algorithms

Once the interactions of particles in the liquid scintillator and surrounding objects – including acrylics and PMT's – have been suitably modelled, there remains to deal with the photo-electrons generated at the photo-cathodes of the PMT's. This is precisely at this point that the Double Chooz Read-out System Simulation (DCRoSS) package takes over DCGLG4sim. Broadly speaking, DCRoSS simulates all that is related to the electronics of the experiment: PMT's, FEE's, triggers, FADC's, DAQ's. Accordingly, for each event, the end result is an array of digitised waveforms, which bears a striking resemblance to actual data.

To this array of waveforms, the Double Chooz Common Trunk (DCCT) of reconstruction algorithms – such as energy and position, to be detailed in 5.2 – is applied nearly exactly as it would be for its data counterpart. DCGLG4Sim, together with DCRoSS and DCCT, form what is commonly referred to as the Double Chooz Offline Group Software: DOGS.

5.2 Event reconstruction

So as to choose the right energy bin for the interaction of the positron in the liquid scintillators, a first-rate energy resolution is needed. In the same manner, to avoid counting unrelated accidental coincidences and ensure that the right number of events is set in the considered energy bin, a spatial coincidence between the prompt energy deposit and the neutron capture, is required. Moreover, to reject correlated backgrounds induced by cosmic muons, reconstructing the tracks of through-going muons is a must. Thus, this section focuses on all the reconstruction algorithms essential for a top-notch analysis of the data recorded by the Double Chooz detectors.

5.2.1 Charge and time

If the light injection systems – described in 4.3.3.1 – allow to match the charge recorded by a PMT to a number of PE's produced at its cathode, there remains to define how the charge is built from the bare FADC counts that make up a waveform.

Out of the 128 2 ns-time bins forming a 256 ns-long waveform, a 112 ns sub-window of $n_w = 56$ bins is selected such that it maximises the sum of the baseline-subtracted FADC counts. The reconstructed charge (in units of the FADC) for channel i consequently reads

$$q_i = \sum_{k=1}^{n_w} {c_i}^{(k)} \,, \tag{5.25}$$

with ${c_i}^{(k)}$ the baseline subtracted FADC counts in the k-th sample.

The baseline is indeed monitored by a fixed rate trigger (1 Hz) for each channel i, thereby providing its mean value and root-mean-square deviation. The length of the sliding integration window was chosen to match the width of single-PE signals, which are the most common for $\overline{\nu_e}$ interactions. Nevertheless, retaining only the largest sum q_i would sometimes amount to reconstructing the largest noise fluctuation; to reject such blunders, at least one of the time bins must have 2 FADC counts, i.e. about a third of a PE [60]. Along these lines, the sum q_i must outweigh the integrated spread of the baseline of the considered channel, i.e.

$$q_i > \sqrt{\sum_{k=1}^{n_w} {\sigma_i^{(k)}}^2} \,, \tag{5.26}$$

where $\sigma_i^{(k)}$ is the standard deviation of the baseline in the k-th time bin of the i-th channel, which is taken independent from all the other time bins of this channel. Assuming that all

the FADC time bins of a given channel i are modelled by an identical baseline spread σ_i^b value – measured via the fixed trigger – the summed deviation due to baseline fluctuations simplifies and (5.26) becomes

$$q_i > \sigma_i^b \sqrt{n_w}\,, \tag{5.27}$$

which is the condition applied in the DCRecoPulse algorithm, from DCCT. If the (5.27) requirement cannot be met, the waveform is disposed of.

As regards the start time of a pulse reconstructed in channel i, it is defined as the time t_i at which the pulse reaches 20% of its maximum value [82]. The start times are subsequently shifted according to the time offset measured by means of the light injection systems.

5.2.2 Position

The charges and start times of the waveforms recorded for an event are the inputs of the position reconstruction algorithm in the ID: DCRecoBAMA, named for Alabama, the birth place of the code.

5.2.2.1 Principle

Assuming a point-like energy deposition of light intensity I (in photons sr^{-1}) located in space-time at x, and implementing the optical model discussed in details in [100], it is possible to uniquely predict the charge q_i^{pre} and start time t_i^{pre} of each channel $i \in [\![1, n]\!]$, with n the number of active PMT's. Trying values for (x, I) until the charges $(q_i)_{i\in[\![1,n]\!]}$ and times $(t_i)_{i\in[\![1,n]\!]}$ recorded for the event form a highly probable set, i.e. a likely fluctuation around the prediction, allows to determine a plausible position and light intensity for this interaction in the liquid scintillator. This does not necessarily mean that we must go in pursuit of the (x, I) which yields $q_i^{pre} = q_i$ and $t_i^{pre} = t_i$ in each channel i, this search may be fruitless, for indeed, the recorded charges $(q_i)_{i\in[\![1,n]\!]}$ and times $(t_i)_{i\in[\![1,n]\!]}$ are subject to fluctuations, and the odds are small that we observed the mere prediction of the purely optical model.

5.2.2.2 Maximum likelihood estimation

In a systematic way, a maximum likelihood estimation (MLE) is performed. Since a likelihood is also used to identify backgrounds in Chapter 9, it deserves a distinctive emphasis. The likelihood of the parameters (x, I), given the observed $(q_i)_{i\in[\![1,n]\!]}$ and $(t_i)_{i\in[\![1,n]\!]}$, is defined as

$$\mathcal{L}\left(x, I \mid q_1, \ldots, q_n, t_1, \ldots, t_n\right) = f_{x,I}\left(q_1, \ldots, q_n, t_1, \ldots, t_n\right), \tag{5.28}$$

where $f_{x,I}$ is the global model[6] describing the charge and time distributions corresponding to the PMT waveforms for a source characterised by (x, I). In statistical terms, $f_{x,I}$ is the

[6] By virtue of the optical model, one (x, I) point predicts a unique set $(q_i^{pre}, t_i^{pre})_{i\in[\![1,n]\!]}$. Hence, f could also be regarded as a purely statistical model parametrised by the q_i^{pre}, t_i^{pre}'s, not only is this parametrisation both cumbersome and to be confused with the fixed q_i, t_i's, but it is usually written in permissive notations.

parametrised joint probability density of the charge and start time variables. Hypothesising that these variables suffer from independent fluctuations, (5.28) simplifies into

$$\mathcal{L}\left(x, I \mid q_1, \ldots, q_n, t_1, \ldots, t_n\right) = \prod_{i=1}^{n} f_{x,I}^{i}\left(q_i\right) h_{x,I}^{i}\left(t_i\right), \tag{5.29}$$

where $f_{x,I}^{i}$ and $h_{x,I}^{i}$ are respectively the charge and start time probability densities for the i-th channel. Elaborating a little, $f_{x,I}^{i}$ allows to access the probability to record, in the i-th channel, a charge infinitely close to q_i , while one should expect from a source characterised by (x, I) a q_i^{pre} value, not necessarily equal to the recorded charge. The time distributions share the same meaning as the charge ones. Software is fundamentally designed to minimise functions, besides, CPU's are shaped to add numbers, in a nutshell, (5.29) is best turned into

$$\mathcal{F}\left(x, I \mid q_1, \ldots, q_n, t_1, \ldots, t_n\right) = -\sum_{i=1}^{n} \ln f_{x,I}^{i}\left(q_i\right) - \sum_{i=1}^{n} \ln h_{x,I}^{i}\left(t_i\right), \tag{5.30}$$

with $\mathcal{F} = -\ln \mathcal{L}$. For each event, $\mathcal{F}$ is but a function of (x, I), which the minimiser within RecoBAMA varies to attribute the most likely position.

Varying both the position x and the intensity I, via the light injection systems (cf. 4.3.3.1), helps to characterise the $f_{x,I}^{i}$ and $h_{x,I}^{i}$ densities. In addition, the effective light attenuation and the PMT angular responses – which allow to predict q_i^{pre} and t_i^{pre} from (x, I), utilising the optical model – are tuned with calibration sources from both the z-axis and the Guide Tube [60].

5.2.2.3 Resolution

Using both the time and charge information yields a reconstruction resolution of about 12 cm with slightly worse results at energies below 1 MeV or in the GC. Figure 5.4 exemplifies the performance of RecoBAMA in the Guide Tube (GT) of the near detector, for a hastily analysed 60 min-run taken during the last calibration campaign, with a ^{60}Co source in the lowest part of the tube (refer to Figure 4.9 if need be), which is defined by $z = 0$. The GT itself is in the yz-plane; a top-performing algorithm should consequently return a mean $x = 0$[7]. Bearing in mind that the theodolite survey yet remains to be fully analysed, and hence, the actual position of the GT frame to be refined, the performance is already gratifying.

5.2.3 Energy

Now that the charge and position reconstruction algorithms are at hand, i.e. the two pre-requisites for energy reconstruction, we can calibrate the charges with familiar sources at established positions.

[7] We hereby relieve from its duties the former x, i.e. the four-dimensional variable introduced in 5.2.2.1, and return to down-to-earth notations.

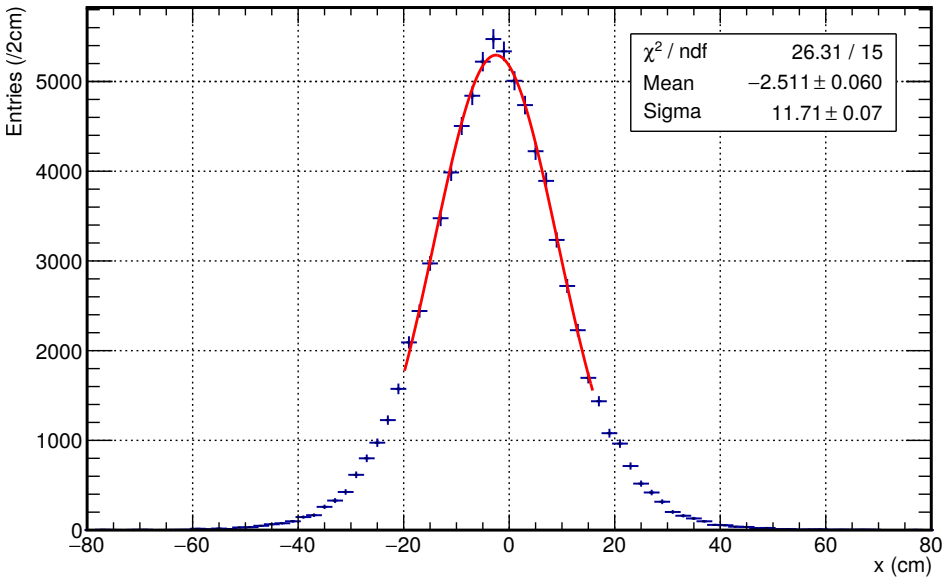

Figure 5.4 – Distribution of the reconstructed x coordinate of a ^{60}Co source positioned along the y-axis. The source was inserted in the GT, in its equatorial part ($x = 0$, $z = 0$). The GT frame itself defines the $x = 0$ plane. A Gaussian fit around the mean reconstructed position, which exhibits a meagre $-2.5\,\text{cm}$ offset, is also presented.

5.2.3.1 Finding the number of photo-electrons

Individual gain

As introduced in 4.3.3.1, the PE's produced on each PMT cathode are the building blocks of the energy reconstruction. For a given light intensity emitted in the detector, the charges reconstructed in 5.2.1 are, in effect, realisations of the random variables $(Q_i)_{i\in[\![1,n]\!]}$, related to the Poisson character of the photo-electron counting $(N_i)_{i\in[\![1,n]\!]}$. Ideally, the charges would be proportional to the number of PE's

$$Q_i = g_i N_i, \tag{5.31}$$

and g_i would be nothing but the PMT gain of the i-th channel, which – once characterised – could be used to inverse all the charges and obtain the number of PE's each waveform represents.

Calibration

To this intent, the sample mean μ_{Q_i} of the charge distribution of data taken with the light injection system (see Figure 5.5) – at a given intensity – can be computed, with i running on all the illuminated channels. The sample mean indeed offers an unbiased estimator of the expectation of Q_i. Similarly, the sample variance, and thence, the standard deviation σ_{Q_i} can be obtained from the LED calibration data. By virtue of the expression of the Poisson

distribution, the expected value μ_{N_i} of N_i is equal to its variance ${\sigma_{N_i}}^2$. In other words,

$$\mu_{Q_i} = g_i {\sigma_{N_i}}^2 \tag{5.32}$$
$$\sigma_{Q_i} = g_i \sigma_{N_i} \,, \tag{5.33}$$

were the estimation subtleties[8] have been looked away. Squaring (5.33) and dividing by (5.32) casts σ_{N_i} into the background to offer a graceful expression – based on the first moments of the charge distributions of the LED data – for the gain of the i-th channel

$$g_i = \frac{{\sigma_{Q_i}}^2}{\mu_{Q_i}} \,. \tag{5.34}$$

Allegedly, then, the gain g_i of each channel i is independent of μ_{N_i}, i.e. of the mean number of expected PE's. Insofar as each light intensity – whether be it from the LED's or scintillation light – fixes the value of μ_{N_i}, we can infer that g_i is independent of the light intensity emitted in the detector and that probing g_i with one light output should suffice.

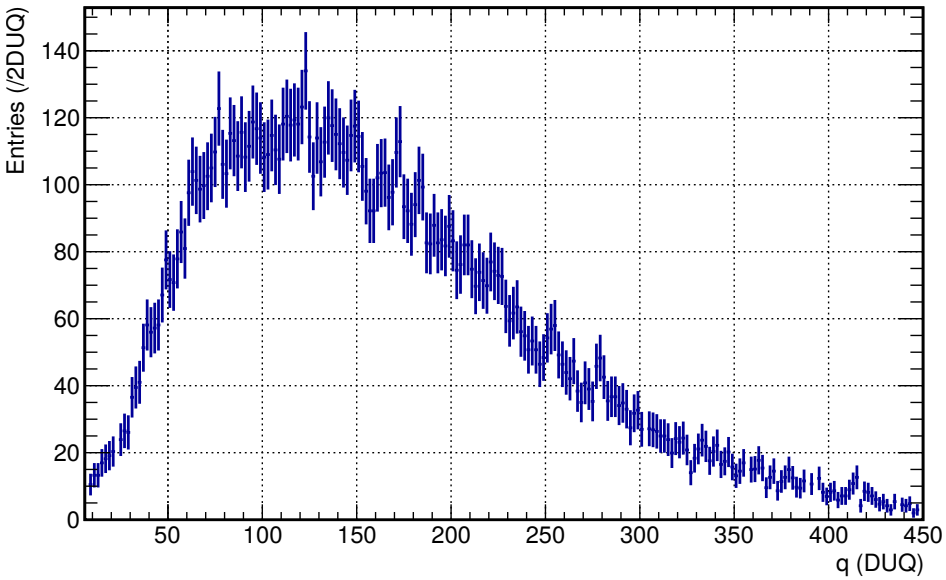

Figure 5.5 – Distribution of the reconstructed charge q, for a typical channel, when taking data with the light injection calibration system. The mean of this distribution, along with its spread, provide a first order estimation of the gain g of this channel. The charge units, coming from the FADC, read "DUQ" for "digital unit of charge". Assuming a single PE contribution, the mean itself already favours a gain $g \sim 120 DUQ/PE$.

[8] Theoretically, there should be some hats on top of the sample means and standard deviations, which are but estimators of the true expectations.

Gain non-linearity

Obviously, there comes a glitch: relation (5.34) does not include the non-linearity of the FADC chips, and on the whole, the electronics [80]. To cut a long story short, g_i is in fact a function, not only of the mean charge μ_{Q_i}, owing to non-linearities, but also of the time since epoch t. Power failures are not extraordinary in the laboratories, and changes in the baselines propagate to the reconstructed charges q_i, and therefore to the mean μ_{Q_i}'s and the gains g_i's.

At a given time, for a given channel i, Figure 5.6 epitomises the non-linearities of the function $\mu_{Q_i} \to g_i(\mu_{Q_i}, t)$ at low charges. For charges past 200 DUQ (on this particular channel), one finds the same gain regardless of the mean charge μ_{Q_i}, and (5.34) holds as if N_i followed Poisson statistics above a certain expectation for the number of PE's. Below this transition point, $\mu_{Q_i} \to g_i(\mu_{Q_i}, t)$ must compensate for the augmented charges.

It is worth underlining that this non-linearity has nothing to do with the quenching of the light output in the liquid scintillator (cf. 5.1.3.3), since all the quantities we are massaging enter in the field at the PMT stage only. To summarise, not only may the scintillation light generated in the liquid by low-energy electrons or heavy particles be dwarfed, but this light may additionally translate into a charge different than expected, owing to non-linearities brought about by the electronics. The three-parameter piecewise function, fitted to the measurements in Figure 5.6, is kept for further purposes as the expression of the gain g_i for the i-th channel.

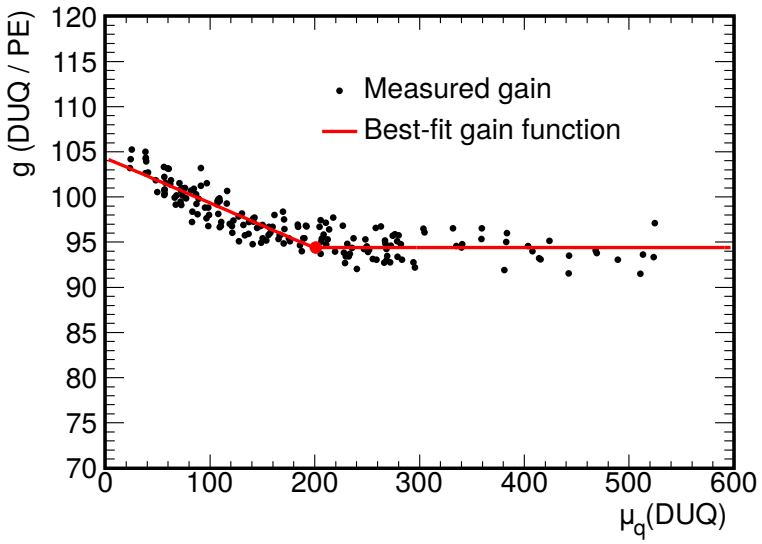

Figure 5.6 – Scatter plot of the gain of a typical channel, obtained from several calibration runs with varying light intensities. Each run provides a unique charge distribution with particular μ_Q and σ_Q values, from which a $g(\mu_Q)$ point can be obtained. A three-parameter piecewise function (red) matches the data. Adapted from [60].

Total number of PE's

The previous paragraphs gave us more than hints at estimating the number of PE's n_i^{PE} seen by each channel: for each recorded charge q_i, we need but divide it by the gain $g_i(q_i, t)$ read on the piecewise function, at this charge, and for this time t of data taking, so that

$$n_i^{PE}(q_i) = \frac{q_i}{g_i(q_i, t)}\,. \tag{5.35}$$

As a result, the total number of PE's detected by a set of channels, recording the charges $(q_1, \ldots, q_n)$ during an event, reads

$$n^{PE}(q_1, \ldots, q_n) = \sum_{i=1}^{n} n_i^{PE}(q_i) = \sum_{i=1}^{n} \frac{q_i}{g_i(q_i, t)}\,, \tag{5.36}$$

since n_{PE} is a realisation of the random variable N, defined as

$$N = \sum_{i=1}^{n} N_i\,. \tag{5.37}$$

To be precise, g_i is a little different for the Monte-Carlo and the data. Evidently, the Monte-Carlo is not plagued by time stability issues.

5.2.3.2 Spatial uniformity

Not only do the PMT's cover less than 15% of the buffer surface, but of all of them also point towards the detector centre (see Figure 4.6). To put it differently, an identical light (or radioactive) source will not yield the same mean total number of PE's μ_N, whether it is in a corner of the detector or at its centre. To be sure, it is scarcely surprising that a given channel i provides a different μ_{N_i} value whenever the source moves, however, if the source has indeed a fixed known intensity, the mean total number μ_N should be conserved regardless of the source's position.

Recalling the set up of all the calibration systems from 4.3.3, it is evident that no complete covering of the inner detector can be achieved via radioactive or light sources. As a matter of fact, in spite of the overburden described in 4.2.3, μ's keep pouring down onto the detector, producing cosmogenic backgrounds and, by and large, releasing a plethora of particles, including neutrons. Although most interact with Gd in the target, about 14% of them are captured on H. In the GC, there are few choices, and the overwhelming majority of neutrons are captured on H. In any case, the statistics are plentiful; bearing in mind that deuterium emits the very same 2.22 MeV γ-rays, without regard to the energy of the captured neutron, the so-called spallation neutrons supply a fine source for testing spatial uniformity in the inner detector.

Analysing events within 1 ms of μ-triggers – which we can spot owing to the vetoes from 4.3 – one can obtain the spatial variations of μ_N, and from there, the correction map m to apply to the observations n^{PE}. The corrected number of PE's n_c^{PE} is expressed as

$$n_c^{PE}(q_1, \ldots, q_n) = m(r, z)\;\; n^{PE}(q_1, \ldots, q_n)\,, \tag{5.38}$$

where (r, z) are the cylindrical coordinates of the event position, determined such that its reconstructed charges $(q_i)_{i\in[\![1,n]\!]}$ and times $(t_i)_{i\in[\![1,n]\!]}$ form a likely set (cf. 5.2.2.2). The map is defined such that $m(0,0) = 1$, which effectively brings all other points to the centre. The latest correction map, for data taken at the near detector, is presented in Figure 5.7. The corresponding Monte-Carlo map m^{MC} exhibits comparable features.

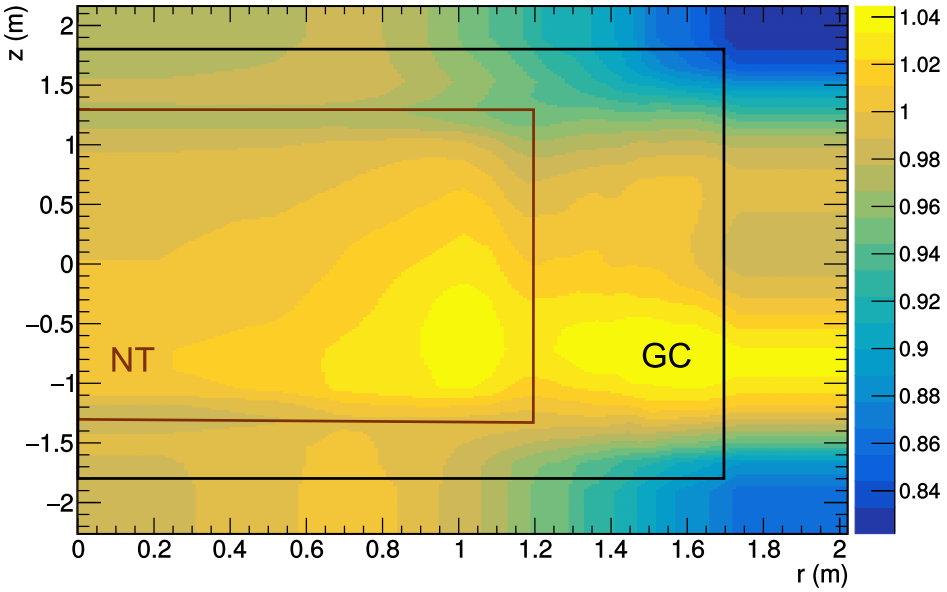

Figure 5.7 – Near detector calibration map correcting for spatial inhomogeneities in the number of detected PE's. The blue areas, close the PMT's signify that more PE's are collected there, when compared to an event reconstructed at the centre, thus, n^{PE} for these locations must be brought down. The approximative positions of the NT (brown) and GC (black) are also indicated.

5.2.3.3 Absolute energy scale

The corrected number of PE's may now be matched to a calibration source of known energy. Again, delayed γ-rays – originating from H captures – are put to practical use considering they offer an anchor in the signal region. The H capture peak is used both in the data and Monte-Carlo; the calibration factors of the two are not necessarily identical since scaling to a 2.224 MeV energy is precisely the first step into equating the data and simulations.

Taking for granted that spatial uniformity is achieved via the correction map from (5.38), anchoring the data and Monte-Carlo – with the high statistics runs from the ^{252}Cf calibration campaign – at the centre of the target, should ensure an identical energy response throughout the rest of the inner detector. The spectrum of H captures induced by a ^{252}Cf source, at the centre of the near detector target, expressed in total number of PE's[9], can be found in Figure

[9]By definition, at the centre $m(0,0) = 1$, hence $n^{PE} = n_c^{PE}$.

5.8. From the mean value of the Gaussian fit to the H peak, i.e. 415 photo-electrons, one can infer that the absolute energy scale conversion factor f_a, at the near detector, evaluates to

$$f_a = \frac{1}{186.7}\ \text{MeV PE}^{-1}\,. \tag{5.39}$$

With regard to the Monte-Carlo, simulations of the ^{252}Cf data at the near detector lead to $f_a^{MC} = 1/182.8\,\text{MeV}$[10], which happens to be stunningly close to the data scale.

As a summary of the progress made so far, once the total number of PE's n^{PE} has been corrected for its incongruous position dependency, one can re-build the energy effectively deposited in the detector E_0^{vis} as a mere linear function of the corrected n_c^{PE}, that is

$$E_0^{vis}\left(n_c^{PE}\right) = f_a\, n_c^{PE}\,. \tag{5.40}$$

Unmistakably, the 0 subscript in (5.40) is more than an decorative ornament for E^{vis} and E_0^{vis} is but a first-order visible energy, whether be it for Monte-Carlo or data (unlike often read), as we will spell out in the following.

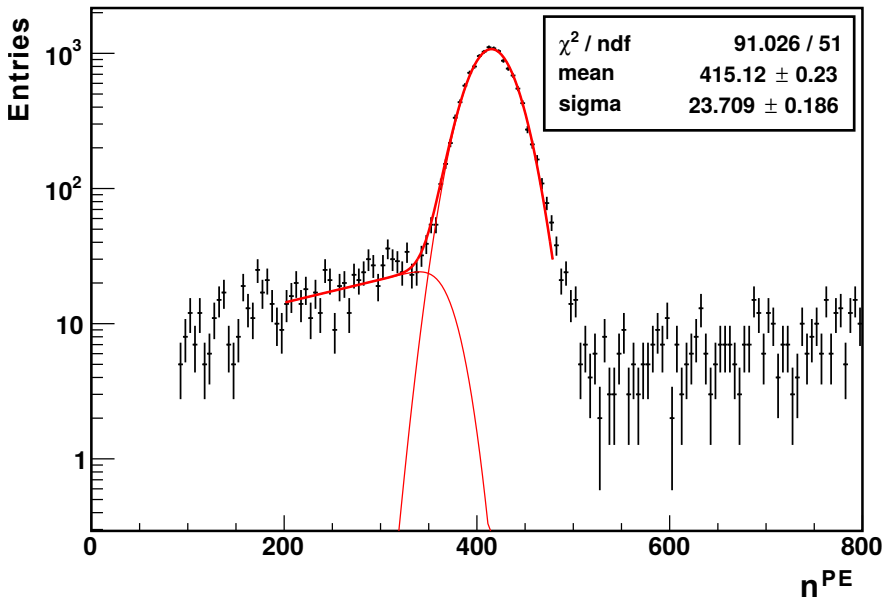

Figure 5.8 – Distribution of the total number of PE's reconstructed for H captures induced by a ^{252}Cf source at the centre of the near detector target. A fit of the 2.224 MeV H capture peak yields an energy scale conversion factor of $f_a = 1/186.7\,\text{MeV}$. Data from [101].

5.2.3.4 Corrections

Charge non-linearity

One can obtain a theoretical expectation for the mean position of the Gd capture peak.

[10] PE's are no actual unit, and they do not quite fit in one-line expressions, so we are leaving them out for now.

Unfortunately, when analysing the IBD delayed events in the upper energy window, i.e. where the Gd peak is expected, one realises that neither the Monte-Carlo nor the data reproduce the theoretical expectation. Worse, the amplitude of the discrepancy with respect to theory is different between the Monte-Carlo and the data (see Figure 5.9), in spite of the forced agreement at $E_0^{vis}[H_{Cf}]$, the H peak of the ^{252}Cf data. In a nutshell, the H anchor is not enough to ensure that the Monte-Carlo and data reproduce the expectations across the energy range of interest.

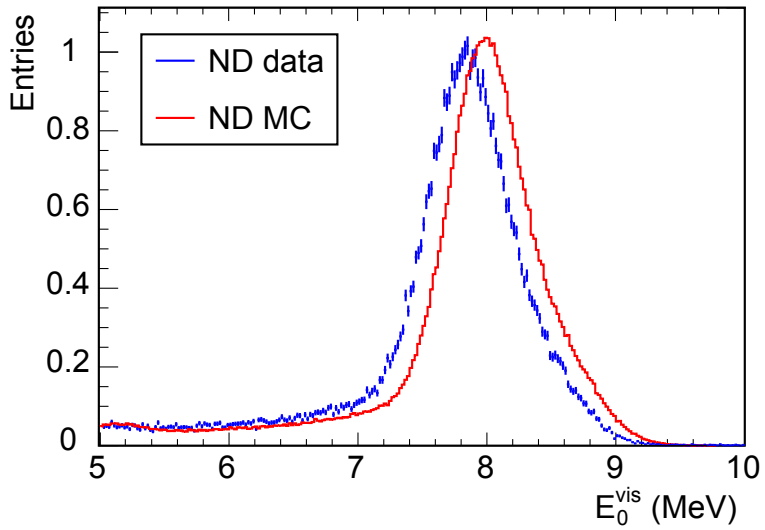

Figure 5.9 – Distribution of E_0^{vis} for the Gd de-excitation peak for the twelve months of near detector data (blue) and corresponding Monte-Carlo simulations (red). Data from [102].

Although this is a topic left for later, the values in the next paragraph would hardly fool the industrious reader: if the mean value of the H peak of the ^{252}Cf campaign at the centre evaluates to – by construction – $E_0^{vis}[H_{Cf}] = 2.224\,\text{MeV}$, the distribution of the twelve months of H captures throughout the entire inner detector[11] actually yields a mean $E_0^{vis}[H_{IBD}] = 2.215\,\text{MeV}$ at the near detector (2.209 MeV at the far detector).

In consequence, so-called charge non-linearities (shortened into QNL) corrections are applied to E_0^{vis} for the Gd peak to sit where it belongs. This correction introduces

$$E_1^{vis}\left(E_0^{vis}\right) = b_{QNL}E_0^{vis} + c_{QNL}{E_0^{vis}}^2\,, \tag{5.41}$$

where $b_{QNL} = 0.996\,522$ and $c_{QNL} = 1.570 \times 10^{-3}\,\text{MeV}^{-1}$ for the near detector data [102], so that indeed $E_1^{vis}\left(E_0^{vis}[H_{IBD}]\right) = 1 \times E_0^{vis}[H_{IBD}]$[12]. Perhaps the QNL correction is better

[11] Yes, despite the calibration map allegedly bringing all areas to the centre, more on that later.

[12] It would make sense to simultaneously shift $E_0^{vis}[H_{IBD}]$ to 2.224 MeV, instead of forcing it to retain its offset, but this is just a matter of steps before that is achieved.

understood as a unitless factor

$$f_{QNL}\left(E_0^{vis}\right) = b_{QNL} + c_{QNL}E_0^{vis}\,, \tag{5.42}$$

depending linearly on the energy, and multiplying the already obtained E_0^{vis} values. In truth, this how the b_{QNL} and c_{QNL} coefficients are obtained, i.e. by finding by how much the reconstructed $E_0^{vis}[Gd] = 7.937\,\text{MeV}$ must be multiplied – bearing in mind that the factor must evaluate to 1 for H – so that the ratio

$$\frac{E_1^{vis}[Gd_{IBD}]}{E_1^{vis}[H_{IBD}]} = f_{QNL}\left(E_0^{vis}[Gd_{IBD}]\right)\frac{E_0^{vis}[Gd_{IBD}]}{E_0^{vis}[H_{IBD}]} \tag{5.43}$$

equates the theoretical expectation $r_E = 3.569$. Assuredly, one could merge the f_a value into the b_{QNL} and c_{QNL} coefficients, thus offering a single function to be applied onto n_c^{PE}, which might baffle the reader less than all these steps. Nonetheless, the former formulation allows to appreciate the minuteness of the QNL correction.

The difference between the QNL-corrected energy and E_0^{vis} is plotted in Figure 5.10, as a function of E_0^{vis}, for the near detector data. From this plot, it stems that the Gd peak from the near detector data must be shifted upwards by about 70 keV. On the contrary, the Gd peak from the Monte-Carlo must be brought down, although by an order of magnitude less than for the data, for indeed $b_{QNL}^{MC} = 1.000\,254$ and $c_{QNL} = -1.14 \times 10^{-4}\,\text{MeV}^{-1}$ [102]. The far detector factors differ, but the reasoning is a carbon copy of what we have presented.

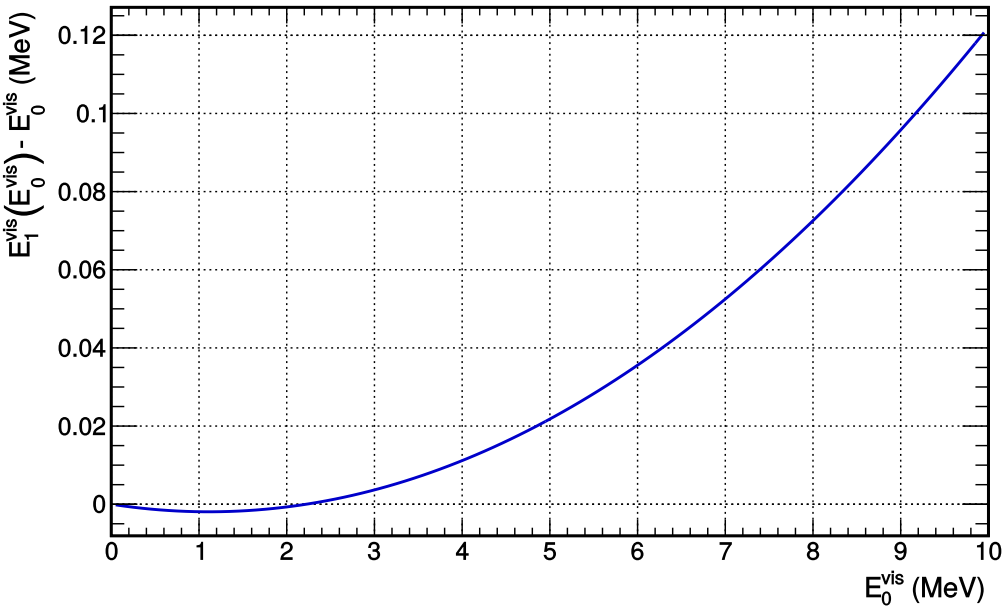

Figure 5.10 – Difference between E_1^{vis} and E_0^{vis} as a function of E_0^{vis}, for the near detector data. The curve epitomises the QNL correction which must be applied to E_0^{vis}.

Whence the "charge-non linearity" naming? Gd emits not one, but several γ-rays when de-exciting (cf. Table 4.1), as a stroke of luck, the average energy of the emitted photons is

2.2 MeV. The latter statement consequently rules out quenching of the light output as the source of QNL. Indeed, quenching applies to the energy of individual γ's, which determine the energies of the Compton electrons (cf. 5.1.3.3). To put it differently, on average, the γ's from Gd capture should be quenched by the same amount as the single γ from Hydrogen, which would leave the ratio of the total energy they deposit in the scintillator untouched. As a result, we are left with the modelling of the PMT gains: for all the efforts in quantifying the non-linearity of the charge-PE conversion, E_0^{vis} still bears the indelible stamp of QNL.

Equalisation factor

The QNL correction guarantees that the ratio of the mean reconstructed energies from the Gd and H IBD peaks match the theoretical expectation r_E. That being said, no constraints on the actual position of the peaks themselves – other than that imposed by the absolute f_a scale, embedded into E_0^{vis} – are forced upon E_0^{vis} by f_{QNL}. Somehow, the QNL correction makes sure that the disagreement of the absolute energy scale is identical for both H and Gd.

Without doubt, it is disheartening to see that using a calibration source at the centre of the detector, and correcting for the spatial uniformity by means of spallation neutrons, falls short of producing an IBD peak at 2.224 MeV. Presumably, the two reasons for the difference between $E_0^{vis}[H_{Cf}]$ and $E_0^{vis}[H_{IBD}]$ are the time stability and the relatively coarse granularity[13] of the correction map m. Surely, the time evolution of the PMT gains is taken into consideration via the light injection system, but a subtle evolution of the liquids may escape this light calibration. All things considered, twelve months is a long period; efforts have been made into incorporating a time stability factor for several years, but it does not bring improvements on all grounds, in fact, it does worsen the resolution of some other sources [102]. As a consequence, no explicit time dependent factor multiplies E_1^{vis} for the fifteen months of near and far detector data used for the Neutrino 2016 conference. Instead, a so-called "equalisation" factor f_e ensures that the central value of the H peak – and by way of QNL, that of Gd as well – agrees with the theoretical expectation.

The value of the equalisation factor for each detector is obtained by gathering all the H delayed events of the $\overline{\nu_e}$ candidates – wherever they were reconstructed and whenever they were recorded – and by isolating f_e in

$$E_2^{vis}[H_{IBD}] = f_e \, E_1^{vis}[H_{IBD}] \,, \tag{5.44}$$

where $E_2^{vis}[H_{IBD}] = 2.224\,\text{MeV}$ and $E_1^{vis}[H_{IBD}] = E_0^{vis}[H_{IBD}]$. For the near detector, this yields $f_e = 1.0039$ and $f_e^{MC} = 0.9941$. Obviously, the time scapegoat cannot justify the Monte-Carlo flaws, the discretisation of m^{MC} however does.

[13] The map originally comes into 12×12 bins to gather enough statistics, which are then split into 3000×3000 smaller bins, by linear interpolation.

Summary of achievements and prospects

All the corrections previously described can be summarised into

$$E_2^{vis}\left(n_c^{PE}\right) = f_e\, f_a\, \left(b_{QNL} + f_a\, c_{QNL}\, n_c^{PE}\right)\, n_c^{PE}\,. \tag{5.45}$$

Without exception, (5.45) is used for all types of events, including μ's, although DCReco-BAMA effectively locates the cosmic particles at the very centre where $m(0,0) = 1$. For [60], there used to be a cut-off in f_{QNL}, past the endpoint of the $\overline{\nu_e}$ spectrum, where the QNL correction skyrockets[14], but this introduced a distasteful kink in the spectra of the backgrounds just above it; beyond doubt, the high energies of μ's are not be trusted.

Undeniably, the (5.45) formulation begs for a refactoring of its parameters, but we are striving not to drift too far away from Double Chooz's conventions, while endeavouring to unscramble them. Be that as it may, these expressions should not deter us from the fact that the final energy reconstruction is excellent for the signal, both in absolute scale and in resolution, as Figure 5.11 illustrates.

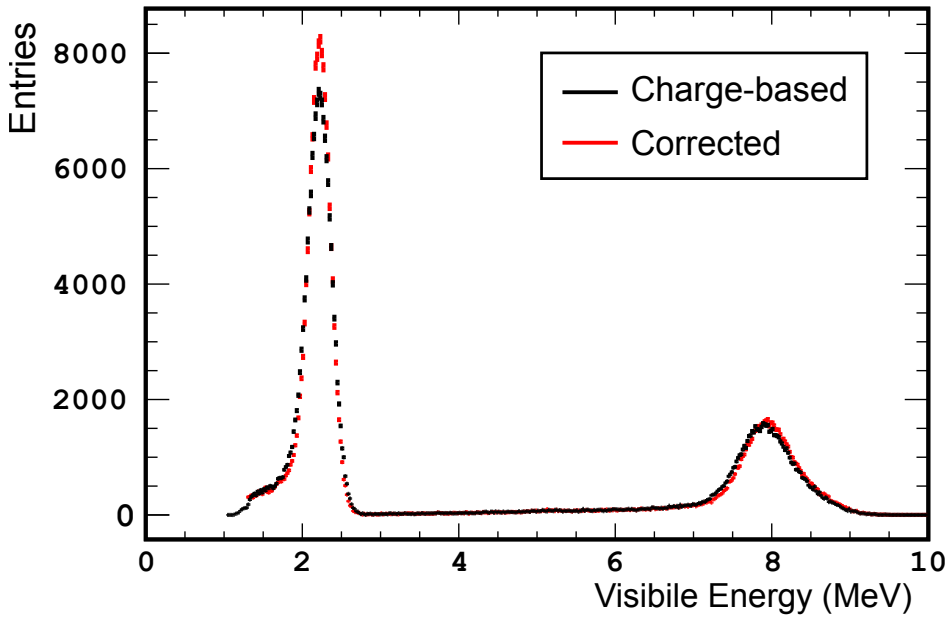

Figure 5.11 – Combined ND spectrum of delayed events from IBD's for both H and Gd captures. The "charge-based" distribution assumes a mere linear correspondence between the total charge recorded and the actual energy deposited in the liquids, anchored on the H peak; on the other hand, the "corrected" one derives from (5.45). The increased height of the H peak on the corrected version goes hand in hand with an improved energy resolution. Adapted from [103].

To complete the picture, one part is actually missing: the so-called Light Non Linearity

[14] Plotting Figure 5.10 up to $E_0^{vis} = 100\,\text{MeV}$ would demonstrate that E_1^{vis} is 15 MeV larger than E_0^{vis} at the near detector, while it would be 22 MeV higher for the current far detector data, large offsets indeed.

(LNL). Contrary to QNL, LNL turns up when normalising the energies by the number of γ-rays they correspond to. In a nutshell, LNL is related to the optical model of the liquids. In spite of the works reviewed in 5.1.3.3, DCGLG4Sim, which includes measured values such as Birk's constant, is not able to perfectly reproduce the quenched light output observed in the data. Thornier still, the Monte-Carlo simulations underestimate the energy in the GC, whereas they produce values slightly above in the NT. These flaws are embodied by Figure 5.12; comparable curves are obtained for all detectors. As we stressed earlier, this quenching is essentially non-existent for electrons past 100 keV, nonetheless, 1 MeV γ-rays scatter off several low-energy electrons, hence the effect. In short, an analysis utilising captures on H – which happen both in the NT and GC – becomes a daring enterprise.

As should be clear from 5.1.3.3, the LNL is particle-dependent and finding by how much the $\overline{\nu_e}$ prediction must be corrected is a task left for the oscillation fit (more follows in 6.3). Needless to say it turns the fit into a quite convoluted endeavour, and DCGLG4Sim craves for a proper re-adjustment of its parameters, although, in itself, such a re-adjustment is also quite the undertaking, considering the optical parameters cannot be varied analytically at the final stage. Instead, a lengthy simulation must be generated for each and every set of optical parameters, hence the LNL plaster.

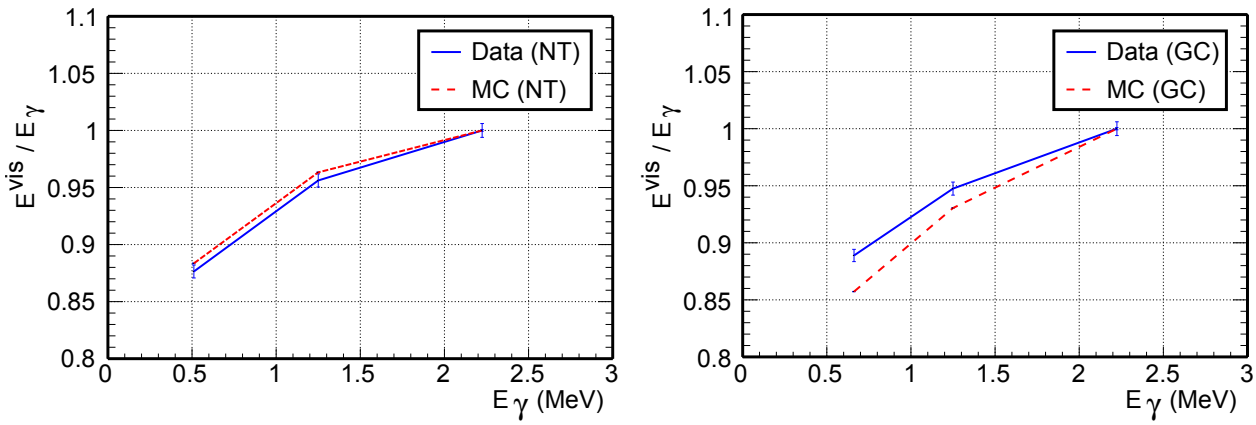

Figure 5.12 – Ratios of the visible energy from (5.45) over the actual energy of a single γ-ray E_γ at the near detector. The ratios for the simulations are represented by red dashed lines, that of data, by blue solid lines. The agreement between the Monte-Carlo and the data is greater in the NT (left) than in the GC (right). The difference between the simulations and the data changes sign between the two volumes. The plots were obtained using calibration sources. For instance, in the near detector target, ^{68}Ge generates two annihilation γ's, which matter as a $E_\gamma^{Ge} = 511$ keV source instead of a 1.022 MeV one. Based on [104].

Thereafter, we shall return to the usual E^{vis} notation to designate quantities originally retrieved as E_2^{vis}.

5.2.4 Muon tracks

The IV (cf. 4.3.2.1) and OV (cf. 4.3.2.3) offer triggers signalling μ-interactions in the detector or its surroundings. These triggers may be used to correlate IBD events to μ's close in time, but being able to support the time component with a spatial information is a must.

5.2.4.1 Making choices

Although an exhaustive reconstruction algorithm for μ-tracks – using all of ID, IV and OV – was used in Double Chooz's latest publication [60], the so-called FIDO reconstruction code (for Fused Inner Detector/veto and Outer veto) is breathtakingly time consuming, to the point that using it for both detectors – with the spring and summer conference deadlines in mind – was but utter fantasy. And if truth be told, the resolution of the former Double Chooz ID-only reconstruction algorithm, DCRecoMuHam – recently put back into fashion on efficiency grounds – is not notably worse than FIDO's. Besides, considering μ-tracks are to be compared to DCRecoBAMA-reconstructed event positions, which have no micro-metric resolutions either (see 5.2.2.3), switching back to DCRecoMuHam – after some necessary adjustments – was a slick move.

5.2.4.2 Estimation method

DCRecoMuHam – named after the whereabouts of its authors at the time of coding, i.e. Hamburg – finds the best-matching μ-track by dint of the Mean Square Weighted Deviation (MSWD) method. As in DCRecoBAMA, both the charge and time information are utilised. Regarding the charges, they are actually first converted into their PE-equivalent, which allows to use the same tweaks and cuts, briefly reviewed in 5.2.4.3, regardless of the detector.

Goodness of fit

The MSWD somehow assesses the goodness of the candidate entry and exit points; hence, minimising the MWSD value σ^2, with respect to these points, can yield a likely μ-track. Assuming an ideal Cherenkov light cone is emitted along the μ-track, each pair of candidate positions can be associated to a set of arrival times $(t_i^{pre})_{i\in[\![1,n]\!]}$, with n the number of active PMT's. Contrary to DCRecoBAMA, which predicts the charges of all the PMT's, DCRecoMuHam makes not attempt at prophesying the number of PE's each channel should observe. In truth, the numbers of PE's actually detected $\left(n_i^{PE}\right)_{i\in[\![1,n]\!]}$ are the weights of the MWSD.

In the following, we denote by x and y the space-time entry and exit points, respectively. The mean sample time shift μ_s, which is used to compute σ^2, is defined as

$$\mu_s\left(x,y\right) = \frac{\sum\limits_{i=1}^{n} n_i^{PE}\ \Delta t_i\left(x,y\right)}{\sum\limits_{i=1}^{n} n_i^{PE}}\,, \tag{5.46}$$

where $\Delta t_i\,(x,y) = t_i^{pre}\,(x,y) - t_i$, with t_i the observed arrival time of channel i. It can be noted that assuming an identical weight $n_i^{PE} = w$ for $i \in [\![1,n]\!]$, yields the usual sample mean formula, with the sample length n at the denominator. An estimator of the variance to minimise is given by

$$\sigma^2\,(x,y) = \frac{\sum_{i=1}^{n} n_i^{PE}\,(\Delta t_i - \mu_s)^2(x,y)}{\sum_{i=1}^{n} n_i^{PE}}\,. \tag{5.47}$$

Minimisation

This minimisation of σ is actually not handed over to MINUIT [105] or any counterpart, rather, its value is manually varied by recomputing the time of flights $(t_i^{pre})_{i\in[\![1,n]\!]}$ for every new guess.

The first x candidate solely corresponds to the PMT p with the earliest arrival time t_p. As for y, its starting value is defined by the PMT q which minimises the deviation Δt_{qp} of $t_q - t_p$ from the travel time one ought to expect from a lightning-fast muon[15]. This deviation reads

$$\Delta t_{qp} = t_q - t_p - \frac{d_{pq}}{c}\,, \tag{5.48}$$

where d_{pq} is distance between the p-th and q-th PMT's.

Let x_c and y_c denote the first candidates. The minimisation procedure computes $\sigma\,(x,y)$ for 9×9 values of (x,y) distributed along two Greek crosses centred around x_c and y_c, whose values are also included. On each one of their equal length arms, the Greek crosses bear two points, 10 cm and 30 cm away from their centres. If there exists a scanned pair (x_n,y_n) satisfying $\sigma\,(x_n,y_n) < \sigma\,(x_c,y_c)$, the procedure starts all over again, with crosses centred on x_n and y_n. Once the off-centre points of the Greek crosses stop offering lower σ values, the cross-centres are saved as entry and exit points, which fully defines the μ-track.

5.2.4.3 Performance

Different detector versions

As far as the near detector goes, neither FIDO nor RecoMuHam were achieving acceptable performance in their stock versions. Whereas the former still struggles with near detector data, rather simple updates could be applied to RecoMuHam by C. Jollet and the author, which ensured a resolution similar to that at the far detector, whose accuracy also improved on this occasion.

The Double Chooz tale is a quaint one; in 5.1.2.3 we underscored that the experiment was somehow a three-detector set-up, including the Bugey constraints, but in all truth, we have to deal with three detector versions, on top of Bugey, which makes it a four-detector hydra.

[15] The kinetic energy reconstruction threshold is effectively of the order the μ mass, so at worst, the Lorentz factor $\gamma \simeq 2$, i.e. the muon is travelling at $\sqrt{3}/2$ the speed of light, which is close enough to lightning-fast.

What is referred to as the "FDI" detector corresponds to the phase of the experiment without the near detector, for which about one year and a half of data have been made available to the analysers. On the other hand, the so called "FDII" detector is more than a shorthand for "far detector during the near operation"; above all, it corresponds to the replacement of certain electronic components, gain changes, and last but not least, to Dynamic Data Reduction (DDR).

Inasmuch as the μ-rate of the near detector is five times higher than the far one – on account of the shallowness of the former – letting the 50 GB far detector one-hour run size inflate was unworkable, which lead to the development of DDR. Furthermore, spurious light emission from the base circuits of the PMT's, known as "light noise" (detailed in 6.1.1), which increases every summer, currently tots up to 300 Hz at FDII, i.e. the order of the useful trigger rate. In other words, a DDR had to be applied to both detectors, to reduce the load of the acquisition chain. With regard to μ's, DDR discards the 256 ns waveforms of each channel, which are replaced by a set of characteristic parameters, computed on-the-fly. Consequently, this irreversible DDR provides slightly different arrival times and charges, which – along with the gain changes aimed at improving the $\overline{\nu_e}$ energy region – made for a poorer performance of all the μ-reconstruction algorithms at FDII and ND (near detector), when compared to FDI. The IV threshold for applying DDR demands that more than 50% of the PMT's be hit.

Benchmark

For the cosmogenic background studies, of particular interest is the lateral distance between background events and μ-tracks, defined as the shortest distance from the DCRecoBAMA-reconstructed prompt position to the line whose direction vector is the μ-track.

An astute way to compute this quantity is to notice that the lateral distance d is always the height of the triangle whose vertices are the position of the event E, the track entry point M and the track exit point N (see Figure 5.13).

Let a, b, c denote the lengths of the line segments MN, EM, EN, respectively. Trivially, the area $\mathcal{A}$ of the triangle is

$$\mathcal{A} = \frac{d \times a}{2} \,. \tag{5.49}$$

Moreover, according to Heron's formula, we also have

$$\mathcal{A} = \sqrt{s\,(s-a)\,(s-b)\,(s-c)}\,, \tag{5.50}$$

where the half-perimeter s reads

$$s = \frac{a+b+c}{2} \,. \tag{5.51}$$

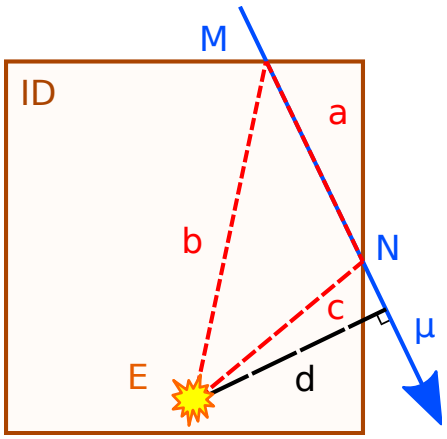

Figure 5.13 – The lateral distance d between an event reconstructed at E and a μ-track (blue) whose ID entry and exit points are M and N, respectively, can easily be expressed in terms of the sole a, b, and c triangle (dashed red) edges, as in (5.52).

Equating (5.49) and (5.50) yields[16]

$$d = \frac{2\sqrt{s\,(s-a)\,(s-b)\,(s-c)}}{a}\,. \tag{5.52}$$

There are only a few background events per day, which are supposed to be located close to μ-tracks; on account of time constraints, it was a pipe dream to reapply DCRecoMuHam – however lighter than FIDO it may be – to all the data to spot any improvement after a series of tweaks. Exactly the way spallation neutrons helped to track spatial uniformities in the energy response, neutron captures – both on Gd and H – provide abundant statistics to measure the performance of DCRecoMuHam. Considering these neutrons are produced by μ's, opening a 1 ms window after each μ, and filling an histogram with the lateral distances, as calculated from (5.52), is an affordable benchmark which demands but a few weeks of data. At the closing of the 2015 year, the collaboration was left with a reconstruction performance exemplified by Figure 5.14. The lateral distance being one the two handles of the cosmogenic background rejection technique (explained in 9.1), the performance at the near detector was deemed unacceptable for the analysis standards of Double Chooz, and FDII yearned for improvements.

Meeting the performance the requirements

Although it went first unnoticed, after having delved into the depths of the code, it became plain that DCRecoMuHam was not using the light injection calibration data – i.e. the

[16] Another canny way is to notice that d matters in the norm of the cross product $\left\|\overrightarrow{EN} \times \overrightarrow{NM}\right\| = d\,a$; from (5.52) it is clear that $\left\|\overrightarrow{EN} \times \overrightarrow{NM}\right\| = 2\,\mathcal{A}$, which is hardly surprising since $\overrightarrow{EN}$ and $\overrightarrow{NM}$ span a parallelogram of area twice that of the red dashed triangle.

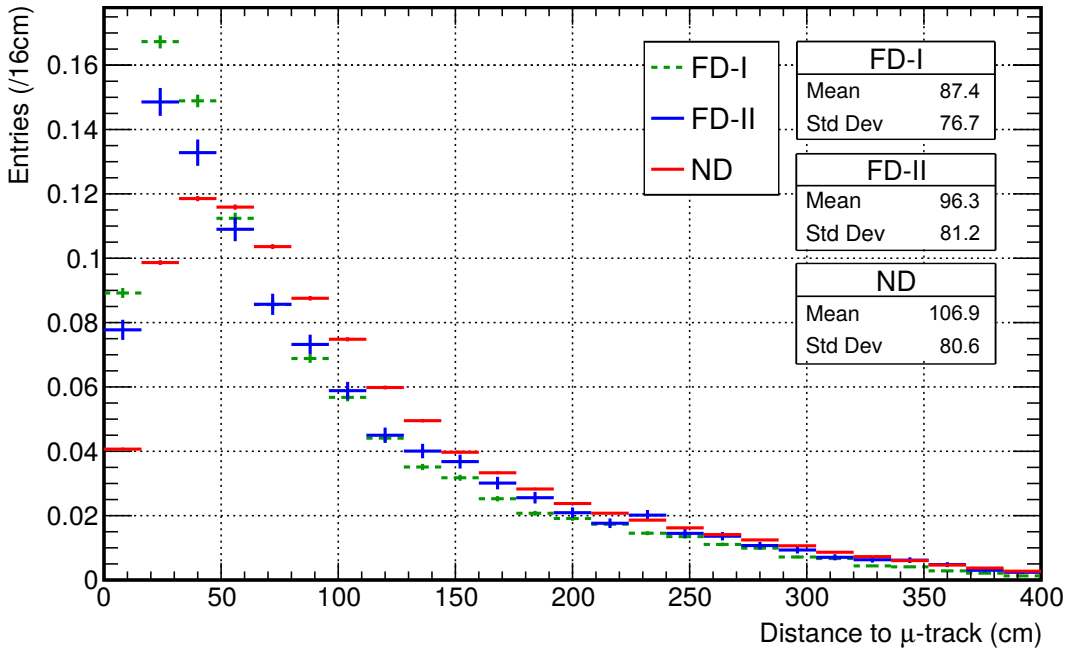

Figure 5.14 – Distribution of the distance between μ-tracks and the capture positions of the spallation neutrons following these tracks within 1 ms. The large statistics of the near detector data (ND) are unambiguous: the μ-track reconstruction performance is worse than at the far detector (FDI and FDII), the performance of which has also decreased in its last version (FDII).

information embedded in the $h_{x,I}$ densities from 5.2.2.2 – for DDR events. More precisely, the DDR arrival time t_i of channel i being defined as

$$t_i = t_i^{50} + t_i^0 \,, \tag{5.53}$$

where t_i^{50} is the DDR time at which the waveform reaches 50% of its maximum value, the calibration data comes into play in the form of the t_i^0 offset. Looking back at (5.47), and considering the non-negligible spread of the t_i^0 values [106], it is evident that such an oversight[17] was worsening the accuracy of DCRecoMuHam.

As regards the gain mismatch across all detector versions, i.e. FDI, FDII and ND, its effect is more subtle. Indeed, a brief review of (5.46) and (5.47) reveals that a global multiplication factor of the gains cancels out in the ratio, namely, only the relative difference between the weights impacts the minimisation procedure. And yet, a channel i such that $n_i^{PE} \leq n_{th}^{PE}$ with $n_{th}^{PE} = 0.2\,\text{PE}$ is excluded from the candidates. Insofar as the PMT gains for FDII and ND have nearly tripled, such gain changes imply a much higher threshold in terms of the recorded charge q_i for these detector versions, which is how DCRecoMuHam

[17]The DCRecoMuHam authors were not any more in the collaboration at the time at which DDR was proposed.

actually applies this condition. Hence, some low-signal channels were unwittingly considered as candidate entry and exit points. There is an additional threshold of 10^4 PE on the total n^{PE} to trigger the minimisation procedure so as to improve the first candidates x_c and y_c; although, considering this is an upper threshold, it could not negatively affect ND and FDII, which had overestimated n^{PE} values at first. More importantly, the μ-tracks yielding the lowest value for σ are subsequently shifted inwards, a procedure tuned on Monte-Carlo simulations [107]. The higher n^{PE}, the larger the inward shifts to x_c and y_c, with n^{PE} values overly boosted, this made for unbalanced shifts.

Correcting for all these misfortunes – in several steps, for the Moriond and Neutrino conference – eventually lead to a wondrous performance of the μ-track reconstruction at ND and FDII, as can be gazed at in Figure 5.15. For comparison, the best FIDO performance, i.e. at FDI, is also presented. Observing Figure 5.15, it may seem as though FDII and ND perform even better than FDI. In all honesty, the μ-energy definition at FDI is quite off that of FDII. Nevertheless, for whimsical reasons, the Double Chooz analysis defines μ's as energy deposits larger than $E_\mu^{th} = 100\,\text{MeV}$, regardless of the detector version[18], hence the lawful plot. For all the efforts in the $\overline{\nu_e}$ region, the science of μ-energies – all the more so that most FADC's are heavily clipped owing to the gain increase at FDII and ND – is an obscure one. Still, matching the peaks of the μ-energy spectra, the current FDII E_μ^{th} seems closer to 170 MeV in FDI; using such a threshold demonstrates an ever so slightly better reconstruction at FDI, when compared to its counterparts.

[18] As we will see, this has little impact on the cosmogenic background studies since most background-producers are anyway well above these thresholds.

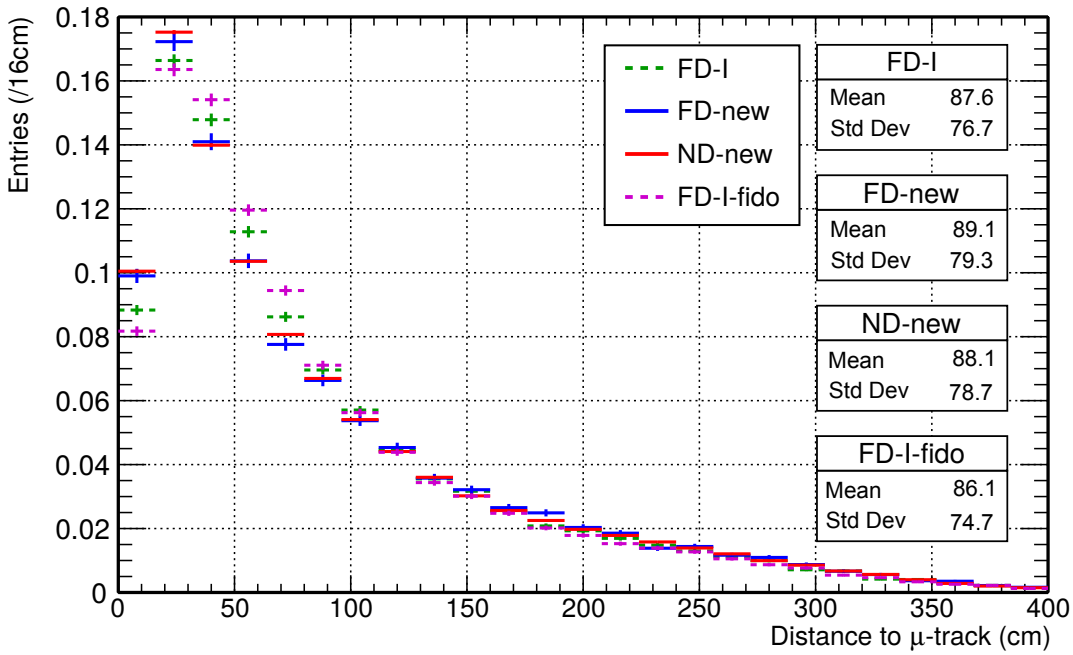

Figure 5.15 – Improved distributions of the distance between μ-tracks and the capture positions of the spallation neutrons following these tracks within 1 ms. FD-new stands for FDII with the corrected algorithm.

Chapter 6

Measuring θ_{13}

By virtue of Chapter 5, we currently have at our disposal all the pre-requisites for a first-class θ_{13} analysis, which entails counting how many $\overline{\nu_e}$'s have had the luck to reach our detectors in mint condition. Without doubt, there must be background hiding in this lot, which must either be hunted down and removed from the precious data sample or studied, and understood, to the point of becoming completely harmless to the sensitivity of the experiment. This sensitivity to the θ_{13} value is also strongly determined by all the uncertainties arising from the limited knowledge of the detectors and the operation of the reactors. Thankfully, these sources of uncertainty bear correlated components, which can be dwarfed by a thoughtful formulation of the analysis, thus making the most of the multi-detector set up.

6.1 Event selection

Dynamic Data Reduction (cf. 5.2.4.3) was promoted for obvious reasons: although they may help to reject correlated backgrounds, most events recorded by Double Chooz have nothing to do with $\overline{\nu_e}$'s. As a case in point, if about 12 $\overline{\nu_e}$'s are detected – on average – during a one-hour run at the near detector, the rate of μ's crossing its ID is about 50 Hz, and this rate including the IV is nearly five times larger, not to mention the by-products of μ-interactions. In other words, before we even contemplate looking for positrons and neutron-captures correlated in space-time, we must needs drastically reject triggers that cannot be related to fission events in the reactor cores, thereby introducing the notion of "valid trigger" or "single".

6.1.1 Singles

6.1.1.1 Muon

By means of the IV, most μ's crossing the detector can be tagged, and subsequently removed from the so-called singles. So to speak, there is no energy calibration in the IV, and thresholds

on the FADC charges q_{IV} (in DUQ's) are applied. The latter are obtained by studying the distribution of the total IV charge q_{IV} versus the total ID charge $q_{ID} = \sum_1^n q_i$ and identifying the concerned area [108]. Due do different gains, the thresholds have different values for the three Double Chooz detectors, but they are believed to approximately correspond to 16 MeV deposits. Events are also considered to be μ's if they deposited more than 100 MeV in the ID, however stealth they might have been in the IV.

For most purposes, all the aforementioned thresholds could be applied using the "mini-charge", computed on-the-fly for all events for DDR, since the agreement between the mini-charge and the actual charge – reconstructed as in 5.2.1 – is satisfactory enough for such high values. Nevertheless, the official Double Chooz selection demands to check whether DDR has effectively been applied to a given event, and cut on the thoroughly reconstructed charge otherwise. It is worth mentioning that OV data – when available – are used to reject IBD pairs themselves, when the latter coincide with OV triggers; the OV intervenes neither in the μ-definition nor in the valid trigger one.

6.1.1.2 After muons

Spallation neutrons – a boon for so many calibration purposes – are rejected from the single sample as well. Indeed, neutron captures, with the 2.2 MeV or 8 MeV γ-ray's they induce, are perfectly valid positron candidates. Furthermore, since several neutrons may be released in a given μ-interaction, a candidate pair may easily be formed out of the neutron shower. As a result, triggers following a muon-event (as defined in 6.1.1.1), within 1 ms, are excluded from the single sample. The duration of this veto is chosen so as to minimise the number of candidates correlated to μ triggers [109]; it exceeds, by a large margin, the mean capture time of neutrons on the Gd nuclei of the target $\tau_{\mathrm{Gd}} \simeq 30\,\mu\mathrm{s}$.

Although they were not used for determining the $\sin^2(\theta_{13})$ value in 6.3, IBD's whose neutron got captured on H are used in III and in the new Double Chooz analysis[1], the results of which should be disclosed before the end of the year. In light of the smaller neutron-capture cross-section of H, a longer 1.25 ms veto is used for such analyses, which is more than six times longer than $\tau_{\mathrm{H}} \simeq 200\,\mu\mathrm{s}$. Notwithstanding the length of the veto, the H sample obtained after having applied this sole after-μ veto bears a sub-percent contamination of 2.2 MeV γ-rays as prompt events. This 1.25 ms veto already induces a tremendous dead-time at the near detector, larger than 25% (about 5% at the far detector); in other words, the veto can only be increased at the costly expense of detection efficiency. As a stroke of luck, this contamination is incidentally removed by another veto (more to follow in 9.2).

[1] Here, we have in mind the analysis combining captures on Gd and H; the H-only measurement was already published as [110].

6.1.1.3 Light noise

Phenomenon

Unforeseen spontaneous light emissions from the PMT's themselves were observed during the commissioning of the far detector. Additional studies of this so-called "light noise" (LN) phenomenon were carried out in Madrid and made plain that the light emission was mainly originating from the transparent epoxy resin covering the electric components of the PMT bases [111]. Regrettably, both the rate and the amplitude of LN increase nearly irreversibly with temperature rises, as long as the photo-cathodes of the PMT's are held at a high voltage (above 1 kV). As a consequence, the far detector eventually received a dedicated air conditioner during this summer. Thanks to the knowledge acquired from the commissioning of the far detector, the near detector had its PMT bases covered with black sheets, and so far, the light noise cuts have not had any event to remove at ND, although they are also applied there for consistency.

Let us emphasise that the PMT bases directly emit light, as the light injection system would; in a nutshell, there is no conversion from energy to light in the case of LN. Still, the amount of light generated by the epoxy bases corresponds to what particles would deposit if they had energies in the range of interest for the fit; this range encompasses the prompt spectrum generated by positrons in IBD's, and the shape of certain backgrounds extending up to 20 MeV.

Rejection

Notwithstanding the amplitude of the energies they correspond to, LN events might not seem so much of a difficulty for they happen on the buffer walls, namely far from the fiducial volume. Unfortunately, DCRecoBAMA performs poorly for such remote light patterns; against all odds, LN events are actually reconstructed in the centre of the detector, thus discarding all hopes of a position-influenced rejection.

Lower-level cuts are utilised instead. The first cut implemented to reject LN is based on the fact that the PMT flasher itself should correspond to a channel k recording the largest charge $q_{max} = q_k$, and that this k-th PMT should have a greater share of the total charge recorded q_{ID}, when compared to actual energy deposits in the detector. It was found that

$$\frac{q_{max}}{q_{ID}} \geqslant 0.2 \tag{6.1}$$

minimises the IBD inefficiency at FDII (below per mille) [112]. The same cut is applied at ND, and FDI uses cut values from [60]. For DDR purposes, a conservative 0.3 threshold discards the waveforms of events unquestionably qualified as LN [113].

Considering the actual location of the light flashes, a great spread in the arrival times of the PMT's σ_T is expected on an event-basis. Similarly, the spread of all the charges σ_Q is demonstrably larger for LN than for singles. Therefore, a second cut excludes events at FDII and ND satisfying

$$\sigma_T \geqslant 36\,\text{ns} \quad \&\& \quad \sigma_Q \geqslant 1680\,\text{DUQ} - 28\,\text{DUQ ns}^{-1}\,\sigma_T\,, \tag{6.2}$$

where the second condition somehow acknowledges that charge spreads above 672 DUQ cannot account for particle interactions. Likewise, it implies that time spreads larger than 60 ns are deemed unacceptable, however small their σ_Q may be. As a matter of fact, for DDR, a safe $\sigma_{T_{DDR}} > 65$ ns cut is applied.

Eventually, a third cut, albeit somewhat redundant with (6.1), makes the most of the local character of the reconstructed charges. Also focusing on the PMT k yielding the largest charge $q_k = q_{max}$, this third cut assesses the weighted deviation of the charges of the neighbours of k with respect to q_{max}, i.e.

$$q_{dev} = \frac{1}{n_k} \sum_{i \in \mathcal{U}_k} \frac{(q_k - q_i)^2}{q_i} , \tag{6.3}$$

where U_k defines the 1 m-radius sphere forming the neighbourhood of k, which tots up n_k PMT's. Events with a weighted deviation considered to be large, that is for which

$$q_{dev} \geqslant 10^5 \,\text{DUQ} , \tag{6.4}$$

are removed from the single sample.

To summarise, any event satisfying any one of (6.1), (6.2) or (6.4), is regarded as an invalid trigger. This defines the global LN cut, whose marvellous cleansing effect is depicted in Figure 6.1, thereby allowing to effortlessly seize the Gd capture peak at 8 MeV, and the highest γ peak from natural radioactivity, at 2.6 MeV, induced by the β-decay of ^{208}Tl into ^{208}Pb.

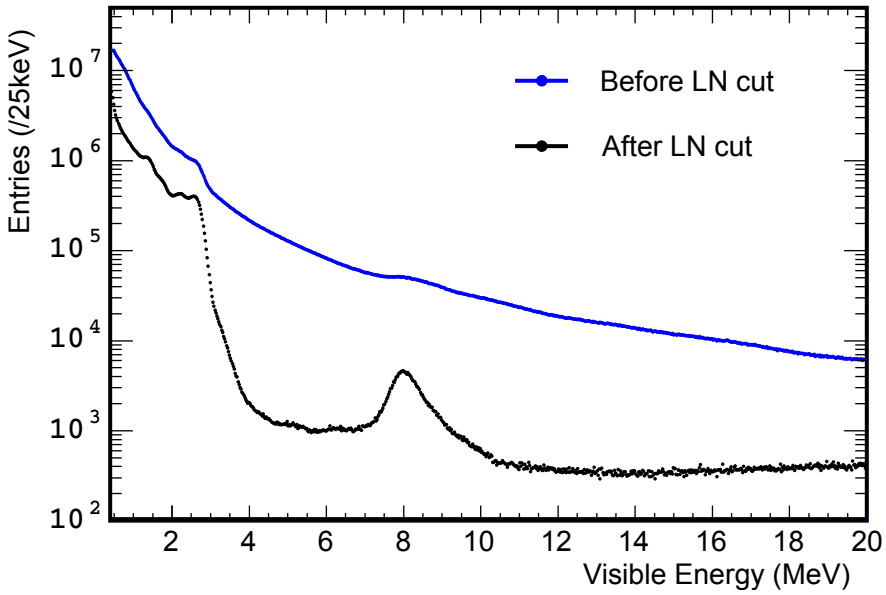

Figure 6.1 – Typical far detector visible energy spectrum of singles, before and after the global LN cut has been applied [60].

