Phenomenology of a Low-Energy Neutrino Factory and Related Experiments

LI, TRACEY,CHUIYEE

Phenomenology of a Low-Energy Neutrino Factory and Related Experiments

Tracey Li

A thesis submitted for the degree of
Doctor of Philosophy

Institute for Particle Physics Phenomenology
Department of Physics
Durham University
November 2010

Abstract

The discovery of neutrino oscillations is one of the most important in the recent history of particle physics, being the first evidence of physics beyond the Standard Model. We describe the theoretical framework of the neutrino oscillation model, motivate the necessity for a new generation of neutrino oscillation experiments and study the phenomenological factors which influence the design of these experiments.

We perform the first detailed study of a European super-beam setup using the CERN to Pyhäsalmi baseline of 2285 km, analysing the physics reach of this setup with a 100 kiloton liquid argon detector and comparing its performance to that of a 50 kiloton liquid scintillator detector and a 440 kiloton water Čerenkov detector. The liquid argon and liquid scintillator detectors are found to perform best, providing sensitivity to θ_{13}, δ and the mass hierarchy for $\sin^2 2\theta_{13} > 10^{-2}$.

A potential successor to super-beam experiments is a neutrino factory. We study a low-energy neutrino factory, a setup which has so far not been analysed in any detail, performing optimisation studies and an analysis of its sensitivity to oscillation parameters and non-standard matter interactions. We show that for $\sin^2 2\theta_{13} > 4 \times 10^{-3}$, a low-energy neutrino factory using a 20 kiloton totally active scintillating detector has 100% CP coverage for hierarchy sensitivity and θ_{13} discovery, and has greater sensitivity to CP violation than the high-energy neutrino factory. We consider the novel concept of including the 'platinum channels' in addition to the 'golden channels', showing that this is a powerful way of resolving the degeneracies between the oscillation and non-standard parameters. This enhances the sensitivity, such that the low-energy neutrino factory can put upper bounds $\gtrsim 10^{-2}$ on the non-standard interaction parameters $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau}$.

Declaration

This thesis is the result of my own work, except where explicit reference is made to the work of others. No part of this thesis has been submitted elsewhere for any other degree or qualification.

The work in Chapter 5 is part of the Large Apparatus for Grand Unification and Neutrino Astrophysics design study and will be submitted for publication at a later date.

The research in Chapter 6 was carried out in collaboration with Alan Bross, Malcolm Ellis, Steve Geer, Enrique Fernández Martínez, Olga Mena and Silvia Pascoli and is based partly on the published work,

'On the improvement of the low energy neutrino factory', Enrique Fernández Martínez, Tracey Li, Olga Mena and Silvia Pascoli, Phys. Rev. D81 (2010), 073010, arXiv:0911.3776.

This work was performed as part of the International Design Study for the Neutrino Factory. The work in Chapter 7 also forms part of this design study and will later be submitted for publication.

Tracey Li
November 2010

Acknowledgements

Above all, I have to thank my supervisor, Silvia Pascoli, for inspiration, motivation, guidance, encouragement and your invaluable criticism of my work. I am equally gratefully for the endless networking (and sight-seeing) opportunities with which you have provided me by posting me off to far-flung places all over the world! I feel extremely lucky to have been your student - thank you.

Thank you to all my collaborators, especially those from the International Design Study for the Neutrino Factory, for many interesting discussions and also for some entertaining times at conferences and meetings. Thank you also to Thomas Schwetz, who managed to finally resolve all my misunderstandings regarding statistical analyses with the contents of a single very enlightening email.

Then I have to thank my colleagues here at the IPPP for contributing to the friendly environment which has made my time here so enjoyable. Thanks especially to my friends and office-mates Aoife Bharucha, Eimear O'Callaghan and David Wilson, for the good company and chat, both work and non-work related, inside and outside the office.

And finally, thank you to a certain someone for some incredibly painstaking and patient proof-reading, and for putting up with me and all this physics.

This thesis was funded by a STFC studentship. Additional funds were provided by EuCARD (European Coordination for Accelerator Research and Development) which is co-funded by the European Commission within the Framework Programme 7 Capacities Specific Programme under Grant Agreement number 227579, and the European Community under the European Commission Framework Programme 7 Design Studies, EUROnu (Project Number 212372) and LAGUNA (Project Number 212343). The EC is not liable for any use that may be made of the information contained herein.

Contents

Chapter 1.

Introduction

Neutrino physics is a diverse and vibrant area of study, spanning the fields of particle physics, cosmology, astrophysics, nuclear physics and geophysics. Yet despite the fact that neutrinos are one of the most abundant particles in the universe, we know relatively little about them. Their tiny interaction cross-sections makes them one of the most elusive particles in spite of their vast natural abundance; to detect and observe them requires the design and construction of very specific experiments and detectors.

We are currently living in one of the most exciting eras of particle physics, thanks to the construction of the Large Hadron Collider (LHC). This is the world's most powerful particle accelerator and the largest physics experiment ever built. The LHC is designed to answer the question of how particles become massive, by creating and detecting the Higgs boson (if it exists) predicted by the Standard Model of particle physics. In addition, the LHC has the capability to confirm and improve upon precision measurements of the electro-weak parameters, enabling us to test the Standard Model more precisely than ever. But first and foremost, the LHC is a discovery machine - it is designed to probe physics at an energy scale higher than any other previous experiment, in a quest to search for physics beyond the Standard Model.

In view of this, it would be easy to forget that there are other ways in which we can search for new physics, and which may be complementary to the physics discovered by the LHC. Neutrinos form part of the Standard Model, and yet we do not have anywhere near the quantity of information about them as we do about the other particles in the model, nor do we have as precise measurements of their properties. In fact, the tau neutrino was the last Standard Model particle to be observed (apart from the Higgs boson), and this was not until as recently as 2000. So the questions we should ask are what neutrino properties are still unknown, which are we able to measure and how, and

what can be deduced from them? These are questions we shall address in this thesis by studying the phenomenon of neutrino oscillations.

The discovery that neutrinos oscillate was announced by the Super-Kamiokande collaboration in 1998. It is considered to be one of the most pivotal discoveries in the recent history of particle physics because the fact that neutrinos oscillate implies that they have non-zero masses, contrary to the predictions of the Standard Model. Thus, neutrino oscillations are the first evidence of physics beyond the Standard Model and that the current theory is incomplete. Although neutrino oscillations are a phenomenon which occurs at low energies relative to the energies of modern hadron colliders, they still give us vital hints and evidence about the existence of high-energy new physics. Therefore it seems wise to pursue this course of investigation and to obtain as much information about neutrino oscillations as possible since in that way we have a good chance of extracting several clues of new physics.

There are three key oscillation parameters which we still need to measure - the mixing angle, θ_{13}, which we know is close to zero, the Dirac phase, δ, which indicates that CP violation is present if not equal to zero or π, and the ordering of the neutrino mass states. Knowing the value of θ_{13} and δ will enable us to piece together the full neutrino mixing matrix. This should give us clues about the physics of the flavour sector of the Standard Model, one of the least understood aspects of the model. We know that the neutrino mixing matrix is very different from the quark mixing matrix, but we need to know why this is so and therefore precisely how it differs. Hence we are required to measure the neutrino mixing parameters to the same precision as the quark mixing parameters. The Dirac phase is also important because a discovery of CP violation in the neutrino sector is a pointer towards leptogenesis, but CP violation is only possible if θ_{13} is non-zero since the phase is only physical if all three mixing angles are non-zero. So it is also vital to gain more knowledge about the value of θ_{13} as soon as possible. In this way we can establish if we will be capable of measuring δ and the mass ordering, and also to decide which is the best experiment to build.

In view of this, it is important to consider all the potential experimental setups which are feasible and are likely to have a good physics reach in order to make an informed decision as to which is the best. The candidate experiments include super-beams and neutrino factories. Super-beams are more powerful versions of conventional neutrino beams and have the advantage of being a well-established and proven technology; therefore these are experiments which can be built some time in the relatively near future. Super-beams have been thoroughly studied in both the US and Japan, but relatively

little effort has been invested so far in considering setups within Europe. But in order to choose the best experimental setup, these options should also be seriously considered as we have in this thesis.

Super-beams can give us hints as to the value of θ_{13}, or if θ_{13} is suitably large, they may also be able to make a measurement of δ and the mass ordering. These measurements will be limited by systematic errors, backgrounds and statistics which prevent these experiments from making measurements if θ_{13} is very small. In this scenario, an experiment such as a β-beam or a neutrino factory will be necessary, both of which have been extensively studied in the literature. A neutrino factory is an experiment designed to produce a very pure and intense beam of neutrinos which propagate over distances of thousands of kilometres before being detected. The necessity for extremely long baselines stems from the need to exploit matter effects which will enable us to determine the neutrino mass ordering. The concept of a low-energy neutrino factory which is a version of a neutrino factory using much lower energy neutrinos and a shorter baseline than the standard neutrino factory setup, has up until now been side-lined in favour of the higher energy setup. Whilst it is true that the high-energy neutrino factory has a very impressive physics reach, it is important to ask whether this is the optimal setup for all scenarios. This is a question which we will answer by showing that a low-energy neutrino factory may sometimes perform better than the high-energy neutrino factory. Thus it is far from obvious that the standard high-energy neutrino factory is always the best option, and we must carefully consider and compare its performance to that of the low-energy neutrino factory.

In the rest of this chapter we shall put the main work of this thesis into context by describing how neutrinos fit into the Standard Model of particle physics and some of the questions which we still need to answer about them. At the end of the chapter we will give an outline of the remainder of the thesis.

1.1. Neutrinos: the story begins

The story of the neutrino (denoted by the Greek letter ν) began in 1930 when Wolfgang Pauli postulated the existence of a neutral particle to explain the 'missing energy' observed in the beta decay of radioactive ions. The process was thought to be a two-body decay, with a neutron decaying into a proton and electron, as only the initial and final nuclei and a free electron were observed in the process. This decay was therefore ex-

pected to produce a single monochromatic electron line, corresponding to the difference in mass between the neutron (or rather the the initial nucleus), and the proton (the final nucleus) and electron. Instead, a continuous spectrum of electron energies was observed. We know now that this 'missing energy' is that carried away by the neutrino, an idea which was first expressed by Wolfgang Pauli in December 1930 in a letter written to a gathering of physicists in Tübingen [1]. However, Pauli considered the work too immature to publish! But the Italian physicist Enrico Fermi took Pauli's idea and went on to develop the theory of weak interactions [2, 3], and it is he who named the particle 'little neutral one' in Italian. However, it was not until 1956 that the neutrino was actually discovered by Reines and Cowan [4], who detected the anti-electron neutrinos produced from a radioactive source, earning Reines the Nobel Prize in 1995.

In 1962 the muon neutrino was discovered at Brookhaven National Laboratory, in a collaboration led by Leon Lederman, Melvin Schwartz and Jack Steinberger [5] - they won the 1988 Nobel prize as a result. They created the first ever neutrino beam by taking protons from the Alternating Gradient Synchrotron, colliding them into a target to produce a shower of pions, and then allowing the pions to decay into muons and muon neutrinos. This is still the method by which neutrino beams are produced, as we will describe in more detail in Chapter 3. A detector was placed several metres away from the target, with a heavy iron wall in between the beam and the detector. The muons were unable to pass through the wall but the neutrinos could. The fact that muons were seen in the detector indicated that these neutrinos were producing muons and not electrons when they interacted, thereby indicating that they were different to the neutrinos detected by Reines and Cowan.

In 1975 the tau lepton was discovered [6], arousing suspicions that there should be a third neutrino, the tau neutrino. The existence of three light neutrino species was later confirmed by the LEP experiment [7] which measured Z^0 decays. By measuring the total width of the Z^0, and the width of all the visible decays into hadrons and charged leptons, the 'invisible' width could be attributed to decays into neutrinos. The data were consistent with the existence of three neutrinos. However, the tau neutrino was not observed until 2000 by the DONUT (Direct Observation of NU Tau) experiment [8] which detected the tau neutrinos produced from the decay of charmed particles. We shall explain in Section 3.3.4 why it took so long for this observation to be made!

1.2. Neutrinos in the Standard Model

The Standard Model of particle physics describes the fundamental particles of nature and their interactions which are encoded in the Standard Model Lagrangian. It is a quantum field theory based upon the concept of mathematical symmetries and groups. The strong, weak and electromagnetic interactions are completely specified by demanding local gauge invariance under the $SU(3)$, $SU(2)$ and $U(1)$ groups, respectively, with the weak and electromagnetic interactions being unified into a single force [9, 10, 11]. There is also a fourth interaction - gravity - which is yet to be consistently incorporated into the Standard Model. However, the strong, weak and electromagnetic interactions are well understood and theoretical predictions have been rigourously tested experimentally. The picture which has emerged is that the strong force is felt only by quarks, which are charged under the $SU(3)$ group ('colour charge'), and not by leptons. It is mediated by massless gluons, and is most relevant for sub-nuclear interactions. The 'weak' force, which will be discussed shortly, is mediated by the heavy $W^{\pm}$ and Z^0 gauge bosons, making it a short-range force. It is felt by both quarks and leptons. The electromagnetic force is mediated by the massless photon and thus has infinite range. Particles which have electric charge interact via this force.

Neutrinos are leptons and appear in the Standard Model as the $SU(2)$ partners of the charged leptons. They are the only electrically neutral matter particle in the Standard Model. Therefore they interact *only* via the weak force. Like all the other matter particles they are fermions (they have $\frac{1}{2}$-integer spin).

The matter particles of the Standard Model are shown in Table 1.1. For our purposes we have divided them simply into $SU(2)$ doublets and $SU(2)$ singlets as we shall only be concerned with weak interactions. The relevant weak interaction terms of the Lagrangian (for neutrinos only) are

$$\mathcal{L}_\nu = - \frac{ig}{\sqrt{2}}\bar{\nu}_L\gamma^\mu e_L W^+_\mu - \frac{ig}{\sqrt{2}}\bar{e}_L\gamma^\mu\nu_L W^-_\mu \tag{1.1a}$$

$$- \frac{ig}{2\cos\theta_W}\bar{\nu}_L\gamma^\mu\nu_L Z^0_\mu + \text{ H.C.}, \tag{1.1b}$$

(H.C. is the hermitian conjugate) where θ_W is the Weinberg mixing angle. Note that only left-handed (L) particles appear - the *weak interaction is maximally parity-violating* [12] which means that it couples only to left-handed particles and right-handed anti-particles. This is why, unlike all other particles, only the left-handed neutrinos and right-handed

anti-neutrinos are present in the Standard Model - the right-handed neutrinos and left-handed anti-neutrinos, if they exist, do not interact with any other Standard Model particle. The discovery that all neutrinos are left-handed was first made by the 1957 Goldhaber experiment [13].

The 'handedness' of the particle refers to its *chirality*, which is property labelling the γ^5 matrix eigenvalue of the particle field (see Appendix A). A general field, ψ, can be decomposed as

$$\begin{aligned} \psi &= \frac{1}{2}(1-\gamma^5)\psi + \frac{1}{2}(1+\gamma^5)\psi \\ &= \psi_L + \psi_R. \end{aligned} \tag{1.2}$$

The factors $\frac{1}{2}(1 \mp \gamma^5)$ are the *left-* and *right-handed projection operators*. We see, using the properties of the γ^5 matrix in Appendix A, that ψ_L has a γ^5 matrix eigenvalue of -1 and ψ_R has an eigenvalue of $+1$. The weak interaction distinguishes between these. A connection to directly observable properties can be made in the relativistic limit, when the chirality of a particle becomes equal to its helicity (the direction of spin with respect to the direction of motion) but in all other cases, this relation does not hold. A fundamental distinction between the two properties is that *helicity is a frame-dependent quantity* (for massive particles), but *chirality is Lorentz-invariant*; for massive particles, it is possible to Lorentz boost to a frame where the particle's velocity is reversed but the spin remains unchanged so that the helicity is reversed. However, a particle which is chirally left-handed in one frame and therefore undergoes weak interactions also interacts in *all* frames, irrespective of its helicity in those frames.

Eq. (1.1a) describes *charged-current* (CC) interactions, mediated by a charged $W^{\pm}$ boson:

$$\nu_\alpha + \ell^-_\beta \rightarrow \nu_\beta + \ell^-_\alpha, \qquad \bar{\nu}_\alpha + \ell^+_\beta \rightarrow \bar{\nu}_\beta + \ell^+_\alpha, \tag{1.3a}$$

$$\nu_\alpha + q_d \rightarrow \ell^-_\alpha + q_u, \qquad \bar{\nu}_\alpha + q_u \rightarrow \ell^+_\alpha + q_d, \tag{1.3b}$$

where q_u is any up-type quark (or down-type anti-quark), q_d is any down-type quark (or up-type anti-quark) and $\alpha = e,\ \mu,\ \tau$. An example is shown in Fig 1.1a.

Eq. (1.1b) describes *neutral-current* (NC) interactions, mediated by a neutral Z^0 boson:

$$\nu_\alpha + X \rightarrow \nu_\alpha + X, \qquad \bar{\nu}_\alpha + X \rightarrow \bar{\nu}_\alpha + X, \tag{1.4}$$

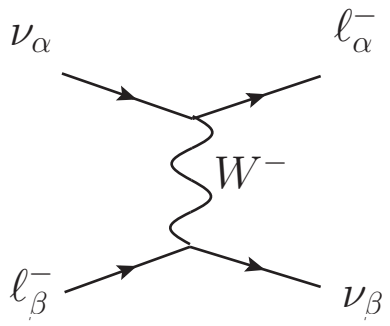

(a) A charged-current interaction.

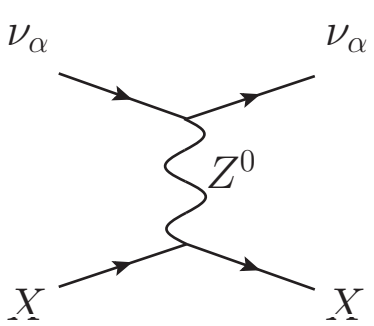

(b) A neutral-current interaction.

Figure 1.1.: Example of a charged-current and a neutral-current neutrino interaction.

an example of which is shown in Fig 1.1b, where X can be any lepton (or anti-lepton) or quark (or anti-quark). All of these interactions are extremely rare relative to the interaction rates of other Standard Model particles because of the tiny interaction cross-sections of neutrinos. For example, the total charged-current cross-section for muon neutrinos to interact with a nucleon at energies of a few GeV is $\sim 10^{-42}\mathrm{m}^2/$ GeV [14], which can be compared with the cross-section for a photon to interact with a proton ($\sim 10^{-32}\mathrm{m}^2/$ GeV) [14]. This is the reason why neutrinos were not discovered until several years after they were theoretically predicted, and for the development of giant kiloton scale detectors in neutrino experiments.

The idea of *currents* in particle physics is directly analagous to the idea in relativistic quantum mechanics where the 4-vector $J^\mu = (\rho, \mathbf{j})$ describes the probability and flux densities. In particle physics, in the case of *charged-currents*, J^μ is associated with the current and flux densities of electromagnetic charge. The charged-currents take the form

$$J^{\mu+}_{CC} = \bar{\nu}_\alpha \gamma^\mu \frac{1}{2}(1-\gamma^5)\ell^-_\alpha, \tag{1.5}$$

for interactions between neutrinos and charged leptons. Quarks also interact weakly and have a similar current. The hermitian conjugates of these currents are the *charge-lowering* currents,

$$(J^{\mu+}_{CC})^\dagger = J^{\mu-}_{CC} = \bar{\ell}^-_\alpha \gamma^\mu \frac{1}{2}(1-\gamma^5)\nu_\alpha, \tag{1.6}$$

for leptonic interactions, and similarly for the quark interactions.

For *neutral-currents* there is no transfer of electric charge. For neutrinos, the current is simply

$$J^{\mu}_{NC} = \bar{\nu}_L \gamma^{\mu} \nu_L. \tag{1.7}$$

Now the Lagrangian Eq. (1.1) can be written more compactly in terms of currents and the amplitudes for weak interaction processes expressed as

$$A_{CC} = \left(\frac{g}{\sqrt{2}} J^{\mu}_{CC}\right) \frac{1}{M_W^2 - q^2} \left(\frac{g}{\sqrt{2}} J^{\dagger}_{CC\mu}\right), \tag{1.8a}$$

$$A_{NC} = \left(\frac{g_{NC}}{\sqrt{2}} J^{\mu}_{NC}\right) \frac{1}{M_Z^2 - q^2} \left(\frac{g_{NC}}{\sqrt{2}} J^{\dagger}_{NC\mu}\right), \tag{1.8b}$$

where g is a fundamental coupling constant which describes the strength of the weak interactions, M_W is the mass of the $W^{\pm}$ boson (~ 80 GeV), $g_{NC} = \frac{g}{\cos\theta_W}$, $M_Z = \frac{M_W}{\cos\theta_W}$ is the mass of the Z^0 boson and q is the 4-momentum transferred by the interaction.

In a low-energy interaction where $q^2 \ll M_W^2$, which is applicable to all the cases we shall consider, we can describe Standard Model neutrino interactions with an *effective Lagrangian*,

$$\mathcal{L}_{\text{eff}} = -2\sqrt{2} G_F (J^{\mu}_{CC} J^{\dagger}_{CC\mu} + J^{\mu}_{NC} J^{\dagger}_{NC\mu}), \tag{1.9}$$

where $G_F = \frac{\sqrt{2} g^2}{8 M_W^2}$ is known as the Fermi coupling constant (the numerical constants are set by convention). In these circumstances, the interaction is essentially point-like and it can be seen that the weak interaction is 'weak' not because g is intrinsically weak, but because M_W is heavy. At high energies such that $q^2 \sim M_W^2$, the weak interaction becomes comparable in strength to the electromagnetic interaction.

1.3. Dirac or Majorana?

One of the most important questions in neutrino physics is whether neutrinos are *Dirac particles* like the other Standard Model fermions, or whether they are *Majorana particles* (named after the Italian physicist Ettore Majorana who first developed the theory [15]). Only electrically neutral particles can be Majorana, which is why neutrinos are the only matter candidate in the Standard Model.

	$SU(2)$ doublets	$SU(2)$ singlets
First generation	$\begin{pmatrix} \nu_e \\ e \end{pmatrix}_L \quad \begin{pmatrix} u \\ d \end{pmatrix}_L$	$e_R \quad u_R \quad d_R$
Second generation	$\begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L \quad \begin{pmatrix} c \\ s \end{pmatrix}_L$	$\mu_R \quad c_R \quad s_R$
Third generation	$\begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L \quad \begin{pmatrix} t \\ b \end{pmatrix}_L$	$\tau_R \quad t_R \quad b_R$

Table 1.1.: The matter content of the Standard Model.

A Dirac fermion is described by a 4-component Dirac spinor which can be written in terms of two 2-component Weyl spinors, ψ_- and ψ_+, if we choose to use the *Weyl representation*:

$$\psi = \begin{pmatrix} \psi_- \\ \psi_+ \end{pmatrix}. \tag{1.10}$$

Then left-handed and right-handed spinors can be written as

$$\psi_L = \begin{pmatrix} \psi_- \\ 0 \end{pmatrix}, \qquad \psi_R = \begin{pmatrix} 0 \\ \psi_+ \end{pmatrix}, \tag{1.11}$$

which have γ^5 eigenvalues of -1 and $+1$ and transform independently under the Lorentz group without mixing. ψ_L and ψ_R correspond to a particular fermion (a 2-component spinor describes a spin-$\frac{1}{2}$ particle) e.g. e_L^- and e_R^- . For every field there is a *conjugate field*, ψ^c, which corresponds to the anti-particle (e_L^+ and e_R^+ in our example). The fields are related via the *charge-conjugation* matrix, C (see Appendix B):

$$\psi^c = C\bar{\psi}^T. \tag{1.12}$$

In the general case, ψ and ψ^c correspond to *different* species, for instance an electron and a positron, which are distinct because they have opposite electric charges. For a

Dirac fermion, such as an electron, there are *four* distinct particles, e_R^-, e_R^+, e_L^- and e_L^+, all of which have the same mass but have different electric charges and chiralities.

Now consider a situation where ψ_+ is related to ψ_- (Appendix B) such that

$$\psi^c = \psi. \tag{1.13}$$

Then ψ is a *self-conjugate field*. This is the Majorana condition which defines a *Majorana spinor*. For a Majorana fermion, there are only *two* distinct particles of opposite chiralities which are otherwise identical. Note that this would imply that *lepton number is violated* (this quantity is otherwise conserved in the Standard Model) because in the Standard Model all leptons (including neutrinos) have lepton number $L = 1$ and all anti-leptons $L = -1$. We have observed ν_L and $\bar{\nu}_R$, but have not seen ν_R and $\bar{\nu}_L$. If neutrinos are Majorana then we could have processes where $\nu_{L,R} \leftrightarrow \bar{\nu}_{L,R}$ transitions occur which would mean that L is not conserved. Then processes such as *neutrinoless double-beta decay* could occur (see Ref. [16] for a review). This process is shown in Fig. 1.2b, together with standard double-beta decay in Fig. 1.2a. The standard double-beta decay occurs in some cases when the ordinary single-beta decay is energetically forbidden; then two neutrons in the nucleus are simultaneously converted to two protons (two down quarks are converted into two up quarks) which is accompanied by the emission of two electrons and two anti-electron neutrinos. The reverse process can also occur, in which case two electrons and two electron neutrinos are emitted. Now if lepton number is *not* conserved, then it is possible for a single (anti-) electron neutrino to be both emitted and absorbed as shown in Fig. 1.2b. This process can be searched for by experiments which measure the energies of the electrons emitted in double-beta decays - if a pair of neutrinos is also emitted in the standard decay then there will be some 'missing energy', but in a neutrinoless double-beta decay there is no missing energy so the energies of the electrons should precisely match the difference between the initial and final nuclei. Information about such experiments can be found in Refs. [17, 18, 19, 20, 21, 22, 23, 24].

Lepton number violation is of the utmost interest because it explains how an excess of leptons over anti-leptons in the early universe could have arisen. This could have been converted into an excess of baryons over anti-baryons (*baryogenesis via leptogenesis* [25, 26]) and thus provide a mechanism to explain the matter/ anti-matter asymmetry of the universe - one of the biggest puzzles of fundamental physics. Hence knowing whether neutrinos are Majorana particles or not is of considerable importance! We shall

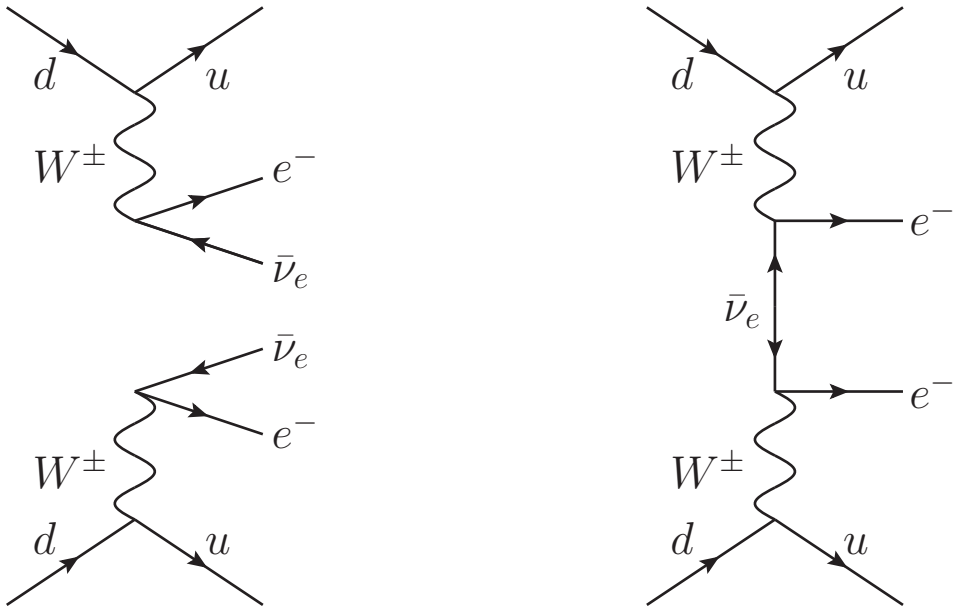

a) Double-beta decay.

b) Neutrinoless double-beta decay.

Figure 1.2.: Double-beta decay and neutrinoless double-beta decay.

mention the topic of leptogenesis again in Chapter 2 when we discuss CP violation in the neutrino sector.

1.4. Neutrino masses

The discussion of Dirac and Majorana particles leads naturally onto the topic of neutrino masses. A general overview of the topic can be found in e.g. Ref. [27]. A key pillar of the Standard Model, formulated by Glashow, Weinberg and Salam in the 1960's [9, 10, 11] is the mechanism of electro-weak symmetry breaking. This symmetry of the model is spontaneously broken, giving rise to a massive 'Nambu-Goldstone boson' [28, 29, 30] - the Higgs boson [31, 32, 33] - the discovery of which is one of the primary goals of the Large Hadron Collider (LHC). It is through the Higgs mechanism [31, 32, 33] that the particles of the Standard Model become massive; a model in which all the symmetries are preserved predicts only the existence of massless particles.

In a sense, 'mass' is a property which arises from the interaction between left- and right-handed fields; ψ_L and ψ_R remain entirely independent of one another except when they are coupled by a mass term. In the Higgs mechanism, the resulting mass is known

as a *Dirac mass* and arises from the coupling between the left- and right-handed components of the *same field*. The corresponding terms in the Lagrangian are

$$\mathcal{L}_D = \frac{1}{2} m_D (\bar{\psi}_L \psi_R + \bar{\psi}_R \psi_L + \bar{\psi}^c_L \psi^c_R + \bar{\psi}^c_R \psi^c_L), \tag{1.14}$$

which can be written as

$$\mathcal{L}_D = \frac{1}{2} m_D (\bar{\psi}\psi + \bar{\psi}^c \psi^c), \tag{1.15}$$

because $\bar{\psi}_L \psi_L = \bar{\psi}_R \psi_R = 0$. By writing the Lagrangian like this, Eq. (1.15) also makes it explicit that the mass eigenstates are $\psi = \psi_R + \psi_L$ and $\psi^c = \psi^c_R + \psi^c_L$.

However, we could also couple the left- and right-handed components of *conjugate fields* e.g. $(\psi)_L$ and $(\psi^c)_R$ to produce *Majorana mass terms*. If we note that

$$(\psi^c)_R = \frac{1}{2}(1+\gamma^5)\psi^c = \left(\frac{1}{2}(1-\gamma^5)\psi\right)^c = (\psi_L)^c = \psi^c_L, \tag{1.16}$$

which means that the right-handed component of the conjugate field is the same as the conjugate of the left-handed component of the field (and vice-versa), then we can also think of Majorana mass terms as coupling left-handed (right-handed) fields to the conjugate of the left-handed (right-handed) fields:

$$\mathcal{L}_M = m_{M1}(\bar{\psi}^c_L \psi_L + \bar{\psi}_L \psi^c_L) + m_{M2}(\bar{\psi}^c_R \psi_R + \bar{\psi}_R \psi^c_R). \tag{1.17}$$

So these terms mix neutrinos and anti-neutrinos and therefore violate lepton number, as was mentioned at the end of Section 1.3. We can also write the Majorana mass terms more compactly as

$$\mathcal{L}_M = m_{M1} \bar{\nu}\nu + m_{M2} \bar{N} N, \tag{1.18}$$

where $\nu = \psi_L + \psi^c_L$ and $N = \psi_R + \psi^c_R$ are the Majorana mass eigenstates *which are self-conjugate*, and m_{M1} and m_{M2} are symmetric mass matrices (see e.g. Ref. [27]). Now if we identify ν as being the Standard Model neutrino pair (ν_L, $\bar{\nu}_R$), then N is a second Majorana pair of neutrinos not predicted by the Standard Model. These may be sterile (non-interacting) and could have much larger masses than ν - these heavy right-handed neutrinos appear frequently in beyond the Standard Model theories.

The most general term mass term which can be written is thus the sum of Eq. (1.14) and (1.18) which can be written in matrix form as

$$\mathcal{L}_{mass} = \begin{pmatrix} \bar{\nu} & \bar{N} \end{pmatrix} \begin{pmatrix} m_{M1} & \frac{1}{2}m_D \\ \frac{1}{2}m_D & m_{M2} \end{pmatrix} \begin{pmatrix} \nu \\ N \end{pmatrix}. \tag{1.19}$$

The mass matrix can be diagonalised to yield the mass eigenvalues

$$M_{\pm} = \frac{1}{2}\left(m_{M1} + m_{M2} \pm \sqrt{(m_{M1} - m_{M2})^2 + m_D^2} \right), \tag{1.20}$$

with corresponding eigenstates

$$\phi_+ = \nu\cos\theta + N\sin\theta, \qquad \phi_- = \nu\sin\theta - N\cos\theta, \tag{1.21}$$

where

$$\tan 2\theta = \frac{m_D}{m_{M1} - m_{M2}}. \tag{1.22}$$

The mass eigenstates are again self-conjugate. So the most general mass term for a four-component fermion field (1.19) describes *two pairs of Majorana particles* with *different masses*. When $m_{M1} = m_{M2} = 0$ then $\theta = \frac{\pi}{4}$ and the mass eigenvalues are $M_{\pm} = \pm\frac{1}{2}m_D$. To obtain the physical mass eigenstates (which have positive masses) we need to perform a chiral transformation of the states (1.21) (see e.g. Section 13.2 of Ref. [34]) to obtain the solutions

$$\chi_+ = \frac{1}{\sqrt{2}}(\psi_L + \psi_L^c + \psi_R + \psi_R^c) = \phi_+, \tag{1.23a}$$

$$\chi_- = \frac{1}{\sqrt{2}}(-\psi_L + \psi_L^c - \psi_R + \psi_R^c), \tag{1.23b}$$

which both have masses of $\frac{1}{2}m_D$ and are both self-conjugate - although $\chi_+^c = \chi_+$ and $\chi_-^c = -\chi_-$, the minus sign can be absorbed into the matrix C since the matrix $-C$ has the same properties (Eq. (B.2)) as C (this indicates that χ^+ and χ^- have opposite CP parities). Since these are degenerate eigenstates, any linear combination of them is also an eigenstate. If we choose the solutions to be $(\chi_+ \pm \chi_-)$ then we recover our Dirac solutions from Eq. (1.15), $\psi = \psi_L + \psi_R$ and $\psi^c = \psi_L^c + \psi_R^c$. Thus we find that *a pair of Majorana states with degenerate masses and opposite CP parities is equivalent to a Dirac state of the same mass.*

The reason for not including right-handed neutrinos and left-handed anti-neutrinos in the Standard Model initially was that they simply had not been observed, and it was assumed that they were Dirac particles. As neutrinos were, for a long time, thought to be massless, this was also consistent with the fact that there should only be one handedness of neutrino - the presence of both chiralities would imply that neutrinos should have a non-zero Dirac mass. Today, we know that neutrinos have non-zero masses (because of neutrino oscillations, as explained in the next chapter) but that they are many orders of magnitude smaller than the mass of the next lightest particle, the electron (the current bound is $m_{\nu_e} \lesssim 2$ eV [35]), although we have yet to measure the *exact* neutrino mass. Thus it is questionable whether neutrinos acquire mass through the Higgs mechanism as this requires an explanation for why their Higgs interactions are so tiny relative to that of the other particles. On the other hand, generating neutrino masses through Majorana couplings requires the presence of terms which violate lepton number, a quantity which is otherwise conserved in the Standard Model.

A popular idea is that neutrinos become massive via a 'see-saw' mechanism [36, 37], which requires the presence of a heavy sterile (N_R, $\bar{N}_L$) pair to generate both Dirac and Majorana mass terms, as in Eq. (1.19), with $m_{M1} = 0$ and $m_D \ll m_{M2}$. In this way, the smallness of the Standard Model neutrino masses is a direct consequence of the heaviness of the sterile neutrinos which are expected to have masses of the order of the grand unification scale. The development of these see-saw models, and others, to explain tiny neutrino masses is currently a very active field of theoretical research (see Ref. [38] for a review) as is the experimental quest to make a measurement of the absolute neutrino mass [35, 39, 40, 41].

The question of whether neutrinos are Majorana or Dirac particles is of fundamental importance but may unfortunately remain unanswered for many years to come! A discovery that neutrinos are Dirac particles would imply that they have Dirac masses, the smallness of which would require an explanation. In addition we would know that lepton number, like baryon number, is conserved in the Standard Model. If, however, we find that neutrinos are Majorana particles, then the see-saw mechanism becomes a likely explanation of the smallness of neutrino masses and additionally we will have found that lepton number is not conserved in the Standard Model. This would make leptogenesis seem like a plausible scenario.

1.5. Outline of the thesis

In the remainder of this thesis we will focus our attention on the phenomena of neutrino oscillations from the theoretical, phenomenological and experimental perspectives. In Chapter 2 we will describe the theory and formalism behind neutrino oscillations. We will discuss neutrino oscillations in vacuum and then oscillations in matter, describing how 'matter effects' alter the oscillations relative to the vacuum case. We also summarise our current knowledge of the neutrino oscillation parameters. Following this we will introduce non-standard interactions which are lepton flavour-changing processes that may arise from physics beyond the Standard Model, explaining why neutrino oscillation experiments are ideal tools for hunting for such processes.

The experiments which are designed to detect these phenomena are discussed in Chapter 3. To begin with, we explain how the design of a neutrino oscillation experiment is linked to the theory described in Chapter 2. The components of past, current and future neutrino oscillation experiments are then described. We discuss their benefits and limitations and introduce the super-beam and neutrino factory which are potential future experiments that we study in this thesis. We then describe the techniques used to simulate and analyse these experiments.

The phenomenology at these experiments is the topic of Chapter 4. The golden channel, the $\nu_e \to \nu_\mu$ channel, which is the channel that a neutrino factory is primarily designed to observe, is discussed in detail; super-beams measure the $\nu_\mu \to \nu_e$ channel which is phenomenologically similar. We show how the unknown mixing parameters, θ_{13}, δ and the mass hierarchy, can be extracted from the spectrum of the golden channel and how the value of θ_{13} and matter effects influence these measurements. We briefly mention the other channels which may be accessible and finish with a discussion about degeneracies, explaining the problem and some possible solutions.

Super-beams, which are more powerful versions of conventional neutrino beams, are studied in Chapter 5. We optimise and assess the potential of a super-beam within Europe as part of the Large Apparatus for Grand Unification and Neutrino Astrophysics (LAGUNA) design study. There are seven possible European baselines and three detector options - liquid argon, liquid scintillator and water Čerenkov - being considered by this design study; up until now, only the shortest baseline with one of the detector options has been studied in detail. In the first half of the chapter we perform studies of the liquid scintillator detector which is so far the least well-developed of the three

detector technologies, in order to ascertain the potential and limiting factors of this detector. In the second half we perform the first detailed study of the potential of the longest LAGUNA baseline - 2285 km from CERN to Pyhäsalmi - which is the baseline that will have the best physics reach, considering all three of the detector options. These are setups which can be stepping stones towards the physics goals of an experiment such as the neutrino factory, the topic of Chapters 6 and 7.

Neutrino factories have, in general, been thoroughly analysed in the literature, although a low-energy version of the setup has not yet been extensively studied as most effort has been concentrated on the standard high-energy neutrino factory. In Chapter 6 we introduce, optimise and analyse the performance of a low-energy neutrino factory as part of the International Design Study for the Neutrino Factory. The low-energy version of the neutrino factory uses much lower energy neutrinos than the standard high-energy neutrino factory and therefore a shorter baseline of around 1000 km rather than several thousands of kilometres is required. For the first time, we consider the possibility of adding the platinum channels, $\nu_\mu \to \nu_e$ and $\bar{\nu}_\mu \to \bar{\nu}_e$, to the setup, as well as considering the effect of altering factors such as the energy threshold of the detector. We show that although the high-energy neutrino factory is optimal for very small values of θ_{13} (such that $\sin^2 2\theta_{13} \lesssim 10^{-3}$), the low-energy setup actually performs better for larger values of θ_{13}, in terms of sensitivity to CP violation. As well as the discovery potential of the experiment, we also consider the precision with which it can measure the oscillation parameters.

We then go on to analyse the ability of the low-energy neutrino factory to detect non-standard interactions in Chapter 7. These are neutrino interactions not predicted by the Standard Model and so their detection would be a direct signal of new physics. It turns out that long-baseline neutrino oscillation experiments are ideal tools for searching for these signals and so we should, if possible, incorporate their detection into the optimisation and analysis of the low-energy neutrino factory. We show how the combination of the golden and platinum channels at a low-energy neutrino factory is an extremely powerful way of resolving the degeneracies between the oscillation and non-standard parameters, thereby increasing the sensitivity of the setup.

Finally, in Chapter 8 we draw our conclusions.

Chapter 2.

Neutrino oscillations and non-standard interactions

In this chapter we will introduce the main topic of this thesis - neutrino oscillations - the phenomenon of neutrinos changing flavour as they propagate through space. The discovery that neutrinos oscillate is one of the most important in the recent history of particle physics as oscillations only occur if *neutrino masses are non-zero*, as will be explained. So neutrino oscillations were the first evidence that neutrinos had non-zero masses and thus that the Standard Model, with its prediction of massless neutrinos, was incomplete.

We will derive the probabilities for neutrino oscillations both in vacuum and in matter, describing the Mikheyev-Smirnov-Wolfenstein effect. We will then go on to introduce *non-standard interactions* which are neutrino interactions other than those predicted by the Standard Model, and which may thus provide clues of new physics. We will describe the formalism and explain why future long-baseline oscillation experiments are powerful tools for searching for these interactions.

2.1. The birth of neutrino oscillation physics

In 1957 the Italian physicist Bruno Pontecorvo first postulated that neutrinos might oscillate [42]. Some years later, in the 1960's, uncertainty about neutrino properties began to arise due to the 'solar neutrino problem': at that time, understanding of solar processes had reached a level that enabled confident predictions to be made about the expected flux of electron neutrinos arriving at the earth from the sun. However,

experiments measuring this flux observed only about a third of that predicted; the first experiment to detect this deficit was Ray Davis' Homestake Experiment [43], for which Davis won the Nobel Prize in 2002.

Suspicions that the problem was due to neutrino oscillations, as predicted by Pontecorvo, culminated in 1998 when the Super-Kamiokande collaboration [44] released results which showed that their detector was seeing a deficit of atmospheric muon neutrinos which was consistent with a model where muon neutrinos oscillated into tau neutrinos. The proof that neutrinos were oscillating between flavours, and not just disappearing, came in 2001 when the SNO experiment [45] became the first experiment to be sensitive to the *total* flux of neutrinos arriving from the sun. In this way they showed that the total flux, of all three flavours combined, was consistent with that predicted by solar models, thereby implying that neutrinos born as electron neutrinos in the sun were oscillating into different flavours en-route to the earth.

Pontecorvo postulated the existence of neutrino oscillations following the observation of particle mixing in the neutral meson sector where K^0 mesons were seen to oscillate into $\bar{K}^0$. Similarly, B^0 and D^0 mesons also oscillate into their anti-particles. These oscillations occur because the flavour (interaction) eigenstates of these particles do not coincide with the mass (propagating) eigenstates. The two sets of states can then be linked by a mixing matrix; in the quark sector this is the Cabbibo-Kobayashi-Maskawa (CKM) matrix [46, 47] whilst in the neutrino sector the mixing matrix is called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [48]. The mass states are conventionally denoted as ν_1, ν_2 and ν_3 whilst the flavour states are ν_e, ν_μ and ν_τ:

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

In the absence of sterile neutrinos, the three Standard Model neutrinos are the only ones to exist and the PMNS matrix is a unitary 3×3 matrix. It can be parameterised by three mixing angles and three complex phases (see Appendix C) and written as the

product of three 2×2 rotation matrices and two phase matrices,

$$
\begin{aligned}
U &= U_{23} U_{13,\delta} U_{12} V_{\zeta} \\
&= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13} e^{\pm i\delta} \\ 0 & 1 & 0 \\ -s_{13} e^{\mp i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{i\zeta_1/2} & 0 & 0 \\ 0 & e^{i\zeta_2/2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \\
&= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13} e^{\pm i\delta} \\ -s_{12}c_{23} - c_{12}s_{13}s_{23} e^{\mp i\delta} & c_{12}c_{23} - s_{12}s_{13}s_{23} e^{\mp i\delta} & c_{13}s_{23} \\ s_{12}s_{23} - c_{12}s_{13}c_{23} e^{\mp i\delta} & -c_{12}s_{23} - s_{12}s_{13}c_{23} e^{\mp i\delta} & c_{13}c_{23} \end{pmatrix} \begin{pmatrix} e^{i\zeta_1/2} & 0 & 0 \\ 0 & e^{i\zeta_2/2} & 0 \\ 0 & 0 & 1 \end{pmatrix},
\end{aligned}
$$

where $s_{jk} = \sin\theta_{jk}$ and $c_{jk} = \cos\theta_{jk}$ and the upper signs on the phase δ refers to neutrinos, and the lower sign to anti-neutrinos. The angles θ_{jk} (or rather, the physically measurable combinations of $\cos\theta_{jk}$ and $\sin\theta_{jk}$) describe the mixing between the j and k sectors. The phase δ is called the Dirac phase and, if not equal to zero or π, and all three mixing angles are non-zero, indicates that CP violation occurs in the neutrino sector (CP violation is only possible if there is mixing between *more than two families* [47]). The phases ζ_1 and ζ_2 are the Majorana phases which are only physical if neutrinos are Majorana particles, and can be detected only via experiments where the Majorana character of neutrinos manifests itself - in lepton number violating processes such as neutrinoless double-beta decay. Since neutrino oscillations violate *lepton flavour* but not *lepton number*, the Majorana character of neutrinos will not be exhibited in oscillations and so the Majorana phases are irrelevant to neutrino oscillation physics.

2.2. Neutrino oscillations in vacuum

The simplest case of neutrino oscillations are those that occur in vacuum. Here we will outline the calculation of the probability for these oscillations. A rigourous derivation requires the consideration and treatment of quantum mechanical wave packets, as described in Ref. [49] and briefly discussed in Appendix D, but for simplicity we shall use the commonly-used *equal momentum approximation* which, although not being a generally true assumption, produces the correct result for our situations.

We must first note that the neutrino *mass states*, $|\nu_1\rangle$, $|\nu_2\rangle$ and $|\nu_3\rangle$, are the physically propagating states which evolve according to the time-dependent Schrödinger equation,

$$-i\frac{\partial}{\partial t}|\nu_j\rangle = \hat{H}|\nu_j\rangle, \tag{2.1}$$

where $j = 1, 2, 3$. The solutions are plane waves of the form $e^{-iE_j t}|\nu_j\rangle$.

We wish to calculate the probability that a neutrino which starts as flavour α oscillates into a neutrino of flavour β after propagating a distance L - neutrino oscillation experiments measure only the distance, and not the time, over which a neutrino propagates. Since the mass states and not the flavour states are those which propagate, we need to write both the initial and final states in terms of the mass states. The initial state is

$$|\psi(x=0)\rangle = |\nu_\alpha\rangle = \sum_{j=1}^{3} U^*_{\alpha j}|\nu_j\rangle, \tag{2.2}$$

which after propagating a distance $x = L$ in time t becomes

$$|\psi(x=L)\rangle = \sum_{j=1}^{3} U^*_{\alpha j} e^{-iE_j t}|\nu_j\rangle, \tag{2.3}$$

so that the amplitude for $\nu_\alpha \to \nu_\beta$ is

$$\langle\nu_\beta|\psi(L)\rangle = \sum_{j=1}^{3} U^*_{\alpha j} U_{\beta j} e^{-iE_j t}, \tag{2.4}$$

where we have used the fact that the mass states are orthonormal: $\langle\nu_k|\nu_j\rangle = \delta_{jk}$. Now we use the equal-momentum assumption,

$$E_j = \sqrt{p_j^2 + m_j^2} \simeq p\left(1 + \frac{m_j^2}{2p^2}\right), \tag{2.5}$$

where p is the common momentum of all the mass states, and the relativistic approximation $t \simeq L$, so that

$$\langle\nu_\beta|\psi(L)\rangle = \sum_{j=1}^{3} U^*_{\alpha j} U_{\beta j} e^{-iL\left(p + \frac{m_j^2}{2p}\right)}. \tag{2.6}$$

(Here we see that one of the problems with the equal-momentum assumption is that the neutrino states only appear to propagate in *time* and not space - the spatial dependence is only introduced by using the $t \simeq L$ approximation). The probability is the square of the amplitude, and to a good approximation we can set $p^{-1} \simeq E^{-1}$. The probability is then

$$P_{\nu_\alpha \to \nu_\beta}(L, E) = \left| \sum_{j=1}^{3} U^*_{\alpha j} U_{\beta j} e^{-iL\left(p + \frac{m_j^2}{2E}\right)} \right|^2 , \tag{2.7}$$

which can be simplified (see Appendix D) by using the unitarity condition $\sum_{j=1}^{3} U^*_{\alpha j} U_{\beta j} = \delta_{\alpha\beta} = |\sum_{j=1}^{3} U^*_{\alpha j} U_{\beta j}|^2$ and the half-angle relation $1 - \cos\theta = 2\sin^2\frac{\theta}{2}$ to obtain the final expression,

$$P_{\nu_\alpha \to \nu_\beta}(L, E) = \delta_{\alpha\beta} \quad - \quad 4 \sum_{j>k} \mathrm{Re}[U^*_{\alpha j} U_{\beta j} U_{\alpha k} U^*_{\beta k}] \sin^2\left(\frac{\Delta m^2_{jk} L}{4E}\right) \tag{2.8a}$$

$$+ \quad 2 \sum_{j>k} \mathrm{Im}[U^*_{\alpha j} U_{\beta j} U_{\alpha k} U^*_{\beta k}] \sin\left(\frac{\Delta m^2_{jk} L}{2E}\right). \tag{2.8b}$$

Thus, from Eq. (2.8) we find that neutrino oscillations depend upon the parameters of the PMNS matrix (the three angles θ_{13}, θ_{23} and θ_{12}, and the CP phase δ), on the mass-squared differences between neutrino mass states ($\Delta m^2_{31} = m^2_3 - m^2_1$, $\Delta m^2_{32} = m^2_3 - m^2_2$ and $\Delta m^2_{21} = m^2_2 - m^2_1$), on the distance between the neutrino source and the point at which it is detected (the 'baseline' L), and on the neutrino energy, E. Neutrino oscillation experiments measure oscillation probabilities as a function of E; usually L is fixed but there is a range of neutrino energies. The fact that oscillations have been observed which depend on two mass-squared splittings is evidence that *at least two neutrino masses are non-zero.*

Eq. (2.8a) is the CP conserving part of the probability, whilst Eq. (2.8b) (the imaginary part) denotes the CP violating part of the probability. The implication of CP violation by non-real numbers stems from the fact that the CP operator transforms $i \to -i$. In neutrino oscillations, the only possible complex quantity is $e^{i\delta}$ and so if δ is equal to 0 or π then CP is conserved, whereas for any other value, CP is violated (if all three mixing angles are non-zero). It is also possible that the Majorana phases mentioned in Section 2.1 may violate CP, but as we mentioned there, these phases are not relevant to neutrino oscillations.

Shortly, in Section 2.3, we will describe in detail our current knowledge of the neutrino mixing parameters. For now, we will only mention that we know that $\Delta m^2_{32} \simeq \Delta m^2_{31}$ and that the magnitude of this quantity is ~ 30 times larger than Δm^2_{21} so that Eq. (2.8) can sometimes be simplified by making the approximation $\Delta m^2_{32} = \Delta m^2_{31}$. Then there are only two, not three, different oscillation frequencies. Further simplifications can be made depending on the particular value of $\frac{L}{E}$ being considered:

- If $\frac{L}{E}$ is very small, then all of the $\frac{\Delta m^2_{jk} L}{4E}$ are much less than unity so $\sin\left(\frac{\Delta m^2_{jk} L}{4E}\right) \simeq 0$. Effectively, there are no oscillations.
- If $\frac{L}{E}$ is very large so all of the $\frac{\Delta m^2_{jk} L}{4E}$ are much greater than unity, then all the oscillation modes are very rapid and, within the energy and spatial ranges visible in a detector, cannot be resolved. The CP conserving terms, $\sin^2\left(\frac{\Delta m^2_{jk} L}{4E}\right)$, average to $\frac{1}{2}$ whilst the CP violating terms, $\sin\left(\frac{\Delta m^2_{jk} L}{2E}\right)$, average to zero.
- If $\frac{\Delta m^2_{21} L}{4E} \sim O(1)$ then $\frac{|\Delta m^2_{31}| L}{4E} \gg 1$, and the rapid Δm^2_{31} oscillations are modulated by the slower Δm^2_{21} oscillations. Only these slower oscillations can be seen. An example is the KamLAND experiment [50] which measured the $\bar{\nu}_e \to \bar{\nu}_e$ channel using $\bar{\nu}_e$'s from nuclear reactors with energies of a few MeV. The baseline was 180 km which means that $\frac{|\Delta m^2_{31}| L}{4E} \simeq 100$ whereas $\frac{\Delta m^2_{21} L}{4E} \simeq 4$. The probability in this case is then (see Appendix E for a detailed derivation)

$$P_{\bar{\nu}_e \to \bar{\nu}_e} \simeq 1 - \sin^2 2\theta_{12} \sin^2\left(\frac{\Delta m^2_{21} L}{4E}\right), \tag{2.9}$$

  which is plotted in Fig. 2.1a as a function of $\frac{L}{E}$ (using the value of θ_{12} given in Section 2.3).
- If $\frac{|\Delta m^2_{31}| L}{4E} \sim O(1)$, implying that $\frac{\Delta m^2_{21} L}{4E} \ll 1$, then effectively there are only Δm^2_{31} oscillations. An example is the MINOS experiment [51] which measures the $\nu_\mu \to \nu_\mu$ channel using neutrinos with $E \sim 1$ GeV and a baseline of $L = 735$ km. In this case, $\frac{|\Delta m^2_{31}| L}{4E} \simeq 2$ whereas $\frac{\Delta m^2_{21} L}{4E} \simeq 0.07$ so that the probability is (Appendix E)

$$P_{\nu_\mu \to \nu_\mu} \simeq 1 - \sin^2 2\theta_{23} \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right), \tag{2.10}$$

  which is plotted in Fig. 2.1b (using the value of θ_{23} given in Section 2.3).

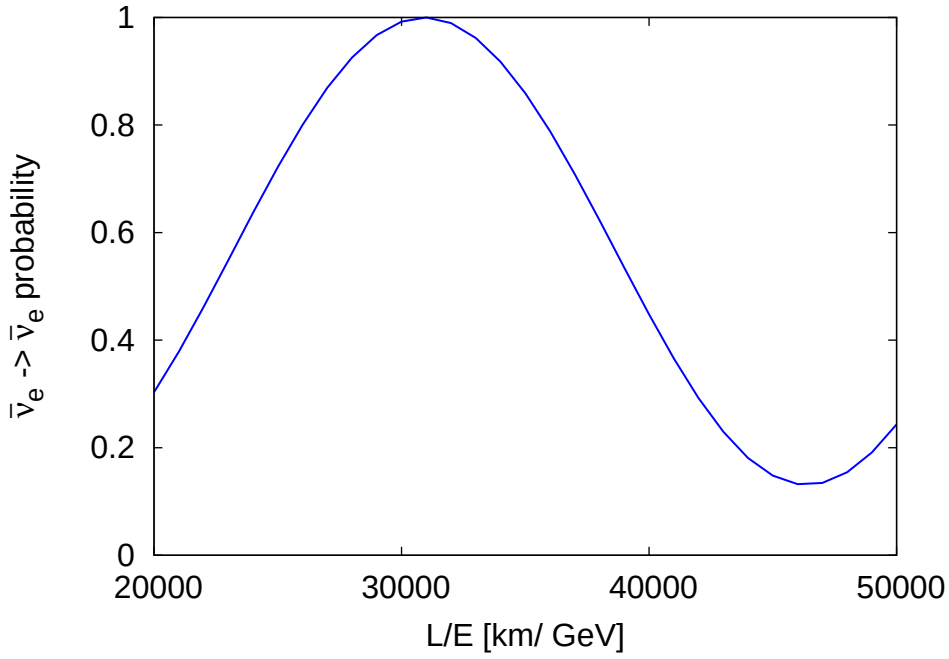

(a) $\frac{\Delta m^2_{21} L}{4E} \sim 1$ and $\frac{\Delta m^2_{31} L}{4E} \gg 1$.

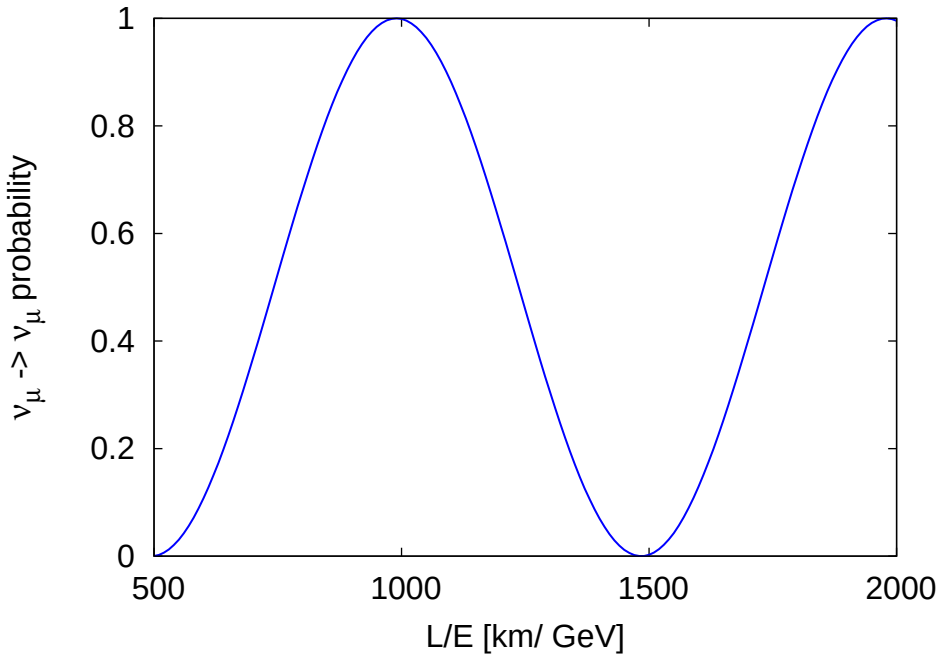

(b) $\frac{\Delta m^2_{21} L}{4E} \gg 1$ and $\frac{\Delta m^2_{31} L}{4E} \sim 1$.

Figure 2.1.: Oscillation probability for the a) KamLAND experiment ($P_{\bar{\nu}_e \to \bar{\nu}_e}$ with $\frac{L}{E} \sim 10^5$ km/ GeV and $\theta_{12} = 34°$) and b) MINOS experiment ($P_{\nu_\mu \to \nu_\mu}$ with $\frac{L}{E} \sim 10^3$ km/ GeV and $\theta_{23} = 42°$).

Thus we see that it is necessary to have different experiments with different baselines and energies in order to be able to measure *all* of the oscillation parameters. In the following section we review the results of past and current experiments to show which of the parameters have already been measured and which of them still need to be measured. Following this, we shall discuss neutrino oscillations in matter.

2.3. Current knowledge of neutrino mixing parameters

The results of past and current neutrino oscillation experiments tell us the following about the neutrino oscillation parameters: one angle, θ_{13}, is approximately zero, whilst another angle, θ_{23}, is approximately maximal ($\frac{\pi}{4}$). The third angle, θ_{12}, has an intermediate value. As yet there is no infomation about the phase δ except for a recent tentative bound from Super-Kamiokande [52]. We know that one mass-squared splitting, Δm^2_{21}, is positive (i.e. $m_2 > m_1$) and that it is very much smaller than the magnitude of the other two splittings. Thus, as was already mentioned in Section 2.2, it is sometimes possible to make the approximation $\Delta m^2_{31} \simeq \Delta m^2_{32}$, which we shall use throughout this chapter. Although the *magnitude* of this quantity has been accurately measured, we do not know its *sign* i.e. whether it is positive or negative, which corresponds to us not knowing the *mass hierarchy* (explained shortly). The most recent values for these parameters at the time of writing, obtained by a global fit to all existing oscillation data, can be found in Ref. [53] which gives the following 3σ ranges:

- $\Delta m^2_{21} = 7.59^{+0.61}_{-0.69} \times 10^{-5}$ eV^2
- $\Delta m^2_{31} = +2.51^{+0.39}_{-0.36} \times 10^{-3}$ eV^2 (normal hierarchy), $-2.40^{+0.37}_{-0.39} \times 10^{-3}$ eV^2 (inverted)
- $\theta_{12} = 34.4^{+3.2}_{-2.9}$ degrees
- $\theta_{23} = 42.3^{+11.4}_{-7.1}$ degrees
- $\theta_{13} \leq 13.2$ degrees
- $\delta \subset [0, 360]$ degrees.

This information can be depicted diagrammatically, as in Fig. 2.2.

Of the mass-squared splittings, Δm^2_{21} is called the 'solar mass-squared difference' as this is the parameter which modulates the frequency of solar ν_e oscillations, whereas Δm^2_{31} is called the 'atmospheric mass-squared difference' as this parameter controls the frequency of atmospheric ν_μ oscillations. The fact that we do not know whether the atmospheric mass-squared difference is positive or negative corresponds to us not knowing the *ordering* of the neutrino mass states. 'Normal ordering' corresponds to a scheme in which ν_1 is the lightest species - this would be consistent with the mass hierarchy of the other particles in the Standard Model where the first generation is the lightest (although the first generation neutrino is the ν_e, which does not have a definite mass, it is predominantly composed of ν_1 and so ν_1 can be associated with the first generation). The opposite, 'inverted ordering', corresponds to a scheme in which ν_3 is the lightest species. The phrases 'normal hierarchy' and 'inverted hierarchy' are often used interchangeably with 'normal ordering' and 'inverted ordering', as they are in this thesis. Strictly speaking, 'ordering' refers only to the ordering of masses as just described, whereas 'hierarchy', of which there are three types, also takes into account the absolute scale of the masses as well as the ordering:

- Normal hierarchical spectrum: $m_1 \ll m_2 \ll m_3$
- Inverted hierarchical spectrum: $m_3 \ll m_1 \simeq m_2$
- Quasi-degenerate spectrum: $m_1 \simeq m_2 \simeq m_3 \gg 0$.

Neutrino oscillation experiments therefore provide information about the *mass ordering*, but not the *mass hierarchy*, as they are sensitive to mass-squared *differences* but not *absolute* masses.

The values of the mixing angles tell us the following (see Fig. 2.2): ν_1 consists predominantly of ν_e with an approximately equal admixture of ν_μ and ν_τ. ν_2 is an approximately equal mix of all three flavours, and ν_3 consists mostly (and possibly entirely) of an exactly or nearly equal mix of ν_μ and ν_τ. It is the value of the angle θ_{13} which determines the size of the ν_e component in ν_3, and θ_{23} the relative amount of ν_μ and ν_τ. The extreme case of this mixing scheme is that described by 'tri-bimaximal mixing' [54] when θ_{13} is exactly zero and θ_{23} exactly maximal, so-called because ν_2 is 'trimaximally' mixed and ν_3 is 'bimaximally' mixed. All data so far are consistent with this scheme, but only precision measurements in the future will be able to tell us whether or not this scheme is adhered to exactly in nature.

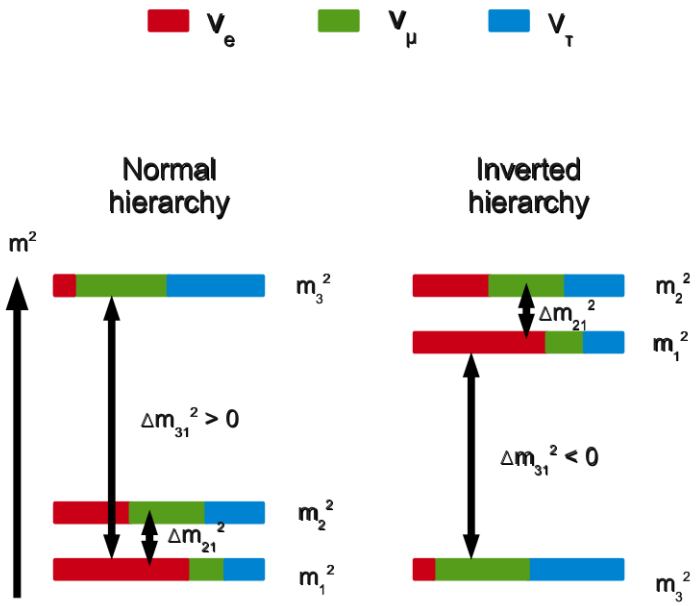

Figure 2.2.: A schematic diagram of the mass spectrum and flavour content of the neutrino mass eigenstates: each mass eigenstate is a linear superposition of the flavour eigenstates (shown by different colours), and the mass-squared differences but not the absolute masses are known. Details are given in the text below.

In recent years there have been several tantalising (if not necessarily statistically significant!) hints that θ_{13} may be non-zero [55]. In summary, different groups obtain best fits values for $\sin^2\theta_{13}$ of between 0.01 and 0.02, with a significance of between 1σ and 2σ, when performing global fits. The data which contributes to favouring a non-zero θ_{13} includes the 2008 KamLAND data [56], the 2009 MINOS data [57] (although this hint has now disappeared following the acquisition of more data since then), and the 2010 Super-Kamiokande data [58]. A non-zero θ_{13} is of extreme interest phenomenologically, because a non-zero value of θ_{13} is required in order for CP violation to be possible in the neutrino sector - CP violation can only occur if there is mixing between more that two families, and hence all three mixing angles must be non-zero. Additionally, a non-zero value of θ_{13} is required in order for us to be able to determine the mass hierarchy from neutrino oscillations. Theoretically, the value of θ_{13} is also of vital importance - even if θ_{13} is precisely zero at some high energy (grand unification) scale which is indicative of an exact flavour symmetry, then at low energies it is expected to deviate from zero. If it is exactly zero even at low energy scales then this will be a hint of some powerful conservation law at work. Either way, the precise value of θ_{13} (and also of θ_{23}) will help to put tight constraints on several models. The hunt for θ_{13} continues with the experiments mentioned in Section 3.4. A summary of the current status and prospects for measuring θ_{13} can be found in Ref. [55].

The final parameter is the CP phase, δ, about which there is currently no information apart from a preliminary result from Super-Kamiokande [52] which obtains $141^\circ < \delta < 297^\circ$ at 68% confidence. Ascertaining whether or not CP violation exists in the neutrino sector is one of the most important goals of future neutrino oscillation experiments, because the discovery of CP violation in the neutrino sector at the low energies of neutrino oscillation experiments could indicate the presence of the high-energy CP violation required for leptogenesis to have occurred in the early universe [25, 26]. Previously it was thought that the existence of low-energy leptonic CP violation did not guarantee the existence of the high-energy CP violation required for leptogenesis, but work in Refs. [59, 60] showed that this was not necessarily the case in all circumstances. Hence a discovery of CP violation in neutrino oscillations would be a hint in favour of leptogenesis!

2.4. Neutrino oscillations in matter

Having described neutrino oscillations in vacuum, we will now move onto examining neutrino oscillations in matter. This is relevant for the long-baseline terrestrial experiments such as those studied in this thesis, where neutrinos propagate through the earth between source and detector. The phenomenon was first discussed by Lincoln Wolfenstein when he realised that *there could be neutrino oscillations in matter even if all neutrinos are massless* [61] due to the interactions of neutrinos with matter. Neutrinos propagating through the earth interact with electrons and up and down quarks in the earth. *All* flavours of neutrino will interact via neutral-current interactions, but *only* ν_e *and* $\bar{\nu}_e$ will interact with the electrons in charged-current interactions. Thus the neutral-current interactions introduce a flavour-symmetric term into the propagation Hamiltonian, whereas the charged-currents distinguish between flavours. The effect of these interactions is to add an effective potential onto the neutrino energy; the neutral-current interactions will contribute equally to all neutrino flavours so for our purposes we only need to calculate the potential induced by the charged-current interactions; in other words we need to calculate the potential induced on an electron neutrino by the electrons in the earth. We start by looking at the terms in the Lagrangian relevant to electron neutrinos, including both the kinetic and 'potential terms' (and ignoring mass terms):

$$\mathcal{L}_{\nu_e} = \bar{\nu}_{eL} i\gamma^\mu \partial_\mu \nu_{eL} - 2\sqrt{2} G_F (\bar{\nu}_{eL}\gamma^\mu \nu_{eL})(\bar{e}_L \gamma_\mu e_L), \tag{2.11}$$

where G_F is the Fermi coupling constant, as in Eq. (1.9), and we have performed a Fierz rearrangement on the charged-current term in that equation to obtain our 'potential term'. The equation of motion for ν_e (using the Euler-Lagrange equations) is then

$$(i\gamma^\mu \partial_\mu + 2\sqrt{2} G_F \gamma^\mu \langle \bar{e}_L \gamma_\mu e_L \rangle)\nu_{eL} = 0. \tag{2.12}$$

Now we can assume that the background through which the neutrinos are propagating is static (this assumption is obviously not Lorentz invariant!) and therefore that all the electrons are at rest. In this case, only the time-like component of the electron potential,

$\langle \bar{e}_L \gamma_\mu e_L \rangle$, is non-zero. Then

$$\begin{aligned} \langle \bar{e}_L \gamma_\mu e_L \rangle &= \delta_{\mu 0} \langle \bar{e}_L \gamma_\mu e_L \rangle \\ &= \delta_{\mu 0} \langle e_L^\dagger e_L \rangle \\ &= \frac{1}{2} \delta_{\mu 0} n_e, \end{aligned} \tag{2.13}$$

where n_e is the number density of electrons in the matter through which the neutrinos are propagating and we require a factor of $\frac{1}{2}$ because neutrinos only interact with the left-handed electrons. Our equation of motion becomes

$$(i\gamma^\mu \partial_\mu + \sqrt{2} G_F n_e \gamma^0) \nu_{eL} = 0 \tag{2.14}$$

which has plane-wave solutions with

$$E = \mathbf{p} + \sqrt{2} G_F n_e. \tag{2.15}$$

The positive sign becomes a negative sign for anti-neutrinos. The 'potential', $\sqrt{2} G_F n_e$, is commonly denoted by A. The calculation of the potential arising from the neutral-current interactions is similar.

