Sensitivity study and first prototype tests for the CHIPS neutrino detector R&D program

Maciej Marek Pfützner
University College London

Submitted to University College London in fulfilment
of the requirements for the award of the
degree of **Doctor of Philosophy**

July 20, 2018

Declaration

I, Maciej Marek Pfützner confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Maciej Pfützner _____

Abstract

CHIPS (CHerenkov detectors In mine PitS) is an R&D project aiming to develop novel cost-effective detectors for long baseline neutrino oscillation experiments. Water Cherenkov detector modules will be submerged in an existing lake in the path of an accelerator neutrino beam, eliminating the need for expensive excavation. In a staged approach, the first detectors will be deployed in a flooded mine pit in northern Minnesota, 7 mrad off-axis from the existing NuMI beam. A small proof-of-principle model (CHIPS-M) has already been tested and the first stage of a fully functional 10 kt module (CHIPS-10) is planned for 2018.

The main physics aim is to measure the CP-violating neutrino mixing phase (δ_{CP}). A sensitivity study was performed with the GLoBES package, using results from a dedicated detector simulation and a preliminary reconstruction algorithm. The predicted physics reach of CHIPS-10 and potential bigger modules is presented and compared with currently running experiments and future projects.

One of the instruments submerged on board CHIPS-M in autumn 2015 was a prototype detection unit, constructed at Nikhef. The unit contains hardware borrowed from the KM3NeT experiment, including 16 3 inch photomultiplier tubes and readout electronics. In addition to testing the mechanical design and data acquisition, the detector was used to record a large sample of cosmic ray muon events. A preliminary analysis of the collected data was performed, in order to measure the cosmic background interaction rates and validate the Monte Carlo simulation used to optimise future designs.

The first in situ measurement of the cosmic muon rate at the bottom of the Wentworth Pit is presented, and extrapolated values for CHIPS-10 show that the dead time due to muons is below 0.3 %.

Acknowledgements

This work could not have happened without the help and contribution of many others. I am very grateful to everyone who aided me in this journey, and I would like to mention the following people in particular.

First and foremost I thank my supervisor, Prof. Jenny Thomas. Her advice and support were invaluable during all stages of my PhD at UCL, from the application process all the way to the viva examination. She always provided constructive comments on any of my research problems and word of encouragement to doubts or worries. Jenny also exercised great leadership of the CHIPS project, keeping all of us focused and driven even during the most hectic times of the testing campaigns; all in addition to after-hours tutorials in table etiquette.

Big thanks to Prof. Paul Kooijman, my host and supervisor during the visit at Nikhef. Working together on the design of CHIPS was a great experience, with Paul always having an abundance of new ideas. I found our long discussions about them very stimulating and enjoyable.

I am thankful to Leigh Whitehead, who acted as a mentor, colleague, and a friend, for all the advice on both technical issues and life philosophy, all delivered with sarcastic humour. Also thanks to the other members of the UCL MINOS/CHIPS group, including Andy Perch and Stefano Germani for many helpful discussions, and Ryan Nichol and Anna Holin for assistance with MINOS-related projects.

The two summers spent in Minnesota working on CHIPS-M were the highlight of my PhD. I wish to thank everyone who contributed to the success of the project and the great friendly atmosphere, particularly the versatile CHIPS heroes: Richard, Josh, Brandon, and Ramon – with whom I spent many hours crimping cables, soldering and testing electronic components. Special thanks go to Jack Zorman who taught

me many useful manual skills (along with colourful shop vernacular), Jerry Meier, and the rest of the Soudan laboratory crew.

Thanks to all the great people from the UCL HEP group, including my various office mates and fellow PhD students from my year, for creating the welcoming and inspiring environment where I felt right at home. Particularly warm thanks to Laura Manenti for the moral support and unwavering friendship, especially precious in the final weeks of writing up.

Last but not least, I thank my parents, Marek and Małgosia, for all their love and encouragement on my path as a physicist and beyond.

Contents

List of figures

List of tables

Introduction

Neutrinos are extremely abundant fundamental particles, yet very challenging to detect and study. From the first hypothesis of their existence, through their role in the discovery of weak neutral currents, to the surprising observation of flavour oscillations, neutrinos provide many insights into the workings of the universe. Despite the great success of the three-flavour model of neutrino oscillations, and the information available by studying them, there are still important questions in neutrino physics which remain unanswered. Long-baseline oscillation experiments with accelerator neutrino beams are in a great position to solve some of those problems, but every new generation of detectors and experiments becomes more and more expensive.

This thesis describes CHIPS, an R&D project aiming to mitigate the issue by developing cost-effective and scalable water Cherenkov detectors. The cost reduction comes from submerging the detector structure in an existing water body instead of an underground cavern, as well as using small photomultiplier tubes with a layout optimised for the beam, and a minimal water filtration system. The first stage of a fully functional detector module, CHIPS-10, will be deployed in summer 2018 in a flooded mine pit in northern Minnesota, in the path of the NuMI neutrino beam. First prototypes have already been tested at that location, when CHIPS-M, a small-scale test-bed platform, was deployed with five IceCube optical modules in 2014, and again in 2015, with two additional prototype detection units (or planes).

Chapter 1 presents an overview of neutrino oscillations and their experimental status. First, the history of neutrino physics and their role in the Standard Model are briefly described. The formalism of neutrino oscillations is introduced and the experimental evidence of this phenomenon presented, including the historical and the latest results. The important open questions in neutrino physics are presented, with focus on neutrino oscillations, and the role of current and future long baseline accelerator experiments in solving them is discussed.

Chapter 2 introduces the CHIPS R&D project. The NuMI neutrino beam is described and the selected CHIPS location in the path of the beam is characterised. Next, the core concepts of the CHIPS water Cherenkov detectors are explained, and the design of the upcoming CHIPS-10 detector module is presented. Finally, the software effort is described, including the Monte Carlo detector simulation and a preliminary track reconstruction algorithm, employing a maximum likelihood method.

Chapter 3 contains the results of a sensitivity study performed with the GLoBES package. It is the first prediction of the physics reach of experiments with the CHIPS detectors that is based on the performance of the new reconstruction algorithm. The chapter describes the GLoBES framework and the used experimental definition of CHIPS. The sensitivity to CP violation and to the mass hierarchy is studied for various exposures and run scenarios. The results are compared with those of the currently running experiments, NOvA and T2K. A possible deployment of CHIPS in the LBNF beam is also considered and contrasted with the early performance of the DUNE experiment.

Chapter 4 describes the prototype testing campaigns conducted in 2014 and 2015 at the Wentworth Pit. First, the 2014 campaign is reported, including the design and construction of the CHIPS-M detector platform, the instrumentation with IceCube DOMs, the deployment and underwater operation, and the recovery in summer 2015. Next, the Nikhef prototype detection unit is described, detailing the KM3NeT readout hardware, and the design, construction and testing of two design iterations of the plane. The second testing campaign, during which the Nikhef plane was tested, is then discussed, outlining the modifications to CHIPS-M, a second prototype unit, and the deployment and operation of the detector in autumn 2015.

Chapter 5 contains the description and results of a preliminary analysis of the cosmic muon data recorded by the prototype plane. The flow of data, from the acquisition of raw hits to event reconstruction, is depicted. The Monte Carlo simulation of cosmic ray events in CHIPS-M is described. Next, the PMT time calibration with an IceCube LED flasher is discussed, along with results and unresolved issues. The following section details the validation and tuning of the simulation ending up with the measurement of the total cosmic muon rate in CHIPS-M. The result is then used to predict the muon rate and the effective dead time in the CHIPS-10 detector.

Author's contribution

Particle physics is a very collaborative field and the CHIPS project is no exception. The final work presented in this thesis would not have been possible without the input of many people. This section lists the individual contributions of the author to the presented work. The main pieces of original work are the use of GLoBES for the sensitivity predictions, an important role in the hardware construction effort, especially of the Nikhef plane, and the full analysis of the cosmic muon data recorded with the plane.

Chapter 2 contains the description of simulation and reconstruction software. Most of this work has been performed by Leigh Whitehead and Andy Perch and is described in detail in Ref. [1]. The first attempt at implementing the time likelihood calculation part (described in Sec. 2.4.3) was an original work of the author, however this part did not end up being included in the main reconstruction code. Additionally, the author was involved in the general effort of developing, testing and maintaining the simulation and reconstruction code base. The preliminary analysis presented at the end of Chapter 2, the results of which are used in the sensitivity study, was performed mainly by Andy Perch. On the other hand, the application of the Monte Carlo simulation during the cosmic muon analysis in Chapter 5, was done by the author.

The sensitivity study presented in Chapter 3 is one of the main pieces of original work by the author. Although it is based on other work presented earlier, and makes use of an external software package, the user code was created by the author from scratch. Some modifications to the previous procedure have also been done, and the study made use of a new experiment definition, compiled by the author.

The author was actively engaged in the hardware work on the construction, testing, deployment and operation of the prototype instruments, including the CHIPS-M structure and the Nikhef prototype detection unit, described in Chapter 4. During the first testing campaign, the main responsibilities included co-creating and supervising the process of installing the detector liner, developing a feed-through for the instrumentation cables and pipes that formed the umbilical cable, and designing and constructing the steel tripod for the structure. In addition, the author

was actively involved in general construction of the CHIPS-M platform, setup and operation of the IceCube DOMs, and various deployment activities.

For the second testing campaign, the author was the main person responsible for coordinating the development of the Nikhef prototype detection unit. The author visited the Nikhef institute in Amsterdam in order to establish collaboration with the KM3NeT experiment. This included engaging in testing work of KM3NeT DOMs, coordination of a student project focussed on constructing the first prototype, as well as active involvement in the final stages of construction and testing. The author was also responsible for organising the transport of the plane to Minnesota, the subsequent leak testing campaign, installation in the detector, and a complete setup of the local data acquisition software and hardware. After the initial deployment attempt, the author contributed to the construction of the second plane version and led the leak testing, also helping with testing other instruments and the deployment of the CHIPS-M. Finally, the author was responsible for recording and storing the data from the plane.

Last but not least, the preliminary analysis of the data recorded with the Nikhef plane during deployment in the Wentworth Pit, presented in Chapter 5 is entirely an original piece of work by the author.

Chapter 1

Theory and status of neutrino oscillations

Neutrinos are one of the fundamental particles of matter, in fact the second most abundant species after the photons. Billions of neutrinos produced in the Sun and Earth's upper atmosphere pass through our bodies every second, without any interaction at all. A single neutrino is extremely unlikely to interact even when passing through the whole diameter of the Earth. In fact, the mean free path of a 100 MeV neutrino is more than a light year in lead [2].

Neutrinos were first hypothesised almost ninety years ago and experimentally discovered more than sixty years ago [3, 4]. Used as probes of the weak interactions and nucleon structure, they played a big role in establishing the Standard Model. They were also studied by astrophysicists, trying to learn more about the Sun and supernovae, and as unavoidable background in nucleon decay experiments. With the surprising discovery of oscillations in the 2000s, proving that despite theoretical predictions neutrinos are massive particles, neutrino physics has gained additional recognition. Now, the occurrence of oscillations is firmly established and the experiments are entering a phase of precision measurements. However, some questions remain open, and their answers are believed to contain deep insights into physics beyond the Standard Model.

1.1 Historical background

The history of neutrino physics starts in 1930 when W. Pauli proposed their existence as a solution to the problem of energy and angular momentum conservation in beta decay [3]. The measured continuous energy spectrum of the electron did not fit one expected in a two-body decay, if the electron and nucleon were the only decay products. The paradox could be solved if a third particle, electrically neutral and with low mass and spin 1/2, was also created in the interaction. This particle is now known as the electron antineutrino.

Soon afterwards, E. Fermi developed the first theory of weak interactions and in 1934 H. Bethe and R. Peierls used it to estimate the expected cross-section for the inverse beta decay process, in which a neutrino could be observed [5]. The very low result, less than $10^{-44}\,\mathrm{cm}^2$, showed the enormous challenge of experimental neutrino physics, requiring very large intensities and detector volumes to observe significant event rates.

The first unambiguous direct observation came more than twenty years later, by C. Cowan and F. Reines, in 1956[1] [4]. The Savannah River experiment detected electron antineutrinos produced in the core of a nuclear reactor of the same name. The detector used 200 litres of cadmium-doped water as the target and 1400 litres of liquid scintillator, instrumented with 100 photomultiplier tubes (PMTs), as the active material. Neutrinos were detected via inverse beta decay and the experimental signature consisted of back-to-back γ emission from the positron annihilation and a delayed signal from the neutron capture. The coincidence of those two signals allowed the rejection of the very large cosmic ray background and the observation of a significant event rate compared to when the reactor was turned off [4].

In 1962 at the Alternating Gradient Synchrotron at Brookhaven, neutrinos created in pion decays, together with muons, were observed to produce only muons, not electrons, in subsequent interactions in the detector [7]. This confirmed the existence of the muon neutrino, distinct from the previously known electron neutrino.

In 1973, the Gargamelle experiment at CERN, using an accelerator-produced neutrino beam and a big bubble chamber detector, discovered the existence of weak

[1] The first hints appeared in their earlier experiment in 1953 [6], but the massive background prevented the result from being statistically significant.

neutral currents [8]. This proved the existence of the neutral Z boson and confirmed the Glashow-Salam-Weinberg theory of electroweak interactions. This theory is now a crucial part of the Standard Model and describes neutrino interactions in excellent agreement with experimental data.

Experiments running at the LEP e^+ e^- collider in the 1990s performed a precision measurement of the Z decay width, $\Gamma_Z = (2.4952 \pm 0.0023)\,\mathrm{GeV}$ [9]. A fit to the data showed that there are exactly three active generations of neutrinos (lighter than half of the Z mass). In 2000, the DONUT experiment at the Tevatron collider in Fermilab performed a direct detection of the tau neutrino [10], completing the three flavour picture.

1.2 Neutrinos in the Standard Model

The best theory describing the behaviour of subatomic particles and incorporating three out of four known fundamental forces is the Standard Model. It is a gauge theory, in which the Langrangian obeys local gauge symmetries described by the groups $U(1) \times SU(2) \times SU(3)$.

According to the model, all matter is made of fermions – particles with spin 1/2. There are six quarks and six leptons, both groups divided into three generations or *flavours*. The particles and their main properties are listed in Fig. 1.1. In addition, all particles have a partner antiparticle, described by the same quantum numbers, but with opposite charge sign.

Each generation of leptons consists of a charged lepton and a neutrino in the corresponding flavour: electron and electron neutrino, muon and muon neutrino, and tau lepton and the tau neutrino. All the charged leptons have the same electric charge and well defined masses, while neutrinos are neutral and assumed to be massless

The six quarks – down, up, strange, charm, bottom and top – carry electric charges being a fraction of the electron charge, as indicated in Fig. 1.1. Their masses span a range of more than four orders of magnitude. In addition to flavour and the electric charge, quarks carry a *colour* charge, and each quark exists in one of the three colours. However, they do not exist in a free state (unless for a very short time) and

bind together to create colourless hadrons. These are typically divided into mesons, consisting of a quark and an antiquark (e.g. pions or kaons) and baryons, made up of three quarks or three antiquarks (e.g. protons, neutrons and their antiparticles).

The three forces – electromagnetic, weak and strong – are carried by boson fields, manifesting as spin 1 particles. The electromagnetic force is propagated by the massless photon and affects all electrically charged particles. It is responsible for almost all phenomena occurring on the atomic scale. The massive W and Z bosons mediate the weak force, which is associated with nuclear decay processes. The strong force is what binds the quarks into hadrons and is responsible for short-range interactions between protons and neutrons in the nucleus. It affects only coloured particles, and is mediated by eight massless gluons, which also carry a colour charge themselves.

The final element of the SM is the scalar Higgs boson, which gives mass to the fundamental particles in the process of spontaneous symmetry breaking to W and Z bosons, and to fermions, via Yukawa couplings.

Gravity is not incorporated into the Standard Model, but it is many orders of magnitude weaker than all the other forces, and has usually no significant effect at the subatomic scale.

Neutrinos are neutral leptons and therefore interact only weakly, exchanging the Z and W boson mediators with other particles or bound states (e.g. atomic nuclei). These interactions are described by the electroweak theory, also known as the Glashow-Salam-Weinberg model.

Interactions with the neutral Z boson are called *neutral current*, or NC, interactions. An example Feynman diagram of a neutrino scattering on a nucleon is shown on Fig. 1.2. Since no charge or flavour are exchanged and the cross section is the same for all types of neutrinos, it is impossible to determine the neutrino flavour in NC interactions. At higher energies, NC interactions are accompanied by a pion or a hadronic shower. NC processes, where a neutral π^0 is created, are a particularly important background for CHIPS, as they can mimic electrons in water Cherenkov detectors.

In contrast, *charged current*, or CC, interactions occur via the exchange of a W boson. At the boson vertex, the neutrino is transformed to a charged lepton

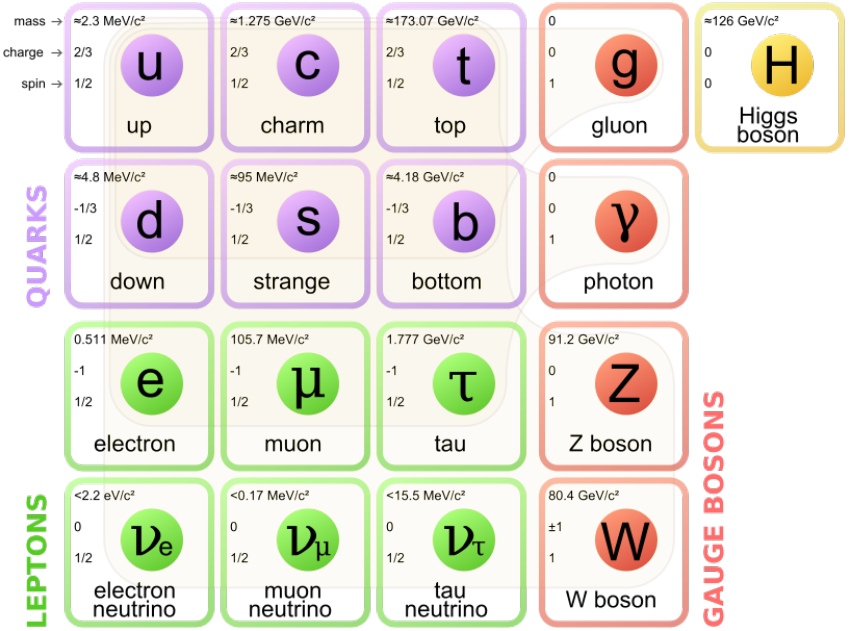

Figure 1.1: Particles of the Standard Model [11]. Leptons and quarks make up the matter, while the gauge bosons mediate forces and the Higgs is responsible for mass.

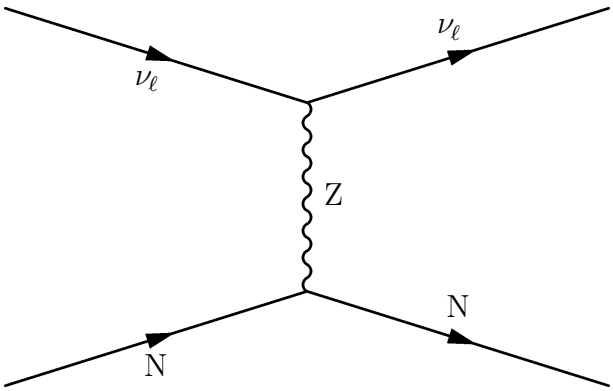

Figure 1.2: Feynman diagram of a neutrino with flavour ℓ undergoing neutral current elastic scattering on a nucleon. Without a charged lepton involved in the interaction, the flavour of the neutrino cannot be identified.

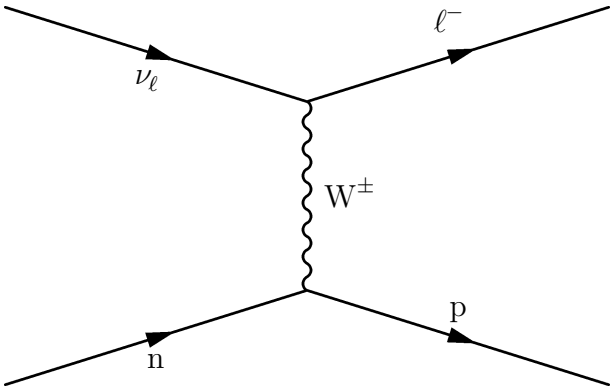

Figure 1.3: Feynman diagram of a charged current quasi-elastic interaction of a ν_ℓ with a neutron. Identification of the charged lepton allows determination of the original neutrino flavour

with the same flavour, conserving charge, flavour and lepton number. Figure 1.3 shows a possible CC interaction of a neutrino with a neutron. This is an example of a *quasi-elastic* (QE) interaction, where no additional particles are created. CC interactions can also be classified as *resonant scattering* (RES), if there is a resonant creation of a hadron, typically a Δ decaying later to neutron and pion, or *deep inelastic scattering* (DIS), where the neutrino interacts directly with a quark and produces a hadron shower accompanying the outgoing lepton.

Figure 1.4 shows measured and predicted CC interaction cross sections of muon neutrinos and antineutrinos for the three processes and all in total. The measurements come from a range of experiments, compiled in Ref. [12]. The predictions are generated with the NUANCE Monte Carlo generator [13]. Quasi-elastic interactions dominate at energies below 0.5 GeV, while at higher energies most interactions are deeply inelastic. Also worth noting is the energy threshold for any CC interaction, associated with the minimal energy required to produce the charged lepton. In case of muon neutrinos, the centre-of-mass energy of the neutrino and nucleon has to be at least 106 MeV in order to create the muon.

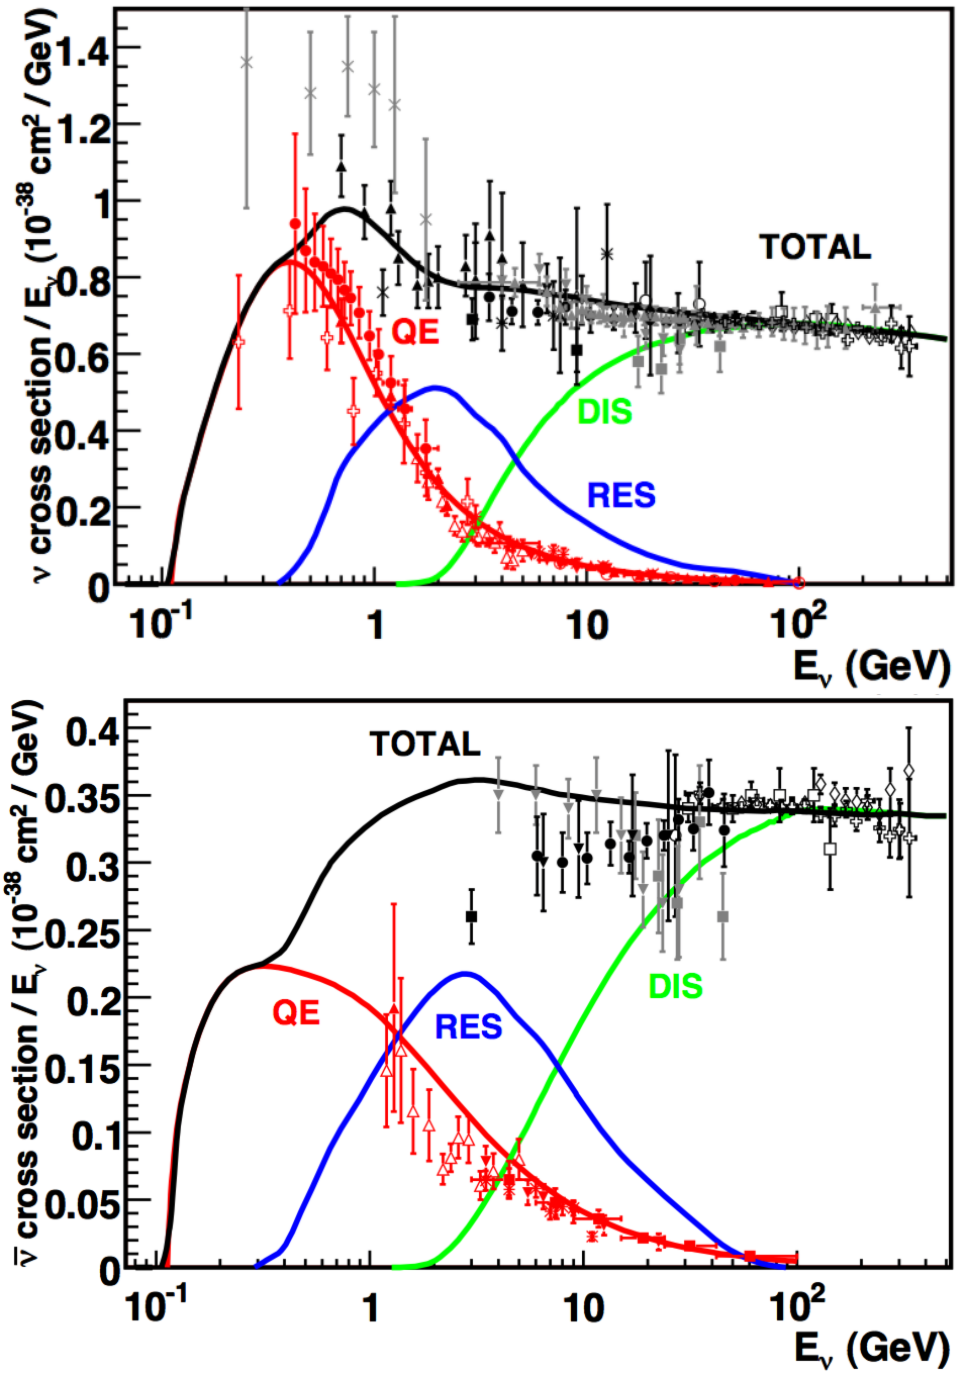

Figure 1.4: Charge current interaction cross sections for muon neutrinos (top) and antineutrinos (bottom) per nucleon for an isoscalar target. Points with error bars show measured values [12] and the solid lines show theoretical predictions [13]. Figure taken from Ref. [14].

1.3 Neutrino oscillations

In the original Standard Model, neutrinos are massless. The observation of neutrino oscillations invalidates this assumption by requiring at least two neutrino mass states to be non-zero. Although it is possible to modify the SM to accommodate this fact without requiring any new physics [15], neutrino oscillations are generally considered to be a significant break in the Standard Model. Precise measurement of the oscillation parameters will provide important clues and place constraints on New Physics theories.

Neutrino oscillations are a quantum phenomenon, analogous to the neutral kaon oscillations observed in the quark sector. If neutrinos are massive, the mass eigenstates are not necessarily equivalent to the weak interaction eigenstates. This is called mixing and it is described by the PMNS matrix U, named after B. Pontecorvo, Z. Maki, M. Nakagawa and S. Sakata. It can be seen as a rotation matrix transforming between the mass and flavour bases and is analogous to the CKM matrix, which describes quark mixing.

When propagating through space-time, the dispersion relation is well defined for the mass eigenstates, which means that the flavour composition of the neutrino changes with time. This change is periodic in time or distance, hence the name *oscillations*. To illustrate, a neutrino created in an interaction with a muon will have a definite flavour state ν_μ. However, after travelling a distance it may be detected as a ν_e or ν_τ.

1.3.1 Oscillations in vacuum

The derivation presented here is a simplified approximation using plane-wave representation of the neutrino states. The necessary assumptions are not fully consistent, however the final result, i.e. the oscillation probability formula, is exactly the same as for a thorough wave-packet treatment [16]. The main flaw of the presented approach is that plane waves are not localised. Since they have a well defined momentum, they spread over all of space-time and therefore cannot describe localised interactions during the production and detection of neutrinos. In addition, different neutrino eigenstates should have different energies and momenta, whereas the presented ap-

proximation assumes a common momentum. Both of those issues are solved when the neutrino states are described by wave packets with a spread of momenta and energies, localised at the interaction points. However, the presented derivation is much simpler while still illustrating well the conceptual points, and the final result is equally correct.

The unitary mixing matrix U describes the relation between flavour eigenstates $|\nu_\alpha\rangle$, where $\alpha = \mathrm{e}, \mu, \tau$, and mass eigenstates $|\nu_k\rangle$, $k = 1, 2, 3$, as follows

$$|\nu_\alpha\rangle = \sum_k U^*_{\alpha k} |\nu_k\rangle, \tag{1.1}$$

$$|\nu_k\rangle = \sum_\alpha U_{\alpha k} |\nu_\alpha\rangle. \tag{1.2}$$

The state $|\nu_k\rangle$ is an eigenstate of the Hamiltonian with energy eigenvalue E_k, and its time evolution follows the Schrödinger equation[2]

$$\frac{d}{dt}|\nu_k\rangle = E_k|\nu_k\rangle \Rightarrow |\nu_k(t)\rangle = e^{-iE_k t}|\nu_k\rangle. \tag{1.3}$$

Neutrinos are only created in weak interactions, and the resulting state (at the time of the interaction) is a pure weak eigenstate. Assume that $|\nu_\alpha\rangle$ is created at time $t = 0$ with four-momentum $\mathbf{p}$ so that $\mathbf{p}^2 = E^2 - m^2$. The evolution of the propagating flavour state can be expressed in the following way

$$|\nu_\alpha(t)\rangle = \sum_k U^*_{\alpha k} e^{-iE_k t} |\nu_k\rangle. \tag{1.4}$$

The neutrinos are detected by taking part in another weak interaction. In a charged current interaction, a neutrino of a definite flavour is absorbed and a charged lepton of the same flavour is produced. Detection of the charged lepton, and particularly measurement of its mass, is what makes it possible to identify the flavour of the incoming neutrino. The probability amplitude of observing the state

[2]Unless explicitly stated differently, the unit system used is such that $c = \hbar = 1$.

$|\nu_\beta\rangle$ is $\langle\nu_\beta|\nu_\alpha(t)\rangle$. The oscillation probability is therefore equal to

$$P(\nu_\alpha \to \nu_\beta; t) = \left|\langle\nu_\beta|\nu_\alpha(t)\rangle\right|^2 = \left|\sum_k U^*_{\alpha k} U_{\beta k} e^{-iE_k t}\right|^2 = \\ = \sum_k \sum_j U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} e^{-i(E_k - E_j)t}. \tag{1.5}$$

Because neutrinos are ultrarelativistic (for any practically considered energy range), some simplifying approximations can be made

$$E_k \simeq E + \frac{m_k^2}{2E} \Rightarrow E_k - E_j = \frac{\Delta m^2_{kj}}{2E}, \tag{1.6}$$

where $\Delta m^2_{kj} = m^2_k - m^2_j$ is the squared mass difference or mass splitting.

Additionally, the propagation time $t = ct \simeq L$, the distance from the creation point to detection. After applying these approximations, the expression inside the exponent becomes $\frac{\Delta m^2_{kj} L}{2E}$. Finally, the probability formula can be presented in the following form

$$P(\nu_\alpha \to \nu_\beta) = \delta_{\alpha\beta} - 4\sum_{k>j} \Re(U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}) \sin^2\left(\frac{\Delta m^2_{kj} L}{4E}\right) + \\ \pm 2\sum_{k>j} \Im(U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}) \sin\left(\frac{\Delta m^2_{kj} L}{4E}\right). \tag{1.7}$$

The last term has a positive sign for neutrinos and a negative sign for antineutrinos (i.e. $P(\bar\nu_\alpha \to \bar\nu_\beta)$).

The form of Eq. 1.7 reveals the periodic dependence on $\frac{\Delta m^2_{kj} L}{4E}$, which can be seen as the phase of oscillations, while the mixing matrix elements determine the amplitude of each term. When the phase is very small, i.e. when $\Delta m^2 \ll E/L$, the particular transition probability term goes to zero. On the other hand, when the phase is large, small differences in energy lead to fast changes in the probability. If the detection energy resolution is not fine enough, the $\sin^2(\Delta m^2 L/E)$ term averages to 1/2. When the phase is close to $\pi/2$, the first oscillation maximum can be observed in the energy spectrum of the detected neutrinos. The squared mass difference is

a constant of nature, while distance and energy depend on the experimental setup. Different values of L/E will therefore probe different mass splitting scales.

It is evident that in order for oscillations to occur, at least one squared mass difference must be different from zero. This in turn requires that at least one of the neutrino mass states has a non-zero mass. With a theoretical justification for all neutrinos to be massless abandoned, it seems likely that all have non-zero masses, but that still needs to be confirmed experimentally.

Since neutrinos have left-handed helicity and antineutrinos are right-handed, the combination of charge conjugation (C) and parity inversion (P), denoted as CP, transforms between $|\nu\rangle$ and $|\bar{\nu}\rangle$ states. Therefore, CP symmetry is conserved if $P(\nu_\alpha \to \nu_\beta) = P(\bar{\nu}_\alpha \to \bar{\nu}_\beta)$ and that is true if the last term of Eq. 1.7 is zero. Since the term contains only imaginary parts of the mixing matrix, neutrino oscillations conserve CP symmetry if the PMNS matrix is real. It is worth noting that only this last term is sensitive to the sign of Δm^2_{kj}; without CP violation the mass state ordering cannot be determined from vacuum oscillations.[3]

The time reversal transformation (T) converts $P(\nu_\alpha \to \nu_\beta)$ to $P(\nu_\beta \to \nu_\alpha)$ and the combined CPT transforms $P(\nu_\alpha \to \nu_\beta)$ into $P(\bar{\nu}_\beta \to \bar{\nu}_\alpha)$. In particular, the CPT symmetry requires that $P(\nu_\alpha \to \nu_\alpha) = P(\bar{\nu}_\alpha \to \bar{\nu}_\alpha)$.

For N generations of neutrinos, the mixing matrix can be parametrised by $\frac{N(N-1)}{2}$ real mixing angles and $\frac{N(N+1)}{2}$ complex phases. However, $2N-1$ phases can be eliminated from the oscillation amplitude by a global $U(1)$ rephasing and are therefore neglected [16]. If neutrinos are Majorana particles, i.e. their own antiparticles with only left-handed neutrinos and right-handed antineutrinos existing (see Sec. 1.5.4), additional phases are introduced, but the same argument applies and they do not appear in oscillation probabilities.

It is instructive to look at the phenomenology of oscillations for just two flavours, e.g. ν_e and ν_μ. In this case the 2 by 2 mixing matrix is parametrised by just one

[3] As explained in Sec. 1.3.2, matter effects can introduce effective CP violation even if the mixing matrix is real.

mixing angle θ (with no complex phases left) and is always real

$$U = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}. \tag{1.8}$$

The oscillation probability depends also on the difference between the two squared masses Δm^2. For example for the transition from ν_e to ν_μ

$$P(\nu_\mathrm{e} \to \nu_\mu) = \sin^2 2\theta \sin^2\left(\frac{\Delta m^2 L}{4E}\right). \tag{1.9}$$

It is clear from the equation that the transition probability is periodic with regards to $\frac{\Delta m^2 L}{4E}$. At the oscillation maximum, i.e. when the phase $\frac{\Delta m^2 L}{4E} = \pi/2$, the transition probability is equal to $\sin^2 2\theta$. If $\theta = \pi/4$, then $\sin^2 2\theta = 1$ and all ν_e transform to ν_μ; this scenario is called maximal mixing.

Current experimental results strongly support the three neutrino framework with three flavours and mass states. In this case, there are two independent squared mass differences and the mixing matrix is parametrised by three mixing angles θ_{12}, θ_{13}, and θ_{23}, and one complex phase δ

$$U = \begin{pmatrix} c_{13}c_{12} & c_{13}s_{12} & s_{13}e^{-i\delta} \\ -c_{23}s_{12} - s_{12}c_{12}s_{23}e^{+i\delta} & c_{23}c_{12} - s_{13}s_{12}s_{23}e^{+i\delta} & c_{13}s_{23} \\ s_{23}s_{12} - s_{13}c_{12}c_{23}e^{+i\delta} & -s_{23}c_{12} - s_{13}s_{12}c_{23}e^{+i\delta} & c_{13}c_{23} \end{pmatrix}, \tag{1.10}$$

where $c_{ij} = \cos\theta_{ij}$ and $s_{ij} = \sin\theta_{ij}$.

As mentioned earlier, CP symmetry is violated in oscillations if the mixing matrix is not real. Since δ is the only complex phase in the matrix, this occurs if and only if its value is different from 0 or π. Because of this, it is often called the CP violating phase and denoted as δ_CP.

The oscillation probability formulas with three flavours are quite complicated, although simplifications can be made for specific transitions in some experimental scenarios. For example, the survival probability of a muon neutrino in a long baseline

accelerator experiment can be approximated closely by the following formula [17]

$$P(\nu_\mu \to \nu_\mu) = 1 - 4\cos^2\theta_{13}\sin^2\theta_{23}\left[1 - \cos^2\theta_{13}\sin^2\theta_{23}\right]\sin^2\left(\frac{\Delta m^2_{32}L}{4E}\right), \quad (1.11)$$

which is dominated by the values of $\sin^2\theta_{23}$ and Δm^2_{32}, since $\cos^2\theta_{13}$ is close to one.

On the other hand, when studying electron neutrino appearance, the vacuum oscillation formulas are not enough. In this case, the probability is driven by terms with the small θ_{13}. However, matter effect are similar in amplitude and must be taken into account as well.

1.3.2 Oscillations in matter

The above formalism assumes that neutrinos are travelling through vacuum, with no particles to interact with on the way. However, when propagating through matter, there is an additional potential from coherent elastic scattering on nucleons and electrons. All flavours of neutrinos can scatter coherently on nucleons via NC interactions, where no charged leptons are created and the exiting neutrino has the same state as the entering one. The NC interaction potential affects the three flavour states in the same way, and does not affect the oscillation probability. In fact, it can be eliminated from the transition amplitude by a global phase change [16].

Electron neutrinos, on the other hand, can scatter coherently on electrons by CC interactions (Fig. 1.5). This introduces a modification of the oscillation probability depending on the baseline and electron density, known as the MSW effect, named after S. Mikheyev, A. Smirnov and L. Wolfenstein.

The formalism for two flavours is significantly simpler than for three flavours, while still illustrating the main points. The following derivation is based on Ref. [18].

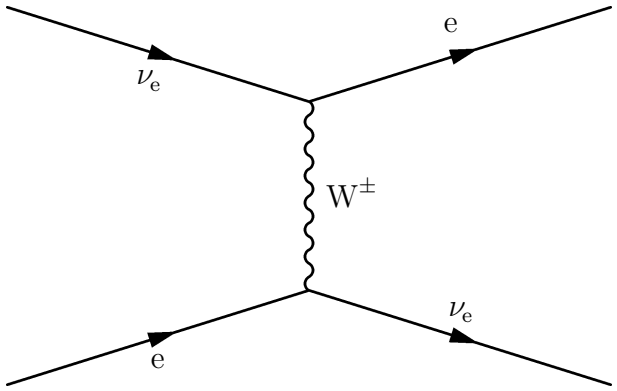

Figure 1.5: Feynman diagram of a charge current elastic scattering of a ν_e on an electron. In a coherent interaction, the spins of the incoming and outgoing electron have to be the same.

Rewriting Eq. 1.3 for an arbitrary state in the mass basis gives

$$i\frac{d}{dt}\begin{pmatrix}\nu_1\\ \nu_2\end{pmatrix} = \begin{pmatrix}E_1 & 0\\ 0 & E_2\end{pmatrix}\begin{pmatrix}\nu_1\\ \nu_2\end{pmatrix} \simeq \begin{pmatrix}E+\frac{m_1^2}{2E} & 0\\ 0 & E+\frac{m_2^2}{2E}\end{pmatrix}\begin{pmatrix}\nu_1\\ \nu_2\end{pmatrix} = \\ = \left[\frac{1}{2E}\begin{pmatrix}0 & 0\\ 0 & \Delta m^2\end{pmatrix} + \begin{pmatrix}E+\frac{m_1^2}{2E} & 0\\ 0 & E+\frac{m_1^2}{2E}\end{pmatrix}\right]\begin{pmatrix}\nu_1\\ \nu_2\end{pmatrix}, \tag{1.12}$$

where $\Delta m^2 \equiv \Delta m^2_{21} \equiv m_2^2 - m_1^2$.

Any multiple of unity can be removed by a global rephasing that does not affect the oscillation amplitude, leaving

$$i\frac{d}{dt}\begin{pmatrix}\nu_1\\ \nu_2\end{pmatrix} = \frac{\Delta m^2}{2E}\begin{pmatrix}0 & 0\\ 0 & 1\end{pmatrix}\begin{pmatrix}\nu_1\\ \nu_2\end{pmatrix}. \tag{1.13}$$

This can be transformed to flavour basis

$$i\frac{d}{dt}U^{\dagger}\begin{pmatrix}\nu_e\\ \nu_\mu\end{pmatrix} = \frac{\Delta m^2}{2E}\begin{pmatrix}0 & 0\\ 0 & 1\end{pmatrix}U^{\dagger}\begin{pmatrix}\nu_e\\ \nu_\mu\end{pmatrix}. \tag{1.14}$$

Multiplying from left side by U (noting that $UU^\dagger = 1$) returns the Schrödinger equation

$$i\frac{d}{dt}\begin{pmatrix}\nu_\mathrm{e}\\ \nu_\mu\end{pmatrix} = H_V\begin{pmatrix}\nu_\mathrm{e}\\ \nu_\mu\end{pmatrix}, \tag{1.15}$$

where the vacuum Hamiltionian equals (after removing multiples of the identity matrix)

$$H_V = \frac{\Delta m^2}{2E}U\begin{pmatrix}0 & 0\\ 0 & 1\end{pmatrix}U^\dagger\begin{pmatrix}\nu_\mathrm{e}\\ \nu_\mu\end{pmatrix} = \frac{\Delta m^2}{4E}\begin{pmatrix}-\cos 2\theta & \sin 2\theta\\ \sin 2\theta & \cos 2\theta\end{pmatrix}. \tag{1.16}$$

Interactions of electron neutrinos (and antineutrinos) with electrons in matter cause an additional potential $V_\mathrm{e} = \pm\sqrt{2}G_F N_\mathrm{e}$ (with the negative sign for antineutrinos). Here, G_F is the Fermi coupling constant and N_e is the electron density in matter. The combined (matter) Hamiltionian has the form

$$H_M = \frac{\Delta m^2}{4E}\begin{pmatrix}-\cos 2\theta & \sin 2\theta\\ \sin 2\theta & \cos 2\theta\end{pmatrix} + \begin{pmatrix}V_\mathrm{e} & 0\\ 0 & 0\end{pmatrix}. \tag{1.17}$$

Removing $V_\mathrm{e}/2$ multiplied by identity, returns

$$H_M = \frac{\Delta m^2}{4E}\begin{pmatrix}-\cos 2\theta + A & \sin 2\theta\\ \sin 2\theta & \cos 2\theta - A\end{pmatrix}, \tag{1.18}$$

where

$$A = \pm\frac{2\sqrt{2}G_F N_\mathrm{e} E}{\Delta m^2}. \tag{1.19}$$

It can be expressed as

$$H_M = \frac{\Delta m_m^2}{4E}\begin{pmatrix}-\cos 2\theta_m & \sin 2\theta_m\\ \sin 2\theta_m & \cos 2\theta_m\end{pmatrix}, \tag{1.20}$$

which has the same form as the vacuum Hamiltonian, except with new parameters θ_m and Δm^2_m, modified to make the equality hold. Substituting the new values in Eq. 1.9 gives the oscillation probability for $\nu_\mathrm{e} \to \nu_\mu$

$$P(\nu_\mathrm{e} \to \nu_\mu) = \sin^2 2\theta_m \sin^2 \left(\frac{\Delta m^2_m L}{4E} \right). \tag{1.21}$$

The effective oscillation parameters in matter are equal to

$$\Delta m^2_m = C\Delta m^2, \tag{1.22}$$

$$\sin 2\theta_m = \frac{\sin 2\theta}{C}, \tag{1.23}$$

$$C = \sqrt{(\cos 2\theta - A) + \sin^2 2\theta}. \tag{1.24}$$

A number of conclusions can be drawn from this result. If $A = \cos 2\theta$, there is a resonance and the oscillation probability becomes maximal regardless of the value of θ. However, the Fermi coupling constant G_F is very small, so for A to reach large enough values, the electron density or the distance need to be large as well. In practice, for the effect to be observable as neutrinos traverse the Earth, the distance has to be on the order of a thousand kilometres.

The sign of A is different for neutrinos and antineutrinos, so matter effects introduce an effective CP violation. It also means that in certain regions of the parameter space, matter effects may diminish or enhance real CP violation driven by the complex phases in the PMNS matrix.

Furthermore, the sign of A depends on the sign of the squared mass difference. In particular, the oscillation probability will be enhanced if A is positive. Therefore, matter effects give sensitivity to the ordering of neutrino mass states (mass hierarchy).

For three flavours, the formalism becomes more complicated, but it is based on the same idea of diagonalising the matter Hamiltonian.

$$H_M = \frac{1}{2E} U \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta m^2_{21} & 0 \\ 0 & 0 & \Delta m^2_{31} \end{pmatrix} U^\dagger + \begin{pmatrix} \pm\sqrt{2} G_F N_\mathrm{e} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{1.25}$$

This work focuses on long-baseline accelerator experiments, where a beam of neutrinos with energies on the order of GeV traverse trough a section of Earth approximately a thousand kilometres long. The transition of interest is from ν_μ to ν_e, and the effective approximated probability, including matter effects, is [19]

$$\begin{aligned} P(\nu_\mu \to \nu_\mathrm{e}) \simeq & \sin^2\theta_{23}\sin^2 2\theta_{13}\frac{\sin^2\Delta(1-A)}{(1-A)^2} \\ & + \alpha\tilde{J}\cos(\Delta \pm \delta_\mathrm{CP})\frac{\sin\Delta A}{A}\frac{\sin\Delta(1-A)}{(1-A)} \\ & + \alpha^2\cos^2\theta_{23}\sin^2 2\theta_{12}\frac{\sin^2\Delta A}{A^2}, \end{aligned} \tag{1.26}$$

where

$$\alpha = \frac{\Delta m^2_{21}}{\Delta m^2_{32}}, \tag{1.27}$$

$$\Delta = \frac{\Delta m^2_{31} L}{4E}, \tag{1.28}$$

$$\tilde{J} = \cos\theta_{13}\sin 2\theta_{13}\sin 2\theta_{12}\sin 2\theta_{23}, \tag{1.29}$$

$$A = \pm\frac{2\sqrt{2}G_F N_\mathrm{e} E}{\Delta m^2_{31}}, \tag{1.30}$$

with the usual convention where the $\pm$ symbols denote a positive sign for neutrinos and a negative sign for antineutrinos. Measurement of this oscillation channel (called electron appearance) gives sensitivity to CP violation via the second term with δ_CP and to the sign of Δm^2_{31} (mass hierarchy) via A.

1.3.3 Conclusions

To summarise, if neutrinos have non-zero masses and their mass eigenstates are not identical to the weak interaction eigenstates, they will exhibit oscillations. A neutrino created in the interaction with a charged lepton may be detected via an interaction with a lepton of a different flavour, proving that a transition occurred. The probability of the flavour transition depends on the difference between squared masses of the mass eigenstates, the neutrino energy and propagation distance, and on the mixing matrix describing the relation of the weak eigenstates to mass eigenstates.

In addition, if neutrinos travel through matter, the oscillation probability is modified by the potential from electrons on the electron neutrinos.

1.4 Status of oscillation experiments

The possibility of oscillations (between neutrinos and antineutrinos) was first hypothesised by B. Pontecorvo in the 1950s [20]. By 2001, the puzzles of solar neutrino deficit and the atmospheric neutrino anomaly were solved and the oscillation mechanism was confirmed as their cause [21, 22]. Since then, neutrino oscillations have been observed by numerous experiments. Almost all of the current data points to the three generation model; some anomalies are mentioned in Sec. 1.5.5.

Oscillation experiments are typically divided by the set of mixing parameters being probed. These were historically associated with different neutrino sources, giving names to the three sectors: the solar, the atmospheric and the reactor sector.

The PMNS matrix can be presented as a product of three independent rotations, each parametrised by one of the mixing angles

$$U = \begin{pmatrix} 1 & 0 & 0\,0 & c_{23} & c_{13} \\ 0 & -s_{23} & c_{13} & & \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{+i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1.31)$$

This representation helps to divide the parameter space into the three sectors. In addition, the squared mass differences have significantly different scales. The value of Δm^2_{21}, sometimes denoted as Δm^2_{sol} or δm^2, is almost two orders of magnitude smaller than $\Delta m^2_{32} \simeq \Delta m^2_{31}$, often denoted as Δm^2_{atm} or Δm^2. Because of that, the two mass splittings manifest at entirely different combinations of baseline and energy.

1.4.1 The solar neutrino sector

The rightmost sub-matrix of Eq. 1.31 is part of the *solar sector*. Experiments detecting neutrinos that had been created in nuclear reactions in the Sun are sensitive to the mixing angle θ_{12} and the mass difference Δm^2_{21}. The history of this channel started

as the *solar anomaly*, since the first observations noted a deficit of electron neutrinos compared to predictions made by the Standard Solar Model (SSM) [23].

This was first observed by the Homestake experiment [24]. The experiment consisted of a large horizontal tank with 400 000 litres of perchloroethylene (a dry-cleaning fluid), containing 520 t of chlorine, placed in the Homestake Mine, 1.5 km underground [24]. Neutrinos interacting with chlorine created radioactive argon atoms (with a half-life of 35 days). Because argon is a noble gas and does not bind to the perchloroethylene, it could be extracted by purging the liquid with gaseous helium and then extracted from the helium with a cooled carbon trap. Later, the trap was heated and the radioactive argon atoms counted with a small proportional counter, giving a handle on the solar neutrino flux [24].

Other radiochemical experiments using gallium also observed the deficit [25,26], as did the water Cherenkov experiment Kamiokande [27]. In addition, fluxes measured in different experiments were not consistent, depending on the probed energy range and whether the detection mechanism was sensitive only to ν_e or to other flavours as well. One considered solution was an error in the SSM, for which there was no substantial evidence; another was the possibility of neutrino flavour transitions.

The solar neutrino deficit was conclusively proven to be an effect of neutrino oscillations by the Sudbury Neutrino Observatory [28]. SNO was a heavy water Cherenkov experiment, located in a mine near Sudbury, Ontario, Canada, at a depth of 2 km. The detector contained 1 kt of D_2O in a transparent spherical tank, surrounded by an outside tank with regular water and 9.5 thousand 8 inch PMTs detecting Cherenkov light from neutrino interactions. The experiment had performed three separate measurements of neutrino interaction rates: NC on deuteron, CC on deuteron and elastic scattering on electrons. Since each of those rates has a different relation between the flux of ν_e and other neutrinos, SNO was able to prove that electron neutrinos are changing flavour, while the total flux remains constant and in agreement with the SSM [28].

The oscillations of electron neutrinos travelling from the Sun to the Earth are dominated by the MSW effect. Because of the very high and smoothly changing electron density in the core of the Sun, there is a resonance point, where matter effects significantly increase the transition probability from ν_e to ν_μ (see Eq. 1.24 and subsequent discussion). Even after confirming that the oscillation mechanism

causes a deficit in the electron neutrino flux, the role of matter effects was uncertain due to a parameter degeneracy. It was considered possible that the values of Δm^2_{21} and θ_{12} were *just right* so that vacuum oscillations could create the observed deficit and its energy dependence.

In the end, only the MSW solution called LMA (large mixing angle) could fit all the available data. The final values of solar sector parameters have been confirmed by the KamLAND experiment, which observed the solar sector oscillations in a terrestrial environment [29]. The Kamioka Liquid scintillator Anti-Neutrino Detector is located in the Kamioka mine in Japan. It is a 13 m diameter transparent nylon balloon filled with 1 kt of liquid scintillator; contained in a steel vessel filled with mineral oil. Inside, 1.9 thousand PMTs are looking for scintillation light created in the inverse beta decay reaction caused by antineutrinos. It measured rates of neutrinos produced in nuclear reactor in Japan (with energies on the order of MeV), over a long baseline (with the mean distance of 180 km), probing the oscillation parameter values in the absence of the MSW effect.

The latest results on the solar neutrino flux, constraining the solar sector parameters, come from a number of current and historic experiments observing neutrinos from the Sun: the radiochemical Homestake [30] (using chlorine), and Gallex, GNO [31] and SAGE [32] (using gallium); water Cherenkov Super-Kamiokande (successor of Kamiokande) [33–36]; heavy water Cherenkov SNO [37]; and the liquid scintillator experiment Borexino [38–40]. The solar sector is also probed by KamLAND [41].

1.4.2 The atmospheric neutrino sector

The leftmost sub-matrix in Eq. 1.31 is part of the *atmospheric sector*. The mixing angle θ_{23} and the mass difference $\Delta m^2_{32} (\simeq \Delta m^2_{31})$ dominate oscillations of neutrinos produced in cosmic ray interactions in the upper atmosphere. These were first studied as background in proton decay experiments Kamiokande and IMB [42, 43]. A deficit in muon neutrino flux, but none for electron neutrinos, was observed; and the puzzle was named the atmospheric neutrino anomaly.

Atmospheric neutrinos reaching a terrestrial detector travel different distances depending on their zenith angle. Neutrinos from directly above have only tens of kilometres from their origin to the detector, while those coming from directly below

must have been produced on the other side of the Earth, six thousand kilometres away.

This fact allowed Super-Kamiokande to prove the role of oscillations in the atmospheric anomaly by measuring the direction of incoming neutrinos. It is a Water Cherenkov detector, located 1 km underground, in the Kamioka mine. The 40 m diameter and 40 m tall cylinder contains 50 kt of pure water, observed by 11 thousand 20 inch photomultipliers, detecting Cherenkov light from neutrino interactions. The experiment measured distributions of ν_μ and ν_e fluxes as functions of the zenith angle, showing a clear deficit of up-going muon neutrinos compared to down-going ones, and no such asymmetry for electron neutrinos [44]. The ν_μ/ν_e flux ratio was consistent with the hypothesis of muon neutrinos oscillating over a distance comparable to the diameter of the Earth. No increase in the ν_e flux meant that ν_μ transition mostly into ν_τ. In fact, the mixing seemed to be maximal, indicating a value of θ_{23} close to $\pi/4$. On the other hand, the value of θ_{13}, which would drive ν_μ to ν_e oscillations, was confirmed to be small.

The results of atmospheric studies were validated by long baseline accelerator experiments. By observing ν_μ disappearance in a high intensity beam over a distance of hundreds of kilometres, the atmospheric sector parameters can be probed in a controlled environment. K2K (KEK to Kamioka) ran in Japan with a beam from the KEK facility to the Super-Kamiokande detector and was the first accelerator experiment to observe neutrino oscillations[4] [46].

The main measurement channel in this sector is the disappearance of muon neutrinos (which almost all transform into ν_τ), as opposed to an appearance measurement of the tau neutrinos. This is because most neutrinos, especially in accelerator experiments, do not have the energy necessary to create the heavy tau lepton in CC interactions. In addition, the short lifetime of τ before it decays to lighter leptons or hadrons makes for a challenging detection. However, the OPERA experiment, using a high energy neutrino beam produced at CERN, managed to detect five ν_τ event candidates over several years of running (with expected background of 0.25 events), consistent with predictions [47].

[4]The LSND experiment has reported a significant observation of $\overline{\nu}_\mu \to \overline{\nu}_\mathrm{e}$ oscillations earlier, but this result has still not been validated [45].

The strongest constraints on the atmospheric parameters come from accelerator disappearance experiments. The recently finished MINOS/MINOS+ has achieved the world's best Δm^2_{32} measurement, using the NuMI beam produced at Fermilab in the US and a far detector in Northern Minnesota, for both neutrinos and antineutrinos [48]. The most precise θ_{23} measurements come from the NOvA experiment, located at the same beam, and T2K, running in Japan with a beam from J-PARC to Super-Kamiokande [49, 50]. Atmospheric neutrino observations by the still running Super-Kamiokande provide additional constraints [51].

1.4.3 The reactor neutrino sector

Electron neutrino appearance in long baseline experiments is driven by Δm^2_{31} and the parameters of the middle sub-matrix in Eq. 1.31, but because θ_{13} is small, it is a subdominant channel. Short baseline reactor experiments are better suited to measure it via $\overline{\nu}_\mathrm{e}$ disappearance, giving the name to the *reactor sector*.

These experiments measure the flux of electron antineutrinos produced in the cores of nuclear reactors, at baselines of hundreds of metres to few kilometres. The antineutrinos are detected via the inverse beta decay process, which allows for a very robust background discrimination due to the coincidence of the prompt electron signal and the delayed one from neutron capture. Comparison of the interaction rates between far and near detectors, corrected for distance effect, returns the $\overline{\nu}_\mathrm{e} \to \overline{\nu}_\mu$ transition probability, which is approximately

$$P(\overline{\nu}_\mathrm{e} \to \overline{\nu}_\mu) \simeq 1 - \sin^2 2\theta_{13} \frac{\Delta m^2_{31} L}{4E}, \tag{1.32}$$

giving a direct handle on $\sin^2 2\theta_{13}$.

Because the value of θ_{13} is very small, it was measured much later than the other two mixing angles. In 2012, the Daya Bay reactor experiment located in China reported observation of $\overline{\nu}_\mathrm{e}$ disappearance, confirming that θ_{13} is greater than zero with 5σ significance [52]. The Korean RENO experiment published a consistent result shortly afterwards [53].

As mentioned earlier, measurement of θ_{13} is also possible by observing electron neutrino appearance in accelerator experiments. In contrast to reactor experiments,

this transition probability depends on many other oscillation parameters, including δ_{CP} and the mass hierarchy (Eq. 1.26). Appearance experiments also have to fight large background from neutral current events, in contrast to the powerful delayed coincidence of the inverse beta decay channel. Because the oscillation probability is small, disappearance experiments also have naturally bigger event statistics and therefore smaller uncertainties. T2K [54] and MINOS [55] have both published results showing a non-zero value of the θ_{13} mixing angle in 2013.

Current best results on the value of θ_{13} come from reactor experiments, including Daya Bay, RENO and Double CHOOZ [56–58]. Appearance results from accelerator experiments – MINOS, NOvA and T2K – place weak constraints on the value of δ_{CP} and sign of Δm^2_{32}. In fact long baseline experiments are the main way to search for CP violation in oscillations and a solution to the mass hierarchy problem via matter effects in the Earth.

1.4.4 Global fit to oscillation parameters

Today, the phenomenon of neutrino oscillations is thoroughly confirmed by multiple experimental results, and most of the parameters are measured with increasing precision. The data strongly supports the model of three neutrino species, by observing two independent mass splittings, driving neutrino oscillations at different scales of L/E.

Global fits to available experimental data provide best estimates of central values and confidence intervals for the mixing parameters. NuFIT is a global fit project, described in Ref. [59], with latest results published on the project webpage [60]. The best fit values and uncertainties obtained in this fit are presented in Fig. 1.6. Figure 1.7 shows 1-dimensional χ^2 projections on the values of $\sin^2$ of each of the three mixing angles, two independent squared mass differences and δ_{CP}. These results are used as the reference true values of the oscillation parameters in the CHIPS sensitivity study with GLoBES in Chapter 3.

NuFIT 3.0 (2016)

	Normal Ordering (best fit)		Inverted Ordering ($\Delta\chi^2 = 0.83$)		Any Ordering
	bfp $\pm 1\sigma$	3σ range	bfp $\pm 1\sigma$	3σ range	3σ range
$\sin^2\theta_{12}$	$0.306^{+0.012}_{-0.012}$	$0.271 \to 0.345$	$0.306^{+0.012}_{-0.012}$	$0.271 \to 0.345$	$0.271 \to 0.345$
$\theta_{12}/^\circ$	$33.56^{+0.77}_{-0.75}$	$31.38 \to 35.99$	$33.56^{+0.77}_{-0.75}$	$31.38 \to 35.99$	$31.38 \to 35.99$
$\sin^2\theta_{23}$	$0.441^{+0.027}_{-0.021}$	$0.385 \to 0.635$	$0.587^{+0.020}_{-0.024}$	$0.393 \to 0.640$	$0.385 \to 0.638$
$\theta_{23}/^\circ$	$41.6^{+1.5}_{-1.2}$	$38.4 \to 52.8$	$50.0^{+1.1}_{-1.4}$	$38.8 \to 53.1$	$38.4 \to 53.0$
$\sin^2\theta_{13}$	$0.02166^{+0.00075}_{-0.00075}$	$0.01934 \to 0.02392$	$0.02179^{+0.00076}_{-0.00076}$	$0.01953 \to 0.02408$	$0.01934 \to 0.02397$
$\theta_{13}/^\circ$	$8.46^{+0.15}_{-0.15}$	$7.99 \to 8.90$	$8.49^{+0.15}_{-0.15}$	$8.03 \to 8.93$	$7.99 \to 8.91$
$\delta_{\rm CP}/^\circ$	261^{+51}_{-59}	$0 \to 360$	277^{+40}_{-46}	$145 \to 391$	$0 \to 360$
$\frac{\Delta m^2_{21}}{10^{-5}\ \text{eV}^2}$	$7.50^{+0.19}_{-0.17}$	$7.03 \to 8.09$	$7.50^{+0.19}_{-0.17}$	$7.03 \to 8.09$	$7.03 \to 8.09$
$\frac{\Delta m^2_{3\ell}}{10^{-3}\ \text{eV}^2}$	$+2.524^{+0.039}_{-0.040}$	$+2.407 \to +2.643$	$-2.514^{+0.038}_{-0.041}$	$-2.635 \to -2.399$	$\begin{bmatrix}+2.407 \to +2.643\\ -2.629 \to -2.405\end{bmatrix}$

Figure 1.6: Results from the NuFIT v3.0 global fit to neutrino oscillation data [59], showing best fit points with 1σ uncertainties and $\pm 3\sigma$ ranges. The first column is for normal hierarchy (m_3 being the heaviest state), the second for inverted hierarchy (m_3 being the lightest), and the third shows the global minima without assuming either scenario. Here $\Delta m^2_{3\ell} = \Delta m^2_{31} > 0$ for the normal hierarchy and $\Delta m^2_{3\ell} = \Delta m^2_{32} < 0$ for the inverted hierarchy.

1.5 Open questions

Although the Standard Model is very successful and in excellent agreement with almost all experimental data, it has a number of weaknesses. It does not explain several cosmological puzzles, such as the existence of dark matter and dark energy, or the matter-antimatter asymmetry. There are structures or coincidences, which seem unlikely to be accidental, yet are not required by any fundamental laws or symmetries. For example, there are three generations of both leptons and quarks, even though almost all matter consists of only the first one. All of the particle masses are free parameters, but they are very hierarchical.

Neutrinos are part of that structure, but they do not quite fit. Their masses are extremely small, orders of magnitude below the mass of electron. The PMNS mixing matrix is analogous to the CKM matrix, but while the latter is almost diagonal, the elements of the former are all of comparable size. Since more complete theories aim to predict SM parameters from more fundamental laws, studying neutrinos can provide invaluable clues for discovering the underlying structure. At the same time,

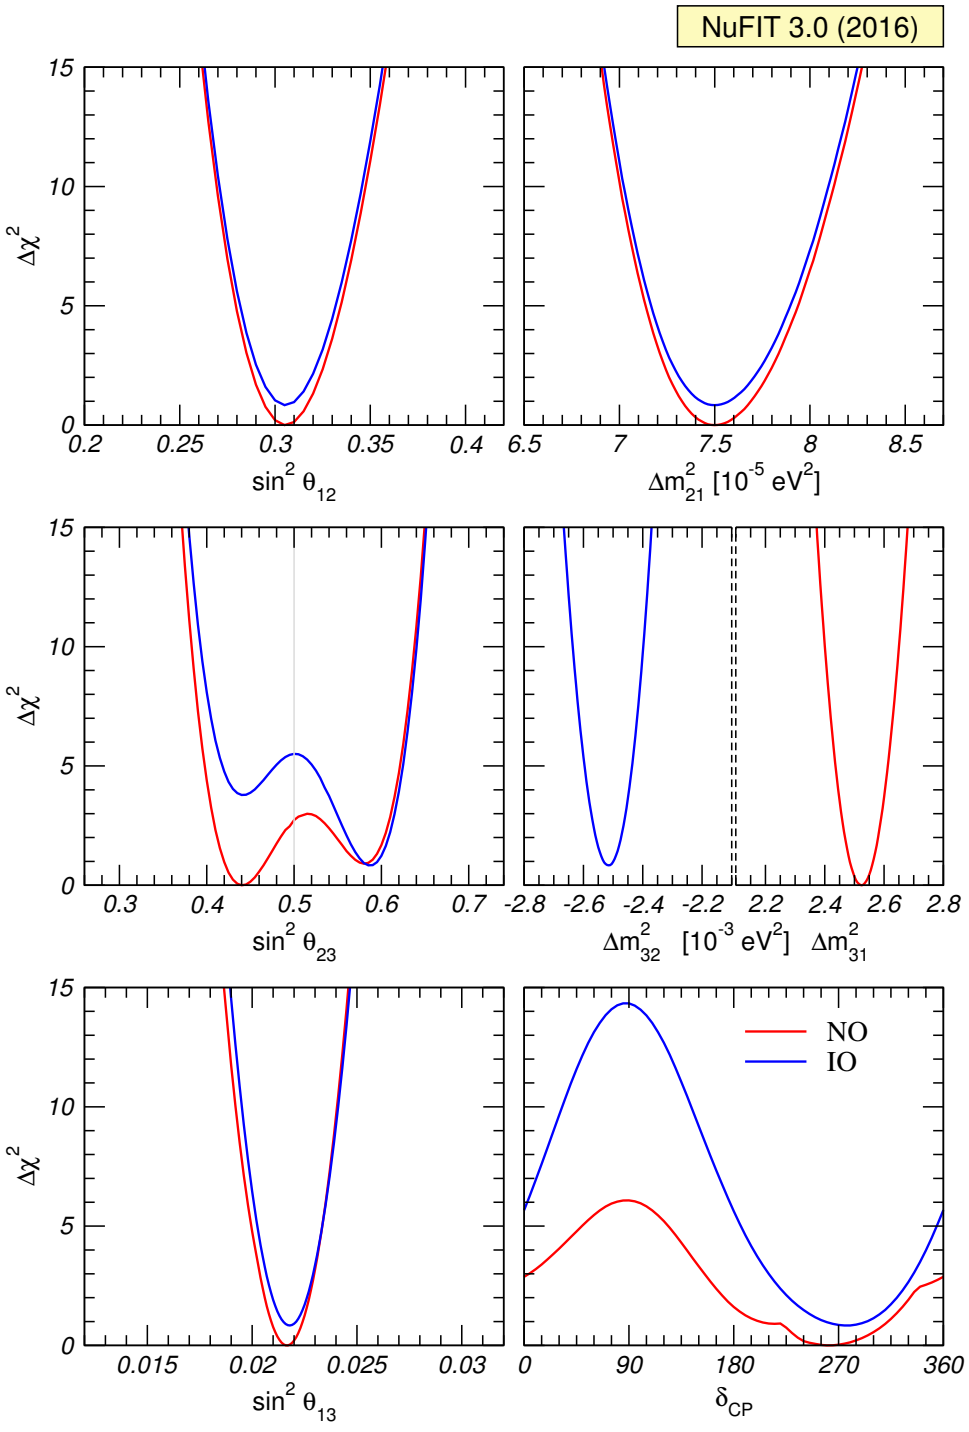

Figure 1.7: Projections of $\Delta\chi^2$ on the six parameters of the three flavour oscillation framework in the NuFIT v3.0 global fit [59]. The red curves show $\Delta\chi^2$ values when normal hierarchy is assumed, and the blue ones are for inverted hierarchy. For the normal hierarchy scenario only Δm^2_{31} projection is plotted and for inverted hierarchy it is Δm^2_{32}.

precise determination of the mixing parameters will make it possible to discriminate between various models.