6.1.1.4 Energy range and summary

Eventually, the single definition embraces considerations about the expected prompt spectrum of IBD positrons and backgrounds. In accordance with (4.5), the prompt signal should lie from 1 MeV upwards, and yet, in order to obtain a 100% trigger efficiency with negligible uncertainty, the lower threshold applied to the energies obtained via (5.45) is $E_{s_1}^{vis} = 0.4\,\text{MeV}$. Aside from ensuring a flawless trigger efficiency, $E_{s_1}^{vis}$ provides a handle on backgrounds and retains a margin to identify any unexpected features when seeking correlated pairs within the singles. Although the oscillation fit does not include bins above 20 MeV, the upper bound $E_{s_2}^{vis} = 100\,\text{MeV}$ for the single definition is complementary to the μ-definition from 6.1.1.1; above all, it allows to characterise backgrounds, such as μ's stopping in the detector and neutrons penetrating it, whose prompt spectra extend up to 60 MeV. In addition, contemplating the relative scarcity of events between 60 MeV and 100 MeV provides trust in the understanding of the backgrounds.

To summarise, the single sample consists of events recorded by the DAQ (cf. 4.4.1) which are neither μ's, nor after-μ's, nor LN, and whose reconstructed visible energy E^{vis} lies within the interval $\left[E_{s_1}^{vis}; E_{s_2}^{vis}\right]$.

6.1.2 Pairs

6.1.2.1 Space-time correlation

With a clean sample of singles, prompt positrons and delayed neutron captures correlated in space-time may be sought after.

In light of the characteristic capture times given in 6.1.1.2, an analysis aimed at retrieving the fast high-energy delayed signal embodied by Gd capture is well-fitted with a time coincidence $\Delta t_{pd} = t_d - t_p$ lower than 150 µs, with t_d the trigger time of the delayed event, and t_p, that of the prompt. If longer ΔT_{pd} windows do integrate over more accidental coincidences, the latter are easily measured; thornier is the phenomenon by which neutrons – with a longer life-expectancy in the GC – break into the NT and are captured on Gd. Looser time coincidence bounds not only increase the probability that neutrons spill into the target, it also increases the discrepancy between the models predicting it [69], hence the 150 µs cut-off.

As far as the lower bound goes, μ-decays happening with a 2.2 µs lifetime, a non-zero Δt_{pd} threshold helps to discards a few Michel electrons, which bear a likeness to γ's following neutron captures. This threshold may not be too large for the neutron-capture cross-section has a strong energy dependence; significantly increasing the Δt_{pd} threshold entails an exceptionally reliable modelling of neutron scattering during thermalisation, a knowledge we could but boldly claim. All things, considered, the Gd analysis demands that

$$0.5\,\mu\text{s} < \Delta t_{pd} < 150\,\mu\text{s}\,, \tag{6.5}$$

whereas the inclusion of H captures extends the upper bound to 800 µs.

Additionally, a space correlation cut reduces the number of accidental coincidences at the expense of a mere per mille signal loss [60]. Denoting by Δr_{pd} the distance between the prompt and delayed DCRecoBAMA-reconstructed positions, the spatial coincidence cut reads

$$\Delta r_{pd} < 1\,\mathrm{m}\,, \tag{6.6}$$

while it extends to 1.2 m for H-based analyses, by reason of the meagre H neutron-capture cross-section.

6.1.2.2 Energy windows

Prompt

The energy windows in which the prompt and delayed signals lie are further reduced compared to that of 6.1.1.4. The prompt energy window is

$$0.5\,\mathrm{MeV} < E_p^{vis} < 20\,\mathrm{MeV}, \tag{6.7}$$

which undoubtedly shrinks the upper part of the window for singles. Doing so provides an independent set of events with $20\,\mathrm{MeV} \leqslant E_p^{vis} \leqslant 100\,\mathrm{MeV}$, from which some backgrounds may be measured, and extrapolated downwards, without introducing any correlations or adding to the complexity of the oscillation fit.

The H prompt threshold is raised to 1 MeV, mainly to reject events exclusive to the GC: IBD's in the buffer deposit no visible energy, and yet, they release annihilation γ's, one of which may easily reach the GC and produce a $\sim 0.5\,\mathrm{MeV}$ trigger. Unmistakably, such a trigger can hardly be matched to the kinetic energy of the positron that gave birth to it. In short, better remove these low-energy accidents than throwing dices.

Delayed

To maximise the detection efficiency and reduce the systematic uncertainties associated to it, the delayed energy window is broad, thus encompassing all the Gd de-excitation energy peaks mentioned in 4.1.1. The visible energy of the delayed event E_d^{vis} must satisfy

$$4\,\mathrm{MeV} < E_d^{vis} < 10\,\mathrm{MeV}\,, \tag{6.8}$$

which epitomises the Gd analysis.

When referring to the pure H analysis, we will have in mind delayed energy bounds of $1.3\,\mathrm{MeV} < E_d^{vis} < 3\,\mathrm{MeV}$. When mentioning the combined or so-called "Gd++" analysis, the delayed window of interest will be taken as that comprised of the two capture peaks, including the signal desert in between (cf. Figure 5.11), which, however full of accidentals, is effectively removed by an artificial neural network (ANN). The Gd++ delayed window $1.3\,\mathrm{MeV} < E_d^{vis} < 10\,\mathrm{MeV}$ is a generous one indeed.

6.1.2.3 Isolation cut

Despite the after-μ veto applied to define single events, flocks of neutrons may still produce phony signals. In consequence, an isolation condition stipulates that there should be no other valid singles except for a unique delayed trigger within a window centred on the prompt candidate.

The oscillation fit results given at the end of the chapter, stemming from the Gd analysis presented at the Moriond 2016 conference, utilise a $[t_{u_1}; t_{u_2}]$ isolation window with $t_{u_1} = -200\,\mu s$ and $t_{u_2} = 600\,\mu s$. After the Moriond conference, it was found out that at ND, the distribution of the number of events rejected by t_{u_1} was not flat until $t_{u_1} = -600\,\mu s$, hence the alleged presence of a correlated background [114]. This larger window only increased the isolation inefficiency to 0.46% while rejecting an additional $0.44\,d^{-1}$ rate of correlated events. This larger unicity window is part of the results in III.

With regard to analyses including H captures, the isolation window is characterised by $t_{u_1} = -800\,\mu s$ and $t_{u_2} = 900\,\mu s$, in accordance with the longer capture-time of neutrons in the GC.

6.2 Backgrounds

Notwithstanding the uniqueness of the IBD signature, unfortunate or correlated combinations of events may mimic it. Therefore, the expected number of $\overline{\nu_e}$ interactions per energy bin, obtained from (5.17), must be corrected for the expected number of background events, provided one wishes to compare the prediction to the observation to try and deduce a $\sin^2(2\theta_{13})$ value. Assuredly, backgrounds may be partially removed by vetoes, but the remaining contamination must always be evaluated.

As from now on, we mostly focus on the so-called Moriond background rejection, i.e. on candidate pairs whose delayed event is a Gd capture, which corresponds to the oscillation fit results presented at the end of this chapter. The so-called Neutrino selection has a slightly different stopping-μ rejection, but all the results in III are virtually immune to these nit-picking changes.

6.2.1 Accidental background

All the more so in the Gd analysis, the lower 4 MeV bound on the delayed energy window cuts down dramatically on accidental coincidences. Similarly, the shielding and radiopurity efforts described in 4.3 help reducing the number of random coincidences happening each day. If the prompt energy can originate from natural radioactivity, a 4 MeV threshold leaves only room for neutron-captures on Gd, ^{12}C (4.4 MeV from the de-excitation γ-ray) or ^{56}Fe, the β-decays of cosmogenic nuclei with a sufficiently high endpoint (13.4 MeV for ^{12}B), and light noise.

Although there exists, so to speak, no veto for accidental events – aside from desperately shrinking all the windows and rising the energy thresholds – what contamination remains, can easily be assessed. Indeed, by definition, the accidental background consists of two random coincidences; building up dummy pairs by choosing singles in the right energy windows and spatially close to one another, albeit remote in time, i.e. uncorrelated, not only provides the number of such unfortunate coincidences but also the energy shape of the accidental background.

More rigorously, a selection procedure identical to that described in 6.1.2 is applied with the only difference being that the time coincidence window is offset by several seconds, ensuring independence of the prompt and delayed candidates. Besides, several so-called off-time windows are used to increase the statistical significance. Subsequently rescaling the numbers and spectra thus found to the length of one 149.5 µs window, the accidental rates r_{acc} from Table 6.1 can be obtained for all three detectors versions. The larger rate at FDII, when compared to FDI, is rooted in a light noise increase and the switching on of certain PMT's whose bases are particularly keen on flashing.

Detector	r_{acc} (d^{-1})
FDI	0.070 ± 0.003
FDII	0.106 ± 0.002
ND	0.344 ± 0.002

Table 6.1 – Daily accidental rates for the three Double Chooz detectors [115].

The ND prompt spectrum obtained from the off-time window method is displayed in Figure 6.2. Its shape is a particular reminder of the low-energy region of the spectrum of singles, dominated by natural radioactivity. After having normalised Figure 6.2 by the 150.76 d effective live-time[2] of the data set dedicated to Moriond, the background expectation to add to the $\overline{\nu_e}$ prediction in each visible energy bin, can be extracted.

6.2.2 Fast neutrons and stopping muon's

The correlated background contamination left after all the vetoes exhibits a flat energy spectrum above 12 MeV; despite being fundamentally different processes, the amplitude of the fast neutrons and stopping μ's backgrounds is conjointly measured. In addition, both backgrounds are produced by cosmic μ's and unlike the cosmogenic decays from 6.2.3, fast neutrons and stopping μ's account for prompt signals less than 1 ms after an IV trigger.

[2]The effective live-time already includes the dead-time induced by the after-μ veto.

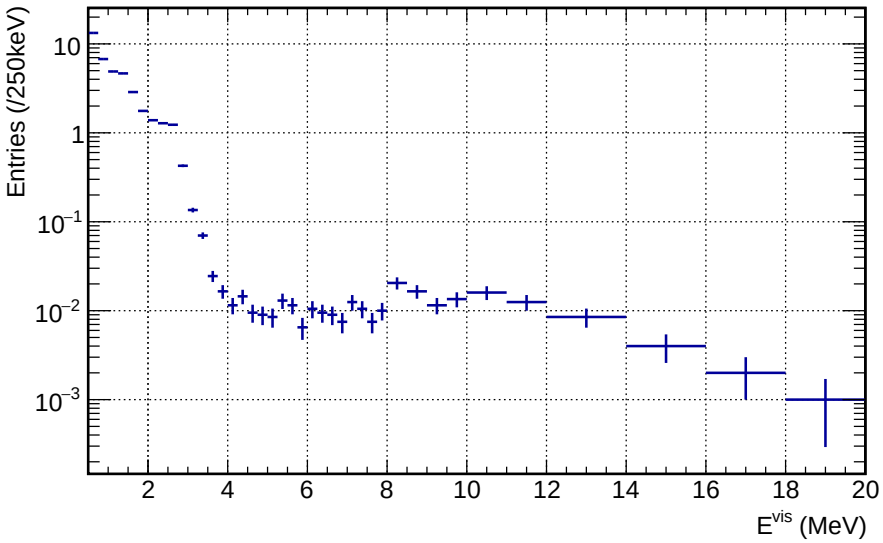

Figure 6.2 – Prompt spectrum of accidental events at the near detector for the Gd analysis. Data from [116].

6.2.2.1 Stopping muons

IV and OV

The energy deposits of μ's are large enough to generate signals in the IV and OV whenever they cross them, which translates into an painless removal of most stopping μ's with negligible dead time. Prompt triggers reconstructed less than 3.7 m away from $\sim$ 0.2 MeV strong IV triggers[3], and whose time coincidences Δt_{IV-p} fall within

$$-10\,\text{ns} < \Delta t_{IV-p} < 60\,\text{ns}\,, \tag{6.9}$$

are rejected. The negative bound on $\Delta t_{IV-p} = t_p - t_{IV}$, with t_{IV} the IV trigger, leaves room for some flaws in the time reconstruction and synchronisation. Regarding the OV, prompt candidates coincident with OV hits within 224 ns are normally vetoed, but not OV data was available to the main analysers for the Moriond conference.

Ideally, the IV and OV would be running all the time and cover a gap-less area above the ID. In practice, the OV failures and the chimney are loopholes in the tagging of stopping μ's, which – when stopping in the upper part of the detector – may deposit energies low enough to fall in the prompt window. Their decays into the lightest charged leptons generate energies for which the overlap with the delayed window is non-negligible.

Delayed reconstruction veto

For each set of recorded charges and times, the DCRecoBAMA $\mathcal{F}$, as defined in (5.30), is

[3]Despite the reflective surfaces mentioned in 4.3.2.1, there is some sort of position reconstruction in the IV .

a continuous function of the position and light intensity (x, I). The minimum value of $\mathcal{F}$ for an event defines its likely position (x_m, I_m), and $D = \mathcal{F}(x_m, I_m)$ provides a measure of the reconstruction inconsistency[4], for indeed, $\mathcal{F}$ is the opposite of the logarithm of the likelihood. Akin to a χ^2, D measures the deviation of the best fit (x_m, I_m) with respect to the expectation, as obtained from the optical and statistical models.

For unusual events such as stopping μ's decaying in the detector, the energy deposits are not point-like and the deviation from the model is large. A rejection of high D values efficiently removes the background. It is worth making note of the fact that light noise events also produce large reconstruction inconsistencies and are, accordingly, further tamed down by this reconstruction veto[5]. So as to prevent any bias in the $\overline{\nu_e}$ spectrum, only the reconstruction deviation of the delayed event D_d, is considered. Instead of ruthlessly throwing meaningless numbers at the reader, let us give some expressive names to the parameters of the veto, which reads

$$E_0^{vis} \exp\left(\frac{D_d}{D_d^{bg}}\right) \geqslant E_d^{vis}\,, \tag{6.10}$$

where D_d^{bg} is the characteristic deviation expected for backgrounds (a value too low would reject all delayed), and E_0^{vis} is the maximum delayed energy that an event with a flawless reconstruction ($D_d = 0$) can have to be rejected. The cut (6.10) can be understood as "if the reconstruction inconsistency D_d is far too large for the small energy deposit E_d^{vis} represents, reject the pair". Studying 2D plots such as Figure 6.3, $E_0^{vis} = 120\,\text{keV}$ and $D_d^{bg} = 1.6$ were accepted for the delayed reconstruction cut at FDII and ND.

Chimney veto

A second veto, albeit with some redundancy with the previous one, increases the stopping μ rejection in a way largely[6] independent of the prompt energy.

Doubtlessly, the stopping μ field in Double Chooz surpasses the impenetrability of the energy reconstruction. The official name of the veto is "Chimney Pulse Shape (veto)"(CPS), and it is designed to tag unusual deposits near the chimney. The "pulse shape" part refers to the distribution of the time-of-flight-corrected arrival times in all the channels, as deduced from the considered position of interaction. The CPS veto is peddled as the normalised likelihood of the pulse shape if the alleged position of the event is in the chimney. It is said that the higher the likelihood at the chimney, the more likely the event is a valid prompt, i.e. not a stopping μ. Could it do with more sense? Unquestionably. Let us rewind the tape: the CPS cut certainly does not measure how likely the distribution of arrival times would be, were the event right in the chimney; as a review of the code unveils, what it actually

[4]Double Chooz presents this as the "functional value" (veto), but what does it tell you about the meaning of the value or the function? Hardly anything to most Double Chooz members, all the less to the laymen.

[5]In fact, for Neutrino 2016, the delayed reconstruction veto, as I call it, was tuned to remove light noise events.

[6]In truth, the use of the pulse distribution of the prompt event introduces a very slight bias, hence the qualification; only the delayed hit times are utilised for the Neutrino 2016 selection.

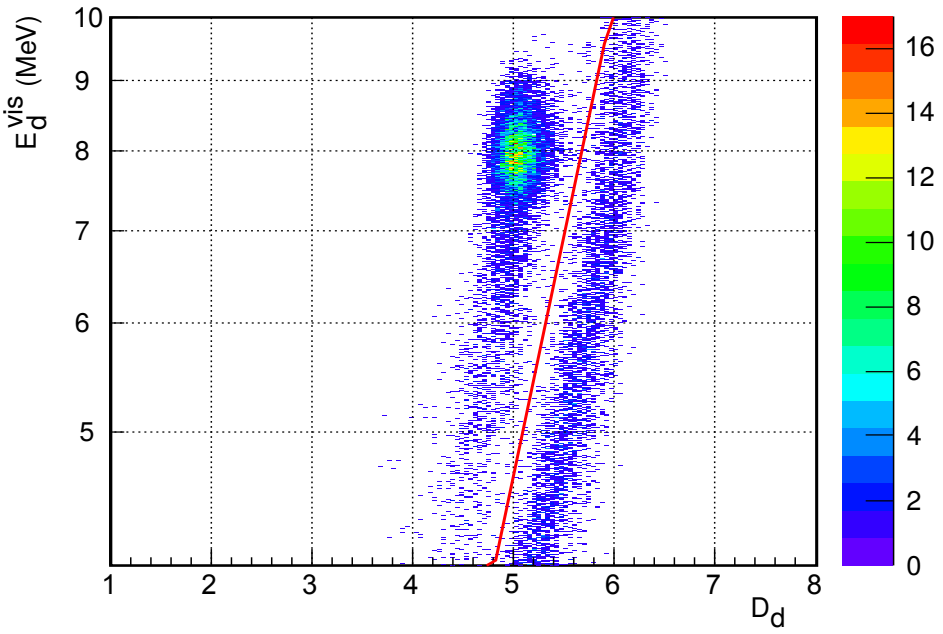

Figure 6.3 – Scatter plot of E_d^{vis} (logarithmic scale) versus the reconstruction deviation from a point-like pattern D_d , for the Gd analysis at the near detector. The Gd capture peak sits on the left of the red solid line while stopping μ decays are located on the right and have higher deviations. Adapted from [117].

measures is how inconsistent this time distribution is for an event at the chimney. The more inconsistent the trigger looks when assuming it happened in the chimney, the more likely the event is not a background.

Denoting by $\mathcal{I}_{ch}$ the inconsistency of the pulse assuming the event is located in the chimney, and by $\mathcal{I}_x$ its inconsistency when taking its DCRecoBAMA position x as the truth, events with large $\mathcal{I}_{ch}/\mathcal{I}_x$ ratios are kept as validate candidate pairs. Applying this veto to the prompt and delayed events, leads to the rejection of events with

$$\frac{\mathcal{I}^p_{ch}}{\mathcal{I}^p_x} + \frac{\mathcal{I}^d_{ch}}{\mathcal{I}^d_x} \leqslant 2\,, \tag{6.11}$$

namely stopping μ's since the latter only exhibit a small inconsistency when reconstructing their prompt and delayed events (p and d subscripts) at the chimney. Since the prompt and delayed events for stopping μ's exhibit a characteristic 2.2 µs life-time, comparing the rejection of events with $\Delta t_{pd} < 10\,\mu$s and those farther apart helps to set the cut value from (6.11).

Buffer stopping μ's

Last but not least, considering ND liquid scintillator is allegedly leaking into the buffer, a high sensitivity of the ND to so-called "buffer stopping μ's" has been observed. Again,

stopping μ's share features with LN, and it is scarcely surprising that a cut using q_{max}/q_{ID} proves powerful. The dedicated cut can be re-written in an intelligible form, such as

$$\frac{q_{max}}{q_{ID}} \geqslant \left(\frac{E_{th}^{vis}}{E_{p}^{vis}}\right)^{g} \tag{6.12}$$

with $g = 0.42$ and $E_{th}^{vis} = 325\,\text{keV}$. Since $0 < \frac{q_{max}}{q_{ID}} < 1$, events with $E_{p}^{vis} \leqslant E_{th}^{vis}$ would certainly pass the veto unscathed. However low, E_{th}^{vis} can consequently be regarded as the threshold as from which events are actively vetoed. Low-energy events produce less homogeneous light; moreover, the stopping μ spectrum extend to high energies, thus, it appears reasonable to be less forgiving of high q_{max}/q_{ID} values as E_{p}^{vis} grows.

Be that as it may, the explicit use of E_{p}^{vis} in a veto raised the hackles of some, and it has now been dropped for a more nebulous cut, relying on reconstruction discrepancies between DCRecoBAMA and an alternative Japanese algorithm [118].

6.2.2.2 Fast neutrons

Fast neutrons are created by μ spallation on nuclei outside the detector; their neutral nature makes them difficult to tag in the IV, and their scattering off protons in the liquid scintillator may lead to proton recoils, whose corresponding energy deposits may satisfy the criteria for a prompt signal. Fast neutrons usually come in numbers, and another neutron can capture on Gd. As they are related to the thermalisation time of neutrons, the delayed signals exhibit undistinguishable time features when compared to neutron captures from IBD interactions.

Nevertheless, some coincidences of prompt-delayed pairs with IV and OV triggers allow to extract the energy shape of this signal, which is flat within uncertainties (see Figure 6.4). By reason of the vetoes from 6.2.2.1, about 95% of the stopping μ's are eliminated from the $\overline{\nu_e}$ candidates [117], thus, this flat shape mainly accounts for proton recoils.

Insofar as the other backgrounds and the signal are essentially absent from the energy region above 12 MeV, the flat contribution of the fast-neutron-dominated background can be estimated from the upper energy region, cross-checked below 20 MeV, and extrapolated downwards to the $\overline{\nu_e}$ region. The rates of fast-neutrons and remaining stopping μ's thus retrieved are presented in Table 6.2. In that case, the FDI and FDII rates are found to agree, thus, they are represented by a single value; the ND rate is higher, in accordance with the shallower site configuration.

Detector	$r_{\mu-n}$ (d^{-1})
FD	0.586 ± 0.061
ND	3.42 ± 0.23

Table 6.2 – Joint stopping μ and fast neutron daily rates for the Double Chooz detectors, after all the vetoes have been applied. FD stands for either FDI or FDII [115].

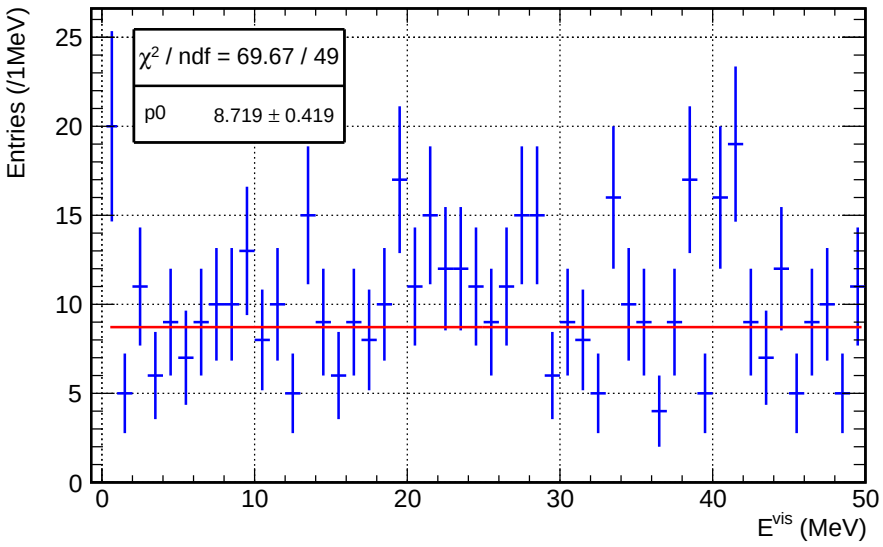

Figure 6.4 – Prompt spectrum of the fast neutron background at ND, extracted from coincidences with the IV. Adding a slope to the fit does not improve the agreement with the data and the slope from the fit is compatible with zero. Data from [119].

6.2.3 Cosmogenic background

The decay of cosmogenic isotopes is a topic covered in depth in III, and we will only give a brief summary here. In this document, what is referred to as a "cosmogenic isotope" is a nuclide created when a high-energy μ interacts with the nucleus of an atom, causing μ spallation. We hereby restrain the meaning of this expression to beta delayed-neutron emitters, which signifies ^{8}He or ^{9}Li for organic scintillators.

Aside from the possible formation of the short-lived ortho-positronium state (cf. 4.1.1), there is hardly any difference between e^- and e^+ deposits[7], and the tagging of the ortho-positronium has a large failure rate in the Double Chooz liquids. Furthermore, the cosmogenic nuclides usually do not only emit an e^- and a neutron: the β-decay is followed by a non-trivial instantaneous decay chain such that the prompt signal eventually extends up to $\sim 12\,\mathrm{MeV}$. From there, it is obvious that these βn-emitters can effortlessly mimic the $\overline{\nu_e}$ signal, the neutron capture acting as the delayed event. The tremendously long lifetimes of ^{8}He and ^{9}Li, 172 ms and 257 ms, respectively, make it impossible to veto every event following the detection of a μ, all the more so with the after-μ dead-times from 6.1.1.2 in mind.

If truth be told, there exists an active veto, based on the fact that candidate pairs close to μ-tracks which produced a plethora of particles are more likely to represent cosmogenic decays, but its vetoing efficiency lies between 25% and 50%. In spite of its limited efficiency, the cosmogenic veto is the only tool allowing to extract the energy shape of this background within the data.

[7] The annihilation γ's part of the e^+ deposit suffer from a larger quenching than the equivalent e^- energy, but the broad decay spectra overlap the IBD e^+ spectrum anyway.

Detector	r_{co} (d^{-1})
FD	$0.97^{+0.41}_{-0.16}$
ND	5.01 ± 1.43

Table 6.3 – Cosmogenic background rates for the Double Chooz detectors after all the vetoes have been applied. FD stands for either FDI or FDII [115].

Not only is the cosmogenic background the largest statistically (see Table 6.3), but it is also the one with the most complicated energy shape (anything but flat, see Figure 6.5), which happens to stand as the most difficult to retrieve (accidentals are mustered with ease, fast neutrons and stopping μ's produce triggers right after IV hits). This background corresponds to the largest rates, as obtained from a fit of the exponential decay of these nuclides, yet these rates are barely enough to provide a statistically accurate spectrum. All this craves for Monte-Carlo simulations, but generating these spectra is no pleasure cruise either, for the decay trees of these isotopes are both convoluted and seriously lacking in nuclear data inputs.

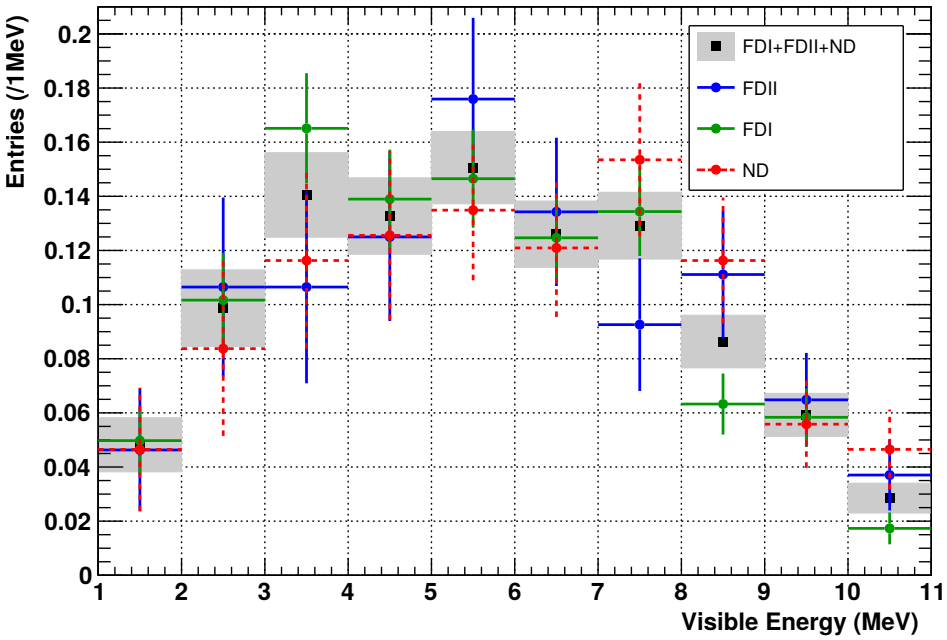

Figure 6.5 – Combined cosmogenic spectrum (FDI+FDII+ND, grey error bands) input of the oscillation fit for the Moriond conference. As confirmed by statistical analyses, all three spectra are not inconsistent with one another and they were consequently summed to increase the statistical significance. ^{8}He and ^{9}Li events cannot be disentangled.

6.3 Oscillation fit

6.3.1 Formulation

6.3.1.1 Strategy

The live-times for the Moriond data sets are relatively small, and the accuracy of the $\sin^2(2\theta_{13})$ measurement is dominated by the statistical uncertainty. Table 6.4 gathers the live-times for the three detectors. The daily rates of $\overline{\nu_e}$ candidates observed – after having applied all the vetoes, yet without subtracting the remaining background contaminations estimated in 6.2 – are presented as well. The lower ND live-time, when compared to FDII, is mostly due to the five times higher μ-rate (resulting in a larger dead-time), and the lower data taking efficiency for a few months after the ND commissioning. The differences between the FDI and FDII rates is accounted for in the prediction as well, it is rooted in a different operation of the power plant, i.e. more time with both reactors on for FDII.

Detector	t_{live} (d)	r_{cand} (d^{-1})
FDI	460.93	37.64
FDII	212.21	40.29
ND	150.76	293.4

Table 6.4 – Moriond live-times and $\overline{\nu_e}$ candidate rates observed in the Gd channel, after all the vetoes have been applied [115].

The scant number of detected events leaves little room for a direct comparison between the ND and FDII data, should one so much as hope to compete with other θ_{13} experiments. Not only do the former FDI data provide a larger set of events, but their comparison to a prediction strongly constrains the backgrounds to add to them, so that they reproduce the observed spectra. Assuredly, heavily relying on Monte-Carlo simulations of the reactor spectra, all the more so with the recent upheaval caused by the so-called 5 MeV bump, has its drawbacks. Fortunately, the bump region is far enough from the oscillation maximum and has little impact on θ_{13} [115]. To maximise the significance of the $\sin^2(2\theta_{13})$ measurement, FDII and ND are also compared to their respective predictions, yet with profitable correlations between the two predictions. Although they boast about 60% more events for the Gd analysis, the Neutrino data can still be qualified as a meagre set, thereby receiving the same treatment. The same goes for Gd++.

6.3.1.2 Least squares

Single parameter

The simultaneous comparison of all the data sets to their predictions is achieved via a method of least squares.

For each detector, to the number of events simulated in accordance with (5.17), and reconstructed in each energy bin ${N_{\overline{\nu}_e}^{MC}}^{(i)}$[8] with $i \in [\![1;40]\!]$, are added the remaining backgrounds ${N_{bg}}^{(i)} = {N_{acc}}^{(i)} + {N_{\mu-n}}^{(i)} + {N_{co}}^{(i)}$ to reproduce the observation ${N_{cand}}^{(i)}$. The simulated spectrum of a given detector may be multiplied by the $\sin^2(2\theta_{13})$-dependent oscillation probability, thereby defining its prediction $s \to N_{\overline{\nu}_e}^{pre}(s)$ with $s = \sin^2(2\theta_{13})$. The total number of background events is obtained from the shapes and rates r reviewed in 6.2, scaled by the live-times t_{live} found in 6.4. The $n_b = 40$ bins, common to all detectors and chosen to overcome the meagre statistics for the backgrounds above 8 MeV, are presented in Table 6.5.

Range (MeV)	n_{rb}	ΔE (MeV)
0.5-8	30	0.25
8-10	4	0.5
10-12	2	1
12-20	4	2

Table 6.5 – Visible energy binning for the Moriond oscillation fit using the Gd channel. The number of bins within each range n_{rb} and their widths ΔE are presented.

By varying a common $\sin^2(2\theta_{13})$ value, the fit concurrently minimises the residuals $(R^i_\delta)_{i\in[\![1;n_b]\!]}$ for each detector $\delta \in \mathcal{D}$ with $\mathcal{D} = \{FDI, FDII, ND\}$. The residuals read

$$R^i_\delta(s) = {N_{\overline{\nu}_e,\,\delta}^{pre}}^{(i)}(s) + {N_{bg,\,\delta}}^{(i)} - {N_{cand,\,\delta}}^{(i)}. \tag{6.13}$$

More precisely, the three vectors of n_b residuals R_δ with $\delta \in \mathcal{D}$, define the combined vector of residuals R, of size $3\,n_b$, whose transpose is presented below

$${}^{t}R(s) = \begin{pmatrix} R^1_{FDI} & \cdots & R^{n_b}_{FDI} & R^1_{FDII} & \cdots & R^{n_b}_{FDII} & R^1_{ND} & \cdots & R^{n_b}_{ND} \end{pmatrix}. \tag{6.14}$$

In the simplest case, the background measurements from 6.2 and the energy from 5.2.3 are taken for granted and their uncertainties find their way into a covariance matrix M of size $3\,n_b \times 3\,n_b$, along with the remaining uncertainties originating from the modelling of the reactor, the detector, the limited statistics of the observations. Within M, 3×3 blocks stand out, each of which representing the correlations between the different detectors $\delta \in \mathcal{D}$. Under all these conditions, the simple χ^2 to minimise with respect to $\sin^2(2\theta_{13})$ is

$$\chi^2(s) = {}^{t}R(s) M^{-1} R(s). \tag{6.15}$$

[8] In (5.17), $N_{\overline{\nu}_e}^{tot}$ refers to the number of counts predicted for a given run of length ΔT, on the other hand, $N_{\overline{\nu}_e}^{MC}$ corresponds to this number simulated for all runs and eventually summed.

Pull terms

In practice, there is a liking for letting some freedom to the 6.2 measurements. The background rates are indeed allowed to vary via so-called pull terms, and so can be the Δm^2_{31} value (cf. 2.3.1.2), along with some terms accounting for the Poisson statistics of the $\sim 7\,\mathrm{d}$ of FDI data with both reactors off. By the same token, the uncertainties on the QNL correction (cf. 5.2.3.4) and the LNL can be re-factored into three parameters [120], which can be varied for each fit. To put it differently, s turns into a vector of several parameters $\overrightarrow{s}$, on which R and χ^2 depend, and that is manually controlled by the addition of $n_p \in \mathbb{N}$ pull terms. The k-th pull, with $k \in [\![1; n_p]\!]$, can be written as

$$ {}^{\mathrm{t}}V_k\left(\overrightarrow{s}\right) {P_k}^{-1} V_k\left(\overrightarrow{s}\right) , \tag{6.16} $$

where V_k is the vector of residuals for the corresponding coefficients of $\overrightarrow{s}$, and P_k the covariance matrix answering for the uncertainties and correlations within the parts of $\overrightarrow{s}$ considered. The χ^2 thus becomes

$$ \chi^2\left(\overrightarrow{s}\right) = {}^{\mathrm{t}}R\left(s\right)M^{-1}R\left(\overrightarrow{s}\right) + \sum_{i=1}^{n_p} {}^{\mathrm{t}}V_k\left(\overrightarrow{s}\right) {P_k}^{-1} V_k\left(\overrightarrow{s}\right) . \tag{6.17} $$

6.3.2 Uncertainties and correlations

6.3.2.1 Backgrounds

Accidentals

Taking the accidental rate as the first pull term, and making explicit the detector indices $\delta \in \mathcal{D}$, its residual is defined as

$$ V_1\left(\overrightarrow{s}\right) = \begin{pmatrix} {V_1}_{\,FDI}\left(\overrightarrow{s}\right) \\ {V_1}_{\,FDII}\left(\overrightarrow{s}\right) \\ {V_1}_{\,ND}\left(\overrightarrow{s}\right) \end{pmatrix} = \begin{pmatrix} r^{acc}_{FDI} - s^{acc}_{FDI} \\ r^{acc}_{FDII} - s^{acc}_{FDII} \\ r^{acc}_{ND} - s^{acc}_{ND} \end{pmatrix} , \tag{6.18} $$

and all the r^{acc}_{δ} values are retrieved from Table 6.1 whereas the s^{acc}_{δ} parameters vary with $\overrightarrow{s}$ changes. The accidental rate variables, could be, for instance, the ones just after $\sin^2(2\theta_{13})$ in $\overrightarrow{s}$, that is $s^{acc}_{FDI} = s_2$, $s^{acc}_{FDII} = s_3$, $s^{acc}_{ND} = s_4$.

The covariance matrix for the pull of accidentals reads

$$ P_1 = \begin{pmatrix} {\sigma_{r^{acc}_{FDI}}}^2 & 0 & 0 \\ 0 & {\sigma_{r^{acc}_{FDII}}}^2 & 0 \\ 0 & 0 & {\sigma_{r^{acc}_{ND}}}^2 \end{pmatrix} , \tag{6.19} $$

which is a mere diagonal matrix. Indeed, the configuration and light noise changes between FDI and FDII trample on correlations. Likewise, the fact that ND and FDII are located at different sites crushes any correlation between their accidental rates. In other words, the pull for accidentals simplifies into a sum of three simple terms, i.e.

$$ {}^{\mathrm{t}}V_1\left(\overrightarrow{s}\right) {P_1}^{-1} V_1\left(\overrightarrow{s}\right) = \sum_{\delta\in\mathcal{D}} \left(\frac{r^{acc}_{\delta} - s^{acc}_{\delta}}{\sigma_{r_\delta}} \right)^2 . \tag{6.20} $$

Fast neutrons and stopping muons

The fast neutron and stopping μ rates are considered fully correlated between FDI and FDII. Indeed, separate a posteriori analyses yield similar values, and these backgrounds are not sensitive to the switching on of particularly flashing PMT's, nor are they particularly affected by upgrades of the electronics. On the other hand, the different overburdens at FDII and ND strictly forbid correlations between their varying rates. The residual for their pull term equals

$$V_2\left(\overrightarrow{s}\right) = \begin{pmatrix} r_{FDI}^{\mu-n} - s_{FDI}^{\mu-n} \\ r_{FDII}^{\mu-n} - s_{FDII}^{\mu-n} \\ r_{ND}^{\mu-n} - s_{ND}^{\mu-n} \end{pmatrix}, \tag{6.21}$$

and its non-diagonal matrix reads

$$P_2 = \begin{pmatrix} {\sigma_{r_{FDI}^{\mu-n}}}^2 & \sigma_{r_{FDI}^{\mu-n}}\sigma_{r_{FDII}^{\mu-n}} & 0 \\ \sigma_{r_{FDI}^{\mu-n}}\sigma_{r_{FDII}^{\mu-n}} & {\sigma_{r_{FDII}^{\mu-n}}}^2 & 0 \\ 0 & 0 & {\sigma_{r_{ND}^{\mu-n}}}^2 \end{pmatrix}. \tag{6.22}$$

Cosmogenic decays

Regarding the cosmogenic rates, their measurement is delicate and not tightly constraining anyway, they however ensure a posteriori that the fit results are not completely off, and vice versa. A single rate parameter is attributed to FDI and FDII, which can be regarded as a full correlation between their individual rates, but its value is not constrained by a pull term, and neither is the ND rate.

Insofar as the cosmogenic isotopes decay at rest, the prompt spectra they correspond to is independent of the process that spawned them, thereby further backing the use of the same energy distribution for all detectors[9]. The uncertainty $(\sigma_{co}^i)_{i\in[\![1;n_b]\!]}$ on the shape of their common data spectrum, labelled "FDI+FDII+ND" in Figure 6.5, defines a diagonal covariance matrix M_{co}^{spec}, of size $n_b \times n_b$, thus acknowledging that all n_b bins are independent from one another. Therefore, the cosmogenic component M_{co} of M is composed of blocks consisting of the diagonal M_{co}^{spec}, scaled by the varying cosmogenic rates $(s_\delta^{co})_{\delta\in\mathcal{D}}$ and fixed live-times $\left(t_\delta^{live}\right)_{\delta\in\mathcal{D}}$, namely

$$M_{co} = \begin{pmatrix} {s_{FDI}^{co}}^2\,{t_{FDI}^{live}}^2\,M_{co}^{spec} & s_{FDI}^{co}\,s_{FDII}^{co}\,t_{FDI}^{live}\,t_{FDII}^{live}\,M_{co}^{spec} & s_{ND}^{co}\,s_{FDI}^{co}\,t_{ND}^{live}\,t_{FDI}^{live}\,M_{co}^{spec} \\ s_{FDI}^{co}\,s_{FDII}^{co}\,t_{FDI}^{live}\,t_{FDII}^{live}\,M_{co}^{spec} & {s_{FDII}^{co}}^2\,{t_{FDII}^{live}}^2\,M_{co}^{spec} & s_{ND}^{co}\,s_{FDII}^{co}\,t_{ND}^{live}\,t_{FDII}^{live}\,M_{co}^{spec} \\ s_{ND}^{co}\,s_{FDI}^{co}\,t_{ND}^{live}\,t_{FDI}^{live}\,M_{co}^{spec} & s_{ND}^{co}\,s_{FDII}^{co}\,t_{ND}^{live}\,t_{FDII}^{live}\,M_{co}^{spec} & {s_{ND}^{co}}^2\,{t_{ND}^{live}}^2\,M_{co}^{spec} \end{pmatrix} \tag{6.23}$$

In (6.23), $s_{FDII}^{co} = s_{FDI}^{co}$ could be enforced. By virtue of the correlations between all the detectors, M_{co} is not diagonal itself, and yet, since M_{co}^{spec} is diagonal, a graphical representation of M_{co} would exhibit a pattern of inclined parallel stripes.

[9] To be fair, the share of ^{8}He and ^{9}Li could differ, but we are yet to reach this level of precision.

6.3.2.2 Reactor

Setting a value for the reactor correlation, which comes into play in the covariance matrix for the $N^{pre}_{\overline{\nu}_e,\,\delta}$'s in 6.13, is a thorny task. The uncertainties are numerous and with potentially different correlations across reactors $R \in \mathcal{R}$. The inaccuracy of the predictions are rooted in the limited knowledge of the fission fractions α^R_k, thermal powers P^R_{th}, baselines L_R, IBD cross-section σ_p, energies per fission E^m_f, reference spectra S_m, with $m \in \mathcal{F}$ (cf. 5.1.2.3). Some inputs are plainly independent from the reactor index, and accordingly correlated across them. As for the others, similar estimation methods might favour correlations, but nothing ought to be taken for granted.

During the single detector phase, a full correlation between both reactors worsens the sensitivity to $\sin^2(2\theta_{13})$. On the contrary, when FDII and ND share an identical period of data taking, the uncertainties on the predictions are all the more constrained by the observations that both reactors are correlated. Were FDII and ND in perfect iso-flux configurations and the two reactors, fully correlated, the flux uncertainty would vanish. For lack of a more daring treatment, an average value of $\rho_{B1,B2} = 0.78$ was selected as the most conservative correlation coefficient, namely the one degrading the most the sensitivity of the θ_{13} fit [120]. Nonetheless, the use of several detectors brings down the 1.7% relative uncertainty on the absolute normalisation of the reactor flux to 0.07% [115].

6.3.2.3 Detection

The validity of the prediction is also questioned by the number of target protons n^δ_p, and the detection efficiency ϵ_δ, in each detector $\delta \in \mathcal{D}$.

Target protons

In a given detector $\delta \in \mathcal{D}$, the number of target protons n^δ_p stems from the mass of H atoms in the target, itself evaluated as the product of the NT liquid mass $M^{liquid}_{NT,\,\delta}$ and the mass fraction of H f^δ_H. The retrieval of the NT masses is detailed in Chapter 11, the differing time scales and sensors employed makes these measurements largely independent from one another[10]. However, the near and far detector liquids are coming from the same production batch and there is only one $f_H = f^\delta_H$ for $\delta \in \mathcal{D}$. The 0.3% uncertainty on f_H prevails over the $< 0.1\%$ one of the mass measurements, and if the latter were flawlessly known, n^{ND}_p and $n^{FDI}_p \simeq n^{FDII}_p$ would be fully correlated. Albeit a little bold, the last hypothesis was retained by the oscillation fit team, i.e. $\rho^{n_p}_{\delta,\gamma} = 1$ for $(\delta, \gamma) \in \mathcal{D}^2$.

Spill effect

The second largest uncertainty in a single-detector analysis is the spill effect, which constrained our upper time coincidence limit in 6.1.2.1. The absence of Gd in the GC extends

[10]Evidently, neglecting the ever so slightly different average temperatures, FDI and FDII share the same mass of liquid in their fiducial volumes.

the life expectancy of neutrons in this volume, which results in a flow of events spilling into the NT larger than the one leaving it to avoid a capture on Gd. As measured from the simulations, the spill current is such that about 2% of the events in the Gd data analysis should not have been selected, but by definition, this spill current is already in the prediction, so there is no need to correct $N^{pre}_{\overline{\nu_e},\,\delta}$ a posteriori. Nevertheless, on account of the discrepancy between the neutron scattering models, $N^{pre}_{\overline{\nu_e},\,\delta}$ bears an additional 0.3% uncertainty [69]. Given that the near and far detector simulations are virtually identical, the spill current is fully correlated across all detector versions, that is $\rho^{spill}_{\delta,\gamma} = 1$ for $(\delta, \gamma) \in \mathcal{D}^2$. Considering that the liquids are identical, and that the acrylics were produced by the same company, with the same materials, the actual spill currents, which we can but hardly measure, should themselves also be identical.

Neutron detection

In light of the figures from 6.3.2.2 and since the aforementioned detection systematics cancel out in the oscillation fit, the Gd fraction contributes the most to the detection uncertainty, via ϵ_δ. The so-called Gd fraction accounts for the fact that about 14% of the $\overline{\nu_e}$'s interacting in the NT are missing from the Gd channel, since the outgoing neutron from the IBD is captured on H. In accordance with the 6.13 definition of the residuals, what matters is to gauge how the simulations and data differ, when it comes to correcting the predictions.

The Gd fraction f^δ_{Gd} is defined as the number of delayed events N^δ_d captured on Gd over that captured either on H or Gd

$$f^\delta_{Gd} = \frac{N^\delta_d\,(3.5\,\text{MeV} < E^{vis}_d < 10\,\text{MeV})}{N^\delta_d\,(1\,\text{MeV} < E^{vis}_d < 10\,\text{MeV})}. \tag{6.24}$$

Undoubtedly, the delayed energy bounds from (6.24) can be questioned, and this is precisely be varying them that the systematic uncertainties on f^δ_{Gd}, and its Monte-Carlo equivalent, can be assessed [120]. From there, a systematic uncertainty for each detector $\delta \in \mathcal{D}$, can be attributed to the ratio of the data and simulated f^δ_{Gd}. In the single detector case, the Gd fraction induces a 0.25% uncertainty on the prediction, out of these 0.25%, a part originates from the moderate statistics, which are obviously not correlated across the detectors. On the other hand, the systematic uncertainty is taken to be correlated [120], such that the correlation coefficients evaluate to

$$\rho^{frac}_{FDI,FDII} = 0.67 \tag{6.25}$$

$$\rho^{frac}_{FDII,ND} = 0.54\,. \tag{6.26}$$

Moreover, the selection cuts presented in 6.1.2 are an additional source of discrepancy between the observation and the prediction, which can be studied in a way similar to the Gd fraction, i.e. by opening the selections.

Global uncertainties

As a conclusion to the detection uncertainties, the detections systematics for all detectors are nearly divided by two on account of the correlation assumptions presented here. For instance, the detection uncertainty on the ND prediction is reduced from 0.38% to 0.15%, as Table 6.6 implies.

Detector	Single-detector (%)	Multi-detector (%)
FDI	0.49	0.26
FDII	0.47	0.22
ND	0.38	0.15

Table 6.6 – Relative detection uncertainties on the normalisation of the energy-integrated prediction – including the errors and correlations on the number of target protons, spill effect, neutron detection – in the case of a single-detector or multi-detector analysis [115].

In comparison, it is enlightening to look back at Tables 6.4 and 6.3: at the ND, which has the highest signal over background ratio, the cosmogenic background accounts for about 1.7% of the total counts, at FDII, for about 2.4%. The standard deviations on these background rates being 30% and 40%, respectively, this translates into normalisation errors on the predictions $N^{pre}_{\overline{\nu}_e,\,\delta}$ of about $\sim 0.5\%$ and $\sim 1\%$. To put it differently, with the reactor uncertainty crushed, and the detector one having shrunk, the Double Chooz experiment turns into a background experiment, and one may dare state into a cosmogenic background analysis.

6.3.2.4 Energy

Correcting parameters

As mentioned in 6.3.1.2, all the energy parameters of each detector $\delta \in \mathcal{D}$ can be summarised into a $(a_\delta, b_\delta, c_\delta)$ set, or vector. Via these three parameters, all the uncertainties on the corrections from 5.2.3 – directly applied to build E^{vis} – can show through, and both the LNL correction itself and its errors affect the $(a_\delta, b_\delta, c_\delta)$ prior. In other words, the magnitude of these parameters is determined by the inability of the partially tweaked simulations to reproduce the LNL observed in the data (cf. 5.2.3.4), while the uncertainties on them encompass all the corrections applied to E^{vis}, including the last-minute LNL.

Pull terms

For simplicity, and to our disadvantage, all three sets of energy parameters were considered independent from one detector another, although the use of identical liquids, for the LNL part of the energy, for instance, would vouch for profitable correlations. As a result, the energy pull term part of (6.17) is the sum of three terms, and it simplifies with respect to the detectors $\delta \in \mathcal{D}$ as for accidentals (cf. (6.20)). However, if each term of the (6.20) sum

for accidentals is a simple ratio of numbers, regarding the energy scale, correlations within these three parameters, on a detector-basis, must be implemented. Indeed, mixing all the corrective terms assuredly calls for correlations.

Taking the energy constraint on FDI as the third ${}^{\mathrm{t}}V_3\left(\overrightarrow{s}\right) {P_3}^{-1} V_3\left(\overrightarrow{s}\right)$ pull, we have the detector-dependent residuals

$$V_3\left(\overrightarrow{s}\right) = \begin{pmatrix} a_{FDI} - s^a_{FDI} \\ b_{FDI} - s^b_{FDI} \\ c_{FDI} - s^c_{FDI} \end{pmatrix} , \tag{6.27}$$

where s^a_{FDI}, s^b_{FDI}, s^c_{FDI} vary with $\overrightarrow{s}$ changes. The corresponding correlation matrix plainly reads

$$P_3 = \begin{pmatrix} {\sigma_{a_{FDI}}}^2 & \rho_{a_{FDI},b_{FDI}}\sigma_{a_{FDI}}\sigma_{b_{FDI}} & \rho_{a_{FDI},c_{FDI}}\sigma_{a_{FDI}}\sigma_{c_{FDI}} \\ \rho_{a_{FDI},b_{FDI}}\sigma_{a_{FDI}}\sigma_{b_{FDI}} & {\sigma_{b_{FDI}}}^2 & \rho_{b_{FDI},c_{FDI}}\sigma_{b_{FDI}}\sigma_{c_{FDI}} \\ \rho_{a_{FDI},c_{FDI}}\sigma_{a_{FDI}}\sigma_{c_{FDI}} & \rho_{b_{FDI},c_{FDI}}\sigma_{b_{FDI}}\sigma_{c_{FDI}} & {\sigma_{c_{FDI}}}^2 \end{pmatrix} , \tag{6.28}$$

where the correlation coefficients ρ can be extracted from [121].

The fourth and fifth pulls would have a straightforward definition as well. Please note that if for each background, there is only one pull potentially correlated across detectors, as far as the energy goes, there is one pull per detector, independent from the others, albeit not trivial itself.

Implementation

To be precise, changes in the energy scale of a prediction are applied on an event-basis, minute changes in the varying s^a_{FDI}, s^b_{FDI}, s^c_{FDI} may leave the prediction spectrum unscathed, slightly less delicate shifts can move some events from one bin to the other. As a consequence, continuous changes in $\overrightarrow{s}$ do not translate into smooth spectral changes. For this reason, discrete grids of these 3×3 variables are generated, and a χ^2 minimisation is performed for each and every point of the staggering grid, hence the ordeal.

6.3.3 Fit results and prospects

6.3.3.1 Best fit

From the minimisation of $\overrightarrow{s} \to \chi^2\left(\overrightarrow{s}\right)$ as defined in (6.17), we find

$$\sin^2(2\theta_{13}) = 0.111 \pm 0.018 \,, \tag{6.29}$$

i.e. the first component of the vector $\overrightarrow{s_{min}}$ yielding the lowest χ^2 value χ^2_{min}. The χ^2 minimum is represented by

$$\frac{\chi^2_{min}}{ndf} = \frac{128.8}{120} \,, \tag{6.30}$$

where *ndf* is the number of degrees of freedom in the simultaneous fit [115]. Consequently, with a multi-detector set up, the Double Chooz experiment observes a $\overline{\nu_e}$ disappearance with a $\sim 6\sigma$ significance.

Along with $\sin^2(2\theta_{13})$ we get the cosmogenic rates r^{co}_{ND} and $r^{co}_{FDI} = r^{co}_{FDII}$ according to the fit. Although the pull terms give them little freedom, the joint rates of fast neutrons and stopping μ's may be retrieved from $\overrightarrow{s_{min}}$ as well. All these values are displayed in Table 6.7.

Detector	r^{fit}_{co} (d^{-1})	$r^{fit}_{\mu-n}$ (d^{-1})
FD	0.75 ± 0.14	0.535 ± 0.035
ND	4.89 ± 0.78	3.53 ± 0.16

Table 6.7 – Cosmogenic background rates r^{fit}_{co} obtained from the simultaneous fit all the data sets. The joint fast neutrons and stopping μ's rates $r^{fit}_{\mu-n}$ are also presented [115].

The relatively large error bar on the r^{co}_{ND} measurement (cf : Table 6.3) guarantees a flawless agreement with the best fit output. The FD measurement is a little less forgiving and a $\sim 1.4\sigma$ difference is observed between the prior and the fit, albeit not yet troubling.

6.3.3.2 Comparison to the $\overline{\nu_e}$ survival probability

Although $\overrightarrow{s_{min}}$ is obtained by minimising the residuals from 6.13, i.e. the difference between the background-augmented prediction and the observation, plotting the ratios of $N_{cand,\,\delta}$ over $N^{pre}_{\overline{\nu_e},\,\delta}(\overrightarrow{s_{no}}) + N_{bg,\,\delta}$ allows for a direct comparison with the oscillation probability at fixed distances. Here, $\overrightarrow{s_{no}}$ corresponds to $\overrightarrow{s_{min}}$ aside from the fact that it represents the no-oscillation hypothesis, i.e. $s_1 = \sin^2(2\theta_{13}) = 0$. The $N_{cand,\,\delta}$ to $N^{pre}_{\overline{\nu_e},\,\delta}(\overrightarrow{s_{no}}) + N_{bg,\,\delta}$ ratios, along with the $\overline{\nu_e}$ survival probability (2.72) using $\sin^2(2\theta_{13})$ from (6.29), at the corresponding weighted distances from the reactor cores, are plotted in Figures 6.6, 6.7, and 6.8.

Expectedly, at an average distance of $\sim 1\,\mathrm{km}$, the ratio of the observation to the prediction for FDI and FDII is quite compatible with the probability we drew in Figure 2.5. The larger error bars at FDII are in accordance with its halved statistics, in comparison to FDI. The ND ratio, with many more $\overline{\nu_e}$'s, looks nearly flat, in agreement with its $\sim 400\,\mathrm{m}$ average distance to the reactor cores.

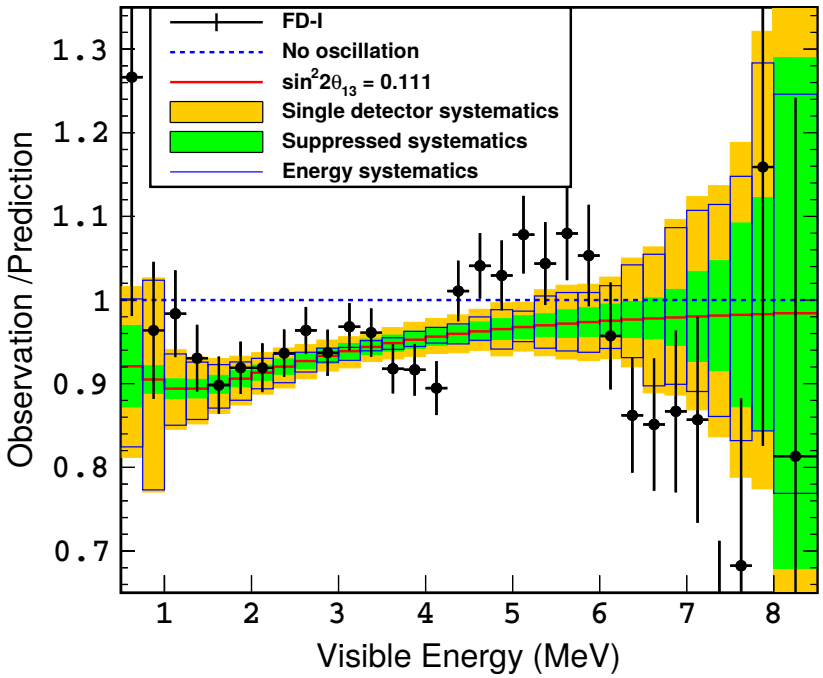

Figure 6.6 – Ratio of the observed to expected spectra for FDI (black points), along with the survival probability (red line) for the $\sin^2(2\theta_{13})$ value obtained from the unique best fit of all detector sets. To let the disappearance show through, the expectation is non-oscillated.

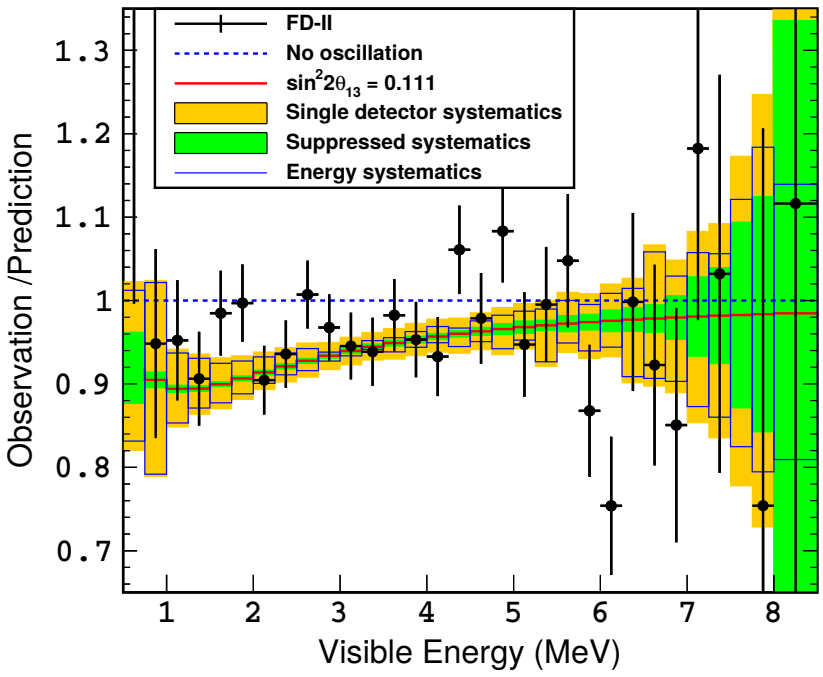

Figure 6.7 – Ratio of the observed to expected spectra for FDII (black points), along with the survival probability (red line) for the $\sin^2(2\theta_{13})$ value obtained from the unique best fit of all detector sets. To let the disappearance show through, the expectation is non-oscillated.

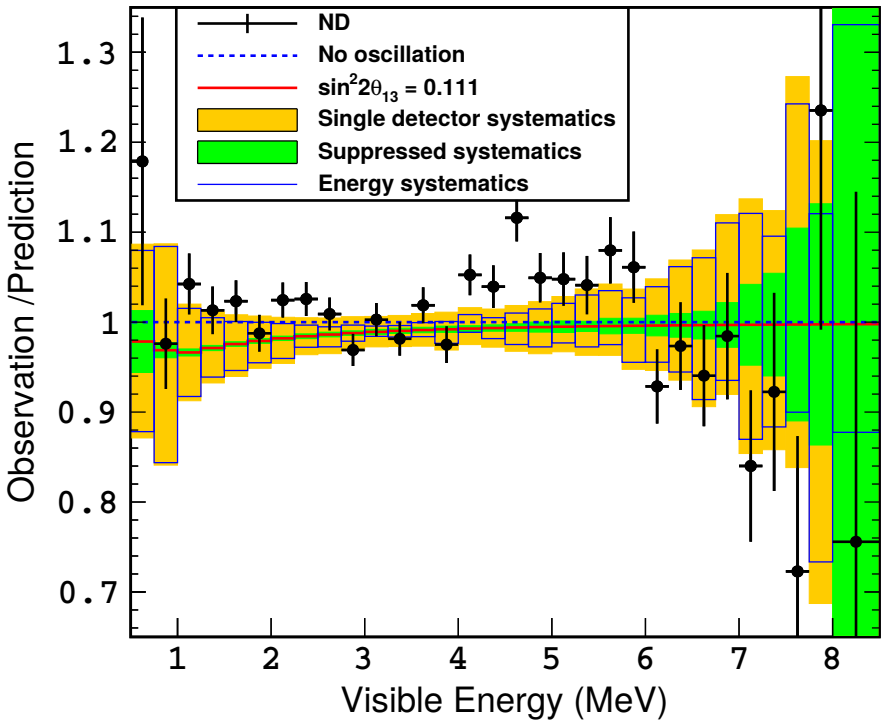

Figure 6.8 – Ratio of the observed to expected spectra for ND (black points), along with the survival probability (red line) for the $\sin^2(2\theta_{13})$ value obtained from the unique best fit of all detector sets. To let the disappearance show through, the expectation is non-oscillated.

Particularly on the FDI and ND ratios, which are not utterly smeared out by the lack of statistics, an excess of events in the candidates – not predicted by the simulations – starts to show in the $4-6$ MeV range. Various studies were carried out [77], although none has been conclusive enough to topple the others. The main convicted are the reactor flux predictions, and the energy scale. In any event, the impact of this bias on the extraction of θ_{13} is deemed negligible, on account of the location of the oscillation minimum. With more statistics, data to data comparison should be beneficial, and strongly suppress the effect of this distortion, as it has been the case for years, in the Daya Bay analyses. The simple Double Chooz layout should also help the near detector to characterise this distortion, as should shorter baseline experiments aimed at finding sterile ν oscillations near nuclear reactors.

6.3.3.3 Double Chooz amongst others

The relatively high $\sin^2(2\theta_{13}) = 0.111 \pm 0.018$ from the Double Chooz fit deserves to be compared to measurements from other reactor and accelerator experiments. As stated in 2.3.2.2, the latter are sensitive to the mass hierarchy when studying ν_e appearance in a ν_μ beam, hence the presence of two values in Figure 6.9.

As can be inferred from the same diagram, the Double Chooz value is more than 1σ larger than the $\sin^2(2\theta_{13}) = 0.084 \pm 0.005$ value from Day Bay. Insofar as the uncertainty

on the Double Chooz result is amply driven by statistics, this 1.5σ disagreement should not cause an outcry in the θ_{13} world. Nonetheless, utterly independent measurements are a key to claiming discoveries and providing the rest of the world with long-standing numerical values.

By all manner of means, even with the Gd++ analysis putting forth $\sim$ 26 t fiducial volumes, the odds are strong that the θ_{13} legacy value will be guided by the Daya Bay measurements. And yet, presenting competitive cross-checks and sharing experiences assuredly helps hunting down systematics biases in the analyses of all teams, may they be in the energy reconstruction, the reactor predictions, or the detection efficiency.

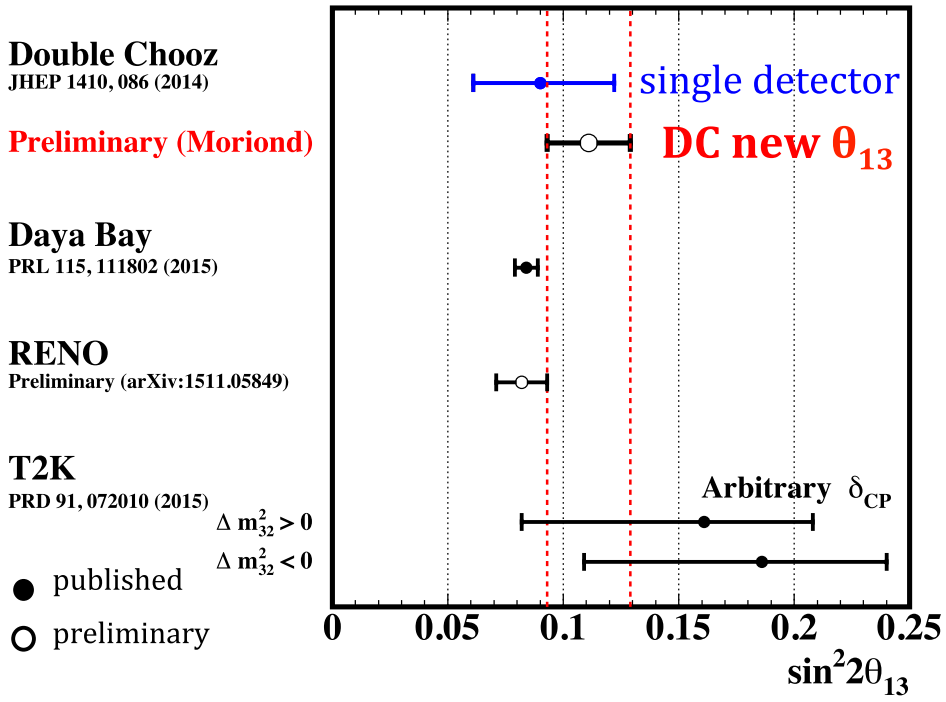

Figure 6.9 – Main $\sin^2(2\theta_{13})$ measurements and their associated error bars, as of March 2016.

Part III

Cosmogenic background studies

Chapter 7

Generating cosmogenic decays

In the light of the previous chapter, the reader should be convinced of the prime importance of the cosmogenic background, brought about by the decay of βn-emitters in the detector itself. If there indeed exists an active cosmogenic veto, comparison with a spectral prediction allows to validate the energy distribution of the relatively limited sample of vetoed events. Indeed, data may be data, but exactly like Monte-Carlo simulations, it does not prevent them from representing a process which is not quite that which the analysers had in mind. Misuse of data spectra – in the oscillation fit extracting θ_{13} – is no lighter blunder than relying on skewed predictions. Somehow, the both of them, rather than compete, are complementary. In addition, the vetoed events are inevitably a mixture of ^{8}He and ^{9}Li decays; we must rely on external knowledge, provided by predictions, to extract their respective contributions. Last but not least, confidence in the Monte-Carlo predictions allows to make the most of the difference between the IBD and cosmogenic endpoints, namely deduce the efficiency of an energy-based cut on the prompt energy.

The cosmogenic nuclides usually do not only emit an e^{-} and a neutron, and the β-decay is followed by a non-trivial decay chain possibly releasing ^{3}H's, ^{4}He's or even larger nuclides, all of which may contribute to the prompt signal. Before we contemplate generating spectra, the first task comes down to being able to process different decay types – β, many-body – an endeavour led in 7.1, and to propagate them consistently through a given decay path, an issue tackled in 7.2. Handling the complete structure of a decay tree, and thereby the selection of decay chains to process, is a thorny topic left for Chapter 8; hence the limited scope suggested by the title of the current chapter.

7.1 Generating each raw decay

In the decay from one state to the next, the computation and propagation of the relevant physical quantities is achieved by means of a `C++11` program, aptly baptised "Spectrum Generator", entirely written by the author. Within this chapter, we will restrain ourselves to the part of the Spectrum Generator aimed at handling a specified chain, although it can

do much and more, thus justifying its plain name. The "raw" qualifier for "decay" means to stress that the four-momenta of the particles have not been run through the detector simulation.

Thereafter, we employ the so-called natural units.

7.1.1 Handling strong decays

In this document, a "strong decay" refers to the decay of an unbound nucleus. As an example, any level in the ^{5}He nucleus falls into this category.

Considering the acute lack of nuclear data for the unbound nuclei involved in the cosmogenic decay chains, only kinematic effects have been implemented. Any possible dynamic effect due to spin or parity has been overlooked. This is also justified by the high statistics of the Monte-Carlo, which smear out any angular dependence on an event by event basis.

However, to take into account the recoil of light nuclei and provide a suitable modelling of electrons, complete relativistic kinematics have been worked out. Considering how close the mass of a ^{5}He nucleus is to that of an α particle, or even to that of a neutron, one understands that neglecting the recoil energy of nuclei may lead to a substantial amount of missing energy (of the order a few MeV's for the ^{9}Li or ^{8}He decay chains). Therefore, after the break up of ^{9}Be into (^{5}He, ^{4}He), the subsequent decay of ^{5}He has to be performed in flight. Accordingly, the decay spectrum of ^{5}He does not consist any longer of two distinct peaks determined by the sole α and neutron masses. Staying within the relativistic kinematics frame ensures that the conservation of energy, and three-momentum, is well taken care of.

7.1.1.1 Two-body decays

Whether because the cosmogenic decay chains involve many two-body decays, or because each many-body decay can be split into several two-body decays, the simple kinematics of the two-body decay form the core of the Spectrum Generator.

Centre of momentum frame

In the centre of momentum frame, the norms of the momenta of the outgoing particles, $|\overrightarrow{p_2}|$ and $|\overrightarrow{p_3}|$, in a $1 \to 2 + 3$ decay, are well determined by

$$|\overrightarrow{p_2}| = |\overrightarrow{p_3}| = \frac{\sqrt{\lambda({m_1}^2, {m_2}^2, {m_3}^2)}}{2m_1}. \tag{7.1}$$

λ is known as the triangle function and its expression reads $\lambda(x, y, z) = x^2 + y^2 + z^2 - 2xy - 2xz - 2yz$.

Since spin effects can hardly be considered, the direction of $\overrightarrow{p_2}$ is randomly generated using spherical coordinates, i.e. by picking $\cos(\theta)$ uniformly in $[-1; 1]$, and ϕ in $[0; 2\pi]$, with θ the polar angle, and ϕ the azimuthal angle. The direction of the other outgoing

three-momentum $\overrightarrow{p_3}$ is then $\overrightarrow{p_3} = -\overrightarrow{p_2}$. The energies follow as $E_i = \sqrt{|\overrightarrow{p_i}|^2 + {m_i}^2}$, with $2 \leqslant i \leqslant 3$.

Lorentz boost

If need be, a Lorentz boost of direction given by $\overrightarrow{p} = -\overrightarrow{p_1}$ is performed over the four-momenta P_2 and P_3 to go back to the lab frame using $\gamma = E_1/m_1$. Only the components along the direction of the boost are affected, hence the change of frame

$$\begin{cases} E' = \gamma(E - \overrightarrow{\beta}.\overrightarrow{p}) \\ \overrightarrow{p}' = \overrightarrow{p_\perp} + \gamma(\overrightarrow{p_\parallel} - E\overrightarrow{\beta}) \\ \quad\quad = (\gamma - 1)\dfrac{\overrightarrow{p}.\overrightarrow{\beta}}{\beta}\dfrac{\overrightarrow{\beta}}{\beta} + \overrightarrow{p} - \gamma E\overrightarrow{\beta} \end{cases}, \tag{7.2}$$

with $\overrightarrow{p_\parallel}$ and $\overrightarrow{p_\perp}$ the parallel and orthogonal components of the decomposition $\overrightarrow{p} = \overrightarrow{p_\parallel} + \overrightarrow{p_\perp}$. More practically, 7.2 defines the matrix

$$\Lambda = \begin{pmatrix} \gamma & -\gamma\beta_x & -\gamma\beta_y & -\gamma\beta_z \\ -\gamma\beta_x & 1 + \frac{(\gamma-1){\beta_x}^2}{\beta^2} & \frac{(\gamma-1)\beta_y\beta_x}{\beta^2} & \frac{(\gamma-1)\beta_z\beta_x}{\beta^2} \\ -\gamma\beta_y & \frac{(\gamma-1)\beta_x\beta_y}{\beta^2} & 1 + \frac{(\gamma-1){\beta_y}^2}{\beta^2} & \frac{(\gamma-1)\beta_z\beta_y}{\beta^2} \\ -\gamma\beta_z & \frac{(\gamma-1)\beta_x\beta_z}{\beta^2} & \frac{(\gamma-1)\beta_y\beta_z}{\beta^2} & 1 + \frac{(\gamma-1){\beta_z}^2}{\beta^2} \end{pmatrix}, \tag{7.3}$$

to be applied to P_2 and P_3 with $\beta = \sqrt{1 - 1/\gamma^2}$ and $\beta_i = \beta\ p_i/\left|\overrightarrow{p}\right|$.