In vacuum, in the mass basis, the Hamiltonian is simply a diagonal matrix, $\hat{H}^{\text{vac}}_{\text{mass}} = \text{diag}(E_1, E_2, E_3)$, which, using the approximation $E_j - E_k \approx \frac{\Delta m^2_{jk}}{2E}$ as in Section 2.2, and defining $\frac{\Delta m^2_{jk}}{2E} = \Delta_{jk}$, becomes

$$\hat{H}^{\text{vac}}_{\text{mass}} = \begin{pmatrix} E_1 & 0 & 0 \\ 0 & E_1 & 0 \\ 0 & 0 & E_1 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta_{21} & 0 \\ 0 & 0 & \Delta_{31} \end{pmatrix}. \tag{2.16}$$

The first term is a diagonal constant which will simply contribute a global phase factor to the solutions and so can be neglected. The matter interaction terms can be converted into the mass basis using the PMNS matrix and added to obtain the Hamiltonian in matter (the neutral-current terms are a diagonal constant which again can be neglected):

$$\hat{H}^{\text{mat}}_{\text{mass}} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta_{21} & 0 \\ 0 & 0 & \Delta_{31} \end{pmatrix} \pm U \begin{pmatrix} A & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} U^\dagger. \tag{2.17}$$

where the upper (+) sign is for neutrinos, the lower (−) sign is for anti-neutrinos and $A = \sqrt{2}G_F n_e$. The propagation of neutrinos in matter is described by the Schrödinger equation, similar to Eq. (2.1), and diagonalisation of Eq. (2.17) will yield the eigenvectors and eigenvalues of the matter Hamiltonian, corresponding to the propagating matter states and their energies. In the case that n_e is constant along the propagation path so that A is constant, Eq. (2.1) can be solved exactly. In the case that A varies along the baseline, such as for long baselines traversing multiple layers of the earth, Eq. (2.1) can only be solved numerically, except for some particular circumstances which we will discuss shortly.

In the limit that $\Delta_{21} = 0$ we have a scenario where there is no $1 \leftrightarrow 2$ mixing - in this limit these two states have degenerate masses and so no oscillations occur between them. Therefore there are only $1 \leftrightarrow 3$ and $2 \leftrightarrow 3$ oscillations. So we can effectively set θ_{12} to zero and have a simplified mixing scenario where there are only three parameters: Δm^2_{31}, θ_{13} and θ_{23} (the Dirac phase, δ, is not physical unless there are three non-zero mixing angles, as explained in Section 2.1).

We will discuss this two-family scenario first, then afterwards consider the situation where we make the approximation that A is constant along the propagation path, which is valid for the experimental setups which we consider in the subsequent chapters. There we perform the matrix diagonalisation perturbatively, rather than exactly, in order to produce a simpler analytic form. We will follow the method used in Appendix C of Ref. [62], which treats the quantity Δ_{21} as the perturbation. This is a valid approximation since the relevant scales involved are Δ_{21}, Δ_{31} and A, with $\frac{\Delta_{21}}{\Delta_{31}} \sim 0.03$ and $\frac{A}{\Delta_{31}} \sim 0.3$ for the energies and value of A which we consider (see Chapter 4), and so it is Δ_{21} and not A which we treat as the perturbation.

2.4.1. A two-family approximation: $\nu_e \leftrightarrow \nu_\tau$ oscillations in matter

For simplicity, let us first consider a scenario where there are only $1 \leftrightarrow 3$ oscillations so that the effective mixing matrix is $U_{\text{eff}} = U_{13}$ (which is effectively a 2×2 rotation matrix) and $\nu = (\nu_e, \nu_\tau)$. The mass eigenstates are then $|\nu_1\rangle = \cos\theta_{13}|\nu_e\rangle - \sin\theta_{13}|\nu_\tau\rangle$ and $|\nu_2\rangle = \sin\theta_{13}|\nu_e\rangle + \cos\theta_{13}|\nu_\tau\rangle$ (and similar for $\bar{\nu}$). The Hamiltonian is given by

Eq. (2.17) with $\Delta_{21} = 0$ and $U = U_{13}$:

$$\hat{H}^{\text{mat}} = \begin{pmatrix} 0 & 0 \\ 0 & \Delta_{31} \end{pmatrix} \pm \begin{pmatrix} A\cos^2\theta_{13} & A\sin\theta_{13}\cos\theta_{13} \\ A\sin\theta_{13}\cos\theta_{13} & A\sin^2\theta_{13} \end{pmatrix}, \tag{2.18}$$

The eigenvectors are

$$|\nu_1^{\text{mat}}\rangle = \cos\theta_{m\mp}|\nu_e\rangle - \sin\theta_{m\mp}|\nu_\tau\rangle, \tag{2.19a}$$
$$|\nu_2^{\text{mat}}\rangle = \sin\theta_{m\mp}|\nu_e\rangle + \cos\theta_{m\mp}|\nu_\tau\rangle, \tag{2.19b}$$

where the mixing angle in matter, $\theta_{m\mp}$, is related to the mixing angle in vacuum, θ_{13}, by [61]

$$\tan 2\theta_{m\mp} = \frac{\Delta_{31}\sin 2\theta_{13}}{\Delta_{31}\cos 2\theta_{13} \mp A}. \tag{2.20}$$

The vacuum limit corresponds to $A = 0$. For $A \neq 0$ and if we consider neutrinos, if the mass hierarchy is normal so $\Delta_{31} > 0$ and there is a potential such that $A = \Delta_{31}\cos 2\theta_{13}$ then θ_m is maximal. So there is maximal mixing between ν_1 and ν_3 in matter even though θ_{13} is very small and there is very little mixing in vacuum - this is the *resonance condition* when neutrino oscillations are maximally enhanced in matter (and anti-neutrino oscillations are suppressed). If the mass hierarchy is inverted, then there is a resonance for anti-neutrinos and the neutrino oscillations are suppressed. This is called the Mikheyev-Smirnov-Wolfenstein (MSW) effect [61, 63]. Finally, if $A \gg \Delta_{31}$ then oscillations are suppressed for both neutrinos and anti-neutrinos.

Now let us consider the case of a non-constant matter potential, such that A is a function of the distance travelled by the neutrino, L. From Eq. (2.20) we see that if A is varying, then the evolution of θ_m is given by

$$\frac{d\theta_m}{dL} = \pm\frac{1}{2}\frac{\Delta_{31}\sin 2\theta_{13}}{(\Delta_{31}\cos 2\theta_{13} - A)^2 + (\Delta_{31}\sin 2\theta_{13})^2}\frac{dA}{dL}, \tag{2.21}$$

where we have used the relations

$$\cos 2\theta_m = \frac{\Delta_{31}\cos\theta_{13} - A}{\sqrt{(\Delta_{31}\cos 2\theta_{13} - A)^2 + (\Delta_{31}\sin 2\theta_{13})^2}}, \tag{2.22a}$$
$$\sin 2\theta_m = \frac{\Delta_{31}\sin 2\theta_{13}}{\sqrt{(\Delta_{31}\cos 2\theta_{13} - A)^2 + (\Delta_{31}\sin 2\theta_{13})^2}}. \tag{2.22b}$$

In general this situation cannot be studied analytically, because the eigenstates $\nu_1^{\rm mat}$ and $\nu_2^{\rm mat}$ are then also functions of L and *mix with one another at each point of propagation* (see, for example, Ref. [27]). However, in an *adiabatic process*, the potential varies sufficiently slowly to enable the propagating neutrino state to adapt to the change, such that the initial state is an eigenstate of the initial Hamiltonian, the final state is an eigenstate of the final Hamiltonian, and there is negligible mixing between eigenstates at each point. This occurs if $\frac{d\theta_m}{dL} \ll 1$ which corresponds to

$$\frac{dA}{dL} \ll \frac{\Delta_{31}\sin 2\theta_{13}}{(\Delta_{31}\cos 2\theta_{13} - A)^2 + (\Delta_{31}\sin 2\theta_{13})^2}. \tag{2.23}$$

Coincidentally, it is just such a situation which caused the 'solar neutrino problem' that first aroused suspicions of neutrino oscillations! Solar neutrinos are created deep within the sun's interior, and are born as ν_e's. The electron density in the sun varies approximately exponentially, decreasing away from the centre of the sun until it drops to zero in empty space. The density at the centre is very high, so that A is very large. This means that $\theta_{m_-}(L=0) \simeq \frac{\pi}{2}$, so that in our two-family approximation the matter state $\nu_2^{\rm mat}(L=0)$ is essentially the same as ν_e (in fact, it is $1 \leftrightarrow 2$ oscillations which are relevant for solar neutrinos and not $1 \leftrightarrow 3$ oscillations, but we will use our two-family approximation to explain the principle). As the neutrino travels outwards towards the sun's surface, because the electron density varies slowly, the state is always an eigenstate of the Hamiltonian at that point, at all points along the path. However, the flavour content of this state evolves as A and hence θ_{m_-} changes. In other words, the flavour content of the eigenstate at a distance L from the centre, $\nu_2^{\rm mat}(L)$, is a function of L (through A). So the neutrino which was born as a ν_e at the centre of the sun *does* evolve as it propagates outwards, *but it evolves such that it is always the 'second' state of the Hamiltonian at that particular point.* Therefore when it reaches the sun's surface, the neutrino is now a pure $\nu_2^{\rm vac}$ state - the eigenstate in vacuum - and this is the state which propagates through space to the earth. So it is ν_2's and not ν_e's which arrive at the earth; therefore the probability that we detect a solar neutrino as a ν_e is $|\langle\nu_e|\nu_2\rangle|^2$ and *not* unity as originally expected!

2.4.2. Three-family oscillations in matter: a perturbative approach

Returning to a full three-family mixing scenario, our task is to find a matrix which transforms from the flavour basis into the matter basis; in this way we can derive oscillation probabilities in matter using an identical method to the derivation in vacuum (Section 2.2) but with an altered mixing matrix. We sketch out the solution here, with more detail being given in Appendix F. In the limit $\Delta_{21} = 0$ we can still set $\theta_{12} = 0$, as in the 2-family approximation, although we now have to consider $2 \leftrightarrow 3$ oscillations in addition to $1 \leftrightarrow 3$ oscillations. The effective mixing matrix is then $U_{\text{eff}} = U_{23}U_{13}$ although only θ_{13} is modified by A, as in the two-family case (because $U_{23}\,[\text{diag}(A,0,0)]\,U_{23}^{\dagger} = [\text{diag}(A,0,0)]$). We start with the Hamiltonian in the *flavour* basis:

$$\hat{H}_{\text{fl}}^{\text{mat}} = U_{\text{eff}}^{\dagger} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \Delta_{31} \end{pmatrix} U_{\text{eff}} \pm \begin{pmatrix} A & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{2.24}$$

which is diagonalised to [64]

$$\hat{H}_{\text{fl}}^{\text{mat}} = (\bar{U}_{\mp}^{(0)})^{\dagger} \begin{pmatrix} \frac{\Delta_{31} \pm A - B_{\mp}}{2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \frac{\Delta_{31} \pm A + B_{\mp}}{2} \end{pmatrix} \bar{U}_{\mp}^{(0)}, \tag{2.25}$$

where

$$\bar{U}_{\mp}^{(0)} = U_{23}(\theta_{23})U_{13}(\theta_{m\mp}), \tag{2.26a}$$

$$B_{\mp} = \sqrt{(\Delta_{31}\cos 2\theta_{13} \mp A)^2 + (\Delta_{31}\sin 2\theta_{13})^2}. \tag{2.26b}$$

$\theta_{m\mp}$ is the same as in Eq. (2.20) and again, the upper signs refer to neutrinos and the lower signs to anti-neutrinos. We use the superscript (0) to denote that this is the mixing matrix in the limit $\Delta_{21} = 0$. The rows of $\bar{U}_{\mp}^{(0)}$ are the eigenvectors, with corresponding eigenvalues given by $\lambda_1 = \frac{\Delta_{31} \pm A - B_{\mp}}{2}$, $\lambda_2 = 0$ and $\lambda_3 = \frac{\Delta_{31} \pm A + B_{\mp}}{2}$. Now we can introduce Δ_{21} as a perturbation, which needs to be written in the matter basis. In the mass basis, the perturbation is $\hat{H}_{\text{mass}}^{(1)} = [\text{diag}(0, \Delta_{21}, 0)]$ and in the flavour basis,

$$\hat{H}_{\text{fl}}^{(1)} = U^{\dagger}\hat{H}_{\text{mass}}^{(1)}U, \tag{2.27}$$

where U is now the full PMNS matrix involving all three mixing angles. The matrix $\bar{U}^{(0)}_{\mp}$ transforms between the flavour and matter bases, such that

$$\hat{H}^{(1)}_{\rm fl} = (\bar{U}^{(0)}_{\mp})^{\dagger} \hat{H}^{(1)}_{\rm mat} \bar{U}^{(0)}_{\mp}, \tag{2.28}$$

and so $\hat{H}^{(1)}_{\rm mat}$ is given by

$$\begin{aligned} \hat{H}^{(1)}_{\rm mat} &= \bar{U}^{(0)}_{\mp} \hat{H}^{(1)}_{\rm fl} (\bar{U}^{(0)}_{\mp})^{\dagger} \\ &= \bar{U}^{(0)}_{\mp} U \hat{H}^{(1)}_{\rm mass} U^{\dagger} (\bar{U}^{(0)}_{\mp})^{\dagger}. \end{aligned} \tag{2.29}$$

Using quantum mechanical perturbation theory, the first order corrections to the eigenvalues are given by the diagonal entries of the perturbation:

$$\lambda^{(1)}_{j} = \hat{H}^{(1)}_{jj}, \tag{2.30}$$

and the first order corrections to the eigenvectors are given by

$$v^{(1)}_{j} = \sum_{k \neq j} \frac{\hat{H}^{(1)}_{jk}}{\lambda^{(0)}_{j} - \lambda^{(0)}_{k}} v^{(0)}_{j}. \tag{2.31}$$

The explicit expressions can be found in Appendix C of Ref. [62] and are also given in Appendix F. Thus we now have a matrix, $\bar{U}_{\mp} = \bar{U}^{(0)}_{\mp} + \bar{U}^{(1)}_{\mp}$ (where $\bar{U}^{(1)}_{\mp}$ is the matrix composed of the vectors $v^{(1)}_{j}$), to first order in Δ_{21}, which transforms from the flavour basis into the matter basis. The derivation of neutrino oscillation probabilities in matter is now identical to that in vacuum, except for the substitution of U with $\bar{U}_{\mp}$.

2.4.3. Matter effects - a brief summary

The presence of matter alters the propagation Hamiltonian for neutrinos relative to the vacuum case. Therefore, the neutrino eigenstates in matter must necessarily be different to the eigenstates in vacuum (which are the mass eigenstates). In vacuum, we know that the flavour states are a linear combination of the mass states with coefficients determined by the mixing angles (and complex phase) of the PMNS matrix; of course we can invert this statement and say that the mass states are a linear combination of the flavour states. In matter, we can think of these relationships as remaining the same, with the only alteration being that the values of the mixing angles are altered. For instance, we

can say that in vacuum the coefficient of ν_τ in ν_1 is $\sin\theta$ where $\theta = \theta_{13}$ (ignoring the CP phase for now); in matter this relation still holds but with $\theta = \theta_m$.

A simplified but intuitive explanation is that interactions with matter cause neutrinos to be refracted, like light, and charged-current interactions mean that ν_e's are refracted more than the other neutrino flavours. This makes ν_e's effectively more heavy relative to the vacuum case. The effect is similar for $\bar{\nu}_e$'s, but since anti-neutrinos have opposite 'weak charge' to neutrinos, the weak interactions have the opposite effect and so $\bar{\nu}_e$'s become effectively lighter. Referring back to Fig. 2.2 we see that the majority of the ν_e state is contained within ν_1 and ν_2, and so making ν_e heavier corresponds to increasing the masses of ν_1 and ν_2 as shown in Fig. 2.3 (although the vacuum mass states are *not* the physically propagating states in matter, the explanation still holds). Thus for a normal hierarchy, Δm^2_{31} is effectively *decreased* which means that *neutrino oscillations are enhanced in matter in the case of a normal hierarchy*, as the energy gap between neutrino states is decreased. In the case of an inverted hierarchy, Δm^2_{31} is effectively *increased* and so *neutrino oscillations are suppressed in matter in the case of an inverted hierarchy*. The opposite is true for anti-neutrinos.

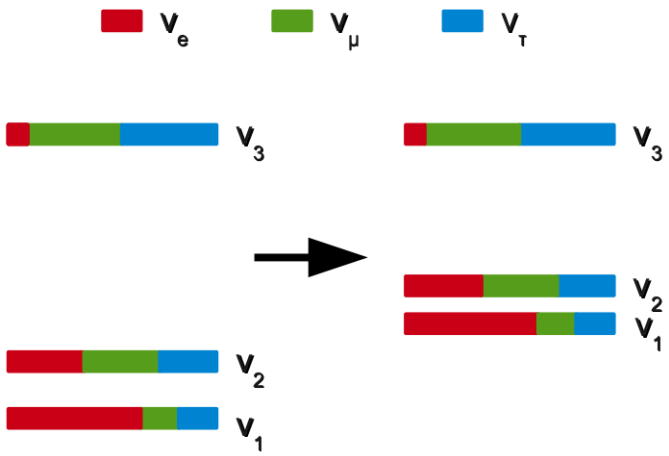

Figure 2.3.: In the case of a normal hierarchy, the propagating vacuum states, ν_1 and ν_2, become heavier in matter, decreasing the mass gap between these eigenstates and ν_3. Thus oscillations are enhanced, whereas the opposite is true for an inverted hierarchy. For anti-neutrinos, oscillations are suppressed for a normal hierarchy and enhanced for an inverted hierarchy.

The fact that the oscillatory behaviour of neutrinos in matter is dependent upon the mass hierarchy is exactly the technique which will be used by future long-baseline oscillation experiments to determine the hierarchy.

2.5. Non-standard interactions

As neutrino interactions are incredibly rare, the chance of a neutrino interacting more than once in the detector is, essentially, zero so that a neutrino will never leave a direct track in a detector. The only way of inferring that a neutrino interaction has occurred is by observing the signatures of the other particles produced in the interaction. The Standard Model tells us that the flavour of a neutrino cannot be determined from a *neutral-current* interaction, but that it can be determined by the flavour of the charged lepton which is produced in a *charged-current* interaction (see Section 1.2). In other words, if, for instance, an e^- is detected then it is *assumed* that the incoming neutrino was a ν_e because that is what the Standard Model predicts. But what if this assumption is wrong?

In addition to oscillation measurements, it is also possible for neutrino oscillation experiments to search for signals from new physics which exhibit themselves as neutrino non-standard interactions. The term 'non-standard interactions' (NSI's) is used in this context to refer to any neutrino interactions which are not described by the Standard Model, but arise from some new physics mechanism. They will therefore be mediated by some heavy new particle, which means that in a low-energy experiment the interactions can be effectively described as a point-like interaction, exactly analogous to the discussion of the weak interaction in Section 1.2. The effective low-energy vertex looks simply like a four-point vertex, which will be suppressed $\sim \frac{1}{M^2_{\text{NSI}}}$. Conventionally, this interaction strength is parameterised by the dimensionless parameter $\varepsilon_{\alpha\beta}$, which is the ratio of the strength of the NSI interaction, $\nu_\alpha \to \nu_\beta$, relative to the strength of the weak interaction, G_F. A rough estimate of the magnitude of $\varepsilon_{\alpha\beta}$ can be made using

$$\varepsilon_{\alpha\beta} \sim \frac{M_W^2}{M^2_{\text{NSI}}} \sim 10^{-2}, \tag{2.32}$$

with $M_{\text{NSI}} \sim 1$ TeV as in Refs. [65, 66].

NSI's can arise at each of the three stages of an oscillation experiment: the point of neutrino *production* ('source effects'), during *propagation* ('matter effects'), or at the

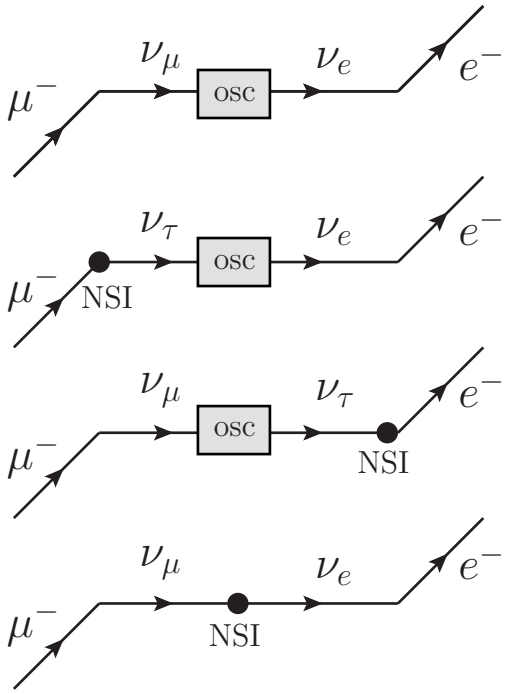

Figure 2.4.: Oscillation and non-standard processes which could lead to a $\nu_\mu \to \nu_e$ event.

point of *detection* ('detector effects'), as discussed in e.g. Refs. [67, 68, 69]. In the case of source effects, the question we wish to answer is whether, for example, the neutrinos produced from the decays of μ^- are *always* ν_μ and $\bar{\nu}_e$, or if sometimes there are some non-standard processes leading to other neutrino flavours being produced? Similarly, for detector effects, we would like to know whether the detection of an e^- is always an indicator that the incoming neutrino was a ν_e. These possibilities are shown in Fig. 2.4 where we show how the apparent observation of a $\nu_\mu \to \nu_e$ oscillation could actually be instigated by a NSI at any of the stages. This has sparked recent interest in the idea of using *near detectors* at oscillation experiments to search for source NSI's [70], by looking for 'zero-point' interactions - flavour changes which occur even when $L = 0$. The most powerful channel for such discoveries turns out to be the $\nu_\mu \to \nu_\tau$ channel, now dubbed the 'discovery channel' [71]. Lepton flavour violating interactions such as $\mu \to e\gamma$ will also be able to provide powerful constraints on these source NSI's [72, 73].

The NSI's in which we will be interested in this thesis are those which arise in the process of neutrino *propagation* - they are a *non-standard matter effect*. Essentially, they describe a neutrino flavour change instigated by some interaction between the original incoming neutrino and other matter particles - either electrons or up or down quarks in the case of terrestrial experiments [74]. They are defined by the addition of the following

terms to the Lagrangian:

$$L_{\rm NSI} = G_F \varepsilon_{\alpha\beta}^{f\mp} [\bar{f}\gamma^\mu(1 \mp \gamma^5)f][\bar{\nu}_\alpha \gamma_\mu (1-\gamma^5)\nu_\beta], \tag{2.33}$$

where $f = e$, u, d and α, $\beta = e$, μ, τ. These terms have a similar form to the Standard Model charged-current terms defined in Section 1.2 except that we also allow for the possibility that the charged leptons may be right- or left-handed. The quantum-mechanical Hamiltonian (Eq. (2.24)) is altered by the addition of the term

$$\hat{H}^{\rm fl} = U^\dagger \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta_{21} & 0 \\ 0 & 0 & \Delta_{31} \end{pmatrix} U \pm A \begin{pmatrix} 1+\varepsilon_{ee} & \varepsilon_{e\mu} & \varepsilon_{e\tau} \\ \varepsilon_{e\mu}^* & \varepsilon_{\mu\mu} & \varepsilon_{\mu\tau} \\ \varepsilon_{e\tau}^* & \varepsilon_{\mu\tau}^* & \varepsilon_{\tau\tau} \end{pmatrix}, \tag{2.34}$$

where $\varepsilon_{\alpha\beta} = \varepsilon_{\beta\alpha}$ because $\hat{H}$ must be Hermitian. This then implies that the diagonal entries must be real, whereas the non-diagonal entries can be complex and thus contain additional, possible CP violating, phases $\phi_{\alpha\beta}$: $\varepsilon_{\alpha\beta} = |\varepsilon_{\alpha\beta}|e^{i\phi_{\alpha\beta}}$ [65]. To solve the Schrödinger equation we can treat the NSI matrix as a perturbation, considering each element in turn, and follow a similar method to Section 2.4.2 to derive relatively simple analytic expressions for oscillation probabilities to leading order in the NSI's. More detail is given in Appendix G.

The recent (statistically insignificant) hints of apparent CPT violation in the MINOS experiment [75], which observed the atmospheric mass-squared splitting to be *different* (larger) for $\bar{\nu}_\mu \to \bar{\nu}_\mu$ oscillations than $\nu_\mu \to \nu_\mu$ oscillations could be explained by matter NSI's instead of CPT violation. In Ref. [76] it was shown how a non-zero $\varepsilon_{\mu\tau}$ could alter the oscillation phase, with opposite signs for neutrinos and anti-neutrinos, thus leading to the apparent discrepancy in Δm^2_{32}.

Interest in studying NSI's at long-baseline oscillation experiments arose after it was recognised that oscillation experiments would be particularly powerful tools for measuring NSI's because of the presence of an *interference* term between the oscillation and non-standard processes in the amplitude for the transition $\nu_\alpha \to \nu_\beta$ [66] - a neutrino flavour change can occur either via a NSI process or via a standard oscillation. The amplitude for the flavour change is the square of this sum: $|\text{NSI } (\nu_\alpha \to \nu_\beta) + \text{osc } (\nu_\alpha \to \nu_\beta)|^2$ and thus the leading order contribution to the NSI term is *linear* rather than quadratic in $\varepsilon_{\alpha\beta}$. This is beneficial if the amplitude of the oscillation is much greater than that of the NSI transition, but this is not necessarily true of the channels which future long-baseline

experiments seek to observe and which are discussed in the next chapter. However, exactly because the signals of some of these channels, in particular the *golden* channel, $\nu_e \to \nu_\mu$, are suppressed by θ_{13}, it may be possible that the NSI effect is actually of a comparable size to the standard oscillation. Thus, these channels may actually have an enhanced sensitivity to NSI's because there is less of an oscillatory 'background'.

The question of how to theoretically generate measurably large matter NSI's whilst respecting the gauge symmetry of the Standard Model has been rigorously studied in the literature e.g. in Refs. [74, 77]. We are interested in those NSI's which have the possibility of being detected in oscillation experiments - those that involve four leptons, such as in Fig. 2.4 and Eq. (2.33), and which may arise from the gauge-invariant dimension-6 operator

$$O^6 = \frac{1}{\Lambda^2}(\bar{L}_\sigma \gamma^\lambda L_\rho)(\bar{L}_\tau \gamma_\lambda L_\zeta), \tag{2.35}$$

where L is the leptonic $SU(2)$ doublet. The problem is that as well as flavour-changing *neutrino* transitions, gauge-invariance demands that there must also necessarily be flavour-changing *charged lepton* transitions of the same strength. However, there already exist strong bounds on several charged lepton flavour violating processes. For instance, for $\sigma = \mu$ and $\rho = \tau = \zeta = e$, there would be a neutrino vertex $\nu_\mu \nu_e e^- e^-$ which must be accompanied by a four charged lepton vertex, $\mu^- e^- e^- e^-$. The neutrino vertex is not bounded but the $\mu \to 3e$ branching ratio is constrained to the level of 10^{-12} [14]; bounds such as these must somehow be evaded if we want to create observable neutrino NSI's. The findings from Ref. [77] were that there is only *one* possible parameter, $\varepsilon_{\tau\tau}$, which may be created at the dimension-6 level and is not yet constrained by bounds on charged lepton processes if we consider mechanisms with one mediator only. Other parameters may be created in more complex schemes with more than one mediator, or at the dimension-8 level; in the latter case several cancellation conditions must be imposed to evade the charged lepton bounds.

So theoretically, the outlook for observable NSI's is not hopeful! However, from an experimental outlook with no theoretical prejudice, there do exist some NSI parameters which are not yet bounded below the observable level ($\sim 10^{-3}$) [78, 79] and which may be within the reach of near-term future experiments. Additionally, we should be open-minded and bear in mind the possibility that NSI's may be created by mechanisms which we have not yet considered. In any case, a rigorous test of our current assumptions and theories can be made by placing the best possible bounds on as many NSI parameters

as possible; a null observation and placing stronger bounds on NSI parameters would reinforce our evidence that certain aspects of the Standard Model are correct, whereas a discovery of non-zero NSI's would be definitive evidence of exciting new physics!

Chapter 3.

Neutrino oscillation experiments

Having presented the theoretical perspective of neutrino oscillations, in this chapter we will introduce the experimental aspects of neutrino oscillation physics. To begin with, we explain how the design of an experiment is linked to the theory described in the previous chapter. Then we will go on to describe the various neutrino sources and detector technologies which are used currently and those that will be used in the future, including a description of the super-beam and neutrino factory experiments that we study in this thesis. We then describe how these experiments can be simulated and how we perform quantitative comparisons between them.

3.1. From oscillation probabilities to oscillation experiments

The design of any experiment is guided by the signal for which it is being designed to measure, and so the starting point for neutrino oscillation experiments is to examine the form of the oscillation spectrum. Primarily, this is a sinusoidal function of the ratio $\frac{L}{E}$ (Eq. (2.8)). For experiments detecting solar or atmospheric neutrinos, both L and E are fixed by the production mechanism and the location of the source (e.g. the sun) so by counting the number of neutrinos observed and comparing it to the number expected to have been produced at the source, it is possible to fit the data, for that fixed value of L and E, to Eq. (2.8) and thereby deduce the values of the oscillation parameters on which that channel depends. As with any experiment, we want to maximise the number of events - this is especially true of neutrino experiments since neutrino interaction cross-sections are so tiny as mentioned in Section 1.2 - by using high fluxes and large detectors

if possible. The total number of events in an experiment is given by

$$\mathrm{N} = \text{flux} \times \text{cross-section} \times \text{detector mass} \times \text{running time}, \tag{3.1}$$

where the flux is the number of neutrinos per unit energy per unit area, at a specific distance away from the neutrino source and the cross-section is specific to the target (detector) material and is the number of interactions per unit of target mass per unit time per unit flux. We should also take into account the *efficiency* of the detector (as we will in Chapters 5, 6 and 7) and therefore multiply by the percent efficiency. Therefore to maximise the number of signal events, we should try to maximise all of these factors.

For man-made neutrino sources, it is possible to *choose* the ratio $\frac{L}{E}$, or to choose a *range* of values by producing neutrinos with a band of energies for a single fixed baseline, and also to choose which oscillation channel(s) to observe. *Disappearance* experiments are those that detect neutrinos which are of the same flavour as those produced at the source - they measure the 'survival probabilities' - the $\nu_\alpha \to \nu_\alpha$ channels. *Appearance* experiments instead detect neutrinos which have oscillated into other flavours during their propagation - $\nu_\alpha \to \nu_\beta$ with $\beta \neq \alpha$. In all cases the aim is to produce neutrinos with an energy spectrum such that it is possible to observe a significant portion of an oscillation wavelength for the baseline being studied. As described in Section 2.1, this is only possible if the argument of the spectral functions is $\sim O(1)$; if it is much less than unity then the oscillation wavelength is too long to be observable within the given energy range and so the spectrum appears to be flat; if it is much greater than one then the oscillations are too rapid to be resolved by the detector. Thus, for a chosen ratio of $\frac{L}{E}$, we can have sensitivity to either the solar oscillations by making $\frac{\Delta m^2_{21} L}{4E} \sim 1$, or to the atmospheric oscillations by choosing $\frac{|\Delta m^2_{31}| L}{4E} \sim 1$, as was shown in Fig 2.1. A precise measurement of Δm^2_{21} or Δm^2_{31} can then be made by making detailed measurements of the spectral dependence, in particular the energy-dependent *position* of the oscillation peak. The *amplitude* of the peak gives information about the mixing angle(s) relevant to the specific channel. As an example, the MINOS experiment [51], as already mentioned in Section 2.2 measures the $\nu_\mu \to \nu_\mu$ survival probability. At their baseline of 735 km and energy of $\sim$ 1 GeV, the probability is given by $P_{\mu\mu} \sim 1 - \sin^2 2\theta_{23} \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right)$ which is plotted in Fig. 3.1, as a function of the neutrino energy E (rather than $\frac{L}{E}$ as in Fig. 2.1). Thus we can see that the amplitude of the oscillation is given by the value of $\sin^2 2\theta_{23}$ whereas the position of the minimum depends on the value of $|\Delta m^2_{31}|$. Specifically, the position of the first minimum (this terminology refers to the minimum which occurs at the highest energy) is where $\frac{|\Delta m^2_{31}| L}{4E} = \frac{\pi}{2}$. Thus by measuring the energy

at which this occurs, E_{min}, it is possible to infer $|\Delta m^2_{31}|$ via $|\Delta m^2_{31}| = \frac{2\pi E_{min}}{L}$. Note that the dependence on the *square* of $\sin\left(\frac{\Delta m^2_{31} L}{4E}\right)$ is the reason why we do not yet know whether Δm^2_{31} is positive or negative.

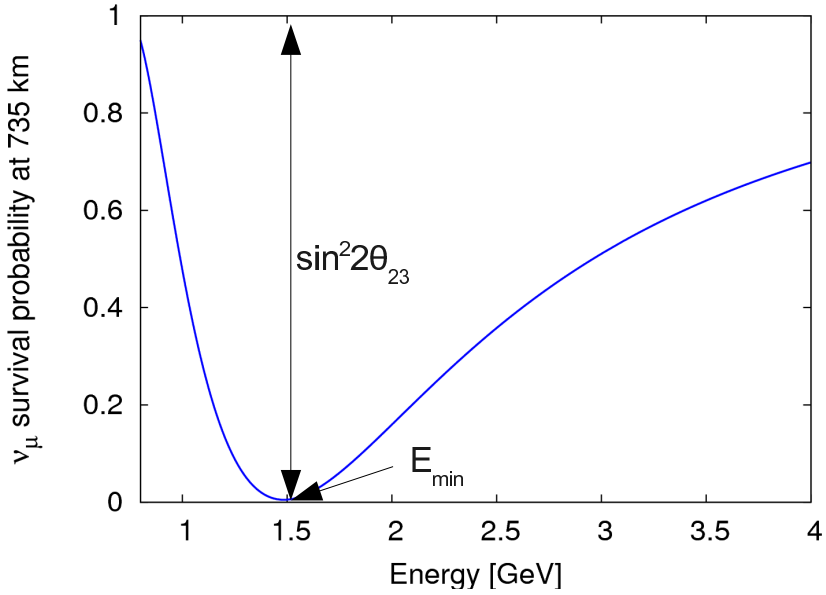

Figure 3.1.: Oscillation probability for $\nu_\mu \to \nu_\mu$ at 735 km (MINOS experiment). The amplitude of the oscillation is determined by the value of $\sin^2 2\theta_{23}$ and the position of the oscillation maxima and minima are determined by the value of $|\Delta m^2_{31}|$.

3.2. Oscillation experiments: neutrino sources

Since the dawn of the solar neutrino problem and suspicions that neutrinos might oscillate, there have been a multitude of experiments designed to test different regions of the neutrino mixing parameter space which have eventually led to the successful measurement of the parameters described in Section 2.3. These experiments have used neutrinos from both natural and man-made sources in conjunction with a diverse array of neutrino detectors, described below.

3.2.1. Natural neutrino sources

Neutrinos are produced in several naturally occurring sources including the sun, the atmosphere, supernovae and from the background radiation of rocks. The first two of

these can be used to study neutrino oscillations, with the advantage that it costs nothing to produce the neutrinos!

In the case of solar neutrinos, ν_e's are produced in multiple solar processes, mostly the *pp* (proton-proton) chain, resulting in a set of fluxes which are predicted by solar models [80]. The energies of solar neutrinos take values up to 10^7 eV. With current detection methods it is not possible to detect the entire energy range within one detector; a detector will only be sensitive to a portion of the energy spectrum. Solar neutrino experiments compare the number of neutrinos detected over a specific energy range at the fixed sun-to-earth baseline of $\sim 10^8$ km and compare it to the number predicted theoretically.

Atmospheric neutrinos are produced as decay products in hadronic showers which result from collisions of cosmic rays with nuclei in the upper atmosphere. Electron and muon neutrinos are produced in the decay chain of charged pions: $\pi^+ \to \mu^+\nu_\mu$, and similarly for π^-. The μ^+ then decay via $\mu^+ \to \bar{\nu}_\mu e^+\nu_e$ (similarly for μ^-). The ratio of atmospheric $\nu_\mu + \bar{\nu_\mu} : \nu_e + \bar{\nu_e}$ is therefore expected to be 2 : 1 for neutrinos with energies up to $\sim$ 2 GeV (higher energy muons reach the earth before decaying); the fact that a *smaller* ratio was observed by experiments such as Soudan-2 [81] was evidence that muon neutrinos were 'disappearing', hinting at neutrino oscillations.

The energies of atmospheric neutrinos are typically in the range 0.1 GeV to 100 GeV. At the distance of $L \sim 10$ km between the upper atmosphere and the earth's surface, neutrinos travelling directly from the atmosphere to the earth's surface ('down-going') do not oscillate; it is the neutrinos with energies of a few GeV that travel through the earth to reach the detector ('up-going') which will oscillate. Atmospheric neutrino experiments thus observe neutrino oscillations as a function of both the neutrino energy, and the angular distribution of the neutrinos.

3.2.2. Man-made neutrino sources

There are several advantages to using artificially produced neutrinos is an experiment: the neutrinos can be directed specifically at a detector, the flux is far greater than that of any natural source, and it is possible to place a second detector very near to the source to measure, rather than theoretically predict, the unoscillated spectrum which can't be done for solar or atmospheric neutrinos! Most importantly, it is possible to *choose* the ratio $\frac{L}{E}$ which we wish to observe by tuning the energy of the beam and

choosing the position of the detector with respect to the neutrino source. When the aim of past experiments was to measure the values of the mass-squared splittings, without any prior hints as to their magnitudes, different regions could be explored by situating the detector at different baselines. Now that these values are known to within a few per cent, the baseline is chosen such that the detector is situated at the position of the first oscillation maximum in order to maximise the signal rate.

The two artificial neutrino sources used today are nuclear reactors, and particle accelerators. Neutrinos from nuclear reactors are produced in the beta-decay of radioactive products, where a $\bar{\nu}_e$ is emitted. The energy is typically a few MeV. Although this is fixed, there is freedom to choose the baseline i.e. the location of the detector.

Neutrino beams created by particle accelerators provide the most intense sources of neutrinos available. An initial beam of protons is brought to collision with a graphite target, producing a secondary shower of mesons - mainly $\pi^{\pm}$ and $K^{\pm}$. Focusing magnets allow either the positive or negative mesons to be selected. The positively-charged mesons decay into a ν_μ beam: $\pi^+ \to \mu^+\nu_\mu$ and $K^+ \to \mu^+\nu_\mu$, whereas $\bar{\nu}_\mu$ are created in a similar process by the decay of π^- and K^-. Thus an intense beam of ν_μ ($\bar{\nu}_\mu$) is produced, with energies centred on a value determined by the energy of the proton beam. The precise energy spectrum is dictated by the kinematics of the meson decay. A small amount of secondary decay, $\mu^+ \to e^+\nu_e\bar{\nu}_\mu$, also occurs (and similarly for μ^-), so that a ν_μ ($\bar{\nu}_\mu$) beam contains intrinsic contamination from ν_e and $\bar{\nu}_\mu$ ($\bar{\nu}_e$ and ν_μ) at the level of a few percent.

3.3. Oscillation experiments: neutrino detectors

As pointed out in Section 2.5, 'neutrino detectors' *do not actually detect neutrinos!* They detect the particles produced or scattered as the result of a neutrino interaction which can be either a charged-current interaction, in which case the flavour of the neutrino can be identified, or neutral-current interactions, which give no information about the neutrino flavour.

Different detectors are sensitive to different lepton energies and flavours, which has led to a huge variety of detectors being used by past and current experiments. Unlike other particle detectors such as those used in collider experiments, neutrino detectors are essentially homogeneous as there is no need for multiple layers to detect multiple

types of particle. The emphasis instead is on maximising the event rate by maximising the volume of the detector, in order to maximise the chance of an interaction.

There are essentially two geometries for neutrino detectors - segmented and unsegmented. A segmented detector is one which is instrumented in sections, which means that these detectors are capable of tracking the path of a particle, of handling multi-source events, and that the detection medium may not be the same as the interaction (target) medium. Unsegmented detectors are instrumented as a single unit which means that for the same quantity and cost of electronics as a segmented detector, a much larger volume can be instrumented. However, multi-source events cannot be analysed. The detection medium is the same as the target medium. The detection medium in all detectors is either a Čerenkov medium or a scintillator, which will now be described.

3.3.1. Čerenkov detectors

Čerenkov radiation is emitted when a charged particle travels through a medium at a speed greater than that with which light travels through the medium. In the case of neutrino detectors, these particles will be the charged leptons produced or scattered in a neutrino interaction. The emitted light can be amplified and detected by photo-multiplier tubes (PMT's) and the information used to reconstruct the particle's energy, direction and flavour, and the interaction point and type (neutral-current or charged-current). Probably the most famous neutrino detector is the water Čerenkov (WC) Super-Kamiokande [44] detector, a steel tank filled with 50 kton of ultra-pure water with $\sim$ 11000 PMT's mounted on the walls (Fig. 3.2). WC detectors are capable of measuring low-energy ($\lesssim$ 1 GeV) ν_e ($\bar{\nu}_e$) and ν_μ ($\bar{\nu}_\mu$) events down to $\sim$ 5 MeV, when most events are *quasi-elastic*, $\nu_\alpha n \to \ell_\alpha^- p$ and $\bar{\nu}_\alpha p \to \ell_\alpha^+ n$.

In addition to water, other materials can be used as the Čerenkov medium, depending on the energies and flavours of the particles that the detector is designed for. For instance, the SNO detector [45] used D_2O instead of ordinary water which enabled additional channels to be accessible at low energies, such as the neutral-current process $\nu_\alpha(\bar{\nu}_\alpha)D \to pn\nu_\alpha(\bar{\nu}_\alpha)$. The MiniBooNE detector [82] uses mineral oil, in which particles emit both Čerenkov and scintillation light, whereas the IceCube neutrino telescope [83] is situated in the ice of the South Pole, allowing a very large volume (1 km^3) of pure ice to be used as the Čerenkov medium.

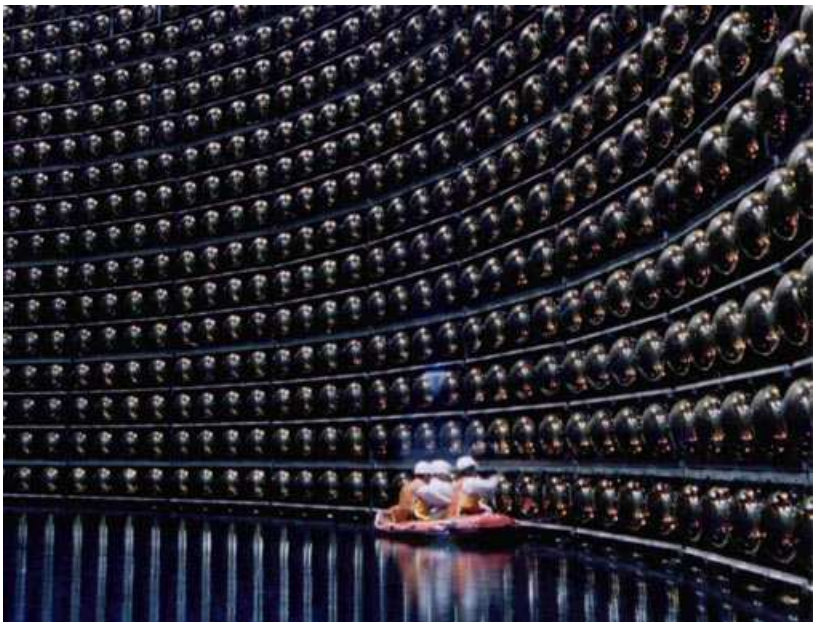

Figure 3.2.: The 50 kton water Čerenkov detector, Super-Kamiokande.

3.3.2. Scintillator detectors

Scintillating materials are those that produce light shortly after absorbing energy from a passing particle. In a neutrino detector, this light can be collected either by PMT's e.g. the future LENA detector [84], as for Čerenkov detectors, or read out by optical fibres as in the MINOS detectors [85]. MINOS is a *tracking calorimeter* - alternating planes of iron target and plastic scintillator enable the path of the lepton to be tracked, from which its kinematics and properties can be reconstructed. The advantages of using iron as a target are that it is dense, enhancing the interaction cross-section, and also that it is easy to magnetise. This is important for MINOS as it enables the momentum of the lepton to be determined very accurately, from its curvature in the magnetic field. Magnetisation will be important at a future *neutrino factory* (Section 3.5.3) when determination of the charge of the lepton is essential.

The next generation of scintillator detectors will be *totally active* scintillator detectors.

Totally Active Scintillating Detector (TASD)

The TASD [86] is a tracking calorimeter detector currently being used in the MINERvA experiment [87], and it will also be used in the forthcoming NOvA experiment [88]. The

phrase 'totally active' or 'active target' refers to the fact that all of the target medium is also the detector medium. This is different from other tracking calorimeters such as the MINOS detectors, where iron sheets constitute the target, and these are interleaven with layers of plastic scintillator which form the detector medium. The MINERvA detector is composed of extruded plastic tubular modules filled with a liquid scintillator such as mineral oil, with wavelength-shifting fibres embedded into each module. The use of liquid scintillator ensures sensitivity to low energy particles and the fine sampling means that this detector has excellent energy and spatial resolution. The optimal geometry for the plastic extrusions in terms of granularity and accurate vertex reconstruction has been determined to be triangular.

3.3.3. Liquid Argon Time Projection Chamber (LArTPC)

Liquid argon detectors were first proposed in Refs. [89, 90]. They are ionisation detectors, which means that when a particle passes through the liquid argon, it leaves a path of ionisation electrons which can be detected, tagging the path of the incoming particle. In a LArTPC, the paths of the ionisation electrons are drifted using an electric field. At one end of the detector there is a set of wires oriented in such a way that the time, magnitude and position of each path can be reconstructed. Snapshots are recorded at a frequency of $\sim$ 40 MHz, of the relative appearance of the ionisation electrons, thus allowing the particles to effectively be tracked in real-time. Put in sequence, the paths can be reconstructed, which results in bubble-chamber-like images such as the one shown in Fig. 3.3 from the ArgoNEUT experiment [91]. Specific interactions can be reconstructed from the track topology and the energy deposited along each track.

Liquid argon detectors enable excellent energy measurement and particle identification to be performed, thanks to their exceptional ability to resolve particle trajectories. The most successful LArTPC to date is the ArgoNEUT detector, which has a mass of 0.24 ton. In order to be effective for neutrino oscillation physics, a mass of the order of 100 kton is required. Much work is needed to prove that this scalability of six orders of magnitude is feasible, and also to address the problem of how to magnetise a detector of this scale. Another particular problem is how to obtain and maintain the purity of such a large quantity of liquid argon - reactive atoms and molecules which contaminate the liquid argon will absorb the ionisation electrons, depleting the signal. However, progress is becoming increasingly rapid with research groups in both the US [92] and Japan [93] working hard to aim towards a kton-scale detector.

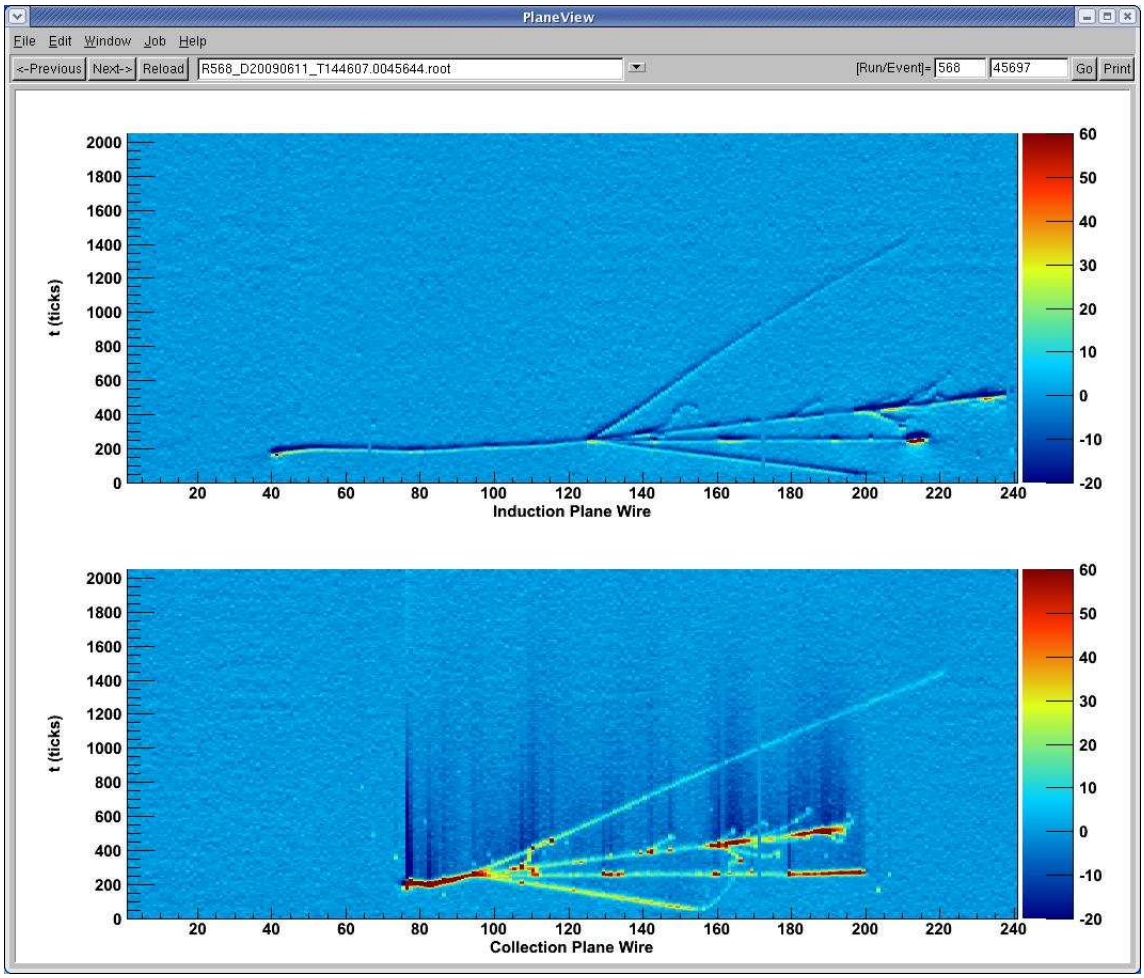

Figure 3.3.: An event from ArgoNEUT, a 0.24 ton LArTPC. From http://t962.fnal.gov/Images.html.

3.3.4. ν_τ detection

Recent interest in technologies to detect ν_τ's has partly been motivated by the desire to experimentally verify the predictions for the ν_τ channels made by the standard three-family oscillation model, which is the goal of the OPERA experiment [94]. A more recently developed motivation is the realisation that the best way to search for and bound the NSI's mentioned in Section 2.5 is by looking for the $\nu_\mu \rightarrow \nu_\tau$ channel with a *near detector* [70]. Unfortunately, the experimental search for ν_τ's turns out to be extremely challenging!

In the first place, it is difficult to make a ν_τ undergo a charged-current interaction - the threshold for τ production is ~ 3.5 GeV which is above the beam energies produced by current accelerators. Secondly, even if this is possible, τ's are very difficult to detect because of their extremely short lifetime (291 femtoseconds) - even if highly boosted, the decay length is only a few millimetres. Thus to see the signature of a τ - a short τ track followed by decay via a kink or to multi-prongs - requires a very fine-grained detector with excellent spatial resolution.

The ν_τ was first detected by the DONUT (Direct Observation of NU Tau) experiment [8] which observed the ν_τ's produced from the decay of charmed particles. Later on, CHORUS [95] was the first experiment to search for ν_τ's produced by the oscillation of ν_μ's and ν_e's (they obtained a null result as the $\frac{L}{E}$ value was too small). Currently running now is the OPERA experiment [94], which is optimised to detect the ν_τ's produced from the oscillation of ν_μ's produced at CERN. All of these experiments used 'emulsion cloud chamber' (ECC) detectors - thin films of photographic emulsion interleaved with thicker layers of target material, such as iron. In order for a τ to be identified, it must interact in the target material, and then the τ-specific signature - a kink or multi-prong decay - must occur in one of the layers of emulsion (Fig. 3.4). Altering the thicknesses of the target and emulsion layers changes the properties of the detector.

This technology has proven to be successful, but incredibly painstaking and inefficient - the chance of an appropriate τ interaction is extremely low, added to which the scanning of the emulsion layers must be done manually. The OPERA collaboration have recently succeeded in detecting their first ν_τ candidate [94], but with an expected event rate of around two per year, this technology as it currently stands is not sufficiently efficient to produce viable statistics! Developments in the forthcoming years may improve upon this, or a better candidate for ν_τ detection may turn out to be liquid argon time projection chambers (Section 3.3.3).

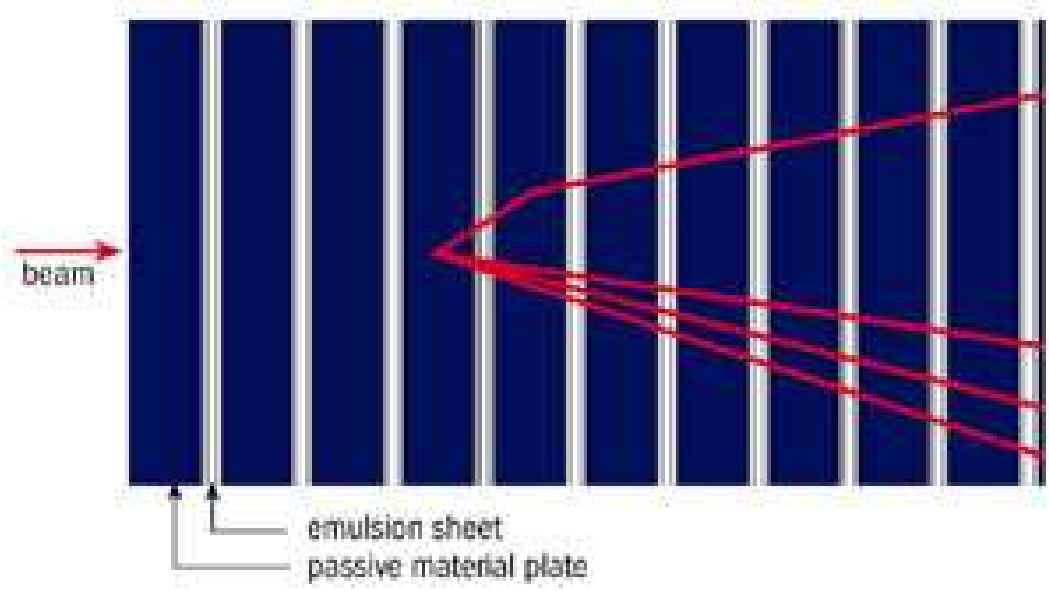

Figure 3.4.: Alternating layers of target and emulsion in the OPERA detector, and a kinked track characteristic of a ν_τ event. From cerncourier.com.

3.4. The future of neutrino oscillation experiments

The first neutrino oscillation experiments sought to verify the existence of neutrino oscillations - Homestake [43], Super-Kamiokande [44], SNO [45], GALLEX [96], SAGE [97] and Kamiokande [98]. Once accomplished, the next goal became to measure the oscillation parameters. Of these, the first to be measured turned out to be the solar mixing angle and mass-squared splitting (θ_{12} and Δm^2_{21}), measured by KamLAND [50], and the atmospheric mixing angle and corresponding mass-squared splitting (θ_{23} and Δm^2_{31}), measured by MINOS and Super-Kamiokande. It is these experiments and others which have consistently given us the measurements we now have, given in Section 2.3, and have enabled us to build the current model of three-family neutrino oscillations.

The exception is the LSND anomaly and, very recently, the latest MiniBooNE data. The LSND (Liquid Scintillator Neutrino Detector) experiment [99, 100] was an accelerator experiment searching for $\bar{\nu}_\mu \to \bar{\nu}_e$ and later $\nu_\mu \to \nu_e$ oscillations at a baseline of 30 m and energies between $\sim$ 20 MeV and $\sim$ 200 MeV. They detected a positive signal at 3.8σ confidence, consistent with neutrino oscillations driven by a mass-squared splitting of between 0.2 eV^2 and 10 eV^2, implying that at least one neutrino has a mass > 0.4 eV. Recall that all other experiments had produced data consistent with a solar mass-squared splitting of $\sim 8 \times 10^{-5}$ eV^2 and an atmospheric mass-squared splitting

of $\sim 2.5 \times 10^{-3}$ eV^2, so that the LSND data seemed to imply the existence of at least one additional neutrino, much heavier than the three Standard Model neutrinos. As an independent test of this result, the MiniBooNE (Mini-Booster Neutrino Experiment) experiment [101] was designed and constructed specifically to independently test the same region of parameter space as LSND. A decade after LSND, the most recent MiniBooNE data [102] are intriguing: the neutrino data are inconsistent with LSND oscillations but the (limited) anti-neutrino data appear to be consistent with the original anomaly (the data is only compatible with a no-oscillation model at 0.5% probability). The story continues...

LSND aside, the goal of future oscillation experiments is to complete and confirm the model of three-family oscillations. The current generation of experiments (T2K [103], DoubleChooz [104], Daya Bay [105], RENO [106]) have been designed to measure the third mixing angle, θ_{13}. Those being developed to start running (hopefully) in the next 10 to 15 years - NOvA [88], LBNE [107], LAGUNA [108] - seek to detect CP violation, if it exists, to search for a still smaller value of θ_{13} if it has not yet been discovered by the current generation of experiments, and to identify the mass hierarchy. These experiments can be successful only if $\sin^2 2\theta_{13} \gtrsim 10^{-2}$. The experiments needed if $\sin^2 2\theta_{13}$ is smaller than this value are described in the next section.

3.5. Next-generation neutrino oscillation experiments

The goal of the next generation of neutrino oscillation experiments is to either make a discovery of θ_{13} and CP violation and determine the mass hierarchy if these have not already been achieved by previous experiments or, if another experiment has already succeeded in making a measurement of one or more of these parameters, then to make a high-precision measurement of θ_{13} and δ, and furthermore, to search for NSI's. In order to do this, they will be designed to measure the $\nu_e \to \nu_\mu$ and $\bar{\nu}_e \to \bar{\nu}_\mu$ or $\nu_\mu \to \nu_e$ and $\bar{\nu}_\mu \to \bar{\nu}_e$ channels. These are subdominant oscillation channels, suppressed by a factor of $\sin^2 2\theta_{13}$, and are chosen because a single channel allows for the possibility to extract a measurement of θ_{13}, δ and the mass hierarchy. These channels will be discussed fully in the next chapter.

Next-generation experiments are designed to make *precision measurements*. This desire for precision in the neutrino sector stems partly from the need to match the precision which has been obtained in the quark sector for the parameters of the CKM matrix. Theoretically, it is expected that at some scale both the quark and the lepton mixing matrices are unified. The test of any model which predicts the way in which this unification is broken at the electro-weak scale can only be tested if the parameters of both matrices are measured to a similar degree of precision. At present, the CKM parameters have experimental errors of $\sim 1\%$ whereas the most precisely known neutrino mixing parameters have errors of $\sim 10\%$.

In order to accomplish these tasks, it will be necessary to use long baselines of over $\sim$1000 km in order to exploit matter effects to make a measurement of the mass hierarchy (Section 2.4). In order to achieve an adequately high flux in spite of the very long baseline, an extremely intense production source is required. There are three proposed candidates: the super-beam, the β-beam and the neutrino factory.

3.5.1. The super-beam

A 'super-beam' is the term used to describe a more powerful version ($>$ 1 MW) of a conventional accelerator neutrino beam (Section 3.2.2). The first of these is T2K [103].

It is possible to construct either a conventional narrow-band beam, like T2K, or a *wide-band beam* as is envisioned for LBNE [107] which has a broader spread of neutrino energies. In addition, the detector(s) can be placed *on-axis* or *off-axis*. *On-axis* means that the detector is placed directly perpendicular to the neutrino beam. *Off-axis* means that the detector is off-set from the line of the neutrino beam (by $\sim 1^\circ$) - the reason for doing this is that the kinematics of the neutrino production mean that in an off-axis position, a high flux of neutrinos with a very narrow spread of energies is obtained (Fig. 3.5). The decision to use a narrow-band or wide-band beam and an on-axis or off-axis detector depends on the specific goals of the experiment.

The advantage of a super-beam is that both the accelerator and detector facilities required are proven technologies. The main drawback is the presence of the intrinsic ν_e ($\bar{\nu}_e$) contamination of the ν_μ ($\bar{\nu}_\mu$) beam which limits the precision of the measurement of the $\nu_\mu \rightarrow \nu_e$ ($\bar{\nu}_\mu \rightarrow \bar{\nu}_e$) channel which is the primary channel of the super-beam. Additionally, there is a lack of precision relative to β-beams and neutrino factories with which the beam flux can be predicted - the beam originates from a hadronic decay and

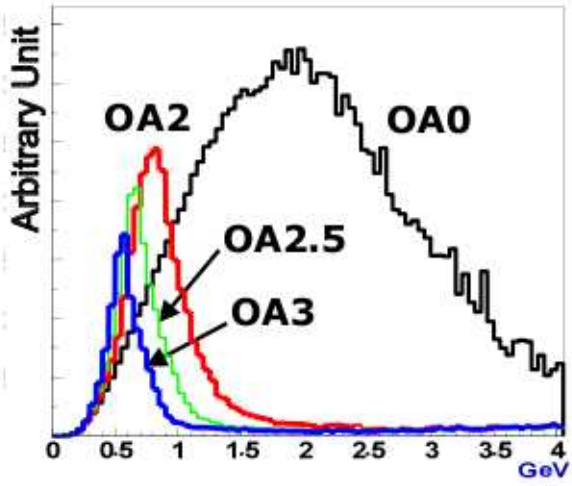

Figure 3.5.: Neutrino beam spectra on-axis (OA0) and off-axis by 2° (OA2), 2.5° (OA2.5) and 3° (OA3). Taken from Ref. [109].

therefore is difficult to predict down to the percent level. The use of a near detector will help to alleviate both of these problems, as will the results from current and future neutrino fixed-target experiments, run specifically to measure neutrino fluxes and cross-sections. We shall study super-beam setups in Chapter 5.

3.5.2. The β-beam

A β-beam [110] is produced from the decay of boosted radioactive ions, resulting in an extremely pure beam of ν_e or $\bar{\nu}_e$. The neutrino flux depends on the end-point kinetic energy of the electron produced in the decay, the γ factor of the ion, and the baseline. As the kinematics of the leptonic decay process are known very well, the neutrino flux can be calculated very precisely once these three factors are fixed. The absolute purity of the beam is an advantage; the disadvantages are the technological and practical difficulties associated with constructing an accelerator which can produce sufficiently boosted ions (the more highly boosted the ions, the more energetic the neutrinos) to enable the β-beam to be used with a very long baseline. The physics potential of a variety of β-beam setups has been extensively analysed in studies such as Refs. [111, 112, 113, 114, 115, 116].

3.5.3. The neutrino factory

A neutrino factory [117] is seen to be the ultimate neutrino oscillation experiment. As for a super-beam, the initial beam is a proton beam which produces a secondary beam of pions. The pions decay into muons and neutrinos, but rather than using the neutrinos from this secondary beam, the neutrino factory captures the muons and accelerates them. The muons decay via $\mu^- \to \nu_\mu e^- \bar{\nu}_e$ and $\mu^+ \to \bar{\nu}_\mu e^+ \nu_e$, with the neutrinos having a range of energies up to the energy of the accelerated muons. Thus this tertiary neutrino beam consists of a pure mix of 50% ν_μ ($\bar{\nu}_\mu$) and 50% $\bar{\nu}_e$ (ν_e), allowing for up to *twelve* oscillation channels to be studied. In addition, an extremely high intensity can be reached, and all beam characteristics can be very precisely predicted as the neutrinos are produced by a leptonic decay which is a lot 'cleaner' than a hadronic process. The design of the neutrino factory is the task of the International Design Study for the Neutrino Factory (IDS-NF) [86], to which part of the work in this thesis contributes.

Fig. 3.6 shows the experimental layout of the accelerator section of the experiment. The conventional neutrino factory setup, as defined by the IDS-NF, accelerates muons to 25 GeV, with an estimated 1.4×10^{21} muons per year per polarity (μ^- and μ^+) being produced. The muons are divided into two racetrack-shaped storage rings; in the straight sections the muons decay to form the neutrino beam. One beam is aimed towards the 'intermediate detector' at a baseline of between 3000 km and 5000 km, and the second beam is aimed at the far detector at the 'magic baseline' (see Section 4.5) of 7000 km to 8000 km. The exact lengths of the baselines are subject to restrictions from geographical locations. Each of the detectors is a magnetised iron neutrino detector (MIND) [118], the closer one having a mass of 100 kton and the farther one a mass of 50 kton. These are essentially larger versions of the MINOS detectors described in Section 3.3.2. A near detector (or possibly more than one), the exact design of which is still under consideration, will be situated close to the straight section(s) of one or both of the muon storage rings to make flux and cross-section measurements, and to search for NSI's as described in Section 2.5. This setup is what we shall refer to from now onwards as the 'high-energy neutrino factory' (HENF) to make a clear distinction from the low-energy neutrino factory (LENF) which is the topic of Chapters 6 and 7.

