Determining the Neutrino Mass Hierarchy with the Precision IceCube Next Generation Upgrade (PINGU)

Dissertation
zur
Erlangung des Doktorgrades (Dr. rer. nat.)
der
Mathematisch-Naturwissenschaftlichen Fakultät
der
Rheinischen Friedrich-Wilhelms-Universität Bonn

von
Lukas Schulte
aus
Mainz

Bonn, April 2015

Dieser Forschungsbericht wurde als Dissertation von der Mathematisch-Naturwissenschaftlichen Fakultät der Universität Bonn angenommen und ist auf dem Hochschulschriftenserver der ULB Bonn `http://hss.ulb.uni-bonn.de/diss_online` elektronisch publiziert.

1. Gutachter:	Prof. Dr. Marek Kowalski
2. Gutachter:	Prof. Dr. Norbert Wermes
Tag der Promotion:	22.09.2015
Erscheinungsjahr:	2015

Abstract

In this thesis, the development of a fast effective simulation for the planned PINGU experiment at the geographic South Pole is described, which will make a precision measurement of the atmospheric neutrino flux at low GeV energies. In this flux, the effects of neutrino oscillations in the matter potential of the Earth are visible, which will be observed by PINGU with unprecedented precision.

Using the aforementioned simulation, PINGU's expected precision in determining the relevant neutrino oscillation parameters and the neutrino mass hierarchy is calculated, incorporating a variety of parameters covering systematic uncertainties in the experimental outcome. The analysis is done in the framework of the Fisher Matrix technique, whose application to a particle physics experiment is novel. It allows for a fast and stable evaluation of the multi-dimensional parameter space and an easy combination of different experiments.

Contents

CHAPTER 1

Introduction

Although the first conclusive observation of neutrino oscillations was made not even twenty years ago, this phenomenon of neutrinos changing their flavour when travelling macroscopic distances has been one of the major areas of research in particle physics and astrophysics ever since. Up to now, it is the only manifestation of so-called "physics beyond the standard model" that has been confirmed experimentally. During the past two decades, many dedicated experiments have mapped out the parameters characterising neutrino oscillations in great detail, leaving only two parameters to be determined.

One of these parameters is the so-called neutrino mass hierarchy. It refers to the sign of another parameter, one of the two independent mass splittings, whose absolute value has already been measured. The fact that the absolutes of parameters can be determined precisely without learning about its sign is one of the peculiarities of the neutrino oscillation formalism, where central parameters enter quadratically in most cases.

A chance to access the neutrino mass hierarchy is to study the differences in the oscillation probabilities of neutrinos and antineutrinos at low GeV energies that are created in the Earth's atmosphere and propagate through its interior. The proposed Precision IceCube Next Generation Upgrade (PINGU) will be a facility apt to observe the small modulations on top of the flux of atmospheric neutrinos with the required precision.

As its name suggests, PINGU is planned as an upgrade to the existing IceCube neutrino telescope at the geographic South Pole in Antarctica. IceCube has been constructed to discover extra-terrestrial neutrinos at TeV to PeV energies. Neutrino oscillation patterns in the atmospheric flux at medium GeV energies, however, have already been observed as well using its DeepCore extension. PINGU is now intended to further lower the energy threshold down to a regime where signatures of the neutrino mass hierarchy appear. This also provides an opportunity to measure the absolute values of the relevant oscillation parameters with high precision.

In this thesis, the development of an effective detector simulation for PINGU, named `PaPA` for "Parametrised PINGU Analysis", is described. The outcome of this simulation is analysed using the Fisher Matrix formalism, a tool that is well established in cosmology, but novel to be applied to a particle physics experiment. In a linear approximation, it allows for a fast construction of the full covariance matrix of the experiment including a large number of systematic uncertainties.

After checking that the prerequisites for the Fisher Matrix are in fact fulfilled, PINGU's expected sensitivity to the mass hierarchy is evaluated, showing its dependence on controlling the relevant systematics. The expected precision in measuring the accessible oscillation parameters is calculated as

well. Following this, the effect of changing various simulation input parameters is studied in terms of the resulting sensitivity to the neutrino mass hierarchy.

Finally, PaPA is modified to simulate the outcome of JUNO, a nuclear reactor based neutrino experiment targeted at the determination of the neutrino mass hierarchy as well, yet exploiting a very different physical effect. JUNO's sensitivity to the mass hierarchy will be calculated, where a Fourier transformation of the observed signal is needed in order to justify the application of the Fisher Matrix. Afterwards, the standalone results of PINGU and JUNO can be combined easily as in the Fisher Matrix formalism this corresponds to a mere matrix addition and inversion. Then the benefits from marginalising over common systematics can be investigated.

The thesis concludes with a summary of the results found using PaPA and analysing its outcome in terms of the Fisher Matrix. An outlook is given on the future of PaPA and its integration into a wider software framework for PINGU detector simulations.

CHAPTER 2

Neutrinos in the Standard Model

In this chapter, the theoretical background of this thesis will be discussed in detail. Here the focus will be on neutrinos and their properties and interactions, while neutrino oscillations will be described in the following chapter.

2.1 Standard Model in a Nutshell

In the Standard Model of Particle Physics, or just Standard Model, the current theories of the electroweak and strong interactions are combined [1–4], for an overview see e.g. [5]. It is a quantum field theory of the fundamental interactions and particles relevant on the scales that are accessible for particle physics experiments.

The particles, listed in Fig. 2.1, can be divided in two classes: fermions with an intrinsic spin of 1/2 make up everything that is usually called "matter", and exchange bosons with integer (one in most cases) spin that convey the interactions and couple to the respective charge. Formally, the bosons are the generators of the gauge symmetry group of the particular interaction.

This means that the strong force, which obeys a *SU(3)* symmetry, has eight generators that are represented by eight gluons g. The gluons couple to the strong charge which is usually referred to as "colour". Since it has the largest coupling constant, the strong interaction is dominant whenever a colour charge is present. However colour is "confined", i.e. free particles must not have a net colour. This means that any coloured particles have to be bound inside a compound object at all times. The range of the strong interaction is limited to about the size of a nucleus since the gluons are coloured themselves and hence self-coupling.

The the electromagnetic interaction is about two orders of magnitude weaker. According to its *U(1)* symmetry, it has only one exchange boson, the photon γ, coupling to the electrical charge. It is massless and electrically neutral, hence the electromagnetic interaction is not restricted in range. This and the fact that there is no confinement on the electrical charge mean that electromagnetic phenomena are dominant on macroscopic scales.

At low energies, the effective coupling constant of the weak interaction is another three orders below the electromagnetic one. From its *SU(2)* symmetry originate three exchange bosons, $W^{\pm}$ and Z^0 with masses of about 90 MeV limiting its range to the subatomic scale. However with increasing energy, the mass of the gauge bosons becomes more and more negligible and the effective coupling rises. Above the electroweak unification at about 100 GeV, the weak and electromagnetic interactions can be described

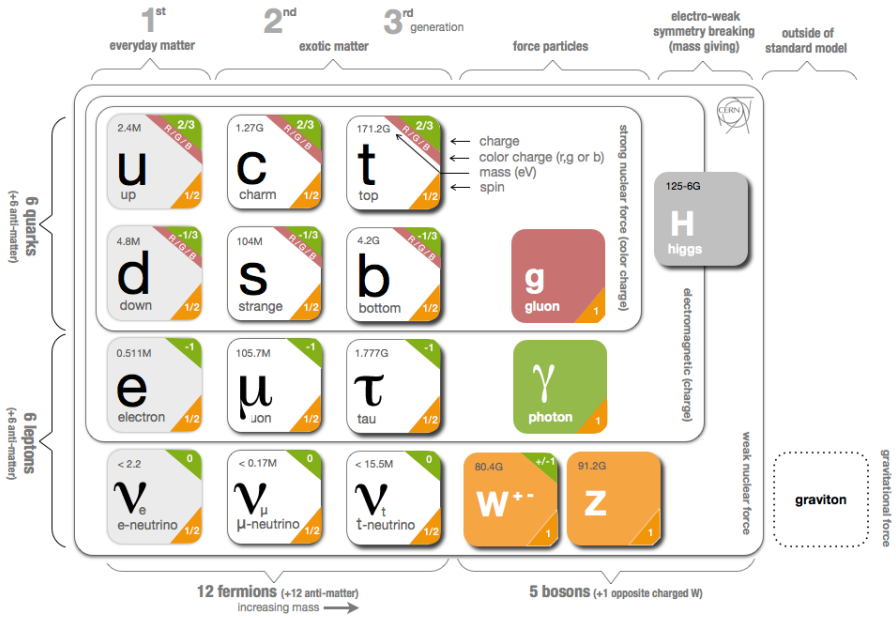

Figure 2.1: The fundamental particles in the Standard Model. Figure taken from [6].

by one unified theory, whose existence is also hinted at by the fact that the weak gauge bosons are electrically charged.

Their masses arise from another spontaneously broken local *SU(2)×U(1)* symmetry of the so-called Higgs[1] field. After breaking, the generators of the *SU(2)* part mix with the weak bosons, giving them mass, while the generator of the remaining *U(1)* can be observed as the only scalar gauge boson, the Higgs boson. The Higgs boson was the last fundamental particle of the standard model to be detected, its discovery was claimed by the ATLAS and CMS collaborations in 2012 [9, 10].

The other group of fundamental particles are the fermions (and their corresponding antiparticles). They can be divided again into two subclasses: the six quarks u, d, c, s, t, and b, which obey all forces and—being coloured—are confined, so that no free quarks can be found in nature. Bound quarks are making up baryons, like protons and neutrons, consisting of three quarks, and unstable mesons like pions, kaons, and many others, which consist of a quark and an antiquark. Baryons and mesons, together called hadrons, are the only free particles participating in the strong interaction, since they contain coloured quarks, although not being coloured themselves.

The second subclass are the leptons, the three charged leptons e, μ, and τ, as well as the corresponding (neutral) neutrinos ν_e, ν_μ, and ν_τ. The charged leptons interact predominantly electromagnetically, most prominently electrons are bound to nuclei via electrical attraction. However the decay of μ and τ is—like every flavour-changing process—a weak interaction. The electron as the lightest charged lepton has to be stable due to conservation of energy and charge.

Since neutrinos are neither coloured nor electrically charged, they only interact weakly. This means that they are very hard to detect directly. In fact, their existence had already been suggested in 1930 by Wolfgang Pauli as a solution for the problem of missing energy in radioactive β decays [11]. However

[1] After Peter Higgs, who, together with others, laid the foundations of this theory in the 1960's [7, 8].

the first direct detection of (electron) neutrinos, ν_e, from a nuclear reactor was achieved only in 1956 in the so-called Cowan-Reines experiment [12]. The existence of a second neutrino, the muon neutrino ν_μ, was established few years later in 1962 from the study of charged pion decays [13]. The third neutrino, the ν_τ, was finally discovered directly by the DONUT experiment in 2001 in the decay of D_S mesons into $\bar{\nu}_\tau$ and τ, which again decay into ν_τ and other leptons [14].

In addition to having neither colour nor electrical charge, the standard model also predicts that neutrinos are massless. Thus the observation of neutrino oscillations by the Super-Kamiokande experiment in 1998[2] [16] gained much attention, this being the first detection of physics beyond the standard model.

The term "neutrino oscillations" describes the phenomenon that neutrinos propagating over macroscopic distances can change their flavour eigenstate on the way between production and detection. The details of this effect will be described in Sec. 3. However it can only occur when there are different mass eigenstates available for the neutrinos, meaning that only one of them—if any—can have zero mass, while the others must correspond to finite mass.

Since their first observation, neutrino oscillations have been a field of intensive research. After establishing all oscillation channels, nowadays' focus is on the precise measurement of the parameters that characterise the oscillation. The planned PINGU experiment (see Sec. 4.2), whose simulation is the main topic of this thesis, is aimed to reach unprecedented accuracy in measuring the parameters ϑ_{23} and Δm^2_{31}.

2.2 Neutrino Sources

As mentioned above, neutrinos do not participate in the strong and electromagnetic interaction, leaving only weak processes for them to be created or detected. On the other hand, neutrinos are produced in nearly every weak interaction, making them a very common particle that can stem from a variety of different sources in very different energy ranges.

2.2.1 Natural Radioactivity

On Earth, the most common source is the β decay of natural radionuclides. Depending on the type of the decay (β^+ or β^-), an electron (anti-) neutrino is emitted along with the charged lepton. The general equations read:

$$\beta^+ : \quad {}^{A}_{Z}\mathrm{X} \quad \rightarrow \quad {}^{A}_{Z-1}\mathrm{Y} + e^+ + \nu_e \tag{2.1}$$

$$\beta^- : \quad {}^{A}_{Z}\mathrm{X} \quad \rightarrow \quad {}^{A}_{Z+1}\mathrm{Y} + e^- + \bar{\nu}_e \tag{2.2}$$

Examples for typical β emitters are ^{40}K (both β^+ and β^-) and intermediate products from the decay chains of ^{232}Th or ^{238}U (β^-), the neutrino energies are usually on the scale of few MeV. These neutrinos, originating from nuclear decays inside the Earth's crust and mantle, are commonly referred to as geoneutrinos.

In fact, the β decay was the original reason to propose the existence of the then undetectable neutrino. Since only the daughter nucleus and the charged lepton were visible as decay products, the process seemed to be a two-body decay. This means that the energies of the decay products are exactly determined from kinematics and hence the emitted electrons or positrons would be mono-energetic. Observations showed, however, a broad spectrum in energy instead of a single line. Without violating

[2] Hints of neutrino oscillations had already been observed in the 1960's [15], but were widely refused by the scientific community.

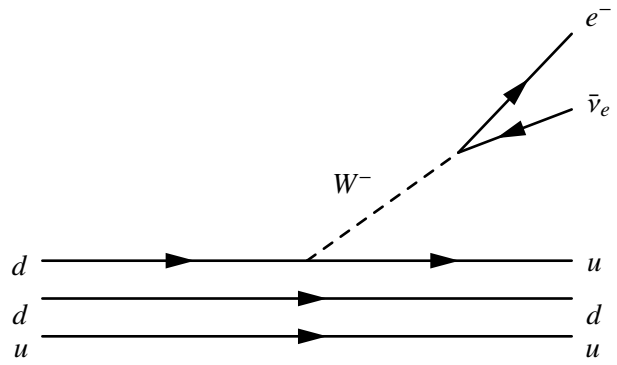

Figure 2.2: Feynman diagram of a β^- decay

the conservation of energy, this can only be achieved if a third particle is produced in the process, which has to be electrically neutral but can carry away energy and momentum.

On the subatomic level, in a β^+ decay one proton inside a nucleus emits a virtual W^+ boson, turning an u into a d quark, and becomes a neutron. During a β^- decay, the opposite happens via the emission of a W^-, as shown in Fig. 2.2. The W boson subsequently decays into an electron and a $\bar{\nu}_e$ or positron and ν_e, respectively[3].

2.2.2 Nuclear Reactors

In nuclear reactors, the controlled fission of heavy elements is used for the generation of electrical power. The intermediate products of these nuclear fissions are unstable isotopes that usually have a large surplus of neutrons compared to a stable configuration. These unstable nuclides undergo a series of β^- decays until they they reach a stable ratio of proton and neutron numbers.

Since in each of those β^- decays a neutrino is emitted, a nuclear reactor provides a strong and steady flux of $\bar{\nu}_e$ in the low MeV range, which can be monitored via the thermal power of the reactor. This makes reactor neutrinos a popular target for experiments, especially for the study of neutrino oscillations. The main challenge in such experiments is the accurate modelling of the neutrino energy spectrum, which is the sum of the spectra of all of the different β decays in the decay chain. Even though there are very elaborate flux models available, there might still be components unaccounted for, resulting in unexpected features in the measured neutrino flux [17].

2.2.3 Neutrino Beams

Another artificial source of neutrinos, but somewhat higher in energy (typically at a few GeV), are neutrino beams. Since the neutral neutrinos cannot be accelerated directly, usually a high-energy, high-intensity proton beam is aimed at a target in which it produces mesons, mostly pions and kaons, which subsequently decay under the emission of neutrinos [18]:

$$\pi^\pm \rightarrow \mu^\pm + \overset{(-)}{\nu}_\mu \tag{2.3}$$

$$K^\pm \rightarrow \mu^\pm + \overset{(-)}{\nu}_\mu \tag{2.4}$$

$$\tag{2.5}$$

[3] In principle, a W boson can decay into any charged lepton and corresponding neutrino. Yet in natural radioactivity the typical decay energies are few MeV, hence the W boson is only virtual and since the overall energy has to be conserved, muons or tau leptons are too massive to be produced.

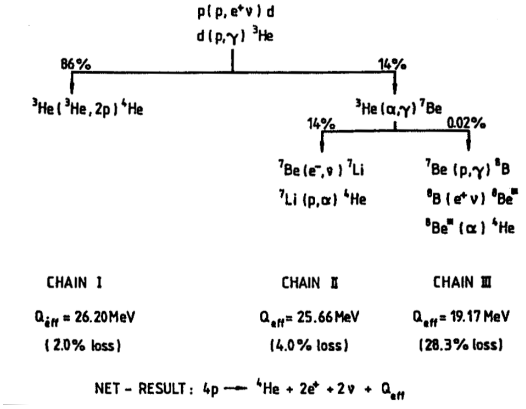

(a) The different branches of the pp chain. "Lost" energy is not dissipated into the Sun, but carried away by neutrinos.

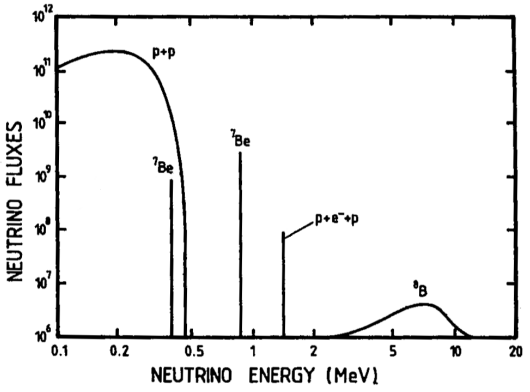

(b) The solar neutrino spectrum from the pp chain. The neutrino fluxes are in units of $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}\,\mathrm{MeV}^{-1}$ for continuous and $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ for discrete components.

Figure 2.3: The reactions and resulting neutrino spectrum of the solar pp chain. Figures adopted from [19].

Such beams are the source of neutrinos that can be controlled best in terms of energy and intensity, making them a preferred choice for precision experiments such as measurements of neutrino cross-sections. However they are very expensive to build and operate in contrast to natural sources or nuclear reactors, the latter usually being operated by commercial power suppliers and thus provide the neutrino flux "for free".

2.2.4 Solar Neutrinos

In terms of total flux, the strongest source of neutrinos on Earth is the Sun. In its interior, hydrogen is fused to helium mostly in the so-called pp chain[4], producing the energy that powers the Sun's radiation[19].

Effectively, this is carried out via the reaction

$$4p \rightarrow {}^{4}_{2}\mathrm{He} + 2e^{+} + 2\nu_e + 26.73\,\mathrm{MeV} \quad . \tag{2.6}$$

In reality, this will not occur in a single step since it is a weak interaction (as neutrinos are produced in its course) with a correspondingly small cross-section which in addition has to overcome the Coulomb repulsion of the four protons. Instead the fusion process involves several intermediate stages, the first of which is the fusion of two protons to a deuteron,

$$p + p \rightarrow\ d + e^{+} + \nu_e \quad , \tag{2.7}$$

releasing 0.42 MeV of energy carried by the (subsequently annihilating) positron and the so-called pp neutrino. The low Q value along with the aforementioned low cross-section and Coulomb repulsion are the reason for the long lifetime of free protons inside the Sun[5]—it takes them an average of about 10^{10} years to fuse to deuterium. A competing, but even more improbable reaction is the pep process, where

[4] Other fusion processes such as the CNO cycle and the production of heavier elements are strongly suppressed since they need extremely high pressures that the Sun cannot supply due to its comparatively low mass.

[5] And hence also the long lifetime of the Sun itself

an electron is involved directly in the fusion and no positron is produced:

$$p + e^- + p \rightarrow \ d + \nu_e \tag{2.8}$$

Since there are only two particles in the final state, the produced neutrinos are monoenergetic at 1.44 MeV. Once a deuteron is produced, it quickly ($\approx$ 1 s) merges with another proton to $^3_2$He and emits a photon.

At this stage, the pp chain divides into three different branches (see Fig. 2.3a). In the main branch, ppI, two $^3_2$He nuclei fuse to $^4_2$He, also called an α particle, and two protons that can then enter the pp chain again. However in terms of neutrinos, the two subdominant branches ppII and ppIII are much more interesting, especially since the pep and pp neutrinos, who make up the main part of the solar neutrinos, are very low in energy and thus difficult to detect.

If a $^3_2$He does not fuse with another $^3_2$He, but with a $^4_2$He instead, $^7_4$Be is formed. In most cases, this will then capture an electron to produce $^7_3$Li under the emission of a neutrino in the ppII branch:

$$^7_4\mathrm{Be} + e^- \rightarrow \ ^7_3\mathrm{Li} + \nu_e \tag{2.9}$$

These so-called ^{7}Be neutrinos have an energy of 0.86 MeV. Due to this high energy and their rather high flux portion of about 14 %, they were targeted in the first detection of solar neutrinos [15]. The $^7_3$Li finally catches another proton to form two $^4_2$He nuclei.

Sometimes, the $^7_4$Be reacts with a proton rather than an electron (ppIII branch) and forms $^8_5$B, which is a β^+ emitter with a half-life of 0.77 s [20]. Although their flux is very low, the very high Q-value of 14.1 MeV of this decay makes the so-called ^{8}B a favourable target for the search for solar neutrinos.

The excited $^8_4$Be* created in the decay

$$^8_5\mathrm{B} \rightarrow \ ^8_4\mathrm{Be}^* + e^+ + \nu_e \tag{2.10}$$

will then further decay into two α particles immediately.

2.2.5 Atmospheric Neutrinos

Neutrinos that are created in the interaction of high-energy charged particles with atoms in the Earth's atmosphere are several orders of magnitude higher in energy. During the past decades, the spectrum of these particles, the so-called cosmic radiation, has been measured in great detail over many orders of magnitude. Although it covers such a wide range in energy and flux, it shows almost no features and can be described by a simple power law with a spectral index of $\gamma = -2.7$, softening to $\gamma = -3.0$ at the so-called "knee" at a few PeV and turning back to $\gamma = -2.7$ at the "ankle" at the highest energies [21].

The origin of the cosmic radiation is not fully established yet. However it is commonly assumed that the particles are accelerated in non-thermal processes, usually moving shock fronts, that develop in extreme astrophysical environments. This mechanism is known as Fermi acceleration [26]. Candidates for the acceleration sites are both galactic sources, such as supernova remnants, as well as extragalactic ones like gamma ray bursts or active galaxies. Due to their small size and rather low energy density, galactic sources are believed to dominate the low-energy part of the spectrum in the GeV to TeV regime, while the extragalactic contribution takes over at the knee region.

When such a high energy particle hits the Earth's atmosphere, it interacts with a nucleus in the air (usually nitrogen or oxygen) in a so-called deep-inelastic scattering process. Since the energy of the incoming particle is far beyond all binding energies in the nucleus, it is completely disrupted. From the fragments of the nucleus that are still highly energetic, a shower of secondary particles develops, which then travels down to Earth.

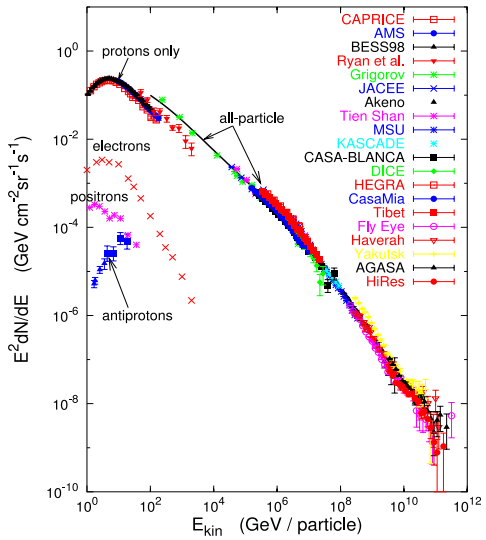

(a) The all particle spectrum of the charged cosmic radiation. Figure taken from [21].

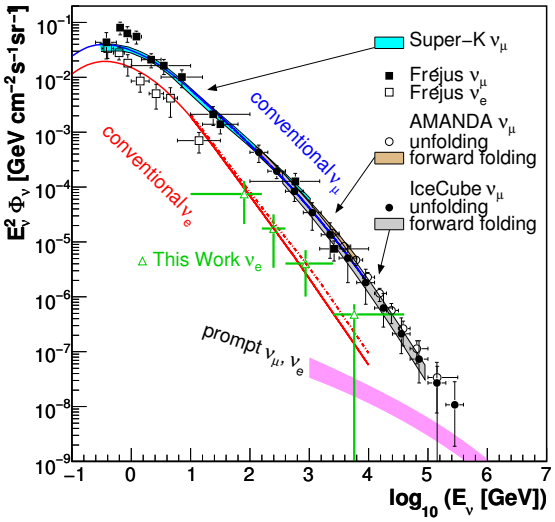

(b) The atmospheric flux of electron and muon neutrinos. Predictions for the conventional flux are from [22] (solid lines) and [23] (dashed), the band for the prompt flux is according to [24]. Figure taken from [25].

Figure 2.4: Spectra of the cosmic radiation at Earth and the resulting atmospheric neutrino spectrum.

Main components of these particle showers are kaons, pions, and muons. All of these particles are relatively short-lived[6] and produce neutrinos in their decay, according to (2.3), (2.4), and

$$\mu^{\pm} \rightarrow e^{\pm} + \overset{(-)}{\nu}_e + \overset{(-)}{\nu}_\mu \tag{2.11}$$

As one can see from the above equations, these so-called conventional atmospheric neutrinos coming mostly from pion and kaon decays have a flavour ratio of $\overset{(-)}{\nu}_\mu : \overset{(-)}{\nu}_e = 2 : 1$.

Their energy spectrum is linked to the primary spectrum of cosmic rays since the secondary pions and muons follow the primary energy distribution directly. Thus being highly relativistic, their lifetime is Lorentz boosted by a factor $\gamma \propto E$. So the higher the particle's energy, the higher is its probability *not* to decay into neutrinos in-flight but reach the Earth's surface and interact there, producing another shower of much less energetic particles. Therefore, the spectrum of conventional atmospheric neutrinos is suppressed roughly by a factor of $1/E$ with respect to the primary cosmic ray distribution.

For a more precise calculation, details of high-energy proton interactions and the geomagnetic field have to be taken into account [22, 23]. As shown in Fig. 2.4b, these predictions show good agreement with measurements.

Another predicted, but not yet observed component of the atmospheric neutrino flux are the so-called prompt neutrinos. They originate from charmed mesons that are rarely produced in cosmic ray induced air showers as well. Since these are so short-lived that they always decay in-flight ("promptly") despite of their relativistic boost, their energy spectrum is the same as the primary one and thus harder than the conventional component, but with a much smaller normalisation.

[6] $\tau_\mu = 2.2 \times 10^{-6}$ s, $\tau_K = 1.2 \times 10^{-8}$ s, and $\tau_\pi = 2.6 \times 10^{-8}$ s [27]

2.2.6 Astrophysical Neutrinos

The highest energy neutrinos are the so-called astrophysical ones. They are assumed to be produced at similar sites as the cosmic radiation, i. e. in highly energetic shock fronts. There, Δ resonances are generated in the collision of protons and high-energy photons, producing pions in their decay:

$$p + \gamma \rightarrow \Delta^+ \rightarrow n + \pi^+ \; (p + \pi^0) \tag{2.12}$$

The pions then decay further into neutrinos as shown in (2.3).

Since astrophysical neutrinos are produced in the same processes as cosmic rays, their fluxes are linked. The flux of cosmic rays has been measured in quite some detail over the recent decades [28], hence an upper limit for the flux of astrophysical neutrinos can be derived that is independent of any model assumptions about the production sites [29].

The first high energy astrophysical neutrino events have been recorded only recently by the IceCube neutrino telescope [30, 31], making them the highest energy neutrinos ever observed. Although event statistics are still low, the flux seems to be very close to the predicted upper bound, meaning that the neutrino production efficiency is close to maximal. The spectral shape of the flux is compatible with a power law with an index of –2 and an exponential cut-off around a few PeV as well as a steeper power law with index –2.3 [32].

2.3 Detection of Neutrinos

As already mentioned, neutrinos only interact with other particles in weak processes where the total cross-sections are typically very low. Thus high fluxes or large target volumes (or both) are needed to detect a sufficient number of neutrinos.

And even if these requirements are fulfilled, in most cases the neutrino signal has to be distinguished from a background of dominant processes, whose rate can be several orders of magnitude higher than the neutrino event rate. Depending on the targeted energy range, the most common background processes are inherent radioactivity of the surroundings and the detector itself at MeV energies, and muons created in cosmic ray induced air showers which can penetrate even strong shielding.

2.3.1 Neutrino cross-sections

The calculation and experimental testing of neutrino cross-sections has been a field of extensive research over the past decades. On the experimental side, the challenge is the smallness of the cross-sections. The key point on the theoretical side is the calculation of the matrix elements associated with the interaction of interest.

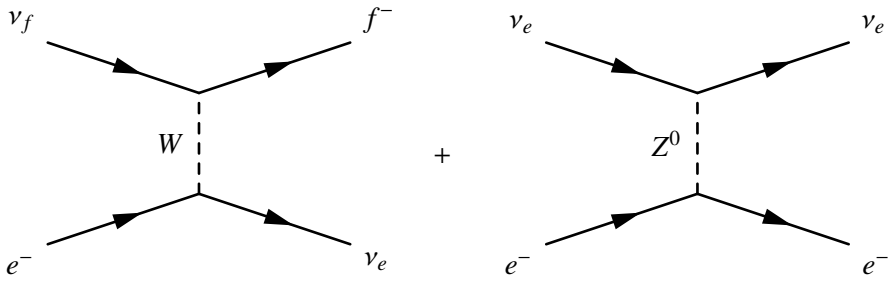

Figure 2.5: Feynman diagrams for the charged (left) and neutral (right) current contributions of $\nu_f e^- \rightarrow \nu_e f^-$ scattering.

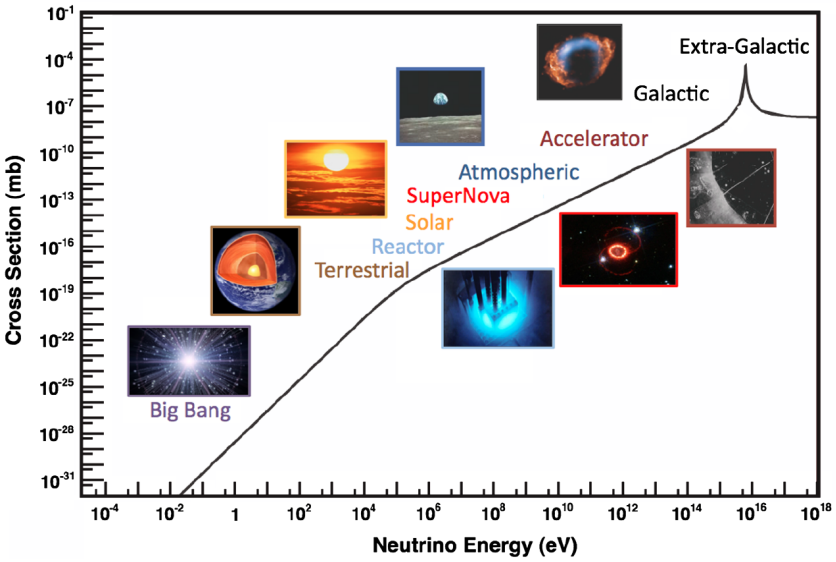

Figure 2.6: Electroweak cross-section for $\nu_e e^- \rightarrow \nu_e e^-$ scattering on free electrons as a function of neutrino energy. Various neutrino sources are also shown at their respective energy scales. [33]

Neutrinos scattering off a free electron will serve as an introductory example. The cross-section is [33]

$$\frac{d\sigma}{dq^2} = \frac{1}{16\pi} \frac{|\mathcal{M}^2|}{(s - m_e^2)^2} \,, \tag{2.13}$$

with s and q^2 being the centre-of-mass energy and the four-momentum transfer, respectively, assuming very small neutrino mass. In this case, calculating the matrix element $|\mathcal{M}^2|$ is rather straightforward, since only weak interactions between fundamental particles have to be considered. The scattering process itself is the sum of a charged-current ($W^\pm$ exchange, CC)

$$\mathcal{M}_{\mathrm{CC}} = -\frac{G_F}{\sqrt{2}} \left\{ \left[\bar{l} \gamma^\mu \left(1 - \gamma^5\right) \nu_l \right] \left[\bar{\nu}_e \gamma_\mu \left(1 - \gamma^5\right) e \right] \right\} \tag{2.14}$$

and a neutral current (Z^0 exchange, NC) contribution,

$$\mathcal{M}_{\mathrm{NC}} = -\sqrt{2} G_F \left\{ \left[\bar{\nu}_l \gamma^\mu \left(g_V^\nu - g_A^\nu \gamma^5\right) \nu_l \right] \left[\bar{e} \gamma_\mu \left(g_V^f - g_A^f \gamma^5\right) e \right] \right\} \quad , \tag{2.15}$$

as shown in Fig. 2.5. The vector and axial couplings of the involved leptons are given by:

$$g_V^\nu = g_A^\nu = +\frac{1}{2} \tag{2.16}$$

$$g_V^f = -\frac{1}{2} + 2\sin^2 \vartheta_W \tag{2.17}$$

$$g_A^f = -\frac{1}{2} \tag{2.18}$$

After converting the momentum transfer to the energy fraction y carried by the outgoing lepton,

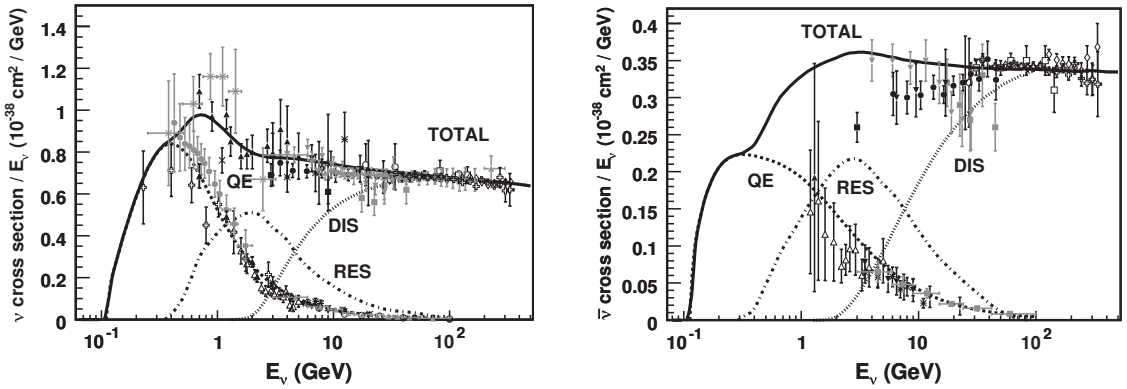

Figure 2.7: Total CC cross-section for neutrino (left) and antineutrino (right) cross-section for an isoscalar nucleon, $N = (p + n)/2$, divided by the neutrino energy and plotted as a function of energy. Shown are data from various experiments and predictions for the quasi-elastic (QE), resonance (RES), and deep inelastic (DIS) contributions [33].

$\frac{dq^2}{dy} = 2m_e E_\nu$, the charged current cross-section for scattering off an electron is given by [33]

$$\left.\frac{d\sigma}{dy}\right|_{\mathrm{CC}} = \frac{2m_e G_F^2 E_\nu}{\pi}\left(1 - \frac{m_f^2 - m_e^2}{2m_e E_\nu}\right), \tag{2.19}$$

where E_ν is the incoming neutrino energy, and m_e and m_f are the masses of the electron and the outgoing fermion. The corresponding total cross-section for $\nu_e e^- \to \nu_e e^-$ scattering is shown in Fig. 2.6. In case that m_f can be neglected compared to the neutrino energy, the above can be integrated to give a simple expression for the total cross-section:

$$\sigma \simeq \frac{2m_e G_F^2 E_\nu}{\pi} = \frac{G_F^2 s}{\pi} \tag{2.20}$$

In Fig. 2.6, this corresponds to the regime above about 10^7 eV, well above the electron mass. The sharp peak at $\approx 10^{16}$ eV is the so-called Glashow resonance, at which the centre-of-mass energy is comparable to the W boson mass, enabling resonant production of real W^- bosons and hence causing a strong enhancement of the cross-section.

When looking at other scattering processes, the principles of deriving the cross-section remain the same. However the kinematic part is subject to change, mostly due to different angular momentum states, and of course the matrix elements depend strongly on the respective target. For non-fundamental targets, form factors describing the internal charge distribution have to be taken into account as well.

In general, the cross-section for hadronic interactions are much larger than the leptonic ones (e. g. neutrino-electron scatting as discussed above) due to the larger target masses. This means that for conventional targets consisting of atoms with a nucleus and an electron hull, a neutrino is much more likely to interact with the nuclei than with the shell electrons.

2.3.2 Neutrino interactions with hadrons at the GeV scale

For the scope of this thesis, the most interesting energy regime is the low GeV scale, especially the range of $1 - 50$ GeV. Here the cross-section for neutrino interactions with nuclei is quite complex to describe,

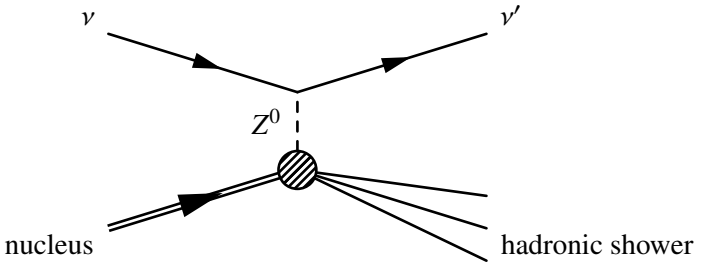

Figure 2.8: Neutral current interaction between a neutrino and a nucleus.

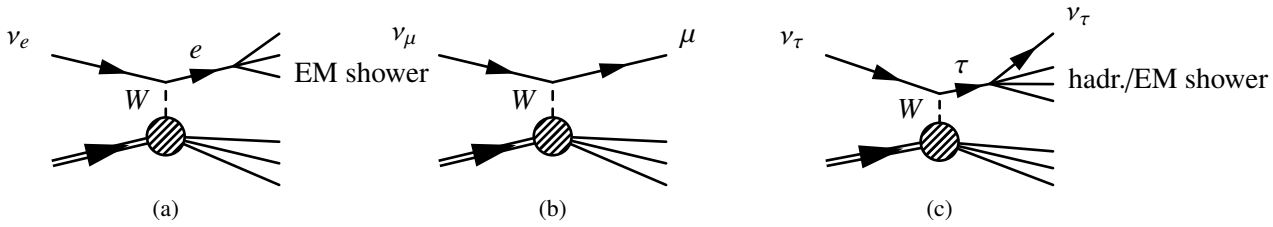

Figure 2.9: Charged current interactions between a ν_e (a), ν_μ (b), and ν_τ (c) and a nucleus.

as several distinct processes (shown in Fig. 2.7) have to be considered for the scattering:

Quasi-elastic scattering: At rather low energies, the neutrino scatters off an entire nucleon, removing it (possibly together with other nucleons) from the nucleus. The free proton(s) or neutron(s) will then propagate through the surrounding medium until they have dissipated all their energy. These interactions range out quickly above about 10 GeV.

Resonance production: At the range of about $1 - 3$ GeV, the dominant process is the excitation of short-lived baryonic resonances (such as Δ^+ or N^*) in the target nucleon. These resonances then decay to various final states producing nucleons and π mesons.

Deep inelastic scattering: Above about 10 GeV, the scattering neutrino has sufficient energy to resolve the quark structure of the nucleons. Then it scatters on a quark constituent rather than the whole nucleon. Thereby the nucleon gets disrupted and an hadronic shower consisting of a variety of mesons forms from its remains.

Common to all these processes is that they have both a CC and a NC contribution—similar to the scattering off electrons shown in Fig. 2.5, only that the target is either a whole (for quasi-elastic and resonant processes) or a quark inside a nucleon (deep inelastic scattering).

This means that four different classes of events can be observed. The first are neutral current interactions of any neutrino flavour, as shown in Fig. 2.8. In this case, the final state consists of the scattered neutrino ν' and a hadronic shower or cascade consisting of a variety of mesons, fermions, and photons developing from the fragments of the stricken nucleus.

The details of the development of such a hadronic shower are not fully understood and have to be modelled in Monte Carlo based simulations. However, since the involved particles are mostly short-lived and strongly interacting, the typical size is on the order of one meter or below.

The outgoing neutrino on the other hand is virtually impossible to detect, meaning that the fraction of the total interaction energy that is carried by this neutrino remains invisible. In beam experiments, in which the energy of the incoming particles is known, this is commonly referred to as “missing energy”.

In charged current interactions, the neutrino scatters off a nucleus by exchanging a W boson. From the fragmented nucleus, a hadronic shower develops similarly to the neutral current case. The outgoing

lepton however is now a charged one, the flavour corresponding to the incoming neutrino's flavour. This charged lepton now dominates the event signature:

Electron: In dense matter, a GeV electron will induce an electromagnetic shower by radiating high-energy photons (bremsstrahlung), which in turn will create electron-positron pairs. The typical scale length X_0 of such a shower is given by

$$\frac{1}{X_0} = 4\alpha r_e^2 \frac{N_A}{A} \left\{ Z^2 \left[L_{\text{rad}} - f(Z) \right] + Z\, L'_{\text{rad}} \right\} \quad , \tag{2.21}$$

where α is the fine structure constant, r_e the classical electron radius, and N_A Avogadro's number. A and Z are atomic mass and charge of the medium, and L_{rad}, L'_{rad}, and $f(Z)$ semi-empirical parameters that have been tabulated [27, 34]. For water, the scale length is $X_{0,\,\text{H}_2\text{O}} = 24.33\,\text{g/cm}^2$, meaning that the typical size of such a shower is below one metre, as water has a density of $\varrho_{\text{H}_2\text{O}} = 1\,\text{g/cm}^3$.

Muon: In contrast to the much lighter electron, a muon of a few GeV is a so-called minimum ionising particle. This means it is in the minimum region of the Bethe-Bloch formula for the energy loss of a charged particle in matter:

$$\left\langle -\frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2 m_e c^2 \beta^2 \gamma^2 W_{\text{max}}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \quad . \tag{2.22}$$

Here, $K = 4\pi N_A r_e^2 m_e c^2$, $z = 1$ is the electrical charge of the muon, I the mean excitation potential of the material, $\delta(\beta\gamma)$ a density effect correction, and W_{max} the maximum energy transfer in a single collision between the muon and an electron in the medium. $\beta = v/c$ and γ are the Lorentz variables [27].

This expression has a wide minimum in the range $\beta\gamma \approx 1 - 100$, corresponding to a muon momentum in the low GeV/c regime. Here, (2.22) evaluates to $\langle dE/dx \rangle \approx 2\,\text{MeV/(g/cm}^2)$, hence for a medium with a density of $1\,\text{g/cm}^3$ a muon has a range of approximately 5 m/GeV before it has deposited all its energy and decays at rest.

Tau: In principle, a tau lepton is a minimum ionising particle as well. But with a lifetime of only 2.9×10^{-13} s and a mass of 1.78 GeV, at GeV energies it can travel at most a few hundred microns before it decays. In its decay, a ν_τ has to be created again, whose energy has to be considered "missing" as well. The remaining decay products form an electromagnetic or hadronic shower, depending on their nature.

In a large detector with detection units spaced widely on a scale of several meters, like the PINGU neutrino telescope described in Sec. 4.2, the four event classes discussed above can only be categorized into two channels: tracks and cascades.

Tracks correspond to ν_μ CC events. A muon track of several meters length is pointing away from the hadronic shower at the event vertex. In principle, this allows for a good directional reconstruction of the event. For the energy reconstruction on the other hand there is the disadvantage that the outgoing muon might not be fully contained in the detector, then only a fraction of its total energy can be recorded.

Cascades encompass all other types of events. A hadronic shower is always present, and another shower of electromagnetic or hadronic nature overlays it, if a CC interaction has occurred. The

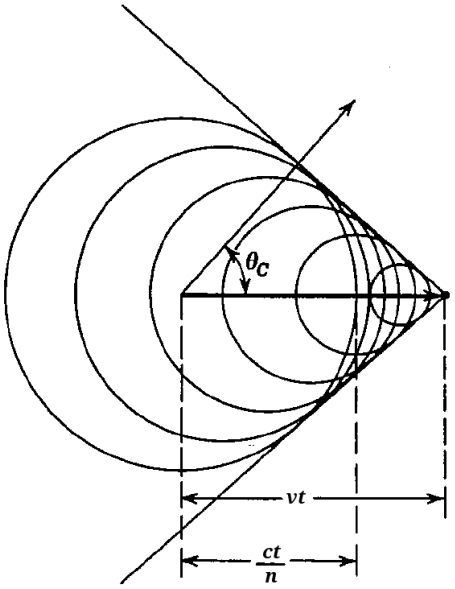

Figure 2.10: Illustration of the Cherenkov effect (schematic). Photons are emitted perpendicular to the surface of the Cherenkov cone at the angle ϑ_C, as indicated by the arrow. Graphics taken from [35].

displacement of the second shower (if it exists) is much smaller than a millimetre and hence not resolvable. Also the showers themselves have to be considered point-like, but have some directionality since the momentum of the incoming neutrino is transferred to the final state.

The achievable directional resolution is worse than for tracks since the events lack a long lever arm in form of an extended muon trajectory. Yet the energy resolution should be better at least for ν_e CC events, since the event is always fully contained in the detector. For NC and ν_τ CC events the energy of the outgoing neutrino in the final state is invisible, causing a reconstruction bias towards low energy.

2.3.3 Cherenkov Effect

As discussed above, high fluxes and/or large target volumes are required for the detection of neutrino interactions. Since the atmospheric flux is rather low compared to focused, high-intensity artificial beams, there is no way around building a megaton scale detector when searching for atmospheric neutrinos. Obviously, a "conventional" detector for subatomic particles[7] is not feasible at these dimensions.

The common choice for measuring natural (i. e. low) neutrino fluxes above MeV energies are water-based Cherenkov detectors. A large volume of water or ice serves as both target and detector: The neutrinos interact with the protons and neutrons in the water molecules, creating relativistic charged particles as described in the previous section. Then the charged particles emit light via the Cherenkov effect, which propagates in the transparent medium and can finally be detected (usually using photomultiplier tubes) to reconstruct the underlying neutrino event.

Whenever a charged particle moves through a dielectric medium, it shortly polarises the atoms along its path, which in turn emit electromagnetic waves. Usually waves from neighbouring atoms cancel, such that no net effect is observable. If, however, the velocity of the charged particle is higher than the

[7] made up from e. g. wire chambers, solid state scintillators, ...

local phase velocity of light given by the index of refraction,

$$v > c_{\text{medium}} = c/n \quad , \tag{2.23}$$

the wave fronts interfere constructively to form a shock front of conical shape, similar to the Mach cone in supersonic motion (see Fig. 2.10). Photons are emitted almost[8] rectangular to the shock front.

The opening angle of the Cherenkov cone is given by the relation

$$\cos\vartheta_{\text{Ch}} = \frac{c_{\text{medium}}}{v} = \frac{1}{\beta n} \quad . \tag{2.24}$$

It depends on the particle's velocity and hence its energy. Usually, the constituents of the particle showers are relativistic, so $\beta \simeq 1$ and hence

$$\vartheta_{\text{Ch}} \simeq \arccos\frac{1}{n_{\text{ice}}} \simeq 41^\circ \tag{2.25}$$

for ice with $n_{\text{ice}} \simeq 1.32$ [36].

The energy loss due to Cherenkov radiation is described by the Frank-Tamm formula that can be found in [35]:

$$\frac{dE}{dx} = \frac{(Ze)^2}{c^2}\int_{n(\omega)>1/\beta^2}\omega\left(1-\frac{1}{(\beta n(\omega))^2}\right)d\omega \quad . \tag{2.26}$$

This integral does not diverge as the refractive index drops below 1 in the ultraviolet region.

Its contribution to the total energy dissipation of the particle[9] is marginal, but from this the rate of photon production can be inferred. Expressed as an energy spectrum, the number of photons generated per track length is [27]

$$\frac{d^2N_\gamma}{dx\,dE} = \frac{\alpha Z^2}{\hbar c}\sin^2\vartheta_{\text{Ch}} = \frac{\alpha Z^2}{\hbar c}\left(1-\frac{1}{(\beta n(E))^2}\right) \quad , \tag{2.27}$$

where for simplicity the refractive index was treated as constant. Evaluating the constants, one finds for particles of unit charge ($Z^2 = 1$) in ice:

$$\frac{d^2N_\gamma}{dx\,dE} \approx 370\,\sin^2\vartheta_{\text{Ch}}\,\text{eV}^{-1}\text{cm}^{-1} \approx 175\,\text{eV}^{-1}\text{cm}^{-1} \quad . \tag{2.28}$$

One also often finds a version of this formula giving a wavelength spectrum [27]:

$$\frac{d^2N_\gamma}{dx\,d\lambda} = \frac{2\pi\alpha Z^2}{\lambda^2}\left(1-\frac{1}{(\beta n(\lambda))^2}\right) \quad . \tag{2.29}$$

Here the $1/\lambda^2$ dependency of the spectrum becomes evident[10]. When using this wavelength dependent equation, one has to bear in mind that λ always means the *vacuum* wavelength and not the wavelength in the medium that differs from that in the vacuum by a factor of n.

As mentioned above, the energy dependent expression can be treated as constant with good accuracy. Then the integration over a certain energy interval becomes trivial and one can easily estimate the

[8] As group and phase velocity differ, the angle is not quite 90°.
[9] Which is described by the Bethe-Bloch formula (2.22).
[10] This is the reason for the bright blue colour in which Cherenkov light appears to the human eye.

number of Cherenkov photons for a certain particle. Again using the values for ice, the calculation yields a number of

$$\frac{dN_\gamma}{dx} \approx 325\,\mathrm{cm}^{-1} \tag{2.30}$$

photons in the visible range between 300 nm and 600 nm of vacuum wavelength.

CHAPTER 3

Neutrino Oscillations

The Standard Model of Particle Physics, as described in the previous chapter, has been one of the most successful theories in the history of physics. It is, however, not fundamental in the sense that it can explain all physical phenomena alone. Its shortcomings are e. g. the missing inclusion of the fourth fundamental force, gravity, and a lack of explanation for the fundamental asymmetry between bosons and fermions. There are theoretical extensions to the Standard Model addressing these questions, such as "Grand Unified Theories", supersymmetry, and many others [37], but all of them lack experimental evidence so far.

Yet there is one effect of so-called "Physics beyond the Standard Model" that has been well established experimentally during the past years: neutrino oscillations. As already mentioned in Sec. 2.1, this term refers to neutrinos changing their flavour when travelling over macroscopic distances, which can be explained by finite neutrino masses, while in the Standard Model they have zero mass.

The theory behind this process will be described in the following. In-depth treatments of this topic can be found in many textbooks, e. g. [37–40]. The notation will follow [38] here.

3.1 Vacuum Oscillations

There are two bases of eigenstates to which a neutrino can be decomposed: the flavour and the mass base. The flavour eigenstates are $|\nu_e\rangle$, $|\nu_\mu\rangle$, and $|\nu_\tau\rangle$, which will be summarised as $|\nu_\alpha\rangle$. These are the eigenstates of the weak interaction, hence neutrinos are always produced as a pure flavour eigenstate and have to be projected back onto these eigenstates whenever they interact.

On the other hand there are the three mass eigenstates $|\nu_1\rangle$, $|\nu_2\rangle$, and $|\nu_3\rangle$, summarised as $|\nu_k\rangle$, corresponding to the three neutrino masses m_k. The absolute values of these masses are yet unknown, but since neutrino oscillations have been observed, at least two of them have to be different from zero. The mass eigenstates have to be considered when describing the propagation of a neutrino in vacuum since they are the eigenstates of the corresponding Hamiltonian

$$\hat{H}\,|\nu_k\rangle = E_k\,|\nu_k\rangle \quad . \tag{3.1}$$

3.1.1 General Case

Changes between the two bases are carried out via the so-called PMNS matrix[1] $\mathcal{U}_{\text{PMNS}}$ that can be parametrised using three Euler angles ϑ_{ij}, also called mixing angles, and one complex phase angle δ that is related to possible CP violation:

$$\mathcal{U}_{\text{PMNS}} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{3.2}$$

Here, s_{ij} and c_{ij} are shorthands for $\sin\vartheta_{ij}$ and $\cos\vartheta_{ij}$, respectively. Transformations between flavour and mass base are then given by

$$|\nu_\alpha\rangle = \sum_k U^*_{\alpha k} |\nu_k\rangle \quad , \tag{3.3}$$

$$|\nu_k\rangle = \sum_\alpha U_{\alpha k} |\nu_\alpha\rangle \quad . \tag{3.4}$$

To ensure lepton number conservation, unitarity of $\mathcal{U}_{\text{PMNS}}$ has to be required. Any deviation from this can be interpreted as a hint for additional neutrino flavours that do not participate in the weak interaction[2]. Such signals have been reported (e. g. [42]), but the overall picture remains inconclusive [27].