7.1.1.2 Many-body decays

Model

As regards many-body decays, a good part of the work has already been accomplished once the two-body decays have been coded. We can indeed make the most of the phase space recurrence formula that allows one to decompose a $n-$body ($n > 2$) decay $m \to \sum_{i=1}^{n} m_i$ into $n-1$ two-body disintegrations. In this particular decomposition, the particle of mass m decays into m_n and an invariant mass M_{n-1} representing all the other particles. Then, M_{n-1} is decayed into m_{n-1} and the next invariant mass M_{n-2} and so on until $M_1 = m_1$. In the following, P_j will denote the four-momentum of the particle of mass m_j with $j \in [\![1; n]\!]$. These four-momenta P_j are obviously the quantities to be retrieved from each one of the two-body disintegrations. An example of such a splitting is given in Figure 7.1 for $n = 4$.

Recurrence formula

How such a decomposition into two-body decays can be valid? How can we choose the values of the invariant masses $M_l = \sqrt{\left(\sum_{i=1}^{l} P_i\right)^2}$ with $l \in [\![1; n]\!]$ in a $n-$body decay, if we are precisely trying to predict the outgoing four-momenta? The answers lie in the recurrence formula for the phase space.

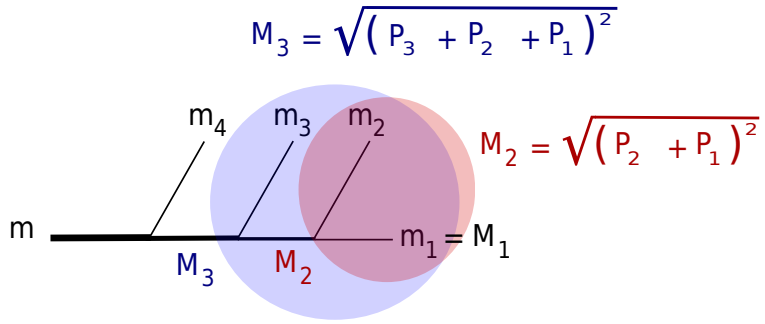

Figure 7.1 – Splitting of a four-body decay into three two-body decays. Firstly m decays into m_4 and M_3, where M_3 is the invariant mass of the system $\{m_3, m_2, m_1\}$. Secondly, M_3 decays into m_3 and M_2, with M_2 being the invariant mass of $\{m_2, m_1\}$. Lastly, M_2 decays into m_2 and $M_1 = m_1$.

Here comes into play a bit of maths, which the uninterested reader may skip; the thoroughgoing description[1] however serves understanding. In a $n-$body decay $m \to \sum_{i=1}^{n} m_i$, the accessible phase space R_n is defined as

$$R_n\left(P; m_{1\leq i\leq n}\right) = \int_{\mathbb{R}^{4n}} (2\pi)^4 \delta^4\left(P - \sum_{i=1}^{n} P_i\right) \prod_{1\leq j\leq n} 2\pi\delta\left(P_j^2 - m_j^2\right)\theta\left(P_j^0\right)\frac{\mathrm{d}^4 P_j}{(2\pi)^4} \tag{7.4}$$

with θ the Heavyside distribution and P the four-momentum of the particle of mass $m = M_n$ (the invariant masses have been extended up to $M_n = \sqrt{\left(\sum_{i=1}^{n} P_i\right)^2} = \sqrt{P^2}$ to simplify the notations). We wish to split it into two simpler parts the way n can be split into two other integers l and $n-l$. Thus, let us introduce a four-momentum we call Q_l over which we perform an additional integration[2]

$$\begin{aligned} R_n\left(P; m_{1\leq i\leq n}\right) = &\int_{\mathbb{R}^{4(n+1)}} (2\pi)^4 \delta^4\left(Q_l - \sum_{i=1}^{l} P_i\right) \prod_{1\leq j\leq l} 2\pi\delta\left(P_j^2 - m_j^2\right)\theta\left(P_j^0\right)\frac{\mathrm{d}^4 P_j}{(2\pi)^4} \\ &(2\pi)^4 \delta^4\left(P - Q_l - \sum_{r=l+1}^{n} P_r\right) \prod_{l+1\leq k\leq n} 2\pi\delta\left(P_k^2 - m_k^2\right)\theta\left(P_k^0\right)\theta\left(Q_l^0\right)\frac{\mathrm{d}^4 P_k}{(2\pi)^4}\frac{\mathrm{d}^4 Q_l}{(2\pi)^4} \end{aligned} . \tag{7.5}$$

The first line is identified with ease as $R_l\left(Q_l; m_{1\leq i\leq l}\right)$ after the first $4l$ integrations have been performed

$$\begin{aligned} R_n\left(P; m_{1\leq i\leq n}\right) = &\int_{\mathbb{R}^{4(n+1-l)}} R_l\left(Q_l; m_{1\leq i\leq l}\right) \\ &(2\pi)^4 \delta^4\left(P - Q_l - \sum_{r=l+1}^{n} P_r\right) \prod_{l+1\leq k\leq n} 2\pi\delta\left(P_k^2 - m_k^2\right)\theta\left(P_k^0\right)\theta\left(Q_l^0\right)\frac{\mathrm{d}^4 P_k}{(2\pi)^4}\frac{\mathrm{d}^4 Q_l}{(2\pi)^4} \end{aligned} . \tag{7.6}$$

[1] Only a sketch of the beginning of the demonstration can be found in [122], or in the oft-quoted reference [123].

[2] The product by $\theta\left(Q_l^0\right)$ can be added as the energy component of $\sum_{i=1}^{l} P_i$ is doubtless positive.

The second line is not so clear yet, it lacks a mass-shell term $2\pi\delta\left(Q_l^2 - M_l^2\right)$ for the newly introduced four-momentum Q_l. Such a term can be added as long we integrate over M_l^2 between ${M_l^2}_{min} = \min\left(\sum_{i=1}^{l} P_i\right)^2 = \left(\sum_{i=1}^{l} m_i\right)^2$ and ${M_l^2}_{max} = \max\left(\sum_{i=1}^{l} P_i\right)^2 = \left(m - \sum_{i=l+1}^{n} m_i\right)^2$, thus

$$\begin{aligned} R_n\left(P; m_{1\leq i\leq n}\right) = & \int_{{M_l^2}_{min}}^{{M_l^2}_{max}} \int_{\mathbb{R}^{4(n+1-l)}} R_l\left(Q_l; m_{1\leq i\leq l}\right) (2\pi)^4 \delta^4\left(P - Q_l - \sum_{r=l+1}^{n} P_r\right) \\ & \prod_{l+1\leq k\leq n} 2\pi\delta\left(P_k^2 - m_k^2\right) \frac{2\pi}{2\pi}\delta\left(Q_l^2 - M_l^2\right)\theta\left(P_k^0\right)\theta\left(Q_l^0\right) \frac{\mathrm{d}^4 P_k}{(2\pi)^4}\frac{\mathrm{d}^4 Q_l}{(2\pi)^4}\,\mathrm{d}M_l^2\,. \end{aligned} \tag{7.7}$$

With the mass-shell term, $R_l\left(Q_l; m_{1\leq i\leq l}\right)$ eventually appears as the decay of a fictitious particle of mass M_l while the other terms describe the decay of a particle of mass m into the fictitious particle of mass M_l and $n-l$ particles of masses $m_{l+1\leq k\leq n}$. Nevertheless, great care must be taken to perform the integration over Q_l since it appears in the first phase space term R_l. The latter is a Lorentz invariant, so it can be evaluated in the frame where $Q_l = \left(M_l, \overrightarrow{0}\right)$. Therefore, we do abide by the law when we write

$$R_n\left(P; m_{1\leq i\leq n}\right) = \frac{1}{2\pi}\int_{{M_l^2}_{min}}^{{M_l^2}_{max}} R_l\left(M_l; m_{1\leq i\leq l}\right) R_{n-l+1}\left(P; M_l, m_{l+1\leq k\leq n}\right)\mathrm{d}M_l^2\,, \tag{7.8}$$

i.e. the $n-$body phase space decomposes into l-body and $n-l+1$-body phase spaces, thereby demonstrating a useful form of the recurrence formula.

Complete decomposition into two-body decays

Let us illustrate and justify the picture we gave in Figure 7.1. By picking $l = 3$ we can firstly decompose a four-body decay into a three-body decay and a two-body decay

$$R_4\left(P; m_{1\leq i\leq 4}\right) = \frac{1}{2\pi}\int_{M_{3\,min}^2}^{M_{3\,max}^2} R_3\left(M_3; m_{1\leq i\leq 3}\right) R_2\left(P; M_3, m_4\right)\mathrm{d}M_3^2\,, \tag{7.9}$$

which is the step corresponding to the first branching in Figure 7.1. Now let us proceed to the next decay of the scheme, i.e. let us decay M_3 into m_3 and the next invariant mass M_2. To do so, we need but use $n = 3$ and $l = 2$ in (7.8) and insert the result in (7.9), leading to the two-body decomposition

$$\begin{aligned} R_4\left(P; m_{1\leq i\leq 4}\right) = & \frac{1}{(2\pi)^2}\int_{M_{3\,min}^2}^{M_{3\,max}^2}\int_{M_{2\,min}^2}^{M_{2\,max}^2} R_2\left(M_2; m_{1\leq i\leq 2}\right) \\ & R_2\left(M_3; M_2, m_3\right) R_2\left(P; M_3, m_4\right)\mathrm{d}M_2^2\,\mathrm{d}M_3^2\,. \end{aligned} \tag{7.10}$$

A simple mathematical induction starting from (7.8) yields the two-body decomposition

for the general case ($n > 2$)

$$R_n\left(P; m_{1\leq i\leq n}\right) = \frac{1}{(2\pi)^{n-2}} \int_{M^2_{n-1\,min}}^{M^2_{n-1\,max}} \int_{M^2_{n-2\,min}}^{M^2_{n-2\,max}} \dots \int_{M^2_{2\,min}}^{M^2_{2\,max}} R_2\left(M_2; M_1 = m_1, m_2\right) \dots R_2\left(M_{n-1}; M_{n-2}, m_{n-1}\right) R_2\left(P; M_{n-1}, m_n\right) \mathrm{d}M_2^2 \dots \mathrm{d}M_{n-2}^2 \, \mathrm{d}M_{n-1}^2 \,. \tag{7.11}$$

When evaluating the $n-$body phase space in the rest frame of the decaying particle[3] where $P = \left(M_n, \overrightarrow{0}\right)$, we obtain the more neatly written formula

$$R_n\left(M_n; m_{1\leq i\leq n}\right) = \frac{1}{(2\pi)^{n-2}} \int_{M^2_{n-1\,min}}^{M^2_{n-1\,max}} \int_{M^2_{n-2\,min}}^{M^2_{n-2\,max}} \dots \int_{M^2_{2\,min}}^{M^2_{2\,max}} \prod_{2\leq i\leq n} R_2\left(M_i; M_{i-1}, m_i\right) \mathrm{d}M_2^2 \dots \mathrm{d}M_{n-2}^2 \, \mathrm{d}M_{n-1}^2 \,. \tag{7.12}$$

which indeed describes $n-1$ two-body decays.

Invariant mass distribution

What arises from the two-body decomposition formula (7.12) is actually the distribution for the set of squared masses M_l^2 with $l \in [\![2; n-1]\!]$. Indeed, $R_n\left(M_n; m_{1\leq i\leq n}\right)$ is constant with respect to the integration variables and we can write

$$1 = \int_{M^2_{n-1\,min}}^{M^2_{n-1\,max}} \int_{M^2_{n-2\,min}}^{M^2_{n-2\,max}} \dots \int_{M^2_{2\,min}}^{M^2_{2\,max}} \frac{1}{(2\pi)^{n-2} R_n\left(M_n; m_{1\leq i\leq n}\right)} \prod_{2\leq i\leq n} R_2\left(M_i; M_{i-1}, m_i\right) \mathrm{d}M_2^2 \dots \mathrm{d}M_{n-2}^2 \, \mathrm{d}M_{n-1}^2 \,. \tag{7.13}$$

Hence the invariant mass distribution f of the set of squared masses M_l^2 with $l \in [\![2; n-1]\!]$ is

$$f(M_2^2, \dots, M_{n-1}^2) = \frac{1}{(2\pi)^{n-2} R_n\left(M_n; m_{1\leq i\leq n}\right)} \prod_{2\leq i\leq n} R_2\left(M_i; M_{i-1}, m_i\right) . \tag{7.14}$$

Invariant mass generation

One may wonder what the point is since the unknown $n-$body phase space R_n still appears in (7.14). However, the latter does not depend on the invariant masses $\overrightarrow{M} = (M_2, \dots, M_{n-1})$ of the intermediate two-body disintegrations, along with the $(2\pi)^{n-2}$, it reduces to a normalisation constant. Thus, for a Monte-Carlo method aimed at obtaining $\overrightarrow{M}$ to perform a

[3] This how the Spectrum Generator works and a final Lorentz boost can always be applied.

many-body decay, we can simply overlook this constant with a rejection sampling method. We can indeed generate $\left(\overrightarrow{M}, y\right)$ in $(\mathbb{R}^+)^{n-1}$ and retain $\overrightarrow{M}$ whenever $y < g\left(\overrightarrow{M}\right)$ with g defined by[4]

$$g\left(\overrightarrow{M}\right) = (2\pi)^{n-2} R_n\left(M_n; m_{1\leq i\leq n}\right) f\left(\overrightarrow{M}\right). \tag{7.15}$$

We may be a little cleverer by generating $\left(\overrightarrow{M}, y\right)$ in $S \times \mathbb{R}^+$ where S is the sensible space for $\overrightarrow{M}$, namely $S = \prod_{2\leq i\leq n-1}\left[{M_i}^{min}, {M_i}^{max}\right]$ with the bounds as in (7.7). For completeness there remains to give the expression of the well-known (see also [124]) two-body phase space

$$R_2\left(m_1; m_2, m_3\right) = \frac{\sqrt{\lambda({m_1}^2, {m_2}^2, {m_3}^2)}}{8\pi {m_1}^2} \tag{7.16}$$

with λ as defined in 7.1.1.1.

Upper bound for the invariant mass distribution

Notwithstanding the elegance of the method, numerically generating points in $S \times \mathbb{R}^+$ can prove extremely inefficient. We need a bound on the $\mathbb{R}^+$ axis. This bound is found by noticing that the total phase space is largest when the available energy is largest in each two-body phase space. Thus,

$$g\left(\overrightarrow{M}\right) < \prod_{2\leq i\leq n} R_2\left(M_i^{max}; M_{i-1}^{min}, m_i\right). \tag{7.17}$$

This bound allows for fast generation of invariant mass sets by generating points in $S \times [0; y_{max}]$ with $y_{max} = \prod_{2\leq i\leq n} R_2\left(M_i^{max}; M_{i-1}^{min}, m_i\right)$. Please note that y_{max} is not associated with a possible decay since the energy available in each one of the two-body phase spaces would already be $\Delta_i = M_i^{max} - M_{i-1}^{min} - m_i = M_n - \sum_{k=1}^{n} m_k = \Delta$. In other words, the total energy available would be $(n-1)\Delta > \Delta$ for $n > 2$.

7.1.2 Processing β - decays

When striving to detect $\overline{\nu_e}$'s in a liquid scintillator, nuclei produced by cosmic μ's are troublesome as soon as their decay chains involve the emission of an electron, whose properties are difficult to distinguish from that of a positron. The electron emission is induced by the β-decay of the cosmogenic isotope itself.

7.1.2.1 Obtaining the kinetic energy of an electron

Fortunately, the code "BESTIOLE"[5], already written by my supervisor D. Lhuillier and T. Mueller, returns the kinetic energy distributions of electrons and $\overline{\nu_e}$'s in the $\beta-$decay of

[4] We hereby gracelessly re-parametrise f so that it depends on $\overrightarrow{M}$ directly.

[5] Beta Energy Spectrum Tool for an Improved Optimal List of Elements.

a given nucleus. For each decay, there remains to pick a random kinetic energy T_e in the electron spectrum. The predicted kinetic energies are accurate and include Weak Magnetism. Fermi, radiative and finite-size corrections are also implemented. More details can be found in [92]. The norms of the momenta are easily obtained as $|\overrightarrow{p_e}| = \sqrt{{T_e}^2 + 2m_e T_e}$ and a random direction of motion is generated to provide a full description.

7.1.2.2 Alleged classification of the β-decays

For all its accuracy, one must bear in mind that BESTIOLE has the most complete output when dealing with allowed transitions. As pointed out in 7.1.1, the spin and parities of the levels fed by the cosmogenic decays are uncertain. Nonetheless, we may put to question the damage done, basing our judgement on the most probable J^π values, where J represents the spin of the level, and π, its parity. For the ^{9}Li and ^{8}He ground states, the spin and parities are

$$\begin{aligned} J^\pi\left(^9\mathrm{Li}^{\,g.s.}\right) &= \frac{3}{2}^- \\ J^\pi\left(^8\mathrm{He}^{\,g.s.}\right) &= 0^+. \end{aligned} \tag{7.18}$$

Considering these values and the alleged J^π values [125] for the levels fed by the decays of these cosmogenic isotopes, one can build tables summarising the probable type of the transitions. Please find in Table 7.1 the probable types of the transitions to the levels fed by the ^{9}Li$^{\,g.s.}$ decays, and in Table 7.2, the probable types of the transitions to the levels fed by the ^{8}He$^{\,g.s.}$ decays. The impatient reader may find decay schemes in 7.2.1.1.

Energy (MeV)	Q_β (MeV)	J^π	Transition from ^{9}Li$^{\,g.s.}$
11.81	1.8	$\frac{5}{2}^-$	allowed
11.28	2.33	$\left(\left(\frac{7}{2}^-\right)\right)$	non-unique second-forbidden
7.94	5.67	$\left(\frac{5}{2}^-\right)$	allowed
2.78	10.83	$\frac{1}{2}^-$	allowed
2.43	11.18	$\frac{5}{2}^-$	allowed

Table 7.1 – Probable types of the transitions from ^{9}Li$^{\,g.s.}$ to the ^{9}Be excited states. The first column lists the energy of the ^{9}Be excited states, with respect to ^{9}Be$^{\,g.s.}$. The second, the maximum energy Q_β which can be attributed to an e^- during a transition to the corresponding level. The third, the most probable J^π values for these levels. The fourth, the most probable transition types. The number of parentheses around the J^π values is an indicator of the likelihood of the value assigned, the more parentheses, the greater the uncertainty on the value, and hence on the transition type.

Energy (MeV)	Q_β (MeV)	J^π	Transition from ^{8}He $^{g.s.}$
9.67	0.981	(1^+)	allowed
5.4	5.251	(1^+)	allowed
3.21	7.441	1^+	allowed

Table 7.2 – Probable types of the transitions from ^{8}He $^{g.s.}$ to the ^{8}Li excited states. The first column lists the energy of the ^{8}Li excited states, with respect to ^{8}Li $^{g.s.}$. The other columns follow the conventions established in Table 7.1.

The main point to notice on these tables is that it is likely that most transitions fall in our ballpark, being of the "allowed" type. The other point that deserves attention is the possible presence of a non-unique second-forbidden transition, occurring when decaying from ^{9}Li $^{g.s.}$ to the 11.28 MeV level in ^{9}Be. And yet, one ought not to worry too obstinately about it. Firstly, the J^π value for the latter is the most uncertain of all. Secondly, the maximum energy available for the electron in the decays to this state is only 2.33 MeV, which amounts to less than 20% of the total energy available[6] when decaying ^{9}Li $^{g.s.}$ to (e, α, α, n). Last but not least, this controversial state is fed by only 2.2% of the βn-decays. To put it differently, the uncertainties related to the modelling of the β-decays to this state do not prevail in the complete cosmogenic spectrum.

7.1.2.3 Accounting for the width of the daughter nucleus

However well-tested for allowed transitions, BESTIOLE does not take into account the width of the daughter nucleus.

For all practical purposes, the β-emitters themselves have no widths. Indeed, ^{9}Li or ^{8}He have lifetimes of the order of a hundred of milliseconds, which would give a width $\Gamma = \frac{\hbar}{\tau}$ (with τ the lifetime) of the order of 10^{-15} eV. On the other hand, the daughter nuclei may well have widths of the order of the MeV, significantly affecting the value of the endpoint in the β-decays, and hence the energy of the electron.

More details concerning the implementation of widths in the Spectrum Generator are to be read in 7.2.2. Nevertheless, let us stress that BESTIOLE, which relies on Evaluated Nuclear Structure Data File (ENSDF) files, had to be run several times with custom database inputs, containing only one branch, and with a different endpoint Q_β at each iteration, thereby accounting for the width of the target level in the daughter nucleus. The goal was to save, for each β-emitter, and for each level in the daughter nucleus, a two-dimensional histogram, as a ROOT file [105]. The x-axis of each two-dimensional histogram thus saved represents the kinetic energy of the electron T_e, in steps of 25 keV, whereas the y-axis represents the Q_β of the branch. The Q_β's have been generated using a variable binning, which

[6] Some may rightfully underline that in a liquid scintillator, electrons and positrons account for most of the visible energy; the third argument awaits them.

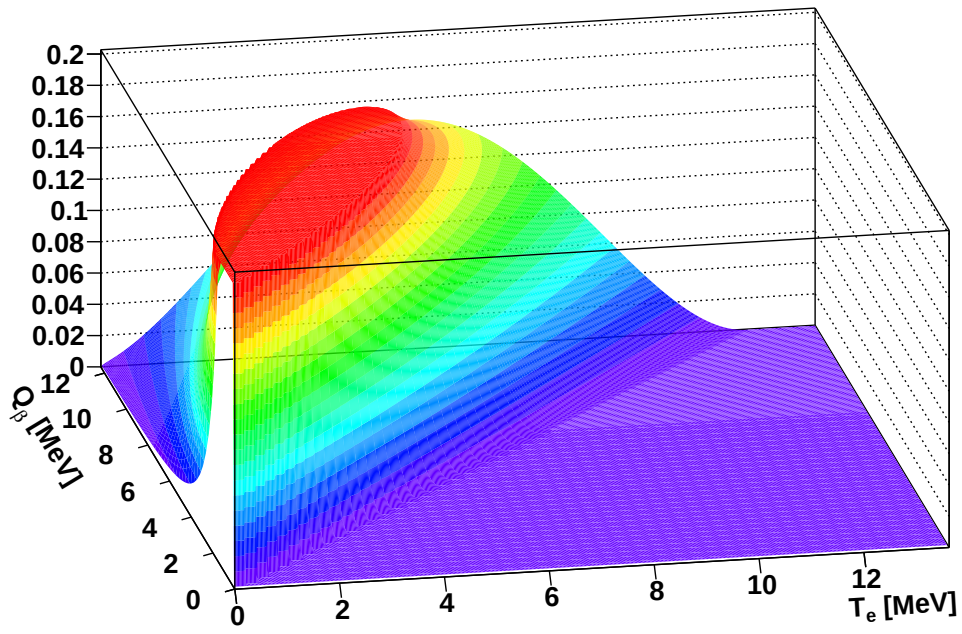

Figure 7.2 – Two-dimensional histogram used to retrieve the electron kinetic energy distribution, for various values of the endpoint, in the β-decay of the ^{9}Li ground state to the 2.78 MeV level in ^{9}Be. The x-axis represents the kinetic energy T_e of the electron and the y-axis the endpoint Q_β of the decay.

ensures that 100 points are available in the range spanning 6Γ, centred around the mass of the resonant state. For the largest resonance in ^{9}Be, the 2.78 MeV level, this leads to a central binning in 66 keV steps. As the distribution used to model resonances flattens, wider steps of 0.25Γ mark the y-axis. An example of such a map is given in Figure 7.2.

Each set of 2D-histograms representing the possible decay branches of a given nucleus is saved once and for all, in a ROOT file, to be used later by the Spectrum Generator. A 2D-dimensional histogram comes into play in the following way: for each simulation of a complete cosmogenic decay, random positions for the levels forming the decay path are pulled, which results in a random

$$Q_\beta \simeq m^{at.}_{mother} - m^{at.}_{daughter}\,, \tag{7.19}$$

with $m^{at.}_{mother}$ and $m^{at.}_{daughter}$, the atomic masses[7] of the decaying nucleus and its daughter, respectively. The 2D-histogram is then projected onto the x-axis at $y = Q_\beta$. As Q_β does not necessarily correspond to the centre of a particular y-bin, a linear interpolation between

[7]When using atomic masses, the electron mass naturally appears in the mass difference and the binding energy of the electron can be safely neglected.

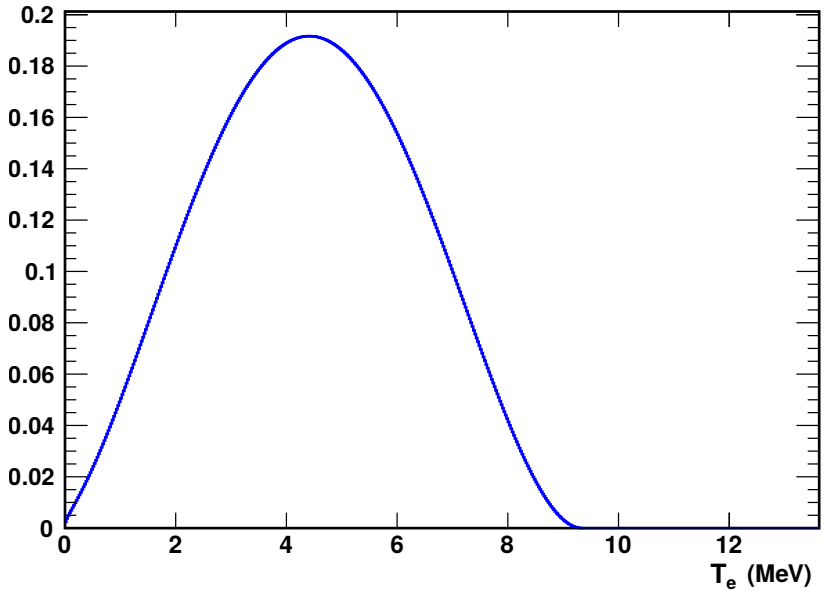

Figure 7.3 – Distribution of the kinetic energy of an electron in the β-decay of the ^{9}Li ground state to the 2.78 MeV level in ^{9}Be. The distribution was obtained by slicing the y-axis of the two-dimensional histogram presented in Figure 7.2 at $Q_\beta = 9.342$ MeV.

the 1D histogram corresponding to the bin centre y_{bin} closest to Q_β , and that of the next bin ($Q_\beta > y_{bin}$) or the previous one ($Q_\beta < y_{bin}$), allows the creation of an interpolated 1D-histogram. With such an interpolated histogram (an example of which is graphed in Figure 7.3) a value for the kinetic energy of the electron in the selected β-decay can be picked. The four-momentum follows as described in 7.1.2.1.

7.2 Simulating a raw event

So much for the kinematics of the disintegrations, now let us put together all these decays into a chain. This section aims at introducing the modelling decisions in the Spectrum Generator allowing to produce what we refer to as a "raw event", i.e. the simultaneous release of all the outgoing particles from a complete decay path.

7.2.1 First decays

7.2.1.1 Decay schemes

Before discussing the assumptions affecting the handling of the triggering events of the background signal, let us first draw the decay schemes of the βn-emitters to get an overview of the chains we have to process.

With the composition of the liquids that fill the detector in mind, only relatively light

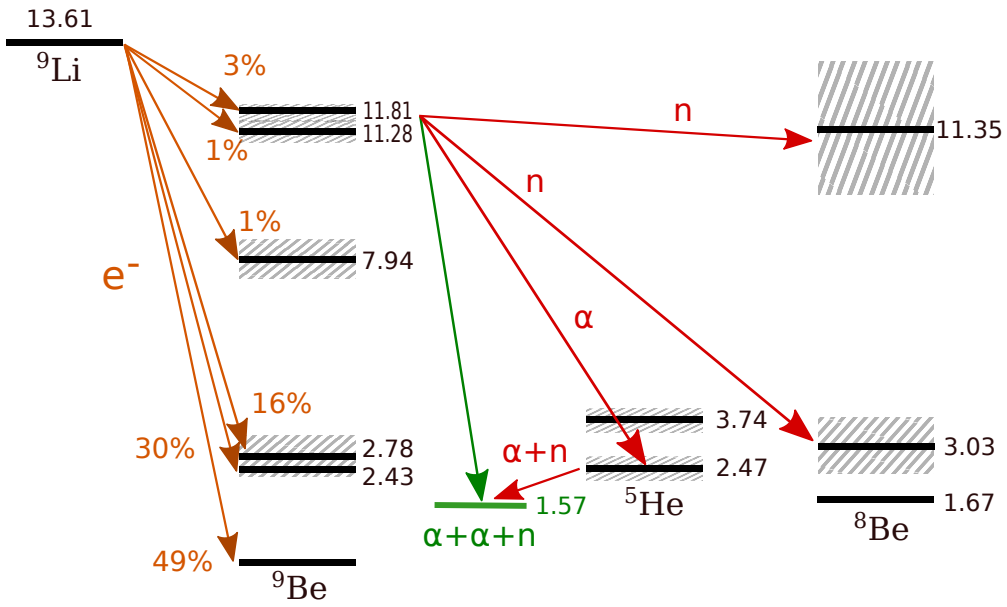

Figure 7.4 – ^{9}Li decay scheme. The energy levels are in MeV and referred to the ^{9}Be ground state. See text for more explanations. Inspired by [125].

isotopes can be produced, most of them from spallation on ^{12}C. Amongst them, only the βn-emitters are to be convicted, which shrinks the list [126] to two outcasts: ^{8}He and ^{9}Li. Consequently, most of the examples and results in what follows will be driven by the decay trees of those two, though the Spectrum Generator code can theoretically eat up any nuclide whose decay tree is given and in which the dynamics are negligible everywhere but for the β-decays.

The ^{9}Li scheme is given in Figure 7.4. Please note that the energy levels, which are given with respect the ^{9}Be ground state, make sense only when corrected for the particles emitted between them. For instance, the 3.03 MeV level in ^{8}Be is found by computing $m^{g.s.}_{^{8}\mathrm{Be}} + m_n - m^{g.s.}_{^{9}\mathrm{Be}}$, with $m^{g.s.}_{^{A}\mathrm{X}}$ the mass of the ground state of ^{A}X. Such conventions allow to gauge the energy released in a decay at a glance, although, as the formulae from 7.1 hint at, the absolute atomic masses of the elements are used within the software.

In Figure 7.4, the energy widths of the states are proportional to the hatched boxes. The $\beta-$decay branching ratios are indicated, as well as a few possible decay paths. The particles released when going from one level to another are indicated on the arrows, using unambiguous notations. Even when not indicated for the sake of clarity, the final state for the ^{9}Be decay, which follows the ^{9}Li decay, must always be understood as $\alpha + \alpha + n$. The green arrow represents the case of a direct three-body decay into these three particles. Roughly speaking, any level kinematically allowed may be reached during a decay.

The ^{8}He scheme is not significantly different from that of ^{9}Li, it also starts with the β-emitter, and leads to exceptionally unstable states in nearly 17% of the cases[8], which are the

[8]The industrious reader might have noticed that the sum of all β-ratios slightly exceed 100%; we are

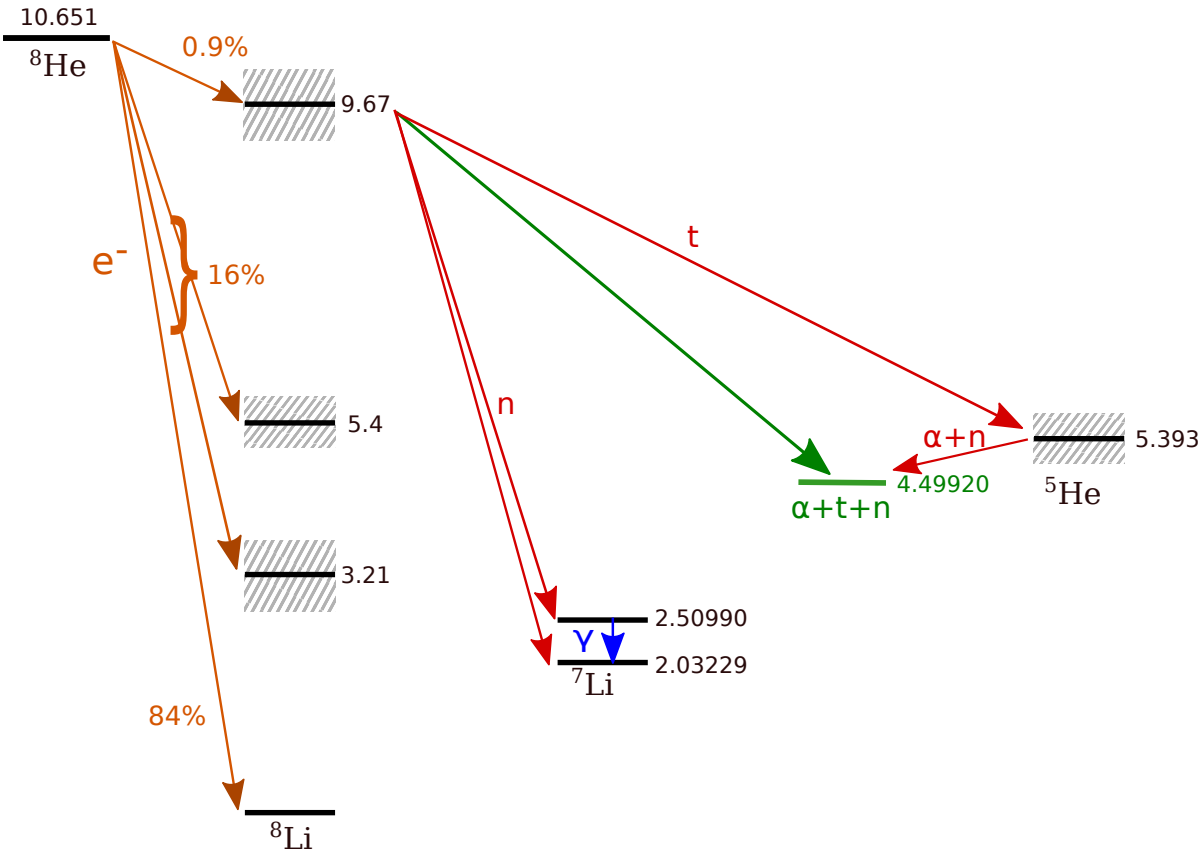

Figure 7.5 – ^{8}He decay scheme. The energy levels are referred to the ^{8}Li ground state. See text for more explanations. Inspired by [127].

ones of interest to our matters, for only they are accompanied by the release of a neutron. In fact, the ^{8}He scheme looks like a less cluttered version of the ^{9}Li scheme, as Figure 7.5 demonstrates.

Except for the fact that the energy levels are referred to the ^{8}Li ground state, the graphical conventions are identical to the ^{9}Li ones. Please note that the ^{7}Li ground state is stable with regard to the strong force. Therefore, $^{7}\mathrm{Li}^{g.s.}$ will be one of the heavy bullets shot into the detector, albeit to little din, in view of the quenching of the light output of heavy particles. The ^{8}He decays may also proceed through the ^{5}He mode, as ^{9}Li, although instead of disposing of an α particle, ^{8}Li will emit the relatively long-lived tritium isotope (t in our scheme).

7.2.1.2 Rest decays

As is clear in Figures 7.4 and 7.5, a cosmogenic decay chain always begins with a β-decay. Cosmic μ's generating the cosmogenic isotopes do have a high kinetic energy usually, but once a cosmogenic nuclide has been spawned, it decays at rest (well, it would if it were not

simply sticking to the conventions from [125].

for its thermal energy). As a consequence, β-decays are considered to happen at rest in the Spectrum Generator.

The electron mass is extremely small in comparison to that of the cosmogenic nuclei, however light they may seem (with an atomic mass number $A = 8$, they are still more than four orders of magnitude heavier than an electron). In accordance, the strong decay happening right after the first decay (in Figure 7.5 it would mean the ^{8}Li break-up) can also be performed at rest, without impairing the accuracy of the final cosmogenic spectra.

7.2.2 Dealing with resonant states

As can be seen in Figures 7.4 and 7.5, the cosmogenic isotopes decay to unstable states with large widths. How are we to handle these when trying to simulate one cosmogenic event, i.e. when trying to predict the four-momenta of all the particles generated by the decay of one βn-emitter ?

7.2.2.1 Non relativistic Breit-Wigner distribution

The Breit–Wigner approach in particle physics is meant to take into account the finite width of a meta-stable particle, and it is the first method to come to one's mind. The validity conditions are nonetheless as various as the number of approaches (see [128]). Besides, several states with broad resonances are close to one another in the cosmogenic decay schemes, this is particularly true for the ^{9}Be levels. In a nutshell, these states overlap and should be affected by hard-to-determine quantum interferences.

For lack of a better knowledge, a simple hypothesis has been considered in the Spectrum Generator to account for the widths of the states, and interferences have been overlooked. This hypothesis comes down to picking – in appropriate distributions – masses for the states that form a complete decay path, and regard them as bound states, i.e. apply all the rules given in 7.1. Insofar as a resonant state is not simply a stable particle with a given mass, such a procedure assuredly is an approximation.

As suggested, the literature bears mention of different "Breit-Wigner" distributions. The first one is known under either of the following names: "non-relativistic Breit-Wigner distribution", "Cauchy distribution", or "Lorentzian"; the dependence of the Cauchy distribution is the square of the energy. The second one goes under the name of "relativistic Breit-Wigner distribution" and its dependence is the energy to the power of four, making it valid for narrower resonances in highly relativistic decays. Hence, to give a mass m_{s_X} to one particular level s of a nuclide X, one could use the following "non-relativistic" Breit-Wigner distribution

$$l\left(m_{s_X}; m^0_{s_X}, \Gamma_{s_X}\right) = \frac{\frac{\Gamma_{s_X}}{2}}{\pi\left[\left(m_{s_X} - m^0_{s_X}\right)^2 + \left(\frac{\Gamma_{s_X}}{2}\right)^2\right]}, \tag{7.20}$$

with Γ_{s_X} the width of the state as given in the nuclear databases [78], and $m^0_{s_X}$ the mass of the ground state of the nuclide X corrected by the mean excitation energy of the level E_{s_X},

i.e. $m^0_{s_X} = m_{g.s._X} + E_{s_X}$. It must be underlined that what is commonly called "the width of a state" actually corresponds to the full width at half-maximum of the non-relativistic Breit-Wigner distribution. Thus, let us introduce the half-width at half-maximum, which simply reads $\gamma_{s_X} = \frac{\Gamma_{s_X}}{2}$.

7.2.2.2 Issues with the non-relativistic Breit-Wigner distribution

Extension of the distribution

A non-relativistic Breit-Wigner, however, wanes slowly. Even more than $3\gamma_{s_X}$ away from the central value, the distribution is still going strong. To give an idea of the slow drop of the distribution, we drew in Figure 7.6 a non-relativistic Breit-Wigner distribution of mean value $m^0_{s_X} = 10$ and width $\Gamma_{s_X} = 1$ (we have dropped units momentarily, but it could well model a 1 MeV resonance for a 10 MeV level as the mass of the ground state induces but a shift).

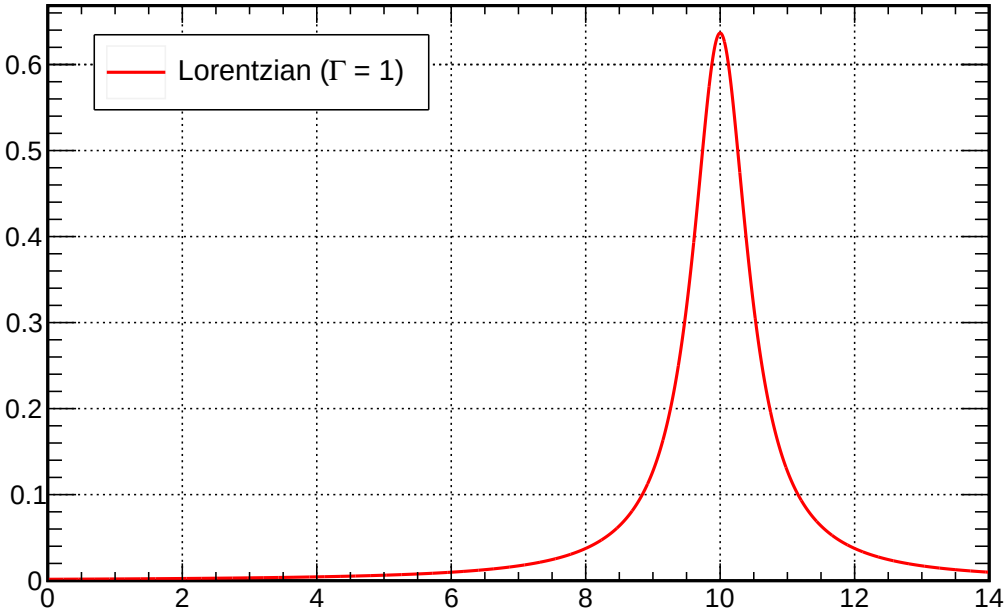

Figure 7.6 – Non-relativistic Breit-Wigner distribution centred at 10. The full width at half-maximum of this non-relativistic Breit-Wigner distribution evaluates to $\Gamma_{s_X} = 1$.

The probability to pick values for the mass m_{s_X} of a state, around its central value $m^0_{s_X}$, neither depends on the largeness of the state, nor on the value of $m^0_{s_X}$. In a nutshell, a probability table, valid for all nuclides, can be built with ease for any such distribution. Table 7.3 summarises the values of the integral of the non-relativistic Breit-Wigner distribution between symmetric bounds, unambiguously emphasising the colossal extension of the function.

Bounds	Integral Value
$\pm 1\gamma_{s_X}$	0.5
$\pm 2\gamma_{s_X}$	0.7
$\pm 3\gamma_{s_X}$	0.8
$\pm 4\gamma_{s_X}$	0.84
$\pm 5\gamma_{s_X}$	0.87
$\pm 10\gamma_{s_X}$	0.94

Table 7.3 – Value of the integrals of a non-relativistic Breit-Wigner distribution between symmetric bounds. The first columns lists the integration bounds around the central value of the distribution, in terms of the half-width at half-maximum γ_{s_X}, and the second one, the value of the integral between these bounds.

Asymetric bounds

When simulating cosmogenic events, the kinematical constraints are strong bounds to the extension of the distribution modelling each resonant state. Without doubt, the mass of a ^{9}Be state cannot go below the sum of the masses of two ^{4}He and a neutron, as the green line in Figure 7.4 epitomises. Similarly, the position of the ^{8}He ground state acts as a 10.651 MeV-high ceiling for the ^{8}Li states (see Figure 7.5): no level above this value can be reached by β-decay. Consequently, all the distributions involved in a decay path are harshly chopped, and in asymmetrical way.

The curve drawn in Figure 7.6 exhibits this asymmetry and the mean[9] value of the Lorentzian is a bit off from the most probable value $m^0_{s_X} = 10$, if we restrict ourselves to the $[0; 14]$ range. This asymmetry means that the mean position of the 2.78 MeV level in ^{9}Be will actually appear a few hundreds of keV higher. As a result, if one saves the kinetic energy of the electron T_e, when generating a raw spectrum for ^{9}Li decays always feeding the 2.78 MeV level in ^{9}Be, the mean β-spectrum will look distorted when compared to the spectrum retrieved when widths are zero. Such a bias in the mean energy share for this dominant ^{9}Li decay mode is embodied by Figure 7.7. Make note of the fact that all the distributions (more about that follows in 7.2.3) of the decay chain contribute to the observed distortion. Please also notice that in the presence of a width, the spectrum extends to higher T_e's and hence higher Q_β's.

7.2.2.3 Gaussian modelling

Considering that the use of non-relativistic Breit-Wigner distributions modifies the energy share between the particles in each raw decay, and straightforwardly the raw spectrum of

[9]Lorentzian provide no expected value since $m\, l\left(m; m^0, \Gamma\right) \sim \frac{1}{m}$ for $m \to +\infty$, which is not integrable there, but for our limited range, there is indeed a varying mean, different from m^0.

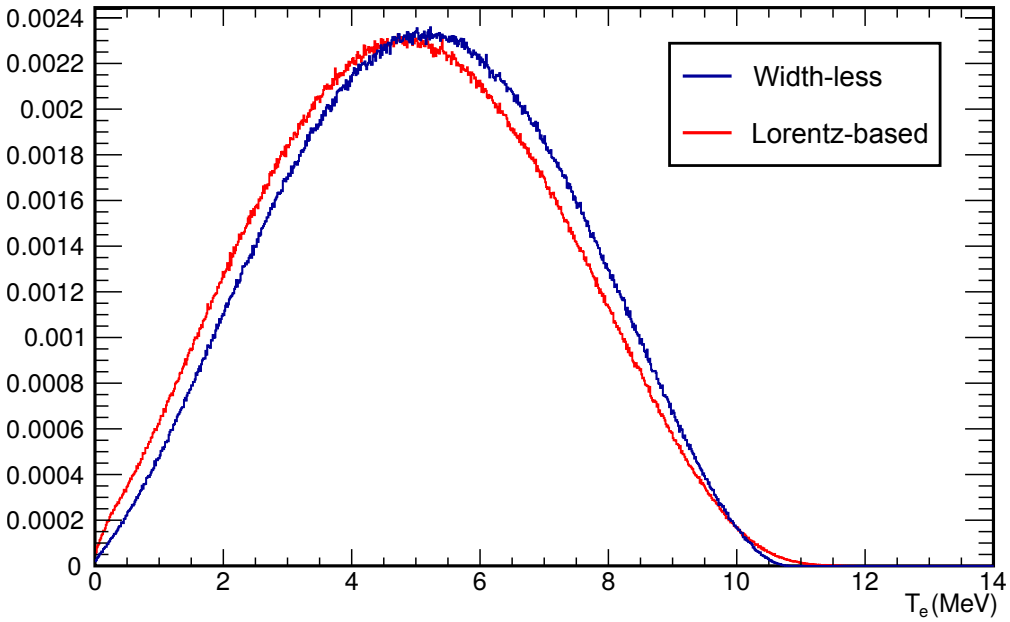

Figure 7.7 – Raw β-spectra obtained when generating ^{9}Li decays always feeding the 2.78 MeV state in ^{9}Be. The modelling of the widths via Lorentzians (red) biases the mean kinetic energy of the electron, as obtained from the width-less generation (blue).

each individual branch, the compromise role befits the normal distribution. The latter allows indeed one to scan the available phase space, while yet remaining at an acceptable distance of the bounds imposed by the kinematical constraints.

Figure 7.8 shows the comparison between a non-relativistic Breit-Wigner distribution of mean value $m^0_{s_X} = 10$ and width $\Gamma_{s_X} = 1$ (we have left the units out as in 7.2.2.2), and a Gaussian with a mean value $\mu_{s_X} = 1$ and a standard deviation $\sigma_{s_X} = \frac{1}{2\sqrt{2\ln(2)}}$. The parameters of the distributions in Figure 7.8 have been chosen so that the full widths at half-maximum are identical and equal to one. Since a Lorentzian does not have finite moments of any order, RMS comparisons would have been intricate, to say the least. Therefore, the standard deviation σ_{s_X} to set, when modelling the width of a level s in a nuclide X with a Gaussian, should relate to the nuclear width Γ_{s_X} according to

$$\sigma_{s_X} = \frac{\Gamma_{s_X}}{2\sqrt{2\ln(2)}}, \tag{7.21}$$

hence defining the normal distribution as

$$f\left(m_{s_X}; m^0_{s_X}, \Gamma_{s_X}\right) = \frac{2}{\Gamma_{s_X}}\sqrt{\frac{\ln(2)}{\pi}}\mathrm{e}^{-4\ln(2)\frac{\left(m_{s_X}-m^0_{s_X}\right)^2}{{\Gamma_{s_X}}^2}} . \tag{7.22}$$

To avoid the issues largely covered in sub-section 7.2.2.2, the Gaussian modelling was retained in the Spectrum Generator and definitive results using the conversion from nuclear widths to standard deviations, as described in equation 7.21, will be presented in section 8.4.

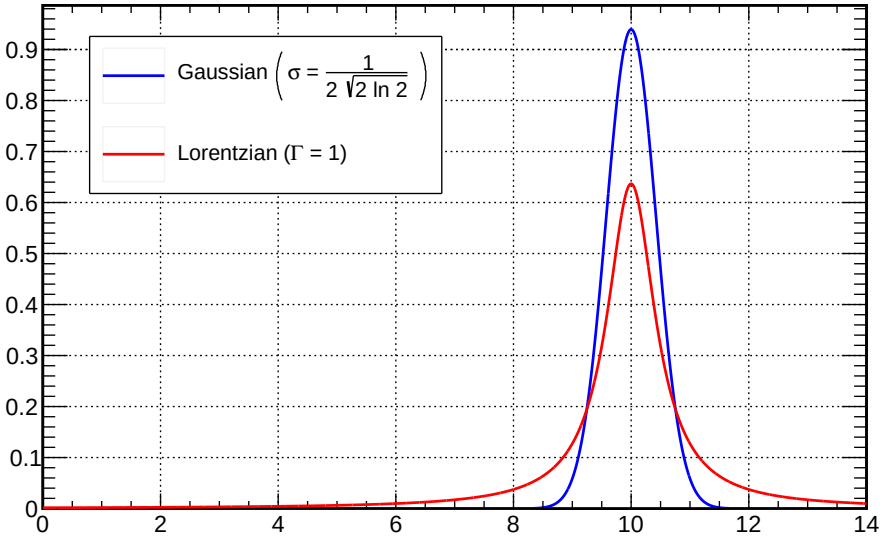

Figure 7.8 – Comparison of a normal distribution (blue) and a non-relativistic Breit-Wigner distribution (red). The distributions are centred at 10. The standard deviation of the Gaussian is $\sigma_{s_X} = \frac{1}{2\sqrt{2\ln(2)}}$, whereas the width of the Lorentzian is $\Gamma_{s_X} = 1$.

7.2.3 The chain constraints

In order to simulate a complete decay path one does not only need to know how to simulate each raw decay (see 7.1), how to deal with the widths of a state (see 7.2.2), one also needs to handle the constraints that bind together all the states involved in a decay chain. After the β-decay, which always comes into play first, the cosmogenic isotopes reach strongly unstable states, whose lifetimes are well below 10 ns. Therefore, all the decays after the first, release particles which appear at the same time in the detector, and any time evolution during the decay chain is virtually non-existent for the cosmogenic break-ups in a reactor experiment. Nonetheless, we must preserve a logical ordering and watch for biases.

7.2.3.1 Sequential processing

Plainly, there is first need to guarantee proper forwarding of the four-momenta of the decaying particles to subsequent decays so that they may be performed in-flight. This part is just an inductive step; only two levels at a time are involved in the forwarding.

As far as the widths are concerned, this is a completely different tale, all the states have to be handled concurrently. Once a decay chain has been chosen, it must not be processed sequentially. In other words, we cannot harmlessly pull a position for the next level, decay to it, and afterwards, pick the position of the new level to process.

Let us exemplify this with the decay of the ^{8}He ground state. The textbook chain for the matter is $^8\mathrm{He}^{\,g.s.} \to {}^8\mathrm{Li}^{\,9.67} \to {}^5\mathrm{He}^{\,g.s.} \to (\alpha, n)$. In the latter chain, we shall not pick the actual position of the $^8\mathrm{Li}^{\,9.67}$ state in a Gaussian centred at 9.67 MeV with $\Gamma_{{}^8\mathrm{Li}^{9.67}} = 1\mathrm{MeV}$,

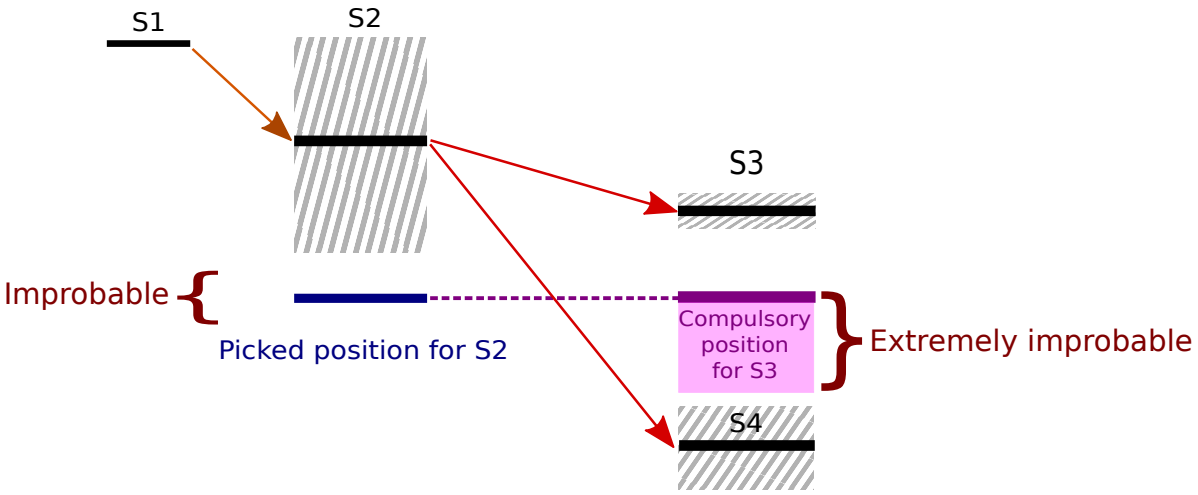

Figure 7.9 – Effect of processing the decays in a sequential chain. What should be an extremely improbable configuration becomes a relatively common one.

we shall not perform the decay to this $^{8}\mathrm{Li}^{9.67}$ state, and we shall not eventually pick – from this state – the mass of $^{5}\mathrm{He}^{g.s.}$ in a second Gaussian. Indeed, should we proceed sequentially, we would simply neglect the Gaussian associated to the next state because if the $^{8}\mathrm{Li}^{9.67}$ were pulled low-lying, we would have to pick random positions for $^{5}\mathrm{He}^{g.s.}$ until one was sufficiently low to make the decay $^{8}\mathrm{Li}^{9.67} \rightarrow {}^{5}\mathrm{He}^{g.s.}$ kinematically possible.

This awkward situation is illustrated by Figure 7.9. When processing the $S_1 \rightarrow S_2 \rightarrow S_3$ chain sequentially, the probability of pulling a low-lying value for the mass of S_3 is simply meagre, if S_2 has a large width. However, picking both S_3 and S_2 so low should seldom happen, in light of the average width of S_3. In a nutshell, a sequential processing amounts to looking away the distribution associated to the position of S_3. The upwards situation is also true, if one state is picked too high in a sequential process, the next one will also be forced to go well above its mean position. By and large, the probability of such misfortunes is driven by the probability of pulling the first state far from its expectation, considering the next one is picked until it fits.

7.2.3.2 Synchronous processing

Although still approximate (because the branching ratios should be dynamically rescaled as the positions of the states drift off their mean positions), a better approach to sequential decays is to split the simulation of a decay path into three stages, each of them simultaneously affecting all the states of the path.

Firstly, pick a decay chain $S_1 \rightarrow \cdots \rightarrow S_n$, where n stands for the number of states S_i in the chain (the numbering for the states is different from the one used in Figure 7.9). Secondly, pick a random position for $\overrightarrow{m} = (m_{S_1}, \ldots, m_{S_n})$ where m_{S_i} is the mass of the i-th state in the decay chain. This random position is pulled using the uncut normal distributions

(as defined in equation 7.22) for each state. Choosing the layout of the states amounts to picking $\overrightarrow{m}$ following

$$g\left(\overrightarrow{m}\right) = \prod_{1 \leq i \leq n} f\left(m_{S_i}; m^0_{S_i}, \Gamma_{S_i}\right), \tag{7.23}$$

and checking whether the decay chain is kinematically possible, i.e. if

$$\forall i \in [\![1; n-1]\!],\ m_i \geq m_{i+1} + \sum_{j=1}^{l_i} m_i^j\,, \tag{7.24}$$

where $\sum_{j=1}^{l_i} m_i^j$ is the sum of the masses of the stable particles (α's, tritons within our time scale and so on) released in the l_i-body decay $S_i \rightarrow S_{i+1}$. Masses $\overrightarrow{m}$ have to be picked following g in 7.23 until one set is kinematically possible. At last, the chosen chain can be processed with the value obtained for $\overrightarrow{m}$ as if the states were stable (cf. 7.1), only with a mass different from their mean masses.

The method reviewed here is the one retained to generate raw decays in the Spectrum Generator.

7.2.4 Position of the event

As highlighted in 7.2.3, all the decays of a chain which make up a "raw event" release particles which will appear as emanating from the same point in space-time. The cosmogenic decays in the Double Chooz experiment have a low rate of a few events per day (cf. Table 6.3). Accordingly, little attention is paid to the time at which events are generated as long as they are far enough from one another in time, which implies that they cannot interact with one another. There remains to choose a spatial location for each event.

7.2.4.1 Carbon share

What drives the production of cosmogenic isotopes? Supposedly, the number of carbon atoms in the detector dictates how many cosmogenic isotopes are likely to be produced, and consequently, how many cosmogenic events are to happen. Generating the right number of events in each volume is a serious matter inasmuch as neutrons spilling in or out (cf. 6.3.2.3) will loose some invisible energy in the acrylics separating those . In order to assign an event to either the neutrino target or the gamma catcher[10], the probability to find an event in the target has been defined as

$$p_{NT} = \frac{m^C_{NT}}{m^C_{NT} + m^C_{GC}}, \tag{7.25}$$

where m^C_{NT} is the mass of carbon in the NT, and m^C_{GC} that of carbon in the GC. Please note that the ratio in (7.25) is equal to the same ratio with the mass swapped for the number of carbon atoms.

[10]The volume occupied by the acrylics themselves was neglected. Cosmogenic production in the buffer was neglected as well.

The mass of carbon in the neutrino target can be swiftly computed using the weight measurement of the NT liquid scintillator – which yielded a slightly dated central value $m_{NT} = 8287\,\mathrm{kg}$ – and its carbon weight density $wd^C_{NT} = 86.0\%$. On the other hand, for the GC, one can but use the volume of the vessel $V_{GC} = 22\,542\,\mathrm{L}$ to gauge an actual carbon mass. The volumetric mass density of the GC liquid is $d_{GC} = 0.804\,\mathrm{kg\,L^{-1}}$, and its carbon weight density is $wd^C_{GC} = 85.3\%$. Plugging in all the numbers leads to

$$p_{NT} = \frac{m_{NT}\; wd^C_{NT}}{m_{NT}\; wd^C_{NT} + V_{GC}\; d_{GC}\; wd^C_{GC}} \simeq 0.316\,, \tag{7.26}$$

which, in view of the similarity of the carbon contents of the NT and the GC, is quite close to the ratio one would obtain by using the NT and GC masses directly. Considering, there is slightly less carbon in the GC than in the NT, the rough mass fraction would yield two per mille lower probabilities.

In (7.26), relatively dated values – from the FDI analysis – were considered. As of 2016, the NT mass and the GC mass have been evaluated thoroughly for both the near and far detector [129], bringing forth

$$p^{ND}_{NT} = 0.315 \tag{7.27}$$
$$p^{FD}_{NT} = 0.317\,, \tag{7.28}$$

which fortunately remains close (per mille difference) to the (7.26) used in the Spectrum Generator. As a stroke of luck, both the ND NT mass, and the ND GC volume – albeit measured completely independently – appear to be higher than their FD counterparts.

7.2.4.2 Position generation

A random number $p \sim \mathcal{U}(0,1)$, with $\mathcal{U}$ the uniform law, is picked for deciding whether the event has occurred in the target ($p \leq p_{NT}$) or the gamma catcher.

Denoting by r the radial distance, θ the azimuth angle, and z the signed height with respect to the NT centre, in cylindrical coordinates, the volume of an elementary element is

$$r\,\mathrm{d}\theta\,\mathrm{d}r\,\mathrm{d}z = \mathrm{d}\left(\frac{r^2}{2}\right)\mathrm{d}\theta\,\mathrm{d}z\ . \tag{7.29}$$

It follows that θ can be pulled uniformly as $\theta \sim \mathcal{U}(0, 2\pi)$ and $z \sim \mathcal{U}\left(-\frac{h_{GC}}{2}, \frac{h_{GC}}{2}\right)$ where h_{GC} is the height of the GC vessel. As demonstrated by (7.29), the radial distance of the position $(r\cos\theta, r\sin\theta, z)$ may not be generated uniformly. According to whether $|z| < h_{NT}$, with h_{NT} the height of NT vessel, r^2 must be generated following $r \sim \mathcal{U}\left({R_{NT}}^2, {R_{GC}}^2\right)$ or $r \sim \mathcal{U}\left(0, {R_{GC}}^2\right)$, where R_{NT} and R_{GC} are the radii of the NT and the GC, respectively. From there, we need but take the square root of r^2 and fill the obtained position, along with the four-momenta of the particles released by the decay chain (cf. 7.2.3.2), into a Geant4-compatible file with which it will be fed.

Chapter 8

Obtaining predicted spectra

Chapter 7 threw light on the computation of the four-momenta appearing in each decay mode of the cosmogenic isotopes. Nevertheless, there is more to spectra generation than accumulating a stack of decay events. Generating characteristic spectra entails putting all the possible decay modes into perspective, and representing the whole tree, rather than a fully determined path. Defining how the decay chains are effectively selected implies connecting with nuclear data, whose limited availability calls for a dedicated uncertainty propagation. Moreover, the predictions must eventually be compared to data. To this end, not only must they be passed through the detector simulation, but the effect of the analysis cuts on them must also be reviewed.

A flow chart of the architecture of the simulation chain, which will be referred to throughout this chapter, is given in Figure 8.1. The reader is already partly acquainted with the first tools of the chain, i.e. the Spectrum Generator and DOGS, the latter being able to process files produced by the former, to simulate detection. The last point is embodied by the Covariance Tool, which simultaneously provides the mean simulated spectrum and its systematic errors.

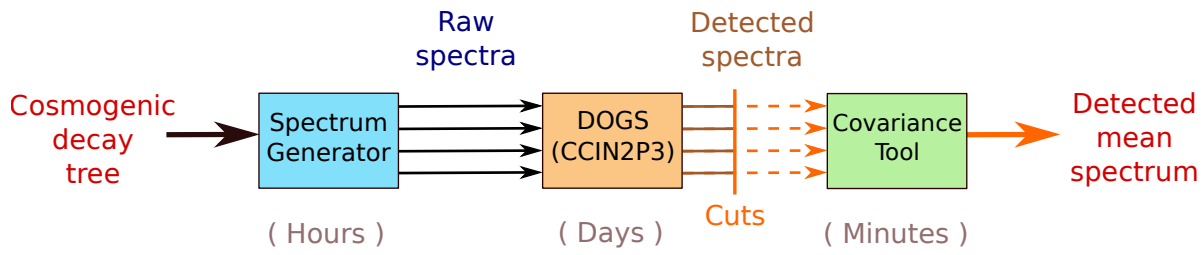

Figure 8.1 – Architecture of the simulation chain for the cosmogenic spectra generation in Double Chooz. The approximate running time for 1 million events is indicated in parenthesis. DOGS is run at Lyon's Computing Centre on hundreds of logical cores. All the other tools can be painlessly run on a notebook.

8.1 Handling the decay trees

Galvanised by our expertise in the processing of a decay chain, it is high time we tackled something more global: the decay trees. We will first describe the ideal operating of the Spectrum Generator in 8.1.1 and shape it for our peculiar needs in 8.1.2.

8.1.1 Generating a raw spectrum

8.1.1.1 Input

At the moment, for every new decay to simulate, the decay tree has to be written in `C++` code (objects for handling text files as input might be written later). Each tree has to be coded in the form of a functor that allows to save all the physical parameters needed to perform the decays. This is useful for multi-threading as well.

The event generation method makes the most of dynamic polymorphism, which amounts to being able to supply an object at runtime with virtual member functions implementing the relevant strategy. This enables the core of the Spectrum Generator to treat different decay types implemented in different classes (for instance `BetaState` for β-decaying states) as simple objects all inheriting from a base class `State`. In a nutshell, `state.process()` redirects to the right strategy whether `state` is actually a β-emitter or a virtual state or some other runt of the base class.

In a like manner, dynamic polymorphism allows to store different types of states in a common container from the Standard Template Library (STL) [130]. Being able to use a unique container proves particularly valuable when storing the daughters of a mother nuclide, which can be registered within it as a vector of pointers – or rather, references – to instances of the base class `State`, regardless of their actual type.

8.1.1.2 Processing

Point 7.2.3 made it clear that all the states in the chosen chain had to be treated simultaneously when it came to setting their positions. Therefore, the tree built in 8.1.1.1 has to be processed three times to create but one event for the raw spectrum.

First, the event generator has to go through the tree using the overridden `Next` method to choose one path. The search stops whenever a state whose daughters are all stable is reached. Secondly, it has to shake the positions of the levels – in the sub-part of the tree representing the chosen path – to account for the resonances (cf. 7.2.3.2). Eventually, the selected path has to be browsed through a third time, to actually decay the states (cf. 7.1), whose masses have been set by the second step. During this final action, the four-momenta of the outgoing particles are saved.

Thankfully, custom and lightweight classes have been written to limit the dependencies to the ROOT library, thereby providing a thread-safe environment allowing to work on several

branches at a time. On four logical[1] cores only, this has divided the execution time of the Spectrum Generator by more than six, somehow turning the $\sim 1\,\mathrm{h}$ from 8.1 into a mere 10 m undertaking.

8.1.1.3 Saving the output

So to speak, this is not really a "spectrum" that feeds the next tool of the simulation chain, DOGS. Rather, standard High Energy Physics (HEP) files, which are but text files, containing all the relevant information to be considered in a detector simulation, are handed over to the next processing tool. These files gather properties such as Particle Data Group (PDG) codes that serve to identify the particle types (electron, ^{3}H and so on), four-momenta and generation positions in the detector.

8.1.2 Generating individual branches

8.1.2.1 Specifying the parameters in the raw generator

The Spectrum Generator is able to treat complete decay trees as long as the branching ratios from one branch to the others are known. There's the rub, most branching ratios in the cosmogenic decay trees are unknown.

The next tool of the simulation chain (cf. Figure 8.1), embodied by DOGS, demands 150 CPU's to work simultaneously for 30 h to propagate 1 million events[2] contained in the HEP files produced by the Spectrum Generator. Hence, there is little freedom in varying the unknown branching ratios within the Spectrum Generator itself; the flexibility is much greater when varying them after detection. Therefore, the Spectrum Generator was run so as to generate one complete decay path (see 7.2.3) per HEP file, regardless of the branching ratios. For instance, one HEP file would contain all the outgoing particles released by the chain $^8\mathrm{He}^{\,g.s.} \rightarrow {}^8\mathrm{Li}^{\,9.67} \rightarrow {}^5\mathrm{He}^{\,g.s.} \rightarrow (\alpha, n)$, as many times as we wish to have events. In the ^{8}He case this amounts to 9 HEP files, and in the ^{9}Li case this number swells to 24. To be able to create these HEP files for the full decay paths, one needs but to lock all the other branches in the tree to allow only the path of interest for a given HEP file. These files form a down-to-earth database of raw "spectra".

8.1.2.2 Simulating the detected branches

Once one HEP file has been written for each unknown branching ratio appearing in the cosmogenic decay tree of interest, all we need is to obtain the corresponding spectra after the

[1] Intel®Hyper-Threading technology somehow emulates two logical cores per physical core, i.e. we do not even demand four actual cores to better the performance.

[2] Of this million, perhaps a fourth only will pass the analysis cuts of the selected channel, itself usually restricted to one volume.

particles have been detected. This achieved by feeding DOGS, i.e. the detector simulation (cf. 5.1.3.4), with the aforementioned HEP files.

In truth, to collect sizeable statistics, each HEP file – accounting for a complete decay path – has to be temporarily split into dozens of smaller HEP files. Indeed, with the processing durations from 8.1.2.1 in mind, one had better launch thousands of asynchronous tasks at Centre de Calcul de l'Institut National de Physique Nucléaire et de Physique des Particules (CCIN2P3). Simple scripts, splitting the files and feeding DOGS with them, have been written in `Python`, emulating, by and large, parallel computing. All the files subsequently produced by DOGS, in a ROOT format identical to actual data, consequently need to be merged.

It is worth underlining that – whatever the complications may be to achieve this – the simulations of all the detectors are shaped to produce identical energy spectra, hence, the predictions do not only compare well to the data they are affiliated with, but also to the other detector versions. As a consequence, all the more so with the limited data samples to which they will be compared to, there is no need to go through several detector versions when simulating the detected branches; currently, the simulations of the cosmogenic spectra have no notion of far or near detector.

8.1.3 Applying the analysis cuts

After the simulation of the raw spectra, and after having turned the number of photo-electrons for each trigger into visible energy, we have at hand what will be referred to as "raw detected spectra". When a detected spectrum models a decay branch which releases highly energetic neutrons, the spatial cuts, as well as the cuts on the delayed energy, which select either of the two neutron-capture peaks in standard analyses, may have a non-negligible effect.

8.1.3.1 Definition

The detail of the analysis cuts that were applied to build the detected spectra database in order to select the neutron capture type are given in Table 8.1. They follow the ones which were presented in 6.1.2 to define the $\overline{\nu_e}$ selection.

The backgrounds vetoes from 6.2 were not applied, but their effect is virtually irrelevant. Firstly, they are no μ's simulated and all the μ-related vetoes are uncorrelated with the energies deposited by the rest decays of the cosmogenic isotopes. Secondly, the remaining light noise and stopping-μ cuts were designed so as to have no or insignificant impact on the prompt spectrum. Moreover, all the delayed events in the simulation are neutron-captures, i.e. they do not generate extraordinary light patterns with which the reconstruction algorithm – DCRecoBAMA – would struggle.

Variable	Cuts	
	Gd	H
E_d^{vis}	$4 - 10$ MeV	$1.3 - 3$ MeV
Δt_{pd}	$0.5 - 150$ µs	$0.5 - 800$ µs
Δr_{pd}	$0 - 100$ cm	$0 - 120$ cm

Table 8.1 – Analysis cuts on the visible energy of the delayed event E_d^{vis}, on the time interval Δt_{pd} between the prompt and delayed event, and on the distance between these two Δr_{pd}, when trying to select either Gd or H captures.

8.1.3.2 Impact

Gd captures

Given that only the liquid scintillator of the target is Gd-loaded, selecting Gd captures boils down to overlooking neutrons spilling out of the target and loosing some invisible energy in the acrylics separating the target and the gamma catcher. Furthermore, selecting only captures on Gd atoms does not allow one to see all the neutrons spilling in since some of them may land on an Hydrogen atom of the target instead. Besides, the time coincidence cut is known to efficiently curb the spill-in effect, which drives the value of its upper bound, as explained in 6.1.2. As a result, the Gd analysis cuts produce spectra with a higher bin content in the higher energy region, when compared to a mere mustering of all events, regardless of their time correlation or capturing nucleus.

The example of such an effect – all the more significant that the energy of the neutron is large – is embodied by Figure 8.2. Although on the cherry-picked raw spectrum in Figure 8.2, the impact of the cuts is substantial, let us underscore that such events form a small minority, for more than 90% of the ^{9}Li decays feed the ^{9}Be $^{2.43}$ and ^{9}Be $^{2.78}$ states (cf. Figure 7.4). Moreover, out of the remaining 10%, a non-negligible proportion should break-up via the ^{5}He and three-body roads, thus further mitigating the results. The same goes for ^{8}He, with an even greater share of relatively low-lying states, close to 95% (cf. Figure 7.5).

When lower energy neutrons are involved, the analysis cuts have a smaller effect, although it must be scaled to the mean energy of the deposits. The discrepancy observed is supported by the measurements of the quenching of the light output in the GC, for α's (cf. Table 5.1). The case of the decay chain $^{9}\mathrm{Li}\,^{g.s.} \rightarrow {}^{9}\mathrm{Be}\,^{11.81} \rightarrow {}^{5}\mathrm{He}\,^{2.47.} \rightarrow (n, \alpha)$ is given for reference in Figure 8.3.