Despite the great success of the three-flavour neutrino mixing model and the improving precision of experiments, there are still many unknowns. The three most important questions directly related to neutrino oscillations are the following:

1. Do neutrino oscillations violate CP symmetry – i.e. what is the value of δ_{CP}?
2. What is the sign of Δm^2_{32} – is the mass hierarchy normal or inverted?
3. What is the exact value of θ_{23} – if it is not maximal, which octant is it in?

The next great puzzle is the origin of neutrino mass. Before answering why are they so small, it must be determined whether neutrinos are Dirac or Majorana particles and what are the actual values of the mass states. Another important problem is the possibility that the three flavour framework is insufficient, with some current experimental results hinting at the existence of additional sterile neutrinos. The following sections explore all those questions in order.

1.5.1 CP violation

Violation of the CP symmetry has been observed experimentally in the quark sector, first in neutral kaon decays [61], but so far there is no evidence for it in lepton interactions. As there are no theoretical requirements for the PMNS matrix to be real, CP violation in neutrino oscillations is not unexpected.

In the three flavour oscillation framework, CP is violated if the complex phase in the mixing matrix is different from 0 or π. Currently running experiments have only placed weak limits on possible values of δ_{CP}, but they suggest that the symmetry is in fact broken, with the best-fit value of $\delta_{\mathrm{CP}} \simeq 3\pi/2$ (or equivalently $-\pi/2$) [49, 50] (see Sec. 1.6.1). If confirmed, this would be the first evidence of CP violation in the lepton sector.

The most important implication of leptonic CP violation is its role in explaining the matter-antimatter asymmetry in the Universe. A possible explanation comes from leptogenesis models. In these models neutrino masses are generated in a see-saw mechanism, where interactions with heavier neutrino partners with masses on the

order of 10^{15} GeV explain the tiny masses of ordinary neutrinos. This requires a Majorana mass term and hence lepton number violation, allowing the very heavy neutrinos to decay to baryons. Violation of the CP symmetry is necessary to ensure that there are more baryons then antibaryons, reproducing the observed asymmetry.

In addition, δ_{CP} is an important parameter of the mixing matrix, and the only one not conclusively shown to be non-zero. Measuring it is the biggest missing step to a complete model of three-flavour neutrino oscillations.

One of the most promising ways to discover CP violation is by observing electron neutrino appearance in muon neutrino beams in accelerator experiments, where the value of δ_{CP} significantly modifies the oscillation probability. Depending on the actual value of the phase, it can enhance or diminish the appearance signal, and the effect is opposite for neutrinos and antineutrinos (Eq. 1.26). Running in both modes maximises the sensitivity and also removes a degeneracy with the mass hierarchy.

1.5.2 Mass hierarchy

Although the sign of Δm^2_{21} is known to be positive, Δm^2_{32} (or equivalently Δm^2_{31}) is measured only up to its absolute value. That means that the state ν_2 is heavier than ν_1, but ν_3 might be heavier or lighter than both of them. This is called the question of mass hierarchy or mass ordering.

The scenario in which Δm^2_{32} is positive, i.e. ν_3 is the most massive state, is known as the *normal hierarchy*. Since the mass states are defined in an order that makes the mixing matrix as diagonal as possible, the state ν_{e} consists mostly of ν_1 and ν_τ of ν_3. Normal hierarchy thus follows the mass ordering of charged leptons and hence its name. This is shown schematically in Fig. 1.8 on the left.

On the other hand, if Δm^2_{32} is negative, the ν_3 would have the lowest mass of all three states, a situation called the *inverted hierarchy* (Fig. 1.8 on the right). Such an ordering would break the hierarchy observed for other particles, but in the Standard Model the masses are all free parameters and no ordering is predicted.

However, this apparent hierarchy of quarks and leptons seems unlikely to be coincidental and many Beyond Standard Model (BSM) theories try to predict it from

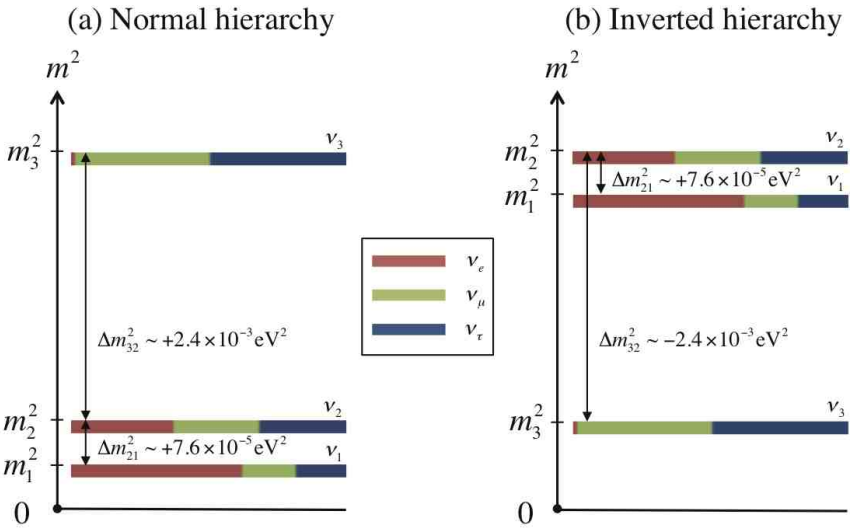

Figure 1.8: Diagram of the neutrino mass hierarchy. The normal hierarchy scenario is shown on the left and the inverted hierarchy one on the right. The colours indicate relative content of flavour states in each mass state. Image taken from Ref. [62].

fewer parameters. Determining whether the neutrino masses follow the normal or inverted hierarchy could therefore exclude approximately half of them [63].

The mass hierarchy question is about more than just the ordering of the mass states. In the case of normal hierarchy, it is possible that the lightest state has an almost vanishing mass, and all three states have masses on very different scales – ν_1 and ν_2 separated by the small mass splitting, and ν_2 and ν_3 separated by the large one. However, if ν_3 is the lightest state, the other two will be quasi-degenerate; a situation which does not occur for charged leptons [14].

It is also possible that in either of those scenarios, all three neutrino masses are quasi-degenerate, i.e. differing by very little compared to their absolute value. To answer this fully, it is important to measure the absolute masses of neutrinos. Unfortunately, this cannot be done in oscillation experiments.

Determining the mass hierarchy is also crucial for neutrinoless double beta decay searches, aiming to discover whether neutrinos are Dirac or Majorana particles. As explained in Sec. 1.5.4, in the case of inverted hierarchy, there is a lower bound on the decay rates measured in these experiments, which would allow for a complete exclusion of the hypothesis if these rates are not observed.

The sign of Δm^2_{21} is measured due to the MSW effect in oscillations of solar neutrinos. Similarly, matter effects in the Earth offer the best opportunity to determine the mass hierarchy in the atmospheric sector. Because the density of the Earth is much smaller than that of the Solar core, the distance covered by neutrinos must be large in order for the effect to become significant. Long baseline accelerator experiments offer the best opportunity, with high intensity beams and energies fine tuned to the oscillation maxima. Another experimental avenue is the study of atmospheric neutrinos, especially at very high energies, such as proposed by PINGU or ORCA [64, 65]. The advantage there is the use of the total diameter of the Earth and its dense core, but the neutrino flux is not controlled and good direction resolution is required.

1.5.3 Octant of θ_{23}

The mixing angle θ_{23} is known to be large, and first results suggested that it might be maximal, i.e. $\theta_{23} = \pi/4$. Results from MINOS and NOvA indicate that it is in fact not maximal, but it is not known whether it is smaller or larger than $\pi/4$.

The most precise results come from the disappearance channel, measuring the muon neutrino survival probability. The dominant term depends on θ_{23} only as $\sin^2 2\theta_{23}$, and there is a degeneracy between the values of θ_{23} and $\pi/2 - \theta_{23}$, meaning that the octant of the angle cannot be determined. However, in the ν_e appearance probability there is the term with $\sin^2 \theta_{23}$, which allows for the resolution of this issue.

A precise measurement of the mixing parameters is important for discriminating between models which predict the shape of the PMNS matrix.

1.5.4 Mass mechanism and absolute values

Another big question in neutrino physics relates to the absolute mass of neutrinos and its origin mechanism. The extreme smallness of neutrino masses (at least 5 orders of magnitude lower than the mass of electron) suggests the existence of another mechanism, as opposed to the Yukawa couplings just being small. Some models solve this problem by employing a see-saw mechanism, where mixing between neutrinos

with masses on the SM scale ($\sim 100\,\mathrm{GeV}$) and New Physics scale ($\sim 10^{15}\,\mathrm{GeV}$) result in the apparent neutrino mass of around $10^{-2}\,\mathrm{eV}$, consistent with current observations.

Massive neutrinos may be Dirac or Majorana particles. In the first case, they acquire mass in the same way as the other fermions, via a Yukawa coupling that mixes left- and right-handed fermion fields. This would require the existence of right-handed neutrinos (and left-handed antineutrinos), which do not interact weakly. In this scenario, neutrinos are different from antineutrinos, producing four distinct states in total, two of which are Standard Model singlets.

On the other hand, a Majorana mass term in the Langrangian is constructed from left-handed states only. It requires that the particle is identical to its own antiparticle, which is possible for neutrinos, as they do not carry any charge (which would need to flip sign under the C transformation). If neutrinos are Majorana particles, they have only two states distinguished by the helicity, left-handed neutrinos and right-handed antineutrinos. Majorana masses are necessary in leptogenesis models explaining baryon-antibaryon asymmetry, as they are an integral part of the see-saw mechanism employing heavy neutrinos and require lepton number violation, allowing lepton to baryon transitions.

Majorana nature of neutrinos could be confirmed by observing a neutrinoless double beta decay ($0\nu\beta\beta$), i.e. a simultaneous double beta decay where the two electron antineutrinos annihilate. This process breaks the lepton number conservation by two and is forbidden in the Standard Model. Numerous experiments are conducting such searches with various target nuclides, including EXO-200, KamLAND-Zen, GERDA and CUORE, but no positive signal has been observed yet [66].

The absolute values of neutrino masses are also unknown, with the most reliable measurements coming from decay at rest experiments. In particular, tritium beta decay experiments probe the effective mass of the electron neutrino, by measuring the tail of the electron spectrum. So far, only upper limits have been placed, with the strongest being $m_{\mathrm{ee}} = \sum_i U_{ei}^2 m_i < 2.05\,\mathrm{eV}$ at 95% CL reported by the Troitsk experiment [67]. The next generation experiment KATRIN will use a very large magnetic spectrometer to bring that limit down by an order of magnitude [68].

Other bounds on the values of neutrino masses come from astrophysics experiments. Many of these experiments are sensitive to the total mass of relic neutrinos, i.e.

neutrinos which had decoupled from matter in the very early universe [69]. The best current constraints come from measurements of the temperature and polarisation anisotropies of the Cosmic Microwave Background (CMB), and observations of the clustering of galaxies. A combined analysis of the results from the Planck experiment and from the Baryon Oscillation Spectroscopic Survey (BOSS) gives an upper bound on the sum of the neutrino mass states, $\Sigma m_\nu < 0.183\,\mathrm{eV}$ at the 95 % confidence level [70].

Although oscillation experiments cannot probe those questions directly, their findings can still prove helpful. The measured squared mass differences provide a lower bound on the two larger neutrino masses – assuming that the lightest state has a zero mass, the mass of the heaviest state must be at least 50 meV, satisfying the requirement that $\Delta m^2_{32} = m^2_3 - m^2_2 \simeq \Delta m^2_{31} > 2.4 \times 10^{-3}\,\mathrm{eV}^2$.

The mass hierarchy is also very important for placing constraints on results from neutrinoless double beta searches. The $0\nu\beta\beta$ decay rates are proportional to the square of the effective electron neutrino mass, m^2_{ee}. If the hierarchy is inverted, the value of m_{ee} has a lower bound of approximately 20 meV coming from the known oscillation parameters. This bound, combined with potential exclusion limits from neutrinoless double beta experiments, could lead to a conclusive rejection of the Majorana neutrino hypothesis. If, on the other hand, the hierarchy is normal, the electron neutrino may be very light and only positive discovery will be definitive.

1.5.5 Sterile neutrinos

Finally, even though most data is consistent with the three flavour model, some results show hints for the existence of sterile neutrinos, i.e. flavour eigenstates that do not interact weakly, but can mix with the other states in oscillations. In particular, findings from short baseline accelerator experiments LSND and MiniBooNE suggest the existence of a third squared mass difference on the order of $1\,\mathrm{eV}^2$ (compared to $\Delta m^2_{21} \simeq 10^{-5}\,\mathrm{eV}^2$ and $\Delta m^2_{32} \simeq 10^{-3}\,\mathrm{eV}^2$) [45, 71]. Since the data from LEP experiments allow only three light neutrinos to couple to Z, the additional state cannot interact weakly.

Also, many reactor experiments have found the antineutrino flux at very short baselines to be lower than predicted by theoretical models, which could be a signature of sterile neutrino oscillations or a shortcoming of the model [72, 73].

Several other experiments have performed searches for sterile neutrinos at various energy scales. MINOS+, Daya Bay, IceCube and other experiments have all reported null results and placed strong limits, excluding almost all of the parameter space allowed by the positive claims of LSND and MiniBooNE [74, 75].

Another promising experimental initiative is the Short Baseline Neutrino Program at Fermilab [76]. It consists of three detectors utilising a novel liquid argon time projection chamber (LAr TPC) technology, located at the Booster Neutrino Beam: MicroBooNE, ICARUS and SBND. The goals of the project are to explain the nature of the event excess observed in MiniBooNE, definitively answer the question of sterile neutrino oscillations at the 1 eV mass splitting range, while providing input to the design of DUNE by measuring neutrino-argon interaction cross sections and prototyping the LAr TPC technology.

The MicroBooNE detector is has a total mass of 170 t and is located 470 m from the beam production target. It began operation in 2015. The ICARUS T600 detector is the current biggest LAr TPC, at 760 t total mass [77]. It ran in the CNGS neutrino beam from CERN to Gran Sasso and is currently being redeployed to the Fermilab site, at a baseline of 600 m. The SBND will act as a near detector, characterising the neutrino beam before potential oscillations. It will be located 110 m away from the proton target and have a total mass of 112 t of argon. ICARUS and SBND are expected to start operation in 2018.

Confirmation of sterile neutrinos' existence would revolutionise neutrino physics, adding a whole new set of parameters to explore [14].

1.6 Long baseline oscillation experiments

Many of the important questions mentioned in the previous section can be answered by long baseline accelerator experiments. The channel of interest is the observation of electron neutrinos at a far detector, located hundreds of kilometres from the ν_μ beam origin. A near detector placed much closer (before any oscillations occur)

is used to constrain the flux, along with the intrinsic ν_e content. The appearance probability (Eq. 1.26) is sensitive to δ_{CP}, the mass hierarchy, and the octant of θ_{23}.

Matter effects in Earth are of similar magnitude as CP violation, and the effects can add up or cancel each other, depending on the mass hierarchy, δ_{CP} value and sign of neutrinos. Measuring both neutrino and antineutrino oscillation probabilities makes it possible to unfold these effects. Accelerator experiments can achieve this thanks to the design of conventional neutrino beams. As explained in Chapter 2, a proton beam hitting a target creates secondary hadrons, which in turn decay to neutrinos. These hadrons, mostly charged π and K mesons, are focused by a device called the magnetic horn. Reversing the current polarity of the horns switches the sign of the focused hadrons, creating a beam enhanced in either neutrinos or antineutrinos. Comparing results from neutrino and antineutrino runs at the same experiment makes it possible to disentangle the effects of CP violation and mass hierarchy on the oscillation probability.

1.6.1 Currently running experiments

Two experiments designed to search for CP violation and to determine the mass hierarchy, which are currently running, are NOvA in the US and T2K in Japan.

These projects share many similarities. Both experiments use conventional neutrino beams, described in more detail in Sec. 2.1.1. They also have detectors located away from the main beam axis – the *off-axis effect* creates a very narrow energy spectrum, maximising the sensitivity to electron appearance (see Sec. 2.1.2).

Both of the experiments use a two-detector setup. The *near* detector located close to the beam production point measures precisely the neutrino flux and its composition before oscillation, as well as neutrino cross sections. The *far* detector observes neutrino interactions at the end of the baseline and searches for signatures of oscillation, either a decrease in the muon neutrino flux or an increase in electron events. The use of a near detector significantly reduces systematic errors appearing in the analysis.

T2K stands for Tokai to Kamioka, as its beam is produced in the J-PARC facility near Tokai, and sent towards the Super-Kamiokande detector located in the Kamioka mine [78]. The water Cherenkov far detector is at a baseline of 295 km, 2.5° (ca.

44 mrad) off-axis. T2K employs a suite of on- and off-axis near detectors, located 280 m away from the proton target. The main near detector (named nd280) is approximately 2° off-axis and uses multiple components, including a tracker, π^0 detector and electromagnetic calorimeter, to precisely measure the properties of neutrino interactions and product particles.

NOvA (NuMI Off-axis ν_e Appearance) is a more recent experiment, utilising the NuMI beam [79, 80]. In contrast to T2K, both of its detectors are segmented liquid scintillator tracking calorimeters, built with the same technology. The far detector has a mass of 14 kt and is located on the surface (with just a small rock overburden) near Ash River in Northern Minnesota. It is 810 km away from the beam production and 14 mrad off-axis. Due to the significantly longer beamline, NOvA is much more sensitive to the matter effects than T2K. The smaller near detector is located in Fermilab, 100 m underground and 1 km away from the target. Using functionally identical detectors allows many systematic uncertainties to cancel out, when considering the ratio between the far and near detector neutrino energy spectra.

Both NOvA and T2K have presented results from appearance analyses, reporting constraints on the values of θ_{13} and δ_{CP}, the mass hierarchy and the octant of θ_{23} [49, 50]. The measured values of θ_{13} are consistent with reactor results. Both experiments slightly favour normal hierarchy, NOvA by around 0.5 units of χ^2 and T2K by about two units of χ^2. If that difference becomes more significant, it will rule out the possibility of excluding the existence of Majorana neutrinos by the neutrinoless double beta experiments.

T2K observes more electron neutrino events than the maximum expected for $\delta_{CP} = 3\pi/2$, and fewer antineutrino events than the minimum for the same value. Because of this, its limit on δ_{CP} is stronger than predicted sensitivity, excluding δ_{CP} between 0 and π at 90 % confidence level. The NOvA result also has a best fit value of $\delta_{CP} = 3\pi/2$, and excludes $\delta_{CP} = \pi/2$ in inverted hierarchy, for the lower octant of θ_{23}, at 3σ.

These results give hints that the task of discovering CP violation in the neutrino sector might be easier than expected. If the true value of δ_{CP} is close to the current best fit point, it is possible that NOvA and T2K alone will be able to confirm that with high significance. On the other hand, a conclusive determination of the correct

mass hierarchy as well as precise a precise measurement of the actual value of δ_{CP}, a new generation of experiments will be necessary.

1.6.2 Future projects

Because the current generation of long baseline experiments is not predicted to unambiguously determine the mass hierarchy and the value of δ_{CP}, there are plans for much more ambitious future projects. The two with most momentum and closest timelines are DUNE in the US and Hyper-Kamiokande in Japan.

The Hyper-Kamiokande proposal includes a new water Cherenkov detector in a mine just next to the one that houses Super-Kamiokande and KamLAND, and an upgrade to the already existing T2K beamline at J-PARC. The experiment would have the same baseline as T2K, i.e. 295 km at 2.5° off-axis. The neutrino energy will also remain the same, but the new beam will have much higher intensity. The flux of neutrinos is proportional to the power of the primary proton beam, which is planned to increase from the current 400 kW to 1.3 MW [81].

The new detector will use almost the same technology as Super-Kamiokande. It will consist of two large cylindrical tanks filled with ultra-pure water, and instrumented with photomultiplier tubes on the walls detecting Cherenkov light from the products of neutrino interactions [82]. However, the fiducial mass will be significantly bigger, 380 kt compared to 22 kt at Super-Kamiokande, increasing the event rates and sensitivity. In addition, the new detector will benefit from higher performance PMTs, the development of which is ongoing. While a near detector already exists, there are several proposals for new detectors at near or intermediate baselines to augment Hyper-Kamiokande.

If the budget proposal to the Japanese government is approved, the construction could start in 2018 and first data taking around 2025 [83]. The total cost is estimated to be around $700 million [84].

The DUNE project (Deep Underground Neutrino Experiment) calls for construction of a completely new beam and a novel liquid argon time projection chamber detector [85]. The beamline project is called the Long Baseline Neutrino Facility (LBNF), and involves a new very high intensity conventional beam from Fermilab to South Dakota. The planned beam power is 1.2 MW with the possibility of a future

upgrade to 2.4 MW. The DUNE far detector will be located 1500 m underground in the Sanford Underground Research Facility (in the Homestake Mine), 1300 km away. The far detector will be a liquid argon TPC, which offers a very precise particle track reconstruction, providing great energy resolution and particle identification capabilities. Because of that, DUNE can use an on-axis wide energy beam and benefit from the observation of two whole oscillation peaks. The near detector design has not yet been finalised, but it is likely that it will include both a liquid argon TPC and an additional multi-purpose tracker [86].

The far detector will consist of 10 kt modules, which will be by far the biggest LAr TPCs ever built; the current largest detector of this type is ICARUS T600 with a total mass of 760 t [77]. DUNE will be constructed in stages, starting with a single 10 kt module, followed shortly by a second one, and a later expansion to the full size of 40 kt. The project aims to accumulate a total exposure of 120 kt $\times$ MW $\times$ year by 2035, which corresponds to approximately ten years of running with a single module. The total project cost is unprecedentedly high for a neutrino experiment, with just the contribution from the US Department of Energy estimated at \$1.5 billion [87].

Both projects have wide and extensive physics programs, including unambiguous determination of mass hierarchy and the octant of θ_{23}, a precise measurement of δ_{CP} and exclusion or confirmation of CP violation for a large range of its possible values.

Chapter 2

The CHIPS project

As shown in the previous chapter, there are interesting and important questions that can be answered by studying oscillations, and the long baseline accelerator experiments are in an excellent position to do so. However, the currently running experiments, NOvA and T2K, will probably not be able to unambiguously determine the mass hierarchy or to measure δ_{CP} with a high precision. Moreover, the significance of CP violation exclusion or confirmation depends on the unknown value of δ_{CP}; although the first results suggest that it is in a favourable phase space location, it might also lie in a region where the sensitivity to CP violation is much smaller.

Currently planned next generation projects, such as DUNE or Hyper-Kamiokande, have much farther physics reach, but due to record-high budgets, on the order of a billion dollars, their future depends heavily on the available funding. Even if built in the expected full configuration, the process will take a long time, with first results at least a decade away. A novel approach could help solve those problems, by creating a detector which is cheaper and can deliver significant physics results in the meantime.

CHIPS (CHerenkov detectors In mine PitS) is an R&D program exploring the viability of such a strategy. The main goal is to develop a water Cherenkov detector for an accelerator beam, which would be cheaper than current technology and easily scalable but without sacrificing too much performance. The target cost is \$300 k per kt of fiducial detector mass, compared to more than \$4 million/kt in Super-Kamiokande or \$2.5 million/kt predicted for the potential water detector at LBNF [88, 89]. Although this is not an ideal metric, since the sensitivity depends on more than just the raw volume, it allows for an easy comparison between similar detectors.

These aims can be met by submerging the main structure in an existing water body as opposed to excavating an underground cavern. Other cost-saving measures include a lightweight mechanical structure, sparse PMT placement (optimised for beam events), inexpensive water purification treatment and novel readout technologies. CHIPS detectors will be modular, allowing more flexibility in construction. A single detector can be recovered and upgraded after deployment, or even disassembled and moved to a different location. If successful, this project can lead the way to future megaton-scale detectors, which might not be achievable otherwise.

A tangible goal of the R&D program is to build and deploy CHIPS-10, a 10 kt module,[1] which will serve as a technology demonstrator and development platform, but also as a fully functional detector helping the global search for CP violation and mass hierarchy. CHIPS-10 will be deployed in a flooded mine pit in northern Minnesota, in the path of the existing NuMI neutrino beam. Small scale prototypes have already been tested at the Wentworth Pit and the first data provides constraints on the cosmic muon background (see Chapters 4 and 5).

2.1 Neutrino beam and detector location

CHIPS is developed as a detector for long baseline beam experiments, where high intensity beams of muon neutrinos with ∼GeV energy are observed at a distance of 300 to 1300 km. There are two such beams currently in operation: one in Japan, produced in J-PARC near Tokai, and one in US, produced in Fermilab near Chicago. The latter, called NuMI, will be used for developing and testing CHIPS.

Both are examples of a conventional neutrino beam, or a superbeam.[2] In such a design, accelerated protons hit a carbon target, creating secondary hadrons (mostly charged pions), which, after focusing, decay in flight into neutrinos. The beam composition is dominated by muon neutrinos. Depending on the current polarity in the focusing system, a beam of neutrinos or antineutrinos can be chosen.

[1] The word "module" is used to describe a single self-contained detector. Potential CHIPS detectors with masses on the order of 100 kton would consist of multiple such modules.

[2] Superbeams operate on the same physics principles as the first accelerator neutrino beams, but have significantly higher intensity, hence the name.

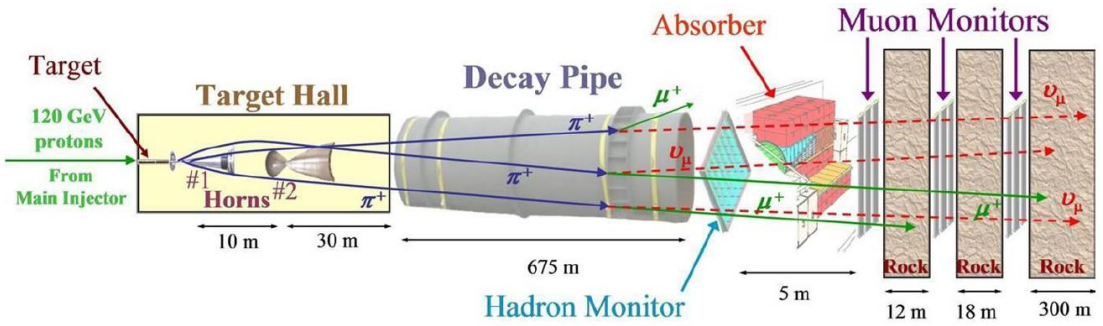

Figure 2.1: Diagram of the major components of the NuMI beamline. Horizontal distances are not to scale. In the neutrino mode, positive hadrons are focused and negative ones ejected. Image taken from Ref. [80].

The following section describes the NuMI beam in greater detail. Other conventional beams, such as the J-PARC beam, have very similar design. The main differences are in target shapes, the number of magnetic horns or the proton beam energy and bunch structure. However, the main principles remain the same.

A new beamline, called LBNF, is planned to be constructed by 2026. It will be used primarily by the DUNE experiment, but the possibility of placing a CHIPS detector in a reservoir in the path of the beam is under consideration. The LBNF will be a conventional beam as well, although with a much higher intensity.

2.1.1 The NuMI beam

NuMI (Neutrinos at the Main Injector) is a muon neutrino beam produced in the Fermilab accelerator facility in the US. It is currently the most powerful neutrino beam in the world, reaching recently a milestone power of 700 kW [90]. A schematic diagram of the beam production is shown on Fig. 2.1. The beamline was built in 2005 for the MINOS experiment. In 2012-2013 it was upgraded to provide a higher intensity and a different neutrino energy for NOvA. The typical neutrino energy at the MINOS far detector was around 3 GeV before the upgrade, and approximately 7 GeV afterwards. This was done in order to tune the energy at NOvA to 2 GeV, close to the oscillation maximum, via the off-axis effect (Sec. 2.1.2).

The primary particles are protons accelerated in the Main Injector synchrotron to the energy of 120 GeV. A spill is extracted every 1.3 s and directed at a stationary carbon target. The spill consists of several bunches and is 10 µs long. This short

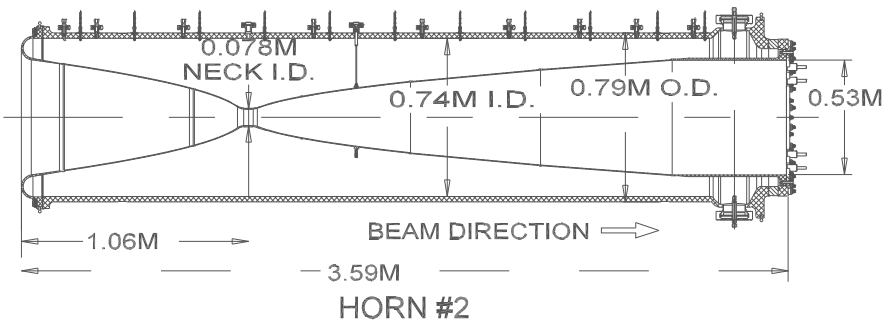

Figure 2.2: Schematic diagram of the second magnetic horn in the NuMI beamline. Image taken from Ref. [80].

duration is essential for CHIPS (and NOvA) to reject the massive cosmic ray muon background. Recording the precise time of the beam spill, corrected for the distance to the detector, allows one to discard all events occurring outside of it.

The target is a graphite rod made of forty-eight fins, each 7.4 mm wide [91]. Interactions of the primary protons with carbon nuclei in the target produce secondary hadrons; mostly charged pions and kaons with a very small mixture of other particles. The narrow profile of the target ensures that most of them exit the target without re-interaction. Because the primary beam is positively charged and the target is electrically neutral, the secondary particles have a net positive charge.

The secondary hadrons propagate downstream through the focusing system consisting of two magnetic horns. The horns are hollow aluminium structures with an inner and outer conductor. The inner conductor has a parabolic profile, with a neck in the middle (Fig. 2.2). When the electric current flows through the horn, a toroidal magnetic field is created between the conductors. Hadrons with very low momentum transverse to the beam axis will travel through the neck and not experience any magnetic field. On the other hand, charged particles with higher transverse momentum will cross the inner conductor and will be focused towards the beam axis or ejected depending on their sign. The parabolic profile ensures that the total effect of the magnetic field is proportional to the transverse momentum of the particle. The typical current flowing through the horns is 200 kA for a pulse time of 2 ms [80].

The polarity of the horn current determines the sign of the charged hadrons which are focused. Because positive pions and kaons decay to produce (predominantly)

neutrinos and negative hadrons produce antineutrinos, the horn current polarity allows for choosing the beam composition. In an on-axis beam, the antineutrino mode flux is still dominated by muon neutrinos. This is because hadrons with low transverse momentum do not pass through the horns at all, and since most of them are positively charged, these will produce neutrinos, typically of higher energy. However, the off-axis effect (Sec. 2.1.2) effectively filters out those high energy wrong-sign neutrinos and dramatically reduces the contamination.

After passing the horns, the secondary hadrons enter the decay pipe. The 675 m long pipe has a diameter of 2 m and is filled with helium, which minimises the number of secondary interactions. Here, the hadrons can decay in flight, and the dominating decay modes are

$$\begin{aligned} \pi^+ &\rightarrow \mu^+ + \nu_\mu, \\ \mathrm{K}^+ &\rightarrow \mu^+ + \nu_\mu, \end{aligned} \tag{2.1}$$

where the former accounts for approximately 95 % of the muon neutrinos. A charge conjugation flipping all signs will give the expressions for antineutrino production.

The muons in turn decay as $\mu^+ \rightarrow \mathrm{e}^+ + \nu_\mathrm{e} + \overline{\nu}_\mu$, producing intrinsic electron neutrino background. However, since the muon lifetime is approximately 100 times longer than that of the pion (and the decay pipe length is optimised for pion decays), their contribution to the flux is proportionally smaller [92]. Additional electron contamination comes from decays of charged and neutral kaons

$$\begin{aligned} \mathrm{K}^+ &\rightarrow \pi^0 + \mathrm{e}^+ + \nu_\mathrm{e}, \\ \mathrm{K}^+ &\rightarrow \mathrm{e}^+ + \nu_\mathrm{e}, \\ \mathrm{K}^0_\mathrm{L} &\rightarrow \pi^\pm + \mathrm{e}^\mp + \overset{(-)}{\nu}_\mathrm{e} \end{aligned}$$

The total electron neutrino and antineutrino content in the beam (before oscillations), as seen at the CHIPS-10 location, is on the order of 1 %, and the wrong-sign muon neutrino contamination is on the order of 5 % (see Figs. 2.14 and 2.15).

The beam absorber is located after the decay pipe. It is a large structure with an aluminium core, surrounded by steel and concrete. It stops all the remaining hadrons, including the large fraction of primary protons which did not interact. In front of the absorber, there is a hadron monitor for beam alignment diagnostics.

Downstream of the absorber, there are also three muon monitors located in alcoves excavated in the rock. Between the absorber and the near detector hall, 240 m of dolomite rock serves to absorb all the muons produced in the decay pipe.

The near detector cavern accommodates several neutrino detectors. Because the beam is directed at a downward small angle of 3.3°, to pass through the Earth and towards the far detectors, the hall is 100 m underground. This provides overburden limiting the cosmic muon background in the detectors.

The most upstream in the beam axis is the MINERvA detector. MINERvA is focused on precise measurements of neutrino interaction cross sections on a variety of target materials, at a distance not affected by standard oscillations. The detector is a finely segmented scintillator, surrounded by electromagnetic and hadron calorimeters. In the front, it contains a target region with a number of different target panels, including one with liquid water.

Further downstream is the MINOS near detector. It is a 5.4 kt magnetised tracking calorimeter, composed of alternating plates of steel and scintillator strips. The MINOS experiment and its extension named MINOS+ have both finished data taking. However, the MINOS near detector is still used by MINERvA as a muon spectrometer.

Located in a nearby cavern, 14 mrad off the beam axis, is the NOvA near detector. It has a mass of 222 t and is a fully active tracking scintillator, built of extruded PVC cells filled with liquid scintillator. These are arranged in planes of strips, with alternating horizontal and vertical orientation. The scintillation light is collected via wavelength-shifting optic fibres and read out with PMTs.

Although a near detector for CHIPS was considered and might be built for use with future modules, there is no plan for a near detector for CHIPS-10. Therefore, data from MINERvA and the NOvA ND will be a very important source of constraints on the beam flux and composition at CHIPS.

The beam travels onwards through the Earth's crust and surfaces in Northern Minnesota. Its original target, the MINOS far detector, was located in the Soudan Underground Laboratory. Now the beam serves primarily NOvA, which is located on the surface in Ash River. The region is rich with lakes and flooded pits, which makes it suitable for a CHIPS location.

2.1.2 The off-axis effect

In contrast to MINOS, the NOvA far detector is located away from the beam axis. The same is true for Super-Kamiokande in the T2K beam, and will be for CHIPS as well. This design takes advantage of the *off-axis* effect, which produces a narrow energy spectrum of the neutrinos observed in the detector, due to the kinematics of the pion decay.

The pion decay (Eq. 2.1) is a two-body process, and the neutrino energy has a well defined value in the rest frame of the pion,

$$E_\nu^* = \frac{m_\pi^2 - m_\mu^2}{2m_\pi} \simeq 29.79\,\mathrm{MeV}. \tag{2.2}$$

The asterisk marks quantities expressed in the centre of mass reference frame, while the variables without it describe the beamline rest frame.

Assuming that the neutrino is ultrarelativistic, $E_\nu \simeq p_\nu$ and the cosine of the neutrino emission angle relative to the pion direction is

$$\cos\theta^* = \frac{p_{\nu l}^*}{p_\nu^*} = \frac{\gamma(p_\nu \cos\theta - \beta E_\nu)}{p_\nu^*} \simeq \frac{\gamma E_\nu(\cos\theta - \beta)}{E_\nu^*}, \tag{2.3}$$

where $\beta = p_\pi/E_\pi$, and $\gamma = 1/\sqrt{1-\beta^2} = E_\pi/m_\pi$.

Neutrino energy in the beam reference frame is

$$E_\nu = \gamma(E_\nu^* + \beta p_\nu^* \cos\theta^*) = \gamma E_\nu^* + \gamma^2 \beta E_\nu(\cos\theta - \beta), \tag{2.4}$$

where the second step uses Eq. 2.3.

Solving Eq. 2.4 for neutrino energy (using Eq. 2.2 after the first step) gives the expression for E_ν as a function of pion energy and the emission angle

$$\begin{aligned} E_\nu = \frac{\gamma E_\nu^*}{1 - \gamma^2\beta(\cos\theta - \beta)} = \frac{E_\pi(m_\pi^2 - m_\mu^2)}{2m_\pi^2[1 - \frac{p_\pi E_\pi}{m_\pi^2}(\cos\theta - \frac{p_\pi}{m_\pi})]} = \\ = \frac{E_\pi(m_\pi^2 - m_\mu^2)}{2(m_\pi^2 + p_\pi^2 - p_\pi E_\pi \cos\theta)} = \frac{m_\pi^2 - m_\mu^2}{2(E_\pi - p_\pi\cos\theta)} \end{aligned} \tag{2.5}$$

Figure 2.3 shows the dependence of the neutrino energy on the pion energy for several off-axis angles. It can be seen that there is a maximal neutrino energy for a given emission angle (larger than zero).

This is shown in the following derivation

$$\tan\theta = \frac{E_\nu \sin\theta}{E_\nu \cos\theta} = \frac{E^*_\nu \sin\theta^*}{E_\nu \cos\theta} \simeq \frac{E^*_\nu \sin\theta^*}{E_\nu}, \tag{2.6}$$

where the last approximation holds for large neutrino and pion energies compared to the mass of pion, equivalent to small values of the emission angle. For non-zero values of θ Eq. 2.6 can be inverted, giving

$$E^*_\nu = \frac{E^*_\nu \sin\theta^*}{\tan\theta}, \tag{2.7}$$

which cannot be larger than $\frac{E^*_\nu}{\tan\theta}$. Since E^*_ν is a constant, that means that the neutrino energy observed in the laboratory frame has a maximal value, determined solely by the emission angle.

Figure 2.3 also shows that at higher off-axis angles, the neutrino energy becomes less dependent on the pion energy and takes values in a small range, particularly in the region where the pion flux peaks. Assuming that the pions are narrowly focused along the beam axis, the angle θ is equivalent to the off-axis angle of the detector relative to the beam. Indeed, neutrinos detected at a larger off-axis angle have a narrower energy spectrum (Fig. 2.4). Although the total beam intensity decreases, the intensity at the specific energy may increase. This can be used to tune the energy peak to the oscillation maximum at the far detector distance.

More importantly, the lack of a high energy tail notably reduces background from NC events. In those interactions the outgoing neutrino carries away significant momentum and the total energy visible in the detector is lower than it would for a CC interaction. A narrow energy spectrum ensures that there are fewer NC events in the analysis energy window.

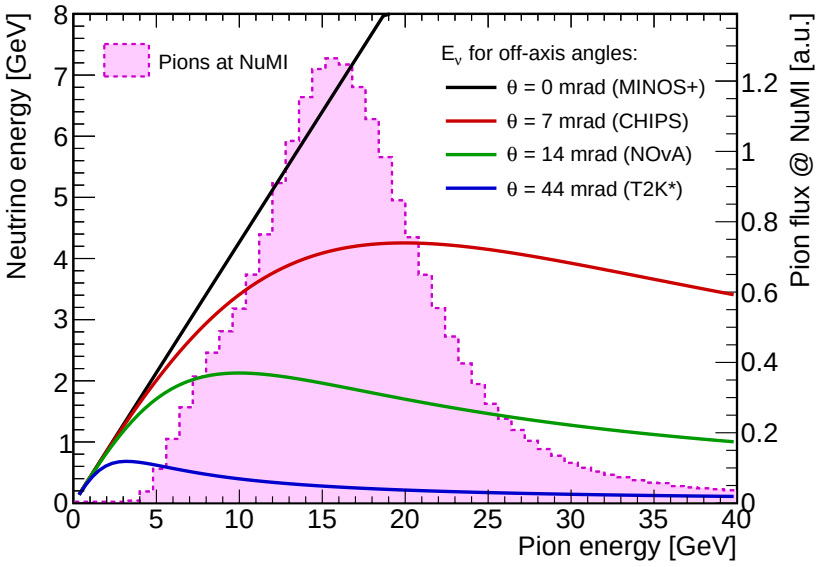

Figure 2.3: Neutrino energy plotted as a function of pion energy for different angles of neutrino emission (Eq. 2.5). The solid lines show this dependence for off-axis angles corresponding to the far detector locations in the following experiments: MINOS+ (black), CHIPS (red), NOvA (green) and T2K (blue). The pink shaded histogram shows the energy distribution of the secondary pions exiting the target in the NuMI beam, which serves all mentioned experiments except T2K.

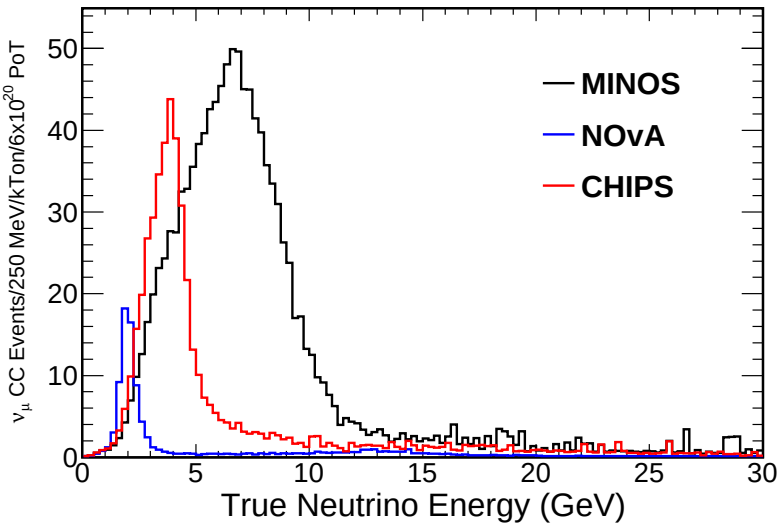

Figure 2.4: Muon neutrino flux from NuMI at different off-axis angles, in the absence of oscillations. Neutrino energy spectra are shown for MINOS (on-axis, black), NOvA (14 mrad, blue) and CHIPS (7 mrad, red), based on the flux simulation. Figure taken from Ref. [93].

2.1.3 CHIPS-10 location

During the search for the location for CHIPS-10 deployment and prototype testing, several factors had to be considered. Most importantly, the baseline and off-axis angle relative to the beam affect the neutrino flux and energy spectrum. The beam simulation (described in Sec. 2.3.2) can be used to efficiently recalculate the flux given the geographic coordinates, which made it possible to study the physics sensitivity for different locations.

The second factor is the maximum depth of the water body. Most neutrino detectors are located deep underground, to shield against cosmic ray muons. For CHIPS, which can only utilise a shallow water overburden, the muons are a source of major background. The reduction in muon flux is exponential with depth in water and early studies have indicated that an overburden of around 40 m is desired for cosmic muon rates below 50 kHz [93, 94].

Finally, there are practical aspects, such as accessibility, legal availability and environmental considerations.

The location chosen ultimately is the Wentworth Pit 2W (Fig. 2.5). It is an unused and flooded iron mine pit, approximately 0.8 km × 1.2 km in size, with 60 m maximum depth (Fig. 2.6). The centre of the pit is located 708 km away from the beam target, at 7 mrad off-axis angle. The energy spectrum of unoscillated muon neutrinos is shown on Fig. 2.4. The neutrino energy is not as finely tuned to the oscillation maximum as in the NOvA experiment, but offers complementarity to it. The higher total event rate and wider energy spectrum make CHIPS-10 more sensitive to δ_{CP} in a different range of allowed values (see Chapter 3).

The Wentworth Pit offers a number of practical advantages. As it is very recent and almost devoid of fish and other macroscopic life, there are few environmental concerns. Because the area is a private property shared between several mining companies, public access is restricted, and the site is guarded. The pit is easily accessible by road, with just a short stretch on dirt. It is only 24 miles away from Soudan, where the surface building of the MINOS experiment can be used as a workshop and manufacturing base. Due to its artificial nature, the pit is much deeper than typical lakes in the area, providing the necessary overburden. The water is also

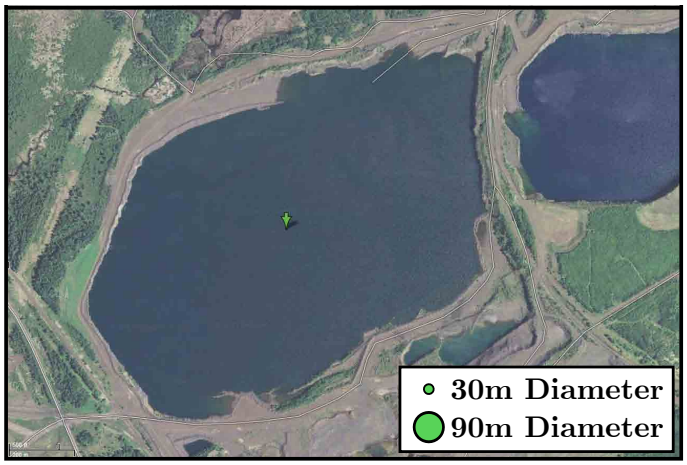

Figure 2.5: Satellite picture of the Wentworth 2W mine pit. The two size markers are shown in the same scale. Image taken from Ref. [93].

very clear, measuring 0.7 ± 0.5 Nephelometric Turbidity Units,[3] and approximately 2 to 3 m of attenuation length at 410 nm. The high clarity of the pit water facilitates the purification process.

2.2 Detector design

2.2.1 Cherenkov radiation

The CHIPS design employs a water Cherenkov detection technique, where neutrino interactions are observed by recording Cherenkov radiation produced by charged particles in water. Since water is a very cheap material and optically transparent, requiring only instrumentation around the volume, this technology has the best prospects to maximise target mass at low cost.

When a charged particle propagates through a dielectric medium faster than the phase velocity of light in that medium, it causes a shock wave of coherent emission of light in the visible and near ultraviolet range (Fig. 2.7). This forms a cone, centred on the direction of the particle, and projected on the downstream detector wall, forms a characteristic ring shape. The emission angle θ_C is related to the velocity of

[3]Turbidity describes the haziness of water due to particulate contamination. According to the World Health Organization, turbidity of drinking water should be below 1 NTU [95].

Figure 2.6: Topographic map of the Wentworth Pit 2W, based on a survey conducted when it was just starting to fill up with water. The areas in the lower centre of the pit are the deepest, with at least 50 m depth. Image taken from Ref. [93].

the particle $\beta = v/c$ and the refraction index of the medium n [96]

$$\cos\theta_C = \frac{1}{\beta n(\lambda)}. \tag{2.8}$$

In general, n depends on the light wavelength λ, and so does the emission angle. The refractive index of water is very close to 1.33 in the relevant wavelength range, which for ultrarelativistic particles with $\beta \simeq 1$ gives a value of $\theta_C \simeq 41°$.

The number of photons emitted by a particle with charge $\pm e$ per unit of distance travelled and per unit of wavelength is [96]

$$\frac{d^2N}{dxd\lambda} = \frac{2\pi\alpha}{\lambda^2}\left(1 - \frac{1}{\beta^2 n^2(\lambda)}\right), \tag{2.9}$$

where $\alpha \simeq 1/137$ is the fine structure constant. The PMTs used for detection of the Cherenkov light are sensitive to photon wavelengths between approximately 350 nm

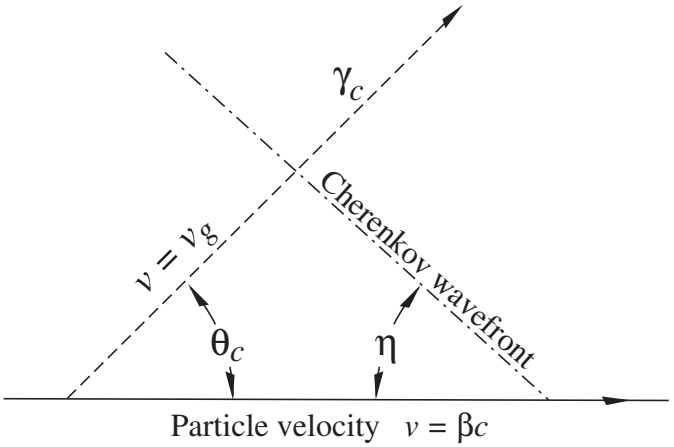

Figure 2.7: Diagram of Cherenkov radiation emission. The charged particle with speed greater than c/n emits a cone of light at an angle θ_C. Figure taken from Ref. [96]

and 650 nm. Integrating Eq. 2.9 in this range with the previous assumptions on β and n, gives an average number of around 240 photons per cm of distance travelled by the particle [1].

With such a relatively low amount of light, care must be taken to minimise losses.[4] Light absorption and scattering[5] in the water put a limit on the absolute detector size and require that the water is purified. This includes the removal of small particles and micro-organisms, and potentially chemical impurities as well.

2.2.2 The core design concepts

A CHIPS detector module contains a large volume of water, on the order of 10 to 30 kt, surrounded by inward-facing photomultiplier tubes. Although the structure is fully submerged in water, the volume has to be kept in a container that is both light- and water-tight. The former requirement is crucial to eliminate any external light sources, which could trigger the PMTs. The latter serves to keep the purified water from leaking or mixing with the dirtier water outside.

[4] In fact, a 3 inch PMT in CHIPS-10 will typically receive only a single photon.

[5] Although scattering does not actually reduce the amount of light detected in the event (at least in the case of uniform PMT coverage), it affects the particle identification ability by washing out the differences in Cherenkov light emission.

However, the tank does not have to withstand large forces. Differences in water density between the inside and outside (e.g. due to filtration) will only cause a small pressure difference. Without large loads, the mechanical structure can be very lightweight and the tank liner made of flexible material. A good candidate is geomembrane, used commercially to line water pools and reservoirs. A possible alternative is to have a lightweight but rigid wall structure built out of fibreglass panels.

The optimal shape for a Cherenkov detector is a sphere, which maximises the ratio of volume to surface area. This in turn minimises the number of PMTs needed to cover the walls for a given detector mass and photocathode coverage and keeps down the total cost, of which the PMTs are a major driver. However, a cylindrical struture is significantly easier do engineer, while still keeping the volume to area ratio high. It is also more suitable for a modular design where a single detector structure can be expanded by adding more horizontal levels.

An important feature of the design is the ability to recover the module after initial deployment and to expand it or add instrumentation. Such approach grants much greater flexibility in funding and construction effort, making it possible to split the assembly of one module over two or more seasons. This is especially important in the climate of the American Midwest, where the lakes freeze over in winter. All construction activities on water must be finished before the surface covers with ice, and no equipment can be left on the surface, as the moving ice in spring can easily destroy it. Therefore, any deployments will be performed during summertime.

Because CHIPS is dedicated to beam measurements, the optimal instrumentation may be very different than for multi-purpose detectors, such as Super-Kamiokande. For example, the photocathode coverage can be lower, since the typical interaction energies are much higher than for solar neutrinos or nucleon decays, and the amount of Cherenkov light depends almost linearly on energy. In addition, the beam neutrinos all come from the same direction, and most CC events preserve the directionality, producing activity in the downstream regions.

This knowledge can be applied to reduce the number of PMTs, which is a large factor in the total cost of Cherenkov detectors. The total photocathode coverage, i.e. the ratio of the combined photocathode area to the total surface area, can be less than 10 %, compared to 20 % at Super-Kamiokande. The upstream walls of

the detector, where little activity is expected from beam neutrino events, can be instrumented even less densely, acting mostly to identify background processes.

Past water Cherenkov detectors aimed to maximise the light yield and used PMTs with large photocathode diameters, e.g. with 8 inch in IMB and SNO, and a whole 20 inch in Super-Kamiokande. However, for the needs of CHIPS, smaller PMTs (e.g. 3 inch) are more suitable. Their cost per photocathode area is the same or cheaper than for bigger tubes, time resolution is typically smaller, and the use of a large number of small PMTs enhances the reconstruction performance due to higher granularity (assuming constant photocathode coverage) [19]. Digital readout systems, which can be housed underwater close to the PMTs, make it feasible to coordinate a large number of channels, without the need to control them individually from the shore.

For the readout, an existing solution can be chosen or a new one developed, and CHIPS pursues both avenues in parallel. The first option is realised in a collaboration with KM3NeT, which has developed a digital readout system for 3 inch PMTs, operating in optical modules underwater. Sharing their existing and tested hardware almost entirely eliminates development costs for CHIPS. The second option incurs additional costs, but allows for a customised solution, which can prove cheaper on a longer timescale. At the same time, innovative technology may help attract additional funding, alleviating the issue.

The water still has to be purified to limit light losses. Absorption and scattering of optical light increase with impurities and reduce the amount of Cherenkov light reaching the PMTs. Historically, water Cherenkov experiments aimed for ultra pure water with Super-Kamiokande reaching an attenuation length of almost 100 m at 420 nm wavelength [97]. To achieve such low levels of impurities, expensive deionisation filters and a reverse osmosis system are necessary. For CHIPS, due of the smaller module size and higher energy scale, the required attenuation length is less than 50 m for blue light and the goal is to minimise the cost of water treatment. Recent studies, performed on the Wentworth Pit water, suggest that this is possible with only mechanical filtering and UV sterilisation [98].

As a measure against the cosmic ray muon background, the detector may have an additional veto volume. It surrounds the main volume, with a light-tight material separating them optically. It can be instrumented with separate detection units

or with PMTs from the main planes, but pointing outwards. Muons entering the detector from the outside would first cross the outer volume before producing activity in the inner volume. If the Cherenkov light is observed by the veto PMTs, such events can be tagged as background and rejected.

2.2.3 Design of the CHIPS-10 detector

The CHIPS-10 module is the next step in the R&D program. It will be a fully functional detector, but simple and inexpensive enough to act as a technology development platform. The deployment of a first stage, consisting of two endcaps and a number of wall levels is planned for summer 2018. This section presents the current state of the module design, which is not yet fully completed; some details may still change.

The main structural components are the top and bottom endcaps of the cylindrical tank, made of a stainless steel truss (Fig. 2.8). The endcaps will be connected with numerous Dyneema[6] ropes around the perimeter, keeping the buoyant top cap anchored to the bottom one, resting on the lake bed. The whole structure will be wrapped in a geomembrane liner and sealed shut. The diameter of the endcaps will be around 25 to 30 m and the full module will be 20 m high. These dimensions correspond to the approximate fiducial volume of 10 kt, for a $\sim$1 m cut on the event vertex distance from a wall.[7]

The PMTs will be installed in units, called Planar Optical Modules (POMs) or planes, each containing multiple small PMTs and first stage readout electronics. Two types of planes will be used. The downstream area of the detector (downstream walls and parts of the endcaps), where most hits from beam events are expected, will be instrumented with high density planes. They will be built at Nikhef in Amsterdam and use 3 inch PMTs and KM3NeT readout hardware. Figure 2.9 shows a proposed design of the Nikhef plane with 30 PMTs, and a single PMT enclosure with a light

[6] Dyneema is a brand name for ultra-high molecular weight polyethylene fibres, which have comparable strength to steel, at much lower weight.

[7] The fiducial volume cut primarily serves to remove events with the vertex too close to the downstream wall, where the Cherenkov ring cannot be resolved. Because of this, the minimum distance to the downstream wall might be greater than 1 m, while for the upstream wall no cut will be necessary. The uniform 1 m cut was chosen as a simplified preliminary measure before studying the problem in more detail.

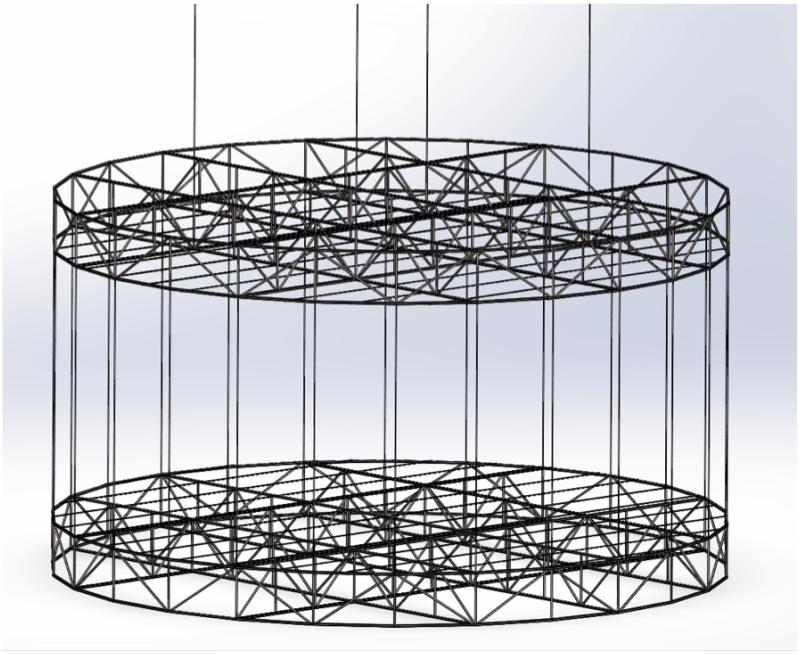

Figure 2.8: Proposed design of the steel truss endcaps connected with Dyneema ropes. The top cap is buoyant, keeping the ropes under tension. Planes with PMTs will be attached to the truss and suspended on the ropes. Image taken from Ref. [99]

collection cone. This design is based on a prototype unit built and tested in 2015 (Chapter 4).

The upstream area, useful for background discrimination, will be equipped with units with a lower photocathode coverage. They will contain old 3 inch PMTs, 400 of which have been donated to CHIPS by the NEMO-3 experiment, and new readout chips, currently developed at the University of Wisconsin in Madison.

Each detection unit will contain readout electronics, aggregating and digitising the signals from PMTs, and storing them before transmitting the data to shore. In the KM3NeT system, this is handled by custom-designed boards, with an FPGA-emulated processor at the core. The Madison planes, on the other hand, will utilise cheap BeagleBone Black single-board computers [101] to collect the PMT signals and send out data. Communication between the shore and PMT planes will use optical fibres and the White Rabbit system for time synchronisation [102]. Power will be transferred as high voltage AC current, with voltage converters placed in the pressure vessels with the electronics. Several switches located in the detector

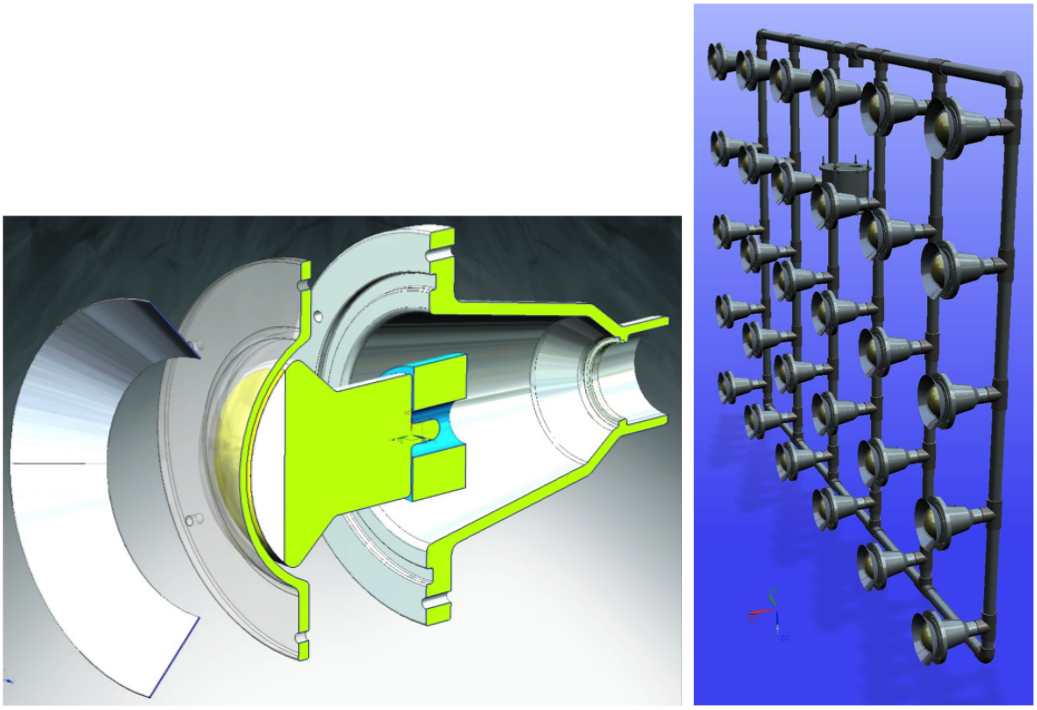

Figure 2.9: Proposed design of the Planar Optical Module using KM3NeT PMTs and readout (right) and an exploded view diagram of a single PMT enclosure in that unit (left). Image taken from Ref. [100].

will group signals from the planes and allow for using fewer fibres in the umbilical connecting the module with the on-shore DAQ station.

There are currently no plans for a veto region around the sides or the bottom of the detector. There will, however, be veto PMTs installed in the top endcap. They will be part of the top instrument planes, except pointing upwards. A light-tight sheet will separate the inner volume from the top veto volume of the detector.

CHIPS-10 will be connected with the shore via an umbilical pipe resting on the bottom of the pit. It will contain a bundle of cables, optic fibres and PVC pipes, providing power, communication and water circulation to the detector. The on-shore station, likely built out of modular segments, will house the data acquisition system and the water treatment plant. Electric power will be supplied from the grid, via a medium voltage power line available on site.

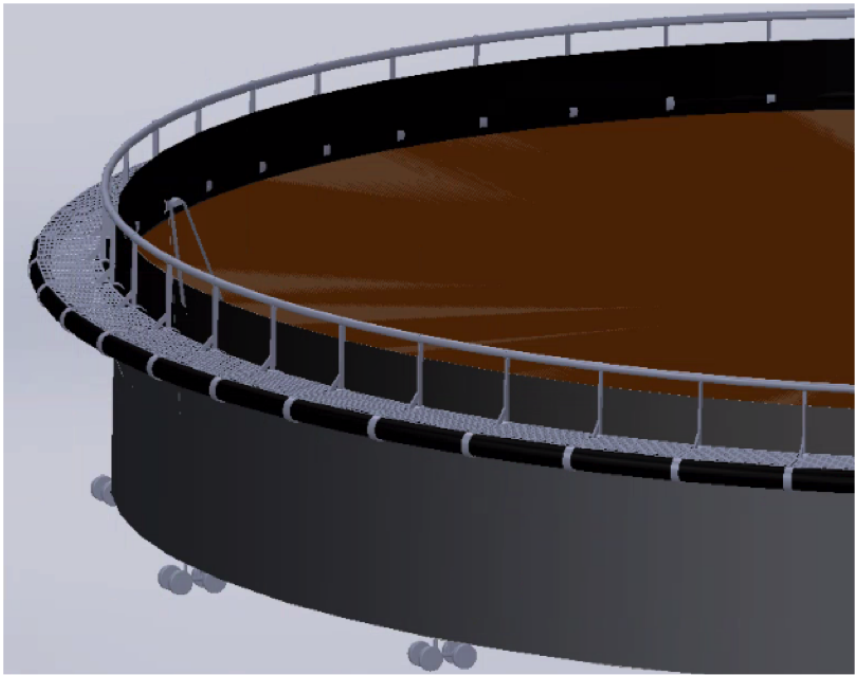

Figure 2.10: Illustration of the CHIPS-10 construction process. The floating ring provides buoyancy and work area, together with the platform in the middle. The submerged part is already completed and surrounded by the liner. Image taken from Ref. [103].

2.2.4 Assembly and deployment procedure

The assembly of the detector is a vital consideration in the design process. The module will be constructed mostly on water, and needs to be modular enough to allow for recovery and refitting with additional instrumentation in the future.

According to the current plan, the two endcaps will be assembled on land and equipped with the PMT planes, cabling and electronics boxes. Then, the top endcap will be placed on the bottom one, and the whole assembly will be moved on wheels, down the ramp into the pit water. A floating ring platform, used commercially for fish farming, will be built around the structure to support the weight and provide working space (Fig. 2.10). A floating deck on the top endcap will allow workers to install the side PMT planes on the Dyneema ropes, held above the water surface with lightweight cranes. On the outside of the ring, a welding team will attach a next layer of the geomembrane liner to the previous one.

Figure 2.11: Diagram of the deployed CHIPS-10 module prior to the disassembly of the floating platform. Image taken from Ref. [103].

After a whole storey is completed, the bottom endcap will be lowered with winches, while the detector is filled with purified water, and the construction of a next level will begin. When the module is finished, the liner will be sealed shut with a top cover, and the complete structure will be lowered to the lake bed (Fig. 2.11). Afterwards, the floating platform will be disassembled, in preparation for the ice covering the pit in winter.

The winch ropes will be tied to underwater buoys, so that the module can be easily recovered. After the recovery, the liner can be cut open and new layers installed following the same procedure. The only requirement is that the communication and power cables for future PMT planes are installed before the first deployment.

2.3 Monte Carlo simulation

Monte Carlo (MC) simulations are an indispensable tool in particle physics, both during the design stage and later for data analysis. Especially for expensive projects,

where iterative design revision is slow or impossible, the ability to simulate different scenarios in advance is crucial. A good model of the detector also allows for the development and testing of event reconstruction techniques, and the study of potential physics sensitivity.

The Monte Carlo method uses pseudo-random number generators to approximate calculations, which would be intractable or impractical to compute analytically. It is also well suited to generate probabilistic data. A typical MC simulation starts with an event generator: using information about the beam and the target it calculates the probability distributions of the possible outcomes, and generates a series of possible events by sampling those distributions. The same method is then used to simulate the propagation of particles through the detector medium. For example, an electron might have a certain probability of interacting with the hydrogen or oxygen nuclei of water in the detector and emitting a bremsstrahlung photon. To account for this, the simulation uses a pseudo-random number generator at each step of the simulated track to determine whether to create a new photon track object and to calculate its energy, direction and other parameters.

A useful MC simulation will provide output that matches closely the observables in a real detector. That makes it possible to compare the registered output to theoretical predictions without the need to unfold all the detector effects, which might not be possible for a sufficiently complex detector. At the same time, simulations must be carefully validated to make sure that they approximate the reality well enough to be used in the analysis of physics results.

As explained in the introduction chapter, the simulation and reconstruction work described in this and the following section is primarily the work of others. The author's contribution to the development of the software was mostly focused on support and maintenance. The following description is presented because of its vital role in the following chapter, detailing the sensitivity study based on the results of the preliminary reconstruction and particle identification analysis, performed with these tools.

2.3.1 Detector simulation

The MC detector simulation is built with the Geant4 framework [104, 105] and is based on the WCSim package [106]. WCSim was initially developed for the LBNE project (now DUNE/LBNF), to study the viability of a water Cherenkov far detector, before the liquid argon option was chosen. It was designed to allow easy creation of generic water Cherenkov models in Geant4. The original code has been extensively modified and in places completely rewritten to adapt it to the needs of the CHIPS project. The main goal was to further simplify the process of creating and loading different geometries, particularly without the need to recompile the source. The physics lists used in the simulation are default to Geant4. Material properties are mostly hardcoded in WCSim and come from the Super-Kamiokande simulation SKDetSim. No additional validation was performed specifically for the needs of the CHIPS project as of yet.

The detector definitions are stored in external configuration files and loaded by the simulation software at runtime. Currently, the only used geometry is a regular polygonal prism, approximating a cylinder.[8] The height, diameter and number of sides can be specified. The shape forms a tank of water with black sheet walls.

The layout of the PMTs is defined as a unit cell. A single cell can contain any number of PMTs of any type, with specified relative positions. The detector walls are then tiled with copies of the unit cell, adjusting the size to produce desired photocathode coverage or to utilise a given number of available PMTs.

To allow for a design with a non-uniform PMT coverage, optimised for the detection of beam neutrino events, the detector is split into zones, which are grouped into regions. Each side of the polygonal prism is a separate zone, and the endcaps can be divided by specifying a set of polar angles (Fig. 2.12). Each region can use a dedicated definition of a unit cell and the desired coverage. A PMT in the cell definition can also be given an angle relative to the wall, different from the default 90°. This may be used for example to have the top and bottom endcap PMTs face more directly upstream.