The neutrino factory exploits the 'golden' signature of the *wrong-sign* muon events [117, 119] - muons with opposite sign to the muons stored in the neutrino factory. If μ^+ are stored, wrong-sign muons (μ^-) result from $\nu_e \to \nu_\mu$ oscillations, and can be used to measure the mixing angle θ_{13}, determine the neutrino mass hierarchy, and search

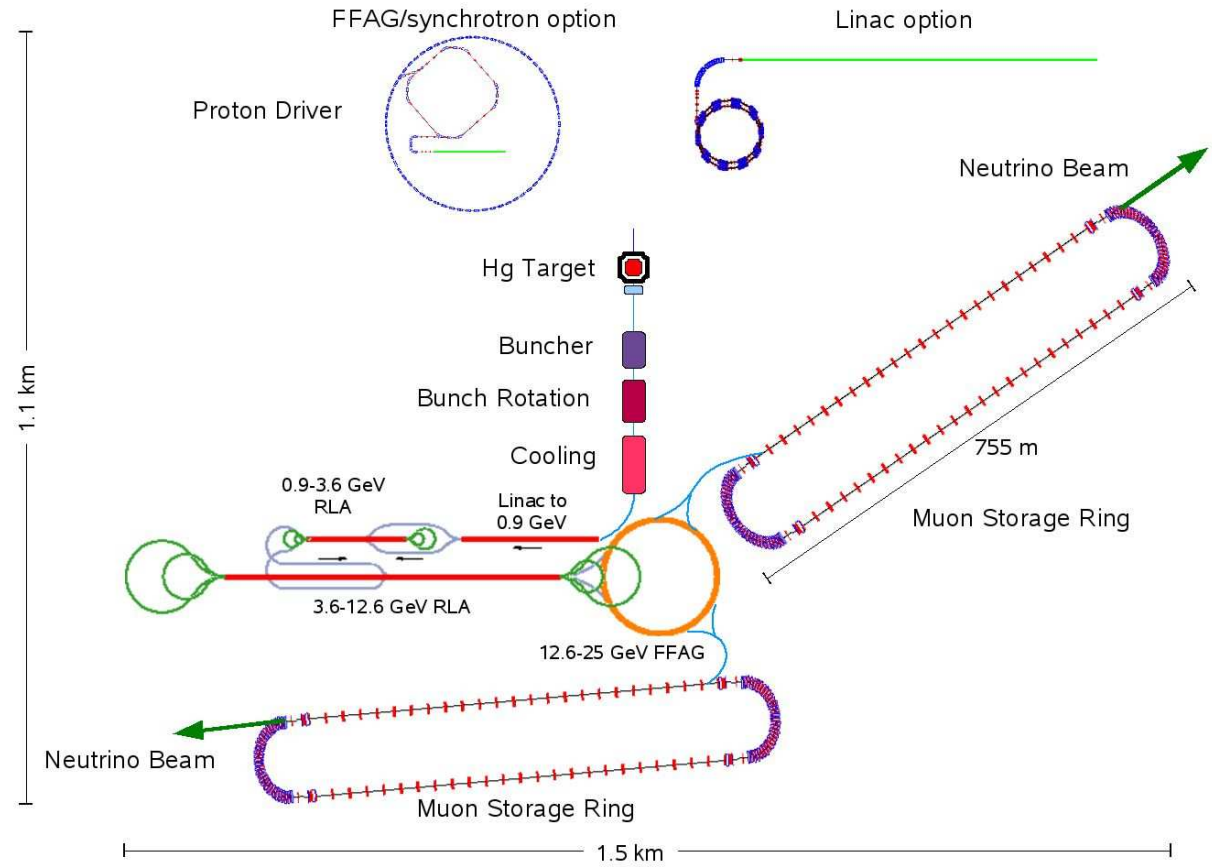

Figure 3.6.: Layout of the neutrino factory accelerator complex as defined by the International Design Study for the Neutrino Factory. From Ref. [86].

for CP violation (see next chapter). In addition to the *wrong-sign* muon signal there will also be *right-sign* muon events. These events come from the disappearance muon neutrino channel, $\bar{\nu}_\mu \to \bar{\nu}_\mu$ ($\nu_\mu \to \nu_\mu$), if μ^+ (μ^-) are stored. The discrimination of the *wrong* and *right* sign muons requires the identification of charged-current ν_μ and $\bar{\nu}_\mu$ interactions, and the measurement of the sign of the produced muon. If the interacting neutrinos have energies of more than a few GeV, standard neutrino detector technology based on large magnetised tracking calorimeters, like those described in Section 3.3, can be used to measure wrong-sign muons with high efficiencies and very low backgrounds. This has been shown to work for neutrino factories with energies of about 20 GeV or more [62, 86, 120].

Neutrino factories have been shown to be extremely sensitive tools for studying neutrino oscillation physics [62, 117, 119, 120, 121, 122, 123, 124, 125, 126, 127, 128, 129, 130, 131, 132, 133, 134, 135, 136, 137, 138] and the HENF setup described above has been found to be optimal both in terms of standard oscillation physics and for NSI searches [67, 139, 140].

In this thesis we shall study a *low-energy* neutrino factory. An obvious question, in view of the excellent performance of the HENF, is why one should consider an alternative, and in particular, lower energy setup? 'Low-energy' in this context means that the muons are accelerated to around 5 GeV, rather than 25 GeV, and the corresponding (single) baseline is $O(1000)$ km [136, 137]. The physics motivation for this idea will be demonstrated at the end of Chapter 6 when we shall see that for a particular range of values of $\sin^2 2\theta_{13}$, a neutrino factory with a shorter baseline and lower energy is *more sensitive to CP violation than the HENF setup*, whilst still retaining a similar sensitivity to θ_{13} and the mass hierarchy as the HENF. Practical advantages of the LENF include cost - an accelerator for 5 GeV muons will cost substantially less than a 25 GeV accelerator - and construction factors.

We will focus on the experimental optimisation of the low-energy neutrino factory and analyse its sensitivity to both oscillation parameters and NSI's in Chapters 6 and 7.

3.6. Simulations of future experiments

3.6.1. GLoBES

Current experiments collect and analyse real data. In order to probe the power of future experiments, we need to *simulate* data, which means simulating firstly the 'true' events (events which actually occur, as created by nature), and then the 'reconstructed' events seen by the experiment. The true event is an incident neutrino with a specific energy. This is transmitted via secondary particles and detector properties into a reconstructed event which will hopefully be a good approximation to the original neutrino event. In general, reconstructed events will not correspond exactly to the true events because no experiment is perfect - there are experimental limitations and finite measurement errors. The question we need to address is how these limitations and errors affect our ability to obtain information about the true neutrino events.

A software package which is designed to do precisely this is the General Long-Baseline Experiment Simulator (GLoBES) [141, 142] which has been used for all the simulations described in this thesis. Each of the three components of an oscillation experiment - production, propagation (oscillation) and detection - are simulated, using information specified by the user. The information used to simulate the source and detector is specified in the 'experiment file'. The beam composition, power, flux and running time must be defined to describe the source. The detector is described in terms of its mass, energy range, energy-dependent energy resolution, energy-dependent efficiencies, energy-dependent backgrounds and systematic errors for each of the channels which it can detect. Cross-section files, as a function of neutrino energy, must also be defined - we use those that come with GLoBES, which originate from Refs. [143, 144]. In addition, the baseline and a matter profile type must be specified. We have used a constant average density (calculated by GLoBES from Refs. [145, 146]) together with a matter density uncertainty of 2% which provides a very good approximation for the baselines which we shall be simulating.

The baseline and matter information are used to simulate the propagation stage. The user must specify true values for all the oscillation parameters from which the oscillation probabilities for all specified channels are calculated. In the case that a baseline-dependent matter density profile is chosen, the profile is divided up into steps of constant density. In our case of constant density, there is just one step. The Hamiltonian in matter (see Eq. (2.17)) with the specified true values of the oscillation parameters

and a matter potential calculated from the matter profile, is diagonalised at each step, propagating the neutrino transition amplitude along that step. The final oscillation probability is the square of the sum of all the transition amplitudes. From these probabilities and the source information, the rates of all the specified oscillation channels are defined - these are the *true* rates. To obtain the *reconstructed* rates, each true event is transformed by the detector properties which maps it onto a reconstructed event.

We want to know if it is possible to infer the *true values* of the oscillation parameters from the *reconstructed* rates. To do this, GLoBES performs a statistical 'parameter estimation' test where the user can choose how many and which parameters are to be constrained. Constraints are performed using a chi-squared (χ^2) minimisation test. Ideally, a minimisation should be performed over *all* parameters except those which are being constrained. GLoBES is designed to perform these multi-parameter minimisations efficiently, which is fortunate since the neutrino oscillation space contains six oscillation parameters and the matter density is also treated in GLoBES as an oscillation parameter, so that there are seven possible parameters over which to minimise!

In reality, a neutrino oscillation experiment, either real or simulated, follows a Poisson distribution: the measurement is of a discrete number of events within an allocated time period, the events are independent of one another (ignoring the dead-time of the detector electronics which is essentially irrelevant considering the tiny neutrino cross-sections) and there is a predicted number of expected events. For a parameter estimation, a comparison is made between the detected data and a set of hypotheses, where each hypothesis is an oscillation model where the oscillation parameters take particular values. By making this comparison, it can be seen which sets of parameters are statistically compatible with the data, and hence particular parameters can be constrained. In a real experiment, the hypotheses are based on the data taken at the near detector which gives the number of unoscillated events expected, taking into account the detector response. A set of oscillation parameters is then chosen and the oscillated data expected at the far detector are inferred from the near detector data and the chosen set of oscillation parameters. In simulations, it is assumed that the unoscillated data and detector response will be known to a certain accuracy specified by the values of the systematic errors, and the expected oscillated data are inferred from these assumptions.

Experiments have a finite energy resolution and so events are 'binned' into energy bins of a finite width. For a given energy bin j, the likelihood of measuring x_j events

given a hypothesised number μ_j is

$$L(x_j, \mu_j) = \frac{\mu_j^{x_j}\ e^{-\mu_j}}{x_j!}. \tag{3.2}$$

However, it is also necessary to include the effects of the systematic errors mentioned above, which can be treated as Gaussian:

$$L(x_j, \mu_j, n_{j,k}, \sigma_{j,k}) = \prod_k \frac{\mu_j^{x_j}\ e^{-\mu_j}}{x_j!} e^{-\frac{n_{j,k}^2}{2\sigma_{j,k}^2}}, \tag{3.3}$$

where $n_{j,k} = s_{j,k} - \rho_{j,k}$, the indices k correspond to each of the individual errors, $s_{j,k}$ is the actual value of the systematic error k in bin j, $\rho_{j,k}$ is the mean value of the error k in bin j and $\sigma_{j,k}$ is its standard deviation. In the case that there are a large number of events, the Poisson distribution can be approximated as a Gaussian distribution with mean and variance μ_j such that the likelihood is

$$L(x_j, \mu_j, n_{j,k}, \sigma_{j,k}) = \prod_k \sqrt{\frac{1}{2\pi\mu_j}}\ e^{-\frac{(x_j-\mu_j)^2}{2\mu_j}}\ e^{-\frac{n_{j,k}^2}{2\sigma_{j,k}^2}}. \tag{3.4}$$

A χ^2 *minimisation* corresponds to finding the value of μ_j and the values of the $n_{j,k}$'s and $\sigma_{j,k}$'s which *maximise* the likelihood $L(x_j, \mu_j, n_{j,k}, \sigma_{j,k})$; this is the 'best-fit' point. Maximising the likelihood corresponds to *minimising* the 'χ^2 statistic',

$$\chi^2 = \sum_{j,k} \frac{(x_j - \mu_j)^2}{\mu_j} + \frac{n_{j,k}^2}{2\sigma_{j,k}^2}. \tag{3.5}$$

In our simulations we have treated the systematic errors as being uncorrelated, which means that each of the $n_{j,k}$'s and $\sigma_{j,k}$'s are minimised independently. More sophisticated treatments are possible but for our current studies these would have a negligible effect on the results. The minimisation is performed for each bin, and the sum of all the minimised χ^2 values for each bin gives the total χ^2_{min}.

A χ^2 *fit* means that the χ^2 statistic is simply evaluated for each given set of hypothesised parameter values, although the systematics are still minimised. Again, the χ^2 value is evaluated for each bin individually and the sum gives the total χ^2. The parameter set which gives the smallest value of χ^2 is the best-fit point. At this point it is important to make the distinction between the *true values* of the parameters and their *test values*. The *true values* are the parameter values which are used to simulate the real data, and

thus are those values which we are supposing that nature has chosen. The *test values* are the parameter values of the hypothesis being tested. In our simulations, the minimum value of χ^2 will always be zero and the best-fit point will always lie at the true values (assuming that we include in our hypotheses the one which actually corresponds to the true values) because the data are simulated according to a model which is identical to one of our hypotheses, the detector response is fully accounted for by minimising over the systematic errors, and statistical fluctuations are not included; therefore there will be a perfect correpondence between the simulated data and one of the hypotheses. If this were *not* the case such that $\chi^2_{min} \neq 0$, then the measure of compatibility would be the value of $\Delta\chi^2$ which is the difference between the χ^2 value at a point and χ^2_{min}. Since $\chi^2_{min} = 0$ in our simulations, $\Delta\chi^2$ is the same as χ^2. This quantity is indicative of how compatible the data are with the hypothesis.

In this thesis, we perform one- and two-parameter fits. For a parameter estimation, this corresponds to having one or two degrees of freedom (d.o.f.) - this is a different definition than that for a 'goodness-of-fit' test. In the case of a one-parameter fit, for instance θ_{13}, we want to know how likely it is that a particular value of θ_{13} is compatible with the data, *regardless of the values of all the other parameters.* Therefore we want to find the minimum value of χ^2 for that particular value of θ_{13} by marginalising over all the other parameters - this means simply that all the other parameters, including the matter density and systematic errors, are allowed to vary so that the smallest χ^2 value can be found. If this value is below a certain confidence level (conventionally 3σ - 99.7% confidence or 5σ - 99.99994% confidence; the corresponding value of χ^2 depends on the number of degrees of freedom) then that θ_{13} value *is* consistent with the data at 3σ (5σ). In the case of two-parameter fits, for instance θ_{13} and δ, we want to know how likely it is that a *pair* of these values is compatible with the data, regardless of the value of all the other parameters. These results are conventionally shown as contour plots with the 68% (1σ), 90% and 95% (2σ) contours displayed. The power of GLoBES is not only that it can perform these computations efficiently, but that by fixing certain parameters to take certain values, degenerate solutions can be located by forcing the minimiser into a local minimum rather than the global minimum. This is most commonly used to obtain the 'fake' solutions arising from the hierarchy degeneracy (see Section 4.5) - if we simulate a true normal hierarchy, with the true Δm^2_{31} being positive, then start the minimiser at a *negative* value of Δm^2_{31}, it is possible to see if there are regions in which an inverted hierarchy is also compatible with the data. In order to do this we need to find the minimum possible χ^2 assuming an inverted hierarchy, relative to the true minimum ($\chi^2 = 0$). To prevent the minimiser from falling into the true minimum, the

allowed values of Δm^2_{31} can be constrained by specifying an input error. Input errors are also always specified for the other parameters - usually the 1σ errors from experiments - in order to ensure that the minimiser locates the correct minimum.

3.6.2. Statistical analysis and standard performance indicators

In order to compare the performances of different experiments, as will be done in Chapters 5 and 6, it is necessary to take into account that because oscillation probabilities and therefore experimental sensitivities are dependent upon the values of θ_{13} and δ (see Section 4.1), especially if comparing experiments which measure different oscillation channels and are therefore optimally sensitive to different regions of the $\theta_{13} - \delta$ plane, the performance of the experiment needs to be assessed for all possible pairs of values of θ_{13} and δ. The merit of an experiment is judged to be the amount of 'coverage' it obtains in this plane - as the values of θ_{13} and δ are unknown, the best experiment is the one which has the best chance of being able to obtain measurements for the largest number of possible (θ_{13}, δ) values. Computationally, this means setting up a two-dimensional grid of points in the $\theta_{13} - \delta$ plane, and at each point setting these values to be the *true values* with which the data are simulated. A χ^2 fit as described in the previous section is then performed at each point, so that it is possible to see in which regions of the $\theta_{13} - \delta$ parameter space the experiment has sensitivity.

There are two conventional presentation methods. The first is to show sensitivities as a function of both $\sin^2 2\theta_{13}$ (this parameter is used rather than θ_{13} itself since the golden channel probability depends on the quantity $\sin^2 2\theta_{13}$ rather than θ_{13} - see Chapter 4) and δ so that one can see explicitly the regions of the parameter space to which a particular experiment is sensitive. In general, the smaller the value of $\sin^2 2\theta_{13}$ the harder it is to make a measurement. But the minimum value of $\sin^2 2\theta_{13}$ for which an experiment has sensitivity depends on the value of δ - these points are those that lie along the displayed contour. All points which lie to the *right-hand side* of the contour (assuming $\sin^2 2\theta_{13}$ is displayed on the horizontal axis) are the points for which the particular measurement *can* be made, and all points to the left are those to which the experiment has no sensitivity.

The second method is more useful for making precise quantitative comparisons between different experiments. Sensitivities are shown as a function of $\sin^2 2\theta_{13}$ but are given in terms of the 'CP-fraction' (also called the 'δ-fraction'). This is a measure of the fraction of points in δ-space for which the measurement is successful, for a fixed value of $\sin^2 2\theta_{13}$. In other words, if the computer simulation is written such that the

δ-space is divided into n steps ($n+1$ points separated by a step-size of $\frac{2\pi}{n+1}$) then the CP-fraction for each value of $\sin^2 2\theta_{13}$ is the number of points for which a measurement was successful divided by $(n+1)$. This fraction is calculated for each value of $\sin^2 2\theta_{13}$ and the points are plotted. A perfect experiment would achieve a CP-fraction of 1 for all values of $\sin^2 2\theta_{13}$; realistically the CP-fraction drops as $\sin^2 2\theta_{13}$ decreases.

The three standard performance indicators which are used are the following:

- **θ_{13} discovery potential** is the ability of an experiment to *exclude a value of* $\theta_{13} = 0$, if the true value of θ_{13} is non-zero. At each point in the $\sin^2 2\theta_{13} - \delta$ plane, those pair of values are set to be the *true values* with which the data are simulated. The question is how compatible the detected data are with a hypothesis where $\theta_{13} = 0$ - this is the *test value*. A χ^2 minimisation of the data is performed with all parameters left free to be minimised over except for θ_{13} (which is fixed at zero). This includes the sign (positive or negative) of Δm^2_{31} i.e. it may be possible that the best fit is obtained with the incorrect hierarchy. The point at which the χ^2 value is minimal is the best-fit point, and will lie somewhere in the $\theta_{13} = 0$ plane. We need to know where this minimum value of χ^2 lies relative to a certain threshold (3σ in this thesis). If it lies *above* the threshold then this indicates that a good fit cannot be obtained to $\theta_{13} = 0$ and thus θ_{13} can be excluded at 3σ confidence, indicating that its value is non-zero. If, on the other hand, the χ^2 value lies *below* the threshold then it is possible to fit the data to $\theta_{13} = 0$ and thus this point cannot be excluded, indicating that θ_{13} is too small to be measured by this experiment. Typically this will be for small true values of $\sin^2 2\theta_{13}$, and maximal values of δ where cancellation can occur between the CP and atmospheric terms (see Section 4.1).

- **CP discovery potential** is the ability to *exclude a value of* $\delta = 0$ *or* π which corresponds to being able to confirm that there *is* CP violation in the neutrino sector. The computational method is similar to that for θ_{13} discovery, but instead of fixing the test value to be $\theta_{13} = 0$, the test points are $\delta = 0$ and $\delta = \pi$. These values are *fixed* and all other parameters are minimised over to find the best-fit point. Again, this marginalisation includes the sign of Δm^2_{31}. Whichever of $\delta = 0$ or π has the lower χ^2 value is taken to be the best-fit point, and whether this χ^2 value lies above or below the 3σ threshold determines the plotted curves. This measurement is difficult for values of δ which lie close to 0 or π and becomes progressively easier as the value of δ approaches maximal CP violation ($\pm\frac{\pi}{2}$).

- **Hierarchy sensitivity** is the ability to *exclude the wrong mass hierarchy.* In the case of simulating a true normal hierarchy, for each point in the $\sin^2 2\theta_{13} - \delta$ plane, a fit is forced to an inverted hierarchy, allowing all other parameters to vary to find the best-fit point. If the χ^2 value at this point lies below the 3σ threshold then it is not possible to identify the mass hierarchy at that confidence level as good fits can be obtained to both hierarchies. If the χ^2 value is above the threshold, then the wrong hierarchy can be excluded at 3σ confidence.

An additional presentation tool which is used for more qualitative purposes are the 'potato plots' (as they are informally known!) which display the results of the two-parameter $\theta_{13} - \delta$ fits. These are used to gauge the sensitivity of an experiment to θ_{13} and δ and are also a useful way of visualising degenerate solutions. These are the plots which appear in, for example, Fig. 6.4. The pairs of values of θ_{13} and δ which are compatible with the detected data at the 68%, 90% and 95% confidence levels are enclosed by the contours. The displaced 'satellite' which appear in Fig. 6.4a (which is for a true value of $\theta_{13} = 1^\circ$), at the $\delta = 90^\circ$ point but at the *wrong* value of θ_{13}, corresponds to the region which is compatible with the data if the *wrong hierarchy* is assumed. The fact that there are no wrong-hierarchy regions in the $\theta_{13} = 5^\circ$ plot (Fig. 6.4b) indicates that the hierarchy can be correctly identified at the 95% confidence level in this case.

Chapter 4.

Phenomenology of future long-baseline oscillation experiments

In this chapter we will describe the oscillation phenomenology at future long-baseline neutrino oscillation experiments, in the context of a neutrino factory. The phenomenology of a super-beam is qualitatively similar. To begin with, we will examine the primary channel at the neutrino factory - the so-called 'golden channel', $\nu_e \to \nu_\mu$, showing how measurements of θ_{13}, δ and the mass hierarchy can be extracted. We will then return to the topic of matter effects, considering their effects at different baselines, then briefly mention the other oscillation channels which may be accessible to a neutrino factory. Finally, we will discuss the problem of degeneracies and explain why the 'magic baseline' is so magical! In this chapter, we shall concentrate only on the standard oscillation parameters, leaving an examination of the non-standard interactions until Chapter 7.

4.1. The golden channel, $\nu_e \to \nu_\mu$

As already mentioned, a neutrino factory will search for the wrong-sign muons from the $\nu_e \to \nu_\mu$ and $\bar{\nu}_e \to \bar{\nu}_\mu$ channels. These are the so-called *golden channels* [62], named because this single channel is rich enough to provide information on all the unknown oscillation parameters - θ_{13}, δ and the mass hierarchy ($\text{sign}(\Delta m^2_{31})$). The fact that this channel is highly suppressed by the small value of θ_{13} and therefore is a subdominant oscillation mode, is the reason that not until recently has there been sufficient technological development to realistically allow for the observation of this channel. T2K is the first experiment to do so [103].

Let us begin by first considering the golden channel probability in vacuum. This is simply an application of Eq. (2.8), and the result is

$$
\begin{aligned}
P^{\text{vac}}_{\nu_e\to\nu_\mu} &= s^2_{213}s^2_{23}\left(c^2_{12}\sin^2\left(\frac{\Delta_{31}L}{2}\right)+s^2_{12}\sin^2\left(\frac{\Delta_{32}L}{2}\right)\right) && \text{(4.1a)}\\
&+ \frac{1}{2}s_{213}c_{13}s_{212}s_{223}\cos\delta\left(\sin^2\left(\frac{\Delta_{31}L}{2}\right)-\sin^2\left(\frac{\Delta_{32}L}{2}\right)\right) && \text{(4.1b)}\\
&+ \frac{1}{4}s_{213}c_{13}s_{212}s_{223}\sin\delta\left(\sin\left(\Delta_{31}L\right)-\sin\left(\Delta_{32}L\right)\right) && \text{(4.1c)}\\
&+ \left(c^2_{13}s^2_{212}c^2_{23}-\frac{1}{2}s^2_{213}s^2_{212}s^2_{23}+\frac{1}{2}s_{213}c_{13}s_{212}c_{212}s_{223}\cos\delta\right)\times\\
&\quad \sin^2\left(\frac{\Delta_{21}L}{2}\right) && \text{(4.1d)}\\
&- \frac{1}{4}s_{213}c_{13}s_{212}s_{223}\sin\delta\sin\left(\Delta_{21}L\right). && \text{(4.1e)}
\end{aligned}
$$

We use a notation where $s_{jk}=\sin\theta_{jk}$, $s_{2jk}=\sin 2\theta_{jk}$, $c_{jk}=\cos\theta_{jk}$, $c_{2jk}=\cos 2\theta_{jk}$, $\Delta_{jk}=\frac{\Delta m^2_{jk}}{2E}$, E is the neutrino energy and L is the baseline. The probability for anti-neutrinos is obtained by exchanging $\delta\to-\delta$ (this is also the probability for the $\nu_\mu\to\nu_e$ channel which super-beams observe). Note that we have not automatically made the approximation $\Delta m^2_{32}=\Delta m^2_{31}$ as it will soon be made apparent that this is not a valid assumption for all the situations which we shall consider.

The terms in line (4.1a) and also those terms which involve $\cos\delta$ - (4.1b) and (4.1d) - are the CP conserving terms. Those which depend on $\sin\delta$ - (4.1c) and (4.1e) - are the CP violating terms. The prefactor of the CP violating terms, $s_{213}c_{13}s_{212}s_{223}\sin\delta$, is the Jarlskog prefactor which forms part of the Jarlskog determinant, J [147], an invariant originally defined as being the commutator of the up-type and down-type quark mass matrices but which applies analogously to the neutrino sector. When written in terms of the lepton mass matrices it becomes apparent that J vanishes if any two masses are identical, which translates into mixing angles as saying that CP is conserved if any of the mixing angles are zero (neutrino states with identical masses will not mix). When written as above in terms of mixing parameters, it can be seen that CP violation vanishes also if $\delta=0$ or π. All CP violating effects in the Standard Model are proportional to J.

The simplest situation to consider first is for an energy and baseline tuned to the atmospheric oscillations only. In this case we *can* use the $\Delta m^2_{32}=\Delta m^2_{31}$ approximation

and Eq. (4.1) reduces to a simple form which is dependent only upon θ_{13}, θ_{23} and Δm^2_{31}:

$$P^{\text{vac,atm}}_{\nu_e \to \nu_\mu} \approx s^2_{213} s^2_{23} \sin^2 \left(\frac{\Delta_{31} L}{2} \right). \tag{4.2}$$

This approximation becomes more precise as the length of the baseline is decreased as this minimises contamination from matter effects, improves the approximation $\Delta_{32} L \approx \Delta_{31} L$, and also suppresses the solar terms, (4.1d) and (4.1e), by minimising the value of $\Delta_{21} L$. It is valid for the T2K experiment [103] which has a baseline of 295 km, and can therefore search for θ_{13} via this appearance channel, independently of the reactor experiments which measure the $\bar{\nu}_e$ disappearance probability given in Section 4.4.2. Note that in this regime there is no sensitivity to CP violation, nor to the mass hierarchy.

Next we consider an intermediate situation where the baseline is long but matter effects are neglected. In this situation we cannot immediately neglect any terms in Eq. (4.1) as the quantity $\sin \Delta_{21} L$ and the difference between $\Delta_{32} L$ and $\Delta_{31} L$ may now be of a comparable magnitude to the term (4.1a), depending on the value of θ_{13}. We write everything in terms of two oscillation frequencies by making use of the relation $\Delta m^2_{32} = \Delta m^2_{31} - \Delta m^2_{21}$, which applies to both hierarchies (refer to Fig. 2.2) with Δm^2_{31} and Δm^2_{32} being positive in the case of a normal hierarchy and negative in an inverted hierarchy. We also make the approximation $\sin \Delta_{21} L \approx \Delta_{21} L$ and consider all terms which are second order in s_{213} and $\Delta_{21} L$:

$$\begin{aligned}
P^{\text{vac,L}}_{\nu_e \to \nu_\mu} &= s^2_{213} s^2_{23} \sin^2 \left(\frac{\Delta_{31} L}{2} \right) && \text{(4.3a)} \\
&+ s_{213} c_{13} s_{212} s_{223} \left(\frac{\Delta_{21} L}{2} \right) \sin \left(\frac{\Delta_{31} L}{2} \right) \cos \left(\delta - \frac{\Delta_{31} L}{2} \right) && \text{(4.3b)} \\
&+ c^2_{13} s^2_{212} c^2_{23} \left(\frac{\Delta_{21} L}{2} \right)^2 . && \text{(4.3c)}
\end{aligned}$$

So in this situation there are three terms and sensitivity to θ_{13}, δ and the mass hierarchy ($\text{sign}(\Delta m^2_{31})$). Again, the anti-neutrino probability is obtained by exchanging $\delta \to -\delta$. The term (4.3a) is known as the atmospheric term and is sensitive to θ_{13}, (4.3b) is the CP term which is sensitive to δ and $\text{sign}(\Delta m^2_{31})$, and (4.3c) is the solar term which is sensitive to none of the above parameters. These will be mentioned again shortly when we include matter effects. The CP term is also called the *interference* term because it arises from mixing between the atmospheric and solar regimes. In the case of maximal CP violation, $\delta = \frac{\pi}{2}$ or $\frac{3\pi}{2}$, the probabilities for normal and inverted hierarchies

are identical and so cannot be distinguished. For all other values of δ these probabilities are different and can, in principle, be distinguished as shown in Fig. 4.1. Here we show the probabilites for both normal hierarchy (NH) and inverted hierarchy (IH) for $\theta_{13} = 3°$ and $\delta = 0$ (solid lines) and $\frac{\pi}{4}$ (dashed lines). We use a baseline of 1300 km which is the standard LENF baseline we shall be using in Chapters 6 and 7. From this plot it can be seen that the IH spectrum is essentially an inversion of the NH spectrum which arises because the hierarchy sensitivity originates from the factor of $\sin\left(\frac{\Delta_{31}L}{2}\right)$ in the CP term; changing the sign of Δm^2_{31} changes the CP term from a positive interference to a negative interference.

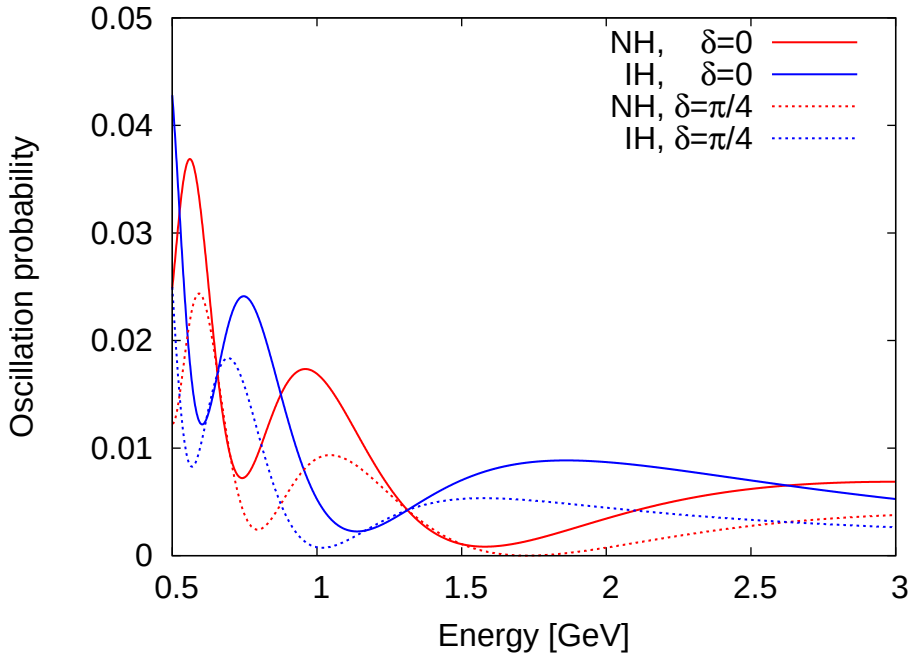

Figure 4.1.: Golden channel probability for a normal hierarchy (NH) and inverted hierarchy (IH) at 1300 km without matter effects, with $\theta_{13} = 3°$ and $\delta = 0$ and $\frac{\pi}{4}$. (The values of the other oscillation parameters are $\sin^2 2\theta_{12} = 0.3$, $\theta_{23} = \frac{\pi}{4}$, $\Delta m^2_{21} = 8.0 \times 10^{-5}$ eV^2 and $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2).

Relevant to the neutrino factory and other long-baseline experiments is the oscillation probability including matter effects. A calculation of the golden channel probability, following the method described in Section 2.4.2 by treating Δ_{21} as a perturbation, is

described in detail in Appendix F. The result is

$$P^{\rm mat}_{\nu_e\to\nu_\mu} = s^2_{213}s^2_{23}\frac{\sin^2\left(\frac{(\Delta_{31}-A)L}{2}\right)}{\left(1-\frac{A}{\Delta_{31}}\right)^2} \quad (4.4a)$$

$$+ \; s_{213}c_{13}s_{212}s_{223}\frac{\Delta_{21}}{A}\sin\left(\frac{AL}{2}\right)\frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}}\cos\left(\frac{\Delta_{31}L}{2}-\delta\right) \quad (4.4b)$$

$$+ \; c^2_{13}s^2_{212}c^2_{23}\left(\frac{\Delta_{21}}{A}\right)^2\sin^2\left(\frac{AL}{2}\right), \quad (4.4c)$$

where $A=\sqrt{2}G_F n_e$ as derived in Section 2.4. To obtain the anti-neutrino probability, exchange $\delta\to-\delta$ and $A\to-A$. The first line of the probability, (4.4a), is the *atmospheric term* which is quadratic in $\sin 2\theta_{13}$ and will be dominant in the scenario that θ_{13} is large ($\sin^2 2\theta_{13}\gtrsim 10^{-2}$), and at high energies. The atmospheric term provides sensitivity to θ_{13}, the mass hierarchy, and is sensitive to the octant of θ_{23} ($\theta_{23}<\frac{\pi}{4}$ or $>\frac{\pi}{4}$). The second line, (4.4b), is the *CP term* which is linear in $\sin 2\theta_{13}$ and dominates for intermediate values of θ_{13} if δ is not close to 0 or π. The dependence on δ enters via the oscillatory cosine term which can take either a positive or negative sign depending on the value of the phase. This can lead to constructive or destructive interference between the atmospheric and CP terms, meaning that sensitivities to θ_{13} and the mass hierarchy are strongly dependent on the value of δ. Due to the inverse dependence on energy, the CP term becomes most visible at *lower* energies; therefore it is important to have access to the second oscillation maximum, by using a detector with a low energy threshold, to establish if CP is violated. Thus a shorter baseline and therefore lower energy is desirable to enable a clean measurement of δ, whereas a higher energy and, especially, a long baseline, guarantees sensitivity to the mass hierarchy, as will be explained shortly. The third line, (4.4c), is the *solar term* which is independent of δ and the mass hierarchy, and is dominant in the case that θ_{13} is very small ($\sin^2 2\theta_{13}\lesssim 10^{-4}$). In this regime, measurements will be extremely challenging.

It is the dependence of the golden channel spectrum on the value of θ_{13} which makes the measurement of θ_{13} so crucial. The phenomenology changes drastically depending on the magnitude of $\sin^2 2\theta_{13}$, and dictates the optimisation of the experiment. In Figs. 4.2 and 4.3 we show how each of the three individual terms varies with the value of $\sin^2 2\theta_{13}$. The largest value of $\sin^2 2\theta_{13}=0.1$ corresponds roughly to the 3σ bound given in [53] ($\theta_{13}\simeq 10^\circ$). The CP term is shown for three values of the CP phase: $\delta=0$, $\frac{\pi}{4}$ and $\frac{\pi}{2}$. In Figs. 4.2a and 4.2b we show the terms in the case of a normal mass hierarchy for an

energy of, respectively, 1 GeV and 5 GeV. In Figs. 4.3a and 4.3b we show the same for an energy of 3 GeV (this is the region of the first oscillation peak for the LENF baseline and so is the most important part of the energy spectrum), and show the terms for both the cases of a normal and an inverted hierarchy.

The atmospheric term (red) is proportional to $\sin^2 2\theta_{13}$ (although here it is shown on a log scale, so does not appear that way!) and thus grows more rapidly with θ_{13} than the other terms. It is proportional to $(1 - \frac{2EA}{\Delta m^2_{31}})^{-2}$ which means that it is enhanced at high energies.

The CP term (green) is proportional to $\sin 2\theta_{13}$ and therefore grows more slowly with θ_{13} than the atmospheric term. Its primary dependence on energy is E^{-1} and as such it is enhanced at low energies. The CP term is the only one which can take a negative or positive sign, depending on the type of hierarchy and the value of δ. It is for this reason that experimental sensitivities will be heavily dependent upon not just the value of θ_{13}, but also on the value of δ and the hierarchy - there are some values for which the CP term interferes destructively with the atmospheric and solar terms, leading to cancellation in certain parts of the spectrum.

Finally, the solar term (blue) is essentially a background term, as it has no dependence on any of the parameters. It is unaffected by the value of θ_{13}, so that in the regime that θ_{13} is very small and the other terms are suppressed, the solar term remains constant and will dominate. It has a E^{-2} dependence, such that it can be 'beaten' by high energies - the atmospheric term overtakes the solar term at a smaller value of $\sin^2 2\theta_{13}$ at higher energies, and although increasing the energy also suppresses the CP term, the solar term is suppressed more quickly.

4.2. θ_{13}, δ and the mass hierarchy from the golden channel

In Fig. 4.4 we show how the values of θ_{13} and δ affect the golden channel spectrum. In Fig. 4.4a we fix a value of $\delta = 0$ and a normal mass hierarchy, and show the spectrum for values of $\theta_{13} = 1^\circ, 3^\circ$ and 10° ($\sin^2 2\theta_{13} \simeq 10^{-3}, 10^{-2}$ and 10^{-1}). As mentioned above, the atmospheric term has the strongest dependence on θ_{13}, and grows with energy. Therefore it is most prominent at the first oscillation peak (the highest energy maximum), and so the effect of θ_{13} is greatest at this first peak - it can be seen in Fig. 4.4a that the change

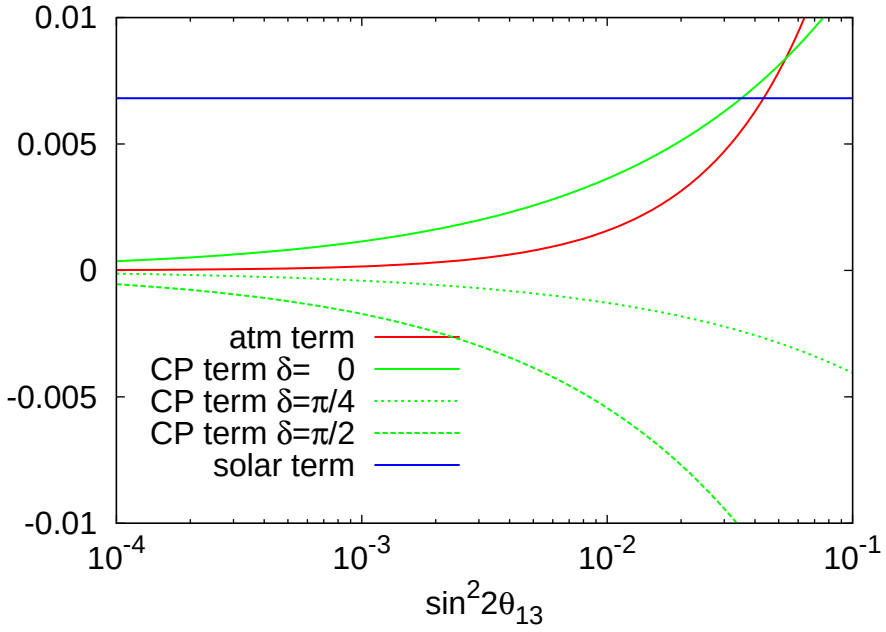

(a) $E = 1$ GeV, normal hierarchy.

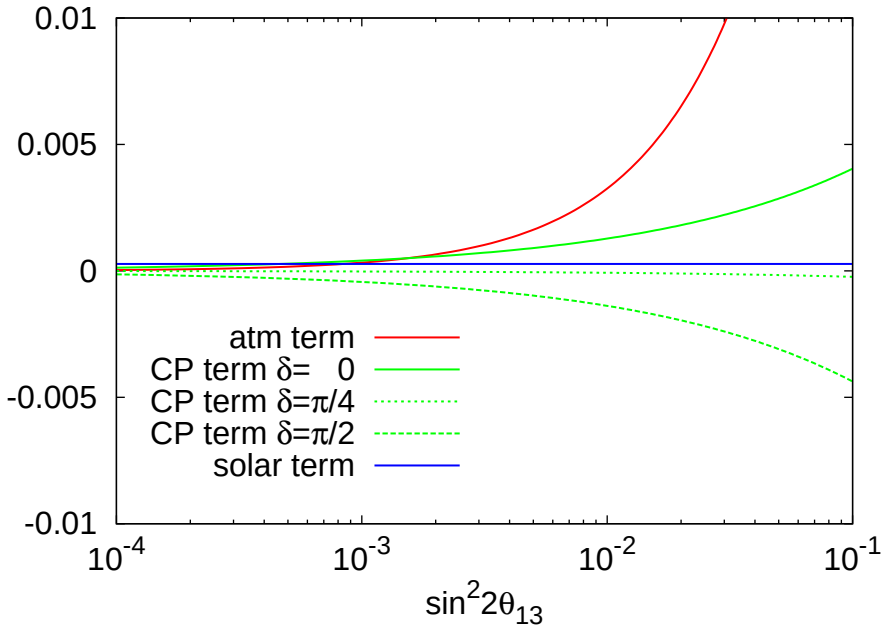

(b) $E = 5$ GeV, normal hierarchy.

Figure 4.2.: Variation of the atmospheric, CP and solar terms of the golden channel at 1300 km, including matter effects, as a function of $\sin^2 2\theta_{13}$, for a) $E = 1$ GeV and b) $E = 5$ GeV, in the case of a normal hierarchy. The y axis shows the partial probability for each of the terms. Values of $\delta = 0$, $\frac{\pi}{4}$ and $\frac{\pi}{2}$ are shown. (The values of the other parameters are $\sin^2 2\theta_{12} = 0.3$, $\theta_{23} = \frac{\pi}{4}$, $\Delta m^2_{21} = 8.0 \times 10^{-5}$ eV^2, $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2 and $A = 1.4 \times 10^{-22}$ GeV).

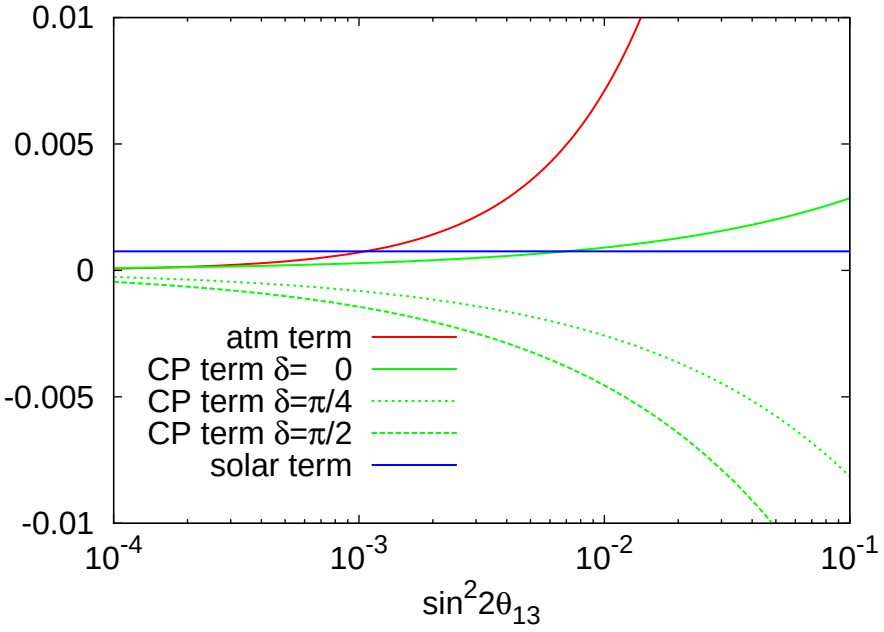

(a) 3 GeV, normal hierarchy.

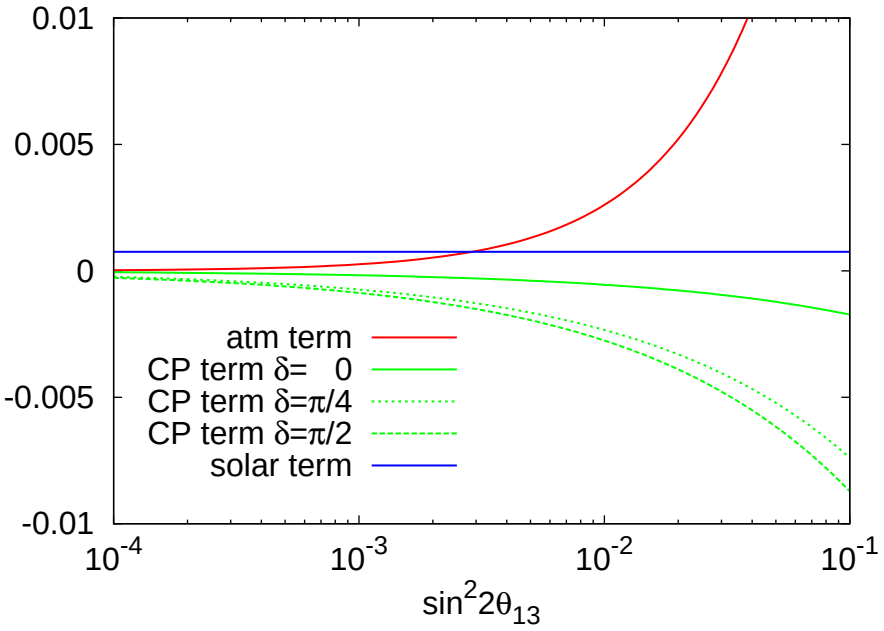

(b) 3 GeV, inverted hierarchy.

Figure 4.3.: Variation of the atmospheric, CP and solar terms of the golden channel at 1300 km, including matter effects, as a function of $\sin^2 2\theta_{13}$ for $E = 3$ GeV, for a) normal hierarchy and b) inverted hierarchy. The y axis shows the partial probability for each of the terms. Values of $\delta = 0$, $\frac{\pi}{4}$ and $\frac{\pi}{2}$ are shown. (The values of the other parameters are $\sin^2 2\theta_{12} = 0.3$, $\theta_{23} = \frac{\pi}{4}$, $\Delta m^2_{21} = 8.0 \times 10^{-5}$ eV^2, $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2 and $A = 1.4 \times 10^{-22}$ GeV).

in size of the first oscillation peak for different values of θ_{13} is more drastic than for the second or third peak.

In Fig. 4.4b we fix a value of $\theta_{13} = 3°$ and a normal mass hierarchy, and show values of $\delta = 0$, $\frac{\pi}{2}$, π and $\frac{3\pi}{2}$. The value of δ affects both the amplitude and the position of the oscillation. As mentioned above, the CP term can interfere either constructively or destructively with the other two terms, hence the reason for the effect on the oscillation amplitude. In addition, the CP term contains spectral sine and cosine factors, and so δ shifts the position of the oscillation peaks. Comparing the lines for $\delta = \frac{\pi}{2}$ (green) and $\delta = \pi$ (blue), we see an example of how it may be difficult to distinguish between these two cases - even though these correspond to the extreme scenarios of maximal CP violation and CP conservation, the spectra look very similar and distinguishing between them may be a difficult task.

In Fig. 4.5 we again fix a value of $\theta_{13} = 3°$ and show how the hierarchy affects the spectrum, for both a CP conserving ($\delta = 0$ - red line) and maximally CP violating ($\delta = \frac{\pi}{2}$ - blue line) case. The solid and dotted lines correspond to, respectively, normal and inverted hierarchies. Here we see an illustration of what was described in Section 2.4 - that for neutrinos (which we are currently considering), oscillations are enhanced for a normal hierarchy, and suppressed for an inverted hierarchy. This is at least always true if we consider the region of the first oscillation maximum, but is not necessarily true at the lower energy peaks, where it can be seen that the amplitude for $\delta = 0$ is actually slightly greater for an inverted than normal hierarchy. The reason for this is that at lower energies the atmospheric term is less relevant. The solar term starts to become relevant, which has no dependence on the hierarchy, and the amplitude of the lower energy peaks is then determined primarily by the sign of the CP term which depends on both δ and the hierarchy. Any combination which results in a CP term with positive sign will enhance the probability (for instance, inverted hierarchy and $\delta = 0$ or normal hierarchy and $\delta = \frac{\pi}{2}$, as shown here) and any combination which results in a negative term will suppress the oscillation.

The optimisation of the neutrino factory is dictated by these considerations and the need to be able to resolve the different spectra shown in Figs. 4.4 and 4.5 to extract the values of θ_{13}, δ and the mass hierarchy. This leads to the following requirements, which also apply to any other experiment measuring the golden channel:

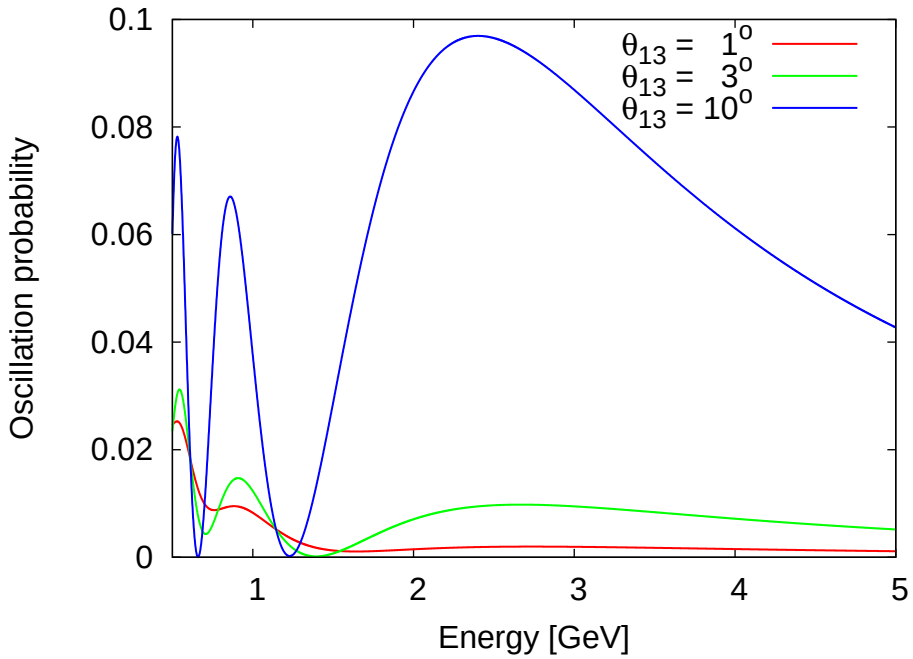

(a) Effect of θ_{13} on the golden channel probability ($\delta = 0$, normal hierarchy).

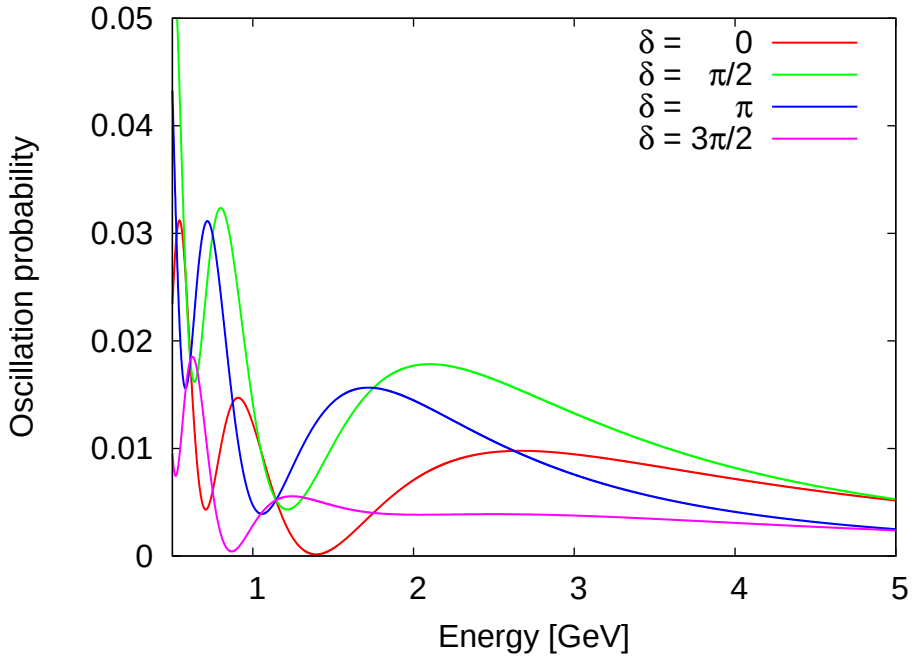

(b) Effect of δ on the golden channel probability ($\theta_{13} = 3°$, normal hierarchy).

Figure 4.4.: Alteration of the golden channel spectrum at 1300 km, including matter effects, by a) θ_{13} ($\theta_{13} = 1°$, $3°$ and $10°$) with $\delta = 0$ and a normal hierarchy and b) δ ($\delta = 0$, $\frac{\pi}{2}$, π and $\frac{3\pi}{2}$) with $\theta_{13} = 3°$ and a normal hierarchy. (The values of the other parameters are $\sin^2 2\theta_{12} = 0.3$, $\theta_{23} = \frac{\pi}{4}$, $\Delta m^2_{21} = 8.0 \times 10^{-5}$ eV^2, $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2 and $A = 1.4 \times 10^{-22}$ GeV).

- θ_{13} determines the *magnitude* of the oscillation peaks. The fact that θ_{13} is small is the reason that the appearance channel is so strongly suppressed relative to the disappearance channels. Therefore *high statistics* are required.

- The CP phase δ manifests itself most prominently at low energies, for instance at the second oscillation maximum, due to the inverse dependence on energy of the CP term. In other words, it is easier to distinguish between CP conservation ($\delta = 0$ or π) and CP violation at the second oscillation peak than at the first. For this reason, it is important that future detectors have a *low energy threshold* so that the second peak can be observed.

- Conversely, the difference between a normal and inverted mass hierarchy is easier to observe at high energies (the first oscillation peak) than at low energies (the second peak). It is driven by the matter effects discussed in Section 2.4 and is the reason that *long baselines* are needed for these experiments, which is discussed further in the next section.

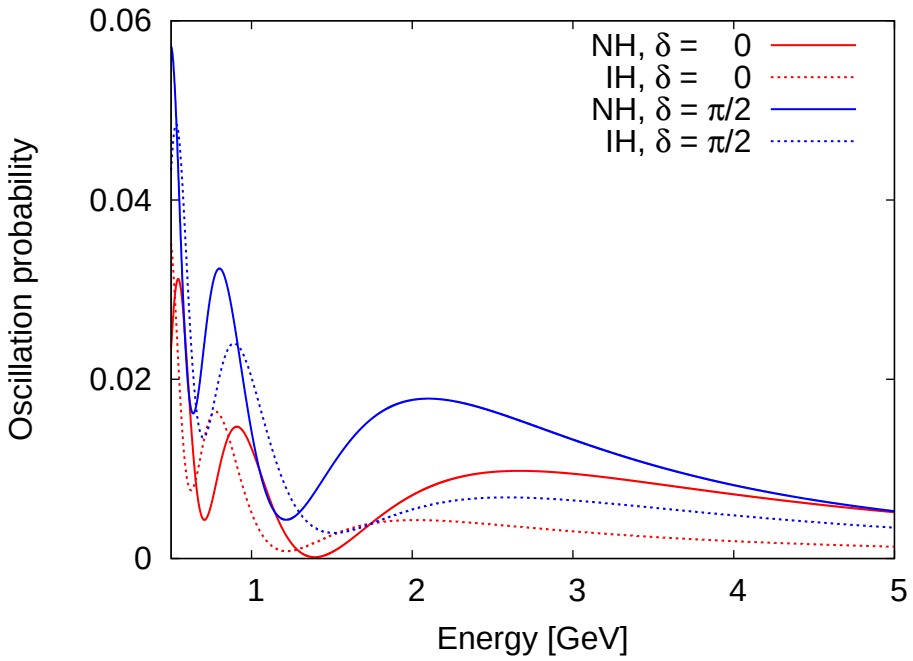

Figure 4.5.: Alteration of the golden channel spectrum at 1300 km, including matter effects, by $\text{sign}(\Delta m^2_{31})$ for $\theta_{13} = 3°$ and $\delta = 0$ and $\frac{\pi}{2}$. (The values of the other parameters are $\sin^2 2\theta_{12} = 0.3$, $\theta_{23} = \frac{\pi}{4}$, $\Delta m^2_{21} = 8.0 \times 10^{-5}$ eV^2, $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2 and $A = 1.4 \times 10^{-22}$ GeV).

In summary, for an experiment to successfully extract the oscillation parameters from the spectrum of the golden channel, the following are required: high statistics, a

detector with a low energy threshold, and a long baseline. It is also advantageous to have a detector which has good energy resolution, such that the *shape* of the spectrum can be accurately resolved as this enables similar shaped spectra arising from different parameter values to be distinguished from one another.

4.3. Matter matters

The importance of a long baseline becomes obvious by comparing Eqs. (4.2) and (4.3) - a long-baseline makes the interference term non-negligible, introducing CP violating dependence and hierarchy dependence into the probability. The effect of a non-zero matter potential can be seen by comparing Eqs. (4.3) and (4.4). We shall discuss two matter-related topics: the effect of matter on determining the mass hierarchy, and the effect on the detection of CP violation.

4.3.1. Matter effects and the mass hierarchy

A short-baseline experiment has no sensitivity to the mass hierarchy as Eq. (4.2) shows. A long baseline with no matter effects does, except for particular values of δ (Eq. (4.3)), as does a long baseline with matter effects included. The baseline at which matter effects become relevant for hierarchy determination can be seen by comparing the oscillation spectra for different baselines with and without matter effects as shown in Figs. 4.6 and 4.7. In Fig. 4.6 we consider baselines of 500 km (Fig. 4.6a) and 1300 km (Fig. 4.6b), which is the LENF baseline. In Fig. 4.7 we consider baselines of 4000 km (Fig. 4.7a) - the HENF near detector baseline, and 8000 km (Fig. 4.7b) - the HENF far detector baseline, for $\theta_{13} = 3^\circ$ and $\delta = 0$. For the two shortest baselines (Fig. 4.6), we take the average matter density to be $\sim 2.7\text{g/cm}^3$ (density of the earth's crust) which converts to $A = 1.4 \times 10^{-22}$ GeV. For the two longer baselines (Fig. 4.7) the beamline will also cross a portion of the earth's mantle, and the average density is taken to be $\sim 4.3\text{g/cm}^3$ which means that $A = 2.2 \times 10^{-22}$ GeV.

At all baselines, the effect described in Section 2.4 can be seen: matter *enhances* the oscillations for a *normal hierarchy* around the first oscillation maximum (so the probability is greater in matter than in vacuum) and *suppresses* the oscillations for an *inverted hierarchy* (so the probability is smaller in matter than in vacuum). This is not necessarily true for the lower energy peaks which are not so susceptible to matter effects

- the matter dependence of each of the terms in Eq. (4.4) roughly follows the energy dependence of each of the terms described in Section 4.1. In addition, there is also an energy-shift of the oscillation since the oscillations now depend on $\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)$ rather than $\sin\left(\frac{\Delta_{31}L}{2}\right)$.

Starting with the shortest baseline of 500 km (Fig. 4.6a), it is actually apparent that in this case matter effects are detrimental to the ability to separate the two hierarchies. The separation between the two spectra is more pronounced in the vacuum case. However, making the distinction in either the vacuum or matter case would be very difficult - take note of the scale of the oscillations. For the 1300 km baseline (Fig. 4.6b), the difference between the NH and IH is altered but not increased significantly relative to the vacuum case. However, the difference in either the vacuum or matter case is a lot more prominent than for the 500 km baseline (again, take note of the scale) - only at baselines greater than $\sim$ 1000 km is it experimentally feasible to differentiate between the hierarchies.

For the longer baselines (Figs. 4.7a and 4.7b), matter effects are far more significant. The effects are clearly visible at the first oscillation maximum, and particularly favourable in the case of a normal hierarchy when the oscillations are enhanced, thus increasing the statistics. It is also obvious to see that at these baselines the splitting between the normal and inverted hierarchy spectra is hugely enhanced at the first oscillation maximum, and this is how the HENF will determine the mass hierarchy.

4.3.2. Matter effects and CP violation

CP violation implies that there is a difference between the way that particles and anti-particles behave. Theoretically, it is possible to make a measurement of CP violation using one channel alone, for instance just the golden channel, by extracting a measurement of δ from the oscillation spectrum. The problem is that there are degeneracies and correlations between the parameters, which will be described later in this chapter and in Appendix H. Also, the accuracy of the measurement is limited by experimental errors. Therefore, to search for such a theoretically important phenomenon, it is preferable to make a *direct* observation of CP violation; that is by directly observing a difference in the behaviours of neutrinos and anti-neutrinos. In the vacuum case, Eq. (4.1) tells us that there is only a difference between neutrino and anti-neutrino probabilities if $\delta \neq 0$ or π. However, from Eq. (4.4) we find that even if $\delta = 0$, *neutrinos and anti-neutrinos*

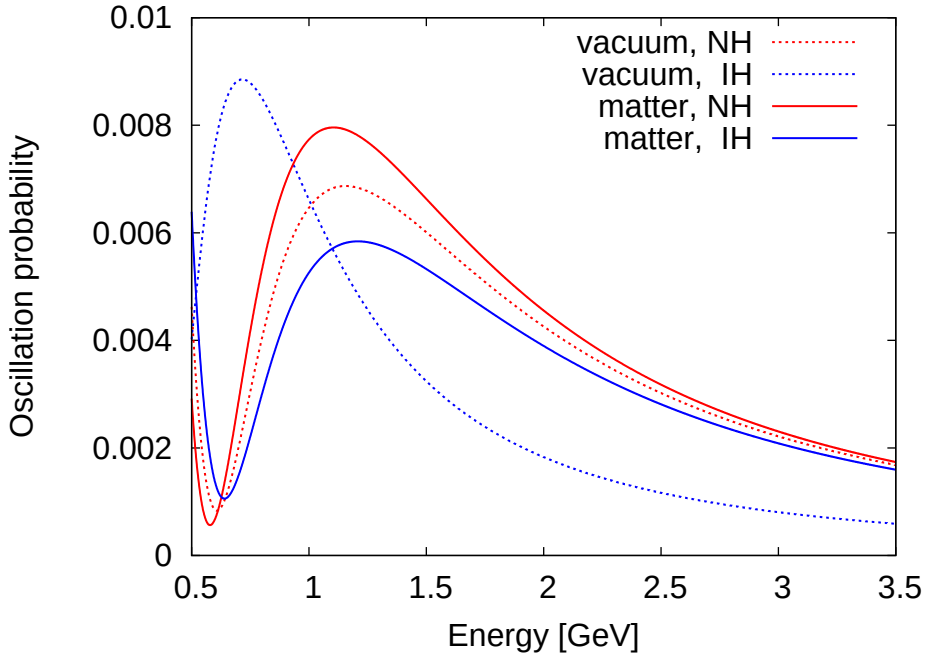

(a) $L = 500$ km.

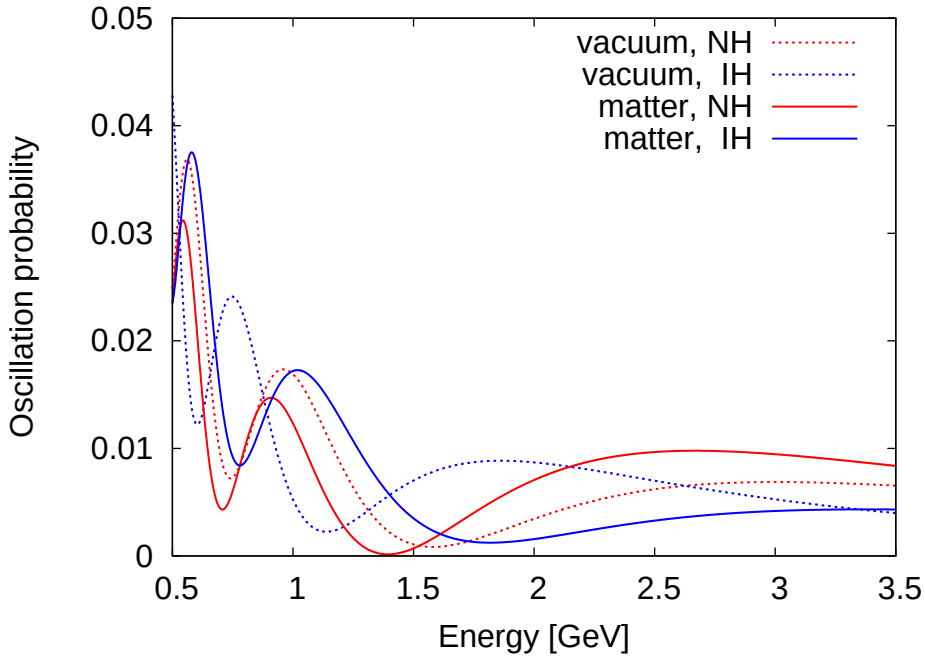

(b) $L = 1300$ km.

Figure 4.6.: Difference between normal hierarchy (NH) and inverted hierarchy (IH) for vacuum and matter oscillations at a baseline of a) 500 km and b) 1300 km. (The values of the other parameters are $\sin^2 2\theta_{12} = 0.3$, $\theta_{23} = \frac{\pi}{4}$, $\Delta m^2_{21} = 8.0 \times 10^{-5}$ eV^2, $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2 and $A = 1.4 \times 10^{-22}$ GeV).

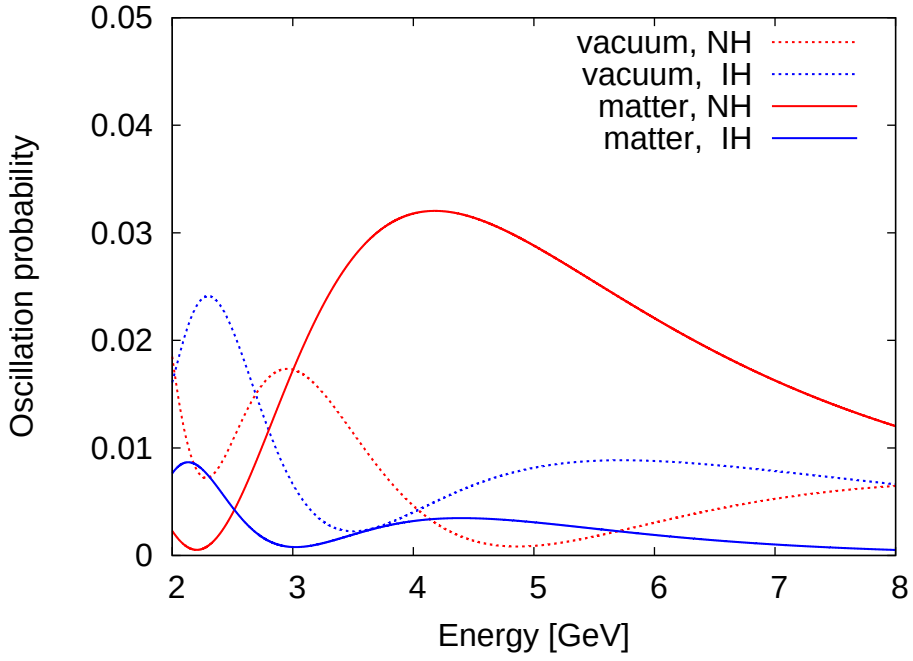

(a) $L = 4000$ km.

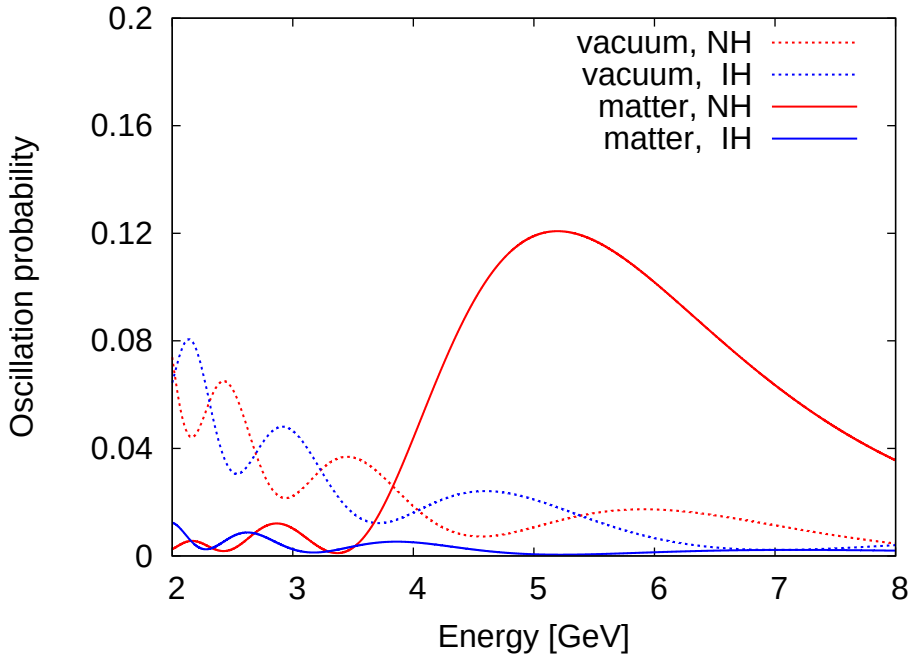

(b) $L = 8000$ km.