GC dominated volume

Perhaps less of a textbook case than Figures 8.2 and 8.3 is the comparison between some branches analysed with the Gd or H cuts. For the impact to have a sizeable effect, we have also chosen to show in Figure 8.4 decays through $^{9}\mathrm{Be}\,^{11.81}$ inasmuch as these assign the lowest

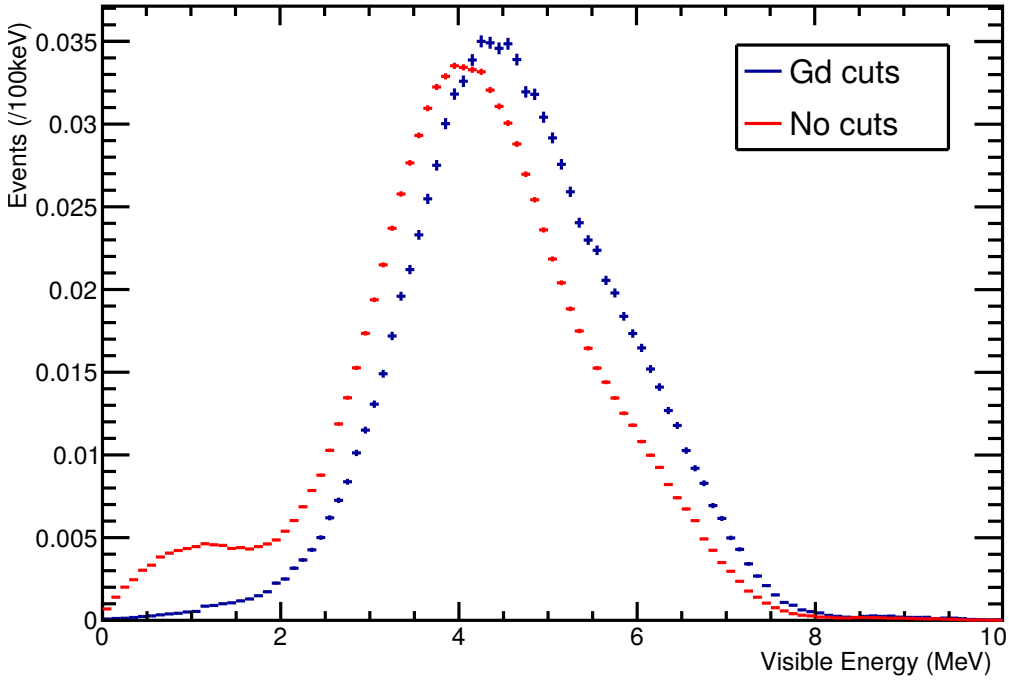

Figure 8.2 – Effect of the analysis cuts on the detected spectrum associated to the decay chain $^{9}\mathrm{Li}^{g.s.} \rightarrow {}^{9}\mathrm{Be}^{11.81} \rightarrow {}^{8}\mathrm{Be}^{1.67.} \rightarrow (\alpha, \alpha)$. Both spectra are normalised to one (hence the higher peak with the Gd cuts). Please note that the errors are purely statistical.

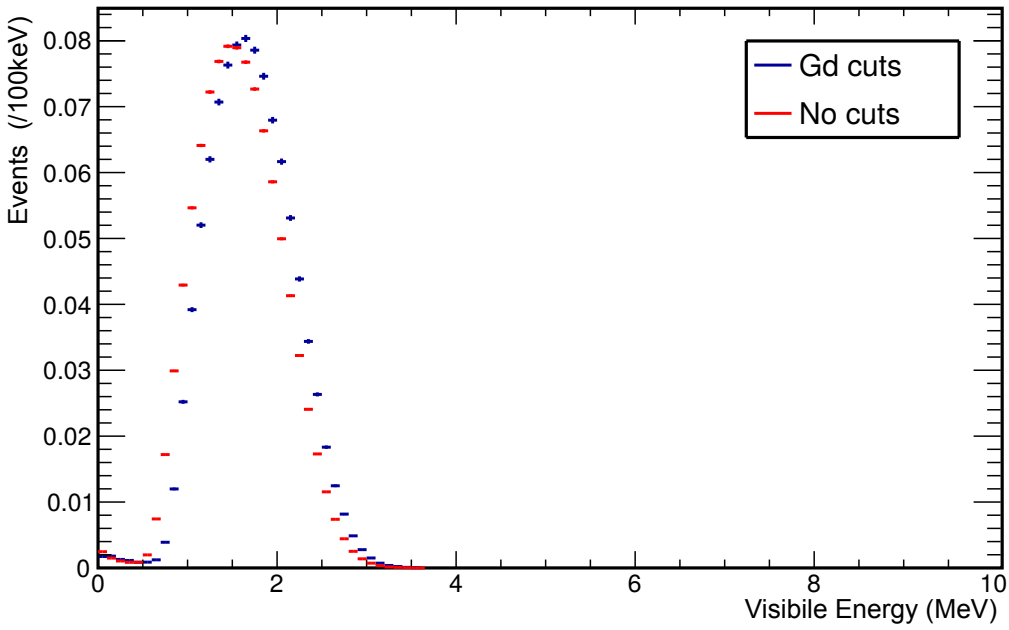

Figure 8.3 – Effect of the analysis cuts on the detected spectrum associated to the decay chain $^{9}\mathrm{Li}^{g.s.} \rightarrow {}^{9}\mathrm{Be}^{11.81} \rightarrow {}^{5}\mathrm{He}^{2.47.} \rightarrow (n, \alpha)$. Please note that the errors are purely statistical, thus accounting for the smaller error bars on the spectrum without analysis cuts.

energy to the electron, thereby emphasising the role of the neutron. When comparing both analyses, it can be observed that the most probable visible energies, in the decays through the $^{8}\mathrm{Be}^{1.67.}$ and $^{8}\mathrm{Be}^{3.03.}$ states, are 400 keV to 500 keV apart.

Again, this hardly comes as surprise, for the H cuts are much more permissive and tolerant to spill effects. If truth be told, the responses of the NT and GC liquids to fast neutrons appear to be considerably different, as the comparison of two spectra, analysed with the H cuts, but whose prompt deposits were reconstructed exclusively in the NT or GC, demonstrates. This discrepancy, compatible with different quenching of the light output in the two volumes, prevails over border effects and accounts for most of the difference we actually observe between the Gd and H spectra.

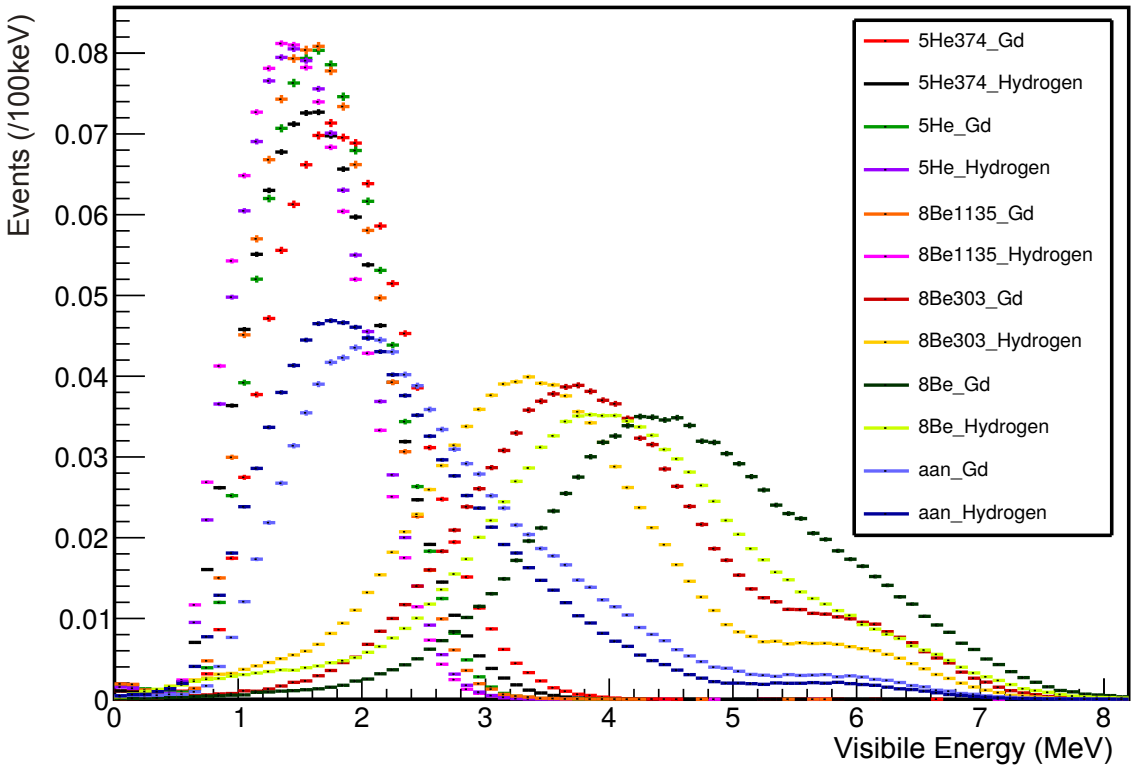

Figure 8.4 – Comparison between the Gd and H analysis cuts for the detected spectra associated to the decay chain $^{9}\mathrm{Li}^{g.s.} \rightarrow {}^{9}\mathrm{Be}^{11.81} \rightarrow X^{s}$ where X^{s} can stand for any level in ^{5}He, ^{8}Be or a three-body compound (α, α, n). The analysis type follows the label of the intermediate state on the legend. The errors are purely statistical.

8.1.3.3 Spectra database

Once the detected spectra have been analysed using the cuts presented in Table 8.1, we have at hand a database of "analysed detected spectra" (and we may drop the "analysed" in the following).

Insofar as the database for ^{9}Li tots up 24 detected spectra, the reader with standard resolving power may find it intricate to make out one histogram from the other in Figure 8.5. Nonetheless, the comparison of all the branches is the only way to comprehend the correlations between the different energy regions, and Figure 8.5 will be referenced several times in 8.4. Likewise, the ^{8}He database is also displayed here and the less bloated Figure 8.6 should draw more attention.

Considering each spectrum in the database is associated with an unknown branching ratio, the last step is to average these spectra and obtain the correlations between the different energy regions.

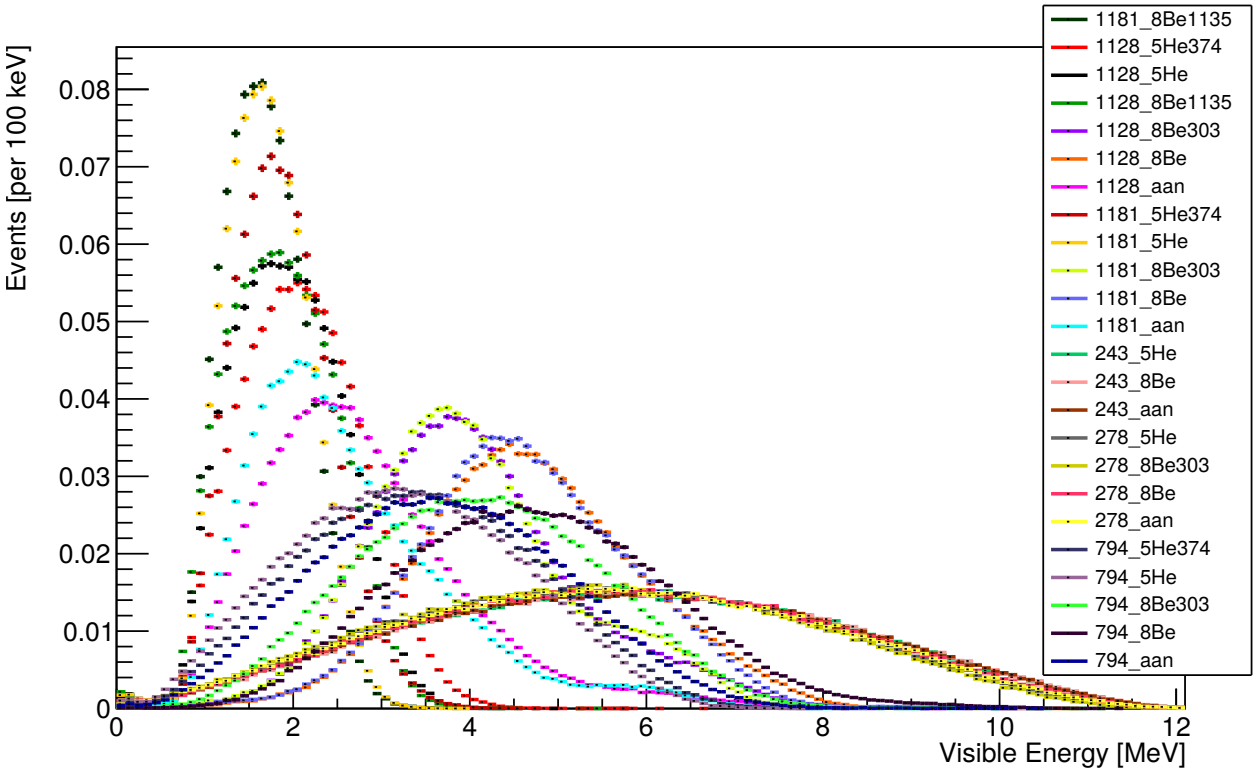

Figure 8.5 – Example of detected branches for ^{9}Li decays after they have been generated with the Spectrum Generator in the whole detection volume, run through DOGS, and selected with Gd analysis cuts. Each spectrum represents a complete decay path. $1181 - {}^{5}$He stands for the decay to the 11.81 MeV level in ^{9}Be and its subsequent decay into ^{5}He, which, in turn, splits into an α and a neutron. The errors on all spectra are purely statistical.

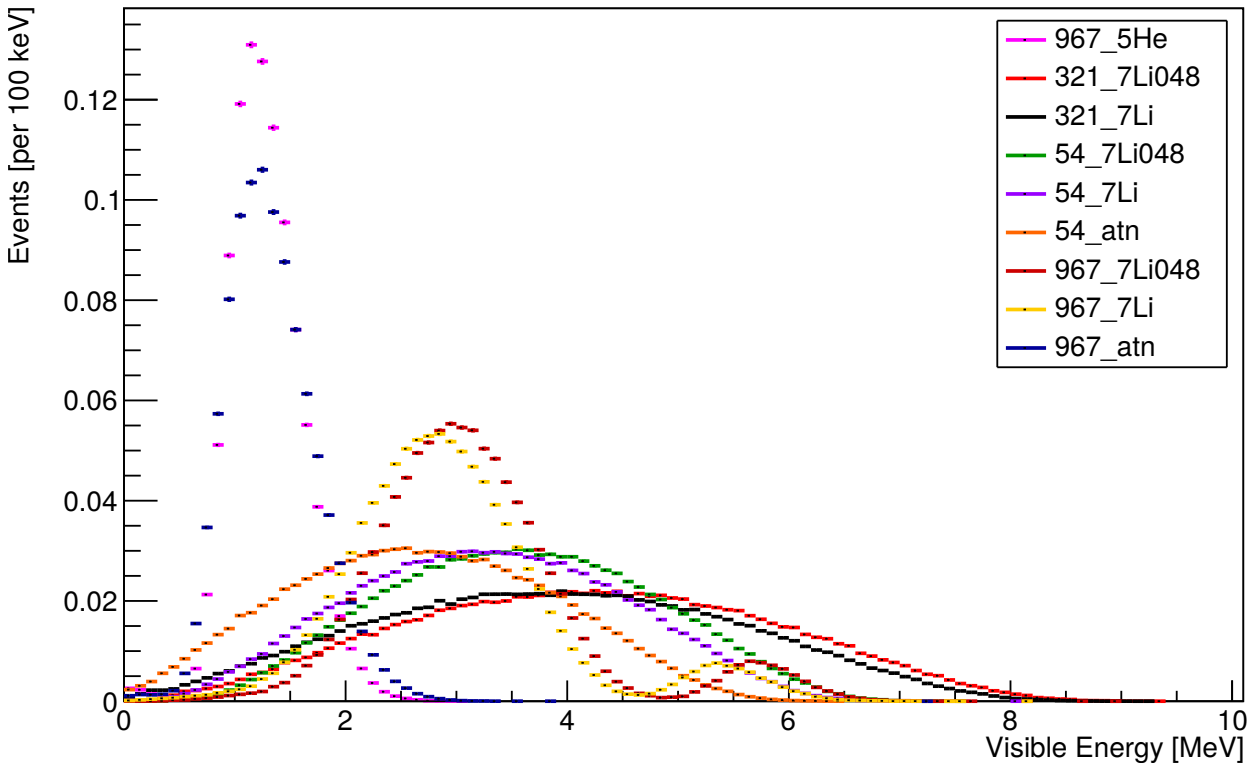

Figure 8.6 – Example of detected branches for ^{8}He decays after they have been generated with the Spectrum Generator in the whole detection volume, run through DOGS, and selected with analysis Gd cuts. Each spectrum represents a complete decay path. $967 - atn$ stands for the decay to the 9.67 MeV level in ^{8}Li and its subsequent decay into an α, a triton, and a neutron. The errors on all spectra are purely statistical.

8.2 Reconstructing a mean spectrum with errors

This part eventually leads us to the last tool introduced in the simulation chain (cf. Fig 8.1), the Covariance Tool. Like the Spectrum Generator, the Covariance Tool, was written by the author in `C++11`. This program allows us to estimate the correlations between the different energy bins of the mean spectrum, which we obtain by varying the unknown branching ratios appearing in the decay tree considered.

The trees are passed as inputs to the Covariance Tool, under the form of XML files. The latter are then read by way of the Property Tree library from Boost [131], aimed at storing and reading arbitrarily deeply nested trees of values. The features of the XML markup language most certainly allow for a straightforward representation of the nuclear decay trees.

8.2.1 Vary that which thou canst not set

8.2.1.1 Ratio modelling

In the literature, there are two types of data regarding the branching ratios of the cosmogenic decay trees, the data that were obtained from fits to measured spectra and published with a symmetric error assumed to be Gaussian (see [125] or [132] for a more detailed view), and the data that are missing. The former type of data mostly concerns the beta decays. The latter type of data overwhelmingly concerns what we referred to as "strong decays" in 7.1.1.

The modelling of the ratios that were measured and qualified by a Gaussian error is unambiguous: they may be randomly pulled in normal distributions centred on the mean value of the ratio and with a standard deviation equal to the error provided in the literature [125, 78]. As for the missing data, they were allowed to vary uniformly in the physically acceptable range. This amounts to picking uniformly distributed random numbers $r \in [\![0;1]\!]$, and rescaling them afterwards so that the sum of all the ratios that describe the connection to a node of the tree is equal to 1. The Gaussian-distributed numbers need rescaling as well for every configuration selected.

The uncertainty carried by the broadness of the levels fed by the β-decays is already included in the raw spectra database (see 7.2.3).

8.2.1.2 Weighting the detected spectra database

For every new set of branching ratios, one obtains a new candidate for the cosmogenic spectrum studied. A given energy spectrum is retrieved by rebuilding the cosmogenic decay tree within the Covariance Tool, and plugging in the set of branching ratios just pulled.

The reconstruction procedure is exemplified by Figure 8.7 for the ^{8}He case. Firstly, the branching ratios for each state decaying through the strong interaction are picked as uniformly distributed numbers (unless otherwise known) and used to weight the corresponding detected spectra of the database. Consequently, for each state decaying through the strong interaction, this provides us with a "sub-spectrum". Secondly, every thus obtained "sub-spectrum", is weighted according to the branching ratio from the literature feeding the level.

Let us consider two successive layers of decay modes as in Figure 8.7. We denote by $i \in \mathcal{B}$ the different states accessible via β-decay from the ground state of the cosmogenic isotope, and by $j \in \mathcal{S}_i$ the subsequent strong decay modes, producing a final spectrum $\mathbf{X}_j^i$[3], as retrieved from the database. Let β_i and γ_j^i stand for the β and strong branching ratios, dynamically rescaled for every iteration and set picked. Therefore, a candidate spectrum $\mathbf{X}$, accounting for the decay paths aforementioned can be expressed as the linear combination

$$\mathbf{X} = \sum_{i \in \mathcal{B}} \beta_i \sum_{j \in \mathcal{S}_i} \gamma_j^i \, \mathbf{X}_j^i \, . \tag{8.1}$$

[3] $\mathbf{X}_j^i$ is itself a vector whose coefficients are the bin contents of the corresponding spectrum.

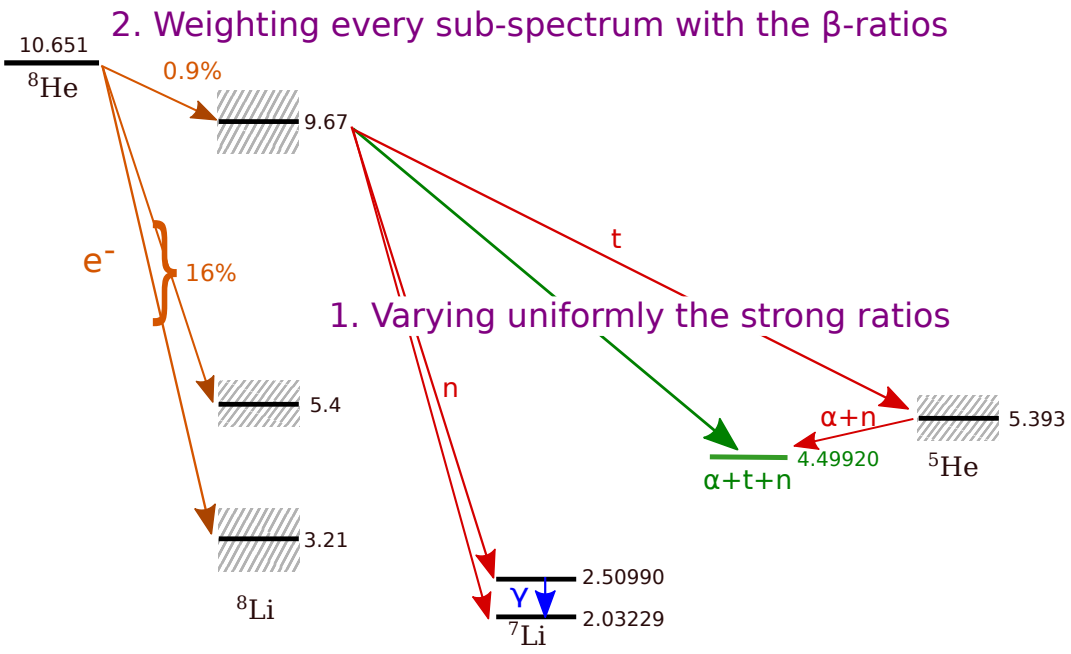

Figure 8.7 – Schematic describing the variation of the branching ratios in the case of the ^{8}He decay tree.

8.2.2 Updating a covariance matrix

The bin contents of each new spectrum – produced as described in 8.2.1.2 – can be seen as a realisation of a multivariate random variable $\mathbf{X} = (X_1, \ldots, X_{nb})$, with nb the number of bins used, and X_i the content of the i-th bin. In order to gain possession of a good estimation of the covariance matrix for $\mathbf{X}$, we ought to observe as many realisations as possible.

8.2.2.1 Recurrence relation for the covariance matrix estimator

Sample mean

To avoid saving millions of observed values of $\mathbf{X}$, the most has to be made of recurrence relations. Obtaining the mean vector is straightforward, a $(n+1)$-th realisation $\mathbf{x}^{(n+1)}$ of the random variable $\mathbf{X}$ updates the sample mean $\overline{\mathbf{x}}_n$ as shown below

$$\overline{\mathbf{x}}_{n+1} = \frac{1}{n+1}\mathbf{x}^{(n+1)} + \frac{n}{n+1}\,\overline{\mathbf{x}}_n\,. \tag{8.2}$$

The mean spectrum from (8.2) undoubtedly coincides with the one obtained when using the mean value of the distribution of each ratio in the decay tree, i.e. when setting all the known ratios to the mean of their Gaussian and most of the strong ratios to a common number rescaled for every-node. Indeed, each bin content within $\mathbf{X}$ is but a linear combination of the random ratios from one layer, as (8.1) exemplifies, and the expectation of $\mathbf{X}$ reads

$$\mathbb{E}\left[\mathbf{X}\right] = \sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{S}_i}\mathbb{E}\left[\beta_i\right]\mathbb{E}\left[\gamma_j^i\right]\;\mathbf{X}_j^i\,, \tag{8.3}$$

since $\mathbb{E}\left[\beta_i \gamma_j^i\right] = \mathbb{E}\left[\beta_i\right] \mathbb{E}\left[\gamma_j^i\right]$, considering they are pulled independently, and $\mathbf{X}_j^i$ is uniquely determined in the database.

Sample covariance matrix

Notwithstanding the reproducibility of the sample mean with plain considerations, exploring the physically accessible space for the ratios provides us with the correlations between the contents X_i of the vector $\mathbf{X}$.

The correlations are contained within the covariance matrix V. After $n+1 > 1$ observations of $\mathbf{X}$, the Bessel correction, which ensures that the estimator $\widehat{V}$ of V is unbiased, introduces the following expression

$$\widehat{V}_{n+1} = \frac{1}{n} \sum_{k=1}^{n+1} \left(\mathbf{x}^{(k)} - \overline{\mathbf{x}}_{n+1}\right)^{\mathrm{t}} \left(\mathbf{x}^{(k)} - \overline{\mathbf{x}}_{n+1}\right). \tag{8.4}$$

In order to find a recurrence relation, it is apposite to develop (8.4) into

$$\widehat{V}_{n+1} = \frac{1}{n} \sum_{k=1}^{n+1} \mathbf{x}^{(k)\,\mathrm{t}} \mathbf{x}^{(k)} - \frac{n+1}{n}\, \overline{\mathbf{x}}_{n+1}{}^{\mathrm{t}} \overline{\mathbf{x}}_{n+1}\,, \tag{8.5}$$

which calls for defining a product matrix P, whose sample value

$$\begin{aligned} P_{n+1} &= \sum_{k=1}^{n+1} \mathbf{x}^{(k)\,\mathrm{t}} \mathbf{x}^{(k)} \\ &= P_n + \mathbf{x}^{(n+1)\,\mathrm{t}} \mathbf{x}^{(n+1)} \end{aligned} \tag{8.6}$$

is effortlessly up-datable with each new realisation $\mathbf{x}^{(n+1)}$ of the bin contents $\mathbf{X}$. Eventually, the estimator of the covariance matrix after $n+1$ realisations can be re-written

$$\widehat{V}_{n+1} = \frac{1}{n} P_{n+1} - \frac{n+1}{n}\, \overline{\mathbf{x}}_{n+1}{}^{\mathrm{t}} \overline{\mathbf{x}}_{n+1}\,. \tag{8.7}$$

As a result, each term within (8.7) need only access its previous value at the n-th iteration, along with the latest spectrum defining $\mathbf{x}^{(n+1)}$.

8.2.2.2 Convergence test

At one point we have to stop pulling random ratios to generate new spectra and update the covariance matrix. Thus, a "convergence" criterion must be defined.

Clearly, there is no alleged limit to refer to in order to test a convergence in probability of the random sequence $\left(\widehat{V}_n\right)_n$. And yet, the matrix space is complete, which entails that if the sequence is Cauchy, it converges in probability (see for instance [133]). In light of the law of large numbers, by construction, the sequence should anyhow converge; accordingly, an "approximate Cauchy test" provides a worthy criterion within the Covariance Tool.

Strictly speaking, we can test that for an arbitrary small precision ϵ, we can find a number of realisations n_0 such that all subsequent matrices of the sequence are close to $\widehat{V}_{n_0}$, i.e.

$$\forall \epsilon > 0,\ \exists\, n_0 \in \mathbb{N},\ \forall k > 0,\ \left\| \widehat{V}_{n_0+k} - \widehat{V}_{n_0} \right\| < \epsilon. \tag{8.8}$$

Certainly, we cannot perform the test on an infinite number of matrices $\widehat{V}_{n_0+k}$ with $k > 0$, so a range for k has to be specified; hence the "approximate" qualifier. A queue of a few matrices is quite reasonable; a few scores are also well within the capabilities of average mobile CPU's. Repeating the procedure several times with a fixed ϵ and a n_0 of choice also ensures that the probability to find $\left\| \widehat{V}_{n_0+k} - \widehat{V}_{n_0} \right\| \geq \epsilon$ is arbitrarily small, which actually is the defining statement of "convergence in probability".

Some may argue that off-diagonal terms are customarily slow to converge, but in a finite-dimensional space, all norms are equivalent. For convenience, we employ the Hilbert-Schmidt norm, which reads for a matrix A

$$\|A\| = \sqrt{\mathrm{Tr}\left(A^* A\right)} \tag{8.9}$$

where A^* denotes the Hermitian conjugate of A.

All the computations and tests within the Covariance Tool are performed using the lightning-fast Eigen `C++` template library [134] for linear algebra. The Eigen library allows relative comparisons for the norm, so $\left\| \widehat{V}_{n_0+k} - \widehat{V}_{n_0} \right\|$ was effectively rescaled with $\left\| \widehat{V}_{n_0} \right\|$. Setting $\epsilon = 10^{-6}$ and demanding a consecutive range of 30 matrices in the Cauchy test ($k < 30$) requires about $n_0 = 3 \times 10^6$ iterations for the intricate ^{9}Li decay tree.

A graphical representation of the estimator $\widehat{V}$ of the covariance matrix of $\mathbf{X}$, after the convergence test has been passed, is displayed in Figure 8.8. The indexing of this matrix and all that will follow is such that each index corresponds to a 100 keV-wide bin, starting with a bin centre at 50 keV for the first index. The square roots of the coefficients of the diagonal of the covariance matrix $\widehat{V}_n$ can be used as estimators of the errors on the bin contents $\overline{\mathbf{x}}_n$ of the mean detected spectrum, for n large enough.

Although we will review the final matrices in depth in 8.4, we may here present the corresponding correlation matrix in its bare form as Figure 8.9. The correlation matrix is much more evocative as it smears out the amplitudes of the bin contents of the spectra and stresses the energy bins bound to one another. Please note that the correlation matrix can only be drawn for the non-zero coefficients of the covariance matrix, which implies disposing of the last bins after 12 MeV, where the detected spectra are all zero.

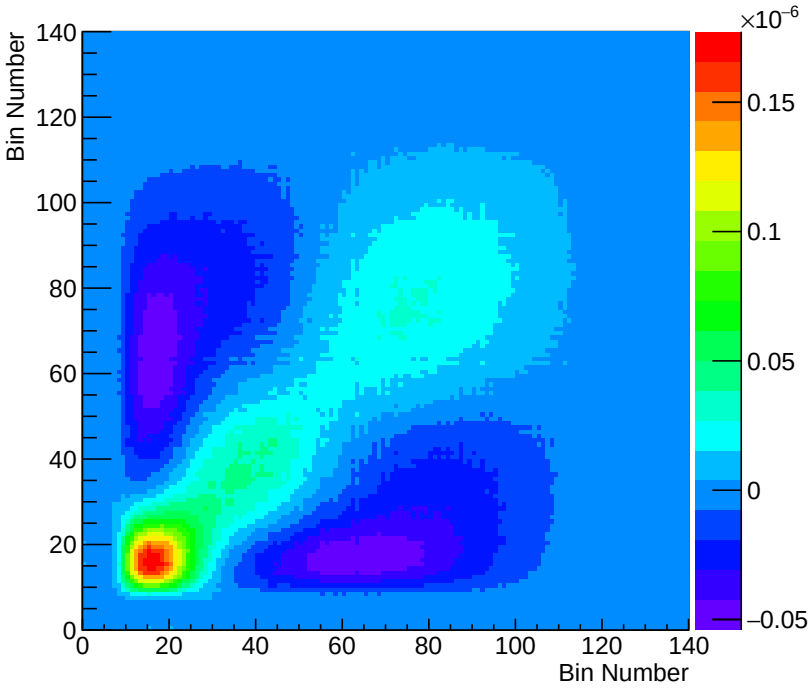

Figure 8.8 – Example of covariance matrix in the ^{9}Li case, after having varied the branching ratios of the decay tree. The matrix contents are represented for a 100 keV binning of the visible energy.

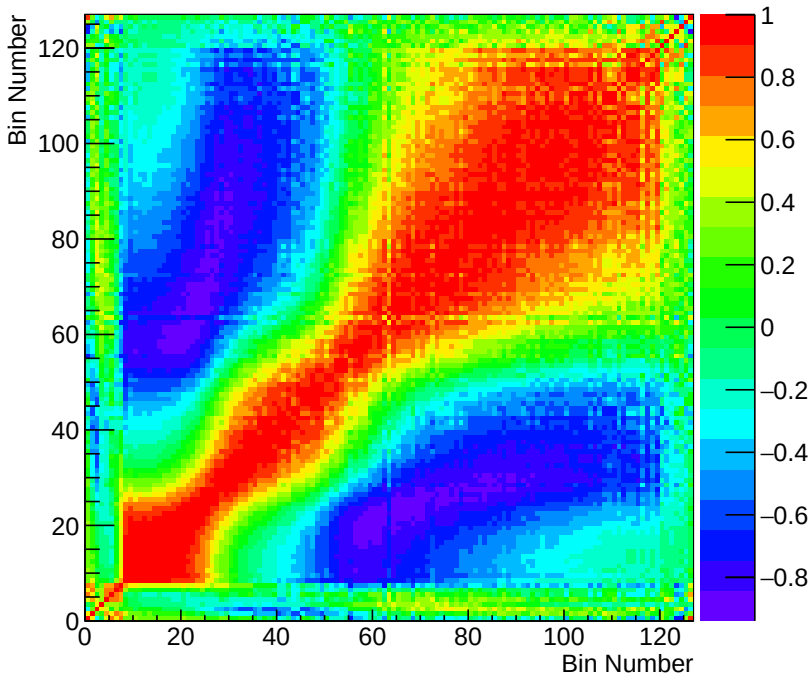

Figure 8.9 – Example of correlation matrix in the ^{9}Li case, after having varied the branching ratios of the decay tree. The matrix contents are represented for a 100 keV binning of the visible energy.

8.2.3 Weak magnetism uncertainty

The poor knowledge of the branching ratios translates into uncertainties across most energy regions of the spectra, although these uncertainties are usually moderated by competitive constraints on the β-ratios. And yet, weak magnetism [135], which embodies the interaction of the β particle with the magnetic moment of the decaying nucleus, is another non-negligible source of uncertainty on the final spectra.

The uncertainty carried by the weak magnetism corrections is assessed to be a linear function of the energy range and proportional to the bin contents of the β-spectrum. It is expressed in the following way

$$\sigma_i = S a_i X_i, \tag{8.10}$$

where S is the slope of the error, a_i the i-th bin centre of the histogram representing the $\beta-$spectrum, and X_i the content of the i-th bin. The value of S was set to $S = 0.005\,\mathrm{MeV}^{-1}$ in the Covariance Tool, thereby acknowledging a 5% uncertainty on the bin contents at 10 MeV. The (8.10) expression represents a fairly conservative 100% percent error on the value of the weak magnetism correction to β-decay for allowed transitions, as a comparison to the linearisation of the correction found in [136] demonstrates.

One should normally apply the corresponding uncertainty to a database of raw β-spectra. As discussed earlier in 8.1.2.1, picking raw spectra within some error bars and running them

through DOGS, hundreds of times, does not fit into a reasonable amount of time. To put it differently, we somehow have to incorporate the weak magnetism uncertainty after detection. Neutrons, α particles, tritons and γ's contribute to each bin content in the mean cosmogenic spectra, but these are not affected by the weak magnetism uncertainty. Therefore, to apply the uncertainty with accuracy, we would at least need to know the average e^- kinetic energy for each bin content in the whole cosmogenic spectrum. This would, again, require lengthy simulations with e^- only, and in truth, a non-trivial treatment for a bin centre at 5 MeV, in the whole spectrum may correspond to one at 3 MeV in the e^- distribution.

At the moment, one has to bear in mind that the uncertainty grows with energy, and that so does the proportion of e^- in the cosmogenic spectrum. Furthermore, e^- produce a visible energy close to their actual kinetic energy, whereas α's, for instance, are detected with less than a tenth of their actual energy, as stressed in 5.1.3.3 . Accordingly, the detected energy has its strongest contribution from e^- and setting the weak magnetism uncertainty onto the mean cosmogenic spectrum is but a safe bet. In other words, this treatment is quite conservative. As a conclusion, after the convergence has been obtained in 8.2.2, we can simply add to the covariance matrix V thus obtained a second matrix W, which represents the systematic uncertainty on weak magnetism and reads

$$W = \sigma\, {}^{t}\sigma, \tag{8.11}$$

with σ the vector of errors whose contents were defined in (8.10). As is clear from (8.11), the errors are fully correlated for all bins. A visual representation of such a matrix is presented in Figure 8.10. By and large, it echoes to the shape of the ^{9}Li spectrum on its diagonal.

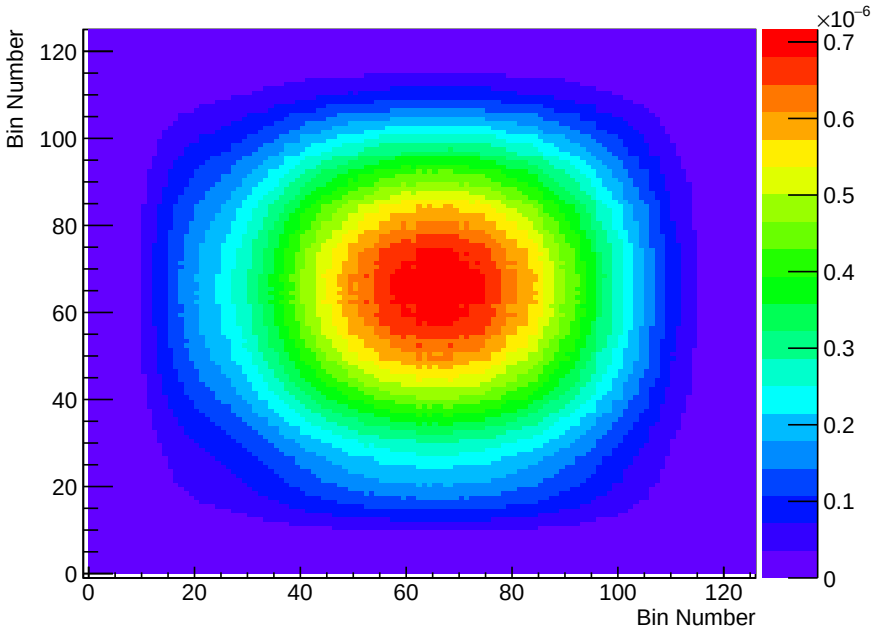

Figure 8.10 – Example of weak magnetism covariance matrix in the ^{9}Li case. The matrix contents are represented for a 100 keV binning of the visible energy.

8.3 Combining spectra from disjoint sets

The Covariance Tool, which has been extensively reviewed in section 8.2, cannot only estimate the correlations between different energy regions within the same spectrum, but it can also compute the covariance matrix between different versions of the spectrum of one cosmogenic isotope. It can also compute the correlations between the spectra of two different cosmogenic isotopes, but that needlessly entails waiting for the program to output rounded zeroes, which it surely does.

One interest of computing the covariance between two versions of one spectrum lies in the ability to provide an accurate covariance matrix, in the case of a fit to an extended data set, which adds up both data spectra. In the following, what we have just introduced as "versions", will simply refer to the analysis cuts applied to the spectra database (cf. 8.1.3), which will leave us with two versions: the "Gd" one, and the "H" one. Please note that if the "Gd" and "H" sets are disjoint, since the delayed energy bounds are chosen so, it only means that the data spectra can be painlessly added, but that is not quite the case for the Monte-Carlo spectra. To put it simply, if the value of the ratio feeding one level is wrong or the weak magnetism correction is off, any spectrum based on this branch, or this nucleus, will carry an error that we have to assess, regardless of further event selection.

8.3.1 Ratio uncertainties

The underlying purpose of this part is to evaluate the extent to which the knowledge of the branching ratios impacts either version of the spectrum, which gives direct information on the correlation between the two spectra.

To this intent, at each iteration n, one must pick the same physical configuration for each decay tree, i.e. the same branching ratios. If the nuclear decay trees are identical, the trees passed to the Covariance Tool are different in that their leaves point to different elements in the spectra database (Gd-analysed spectra or H-analysed spectra).

The multivariate random variables which will represent the bin contents of the two spectra are denoted $\mathbf{X}$ and $\mathbf{Y}$. The sample means $\overline{\mathbf{x}}_n$ and $\overline{\mathbf{y}}_n$ can be updated for each set of the branching ratios as in (8.2). The product matrix P_n, is naturally updated by the $(n+1)$-th realisations $\mathbf{x}^{(n+1)}$ and $\mathbf{y}^{(n+1)}$ of $\mathbf{X}$ and $\mathbf{Y}$, respectively, according to

$$\begin{aligned} P_{n+1} &= \sum_{k=1}^{n+1} \mathbf{x}^{(k)\ \mathrm{t}}\mathbf{y}^{(k)} \\ &= P_n + \mathbf{x}^{(n+1)\ \mathrm{t}}\mathbf{y}^{(n+1)} \end{aligned} \tag{8.12}$$

In agreement with 8.2.2, the estimator $\widehat{C}$ of the covariance matrix between the two spectra, after $n+1$ realisations, reads

$$\widehat{C}_{n+1} = \frac{1}{n} P_{n+1} - \frac{n+1}{n}\, \overline{\mathbf{x}}_{n+1}\ {}^{\mathrm{t}}\overline{\mathbf{y}}_{n+1}. \tag{8.13}$$

As when computing the covariance between the bins of the same spectrum, $\widehat{C}$ must pass the convergence test from 8.2.2.2 for the Covariance Tool to proceed with the next step, i.e. handle the weak magnetism uncertainty.

8.3.2 Weak magnetism uncertainty

With respect to the weak magnetism uncertainty, there is little additional work to do on top of sub-section 8.2.3.

We simply wish to combine spectra of the same isotope, thus, the ${\sigma_{\mathbf{X}}}^{(i)} = Sa_iX_i$ and ${\sigma_{\mathbf{Y}}}^{(j)} = Sb_jY_j$ uncertainties from (8.10) – where we have extended the notation to the j-th bin centre b_j of the histogram whose contents embody $\mathbf{Y}$ – hold for $\mathbf{X}$ and $\mathbf{Y}$, respectively.

The weak magnetism correction is implemented as a slope, and if this slope is incorrect for the studied isotope, then its value is as much incorrect for $\mathbf{X}$ as it is for $\mathbf{Y}$. In accordance, any bin from one spectrum is fully correlated to any bin from the other spectrum. It follows that the matrix to add to the covariance matrix answering for the ratio uncertainties $\widehat{C}$ is

$$W = \sigma_{\mathbf{X}}\,{}^{\mathrm{t}}\sigma_{\mathbf{Y}}, \tag{8.14}$$

with $\sigma_{\mathbf{X}}$ and $\sigma_{\mathbf{Y}}$ the vector of errors whose contents were defined in the paragraph above.

8.3.3 Linear combination of spectra

Let us here underline that $\widehat{C}$ exhibits no apparent symmetry whereas the covariance matrix $\widehat{V}$ from 8.2.2, which represents the variance of a spectrum, or the covariance between its different energy regions, is symmetric. Accordingly, when building a linear combination of the spectra represented by $\mathbf{X}$ and $\mathbf{Y}$, weighted by α and β, respectively, the matrix below must be used

$$\widehat{V}\left(\alpha\mathbf{X} + \beta\mathbf{Y}\right) = \alpha^2\,\widehat{V}\left(\mathbf{X}\right) + \alpha\beta\left(\widehat{C}\left(\mathbf{X},\mathbf{Y}\right) + {}^{\mathrm{t}}\widehat{C}\left(\mathbf{X},\mathbf{Y}\right)\right) + \beta^2\,\widehat{V}\left(\mathbf{Y}\right), \tag{8.15}$$

with $\widehat{V}\left(\alpha\mathbf{X} + \beta\mathbf{Y}\right)$ the covariance matrix of the linear combination, $\widehat{V}\left(\mathbf{X}\right)$ that of the first spectrum as obtained in 8.2.2.1, $\widehat{V}\left(\mathbf{Y}\right)$ that of the second, and $\widehat{C}\left(\mathbf{X},\mathbf{Y}\right)$ the covariance between the two spectra, as computed in 8.3.1. In (8.15), one is free to change all the $\widehat{V}$'s and $\widehat{C}$'s with their counterparts including the weak magnetism matrices. In our particular case, α and β are the number of Gd and H events the combined spectrum must account for. Do bear in mind that these events are decays of the very same isotope (α and β have nothing to do with the relative abundance of one cosmogenic isotope or the other), whose neutrons have been captured by either a Gd or a H nucleus, exclusively.

Notwithstanding the diligence of the linear combination (8.15), and slightly anticipating the results section, let us shed the light on the fact that all the matrices appearing in (8.15) are acutely similar. The similarity between all these matrices is scarcely surprising when comparing the Gd and H databases, they themselves show small differences between each

other (recall that Figure 8.4 was a worst-case scenario), as a consequence, the ratios drive about the same energy regions in each spectrum, thereby yielding the same correlations between the different energy bins.

8.4 Results and spectra comparison for ^{8}He and ^{9}Li

8.4.1 Mean spectra

The mean spectrum for each isotope – regardless of the analysis cuts that were applied – is simply the sample mean of the vector of bin contents as defined in (8.2), at the step at which the convergence criterion from 8.2.2.2 was fulfilled. The sample means are indexed by numbers corresponding to bins in visible energy, which implies that they may be compared directly to the data spectra, to validate the shapes of either, and possibly extract the contribution of each cosmogenic isotope, an effort left for 9.3.2.

8.4.1.1 Gd analysis

The so-called "^{8}He mean spectrum" may be found in Figure 8.11. This mean spectrum corresponds to raw events generated in the whole detection volume, from which a spectra database was built with the Gd cuts as found in Table 8.1.

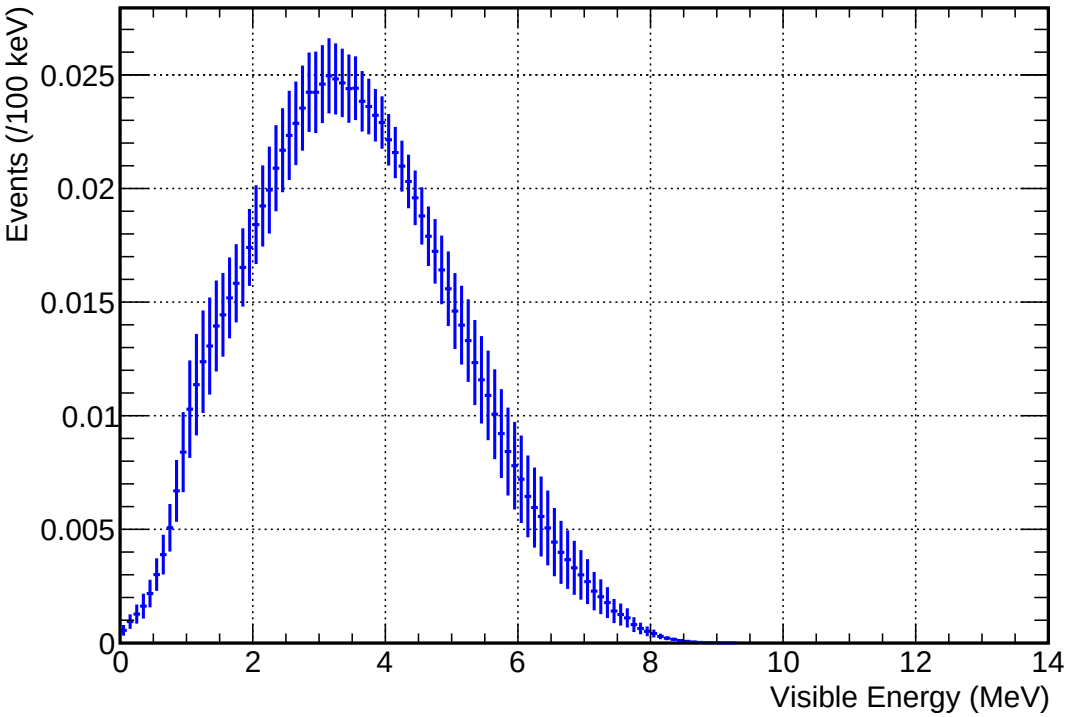

Figure 8.11 – ^{8}He mean spectrum computed by means of a spectra database analysed with the Gd cuts; the raw events have been generated in the whole detection volume.

The maximum energy released in the ^{8}He disintegrations comes from the decay chain $^{8}\mathrm{He}\,^{g.s.} \rightarrow {}^{8}\mathrm{Li}\,^{3.21} \rightarrow {}^{7}\mathrm{Li}\,^{2.51} \rightarrow {}^{7}\mathrm{Li}\,^{g.s.}$. Indeed, the latter distributes most of its total energy to an electron and a photon, and its "endpoint" is higher than the modes going through ^{5}He (you may also refer to the spectra database from Figure 8.6). The energy released in this chain amounts to 8.62 MeV, hence the slow fade after this value of the visible energy. As a side note, this upper bound is not dissimilar to the IBD endpoint.

With regard to errors, particularly at high energies, the uncertainty mainly comes from the weak magnetism covariance matrix. The errors are still far from negligible at lower energies; they result from the poorly determined values of the β-ratios feeding the $^{8}\mathrm{Li}\,^{5.4}$ and $^{8}\mathrm{Li}\,^{3.21}$ levels. The uncertainties for these two are 50% (see [125, 78]); bearing in mind that due to "quenching" (refer to 8.2.3), electrons matter most in the visible spectrum, the sizeable length of the error bars hardly comes as a surprise.

In Figure 8.12, the ^{9}Li mean spectrum extends to higher energies; unlike for ^{8}He, the available energy is constant for all decay paths and it totals 12.04 MeV, well-past the prompt endpoint for $\overline{\nu_e}$ interactions, thereby providing a strong handle on the ^{9}Li background. The uncertainty on the bin contents increases with energy, as dictated by the weak magnetism estimation; at lower energies, it is not outrageously large, by virtue of the greater constraints on the β-ratios for the ^{9}Li decays, in comparison to that of ^{8}He. More details about the emerging structures, in relation to the correlation matrices, will be given in 8.4.2.

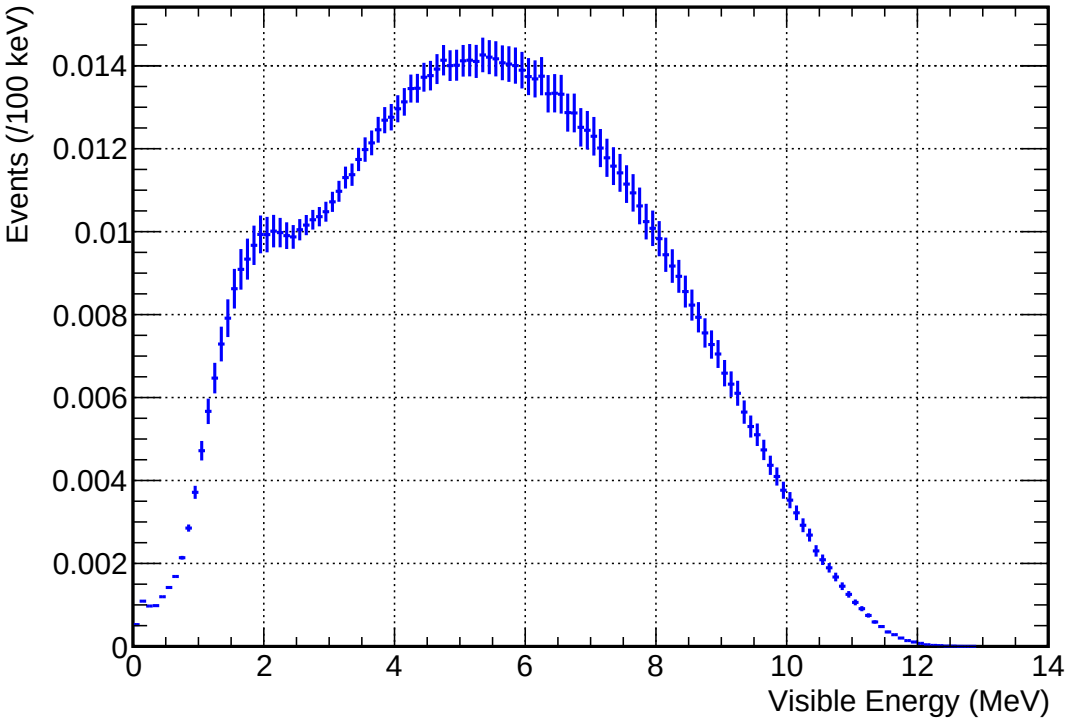

Figure 8.12 – ^{9}Li mean spectrum computed by means of a spectra database analysed with the Gd cuts; the raw events have been generated in the whole detection volume.

8.4.1.2 H analysis

In this paragraph, the mean spectra are given after the H cuts have been applied onto the corresponding database, in accordance with Table 8.1. The Hydrogen ^{8}He mean spectrum can be found in Figure 8.13 and the ^{9}Li one in Figure 8.14.

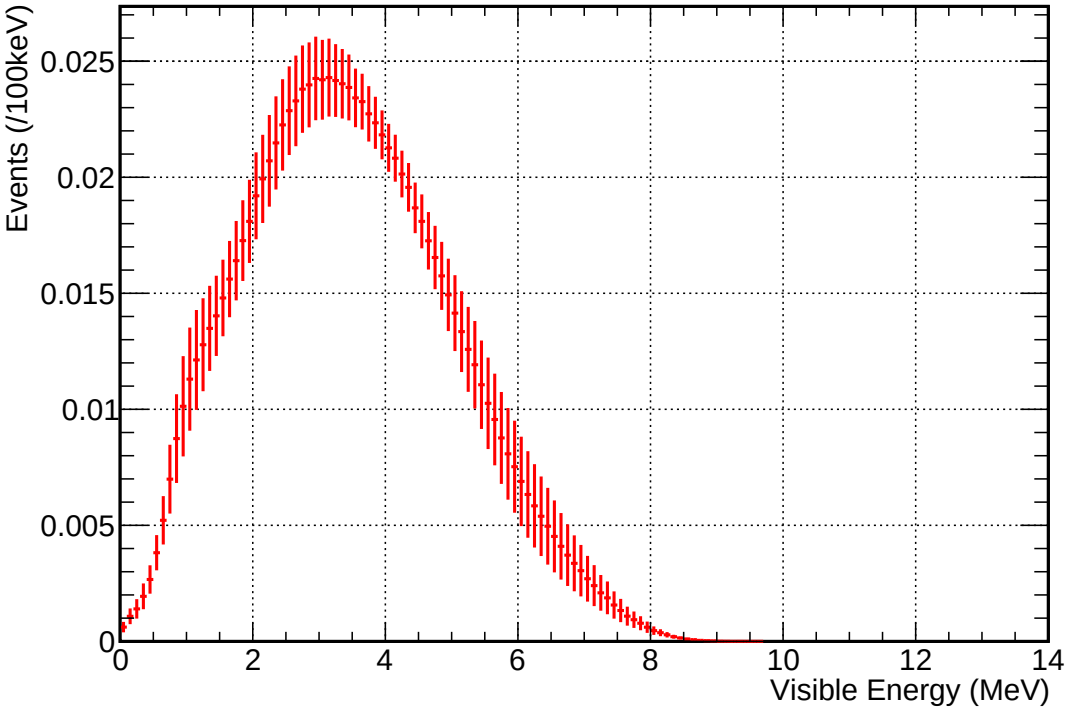

Figure 8.13 – ^{8}He mean spectrum computed by means of a spectra database analysed with the H cuts; the raw events have been generated in the whole detection volume.

Unsurprisingly, the fast-neutron-dominated modes forming a small minority – as discussed in 8.1.3.2 – Figure 8.13 and Figure 8.11 look alike to the naked eye. Global and more quantitative comparison results follow.

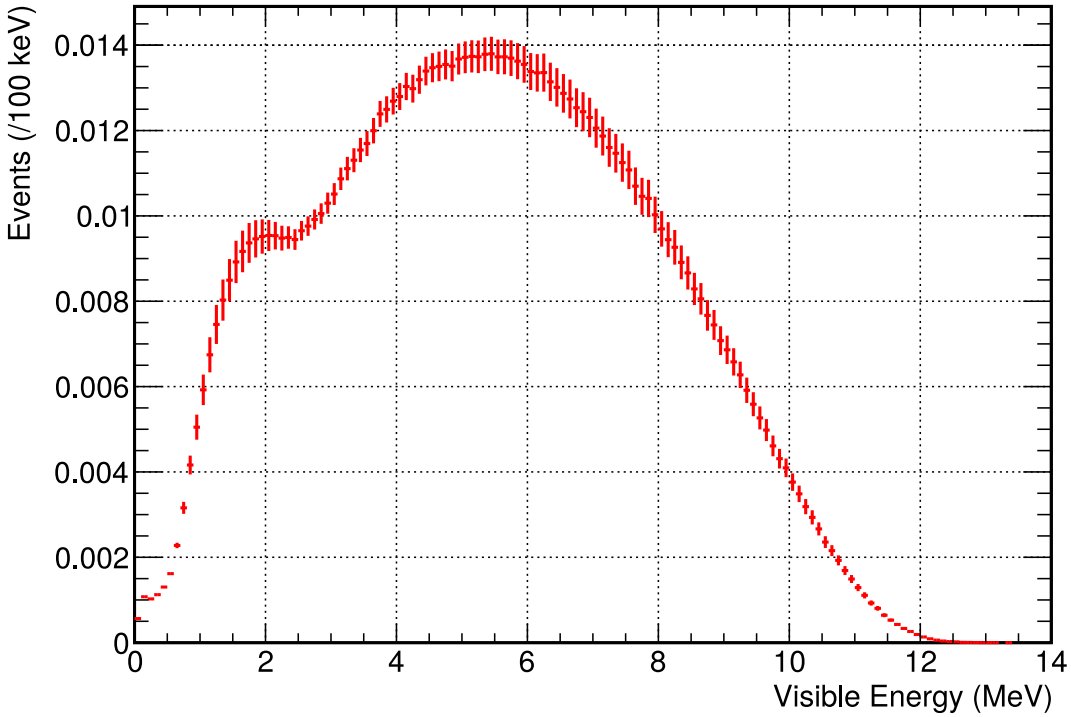

Figure 8.14 – ^{9}Li mean spectrum computed by means of a spectra database analysed with the H cuts; the raw events have been generated in the whole detection volume.

The histogram from Figure 8.14 may not be most talkative when compared to that of Figure 8.12, thus, let us draw the ratio of the H mean spectrum to the Gd mean spectrum and plot it in Figure 8.15.

As expected from the remarks in 8.1.3, the H spectrum has a higher bin content at lower energies (below 2 MeV). With the 1 MeV prompt energy threshold in the H-related analyses, including the Gd++ one, the effective difference between both spectra are further curbed. Since both spectra are normalised to one, the effect conversely spreads at middle energies (between 4 MeV and 8 MeV), though it may not be so obvious without magnifying the graph. At higher energies (above 10 MeV), the bin contents are small and the relative differences are larger than meets the eye when looking at the bare histograms.

Regarding the errors when computing the ratio as an illustrative example, no heed was paid to the correlations between the H and Gd bin contents. In view of the comments from 8.3.3, the correlations are substantial, which should translate into a ratio with smaller errors than what is shown in Figure 8.15. Indeed, and up to minute energy shifts, should the i-th bin content in the H spectrum be a little larger, the i-th bin in the Gd one would also, in most cases, follow the upwards trend, thereby shrinking the uncertainty on the ratio.

The ratio of the ^{8}He spectra has not been included here for brevity; it exhibits features similar to the ^{9}Li case, its largest 38% deviation being also located below the analysis threshold, at around 750 keV.

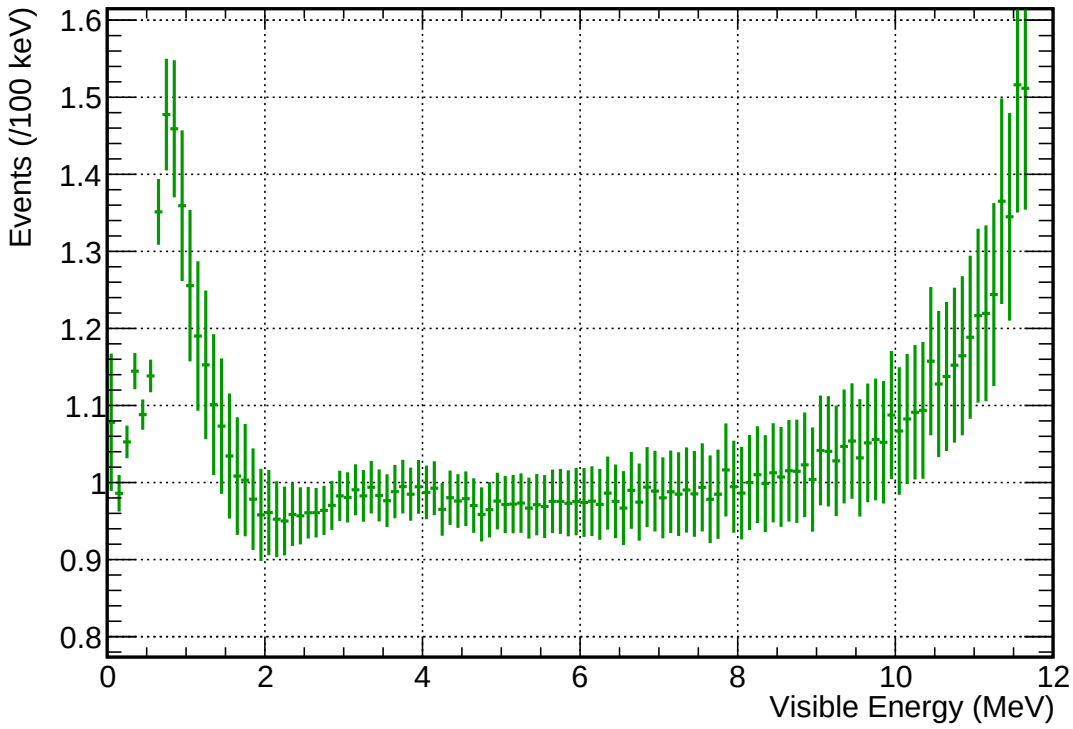

Figure 8.15 – Ratio of the ^{9}Li mean spectrum obtained from the H cuts to the one obtained from the Gd cuts.

8.4.2 Covariance and correlation matrices

The very purpose of the Covariance Tool is – as its name suggests – to produce the covariance matrices for the cosmogenic spectra. In sections 8.2.2 and 8.2.3, we exhibited examples of covariance and correlation matrices obtained when either varying the unknown branching ratios, or accounting for the weak magnetism uncertainty. In this section, the matrices have been turned into a form most useful for defining a χ^2 with the data: they include both systematics.

8.4.2.1 Gd analysis

The reader may find in Figure 8.16 a graphical representation of the ^{8}He covariance matrix, and in Figure 8.17, that of the ^{8}He correlation matrix. Keeping this ordering, the ^{9}Li covariance and correlation matrices are shown in Figure 8.18 and Figure 8.19. All of them have been analysed with the Gd cuts as presented in Table 8.1.

Thereafter, the correlation matrices are in the limelight by reason of their intrinsic normalisation, which helps to draw physical conclusions. Nonetheless, this ought not to deter us from the fact that covariance matrices must be used when comparing data to Monte-Carlo, for the scale must not be overlooked.

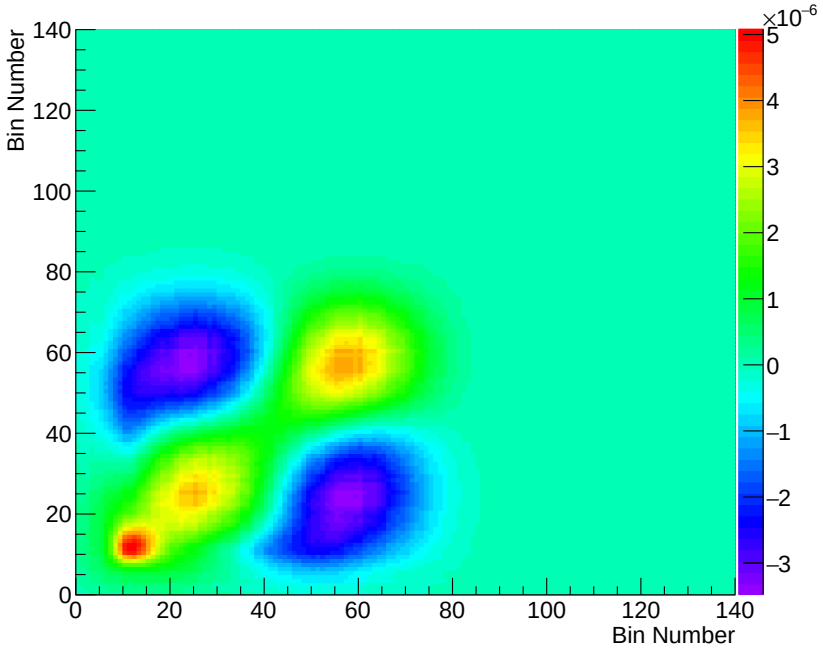

Figure 8.16 – ^{8}He covariance matrix computed by means of a spectra database analysed with the Gd cuts; the raw events have been generated in the whole detection volume. The matrix contents are represented for a 100 keV binning of the visible energy.

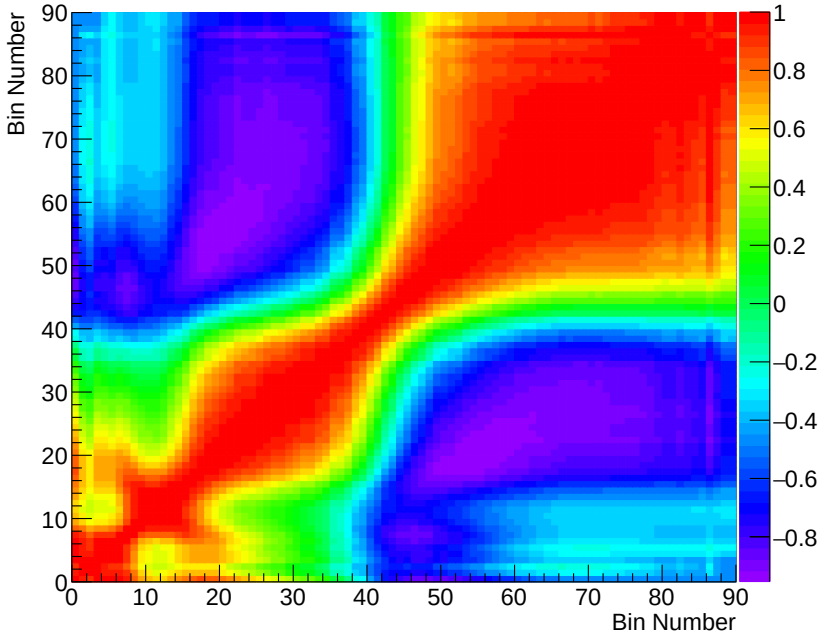

Figure 8.17 – ^{8}He correlation matrix computed by means of a spectra database analysed with the Gd cuts; the raw events have been generated in the whole detection volume. The matrix contents are represented for a 100 keV binning of the visible energy.

In Figure 8.17, the shift from a positive correlation to a negative one confirms that the maximum of the ^{8}He mean spectrum is a little below 4 MeV, recalling that the mean spectrum is normalised. One can also spot an inflection point a little above 1 MeV, whose presence could already be observed by a good eyesight on the mean spectrum from Figure 8.11. This point mainly comes from the low energy decays of the 9.67 MeV level in ^{8}Li – namely the threebody break-up and the ^{5}He channel – who are the only candidates for this energy region, as is testified by the ^{8}He spectra database in Figure 8.6. To put if differently, whenever the ratios driving these channels are rebalanced, all the bins of the corresponding spectra, packed into this narrow energy region, have to move together, hence the correlations across the affected energy range.

At visible energies larger than about 5 MeV, the fully correlated weak magnetism uncertainty kicks in and reddens the graphical representation. A quick glance at Figure 8.16 confirms how the errors on the bin contents of the mean ^{8}He spectrum are plummeting above 8 MeV, which goes hand in hand with the drop in the bin contents themselves. Plainly, there is little point in stressing the numerical computations into taking the inverses of the square roots of zeroes so as to multiply them by whatnots. In other words, the correlation matrices are drawn up to sensible bin numbers only.

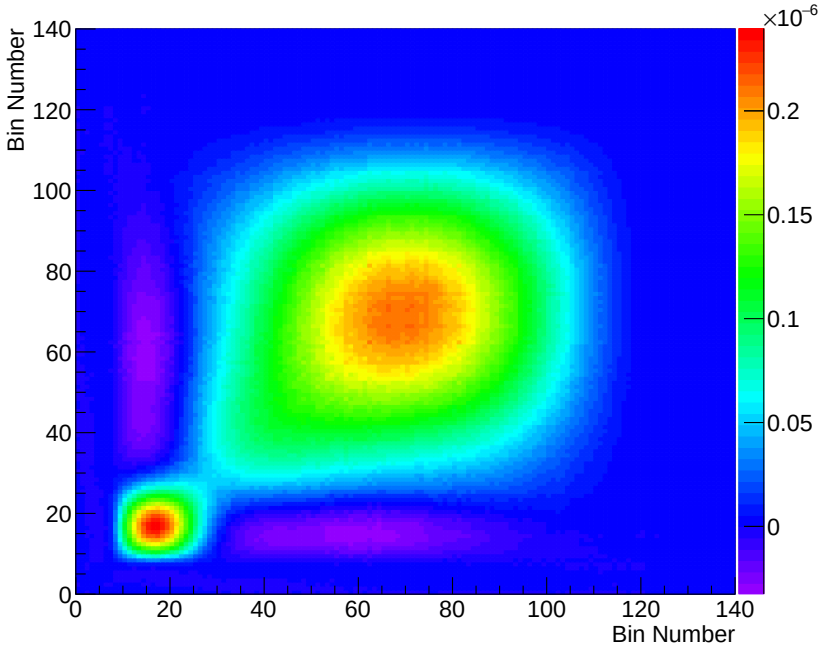

Figure 8.18 – ^{9}Li covariance matrix computed by means of a spectra database analysed with the Gd cuts; the raw events have been generated in the whole detection volume. The matrix contents are represented for a 100 keV binning of the visible energy.

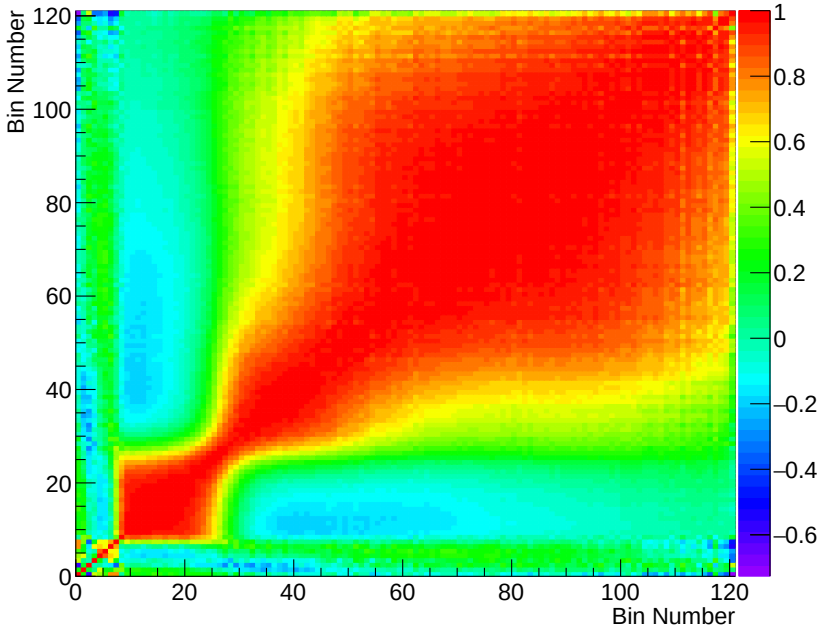

Figure 8.19 – ^{9}Li correlation matrix computed by means of a spectra database analysed with the Gd cuts; the raw events have been generated in the whole detection volume. The matrix contents are represented for a 100 keV binning of the visible energy.

When comparing the ^{9}Li correlation matrix from Figure 8.9, which is purely based on the variations of the branching ratios, with the one shown here as Figure 8.19, one effortlessly grasps how fiercely the weak magnetism uncertainty distorts the ^{9}Li correlation matrix. The latter literally skyrockets past the bins corresponding to about 5 MeV; energy regions far apart exhibit correlations between one another greater than 0.5.

Although the weak magnetism uncertainty for ^{8}He binds equally firmly the same energy regions, i.e. that above 5 MeV, the spectrum itself collapses after 8 MeV. For this reason, the correlations are not so blindingly obvious. This correlation is all the greater for ^{9}Li that its β-ratios are reasonably bound, and the weak magnetism error readily prevails under these circumstances. In short, the ^{8}He case is a little more "balanced" and each uncertainty has its share.

Somewhat reminding of the ^{8}He case, there is a tightly correlated group of bins between 1 MeV and 2 MeV in the ^{9}Li correlation matrix. A swift review of the ^{9}Li database – in Figure 8.5 – reveals that this group is rooted in the ^{5}He and ^{8}Be $^{11.35}$ channels, accessible by the decays of ^{9}Be $^{11.81}$ and ^{9}Be $^{11.28}$. Above this moderately sized correlated area, for bins corresponding to about 2.5 MeV, the local minimum from Figure 8.12 is spotted with ease.