[8] A polygon was chosen over a perfectly cylindrical shape mainly due to simplicity of implementing the algorithm tiling walls with PMT units. In addition, since the walls of the actual detector will be composed of flat planes suspended between Dyneema ropes, a polygon with appropriate number of walls is actually a better representation of reality

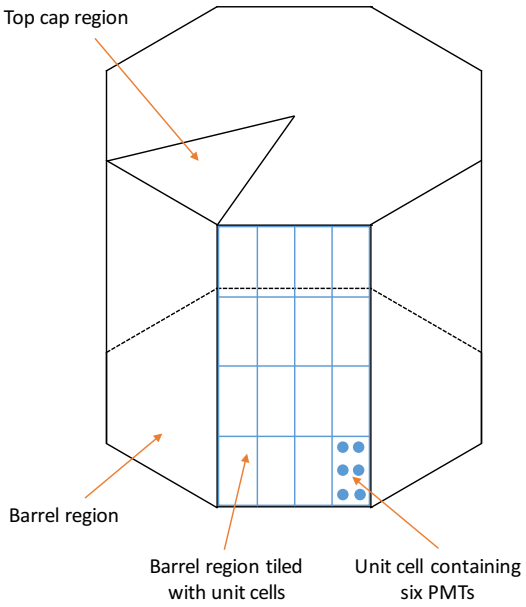

Figure 2.12: Schematic diagram of the detector geometry definition, illustrating the shape, regions and a single PMT cell. Image taken from Ref. [19].

The available types of PMTs are defined in another configuration file. The stored information includes the photocathode diameter, glass thickness, the distance it protrudes from the black sheet, time resolution constant and quantum efficiency, defined for a number of photon wavelengths. It is also possible to add a Winston cone to increase light collection. The geometries of the cones are defined in a separate file, as a list of points describing the radial profile.

There is also a configuration file with material properties affecting the propagation of optical (e.g. Cherenkov) light. It contains the reflectivity of the black sheet material covering the walls (which is around 4 % for the wavelength range of interest), refraction index of the PMT glass, as well as absorption and Rayleigh scattering attenuation lengths in water. It is possible to enable Mie scattering, but it is expected to be negligible in sufficiently pure water and is therefore switched off.

In addition to the inner volume of the detector, a veto volume can be specified. It forms a concentric shell around the main detector volume on all sides, parametrised by its thickness. The veto volume can be instrumented by designating PMTs to point outwards instead of inwards in the PMT cell definition. This mirrors the design concept, in which inner volume and veto PMTs would be located on the same plane, and a light-tight sheet would separate the two. The inner walls of the veto region

are covered with a white sheet, which has high reflectivity (90 %) to maximise the light yield with few PMTs.

The main code first loads and builds the detector geometry with Geant4. For each event to be simulated, it takes a list of primary particles, with their type, energy, position and direction. These can come from an underlying neutrino event generator, cosmic ray muon generator or a built-in particle gun (used mainly for testing). Then, the program simulates the passage of the particles through the detector material, including possible interactions and decays as well as the emission and propagation of Cherenkov photons. When a photon hits the photocathode of a PMT, a hit is recorded and sent to the digitiser simulation.

The digitiser converts the true photon hits into information resembling the final output of a readout system, with a timestamp and total charge. All hits within a time window of 200 ns are grouped together, contributing to a single digitised hit.

The total charge in photoelectrons can be computed in two ways. The original method included in WCSim is based on the Super-Kamiokande simulation SKDetSim. For each incident photon, it probes the single photoelectron charge distribution, and returns the sum of the individual results. A new PMT simulation was also written; it estimates the electron current on the dynode level, and accounts for non-linearity effects at high charges [19].

The digitised time is taken as the time of the first hit, smeared with a Gaussian to account for the PMT resolution. The Gaussian width is $\sigma = (0.33 + \sqrt{C_t/q})$ns, where C_t is a constant defined for each PMT type, and q is the total digitised charge.

The list of digitised hits for each simulated event is stored in an output file as a ROOT tree structure [107], alongside truth information describing the actual properties of the primary particles, such as vertex coordinates, directions, and energies.

PMT dark rate and electronic noise are not yet simulated. However, analyses performed on the simulated data include cuts on the hit charge, such as those that would be used to reject hits caused by noise.

2.3.2 Neutrino event and flux simulation

The calculation of the expected beam flux at a given location uses the existing beam simulation written for the needs of NuMI experiments. The simulation stores the kinematic of neutrino parents, i.e. particles that decayed to create the beam neutrinos. Because the decays are isotropic in the centre of mass frame, it is possible to compute the energy of a neutrino emitted in any given solid angle and the probability of it, applied as an effective weight to the entry in the energy distribution histogram.

The location of the detector expressed in geographic coordinates, specifying the location on the WGS84 ellipsoid,[9] can be transformed to coordinates in the beam reference frame. A set of weights and energies for each entry in the simulated flux is then computed and applied to produce the flux prediction for the new location [93]. Figure 2.13 shows the expected muon event rates (without oscillations) superimposed on a map of northern Minnesota, and L/E contours for the peak beam energy. Figure 2.4 shows the calculated ν_μ energy spectrum for the three locations highlighted on the map, while Figs. 2.14 and 2.15 show all flavour components at Wentworth Pit, for both the neutrino and antineutrino beam mode. This data is also used as input to the GLoBES simulation of CHIPS-10, described in Chapter 3.

To generate muon and electron neutrino events in a CHIPS detector located in the Wentworth Pit 2W, the oscillation probability is applied to the muon neutrino flux spectrum. The $\nu_\mu \to \nu_\mathrm{e}$ formula is applied, with the following parameters

$$\sin^2\theta_{23} = 0.43, \qquad \sin^2\theta_{13} = 0.0945,$$
$$\Delta m^2_{32} = 2.39 \times 10^{-3}\,\mathrm{eV}^2, \qquad \delta_\mathrm{CP} = 0,$$

and the solar scale terms neglected [19]. This calculation does not need to be very precise, as its role is only to estimate the correct energy distributions and relative sample sizes of the simulated events.

The oscillated fluxes are used as input to the Genie neutrino event generator [108]. The primary particle tracks generated by Genie are stored in a NUANCE-formatted file and passed to the Geant4 simulation. The event vertices are placed randomly

[9] The WGS84 is a geodetic standard defining reference frames allowing for conversion between coordinate systems. The ellipsoid is an Earth centred reference shape, differing from the more precise WGS84 geoid by ±100 m.

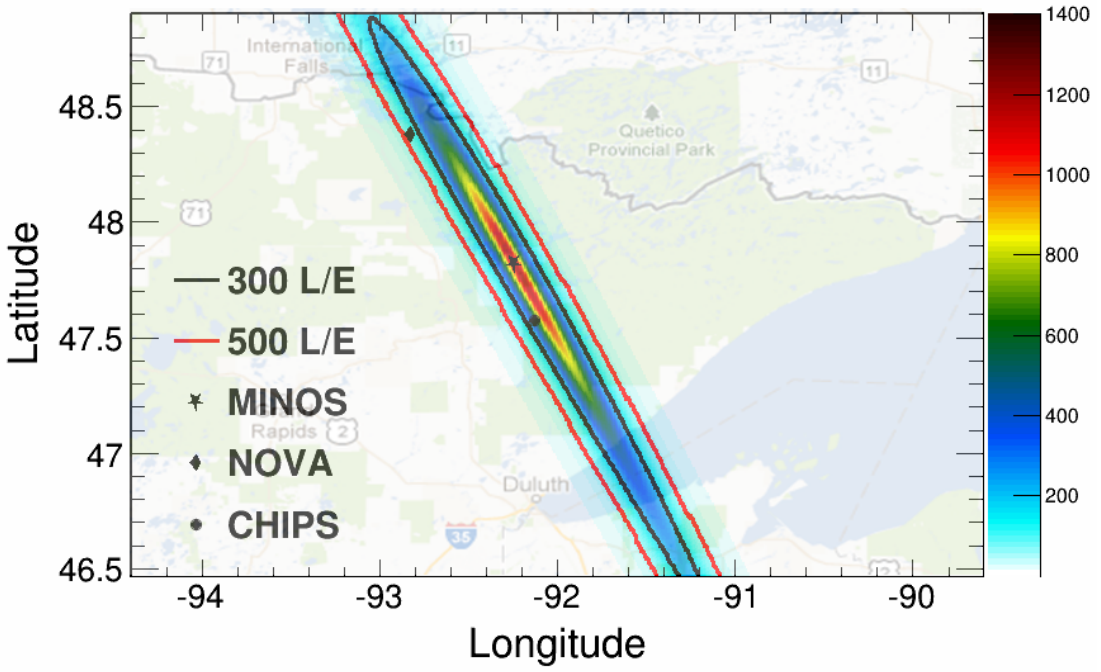

Figure 2.13: Map of the expected CC ν_μ event rate from NuMI in northern Minnesota, in case of no oscillations, at 1 kt × year exposure. Markers show locations of MINOS and NOvA detectors, as well as the Wentworth Pit for CHIPS. Contour lines connect points of constant L/E in km/GeV for the peak beam energy. Image taken from Ref. [93].

within the detector volume, and the initial neutrino direction is aligned with the beam axis.

Figure 2.16 presents an example event display of a CC ν_e interaction. It shows the unrolled walls and endcaps of the detector, with coloured dots for each PMT that registered a hit. The colour can indicate the total charge (as in this plot) or the hit time. The event display can also superimpose outlines marking the Cherenkov cones for true or reconstructed tracks of primary particles.

It is also possible to simulate detector events caused by incoming cosmic ray muons, in order to study their impact and veto capabilities. Muon events are generated with the CRY package [109]. The code uses CRY libraries to generate muon events on the surface of the lake, and propagates them towards the underwater detector employing a simplified linear energy loss of 220 MeV/m. The muon tracks

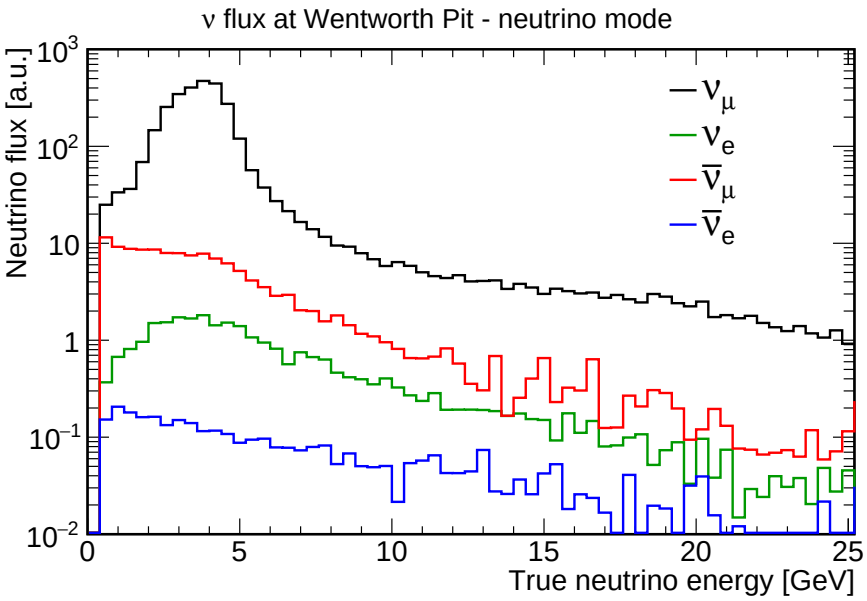

Figure 2.14: Energy distributions of the NuMI neutrino flux at Wentworth Pit 2W, in case of no oscillations. Lines are shown for different components of the neutrino mode beam. The ν_τ component is negligible and not predicted by the simulation.

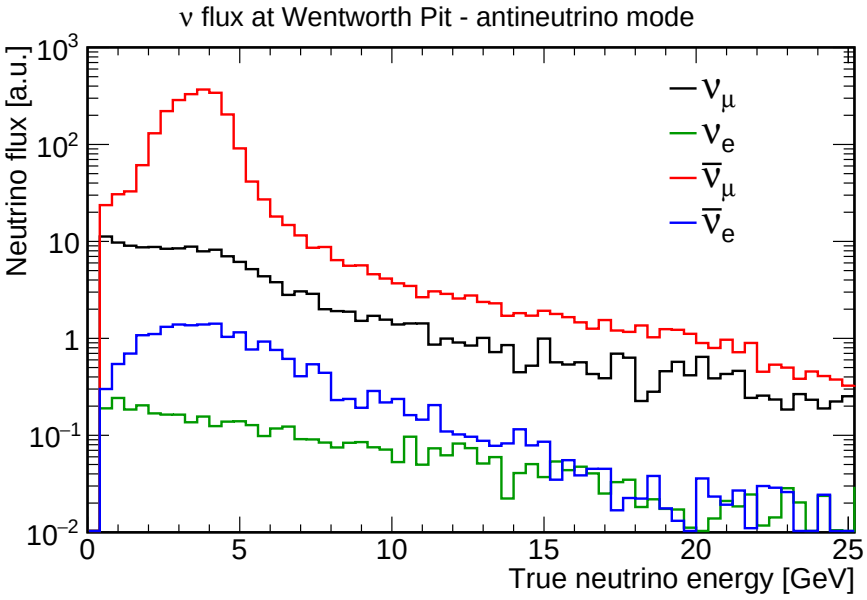

Figure 2.15: Energy distributions of the NuMI neutrino flux at Wentworth Pit 2W, in case of no oscillations. Lines are shown for different components of the antineutrino mode beam. The ν_τ component is negligible and not predicted by the simulation.

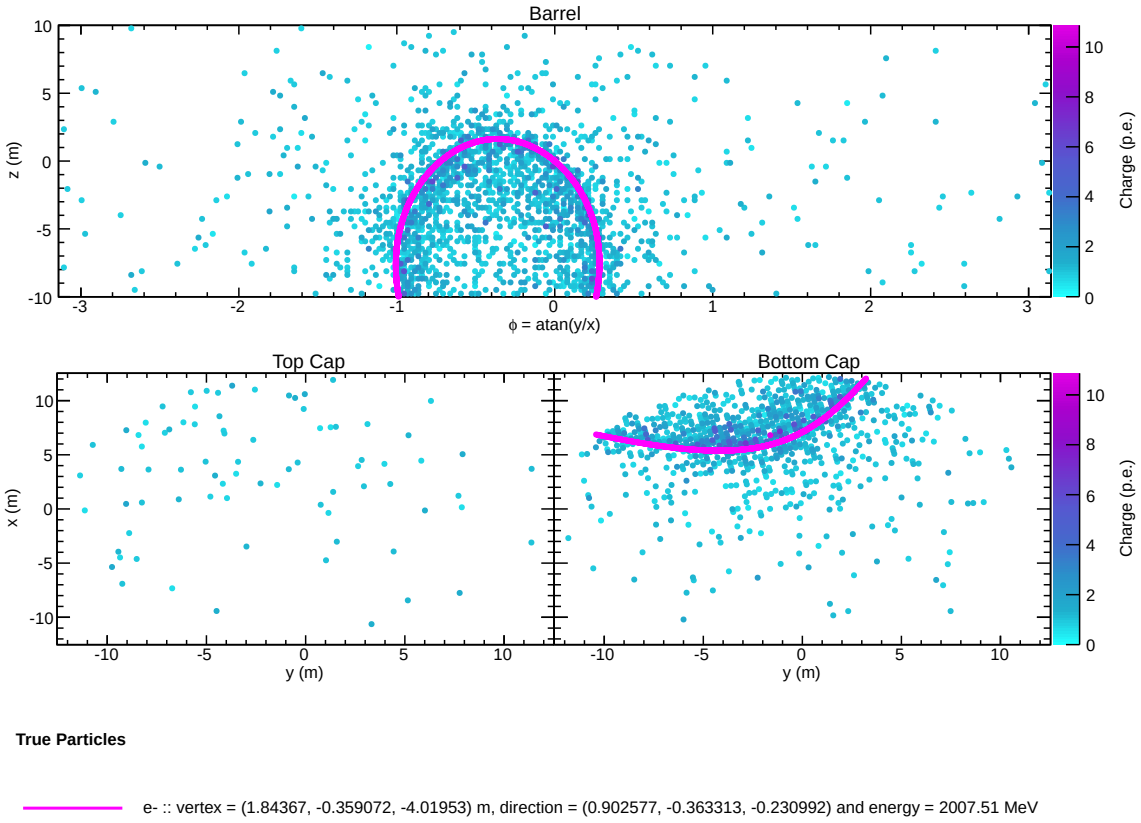

Figure 2.16: Event display of a simulated CC ν_e event in a CHIPS-10 detector instrumented with 3 inch PMTs, with 6 % photocathode coverage. The display shows the unrolled walls and the two endcaps separately. Every coloured dot represents a PMT with a registered hit, and the colour signifies the total charge in photoelectrons. The pink outline shows a projection of the Cherenkov light cone associated with the true electron track.

are saved in the same input file format as the Genie events, with vertices placed 5 m above the top cap of the detector.

2.4 Reconstruction software

Alongside the simulation, a preliminary reconstruction package was developed [19]. Its main purpose is to explore the capabilities of different detector designs, especially with low photocathode coverage and small PMTs. It also serves to evaluate a novel reconstruction algorithm, well suited for ν_e identification. The goal of the

reconstruction is to recreate tracks and types of the primary particles from the digitised PMT hit information. This makes it possible to determine the flavour and kinematics of the initial neutrino.

The reconstruction algorithm is based on a maximum likelihood method used by the MiniBooNE experiment [110, 111]. The pattern of PMT hits, including the time and charge information, is predicted from a set of parameters describing the primary particle tracks. Comparison of the prediction with observed values returns the likelihood of the track hypothesis, which can be minimised by varying the parameters. The procedure is described in Ref. [19] and in greater detail in Ref. [1].

An alternative reconstruction method used in water Cherenkov detectors involves using the Hough transform to find ring shapes created by the hit PMTs on the detector walls. Muon tracks and electromagnetic showers induced by electrons or gamma photons can be identified by measuring the fuzziness of the ring, i.e. the fraction of hits outside the ring boundary; the showers produce much fuzzier shapes than single particle tracks.

However, this method performs poorly for events with multiple rings, especially if the separation between them is small. At the same time, two electromagnetic rings with a small angular separation are a signature of the NC interaction with an outgoing neutral pion. The pion almost always decays into two photons ($\pi^0 \to \gamma\gamma$), which, for GeV energies of the π^0, are boosted forward. When the two rings almost merge, or if one of the photons has much lower energy and is less visible, the event can be easily misidentified as a single electron signal. The maximum likelihood algorithm is reported to handle the π^0 events much better [112], which is important for minimising background from that channel.

The detector simulation is used to evaluate the reconstruction performance, by comparing the reconstructed track parameters with the MC truth information. On the particle identification level, multiple hypotheses are fitted to different simulated event samples and a selection procedure is created to optimise signal efficiency and background rejection.

Before running the maximum likelihood fit, the reconstruction begins with a track seeding algorithm, meant to produce approximate values of the starting track parameters. First, the list of PMT hits is divided into slices of activations close in time and space, to separate hits associated with multiple tracks. A set of vertex-finding

and Hough-transform-based algorithms is applied separately to each of the slices, returning vertex position and time, as well as directions for the primary and potential secondary tracks. While not strictly necessary, this step dramatically improves the speed and performance of the main reconstruction algorithm.

2.4.1 Overview of the maximum likelihood method

A single muon or electron track is described by a vector of parameters $\vec{x}$, containing:

- Vertex positions x_0, y_0, z_0 and time t_0;
- Track direction, parametrised by two angles θ_0, ϕ_0;
- The initial kinetic energy E_0;

and, implicitly, the type of the particle. Since the electron track actually describes an electromagnetic shower, a photon hypothesis is identical, except for the conversion distance – distance from the vertex of the neutral γ to the beginning of the shower – which is added as an additional parameter.

A multi-track hypothesis contains a parameter set for each track, with possible constraints reducing the number of dimensions. For example, a π^0 hypothesis consists of two photon tracks, but with the same vertex coordinates and the energies and directions constrained by the invariant mass of the π^0.

The total likelihood of observing the hits in the detector with measured values of time and charge, assuming that hits in individual PMTs are conditionally independent, is a simple product

$$\mathcal{L}(\vec{x}) = \prod_{\text{unhit}} P_{\text{unhit}}(\vec{x}) \times \prod_{\text{hit}} P_{\text{charge}}(\vec{x}) P_{\text{time}}(\vec{x}), \tag{2.10}$$

where the first term gives the likelihood that the PMTs without observed hits would not receive any photons from the hypothesis described by $\vec{x}$, and the second term tells how likely it was that the hypothesised track(s) produced the observed values of hit times and charges.

Considering the negative log likelihood allows one to transform the product into a sum over the PMTs, simplifying the calculations

$$-\log \mathcal{L}(\vec{x}) = -\sum_{\text{unhit}} \log(P_{\text{unhit}}(\vec{x})) - \sum_{\text{hit}} \log(P_{\text{charge}}(\vec{x})) - \sum_{\text{hit}} \log(P_{\text{time}}(\vec{x})). \quad (2.11)$$

In addition, the charge and time parts can be separated. In practice, the probability of not registering a hit $P_{\text{unhit}}(\vec{x})$ is covered by the charge term $P_{\text{charge}}(\vec{x})$, where the observed charge, $q = 0$. This splits the likelihood computation into the charge part, calculated for each PMT, and the time part, applied only to PMTs with an observed hit.

In a single pass of the fitter, the total negative log likelihood of the track hypothesis is obtained, by summing the charge and time contributions for all PMTs. Then, the track parameters are varied in order to minimise it. After the fit converges, the output consists of the final parameter values, as well as the corresponding total charge and time likelihoods.

2.4.2 Charge likelihood

The charge part of the likelihood method is similar to the original one used in MiniBooNE [110]. However, it has been implemented from scratch and adapted to the cylindrical geometry of a water Cherenkov detector as opposed to the spherical MiniBooNE detector filled with mineral oil. The underlying principle is to use low-level information about the emission and propagation of Cherenkov light to predict the mean number of photons hitting the photocathode of a given PMT (μ). The probability of observing a photoelectron charge q, given μ, and hence the likelihood of the predicted μ to cause the observation, is given by a Poisson distribution, which can be expanded to non-integer values of q by the Gamma function.

The mean photon prediction is evaluated as a sum of approximately point source emissions of Cherenkov photons in the direction of the PMT, over the length of the track. The number of photons reaching the PMT from a short segment of the track, indexed by i, can be expressed as [19]

$$\mu_i = \Phi_i(E, s_i, \cos\theta(s_i))T(R_i)\epsilon(\varphi_i)\frac{\Omega(R_i)}{4\pi} \quad (2.12)$$

and it depends on:

- The number of Cherenkov photons emitted across that segment, in the direction of the PMT, Φ_i;
- The survival probability of the photon to reach the PMT (i.e. probability of not being absorbed or scattered away), T, depending on the distance to the PMT R_i;
- The angular efficiency of the PMT from the direction of the track segment ϵ, depending on the incident angle on the PMT φ_i;
- The fractional solid angle subtended by the PMT as seen from the track segment, Ω, depending on the distance R_i.

The first value Φ_i, can be expressed as a product of three parts

$$\Phi_i = \Phi(E)\rho(E, s_i)g(E, s_i, \cos\theta(s_i)), \tag{2.13}$$

which are called emission profiles and represent:

- The total number of Cherenkov photons emitted by a particle, Φ, depending only on its energy E;
- The fraction of light emitted in that track segment ρ, depending on the energy and distance along the track s_i;
- The fraction of light emitted along the track segment in the angular direction $\cos\theta + d\cos\theta$, g, normalised to the total number of photons emitted along the segment, depending on the energy, distance s_i and the angle θ of the PMT relative to the track (itself a function of s_i).

These quantities are evaluated from large samples of the Monte Carlo simulation, binned in the parameters and stored in tables. Importantly, most of them depend only on a small part of the simulation or geometry, and only have to be recomputed if that component changes, making the prediction very modular.

To save computation resources when evaluating the final sum for each PMT and every track hypothesis, the term $T(s_i)\epsilon(s_i)\Omega(s_i)/4\pi$ (R_i and φ_i can be calculated from s_i) is approximated by a quadratic function. The remaining part is precomputed in bins of track energy and zenith angle, and the distance to the PMT, and stored

in look-up tables. During runtime, the coefficients of the quadratic functions are evaluated from three point on the track, and in combination with the look-up tables, the photon number prediction is computed, without the need to perform costly integrations [1, 19].

2.4.3 Time likelihood

The first attempted time likelihood implementation imitated the one described in Ref. [110], where PMT hit times, corrected for the time of flight, are parametrised in terms of particle energy and predicted charge in a series of incremental fits. A preliminary version of the method, adapted to CHIPS by the author, is described in Ref. [113].

The Monte Carlo simulation is used to produce a large sample of events with primitive muon or electron tracks, with random direction and vertices. For each PMT with a hit, the average time of light and particle propagation is subtracted. Distributions of the resulting corrected time are stored in bins of track energy and the predicted charge at the PMT.

The distributions are parametrised in three stages of function fits: first, they are fitted with a Gaussian distribution convoluted with exponential decay; next, parameters of the first fit are fitted with a polynomial to capture their dependence on the charge; and finally, the secondary parameters are fitted with polynomials to parametrise their dependence on the particle energy. The corrected time distributions can be reconstructed from the set of tertiary parameters, and used as the likelihood function of a track hypothesis (with energy E_0 and predicted charge μ) to cause the observed corrected time t_c (which in turn depends on the vertex coordinates and the track direction).

Unfortunately, this method presents multiple problems. Its inputs contain the final results of the full detector simulation and the charge prediction algorithm, which makes it difficult to validate the underlying assumptions and requires a complete reprocessing if any of the components changes. Furthermore, the numerous fits required to parametrise the distributions in three stages pose practical problems of instability, which prevents the desired level of automation. Due to all those issues, the approach was deemed impractical and a new method was developed.

The new algorithm, developed from scratch by A. Perch, takes a similar approach as the charge component, and aims to predict the time of a hit from first principles of Cherenkov photon emission [1]. For each short track segment, it predicts the time it will take the photons to reach the PMT with a simple calculation. Then, it uses the Cherenkov photon emission profiles (similar to the ones used for charge prediction), to get the fractional amount of light reaching the PMT from this track segment, which is applied as a weight in the predicted mean arrival time of a single photon.

To account for multiple photon hits, where the registered time is that of the first hit, the distribution of the mean arrival time is approximated by a Gaussian. Then, the first hit time is predicted by an analytically derived distribution $P(t_{\min}; n)$, which describes $\min(t)$ for n samples of t, each distributed according to these Gaussian parameters. The likelihood function is taken as the overlap between the $P(t_{\min})$ distribution and the observed hit time, smeared by the time resolution of the PMT [1].

2.4.4 Particle identification and first results

A preliminary particle identification classifier, for detecting electron neutrino charged current events, was built using the reconstruction method outlined above [19].

The studies were performed for a CHIPS-10 detector geometry, a 20-sided prism, with inner height of 20 m and 25 m diameter. Three instrumentation versions were evaluated: with 10 inch PMTs and 10 % photocathode coverage, with 3 inch PMTs and 10 % coverage, and with 3 inch PMTs and 6 % coverage.

In each instrumentation option, the performance of the track reconstruction was tested by fitting single tracks to a sample of CCQE ν_e and CCQE ν_μ events, and comparing the results with the Monte Carlo truth [19]. In each case, the track hypothesis was set to correspond to the flavour of the leading lepton. For each CHIPS version and each of the two flavours, the resolution for the track parameters was computed by taking the width of the distribution of the difference between the reconstructed value and the true one. This was done for the total difference in position and direction, as well as time and energy.

Table 2.1 shows the comparison of the resolution on different parameters for the three instrumentation options. In general, the 10 % photocathode coverage version with 3 inch PMTs is at least as good or slightly better than the one with 10 inch

Sample	Geometry	Reconstruction Resolution			
		Vertex Position (cm)	Vertex Time (ns)	Direction (°)	Energy (MeV)
CCQE ν_e	10 inch, 10%	35	0.9	2.1	208
	3 inch, 10%	35	0.84	1.9	210
	3 inch, 6%	38	0.89	2.1	211
CCQE ν_μ	10 inch, 10%	47	1.35	2.6	113
	3 inch, 10%	44	1.14	2.7	110
	3 inch, 6%	51	1.28	3.0	113

Table 2.1: Comparison of the track reconstruction performance for three instrumentation options of CHIPS-10, described by the diameter of the PMT and the photocathode coverage. The resolution values correspond to the widths of the reconstructed minus true distributions after fitting a track with the correct particle hypothesis (e.g. electron tracks to ν_e events). Table taken from Ref. [19]

PMTs, due to increased granularity. The 6 % coverage version is only minimally worse then the former two, while offering 40 % savings in the total detector cost [19]. The results presented in this section are for that last version, which corresponds to the most likely final design for CHIPS-10.

The following event samples were generated to train and test the classifier:

- CC ν_e and CCQE ν_e interactions as signal,
- CC ν_μ and CCQE ν_μ interactions as background due to misidentified muons,
- NC interactions as background due to misidentified hadronic activity.

First, all events in the samples were tested under both a muon and an electron single-track hypothesis, where the combined charge and time likelihood is minimised. Because the likelihood surface is multi-dimensional and has many local minima, the actual procedure involves multiple passes of the fitter, alternately fixing and freeing different parameters [19].

The maximum likelihood method returns the most likely parameters for every tested track hypothesis, as well as the charge and time likelihood values corresponding to those parameters. The simplest way of discriminating between muon and electron events is to compare the total likelihood values in the muon and electron fits. However, the classification can be enhanced by using the charge and time parts separately, and employing additional variables. They can be processed by a multivariate analysis method returning a single classification variable as output. This approach also allows

for discrimination between CC ν_e and NC interactions, without explicitly considering a neutral current hypothesis.

Two artificial neural networks were created, one for ν_μ–ν_e classification and one for NC–CC ν_e classification. Both use almost the same sets of input variables, which include the minimised charge and time likelihood values when fitted as electron or muon, the total number of PMT hits and several other quantities describing how well the fit reflects the observed energy deposits in the PMTs [19].

The neural network discriminating between muon and electron tracks was trained on the reconstructed events from the CCQE ν_μ and CCQE ν_e samples. The quasi-elastic interactions, with no particles beside the lepton and recoil nucleon, have simple topologies. This makes it easier to train the selection, simplifying the first particle identification attempt. The network returns a single number, $a_{e\mu}$, which takes values close to 0 for ν_μ-like (background) events and close to 1 for ν_e-like (signal) events. The second network is trained on NC events as background, and CCQE ν_e events with $a_{e\mu} > 0.8$ as signal. Its output value, a_{NC}, follows the same convention as $a_{e\mu}$.

The final selection cuts on the neural network outputs must be chosen to optimise the signal efficiency and background rejection, maximising the physics sensitivity. These quantities were evaluated on a larger sample including non-quasi-elastic CC ν_μ and ν_e events, where all CC ν_e events count as signal. The assumed composition of events recorded in the CHIPS-10 detector, based on Genie studies, is 30 % NC and 70 % CC interactions, out of which 20 % are CCQE interactions. Muon neutrino events constitute 95 % of the total, and electron neutrinos make up 5 %. Note that this calculation takes into account the oscillation probability and interaction cross sections, on top of the unoscillated neutrino flux shown in Figs. 2.14 and 2.15.

The figure of merit being optimised is the product of signal purity and efficiency, shown on Fig. 2.17 as a function of cuts on $a_{e\mu}$ and a_{NC}. Figure 2.18 shows the final selection efficiencies for different samples, which are used as input to the GLoBES simulation in the sensitivity study (Chapter 3). The total efficiency of the CC ν_e event selection is 30 %, and the purity of the signal sample is 58 % [19].

These results are the very first attempt at predicting the capabilities of a CHIPS detector, based on a dedicated MC simulation and an actual implementation of a

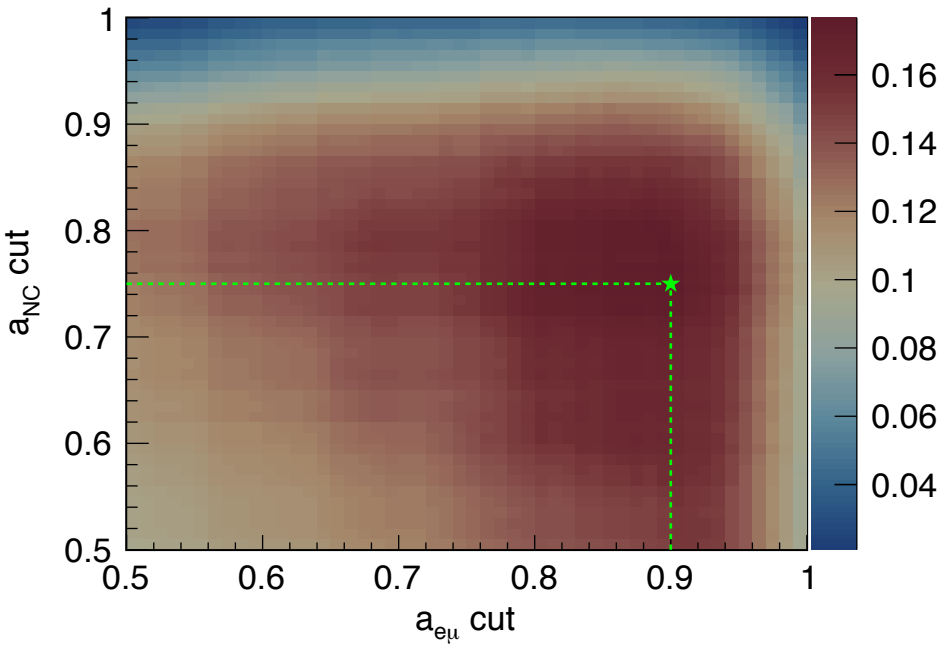

Figure 2.17: Efficiency times purity of the final ν_e CC selection as a function of cuts on the output values of the two selection neural networks. The green marker indicates optimal cut values. Figure taken from Ref. [19]

full reconstruction chain. While very preliminary, they are encouraging, and provide valuable data for a realistic examination of future physics sensitivity.

However, there are many areas of improvement left. The primary next step is the implementation of a dedicated π^0 hypothesis and fit procedure. It will enable the utilisation of the multi-track capabilities of the maximum likelihood reconstruction and reject a significant amount of NC background. A preliminary version of a π^0 fitter is already finished and will be soon incorporated into the full selection [19].

Next, the particle identification procedure will be expanded to better distinguish non-quasi-elastic CC interactions of muon and electron neutrinos, which have more complex topologies. Although the classifier trained on the CCQE samples still performs well for non-QE events, a better optimised method should further increase accuracy.

In addition, cosmic muon background identification has to be studied in detail. With the very high overall rate, even a very small fraction of muons misidentified as signal events can form a non-negligible background. A special track seeding procedure for tracks originating outside of the detector, using hit information from

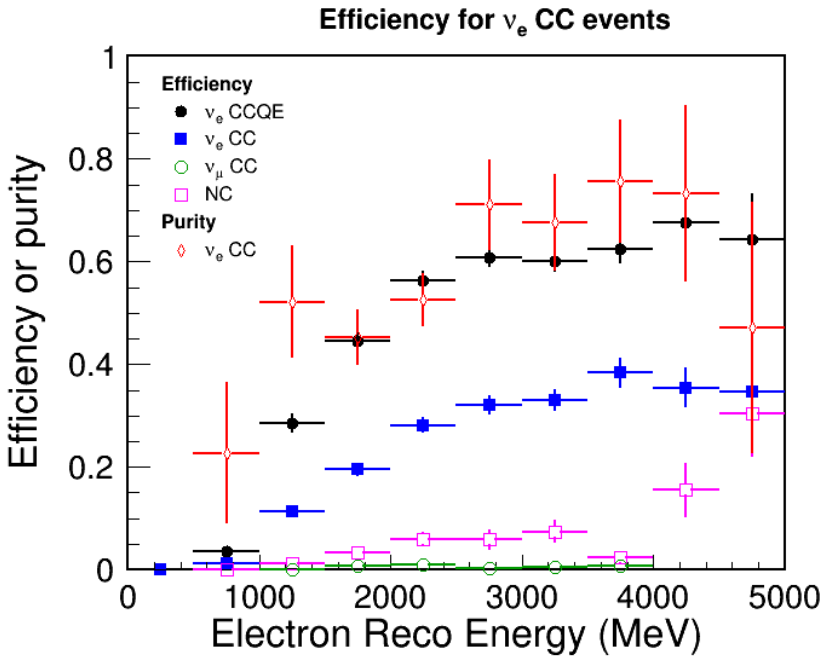

Figure 2.18: Final CC ν_e selection efficiencies for various classes of signal and background events. The purity of signal events in the selected sample is also shown. Figure taken from Ref. [1]

the veto region, has been developed, and a technique to fit beam neutrino interactions overlapping with cosmic muon events is being tested [19].

Finally, the performance of the reconstruction and classification will be tested for a non-uniform PMT layout, in line with the proposed design of the CHIPS-10 module. This will validate the design principles and guide further optimisation. Since all the necessary simulation tools are already in place, this will be completed soon.

Chapter 3

Physics sensitivity of CHIPS

Predicting the physics reach of CHIPS is crucial for evaluating the ability to provide cost effective detectors for beam experiments. This chapter presents an analysis of potential physics sensitivity of CHIPS, which is here used to mean the whole experiment with a CHIPS detector (e.g. CHIPS-10) in the NuMI beam (or, if specified, LBNF).

The study is performed with the use of GLoBES, a software framework for simulating the sensitivity of long baseline accelerator neutrino oscillation experiments, using an abstract definition of the beam and the detector. The experiment definition for CHIPS detectors is based on the results of the preliminary reconstruction method described in Chapter 2. For comparison with other experiments, definitions provided by GLoBES authors are adopted.

3.1 The GLoBES framework

The GLoBES (General Long Baseline Experiment Simulator) package provides code to simulate event rates and the χ^2 value for a test point in the oscillation parameter phase space, given a parametrised description of the experiment performance, beam flux and the true values of the mixing parameters [114, 115]. A built-in minimiser handles the treatment of systematics and enables marginalisation over any of the oscillation parameters.

The framework consists of a definition language for providing experiment descriptions and a C library with the necessary functions. The library is used in a

stand-alone application written by the author in C++ for the purpose of this work. It is based on code used in the initial studies [93], but almost entirely rewritten, along with the macros for preparing input and displaying results.

The experiment definition is composed in the Abstract Experiment Definition Language (AEDL). The following components have to be provided: [116]

- Beam flux description: Arrays of fluxes of the three neutrino and antineutrino flavours as a function of true neutrino energy; the binning is not crucial, as they are later interpolated. Different beam modes, e.g. a neutrino and an antineutrino beam, can be provided as separate flux definitions.
- Baseline definition: Distance from the beam source to the detector, including matter densities (either as a single value or a table of distances and densities).
- Exposure variables: Detector target mass and total run time in years for each defined beam type.
- Neutrino cross sections: Provided in arrays for all neutrino and antineutrino flavours as a function of the logarithm of true neutrino energy. Each interaction type (e.g. CCQE) requires a new cross section definition.
- Energy binning specification: A list of energy bin sizes at each of two levels – pre-smearing (true energy), and post-smearing (reconstructed energy).
- Energy smearing matrices: Arrays transforming the true energy into reconstructed energy (and describing the energy resolution), with the number of rows and columns corresponding to the binning schemes. One array is typically defined for each detector process (e.g. CCQE ν_e).
- Channel definitions: Each channel corresponds to a neutrino flavour observed in a specific interaction (e.g. CCQE ν_e). The specific oscillation transition may also be specified. The channel description includes references to cross sections, beam modes and smearing matrices. It also provides multiplicative efficiencies and additive backgrounds, defined as a function of energy, to be applied at the pre- or post-smearing stages.
- Rule definitions: A rule describes a single analysis, e.g. the observation of electron neutrino appearance. The definition combines all relevant signal and

background channels with optional coefficients, as well as normalisation and energy calibration systematics.

A typical use of the GLoBES features proceeds as follows. A configuration file with the experiments and true oscillation parameters is loaded and parsed by the user code. The GLoBES library is initialised, and for each configuration (which can consist of a single experiment or a combination to simulate a joint analysis) the experiment definitions, including run times, are loaded. The true values of oscillation parameters are set, typically in a loop, e.g. over possible true values of δ_{CP}.

Using the abstract experiment definition and mixing parameters, GLoBES can calculate observed event rates for signal and background. The test values of oscillation parameters are set, along with their uncertainties, and GLoBES is called to compute the minimum $\Delta\chi^2$, marginalising over chosen parameters. The returned χ^2 values are stored in an output file for later processing and presentation.

Event rates are predicted for each channel, combining the neutrino flux information, oscillation probability, interaction cross section, target mass and total run time. The transition probability is calculated numerically using evolution operators after diagonalising the matter Hamiltonian (see Eq. 1.25) in every matter density layer, based on the baseline definition and oscillation parameter values. All these steps are performed in bins of neutrino true energy. Next, the pre-smearing efficiencies are applied, the true energy spectrum is converted to reconstructed energy with the smearing matrix, and finally the post-smearing efficiencies and backgrounds are applied. The rule rates are then calculated by summing over all relevant channels, keeping signal and background separate.

For a given test point in the mixing parameters phase space, the $\Delta\chi^2$ value can be computed,[1] by comparing the signal and background rates for true and test values. The total χ^2 is then summed over each rule in each experiment. In the simplest case with no systematic errors the χ^2comes only from Poissonian statistics of the event rates.

Systematics can be included via the "pull method" [117], where total normalisation and energy scale shift are treated as nuisance parameters and added to the χ^2 with penalties described by Gaussian uncertainties. The χ^2 is then minimised over those

[1]Since no statistical fluctuations are simulated, the χ^2 value at best-fit point is always zero and $\Delta\chi^2 = \chi^2$.

parameters. All systematics are assumed to be uncorrelated between rules and experiments.

GLoBES has also the ability to marginalise over chosen oscillation parameters, taking into account correlations between them. This is accomplished by adding Gaussian priors to the χ^2 for each marginalised parameter, and minimising over the parameter values. This step is performed after the systematics minimisation has been finished. The obtained minimum is local only, and care must be taken in order to avoid degeneracies.

The best fit points of the NuFIT 2016 global fit [59] (Fig. 1.6) are used as the true values of the oscillation parameters. The values of θ_{23} and Δm^2_{31} are defined separately for each mass hierarchy.[2] The same numbers are also used for the test values when marginalised over, with uncertainties taken as 1/6 of the listed 3σ range, to account for non-parabolic shapes of the $\Delta\chi^2$ curves. For θ_{23}, which has two local minima, even this approximation is not correct, and in the future a custom χ^2 prior will be implemented. The exception to the above is the value of δ_{CP}; the sensitivity is always tested for a full range of true values of δ_{CP}, and it is set completely free in the fits.

3.2 CHIPS experiment definition

Previous studies of CHIPS sensitivity with GLoBES were done with an experiment definition based on Super-Kamiokande performance, scaled to different mass and used with simulated NuMI fluxes [93]. For this work, a completely new definition was created, using the results of the preliminary reconstruction algorithm, and it is described in detail in Sec. 3.2.1. The old SK-based definition is presented and compared with the new one in Sec. 3.2.3.

3.2.1 Definition based on reconstruction results

The default abstract experiment definition for CHIPS was created based on findings from the developed MC simulation and the preliminary reconstruction (Sec. 2.4).

[2] NuFIT lists the value of Δm^2_{32} for the inverted hierarchy case, however GLoBES always takes as input $\Delta m^2_{31} = \Delta m^2_{32} + \Delta m^2_{21}$, which is accounted for in the code.

The results obtained with this definition can be considered a lower bound on the predicted performance, as the reconstruction algorithm is still in the early stages of development, and the CC ν_e selection is not yet fully optimised.

The rule and channel structure is implemented in the same way as in the old definition, with only two rules, one for ν_e appearance (used only in the neutrino beam mode), and one for $\overline{\nu}_e$ appearance (used in the antineutrino mode). The signal channels in the neutrino rule are CCQE ν_e and CC non-QE ν_e interactions, where the ν_e come from oscillated ν_μ. The background channels are: CCQE and CC non-QE ν_e with the intrinsic beam ν_e component,[3] CC ν_μ, and NC interactions for all flavours.

The antineutrino rule contains analogous channels with opposite signs describing the $\overline{\nu}_e$ appearance, but also ν_e channels representing the wrong-sign neutrino contamination. This contribution is more significant in the antineutrino mode because the antineutrino interaction cross-sections are smaller. The wrong-sign ν_e which come from oscillation are added to the antineutrino signal, not background, as they still depend on oscillation parameters – their inclusion just decreases the effects of CP violation and mass hierarchy.

The neutrino cross sections used are the standard ones provided with GLoBES, and are defined for CCQE, CC, CC non-QE and NC interactions. The cross sections for low energies come from Ref. [118], and for high energies from Ref. [119].

The unoscillated neutrino fluxes are calculated by reweighting the NuMI beam simulation, as described in Section 2.3.2. The energy spectra for muon and electron (anti)neutrinos in both neutrino and antineutrino modes are shown on Fig. 2.14 and 2.15. The negligible intrinsic tau neutrino component is left out.

The baseline geometry is based on a simple model with a constant matter density of $2.8\,\mathrm{g/cm^3}$ along the whole distance to the far detector, which is 708 km long. The default target mass of the detector is 10 kt, but it can be trivially scaled later when studying bigger or multiple CHIPS modules. Similarly, the beam run times are set to desired values at program runtime.

True neutrino energy is binned in 250 MeV wide bins from 0 to 5 GeV and in 5 GeV wide bins until 20 GeV. The reconstructed energy binning scheme is taken

[3] These are the electron neutrinos which are already present in the beam at the near detector and reach the far detector without oscillating into another flavour.

from the graphs of the final selection efficiency and uses 0.5 GeV wide bins from 0.5 to 5 GeV. The energy smearing matrices are extracted directly from the results of the reconstruction. A single-track electron hypothesis is fitted to events from CC ν_e, CC ν_μ and NC ν_μ Monte Carlo samples. For each event, the true neutrino energy and the reconstructed electron energy[4] (treated as a proxy for the reconstructed neutrino energy) are stored in a two-dimensional histogram, which is equivalent to the matrix. Then, for each bin of true energy, the projected histogram of reconstructed energy (corresponding to columns in the smearing matrix) is normalised, as advised by the GLoBES manual [116]. That way each neutrino is assigned some value of reconstructed energy, unless the whole column is empty. Figure 3.1 shows the resulting smearing matrices for the four relevant channel types: CCQE and CC non-QE ν_e, CC ν_μ and NC. Since the reconstruction has not yet been evaluated on the antineutrino samples, the same matrices are used for antineutrino channels. It can be seen that for quasi-elastic interactions the reconstructed energy is very closely related to the neutrino energy (since the outgoing charged lepton carries away most of the energy). For non-QE interactions, the relationship is weaker due to the hadronic products not being measured. In NC interactions, the outgoing neutrino carries away significant share of the total energy making the measurement very unreliable.

The selection efficiencies, defined as the ratio of events classified as CC ν_e to all events in a given sample, are provided as the final output of the particle identification procedure described in Section 2.4.4, and shown in Fig. 2.18. The efficiency for the CC non-QE ν_e sample is not displayed in the figure (taken from Ref. [1]), but was generated as well. The same binning scheme is preserved for the reconstructed energy in the experiment definition and the smearing matrices shown earlier. The final efficiency values, used for both neutrino and antineutrino events, are shown in Fig. 3.2. In this definition, the efficiencies are stored as a function of reconstructed energy and only applied once, after the energy smearing step.

Background from cosmic ray muons and tau neutrinos is not taken into account. This study was performed under the assumption that the CHIPS-10 module would have a veto volume surrounding the top and sides. High efficiency of the veto, combined with a low rate of muons with the appropriate direction and energy to

[4]This is correct even for the muon events, since the background channels contribute as events misidentified as electrons, and their reconstructed energy would be assumed to come from the electron fit.

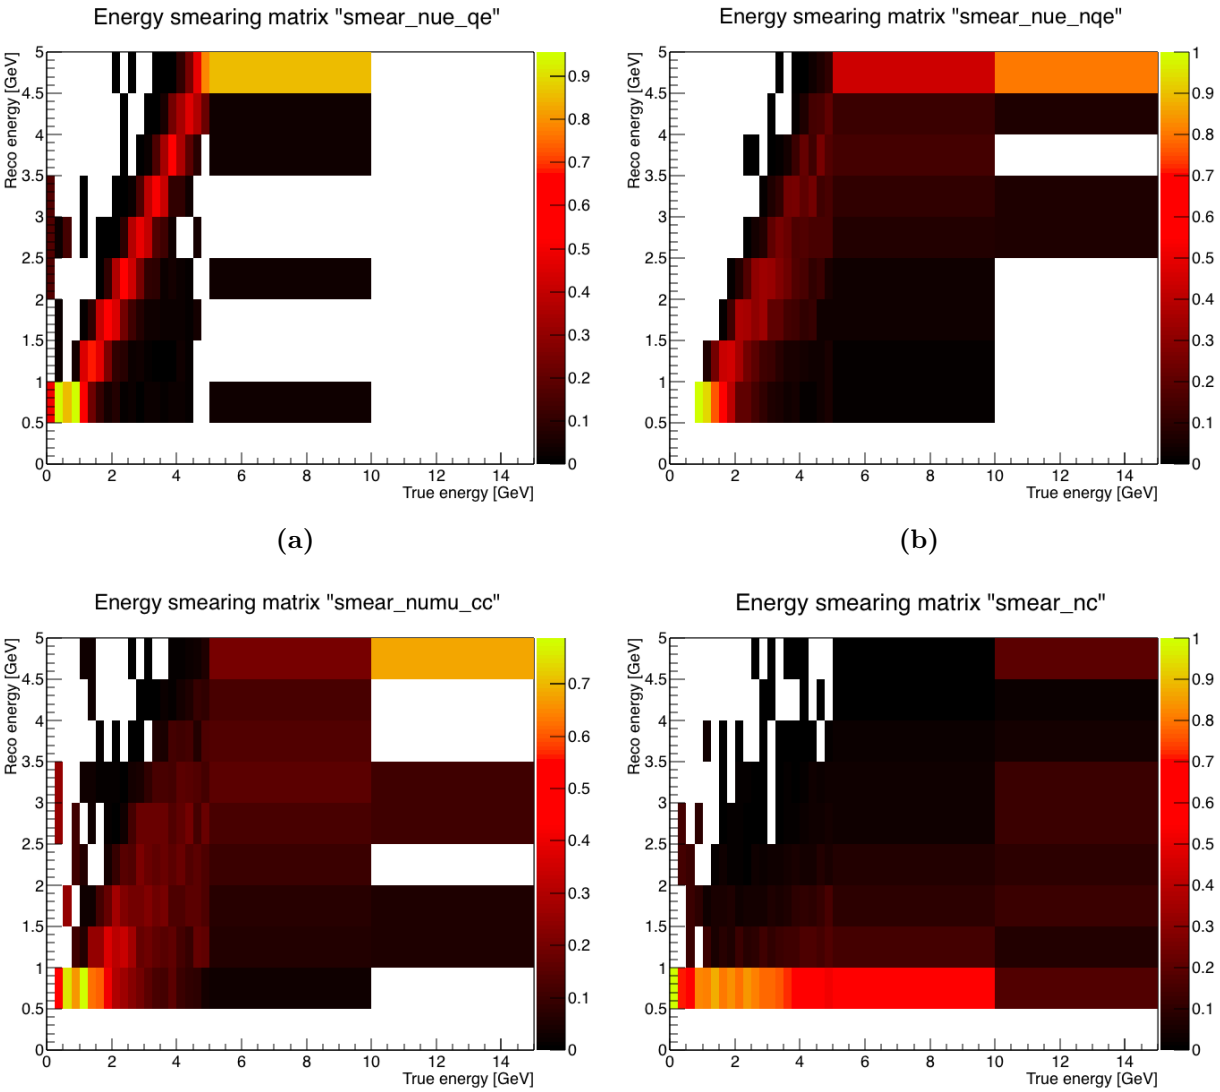

Figure 3.1: Energy smearing matrices for CHIPS, extracted from reconstruction outputs for the following event samples: (a) CCQE ν_e, (b) CC non-QE ν_e, (c) CC ν_μ, and (d) NC. The reconstructed energy scale is constrained by the binning of the selection efficiencies, while the true energy binning was kept at a more granular level and over a wider range.

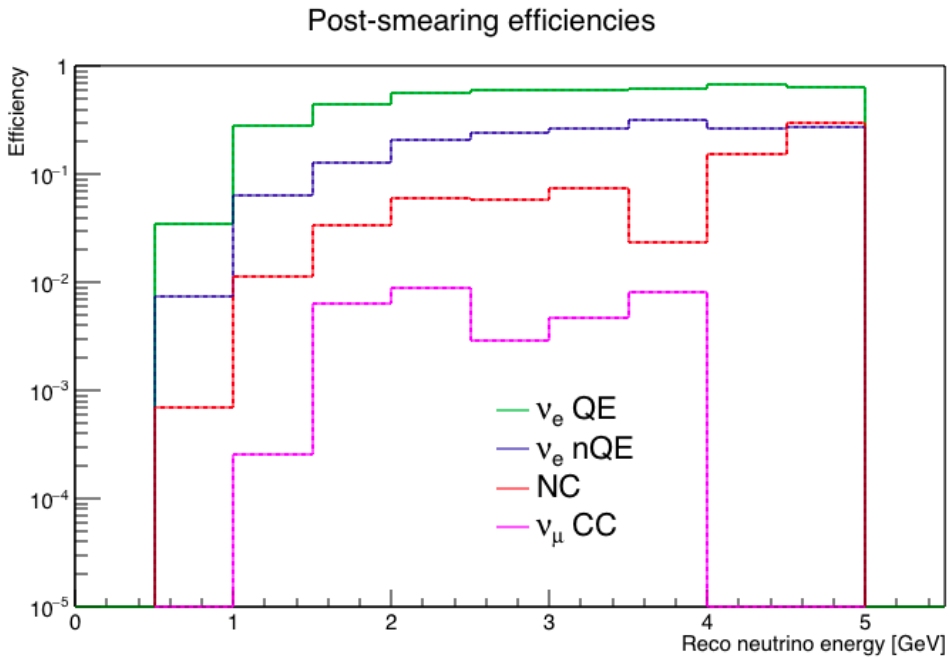

Figure 3.2: Selection efficiencies for different neutrino interaction samples. The values shown are identical to those in Fig. 2.18, except for the non-QE sample. The same efficiencies are used for antineutrino channels.

mimic a beam CC ν_μ interaction and the short beam spill time all indicated that the cosmic background is a subdominant component in the mis-identified muon sample. With the decision to keep only the top cap veto, these assumptions may need to be revisited.

The background due to CC ν_τ interactions was estimated to be small, around 0.35 events per kiloton per year in the neutrino mode [93]. Since total background event rates are on the order of 10 /kt/year, this contribution is a second order effect and is neglected.

To simulate the event rates it is also necessary to provide detector mass and running time. The default value of 10 kt of fiducial mass, corresponding to a full CHIPS-10 module, is used for most of the following results. For simplification, equal run time in both beam modes is assumed, which is not always realistic. At the NuMI beam, the schedule is optimised for the needs of NOvA, and CHIPS will just passively record data. However, NOvA plans to run alternately with neutrinos and antineutrinos to collect similar exposure with both modes. The beam is assumed to

Event type	ν mode		$\bar{\nu}$ mode	
	NH	IH	NH	IH
Appeared ν_e	77	64	23	27
Wrong-sign ν_e			2.8	2.4
Beam ν_e bkg	63	64	34	34
CC ν_μ bkg	50	49	22	22
NC bkg	96	96	54	54

Table 3.1: Total numbers of events for each channel observed in CHIPS-10 during 3 years of running in neutrino and antineutrino beam each, as predicted by GLoBES for $\delta_{CP} = -\pi/2$.

run at the design power of 700 kW, accumulating 6×10^{20} protons on target (POT) every year.[5]

Figure 3.3 shows signal and background event rates in CHIPS-10 for 3 years of running in neutrino and 3 years in antineutrino mode (denoted also as $3 + 3$ years), in case of two different mass hierarchies and four chosen values of δ_{CP}. It can be observed that by measuring both neutrino and antineutrino appearance, the degeneracy between different values of δ_{CP} and the two mass hierarchies can be lifted. Total event counts for $\delta_{CP} = -\pi/2$ and with background breakdown are shown in Table 3.1.

The final element of the experiment definition are systematic uncertainties for total normalisation and energy scale. The first is applied as a global rescaling of the event rates, and the second as a linear distortion of the energy spectrum [115]. The energy calibration error is taken as 2 % for both signal and background, which is possible to achieve with stopping cosmic muons [120]. The normalisation uncertainty for both signal and background is assumed to be 5 %. This level has been reached by T2K with the use of the near detector [50], and should be achievable at CHIPS, utilising cross section and flux measurements from MINERvA and the NOvA ND.

[5] This figure includes planned beam shutdowns, but otherwise nominal performance of the beamline. Detector dead time is not taken into account, but it is below 0.3 % (see Chapter 5).

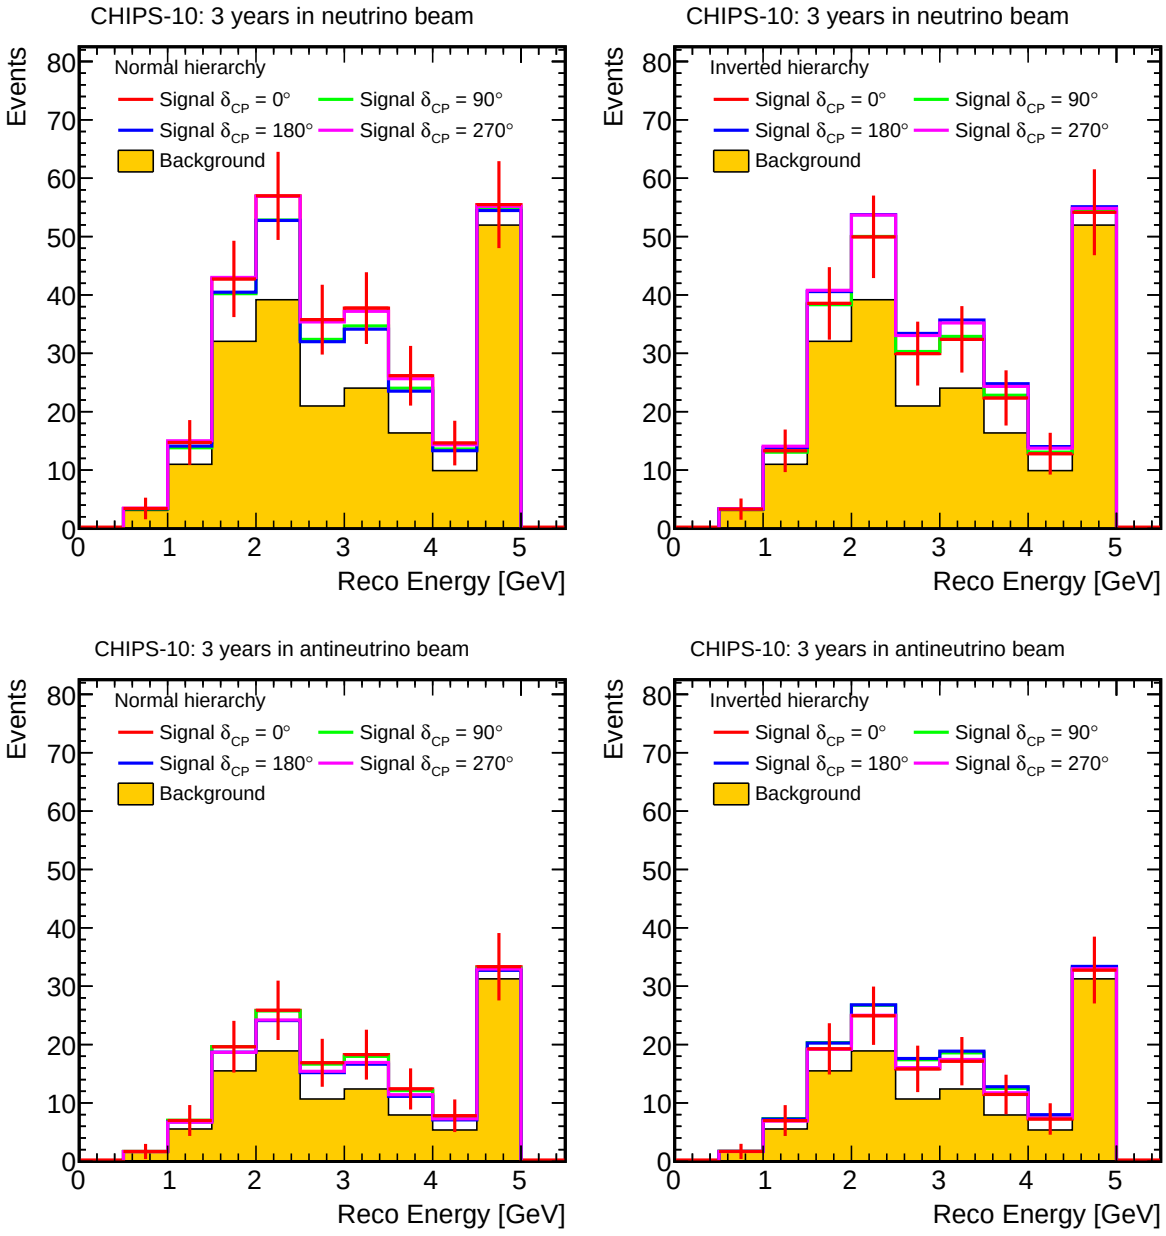

Figure 3.3: Event rates in CHIPS-10 for 3 years running in the neutrino beam mode (top) and 3 years in the antineutrino mode (bottom), in bins of reconstructed energy, in case of normal (left) and inverted hierarchy (right). The filled histograms represent background events. Signal rates, shown with solid lines for four different values of δ_{CP}, are stacked on top of the background. Statistical error bars are shown for one of the signal cases.

3.2.2 Sensitivity calculation

In order to study the physics reach of experiments with CHIPS detectors, and to compare it with other projects, some figures of merit must be chosen. In this work, these are: significance of mass hierarchy determination, significance of CP violation discovery and the 1σ resolution of the δ_{CP} measurement.

Mass hierarchy determination significance

The determination of mass hierarchy is equivalent to the exclusion of the wrong hierarchy. The computation is performed as follows in the main program code:

1. Set true values of oscillation parameters, including δ_{CP} and the choice of mass hierarchy; compute the event rates.
2. Set the test parameter values to be identical, except for Δm^2_{31} and θ_{23} - set those to correspond to the wrong hierarchy.
3. Call GLoBES to calculate the χ^2 of test parameters without any marginalisation in a loop over test values of δ_{CP}; keep the lowest value and the corresponding $\delta^{\mathrm{min}}_{\mathrm{CP}}$.
4. Set the $\delta^{\mathrm{min}}_{\mathrm{CP}}$ from previous step as the test value of δ_{CP}; call GLoBES function to calculate the χ^2 of the test parameters, marginalising over all of them.
5. Repeat previous step, but with $\delta^{\mathrm{min}}_{\mathrm{CP}} + \pi$ as the test value of δ_{CP}.
6. Take the lower χ^2 value of the previous steps and store it for the true δ_{CP} and mass hierarchy.

These steps are repeated for both mass hierarchies and 100 bins of true values of δ_{CP}, from $-\pi$ to π. Steps 3 to 5 ensure that the final result is a global minimum and not just a local one. The significance is expressed as $\sqrt{\Delta\chi^2}$, which corresponds to number of standard deviations σ for a Gaussian-distributed variable with one degree of freedom.

Figure 3.4 shows the mass hierarchy determination sensitivity at CHIPS-10, for three different exposures: 6 years with the neutrino beam only, and 6 years and 12 years with neutrino and antineutrino run time split in half. It is evident that

running in both beam modes gives superior sensitivity, as it enables the unfolding of the matter effects and CP violation, even though the total event rates are lower.

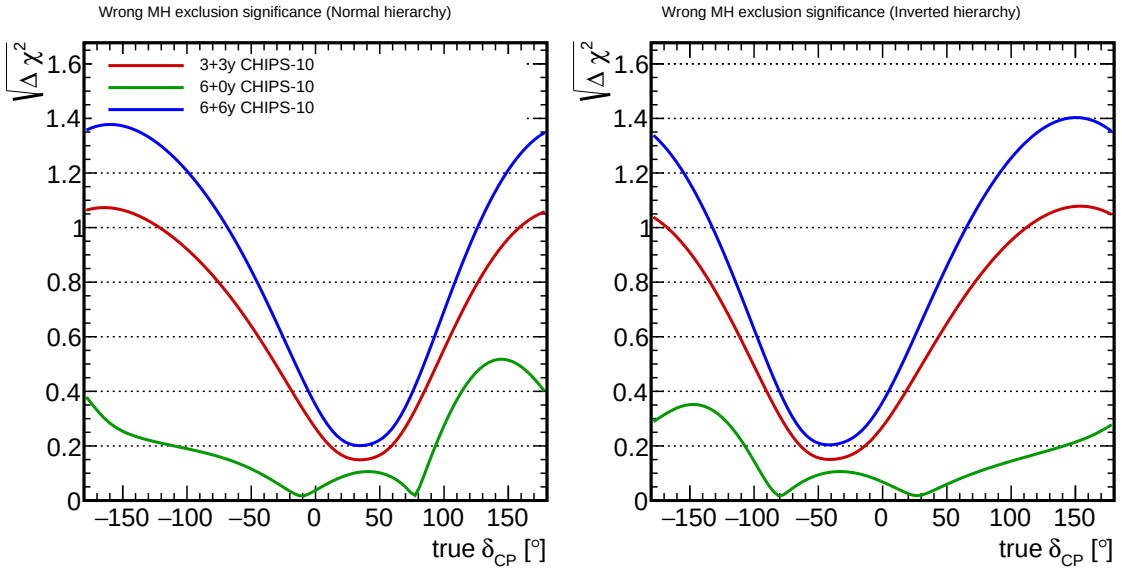

Figure 3.4: Mass hierarchy determination significance at CHIPS-10 as a function of the true value of δ_{CP}, for normal (left) and inverted hierarchy (right). Solid lines correspond to three run scenarios of CHIPS-10: 3 years in neutrino mode and 3 years in antineutrino mode (red), 6 years in neutrino mode only (green) and $6 + 6$ years (blue).

Figure 3.5 shows the sensitivity for three ranges of possible δ_{CP} values as a function of total CHIPS exposure, in the normal hierarchy. For example, the line labelled 50 % shows the minimal mass hierarchy determination significance, if the true δ_{CP} is in the most favourable 50 % of the phase space. The exposure assumes equal neutrino and antineutrino beam time, and could be increased either by running longer, or by adding more modules.

For the most favourable values of δ_{CP}, close to the best-fit point at $-99°$, CHIPS-10 alone can realistically reach a 1σ significance of mass hierarchy determination. Adding exposure by deploying additional modules can increase that further, with a CHIPS-100 running for $3 + 3$ years reaching a 2σ level.