Figure 4.7.: Difference between normal hierarchy (NH) and inverted hierarchy (IH) for vacuum and matter oscillations at a baseline of a) 4000 km and b) 8000 km. (The values of the other parameters are $\sin^2 2\theta_{12} = 0.3$, $\theta_{23} = \frac{\pi}{4}$, $\Delta m^2_{21} = 8.0 \times 10^{-5}$ eV^2, $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2 and $A = 2.2 \times 10^{-22}$ GeV).

have different oscillation probabilities in matter because they interact differently with matter - the matter potential A is positive for neutrinos and negative for anti-neutrinos (see Section 2.4). The question is then how we know whether an observed asymmetry between neutrinos and anti-neutrinos is due to genuine CP violation, described by δ, or simply due to the CP asymmetry of the earth. This problem becomes worse for longer baselines when the matter effects are most significant. Unfortunately, as just shown above, long baselines are necessary to make a measurement of the mass hierarchy!

In Figs. 4.8 and 4.9 we compare the effects of true and 'fake' (matter-induced) CP violation for different baselines. The value of θ_{13} is taken to be 3°, a normal hierarchy is considered (as this is when matter effects are most prominent), and we show the difference between the neutrino and anti-neutrino spectra for $\delta = \frac{\pi}{2}$ in vacuum (maximal genuine CP violation, no matter effects) and $\delta = 0$ with matter effects (no genuine CP violation, only matter effects). We use the same baselines as earlier: 500 km, 1300 km, 4000km and 8000 km.

For $\delta = \frac{\pi}{2}$, it can be seen that matter effects (dotted lines) look very different to genuine CP violation (solid lines) - genuine CP violation and matter effects affect the neutrino and anti-neutrino spectra in very different ways. However, this is not necessarily true for all the other values of δ. At the shortest baseline of 500 km (Fig. 4.8a) we see that the effect of genuine CP violation is much greater than that due to matter effects; therefore there is no danger of the two effects being confused at this baseline. At a baseline of 1300 km (Fig. 4.8b), the effect of genuine CP violation is still greater than that from matter effects so that the two effects are still distinguishable. At 4000 km (Fig. 4.9a) the effects are roughly of the same magnitude, and at 8000 km (Fig. 4.9b) the matter effects are clearly dominant. So potentially, for these very long baselines, this could pose a problem. The solution used by the HENF is to use the *magic baseline* which will be described in Section 4.5.

4.4. Other oscillation channels

The other oscillation channels which may be accessible to a neutrino factory are listed below.

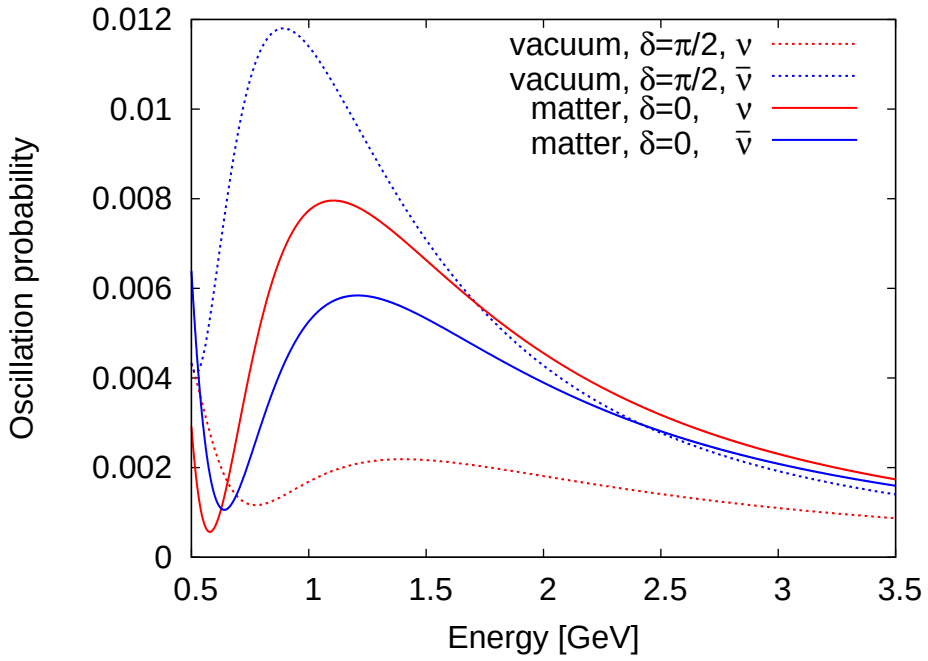

(a) $L = 500$ km.

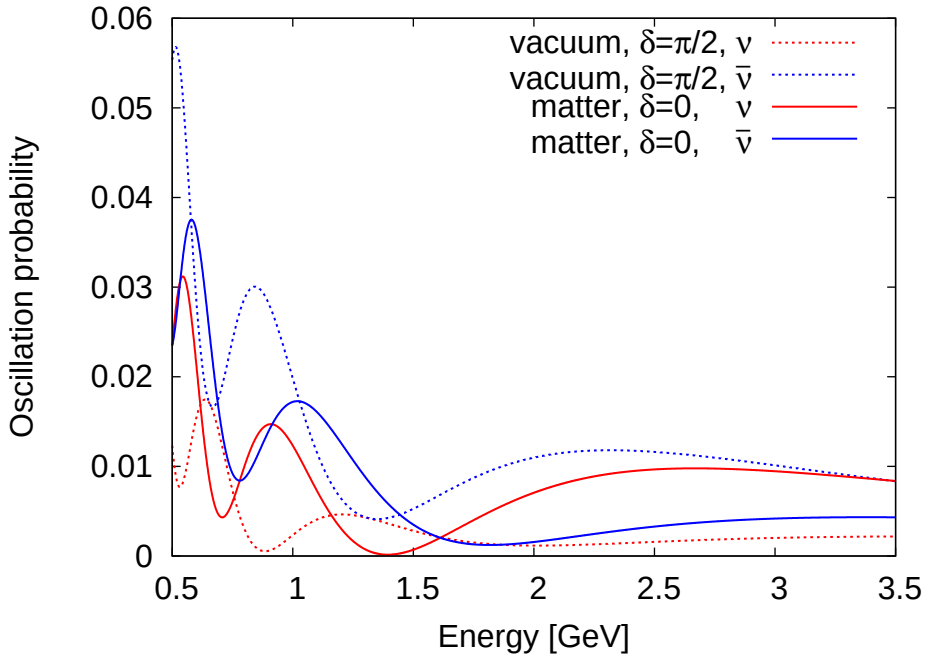

(b) $L = 1300$ km.

Figure 4.8.: Difference between genuine CP violation and fake CP violation from matter effects (with $A = 1.4 \times 10^{-22}$ GeV) at a baseline of a) 500 km and b) 1300 km, for $\theta_{13} = 3°$, $\delta = 0$ and $\frac{\pi}{2}$ and a normal hierarchy. (The values of the other parameters are $\sin^2 2\theta_{12} = 0.3$, $\theta_{23} = \frac{\pi}{4}$, $\Delta m^2_{21} = 8.0 \times 10^{-5}$ eV^2 and $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2).

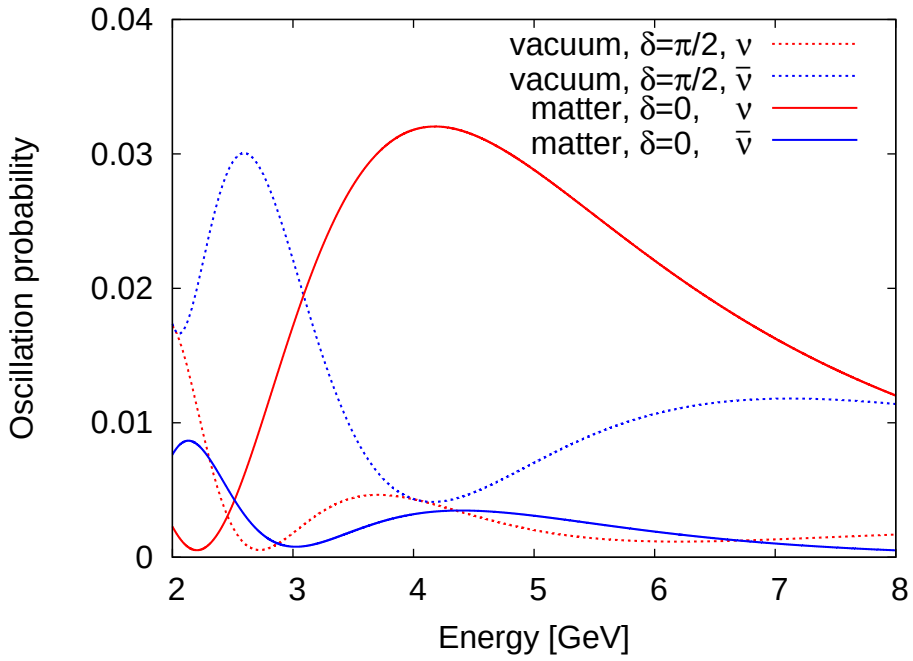

(a) $L = 4000$ km.

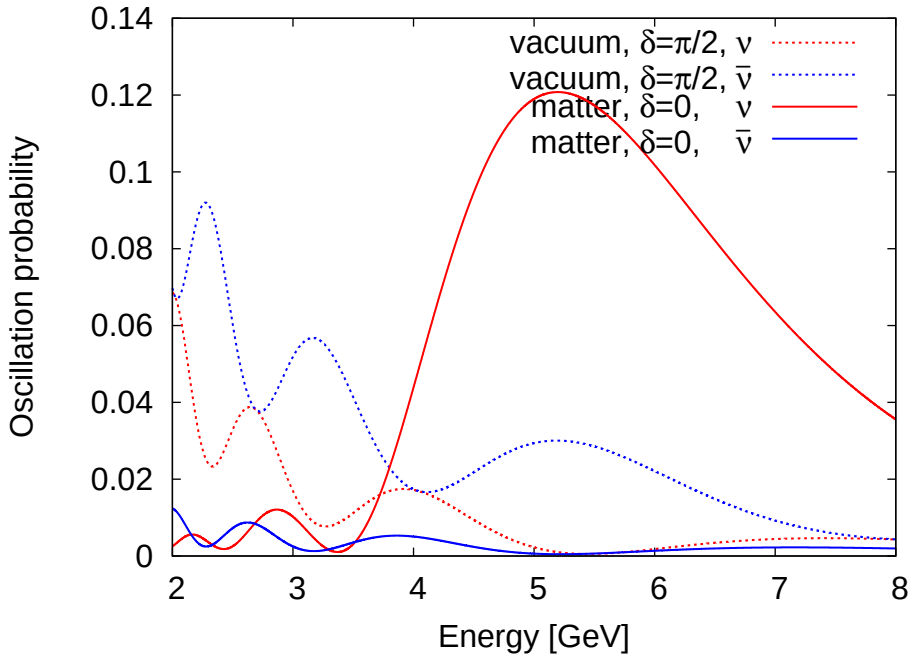

(b) $L = 8000$ km.

Figure 4.9.: Difference between genuine CP violation and fake CP violation from matter effects (with $A = 2.2 \times 10^{-22}$ GeV) at a baseline of a) 4000 km and b) 8000 km, for $\theta_{13} = 3°$, $\delta = 0$ and $\frac{\pi}{2}$ and a normal hierarchy. (The values of the other parameters are $\sin^2 2\theta_{12} = 0.3$, $\theta_{23} = \frac{\pi}{4}$, $\Delta m^2_{21} = 8.0 \times 10^{-5}$ eV^2 and $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2).

4.4.1. The ν_μ disappearance channel, $\nu_\mu \to \nu_\mu$

The probability for this channel is:

$$\begin{aligned} P_{\nu_\mu \to \nu_\mu} &\simeq 1 - s^2_{223} \sin^2 \left(\frac{\Delta_{21} L}{2} \right) + \frac{\Delta_{31} L}{2} c^2_{12} s^2_{223} \sin \left(\Delta_{31} L \right) \\ &+ O(\Delta_{21} s_{213}, \Delta^2_{21}, s^2_{213}). \end{aligned} \tag{4.5}$$

It is easy to gain very high statistics for this channel, as the probability is of order unity and an extremely high ν_μ flux is possible from a neutrino factory. This channel does not contain any information on new parameters, but it is from here that the most precise measurements of θ_{23} (but not the octant) and $|\Delta m^2_{31}|$ can be obtained. The limiting factors are the systematic errors associated with the neutrino fluxes and interaction cross-sections, which is *not* necessarily the case for *appearance channels*, such as the golden channel. The reason is that if, for instance, we run the neutrino factory in μ^- mode so that we produce ν_μ's and $\bar{\nu}_e$'s, then if we try to measure θ_{13} via the ν_μ *disappearance* channel, we are searching for a deficit of ν_μ's. This is going to be an effect which is at or below the percent level, which means that the *initial* flux of ν_μ needs to be known to a precision of greater than a percent for this to be a precise measurement. However, if we instead search for an *appearance* channel such as the golden channel ($\bar{\nu}_e \to \bar{\nu}_\mu$ if we run in μ^- mode), then we are not limited by our knowledge of the initial $\bar{\nu}_\mu$ content of the beam, as we know that this is precisely zero. Hence, if we detect a $\bar{\nu}_\mu$ then it is known with certainty (particle identification and other detector uncertainties aside) that this has arisen from the oscillation of a $\bar{\nu}_e$.

4.4.2. The ν_e disappearance channel, $\nu_e \to \nu_e$

If electron detection is possible, as in a TASD or LArTPC for example, it will be feasible to look for the ν_e and $\bar{\nu}_e$ disappearance channels, $\nu_e \to \nu_e$ and $\bar{\nu}_e \to \bar{\nu}_e$:

$$P_{\nu_e \to \nu_e} \simeq 1 - s^2_{213} \sin^2 \left(\frac{(\Delta_{31} - A) L}{2} \right). \tag{4.6}$$

This is the channel used in reactor experiments to obtain bounds on θ_{13}, but contains no information on either δ or the mass hierarchy. It cannot be used to make a *precision*

measurement of θ_{13}, as can the golden channel, because it is a disappearance channel and so again is limited by the systematic errors of the flux and cross-section measurements.

4.4.3. The ν_τ appearance channel, $\nu_\mu \to \nu_\tau$

For standard oscillations this channel adds nothing to that gained by a precision measurement of the ν_μ channel - the $\nu_\mu - \nu_\tau$ symmetry means that the probabilites for these two channels are essentially identical. To obtain a probability with ν_τ instead of ν_μ, $\nu_\mu \leftrightarrow \nu_\tau$, simply requires the interchanges $s_{23} \to c_{23}$ and $c_{23} \to -s_{23}$.

The fact that muons are very much easier to detect experimentally than taus is the reason that the ν_μ disappearance channel is used and not the ν_τ channel. In addition, for a LENF or a super-beam, the beam energy is not sufficiently high to enable the production of a significant number of taus. However, the true power of this channel lies in the detection of non-standard interactions, as mentioned in Section 2.5.

4.4.4. The platinum channel, $\nu_\mu \to \nu_e$

The platinum channel is the primary channel of a super-beam experiment. It is the T-conjugate of the golden channel, and so its probability is identical to Eq. (4.4) but with the exchange $\delta \to -\delta$. The phenomenology of this channel is therefore very similar to the golden channel, but we will illustrate the power of the synergy between these two channels in the next section.

4.5. Degeneracies

Neutrino oscillation experiments suffer from the problem of *degeneracies* [148, 149, 150, 151, 152, 153, 154, 155, 156, 157] - that is that the oscillation probability for a particular channel may be invariant under the transformation of one or more parameters, making it impossible to identify which are the true values. In the simplest cases, the degeneracy of a single parameter does not affect the measurement of any other parameters. For example, returning to the example of MINOS from Section 3.1, the disappearance probability is $P_{\mu\mu} \simeq 1 - \sin^2 2\theta_{23} \sin^2\left(\frac{\Delta_{31}L}{2}\right)$. This probability is invariant under the transformations $\theta_{23} \leftrightarrow \frac{\pi}{2} - \theta_{23}$ and $\Delta m^2_{31} \leftrightarrow -\Delta m^2_{31}$, either singly or simultaneously. Here, θ_{23} and Δm^2_{31}

are the *true* values of the parameters i.e. those which appear in nature, rather than the values which are produced by an experiment. Thus, from ν_μ disappearance experiments alone, we cannot identify the octant of θ_{23} if its value is not precisely $\frac{\pi}{4}$, and neither can we identify the mass hierarchy. However, the lack of knowledge about either of these parameters *does not affect a measurement of the other parameter.*

The situation becomes more complicated when considering the golden channel, from which we want to extract three unknown parameters - θ_{13}, δ and sign(Δm^2_{31}) from an oscillation probability which is dependent not only upon these parameters, but also on several others (Eq. (4.4)). We then need to consider to what extent the lack of precision on the 'known' parameters, for instance the size of the error bars and the lack of knowledge on the octant of θ_{23}, will affect our measurement of the unknown parameters. Additionally, in trying to determine several parameters from a single experiment, even a very precise measurement allows for different combinations of (θ_{13}, δ, sign(Δm^2_{31})) to be fitted to the data, severely weakening the sensitivity to these parameters. Many strategies have been advocated to resolve this problem which in general involve another detector [158, 159, 160, 161, 162, 163, 164], the combination with other experiments [134, 135, 165, 166, 167, 168, 169, 170, 171, 172, 173, 174, 175, 176, 177] and/ or the addition of new channels [129, 130, 178]. An algebraic treatment of degeneracies can be found in Refs. [127, 136] and in Appendix H.

The strategies mentioned above for resolving degeneracies all work on the basic principle that the data sets obtained from different channels and/ or different $\frac{L}{E}$ will have degenerate solutions which appear in *different* parts of the parameter space. So the fake solutions from one data set can be eliminated by the combination with another data set. Only the true solution appears in both data sets. This is illustrated in Fig. 4.10 which shows simulated data from two different setups fitted to simulated true values of (θ_{13}, δ, sign(Δm^2_{31})) = (1°, $\frac{\pi}{2}$, +1). Fig. 4.10a shows data from only a single channel - the golden channel, $\nu_e \to \nu_\mu$, whereas Fig. 4.10b shows data from both this channel and the platinum channel, $\nu_\mu \to \nu_e$. In both cases, the 68%, 90% and 95% contours are shown, with the solid lines corresponding to the regions obtained when the data are fitted to the correct (normal) hierarchy, and the dashed lines corresponding to a fit to the incorrect (inverted) hierarchy. In this particular case, it can be seen that if we have data from only the golden channel, then the effect of not knowing the hierarchy is to decrease the precision of the θ_{13} measurement. However, with two channels, the wrong-hierarchy solutions occur in different regions of the parameter space and thus can be eliminated, leaving only the true solution which is identified by both data sets. In this particular

example one can see that if we had instead started off with only the platinum channel, then for this particular value of θ_{13} and δ, the wrong-hierarchy solution is not displaced significantly from the true solution and so would not affect the measurement. However, there are other points in the parameter space for which the same problem occurs as for the golden channel; hence the synergy of using both channels together.

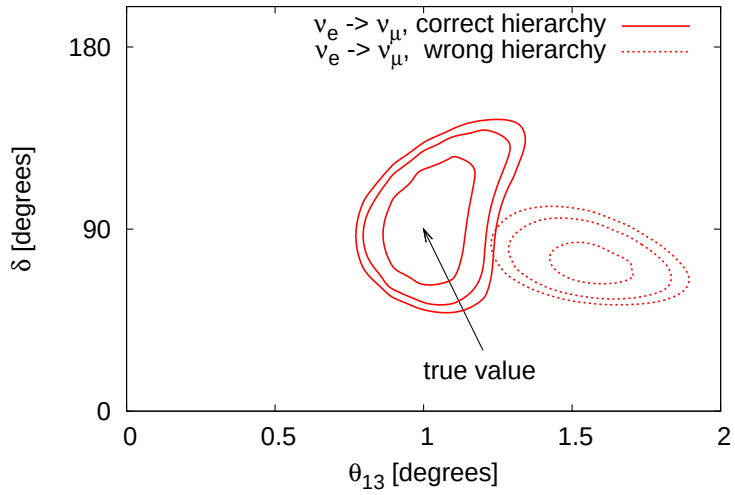

(a) Data from one channel ($\nu_e \to \nu_\mu$ only).

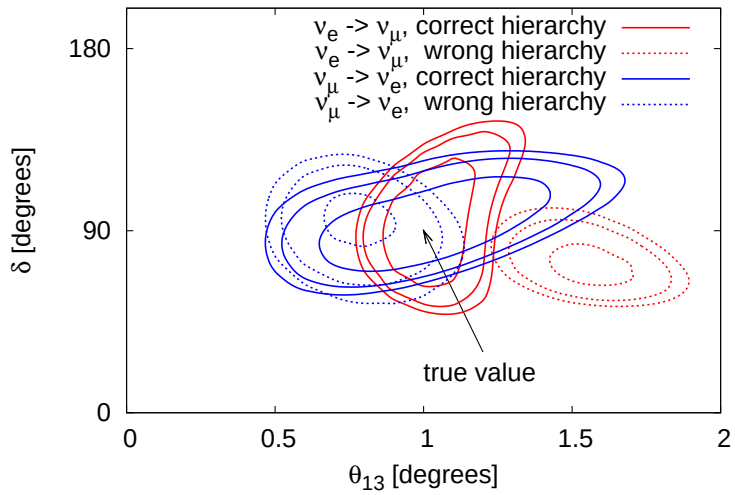

(b) Data from two channels ($\nu_e \to \nu_\mu$ and $\nu_\mu \to \nu_e$).

Figure 4.10.: Eliminating degenerate solutions by combining data from complementary channels.

In Chapters 6 and 7 we shall target the problem of degeneracies by considering detectors which may be able to detect the platinum channel. If we compare the forms of the probabilities, the golden channel, $\nu_e \to \nu_\mu$, is given by Eq. (4.4), the CP-conjugated golden channel, $\bar{\nu}_e \to \bar{\nu}_\mu$, is obtained by exchanging $\delta \to -\delta$ and $A \to -A$, the platinum channel, $\nu_\mu \to \nu_e$, is the T-conjugate of the golden channel and is obtained by exchanging $\delta \to -\delta$, and the CP-conjugated platinum channel, $\bar{\nu}_\mu \to \bar{\nu}_e$ (which is the CPT-conjugate of the golden channel), is obtained by exchanging $A \to -A$.

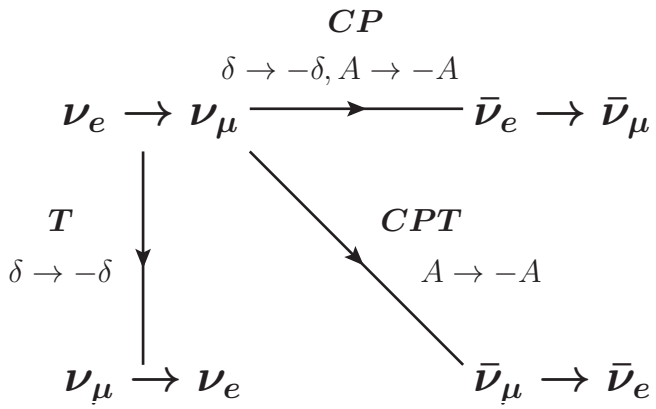

Now we can try to understand the degeneracy shown in Fig. 4.10. For the golden channel, we are interested only in those terms in Eq. (4.4) which have dependence upon θ_{13}, δ and $\text{sign}(\Delta m^2_{31})$. Therefore we shall neglect the solar term and look only at the atmospheric and CP terms, considering a value of $\delta = \frac{\pi}{2}$ as used in Fig. 4.10:

$$P_{\nu_e \to \nu_\mu} = s^2_{213} s^2_{23} \frac{\sin^2\left(\frac{(\Delta_{31}-A)L}{2}\right)}{\left(1-\frac{A}{\Delta_{31}}\right)^2} \tag{4.7a}$$

$$+ \; s_{213} c_{13} s_{212} s_{223} \frac{\Delta_{21}}{A} \sin\left(\frac{AL}{2}\right) \frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}} \cos\left(\frac{\Delta_{31}L}{2} - \frac{\pi}{2}\right). \tag{4.7b}$$

In order to understand how the $\text{sign}(\Delta m^2_{31})$ degeneracy arises, consider the effect of making the transformation $\Delta m^2_{31} \to -\Delta m^2_{31}$ on Equation 4.7. For both the atmospheric and CP terms, we have to look at the fraction

$$\frac{\sin^2\left(\frac{(\Delta_{31}-A)L}{2}\right)}{\left(1-\frac{A}{\Delta_{31}}\right)^2}. \tag{4.8}$$

Expanding out the numerator we find that this term can be written as

$$\sin^2\left(\frac{|\Delta_{31}|L}{2}\right)\cos^2\left(\frac{AL}{2}\right) + \cos^2\left(\frac{|\Delta_{31}|L}{2}\right)\sin^2\left(\frac{AL}{2}\right) \pm 2\sin\left(\frac{|\Delta_{31}|L}{2}\right)\cos\left(\frac{|\Delta_{31}|L}{2}\right)\sin\left(\frac{AL}{2}\right)\cos\left(\frac{AL}{2}\right), \tag{4.9}$$

where the upper + sign refers to a normal hierarchy, and the lower − sign to an inverted hierarchy. So the only difference is in the bottom line which, over the spectrum, will on average be equal for both hierarchies - therefore this particular term is not greatly sensitive to the hierarchy.

The denominator is $\left(1 \mp \frac{A}{|\Delta_{31}|}\right)^2$. The quantity $\frac{A}{|\Delta_{31}|}$ takes the value ~ 0.3 at the first oscillation peak at ~ 3 GeV. So if the data are fitted (incorrectly) to an inverted hierarchy, then the denominator is larger than it should truly be for a normal hierarchy and therefore the fraction (4.8) takes a smaller value than the true value. To compensate for this when fitting the data, a *larger* value of θ_{13} than the true value must be chosen, and hence in Fig. 4.10 we see that the wrong-hierarchy fit lies at a larger value of θ_{13} than the true value. This effect will be greater for the atmospheric term than the CP term, as the atmospheric term has a quadratic dependence on the fraction whereas the CP term has only a linear dependence. The hierarchy also affects the term $\cos\left(\frac{\Delta_{31}L}{2} - \frac{\pi}{2}\right)$, which is simply $\pm\sin\left(\frac{|\Delta_{31}|L}{2}\right)$. Around the first oscillation peak (and also the second), $\sin\left(\frac{|\Delta_{31}|L}{2}\right)$ is always positive, so in the case of a normal hierarchy the CP term takes the same sign as the atmospheric term and adds constructively, whereas for an inverted hierarchy, the CP term takes the opposite sign and so adds destructively. Therefore the same effect occurs here as for the fraction (4.8) and the probability appears to be smaller than the true value. So once again, to compensate, it is necessary to fit to a larger value of θ_{13} which quadratically increases the atmospheric term.

If we consider instead the platinum channel, $\nu_\mu \to \nu_e$, then the probability is identical except for the exchange $\delta \to -\delta$ in the CP term, which is then dependent on $\cos\left(\frac{\Delta_{31}L}{2} + \frac{\pi}{2}\right) = \mp\sin\left(\frac{|\Delta_{31}|L}{2}\right)$. The positive and minus signs are exchanged for normal and inverted hierarchies relative to the golden channel, so that in the case of a normal hierarchy, the CP term adds *destructively* to the atmospheric term whereas for an inverted hierarchy, the CP term adds *constructively*. So if we incorrectly fit to the IH then the probability looks larger than it really is, and this counteracts the effect of fitting the wrong hierarchy to the factor (4.8) which means that the measurement of θ_{13} is not significantly affected.

Thus we demonstrate how different channels are affected differently by parameter degeneracies, such that multiple channels can provide complementary information in different regions of parameter space, enabling degeneracies to be resolved. We shall return to this topic in Chapter 7 when we discuss degeneracies in the context of non-standard interactions.

The magic baseline

The ultimate weapon for resolving degeneracies in neutrino oscillation physics is the so-called 'magic baseline' [179] which forms part of the standard HENF setup. It corresponds to the baseline at which $\frac{AL}{2} = \pi$. For earth matter densities, the magic baseline is at around 7500 km assuming an average earth density of $\sim 4.3\text{g/cm}^3$. At this baseline, terms (4.4b) and (4.4c) in Eq. (4.4) are zero. The only non-zero term is the atmospheric term, and so at this baseline there are only two parameters on which the spectrum depends - θ_{13} and $\text{sign}(\Delta m^2_{31})$ - as there is no dependence on δ. Fitting only two parameters instead of three significantly reduces the problem of degeneracies, enabling a high-precision measurement of θ_{13} to be made. Additionally, as there are no CP violating effects, it must be the case that any observed asymmetry between neutrinos and anti-neutrinos is due purely to matter effects. Therefore at the magic baseline it is also possible to obtain a good estimate of these effects which is fed back into the measurement made at the second detector which measures the full oscillation probability, including the CP term.

Clearly, the magic baseline is a highly desirable feature to incorporate into an experiment, but there are severe technical challenges associated with the construction of a 7500 km baseline experiment!

Chapter 5.

Long-baseline experiments within Europe: the LAGUNA project

5.1. Introduction to LAGUNA

The Large Apparatus for Grand Unification and Neutrino Astrophysics (LAGUNA) project [108, 180, 181] is a European design study for the development of a kiloton-scale underground particle detector. The detector will be a multi-purpose facility with a broad physics reach - searching for proton decay, detecting astrophysical neutrinos, and it will also be used as part of a long-baseline neutrino oscillation experiment.

The LAGUNA study focuses on assessing the potential of three different detector technologies (listed below) and seven possible baselines (Table 5.1) within Europe. The neutrino beam is assumed to originate from CERN, Geneva, and the baseline will be determined by the position of the detector. For the purposes of minimising backgrounds, the detector must be situated underground and so the choice of locations is limited to locations where there is already an existing mine which could accommodate a large-scale detector, or to locations where it would be possible to expand an existing mine. The practical considerations of building a gigantic detector deep underground include engineering, construction, safety and transportation issues, all of which are considered in the design study. We shall be evaluating the physics potential of the project. The possible locations of the detectors, their distances from CERN and the energies of the first oscillation maximum (ignoring matter effects) are shown in Table 5.1.

The options for the detector are:

Location	Distance from CERN [km]	1st osc max [GeV]
Fréjus (France)	130	0.26
Canfranc (Spain)	630	1.27
Umbria (Italy)	665	1.34
Sierozsowice (Poland)	950	1.92
Boulby (UK)	1050	2.12
Slanic (Romania)	1570	3.18
Pyhäsalmi (Finland)	2300	4.65

Table 5.1.: The seven potential sites being studied by the LAGUNA design study. From Ref. [181].

- LENA (a multipurpose detector for Low Energy Neutrino Astronomy and proton decay) [84] - 50 kton liquid scintillator detector.
- MEMPHYS (MEgaton Mass PHYSics) [182] - 440 kton water Čerenkov detector.
- GLACIER (Giant Liquid Argon Charge Imaging ExpeRiment) [183] - 100 kton liquid argon detector.

The physics reach will be determined primarily by the length of the baseline as has been discussed in the previous chapters. We also need to consider the potential of each of the three possible detectors when used with each of the baselines - as mentioned in Section 3.3, different detectors perform optimally for different particle energies and types. Also, for practical reasons, not all the sites would be able to accommodate all of the detectors - for instance, it is not possible to build a water Čerenkov detector in a salt mine as salt is water-soluble!

The final component of the experiment is the beam. Several β-beam configurations have been studied, spanning a large range of γ-factors [114, 115, 116]. The higher the γ-factor, the higher the neutrino energy, and so a high γ is necessary if one of the longer baselines is to be used. However, recent developments in the accelerator schedule at CERN have since limited these options to only a low-γ β-beam as studied in Ref. [114] which would be appropriate for the CERN-Fréjus baseline of 130 km, but none of the other baselines. The best option at present therefore seems to be a super-beam, as studied in Ref. [181]. This is a realistic possibility for which the technology is already well established, and can be applied within the intended time period - the intention is for LAGUNA to begin running in roughly a decade. This is in contrast to a neutrino

factory, which is the topic of the next two chapters, which is a far-future option intended to follow after experiments such as LAGUNA.

5.2. Baseline studies

The baseline-dependence of oscillation physics has already been covered in Chapter 4. In essence, for the seven sites which we are considering, the longer the baseline the better, as this makes it more likely that a measurement of the mass hierarchy will be successful. None of the baselines is sufficiently long for matter effects to be significantly confused with genuine CP violation, as for the HENF. However, it may be possible that a shorter baseline combined with a particular detector can give similar or better results to a longer baseline, which would mean that the energy requirements of the beam are less aggressive and would also provide more flexibility for choosing the detector site.

5.3. The LENA detector for use with a super-beam

Water Čerenkov detectors are already very well-understood following the success of detectors such as Super-Kamiokande [44]. The development of giant liquid argon detectors is in progress, as mentioned in Section 3.3 [92, 93]. However, giant liquid scintillator detectors have been less well studied and so it is useful to perform some preliminary studies to gauge the potential and limitations of such a detector.

The LENA (Low Energy Neutrino Astronomy) detector [108] is a proposed large volume liquid scintillator detector, primarily for the observation of proton decay and low energy neutrinos from natural sources (the earth and astrophysical sources). The detector has the benefit of having excellent particle identification and energy resolution. Here we will study how this detector performs in a neutrino oscillation experiment.

5.3.1. The beam

The use of the LENA detector in a β-beam experiment has been studied in Ref. [184], as part of a LENF in Ref. [185], and in a wide-band super-beam in Ref. [186]. We consider a standard super-beam, using a simulated beam optimised for the NOvA experiment [88]. NOvA has a baseline of 812 km so we will use a similar baseline of 950 km - CERN to

Sierozsowice. The beam composition is predominantly ν_μ ($\bar{\nu}_\mu$) with $\sim 1\%$ contamination from ν_e ($\bar{\nu}_e$) and an energy peak at ~ 2 GeV. The source power is 10^{21} protons on target per year (1.12 MW target power). We consider 5 years' running time in neutrino mode, and 5 years in anti-neutrino mode. Although the beam configuration will obviously affect the sensitivity of the experiment, and the NOvA beam is not precisely optimal for our setup, the purpose of this study is to analyse the performance of the detector.

5.3.2. Detector design

The factors we consider with respect to the detector design are listed below. We set all these parameters to their 'reference' values (also shown below) and alter each one in turn to observe its effect on the experimental sensitivity. We assume that the detector is capable of measuring electrons and muons, enabling us to observe the $\nu_\mu \to \nu_\mu$ and $\overline{\nu}_\mu \to \overline{\nu}_\mu$ channels as well as the $\nu_\mu \to \nu_e$ and $\bar{\nu}_\mu \to \overline{\nu}_e$ channels. As mentioned in Section 3.5.1, the main disadvantage of a super-beam is the intrinsic contamination from ν_e ($\bar{\nu}_e$) which acts as a background to the ν_e ($\bar{\nu}_e$) appearance channels i.e. it is not possible to distinguish an electron that originates from an intrinsic ν_e from the beam (background event) from a ν_μ which has oscillated into a ν_e (signal event).

- **Energy resolution**: reference value - 5% of the particle's true energy, for all particles [187, 188].
- **Energy threshold**: reference value - 1.0 GeV for all particles.
- **Fiducial mass**: reference value - 50 kton [180].
- **Detection efficiency**: reference value - 90% for all particles [188].
- **Background level**: this includes misidentification of electrons and muons (almost negligible for LENA [188] but we include, conservatively, a misidentification rate of 10^{-4}). However, the dominant factor for a super-beam is the intrinsic ν_e contamination. We will use a reference background of 45% of the intrinsic ν_e content of the beam, which is half that of the signal. The same background estimate (half the signal efficiency) is used by NOvA, which will have a near detector to measure the neutrino fluxes prior to oscillation, thereby giving an estimate of the intrinsic ν_e ($\bar{\nu}_e$) spectrum. In addition, cross-section and fixed-target experiments currently being conducted will help to improve the estimates on these predictions in the near future.

- **Systematic errors**: reference value - 5%, uncorrelated (see Section 3.6.1).

5.3.3. Simulations and results

We have used the GLoBES software package (Section 3.6.1) to simulate the experiment and use the same true values of the oscillation parameters as those in Ref. [189]: $\sin^2\theta_{12} = 0.3$, $\theta_{23} = \frac{\pi}{4}$, $\Delta m^2_{21} = 8.0 \times 10^{-5}$ eV^2, and $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2 with a 10% uncertainty on the atmospheric parameters, 4% uncertainty on the solar parameters, and 2% uncertainty on the matter density (these 'uncertainty' values are the ones which we use as input errors for each of the parameters). These are not the most recent best-fit values but are the values used by other groups to perform studies with which we shall be comparing. We use values of $\theta_{13} = 2^\circ$ and 5° ($\sin^2 2\theta_{13} \sim 5 \times 10^{-3}$ and $\sim 10^{-1}$), and $\delta = 0^\circ$, $\pm 90^\circ$ and $\pm 180^\circ$. We use $\theta_{13} = 2^\circ$ rather than 1° as we will in Chapters 6 and 7, as this setup does not have as powerful a reach as the neutrino factory. We show the 68%, 90% and 95% confidence level regions obtained in the $\theta_{13} - \delta$ plane from each of the detector configurations, and try to fit the simulated data to both the correct true hierarchy (normal) and the wrong (inverted) hierarchy. The fits to the correct hierarchy are shown by the red solid lines, and those to the wrong hierarchy by the dashed blue lines. If no fits to the wrong hierarchy are shown, this indicates that the mass hierarchy can be correctly identified at the 95% confidence level by that particular configuration.

Energy resolution

We simulate values of $\frac{\sigma(E)}{E} = 5\%$ and 1% for the energy resolution (more details on the energy resolution will be given in Section 6.2.2) and show the results in Fig. 5.1. Note that Figs. 5.1a and Figs. 5.1b correspond to the results obtained when all the parameters are set to their reference values, so these are the results of the reference setup. Current estimates indicate that a resolution better than 5% should be feasible [187]. We find that the results obtained with an energy resolution of 10% are very similar to those with 5%, but improving to 1% (bottom row) does have some benefit, most notably for large θ_{13} (5°). Therefore, although it is beneficial to optimise the energy resolution as much as possible, it is not a crucial factor and a value of $\sim 10\%$ is adequate.

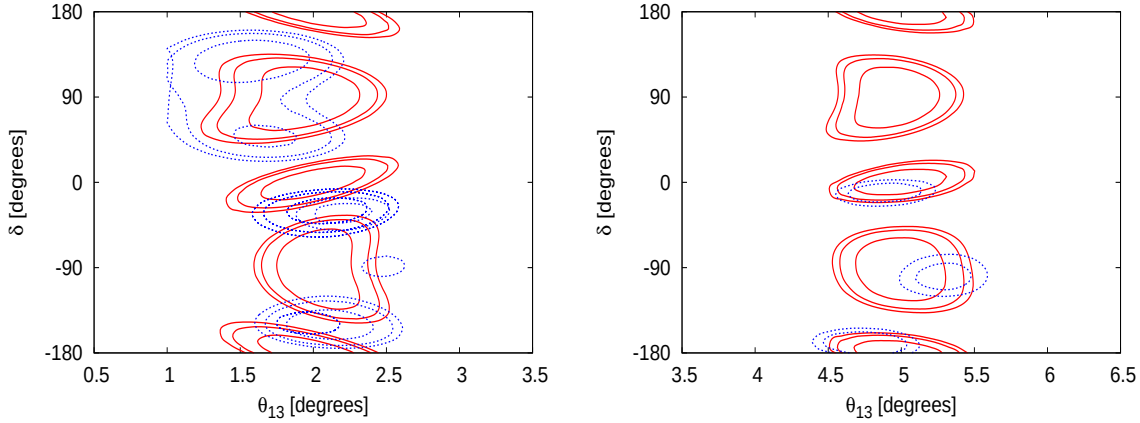

(a) $\theta_{13} = 2°$, $\frac{\sigma(E)}{E} = 5\%$.

(b) $\theta_{13} = 5°$, $\frac{\sigma(E)}{E} = 5\%$.

(c) $\theta_{13} = 2°$, $\frac{\sigma(E)}{E} = 1\%$.

(d) $\theta_{13} = 5°$, $\frac{\sigma(E)}{E} = 1\%$.

Figure 5.1.: $\theta_{13} - \delta$ precision with an energy resolution of $\frac{\sigma(E)}{E} = 5\%$ (top row) and 1% (bottom row), and $\theta_{13} = 2°$ (left column) and 5° (right column). The 68%, 90% and 95% contours are shown for $\delta = 0°$, $\pm 90°$ and $\pm 180°$.

Energy threshold

We use values of 0.5 GeV and 1.5 GeV for the detection threshold. The results for a 1.0 GeV threshold are shown in Figs. 5.2a and 5.2b. We assume the detector efficiencies to be energy independent, such that the efficiency is the same (90%) for all energies above the threshold. The first oscillation maximum at a baseline of 950 km is at $\sim$ 1.9 GeV and so there is little benefit in reducing the threshold significantly below this value. In particular, for small θ_{13}, a very low energy threshold has very little benefit. But it is essential to try to detect the entire oscillation peak - the results for a threshold of 1.5 GeV (bottom row) are worse than those for the reference value of 1.0 Gev (Fig. 5.2). The results are significantly worse if the energy threshold is 2.0 GeV (not shown).

Statistics - detector mass, detection efficiency, running time

To assess the impact of statistics on this setup we have simulated a detector with a mass of 100 kton (twice that of the reference setup) and 150 kton (Fig. 5.3). Roughly, this is equivalent to doubling and trebling the beam flux or the running time. Increasing the detector efficiency will also increase the statistics, although with an anticipated detection efficiency of around 90%, increasing or decreasing the efficiency by $\sim$ 10% has little effect. It is clear to see that roughly doubling the statistics (100 kton instead of 50 kton) will significantly improve the experimental performance, for all values of θ_{13} - compare Figs. 5.3a and 5.3b with Figs. 5.1a and 5.1b. The improvement gained by increasing to 150 kton from 100 kton is less pronounced

Background levels

The dominant background comes from the intrinsic ν_e ($\bar{\nu}_e$) content of the beam. We also include backgrounds from other channels: for the ν_μ ($\bar{\nu}_\mu$) disappearance channels we include 10^{-4} of the neutral-current ν_μ ($\bar{\nu}_\mu$) rates, and for the ν_e ($\bar{\nu}_e$) appearance channels we include 10^{-4} of the charged-current and neutral-current ν_μ ($\bar{\nu}_\mu$) rates i.e. we assume a particle misidentification rate of 10^{-4}. For the ν_e ($\bar{\nu}_e$) background we study values of 90% and 10% of these events (45% is the reference value). This is a very large range, but we see from Fig. 5.4 that these values appear to be irrelevant for the case of large θ_{13} (5°). However, for small θ_{13}, the backgrounds have a large effect as they significantly reduce the signal to background ratio

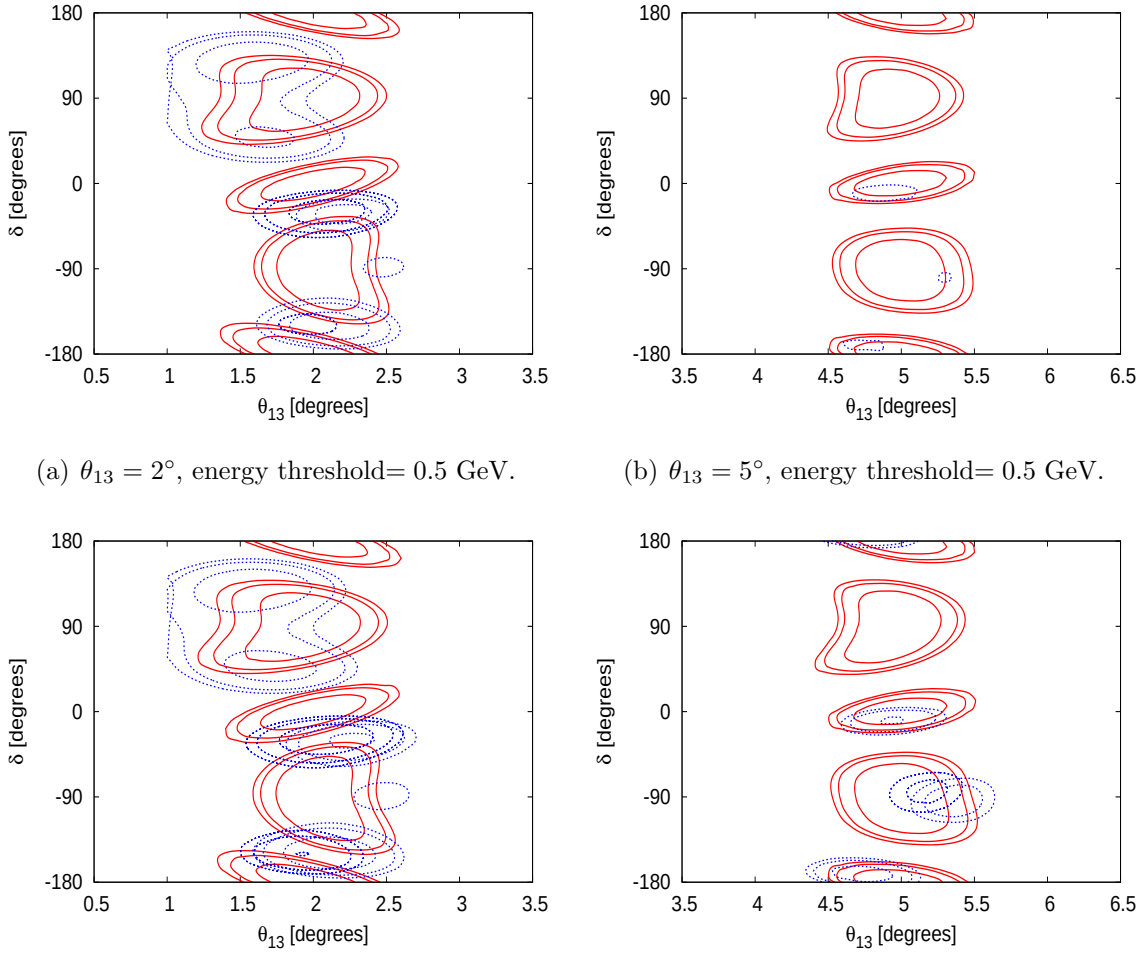

(a) $\theta_{13} = 2°$, energy threshold= 0.5 GeV.

(b) $\theta_{13} = 5°$, energy threshold= 0.5 GeV.

(c) $\theta_{13} = 2°$, energy threshold= 1.5 GeV.

(d) $\theta_{13} = 5°$, energy threshold= 1.5 GeV.

Figure 5.2.: $\theta_{13} - \delta$ precision with an energy threshold of 0.5 GeV (top row) and 1.5 GeV (bottom row), and $\theta_{13} = 2°$ (left column) and 5° (right column). The 68%, 90% and 95% contours are shown for $\delta = 0°$, $\pm 90°$ and $\pm 180°$.

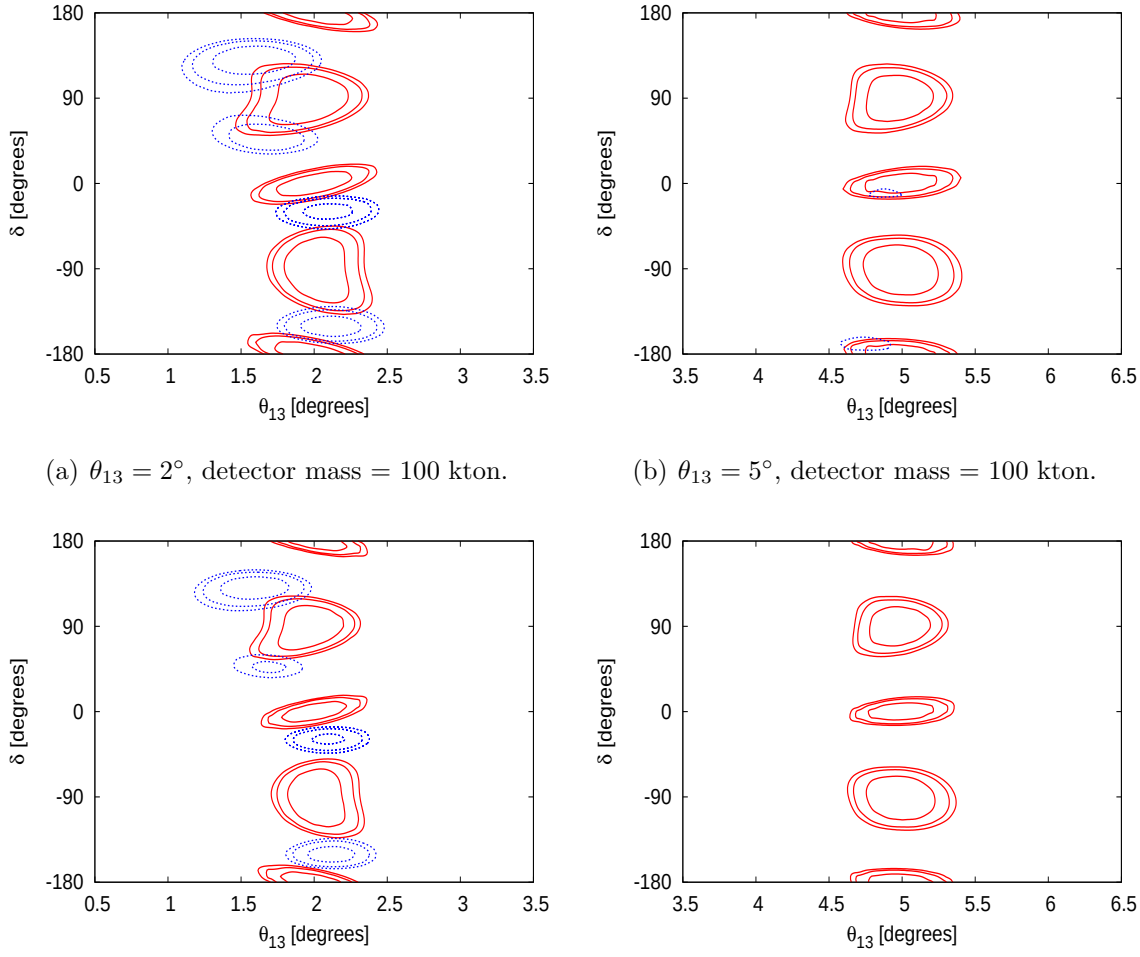

(a) $\theta_{13} = 2°$, detector mass = 100 kton.

(b) $\theta_{13} = 5°$, detector mass = 100 kton.

(c) $\theta_{13} = 2°$, detector mass = 150 kton.

(d) $\theta_{13} = 5°$, detector mass = 150 kton.

Figure 5.3.: $\theta_{13}-\delta$ precision with a detector mass of 100 kton (top row) and 150 kton (bottom row), and $\theta_{13} = 2°$ (left column) and 5° (right column). The 68%, 90% and 95% contours are shown for $\delta = 0°$, $\pm 90°$ and $\pm 180°$.

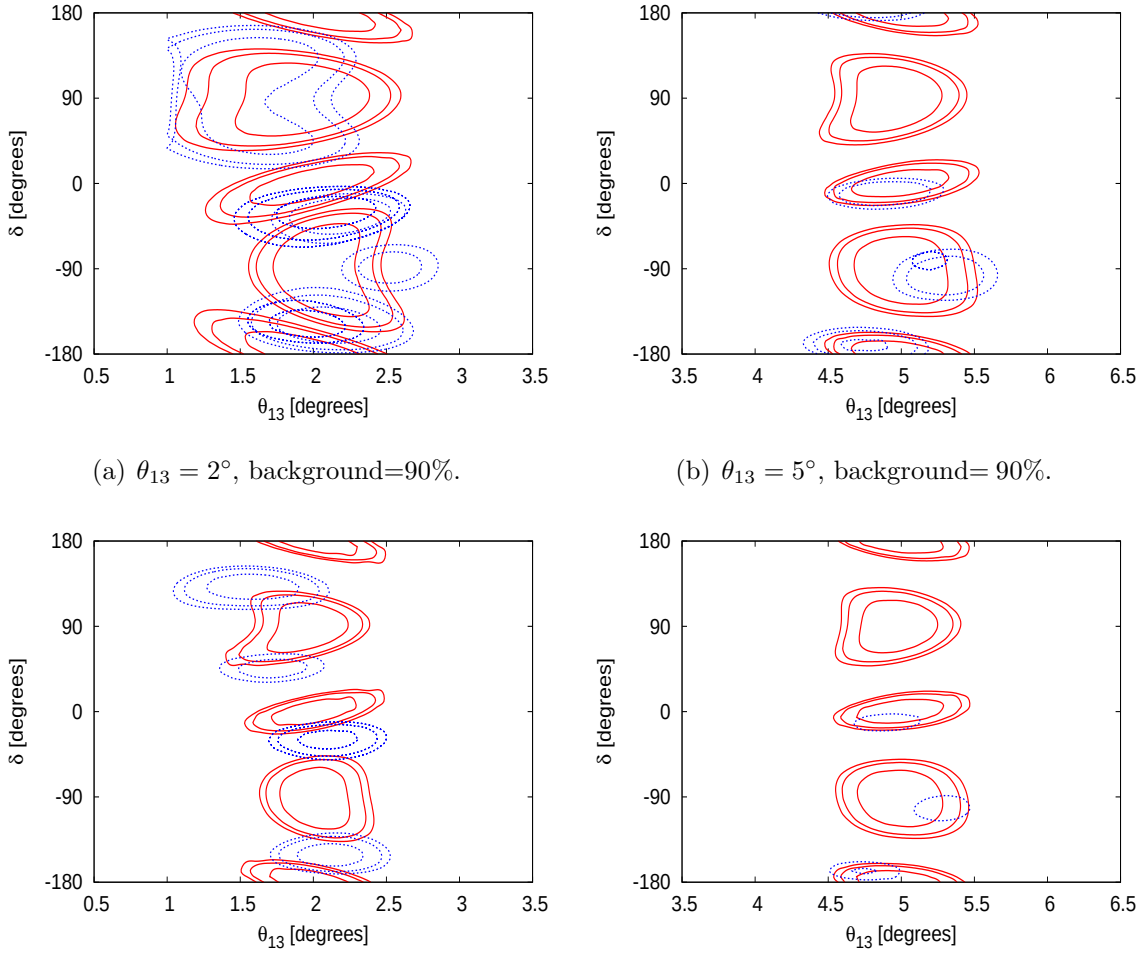

(a) $\theta_{13} = 2°$, background=90%.

(b) $\theta_{13} = 5°$, background= 90%.

(c) $\theta_{13} = 2°$, background= 10%.

(d) $\theta_{13} = 5°$, background= 10%.

Figure 5.4.: $\theta_{13} - \delta$ precision with a background level of 90% (top row) and 10% (bottom row), and $\theta_{13} = 2°$ (left column) and 5° (right column). The 68%, 90% and 95% contours are shown for $\delta = 0°$, $\pm 90°$ and $\pm 180°$.

Systematic errors

We simulate systematic errors of 10% and 2% (5% is the reference value) showing the results in Fig. 5.5. We find that systematic errors have a large effect on the sensitivity, for all values of θ_{13}, and thus that considerable effort should be made to minimise them.

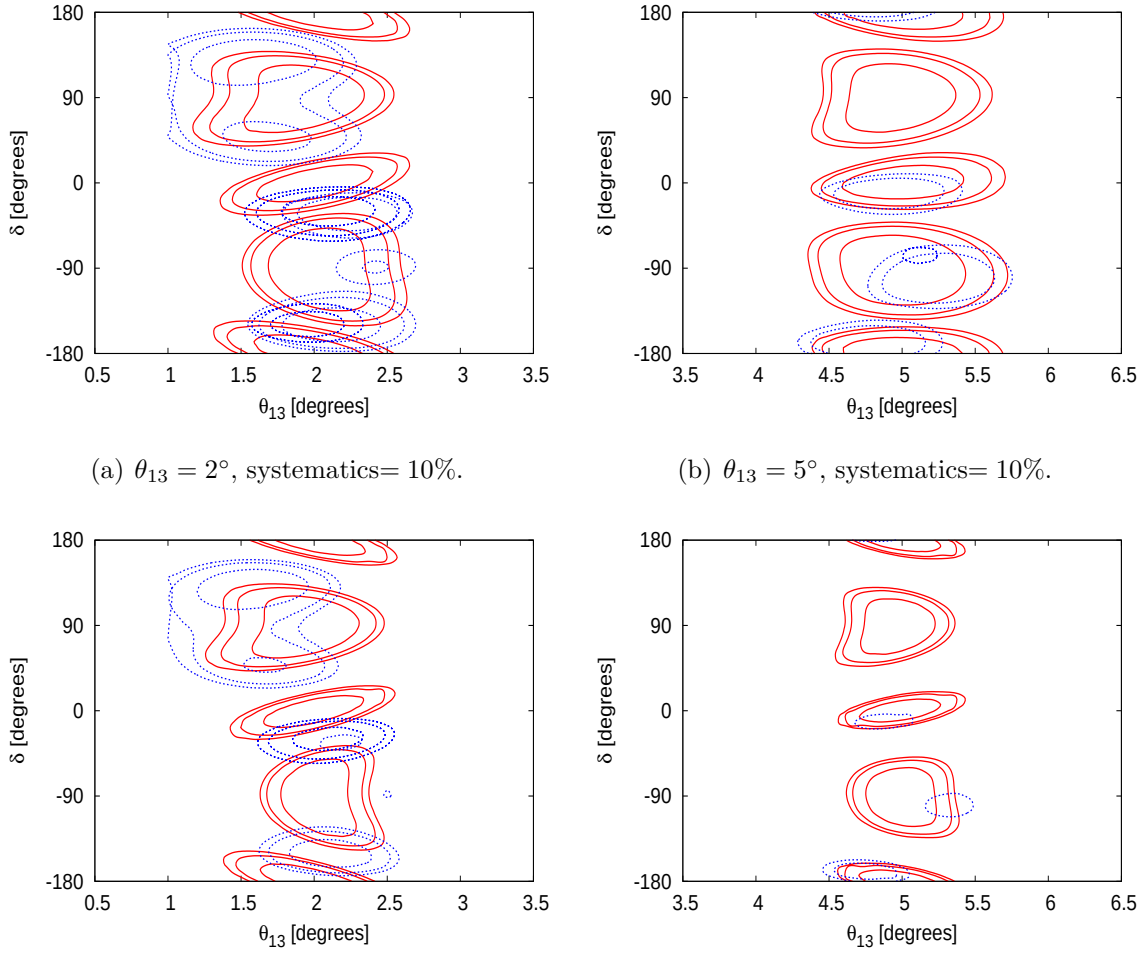

(a) $\theta_{13} = 2°$, systematics= 10%.

(b) $\theta_{13} = 5°$, systematics= 10%.

(c) $\theta_{13} = 2°$, systematics= 2%.

(d) $\theta_{13} = 5°$, systematics= 2%.

Figure 5.5.: $\theta_{13} - \delta$ precision with a systematic error of 10% (top row) and 2% (bottom row), and $\theta_{13} = 2°$ (left column) and 5° (right column). The 68%, 90% and 95% contours are shown for $\delta = 0°$, $\pm 90°$ and $\pm 180°$.

5.4. Super-beam with a 2285 km baseline

We expect that the CERN-Pyhäsalmi baseline of 2285 km will give the best physics reach; therefore we will assess how well a super-beam with this baseline can perform

and compare it to the results of other similar experiments. Again, we take the true values of the oscillation parameters to be those given in Ref. [189].

5.4.1. Beam and baseline setup

We use a simulated beam flux from A. Longhin [190] optimised for a baseline of 2285 km - the peak is at $\sim$ 4.6 GeV which is the energy of the first oscillation peak. The spectra of the ν and $\bar{\nu}$ beams are shown in Fig. 5.6, which shows that in addition to the main ν_μ ($\bar{\nu}_\mu$) content, the beam also contains $\sim 1\%$ contamination from ν_e ($\bar{\nu}_e$) and $\sim 10\%$ contamination from $\bar{\nu}_\mu$ and $\bar{\nu}_e$ (ν_μ and ν_e). The fluxes correspond to the CERN high-power PS2 (HP-PS2) configuration [181]: 50 GeV protons with 3×10^{21} protons on target (PoT) per year. 2 years' of ν running and 8 years' of $\bar{\nu}$ running is assumed; the running time is asymmetric because the flux of $\bar{\nu}$'s is much less that the flux of ν's so a time of $2 + 8$ years gives an approximately equal number of ν's and $\bar{\nu}$'s.

5.4.2. 100 kton liquid argon detector

In this study we primarily consider the performance of a 100 kton liquid argon time-projection chamber (Section 3.3.3). We perform some very detailed simulations using the most recent information available. We simulate the following features:

- Energy resolution: we use the migration matrices from L. Esposito and A. Rubbia [191]. These describe the energy-dependent ability of the detector to reconstruct the energy of the incident particle.
- Efficiency: we assume 80% efficiency for all particles and all interaction types, based on information from B. Fleming reported in Ref. [192]; this is also the value used in Ref. [107].
- Detector backgrounds: primarily particle misidentification and neutral-current events being mistaken for charged-current events. We assume the NC background to be 0.5% of all NC channels, based on the estimates from Ref. [107] and assume that the particle misidentification rate is negligible.
- Intrinsic ν_e ($\bar{\nu}_e$) beam background: the rejection of this background stems purely from the ability to predict or measure the intrinsic ν_e ($\bar{\nu}_e$) component of the beam. Theoretical predictions can be made based on measurements from fixed-target ex-

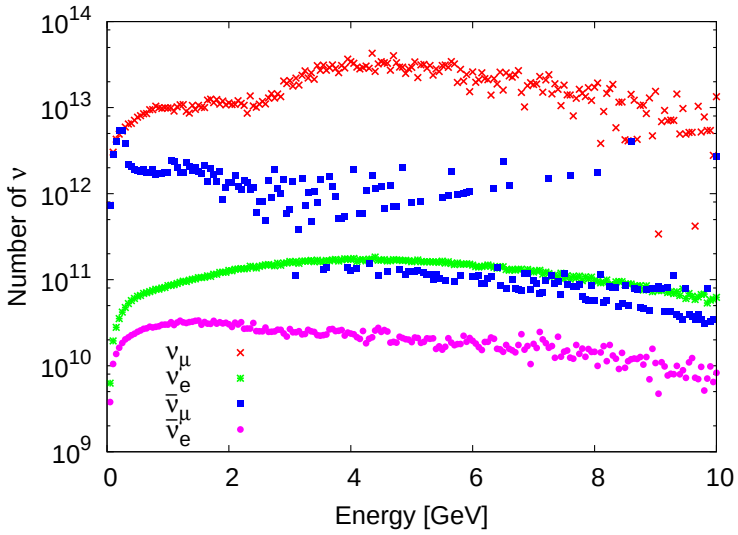

(a) ν beam spectrum.

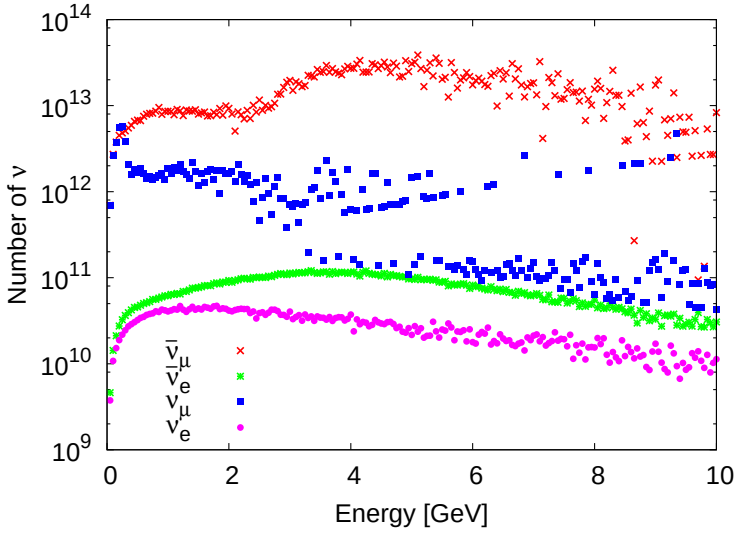

(b) $\bar{\nu}$ beam spectrum.

Figure 5.6.: Neutrino content of the a) ν_μ beam and b) $\bar{\nu}_\mu$ beam.

periments and, ideally, from the use of a near-detector which measures the unoscillated beam spectrum. NOvA [88] and T2K [103] estimate that they can reduce this component of their background to $\sim 50\%$ of the total ν_e ($\bar{\nu}_e$) content using these methods; we will also assume this value in most of our simulations.

- Intrinsic $\bar{\nu}$ (ν) background in ν ($\bar{\nu}$) beam: this forms $\sim 10\%$ of the beam as can be seen from Fig. 5.6. As the detector is not magnetised there is no way of distinguishing between positively and negatively charged leptons, and therefore it is not possible to distinguish between ν and $\bar{\nu}$. In the case that CP is conserved, this is not relevant as ν and $\bar{\nu}$ will behave the same way. However, since one of the primary aims of the experiment is to detect the presence of CP violation, it is crucial to distinguish between ν and $\bar{\nu}$ and so this background needs to be minimised. We estimate that $\sim 5\%$ should be realistic (half of the total content, similar to the way we estimate that the $\bar{\nu}_e$ (ν_e) beam background will be half of the total $\bar{\nu}_e$ (ν_e) content), given current knowledge of neutrino beam fluxes and planned future experiments.

- Systematic errors: these apply separately to errors on the signal and background events. We estimate 5% on both signal and background, uncorrelated, as in Ref. [193].

- Uncertainty in matter density: based on Ref. [194] we use a 2% uncertainty on the matter profile of the baseline.

Quasi-elastic events and τ detection

Simulations of liquid argon detectors distinguish between the detection of quasi-elastic (QE) and non-quasi-elastic (nQE) events. The typical energy of a QE event is $\lesssim 1.5$ GeV although there are still a few QE events at higher energies. We are studying a baseline of 2285 km for which the oscillation maximum is at ~ 4.6 GeV, and so we expect that a beam optimised for this baseline will have relatively few events in the QE region. The information from low-energy events contributes mainly to the sensitivity to CP violation, and so the QE events will be most valuable for CP discovery.

We also consider the possibility of being able to detect the τ appearance channels, $\nu_\mu \to \nu_\tau$ and $\bar{\nu}_\mu \to \bar{\nu}_\tau$. τ detection is experimentally very challenging (Section 3.3.4) and is therefore only of benefit if the additional events produce a significant improvement to the performance. Our expectation is that because the peak energy of the beam is

only slightly above the τ detection threshold of 3.5 GeV, only a very small amount of τ production is kinematically feasible which, conbined with the difficulties of τ detection, will not contribute significantly. We have simulated a configuration including the $\nu_\mu \to \nu_\tau$ and $\bar{\nu}_\mu \to \bar{\nu}_\tau$ channels, using an optimistic efficiency of 50% for τ detection and assuming that the background is 10^{-3} of the $\nu_\mu \to \nu_\mu$ channel rate.

The effects of the QE events, and of including τ detection, are shown in Fig. 5.7. We show how the omission of QE events (green lines) and the addition of τ detection (blue dotted lines) affects the sensitivity of the experiment. The omission of QE events affects only the CP discovery, as expected. Therefore we should consider how the sensitivity to CP violation changes if the efficiency and background for the QE events is different from the nQE events.

As expected, the addition of the τ channels has a negligible effect on all sensitivities.

Systematic errors

We find that the magnitude of the systematic errors is crucial to the performance of the setup. In Fig. 5.8 we show how a pessimistic estimate of a 20% systematic error on the backgrounds (with 5% error on the signal) gives much worse results than a 5% error on both signal and background. Significant improvement is obtained if the errors can be further reduced to 2% (on both signal and background). Also we show how changing the detection efficiency affects the results - the blue lines are for 100% efficiency (rather than 80%) and 20% systematics - from which we find that the effect is significant, especially for CP discovery.

5.4.3. Comparison of liquid argon detector with other detectors and experiments

We will now compare the performance of the LAGUNA experiment with a liquid argon detector to that of the other LAGUNA detectors.

For the MEMPHYS detector (440 kton water Čerenkov), from Ref. [107] we take the efficiency of all channels to be 40% and the beam ν_e ($\bar{\nu}_e$) background to be 40%. We assume that the background from the $\bar{\nu}$ (ν) component of the ν ($\bar{\nu}$) beam is $\sim 5\%$. The energy resolution is taken from Ref. [195] (page 83): $\sigma(E) = 0.017 + 0.007\sqrt{E_\mu}$ for muons, and $\sigma(E) = 0.006 + 0.026\sqrt{E_e}$, which we assume (naively) to be true at all

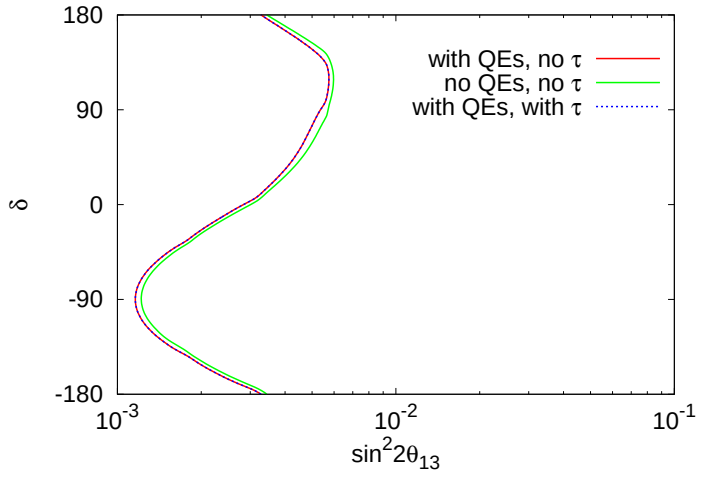

(a) θ_{13} discovery potential.