There is also a similarly looking area centred at 4 MeV, primarily originating from the decays of ^{9}Be $^{11.81}$ and ^{9}Be $^{11.28}$ through the ^{8}Be $^{g.s.}$ and ^{8}Be $^{3.03}$ channels; the squared structure is however asymmetrically distorted by the increasing weak magnetism correlations.

Last but no least, one may have spotted the somewhat sharp green lines on the ^{9}Li correlation matrix, at around 800 keV. These lines cut the pattern below the 2.5 MeV local

minimum described above; they mirror the abrupt fall of the low-energy spectra of the ^{9}Be $^{11.81}$ and ^{9}Be $^{11.28}$ channels. For ^{8}He, the decays of the well-fed 5.4 MeV level in ^{8}Li – which release tritons, whose quenching is substantial – ensure a non-negligible bin content below 1 MeV.

8.4.2.2 H analysis

The features of the covariance and correlation matrices for the H analysis are extremely similar to that of the Gd analysis, a point which could already be noted on the mean spectra of the previous section, at least as far as the diagonals are concerned. Therefore, Figures 8.21 and 8.23 will not be commented much.

The prime characteristic of the H matrices is to have slightly downwards shifted structures, in agreement with the larger quenching and spill effects in the GC. The spectra of the H databases (not displayed in this document, but Figure 8.4 provides the essentials) are sometimes pushed 100 keV or 200 keV down the energy line, and so are the patterns in the matrices. A case in point is the green line edge presented in 8.4.2.1. Standing at around 800 keV on the ^{9}Li Gd correlation matrix, it marks roughly 600 keV in Figure 8.23, in prefect agreement with the spectra from Figure 8.4.

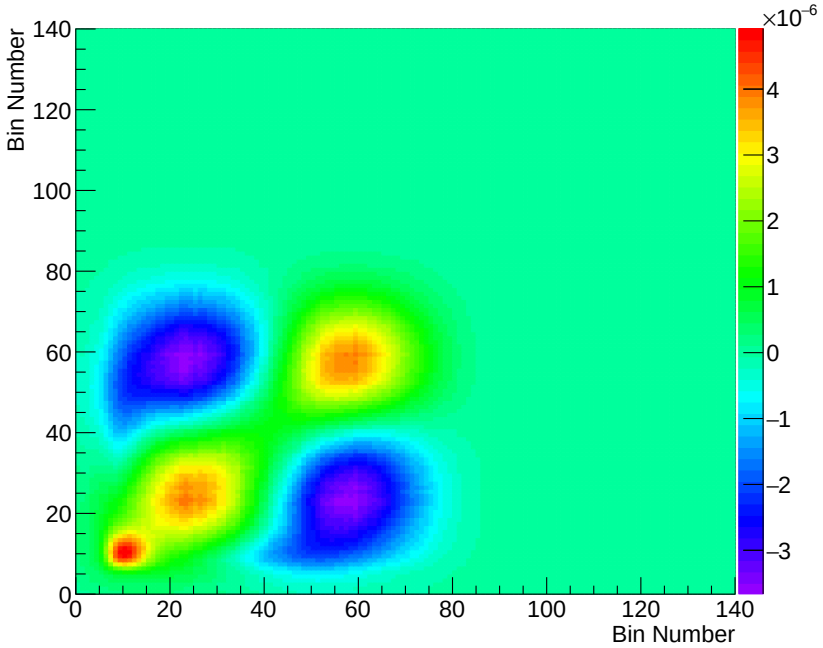

Figure 8.20 – ^{8}He covariance matrix computed by means of a spectra database analysed with the H cuts; the raw events have been generated in the whole detection volume. The matrix contents are represented for a 100 keV binning of the visible energy.

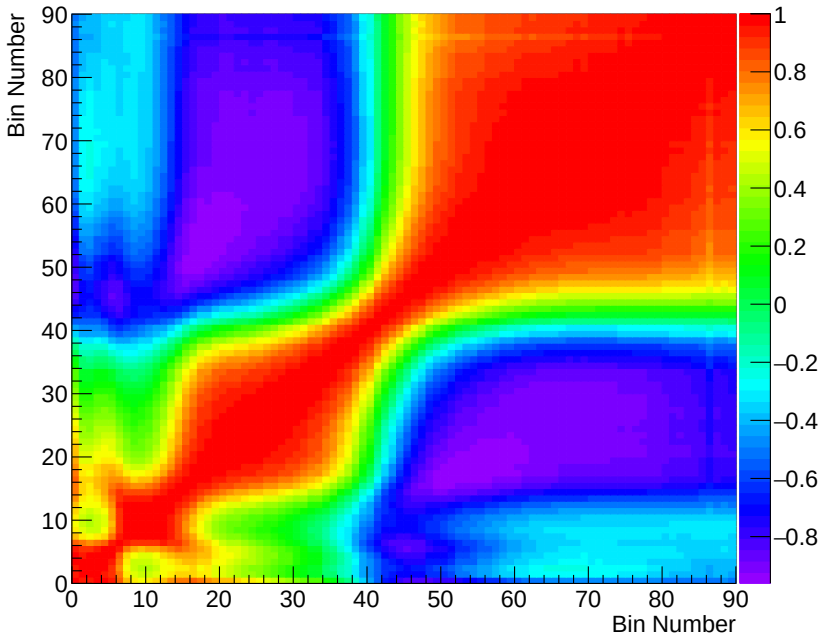

Figure 8.21 – ^{8}He correlation matrix computed by means of a spectra database analysed with the H cuts; the raw events have been generated in the whole detection volume. The matrix contents are represented for a 100 keV binning of the visible energy.

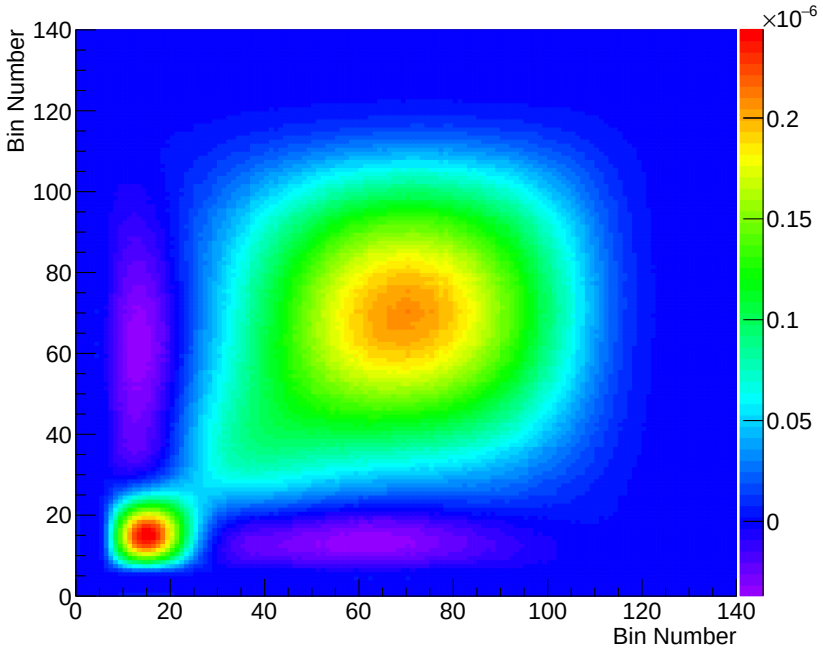

Figure 8.22 – ^{9}Li covariance matrix computed by means of a spectra database analysed with the H cuts; the raw events have been generated in the whole detection volume. The matrix contents are represented for a 100 keV binning of the visible energy.

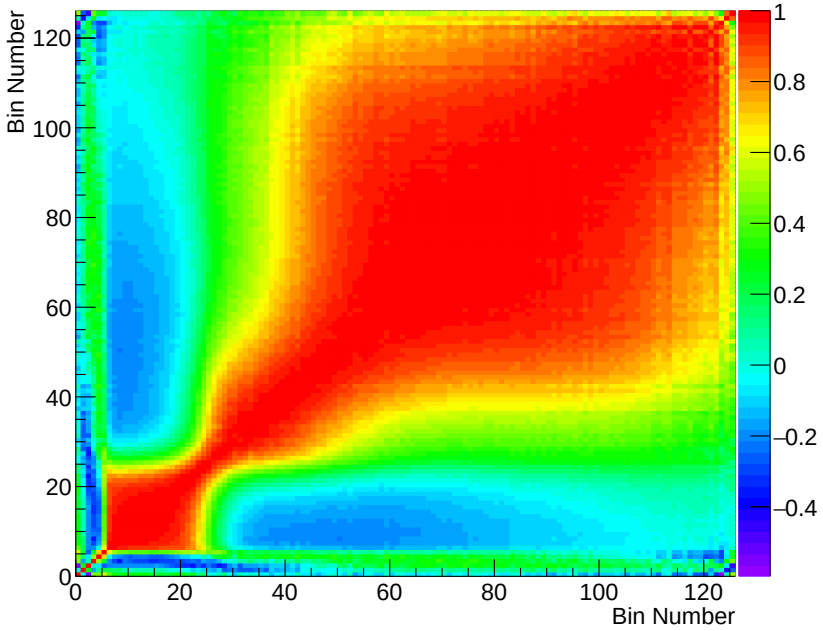

Figure 8.23 – ^{9}Li correlation matrix computed by means of a spectra database analysed with the H cuts; the raw events have been generated in the whole detection volume. The matrix contents are represented for a 100 keV binning of the visible energy.

8.4.3 Conclusion

Throughout this chapter and Chapter 7, several tools used in the production of the simulated cosmogenic spectra for the Double Chooz experiment have been presented. These tools are generic and could be run for simulating the decays of other nuclei with properties akin to that of the cosmogenic isotopes.

The techniques at play in these tools are purely based on Monte-Carlo methods; they allow the production of spectra with associated errors. The errors are estimated whilst studying the correlations between the different energy regions of the spectra. The latter show through when varying one of the largest source of uncertainty: the poorly constrained branching ratios within the decay trees of these nuclei. Weak magnetism, which accounts for a good length of the error bars, is also included in the uncertainty treatment. The widths of the states, which are commonly large, have seen their modelling discussed within these chapters and implemented in the generator.

Two cosmogenic isotopes – ^{8}He and ^{9}Li – which are beta delayed-neutron emitters, have had their decays generated with the simulation chain and analysed as data with various energy and position cuts. Their spectra and corresponding covariance and correlation matrices have been displayed and extensively reviewed in this document.

Chapter 9

Extracting spectra from data

Vetoing the decays of cosmogenic isotopes is not so much important for it reduces the background contamination, but rather because the vetoed events form an independent set from which the cosmogenic data spectra can be extracted. The larger the statistical sample, the greater the background constraints on the contamination remaining in the oscillation fit. Particularly so with a multi-detector set-up, which crushed most systematics, the authenticity of the shape of the cosmogenic spectra is in the limelight. Therefore, a comparison to Monte-Carlo predictions is most appropriate.

9.1 Cosmogenic veto

9.1.1 Targeted events

For the latest single-detector publication [60], an active veto aimed at tagging the decays of cosmogenic isotopes was developed [83]. The vetoing procedure has neither notion of βn nor β decays; by and large, it bases its assumptions on the showering capabilities and proximity of past μ-tracks to a valid trigger. Consequently, it naturally identifies any sort of signal that can be correlated to μ-tracks crossing the ID, which includes the β-decays of ^{12}B, for instance[1].

Selecting pairs as in 6.1.2.1 and applying all the backgrounds vetoes (cf. 6.2) but the cosmogenic one ensures that mostly βn-emitters are tagged. There is, a priori, no reason that the cosmogenic veto should hunt ^{9}Li more efficiently than ^{8}He. In all honesty, it might even track ^{8}He more easily, for the production of the latter ought to correspond to even more energetic μ's. Nonetheless, as suggested by 7.2.1.1, whereas ^{9}Li βn-decays in 51% of the cases, ^{8}He chooses a similar path only 17% of the time. Moreover, the larger mass difference between ^{8}He and ^{12}C, from which it is produced, when compared to that between ^{9}Li and

[1]The cosmogenic ^{12}B is also produced by μ-spallation on ^{12}C, in copious quantities, as the identicalness of their mass numbers hints at.

^{12}C, alludes to an even smaller production of ^{8}He. In a nutshell, ^{8}He is a small component of the cosmogenic spectra, which we shall strive to evaluate.

9.1.2 Formulation

9.1.2.1 Likelihood

As is customary in Double Chooz, the official name of the cosmogenic veto, i.e. "Lithium likelihood (veto)", is a bit of a misnomer[2].

Let us consider a μ which deposited more than 100 MeV in the ID, up to a few ^{9}Li lifetimes $\tau_{\mathrm{Li}} = 257\,\mathrm{ms}$ (larger than $\tau_{\mathrm{He}} = 172\,\mathrm{ms}$) before a prompt trigger, and whose track was reconstructed at d (cf. 5.2.4.3) from the prompt event. We denote by $n \in \mathbb{N}$ the number of neutron-captures detected within the 1 ms following the μ passage, be it on Gd or H. Then, the likelihood of the prompt trigger to represent the decay of a cosmogenic isotope reads

$$\mathcal{L}\left(co \mid n, d\right) = f_{co}(n, d), \tag{9.1}$$

with f_{co} the joint probability density of n and d for cosmogenics.

The (9.1) formulation should be quite reminiscent of (5.28); it states that the likelihood of the prompt, to be a cosmogenic isotope, given that there exists a μ which generated n spallation neutrons at a distance d from it, within a pre-defined time window, is equal to the "probability" to have observed (n, d) were the prompt indeed the sought-after background.

9.1.2.2 Posterior probability

Bayes' legacy

Now that we have defined what the likelihood is, let us focus on the quantity actually utilised by the cosmogenic veto, i.e. the posterior probability.

The latter goes by such a fancy name on account of the priors it must rely on, as a genuine Bayesian method. By reason of Bayes' theorem, the posterior probability $P\left(co \mid n, d\right)$ for an event to be a cosmogenic decay, given that there is a μ at d which produced n neutrons, can be expressed in terms of $f_{co}(n, d)$ and the prior $P(co)$. This prior is the initial degree of belief that a random prompt-μ pair represents the production and decay of a cosmogenic isotope. With $f(n, d)$ the joint probability density of n and d, regardless of the prompt type, we have

$$P\left(co \mid n, d\right) f(n, d) = f_{co}(n, d) P(co)\,. \tag{9.2}$$

The queer mix of probabilities and densities is rooted in the discrete character of the event type: either the event is a true cosmogenic decay or it is a valid $\overline{\nu_e}$ interaction, albeit

[2]The reader might have noticed that although most quantities computed for most analyses are not likelihoods, appending the magic word seems to bestow some sort of ethereal validity upon the methods.

unfortunate enough to happen near an energetic μ-track. Labelling these ill-fortuned events *acc*, we can split the joint probability following

$$f(n,d) = f_{co}(n,d)P(co) + f_{acc}(n,d)P(acc)\,, \tag{9.3}$$

with $P(acc) = 1 - P(co)$, the prior on the proportion of accidental coincidences between prompts (dominated by $\overline{\nu_e}$'s) and μ's. As a result, the posterior probability for a μ-candidate pair is

$$P\left(co \mid n,d\right) = \frac{f_{co}(n,d)P(co)}{f_{co}(n,d)P(co) + f_{acc}(n,d)P(acc)}\,. \tag{9.4}$$

Comparing (9.4) and (9.1), it is clear that although the latter appears in the former, they are not alike, all the more so with the newly-added notion of priors.

Maximal posterior probability

It is worth stressing that a given prompt can be associated to nearly infinitely many posterior probabilities, one for each (n,d), i.e. for each μ. A 700 ms-long time windows was deemed optimal to look for μ's initiating cosmogenic production [83]. Surely, a 700 ms-long window offset 1 ms before the prompt (see after-μ veto in 6.1.1.2) encompasses 93% of the dominant ^{9}Li decays, without letting in too much noise. Consequently, we will often have in mind the maximal posterior probability value that a μ can propose to a given trigger, i.e.

$$P_{max}\left(co\right) = \max_{\mu\in\mathcal{W}_p} P\left(co \mid n_\mu, d_{p-\mu}\right)\,, \tag{9.5}$$

where $\mathcal{W}_p$ denotes the window of μ's preceding the trigger p. A graphical illustration of the maximisation of $P\left(co \mid n,d\right)$ is presented in Figure 9.1.

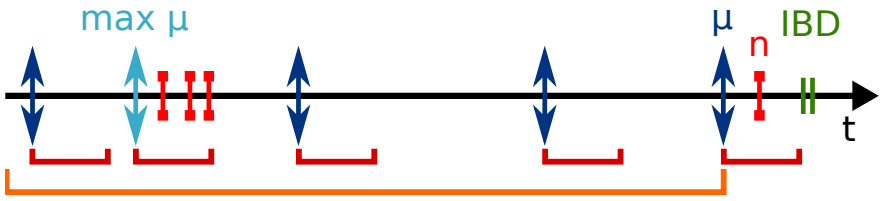

Figure 9.1 – Selection of the μ yielding the maximum posterior probability that the prompt from the IBD candidate (green) is a cosmogenic decay. The $n = 3$ neutrons (red) generated by the maximal μ (light blue), certainly help maximising $P\left(co \mid n,d\right)$. Note that μ's closer in time do not have more chances to be selected. Due to the after-μ veto (red), $\mathcal{W}_p$ (orange) must be offset.

9.1.2.3 Usage

Making the most of the so-called 700 ms "on-time" window $\mathcal{W}_p$ before each prompt p, we may veto the cosmogenic background. Prompt events satisfying

$$P_{max}\left(co\right) > P^{th}_{max}\,, \tag{9.6}$$

with $P^{th}_{max} \in [0;1]$ the cut value to be determined in 9.2.3, may be removed from the $\overline{\nu_e}$ candidates, and qualified as cosmogenic βn-decays.

Undoubtedly, some IBD inefficiency is expected from the (9.6) cut, and applying the aforementioned criterion will not yield cosmogenic data spectra directly comparable to the Monte-Carlo predictions from Chapter 8. The performance of the veto is tightly bound to the extent to which the probability densities of the cosmogenics and accidentals differ, which unequivocally drives our ability to identify the most likely parent μ for a cosmogenic isotope.

9.1.3 Priors

9.1.3.1 Cosmogenic prior

The posterior probability, as its name suggests, relies on the initial knowledge we can claim about the chances to observe a cosmogenic decay or an accidental coincidence between a $\overline{\nu_e}$ and a μ. With regard to μ's, there is no distinction to make at that point on the neutron multiplicity n; thus, we simply need the μ-rate in the ID, according to the definition we gave in 6.1.1.1, to compute the rate of accidental coincidences. The rates of μ satisfying $E^{vis} > 100\,\text{MeV}$ are shown in Table 9.1. The slight decrease in the FDII rate, with respect to FDI's, is consistent with the disagreements in the energy thresholds mentioned in 5.2.4.3.

Detector	r_μ (Hz)
FDI	10.21
FDII	10.05
ND	52.76

Table 9.1 – ID Muon rates for the three Double Chooz detectors.

The total number N_{co} of prompt-μ pairs, expected to correspond to cosmogenic decays, is merely the product of the independently measured rate r_{co} (cf. values from Chapter 10)[3] and the live-time t_{live} (cf. Table 6.4). Regarding accidentals, they are dominated by the $\overline{\nu_e}$ signal and the μ-rate. The rate of IBD candidates r_{cand} is taken from Table 6.4. On average, for each $\overline{\nu_e}$ candidate, there are $r_\mu\, t_{\mathcal{W}_p}$ muons in its window $\mathcal{W}_p$ of length $t_{\mathcal{W}_p}$, and hence, accidental coincidences. Therefore, the total number of accidentals is quite accurately approximated by

$$N_{acc} = r_{cand}\, t_{live}\, r_\mu\, t_{\mathcal{W}_p}\,, \tag{9.7}$$

and the cosmogenic prior reads

$$P(co) = \frac{N_{co}}{N_{co} + N_{acc}} = \frac{r_{co}}{r_{co} + r_{cand}\, r_\mu\, t_{\mathcal{W}_p}}\,. \tag{9.8}$$

[3]Table 6.3 and Table 6.7 only show the remaining background contamination, the total rate can be obtained by adding the vetoed rate.

9.1.3.2 Prior ratios

Dividing the numerator and denominator of (9.4) by $P(acc)$, we can actually write

$$P\left(co \mid n, d\right) = \frac{\pi_r \; f_{co}(n,d)}{\pi_r \; f_{co}(n,d) + f_{acc}(n,d)} \,, \tag{9.9}$$

with the prior ratio π_r being defined as

$$\pi_r = \frac{P(co)}{P(acc)} = \frac{r_{co}}{r_{cand} \; r_\mu \; t_{\mathcal{W}_p}} \,. \tag{9.10}$$

The prior ratios still in use in the code are given in Table 9.2. As can be noticed, a common value was deemed sufficient for both detectors. The cosmogenic rate does not exactly scale with the μ-rate, in accordance with the lower mean μ energy at shallower depths. Above all, the ND $\overline{\nu_e}$ rate is $\sim$ 7.3 that of FD. Consequently, when striving to hunt down cosmogenic decays at ND, however numerous they may be, the odds of success are much worse than when poaching ^{9}Li at FD. Current knowledge – using the background estimations from the oscillation fit, along with the latest veto efficiencies – would rather favour $\pi_r^{ND} = 6.2 \times 10^{-4}$, $\pi_r^{FDI} = 7.0 \times 10^{-3}$ and $\pi_r^{FDII} = 6.6 \times 10^{-3}$; it neither changes the argument nor the performance, as we will discuss.

Detector	π_r
FD	7.7×10^{-3}
ND	5.5×10^{-4}

Table 9.2 – Prior ratios used in the cosmogenic veto.

The prior ratios π_r thus obtained are also valid for the Gd++ analyses. Indeed, when performing the Gd++ analysis, for each detector $\delta \in \{FD, ND\}$, both the cosmogenic rates and the $\overline{\nu_e}$ rates, are multiplied by r^δ_{++}. Temporarily subtracting all background contaminations from the $\overline{\nu_e}$ candidates, we would find the IBD ratios

$$r^{FD}_{++} \simeq 2.7 \,, \tag{9.11}$$

$$r^{ND}_{++} \simeq 2.4 \,. \tag{9.12}$$

The smallness of the ND ratio arises from the longer after-μ veto for H-based analyses, which weighs all the more on the near site that its μ-rate is great.

As discussed in 7.2.4.1, the production of cosmogenics is driven by the ^{12}C densities, whereas the $\overline{\nu_e}$ cross-section is weighted by the number of H nuclei (cf. 5.1.2.1); we have yet to reach that level of nit-picking. Last but not least, since the ID already encompasses both the NT and the GC, the μ rate is identical in the Gd and Gd++ analyses, thereby vouching for identical prior ratios.

9.1.4 Reference probability densities

9.1.4.1 Independent variables

For both the cosmogenics and the accidentals, the joint density was decomposed into the product of the mass function for the neutron multiplicity and the density for the lateral distance, namely,

$$f_{acc}\,(n,d) = P^N_{acc}(n)\; f^D_{acc}(d) \tag{9.13}$$

$$f_{co}\,(n,d) = P^N_{co}(n)\; f^D_{co}(d)\,. \tag{9.14}$$

If (9.13) is indisputable – considering $\overline{\nu_e}$'s are located uniformly in the fiducial volumes, regardless of crossing-μ's – (9.14) is a little more questionable, for indeed, one would expect large μ-showers to generate cosmogenic isotopes a little farther from their tracks. In practice, comparison of the Double Chooz mean neutron path from the μ-tracks with that of KamLAND [137] favours a minimal dependency on the overburden, and hence on the μ energy. Like the KamLAND profiles, the Double Chooz lateral distance distributions – presented in 5.2.4.3 – exhibit a mean free path close to 80 cm. The distance between the cosmogenic decays and the μ-tracks presents no different properties. All the more so with the limited number of cosmogenic decays we can analyse, the assumptions laid out by (9.14) seem quite reasonable.

9.1.4.2 Distance to μ-tracks

To obtain the distance probability density for accidentals f^D_{acc}, we simply need to look at alleged $\overline{\nu_e}$'s and compute their distance to all μ-tracks. Although the cosmogenic decays represent but a few percent of the $\overline{\nu_e}$ candidates, it is no harder to dispose of their correlations by selecting only μ triggers after the candidates, or more than $2\,\mathrm{s} > 7\tau_{\mathrm{Li}}$ before them.

As regards the distribution for correlated prompt-μ pairs, we may start by plotting the distance between IBD candidates and μ's which are close in time, utilising a 700 ms on-time window, but that will not suffice. Indeed, even when mustering prompt-μ pairs near in time, the accidental contamination induced by the $\overline{\nu_e}$ background is overwhelming. In addition, for each βn-decay, all μ's, except for the actual parent, amplify the accidental background in our cosmogenic studies. However, these uncorrelated μ's have no reason to exhibit features incompatible with the accidental probability density aforementioned. Accordingly, the uncorrelated component of the on-time window can be removed by subtracting the accidental distribution, appropriately scaled.

By multiplying the number of time windows, large statistics are achievable for accidentals. Unfortunately, with a few hundreds of cosmogenic decays at hand, the on-time signal is plagued by statistical fluctuations. Instead, somewhat reminiscent of KamLAND' strategy, we can rely on the similarity of the lateral distance profiles across different cosmogenic populations, and focus on the sole β-decays of ^{12}B [83]. The latter is dominantly produced by (n,p) reactions on ^{12}C whereas ^{9}Li is produced by the shorter-range $(\pi^-, {}^3\mathrm{He})$ and the

analogous $(n, n3p)$ reactions. The charged nature of pions goes hand in hand with a smaller distance travelled from the initiating μ-track; still, the high statistics provided by the ^{12}B decays – two orders of magnitude larger than that of the βn-emitters – are tantalising. For generating the cosmogenic profile at FDI, ^{12}B was retained. Both the cosmogenic and accidental lateral distance densities can be studied in Figure 9.2.

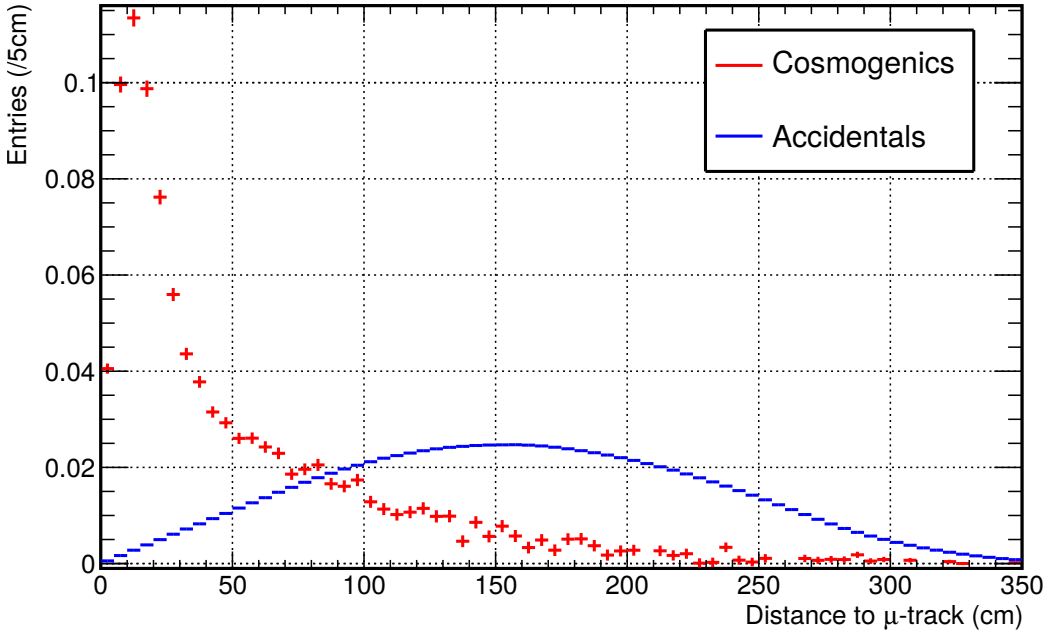

Figure 9.2 – Binned representations of the cosmogenic (red) and accidental (blue) probability densities for the distance between prompt candidates and μ-tracks.

The two distributions depicted in Figure 9.2 were obtained by means of the FIDO reconstruction algorithm (cf. 5.2.4.1). Nevertheless, as epitomised by Figure 5.15, not only does DCRecoMuHam perform as well as FIDO used to at FDI, but the ND and FD neutron lateral distance profiles also back the negligible overburden dependency previously discussed[4]. In other words, provided that ^{12}B is a faithful representative of ^{9}Li, the current f_{co}^{D} and f_{acc}^{D} may be kept without impunity for all detectors.

9.1.4.3 Neutron multiplicity

Accidentals

Like for the lateral distance, the distribution of the number of spallation neutrons generated by μ's uncorrelated to prompt triggers can be effortlessly obtained from offset windows. By virtue of the lower mean μ energy at the ND – itself rooted in its smaller overburden – the neutron multiplicity of all μ's should fall rapidly. Hence, the shape of the accidental mass

[4]Unless the ND μ-track reconstruction has room for improvements and ought to be narrower.

function at ND, presented in Figure 9.3, and compared to the FDI reference retrieved from ^{12}B β-decays, is not unexpected.

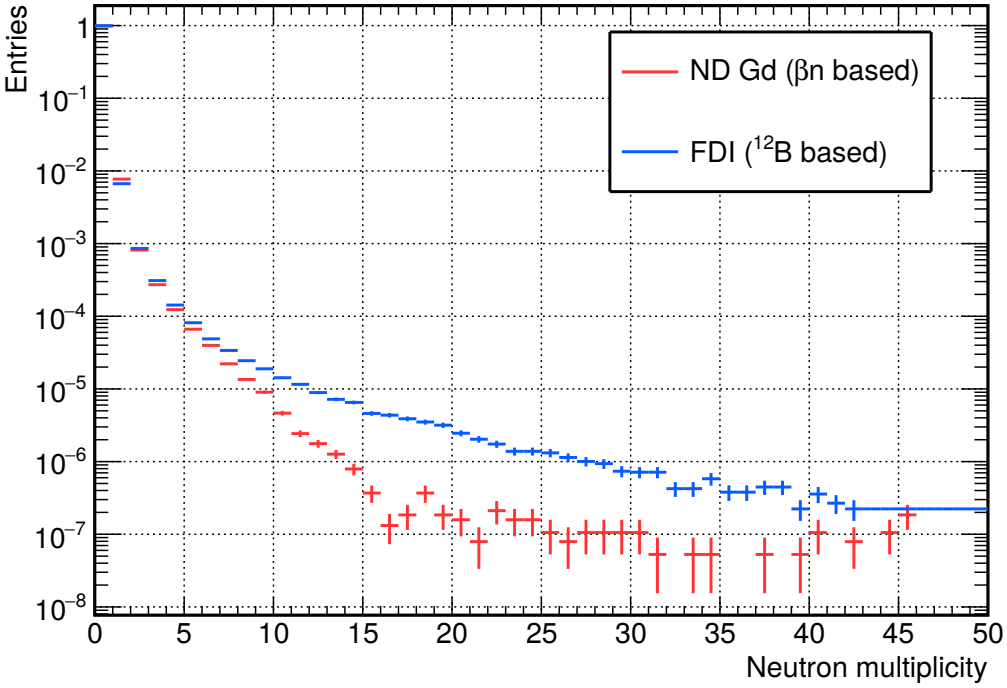

Figure 9.3 – Neutron multiplicity mass functions for μ's uncorrelated to prompt candidates (ND, red) or ^{12}B candidates (FDI, blue). Both distributions, which represent the accidental background in the cosmogenic studies at either detector, overlap within the line thickness at $n = 0$.

The lack of manpower left the ^{12}B seat vacant in the Double Chooz collaboration, and one may rightfully wonder how the neutron multiplicity mass function would look at FDI with IBD candidates, instead of ^{12}B candidates. For that matter, the accidental case is trivial; it is high time we switched to the correlated mass function.

Cosmogenics

Using a 700 ms on-time window naturally selects correlated prompt-μ pairs, but also, inevitably, many non-showering μ's which have nothing to do with βn-decays. From studies presented in Chapter 10, it appears that much more than half of the cosmogenic decays are produced by showering μ's. To put if differently, the cosmogenic signal – after which we may look in the $n = 0$ bin of the on-time window distribution – is akin to a needle in a haystack. For all but one FD data sample, the Neutrino FDI Gd++ one[5], the background-subtracted on-time mass function has a negative bin content for $n = 0$, with a worthless gigantic error

[5]The Neutrino FDI Gd++ data set boasts about 455.21 d of live-time and a $\overline{\nu_e}$ rate ~ 2.7 that of the Moriond FDI Gd set described in Table 6.4.

bar. Normalising the chosen FDI Gd++ cosmogenic mass function P_{co}^{N} between 1 and 50 to $\sum_{n=1}^{50} P_{bo}^{N}(n)$, with P_{bo}^{N} the ^{12}B reference mass function, one obtains promising results.

The graphical comparison of the prompt-based and ^{12}B-based FDI mass functions can be gazed at in Figure 9.4. In spite of its reasonable amplitude, $P_{co}^{N}(0)$ makes P_{co}^{N} close to useless for trumping P_{bo}^{N} in the cosmogenic veto[6]: at the moment, the uncertainty on the $n = 0$ key point is too large. Nonetheless, Figure 9.4 puts on the table the first factual argument vouching for the use of ^{12}B.

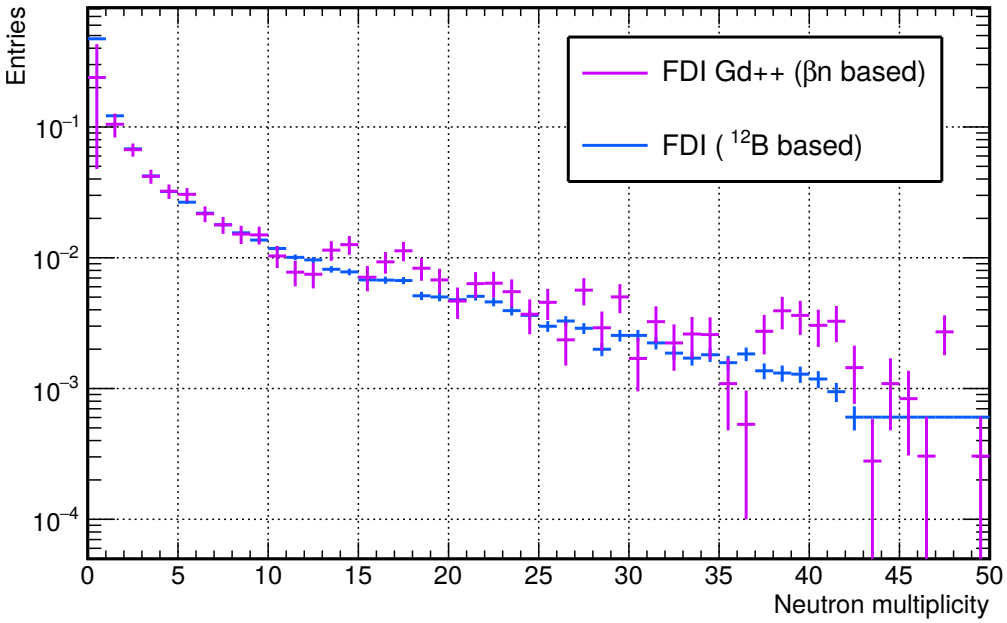

Figure 9.4 – Neutron multiplicity mass functions for μ's correlated to FDI Gd++ prompt candidates (purple) or ^{12}B candidates (blue).

What about the ND ? From 9.3, it is clear that the overall neutron multiplicity is lower at ND than at FD. On the other hand, producing cosmogenic isotopes may well demand "what it takes". It is not unlikely that lower-energy μ's at ND, rather than producing ^{9}Li and fewer fellow neutrons, simply do not produce ^{9}Li at all. Therefore, looking for cosmogenic producers, at ND, necessarily narrows down the number of μ candidates to the few ones having twins at FD. By all means, quite a few μ's are energetic enough to generate more than 10 neutrons, and βn-emitters, at the same time; these are seen more seldom at ND, and one should not expect to have identical signal distributions past a certain n value.

The FDI reference and ND mass functions can be found in Figure 9.5. As when comparing the FDI βn and ^{12}B based functions, all distributions were jointly normalised on the $1-50$ interval, which minimises the impact of statistical fluctuations in the first bin.

[6] As we are currently discussing, if the posterior probability (9.9) craves for actual βn-emitters, on account of its "co" subscripts on the distributions, the statistics have us rely on boron.

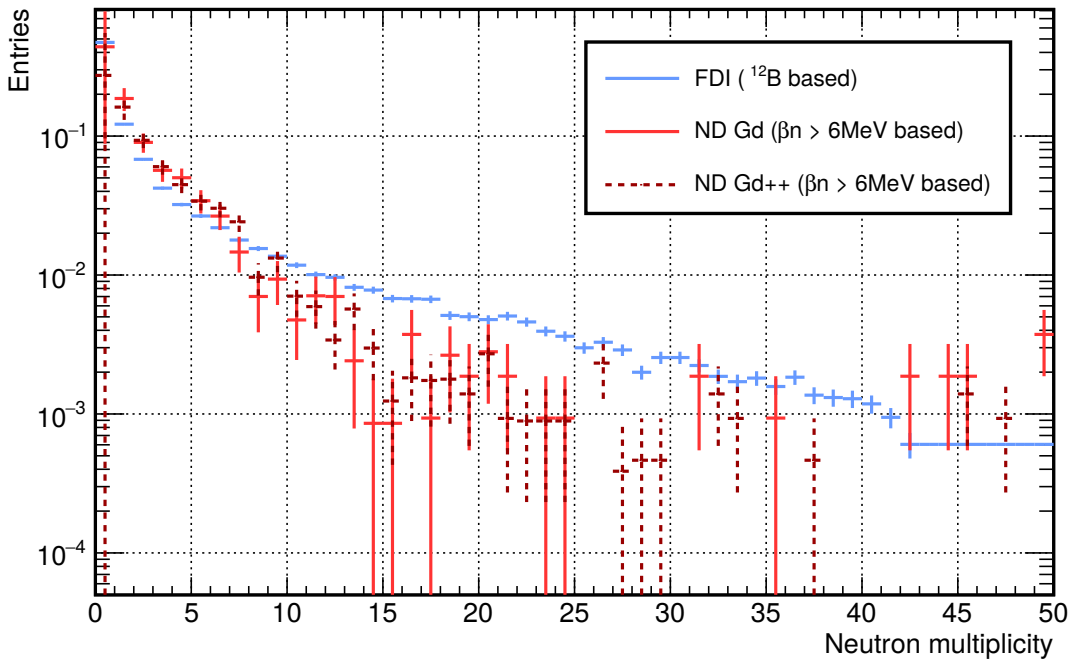

Figure 9.5 – Neutron multiplicity mass functions for μ's correlated to high-energy prompt candidates (Gd, light red; Gd++, dashed dark red) or ^{12}B candidates (blue). The ND distributions were obtained by selecting only prompts with $E^{vis} > 6\,\text{MeV}$.

Although this technique will be reviewed extensively in 9.2, it is worthwhile noting that the only way to obtain a vaguely acceptable first bin is to increase the signal over background ratio in the on-time window. To this end, particularly for selecting ^{9}Li, raising the prompt energy threshold profitably exploits the 12.04 MeV-endpoint of this cosmogenic isotope, whilst trampling on the $\overline{\nu_e}$ heap. Oddly enough, even though the Gd++ data offer a more fitting representation for $n > 0$, the sole Gd data are exemplary for $n = 0$.

Using a prompt energy cut $E^{vis} > 6\,\text{MeV}$ would also improve P^{N}_{co} in 9.4, turning its first $P^{N}_{co}(0) = 0.25 \pm 0.20$ bin into $P^{N'}_{co}(0) = 0.33 \pm 0.14$, yet not rivalling the $P^{N}_{bo}(0) = 0.47 \pm 0.01$ from ^{12}B, albeit no less enlightening.

Alleged impact of the FDI distributions at ND

Utilising the FDI accidental distribution at ND, when computing the posterior probability from (9.9), should lead to a lower vetoing efficiency of cosmogenic decays, inasmuch as the larger bin contents of the FDI mass function, for $n > 1$, would bring $P(co \mid n, d)$ down. Similarly, the lower weight of the $n = 1$ bin at FDI, could translate into more collateral damage, i.e. unfortunate vetoing of actual $\overline{\nu_e}$'s.

By reason of the limited statistical significance of the candidate-based signal, drawing conclusions for the correlated mass functions is more delicate. If the ND bin contents are systematically lower than FDI's for $n > 10$, using the FDI correlated reference would somewhat mitigate the joint use of the FDI distribution for the accidentals. Should the ND cosmogenic

mass function be comparable to the FDI one, we would gain by replacing the FDI ^{12}B-based accidental distribution by the ND one.

9.2 Vetoed signal

In 9.1.4, each prompt-μ pair lead to different variables for the μ themselves would provide changing (n, d) values. On the other hand, with regard to the energy of the prompt E_p^{vis}, be it a rest decay or a random $\overline{\nu_e}$ interaction, μ's take no part in the determination of its value. As a consequence, a naive analysis would assign to tenths of prompt-μ pairs (see the μ-rates in Table 9.1) the very same E_p^{vis}; within these pairs, only one would embody a true correlation[7]. For this reason, a one-to-one correspondence, between a prompt and a μ, is a must-have. To do so, we have at hand the best candidate: the maximum posterior probability. All the subtleties lie within the "maximum" adjective, which ought to provide us the most plausible parent μ for a cosmogenic isotope.

9.2.1 Background subtraction

9.2.1.1 Principle

The maximum posterior probability $P_{max}^{on}(co)$ for the on-time window is provided by CT (cf. 5.1.3.4), for all singles, although it only makes sense to look at its value for prompt candidates. However, the CT value is only useful to other analysers, to veto a part of the cosmogenic background in their studies.

When it comes to extracting the energy distribution of the βn-decays, the accidental contamination within the vetoed sample must be evaluated and subsequently subtracted. To this end, we must seek μ's and prompt candidates remote in time, and compute their posterior probabilities, regardless of the time that separates them. That is, we must compute the distance $d_{p-\mu}$ between them, and count the number of neutrons n_μ in the μ shower. For each prompt candidate p, we can define a 700 ms off-time window $\mathcal{W}_p^{off}$, and retain, within it, the muon μ_m yielding the maximum posterior probability

$$P_{max}^{off}(co) = P(co \mid n_{\mu_m}, d_{p-\mu_m}) = \max_{\mu \in \mathcal{W}_p^{off}} P(co \mid n_\mu, d_{p-\mu}) \tag{9.15}$$

for the uncorrelated prompt. Via this virtual parent μ_m, we can consequently assign to each E_p^{vis} one uncorrelated posterior probability.

Along with the on-time posterior probability $P_{max}^{on}(co)$, we thus have two values per prompt candidate, which will lead to two energy distributions S^{on} and S^{off}. On these

[7]Regarding the neutron multiplicity and the distance, there is also a sole correlated pair, at best, in the on-time window, which accounts for a substantial background, but at least, we get different values for each prompt-μ pair.

two probabilities, we should subsequently apply an identical threshold P^{th}_{max}, as if we were applying the veto from 9.1.2.3 to both sets. Thus, we must check whether

$$P^{on}_{max}\,(co) > P^{th}_{max} \tag{9.16}$$
$$\text{or } P^{off}_{max}\,(co) > P^{th}_{max}\,, \tag{9.17}$$

and augment the bin contents corresponding to E^{vis}_p in S^{on} or S^{off} by one unit, according to which posterior probabilities pass the threshold. Conditions (9.16) and (9.17) may be true concurrently, in which case, S^{on} and S^{off} may be simultaneously filled[8]. After this procedure has been applied to all candidates, we can claim that the cosmogenic spectrum S, exemplifying the correlations between prompts and μ's reads

$$S = S^{on} - S^{off}\,. \tag{9.18}$$

9.2.1.2 Practical constraints

A burden

The procedure from 9.2.1.1 is neat, however, it misses out some key numbers: 20% to 30% of the few hundreds of tagged events are accidental coincidences. That is, a large portion of events must be subtracted from the precious spectrum, input of the oscillation fit. The cosmogenic signal is already minute, thus, so as not to worsen its significance, the energy distribution of accidental events must be known to a satisfactory degree; one off-time window will not suffice.

As such, already using one off-time window to obtain S^{off} is inconvenient, to put it mildly. Indeed, if most analysers do not show the slightest interest in μ's, and thus, can make do with the few-MB reduced files containing singles (cf. 6.1.1), we must cope with the ponderous CT files (more than 10 GB per one-hour run), containing all events. Equally deterring, DOGS must hacked by changing the hard-coded position of its unique veto window, rebuilt, and run over all the CT files for every tweak which might be contemplated, such as changing the priors and the reference distributions. Processing all the FDI CT files for one off-time window, as used to be done, already takes several days at CCIN2P3. With the considerable μ-rate at ND, this turns doing physics into a dull long-winded undertaking, a path some tried to take nonetheless.

A boon

However, there is a way around relying on bloated files. The so-called Japanese light trees not only contain all μ's and singles, but their average file-size is a mere 100 MB. In addition, this Japanese species is only composed of plain ROOT files; in a nutshell, it merrily wanders in sunlit meadows, free of the antediluvian shackles of DOGS.

[8]Not filling them at all is analogous since the subtraction to follow will effectively remove E^{vis}_p; however, that would bias our inefficiency counting.

From these ROOT files, binary files, containing 10 s-long windows of events before each candidate prompt, were saved. To this intent, a deftly-crafted header-only serialisation[9] library was utilised. This fabulous piece of work, developed by two proficient computer scientists of the iLab of the University of Southern California, goes by the aptly-chosen name "Cereal", which also echoes its expressive logo [138]. Cereal is both extremely fast and efficient when it comes to storing bytes; above all, it allows to painlessly write any class instance to a file, and it supports virtually the whole STL, from `std::unordered_map` to all sorts of smart pointers. Unlike ROOT, Cereal does not require you to use distasteful C-style macros, hand-written `LinkDef.h` files, rootmaps, dictionaries, whenever you so much as dare trying to write a simple class, let alone a class composed of custom classes. As a matter of fact, Cereal relies on composition: any class composed of serialisable members (built-in types such as `double`, STL objects) is de facto also serialisable itself.

Beyond question, liberty always comes at a price: none of the CT algorithms are available outside DOGS, including the cosmogenic veto. As it turns out, we precisely had in mind writing a flexible version of the cosmogenic veto, using several off-time windows, and allowing command-line changes of the priors and densities, along with the addition of other vetoes. These requirements were met by conceiving – from scratch – a flexible template framework, including in its bestiary `Window<T>`, `Shower<T, K>` and `Veto<T>` template classes, to name but a few, where `T` and `K` themselves stand for any type or template class. A shower of neutrons following a `Muon`, stored as `Single`'s in `double`-precision, would for instance read `Shower<Muon<float>, Single<double>>`; we may be interested in storing a `Window` of these showers as well.

Producing these Cereal binary files for the fifteen months of Neutrino ND Gd data does not demand weeks on hundreds of cores at CCIN2P3, rather, about 15 min using 50 logical cores. Analysing the resulting 3 GB binary file for the near detector – holding ~ 50 billion μ's and 85×10^3 candidates – on the average notebook, takes about 20s. The meaning of "analyse" must be understood as "produce all the distributions, i.e. the neutron multiplicities, lateral distances, and spectra based on the demanding posterior probability for a given threshold". With far fewer μ's and candidates, all the useful FDI data hold in a 150 MB file, analysed locally in less than a second. All of a sudden, we've just got our hands on the defibrillator allowing us to reach the next stage.

9.2.1.3 Improved accidental background removal

By way of our lightweight software, multiplying the number of off-time windows by $n_{\mathcal{W}} = 12$ – without mustering all computing units at CCIN2P3 – is effortless. The first off-time window starts 10 s before the prompt, while the twelfth ends 1.6 s before it. A diagram, representing the selection of the μ's yielding the maximum posterior probability within each on and off-time window, is displayed in Figure 9.6.

[9] Serialisation consists in writing object to files or buffers and reconstructing them later in their original state.

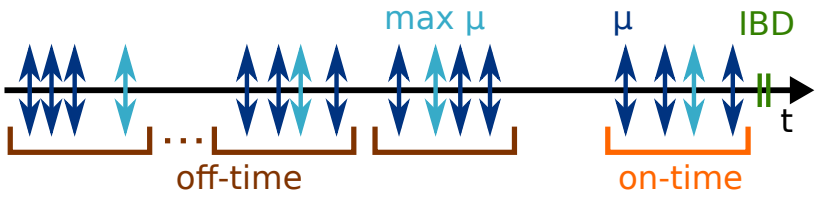

Figure 9.6 – Selection of the μ's (light blue) yielding the maximum posterior probability that the prompt from the IBD candidate (green) is a cosmogenic decay, within on and off-time windows. For simplicity, spallation neutrons were omitted.

Each maximal μ within each off-time window $\mathcal{W}_p^i$, with $i \in [\![1; n_{\mathcal{W}}]\!]$, provides another $P_{max}^i\,(co)$ value to save for the prompt candidate considered. The handling of the on-time window $\mathcal{W}_p^{on}$ is identical to the one presented in 9.2.1.1. Regarding the off-time windows, the only subtlety is that whenever $P_{max}^i\,(co) > P_{max}^{th}$, the bin corresponding to E_p^{vis} in S^{off} must be augmented by $1/n_{\mathcal{W}}$ instead of 1. Equivalently, and for better performance, S^{off} may be filled by a weight equal to the number of maximal μ's providing a posterior probability for the prompt p larger than the threshold P_{max}^{th}, and eventually scaled by $1/n_{\mathcal{W}}$.

Cases in point of S^{on}, S^{off}, and S, for the ND Gd++ Neutrino data, are presented in Figure 9.7. For informative purposes, the ND Gd++ Neutrino data set boasts 257.96 d of live-time, for about 10^3 vetoed events, according to the latest estimations. If it were not for a few differences below 3 MeV, due to natural radioactivity, S^{off} would be fully consistent with the actual $\overline{\nu_e}$ prompt spectrum.

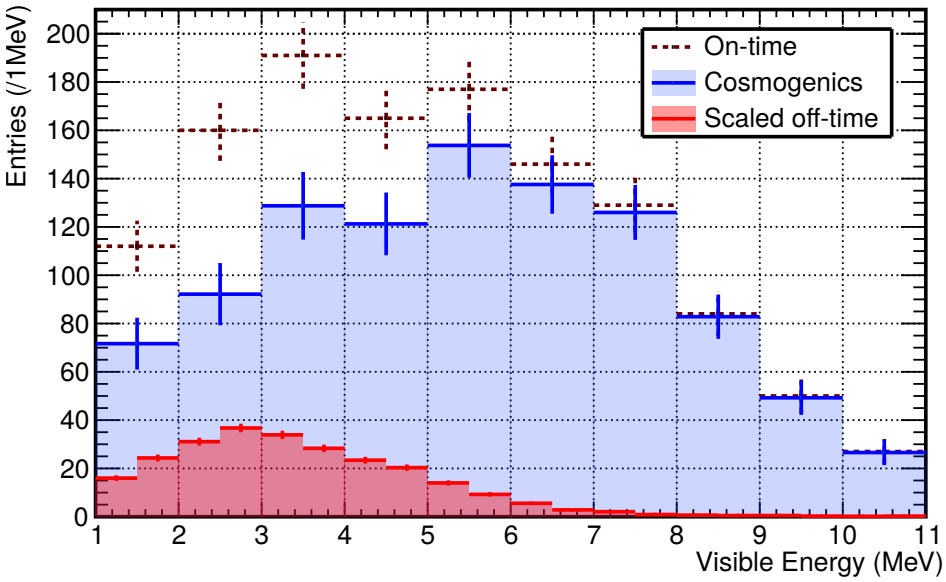

Figure 9.7 – On-time S^{on} (dashed dark red), off-time S^{off} (red), and background subtracted $S = S^{on} - S^{off}$ (blue) spectra. The distributions are based on the fifteen months of ND Gd++ Neutrino data, with a $P_{max}^{th} = 0.4$ threshold.

9.2.1.4 Correlated background

Unexpected features

By definition, the procedure explained in 9.2.1.3 can only remove accidental coincidences of prompts and μ's, within the vetoed sample. A collaboration oversight in the joint analysis of the Gd and H captures was noticed thanks to the cosmogenic veto. The Gd++ naming is, again, a bit of a misnomer, considering that more than half of its delayed events consist of H-captures. At first, a 1 ms after-μ veto was used – as in the Gd analysis – which resulted in the cosmogenic veto tagging prompt events exhibiting a quaint ~ 2 MeV peak, in spite of the aforementioned background subtraction.

With the benefit of hindsight, neutron captures are ideal cosmogenic candidates. Not only are they close to μ-tracks, but they usually have a herd mentality, that is, they are accompanied by other spallation neutrons. How better to boost $P^{on}_{max}\,(co)$ than by providing a short d and a large n ? As far as Gd goes, its large capture cross-section already ensures that all spallation neutrons have been captured 1 ms after a μ-trigger (cf. 6.1.1.2), which leaves the neutron correlated background to the slower H-captures in the GC.

Contamination removal

To dispose of these prompt H-captures, which are thankfully tagged by the cosmogenic veto, we may first increase the Gd++ after-μ veto, thus matching the former H analyses, which had it set to 1.25 ms.

We could also perform a less-intrusive action, which would only concern our tagged sample, without affecting anybody else's analysis, namely offset the on-time window backwards by 2 ms, instead of 1 ms. It was found that a slight correlated contamination remained for a 1.25 ms offset, whose amplitude eventually became the baseline after-μ veto value for the analysis group. An offset larger than 2 ms only served to decrease uniformly the bin contents of the cosmogenic spectra, hence the 2 ms limit.

The comparison of the FDII Gd++ Neutrino data (362.99 d live-time) with 1 ms or 2 ms offsets is shown in Figure 9.8. It is worth underlining that the FDII data set offered the largest statistical deviation, and made for an enlightening illustration, but both FDI and ND exhibit an excess at 2.2 MeV. Let us also emphasise that offsetting the on-time window acts a little differently than increasing the after-μ veto. Whereas the latter truly removes the H-captures from the prompt candidates, offsetting the on-time window prevents the maximisation of $P\,(co \mid n_\mu, d_{p-\mu})$ from finding the true parent μ's of these captures, thereby casting them into the accidental mould.

Insofar as the cosmogenic veto is particularly keen when it comes to tagging these neutron-captures, fitting a H-free spectrum of $\overline{\nu_e}$ candidates with a contaminated cosmogenic data spectrum could have had disastrous consequences in the oscillation fit, all the more so that the oscillation extremum lies near the H peak. In a nutshell, Monte-Carlo predictions may well be biased, but data are no gospel truth either, however much some would have it.

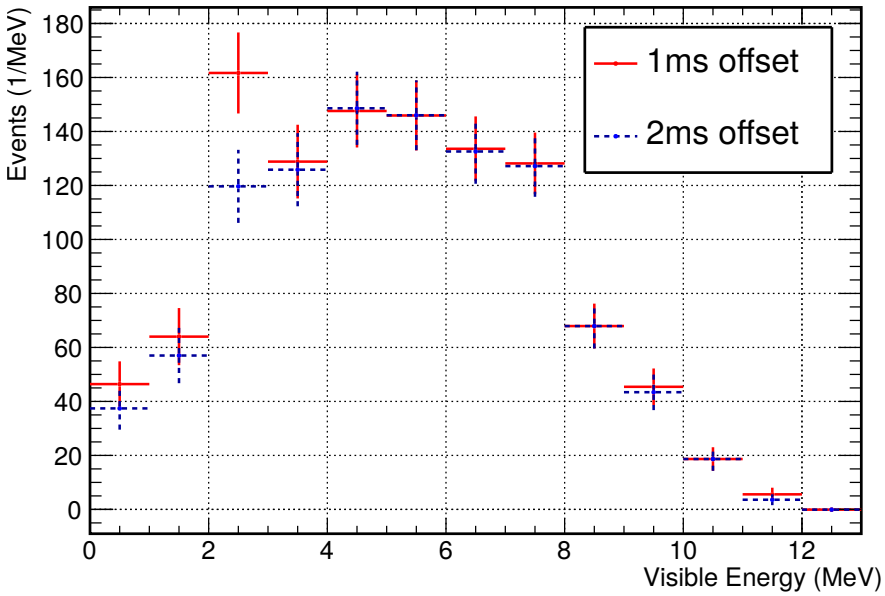

Figure 9.8 – Background subtracted $S = S^{on} - S^{off}$ cosmogenic spectra with 1 ms (red) or 2 ms (dashed blue) offsets. Based on the fifteen months of FDII Gd++ Neutrino data, with a $P^{th}_{max} = 0.4$ threshold.

9.2.2 Veto performance

9.2.2.1 Retrieving characteristic quantities

In order to assess the performance of the veto, we may count the number of correlated events tagged by the veto, and put this efficiency into perspective with the collateral damage caused to the $\overline{\nu_e}$ candidates, i.e. the number of accidental coincidences.

To this end, the distribution of the time intervals $\Delta t_{\mu_m - p} = t_p - t_{\mu_m}$, with p the vetoed prompt and μ_m the muon from the on-time window providing the largest posterior probability, may be analysed. Exactly like when striving to obtain a cosmogenic spectrum, plotting the time difference with the sole maximal μ, rather than with several tenths of them, dramatically curbs the number of accidental coincidences. A combined exponential and first order polynomial fit to the time distribution ought to provide its correlated and accidental components. The method is well-proven, but the errors from the fit are not that small. In the Gd Moriond configuration, for instance, the ND vetoed rate from the fit was $r^{vet}_{co} = 1.37 \pm 0.26$, whereas the down-to-earth method we are about to put forward provided $r^{vet}_{co} = 1.42 \pm 0.11$.

The number of correlated events, and therefore of βn-decays, is directly obtained from the integral and error of the background subtracted spectrum S. As for the $\overline{\nu_e}$ casualties, their number is retrieved from the integral and error of S^{off}. This uncomplicated procedure has the advantage of not relying on a fit, which must, itself, rely on assumptions on the time distribution. Additionally, the use of many off-time windows leads to a well-controlled

subtraction of the accidental background within the vetoed sample, thereby reducing the error on r_{co}^{vet}. Although correlated backgrounds might be present in the first bin of the $\Delta t_{\mu_m - p}$ distribution, their significance would be dwarfed by other bin contents; the fit is partially immune to this type of contamination. On the other hand, unexpected peaks, such as the one found in Figure 9.8, can slyly affect the vetoed rate. This is precisely where the predictions prove useful and vouch for the energy shape extracted by the veto, as is currently the case.

9.2.2.2 Preparing for the near site

To fathom the behaviour of the veto, it is apposite to compute the aforementioned integrals for different P_{max}^{th} threshold values, i.e. for different pairs of spectra $\left(S, S^{off}\right)$. By means of our lightweight software, we can offer a fine-grained study of the evolution of the number of vetoed cosmogenics and collateral $\overline{\nu_e}$'s as a function of P_{max}^{th}.

FD benchmark

The results of the FD case should prove insightful and yet come as no surprise, for the veto was firstly shaped onto the FDI set.

The evolution of the number of cosmogenics and accidentals tagged by the veto, epitomising the efficiency and inefficiency of the latter, is presented for the FD Gd Neutrino data sets in Figure 9.9. The FD Gd++ data sets are identically shaped. Both FDI and FDII are drawn, although their points are reassuringly shifted in accordance with their different live-times (460.35 d for FDI versus 367.10 d for FDII). A signal over background ratio larger than unity is achieved with a threshold as high as $P_{max}^{th} = 0.2$, denoting by "signal" the cosmogenics and "background" the $\overline{\nu_e}$'s, as is customary in this part. This particular value also happens to correspond to a kink in the accidental curve. Thoroughly studying the FD(I) posterior probability reveals that $P_{max}^{FD}(co)$ cannot exceed 0.22 when provided with zero-multiplicity μ's, that is

$$\max_{d \in \mathbb{R}^+} P^{FD}(co \mid 0, d) < 0.22\,. \tag{9.19}$$

One key feature of the FD data is the relatively slow drop of the number of cosmogenics tagged by the veto $N_{co}^{vet}\left(P_{max}^{th}\right)$, as P_{max}^{th} increases. To be sure, Figure 9.9 is presented in logarithmic scale, and yet, let us make our point by quoting some FDI figures. Whereas $N_{co}^{vet}(0.25) = 533 \pm 27$, the cosmogenics do not plummet even with a 0.8 hard cut, that is $N_{co}^{vet}(0.8) = 428 \pm 22$. In other words, the cosmogenics and accidentals are well-separated populations, and most βn-decays are successfully assigned $P_{max}^{on}(co)$ values above 0.8.

Posterior probabilities

Before disclosing the veto performance at ND – via graphs akin to Figure 9.9 – plotting $d \to P(co \mid n, d)$ for fixed values of n, will provide us with better insight on which P_{max}^{th}

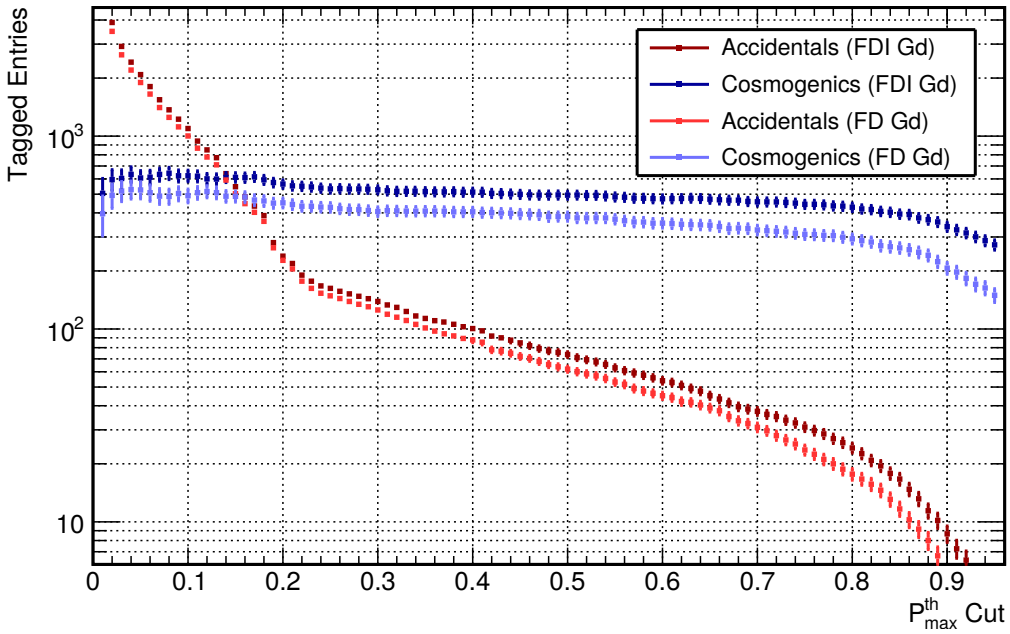

Figure 9.9 – Number of cosmogenics and accidentals tagged by the cosmogenic veto at FDI (dark colours) and FDII=FD (light colours), for the Neutrino Gd data set.

value rules out certain neutron multiplicities and spatial correlations at ND. Inasmuch as the only difference between the ND and FD vetoes is currently the value of the prior ratio π_r (cf. values in Table 9.2), comparing $d \to P\left(co \mid n, d\right)$ for ND and FD emphasises the impact of the former.

Figure 9.10 shows $d \to P^{ND}\left(co \mid n, d\right)$ and $d \to P^{FD}\left(co \mid n, d\right)$ for $n \in [\![0; 2]\!]$. It stems from these graphs that decreasing π_r brings down all $P\left(co \mid n, d\right)$ values, which is not unexpected considering for all $(c, a) \in \mathbb{R}^{+2}$, the function

$$x \to \frac{xc}{xc + a}\,, \tag{9.20}$$

decreases with decreasing x values. It may be stating the obvious, but regardless of the values of the probability densities for accidentals or cosmogenics, which indeed belong to $\mathbb{R}^+$, a lower prior ratio entails smaller posterior probabilities.

All curves in Figure 9.10 exhibit some glitches which are rooted in the binned character of the distance probability densities f^D_{co} and f^D_{acc}. Nonetheless, it is apposite to consider the effect of a baseline $P^{th}_{max} = 0.4$ cut on the maximum distance d it allows. By construction, prompt candidates for which only zero-multiplicity μ's are available in the on-time window cannot be tagged at either site. Likewise, the ND prior ratio implies that only candidates reconstructed less than 4 cm from a μ-track generating one spallation neutron can be vetoed, i.e. essentially none. The FD prior is more forgiving; prompts may be located as far 36 cm from a $n = 1$ μ-track to be removed from the $\overline{\nu_e}$ candidates.

With the benefit of hindsight, all these observations are scarcely surprising: even with

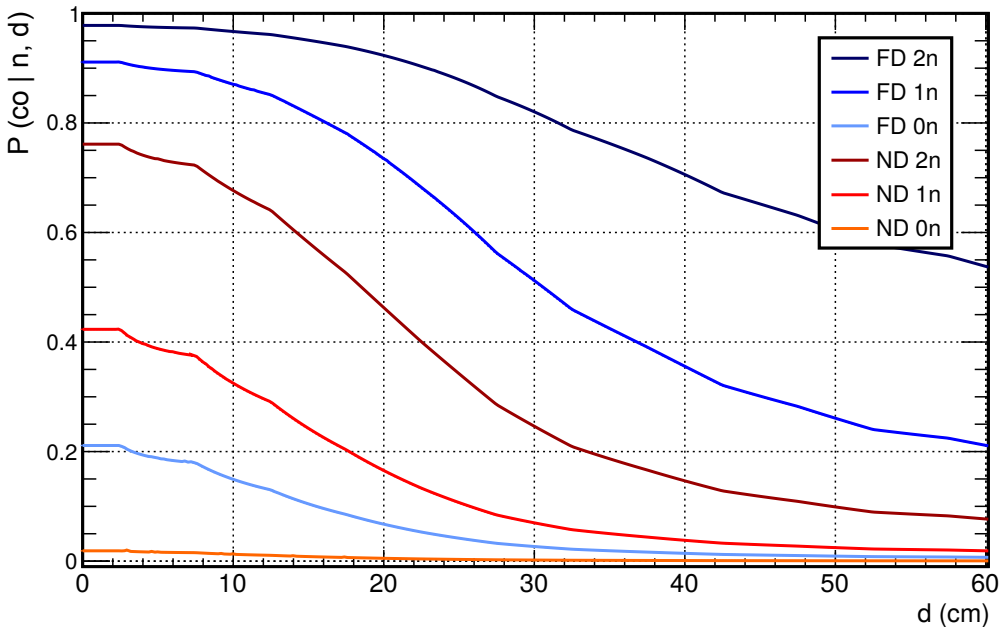

Figure 9.10 – Posterior probability $d \to P\,(co \mid n, d)$ for fixed neutron multiplicities (refer to labels) at ND and FD.

identical (n, d) values in the on-time window, the posterior probability of being a cosmogenic at ND must be lower than it is at FD, because the ratio of accidental coincidences to correlated pairs is indeed substantially larger at ND.

9.2.2.3 Cosmogenics and accidentals at ND

As discussed in 9.1.4.2, the distance part of the cosmogenic veto should be transposable to ND without difficulty, although the appreciably peaked ^{12}B distribution could be questioned. With regard to the neutron multiplicity, the correlated component is difficult to obtain, while the off-time distribution calls for corrections (cf. Figure 9.3). The priors were adapted to the ND case and their ratio is unequivocal: the odds are ~ 15 times worse at ND than at FD. In short, we ought not to expect stellar performance of the cosmogenic veto at ND.

ND walloping

On top of presenting the Neutrino ND Gd points (273.61 d live-time, rates comparable to the Moriond analysis from Chapter 6) with the standard prior ratio $\pi_r^{ND} = 5.5 \times 10^{-4}$, other odds are explored in Figure 9.11. The first striking feature is the rapid drop of all the graphs for the number of cosmogenics, which implies that the βn-decays and the $\overline{\nu_e}$'s are not well-separated populations; the distribution of the maximal posterior probabilities overlap in such a way that letting one in also brings forth the other.

In view of the trickiness of the cosmogenic rate measurement, the ND prior ratio π_r^{ND} is likely to sit a little off from where it should. As a result, it is worthwhile scanning a few other

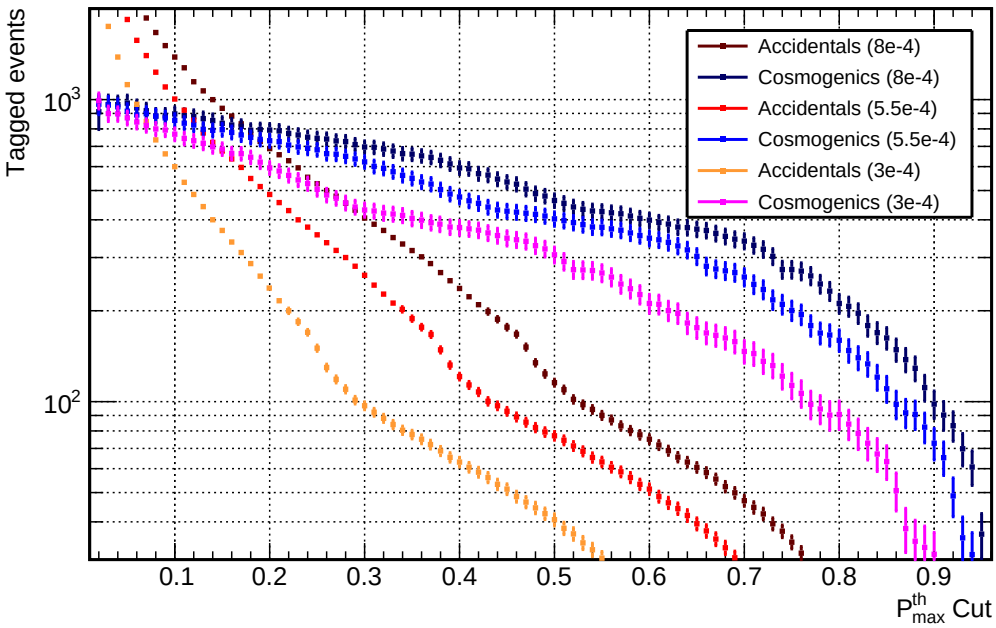

Figure 9.11 – Number of cosmogenics and accidentals tagged by the cosmogenic veto at ND, for the Neutrino Gd data set. Several prior ratio values are tested in addition to the standard $\pi_r = 5.5 \times 10^{-4}$ (refer to legend labels).

π_r^{ND} values, if only to gauge the effect of a prior change on the graphs. From Figure 9.11, it appears that if increasing π_r^{ND} does increase the vetoing efficiency, this improvement does not come without a price: the collateral damage skyrockets for an identical P_{max}^{th} cut. Even more dismaying is that the results achieved for a $P_{max}^{th} = 0.4$ threshold, with $\pi_r = 8 \times 10^{-4}$, are identical to the ones with $P_{max}^{th} \simeq 0.31$ for the default prior. In fact, all graphs are but shifted to the right as π_r^{ND} increases.

Although slightly anticipating the total rate measurements from Chapter 10, let us put forward some efficiency numbers. If $P_{max}^{th} = 0.4$ allows to remove about half of the cosmogenic background at FD, the same cut removes between 23% and 35% (68% C.L.) of the ND βn decays, the uncertainty on the vetoing efficiency being driven by that on the total cosmogenic rate. In light of the neutron multiplicity and distance cuts $P_{max}^{th} = 0.4$ effectively implies at ND and FD (cf. Figure 9.10), this discrepancy in the efficiency of the veto at either detector was not unexpected.