If now a pure flavour eigenstate $|\nu_\alpha\rangle$ is produced at an energy E, the probability to detect it as $|\nu_\beta\rangle$ after propagating over a distance L has to be calculated according to

$$P_{\alpha\to\beta} = \left| \left\langle \nu_\beta \left| \mathcal{U}^\dagger_{\text{PMNS}} \left| \hat{H} \right| \mathcal{U}_{\text{PMNS}} \right| \nu_\alpha \right\rangle \right|^2 \quad . \tag{3.5}$$

In the mass base, the propagation of the neutrinos can be described as plain waves,

$$|\nu_k(t)\rangle = \exp\left(-i(E_k t - \vec{p}_k \cdot \vec{x})\right) |\nu_k(0)\rangle \quad , \tag{3.6}$$

and assuming relativistic neutrinos ($m_k \ll E_k \Rightarrow v \approx c$) one can approximate in natural units[3]

$$E_k = \sqrt{\vec{p}_k^2 - m_k^2} \approx p_k + \frac{m_k^2}{2p_k} \quad , \tag{3.7}$$

$$E_k t - \vec{p}_k \cdot \vec{x} = \left(p_k + \frac{m_k^2}{2p_k} \right) L - p_k L \approx \frac{m_k^2}{2E} L \tag{3.8}$$

With this, (3.5) reduces to

$$P_{\alpha\to\beta} = \sum_{k,j} U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} \exp\left(-i \frac{\Delta m^2_{kj} L}{2E} \right) \tag{3.9}$$

[1] After Bruno Pontecorvo, Ziro Maki, Masami Nakagawa, and Shoichi Sakata.

[2] Measurements of the Z^0 decay width have shown that only three weakly interacting neutrino flavours exist [41]—at least at masses up to half the Z^0 mass, $m_\nu < m_{Z^0}/2 = 45.6\,\text{GeV}$.

[3] $\hbar = c = 1$

with the squared mass differences

$$\Delta m^2_{kj} \equiv m^2_k - m^2_j \quad . \tag{3.10}$$

In the mass base, the Hamiltonian can be replaced by an effective one containing the squared masses:

$$\hat{H}^{\text{eff}} = \frac{1}{2E}\,\text{diag}\left(m_1^2,\, m_2^2,\, m_3^2\right) = \frac{m_1^2}{2E}\,\mathbb{1} + \frac{1}{2E}\,\text{diag}\left(0,\, \Delta m^2_{21},\, \Delta m^2_{31}\right) \tag{3.11}$$

The first summand on the r. h. s. of the above equation can even be ignored since it will only introduce an unobservable global phase shift. This simplification will prove handy when discussing oscillations in matter in Sec. 3.3.

Using the unitarity of $\mathcal{U}_{\text{PMNS}}$, (3.9) can now be rewritten as

$$\begin{aligned} P_{\alpha\to\beta} = \delta_{\alpha\beta} \quad &- \quad 2\sum_{k>j} \Re\left[U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}\right]\left[1 - \cos\frac{\Delta m^2_{kj} L}{2E}\right] \\ &+ \quad 2\sum_{k>j} \Im\left[U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}\right]\sin\frac{\Delta m^2_{kj} L}{2E} \quad . \end{aligned} \tag{3.12}$$

Obviously, this oscillation probability collapses to $P_{\alpha\to\beta} = \delta_{\alpha\beta}$ if all m_i are equal. Hence the observation of actual flavour conversion means that the m_i are different from each other and in particular different from zero (at least two of them), contradicting the standard model prediction of vanishing neutrino masses. Additionally, the second oscillatory term only contributes if there is CP violation in the neutrino sector—otherwise the mixing matrix (3.2) is real.

For antineutrinos, the oscillation probability is derived analogously, only with $\mathcal{U}_{\text{PMNS}}$ being replaced by its complex conjugate. Hence differing vacuum oscillation probabilities for neutrinos and antineutrinos are a proof of CP violation.

3.1.2 Two Flavour Case

In many cases[4] it is sufficient to consider only two neutrino flavours in the oscillation. Then there is only one mass splitting

$$\Delta m^2 \equiv \Delta m^2_{21} \equiv m^2_2 - m^2_1 \tag{3.13}$$

and the mixing matrix $\mathcal{U}$ can be parametrised by one effective mixing angle

$$\mathcal{U} = \begin{pmatrix} \cos\vartheta & \sin\vartheta \\ -\sin\vartheta & \cos\vartheta \end{pmatrix} \quad . \tag{3.14}$$

The expression for the transition probability simplifies to

$$P_{\alpha\to\beta} = \sin^2 2\vartheta \sin^2\left(\frac{\Delta m^2 L}{4E}\right) = \sin^2 2\vartheta \sin^2\left(\pi\frac{L}{L^{\text{osc}}}\right) \qquad (\alpha \neq \beta) \quad , \tag{3.15}$$

introducing the oscillation length

$$L^{\text{osc}} \equiv \frac{4\pi E}{\Delta m^2} \approx 2.47\frac{E\,[\text{GeV}]}{\Delta m^2\,[\text{eV}^2]}\text{km} \quad . \tag{3.16}$$

[4] In particular, if the survival probability of a certain flavour is measured.

From (3.15), the two different groups of parameters in neutrino oscillation phenomenology and how they influence the oscillation probabilities, become obvious: The mixing angles define the amplitude of the oscillation, with $\vartheta = 45°$ giving rise to so-called "maximum mixing" where a full transition from one flavour to another is possible. The mass splittings determine the frequency at a given neutrino energy, expressed through the oscillation length at which the first full oscillation cycle is completed.

So from an experimental point of view, placing a detector at a distance $L = L^{\text{osc}}/2$ from the neutrino source is preferential, since here the oscillation effects are strongest. If $L \ll L^{\text{osc}}$, the flavour transition has not yet happened while at $L \gg L^{\text{osc}}$ only the average transition probability

$$\left\langle P_{\alpha\to\beta} \right\rangle = \frac{1}{2} \sin^2 2\vartheta \tag{3.17}$$

can be measured and no information on Δm^2 can be obtained[5].

3.2 Absolute Neutrino Masses and Mass Hierarchy

Since the existence of neutrino oscillations has unambiguously shown that neutrinos have non-zero masses, the question is what the absolute values of these masses are. Although this question seems to be very simple, it turns out to be experimentally challenging.

To establish absolute neutrino masses one has to consider effects such as distortions at the upper end of the energy spectrum of nuclear β decays (as described in Sec. 2.2.1). Here the decay of tritium is promising due to its small decay energy. In fact, currently the most stringent upper limits for the mass of the electron (anti-) neutrino [6] set by the Mainz and Troitsk experiments [44, 45] stem from this very decay. Their limit of

$$m_{\bar{\nu}_e} \lesssim 2.1\,\text{eV} \tag{3.18}$$

is expected to be improved by one order of magnitude in the KATRIN experiment that targets the tritium decay spectrum as well [46].

However these experiments can only directly measure the superposition of mass eigenstates that corresponds to the ν_e or $\bar{\nu}_e$ flavour eigenstate. A direct measurement of the ν_μ or ν_τ mass (or their antiparticles) is by far more difficult since they cannot be created in a controllable source like a nuclear decay. The only not completely unrealistic options here would be time-of-flight measurements with an extremely long baseline[7].

So the most promising overall approach is to fix the neutrino mass scale at one point by measuring the ν_e mass directly and then derive the other mass eigenstates via the mass differences that are accessible in neutrino oscillations. On the other hand, apart from CP violating effects, which have not yet been observed, all oscillatory terms above are proportional to either cos or $\sin^2$ and hence insensitive to the sign of their argument. Thus, in vacuum oscillations, only information on the distances of the mass eigenstates can be collected, but not on their relative ordering or *hierarchy*.

This could be resolved, however, if one would measure all three mass splittings separately and then use that obviously

$$\Delta m^2_{31} = \Delta m^2_{32} + \Delta m^2_{21} \quad . \tag{3.19}$$

[5] Here it is assumed that in a real experiment one will always have a continuous neutrino energy spectrum and hence a distribution of oscillation lengths. Hence the fast oscillations will smear out on a distance $L \gg L^{\text{osc}}$.

[6] Defined as $m^2_{\nu_e} = \sum_i |U^2_{ei}|\, m^2_i$ [43].

[7] E. g. astrophysical neutrinos that can be associated with a transient optical event such as a supernova or a gamma-ray burst.

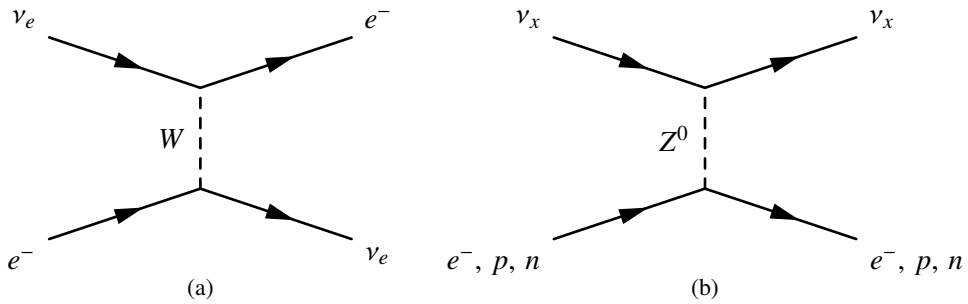

Figure 3.1: Feynman diagrams for the charged (a) and neutral (b) current contributions to coherent forward scattering in matter.

to derive the ordering. Unfortunately, it turns out [47, 48] that

$$\Delta m^2_{32} \simeq \Delta m^2_{31} \gg \Delta m^2_{21} \quad , \tag{3.20}$$

such that with current experiments' precision it is impossible to disentangle Δm^2_{32} and Δm^2_{31}. So what can be done to learn about the ordering of the neutrino mass eigenstates?

There is a possibility to access the neutrino mass hierarchy in oscillation experiments, if the neutrinos in question pass through a sufficient amount of matter along their path. In this case, additional resonances appear that depend on the sign of the mass splittings. These so-called matter effects will be discussed in the following section.

3.3 Oscillations in Matter

When neutrinos are passing through matter, they will undergo coherent forward scattering off the electrons and nucleons in their path. Since the matter distribution is continuous, this can be interpreted as a matter potential which the neutrinos are experiencing, leading to a change of their effective mass similar to photons passing through a transparent medium. The implications of this scenario were first described by L. Wolfenstein in 1978 [49]

Neutral current interactions, as shown in Fig. 3.1b, are open to all neutrino flavours in a similar way. Here, the contributions from electrons and protons cancel since their associated weak currents have opposite sign[8] and only the neutron potential remains:

$$V_{\mathrm{NC}} = -\frac{1}{2}\sqrt{2}\,G_F\,N_n \tag{3.21}$$

Yet since this potential affects all flavours and hence all mass eigenstates in the same way, it also changes all effective masses by the same amount, but leaves the mass splittings unaffected. Charged current scattering (Fig. 3.1a) on the other hand is only possible for ν_e as electrons are the only charged leptons present in ordinary matter. The CC matter potential can be expressed as

$$V_{\mathrm{CC}} = \sqrt{2}\,G_F\,N_e \quad . \tag{3.22}$$

This means that the eigenvalue equation (3.1) has to be modified to include the matter potential. Since the potential acts on the flavour eigenstates, the effective vacuum Hamiltonian (3.11) has to be

[8] Assuming that the matter is macroscopically neutral, hence electrons and protons have equal number densities.

transferred into the flavour base first, then the potential can be added:

$$\hat{H}_{\text{matter}} = \mathcal{U}^{\dagger} \hat{H}_{\text{eff}} \, \mathcal{U} + \text{diag}\left(V_{\text{NC}} + V_{\text{CC}},\, V_{\text{NC}},\, V_{\text{NC}}\right) \tag{3.23}$$

Here and in the following, the index has been dropped from $\mathcal{U}_{\text{PMNS}}$ to improve readability. Constant contributions can now be neglected again, which results in the following effective Hamiltonian in matter:

$$\hat{H}^{\text{eff}}_{\text{matter}} = \frac{1}{2E}\left[\mathcal{U}^{\dagger} \text{diag}\left(0,\, \Delta m^2_{21},\, \Delta m^2_{31}\right) \mathcal{U} + \text{diag}\left(A_{\text{CC}},\, 0,\, 0\right)\right] \quad , \tag{3.24}$$

where

$$A_{\text{CC}} \equiv 2E\, V_{\text{CC}} = 2\sqrt{2}\, E\, G_F\, N_e \quad . \tag{3.25}$$

Here it is very important to note that for antineutrinos the matter potential has the opposite sign. In particular, $\bar{\nu}_e$ will feel a charged current potential of

$$\bar{A}_{\text{CC}} = -A_{\text{CC}} \quad . \tag{3.26}$$

3.3.1 MSW Effect

For simplicity, the effects of a matter potential will be discussed for only two neutrino flavours here, without loss of generality. In this case, (3.24) reduces to

$$\hat{H}^{\text{eff}}_{\text{matter}} = \frac{1}{2E}\left[\mathcal{U}^{\dagger} \text{diag}\left(0,\, \Delta m^2\right) \mathcal{U} + \text{diag}\left(A_{\text{CC}},\, 0\right)\right] \quad . \tag{3.27}$$

Diagonalising this matrix now gives the effective mixing matrix and the mass splitting in matter:

$$\mathcal{U}^{\dagger}_{\text{matter}} \hat{H}^{\text{eff}}_{\text{matter}} \mathcal{U}_{\text{matter}} = \hat{H}^{\text{diag}}_{\text{matter}} \quad , \tag{3.28}$$

where

$$\hat{H}^{\text{diag}}_{\text{matter}} = \text{diag}\left(-\Delta m^2_{\text{M}},\, \Delta m^2_{\text{M}}\right) \tag{3.29}$$

and

$$\mathcal{U}_{\text{matter}} = \begin{pmatrix} \cos\vartheta_{\text{M}} & \sin\vartheta_{\text{M}} \\ -\sin\vartheta_{\text{M}} & \cos\vartheta_{\text{M}} \end{pmatrix} \quad . \tag{3.30}$$

The effective mass splitting is then given by

$$\Delta m^2_{\text{M}} = \sqrt{\left(\Delta m^2 \cos 2\vartheta - A_{\text{CC}}\right)^2 + \left(\Delta m^2 \sin 2\vartheta\right)^2} \quad , \tag{3.31}$$

while

$$\tan 2\vartheta_{\text{M}} = \frac{\tan 2\vartheta}{1 - \frac{A_{\text{CC}}}{\Delta m^2 \cos 2\vartheta}} \tag{3.32}$$

is the effective mixing angle in matter.

In 1985, S. P. Mikheyev and A. Y. Smirnov discovered [50, 51] that for

$$A^{\text{res}}_{\text{CC}} = \Delta m^2 \cos 2\vartheta \Leftrightarrow N^{\text{res}}_e = \frac{\Delta m^2 \cos 2\vartheta}{2\sqrt{2}\, E\, G_F} \tag{3.33}$$

a resonance exists where the effective mixing angle approaches $\pi/4$, meaning that the mixing amplitude

$\sin^2 2\vartheta$ in (3.15) becomes equal to one and a full transition from one flavour to another is possible. At the same time, the mass splitting becomes minimal.

Another thing to note is that in (3.32), both A_{CC} (by being replaced by $\bar{A}_{\mathrm{CC}} = -A_{\mathrm{CC}}$ for antineutrinos) and Δm^2 can carry a negative sign. Thus one can infer the sign of Δm^2 by observing the so-called MSW resonance[9] in either the neutrino or the antineutrino channel:

If, e. g., the hierarchy is normal—meaning that Δm^2 is positive—for antineutrinos, where the matter potential $\bar{A}_{\mathrm{CC}}$ is always negative, the denominator of (3.32) will always be larger than one and no resonance can occur. For neutrinos on the other hand, A_{CC} is positive and can cause a MSW resonance if it has the appropriate value.

This measurement has been done by the SNO experiment, which measured both the flux of solar ν_e and the all-flavour flux of solar neutrinos independently from each other with high precision [52]. Clear signs for a MSW resonance were detected for the solar ν_e[10] travelling through the high electron density of the inner Sun, hence one could conclude that the relevant mass splitting for the oscillation of solar neutrinos, Δm^2_{21}, has a positive sign.

In the planned PINGU experiment (for details, see Sec. 4.2), a similar measurement is envisaged to determine the sign of the other mass splitting Δm^2_{31}, which will be referred to as "Neutrino Mass Hierarchy" (NMH) in the following. Here the matter potential of the Earth will be used to observe a MSW resonance either in the atmospheric neutrino or antineutrino channel [53, 54]. The details of this measurement will be discussed in Sec. 3.5.

3.3.2 Parametric Enhancement

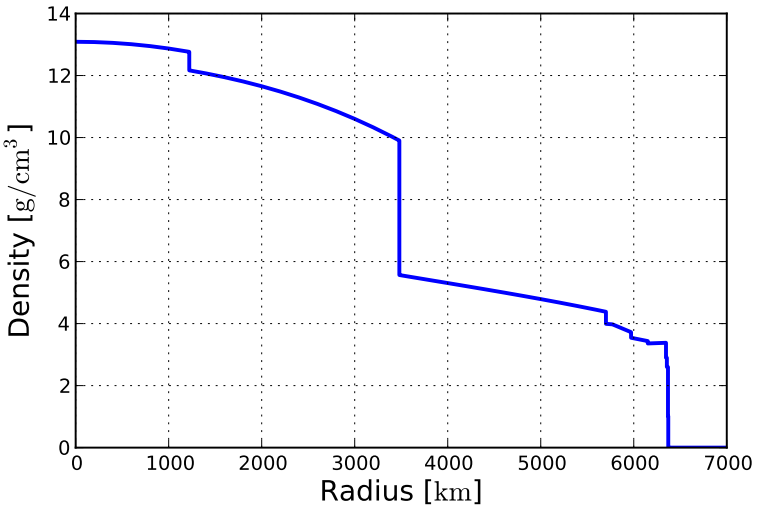

Figure 3.2: The PREM Earth density profile [55].

Another feature of neutrino oscillations in matter are parametric enhancements. In contrast to MSW resonances, they depend on the shape of the matter potential rather than its actual value.

As shown in Fig. 3.2, the density profile and hence the matter potential of the Earth is characterised

[9] After Mikheyev, Smirnov, and Wolfenstein.
[10] No antineutrinos are produced in the Sun, cf. Sec 2.2.4

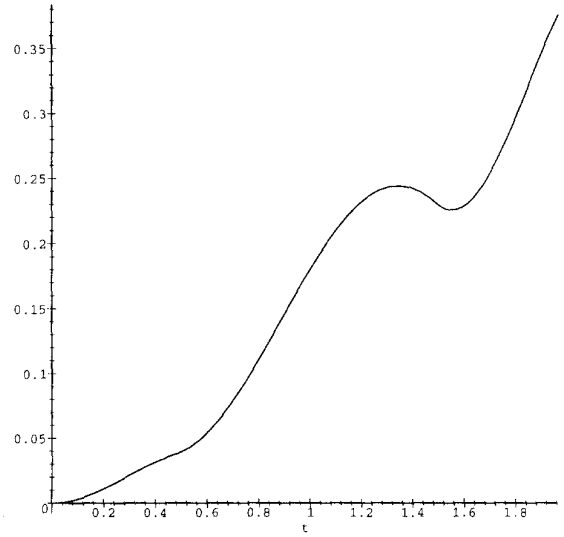

Figure 3.3: ν_e disappearance probability (qualitatively[11]) as a function of the distance t (measured in units of the Earth's radius) along the neutrino trajectory. Figure adopted from [56].

by two major regions: the core in the centre, with a radius of $\approx 3500\,\mathrm{km}$, surrounded by the mantle[12]. Within these regions, the density is rather constant at $\approx 11\,\mathrm{g/cm^3}$ and $\approx 5\,\mathrm{g/cm^3}$, respectively, while at the transition between the regions the density changes by almost a factor of two quasi instantaneously.

Thus for atmospheric neutrinos crossing the Earth, the matter potential can be approximated by a so-called "castle-wall" profile with good precision: First the neutrinos will experience a constant potential A^m_{CC} for a distance l_m while passing through the mantle, with A^m_{CC} given by (3.25) inserting the mean electron density in the mantle. Thereafter, they traverse the core, represented by another constant potential A^c_{CC} with a width of l_c, followed by a second crossing of the mantle that is symmetric to the first one.

The principal resonance occurs if both l_m and l_c are close to half the oscillation length (3.16) for the effective mass splittings (3.31) in the respective regions [56, 57]:

$$\frac{l_m}{L^{\mathrm{osc}}_m} = \frac{l_c}{L^{\mathrm{osc}}_c} = \frac{1}{2} \tag{3.34}$$

This can be achieved by choosing an appropriate zenith angle at which the neutrinos pass through the Earth.

Although the actual calculation of the oscillation probability is rather involved [56], the principal effect can be understood qualitatively: The transition probability is close to its maximum when the neutrinos cross the border between mantle and core, as shown in Fig. 3.3. At this point, the matter potential changes, thus the neutrino state has to be projected into the flavour base in which interactions have to be evaluated. This re-sets the effective propagation length the neutrinos have travelled to zero, such that the oscillation probability does not decrease after reaching its maximum in the first region. Instead, the oscillation is "restarted" and the oscillation probability in the second region is added to the one at the transition between the regions. The same happens at the second region transition, back from

[11] As [56] has been published already in 1999, the values assumed for the mixing parameters do not coincide with the current best fits. The effect in general, however, is the same for more up-to-date values.

[12] The Earth crust has a thickness of only a few tens of km and can thus be neglected.

the core to the mantle, hence the final oscillation probability after passing through the Earth is vastly enhanced w. r. t. the non-resonance case.

These parametric enhancements depend on the neutrino mass hierarchy in the same way as the MSW resonance, so they will also contribute to the NMH signature to observe with PINGU, yet on a smaller scale than said resonance as it will be shown in Sec. 3.5.

3.4 Oscillation Experiments

Over the course of the last decades, a variety of experiments targeting most of the possible neutrino sources listed in Sec. 2.2 have observed oscillations. Consequently, measured values for all of the mixing angles and mass splittings have been published, leading to a fairly consistent global picture [47, 48]. The most prominent of these experiments will be presented in this section, followed by an overview of the current best-fit values of the oscillation parameters.

3.4.1 Solar Neutrinos

As mentioned previously, already in the first detection of solar electron neutrinos from the decay of ^{8}B in the Homestake experiment in the 1960s [15], neutrino oscillations had been observed in the form of a lower than expected event rate. Although oscillations had been considered as the cause for the deficit [58], a measurement of the all-flavour solar neutrino flux was needed to exclude possible errors in the flux calculation. This was provided later by the Kamiokande experiment [59], however not in sufficient precision.

The required precision was finally reached by the SNO experiment, thereby making the first definite observation of solar electron neutrinos oscillating to other flavours [52, 60]. In a two-flavour approximation, values for Δm^2_{21} and ϑ_{12} could be published as well [61].

3.4.2 Atmospheric Neutrinos

The very first conclusive observation of neutrino oscillations was the disappearance of atmospheric muon neutrinos at energies around 1 GeV, reported by the Super-Kamiokande collaboration in 1998 [16]. This detection was facilitated by a large value of the relevant mixing angle ϑ_{23}, which is close to the maximum mixing value of $\pi/4$. Over the course of the last years, more data have been added to this analysis, improving its precision on the measured parameters ϑ_{23} and $\Delta m^2 = (\Delta m^2_{31} + \Delta m^2_{32})/2$.

Additionally, a similar measurement has been done with the IceCube DeepCore neutrino telescope at energies of several tens of GeV [62]. Reaching a comparable accuracy after a much shorter livetime, this result demonstrated that a precision oscillation measurement is possible with a detector using a natural target at which experimenters have much less control than over an artificial one. DeepCore's success has paved the way for PINGU [54], which will map atmospheric oscillations in the full three flavour picture with unprecedented accuracy[13].

3.4.3 Neutrino Beams

Experiments aiming at atmospheric neutrino oscillations have both the problem and the benefit that they will observe events over a large energy spectrum that come from all directions. Although it provides a

[13] With ORCA [63], a very similar experiment has been proposed that is supposed to do the same measurement, only using sea water as detection medium instead of ice.

large lever arm for fitting oscillation parameters, this also means that there is a fairly large uncertainty from event reconstruction.

A way to overcome that problem is to use a controllable neutrino source, such as a neutrino beam from an accelerator pointing towards a dedicated detector. Then, the total flux as well as the energy and arrival direction of the neutrinos is known, and one can concentrate on determining the oscillation parameters—especially the mass splittings, which depend strongly on a precise knowledge of the neutrino energy. If the value of the mixing parameter one is about to constrain is roughly known beforehand, one can even fine-tune the accelerator settings to reach maximal precision.

Examples are the MINOS [64] and T2K [65] experiments, both using a ν_μ beam from Fermilab or J-PARC, respectively, to measure Δm^2_{31} and ϑ_{23}—the same parameters accessible for atmospheric oscillations, thus providing an uncorrelated measurement.

Another goal that can be achieved with neutrino beams due to their high and well-known flux and clean flavour composition is the search for the appearance of neutrinos that have oscillated to other flavours. This has been done in T2K, where the appearance of ν_e has been observed [66], and also in OPERA, an experiment dedicated to search for ν_τ events appearing in a ν_μ beam from CERN [67].

3.4.4 Reactor Neutrinos

Finally another class of experiments uses the strong flux of $\bar{\nu}_e$ at low MeV energies produced by commercial nuclear reactors to study neutrino oscillations. Typically, they consist of one "near" detector as close to the reactor core as possible to achieve a precise normalisation of the un-oscillated neutrino flux and the detector acceptance, and one identical "far" detector in the first oscillation minimum, usually at several kilometres distance. The relevant mixing parameter in this regime is ϑ_{13}, the last one to be measured. The first measurement of ϑ_{13} was published in 2012 by the Daya Bay experiment [68], subsequently confirmed by RENO and Double Chooz [69, 70].

With reactor neutrinos, also the mass hierarchy is accessible. If a far detector is placed at a distance of $\approx 50\,\mathrm{km}$, the small difference between Δm^2_{31} and Δm^2_{32} will cause a fastly oscillating interference pattern on top of the principal oscillation probability, whose exact shape depends on the mass hierarchy. This measurement employs a different effect than the hierarchy determination in atmospheric oscillations, thereby providing a completely independent confirmation with different systematic effects. Even if neither of the two experiments achieves a conclusive significance on its own, the combination of both can vastly enhance the results [71]. This will be shown for the combination of PINGU with JUNO[14] [72, 73] in Sec. 6.5.

3.4.5 Current Status of Neutrino Mixing Parameters

Global fits to neutrino oscillation results from different experiments are available from various authors and usually constantly updated once new results are published. For this thesis, the best fit values and uncertainties from Fogli et al. released in 2012 [74] are used, which include the results on ϑ_{13} from Daya Bay and RENO that were published shortly before. There are more recent analyses available (e. g. [75, 76]), but the differences are small as no major results have been released in the meantime.

Since one of the neutrino mass splittings has a much smaller value than the others, the convention is to label the two mass eigenstates that are close to each other as m_1 and m_2, with $m_1 < m_2$ since the sign of the small splitting has been determined using solar neutrinos (cf. Sec. 3.3.1). The third eigenstate m_3 is then separated from the first two, either above—in the "normal" mass hierarchy (NH)—or below in the "inverted" hierarchy (IH). This is illustrated in Fig. 3.4.

[14] Initially named Daya Bay II.

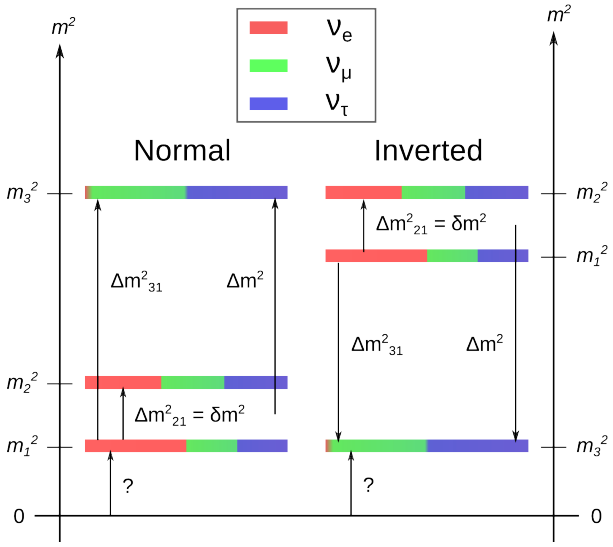

Figure 3.4: Schematic depiction of the ordering of neutrino mass eigenstates in both normal and inverted mass hierarchy. The definition of Δm^2 and δm^2 according to Fogli et al. [74] is indicated as well.

Since only two of the neutrino mass splittings are independent, Fogli et al. choose to do their analysis in terms of one small and one large mass splitting:

$$\delta m^2 \equiv m_2^2 - m_1^2 = \Delta m^2_{21} > 0 \tag{3.35}$$

$$\Delta m^2 \equiv m_3^2 - \left(m_2^2 + m_1^2\right)/2 = \left(\Delta m^2_{31} + \Delta m^2_{32}\right)/2 = \Delta m^2_{31} - \Delta m^2_{21}/2 \begin{cases} > 0 & \text{for NH} \\ < 0 & \text{for IH} \end{cases} \tag{3.36}$$

The values and 1σ ranges given in [74] for δm^2, Δm^2, $\sin^2(\vartheta_{12})$, $\sin^2(\vartheta_{23})$. and $\sin^2(\vartheta_{13})$ have to be converted to Δm^2_{21} and Δm^2_{31} and the values of the mixing angles in degrees in order to be processed by the oscillation code calculating the transition and survival probabilities. These values, which will be used as the *fiducial* oscillation parameters, are listed in Tab. 3.1. The CP violating phase δ_{CP} is set to zero.

Table 3.1: Fiducial values of the oscillation parameters, according to Fogli et al. [47], used throughout this thesis.

Parameter	Best Fit	1σ Range
Δm^2_{31} [10^{-3}eV^2]	2.46 (NH) -2.38 (IH)	0.08
Δm^2_{21} [10^{-5}eV^2]	7.54	0.24
ϑ_{12} [°]	33.6	1.1
ϑ_{23} [°]	38.6	1.3
ϑ_{13} [°]	8.93	0.47

As one can see, the main unknown is the sign of Δm^2_{31}, which will be assessed by PINGU. Another remaining question is the octant of ϑ_{23}, i. e. whether its value is below or above 45°. This unresolved as of now since most oscillation experiments cannot measure the angle directly, but rather $\sin^2(2\vartheta_{23})$,

which is symmetric about 45°. Yet PINGU has good sensitivity to the octant of ϑ_{23}, and in addition the significance of its hierarchy measurement is enhanced if $\vartheta_{23} > 45°$, as one will see in Sec. 6.3.2.

3.5 Mass Hierarchy Signature in PINGU

In the previous sections, all ingredients needed to understand the measurement of the neutrino mass hierarchy PINGU is supposed to perform have been discussed. These are

- the atmospheric flux of ν_e, ν_μ, and their antiparticles (Sec. 2.2.5);
- their probabilities to convert to another flavour via neutrino oscillations in matter, especially the occurrence of the MSW resonance and parametric enhancement (Sec. 3.3);
- and the cross-sections and peculiarities for the interaction of different neutrino flavours with the target material (Sec. 2.3).

The atmospheric neutrino flux as a starting point, follows a steeply falling power law spectrum with an index of $\gamma \approx -3.7$ in all flavours. The normalisation for ν_μ is about twice the ν_e normalisation, while neutrinos and antineutrinos of the same flavour have about the same flux. This flux has to be multiplied by the oscillation probabilities, which depend not only on the neutrinos' energies, but also on the zenith angle at which they arrive—this determines the distance travelled since their production in the Earth's atmosphere as well as the amount of matter traversed on their way through the Earth.

In Figs. 3.5 and 3.6, the oscillation probabilities for ν_e to ν_μ and ν_μ to ν_μ are shown as examples. They have been calculated using the AtmoWeights package that was developed by the IceCube collaboration [77] using the "Preliminary Reference Earth Model" (PREM) [55] for the Earth's density profile, shown in Fig. 3.2. It is easily recognisable that the oscillation probabilities for $\nu_\alpha \to \nu_\beta$ are in principle equal to those for $\bar{\nu}_\alpha \to \bar{\nu}_\beta$. Differences arise from the MSW resonance, easy to spot in the energy range from 2 to 10 GeV for zenith angles between $\cos\vartheta_{\text{zenith}} \approx -0.9$ and -0.4, that appears in neutrinos for the normal and in antineutrinos for the inverted mass hierarchy (cf. Sec. 3.3). In addition, parametric enhancements occur for neutrinos passing through the Earth's core, corresponding to the steepest zenith angles values below $\cos\vartheta_{\text{zenith}} \approx -0.9$. These change between neutrinos and antineutrinos in the same way as the MSW resonance, yet the latter covers a larger $\cos\vartheta_{\text{zenith}}$ range and hence is responsible for most of the mass hierarchy asymmetry.

Yet since the fluxes of neutrinos and antineutrinos of the same flavour are essentially equal and PINGU cannot discriminate between them[15], the question is how the different hierarchies can show up in the recorded data. Here the different cross-sections for neutrinos and antineutrinos come into effect. As one can see from Fig. 2.7, at the relevant energies just below 10 GeV the cross section for antineutrino interactions with a hadronic target material is about a factor of two lower than the one for neutrinos. Thus, the MSW resonance will appear in the data in any case, but much more prominent if the hierarchy is normal. In this case roughly ⅔ of the events—the neutrino-induced ones—are affected by the resonance, while in the inverted hierarchy case only the remaining third caused by antineutrinos is.

So the quantities of interest are the sum of neutrino and antineutrino events for the different flavours at the energy and $\cos\vartheta_{\text{zenith}}$ range where the MSW resonance is expected and how they differ assuming normal and inverted mass hierarchy. To assess how significant the difference is in a given bin in the (E,

[15] The only method to do this would be to identify the sign of the electrical charge of the lepton produced in CC interactions. However it is unrealistic to generate the required magnetic field in the antarctic glacier, where PINGU will be located (see Secs. 4.1 and 4.2).

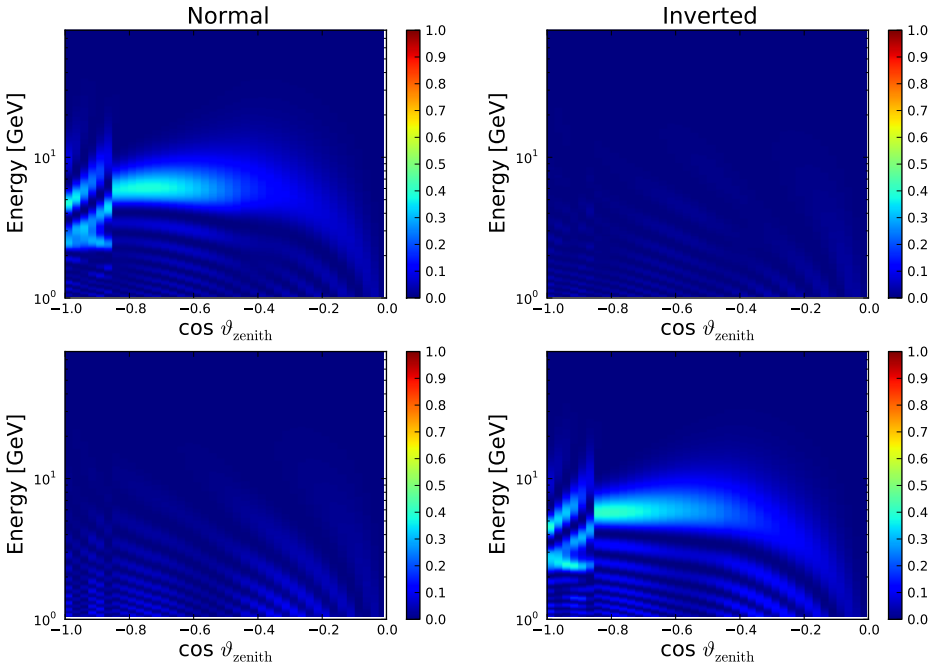

Figure 3.5: Oscillation probabilities for $\nu_e \rightarrow \nu_\mu$ (top) and $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ (bottom) for normal and inverted hierarchy.

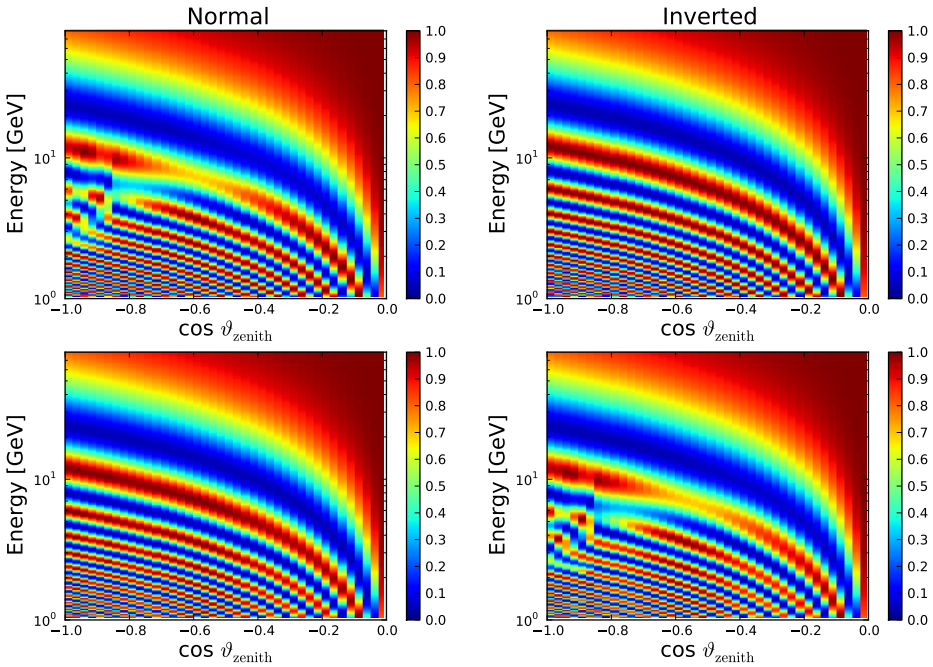

Figure 3.6: Oscillation probabilities for $\nu_\mu \rightarrow \nu_\mu$ (top) and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ (bottom) for normal and inverted hierarchy.

$\cos\vartheta_{\text{zenith}}$) plane, a bin-wise $\Delta\chi$ similar to [53] can be defined as

$$\Delta\chi = \frac{N_{\text{NH}} - N_{\text{IH}}}{\sqrt{N_{\text{NH}}}} \quad , \tag{3.37}$$

where N_{NH} and N_{IH} are the expected number of events in the normal and inverted hierarchy case, respectively. Plots of the event rates and their weighted difference $\Delta\chi$ are shown in Figs. 3.7 to 3.9. In this calculation not the bare cross-sections, but rather the *effective areas* for the different flavours enter, which also include the detection threshold and selection efficiency of the detector. This will be discussed in detail in Sec. 6.2.

Comparing the different flavours, one notes that the largest overall scale of the $\Delta\chi$ values appears in the ν_μ channel, but the contiguous regions of either positive or negative $\Delta\chi$ are rather small and alternate rapidly. Since a realistic detector has a limited resolution in reconstructing individual events, it will be challenging to resolve these fine structures in the data. In the ν_e channel, the features are less pronounced, but more extended than for ν_μ, offering a more robust measurement.

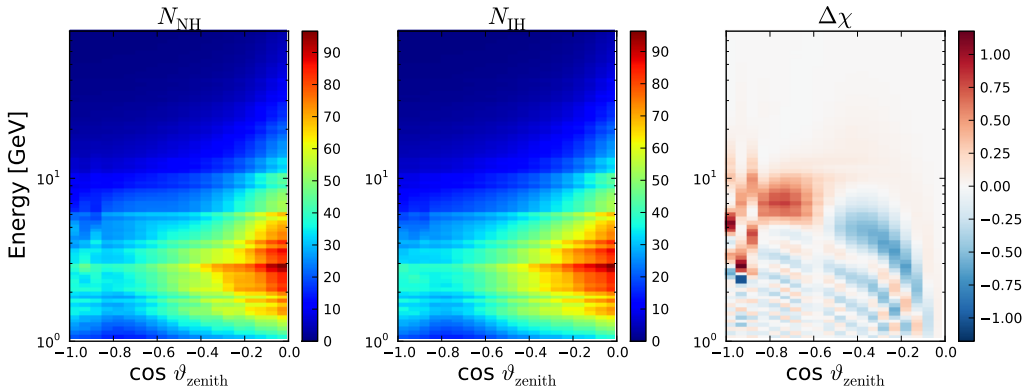

Figure 3.7: Expected $\nu_e + \bar{\nu}_e$ CC event rates in PINGU (arbitrary units) for normal and inverted mass hierarchy and their weighted difference $\Delta\chi$.

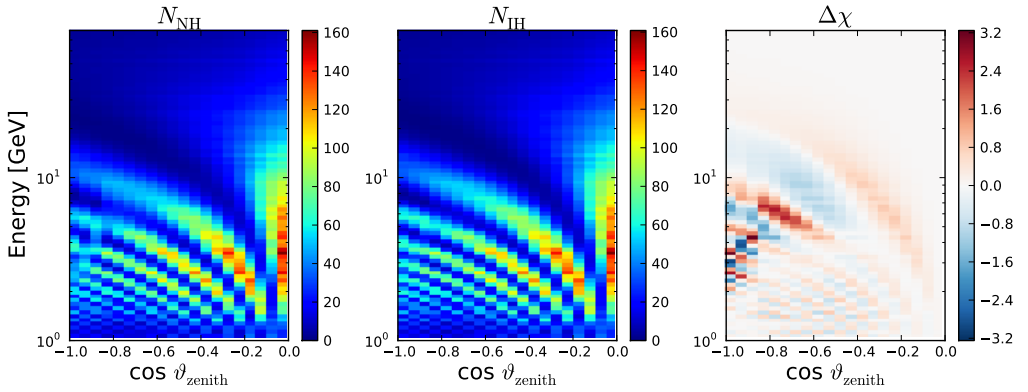

Figure 3.8: Same as Fig. 3.7, but for $\nu_\mu + \bar{\nu}_\mu$ events

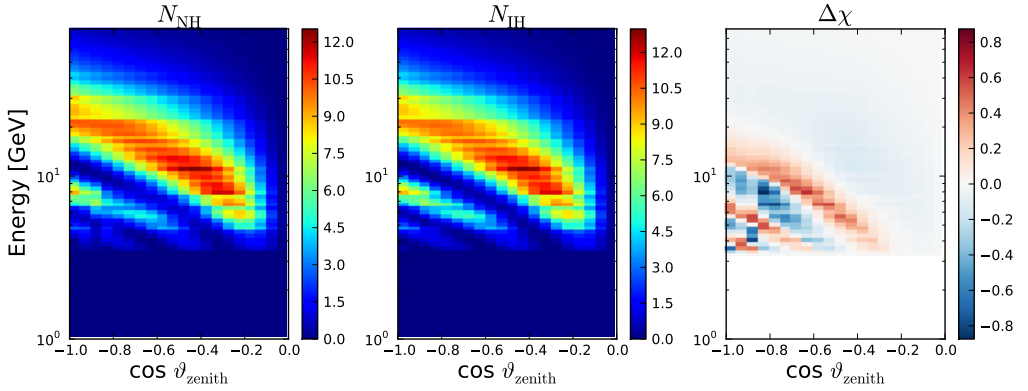

Figure 3.9: Same as Fig. 3.7, but for $\nu_\tau + \bar{\nu}_\tau$ events

CHAPTER 4

Detector

In this chapter, a detailed description of the proposed PINGU neutrino telescope and its predecessor IceCube will be given. First the concept of ice or water based neutrino telescopes based on the detection of Cherenkov radiation is introduced with IceCube/DeepCore as example, followed by a characterisation of its upcoming PINGU upgrade. Thereafter the selection and reconstruction of physics events for an analysis targeting the determination of the neutrino mass hierarchy will be described, and finally how conceptually new hardware might improve the results.

4.1 IceCube/DeepCore

4.1.1 Location

As already mentioned (Sec. 2.3), the natural choice for observing the low natural fluxes of high-energy neutrinos are water-based Cherenkov detectors. Although the basic requirement—a sufficiently large amount of water or ice—seems not very difficult to meet, there are additional constraints that have to be addressed as well:

Size: Depending on the energy range one is interested in, the size of the detector has to be adjusted accordingly. Since the atmospheric flux decreases rapidly with increasing energy while the size of the expected event signatures extend, one needs larger detector volumes to study higher fluxes. Roughly from the GeV scale upwards, the required dimensions are so big (several hundred metres) that artificial structures like the underground caverns of Kamiokande and Super-Kamiokande [16] are not feasible any more and one has to look for suitable natural locations.

Transparency: Since the detection of neutrinos is based on recording Cherenkov radiation, i. e. photons in the optical and near UV regime, obviously the chosen medium has to be transparent for these photons. Here ice has an advantage over fluid water as it has very low absorption down to wavelengths of 300 nm and below [78], while the fluid starts to absorb significantly below 400 nm [79].

Purity: The goal of all experiments is to reconstruct the neutrino events as accurately as possible. Therefore it is desirable to record a large number of unscattered photons, for which a very clear environment is needed. There might be, however, situations where scattering is eligible, e. g. when only the neutrino energy is of interest: then strong scattering keeps the photons inside the detector

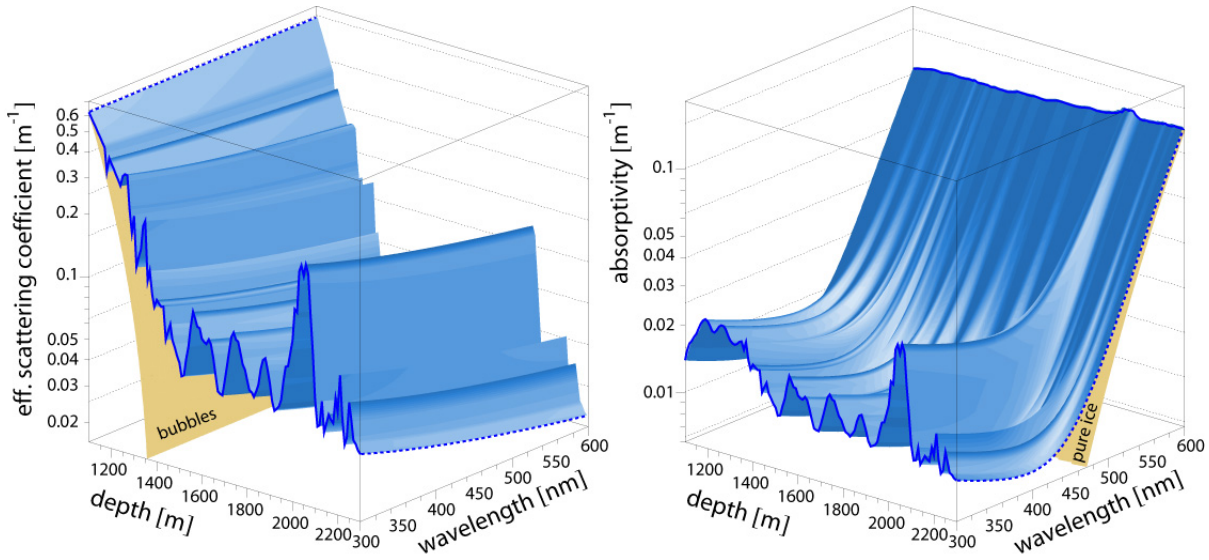

Figure 4.1: Effective scattering and absorption of light in the polar ice. Plot taken from [78].

for a longer time and hence increases the total number of detected photons, thereby improving the energy resolution. Absorption on the other hand is always a downside since it reduces the total number of photons and hence the information available for event reconstruction.

Shielding: In high-energy neutrino experiments, muons from atmospheric showers created by cosmic radiation (cf. Sec. 2.2.5) are a background whose rate is several orders of magnitude higher than the neutrino signal. In order to suppress those muons, detectors have to be placed deep underground so that there is a shielding of several hundred meters thickness, comparable to the penetration depth of $O(100\,\mathrm{GeV})$ muons.

Choosing an environment optimising all these factors, the IceCube neutrino observatory has been constructed in the deep glacial ice at the geographical South Pole in Antarctica. The Antarctic glacier with its thickness of $\approx 2500\,\mathrm{m}$ is a pristine environment of substantial size. In contrast to other natural resources like lakes or the deep sea, it inherently provides a solid holding structure for the instrumentation, is free of lifeforms that might disturb or even destroy the detector, and has a much lower content of radioactive ^{40}K than sea water. Especially at depths more than $\approx 2000\,\mathrm{m}$ below the surface, age and high pressure have facilitated the hydratisation of enclosed air bubbles, leaving an extremely clear ice with scattering and absorption lengths of several tens of metres even in the UV range, see Fig. 4.1. Instrumenting only the deepest ice below 1500 m guarantees a sufficient shielding of atmospheric muons.

The nearby Amundsen-Scott South Pole Station operated by the United States Antarctic Program provides the infrastructure needed for such a large scale experiment. This incorporates the supply of electrical power for the detector itself and the computing farm processing the raw data, satellite communications for transmitting science data, general technical support as well as accommodations for the visiting scientists.

4.1.2 Detector Geometry

A total of 86 strings, each instrumented with 60 Digital Optical Modules (DOMs, see Sec. 4.1.3), have been installed during IceCube's deployment phase from 2005 until December 18, 2010. A hot water

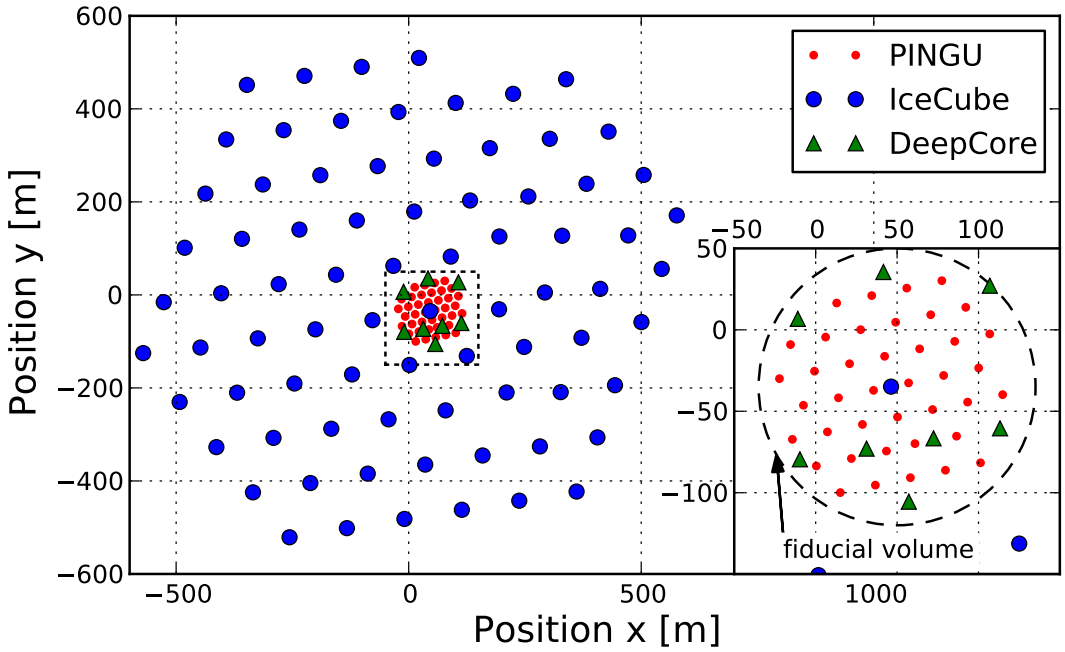

Figure 4.2: Top view of the IceCube string layout, including the DeepCore and planned PINGU (geometry V36) sub-arrays.

drill was used to melt holes of 60 cm diameter into the ice, reaching down to 2450 m, shortly above the underlying bedrock. Then the strings were lowered into the holes still filled with water which then refroze and now firmly encloses the strings.

Eighty of those strings form the hexagonal main array with an inter-string distance of 125 m, while the remaining six are placed at additional positions near the centre of the array, forming a dense sub-array called DeepCore with a string spacing of only 72 m, lowering the threshold energy for neutrino detection from $\approx$ 100 GeV to $\approx$ 20 GeV. A top view of the string layout is shown in Fig. 4.2.

To the lowest 1000 m of the main array strings, sixty DOMs each are attached, evenly spaced with a distance of 17 m. In DeepCore, the vertical DOM spacing is only 7 m below the dust layer found between 1950 m and 2100 m depth[1]. Above it, ten DOMs with 10 m spacing are attached to each DeepCore string to serve as a veto against atmospheric muons [80, 81].

4.1.3 Digital Optical Modules

The 5160 Digital Optical Modules [83] are the basic detection units of IceCube. They are attached to the string and connected to the main string cable during deployment and work autonomously except for the low voltage power supply. This modular design has the convenience that if one DOM fails to work and cannot be fixed as it is frozen in the deep ice, the others are not affected.