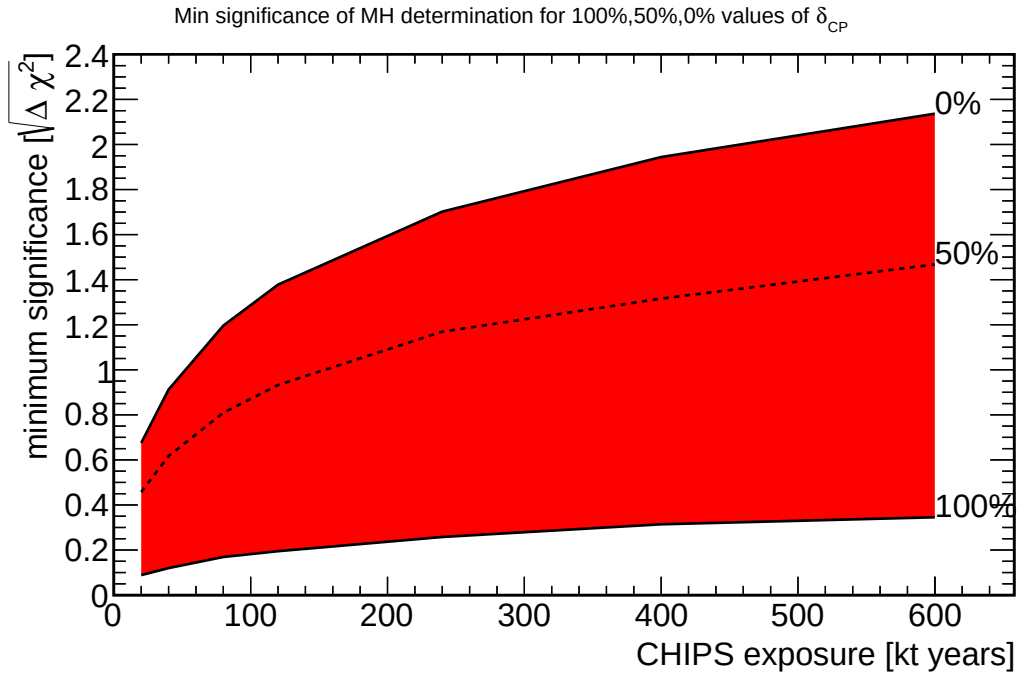

Figure 3.5: Minimum mass hierarchy determination significance as a function of CHIPS exposure for the normal hierarchy and various ranges of δ_{CP} values. The upper line represents the sensitivity if δ_{CP} has the single most favourable value, and the lower one holds for any value of δ_{CP}.

CP violation confirmation significance

Similar to the case of the mass hierarchy, the discovery of CP violation is equivalent to excluding CP conservation, i.e. δ_{CP} values of 0 and π. The following steps are executed:

1. Set true values of oscillation parameters, including δ_{CP} and the choice of mass hierarchy; compute the event rates.
2. Set the test parameter values to be identical to the true ones.
3. Set the test value of δ_{CP} to 0; call GLoBES to calculate the χ^2, marginalising over all parameters except δ_{CP}.
4. Repeat the previous step, but with test value of δ_{CP} set to π; keep the lower value of χ^2.

5. Repeat the two previous steps, but with the test values of Δm^2_{31} and θ_{23} corresponding to the wrong hierarchy.

6. Take the lowest resulting χ^2 value and store it for the true δ_{CP} and mass hierarchy.

The procedure is repeated with both mass hierarchies set as correct and in bins of true δ_{CP}. Step 5 may be skipped if the correct hierarchy is assumed to be already known, however the results presented here do not obey that constraint.

Figure 3.6 shows the CP violation discovery sensitivity at CHIPS-10, for three different exposures. The sensitivity is not symmetric around $\delta_{\mathrm{CP}} = 0$, due to degeneracies from unknown mass hierarchy. If the hierarchy is known from other experiments, the curves become almost exactly periodic. It can also be seen that with only one beam mode (e.g. no antineutrino run), the degeneracies have an even stronger effect, causing the irregular shape of the sensitivity curve. Using both beam modes is therefore crucial for the determination of CP violation.

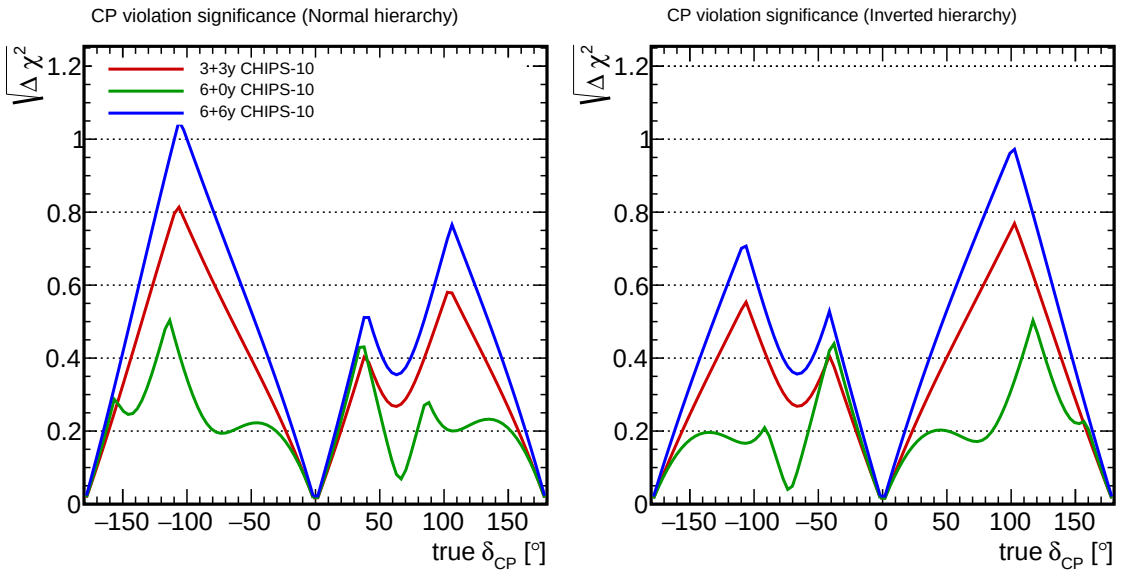

Figure 3.6: CP violation discovery significance at CHIPS-10 as a function of the true value of δ_{CP}, for normal (left) and inverted hierarchy (right). Solid lines correspond to three run scenarios of CHIPS-10: 3 years in neutrino mode and 3 years in antineutrino mode (red), 6 years in neutrino mode only (green) and $6 + 6$ years (blue).

Figure 3.7 shows the sensitivity for three ranges of possible δ_{CP} values as function of total CHIPS exposure. As opposed to the diagram for mass hierarchy, 25 % and 75 % ranges are used, instead of 0 % and 100 %, as in Ref. [85]. There is no guarantee to discover CP violation regardless of the true value of δ_{CP}, because for $\delta_{\mathrm{CP}} = 0$ or π it is not violated at all.

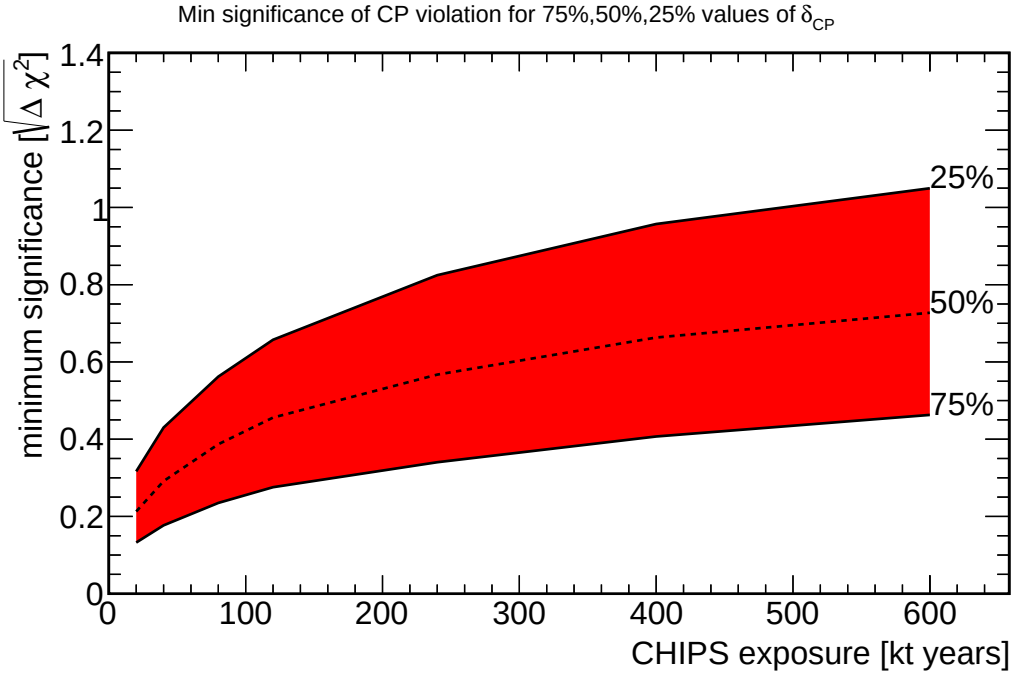

Figure 3.7: Minimum CP violation discovery significance as a function of CHIPS exposure for the normal hierarchy and various ranges of δ_{CP} values. For example, the upper line represents the sensitivity if δ_{CP} is within 25 % of the most favourable values. Since for δ_{CP} close to 0 or π, CP violation can be very small, a significant discovery in all cases cannot be guaranteed.

Even for the most favourable value of true δ_{CP}, CHIPS-10 can barely reach 1σ CP violation significance after 12 years of operation. The sensitivity improves with additional exposure, e.g. by adding more chips modules, however not dramatically.

Resolution of δ_{CP}

In addition to determining if the CP symmetry is violated, i.e. whether δ_{CP} is different from 0 or π, it is important to measure the actual value of the phase. The 1σ resolution of that measurement is the distance from the best fit point to where the $\Delta\chi^2 = 1$ (or $\chi^2 = \chi^2_{\mathrm{min}} + 1$). It is calculated using an approximation.

Assuming that the χ^2 curve is a symmetric parabola with value $\chi^2_{\rm min}$ at $\delta_{\rm min}$

$$\chi^2(\delta) = a(\delta - \delta_{\rm min})^2 + \chi^2_{\rm min} \tag{3.1}$$

and $\chi^2(\delta_+) = \chi^2_{\rm min} + 1$, then

$$\delta_+ - \delta_{\rm min} = \frac{1}{\sqrt{a}}. \tag{3.2}$$

The coefficient a can be calculated from the second derivative of χ^2

$$2a = \frac{d^2\chi^2}{d\delta^2} \simeq \frac{\chi^2(\delta + d\delta) + \chi^2(\delta - d\delta) - 2\chi^2(\delta)}{d\delta^2} \tag{3.3}$$

$$\simeq \chi^2(\delta + 1) + \chi^2(\delta - 1) - 2\chi^2(\delta). \tag{3.4}$$

Then the resolution

$$\delta_+ - \delta_{\rm min} = \sqrt{\frac{2 \times 1\,^{\circ 2}}{\chi^2(\delta_{\rm min} + 1^\circ) + \chi^2(\delta_{\rm min} - 1^\circ) - 2\chi^2(\delta_{\rm min})}}. \tag{3.5}$$

In the code, it is therefore calculated as follows:

1. Set true values of oscillation parameters, including $\delta_{\rm CP}$ and the choice of mass hierarchy; compute the event rates.
2. Set the test parameter values to be identical to the true ones.
3. Call GLoBES to calculate the χ^2, marginalising over all parameters; the result is $\chi^2_{\rm min}$ with the corresponding $\delta_{\rm min}$.
4. Set the test value $\delta_{\rm CP} = \delta_{\rm min} + 1^\circ$; call GLoBES to calculate the χ^2, marginalising over all parameters except $\delta_{\rm CP}$; the result is χ^2_+.
5. Repeat previous step, but with the test value $\delta_{\rm CP} = \delta_{\rm min} - 1^\circ$; the result is χ^2_-.
6. Calculate the resolution $\delta_{\rm res} = \sqrt{\dfrac{2}{\chi^2_+ + \chi^2_- - 2\chi^2_{\rm min}}}$, and store it for the true $\delta_{\rm CP}$ and mass hierarchy.

There are two main consequences of using an approximation. First, the actual shape of the χ^2 curve is not symmetric, and the 1σ uncertainty should be different

on both sides of the central value. However, the approximate result can be seen as an average resolution, which is easier to display as a single value.

Second, the χ^2 shape is periodic in δ_{CP} and at some point the parabolic approximation must break down. Indeed, for results above 100°, the resolution is underestimated compared to the actual value. However, the relation is still monotonic and the returned value can be used for strict comparison.

Figure 3.8 shows δ_{CP} resolution as a function of its true value for both hierarchies, and three different run scenarios, and Fig. 3.9 shows the resolution as a function of CHIPS exposure in the normal hierarchy, for a range of true δ_{CP} values. Differently from the previous significance plots, here a lower resolution means a better measurement. When running in both neutrino and antineutrino beam modes, the dependence of measurement resolution on the actual value of δ_{CP} in CHIPS is relatively flat, creating the narrow band in Fig. 3.9. The resolution also improves rapidly with exposure until approximately 300 kt years.

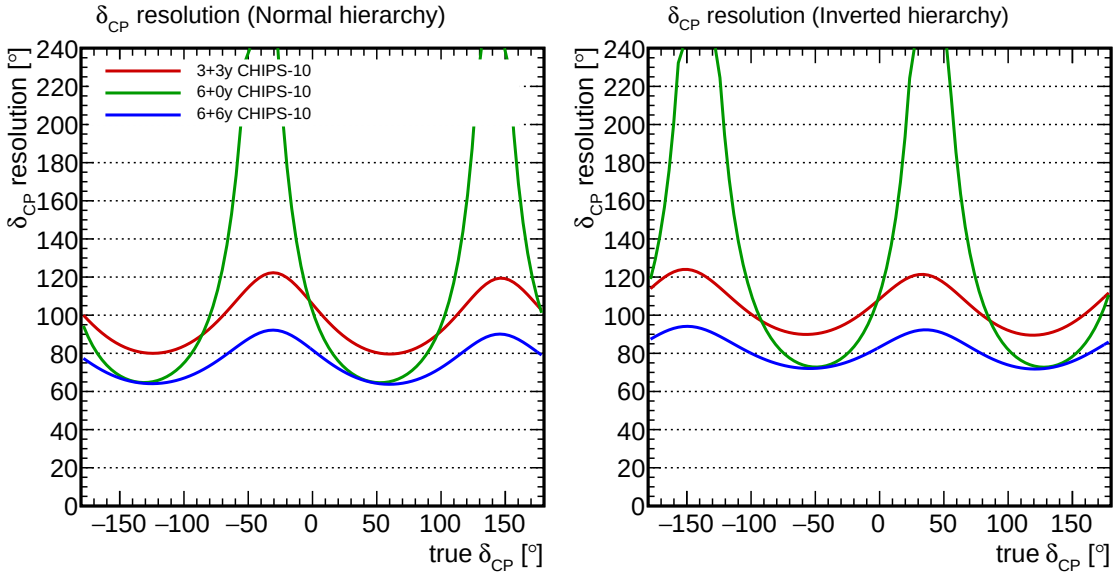

Figure 3.8: Resolution of δ_{CP} measurement at CHIPS-10 as a function of the true value of δ_{CP}, for normal (left) and inverted hierarchy (right). Solid lines correspond to three run scenarios of CHIPS-10: 3 years in neutrino mode and 3 years in antineutrino mode (red), 6 years in neutrino mode only (green) and $6 + 6$ years (blue).

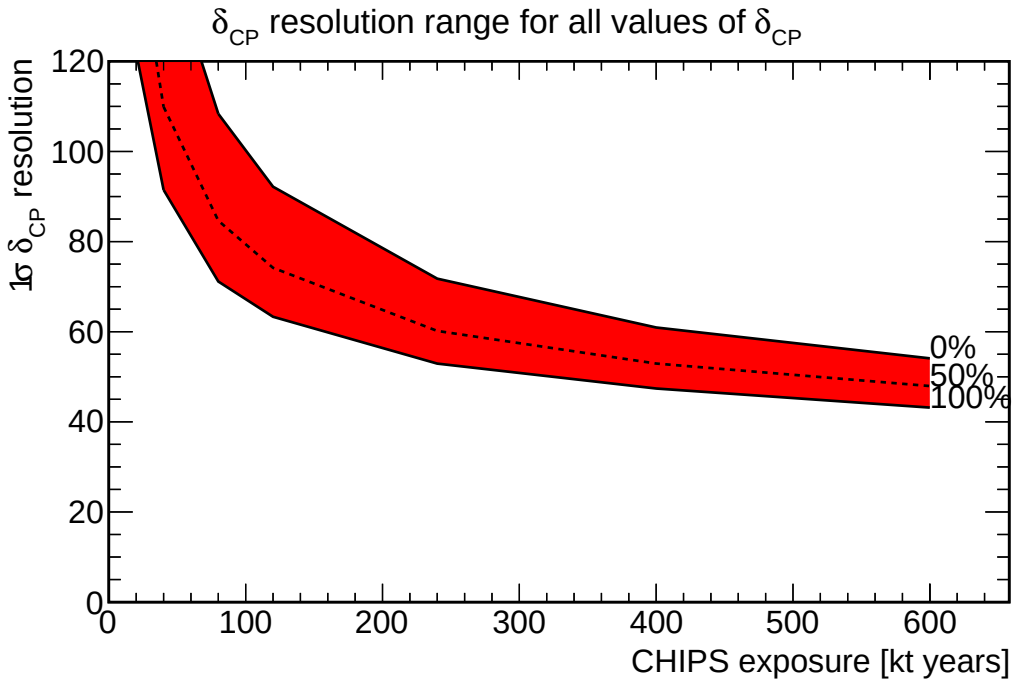

Figure 3.9: Worst resolution of δ_{CP} measurement as a function of CHIPS exposure for the normal hierarchy and various ranges of true δ_{CP} values. The lower line represents the 1σ resolution if δ_{CP} has the most favourable value, and the upper one holds for any value of δ_{CP}.

3.2.3 SK-based water Cherenkov definition

The first AEDL definition used for CHIPS was built from files received in private correspondence from G. Zeller, evolved from the work presented in Ref. [121]. They define a water Cherenkov detector template, based on the performance of Super-Kamiokande on the atmospheric neutrino sample, with added beam backgrounds scaled from T2K. The only change made to adapt it for CHIPS was the use of the NuMI fluxes, and a modified target mass. The outcome of a study performed with this definition was presented in Ref. [93]. Those results can be seen as an upper bound on CHIPS performance, since the definition is for a detector with an optimised particle selection and virtually no cosmic background; although the improved reconstruction could reduce the shortcomings.

As in the case of the default CHIPS definition, only two GLoBES rules are used, describing electron (anti-)neutrino appearance, with the signal and background

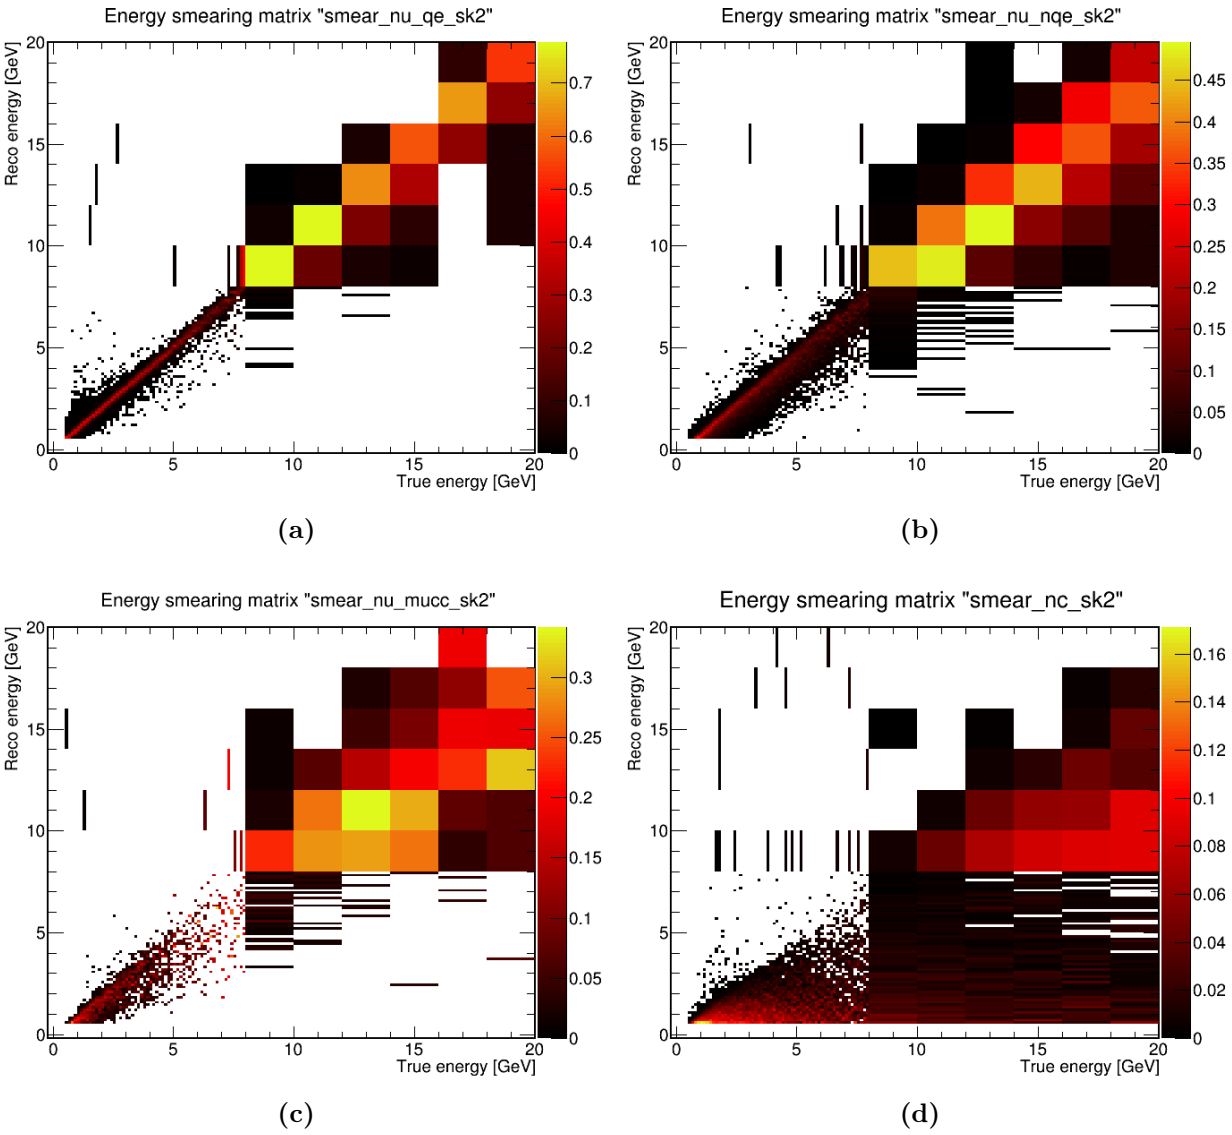

Figure 3.10: Energy smearing matrices for the SK-based definition of CHIPS, for neutrino samples: (a) CCQE ν_e, (b) CC non-QE ν_e, (c) CC ν_μ, and (d) NC. Both true and reconstructed energy scales are cut off at 20 GeV for clarity.

channels defined for the same interaction samples: CCQE and CC non-QE for ν_e (oscillated signal and intrinsic background), CC ν_μ and NC.

True and reconstructed energy are both finely binned in 125 MeV wide bins from 0.5 to 8 GeV and 2 GeV wide bins to 60 GeV. The energy smearing matrices are shown in Fig. 3.10 for neutrinos and in Fig. 3.11 for antineutrinos (the same matrix is used for both signs in the NC sample).

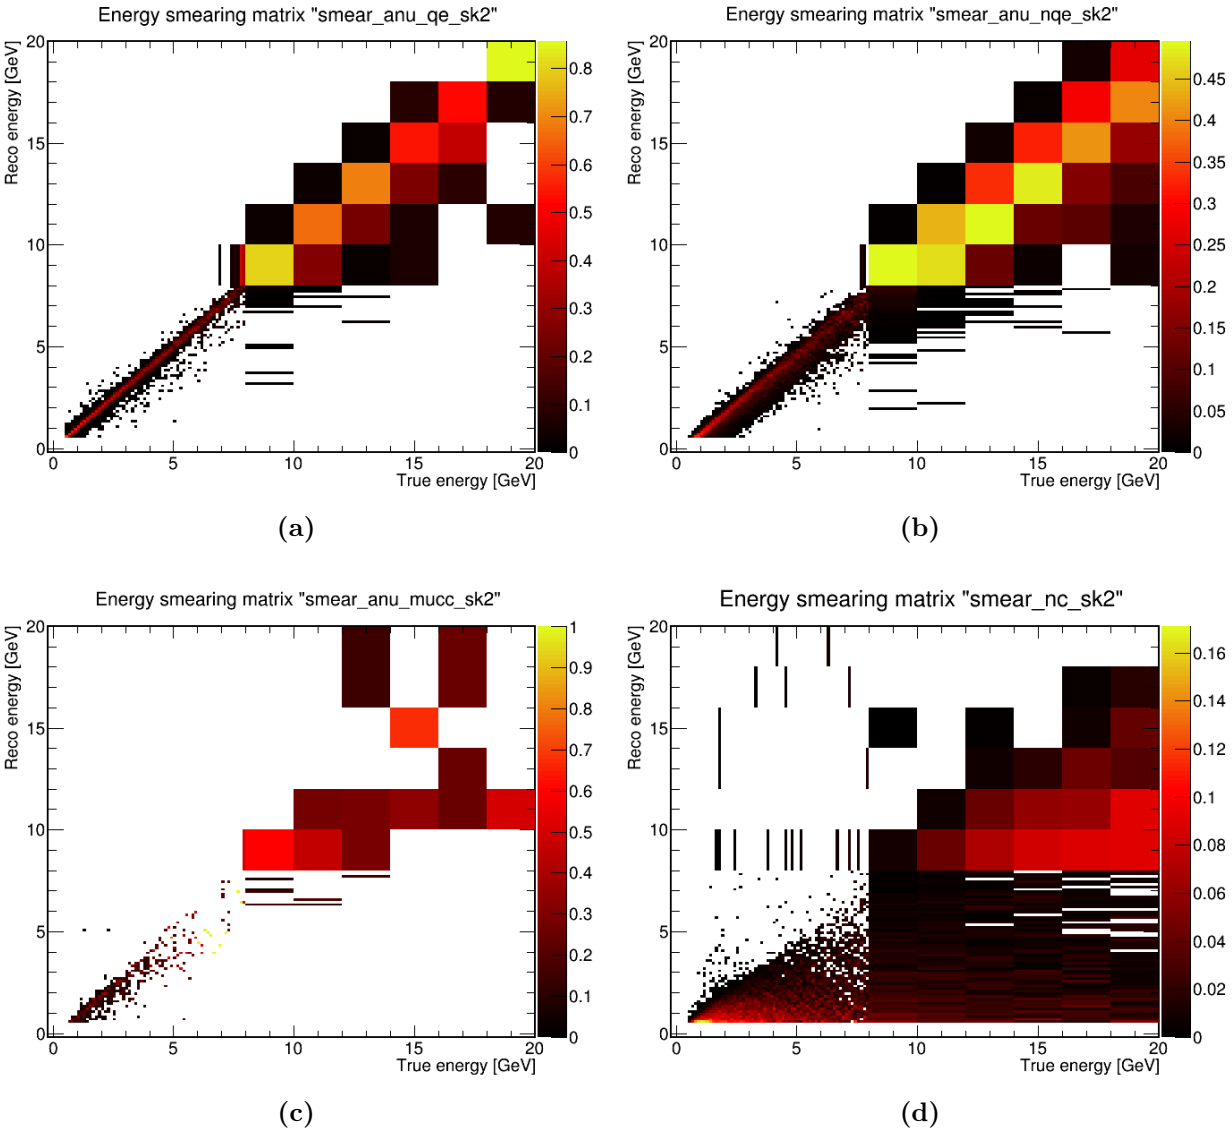

Figure 3.11: Energy smearing matrices for the SK-based definition of CHIPS, for antineutrino samples: (a) CCQE $\overline{\nu}_e$, (b) CC non-QE $\overline{\nu}_e$, (c) CC $\overline{\nu}_\mu$, and (d) NC. The NC smearing matrix for neutrinos and antineutrinos is identical. Both true and reconstructed energy scales are cut off at 20 GeV for clarity.

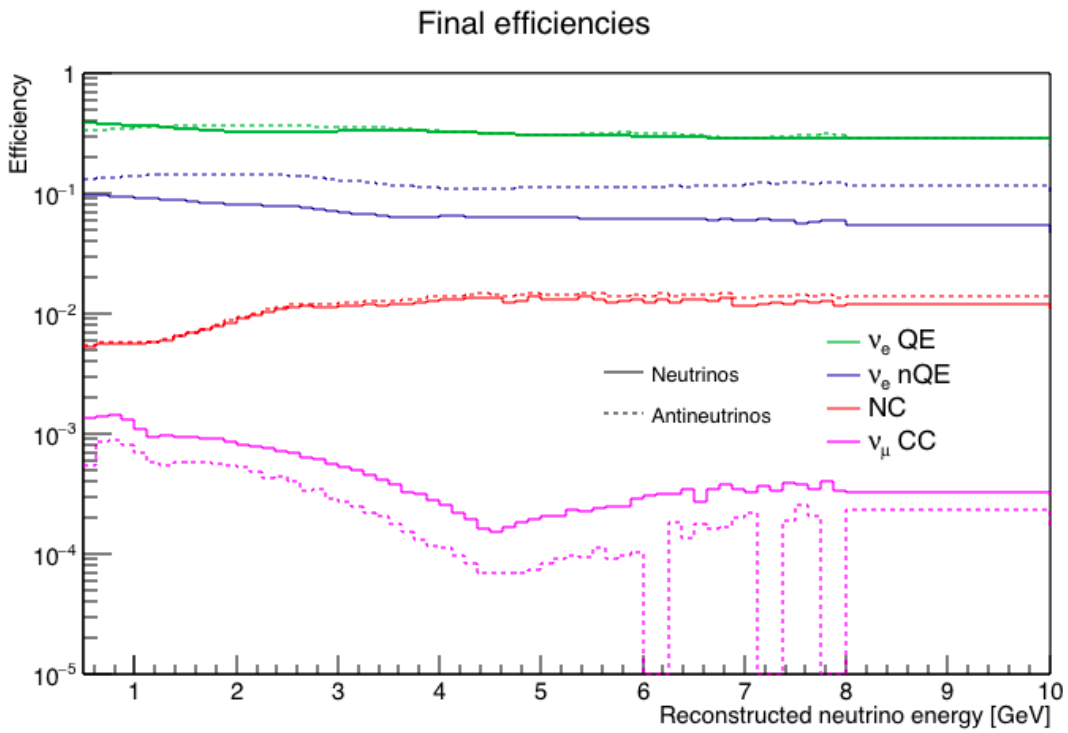

Figure 3.12: Final selection efficiencies in the SK-based definition for different neutrino (solid lines) and antineutrino samples (dotted lines), as a function of reconstructed neutrino energy.

The selection efficiencies are applied in two steps. Pre-smearing efficiencies describe the event pre-selection and are expressed in true neutrino energy, and post-smearing efficiencies correspond to cuts on a maximum likelihood selection. Figure 3.12 displays the final efficiencies for neutrinos and antineutrinos as a function of reconstructed energy. In order to calculate them, the pre-smearing event rates are estimated with GLoBES. The raw rates and rates with pre-smearing efficiencies applied are multiplied by the energy smearing matrices, and their ratio defines the pre-smearing efficiencies as a function of reconstructed energy. Those are then multiplied by the post-smearing efficiencies, giving the final result.

Predicted event rates for $3 + 3$ years of running at the Wentworth Pit are shown in Fig. 3.13, for both hierarchies and four values of δ_{CP}. The binning is reduced to increase visibility and decrease the error bars, and only the relevant energy range is displayed. The event counts are smaller than for the new experiment definition, because of the smaller efficiencies. However, the purity of signal samples is higher, leading to better sensitivity overall.

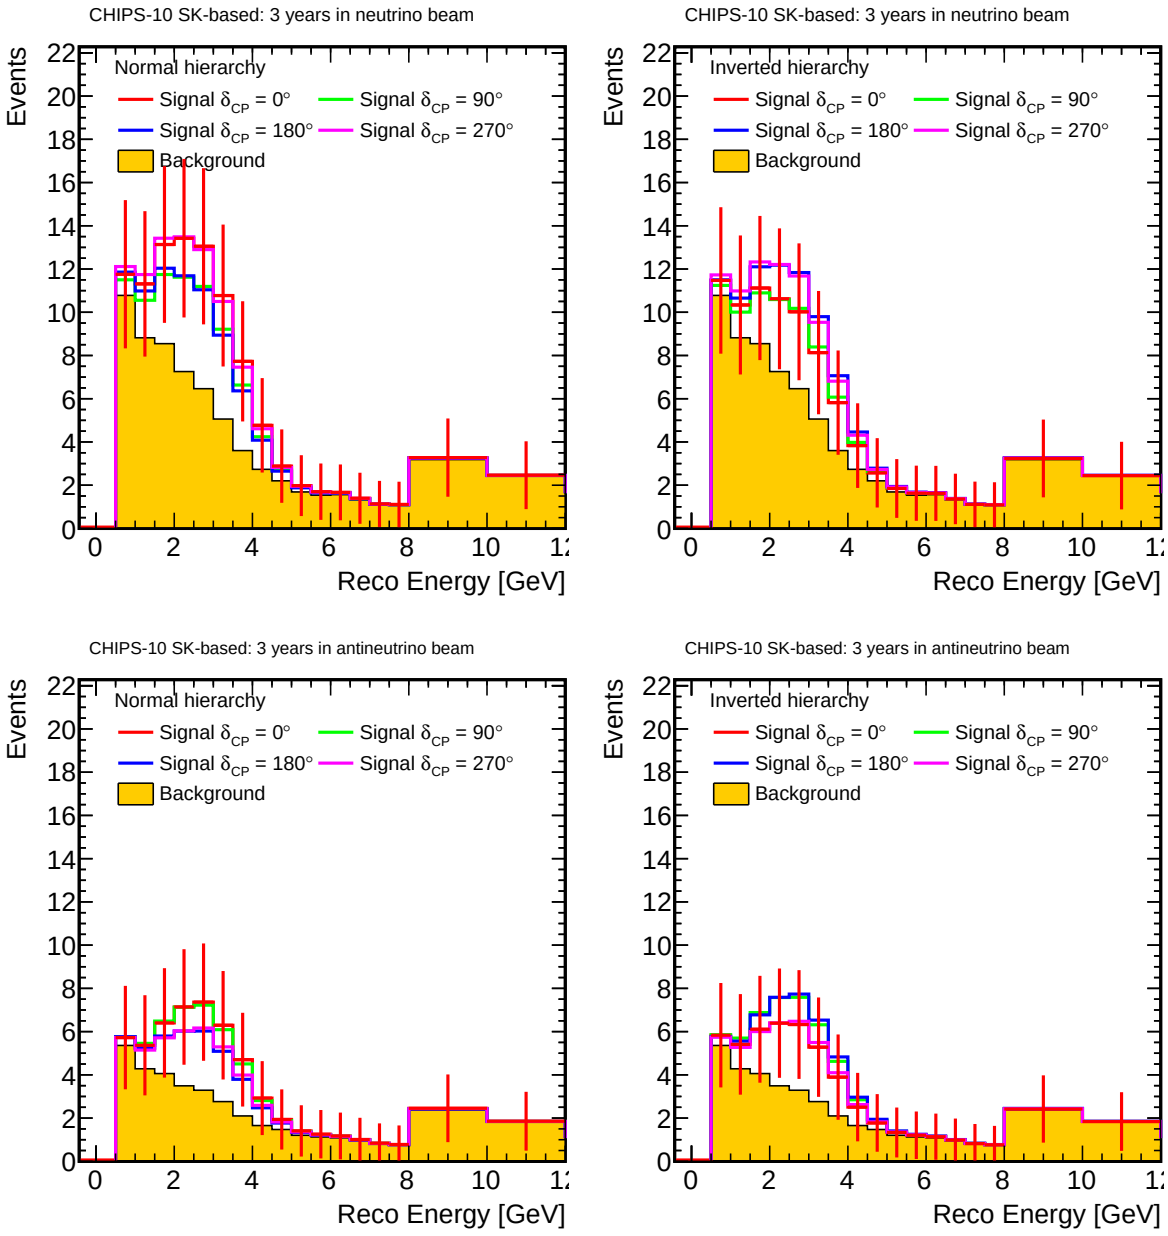

Figure 3.13: Event rates in CHIPS-10 with the SK-based definition for 3 years running in the neutrino beam mode (top) and 3 years in the antineutrino mode (bottom), in case of normal (left) and inverted hierarchy (right). The filled histograms represent background events. Signal rates, shown with solid lines for four different values of δ_{CP}, are stacked on top of the background. Statistical error bars are shown for one of the signal cases.

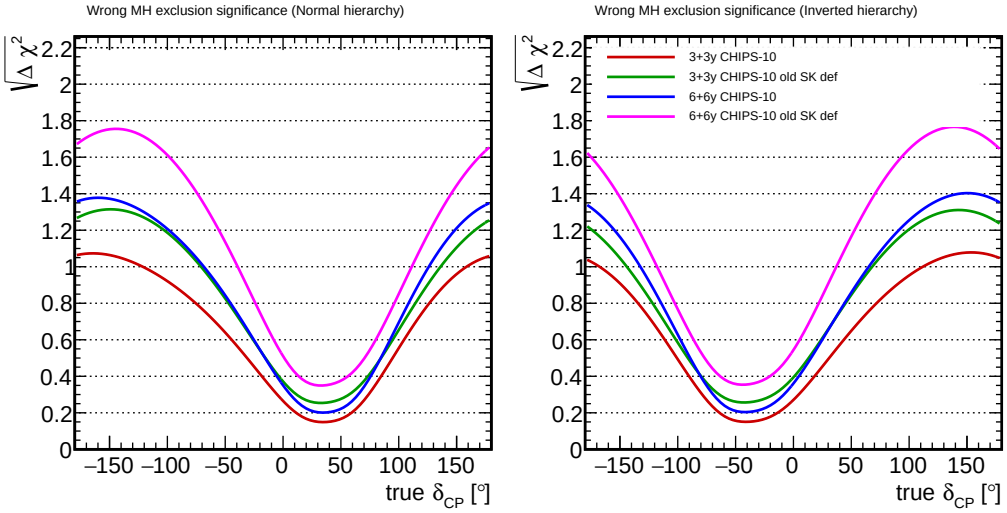

Figure 3.14: Mass hierarchy determination significance at CHIPS-10 as a function of the true value of δ_{CP}, in case of normal (left) and inverted hierarchy (right). Solid lines correspond to two run scenarios with two experiment definitions: $3+3$ years with the default definition (red) and the old one (green), and $6+6$ years with new (blue) and old definition (pink).

The systematic uncertainty of the signal normalisation is 1 % and 5 % for background. There is also a 0.01 % error on the energy scale for both signal and background. This corresponds to practically no calibration error at all, although the final results do not differ significantly if 1 % uncertainty is used instead.

Figures 3.14 to 3.16 show sensitivity of CHIPS-10 with the old SK-based definition and the new one, for two different exposure scenarios. It can be seen that the old definition achieves similar sensitivity as the default one at half the exposure. This suggests that there is significant room for improvement, even if the exact sensitivity level might not be actually reached.

The figures are not comparable with those presented in Ref. [93] due to changes in the code. The χ^2 values are now marginalised over the oscillation parameters, θ_{23} value depends on the mass hierarchy, and the true values of the mixing parameters were updated to match the current global fit results.

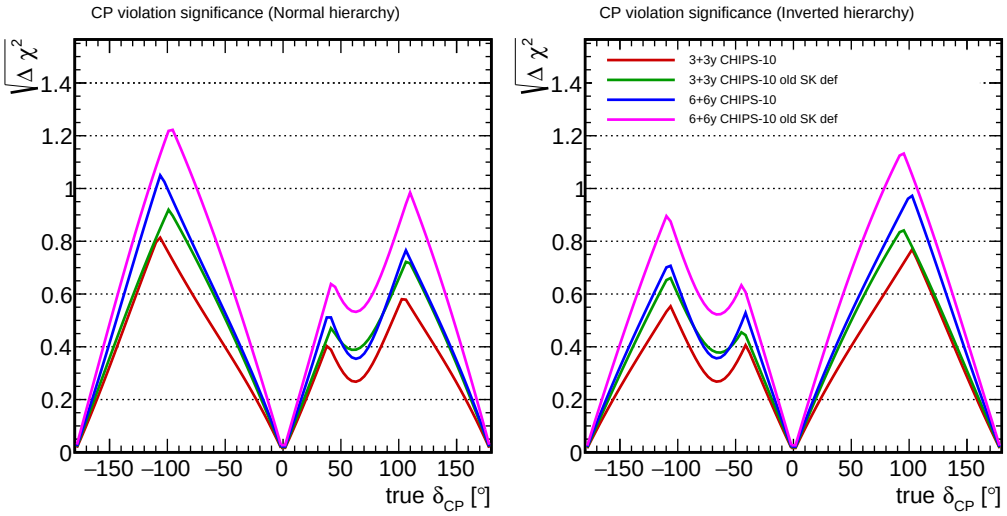

Figure 3.15: CP violation discovery significance as at CHIPS-10 as a function of the true value of δ_{CP}, in case of normal (left) and inverted hierarchy (right). Solid lines correspond to two run scenarios with two experiment definitions: $3+3$ years with the default definition (red) and the old one (green), and $6+6$ years with new (blue) and old definition (pink).

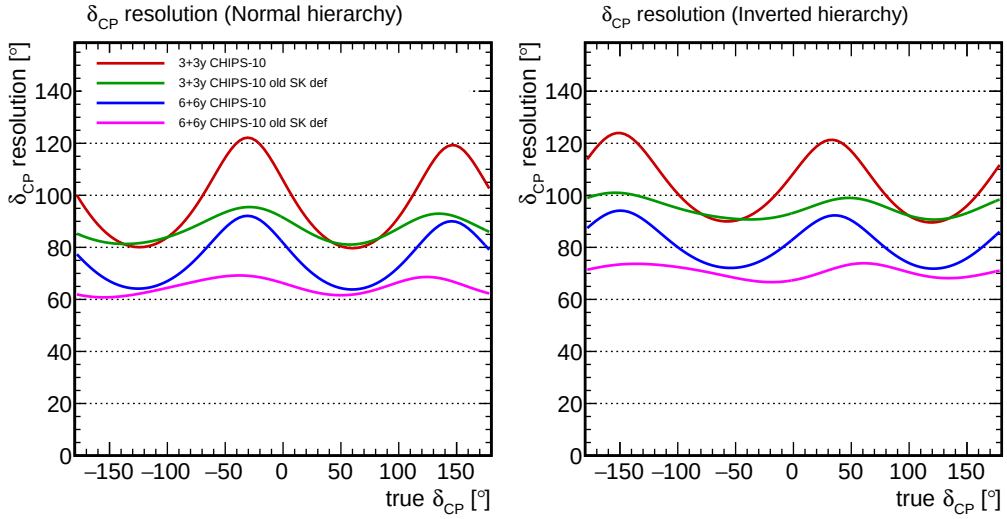

Figure 3.16: Resolution of δ_{CP} measurement at CHIPS-10 as a function of the true value of δ_{CP}, in case of normal (left) and inverted hierarchy (right). Solid lines correspond to two run scenarios with two experiment definitions: $3+3$ years with the default definition (red) and the old one (green), and $6+6$ years with new (blue) and old definition (pink).

3.3 Comparison with current experiments

The motivation behind CHIPS is to disrupt the current environment of neutrino oscillation experiments by creating a more cost-effective detector technology. Therefore, the physics potential of experiments with CHIPS detectors has to be compared to the existing and planned competitors. In this section, CHIPS-10 is contrasted with already running T2K and NOvA experiments, and their combination. Since CHIPS-10 is only a proof-of-principle prototype, a potential expansion to CHIPS-100 is also considered.

3.3.1 Experiment definitions for NOvA and T2K

Both NOvA and T2K experiments have been accumulating data with initial beam powers lower than the design reference or predicted future values. By May 2016, NOvA has collected data corresponding to 6×10^{20} POT full-mass equivalent in neutrino mode, and T2K has accumulated 7.5×10^{20} POT in neutrino and antineutrino mode each [49, 50]. In the case of both experiments, this is equal to approximately one year of planned beam exposure. To simplify the GLoBES definitions setup, the experiments are assumed to gather $1.5 + 1.5$ years of beam time each by the middle of 2018 and run at constant beam power onwards, with equal share of neutrino and antineutrino modes. The same time is used as the deployment and data-taking start of the full CHIPS-10 module.

The AEDL files used for simulating the NOvA experiment are the ones supplied with the GLoBES package, and are based on information provided in Refs. [122–124]. Figure 3.17 shows the official sensitivity predictions, including mass hierarchy determination and CP violation discovery significance, for normal hierarchy and $\delta_{\mathrm{CP}} = -\pi/2$ [125]. Figure 3.18 shows these sensitivities recreated with GLoBES using the default NOvA experiment definition, with $1.5 + 1.5$ years of exposure accumulated by 2018. The two predictions agree quite well, validating the simulation setup.

The T2K definition is an old one supplied with GLoBES before 2007, and based on Ref. [127–129]. A newer standard definition specifies the rules and channels with additive backgrounds and without additional modifications it is only valid

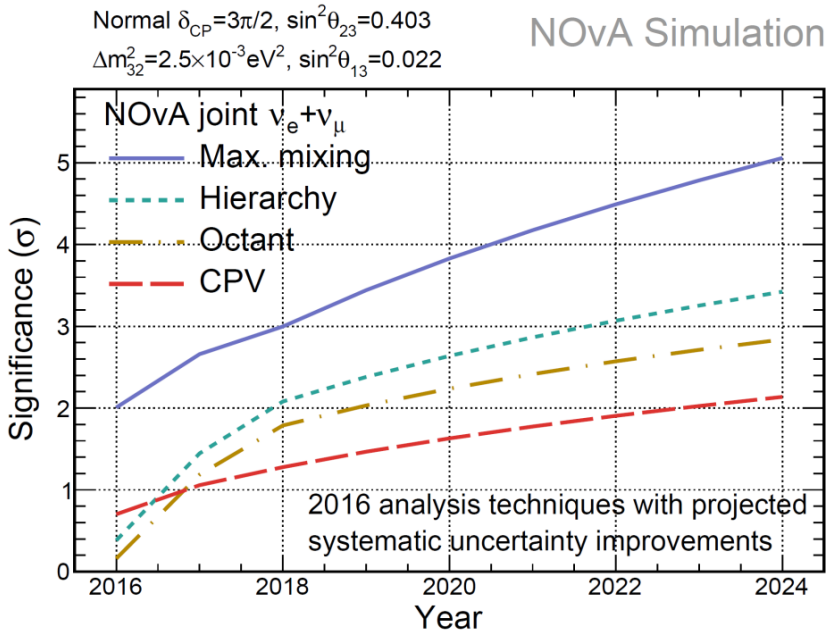

Figure 3.17: Predicted physics sensitivity at NOvA as a function of time, for $\delta_{\mathrm{CP}} = -\pi/2$ and normal hierarchy. The lines show significance of: maximal mixing ($\theta_{23} = \pi/4$) exclusion (solid blue), mass hierarchy (dotted cyan) and θ_{23} octant determination (dash-dotted brown), and CP violation discovery (dashed red). Figure taken from Ref. [125].

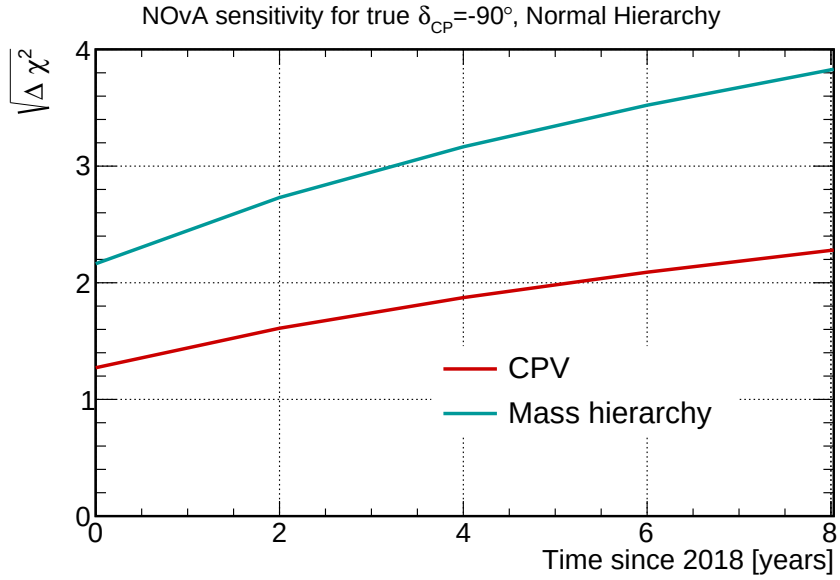

Figure 3.18: NOvA sensitivity as predicted with GLoBES, for normal hierarchy and $\delta_{\mathrm{CP}} = -\pi/2$. Solid lines show significance of: CP violation discovery (red) and mass hierarchy determination (cyan). Compare with Fig. 3.17. Note that here the X axis starts in 2018 and ends in 2026.

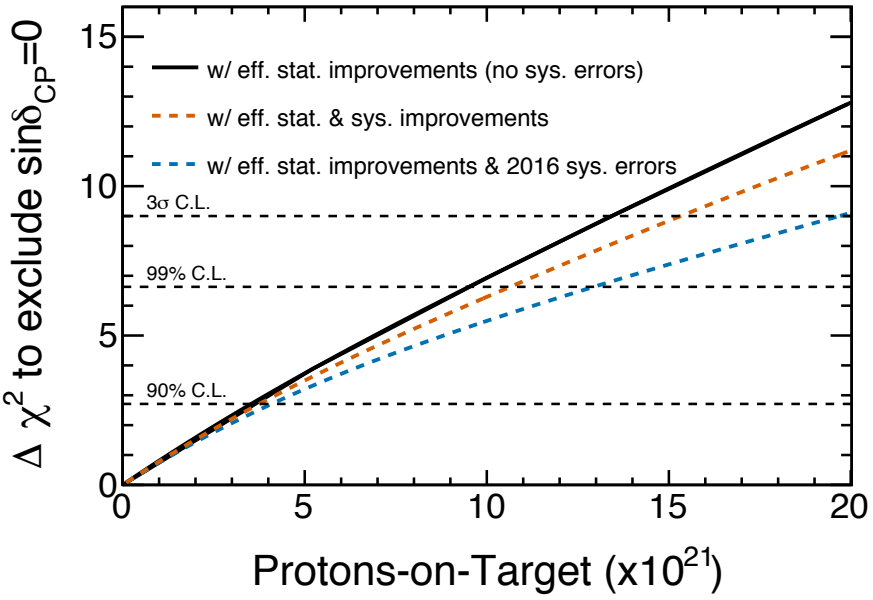

Figure 3.19: T2K sensitivity to CP violation as a function of beam exposure in protons on target (POT) for normal hierarchy and $\delta_{\mathrm{CP}} = -\pi/2$. Figure taken from Ref. [126].

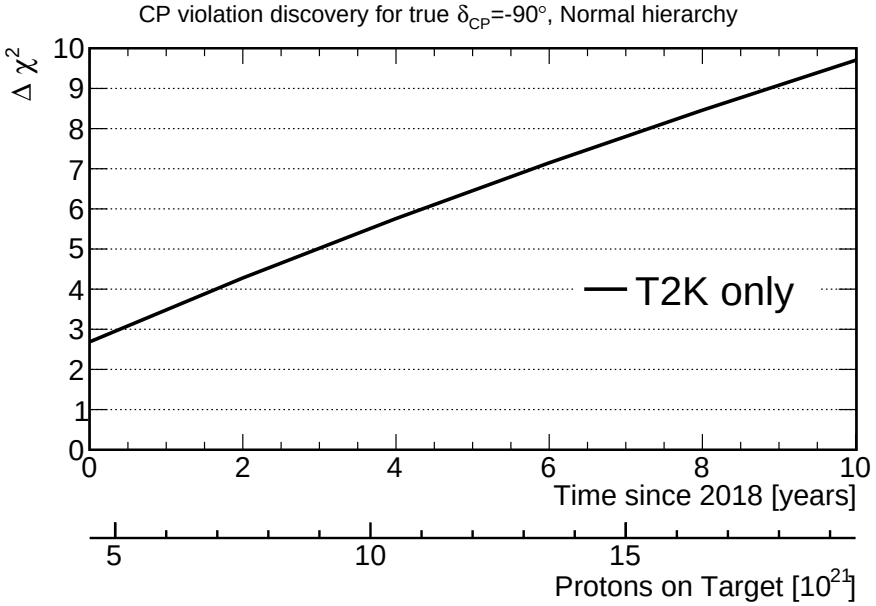

Figure 3.20: GLoBES prediction of T2K sensitivity to CP violation as function of beam exposure for normal hierarchy and $\delta_{\mathrm{CP}} = -\pi/2$. Compare with Fig. 3.19. The POT scale corresponds to 1.5×10^{21} POT a year, with 4.5×10^{12} POT accumulated by 2018.

for the default run time of 5 years in one beam mode, whereas the old one is fully scalable. Efforts were made to use the newest beam fluxes with the old definition, but without updating the channel efficiencies, the sensitivity was severely decreased. In the end, the only change from the default is a rescaling of the beam power by a factor of two. This was chosen so that the predicted sensitivity is similar to that reported by T2K [126], under assumption that a year of operation corresponds to 1.5×10^{21} POT. This number takes a middle ground between the 1×10^{21} POT expected in the next few years, but lower than predicted in 5 years and onwards, keeping general agreement. Figure 3.19 shows the official sensitivity to CP violation as a function of beam exposure, in case of normal hierarchy and for $\delta_{\mathrm{CP}} = -\pi/2$ [126]. The GLoBES reproduction is shown in Fig. 3.20, as a function of time assuming 1.5×10^{21} POT/year, in quite good agreement with the curve representing no improvement in systematic errors. The experiment is expected to accumulate exposure equivalent to 7.8×10^{21} POT by the end of 2021 [50].

3.3.2 Impact of CHIPS-10 in the current landscape

CHIPS-10 is primarily a development prototype to test the technology, but it will also perform actual measurements of electron neutrino appearance, adding physics results to other experiments in order accelerate the global searches. To predict the impact of CHIPS-10 on the combined sensitivity of NOvA and T2K, GLoBES is used to calculate the test statistics for NOvA alone, a combination of NOvA and T2K, and both combined with CHIPS-10; for total exposures corresponding to points in time after CHIPS-10 deployment in 2018.

Figures 3.21 to 3.23 show mass hierarchy determination significance, CP violation discovery significance and δ_{CP} resolution for those three timelines, in the case of normal hierarchy and the true value of δ_{CP} being at its current best-fit point $\delta_{\mathrm{CP}} = -99°$. The shading reflects current 1σ uncertainty around the central value, forming an interval of $-158°$ to $-48°$.

It can be seen from Figs. 3.21 and 3.22 that CHIPS-10 has a very small impact on the combined sensitivity to mass hierarchy or CP violation, although it is not completely negligible. However, there is much more visible improvement for the absolute δ_{CP} measurement (Fig. 3.23), where the addition of CHIPS-10 decreases the measurement error for the best-fit value of δ_{CP} and reduces the range caused by the

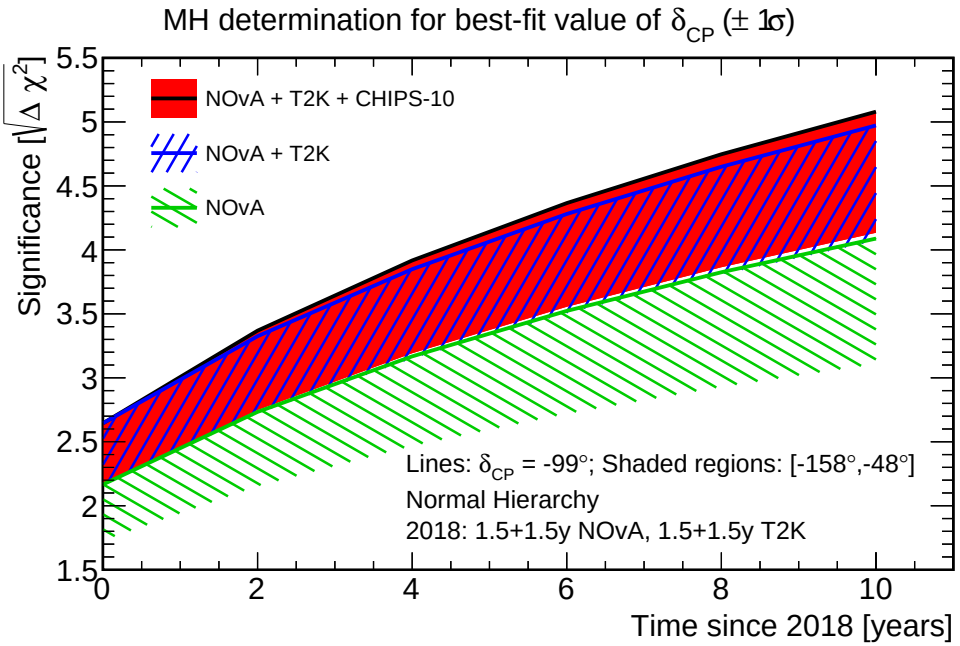

Figure 3.21: Mass hierarchy determination significance as a function of time for a combination of CHIPS-10 and current experiments, in the case of normal hierarchy. The lines show sensitivity for the current best-fit value of δ_{CP} being true and the shaded regions cover the 1σ uncertainty on that number. Displayed experimental combinations are: NOvA alone (green), NOvA and T2K (blue) and both combined with CHIPS-10 (black and red). Equal neutrino and antineutrino beam time is assumed.

uncertainty. This is because the dependence of resolution on the true value of δ_{CP} is much more flat at CHIPS than the two other experiments (see Fig. 3.26). This effect comes from the wider beam energy spectrum, which gives additional shape information to distinguish similar δ_{CP} values.

3.3.3 Staged construction of CHIPS-100

Since CHIPS-10 is a relatively small detector module and at the cost of approximately \$5 million, more than an order of magnitude cheaper than NOvA or T2K, it is worth considering the possibility of deploying a bigger detector.

An example scenario was simulated where after 2018 new modules are deployed in the Wentworth Pit, until the total fiducial detector mass reaches 100 kt. The first

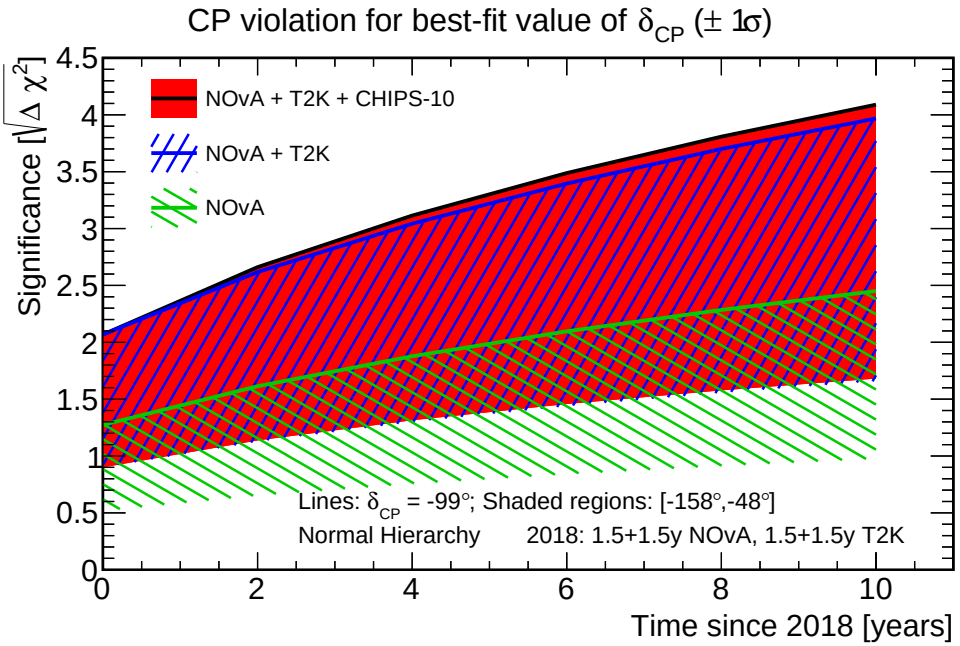

Figure 3.22: CP violation discovery significance as a function of time for a combination of CHIPS-10 and current experiments, in the case of normal hierarchy. The lines show sensitivity for the current best-fit value of δ_{CP} being true and the shaded regions cover the 1σ uncertainty on that number. Displayed experimental combinations are: NOvA alone (green), NOvA and T2K (blue) and both combined with CHIPS-10 (black and red). Equal neutrino and antineutrino beam time is assumed.

module would be CHIPS-10 in 2018, then a 15 kt module in 2020, and three 25 kton modules in 2021, 2022 and 2023.

Figures 3.24 to 3.26 show sensitivity of NOvA, T2K and the expanded CHIPS in 2026, which is the earliest date when DUNE is expected to start operation. The staging plan is equivalent to a CHIPS detector with 58.75 kt running from 2018 until 2026, with total exposure of 470 kt years.

The inclusion of CHIPS offers incremental improvement in the sensitivity to CP violation and mass hierarchy, bringing the significance up to 5σ and 4σ, respectively, for $\delta_{\mathrm{CP}} = -99^\circ$ and normal hierarchy. However, the biggest change is in absolute δ_{CP} measurement, where close to the current best-fit value, CHIPS has the best resolution of all experiments and decreases the global error by 14°.

It is worth noting that even a 100 kt CHIPS is still expected to cost only around $30 million, significantly less than the NOvA or Super-Kamiokande detectors. Ad-

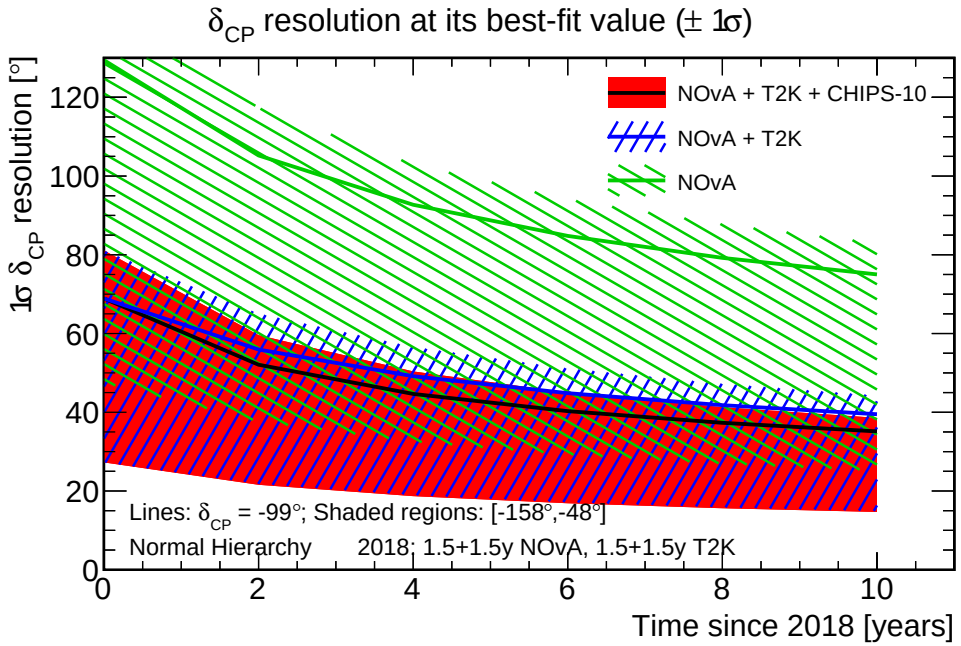

Figure 3.23: Resolution of δ_{CP} measurement as a function of time for a combination of CHIPS-10 and current experiments, in the case of normal hierarchy. The lines show sensitivity for the current best-fit value of δ_{CP} being true and the shaded regions cover the 1σ uncertainty on that number. Displayed experimental combinations are: NOvA alone (green), NOvA and T2K (blue) and both combined with CHIPS-10 (black and red). Equal neutrino and antineutrino beam time is assumed.

ditionally, since the sensitivity prediction for CHIPS likely underestimates actual potential, as explained in Sec. 3.5, the real impact might be even higher.

3.4 Comparison with DUNE

3.4.1 CHIPS in the LBNF beam

An important feature of CHIPS design is the modular construction, making it possible to relatively cheaply disassemble the whole detector and deploy it in a different location. When the new LBNF beamline starts operation, which is predicted to happen in 2026 at the earliest, the NuMI beam will be shut down. CHIPS modules could be relocated to a lake in the path of the new beam to continue data taking.

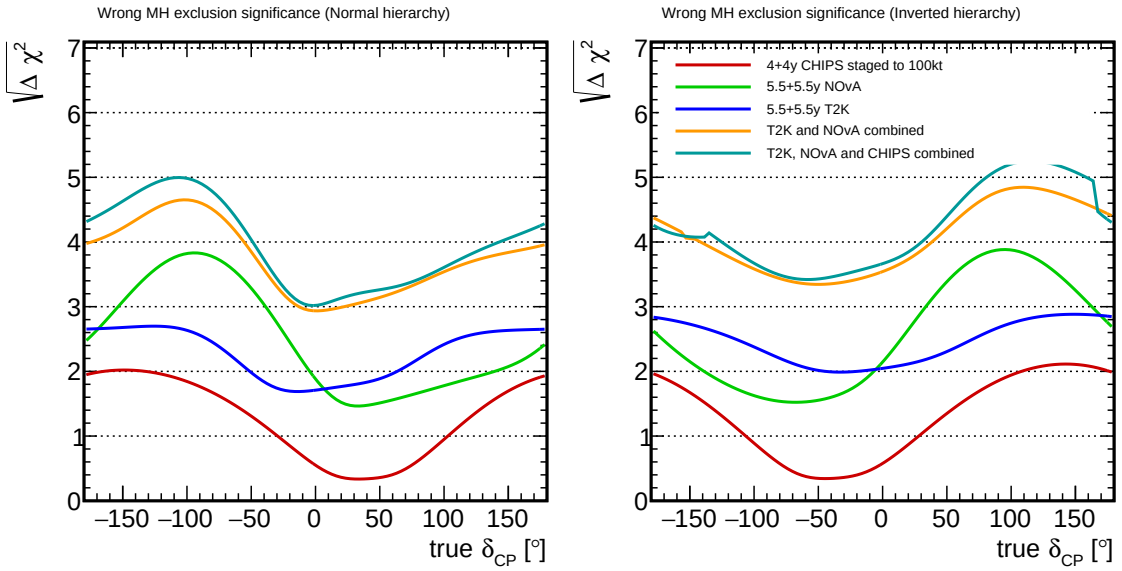

Figure 3.24: Mass hierarchy determination significance as a function of the true value of δ_{CP}, for normal (left) and inverted hierarchy (right). Solid lines correspond to different experiments running until 2026: staged deployment of CHIPS-100 (red), NOvA (green), T2K (blue), NOvA and T2K combined (orange) and all three combined (cyan).