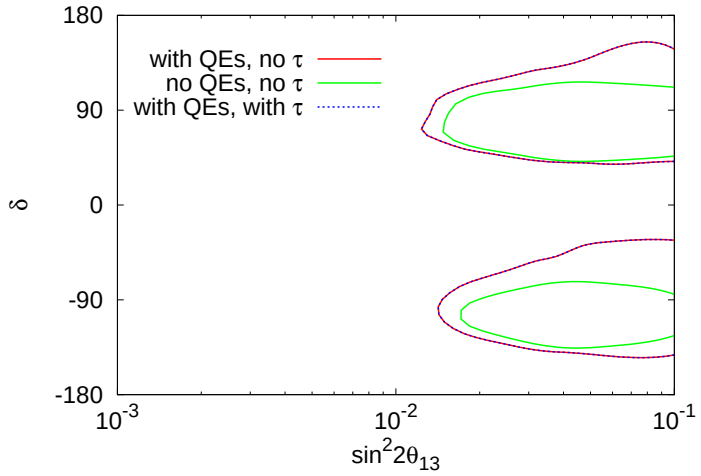

(b) CP discovery potential.

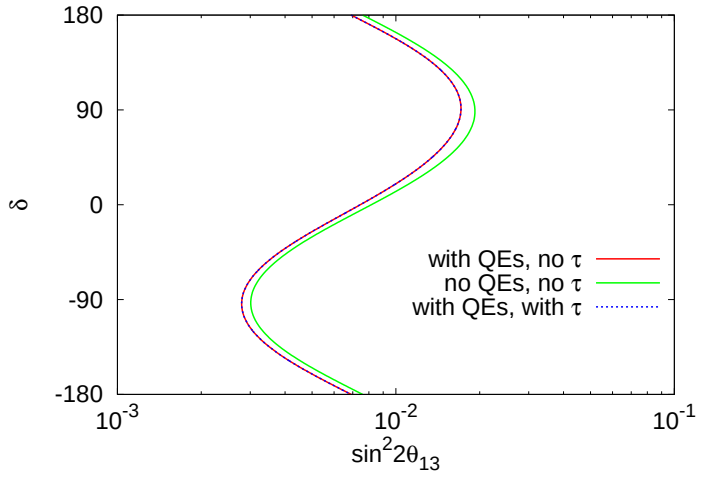

(c) Hierarchy sensitivity.

Figure 5.7.: Configurations with QE events and no τ detection, no QE events and no τ detection, and QE events and τ detection for a) θ_{13} discovery potential, b) CP discovery potential and c) hierarchy sensitivity.

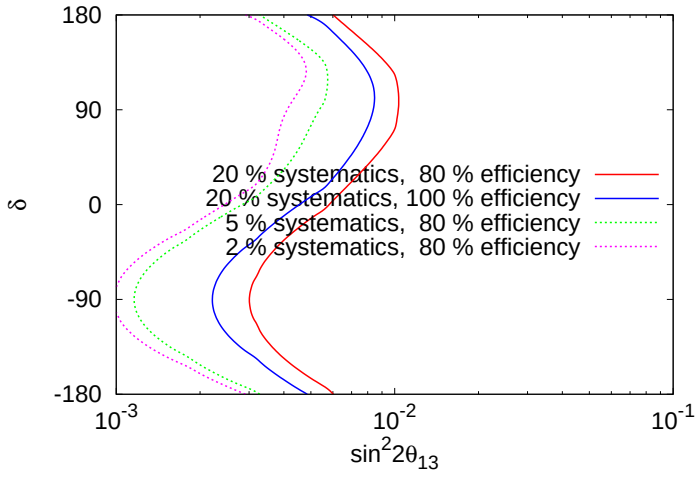

(a) θ_{13} discovery potential.

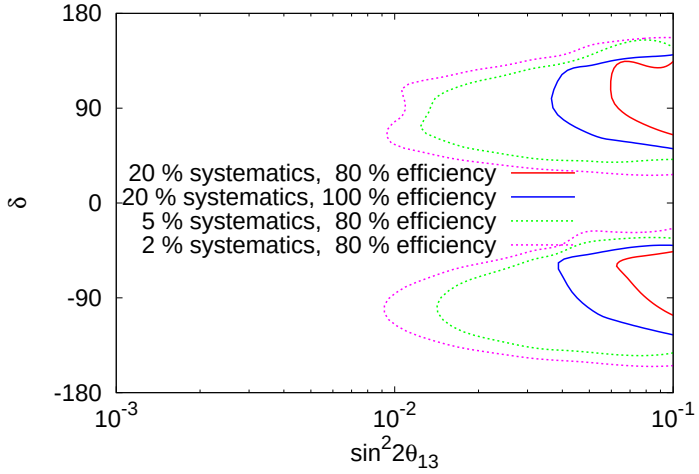

(b) CP discovery potential.

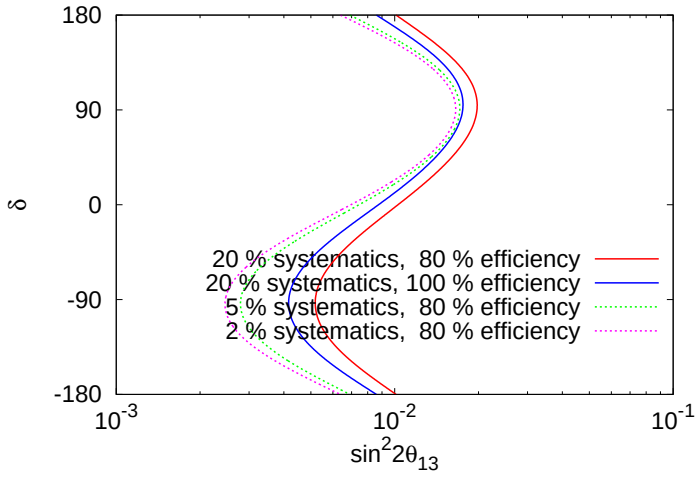

(c) Hierarchy sensitivity.

Figure 5.8.: Effect of systematic errors (20%, 5% and 2%) and efficiency (80% and 100%) on a) θ_{13} discovery potential, b) CP discovery potential and c) hierarchy sensitivity.

energies. We assume an energy threshold of 0.1 GeV and a maximum energy of 10 GeV and use similar binning as in Ref. [107] (0-0.5 Gev, 0.5-1.0 GeV, 1.5-2.0 GeV, 2.0-3.0 GeV, 3.0-10.0 GeV). From Ref. [107], the NC background is 5%.

For the LENA detector (50 kton liquid scintillator), we assume the fiducial mass to be 50 kton from Ref. [180], the efficiency (all channels) to be 90% for all particles with negligible misidentification rate (from Ref. [188]), and the energy resolution to be 5% of the particle's true energy, for all particles from Refs. [187, 188]. We assume an energy threshold of 1.0 GeV, systematic errors of 5% as for the other detectors, the beam ν_e ($\bar{\nu}_e$) background to be 50%, the background from the $\bar{\nu}$ (ν) component of the ν ($\bar{\nu}$) beam is $\sim 5\%$, and the NC background is 0.5% for all channels, as for the liquid argon detector.

We also compare the performance of the LAGUNA setup to that of other experiments which may be running at around the same time. In Fig. 5.9 we show the θ_{13} discovery potential, CP discovery potential and hierarchy sensitivity, as a function of $\sin^2 2\theta_{13}$ and δ for the LAGUNA setup with a liquid argon detector (solid red lines), WC detector both with a realistic 5% NC background (solid green lines) and when this background is hypothetically removed (dashed green lines), and the liquid scintillator detector with a pessimistic 5% NC background (solid blue lines) and with the estimated 0.5% NC background (dashed blue lines). We compare these results with the estimated results of the LBNE experiment [107] with a WC detector (solid grey lines) and a LAr detector (solid pink lines), and to the Super Proton Linac (SPL) setup (solid light blue lines) as described in Ref. [114].

The LAGUNA setup with a WC or liquid scintillator detector, when a 5% NC background is included, has *no* sensitivity to CP violation. The effect of the NC backgrounds on the WC detector is more severe than for the LENA detector. With a 0.5% NC background for LENA, the performance is roughly identical to that of the LAr detector, despite the fact that LENA only has half the mass of the LAr detector; this must be due in part to the assumption of very good energy resolution (5%) at *all* energies. In this case, the LAr detector and LENA begin to have sensitivity to θ_{13} for $\sin^2 2\theta \gtrsim 10^{-3}$ and to δ and the mass hierarchy for $\sin^2 2\theta \gtrsim 10^{-2}$. The WC detector has roughly an order of magnitude less sensitivity to δ and the mass hierarchy. When the NC background on the WC detector is (hypothetically!) removed, the WC detector performs better than either the LAr detector or LENA, especially for CP violation (because WC detectors are optimal for detecting low-energy QE events which provide the most sensitivity to CP violation).

The SPL setup, with a baseline of 130 km, has no sensitivity to the mass hierarchy although it has a better performance for θ_{13} and CP discovery (down to $\sin^2 2\theta_{13} \simeq 10^{-3}$) than any of the LAGUNA setups or LBNE - the SPL beam is used with a WC detector which is an ideal detector to use at this baseline and energy. LBNE has a similar sensitivity to θ_{13} and the mass hierarchy as the LAGUNA setup, although it has better CP sensitivity partly because the beam setup is more aggressive. As for LAGUNA, LBNE performs better with a LAr detector than with a WC detector (when including realistic NC backgrounds).

5.5. Summary

In the first half of this chapter we performed some preliminary studies of a 50 kton liquid scintillator detector as there has been relatively little work carried out on this technology so far. We considered a scenario where θ_{13} is small (2°) and one in which θ_{13} is large (5°). In the case of small θ_{13} we found that the sensitivity is limited by backgrounds and statistics. In the case of large θ_{13} we found that the limiting factors are statistics and systematics. In either case it is essential to have an energy threshold below the energy of the first oscillation maximum. The values of the efficiency and energy resolution are not crucial.

In the second half we assessed the potential of a 100 kton liquid argon detector when used with the HP-PS2 beam from CERN and the CERN to Pyhäsalmi baseline of 2285 km. We performed a sophisticated detailed simulation, including migration matrices for the energy resolution. We found that even though there is a possibility that a liquid argon detector could detect τ's and therefore the τ appearance channels might be accessible, the beam energy at this baseline is not sufficiently above the τ detection threshold to produce a useful number of τ events so that τ detection is essentially useless. We also found that quasi-elastic events are very important for CP sensitivity which means that if a liquid argon detector has different detection properties for quasi-elastic and non-quasi-elastic events, then we should consider how the differences impact upon the sensitivity to CP violation. We then showed that the values of the systematic errors have a large effect on the performance of the experiment which means that a lot of effort should be invested into minimising them.

Finally, we compared the performance of the LAGUNA beam with a 2285 km baseline and each of the three LAGUNA detectors - 100 kton liquid argon, 440 kton water

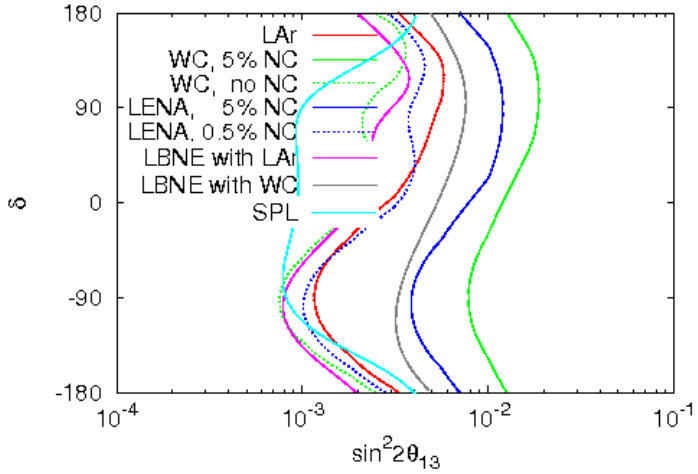

(a) θ_{13} discovery potential.

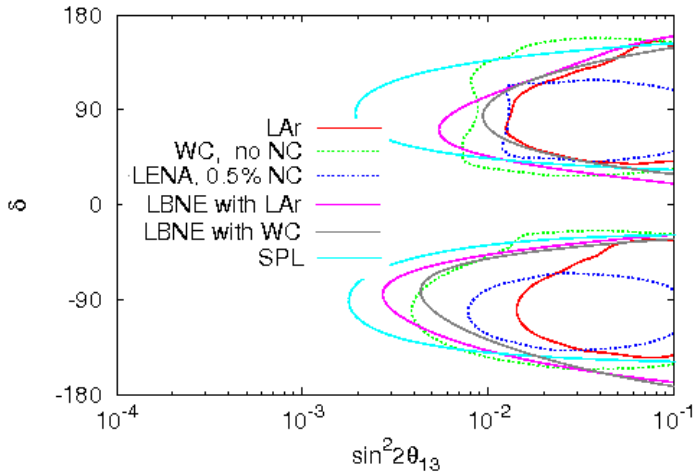

(b) CP discovery potential.

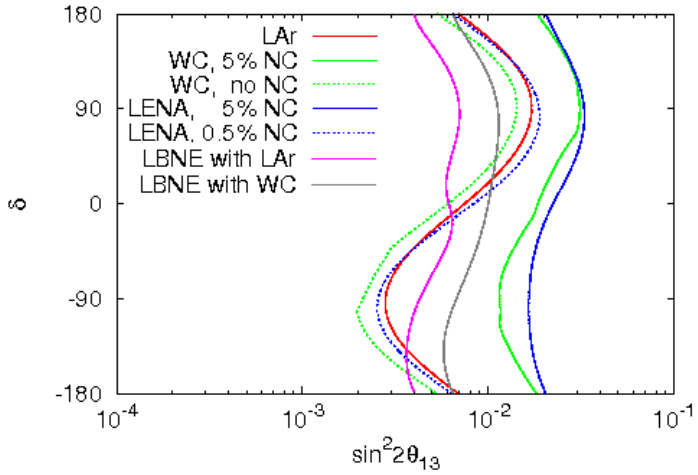

(c) Hierarchy sensitivity.

Figure 5.9.: Comparison of the LAGUNA liquid argon, water Čerenkov and liquid scintillator detectors, with the LBNE and SPL experiments, showing a) θ_{13} discovery potential, b) CP discovery potential and c) hierarchy sensitivity.

Čerenkov and 50 kton liquid scintillator - with that of the proposed SPL experiment and the proposed LBNE setup. We found that the SPL with its 130 km baseline and water Čerenkov detector (which is the optimal detector for this baseline) has the best sensitivity to θ_{13} and δ (down to $\sin^2 2\theta_{13} \simeq 10^{-3}$); however it has no sensitivity to the mass hierarchy because of the short baseline. LBNE with either a liquid argon or water Čerenkov detector also performs very well, partly due to the aggressive beam setup. Of the LAGUNA detectors, the liquid argon and liquid scintillator detectors have nearly identical performances, in spite of the liquid scintillator detector having only half the mass of the liquid argon detector - this is due to the assumption of its excellent background rejection and energy resolution. Their performance is competitive with that of LBNE in terms of sensitivity to θ_{13} and the mass hierarchy. We find that the water Čerenkov detector with a realistic 5% neutral-current background has no sensitivity to CP violation, but when this background is (hypothetically) removed then it does, and performs better than the other LAGUNA detectors. Similarly we find that if a 5% neutral-current background is included in the liquid scintillator detector then its performance is much diminished, although not quite as dramatically as the water Čerenkov detector. Therefore it is crucial to minimise this background and to establish a reliable estimate of its value in order for realistic simulations and comparisons to be performed.

Chapter 6.

Oscillation measurements at a low-energy neutrino factory

In this chapter we present work regarding the measurement of standard oscillation parameters at a low-energy neutrino factory. We will describe the experimental setup and the optimisation studies performed in order to refine the experimental design of the low-energy neutrino factory. Then we investigate the physics performance of the optimised setup, both in terms of its discovery potential and its precision, and compare its discovery potential to that of the high-energy neutrino factory and other long-baseline experiments.

A large proportion of the material in Sections 6.2, 6.3 and 6.5 is reproduced from Ref. [196].

6.1. Experimental setup

In our standard setup for the LENF we will assume a baseline of 1300 km and a 20 kton TASD (Section 3.3.2). The baseline corresponds to the US baseline from Fermi National Accelerator Laboratory (FNAL), Illinois, to the Deep Underground Science and Engineering Laboratory (DUSEL) [197], Soudan. However, the results of these studies are not completely specific to this baseline and will remain qualitatively similar for other baselines of the same scale - this is discussed further in Section 6.6.

The accelerator section of the LENF is shown in Fig. 6.1 from where it can be compared to the HENF accelerator in Fig. 3.6. The practical advantages of using a

much lower muon energy are that the accelerator is smaller and simpler to construct, and there is only one storage ring pointing to a single baseline.

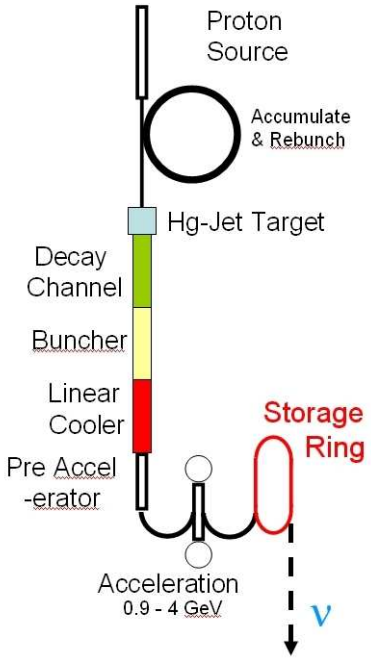

Figure 6.1.: Layout of the low-energy neutrino factory accelerator complex. From A. Bross.

The development of the TASD for the LENF is being led by A. Bross at Fermilab National Accelerator Laboratory. The current design has $\sim 10^5$ modules of plastic scintillator bars with a total mass of 30 kton and fiducial mass of 20 kton and uses the same triangular geometry as MINERvA, which was described in Section 3.3.2. The dimensions of the detector are shown in Fig. 6.2a. The entire detector will be placed in a 0.5 tesla magnetic field by constructing a 'magnetic cavern' out of superconducting transmission lines originally developed for the Very Large Hadron Collider superferric magnets [198], an idea conceived by A. Bross. As will be shown, the TASD is an extremely powerful detector for a LENF. Research and development is ongoing, with one of the main challenges being to find a practical and affordable method of magnetising the TASD.

Studies carried out prior to this thesis have shown that a neutrino factory with a muon energy of around 4 GeV enables very precise measurements of the unknown neutrino oscillation parameters to be made [136, 137]. The setup exploits a fully active calorimeter within a magnet, such as the TASD, which ensures the detection of lower energy muons. A magnetised LArTPC (Section 3.3.3) may also be a possibility. The possibility of a LENF with non-magnetic detectors has also been explored by the au-

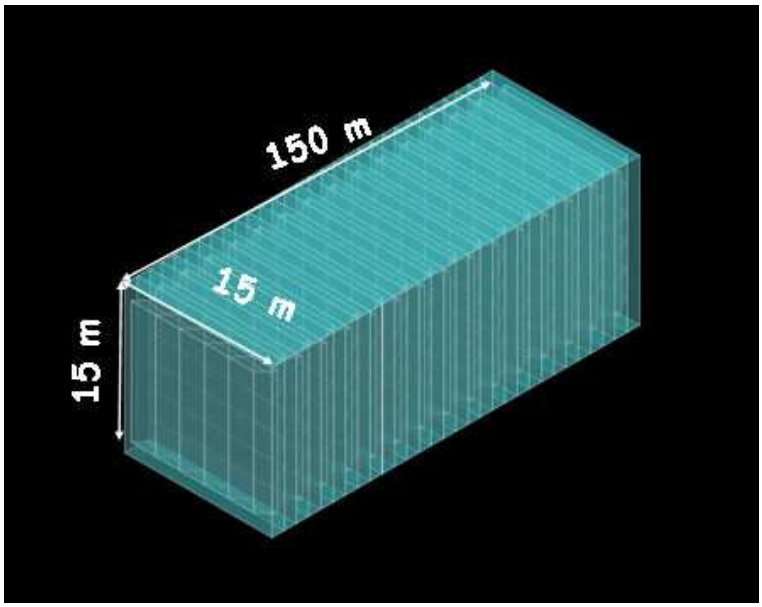

(a) Dimensions of the TASD.

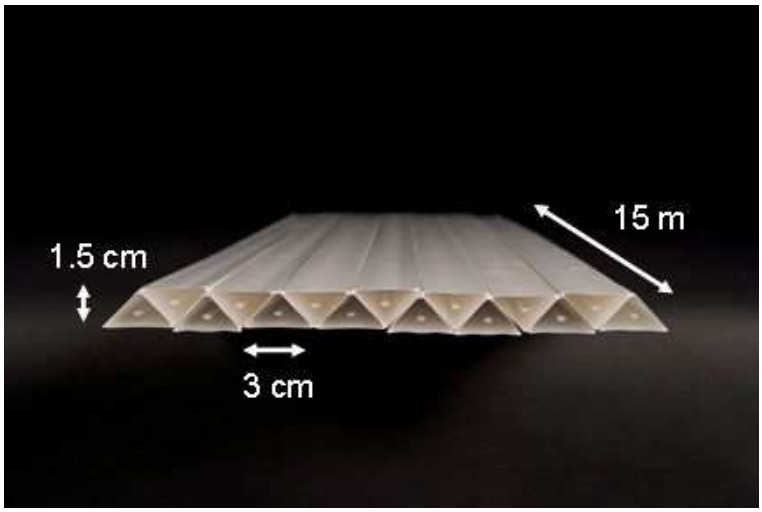

(b) Triangular modules provide the optimal geometry for spatial resolution.

Figure 6.2.: The totally active scintillating detector (TASD). Images courtesy of A. Bross.

thors of Ref. [199] although we will only be considering setups with magnetised detectors. Electron charge identification may also become possible in a LENF equipped with either a TASD or LArTPC - therefore, in addition to the wrong and right-sign *muons*, there will also be wrong and right-sign *electrons* from the $\bar{\nu}_e$ (ν_e) appearance channel (the platinum channel) and the ν_e ($\bar{\nu}_e$) disappearance channel if μ^+ (μ^-) are stored in the decay ring. However, distinguishing the electron signature from the neutral-current events will present a very difficult task. We shall discuss the platinum channel in Section 6.2.6.

6.2. Optimisation studies

In this section we will describe optimisation studies of the LENF, assuming the setup just described. The aim of these studies is to ascertain how each experimental variable affects the performance of the experiment, and to identify the optimal value if there is one. This will give an indication as to which aspects of the experiment design are the most important in order to produce the optimal configuration for maximal physics performance. For the purposes of these studies, it is sufficient to display the results in terms of qualitative 'potato plots' which give an indication as to the *precision* of the measurements which can be made on θ_{13}, δ and sign(Δm^2_{31}). The results are always shown for two different values of θ_{13} - 1° and 5° ($\sin^2 2\theta_{13} \sim 10^{-3}$ and $\sim 10^{-2}$) - corresponding to 'small' and 'large' values of θ_{13}, respectively, and for four different values of δ - 0, $\pm 90°$ and 180°. These are the extreme values of δ and any degeneracies which may arise are likely to be located at one or more of these points (see Sections 4.1 and 4.5).

This work is a continuation of the studies in Refs. [136, 137]. Since these initial studies were performed, advances in the accelerator and detector technologies have enabled improved estimates of the experimental parameters to be made, for which we assess the effects. A summary of the assumptions made for the initial studies, and the refined assumptions used for this present work are given in Table 6.1. $e^{\pm}$ detection was not considered in the original studies; it will be discussed in Section 6.2.6. Backgrounds are defined in Section 6.2.5.

In all simulations, we take the 'true' value of all the known oscillation parameters to be the same as in Ref. [189], as in the previous chapter: $\sin^2\theta_{12} = 0.3$, $\theta_{23} = \frac{\pi}{4}$, $\Delta m^2_{21} = 8.0 \times 10^{-5}$ eV^2, and $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2 with a 10% uncertainty on the atmospheric parameters, 4% uncertainty on the solar parameters, and 2% uncertainty on the matter density (these 'uncertainty' values are the ones which we use as input

Parameter	Previous assumption	Current assumption
Baseline	1480 km	1300 km
Muon energy	4.12 GeV	4.5 GeV
Muon dcays/ year/ polarity	5×10^{20}	1.4×10^{21} decays
Running time/ polarity	10 years	10 years
Detector fiducial mass	20 kton	20 kton
Energy threshold	0.5 GeV	0.5 GeV
Energy resolution	30%	10%
Efficiency for $\mu^{\pm}$ (dis)appearance	73% (all energies)	$73\% \leq 1$ GeV
		$94\% > 1$ GeV
Background on $\mu^{\pm}$ (dis)appearance	10^{-3}	10^{-3}
Efficiency for $e^{\pm}$ appearance	-	$37\% \leq 1$ GeV
		$47\% > 1$ GeV
Background on $e^{\pm}$ appearance	-	10^{-2}
Systematics	2% (uncorrelated)	2% (uncorrelated)

Table 6.1.: Assumptions used in the initial LENF studies [136, 137], and the refined assumptions used in the present work.

errors for each of the parameters). As we mentioned in Chapter 5, although these values are not the most recent best-fit values, the purpose of using them is to be consistent with studies performed by other groups, thus allowing for a fair quantitative comparison to be made. In any case, as the values have not changed significantly and all the parameters are marginalised over in the simulations, the precise value used should have virtually no effect on the results. We use the exact oscillation probabilities including matter effects (calculated by GLoBES), and all parameters are marginalised over unless otherwise specified. We have also assumed a true normal hierarchy in all simulations unless otherwise stated, and have verified that the results are qualitatively similar for a true inverted hierarchy.

6.2.1. Muon energy

In order to determine the value of the optimal muon energy, it is necessary to look at the neutrino spectrum produced from the decay of muons of a particular energy. The 'muon energy' refers to the energy of the muons in the storage ring which decay into

the neutrino beam. There is some spread of muon energies about the nominal value. The maximum energy of any neutrino is the energy of the muon from which it was produced. The aim of the optimisation is to maximise the number of signal (oscillating) events whilst simultaneously minimising the background. Backgrounds arise from non-oscillating events at an energy above that of the oscillation region which produces higher-energy charged-current and neutral-current events. Pions can also be produced by high-energy neutrinos in processes such as $\nu_\mu p \to \mu^- p \pi^+$; this is a particular problem for the TASD as distinguishing between the electron and pion signatures is very difficult.

The optimal muon energy depends upon the baseline used as this determines the position of the oscillation region. Here we consider a baseline of 1300 km, with a discussion about alternative baselines and energies given in Section 6.6. In Fig. 6.3 we show the spectra of neutrinos produced from the decay of muons of different energies - 3.0 GeV (red), 4.5 GeV (green), 6.0 GeV (blue) and 7.5 GeV (purple). The number of neutrinos corresponds to the total number of ν_μ and $\bar{\nu}_e$ if μ^- are stored, or $\bar{\nu}_\mu$ and ν_e if μ^+ are stored. Also shown are the golden channel probabilities (for which the y-axis scale is not relevant) for both normal and inverted hierarchies (black solid and dotted lines, respectively) so that the oscillation region can be seen. The region of the first peak is around 2 to 3 GeV, so we are interested in all events at and below this energy. Using a muon energy of 3 GeV, the neutrino spectrum peaks at 2 to 2.5 GeV. There are no events above ~ 4.5 GeV. Using a muon energy of 4.5 GeV, there is a significant increase in the number of events in the range 1.5 to 3 GeV, where oscillations occur. Moving up to 6 GeV gives a slight increase in the number of higher-energy oscillating events although there is also a slight decrease of low-energy events, and a muon energy of 7 GeV gives no increase in the number of oscillating events (in fact, there is a slight decrease) but there is a large number of high-energy background events. From this it is expected that the optimal muon energy is around 4.5 to 5 GeV.

In Fig. 6.4 we show how altering the muon energy affects the precision of the $\theta_{13} - \delta$ measurement. Results are shown for an energy of 4.12 GeV as in the original studies (solid red lines), 5.0 GeV (solid green lines) and 6.0 GeV (dotted blue lines). There is a visible improvement in increasing the energy from 4.12 GeV to 5 GeV which is a lot more evident for small θ_{13}. A further increase to 6.0 GeV does not significantly change the results in either case, which is consistent with the expectations of the previous paragraph. Taking into account this, and practical considerations for the accelerator design, the reference energy for the LENF setup is taken to be 4.5 GeV.

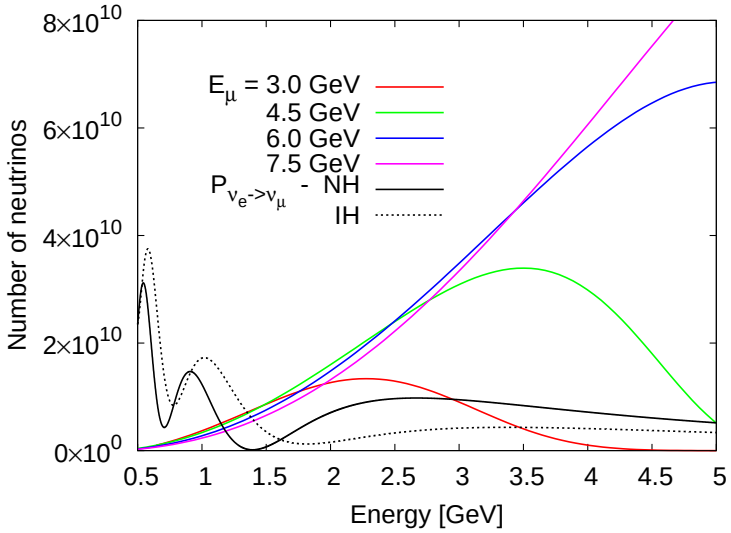

Figure 6.3.: Neutrino spectra arising from different muon energies. Shown in black is the golden channel probability for normal hierarchy (solid line) and inverted hierarchy (dotted line).

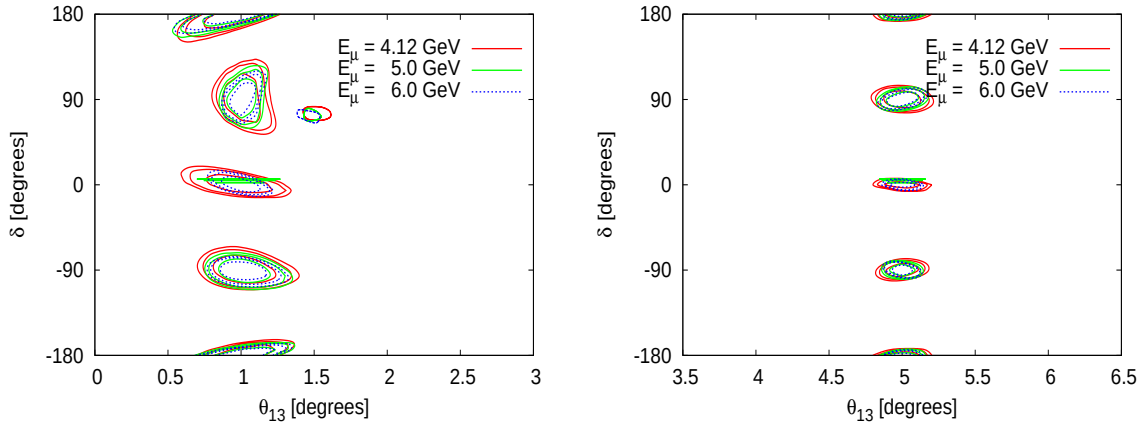

(a) $\theta_{13} = 1°$. (b) $\theta_{13} = 5°$.

Figure 6.4.: 68%, 90% and 95% contours in the $\theta_{13} - \delta$ plane for muon energies of 4.12 GeV (solid red lines), 5.0 GeV (solid green lines) and 6.0 GeV (dotted blue lines) for a) $\theta_{13} = 1°$ and b) $\theta_{13} = 5°$ and $\delta = 0°$, $\pm 90°$ and $180°$.

6.2.2. Energy resolution

The energy resolution is the accuracy with which a particle's energy is measured and reconstructed by the detector and analysis software. The energy resolution is important as it enables the *shape* of the spectrum to be reconstructed accurately. This is vital in order to be able to distinguish between similarly-shaped spectra which arise from different parameter combinations. Essentially all information about the oscillation parameters is encoded in the shape of the spectrum. So theoretically, the more perfect the energy resolution, the more accurate the measurement. In practice, the performance of the experiment will be limited by other factors as well so that there will be some threshold value below which no further gain is obtained.

In GLoBES, the energy resolution, $\sigma(E)$, is parameterised as the function $\sigma(E) = \alpha E + \beta\sqrt{E} + \gamma$. This can be understood roughly as follows: the term αE will arise due to measurements of the energy from a track length, where the percentage error is usually roughly constant; $\beta\sqrt{E}$ is a Poisson term which takes into account the fact that the sampler is finite and so has an error proportional to the square root of the event number $\sim \sqrt{E}$; and the constant γ describes energy-independent effects such as the dark count of the electronics. For the TASD, we consider only the α term i.e. assume a constant percentage error.

In Fig. 6.5 we show how the revised estimate of $\sigma(E) = 10\%$ (solid red lines) improves upon the performance of the original estimate of 30% (dotted blue lines). The measurement of all the parameters θ_{13}, δ and the mass hierarchy is improved. This is visible for large as well as small θ_{13}. More optimistic values for the energy resolution have been simulated and it has been found that they offer no improvement beyond the value of 10%. As this is a value which appears to be experimentally feasible, 10% is taken to be the reference value.

6.2.3. Energy detection threshold

The most stringent requirement on the detection threshold is that it must lie below the value of the first oscillation maximum - this is the region from which most information is obtained, and is the signal region with the highest statistics. It may also be preferable for the detection threshold to lie below the *second* oscillation maximum because, as explained in Section 4.1, CP violation is most visible at *low* energies such as at the second oscillation maximum. In Fig. 6.6 we investigate how vital this property is, by

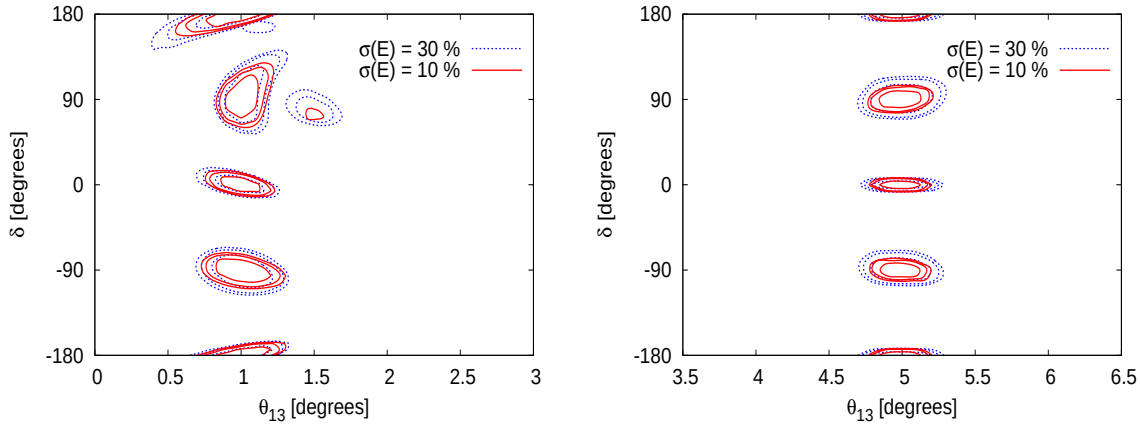

(a) $\theta_{13} = 1°$. (b) $\theta_{13} = 5°$.

Figure 6.5.: 68%, 90% and 95% contours in the $\theta_{13} - \delta$ plane for an energy resolution of 30% (dotted blue lines) and 10% (solid red lines) for a) $\theta_{13} = 1°$ and b) $\theta_{13} = 5°$ and $\delta = 0°$, $\pm 90°$ and $180°$.

comparing our current assumption of a 0.5 GeV threshold (detection is possible at 73% efficiency from 0.5 GeV; maximum efficiency of 94% is reached at 1.0 GeV), shown by the solid red lines, to a scenario where the threshold is at 1.5 GeV (maximum efficiency reached at 2 GeV), shown by the dotted blue lines. It can be seen that there is actually very little difference between the two scenarios - only a slight gain in the resolution of the mass hierarchy for small θ_{13}. So although, in theory, the second oscillation maximum should help, we have found that in practice there aren't sufficient events in that region of the energy spectrum for much useful information to be extracted - this is in agreement with the findings recently published in Ref. [200]. The statistics at the first oscillation maximum, combined with good energy resolution, ensure that the shape of the spectrum around the first maximum can be probed sufficiently accurately to enable the oscillation parameters to be determined.

6.2.4. Statistics

Increasing the number of events in an experiment decreases the statistical error on the measurement and therefore the more events, the better. This number can be increased in a number of ways: by increasing the initial flux of neutrinos (by increasing the initial proton intensity), by running the beam for a longer period of time, or by increasing the mass of the detector. All of these are limited by economic and practical factors!

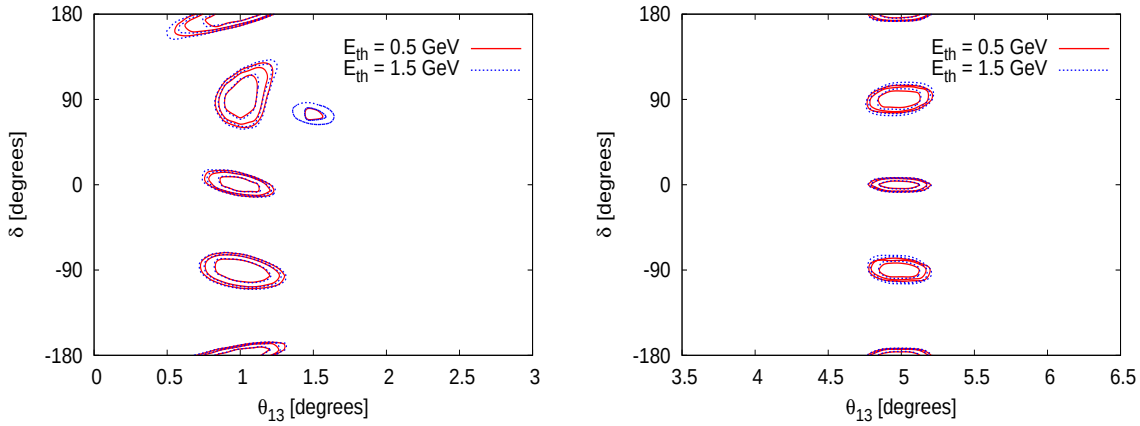

(a) $\theta_{13} = 1°$. (b) $\theta_{13} = 5°$.

Figure 6.6.: 68%, 90% and 95% contours in the $\theta_{13} - \delta$ plane for an energy threshold of 0.5 GeV (solid red lines) and 1.5 GeV (dotted blue lines) for a) $\theta_{13} = 1°$ and b) $\theta_{13} = 5°$ and $\delta = 0°$, $\pm 90°$ and $180°$.

The *efficiency* of a detector is the proportion of incident events it detects relative to the actual number of events which occur in the detector. It is limited by both hardware and software factors, such as the finite spacing of detector components and electronics, electronic dead-time, and software reconstruction efficiencies. In essence, the effect of a higher efficiency is to increase the statistics and so the performance will always be improved by a higher efficiency.

6.2.5. Backgrounds

For simulation purposes, backgrounds are assumed to arise primarily from neutral-current events which are wrongly identified as charged-current events, and from charge misidentification, such that a fixed percentage of events from any channel which produces a lepton of the same flavour but opposite charge to that of the channel signal is taken to be the background. The signal channels and their associated backgrounds which are used in our simulations are shown in Table 6.2.

Polarity mode	Signal channel	Background channels
μ^-	$\nu_\mu \to \nu_\mu$	NC $\nu_\mu \to \nu_\mu$, CC $\bar{\nu}_e \to \bar{\nu}_\mu$
	$\bar{\nu}_e \to \bar{\nu}_\mu$	NC $\nu_\mu \to \nu_\mu$, CC $\nu_\mu \to \nu_\mu$
	$\nu_\mu \to \nu_e$	NC $\bar{\nu}_e \to \bar{\nu}_e$, CC $\bar{\nu}_e \to \bar{\nu}_e$
μ^+	$\bar{\nu}_\mu \to \bar{\nu}_\mu$	NC $\bar{\nu}_\mu \to \bar{\nu}_\mu$, CC $\nu_e \to \nu_\mu$
	$\nu_e \to \nu_\mu$	NC $\bar{\nu}_\mu \to \bar{\nu}_\mu$, CC $\bar{\nu}_\mu \to \bar{\nu}_\mu$
	$\bar{\nu}_\mu \to \bar{\nu}_e$	NC $\nu_e \to \nu_e$, CC $\nu_e \to \nu_e$

Table 6.2.: Backgrounds assumed on each of the LENF channels.

6.2.6. The platinum channel

The synergy between the golden and platinum channels was described in Section 4.5. As explained, theoretically the combination will work to resolve degeneracies. However we find that practically, when the difficulty of detecting the platinum channel signal is taken into account and realistic values for the efficiency and backgrounds are used, the power of the platinum channel may be diminished to the extent that more can be gained by improving other aspects of the detector and experiment, than by inclusion of the platinum channel.

We define *Scenario 1* to be the one in which only $\mu^\pm$ detection is possible, giving us access to only the ν_μ and $\bar{\nu}_\mu$ appearance and disappearance channels. In *Scenario 2* it is also possible to detect $e^\pm$ and hence exploit the additional information from the ν_e and $\bar{\nu}_e$ appearance (platinum) channels. To illustrate the impact of the addition of these channels to our setup, in Fig. 6.7 we compare the sensitivities of the two scenarios when using a muon decay rate of 5.0×10^{20} (left column) as used in Refs. [136, 137] and 1.4×10^{21} (right column) per year, varying the background level of the ν_e ($\bar{\nu}_e$) appearance channel from a hypothetical zero (top row) to 10^{-2} (bottom row). Thus we show how the value of the platinum channel depends on the statistics and background level assumed, and also how an increase in statistics affects the experimental performance.

In the case of the lower statistics, we observe that the addition of the platinum channel with zero background produces a drastic improvement in sensitivity to all parameters. For a background of 10^{-2} the improvement is much smaller but can still help to alleviate the hierarchy degeneracy (see Ref. [178]). At higher backgrounds we find that this gain is lost. In the case of the high statistics, we already observe a smaller improvement for zero background, which becomes insignificant at a background level of 10^{-2}. Thus we

conclude that since the estimated background on the ν_e ($\bar{\nu}_e$) appearance channels will be at best $\sim 10^{-2}$, the platinum channel could help in the measurement of the mass hierarchy if statistics are limited to 5.0×10^{20} useful muon decays per year, whereas it will be almost irrelevant for the higher statistics scenario. An increase in statistics in the golden channel provides a much larger improvement in the performance than the addition of the platinum channel if background levels below 10^{-2} cannot be achieved.

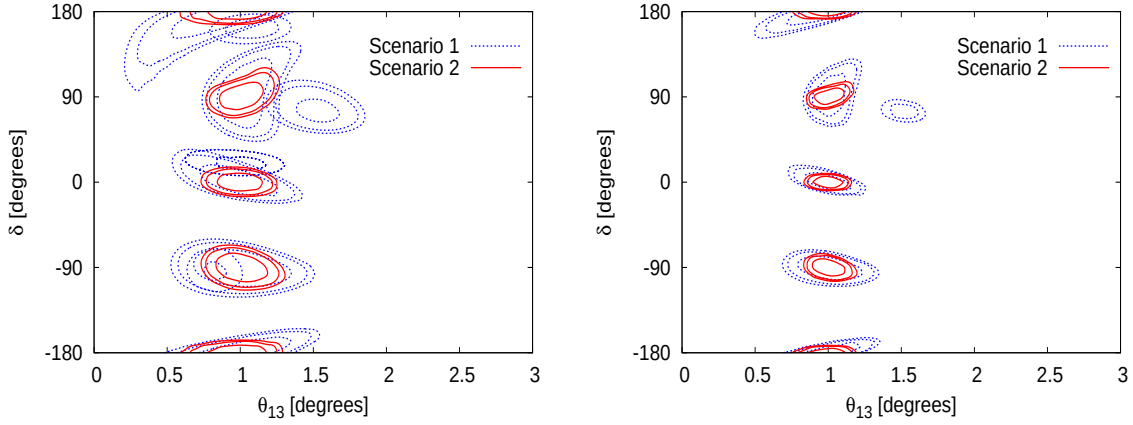

(a) 5.0×10^{20} decays per year, zero background. (b) 1.4×10^{21} decays per year, zero background.

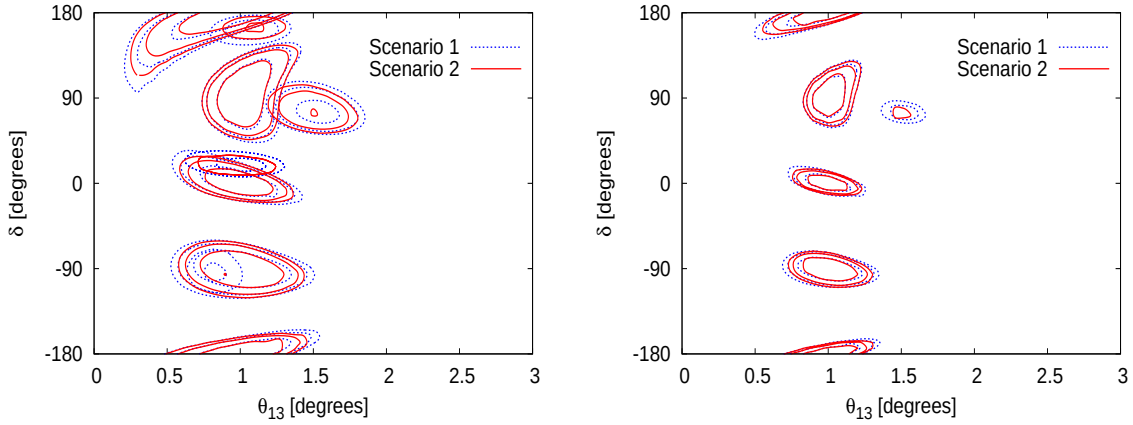

(c) 5.0×10^{20} decays per year, 10^{-2} background. (d) 1.4×10^{21} decays per year, 10^{-2} background.

Figure 6.7.: Comparison of Scenario 1 (ν_μ and $\bar{\nu}_\mu$ appearance and disappearance only - dotted blue lines), and Scenario 2 (ν_e and $\bar{\nu}_e$ appearance included - solid red lines) when using 5.0×10^{20} $\mu^\pm$ decays per year (left) and 1.4×10^{21} decays per year (right), and a background of zero (top row) or 10^{-2} (bottom row) on the ν_e ($\bar{\nu}_e$) channels. The 68%, 90% and 95% confidence level contours in the $\theta_{13} - \delta$ plane are shown, for $\theta_{13} = 1°$ and $\delta = 0°$, $\pm 90°$ and $\pm 180°$.

6.2.7. Systematic errors

'Systematic errors' is the term used to encompass all non-random errors in the experiment. The dominant source of systematic errors in any neutrino oscillation experiment is from the flux and cross-section uncertainties. One of the strengths of the neutrino factory is that the initial flux can be predicted to a much greater precision than for any other experiment, and so the systematic error associated with this measurement is relatively small. The cross-section uncertainties are specific to the detector being used - events have to be modelled and predicted for interactions within the specific detector material and configuration. Other sources of systematic errors are calibration errors, the error on the mass of the detector, and there can also be systematic errors specific to the detector being used. The effect of making a more conservative estimate of 5% for the systematic errors, rather than 2%, is found to be almost negligible.

6.3. Physics performance

Here we present the results which can be obtained by the reference setup defined in the previous section (Table 6.1), in terms of 3σ θ_{13} discovery potential, CP discovery potential, and sensitivity to the mass hierarchy in the $\sin^2 2\theta_{13} - \delta$ plane (Fig. 6.8). In addition we also consider the 3σ sensitivity to θ_{23} in the $\sin^2 2\theta_{13} - \sin\theta_{23}$ plane, both in terms of the ability to exclude a maximal value of θ_{23} (Fig. 6.9a) and to identify the octant of θ_{23} (Fig. 6.9b). The results from our optimised setup described in Section 6.2 are shown by the solid green lines. We have also considered a setup where only the statistics are altered, to 2.8×10^{21} decays per year (solid red lines), and a setup where only the muon energy is increased to 6.0 GeV (dashed blue lines). From this we demonstrate that for all the observables considered, doubling the flux is always preferable to an increase in energy.

For θ_{13} discovery potential, CP discovery potential and θ_{23} sensitivity we only show the results for a normal hierarchy, having verified that similar results are obtained for an inverted hierarchy. We have assumed in Fig. 6.9 (θ_{23} sensitivity) a value of $\delta = 90°$ although we have also studied other values of δ and find no strong dependence on the CP phase, since sensitivity to θ_{23} is mainly obtained from terms with no dependence on δ in the oscillation probabilities discussed in Section 4.1. For the exclusion of $\theta_{23} = 45°$, an upward curve is seen for large θ_{13}. This can be understood because the addition of a

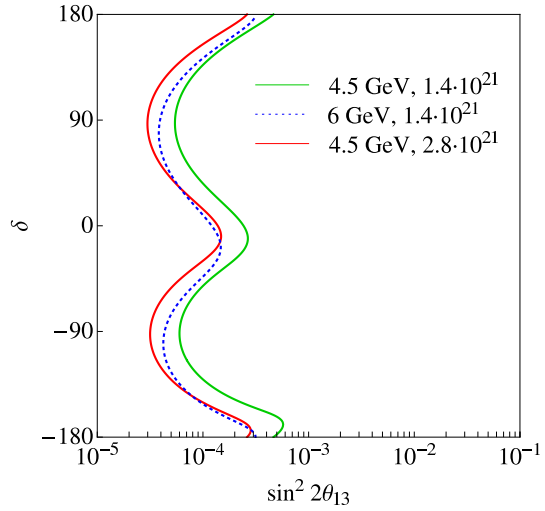

(a) θ_{13} discovery potential.

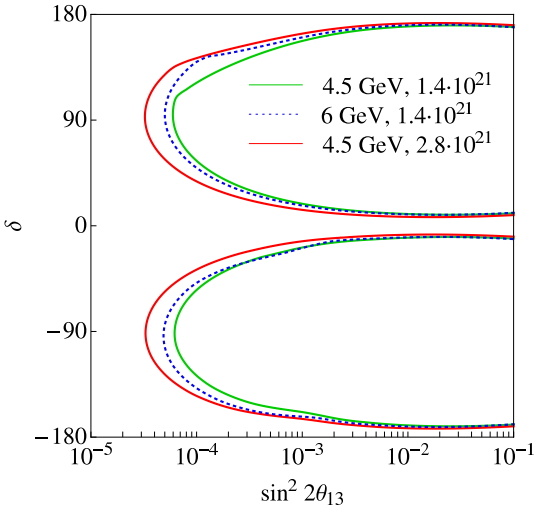

(b) CP discovery potential.

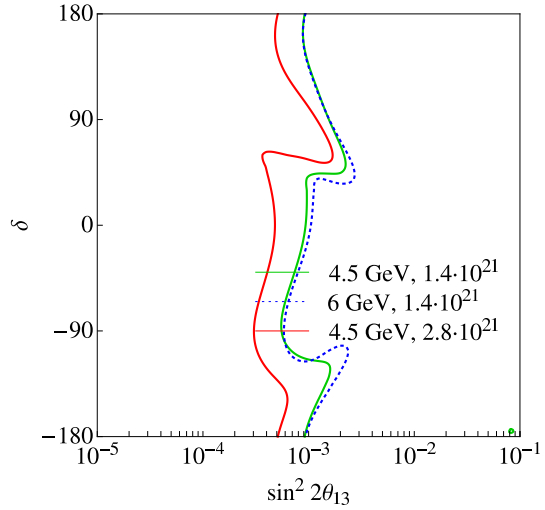

(c) Hierarchy sensitivity (normal hierarchy).

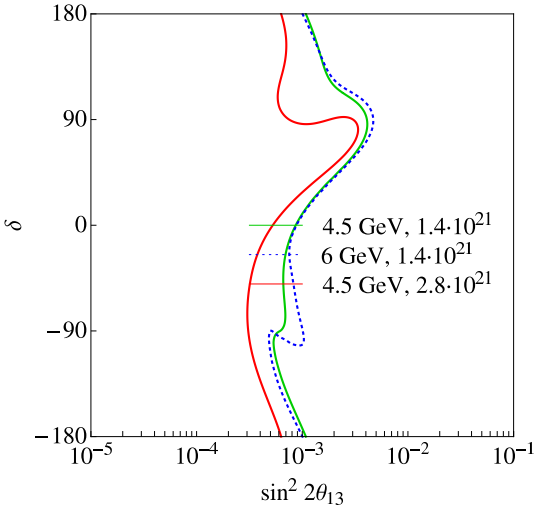

(d) Hierarchy sensitivity (inverted hierarchy).

Figure 6.8.: 3σ confidence level contours in the $\sin^2 2\theta_{13} - \delta$ plane for a) θ_{13} discovery potential, b) CP discovery potential, c) hierarchy sensitivity (for true normal hierarchy) and d) hierarchy sensitivity (for true inverted hierarchy), for muon energies of 4.5 GeV and 6 GeV, and fluxes of 1.4×10^{21} muon decays per year per polarity and 2.8×10^{21} muon decays per year per polarity.

large θ_{13} to the ν_μ disappearance probability introduces an asymmetry in θ_{23} that shifts the contours to larger values (see Eq. 1 and Fig. 8 of Ref. [201]).

We note that this setup has remarkable sensitivity to θ_{13} and δ for values of $\sin^2 2\theta_{13} > 10^{-4}$, and that its sensitivity to the mass hierarchy is an order of magnitude better that that of other proposed experiments exploiting the same baseline e.g. the wide-

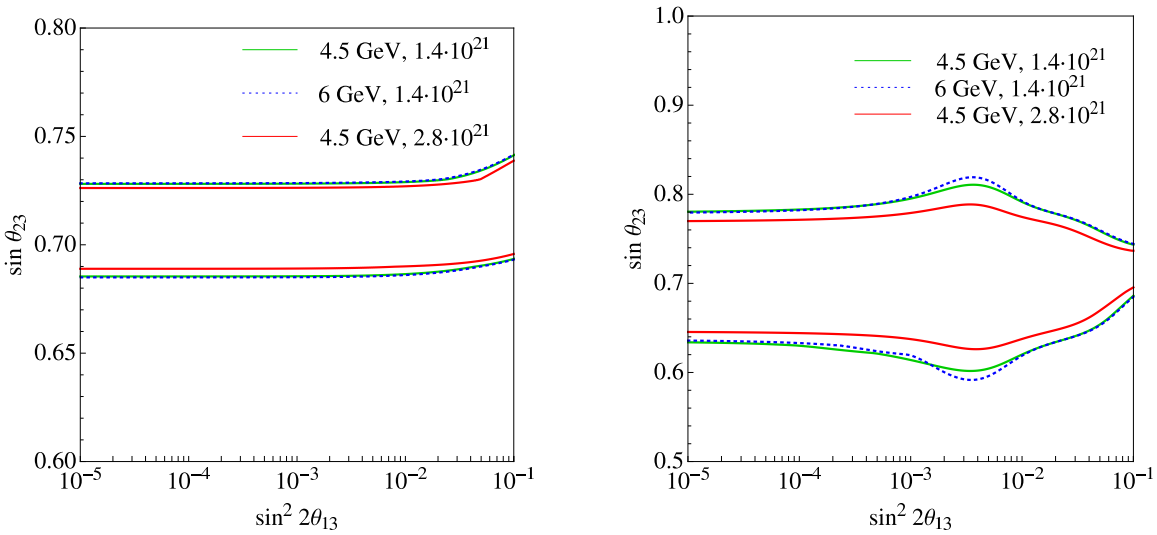

(a) Sensitivity to $\theta_{23} \neq 45°$.

(b) Sensitivity to the θ_{23} octant.

Figure 6.9.: 3σ allowed regions in the $\sin^2 2\theta_{13} - \sin\theta_{23}$ plane for a) potential to exclude $\theta_{23} = 45°$ and b) sensitivity to the θ_{23} octant, for $\delta = 90°$ and muon energies of 4.5 GeV and 6 GeV, and fluxes of 1.4×10^{21} muon decays per year per polarity and 2.8×10^{21} muon decays per year per polarity.

band beam experiment in Refs. [107, 202, 203]. We can attribute these qualities to the unique combination of high statistics and good background rejection coupled with an intermediate baseline, allowing for a clean measurement of the CP phase whilst also allowing for the mass hierarchy to be determined for $\sin^2 2\theta_{13} > 10^{-3}$.

We have also explored how the precision with which θ_{13}, δ and the deviation from maximal θ_{23} could eventually be measured at the LENF, varies as a function of exposure (detector mass $\times$ decays) per polarity. Our standard setup corresponds to 20 kton $\times$ 1.4×10^{21} decays/ year $\times$ 10 years $= 2.8 \times 10^{23}$ kton $\times$ decays per polarity. The gain in precision is much less pronounced for values larger than 6×10^{23} kton $\times$ decays per polarity, hence it may not be worth trying to increase the exposure beyond this value.

Fig. 6.10a shows the 1σ error expected in the measurement of the mixing angle θ_{13} at the LENF as a function of the exposure (in kton $\times$ decays) per polarity, assuming that nature has chosen $\theta_{13} = 5°$. The dependence of these results on the value of the CP violating phase is very mild. The 1σ error in the extraction of θ_{13} when no backgrounds and no systematic errors are included in the analysis is illustrated in Fig. 6.10b. Comparing the two panels we observe that non-zero systematics and backgrounds effectively halve the exposure.

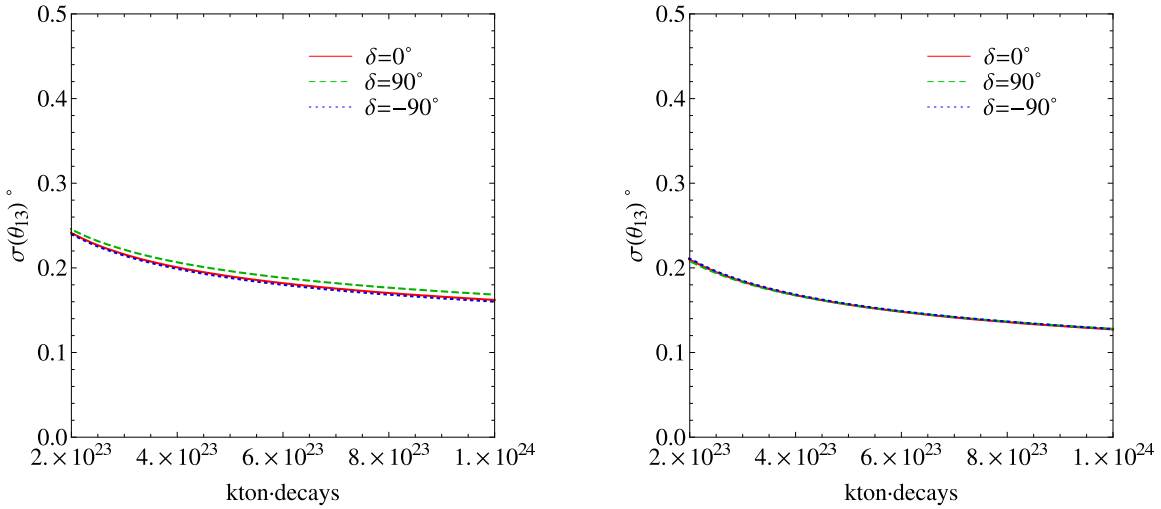

(a) Sensitivity to θ_{13}, including systematic errors and backgrounds.

(b) Sensitivity to θ_{13}, no systematic errors and backgrounds.

Figure 6.10.: 1σ error in the measurement of the θ_{13} mixing angle for a simulated value of $\theta_{13} = 5°$ and $\delta = 0°$ and $\pm 90°$ when a) including systematic errors and backgrounds and b) no systematic errors and backgrounds are included.

Fig. 6.11a shows the 1σ error expected in the measurement of the CP phase δ as a function of the exposure for a simulated value of $\theta_{13} = 5°$, for different values of δ. The results are highly dependent on the value of the CP violating phase, as expected. For $\delta = 90°$, there are strong correlations with θ_{13}, as can be seen from Fig. 6.7, and therefore the error in the measurement of δ is larger. The 1σ error in the extraction of δ when no backgrounds and no systematic errors are included in the analysis is illustrated in Fig. 6.11b. Switching off systematic errors and backgrounds has a larger impact for the $\delta = 0°$ case, again effectively halving the exposure, since correlations among δ and θ_{13} are negligible when $\delta = 0°$ and the precision is more limited by the background and systematic errors instead.

We also explore the sensitivity to maximal mixing, i.e. the ability to exclude $\theta_{23} = 45°$, versus the exposure. We present the 3σ results in Fig. 6.12. We have used a simulated value of $\theta_{13} = 0°$ here (so that δ is irrelevant) as the sensitivity to θ_{23} maximality comes primarily from the ν_μ ($\bar{\nu}_\mu$) disappearance channels which are not dependent on θ_{13}.

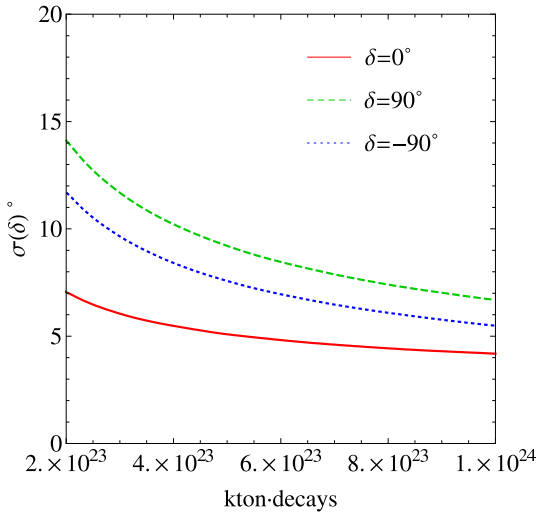

(a) Sensitivity to δ, including systematic errors and backgrounds.

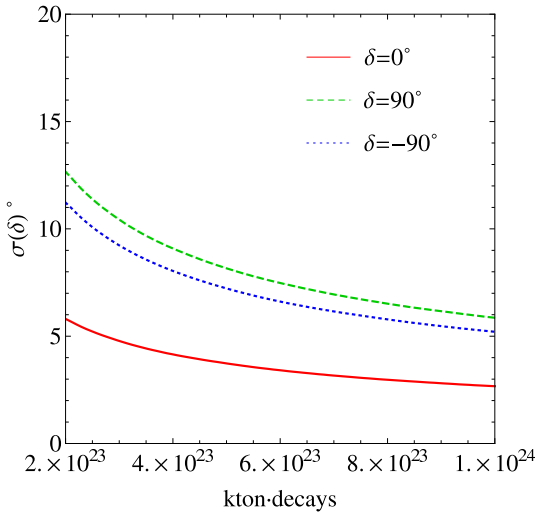

(b) Sensitivity to δ, no systematic errors and backgrounds.

Figure 6.11.: 1σ error in the measurement of δ for a simulated value of $\theta_{13} = 5°$ and $\delta = 0°$ and $\pm 90°$ when a) including systematic errors and backgrounds and b) no systematic errors and backgrounds are included.

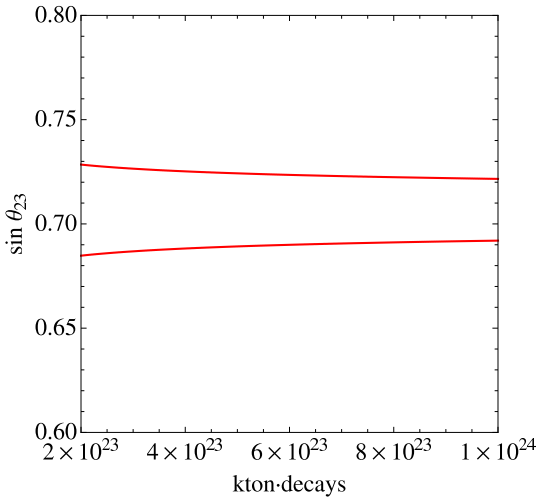

(a) Sensitivity to $\theta_{23} \neq 45°$, including systematic errors and backgrounds.

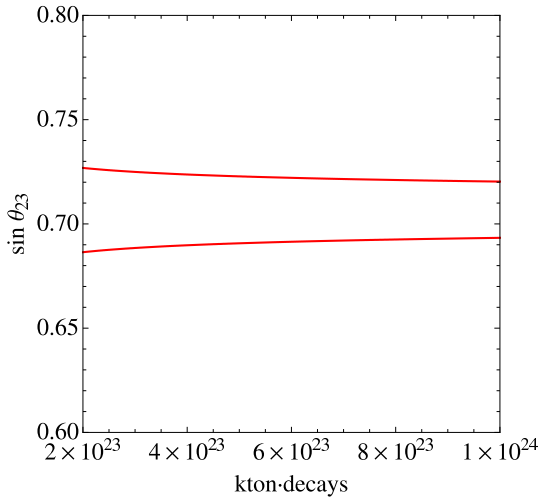

(b) Sensitivity to $\theta_{23} \neq 45°$, no systematic errors and backgrounds.

Figure 6.12.: 3σ regions for which maximal θ_{23} can be excluded, using a simulated value of $\theta_{13} = 0°$ when a) including systematic errors and backgrounds and b) no systematic errors and backgrounds are included.

6.4. Measurement precision as a function of θ_{13}

In Figs. 6.10 and 6.11 the 1σ errors on the measurement of θ_{13} and δ as a function of exposure were displayed. It is also interesting to observe how the errors vary as a function of θ_{13} to provide some indication of the kind of precision which can be expected for particular values of θ_{13}.

In Fig. 6.13 we show the *fractional* error on θ_{13} as a function of the actual true value of θ_{13}. For $\theta_{13} \lesssim 1^\circ$ this error is large because the solar term is dominant. For large values of θ_{13} when the atmospheric term is dominant, the number of oscillation events scales roughly as $N \sim \sin^2 2\theta_{13} \sim \theta_{13}^2$. The error on this number scales roughly as $\sqrt{N} \sim \theta_{13}$. Therefore the fractional error, $\frac{\sigma(\theta_{13})}{\theta_{13}}$, is constant for large θ_{13}. Our LENF setup can achieve an error of $\sim 5\%$ in this region, and roughly a $\sim 20\%$ error when $\theta_{13} = 1^\circ$. These results are for a value of $\delta = 0$, but the value of δ has no significant effect on the precision.

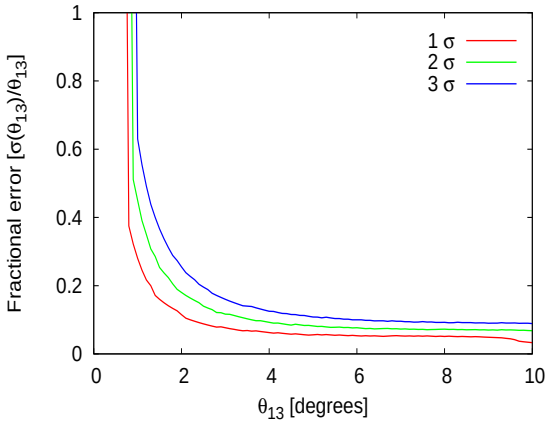

Figure 6.13.: θ_{13}-dependence of the 1, 2 and 3σ errors on the measurement of θ_{13} for $\delta = 0^\circ$.