As shown in Figure 4.3, the DOM is housed by a borosilicate glass sphere of 13" diameter and 0.5" thickness to withstand the pressure arising from the refreezing water in the drill holes. This glass sphere contributes to the noise rate of about 540 Hz per DOM as it contains isotopes of the uranium

[1] The dust layer can be recognised easily in Fig. 4.1 as the region of increased scattering and absorption.

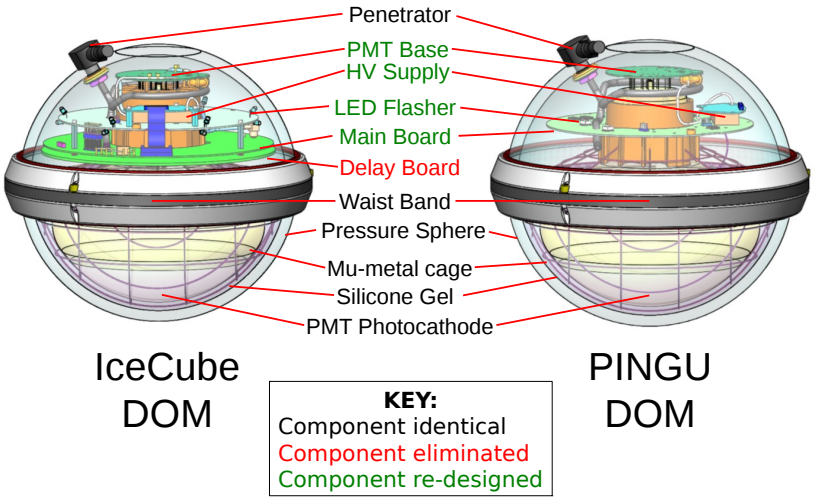

Figure 4.3: Comparing an IceCube/DeepCore DOM to the PDOM used in PINGU. Graphics taken from [82].

and thorium decay chains. The content of natural ^{40}K, that undergoes beta decay where the emitted β particle generates Cherenkov light, has been reduced in the glass.

The main component of the DOM is a 10" photomultiplier tube (Hamamatsu R7081-02 [84, 85]). In DeepCore, an improved version of this PMT with a 35 % higher quantum efficiency was used [81]. The PMT is oriented downwards as the main focus of IceCube is on extraterrestrial neutrinos from the northern hemisphere that have travelled through the Earth and hence arrive at the South Pole from below. The coupling to the glass sphere is provided by an optical gel which also provides protection and fixation to the PMT, while a surrounding Mu-metal grid guarantees the shielding of external magnetic fields.

The upper half of the DOM is filled with an high voltage divider that locally transforms the 96 V voltage that is provided by the string main cable into the high voltage of 1.3–1.5 kV that is needed to fuse the PMT, thus making it independent from possible voltage fluctuations of the pole station's power supply. Around the HV divider, the circular flasher board and DOM mainboard are mounted. The flasher board is populated with LEDs that can be used to produce a standardised signal for calibration. On the mainboard the electronics are located that are needed to read out the PMT signal and digitise the recorded data in situ, after which they are sent to the surface.

The digitisation is done with two di-channel analogue transient waveform digitisers (ATWDs) and one fast analogue-to-digital converter (fADC) operated in parallel. The ATWDs read out the PMT waveforms with three different gain factors (one channel is reserved for internal calibration) for maximum dynamic range at a sampling rate of 300 MHz with 10 bit resolution each, resulting in a total readout window of 422.4 ns with 128 samples. In order to enhance this readout time to a total of 6.4 µs, the fADC has a sampling rate of only 40 MHz with 256 samples at 10 bit resolution to also record late photon hits.

4.2 PINGU

PINGU, the Precision IceCube Next Generation Upgrade, is planned as a further infill to the IceCube/DeepCore array, lowering the energy threshold to few GeV [54]. The current baseline geometry, V36, consists of forty additional strings with 96 PINGU-DOMs (PDOMs, see below) each. In this lay-

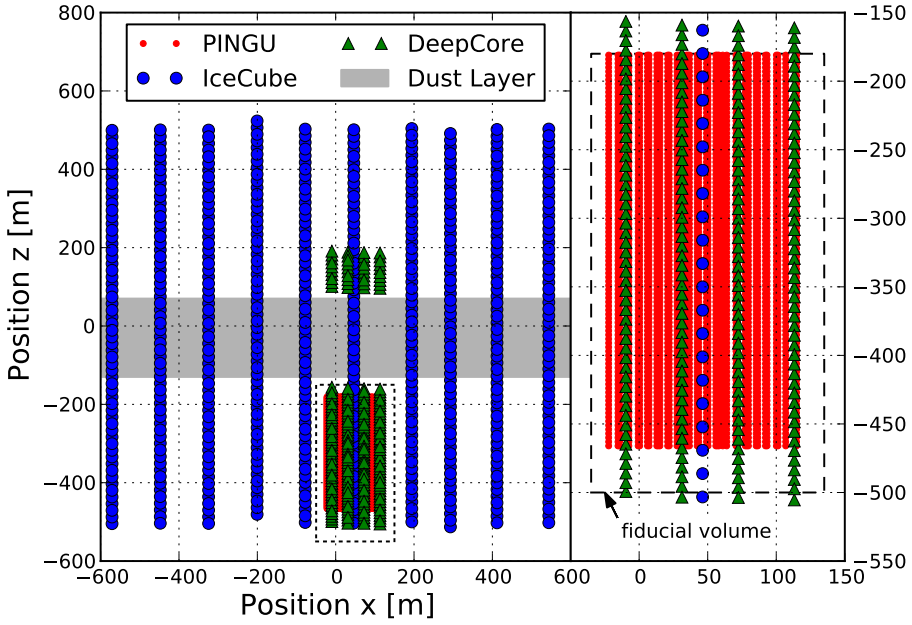

Figure 4.4: Side view of the IceCube string layout, including the DeepCore and planned PINGU (geometry V36) sub-arrays. The approximate position of the dust layer is shown for reference.

out, shown in Figs. 4.2 and 4.4, the string spacing is 22 m while the PDOMs are located at the lowest 300 m of IceCube with a vertical distance of 3 m. In this depth, the same as the main part of DeepCore, the ice has the best optical properties.

This detector geometry has been optimised to yield a maximum sensitivity for the neutrino mass hierarchy. It will be used for all studies described in Chapter 6 unless explicitly stated otherwise.

In terms of hardware and software infrastructure, many things can be adopted from IceCube, redesigning parts where a potential for improvement or simplification has been discovered. Several components of the electronics part of the module, like the main board, the PMT base, and the high voltage supply, have been redesigned or replaced by more recent versions. In particular the PMTs will all be the high quantum efficiency models already installed in DeepCore. Also the rather complicated digitisation unit with three components operating in parallel (cf. Sec. 4.1.3) will be replaced by only one ADC with high dynamic range that can be read out continuously [82]—a component that was not yet available at the construction of IceCube. A delay board has become obsolete with the more recent signal digitisers and will not be part of the PINGU DOMs. Additional components, such as cameras to monitor the freeze-in process, are in discussion.

Although the bulk of PDOMs are an improved version of the technology that has proven to work reliably, prototypes of novel optical modules will be deployed with PINGU as well. As those might be the baseline technology of future neutrino telescopes, they have to be tested under realistic conditions before they can be considered for large-scale use. Two options for these next-generation optical modules are described in Sec. 4.5, studies showing how they will impact the NMH determination are presented in Sec. 6.4.

Another thing to be improved upon in PINGU is the so-called “hole ice”. In IceCube it was observed that the refrozen ice directly around the strings that had been molten during deployment contains lots of

small air bubbles[2]. These inclusions lead to a dramatically reduced scattering length around the DOMs, significantly deteriorating the quality of event reconstruction that strongly profits from a large number of direct (i. e. unscattered) photons being recorded. In order to reduce the impact of the hole ice, the molten water will be degassed as a part of the drilling process for PINGU, thus reducing its air content and hence the number of air bubbles remaining after refreezing.

4.3 Event Reconstruction

As shown in Sec. 3.5, the neutrino mass hierarchy imprints fine features onto the the distribution of neutrino events recorded by PINGU. To resolve this pattern, it is crucial to reconstruct the events as exactly as possible. Although the interesting neutrino properties are only its energy and the zenith angle of its arrival direction[3] as those are determining the oscillation probability, an actual neutrino interaction event in PINGU has far more properties that need to be reconstructed as well. These are:

- The position of the interaction vertex in both space and time, giving rise to four variables in total (x, y, z, t).
- The energy of the hadronic cascade caused by the fragmented target nucleus, E_{cscd}.
- The direction of the hadronic cascade, described by its zenith and azimuth angles (ϑ, φ)
- The energy E_μ of the outgoing muon (if present), which can be replaced by its length l since muons can be considered as minimum ionising particles (see Sec. 2.3.2).
- The direction of the muon, represented by zenith and azimuth angles (ϑ_μ, φ_μ) as well.

This amounts to a total of ten variables (seven if no muon is created in the interaction) whose values have to be found in the reconstruction process. The neutrino energy is then given by the sum of cascade and muon energy, and its direction has to be calculated from their momenta using conservation of the total momentum.

4.3.1 Triggering

Before an event can be reconstructed, it has to trigger the detector to read out the signals of the optical modules. In the simulations that are the basis for this thesis, the preliminary PINGU trigger is a so-called SMT3, standing for triple simple multiplicity trigger, however this is likely to change in the actual implementation of PINGU [86].

This SMT3 trigger requires at least three DOMs registering a hard local coincidence (HLC) within a time window of 2.5 µs. A HLC means that two DOMs on the same string that are not more than two vertical positions (i. e. 6 m) apart have launched within 1 µs. For the PINGU trigger, all PINGU DOMs plus the DeepCore DOMs and the ones one the neighbouring IceCube strings are taken into account.

If an event satisfies the PINGU trigger, the full detector including IceCube and DeepCore is read out. This is necessary for background suppression, since dim atmospheric muons might trigger only PINGU, which has a lower threshold than IceCube[4], but still cause single sub-threshold hits in IceCube due to which it can be recognised and vetoed.

[2] The bulk glacial ice contains air as well, which has been integrated into the crystal structure of the ice. However this hydratisation needs thousands of years under high pressure, such that the air bubbles in the hole ice had not enough time to dissolve yet.

[3] Which determines the length of propagation since the creation in the Earth's atmosphere and the matter potential passed on the way.

[4] In IceCube, the standard trigger is SMT8 with a much larger DOM spacing.

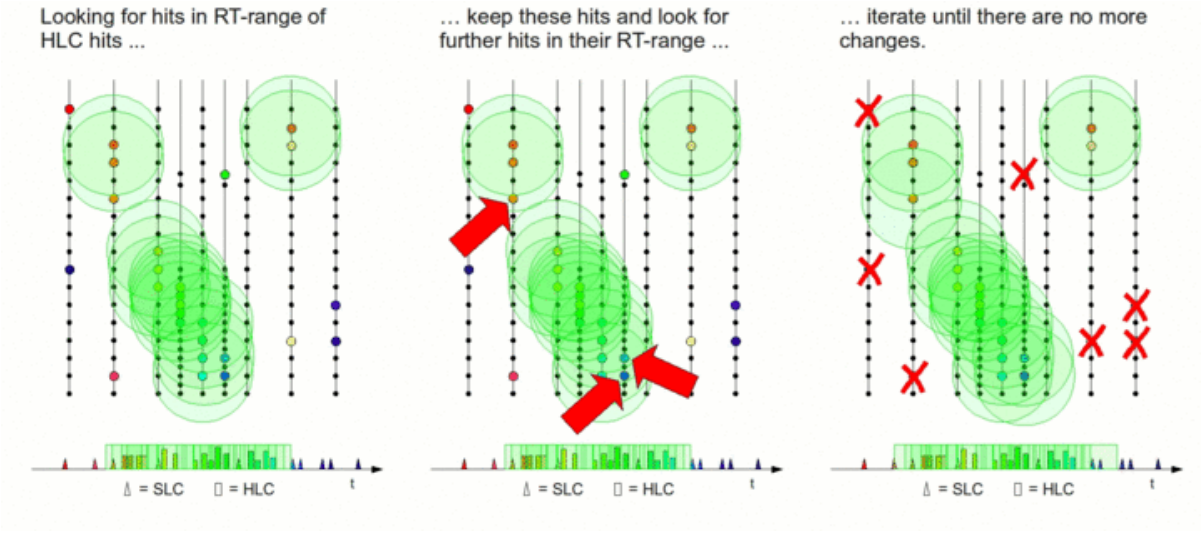

Figure 4.5: Illustration of the seeded RT cleaning algorithm (see text). Figure taken from [89].

4.3.2 Feature Extraction

After the DOMs have been read out, so-called feature extraction takes place. In this process, the raw digitised PMT waveforms that are the DOM output are deconvolved into a time series of single photon hits. As the digitisation unit for PINGU is not fully developed yet, the simulation falls back to the current IceCube technology, where this decomposition is done by the Wavedeform module [87] using the Lawson-Hanson least squares algorithm [88].

The result is a so-called pulse series in time of photon hits on DOMs, which then can be processed by the actual event reconstruction algorithms.

4.3.3 Noise Cleaning

The pulse map returned by Wavedeform still contains noise hits caused by electronic thermal noise or radioactive decays inside the DOM that hold no information on the neutrino event, but might confuse the reconstructions. Yet the majority of these hits can be removed by noise cleaning algorithms. In PINGU, two different cleanings are applied consecutively: first a seeded radius-time (SRT) cleaning, followed by a time-window (TW) cleaning.

SRT cleaning: Starting from the cluster of HLC hits that caused the event to trigger (see Sec. 4.3.1), the *seed*, the remaining hits are searched for ones that fall within a radius r and time window t w. r. t. one of the already selected. If such a hit is found, it is added to the cluster and the procedure is repeated. In PINGU, the boundaries are $r = 100\,\mathrm{m}$ and $t = 200\,\mathrm{ns}$, the same values commonly used in DeepCore. An illustration of the algorithm is shown in Fig. 4.5.

TW cleaning: A time window of $6\,\mu\mathrm{s}$ width is moved over the pulse series, searching for the position containing the largest amount of charge. Based on this position, all hits outside the window are discarded.

4.3.4 CLast

The start of the actual reconstruction chain is CLast [90], a first-guess algorithm for cascade, i. e. point-like events. With $\vec{r}_i$ being the positions of the DOMs that have registered charges a_i at times t_i in the event, first the event vertex is determined a the centre of gravity (COG) of the hit DOMs:

$$\vec{x} = \sum_i a_i \vec{r}_i \tag{4.1}$$

Following that, the time t of the vertex is calculated. For this, the residual time is defined as

$$\tau_i(t) = t_i - (t + d_i/c_{\text{ice}}) \quad , \tag{4.2}$$

where $d_i = |\vec{x} - \vec{r}_i|$ and $c_{\text{ice}} = c/n_{\text{ice}}$ is the speed of light in ice. After that, a direct hit is defined as a DOM hit with $0 < \tau_i < t_w$, with a trigger window of $t_w = 200\,\text{ns}$. Now for each DOM, a test vertex time is chosen to be $t_i - d_i/c_{\text{ice}}$ and the number of direct hits on all other DOMs is calculated. From all those test vertex times, the earliest one resulting in more than four direct hits on other DOMs is chosen as the vertex time t.

After the vertex position has been established, the cascade direction is determined from the "tensor of inertia" of the hit pattern. In a reference frame centred at the COG, $\vec{r}'_i = \vec{r} - \vec{x}$, this is calculated analogue to the quantity in classical mechanics with the mass replaced by the DOM charge:

$$I^{k,l} = \sum_i a_i \left(\delta^{kl} \vec{r}_i'^2 - r_i'^k r_i'^l\right) \tag{4.3}$$

The cascade direction is then given by the eigenvector corresponding to the smallest eigenvalue of the tensor of inertia, which points along the strongest elongation of the hit pattern.

Finally, a first estimate for the cascade energy is calculated using an empirical polynomial fit [91] to the number of hit DOMs in the event, N_{Ch}:

$$\log_{10}(E_{\text{cscd}}[\text{GeV}]) = -3.3 + x\,(9.2 + x\,(-9.7 + x\,(5.3 + x\,(-1.4 + x \cdot 0.134)))) \tag{4.4}$$

with $x = \log_{10}(N_{\text{Ch}})$.

4.3.5 Photonics

The Photonics software [92] is a tool widely used in IceCube to get a fast description of the propagation of Cherenkov photons in the ice. Its main functionality is to return the expected number of registered photons $B(\vec{r}, \vec{x}, t, \vartheta, \varphi)$ for any combination of source $(\vec{x}, t, \vartheta, \varphi)$ and sensor position $\vec{r}$.

Since photon propagation needs extensive computing, B cannot be directly simulated for each request. The number is retrieved from a table look-up instead. The respective tables have to be generated in advance by simulating a large number of sources in all possible depths—since the optical properties of the Antarctic ice sheet depend strongly on the vertical position, see Fig. 4.1—and all possible zenith angles, using data on the ice properties that have been retrieved from onboard LED flasher signals recorded by the IceCube DOMs [93]. The resulting photon density distribution is then stored in six-dimensional tables, the relevant quantities being three spatial dimensions, the relative time between emission and detection, the emission angle at the source, and the angle of incidence on the detector since the DOMs have a zenith dependent detection efficiency, as clearly visible from their layout (Fig. 4.3).

Finally these tables are fitted with multidimensional spline functions, which leads to smoother results and significantly reduces the amount of data that has to be stored [94]. This is an important issue, since

one of those six-dimensional tables has to be loaded into the CPU memory for every source depth, orientation, and type (point-like and track-like, as these have different light output distributions).

4.3.6 Monopod

The Monopod algorithm is a likelihood reconstruction for cascade-like events using detailed photon timing information. It is the single source case of the more general Millipede algorithm [95] that is used for high-energy track reconstruction, where it segments an extended track into a series of point-like energy depositions along the track.

Assuming a cascade-like energy deposition E at the vertex position $\vec{x}$ and time t, pointing into the direction (ϑ, φ), the expected number of photons $B_j^{\text{cscd}}(\vec{r}_i, \vec{x}, t, \vartheta, \varphi)$ per unit energy is retrieved from Photonics (see Sec. 4.3.5) for every DOM i that registered photons in the event and for every time bin j. The total number of expected photons is then

$$\mu_{ij} = B_j^{\text{cscd}}(\vec{r}_i, \vec{x}, t, \vartheta, \varphi)\, E_{\text{cscd}} + \nu_j \quad , \tag{4.5}$$

where ν_j is the expectation of noise hits.

From this, a Poisson likelihood is constructed according to

$$\mathcal{L} = \prod_{i}^{\#\text{ hit DOMs}} \prod_{j}^{\#\text{ time bins}} \frac{\mu_{ij}^{n_{ij}}}{\Gamma\left(n_{ij}+1\right)}\, e^{-\mu_{ij}} \quad , \tag{4.6}$$

with n_{ij} being the actual number of photons registered on DOM i in the time bin j. A fitting routine can now search for the values of $(\vec{x}, t, \vartheta, \varphi)$ that maximise $\mathcal{L}$, corresponding to the most probable values of the true cascade parameters.

Those fitting routines search for the maximum value of $\mathcal{L}$[5] in a multidimensional parameter space in an iterative process. In general, the procedure is to choose a test point in the parameter space, where the likelihood and its gradient vector are evaluated. After that, the test point is moved into the direction of the gradient by a distance depending on its magnitude. This procedure is repeated until the likelihood difference between two subsequent test points is below a predefined convergence value, indicating that the minimum has been found, or the maximum number of iterations has been reached.

In PINGU, Monopod is run with four iterations using the output of CLast as seed, i. e. the first test point. Yet obviously it cannot fully describe all of the expected events since its assumption about the event topology does not contain the possibility of an outgoing muon track. But since muons at PINGU energies of few GeV have a range at the order of ten or twenty metres, comparable to the string spacing, they distort the point-like event topology only slightly. Thus Monopod still gives a reliable estimate of the vertex position and direction, sufficiently precise for the first loose vertex containment cut (Sec. 4.4.1).

4.3.7 HybridReco/MultiNest

The final, most accurate and involved reconstruction currently existing for PINGU events is called HybridReco/MultiNest, or just MultiNest. It consists of two different ingredients: the HybridReco event hypothesis and the MultiNest minimiser.

[5] Usually, and computationally more conveniently, they search for the minimum of $-\log \mathcal{L}$. Hence they are commonly called "minimisers".

HybridReco

To conform to the actual event topology in PINGU, it is necessary to add the possibility of an outgoing muon of finite length to the pure cascade hypothesis that was used for the previous reconstructions. Since its starting point is fixed at the cascade vertex, the muon adds only three variables to the problem: its direction, again defined by two angles, and its length. However, at PINGU energies the kinematic angle between the hadronic cascade and the outgoing lepton is only a few degrees [96], far below the directional resolution of the reconstruction (see Sec. 6.2). Thus the directions of muon and cascade are assumed to be aligned to reduce complexity. In addition, this alignment will cause the reconstructed direction to be a weighted average of the cascade and track direction, which is likely to be closer to the true neutrino direction—which is the relevant quantity—than the two separate directions.

Hence, only the muon energy E_μ, which is proportional to its length as the muon is a minimum ionising particle, is added to the set of parameters. This means that the expression for the expected number of photons on a certain DOM (4.5) has to be extended by a term representing the muon track:

$$\mu_{ij} = B_j^{\text{cscd}}(\vec{r}_i, \vec{x}, t, \vartheta, \varphi)\, E_{\text{cscd}} + B_j^{\text{track}}(\vec{r}_i, \vec{x}, t, \vartheta, \varphi)\, E_\mu + \nu_j \quad , \tag{4.7}$$

Although only one variable is added with respect to the cascade reconstruction, it turns out that finding the maximum likelihood gets much more difficult than before. The reason for this is the fact that two different photon emission patterns, B^{cscd} for the hadronic cascade[6] and B^{track} for the muon, have to be considered in the hypothesis where it is unclear whether a given photon originates from the cascade or the muon. Additionally the typical distances in PINGU are so small that photon scattering plays only a small role and most of the photons reach the DOMs on an undisturbed path, so-called "direct hits" or "direct photons".

As an illustration, consider a photon that is detected close to the muon track in its direction of travel. Since the muon moves with the speed of light in vacuum, but the photons only with the speed of light in ice, $c_{\text{ice}} = c_{\text{vac}}/n_{\text{ice}}$, a direct photon from the muon would arrive earlier than one from the vertex, although it was emitted later. On the other hand, a photon whose arrival time is compatible with a direct hit from the vertex could also originate from the muon if it was scattered on the way and hence took a longer path to its detection. This leads to ambiguities in the photon distributions.

Even more importantly, the large number of direct hits causes strong discontinuities in the likelihood space, as it is very probable that a given photon arrives undisturbed at the earliest time allowed by causality, but forbidden to arrive even a short time earlier. These sharp edges are problematic for conventional minimisers as described in Sec. 4.3.6, which tend to "get stuck" in local minima of the likelihood space during their iterations.

MultiNest

A robust way to handle such a challenging likelihood profile and reliably find its global maximum is offered by the MultiNest algorithm [97, 98]. Instead of starting with a seed point and trying to move it to the maximum along a path following the direction of the local gradient, MultiNest uses an ensemble of active points distributed over an active region in the parameter space to infer the shape of the likelihood

[6] The value of B^{cscd} returned by Photonics always assumes a fully electromagnetic cascade, however the event signatures in PINGU always have a hadronic cascade that is overlaid with an electromagnetic one in $\nu_e/\bar{\nu}_e$ and $\nu_\tau/\bar{\nu}_\tau$ CC events. In HybridReco, the reconstructed energy is corrected to correspond to a fully hadronic cascade. Since hadronic cascades contain more neutral particles and hence have less light output per energy, this leads to a slight overestimation of the neutrino energy in $\nu_e/\bar{\nu}_e$ CC events. In $\nu_\tau/\bar{\nu}_\tau$ CC events, this effect is overcompensated by the energy carried away by the outgoing neutrino.

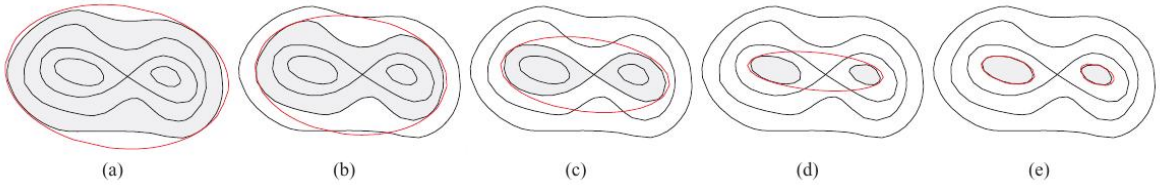

Figure 4.6: Cartoon illustrating the MultiNest sampling process. (a) – (d) show the shrinking sampling region with the iso-likelihood contours approximated by ellipses. In (e) the sampling region has been separated into two distinct active regions. Figure taken from [97].

and then tries to shrink this region down to the maximum.

The algorithm starts off with evaluating the likelihood values at the initial ensemble of points in parameter space. From this, an elliptical iso-likelihood contour is estimated from the points with the lowest likelihoods, as shown in Fig. 4.6a. Then the contour is scaled down by a certain factor, Fig. 4.6b, and the points which are now outside the contour are replaced by new points distributed evenly within the smaller contour. This procedure is repeated until the likelihood values of all active points are inside a sufficiently small range, indicating that the maximum has been found, or the number of iterations has reached the allowed bound.

A powerful feature of the MultiNest algorithm is that it can combine several overlapping ellipses into one iso-likelihood contour, and hence efficiently map out irregularly shaped maxima, and even identify sub-samples of active points belonging to different local maxima. If such sub-samples are detected, the active region is split into several distinct regions, as shown in Fig. 4.6e. Then the different regions are optimised in parallel, and if at some iteration it turns out that one of the local maxima is significantly lower than the others, the corresponding region gets removed again.

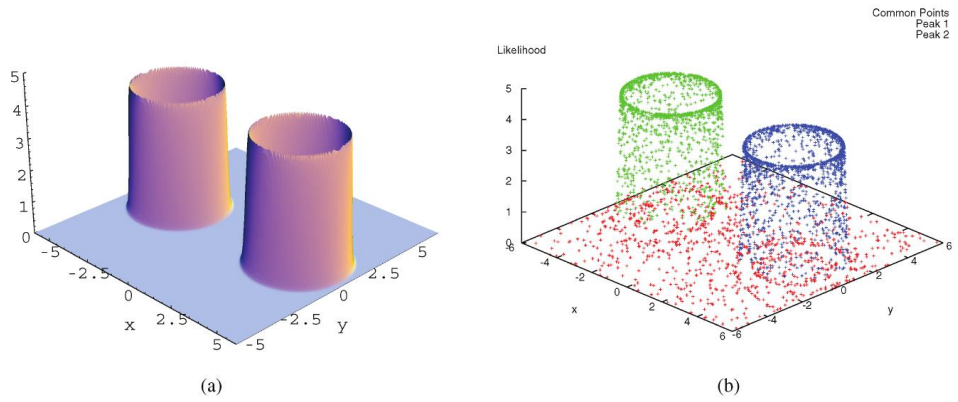

Figure 4.7: Example of MultiNest treating a likelihood landscape with two distinct, sharp maxima (a). In (b), the starting ensemble of test points is shown in red and the maximum likelihood points from the two isolated sub-samples evolving during successive iterations of MultiNest in green and blue, respectively. Figure taken from [97].

Due to these possibilities, MultiNest is suited to find the maximum likelihood solution of the HybridReco hypothesis for events in PINGU, where sharp and irregularly shaped maxima commonly occur. Also the ambiguities described above can be handled reliably via the technique of disjoint samples following two possible solutions of the problem in parallel. An example showing MultiNest operating on an artificial likelihood landscape with sharp edges and two separated maxima is shown in Fig. 4.7.

4.4 Event Selection

Since in PINGU the expected rate of neutrino events is on the order of 1 – 2 mHz, corresponding to about five neutrinos per hour, while background events, the vast majority caused by atmospheric muons, will be triggered at several kHz[7], an efficient background rejection algorithm is essential. At the same time, one has to make sure that a large fraction of the neutrinos carrying the neutrino mass hierarchy signal pass the imposed cuts—since it is only a second-order effect event statistics are of great importance—and that the signal in the data is not destroyed or changed by the cuts.

For the data samples used in this study, a two step selection has been applied [91]. The rationale for this approach is the efficient use of computing resources. At least the first part of the data reduction has to be run directly at the South Pole due to the limited bandwidth for data transmission to the northern hemisphere. Hence the computing for this part has to be lightweight, as computing resources in Antarctica are scarce as well. After the easily identifiable background events have been removed, more involved reconstructions can be run on the remaining data, whose results will be used for a refined second event selection and eventually in the analysis of the final data sample.

4.4.1 Step 1

The first set of cuts is based on the rather fast CLast and Monopod reconstructions and on the topology of the charge distribution in the detector itself. The cut variables have demonstrated their discrimination power in previous DeepCore analyses, in which atmospheric neutrino oscillations have already been measured at higher energy [62].

$z_{\mathrm{CLast}} < -200$ m: The very first cut requires the event vertex reconstructed by the CLast algorithm to be below – 200 m in IceCube coordinates, which is 20 m below the topmost PINGU DOMs. Since atmospheric muons enter the detector from above moving downwards, their vertices tend to be reconstructed at the highest possible location while the neutrinos that PINGU is interested in have travelled through the Earth and interact inside the detector volume, pointing upwards.

C2QR6 > 0.5: The next cut is the so-called "C2QR6" quantity being larger than 0.5. It is defined as the fraction of the total PMT charge in the event that has been accumulated within the first 600 ns, while the two very first hits are ignored. Muons are usually travelling through the detector on an extended path and generate Cherenkov light evenly along it, so that usually only a small fraction of it is recorded during the first 600 ns[8]. Neutrinos on the other hand deposit a large fraction of their energy in the hadronic cascade that only lasts for few ns.

$z_{\mathrm{travel}} > -30$ m: For the third cut, the mean spread in z of all hits is calculated relative to the mean depth of the first quartile of hits in the event. The cut is passed if this value is larger than – 30 m, meaning that the the topology of the event is not too strongly pointing downwards. Again this disfavours muons travelling through the detector from top to bottom while retaining neutrinos coming from below.

$t_{90\,\%} < 2\,\mu$s: This cut requires the time from the start of the event after which 90 % of the total charge has been accumulated to be less than 2 µs. Again the reasoning is that background muons deposit their energy evenly on their rather long path through the detector, while neutrinos are almost point-like sources in both location and time.

[7] In IceCube, the trigger rate is ≈ 3000 Hz.
[8] During this time, a muon propagates $ct \approx 180$ m.

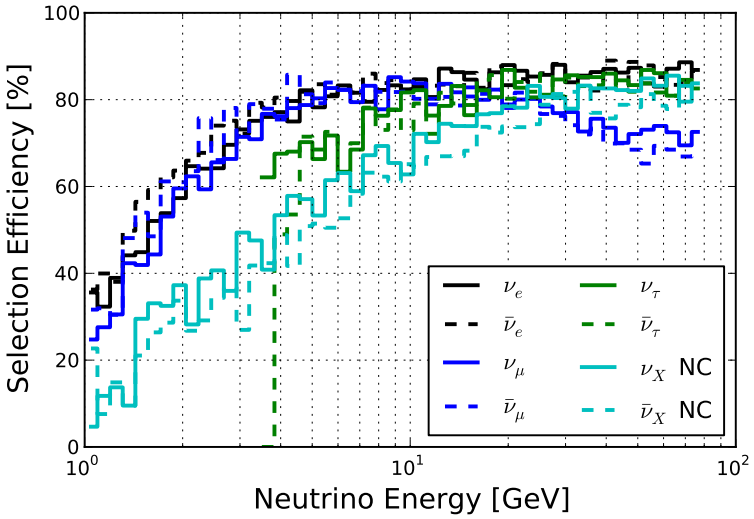

Figure 4.8: Selection efficiency for events of all interaction channels with contained true vertices.

$r_{Monopod} < 95$ m: Looking at the secondary Monopod reconstruction, the final cut at the first level secures that the event is contained in the detector volume not only vertically, but also horizontally. This horizontal containment is considered to be fulfilled if the reconstructed vertex is not farther than 95 m from the detector centre at (50 m, −35 m) in the *xy* direction. The containment assures that the event has the possibility to be successfully reconstructed by the MultiNest algorithm, which has to be run prior to the second cut level.

4.4.2 Step 2

The intention of the second set of cuts is not so much to reject background events, which is the focus of the first cut level, but rather to select well-reconstructed events whose quality is sufficient for the final analysis. Therefore, all the cuts in the second step are merely containment cuts, assuming any non-neutrino event fulfilling the strict containment has already been recognised and removed in the first step. The reference reconstruction for the following cuts is the MultiNest reconstruction (Sec. 4.3.7) that has been run on all events passing the first stage of cuts.

-500 m $< z_{vertex} < -180$ m: The upper boundary of the vertical containment, as shown in Fig. 4.4, has been loosened with respect to the $z_{vertex} < -200$ m cut in step one, where the stricter bound was necessary due to the strong downgoing background contamination. The lower boundary has been introduced since events originating from far below the instrumented volume generally have poor reconstruction quality: only photons that have travelled far through the ice and hence are subject to multiple scattering contribute to the signal recorded in the detector.

$r_{vertex} < 85$ m: The horizontal containment, shown in Fig. 4.2, has been tightened by 10 m compared to the first step. Since the reconstructed zenith angle of the neutrino is a key variable in the mass hierarchy analysis, one needs to make sure that this quantity is reconstructed as accurately as possible. This can not be guaranteed for events that are so close to the edge of the detector that they are recorded only partly.

$\vartheta_{\text{zenith}} \geq \pi/2$ **(optional):** For the mass hierarchy analysis, neutrinos that have undergone oscillations in a matter potential have to be selected. This means that only events that are reconstructed as upgoing and hence have passed through the Earth will be selected for the final data sample. However, the downgoing neutrino sample can still be used as a control region, e. g. for the normalisation of the atmospheric neutrino flux or to estimate the remaining background contamination at the final cut level.

The probability of passing all these cuts apart from the final directional cut is shown as a function of neutrino energy in Fig. 4.8 for events whose true interaction vertex is inside PINGU's fiducial volume as defined by the containment cuts above. ν_e and $\bar{\nu}_e$ events show the typical behaviour with a steady rise at low energies that reaches a plateau of $\approx 85\,\%$ efficiency slightly below 10 GeV. ν_μ and $\bar{\nu}_\mu$ events behave similarly at low energies, but above $\approx$ 20 GeV the selection efficiency starts to decrease again as the event topology becomes less and less point-like with increasing energy of the outgoing muon. The prominent muon track increases the event's probability to fail especially the cuts on C2QR6 and $t_{90\,\%}$ described in Sec. 4.4.1. For ν_τ and $\bar{\nu}_\tau$ events, the most prominent feature is the sharp cutoff at $\approx$ 3.5 GeV. This is of kinematic nature, as a tau lepton with a rest mass of 1.78 GeV [27] has to be produced in the interaction between a massless neutrino and a nucleon at rest. Due to baryon number conservation, a nucleon has to be present in the final state as well, which leads to the following threshold energy (assuming a proton in both initial and final state):

$$s = m_p^2 + 2E_\nu m_p \stackrel{!}{=} (m_\tau + m_p)^2 \tag{4.8}$$

$$\Rightarrow E_\nu = \frac{(m_\tau + m_p)^2 - m_p^2}{2m_p} = 3.46\,\text{GeV} \tag{4.9}$$

Neutral current events of all flavours finally show a behaviour similar to ν_e and $\bar{\nu}_e$ CC events as they correspond to point-like emissions of light as well, yet their turn-on until the efficiency plateau is reached is much slower due to the missing energy carried away by the outgoing neutrino.

4.4.3 Particle Flavour Identification

Since in PINGU the main oscillation channels are muon neutrinos converting into ν_e or ν_τ, which are all selected by the event selection strategy described above, it is necessary to reconstruct the flavours of the neutrino events in the final data sample. As shown in Sec. 2.3.2, however, PINGU can only distinguish between track-like events corresponding to ν_μ CC interactions with a muon in the final state and cascade-like events, made up from of ν_e and ν_τ CC events and NC events of all flavours, which lack the outgoing muon and consist of a cascade only, and possibly missing energy in form of an outgoing neutrino.

To make this discrimination as efficient as possible, a boosted decision tree [99] has been set up to generate a single cut variable from the combination of six input variables, all carrying information on the topology of the underlying event [100]. Three of these variables are based on the distribution of the reduced arrival times t_i^{red} of the DOM hits $(\vec{r}_i,\, t_i)$ in the event with respect to the reconstructed vertex $(\vec{x},\, t)$,

$$t_i^{\text{red}} = t_i - t - |\vec{r}_i - \vec{x}|/c \quad , \tag{4.10}$$

while the remaining three refer to the result of the MultiNest reconstruction:

Early hit charge: The number of DOM hits in the time window $[-200,\, -6]$ ns (i. e. before) relative to the reconstructed vertex time. Track-like events tend to have a larger number of these early hits

since the muon travelling through the detector and hence emitting light over a longer time causes a bias towards late vertex times. Then direct hits from the cascade can be registered prior to the reconstructed interaction vertex (cf. Sec. 4.3.7).

Very late hit charge: The number of DOM hits in the time window [200, 20000] ns relative to the reconstructed vertex time. Track-like events tend to have a larger number of very late hits as well, as the total light deposition is spread out over a longer time span due to the muon.

Time to accumulate 10 % of the total charge: In cascade-like events, light deposition happens at a more or less singular point in space and time, such that the first 10 % of the integrated PMT charge are accumulated faster than in spread-out track events.

Reconstructed muon length: The propagation length of the muon reconstructed by MultiNest before it has deposited all its energy and decays. Obviously, this value should be larger for true track events while cascades should ideally contain a muon of zero length.

Muon fractional energy: The fraction of the total reconstructed energy that is accounted to the muon, $\frac{E_\mu}{E_\mu+E_{\mathrm{cscd}}}$. Similar to the above, this is a measure of how prominent the muon features are in the event and is close to zero for cascades.

Best vs. cascade-only likelihood: For the best fit parameters found by MultiNest, the likelihood of the cascade-only hypothesis is evaluated by setting the muon energy to zero and scaling up the cascade energy accordingly. The larger the difference to the actual value of the best-fit likelihood including a muon, the more probable that a muon is contained in the event.

4.5 Next-Generation Optical Modules

Currently two different prototypes of optical modules are supposed to be deployed in PINGU. The first one, called WOM (Sec. 4.5.1), is a novel approach to enhance the photon collection efficiency by using passive components. The other one, called mDOM (Sec. 4.5.2), is an adaptation of the Km3NeT optical module [101] whose shape has been changed from spherical to cylindrical in order to fit into the holes drilled for the PINGU strings.

Both types of modules will be described in more detail below. In Sec. 6.4, the performance of a PINGU detector consisting fully of these next-generation modules will be investigated.

4.5.1 Wavelength-shifting Optical Module (WOM)

A description of the wavelength-shifting optical module including estimates of its photon detection capabilities has already been published as a conference contribution [103], with the author of this thesis as main author. Thus this section will closely follow said publication, with updated plots and figures reflecting recent findings.

The main idea of the WOM is to increase the sensitive area of a PMT by using passive components that act as light collectors and concentrators (a sketch of the module design is shown in Fig. 4.9b). Due to Liouville's theorem, mirrors or lenses can not be used for this purpose. However, by introducing a wavelength shift (at the cost of a short time delay), this limitation can be overcome.

The main design component are tubes that have wavelength-shifting properties. These tubes collect Cherenkov photons on their outer surface. The Cherenkov photons, which are mainly in the UV regime, are absorbed and then re-emitted isotropically at a larger wavelength within the tube. The isotropic

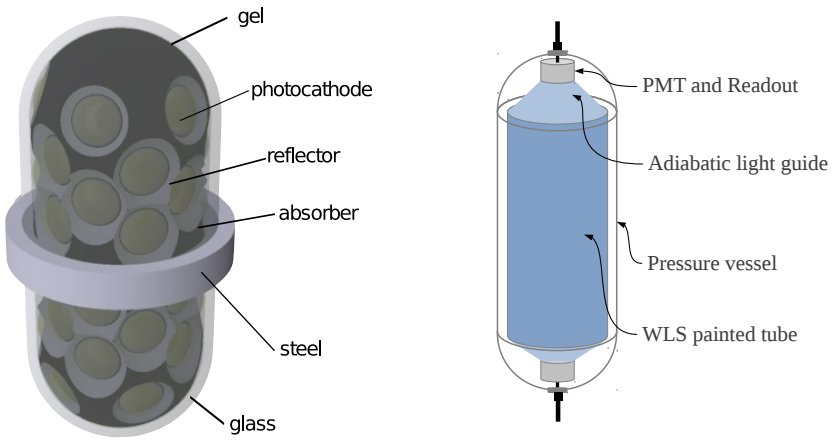

(a) (b)

Figure 4.9: The (a) mDOM and (b) WOM module concepts. Graphics taken from [102] and [103], respectively.

emission ensures that a large fraction of the photons, which were incident roughly perpendicular to the surface, will now be captured inside the tube and then guided towards the end via multiple total internal reflection.

At both ends of the tube, small high-efficiency PMTs will be placed that read out the incoming photons. Since their spectrum has been shifted away from UV towards the optical blue, it is now better suited for readout by conventional PMTs as those are usually most sensitive in the optical blue and green.

The whole assembly will be enclosed by a transparent pressure vessel that protects the components from physical damage. It will also ensure that the wavelength-shifting tube is surrounded by a small gap of air and not in direct optical contact with the outside (e. g. the glacial ice). Although an optical contact to the surroundings usually is desirable to gather as many photons as possible, in this case a large difference in refractive index is needed to achieve efficient total internal reflection.

The three most important components of the WOM, for which details of their properties will be discussed, are:

- the pressure housing
- the PMT
- and the wavelength-shifting tube.

For the pressure housing, apart from mechanical stability[9], two requirements have to be met. Firstly, optical transparency has to be guaranteed not only in the optical but also in the UV regime down to $\approx 250\,\text{nm}$. Secondly, the material needs to have a high radio purity. In regular glass, ^{40}K is a common contamination and its decay would give rise to a very high module noise rate. It turns out that fused quartz glass is a very good choice, having a transparency of 90 % at 250 nm and an activity of only 0.02 Bq/kg [104].

A possible choice for the PMT is a prototype manufactured by the Hamamatsu corporation that has a novel photocathode with enhanced green sensitivity [105]. Although for this prototype no noise mea-

[9] Static pressure of up to 10000 psi can occur during the freeze-in process when deploying in deep Antarctic ice.

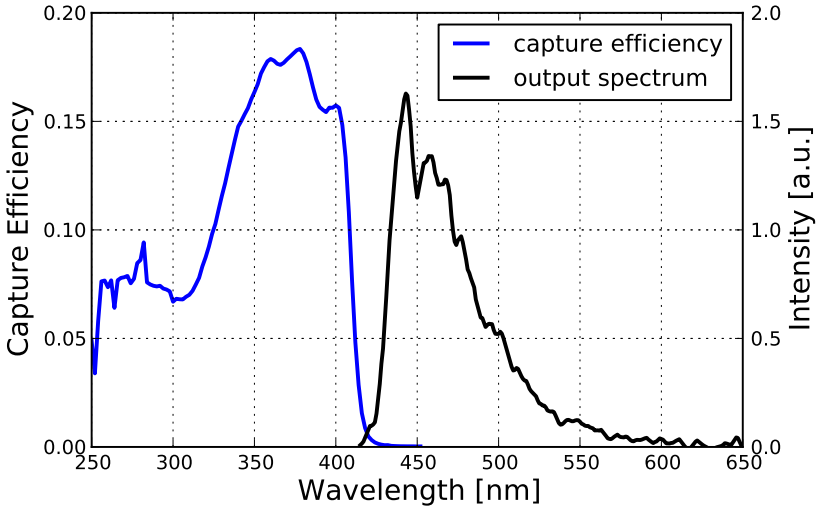

Figure 4.10: Capture efficiency and output spectrum of a 20 mm diameter fused quartz glass tube coated with EJ-298 wavelength-shifting paint.

surements at low temperature are available, for similar models typical noise rates are reported to be below 1 Hz per cm^2 of cathode area at a temperature of −30 °C [106].

The wavelength-shifting tube can be made of quartz glass as well although the wall thickness can be much thinner than for the outer cylinder, since there is no pressure inside. This inner cylinder will be covered by a thin film of wavelength-shifting material. First tests have been done with samples produced by dip-coating 20 mm and 5 mm diameter glass tubes with EJ-298 wavelength-shifting paint, which consists of a fluorescent dopant and a PVT base [107]. The film thickness is ≈ 50 µm, its surface roughness RMS has been measured with an AFM to be 3 nm. The wavelength shifter itself is an organic material with an inherent activity of 0.4 Bq/kg [108], yet its contribution is negligible due to the thinness of the layer. Hence the noise rate of this component is dominated by the quartz glass tube.

To calculate the performance of a fully assembled WOM, the efficiencies of the individual parts have to be combined:

Photon capture efficiency The photon capture efficiency (CE), is defined as

$$CE(\lambda) = \frac{\sharp \text{ detectable photons}}{\sharp \text{ photons injected at outer surface}} \quad . \tag{4.11}$$

This quantity has been measured in a lab setup using a monochromator with wavelength range 250 – 1200 nm and two identical photo-sensors, read out via lock-in amplifiers. Details of the experimental setup for this measurement and the extraction of the capture efficiency are given in App. A.1 and A.2, respectively. The emission spectrum and capture efficiency as a function of wavelength are shown in Fig. 4.10 for the 20 mm sample described in the previous section. For the 5 mm sample, the results are identical within the precision of the measurement.

The peak efficiency of ≈ 23 % has to be compared to the theoretical maximum, which, assuming that all incoming photons are absorbed by the active dye and then re-emitted isotropically with 100 % quantum efficiency, and ignoring any transport loss, is given by the fraction of solid angle

where total internal reflection occurs:

$$CE_{\max} = \cos(\vartheta_c) = \cos\left(\arcsin\left(\frac{n_{\text{air}}}{n_{\text{glass}}}\right)\right) \approx 75\,\% \quad . \tag{4.12}$$

The decrease in efficiency down to $\approx 8\,\%$ below 350 nm is assumed to be due to re-absorption and/or UV absorption inside the PVT base of the paint. This should be reducible by thinning the paint layer and future R&D for the dye. An improved production scheme allowing for thinner layers is under development, and first studies with custom-made paint mixtures combining several dyes that are active in different wavelength regimes have been done [109].

PMT readout efficiency To account for the fact that the captured photons coming out at either end of the wavelength-shifting glass tube will be read out by PMTs, its (normalised) output spectrum (also shown in Fig. 4.10) has to be convolved with the quantum efficiency of a realistic PMT. This gives the fraction ε_{PMT} of photons coming out of the wavelength-shifting tube that are actually detected by the PMT. Using the data provided by Hamamatsu for their R7600-EG prototype [105], the readout efficiency can be calculated to be $\varepsilon_{\text{PMT}} = 30.8\,\%$.

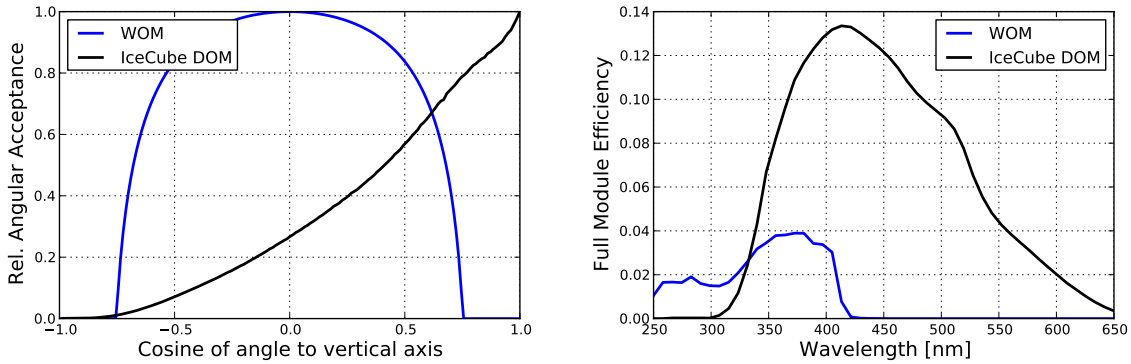

(a) (b)

Figure 4.11: Relative angular acceptance (a) and full module photo-detection efficiency at optimal illumination angle (b) of WOM and IceCube DOM.

Angular efficiency The angular acceptance ε_Ω of the WOM can be calculated with Fresnel's equations. For this a homogeneous, parallel bundle of light incident onto the WOM at an angle ϑ w. r. t. its longitudinal (i. e. vertical) axis is assumed. The acceptance at this angle can be calculated by integrating the transmission probability from ice ($n = 1.33$) through glass ($n = 1.48$) into air ($n = 1.0$) over the whole visible surface of the WOM, including the fact that the WOM appears shortened by a factor of $\sin(\vartheta)$ at angles different from 90°. This calculation is carried out in detail in App. A.3. Note that ε_Ω only covers the propagation of photons from the surrounding ice into the air gap, since the penetration from air into the actual wavelength-shifting tube is included in the measured capture efficiency described above.

In Fig. 4.11a the relative acceptance is shown w. r. t. the maximum for the WOM and an IceCube Digital Optical Module (DOM), the photosensor module used in IceCube [93]. For the WOM, a maximum acceptance $\varepsilon_\Omega(\max) = 70.1\,\%$ is achieved at $\vartheta = 90°$. Averaging the relative angular

acceptance shown in Fig. 4.11a over all angles of incidence, the mean angular acceptances are $\bar{\varepsilon}_\Omega(\mathrm{WOM}) = 57.5\,\%$ and $\bar{\varepsilon}_\Omega(\mathrm{DOM}) = 34.1\,\%$, respectively.

To obtain the detection efficiency of the fully assembled module, the three efficiencies discussed above, CE, $\varepsilon_{\mathrm{PMT}}$, and ε_Ω, have to be multiplied:

$$ME(\lambda) = \varepsilon_\Omega(\max) \cdot \varepsilon_{\mathrm{PMT}} \cdot CE(\lambda) \quad . \tag{4.13}$$

The result, assuming optimal illumination (i.e. $\vartheta = 90°$), is shown in Fig. 4.11b, compared to the efficiency of an IceCube DOM at its optimal illumination angle ($\vartheta = 0°$).

Comparing the module efficiencies of WOM and IceCube DOM in Figs. 4.11a and 4.11b, at first sight the WOM does not provide an improvement in photon collection efficiency. However, one has to keep two things in mind:

First, the WOM can be produced in a very large size since its size-critical components are just two glass tubes (pressure vessel and WLS painted inner tube), that are easily scalable. The size of the DOM on the other hand is given by the size of the enclosed PMT which can hardly be pushed much further at reasonable cost[10].

Second, the WOM is more sensitive in the UV below $\approx 370\,\mathrm{nm}$. As the photons that are to be detected stem from Cherenkov radiation, their spectrum is proportional to $1/\lambda^2$. Accounting for this initial spectrum when calculating the mean module efficiency

$$\overline{ME} = \frac{\int_{250\,\mathrm{nm}}^{600\,\mathrm{nm}} ME(\lambda)\frac{\mathrm{d}\lambda}{\lambda^2}}{\int_{250\,\mathrm{nm}}^{600\,\mathrm{nm}} \frac{\mathrm{d}\lambda}{\lambda^2}} \quad , \tag{4.14}$$

the performance of the WOM improves significantly w. r. t. the IceCube DOM.

Assuming a WOM with a diameter of 20 cm and a length of 2 m as above and, for comparison, a DOM with 30 cm diameter [83], one can calculate the full effective area of WOM and DOM for an isotropic Cherenkov spectrum between 250 and 600 nm. The results are shown in Tab. 4.1.

Table 4.1: Comparison of WOM and IceCube DOM properties for a Cherenkov spectrum between 250 and 600 nm.

Module	$\overline{ME}$	$\bar{\varepsilon}_\Omega$	Eff. Area	Noise
WOM	1.57 %	57.5 %	36.1 cm^2	≈ 10 Hz
DOM	5.36 %	34.1 %	12.9 cm^2	800 Hz

Hence in terms of total effective area,

$$A_{\mathrm{eff}} = \overline{ME} \cdot \bar{\varepsilon}_\Omega \cdot A_{\mathrm{xsec}} \quad , \tag{4.15}$$

one WOM with a cross-section of $A_{\mathrm{xsec}} = 2RL = 4000\,\mathrm{cm}^2$ is the equivalent of about three IceCube DOMs ($A_{\mathrm{xsec}} = 707\,\mathrm{cm}^2$).

In summary, the WOM concept offers a technological solution to obtain a large area single photon sensor with exceptionally low noise rate. First measurements with off-the-shelf components already

[10] This is mostly due to the deployment in the antarctic glacier: Larger spherical DOMs need larger diameter drill holes, dramatically increasing the costs for drilling.

show promising results, yet the photon sensitivity of the module should be improvable especially in the UV range where the Cherenkov spectrum is peaked. This can be achieved by adding UV-active dyes to the paint mixture and improving the overall sample quality, avoiding the loss of already captured photons. Thus, future work will concentrate on the optimisation of the film coating, the choice of dyes, as well as on the proof of long term stability and large scale expandability. In the long term, the WOM technology is expected to enhance the number of detected photons per event by a factor of a few compared to the DOMs used in IceCube and PINGU.