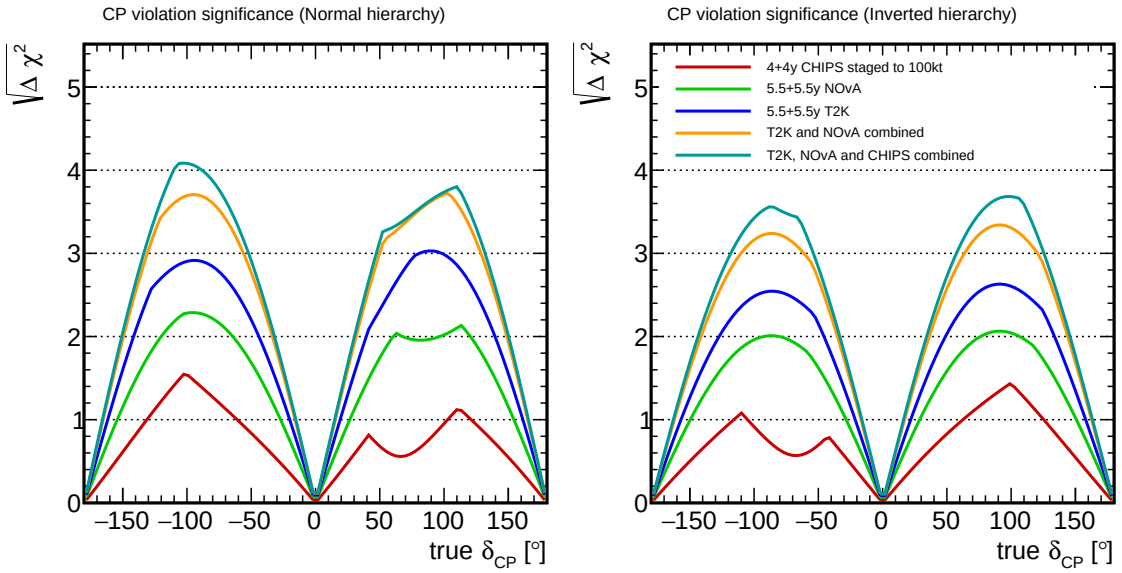

Figure 3.25: CP violation discovery significance as a function of the true value of δ_{CP}, for normal (left) and inverted hierarchy (right). Solid lines correspond to different experiments running until 2026: staged deployment of CHIPS-100 (red), NOvA (green), T2K (blue), NOvA and T2K combined (orange) and all three combined (cyan).

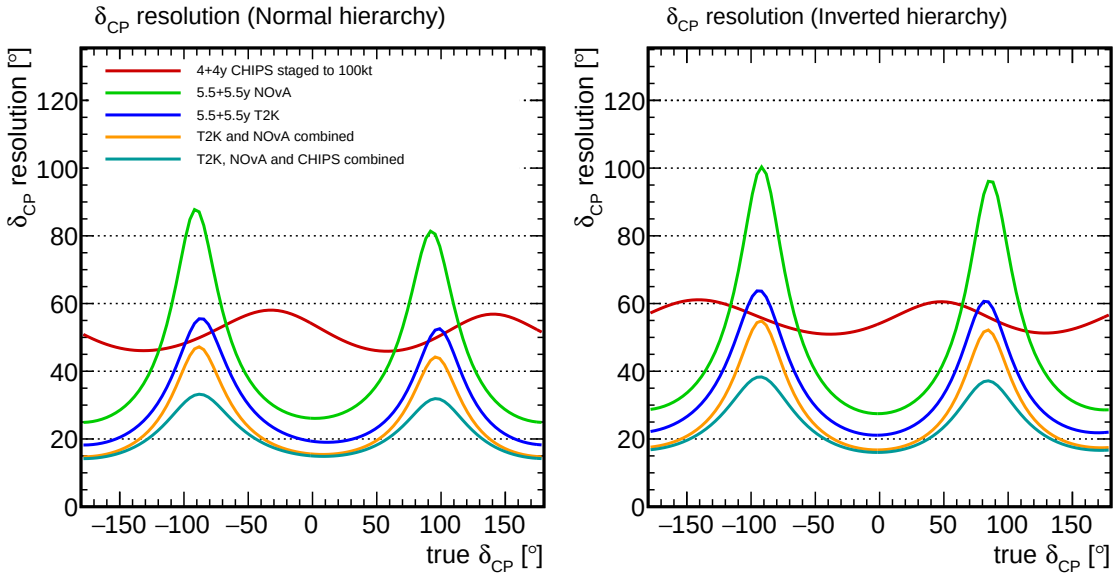

Figure 3.26: Resolution of δ_{CP} measurement as a function of the true value of δ_{CP}, for normal (left) and inverted hierarchy (right). Solid lines correspond to different experiments running until 2026: staged deployment of CHIPS-100 (red), NOvA (green), T2K (blue), NOvA and T2K combined (orange) and all three combined (cyan).

The Pactola Lake is a candidate location for such a scenario. It is a large reservoir in South Dakota with maximum depth of around 50 m, located 1250 km away from the planned LBNF production point, at 20 mrad off-axis. The flux simulation, developed for DUNE studies, is reweighted to predict the neutrino fluxes at the location, shown in Fig. 3.27 and 3.28 for the neutrino and antineutrino beam modes, respectively [130]. All the other elements of the experiment definition are kept unchanged with respect to CHIPS at NuMI.

3.4.2 CHIPS-100 compared to early DUNE performance

It is interesting to see how a CHIPS-100 detector first deployed in the NuMI beam and later moved to LBNF can accelerate physics results in the beginning of the DUNE era. A single 10 kt DUNE module is assumed to begin data taking with the new beam in 2026. The abstract experiment definition for DUNE, including efficiencies, cross sections and the simulated neutrino flux, was received in private correspondence from P. Guzowski [130].

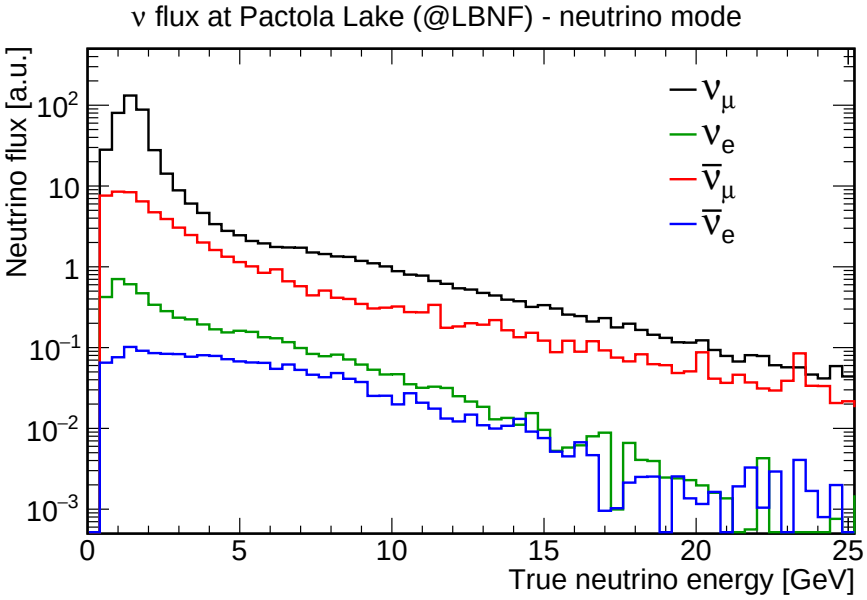

Figure 3.27: Energy distributions of the LBNF neutrino flux at the Pactola Lake, in case of no oscillations. Lines shown for different components of the neutrino mode beam. The ν_τ component is negligible and not predicted by the simulation.

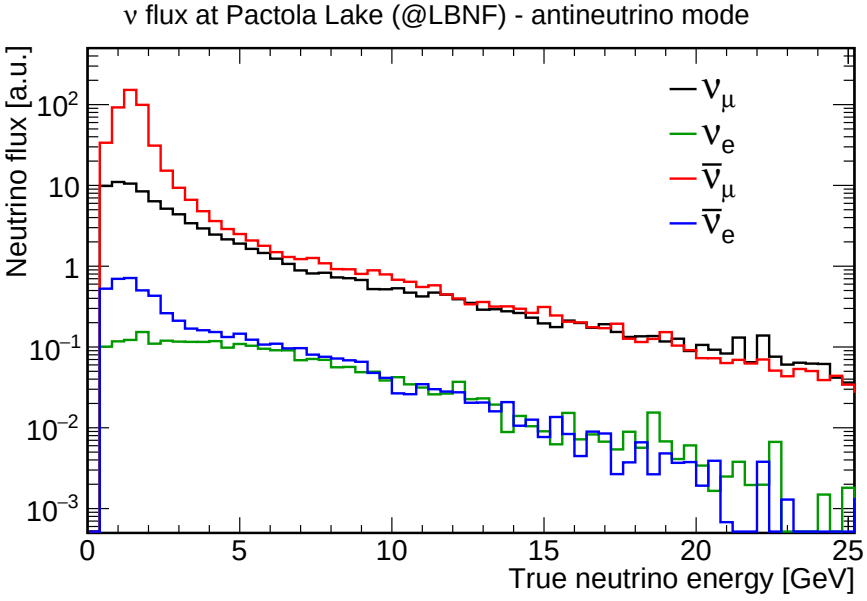

Figure 3.28: Energy distributions of the LBNF neutrino flux at the Pactola Lake, in case of no oscillations. Lines shown for different components of the antineutrino mode beam. The ν_τ component is negligible and not predicted by the simulation.

The following plots present the physics reach of NOvA, T2K, CHIPS and DUNE and their combinations in 2030, with the following assumptions:

- NOvA runs until 2026, when the NuMI is shut down;
- T2K runs constantly until 2030;
- a 10 kt module of DUNE runs from 2026 until 2030;
- CHIPS-100 is built in stages as described in Sec. 3.3.3, and runs in the NuMI beam until 2026, then is redeployed in the LBNF beam where it runs until 2030.

Figures 3.29 to 3.31 show mass hierarchy determination significance, CP violation discovery significance and the resolution of δ_{CP} with NOvA, T2K, CHIPS and DUNE alone, and combinations of all experiments and all except CHIPS.

The inclusion of CHIPS provides a small but non-negligible boost to CP violation sensitivity, raising the combined sensitivity to 5σ in case of normal hierarchy and $\delta_{\mathrm{CP}} = -99°$. The impact on mass hierarchy determination significance is even bigger, due to the long baseline in the LBNF location, and hence high sensitivity to matter effects. In this scenario, CHIPS has again the best resolution of the δ_{CP} measurement in case of maximal CP violation. For the current best fit value of δ_{CP}, including CHIPS in the global combination reduces the measurement error by more than 20 %.

3.5 Conclusions

3.5.1 Comparison with previously published results

The predicted sensitivities for CHIPS experiments and their contribution to global combination results are definitely not negligible, but may also not seem very impressive. In particular, the impact of CHIPS appears to be significantly smaller than in previous studies [93]. This difference is illustrated in Fig. 3.32, which compares CP violation significance and δ_{CP} measurement resolution as presented in [93] with new versions updated with current GLoBES definitions.

The dashed lines on these plots show sensitivities recreated with the new code, but same settings as the original results. This includes using oscillation parameter

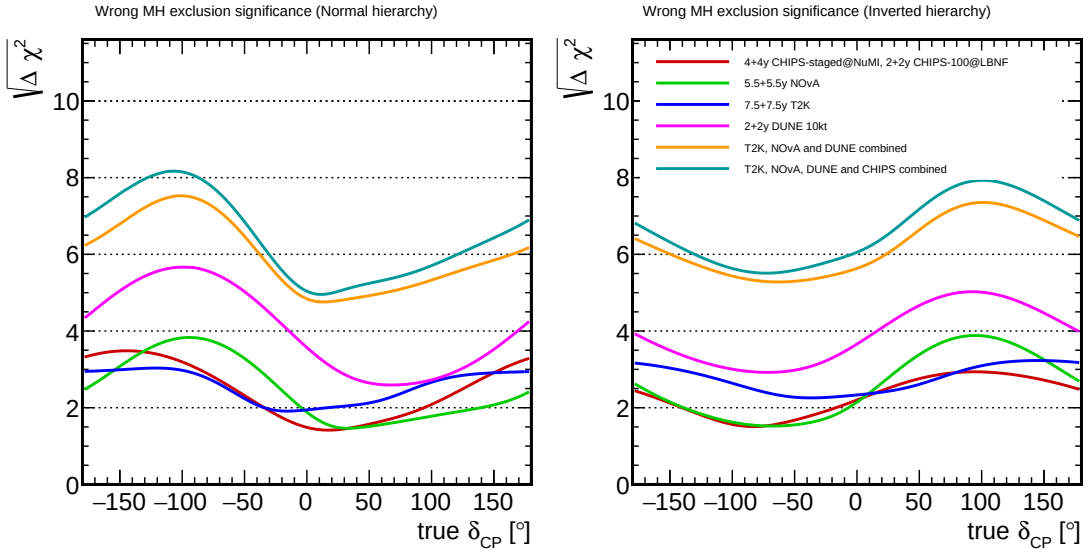

Figure 3.29: Mass hierarchy determination significance as a function of the true value of δ_{CP}, for normal (left) and inverted hierarchy (right). Solid lines correspond to different experiment combinations in 2030: staged deployment of CHIPS-100 in NuMI and redeployment in LBNF (red), NOvA running until 2026 (green), T2K (blue), 10 kt DUNE running from 2026 (dark blue), combination of all except CHIPS (orange), and all with CHIPS (cyan).

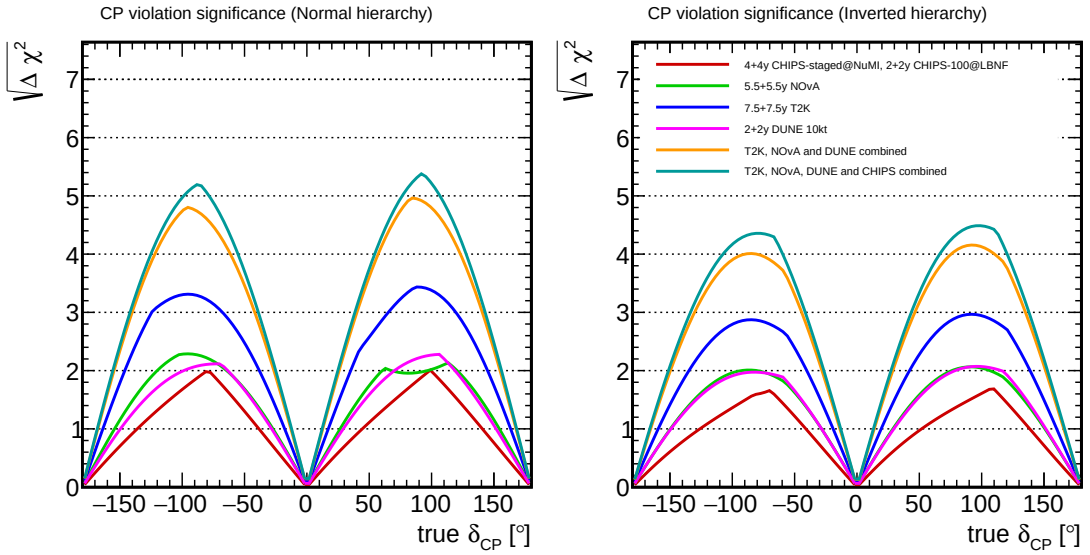

Figure 3.30: CP violation discovery significance as a function of the true value of δ_{CP}, for normal (left) and inverted hierarchy (right). Solid lines correspond to different experiment combinations in 2030: staged deployment of CHIPS-100 in NuMI and redeployment in LBNF (red), NOvA running until 2026 (green), T2K (blue), 10 kt DUNE running from 2026 (dark blue), combination of all except CHIPS (orange), and all with CHIPS (cyan).

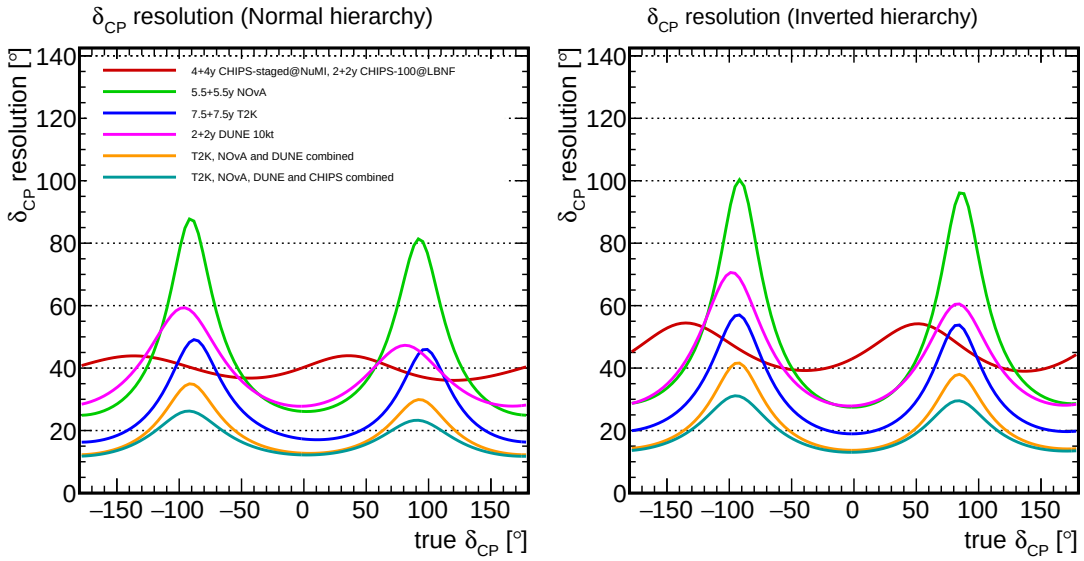

Figure 3.31: Resolution of δ_{CP} measurement as a function of the true value of δ_{CP}, for normal (left) and inverted hierarchy (right). Solid lines correspond to different experiment combinations in 2030: staged deployment of CHIPS-100 in NuMI and redeployment in LBNF (red), NOvA running until 2026 (green), T2K (blue), 10 kt DUNE running from 2026 (pink), combination of all except CHIPS (orange), and all with CHIPS (cyan).

values from June 2012 [93, 131], lack of marginalisation over the parameters and a T2K definition with 6 years of neutrino beam only. The presented scenario assumed that CHIPS-100 will start operation 4 years after the start of NOvA and used the SK-based GLoBES definition (Sec. 3.2.3). The resulting figures are virtually identical to the ones shown in Ref. [93], confirming the validity of the new code.

The solid lines in the plots in Fig. 3.32 show sensitivity values calculated using all the same settings except updated definitions of CHIPS-100 (Sec. 3.2.1) and T2K (Sec. 3.3.1). In the latter case, the new definition includes a longer running time (10 years as opposed to 6) and the use of an antineutrino beam for half of the run.

The figure shows that the lower relative contribution of CHIPS to the global physics reach compared to previous results is caused by two main effects. The first one is a significant boost in performance of T2K, coming from the inclusion of both beam modes, increased beam intensity and the longer running time. The second effect is the decrease in CHIPS effectiveness as described by the new GLoBES definition compared to the old one (see Sec. 3.2.3).

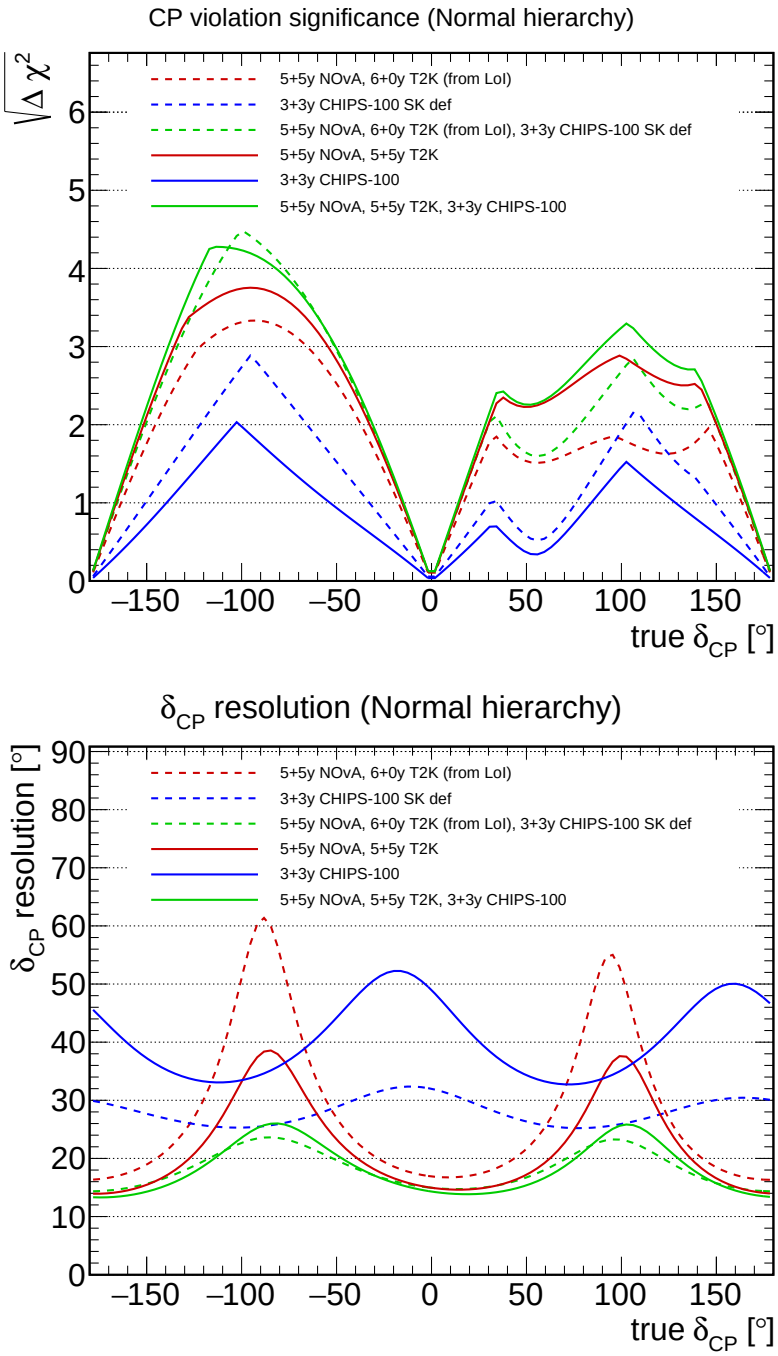

Figure 3.32: CP violation discovery significance (top) and resolution of δ_{CP} measurement (bottom) as a function of the true value of δ_{CP}, for normal hierarchy, comparing the results presented in [93] and updated with new GLoBES definitions. Dashed lines correspond to different experiment combinations, defined as in [93]: 6 years of CHIPS-100 (blue), 10 years of NOvA combined with 6 years of T2K in neutrino mode only (green), and all three combined (green). Solid lines show the values for updated GLoBES definitions including a longer exposure of T2K with an antineutrino beam ($5+5$ years), but same methodology (different to the one employed in the rest of this work).

In addition, the sensitivities presented earlier in this work are calculated with marginalisation over the oscillation parameters and updated best-fit central values. Although these changes have the same effect on the physics reach of CHIPS and other experiments, they may create the impression of an even further decrease of CHIPS experiment sensitivity relative to the predictions shown in Ref. [93].

3.5.2 Future improvements

Some of the shortcomings of the CHIPS sensitivity predictions can be attributed to the fact that the experimental definition is based on the performance of a very preliminary reconstruction implementation. As the work on its optimisation continues, it will be possible to directly input the new results in the existing framework to recalculate more robust sensitivity predictions.

The event sample used for the evaluation of the reconstruction and selection performance is relatively small. The statistical fluctuations propagate to the smearing matrices and the selection efficiencies. In addition, the process was only performed for neutrinos. Reproducing the reconstruction on the antineutrino sample, and on the neutrino sample with significantly higher statistics, are the two important next steps to be performed. This will increase the reliability of the efficiencies and smearing matrices, and provide valid quantities for antineutrinos.

Even more importantly, the selection is not yet fully optimised. Current CHIPS performance is mainly limited by the relatively poor NC background rejection. When a proper π^0 fit is implemented in the reconstruction, the NC interaction identification and rejection should improve considerably. A simple study of the effect of improved particle identification on CHIPS sensitivity was performed by artificially reducing the NC background component in the GLoBES definition by 80 %. This results in CP violation discovery and the mass hierarchy determination significance levels improved by 20 % to 30 % for an exposure of 600 ktyears. In addition, selection training on non-QE CC ν_μ events will likely further improve the rejection of misidentified muon tracks.

On the GLoBES code side, the non-Gaussian distribution of the best-fit θ_{23} value has to be implemented in order to study the sensitivity to the octant of the mixing angle. With the current Gaussian errors, sensitivity would be significantly overesti-

mated. This modification will also increase the fidelity of sensitivity predictions for mass hierarchy and CP violation.

As observed in the previous chapter, increasing the photocathode coverage does not significantly improve the reconstruction performance. However, the reconstruction and selection procedure can still be optimised, reducing the background contamination and increasing sensitivity. In addition, adding more modules will increase the total mass and exposure of the detector. Thanks to the modular design and the lack of need for excavating underground caverns, CHIPS is well suited for such ongoing expansion.

At the same time, the performance of NOvA and T2K appears to be better than predicted in previous studies, and CHIPS would have to be built more rapidly to contribute significantly to the global results. However, it is possible that the start of DUNE operation gets delayed and in this case deploying new CHIPS detectors, either in the NuMI or LBNF beamlines, may prove to be a cost-effective strategy for bridging the gap between the current and next-generation experiments.

Chapter 4

R&D status and prototype testing

The CHIPS project employs a staging philosophy. The development starts with small proof-of-principle prototypes, followed by modules with increasing functionality, allowing for the construction of working detectors while still testing and improving the design. This approach is flexible to funding and able to produce physics results contributing to the global combination in a short timescale. The first prototypes are deployed in the Wentworth Pit 2W in northern Minnesota, which is located in the path of the NuMI beam (Sec. 2.1). Two testing campaigns with CHIPS-M (CHIPS-Model), a small scale test-bed setup, have already been performed.

During summer 2014, the CHIPS-M structure was built and submerged in the pit, instrumented with five IceCube DOMs, containing 10 inch PMTs. This first test run was focused on testing the basic principles of the mechanical design and procedures during construction, deployment and operation under water. The CHIPS-M detector was successfully recovered in June 2015, after spending nearly a year at the bottom of the lake. It underwent several small improvements and was equipped with two prototype detection units. The main goals of the second deployment were the testing and validation of the prototype planes design and readout operation.

This chapter describes the goals and timelines of the two test runs, the CHIPS-M structure and infrastructure, and the tested instruments, with focus on the Nikhef prototype unit. Analysis of the data recorded by the Nikhef plane is presented in the next chapter.

4.1 The first testing campaign

The first testing campaign was centred around the deployment of CHIPS-M detector in the Wentworth Pit. It started in early summer 2014, when the construction of CHIPS-M began in the MINOS surface building in Soudan, Minnesota. The module, instrumented with five optical modules borrowed from the IceCube experiment, was successfully deployed in the end of July and operated under water until June 2015, when it was resurfaced, inspected and prepared for the second deployment.

The main goals of the first test run were:

- Validation of the basic mechanical design principles, in particular the suitability of geomembrane as the tank liner material and of the sealing method.
- Testing of the performance of a simple water filtration system.
- Monitoring of the stability of the detector instruments and conditions at the bottom of the lake.
- Inspection of the structure and electronics after a long stay under water.
- Practice with deployment and recovery operations on water.
- Use of simple instrumentation to gather data on the underwater cosmic muon flux (and a potential observation of beam neutrino interactions).

4.1.1 The CHIPS-M platform

CHIPS-M can be used to mean just the mechanical structure, the whole detector module including instrumentation, or even the whole test run, taking data with cosmic muons. In this work, it is used to describe only the structural platform, without the tested instruments. The name can be expanded as CHIPS-Model, in analogy to CHIPS-10, as it is a small scale model of the full detector. CHIPS-M was designed primarily to test the viability of submerging a lightweight tank structure sealed with flexible liner and operating it at the bottom of a lake as a water Cherenkov detector.

The mechanical structure (Fig. 4.1) has the shape of a regular octagonal prism, with two truss endcaps joined by stainless steel Unistrut columns. The liner is

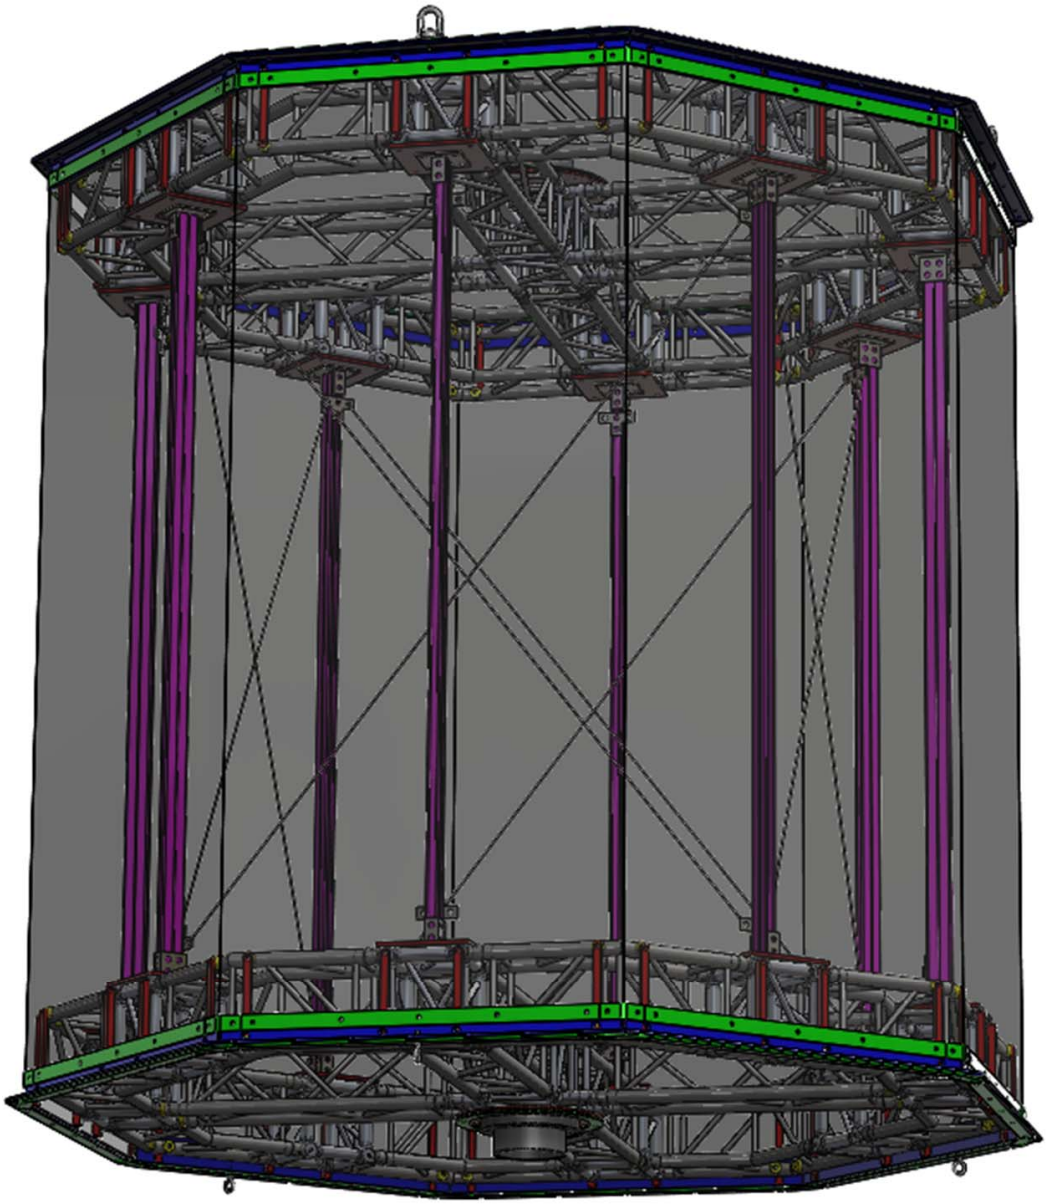

Figure 4.1: Design diagram of CHIPS-M, as viewed from below. The main structure consists of two truss endcaps, connected with steel columns (in purple). The module is wrapped with a liner, here shown as semi-transparent for clarity. Also depicted are the cross-link cables, the batten bars and rims used for clamping the liner (in green and blue), and the feedthrough flanges in the endcap centres. Image taken from Ref. [132].

bolted to the frame with aluminium batten bars. Two flanges in the centres of the top and bottom caps contain feedthroughs for electrical and water connections. Inside, mounts, attached to the Unistrut columns or to the truss directly, hold the instrumentation. The structure sits atop a wide steel tripod.

The endcaps are made of aluminium truss, assembled like a stage frame from smaller components, custom-welded by a commercial company. The caps are octagonal in shape, with a side length of 1.32 m, total length from side to side of 3.32 m, and 25 cm height. Main nodes of the truss consist of aluminium cubes with tapped holes in each free side. This allows the attachment of fixtures for the Unistrut columns, the inner batten rims for the liner, and the mounts for the IceCube Optical Modules.

The two caps are connected and supported by eight stainless steel columns, made from H-channel Unistrut beams. To preserve the rigidity of the structure, particularly against twisting, there are eight steel cross-link cables, tightened with turnbuckles. They connect the top and bottom of every other column (Fig. 4.1). The total height of the structure is 3.318 m.

The liner is made of a Seaman 8130 XR-5 geomembrane, a chemically resistant polymer-coated polyester fabric, used for lining water pools and reservoirs. The material is very resilient to wear and tear, and can be joined by welding two sheets together. The liner is attached to aluminium rims surrounding the circumference of the endcaps, and clamped with bolted aluminium batten bars (Figs. 4.1 and 4.2c). Sealing is achieved by placing a tacky butyl tape between the liner and the inner rim (Fig. 4.2b). The L-shaped cross-section of the rim allows for the attachment of both the endcap and the side liner. The wall liner is closed up with a vertical seam, with the batten bars attached to one of the Unistrut columns. The supplied liner was black on one side and white on the other. It was installed with the black side facing inside, to minimise light reflection inside the detector and increase timing precision of the PMT hits.

Because the conditions at the bottom of the pit were unknown, a stand was designed to protect the bottom liner from possible sharp rocks. It consists of a triangular steel frame with three 1.22 m long legs (Fig. 4.3a), ending with flat plates to act as feet in case of silt or mud. The tripod is bolted to the structure using some of the holes in the bottom rim.

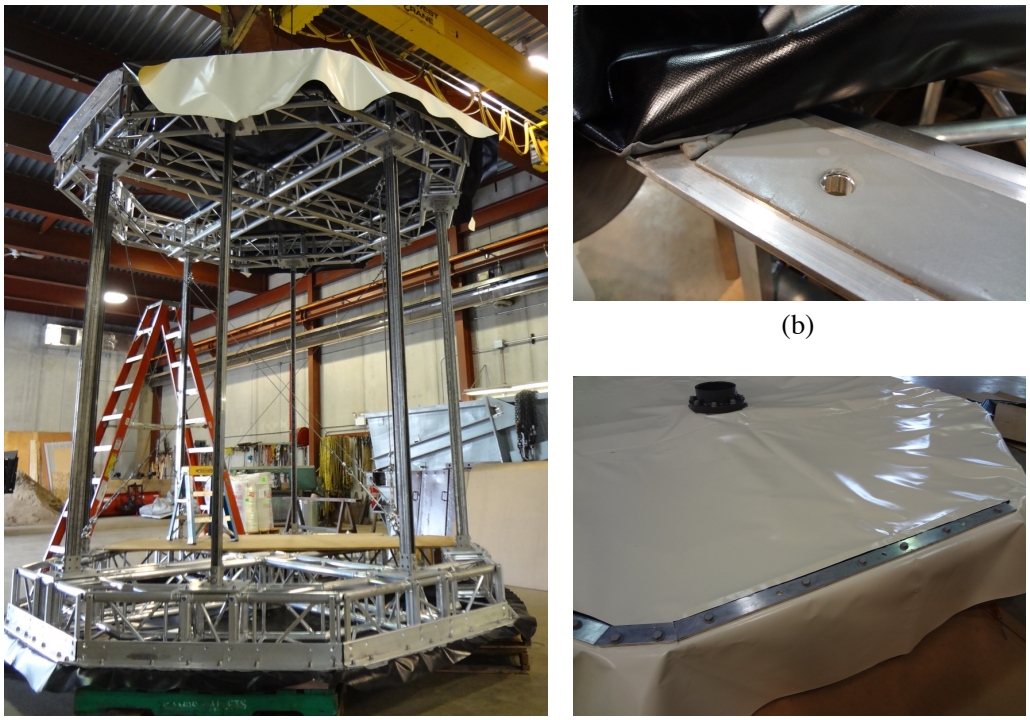

(a) (b) (c)

Figure 4.2: Early stages of the construction of CHIPS-M. (a) The main structure with two truss endcaps, connected with Unistrut columns, with cross-link cables installed. (b) Butyl tape for sealing the liner. (c) Batten bar bolted to the rim beneath, clamping the liner.

In the centre of each endcap, there is an opening for connecting cables and pipes. PVC pipe flanges are bolted to the truss clamping the liner, sealed with the butyl tape. Flat PVC disks are glued inside, with feedthroughs and fittings embedded in them. In the top flange, an air return valve and a manual venting valve are fixed with a threaded connection. In the bottom one, power and communication cables are fed through holes and sealed with RTV silicone and epoxy resin (Figs. 4.3b and c). In addition, water circulation pipes are glued in standard PVC couplings. Inside the detector, the water supply pipe is routed to the top to ensure an even water circulation. The cables and pipes exiting the bottom flange form an umbilical connecting the detector with the shore. They are bundled together and protected with a corrugated conduit.

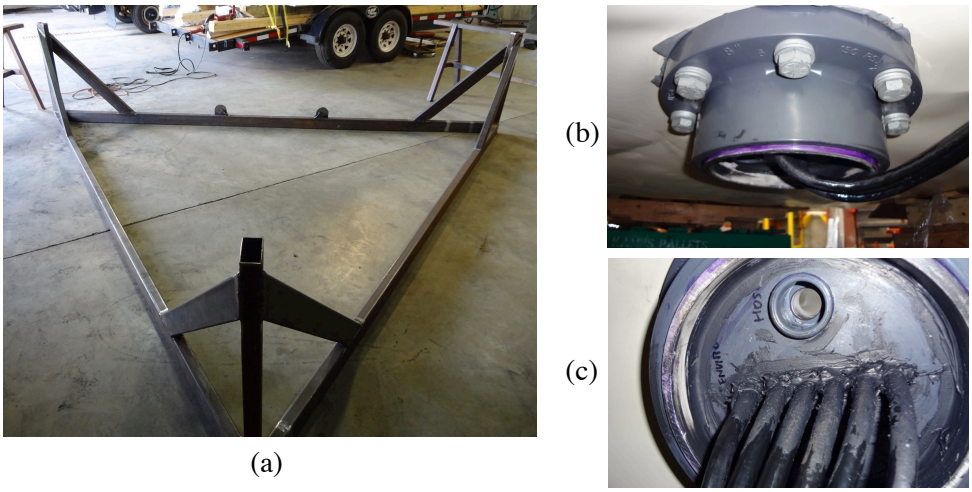

(a) (b) (c)

Figure 4.3: (a) The tripod stand during construction, before welding on the feet plates. (b) and (c) The bottom feedthrough flange with a PVC disk and electric cables sealed with RTV. The water circulation pipes were glued glued into two standard PVC couplings.

4.1.2 IceCube instrumentation

The first test campaign, was focused on testing the engineering and operational principles, and no detection hardware had yet been developed. Five Digital Optical Modules (DOMs) with 10 inch PMTs were borrowed from the IceCube experiment, along with the necessary data acquisition (DAQ) hardware, to install on board of CHIPS-M as instrumentation.

IceCube is an astrophysical experiment at the South Pole, with strings of spherical modules suspended deep in the Antarctic ice, effectively creating a Cherenkov detector with a volume of about a cubic kilometre, sensitive to ultra high energy neutrinos [133].

Each DOM contains a single PMT with 25 cm diameter and full readout electronics, encased in a 17 inch diameter glass sphere (Fig. 4.4). The casing is made of two hemispheres, held together by a metal ring around the equator, which also serves as an attachment point for ropes (or mount structures in CHIPS-M).

The electronics boards supply the high voltage to the PMT, read out and digitise the signal, and store it before sending to the surface. Both the power supply and communication use a single copper wire pair, contained in a cable which is fed through a high pressure penetrator. The DOM contains also a flasher system with 12 light

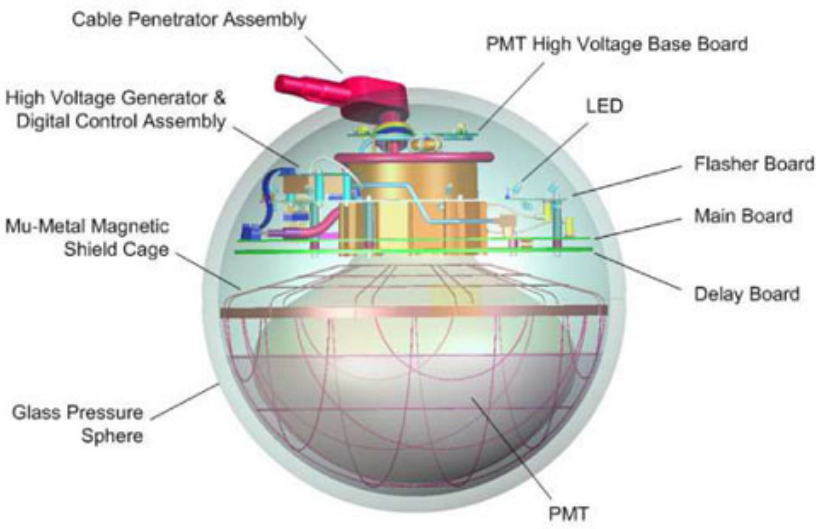

Figure 4.4: Schematic cross-section of the IceCube Digital Optical Module with main components labelled. Image taken from Ref. [133].

emitting diodes (LEDs), used for time calibration and monitoring the transparency of the surrounding ice (or water).

The five borrowed DOMs were identified by their names: *Haifa*, *Melbourne 1956*, *Munich 1972*, *Filea* and *Angkor Wat*. The first four were installed in an upper corner of CHIPS-M facing the direction of the NuMI beam. Because the optical gel coupling the PMT to the glass hemisphere cannot hold if the DOM is not pointing downwards, they were angled at 45° (Fig. 4.5a). To increase the chances of observing beam neutrino events, the final DOM was assigned the role of a veto PMT covering the opposite direction. Since it could not be placed facing upwards, two mirrors were installed in the lower upstream corner of CHIPS-M, and *Angkor Wat* was mounted there, pointing downwards (Fig. 4.5b).

The data acquisition system consists of a single DOMHub PC, which includes the DOM power supply units and two boards responsible for DOM communication, called DOR cards. There is also a filter box that emulates the properties of IceCube cables, which are an order of magnitude longer than the ones in CHIPS-M. A timing receiver, using the Global Positioning System (GPS), provides the time synchronisation signal. Additionally, a box called the StarDOM emulates a sixth DOM and allows for triggering on the NuMI beam spill time-stamp relayed from Fermilab.

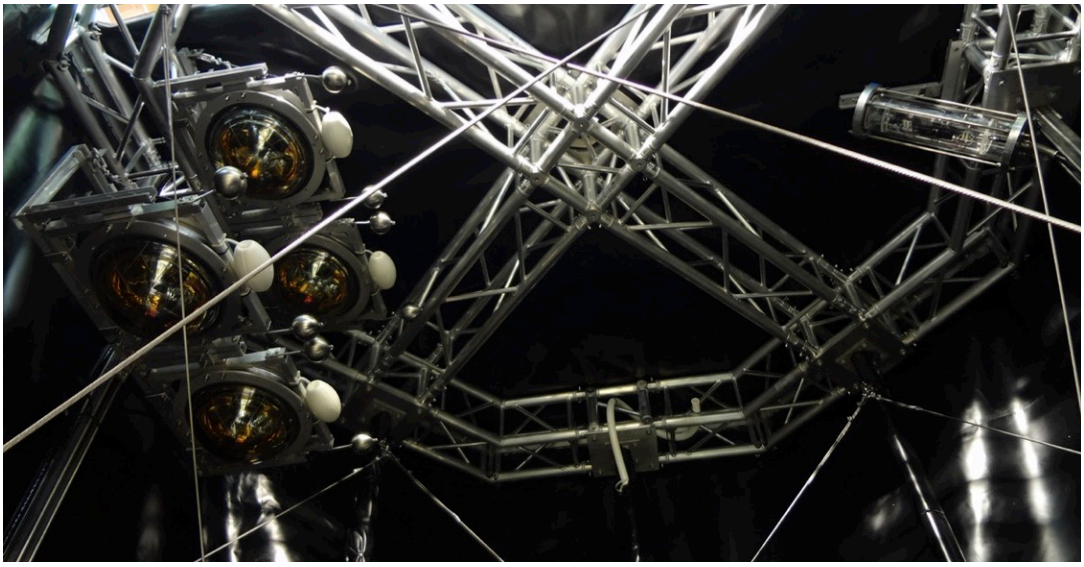

(a)

(b)

(c)

Figure 4.5: (a) Inside CHIPS-M before the final sealing, looking up. The four DOMs visible on the left are oriented downwards for the transport, but floats on the mounts lift them to a 45° angle when submerged. The inside environmental vessel with a wide-angle camera covering all five DOMs is visible to the right. (b) The fifth veto DOM located directly beneath the environmental vessel and the two mirrors (one vertical and one horizontal). (c) The shed housing water filtration system and the on-shore DAQ hardware at the pit shore. Also visible is the gas-powered generator.

4.1.3 Environmental sensors

Apart from detector instrumentation, CHIPS-M was fitted with two environmental sensor vessels. The electronics were housed in pressure resistant containers, consisting of a glass cylinder with aluminium endcaps, held together with three threaded rods and sealed with rubber gaskets. One of the two identical vessels was placed inside the detector (Fig. 4.5a) and one on the outside. Each contained a BeagleBone black single-board computer, a web camera with white LED lights, and external temperature and pressure sensors embedded in the endcaps. A common triple Cat5 copper cable, connecting both vessels with each other and the outer one to the shore, provided Ethernet communication with the BeagleBone Black computers, and separate power for the LED lights and the remaining components. The outside vessel contained a network switch and served as a communication hub to the inner one. The main purpose of the sensors was to monitor the conditions inside and outside of the detector during the deployment, particularly when filling CHIPS-M with water and lowering it to the bottom of the pit.

4.1.4 On-shore infrastructure

The umbilical bundle, protected in a conduit, is around 300 m in length, and connects the deployed detector with the on-shore operations centre. It contains five cables for the DOMs, the triple Cat5 cable for the environmental vessels and two 3/4 inch water pipes for the filtration system.

The on-shore infrastructure included the water filtration circuit and data acquisition hardware, both housed in a small shed situated next to the pit shore. Electric power was provided by a small gas generator. The other components were a gas heater, a GPS receiver for timing, a long-range Wi-Fi antenna for internet connection, and a UPS battery pack in case of power outage.

For the small CHIPS-M module with its limited instrumentation the water clarity is not such a critical issue. However, in a bigger detector limiting light losses is very important and water purification is necessary. CHIPS-M provided a great opportunity to test the performance of a simple system with mechanical filters. To limit the costs, a small commercial installation was ordered. It contained a set of replaceable cartridge filters with the following ratings: one 5 µm, two 1 µm and two

0.222 µm; as well as an activated carbon filter and a UV lamp for water sterilisation. More advanced processes, such as deionisation or reverse osmosis were not used.

A centrifugal water pump was used to fill the detector during deployment and circulate it during operation. Later, a self-priming jet pump was also installed, to help with issues due to air rising from the detector (likely trapped since deployment).

An electronic rack housed the DAQ and control hardware. This included the DOMHub computer with the peripheral boxes, the GPS timing system, power supplies and a second PC for with the environmental sensors, and a gateway machine through which the system was connected to the internet.

4.1.5 Construction and deployment

The construction and assembly of CHIPS-M took place in the MINOS Surface Building, which was part of the Soudan Underground Laboratory, in Soudan, Minnesota.

First, the liner was attached to the two endcaps, along with the pipe flanges in the centres. The two endcaps were then connected with the steel Unistrut columns, and the cross-link cables were installed. The side liner was attached to the top and bottom rims around the structure, leaving a gap before closing the final vertical seam. At the same time, the IceCube DOMs, the two mirrors and the environmental vessel were installed inside, and the cables were sealed in the feedthrough flange. Finally, the side liner was sealed with the vertical batten bars. Because the clamping was done with bolts and nuts, a person had to stay inside and exit through an unsealed part of the top liner. This made it difficult to open the detector again, if a need would arise, and was changed in the following year. After the final seal, CHIPS-M has not been opened again before deployment.

A light-tightness test was performed at night-time, by switching on the PMTs and directing a source of light at different parts of the structure. The liner and seals were confirmed to be light-tight, but the white PVC pipes fed through the bottom flange were found to introduce a light leak. This was fixed by wrapping spare liner around the umbilical bundle where it exits the detector.

The detector was transported to the Wentworth Pit on a low loader trailer with air suspension, at a low speed of 30 mph (Fig. 4.6b). There, it was mated with the

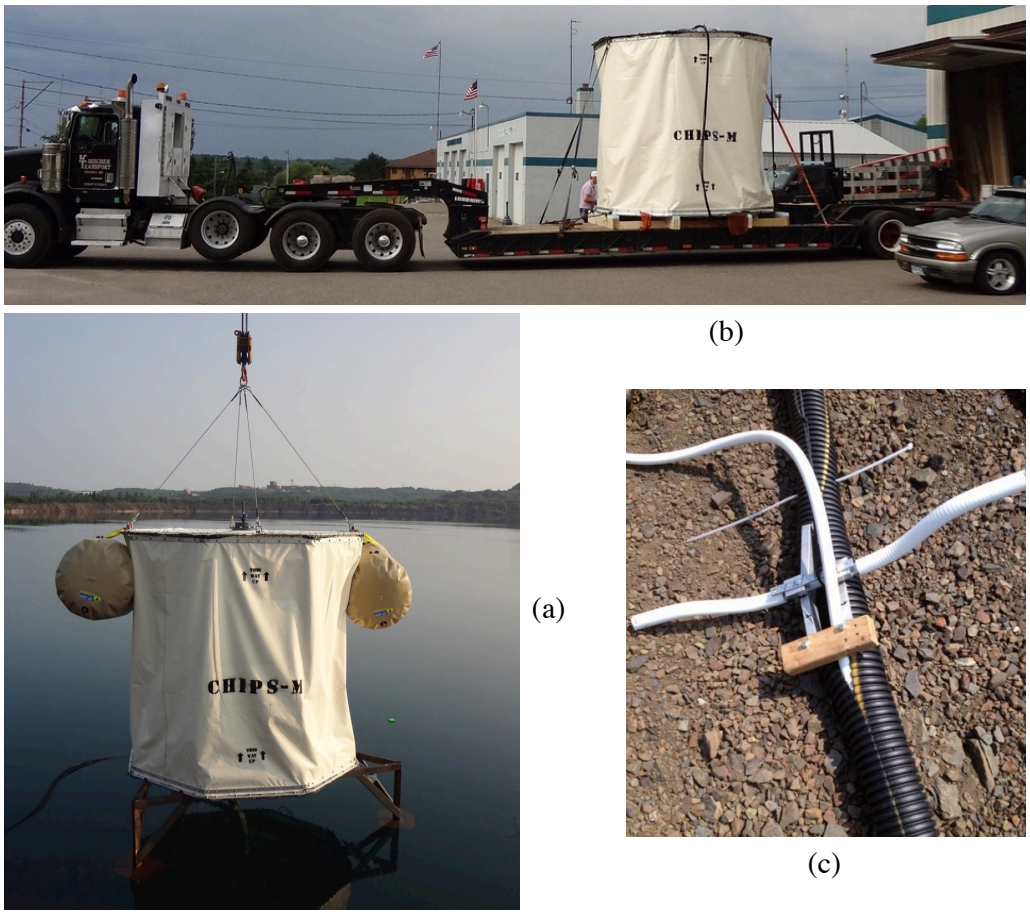

Figure 4.6: (a) The deployment of CHIPS-M. The detector is lowered into the pit from a crane and slowly filled with pure water. (b) Transport of the sealed structure to the pit site. (c) The *longhorn* tool used to place the umbilical bundle inside the conduit, here shown with a single PVC hose. It creates an opening in the conduit, which can be easily moved along.

tripod stand, which had been transported separately. The IceCube DOM pigtails were coupled with the 300 m long umbilical cables using underwater connectors. The conduit was cut open lengthwise and the umbilical bundle was inserted with the use of a custom-made tool, named the *longhorn* due to the shape of the handles (Fig. 4.6c). Cable ties were used to close the conduit and secure the bundle inside.

Inflatable floating bags were attached to the top of CHIPS-M to provide buoyancy after submerging. The detector was lifted with a crane and slowly lowered into the pit water, while being filled with pure water passing through the filtration system (Fig. 4.6a). The pigtail connectors were used to temporarily attach a short water pipe, to avoid the resistance of the 300 m long one and speed up the filling process.

Despite that measure, the filling took almost 10.5 hours. The detector was left in the water overnight, suspended from the crane.

The next morning (on the 1st of August 2014) CHIPS-M was disconnected from the crane. The support cable was attached to the winch cable, and after deflating the air bags, the weight of the detector was transferred to the winch assembly located on a barge built out of floating dock segments. The barge and the attached detector were towed with a boat to the final location on the lake. The location had been identified earlier, using a fishing sonar to find a place with sufficient depth and a flat and featureless lake bed.

Plastic 1 gallon water bottles were tied to the umbilical to act as buoys. It was pushed on water and it stretched across the pit surface as the detector was towed to the location. As the detector was lowered with a hydraulic winch from the floating platform, the buoys were systematically cut off from a boat, letting the umbilical sink to the bottom.

As CHIPS-M reached the depth of 6 m above the bottom, a buoy was attached to the winch cable to allow for a later capture and retrieval. Then, the detector was lowered again until it finally rested at the bottom of the pit. A rope, attached to one of the tripod legs and tensioned from the shore was used to ensure the correct orientation of the module, with the four DOMs facing upstream in the NuMI beam.

4.1.6 Operation and recovery

After the successful deployment it was found that the veto DOM *Angkor Wat* would not enter data-taking mode. This was later confirmed to have been caused by issues with the underwater connector. The other four DOMs functioned nominally, but revealed that there was a light leak in the detector. Because of this, data taking was limited to night-time.

No special water tests before the deployment were performed, mainly due to tight time constraints.

CHIPS-M operated with the four IceCube DOMs for almost a year, recording interactions in the detector every night. Figure 4.7 shows the rates of single PMT hits with a charge above 2 photoelectrons and the rates of events defined as a coincidence

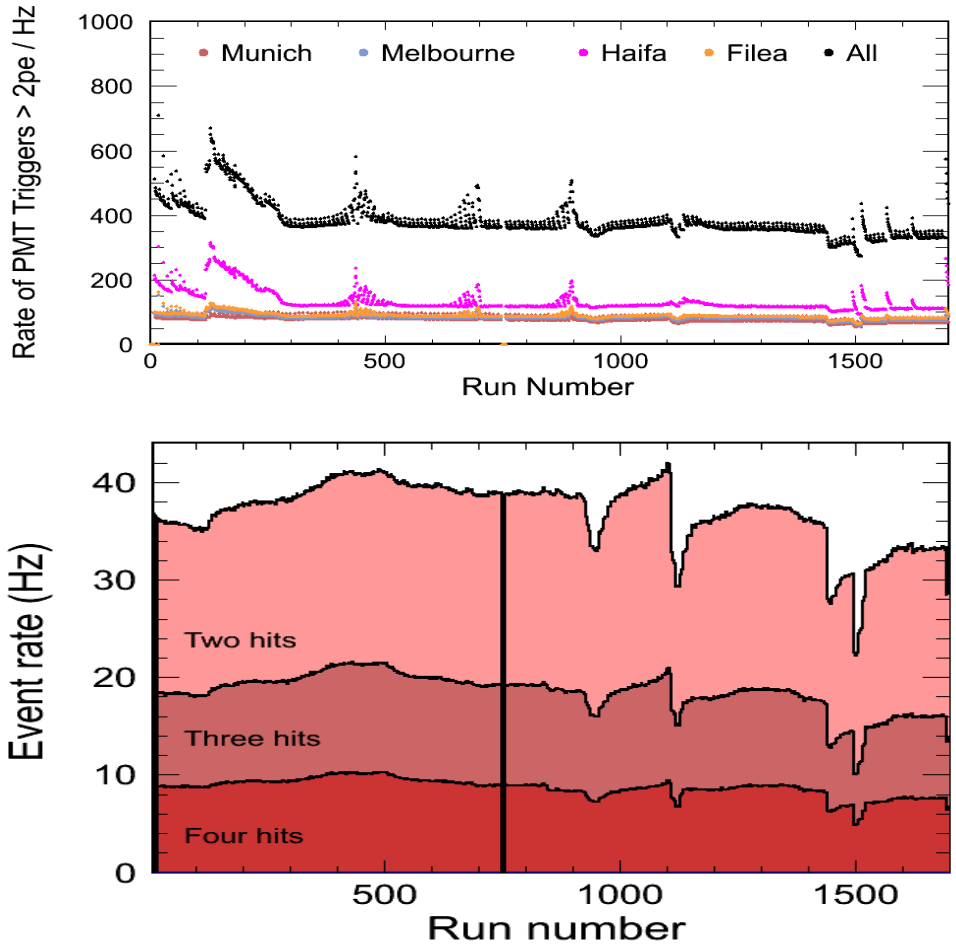

Figure 4.7: Results from four IceCube DOMs running in the CHIPS-M during the first test campaign. The plots show the single PMT hit rate above a 2 PE threshold (top), and the rate of coincident hits above 5 PE within 15 ns. Visible spikes of the single rates seem to correlate with full Moons, while the sudden drops in hit and event rate correspond to interruptions in the water circulation. Figure taken from Ref. [134].

of n hits with at least 5 PE and within a 15 ns time window. The time is expressed in run numbers, starting on 20th of August, and with 9 1 hour long runs taken every night. The hit rates exhibit periodic spikes, which correspond roughly to the occurrence of the full Moon. The dependence is not exact, but cloud cover, which was not tracked, could also affect this. The spikes in single hit activity have little effect on the multi-hit coincidences. However, the event rates are not stable and change over time. In the second half of the shown run period, there are visible several sharp drops and a general decrease of the coincidence rate.

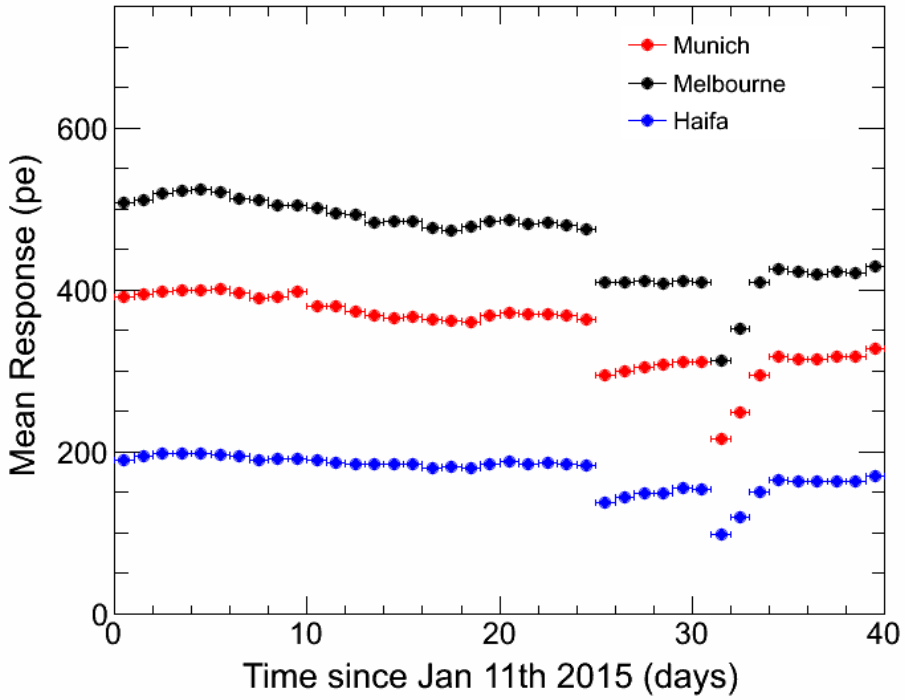

Figure 4.8: Mean response of three IceCube DOMs to the LED flasher on *Filea* during the water monitoring runs, as a function of time. Sharp drops in water clarity, followed by slow recovery, can be seen. Figure taken from Ref. [134].

The flasher LEDs contained by the IceCube DOMs allowed for monitoring of the water clarity. During the last hour every night, starting from January 2015, the flasher on the *Filea* DOM was activated for a monitoring run. The mean response of the other DOMs is displayed in Fig. 4.8. The sharp drops correspond to the last two visible in Fig. 4.7.

The drops correspond to times when the water pump stopped, e.g. due to interruption in the work of the generator. At the same time, the filters were getting clogged with black sludge. The sludge was confirmed to be the product of sulphate-reducing bacteria, dwelling in the sulphate-rich pit water and producing metal sulphides, which precipitate as black sediment, impacting water clarity.

The recovery of CHIPS-M took place in the beginning of June 2015. The underwater buoy holding the winch cable was localised, and the cable recovered and attached again to the winch assembly on the floating platform. The detector was then lifted to the surface, where the floating bags were attached, and then towed back to the shore. There, it was attached to the crane and slowly lifted, while the

water inside was drained. After returning to the shore, the structure was closely inspected.

The most important discovery was a tear in the liner at the vertical seam. It is suspected to have been caused by the sudden impact on the lake bottom during deployment, which would result in a sharp rise in the pressure inside. The hole is believed to be the sole cause of the light leak and the main reason for the deteriorating water clarity. Although the bacteria were being killed by the sterilising system and the black sludge was removed by the filters, the opening in the liner constantly allowed new organisms inside.

4.2 The Nikhef prototype unit

The success of the first CHIPS-M deployment opened the possibility to use it for testing prototype hardware, particularly PMTs and readout systems. At the same time, collaboration with KM3NeT was established, to probe the idea of using the components from the KM3NeT optical modules in future CHIPS detectors. The prototype detection unit was developed at the Dutch National Institute for Subatomic Physics (Nikhef) in Amsterdam. The desired form was a flat panel with multiple PMTs and a central readout, which could be used to tile the walls of a CHIPS detector. Due to the planar geometry, it was called a plane; or the Nikhef plane to distinguish it from a second prototype developed at University of Wisconsin in Madison.

The design of the plane underwent two major iterations. The first version utilised 31 PMTs enclosed between two $1\,\mathrm{m} \times 2\,\mathrm{m}$ acrylic sheets with acrylic spheres housing the PMTs. However, this design proved later to have major flaws and be inherently susceptible to leaks. The second iteration used PVC piping instead and was much more robust. Due to time and logistic constraints, it only contained 16 PMTs.

The following section briefly describes the KM3NeT experiment and the components of a Digital Optical Module (DOM), which were used in the prototype unit. Next the mechanical design of the two iterations of planes is presented, along with a report on the construction and testing process.

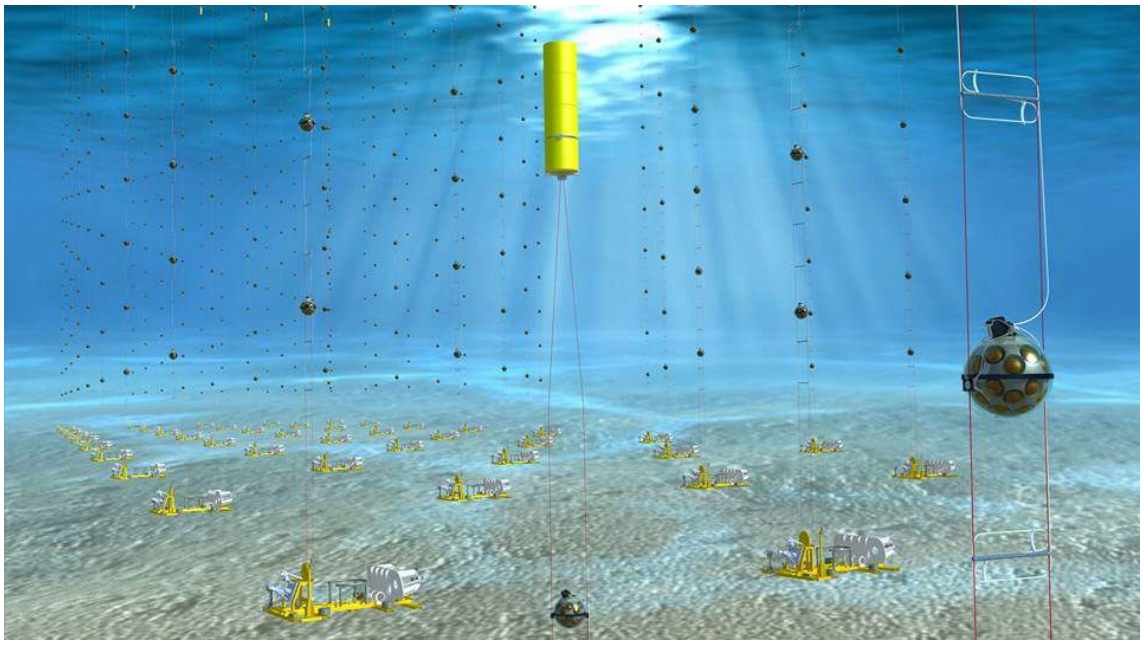

Figure 4.9: An artistic depiction of the KM3NeT detector at the bottom of the Mediterranean Sea. Image taken from Ref. [137].

4.2.1 The KM3NeT DOM

KM3NeT (Cubic Kilometre Neutrino Telescope) is a neutrino astrophysics and oscillation experiment [135, 136]. The experimental design is similar to that of IceCube, with strings of optical modules, however instead of in the Antarctic ice, it is located at the bottom of the Mediterranean Sea. It consists of two sub-experiments, organised in a single collaboration: KM3NeT/ARCA (Astroparticle Research with Cosmics in the Abyss), which is focused on detecting galactic and extragalactig ultra-high energy neutrinos, with a complementary sky view to IceCube; and KM3NeT/ORCA (Oscillation Research with Cosmics in the Abyss), which aims to measure the neutrino mass hierarchy through studies of atmospheric neutrino oscillations [136].

The detectors will consist of building blocks of 115 Detection Units (DUs), each of which contains a string of 18 optical modules (DOMs),[1] and a base container anchored at the seabed (Fig. 4.9). The DU base containers connect to an underwater

[1]This is in contrast to how the expression "detection unit" is used in the rest of this thesis to refer to a single planar optical module, or a plane.

electro-optic cable network, which provides power supply and communication with the shore.

The KM3NeT/ARCA detector will consist of two building blocks, deployed 100 km offshore from Portopalo di Capo Passero, Sicily, at the depth of 3500 m. The strings will be placed with a horizontal separation of 90 m, and the DOMs will be separated vertically by 36 m. This creates a sparsely instrumented Cherenkov detector, but with a very large volume, which is optimal for very high energy neutrino events that produce large amounts of light, but occur very rarely.

KM3NeT/ORCA will have one building block with strings placed 20 m apart and the DOMs every 9 m. It will be deployed at a location 40 km from the shore of Toulon, France, at a depth of 2500 m.

Currently the experiment is in Phase 1.0, during which 24 string DUs will be deployed at the ARCA site near Sicily (with 8 additional DUs with a different design and operated separately), and 7 denser strings at the ORCA site in France. The development of the mechanical structures and the readout systems has finished and the first strings are being deployed. In Phase 2.0, the collaboration aims to complete three full building blocks by 2020.

DOM components

The DOMs have a similar design to the IceCube ones, being encased in the same 17 inch glass spheres and containing a complete readout system. However, instead of using one large PMT, each KM3NeT DOM contains 31 smaller 3 inch PMTs. The use of multiple PMTs in a single module provides many advantages. With the inclusion of light collecting rings, the total photocathode area is three to four times higher compared to a single PMT [136]. The ability to record multiple hits gives also better track direction and time resolution.