Other mending attempts

We can also think about actually augmenting the prior ratio by reducing the length of the window $t_{\mathcal{W}_p}$ in (9.10). Indeed, dividing $t_{\mathcal{W}_p}$ by two naturally divides the rate of accidental coincidences by two as well. Meanwhile, a 350 ms-long on-time window reduces further the ^{9}Li signal by only 20%. Please note that the total rates r_{co}, r_{cand}, and r_μ in (9.10) are insensitive to $t_{\mathcal{W}_p}$ changes. The results brought about by the shortening of the on-time

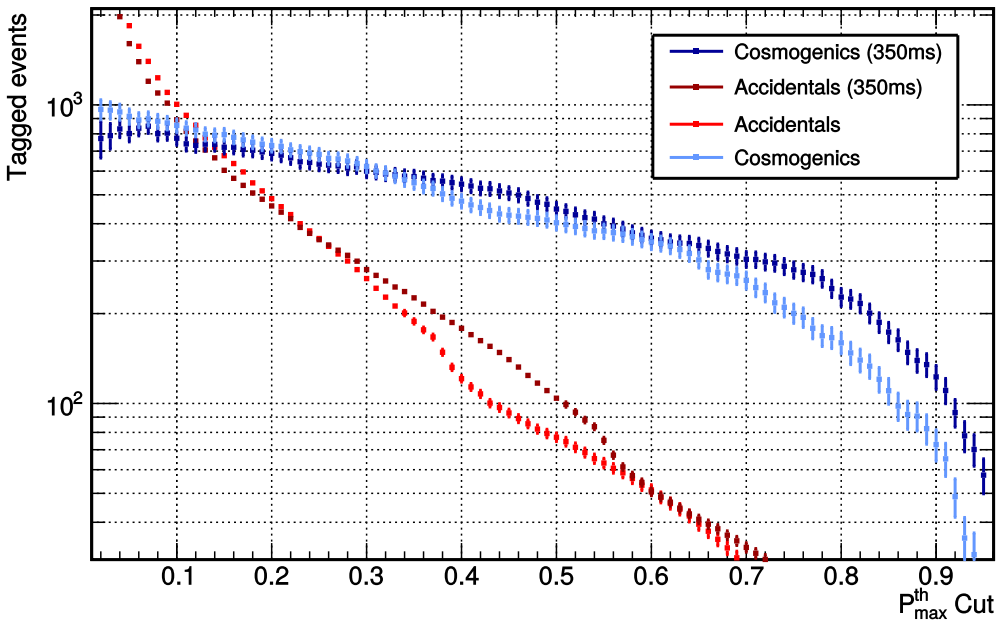

Figure 9.12 – Number of cosmogenics and accidentals tagged by the cosmogenic veto at ND, for the Neutrino Gd data set. The standard 700 ms-long window is compared to a shorter 350 ms version aimed at improving the cosmogenic tagging odds.

window are presented in Figure 9.12. Plainly, the induced prior increase works no miracles. Let us underscore that comparisons at fixed P^{th}_{max} must be put into perspective with the increase of the prior as the window shrinks.

By the same token, one could try and reduce the ND 52.76 Hz μ-rate down to 10 Hz, like at FD. This can be achieved by requiring $E^{vis} > 300$ MeV in the ND μ-definition. Regrettably, the resulting graphs of tagged events are nearly superimposed onto their standard counterparts. In retrospect, these observations are not startling: cosmogenics isotopes are produced by high-energy μ's and disposing of μ's that were not maximising the posterior probability anyway, could not work wonders.

Last but not least, echoing back to the opening paragraph of 9.2.2.3, we can try to shape the neutron multiplicity mass functions for the ND. Even without input from ^{12}B studies, the ND mass function from Figure 9.3, retrieved from uncorrelated prompt-μ pairs, will serve as a P^N_{acc} upgrade, after the missing bin contents have been filled based on that of their neighbours. As for the correlated part, the ND plots from Figure 9.5 are a good starting point, but the first bin content must definitely be worked on. Inspired by the rate measurements from Chapter 10, $P^N_{co}(0) = 0.2$ was set, while the second bin was increased so that $P^N_{co}(1) = 0.45$; the rest of the bins were retrieved from the ND plots, only scaled such that $\sum_{n=0}^{50} P^N_{co}(n) = 1$. Without more statistics or ^{12}B data, we are but blinded, and the P^N_{co} candidates are infinitely many; far be it from the author to claim any sort of flawless systematic study in the building of this custom P^N_{co}, which only turns out to be an educational

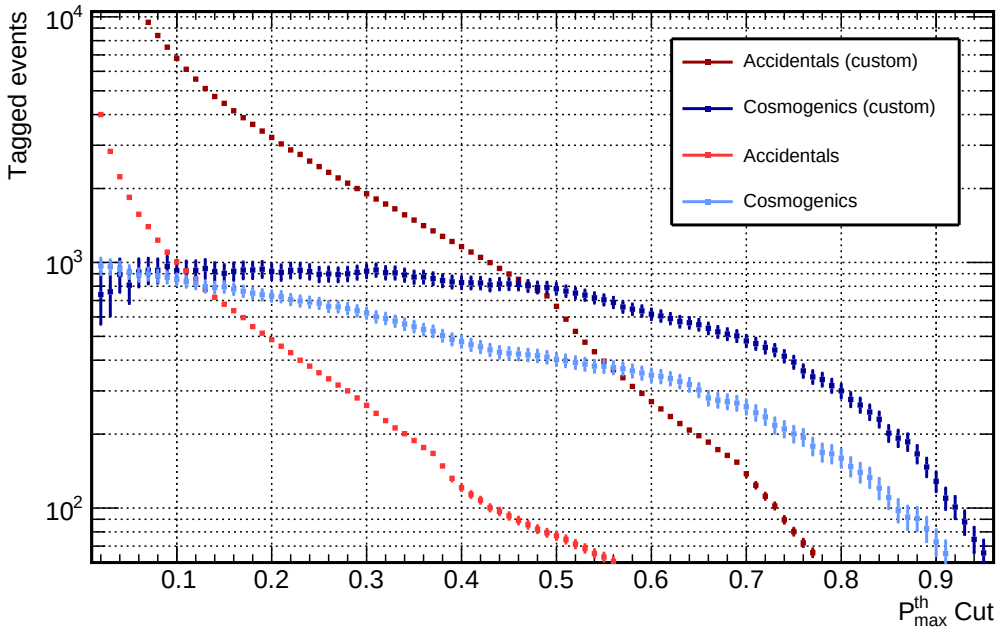

Figure 9.13 – Number of cosmogenics and accidentals tagged by the cosmogenic veto at ND, for the Neutrino Gd data set. The veto performance, when using the default FDI ^{12}B neutron multiplicity mass functions or custom ones, is shown.

mass function.

The results produced by this custom P^N_{co} can be found in Figure 9.13. On its own, the cosmogenic graph is encouraging, it dwindles less readily than the standard version of the veto. Should we graph the posterior probability with these custom mass functions – as we did in Figure 9.10, for fixed multiplicities – we would observe that a $P^{th}_{max} = 0.4$ threshold at ND allows to tag prompts for which a $n = 1$ μ-track is available less than 32 cm away, instead of 4 cm (see 9.2.2.2).

Unfortunately, the custom distributions are not without consequences on the $\overline{\nu_e}$ tagging. Were we to tolerate as few accidentals with the custom veto as with the FDI mass functions, we would need to set a $P^{th}_{max} \simeq 0.71$ cut, for which the veto efficiency is identical to the $P^{th}_{max} = 0.4$ one with the default veto... Fighting against the odds seems vain for all we ever earn is stolen from somewhere else.

9.2.3 Selected threshold

Points 9.2.2.2 and 9.2.2.3 gave us insight into the behaviour of the veto. However, we did not state any criterion to set the P^{th}_{max} value, although we frequently referred to the $P^{th}_{max} = 0.4$ FDI baseline. If truth be told, this baseline was never motivated by any plot, although the corresponding fraction of the $\overline{\nu_e}$ signal wrongly tagged – $I_\nu(0.4) = 0.569 \pm 0.016\%$ – counts very little in the final 92.22% $\overline{\nu_e}$ detection inefficiency due to all the FDI vetoes.

9.2.3.1 Providing the best constraining spectrum

Whilst trying not to increase the minute IBD inefficiency, we may focus on the good that the cosmogenics spectra S can do. To this end, we may recall the usage of the cosmogenic data passed on to the oscillation fit (cf. 6.3.2.1): if the rate is unconstrained, the energy shape of the background is added to the prediction to reproduce the contamination remaining in the $\overline{\nu_e}$ candidates. As a consequence, we wish to maximise the accuracy of the data spectra provided to the oscillation fit.

A high vetoing threshold P^{th}_{max} will tag a meagre, albeit pure, cosmogenic data set. On the other hand, if a low P^{th}_{max} cut will surely boost the number of cosmogenics, this increase will come at the price of a worsened signal over background ratio, denoting by "signal" the cosmogenics and "background" the $\overline{\nu_e}$'s, as usual. A low signal over background ratio makes for a substantial accidental subtraction; however small the error on S^{off} may be with twelve off-time windows, should the accidental coincidences make up 80% of the vetoed events – as is the case at FDI for $P^{th}_{max} = 0.05$ – the error bars of S would inflate. By plotting the relative uncertainty on the integral of S, we may choose a fitting P^{th}_{max} value. Figure 9.14 shows such integrals for all three detectors and the two Neutrino Gd and Gd++ data sets.

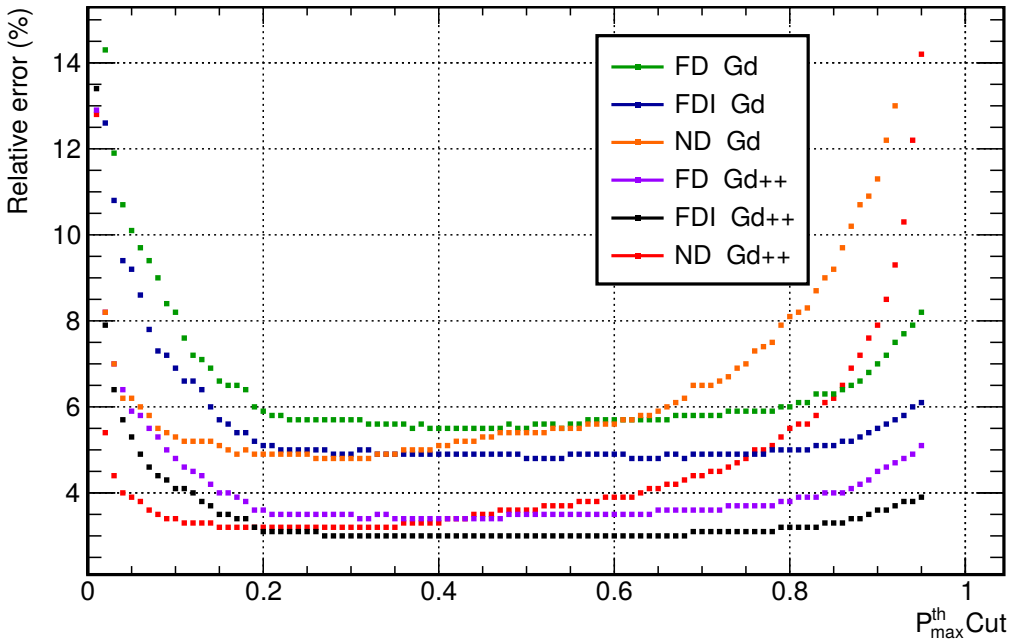

Figure 9.14 – Relative error on the integral of the cosmogenic spectrum $S = S^{on} - S^{off}$ as a function of the veto cut P^{th}_{max}. All three detectors and data sets are presented (see labels; N.B. FDII = FD).

In Figure 9.14, the four FD data sets exhibit a reassuringly broad plateau, for $P^{th}_{max} \in [0.2; 0.8]$; above or below these approximate bounds, the significance of the oscillation fit input worsens, in accordance with our previous statements. Therefore, a $P^{th}_{max} = 0.4$ cut

works flawlessly, but the 0.3 and 0.5 neighbours would do no worse; a 0.5 threshold would decrease even further the IBD inefficiency. Nevertheless, the historical $P_{max}^{th} = 0.4$ value was kept.

Not without reminding the poorer performance of the veto on noisy data, the ND sets are less forgiving: only a narrower $[0.1; 0.4]$ interval offers an absolute variation in the relative uncertainty of less than 0.5%. Although the $I_\nu\,(0.4) = 0.140 \pm 0.004\%$ inefficiency is already well below FD's, thereby allowing to choose a 0.3 cut, for simplicity, $P_{max}^{th} = 0.4$ was also deemed worthy of the ND analyses.

9.2.3.2 Vetoed rates

The rates of cosmogenic decays removed by the veto at each detector, for $P_{max}^{th} = 0.4$ cuts on the Gd Neutrino data sets, are shown in Table 9.3. The Moriond set up yields identical results. Satisfactorily, despite the different live-times and configurations, the FDI and FDII data sets end up presenting identical rounded vetoed rates. Despite the allegedly larger total cosmogenic rate at ND (between 2.5 and 3.2 times higher at a 68% C.L., based on Chapter 10), the vetoing performance is such that only $\times 1.6$ more events are daily tagged at the near site.

Version	N_{co}^{vet}	r_{co}^{vet}
FDI	520 (4.8%)	$(1.12 \pm 0.06)\,\mathrm{d}^{-1}$
FDII	409 (5.5%)	$(1.12 \pm 0.06)\,\mathrm{d}^{-1}$
ND	483 (5.1%	$(1.77 \pm 0.09)\,\mathrm{d}^{-1}$

Table 9.3 – Number of cosmogenics tagged by the cosmogenic veto N_{co}^{vet} (relative uncertainty in parentheses) and corresponding daily rate r_{co}^{vet}, for the Neutrino Gd data sets.

The Gd++ counterpart of Table 9.3 is given in Table 9.4. The ratios of the FD Gd++ rates to their Gd equivalent are not inconsistent with the IBD ratios from (9.11), indeed

$$r_{++}^{co,FDI} = 2.50 \pm 0.14\,, \tag{9.21}$$

$$r_{++}^{co,FDII} = 2.50 \pm 0.16\,. \tag{9.22}$$

Similarly, the ND Gd++ to Gd rate ratio

$$r_{++}^{co,ND} = 2.24 \pm 0.14\,, \tag{9.23}$$

is quite compatible with (9.12). All these ratios support the hypothesis of identical prior ratios for the Gd and Gd++ data sets, although the priors are defined with the total rates.

Version	N_{co}^{vet}	r_{co}^{vet}
FDI	1276 (3.1%)	$(2.80 \pm 0.09)\,\mathrm{d}^{-1}$
FDII	1018 (3.5%)	$(2.80 \pm 0.10)\,\mathrm{d}^{-1}$
ND	1025 (3.5%)	$(3.97 \pm 0.14)\,\mathrm{d}^{-1}$

Table 9.4 – Number of cosmogenics tagged by the cosmogenic veto N_{co}^{vet} (relative uncertainty in parentheses) and corresponding daily rate r_{co}^{vet}, for the Neutrino Gd++ data sets.

9.2.3.3 IBD inefficiencies

Regarding the accidental coincidences of prompt-μ pairs tagged by the veto, they are best presented in terms of IBD inefficiencies, which represent the fraction of true $\overline{\nu_e}$'s removed from the IBD candidates for each detector.

By selecting the μ giving the largest posterior probability, we build a unique pair for each prompt candidate, but it does not prevent the βn-decays from acting as prompts artificially amplifying the number of accidental coincidences. As a result, the number of accidentals must be scaled down by 2.1% and 4.5%, at ND and FD, respectively, to deduce the number of $\overline{\nu_e}$ casualties. It is worth making note of the fact that these corrections, comparable to the uncertainties on the number of accidentals themselves, must rely on the total cosmogenic rate measurements, since the cosmogenic veto has yet to be applied.

The Gd inefficiencies can be found in Table 9.5 while Table 9.6 gives the Gd++ ones. Computing the ratios of the Gd++ to Gd numbers shows nothing alarming.

Version	N_{acc}^{vet}	I_ν
FDI	102 (2.9%)	$0.569 \pm 0.016\%$
FDII	88 (3.1%)	$0.566 \pm 0.018\%$
ND	121 (2.6%)	$0.140 \pm 0.004\%$

Table 9.5 – Number of accidentals tagged by the cosmogenic veto N_{acc}^{vet} (relative uncertainty in parentheses) and corresponding IBD inefficiency I_ν, for the Neutrino Gd data sets.

Version	N_{acc}^{vet}	I_ν
FDI	260 (1.8%)	$0.522 \pm 0.009\%$
FDII	221 (2.0%)	$0.526 \pm 0.011\%$
ND	261 (1.8%)	$0.125 \pm 0.002\%$

Table 9.6 – Number of accidentals tagged by the cosmogenic veto N_{acc}^{vet} (relative uncertainty in parentheses) and corresponding IBD inefficiency I_ν, for the Neutrino Gd++ data sets.

9.3 Spectral analysis

By means of the cosmogenic veto, we have at hand spectra of prompt candidates which should account for the βn-decays of ^{9}Li and ^{8}He inside each detector, for we strived to remove the accidental and correlated backgrounds within the tagged samples. It is high time we put forth arguments stronger than the graphical visualisation of Figure 6.5 to combine all the data spectra, and eventually compare them to the predictions from Chapter 8. Thereafter, all results rely on the blessed $P^{th}_{max} = 0.4$ threshold.

9.3.1 Compatibility across detectors

9.3.1.1 Graphical appetiser

Although we showed the Moriond Gd data spectra in Chapter 6, because these formed the fit input of the only multi-detector result from Double Chooz to this date, we can here present the Neutrino Gd data, in accordance with the previous sections of the current chapter.

The reader can find the Gd spectra, normalised to unit area for clear-sighted comparisons, in Figure 9.15. So far, there are no visual elements demonstrating significant discrepancies between these three independent data sets. In spite of the $\times 1.81$ and $\times 1.73$ live-time increases at FD and ND, respectively, with respect to the Moriond Gd data set, the hypothesis that all spectra are compatible with one another still holds.

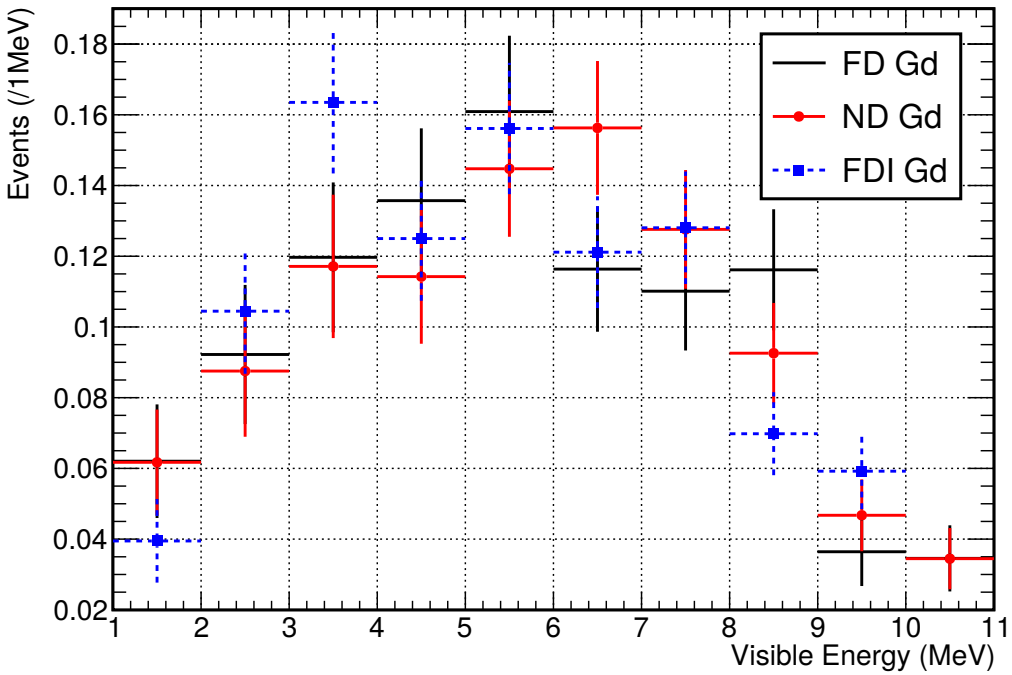

Figure 9.15 – Cosmogenic data spectra built from the Neutrino Gd data sets. All the spectra are normalised to unit area. N.B. FDII = FD.

The Gd++ counter part of Figure 9.15 is presented in Figure 9.16. Regardless of the augmented statistics, with respect to the Gd sets, (cf. 9.2.3.2), the Gd++ spectra seem to agree with one another even better than their Gd equivalent.

Let us underscore that, strictly speaking, the errors on all the spectra are different from the mere square roots of their original bin contents, for indeed, below 8 MeV, the $\overline{\nu_e}$ spectrum within S^{off} matters. Nonetheless, we did our best to dwarf the impact of the accidental background subtraction on the accuracy of the final S.

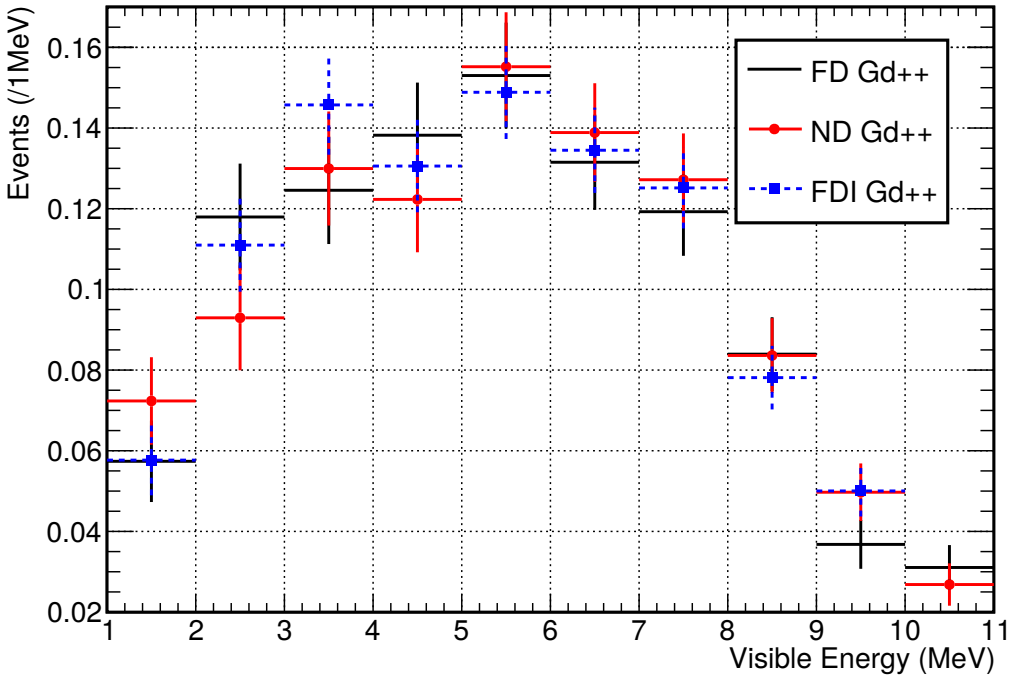

Figure 9.16 – Cosmogenic data spectra built from the Neutrino Gd++ data sets. All the spectra are normalised to unit area. N.B. FDII = FD.

9.3.1.2 χ^2 test for homogeneity

Perhaps less disputable than comparisons to the naked eye is the χ^2 test for homogeneity, which is a statistical test for probing the compatibility between two spectra. To perform it, a coarse 1 MeV binning – as in Figures 9.15 and 9.16 – was utilised, thereby ensuring Gaussian statistics above 9 MeV, where the meagre FDII Gd set counts few events in S^{on}, and thus, in S.

The χ^2 test for homogeneity can be applied to two spectra from different data sets, preferably from different detectors, yet analysed with the same selection, i.e. either Gd or Gd++. The null hypothesis H_0, on which the test relies, states that the two binned spectra represent random values with identical distributions. The χ^2 test returns the so-called "P-value".

However more many would have it represent, the P-value is but the probability of obtaining a χ^2 value larger than observed, provided that the null hypothesis is true. Traditionally, if the P-value is lower than 0.05, the odds of obtaining more extreme results, considering H_0 holds, are deemed too low[10] for our observation to be qualified as a mere stroke of bad luck. In other words, an a priori threshold on the acceptable P-values offers a criterion for discarding H_0.

Be that as it may, we must needs emphasise that the P-value cannot validate the null hypothesis, however high the P-value may be. Exactly like the actual likelihood from 9.1.2.1, a frequentist approach only reports the plausibility of one's observation, assuming a certain hypothesis to be true. In order to actually reverse the conditions, that is, compute the probability of one's hypothesis – as we did with the posterior probability[11] in 9.1.2.2 – a degree of belief, embodied by Bayes' priors, cannot be avoided. To put the final nail into this coffin of misconceptions, let us stress that a high P-value, from a χ^2 test, only signifies that we have failed to reject the null hypothesis. In our case, it implies that we could not prove that the two spectra were pulled from different distributions, no more, no less, which is not disheartening per se.

The P-values from the χ^2 tests for homogeneity can be found in Table 9.7. The off-diagonal blocks, testing the compatibility between the Gd++ and Gd sets are only shown for completeness. As can be noticed, the two blocks on the diagonal display P-values well-above the a priori 0.05 threshold for rejecting H_0.

Version	ND_{++}	FD_{++}	FDI_{++}	ND	FD	FDI
ND_{++}	1	0.53	0.74	0.61	0.43	0.089
FD_{++}	0.53	1	0.31	0.089	0.54	0.019
FDI_{++}	0.74	0.31	1	0.065	0.1	0.18
ND	0.61	0.089	0.065	1	0.85	0.16
FD	0.43	0.54	0.1	0.85	1	0.17
FDI	0.089	0.019	0.18	0.16	0.17	1

Table 9.7 – P-value from χ^2 tests between all six data spectra.

9.3.1.3 Kolmogorov-Smirnov test

To compare two empirical samples, the Kolmogorov-Smirnov test defines the largest difference between the two cumulative distribution functions of the sets as a measure of disagreement. Because it relies on the largest distance between the two cumulative distribution

[10] Other thresholds, such as 0.1 or 0.01 may be used for rejecting H_0.

[11] For the cosmogenic veto, reversing the conditions boils down to translating "the probability of observing (n, d) if the event is indeed a βn-decay" into "the probability that the event is a cosmogenic decay given the observed (n, d)".

functions, i.e. the so-called supremum, the Kolmogorov-Smirnov test is sensitive to both normalisation issues and shape distortions. Although the Kolmogorov-Smirnov test is fundamentally designed to be applied to continuous unbinned data, the only way to remove the accidental contribution from the vetoed sample is to perform a statistical subtraction, as opposed to an event-by-event discrimination. Nevertheless, by testing several binnings, we can somehow put trust into the resulting P-values, which challenge the same null hypothesis as the χ^2 from 9.3.1.2.

Table 9.8 discloses the most pessimistic values we found, with a 100 keV-binning. Other bin widths, such as 250 keV, 500 keV and 1 MeV, were also tested; they all yielded larger P-values than those in Table 9.8. Please note that the belief according to which P-values increase whilst the number of bins decreases is – more oft than not – an erroneous rule of thumb [139].

Version	ND_{++}	FD_{++}	FDI_{++}	ND	FD	FDI
ND_{++}	1	0.78	0.69	0.77	0.91	0.68
FD_{++}	0.78	1	0.99	0.17	0.92	0.84
FDI_{++}	0.69	0.99	1	0.19	0.51	0.98
ND	0.77	0.17	0.19	1	0.83	0.26
FD	0.91	0.92	0.51	0.83	1	0.81
FDI	0.68	0.84	0.98	0.26	0.81	1

Table 9.8 – P-value from Kolmogorov tests between all six data spectra.

Like its χ^2 analogue, Table 9.8 shows P-values well above the most common thresholds. In accordance with our pet peeve, currently, there is no evidence against combining all the data spectra for each selection. With more statistics, we might reach a different conclusion. Obviously, with infinitely many events, we would reject the null hypothesis, because for all our endeavours into shaping the three detectors alike, some discrepancies would eventually peer through (enery scale, overburden), but this time is nowhere near.

9.3.2 Comparison to Monte-Carlo predictions

Based on the alleged compatibility of all the spectra for one analysis channel, as discussed in 9.3.1, S will henceforth denote the sum of all three data spectra, i.e. $S = S^{ND}+S^{FDII}+S^{FDI}$, be it in the Gd (defining S^{Gd}) or Gd++ (S^{++}) analysis. These combined spectra – which must be seen as vectors of bin contents, as emphasised in 8.2.1.2 – will be compared to the Monte-Carlo predictions from 8.4.

9.3.2.1 χ^2 fit

Bestiary

Owing to 8.4, we have in our possession four building blocks, namely, one simulated spectrum per cosmogenic isotope (^{8}Li or ^{9}Li) and neutron-capture selection (Gd or H), which form the following set

$$\mathcal{S} = \left\{ S^{\mathrm{Gd}}_{\mathrm{He}}, S^{\mathrm{H}}_{\mathrm{He}}, S^{\mathrm{Gd}}_{\mathrm{Li}}, S^{\mathrm{H}}_{\mathrm{Li}} \right\} . \tag{9.24}$$

By the same token, we have access to four covariance matrices

$$\mathcal{V} = \left\{ V^{\mathrm{Gd}}_{\mathrm{He}}, V^{\mathrm{H}}_{\mathrm{He}}, V^{\mathrm{Gd}}_{\mathrm{Li}}, V^{\mathrm{H}}_{\mathrm{Li}} \right\} , \tag{9.25}$$

accounting for the correlations and uncertainties within each spectrum. In addition, there are two covariance matrices

$$\mathcal{C} = \left\{ C^{\mathrm{Gd,H}}_{\mathrm{He}}, C^{\mathrm{Gd,H}}_{\mathrm{Li}} \right\} , \tag{9.26}$$

providing the correlations of the energy regions across capture channels. Please note that, by definition of the matrix elements in $C \in \mathcal{C}$, we have

$$C^{\mathrm{H,Gd}}_{\mathrm{He}} = {}^{\mathrm{t}}C^{\mathrm{Gd,H}}_{\mathrm{He}} , \tag{9.27}$$

$$C^{\mathrm{H,Gd}}_{\mathrm{Li}} = {}^{\mathrm{t}}C^{\mathrm{Gd,H}}_{\mathrm{Li}} . \tag{9.28}$$

Gd χ^2 expression

The Gd data spectrum S^{Gd} is fitted by a linear combination of the $S^{\mathrm{Gd}}_{\mathrm{He}}$ and $S^{\mathrm{Gd}}_{\mathrm{Li}}$ predicted spectra, weighted by the number of ^{8}He and ^{9}Li decays, $\overrightarrow{n} = (n_{\mathrm{He}}, n_{\mathrm{Li}})$, respectively, which are the two free parameters of the total Monte-Carlo spectrum S^{Gd}_{MC}. With the current statistics, it does not make sense to distinguish the ^{8}He fraction between the Gd and Gd++ channels[12]. We have

$$S^{\mathrm{Gd}}_{MC}\left(\overrightarrow{n}\right) = n_{\mathrm{He}}\, S^{\mathrm{Gd}}_{\mathrm{He}} + n_{\mathrm{Li}}\, S^{\mathrm{Gd}}_{\mathrm{Li}} , \tag{9.29}$$

which defines the Gd residuals

$$R_{\mathrm{Gd}}\left(\overrightarrow{n}\right) = S^{\mathrm{Gd}} - S^{\mathrm{Gd}}_{MC}\left(\overrightarrow{n}\right) , \tag{9.30}$$

and the Gd χ^2

$$\chi^2_{\mathrm{Gd}}\left(\overrightarrow{n}\right) = {}^{\mathrm{t}}R_{\mathrm{Gd}}\left(\overrightarrow{n}\right) V_{\mathrm{Gd}}\left(\overrightarrow{n}\right)^{-1} R_{\mathrm{Gd}}\left(\overrightarrow{n}\right) . \tag{9.31}$$

The covariance matrix $V_{\mathrm{Gd}}\left(\overrightarrow{n}\right)$ is the variance of the multivariate residual $R_{\mathrm{Gd}}\left(\overrightarrow{n}\right)$, which de facto changes with every new $\overrightarrow{n}$ value. In accordance, denoting by V^{Gd}_{data} the diagonal matrix, whose coefficients are the squares of the errors on the bin contents of the data spectrum S^{Gd}, we find

$$V_{\mathrm{Gd}}\left(\overrightarrow{n}\right) = V^{\mathrm{Gd}}_{data} + {n_{\mathrm{He}}}^2\, V^{\mathrm{Gd}}_{\mathrm{He}} + {n_{\mathrm{Li}}}^2\, V^{\mathrm{Gd}}_{\mathrm{Li}} , \tag{9.32}$$

[12] In view of the presence of fast-neutrons in the βn decay trees, there could be some subtle differences in the sensitivity to ^{8}He in one channel or the other; that is well beyond the scope of the experiment.

since the ^{8}He and ^{9}Li simulations have independent errors[13].

The fit is performed by minimising χ^2_{Gd} with respect to $\overrightarrow{n}$, whose two elements are free positive real numbers. Practically, a `ROOT::Minuit2::Minuit2Minimizer` is utilised on a `C++` functor encapsulated inside a `ROOT::Math::Functor` [105].

Gd++ χ^2 expression

With the spectra and matrices we have at our disposal, the differences between the Gd++ and Gd fits lie in the inclusion of the Monte-Carlo simulations. When performing the Gd++ analysis on the data, we can compute the fraction f_{Gd} of events captured on Gd; this fraction evaluates to $f_{\mathrm{Gd}} \simeq 0.36$. Thus, we can write

$$S^{++}_{MC}\left(\overrightarrow{n}\right) = n_{\mathrm{He}}\left(f_{\mathrm{Gd}}\ S^{\mathrm{Gd}}_{\mathrm{He}} + (1-f_{\mathrm{Gd}})\ S^{\mathrm{H}}_{\mathrm{He}}\right) + n_{\mathrm{Li}}\left(f_{\mathrm{Gd}}\ S^{\mathrm{Gd}}_{\mathrm{Li}} + (1-f_{\mathrm{Gd}})\ S^{\mathrm{H}}_{\mathrm{Li}}\right), \tag{9.33}$$

considering we make no distinction between the ^{8}He fractions of the two channels.

The Gd++ residuals are a carbon copy of (9.30), only swapping the Gd subscripts and superscripts for ++. The covariance matrix $V_{++}\left(\overrightarrow{n}\right)$ is slightly thornier, since correlations between Gd and H appear, therefore

$$\begin{aligned} V_{++}\left(\overrightarrow{n}\right) = V^{++}_{data} \\ &+ {n_{\mathrm{He}}}^2\left[{f_{\mathrm{Gd}}}^2\ V^{\mathrm{Gd}}_{\mathrm{He}} + f_{\mathrm{Gd}}\left(1-f_{\mathrm{Gd}}\right)\left(C^{\mathrm{Gd,H}}_{\mathrm{He}} + {}^{\mathrm{t}}C^{\mathrm{Gd,H}}_{\mathrm{He}}\right) + \left(1-f_{\mathrm{Gd}}\right)^2\ V^{\mathrm{H}}_{\mathrm{He}}\right] \\ &+ {n_{\mathrm{Li}}}^2\left[{f_{\mathrm{Gd}}}^2\ V^{\mathrm{Gd}}_{\mathrm{Li}} + f_{\mathrm{Gd}}\left(1-f_{\mathrm{Gd}}\right)\left(C^{\mathrm{Gd,H}}_{\mathrm{Li}} + {}^{\mathrm{t}}C^{\mathrm{Gd,H}}_{\mathrm{Li}}\right) + \left(1-f_{\mathrm{Gd}}\right)^2\ V^{\mathrm{H}}_{\mathrm{Li}}\right]. \end{aligned} \tag{9.34}$$

The χ^2 is straightforward.

9.3.2.2 Fit results

Gd results

The combined data spectrum S^{Gd}, along with the prediction at the best fit $S^{\mathrm{Gd}}_{MC}\left(\overrightarrow{n_m}\right)$, are presented in Figure 9.17.

As a matter of fact, the best fit corresponds to a zero ^{8}He fraction and a reduced chi-square

$$\frac{\chi^2_{\mathrm{Gd}}\left(\overrightarrow{n_m}\right)}{n_{df}} = \frac{20.1}{18} \tag{9.35}$$

with n_{df} the number of degrees of freedom in the fit. The (9.35) value corresponds to a P-value of 0.33, i.e. there is no evidence that the prediction and the data come from different distributions. The norm of $S^{\mathrm{Gd}}_{MC}\left(\overrightarrow{n_m}\right)$ is found to be $98.5 \pm 3.9\%$ that of the data spectrum, via the n_{Li} value. The overall excellent agreement between the data and the Monte-Carlo is rooted both in the limited statistics of the sample and (mayhaps) in the quality of the predictions and analyses.

[13] Had we included some systematic uncertainty on the energy scale or quenching, the latter statement would only approximately hold.

Had we not constrained the number of events to positive values, we would have found a ^{8}He fraction $\alpha_{\mathrm{He}} = -5.0 \pm 4.4\%$ at the best fit, that is, a value perfectly compatible with zero.

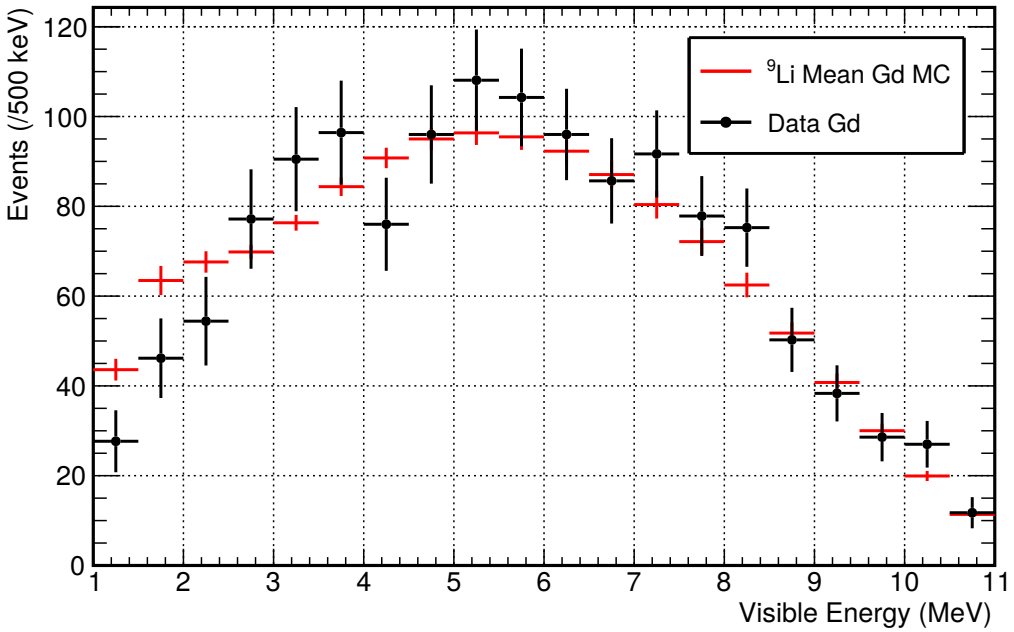

Figure 9.17 – Combined Gd data spectrum S^{Gd} (black) and best-fit prediction $S^{\mathrm{Gd}}_{MC}(\overrightarrow{n_m})$ (red). No evidence of ^{8}He to $> 1\sigma$.

Gd++ results

The standard Gd++ fit, as described in 9.3.2.1, returns a reduced chi-square of 33.7/18, thereby corresponding to a P-value of 0.013, which is below the common 0.05 threshold.

For all the H captures within the Gd++ set, a fit of the Gd++ data with the sole H simulations yields an even worse reduced chi-square of 38.3/18. On the contrary, a fit of the Gd++ data spectrum with the sole Gd prediction seems a better match, returning

$$\frac{\chi^2_{++}(\overrightarrow{n_m})}{n_{df}} = \frac{28.5}{18}, \tag{9.36}$$

and accordingly, a more suitable, albeit still low, P-value of 0.055. The best-fit predictions and the Gd++ spectrum from data are shown in Figure 9.18. The value of the ^{8}He fraction will be discussed in 9.3.2.3.

Let us try and justify the poorer plausibility of the H fit. First and foremost, as reviewed in 5.2.3.4, the agreement between the data and the Monte Carlo of the calibration sources is worse in the GC – where H-captures are the rule – than in the NT. This discrepancy originates from slightly ill-suited optical parameters for the liquids, and the quaint particles we throw at the detector simulation – all the more so with high energies – might not be modelled as

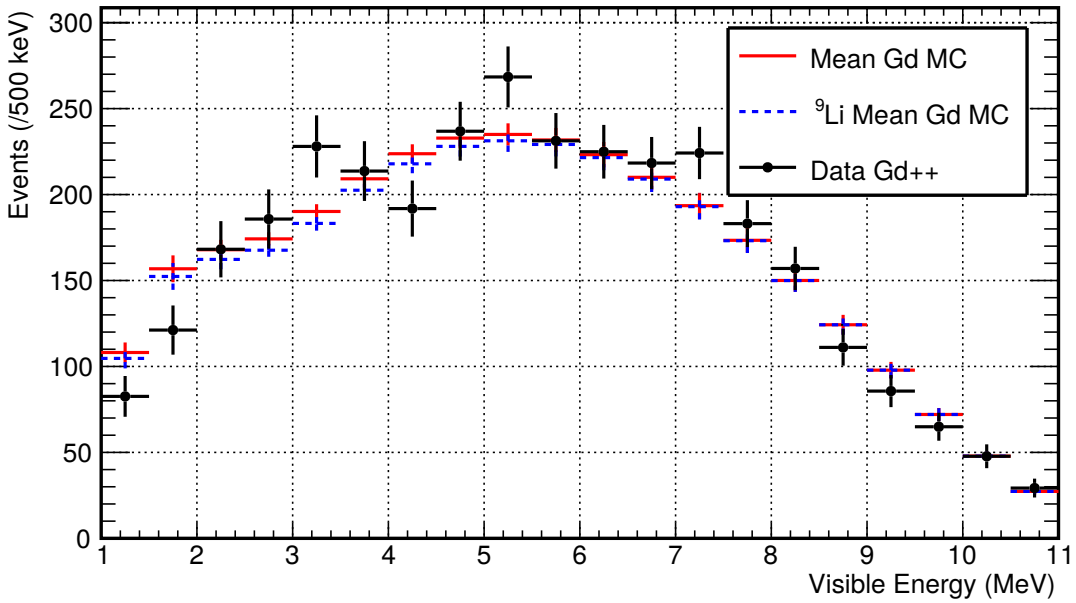

Figure 9.18 – Combined Gd++ data spectrum S^{++} (black), best-fit prediction $S^{\mathrm{Gd}}_{MC}(\overrightarrow{n_m})$ (red) and ^{9}Li component (dashed blue).

flawlessly as standard IBD interactions. Not only may they be modelled approximately, but the quality of the modelling might be lesser in the GC.

As a matter of fact, the largest residual in Figure 9.18 is a little below 2 MeV, which is precisely where several decays modes of the 11.28 MeV and 11.81 MeV levels in ^{9}Be emerge, as may be verified via the spectra database from Figure 8.5. Presumably, it is too soon to fervently bet on this; still, the fact that these high energy-levels minimise the energy attributed to the e^-, whilst offering a greater share to strongly quenched particles, such as α's (cf. 5.1.3.3), is arousing. Let us emphasise that this paragraph applies equally well to the Gd spectra from Figure 9.17.

Moreover, it is also for the high-lying states that the discrepancies between the Gd and H cuts come out in a most striking way, bordering on the dubious. Notwithstanding the presence of a few average residuals, the agreement of the predictions, with a data set nearly four times larger than its FDI predecessor, is remarkable.

9.3.2.3 ^{8}He fraction

Whilst the agreement of the cosmogenic data spectra and the simulations is tested, a ^{8}He fraction is de facto estimated. Measuring the ^{8}He fraction, or setting limits on it, would help to characterise the dependency of the cosmogenic production on the overburden.

Definition and uncertainty

Since the unconstrained fit yielded a negative number of ^{8}He decays, we relied on the integral

of the data spectrum to assess the denominator in the definition of the ^{8}He fraction α_{He} for the Gd fit. However, the ^{8}He counts n_{He} for the Gd++ set are positive, thereby allowing to write

$$\alpha_{\mathrm{He}} = \frac{n_{\mathrm{He}}}{n_{\mathrm{He}} + n_{\mathrm{Li}}} \,. \tag{9.37}$$

Computing the error on α_{He} is not as trivial as it may seem. To this end, let us denote by f a function of two random variables X, Y whose expression is defined below,

$$f\left(X, Y\right) = \frac{X}{X + Y} \,. \tag{9.38}$$

A Taylor expansion around the mean expectation $\mathbb{E}\left[f\left(X, Y\right)\right]$ of f allows to approximate its variance $\mathbb{V}$ as

$$\mathbb{V}\left[f\left(X, Y\right)\right] \simeq f'_x\left(e\right)^2 \mathbb{V}\left[X\right] + f'_y\left(e\right)^2 \mathbb{V}\left[Y\right] + 2 f'_x\left(e\right) f'_y\left(e\right) \mathbb{C}\left[X, Y\right] , \tag{9.39}$$

where e symbolises $\left(\mathbb{E}\left[X\right], \mathbb{E}\left[Y\right]\right)$, $\mathbb{C}$ the covariance, f'_x and f'_y the partial derivatives of f with respect to X or Y. Substituting the expressions of the derivatives into (9.39), we find

$$\mathbb{V}\left[f\left(X, Y\right)\right] \simeq \frac{\mathbb{E}\left[X\right]^2 \mathbb{E}\left[Y\right]^2}{\left(\mathbb{E}\left[X\right] + \mathbb{E}\left[Y\right]\right)^4} \left(\frac{\mathbb{V}\left[X\right]}{\mathbb{E}\left[X\right]^2} + \frac{\mathbb{V}\left[Y\right]}{\mathbb{E}\left[Y\right]^2} - \frac{2\,\mathbb{C}\left[X, Y\right]}{\mathbb{E}\left[X\right]\mathbb{E}\left[Y\right]} \right) . \tag{9.40}$$

Insofar as the correlation coefficient between n_{He} and n_{Li} is estimated to be $\rho = -0.78$ by the Gd++ fit (and $\rho = -0.76$ by the Gd one), it is apt to use (9.40) instead of plain derivatives for the error estimation, which would otherwise largely underestimate[14] the error on α_{He}. Taking the square root of 9.40 and adapting the notations, the uncertainty on α_{He} eventually reads

$$\sigma_{\alpha_{\mathrm{He}}} = \frac{\sqrt{{n_{\mathrm{Li}}}^2\, {\sigma_{n_{\mathrm{He}}}}^2 + {n_{\mathrm{He}}}^2\, {\sigma_{n_{\mathrm{Li}}}}^2 - 2\rho\, n_{\mathrm{He}}\, n_{\mathrm{Li}}\, \sigma_{n_{\mathrm{He}}}\, \sigma_{n_{\mathrm{Li}}}}}{\left(n_{\mathrm{He}} + n_{\mathrm{Li}}\right)^2} \,. \tag{9.41}$$

From (9.41), we can state that the ^{8}He fraction corresponding to the (9.36) chi-square value is

$$\alpha_{\mathrm{He}} = 1.7 \pm 3.2\% \,, \tag{9.42}$$

i.e. a fraction fully consistent with zero. Input from other experiments exist, but they were obtained at different depths, and in the case of Borexino, only a lower limit exists. This limit asserts that the ^{8}He fraction must be lower than 13% at a 99.7% confidence level [140]. From KamLAND, an extrapolation to the Double Chooz depth was performed in [141], resulting in $\alpha_{\mathrm{He}} = 7.9 \pm 6.5\%$. To cut the matter short, at the present time, ^{8}He keeps a low profile.

[14] Using plain derivatives instead of starting from the definition of the variance amounts to setting $\rho = 1$, which usually overestimates errors; this is quite the opposite in our case.

Sensitivity study

The amount of cosmogenic data available has more than doubled by means of the novel Gd++ selection; the combined Gd++ spectrum boasts about 3.3×10^3 events (cf. Table 9.4). To grow the sample larger, there is probably not much to be done but wait, a grim prospect. To be sure, there might exist state of the art techniques to tag more efficiently cosmogenic decays, but the tagged rates cannot be increased tenfold. In any event, it is enlightening to gauge how many more βn decays we would need in the combined Gd++ cosmogenic spectrum to measure a ^{8}He fraction with a certain significance. In a nutshell, we here present a sensitivity study.

This sensitivity study relies on the fit methods presented in 9.3.2.1. The main difference is that we purposefully choose a value $\alpha \in]0;1[$ for the ^{8}He fraction, to build, from the simulations, a fake data spectrum[15]

$$S_{toy} = \alpha \; S^{\mathrm{Gd}}_{\mathrm{He}} + (1-\alpha) \; S^{\mathrm{Gd}}_{\mathrm{Li}} \, . \tag{9.43}$$

This toy spectrum S_{toy} is then fitted by $\overrightarrow{n} \to S^{\mathrm{Gd}}_{MC}(\overrightarrow{n})$, as defined in (9.29). Unmistakably, the best-fit parameters $\overrightarrow{n_m}$ always satisfy

$$\overrightarrow{n_m} = (\alpha, 1-\alpha) \; , \tag{9.44}$$

and the chi-square is always zero at $\overrightarrow{n_m}$. Nevertheless, from the errors on $\overrightarrow{n_m}$, we can evaluate what margin we have from a zero ^{8}He fraction, namely, check whether

$$\alpha - k\,\sigma_{n_{\mathrm{He}}} > 0 \, , \tag{9.45}$$

with k the desired significance in terms of σ's. If (9.45) cannot be met, we seek how larger a data set we should need to observe α with a k significance, i.e. by which factor $\sqrt{d_k}$ the errors bars within V^{++}_{data} must be divided. We here assume that multiplying our number of events by d_k will divide all the data error bars by $\sqrt{d_k}$, which holds to a good approximation. Undoubtedly, the Monte-Carlo errors will always remain within $V_{++}(\overrightarrow{n})$, and there always is a d_k value beyond which the sensitivity stops noticeably improving.

Many α's and d_k values for $k \in \{2,3,4\}$ were probed to build Figure 9.19. To illustrate one case in point, Double Chooz would need twice as much data as what the current Gd++ set represents, to rule out KamLAND's extrapolation to 3σ. Ampler data samples are displayed for indicative purposes.

[15] By construction, this data spectrum is normalised to one.

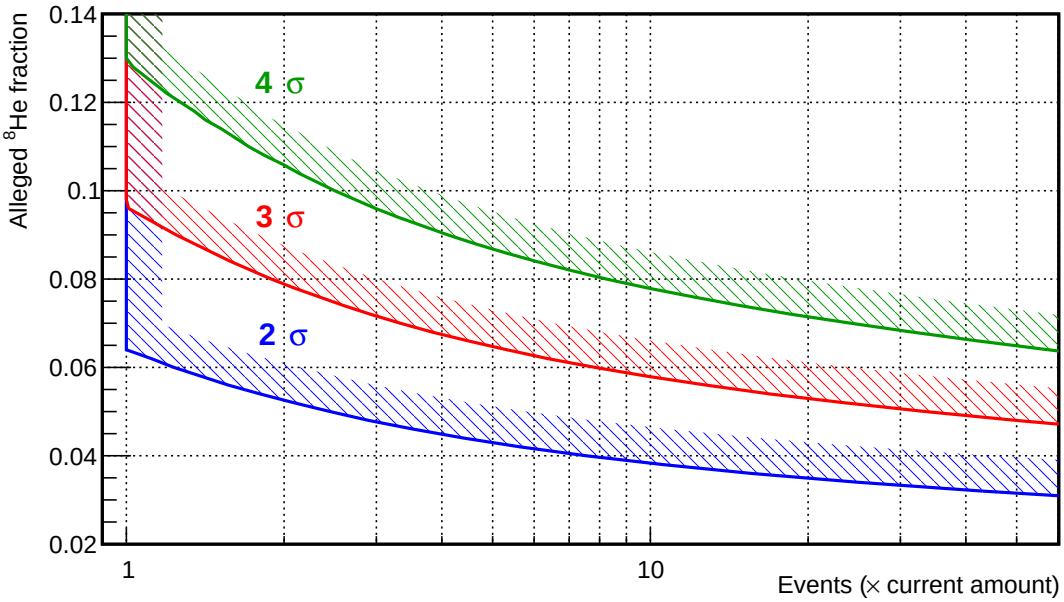

Figure 9.19 – ^{8}He fraction exclusion contours for different significance levels (see annotations). The x-axis represents the number of events required to observe the ^{8}He fractions on the y-axis. The number of events is presented in the "current amount" basis, i.e. it indicates how many more times we need the current Neutrino Gd++ combined statistics to make out the considered α.

Chapter 10

Rate estimations

After all the vetoes have been applied to the selected IBD pairs, a handful of βn decays still remains within the sample of candidates. Via the cosmogenic veto, or the Monte-Carlo predictions, we have a measure of the energy shape of this background, which can be utilised to evaluate the contamination remaining in the $\overline{\nu_e}$ candidate set by means of the oscillation fit. Nonetheless, biases in the spectra, be it in the cosmogenic shapes or the $\overline{\nu_e}$ one, could hinder the reliability of this estimation. As a consequence, it is apposite to challenge the remaining cosmogenic rate estimation by a dedicated analysis, relying on the study of the decay time profiles.

10.1 Some estimation methods

10.1.1 Principle

To count the number of βn decays hiding in the $\overline{\nu_e}$ candidates, we can take advantage of the characteristic exponential decay of the cosmogenic isotopes; for indeed, the lifetimes of ^{8}He and ^{9}Li, $\tau_{\mathrm{He}} = 172\,\mathrm{ms}$ and $\tau_{\mathrm{Li}} = 257\,\mathrm{ms}$, are well-known.

Although we do not start off with a radioactive sample of ^{9}Li nuclei to put in front of a counter, all the βn decays throughout the candidate set share the same production mechanism, i.e. their birth date corresponds to the μ-trigger that spawned them. By collecting all the βn prompt triggers t_p and the time intervals $\Delta t_{\mu-p} = t_p - t_\mu$ that separate them from their parent μ's, we define, in effect, a sample obeying the law of exponential decay. A simple integration of the exponential matching best the data thus collected should consequently provide the number of βn-emitters N_{co} in the considered sample. From there, it is essentially a matter of dividing N_{co} by the corresponding live-time to find the rate of βn-decays.

Assuredly, the cosmogenic decays are not easily spotted, neither are the parent μ's. Therefore, so as not to miss any decay, we must plot the time differences between all prompts and all μ's. In all truth, there is, a priori, no need to look for parent μ's such that $t_{\mu-p} < 0$.

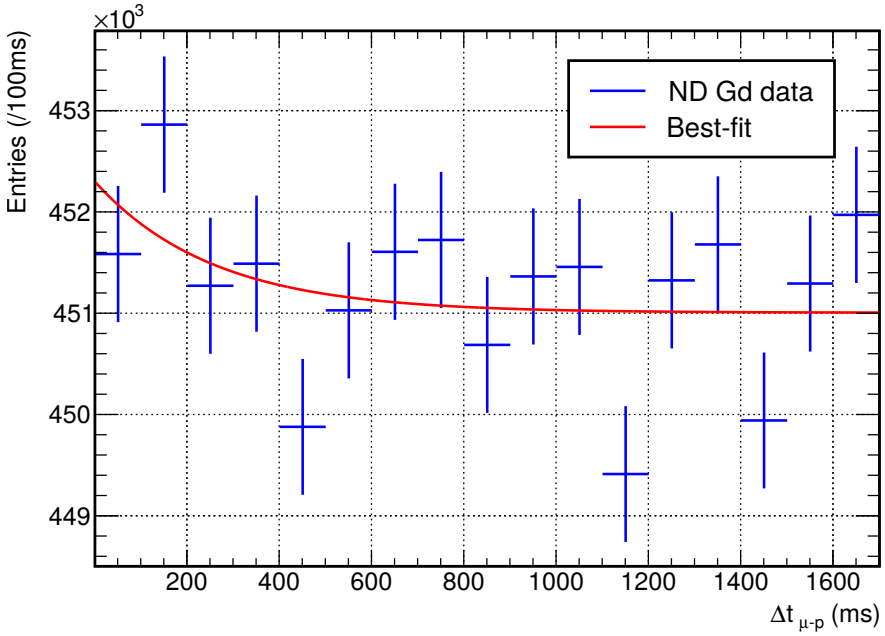

Figure 10.1 – Distribution of the time intervals $\Delta t_{\mu-p}$ for the ND Gd Neutrino data set (blue). The sum of an exponential with a fixed lifetime τ_{Li} and a first degree polynomial, at best-fit, is also shown (red).

By the same token, μ's preceding prompts by more than $5\tau_{\mathrm{Li}}$ cannot account for more than 0.7% of the dominating ^{9}Li decays (see 9.3.2.3). With these requirements, all the more so at the ND, overwhelmed by both cosmic μ's and $\overline{\nu_e}$'s, there are still many combinations, and thus, a substantial amount of noise. If the FDI Gd data set was somewhat clean and large enough to return a vague number of cosmogenics when considering all events, the brute-force approach is doomed at ND, as Figure 10.1 exemplifies. One way or another, we must reduce the number of combinations, whilst retaining the ability to tell the number of decays missed by the reduction.

10.1.2 Maximum posterior probability approach

To cull the number of combinations, we benefit from what appears to be an appropriate tool: the cosmogenic veto. More precisely, the maximum posterior probability provides a unique prompt - μ association, of which we took advantage to retrieve the energy distribution of the βn decays (cf. 9.2.1.3). As with all the cuts that might be contemplated, the thorniest part amounts to assessing the number of cosmogenic decays overlooked by the selection.

How better to assess the efficiency corresponding to a $P_{max}\,(co)$ cut (cf. 9.1.2.2) than by building up mock data, containing $\overline{\nu_e}$'s and cosmogenics, in proportions realistically chosen? The mock data must also contain μ's, or rather $t_{\mu-p}$'s values for each prompt, drawn from well-suited distributions (exponential for the signal, approximately flat for the background).

Furthermore, for every prompt, a maximum posterior probability must be pulled.

There is the rub, below $P^{th}_{max} = 0.02$, the number of accidental coincidences skyrockets; performing a background subtraction to retrieve the signal from an on-time window, as we did for the neutron multiplicity in 9.1.4.3, is fruitless. In fact, even at FDI, where the signal over background ratio is encouraging, a negative content – with a gigantic error bar – is retrieved for the bin corresponding to $P_{max}\,(co) \in \left[0; P^{th}_{max}\right[$. Even when picking bin contents within this error bar and retaining only the positive ones, the toy Monte-Carlo suggests that up to 40% of the cosmogenics could have a maximal posterior probability lower than 0.02 [83]. To be sure, the toy Monte-Carlo also claims that, on average, only 6% of the βn decays should lie below the 0.02 threshold.

To cut the matter short, utilising the cosmogenic veto to perform a rate estimation is delicate, and it does requires a leap of faith, to some extent. Besides, at ND, the performance of the veto is even more challenging. Nonetheless, it provides an interesting cross-check of the other method used at FDI for [60], summarised in 10.1.3. The posterior probability approach yields a total cosmogenic rate at FDI for the Gd channel [83]

$$R^{FDI}_{prob} = 2.28^{+0.38}_{-0.28}\,. \tag{10.1}$$

Please note that all the cosmogenic rate estimations are better performed before the application of the cosmogenic veto, in such a way that the sought-after needle stands out of the haystack a little more, hence the "total" qualifier. If need be, the vetoed rate can be effortlessly subtracted from the total rate, to allow comparisons with the output of the oscillation fit, for instance.

10.1.3 Lateral distance approach

To increase the signal over background ratio when estimating the cosmogenic rate in the data, we may also cut on the distance between prompts and μ-tracks $d_{p-\mu}$.

From FLUKA simulations, two characteristic production lengths are expected for each cosmogenic isotope, one of the order of a few centimetres, the other larger than 50 cm [141]. The larger component is supposed to account for more than 85% of the decays. In accordance, and for simplicity, a single exponential decaying radially from the parent μ-track, with a characteristic distance λ, was used for [60].

10.1.3.1 Envelope functions

Asserting that the distance between cosmogenic decays and parent μ's follows the sum of exponentials is a valid statement only if the detector is infinite. Indeed, in finite detectors designed like the Double Chooz ones, low energy μ-deposits correspond to tracks passing through the buffer and clipping the GC, for which βn-emitters might be produced though few may actually be observed, since the buffer is inert. On the contrary, high energies are deposited by μ's passing right through the NT. Therefore, the average distance at which

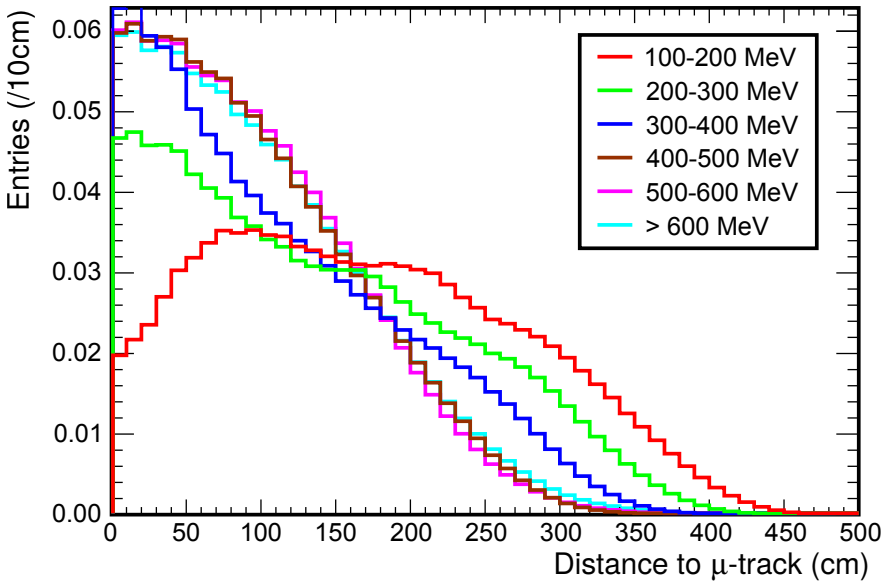

Figure 10.2 – Envelope functions for different FDI μ energies (see labels), obtained from uncorrelated prompt - μ pairs in the H analysis. Curves from [141].

cosmogenics are detected depends on the energy deposited by μ's. There is more to the latter statement than meets the eye, the exponential part of the distribution – valid in an infinite detector – is assumed not to depend on the μ energy, at all, an hypothesis supported by the weak overburden dependency observed in 5.2.4.3. As a result, the only energy dependency of the model is embedded in the so-called envelope functions, which allegedly account for the geometry of the detector.

Considering the envelope functions are energy-dependent, several sub-samples of μ's were considered for the FDI data. At this stage, λ has not yet been evaluated from the data, thus, we cannot estimate the envelope functions from correlated prompt-μ pairs. On the other hand, when it comes to prompts uncorrelated to μ's, such as IBD's, their number has a linear radial dependency around each μ-track, in an infinite detector. For all types of events, it is presupposed that the distance distribution observed is the product of the envelope function and the distribution for an infinite detector. Dividing the observation by the later provides the envelopes. Examples of them, taken from the FDI H analysis, are shown for informative purposes in Figure 10.2.

10.1.3.2 Characteristic production length estimation

Assuming that these envelopes only represent the finite size of the detector, they should also be contained within the distribution of the distance between cosmogenic decays and μ-tracks.

The highest energy sample of μ's has a sufficiently high signal over background ratio to successfully perform a subtraction of the accidental background within it, thereby yielding

the distance distribution for the cosmogenics at this μ energy. This distance profile can be compared to several distributions, obtained by multiplying the envelope of the highest energy sample by exponentials with varying λ values. This allows to deduce the best matching λ, which unequivocally represents the distribution in an infinite detector. Additionally, the Gaussian resolution of the μ-track reconstruction may be taken into account as a free σ parameter. The comparison between the prediction and the background subtracted distance distribution, from an on-time window, for the highest μ energies, is presented in Figure 10.3. By means of this best matching prediction, a characteristic production length $\lambda = (49.1 \pm 10.5)$ cm, supposedly valid for all μ energies, is estimated.

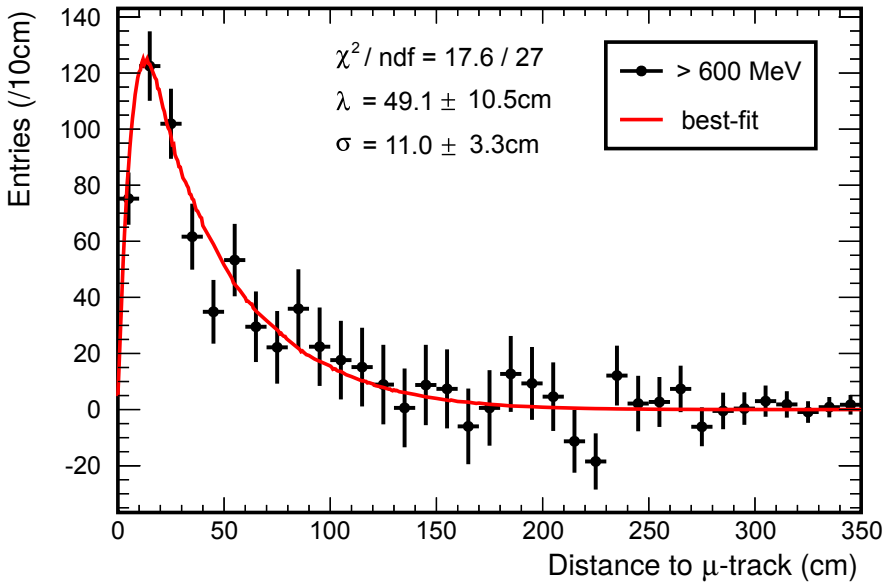

Figure 10.3 – Cosmogenic distance distribution for the FDI H μ sample with the highest energy (black) and best matching prediction (red). Curves from [141].

10.1.3.3 Total cosmogenic rate and prospects

FDI textbook case

Once the envelopes, as well as the characteristic production length λ, are known, the signal for lower energy μ's may be deduced. Let us emphasise that in Figures 10.2 and 10.3, the μ energies are not transposable to current studies, for the energy reconstruction now applied to the FDI data set has dramatically changed. Besides, the LNL correction (cf. 5.2.3.4) distorts the shapes of the spectra, hence, a simple matching of the most probable μ energies would not suffice.

From the cosmogenic distance distributions of all the FDI Gd μ sub-samples, the efficiency of distance-based cuts can be estimated with ease. This, however, entails estimating a cosmogenic rate for each μ sub-sample, before summing them all. To cut the matter short,

applying distance cuts allows to infer a total cosmogenic rate [60]

$$R_{dist}^{FDI} = 2.08^{+0.41}_{-0.15} \tag{10.2}$$

for the FDI Gd data, which is consistent with (10.1). The strongly asymmetric errors in (10.2) stem from the fact that it is simpler to apply a cut and obtain a precise rate measurement, than estimate how much you are missing.

ND extrapolation

With regard to the ND rate estimation, for the Moriond conference, the μ-track reconstruction issues were only partially solved (the time calibration was missing), which made relying on distance cuts unattractive. This lead to the development of alternative methods, reviewed in 10.2 and 10.3. Even in between the Moriond and Neutrino 2016 conferences, disregarding the stringent time constraints, the considerable discrepancies between the μ energies of the three detector versions would have complicated comparisons of dedicated analyses, even between FDI and FDII.

For all the validity of these reasons, the actual show-stopper at ND was the tremendous amount of noise, i.e. of $\overline{\nu_e}$'s; even with the plentiful ND Gd++ Neutrino set and shrewd cuts, the ND background subtracted distance density is currently worthless, no λ value may be extracted from it. Nevertheless, let us recall that after Moriond, DCRecoMuHam's performance hinted at a moderate overburden dependency of the mean spallation neutron path. As a result, the former FDI λ value might be good enough to correct distance-based cuts at ND, thereby removing the need for obtaining signal distributions. Although they are much more accessible via IBD's, we would still need to find the envelopes for μ sub-samples to define. In any case, if they do provide meaningful results at FD, distance-based cuts do not shrink enough the number of prompt-μ pairs to produce usable rates at ND, even with the Neutrino data sets considered throughout this chapter.

10.2 Muon sample cleansing

The two methods from 10.1.2 and 10.1.3 focused on reducing the number of μ's associated to each prompt candidate; it is not without reminding them that we chose to select potential βn producers via neutron multiplicity cuts.

10.2.1 Analysis method

10.2.1.1 Neutron multiplicity threshold

The chosen multiplicity threshold is minimal, when plotting the distribution of $\Delta t_{\mu-p}$'s, we require μ's satisfying $n_\mu \geqslant 1$, with n_μ the number of spallation neutrons generated by each μ, following the multiplicity definition from 9.1.2.1. The goal of this threshold is to

curb the number of accidental coincidences, yet without introducing too hard a cut, whose efficiency would be difficult to correct. Throughout this section, we will not try to correct the efficiency introduced by the multiplicity cut, this step will be performed by means of independent measurements and discussed in 10.3.

Not only requiring at least one neutron to be produced by all μ's we take into account shrinks the number of combinations by about 10^2, in accordance with the multiplicity mass functions from Figure 9.3, but it is also a threshold immune to any miscalibration of the μ energies.

10.2.1.2 Fitting the time intervals

We can find the number of cosmogenic decays N_{co} brought forth by our cut by computing the integral of the exponential law, which the time intervals $\Delta t_{\mu-p}$'s for correlated pairs must follow. The remaining accidental coincidences N_{acc} are estimated by letting their number multiply the μ-rate r_μ and the time bin width $T = 100\,\text{ms}$, in accordance with (9.7). The μ-rate is recomputed within every analysis, thus taking into account possible purifications of the μ-sample.

The fit function used throughout this chapter is

$$f\left(\Delta t_{\mu-p}\right) = T\left(\frac{N_{co}}{\tau_{Li}}\,\exp\!\left(-\frac{\Delta t_{\mu-p}}{\tau_{\text{Li}}}\right) + N_{acc}\,r_\mu - S^2\,\Delta t_{\mu-p}\right), \tag{10.3}$$

where S – expressed in Hz – accounts for the mild inclination of the accidental contribution. Its value, returned by the fit, is usually of the order of 1 Hz.

The sole presence of τ_{Li} in (10.3) is motivated by the results from 9.3.2.3, which excluded any significant presence of ^{8}He. Inasmuch as the β-decays of the abundant ^{12}B are also correlated to μ-triggers, the $\tau_{\text{B}} = 29.1\,\text{ms}$ life-time of the β-emitter could show through. However, as can be checked on the vetoed spectrum from Figure 9.8, in spite of its $Q_\beta = 13.4$ MeV endpoint, no events are observed above the 12.0 MeV ^{9}Li endpoint. This is all the more convincing that the cosmogenic veto was built onto the ^{12}B probability densities, thereby augmenting the chances to tag the β-emitter. Although Figure 9.8 was a FD example, ND shows no sign of prompt ^{12}B events either.

The aforementioned arguments have only excluded ^{12}B as a prompt event, and a β-decay might find itself in random coincidence with a prompt single. In such a case, although the single would not be truly correlated to the parent μ of the ^{12}B nucleus, the smallness of the neutron-capture times $\tau_{\text{H}} > \tau_{\text{Gd}}$, with respect to τ_{B}, would effectively correlate the prompt single to the μ. This background is evaluated to be smaller than 2% of the βn signal [142], i.e. well below the error bars we shall claim.

10.2.1.3 After-μ veto

By definition, the number of cosmogenics N_{co} represents the sum of all the decays that may happen between $\Delta t_{\mu-p} = 0\,\text{ms}$ and $\Delta t_{\mu-p} = k\ \tau_{Li}$ with $k \to +\infty$. However, by virtue of

the after-μ veto, a small fraction of cosmogenic decays is naturally removed from the IBD candidates.

More practically, about 0.39% of the ^{9}Li decays in the Gd analysis, and 0.49% in analyses utilising H captures, are discarded by the after-μ veto. Consequently, it is apposite to lower the N_{co} values by these minute corrections. Certainly, these are not paramount, yet they are easily applied.

10.2.2 Results

10.2.2.1 Far detector

Utilising the neutron multiplicity threshold introduced in 10.2.1.1, along with the fit function from (10.3), we can quantify the exponentially decaying part of the $\Delta t_{\mu-p}$ distribution.

As an illustration, the $\Delta t_{\mu-p}$ distribution for the FDI Gd++ data sub-set including showering μ's ($n_\mu \geqslant 1$) only, is shown and fitted in Figure 10.4. The best fit corresponds to a number of cosmogenic decays $N_{co} = 1955 \pm 75$ and a reduced chi-square $\chi^2/n_{df} = 105.6/97$, which does not cast notable doubts on the fit model (10.3).