4.5.2 Multi-PMT Optical Module (mDOM)

Comparing the mDOM, short for Multi-PMT Optical Module, to the standard PINGU DOMs, the main difference is that instead of one single large PMT, a total of 41 small PMTs with 3" diameter will be used for photon detection, see Fig. 4.9a. The advantages from this layout are that the angular acceptance covers almost every direction, in contrast to the downwards-pointing single PMT that has no sensitivity for photons arriving from above [102].

In addition, the use of more than one PMT per module allows for a very effective noise reduction. If one only counts module hits where at least two different PMTs on the same module have registered a photon within a very short coincidence time, the most important sources of module noise can be strongly suppressed: Both radioactive decays inside the PMT glass housing and random electronic noise are restricted to one PMT at a time and hence vetoed with close to 100 % efficiency.

CHAPTER 5

Simulation

As this thesis is about estimating the performance of the PINGU detector before it is actually built, the estimate has to be based on simulations. In this chapter, the simulations used to generate the results reported in Chapter 6 will be described in detail.

The chapter, as the simulation process, is divided into two sections. In Sec. 5.1 the existing and well-established IceCube Monte Carlo (MC) chain will be discussed, which has been adopted for PINGU simulations. Here, individual neutrino events are generated and their output of Cherenkov photons is modelled. After propagating the photons through the ice, the resulting pattern of triggered optical modules is processed through the standard reconstruction and event selection specified in Secs. 4.3 and 4.4, respectively. The outcome of this event-by-event MC, i. e. the effective areas, reconstruction resolutions, and particle identification efficiencies for all neutrino flavours, are then used as input for the second part of the detector simulation.

The Parametric PINGU Analysis, `PaPA` in short, was written specifically for the rapid analysis of PINGU's neutrino mass hierarchy sensitivity including a variety of systematic parameters. Since propagating these through the full MC chain would be too time-consuming, an effective detector simulation was implemented. Instead of operating on individual events, the expected event distributions are generated directly, based on the detector performance retrieved from the MC data. `PaPA` will be described in detail in Sec. 5.2.

5.1 The IceCube/PINGU Simulation Chain

5.1.1 Event Generation

The first step in the MC chain is to model the interaction of an incoming neutrino with a target nucleus in the ice and the resulting final state, the so-called event generation. In the dedicated PINGU MC, this is carried out using the GENIE (Generates Events for Neutrino Interaction Experiments [110]) software package. This is already the first modification of the standard IceCube MC chain, where NuGen [111], an IceCube-specific neutrino generator, is the default. NuGen is laid out for high-energy neutrino events where only deep inelastic scattering has to be considered as an interaction process. In PINGU, however, the low GeV energy range contains the interesting signal for an oscillation measurement, and here the complex interplay between quasi-elastic and deep-inelastic scattering as well as resonant processes has to be taken into account (see Sec. 2.3.2). Since GENIE puts much effort into modelling especially this

energy range with great care and validating it against experimental results, it is the natural choice for generating PINGU events.

GENIE starts off with an isotropic flux of neutrinos of a given flavour following a user-defined power-law distribution in energy (usually $\propto E^{-1}$ or E^{-2} for PINGU MC [112]) on the surface of a cylindrical generation volume well encompassing the full IceCube detector. Any generated neutrino passing through the interaction volume, which is fully inside the generation volume but still contains the detector as a whole, is forced to interact inside this volume. The interaction type is chosen randomly from the ones that are allowed and the event is assigned a weight $\mathcal{W}_i$ proportional to the particular interaction probability, taking into account the generated energy spectrum. This weighting strategy makes it possible to re-weight the generated events to any desired incoming flux $\Phi(E, \vartheta, \varphi)$ later on. Then the actual weight is simply given by

$$w_i = \frac{\Phi(E_i, \vartheta_i, \varphi_i)\,\mathcal{W}_i}{N_{\text{evts}}} \quad , \tag{5.1}$$

where N_{evts} is the total number of simulated events.

After the interaction mechanism has been determined, the interaction itself is modelled in detail and all involved particles, from the initial neutrino and nucleus over possible intermediate states to the final (meta-)stable particles like pions or muons, are stored inside an `I3MCTree` object for further processing. The reference to a tree comes from the fact that this object has the structure of a multiply nested list, where every particle is the root of a sub-tree (or branch) holding the particles created in its decay. The particles are characterised by their identities, positions, four-momenta, and state (such as 'initial', 'intermediate', or 'final'). Additional GENIE-specific information such as the number of generated events, N_{evts}, the size of the interaction volume, and others, are kept as an `I3MCWeightDict` object.

5.1.2 Particle Propagation

The `I3MCTree` generated by GENIE is handed off to the mmc module [113], which propagates the final state particles in the tree as well as possible secondaries created in their decay through the ice until they have deposited all their energy. The particles involved here are mainly produced in the electromagnetic and hadronic showers, as discussed in Sec. 2.3.2. These are electrons, photons, and pions as well as muons and taus. The `I3MCTree` is extended with the outcome of mmc, additional information being stored as `MMCTrackList` and passed on to another module called clsim [114].

clsim is then used to generate the Cherenkov photons produced by the particles propagating through the ice. Therefore, every particle is converted into a series of steps of constant velocity $\beta = v/c$, over which Cherenkov photons are emitted according to (2.29). Usually this process of photon generation is handled by the Geant4 package [115, 116], but can also be done using an effective parametrisation as well.

These photons are then propagated through the ice until they either get absorbed or hit a DOM. Since photon propagation is a process that can well be run in parallel by multiple computation threads, clsim uses the publicly available OpenCL library [117] to outsource the calculations to GPUs, resulting in a significant speedup compared to a simulation on CPUs. The photons that have hit and launched a DOM, taking into account its angle- and wavelength-dependent quantum efficiency (see Fig. 5.1), are stored in a `I3MCHitSeriesMap`. This object contains the registered photons and the corresponding DOM IDs; other information such as the parent particle, wavelength, and position and angle of incidence on the hit DOM can be stored as well. This information gets passed on to emulate the response of the actual detector.

5.1.3 Detector Response

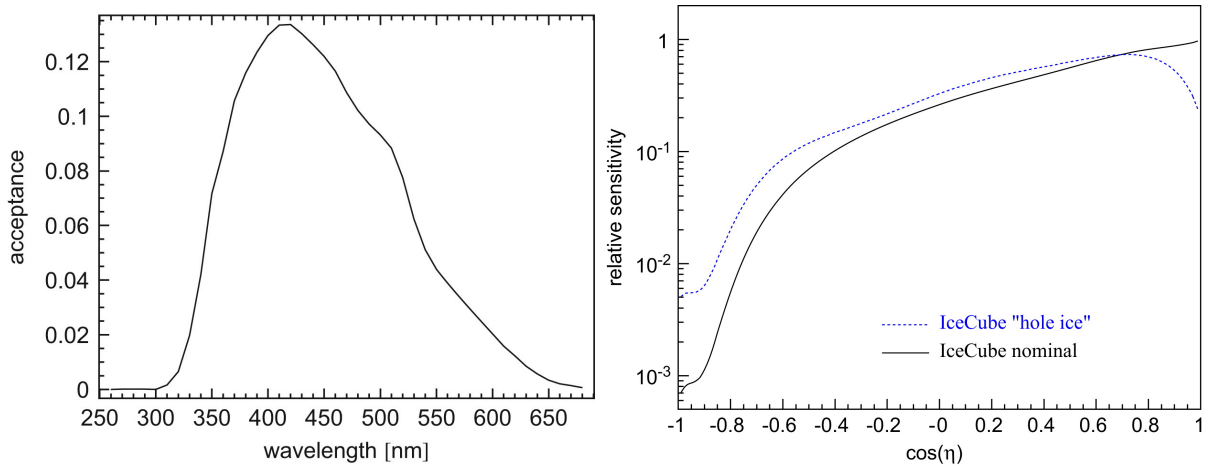

(a) (b)

Figure 5.1: (a) Quantum efficiency of the IceCube DOM as a function of the photon wavelength for head-on illumination. (b) Normalised angular dependence of the acceptance for a "bare" DOM and a DOM inside "hole ice", with η denoting the angle towards the centre of the PMT front. Plots taken from [93].

In PINGU simulations, all DOMs are represented by an identical copy of the standard DeepCore DOM, having a 35 % higher quantum efficiency than the IceCube DOM (see Sec. 4.1.3). This is only an approximation of the actual PINGU DOM, which currently only exists as a prototype, however it has already become clear that especially the digitisation process will be simplified. Yet as the PDOM design is not finalised, the DeepCore DOM is the closest approximation at hand.

Before the response of the DOMs gets evaluated, noise hits from both thermal electronic noise as well as radioactive decays inside the DOM and the accompanying scintillation and fluorescence light are added to the `HitSeries` with the vuvuzela module [118]. Then the DOMLauncher module [119] is called to generate the actual DOM output.

The DOMLauncher first calls the PMTResponseSimulator submodule [120] to convert the single photo-electron produced at the PMT cathode by a DOM hit to a charge pulse entering the DOM electronics. Here, the PMT transit time jitter is applied, which takes into account that there is a spread in the time needed by the electron avalanche developing on the PMT dynodes to pass through all amplification stages. This distribution is shown in Fig. 5.2a. The distribution of the amount of charge generated by a single photo-electron is dominated by a Gaussian, per construction centred at the charge equivalent of one photo-electron, but also contains a exponentially decreasing component of small-amplitude pulses, as shown in Fig. 5.2b.

Once the main photon pulse has been processed, secondary effects like pre-, late, and after-pulses are added. These result from photons hitting the first dynode instead of the photocathode, scattered avalanche electrons hitting the same dynode twice, and ionised residual gas atoms drifting onto the photocathode, respectively, and are offset by a specific time window from the main bunch of photo-electrons, but causally connected. Finally, saturation effects are taken into account, which have to be considered for events of very high energy or with an interaction vertex in the close proximity of a single DOM.

The full PMT charge output as a function of time, or "waveform", is then passed to the main DOMLauncher module, which simulates the processing and digitisation of the raw PMT signal on the DOM

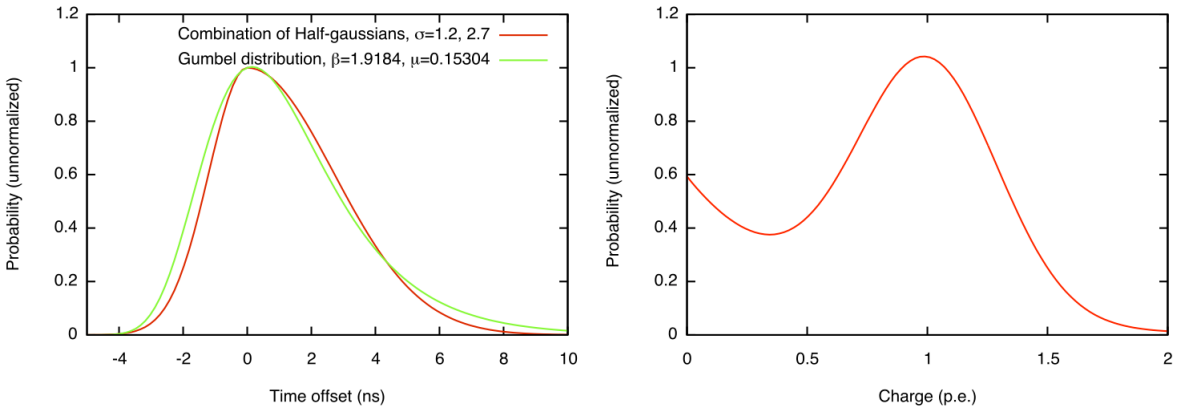

(a) (b)

Figure 5.2: Parametrisation of (a) the PMT transit time jitter distribution (in green) and (b) the single photoelectron charge distribution as used by the `PMTResponseSimulator`. Plots taken from [120].

main board [119]. First a discriminator threshold and local coincidence logic are applied, deciding whether a waveform gets digitised based on its strength and coincidence with a hit on a neighbouring DOM. These steps will be removed in the actual PDOM since advances in technology allow a continuous readout of the PMT waveform by a single ADC instead of the multiple parallel ATWDs [82]. Finally, electronic noise in the digitisers and uncertainties in the time calibration are added and a digitised representation of the waveform is created, which can then be injected in the actual reconstruction chain described in Secs. 4.3 and 4.4.

5.2 The PaPA Code

5.2.1 Idea

As already discussed in Sec. 3.5, the primary science goal of PINGU is the determination of the neutrino mass hierarchy, imprinted on the oscillation probability pattern of the atmospheric electron and muon neutrinos travelling through the Earth's matter potential. Observing such a small effect requires a precise high-level analysis and a detailed knowledge of the detector performance.

The observable in the mass hierarchy analysis is the distribution of arriving neutrinos in the (E, $\cos\vartheta_{\text{zenith}}$) plane. The shape of this distribution is determined not only by the mass hierarchy, but also by the true values of the other neutrino mixing parameters as well as the reconstruction efficiency and precision of the detector. All those quantities—and especially their uncertainties—have to be taken into account estimating the significance which PINGU can determine the mass hierarchy with.

The standard procedure to account for systematic uncertainties in IceCube is a brute force approach, where the parameters in question get increased and decreased by an amount corresponding to the estimated uncertainty. Then the simulation is re-run for each setting and the sensitivity re-evaluated. For the mass hierarchy measurement with PINGU, however, this strategy is not applicable. The reason for this is that in IceCube analyses usually only search for the existence of events passing a certain set of selection criteria—it is a detection experiment, where systematics in general have a much less severe impact compared to the search for a pattern in a large amount of data. Here, correlations between the systematics have to be taken into account as well, while in IceCube they can be considered indepen-

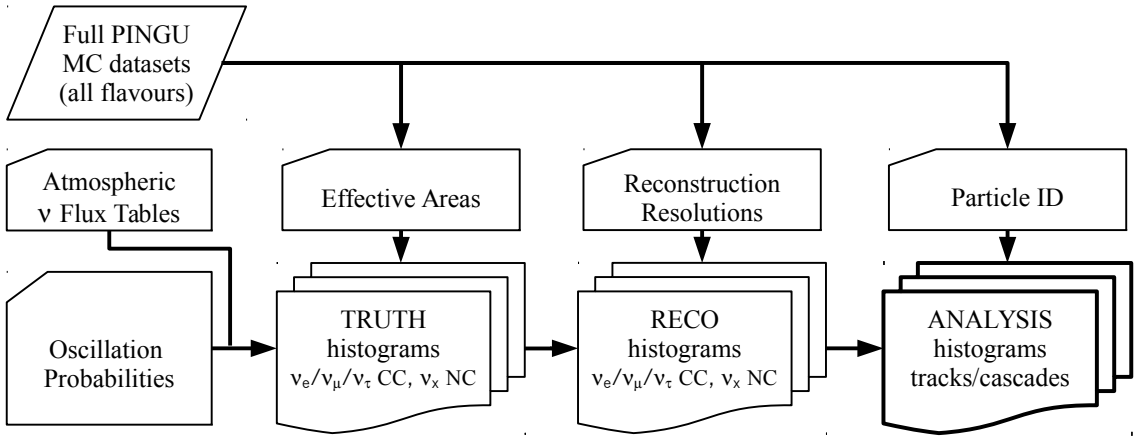

Figure 5.3: Flow chart of the PaPA simulation chain

dent. Additionally, the number of relevant systematic parameters is much higher for PINGU, since in IceCube oscillations are not measured in such detail, and the detector itself is much better understood as it is already operating and hence could be studied and modelled in great detail.

Thus, re-running the whole simulation chain for all possible combinations of these parameters—which will be treated in detail in Sec. 5.2.3—is not feasible in a reasonable amount of time. This is especially true since a comparatively large number of MC events—in the order of 10^6—is needed to make a reliable estimate of the sensitivity. In case of insufficient MC events, statistical fluctuations will exceed the strength of the imprinted mass hierarchy pattern and bias the calculated significance towards high values, as it will be demonstrated in Sec. 6.3.5.

The basic idea how to overcome this fundamental problem is to move away from an event-by-event MC simulation and simulate the event histograms in $(E, \cos\vartheta_{\text{zenith}})$ directly, based on parametrisations of the relevant detector quantities that are extracted from MC data. This is much faster than filling the histograms with individual events and has the benefit of quasi-infinite statistics since the parametrisations will in general be smooth functions, resulting in inherently smooth distributions. This strategy has been implemented in the PaPA code, which will be described in detail in the following section.

A certain amount of MC events is still needed to produce the parametrisations that are the input to PaPA, however it has been found during the parametrisation process that about 20,000 events per neutrino flavour are sufficient to get stable fit results. This is at least a factor of ten less than the number of events required to reach a stable result in Sec. 6.3.5. Also, the MC simulation does not need to be re-run for the different settings of the systematic parameters, as they are tuned within PaPA.

5.2.2 Implementation

The PaPA source code is available in the IceCube subversion repository [121], with a detailed manual and version history maintained at the IceCube wiki pages [122]. It is written entirely in the python programming language [123] and relies heavily on the numpy and scipy packages for numerical and scientific computing [124, 125], with graphical output functionality based on the matplotlib library [126].

PaPA employs a staged processing procedure following the logical ordering and emulating the effects of the multiple physics processes that are involved in the actual measurement, as illustrated in

Fig. 5.3. The code itself is split into two main parts that are run separately. The first part is the `PhysicsSimulation`, responsible only for the calculation of neutrino oscillation probabilities. In the second part of `PaPA`—the `DetectorSimulation`—the actual event histograms as observed in PINGU are calculated from the oscillation probabilities and detector settings.

`PhysicsSimulation`

In the physics simulation the neutrino oscillation probabilities in the full three-flavour mode including matter effects are calculated, using either NuCraft [127] or AtmoWeights [77] and assuming the Preliminary Reference Earth Model (PREM) [55] for the Earth's density profile. The difference between the two software modules is that NuCraft is more versatile in the sense that it can include the oscillation into sterile neutrino flavours, model the generation heights of the neutrinos in the Earth's atmosphere in more detail, and handle a varying electron density in the Earth while in AtmoWeights this value is fixed to 0.5 electrons per nucleon. On the other hand, AtmoWeights is much faster. After checking consistency of the two softwares to the sub-permill level given the same input, AtmoWeights has been selected for the baseline analysis due to its higher speed. The additional options provided by NuCraft are not essential for the study of the neutrino mass hierarchy.

The oscillation probabilities are calculated for ν_e, $\bar{\nu}_e$, ν_μ and $\bar{\nu}_\mu$ injection, such that twelve histograms of oscillation probabilities are obtained:

- $P(\nu_\alpha \to \nu_e/\bar{\nu}_e)$
- $P(\nu_\alpha \to \nu_\mu/\bar{\nu}_\mu)$
- $P(\nu_\alpha \to \nu_\tau/\bar{\nu}_\tau)$

with $\nu_\alpha \in \left[\nu_e, \bar{\nu}_e, \nu_\mu, \bar{\nu}_\mu\right]$. Since no CP violation is assumed[1], there is no oscillation between neutrinos and antineutrinos.

To properly sample the fast oscillations at low energies, `PaPA` can be set up to oversample by given factors in energy and zenith. This means that every point at which the oscillation probabilities are calculated is replaced by a number of evenly spaced points in the vicinity, and the resulting probabilities are averaged over. If no oversampling is applied, the values at the bin centres are used.

As mentioned above, the calculation of the oscillation probabilities is the most time-consuming part of the simulation as here the Schrödinger equations for full three-flavour oscillations in a varying matter potential (see Sec. 3.3) have to be solved numerically.

To get all oscillation probabilities needed for a full analysis, the three relevant oscillation parameters Δm^2_{31}, ϑ_{23}, and ϑ_{13} as well as the energy scale (see Sec. 5.2.3) have to be varied independently at 15 test points each. This typically takes several hours on a single CPU whereas with NuCraft this goes up to a few days. Fortunately, the fiducial settings for the oscillation parameters do not change when the detector parametrisation changes, meaning that the oscillation probabilities only have to be calculated once and can then be re-used for different detector settings.

`DetectorSimulation`

In the detector simulation the pure oscillation probabilities calculated in the physics simulation are converted into an actual detector response, taking into account all detector related parameters. It consists of three major steps, as depicted by the flow chart in Fig. 5.3. At each step, a histogram in $(E, \cos\vartheta_{\text{zenith}})$

[1] It has been checked that if CP violation existed, it would not be detectable with PINGU in any case.

will be produced for every event signature that has to be considered at the respective stage. Event histograms at all simulation stages are shown in Fig. 5.4.

Truth Histograms The first step is to calculate the true rates of events triggering the detector. Therefore, spline fits are generated for the azimuth-averaged atmospheric neutrino flux tables by Honda et al. [128, 129], from which then the fluxes for $\nu_e/\bar{\nu}_e$ and $\nu_\mu/\bar{\nu}_\mu$ can be retrieved. These fluxes are multiplied by the oscillation probabilities $P(\nu_\alpha \to \nu_\beta)$ calculated in the physics simulation to obtain the flux for each neutrino flavour at the detector:

$$\Phi_{\beta,\,\mathrm{det}} = \sum_\alpha \Phi_{\alpha,\,\mathrm{atm}}\; P(\nu_\alpha \to \nu_\beta) \tag{5.2}$$

To obtain event rates, the neutrino fluxes are then multiplied by the effective areas of PINGU for $\nu_e/\bar{\nu}_e$, $\nu_\mu/\bar{\nu}_\mu$ and $\nu_\tau/\bar{\nu}_\tau$ CC interactions as well as $\nu_X/\bar{\nu}_X$ NC—the four different interaction channels—after background rejection cuts (cf. Sec. 4.4). The effective area, a function of neutrino energy and zenith direction, is defined as the collection area an ideal detector identifying every neutrino passing through it would need to have in order to achieve the same event rate as PINGU. For a given interaction channel, it can be calculated according to

$$A_{\mathrm{eff}}(\mathrm{channel}) = \sum_{\mathrm{target}} N_{\mathrm{target}}\; \sigma_{\mathrm{target,\,channel}}\; \varepsilon_{\mathrm{channel}} \quad , \tag{5.3}$$

summing over the possible targets, protons and neutrons. Here, N_{target} stands for the total number of the respective targets inside the detector volume, $\sigma_{\mathrm{target,\,channel}}$ is the interaction cross section in the given channel, and $\varepsilon_{\mathrm{channel}}$ is the efficiency which interactions occurring inside the detector actually pass the triggering, reconstruction, and event selection process with. Although the selection efficiency is usually of the order of 80 % [130], see Fig. 4.8, the effective area is much smaller than the geometrical size of the detector as the interaction cross sections for neutrinos are extremely small (see Sec. 2.3.1). Hence, most neutrinos pass the detector without any interaction.

The number of expected events per lifetime t in a given channel can, thus, be calculated by

$$N_{\nu_\alpha\ \mathrm{CC}} = A_{\mathrm{eff}}(\nu_\alpha\ \mathrm{CC})\; \Phi_{\alpha,\,\mathrm{det}}\; t \tag{5.4}$$

$$N_{\nu_X\ \mathrm{NC}} = A_{\mathrm{eff}}(\nu_X\ \mathrm{NC}) \sum_\alpha \Phi_{\alpha,\,\mathrm{det}}\; t \tag{5.5}$$

for the three charged and one neutral current channels, respectively. Finally, the histograms for neutrinos and anti-neutrinos in the same interaction channels are added since PINGU will not be able to distinguish between them. Also in terms of reconstruction quality and flavour identification (see below) there is no difference between neutrinos and anti-neutrinos.

Since both the fluxes and the effective areas are functions of E and $\cos\vartheta_{\mathrm{zenith}}$, the result of this simulation step are four 2D histograms, one each for the four channels $\nu_e + \bar{\nu}_e$ CC, $\nu_\mu + \bar{\nu}_\mu$ CC, $\nu_\tau + \bar{\nu}_\tau$ CC, and $\nu_X + \bar{\nu}_X$ NC, examples of which are shown in the second row of Fig. 5.4.

Reconstructed Histograms In the next step, event reconstruction is applied to the histograms that up to now contain *true* neutrino energies and directions. This is done by smearing the true histograms with a kernel[2] whose shape depends on the true energy and zenith angle of the event, i. e. the position in the histogram. There are two different ways how these kernels can be handed over to `PaPA`:

[2] A *smearing kernel* can be interpreted as a binned point spread function.

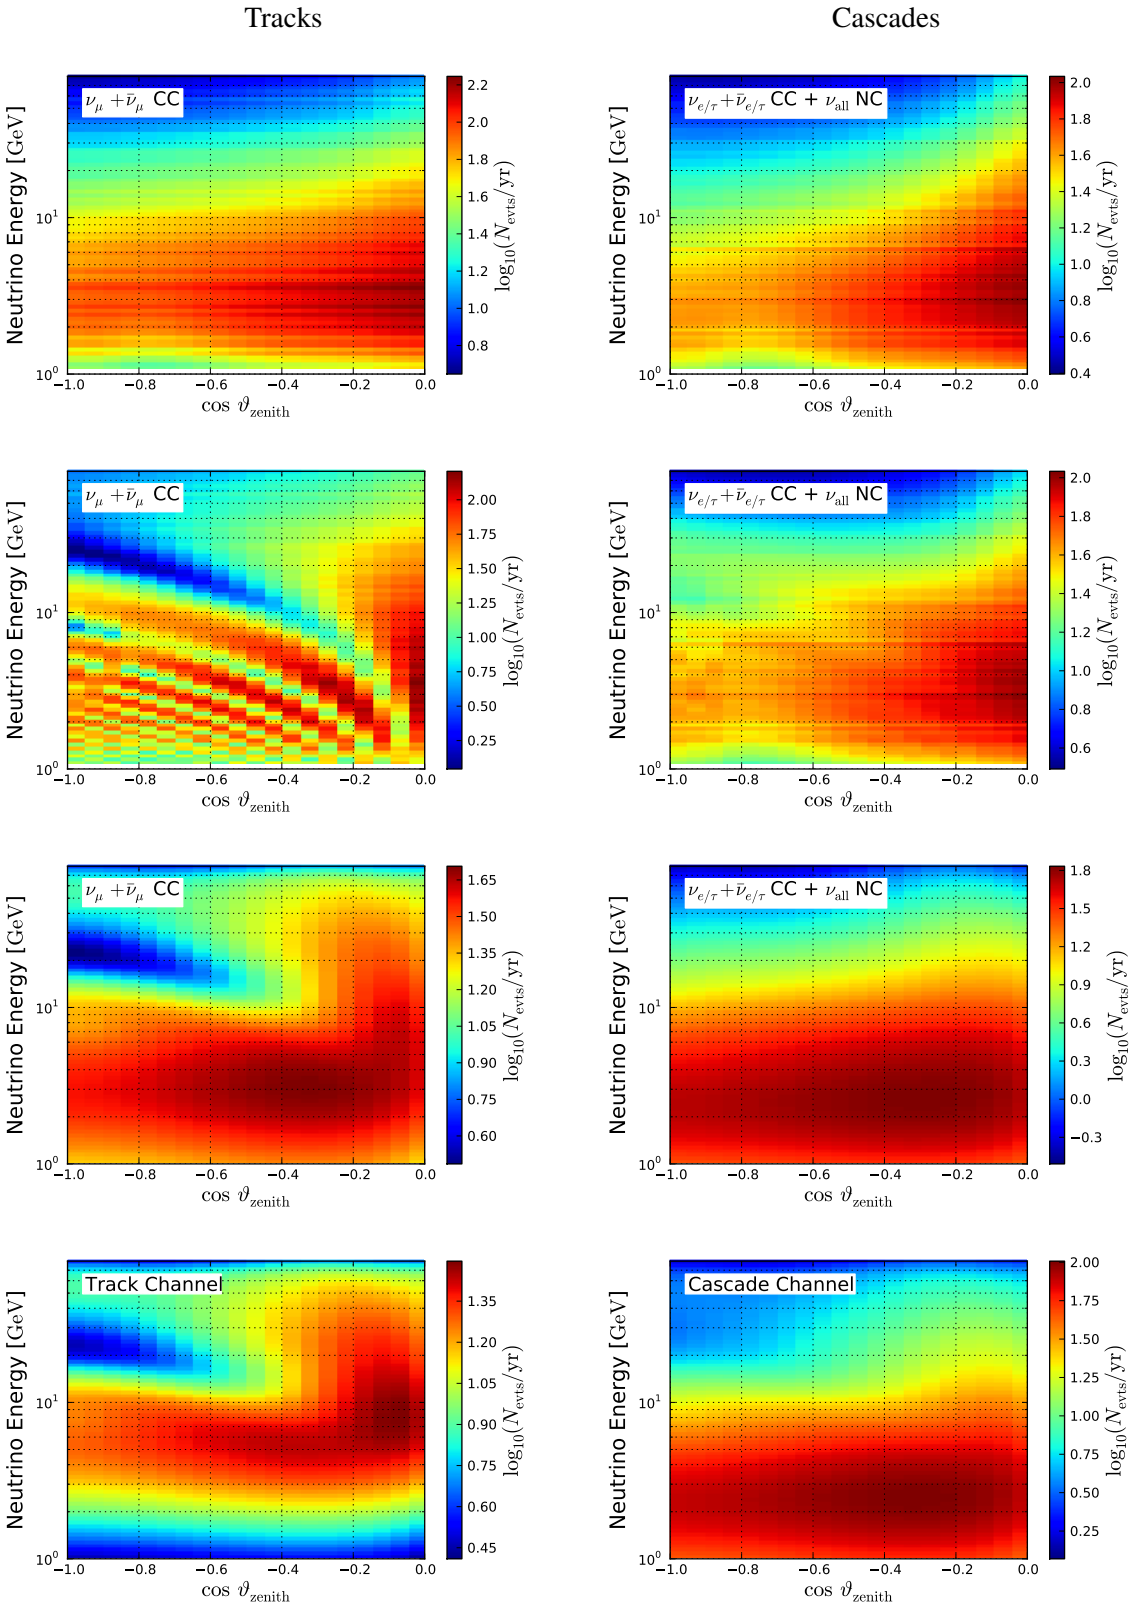

Figure 5.4: Event counts in one year of PINGU lifetime at the different simulation stages, assuming normal mass hierarchy. From top to bottom: Truth Histograms without and with oscillations, Reconstructed and Analysis Histograms. The variables are true neutrino energy and direction for the upper two rows of histograms and reconstructed energy and direction for the bottom ones. For details, refer to the text.

(a) Parametrised point spread functions.
These have to be fitted from MC data and generally take the form of a double Gaussian for both energy and $\cos\vartheta$ to properly account for the tails of the reconstruction distributions:

$$\mathrm{PSF} = (1-f)\cdot\exp\left(-\frac{(x-\mu_1)^2}{\sigma_1^2}\right) + f\cdot\exp\left(-\frac{(x-\mu_2)^2}{\sigma_2^2}\right) \tag{5.6}$$

The five parameters (two mean values μ_i, two widths σ_i, and the relative normalisation f) are supplied to the code as a function of true neutrino energy.
Here it can happen that events are "leaking" out of the histograms. At the edges of the energy range this effect can be neglected since the event rates are small anyways (due to low flux/small effective area). In the directional dimension, losing events at the horizon is accepted since mis-reconstructed events will be also be lost in the real experiment due to cutting the analysis at the horizon. Events that would be lost at the zenith are reflected back into the histogram, as events migrating "over the zenith" will flip the sign of their azimuth angle (which is irrelevant) and move back to larger values of the zenith angle.

(b) Tabulated point spread functions.
If the MC statistics is sufficient, i. e. in the order of 10^6 events per flavour, the resolutions can be retrieved directly from the MC data. This means creating a 4D histogram of true vs. reconstructed (E, $\cos\vartheta_{\mathrm{zenith}}$), which is equivalent to the individual smearing kernel for each bin in true (E, $\cos\vartheta_{\mathrm{zenith}}$).

Option (a) is used primarily in this study, since for the PINGU geometry V36 that is the baseline for all analyses, the amount of MC events is insufficient for option (b). Also, the parametrised reconstruction resolutions make it possible to investigate the effects of an improved resolution within a well-defined metric in Sec. 6.4.1. For the earlier PINGU geometry V15, however, enough MC events exists to follow option (b) for the event reconstruction stage. This is reported on in Sec. 6.3.5, where also the effect that an insufficient MC sample has on the calculated NMH sensitivity is demonstrated.

Example histograms after the reconstruction stage—still in the four interaction channels already present after the previous stage—are shown in the third row of Fig. 5.4. However one has to keep in mind that now the variables have changed from *true* to *reconstructed* neutrino energies and directions. The reconstructed neutrino energy is given by the sum of the deposited energies of the cascade and track as fitted by HybridReco/MultiNest. The reconstructed direction comes from the same algorithm, where track and cascade are assumed to be aligned (cf. Sec. 4.3.7).

Analysis Histograms Up to now, the histograms are still divided into the four "effective flavour" channels, ν_e, ν_μ, ν_τ CC and ν_X NC, each including anti-neutrinos. The final step in the detector simulation is to apply the particle flavour identification (PID). However, current tools for particle ID can only distinguish events with an outgoing muon signature (only present in ν_μ CC interactions), which are classified as tracks, from other events (cascades), as already described above (cf. Sec. 4.4.3).

So for every interaction channel, two analytic functions of the neutrino energy have to be supplied to `PaPA`, giving the probabilities of identifying a neutrino of that channel and energy as either track or cascade. If those two probabilities add up to less than one, the remaining events are assigned to an "unidentified" channel that can be used for the analysis as well.

For the baseline settings (Sec. 6.2), the PID is a binary decision, meaning that there are no unidentified events. In Sec. 6.3.4, the effects of an event selection producing cascade and track samples of higher purity plus a set of unidentified events will be discussed.

After this final stage, all key features of a real experiment (triggering, event selection, reconstruction, and flavour identification) have been taken into account, and the analysis histograms are the best estimate of the outcome of an actual PINGU physics run. Example histograms after this final step, also referred to as analysis histograms, are shown in the bottom row of Fig. 5.4 and will also be used throughout Chapter 6 to illustrate the various effects studied.

5.2.3 Systematic Parameters

The main benefit of the fast simulation provided by `PaPA` is that it makes the inclusion of the many systematic parameters feasible, which are needed to correctly evaluate the mass hierarchy significance. The parameters that were considered in this thesis can be divided into two groups. The first one contains the *physics* parameters, i. e. the neutrino oscillation parameters: three mixing angles, two mass splittings, the CP violating phase and of course the sign of the NMH. The best fit values and priors for the mixing angles and mass splittings as listed in Tab. 3.1 are taken from [47], assuming inverted NMH as fiducial value. The CP violating phase δ_{CP} is set to a fiducial value of zero. After having checked that their impact is negligible, the so-called solar mixing parameters Δm^2_{21} and ϑ_{12} as well as the CP phase will not be left free to vary any more in order to save computation time.

The mass hierarchy needs special treatment since it is a binary quantity, yet the Fisher matrix formalism (see Sec. 6.1), which will be used to quickly evaluate the simulation outcome, can only treat continuous parameters. Thus the NMH needs to be converted into a continuous variable. This is achieved by calculating the oscillation probabilities assuming normal and inverted NMH once each and propagating both through the full `DetectorSimulation` chain as described in the previous section. Then a hierarchy parameter $0 \leq h \leq 1$ is introduced that allows for a continuous transition between normal (NH) and inverted (IH) hierarchy:

$$N(h) = hN_{\mathrm{NH}} + (1-h)N_{\mathrm{IH}} \quad , \tag{5.7}$$

where N_{NH} and N_{IH} are the expected number of events in a given $(E, \cos\vartheta_{\mathrm{zenith}})$ bin of the analysis histograms for normal and inverted mass hierarchy, respectively. All other oscillation parameters are kept at their fiducial values.

The rationale is that while in theory any value of h other than 0 or 1 is unphysical, in practice this definition allows to assess the deviation from the physical points, as long as these are close to the fiducial values. It is such a measure of the hierarchy *likeness*, similar to a likelihood[3]. If no other parameters are present, this description is mathematically equivalent to a bin-wise $\Delta\chi^2$ sum as defined in [53], as it will be shown in Sec. 6.1.3.

The second set of parameters are related to the detector, which is commonly referred to as nuisance parameters. Most of them scale properties of the detector up or down, reflecting uncertainties in the knowledge of the detector performance. In several cases, however, they can be interpreted as theoretical uncertainties on, e. g., the neutrino cross sections or atmospheric flux tables. The systematic parameters belonging to this class are, ordered according to their entry point into the simulation chain:

The energy scale $\mathbf{s_E}$ describes a potential scaling error in the assumed relation between true and reconstructed neutrino energy. Such a mis-scaling, i. e. a systematic under- or overestimation of the neutrino energy can be the result of a wrong value for e. g. PMT efficiency, absorption effects in the ice model etc. that are used to reconstruct the neutrino events. It is implemented in a way that the true values of the energy bins that are fed to the oscillation code are scaled w. r. t. the energy

[3] It can also be interpreted as a measure of the orthogonal distance of the hyperplanes in parameter spaces defined by normal and inverted hierarchy

bins in the analysis according to

$$E_{\text{true}} = s_E \cdot E_{\text{ana}} \quad . \tag{5.8}$$

So the energy bins in the actual histograms remain unchanged, but the oscillation minima and maxima do not appear at the expected energies. This mimics the effect of a wrong energy scale in an actual analysis, since one would of course assume the energy scale to be correct. Note that in PaPA, the energy scale does affect neither the atmospheric flux calculations for the South Pole, as they are usually validated with existing IceCube data that would likely have the same energy mis-scaling as PINGU. Nor does it change the effective areas, since their main feature, the energy threshold for detecting a neutrino interaction, would scale in parallel with the energy bin edges. Other accompanying effects, such as the overall normalisation of the effective area, are absorbed in different systematic parameters.
Although the energy scale is a detector-related parameter, it has to enter at the `PhysicsSimulation` stage already as the rescaled energy bins have to be passed on to the calculation of the oscillation probabilities. Thus, it is the very first detector parameter in the simulation chain to be accounted for.

The relative flux normalisation $\mathbf{r}_{\mathbf{\Phi},\nu_e-\nu_\mu}$ represents the uncertainty on the relative contributions of ν_e and $\bar{\nu}_e$ vs. ν_μ and $\bar{\nu}_\mu$ to the total flux of atmospheric neutrinos, which is approximately 1:2. It is applied to the un-oscillated atmospheric neutrino fluxes (cf. (5.2)) such that

$$\Phi'_{\alpha,\,\text{atm}} = \Phi_{\alpha,\,\text{atm}} \cdot (1 \pm r_{\Phi,\,\nu_e-\nu_\mu}) \quad , \tag{5.9}$$

with positive sign for $\alpha = \nu_e$ or $\bar{\nu}_e$ and negative sign for $\alpha = \nu_\mu$ or $\bar{\nu}_\mu$.

The effective area scale $\mathbf{s}_{\mathbf{A}_{\text{eff}}}$ opens the possibility that the effective areas for all interaction channels scale differently with energy than assumed in the generation of the MC from where they are extracted. It also partially compensates for a potential error in the spectral index of the atmospheric flux, since it effectively alters the total number of events per energy bin as a function of energy according to

$$A_{\text{eff}}{}' = A_{\text{eff}} \cdot (1 + s_{A_{\text{eff}}} \cdot E) \quad . \tag{5.10}$$

The relative effective area normalisation $\mathbf{r}_{\mathbf{A}_{\text{eff}},\nu-\bar{\nu}}$ for neutrinos and anti-neutrinos accounts for a possible error in the relative normalisation of the neutrino-nucleon cross sections for neutrinos vs. antineutrinos, which effectively alters the effective areas[4]. The effective areas are then given by

$$A_{\text{eff}}{}' = A_{\text{eff}} \cdot (1 \pm r_{A_{\text{eff}},\,\nu-\bar{\nu}}) \quad , \tag{5.11}$$

with the positive sign for neutrinos and the negative sign for antineutrinos. Since the relative strength of neutrino and antineutrino interactions governs the size of the NMH signal that can be observed with PINGU (see Sec. 3.5), this parameter is of particular importance.

The overall effective area normalisation $\mathbf{n}_{\mathbf{A}_{\text{eff}}}$ scales the effective areas for all channels independently of energy:

$$A_{\text{eff}}{}' = A_{\text{eff}} \cdot n_{A_{\text{eff}}} \tag{5.12}$$

[4] Inflicting this re-normalisation on the atmospheric fluxes instead of the effective areas would have the same effect. Since the effects are completely degenerate and uncertainties on the relative cross sections are larger than on the relative fluxes, the flux normalisation is not treated as a parameter.

Since it just evenly increases the total number of events in all channels and, hence, is fully degenerate with the detector lifetime, it is not that interesting by itself. However, it is needed in combination with the relative scalings mentioned above to effectively allow for separate scaling of one single channel.
Obviously, all three parameters related to the effective area enter the calculation at the same time during the calculation of the true event rates, represented by (5.4) and (5.5), i. e. after the oscillation probabilities have been applied.

The PID scaling s_{PID} scales the probabilities to detect a track or a cascade signature up or down, accounting for the possibility to mis-estimate the overall effectiveness of the particle ID algorithm.

The PID offset Δ_{PID} shifts the whole particle ID function in energy to open the possibility that PID might become effective[5] at higher or lower energies than expected from Monte Carlo data. The actual PID function is then

$$P'_{\mathrm{ID}}(E) = s_{\mathrm{PID}} \cdot P_{\mathrm{ID}}(E - \Delta_{\mathrm{PID}}) \quad . \tag{5.13}$$

[5] In particular for the correct identification of ν_μ CC events as track-like, the PID function has a step-like shape, with a steep increase of the selection efficiency between 5 and 10 GeV, cf. Sec. 6.2.

CHAPTER 6

Analysis

Now that the tools and theoretical foundations needed for the simulation of an PINGU physics run have been presented and described, the next step is the actual calculation of PINGU's physics capabilities, in particular its sensitivity for the neutrino mass hierarchy. This will be done using the Fisher information matrix, a tool to quickly evaluate the covariance matrix of an experiment where a large number of systematic uncertainties has to be taken into account, which will be introduced in Sec. 6.1.

Thereafter, the simulation input that was used for the generation of the results is documented (Sec. 6.2) before finally the results themselves can be presented and discussed in Sec. 6.3. Looking even more into the future, in Sec. 6.4 the possible benefits of instrumenting PINGU with the next generation optical modules described in Sec. 4.5 will be estimated. Finally, one can make use of the Fisher matrix's ability to easily combine the results of different experiments measuring the same physical effect to evaluate the benefits gained from the joint analysis of PINGU and another neutrino oscillation experiment, JUNO [72] (Sec. 6.5). The chapter concludes with a summary of the findings in Sec. 6.6.

6.1 Fisher Information Matrix

The Fisher information matrix, or just Fisher matrix in the following, provides a way to estimate the full covariance matrix of an experiment and therefore the accuracy of its intended measurement in advance. It has been widely used, especially in cosmology [131], but also in neutrino astrophysics [132]. For a detailed discussion see e. g. [131, 133, 134].

The experiment that is to be modelled is characterised by two sets of variables:

Observables are the variables f_n that will actually be measured by the experiment. In the case of PINGU, these are the expected event counts binned in energy and zenith angle. Since the partial derivatives of the observables w. r. t. the parameters enter the calculation of the Fisher matrix, their dependence on the parameters has to be known either analytically or (as in this analysis) from simulations, such that these derivatives can be determined in an analytical or numerical way.

Parameters are the variables p_i that will be extracted from the measurement (i. e. from the observables). These are the physical parameters which are the actual target of the experiment as well as nuisance parameters required to account for systematic uncertainties.

Then the Fisher matrix is defined as:

$$\mathcal{F}_{ij} = \sum_n \frac{1}{\sigma_n^2} \frac{\partial f_n}{\partial p_i} \frac{\partial f_n}{\partial p_j}\bigg|_{\text{fid. model}} \tag{6.1}$$

Here the σ_n denote the errors on the measurement of the observables. Since in this analysis, these are the expected numbers of events in the bins n of the analysis histograms (cf. Sec. 5.2.2), $f_n = N_n$, one can apply Poissonian statistics where the errors are simply given by the square root of the number of events:

$$\sigma_n = \sqrt{f_n} = \sqrt{N_n} \tag{6.2}$$

The derivatives in (6.1) are evaluated at the set of "true" values of p_i that are used as an input for the simulation. This set of parameters p_i is called the *fiducial model* and is chosen with the best existing knowledge.

6.1.1 Properties

Once the Fisher matrix has been constructed, the covariance matrix of the experiment is obtained by inverting the Fisher matrix. Then the full errors of the parameters σ_i^{full} and their correlations can be read from the covariance matrix $\mathcal{S}$:

$$\sigma_i = \sigma_i^{\text{full}} = \sqrt{\mathcal{S}_{ii}} = \sqrt{(\mathcal{F}^{-1})_{ii}} \tag{6.3}$$

The purely statistical error of a given parameter σ_i^{stat} (i. e. the error that is obtained assuming all other parameters are fixed) can be calculated from the corresponding diagonal element of the Fisher matrix:

$$\sigma_i^{\text{stat}} = \sqrt{(\mathcal{F}_{ii})^{-1}} \tag{6.4}$$

External constraints on parameters, so-called *priors*, can easily be incorporated as well. If parameter i has an prior σ_i^{ext}, it can simply be added to the corresponding diagonal element of the Fisher matrix:

$$\mathcal{F}_{ii}^{\text{prior}} = \mathcal{F}_{ii} + \left(\frac{1}{\sigma_i^{\text{ext}}}\right)^2 \tag{6.5}$$

This assumes that the prior is *Gaussian*, i. e. the underlying probability density function for the true value of parameter p_i is a Gaussian distribution with the width σ_i^{ext} centred at the fiducial value of p_i. To fix one parameter completely, the corresponding row and column are removed from the Fisher matrix before inversion.

With this, the full error on a certain parameter can be decomposed into the purely statistical part (6.4) and the contribution arising from the uncertainties of the other parameters, denoted as *systematic* error σ_i^{syst} below:

$$\sigma_i^{\text{syst}} = \sqrt{\sigma_i^2 - \left(\sigma_i^{\text{stat}}\right)^2} \quad , \tag{6.6}$$

where σ_i is calculated without any possible priors on parameter i.

The correlation coefficient between two parameters i and j is given by:

$$c_{ij} = \frac{\left(\mathcal{F}^{-1}\right)_{ij}}{\sigma_i \sigma_j} = \frac{\sigma_{ij}}{\sigma_i \sigma_j} \tag{6.7}$$

With this, the error ellipse in the (p_i, p_j) plane can be calculated as well. Its semi-major and semi-minor axes a and b are given by the eigenvalues of the 2×2 sub-matrix of the covariance matrix corresponding to the two parameters,

$$\mathcal{S}^{(i,j)} = \begin{pmatrix} \sigma_i^2 & \sigma_{ij} \\ \sigma_{ij} & \sigma_j^2 \end{pmatrix} \quad , \tag{6.8}$$

which are equal to [135]

$$a^2 = \frac{\sigma_i^2 + \sigma_j^2}{2} + \sqrt{\frac{\left(\sigma_i^2 + \sigma_j^2\right)^2}{4} + \sigma_{ij}^2} \tag{6.9}$$

$$b^2 = \frac{\sigma_i^2 + \sigma_j^2}{2} - \sqrt{\frac{\left(\sigma_i^2 + \sigma_j^2\right)^2}{4} + \sigma_{ij}^2} \tag{6.10}$$

for $\sigma_i > \sigma_j$. Otherwise, the roles of semi-major and semi-minor axis have to be exchanged. These values have to be scaled by a factor α reflecting the confidence level the ellipse is supposed to represent, e. g. CL = 68 % for a 1σ ellipse. It can be determined with

$$\alpha = \sqrt{\mathrm{PPF}_{\chi^2,2}(\mathrm{CL})} \quad , \tag{6.11}$$

where $\mathrm{PPF}_{\chi^2,2}$ is the *percent point function*—the inverse of the cumulative distribution function—for the χ^2 distribution with two degrees of freedom. The rotation of the ellipse corresponds to the angle φ between the p_i axis and the eigenvector of the reduced covariance matrix (6.8):

$$\tan 2\varphi = \frac{2\sigma_{ij}}{\sigma_i^2 - \sigma_j^2} \tag{6.12}$$

Since the parameters p_i and p_j do not necessarily have the same dimension, all values in the equations above have to taken as dimensionless, i. e. in units of actual size on a plot. For this reason, (6.12) has to be divided by the aspect ratio of the plot before taking the arc tangent.

Another important property of the Fisher matrix is its additivity. As one can see from (6.1), where the total Fisher matrix is given by summing the contributions from all observables, one can always add a new measurement to the ones already accounted for just by adding the respective Fisher matrices. If there are parameters that only appear in one of the measurements—e. g. some nuisance parameter for detector A that does not affect detector B—the matrix of the experiment not yet having said parameter is expanded by a corresponding row and column filled with zeroes: As the measurement does not depend on the new parameter, the uncertainty on its value as extracted from the measurement (6.4) is $1/\sqrt{0} \to \infty$. The full covariance matrix of the combination of the two experiments is then simply given by the inverse of the added Fisher matrices.

6.1.2 Prerequisites

As the simplicity of the Fisher matrix formalism suggests, it is not universally applicable, but an approximation that is only valid under certain circumstances. The basic requirement that has to be met is that the fiducial model, in which the Fisher matrix is constructed, already gives a good description of the actual "truth". There are several criteria to judge whether this requirement is in fact fulfilled, which turn out to all be equivalent:

(a) In parameter space, the fiducial and true points have to be close enough, such that the dependence of the observables on the parameters, $f_n(p_i)$, can be approximated by linear functions. I. e. the partial derivatives $\frac{\partial f_n}{\partial p_i}$ entering in (6.1) are the same at both points. If this was not the case, the experiment's "response" to the parameters would be different at the two points and hence the fiducial point was not appropriate to make predictions about the actual truth.

(b) When constructing the likelihood landscape over the parameter space, such that

$$\frac{\partial^2 \mathcal{L}}{\partial p_i^2} = \sum_n \frac{\partial f_n}{\partial p_i} \tag{6.13}$$

with a minimum of $\mathcal{L}$ at the true point, at the fiducial point the shape of $\mathcal{L}$ can still be approximated by a parabola with the apex at the true point. This ensures that the error ellipses, which are nothing else than iso-likelihood contours, are in fact elliptical.

(c) The test statistics of the experiment is Gaussian. This means that the likelihood values of an ensemble of pseudo-experiments follow a Gaussian distribution, no matter whether they are thrown at the fiducial or at the true point. *Throwing a pseudo-experiment* means in this case that a possible "real" experimental outcome is simulated by generating a 2D histogram of events filled with random numbers following a Poisson distribution with the means taken from the expected event numbers generated with `PaPA`.

That these conditions are, in fact, fulfilled for the baseline settings of PINGU is demonstrated in App. B. As for conditions (b) and (c) this would require a likelihood evaluation of the full parameter space[1], here the validation is restricted to a case with only two systematic parameters, Δm^2_{31} and ϑ_{23}.

6.1.3 The Hierarchy Parameter

The Fisher matrix can only handle continuous parameters p_i, as the derivatives of the observables w. r. t. these parameters enter in (6.1). For this reason the hierarchy parameter h was introduced in (5.7), making the intrinsically binary[2] neutrino mass hierarchy continuous. However it still needs to be demonstrated that this definition leads to a result that can be interpreted in a meaningful way.

In [53], a metric is defined that allows to evaluate the sensitivity of an atmospheric neutrino experiment to the mass hierarchy in terms of standard $\Delta\chi^2$ statistics. For every bin of the final analysis histograms in $(E, \cos\vartheta_{\text{zenith}})$, a $\Delta\chi$ is defined as

$$\Delta\chi = \frac{N_{\text{NH}} - N_{\text{IH}}}{\sqrt{N_{\text{NH}}}} \quad . \tag{6.14}$$

The expected significance of the experiment's mass hierarchy measurement is then given by

$$\mathcal{S}_{\chi^2} = \sqrt{\sum_{\text{bins}} \Delta\chi^2} = \sqrt{\sum_{\text{bins}} \frac{(N_{\text{NH}} - N_{\text{IH}})^2}{N_{\text{NH}}}} \quad . \tag{6.15}$$

With the definition of h according to (5.7), from the Fisher Matrix the statistical error on h can be

[1] This is not possible within reasonable computing time, which is why the Fisher Matrix approach has been chosen in the first place.

[2] By construction it can only be either normal or inverted.

retrieved:

$$\sigma_h^{\text{stat}} = \sqrt{(\mathcal{F}_{hh})^{-1}} \quad , \tag{6.16}$$

with

$$\mathcal{F}_{hh} = \sum_{\text{bins}} \frac{1}{\sigma_b^2} \frac{\partial N_b}{\partial h} \frac{\partial N_b}{\partial h} = \sum_{\text{bins}} \frac{1}{N_{\text{NH}}} (N_{\text{NH}} - N_{\text{IH}})^2 \quad . \tag{6.17}$$

As the pure normal and pure inverted hierarchy case correspond to $h = 1$ and $h = 0$, respectively, i. e. the two cases that have to be distinguished differ by 1, the statistical significance (without including any systematic parameters) retrieved from the Fisher matrix is

$$\mathcal{S}_{\text{Fisher}} = \frac{1}{\sigma_h^{\text{stat}}} = \sqrt{\mathcal{F}_{hh}} = \sqrt{\sum_{\text{bins}} \frac{(N_{\text{NH}} - N_{\text{IH}})^2}{N_{\text{NH}}}} = \mathcal{S}_{\chi^2} \quad . \tag{6.18}$$

Thus the significances calculated according to $\Delta\chi^2$ statistics, where by construction no systematic parameters are included, and from the Fisher matrix with explicitly excluding systematics are, in fact, identical.