Apart from the PMTs, a single DOM contains a set of electronics boards providing the necessary power, readout and communication capabilities, a calibration LED (the *Nanobeacon*) and a piezoelectric acoustic sensor acting as a sonar hydrophone to determine the position of the DOM (Fig. 4.10). The PMTs are arranged in a pattern of 5 rings and a single PMT pointing directly downwards. There are 19 PMTs in the lower hemisphere and 12 in the upper one, installed in a lightweight support

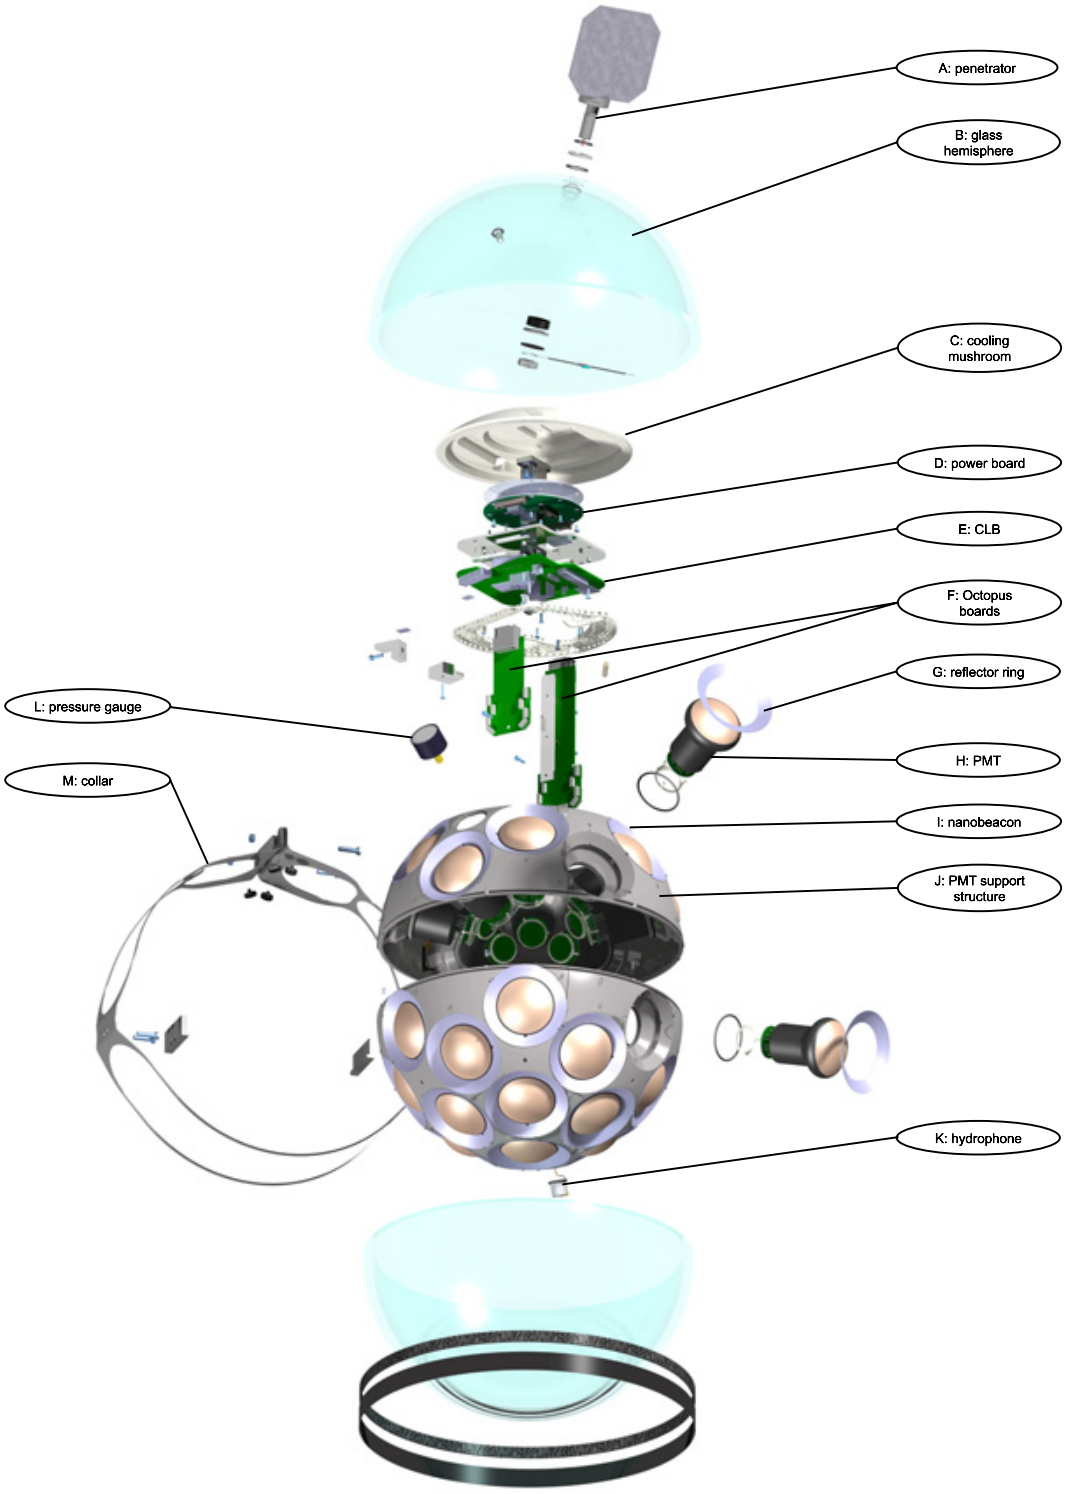

Figure 4.10: Design diagram showing the KM3NeT optical module in an exploded view. All the major components are shown and labelled. Image taken from Ref. [138].

structure manufactured with 3D printing. Reflective rings placed around each PMT face increase the effective photocathode area by 40 % [136]. The PMTs are coupled to the inner glass surface with an optical silicone gel.

The main electronics boards are mounted on the cooling mushroom – an aluminium structure in the top of the DOM, providing a heat sink and a large surface to exchange the heat through the glass with the surrounding water. Placed closest to the mushroom is the Power Board, which contains a suite of DC/DC converters, transforming the 12 V DC externally supplied power to the voltage levels required by all DOM components.

Mounted directly on top of the Power Board is the Central Logic Board (CLB), which is the heart of the DOM. It collects the output signals from the PMTs and converts them to digital hit information, controls the operation of all DOM components and communicates with the shore, sending data and the receiving control commands. The main chip on the CLB is a field programmable gate array (FPGA), which emulates a microprocessor, a set of Time-to-Digital Converters (TDCs) and the White Rabbit Ethernet protocol hardware.

Two Octopus boards are plugged in vertically to the CLB, using a central aluminium pillar extending from the cooling mushroom as mechanical support (Fig. 4.10). They act as a bridge between the CLB and the PMT base boards, with one serving each hemisphere.

Every PMT has a base board, which contains a Cockroft-Walton voltage multiplier and an integrated readout and a control circuit. The bases connect to the Octopus boards with cables which are printed copper paths on a short elastic kapton tape.

The DOM connects to the base container with two copper wires for power and two optic fibres for communication. They are contained in an oil-filled pressure-balanced cable and enter the glass sphere through a high-pressure penetrator.

Readout and operation

A photomultiplier tube is an instrument for detecting very low levels of light, down to a single photon. It consists of an evacuated glass tube, with photocathode, usually

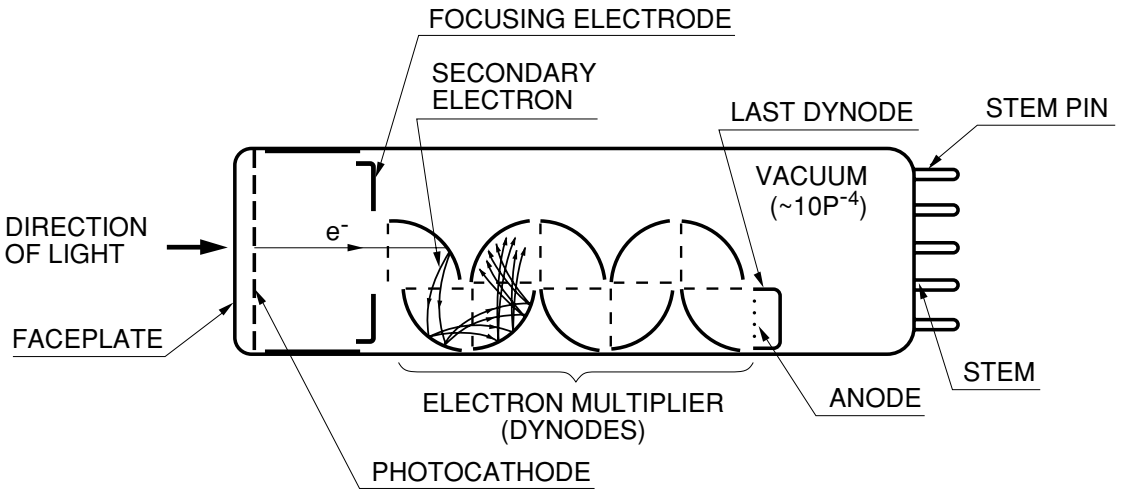

Figure 4.11: Schematic diagram illustrating the operation of a photomultiplier tube. Photons hitting the photocathode eject photoelectrons, which are multiplied by the dynode chain and collected by the anode. Image taken from Ref. [139].

deposited on the inside of the front window, a series of electrodes called dynodes and an anode (Fig. 4.11).

When a photon strikes the photocathode, it may cause the ejection of an electron from the photocathode, via the photoelectric effect. The electrodes are held at different electric potentials, so that there exist an electric field between each pair, from the photocathode all the way to the anode. The initial photoelectron will drift towards the first dynode, usually directed by an additional focusing electrode. As it accelerates in the electric field, it will reach a kinetic energy corresponding to the voltage difference, typically at the level of 100 eV and, after striking the first dynode, produce around 5 secondary electrons. The secondary electrons are accelerated towards the second dynode and the whole process repeats. The whole chain forms an electron multiplier, with a typical gain of 10^6 to 10^7. The final charge collected at the anode creates a pulse, which can be detected and measured.

KM3NeT uses PMTs from different manufacturers, but the ones borrowed for the prototype plane were the Hamamatsu R12199-02. This model has a 72 mm photocathode diameter, 10 dynodes, and a typical gain of 5×10^6 at a voltage difference between the anode and the photocathode of 1000 to 1400 V [140].

The PMT base contains a readout and control chip, which includes a pre-amplifier, and a comparator which determines when the amplified PMT signal is above a

predefined threshold voltage [141]. This information is sent out as a binary low voltage differential signal (LVDS). The control circuit communicates with the CLB via the I^2C protocol, and allows setting the high voltage and readout threshold levels, as well as switching the PMT on or off.

When the PMT registers a hit and the output is higher than the threshold, a logical one signal is sent on the LVDS lines to the Octopus board, for the duration of the pulse. The output is directed by the Octopus board to the CLB, where a time to digital converter (TDC) time-stamps the start of the pulse and records the duration, with a precision of 1 ns. The initial time is used as the hit time, and the duration, called time over threshold (ToT), reflects the hit pulse height. The ToT value is stored as an unsigned integer in a single byte of memory and therefore cannot be larger than 255 ns. In case of a longer pulse, it is automatically split into multiple entries.

The CLB stores this information for all PMT hits in local memory and sends the accumulated package at regular intervals, called timeslices, typically lasting for 100 ms. All hit data is sent to the shore, where the triggering to find coincidences and reduce the data rate is performed on a PC farm. Using multiple small PMTs with the ToT method affects the way the data is processed. Instead of trying to measure the number of photoelectrons at each PMT, hits are treated as binary events, and the amount of light is determined from hit multiplicities.

The CLB communicates with the on-shore DAQ via the Ethernet protocol using single mode optical fibres. To allow for a maximum bandwidth for a large number of DOMs and a minimal number of fibres in the long underwater cable, wavelength division multiplexing is used. Each DOM on a single string contains an optical transmitter using a different wavelength. At the base container of each string, all the signals are collected onto a single fibre. At the shore, a wavelength splitter is used to demultiplex the signals. The incoming slow control data is broadcast to all the DOMs using a common wavelength on a single fibre. Time synchronisation with the shore station and other DOMs is achieved with the modified White Rabbit protocol [102, 136].

4.2.2 First design iteration of the prototype plane

The first prototype plane was designed and constructed as part a semester-long Master's student project at Nikhef. The project was coordinated by P. Kooijman and technical advice was provided by electrical and mechanical engineers working at Nikhef. This section describes the first design iteration, but many of the solutions were also used in the second version.

The basic principle behind the prototype plane with KM3NeT hardware is very simple. A single DOM contains all the main components necessary for a self-contained unit of PMTs, which would be used to tile the walls of a CHIPS detector. The only important required modification is to place the PMTs and readout boards in a planar structure instead of the glass sphere and it can be operated just like the KM3NeT DOMs

Pressure vessel and the mechanical structure

The main design challenge is to create a large and cost-effective pressure vessel. The structure cannot be as compact as a KM3NeT DOM, because the PMTs have to form a plane. In addition, the optimal photocathode coverage is less than 10 %, meaning that a unit containing 31 3 inch PMTs should cover an area on the order of 2 m^2. The PMT bases and readout boards cannot be exposed to water, but using multiple pressure vessels (e.g. one for each PMT) would require many pressure feedthroughs for the cables, which creates more failure points, while adding to the cost and complexity of the design. Although the pressure at the bottom of the pit is only around 5 to 6 bar (corresponding to the 50 to 60 m depth), which is significantly smaller than at the bottom of the sea, water leaks are still a dangerous risk.

The glass surface of the PMTs could potentially be exposed to the water, as it should be able to withstand the pressure, provided there is a seal protecting the base board. However, in the KM3NeT readout system, the photocathode is at high voltage (while the anode is grounded), and it was believed that the high electric field could increase the dark rates. Initial tests showed that PMTs exposed to water at ground potential did indeed exhibit higher dark rates than ones that were insulated. Complete housing with transparent windows was chosen instead, which also simplified the design.

In the first version, the pressure vessel consisted of two acrylic sheets,[2] 2 m long and 1 m wide, with 32 acrylic hemispheres glued on each side to provide spherical housings for the PMTs (Fig. 4.12). The sheets were bolted together, with an o-ring along the perimeter providing a waterproof seal. The approximately 6 m long o-ring was made from a rubber cord, which was cut to size and glued with cyanoacrylate. Holes for the PMT housings and bolts as well as the grooves for the o-ring and PMT cables were milled out to order by a commercial company. The back side sheet was 10 mm thick, while the front one was 12 mm thick, to accommodate the o-ring and cable grooves.

The main readout boards were housed in an aluminium container, attached in the centre of the plane. A flat front cap (Fig. 4.12) and a ring in the back were connected with bolts, with o-rings outside the perimeter of the bolt holes providing a seal to the acrylic. A third tophat-shaped cap was bolted to tapped holes in the backside ring with another o-ring, completing the water-tight container (Fig. 4.14 bottom).

Acrylic was chosen as the material for the transparent windows because it is cheap and strong under compression (e.g. from outside pressure). It is also easy to glue to other acrylic components, creating watertight joints. In addition, acrylic has an index of refraction close to that of glass and water, minimising reflection losses. It was also considered a good choice for the flat sheets, due to being structurally strong compared to its weight, easy to glue with the hemispheres and relatively easy to machine. The 32 spheres were chosen for symmetry, to make the design simpler. Since there were only 31 PMTs, one of the spheres was used to place the Nanobeacon LED.

A frame made of aluminium L-beams surrounded the plane. Its function was to evenly distribute the load from the bolts along the o-ring, provide stiffness and mechanical support for the acrylic sheets and to offer attachment points for mounting in the detector.

[2]Also known by the brand names Plexiglas of Perspex

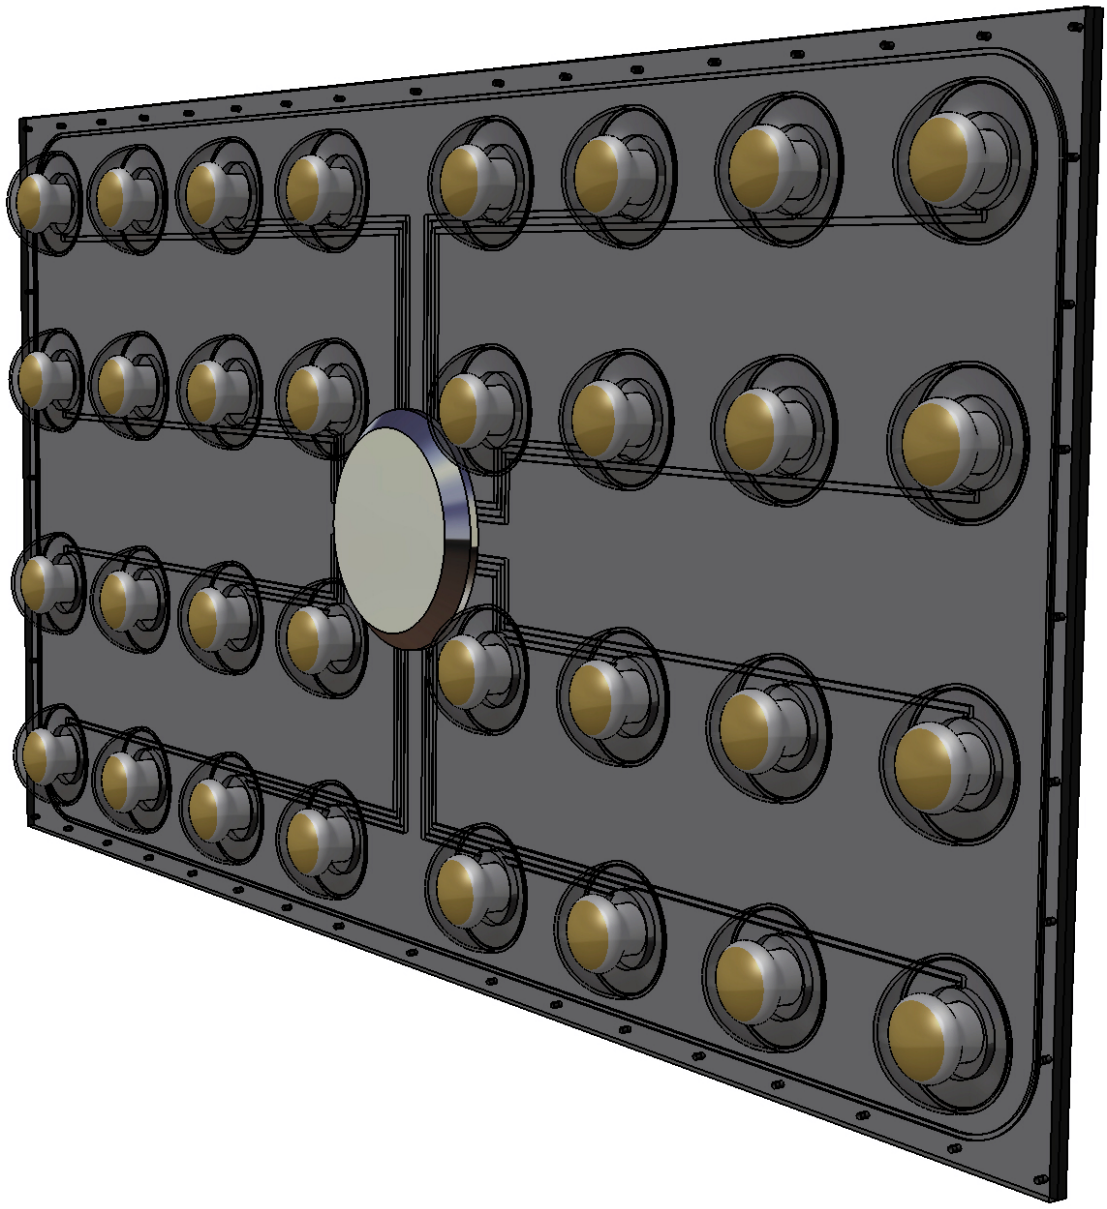

Figure 4.12: Design diagram of the first version of the prototype detection unit [142]. Front side of the plane is shown, with PMTs housed in acrylic spheres. Also visible are: the front cap of the aluminium electronics container, the cable grooves, and the o-ring groove around the perimeter of the plane. Missing from the diagram is the aluminium support frame. Although 32 PMTs are depicted, only 31 were used, with one of the hemispheres occupied by the calibration LED.

Accommodating the KM3NeT hardware

Wacker SILGEL 612 optical gel was used to couple the PMTs to the acrylic hemispheres. The diameter of the hemispheres was chosen to match the curvature of the PMT front face in order to minimise the amount of gel (Figs. 4.13a and b).

To accommodate the readout boards, the front cap of the aluminium container was milled out in a shape fitting the power board and the CLB, the same as in the KM3NeT cooling mushroom. The long Octopus boards extending vertically from the CLB are designed to fit in the limited space in the spherical DOM and to reach close to the PMTs. However, in a standalone container, they require too much space and force higher mass. For use in future planes, the boards will be redesigned to fit more compactly in a single cylinder.

The container had three openings: two 1/2 inch NPT tapped holes for attaching the cable glands acting as feedthroughs for a Cat5 cable delivering power and a fibre optic cable for communication; and an additional hole with fixings for nail clamps to attach a KF vacuum flange for leak testing.

The short kapton tape cables attached to the PMT bases are optimal for the compact assembly of the KM3NeT DOM, but much too short for the larger plane. Extension cables were initially planned to be made with twisted pair ribbon cable. However, the connectors used to connect the PMTs to the Octopus board are very small and soldering them to the wires was found to be very difficult. A decision was made to use the same technology of wire paths printed on kapton foil. This made the cables very compact and allowed the connectors to be surface mounted to pads on the printed circuit. Cables from two PMTs were placed in a single groove, facing each other with a grounded backside pad to avoid interference and channel cross-talk.

While the kapton cables worked well, their cost was very significant. Because they were printed on the original foil in a maze-like pattern (Fig. 4.2c), they had to be folded to form long straight sections, and the bends proved to be vulnerable to cracking, which destroyed several of the cables. For future models, regular cables or wires with bigger connectors are preferred.

The optical transceiver is not a fixed part of the CLB; a Small Format Pluggable (SFP) is used instead, with the appropriate slot mounted on the CLB. This makes

it possible to assign different wavelength transmitters on different DOMs, but also to use simpler components in CHIPS. In the prototype plane, a cheap SFP was installed, using two wavelengths for bidirectional communication on a single optic fibre.

The Nanobeacon circuit with the calibration LED was included in the plane, taking one of the free acrylic spheres. In a detector with multiple units, the LED would be used to illuminate PMTs in other planes, but it could still be used for the calibration of one plane, provided the geometry was taken into account properly. The hydrophone was not included in the prototype plane, since it would not work without an acoustic emitter (which in KM3NeT is located in the base containers).

4.2.3 Construction and testing of the first plane

Before the assembly started, techniques for gluing the hemispheres to the acrylic sheet in a way that creates a tight and structurally strong bond had been tested on small demonstrators in a pressure chamber. A combination of groove depth and width was found that withstood a pressure of more than 12 bar.

After finishing gluing the hemispheres to the two sheets, but before placing the PMTs, the plane was clamped with the o-ring in place and tested for leaks with a helium sniffer. A roughing pump and the sniffer were connected with a KF vacuum flange. Helium from a bottle was then sprayed around potential leak points and the leak rate read out from the sniffing machine. It was found that the acrylic hemispheres joints were very leak-tight, but the o-ring on the perimeter of the plane was not. A bigger diameter o-ring was ordered and the groove polished to remove ridges left after the milling process.

Time constraints required further assembly without testing and no more helium tests were performed with the PMTs in place to avoid helium poisoning. This occurs when helium atoms penetrate the PMT glass and contaminate the vacuum inside. When hit by a stream of multiplied electrons, the helium can get ionised. The helium ion is then accelerated back towards the photocathode, where it originates another electron multiplication, manifested as an afterpulse. In the extreme case when the helium concentration is high enough, this can create a self-sustaining reaction and

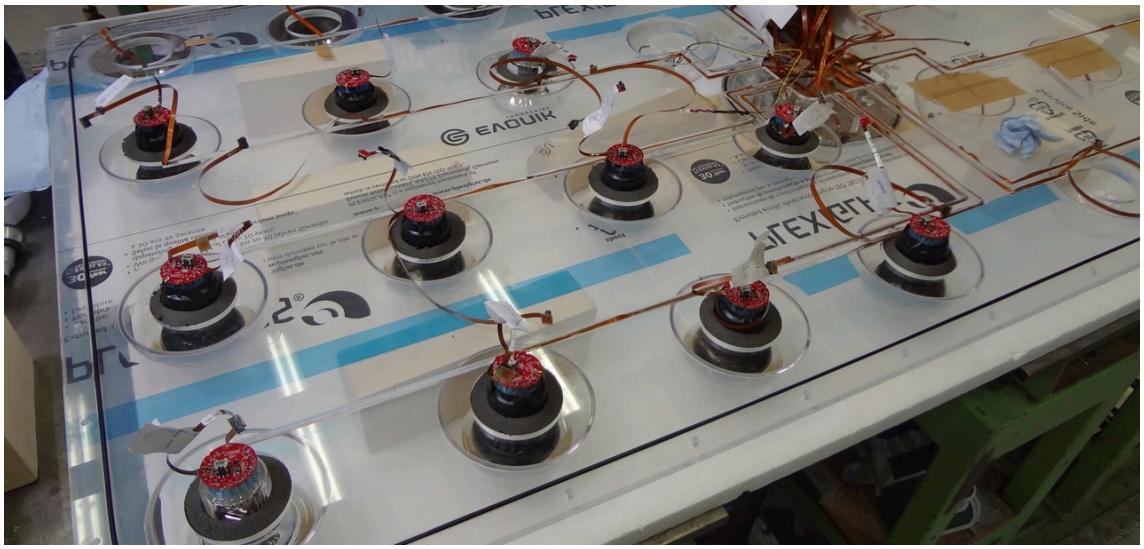

(a)

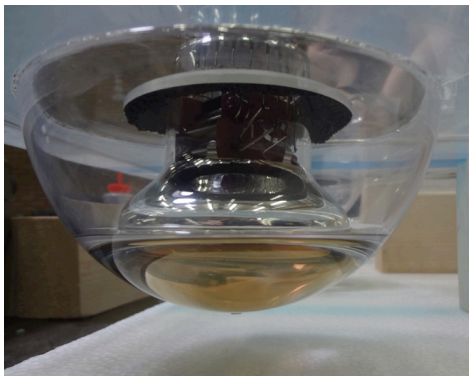

(b)

(c)

Figure 4.13: Construction of the first version of the Nikhef plane. (a) PMTs being installed in the front sheet hemispheres. Kapton cables already glued in the grooves are visible, as is the o-ring. (b) An installed PMT, coupled to the acrylic with optical gel. (c) The kapton extension cables after attaching connectors but before the folding.

the PMT is unusable. Since the PMTs used in the plane were on loan, the possibility of helium poisoning was highly undesirable.

The PMTs were installed in the front part sheet hemispheres, placed one by one right after the appropriate amount of optical gel was poured. After curing, the gel works as adhesive to keep the PMTs in place, but to limit potential lateral movement, foam rings were added to keep the tubes in a tight fit in the holes (Figs. 4.13a and b).

A wooden shipping crate, built for the plane, was lined with black neoprene foam to double as a dark box. Neoprene is very good at absorbing light and limiting reflections, especially at an angle. The plane was tested in the dark box to ensure that all the PMTs perform nominally.

The plane was shipped with air freight to Minnesota, where extensive testing was performed before installation in CHIPS-M. In particular, the new o-ring has been installed and the structure was tested for leaks. Without the helium sniffer, this was done by pumping out the air inside and monitoring the pressure to observe vacuum decay. Only a pressure difference of approximately half an atmosphere (0.5 bar) was created, to avoid potential small bubbles of air trapped in the cured optical gel from expanding and spoiling the transparency. A digital sensor with a thermometer and barometer was installed, allowing for a precise online pressure measurement, corrected for temperature changes (Fig. 4.14 bottom).

During testing it was discovered that despite the new o-ring, the plane had a significant leak rate. After repeating the test in a water tub, the leak rate of water was measured to be 2.5×10^{-4} mbar l/s. With the pressure difference at the bottom of the pit being approximately ten times higher, such a leak would lead to macroscopic amounts of water flooding the plane within hours, and was not acceptable.

It was suspected that cracks may have developed in the joints between the hemispheres and the acrylic sheet. The pressure decay test was repeated in the air, with only one surface covered with a layer of water. Small leak rates are proportional to viscosity and water is approximately 100 times more viscous than air. Covering only some parts in water allows determining if the leak is there by observing the rate of pressure decay. Indeed, the observed leak rate had fallen drastically compared to readings without the water layer, confirming that the new o-ring performed well and there was a leak in one of the hemispheres. By dividing the area covered in water with small levees made out of tacky tape, the main culprits were identified. The leaks were patched with an adhesive for acrylic, containing methyl methacrylate, which dissolves the acrylic and creates a continuous bond.

However, as the tests were repeated after moving, reorienting, or finally installing the plane in the detector, new leaks were discovered. The root cause of the problem was judged to be the inherent flexibility of the acrylic sheets, which could not be

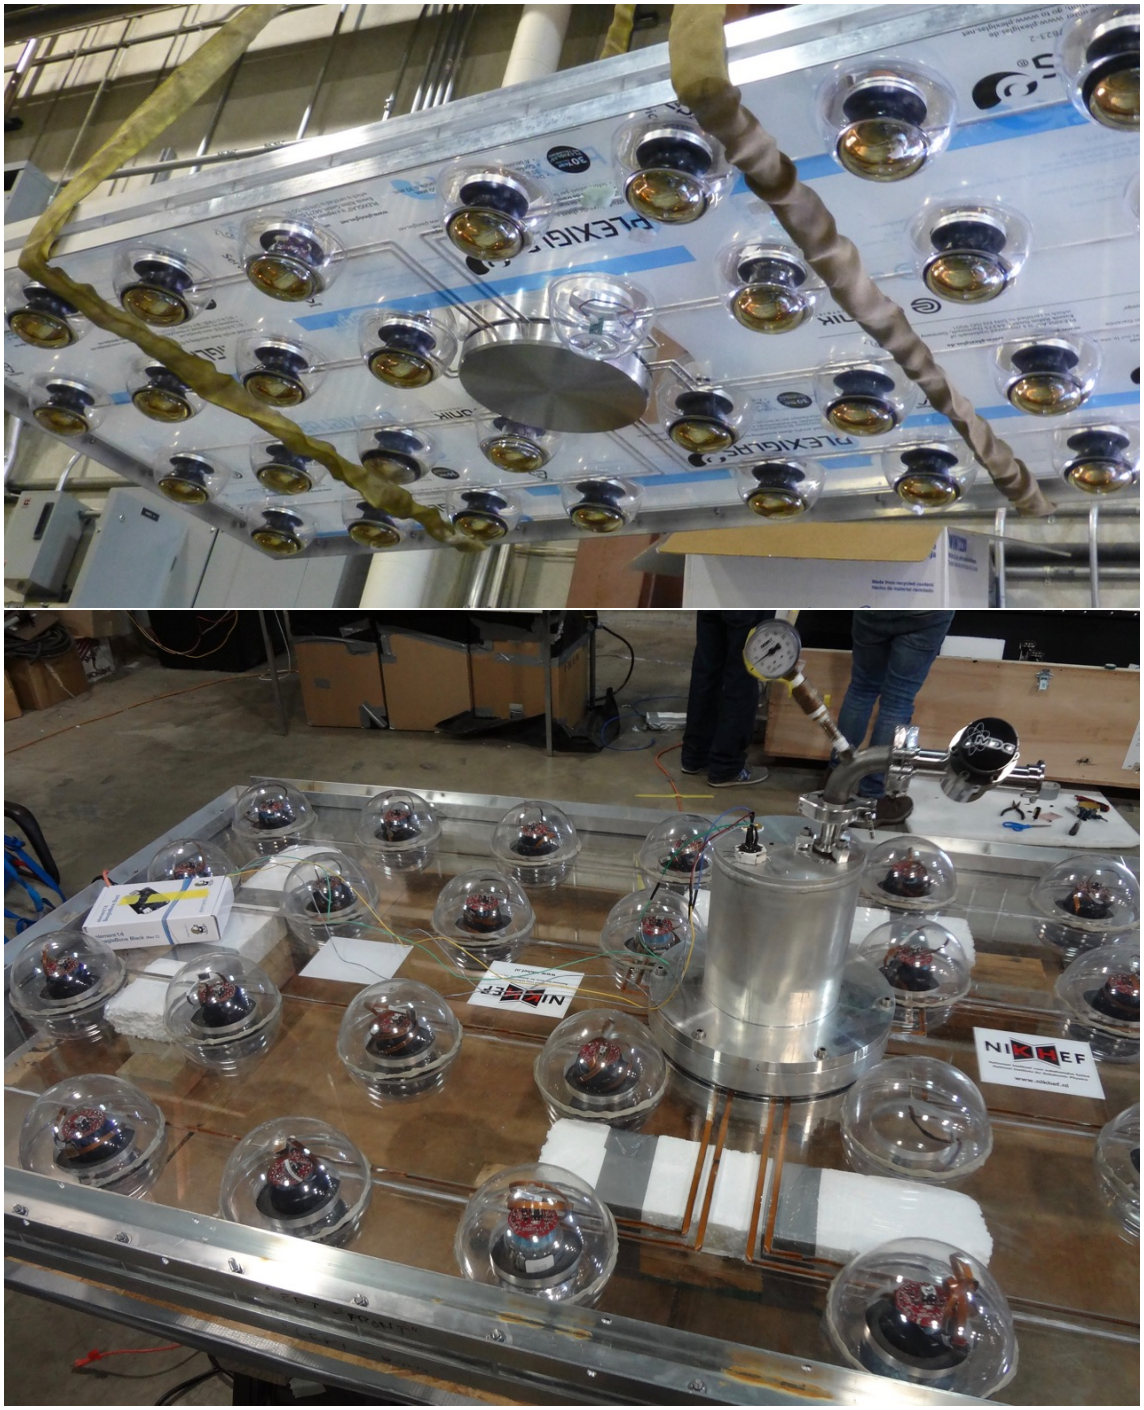

Figure 4.14: Top: The assembled plane arriving in Minnesota. Bottom: Leak testing by observing pressure decay. The plane was partially evacuated, and a digital sensor reads out the pressure and temperature inside the electronics cylinder.

(a) (b)

Figure 4.15: Second version of the Nikhef plane. (a) The finished plane with 16 PMTs and a calibration LED. (b) The black PVC ring made to adapt the electronics container to the pipe structure. The CLB mounted in the aluminium cap, with two Octopus boards plugged in, is visible.

remedied by the aluminium frame and was causing large stress on the hemisphere joints every time the unit was moved.

4.2.4 Second version of the plane

Since other CHIPS-M instruments were also suffering from delays due to unexpected issues, the deployment attempt was postponed by a month. In this time a new design was completed and the second version of the plane constructed. The internal components were the same, but the pressure vessel concept was completely overhauled, replacing the acrylic sheets with high pressure PVC piping. Due to time limits and difficulties with transport from Amsterdam to Minnesota on short notice, the second version contained only 16 PMTs.

The main problem with the first plane design was the inherent flexibility of the acrylic sheets, which put large stresses on the glued hemisphere joints. In addition, after installation of the PMTs it could only be leak tested by observing low

underpressure decay, which makes it difficult and time consuming to find specific leaks.

Constructing the pressure vessel out of high pressure PVC piping offers many advantages. First, it solves the flexibility problem by placing each PMT in a separate pipe stub with a window covering one end, and hence virtually decoupled from the stresses acting on the whole structure. Because the components are commonly used in the industry, they are significantly cheaper than the custom-cut acrylic sheet. The structure is easy to assemble, with glued joints becoming essentially fused and very leak-tight. The design is also much more flexible, allowing for adjustment of the size or shape of individual planes, e.g. to fit in a corner on the circular endcap. The entire plane is also lighter and does not require a metal support frame. Finally, the pipes make it easier to route signal cables inside, removing the limitation of shallow grooves.

If the attachment of the window to the PMT housing is secure enough, the pipe design also makes it possible to test for leaks with overpressure inside, since the pipes are designed to withstand it. In this case, the pressure difference between inside and outside can be made higher, and leaks are easier to find by using a soapy liquid producing visible bubbles. In the prototype plane, the acrylic windows were attached with the PVC glue, forming a bond not strong enough for very high overpressure. However, the bubble detection method was used successfully to eliminate all leaks.

The new design has also some disadvantages. The PVC pipes are joined with glue, meaning that they cannot be disassembled non-destructively or even adjusted after the assembly. The method of attaching the windows to the pipes housing the PMTs becomes a crucial design challenge. It should be reversible to allow inspecting or replacing the PMTs, but robust enough not to introduce potential leak points. Also, the pipes are designed for high pressure on the inside, not from the outside, which presents the risk of buckling. In practice, this should not be a concern for short pipe sections with small diameter and thick walls, but it must be taken into account during design by making sure that the dimensions satisfy design safety margins.

The second version of the Nikhef plane kept most of the original components. All the electronics boards, the kapton cables and the electronics container were reused. To simplify the logistics, 16 new PMTs were used, and the 31 old ones from the first plane version were returned to Nikhef. The aluminium cylinder was accommodated

into the pipe structure by placing it on a PVC ring with four holes, where the pipe connections to the structure had been glued (Fig. 4.15b).

The plane construction started at Nikhef. The acrylic windows were first glued to wide pipe couplings, and then PMTs were placed and coupled to the windows with the optical gel. Straight sections of the main structure, corresponding to rows of PMTs in the 4×4 array, were glued together. The components were then transported to Soudan in regular checked luggage. There, the full structure was assembled, the cables were routed through the pipes and connected to PMTs, and the PMT assemblies were glued into the main structure. Finally, the signal cables were connected to the Octopus boards and electronics the container was bolted, clamping the PVC ring.

The new plane was tested for leaks by pressurising the inside volume and applying leak detection liquid to the glued joints, which are the most likely leak sites. Some of the connections between the PMT windows and the PVC were indeed found to leak and were sealed with a bitumen tape. The tape, which is also used in KM3NeT DOMs, is very adhesive and viscous, creating a water- and air-tight seal. Further testing did not reveal any more leaks.

After installing the plane in the detector, it was discovered that two PMTs became disconnected from the CLB. After inspection of the electronics container, it was determined that this must have been caused by the decoupling of the connectors between the PMT base and the extension cable. It is a known weakness of this connector type that they can disengage easily, and the glued attachment of the acrylic windows made it impossible to open the plane and reconnect the cables. In the final plane design, this issue will be solved by incorporating a detachable fastening method for the acrylic windows and by using more robust 8P8C connectors with regular Cat5 cables.

4.3 The second testing campaign

The second deployment of CHIPS-M included the Nikhef plane, but also a second prototype unit, the Madison plane, as well as several upgrades to the detector structure and operation procedures.

The main goals of the second test run were:

- Testing the two prototype detection units, in particular:
  - The robustness of the mechanical design in underwater conditions;
  - The ability to use the on-board readout system to run and record data with the PMTs;
  - On-shore data acquisition, storage and processing;
  - Use of led flashers for time calibration, between PMTs and different modules.
- Validation of the fixes to issues encountered during the first deployment, especially maintaining the light tightness.
- Gathering cosmic muon data without a light leak and water clarity issues.

4.3.1 Changes from the first test

Tested instruments

In addition to the Nikhef plane, a second prototype unit, constructed at University of Wisconsin in Madison, was tested during the second test campaign of CHIPS-M (Fig. 4.16b). The Madison plane utilised a readout chip called PARISROC (PMT Array Read-Out Chip), which can collect signals from up to 16 PMTs and digitise the waveforms [143]. A Raspberry Pi single-board computer controlled the chip, read out the digitised data and sent it to the on-shore DAQ PC. The plane used 16 3 inch PMTs, which were donated by the decommissioned NEMO-3 experiment. CHIPS has received 400 of those PMTs in total, therefore it is worthwhile to try to incorporate them into an instrument for CHIPS-10.

The environmental sensors had also been upgraded and featured a new design, with aluminium cylinder bodies and an acrylic hemisphere on one end of the cylinder, acting as a window for the camera and LED lights. This time, three vessels were installed: one remained outside, and the other two monitored the inside of the detector; one was located in the middle of the bottom truss and pointed upwards, while the second one was placed on one of the Unistrut column to overlook the Nikhef plane PMTs and the vertical liner seam.

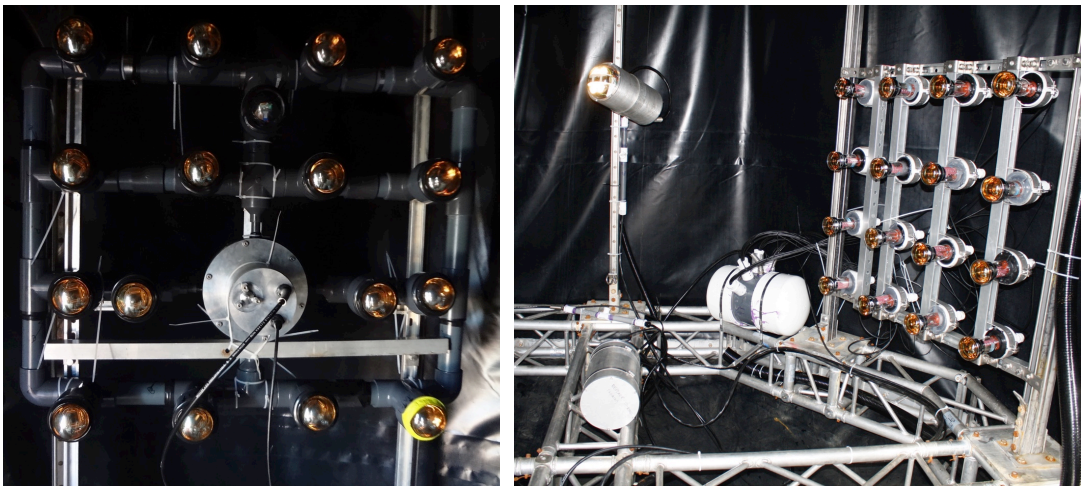

(a) (b)

Figure 4.16: The prototype detection units in CHIPS-M. (a) The second version of the Nikhef plane, facing upstream in the beam. (b) The Madison plane with NEMO-3 PMTs and the PARISROC readout contained in a PVC pressure vessel. In the lower left, the communications hub container, and in the upper left, one of the new environmental vessels, are visible.

The Madison plane utilised one of the DOM cables for high voltage power supply, while the Nikhef plane and the environmental vessels shared the triple Cat5 cable still left in the umbilical. All three instruments used optic fibres for communication. Since the fibres were delivered on a single cable, a hub container was added inside CHIPS-M (Fig. 4.16b). It contained the fibre fanout, two media converters to pass the optical signal to copper Ethernet cables for the Madison plane and environmental vessels, as well as DC/DC power converters to supply the correct voltage to the Nikhef plane and the sensors.

With one of the IceCube DOM cables used for powering the Madison plane, only four DOMs were installed in CHIPS-M. These were the same four that had continued to collect data during the first testing run, and they were mounted in the same spot inside the detector. The mirrors, which had been used previously with the fifth DOM, were removed to make space for the new prototype units. The cable previously connecting the *Angkor Wat* and was now used for the *Filea* DOM. While it seemed to had caused communication issues before, it performed well in later tests, and was still used.

Upgrades to CHIPS-M

The main CHIPS-M structure remained the same, but several modifications have been implemented.

The vertical seam was made possible to bolt from the outside, by introducing a fixed inside batten bar with tapped holes. This allowed for easy access to the inside of the detector, even moments before the final deployment.

Instead of filtering the water during filling, which limited the filling rate, the water has been preprocessed and stored in an outdoor swimming pool, acting as the pure water tank. During deployment, CHIPS-M was filled straight from the tank, which sped up the process significantly. The small water filling pipes (3/4 inch diameter) have been replaced with 3 inch hoses, additionally speeding up the filling process. Additional air venting caps in the top of the detector were added, to aid with the fast filling. They were left open during deployment and then closed when the top of the detector reached the water line.

The cross-link cables inside the structure, whose function was to prevent twisting have been removed, as they impacted the installation of new instruments.

All the pipes and cables in the umbilical were kept, with one of the IceCube DOM cables serving as power supply to the new Madison prototype plane. The only addition was an optical fibre cable. It contained six single-mode optical fibres, and was armoured with a corrugated steel casing. Three of them were used for communication with the Nikhef plane, the Madison plane and the environmental sensors. However, the fibre cable was not added to the umbilical bundle, already well contained in the conduit, and was submerged separately.

4.3.2 Deployment and operation

The second deployment of CHIPS-M had initially been planned for late August 2015. However, delays with assembly and testing of all of the instruments made it impossible to reach that goal. The deployment was postponed for a month, to the end of September, giving more time to fix the issues and reorganise schedules.

As the Nikhef plane design had been judged inherently faulty, it was abandoned and a new version was developed over that time, as described in Sec. 4.2.4. Other components, including the environmental sensor vessels, the electronics container for the Madison plane and the central communications hub container, were extensively tested for leak-tightness.

In contrast to the first deployment in 2014, the CHIPS-M structure was transported to the Wentworth Pit empty, before installing the instrumentation. After the final closing of the detector, the Nikhef plane and the IceCube DOMs were then used for light tightness testing. During tests at night with a flashlight and during day with a portable cover, a few small light leaks in the liner were discovered. Most of them were visually imperceptible and contributed very little to PMT rates inside CHIPS-M. All the leaks were patched with light-tight adhesive tape, until no difference between night and daylight rates could be observed.

The deployment proceeded on the 30th of August 2015. The procedure was very similar to the first deployment: the detector was lifted with a crane and slowly lowered into the lake, while being filled with water. The main difference was the use of stored water and large hoses, which resulted in a significant time decrease in filling, from almost 10.5 hours to barely over one hour. Afterwards, CHIPS-M was towed to the deployment location, and submerged from the floating platform until it reached the bottom. The lowering was conducted much slower than previously, in order to avoid an impact on the bottom of the pit, which is believed to have caused the liner tear a year earlier.

Soon after the operation was finished, the prototype planes were switched on and started recording data. The observed rates, especially during sunset and the following sunrise, confirmed that there were no observable light leaks. By the second day after the deployment, the Madison plane started experiencing connectivity issues and later stopped responding at all. Also attempts to communicate with the environmental sensors (which were nominally switched off during data taking) proved unsuccessful. Finally, after two days of running, the Nikhef plane ceased communication as well. The culprit was believed to be the fanout box, which was the component most vulnerable to leaks as it contained six separate cable glands and it also controlled access to three separate instruments.

At the same time, two of the DOMs were experiencing unrelated problems, likely caused by the flooding of underwater connectors, as they exhibited the same symptoms as the veto DOM during the first deployment.

The detector stayed under water for a month, with the remaining DOMs monitoring the water clarity, and then it was recovered again. The Nikhef plane was confirmed to have remained dry and was transported to Madison. Later test runs in a dark box reaffirmed that all the PMTs and readout were fully functional.

Chapter 5

Analysis of the cosmic muon data

The deployed prototype plane ran and recorded data for two days, detecting Cherenkov light from secondary cosmic ray muons, passing through CHIPS-M. Analysis of this data allows for an almost direct measurement of the muon intensity at the bottom of the Wentworth Pit. Such information is very valuable for the design of future detector modules such as CHIPS-10, particularly for determining the cosmic veto requirements. In addition, the analysis serves as a validation of the simulation software. The process also highlights issues and difficulties with interpreting the data and helps to prepare future solutions.

This chapter first describes the processing of raw data and gives an overview of the event reconstruction procedure. Next, a Monte Carlo simulation of CHIPS-M is presented, along with the cosmic muon ray generation, and a discussion of water attenuation properties. Subsequent sections explore the specific points of the analysis, comparing the data with simulation predictions, in an attempt to limit and understand the disagreements. The final selection process is used to estimate the total rate of cosmic muons passing through CHIPS-M and, finally, to predict the rate and effective dead time in CHIPS-10.

5.1 Data acquisition and event reconstruction

5.1.1 Raw data acquisition and storage

The Nikhef plane functions in a similar way to the KM3NeT DOMs. While it is running, the CLB on board collects PMT hit data, and sends them to the on-shore data acquisition centre at regular intervals. All hits are recorded and any triggering can be performed on shore.

The setup with a single plane in CHIPS-M is much simpler than a whole KM3NeT string or a full CHIPS-10 module, since it only requires communication with one CLB. The on-shore DAQ setup for CHIPS-M consisted of a single Linux PC with a network card, which contained an SFP slot with the optical transceiver connected to the optic fibre.

One of the programs running on the PC was a slow control GUI (graphical user interface). It communicated with the CLB allowing the user to set up and run the plane. Each PMT can be configured with individual values of high voltage and readout threshold, and switched on independently from the other ones. The LED flasher can also be set up from this program, and the plane can be switched to running mode, where all the PMTs are on and the CLB relays the data to shore.

The fast data with hit information comes from the CLB in UDP packets and is handled by a chain of DAQ processes. They are set up and started by a test script created for single DOM testing, which is suitable for running with a single plane. The data acquisition chain was designed to operate with the whole KM3NeT detector, containing many strings of numerous DOMs. In such operation, the processes are distributed among multiple computers. With only one CLB, this architecture is redundant, but for simplicity the data flow was unchanged, with the individual components running all on a single PC.

The data packets arrive periodically and contain the readout from the plane during one timeslice, which lasts for 100 ms. The timeslices are collected by the *DataQueue* process. In a run with multiple DOMs, the *DataQueues* are responsible for passing the data to the next stage (the *DataFilter*) synchronously, so that the whole detector state during one timeslice can be analysed. The *DataFilter* collects the timeslices from all *DataQueue* instances, and performs triggering to select hit

and DOM coincidences. When running with the Nikhef plane, there was only one *DataQueue* and the triggering was skipped, so the process is almost trivial and all hits are passed on. The last step is the *DataWriter*, which stores the information received from the *DataFilter* to ROOT output files, organised in a data structure specific to the KM3NeT software framework.

The run sequence was set up to collect 1 hour long data-taking runs continuously, almost immediately after the deployment of CHIPS-M.

In addition to the actual hit data, the CLB sends packets with summary information about the timeslices, including the CLB clock time and hit counts in the PMTs. These summary frames can be captured with a tool called *SwissKnife*, which was used to display the monitoring data in real time and store it in a text file.

5.1.2 Recorded data sample

Overall, 37 separate runs were recorded, with the total run time of 35.6 hours. The sample is divided into two superruns, between which the plane was switched off and the CLB rebooted. No differences due to the rebooting are expected, and the superruns are kept separate only because their relative hit times are not synchronised. This is because in the absence of the White Rabbit system, the CLB clock was synchronised with the DAQ PC manually.

The first superrun includes 15.3 h of data, with a few gaps between individual runs due to a small oversight. The second superrun lasts for 20.4 hours and includes three calibration runs, where an LED located in one of the IceCube DOMs was active, flashing with a frequency of 4.8 Hz. Two runs were recorded with the *Haifa* DOM flashing and one with the *Munich 1972* DOM. Each of these runs was 1 hour long.

Some tests performed during the data taking required connection to one of the environmental sensors. A network switch on board of its vessel had non-masked LED indicators, whose light shone directly onto the plane PMTs. At the first time the environmental suite was powered the plane was switched off, but during the second attempt, the PMTs were left on. For the approximately 5 minutes the LED was on, the plane events are dominated by 2 hit coincidences and the data from this period, occurring at the end of the second calibration run, is discarded.

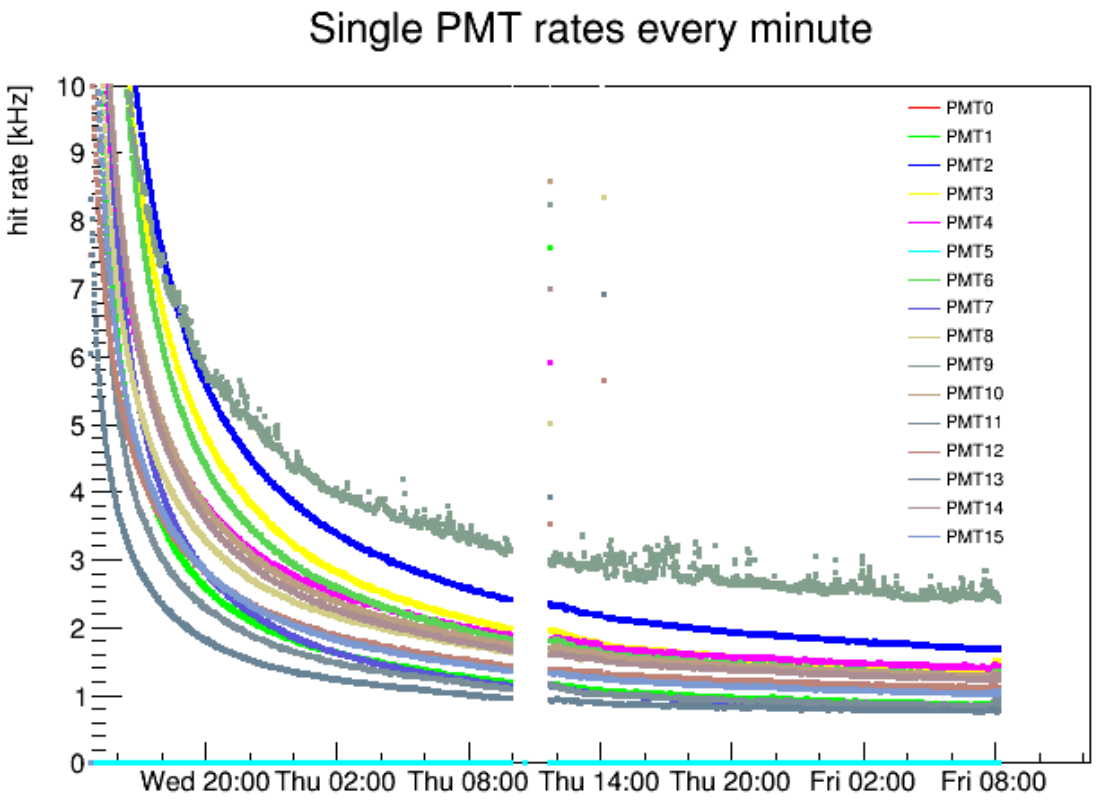

Figure 5.1: Raw hit rates of individual PMTs in the Nikhef plane, averaged every minute. The recording starts almost immediately after the deployment, at 14:48 on the 30th September 2015 (Wednesday).

Figure 5.1 shows the raw hit rates, averaged every minute, for individual PMTs, extracted from the monitoring data. Just after deployment and switching on, the PMT dark rates were extremely high. This was caused by accidental exposure of the PMTs to the sunlight, after opening small air-venting caps in the top of CHIPS-M, while the plane was performing a last-minute check of the light-tightness of the detector. Such an exposure causes the excitation of the photocathode, which then relaxes while in the darkness. As can be seen in the figure, the initial drop in single rates was very rapid, and by the end of the two day period, most PMTs achieved a rate of approximately 0.8 to 1.5 kHz. Later tests, performed with the recovered plane in a dark box in Madison, suggest that the dark rate can decrease to 300 Hz if left in the darkness for a sufficient time.

The smooth decrease of PMT rates (except for PMT 9, which was exceptionally noisy), particularly during sunset and sunrise, confirms the light-tightness of the

detector. The few outlier points are artefacts of the averaging procedure or correspond to the LED light from the network switch.

5.1.3 Pre-event file conversion

The readout system used for recording Nikhef plane data stores the information on every recorded PMT hit. While most of the single hits correspond to dark current noise in the PMTs, some of them come from the Cherenkov light produced by cosmic ray muons passing through the detector. The latter ones are found by requiring coincidences, i.e. multiple hits occurring in a short time interval. These collections of hits are called events, with the assumption that a single event corresponds to a single charged track in the detector.

The first step is the extraction of raw hit data from the output files. The hit information stored by the DAQ run scripts is organised into data structures designed to work with the KM3NeT software framework. Because this analysis was developed independently and deals with only 14 PMTs in a single plane, the information is converted to a more accessible data format.

The code to extract the hit information is based on the *MRunAnalyzer* program. This tool was developed at KM3NeT to produce summaries of test runs for single DOMs. The code was modified by the author to extract the raw hit data and store it in a flat ROOT tree format.

This step also includes initial filtering, where only hits which are within 100 ns of another one are saved. This ensures that all hits which could be part of an event shorter than 100 ns (significantly more than the typical duration of 10 to 20 ns) are kept, while dramatically decreasing the total file size.

Another function performed at this stage is the processing of very long hits. Pulses with time over threshold longer than 255 ns are split in parts by the CLB, to save disk space by storing the ToT as an 8 bit value. The code looks for continuations of hits with ToT equal to 255 ns and combines them into a single entry.

The output files, called *pre-event* files, are used as input to the coincidence finding code. Each entry in the chronologically ordered tree contains the hit time (stored as

a timeslice time-stamp and the hit time during the timeslice in ns), the time over threshold (ToT) in ns, and the PMT index.

5.1.4 Event reconstruction

The formation of events consists of multiple steps, whose parameters were adjusted iteratively. This section gives an overview of the entire procedure, while later sections provide the justification, by comparing data to the output of a detector simulation. All the following code has been developed from scratch for the needs of this analysis.

First, a correction is applied to the hit times to account for timing offsets between PMTs. The corrections are evaluated by extracting the intrinsic offsets from events caused by the LED flasher in calibration runs, and the full procedure is described in Sec. 5.3.

Then, a time over threshold cut is applied to reduce the electronic noise background. Hits with ToT below the cut value are discarded. The typical ToT value at the 1 photoelectron (PE) peak in the plane PMTs is 32 ns, and the ToT cut value is nominally set at 12 ns (Sec. 5.4.1).

Next, the algorithm looks for hits that satisfy the event definition. First, all hits fitting inside a fixed time window are collected into a single event. Specifically, the algorithm loops over the hits and checks if the time difference between the current and previous hit is smaller than the window size. If this is true, then those two hits are considered part of an event, and every subsequent hit falling inside the window is added as well.

After all such hits are identified, the event is saved to the output file, and the whole process repeats. Consecutive hits from the same PMTs are discarded to avoid including fast after-pulses. A number of useful event variables, such as the average hit time, total ToT or the event duration, are calculated and stored with the event.

The list of hits in the pre-event files is ordered according to the original hit time, so the chronological order it is not preserved after applying the time corrections. Because of that, it is possible that some hits which should be included in the event will be omitted. However, the time corrections are on the order of 1 to 2 ns and a

reasonable time window is around 30 to 50 ns, with most hits occurring in the first 15 ns, so this effect is negligible. The event duration is discussed in Sec. 5.4.2.

In addition to the fixed time window, a simple clustering algorithm is applied to find shorter groups of hits within the event. For a given maximum time distance between two hits (e.g. 10 ns), the first cluster of at least two hits satisfying this constraint is identified, and all the other hits are discarded. In practice, this treatment excludes mostly single hits occurring a significant time before or after all the other ones.

The clustering could be used as an alternative event definition instead of the fixed window. However, for historical reasons, it is only applied afterwards and the included hits are marked with a flag, while all original hits are kept in the event.

The event forming program can convert one or more pre-event files into a single *event file*, with correct hit time-stamps relative to the run start and a variable counting the total run time.

Figure 5.2 shows the rate of events with at least 3 hits in a 50 ns long time window and in clusters with maximum separation of 10 ns, during the whole data taking period. The spikes in the event rate at the beginning of the second superrun correspond to the calibration runs with the *Haifa* DOM and the short period of the environmental sensor LED flooding the plane with light. The flashing pulses from the calibration LED are only occurring every 0.21 s, but they are each followed by a barrage of low-multiplicity events, possibly caused by afterpulses, reflected light and photocathode de-excitation. The smaller increase in rates halfway through the second superrun corresponds to the calibration run with the *Munich 1972* DOM.

The event rate is slightly higher in the very beginning of the run, due to the very large dark rates contributing to random coincidences, even of three or more PMTs. This contribution diminishes quickly, and is much less pronounced when imposing the additional clustering requirement. The total event rate continues to decrease slowly during the run period, but is relatively stable overall, at the level of around 70 Hz.

A simple event display was written to visualise single events and allow for manual scanning of the data with arbitrary cuts. An example event is shown on Fig. 5.3.

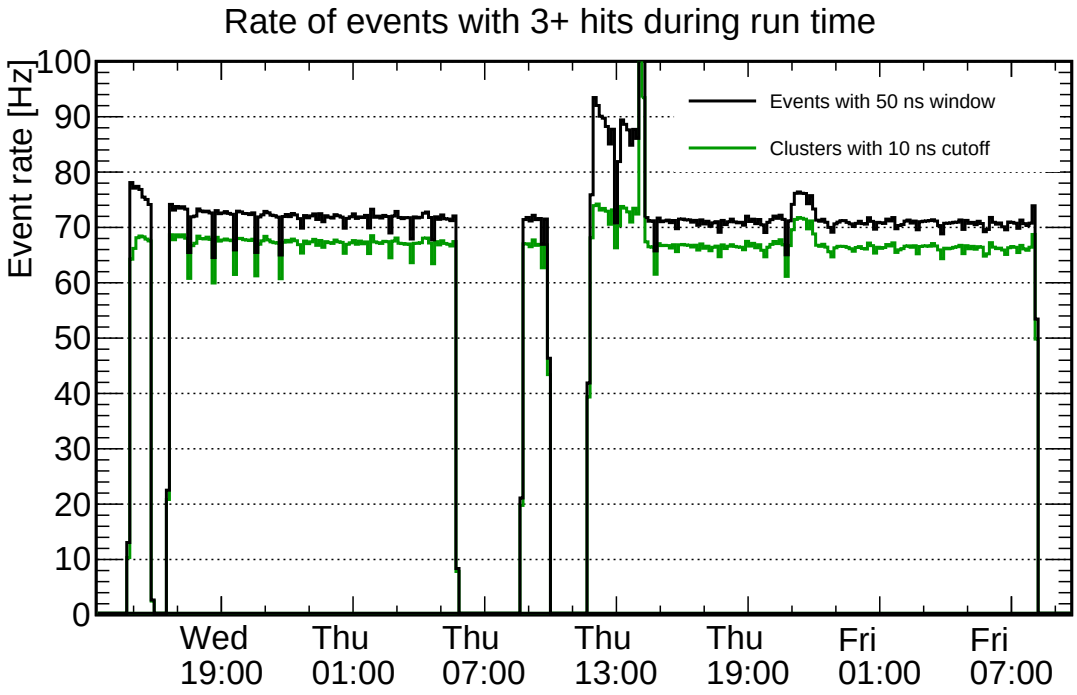

Figure 5.2: Rate of events registered by the Nikhef plane as a function of run time, starting from 2015-09-30 (Wed) at 14:48. Solid lines show the rates for all events with at least three hits within a 50 ns window (black) and with an additional requirement of maximum 10 ns separation between hits. Time calibration with the *Haifa* flasher was applied during the event formation, as well as a ToT cut at 12 ns.

The coloured circles represent hit PMTs, with the hit time signified by the colour and time over threshold of the hit by the circle size.

5.2 Monte Carlo simulation of the cosmic data

The CHIPS detector Monte Carlo simulation was used with the CRY package to simulate the cosmic muon events recorded by the plane. The results were compared to the data, and the conclusions served to refine the event reconstruction process as well as to tune the simulation parameters.

The simulation software used is the same one as described in Sec. 2.3, developed for studying CHIPS-10 and larger modules. The CHIPS-M geometry is represented by an octagonal prism, with dimensions corresponding to the shape encompassed by

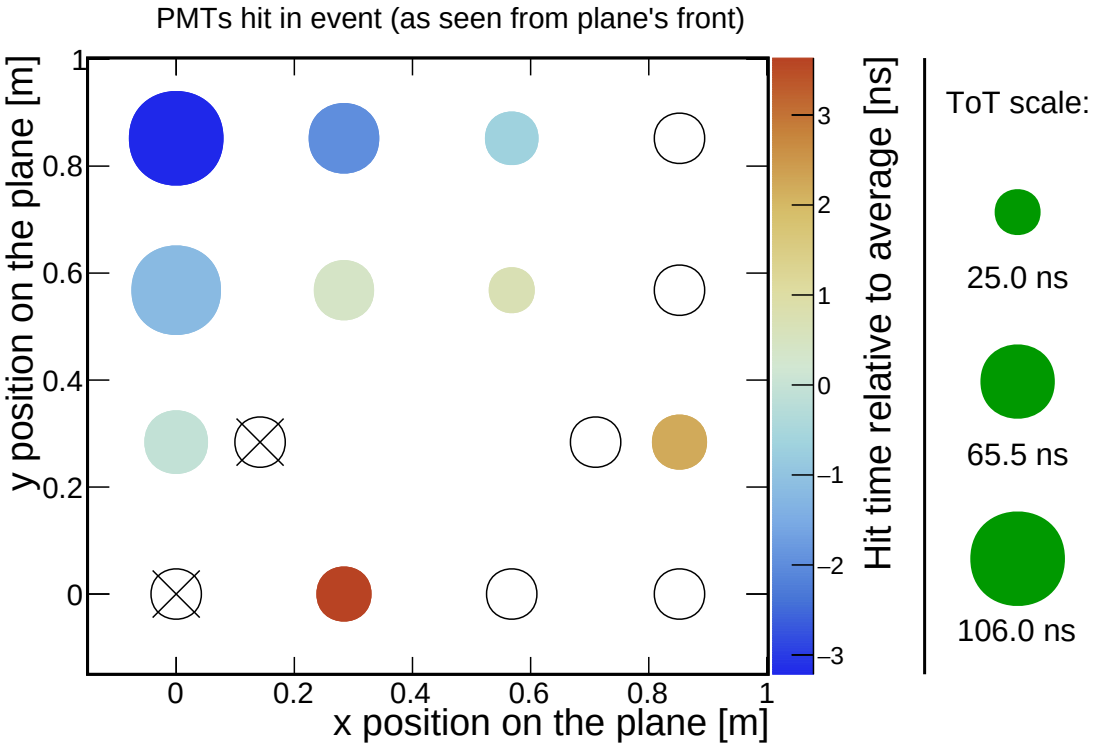

Figure 5.3: Event display of an example event registered by the plane. Each coloured marker corresponds to a PMT in its correct position on the plane, with the colour representing hit time relative to the average and the area proportional to the ToT. The empty circles represent PMTs without a hit, and the crosses signify the two disabled PMTs 0 and 5.

the liner: 3.3 m height and 1.72 m distance from the centre to the middle of a wall. The inside and outside volumes are filled with water, with light attenuation properties adjusted to match the lower purity found in CHIPS-M (Sec. 5.2.1). The detector is instrumented with 3 inch PMTs, at 6 % photocathode coverage, which corresponds to tiling the inside with copies of the Nikhef plane. The PMT definition, including photocathode size and quantum efficiency, is based on the actual Hamamatsu R12199-02 model used in the plane. A logical selection of 16 PMTs in the middle of one wall (close to the original location in CHIPS-M) is treated as the Nikhef plane to extract the plane hits. Such setup allows for the comparison of events recorded by the plane PMTs to those that could potentially be observed with a fully instrumented detector.

The cosmic muon tracks are generated by CRY, which is set up with the correct date and latitude, to take into account seasonal and geographic variations in the muon flux. The simulation only generates muons on the surface, so a simple transport

code propagates the surface tracks to the depth of 5 m above the detector. This calculation employs a linear energy loss with the average value of 220 MeV/m and assumes a straight track, with no Coulomb scattering. The muons are produced over an area of 1 km × 1 km to include zenith angles up to 5 degrees from horizontal, which is close to the real limit due to the rock pit walls. Muons that would miss CHIPS-M are discarded, but taken into account when calculating total exposure time and hence event rate in CHIPS-M.