In Fig. 6.14 we show the *absolute* error on the measurement of θ_{23} for $\theta_{23} = 45^\circ$. In Fig. 6.14a we show, separately, the positive (solid lines) and negative (dotted lines) errors on this measurement, whereas Fig. 6.14b shows the *total* error i.e. the sum of the positive and negative errors. Most of the information on θ_{23} comes from the ν_μ disappearance channel $\sim \sin^2 2\theta_{23}$; as this is insensitive to the octant of θ_{23} the errors obtained from this channel should be symmetric. The discrepancy in the magnitudes of the positive and negative errors therefore must arise from information from the golden channel and can be explained by considering the θ_{23}-dependence of the terms which are dominant in each particular region. For small θ_{13} when the solar term is dominant, the

contribution from θ_{23} comes from $\cos^2\theta_{23}$ (see Eq. (4.4)). At $\theta_{23} = 45°$, the gradient of this function is steeper for the region above 45° than for below and so it is easier to make a measurement in the upper quadrant. Therefore the negative error should be larger as this is the harder measurement. For large θ_{13}, the atmospheric term depends on $\sin^2\theta_{23}$ for which the gradient is steeper in the lower than the upper quadrant; hence the positive error is larger. A sharp dip can be seen in the positive 1σ error; this occurs because of the contribution to the ν_μ disappearance probability for large θ_{13} (see the comment about Fig. 6.9 in Section 6.3). The *total* error (Fig. 6.14b) is roughly independent of θ_{13} and is $\sim 2.5°$ at 1σ for the LENF. Again, these results are for $\delta = 0$ and the value of the CP phase does not affect the sensitivity to θ_{23}.

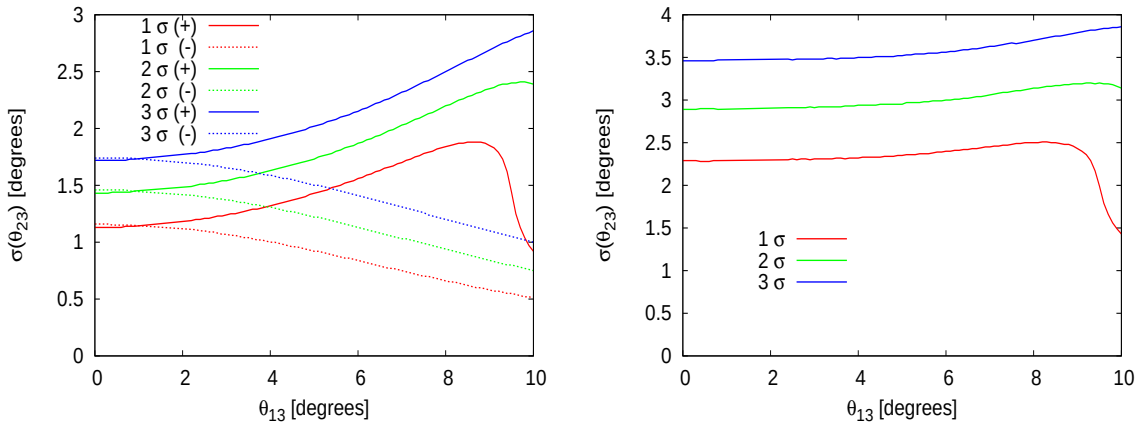

(a) Separate positive and negative errors.

(b) Total error.

Figure 6.14.: θ_{13}-dependence of the 1, 2 and 3σ errors on the measurement of θ_{23} for $\delta = 0°$. The positive and negative errors are shown separately in a), and the total errors in b).

In Fig. 6.15 we show the absolute errors on the measurement of δ, for four different values of δ. For $\theta_{13} \gtrsim 2°$ the error is roughly constant - $\sim 10° - 20°$ at 1σ depending on the value of δ. Again, the correlations between θ_{13} and $\delta = 90°$ are apparent from the way that the error aymptotes more slowly for $\delta = 90°$, in addition to the fact that the errors are larger.

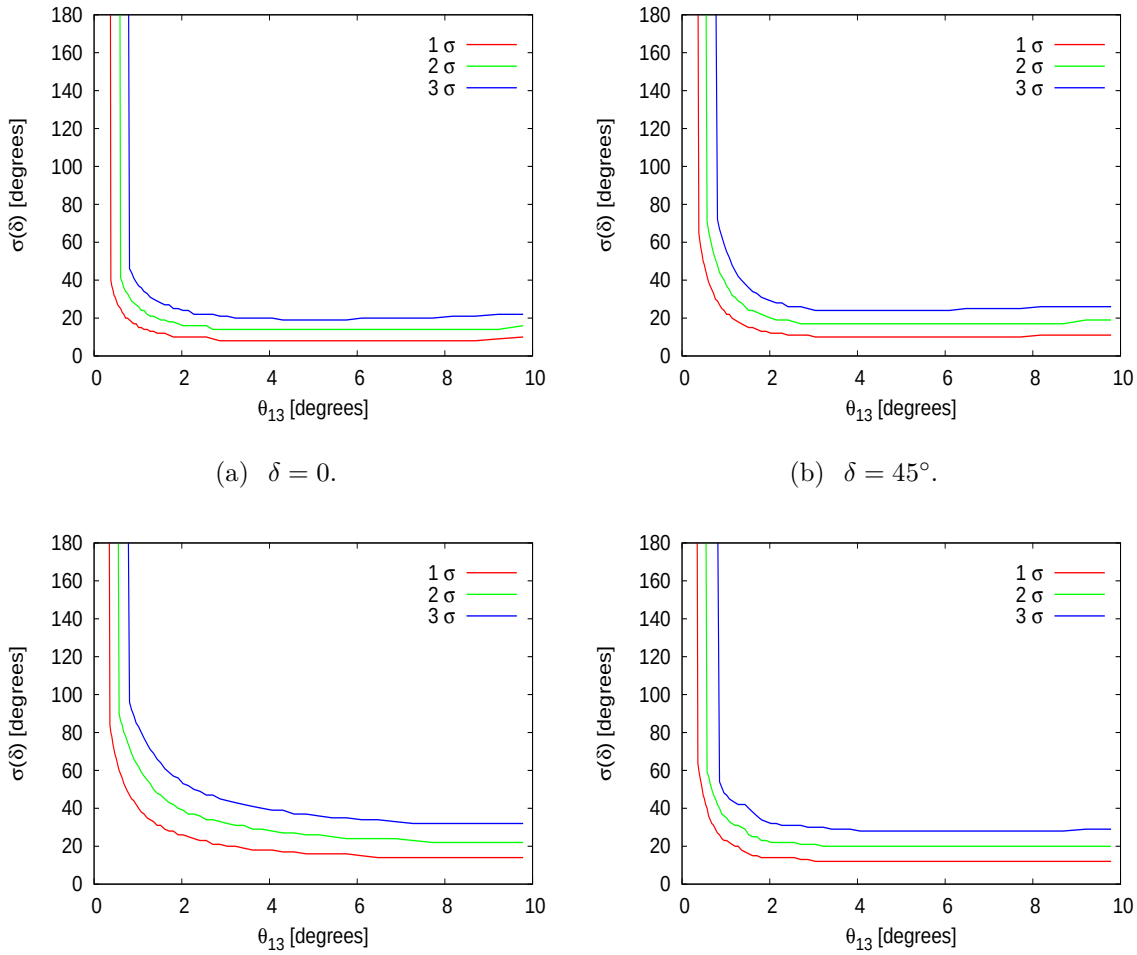

(a) $\delta = 0$.

(b) $\delta = 45°$.

(c) $\delta = 90°$.

(d) $\delta = -90°$.

Figure 6.15.: θ_{13}-dependence of the 1, 2 and 3σ errors on the measurement of δ, for a) $\delta = 0°$, b) $\delta = 45°$, c) $\delta = 90°$ and d) $\delta = -90°$.

6.5. Liquid argon detector and comparison with other experiments

Recently there has been much interest in the possibility of constructing a kiloton-scale liquid argon (LAr) detector [89, 91, 92, 93, 183]. If such a detector can be magnetised, it could be utilised in combination with a LENF and we have performed some preliminary studies to assess the potential of a 100 kton LAr detector for this experiment. As the design of large LAr detectors is still in the early stages, there are large uncertainties in the estimates for the detector performance. We assume an efficiency of 80% on all channels and 5% energy resolution for quasi-elastic events, then consider a range of values for the other parameters. In the most conservative scenario, we assume 5% systematics, 20% energy resolution for non quasi-elastic events, and backgrounds of 5×10^{-3} on the ν_μ ($\bar{\nu}_\mu$) (dis)appearance channels and 0.8 on the ν_e ($\bar{\nu}_e$) appearance channels (private communication reported in Ref. [192]). For the optimistic scenario we use values identical to the TASD: 2% systematics, 10% energy resolution for non quasi-elastic events, and backgrounds of 1×10^{-3} on the ν_μ ($\bar{\nu}_\mu$) (dis)appearance channels and 1×10^{-2} on the ν_e ($\bar{\nu}_e$) appearance channels. We find that varying the systematics, energy resolution and ν_e ($\bar{\nu}_e$) background do not play a large role in altering the results; the dominant effect comes from the variation of the ν_μ ($\bar{\nu}_\mu$) background.

In Fig. 6.16 the results of the LENF with both the TASD and the two assumptions for the LAr detector are compared with other long-baseline experiments. We show the 3σ results for θ_{13} discovery, CP discovery potential, and hierarchy sensitivity (for normal hierarchy only) as a function of $\sin^2 2\theta_{13}$ in terms of the CP fraction. In order to make a fair comparison, we have used half the flux stated in Section 6.2 for the LENF, to make it consistent with the other experiments which assume only 5 years per polarity of observation. However, we believe that the fluxes in Section 6.2 are feasible. The results from the TASD are shown by the red line and those from the LAr detector are shown by the blue band. The right-hand edge of the band corresponds to the conservative estimate of the detector performance, and the left-hand edge to the most optimistic estimate. As the optimistic scenario assumes an almost identical performance to the TASD, the left-hand edge of the blue band also corresponds to the results obtainable from a 100 kton TASD. Results from the HENF [86], wide-band beam [107, 202, 203], T2HK [103], 100γ β-beam [110], 350γ β-beam [111, 112] and four-ion β-beam [113] are also shown.

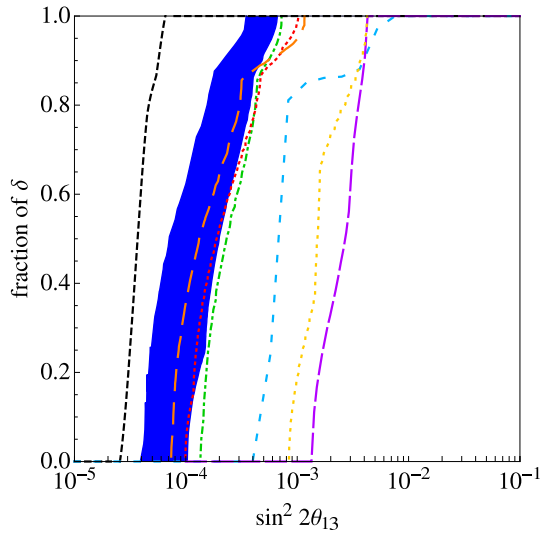

(a) θ_{13} discovery potential.

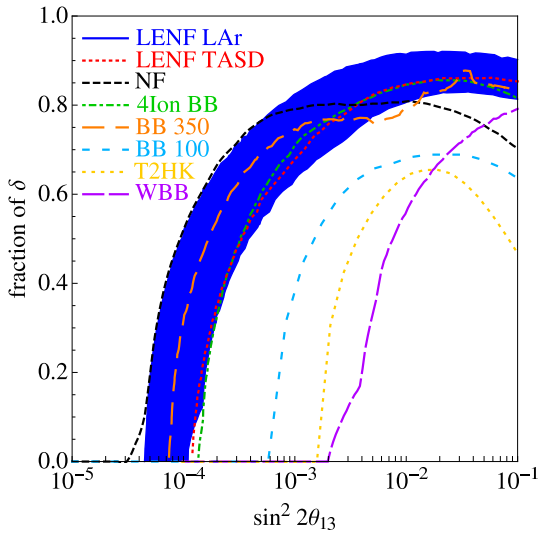

(b) CP discovery potential.

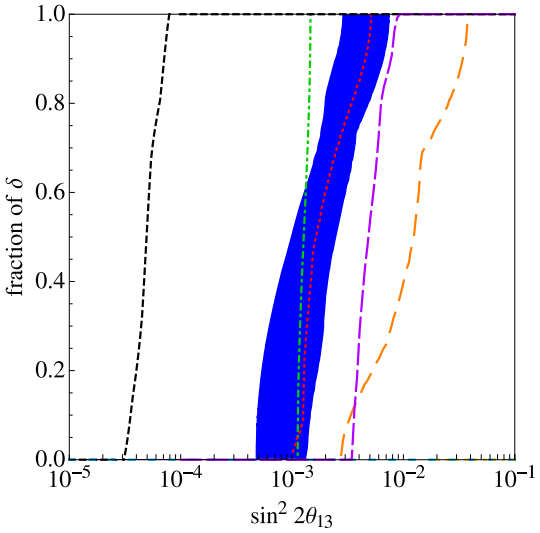

(c) Hierarchy sensitivity.

Figure 6.16.: Comparison of 3σ allowed contours for the LENF with a 20 kton TASD (red line) and 100 kton LAr detector (blue band), the HENF (black line), the wide-band beam (purple line), T2HK (yellow line) and three β-beams (green, orange, light blue lines) for a) θ_{13} discovery potential, b) CP discovery potential and c) hierarchy sensitivity.

In terms of sensitivity to θ_{13}, a conservative LENF is an order of magnitude less sensitive than the HENF, but is still competitive with the β-beam experiments, giving an approximately equal performance to the four-ion β-beam (which requires two baselines to resolve the degeneracy problem, as for the HENF). However, the performance of an

aggressive LENF setup surpasses that of all other experiments except for the HENF. For CP violation, the LENF gives remarkable results: the most optimistic setup outperforms the HENF for all values of θ_{13}, and even the most conservative setup gives a superior performance for $\sin^2 2\theta_{13} > 2 \times 10^{-3}$, again equaling that of the four-ion β-beam. For sensitivity to the mass hierarchy, the LENF gives an improvement over all other experiments apart from the higher energy setup and the four-ion β-beam with their challenging 7000 km baseline.

6.6. Baseline optimisation

The study in Section 6.2.1 indicates that for a 1300 km baseline, an energy of around 4.5 GeV is optimal. However, it is possible to be more quantitative about this, and one should also consider the possibility of using different baselines as there are several possible combinations for accelerator and detector sites over the world. Therefore the experimental sensitivities can be analysed as a function of the baseline, L, and the muon energy, E. The results for θ_{13} discovery, CP discovery and hierarchy sensitivity are shown in Fig. 6.17. The results are shown for $\theta_{13} = 2^\circ$ and the contour numbers refer to the CP-fraction that can be measured at each value of L and E (the contour at the highest energy gives the highest CP-fraction in all cases).

For θ_{13} discovery (Fig. 6.17a), it can be seen that the baseline is unimportant and that 100% coverage can be obtained for all baselines, provided that the energy is above that of the first oscillation maximum. The LENF is in the optimal region.

CP discovery (Fig. 6.17b) is more complicated as a sufficiently long baseline and high energy are required in order to determine the mass hierarchy and therefore eliminate any degeneracies, whereas low energies enhance the effect of CP violation (Section 4.1) and a shorter baseline decreases the uncertainty on matter effects (Section 4.3.2). Therefore the absolute optimal point is at a baseline of around 1800 km and 15 GeV. However there is a large region which is nearly optimal - the variation shown by the contours is only $\sim 10\%$ - and the defined LENF setup lies within this region.

The sensitivity to the hierarchy (Fig. 6.17c), as explained in Section 4.3.1, increases for longer baselines and higher energies. For this particular value of $\theta_{13} = 2^\circ$, the LENF is in the optimal region where 100% coverage can be obtained.

The results are qualitatively similar for other values of θ_{13} - the values of the displayed contours will change but the general shape remains the same.

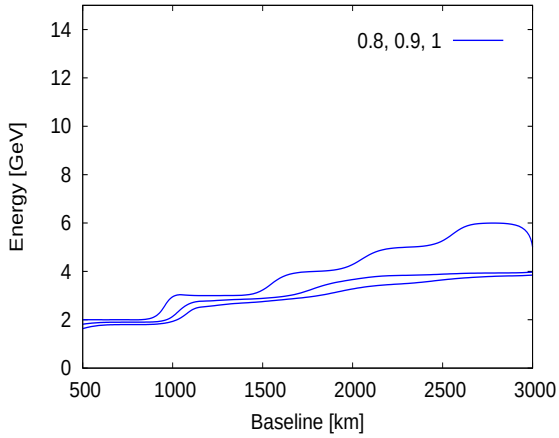

(a) θ_{13} discovery potential.

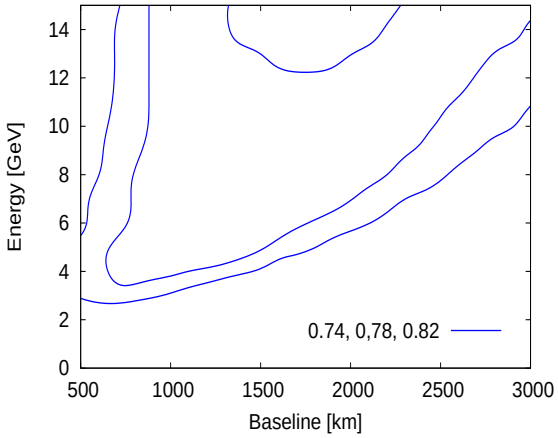

(b) CP discovery potential.

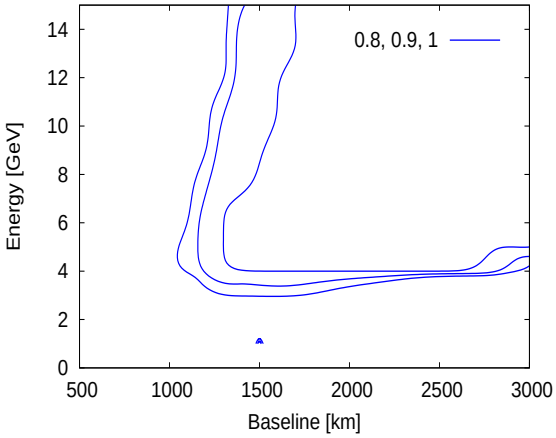

(c) Hierarchy sensitivity.

Figure 6.17.: a) θ_{13} discovery potential, b) CP discovery potential and c) hierarchy sensitivity as a function of baseline, L, and muon energy, E, for $\theta_{13} = 2°$.

6.7. Summary

We have optimised a low-energy neutrino factory setup with a baseline of 1300 km, defining a reference setup to be one with a muon energy of 4.5 GeV and 1.4×10^{21} useful muon decays per year, per polarity, running for ten years per polarity. For the detector we assume a totally active scintillating detector (TASD) with a fiducial mass of 20 kton,

energy threshold of 0.5 GeV, energy resolution of 10%, efficiency for $\mu^{\pm}$ detection of 73% below 1 GeV and 94% above, efficiency for $e^{\pm}$ detection of 37% below 1 GeV and 47% above, and a background level of 10^{-3} on the $\nu_e \to \nu_\mu$ ($\bar{\nu}_e \to \bar{\nu}_\mu$) and $\nu_\mu \to \nu_\mu$ ($\bar{\nu}_\mu \to \bar{\nu}_\mu$) channels and 10^{-2} on the $\nu_\mu \to \nu_e$ ($\bar{\nu}_\mu \to \bar{\nu}_e$) channels. We have also considered a 100 kton liquid argon detector and found that its performance would equal or surpass that of the 20 kton TASD.

We have demonstrated how improving the energy resolution and statistics improves the performance of the setup, showing that in particular high statistics play a vital role. For the defined setup, an energy threshold below the energy of the second oscillation maximum is not crucial; there are insufficient events at this energy to yield any useful information. We have also shown how the combination of golden and platinum channels could be a powerful way of resolving degeneracies, especially in the case of limited statistics. However, once realistic background levels of at least 10^{-2} are considered for the platinum channel, the improvement achieved by adding this channel is negligible. Therefore, more effort should be invested into achieving larger statistics for the golden channel than in improving the platinum channel, at least for standard physics searches.

Using our optimised setup, the LENF can have sensitivity to θ_{13} and δ for $\sin^2 2\theta_{13} > 10^{-4}$, competitive with the HENF. Sensitivity to the mass hierarchy is accessible for $\sin^2 2\theta_{13} > 10^{-3}$, better than other experiments using the same baseline. Even if the flux is halved to equal that of other long-baseline experiments, the LENF is still competitive, performing especially well for CP discovery at large values of θ_{13}. For $\sin^2 2\theta_{13} > 4\times10^{-3}$, the LENF has 100% CP coverage for hierarchy sensitivity and θ_{13} discovery, and has greater sensitivity to CP violation than the HENF. We have also studied the sensitivity to θ_{23}, finding that it is possible to exclude maximal θ_{23} at 3σ for $\theta_{23} \lesssim 43^\circ$ and $\theta_{23} \gtrsim 47^\circ$, roughly independent of θ_{13}, and to identify the octant for $\theta_{23} \lesssim 37^\circ$ and $\theta_{23} \gtrsim 53^\circ$.

We also studied the *precision* of the setup, in terms of the size of the 1σ errors on each of the parameters. We found that for $\sin^2 2\theta \gtrsim 10^{-3}$, our setup can measure θ_{13} with a $\sim 5\%$ 1σ error, θ_{23} to within $\pm \sim 2.5^\circ$ and δ to within $\sim \pm 10^\circ$. For exposures up to $\sim 6 \times 10^{23}$ kton $\times$ decays per polarity, this precision is significantly improved. The effect of non-zero systematic errors and backgrounds is to effectively halve the exposure.

Finally, we showed how setups with different values of L and E can have similar performances, finding that our current LENF setup is one that is nearly optimal.

Chapter 7.

Non-standard interactions at a low-energy neutrino factory

In the previous chapter we performed a detailed analysis of the sensitivity of the low-energy neutrino factory to the oscillation parameters θ_{13}, δ, θ_{23} and the mass hierarchy. A reference setup to measure these parameters was defined and summarised in Section 6.7. In this chapter we shall examine the potential of this setup to measure the non-standard matter interactions described in Section 2.5, with the only alteration being that we use a setup running for 5 years per polarity, as for the high-energy neutrino factory, rather than the anticipated 10 years per polarity. In particular, we will focus on the synergy between the golden and platinum channels which we will show is particularly powerful for resolving degeneracies and hence increases the sensitivity to non-standard interactions. We also consider the effect of adding a second baseline to resolve these degeneracies.

7.1. Non-standard interactions in the golden channel

To examine NSI's from the phenomenological viewpoint, the starting point is to examine the golden channel probability, including NSI's. The NSI parameters, like the the ratio $\frac{\Delta_{21}}{\Delta_{31}}$, are treated as perturbations to the Hamiltonian; in this case we take the unperturbed Hamiltonian to be that which is first order in Δ_{21}. We then follow a similar method to the one described in Section 2.4.2. A detailed derivation is described in Appendix G. After expanding the probability to second order in the small parameters s_{213}, Δ_{21} and all the NSI parameters $\varepsilon_{\alpha\beta}$ (where α, $\beta = e$, μ or τ), the only NSI terms which

remain are those proportional to $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau}$:

$$
\begin{aligned}
P_{\nu_e\to\nu_\mu} &= s^2_{213}s^2_{23}\frac{\sin^2\left(\frac{(\Delta_{31}-A)L}{2}\right)}{\left(1-\frac{A}{\Delta_{31}}\right)^2} && \text{(7.1a)}\\
&+ s_{213}c_{13}s_{212}s_{223}\frac{\Delta_{21}}{A}\sin\left(\frac{AL}{2}\right)\frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}}\cos\left(\frac{\Delta_{31}L}{2}-\delta\right) && \text{(7.1b)}\\
&+ c^2_{13}s^2_{212}c^2_{23}\left(\frac{\Delta_{21}}{A}\right)^2\sin^2\left(\frac{AL}{2}\right) && \text{(7.1c)}\\
&- 4\varepsilon_{e\mu}s_{213}c_{23}s^2_{23}\sin\left(\frac{AL}{2}\right)\frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}}\cos\left(\frac{\Delta_{31}L}{2}-\delta-\phi_{e\mu}\right) && \text{(7.1d)}\\
&- 4\varepsilon_{e\mu}\frac{\Delta_{21}}{\Delta_{31}}s_{212}c^2_{23}s^2_{23}\sin\left(\frac{AL}{2}\right)\frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}}\cos\left(\frac{\Delta_{31}L}{2}+\phi_{e\mu}\right) && \text{(7.1e)}\\
&+ 4\varepsilon^2_{e\mu}c^2_{23}s^2_{23}\sin^2\left(\frac{AL}{2}\right) && \text{(7.1f)}\\
&- 4\varepsilon_{e\tau}s_{213}c_{23}s^2_{23}\sin\left(\frac{AL}{2}\right)\frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}}\cos\left(\frac{\Delta_{31}L}{2}-\delta-\phi_{e\tau}\right) && \text{(7.1g)}\\
&+ 4\varepsilon_{e\tau}\frac{\Delta_{21}}{\Delta_{31}}s_{212}c^2_{23}s^2_{23}\sin\left(\frac{AL}{2}\right)\frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}}\cos\left(\frac{\Delta_{31}L}{2}+\phi_{e\tau}\right) && \text{(7.1h)}\\
&+ 4\varepsilon^2_{e\tau}c^2_{23}s^2_{23}\sin^2\left(\frac{AL}{2}\right), && \text{(7.1i)}
\end{aligned}
$$

where $s_{jk}=\sin\theta_{jk}$, $c_{jk}=\cos\theta_{jk}$ and $\Delta_{jk}=\frac{\Delta m^2_{jk}}{2E}$ as before. The CP conjugate channel, $P_{\bar\nu_e\to\bar\nu_\mu}$, takes a similar form with the transformations $\delta\to-\delta$, $A\to-A$, $\phi_{e\mu}\to-\phi_{e\mu}$ and $\phi_{e\tau}\to-\phi_{e\tau}$. The platinum channel, $P_{\nu_\mu\to\nu_e}$, is obtained by interchanging $\delta\to-\delta$, and the CP conjugate platinum channel, $P_{\bar\nu_\mu\to\bar\nu_e}$, by interchanging $A\to-A$, $\phi_{e\mu}\to-\phi_{e\mu}$ and $\phi_{e\tau}\to-\phi_{e\tau}$.

The fact that we are considering *matter-induced* NSI's is apparent from the dependence of all the NSI terms on $\sin\left(\frac{AL}{2}\right)$. For the LENF with its baseline of 1300 km, this factor is ~ 0.4. If it were possible to use a baseline such that $AL=\pi$ ($L=4400$ km), the NSI terms would be maximised. Incidentally, the HENF near detector baseline nearly coincides with this value. Conversely, if L is chosen such that $AL=2\pi$ - the magic baseline - then all the NSI terms vanish. Thus the magic baseline is not just a powerful tool for resolving the degeneracies between oscillation parameters, but additionally it

has the power to minimise the confusion between oscillation and NSI parameters. This, and the fact that the HENF uses a high energy which enhances matter effects, including non-standard matter effects, makes it a near-optimal setup for studying NSI's. This has been studied in depth by several authors [67, 139, 140, 204, 205, 206, 207].

From the point of view of detecting NSI's (assuming that they exist!) in neutrino oscillation experiments, the most favourable condition will be if the $\varepsilon_{\alpha\beta}$'s are of roughly the same magnitude as $\sin 2\theta_{13}$, which is the parameter that controls the amplitude of the oscillations. This would correspond to, for example, $\varepsilon_{\alpha\beta} \sim 10^{-2}$ (very large) and $\theta_{13} \sim 0.3^\circ$ (very small). This would put θ_{13} beyond the reach of the LENF, but maybe just accessible to the HENF. Were this the case, then the unusual situation would arise that the non-standard effects are of a comparable magnitude to the standard oscillation effects. The problem then is whether it is possible to distinguish the two effects from one another. In the scenario that the NSI's are very much smaller than $\sin 2\theta_{13}$ then detecting them will be very difficult, but on the bright side it means that they will not interfere with the oscillation measurements! In the scenario that the NSI's are large whilst θ_{13} is very small such that the non-standard effects are *larger* than the golden channel oscillations, distinguishing between the oscillations and NSI effects may be problematic.

Referring to Eq. (7.1) we can see that if all the phases (δ, $\phi_{e\mu}$, $\phi_{e\tau}$) are zero, optimal sensitivity is obtained via the $\bar{\nu}_e \to \bar{\nu}_\mu$ and $\bar{\nu}_\mu \to \bar{\nu}_e$ channels, when all three of the terms dependent on $\varepsilon_{e\mu}$ ((7.1d), (7.1e), (7.1f)) take the same sign. However, in the general case, the degree of sensitivity will be heavily dependent on the values of the phases. If either of the NSI parameters is exceptionally large, such that terms (7.1f) and (7.1i) are dominant, it will not be possible to distinguish between them as they look identical.

In Fig. 7.1 we show how the presence of $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau}$ modify the oscillation probability, for $\theta_{13} = 3^\circ$ ($\sin^2 2\theta_{13} \sim 10^{-2}$), $\delta = 0$ and $\varepsilon_{e\mu} = \varepsilon_{e\tau} = 0.01$. Even in this scenario where the NSI's are very large, distinguishing them from the standard oscillation will not be easy as all the spectra are very similar. The largest discrepancy occurs at the first oscillation peak so this will be the most important region for the detection of NSI's, as for standard oscillations. The second peak cannot realistically be used, especially when considering that the flux at these energies is very low. Therefore we shall consider only the first oscillation peak for our discussion.

If both the NSI phases, $\phi_{e\mu}$ and $\phi_{e\tau}$, are zero, $\varepsilon_{e\mu}$ is more prominent than $\varepsilon_{e\tau}$ - compare the red and green lines to see that the green line ($\varepsilon_{e\mu} = 0.01$) is slightly more

distinct from the black line than the red line ($\varepsilon_{e\tau} = 0.01$). This is because for our chosen values of $\theta_{13} = 3°$, $\varepsilon_{e\mu} = 0.01$ and $\varepsilon_{e\tau} = 0.01$, the NSI terms (7.1d), (7.1e), (7.1g) and (7.1h) are larger than (7.1f) and (7.1i). (7.1d) and (7.1e) interfere *constructively* whilst (7.1g) and (7.1h) interfere *destructively*. The opposite effect occurs if we choose $\phi_{e\tau}$ and $\phi_{e\mu} = \frac{\pi}{2}$ (blue and purple) with $\delta = 0$, because the cosine factor in (7.1d) and (7.1g) is simply $\sin\left(\frac{\Delta_{31}L}{2}\right)$ whereas the factor in (7.1e) and (7.1h) is $-\sin\left(\frac{\Delta_{31}L}{2}\right)$. So in this case, (7.1g) and (7.1h) interfere constructively whereas (7.1d) and (7.1e) interfere destructively. Therefore $\varepsilon_{e\tau}$ will be easier to detect in this case.

Clearly, the situation is very complicated, especially if *both* $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau}$ are non-zero and have non-zero phases. Here we have only considered a simple case where all the phases are set to zero and we study the correlations between the NSI and oscillation parameters for this case.

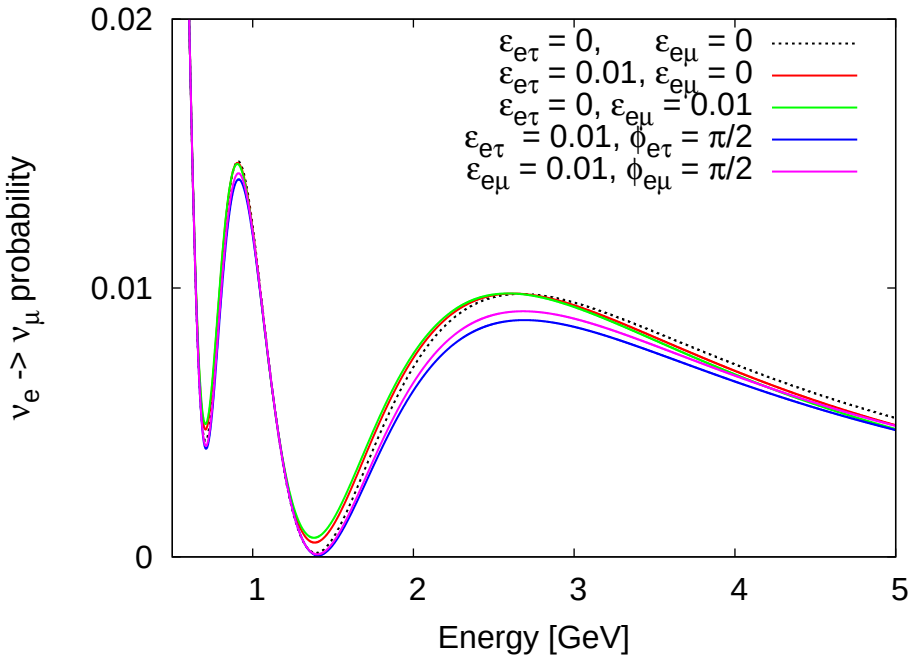

Figure 7.1.: Golden channel probability including the NSI parameters $\varepsilon_{e\tau}$ and $\varepsilon_{e\mu}$, for $\theta_{13} = 3°$, $\delta = 0$, $\varepsilon_{e\tau}$ and $\varepsilon_{e\mu} = 0$ and 0.01, and $\phi_{e\tau}$ and $\phi_{e\mu} = 0$ and $\frac{\pi}{2}$.

From a phenomenological perspective, there are two main items to address:

- The NSI-oscillation degeneracy: the problem of degeneracies in neutrino oscillation experiments has been well established and was described in Section 4.5. When taking into account the possibility of NSI's, the parameter space is vastly expanded and so the problem is magnified. Of particular concern is the question of whether NSI's can interfere with the measurement of the standard oscillation parameters.

For example, supposing that $\delta = 0$ so CP is conserved, then we want to know whether it is possible for NSI's to mimic CP violation and fake a non-zero value of δ? More precisely, how does expanding the parameter space by marginalising over NSI parameters affect the sensitivity to the oscillation parameters?

Note that this is not a problem for the HENF where the magic baseline comes to the rescue - recall that the magic baseline is such that $\sin\left(\frac{AL}{2}\right) = 0$. This not only eliminates the CP and solar terms from the standard oscillation probability, but also *all* the NSI terms vanish at this baseline. This then eliminates any confusion between θ_{13}, the mass hierarchy and the NSI parameters, making the oscillation sensitivity of the HENF setup virtually immune to interference from NSI's.

- How precisely can the NSI parameters be measured? Current bounds on $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau}$ are $O(1)$ [78, 79]. Any next-generation experiment should aim to improve on these bounds by at least an order of magnitude.

We will consider both these points with particular emphasis on how the platinum channels significantly affect the results. It was found in the previous chapter that the platinum channel is redundant for oscillation measurements if the statistics are sufficiently high, which indicates that the platinum channel does not add anything that cannot be accomplished by other design optimisations. However, we shall see that this is not the case for NSI's.

7.2. Simulation details - MonteCUBES

The computational techniques used in the simulations of standard neutrino oscillation were described in Section 3.6.1. For the purposes of including NSI's into the simulations, GLoBES by itself is not the best tool. The reason for this is that the minimisation algorithm used by GLoBES becomes more inefficient as the dimension of the parameter space is increased, such as by the inclusion of NSI's. If the entire NSI matrix (see Eq. 2.34) were to be included into the simulations, this would add another nine parameters (six moduli and three phases) into the parameter space, on top of the existing six oscillation parameters and the matter density. The minimisation technique used by GLoBES is based on setting up an N-dimensional grid for an N-dimensional parameter space, and making n samplings per parameter. Therefore $O(n^N)$ evaluations are performed to establish the number of events and χ^2 value at each point, which means that the computation time

scales exponentially with the number of parameters, N. So adding an additional nine NSI parameters to the seven oscillation parameters is not very time-efficient! The only option to perform an analysis in a reasonable amount of time is to only include one, or at most two, NSI parameters at a time into the simulations. In most cases this gives a good approximation to the real situation, but in order to be really rigourous it is necessary to include correlations between *all* the NSI and oscillation parameters.

The MonteCUBES (Monte Carlo Utility Based Experiment Simulator) software package [208] was designed precisely for this purpose. It is a plug-in to GLoBES which allows for the GLoBES experiment files to be used, whilst implementing a more efficient minimisation method by using Markov Chain Monte Carlo algorithms. These scale polynomially, rather than exponentially, with the number of parameters. Additionally it enables the modified oscillation probabilities, including NSI parameters, to be automatically defined and calculated exactly, rather than having to manually reset and specify them in GLoBES.

Although the computational technique is different, the statistical analysis used in these NSI studies is essentially the same as that described in Section 3.6.2. We have performed only two-parameter fits in this chapter in order to illustrate the correlations between the NSI and oscillation parameters. The main focus is the impact of the platinum channel, so we once again use the notation where *Scenario 1* denotes a setup which measures only the ν_μ and $\bar{\nu}_\mu$ appearance (golden) and disappearance channels, whereas *Scenario 2* is the setup which can also detect the platinum channels. The usual 68%, 90% and 95% contours are shown.

7.3. Degeneracies between oscillation and NSI parameters

7.3.1. $\varepsilon_{e\mu} = 0$, $\varepsilon_{e\tau} = 0$

In Figs. 7.2a and 7.2b we show the 68%, 90% and 95% allowed regions in the $\theta_{13} - \delta$ plane for true values of $\theta_{13} = 5^\circ$ and $\delta = 0$, for the case when we fit to a *standard oscillation model only* i.e. we set all NSI parameters to zero and *do not marginalise over them*. In Fig. 7.2b we show the results when marginalisation over all oscillation *and* NSI parameters is performed. We find that for Scenario 1 (blue dashed lines), the precision

of the measurement of θ_{13} is significantly weakened. The width of the 95% contour is approximately doubled by including a marginalisation over the NSI parameters. The precision of the CP measurement is also weakened, but to a lesser extent. However we find that the inclusion of the platinum channels makes this setup far more robust, with the inclusion of NSI's having barely any effect.

We can understand this by comparing the oscillation probabilities for the golden and platinum channels around the region of the first oscillation maximum where $(\Delta_{31}-A)L = \pi$. In the case of large θ_{13}, the golden channel probability is given approximately by:

$$P_{\nu_e\to\nu_\mu} \simeq \frac{s^2_{213}s^2_{23}}{\left(1-\frac{A}{\Delta_{31}}\right)^2} \tag{7.2a}$$

$$+ \frac{s_{213}c_{13}s_{212}s_{223}}{1-\frac{A}{\Delta_{31}}}\frac{\Delta_{21}}{A}\sin\left(\frac{AL}{2}\right)\cos\left(\frac{\pi+AL}{2}-\delta\right) \tag{7.2b}$$

$$- 4\varepsilon_{e\mu}\frac{s_{213}c_{23}s^2_{23}}{1-\frac{A}{\Delta_{31}}}\sin\left(\frac{AL}{2}\right)\cos\left(\frac{\pi+AL}{2}-\delta+\phi_{e\mu}\right) \tag{7.2c}$$

$$- 4\varepsilon_{e\tau}\frac{s_{213}c_{23}s^2_{23}}{1-\frac{A}{\Delta_{31}}}\sin\left(\frac{AL}{2}\right)\cos\left(\frac{\pi+AL}{2}-\delta+\phi_{e\tau}\right). \tag{7.2d}$$

The confusion between the NSI's and θ_{13} arises because the NSI terms, (7.2c) and (7.2d), have a similar energy dependence to the atmospheric term (since $\frac{A}{\Delta_{31}}$ is small in our setup, $\left(1-\frac{A}{\Delta_{31}}\right)^2 \simeq \left(1-\frac{2A}{\Delta_{31}}\right)$). Additionally, our setup has *no sensitivity* to the NSI phases (see Fig. 7.3); therefore it is possible to fit to a larger than true value of θ_{13} by choosing $\varepsilon_{e\mu}$ and/ or $\varepsilon_{e\tau}$ to be non-zero, and the NSI phases to take values such that terms (7.2c) and (7.2d) are negative. Alternatively, it is possible to fit to a smaller than true value of θ_{13} by choosing non-zero NSI's with phases that make terms (7.2c) and (7.2d) positive. δ is affected less because the CP term has a *different* energy dependence to the NSI terms which therefore makes it harder to mimic.

Consider now the CP conjugate platinum channel:

$$
\begin{aligned}
P_{\bar{\nu}_\mu \to \bar{\nu}_e} &\simeq \frac{s_{213}^2 s_{23}^2}{\left(1 - \frac{A}{\Delta_{31}}\right)^2} && \text{(7.3a)} \\
&+ \frac{s_{213} c_{13} s_{212} s_{223}}{1 - \frac{A}{\Delta_{31}}} \frac{\Delta_{21}}{A} \sin\left(\frac{AL}{2}\right) \cos\left(\frac{\pi + AL}{2} - \delta\right) && \text{(7.3b)} \\
&+ 4\varepsilon_{e\mu} \frac{s_{213} c_{23} s_{23}^2}{1 - \frac{A}{\Delta_{31}}} \sin\left(\frac{AL}{2}\right) \cos\left(\frac{\pi + AL}{2} - \delta - \phi_{e\mu}\right) && \text{(7.3c)} \\
&- 4\varepsilon_{e\tau} \frac{s_{213} c_{23} s_{23}^2}{1 - \frac{A}{\Delta_{31}}} \sin\left(\frac{AL}{2}\right) \cos\left(\frac{\pi + AL}{2} - \delta - \phi_{e\tau}\right). && \text{(7.3d)}
\end{aligned}
$$

Note the sign differences in terms (7.3c) and (7.3d) compared to the golden channel. Algebraically, if the two probabilities (7.2) and (7.3) are added together, then for $\phi_{e\mu} = \phi_{e\tau} = 0$ the sum is simply the atmospheric and CP terms, identical to the situation where there are no NSI's. So, assuming a sufficient number of platinum events, this helps to explain why the golden and platinum channels together give a sensitivity to θ_{13} and δ which is identical to that obtained in the case where there are no NSI's. The situation is different if θ_{13} is small. In this case different terms are dominant and the golden channel probability is

$$
\begin{aligned}
P_{\nu_e \to \nu_\mu} &\simeq \frac{s_{213} c_{13} s_{212} s_{223}}{1 - \frac{A}{\Delta_{31}}} \frac{\Delta_{21}}{A} \sin\left(\frac{AL}{2}\right) \cos\left(\frac{\pi + AL}{2} - \delta\right) && \text{(7.4a)} \\
&+ c_{13} s_{212}^2 c_{23}^2 \left(\frac{\Delta_{21}}{A}\right)^2 \sin^2\left(\frac{AL}{2}\right) && \text{(7.4b)} \\
&- 4\varepsilon_{e\mu} \frac{s_{213} c_{23} s_{23}^2}{1 - \frac{A}{\Delta_{31}}} \sin\left(\frac{AL}{2}\right) \cos\left(\frac{\pi + AL}{2} - \delta + \phi_{e\mu}\right) && \text{(7.4c)} \\
&- 4\varepsilon_{e\mu} \frac{\Delta_{21}}{\Delta_{31}} \frac{s_{212} c_{23}^2 s_{23}^2}{1 - \frac{A}{\Delta_{31}}} \sin\left(\frac{AL}{2}\right) \cos\left(\frac{\pi + AL}{2} + \phi_{e\tau}\right) && \text{(7.4d)} \\
&- 4\varepsilon_{e\tau} \frac{s_{213} c_{23} s_{23}^2}{1 - \frac{A}{\Delta_{31}}} \sin\left(\frac{AL}{2}\right) \cos\left(\frac{\pi + AL}{2} - \delta + \phi_{e\tau}\right) && \text{(7.4e)} \\
&+ 4\varepsilon_{e\tau} \frac{\Delta_{21}}{\Delta_{31}} \frac{s_{212} c_{23}^2 s_{23}^2}{1 - \frac{A}{\Delta_{31}}} \sin\left(\frac{AL}{2}\right) \cos\left(\frac{\pi + AL}{2} + \phi_{e\tau}\right), && \text{(7.4f)}
\end{aligned}
$$

whereas the CP conjugate platinum channel is

$$
\begin{aligned}
P_{\bar{\nu}_\mu\to\bar{\nu}_e} &\simeq \frac{s_{213}c_{13}s_{212}s_{223}}{1-\frac{A}{\Delta_{31}}}\frac{\Delta_{21}}{A}\sin\left(\frac{AL}{2}\right)\cos\left(\frac{\pi+AL}{2}-\delta\right) && \text{(7.5a)}\\
&+ c_{13}s^2_{212}c^2_{23}\left(\frac{\Delta_{21}}{A}\right)^2\sin^2\left(\frac{AL}{2}\right) && \text{(7.5b)}\\
&- 4\varepsilon_{e\mu}\frac{s_{213}c_{23}s^2_{23}}{1-\frac{A}{\Delta_{31}}}\sin\left(\frac{AL}{2}\right)\cos\left(\frac{\pi+AL}{2}-\delta-\phi_{e\mu}\right) && \text{(7.5c)}\\
&+ 4\varepsilon_{e\mu}\frac{\Delta_{21}}{\Delta_{31}}\frac{s_{212}c^2_{23}s^2_{23}}{1-\frac{A}{\Delta_{31}}}\sin\left(\frac{AL}{2}\right)\cos\left(\frac{\pi+AL}{2}-\phi_{e\tau}\right) && \text{(7.5d)}\\
&+ 4\varepsilon_{e\tau}\frac{s_{213}c_{23}s^2_{23}}{1-\frac{A}{\Delta_{31}}}\sin\left(\frac{AL}{2}\right)\cos\left(\frac{\pi+AL}{2}-\delta-\phi_{e\tau}\right) && \text{(7.5e)}\\
&- 4\varepsilon_{e\tau}\frac{\Delta_{21}}{\Delta_{31}}\frac{s_{212}c^2_{23}s^2_{23}}{1-\frac{A}{\Delta_{31}}}\sin\left(\frac{AL}{2}\right)\cos\left(\frac{\pi+AL}{2}-\phi_{e\tau}\right). && \text{(7.5f)}
\end{aligned}
$$

There are many more relevant terms which make the spectrum more complicated. Once again, algebraically summing the probabilities results in only the oscillation terms if $\phi_{e\mu} = \phi_{e\tau} = 0$, so theoretically the combination of golden and platinum channels should still help to resolve degeneracies. However, in Figs. 7.2c and 7.2d we show the same as Figs. 7.2a and 7.2b but for a true value of $\theta_{13} = 1^\circ$. From here it can be seen that the impact of the platinum channel is much reduced. This is in direct contrast to the standard oscillation case where, from Fig. 6.7, we saw that for $\theta_{13} = 1^\circ$, the platinum channel has a large effect (for low statistics). The plots for larger values of θ_{13} were not shown but it was found that the effect of the platinum channels *decreases* for large values of θ_{13}. This is because statistics are the limiting factor for oscillation measurements, and this is automatically increased by a large value of θ_{13} (larger oscillations and therefore more signal events). This sensitivity to statistics was quantitatively studied by the authors of Ref. [209], for standard oscillation measurements. Whilst statistics are crucial in this case, it is not so for NSI's. We have compared the performance of the LENF when running for 5 years per polarity (the running time used in all the simulations in this chapter) and when running for 10 years per polarity (the original estimated flux), and find that *the sensitivity to NSI's is barely affected by a doubling of statistics.* This implies that, running for 5 years per polarity, the LENF sensitivity to NSI's is *not* limited by the statistical error but by the correlations between the oscillation and NSI parameters.

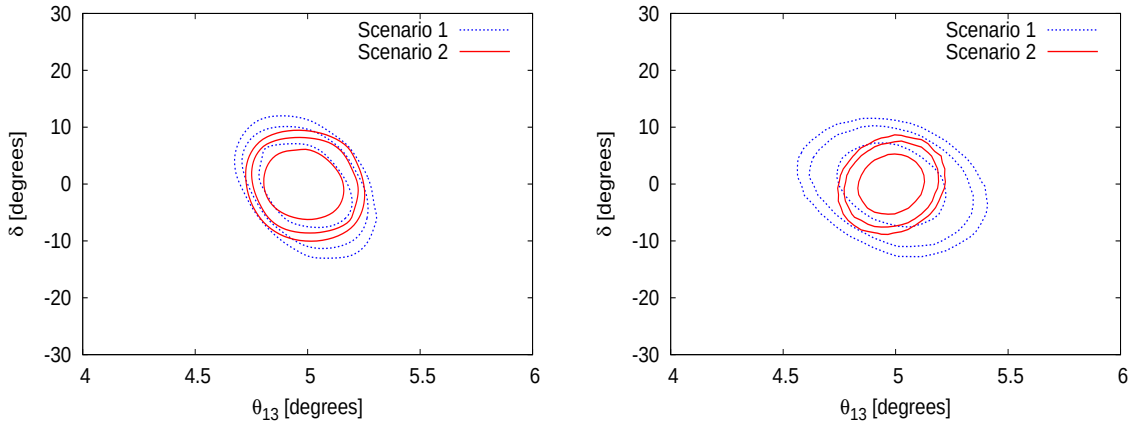

(a) $\theta_{13} = 5°$, marginalisation over oscillation parameters only.

(b) $\theta_{13} = 5°$, marginalisation over all oscillation and NSI parameters.

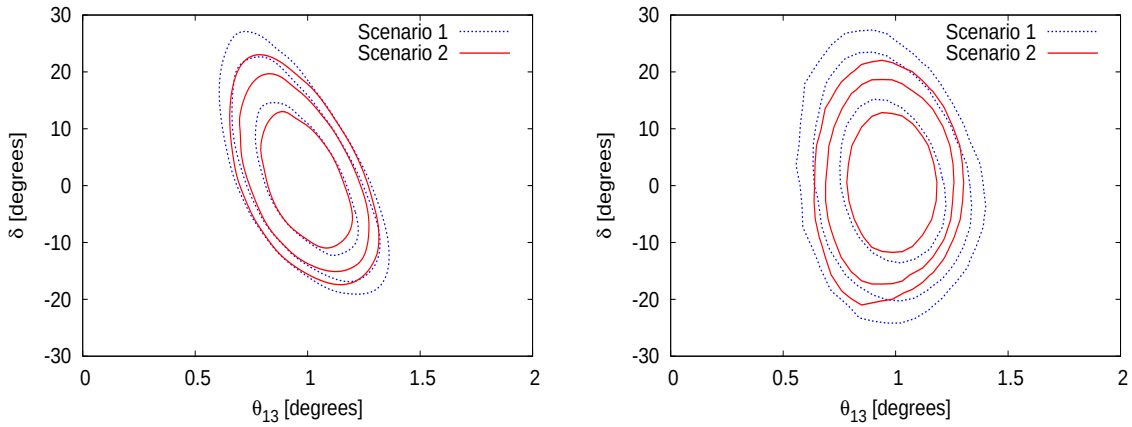

(c) $\theta_{13} = 1°$, marginalisation over oscillation parameters only.

(d) $\theta_{13} = 1°$, marginalisation over all oscillation and NSI parameters.

Figure 7.2.: Sensitivity to θ_{13} and δ when marginalising over only oscillation parameters (left column) and all oscillation and NSI parameters (right column), for true values of $\delta = 0$, $\varepsilon_{e\mu} = \varepsilon_{e\tau} = 0$ and $\theta_{13} = 5°$ (top row) and $1°$ (bottom row).

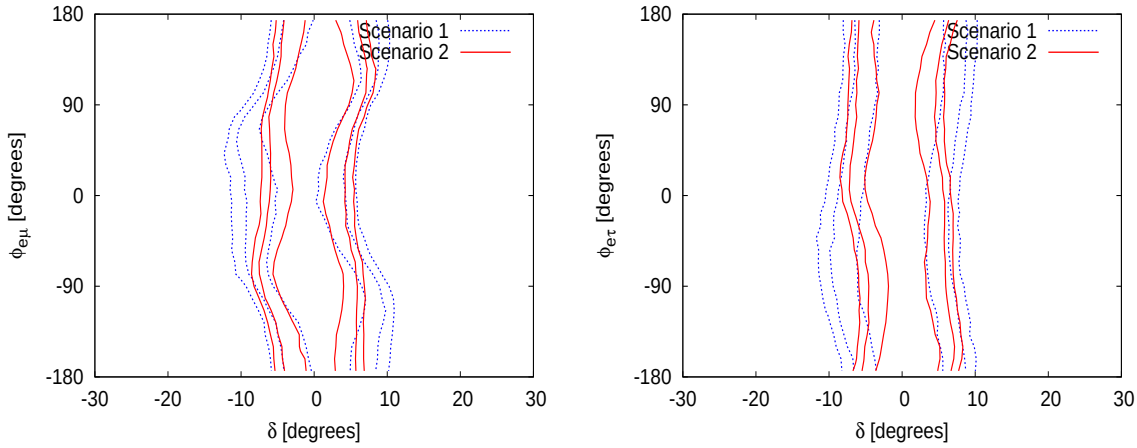

(a) Sensitivity to $\phi_{e\mu}$. (b) Sensitivity to $\phi_{e\tau}$.

Figure 7.3.: Sensitivity to $\phi_{e\mu}$ (left) and $\phi_{e\tau}$ (right) for true values of $\theta_{13} = 5°$, $\delta = 0$ and $\phi_{e\mu} = \phi_{e\tau} = 0$. There is no sensitivity to either of the phases.

7.3.2. $\varepsilon_{e\mu} \neq 0$, $\varepsilon_{e\tau} \neq 0$

Next we show how *non-zero* NSIs impact upon the measurement of θ_{13} and δ. In Fig. 7.4 we compare the results obtained when $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau}$ are both zero (blue dotted lines), to the case when either or both are non-zero (red solid lines), for Scenario 1 (left-hand column) and Scenario 2 (right-hand column). Once again we are marginalising over all oscillation and NSI parameters. We choose $\varepsilon_{e\mu} = 0.01$ and $\varepsilon_{e\tau} = 0$ with $\theta_{13} = 5°$. Similar results are obtained for $\varepsilon_{e\mu} = 0$ and $\varepsilon_{e\tau} = 0.01$. In this case, the effect of non-zero NSI's is minimal, with only a $\sim 10\%$ decrease in precision on the θ_{13} and δ measurements for Scenario 1, and an even smaller change for Scenario 2. In Fig. 7.5, we show a more challenging case, where both $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau} = 0.01$ and θ_{13} is small (1°). Now there is a much larger impact on the precision of the measurements, for both scenarios. Interestingly, only the *upper* error on θ_{13} is affected which at least means that there is no danger of faking a zero value of θ_{13} if θ_{13} is really non-zero. This occurs because for small θ_{13}, the dominant θ_{13}-dependent term is the CP term. To fake a smaller value of θ_{13} requires cancellation of this term, which is difficult because of its unique energy dependence. However, it is still possible to fake a larger value of θ_{13} because the NSI terms can mimic the atmospheric term.

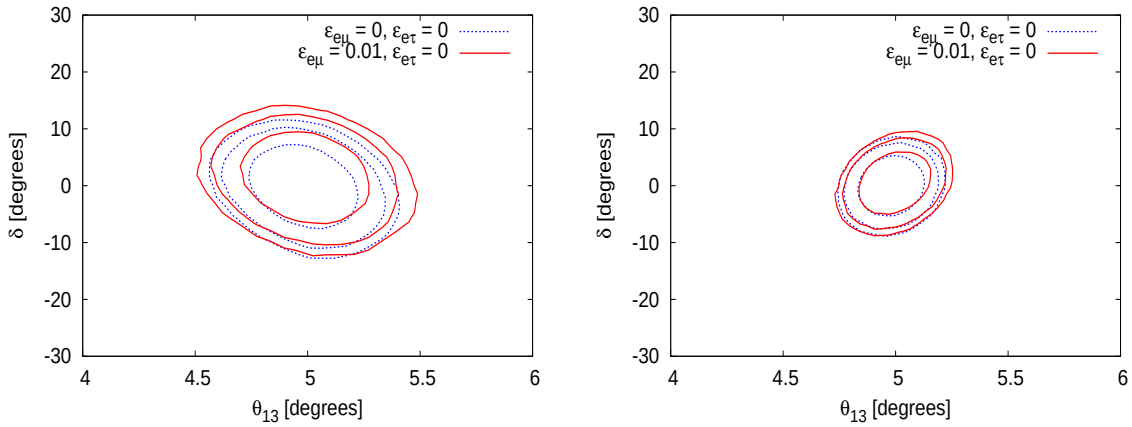

(a) Scenario 1. (b) Scenario 2.

Figure 7.4.: Effect of non-zero $\varepsilon_{e\mu}$ (red solid lines) on the measurement of θ_{13} and δ, compared to $\varepsilon_{e\mu} = 0$ (blue dotted lines), for a) Scenario 1 and b) Scenario 2, for $\theta_{13} = 5°$ and $\delta = 0°$.

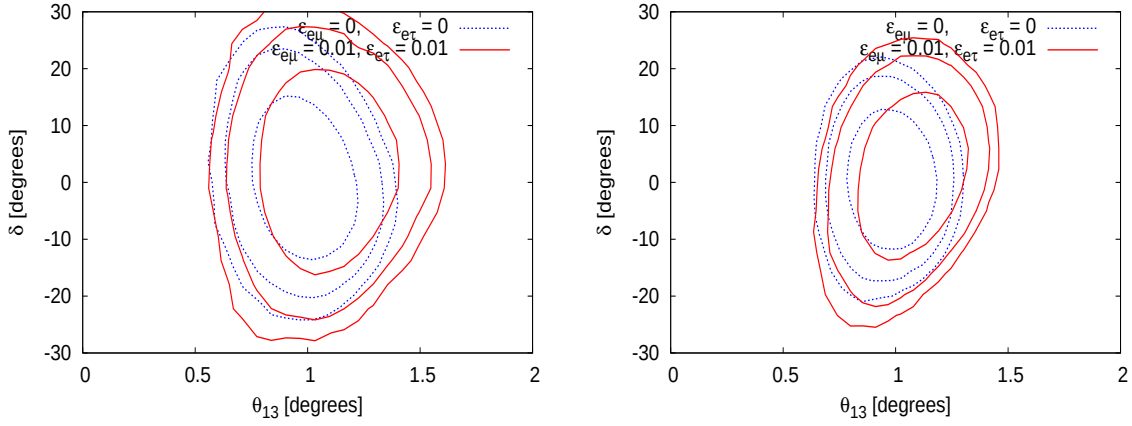

(a) Scenario 1. (b) Scenario 2.

Figure 7.5.: Effect of non-zero $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau}$ (red solid lines) on the measurement of θ_{13} and δ, compared to $\varepsilon_{e\mu} = \varepsilon_{e\tau} = 0$ (blue dotted lines), for a) Scenario 1 and b) Scenario 2, for $\theta_{13} = 5°$ and $\delta = 0°$.

7.3.3. Precision measurement of NSI's

Here we study the sensitivity to the NSI parameters themselves, rather than how their presence affects the sensitivity to the oscillation parameters. We simulate values of $\theta_{13} = 1^\circ$ and 5°, and $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau} = 0$ or 0.01, comparing the performances of Scenarios 1 and 2.

In Fig. 7.6 we have used $\varepsilon_{e\mu} = 0$ and $\varepsilon_{e\tau} = 0$ for $\theta_{13} = 1^\circ$ (top row) and 5° (bottom row). We can obtain bounds of the order of 10^{-2} on both NSI parameters. For $\theta_{13} = 1^\circ$, Scenario 2 (red solid lines) gives a bound which is $\sim 10\%$ better than that of Scenario 1 (blue dotted lines) at the 95% CL. For $\theta_{13} = 5^\circ$, the difference is $\sim 30\%$. For $\theta_{13} = 5^\circ$, the sensitivity to $\varepsilon_{e\mu}$ is greater than for $\varepsilon_{e\tau}$ because of the cancellation between the $\varepsilon_{e\tau}$ terms mentioned earlier, which is most prominent at this value of θ_{13}. For $\theta_{13} = 1^\circ$ the cancellation is less severe and so the sensitivity to both NSI parameters is roughly equal.

In Fig. 7.7 we use $\varepsilon_{e\mu} = 0.01$ and $\varepsilon_{e\tau} = 0$ and show the sensitivity to θ_{13} and $\varepsilon_{e\mu}$. We find (not shown) that the sensitivity to $\varepsilon_{e\tau}$ is not significantly affected i.e. the sensitivity to $\varepsilon_{e\tau}$ is not altered by a non-zero value of $\varepsilon_{e\mu}$. In Fig. 7.8 we use $\varepsilon_{e\mu} = 0$ and $\varepsilon_{e\tau} = 0.01$ and show the sensitivity to $\varepsilon_{e\tau}$. Similarly, we have found that the sensitivity to $\varepsilon_{e\mu}$ is not affected.

For $\theta_{13} = 1^\circ$, $\varepsilon_{e\mu} = 0.01$ and $\varepsilon_{e\tau} = 0$, we can nearly exclude $\varepsilon_{e\mu} = 0$ at 68% confidence; for $\theta_{13} = 5^\circ$ we can do this at 90% confidence. However the sensitivity to non-zero $\varepsilon_{e\tau}$ is not as good, and it is not possible to exclude $\varepsilon_{e\tau} = 0$.

In Fig. 7.9 we now consider a similar situation to Fig. 6.7, when we simulate a hypothetically perfect platinum channel (same efficiency as for the golden channel, negligible background). We use $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau} = 0$ and show the sensitivity to the NSI parameters, as in Fig. 7.6. In this case we find that the platinum channel helps even for small values of θ_{13}. Since the performance is enhanced by a higher efficiency *and* minimal background, but *not* by an increase in statistics alone, this implies that the platinum channel requires a *critical signal to background ratio* to be effective.

7.4. Alternative baselines

Now we consider how a different baseline affects the sensitivity to NSI's. Matter NSI's are enhanced by high energies - their energy dependence is the same as the atmospheric

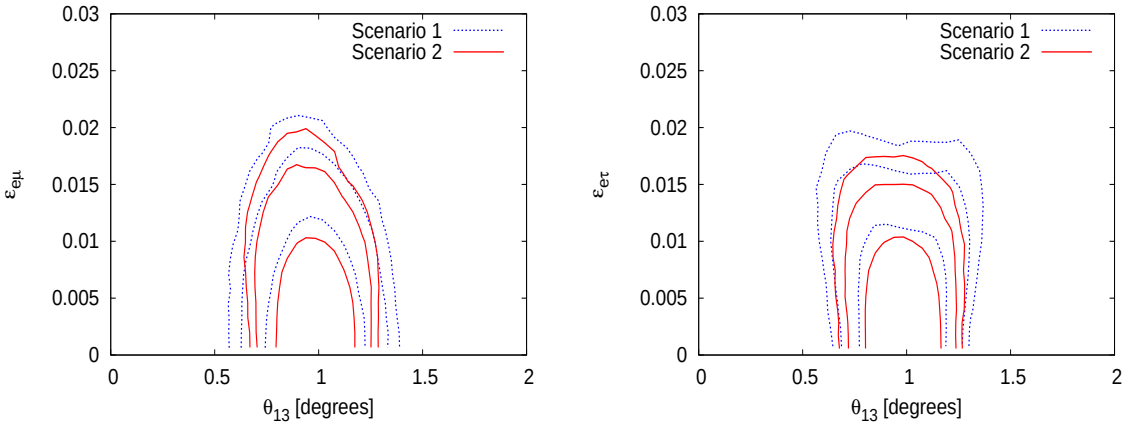

(a) $\theta_{13} = 1°$, sensitivity to $\varepsilon_{e\mu}$.

(b) $\theta_{13} = 1°$, sensitivity to $\varepsilon_{e\tau}$.

(c) $\theta_{13} = 5°$, sensitivity to $\varepsilon_{e\mu}$.

(d) $\theta_{13} = 5°$, sensitivity to $\varepsilon_{e\tau}$.

Figure 7.6.: Sensitivity to $\varepsilon_{e\mu}$ (left) and $\varepsilon_{e\tau}$ (right) for true values of $\varepsilon_{e\mu} = \varepsilon_{e\tau} = 0$ and $\theta_{13} = 1°$ (top) and $5°$ (bottom) and $\delta = 0°$.

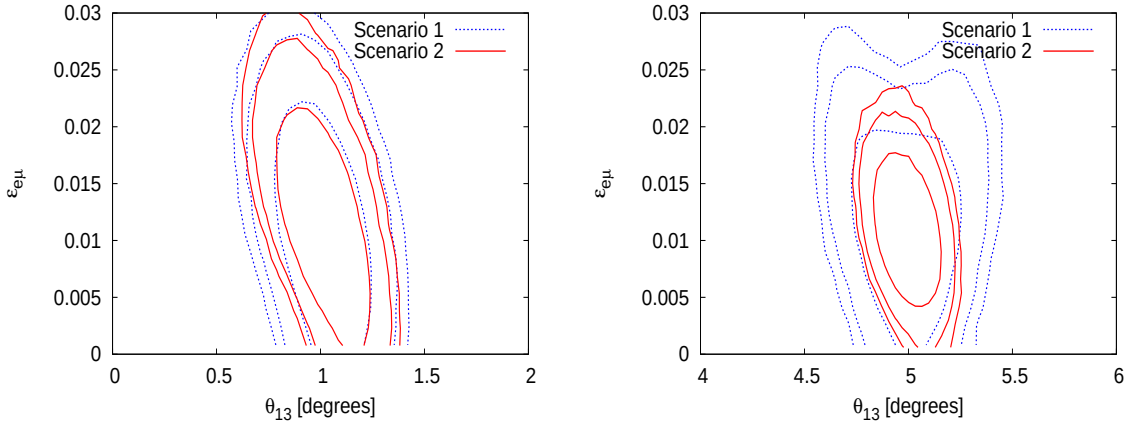

(a) $\theta_{13} = 1°$, sensitivity to $\varepsilon_{e\mu}$. (b) $\theta_{13} = 5°$, sensitivity to $\varepsilon_{e\mu}$.

Figure 7.7.: Sensitivity to $\varepsilon_{e\mu}$ for true values of $\varepsilon_{e\mu} = 0.01$, $\varepsilon_{e\tau} = 0$, $\delta = 0°$ and a) $\theta_{13} = 1°$ and b) 5°.

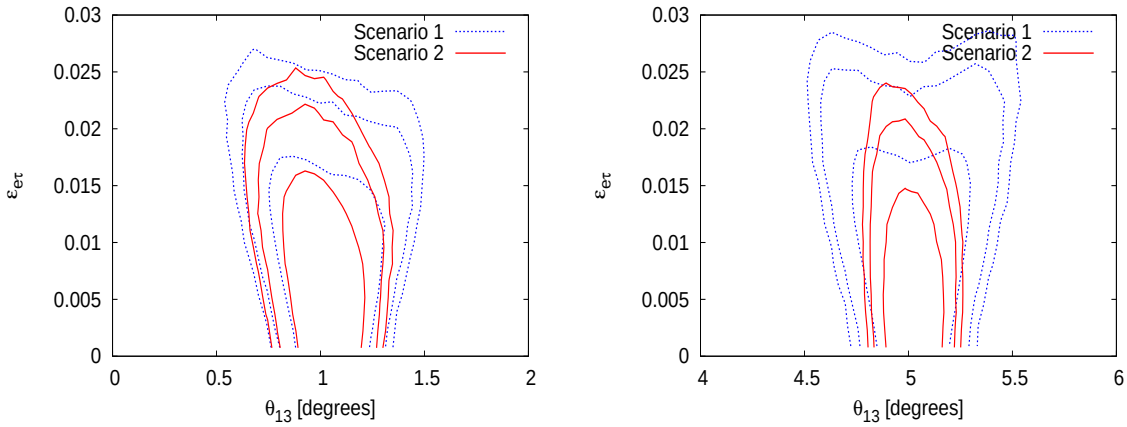

(a) $\theta_{13} = 1°$, sensitivity to $\varepsilon_{e\tau}$. (b) $\theta_{13} = 5°$, sensitivity to $\varepsilon_{e\tau}$.

Figure 7.8.: Sensitivity to $\varepsilon_{e\tau}$ for true values of $\varepsilon_{e\mu} = 0$, $\varepsilon_{e\tau} = 0.01$, $\delta = 0°$ and a) $\theta_{13} = 1°$ and b) 5°.