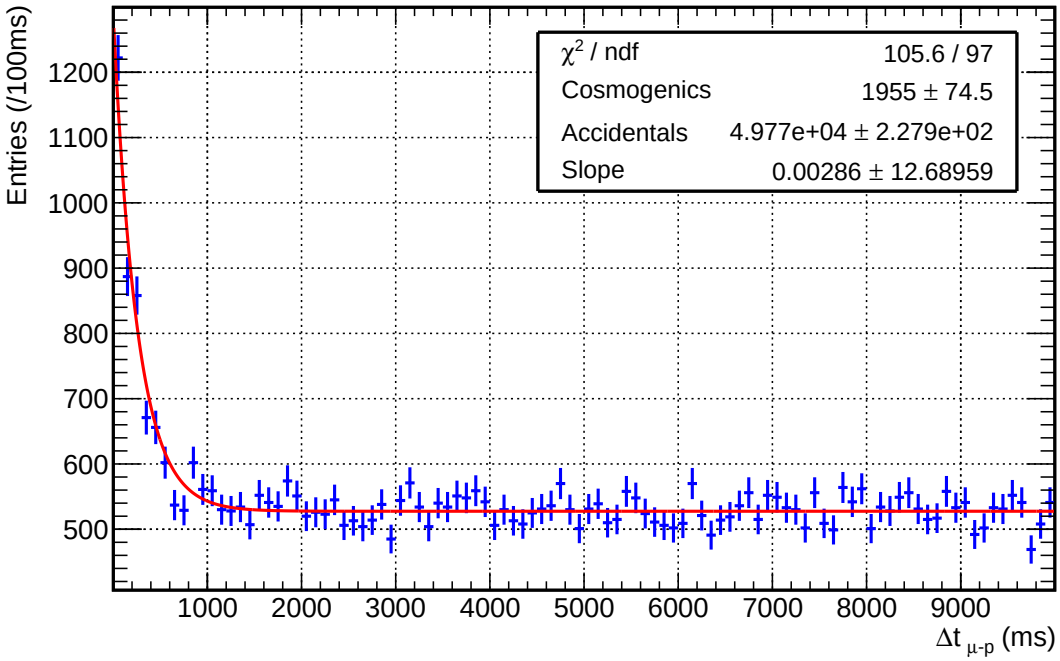

Figure 10.4 – Distribution of the time intervals for the FDI Gd++ data. A fit of this binned distribution yields $N_{co} = 1955 \pm 75$ cosmogenic decays. The number of accidentals N_{acc} and the slope S (in Hz) are also indicated.

In particular for the plentiful Gd++ set, finer binnings, such $T = 20\,\text{ms}$ or $T = 50\,\text{ms}$ were tested; these did not seriously call into question the plausibility of the fit model, nor the accuracy of the measured rate. For instance, keeping the FDI Gd++ data set as an example,

a $T = 20\,\text{ms}$ binning yields $N_{co} = 1928 \pm 73$ for a reduced chi-square $\chi^2/n_{df} = 501/497$, which is in satisfactory agreement with the results obtained with a $T = 100\,\text{ms}$ binning in Figure 10.4.

Likewise, excluding the first 60 ms of the fit, which would discard 87% of a potential ^{12}B component, returns a $N_{co} = 1792 \pm 93$ value, along with a $\chi^2/n_{df} = 492/494$. Let us underscore that all the N_{co} values we have presented so far are directly comparable. Indeed, however small a portion of time intervals we consider, by definition, N_{co} corresponds to the integral of the exponential over all $\Delta t_{\mu-p}$ values.

10.2.2.2 Near detector

In accordance with our preliminary remarks from 10.1.1, the ND data are much noisier; demanding μ's satisfying $n_\mu \geqslant 1$ will not produce a distribution as clean the FD histograms. The ND distribution and the corresponding best fit are given in Figure 10.5.

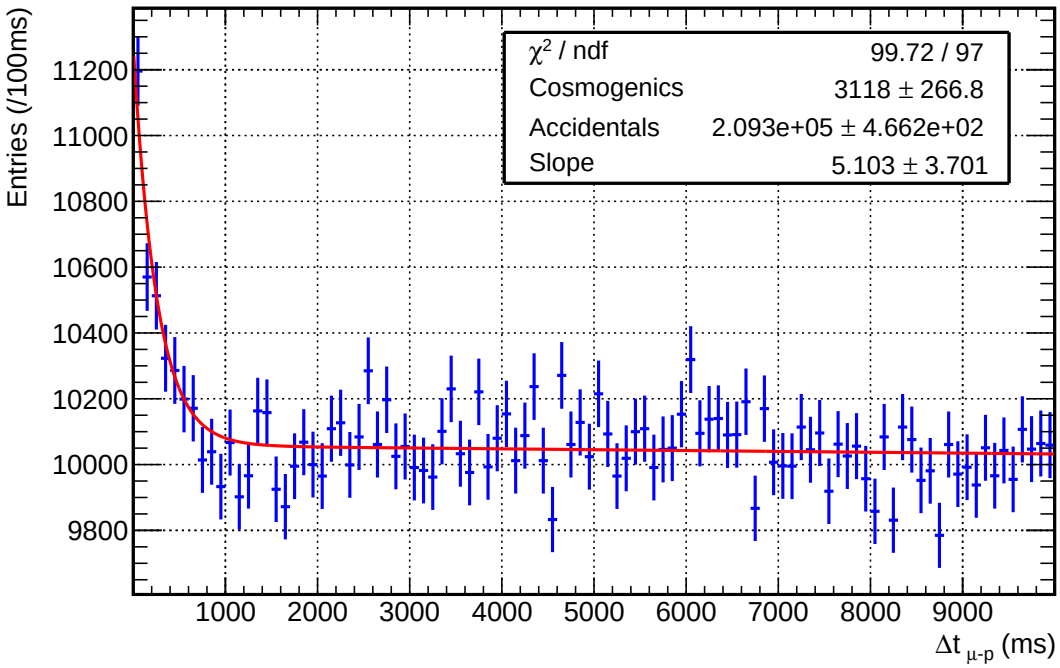

Figure 10.5 – Distribution of the time intervals for the ND Gd++ data. Fitting the distribution with (10.3) yields $N_{co} = 3118 \pm 267$ cosmogenic decays.

Compared to the nine-month Moriond configuration, for which a 17.2% relative uncertainty on the cosmogenic rate for $n_\mu \geqslant 1$ was obtained, the fifteen months of ND Gd++ data for the Neutrino conference already provide 8.4% relative error on the same rate. This improvement is scarcely surprising considering the live-time has increased by $\times 1.71$ between the Moriond and Neutrino conferences (cf. 9.2.1.3) whilst the Gd++ rate is $r_{++}^{ND} \simeq 2.4$ that of the Gd one (cf. 9.12). To put it differently, the relative error should have decreased by

about $\sqrt{4.1} \simeq 2$, which flawlessly matches the observation.

10.2.2.3 Rates summary

Rather than including three plots per analysis type, we summarise the so-called $1n$ rates, i.e. for which only μ's satisfying $n_\mu \geqslant 1$ are included, for all six configurations under Tables 10.1 and 10.2.

All the numbers within Table 10.1 are in accordance with the amounts of statistics they represent and the corresponding signal over background ratios. In other words, the FDII rate is a little less accurate than its FDI counterpart while the ND itself fares considerably worse. By the same token, the FDII and FDI rates are convincingly close.

Version	N_{co}^{1n}	r_{co}^{1n} (d^{-1})
FDI	783 (6.1%)	1.69 ± 0.10
FDII	611 (6.7%)	1.66 ± 0.11
ND	1208 (14.1%)	4.40 ± 0.62

Table 10.1 – Number of cosmogenics N_{co}^{1n} (relative uncertainty in parentheses) produced by showering μ's and corresponding daily rate r_{co}^{1n}, for the Neutrino Gd data sets.

The comments made for the Gd data also apply to the Gd++ data from Table 10.2. The FDII and FDI rates are slightly more apart but nothing of statistical significance is to be spotted yet. Similarly, the ratios of the Gd++ rates to the Gd ones are not incompatible with the ratios of vetoed rates from 9.2.3.2.

Version	N_{co}^{1n}	r_{co}^{1n} (d^{-1})
FDI	1891 (3.8%)	4.13 ± 0.16
FDII	1588 (4.3%)	4.35 ± 0.19
ND	3118 (8.6%)	12.03 ± 1.03

Table 10.2 – Number of cosmogenics N_{co}^{1n} (relative uncertainty in parentheses) produced by showering μ's and corresponding daily rate r_{co}^{1n}, for the Neutrino Gd++ data sets.

10.3 Candidate sample cleansing

Aside, perhaps, from the ND (Gd) case, all the rates obtained by way of the multiplicity cut are quite competitive in terms of accuracy. However, we did purposefully overlook the stumbling block the efficiency assessment embodies, which we propose to address throughout this section, via an independent method relying on an increase of the cosmogenics over $\overline{\nu_e}$'s ratio.

10.3.1 Method

10.3.1.1 Motivations and past achievements

In order to evaluate how many events we are missing by requiring showering μ's at FD and ND, we must perform independent measurements of the total cosmogenic rates. In all truth, these independent measurements will stand as the final FD rates while providing the alleged efficiency correction for ND. Let us stress that simulations are of no help to estimate the efficiency correction since the various simulation models produce substantially different cosmogenic production yields [83].

Theoretically, we could use the cosmogenic neutron multiplicity mass functions P_{co}^{N} from 9.1.4.3 and focus on their $P_{co}^{N}(0)$ values. If relying on the ^{12}B reference densities in the cosmogenic veto could only lead to poorer vetoing efficiencies or inefficiencies, which were always measurable, it seems wiser to derive corrections from the βn-emitters themselves, to assess the total cosmogenic rates. Unfortunately, in the same way that producing a mass function from an on-time window with a meaningful $n_\mu = 0$ bin content is delicate, estimating the efficiency of a $n_\mu \geqslant 1$ cut is intricate.

As a matter of fact, this was precisely by performing a so-called candidate sample purification, namely an increase of the cosmogenics over $\overline{\nu_e}$'s ratio, that we could get the only usable non-negative $P_{co}^{N}(0)$ value from βn-emitters. The latter, obtained from the largest and most favourable (see prior ratios from 9.1.3.2) FDI Gd++ set, measured $P_{co}^{N}(0) = 0.33 \pm 0.14$. It follows that a preliminary efficiency correction for the multiplicity cut is

$$C_{pre} = 1.49 \pm 0.31\,. \tag{10.4}$$

The (10.4) correction could be obtained by means of a prompt energy cut $E^{vis} > 6\,\mathrm{MeV}$. Not only do prompt energy cuts increase the signal over background ratio in our studies – because there are essentially no $\overline{\nu_e}$'s above 8 MeV, as the spectra from Figure 9.7 epitomise – but since the βn-decays are performed at rest, they do not bear the neutron multiplicity stamp of their parent μ.

10.3.1.2 Monte-Carlo correction

In light of the robustness of the Monte-Carlo predictions of the cosmogenic spectra from 8.4, it is sensible to consider a prompt energy cut as a means of obtaining the total cosmogenic rate, which would benchmark further the Monte-Carlo against the vetoed spectrum. Besides, the only helpful value, the (10.4) correction, came from a Gd++ set. Although there is no evidence that the Gd and Gd++ corrections significantly differ, it is apposite to validate this assumption. Therefore, we set forth an alternative method, providing the total cosmogenic rates, by way of well-suited efficiency corrections taken from the Monte-Carlo spectra.

Total rate

Inasmuch as there is no evidence of ^{8}He yet, we only consider the ^{9}Li Gd and H spectra. To

motivate this decision further, let us emphasise that the statistical uncertainty still largely prevails over systematics.

To obtain the total rate r_{co}, we divide the reduced rate $r_{co}(E_{th})$, obtained by demanding only prompts satisfying $E^{vis} > E_{th}$, by the integral of the prediction S between E_{th} and a value well past the ^{9}Li endpoint, i.e. $E_{max} = 14\,\text{MeV}$. Considering that S is normalised to unity between 0 MeV and E_{max}, dividing $r_{co}(E_{th})$ by $\int_{E_{th}}^{E_{max}} S(E)\,\mathrm{d}E$ effectively brings the reduced rate to the full energy range. Since the Gd and Gd++ analyses have themselves an energy threshold, denoted by $E_{min} = 0.5\,\text{MeV}$ (Gd) or $E_{min} = 1\,\text{MeV}$ (Gd++), an additional correction[1] – of a mere 0.48% for Gd and 1.66% for Gd++ – must be made. As a result, the total cosmogenic rate reads

$$r_{co} = \frac{r_{co}(E_{th}) \displaystyle\int_{E_{min}}^{E_{max}} S(E')\,\mathrm{d}E'}{\displaystyle\int_{E_{th}}^{E_{max}} S(E')\,\mathrm{d}E'}\,. \tag{10.5}$$

Sliding integration

The values of all the integrals appearing in (10.5) can be obtained from a function $E \to \int_{E}^{E_{max}} S(E')\,\mathrm{d}E'$. Such functions, for the Gd and H simulations are shown in Figure 10.6.

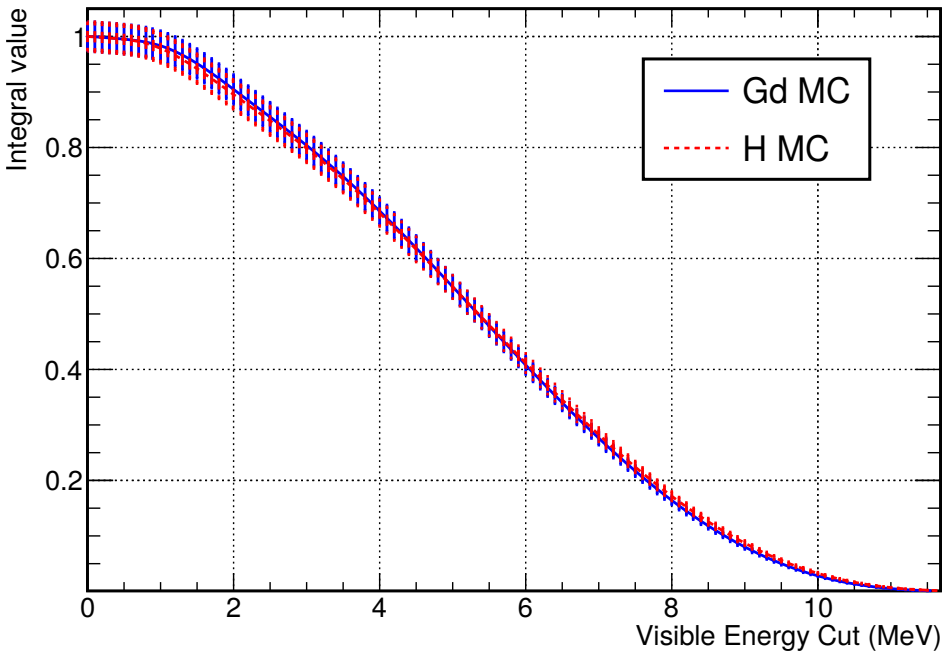

Figure 10.6 – Graphs of the functions $E \to \int_{E}^{E_{max}} S(E')\,\mathrm{d}E'$ with S being either the Gd Monte-Carlo spectrum (blue) or the H one (dashed red).

[1] To be sure, we could normalise the predictions between 0.5 MeV and E_{max} to 1 but we would still need to account for the different Gd++ threshold.

Obviously, when integrating the Gd or H spectra, only minute discrepancies arise; visually, both graphs from Figure 10.6 are nearly superimposed. Moreover, considering the Gd simulations appear slightly more reliable (see 9.3.2.2), the Gd spectrum remains the chosen one as from now on.

Although the uncertainties on the integrals of the Gd spectrum S are yet to be dominant, they can be included with ease via the Gd covariance matrix V from 8.4.2. Evidently, both S and V are discrete, thus, it is more fitting to use prompt energy cuts E_{th} based on bin boundaries. Let us denote by l the index of the bin whose left boundary corresponds to E_{th}, and by n the last bin of the discretised spectrum $(S_i)_{i\in[\![1;n]\!]}$. It follows that

$$\int_{E_{th}}^{E_{max}} S\left(E'\right) \mathrm{d}E' \simeq \sum_{i=l}^{n} S_i \, \Delta E_i \,, \tag{10.6}$$

with ΔE_i the width of the i-th bin[2]. The error on (10.6), to include when computing (10.5), is consequently

$$\sqrt{\mathbb{V}\left(\sum_{i=l}^{n} S_i \, \Delta E_i\right)} = \sqrt{\sum_{i=l}^{n}\sum_{j=l}^{n} \Delta E_i \, \Delta E_j \, V_{i,j}} \,, \tag{10.7}$$

where $\mathbb{V}$ denotes the variance operator. The sum in (10.7) is nothing more than a summation over all the matrix coefficients contained within the block spanning the integrated bins $l \to n$. As for the small correction $\int_{E_{min}}^{E_{max}} S\left(E'\right) \mathrm{d}E'$, we deem its uncertainty marginal, indeed, we may write

$$\int_{E_{min}}^{E_{max}} S\left(E'\right) \mathrm{d}E' = 1 - \int_{0}^{E_{min}} S\left(E'\right) \mathrm{d}E' \simeq 1 - \sum_{i=1}^{p} S_i \, \Delta E_i \,, \tag{10.8}$$

where the sum over the first few p bins bears an insignificant uncertainty.

10.3.2 Application

10.3.2.1 Total far detector rates

FDI Gd benchmark

We can start tinkering with prompt energy cuts with the least noisy set, namely the FDI one. Plainly, low energy thresholds will offer no improvement over a conventional estimation including all $\Delta t_{\mu-p}$ pairs. On the other hand, high energy cuts will leave very few βn-emitters. For a start, it is pointless to try cuts well-above 8 MeV's since past this value, $\overline{\nu_e}$'s are already a rarity. To systematically find a sweet-spot for the E_{th} cut, we can scan bin boundaries and retrieve the relative uncertainty on the rate estimation.

The relative uncertainty on the estimation of the Monte-Carlo corrected rate r_{co}, as a function of the prompt energy cut E_{th} applied to the FDI Gd $\overline{\nu_e}$ candidates, is presented in Figure 10.7. Only the energy range yielding acceptable uncertainties was kept for the plot.

[2] A bin content is usually expressed as a number over the corresponding ΔE_i, thus, in practice, we need but read the number itself since $1/\Delta E_i$ cancels out in the product.

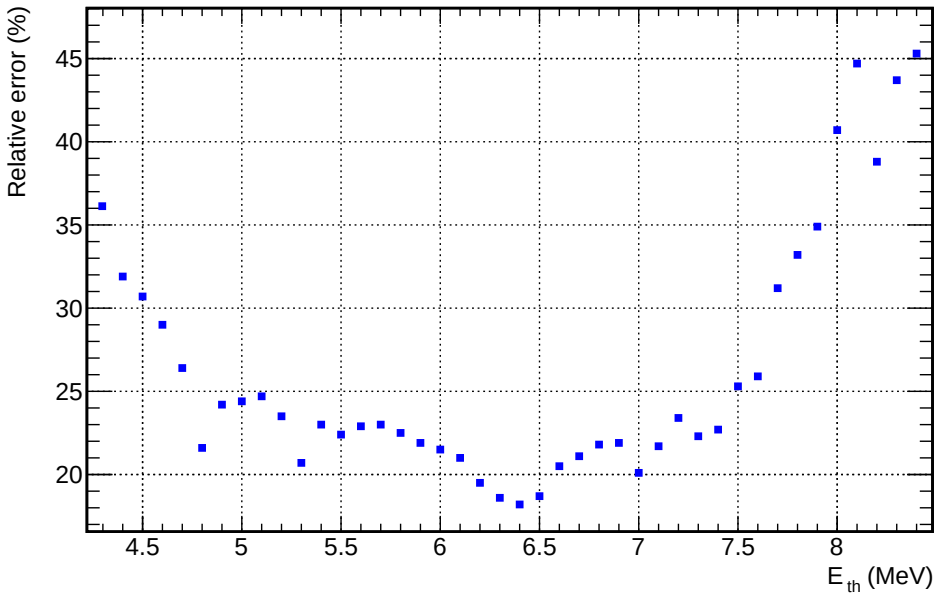

Figure 10.7 – Relative uncertainty on the Monte-Carlo corrected FDI Gd rate r_{co} as a function of the E_{th} cut applied to produce the reduced sample..

It appears that the best compromise between a good signal over background ratio and acceptable statistics is around 6.5 MeV, at least for the FDI Gd set. The few odd local minima, visible on Figure 10.7, come forth when the (10.3) fit returns N_{co} values slightly higher than the neighbouring points whilst delivering a common absolute error, hence the misleading increase of apparent significance. Lest we be biased towards high rates, it is well-advised to also draw the corrected cosmogenic rate from (10.5), as a function of the energy cut. To this intent, we select a range where the relative uncertainty is minimal, i.e. between $5 - 7.5$ MeV. Figure 10.8 reflects this choice.

Although the Monte-Carlo corrections vary by a factor larger than 2 across the energy range presented (cf. Figure 10.6), when the fit is stable, the corrected rate is remarkably flat. In a nutshell, however much E_{th} takes away, the Monte-Carlo brings it back. Please note that neighbouring data points in Figure 10.8 are strongly correlated, consequently, they have no reason to be randomly distributed above and below their mean. By the same token, trying to combine their relative uncertainties, which are of comparable order across the range, will not reduce the error on the rate measurement via this method. Nonetheless, we may take the plain averages of the rate and relative uncertainty over this range, a quite conservative approach, to estimate $r_{co}^{FDI,Gd} = (2.32 \pm 0.50)\,\mathrm{d}^{-1}$.

All FD sets

We can proceed similarly for the other far detector sets. Regarding the FDII Gd one, its meagre statistics are best utilised with a prompt energy cut between 7.4 MeV and 7.8 MeV. For the FDI Gd++ set, we retain the $6 - 6.5$ MeV range, while the $6.5 - 6.9$ MeV one is

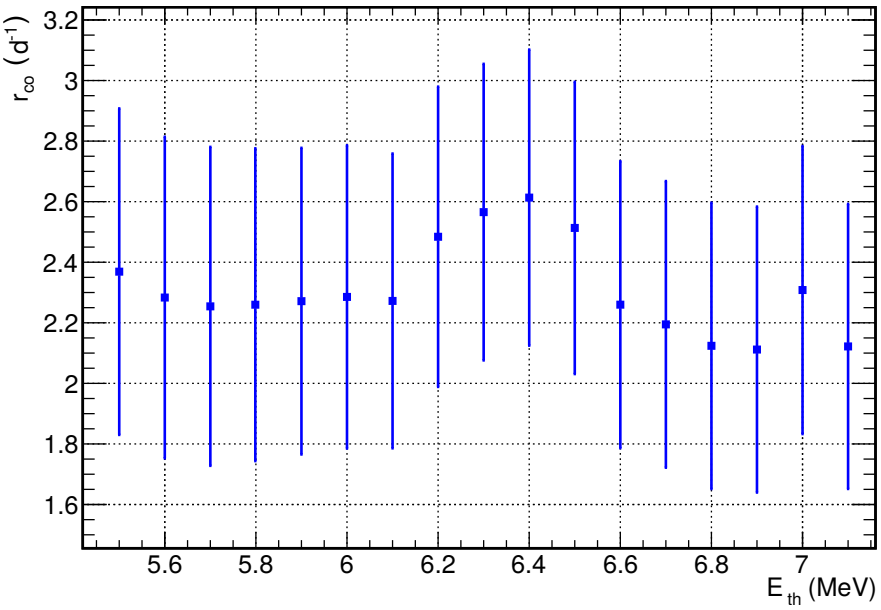

Figure 10.8 – Efficiency corrected cosmogenic rate r_{co} as a function of the E_{th} cut applied to produce the reduced FDI Gd sample.

more apt for the FDII Gd++ data. Not unexpectedly, the larger the sample, the smaller the energy cut to be able to sort the wheat from the chaff.

All the far detector rates are summarised in Table 10.3. Although the systematic errors are included in Table 10.3, the rate uncertainties are practically dominated by statistical fluctuations. As is visible, the FDII Gd rate is lower than the FDI Gd one, but both bear standard deviations compatible with the observed discrepancy. Please note that considering another energy threshold for the FDII Gd set would not have lead to different conclusions. Furthermore, the FDII Gd++ rate, which effectively contains the FDII Gd one, is reassuringly closer to the FDI Gd++ result, thereby supporting a downwards statistical fluctuation in the scant FDII Gd data. Besides, we can compute the ratios of the Gd++ to Gd rates, as we did for the number of IBD's and vetoed rates. So as not to overestimate the errors and brag about the compatibility of the measurements, we took a $\rho = \sqrt{f_{Gd}} = \sqrt{0.36} = 0.6$ correlation coefficient[3] (as a reminder, f_{Gd} is the Gd fraction within the Gd++ set) when computing the uncertainty[4] on the following values

$$r_{++}^{co,FDI} = 2.67 \pm 0.46\,, \tag{10.9}$$

$$r_{++}^{co,FDII} = 3.70 \pm 1.07\,. \tag{10.10}$$

Whereas the central value from (10.9) is in flawless agreement with the IBD and vetoed rate ratios from 9.1.3.2 and 9.2.3.2, (10.10) is definitely high, albeit consistent within 1σ

[3] Denoting by X and Y the Gd and H Poisson counts, $\rho_{X,Y} = \frac{\mathbb{C}[X+Y,Y]}{\sqrt{\mathbb{V}[X+Y]\mathbb{V}[Y]}} = \sqrt{\frac{\mathbb{V}[Y]}{\mathbb{V}[X]+\mathbb{V}[Y]}} = \sqrt{\frac{\mathbb{E}[Y]}{\mathbb{E}[X]+\mathbb{E}[Y]}}$.

[4] The formula for the uncertainty derives from the (9.39) general Taylor expansion with $f(X,Y) = X/Y$.

Version	r_{co} (d^{-1})
FDI Gd	2.32 ± 0.50
FDII Gd	1.46 ± 0.52
FDI Gd++	6.19 ± 0.80
FDII Gd++	5.40 ± 0.87

Table 10.3 – Efficiency corrected cosmogenic rate r_{co} for the FD samples indicated in the first column.

with its counterparts.

Last but not least, the two far detector Gd r_{co} values are in good agreement with past studies, whose results were reviewed in 10.1.2 and 10.1.3, thus corroborating the reliability of our energy corrections.

10.3.2.2 Efficiency of the multiplicity cut

Strategy

By comparing the rates from 10.2.2.3 and 10.3.2.1, we may deduce an efficiency correction for the multiplicity cut.

The reader would not have missed to notice that there are no ND rates for the energy based studies; indeed, the raw ND data are too overwhelmed by $\overline{\nu_e}$'s to obtain meaningful values by way of a sole E_{th} cut. Accordingly, to boost the so-called 1n ND rates, we will use an efficiency correction based on the independent sample purification performed at FD. As stated in 9.1.4.3, the neutron multiplicity mass functions for the cosmogenics are acutely difficult to bring out. Unquestionably, with infinite statistics, we would be confronted with different FD and ND mass functions, and thus, efficiency corrections for the $n_\mu \geqslant 1$ threshold, yet at the moment, there is no evidence that they differ in a relevant way. In consequence, and for lack of a better alternative, we shall rely on the FD correction, whose own uncertainty ought to encompass our ignorance.

Correction values

The values of the multiplicity corrections are literally the ratios of the total FD rates, for the four configurations, over their corresponding 1n rates from 10.2.2.3. Theoretically, the total rates and the 1n rates share a small fraction of IBD candidates[5], as well as a few μ's. Be that as it may, regarding these measurements as independent values can only overestimate the uncertainty on the correction, a particularly fitting approach.

The multiplicity corrections C, by which to multiply the 1n rates, are presented in Table 10.4. Not unexpectedly, the FDII Gd correction factor is smaller than 1, which is not physical,

[5]To obtain the total rate, we discard more than half of the βn-emitters and nearly all $\overline{\nu_e}$'s, so the overlap of candidates is by no means huge.

Version	C
FDI Gd	1.38 ± 0.31
FDII Gd	0.88 ± 0.32
FDI Gd++	1.50 ± 0.20
FDII Gd++	1.24 ± 0.21

Table 10.4 – Efficiency corrections C for the multiplicity cut $n_\mu \geqslant 1$, based on the different data sets indicated in the first column.

although the error bar does not prevent the value from being more sensible. Otherwise, all the correction factors are consistent with one another. Due to bordering detection effects for neutrons[6], there might be some subtle differences between the Gd++ and Gd mass functions, although as Table 10.4 and Figure 9.4 testify, they are currently beyond our statistical power. As a consequence, one correction factor will be derived, for all analyses, based on the two independent FDI and FDII Gd++ data sets. It is worth making note of the fact that combining the Gd measurements with the Gd++ ones would but introduce correlations while adding practically no value.

Combination

To combine the Gd++ far measurements, the different significance of the FDI and FDII sets must be taken into account. Not without reminding DCRecoMuHam, which weights its PMT time shifts by the number of observed photo-electrons (cf. 5.2.4.2), we shall weight the corrections C_δ by the inverse of their variances $\sigma^2_{C_\delta}$, with δ running over the considered sets. Additionally, choosing these weights ensures that the computed value is a maximum likelihood estimate of the mean. The average multiplicity efficiency correction reads

$$C = \frac{C_1\ \sigma_{C_1}^{-2} + C_2\ \sigma_{C_2}^{-2}}{\sigma_{C_1}^{-2} + \sigma_{C_2}^{-2}}\,, \tag{10.11}$$

with the 1 and 2 indexes standing for the FDI Gd++ and FDII Gd++ samples, in no particular order. The uncertainty associated to this mean correction is

$$\sigma_C = \frac{1}{\sqrt{\sigma_{C_1}^{-2} + \sigma_{C_2}^{-2}}}\,. \tag{10.12}$$

A numerical applications yields

$$C = 1.38 \pm 0.14\,. \tag{10.13}$$

[6]Neutrons produced by μ's clipping the GC may leave unoticed whereas the parent μ of a prompt in the NT should see its neutrons efficiently captured on H or Gd, both being considered in the multiplicity definition.

Insofar as $\sigma_1 \simeq \sigma_2$, the industrious reader may have noticed that $\sigma_C \simeq \sigma_1/\sqrt{2}$ and the mean C is only ever so slightly different from a plain average. The efficiency correction (10.13) is convincingly close to the preliminary (10.4) value, derived from the sole FDI Gd++ multiplicity mass function.

As a side note, let us shed light on the fact that it is possible to deduce all these efficiency corrections without relying on the simulated spectra. Indeed, by using the very same E_{th} cut on the 1n sample, i.e. the one for which only showering μ's are considered, and on the full data set, Monte-Carlo corrections are bypassed. If that provides results similar to the ones presented here, they do suffer from larger uncertainties, for indeed, the optimal E_{th} cut for the full set is never the most appropriate for the 1n sample. In fact, the latter seldom benefits from an additional purification; to make the most of this complementary method, the E_{th} cut should be chosen according to the full set exclusively.

10.3.2.3 Near detector

As mentioned earlier, only applying prompt energy cuts to the full ND samples does not yield values able to compete with the oscillation fit results. At best, the ND Gd++ set can deliver $r_{co}^{ND++} = (22 \pm 7)\,\mathrm{d}^{-1}$ while the ND Gd values can hardly be trusted for anything more precise than $r_{co}^{ND,Gd} = (10 \pm 4)\,\mathrm{d}^{-1}$. Should we contemplate voicing concerns whilst comparing these independent rate measurements with the fit results, we must needs adopt a cannier strategy.

The 1n ND rates from 10.2.2.3 are already an achievement, and we could simply correct them with the C from (10.13). However, the 14.1% relative error on the ND Gd rate from Table 10.1 craves for dwindling further. In Figure 10.9, the evolution of the relative uncertainty on the energy-corrected $r_{co}^{1n}(E_{th})$ rate, namely that obtained by demanding both $n_\mu \geqslant 1$ and $E^{vis} > E_{th}$ [7], may be studied.

Unquestionably, we had better focus on a $\sim 1\,\mathrm{MeV}$ range around 6.5 MeV to provide the most precise measurement. As we did for the far detector, we may verify our assumptions by plotting the energy-corrected rate into Figure 10.10, which simply corresponds to r_{co}^{1n}, and which is not inconsistent with the plain 1n ND Gd value from Table 10.1. The distribution of the corrected rate is convincingly flat, by a correlated conservative average over this range, we obtain $r_{co}^{1n,ND} = (5.36 \pm 0.46)\,\mathrm{d}^{-1}$, which is 40% more precise than the sole 1n analysis.

We proceed identically for the ND Gd++ data, which already benefited from a 8.6% relative uncertainty with the sole multiplicity cut. The most precise corrected rates, obtained by dint of both a prompt energy cut and the multiplicity threshold, may only bring forth a 6.9% relative uncertainty, including the Monte-Carlo systematics. This 20% improvement is modest, yet not worthless. Above all, the flatness of the distribution of the ND Gd++ 1n corrected rates – all the more so considering we are only relying on the Gd ^{9}Li simulation – is staggering (see Figure 10.11). It is worth stressing that the ^{9}Li efficiency of a 4 MeV cut is

[7] As usual, we apply $E^{vis} > E_{th}$ to produce the reduced sample defining $r_{co}^{1n}(E_{th})$, yet we present plots including the Monte-Carlo correction and its uncertainty.

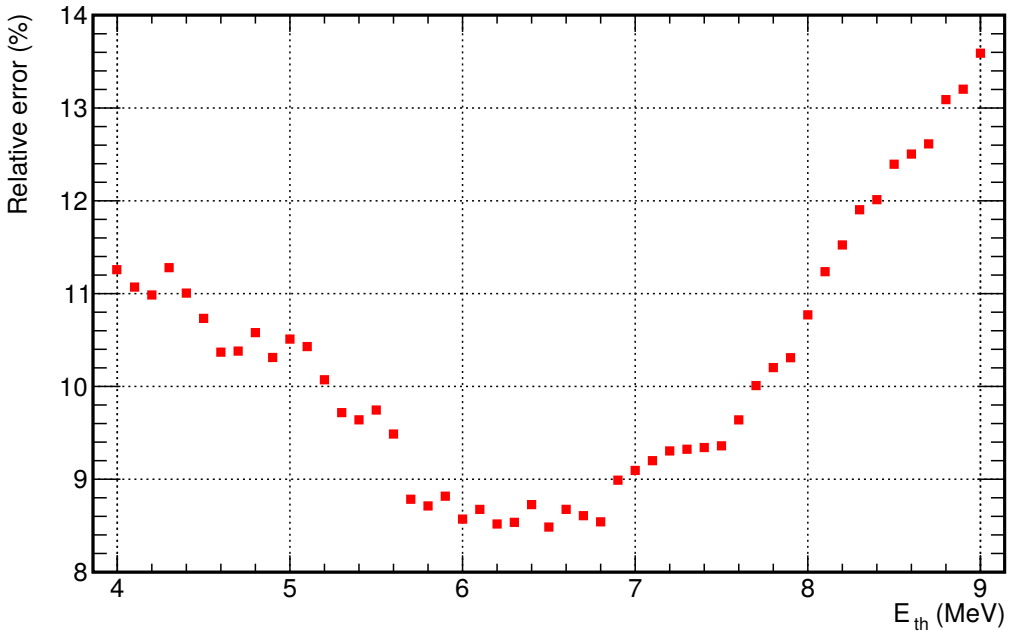

Figure 10.9 – Relative uncertainty on the Monte-Carlo corrected ND Gd rate r_{co}^{1n} as a function of the E_{th} cut applied to produce the reduced sample.

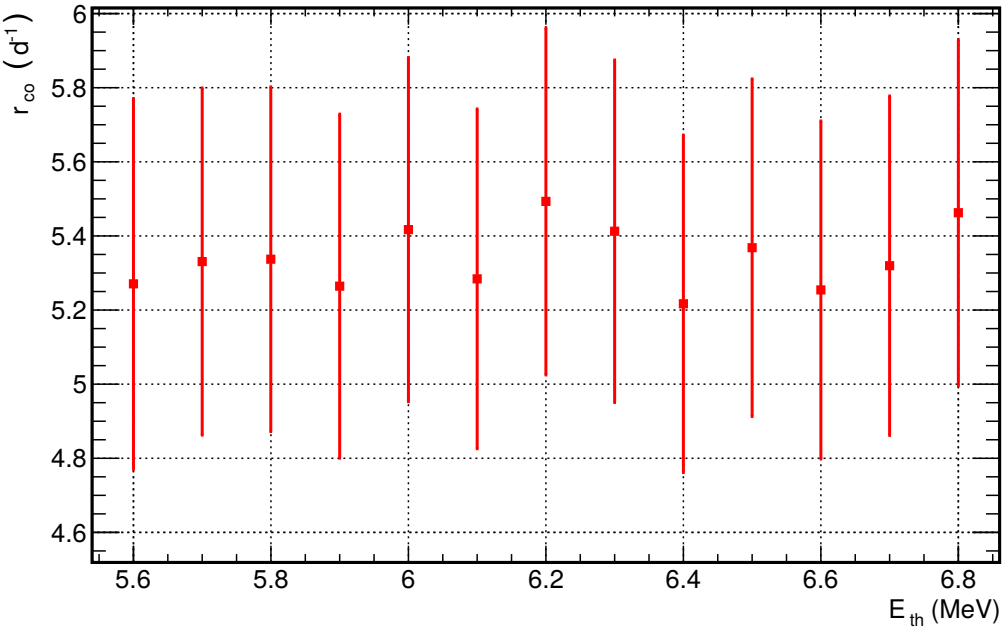

Figure 10.10 – Efficiency corrected cosmogenic rate r_{co}^{1n} as a function of the E_{th} cut applied to produce the reduced ND Gd sample.

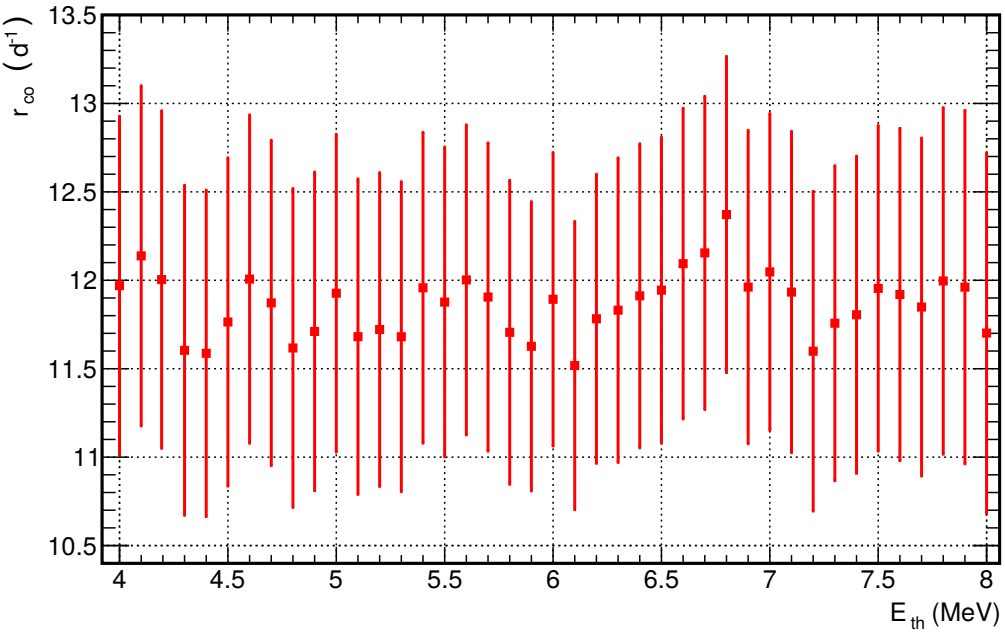

Figure 10.11 – Efficiency corrected cosmogenic rate r_{co}^{1n} as a function of the E_{th} cut applied to produce the reduced ND Gd++ sample.

around 70% whereas a 8 MeV one retains only $\sim 15\%$ of the βn-decays; and yet, for all these variations, the candidate rates for this lavish and relatively pure ($n_\mu \geqslant 1$ has been applied) sample are unfaltering. If truth be told, this astounding stability was somewhat heralded by the Monte-Carlo to vetoed events fits (cf. 9.3.2).

Making the most of the finest $5.9 - 6.4$ MeV interval from Figure 10.11, we deduce the 1n cosmogenic rate for the ND Gd++ data $r_{co}^{1n,ND++} = (11.75 \pm 0.81)\,\mathrm{d}^{-1}$. By applying the multiplicity correction (10.13) to both ND data sets, we find the total cosmogenics rates for the near site, along with their associated uncertainties. In view of the data sets used to compute the C factor, namely the FD samples, C is strictly independent from the corrected 1n ND rates, thereby ensuring a straightforward error propagation. The 1n ND rates, along with the total rates, are summarised in Table 10.5.

Version	r_{co}^{1n} (d^{-1})	r_{co} (d^{-1})
ND Gd	5.36 ± 0.46	7.40 ± 0.98
ND Gd++	11.75 ± 0.81	16.22 ± 1.99

Table 10.5 – Energy corrected 1n cosmogenic rates r_{co}^{1n} and multiplicity corrected total cosmogenic rates r_{co} for the two ND samples.

As a means of supporting our results – assuming a correlation coefficient $\rho = \sqrt{f_{Gd}} = 0.6$, as we did for FD – we can compute the ratio of the Gd++ to Gd rates $r_{++}^{co,ND}$. To this

intent, taking the uncertainties on the total rates is too forgiving, because C is a correlated component of r_{co}^{ND++} and r_{co}^{ND}, which would artificially conceal disagreements in its error bar. Fortunately, C cancels in the ratio $r_{++}^{co,ND}$, so that

$$r_{++}^{co,ND} = \frac{r_{co}^{ND++}}{r_{co}^{ND}} = \frac{r_{co}^{1n,ND++}}{r_{co}^{1n,ND}} = 2.19 \pm 0.16, \tag{10.14}$$

which is the most stringent uncertainty that could dispute our results. Despite its tight constraints, this ratio is in striking agreement with the corresponding IBD and vetoed rate ratios from (9.12), and (9.23), respectively.

10.4 Final rates

10.4.1 Remaining rates

The $\overline{\nu_e}$ candidates considered in the final oscillation fit have already been stripped of most backgrounds; $\sim 25\%$ (ND) to $\sim 50\%$ (FD) of the βn-emitters were cast aside by the cosmogenic veto. Accordingly, the vetoed rates – found in Tables 9.3 and 9.4 – must be subtracted from the total cosmogenic rates computed in 10.3.2.1 and 10.3.2.3. The resulting rates r_{co}^{re} may be compared to the oscillation fit results, a task left for 10.4.2.

The total and remaining rates, i.e. the ones obtained after the veto has been applied, are summarised in Table 10.6. To compute the remaining rates, the total and vetoed rates were considered as independent variables, which is the most conservative approach when handling a difference. The final errors are largely dominated by the errors on the total rates, thus, even with a $\rho = \sqrt{0.5} \simeq 0.71$ correlation coefficient at FD – based on the observed 50% vetoing efficiency – the uncertainties would dwindle but by 10%. The error reduction at ND would be less than 5% with the calculated 24% vetoing efficiency. In light of 9.2.2, it is worth noting that these efficiency values were not unannounced.

Version	r_{co} (d^{-1})	r_{co}^{re} (d^{-1})
FDI Gd	2.32 ± 0.50	1.20 ± 0.50
FDII Gd	1.46 ± 0.52	0.34 ± 0.52
ND Gd	7.40 ± 0.98	5.63 ± 0.98
FDI Gd++	6.19 ± 0.80	3.39 ± 0.81
FDII Gd++	5.40 ± 0.87	2.60 ± 0.88
ND Gd++	16.22 ± 1.99	12.25 ± 1.99

Table 10.6 – Total r_{co} and remaining cosmogenic r_{co}^{re} rates for all six Neutrino configurations.

For all our efforts to obtain accurate total rates, the simple subtraction of the vetoed rates, which are thankfully known with a great accuracy, substantially increases the relative

uncertainties on the remaining rates. The relative errors are all the greater that the veto is efficient. For this reason, the FDII Gd remaining rate looks particularly queer. The latter would certainly benefit from asymmetric uncertainties, with a lower bound driven by the corresponding 1n rate; it does convolute the treatment. If truth be told, the cosmogenic veto somewhat reduces the significance of the oscillation fit results as well; an alternative approach would leave all the cosmogenic background within the $\overline{\nu_e}$ candidates, to increase the statistical power of all analyses, but that would entail relying on the predicted spectra. In other words, weaker constraints on the remaining rates are the price to pay to use data – as opposed to Monte-Carlo – spectra in the oscillation fit.

10.4.2 Comparison

As could be noticed on the prior ratios for the cosmogenic veto in 9.1.3.2, the total number of cosmogenic events does not directly scale with the μ-rate. If the near μ-rate is about $\times 5.2$ that of the far site, the Gd++ data claim a cosmogenic rate ratio of 2.82 ± 0.44, averaging both FD measurements, while the ND-FDI Gd pair favours 3.34 ± 0.82. This discrepancy between the μ-rate and the cosmogenic one backs the toughness of the ND analyses. Averaging both FD Gd values, we do get a cosmogenic rate ratio of 4.06 ± 0.91, closer to the μ-rate ratio. Nevertheless, there is some evidence (cf. 10.3.2.1) that this value is an upwards fluctuation, on account of the smallness of the FDII Gd rate.

With regard to the Gd rates, there has been no Gd oscillation fit since the Moriond conference. Nonetheless, its results ought to be comparable to our independent analyses of the Neutrino Gd data. These fit results are recalled in Table 10.7. Additionally, extremely up-to-date results have been presented at CERN on 20th September with the Gd++ captures; these oscillation fit values have been included in the table as well.

Version	r_{co}^{fit} (d^{-1})
FD Gd	0.75 ± 0.14
ND Gd	4.89 ± 0.78
FD Gd++	2.55 ± 0.23
ND Gd++	14.4 ± 1.2

Table 10.7 – Cosmogenic background rates r_{co}^{fit} obtained from the simultaneous fit all the Gd (Moriond) or Gd++ (Neutrino sample eventually presented at CERN) data sets.

Let us recall that the FDI and FDII rates are considered fully correlated in the oscillation fits, hence the presence of a sole far detector number per analysis in Table 10.7. If we combine our far Gd results into a weighted average, we find $r_{co}^{FD,Gd} = (0.79 \pm 0.36)\,\mathrm{d}^{-1}$, which is in stunning agreement with the Moriond fit results, as are the individual FDI and FDII rates themselves. In a like manner, we obtain $r_{co}^{FD++} = (3.03 \pm 0.60)\,\mathrm{d}^{-1}$, i.e. a value less than

1σ larger than the corresponding fit output. When it comes to the our somewhat bold ND measurements, they are also perfectly compatible with the final fit estimations.

The compatibility of the fit and independent rates not only supports the accuracy of the Monte-Carlo corrections for all the decays, but also the validity of the vetoed spectrum, used in the fit. Indeed, were the vetoed shape biased, the fit might return over or under estimates of the remaining rates. Indirectly, this harmony confirms that the vetoed shape, the untagged one, and the simulations are all in the same ballpark. Furthermore, entirely overlooking the potential ^{8}He component in the $\Delta t_{\mu-p}$ analysis seems not hamper the reproducibility of the oscillation fit results, which genuinely include the ^{8}He spectrum by virtue of the veto.

By all manner of means, the rightfulness of our assumptions is rooted in the limited statistics; unquestionably, in an infinitely remote future blessed with unlimited statistics, we would reject the null hypotheses. This final comparison with the fit results marks the end of our cosmogenic journey.

Part IV

Relative normalisation of the $\overline{\nu_e}$ rates

Chapter 11

Weight measurements

If the production of cosmogenic isotopes in a given detector is related to the number of carbon nuclei inside this detector, we ought not to forget that the main signal of the Double Chooz experiment – consisting of $\overline{\nu_e}$ interactions – has a rate directly proportional to the number of free protons in the considered target.

Computing the mass of liquid scintillator poured into a target gives access to the number of protons within it, via the chemical composition of the liquid. Precisely evaluating the mass of scintillator within each detector is all the more paramount that the two targets have been filled with liquid coming from the same production batch, whose associated composition uncertainty cancels out in a relative measurement. In order to accurately quantify these masses, plain geometrical calculations will not meet the high standards of the Double Chooz analyses. On the other hand, extreme diligence when analysing the data of the so-called weight measurements can provide results of astonishing accuracy.

11.1 Performing a weight measurement

11.1.1 Principle

The principle of the weight measurement itself – which consists in monitoring in real-time the mass of liquid that remains to be poured into the target – is fairly uncomplicated, at first sight. Nonetheless, when it comes to setting errors on the difference between the mass of liquid originally meant for the target and that which remained at the end of the filling, great care must be taken.

In truth, by setting up mass sensors beneath a weighing tank from which liquid is removed, we effectively measure its mass, augmented by that of the liquid. Obviously, the former cancels out in the difference between the starting and remaining masses, which provides us with a first estimate of the mass actually poured into the target. This value must indeed be qualified as a "first estimate", for as we shall see, other non-negligible corrections do not cancel as the mass of the weighing tank does.

11.1.2 Standards

For the reader to grasp why every single detail will be considered throughout this chapter, we must emphasise how small of a margin was left for carelessness. Indeed, the Double Chooz group was bent on obtaining a 0.2% relative uncertainty on the mass of liquid contained within the volume of each target ; this mass ought to stand as a quantity independent from one detector to the other. Considering an approximative mass $M_{NT}^{liquid} \simeq 8\,\mathrm{t}$, such a demanding precision implies a $\sigma_{M_{NT}^{liquid}} = 16\,\mathrm{kg}$ error, at worst, on M_{NT}^{liquid}, that is to say a trifle.

Incidentally, geometrical calculations can hardly claim a relative uncertainty below 0.6% [129]. Originally, the Double Chooz experiment was not designed to use its gamma catcher as a detection vessel, and this additional volume comes with a comparably high uncertainty. Although the Neutrino Gd++ analysis has divided all the statistical error bars by a factor of two, with respect to the Moriond Gd endeavour, this improvement came with strings attached: the control of the systematics on the normalisation of the $\overline{\nu_e}$ rates has dramatically suffered. If truth be told, it worsened so staggeringly that it trumped, albeit by a small margin, backgrounds in the race for the largest uncertainty share.

Within this chapter, we focus on the sole Gd analysis. In particular, we will be interested in the data relevant for the Moriond oscillation fit (cf. 6.3), which the numbers of protons obtained by means of target weight measurements embody. We shall start with the near detector, as the latter is a perfect textbook case, which took advantage of the mistakes made during the filling of the older detector.

11.2 Near detector data analysis

11.2.1 Weight measurement

The goal of this section is not to comment on the well-designed weight measurement system [143], but rather to briefly review the characteristics of the stages at which the mass of the weighing tank is recorded. Understanding the key features of these stages will help us to define the corrections that must be made to the bare mass difference.

11.2.1.1 Full weighing tank

For the first weight measurement, the weighing tank is filled with liquid scintillator and flushed with nitrogen. All the tubes and other tanks are empty so that all the liquid scintillator that may end up in the target of the near detector itself can be weighed by the first measurement. The reader is invited to take a look at Figure 11.1 to picture the configuration.

A first measurement yields $M_0 = 10\,638\,\mathrm{kg}$ for the mass of the full weighing tank. This value is obtained just after the weighing has been heaved and before the drift of the mass sensors, which go as low as to display a 10 634 kg value, as Figure 11.2 hints.

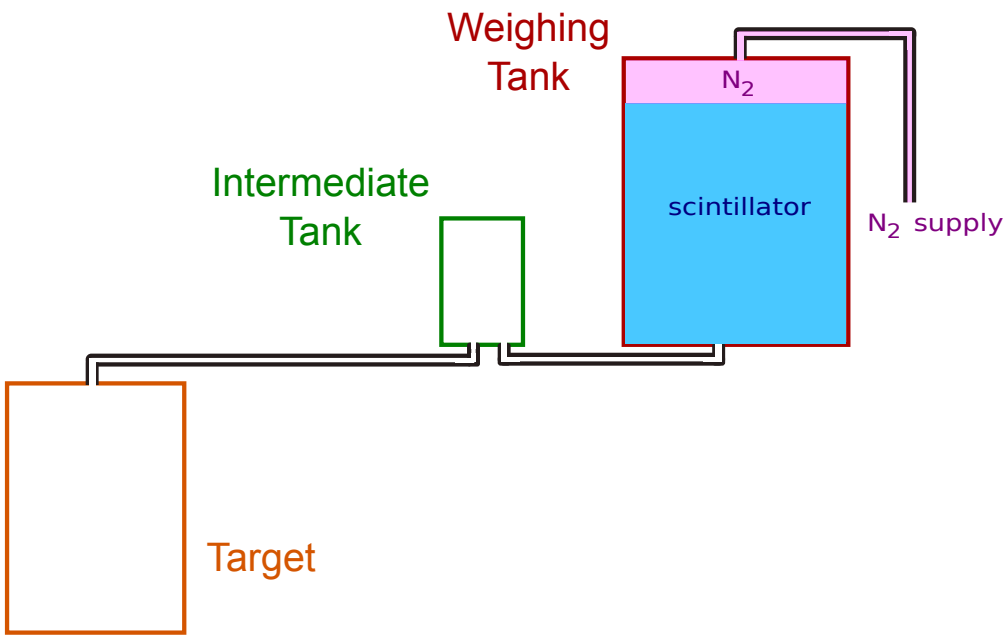

Figure 11.1 – Simplified diagram of the filling system for the first weight measurement.

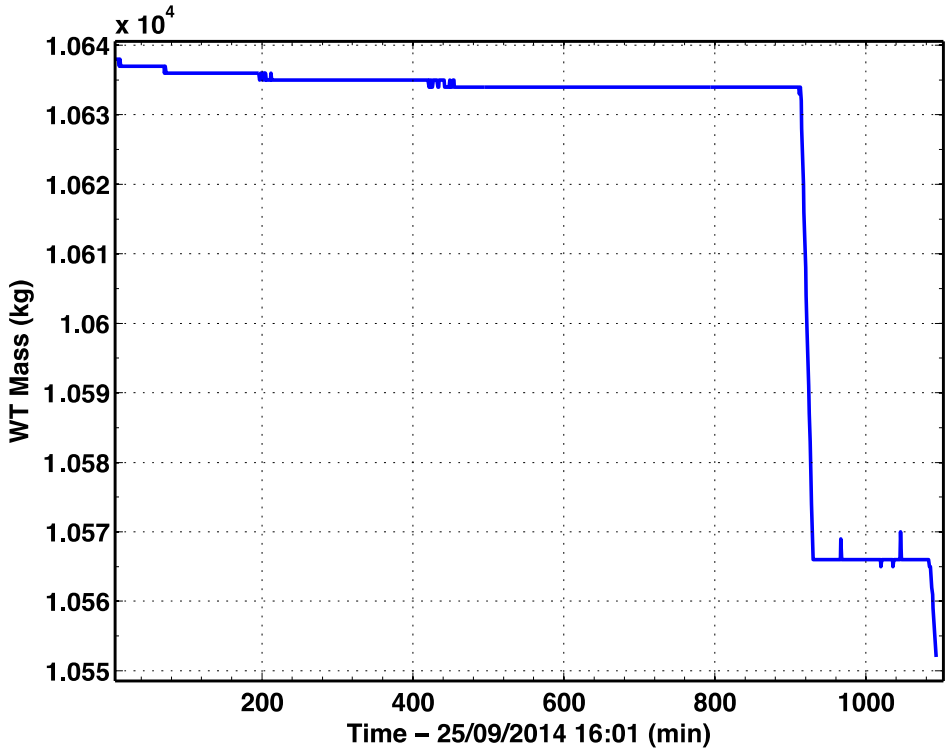

Figure 11.2 – Measured mass of the weighing tank at the filling start. A drift of the mass of the weighing tank – from $M_0 = 10\,638\,\text{kg}$ to $10\,634\,\text{kg}$ – can be observed over the first 15 h of acquisition, and before the actual filling starts.

11.2.1.2 Target filled

During the filling, the intermediate tank is being filled with scintillator and the target in turn. It takes about 100 h of filling to reach the first position sensor (CPS) 3 cm below the start of the chimney. An overview of the mass of the weighing tank during the filling and subsequent thermalisation stage can be found in Figure 11.3.

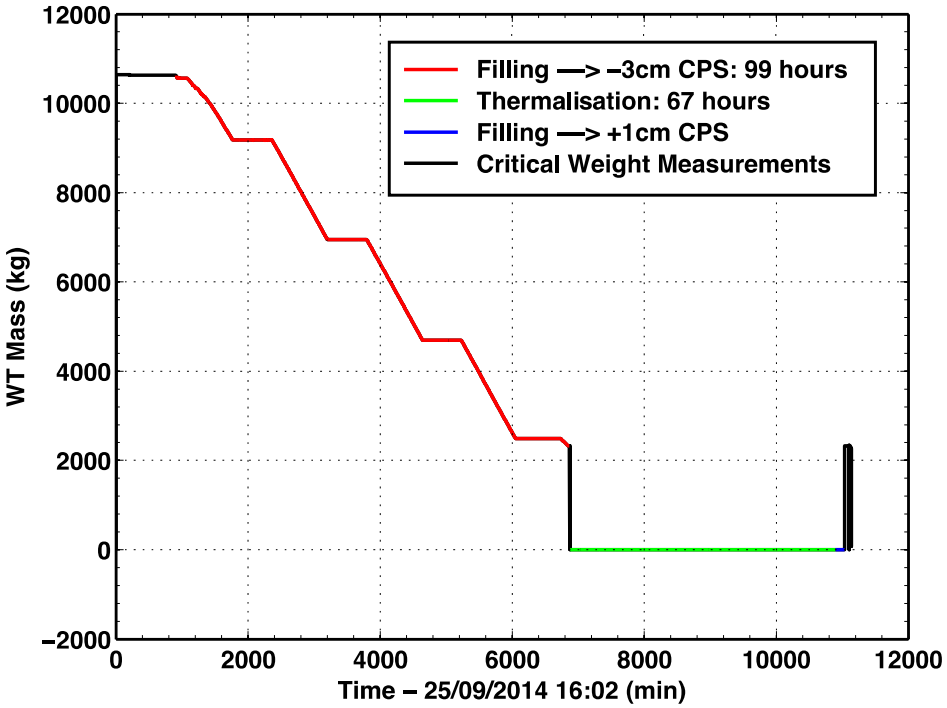

Figure 11.3 – Overview of the measured mass of the weighing tank over the whole filling period.

To ensure that the temperature of the liquid is homogeneous in the target, and that it does not expand into the chimney, the filling stops for about three days, and the temperature evolution is recorded by five sensors distributed along a rod. The temperatures of all the sensors during the thermalisation stage are gathered in Figure 11.4.

Once the level in the target reaches the start of the chimney, namely when the scintillator is in contact with the second CPS about 1 cm up in the chimney (see [144] for more details), the second weight measurement is performed. At first, the intermediate tank is not empty and the sensors indicate 2327 kg for the mass of the weighing tank. A close-up look of Figure 11.3 over this second critical weight measurement is given in Figure 11.5. After the liquids in the intermediate tank and in the tube connecting the weighing tank to the intermediate tank have been drained back into the weighing tank, the sensors display $M_1 = 2334\,\text{kg}$ as Figure 11.5 highlights. For this measurement, the weighing tank is lifted to cancel a possible sensor drift. Studies carried out by J. - C. Barrière in Saclay with known masses

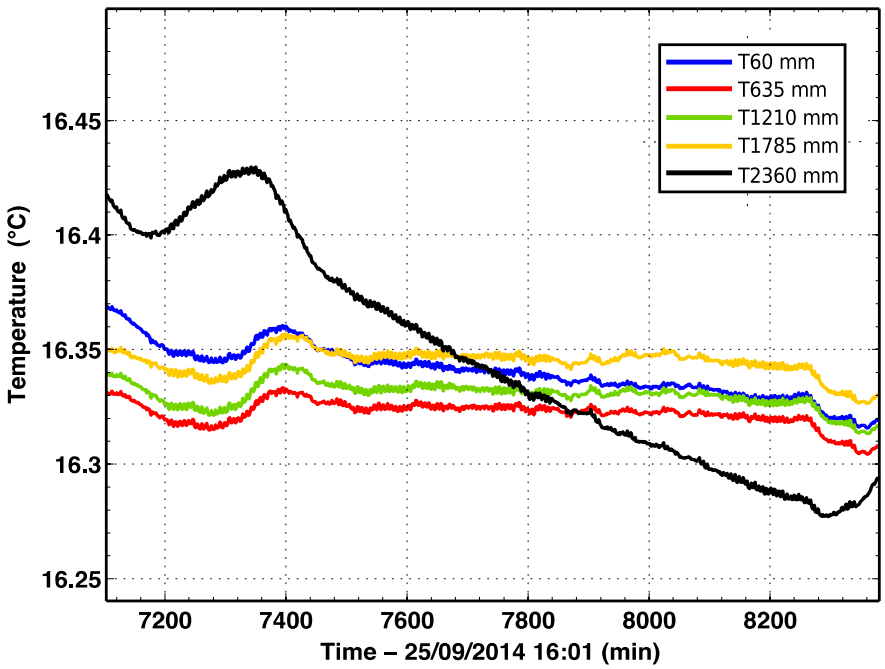

Figure 11.4 – Temperature evolution of the target sensors during the thermalisation stage. The labels of the sensors carry their position in the target along the upwards z-axis, whose origin is located at the bottom of the target. The sensor at $z = 2360\,\text{mm}$ behaves as if a little outside the liquid.

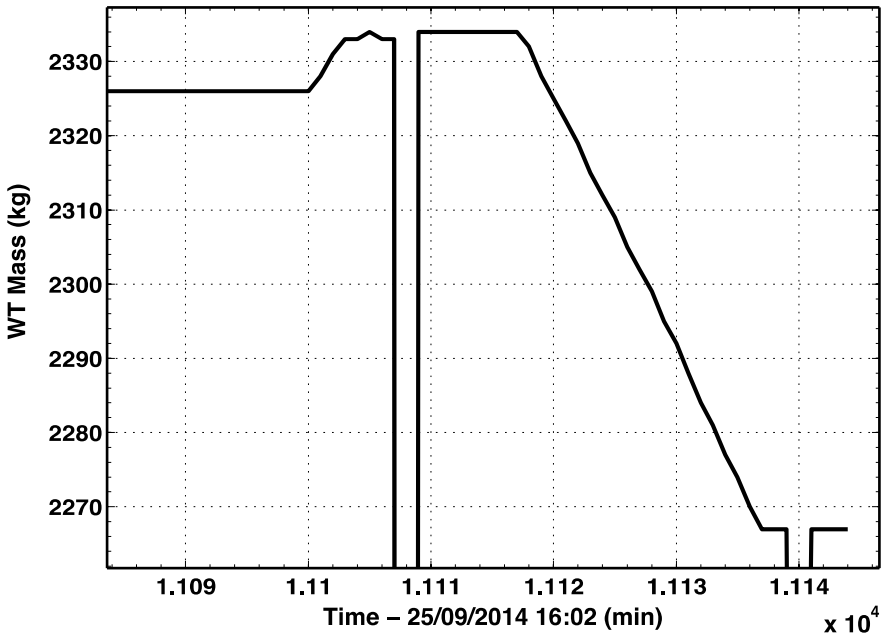

Figure 11.5 – Measured mass of the weighing tank during the final weight measurement. The contents of the intermediate tank are drained back into the weighing tank, yielding a final mass $M_1 = 2334\,\text{kg}$.

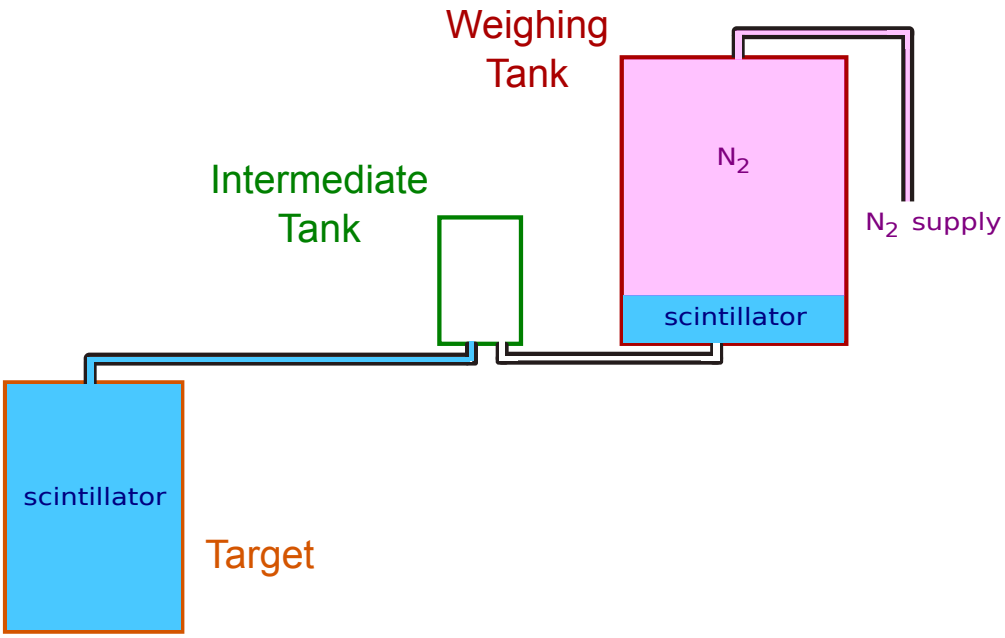

Figure 11.6 – Simplified diagram of the filling system for the second weight measurement.

demonstrate that the sensors do not drift when the measurements are recorded up to 10 min after the mass has been heaved (see [145] for more details). A 500 g tare has also been added on the weighing tank to make sure that the last figure on the reading system did not change. The tube connecting the intermediate tank to the target, however, had not been drained into the weighing tank and must therefore be taken into account in the target mass estimation. A diagram summarising the state of each component of the filling chain at this step can be found in Figure 11.6. The temperature in the target was measured to be $T = (16.3 \pm 0.2)\,°\mathrm{C}$ with great confidence, as vouched for by Figure 11.4. The pressure was not recorded for any of the weight measurements but can be assigned the atmospheric value $P = (1013 \pm 20)\,\mathrm{mbar}$.

11.2.2 Target mass estimation

The mass of liquid scintillator in the target M_{NT}^{liquid} is the difference between the mass that was in the weighing tank before the filling started, i.e. M_0 and the one that was not poured into the target, i.e. M_1. Nevertheless, because of the N_2 supply, the sensors for the second measurement have weighed something they should not, the N_2 mass increase $M_{WT}^{N_2}$. Similarly, as the tube between the intermediate tank and the target was not empty, the corresponding mass M_{tube}^{liquid} should be gauged and added to the value of M_1. To conclude, the mass of liquid scintillator in the target of the near detector reads

$$M_{NT}^{liquid} = M_0 - M_1 + M_{WT}^{N_2} - M_{tube}^{liquid}. \tag{11.1}$$

11.2.2.1 Sensors

The configuration of the three sensors under the weighing tank is such that each of them supports about a third of the total weight. The measuring system displays the sum of the three weights measured by the three sensors. Each sensor has a $\sigma_{sensor} = 1\,\text{kg}$ uncertainty in our range. The measurements made by the three sensors are independent from one another, consequently, the uncertainty of a weight measurement is $\sigma_{sensors} = \sqrt{3\,\sigma_{sensor}} \simeq 1.7\,\text{kg}$. To this uncertainty must be added the unrelated "reading" uncertainty, originally of 1 kg but which was cut down in half by the use of the 0.5 kg tare for the measurement of M_1 (see 11.2.1.2). As a result

$$M_0 = (10\,638.0 \pm 2.0)\,\text{kg} \tag{11.2}$$

$$M_1 = \quad (2334.0 \pm 1.8)\,\text{kg}. \tag{11.3}$$

Please note that unlike for the weight measurement at the far detector, the mass sensors were calibrated in the near detector itself. Thus, no gravity correction need be made.

11.2.2.2 Nitrogen

How can we estimate the N_2 mass increase in the weighing tank ? Using the dimensions of the weighing tank to assess the value of $M_{WT}^{N_2}$ would lead to larger uncertainties than when using the great accuracy of the weight measurement. Thus, let us express $M_{WT}^{N_2}$ in terms of the volumetric mass densities of both N_2 and the scintillator. One can straightforwardly derive that

$$M_{WT}^{N_2}(T,P) = \frac{\rho_{N_2}(T,P)\,(M_0 - M_1)}{\rho_{liquid}(T,P)}. \tag{11.4}$$

Considering the central values of (T,P) given in 11.2.1.2 and using for instance [146], we find $\rho_{N_2}(T,P) = (1.18 \pm 0.03)\,\text{kg}\,\text{m}^{-3}$. It must be underlined that although the temperature in the liquid may not quite be that of N_2 in the weighing tank, the uncertainty evaluated for $\rho_{N_2}(T,P)$ easily covers the $10-20°\text{C}$ range and includes the uncertainty on the atmospheric pressure P. As for the volumetric mass density of the scintillator $\rho_{liquid}(T,P) = \rho_{liquid}$, the reference value at $T_{ref} = 14\,°\text{C}$ is $\rho_{liquid}(T_{ref},P) = \rho^0_{liquid} = (804 \pm 1)\,\text{kg}\,\text{m}^{-3}$ and the coefficient of thermal expansion for the scintillator is $\kappa = 7 \times 10^{-4}\,°\text{C}^{-1}$. In accordance, we may write

$$\rho_{liquid} = \rho^0_{liquid} + \rho^0_{liquid}\,\kappa\,(T_{ref} - T)\,. \tag{11.5}$$

Bearing in mind that the error on ρ^0_{liquid} and that on the temperature in the near detector are uncorrelated, we find that

$$\sigma_{\rho_{liquid}} = \sigma_{\rho^0_{liquid}} + \kappa\,\sqrt{\left(\sigma_{\rho^0_{liquid}}(T_{ref} - T)\right)^2 + \left(\rho^0_{liquid}\sigma_T\right)^2}, \tag{11.6}$$

which in turn yields[1]

$$\rho_{liquid} = (802.7 \pm 1.1)\,\mathrm{kg\,m^{-3}}. \tag{11.7}$$

We can take the differentiate form of (11.4) to estimate the uncertainty $\sigma_{M_{WT}^{\mathrm{N_2}}}$ on the computed value of $M_{WT}^{\mathrm{N_2}}$

$$\sigma_{M_{WT}^{\mathrm{N_2}}} \simeq \frac{\rho_{\mathrm{N_2}}\sqrt{\sigma_{M_0}^2 + \sigma_{M_1}^2}}{\rho_{liquid}} + \frac{(M_0 - M_1)\,\sigma_{\rho_{\mathrm{N_2}}}}{\rho_{liquid}} + \frac{\rho_{\mathrm{N_2}}\,(M_0 - M_1)\,\sigma_{\rho_{liquid}}}{{\rho_{liquid}}^2}, \tag{11.8}$$

where we have dropped some dependencies for brevity and used the uncertainty on the mass difference $\sqrt{\sigma_{M_0}^2 + \sigma_{M_1}^2}$. Such an estimation is unquestionably conservative since $(M_0 - M_1)$ and $\rho_{\mathrm{N_2}}/\rho_{liquid}\,(T, P)$ are actually independent variables and could otherwise benefit from an even finer treatment. This eventually leads to

$$M_{WT}^{\mathrm{N_2}}\,(T, P) = (12.21 \pm 0.33)\,\mathrm{kg}. \tag{11.9}$$

11.2.2.3 Filling tube

The mass M_{tube}^{liquid} is assessed using the geometrical dimensions of the tube. From the manufacturer, the diameter of the tube is $d = (9.55 \pm 0.13)\,\mathrm{mm}$. The length of the pipe to the top edge of the chimney is $l_{edge} = (23.68 \pm 0.05)\,\mathrm{m}$ but there remains to add to this value the length of the part that goes down from the chimney flange to the target. Utilising the plan of the detector, this additional part is conservatively estimated to $l_{chim} = (2.22 \pm 0.25)\,\mathrm{m}$. Therefore, the total length of the tube is $l = l_{edge} + l_{chim} = (25.90 \pm 0.25)\,\mathrm{m}$ for the lengths are uncorrelated. The uncertainty on ρ_{liquid} is clearly negligible in this case but it has been included for completeness without much effort. The mass reads

$$M_{tube}^{liquid} = \rho_{liquid}\,\pi\,\frac{d^2}{4}\,l, \tag{11.10}$$

and

$$\sigma_{M_{tube}^{liquid}} \simeq \pi\,\frac{d}{4}\,\sqrt{{\rho_{liquid}}^2\left[(2l\sigma_d)^2 + (d\sigma_l)^2\right] + \left(d\,l\,\sigma_{\rho_{liquid}}\right)^2}, \tag{11.11}$$

insofar as the errors on d^2, l and ρ_{liquid} are uncorrelated. Please note that the error on d^2 has been derived using a simple differentiate form[2]. Plugging in the numbers yields

$$M_{tube}^{liquid} = (1.49 \pm 0.05)\,\mathrm{kg}. \tag{11.12}$$

[1] This intermediate numerical value is written here only to allow the reader to gauge the different uncertainties.