The generated cosmic muons are passed to the Geant4 code as primitive muon tracks with vertices 5 m above the top cap. Then, the event is fully simulated, including muon propagation, Cherenkov light production, and photomultiplier hits with subsequent digitisation. The resulting files are processed by event reconstruction code with output saved in the same format as for data event files. To simplify the comparisons with data, hits in PMTs 0 and 5, corresponding to the ones disconnected in the original plane, are discarded. The total exposure time and the number of full detector events are stored as well. The hit digitisation algorithm, based on SKDetSim, outputs PMT charge in number of photoelectrons. To obtain the time over threshold value comparable to the data, a simple linear conversion is performed, with 32 ns ToT corresponding to 1 PE (Sec. 5.4.1). Figure 5.4 shows the event display of an example simulated cosmic muon event in CHIPS-M.

One of the main limitations of the generated Monte Carlo samples is the low number of events, leading to large statistical errors. In a typical particle physics analysis, the MC statistics are usually much higher than in the data, since additional CPU time is generally a cheaper resource than detector exposure. However in this case, the iterative nature of the analysis required frequent production of new Monte Carlo samples with different settings, each of which was computationally intensive. On the other hand, the sample size in the data is virtually infinite, due to the high intensity of cosmic muons.

5.2.1 Water attenuation properties

The detector simulation is based heavily on WCSim, and has inherited the description of water attenuation properties. Absorption and Rayleigh scattering lengths are hard-coded as tables of values for a number of photon energies ranging from approximately 1.5 to 6.2 eV, corresponding to wavelengths of 200 to 800 nm. The curves are

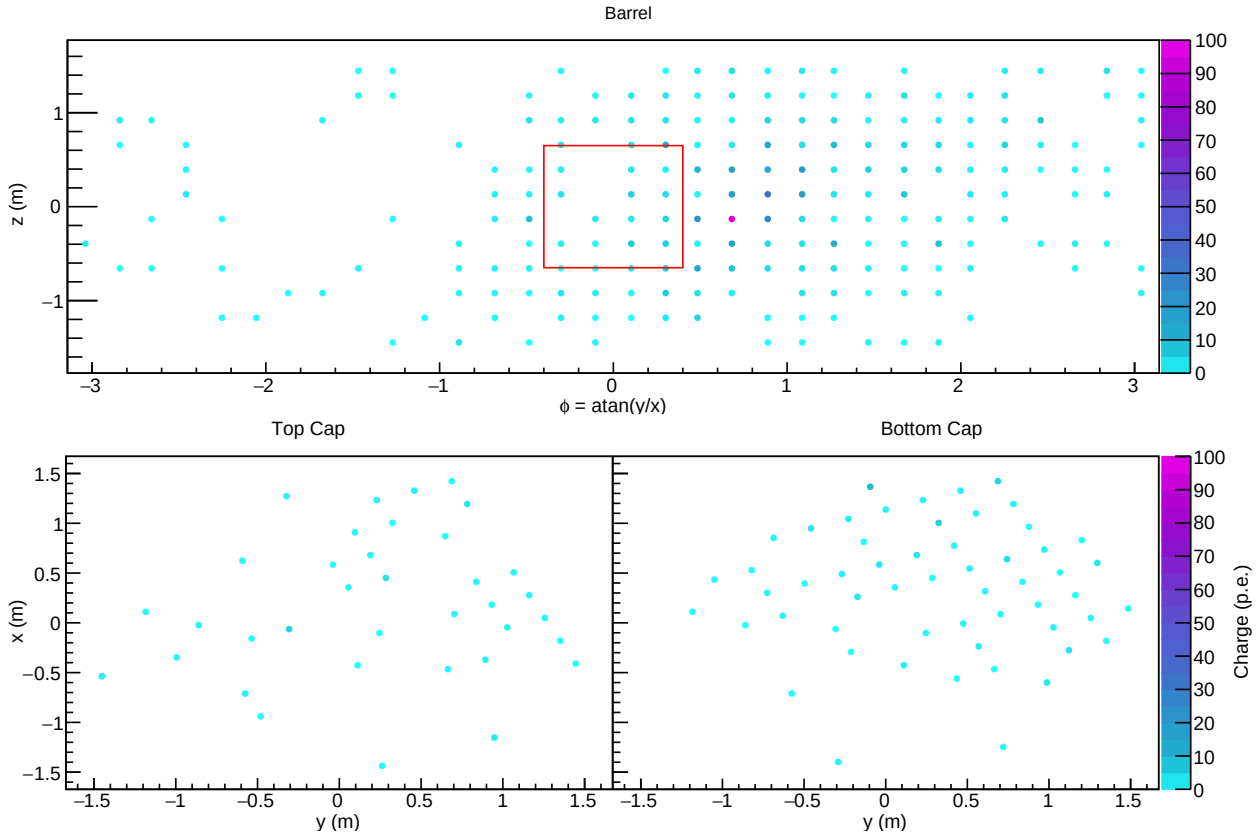

Figure 5.4: An example cosmic muon event, simulated in a hypothetical fully instrumented CHIPS-M detector. The event display shows the unrolled detector wall and two endcaps, with the coloured markers representing hit PMTs. The colour corresponds to the registered hit charge. The highlighted selection of 16 PMTs is treated as the Nikhef plane.

then scaled by multiplicative coefficients set in the configuration script and loaded at runtime. In the original WCSim code the total attenuation length[1] reaches a maximum value of 127 m at 410 nm. The values come from the Super-Kamiokande simulation SKDetSim, according to a comment in the code [106]. The combined attenuation length fits quite well other measurements of ultra pure water (Fig. 5.5) [144, 145].

The water in CHIPS-10 is expected to reach 30 to 50 m attenuation length at the same wavelength, but in CHIPS-M it was much less pure. For the initial deployment, the water for filling the detector was pre-filtered and deposited in a swimming

[1]The attenuation length is the inverse of the attenuation coefficient, which is calculated as a sum of the absorption and scattering coefficients. The total attenuation length is therefore the inverse of the sum of inverses of the absorption and attenuation lengths.

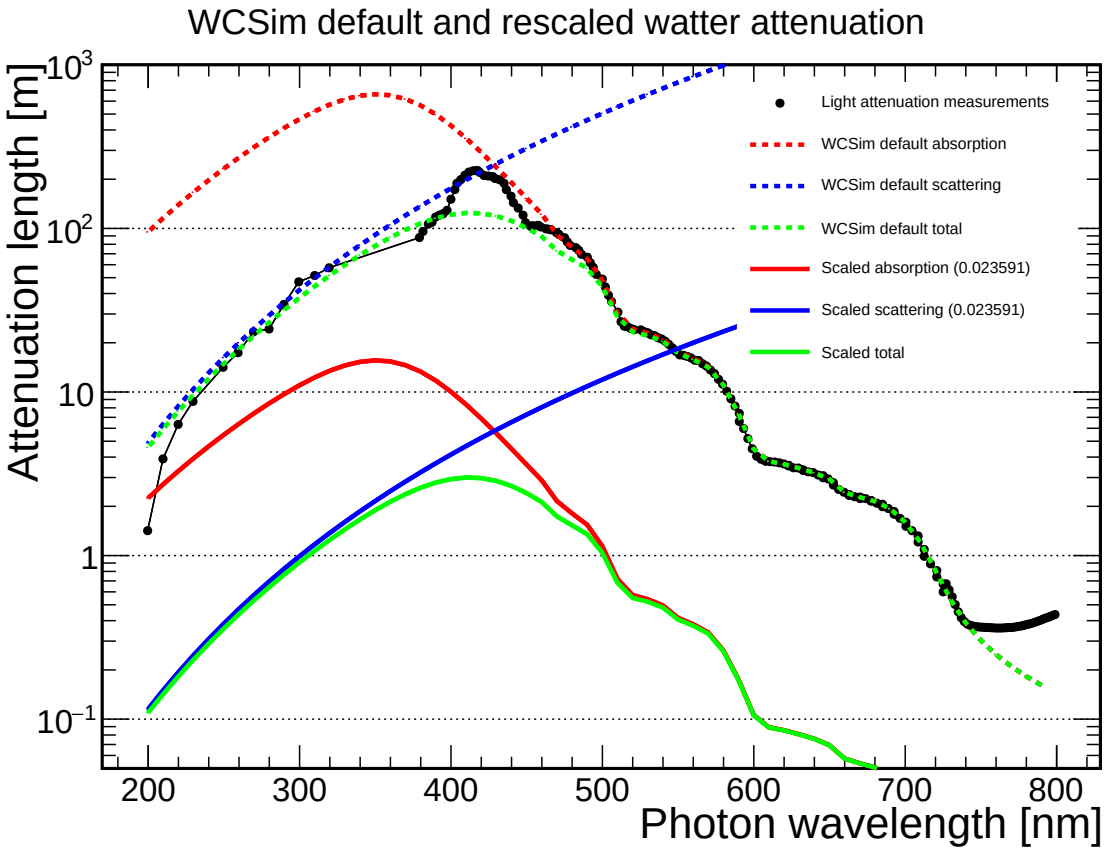

Figure 5.5: Attenuation length in water as a function of photon wavelength, including absorption (red), Rayleigh scattering (blue) and total attenuation (green). The black points show data from measurements [144, 145], the dotted lines show default values in WCSim code and the solid lines show the result of a rescaling in order to achieve 3 m total attenuation at 410 nm with relative weights of absorption and scattering the same as in default settings.

pool acting as a tank. However, the deployment had to be postponed for a month and logistic constrains prevented any reprocessing of the stored water. Because of insufficient cover it has likely kept accumulating dust. In addition, the detector itself was contaminated with particulates, which could not be entirely eliminated, despite extensive cleaning before deployment. This contamination would normally be filtered out by the water purification system, but the process takes time and the recorded data comes only from the first two days.

The attenuation of a water sample from the detector was measured with a dedicated test setup. The same instrument was later used to determine the possible water purity in future detectors achievable with a cheap filtering system [98]. However,

at that time the instrument performance and systematics were not completely understood, and only an approximate result was obtained, of around 2 to 3 m attenuation at 405 nm.

The employed technique of measuring the light intensity of a laser beam at the end of a water-filled pipe can only provide the total attenuation, without separate information about absorption or scattering. However, those two phenomena have very different effects on the propagation of Cherenkov light in the detector. Absorption decreases the light yield, while scattering redirects it, washing out the Cherenkov ring. In CHIPS-M, where the plane covers only a small area of the detector, scattered light can even lead to an increase in observed hits and events. Since the actual contributions of those two effects are unknown, different combinations were tested for their agreement with data (Sec. 5.4.3).

The functional dependence of absorption and Rayleigh scattering length on photon wavelength was kept the same as in the original code. Scale factors were applied to reach a desired total attenuation length and relative weights of absorption and scattering. Figure 5.5 shows an example of such rescaling, where the relative weights are kept the same as in the original code, but the combined attenuation length is 3 m at 410 nm.

The reference Monte Carlo used throughout this chapter is set with the total attenuation of 4.1 m, where the relative weight of absorption to scattering compared to the original code is 3:1, or 7.2:1 in terms of actual attenuation lengths.

In order to avoid the issues with disentangling these two effects in the future, an upgraded water system is currently tested. By observing the laser light from an angle, it will measure the scattering length, making it possible to separate it from absorption.

5.3 Time calibration

The time elapsing between a photon hitting the PMT photocathode and the signal being time-stamped at the CLB depends on many factors. First, the electron cascade propagates down the dynode chain in the PMT, then the current is amplified and processed by the readout chip, and the signal sent through the cable to the CLB,

with the delay at each step subject to variability. The main contribution to the total time spread is introduced at the first stage of electron multiplication. The initial photoelectron travelling from the photocathode to the first dynode has the slowest speed, and small distortions of its propagation time have the biggest effect on the final result.

The variability in the timing between different hits in the same PMT is described as the hit time resolution, and is mostly affected by the nonuniformities of the electric field between the photocathode and first dynode. The typical time spread of the Hamamatsu PMTs used in the Nikhef plane is estimated to be less than 2 ns [140].

In addition to that effect, there are systematic differences between various PMTs, which can be caused e.g. by small differences in the distances between the dynodes, or deviations in the timing of the electronic components, including the cable length. These differences appear to be on the similar scale as the resolution, 1 to 3 ns. The goal of the time calibration is to correct these offsets, so that the difference in corrected hit times reflects the actual photon hit time and makes it possible to reconstruct the muon track direction.

The corrections are determined by analysing events recorded during calibration runs, when a bright LED was periodically illuminating the detector in short pulses. While the plane has its own calibration LED, due to the time constraints it was not used. Instead, LED flashers in the IceCube DOMs were employed, during three 1 hour long calibration runs.

5.3.1 Calibration with the IceCube flasher

First, the runs during which a flasher was active were processed by the event finding procedure, but without the initial time correction step. In the runs with the *Haifa* DOM, which contain 2 hours of data, the LED light was very bright, making it easy to identify events caused by the pulses. The distribution of the combined time over threshold of all hits exhibits a clear peak at around 2700 ns (Fig. 5.6), corresponding to an average ToT of 190 ns. The candidate events were selected by requiring a total ToT of more than 2000 ns and also that all PMTs in the plane register a hit, in order to avoid skewing the average time.

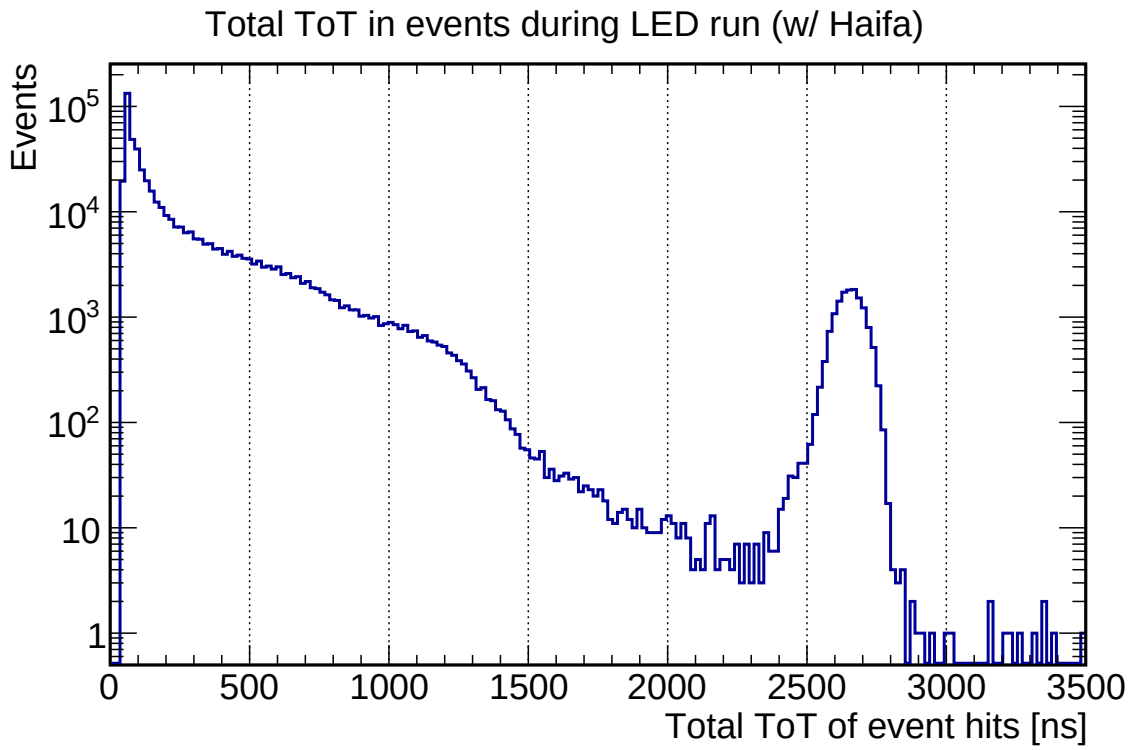

Figure 5.6: Distribution of the total time over threshold in the event, during a calibration run with an LED flasher on the *Haifa* DOM. The events caused by the LED light form a clear peak around 2700 ns.

In the run with a flasher active on the *Munich 1972* DOM, the LED brightness is significantly lower, with the average total ToT around 1150 ns. This requires a lower cut, at 1000 ns, which results in a sample with many more events coming from cosmic muons instead of LED light. Because of that, only the *Haifa* runs are used in the calibration; but the *Munich* events give consistent results.

The selected events occur at intervals consistent with the period of the flasher (0.21 s). An example can be seen on Fig. 5.7, which shows the number of hits in the event as a function of time for a 5 s long fragment of the first calibration run with the *Haifa* DOM.

For each selected event, the difference between the PMT hit times and the average time is calculated and stored in a histogram, separately for every PMT. The average time is assumed to most accurately reflect the true event time, up to an additive factor. If the PMT resolution, and hence the variation in hit times between events, is described by a Gaussian distribution, then the average time also forms a Gaussian

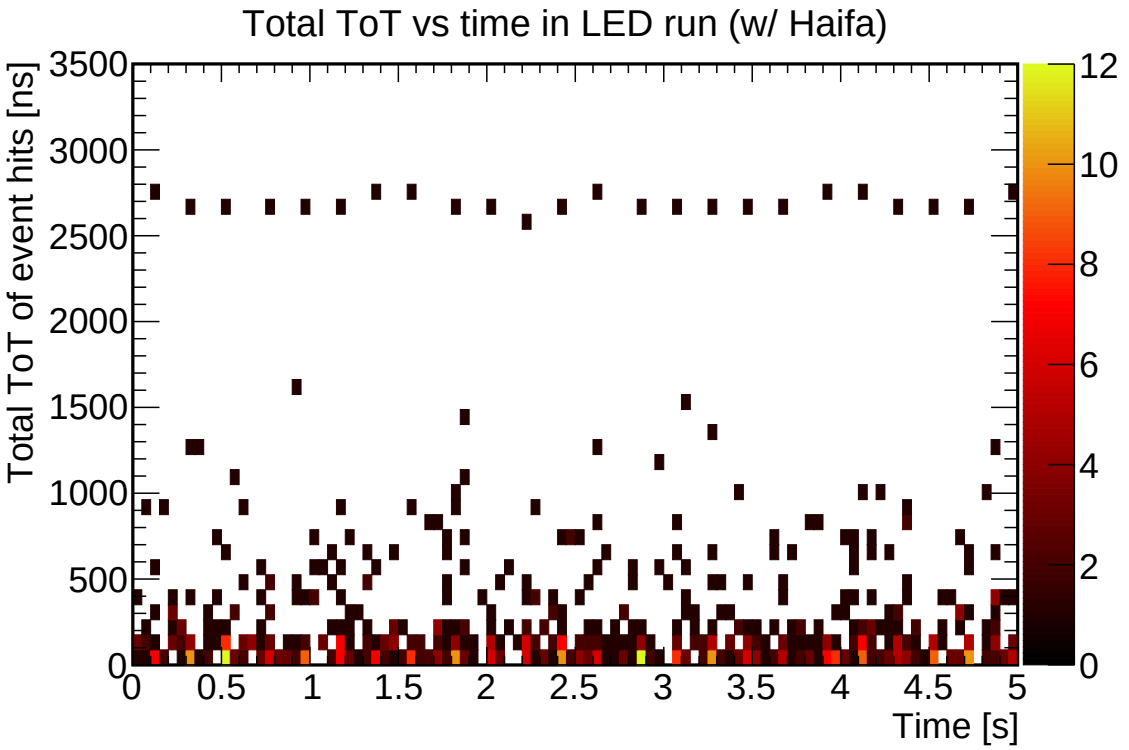

Figure 5.7: Total time over threshold in the event as a function of the first hit time, for an example period during a calibration run with an LED flasher on the *Haifa* DOM. Each entry in the 2D histogram represents a single event. The LED candidate events can be clearly seen, consistent with the frequency of 4.8 Hz.

around the true time, and so does the difference between hit time and the average. In practice, the observed distributions are not entirely Gaussian, but the agreement is quite good.

Figure 5.8 shows the distributions of the hit times relative to the average for every PMT. The widths of the distributions correspond to the resolutions of the individual PMTs, at the observed light level. The differences between the peak positions come from the intrinsic offsets and the different propagation times from the light source to the PMTs.

The distances from the flashing DOM to the individual plane PMTs are calculated based on a simplified geometric model of CHIPS-M. Since precise measurements of the positions of the IceCube DOMs and the Nikhef plane have not been performed, the dimensions are reconstructed from available technical drawings of the CHIPS-M structure and the DOM mounts. However, the distance from the centre of the *Haifa*

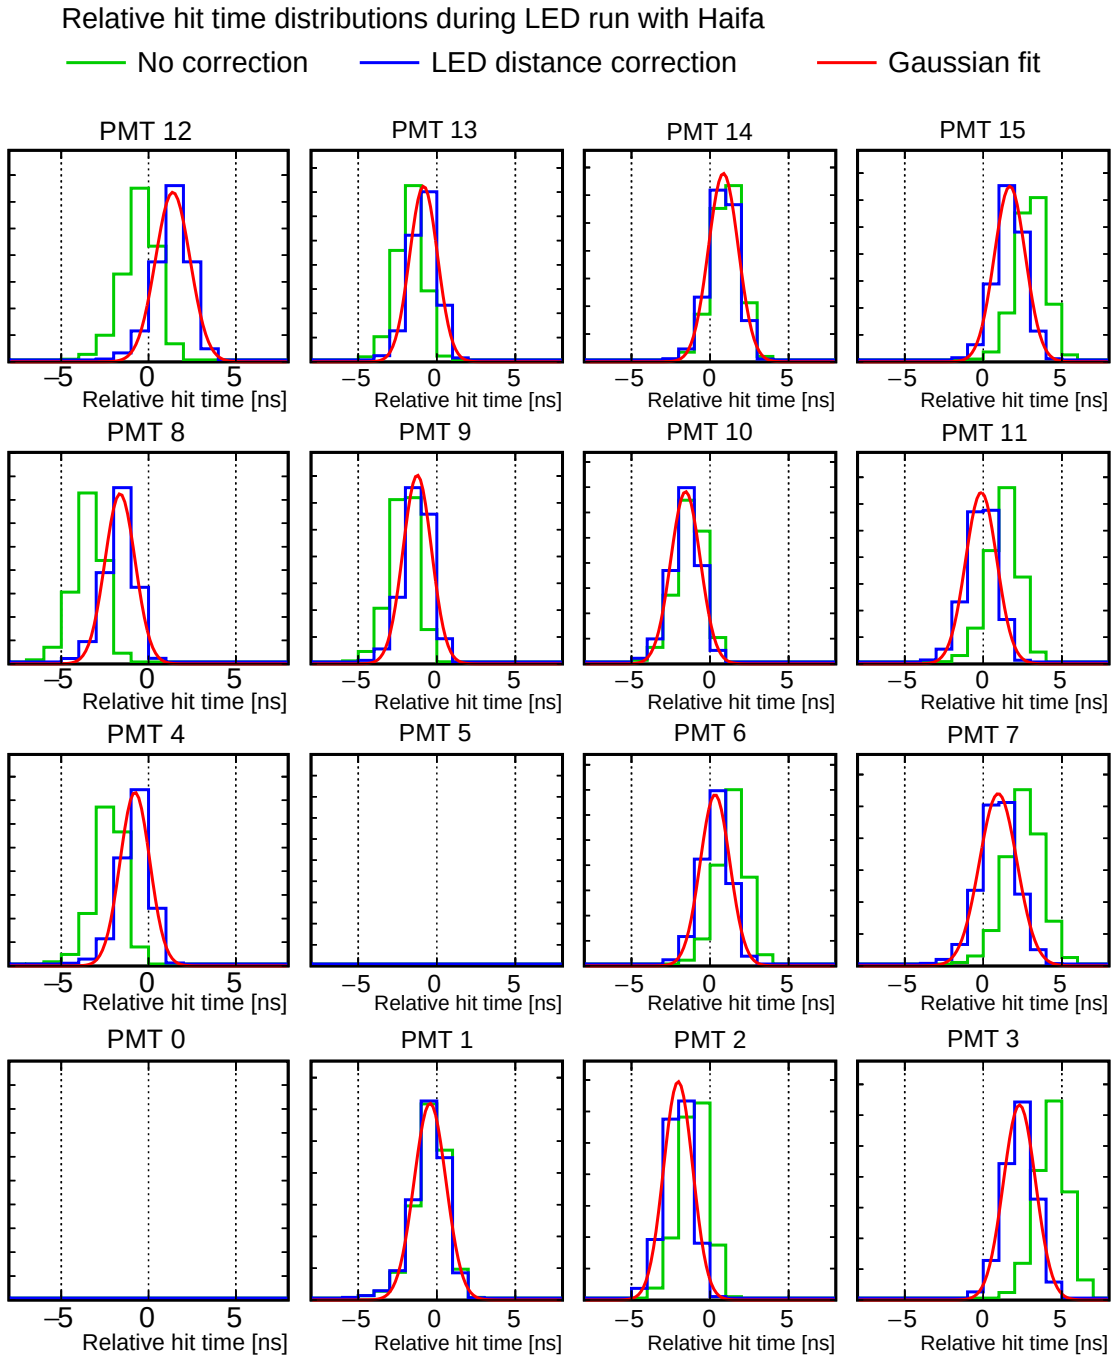

Figure 5.8: Distributions of the relative hit time (difference between hit time and average hit time in the event) for selected LED events, in plane PMTs. The positioning of the plots reflects the PMT positions in the plane, and the LED light comes from top left. The histograms show the hit times before (green) and after a correction for the light propagation time (blue). The red lines show the results of the Gaussian fit.

DOM to two PMTs on the plane (number 12 in the top left and number 3 in the bottom right) were measured, providing constraints on the model. The final geometry is consistent with the known dimensions and the measured distances, and small variations were found not to impact the calibration significantly.

The distances from the source to the PMTs are used to calculate the propagation time of the light pulse, using c/n as the speed of light in water, where the refraction index, $n = 1.34$, is taken for the range around 450 nm. The relative time of flight corrections are calculated by subtracting the average time of flight from the individual values of the recorded hit times during the step of forming relative time distributions. The corrected relative time distributions, together with fitted Gaussian functions are shown in Fig. 5.8.

The extracted fit parameters are stored in a file, indexed by the PMT ID. The mean positions, μ, are assumed to represent the remaining systematic PMT offsets. During the time calibration step of event reconstruction, the offset for the appropriate PMT is subtracted from the observed hit time.

5.3.2 Results and validation

Figure 5.9 shows the duration distribution of events with at least three hits in a 50 ns window, for the uncalibrated data and with time calibration using the *Haifa* DOM flasher. Applying the calibration reduces the mean timespan of the events, which suggests that it is working as intended, since adding random systematic offsets would on average increase the distance between the first and last hit.

The event display can be used to visualise the timing of hits in cosmic muon events. Figure 5.10 shows the average of all events with 14 hits from a single 1 hour run. Counting only events with all the PMTs being hit ensures that averaging the relative hit times is not biased towards PMTs which may receive more Cherenkov light than others. Those events are also expected to be caused mainly by the direct light and have the hit times reflect the muon direction.

Since the angular distribution of the muon flux is symmetric with respect to the vertical axis, the average hit times in the PMTs are expected to exhibit a bilateral (left-right) symmetry. However, this is not observed. The top display in Fig. 5.10 shows the mostly random times due to the uncalibrated offsets, and while the bottom

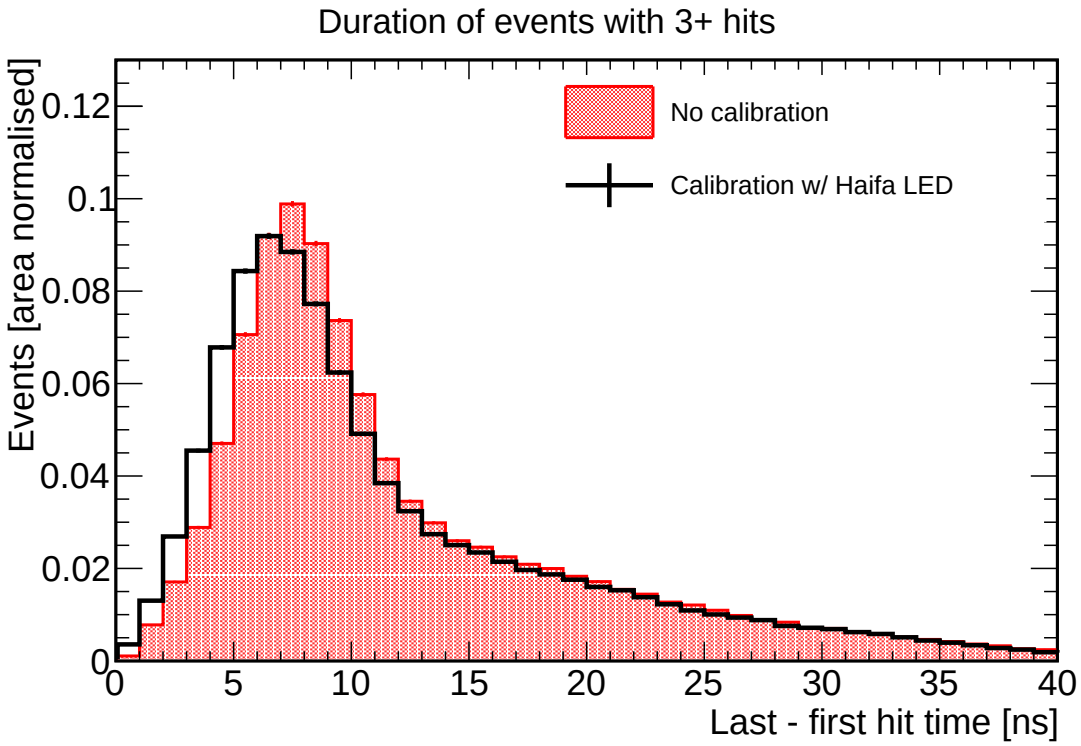

Figure 5.9: Distributions of event duration (difference between the last and first hit time) for events with at least three hits within a 50 ns time window, with ToT above 12 ns, in a single 1 hour run. The filled red histogram shows the duration for events without time calibration and the solid black line is for events with calibration using the *Haifa* DOM flasher.

one exhibits a clear pattern, it is biased diagonally. A time difference between the top and bottom PMTs is expected due to the muons coming from upwards, so the bias seems to be predominantly in the horizontal direction.

Figure 5.11 presents the average event with 14 hits in the plane as predicted by the detector simulation. One striking feature is that the average time over threshold is significantly higher than observed in the data. This issue is discussed in Sec. 5.4.1, but the timing of the simulated events is assumed to be unaffected. Indeed, it confirms the basic expectation that the top hits are on average earlier than the bottom ones, and that the hit times are symmetric with respect to the vertical axis of the plane.

The IceCube DOMs are located to the left of the Nikhef plane (when viewed from the front), with *Haifa* placed around the height of the two top rows of the

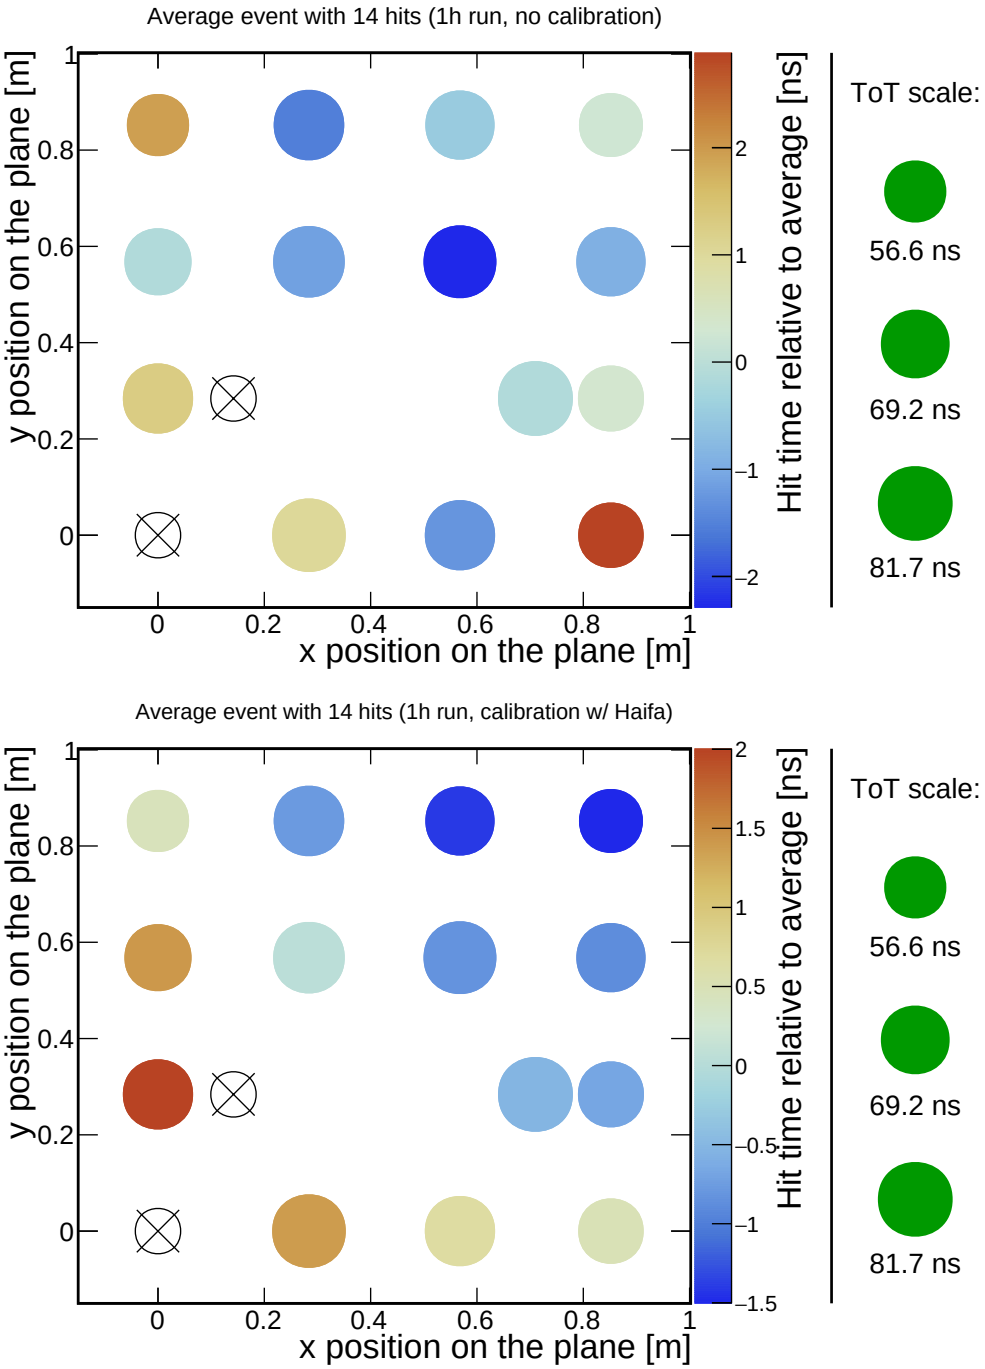

Figure 5.10: Event display of an averaged event with 14 hits within 50 ns in a single 1 hour run, without time calibration (top), and after calibration with the *Haifa* DOM LED (bottom). The coloured markers corresponds to hit PMTs, with the colour representing hit time relative to the average and the area being proportional to the ToT. Assuming all events are caused by cosmic muons, a bilateral symmetry would be expected. The LED light during calibration runs comes from the top left side of the plane.

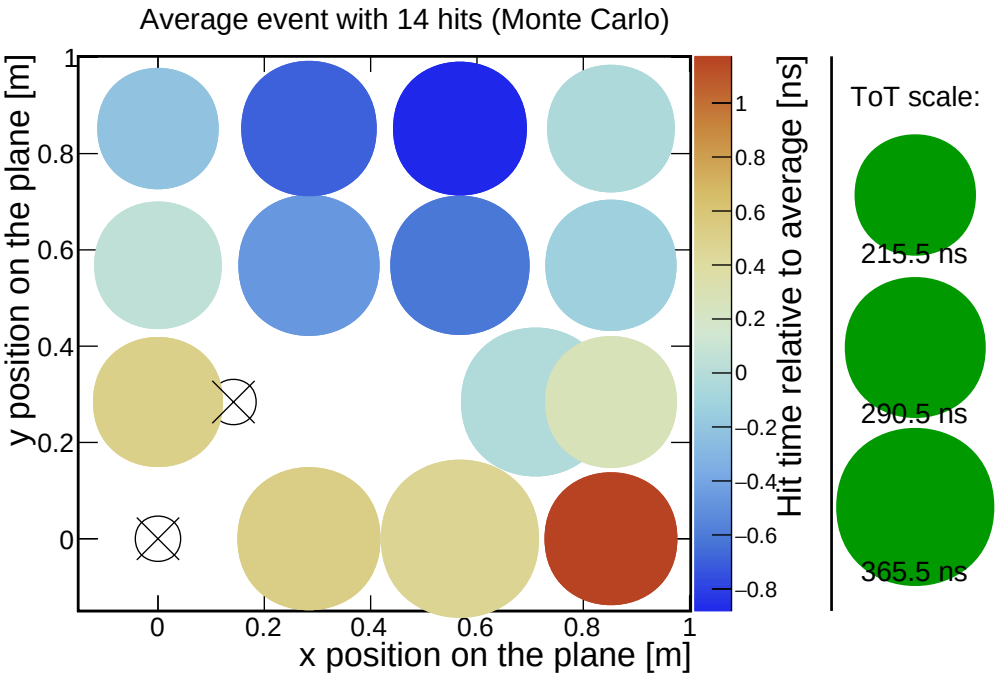

Figure 5.11: Event display of an averaged event with 14 hits within 50 ns as predicted by the simulation with reference values of the water attenuation parameters. The coloured markers corresponds to hit PMTs, with the colour representing hit time relative to the average and the area being proportional to the ToT. A bilateral symmetry can be observed, confirming the expectation.

plane PMTs. The flashing LED is therefore much closer to the leftmost PMTs than it is to the rightmost ones. This can be seen by looking at the effect of the distance correction in Fig. 5.8.

The apparent bias in the calibration could be explained by the fact that light intensity is not taken into account when calculating the expected differences between hit times in LED events. For example, the distance from the *Haifa* DOM to the bottom-right PMT (number 3) is almost two times longer than to the top-left one (12). The number of photons reaching PMT 3 should be on average 3.5 times smaller, according to the inverse square law.

The hit time is counted as the moment when the PMT signal first rises above the threshold. For each photon, the arrival time and the time of the propagation of the electron cascade are subject to variability. The more photons hit the PMT, the smaller on average will be the earliest of those times. If the hit time, t, for each of n photons is described by a Gaussian distribution with the same parameters μ and σ,

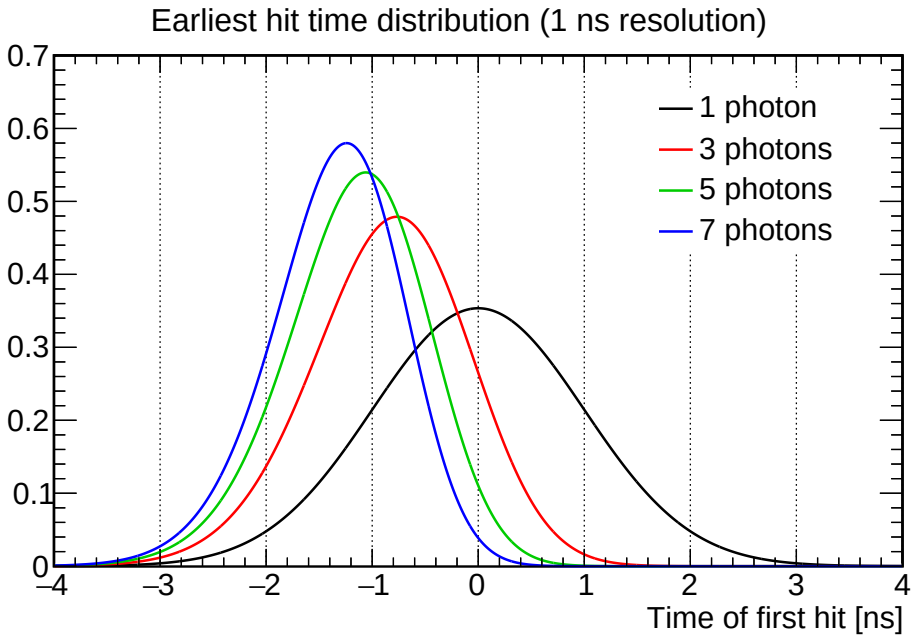

Figure 5.12: Distribution $P(t_{\text{min}})$ of the earliest of n hit times (Eq. 5.1), each of which is described by a Gaussian with $\mu = 0\,\text{ns}$ and $\sigma = 1\,\text{ns}$. Solid lines show the distribution for various numbers of hits: one (black), three (red), five (green) and seven (blue).

the earliest hit time will be distributed according to [1]

$$P(t_{\text{min}}) = \frac{n}{2^n \sigma\sqrt{2}} \left(1 - \text{erf}\left(\frac{t-\mu}{\sqrt{2}\sigma}\right)\right)^{n-1} \exp\left(-\frac{(t-\mu)^2}{2\sigma^2}\right). \tag{5.1}$$

Figure 5.12 shows the distribution of the earliest hit time for several values of n, to illustrate the effect. As the number of photons increases, the width of the distribution decreases and the mean shifts towards earlier times. The difference between means for $n = 1$ and $n = 5$ is comparable to the base resolution (in this case, 1 ns). The mean shift is not linear with n, and cannot be calculated just based on the ratios of light intensity.

Unfortunately, the absolute number of photons from the calibration LED is unknown, since its relation to the observed time over threshold is not well understood (Sec. 5.4.1). Because the earliest time distribution is very similar to a Gaussian, it is virtually impossible to obtain the parameter n from fits to the observed hit times.

Ultimately, the time calibration has a negligible effect on the event rates, and is only necessary for determining the muon directions. As the latter objective ended up out of scope of this work due to time constraints, further investigation of the calibration has been left for future work. Since the existing calibration with the *Haifa* flasher appears to remove at least some of the bias due to PMT offsets, it is still applied in the other stages of the analysis.

One of the promising avenues of solving the issues in the near future is the use of the on-board calibration LED (Nanobeacon). The plane, now located at Madison, was tested in a dark-box, and in fact several LED runs have been recorded. Because the plane LED is much closer to the PMTs, intensity effects are even stronger than with the IceCube flashers. However, finer control over the light intensity and the ability to collect more runs may allow for a better investigation of the PMT response.

5.4 Muon rate measurement

The main goal of the analysis presented here is the measurement of the total rate of cosmic muons entering the CHIPS-M detector. Because the plane occupies only a small area of the wall, it cannot detect every single muon. The detector simulation is required to obtain the geometric acceptance of the plane and apply a correction. For this reason, predictions of the simulation are compared with data to check how well does it describe the reality.

In order to reach agreement between data and Monte Carlo, the event definition is slightly adjusted, and the MC properties are tuned. To ensure that the result is still meaningful, such modifications should be well motivated and ideally performed blindly, i.e. without the knowledge of the final result. This analysis is very preliminary, with many parameters poorly constrained, so these requirements were not always respected. For example, the water attenuation properties were not measured precisely. Therefore different combinations were tested for agreement, and in the initial stages it was done without obscuring the total event rate, making it much easier to spot mistakes or disagreements. Hence, the final result should only be treated as a first estimate, with future work employing more rigour.

The following points are not presented in any particular order. In fact, the analysis was performed in an iterative fashion, sometimes in an attempt to fix problems that disappeared at later stages.

The final event selection, assumed unless specified otherwise, includes rejection of hits with the ToT below 12 ns, time calibration with the *Haifa* DOM, a time window of 50 ns and, optionally, a clustering requirement with maximum separation of 10 ns. Only events with at least three hits are considered. The reference Monte Carlo simulation assumes the water absorption length of 33.6 m at 410 nm and Rayleigh scattering length of 4.7 m at the same wavelength (with total attenuation length of 4.1 m). The wavelength dependence is the same as originally coded in WCSim. However, the final result is presented for Monte Carlo samples with different settings as well.

5.4.1 Time over threshold distribution

A significant issue with the time over threshold processing method is the difficulty in reliably reconstructing the number of registered photoelectrons, i.e. the hit charge [141]. The pulse duration and hence the ToT value do not depend linearly on the total charge. Moreover, two photons hitting the photocathode some time apart can end up creating a long but shallow pulse, mimicking the PMT response to a higher charge. At the same time, the detector simulation does not yet contain a description of the KM3NeT readout electronics and only provides the estimated number of photoelectrons, based on the Super-Kamiokande PMT simulation. Therefore, it is challenging, if not impossible, to consistently compare the hit charge as simulated by the Monte Carlo and registered in data.

Figure 5.13 shows the time over threshold distributions in data (before and after event forming) and in Monte Carlo for two example PMTs. The data is taken from a single 1 hour run taken by the end of the second day of deployment, which is consistent with all the other runs, up to the total rate. The raw data plot shows the ToT of all registered hits, as stored in a histogram during the execution of *MRunAnalyzer*. The line labelled as event data contains hits that are part of an event of at least two hits, with no ToT cut, and 50 ns maximum duration.

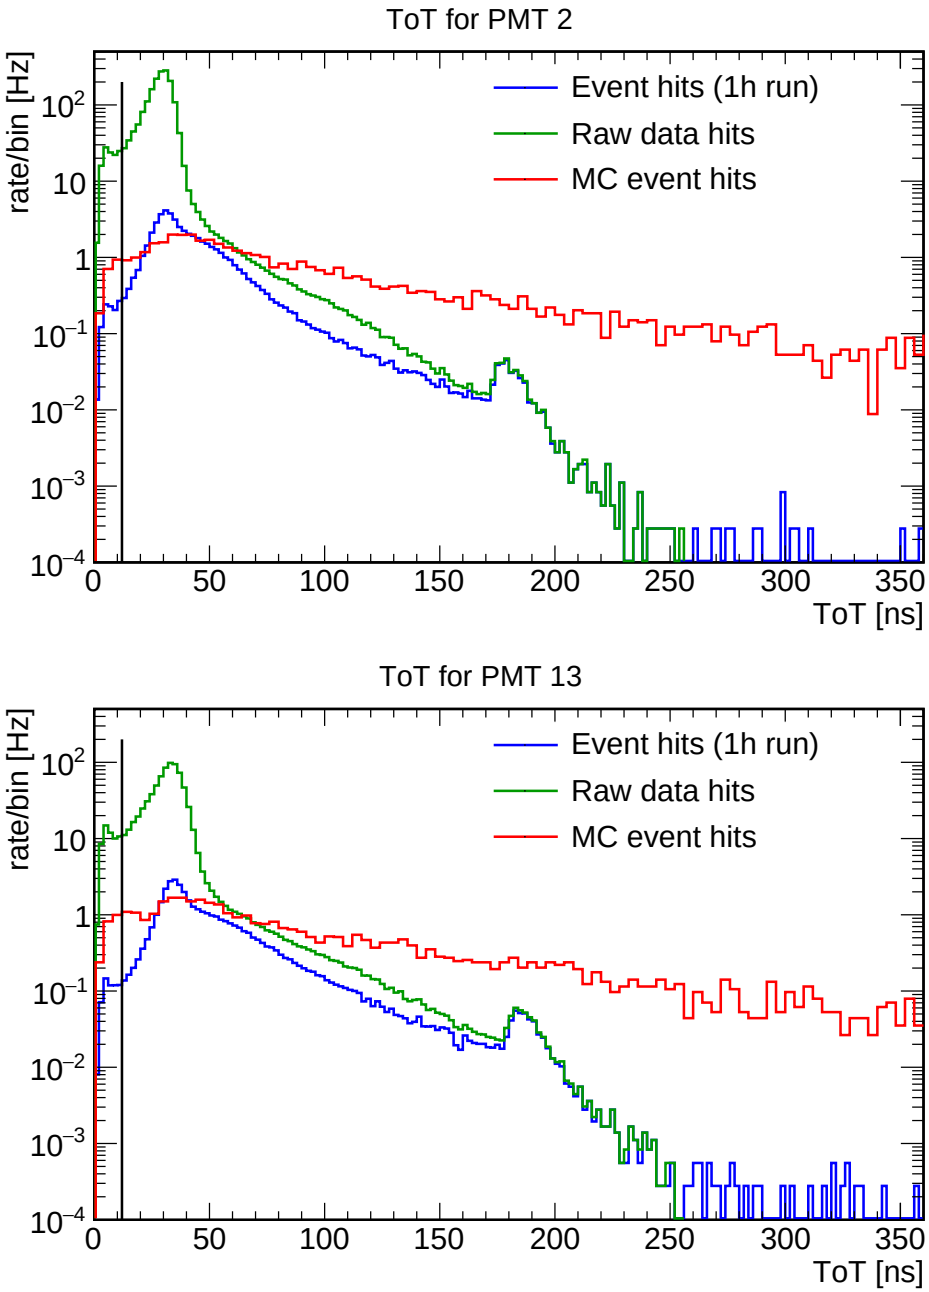

Figure 5.13: Time over threshold distributions of hits in two example PMTs, during a single 1 hour run. Solid lines show rates for all registered hits (green), hits that are part of an event (blue) and the Monte Carlo prediction with a linear conversion from PE to ToT (red). The vertical black line depicts the ToT cut value of 12 ns.

	1 PE peak in the ToT distribution	
PMT ID	Position (μ) [ns]	Width (σ) [ns]
1	33.4	4.2
2	29.6	3.7
3	31.1	3.7
4	33.0	3.8
6	32.9	3.7
7	32.4	4.3
8	32.0	3.7
9	32.0	3.6
10	32.4	3.6
11	32.0	3.3
12	31.3	3.5
13	33.2	4.6
14	31.7	3.7
15	30.1	3.7

Table 5.1: The fitted mean and width parameters of the ToT peaks corresponding to 1 photoelectron, for all available PMTs in the plane. The values are extracted by fitting a Gaussian in the region of ±5 ns around the highest bin in the ToT distribution of raw hits during the whole deployment period (except the LED calibration runs). The uncertainties on the fit parameters are on the order of 0.2 % or smaller and are not displayed.

The raw rate is dominated by dark current hits, mainly caused by thermionic emission, i.e. electrons ejected from the photocathode due to thermal fluctuation. These form the clearly visible 1 PE peak at 30 to 35 ns. The position of this peak is an important part of the ToT calibration. The *MRunAnalyzer* program performs automatic Gaussian fits in the area around the peak, to extract the mean position and width for each PMT. The results of these fits for all runs combined are presented in Table 5.1. The actual position of the 1 PE peak varies slightly between different PMTs, with the average located at 31.9 ns ToT. It is possible to adjust the high voltage or detection threshold levels of individual PMTs to calibrate the response to a common value, but this had not been done before the deployment due to time constraints.

Other sources of dark current and electronic noise contribute to the pedestal of the ToT distribution, with a smaller peak around 0 to 5 ns. The ToT cut applied during event reconstruction is aimed to reject most of these artefacts and reduce the number of events due to random coincidences. The nominal cut value of 12 ns is chosen such as to contain the entire pedestal peak for all PMTs.

The gap between the raw hits and event hits at higher ToT values is likely due to real Cherenkov light that only reached one of the plane PMTs and therefore did not cause any coincidences. The raw hit distribution is also cut off at 255 ns, which is the maximum value before the concatenation of longer hits, performed at a later stage in the code (Sec. 5.1.3).

The ToT distribution observed in data features also another peak at around 190 ns. The origin of this bump is not fully understood, but it is believed to be related to muons passing through, or very close to, the PMT photocathode.[2]

The detector simulation output provides the digitised charge registered at the PMT, expressed in number of photoelectrons. The simplest conversion to time over threshold is a linear one, with 0 ns ToT corresponding to 0 PE and 32 ns to 1 PE. In principle, this conversion should be adjusted to the actual 1 PE peak position for each PMT. However, the variability between PMTs is at best a second order effect in the mismatch between data and MC, so the same parameters are used for simplicity.

It is clear from Fig. 5.13 that the linear conversion of the simulated charge does not agree with observed distributions. The depicted Monte Carlo sample uses the reference values of the water attenuation parameters, with 4.1 m attenuation length at 410 nm and 7.15 ratio between absorption and scattering lengths. However, different settings do not significantly affect the shape of the ToT distribution.

Other PE→ToT transformations were considered, in particular one of the form $\mathrm{ToT} = C_1 \log(C_2 n\mathrm{PE})$, where C_i are the fitted parameters. This conversion is based on the assumption that the PMT pulse follows an exponential decay, and hence its duration is roughly a logarithm of the pulse height. It is very difficult to perform a stable numeric fit, because of the low statistics in the MC and a trade-off between fitting to the slope or peak position. The parameters were therefore adjusted manually, keeping the constraint of 32.5 ns $\sim$ 1 PE.

[2]Other tests conducted by the KM3NeT collaboration found a correlation between similar peaks appearing in the data and muons identified by hodoscope detectors to pass closely to the DOM [146].

While it is possible to choose the parameters such that there appears to be a closer agreement between the data and simulated ToT, this assessment is very subjective and the parameter values are not physically motivated. In addition, the resulting transformations are highly nonlinear, with e.g. the ToT value of 250 ns corresponding to at least 50 PE. At the same time, in a KM3NeT test setup with the same readout and similar gain, this value is reported to be equivalent to 8 to 10 PE instead [147].

These findings suggest that either the time over threshold is much less linear in photoelectrons than expected, or that the modelling of the PMT charge implemented in the simulation is not correct. Unfortunately, without a good understanding of the observed values, it is difficult to investigate and quantify the disagreement. With future measurements involving other readout systems or a dedicated setup with full waveform output, it will be possible to compare the simulated quantities to data directly. In this analysis, it is assumed that the potential charge mis-modelling does not severely affect the hit multiplicities, which are used to tune the water properties in the simulation.

With the unresolved issue of ToT mismatch, it is important to examine the impact of the ToT cut on total event rates. Because the cut affects data and Monte Carlo differently, the geometric acceptance correction needed for a total muon rate measurement, and evaluated directly from the simulation, will not be consistent. To test this, samples from data and MC were processed with three ToT cut values: 0 ns, 12 ns and 18 ns, corresponding loosely to 0 PE, 0.38 PE and 0.56 PE, respectively. The ratio of events observed by the plane in data and Monte Carlo was compared between these three values. For events with at least three hits, both in a window of 50 ns and within a cluster with a maximum separation of 10 ns, the ratio varies by about 2 % between the nominal cut value and either of the extremes. The relative stability of the ratio suggests that the mismatch between the ToT distributions does not affect the final acceptance correction by more than a few percent.

5.4.2 Event duration and PMT resolution

The main parameter affecting the search for hit coincidences is the length of the time window. If the maximum time period during which all hits are considered part of an event is extended too far, there will be more random coincidences and dark rate hits included in real events. On the other hand, if the time window is too small, relevant

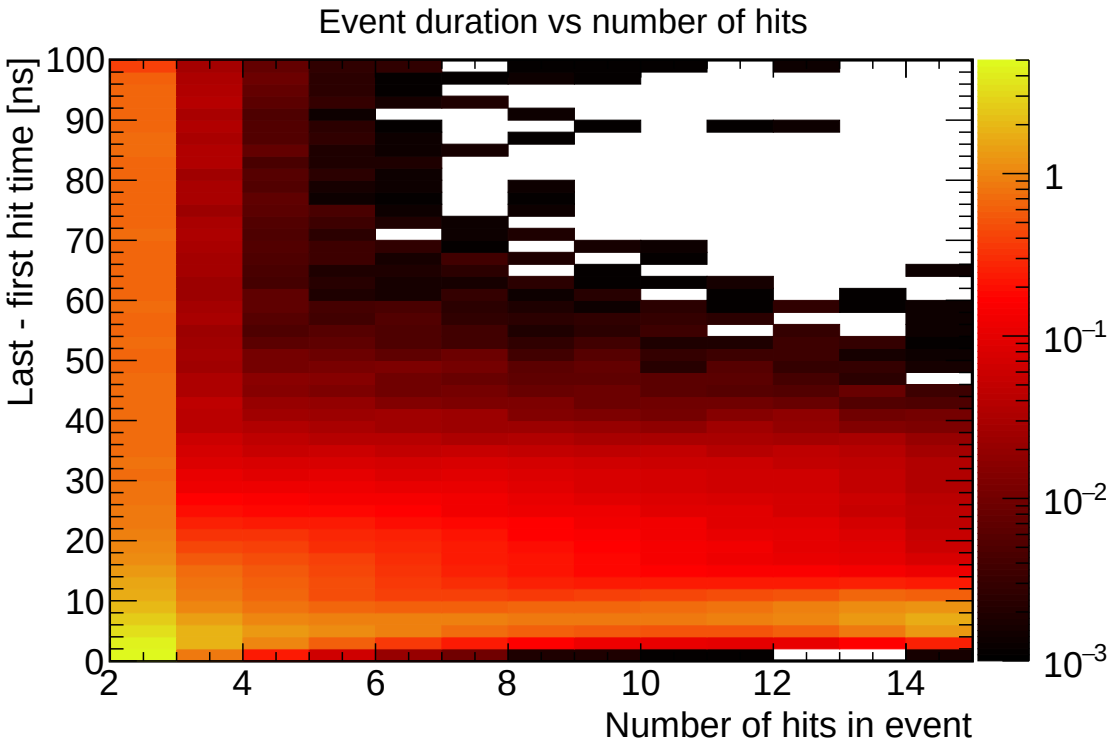

Figure 5.14: Distribution of the total event duration in bins of the number of hits in the event, for a single 1 hour run. The event processing includes time calibration and a ToT cut at 12 ns, but the time window is 100 ns long.

hits may be excluded and the total rate will be underestimated. For the purpose of validating the acceptance correction, the most relevant question is whether the effect of the timing cut on the observables is reflected in the simulation.

The initial value of the event window was chosen based on the timing observed in the data. Figure 5.14 shows the distribution of event duration, measured as the difference in time between the last and the first hit, in bins of hit multiplicity. The events come from a 1 hour run, with time calibration and a ToT cut of 12 ns, but with a 100 ns long window. No additional constraints (e.g. clustering) are imposed.

The majority of events, especially at higher multiplicities, are shorter than 20 ns. The distributions for events with few hits, particularly those with only two hit PMTs, are relatively flat, mainly due to the contribution from random noise. Based on these observations, a conservative time window of 50 ns is chosen to include all reasonable coincidences. Events with only two hits are rejected to limit the impact of the dark current. Events with only three hits are kept, in order to avoid further reduction in

the number of bins (and statistics in the MC samples) for the adjustment of water properties.

When comparing the observed event duration with the Monte Carlo predictions, the PMT resolution must be taken into account. The reported resolution for the Hamamatsu PMTs is less than 2 ns, but a precise value is unknown [140]. The default time spread in the simulation was switched off in favour of applying a flexible smearing at the histogram creation step, to allow for testing of multiple values. This is accomplished by altering the simulated hit time by a value drawn from the normal distribution, with zero mean and the σ parameter equal to the desired resolution. A more correct approach would be to also take into account the hit charge. However, this is not trivial without a good understanding of the ToT behaviour, and unlikely to dramatically affect the result.

Figure 5.15 (top) shows the event duration distribution for events with at least three hits in data and Monte Carlo, for three values of the PMT resolution: 1 ns, 1.5 ns and 2 ns. The distributions are area normalised in order to eliminate the absolute event rate from the shape comparison. From a visual evaluation of the peak region, the correct resolution value seems to be located somewhere between 1.5 and 2 ns. However, the mean of the data distribution is in fact smaller than for the simulation with $\sigma = 1.5$ ns, due to the mismatch in the tail. In addition, the mean values depend on the hit multiplicity, with the MC typically overestimating the duration of events with few hits and underestimating it at high multiplicities, relative to the data. Since the precise value of the time resolution does not impact the event rates significantly, the value of 1.5 ns was kept as a simple and sensible choice.

Another observation from Fig. 5.15 is that the simulated event duration distributions are narrower than the ones for data. This suggests that the event duration is affected by some additional phenomena, not modelled by the Monte Carlo. One such possibility could be the leftover calibration bias. However, implementing systematic PMT offsets into the simulated event hit times produced even narrower distributions, further increasing the disagreement.

During the early stages of the analysis, the water attenuation in the simulation was dominated by absorption, and the predicted event duration distributions were missing the long tail. This served as the motivation to implement the clustering

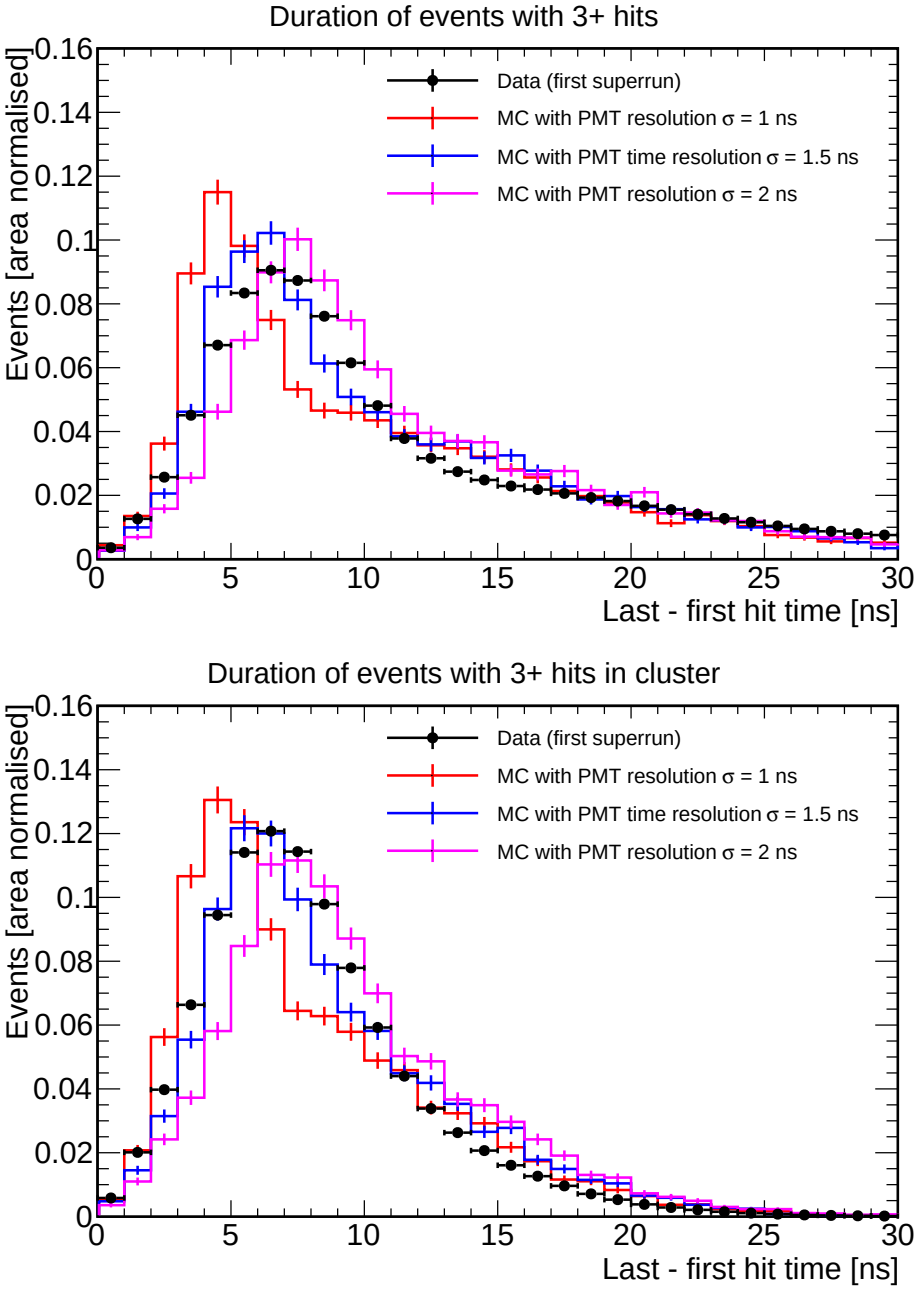

Figure 5.15: Distribution of event duration with a 50 ns time window (top) and an additional clustering with a 10 ns maximal separation (bottom). The black markers show data points from the first superrun, processed with time calibration, and a 12 ns ToT cut. The solid lines correspond to the simulation with reference values of water attenuation parameters and three different PMT resolutions applied at histogram creation: 1 ns (red), 1.5 ns (blue) and 2 ns (pink). All distributions are area normalised.

algorithm (described in Sec. 5.1.4), which aims to reject mostly single hits occurring significantly before or after the main event.

The clustering can be justified by the fact that the simulation does not model the dark rate or PMT effects resulting in late hits. The latter can come from a hit in the peripheral region of the photocathode, where the electric field is weaker and the first photoelectron takes a longer time to reach the first dynode [148]. Another possible effect is that the initial photoelectron gets scattered back from the first dynode and delayed until it turns around and starts the typical cascade [149].

Further investigation proved that the event duration tail can also be caused by the scattered light hits and is much better predicted in MC samples with lower scattering length. Nevertheless, imposing the clustering condition with the initial cutoff value of 10 ns affects the agreement between data and Monte Carlo differently for various MC samples. For this reason, event definitions both with and without clustering are employed when comparing data with simulation, calculating the acceptance correction and presenting the final result. Figure 5.15 (bottom) shows that adding the clustering requirement does not significantly alter the determination of the PMT resolution, and, in case of the reference MC sample, may slightly improve the agreement with data.

5.4.3 Tuning water properties

Absorption and Rayleigh scattering in water have a very significant effect on the rates of events observed in the Nikhef plane, and on the ratio of the observed event rate to the total muon rate. Since the actual values of the absorption and scattering lengths are unknown, and the total attenuation is known only approximately, these parameters have to be inferred from the available data.

The distribution of hit multiplicity is expected to be well modelled by the simulation, since it reflects a very fundamental stage of the detection: determining whether a PMT was hit by a Cherenkov photon or not. Moreover, counting hits (and not e.g. charge) is the preferred method of characterising an event in a detector with many small PMTs and the time over threshold readout.

The distribution of the number of hits in the events detected by the plane during the first superrun is presented in Fig. 5.16. The bin with only two hits is discarded

due to the high content of random coincidences, not simulated in the Monte Carlo. The distribution is bimodal, with peaks around the lowest and the highest number of hits in an event. The broader peak at low multiplicities is probably caused by indirect light or muons whose Cherenkov cones only cover a few PMTs. The second peak, at the maximum value of 14 hits, likely corresponds to events with a muon travelling in the direction of the plane and the PMTs receiving large amounts of direct Cherenkov light.

The attenuation properties of water have a big effect on the shape of the hit multiplicity distribution, which allows for using the latter to test and adjust the simulation parameters. While an automated fit would be ideal, the generation of MC samples is a very computationally intensive process and cannot be realistically incorporated into an automated loop. On the other hand, a faster reweighting method would require a good parametrisation of the effects of absorption and scattering on the distribution. Therefore, as a first step, the tuning was performed manually. Monte Carlo samples with various combinations of absorption and scattering lengths were produced; all with the same muon inputs, generated by CRY, and with the outputs processed in the same way. The agreement of the hit multiplicity distributions with data was evaluated visually.

The findings suggest that low absorption length causes a drop of events in the tail of the distribution, i.e. for high multiplicities. This makes sense, as absorption decreases the total light yield and makes it less likely for all PMTs to be hit.

On the other hand, too high scattering length causes overall lower rates, especially in the low multiplicity part of the distribution. This can be explained by the fact that the plane only covers a small area of the wall, so there are more muons passing through the detector that do not emit Cherenkov light in its direction. Therefore, scattering may cause more indirect light to be scattered towards the plane PMTs than direct light to be scattered away. This would contribute particularly to the events with few hits, consistent with observation.

Finally, the total light attenuation should agree with the measured value of less than 3 m at 405 nm. The best fitting samples actually prefer higher attenuation lengths (i.e. cleaner water), but it is difficult to quantify the tension without a proper uncertainty on the measurement.