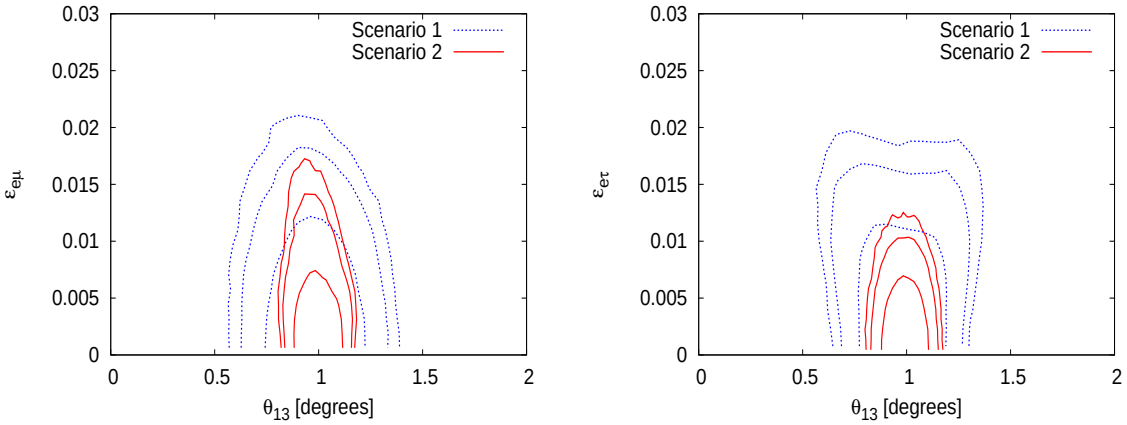

(a) $\theta_{13} = 1°$, sensitivity to $\varepsilon_{e\mu}$ (b) $\theta_{13} = 1°$, sensitivity to $\varepsilon_{e\tau}$

Figure 7.9.: Sensitivity to a) $\varepsilon_{e\mu}$ and b) $\varepsilon_{e\tau}$ for true values of $\varepsilon_{e\mu} = \varepsilon_{e\tau} = 0$, $\theta_{13} = 1°$ and $\delta = 0°$, with a platinum channel with hypothetically perfect performance.

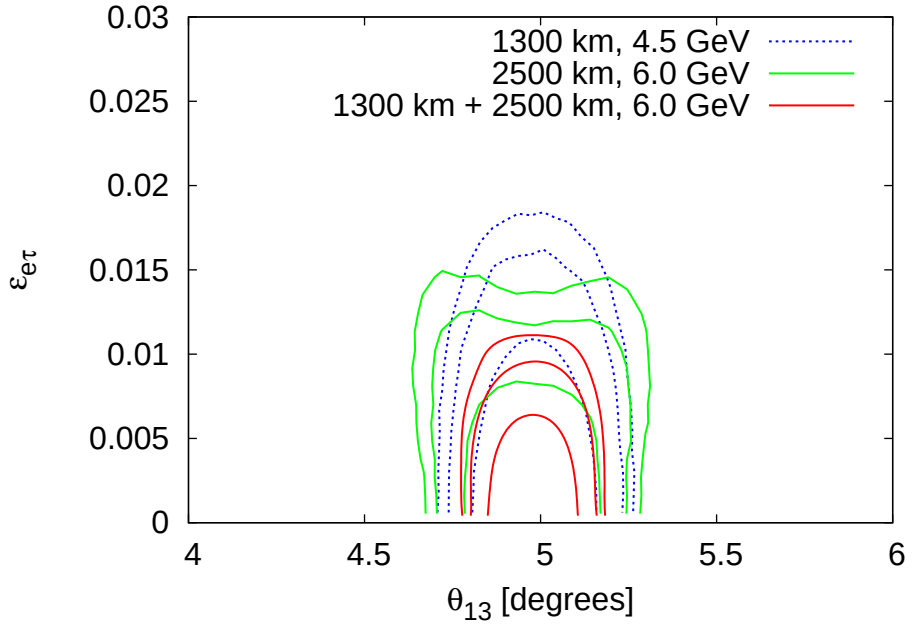

Figure 7.10.: Sensitivity to $\varepsilon_{e\tau}$ from different combinations of baselines and muon energies (1300km and 4.5 GeV, 2500 km and 6.0 GeV, two baselines at 1300 km and 2500 km and 6.0 GeV energy), using true values of $\theta_{13} = 5°$, $\delta = 0°$ and $\varepsilon_{e\mu} = \varepsilon_{e\tau} = 0$.

term (see Equation (7.1)) - and by long baselines, since the more matter there is, the larger the effect of the NSI's. In Fig. 7.10 we show the sensitivity to θ_{13} and $\varepsilon_{e\tau}$ for $\theta_{13} = 5°$ and $\varepsilon_{e\tau} = 0$ for our original LENF baseline of 1300 km and 4.5 GeV (dotted blue lines), for a 2500 km baseline with 6 GeV muons (green lines), and for a setup with two detectors at 1300 km and 2500 km with 6 GeV muons (red lines).

The 2500 km baseline has better sensitivity to the NSI's, as expected, but is less sensitive to θ_{13} than the 1300 km baseline - this is indicative that a shorter baseline suppresses NSI's and therefore maximises sensitivity to the oscillation parameters. We have also simulated a baseline of 800 km and find that the combination with the 1300 km baseline slightly improves the sensitivity to the oscillation parameters but has no effect on the NSI sensitivity. However, the combination of the 1300 km and 2500 km baselines gives a drastic improvement upon either of the single baselines. Remembering that this is *not* just down to doubling the statistics, this demonstrates that the different energy dependencies of each of the baselines helps to resolve the degeneracies between the oscillation and NSI parameters.

7.5. The role of the platinum channel

The platinum channel has been discussed extensively in this chapter and the previous one. From a practical point of view, we need to know precisely how valuable the platinum channel is, and whether its benefits outweigh the technical difficulties of incorporating it into the experimental design.

For standard oscillation measurements, we found that the most important factor for the LENF is statistics, and that increasing the number of golden channel events improves the performance of the experiment more than anything else. There are no severe inherent degeneracies which cannot be resolved given sufficiently good statistics and detector performance. However, in the case that the actual flux turns out to be lower than expected, then the platinum channel could be a powerful way of resolving degeneracies. Here we briefly return to our analysis of the oscillation parameters to quantify how the ability to detect electrons can compensate for a lower than expected flux. The benefits of the platinum channel here are two-fold: firstly it increases the number of events observed, thus increasing the statistics, and secondly the complementarity of the platinum channels and the golden channels means that the benefit is greater than that obtained by simply increasing the signal of the golden channel by the number of

platinum events. We compare the 3σ CP discovery potential and hierarchy sensitivity obtained when using the estimated flux of 1.4×10^{21} muon decays per year per polarity and running for 10 years per polarity (a total of 2.8×10^{22} decays), shown by the solid red lines, with the results obtained by using 1×10^{21} muon decays per year per polarity, running for 5 years per polarity (1.0×10^{22} decays, corresponding to the standard HENF setup but where all the muons go to a *single* LENF detector, rather than being split between *two* HENF detectors), both with the platinum channel (dotted blue lines) and without (dotted green lines). This is assuming that the ν_e detection efficiency is 47% and that the background level is 10^{-2} as in the previous chapter.

The benefit of the platinum channel is most prominent for $10^{-3} \lesssim \sin^2 2\theta_{13} \lesssim 10^{-2}$. Recall that the LENF performance is best, relative to the other options including the HENF, for $\sin^2 2\theta_{13} \gtrsim 10^{-3}$ (see Fig. 6.16). For 1.0×10^{22} decays and $\sin^2 2\theta_{13} \sim 10^{-3}$, the platinum channel increases CP sensitivity by $\sim 10\%$ although higher statistics are always better. For hierarchy sensitivity, if $\sin^2 2\theta_{13} \lesssim 4 \times 10^{-3}$ then higher statistics are better, for $4 \times 10^{-3} \lesssim \sin^2 2\theta_{13} \lesssim 10^{-2}$ the platinum channel gives $\sim 20\%$ more coverage, and for $\sin^2 2\theta_{13} > 10^{-2}$ the platinum channel is unnecessary.

Therefore the best option obviously depends on the value of θ_{13} (as measured by the current generation of experiments). In the case that $\sin^2 2\theta_{13} \lesssim 10^{-3}$, the best option is to use the HENF as no other setup has sensitivity in this region. In the case that $10^{-3} \lesssim \sin^2 2\theta_{13} \lesssim 10^{-2}$ the LENF is a viable option, with the addition of the platinum channel giving a $\sim 10 - 20\%$ increase in CP coverage for CP discovery and hierarchy sensitivity. The coverage for θ_{13} discovery is already nearly maximal in this range, but will also be enhanced. For $\sin^2 2\theta_{13} \gtrsim 10^{-2}$ the LENF is still a good option, but the platinum channel adds very little to the oscillation sensitivity.

For NSI measurements, the LENF sensitivity is limited by the degeneracies between the oscillation and NSI parameters. To resolve this, complementary information from an additional channel or baseline is required. Statistics are not important for NSI measurements as they do not help to disentangle the degeneracies, and increasing the total number of muon decays above 1.4×10^{22} does not have any benefit. For the efficiency and background level which we assume for the platinum channel (47%, 10^{-2}), the platinum channel is only helpful if $\theta_{13} \gtrsim 5^\circ$, but if the performance of this channel can be sufficiently improved (technically difficult!) the platinum channel will also be useful for small θ_{13}.

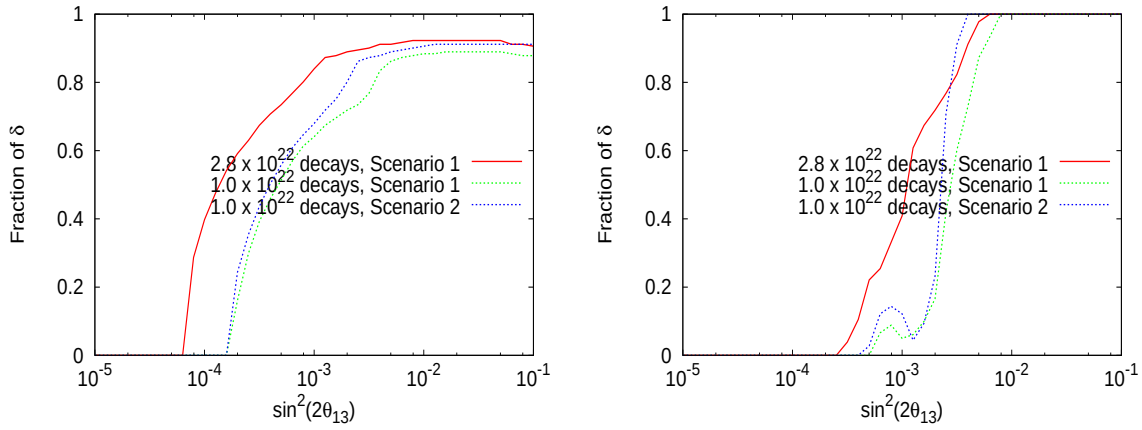

(a) CP discovery potential. (b) Hierarchy sensitivity.

Figure 7.11.: 3σ a) CP discovery potential and b) hierarchy sensitivity for 1.0×10^{22} decays and 2.8×10^{22} decays, comparing Scenarios 1 and 2.

7.6. Summary

We have analysed the sensitivity of the LENF to the NSI parameters $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau}$. We have found that the sensitivity of the setup is limited by the degeneracies between the standard oscillation and NSI parameters, and that therefore increasing the statistics, by itself, does *not* increase the sensitivity of the setup. This is in direct contrast to the results we obtained in the previous chapter for standard oscillation parameters. Instead, it is necessary to include information from a complementary channel, such as the platinum channel, or a second baseline. For an assumed efficiency of 47% for the platinum channel and background of 10^{-2}, its impact is greatest for *large* values of θ_{13} - again, this is in direct contrast to the oscillation parameter measurements. However, the fact that a platinum channel with negligible background has an impact for *all* values of θ_{13} indicates that there is a critical signal to background threshold which must be overcome in order for this channel to become effective.

If including a second baseline, adding a longer baseline gives the best results since matter NSI's are enhanced by high energies and long baselines.

Using the reference LENF setup defined in Section 6.7, including the platinum channel but using a running time of 5 years per polarity rather than 10 years, the LENF has sensitivity to $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau}$ down to $\sim 10^{-2}$. This is a significant improvement upon the current bounds, but it is not as strong as the bounds which the HENF can obtain

(10^{-3} or below). Unfortunately, there is no way to reach this level of sensitivity with the energy and baseline of the LENF.

The general value of the platinum channel was then discussed. With the current estimated performance, from the point of view of oscillation measurements, it is only advantageous to incorporate these channels if statistics are lower than the anticipated value, or if $\sin^2 2\theta_{13} \sim 10^{-3}$. If $\sin^2 2\theta_{13}$ is smaller than this then the HENF is a better option, whereas if $\sin^2 2\theta_{13}$ is very large, then the LENF performance is already optimal and the platinum channel adds nothing. However, for NSI measurements, it is advantageous to incorporate these channels regardless of the flux, and it is most effective for large values of θ_{13}. So if optimising the LENF for oscillation measurements, which we find to be statistics-limited, the priority should be to increase the statistics of the golden channel before incorporating the platinum channel. However, if the optimisation is for NSI's, then the improvement of the platinum channel performance should be prioritised so that it also becomes effective for small values of θ_{13}.

Chapter 8.

Conclusions

8.1. Summary and conclusions

In the era of the Large Hadron Collider, the largest and most ambitious physics experiment ever built, it is easy to think that collider experiments such as these will provide us with the answers to all the unsolved problems in particle physics. Whilst it is true that hadron colliders are the perfect arena in which to study topics such as the physics of electro-weak symmetry breaking, are they necessarily the best tools for studying aspects such as flavour physics, or the physics of the leptonic sector of the Standard Model? The fact that our first evidence of physics beyond the Standard Model came from the discovery of neutrino oscillations indicates that the neutrino sector can provide us with many clues of physics at high energies, from a low-energy perspective.

Neutrinos are the Standard Model particles about which we have the least information, as we found in Chapters 1 and 2 - we still have to establish whether they are Dirac or Majorana particles, we do not yet know whether there is CP violation in the neutrino sector, the mixing matrix which controls the three-family oscillations is not yet complete, and we have yet to measure the masses of the neutrinos and the ordering of the masses. The first two of these questions have resounding implications not just for particle physics, but also from a cosmological perspective. For the first, it is interesting in itself to discover whether neutrinos are Dirac particles like the other Standard Model fermions, or whether they are Majorana particles and thus fundamentally different. But one of the consequences of Majorana neutrinos is that it leads to the possibility of *lepton number violation*, a quantity which is otherwise conserved in the Standard Model. If lepton number is violated then it means that *leptogenesis* is possible (Section 1.3). This

is the process whereby a net lepton/ anti-lepton asymmetry is generated in the early universe. This can be converted into a baryon asymmetry and thus provides a mechanism to explain the matter/ anti-matter asymmetry of the universe - one of the biggest unsolved mysteries in fundamental physics today. The question of whether neutrinos are Dirac or Majorana can only be answered by dedicated neutrino experiments which are sensitive to the lepton-number violating processes characteristic of Majorana neutrinos, such as neutrinoless double-beta decay.

The answer to the second question regarding CP violation has similar implications to the first; if there *is* CP violation in low-energy neutrino oscillations then this may be an indicator towards the existence of the high-energy CP violation required for leptogenesis. The possibility for CP violation in the neutrino sector originates from the three-family mixing of neutrinos - a complex 3×3 matrix is parameterised by three mixing angles and one complex phase; the phase is physical only if all three mixing angles are non-zero. So to measure CP violation we need to measure this phase, δ, but before we do so we need to establish if all the mixing angles are indeed non-zero.

The mixing angles, and the CP phase, δ, can be measured by observing neutrino oscillations, described in Chapter 2. Particle oscillations are a quantum-mechanical phenomena which occur when the particle mass eigenstates (the propagating states) and the flavour states (the interaction states) do not coincide. For neutrinos, we showed that the oscillations are dependent not only upon the three mixing angles (θ_{12}, θ_{23}, θ_{13}) and complex phase (δ) of the neutrino mixing matrix (the PMNS matrix) but also on the mass-squared differences between the neutrino mass eigenstates (Δm^2_{21}, Δm^2_{32} and Δm^2_{31} where $\Delta m^2_{jk} = m^2_j - m^2_k$), and the ratio $\frac{L}{E}$ where L is the distance over which the neutrino has propagated (the 'baseline') and E is its energy. We described how, by varying the ratio $\frac{L}{E}$ and by choosing to observe different oscillation channels, an experiment could be tuned to be sensitive to different oscillation parameters. In this way, past experiments have successfully measured the so-called solar parameters, θ_{12} and Δm^2_{21}, and the so-called atmospheric parameters, θ_{23} (although experiments have so far only actually measured $\sin^2 2\theta_{23}$ which means that we do not know whether θ_{23} is greater or less than 45°) and $\Delta m^2_{32} \simeq \Delta m^2_{31}$. The emerging picture of the mixing angles is of one which follows the pattern of 'tri-bimaximal mixing'. This pattern predicts the value of θ_{13} to be exactly zero; our current experimental bound on this angle is at present very weak ($\lesssim 13°$ at 3σ). The picture of the neutrino masses is one where $\Delta m^2_{21} > |\Delta m^2_{32}|$ and $\Delta m^2_{32} \simeq \Delta m^2_{31}$ although we do not know whether Δm^2_{32} and Δm^2_{31} are positive or negative; in other words we do not know the *mass hierarchy* - which mass eigenstate is

the lightest or heaviest. At present there has also not been any significant measurement of the CP phase, δ.

It is the primary goal of future neutrino oscillation experiments to measure these unknown oscillation parameters. By doing so they will enable us to put rigourous constraints on theoretical models of new physics. By measuring δ they will hint at whether or not leptogenesis is a viable possibility, and they will give us strong clues about the physics of flavour in the Standard Model. This is one of the least well-understood aspects of the model. The emerging pattern of neutrino mixing is very different to that of mixing in the quark sector; an obvious question is why this should be so? To answer this question definitively requires *precision* measurements of the neutrino mixing parameters. Past and current experiments (Chapter 3) have done well so far to measure the atmospheric and solar parameters to $\sim 10\%$ accuracy, but to obtain the desired precision down to the $\sim 1\%$ level, to match the accuracy obtained in the quark sector, necessitates a new generation of long-baseline neutrino oscillation experiments (Section 3.5). These experiments also turn out to be a powerful tool for searching for *non-standard interactions* (Section 2.5) - lepton flavour-changing interactions arising from new physics processes.

It is important not just to build a future experiment which is capable of making the desired measurements, but to choose and build the best one. We discussed in Chapter 4 how the design of a next-generation experiment depends crucially on the value of θ_{13} - we saw that the value of θ_{13} dictates which of the terms in the $\nu_e \rightarrow \nu_\mu$ channel probability are dominant. This is the so-called 'golden channel' which future experiments are designed to measure, chosen because measurements of all the unknown oscillation parameters can be made from this single channel. The reason why no experiments to date have been able to observe the golden channel is because it is a sub-dominant channel, suppressed by θ_{13}, and current experiments do not yet have the power to make such precise measurements. The smaller the value of θ_{13}, the smaller the amplitude of the oscillations and the more difficult the task will be.

In the scenario that $\sin^2 2\theta_{13} \gtrsim 10^{-2}$ then super-beam experiments should be able to make some preliminary measurements of the oscillation parameters in the relatively near future since the technology required is very similar to that used by current conventional neutrino beams. Even at this stage, we still want to build the optimal experiment which is why it is important to consider all potential setups which have a good physics reach. We studied European super-beam setups, in the context of the European Large Apparatus for Grand Unification and Neutrino Astrophysics (LAGUNA) design study,

which as yet have not been seriously considered. Three detectors are being considered by the design study: a 100 kton liquid argon time projection chamber, a 440 kton water Čerenkov detector, and a 50 kton liquid scintillator detector. Of these, liquid scintillator detectors are the least well developed. We therefore performed some preliminary studies (Section 5.3) to ascertain the potential of such a detector and to determine the factors which most affect its performance. We found that for small θ_{13} it is the background level which limits the performance, whereas for large θ_{13} it is the systematic errors. We then went on to study the potential of a super-beam using the 2285 km baseline from CERN to Pyhäsalmi, finding that if a 100 kton liquid argon or 50 kton liquid scintillator detector is used, this setup can have sensitivity to θ_{13}, δ and the mass hierarchy for $\sin^2 2\theta_{13} \gtrsim 10^{-2}$ which is competitive with experiments such as the US LBNE experiment. This limit is set mostly by the intrinsic beam background of a super-beam - this is the $\sim 1\%$ $\nu_e(\bar{\nu}_e)$ content of the ν_μ ($\bar{\nu}_\mu$) beam which acts as a background to the $\nu_\mu \to \nu_e$ ($\bar{\nu}_\mu \to \bar{\nu}_e$) channels that super-beam experiments detect. Therefore, whilst super-beams such as the LAGUNA setup can measure the oscillation parameters if θ_{13} is large, thus helping us to choose which experiment we should build next, it cannot make truly *precise* measurements and neither will it have any sensitivity if θ_{13} is small.

If θ_{13} is very small such that $\sin^2 2\theta_{13} \lesssim 10^{-4}$, then the only experiment capable of making the required measurements is one which has a very long baseline, of the order of several thousands of kilometres, with a corresponding neutrino energy of several tens of GeV. This is because, as we showed in Section 4.3, 'matter effects' enhance the difference in the oscillation probabilities for normal and inverted mass hierarchies, thus making it easier to distinguish between the two hierarchies especially if θ_{13} is very small. Measuring the hierarchy is important in itself, but also because if we are unable to determine the hierarchy, then this is detrimental to our measurement of the other parameters because of the *degeneracies* between the parameters (Section 4.5). Matter effects are enhanced by a long baseline and high energy and so it is natural to expect that a longer baseline is always preferable. Comparison studies between various long-baseline experiments have shown that this is indeed the case if θ_{13} is very small; however we have to address the question of whether this is true for other values of θ_{13}.

One such long-baseline experiment is the *neutrino factory* (Section 3.5.3), an experiment capable of producing an extremely intense and pure neutrino beam consisting of an equal mix of ν_μ and $\bar{\nu}_e$ (or $\bar{\nu}_\mu$ and ν_e) from the decay of muons. It is because of this high flux and purity that the neutrino factory has been found to be an incredibly powerful neutrino oscillation experiment with a very impressive physics reach. This is aided by

the fact that the beam properties can be predicted to a high degree of accuracy, because of the leptonic nature of the neutrino production which is very 'clean'. Additionally, a large number of oscillation channels are accessible because of the initial mixed flavour content of the beam. The standard neutrino factory setup, as defined by the International Design Study for the Neutrino Factory, uses a neutrino beam of energy ∼25 GeV and two detectors - one at ∼4000 km and a second at the *magic baseline* of ∼7500 km (Section 4.5). This combination of two baselines enables parameter degeneracies to be resolved, making it an optimal setup for measuring the neutrino oscillation parameters and also for searching for non-standard interactions.

Prior to the work of this thesis, no detailed studies of a *low-energy neutrino factory* had been performed. In Chapter 6 we performed detailed simulations of this setup. The low-energy neutrino factory is a neutrino factory which uses a neutrino beam of energy ∼5 GeV and a corresponding single baseline of ∼1000 km together with a magnetised totally active scintillating detector. It was originally proposed because of the realisation that at this energy and baseline, the oscillation spectrum is very rich, potentially enabling for all the desired oscillation parameters to be measured if $\sin^2 2\theta_{13} \gtrsim 10^{-3}$ without the need for very high energies and an extremely long baseline. We performed optimisation studies of this setup to gauge which factors should most influence the experimental design.

We considered, for the first time, the addition of the *platinum channels* (Section 6.2.6), $\nu_\mu \to \nu_e$ and $\bar{\nu}_\mu \to \bar{\nu}_e$, to the setup. In theory, these channels are complementary to the golden channels and thus help to resolve the degeneracies which otherwise limit the precision of the measurements. However we found that in practice, the technical difficulties associated with detecting the platinum channel (in particular, distinguishing the electron signal from the neutral-current and pion backgrounds, and correctly identifying the charge of the electron) mean that the theoretical gain is lost; for an experimentally realistic efficiency of 47% and background level of 10^{-2} it is not possible to obtain any useful information (for comparison, the efficiency for muon detection is 94% with a background of 10^{-3}). In certain circumstances, such as if only around a third of our predicted exposure can be achieved, the platinum channel *is* useful. However, in general, the small benefit does not appear to be worthy of the technical effort required to optimise the detector for electron detection.

An additional finding was that, contrary to popular belief amongst the long-baseline community, being able to measure the second oscillation maximum in addition to the first maximum does not significantly improve the performance of the experiment (Sec-

tion 6.2.3). Theoretically, the observation of the second maximum should aid our measurement of CP violation because this is an effect which is most prominent at low energies. Also the information from the second maximum should be complementary to that from the first maximum, which again should help to resolve degeneracies. But we find once again that in practice, this theoretical gain is lost. The reason is that a beam optimised to peak around the region of the first oscillation maximum has relatively few events around the second peak; the difference in statistics in these two regions means that essentially no useful information can be extracted from the second peak. In our setup, there is such a high flux at the first maximum, and the energy resolution of the detector is so good that it is possible to probe the shape of the spectrum around the first oscillation peak sufficiently accurately to negate the necessity for complementary information from the second peak.

The main message which emerged from our optimisation studies was that maximum effort should be put into obtaining as many golden channel events as possible, in the energy region of the first oscillation peak. We found that the benefit of increasing the statistics of the experiment always outweighs any other optimisation. If we can obtain an exposure of 2.8×10^{23} kton $\times$ decays per polarity (1.4×10^{21} muon decays per year per polarity, running for 5 years per polarity with a 20 kton detector) then a low-energy neutrino factory with a baseline of 1300 km, muon energy of 4.5 GeV and a 20 kton magnetised totally active scintillating detector, or liquid argon detector, has sensitivity to θ_{13} and the mass hierarchy, at 3σ confidence, for $\sin^2 2\theta_{13} \gtrsim 10^{-3}$. We can also determine whether θ_{23} is greater or less than 45° if the deviation is $\gtrsim 8^\circ$, and exclude maximal θ_{23} if the deviation is $\gtrsim 2^\circ$. In addition, we find that for $\sin^2 2\theta_{13} \gtrsim 10^{-3}$, this setup has *better sensitivity to CP violation than the high-energy neutrino factory* (and, for $\sin^2 2\theta_{13} > 4 \times 10^{-3}$ also has 100% CP coverage for the mass hierarchy and θ_{13}). This occurs because of the difficulty in distinguishing between genuine CP violation, as described by the phase δ, and apparent 'CP violation' due to the CP asymmetry of the baseline: we explained in Section 4.3.2 how matter effects fake CP violation - the fact that the earth contains only matter and not anti-matter means that neutrinos and anti-neutrinos interact differently as they propagate through the earth. However, this is exactly the effect of CP violation - a distinction between the oscillation properties of neutrinos and anti-neutrinos. Therefore it is imperative that we can distinguish between genuine CP violation and fake CP violation due to matter effects if we want to make a precision measurement of δ. For large values of θ_{13}, the sensitivity of the high-energy neutrino factory to CP violation becomes limited by the systematic uncertainty of the matter density of the baseline; this limits the accuracy to which the matter effects can

be predicted. In contrast, the low-energy neutrino factory with its shorter baseline does not suffer from this problem, enabling it to maintain its CP sensitivity in this region. Yet the baseline is sufficiently long such that when combined with the high statistics of the setup, the mass hierarchy can still be determined - this is a unique feature of the low-energy neutrino factory.

We also studied the *precision* of the setup, in terms of the size of the 1σ errors on each of the parameters. We analysed the errors as a function of θ_{13} (Section 6.4), finding that for $\sin^2 2\theta \gtrsim 10^{-3}$, our setup can measure θ_{13} with a $\sim 5\%$ 1σ error, θ_{23} to within $\pm \sim 2.5^\circ$ and δ to within $\sim \pm 10^\circ$. We found that for exposures up to $\sim 6 \times 10^{23}$ kton $\times$ decays per polarity, the precision is significantly improved. The effect of non-zero systematic errors and backgrounds was found to be to effectively halve the exposure. Finally, in Section 6.6 we showed how setups with *different* values of L and E can have similar performances, finding that our current setup is indeed one that is nearly optimal.

We assessed the sensitivity of the low-energy neutrino factory to non-standard matter interactions in Chapter 7. These are flavour-changing non-standard matter effects, not predicted by the Standard Model, and so they are a direct signal of new physics. They can be modelled in our experiments as effective four-point interactions with strength $\varepsilon_{\alpha\beta}$, relative to the weak coupling constant. This parameter describes the rate of the process $\nu_\alpha \to \nu_\beta$. Although measurably large non-standard interactions ($\gtrsim 10^{-4}$) are theoretically difficult to generate at scales above the electro-weak scale, it is still important to search for them both as a means of confirming our Standard Model predictions and because of the possibility that they may arise from mechanisms which we have not yet considered - it is good to be open-minded!

The low-energy neutrino factory has sensitivity to the non-standard matter parameters $\varepsilon_{e\mu}$ and $\varepsilon_{e\tau}$. We found that our setup could put bounds on these parameters down to the level of $\sim 10^{-2}$. This is roughly an order of magnitude better than current bounds, but an order of magnitude worse than the bounds which a high-energy neutrino factory could obtain - non-standard matter interactions are enhanced by a long baseline and high energy and so the high-energy setup is optimal for these non-standard searches. We found that the sensitivity of our setup is limited by the degeneracies between the standard oscillation parameters and non-standard parameters. This is not something which can be resolved by simply increasing statistics, but requires the addition of complementary information from either a second baseline or complementary channels. So in this situation the platinum channels *are* necessary to optimise the performance of the experiment. With the predicted efficiency and background, the platinum channels

have most effect for large θ_{13}. We found that if this performance could be hypothetically improved then the platinum channels become useful also for small values of θ_{13} - this indicates that there is a threshold signal to background ratio which must be exceeded for these channels to be effective. The platinum channels help not only to measure the non-standard parameters, but also to *maintain* our sensitivity to the oscillation parameters which is otherwise severely inhibited when we marginalise over the non-standard parameters, even if they are all zero, as well as the oscillation parameters.

8.2. Outlook

Rapid progress is being made with respect to the design of the neutrino factory. The technological, logistical and economic problems are also becoming apparent in the process! A recent development is that the distinction between the high-energy and low-energy neutrino factory setups is far from clear - whereas it was previously thought (at the time that the work of this thesis was being performed) that these were two distinct setups using two distinct detector technologies, we now realise that this is not necessarily the best way to approach the experimental design. This development is due in part to progress in the work on the magnetised iron neutrino detector which is the detector being designed for the high-energy neutrino factory. New selection algorithms mean that the efficiencies of the lower energy bins have recently been increased so that it is possible to obtain a reasonable number of events around the ~ 5 GeV peak of the low-energy setup [118]. Coupled with the finding that contrary to popular belief, the second oscillation maximum is not necessary if we have sufficient events at the first oscillation maximum, this means that we could also consider using the iron detector with the low-energy neutrino factory, especially if we consider slightly longer baselines than the 1300 km baseline of our reference setup (so that the oscillation region is at a slightly higher energy).

There is a strong case to build a neutrino factory rather than a super-beam, even if our current generation of experiments discover that θ_{13} is large, from the point of view of *precision* and also in the search for non-standard interactions. The extremely high fluxes which a neutrino factory can achieve, in conjunction with the absolute purity of the beam, aid not only the discovery potential of the neutrino factory, but also its precision. In the introductory chapters of this thesis we highlighted the physics motivation for making precise measurements of the oscillation parameters and for searching for non-standard interactions - neutrino oscillations provide a window to high-energy

physics from a low-energy perspective. The more precisely we can measure the oscillation parameters, the more strongly we can constrain models of new physics, and the more precise the clues are that point us towards the correct theory. The discovery of neutrino oscillations was the first evidence that there is physics beyond the Standard Model; it is not unlikely that neutrino oscillations will bring us yet more evidence in the future.

Appendix A.

γ-matrices

The Lorentz group can be represented by any set of matrices which obey the same commutation rules as the group algebra. We are interested in representations which correspond to spin-$\frac{1}{2}$ particles (fermions). A set of matrices which form such a representation are

$$S^{\mu\nu} = \frac{i}{4}[\gamma^\mu, \gamma^\nu], \tag{A.1}$$

where γ^μ and γ^ν are the *γ-matrices*. A four-component field that Lorentz transforms according to $S^{\mu\nu}$ is called a *Dirac spinor*. The number of γ matrices corresponds to the number of space-time dimensions; therefore we have a set of four matrices, $\{\gamma^0, \gamma^1, \gamma^2, \gamma^3\}$. They satisfy the anti-commutation relations

$$\{\gamma^\mu, \gamma^\nu\} = 2\eta^{\mu\nu} I, \tag{A.2}$$

where $\eta^{\mu\nu}$ is the metric tensor and I is the 4×4 identity matrix. The explicit form of the matrices is representation-dependent but their algebra is not. There is a fifth matrix, γ^5, which does *not* form part of the representation but is useful to define as it has some important properties. It is related to the other four γ-matrices via

$$\begin{aligned} \gamma^5 &= \frac{i}{4!}\epsilon_{\nu\mu\sigma\rho}\gamma^\nu\gamma^\mu\gamma^\sigma\gamma^\rho \\ &= i\gamma^0\gamma^1\gamma^2\gamma^3, \end{aligned} \tag{A.3}$$

and has the properties

$$\{\gamma^5, \gamma^\mu\} = 0, \tag{A.4a}$$
$$(\gamma^5)^2 = I, \tag{A.4b}$$
$$(\gamma^5)^\dagger = I. \tag{A.4c}$$

If we choose to use the Weyl representation (also called the chiral representation), as in Section 1.2, then the γ-matrices take the block-diagonal form

$$\gamma^0 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \gamma^j = \begin{pmatrix} 0 & \sigma^j \\ -\sigma^j & 0 \end{pmatrix}, \qquad \gamma^5 = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \tag{A.5}$$

where σ^j are the 2×2 Pauli matrices. Dirac spinors can be written in this representation in terms of two two-component *Weyl spinors*, ψ_- and ψ_+:

$$\begin{aligned} \psi &= \begin{pmatrix} \psi_- \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ \psi_+ \end{pmatrix} \\ &= \psi_L + \psi_R, \end{aligned} \tag{A.6}$$

where L and R denote left-handed and right-handed. Now it is easy to see that $\gamma^5\psi_{L,R} = \mp\psi_{L,R}$.

Appendix B.

Charge-conjugation and the Majorana condition

The Standard Model Lagrangian is invariant under continuous Lorentz transformations, but there are also three *discrete* transformations which we can perform. These are charge-conjugation, parity reversal and time reversal. Time reversal corresponds to the transformation $(t, \mathbf{x}) \to (-t, \mathbf{x})$; for instance $(\nu_e \to \nu_\mu) \to (\nu_\mu \to \nu_e)$. Parity reversal is a spatial reflection, $(t, \mathbf{x}) \to (t, -\mathbf{x})$; for example this reverses the helicity of a particle. Charge-conjugation transforms a particle of a given helicity to an anti-particle with the same helicity (by convention). All experimental observations so far indicate that the combination CPT is a complete symmetry. In other words, the Standard Model is invariant under this combination of transformations.

For our discussion of Majorana particles, we are interested in charge-conjugation matrix. A spinor corresponding to a particle with momentum $\mathbf{p}$ and spin s, $\psi^s(\mathbf{p})$, is related to the spinor of the anti-particle with the same momentum and spin, $(\psi^s(\mathbf{p}))^c$, via

$$\psi^c = C\bar{\psi}^T \tag{B.1}$$

where C is the charge-conjugation matrix which has the properties

$$C^\dagger C = I, \tag{B.2a}$$
$$C^T = -C, \tag{B.2b}$$
$$C\gamma^\mu C^{-1} = -\gamma^\mu. \tag{B.2c}$$

In the Weyl representation, C takes the explicit form

$$C = -i\gamma^2\gamma^0, \tag{B.3}$$

where the γ-matrices are defined in Appendix A. Then Eq. (B.1) becomes

$$\begin{pmatrix} \psi_- \\ \psi_+ \end{pmatrix}^c = \begin{pmatrix} 0 & -i\sigma^2 \\ i\sigma^2 & 0 \end{pmatrix} \begin{pmatrix} \psi_- \\ \psi_+ \end{pmatrix}^* . \tag{B.4}$$

Now we can see that the Majorana condition (Eq. (1.13)), $\psi^c = \psi$, corresponds to the condition

$$\psi_- = -i\sigma^2\psi_+^* \qquad \text{or equivalently} \qquad \psi_+ = i\sigma^2\psi_-^*. \tag{B.5}$$

Appendix C.

Parameterisation of the PMNS matrix

To count the number of independent parameters in the PMNS matrix, first consider a generic $n \times n$ complex matrix. This has $2n^2$ degrees of freedom (n^2 real and n^2 imaginary). Imposing the constraint of unitarity, which is required in order for transition probabilities to be unitary, imposes $\frac{1}{2}n(n+1)$ constraints on the real components and $\frac{1}{2}n(n+1)$ on the imaginary components - a total of $n(n+1)$ constraints. This leaves $n^2 - n$ parameters. If we first consider the CKM matrix, V^{CKM}, which describes mixing in the quark sector, this appears in the Standard Model Lagrangian in terms such as

$$\mathcal{L}^{CKM} \sim W_\mu^+ V_{jk}^{CKM} \bar{\ell}_{uj} \gamma^\mu (1+\gamma_5) \ell_{dk} + W_\mu^- V_{jk}^{CKM\dagger} \bar{\ell}_{dj} \gamma^\mu (1+\gamma_5) \ell_{uk}, \tag{C.1}$$

where ℓ_{uj} are the up-type quarks and ℓ_{dj} are the down-type quarks. Now we can remove $n-1$ phases from the CKM matrix by redefining the quark fields: $|\ell_{uj}\rangle \rightarrow e^{i\theta_{uj}} |\ell_{uj}\rangle$ and similarly for the down-type quarks. Amplitudes such as $\langle \ell_{uj} | \ell_{uj} \rangle$ are invariant. As all the other quark terms in the Standard Model are flavour-diagonal, they also remain invariant under this redefinition.

Only $n-1$ and not n phases can be removed because one phase is an overall phase (rather than just a relative phase like the others). So the CKM matrix has $n^2 - n - (n-1) = n^2 - 2n + 1$ free parameters. To ascertain how many of these are real angles and how many are complex phases, we can compare it to a real orthogonal matrix. This has n^2 real parameters and orthogonality imposes $\frac{1}{2}n(n+1)$ constraints, leaving $\frac{1}{2}n(n-1)$ independent parameters which are all real. Therefore the orthogonal complex quark mixing matrix is parameterised by $\frac{1}{2}n(n-1)$ real mixing angles and so the remaining

$\frac{1}{2}(n-1)(n-2)$ parameters must be complex phases. Hence in the Standard Model where $n = 3$, the CKM matrix contains three mixing angles and one complex phase.

Returning now to neutrinos and the PMNS matrix, U^{PMNS}, this appears in the Standard Model Lagrangian in the terms

$$\mathcal{L}^{PMNS} \sim W^+_\mu U^{PMNS}_{jk} \bar{\nu}_j \gamma^\mu (1+\gamma_5) \ell_k + W^-_\mu U^{PMNS\dagger}_{jk} \bar{\ell}_j \gamma^\mu (1+\gamma_5) \nu_k, \qquad \text{(C.2)}$$

where ℓ_j is any up or down-type quark. Three of the phases can still be removed by redefining the charged lepton fields, but the phases of the neutrino fields *cannot* automatically be redefined because of the possibility that they are Majorana fields; if $\nu_j = \nu_j^c$ and we redefine $|\nu_j\rangle \to e^{i\theta_j}|\nu_j\rangle$, then $\langle\nu_j|\nu_j\rangle \to \langle\nu_j|\nu_j\rangle e^{2i\theta_j}$; this means that these phases *are* physical. Thus the additional $n-1$ phases must be retained, and these are the *Majorana phases*. The PMNS matrix can then be written as a product of three 2×2 rotation matrices and two phase matrices as given in Eq. (2.1).

Appendix D.

Neutrino oscillations in vacuum

Derivation of oscillation probabilities using the equal momenta assumption

The probability for the oscillation $\nu_\alpha \to \nu_\beta$, in vacuum, is given by

$$P_{\nu_\alpha \to \nu_\beta}(L,E) = \left| \sum_{j=1}^{3} U^*_{\alpha j} U_{\beta j} e^{-iL\left(p + \frac{m_j^2}{2p}\right)} \right|^2 , \tag{D.1}$$

as given in Eq. (2.7). Writing this out explicitly, we have

$$\begin{aligned} P_{\nu_\alpha \to \nu_\beta}(L,E) &= \left(U^*_{\alpha 1} U_{\beta 1} e^{-iL\left(p+\frac{m_1^2}{2p}\right)} + U^*_{\alpha 2} U_{\beta 2} e^{-iL\left(p+\frac{m_2^2}{2p}\right)} + U^*_{\alpha 3} U_{\beta 3} e^{-iL\left(p+\frac{m_3^2}{2p}\right)} \right) \times \\ &\quad \left(U_{\alpha 1} U^*_{\beta 1} e^{iL\left(p+\frac{m_1^2}{2p}\right)} + U_{\alpha 2} U^*_{\beta 2} e^{iL\left(p+\frac{m_2^2}{2p}\right)} + U_{\alpha 3} U^*_{\beta 3} e^{iL\left(p+\frac{m_j^3}{2p}\right)} \right) . \end{aligned} \tag{D.2}$$

The diagonal terms are

$$\begin{aligned} & |U_{\alpha 1}|^2 |U_{\beta 1}|^2 + |U_{\alpha 2}|^2 |U_{\beta 2}|^2 + |U_{\alpha 3}|^2 |U_{\beta 3}|^2 \\ = & \left| \sum_{j=1}^{3} U^*_{\alpha j} U_{\beta j} \right|^2 - (U^*_{\alpha 1} U_{\alpha 2} U_{\beta 1} U^*_{\beta 2} + U^*_{\alpha 1} U_{\alpha 3} U_{\beta 1} U^*_{\beta 3} + U^*_{\alpha 2} U_{\alpha 3} U_{\beta 2} U^*_{\beta 3} + \text{ C.C.}) \\ = & \delta_{\alpha\beta} - 2\mathrm{Re}[U^*_{\alpha 1} U_{\alpha 2} U_{\beta 1} U^*_{\beta 2}] - 2\mathrm{Re}[U^*_{\alpha 1} U_{\alpha 3} U_{\beta 1} U^*_{\beta 3}] - 2\mathrm{Re}[U^*_{\alpha 2} U_{\alpha 3} U_{\beta 2} U^*_{\beta 3}], \end{aligned} \tag{D.3}$$

where C.C. denotes the complex conjugate and we have used the unitarity condition $\sum_{j=1}^{3} U^*_{\alpha j} U_{\beta j} = \delta_{\alpha\beta} = |\sum_{j=1}^{3} U^*_{\alpha j} U_{\beta j}|^2$.

The cross-terms are

$$\begin{aligned}
& (U^*_{\alpha1}U_{\alpha2}U_{\beta1}U^*_{\beta2}e^{-i\frac{\Delta m^2_{12}L}{2E}} + U^*_{\alpha1}U_{\alpha3}U_{\beta1}U^*_{\beta3}e^{-i\frac{\Delta m^2_{13}L}{2E}} + U^*_{\alpha2}U_{\alpha3}U_{\beta2}U^*_{\beta3}e^{-i\frac{\Delta m^2_{23}L}{2E}} + \text{ C.C.}) \\
= \; & 2\text{Re}[U^*_{\alpha1}U_{\alpha2}U_{\beta1}U^*_{\beta2}]\cos\left(\frac{\Delta m^2_{21}L}{2E}\right) - 2\text{Im}[U^*_{\alpha1}U_{\alpha2}U_{\beta1}U^*_{\beta2}]\sin\left(\frac{\Delta m^2_{21}L}{2E}\right) \\
+ \; & 2\text{Re}[U^*_{\alpha1}U_{\alpha3}U_{\beta1}U^*_{\beta3}]\cos\left(\frac{\Delta m^2_{31}L}{2E}\right) - 2\text{Im}[U^*_{\alpha1}U_{\alpha3}U_{\beta1}U^*_{\beta3}]\sin\left(\frac{\Delta m^2_{31}L}{2E}\right) \\
+ \; & 2\text{Re}[U^*_{\alpha2}U_{\alpha3}U_{\beta2}U^*_{\beta3}]\cos\left(\frac{\Delta m^2_{32}L}{2E}\right) - 2\text{Im}[U^*_{\alpha2}U_{\alpha3}U_{\beta2}U^*_{\beta3}]\sin\left(\frac{\Delta m^2_{32}L}{2E}\right). \qquad \text{(D.4)}
\end{aligned}$$

Adding all the terms together and using the half-angle relation $1 - \cos\theta = 2\sin^2\frac{\theta}{2}$, we obtain the final expression as given in Eq. (2.8):

$$\begin{aligned}
P_{\nu_\alpha\to\nu_\beta}(L,E) = \delta_{\alpha\beta} \; & - \; 4\sum_{j>k}\text{Re}[U^*_{\alpha j}U_{\beta j}U_{\alpha k}U^*_{\beta k}]\sin^2\left(\frac{\Delta m^2_{jk}L}{4E}\right) \\
& + \; 2\sum_{j>k}\text{Im}[U^*_{\alpha j}U_{\beta j}U_{\alpha k}U^*_{\beta k}]\sin\left(\frac{\Delta m^2_{jk}L}{2E}\right). \qquad \text{(D.5)}
\end{aligned}$$

In this derivation we have assumed that all the neutrino mass states carry *equal momenta*, as described in Section 2.2, and that after time t, the mass eigenstate j picks up a phase factor e^{-iE_jt}. This is equivalent to the assumption that the neutrino states are *plane waves* which is the origin of the problem mentioned in Section 2.2 - the probability of finding a particle described by a plane wave does not depend on the spatial coordinate, implying that oscillation probabilities are a function only of time and not space.

We could instead have used the assumption that all the neutrino mass eigenstates have *equal energies* and performed a slightly different derivation, assuming that after a displacement $\mathbf{x}$ the mass eigenstate j picks up a factor $e^{-i\mathbf{p}\cdot\mathbf{x}}$; this would result in oscillation probabilities that do evolve in space. However, neither this assumption of equal energies, nor the assumption of equal momenta, is Lorentz invariant (the eigenstates have different masses and thus different velocities, so they cannot have equal energies or momenta in *all* frames). In addition there is no reason to assume that the neutrino mass states produced in a weak interaction have either the same momenta or energies (see, for example, Ref. [210]). The exception is for hypothetical 'Mössbauer neutrinos' [211] - neutrinos emitted from recoilless electron capture processes - when the mass eigenstates *are* expected to have equal energies - however this is a highly specialised case!

The wave packet approach

Here we will give a brief summary of the derivation given in Ref. [49], using wave packets, which does not require the use of the equal momenta or equal energies assumptions.

A wave-packet is basically described by a plane wave multiplied by some shape function (often Gaussian) which localises the wave. If we assume that the shape function is sharply peaked about $\mathbf{p} = \mathbf{p}_j$ with width $\sigma_p \ll p_j$, then the wave packet, $\boldsymbol{\psi}_j(t, \mathbf{x})$, describing a propagating neutrino mass eigenstate ν_j, can be written as

$$\boldsymbol{\psi}_j(t, \mathbf{x}) \simeq e^{i(\mathbf{p}\cdot\mathbf{x} - E_j(p)t)} f_j^s(\mathbf{x} - \mathbf{v}_j^g t), \tag{D.6}$$

where $f_j^s(\mathbf{x} - \mathbf{v}_j^g t)$ is the shape factor (the superscript s denotes that this is the neutrino produced at the *source*) and $\mathbf{v}_{gj}$ is the group velocity of the wave packet for ν_j. In general, different mass eigenstates will have different shape factors. A neutrino of flavour α, emitted at the source, is then

$$|\nu_\alpha(t, \mathbf{x})\rangle = \sum_{j=1}^{3} U_{\alpha j}^* \boldsymbol{\psi}_j^s(t, \mathbf{x}) |\nu_j\rangle, \tag{D.7}$$

whereas a neutrino of flavour β which is *detected* is

$$|\nu_\beta(\mathbf{x} - \mathbf{L})\rangle = \sum_{j=1}^{3} U_{\beta j}^* \boldsymbol{\psi}_j^D(\mathbf{x} - \mathbf{L}) |\nu_j\rangle. \tag{D.8}$$

Note that there is no time dependence in this detected state since the detection process is essentially time independent. Also the average momenta of the source and detected states of the *same* neutrino mass eigenstate may be different because the detector may not be sensitive to the exact energy spectrum of the emitted state. The transition amplitude for $\nu_\alpha \to \nu_\beta$ is

$$\begin{aligned} A_{\alpha\beta}(t, L) &= \int d^3x \langle \nu_\beta(\mathbf{x} - \mathbf{L}) | \nu_\alpha(t, \mathbf{x}) \rangle \\ &= \sum_{j=1}^{3} U_{\alpha j}^* U_{\beta j} \int d^3x (\boldsymbol{\psi}_j^D(\mathbf{x} - \mathbf{L}))^* \boldsymbol{\psi}_j^s(t, \mathbf{x}), \end{aligned} \tag{D.9}$$

and the probability, $P_{\alpha\beta}$, is the square of this amplitude. One must also integrate over time, taking into account the temporal response function of the detector. Now neutrino oscillations are caused by the evolving differences in *phases* between the mass eigenstates,

so this is the part of $P_{\alpha\beta}$ in which we are interested. The phase difference between the states j and k is

$$\begin{aligned} \Delta\phi_{jk} &= (E_j - E_k)t - (p_j - p_k)L \\ &= \Delta E_{jk}t - \Delta p_{jk}L, \end{aligned} \tag{D.10}$$

where $E_j = \sqrt{p_j^2 + m_j^2}$. This phase difference is Lorentz invariant. Now let us consider the situation where $|\Delta E_{jk}| \ll E$ (where E is the average energy). This corresponds to relativistic or quasi-degenerate neutrinos. In this case we can correctly use the approximation

$$\begin{aligned} \Delta E_{jk} &= \frac{\partial E}{\partial p}\Delta p_{jk} + \frac{\partial E}{\partial m^2}\Delta m^2_{jk} \\ &= v_g \Delta p_{jk} + \frac{1}{2E}\Delta m^2_{jk}, \end{aligned} \tag{D.11}$$

so that

$$\Delta\phi_{jk} = \frac{\Delta m^2_{jk} t}{2E} - (L - v_g t)\Delta p_{jk}. \tag{D.12}$$

The first term is the standard oscillation phase which is recovered if we make the equal momenta assumption so that $\Delta p = 0$ and the second term vanishes. But note that the second term also vanishes at the centre of the wave packet when $L = v^g t$ but is non-zero at all other points in the wave packet. Thus we can identify the first term as being the phase accumulated by a *point-like* neutrino over the distance L, whereas the second term takes into account the finite size of the wave packet. In most oscillation experiments, the neutrinos which are produced are highly relativistic so that the condition $|\Delta E_{jk}| \ll E$ is satisfied. In addition, the neutrinos are produced *coherently* - that is that the spatial extent of their wave packets is much less than the oscillation wavelength; if this condition is not met then oscillations will not occur. This means that the second term in Eq. (D.12) vanishes. Using $t \simeq L$ for our ultra-relativistic neutrinos, we then recover the familiar oscillation phase $\frac{\Delta m^2_{jk} L}{2E}$, the same answer which we obtained by using the equal momenta assumption. However, although this assumption leads to the correct answer by means of a simple derivation, it is *not*, in general, the correct reasoning. The correct reasons are those outlined above and explained in more detail in Ref. [49].

Appendix E.

Oscillation probabilities for KamLAND and MINOS

We will derive the oscillation probabilities for the KamLAND and MINOS experiments, neglecting matter effects. This is a valid assumption for both the KamLAND baseline of 180 km and the MINOS baseline of 735 km since matter effects do not become significant until the baseline exceeds $\sim$ 1000 km. We can therefore start from the general vacuum oscillation probability, Eq. (2.8),

$$
\begin{aligned}
P_{\nu_\alpha \to \nu_\beta}(L,E) &= \delta_{\alpha\beta} && \text{(E.1a)}\\
&- 4\sum_{j>k} \mathrm{Re}[U^*_{\alpha j} U_{\beta j} U_{\alpha k} U^*_{\beta k}] \sin^2\left(\frac{\Delta m^2_{jk} L}{4E}\right) && \text{(E.1b)}\\
&+ 2\sum_{j>k} \mathrm{Im}[U^*_{\alpha j} U_{\beta j} U_{\alpha k} U^*_{\beta k}] \sin\left(\frac{\Delta m^2_{jk} L}{2E}\right). && \text{(E.1c)}
\end{aligned}
$$

In the case of KamLAND we need to calculate the $\bar{\nu}_e \to \bar{\nu}_e$ probability so in Eq. (E.1) we have $\alpha = \beta = e$; therefore (E.1a) is equal to 1 and (E.1c) is zero. KamLAND has

$\frac{L}{E} \sim \frac{180\text{km}}{1\text{MeV}}$ so that $\frac{\Delta m^2_{31} L}{4E} \simeq 100$ and $\frac{\Delta m^2_{21} L}{4E} \simeq 4$. Therefore the sum in (E.1b) is equal to

$$\begin{aligned}
& \sum_{j>k} \text{Re}[U^*_{\alpha j} U_{\beta j} U_{\alpha k} U^*_{\beta k}] \sin^2\left(\frac{\Delta m^2_{jk} L}{4E}\right) \\
= \; & |U_{e2}|^2 |U_{e1}|^2 \sin^2\left(\frac{\Delta m^2_{21} L}{4E}\right) + |U_{e3}|^2 (|U_{e1}|^2 + |U_{e2}|^2) \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right) && \text{(E.2)} \\
= \; & c^4_{13} s^2_{12} c^2_{12} \sin^2\left(\frac{\Delta m^2_{21} L}{4E}\right) + s^2_{13} c^2_{13} \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right) && \text{(E.3)} \\
\simeq \; & s^2_{12} c^2_{12} \sin^2\left(\frac{\Delta m^2_{21} L}{4E}\right) && \text{(E.4)} \\
& \text{using the approximation } s^2_{13} \simeq 0 \text{ and } c^2_{13} \simeq 1 \text{ because } \theta_{13} \text{ is small} \\
= \; & \frac{1}{4} s^2_{212} \sin^2\left(\frac{\Delta m^2_{21} L}{4E}\right), && \text{(E.5)}
\end{aligned}$$

and the total probability is then

$$P_{\bar{\nu}_e \to \bar{\nu}_e} \simeq 1 - \sin^2 2\theta_{12} \sin^2\left(\frac{\Delta m^2_{21} L}{4E}\right). \qquad \text{(E.6)}$$

For MINOS, we need to calculate the $\nu_\mu \to \nu_\mu$ probability; this is again a disappearance channel like the KamLAND channel so that (E.1a) is equal to 1 and (E.1c) is zero. MINOS has $\frac{L}{E} \sim \frac{735 \text{ km}}{1 \text{ GeV}}$ so that $\frac{\Delta m^2_{31} L}{4E} \simeq 2$ and $\frac{\Delta m^2_{21} L}{4E} \simeq 0.07$. (E.1b) is equal to

$$\begin{aligned}
& |U_{\mu 2}|^2 |U_{\mu 1}|^2 \sin^2\left(\frac{\Delta m^2_{21} L}{4E}\right) + |U_{\mu 3}|^2 (|U_{\mu 1}|^2 + |U_{\mu 2}|^2) \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right) \\
\simeq \; & |U_{\mu 3}|^2 (|U_{\mu 1}|^2 + |U_{\mu 2}|^2) \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right) && \text{(E.7)} \\
& \text{because} \sin^2\left(\frac{\Delta m^2_{21} L}{4E}\right) \simeq 0 \\
= \; & |U_{\mu 3}|^2 (1 - |U_{\mu 3}|^2) \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right) && \text{(E.8)} \\
& \text{using the unitarity condition } \textstyle\sum_{j=1}^3 U^*_{\alpha j} U_{\beta j} = |\sum_{j=1}^3 U^*_{\alpha j} U_{\beta j}|^2 = \delta_{\alpha\beta} \\
= \; & c^2_{13} s^2_{23} (1 - c^2_{13} s^2_{23}) \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right) && \text{(E.9)} \\
\simeq \; & s^2_{23} c^2_{23} \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right) && \text{(E.10)} \\
& \text{using } c^2_{13} \simeq 1 \\
= \; & \frac{1}{4} s^2_{223} \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right), && \text{(E.11)}
\end{aligned}$$

so that the probability is

$$P_{\nu_\mu \to \nu_\mu} \simeq 1 - \sin^2 2\theta_{23} \sin^2 \left(\frac{\Delta m^2_{31} L}{4E} \right). \tag{E.12}$$

Appendix F.

Calculation of the golden channel probability

The probability for the golden channel, $\nu_e \to \nu_\mu$, in vacuum, is an application of Eq. (2.8),

$$\begin{aligned} P_{\nu_e \to \nu_\mu} = & - 4 \sum_{j>k} \mathrm{Re}[U^*_{ej} U_{\mu j} U_{ek} U^*_{\mu k}] \sin^2 \left(\frac{\Delta m^2_{jk} L}{4E} \right) \\ & + 2 \sum_{j>k} \mathrm{Im}[U^*_{ej} U_{\mu j} U_{ek} U^*_{\mu k}] \sin \left(\frac{\Delta m^2_{jk} L}{2E} \right), \end{aligned} \tag{F.1}$$

where $U_{\alpha j}$ are the elements of the usual PMNS matrix given in Eq. (2.1). The final expression is given by Eq. (4.1),

$$\begin{aligned} P^{\mathrm{vac}}_{\nu_e \to \nu_\mu} = & \; s^2_{213} s^2_{23} \left(\sin^2 \left(\frac{\Delta_{31} L}{2} \right) + \sin^2 \left(\frac{\Delta_{32} L}{2} \right) \right) \\ & + \frac{1}{2} s_{213} c_{13} s_{212} s_{223} \cos\delta \left(\sin^2 \left(\frac{\Delta_{31} L}{2} \right) - \sin^2 \left(\frac{\Delta_{32} L}{2} \right) \right) \\ & + \frac{1}{4} s_{213} c_{13} s_{212} s_{223} \sin\delta \left(\sin\left(\Delta_{31} L\right) - \sin\left(\Delta_{32} L\right) \right) \\ & + \left(c^2_{13} s^2_{212} c^2_{23} - \frac{1}{4} s^2_{213} s^2_{212} s^2_{23} + s_{213} c_{13} s_{212} c_{212} s_{223} \cos\delta \right) \sin^2 \left(\frac{\Delta_{21} L}{2} \right) \\ & - \frac{1}{4} s_{213} c_{13} s_{212} s_{223} \sin\delta \sin\left(\Delta_{21} L\right), \end{aligned} \tag{F.2}$$

with the expression for the anti-neutrino probability obtained by setting $\delta \to -\delta$.

To calculate the probability in matter we use the method described in Section 2.4.2 to calculate a mixing matrix in matter, which transforms from the neutrino flavour states to states which propagate in matter. We start with the Hamiltonian in matter which

we will write in the flavour basis because the matter potential takes a simple diagonal form in this basis (only the electron component is non-zero):

$$\hat{H}_{\text{fl}} = U_{\text{eff}}^{\dagger} \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta_{21} & 0 \\ 0 & 0 & \Delta_{31} \end{pmatrix} U_{\text{eff}} \pm \begin{pmatrix} A & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{F.3}$$

where $A = \sqrt{2} G_F n_e$ (see Section 2.4), $U_{\text{eff}} = U_{23} U_{13}$ as explained in Section 2.4.2 and the upper (+) sign applies to neutrinos and the lower (−) sign to anti-neutrinos. In the limit $\Delta_{21} = 0$ this matrix is diagonalised to (see Ref. [64])

$$\hat{H}_{\text{fl}}^{(0)} = \bar{U}_{\mp}^{\dagger} \begin{pmatrix} \frac{\Delta_{31} \pm A - B_{\mp}}{2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \frac{\Delta_{31} \pm A + B_{\mp}}{2} \end{pmatrix} \bar{U}_{\mp}, \tag{F.4}$$

where

$$\bar{U}_{\mp} = U_{23}(\theta_{23}) U_{13}(\theta_{m\mp}), \tag{F.5a}$$

$$B_{\mp} = \sqrt{(\Delta_{31} \cos 2\theta_{13} \mp A)^2 + (\Delta_{31} \sin 2\theta_{13})^2}, \tag{F.5b}$$

$$\tan 2\theta_{m\mp} = \frac{\Delta_{31} \sin 2\theta_{13}}{\Delta_{31} \cos 2\theta_{13} \mp A}, \tag{F.5c}$$

and all the upper signs apply to neutrinos and the lower signs to anti-neutrinos. We use the superscript (0) to denote that this is the Hamiltonian in the limit $\Delta_{21} = 0$. Thus the mixing matrix in matter, in the limit $\Delta_{21} = 0$, is $\bar{U}_{\mp} = U_{23}(\theta_{23}) U_{13}(\theta_{m\mp})$. Explicitly, this is

$$\bar{U}_{\mp} = \begin{pmatrix} c_{\theta_{m\mp}} & 0 & s_{\theta_{m\mp}} e^{\pm i\delta} \\ -s_{\theta_{m\mp}} s_{23} e^{\mp i\delta} & c_{23} & s_{23} c_{\theta_{m\mp}} \\ -s_{\theta_{m\mp}} c_{23} e^{\mp i\delta} & -s_{23} & c_{23} c_{\theta_{m\mp}} \end{pmatrix}, \tag{F.6}$$

which we will now call $\bar{U}_{\mp}^{(0)}$. If we rewrite Eq. (F.4) in the conventional form for matrix diagonalisation, $\Lambda = S^{-1} M S$, where S is a matrix whose columns are the eigenvectors

of M and Λ is the diagonal matrix of eigenvalues, and use the fact that $\bar{U}_{\mp}^{\dagger} = \bar{U}_{\mp}^{-1}$, then

$$\begin{pmatrix} \frac{\Delta_{31} \pm A - B_{\mp}}{2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \frac{\Delta_{31} \pm A + B_{\mp}}{2} \end{pmatrix} = \bar{U}_{\mp} H_{\text{fl}}^{(0)} \bar{U}_{\mp}^{\dagger}. \tag{F.7}$$

Now it is clear that the columns of $\bar{U}_{\mp}$, which we will call $(\mathbf{v}_1^{(0)}, \mathbf{v}_2^{(0)}, \mathbf{v}_3^{(0)})$, are the eigenvectors in matter in the flavour basis with corresponding eigenvalues given by

$$\lambda_1^{(0)} = \frac{\Delta_{31} \pm A - B_{\mp}}{2}, \tag{F.8a}$$

$$\lambda_2^{(0)} = 0, \tag{F.8b}$$

$$\lambda_3^{(0)} = \frac{\Delta_{31} \pm A + B_{\mp}}{2}. \tag{F.8c}$$

We now want to find the first order corrections to $\bar{U}_{\mp}$ in the case of non-zero Δ_{21}. The first-order corrections to the eigenvectors are given by quantum-mechanical perturbation theory as

$$\mathbf{v}_j^{(1)} = \sum_{k \neq j} \frac{\hat{H}_{jk}^{(1)}}{\lambda_j^{(0)} - \lambda_k^{(0)}} \mathbf{v}_k^{(0)}, \tag{F.9}$$

and the first order corrections to the eigenvalues are

$$\lambda_j^{(1)} = \hat{H}_{jj}^{(1)}, \tag{F.10}$$

where $\hat{H}_{jk}^{(1)}$ are the elements of the first order perturbation to the Hamiltonian. In the vacuum (mass) basis, this perturbation is

$$\hat{H}_{\text{mass}}^{(1)} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta_{21} & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{F.11}$$

and in the flavour basis,

$$\hat{H}_{\text{fl}}^{(1)} = U^{\dagger} \hat{H}_{\text{mass}}^{(1)} U, \tag{F.12}$$

where U is the full PMNS matrix involving all three mixing angles. The matrix $\bar{U}^{(0)}_{\mp}$ transforms between the flavour and matter bases, such that

$$\hat{H}^{(1)}_{\rm fl} = (\bar{U}^{(0)}_{\mp})^{\dagger} \hat{H}^{(1)}_{\rm mat} \bar{U}^{(0)}_{\mp}, \tag{F.13}$$

and so $\hat{H}^{(1)}_{\rm mat}$ is given by

$$\begin{aligned} \hat{H}^{(1)}_{\rm mat} &= \bar{U}^{(0)}_{\mp} \hat{H}^{(1)}_{\rm fl} (\bar{U}^{(0)}_{\mp})^{\dagger} \\ &= \bar{U}^{(0)}_{\mp} U^{\dagger} \hat{H}^{(1)}_{\rm mass} U (\bar{U}^{(0)}_{\mp})^{\dagger}. \end{aligned} \tag{F.14}$$

Defining

$$\bar{\theta}_{\mp} = \theta_{13} - \theta_{m\mp}, \tag{F.15}$$

$\hat{H}^{(1)}_{\rm mat}$ is given by

$$\begin{aligned} \hat{H}^{(1)}_{\rm mat} &= \begin{pmatrix} c_{12}c_{\bar{\theta}_{\mp}} & s_{12}c_{\bar{\theta}_{\mp}} & -s_{\bar{\theta}_{\mp}}e^{\pm i\delta} \\ -s_{12} & c_{12} & 0 \\ -c_{12}s_{\bar{\theta}_{\mp}}e^{\mp i\delta} & -s_{12}s_{\bar{\theta}_{\mp}}e^{\mp i\delta} & c_{\bar{\theta}_{\mp}} \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta_{21} & 0 \\ 0 & 0 & 0 \end{pmatrix} \times \\ &\quad \begin{pmatrix} c_{12}c_{\bar{\theta}_{\mp}} & -s_{12} & -c_{12}s_{\bar{\theta}_{\mp}}e^{\pm i\delta} \\ s_{12}c_{\bar{\theta}_{\mp}} & c_{12} & -s_{12}s_{\bar{\theta}_{\mp}}e^{\pm i\delta} \\ -s_{\bar{\theta}_{\mp}}e^{\mp i\delta} & 0 & c_{\bar{\theta}_{\mp}} \end{pmatrix} \\ &= \begin{pmatrix} \Delta_{21}s^2_{12}c^2_{\bar{\theta}_{\mp}} & \frac{1}{2}\Delta_{21}s_{212}c_{\bar{\theta}_{\mp}} & -\frac{1}{2}\Delta_{21}s^2_{12}s_{2\bar{\theta}_{\mp}}e^{\pm i\delta} \\ \frac{1}{2}\Delta_{21}s_{212}c_{\bar{\theta}_{\mp}} & \Delta_{21}c^2_{12} & -\frac{1}{2}\Delta_{21}s_{212}s_{\bar{\theta}_{\mp}}e^{\pm i\delta} \\ -\frac{1}{2}\Delta_{21}s^2_{12}s_{2\bar{\theta}_{\mp}}e^{\pm i\delta} & -\frac{1}{2}\Delta_{21}s_{212}s_{\bar{\theta}_{\mp}}e^{\mp i\delta} & \Delta_{21}s^2_{12}s^2_{\bar{\theta}_{\mp}} \end{pmatrix}. \end{aligned} \tag{F.16}$$

Using Eq. (F.9), the first order corrections to the eigenvectors (columns of $\bar{U}_{\mp}$) are

$$\mathbf{v}_1^{(1)} = \frac{\Delta_{21} s_{212} c_{\bar{\theta}_{\mp}}}{\Delta_{31} \pm A - B} \begin{pmatrix} 0 \\ c_{23} \\ -s_{23} \end{pmatrix} + \frac{\Delta_{21} s_{12}^2 s_{2\bar{\theta}_{\mp}} e^{\pm i\delta}}{2B} \begin{pmatrix} s_{\bar{\theta}_{\mp}} e^{\pm i\delta} \\ s_{23} c_{\bar{\theta}_{\mp}} \\ c_{23} c_{\bar{\theta}_{\mp}} \end{pmatrix}, \tag{F.17a}$$

$$\mathbf{v}_2^{(1)} = \frac{\Delta_{21} s_{212} c_{\bar{\theta}_{\mp}}}{-\Delta_{31} \mp A + B} \begin{pmatrix} c_{\bar{\theta}_{\mp}} \\ -s_{23} s_{\bar{\theta}_{\mp}} e^{\mp i\delta} \\ -c_{23} s_{\bar{\theta}_{\mp}} e^{\mp i\delta} \end{pmatrix} + \frac{\Delta_{21} s_{212} s_{\bar{\theta}_{\mp}} e^{\mp i\delta}}{\Delta_{31} \pm A + B} \begin{pmatrix} s_{\bar{\theta}_{\mp}} e^{\pm i\delta} \\ s_{23} c_{\bar{\theta}_{\mp}} \\ c_{23} c_{\bar{\theta}_{\mp}} \end{pmatrix} \tag{F.17b}$$

$$\mathbf{v}_3^{(1)} = \frac{-\Delta_{21} s_{12}^2 s_{2\bar{\theta}_{\mp}} e^{\pm i\delta}}{2B} \begin{pmatrix} c_{\bar{\theta}_{\mp}} \\ -s_{23} s_{\bar{\theta}_{\mp}} e^{\mp i\delta} \\ -c_{23} s_{\bar{\theta}_{\mp}} e^{\mp i\delta} \end{pmatrix} - \frac{\Delta_{21} s_{212} s_{\bar{\theta}_{\mp}} e^{\pm i\delta}}{\Delta_{31} \pm A + B} \begin{pmatrix} 0 \\ c_{23} \\ -s_{23} \end{pmatrix}, \tag{F.17c}$$

so that the mixing matrix, to first order in Δ_{21}, is

$$\bar{U}_{\mp} = \begin{pmatrix} c_{\bar{\theta}_{\mp}} & 0 & s_{\bar{\theta}_{\mp}} e^{\mp i\delta} \\ -s_{\bar{\theta}_{\mp}} s_{23} e^{\pm i\delta} & c_{23} & s_{23} c_{\bar{\theta}_{\mp}} \\ -s_{\bar{\theta}_{\mp}} c_{23} e^{\pm i\delta} & -s_{23} & c_{23} c_{\bar{\theta}_{\mp}} \end{pmatrix} + \begin{pmatrix} | & | & | \\ \mathbf{v}_1^{(1)} & \mathbf{v}_2^{(1)} & \mathbf{v}_3^{(1)} \\ | & | & | \end{pmatrix}, \tag{F.18}$$

which we will denote by

$$\bar{U}_{\mp} = \begin{pmatrix} | & | & | \\ \mathbf{V}_1^{\Delta_{21}} & \mathbf{V}_2^{\Delta_{21}} & \mathbf{V}_3^{\Delta_{21}} \\ | & | & | \end{pmatrix}. \tag{F.19}$$

The probability for $\nu_e \to \nu_\mu$ in matter is given by

$$\begin{aligned} P_{\nu_e \to \nu_\mu} = & - 4 \sum_{j>k} \mathrm{Re}[(\bar{U}_{\mp})^*_{ej} (\bar{U}_{\mp})_{\mu j} (\bar{U}_{\mp})_{ek} (\bar{U}_{\mp})^*_{\mu k}] \sin^2 \left(\frac{\Delta m^2_{jk} L}{4E} \right) \\ & + 2 \sum_{j>k} \mathrm{Im}[(\bar{U}_{\mp})^*_{ej} (\bar{U}_{\mp})_{\mu j} (\bar{U}_{\mp})_{ek} (\bar{U}_{\mp})^*_{\mu k}] \sin \left(\frac{\Delta m^2_{jk} L}{2E} \right), \end{aligned} \tag{F.20}$$

which is exactly analagous to the vacuum solution, Eq. (F.1), but with the replacement of U with $\bar{U}_{\mp}$. Keeping only the terms up to first order in Δ_{21}, the explicit form of the

solution is

$$
\begin{aligned}
P_{\nu_e\to\nu_\mu} &= s^2_{23}s^2_{2\theta_{m\mp}}\sin^2\left(\frac{B_\mp L}{2}\right) \\
&- s^2_{23}s^2_{12}\left(s_{4\theta_{m\mp}}s_{2\bar{\theta}_\mp}\sin^2\left(\frac{B_\mp L}{2}\frac{\Delta_{21}}{B_\mp}+s^2_{2\theta_{m\mp}}c_{2\bar{\theta}_\mp}\sin(B_\mp L)\frac{\Delta_{21}L}{2}\right)\right) \\
&+ s_{212}s_{223}s_{2\theta_{m\mp}}\sin\left(\frac{B_\mp L}{2}\right)\Delta_{21}\times \\
&\quad\left[\sin\left(\frac{\lambda^{(0)}_1 L}{2}\right)\cos\left(\pm\delta-\frac{\lambda^{(0)}_3 L}{2}\right)\left(\frac{c_{\theta_{m\mp}}c_{\bar{\theta}_\mp}}{\lambda^{(0)}_1}-\frac{s_{\theta_{m\mp}}s_{\bar{\theta}_\mp}}{\lambda^{(0)}_3}\right)-s_{\theta_{m\mp}}s_{\bar{\theta}_\mp}\cos\delta\sin\left(\frac{B_\mp L}{2}\right)\frac{1}{\lambda^{(0)}_3}\right].
\end{aligned}
\tag{F.21}
$$

If we expand $\theta_{m\mp}$ and $B_\mp$ in terms of θ_{13} and keep only terms up to second order in θ_{13} (or rather, s_{13} and s_{213}), then we recover the expression given in Eq. (4.4).