6.1.4 Constructing the Fisher Matrix with `PaPA`

To construct the Fisher matrix for PINGU primarily means to calculate the partial derivatives entering (6.1), i. e. the derivatives of the expected number of events w. r. t. to all systematic parameters p_i introduced in Sec. 5.2.3 for each bin in the final analysis histograms. In order to do this, for every parameter a set of test points[3] is selected in the vicinity of its fiducial value and `PaPA` is run for each of these test points while all other parameters are kept at their fiducial values.

Then for every bin b in the analysis histograms, the expected bin count N_b is fitted as a function of p_i with a parabola, i. e.

$$N_b^{\text{fit}}(p_i) = c_0 + c_1 \cdot p_i + c_2 \cdot p_i^2 \quad . \tag{6.19}$$

The sought-after derivative is then simply given by

$$\frac{\partial N_b^{\text{fit}}}{\partial p_i} = c_1 + 2c_2 p_i \tag{6.20}$$

with p_i at the fiducial value. As this procedure is repeated for ever parameter and every bin, it produces all ingredients needed to finally construct the Fisher matrix.

Here another advantage of the Fisher matrix formalism becomes obvious: since all systematic parameters are treated independently from each other, adding a new parameter results only in a small relative increase in computing time. The parameter space is sampled only along one-dimensional paths aligned with the parameter axes, thus adding an extra dimension only means adding one path to the sampling. Hence the time needed to get a result scales only linearly with the number of parameters taken into account while for algorithms sampling the full parameter space the scaling is exponential. Using a more graphic explanation, the Fisher matrix infers the shape of the likelihood landscape by retrieving a cross-section in each coordinate direction and then "rotates" them assuming a parabolic shape.

[3] Typically $\approx$ 10 different values.

6.2 Simulation Input

In this section, the actual values of the inputs required by PaPA as described in Sec. 5.2.2 will be presented as they were used to generate the results reported in Sec. 6.3. The detector-related inputs were extracted from the official PINGU Monte Carlo datasets for geometry V36[4].

6.2.1 Atmospheric Neutrino Flux

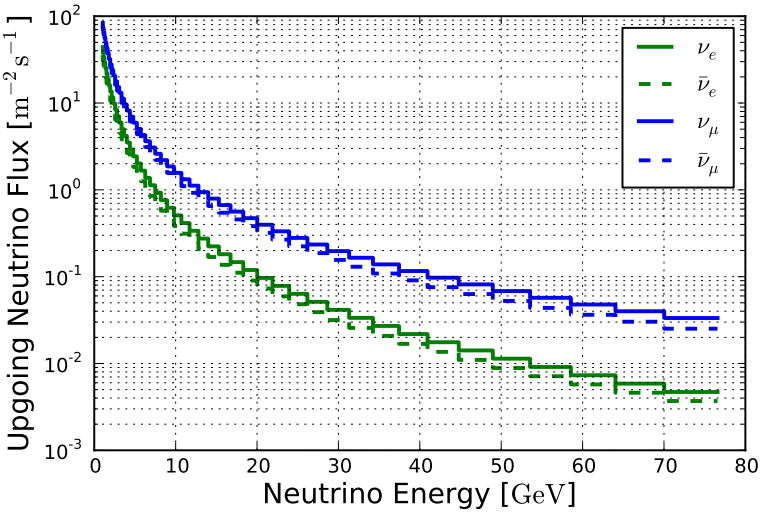

Figure 6.1: The atmospheric neutrino flux at the South Pole integrated over all upgoing ($\cos\vartheta_{\text{zenith}}$ in $[-1,0]$) directions. Based on the azimuth-averaged neutrino flux tables from [129].

As already mentioned, the incoming atmospheric neutrino flux without oscillations is calculated from the 2014 re-calculation of the flux tables published by Honda et al. [128, 129]. A plot of the flux in the energy range covered by the PINGU simulation ($1-80\,\text{GeV}$) is shown in Fig. 6.1. The flux has been integrated over all upgoing ($\cos\vartheta_{\text{zenith}}$ in $[-1,0]$) directions, as downgoing neutrinos arriving from above the detector are not included in the mass hierarchy analysis. Since they do not pass through a significant amount of matter, they do not bear any information on the neutrino mass hierarchy.

6.2.2 Oscillation Probabilities

The neutrino oscillation probabilities have been calculated using the AtmoWeights code for the PREM Earth density profile as described in Sec. 3.5. The fiducial values of the mixing parameters used for the calculation follow the global fit of Fogli et al. [47], listed in Tab. 3.1, with priors according to the stated uncertainties. Example plots of the probabilities demonstrating the characteristic signature of the mass hierarchy are shown in high resolution in that section as well, the full set of all possible oscillation channels can be found in App. C. The binning of those plots is the same as used for the actual analysis, which is 79 logarithmic bins in energy between 1 GeV and 80 GeV and 20 equally sized bins between -1 and 0 in $\cos\vartheta_{\text{zenith}}$. In the calculation of the probabilities, however, the binning has been oversampled (cf. Sec. 5.2.2) by a factor of eight in energy and two in $\cos\vartheta_{\text{zenith}}$.

[4] These are the PINGU Monte Carlo runs 390 for ν_e and $\bar{\nu}_e$, 389 for ν_μ and $\bar{\nu}_\mu$, and 390 for ν_τ and $\bar{\nu}_\tau$.

6.2.3 Effective Areas

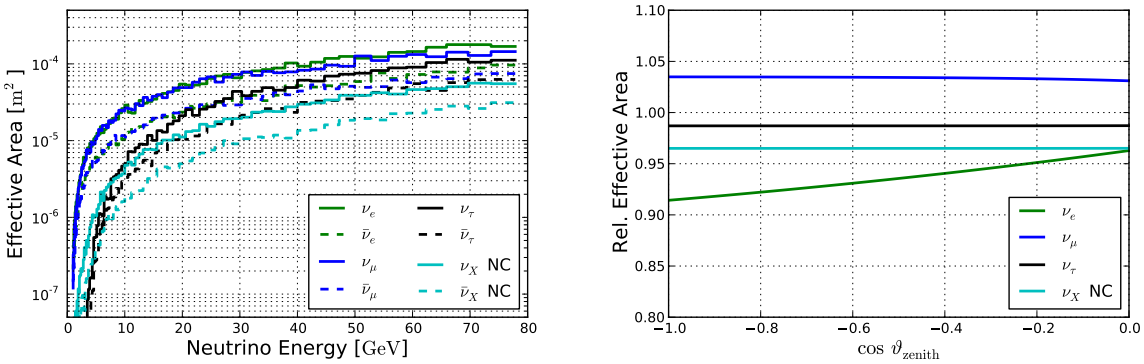

Figure 6.2: Effective areas for all relevant neutrino interactions. Shown are energy (left) and zenith (right) dependence.

The effective areas are extracted from PINGU Monte Carlo datasets via the OneWeight quantity multiplied by 4π, which is equivalent to a per-event effective area [136]. The zenith dependence of the effective area is modelled by an analytic fit to the MC data, assuming that energy and zenith dependence can be handled separately. A plot of the effective areas and their zenith dependence is shown in Fig. 6.2.

Note that the absolute values of the functions describing the zenith dependence are chosen to give a good visual separation between the channels. In `PaPA` the functions get internally normalised.

6.2.4 Reconstruction Resolutions

The event reconstruction for the baseline model will be simulated by smearing the 2D event histograms with kernels represented by the sum of two Gaussian distributions (cf. Sec. 5.2.2). The five parameters needed to describe these Gaussians, see (5.6), are given as functions of the neutrino energy. To get the energy dependence of these five parameters, a two-stage fitting procedure is applied.

In the first stage, all events of a given interaction channel are divided into subsets according to their true neutrino energy, where each subset covers an energy range of 2 GeV. For each of the subsets up to an energy of 20 GeV[5], the difference between true and reconstructed energy is histogrammed and fitted with the aforementioned kernel function. This results in values for the five parameters of (5.6) as a function of energy. These are fitted again, now as (piecewise) linear functions of the true neutrino energy.

After repeating the procedure for the resolution in $\cos\vartheta_{\mathrm{zenith}}$, the fit function definitions are stored in a dictionary that serves as input to `PaPA`. This dictionary can be found in App. D. For the actual fits to the resolutions, examples are shown in Fig. 6.3. In the ν_τ CC and ν_X NC interaction channels, the reconstructed energies are biased towards too low values. This is a result of the neutrinos in the respective final states which is not detected and hence carries away "missing energy" (cf. Sec. 2.3.2).

Although the fits are only anchored to data up to 20 GeV, they are extrapolated to higher energies. This approach is valid as events beyond 20 GeV are so far above PINGU's energy threshold that new features in the description of the reconstruction resolutions are not to be expected. In addition, the mass hierarchy signature is located at energies below 10 GeV (see Figs. 3.7 – 3.9), such that inaccuracies

[5] At higher energies, the event statistics are too low for the fit to converge.

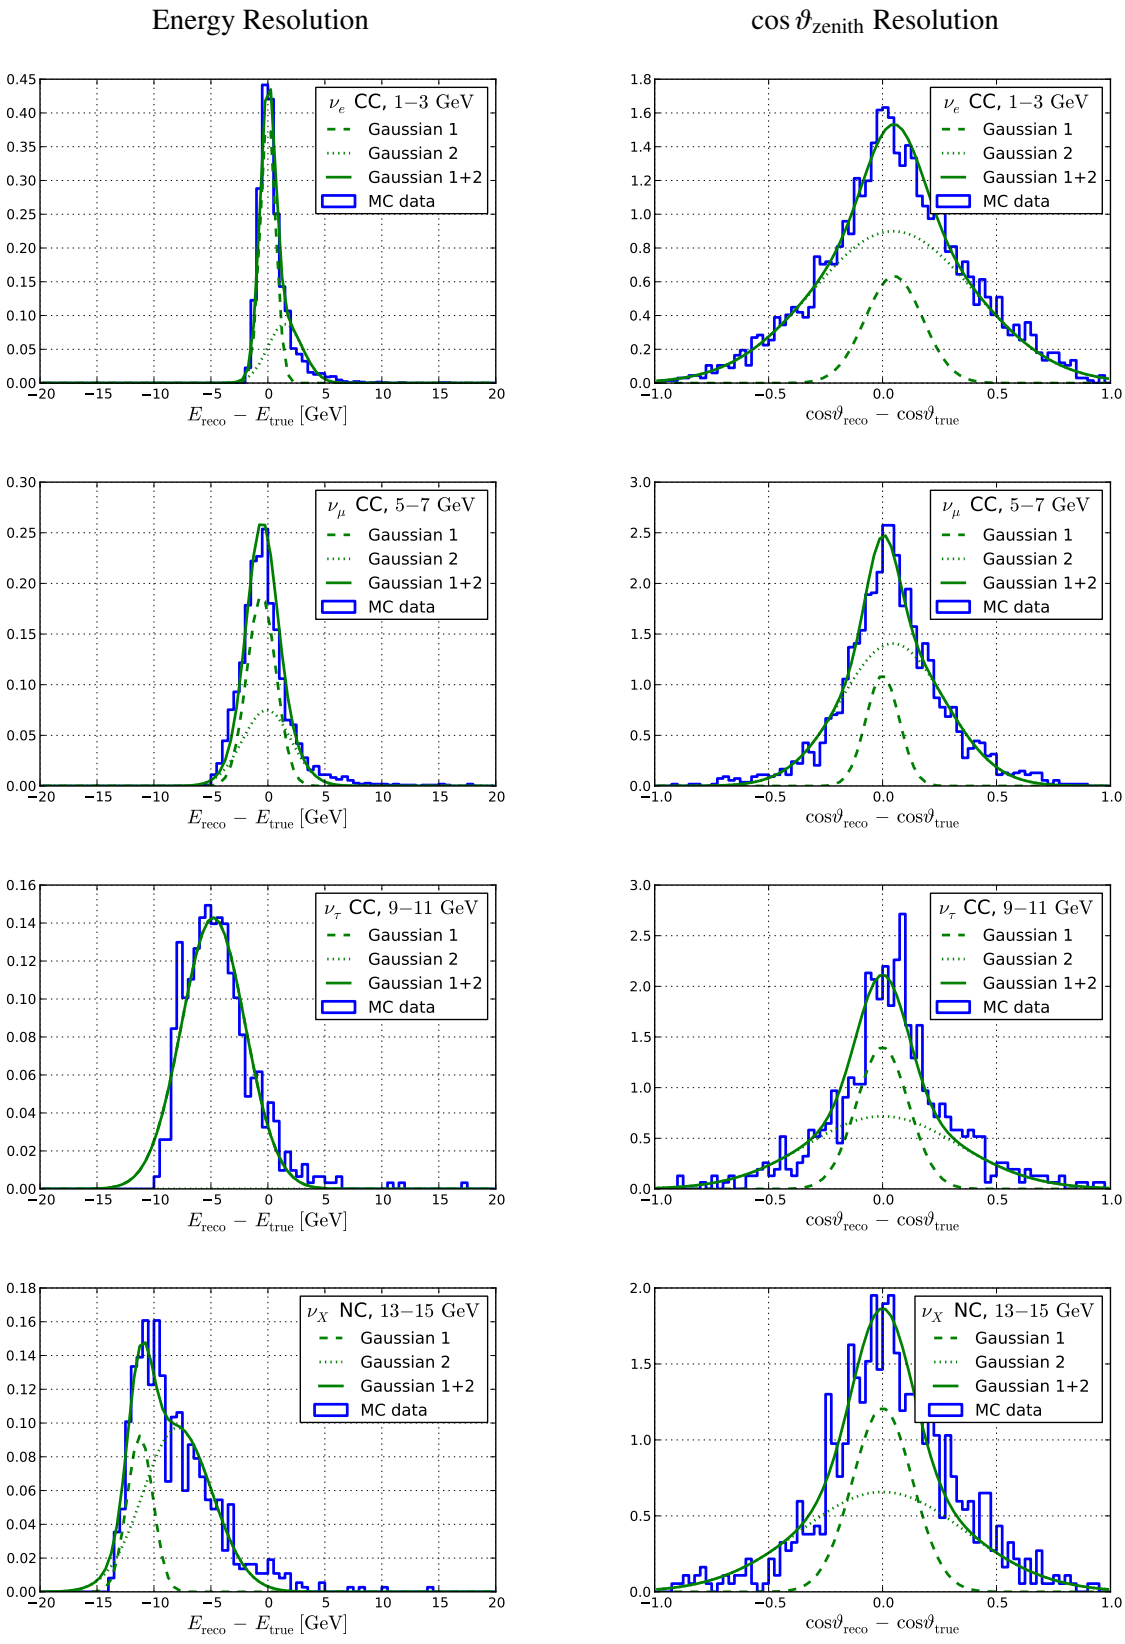

Figure 6.3: Examples for the parametrisations of the energy (left) and cos $\vartheta_{\text{zenith}}$ (right) reconstruction resolutions for (from top to bottom) ν_e, ν_μ, and ν_τ CC and ν_X NC events. Note the bias towards low reconstructed energies for ν_τ and NC.

in the resolution parametrisations above 20 GeV do not influence the calculation of the mass hierarchy sensitivity.

6.2.5 Particle Flavour Identification

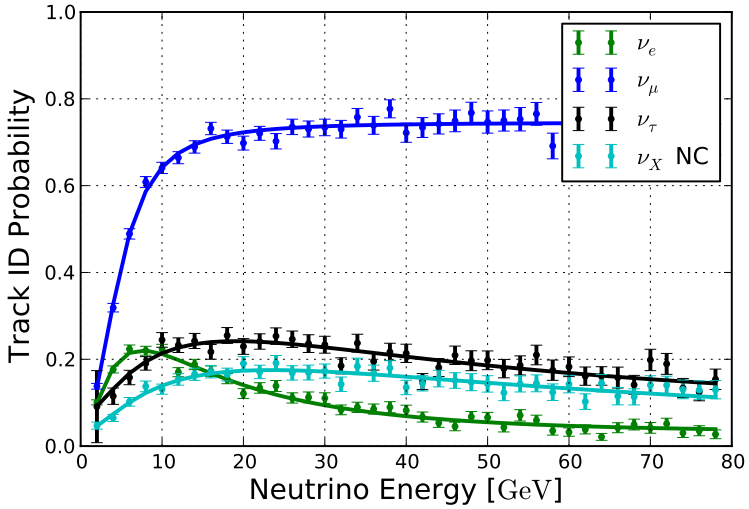

Figure 6.4: Track identification probability as function of energy in all four interaction channels. The straight lines show fits to the data.

The classification of the neutrino events into track-like and cascade-like events is according to their final score in the boosted decision tree described in Sec. 4.4.3. In the baseline detector settings, the decision is of binary nature, meaning that the probability to classify a given event as cascade is one minus the probability to classify it as a track.

Data points for the track identification probabilities in all channels as a function of the neutrino energy have been provided by [137]. These data were fitted with analytic functions, the fits are shown together with the data points in Fig. 6.4. The function definitions themselves are listed in App. E.

6.3 Results for the Baseline Geometry

Table 6.1: Uncertainties on all systematic parameters for the baseline detector model with three years of lifetime, ranked according to their impact on the mass hierarchy parameter h.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	3.43×10^{-1}	2.33×10^{-1}	2.52×10^{-1}	free
$r_{A_{\text{eff}},\,\nu-\bar{\nu}}$	8.9	0.00	4.73×10^{-2}	4.76×10^{-3}	1.47×10^{-1}	5.00×10^{-2}
ϑ_{13} [°]	8.0	8.93	4.64×10^{-1}	8.24×10^{-1}	3.67	4.68×10^{-1}
$n_{A_{\text{eff}}}$	7.4	0.00	1.94×10^{-2}	1.99×10^{-3}	1.94×10^{-2}	2.00×10^{-1}
ϑ_{23} [°]	3.2	3.86×10^{1}	4.67×10^{-1}	3.02×10^{-1}	3.97×10^{-1}	1.32
Δm^2_{31} [eV^2]	2.7	2.46×10^{-3}	6.49×10^{-5}	1.77×10^{-5}	1.10×10^{-4}	8.00×10^{-5}
$r_{\Phi,\,\nu_e-\nu_\mu}$	2.4	0.00	1.10×10^{-2}	5.12×10^{-3}	1.01×10^{-2}	5.00×10^{-2}
$s_{A_{\text{eff}}}$ [m^2/GeV]	1.1	0.00	2.19×10^{-4}	1.23×10^{-4}	1.82×10^{-4}	free
Δ_{PID} [GeV]	0.9	0.00	1.70×10^{-2}	1.58×10^{-2}	6.23×10^{-3}	5.00×10^{-1}
s_E	0.1	1.00	2.81×10^{-2}	7.66×10^{-3}	3.31×10^{-2}	5.00×10^{-2}

Using the settings described in the previous section, the Fisher matrix for PINGU can now be constructed with `PaPA`. The full, statistical, and systematic errors are for all parameters are listed in Tab. 6.1 for a nominal PINGU lifetime of three years. The parameters are ordered after their *impact* on the mass hierarchy parameter h, which is defined as the square of their correlation coefficient c_{ih} (6.7) with h. Note that for the baseline settings, the systematic parameter s_{PID} has been excluded as it is fully degenerate with $n_{A_{\text{eff}}}$: since the PID decision is binary, no channel of unidentified events exists and hence both parameters just evenly increase the overall number of events, effectively.

From the first line of Tab. 6.1, one can read off the expected significance of PINGU's mass hierarchy measurement by inverting the full error (see Sec. 6.1.3). This gives an expected significance of $2.9\,\sigma$ after three years. Looking at the statistical error alone, the significance increases to $4.3\,\sigma$, emphasising the important role systematic parameters are playing in the determination of the mass hierarchy.

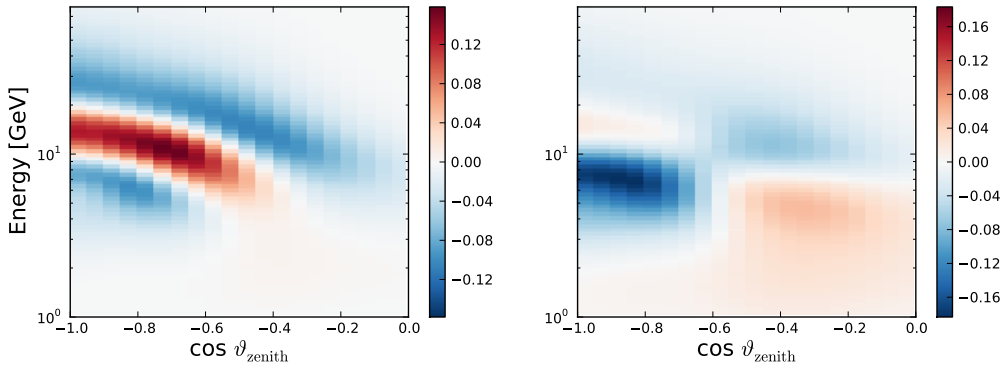

Figure 6.5: $\Delta\chi$ distribution in the track (left) and cascade (right) channels for the baseline settings.

Treating the track and cascade channels separately, the expected significances are $1.9\,\sigma$ with and $3.0\,\sigma$ without systematics for the cascade and $1.3\,\sigma$ ($3.1\,\sigma$) for the track channel, respectively[6]. Although the

[6] The full error listings corresponding to Tab. 6.1 can be found in App. F.1

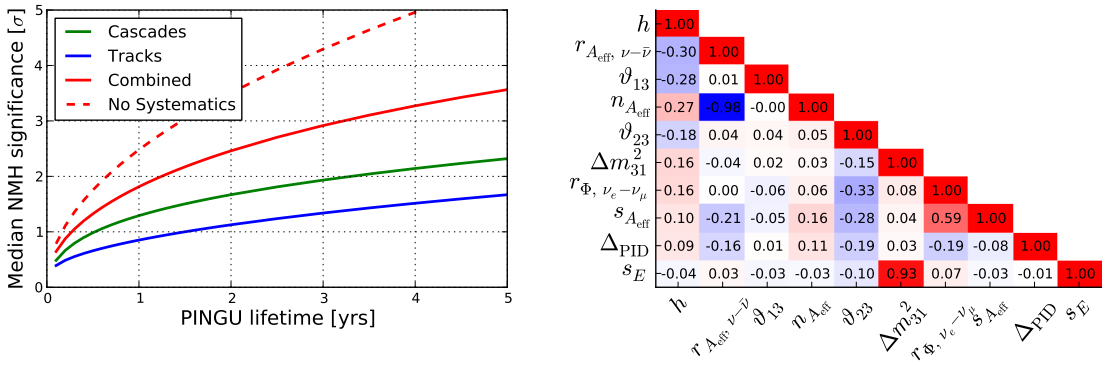

(a) (b)

Figure 6.6: (a) Evolution of PINGU's expected mass hierarchy significance with time and (b) full correlation matrix for PINGU for the baseline settings.

purely statistical significances are comparable, when taking systematics into account the significance for the cascades remains much higher than the track significance. The reason for this becomes obvious when looking at the $\Delta\chi$ distributions for both channels in Fig. 6.5. In the track channel, there are three distinct regions driving the expected significance. These regions are separated only by small margins and have alternating sign, while in the cascade channel most of the significance comes from one contiguous region. Together with the fact that the cascade channel has roughly three times higher event statistics than the track channel—62,000 vs. 22,000 neutrino events per year—this makes the cascade channel much more robust against the impact of systematic parameters.

In Fig. 6.6a, the significances for the individual channels and for their combination are shown as a function of time. The purely statistical significance is plotted as well, exhibiting a scaling with the square root of the lifetime as one would expect for a counting experiment where the relative error is proportional to $1/\sqrt{N}$. The actual significances including systematics increase much slower with time as a part of the accumulating statistics has to be "spent" in order to better constrain the systematic parameters. However the combined analysis of the two channels still gives a much higher sensitivity than simply adding the two individual channels in quadrature as one would do for two completely unrelated measurements of the same quantity. E. g. for a lifetime of three years as above, the quadratic sum of the track and cascade significances is $\sqrt{(1.3\,\sigma)^2 + (1.9\,\sigma)^2} \approx 2.3\,\sigma$, considerably lower than the $2.9\,\sigma$ for the combined analysis.

Finally one can look at the correlation matrix of PINGU. In contrast to the error listing for all parameters as in Tab. 6.1, here the interdependences between the parameters are in the focus. The graphical representation in Fig. 6.6b shows, for every combination of parameters, their correlation coefficient c_{ij}. Using these quantities instead of the entries σ_{ij} of the covariance matrix themselves has the benefit that due to their normalisation, cf. (6.7), the entries of the matrix are dimensionless and restricted to the range $[-1, +1]$, thus making them easier to interpret.

From the correlation matrix itself, several things can be learned. First, the hierarchy parameter is the one with strongest overall correlations. This emphasises the fact that the determination of the neutrino mass hierarchy is a very delicate measurement relying on a small effect, and that the inclusion of so many systematic parameters is indeed necessary to get a robust result.

Furthermore, two combinations of parameters stick out due to their very strong correlation. The first one is the relative normalisation of the effective areas for neutrinos and antineutrinos, $r_{A_{\text{eff}},\,\nu-\bar{\nu}}$,

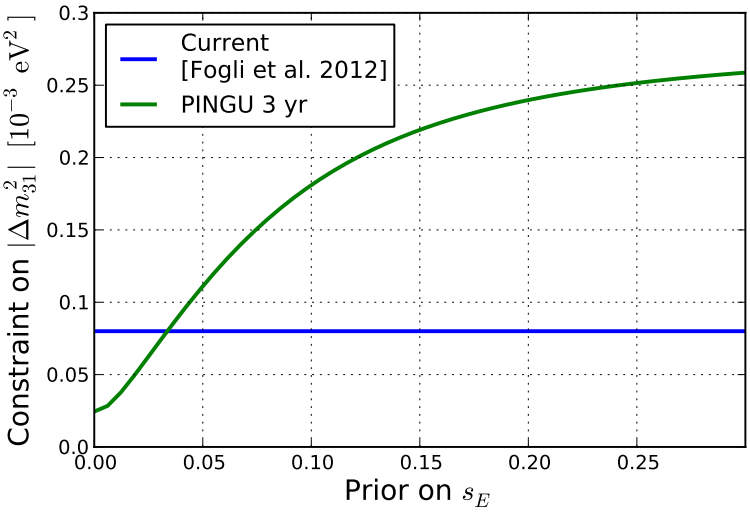

Figure 6.7: PINGU constraint on Δm^2_{31} as a function of the prior on the energy scale. No prior knowledge about Δm^2_{31} is assumed.

and the overall normalisation of all effective areas, $n_{A_{\text{eff}}}$. Their anticorrelation is obvious from their definition: $r_{A_{\text{eff}},\nu-\bar{\nu}}$ increases the number of neutrino events while decreasing the number of antineutrino events. Since PINGU cannot distinguish between those and the cross-section for neutrinos is higher approximately by a factor of two (see Fig. 2.7), the total number of events is increased, which can be compensated by decreasing $n_{A_{\text{eff}}}$. Only the MSW resonance oscillation causes an asymmetry between neutrinos and antineutrinos, such that the anticorrelation is not exactly -1.

The second strong correlation can be observed between the absolute value of the mass splitting Δm^2_{31} and the energy scale s_E. The value of Δm^2_{31} is determined from the position of the oscillation minimum in the track channel that is easy to spot in Fig. 5.4. In a two-flavour approximation, this can be described by equation (3.16), where the oscillation length is determined by the zenith angle. Then the size of Δm^2_{31} is inversely proportional to the neutrino energy at which the oscillation minimum appears. Thus a larger value of Δm^2_{31} can be compensated by increasing the energy scale accordingly.

This means that if PINGU is supposed to make a precise measurement of Δm^2_{31}, the energy scale has to be known very accurately. As the energy of a neutrino event is determined primarily from the number of detected photons[7], this means that the photon detection efficiency of the optical modules has to be well calibrated. This behaviour is illustrated in Fig. 6.7, where PINGU's self-contained constraint on the value of Δm^2_{31} is plotted against the prior put on s_E. To improve on the current limits, the energy scale has to be known with an accuracy of at least 3 %.

6.3.1 Measuring the Atmospheric Mixing Parameters

To determine the neutrino mass hierarchy, PINGU makes a precision measurement of the oscillations of atmospheric neutrinos. In fact, with more than 80,000 events recorded per year it will collect the largest sample of atmospheric neutrinos so far. After DeepCore has been established as a serious contributor to the global effort of characterising neutrino oscillations, PINGU will go even further into that direction and provide tight constraints on the values of Δm^2_{31} and in particular ϑ_{23}.

[7] The Cherenkov light output is directly proportional to the deposited energy.

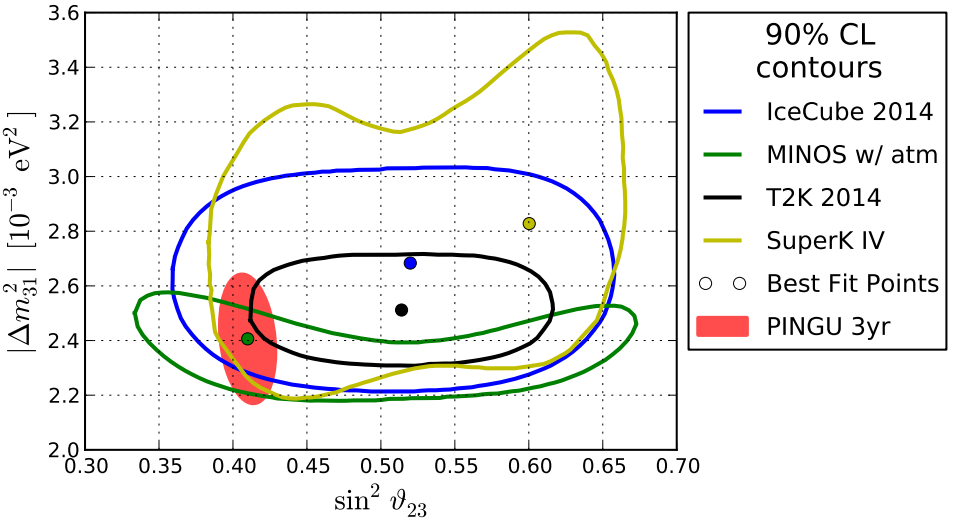

Figure 6.8: PINGU constraint on Δm^2_{31} and $\sin^2 \vartheta_{23}$ for the baseline settings. No prior knowledge about the two parameters is assumed. The latest constraints from the IceCube/DeepCore [62], MINOS [64], T2K [65], and SuperKamiokande [138] are shown for reference.

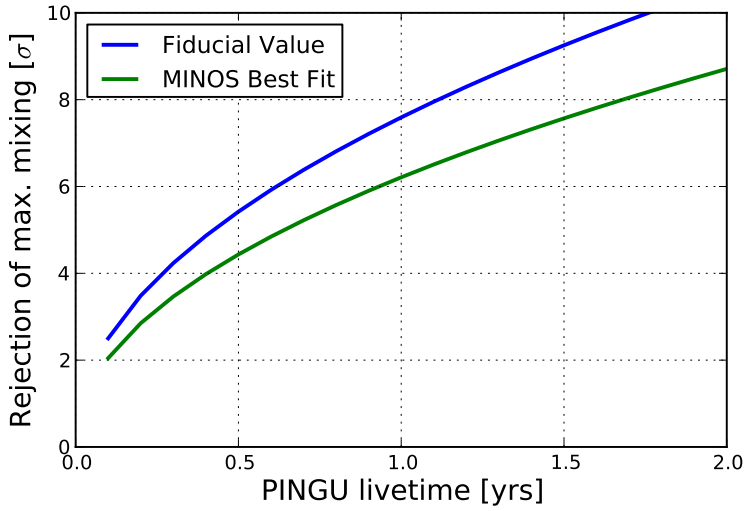

Figure 6.9: Confidence level of rejecting the maximal mixing case as a function of PINGU's lifetime for different true values of ϑ_{23}.

These constraints are shown in Fig. 6.8, along with the most recent confidence regions of current oscillation experiments. For PINGU, no priors have been put on Δm^2_{31} or ϑ_{23}, meaning that the displayed confidence ellipse comes from PINGU data alone and does not profit from external knowledge.

The most obvious feature of PINGU's confidence ellipse is its orientation, which is different from the other experiments: While it cannot constrain Δm^2_{31} any better than current experiments, the value of ϑ_{23} will be much more precise than any measurement available today. The reason for this is that the value of Δm^2_{31} is extracted from the position of the oscillation minimum in energy, which needs a precise calibration of the energy scale as shown above. This is much easier to achieve in experiments on a neutrino beam like MINOS or T2K where the beam energy is well-known. ϑ_{23} on the other hand is determined from the relative depth of the oscillation minimum as one can see from equation (3.15). Here PINGU profits from its wide energy range including control regions without oscillations and the large event statistics, such that especially the overall detector efficiency[8], which usually is difficult to constrain in beam experiments, only has minor impact.

From the error contours in Fig. 6.8 one can also see that the value of the mixing angle ϑ_{23} is not very tightly constrained yet. So far, all experiments' results are compatible with the maximal mixing value of $\vartheta_{23} = 45°$[9], however MINOS's and SuperKamiokande's best fit points are rather distant from maximal mixing.

Since PINGU will make a much more precise measurement of ϑ_{23} than any current experiment, it will likely be able to exclude the maximal mixing case with high confidence if the true value of ϑ_{23} is sufficiently offset from 45°. In Fig. 6.9 the confidence which PINGU can reject the maximal mixing hypothesis with is shown as a function of the experimental lifetime. Although this depends on the assumed true value of ϑ_{23}, it is obvious that PINGU is likely to exclude maximal mixing at the $5\,\sigma$ level already after the first year of data taking.

6.3.2 Impact of the Octant of ϑ_{23}

As shown in the previous section, currently it is still unclear whether ϑ_{23} is smaller or larger than 45°. The fiducial value chosen for the baseline analysis was the best fit point from the Fogli global fit [47], which is close to the MINOS best fit point and such gives a comparably small mixing angle.

Thus it is important to test whether the fiducial value of ϑ_{23} that is chosen in `PaPA` has any impact on the expected sensitivity to the mass hierarchy. The `PhysicsSimulation` step, i. e. the calculation of the oscillation probabilities, has been repeated for several different fiducial values of ϑ_{23}, all larger than the baseline one. The resulting oscillation probabilities were consecutively processed through the `DetectorSimulation`. For each of these settings, the sensitivity to the neutrino mass hierarchy has been evaluated for a lifetime of three years and plotted against the injected fiducial value for ϑ_{23} in Fig.6.10a.

As it turns out, the sensitivity increases considerably as ϑ_{23} takes larger values and even exceeds $5\,\sigma$ in the second octant, i. e. at values above 45°. This increase in sensitivity comes exclusively from the cascade channel, which already points to its origin: the oscillation of ν_μ to ν_e, which is mainly defined by this mixing angle. However the neutrinos propagate through the Earth, so this angle has to be replaced by the corresponding effective mixing angle in matter (3.32), which is further from maximal mixing than the vacuum value.

The relation between matter and vacuum value of ϑ_{23} is shown qualitatively in Fig. 6.10b. The choice of $A_{\mathrm{CC}}/\Delta m^2 = -0.5$ corresponds to a neutrino energy of $\approx 3.2\,\mathrm{GeV}$ at a matter density of $5\,\mathrm{g/cm^{-3}}$, the negative sign is due to A_{CC} being below zero as it caused by electrons with negative charge, see

[8] Represented by $n_{A_{\mathrm{eff}}}$ in `PaPA`.

[9] Since the "strength" of the mixing is proportional to $\sin^2 2\vartheta$, see (3.15), it becomes maximal for $\vartheta = 45°$.

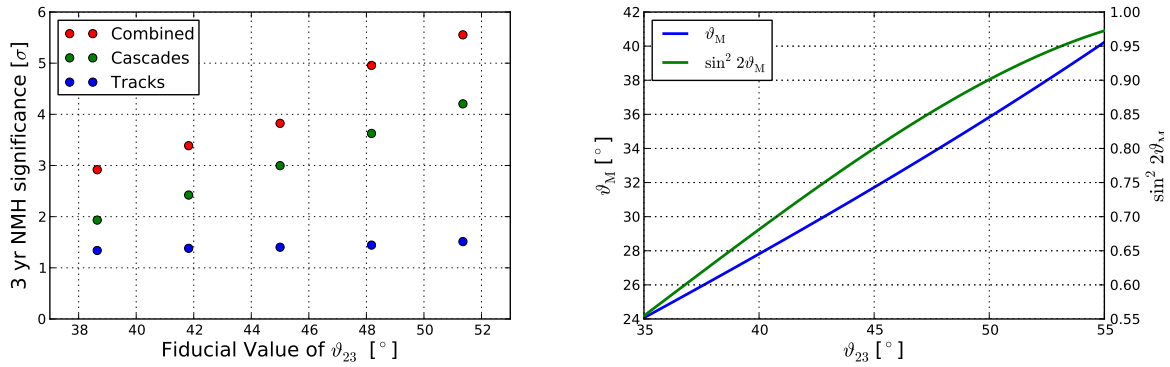

(a) (b)

Figure 6.10: PINGU three-year sensitivity to the neutrino mass hierarchy (a) and effective mixing angle in matter for $A_{\mathrm{CC}}/\Delta m^2 = -0.5$ (b) as a function of the fiducial value of ϑ_{23}.

(3.25). So due to the effective reduction of ϑ_{23} in matter, the amplitude of the oscillation, proportional to $\sin^2(2\vartheta_{\mathrm{M}})$, keeps growing even beyond $\vartheta_{23} = 45°$, where for oscillations in vacuum the amplitude would decrease again. Consequently, increasing the overall scale of $\nu_\mu \rightarrow \nu_e$ oscillations results in an increased absolute difference between the cascade event histograms for normal and inverted mass hierarchy and hence a better sensitivity for the hierarchy determination.

A beneficial side effect of the effective reduction of ϑ_{23} in matter is that it causes the relation of the parameter, ϑ_{23}, and the observable, the oscillation probability proportional to $\sin^2(2\vartheta_{\mathrm{M}})$, to be in the linear regime, as shown in Fig. 6.10b as well. This makes the evaluation in the Fisher matrix formalism well justified (cf. Sec. 6.1.2), whereas for vacuum oscillations the oscillation probability $\sin^2(2\vartheta_{23})$ becomes an increasingly non-linear function close to the maximal mixing value of ϑ_{23}.

In addition, the symmetry between values of ϑ_{23} in the first and second octant, mirrored at 45°, gets broken. This symmetry can be found in most of the confidence contours in Fig. 6.8 and affects all experiments that observe vacuum oscillations, making it difficult to determine the actual octant of ϑ_{23} even if the best fit point does not correspond to maximal mixing. Yet PINGU does not see said symmetry, thus its exclusion of maximal mixing discussed in the previous section does actually include a definite measurement of the octant. In other words, the likelihood maximum indicated by the PINGU ellipse in Fig. 6.8 does *not* have a counterpart at a corresponding value of ϑ_{23} in the second octant.

6.3.3 Fiducial Value of the Mass Hierarchy

Thus far, the *true* neutrino mass hierarchy was assumed to be inverted. This means that the fiducial value for the mass hierarchy used as an input to `PaPA` was normal. This may sound paradoxical at first, yet it corresponds to the actual analysis that will be done in PINGU:

In order to "detect" e. g. an inverted hierarchy, the normal hierarchy has to be excluded. This means to compare the data that have been measured, represented by the histograms resulting from the best fit IH oscillations and the fiducial detector settings, to all possible realisations of the NH scenario. These are simulated by varying all systematic parameters around the best fit NH model, which is what happens when the fiducial value of the hierarchy parameter in `PaPA` corresponds to the NH case.

Now it is of course important to know how the assumed true mass hierarchy affects the significance of the mass hierarchy determination. If it would drop considerably when assuming a normal ordering,

PINGU could not claim to provide a definite measurement of the mass hierarchy since the inverted hierarchy case could never be excluded.

Calculating the significance for a true normal mass hierarchy, i. e. choosing the IH value of Δm^2_{31} from Tab. 3.1 as fiducial value in PaPA, the expected three-year significance for the baseline detector model is $3.1\,\sigma$. This moderate increase by $0.2\,\sigma$ w. r. t. the IH case originates predominantly in the cascade channel, whose individual significance increases from $1.9\,\sigma$ to $2.2\,\sigma$. The full error listings are shown in App. F.2. The purely statistical significances do not change w. r. t. to the IH true case as they only depend on the difference between both models.

The reason for the increased significance in the cascade channel is that the most significant feature in the analysis histograms is the region of negative $\Delta\chi$ values in the cascade channel (Fig. 6.5 right) at energies between 5 and 10 GeV and $\cos\vartheta_{\text{zenith}} < -0.6$. According to the definition of $\Delta\chi$ (6.14), in this region more events are expected for true IH than for true NH, $N_{\text{IH}} > N_{\text{NH}}$. Since the absolute difference between N_{IH} and N_{NH} is the same regardless of which hierarchy is true, the relative difference to the IH expectation is larger if NH is true than vice versa.

So assuming true inverted hierarchy as well as ϑ_{23} in the first octant, i. e. being smaller than 45°, is in fact a conservative choice. As shown in this and the previous section, PINGU's expected sensitivity to the neutrino mass hierarchy is higher if the mass hierarchy is in truth normal and especially for larger values of ϑ_{23}. Consequently, those settings will be used as well for the various scenarios studied in the following sections.

6.3.4 High-Purity Event Classification

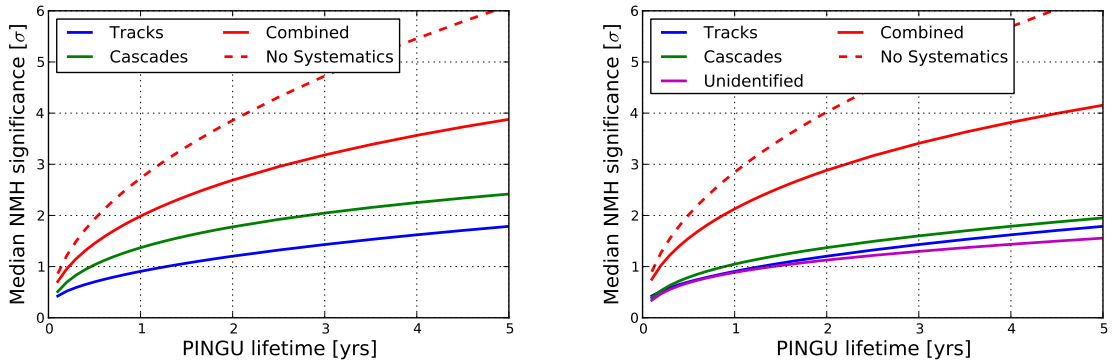

(a) (b)

Figure 6.11: Individual and combined mass hierarchy sensitivity for the baseline model with V15 particle identification with two (a) and three (b) channels.

Thus far, the particle identification has been considered as a binary decision, classifying each event as either track-like or cascade-like, based on a BDT score (cf. Sec. 4.4.3). Yet PaPA offers the option to have separate event selections for tracks and cascades, such that e. g. the cascade channel is populated only by events that actually "look like cascades" rather than events that just "do not look like tracks". To make sure that a given event does not end up in both samples, the decision should be based on the same BDT, but now with two cut values. Then the scores below the lower cut would then e. g. correspond to cascades while the ones above the higher cut would be classified as tracks. All events with scores between the two cut values would end up in the "unidentified" sample.

Such an event selection does only exist for an anterior geometry version of PINGU, V15, that was used for intense systematic and reconstruction studies before optimising the geometry for maximal sensitivity to the mass hierarchy. The identification efficiencies are plotted in Fig. E.1, with the corresponding function definitions listed in App. E.2. Note that the track identification efficiency is much better than for the baseline geometry V36, shown in Fig. 6.4, since due to an error in the simulation configuration all Monte Carlo for V15 was produced without generating noise hits on the PINGU optical modules, resulting in very clear event signatures.

The resulting significances are shown as a function of time in Fig. 6.11. As expected from the improved track identification w. r. t. the baseline settings, already the two-channel significance (Fig. 6.11a) is enhanced compared to the baseline. When opening the third channel of unidentified events (Fig. 6.11b), the cascade channel loses sensitivity due to the reduced statistics and is now at approximately the same level as the track channel, which remains unchanged. The unidentified events alone yield a significance on the scale of tracks and cascades individually as well, yet they suffer slightly stronger from systematic effects[10]. The full error listings are tabulated in App. F.3.

This behaviour is expected as the event selection for tracks stays the same as in the two-channel settings while the cascade sample gets split into two subsamples. Comparing the two corresponding error listings—Tabs. F.7 for the case of all "non-track" events being combined into one channel as in the baseline model and F.10 for having separate cascade and unidentified channels and combining the respective Fisher matrices—the statistical error on the mass hierarchy parameter h is slightly higher when the cascade channel is split due to larger relative statistical errors. The systematic uncertainty on h on the other hand is reduced for the separate analysis, leading to the expected overall increase in sensitivity. Yet this increase is rather moderate, indicating that there is not much gain to be expected from moving further into this direction[11].

For the remaining parameters one can see that e. g. the impact of the relative flux scaling $r_{\Phi,\,\nu_e-\nu_\mu}$ is reduced for the separated analysis, driven by the high-purity cascade sample that has only a small contamination of ν_μ events. Additionally, now the parameter s_{PID}, scaling the overall cascade and track identification efficiencies, comes into play for the three-channel analysis, as it is no longer degenerate with the overall normalisation $n_{A_{\mathrm{eff}}}$ (at least for the unidentified channel), yet it has no noticeable impact on the expected mass hierarchy significance.

6.3.5 The Missing Monte Carlo Effect

In Sec. 5.2.1 it was pointed out that the "conventional" approach to evaluate PINGU's sensitivity to the neutrino mass hierarchy, i. e. to propagate the full Monte Carlo detector simulation through the analysis, is not only computationally challenging, but introduces a bias towards too high significances if the amount of available MC statistics is insufficient. This effect can be demonstrated using an event reconstruction based on individual Monte Carlo events processed for the earlier PINGU geometry V15.

For this geometry, enough MC events have been generated and processed[12] to use them for the reconstruction in `PaPA` directly. The the reconstruction kernels were created by histogramming the MC events as described in Sec. 5.2.2, i. e. now option (b) is used for calculating the reconstructed histograms. This is repeated multiple times, reducing the fraction of the total number of events that is actually filled into

[10] This is apparent from the fact that the curves for tracks and unidentified events lie on top of each other for lifetimes below one year, but the track significance keeps to increase stronger as more statistics gets accumulated.

[11] As an extreme, one could imagine to open a third axis on the event histogram by introducing a Björken-y-like variable, e. g. the ratio of cascade and track energy reconstructed by HybridReco/MultiNest.

[12] For the most important channels, ν_e and ν_μ, $\approx$ 270,000 and $\approx$ 340,000 events are available, respectively. In addition, $\approx$ 40,000 ν_τ and $\approx$ 80,000 ν_X NC events could be used for this study.

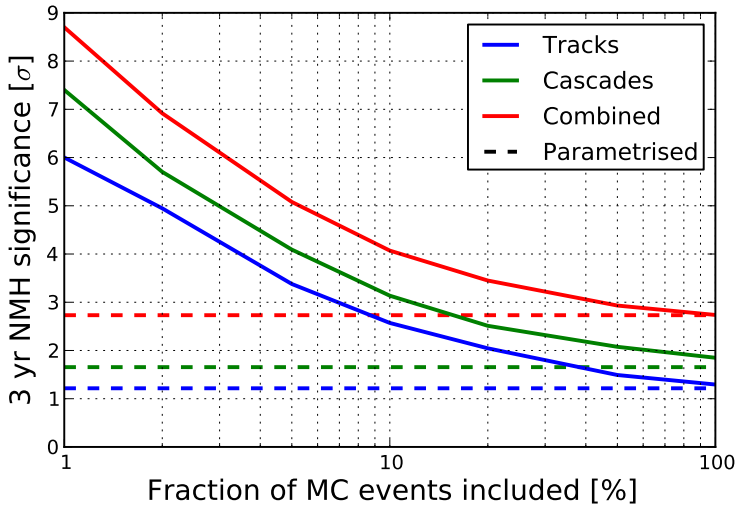

Figure 6.12: PINGU's sensitivity to the neutrino mass hierarchy with reconstruction from MC data for geometry V15. The result for a reconstruction parametrisation from the same data is shown for reference.

the kernels from 100 % to 1 %. Additionally, a reconstruction parametrisation is fitted to the full dataset, which is given in App. D.2.

In Fig. 6.12, the resulting mass hierarchy significances for three years of PINGU lifetime are shown as a function of the fraction of MC data used for the reconstruction. Only for the full amount of statistics a relatively stable result is achieved, which is also consistent with the parametrised reconstruction.

The reason for this behaviour can be found when inspecting the $\Delta\chi$ distributions for the different settings, shown in Fig. 6.13. When using the full sample of MC events (middle panel), the pattern generated with the parametrised reconstruction is reproduced quite accurately, reflected in the resulting significance values being in agreement. For the extreme case of populating the reconstruction kernels with only 1 % of the available events, the pattern imprinted by the mass hierarchy difference has almost completely vanished and is replaced by random fluctuations.

However the scale of the $\Delta\chi$ values is much higher in the histograms for very low MC statistics, as one can see from the range of the colour scale, which is linked to the highest individual bin value. The reason for this is that if the kernel tables are underpopulated, the reconstruction resolution effectively gets enhanced: When the statistics is sufficient, the hierarchy sensitivity of a single bin in the histogram of true energy and zenith angle is smeared out into the surrounding bins. If, in the extreme, only one MC event falls into this bin, its sensitivity is not smeared out, but only moved to a single other bin in the reconstructed histogram (the one where the single event is reconstructed in). The full sensitivity from the true histogram being retained at one point now increases the total significance as it is given by quadratically summing all individual bins.

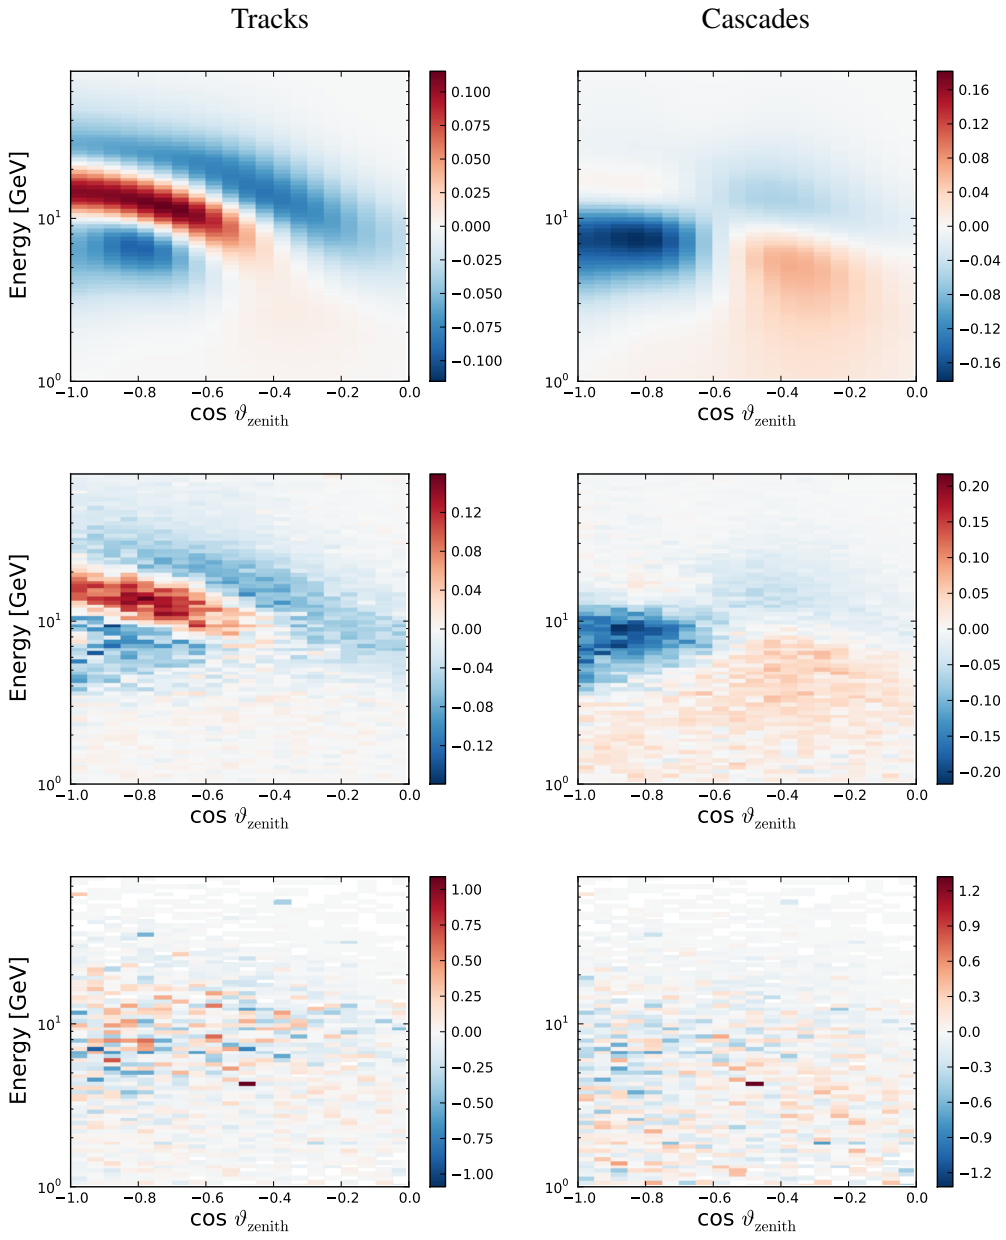

Figure 6.13: $\Delta\chi$ distribution in the track (left) and cascade (right) channels for (top to bottom) the reconstruction parametrisation based on geometry V15 and the reconstruction directly from 100 % and 1 % of the Monte Carlo events available for V15.