[2] The errors are small enough anyway.

11.2.2.4 Target mass value after filling

Using (11.1), and considering that the errors from sections 11.2.2.1, 11.2.2.2, 11.2.2.3 are fully uncorrelated, one can promptly exhibit a M_{NT}^{liquid} value with

$$\sigma_{M_{NT}^{liquid}} = \sqrt{{\sigma_{M_0}}^2 + {\sigma_{M_1}}^2 + {\sigma_{M_{WT}^{N_2}}}^2 + {\sigma_{M_{tube}^{liquid}}}^2} \,, \tag{11.13}$$

namely

$$M_{NT}^{liquid} = (8314.7 \pm 2.7)\,\text{kg}. \tag{11.14}$$

Please keep in mind that this result has been obtained for $T = 16.3\,°\text{C}$ and $P = 1013\,\text{mbar}$. The mass property of an object in Physics neither depends on the temperature nor the pressure, but the mass of liquid in a given volume does depend on these variables for the liquid may expand with them.

11.2.3 Target mass evolution and number of protons

After having determined the mass of the target $M_{NT}^{liquid}(T, P)$ at an homogeneous temperature $T = 16.3\,°\text{C}$ and at the atmospheric pressure $P = 1013\,\text{mbar}$, we can scale that value as the temperature changes, as long as the expansion of the target acrylic vessel and that of the liquid are understood.

11.2.3.1 Vessel expansion

The Polymethyl methacrylate (PMMA) composing the target vessel and the glue holding it together do suffer from thermal expansion. Let $\alpha = 7 \times 10^{-5}\,°\text{C}^{-1}$ be the linear thermal expansion coefficient of PMMA, as provided by the manufacturer (Neotec). The uncertainty σ_α did not come along but an overview of the literature [147] and the global polymer market [148] allows $\sigma_\alpha = 2 \times 10^{-5}\,°\text{C}^{-1}$ to span the range of PMMA available in the world.

The volume of the target at T_{use} evolves according to

$$V_{NT}^{liquid}(T_{use}, P) = (1 + \alpha\,(T_{use} - T))^3\, V_{NT}^{liquid}(T, P)\,, \tag{11.15}$$

which is to say that the target shrinks if the operating temperature is lower than that of the weight measurement.

11.2.3.2 Liquid expansion

Should the target shrink or expand, the liquid would concurrently see its density increase or decrease. Although anti-correlated, both effects do not completely cancel each other. Expressing $\rho_{liquid}(T_{use}, P)$ in terms of κ as in (11.5), we find

$$\rho_{liquid}(T_{use}, P) = (1 + \kappa\,(T - T_{use}))\,\rho_{liquid}(T, P)\,. \tag{11.16}$$

11.2.3.3 Overall evolution

Evidently, we have

$$M_{NT}^{liquid}\left(T_{use},P\right)=\rho_{liquid}\left(T_{use},P\right)V_{NT}^{liquid}\left(T_{use},P\right). \tag{11.17}$$

Substituting (11.15) and (11.16) into (11.17), we obtain

$$M_{NT}^{liquid}\left(T_{use},P\right)=\left(1+\kappa\left(T-T_{use}\right)\right)\left(1+\alpha\left(T_{use}-T\right)\right)^{3}M_{NT}^{liquid}\left(T,P\right). \tag{11.18}$$

With regard to the error treatment, the diligence of (11.18) is excessive and a first order expansion will suffice[3]

$$M_{NT}^{liquid}\left(T_{use},P\right)\simeq\left(1+\left(\kappa-3\alpha\right)\left(T-T_{use}\right)\right)M_{NT}^{liquid}\left(T,P\right). \tag{11.19}$$

Insofar as $M_{NT}^{liquid}\left(T,P\right)$ is mostly correlated to the weighing tank mass difference M_0-M_1, and in accordance largely independent of $\Delta T=T-T_{use}$, we have

$$\sigma_{M_{NT}^{liquid}\left(T_{use}\right)}\simeq\sqrt{\left(\sigma_{exp}M_{NT}^{liquid}\left(T\right)\right)^{2}+\left(\left(1+\left(\kappa-3\alpha\right)\Delta T\right)\sigma_{M_{NT}^{liquid}\left(T\right)}\right)^{2}}, \tag{11.20}$$

where $\sigma_{exp}=\sqrt{\left(3\sigma_{\alpha}\Delta T\right)^{2}+\left(\left(\kappa-3\alpha\right)\sigma_{\Delta T}\right)^{2}}$ and $\sigma_{\Delta T}=\sqrt{\sigma_{T}^{2}+\sigma_{T_{use}}^{2}}$.

The average temperature of the inner near detector over the so-called nine-month Moriond data run list can be confidently summarised[4] by $T_{use}=\left(13.4\pm0.5\right)\,^{\circ}\mathrm{C}$, which translates into a scaled mass

$$M_{NT}^{liquid}\left(T_{use},P\right)=\left(8326.5\pm3.8\right)\,\mathrm{kg}. \tag{11.21}$$

The reader may find of interest to know that such temperature variations embody an increase of the density of the liquid by 0.203% and a volume shrinkage of about 0.061%.

11.2.3.4 Number of target protons

One can go a little further for the experiment and compute the number of protons $n_H\left(T_{use},P\right)$ corresponding to the target mass from (11.21). This final step requires the knowledge of the mass fraction of hydrogen in the liquid scintillator f_H to be substituted into the following equation

$$n_H\left(T_{use},P\right)=\frac{f_H M_{NT}^{liquid}\left(T_{use},P\right)}{m_H}. \tag{11.22}$$

The value of the hydrogen fraction was provided by C. Buck and published in [73]

$$f_H=0.1360\pm0.0004. \tag{11.23}$$

[3] The central value changes but for 10 g, which we surely can afford.

[4] This also accounts for potential calibration biases of the sensors.

This mass fraction is a chemical property of the liquid and the atomic mass $m_H = 1.007\,825\,\mathrm{u}$ [78] must be used[5]; we re-write it as $m_H = 1.673\,533 \times 10^{-27}\,\mathrm{kg}$ to obtain the number of protons in the target. The uncertainty that prevails in (11.22) comes from the product $f_H M_{NT}^{liquid}(T_{use}, P)$ and we have

$$\begin{aligned}\sigma_{n_H} &\simeq \frac{\sqrt{\left(f_H \sigma_{M_{NT}^{liquid}}\right)^2 + \left(M_{NT}^{liquid}\sigma_{f_H}\right)^2 + \left(\sigma_{M_{NT}^{liquid}}\sigma_{f_H}\right)^2}}{m_H} \\ &\simeq \frac{\sqrt{\left(f_H \sigma_{M_{NT}^{liquid}}\right)^2 + \left(M_{NT}^{liquid}\sigma_{f_H}\right)^2}}{m_H},\end{aligned} \tag{11.24}$$

with which we get

$$n_H^{ND}\left(T_{use}^{ND}, P^{ND}\right) = (6.767 \pm 0.020) \times 10^{29}, \tag{11.25}$$

where we have explicitly written the ND upper-script for the near detector. It is worth making note of the fact that even with a relative uncertainty on M_{NT}^{liquid} lower than 0.05%, the relative uncertainty on n_H in one detector cannot be lower than 0.3%, on account of the poorer accuracy on the chemical composition of the liquid.

11.3 Far detector data re-analysis

This section is set on re-analysing the far detector data with an unprecedented error treatment. It assumes a good knowledge and understanding of the weight measurement stages and corrections to the mass of the target as detailed in 11.2.

The former far detector mass $M_{NT}^{FD,old} = (8284.5 \pm 2.0)\,\mathrm{kg}$ [145] is manifestly unreliable. Not only is its central value off by more than 1σ, as we shall see, but the error on it is largely underestimated, as agreed with the authors of [145]. Far be it from us to blame the latter, for we gained much from their knowledge. In fact, their approach made all the sense it could a few years ago: without the near detector and relative measurements of the $\overline{\nu_e}$ signal, the uncertainty on the $\overline{\nu_e}$ rate was overwhelmed by that on the hydrogen mass fraction from (11.23), which naturally discouraged efforts at painstakingly computing the mass.

11.3.1 Weight measurement

The overall weight measurement at the far detector did not go as smoothly as it did at the near detector, and much was learned to perform a weight measurement with the best accuracy in the later case.

[5] Though more significant digits and the uncertainty on this most accurate quantity are used in the code at no cost, their presence is irrelevant.

11.3.1.1 Weighing tank loading

Whereas the weighing tank was thoughtfully lifted before measuring M_0 in the recent weight measurement, the first iteration was not blessed with such care and we must review the filling of the weighing tank itself first, if we are to understand the essential sensor drift.

The loading of the weighing tank lasted for about seven days, and was split into two steps, with a five-day break in between. During that break, the first evidence of a drift of the mass sensors was observed. Indeed, over five days, the value read for the mass of the weighing tank had decreased by about $M_{break}^{drift} = 3\,\text{kg}$, as Figure 11.7 emphasises. Please note that in Figure 11.7, as in all the figures of the same type throughout these studies, random spikes are caused by careless activity in the laboratory; they ought not to draw the attention of the reader.

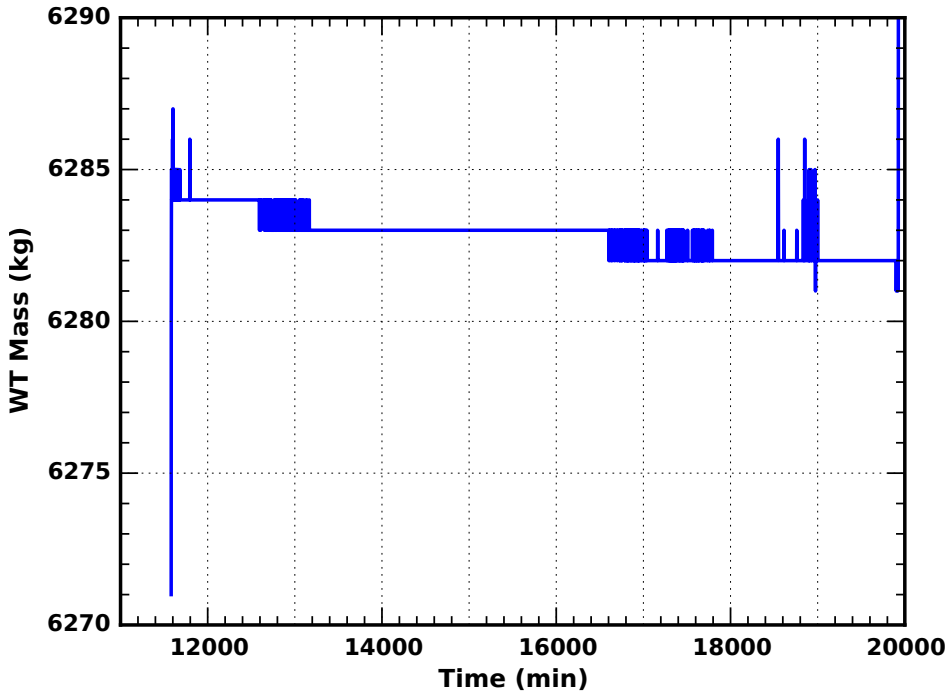

Figure 11.7 – Measured mass of the weighing tank during the loading break. The mass sensors exhibit a drift of about $M_{break}^{drift} = 3\,\text{kg}$ before the loading resumes.

11.3.1.2 Full weighing tank

Expectedly, the mass sensors do not only drift during the loading of the weighing tank itself. In fact, the drift of the sensors during the loading of the weighing tank must be added to the drift happening between the end of the loading and the actual start of the target filling. At the end of the loading, the mass of the weighing tank read $M_0 = 10\,634\,\text{kg}$, and after a thirty-five-day gap, the reading had lost $M_{start}^{drift} = 9\,\text{kg}$, as Figure 11.8 demonstrates.

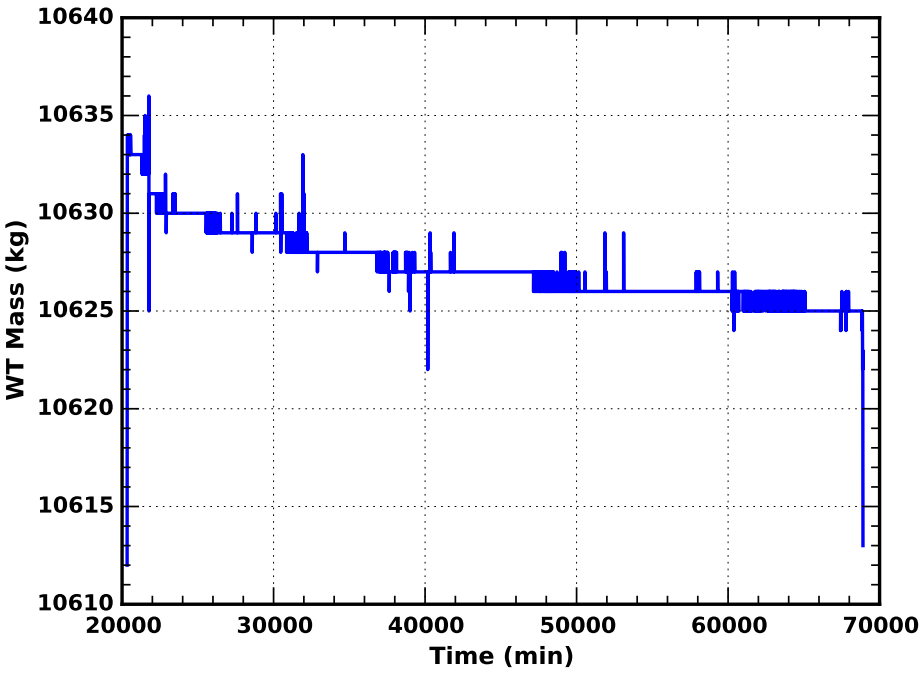

Figure 11.8 – Measured mass of the weighing tank before the filling start. A drift of the mass of the weighing tank from $M_0 = 10\,634\,\text{kg}$ to $10\,625\,\text{kg}$ can be observed over the first 35 d of acquisition, and before the actual filling starts.

11.3.1.3 Target filled

The filling lasted considerably longer at the far detector than it did at the near site; it took about ten days to reach the first CPS 3 cm below the start of the chimney. An overview of the mass of the weighing tank during the filling and subsequent thermalisation stage can be found in Figure 11.9.

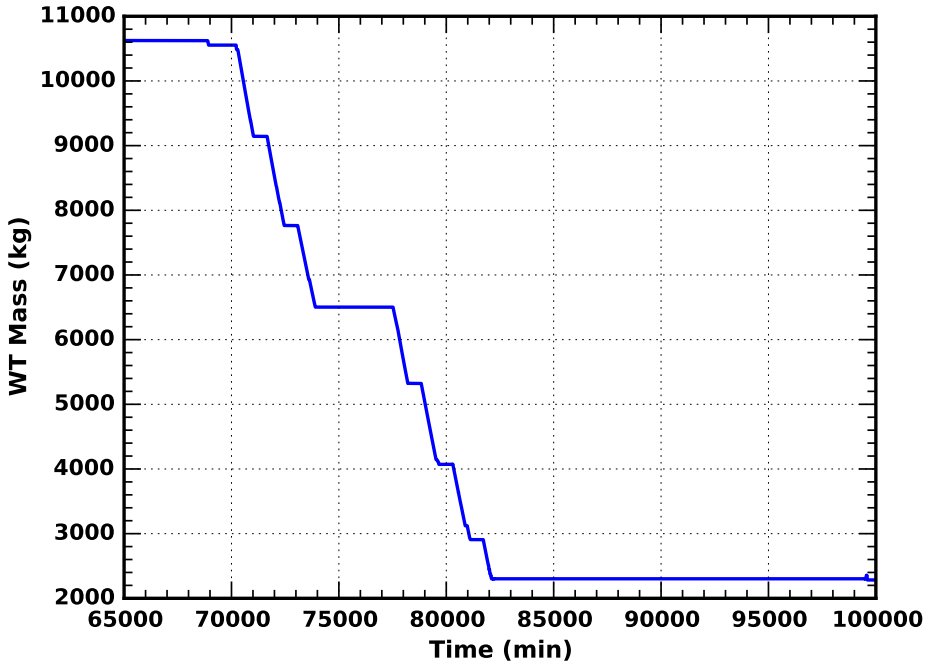

Figure 11.9 – Overview of the measured mass of the weighing tank over the whole filling period.

The thermalisation phase approximately started at $t_1 = 82 \times 10^3$ min and ended at $t_2 = 100 \times 10^3$ min, spanning around 12 d. The temperature sensors suffered a wider spread than during the filling of the near detector, which can be attributed to miscalibrations and a considerably more intricate acquisition chain, with several electric converters and tremendous cable lengths, bypassed altogether four years later.

The temperatures of the working sensors[6] along the rod, during the thermalisation stage, are plotted in Figure 11.10. In view of the non-negligible spread of the sensors, an average temperature $T = (14.0 \pm 0.5)\,°\mathrm{C}$ is inferred as the temperature of the target. The evolution of this average can be monitored in Figure 11.11. The pressure was recorded as well, and it is here condensed into $P = (1003 \pm 20)$ mbar.

Exactly like for the near detector, the contents of the intermediate tank were large enough to reach the second CPS in the chimney, without the need to empty further the weighing

[6]The sensor at $z = 2360$ mm had a defect.

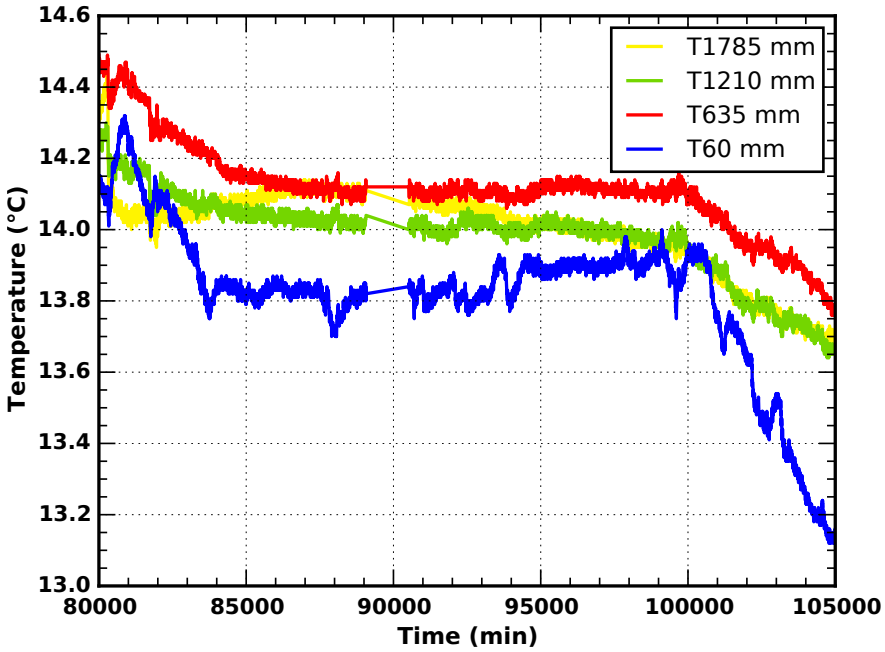

Figure 11.10 – Evolution of the temperature of the target sensors during the thermalisation stage. The labels of the sensors carry their position in the target along the upwards z-axis, whose origin is located at the bottom of the target. A non-negligible spread can be observed despite the long thermalisation time.

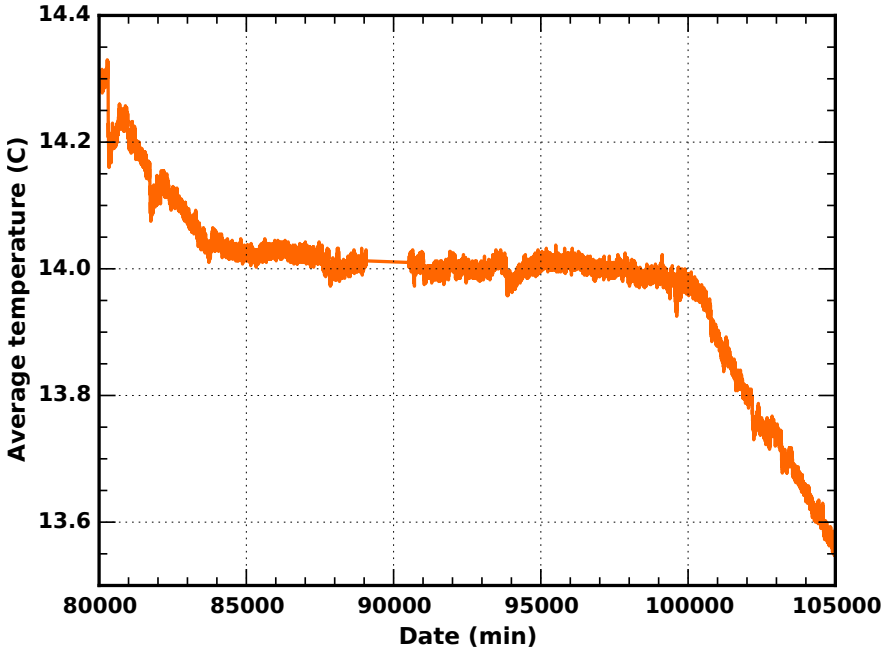

Figure 11.11 – Average temperature of the four working target sensors during the thermalisation stage.

tank itself. The leftovers were subsequently drained back into the weighing tank not to be counted as effectively in the target. Such an operation is embodied by Figure 11.12. Even though the tank was not lifted at that time, and accordingly bears the mark of a positive sensor drift, the 500 g tare was added to confirm the $M_1 = 2354\,\text{kg}$ value. The drift correction to this measurement will come from a later observation, namely that of the negative reading once the tank has been heaved at the end of the enterprise.

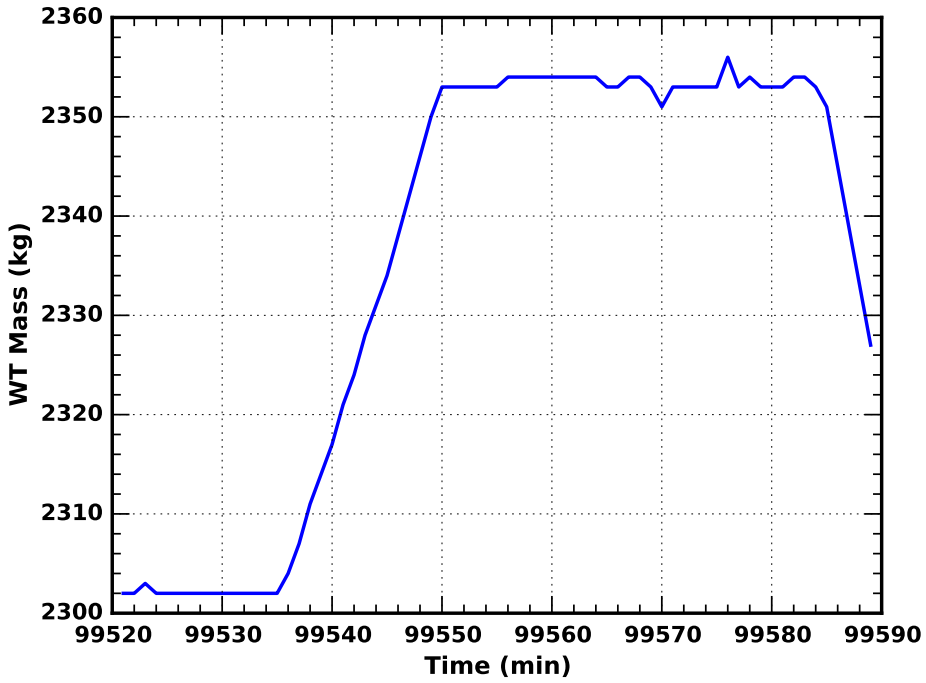

Figure 11.12 – Measured mass of the weighing tank during the final weight measurement. The contents of the intermediate tank are drained back into the weighing tank, yielding a final mass $M_1 = 2354\,\text{kg}$.

11.3.2 Target mass estimation

11.3.2.1 Sensor drift

Beyond question, correcting the measured masses M_0 and M_1 for the drift of the mass sensors is the thorniest task.

If the M_{break}^{drift} value from 11.3.1.1 is a good starting point for the correction on M_0, there remains to account for the unmeasured drifts before and after the loading breaks. The loading before the break lasts about 395 min and the one after, 415 min. These values are to be compared to the 8330 min duration of the break itself, or more importantly, to the first $\simeq 400\,\text{min}$ over which a $M_{400}^{drift} = (1 \pm 2)\,\text{kg}$ drift was observed. The conservative approach supposes that the read value bears a 2 kg uncertainty as in 11.2.2.1, effectively acknowledging

that the drifts can be overestimated. Inasmuch as the drift is smaller for lighter charges, yet larger for shorter period of times, the drift during the first load can be assigned the $M_{load_1}^{drift} = M_{400}^{drift}$ value. Likewise, the second load $M_{load_2}^{drift}$ can be attributed the same value. Therefore, M_0 must be corrected by

$$\begin{aligned} M_0^{drift} &= M_{break}^{drift} + M_{load_1}^{drift} + M_{load_2}^{drift} \\ &= (5.0 \pm 3.5)\,\text{kg}. \end{aligned} \tag{11.26}$$

As far as M_1 is concerned, we may start by stressing that at the very end of the filling (when the chimney and the expansion tank had been themselves filled), after the weighing tank had been heaved, the sensors exhibited a negative drift of amplitude $M_{lift}^{drift} = (7 \pm 2)\,\text{kg}$, thereby suggesting that a part of the $M_0^{drift} + M_{start}^{drift}$ had been absorbed during the filling. Indeed, in addition of the drift during the loading of the tank M_0^{drift}, we observed a 9 kg drift in Figure 11.8; had the drift retained its magnitude until the end of the filling, we would have expected a larger M_{lift}^{drift} value.

From this final measurement, we can go backwards to the measurement of M_1. Figure 11.13 tends to show that a $M_{end_1}^{drift} = (2 \pm 2)\,\text{kg}$ drift happens between 99.5×10^3 min and 102×10^3 min, boosting the eventual negative display. On the other hand, for over more than 5000 min, between 104×10^3 min and 109×10^3 min, a positive drift of $M_{end_2}^{drift} = (1 \pm 2)\,\text{kg}$ partly compensates the first negative drift.

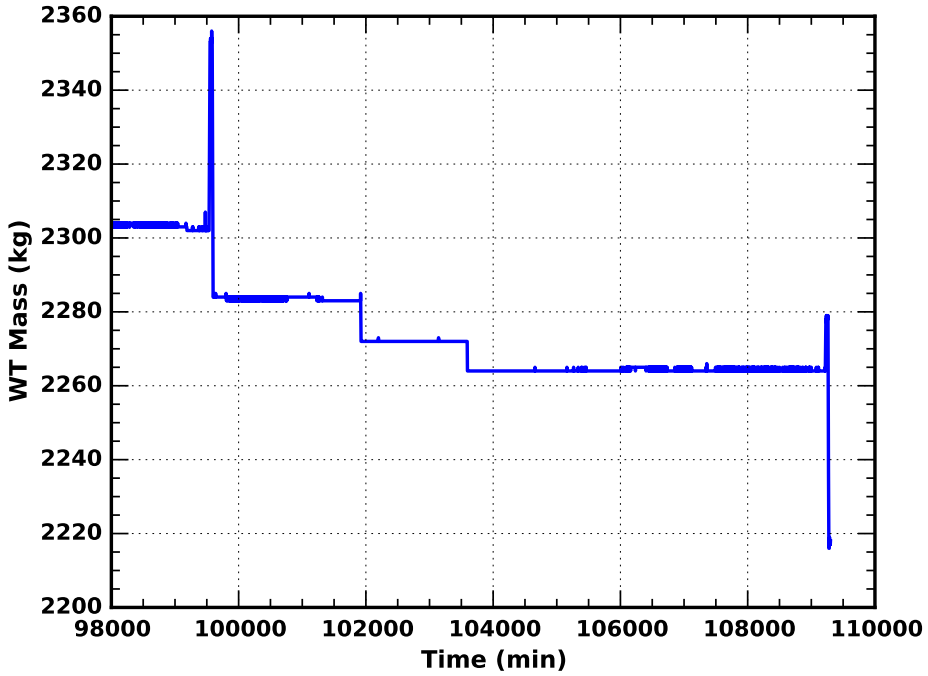

Figure 11.13 – Measured mass of the weighing tank during the filling of the chimney and the expansion tank. Partly compensating drifts can be observed whilst no liquid is being poured.

It follows from these statements that the correction to M_1 can be written

$$\begin{aligned} M_1^{drift} &= M_{lift}^{drift} - M_{end_1}^{drift} + M_{end_2}^{drift} \\ &= (6.0 \pm 3.5)\,\mathrm{kg}. \end{aligned} \tag{11.27}$$

It may seem as though the drifts can but lower the final target mass, this is too hasty a reasoning. Indeed, should both M_0 and M_1 be underestimated, the difference between the two would have little reason to be underestimated itself. Ultimately,

$$M_0 + M_0^{drift} = (10\,639.0 \pm 4.0)\,\mathrm{kg} \tag{11.28}$$

$$M_1 + M_1^{drift} = \quad (2360.0 \pm 4.0)\,\mathrm{kg}. \tag{11.29}$$

11.3.2.2 Gravity correction

The mass sensors were overhasty calibrated at Heidelberg, where the gravitational acceleration $g_{Hb} = (9.8093 \pm 0.0002)\,\mathrm{m\,s^{-2}}$ is slightly lower than at Chooz [149], for which $g_{Cz} = (9.8104 \pm 0.0002)\,\mathrm{m\,s^{-2}}$ was evaluated [150]. Consequently, the sensors at Chooz are slightly overestimating the weights they display. The appropriate correction factor to all the masses given previously is the ratio of the gravitational accelerations

$$r_g = \frac{g_{Hb}}{g_{Cz}} \tag{11.30}$$

$$= 0.999\,888 \pm 0.000\,029. \tag{11.31}$$

With the correction applied,

$$\left(M_0 + M_0^{drift}\right) r_g = (10\,637.8 \pm 4.0)\,\mathrm{kg} \tag{11.32}$$

$$\left(M_1 + M_1^{drift}\right) r_g = \quad (2359.7 \pm 4.0)\,\mathrm{kg}. \tag{11.33}$$

The reader shall notice that the corrected M_0 value is reassuringly similar to the one found for the near detector. As for M_1, the higher value at the far detector is unreservedly consistent with the 0.6% uncertainty on the geometrical dimensions of the target acrylic vessel. In a nutshell, the target volume of the far detector may well be smaller than that of the near detector.

11.3.2.3 Nitrogen

Once the unpleasant M_0 and M_1 values have been corrected for the sensor drift, the mass of the nitrogen excess comes unambiguously. For that matter, we need but swap M_0 for $M_0 + M_0^{drift}$ and M_1 for $M_1 + M_1^{drift}$ in (11.4) and (11.8).

As a matter of fact, the thermalisation temperature $T = (14.0 \pm 0.5)\,^{\circ}\mathrm{C}$ (cf: Figure 11.11) at the far detector corresponds to the temperature reference of ρ_{liquid}^0. Including σ_T in (11.6) slightly increases the uncertainty on ρ_{liquid} and we may confidently write

$$\rho_{liquid} = (804.0 \pm 1.3)\,\mathrm{kg\,m^{-3}}. \tag{11.34}$$

With the scant number of significant figures we need, the nitrogen density cautiously holds the same value as during the weight measurement at the near detector, that is to say

$$\rho_{\mathrm{N}_2}(T,P) = (1.18 \pm 0.03)\,\mathrm{kg\,m^{-3}}. \tag{11.35}$$

As a consequence, the nitrogen mass at the far detector is

$$M_{WT}^{\mathrm{N}_2}(T,P) = (12.15 \pm 0.34)\,\mathrm{kg}. \tag{11.36}$$

11.3.2.4 Filling tube

As is the case for the near detector, the tube connecting the IMT to the detector is the only part of the system that is empty when we measure M_0, yet full when M_1 is being measured. All the connections within the IMT itself are kept full throughout the procedure. If the diameter of the tube $d = (9.55 \pm 0.13)\,\mathrm{mm}$ is identical in both cases, the length at the far detector was certainly shorter, and probably by more than half, though its exact length on the deck cannot be recalled as of now. In the meantime, we shall take $l_{edge} = (10 \pm 8)\,\mathrm{m}$. The length from the chimney flange to the target is however identical in both detectors, or at least unmistakably within the uncertainty. Thus, the mass of the tube ought to be easily covered by

$$M_{tube}^{liquid} = (0.70 \pm 0.46)\,\mathrm{kg}. \tag{11.37}$$

11.3.2.5 Target mass value after filling

Combining results from previous subsections 11.3.2.1, 11.3.2.3 and 11.3.2.4—with no truncations in the analysis code—yields the mass of the target at $T = (14.0 \pm 0.5)\,^{\circ}\mathrm{C}$ and $P = (1003 \pm 20)\,\mathrm{mbar}$

$$M_{NT}^{liquid} = (8289.5 \pm 5.7)\,\mathrm{kg}. \tag{11.38}$$

11.3.3 Target mass evolution and number of protons

11.3.3.1 Overall evolution

The average temperature of the inner far detector over the nine-month data run list is largely covered by $T_{use} = (13.5 \pm 1.0)\,^{\circ}\mathrm{C}$, allowing one to scale the mass after the filling to the mass at which the detector has been operating, i.e.

$$M_{NT}^{liquid}(T_{use},P) = (8291.5 \pm 7.3)\,\mathrm{kg}. \tag{11.39}$$

11.3.3.2 Number of protons

The liquids in either detector being identical, the same hydrogen mass fraction $f_H = 0.1360 \pm 0.0004$ is used here. In spite of the poorer accuracy on the target mass of the far detector, it must be emphasised that its relative uncertainty remains lower than 0.1%, which ensures that

the absolute normalisation at the far detector still approaches a 0.3% relative uncertainty. Eventually,

$$n_H^{FD}\left(T_{use}^{FD}, P^{FD}\right) = (6.739 \pm 0.021) \times 10^{29}. \tag{11.40}$$

The ratio of the number of protons at the far detector over that at the near detector is a particularly useful quantity as it rules out the otherwise dominating uncertainty on the chemical composition of the liquid, hence leaving a ratio driven by the weight measurements. The ratio reads

$$r_H\left(T_{use}^{FD}, T_{use}^{ND}, P^{FD}, P^{ND}\right) = \frac{n_H^{FD}}{n_H^{ND}} \tag{11.41}$$

$$= \frac{M_{NT}^{FD}}{M_{NT}^{ND}}, \tag{11.42}$$

where we have explicitly dropped the dependencies and marked the target masses of each detector with an upper-script. The uncertainty on this ratio can be expressed as

$$\sigma_{r_H} = r_H \sqrt{\left(\frac{\sigma_{M_{NT}^{FD}}}{M_{NT}^{FD}}\right)^2 + \left(\frac{\sigma_{M_{NT}^{ND}}}{M_{NT}^{ND}}\right)^2}, \tag{11.43}$$

such that

$$r_H = 0.9958 \pm 0.0010. \tag{11.44}$$

11.4 Main achievements

In this chapter, the weight measurement principle has been quickly reviewed, the values measured have been given and exploited for the subsequent and thorough error treatment. A final value for the mass of liquid scintillator in the neutrino target of the Double Chooz near detector has been estimated at the filling temperature and atmospheric pressure. This value has eventually been scaled up to account for the temperature evolution inside the inner detector over the year 2015. All the data from the weight measurement at the far laboratory have also been re-evaluated from scratch with an unexampled analysis of the uncertainties.

Table 11.1 summarises the key results of this work. In all truth, the later exceed the Double Chooz expectations of a 0.2% relative uncertainty on the target masses (cf. 11.1.2), the ratio of the independent near and far detector masses being known to 0.1%.

Target	M^{weight} (kg)	M^{use} (kg)	n_H^{use}	r_H
ND	8314.7 ± 2.7	8326.5 ± 3.8	$(6.767 \pm 0.020) \times 10^{29}$	0.9958 ± 0.0010
FD	8289.5 ± 5.7	8291.5 ± 7.3	$(6.739 \pm 0.021) \times 10^{29}$	

Table 11.1 – Neutrino target masses and number of protons. The first column represents the weight measurement masses, the second one corresponds to the masses for the data taking over the year 2015 (the temperature uncertainties also encompass the single-detector phase). The third column gives the corresponding number of target protons and the fourth shows their ratio. The relevant temperatures and pressures for each detector can be found in 11.2.3 and 11.3.3.

Conclusion

However shielded, neutrino experiments are fundamentally limited by backgrounds. Be they carried out inside the containment building of a reactor radiating neutrons and γ-rays or deep underground, studies of the neutral leptons go hand in hand with background rejection techniques. Should they fail at flawlessly tagging all their background events, experiments must model or quantify the remaining contamination.

When it comes to the Double Chooz experiment, the cosmogenic background – generated via muon spallation – dominates and hinders the precision of the θ_{13} measurement. With the scant statistics available three years ago for the sole far detector set, simulating the decays of the ^{9}Li and ^{8}He cosmogenic isotopes polluting the $\overline{\nu_e}$ signal was a natural undertaking. For all the statistics accumulated since these first endeavours, the simulated spectra still compare remarkably well with the data. These achievements were not taken for granted because the nuclear data themselves are scarce, and the intermediate states involved in the decay chains, peculiar. To cope with the nuclear uncertainties, an industrious error treatment was performed, yielding covariance matrices embodying the correlations between the different energy regions of the simulated spectra, information particularly handy when fitting the predictions to the cosmogenic data spectra.

By way of more permissive analysis techniques – combining neutron-captures on both Gd and H – and a near detector, the number of events that could be confidently tagged as cosmogenic decays has increased sixfold with respect to the latest Double Chooz publication. This amount of data was retrieved by the author, who adapted and systematically studied the performance of the veto aimed at tagging the βn-emitters, a task all the more delicate at the near detector that the neutral leptons account for the overwhelming majority of the correlated pairs observed at the near site, thereby overshadowing important background events. The energy distribution of the vetoed events does not strictly correspond to βn-decays, the $\overline{\nu_e}$ contamination need first be removed; the background subtracted vetoed spectra of both detectors were tested for compatibility and subsequently combined. Via a fit of the Monte-Carlo predictions to this sum, no evidence for the production and decay of ^{8}He was found, which is in agreement with observations from other experiments at deeper depths.

Although the cosmogenic veto provides events that may be used to build the precious data spectra, particularly at the near detector, its tagging efficiency is limited. In consequence, the rate of background events that remained within the $\overline{\nu_e}$ candidates must be evaluated for all detectors. To this end, the Monte-Carlo spectra were used to deduce the efficiency of

prompt energy cuts, highly increasing the cosmogenics over $\overline{\nu_e}$ ratio. In view of the large μ-rate at the near detector, this cut had to be combined with a selection of muons more likely to have produced βn-emitters, i.e. showering muons. The efficiency of this selection was gauged via the cleaner far sample. The results thus obtained were successfully cross-checked against former far detector results, and against independent estimations from the oscillation fit itself.

For an analysis based on the Gd channel, the uncertainty on the background determination (itself dominated by ^{9}Li decays) is the most significant impairment of the experiment, but only because the normalisation of the $\overline{\nu_e}$-rate detected at each site is well-controlled. When binding both detectors via correlations, the uncertainty dominating each normalisation is the number of target protons. As both targets are filled with identical liquids, the correlation between the two is maximal, if and only if their volumes are perfectly known. There lies the stumbling block dramatically limiting the accuracy of the Gd++ analysis, which simultaneously relies on the target and gamma catcher. On the other hand, for the Gd analysis, the detections systematics on the masses of liquid scintillator poured into the targets have been estimated with utmost diligence in this thesis, by means of weight measurements, yielding uncertainties thrice smaller than those of the chemical composition of the liquids, albeit conservative.

All this work has been part of the first multi-detector analysis, whose convoluted framework was thoroughly reviewed throughout this document, with up-to-date information on all fronts. This framework enabled the collaboration to disclose a puzzling $\sin^2(2\theta_{13}) = 0.111 \pm 0.018$ value in March 2016, at the Moriond conference. In less than five years, the status of the smallest mixing angle went from unknown to the best measured of the three, the Daya Bay competitors boasting a $\sin^2(2\theta_{13}) = 0.0841 \pm 0.0033$ value since July 2016, i.e. a relative uncertainty smaller than 5%.

The field of neutrino physics is an ever-evolving one. In fact, its face changes with such a breathtaking liveliness that in between the beginning of the writing of this thesis and the relief setting these final words in the digital stone embodies, exclusion contours have closed in on the space of physical parameters. Beyond doubt, the pressure of the famed Neutrino 2016 conference enticed analysers to shed more light onto the CP violation phase δ_{CP} and the mass hierarchy. For all the fanciness of many theories, it seems more and more likely that the ordering of the neutrino masses follows the plain one that their charged partners exhibit. In a like manner, δ_{CP} seemingly points to anything but zero, and it could actually be maximal, which would surely pave the way for explaining the staggering asymmetry between the amount of matter and antimatter in the world surrounding us.

Nevertheless, some doubts remain as to which scenario is preferred, all the more so that a larger θ_{13} value leaves more room for the inverted mass hierarchy and fickle δ_{CP} values. As the error bars dwindle, discrepancies between experiments measuring $\sin^2(2\theta_{13})$ are scarcely startling. Perhaps now more than ever, redundancy and cooperation for refining the value of the smallest mixing angle are of paramount importance.

Bibliography

[1] A. H. Becquerel. 'On the rays emitted by phosphorescence'. In: *Comptes Rendus* 122 (1896), pp. 420–421.

[2] J. Chadwick. 'The intensity distribution in the magnetic spectrum of β particles from radium (B + C)'. In: *Verh. Phys. Gesell.* 16 (1914), pp. 383–391.

[3] W. Pauli. 'Dear radioactive ladies and gentlemen'. In: *Phys. Today* 31N9 (1978), p. 27.

[4] J. Chadwick. 'Possible existence of a neutron'. In: *Nature* 129 (1932), p. 312.

[5] E. Fermi. 'Tentativo di una teoria dell'emissione dei raggi beta'. In: *La Ricerca Scientifica* II (1933), p. 12.

[6] C.M.G Lattes et al. 'Observation on the Tracks of Slow Mesons in Photographic Emulsions'. In: *Nature* 160 (1947), p. 453.

[7] R. B. Leighton et al. 'The Energy Spectrum of the Decay Particles and the Mass and Spin of the Mesotron'. In: *Phys. Rev.* 75 (9 May 1949), pp. 1432–1437.

[8] R. Davis. 'Attempt to Detect the Antineutrinos from a Nuclear Reactor by the $Cl^{37}(\overline{\nu}, e^-)A^{37}$ Reaction'. In: *Phys. Rev.* 97 (3 Feb. 1955), pp. 766–769.

[9] G. Fidecaro. 'Bruno Pontecorvo: from Rome to Dubna'. In: (Dec. 1996), pp. 472–486.

[10] F. Reines and C. L. Cowan. 'Detection of the Free Neutrino'. In: *Phys. Rev.* 92 (3 Nov. 1953), pp. 830–831.

[11] C. L. Cowan et al. 'Detection of the Free Neutrino: a Confirmation'. In: *Science* 124.3212 (1956), pp. 103–104.

[12] E. J. Konopinski and H. M. Mahmoud. 'The Universal Fermi Interaction'. In: *Phys. Rev.* 92 (4 Nov. 1953), pp. 1045–1049.

[13] E. P. Hincks and B. Pontecorvo. 'Search for gamma-radiation in the 2.2-microsecond meson decay process'. In: *Phys. Rev.* 73 (1948), pp. 257–258.

[14] B. Pontecorvo. 'Electron and Muon Neutrinos'. In: *Sov. Phys. JETP* 10 (1960). [Zh. Eksp. Teor. Fiz.37,1751(1959)], pp. 1236–1240.

[15] G. Danby et al. 'Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos'. In: *Phys. Rev. Lett.* 9 (1962), pp. 36–44.

[16] K. Kodama et al. 'Observation of tau-neutrino interactions'. In: *Phys. Lett.* B504 (2001), pp. 218–224.

[17] T. D. Lee and C. N. Yang. 'Question of Parity Conservation in Weak Interactions'. In: *Phys. Rev.* 104 (1 Oct. 1956), pp. 254–258.

[18] C. S. Wu et al. 'Experimental Test of Parity Conservation in Beta Decay'. In: *Phys. Rev.* 105 (1957), pp. 1413–1414.

[19] M. Goldhaber, L. Grodzins and A. W. Sunyar. 'Helicity of Neutrinos'. In: *Phys. Rev.* 109 (3 Feb. 1958), pp. 1015–1017.

[20] J. H. Christenson et al. 'Evidence for the 2π Decay of the K_2^0 Meson'. In: *Phys. Rev. Lett.* 13 (4 July 1964), pp. 138–140.

[21] L. D. Landau. 'On the conservation laws for weak interactions'. In: *Nucl. Phys.* 3 (1957), pp. 127–131.

[22] R. P. Feynman and M. Gell-Mann. 'Theory of the Fermi Interaction'. In: *Phys. Rev.* 109 (1 Jan. 1958), pp. 193–198.

[23] R. Davis. 'Solar neutrinos. II: Experimental'. In: *Phys. Rev. Lett.* 12 (1964), pp. 303–305.

[24] R. Davis et al. 'Search for neutrinos from the sun'. In: *Phys. Rev. Lett.* 20 (1968), pp. 1205–1209.

[25] J. C. Evans et al. 'Brookhaven solar neutrino detector and collapsing stars'. In: *Nature* 251 (1974), pp. 486–488.

[26] J.N. Bahcall et al. 'The Rate Of The Proton-proton Reaction And Some Related Reactions'. In: *Astrophys. J.* 155 (1969), pp. 501–510.

[27] J.N. Bahcall et al. 'Sensitivity of the Solar-Neutrino Fluxes'. In: *Astrophys. J.* 156 (1969), pp. 559–568.

[28] John N. Bahcall et al. 'Solar Neutrinos - a Scientific Puzzle'. In: *Science* 191 (1976), pp. 264–267.

[29] K. S. Hirata et al. 'Results from one thousand days of real-time, directional solar-neutrino data'. In: *Phys. Rev. Lett.* 65 (11 Sept. 1990), pp. 1297–1300.

[30] K. Lande et al. 'Results from the Homestake solar neutrino observatory'. In: *Conf. Proc.* C900802 (1990), pp. 867–675.

[31] V. N. Gavrin et al. 'First measurement of the integral solar neutrino flux by the Soviet/American Gallium Experiment'. In: *15th Texas Symposium on Relativistic Astrophysics and 4th ESO-CERN Symposium Brighton, England, December 16-21, 1990*. 1990.

[32] P. Anselmann et al. 'Solar neutrinos observed by GALLEX at Gran Sasso.' In: *Phys. Lett.* B285 (1992), pp. 376–389.

[33] J. Boger et al. 'The Sudbury neutrino observatory'. In: *Nucl. Instrum. Meth.* A449 (2000), pp. 172–207.

[34] Q. R. Ahmad et al. 'Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory'. In: *Phys. Rev. Lett.* 89 (2002), p. 011301.

[35] Y. Fukuda et al. 'Evidence for Oscillation of Atmospheric Neutrinos'. In: *Phys. Rev. Lett.* 81 (8 Aug. 1998), pp. 1562–1567.

[36] K. Eguchi et al. 'First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance'. In: *Phys. Rev. Lett.* 90 (2 Jan. 2003), p. 021802.

[37] T. Araki et al. 'Measurement of Neutrino Oscillation with KamLAND: Evidence of Spectral Distortion'. In: *Phys. Rev. Lett.* 94 (8 Mar. 2005), p. 081801.

[38] A. Gando et al. 'Reactor on-off antineutrino measurement with KamLAND'. In: *Phys. Rev. D* 88 (3 Aug. 2013), p. 033001.

[39] K. Abe et al. 'Indication of Electron Neutrino Appearance from an Accelerator-Produced Off-Axis Muon Neutrino Beam'. In: *Phys. Rev. Lett.* 107 (4 July 2011), p. 041801.

[40] P. Adamson et al. 'Combined Analysis of ν_μ Disappearance and $\nu_\mu \to \nu_e$ Appearance in MINOS Using Accelerator and Atmospheric Neutrinos'. In: *Phys. Rev. Lett.* 112 (19 May 2014), p. 191801.

[41] B. Pontecorvo. 'Mesonium and anti-mesonium'. In: *Sov. Phys. JETP* 6 (1957), p. 429.

[42] B. Pontecorvo. 'Inverse beta processes and non-conservation of lepton charge'. In: *Sov. Phys. JETP* 7 (1958), pp. 172–173.

[43] Z. Maki et al. 'Remarks on the unified model of elementary particles'. In: *Prog. Theor. Phys.* 28 (1962).

[44] S. Eliezer and A. Swift. 'Experimental consequences of $\nu_\mu \leftrightarrows \nu_e$ mixing in neutrino beams'. In: *Nuclear Physics B* 105.1 (1976), pp. 45–51.

[45] H. Fritsch and P. Minkowski. 'Vector-like weak currents, massive neutrinos, and neutrino beam oscillations'. In: *Physics Letters B* 62.1 (1976), pp. 72–76.

[46] Haim Harari and Miriam Leurer. 'Recommending a standard choice of Cabibbo angles and KM phases for any number of generations'. In: *Physics Letters B* 181.1 (1986), pp. 123–128.

[47] E. Kh. Akhmedov and J. Kopp. 'Neutrino oscillations: Quantum mechanics vs. quantum field theory'. In: *JHEP* 04 (2010), p. 008.

[48] C. Giunti. 'Theory of neutrino oscillations'. In: *Particle physics in laboratory, space and universe. Intelligentsia and education. Proceedings, 11th Lomonosov Conference on elementary particle physics, Moscow, Russia, August 21-27, 2003, and 5th International Meeting on problems of intelligentsia, Moscow, Russia, August 27, 2003*. 2004, pp. 35–44. arXiv: hep-ph/0401244.

[49] Ch. Kraus et al. 'Final results from phase II of the Mainz neutrino mass search in tritium β decay'. In: *The European Physical Journal C - Particles and Fields* 40.4 (2005), pp. 447–468.

[50] E. Di Valentino et al. 'Cosmological axion and neutrino mass constraints from Planck 2015 temperature and polarization data'. In: *Physics Letters B* 752 (2016), pp. 182–185.

[51] K. A. et al. Olive. 'Review of Particle Physics'. In: *Chin. Phys.* C38 (2015), p. 090001.

[52] X. Qian and P. Vogel. 'Neutrino mass hierarchy'. In: *Progress in Particle and Nuclear Physics* 83 (2015), pp. 1–30.

[53] S. F. King. 'Neutrino Mass and Mixing in the Seesaw Playground'. In: *Nucl. Phys.* B908 (2016), pp. 456–466. arXiv: 1511.03831.

[54] M. Apollonio et al. 'Search for neutrino oscillations on a long baseline at the CHOOZ nuclear power station'. In: *Eur. Phys. J.* C27 (2003), pp. 331–374.

[55] Y. Abe et al. 'Indication of Reactor $\overline{\nu}_e$ Disappearance in the Double Chooz Experiment'. In: *Phys. Rev. Lett.* 108 (13 Mar. 2012), p. 131801.

[56] F. P. An et al. 'Observation of Electron-Antineutrino Disappearance at Daya Bay'. In: *Phys. Rev. Lett.* 108 (17 Apr. 2012), p. 171803.

[57] K. Anderson et al. 'White paper report on Using Nuclear Reactors to Search for a value of θ_{13}'. In: (2004). arXiv: hep-ex/0402041 [hep-ex].

[58] F. P. An et al. 'New Measurement of Antineutrino Oscillation with the Full Detector Configuration at Daya Bay'. In: *Phys. Rev. Lett.* 115 (11 Sept. 2015), p. 111802.

[59] J. K. Ahn. 'Observation of Reactor Electron Antineutrinos Disappearance in the RENO Experiment'. In: *Phys. Rev. Lett.* 108 (19 May 2012), p. 191802.

[60] Y. Abe et al. 'Improved measurements of the neutrino mixing angle θ_{13} with the Double Chooz detector'. In: *Journal of High Energy Physics* 2014.10 (2014), pp. 1–44.

[61] P. Adamson et al. 'First Measurement of Electron Neutrino Appearance in NOvA'. In: *Phys. Rev. Lett.* 116 (15 Apr. 2016), p. 151806.

[62] V. I. Kopeikin et al. 'Reactor as a source of antineutrinos: Thermal fission energy'. In: *Physics of Atomic Nuclei* 67.10 (2004), pp. 1892–1899.

[63] A. Onillon. 'Prédiction des taux de fission des coeurs de Chooz et estimation des incertitudes associées dans le cadre de l'expérience Double Chooz'. PhD thesis. Ecole des Mines de Nantes, Nov. 2014, pp. 103–106, 57–59.

[64] M.B. Chadwick et al. 'Special Issue on ENDF/B-VII.1 Library ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data'. In: *Nuclear Data Sheets* 112.12 (2011), pp. 2887–2996. ISSN: 0090-3752.

[65] Krzysztof Miernik. *Chart of nuclides drawer*. 2016. URL: https://github.com/kmiernik/Chart-of-nuclides-drawer (visited on 25th June 2016).

[66] X. B. Ma et al. 'Improved calculation of the energy release in neutron-induced fission'. In: *Phys. Rev. C* 88 (1 July 2013), p. 014605.

[67] Y. Abe et al. 'Ortho-positronium observation in the Double Chooz experiment'. In: *Journal of High Energy Physics* 2014.10 (2014), pp. 1–17.

[68] Huang Ming-Yang, Guo Xin-Heng and Young Bing-Lin. 'Detection of supernova neutrinos at spallation neutron sources'. In: *Chin. Phys.* C40 (2016), p. 073102.

[69] A. Collin. 'Étude des antineutrinos de réacteurs : mesure de l'angle de mélange leptonique θ_{13} et recherche d'éventuels neutrinos stériles'. PhD thesis. Paris XI, Jan. 2014, pp. 127–144.

[70] P. Vogel et al. 'Angular distribution of neutron inverse beta decay, $\overline{\nu}_e + \vec{p} \; e^+ + n$'. In: *Phys. Rev. D* 60 (5 July 1999), p. 053003.

[71] C. Bemporad et al. 'Reactor-based neutrino oscillation experiments'. In: *Rev. Mod. Phys.* 74 (2 Mar. 2002), pp. 297–328.

[72] D. Dietrich et al. 'Monte Carlo aided design of the inner muon veto detectors for the Double Chooz experiment'. In: *Journal of Instrumentation* 7.08 (2012), P08012.

[73] C. Aberle. 'Large scale Gd-beta-diketonate based organic liquid scintillator production for antineutrino detection'. In: *Journal of Instrumentation* 7.06 (2012), P06008.

[74] H.O. Back et al. 'Study of phenylxylylethane (PXE) as scintillator for low energy neutrino experiments'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 585.1–2 (2008), pp. 48–60.

[75] A. Minotti. 'Exploitation of pulse shape analysis for correlated background rejection and ortho-positronium identification in the Double Chooz experiment'. PhD thesis. Université de Strasbourg, Oct. 2015, pp. 49–50.

[76] K. Zbiri. 'Note on Drexel tests of the IMB R1408 PMTs used in the inner veto of both far and near detectors of the Double Chooz experiment'. In: (2011). arXiv: 1104.4045.

[77] R. Carr. 'Measurements of Electron Antineutrino Disappearance in the Double Chooz Experiment'. PhD thesis. Columbia University, Aug. 2015, pp. 62–63, 178–183.

[78] Brookhaven National Laboratory. *National Nuclear Data Center*. 2016. URL: http://www.nndc.bnl.gov/ (visited on 1st Sept. 2016).

[79] E. J. Axton and A. G. Bardell. 'Neutron Yield from the Spontaneous Fission of 252 Cf'. In: *Metrologia* 21.2 (1985), p. 59.

[80] Y. Abe et al. 'The waveform digitiser of the Double Chooz experiment: performance and quantisation effects on photomultiplier tube signals'. In: *Journal of Instrumentation* 8.08 (2013), P08015.

[81] F. Beissel et al. 'The trigger and timing system of the Double Chooz experiment'. In: *Journal of Instrumentation* 8.01 (2013), T01003.

[82] Y. Abe et al. 'Reactor $\overline{\nu}_e$ disappearance in the Double Chooz experiment'. In: *Phys. Rev. D* 86 (5 Sept. 2012), p. 052008.

[83] E. Conover. 'Muon-Induced Backgrounds in the Double Chooz Neutrino Oscillation Experiment'. PhD thesis. University of Chicago, June 2014, pp. 33–34, 78–117.

[84] O. Méplan et al. 'MCNP Utility for Reactor Evolution - Description of the methods, first applications and results'. In: *European Nuclear Society* (2005), pp. 1–7.

[85] Nuclear Energy Agency. *MURE, MCNP Utility for Reactor Evolution: couples Monte-Carlo transport with fuel burnup calculations*. 2016. URL: http://www.oecd-nea.org/tools/abstract/detail/nea-1845 (visited on 25th July 2016).

[86] C. L. Jones et al. 'Reactor simulation for antineutrino experiments using DRAGON and MURE'. In: *Phys. Rev. D* 86 (1 July 2012), p. 012001.

[87] A. Onillon. *Reactor data status and effective prediction of DC reactor fission fraction, DocDB 6357*. Talk. Double Chooz, Nov. 2015.

[88] K. Schreckenbach et al. 'Determination of the antineutrino spectrum from 235U thermal neutron fission products up to 9.5 MeV'. In: *Physics Letters B* 160.4 (1985), pp. 325–330.

[89] A.A. Hahn et al. 'Antineutrino spectra from 241Pu and 239Pu thermal neutron fission products'. In: *Physics Letters B* 218.3 (1989), pp. 365–368.

[90] N. Haag et al. 'Experimental Determination of the Antineutrino Spectrum of the Fission Products of ^{238}U'. In: *Phys. Rev. Lett.* 112 (12 Mar. 2014), p. 122501.

[91] Patrick Huber. 'Determination of antineutrino spectra from nuclear reactors'. In: *Phys. Rev. C* 84 (2 Aug. 2011), p. 024617.

[92] T. Mueller. 'Expérience Double Chooz : simulation des spectres antineutrinos issus de réacteurs'. PhD thesis. Paris-Sud (Paris XI), Sept. 2010, pp. 114–127.

[93] Y. Declais et al. 'Study of reactor anti-neutrino interaction with proton at Bugey nuclear power plant'. In: *Phys. Lett.* B338 (1994), pp. 383–389.

[94] J. Allison et al. 'GEANT4 developments and applications'. In: *Nuclear Science, IEEE Transaction* 51 (2006), pp. 270–278.

[95] JS. Agostinelli et al. 'Geant4 – A Simulation Toolkit'. In: *Nuclear Instruments and Methods* A506 (2003), pp. 250–303.

[96] G. Horton-Smith. *An introduction to GLG4sim features*. 2006. URL: http://neutrino.phys.ksu.edu/~GLG4sim/ (visited on 1st Sept. 2016).

[97] C. Aberle. ' Optimization, simulation and analysis of the scintillation signals in the Double Chooz experiment'. PhD thesis. Ruprecht-Karls-Universität, Dec. 2011, pp. 58–108, 229–239.

[98] C. Aberle et al. 'Light output of Double Chooz scintillators for low energy electrons'. In: *Journal of Instrumentation* 6.11 (2011), P11006.

[99] S. M. Seltzer and M. J. Berger. 'Evaluation of the collision stopping power of elements and compounds for electrons and positrons'. In: *The International Journal of Applied Radiation and Isotopes* 33.11 (1982), pp. 1189–1218.

[100] I. Stancu. *The Double Chooz Optical Model (MC), DocDB 2897*. Note. Double Chooz, Aug. 2011.

[101] E. Chauveau. *Energy scale status, DocDB 6646*. Talk. Double Chooz, Apr. 2016.

[102] E. Chauveau. *Energy Scale v11, DocDB 6741*. Talk. Double Chooz, May 2016.

[103] E. Chauveau. *ESv11 Blessing package, DocDB 6827*. Talk. Double Chooz, July 2016.

[104] E. Chauveau. *Update of DC-IV Energy Model, DocDB 6850*. Talk. Double Chooz, July 2016.

[105] I. Antcheva et al. 'ROOT — A C++ framework for petabyte data storage, statistical analysis and visualization'. In: *Computer Physics Communications* 180 (2009), pp. 2499–2512.

[106] R. Sharankova. *Status of T0 calibration constants w/ IDLI, DocDB 6452*. Talk. Double Chooz, Jan. 2016.

[107] P.-J. Chang et al. *Muon measurements in Double Chooz, DocDB 4540*. Note. Double Chooz, May 2013.

[108] T. Brugière. *IV Tag and muon definition, DocDB 6363*. Talk. Double Chooz, Nov. 2015.

[109] E. Chauveau. *Muon veto for FD+ND, DocDB 6397*. Talk. Double Chooz, Sept. 2015.

[110] Y. Abe et al. 'Measurement of θ_{13} in Double Chooz using neutron captures on hydrogen with novel background rejection techniques'. In: *Journal of High Energy Physics* 2016.1 (2016), pp. 1–29.

[111] Y. Abe et al. 'Characterization of the spontaneous light emission of the PMTs used in the Double Chooz experiment'. In: *Journal of Instrumentation* 11.08 (2016), P08001.

[112] E. Chauveau. *Light Noise rejection for FD+ND on IBD candidates and ^{252} Cf calibration data, DocDB 6390*. Talk. Double Chooz, Dec. 2015.

[113] E. Chauveau. *Mini Data summary, DocDB 5908*. Talk. Double Chooz, Dec. 2014.

[114] C. Palomares. *Unicity Condition, DocDB 6692*. Talk. Double Chooz, May 2016.

[115] M. Ishitsuka. *New Results of Double Chooz*. Talk. 51st Rencontres de Moriond. Mar. 2016.

[116] M. Ishitsuka. *DC-IV Rate+Shape fit inputs for Moriond, DocDB 6495*. Files. Double Chooz, Mar. 2016.

[117] A. Hourlier. *DC IV FVV tuning, DocDB 6415*. Talk. Double Chooz, Dec. 2015.

[118] A. Hourlier. *DCIV SM rejection, DocDB 6704*. Talk. Double Chooz, May 2016.

[119] A. Meregaglia. *Correlated BG fits and rates, DocDB 6508*. Files. Double Chooz, Feb. 2016.

[120] G. Yang. 'Measurement of θ_{13} in the Double Chooz experiment'. PhD thesis. Illinois Institute of Technology, July 2016, pp. 122–182.

[121] E. Chauveau. *DC-IV Energy Systematics and R+S Treatment, DocDB 6796*. Talk. Double Chooz, June 2016.

[122] R. Hagedorn. *Relativistic Kinematics*. W. A. Benjamin, 1963, pp. 89–97.

[123] F. James. *Monte Carlo Phase Space*. CERN, 1968.

[124] D. Griffiths. *Introduction to Elementary Particles*. Wiley-VCH, 2008.

[125] D.R. Tilley et al. 'Energy levels of light nuclei $A = 8, 9, 10$'. In: *Nuclear Physics* A487 (2004), pp. 152–241.

[126] L. Winslow. *Cosmic Muon Spallation and Production of Radioactive Isotopes in KamLAND*. Tech. rep. UCLA.

[127] F.C. Barker and E.K. Warburton. 'The beta-decay of ^{8}He'. In: *Nuclear Physics* A487 (1988), pp. 269–278.

[128] G. Cacciapaglia et al. 'Nearby resonances beyond the Breit–Wigner approximation'. In: *Physics* B682 (2009), pp. 43–49.

[129] M. Vivier. *Estimation of GC and NT far/near detector volumes, DocDB 6768*. Talk. Double Chooz, June 2016.

[130] P.J. Plauger et al. *C++ Standard Template Library*. 1st. Upper Saddle River, NJ, USA: Prentice Hall PTR, 2000. ISBN: 0134376331.

[131] B. Schling. *The Boost C++ Libraries*. XML Press, 2011. ISBN: 0982219199.

[132] G. Nyman et al. '"The beta decay of ^{9}Li to levels in ^{9}Be": A new look"'. In: *Nuclear Physics* A510 (1990), pp. 189–208.

[133] V. V. Buldygin and Yu. V. Kozachenko. *Metric characterization of random variables and random processes (Translations of Mathematical Monographs)*. American Mathematical Society, 2000.

[134] B. Jacob and G. Guennebaud. *Eigen C++ template library for linear algebra*. 2016. URL: http://eigen.tuxfamily.org/dox/ (visited on 2nd Oct. 2016).

[135] E. Greuling and N. Huffaker. 'Section Of Physical Sciences: Effect Of "Weak Magnetism" On Beta-Radioactivity'. In: *Transactions of the New York Academy of Sciences* 24.5 Series II (1962), pp. 591–602. ISSN: 2164-0947.

[136] A.C. Hayes et al. '"Systematic Uncertainties in the Analysis of the Reactor Neutrino Anomaly"'. In: *Phys.Rev.Lett.* 112 (2014), p. 202501.

[137] S. Abe et al. 'Production of radioactive isotopes through cosmic muon spallation in KamLAND'. In: *Phys. Rev. C* 81 (2 Feb. 2010), p. 025807.

[138] Sh. Grant and R. Voorhies. *Cereal - a C++11 library for serialisation*. 2016. URL: http://uscilab.github.io/cereal/index.html (visited on 23rd Sept. 2016).

[139] R.H.C. Lopes et al. 'The two-dimensional Kolmogorov-Smirnov test'. In: *XI International Workshop on Advanced Computing and Analysis Techniques in Physics Research, Nikhef, Amsterdam, the Netherlands, April 23-27*. 2007.

[140] G. Bellini et al. 'Cosmogenic Backgrounds in Borexino at 3800 m water-equivalent depth'. In: *Journal of Cosmology and Astroparticle Physics* 2013.08 (2013), p. 049.

[141] L. F. F. Stokes. 'Cosmogenic Radioisotopes in the Double Chooz Far Detector'. PhD thesis. Eberhard Karls Universität Tübingen, Oct. 2015, pp. 67–78.

[142] R. Sharankova. *Gd-III: 12B background estimation, DocDB 5619*. Talk. Double Chooz, May 2014.

[143] T. Lasserre T. Mueller J. - C. Barrière. *Target H Measurement Physics, DocDB 1730*. Tech. rep. Double Chooz, July 2010.

[144] M. Göger C. Buck M.Franke. *Filling procedure for the Double Chooz near detector, DocDB 5740*. Tech. rep. Double Chooz, Sept. 2014.

[145] T. Lasserre J. - C. Barrière. *Target H Measurement, DocDB 3281*. Talk. Double Chooz, Oct. 2011.

[146] B. Wischnewski. *Peace Software*. 2016. URL: http://www.peacesoftware.de/einigewerte/stickstoff_e.html (visited on 25th Sept. 2016).

[147] H. Warlimont W. Martienssen. *Handbook of Condensed Matter and Materials Data*. Springer, 2005, pp. 497–499.

[148] SpecialChem. *Omnexus*. 2016. URL: http://omnexus.specialchem.com/polymer-properties/properties/coefficient-of-linear-thermal-expansion (visited on 7th Sept. 2016).

[149] Bureau de Recherche en Géologie Minière. 'Chooz gravity'. Contact. 2011.

[150] Bundesamt fuer Kartographie und Geodaesie. 'Heidelberg gravity'. Contact. 2011.

[151] A. J. Franke. 'Searching for Reactor Antineutrino Flavor Oscillations with the Double Chooz Far Detector'. PhD thesis. Columbia University, Dec. 2012, pp. 229–239.

[152] T. Mueller et al. 'Improved Predictions of Reactor Antineutrino Spectra'. In: *Phys. Rev.* C83 (2011), p. 054615.

[153] C. Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford, 2007.

[154] T. Hagner. 'Muon-induced production of radioactive isotopes in scintillation detectors'. In: *Astroparticle Physics* 14 (2000), pp. 33–47.

[155] Y. Prezado et al. 'Large asymmetry in the strongest -transition for A=9'. In: *Phys. Lett.* B576 (2003), pp. 55–61.

[156] L. Winslow and M. Elnimir. *DC2ndPub: The ^{9}Li Spectrum, DocDB 3986*. Talk. May 2012.

[157] T. Mueller J. - C. Barrière T. Lasserre. *Target liquid mass determination for the Double Chooz far detector*. Tech. rep. Double Chooz, Oct. 2011.

université PARIS-SACLAY | ÉCOLE DOCTORALE PHENIICS

Title: Measuring the θ_{13} mixing angle with the two Double Chooz detectors

Keywords: Double Chooz, neutrino oscillations, θ_{13} mixing angle, cosmogenics, ^{9}Li

Abstract: The Double Chooz experiment aims at accurately measuring the value of the θ_{13} leptonic mixing angle. To this intent, the experiment makes the most of two identical detectors – filled with gadolinium-loaded liquid scintillator – observing $\overline{\nu_e}$'s released by the two 4.25 GW_{th} nuclear reactors of the French Chooz power plant. The so-called "far detector" – located at an average distance of 1050 m from the two nuclear cores – has been taking data since April 2011. The "near detector" – at an average distance of 400 m from the cores – has monitored the reactor since December 2014. The θ_{13} mixing parameter leads to an energy dependent disappearance of $\overline{\nu_e}$'s as they propagate from the nuclear cores to the detection sites, which allows for a fit of the $\sin^2 2\theta_{13}$ value. By reason of correlations between the detectors and an iso-flux layout, the detection systematics and the $\overline{\nu_e}$ flux uncertainty impairing the θ_{13} measurement are dramatically suppressed. In consequence, the precision of the θ_{13} measurement is dominated by the uncertainty on the backgrounds and the relative normalisation of the $\overline{\nu_e}$-rates. The main background originates from the decay of βn-emitters – generated by μ-spallation – within the detector itself. The energy spectra of these cosmogenic isotopes have been simulated and complemented by a diligent error treatment. These predictions have been successfully compared to the corresponding data spectra, extracted by means of an active veto, whose performance has been studied at both sites. The rate of cosmogenic background remaining within the $\overline{\nu_e}$ candidates has also been assessed. Additionally, the normalisation of the $\overline{\nu_e}$ rates, bound to the number of target protons within each detector, has been evaluated. All this work was part of the first Double Chooz multi-detector results, yielding $\sin^2(2\theta_{13}) = 0.111 \pm 0.018$.

Titre : Mesure de l'angle de mélange θ_{13} avec les deux détecteurs de Double Chooz

Mots-clés : Double Chooz, oscillations de neutrinos, angle de mélange θ_{13}, cosmogéniques, ^{9}Li

Résumé : L'expérience Double Chooz a pour but de mesurer l'angle de mélange leptonique θ_{13} avec précision. Pour ce faire, l'expérience met à profit deux détecteurs identiques – remplis de liquide scintillant dopé au gadolinium – afin d'étudier les $\overline{\nu_e}$ produits par les deux réacteurs nucléaires de 4.25 GW_{th} de la centrale de Chooz. Le détecteur lointain – situé à une distance moyenne de 1050 m des cœurs – prend des données depuis Avril 2011. Le détecteur proche – à une distance moyenne de 400 m des cœurs – observe les réacteurs depuis Décembre 2014. Le paramètre de mélange θ_{13} conduit à une disparition d'$\overline{\nu_e}$ lorsque ceux-ci voyagent des cœurs jusqu'aux sites de détection ; la dépendance en énergie de ce déficit permet d'extraire la valeur de $\sin^2 2\theta_{13}$, par ajustement. Les systématiques de détection, ainsi que l'incertitude sur la prédiction du flux $\overline{\nu_e}$, sont formidablement réduites grâce aux corrélations entre les détecteurs et à la configuration iso-flux du site. Par conséquent, la précision sur la mesure de θ_{13} est dominée par l'incertitude sur les bruits de fond et par celle sur la normalisation relative des taux d'$\overline{\nu_e}$. Le bruit de fond principal provient de la désintégration d'émetteurs βn – produits par spallation des muons – dans le détecteur lui-même. Les spectres de ces isotopes cosmogéniques ont été simulés et complétés par un traitement d'erreur rigoureux. Ces prédictions sont en bon accord avec les données, elles-mêmes extraites à l'aide d'un veto actif, dont la performance a été étudiée pour les deux sites. Le taux d'évènements cosmogéniques restant parmi les candidats $\overline{\nu_e}$ a également été estimé. En outre, la normalisation relative des taux d'$\overline{\nu_e}$, liée aux nombres de protons dans les cibles de chaque détecteur, a été évaluée. Tous ces travaux se sont inscrits au sein des premières analyses Double Chooz à l'aide de deux détecteurs, aboutissant à $\sin^2(2\theta_{13}) = 0.111 \pm 0.018$.