Appendix G.

The golden channel probability including non-standard interactions

Let us first discuss the possible processes which can lead to the apparent detection of a golden channel event ($\nu_e \to \nu_\mu$, and similarly for anti-neutrinos). For example, in a neutrino factory, it is *assumed* that the neutrino produced at the source from the decay of a μ^+, together with $\bar{\nu}_\mu$ and e^+, is a ν_e, and that when a μ^- is detected it was accompanied by a ν_μ. However, non-standard interactions (NSI's) at the source or detector can mean that these assumptions may not always be true. We shall consider these source and detector NSI's first, and discuss matter (propagation) NSI's afterwards. There are several transitions which can lead to a $\nu_e \to \nu_\mu$ transition with only source and detector NSI's. We will use the notation $\Gamma_{\alpha\beta}$ to denote the amplitude for the standard oscillation $\nu_\alpha \to \nu_\beta$ at a baseline L so that

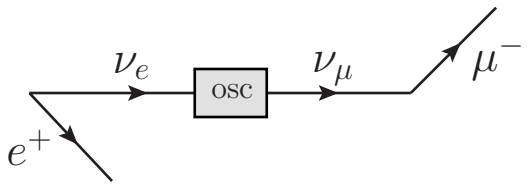

$$= \Gamma_{e\mu}.$$

Now we could have source NSI's (ε^s) which contribute to the transition amplitude in processes such as

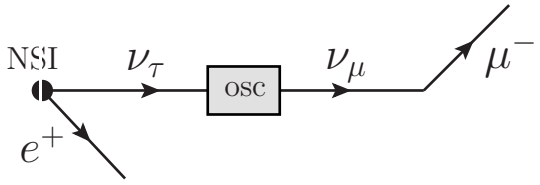

$$= \varepsilon^s_{e\tau}\Gamma_{\tau\mu},$$

and detector NSI's (ε^d) in processes such as

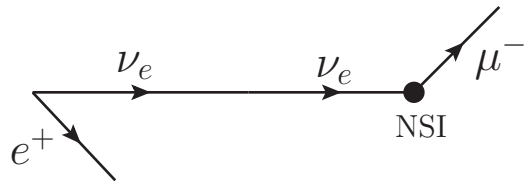

$= \varepsilon^d_{e\mu}$.

To obtain the total probability we have to add and square all the amplitudes:

$$\begin{aligned} P_{\nu_e\to\nu_\mu}(\varepsilon^s,\,\varepsilon^d) &= |\Gamma_{e\mu} + \varepsilon^s_{ee}\Gamma_{e\mu} + \varepsilon^s_{e\mu}\Gamma_{\mu\mu} + \varepsilon^s_{e\tau}\Gamma_{\tau\mu} + \Gamma_{ee}\varepsilon^d_{e\mu} + \Gamma_{e\mu}\varepsilon^d_{\mu\mu} + \Gamma_{e\tau}\varepsilon^d_{\tau\mu} + O(\Gamma^2,\,\varepsilon^2)|^2 \\ &= |\Gamma_{e\mu}|^2 + 2\mathrm{Re}[\Gamma^*_{e\mu}(\varepsilon^s_{ee}\Gamma_{e\mu} + \varepsilon^s_{e\mu}\Gamma_{\mu\mu} + \varepsilon^s_{e\tau}\Gamma_{\tau\mu} + \Gamma_{ee}\varepsilon^d_{e\mu} + \Gamma_{e\mu}\varepsilon^d_{\mu\mu} + \Gamma_{e\tau}\varepsilon^d_{\tau\mu})] \\ &\quad + O(\Gamma^3,\,\varepsilon^2). \end{aligned} \tag{G.1}$$

To consider the effect of each of the NSI parameters, we need to remember that the golden channel oscillation probability is suppressed by $\sin^2 2\theta_{13}$ (see Section 4.1) whereas the disappearance probabilities, $\nu_e \to \nu_e$ and $\nu_\mu \to \nu_\mu$, are $O(1)$, as is $\nu_\mu \to \nu_\tau$ and $\nu_\tau \to \nu_\mu$ (at the oscillation maximum, $\Gamma_{ee} \sim 1$ and $\Gamma_{\mu\mu} \sim \Gamma_{\mu\tau} \sim \Gamma_{\tau\mu} \sim \frac{1}{2}$). So the NSI's which are linked to these channels (all of those listed above except for ε^s_{ee}) may significantly enhance the transition rate if $\varepsilon > \Gamma_{e\mu}$.

However, we are interested in matter NSI's which cannot be calculated in the same way as above, because matter NSI's alter the propagation Hamiltonian (source and detector NSI's do not) which means that they cannot be decoupled from and treated separately to oscillations. Matter NSI's add additional terms on to the propagation Hamiltonian which we can treat as a perturbation as we will assume that they are small ($\varepsilon_{\alpha\beta} \lesssim \frac{\Delta_{21}}{\Delta_{31}}$). We will then use an identical method to that used in Appendix F, when Δ_{21} was treated as a perturbation. In this case we take the unperturbed Hamiltonian to be the Hamiltonian in matter to first order in Δ_{21} ($\hat{H}^{(0)} + \hat{H}^{(1)}$ as given by Eqs. (F.4) and (F.12)) and the unperturbed mixing matrix to be the matrix which is first order in Δ_{21} (Eq. (F.18)).

The perturbation from matter NSI's, in the flavour basis, is

$$\hat{H}^{\mathrm{fl}}_{\mathrm{NSI}} = A \begin{pmatrix} \varepsilon_{ee} & \varepsilon_{e\mu} & \varepsilon_{e\tau} \\ \varepsilon_{\mu e} & \varepsilon_{\mu\mu} & \varepsilon_{\mu\tau} \\ \varepsilon_{\tau e} & \varepsilon_{\mu\tau} & \varepsilon_{\tau\tau} \end{pmatrix}, \tag{G.2}$$

which must necessarily be Hermitian as the Hamiltonian must be Hermitian. Thus the diagonal entries must be real whereas the off-diagonal entries can be complex with possible CP violating phases. So we can rewrite this as

$$\hat{H}_{\rm fl}^{\rm NSI} = A \begin{pmatrix} \varepsilon_{ee} & \varepsilon_{e\mu}e^{\pm i\phi_{e\mu}} & \varepsilon_{e\tau}e^{\pm i\phi_{e\tau}} \\ \varepsilon_{e\mu}e^{\mp i\phi_{e\mu}} & \varepsilon_{\mu\mu} & \varepsilon_{\mu\tau}e^{\pm i\phi_{\mu\tau}} \\ \varepsilon_{e\tau}e^{\mp i\phi_{e\tau}} & \varepsilon_{\mu\tau}e^{\pm i\phi_{\mu\tau}} & \varepsilon_{\tau\tau} \end{pmatrix}, \tag{G.3}$$

where here the $\varepsilon_{\alpha\beta}$ are real parameters which are the moduli of the NSI's. So there are nine parameters (six moduli and three phases) in the perturbation. To transform to the matter basis, we use the relation

$$\hat{H}_{\rm NSI}^{\rm mat} = \bar{U}_{\mp}\hat{H}_{fl}^{\rm NSI}\bar{U}_{\mp}^{\dagger}. \tag{G.4}$$

We will only consider the mixing matrix to zeroth order in Δ_{21}, $\bar{U}_{\mp} = \bar{U}_{\mp}^{(0)}$, as given in Eq. (F.6) because ultimately, in the final probability, we shall only be keeping terms which are second order in (θ_{13}, Δ_{21}, $\varepsilon_{\alpha\beta}$) so at this stage it is adequate to keep only the perturbative terms which are first order in these quantities. For simplicity, we can 'switch on' one parameter at a time, setting all the others to zero, then add all the terms together at the end. As an example, we shall consider the case when only $\varepsilon_{e\tau} \neq 0$. In this case the perturbation is (we shall suppress the subscripts $e\tau$ to simplify the notation so that $\varepsilon = \varepsilon_{e\tau}$ and $\phi = \phi_{e\tau}$)

$$\begin{aligned} \hat{H}_{\rm fl}^{\rm NSI} &= \bar{U}_{\mp}^{(0)} A \begin{pmatrix} 0 & 0 & \varepsilon e^{\pm i\phi} \\ 0 & 0 & 0 \\ \varepsilon e^{\mp i\phi} & 0 & 0 \end{pmatrix} (\bar{U}_{\mp}^{(0)})^{\dagger} \qquad \text{(G.5)} \\ &= A\varepsilon \begin{pmatrix} -2s_{\bar{\theta}_{\mp}}c_{\bar{\theta}_{\mp}}c_{23}\cos(\phi-\delta) & -c_{\bar{\theta}_{\mp}}s_{23}e^{\pm i\phi} & c_{23}(c^2_{\bar{\theta}_{\mp}}e^{\pm i\phi} - s^2_{\bar{\theta}_{\mp}}e^{\mp i(\phi-2\delta)}) \\ -c_{\bar{\theta}_{\mp}}s_{23}e^{\mp i\phi} & 0 & -s_{\bar{\theta}_{\mp}}s_{23}e^{\mp i(\phi-\delta)} \\ c_{23}(c^2_{\bar{\theta}_{\mp}}e^{\mp i\phi} - s^2_{\bar{\theta}_{\mp}}e^{\pm i(\phi-2\delta)}) & -s_{\bar{\theta}_{\mp}}s_{23}e^{\pm i(\phi-\delta)} & 2s_{\bar{\theta}_{\mp}}c_{\bar{\theta}_{\mp}}c_{23}\cos(\phi-\delta) \end{pmatrix}. \end{aligned}$$

We will again use Eq. (F.9) to find the corrections to the eigenvectors of the mixing matrix, Eq. (F.18), to first order in ε. We will then be able to write a mixing matrix which is first order both in Δ_{21} and in ε.

We will be calculating the transition probability for the process $\nu_e \to \nu_\mu$, up to second order in the quantities (s_{13}, Δ_{21}, ε). We know from Eq. (2.8) that the oscillation

probability is fourth order in the matrix elements; therefore it is sufficient to keep only terms which are first order in (θ_{13}, Δ_{21}, ε) when calculating the corrections to the mixing matrix. All the elements of $\hat{H}^{\rm NSI}_{fl}$ are already first order in ε, so for the denominator of Eq. (F.9) we only need to include the leading order contribution of the eigenvectors i.e. we only need $\Gamma^{(0)}_j$ and can neglect $\Gamma^{(1)}_j$. Similarly, we can neglect factors $\sim s_{13}$ so that $B_\mp \simeq \Delta_{31} \mp A$. With these simplifications, we then have the following first-order corrections to the eigenvectors, where the vectors $\mathbf{V}^{\Delta_{21}}_j$ denote the columns of the matrix Eq. (F.18):

$$\mathbf{V}^\varepsilon_1 = v^\varepsilon_{12}\mathbf{V}^{\Delta_{21}}_2 + v^\varepsilon_{13}\mathbf{V}^{\Delta_{21}}_3, \tag{G.6a}$$
$$\mathbf{V}^\varepsilon_2 = v^\varepsilon_{21}\mathbf{V}^{\Delta_{21}}_1 + v^\varepsilon_{23}\mathbf{V}^{\Delta_{21}}_3, \tag{G.6b}$$
$$\mathbf{V}^\varepsilon_3 = v^\varepsilon_{31}\mathbf{V}^{\Delta_{21}}_1 + v^\varepsilon_{32}\mathbf{V}^{\Delta_{21}}_2, \tag{G.6c}$$

where

$$v^\varepsilon_{12} = -\varepsilon c_{\bar{\theta}_\mp} s_{23} e^{\mp i\phi}, \tag{G.7a}$$
$$v^\varepsilon_{13} = \frac{\pm A\varepsilon c^2_{\bar{\theta}_\mp} c_{23} e^{\mp i\phi}}{\pm A - \Delta_{31}}, \tag{G.7b}$$
$$v^\varepsilon_{21} = \varepsilon c_{\bar{\theta}_\mp} s_{23} e^{\pm i\phi}, \tag{G.7c}$$
$$v^\varepsilon_{23} = \frac{\pm A\varepsilon s_{\bar{\theta}_\mp} s_{23} e^{\pm i(\phi-\delta)}}{\Delta_{31}}, \tag{G.7d}$$
$$v^\varepsilon_{31} = \frac{\pm A\varepsilon c^2_{\bar{\theta}_\mp} c_{23} e^{\pm i\phi}}{\Delta_{31} \mp A}, \tag{G.7e}$$
$$v^\varepsilon_{32} = \frac{\mp A\varepsilon s_{\bar{\theta}_\mp} s_{23} e^{\mp i(\phi-\delta)}}{\Delta_{31}}. \tag{G.7f}$$

Our mixing matrix in matter, to first order in Δ_{21} and ε, is then

$$\bar{U}_\mp(\Delta_{21}, \varepsilon) = \begin{pmatrix} | & | & | \\ \mathbf{V}^{\Delta_{21}}_1 & \mathbf{V}^{\Delta_{21}}_2 & \mathbf{V}^{\Delta_{21}}_3 \\ | & | & | \end{pmatrix} + \begin{pmatrix} | & | & | \\ \mathbf{V}^\varepsilon_1 & \mathbf{V}^\varepsilon_2 & \mathbf{V}^\varepsilon_3 \\ | & | & | \end{pmatrix}. \tag{G.8}$$

Keeping only terms which are at most first order in $s_{\bar{\theta}_\mp}$ and Δ_{21} the first term is

$$\bar{U}_\mp(\Delta_{21}) = \begin{pmatrix} c_{\bar{\theta}_\mp} & \mp\frac{c_{\bar{\theta}_\mp}\Delta_{21}s_{212}}{2A} & s_{\bar{\theta}_\mp}e^{\mp i\delta} \\ \pm\frac{\Delta_{21}c_{23}s_{212}}{2A} - s_{\bar{\theta}_\mp}s_{23}e^{\pm i\delta} & c_{23} & c_{\bar{\theta}_\mp}s_{23} \\ \mp\frac{\Delta_{21}s_{23}s_{212}}{2A} - s_{\bar{\theta}_\mp}c_{23}e^{\pm i\delta} & -s_{23} & c_{\bar{\theta}_\mp}c_{23} \end{pmatrix}, \tag{G.9}$$

and the second term, $\bar{U}_\mp(\varepsilon)$, will be proportional to ε. In other words we can write

$$[\bar{U}_\mp(\Delta_{21},\,\varepsilon)]_{\alpha j} = [\bar{U}_\mp^{\Delta_{21}}]_{\alpha j} + v^\varepsilon_{jk}[\bar{U}_\mp^{\Delta_{21}}]_{\alpha k}. \tag{G.10}$$

This will simplify the final step, which is to calculate the transition probability for $\nu_e \to \nu_\mu$. In this situation it will be easier to start from the form of the probability given in Eq. (D.2), with $U_{\alpha\beta}$ replaced by the matrix $\bar{U}_\mp(\Delta_{21},\,\varepsilon)$, and E_j replaced with the eigenvalues of $\bar{U}_\mp(\Delta_{21},\,\varepsilon)$ which we shall call Λ_j. These eigenvalues are the sum of the eigenvalues of $\bar{U}_\mp(\Delta_{21})$, to first order in Δ_{21} (the sum of Eq. (F.8) and the diagonal elements of Eq. (F.16)), which we shall call λ_j, and the perturbation $(\hat{H}^{\rm NSI})_{jj}$:

$$\begin{aligned}
\lambda_1 &= \pm A + \Delta_{21}s^2_{12}c^2_{\bar{\theta}_\mp}, && \text{(G.11a)}\\
\lambda_2 &= \Delta_{21}c^2_{12}, && \text{(G.11b)}\\
\lambda_3 &= \Delta_{31} + \Delta_{21}s^2_{12}s^2_{\bar{\theta}_\mp}, && \text{(G.11c)}\\
\Lambda_1 &= \pm A + \Delta_{21}s^2_{12}c^2_{\bar{\theta}_\mp} \mp 2A\varepsilon s_{\bar{\theta}_\mp}c_{\bar{\theta}_\mp}c_{23}\cos(\phi-\delta), && \text{(G.11d)}\\
\Lambda_2 &= \Delta_{21}c^2_{12}, && \text{(G.11e)}\\
\Lambda_3 &= \Delta_{31} + \Delta_{21}s^2_{12}s^2_{\bar{\theta}_\mp} \pm 2A\varepsilon s_{\bar{\theta}_\mp}c_{\bar{\theta}_\mp}c_{23}\cos(\phi-\delta), && \text{(G.11f)}
\end{aligned}$$

where we have used $B_\mp \simeq \Delta_{31} \mp A$. If we call $\bar{U}_\mp(\Delta_{21},\,\varepsilon)$ simply $\tilde{U}$ for simplicity then the transition probability is

$$\begin{aligned}
P_{\nu_e\to\nu_\mu}(\Delta_{21},\,\varepsilon) &= (\tilde{U}^*_{e1}\tilde{U}_{\mu1}e^{i\Lambda_1 L} + \tilde{U}^*_{e2}\tilde{U}_{\mu2}e^{i\Lambda_2 L} + \tilde{U}^*_{e3}\tilde{U}_{\mu3}e^{i\Lambda_3 L}) \times \\
&\quad (\tilde{U}_{e1}\tilde{U}^*_{\mu1}e^{-i\Lambda_1 L} + \tilde{U}_{e2}\tilde{U}^*_{\mu2}e^{-i\Lambda_2 L} + \tilde{U}_{e3}\tilde{U}^*_{\mu3}e^{-i\Lambda_3 L}). \qquad \text{(G.12)}
\end{aligned}$$

We have already calculated the probability in matter, to second order in Δ_{21}, in the absence of NSI's - this is the standard oscillation probability given in Eq. (4.4). Therefore we wish to calculate the additional terms $\sim \varepsilon$. As we are working up to second order in the small quantities, these will be terms $\sim \varepsilon s_{\bar{\theta}_\mp}$, $\sim \varepsilon\Delta_{21}$ or $\sim \varepsilon^2$. Using Eq. (G.10) and

calling $\bar{U}_{\mp}^{\Delta_{21}}$ simply $\dot{U}$, the relevant terms are

$$\begin{aligned}
P_{\nu_e\to\nu_\mu}(\varepsilon) &= 2\mathrm{Re}[|\dot{U}_{e1}|^2\dot{U}^*_{\mu1}(v^\varepsilon_{12}\dot{U}_{\mu2}+v^\varepsilon_{13}\dot{U}_{\mu3}) \\
&+ |\dot{U}_{\mu2}|^2\dot{U}^*_{e2}v^\varepsilon_{21}\dot{U}_{e1} \\
&+ |\dot{U}_{\mu3}|^2\dot{U}^*_{e3}v^\varepsilon_{31}\dot{U}_{e1} \\
&+ \dot{U}^*_{e1}\dot{U}^*_{\mu2}[\dot{U}_{\mu1}v^\varepsilon_{21}\dot{U}_{e1}+\dot{U}_{e2}(v^\varepsilon_{12}\dot{U}_{\mu2}+v^\varepsilon_{13}\dot{U}_{\mu3})]e^{i(\Lambda_1-\Lambda_2)L} \\
&+ \dot{U}^*_{e1}\dot{U}^*_{\mu3}[\dot{U}_{\mu1}v^\varepsilon_{31}\dot{U}_{e1}+\dot{U}_{e3}(v^\varepsilon_{12}\dot{U}_{\mu2}+v^\varepsilon_{13}\dot{U}_{\mu3})]e^{i(\Lambda_1-\Lambda_3)L} \\
&+ \dot{U}_{\mu2}\dot{U}^*_{\mu3}[\dot{U}^*_{e2}v^\varepsilon_{31}\dot{U}_{e1}+\dot{U}_{e3}v^\varepsilon_{21}\dot{U}_{e1}]e^{i(\Lambda_2-\Lambda_3)L}]. \qquad \text{(G.13)}
\end{aligned}$$

We can use the same method to calculate the contributions from ε^m_{ee}, $\varepsilon^m_{e\mu}$, $\varepsilon^m_{\mu\mu}$, $\varepsilon^m_{\mu\tau}$ and $\varepsilon^m_{\tau\tau}$. We find that, to second order in θ_{13}, Δ_{21} and ε, the only NSI parameters which contribute to the golden channel are $\varepsilon^m_{e\mu}$ and $\varepsilon^m_{e\tau}$, which contribute the terms

$$\begin{aligned}
P^\varepsilon_{\nu_e\to\nu_\mu} &= 4\varepsilon_{e\mu}s_{213}c_{23}s^2_{23}\sin\left(\frac{AL}{2}\right)\frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}}\cos\left(\frac{\Delta_{31}L}{2}-\delta-\phi_{e\mu}\right) \\
&- 4\varepsilon_{e\mu}\frac{\Delta_{21}}{\Delta_{31}}s_{212}c^2_{23}s^2_{23}\sin\left(\frac{AL}{2}\right)\frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}}\cos\left(\frac{\Delta_{31}L}{2}+\phi_{e\mu}\right) \\
&+ 4\varepsilon^2_{e\mu}c^2_{23}s^2_{23}\sin^2\left(\frac{AL}{2}\right) \\
&- 4\varepsilon_{e\tau}s_{213}c_{23}s^2_{23}\sin\left(\frac{AL}{2}\right)\frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}}\cos\left(\frac{\Delta_{31}L}{2}-\delta-\phi_{e\tau}\right) \\
&+ 4\varepsilon_{e\tau}\frac{\Delta_{21}}{\Delta_{31}}s_{212}c^2_{23}s^2_{23}\sin\left(\frac{AL}{2}\right)\frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}}\cos\left(\frac{\Delta_{31}L}{2}+\phi_{e\tau}\right) \\
&+ 4\varepsilon^2_{e\tau}c^2_{23}s^2_{23}\sin^2\left(\frac{AL}{2}\right), \qquad \text{(G.14)}
\end{aligned}$$

which is in agreement with the solutions obtained by Refs. [139] and [140] even though we have used a different method.

Appendix H.

An algebraic treatment of degeneracies

Future neutrino oscillation experiments seek to measure θ_{13}, δ and the mass hierarchy. In addition, we would like to measure other unknown quantities e.g. the quadrant of θ_{23}, and non-standard interactions. We need to consider how these unknown quantities are correlated in order to establish how the measurement of one parameter affects the measurement of another, and so that we can find solutions in the case that these correlations destroy the experimental sensitivity. As an example, we will look at the correlation between θ_{13} and δ. The observation of CP violation in the neutrino sector is one of the main goals of a neutrino factory and related experiments, and so we need to ask how, if θ_{13} is unknown, this will affect our sensitivity to CP violation. In other words, we need to know how many values of (θ_{13}, δ), in addition to the true values, which we will call ($\bar{\theta}_{13}$, $\bar{\delta}$), the data can be fitted to. In order to do this we will return once again to the golden channel probability in matter, Eq. (4.4),

$$\begin{aligned} P_{\nu_e\to\nu_\mu} &= s^2_{213}s^2_{23}\frac{\sin^2\left(\frac{(\Delta_{31}-A)L}{2}\right)}{\left(1-\frac{A}{\Delta_{31}}\right)^2} \\ &+ s_{213}c_{13}s_{212}s_{223}\frac{\Delta_{21}}{A}\sin\left(\frac{AL}{2}\right)\frac{\sin\left(\frac{(\Delta_{31}-A)L}{2}\right)}{1-\frac{A}{\Delta_{31}}}\cos\left(\frac{\Delta_{31}L}{2}-\delta\right) \\ &+ c^2_{13}s^2_{212}c^2_{23}\left(\frac{\Delta_{21}}{A}\right)^2\sin^2\left(\frac{AL}{2}\right), \end{aligned} \tag{H.1}$$

where $A = \sqrt{2}G_F n_e$. We are interested only in the parameters θ_{13} and δ so let us rewrite this equation more compactly as

$$P_{\nu_e \to \nu_\mu} = X s_{213}^2 + Y s_{213} \cos\left(\frac{\Delta_{31} L}{2} - \delta\right) + Z. \tag{H.2}$$

We want to find s_{213} as a function of δ by solving

$$\begin{aligned} & P(s_{213},\, \delta) = P(\bar{s}_{213},\, \bar{\delta}) \qquad (H.3) \\ \rightarrow \quad & X s_{213}^2 + Y s_{213} \cos\left(\frac{\Delta_{31} L}{2} - \delta\right) + Z = P(\bar{s}_{213},\, \bar{\delta}) \\ \rightarrow \quad & s_{213} = \frac{-Y}{2X} \cos\left(\frac{\Delta_{31} L}{2} - \delta\right) \pm \sqrt{\frac{Y^2}{4X^2} \cos^2\left(\frac{\Delta_{31} L}{2} - \delta\right) + \frac{1}{X}(P(\bar{s}_{213},\, \bar{\delta}) - Z)}. \end{aligned}$$

From this we see that it is possible to fit to *any* value of δ if there are no constraints on θ_{13} - Eq. (H.3) is a continuous locus of points in the $s_{213} - \delta$ plane. The accuracy to which δ can be constrained is thus correlated to the accuracy with which θ_{13} is constrained, *if we have only a single channel and baseline.* However, in all the experimental setups we have considered in this thesis, we have assumed that there will be both neutrino and anti-neutrino running. For anti-neutrinos we can obtain a similar expression to Eq. (H.3) by exchanging $A \leftrightarrow -A$ and $\delta \leftrightarrow -\delta$. The question then is whether the neutrino and anti-neutrino curves *intersect at more than one point?* If they do, then this means that there is at least one other pair of solutions $(s_{213},\, \delta)$, other than the true solution, $(\bar{s}_{213},\, \bar{\delta})$, and that it is impossible to determine which is the true solution if we do not have additional information. If, however, the curves only intersect at a single point (the true solution) then there are no degenerate solutions, and the combination of neutrino and anti-neutrino data is sufficient to enable us to determine the true values of θ_{13} and δ.

This method can also be applied to the other unknown parameters, for example to see how the sign of Δm_{31}^2 or the octant of θ_{23} is correlated to θ_{13} and δ, and also to examine the correlations between NSI and standard oscillation parameters.

Appendix I.

Simulation details

In this appendix we provide a summary of the AEDL (Abstract Experiment Definition Language) files which were used with GLoBES to implement the simulations described in Chapter 5 (Table I.1) and in Chapters 6 and 7 (Table I.2). We present the reference setups only; refer to the main text to obtain the details of each of the individual simulations.

We will briefly describe the 'background' entries in the tables: for the super-beam, '$\nu_e \to \nu_e$' refers to the intrinsic beam background and the 'wrong polarity' background refers to the anti-neutrino component of the neutrino beam, and neutrino component of the anti-neutrino beam. For the LENF, the backgrounds on each channel are estimated to arise primarily from charge misidentification and neutral-current events; the entries are a fraction of the detected channel rates (channel rates including detector efficiencies) which result in leptons of the wrong sign to that of the signal channel.

	Liquid argon	**Water Čerenkov**	**Liquid scintillator**
Flux	from A. Longhin	from A. Longhin	from A. Longhin
Running time/ polarity	2 years ν 8 years $\bar{\nu}$	2 years ν 8 years $\bar{\nu}$	2 years ν 8 years $\bar{\nu}$
Detector mass	100 kton	440 kton	50 kton
Energy threshold	0.1 GeV	0.1 GeV	1 GeV
Maximum energy	10 GeV	10 GeV	7 GeV
Bin widths/ GeV	0.15 for first bin 0.25 for all others	0.4,0.5,0.5, 0.5,1.0,7.0	0.05 for $1 < E < 3$ 0.1 for $3 < E < 5$ 0.25 for $5 < E < 7$
Matter profile	constant average	constant average	constant average
Baseline	2285 km	2285 km	2285 km
Energy resolution	matrices from L. Esposito	$0.007\sqrt{E} + 0.017$ for $\mu^{\pm}$ $0.026\sqrt{E} + 0.006$ for $e^{\pm}$	$0.05E$ for $\mu^{\pm}$ $0.05E$ for $e^{\pm}$
Cross-sections	From Refs. [143, 144]	From Refs. [143, 144]	From Refs. [143, 144]
Efficiency for $\mu^{\pm}$ disappearance	80%	40%	90%
Background on $\mu^{\pm}$ disappearance	50% $\nu_e \to \nu_\mu$ 5% wrong polarity 0.5% NC	40% $\nu_e \to \nu_\mu$ 5% wrong polarity 5% NC	50% $\nu_e \to \nu_\mu$ 5% wrong polarity 0.5% NC
Efficiency for $e^{\pm}$ appearance	80%	40%	90%
Background on $e^{\pm}$ appearance	50% $\nu_e \to \nu_e$ 5% wrong polarity 0.5% NC	40% $\nu_e \to \nu_e$ 5% wrong polarity 5% NC	50% $\nu_e \to \nu_e$ 5% wrong polarity 0.5% NC
Systematics on signal	5%	5%	5%
Systematics on background	5%	5%	5%

Table I.1.: Description of the super-beam reference setups used in Chapter 5.

	Totally active scintillating detector	**Liquid argon**
Flux	built-in GLoBES flux	built-in GLoBES flux
Muon energy	4.5 GeV	4.5 GeV
Muon decays/ year/ polarity	1.4×10^{21}	1.4×10^{21}
Running time/ polarity	10 (5) years	10 (5) years
Detector fiducial mass	20 kton	100 kton
Energy threshold	0.5 GeV	0.5 GeV
Maximum energy	6 GeV	6 GeV
Bin widths/ GeV	0.1,0.2,0.2,0.2,0.2,0.2,0.2, 0.25,0.25,0.25,0.25,0.3,0.3, 0.4,0.5,0.5,0.5,0.5	0.1,0.2,0.2,0.2,0.2,0.2,0.2, 0.25,0.25,0.25,0.25,0.3,0.3, 0.4,0.5,0.5,0.5,0.5
Matter profile	constant average	constant average
Baseline	1300 km	1300 km
Energy resolution	$0.1E_\nu$	$0.05E$ for QE events $(0.1-0.2)E$ for non-QE events
Cross-sections	From Refs. [143, 144]	From Refs. [143, 144]
Efficiency for $\mu^\pm$ (dis)appearance	94%	80%
Background on $\mu^\pm$ (dis)appearance	1×10^{-3}	$(1-5) \times 10^{-3}$
Efficiency for $e^\pm$ appearance	0 (Scenario 1) 47% (Scenario 2)	80%
Background on $e^\pm$ appearance	0.01	$0.01-0.8$
Systematics on signal	2%	2%
Systematics on background	2%	2%

Table I.2.: Description of the low-energy neutrino factory setups used in Chapters 6 and 7.

List of figures

List of tables

Bibliography

[1] W. Pauli. Letter to L. Meitner and colleagues, 4th December 1930.

[2] E. Fermi. An attempt of a theory of beta radiation. 1. *Z. Phys.*, 88:161–177, 1934.

[3] E. Fermi. Trends to a Theory of beta Radiation. (In Italian). *Nuovo Cim.*, 11:1–19, 1934.

[4] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the free neutrino: A Confirmation. *Science*, 124:103–104, 1956.

[5] G. Danby et al. Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos. *Phys. Rev. Lett.*, 9:36–44, 1962.

[6] Martin L. Perl et al. Evidence for anomalous lepton production in e+ e- annihilation. *Phys. Rev. Lett.*, 35:1489–1492, 1975.

[7] M. Acciarri et al. Determination of the number of light neutrino species from single photon production at LEP. *Phys. Lett.*, B431:199–208, 1998.

[8] K. Kodama et al. Detection and analysis of tau neutrino interactions in DONUT emulsion target. *Nucl. Instrum. Meth.*, A493:45–66, 2002.

[9] S. L. Glashow, J. Iliopoulos, and L. Maiani. Weak Interactions with Lepton-Hadron Symmetry. *Phys. Rev.*, D2:1285–1292, 1970.

[10] Jeffrey Goldstone, Abdus Salam, and Steven Weinberg. Broken Symmetries. *Phys. Rev.*, 127:965–970, 1962.

[11] Abdus Salam and John Clive Ward. Electromagnetic and weak interactions. *Phys. Lett.*, 13:168–171, 1964.

[12] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105:1413–1414, 1957.

[13] M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of neutrinos. *Phys. Rev.*, 109:1015–1017, 1958.

[14] K. Nakamura et al. Review of particle physics. *J. Phys.*, G37:075021, 2010.

[15] Ettore Majorana. Theory of the Symmetry of Electrons and Positrons. *Nuovo Cim.*, 14:171–184, 1937.

[16] S. M. Bilenky. Neutrinoless Double Beta-Decay. 2010.

[17] K. Zuber. COBRA: Double beta decay searches using CdTe detectors. *Phys. Lett.*, B519:1–7, 2001.

[18] R. Ardito et al. The CUORICINO and CUORE double beta decay experiments. *Prog. Part. Nucl. Phys.*, 57:203–216, 2006.

[19] Razvan Gornea. Double beta decay in liquid xenon. *J. Phys. Conf. Ser.*, 179:012004, 2009.

[20] Stefan. Schonert et al. The GERmanium Detector Array (GERDA) for the search of neutrinoless beta beta decays of Ge-76 at LNGS. *Nucl. Phys. Proc. Suppl.*, 145:242–245, 2005.

[21] Richard Gaitskell et al. White paper on the Majorana zero-neutrino double-beta decay experiment. 2003.

[22] H. Ejiri. MOON for neutrino-less beta beta decays and beta beta nuclear matrix elements. *AIP Conf. Proc.*, 1180:35–39, 2009.

[23] L. Simard. The NEMO-3 experiment and the SuperNEMO project. *Prog. Part. Nucl. Phys.*, 64:270–272, 2010.

[24] Christine Kraus and Simon J. M. Peeters. The rich neutrino programme of the SNO+ experiment. *Prog. Part. Nucl. Phys.*, 64:273–277, 2010.

[25] A. D. Sakharov. Violation of CP Invariance, c Asymmetry, and Baryon Asymmetry of the Universe. *Pisma Zh. Eksp. Teor. Fiz.*, 5:32–35, 1967.

[26] M. A. Luty. Baryogenesis via leptogenesis. *Phys. Rev.*, D45:455–465, 1992.

[27] Samoil M. Bilenky and S. T. Petcov. Massive Neutrinos and Neutrino Oscillations. *Rev. Mod. Phys.*, 59:671, 1987.

[28] Yoichiro Nambu and G. Jona-Lasinio. Dynamical model of elementary particles

based on an analogy with superconductivity. I. *Phys. Rev.*, 122:345–358, 1961.

[29] Yoichiro Nambu and G. Jona-Lasinio. Dynamical model of elementary particles based on an analogy with superconductivity. II. *Phys. Rev.*, 124:246–254, 1961.

[30] J. Goldstone. Field Theories with Superconductor Solutions. *Nuovo Cim.*, 19:154–164, 1961.

[31] Peter W. Higgs. Broken symmetries, massless particles and gauge fields. *Phys. Lett.*, 12:132–133, 1964.

[32] Peter W. Higgs. Spontaneous Symmetry Breakdown without Massless Bosons. *Phys. Rev.*, 145:1156–1163, 1966.

[33] Peter W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.*, 13:508–509, 1964.

[34] Ta-Pei Cheng and Ling-Fong Li. Gauge theory of elementary particle physics, published by Oxford University Press.

[35] Ch. Kraus et al. Final Results from phase II of the Mainz Neutrino Mass Search in Tritium β Decay. *Eur. Phys. J.*, C40:447–468, 2005.

[36] M. Gell-Mann, P. Ramond, and R. Slansky. *Supergravity, ed. by D. Freedman et al., North Holland*, 1979.

[37] Tsutomu Yanagida. Horizontal gauge symmetry and masses of neutrinos. *Prog. Theor. Phys.*, 64:1103, 1980.

[38] S. F. King. Neutrino mass models. *Rept. Prog. Phys.*, 67:107–158, 2004.

[39] A. Osipowicz et al. KATRIN: A next generation tritium beta decay experiment with sub-eV sensitivity for the electron neutrino mass. 2001.

[40] A. Monfardini et al. The microcalorimeter arrays for a rhenium experiment (MARE): A next-generation calorimetric neutrino mass experiment. *Prog. Part. Nucl. Phys.*, 57:68–70, 2006.

[41] V. M. Lobashev et al. Direct search for neutrino mass and anomaly in the tritium beta-spectrum: Status of 'Troitsk neutrino mass' experiment. *Nucl. Phys. Proc. Suppl.*, 91:280–286, 2001.

[42] B. Pontecorvo. Inverse processes and nonconservation of lepton charge. *JINR*

Preprint, Dubna, page 95, 1957.

[43] Raymond Davis, Jr., Don S. Harmer, and Kenneth C. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205–1209, 1968.

[44] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.

[45] Q. R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002.

[46] Nicola Cabibbo. Unitary Symmetry and Leptonic Decays. *Phys. Rev. Lett.*, 10:531–533, 1963.

[47] Makoto Kobayashi and Toshihide Maskawa. CP Violation in the Renormalizable Theory of Weak Interaction. *Prog. Theor. Phys.*, 49:652–657, 1973.

[48] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870–880, 1962.

[49] Evgeny Kh. Akhmedov and Alexei Yu. Smirnov. Paradoxes of neutrino oscillations. *Phys. Atom. Nucl.*, 72:1363–1381, 2009.

[50] K. Eguchi et al. First results from KamLAND: Evidence for reactor anti- neutrino disappearance. *Phys. Rev. Lett.*, 90:021802, 2003.

[51] D. G. Michael et al. Observation of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam. *Phys. Rev. Lett.*, 97:191801, 2006.

[52] Yasuo Takeuchi on behalf of the Super-Kamiokande collaboration. Plenary talk at Neutrino 2010 conference, to be published in the proceedings. 15th June 2010.

[53] M. C. Gonzalez-Garcia, Michele Maltoni, and Jordi Salvado. Updated global fit to three neutrino mixing: status of the hints of $\theta_{13} > 0$. *JHEP*, 04:056, 2010.

[54] P. F. Harrison, D. H. Perkins, and W. G. Scott. Tri-bimaximal mixing and the neutrino oscillation data. *Phys. Lett.*, B530:167, 2002.

[55] Mauro Mezzetto and Thomas Schwetz. θ_{13}: phenomenology, present status and prospect. 2010.

[56] S. Abe et al. Precision Measurement of Neutrino Oscillation Parameters with

KamLAND. *Phys. Rev. Lett.*, 100:221803, 2008.

[57] P. Adamson et al. Search for muon-neutrino to electron-neutrino transitions in MINOS. *Phys. Rev. Lett.*, 103:261802, 2009.

[58] R. Wendell et al. Atmospheric neutrino oscillation analysis with sub-leading effects in Super-Kamiokande I, II, and III. *Phys. Rev.*, D81:092004, 2010.

[59] S. Pascoli, S. T. Petcov, and Antonio Riotto. Connecting Low Energy Leptonic CP-violation to Leptogenesis. *Phys. Rev.*, D75:083511, 2007.

[60] S. Pascoli, S. T. Petcov, and Antonio Riotto. Leptogenesis and low energy CP violation in neutrino physics. *Nucl. Phys.*, B774:1–52, 2007.

[61] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev.*, D17:2369–2374, 1978.

[62] A. Cervera et al. Golden measurements at a neutrino factory. *Nucl. Phys.*, B579:17–55, 2000.

[63] S. P. Mikheyev and A. Yu. Smirnov. Resonant neutrino oscillations in matter. *Prog. Part. Nucl. Phys.*, 23:41–136, 1989.

[64] Osamu Yasuda. Three flavor neutrino oscillation analysis of the Superkamiokande atmospheric neutrino data. 1998.

[65] M. C. Gonzalez-Garcia, Y. Grossman, A. Gusso, and Y. Nir. New CP violation in neutrino oscillations. *Phys. Rev.*, D64:096006, 2001.

[66] Toshihiko Ota, Joe Sato, and Nao-aki Yamashita. Oscillation enhanced search for new interaction with neutrinos. *Phys. Rev.*, D65:093015, 2002.

[67] Joachim Kopp, Manfred Lindner, and Toshihiko Ota. Discovery reach for non-standard interactions in a neutrino factory. *Phys. Rev.*, D76:013001, 2007.

[68] Joachim Kopp, Manfred Lindner, Toshihiko Ota, and Joe Sato. Non-standard neutrino interactions in reactor and superbeam experiments. *Phys. Rev.*, D77:013007, 2008.

[69] Tommy Ohlsson and He Zhang. Non-Standard Interaction Effects at Reactor Neutrino Experiments. *Phys. Lett.*, B671:99–104, 2009.

[70] Stefan Antusch, Mattias Blennow, Enrique Fernandez-Martinez, and Toshihiko Ota. New physics searches at near detectors of neutrino oscillation experiments.

JHEP, 06:068, 2010.

[71] Andrea Donini, Ken-ichi Fuki, J. Lopez-Pavon, Davide Meloni, and Osamu Yasuda. The discovery channel at the Neutrino Factory: $\nu_\mu \to \nu_\tau$ pointing to sterile neutrinos. *JHEP*, 08:041, 2009.

[72] Elisabetta Baracchini. First Result From The MEG Experiment. 2010.

[73] Yoshitaka Kuno. Lepton flavor violation: Muon to electron conversion, COMET and PRISM/PRIME at J-PARC. *PoS*, NUFACT08:111, 2008.

[74] Stefan Antusch, Jochen P. Baumann, and Enrique Fernandez-Martinez. Non-Standard Neutrino Interactions with Matter from Physics Beyond the Standard Model. *Nucl. Phys.*, B810:369–388, 2009.

[75] Patricia Vahle on behalf of the MINOS collaboration. Plenary talk at Neutrino 2010 conference, to be published in the proceedings. 14th June 2010.

[76] W. Anthony Mann, Daniel Cherdack, Wojciech Musial, and Tomas Kafka. Apparent multiple Δm^2_{32} in muon anti-neutrino and muon neutrino survival oscillations from non-standard interaction matter effect. 2010.

[77] M. B. Gavela, D. Hernandez, T. Ota, and W. Winter. Large gauge invariant non-standard neutrino interactions. *Phys. Rev.*, D79:013007, 2009.

[78] Carla Biggio, Mattias Blennow, and Enrique Fernandez-Martinez. Loop bounds on non-standard neutrino interactions. *JHEP*, 03:139, 2009.

[79] Carla Biggio, Mattias Blennow, and Enrique Fernandez-Martinez. General bounds on non-standard neutrino interactions. *JHEP*, 08:090, 2009.

[80] John N. Bahcall and Carlos Pena-Garay. Solar models and solar neutrino oscillations. *New J. Phys.*, 6:63, 2004.

[81] W. W. M. Allison et al. The atmospheric neutrino flavor ratio from a 3.9 fiducial kiloton-year exposure of Soudan 2. *Phys. Lett.*, B449:137–144, 1999.

[82] A. A. Aguilar-Arevalo et al. The MiniBooNE Detector. *Nucl. Instrum. Meth.*, A599:28–46, 2009.

[83] J. Ahrens et al. Sensitivity of the IceCube detector to astrophysical sources of high energy muon neutrinos. *Astropart. Phys.*, 20:507–532, 2004.

[84] T. Marrodan Undagoitia et al. LENA: A multipurpose detector for low energy neutrino astronomy and proton decay. *J. Phys. Conf. Ser.*, 120:052018, 2008.

[85] I. Ambats et al. The MINOS Detectors Technical Design Report. NUMI-L-337.

[86] A. Bandyopadhyay et al. Physics at a future Neutrino Factory and super-beam facility. *Rept. Prog. Phys.*, 72:106201, 2009.

[87] D. Drakoulakos et al. Proposal to perform a high-statistics neutrino scattering experiment using a fine-grained detector in the NuMI beam. 2004.

[88] D. S. Ayres et al. NOvA proposal to build a 30-kiloton off-axis detector to study neutrino oscillations in the Fermilab NuMI beamline. 2004.

[89] C. Rubbia. The Liquid Argon Time Projection Chamber: A New Concept for Neutrino Detectors. CERN-EP-INT-77-08.

[90] John N. Bahcall, M. Baldo-Ceolin, D. B. Cline, and C. Rubbia. Predictions for a liquid argon solar neutrino detector. *Phys. Lett.*, B178:324, 1986.

[91] M. Soderberg. ArgoNeuT: A Liquid Argon Time Projection Chamber Test in the NuMI Beamline. 2009.

[92] Mitch Soderberg. Plenary talk at Neutrino 2010 conference, to be published in the proceedings. 15th June 2010.

[93] Takuya Hasegawa. Plenary talk at Neutrino 2010 conference, to be published in the proceedings. 15th June 2010.

[94] N. Agafonova et al. Observation of a first ν_τ candidate in the OPERA experiment in the CNGS beam. 2010.

[95] E. Eskut et al. The CHORUS experiment to search for $\nu_\mu \to \nu_\tau$ oscillation. *Nucl. Instrum. Meth.*, A401:7–44, 1997.

[96] M. Altmann et al. GNO solar neutrino observations: Results for GNO I. *Phys. Lett.*, B490:16–26, 2000.

[97] J. N. Abdurashitov et al. Measurement of the solar neutrino capture rate with gallium metal. *Phys. Rev.*, C60:055801, 1999.

[98] Y. Fukuda et al. Atmospheric muon-neutrino / electron-neutrino ratio in the multiGeV energy range. *Phys. Lett.*, B335:237–245, 1994.

[99] C. Athanassopoulos et al. Evidence for $\bar{\nu}_\mu \to \bar{\nu}_e$ oscillation from the LSND experiment at the Los Alamos Meson Physics Facility. *Phys. Rev. Lett.*, 77:3082–3085, 1996.

[100] C. Athanassopoulos et al. Evidence for $\nu_\mu \to \nu_e$ neutrino oscillations from LSND. *Phys. Rev. Lett.*, 81:1774–1777, 1998.

[101] A. A. Aguilar-Arevalo et al. A Search for electron neutrino appearance at the $\Delta m^2 \sim 1\mathrm{eV}^2$ scale. *Phys. Rev. Lett.*, 98:231801, 2007.

[102] A. A. Aguilar-Arevalo et al. Observed Event Excess in the MiniBooNE Search for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations. 2010.

[103] Y. Itow et al. The JHF-Kamioka neutrino project. 2001.

[104] F. Ardellier et al. Double Chooz: A search for the neutrino mixing angle θ_{13}. 2006.

[105] Herbert Steiner. The Daya Bay Experiment to measure θ_{13}. *Prog. Part. Nucl. Phys.*, 64:342–345, 2010.

[106] J. K. Ahn et al. RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang. 2010.

[107] V. Barger et al. Report of the US long baseline neutrino experiment study. 2007.

[108] D. Angus et al. The LAGUNA design study - towards giant liquid based underground detectors for neutrino physics and astrophysics and proton decay searches. 2010.

[109] Marco Zito. The T2K neutrino oscillation experiment. *J. Phys. Conf. Ser.*, 110:082023, 2008.

[110] P. Zucchelli. A novel concept for a anti-ν_e / ν_e neutrino factory: The beta beam. *Phys. Lett.*, B532:166–172, 2002.

[111] J. Burguet-Castell, D. Casper, E. Couce, J. J. Gomez-Cadenas, and P. Hernandez. Optimal beta-beam at the CERN-SPS. *Nucl. Phys.*, B725:306–326, 2005.

[112] J. Burguet-Castell, D. Casper, J. J. Gomez-Cadenas, P. Hernandez, and F. Sanchez. Neutrino oscillation physics with a higher gamma beta- beam. *Nucl. Phys.*, B695:217–240, 2004.

[113] Sandhya Choubey, Pilar Coloma, Andrea Donini, and Enrique Fernandez-

Martinez. Optimized Two-Baseline Beta-Beam Experiment. *JHEP*, 12:020, 2009.

[114] Jean-Eric Campagne, M. Maltoni, M. Mezzetto, and T. Schwetz. Physics potential of the CERN-MEMPHYS neutrino oscillation project. *JHEP*, 04:003, 2007.

[115] Davide Meloni, Olga Mena, Christopher Orme, Sergio Palomares-Ruiz, and Silvia Pascoli. An intermediate gamma beta-beam neutrino experiment with long baseline. *JHEP*, 07:115, 2008.

[116] Christopher Orme. High-γ Beta Beams within the LAGUNA design study. 2010.

[117] S. Geer. Neutrino beams from muon storage rings: Characteristics and physics potential. *Phys. Rev.*, D57:6989–6997, 1998.

[118] A. Cervera, A. Laing, J. Martin-Albo, and F. J. P. Soler. Performance of the MIND detector at a Neutrino Factory using realistic muon reconstruction. 2010.

[119] A. De Rujula, M. B. Gavela, and P. Hernandez. Neutrino oscillation physics with a neutrino factory. *Nucl. Phys.*, B547:21–38, 1999.

[120] Carl H. Albright et al. Physics at a neutrino factory. 2000.

[121] Vernon D. Barger, S. Geer, and K. Whisnant. Long baseline neutrino physics with a muon storage ring neutrino source. *Phys. Rev.*, D61:053004, 2000.

[122] A. Donini, M. B. Gavela, P. Hernandez, and S. Rigolin. Neutrino mixing and CP-violation. *Nucl. Phys.*, B574:23–42, 2000.

[123] Vernon D. Barger, S. Geer, R. Raja, and K. Whisnant. Neutrino oscillations at an entry-level neutrino factory and beyond. *Phys. Rev.*, D62:073002, 2000.

[124] Vernon D. Barger, S. Geer, R. Raja, and K. Whisnant. Long-baseline study of the leading neutrino oscillation at a neutrino factory. *Phys. Rev.*, D62:013004, 2000.

[125] M. Freund, Patrick Huber, and M. Lindner. Extracting matter effects, masses and mixings at a neutrino factory. *Nucl. Phys.*, B585:105–123, 2000.

[126] Vernon D. Barger, S. Geer, R. Raja, and K. Whisnant. Determination of the pattern of neutrino masses at a neutrino factory. *Phys. Lett.*, B485:379–387, 2000.

[127] J. Burguet-Castell, M. B. Gavela, J. J. Gomez-Cadenas, P. Hernandez, and Olga Mena. On the measurement of leptonic CP violation. *Nucl. Phys.*, B608:301–318, 2001.

[128] M. Freund, Patrick Huber, and M. Lindner. Systematic exploration of the neutrino factory parameter space including errors and correlations. *Nucl. Phys.*, B615:331–357, 2001.

[129] A. Donini, D. Meloni, and P. Migliozzi. The silver channel at the Neutrino Factory. *Nucl. Phys.*, B646:321–349, 2002.

[130] D. Autiero et al. The synergy of the golden and silver channels at the Neutrino Factory. *Eur. Phys. J.*, C33:243–260, 2004.

[131] A. Blondel et al. The neutrino factory: Beam and experiments. *Nucl. Instrum. Meth.*, A451:102–122, 2000.

[132] M. Apollonio et al. Oscillation physics with a neutrino factory. ((G)) ((U)). 2002.

[133] Carl H. Albright et al. The neutrino factory and beta beam experiments and development. 2004.

[134] Olga Mena. Unveiling Neutrino Mixing and Leptonic CP Violation. *Mod. Phys. Lett.*, A20:1–17, 2005.

[135] P. Huber, M. Lindner, M. Rolinec, and W. Winter. Optimization of a neutrino factory oscillation experiment. *Phys. Rev.*, D74:073003, 2006.

[136] Steve Geer, Olga Mena, and Silvia Pascoli. A Low energy neutrino factory for large θ_{13}. *Phys. Rev.*, D75:093001, 2007.

[137] Alan D. Bross, Malcolm Ellis, Steve Geer, Olga Mena, and Silvia Pascoli. A Neutrino factory for both large and small θ_{13}. *Phys. Rev.*, D77:093012, 2008.

[138] Patrick Huber and Walter Winter. Neutrino Factory Superbeam. *Phys. Lett.*, B655:251–256, 2007.

[139] Joachim Kopp, Toshihiko Ota, and Walter Winter. Neutrino factory optimization for non-standard interactions. *Phys. Rev.*, D78:053007, 2008.

[140] N. C. Ribeiro, H. Minakata, H. Nunokawa, S. Uchinami, and R. Zukanovich-Funchal. Probing Non-Standard Neutrino Interactions with Neutrino Factories. *JHEP*, 12:002, 2007.

[141] Patrick Huber, M. Lindner, and W. Winter. Simulation of long-baseline neutrino oscillation experiments with GLoBES. *Comput. Phys. Commun.*, 167:195, 2005.

[142] Patrick Huber, Joachim Kopp, Manfred Lindner, Mark Rolinec, and Walter Winter. New features in the simulation of neutrino oscillation experiments with GLoBES 3.0. *Comput. Phys. Commun.*, 177:432–438, 2007.

[143] Mark D. Messier. Evidence for neutrino mass from observations of atmospheric neutrinos with super-kamiokande. UMI-99-23965.

[144] E. A. Paschos and J. Y. Yu. Neutrino interactions in oscillation experiments. *Phys. Rev.*, D65:033002, 2002.

[145] Dziewonski. A. M. and D. L. Anderson. Preliminary reference earth model. *Phys. Earth Planet Interiors*, 25:297–356, 1981.

[146] F. D. Stacey. Physics of the earth. *2nd edition*, Wiley, 1977.

[147] C. Jarlskog. Commutator of the Quark Mass Matrices in the Standard Electroweak Model and a Measure of Maximal CP Violation. *Phys. Rev. Lett.*, 55:1039, 1985.

[148] Gian Luigi Fogli and E. Lisi. Tests of three-flavor mixing in long-baseline neutrino oscillation experiments. *Phys. Rev.*, D54:3667–3670, 1996.

[149] Hisakazu Minakata and Hiroshi Nunokawa. Exploring neutrino mixing with low energy superbeams. *JHEP*, 10:001, 2001.

[150] Vernon D. Barger, S. Geer, R. Raja, and K. Whisnant. Exploring neutrino oscillations with superbeams. *Phys. Rev.*, D63:113011, 2001.

[151] Takaaki Kajita, Hisakazu Minakata, and Hiroshi Nunokawa. Method for determination of $|U(e3)|$ in neutrino oscillation appearance experiments. *Phys. Lett.*, B528:245–252, 2002.

[152] Hisakazu Minakata, Hiroshi Nunokawa, and Stephen J. Parke. Parameter degeneracies in neutrino oscillation measurement of leptonic CP and T violation. *Phys. Rev.*, D66:093012, 2002.

[153] Patrick Huber, Manfred Lindner, and Walter Winter. Superbeams versus neutrino factories. *Nucl. Phys.*, B645:3–48, 2002.

[154] A. Donini, D. Meloni, and S. Rigolin. Clone flow analysis for a theory inspired neutrino experiment planning. *JHEP*, 06:011, 2004.

[155] Mayumi Aoki, Kaoru Hagiwara, and Naotoshi Okamura. Lifting degeneracies in the oscillation parameters by a neutrino factory. *Phys. Lett.*, B606:371–383, 2005.

[156] Osamu Yasuda. New plots and parameter degeneracies in neutrino oscillations. *New J. Phys.*, 6:83, 2004.

[157] Olga Mena and Stephen J. Parke. Physics potential of the Fermilab NuMI beamline. *Phys. Rev.*, D72:053003, 2005.

[158] Hisakazu Minakata and Hiroshi Nunokawa. How to measure CP violation in neutrino oscillation experiments? *Phys. Lett.*, B413:369–377, 1997.

[159] V. Barger, D. Marfatia, and K. Whisnant. Off–axis Beams and Detector Clusters: Resolving Neutrino Parameter Degeneracies. *Phys. Rev.*, D66:053007, 2002.

[160] Olga Mena Requejo, Sergio Palomares-Ruiz, and Silvia Pascoli. Super-NOvA: A long-baseline neutrino experiment with two off-axis detectors. *Phys. Rev.*, D72:053002, 2005.

[161] Masaki Ishitsuka, Takaaki Kajita, Hisakazu Minakata, and Hiroshi Nunokawa. Resolving Neutrino Mass Hierarchy and CP Degeneracy by Two Identical Detectors with Different Baselines. *Phys. Rev.*, D72:033003, 2005.

[162] Kaoru Hagiwara, Naotoshi Okamura, and Ken-ichi Senda. Solving the neutrino parameter degeneracy by measuring the T2K off-axis beam in Korea. *Phys. Lett.*, B637:266–273, 2006.

[163] Olga Mena, Sergio Palomares-Ruiz, and Silvia Pascoli. Determining the neutrino mass hierarchy and CP violation in NOnuA with a second off-axis detector. *Phys. Rev.*, D73:073007, 2006.

[164] Takaaki Kajita, Hisakazu Minakata, Shoei Nakayama, and Hiroshi Nunokawa. Resolving eight-fold neutrino parameter degeneracy by two identical detectors with different baselines. *Phys. Rev.*, D75:013006, 2007.

[165] J. Burguet-Castell, M. B. Gavela, J. J. Gomez-Cadenas, P. Hernandez, and Olga Mena. Superbeams plus neutrino factory: The golden path to leptonic CP violation. *Nucl. Phys.*, B646:301–320, 2002.

[166] V. Barger, D. Marfatia, and K. Whisnant. How two neutrino superbeam experiments do better than one. *Phys. Lett.*, B560:75–86, 2003.

[167] Patrick Huber, M. Lindner, and W. Winter. Synergies between the first-generation JHF-SK and NuMI superbeam experiments. *Nucl. Phys.*, B654:3–29, 2003.

[168] Hisakazu Minakata, Hiroshi Nunokawa, and Stephen J. Parke. The complementarity of eastern and western hemisphere long-baseline neutrino oscillation experiments. *Phys. Rev.*, D68:013010, 2003.

[169] K. Whisnant, Jin Min Yang, and Bing-Lin Young. Measuring CP violation and mass ordering in joint long baseline experiments with superbeams. *Phys. Rev.*, D67:013004, 2003.

[170] Patrick Huber, M. Lindner, T. Schwetz, and W. Winter. Reactor Neutrino Experiments Compared to Superbeams. *Nucl. Phys.*, B665:487–519, 2003.

[171] Patrick Huber, M. Lindner, M. Rolinec, T. Schwetz, and W. Winter. Prospects of accelerator and reactor neutrino oscillation experiments for the coming ten years. *Phys. Rev.*, D70:073014, 2004.

[172] A. Donini, E. Fernandez-Martinez, and S. Rigolin. Appearance and disappearance signals at a beta-beam and a super-beam facility. *Phys. Lett.*, B621:276–287, 2005.

[173] Olga Mena and Stephen J. Parke. Untangling CP violation and the mass hierarchy in long baseline experiments. *Phys. Rev.*, D70:093011, 2004.

[174] Patrick Huber, Michele Maltoni, and Thomas Schwetz. Resolving parameter degeneracies in long-baseline experiments by atmospheric neutrino data. *Phys. Rev.*, D71:053006, 2005.

[175] Sandhya Choubey and Probir Roy. Probing the deviation from maximal mixing of atmospheric neutrinos. *Phys. Rev.*, D73:013006, 2006.

[176] A. Blondel et al. Future neutrino oscillation facilities. *Acta Phys. Polon.*, B37:2077–2113, 2006.

[177] Olga Mena, Hiroshi Nunokawa, and Stephen J. Parke. NOvA and T2K: The race for the neutrino mass hierarchy. *Phys. Rev.*, D75:033002, 2007.

[178] Andreas Jansson, Olga Mena, Stephen J. Parke, and Niki Saoulidou. Combining CPT-conjugate Neutrino channels at Fermilab. *Phys. Rev.*, D78:053002, 2008.

[179] Patrick Huber and Walter Winter. Neutrino factories and the 'magic' baseline. *Phys. Rev.*, D68:037301, 2003.

[180] D. Autiero et al. Large underground, liquid based detectors for astro- particle physics in Europe: scientific case and prospects. *JCAP*, 0711:011, 2007.

[181] A. Rubbia. A CERN-based high-intensity high-energy proton source for long baseline neutrino oscillation experiments with next- generation large underground detectors for proton decay searches and neutrino physics and astrophysics. 2010.

[182] A. de Bellefon et al. MEMPHYS: A large scale water Cerenkov detector at Frejus. 2006.

[183] A. Rubbia. Underground Neutrino Detectors for Particle and Astroparticle Science: the Giant Liquid Argon Charge Imaging ExpeRiment (GLACIER). *J. Phys. Conf. Ser.*, 171:012020, 2009.

[184] Juha Peltoniemi. Simulations of neutrino oscillations with a high-energy beta beam from CERN to LENA at Pyhasalmi Mine. 2009.

[185] Juha Peltoniemi. Simulations of neutrino oscillations for a low-energy neutrino factory with a magnetized large-volume liquid scintillator at 2288 km baseline. 2009.

[186] Juha Peltoniemi. Simulations of neutrino oscillations for a wide band beam from CERN to LENA. 2009.

[187] Juha Peltoniemi. Accuracy and energy resolution of a large volume liquid scintillator experiment at high energies. *internal note*, 2010.

[188] Juha Peltoniemi. Liquid scintillator as tracking detector for high-energy events. 2009.

[189] IDS-NF Steering Group. Neutrino Factory: specification of baseline for the accelerator complex and detector systems. *IDS-NF-Baseline-2007/1.0*, 2008.

[190] A. Longhin. http://irfu.cea.fr/en/Phocea/Pisp/index.php?id=72.

[191] L. Esposito and A. Rubbia. Private communication.

[192] V. Barger, P. Huber, D. Marfatia, and W. Winter. Which long-baseline neutrino experiments are preferable? *Phys. Rev.*, D76:053005, 2007.

[193] A. Meregaglia and A. Rubbia. Neutrino oscillation physics at an upgraded CNGS with large next generation liquid argon TPC detectors. *JHEP*, 11:032, 2006.

[194] E. Kozlovskaya, J. Peltoniemi, and J. Sarkamo. The density distribution in the Earth along the CERN- Pyhasalmi baseline and its effect on neutrino oscillations. 2003.

[195] S. Kasuga. Small ν_μ / ν_e Ratio of Atmospheric Neutrinos in Super-Kamiokande by the Method of Particle Identification. *PhD thesis, University of Tokyo*, 1998.

[196] Enrique Fernandez Martinez, Tracey Li, Silvia Pascoli, and Olga Mena. Improvement of the low energy neutrino factory. *Phys. Rev.*, D81:073010, 2010.

[197] S. Raby et al. DUSEL Theory White Paper. 2008.

[198] Ambrosio et. al. *Report No. Fermilab-TM-2018*, 1997.

[199] Patrick Huber and Thomas Schwetz. A low energy neutrino factory with non-magnetic detectors. *Phys. Lett.*, B669:294–300, 2008.

[200] Patrick Huber and Joachim Kopp. Two experiments for the price of one? – The role of the second oscillation maximum in long baseline neutrino experiments. 2010.

[201] A. Donini, E. Fernandez-Martinez, D. Meloni, and S. Rigolin. nu/mu disappearance at the SPL, T2K-I, NOnuA and the neutrino factory. *Nucl. Phys.*, B743:41–73, 2006.

[202] M. Diwan et al. Proposal for an experimental program in neutrino physics and proton decay in the homestake laboratory. 2006.

[203] V. Barger et al. Precision physics with a wide band super neutrino beam. *Phys. Rev.*, D74:073004, 2006.

[204] Patrick Huber, T. Schwetz, and J. W. F. Valle. Confusing non-standard neutrino interactions with oscillations at a neutrino factory. *Phys. Rev.*, D66:013006, 2002.

[205] Stefan Antusch, Mattias Blennow, Enrique Fernandez-Martinez, and Jacobo Lopez-Pavon. Probing non-unitary mixing and CP-violation at a Neutrino Factory. *Phys. Rev.*, D80:033002, 2009.

[206] Patrick Huber and J. W. F. Valle. Non-standard interactions: Atmospheric versus neutrino factory experiments. *Phys. Lett.*, B523:151–160, 2001.

[207] Davide Meloni, Tommy Ohlsson, Walter Winter, and He Zhang. Non-standard interactions versus non-unitary lepton flavor mixing at a neutrino factory. *JHEP*, 04:041, 2010.

[208] Mattias Blennow and Enrique Fernandez-Martinez. Neutrino oscillation parameter sampling with MonteCUBES. *Comput. Phys. Commun.*, 181:227–231, 2010.

[209] Jian Tang and Walter Winter. Neutrino factory in stages: Low energy, high energy, off- axis. *Phys. Rev.*, D81:033005, 2010.

[210] C. Giunti, C. W. Kim, and U. W. Lee. When do neutrinos really oscillate?: Quantum mechanics of neutrino oscillations. *Phys. Rev.*, D44:3635–3640, 1991.

[211] V. M. Visscher. Neutrino Detection by Resonance Absorption in Crystals at Low Temperatures. *Phys. Rev.*, 116:1581–1582, 1959.