6.4 Effects of Advanced Optical Modules

In Sec. 4.5, two different concepts to substantially improve the optical modules have been introduced, prototypes of both are supposed to be deployed in PINGU: the mDOM that has several small PMTs instead of a single large one, and the WOM, where the effective area of a single small PMT is enhanced drastically by passive components. In this chapter, PINGU will be modelled as it was built completely from these next-generation optical modules. Then the possible benefits for the experiment's outcome can be analysed.

6.4.1 WOM: Increasing the Photon Statistics

With its enhanced photosensitive area, the WOM will primarily increase the number of photons detected per event. Since this is the driving parameter for event reconstruction quality—the more photons are detected, the more information is available for the reconstruction algorithm—more precise reconstructions are expected especially at low energies, where photon statistics are scarce.

To simulate the improvement resulting from a photon count increased by a factor f_{ph} in `PaPA`, the *absolute* widths of the energy reconstruction functions that are usually given by linear functions of the form

$$\sigma_E(E) = aE + b \quad , \tag{6.21}$$

where a and b are constants determined by a fit (see App. D), are replaced by

$$\sigma'_E(E) = aE + b/f_{\text{ph}} \quad . \tag{6.22}$$

The rationale behind this is that the *relative* energy resolution is assumed to be a function of the number of detected photons, which in turn is directly proportional to the neutrino energy[13]. Thus increasing the number of detected photons by f_{ph} means that an an event with the energy E will now have the relative energy resolution of an event of energy $f_{\text{ph}} \cdot E$:

$$\frac{\sigma'_E(E)}{E} = \frac{\sigma_E(f_{\text{ph}}E)}{f_{\text{ph}}E} = \frac{af_{\text{ph}}E + b}{f_{\text{ph}}E} \tag{6.23}$$

$$\Rightarrow \sigma'_E(E) = aE + b/f_{\text{ph}} \tag{6.24}$$

For the $\cos\vartheta_{\text{zenith}}$ resolution, the modification is more straightforward as one can simply read out the parametrisations for the absolute widths at an energy increased by the factor f_{ph}:

$$\sigma'_{\cos\vartheta}(E) = \sigma_{\cos\vartheta}(f_{\text{ph}} \cdot E) \quad . \tag{6.25}$$

In addition, the threshold for passing PINGU's event selection is expected to become lower with an increasing number of photons per neutrino energy, enhancing the event statistics at low energy. This effect is mimicked by scaling the effective area at a given energy E by the ratio of the selection efficiencies ε_{sel} (as shown in Fig. 4.8) at E and $f_{\text{ph}} \cdot E$:

$$A_{\text{eff}}{}'(E) = A_{\text{eff}}(E) \cdot \frac{\varepsilon_{\text{sel}}(f_{\text{ph}} \cdot E)}{\varepsilon_{\text{sel}}(E)} \tag{6.26}$$

The scaling is applied for all flavours except from ν_τ and $\bar{\nu}_\tau$ CC events, as there the main feature is the kinematic cutoff due to the large mass of the tau lepton (cf. Sec. 4.4.2), which of course is independent

[13] The number of Cherenkov photons is proportional to the total deposited energy, see Sec. 2.3.3.

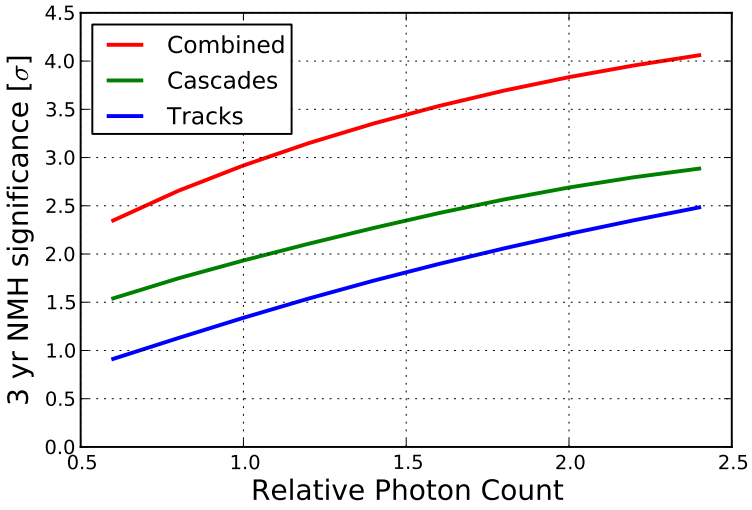

Figure 6.14: Relative expected three-year significance for the mass hierarchy as a function of f_{ph}.

from the performance of the optical modules used in the detector.

This gives a conservative estimate as only effects after the actual triggering of the detector are taken into account. Since the trigger condition essentially is a certain number of photon hits being registered within a short time window, the trigger threshold itself does sink as well, enhancing the overall effect. Yet as the available amount of MC events before triggering is insufficient to quantify this, the reduced trigger threshold will be neglected.

The median three-year significance for the mass hierarchy is shown as a function of f_{ph} in Fig. 6.14. As expected, the mass hierarchy significance increases with the photon statistics as a result of the improved resolutions and enhanced effective area. When the photon statistics are doubled, i. e. $f_{\text{ph}} = 2$, it reaches $3.8\,\sigma$, an increase of about 30 %. Looking at the track and cascade channels separately, the tracks are profiting slightly more from the higher number of photons. This emphasises the fact that the track channel is more affected by systematics (cf. Sec. 6.3), which can be resolved better if the events are reconstructed more precisely.

6.4.2 mDOM: Eliminating the Noise

The segmented photosensitive area of the mDOM will allow for an almost perfect rejection of noise hits. Thus the event signatures will be clearer, leading to an improved reconstruction and particle identification. In `PaPA`, this can be modelled by using the parametrised reconstruction functions and event classification extracted from the Monte Carlo data for geometry V15, which was erroneously produced without simulating noise. As the difference to geometry V36 in terms of module density—which is crucial for the reconstruction performance—is rather small[14], the resulting parametrisations, written out in App. D.2 and E.2, respectively, should give a good approximation of what to expect for the baseline geometry V36 with fully suppressed noise.

The resulting significance as function of lifetime and the correlation matrix are shown in Fig. 6.15. Comparing this to the results for the baseline settings in Fig. 6.6, only a very small improvement can

[14] Geometry V36 has a horizontal string spacing of 22 m and a vertical module spacing of 3 m, while in V15 string and module spacing are 20 m and 5 m, respectively.

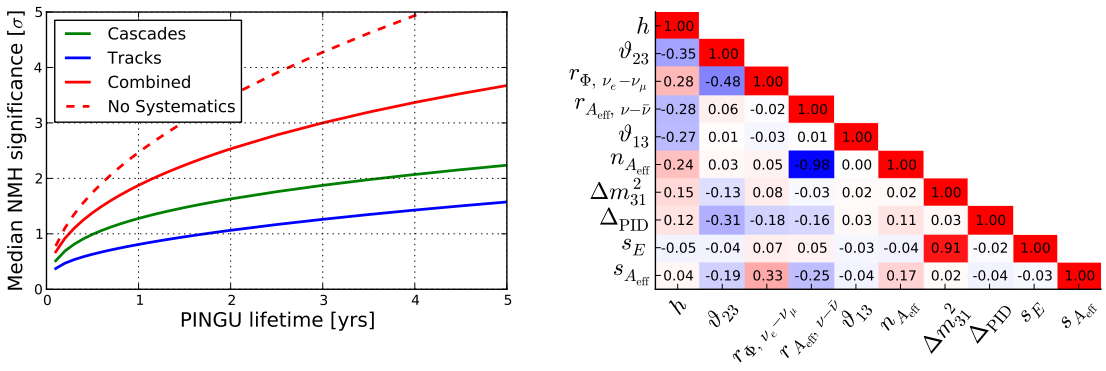

(a) (b)

Figure 6.15: (a) Evolution of PINGU's expected mass hierarchy significance with time and (b) full correlation matrix for PINGU assuming reconstruction and particle identification as in geometry V15, i. e. no module noise.

be observed. Looking at the correlation matrix, the most obvious change is the increased impact of the mixing angle ϑ_{23} and the flux ratio $r_{\Phi,\nu_e-\nu_\mu}$ on the mass hierarchy.

The reason for this is that especially the cascade channel profits from the improved particle identification, increasing its significance without systematics from $3.0\,\sigma$ to $3.3\,\sigma$[15], which is reflected by the increased significance scale of the $\Delta\chi$ distribution in the cascade channel, shown in Fig. 6.16. As such a feature basically means more ν_e-like events being detected, it can of course be mimicked by a wrong relative normalisation of the primary ν_e and ν_μ fluxes or a stronger oscillation from ν_μ to ν_e, determined by the value of ϑ_{23}. Unfortunately, these two parameters are insufficiently constrained by themselves, thus they can partly "absorb" the gain in sensitivity to the mass hierarchy.

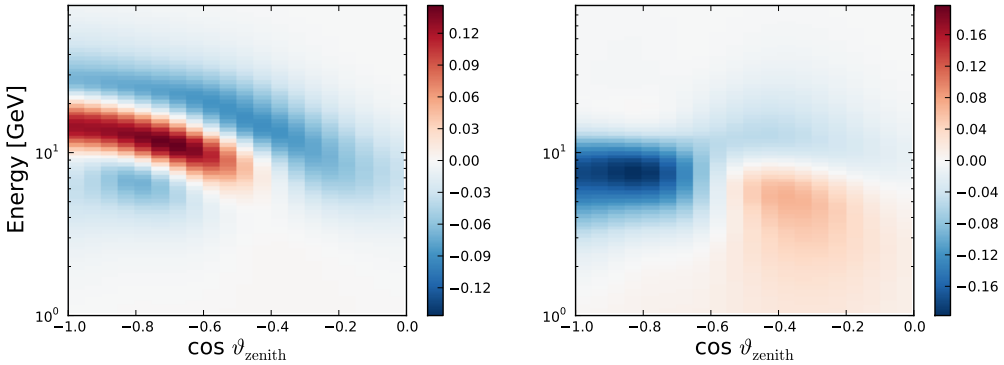

Figure 6.16: $\Delta\chi$ distribution in the track (left) and cascade (right) channels assuming reconstruction and particle identification as in geometry V15, i. e. no module noise.

[15] The full error listings are shown in App. F.4

6.5 Combining PINGU with JUNO

6.5.1 The JUNO Experiment

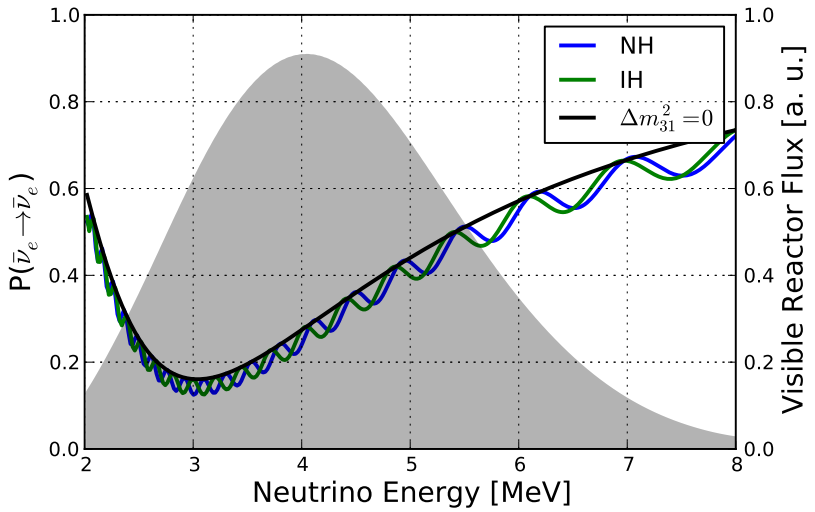

Figure 6.17: $\bar{\nu}_e$ survival probability for a baseline of 50 km, overlaid with the un-oscillated nuclear reactor spectrum as it would be detected by JUNO (including detector acceptance).

As already briefly discussed in Sec. 3.4.4, JUNO[16] is a neutrino experiment under construction in China's Guangdong province. With its sensitive region at low MeV energies, it is able to detect supernova and geoneutrinos, yet its main target are reactor neutrinos from the Taishan and Yangjiang nuclear power plants to be erected in ≈ 50 km distance each. These $\bar{\nu}_e$'s can be detected via inverse beta decay:

$$ {}^{A}_{Z}\mathrm{X} + \bar{\nu}_e \quad \rightarrow \quad {}^{A}_{Z-1}\mathrm{Y} + e^+ \tag{6.27} $$

The $\bar{\nu}_e$ undergo oscillations that are mostly due to the smaller mass splitting Δm^2_{21}, but have a fast modulation due to Δm^2_{31}. The vacuum $\bar{\nu}_e$ survival probability[17] can be expressed analytically:

$$ \begin{aligned} P(\bar{\nu}_e \to \bar{\nu}_e) \quad = \quad 1 \quad &- \quad \cos^4(\vartheta_{13})\,\sin^2(2\vartheta_{12})\,\sin^2\left(\Delta m^2_{21}\frac{L}{4E}\right) \\ &- \quad \cos^2(\vartheta_{12})\,\sin^2(2\vartheta_{13})\,\sin^2\left(\Delta m^2_{31}\frac{L}{4E}\right) \\ &- \quad \sin^2(\vartheta_{12})\,\sin^2(2\vartheta_{13})\,\sin^2\left(\Delta m^2_{32}\frac{L}{4E}\right) \end{aligned} \tag{6.28} $$

with $\Delta m^2_{32} = \Delta m^2_{31} - \Delta m^2_{21}$. As shown in Fig. 6.17, for a baseline of $L = 50$ km this is dominated by the slow Δm^2_{21} oscillation, corresponding to the first term in (6.28). On top of that is a pattern of rapid oscillation originating from the interference of the second and third term of (6.28). The exact shape of this pattern, especially the position of the local minima and maxima, depend on the mass hierarchy which changes the relative sizes of Δm^2_{31} and Δm^2_{32}[18].

[16] Short for Jiangmen Underground Neutrino Observatory.
[17] Matter effects can be neglected here.
[18] Their signs are irrelevant as the outer $\sin^2$ is symmetric about zero.

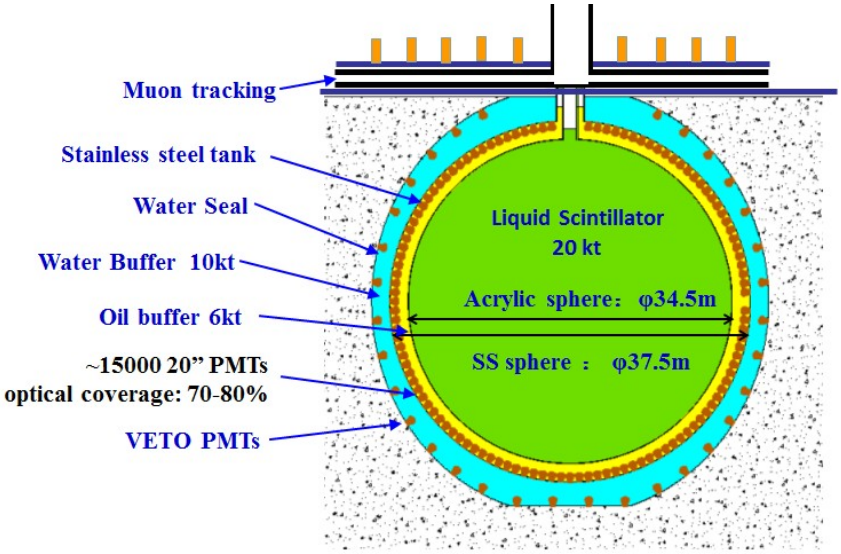

Figure 6.18: Layout of the JUNO detector. Figure taken from [73].

Taking the detector acceptance into account, the expected reactor neutrino spectrum [139] is very suitable to observe this signal, as one can see in Fig. 6.17 as well. However, a very precise energy reconstruction is needed to actually resolve the spectral features associated with the neutrino mass hierarchy.

With the detector setup shown in Fig. 6.18, a relative resolution of $3\,\%/\sqrt{E[\mathrm{MeV}]}$ is aimed for. The target volume is filled with 20 kt of liquid scintillator to enhance the photon output. This essentially destroys any directionality of the event signatures, however, no information about the arrival direction of the neutrinos is needed as the baseline for the oscillation is known[19]. The inner surface of the volume is covered with ≈ 15000 20" PMTs, resulting in a photocoverage of at least 70 %. Tagging and vetoing of atmospheric muons is ensured by scintillator tiles on top of the detector and a 10 kt water Cherenkov detector surrounding the target volume. Natural radioactivity is buffered in a layer between muon veto and active region filled with 6 kt of mineral oil [73].

If this precision can indeed be realised, JUNO claims to be able to determine the neutrino mass hierarchy with a significance of $\approx 3.3\,\sigma$ with a total of 10^5 recorded neutrino events, corresponding to six years of data taking [72, 73, 139]. In the following sections a modification of `PaPA` is described, aiming at a detector simulation for JUNO instead of PINGU in order to reproduce the reported sensitivity.

6.5.2 Simulating JUNO with `PaPA`

The major difference between the observable signals in PINGU and JUNO is that PINGU will record a two-dimensional histogram in $(E, \cos\vartheta_{\mathrm{zenith}})$ while in JUNO only the energy spectrum is measured. Thus only one very narrow zenith bin with edges at $\cos\vartheta_{\mathrm{zenith}} = [-0.0039245, -0.0039235]$ is simulated in `PaPA`, corresponding to a baseline of 50 km. In energy, 300 bins between 2 MeV and 8 MeV are used. No oversampling is applied as the oscillation probabilities are already sufficiently smooth (see Fig. 6.17).

Since only the $\bar{\nu}_e$ survival probability is relevant for JUNO, which can be calculated analytically ac-

[19] In PINGU, the $\cos\vartheta_{\mathrm{zenith}}$ information is needed to infer the location where the neutrinos were generated in the Earth's atmosphere and hence the distance they have travelled.

cording to (6.28), the `PhysicsSimulation` has been extended by a module doing exactly this analytical calculation. Avoiding the numerical solution of the Schrödinger equation, the `PhysicsSimulation` is sped up dramatically.

For the `DetectorSimulation`, the software itself is not changed, however the inputs of course have to model the JUNO detector. The neutrino flux is adopted from [139] and stored in a table of the same format as the atmospheric flux tables provided by [129]. This flux, shown as an overlay in Fig. 6.17, has the detector acceptance already folded in, but is only given in arbitrary units. Thus, the effective area is set to a constant value of $1\,\mathrm{m}^2$ for $\bar{\nu}_e$ CC and zero for all other interaction channels while the analysis histograms will be normalised to 10^5 $\bar{\nu}_e$ events before analysis.

The directional reconstruction is parametrised by a single Gaussian with a fixed width of 10^{-5} in $\cos\vartheta_{\mathrm{zenith}}$, meaning that all events will stay in place. As there is only one bin in $\cos\vartheta_{\mathrm{zenith}}$, migration of events is impossible in any case.

The energy reconstruction is represented by a single Gaussian as well, with mean and width given by

$$\mu(E) = E - 0.8\,\mathrm{MeV} \tag{6.29}$$

$$\sigma(E)/E = 4\,\%/\sqrt{E[\mathrm{MeV}] - 0.8}\quad . \tag{6.30}$$

The shift in energy reflects the fact that the visible energy in an inverse beta decay is smaller than the initial neutrino energy. 511 keV are needed to create the e^+, which in turn stops to emit Cherenkov light once it falls below the Cherenkov threshold (2.23), corresponding to a kinetic energy of $\approx 280\,\mathrm{keV}$ depending on the optical medium. These two contributions sum up to $\approx 0.8\,\mathrm{MeV}$. The relative energy resolution refers to the visible energy as well. Additionally, it has been deteriorated w. r. t. the published specifications to be $4\,\%/\sqrt{E_{\mathrm{vis}}[\mathrm{MeV}]}$. This reflects the fact that the nuclear power plants used as neutrino sources have several reactor cores which are up to 500 m apart from each other. This introduces an uncertainty of 1 % on the baseline L, which is equivalent to an additional uncertainty of 1 % in the energy reconstruction as the relevant quantity is L/E.

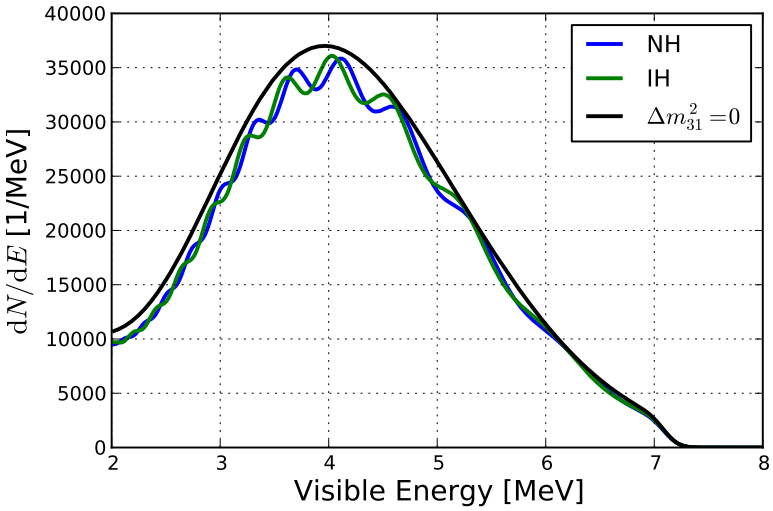

Figure 6.19: Expected event spectrum in JUNO including all detector effects.

Finally, in JUNO there is no discrimination between cascade-like and track-like events as only $\bar{\nu}_e$-induced inverse beta decays are expected. Hence, all events are classified as cascades and the track

channel is not considered for the analysis. The resulting observed event distribution for the normal and inverted hierarchy case as well as for oscillations with $\Delta m^2_{31} = 0$ is shown in Fig. 6.19.

6.5.3 Preparing the JUNO Signal for Fisher Matrix Analysis

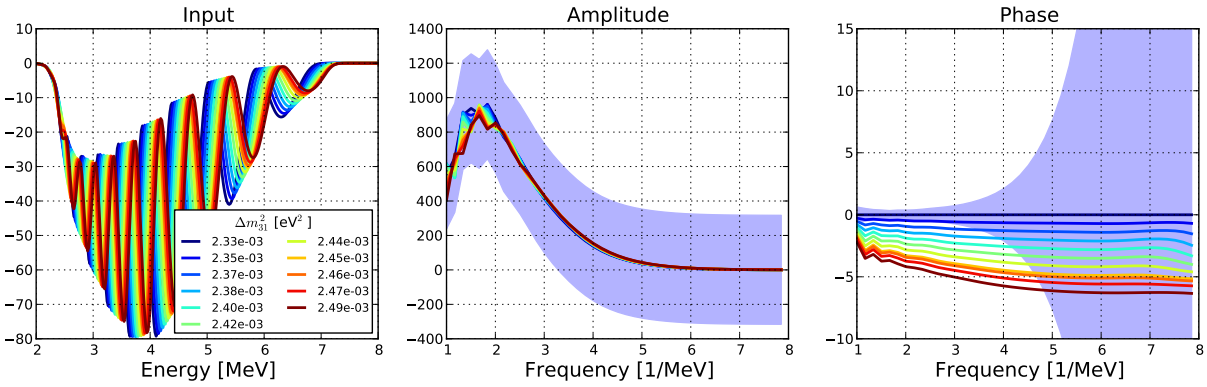

Figure 6.20: Linearisation of the detector response to Δm^2_{31} via Fourier transformation. For details, refer to the text.

In contrast to PINGU, the analysis histograms for JUNO as shown in Fig. 6.19 cannot be used for the Fisher matrix analysis directly. The reason for this are the rapid oscillations of the event count as a function of energy that are the signal to be observed. This means that for a given bin in energy, the expected number of events is usually not a linear function of the systematic parameters.

In order to apply the Fisher matrix formalism, JUNO's response to the systematic parameters has to be linearised. As the signal itself is oscillatory, the natural choice is to apply a Fourier transformation before analysis. To assure an input signal that smoothly fades out to zero at the edges of the considered energy interval, the expected signal for $\Delta m^2_{31} = 0$ (black line in Fig. 6.19) is subtracted from the actual event rates and an exponential cut-off at low energies applied to the result.

This quantity is shown in the left panel of Fig. 6.20, the non-linearity of the individual bin counts as a function of Δm^2_{31} can clearly be seen. It is then fed to the implementation of the fast Fourier transform (FFT) algorithm [140] provided by the `numpy` package.

As the input spectrum is real, only positive frequencies have to be considered in the FFT result. In addition, frequencies below 1/MeV are ignored as they are subject to large fluctuations. Thus the output of the FFT is a complex frequency spectrum in the range from 1/MeV to 8/MeV, given as its real and imaginary part. As the FFT is a linear transform, the error on both is given by

$$\mathrm{FFT}(\Sigma) = \mathcal{F}\Sigma\mathcal{F}^H \quad , \tag{6.31}$$

where $\mathcal{F}$ is the Fourier matrix, i. e. the Fourier transform of the n-dimensional unit matrix where n is the number of bins in the input series, and $\mathcal{F}^H$ its conjugate transpose. Σ is the covariance matrix of the input data, see e. g. [141].

In this case, the input data are independent bin counts N_i with statistical errors $\sigma_i = \sqrt{N_i}$, hence Σ is diagonal:

$$\Sigma = \mathrm{diag}(\vec{\sigma}) = \mathrm{diag}(\sigma_1^2, \sigma_2^2, \ldots, \sigma_n^2) = \mathrm{diag}(N_1, N_2, \ldots, N_n) \tag{6.32}$$

Propagating these properties through (6.31), one finds that the error on both the real and the imaginary

part of the FFT output is constant and proportional to the square root of the total number of bin entries[20]:

$$\Re[\mathrm{FFT}(\Sigma)] = \Im[\mathrm{FFT}(\Sigma)] = \mathbb{1}\sum_{i=1}^{n}\sigma_i^2 = \mathbb{1}\sum_{i=1}^{n}N_i = \mathbb{1}N_{\mathrm{tot}} \tag{6.33}$$

Yet for the analysis of the JUNO oscillation signal, amplitude and phase of the complex spectrum are better suited than its real and imaginary part. These quantities are calculated according to

$$A = \sqrt{\Re^2 + \Im^2} \tag{6.34}$$

$$\varphi = \arctan(\Im/\Re) + 2\pi k \tag{6.35}$$

with $k \in \mathbb{N}$ assuring that φ increases monotonically with increasing frequency, removing discontinuities. Their errors are then given by

$$\Delta A = \sqrt{\left(\frac{\Re}{A}\Delta\Re\right)^2 + \left(\frac{\Im}{A}\Delta\Im\right)^2} = \sqrt{N_{\mathrm{tot}}} \tag{6.36}$$

$$\Delta\varphi = \Delta A/A \quad . \tag{6.37}$$

The FFT result in the amplitude-phase base is shown in the middle and right panel of Fig. 6.20. One can see that the amplitude of the spectrum is virtually independent from Δm^2_{31} while the phase is approximately proportional to the mass splitting[21]. The shaded region marks the error range for the spectra for $\Delta m^2_{31} = 2.33 \cdot 10^{-3}\,\mathrm{eV}^2$, for the other spectra the errors are very similar.

6.5.4 Results for JUNO

Table 6.2: Uncertainties on all systematic parameters expected for JUNO with 10^5 detected events, ranked according to their impact on the mass hierarchy parameter h.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	3.32×10^{-1}	7.99×10^{-2}	3.22×10^{-1}	free
ϑ_{13} [°]	54.9	8.93	2.34	5.62×10^{-1}	2.27	free
Δm^2_{31} [eV^2]	38.2	2.46×10^{-3}	1.71×10^{-5}	4.74×10^{-6}	1.64×10^{-5}	free
$s_{A_{\mathrm{eff}},\mathrm{JUNO}}$	17.9	0.00	8.26×10^{1}	2.05×10^{1}	8.00×10^{1}	free
$s_{E,\mathrm{JUNO}}$	15.0	1.00	1.52×10^{-2}	5.14×10^{-3}	2.26×10^{-2}	2.00×10^{-2}
ϑ_{12}	10.5	3.36×10^{1}	1.87	8.65×10^{-1}	1.65	free
Δm^2_{21}	9.8	7.54×10^{-5}	6.96×10^{-6}	3.86×10^{-6}	5.79×10^{-6}	free
$n_{A_{\mathrm{eff}},\mathrm{JUNO}}$	0.0	0.00	2.00×10^{-2}	9.43×10^{-2}	2.79	2.00×10^{-2}

Now that an observable has been constructed which is sufficiently linear as a function of the parameters that JUNO will measure, the Fisher Matrix can be applied.

[20] The full error on any observable, i. e. entry in the FFT output spectrum, is given by the square root of the corresponding diagonal element of the covariance matrix, cf. (6.3).

[21] Note that all phases are shown relative to the phase spectrum for $\Delta m^2_{31} = 2.33 \cdot 10^{-3}\,\mathrm{eV}^2$ to visualise the behaviour more clearly.

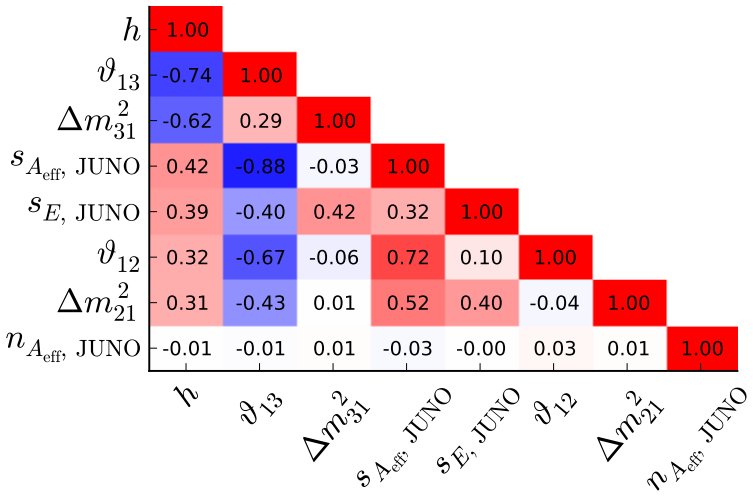

Figure 6.21: Covariance matrix for JUNO, without external constraints on the oscillation parameters.

The full error listing is given in Tab. 6.2, where one can read off a NMH sensitivity of $3.0\,\sigma$ for an event sample of $N_{\text{tot}} = 10^5$. Yet in contrast to the PINGU analysis, no priors have been put on any of the oscillation parameters as the resulting significance is to be compared to the one quoted by the JUNO collaboration themselves. In [73], for a "reactor-only analysis", i. e. no without external knowledge, a confidence level of $\Delta\chi^2 \approx 11$ is reported, corresponding to a significance of $\mathcal{S} = \sqrt{\Delta\chi^2} \approx 3.3\,\sigma$. This does not match up perfectly with the figure retrieved from the modified `PaPA` simulation, however the agreement is fairly good given that the two calculations are completely independent in simulation and analysis method, and in addition `PaPA` was intended for a very different experiment.

Another claim from [73] is that the NMH significance can be increased to $\Delta\chi^2 \approx 19 \mathrel{\hat{=}} \mathcal{S} \approx 4.4\,\sigma$ when using a prior measurement of $\Delta m^2_{\mu\mu}$ with 1 % precision. This parameter can be measured in ν_μ disappearance experiments, usually with a neutrino beam from an accelerator travelling over a baseline of $O(100\,\text{km})$. But since $\Delta m^2_{\mu\mu}$ is a rather complex combination of all mixing parameters [142], it is not possible to put the corresponding prior in the base of mixing parameters chosen in this thesis.

One can, however, include priors on all mixing parameters according to a current global fit [47] as it was done in the analysis for PINGU (cf. Sec 6.2.2). This increases the significance retrieved from the Fisher Matrix to $4.6\,\sigma$, a gain of $1.6\,\sigma$ w. r. t. the value without any priors. The full errors are listed in Tab. F.14. In JUNO's own analysis, only $1.1\,\sigma$ are added by the inclusion of external constraints, however there only one prior is put on an composite parameter, which is a weaker statement than putting individual priors on four underlying parameters.

The strongest impact on JUNO's expected sensitivity to the neutrino mass hierarchy has the uncertainty of the mixing angle ϑ_{13}, as one can read off from Tab. 6.2 and the covariance matrix, shown in Fig. 6.21. Including the current knowledge of this parameter alone enhances the significance to $4.4\,\sigma$.

Next in size is the impact of Δm^2_{31} on the mass hierarchy. For this parameter, however, including the current knowledge as a prior does not significantly improve the NMH measurement. The reason is that the current uncertainty on Δm^2_{31} is $8 \cdot 10^{-5}\,\text{eV}^2$ [47], while JUNO will already provide a constraint of $1.7 \cdot 10^{-5}\,\text{eV}^2$ by itself (see Tab. 6.2), i. e. there is no other experiment that can make a more precise measurement of Δm^2_{31} than JUNO.

This sub-percent precision on the value of Δm^2_{31} is stated by [73], too. A similar precision is claimed

for $\sin^2(\vartheta_{12})$ and Δm^2_{21}. Looking at Tab. 6.2, however, the relative uncertainties from the Fisher matrix are on the order of 10 %[22]. Yet in the given reference it remains unclear whether the stated uncertainties include the prior on $\Delta m^2_{\mu\mu}$. After adding priors on all oscillation parameters, the Fisher matrix lists a precision of $\approx 3\,\%$ for both parameters (see Tab. F.14).

6.5.5 Joint Analysis of JUNO and PINGU

Finally, the results of PINGU and JUNO can be combined by adding the respective Fisher matrices. To start off on common grounds, only the oscillation parameters will be considered as systematics without any prior knowledge. This severely decreases PINGU's NMH significance to only $1.5\,\sigma$ after three years of lifetime as it is strongly correlated with ϑ_{13}, see Tab. F.16. Here a fundamental difference in the approaches taken by PINGU and JUNO becomes evident: JUNO is able to determine the mass hierarchy with appreciable significance all by itself, while in PINGU at least the value of ϑ_{13} has to be retrieved from an independent experiment, e. g. Daya Bay [68]. JUNO on the other hand still has a NMH significance of $3.5\,\sigma$ considering only physics systematics without priors (Tab. F.15).

Table 6.3: Error listings for the combination of PINGU and JUNO, including only oscillation parameters without any priors.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	2.33×10^{-1}	7.56×10^{-2}	2.21×10^{-1}	free
Δm^2_{31} [eV^2]	83.4	2.46×10^{-3}	1.21×10^{-5}	4.58×10^{-6}	1.12×10^{-5}	free
ϑ_{13} [°]	76.8	8.93	1.01	4.64×10^{-1}	8.99×10^{-1}	free
ϑ_{12}	2.5	3.36×10^{1}	1.25	8.65×10^{-1}	9.08×10^{-1}	free
ϑ_{23} [°]	1.4	3.86×10^{1}	3.05×10^{-1}	3.02×10^{-1}	3.79×10^{-2}	free
Δm^2_{21}	0.7	7.54×10^{-5}	5.52×10^{-6}	3.86×10^{-6}	3.94×10^{-6}	free

If now the Fisher matrices of PINGU and JUNO are added, the combined sensitivity to the NMH increases to $4.3\,\sigma$, much larger than the square sum of the individual sensitivities, which is $3.8\,\sigma$. This is mainly due to the good constraint on ϑ_{13} from JUNO being passed on to PINGU and hence enhancing its NMH sensitivity. JUNO on the other hand can not profit much from PINGU since the only parameter that it can measure with greater precision than JUNO is ϑ_{23}, which does not affect JUNO at all.

A second outcome of the joint analysis is that JUNO's extremely precise measurement of Δm^2_{31} can be used to shrink down PINGU's error ellipse in the $(\Delta m^2_{31}, \sin^2(\vartheta_{23}))$ parameter space, which was shown above in Fig. 6.8. Fig. 6.22 is an update of the former including the expected result of JUNO. The combined ellipse of PINGU and JUNO demonstrates the enormous improvement in the knowledge about these two parameters that is to be expected over the course of the next decade—the time needed for the two experiments to be constructed and take their nominal amount of data.

Since Δm^2_{31} and ϑ_{23} are virtually uncorrelated in PINGU—as one can recognise from the semi-major axes of PINGU's error ellipse that are almost aligned to the coordinate axes—JUNO's strict constraint on Δm^2_{31} does not affect PINGU's precision on ϑ_{23}. Still the uncertainties on both parameters will be reduced by almost one order of magnitude w. r. t. the current precision.

[22] For the relative error on $\sin^2(\vartheta_{12})$, the errors listed in Tabs. 6.2 and F.14 for ϑ_{12} in degrees have to be converted: $\sin^2(\vartheta_{12}) \approx 0.306$ and $\sigma^{\text{full}}_{\sin^2(\vartheta_{12})} = \frac{\partial \sin^2(\vartheta_{12})}{\partial \vartheta_{12}} \sigma^{\text{full}}_{\vartheta_{12}} = \sin(2\vartheta_{12})\, \sigma^{\text{full}}_{\vartheta_{12}}$ with $\sigma^{\text{full}}_{\vartheta_{12}}$ in radians.

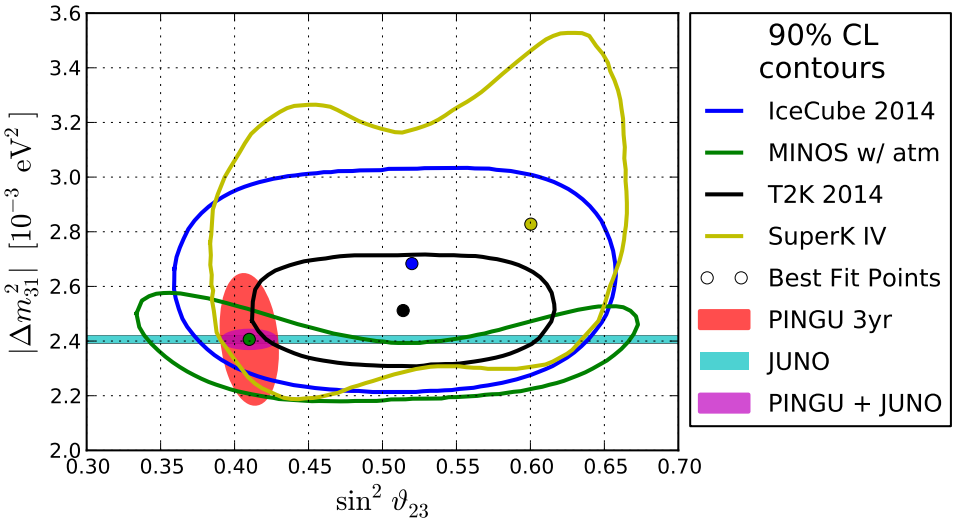

Figure 6.22: Same as Fig. 6.8, adding JUNO's expected constraint on Δm^2_{31} and the combined error ellipse of PINGU and JUNO.

6.6 Summary

In this chapter, PINGU's sensitivity to the neutrino oscillation parameters, especially the neutrino mass hierarchy, has been studied in depth. The baseline detector parametrisation leads to a median significance of $2.9\,\sigma$ after three years of data taking. Looking at the track and cascade channels separately, the individual significances are $1.3\,\sigma$ and $1.9\,\sigma$, respectively. This is a somewhat conservative estimate, and various effects increasing the expected significance have been explored in this chapter. These are summarised in Tab. 6.4.

The choice of the fiducial values for the oscillation parameters puts PINGU's baseline sensitivity to the neutrino mass hierarchy at the lower bound of the range of possible values: As one can see from the second set of rows in Tab. 6.4, PINGU will determine the NMH with higher significance if the true mass hierarchy is normal or ϑ_{23} is in the second octant. The latter effect is much more dramatic, possibly almost doubling the expected significance. In both cases, the increase is primarily due to the cascade channel.

Technical improvements to the detector itself are covered in the third section of the table. A more evolved event classification with more channels than just tracks and cascades only leads to a moderate increase of the NMH sensitivity, thus efforts to modify the PID in this direction seem not too promising. In terms of sensor technology, a further reduction of the module noise has virtually no positive impact on the significance. This points to the conclusion that already with current optical modules random noise does not notably affect the event reconstruction—the noise cleaning algorithms (Sec. 4.3.3) are very effective and the reconstruction results stable against small perturbations. Increasing the number of detected photons per neutrino energy by enhancing the photosensitive area per sensor on the other hand leads to improved reconstruction results. From these, PINGU's NMH measurement can indeed profit, as the features of the mass hierarchy pattern can be resolved more precisely.

The last row of Tab. 6.4 covers the combination of PINGU and JUNO. For this, all priors and detector parameters are removed, strongly reducing PINGU's sensitivity, especially in the track channel. The

Table 6.4: Summary of the effects studied in this chapter and their impact on PINGU's sensitivity to the NMH, relative to the baseline settings.

Model	Combined	Tracks	Cascades	Note	Section
Baseline	$2.9\,\sigma$	$1.3\,\sigma$	$1.9\,\sigma$	Absolute significance, reference for following rows	6.3
Remove priors on oscillation parameters, no detector nuisance parameters	$-1.4\,\sigma$	$-1.0\,\sigma$	$-0.6\,\sigma$	To be compatible with JUNO's analysis	6.5.5
ϑ_{23} in 2nd octant	$+2.7\,\sigma$	$+0.2\,\sigma$	$+2.3\,\sigma$	Fiducial value of ϑ_{23} mirrored at 45°	6.3.2
True normal NMH	$+0.2\,\sigma$	$+0.0\,\sigma$	$+0.2\,\sigma$		6.3.3
Three-channel PID	$+0.2\,\sigma$	$\pm0.0\,\sigma$	$-0.4\,\sigma$	Unidentified events contribute $1.3\,\sigma$ Reference uses PID from geometry V15	6.3.4
No module noise	$+0.1\,\sigma$	$+0.0\,\sigma$	$+0.0\,\sigma$	PINGU built from mDOMs	6.4.2
Photon statistics doubled	$+0.9\,\sigma$	$+0.9\,\sigma$	$+0.8\,\sigma$	PINGU built from WOMs	6.4.1
Combining with JUNO	$+2.8\,\sigma$	n/a	n/a	JUNO contributes $3.5\,\sigma$ Reference is second row	6.5.5

joint analysis with JUNO then leads to an expected NMH significance exceeding all values before, however the signal is mostly driven by JUNO.

Through all the analyses the picture holds that the cascade channel is driving PINGU's sensitivity to the neutrino mass hierarchy. Not only is its intrinsic significance higher than the track channel's, it also contributes most of the possible enhancement due to varying fiducial settings of the oscillation parameters and refined particle identification. In addition it is more stable if external constraints are loosened—the track channel on the other hand loses its discrimination power almost completely in this case.

CHAPTER 7

Conclusion

In this thesis, the development of PaPA, a standalone simulation for the planned PINGU atmospheric neutrino experiment, has been described. With PaPA, the expected event histograms in $(E, \cos\vartheta_{\text{zenith}})$ can be generated directly from inputs like the atmospheric neutrino flux, oscillation probabilities, effective areas of the detector and a parametrisation of the detector resolution, taking into account a variety of systematic parameters, cf. Sec. 5.2. Since most of the inputs are extracted from Monte Carlo data in advance, the time-consuming event-by-event simulation can be avoided in PaPA, making it well suited to explore PINGU's multi-dimensional parameter space in short time and robustly even in case of rather low Monte Carlo statistics.

Using the results of PaPA, PINGU's capability to measure the oscillation parameters ϑ_{23} and Δm^2_{31} and especially to determine the neutrino mass hierarchy, i. e. the sign of Δm^2_{31}, was assessed with the Fisher matrix technique. This analysis method, whose application to a particle physics experiment is novel, is especially suited for problems with a large number of parameters and essentially linear detector response to these parameters.

After checking that the requirements for applying the Fisher matrix are met (Sec. 6.1.2), PINGU's performance was evaluated for its current baseline geometry V36 with the most up-to-date estimates of reconstruction resolution. In Sec. 6.3 it was shown that for a nominal lifetime of three years, PINGU is expected to determine the neutrino mass hierarchy with a confidence level of $2.9\,\sigma$, with the major contribution coming from the analysis of cascade-like events[1]. In addition, PINGU will provide precision measurements of the oscillation parameters Δm^2_{31} and ϑ_{23}, and in particular resolve the octant of ϑ_{23}. The latter is possible as PINGU observes neutrino oscillations in matter, where the symmetry between the two octants is broken. This also means that the asymmetry between the normal and inverted mass hierarchy cases grows with an increasing value of ϑ_{23}, such that the expected significance of the mass hierarchy measurement can reach $5\,\sigma$ and more if the true value of ϑ_{23} is in the second octant. The use of next-generation optical sensors—like the WOM concept, for which basic work has been done in the context of this thesis—can further enhance PINGU's performance by increasing the expected photon statistics per event and hence improving the detection threshold and reconstruction precision (Sec. 6.4).

Finally, in Sec. 6.5 a combined analysis of PINGU and JUNO, a medium baseline reactor neutrino experiment currently under construction, was done to explore possible synergies. Given the respective settings, JUNO's expected neutrino spectrum could be simulated with PaPA without any modifications. Yet for the analysis with the Fisher matrix method the spectrum had to be Fourier transformed to achieve

[1] All events that are not caused by a ν_μ or $\bar{\nu}_\mu$ CC interaction, i. e. without an outgoing muon.

linear relations between the observables and the underlying physical parameters. After this, JUNO's sensitivity to the neutrino mass hierarchy as reported by the JUNO collaboration itself could be reproduced. The combined analysis of both experiments, i. e. adding the respective Fisher matrices, results in a NMH significance exceeding the squared sum of the individual ones by $0.5\,\sigma$, mainly due to PINGU profiting from JUNO's constraint on the oscillation parameter ϑ_{13}[2].

In a broader picture, intermediate results of this work have been published as official sensitivity estimates by the PINGU collaboration. The current version of the PINGU Letter of Intent [54] relies on the `PaPA` simulation of the previous baseline geometry V15, evaluated with the Fisher matrix technique. An update of this report is planned in the near future, referring to the current baseline geometry V36, reflecting all improvements on event selection, reconstruction etc. achieved in the meantime, and covering a wider range of systematic parameters.

In this update, the analysis will only partly depend on `PaPA`. The main calculations will be done with `pisa` [143], the software framework `PaPA` has evolved into. It keeps the idea of a direct simulation of event histograms in a staged way and develops it further. Separating the individual stages even more than `PaPA`, `pisa` offers the option to add alternative implementations of particular stages[3] and easily apply different analysis methods[4] to exactly the same data, as well as better maintainability of the code itself.

Since `PaPA` and the Fisher Matrix analysis have been fully re-implemented within `pisa` and both input and output format are compatible, the results presented in this thesis can always be reproduced. Vice versa, the concepts first introduced here, namely the staged simulation of event histograms using parametrisations fitted to distributions of Monte Carlo events and the Fisher matrix analysis, leading to a very fast evaluation of systematic effects, will stay present within PINGU and might even be adopted by other collaborations.

[2] For this study, no detector systematics were considered and no external constraints were put on the oscillation parameters.

[3] E. g. other oscillation codes, event reconstruction from parametrised reconstruction functions or directly from Monte Carlo data, ...

[4] E. g. the Fisher Matrix and the Likelihood Ratio techniques.

APPENDIX A

Details of the WOM Efficiency Calculation

A.1 Measurement of the Capture Efficiency

A fixed setup (see Fig. A.1) has been installed, which is specifically designed to allow for the quick and easy testing of different samples with wavelength-shifting (WLS) properties. As input radiation, the monochromated light of a xenon arc lamp is used, which allows for scanning wavelengths down to 200 nm. After passing through an optical chopper, the incident beam is split, with one beam used for monitoring the intensity of the incoming light with a reference photo diode. The other beam is coupled to a UV-transmitting light guide which then irradiates the sample.

The WLS sample is fixed in a mount that permits an easy exchange of samples of different size. Therefore also the irradiating end of the light guide is installed in a way that its position can be varied so that—apart from accounting for different sample sizes—one is able to vary the distance between light incidence and readout. This setup provides a handle on determining the absorption length at the emitted wavelength.

The WLS read-out is done via a second photo diode identical to the reference diode that is coupled to the WLS by an index-matching gel. Its output current, as well as the one of the reference diode, is measured with a lock-in amplifier. Using the photo diode calibration curve (response function), along

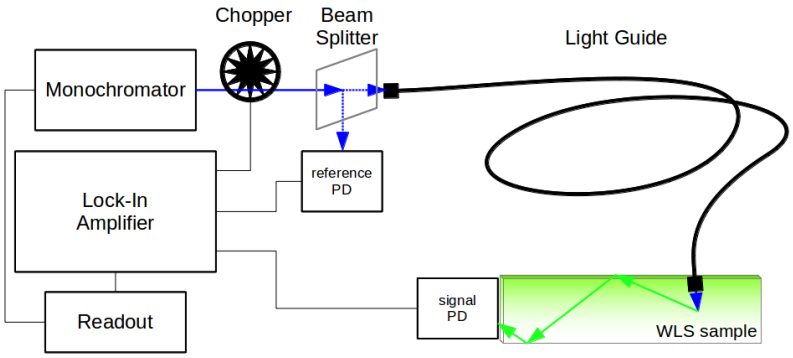

Figure A.1: Experimental setup for measuring the wavelength-dependent capture efficiency of the WLS samples.

with the WLS emission spectra (Fig. 4.10), the measured current is transformed into a photon flux.

For a typical measurement, the input wavelength is tuned with a resolution of 5 nm or less in the relevant regime. At each input wavelength, the current wavelength and the output from both diodes are stored. From this, the capture efficiency of the WLS sample is calculated as described in App. A.2. Note that this capture efficiency includes not only the light conversion and propagation efficiency, but also reflection effects at the surface occurring in the measurement. The results are hence somewhat dependant on the geometry of the WLS that is probed. Yet as the samples match the cylindrical shape of the intended final sensor geometry, those losses will also be present in the fully assembled module.

A.2 Calculating the WLS Capture Efficiency

The quantity that is to be extracted from the measurements is the wavelength dependent capture efficiency of a given WLS sample, $CE(\lambda)$.

In the data recorded are only the photo currents of signal and reference diode, $I_{\mathrm{sig}}(\lambda)$ and $I_{\mathrm{ref}}(\lambda)$, respectively. As additional input we need a calibration measurement, which is a normal wavelength scan without any sample installed so that the full output of the light guide is recorded. From this we get the output fraction

$$F(\lambda) = \frac{I_{\mathrm{sig,\,cal}}(\lambda)}{I_{\mathrm{ref,\,cal}}(\lambda)} \quad . \tag{A.1}$$

Furthermore, the response function

$$R(\lambda) = I_{\mathrm{out}}/P_{\mathrm{in}} \tag{A.2}$$

of the signal diode, which can be taken from the data sheet [144], has to be known and the (normalized) emission spectrum of the WLS $S_{\mathrm{WLS}}(\lambda)$ needs to be measured. The finite width of the monochromator spectrum (FWHM < 3 nm) is neglected.

To determine the capture efficiency, the number of photons detected when coming out of the sample has to be divided by the number of photons radiated into the sample. The incoming photon flux is calculated via the power emitted from the light guide:

$$P_{\mathrm{in}} = I_{\mathrm{ref}} \cdot F/R \tag{A.3}$$

Then the photon rate is then given by

$$\dot{N}_{\gamma,\,\mathrm{in}} = P_{\mathrm{in}}/\tfrac{hc}{\lambda} \quad . \tag{A.4}$$

When calculating the detected photon rate one has to take into account that the emitted spectrum is not monochromatic, hence one needs to average over the relevant wavelength regime:

$$\dot{N}_{\gamma,\,\mathrm{out}} = I_{\mathrm{sig}}/\langle R \cdot E_\gamma \rangle \quad , \tag{A.5}$$

where

$$\langle R \cdot E_\gamma \rangle = \int S_{\mathrm{WLS}}(\lambda)\, R(\lambda)\, \tfrac{hc}{\lambda}\, \mathrm{d}\lambda \quad . \tag{A.6}$$

The capture efficiency is finally given by

$$CE(\lambda) = \frac{\dot{N}_{\gamma,\,\mathrm{out}}(\lambda)}{\dot{N}_{\gamma,\,\mathrm{in}}(\lambda)} \cdot \Omega \quad , \tag{A.7}$$

where Ω is a factor accounting for the fact that the readout area in this setup is smaller than it will be in

a real application. Its value is given by

$$\Omega = \frac{\text{WLS front area}}{\text{photo diode area}} \cdot 2 \quad , \tag{A.8}$$

assuming that ultimately the WLS will be read over the full front area on both ends.

A.3 Transmission into the WOM

For light inciding into the WOM, its effective area A_{eff} is to be calculated given the the WOM length L (without end caps) and radius R. For simplicity, we will first handle the case of frontal illumination and then extend the result for arbitrary angles.

For the incoming light, the WOM has a rectangular cross-section of

$$A_{\text{tot}} = 2RL \quad . \tag{A.9}$$

The effective area is then given by

$$A_{\text{eff}} = \int_0^L \mathrm{d}l \cdot 2 \int_0^R \mathrm{d}r\, T(r) = A_{\text{tot}} \cdot \varepsilon_\Omega \tag{A.10}$$

with

$$\varepsilon_\Omega = \frac{1}{R} \int_0^R \mathrm{d}r\, T(r) = \int_0^{\frac{\pi}{2}} \cos\vartheta\, \mathrm{d}\vartheta\, T(\vartheta), \quad \sin\vartheta = r/R \tag{A.11}$$

and T being the transmission through two consecutive optical surfaces according to Fresnel's formulae. Since the transmission formulae are given as functions of the angle of incidence ϑ, it is convenient to transform the integral in (A.11) to depend on this angle (for an illustration see Fig. A.2b). Assuming the permeability is $\mu = 1$, Fresnel's formulae for the transmission of light polarised parallel and perpendicular to the optical surface read [35]:

$$T_\parallel = 1 - \left(\frac{n_t \cos\vartheta_i - n_i \cos\vartheta_t}{n_t \cos\vartheta_i + n_i \cos\vartheta_t} \right)^2 \tag{A.12}$$

$$T_\perp = 1 - \left(\frac{n_i \cos\vartheta_i - n_t \cos\vartheta_t}{n_i \cos\vartheta_i + n_t \cos\vartheta_t} \right)^2 \tag{A.13}$$

and accordingly for unpolarised light:

$$T_{\text{mean}}(\vartheta_i,\, n_i,\, n_t) = \left(T_\parallel(\vartheta_i,\, n_i,\, n_t) + T_\perp(\vartheta_i,\, n_i,\, n_t)\right)/2 \quad . \tag{A.14}$$

The angle of the refracted photon ϑ_t is given by Snell's law:

$$\vartheta_t(\vartheta_i,\, n_i,\, n_t) = \arcsin\left(n_i/n_t \sin\vartheta_i\right) \quad . \tag{A.15}$$

The combined transmission of two surfaces with a transition from medium 1 via medium 2 to medium 3, with refractive indices n_1, n_2, n_3, can be written as the product of two transmissions of a single surface where the angle of incidence on the second surface is given by the refraction at the first one:

$$T(\vartheta,\, n_1,\, n_2,\, n_3) = T_{\text{mean}}(\vartheta,\, n_1,\, n_2) \cdot T_{\text{mean}}(\vartheta_t(\vartheta,\, n_1,\, n_2),\, n_2,\, n_3) \quad . \tag{A.16}$$

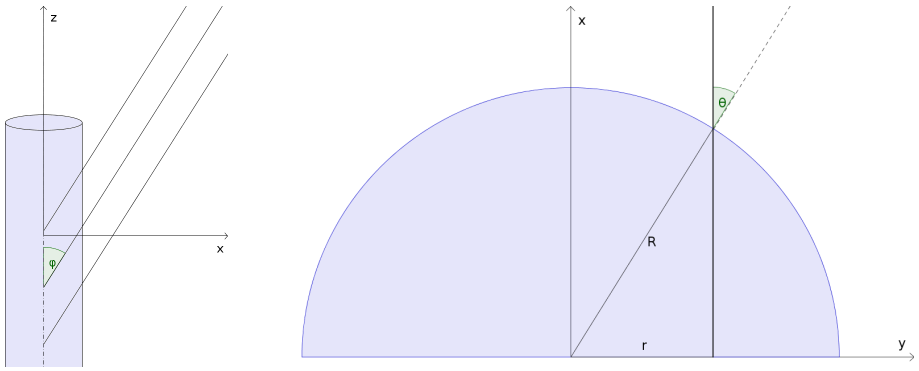

(a) (b)

Figure A.2: Definition of the angle of incidence on the WOM φ and the surface normal ϑ.

Putting this into (A.11), the integral can be solved numerically.

If now the incident light has an angle $\varphi \neq \pi/2$ with the WOM axis (see Fig. A.2a), the calculation has to be adjusted. If we choose the coordinates in a way that the WOM axis coincides with the z axis and the x-z-plane is defined by the incident light $\hat{i}$, ϑ is still embedded in the x-y-plane, but no longer equal to the angle between incident light and the surface normal $\hat{n}$ of the WOM. The latter, however, which we will call α, is needed for the Fresnel transmission. It is determined by

$$\cos\alpha = \hat{i} \cdot \hat{n} = \begin{pmatrix} \sin\varphi \\ 0 \\ \cos\varphi \end{pmatrix} \cdot \begin{pmatrix} \cos\vartheta \\ \sin\vartheta \\ 0 \end{pmatrix} = \sin\varphi \, \cos\vartheta \quad . \tag{A.17}$$

Additionally one has to account for the fact that for light inciding at an angle φ the length L of the WOM is jolted by a factor $\sin\varphi$. Including all this in (A.11) one ends up with

$$\varepsilon_\Omega(\varphi) = \sin\varphi \int_0^{\frac{\pi}{2}} \cos\vartheta \, \mathrm{d}\vartheta \, T\left(\alpha(\vartheta, \varphi)\right) \quad . \tag{A.18}$$

The full transmission efficiency ε_Ω is plotted as a function of $\cos\varphi$ in Fig. 4.11a.

APPENDIX **B**

Validation of the Fisher Matrix Approach

B.1 Linearity

As established in Sec. 6.1.2, one requirement for the Fisher Matrix to be valid is that the observables (i. e. the bin counts for PINGU) are described by linear functions of the parameters, hence the partial derivatives entering (6.1) can be considered to be constant within the considered range. To quantify the non-linearity, one can define the quantity

$$\Upsilon_n(p_i) = \left(N_n(p_{i,\,\mathrm{fid}} + \sigma_{p_i}) - N_n(p_{i,\,\mathrm{fid}}) + \frac{\partial N_n}{\partial p_i}\sigma_{p_i} \right) \Big/ N_n(p_{i,\,\mathrm{fid}}) \quad , \tag{B.1}$$

describing the relative contribution of the non-linear terms to the variation of the event count N_b in bin b as a function of the parameter p_i over its uncertainty range $p_{i,\,\mathrm{fid}} \pm \sigma_{p_i}$.

As one can see from Figs. B.1 and B.2, for the simple scaling parameters like the mass hierarchy h or the overall normalisation $n_{A_{\mathrm{eff}}}$ non-linearities are on the order of the machine precision. But also for the physics parameters they are typically smaller than 10^{-3}, hence the requirement of linearity (condition (a) in Sec. 6.1.2) is well fulfilled.

For means of illustration, the actual dependence of all bin entries on the systematic parameters is shown as well. By eye, deviations from linearity can only be recognised for Δm^2_{31} and the energy scale. Since both shift the sinusoidal oscillation pattern in energy non-linearities have to be expected, however they only occur on a scale conceivably larger than the actual uncertainty range.

B.2 Comparison to a Log-Likelihood Ratio Method

On the other hand, the test statistics of the experimental outcome is required to be Gaussian (condition (c) in Sec. 6.1.2). To demonstrate that this is in fact fulfilled, templates of the expected bin counts for one year of lifetime[1] have been created on a full 2D grid in Δm^2_{31} and ϑ_{23} for both normal and inverted hierarchy. One of the templates was chosen as "truth" to draw Poissonian pseudo experiments from. For each of the pseudo experiments X, the Log-Likelihood Ratio between normal and inverted hierarchy was

[1] This study has been done using the detector parametrisation for the outdated geometry V15. However, here the goal is to verify the method as such, thus the exact description of the detector is not relevant as long as it is not completely unrealistic.

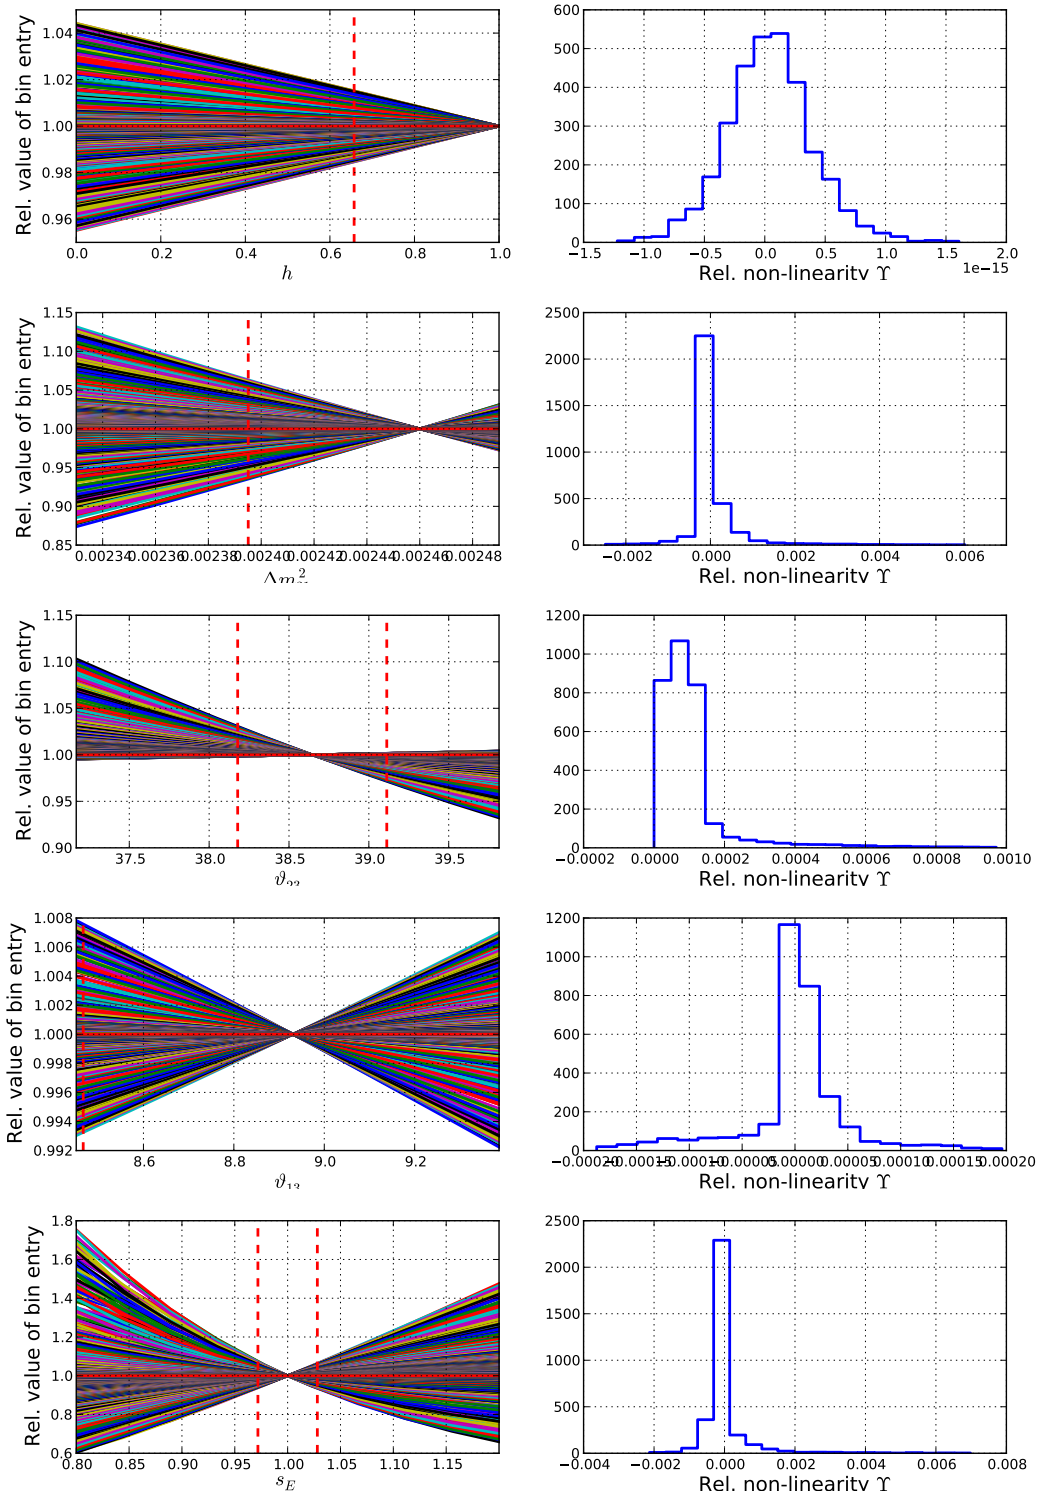

Figure B.1: *Left:* Relative values of all bin entries in the analysis histograms as functions of the systematic parameters. The the entries at the fiducial parameter values are set to one. The vertical lines indicate the full error range as listed in Tab. 6.1.
Right: Histograms of the non-linearities ϒ of the bin counts as functions of different systematic parameters (see equation (B.1)).

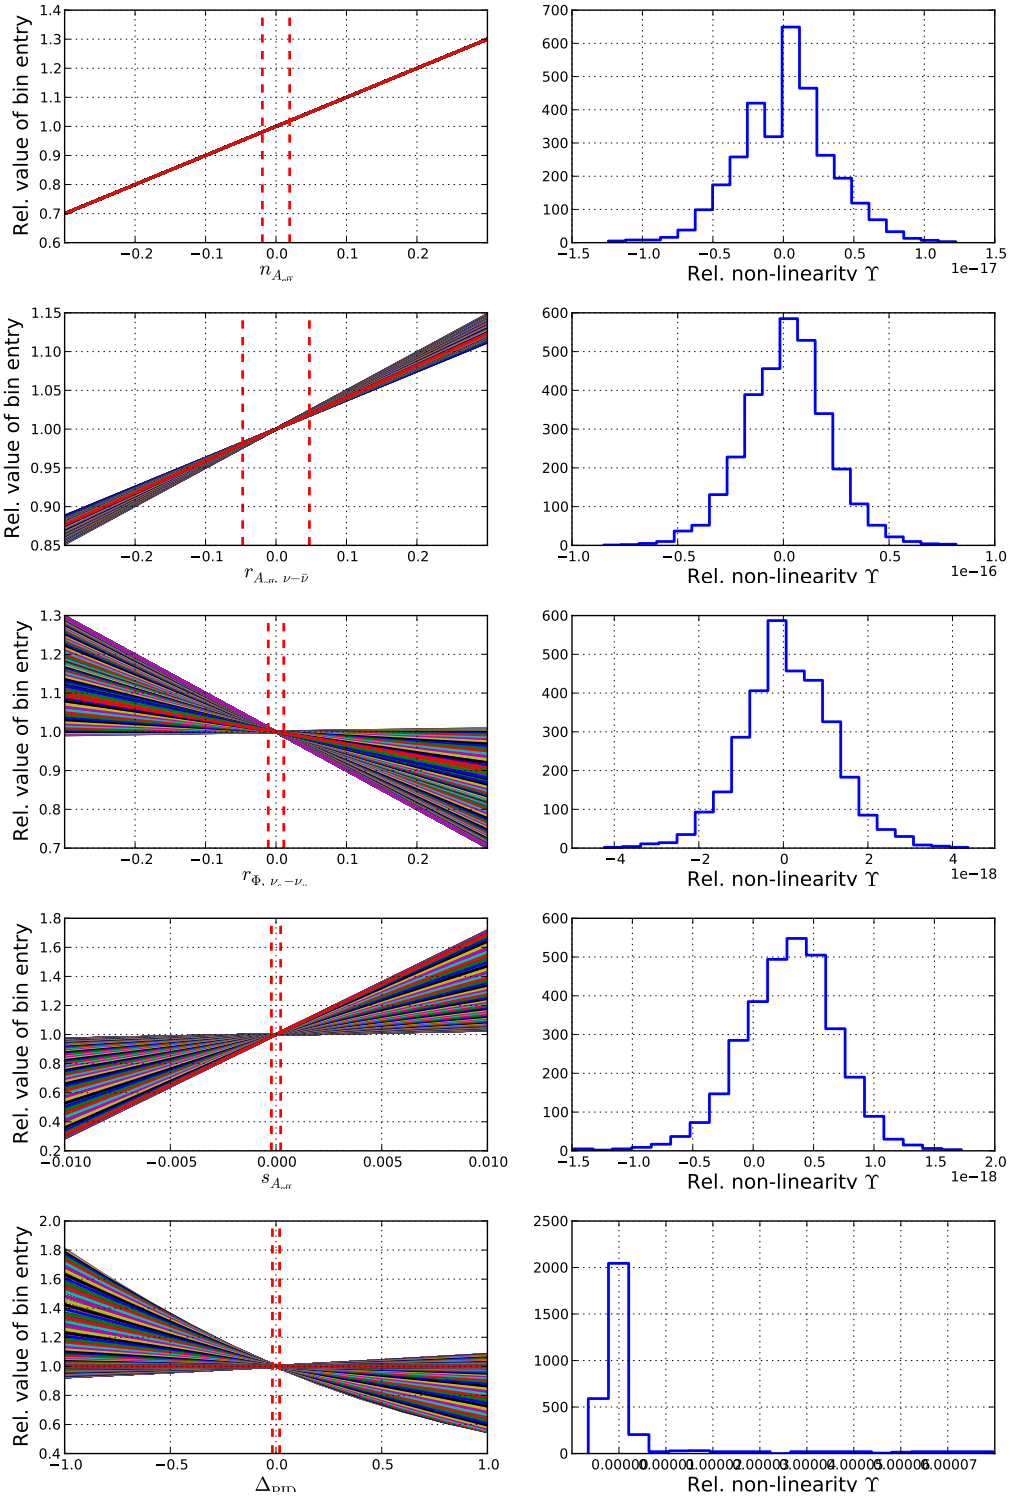

Figure B.2: Same as Fig. B.1 for the remaining parameters.

calculated according to

$$\mathrm{LLR} = -\log_{10} \frac{\mathcal{L}(X\,|\,\mathrm{IH})}{\mathcal{L}(X\,|\,\mathrm{NH})} \quad . \tag{B.2}$$

Here, $\mathcal{L}(X\,|\,\mathrm{IH})$ is the maximum Poissonian likelihood for comparing the pseudo experiment $X = (n_1, n_2, \dots)$ to all templates $T = (\mu_1, \mu_2, \dots)$ generated for the normal mass hierarchy:

$$\mathcal{L}(X\,|\,\mathrm{IH}) = \max_{T \in \mathrm{IH}} \prod_{i \in \mathrm{bins}} \frac{\mu_i^{n_i}}{n_i!} e^{-\mu_i} \tag{B.3}$$

$\mathcal{L}(X\,|\,\mathrm{NH})$ is calculated equivalently for the normal hierarchy.

As shown in Fig. B.3a, the test statistics is well described by a Gaussian distribution for both hierarchies.

Additionally, a median significance of the NMH measurement can be derived from the LLR distributions: Taking the median LLR value for NH (dashed vertical line in Fig. B.3a) and integrating over the IH distribution from this value upwards (hatched area) results in a one-sided p-value of 2.7 %, corresponding to a significance of $1.922\,\sigma$. Evaluating the Fisher Matrix on the same data (see below), with only Δm^2_{31} and ϑ_{23} as systematic parameters and without any priors, gives a consistent significance of $1.923\,\sigma$.

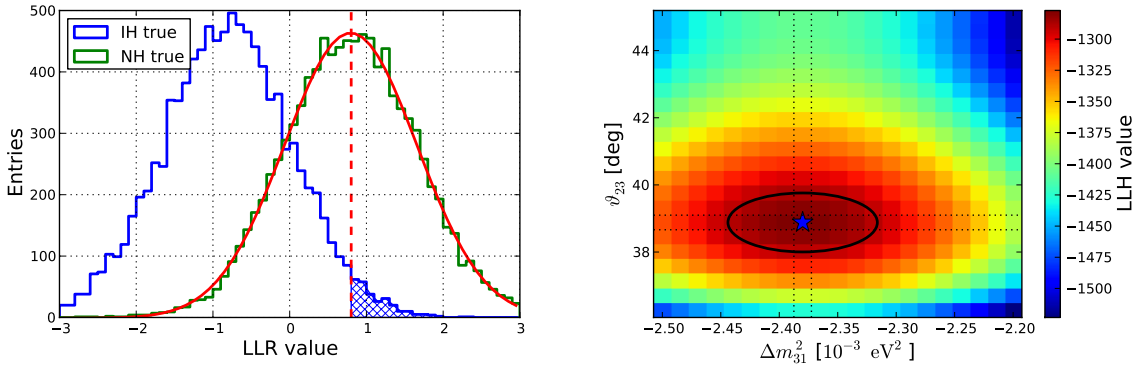

(a) (b)

Figure B.3: (a) Test statistic for PINGU with 10^4 pseudo experiments thrown for each NH and IH as assumed truth. A Gaussian fit to the NH distribution and its median value are indicated by the solid and dashed lines, respectively.
(b) Grid scan of the log-likelihood landscape for Δm^2_{31} (IH) and ϑ_{23}, overlaid with the 68 % CL ellipse calculated with the Fisher Matrix from the highlighted row and column. The star marks the injected truth.

When looking at the likelihood landscape itself (Fig. B.3b), the shape of the minimum in the (Δm^2_{31}, ϑ_{23}) plane is approximately elliptical, so condition (b) in Sec. 6.1.2 is fulfilled as well. From the slices through the parameter space corresponding to the injected true values of Δm^2_{31} and ϑ_{23}, highlighted in the figure, and the corresponding best fit point for inverted hierarchy, a Fisher matrix can be constructed. The $1\,\sigma$ error ellipse retrieved from this matrix aligns well with the true shape of the likelihood landscape, hence the Fisher matrix gives a precise description of the problem.

APPENDIX C

Oscillation Probabilities

C.1 Baseline Settings

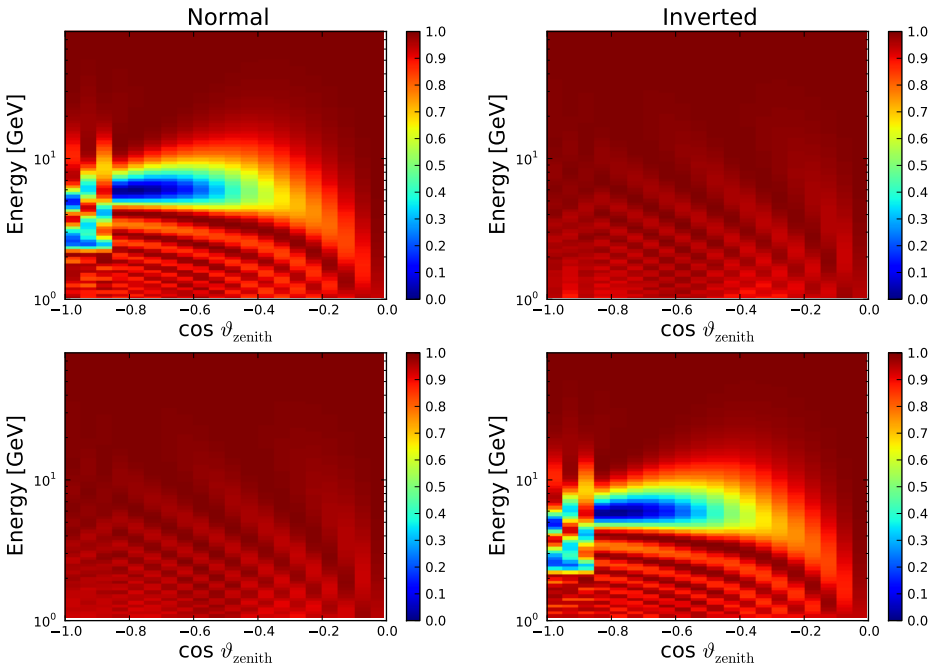

Figure C.1: Oscillation probabilities for $\nu_e \to \nu_e$ (top) and $\bar{\nu}_e \to \bar{\nu}_e$ (bottom) for normal and inverted hierarchy.

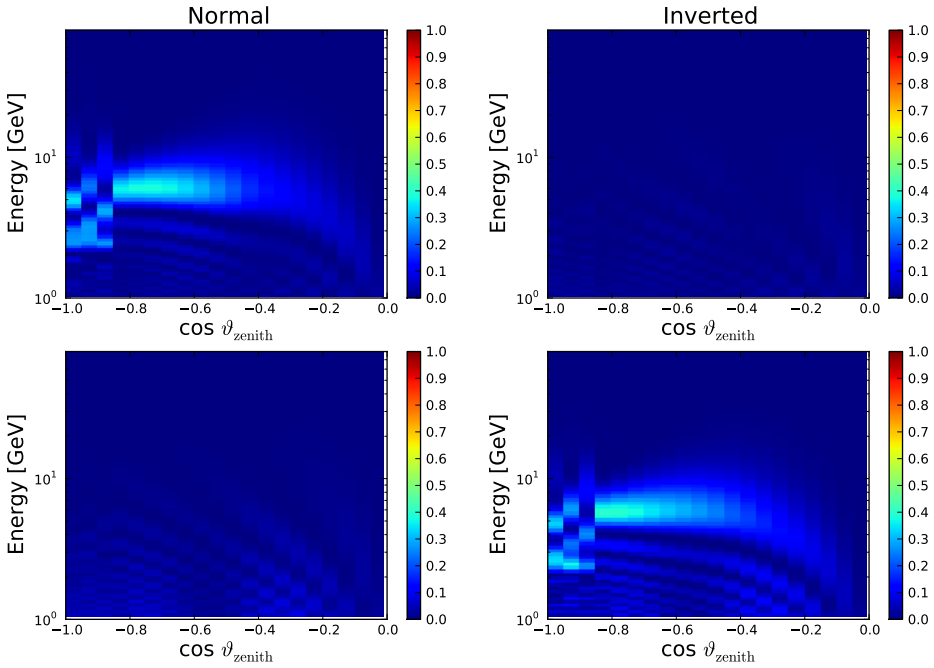

Figure C.2: Oscillation probabilities for $\nu_e \to \nu_\mu$ (top) and $\bar{\nu}_e \to \bar{\nu}_\mu$ (bottom) for normal and inverted hierarchy.

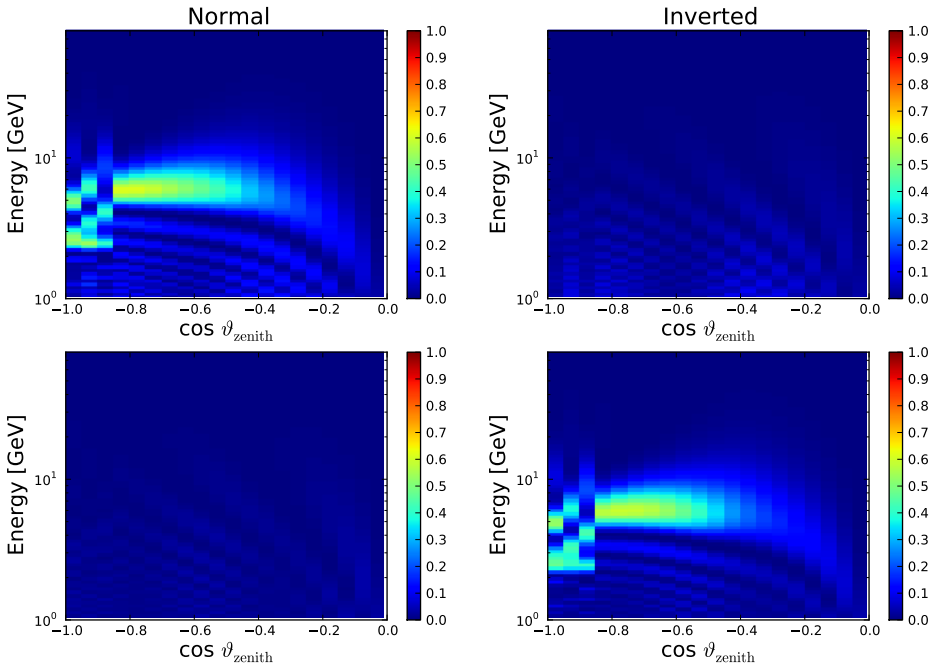

Figure C.3: Oscillation probabilities for $\nu_e \to \nu_\tau$ (top) and $\bar{\nu}_e \to \bar{\nu}_\tau$ (bottom) for normal and inverted hierarchy.

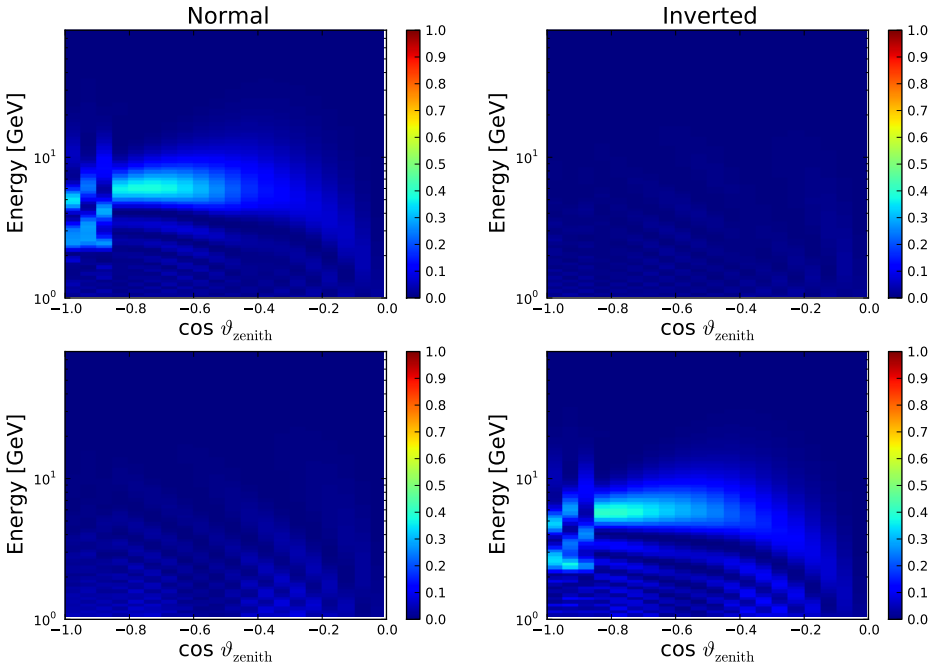

Figure C.4: Oscillation probabilities for $\nu_\mu \to \nu_e$ (top) and $\bar{\nu}_\mu \to \bar{\nu}_e$ (bottom) for normal and inverted hierarchy.

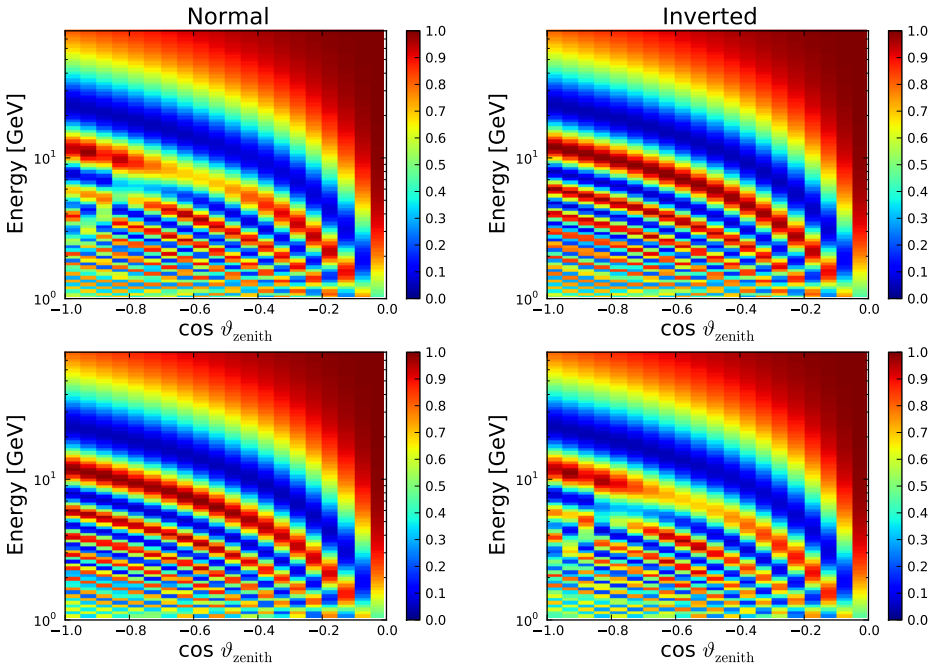

Figure C.5: Oscillation probabilities for $\nu_\mu \to \nu_\mu$ (top) and $\bar{\nu}_\mu \to \bar{\nu}_\mu$ (bottom) for normal and inverted hierarchy.

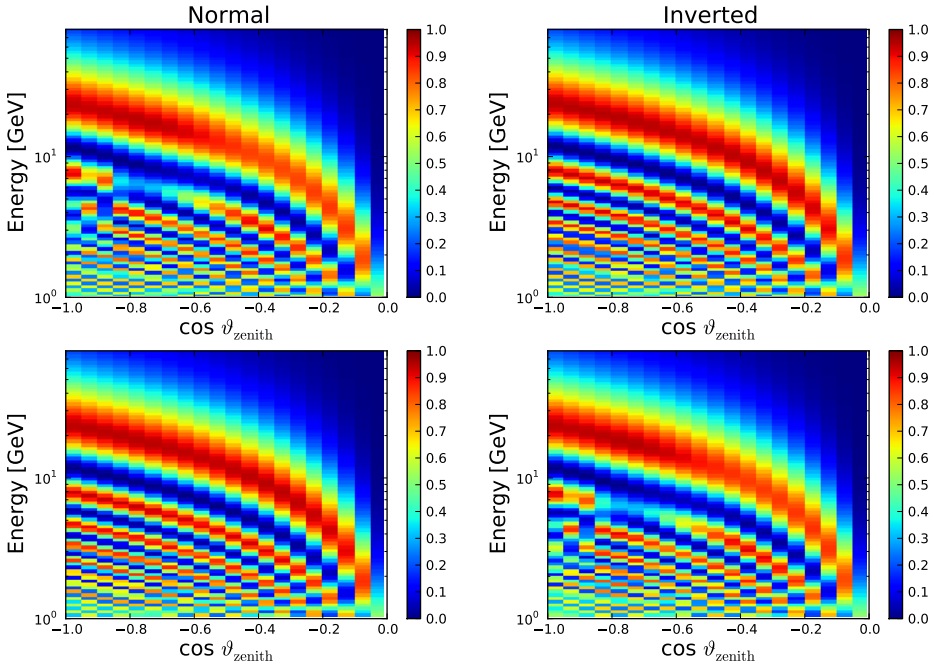

Figure C.6: Oscillation probabilities for $\nu_\mu \to \nu_\tau$ (top) and $\bar{\nu}_\mu \to \bar{\nu}_\tau$ (bottom) for normal and inverted hierarchy.

C.2 Scanning the fiducial value of ϑ_{23}

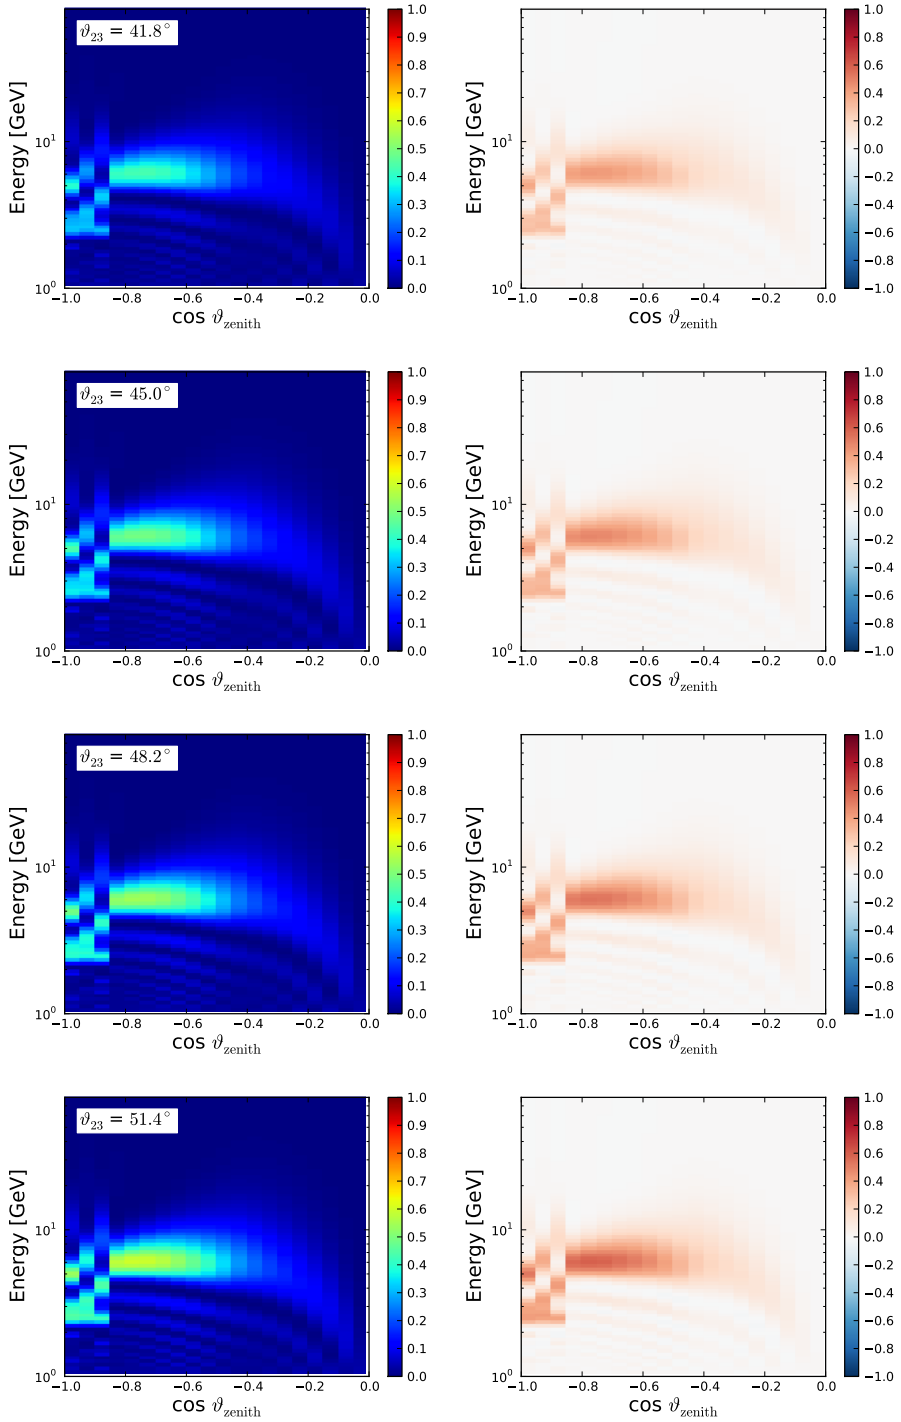

Figure C.7: Probabilities for $\nu_\mu \rightarrow \nu_e$ oscillations in normal mass hierarchy (left column) and difference between normal and inverted hierarchy oscillation probabilities (right column) for different values of ϑ_{23} in both octants.

APPENDIX D

Parametrisations of the Detector Resolutions

The full parametrisations of the reconstruction performances will be listed in form of the actual PaPA input. These are nested dictionaries giving the resolutions in energy (`'e'`) and $\cos\vartheta_{\text{zenith}}$ (`'coszen'`) for all four interaction channels: ν_e, ν_μ, and ν_τ CC (`'nue'`, `'numu'`, and `'nutau'`, respectively) and ν_X NC (`'NC'`). For each of those the resolution is given by the five parameters `'fraction'`, `'loc1'`, `'loc2'`, `'width1'`, and `'width2'`, corresponding to f, μ_1, μ_2, σ_1 and σ_2 in (5.6).

The actual function definitions for the five parameters is then supplied as a text string that can be interpreted as a python function by python's `eval()` function. In those definitions, `'n'` is a shorthand for the numpy library [124] used for most of the numerical operations in PaPA.

D.1 Baseline Settings

```
{
  "NC": {
    "coszen": {
      "fraction": "lambda E: 0.6*n.ones_like(E)",
      "loc1": "lambda E: -0.0*E + 0.0",
      "loc2": "lambda E: -0.0*n.ones_like(E)",
      "width1": "lambda E: 0.189 * n.exp(-E / 14.9) + 0.0581",
      "width2": "lambda E: 0.236 * n.exp(-E / 9.73) + 0.308"
    },
    "e": {
      "fraction": "lambda E: n.min([0.374 - 0.026*E, n.ones_like(E)], axis=0)",
      "loc1": "lambda E: -0.905*E + 1.45",
      "loc2": "lambda E: -0.686*E + 1.68",
      "width1": "lambda E: 0.0316*E + 0.688",
      "width2": "lambda E: 0.170*E + 0.650"
    }
  },
  "nue": {
    "coszen": {
      "fraction": "lambda E: 0.8*n.ones_like(E)",
      "loc1": "lambda E: 0.05*n.ones_like(E)",
      "loc2": "lambda E: -0.00240*E + 0.0509",
      "width1": "lambda E: 0.121 * n.exp(-E / 5.57) + 0.0414",
```

```
      "width2": "lambda E: 0.308 * n.exp(-E / 6.63) + 0.127"
    },
    "e": {
      "fraction": "lambda E: n.max([0.4 - 0.034*E, n.zeros_like(E)], axis=0)",
      "loc1": "lambda E: -0.0*E - 0.0",
      "loc2": "lambda E: 0.0*E + 1.5",
      "width1": "lambda E: 0.240*E + 0.187",
      "width2": "lambda E: 0.0*E + 1.5"
    }
  },
  "numu": {
    "coszen": {
      "fraction": "lambda E: 0.8*n.ones_like(E)",
      "loc1": "lambda E: n.zeros_like(E)",
      "loc2": "lambda E: -0.00331*E + 0.0641",
      "width1": "lambda E: 0.135 * n.exp(-E / 8.99) + 0.00369",
      "width2": "lambda E: 0.335 * n.exp(-E / 6.32) + 0.097"
    },
    "e": {
      "fraction": "lambda E: n.min([0.1 + 0.05*E, 0.6*n.ones_like(E)], axis=0)",
      "loc1": "lambda E: -0.0732*E - 0.125",
      "loc2": "lambda E: n.max([1.15 - 0.208*E, -2.*n.ones_like(E)], axis=0)",
      "width1": "lambda E: 0.0842*E + 0.467 + 0.3",
      "width2": "lambda E: 0.185*E + 1.01"
    }
  },
  "nutau": {
    "coszen": {
      "fraction": "lambda E: 0.6*n.ones_like(E)",
      "loc1": "lambda E: n.zeros_like(E)",
      "loc2": "lambda E: n.zeros_like(E)",
      "width1": "lambda E: n.max([0.151 - 0.00369*E, 0.6*n.ones_like(E)], axis=0)",
      "width2": "lambda E: 0.326 * n.exp(-E / 16.3) + 0.157"
    },
    "e": {
      "fraction": "lambda E: 0.0*n.ones_like(E)",
      "loc1": "lambda E: -0.457*E - 0.216",
      "loc2": "lambda E: -0.0*E + 0.0",
      "width1": "lambda E: 0.282*E - 0.0287",
      "width2": "lambda E: 0.0*E + 1."
    }
  }
}
```

D.2 Geometry V15 (Noiseless)

```
{
  "NC": {
    "coszen": {
      "fraction": "lambda E: 0.6*n.ones_like(E)",
      "loc1": "lambda E: 0.*n.ones_like(E)",
      "loc2": "lambda E: 0.*n.ones_like(E)",
```

```
      "width1": "lambda E: 0.131 * n.exp(-E / 12.0) + 0.0757",
      "width2": "lambda E: 0.273 * n.exp(-E / 13.7) + 0.227"
    },
    "e": {
      "fraction": "lambda E: 0.5*n.ones_like(E)",
      "loc1": "lambda E: -0.798*E + 1.17",
      "loc2": "lambda E: -0.522*E + 1.33",
      "width1": "lambda E: 0.0974*E + 0.499",
      "width2": "lambda E: 0.149*E + 0.789"
    }
  },
  "nue": {
    "coszen": {
      "fraction": "lambda E: n.max([-0.007*E + 0.78, n.zeros_like(E)], axis=0)",
      "loc1": "lambda E: 0.*n.ones_like(E)",
      "loc2": "lambda E: 0.*n.ones_like(E)",
      "width1": "lambda E: 0.103 * n.exp(-E / 4.64) + 0.0508",
      "width2": "lambda E: 0.319 * n.exp(-E / 4.94) + 0.137"
    },
    "e": {
      "fraction": "lambda E: n.min([0.0117*E + 0.0679, 0.4*n.ones_like(E)], axis=0)",
      "loc1": "lambda E: 0.331*n.ones_like(E)",
      "loc2": "lambda E: 0.16*E + 0.492",
      "width1": "lambda E: 0.241*E + 0.514",
      "width2": "lambda E: 0.0752*E + 0.623"
    }
  },
  "numu": {
    "coszen": {
      "fraction": "lambda E: n.max([-0.0117*E + 0.785, 0.4*n.ones_like(E)], axis=0)",
      "loc1": "lambda E: 0.*n.ones_like(E)",
      "loc2": "lambda E: 0.*n.ones_like(E)",
      "width1": "lambda E: 0.0918 * n.exp(-E / 9.9) + 0.0387",
      "width2": "lambda E: 0.278 * n.exp(-E / 14.5) + 0.0795"
    },
    "e": {
      "fraction": "lambda E: 0.15*n.ones_like(E)",
      "loc1": "lambda E: n.array([0.0748*e + 0.142 if e < 6 else -0.0809*e + 0.975 for e in E])",
      "loc2": "lambda E: n.array([2.10 if e < 6 else -0.876*e + 9.49 for e in E])",
      "width1": "lambda E: n.array([0.285*e + 0.269 if e < 6 else 0.116*e + 1.26 for e in E])",
      "width2": "lambda E: n.array([0.617*e - 0.346 if e < 15 else 0.008*e + 2.62 for e in E])"
    }
  },
  "nutau": {
    "coszen": {
      "fraction": "lambda E: 0.7*n.ones_like(E)",
      "loc1": "lambda E: 0.*n.ones_like(E)",
      "loc2": "lambda E: 0.*n.ones_like(E)",
      "width1": "lambda E: n.max([-0.00379*E + 0.363, 0.2*n.ones_like(E)], axis=0)",
      "width2": "lambda E: n.max([-0.00907*E + 0.133, 0.05*n.ones_like(E)], axis=0)"
    },
    "e": {
      "fraction": "lambda E: 0.3*n.ones_like(E)",
```

```
      "loc1": "lambda E: -0.341*E + 0.588",
      "loc2": "lambda E: -0.599*E + 0.148",
      "width1": "lambda E: 0.233*E + 0.373",
      "width2": "lambda E: 0.144*E + 0.255"
    }
  }
}
```

APPENDIX E

PID Functions

E.1 Baseline Settings

The particle identification is a binary decision, thus only the track identification probabilities $P_{\text{channel}\to\text{track}}$ are listed:

$$P_{\nu_e\to\text{track}}(E) = 0.192\exp\left(-\frac{(\log_{10}(E[\text{GeV}])-0.878)^2}{0.404^2}\right)+0.0309 \tag{E.1}$$

$$P_{\nu_\mu\to\text{track}}(E) = \frac{0.687}{\exp\left((0.683-\log_{10}(E[\text{GeV}]))/0.183\right)+1}+0.0585 \tag{E.2}$$

$$P_{\nu_\tau\to\text{track}}(E) = 0.197\exp\left(-\frac{(\log_{10}(E[\text{GeV}])-1.28)^2}{0.466^2}\right)+0.0732 \tag{E.3}$$

$$P_{\nu_X\,\text{NC}\to\text{track}}(E) = 0.171\exp\left(-\frac{(\log_{10}(E[\text{GeV}])-1.37)^2}{0.483^2}\right)+0.0339 \tag{E.4}$$

The cascade identification probabilities are given by $P_{\text{channel}\to\text{cascade}} = 1 - P_{\text{channel}\to\text{track}}$.

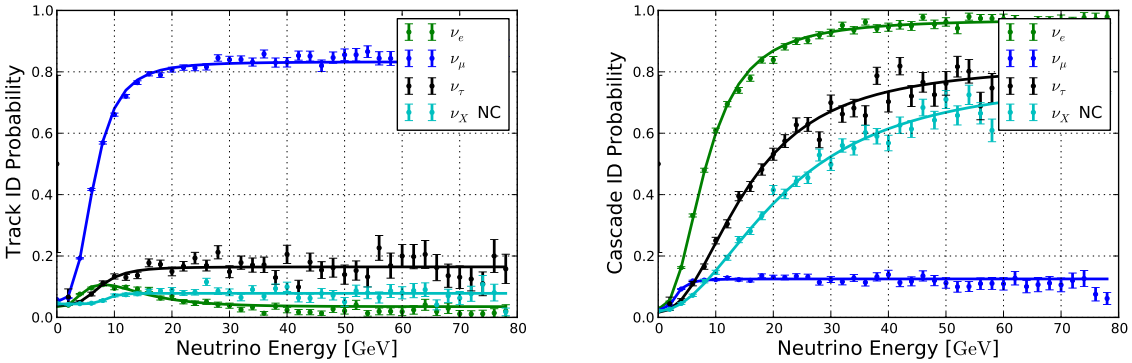

Figure E.1: Classification efficiencies for the high-purity track (left) and cascade (right) selections for PINGU V15, similar to Fig. 6.4.

E.2 High-Purity Event Classification

The particle identification is done by selecting high-purity track and cascade samples, events ending up in neither sample are classified as "unidentified":

$$P_{\nu_e\to\text{track}}(E) = 0.0481\exp\left(-\frac{(\log_{10}(E[\text{GeV}])-0.904)^2}{0.265^2}\right)+0.0350 \tag{E.5}$$

$$P_{\nu_e\to\text{cascade}}(E) = \frac{0.945}{\exp\left((0.922-\log_{10}(E[\text{GeV}]))/0.180\right)+1}+0.0264 \tag{E.6}$$

$$P_{\nu_\mu\to\text{track}}(E) = \frac{0.782}{\exp\left((0.808-\log_{10}(E[\text{GeV}]))/0.135\right)+1}+0.0507 \tag{E.7}$$

$$P_{\nu_\mu\to\text{cascade}}(E) = \frac{0.103}{\exp\left((0.499-\log_{10}(E[\text{GeV}]))/0.127\right)+1}+0.0223 \tag{E.8}$$

$$P_{\nu_\tau\to\text{track}}(E) = \frac{0.125}{\exp\left((0.882-\log_{10}(E[\text{GeV}]))/0.120\right)+1}+0.0396 \tag{E.9}$$

$$P_{\nu_\tau\to\text{cascade}}(E) = \frac{0.811}{\exp\left((1.19-\log_{10}(E[\text{GeV}]))/0.200\right)+1}+0.0163 \tag{E.10}$$

$$P_{\nu_X\,\text{NC}\to\text{track}}(E) = \frac{0.0344}{\exp\left((0.939-\log_{10}(E[\text{GeV}]))/0.0516\right)+1}+0.0439 \tag{E.11}$$

$$P_{\nu_X\,\text{NC}\to\text{cascade}}(E) = \frac{0.771}{\exp\left((1.34-\log_{10}(E[\text{GeV}]))/0.216\right)+1}+0.0202 \tag{E.12}$$

The fit of these functions to data is shown in Fig. E.1.

APPENDIX F

Full Error Listings

Here the full error lists, similar to Tab. 6.1, for various detector settings and sub-channels will be collected. These have not been included in the main text for the sake of better readability. For a description how to read the tables, refer to the explanations given for Tab. 6.1 in Sec. 6.3.

F.1 Baseline Settings

Table F.1: Same as Tab. 6.1, but for the track channel only

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	7.47×10^{-1}	3.20×10^{-1}	6.75×10^{-1}	free
s_E	8.4	1.00	2.98×10^{-2}	8.36×10^{-3}	3.62×10^{-2}	5.00×10^{-2}
Δm^2_{31} [eV^2]	7.2	2.46×10^{-3}	6.70×10^{-5}	1.93×10^{-5}	1.21×10^{-4}	8.00×10^{-5}
ϑ_{23} [°]	6.0	3.86×10^{1}	5.83×10^{-1}	3.55×10^{-1}	5.43×10^{-1}	1.32
ϑ_{13} [°]	3.6	8.93	4.67×10^{-1}	9.47×10^{-1}	1.01×10^{1}	4.68×10^{-1}
$r_{\Phi,\,\nu_e-\nu_\mu}$	2.4	0.00	3.59×10^{-2}	5.84×10^{-3}	5.12×10^{-2}	5.00×10^{-2}
Δ_{PID} [GeV]	2.3	0.00	3.80×10^{-2}	1.71×10^{-2}	3.41×10^{-2}	5.00×10^{-1}
$n_{A_{\text{eff}}}$	1.3	0.00	3.00×10^{-2}	3.87×10^{-3}	3.01×10^{-2}	2.00×10^{-1}
$s_{A_{\text{eff}}}$ [m^2/GeV]	0.3	0.00	4.13×10^{-4}	1.64×10^{-4}	3.79×10^{-4}	free
$r_{A_{\text{eff}},\,\nu-\bar{\nu}}$	0.2	0.00	4.86×10^{-2}	9.29×10^{-3}	2.06×10^{-1}	5.00×10^{-2}

Table F.2: Same as Tab. 6.1, but for the cascade channel only

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	5.17×10^{-1}	3.39×10^{-1}	3.91×10^{-1}	free
$r_{\Phi,\,\nu_e-\nu_\mu}$	24.9	0.00	2.44×10^{-2}	1.06×10^{-2}	2.58×10^{-2}	5.00×10^{-2}
ϑ_{23} [°]	20.7	3.86×10^{1}	1.05	5.78×10^{-1}	1.61	1.32
Δ_{PID} [GeV]	11.4	0.00	1.53×10^{-1}	4.14×10^{-2}	1.55×10^{-1}	5.00×10^{-1}
$s_{A_{\text{eff}}}$ [m^2/GeV]	4.5	0.00	3.96×10^{-4}	1.85×10^{-4}	3.50×10^{-4}	free
$r_{A_{\text{eff}},\,\nu-\bar{\nu}}$	4.2	0.00	4.91×10^{-2}	5.55×10^{-3}	2.60×10^{-1}	5.00×10^{-2}
ϑ_{13} [°]	3.0	8.93	4.66×10^{-1}	1.67	5.53	4.68×10^{-1}
s_E	1.5	1.00	3.14×10^{-2}	1.92×10^{-2}	3.55×10^{-2}	5.00×10^{-2}
Δm^2_{31} [eV^2]	1.2	2.46×10^{-3}	6.73×10^{-5}	4.44×10^{-5}	1.16×10^{-4}	8.00×10^{-5}
$n_{A_{\text{eff}}}$	0.1	0.00	2.29×10^{-2}	2.32×10^{-3}	2.29×10^{-2}	2.00×10^{-1}

F.2 Normal Mass Hierarchy as Truth

Table F.3: Full error listings for the combined analysis of tracks and cascades when assuming true normal mass hierarchy.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	0.00	3.22×10^{-1}	2.32×10^{-1}	2.23×10^{-1}	free
ϑ_{13} [°]	8.5	8.93	4.64×10^{-1}	8.23×10^{-1}	3.64	4.68×10^{-1}
Δm^2_{31} [eV^2]	2.4	2.46×10^{-3}	6.49×10^{-5}	1.77×10^{-5}	1.09×10^{-4}	8.00×10^{-5}
ϑ_{23} [°]	2.1	3.86×10^{1}	4.74×10^{-1}	3.02×10^{-1}	4.08×10^{-1}	1.32
Δ_{PID} [GeV]	0.5	0.00	1.70×10^{-2}	1.58×10^{-2}	6.27×10^{-3}	5.00×10^{-1}
$r_{\Phi,\,\nu_e-\nu_\mu}$	0.5	0.00	1.09×10^{-2}	5.33×10^{-3}	9.79×10^{-3}	5.00×10^{-2}
$r_{A_{\text{eff}},\,\nu-\bar{\nu}}$	0.5	0.00	4.74×10^{-2}	4.83×10^{-3}	1.49×10^{-1}	5.00×10^{-2}
$n_{A_{\text{eff}}}$	0.1	0.00	1.90×10^{-2}	1.99×10^{-3}	1.90×10^{-2}	2.00×10^{-1}
s_E	0.1	1.00	2.81×10^{-2}	7.65×10^{-3}	3.31×10^{-2}	5.00×10^{-2}
$s_{A_{\text{eff}}}$ [m^2/GeV]	0.0	0.00	2.25×10^{-4}	1.23×10^{-4}	1.88×10^{-4}	free

Table F.4: Same as Tab. F.3, but for the track channel only

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	0.00	7.51×10^{-1}	3.20×10^{-1}	6.79×10^{-1}	free
s_E	8.6	1.00	2.98×10^{-2}	8.35×10^{-3}	3.62×10^{-2}	5.00×10^{-2}
Δm^2_{31} [eV^2]	7.0	2.46×10^{-3}	6.70×10^{-5}	1.93×10^{-5}	1.21×10^{-4}	8.00×10^{-5}
ϑ_{23} [°]	7.0	3.86×10^{1}	5.83×10^{-1}	3.53×10^{-1}	5.45×10^{-1}	1.32
$n_{A_{\text{eff}}}$	3.8	0.00	2.95×10^{-2}	3.87×10^{-3}	2.96×10^{-2}	2.00×10^{-1}
ϑ_{13} [°]	3.5	8.93	4.67×10^{-1}	9.48×10^{-1}	1.00×10^{1}	4.68×10^{-1}
$r_{\Phi,\,\nu_e-\nu_\mu}$	3.0	0.00	3.55×10^{-2}	5.90×10^{-3}	5.02×10^{-2}	5.00×10^{-2}
Δ_{PID} [GeV]	2.9	0.00	3.82×10^{-2}	1.71×10^{-2}	3.42×10^{-2}	5.00×10^{-1}
$r_{A_{\text{eff}},\,\nu-\bar{\nu}}$	0.4	0.00	4.86×10^{-2}	9.38×10^{-3}	2.10×10^{-1}	5.00×10^{-2}
$s_{A_{\text{eff}}}$ [m^2/GeV]	0.0	0.00	4.17×10^{-4}	1.64×10^{-4}	3.83×10^{-4}	free

Table F.5: Same as Tab. F.3, but for the cascade channel only

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	0.00	4.64×10^{-1}	3.36×10^{-1}	3.21×10^{-1}	free
ϑ_{23} [°]	25.8	3.86×10^{1}	1.06	5.79×10^{-1}	1.67	1.32
$r_{\Phi,\,\nu_e-\nu_\mu}$	14.5	0.00	2.34×10^{-2}	1.24×10^{-2}	2.33×10^{-2}	5.00×10^{-2}
Δ_{PID} [GeV]	7.7	0.00	1.51×10^{-1}	4.15×10^{-2}	1.53×10^{-1}	5.00×10^{-1}
ϑ_{13} [°]	3.4	8.93	4.66×10^{-1}	1.66	5.61	4.68×10^{-1}
$s_{A_{\text{eff}}}$ [m^2/GeV]	2.3	0.00	3.97×10^{-4}	1.85×10^{-4}	3.51×10^{-4}	free
$n_{A_{\text{eff}}}$	1.6	0.00	2.29×10^{-2}	2.32×10^{-3}	2.29×10^{-2}	2.00×10^{-1}
Δm^2_{31} [eV^2]	0.9	2.46×10^{-3}	6.72×10^{-5}	4.43×10^{-5}	1.15×10^{-4}	8.00×10^{-5}
s_E	0.8	1.00	3.12×10^{-2}	1.92×10^{-2}	3.51×10^{-2}	5.00×10^{-2}
$r_{A_{\text{eff}},\,\nu-\bar{\nu}}$	0.2	0.00	4.91×10^{-2}	5.64×10^{-3}	2.62×10^{-1}	5.00×10^{-2}

F.3 High-Purity Event Classification

Table F.6: Full error listings for the high-purity track channel, see Sec. 6.3.4.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	1.10	4.93×10^{-1}	9.86×10^{-1}	free
s_E	21.3	1.00	3.36×10^{-2}	1.31×10^{-2}	4.34×10^{-2}	5.00×10^{-2}
Δm^2_{31} [eV^2]	10.7	2.46×10^{-3}	6.97×10^{-5}	3.03×10^{-5}	1.38×10^{-4}	8.00×10^{-5}
ϑ_{23} [°]	6.1	3.86×10^{1}	7.53×10^{-1}	5.50×10^{-1}	7.32×10^{-1}	1.32
Δ_{PID} [GeV]	3.6	0.00	8.20×10^{-2}	3.98×10^{-2}	7.29×10^{-2}	5.00×10^{-1}
ϑ_{13} [°]	1.7	8.93	4.67×10^{-1}	1.46	1.83×10^{1}	4.68×10^{-1}
$s_{A_{\text{eff}}}$ [m^2/GeV]	0.4	0.00	6.16×10^{-4}	2.74×10^{-4}	5.52×10^{-4}	free
$r_{\Phi,\, \nu_e - \nu_\mu}$	0.3	0.00	4.37×10^{-2}	9.35×10^{-3}	8.96×10^{-2}	5.00×10^{-2}
s_{PID}	0.0	1.00	2.05×10^{-1}	7.47×10^{-3}	2.05×10^{-1}	free
$r_{A_{\text{eff}},\, \nu - \bar{\nu}}$	0.0	0.00	4.95×10^{-2}	1.78×10^{-2}	3.60×10^{-1}	5.00×10^{-2}
$n_{A_{\text{eff}}}$	0.0	0.00	2.00×10^{-1}	7.47×10^{-3}	NaN	2.00×10^{-1}

Table F.7: Full error listings for the standard cascade channel (i. e. cascades are all events not identified as tracks), see Sec. 6.3.4.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	7.30×10^{-1}	5.47×10^{-1}	4.83×10^{-1}	free
$r_{\Phi,\, \nu_e - \nu_\mu}$	17.1	0.00	3.12×10^{-2}	1.83×10^{-2}	3.54×10^{-2}	5.00×10^{-2}
ϑ_{23} [°]	12.5	3.86×10^{1}	1.26	9.53×10^{-1}	4.03	1.32
$s_{A_{\text{eff}}}$ [m^2/GeV]	7.7	0.00	5.53×10^{-4}	3.37×10^{-4}	4.39×10^{-4}	free
$n_{A_{\text{eff}}}$	3.8	0.00	2.33×10^{-2}	3.89×10^{-3}	2.32×10^{-2}	2.00×10^{-1}
Δ_{PID} [GeV]	1.9	0.00	1.78×10^{-1}	9.95×10^{-2}	1.62×10^{-1}	5.00×10^{-1}
$r_{A_{\text{eff}},\, \nu - \bar{\nu}}$	1.9	0.00	4.96×10^{-2}	9.31×10^{-3}	3.74×10^{-1}	5.00×10^{-2}
ϑ_{13} [°]	1.3	8.93	4.67×10^{-1}	2.92	8.26	4.68×10^{-1}
s_E	0.4	1.00	3.64×10^{-2}	3.88×10^{-2}	3.65×10^{-2}	5.00×10^{-2}
Δm^2_{31} [eV^2]	0.3	2.46×10^{-3}	7.03×10^{-5}	8.96×10^{-5}	1.17×10^{-4}	8.00×10^{-5}

Table F.8: Full error listings for the high-purity cascade channel, see Sec. 6.3.4.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	9.54×10^{-1}	6.35×10^{-1}	7.12×10^{-1}	free
Δ_{PID} [GeV]	29.0	0.00	9.77×10^{-2}	3.74×10^{-2}	9.24×10^{-2}	5.00×10^{-1}
$r_{\Phi,\nu_e-\nu_\mu}$	21.4	0.00	4.34×10^{-2}	4.50×10^{-2}	7.45×10^{-2}	5.00×10^{-2}
$s_{A_{\text{eff}}}$ [m^2/GeV]	16.6	0.00	1.06×10^{-3}	4.43×10^{-4}	9.64×10^{-4}	free
ϑ_{23} [°]	2.8	3.86×10^{1}	1.29	3.02	5.08	1.32
s_{PID}	1.2	1.00	2.05×10^{-1}	9.07×10^{-3}	2.05×10^{-1}	free
s_E	0.5	1.00	4.12×10^{-2}	5.85×10^{-2}	4.30×10^{-2}	5.00×10^{-2}
ϑ_{13} [°]	0.5	8.93	4.67×10^{-1}	3.82	1.20×10^{1}	4.68×10^{-1}
$r_{A_{\text{eff}},\nu-\bar{\nu}}$	0.4	0.00	5.00×10^{-2}	2.07×10^{-2}	1.41	5.00×10^{-2}
Δm^2_{31} [eV^2]	0.1	2.46×10^{-3}	7.38×10^{-5}	1.38×10^{-4}	1.32×10^{-4}	8.00×10^{-5}
$n_{A_{\text{eff}}}$	0.0	0.00	2.00×10^{-1}	9.07×10^{-3}	NaN	2.00×10^{-1}

Table F.9: Full error listings for the sample of unidentified events, see Sec. 6.3.4.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	1.13	8.16×10^{-1}	7.80×10^{-1}	free
$r_{\Phi,\nu_e-\nu_\mu}$	21.8	0.00	4.14×10^{-2}	1.64×10^{-2}	7.21×10^{-2}	5.00×10^{-2}
$n_{A_{\text{eff}}}$	10.6	0.00	2.49×10^{-2}	4.30×10^{-3}	2.47×10^{-2}	2.00×10^{-1}
s_E	9.7	1.00	4.22×10^{-2}	5.06×10^{-2}	6.06×10^{-2}	5.00×10^{-2}
s_{PID}	4.8	1.00	6.42×10^{-3}	2.30×10^{-3}	6.00×10^{-3}	free
Δm^2_{31} [eV^2]	4.1	2.46×10^{-3}	7.39×10^{-5}	1.15×10^{-4}	1.54×10^{-4}	8.00×10^{-5}
Δ_{PID} [GeV]	3.9	0.00	1.15×10^{-1}	4.72×10^{-2}	1.09×10^{-1}	5.00×10^{-1}
$s_{A_{\text{eff}}}$ [m^2/GeV]	2.3	0.00	2.32×10^{-3}	5.12×10^{-4}	2.26×10^{-3}	free
ϑ_{23} [°]	1.8	3.86×10^{1}	1.29	9.74×10^{-1}	6.12	1.32
ϑ_{13} [°]	0.6	8.93	4.67×10^{-1}	3.74	9.99	4.68×10^{-1}
$r_{A_{\text{eff}},\nu-\bar{\nu}}$	0.5	0.00	4.97×10^{-2}	1.04×10^{-2}	4.26×10^{-1}	5.00×10^{-2}

Table F.10: Full error listings for the high-purity cascade channel combined with the unidentified sample after Fisher matrix evaluation.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	6.19×10^{-1}	5.01×10^{-1}	3.63×10^{-1}	free
ϑ_{23} [°]	18.1	3.86×10^{1}	1.20	9.27×10^{-1}	2.74	1.32
$r_{\Phi,\nu_e-\nu_\mu}$	6.7	0.00	2.63×10^{-2}	1.54×10^{-2}	2.68×10^{-2}	5.00×10^{-2}
s_{PID}	4.3	1.00	3.00×10^{-3}	2.23×10^{-3}	2.01×10^{-3}	free
$s_{A_{\text{eff}}}$ [m^2/GeV]	3.3	0.00	5.53×10^{-4}	3.35×10^{-4}	4.40×10^{-4}	free
$r_{A_{\text{eff}},\nu-\bar{\nu}}$	2.1	0.00	4.94×10^{-2}	9.31×10^{-3}	3.31×10^{-1}	5.00×10^{-2}
ϑ_{13} [°]	1.8	8.93	4.67×10^{-1}	2.67	7.62	4.68×10^{-1}
$n_{A_{\text{eff}}}$	0.9	0.00	2.19×10^{-2}	3.89×10^{-3}	2.16×10^{-2}	2.00×10^{-1}
s_E	0.3	1.00	3.63×10^{-2}	3.83×10^{-2}	3.64×10^{-2}	5.00×10^{-2}
Δ_{PID} [GeV]	0.3	0.00	3.37×10^{-2}	2.93×10^{-2}	1.69×10^{-2}	5.00×10^{-1}
Δm^2_{31} [eV^2]	0.2	2.46×10^{-3}	7.02×10^{-5}	8.84×10^{-5}	1.17×10^{-4}	8.00×10^{-5}

F.4 PINGU equipped with mDOMs

Table F.11: Full error listings for PINGU assuming reconstruction and particle identification as in geometry V15, i. e. no module noise.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	3.33×10^{-1}	2.34×10^{-1}	2.37×10^{-1}	free
ϑ_{23} [°]	12.5	3.86×10^{1}	5.06×10^{-1}	2.81×10^{-1}	4.69×10^{-1}	1.32
$r_{\Phi,\nu_e-\nu_\mu}$	7.7	0.00	9.32×10^{-3}	4.85×10^{-3}	8.15×10^{-3}	5.00×10^{-2}
$r_{A_{\text{eff}},\nu-\bar{\nu}}$	7.6	0.00	4.78×10^{-2}	4.73×10^{-3}	1.62×10^{-1}	5.00×10^{-2}
ϑ_{13} [°]	7.4	8.93	4.64×10^{-1}	9.29×10^{-1}	3.46	4.68×10^{-1}
$n_{A_{\text{eff}}}$	6.0	0.00	1.93×10^{-2}	1.96×10^{-3}	1.93×10^{-2}	2.00×10^{-1}
Δm^2_{31} [eV^2]	2.2	2.46×10^{-3}	6.59×10^{-5}	2.28×10^{-5}	1.14×10^{-4}	8.00×10^{-5}
Δ_{PID} [GeV]	1.3	0.00	2.31×10^{-2}	2.00×10^{-2}	1.16×10^{-2}	5.00×10^{-1}
s_E	0.2	1.00	2.88×10^{-2}	9.83×10^{-3}	3.39×10^{-2}	5.00×10^{-2}
$s_{A_{\text{eff}}}$ [m^2/GeV]	0.1	0.00	2.33×10^{-4}	1.47×10^{-4}	1.81×10^{-4}	free

Table F.12: Same as Tab. F.11, but for the track channel only

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	7.93×10^{-1}	3.65×10^{-1}	7.03×10^{-1}	free
s_E	13.1	1.00	3.14×10^{-2}	1.05×10^{-2}	3.91×10^{-2}	5.00×10^{-2}
Δm^2_{31} [eV^2]	7.4	2.46×10^{-3}	6.80×10^{-5}	2.42×10^{-5}	1.26×10^{-4}	8.00×10^{-5}
Δ_{PID} [GeV]	5.7	0.00	5.21×10^{-2}	2.16×10^{-2}	4.78×10^{-2}	5.00×10^{-1}
ϑ_{13} [°]	3.1	8.93	4.67×10^{-1}	1.11	1.08×10^{1}	4.68×10^{-1}
ϑ_{23} [°]	3.1	3.86×10^{1}	5.85×10^{-1}	3.26×10^{-1}	5.65×10^{-1}	1.32
$n_{A_{\text{eff}}}$	0.3	0.00	3.58×10^{-2}	4.34×10^{-3}	3.61×10^{-2}	2.00×10^{-1}
$r_{A_{\text{eff}},\nu-\bar{\nu}}$	0.1	0.00	4.88×10^{-2}	1.05×10^{-2}	2.28×10^{-1}	5.00×10^{-2}
$r_{\Phi,\nu_e-\nu_\mu}$	0.0	0.00	3.41×10^{-2}	5.59×10^{-3}	4.62×10^{-2}	5.00×10^{-2}
$s_{A_{\text{eff}}}$ [m^2/GeV]	0.0	0.00	4.99×10^{-4}	2.10×10^{-4}	4.53×10^{-4}	free

Table F.13: Same as Tab. F.11, but for the cascade channel only

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	5.34×10^{-1}	3.04×10^{-1}	4.38×10^{-1}	free
$r_{\Phi,\nu_e-\nu_\mu}$	43.4	0.00	2.69×10^{-2}	9.78×10^{-3}	3.03×10^{-2}	5.00×10^{-2}
$s_{A_{\text{eff}}}$ [m^2/GeV]	18.8	0.00	4.92×10^{-4}	2.06×10^{-4}	4.46×10^{-4}	free
ϑ_{23} [°]	10.4	3.86×10^{1}	1.19	5.56×10^{-1}	2.69	1.32
s_E	6.7	1.00	3.58×10^{-2}	2.83×10^{-2}	4.26×10^{-2}	5.00×10^{-2}
$n_{A_{\text{eff}}}$	5.1	0.00	2.18×10^{-2}	2.20×10^{-3}	2.18×10^{-2}	2.00×10^{-1}
Δm^2_{31} [eV^2]	2.5	2.46×10^{-3}	6.98×10^{-5}	6.61×10^{-5}	1.27×10^{-4}	8.00×10^{-5}
Δ_{PID} [GeV]	2.5	0.00	1.13×10^{-1}	5.25×10^{-2}	1.03×10^{-1}	5.00×10^{-1}
ϑ_{13} [°]	2.3	8.93	4.66×10^{-1}	1.71	5.49	4.68×10^{-1}
$r_{A_{\text{eff}},\nu-\bar{\nu}}$	2.2	0.00	4.95×10^{-2}	5.29×10^{-3}	3.37×10^{-1}	5.00×10^{-2}

F.5 JUNO

Table F.14: Full error listings for JUNO including prior knowledge on the oscillation parameters as given in [47].

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	2.19×10^{-1}	7.99×10^{-2}	2.04×10^{-1}	free
Δm^2_{31} [eV^2]	37.1	2.46×10^{-3}	1.56×10^{-5}	4.74×10^{-6}	1.52×10^{-5}	8.00×10^{-5}
$s_{A_{\text{eff}},\text{JUNO}}$	32.0	0.00	3.49×10^{1}	2.05×10^{1}	2.82×10^{1}	free
ϑ_{13} [°]	7.2	8.93	4.49×10^{-1}	5.62×10^{-1}	1.51	4.68×10^{-1}
ϑ_{12}	4.4	3.36×10^{1}	8.12×10^{-1}	8.65×10^{-1}	9.32×10^{-1}	1.06
$s_{E,\text{JUNO}}$	2.5	1.00	1.33×10^{-2}	5.14×10^{-3}	1.70×10^{-2}	2.00×10^{-2}
Δm^2_{21}	0.1	7.54×10^{-5}	2.22×10^{-6}	3.86×10^{-6}	4.41×10^{-6}	2.40×10^{-6}
$n_{A_{\text{eff}},\text{JUNO}}$	0.0	0.00	2.00×10^{-2}	9.43×10^{-2}	5.29×10^{-1}	2.00×10^{-2}

Table F.15: Error listings for JUNO, including only oscillation parameters without any priors.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	2.87×10^{-1}	7.99×10^{-2}	2.76×10^{-1}	free
Δm^2_{31} [eV^2]	88.8	2.46×10^{-3}	1.52×10^{-5}	4.74×10^{-6}	1.45×10^{-5}	free
ϑ_{13} [°]	69.3	8.93	1.08	5.62×10^{-1}	9.27×10^{-1}	free
ϑ_{12}	0.9	3.36×10^{1}	1.26	8.65×10^{-1}	9.21×10^{-1}	free
Δm^2_{21}	0.1	7.54×10^{-5}	5.68×10^{-6}	3.86×10^{-6}	4.16×10^{-6}	free

Table F.16: Same as Tab. F.15, but for PINGU.

Parameter	Impact [%]	Best Fit	σ^{full}	σ^{stat}	σ^{syst}	Prior
h	100.0	1.00	6.72×10^{-1}	2.33×10^{-1}	6.30×10^{-1}	free
ϑ_{13} [°]	79.8	8.93	3.18	8.24×10^{-1}	3.07	free
Δm^2_{31} [eV^2]	25.1	2.46×10^{-3}	3.56×10^{-5}	1.77×10^{-5}	3.09×10^{-5}	free
ϑ_{23} [°]	5.4	3.86×10^{1}	3.13×10^{-1}	3.02×10^{-1}	8.05×10^{-2}	free

Bibliography

[1] S. L. Glashow, "Partial-symmetries of weak interactions", *Nuclear Physics* 22.4 (1961) 579–588, ISSN: 0029-5582, DOI: 10.1016/0029-5582(61)90469-2.

[2] S. Weinberg, "A Model of Leptons", *Phys. Rev. Lett.* 19 (21 1967) 1264–1266, DOI: 10.1103/PhysRevLett.19.1264.

[3] A. Salam, "Weak and electromagnetic interactions", *Proceedings of 8th Nobel Symposium, Lerum, Sweden*, ed. by N. Svartholm, Almqvist & Wiksell, 1968 367.

[4] G. 't Hooft and M. Veltman, "Regularization and renormalization of gauge fields", *Nuclear Physics B* 44.1 (1972) 189–213, ISSN: 0550-3213, DOI: 10.1016/0550-3213(72)90279-9.

[5] M. Böhm, A. Denner and H. Joos, *Gauge theories of the strong and electroweak interaction*, 3rd ed., Teubner, 2001, ISBN: 3-519-23045-3.

[6] 27th Oct. 2014, URL: http://www.isgtw.org/sites/default/files/Standard_model_infographic.png.

[7] P. W. Higgs, "Broken Symmetries and the Masses of Gauge Bosons", *Phys. Rev. Lett.* 13 (16 1964) 508–509, DOI: 10.1103/PhysRevLett.13.508.

[8] F. Englert and R. Brout, "Broken Symmetry and the Mass of Gauge Vector Mesons", *Phys. Rev. Lett.* 13 (9 1964) 321–323, DOI: 10.1103/PhysRevLett.13.321.

[9] G. Aad et al., "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC", *Physics Letters B* 716.1 (2012) 1–29, ISSN: 0370-2693, DOI: 10.1016/j.physletb.2012.08.020.

[10] S. Chatrchyan et al., "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC", *Physics Letters B* 716.1 (2012) 30–61, ISSN: 0370-2693, DOI: 10.1016/j.physletb.2012.08.021.

[11] L. M. Brown, "The idea of the neutrino", *Physics Today* 31 (1978) 23–28.

[12] C. L. Cowan et al., "Detection of the Free Neutrino: a Confirmation", *Science* 124.3212 (1956) 103–104, DOI: 10.1126/science.124.3212.103.

[13] G. Danby et al., "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos", *Phys. Rev. Lett.* 9 (1 1962) 36–44, DOI: 10.1103/PhysRevLett.9.36.

[14] K. Kodama et al., "Observation of tau neutrino interactions", *Physics Letters B* 504.3 (2001) 218–224, ISSN: 0370-2693, DOI: 10.1016/S0370-2693(01)00307-0.

[15] R. Davis, D. S. Harmer and K. C. Hoffman, "Search for Neutrinos from the Sun", *Phys. Rev. Lett.* 20 (21 1968) 1205–1209, DOI: 10.1103/PhysRevLett.20.1205.

[16] Y. Fukuda et al., "Evidence for Oscillation of Atmospheric Neutrinos", *Phys. Rev. Lett.* 81 (8 1998) 1562–1567, DOI: 10.1103/PhysRevLett.81.1562.

[17] S.-H. Seo, "New Results from RENO and The 5 MeV Excess" (2014), arXiv: 1410.7987 [hep-ex].

[18] S. E. Kopp, "Accelerator neutrino beams", *Physics Reports* 439.3 (2007) 101–159, ISSN: 0370-1573, DOI: 10.1016/j.physrep.2006.11.004.

[19] C. E. Rolfs, W. S. Rodney and W. A. Fowler, *Cauldrons in the cosmos: nuclear astrophysics*, Theoretical astrophysics, University of Chicago Press, 1988, ISBN: 0-226-72456-5.

[20] G. Pfennig, H. Klewe-Nebenius and W. Seelmann-Eggebert, *Karlsruher Nuklidkarte*, 6th ed., Forschungszentrum Karlsruhe, Institut für Radiochemie, 1995.

[21] A. M. Hillas, "Cosmic Rays: Recent Progress and some Current Questions" (2006), arXiv: astro-ph/0607109 [astro-ph].

[22] M. Honda et al., "Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data", *Phys. Rev. D* 75 (4 2007) 043006, DOI: 10.1103/PhysRevD.75.043006.

[23] G. D. Barr et al., "Three-dimensional calculation of atmospheric neutrinos", *Phys. Rev. D* 70 (2 2004) 023006, DOI: 10.1103/PhysRevD.70.023006.

[24] R. Enberg, M. H. Reno and I. Sarcevic, "Prompt neutrino fluxes from atmospheric charm", *Phys. Rev. D* 78 (4 2008) 043005, DOI: 10.1103/PhysRevD.78.043005.

[25] M. G. Aartsen et al., "Measurement of the Atmospheric ν_e Flux in IceCube", *Phys. Rev. Lett.* 110 (15 2013) 151105, DOI: 10.1103/PhysRevLett.110.151105.

[26] E. Fermi, "On the Origin of the Cosmic Radiation", *Phys. Rev.* 75 (1949) 1169–1174, DOI: 10.1103/PhysRev.75.1169.

[27] J. Beringer et al., "Review of Particle Physics", *Phys. Rev. D* 86 (1 2012) 010001, DOI: 10.1103/PhysRevD.86.010001.

[28] S. Swordy, "The Energy Spectra and Anisotropies of Cosmic Rays", English, *Space Science Reviews* 99.1-4 (2001) 85–94, ISSN: 0038-6308, DOI: 10.1023/A:1013828611730.

[29] E. Waxman and J. Bahcall, "High energy neutrinos from astrophysical sources: An upper bound", *Phys. Rev. D* 59 (2 1998) 023002, DOI: 10.1103/PhysRevD.59.023002.

[30] M. G. Aartsen et al., "Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector", *Science* 342.6161 (2013), DOI: 10.1126/science.1242856.

[31] M. G. Aartsen et al., "Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data", *Phys. Rev. Lett.* 113 (10 2014) 101101, DOI: 10.1103/PhysRevLett.113.101101.

[32] M. G. Aartsen et al., "A Combined Maximum-Likelihood Analysis of the High-Energy Astrophysical Neutrino Flux Measured with IceCube", to be published, 2015.

[33] J. A. Formaggio and G. P. Zeller,
"From eV to EeV: Neutrino cross sections across energy scales",
Rev. Mod. Phys. 84 (3 2012) 1307–1341, DOI: 10.1103/RevModPhys.84.1307.

[34] Y.-S. Tsai, "Pair production and bremsstrahlung of charged leptons",
Rev. Mod. Phys. 46 (4 1974) 815–851, DOI: 10.1103/RevModPhys.46.815.

[35] J. D. Jackson, *Klassische Elektrodynamik*, 3rd ed., de Gruyter, 2002.

[36] P. Price and K. Woschnagg,
"Role of group and phase velocity in high-energy neutrino observatories",
Astropart. Phys. 15.1 (2001) 97–100, ISSN: 0927-6505,
DOI: 10.1016/S0927-6505(00)00142-0.

[37] Y. Nagashima, *Beyond the Standard Model of Elementary Particle Physics*, Wiley-VCH, 2014,
ISBN: 978-3-527-41177-1.

[38] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics*,
Oxford University Press, 2007, ISBN: 978-0-19-850871-7.

[39] K. Zuber, *Neutrino Physics (Series in High Energy Physics, Cosmology and Gravitation)*,
Taylor & Francis, 2003, ISBN: 978-0-750-30750-5.

[40] Z.-Z. Xing and S. Zhou, *Neutrinos in Particle Physics, Astronomy and Cosmology*,
Springer, 2011.

[41] "Precision electroweak measurements on the Z resonance",
Physics Reports 427.5–6 (2006) 257–454, ISSN: 0370-1573,
DOI: 10.1016/j.physrep.2005.12.006.

[42] A. A. Aguilar-Arevalo et al.,
"Event Excess in the MiniBooNE Search for $\bar{\nu}_\mu$ to $\bar{\nu}_e$ Oscillations",
Phys. Rev. Lett. 105 (18 2010) 181801, DOI: 10.1103/PhysRevLett.105.181801.

[43] G. Drexlin et al., "Current direct neutrino mass experiments",
Adv. High Energy Phys. 2013 (2013) 293986, DOI: 10.1155/2013/293986.

[44] C. Kraus et al.,
"Final results from phase II of the Mainz neutrino mass searchin tritium β decay", English,
EPJ C 40.4 (2005) 447–468, ISSN: 1434-6044, DOI: 10.1140/epjc/s2005-02139-7.

[45] V. Aseev et al., "Upper limit on the electron antineutrino mass from the Troitsk experiment",
Phys. Rev. D 84 (11 2011) 112003, DOI: 10.1103/PhysRevD.84.112003.

[46] J. Wolf et al., "The KATRIN neutrino mass experiment",
Nucl. Instr. Meth. A 623 (2010) 442–444, DOI: 10.1016/j.nima.2010.03.030.

[47] G. Fogli et al., "Global analysis of neutrino masses, mixings, and phases: Entering the era of leptonic *CP* violation searches", *Phys. Rev. D* 86 (1 2012) 013012,
DOI: 10.1103/PhysRevD.86.013012.

[48] M. Gonzalez-Garcia, M. Maltoni and T. Schwetz,
"Updated fit to three neutrino mixing: status of leptonic CP violation", English,
Journal of High Energy Physics 2014.11, 52 (2014), DOI: 10.1007/JHEP11(2014)052.

[49] L. Wolfenstein, "Neutrino oscillations in matter", *Phys. Rev. D* 17 (9 1978) 2369–2374, DOI: 10.1103/PhysRevD.17.2369.

[50] S. Mikheyev and A. Smirnov, *Sov. J. Nucl. Phys.* 42 (1985) 913–917.

[51] S. Mikheyev and A. Smirnov,
"Resonant amplification of ν oscillations in matter and solar-neutrino spectroscopy",
Nuovo Cim. C 9.1 (1986) 17–26, ISSN: 0390-5551, DOI: 10.1007/BF02508049.

[52] Q. Ahmad et al., "Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory", *Phys. Rev. Lett.* 89 (1 2002) 011301, DOI: 10.1103/PhysRevLett.89.011301.

[53] E. Akhmedov, S. Razzaque and A. Smirnov,
"Mass hierarchy, 2-3 mixing and CP-phase with huge atmospheric neutrino detectors", English, *JHEP* 2013.2, 82 (2013), DOI: 10.1007/JHEP02(2013)082.

[54] M. Aartsen et al.,
"Letter of Intent: The Precision IceCube Next Generation Upgrade (PINGU)" (2014),
arXiv: 1401.2046 [physics.ins-det].

[55] A. M. Dziewonski and D. L. Anderson, "Preliminary reference Earth model",
Physics of the Earth and Planetary Interiors 25.4 (1981) 297–356, ISSN: 0031-9201,
DOI: 10.1016/0031-9201(81)90046-7.

[56] E. K. Akhmedov, "Parametric resonance of neutrino oscillations and passage of solar and atmospheric neutrinos through the earth", *Nucl. Phys. B* 538.1–2 (1999) 25–51,
DOI: 10.1016/S0550-3213(98)00723-8.

[57] E. K. Akhmedov et al., "Atmospheric neutrinos at super-Kamiokande and parametric resonance in neutrino oscillations", *Nucl. Phys. B* 542.1–2 (1999) 3–30,
DOI: 10.1016/S0550-3213(98)00825-6.

[58] B. T. Cleveland et al.,
"Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector",
The Astrophysical Journal 496.1 (1998) 505.

[59] Y. Fukuda et al., "Solar Neutrino Data Covering Solar Cycle 22",
Phys. Rev. Lett. 77 (9 1996) 1683–1686, DOI: 10.1103/PhysRevLett.77.1683.

[60] Q. Ahmad et al., "Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by ^{8}B Solar Neutrinos at the Sudbury Neutrino Observatory", *Phys. Rev. Lett.* 87 (7 2001) 071301,
DOI: 10.1103/PhysRevLett.87.071301.

[61] B. Aharmim et al., "Combined analysis of all three phases of solar neutrino data from the Sudbury Neutrino Observatory", *Phys. Rev. C* 88 (2 2013) 025501,
DOI: 10.1103/PhysRevC.88.025501.

[62] M. Aartsen et al., "Determining neutrino oscillation parameters from atmospheric muon neutrino disappearance with three years of IceCube DeepCore data" (2014),
arXiv: 1410.7227 [hep-ex].

[63] U. F. Katz, "The ORCA Option for KM3NeT", *ArXiv e-prints* (2014),
arXiv: 1402.1022 [astro-ph.IM].

[64] P. Adamson et al., "Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS", *Phys. Rev. Lett.* 110 (25 2013) 251801,
DOI: 10.1103/PhysRevLett.110.251801.

[65] K. Abe et al., "Precise Measurement of the Neutrino Mixing Parameter ϑ_{23} from Muon Neutrino Disappearance in an Off-Axis Beam", *Phys. Rev. Lett.* 112 (18 2014) 181801, DOI: 10.1103/PhysRevLett.112.181801.

[66] K. Abe et al., "Observation of Electron Neutrino Appearance in a Muon Neutrino Beam", *Phys. Rev. Lett.* 112 (6 2014) 061802, DOI: 10.1103/PhysRevLett.112.061802.

[67] N. Agafonova et al., "Evidence for $\nu_\mu \to \nu_\tau$ appearance in the CNGS neutrino beam with the OPERA experiment", *Phys. Rev.* D89 (2014) 051102, arXiv: 1401.2079 [hep-ex].

[68] F. An et al., "Observation of Electron-Antineutrino Disappearance at Daya Bay", *Phys. Rev. Lett.* 108 (17 2012) 171803, DOI: 10.1103/PhysRevLett.108.171803.

[69] J. Ahn et al., "Observation of Reactor Electron Antineutrinos Disappearance in the RENO Experiment", *Phys. Rev. Lett.* 108 (19 2012) 191802, DOI: 10.1103/PhysRevLett.108.191802.

[70] "Background-independent measurement of in Double Chooz", *Physics Letters B* 735 (2014) 51–56, ISSN: 0370-2693, DOI: 10.1016/j.physletb.2014.04.045.

[71] M. Blennow and T. Schwetz, "Determination of the neutrino mass ordering by combining PINGU and Daya Bay II", *JHEP* 1309 (2013) 089, DOI: 10.1007/JHEP09(2013)089.

[72] Y.-F. Li et al., "Unambiguous determination of the neutrino mass hierarchy using reactor neutrinos", *Phys. Rev. D* 88 (1 2013) 013008, DOI: 10.1103/PhysRevD.88.013008.

[73] Y.-F. Li, "Overview of the Jiangmen Underground Neutrino Observatory (JUNO)", *Int.J.Mod.Phys.Conf.Ser.* 31 (2014) 1460300, DOI: 10.1142/S2010194514603007.

[74] G. Fogli et al., "Global analysis of neutrino masses, mixings and phases: entering the era of leptonic CP violation searches", *Phys. Rev.* D86 (2012) 013012, DOI: 10.1103/PhysRevD.86.013012.

[75] F. Capozzi et al., "Status of three-neutrino oscillation parameters, circa 2013", *Phys. Rev.* D89 (2014) 093018, DOI: 10.1103/PhysRevD.89.093018.

[76] M. Gonzalez-Garcia, "Global analyses of oscillation neutrino experiments", *Physics of the Dark Universe* 4 (2014) 1–5, ISSN: 2212-6864, DOI: 10.1016/j.dark.2014.04.002.

[77] 19th Dec. 2014, URL: http://code.icecube.wisc.edu/svn/sandbox/atmo-weights/branches/PINGU_processing%20/.

[78] M. Ackermann et al., "Optical properties of deep glacial ice at the South Pole", *Journal of Geophysical Research: Atmospheres* 111.D13 (2006), ISSN: 2156-2202, DOI: 10.1029/2005JD006687.

[79] G. M. Hale and M. R. Querry, "Optical Constants of Water in the 200-nm to 200-μm Wavelength Region", *Appl. Opt.* 12.3 (1973) 555–563, DOI: 10.1364/AO.12.000555.

[80] J. Ahrens et al., "IceCube Preliminary Design Report", tech. rep. 1.24, The IceCube Collaboration, Oct. 2001.

[81] C. Wiebusch, "Physics Capabilities of the IceCube DeepCore Detector", *Proceedings of the 31st ICRC* (2009), arXiv: 0907.2263 [astro-ph.IM].

[82] P. Sandstrom, *Generation 2 Digital Optical Module*, Talk at Aachen workshop on next-generation hardware, Dec. 2014.

[83] K. Hanson and O. Tarasova, "Design and production of the IceCube digital optical module", *Nucl. Instr. Meth.* A567 (2006) 214–217, DOI: 10.1016/j.nima.2006.05.091.

[84] R. Abbasi et al., "Calibration and characterization of the IceCube photomultiplier tube", *Nucl. Instr. Meth. A* 618.1-3 (2010) 139–152, ISSN: 0168-9002, DOI: 10.1016/j.nima.2010.03.102.

[85] Hamamatsu Corporation, *Hamamatsu R7081-02 data sheet*.

[86] J. Fordyce, K. Clark and D. Grant, *Trigger Efficiencies*, Presentation at PINGU Analysis Call, 1st Oct. 2014.

[87] *Wavedeform Documentation*, 19th Feb. 2015, URL: http://software.icecube.wisc.edu/simulation_trunk/projects/wavedeform/index.h%20tml.

[88] C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems (Classics in Applied Mathematics)*, Society for Industrial and Applied Mathematics, 1987, ISBN: 0898713560.

[89] *SLC hit cleaning*, 19th Feb. 2015, URL: https://wiki.icecube.wisc.edu/index.php/SLC_hit_cleaning.

[90] Pat Toale, *CLast Documentation*, 28th Jan. 2015, URL: http://software.icecube.wisc.edu/icerec_trunk/doxygen/clast/index.html.

[91] *PINGU Standard Processing*, 29th Jan. 2015, URL: http://code.icecube.wisc.edu/svn/sandbox/pingu-std-proc/trunk/scripts/.

[92] J. Lundberg et al., "Light tracking through ice and water—Scattering and absorption in heterogeneous media with Photonics", *Nucl. Instr. Meth. A* 581.3 (2007) 619–631, ISSN: 0168-9002, DOI: 10.1016/j.nima.2007.07.143.

[93] M. Aartsen et al., "Measurement of South Pole ice transparency with the IceCube LED calibration system", *Nucl. Instr. Meth. A* 711 (2013) 73–89, ISSN: 0168-9002, DOI: 10.1016/j.nima.2013.01.054.

[94] N. Whitehorn, J. van Santen and S. Lafebre, "Penalized Splines for Smooth Representation of High-dimensional Monte Carlo Datasets", *Comput.Phys.Commun.* 184 (2013) 2214–2220, DOI: 10.1016/j.cpc.2013.04.008.

[95] *Millipede Documentation*, 28th Jan. 2015, URL: http://software.icecube.wisc.edu/icerec_trunk/projects/millipede/index.html.

[96] A. A. Aguilar-Arevalo et al., "First measurement of the muon neutrino charged current quasielastic double differential cross section", *Phys. Rev. D* 81 (9 2010) 092005, DOI: 10.1103/PhysRevD.81.092005.

[97] F. Feroz, M. P. Hobson and M. Bridges, "MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics", *Monthly Notices of the Royal Astronomical Society* 398.4 (2009) 1601–1614, DOI: 10.1111/j.1365-2966.2009.14548.x.

[98] F. Feroz et al., "Importance Nested Sampling and the MultiNest Algorithm", *ArXiv e-prints* (2013), arXiv: 1306.2144 [astro-ph.IM].

[99] M. J. Schervish, "A Review of Multivariate Analysis", *Statist. Sci.* 2.4 (1987) 396–413, DOI: 10.1214/ss/1177013111.

[100] J. P. A. M. de André, *Reconstruction methods & performances (tracks, cascades, PID) in ice*, Talk at MANTS-GNN meeting, Sept. 2014.

[101] S. Adrián-Martínez et al., "Deep sea tests of a prototype of the KM3NeT digital optical module", *Eur. Phys. J. C* 74 (2014) 3056, DOI: 10.1140/epjc/s10052-014-3056-3.

[102] A. Kappes, *mDOM for Next Generation IceCube*, Talk at MANTS-GNN meeting, Sept. 2014.

[103] L. Schulte et al., "A large-area single photon sensor employing wavelength-shifting and light-guiding technology", *Proceedings of the 33rd ICRC* (2013), arXiv: 1307.6713 [astro-ph.IM].

[104] T. G. P. Inc., *Fused Quartz Properties*, 30th Mar. 2015, URL: http://www.technicalglass.com/index.html.

[105] H. Corporation, private communication.

[106] H. O. Meyer, "Spontaneous electron emission from a cold surface", *EPL* 89.5 (2010) 58001, DOI: 10.1209/0295-5075/89/58001.

[107] E. Technology, *EJ-298 data sheet*, 30th Mar. 2015, URL: http://www.eljentechnology.com/images/stories/Data_Sheets/Paints_Reflectors/EJ2%2098#2%20data%20sheet.pdf.

[108] S. G. Crystals, *Low level/low background counting*, 30th Mar. 2015, URL: http://www.detectors.saint-gobain.com/document.aspx?docId=201358.

[109] D. Hebecker, "Development of a Single Photon Detector with Wavelength Shifting and Light Guiding Technology", MA thesis: Universität Bonn, 2014.

[110] C. Andreopoulos et al., "The GENIE neutrino Monte Carlo generator", *Nucl. Instr. Meth. A* 614.1 (2010) 87–104, ISSN: 0168-9002, DOI: 10.1016/j.nima.2009.12.009.

[111] *Neutrino Generator*, 18th Feb. 2015, URL: http://code.icecube.wisc.edu/svn/projects/neutrino-generator/.

[112] *PINGU Simulation File Summary*, 18th Feb. 2015, URL: https://wikispaces.psu.edu/display/PINGU/PINGU+Simulation+File+Summary.

[113] D. Chirkin and W. Rhode, "Muon Monte Carlo: A High-precision tool for muon propagation through matter" (2004), arXiv: hep-ph/0407075 [hep-ph].

[114] *clsim Documentation*, 18th Feb. 2015, URL: http://software.icecube.wisc.edu/simulation_trunk/projects/clsim/.

[115] S. Agostinelli et al., "Geant4—a simulation toolkit", *Nucl. Instr. Meth. A* 506.3 (2003) 250–303, ISSN: 0168-9002, DOI: 10.1016/S0168-9002(03)01368-8.

[116] J. Allison et al., "Geant4 developments and applications",
IEEE Trans. Nucl. Sci. 53.1 (2006) 270–278, DOI: 10.1109/tns.2006.869826.

[117] J. E. Stone, D. Gohara and G. Shi,
"OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems",
IEEE Des. Test 12.3 (2010) 66–73, ISSN: 0740-7475, DOI: 10.1109/MCSE.2010.69.

[118] M. Larsen, "Simulation and Identification of Non-Poissonian Noise Triggers in the IceCube Neutrino Detector",
MA thesis: Department of Physics in the Graduate School of The University of Alabama, 2013.

[119] *DOMLauncher Documentation*, 18th Feb. 2015, URL: http://software.icecube.wisc.edu/simulation_trunk/projects/DOMLauncher/index.ht%20ml.

[120] *PMTResponseSimulator Documentation*, 18th Feb. 2015, URL: http://software.icecube.wisc.edu/simulation_trunk/projects/DOMLauncher/PMTRes.h%20tml.

[121] *PaPA Source Code*, 19th Feb. 2015,
URL: http://code.icecube.wisc.edu/svn/sandbox/PaPA/.

[122] *PaPA Webpage*, 19th Feb. 2015,
URL: https://wiki.icecube.wisc.edu/index.php/PaPA.

[123] G. Rossum, "Python Reference Manual", tech. rep., 1995.

[124] S. v. d. Walt, S. C. Colbert and G. Varoquaux,
"The NumPy Array: A Structure for Efficient Numerical Computation",
Computing in Science & Engineering 13.2 (2011) 22–30, DOI: 10.1109/MCSE.2011.37.

[125] E. Jones, T. Oliphant, P. Peterson et al., *SciPy: Open source scientific tools for Python*, [Online; accessed 2015-02-19], 2001–.

[126] J. D. Hunter, "Matplotlib: A 2D Graphics Environment",
Computing in Science & Engineering 9.3 (2007) 90–95, DOI: 10.1109/MCSE.2007.55.

[127] M. Wallraff and C. Wiebusch,
"Calculation of oscillation probabilities of atmospheric neutrinos using nuCraft",
ArXiv e-prints (2014), arXiv: 1409.1387 [astro-ph.IM].

[128] M. Honda et al., "Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data", *Phys. Rev. D* 75 (4 2007) 043006,
DOI: 10.1103/PhysRevD.75.043006.

[129] M. Honda, *HKKM2014 flux*, 24th Feb. 2015,
URL: http://www.icrr.u-tokyo.ac.jp/~mhonda/nflx2014/index.html.

[130] A. Groß, *Update of PINGU background studies*, Presentation at PINGU Analysis Call, Oct. 2014.

[131] M. Tegmark, A. N. Taylor and A. F. Heavens, "Karhunen-Loève Eigenvalue Problems in Cosmology: How Should We Tackle Large Data Sets?", *ApJ* 480.1 (1997) 22.

[132] M. Kowalski, "Measuring diffuse neutrino fluxes with IceCube", *JCAP* 0505 (2005) 010,
DOI: 10.1088/1475-7516/2005/05/010.

[133] A. Albrecht et al., "Report of the Dark Energy Task Force",
ArXiv Astrophysics e-prints (2006), eprint: astro-ph/0609591.

[134] A. Albrecht et al.,
"Findings of the Joint Dark Energy Mission Figure of Merit Science Working Group",
ArXiv e-prints (2009), arXiv: 0901.0721 [astro-ph.IM].

[135] D. Coe, "Fisher Matrices and Confidence Ellipses: A Quick-Start Guide and Software",
ArXiv e-prints (2009), arXiv: 0906.4123 [astro-ph.IM].

[136] C. Finley, "Weighting NeutrinoGenerator Events with OneWeight in IceSim 2.0", tech. rep.,
IceCube, Oct. 2007.

[137] J. P. A. M. de André, *PINGU V36 Particle Flavour Identification Efficiencies*,
private communication.

[138] A. Himmel, "Recent results from Super-Kamiokande", *AIP Conf.Proc.* 1604 (2014) 345–352,
DOI: 10.1063/1.4883450.

[139] JUNO collaboration, "JUNO Technical Design Report",
by private communication with M. Wurm, 2014.

[140] J. W. Cooley and J. W. Tukey,
"An algorithm for the machine calculation of complex Fourier series",
Math. Comp. 19.90 (1965) 297–297, DOI: 10.1090/s0025-5718-1965-0178586-1.

[141] N. A. Thacker, *Error Propagation and the Fourier Transform*, 10th Feb. 2015,
URL: http://www.tina-vision.net/teaching/cvmsc/docs/fourier.ps.

[142] H. Minakata et al., "Determining neutrino mass hierarchy by precision measurements in
electron and muon neutrino disappearance experiments", *Phys. Rev. D* 74 (5 2006) 053008,
DOI: 10.1103/PhysRevD.74.053008.

[143] *The common PINGU simulation and analysis code for the neutrino mass hierarchy*,
26th Mar. 2015, URL: https://github.com/tarlen5/pisa.

[144] OSI Optoelectronics, *UV-035EQC data sheet*.

List of Figures

List of Tables

Danksagung

Und das war es jetzt... Fehlt nur noch ein ganz großes Dankeschön an alle, ohne die diese Arbeit so nicht möglich und die Zeit nicht so schön gewesen wäre.

Zuerst natürlich an Marek Kowalski, der mir schon früh die Promotionsstelle zugesagt und mir sogar – lang, lang ists her – zu meiner Diplomarbeit schon ein paar nützliche Tips gegeben hat. Auch später war trotz vollen Terminplans und am Schluss auch großer Entfernung (fast) immer Zeit für ein paar hilfreiche Worte. Danke auch an Norbert Wermes für die Übernahme der Zweitkorrektur.

Dann der Bonner Doppel-Arbeitskreis. Sehr spannende Konstruktion, zwei doch ziemlich unterschiedliche Themen unter einem Dach. Erst mal vielen Dank an alle für die schöne Arbeitsatmosphäre, den vielen leckeren Kuchen und die allmittägliche Tee- und Kaffeepause. Besonderen Dank an meine fleißigen Korrekturleser Andreas, Markus, Alex und Marcel. Und weil Ruth auch was gelesen hat, sind meine Ergüsse hoffentlich auch für Nicht-IceCuber verständlich.

Sebastian hat zwar nichts gelesen, aber trotzdem wohl am meisten beigetragen. Ob beim Programmieren, Debuggen, bei der Analyse oder bei Präsentationen: Immer gab es hilfreiche Kommentare und Anregungen. Unter denen mag ich manchmal auch etwas gestöhnt haben, aber ein Faultier wie ich braucht den gelegentlichen Tritt in den Hintern. Und zur Belohnung ist er ja jetzt Prof im goldische Meenz.

Womit wir auch schon bei der Family wären. Allen voran meine Schwestern Judith und Johanna, die zwar mittlerweile schon richtig erwachsen, aber trotzdem noch die lustigsten und liebsten Schwestern sind, die man haben kann. Außerdem können sie mich vermutlich beide unter den Tisch trinken – Studenten halt. Dazu mein Onkel Mani, der tatsächlich fast gleichzeitig mit mir aus Mainz gen Norden gezogen ist, allerdings nicht ganz so weit. Und meine Eltern: Danke für alles, trotz Allem.

Und zum Schluss noch einmal nach Bonn, zu den Polyphonikern: Ein toller Chor, nicht nur musikalisch, mit dem ich mich hier direkt heimisch gefühlt habe. Dort hinzugehen, war eine der, wenn nicht die beste und wichtigste Entscheidung meines Lebens.

Denn jetzt habe ich die beste Michi der Welt. Du hast mir das verrückteste und schönste Jahr meines Lebens geschenkt, dem noch bessere folgen werden. Und den besten Grund überhaupt, endlich mal fertig zu werden.

Das bin ich hiermit.