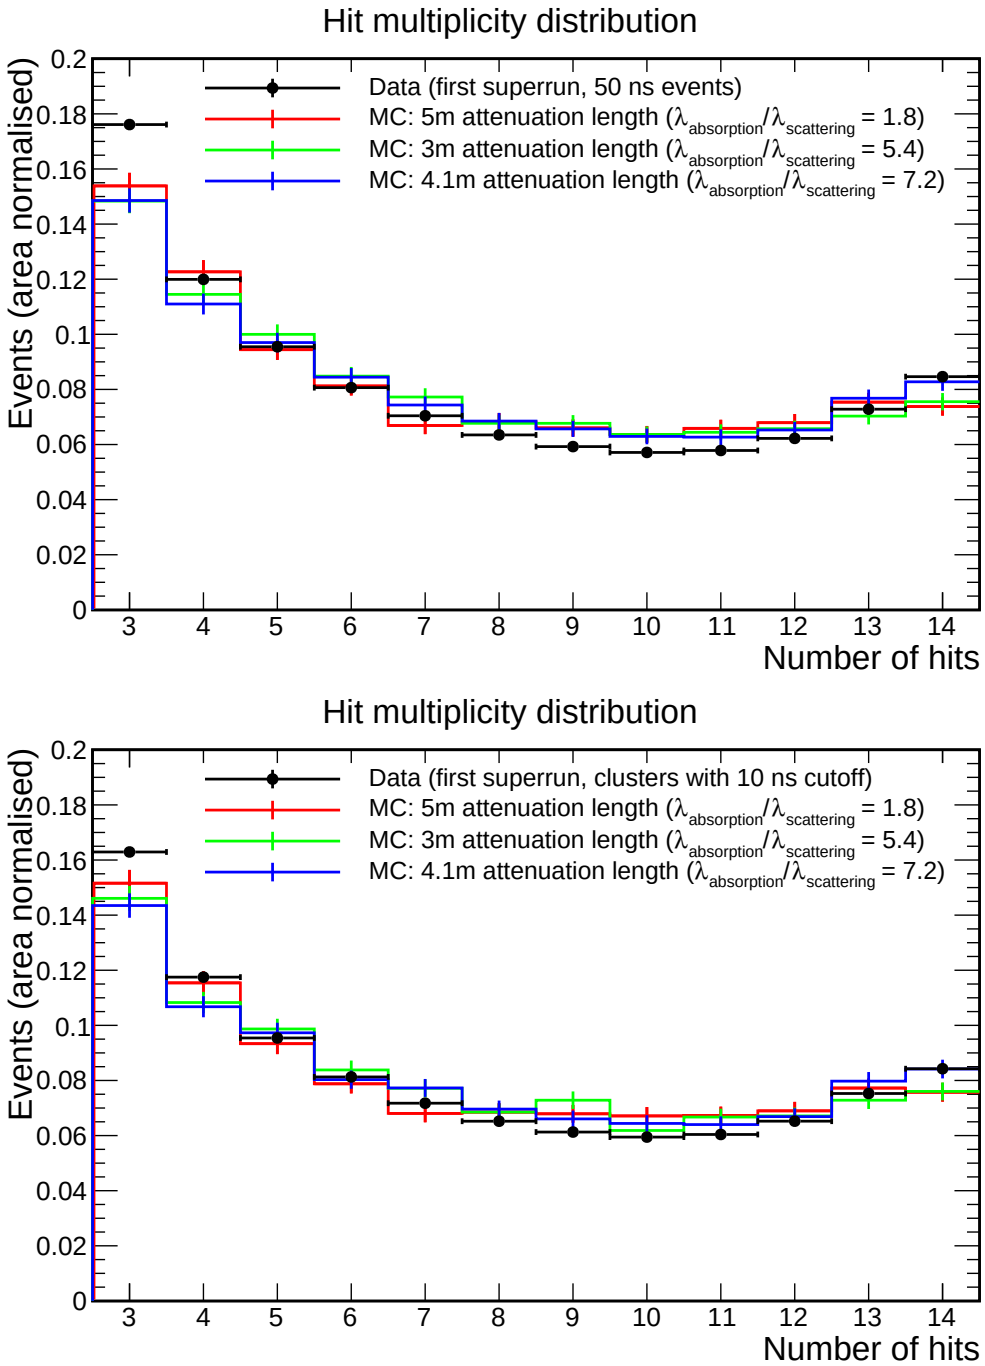

Figure 5.16: Distributions of the number of hits in the events with a 50 ns time window (top) and additional clustering with a 10 ns cutoff (bottom). The black markers show data points from the first superrun, with the default event definition. The solid lines correspond to the simulation samples with different water attenuation properties, listed in Table 5.2 with the indices: 3 (green), 5 (blue) and 6 (red). All distributions are area normalised, after rejecting bins with fewer than three hits.

	Att. length at 410 nm [m]			50 ns events		Clusters, 10 ns cutoff	
#	Total	Absorption	Scattering	χ^2/ndf	Rate [Hz]	χ^2/ndf	Rate [Hz]
1	3.0	8.4	4.7	10.98	144.2 ± 1.2	8.11	141.0 ± 1.3
2	3.0	13.8	3.8	11.93	134.3 ± 1.1	8.02	132.3 ± 1.1
3	3.0	19.1	3.6	7.21	130.5 ± 1.0	4.54	129.1 ± 1.1
4	3.0	24.5	3.4	12.27	128.0 ± 1.0	6.73	126.0 ± 1.0
5	4.1	33.6	4.7	5.93	129.8 ± 1.0	3.73	127.8 ± 1.0
6	5.0	14.0	7.8	5.12	145.0 ± 1.2	2.84	141.2 ± 1.3

Table 5.2: Total muon rate in CHIPS-M for six Monte Carlo samples considered to offer the best fit to data. The table lists the water attenuation properties used in each sample, the reduced χ^2 of the fit of hit multiplicity distributions to data (with free normalisation) and the corrected cosmic muon event rate. The last two values are shown both for the case of events with a 50 ns long window and clusters with a 10 ns cutoff. The uncertainties on the corrected rates are statistical only, coming mostly from the MC samples.

The attenuation properties for six Monte Carlo samples, which were considered to offer the best fit, are listed in Table 5.2. Since a clear labelling of the attenuation settings is not obvious, the numerical indices from the table are used to describe the samples. The hit multiplicity distributions for three of them, indexed as number 3, 5 and 6 are displayed in Fig. 5.16, shown in green, blue and red lines, respectively. The distributions are area normalised to eliminate the total rate from the comparison.

The error bars in the figures represent statistical uncertainty, which are only significant for the Monte Carlo. The exposure necessary to accumulate the cosmic muons in any of the MC samples, as provided by CRY, is 113 s or 1.6 minutes. In comparison, the data sample presented here contains 15 hours of events.

To evaluate the agreement between the observed and simulated hit multiplicities, the ratio between the MC and data distributions is calculated, and a straight horizontal line (with one free parameter) is fitted to the ratio. The χ^2 of the fit can be interpreted as a goodness of fit of the MC distribution to data, with a floating total normalisation. The values of the reduced χ^2 (with 11 degrees of freedom) are listed in Table 5.2 for each of the six samples, both for regular events and with an additional clustering requirement.

Based on the χ^2 values, sample 6 has has the best agreement with data, followed closely by sample 5, and sample 3. However, the χ^2 is not a fully reliable measure. This is because it is intended for comparing data points (which may have uncertainties) with a perfect model. But in this case, it is the model that has significant statistical errors, while the data sample is practically infinite. Therefore, samples with lower statistics can achieve a better χ^2 value even if the underlying model is in worse agreement. On the other hand the statistical errors cannot be neglected, because expected statistical fluctuations should not be counted against a good fit. Since all of the presented samples were produced with the same muon input files, the statistics are very similar, but there are still differences which can skew the comparison.

Sample number 5, with 4.1 m total attenuation length, was chosen as the reference for the other plots in this chapter, as it offers a good fit to data, while not deviating too far from the experimental measurement.

Sample number 6 has the largest values of the scattering and total attenuation lengths at 410 nm. In fact, it keeps the same relative weights of scattering to absorption as the original WCSim code. In contrast, samples 3 and 5 are dominated by a short scattering length. These differences manifest in the comparison of event duration distributions, shown in Fig. 5.17. Sample 6, with limited scattering, containts shorter events than the other two. Although it does not fully reproduce the tail seen in the data distribution for regular events, after clustering it offers the best agreement. For the samples controled by scattering, the situation is reversed; they reproduce the data distribution better if no clustering is required.

As seen in Fig. 5.16, the shapes of the hit multiplicity distributions are not strongly affected by the addition of clustering in the event definition. Nevertheless, the values in Table 5.2 suggest that it leads to better agreement between data and Monte Carlo and can slightly impact the final event rates.

It is worth noting that this search does not exhaust the parameter phase space and other settings might potentially produce even better agreement. The fit is strongly limited by the small number of MC samples, each with relatively low statistics.

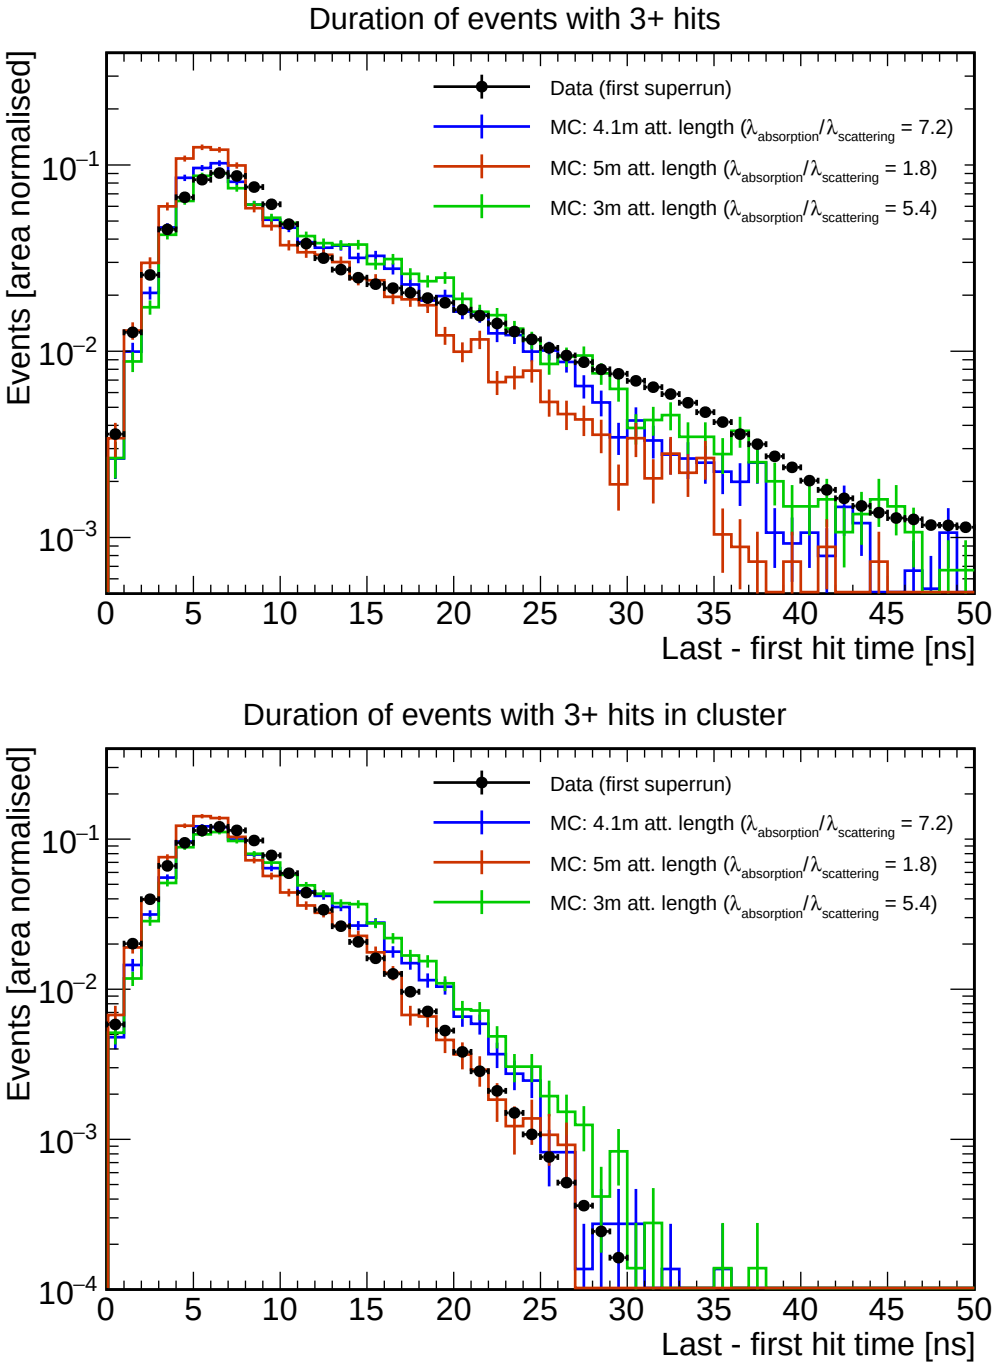

Figure 5.17: Distribution of event duration with a 50 ns time window (top) and an additional clustering with a 10 ns maximal separation (bottom). The black markers show data points from the first superrun, processed with time calibration, and a 12 ns ToT cut. The solid lines correspond to the simulation samples with different water attenuation properties, listed in Table 5.2 with the indices: 3 (green), 5 (blue) and 6 (red). All distributions are area normalised.

5.4.4 Acceptance correction and total event rate

After finding a combination of an event definition and Monte Carlo parameters that produce a reasonable agreement, the simulation can be used to estimate the geometric acceptance of the Nikhef plane. The acceptance is defined as the ratio of the number of observed plane events to the total number of events in the (hypothetical) fully instrumented CHIPS-M detector, as predicted by the simulation.

The detector events do not correspond directly to cosmic muons entering the volume. In fact, the simulation requires at least ten PMT hits in order to register an event at all. This limitation is related to how the software emulates an event trigger and was technically challenging to remove. A correction for this effect could be estimated from the distribution of the hit multiplicity (in the entire detector), but this was not done. The reason is that CHIPS-M instrumented with 6 % photocathode coverage contains 480 PMTs, and 10 hits represents only 2 % of them being active. Since the cosmic muons which cause only a few PMT hits (e.g. by clipping the detector edge) are not really relevant, this threshold is left in place.

The calculation is performed for each of the six final MC samples, where the plane events have to contain at least three hits with ToT higher than 12 ns (using the linear transformation from PE), and within a 50 ns time window; and separately with an additional requirement that the hits are part of a cluster with maximum 10 ns separation between hits. To obtain the total corrected rate measurement, the rate of observed data events, under the same definition and including time calibration, is divided by the acceptance. The final values are shown in Table 5.2. The uncertainties come from statistical errors of the Monte Carlo event counts. The uncertainty of the acceptance is calculated according to the binomial model, which is suitable since the efficiency is around 50 %, far from either 0 or 1.

The total rate values in Table 5.2 vary from 126 to 145 Hz, which is a relative difference of 15 %. Even more interestingly, almost all of this variability is covered between the two "best-fit" samples 5 and 6. There is no clear indication which of these results is more correct, and it is not obvious how to incorporate the difference into a measure of uncertainty. Therefore, the two extreme rate values for samples 5 and 6 are considered to represent an optimistic and a conservative estimate of the muon rate.

The optimistic result is the total rate obtained when using sample 5 without clustering and is equal to (127.8 ± 2.7) Hz, where the uncertainty includes the estimated 2 % shift due to the mis-modelling of the ToT cut (Sec. 5.4.1). The conservative rate measurement comes from sample 6 with clustering and is equal to (145.0 ± 3.1) Hz.

For comparison, the full detector muon rate in Monte Carlo, calculated from the number of detector events and the exposure time, varies from 120 to 122 Hz, depending on the attenuation parameters. Although it is a little smaller than the more optimistic measurement, it is definitely not a stark disagreement. One reason for this prediction to underestimate the data might be the simplified propagation of muon tracks through the 50 m or more of pit water.

5.5 Prediction of the cosmic muon rate in CHIPS-10

The result of the muon rate measurement in CHIPS-M can be used to predict the expected rate in a bigger module, such as CHIPS-10. This provides an experimental result to compare with other predictions, coming from the Monte Carlo simulation or calculations based on muon flux parametrisation.

To obtain the muon rate in CHIPS-10 from the CHIPS-M result, two corrections are required. The first one accounts for the size difference between the modules, particularly their surface areas. The second one corrects for the depth difference: the entire volume of CHIPS-M is at practically the same depth of 50 m, while CHIPS-10 extends from 40 to 60 m depth. Both corrections are based on the calculations presented in Ref. [94], which use a parametrisation of the muon intensity.

The difference in the rates of muons passing through two detector shapes depends on the angular distribution of the muon flux. At sea level, the total vertical intensity of muons with momenta above 1 GeV is approximately equal to $I_{SV} = 70\,/\mathrm{m}^2/\mathrm{s/sr}$ [96]. The flux is symmetric with respect to the azimuthal angle ϕ and the energy-averaged dependence on the zenith angle θ follows closely

$$I_S(\theta) = I_{SV} \cos^2 \theta \tag{5.2}$$

for θ between 0 and $\pi/2$ [96]. For higher zenith angles, the Earth provides a practically perfect shield, reducing the flux of muons to zero.

The zenith angle dependence of the flux varies with muon energy, and is flatter for higher energy particles (i.e. they are on average more horizontal). However, for shallow depths relevant to CHIPS, the dependence stays virtually unchanged [150].

To find the total rate of muons entering the detector at the sea level, C^S, the muon intensity must be integrated over the total surface area and the solid angle.

$$C^S = \int_S \int_\Omega I_S(\theta)(\vec{n} \cdot \vec{r}) d\Omega d\sigma, \tag{5.3}$$

where $\vec{r}$ is unit vector in the direction of the muon and $\vec{n}$ is a unit vector normal to the surface element $d\sigma$. The integral can be separated into two parts, for the top cap and the detector walls. The rate through the bottom cap is zero, because only muons passing from outside to inside are considered, which would limit the zenith angle to values above $\pi/2$, where the muon flux vanishes.

The top cap is assumed to be a flat surface with a total area A_{top}, and the normal vector is always equal to $\vec{n} = (0, 0, 1)$ in the coordinate system where the Z axis is vertical. The muon direction can be described by the polar angles θ and ϕ, which translated to Cartesian coordinates gives $\vec{r} = (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta)$. The inner product is therefore $\vec{n} \cdot \vec{r} = \cos\theta$. The solid angle element expressed in the polar coordinates is $d\Omega = \sin\theta d\theta d\phi$. The integral covers ϕ in the range from 0 to 2π and θ from 0 to $\pi/2$.

$$\begin{aligned} C^S_{\text{top}} &= A_{\text{top}} \int_0^{2\pi} d\phi \int_0^{\pi/2} d\theta I_S(\theta) \cos\theta \sin\theta = \\ &= A_{\text{top}} I_{SV} \int_0^{2\pi} d\phi \int_0^{\pi/2} \cos^3\theta \sin\theta d\theta = \\ &= A_{\text{top}} I_{SV} 2\pi \frac{1}{4} = \frac{\pi}{2} A_{\text{top}} I_{SV}. \end{aligned} \tag{5.4}$$

The integral over the side surface assumes that the detector shape is convex and that the sides are perpendicular to the base, which is satisfied for both the cylindrical CHIPS-10 and the octagonal prism of CHIPS-M. Because the muon flux is symmetric with respect to ϕ, the integral can be simplified by assuming that the normal surface vector always points in the same direction $\vec{n} = (1, 0, 0)$. The product of $\vec{n}$ and the

muon direction is then $\vec{n} \cdot \vec{r} = \sin\theta\cos\phi$. The azimuthal angle ϕ is integrated from $-\pi/2$ to $\pi/2$, in order to only count the entering muons. The muon rate through the walls with surface area A_{side} is

$$\begin{aligned} C^S_{\text{side}} &= A_{\text{side}} \int_{-\pi/2}^{\pi/2} d\phi \int_0^{\pi/2} d\theta I_S(\theta) \sin^2\theta \cos\phi = \\ &= A_{\text{side}} I_{SV} \int_{-\pi/2}^{\pi/2} \cos\phi d\phi \int_0^{\pi/2} \cos^2\theta \sin\theta d\theta = \\ &= A_{\text{side}} I_{SV} 2\frac{\pi}{16} = \frac{\pi}{8} I_{SV} A_{\text{side}}. \end{aligned} \tag{5.5}$$

The total muon rate at sea level is then

$$C^S = C^S_{\text{top}} + C^S_{\text{side}} = I_{SV}\frac{\pi}{8}(4A_{\text{top}} + A_{\text{side}}), \tag{5.6}$$

and the ratio between two detectors with surface areas $A^a_{\text{top}}, A^a_{\text{side}}$ and $A^b_{\text{top}}, A^b_{\text{side}}$, respectively, is

$$R^S_{a/b} = \frac{A^a_{\text{side}} + 4A^a_{\text{top}}}{A^b_{\text{side}} + 4A^b_{\text{top}}}. \tag{5.7}$$

The shape encompassed by the CHIPS-M liner is a regular octagonal prism with height $h = (3.32 \pm 0.10)\,\text{m}$ and side edge length $s = (1.32 \pm 0.10)\,\text{m}$, where the uncertainties represent an estimate of the possible deformation of the flexible liner. The top surface area is

$$A^{\text{CHIPS-M}}_{\text{top}} = 2(1+\sqrt{2})s^2 = (8.41 \pm 0.64)\,\text{m}^2, \tag{5.8}$$

and the wall surface area is

$$A^{\text{CHIPS-M}}_{\text{side}} = 8sh = (35.06 \pm 2.86)\,\text{m}^2. \tag{5.9}$$

CHIPS-10 is assumed to have a cylindrical shape with a height of 20 m and nominal diameter of 30 m. Because the design is not final and the possibility of a 25 m diameter is also considered, the following results are presented for both cases. Table 5.3 contains the top and side surface areas of CHIPS-10 and the geometric correction for the surface muon rate from CHIPS-M.

CHIPS-10 diameter	30 m	25 m
$A_{\text{top}}^{\text{CHIPS}-10}$	$706.9\,\text{m}^2$	$490.9\,\text{m}^2$
$A_{\text{side}}^{\text{CHIPS}-10}$	$1885\,\text{m}^2$	$1571\,\text{m}^2$
$R^S_{\text{CHIPS}-10/-\text{M}}$	68.6 ± 3.8	51.4 ± 2.9

Table 5.3: The CHIPS-10 top cap and side wall surface area, and the ratio of the muon rate at sea level in CHIPS-10 to CHIPS-M. The numbers are shown for two design options of CHIPS-10, with 30 m and 25 m diameter, and a common height of 20 m.

The second correction requires an expression for the dependence of muon flux on the water depth. Reference [151] presents the results of various experimental measurements and a precise parametrisation of the depth-intensity relation for muons with energies above 1 GeV (Fig. 5.18). In Ref. [94], a double exponential function is used to approximate this parametrisation for water depths up to 100 m. The obtained formula for the vertical muon intensity as a function of depth in metres, d, is

$$I_V^{1\,\text{GeV}}(d) = 51.32e^{-d/15.56} + 13.03e^{-d/60.67}. \tag{5.10}$$

Assuming an average muon energy loss in water of 220 MeV/m, the 1 GeV threshold corresponds to an additional 4.55 m of water. Therefore, the expression for intensity of all muons as a function of depth is $I_V(d) = I_V^{1\,\text{GeV}}(d - 4.55\,)$.

After adjusting Eq. 5.4 to account for depth, the rate of muons entering the detector through the top cap is equal to

$$C_{\text{top}}(z) = \frac{\pi}{2} A_{\text{top}} I_V(d), \tag{5.11}$$

where d is the total overburden, or the depth of the top cap. The rate of muons entering through the side walls is then (from Eq. 5.5)

$$C_{\text{side}} = \frac{\pi}{8} A_{\text{side}} \frac{1}{H} \int_d^{d+H} I_V(z) dz, \tag{5.12}$$

where H is the height of the detector, and the term $\frac{1}{H}\int I_V(z)dz$ represents the mean vertical intensity at the wall. The final ratio of the muon rates in two detectors,

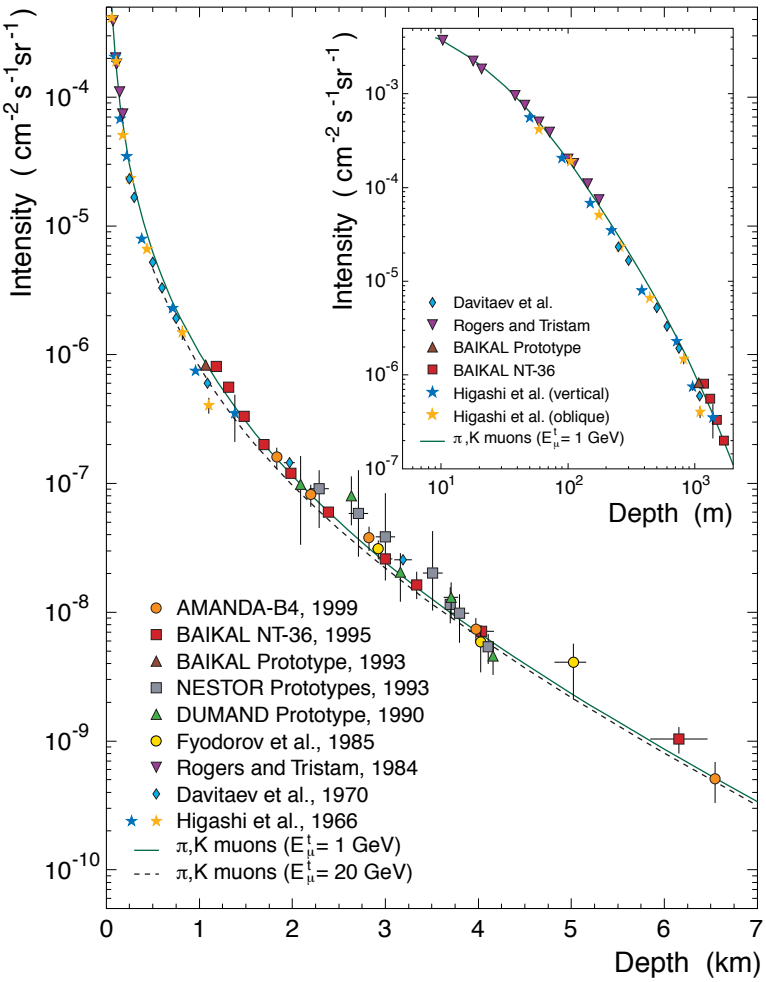

Figure 5.18: Intensity of vertical cosmic ray muons as a function of depth in water. The markers show experimental measurements and the solid line shows the functional parametrisation with muon energy threshold of 1 GeV. The inset focuses on shallow depths from 10 to 1000 m, relevant to this work. Figure taken from Ref. [151].

taking into account both the surface area and the depth, is

$$R_{a/b} = \frac{A^a_{\text{side}} \frac{1}{H^a} \int_{d^a}^{d^a+H^a} I_V(z)dz + 4A^a_{\text{top}} I_V(d^a)}{A^b_{\text{side}} \frac{1}{H^b} \int_{d^b}^{d^b+H^b} I_V(z)dz + 4A^b_{\text{top}} I_V(d^b)}, \tag{5.13}$$

where the superscripts a and b denote the parameters of the respective detectors.

CHIPS-M was deployed in a location in the Wentworth Pit, where the bottom is at a depth of (175 ± 2) feet or (53.5 ± 0.6) m. After accounting for the 1.2 m

high tripod and the 3.3 m height of the detector, the depth of the top cap is $d^{\text{CHIPS-M}} = (48.8 \pm 0.6)$ m. Using those numbers, the relevant muon intensity terms appearing in Eq. 5.13 are

$$I_V(d)^{\text{CHIPS-M}} = (9.27 \pm 0.18)\,/\text{m}^2/\text{s/sr} \tag{5.14}$$

$$\left(\frac{1}{H}\int I_V(z)dz\right)^{\text{CHIPS-M}} = (8.80 \pm 0.16)\,/\text{m}^2/\text{s/sr} \tag{5.15}$$

CHIPS-10 will be deployed in a deeper part of the pit, where it should extend from 60 m to 40 m depth. The muon intensity terms for CHIPS-10 are therefore

$$I_V(d)^{\text{CHIPS-10}} = 12.52\,/\text{m}^2/\text{s/sr} \tag{5.16}$$

$$\left(\frac{1}{H}\int I_V(z)dz\right)^{\text{CHIPS-10}} = 9.15\,/\text{m}^2/\text{s/sr}. \tag{5.17}$$

Table 5.4 shows the total correction factors from CHIPS-M and the final muon rate predictions, based on the optimistic and conservative results from Sec. 5.4.4, and for the two considered diameters.

CHIPS-10 diameter	30 m	25 m
$R_{\text{CHIPS-10/-M}}$	84.9 ± 4.9	62.8 ± 3.6
Optimistic rate in CHIPS-10	(10.85 ± 0.66) kHz	(8.03 ± 0.49) kHz
Conservative rate in CHIPS-10	(12.31 ± 0.75) kHz	(9.11 ± 0.56) kHz

Table 5.4: The ratio of the muon event rate in CHIPS-10 to CHIPS-M, taking into account the depth and size differences, and the final predicted values of the muon event rate in CHIPS-10, both in the optimistic and conservative scenario. The numbers are shown for two design options of CHIPS-10, with 30 m and 25 m diameter, and a common height of 20 m.

The values in Table 5.4 are the first ever prediction of the expected cosmic muon event rate in CHIPS-10 that are based on experimantal measurement. A preliminary estimate of the muon rate in CHIPS-10, using only the CRY package, provides the value of 10.6 kHz (8.2 kHz) for the 30 m (25 m) diameter option. This assumes that the fraction of muons causing detector events among those expected to hit the detector by the simple transport code is the same as in CHIPS-M, around 92 %. These values are very similar to the experimental results in the optimistic case.

On the other hand, a direct use of the numeric parametrisation of the muon intensity provides the values of 21 kHz (15 kHz) for the 30 m (25 m) diameter CHIPS-10. These numbers are significantly higher than either the data or CRY predictions, even the conservative ones. Part of this is due to the fact that the numeric calculation gives the total rate of muons entering the detector volume, while the above results only count events with observable PMT activity, with approximately 2 % of the PMTs being hit. However, most of the discrepancy is likely because the parametrisation is good at describing the depth dependence of the muon intensity, but not necessarily the absolute value. For the same reason, the overestimation of the total rate does not invalidate using the parametrised intensity in the geometric corrections from CHIPS-M to CHIPS-10, since it only appears in ratios, where most of the discrepancy should cancel out.

5.5.1 Event duration and dead time predictions

The main defence measure against the very large background from cosmic muons in CHIPS is the short beam spill time. Beam neutrino events can only occur in a 10 µs long window, opening every 1.3 s, and all other events are known to be background. Although the overall muon rate is very high, the chance of such an event occurring during a beam spill is relatively small. On average 0.123 ± 0.008 (0.091 ± 0.006) cosmic ray events per beam spill are expected for the 30 m (25 m) diameter option, using the conservative rate estimate.

Most muon events are easy to identify, since they have typically low zenith angles (coming from upwards) and at high energies, they usually pass through the entire detector. They can also be extensively studied with detector data collected outside of the beam spill time, making it possible to characterise the background very precisely. Nevertheless, the actual impact of cosmic muon events occurring during beam spill time on the electron appearance analysis needs to be studied further.

At the same time, even perfectly identified and vetoed cosmic events contribute to the dead time, by causing activity in the detector and potentially contaminating concurrent beam events. This is a very important consequence of the shallow overburden of the CHIPS detector. If the muon flux is too high, the muon events can completely swamp out any potential signal.

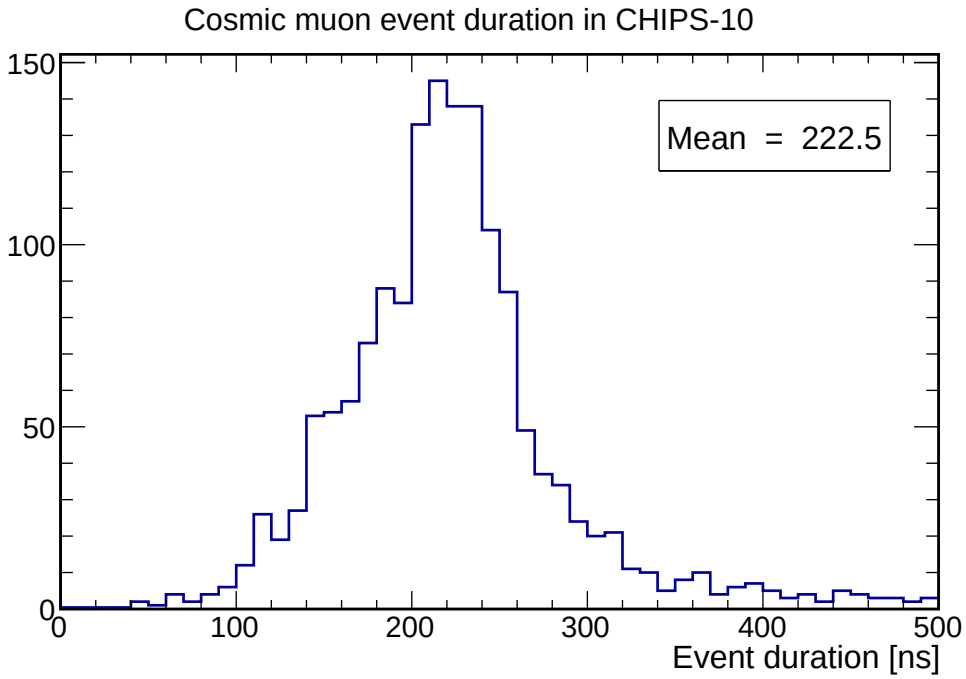

Figure 5.19: Duration of cosmic muon events in CHIPS-10, as predicted by the MC simulation with water attenuation set to 30 m.

To estimate the effective dead time, the average duration of a cosmic event in the CHIPS-10 detector was studied with the MC simulation. Figure 5.19 shows the distribution of the event duration for cosmic muons generated with CRY, where all simulated hits are included in the event.

With the average event duration of 223 ns and the typical frequency of cosmic muon events of 0.128 ± 0.008 per beam spill of 10 µs (for the conservative prediction with 30 m diameter), the effective dead time is (0.27 ± 0.02) %. This is a very small number, and suggests that the dead time due to cosmic ray muons will not be a significant issue in CHIPS-10. In fact, the result obtained here is significantly lower than the 2.5 % assumed conservatively in previous studies [93].

In addition, the reconstruction effort proved that in many cases it is still possible to fit and reconstruct both the through-going cosmic muon and the original beam event lepton, assuming they do not overlap.

Chapter 6

Summary and conclusions

This thesis presented CHIPS, an exciting and ambitious R&D detector project, which aims to develop water Cherenkov detectors for long baseline neutrino oscillation experiments, with a target cost of \$300 k per kt of fiducial mass. Such detectors could help with the ongoing search for CP violation in neutrino oscillations and the determination of neutrino mass hierarchy.

A location for prototype testing and detector deployment in the NuMI beam, the Wentworth Pit 2W, has been identified. The CHIPS collaboration is currently finalising the design of the CHIPS-10 detector, with the construction starting soon and the deployment of a first stage planned for summer 2018. The design is aided by a Monte Carlo detector simulation and a preliminary reconstruction algorithm, which allow the study of the performance of different configurations.

A physics sensitivity study using the GLoBES framework is presented. In contrast to previous work [93], this is the first prediction of the CHIPS physics reach based on the output of the reconstruction implementation. Although the results show that CHIPS-10 alone is too small to have a significant impact on the global sensitivity to CP violation and the mass hierarchy in the short term, which are dominated by NOvA and T2K experiments, larger detectors have the potential to make a bigger difference. This is especially true for the measurement of the actual value of the CP-violating mixing phase, δ_{CP}, where the wide energy spectrum of the NuMI beam at the Wentworth Pit allows CHIPS to reach a better resolution than other experiments. It is important to note that the predictions shown in this thesis are likely a lower bound on the actual potential of CHIPS, as the reconstruction and selection are not

yet fully optimised. Further improvements, in particular the enhanced rejection of neutral current events, may increase the sensitivity by as much as 30 % to 40 %.

Two testing campaigns at the Wentworth Pit have been completed, using a proof-of-principle detector model, CHIPS-M. The first test run in 2014 validated the basic design principles, including the use of a geomembrane liner and the sealing method, the water circulation system, as well as the remote operation of an underwater detector. CHIPS-M was recovered after spending a year at the bottom of the pit and redeployed with two prototype detection units, testing the operation of multiple PMTs with self-contained readout systems.

This thesis focuses on one of those prototypes, the Nikhef plane, which was constructed in collaboration with the KM3NeT experiment. The plane, which went through two major design iterations, contains the PMTs and readout hardware from one KM3NeT optical module. In the final version, the components are housed in a structure made out of inexpensive and robust PVC piping. The current design for the main type of planes to be installed in CHIPS-10, also utilising the KM3NeT-developed readout system, is based directly on the prototype Nikhef plane, with several improvements influenced by the issues encountered during testing.

During the second testing campaign, the prototype plane recorded a large sample of almost 9 million cosmic ray muon events. A preliminary analysis of this data is presented, describing the event reconstruction procedure, time calibration with an LED flasher, and the comparison between data and the predictions of a Monte Carlo simulation. The result is the first measurement of the total cosmic muon rate in CHIPS-M, at the bottom of the Wentworth Pit. Due to the uncertainty about the attenuation properties of the water in CHIPS-M, two results are presented (corresponding to different combinations of the absorption and scattering lengths), representing a lower and an upper bound on the final rate:

$$\mu \text{ rate}^{\,\text{CHIPS-M}}_{\text{optimistic}} = (127.8 \pm 2.7)\,\text{Hz}, \tag{6.1}$$

$$\mu \text{ rate}^{\,\text{CHIPS-M}}_{\text{conservative}} = (145.0 \pm 3.1)\,\text{Hz}. \tag{6.2}$$

The measured muon rate in CHIPS-M is used to make a prediction for the upcoming CHIPS-10 detector. A correction based on a parametrisation of the angular distribution of the muon flux and its depth dependence, taking into account

the differences in geometry and location between the two detectors, is applied. The conservative prediction, for two considered diameters of the CHIPS-10 module, is:

$$\mu \text{ rate } {}^{\text{CHIPS}-10}_{30\,\text{m}} = (12.31 \pm 0.75)\,\text{kHz}, \tag{6.3}$$

$$\mu \text{ rate } {}^{\text{CHIPS}-10}_{25\,\text{m}} = (9.11 \pm 0.56)\,\text{kHz}. \tag{6.4}$$

The results are slightly higher than predictions obtained with the CRY package, and significantly lower than previous calculations using only the muon intensity parametrisation.

Finally, the effective dead time due to identified cosmic muons occurring during beam spill time is calculated to be no more than $(0.27 \pm 0.02)\,\%$, based on the simulated average duration of a cosmic event in CHIPS-10.

Bibliography

[1] A. J. Perch, *Three-flavour neutrino oscillations with MINOS and CHIPS*. PhD thesis, University College London, 2017.

[2] D. Griffiths, *Introduction to Elementary Particles*. Physics textbook. Wiley, 2008.

[3] W. Pauli, *Dear Radioactive Ladies and Gentlemen*, Letter to the physical society of Tübingen, 1930.

[4] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, *Detection of the free neutrino: A Confirmation*, Science **124** (1956) 103–104.

[5] H. Bethe and R. Peierls, *The 'Neutrino'*, Nature **133** (1934) 532.

[6] F. Reines and C. L. Cowan, *Detection of the Free Neutrino*, Phys. Rev. **92** (Nov, 1953) 830–831.

[7] G. Danby, J.-M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger, *Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos*, Phys. Rev. Lett. **9** (Jul, 1962) 36–44. `http://link.aps.org/doi/10.1103/PhysRevLett.9.36`.

[8] Gargamelle Neutrino Collaboration, F. J. Hasert et al., *Observation of Neutrino Like Interactions Without Muon Or Electron in the Gargamelle Neutrino Experiment*, Phys. Lett. **46B** (1973) 138–140.

[9] Line Shape Sub-Group of the LEP Electroweak Working Group, DELPHI, LEP, ALEPH, OPAL, L3 Collaboration, *Combination procedure for the precise determination of Z boson parameters from results of the LEP experiments*, `arXiv:hep-ex/0101027`.

[10] DONUT Collaboration, K. Kodama et al., *Observation of tau neutrino*

interactions, Phys. Lett. **B504** (2001) 218–224, arXiv:hep-ex/0012035.

[11] Wikipedia, *Standard Model of Elementary Particles*, 2006. https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg. [Online; accessed 01-March-2017].

[12] J. A. Formaggio and G. P. Zeller, *From eV to EeV: Neutrino Cross Sections Across Energy Scales*, Rev. Mod. Phys. **84** (2012) 1307–1341, arXiv:1305.7513.

[13] D. Casper, *The Nuance neutrino physics simulation, and the future*, Nucl. Phys. Proc. Suppl. **112** (2002) 161–170, arXiv:hep-ph/0208030. [,161(2002)].

[14] Intensity Frontier Neutrino Working Group Collaboration, A. de Gouvea et al., *Working Group Report: Neutrinos*, in *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*. 2013. arXiv:1310.4340. https://inspirehep.net/record/1260555/files/arXiv:1310.4340.pdf.

[15] M. Fukugita and T. Yanagida, *Physics of Neutrinos: And Applications to Astrophysics*. Physics and astronomy online library. Springer, 2003.

[16] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics*. OUP Oxford, 2007.

[17] T. Dealtry and T. collaboration, *Muon neutrino disappearance at T2K*, Journal of Physics: Conference Series **598** (2015) no. 1, 012018. http://stacks.iop.org/1742-6596/598/i=1/a=012018.

[18] D. H. Perkins, *Particle Astrophysics, Second Edition*. Oxford Master Series in Physics. OUP Oxford, 2009.

[19] A. Blake, S. Germani, Y. B. Pan, A. J. Perch, M. M. Pfützner, J. Thomas, and L. H. Whitehead, *CHIPS Event Reconstruction and Design Optimisation*, arXiv:1612.04604.

[20] B. Pontecorvo, *Mesonium and anti-mesonium*, Sov. Phys. JETP **6** (1957) 429.

[Zh. Eksp. Teor. Fiz.33,549(1957)].

[21] J. N. Bahcall, *Solving the mystery of the missing neutrinos*, arXiv:physics/0406040 [physics].

[22] J. M. LoSecco, *The History of "Anomalous" Atmospheric Neutrino Events: A First Person Account*, Phys. Perspect. **18** (2016) no. 2, 209–241, arXiv:1606.00665 [physics.hist-ph].

[23] J. N. Bahcall, *Solar Neutrinos. I. Theoretical*, Phys. Rev. Lett. **12** (Mar, 1964) 300–302. http://link.aps.org/doi/10.1103/PhysRevLett.12.300.

[24] R. Davis, D. S. Harmer, and K. C. Hoffman, *Search for Neutrinos from the Sun*, Phys. Rev. Lett. **20** (May, 1968) 1205–1209. http://link.aps.org/doi/10.1103/PhysRevLett.20.1205.

[25] A. I. Abazov, O. L. Anosov, E. L. Faizov, V. N. Gavrin, A. V. Kalikhov, T. V. Knodel, I. I. Knyshenko, V. N. Kornoukhov, S. A. Mezentseva, I. N. Mirmov, A. V. Ostrinsky, A. M. Pshukov, N. E. Revzin, A. A. Shikhin, P. V. Timofeyev, E. P. Veretenkin, V. M. Vermul, G. T. Zatsepin, T. J. Bowles, B. T. Cleveland, S. R. Elliott, H. A. O'Brien, D. L. Wark, J. F. Wilkerson, R. Davis, K. Lande, M. L. Cherry, and R. T. Kouzes, *Search for neutrinos from the Sun using the reaction* $^{71}\mathrm{Ga}(\nu_e, e^-)^{71}\mathrm{Ge}$, Phys. Rev. Lett. **67** (Dec, 1991) 3332–3335. https://link.aps.org/doi/10.1103/PhysRevLett.67.3332.

[26] GALLEX Collaboration, P. Anselmann et al., *Solar neutrinos observed by GALLEX at Gran Sasso.*, Phys. Lett. **B285** (1992) 376–389.

[27] K. S. Hirata et al., *Results from one thousand days of real-time, directional solar-neutrino data*, Phys. Rev. Lett. **65** (Sep, 1990) 1297–1300. https://link.aps.org/doi/10.1103/PhysRevLett.65.1297.

[28] SNO Collaboration, B. Aharmim et al., *Electron energy spectra, fluxes, and day-night asymmetries of B-8 solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory*, Phys. Rev. **C72** (2005) 055502, arXiv:nucl-ex/0502021.

[29] KamLAND Collaboration, K. Eguchi et al., *First results from KamLAND: Evidence for reactor anti-neutrino disappearance*, Phys. Rev. Lett. **90** (2003) 021802, arXiv:hep-ex/0212021.

[30] B. T. Cleveland, T. Daily, J. Raymond Davis, J. R. Distel, K. Lande, C. K. Lee, P. S. Wildenhain, and J. Ullman, *Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector*, The Astrophysical Journal **496** (1998) no. 1, 505. http://stacks.iop.org/0004-637X/496/i=1/a=505.

[31] F. Kaether, W. Hampel, G. Heusser, J. Kiko, and T. Kirsten, *Reanalysis of the Gallex solar neutrino flux and source experiments*, Physics Letters B **685** (2010) no. 1, 47 – 54. http://www.sciencedirect.com/science/article/pii/S0370269310000729.

[32] SAGE Collaboration, J. N. Abdurashitov et al., *Measurement of the solar neutrino capture rate with gallium metal. III: Results for the 2002–2007 data-taking period*, Phys. Rev. **C80** (2009) 015807, arXiv:0901.2200.

[33] Super-Kamiokande Collaboration, J. Hosaka et al., *Solar neutrino measurements in super-Kamiokande-I*, Phys. Rev. **D73** (2006) 112001, arXiv:hep-ex/0508053.

[34] Super-Kamiokande Collaboration, J. P. Cravens et al., *Solar neutrino measurements in Super-Kamiokande-II*, Phys. Rev. **D78** (2008) 032002, arXiv:0803.4312.

[35] Super-Kamiokande Collaboration, K. Abe et al., *Solar neutrino results in Super-Kamiokande-III*, Phys. Rev. **D83** (2011) 052010, arXiv:1010.0118.

[36] Y. Nakano, *Solar neutrino results from Super-Kamiokande.*, in *38th International Conference on High Energy Physics, Chicago, USA, 3-10 August 2016*. 2016. https://indico.cern.ch/event/432527/contributions/1072129/attachments/1319706/1979670/ichep2016_nakano_ver5.pdf.

[37] SNO Collaboration, B. Aharmim et al., *Combined Analysis of all Three Phases of Solar Neutrino Data from the Sudbury Neutrino Observatory*, Phys. Rev. **C88** (2011) 025501, arXiv:1109.0763.

[38] Borexino Collaboration, G. Bellini et al., *Measurement of the solar 8B neutrino rate with a liquid scintillator target and 3 MeV energy threshold in the Borexino detector*, Phys. Rev. **D82** (2010) 033006, arXiv:0808.2868.

[39] Borexino Collaboration, G. Bellini et al., *Precision measurement of the 7Be*

solar neutrino interaction rate in Borexino, Phys. Rev. Lett. **107** (2011) 141302, arXiv:1104.1816.

[40] Borexino Collaboration, G. Bellini et al., *Neutrinos from the primary proton-proton fusion process in the Sun*, Nature **512** (08, 2014) 383–386. http://dx.doi.org/10.1038/nature13702.

[41] KamLAND Collaboration, A. Gando et al., *Constraints on θ_{13} from A Three-Flavor Oscillation Analysis of Reactor Antineutrinos at KamLAND*, Phys. Rev. **D83** (2011) 052002, arXiv:1009.4771.

[42] K. Hirata et al., *Experimental study of the atmospheric neutrino flux*, Physics Letters B **205** (1988) no. 2, 416 – 420. http://www.sciencedirect.com/science/article/pii/0370269388916905.

[43] R. Becker-Szendy et al., *Electron- and muon-neutrino content of the atmospheric flux*, Phys. Rev. D **46** (Nov, 1992) 3720–3724. https://link.aps.org/doi/10.1103/PhysRevD.46.3720.

[44] T. Kajita, *Atmospheric neutrino results from Super-Kamiokande and Kamiokande – Evidence for ν_μ oscillations*, Nuclear Physics B - Proceedings Supplements **77** (1999) no. 1, 123 – 132. http://www.sciencedirect.com/science/article/pii/S0920563299004077.

[45] LSND Collaboration Collaboration, A. Aguilar et al., *Evidence for neutrino oscillations from the observation of $\overline{\nu}_e$ appearance in a $\overline{\nu}_\mu$ beam*, Phys. Rev. D **64** (Nov, 2001) 112007. https://link.aps.org/doi/10.1103/PhysRevD.64.112007.

[46] K2K Collaboration, T. Ishida, *The First results of K2K long baseline neutrino oscillation experiment*, AIP Conf. Proc. **549** (2000) 772–776, arXiv:hep-ex/0008047. [,772(2000)].

[47] OPERA Collaboration, N. Agafonova et al., *Discovery of τ Neutrino Appearance in the CNGS Neutrino Beam with the OPERA Experiment*, Phys. Rev. Lett. **115** (Sep, 2015) 121802. http://link.aps.org/doi/10.1103/PhysRevLett.115.121802.

[48] MINOS Collaboration, P. Adamson et al., *Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS*,

Phys. Rev. Lett. **110** (2013) no. 25, 251801, arXiv:1304.6335.

[49] P. Vahle, *New results from NOvA*, in *XXVII International Conference on Neutrino Physics and Astrophysics (Neutrino 2016), London, 4-9 July 2016*. 2016. http://neutrino2016.iopconfs.org/IOP/media/uploaded/EVIOP/event_948/11.30_-_has_to_be_PDF.pdf.

[50] H. Tanaka, *Status, recent results and plans for T2K*, in *XXVII International Conference on Neutrino Physics and Astrophysics (Neutrino 2016), London, 4-9 July 2016*. 2016. http://neutrino2016.iopconfs.org/IOP/media/uploaded/EVIOP/event_582/T2K_Neutrino_2016.pdf.

[51] S. Moriyama, *New atmospheric and solar results from Super-Kamiokande*, in *XXVII International Conference on Neutrino Physics and Astrophysics (Neutrino 2016), London, 4-9 July 2016*. 2016. http://neutrino2016.iopconfs.org/IOP/media/uploaded/EVIOP/event_948/neutrino2016-moriyama-pub-2.pdf.

[52] Daya Bay Collaboration, F. P. An et al., *Observation of electron-antineutrino disappearance at Daya Bay*, Phys. Rev. Lett. **108** (2012) 171803, arXiv:1203.1669.

[53] RENO Collaboration, J. K. Ahn et al., *Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment*, Phys. Rev. Lett. **108** (2012) 191802, arXiv:1204.0626.

[54] T2K Collaboration, K. Abe et al., *Evidence of Electron Neutrino Appearance in a Muon Neutrino Beam*, Phys. Rev. **D88** (2013) no. 3, 032002, arXiv:1304.0841.

[55] MINOS Collaboration, P. Adamson et al., *Electron neutrino and antineutrino appearance in the full MINOS data sample*, Phys. Rev. Lett. **110** (2013) no. 17, 171801, arXiv:1301.4581.

[56] Z. Yu, *Recent Results from the Daya Bay Experiment*, in *XXVII International Conference on Neutrino Physics and Astrophysics (Neutrino 2016), London, 4-9 July 2016*. 2016. http://neutrino2016.iopconfs.org/IOP/media/uploaded/EVIOP/event_948/09.20___2_.pdf.

[57] K. K. Joo, *Results from RENO and prospects with RENO-50*, in *XXVII*

International Conference on Neutrino Physics and Astrophysics (Neutrino 2016), London, 4-9 July 2016. 2016. http://neutrino2016.iopconfs.org/IOP/media/uploaded/EVIOP/event_948/09.45___2_.ppt.

[58] M. Ishitsuka, *New Results of Double Chooz*, in *51st Rencontres de Moriond conference, La Thuile, 12-19 March 2016*. 2016. https://indico.in2p3.fr/event/12279/session/3/contribution/173/material/slides/0.pdf.

[59] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, and T. Schwetz, *Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity*, JHEP **01** (2017) 087, arXiv:1611.01514.

[60] Nu-Fit, *v3.0: Three-neutrino fit based on data available in November 2016*, 2016. http://www.nu-fit.org.

[61] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, *Evidence for the 2π Decay of the K_2^0 Meson*, Phys. Rev. Lett. **13** (Jul, 1964) 138–140. https://link.aps.org/doi/10.1103/PhysRevLett.13.138.

[62] K. Abe et al., *Letter of Intent: The Hyper-Kamiokande Experiment — Detector Design and Physics Potential —*, arXiv:1109.3262.

[63] X. Qian and P. Vogel, *Neutrino Mass Hierarchy*, Prog. Part. Nucl. Phys. **83** (2015) 1–30, arXiv:1505.01891.

[64] IceCube Collaboration, M. G. Aartsen et al., *PINGU: A Vision for Neutrino and Particle Physics at the South Pole*, arXiv:1607.02671.

[65] KM3NeT Collaboration, U. F. Katz, *The ORCA Option for KM3NeT*, in *Proceedings of the 15th International Workshop on Neutrino Telescopes (Neutel 2013): Venice, March 11-15, 2013*. 2014. arXiv:1402.1022. https://inspirehep.net/record/1280181/files/arXiv:1402.1022.pdf.

[66] I. Ostrovskiy and K. O'Sullivan, *Search for neutrinoless double beta decay*, Mod. Phys. Lett. **A31** (2016) no. 18, 1630017, arXiv:1605.00631. [Erratum: Mod. Phys. Lett.A31,no.23,1692004(2016)].

[67] V. N. Aseev et al., *Upper limit on the electron antineutrino mass from the Troitsk experiment*, Phys. Rev. D **84** (Dec, 2011) 112003.

http://link.aps.org/doi/10.1103/PhysRevD.84.112003.

[68] F. Fraenkle, *Status of the neutrino mass experiments KATRIN and Project 8*, in *The European Physical Society Conference on High Energy Physics, Vienna, 22-29 July 2015*. 2015. https://pos.sissa.it/archive/conferences/234/084/EPS-HEP2015_084.pdf.

[69] K. N. Abazajian et al., *Cosmological and Astrophysical Neutrino Mass Measurements*, Astropart. Phys. **35** (2011) 177–184, arXiv:1103.5083 [astro-ph.CO].

[70] E. Giusarma, M. Gerbino, O. Mena, S. Vagnozzi, S. Ho, and K. Freese, *Improvement of cosmological neutrino mass bounds*, Phys. Rev. **D94** (2016) no. 8, 083522, arXiv:1605.04320 [astro-ph.CO].

[71] MiniBooNE Collaboration, A. A. Aguilar-Arevalo et al., *Improved Search for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations in the MiniBooNE Experiment*, Phys. Rev. Lett. **110** (2013) 161801, arXiv:1303.2588.

[72] G. Mention, M. Fechner, T. Lasserre, T. A. Mueller, D. Lhuillier, M. Cribier, and A. Letourneau, *The Reactor Antineutrino Anomaly*, Phys. Rev. **D83** (2011) 073006, arXiv:1101.2755 [hep-ex].

[73] P. Huber, *Reactor antineutrino fluxes – Status and challenges*, Nuclear Physics B **908** (2016) no. Supplement C, 268 – 278. http://www.sciencedirect.com/science/article/pii/S0550321316300505. Neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015.

[74] MINOS, Daya Bay Collaboration, P. Adamson et al., *Limits on Active to Sterile Neutrino Oscillations from Disappearance Searches in the MINOS, Daya Bay, and Bugey-3 Experiments*, Phys. Rev. Lett. **117** (2016) no. 15, 151801, arXiv:1607.01177. [Addendum: Phys. Rev. Lett.117,no.20,209901(2016)].

[75] IceCube Collaboration, M. G. Aartsen et al., *Searches for Sterile Neutrinos with the IceCube Detector*, Phys. Rev. Lett. **117** (2016) no. 7, 071801, arXiv:1605.01990.

[76] M. Bass, *The Short Baseline Neutrino Oscillation Program at Fermilab*, PoS **ICHEP2016** (2016) 481, arXiv:1702.00990 [physics.ins-det].

[77] S. Amerio et al., *Design, construction and tests of the ICARUS T600 detector*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **527** (2004) no. 3, 329 – 410. http://www.sciencedirect.com/science/article/pii/S0168900204004966.

[78] T2K Collaboration, K. Abe et al., *The T2K Experiment*, Nucl. Instrum. Meth. **A659** (2011) 106–135, arXiv:1106.1238.

[79] NOvA Collaboration, D. S. Ayres et al., *The NOvA Technical Design Report*, .

[80] P. Adamson et al., *The NuMI Neutrino Beam*, Nucl. Instrum. Meth. **A806** (2016) 279–306, arXiv:1507.06690.

[81] Hyper-Kamiokande Working Group Collaboration, K. Abe et al., *A Long Baseline Neutrino Oscillation Experiment Using J-PARC Neutrino Beam and Hyper-Kamiokande*, 2014. arXiv:1412.4673. https://inspirehep.net/record/1334360/files/arXiv:1412.4673.pdf.

[82] Hyper-Kamiokande proto-collaboration Collaboration, K. Abe et al., *Hyper-Kamiokande Design Report*, tech. rep., KEK preprint, 2016. https://lib-extopc.kek.jp/preprints/PDF/2016/1627/1627021.pdf.

[83] R. Wendell, *Status of Hyper-Kamiokande*, in *1st Neutrino Oscillation on Tomography Workshop: Earthquake Research Institute, the University of Tokyo, 7-8 January 2016*. 2016. https://indico.cern.ch/event/442108/contributions/1946160/attachments/1208832/1762436/wendell_hk_status_20160108_upload.pdf.

[84] M. Shiozawa, *Hyper-Kamiokande design*, in *11th International Workshop on Next generation Nucleon Decay and Neutrino Detectors: Toyama, Japan, December 13 - 16, 2010*. 2010. http://www-sk.icrr.u-tokyo.ac.jp/NNN10/slides/15pm-Shiozawa.pdf.

[85] DUNE Collaboration, R. Acciarri et al., *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE)*, arXiv:1512.06148.

[86] A. Weber, *Hybrid Detector Concepts*, in *DUNE ND Workshop, Fermilab, 9-10 June 2017*. https://indico.fnal.gov/getFile.py/access?contribId=6&

sessionId=0&resId=0&materialId=slides&confId=14029.

[87] U.S. Department Of Energy Office Of Project Management Oversight And Assessments, *Independent Cost Review of the Long Baseline Neutrino Facility/Deep Underground Neutrino Experiment Project at Fermi National Accelerator Laboratory*, tech. rep., August, 2015. https://web.fnal.gov/organization/OPSS/Projects/LBNFDUNE/LBNF%20APM%20Jul%202015/Review%20Documents/LBNF%20DUNE%20ICR_Final%20Report.pdf.

[88] D. Normile, *Neutrino Observatory Suffers Accident*, Science (2001) . http://www.sciencemag.org/news/2001/11/neutrino-observatory-suffers-accident.

[89] U.S. Department of Energy Office of Science, *Review of Options for Underground Science*, tech. rep., June, 2011. https://science.energy.gov/~/media/np/pdf/review_of_underground_science_report_final.pdf.

[90] R. Laasch, *Fermilab achieves milestone beam power for neutrino experiments*, Fermilab News (2017) . http://news.fnal.gov/2017/01/fermilab-achieves-milestone-beam-power-neutrino-experiments/.

[91] NOvA Collaboration, M. W. McGee, C. Ader, K. Anderson, J. Hylen, and M. Martens, *Thermal and structural stability of medium energy target carrier assembly for NOvA at Fermilab*, Conf. Proc. **C100523** (2010) THPEC042, arXiv:1207.4460.

[92] K. T. McDonald, *An Off-axis neutrino beam*, arXiv:hep-ex/0111033.

[93] P. Adamson et al., *CHerenkov detectors In mine PitS (CHIPS) Letter of Intent to FNAL*, in *Proceedings, Community Summer Study 2013: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*. 2013. arXiv:1307.5918. https://arxiv.org/pdf/1307.5918.pdf.

[94] S. V. Cao, J. Huan, K. Lang, and F. Nova, *Cosmic Ray Muon Rates in Water Cherenkov Detectors at Shallow Depths*, Internal document, CHIPS-doc-1, 2013. http://nile.hep.utexas.edu/DocDB/ut-numi/docs/0000/000001/001/CHIPS-CR-rate-v3.pdf.

[95] WHO, *Guidelines for drinking-water quality, fourth edition*, 2011. http://www.who.int/water_sanitation_health/publications/2011/

dwq_guidelines/en/.

[96] Particle Data Group Collaboration, C. Patrignani et al., *Review of Particle Physics*, Chin. Phys. **C40** (2016) no. 10, 100001.

[97] Super-Kamiokande Collaboration, Y. Fukuda et al., *The Super-Kamiokande detector*, Nucl. Instrum. Meth. **A501** (2003) 418–462.

[98] F. Amat, P. Bizouard, J. Bryant, T. Carroll, S. D. Rijck, S. Germani, T. Joyce, B. Kriesten, M. Marshak, J. Meier, J. Nelson, A. Perch, M. Pfützner, R. Salazar, J. Thomas, J. Trokan-Tenorio, P. Vahle, R. Wade, C. Wendt, L. Whitehead, and M. Whitney, *Measuring the attenuation length of water in the CHIPS-M water Cherenkov detector*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **844** (2017) 108 – 115. http://www.sciencedirect.com/science/article/pii/S016890021631186X.

[99] P. Dills, *CHIPS frame update*, Internal presentation, CHIPS-doc-490, 2017.

[100] P. M. Kooijman, *Planar Optical Module*, Internal presentation, CHIPS-doc-500, 2017.

[101] BeagleBoard.org website. https://beagleboard.org/black. accessed 29 August 2017.

[102] White Rabbit Project. http://www.ohwr.org/projects/white-rabbit/wiki. accessed 2 April 2017.

[103] S. Wolcott, *CHIPS Detector Design*, Internal presentation, CHIPS-doc-505, 2017.

[104] GEANT4 Collaboration, S. Agostinelli et al., *GEANT4: A simulation toolkit*, Nucl. Instrum. Meth. **A506** (2003) 250–303.

[105] J. Allison et al., *Geant4 developments and applications*, IEEE Trans. Nucl. Sci. **53** (2006) 270.

[106] WCSim. https://github.com/WCSim/WCSim. accessed 18 March 2017.

[107] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis*

framework, Nucl. Instrum. Meth. **A389** (1997) 81–86.

[108] C. Andreopoulos et al., *The GENIE Neutrino Monte Carlo Generator*, Nucl. Instrum. Meth. **A614** (2010) 87–104, arXiv:0905.2517.

[109] C. Hagmann, D. Lange, and D. Wright, *Cosmic-ray shower generator (CRY) for Monte Carlo transport codes*, in *2007 IEEE Nuclear Science Symposium Conference Record*, vol. 2, pp. 1143–1146. Oct, 2007.

[110] R. B. Patterson, E. M. Laird, Y. Liu, P. D. Meyers, I. Stancu, and H. A. Tanaka, *The Extended-track reconstruction for MiniBooNE*, Nucl. Instrum. Meth. **A608** (2009) 206–224, arXiv:0902.2222.

[111] R. B. Patterson, *A Search for Muon Neutrino to Electron Neutrino Oscillations at* $\Delta m^2 > 0.1\ eV^2$. PhD thesis, Princeton U., 2007. http://lss.fnal.gov/cgi-bin/find_paper.pl?thesis-2007-19.

[112] T2K Collaboration, S. Cartwright, *Energy Measurement in the T2K Oscillation Analysis*, PoS **NUFACT2014** (2015) 103.

[113] M. M. Pfützner, *Maximum likelihood method in track reconstruction for CHIPS*, First year report, University College London, 2014. https://www.hep.ucl.ac.uk/twiki/pub/HEPGroup/PhD1stYearTalks2014/report-mp.pdf.

[114] P. Huber, M. Lindner, and W. Winter, *Simulation of long-baseline neutrino oscillation experiments with GLoBES (General Long Baseline Experiment Simulator)*, Comput. Phys. Commun. **167** (2005) 195, arXiv:hep-ph/0407333.

[115] P. Huber, J. Kopp, M. Lindner, M. Rolinec, and W. Winter, *New features in the simulation of neutrino oscillation experiments with GLoBES 3.0: General Long Baseline Experiment Simulator*, Comput. Phys. Commun. **177** (2007) 432–438, arXiv:hep-ph/0701187.

[116] P. Huber, J. Kopp, M. Lindner, M. Rolinec, and W. Winter, *General Long Baseline Experiment Simulator. User's and experiment definition manual*, 2007. https://www.mpi-hd.mpg.de/personalhomes/globes/documentation/globes-manual-3.0.8.pdf.

[117] G. L. Fogli, E. Lisi, A. Marrone, D. Montanino, and A. Palazzo, *Getting the most from the statistical analysis of solar neutrino oscillations*, Phys. Rev. **D66** (2002) 053010, arXiv:hep-ph/0206162.

[118] M. D. Messier, *Evidence for neutrino mass from observations of atmospheric neutrinos with Super-Kamiokande*, . UMI-99-23965.

[119] E. A. Paschos and J. Y. Yu, *Neutrino interactions in oscillation experiments*, Phys. Rev. **D65** (2002) 033002, hep-ph/0107261.

[120] T. Kato, *Tau neutrino appearance via neutrino oscillations in atmospheric neutrinos*. PhD thesis, SUNY, Stony Brook, 2007. http://proquest.umi.com/pqdweb?sid=1&Fmt=2&clientId=79356&RQT=309&VName=PQD&did=1490074251.

[121] F. Dufour, T. Kajita, E. Kearns, and K. Okumura, *Further study of neutrino oscillation with two detectors in Kamioka and Korea*, Phys. Rev. **D81** (2010) 093001, arXiv:1001.5165.

[122] P. Huber, M. Lindner, T. Schwetz, and W. Winter, *First hint for CP violation in neutrino oscillations from upcoming superbeam and reactor experiments*, arXiv:0907.1896.

[123] NOvA Collaboration, I. Ambats et al., *NOvA proposal to build a 30-kiloton off-axis detector to study neutrino oscillations in the Fermilab NuMI beamline*, hep-ex/0503053.

[124] NOvA Collaboration, T. Yang and S. Woijcicki, *Study of physics sensitivity of $\nu_m u$ disappearance in a totally active version of NoVA detector*, Off-Axis-Note-SIM-30.

[125] R. Patterson, *Recent progress of NOvA*, in *International Workshop on Next Generation Nucleon Decay and Neutrino Detectors (NNN16), Institute of High Energy Physics (IHEP), Chinese Academy of Sciences in Beijing, 3-5 November 2016*. 2016. http://indico.ihep.ac.cn/event/6156/session/5/contribution/9/material/slides/0.pdf.

[126] K. Abe et al., *Proposal for an Extended Run of T2K to 20×10^{21} POT*, arXiv:1609.04111.

[127] P. Huber, M. Lindner, and W. Winter, *Superbeams versus neutrino factories*, Nucl. Phys. **B645** (2002) 3–48, hep-ph/0204352.

[128] Y. Itow et al., *The JHF-Kamioka neutrino project*, hep-ex/0106019.

[129] M. Ishitsuka, T. Kajita, H. Minakata, and H. Nunokawa, *Resolving neutrino mass hierarchy and CP degeneracy by two identical detectors with different baselines*, Phys. Rev. **D72** (2005) 033003, hep-ph/0504026.

[130] P. Guzowski. Private communication, 2013.

[131] G. L. Fogli, E. Lisi, A. Marrone, D. Montanino, A. Palazzo, and A. M. Rotunno, *Global analysis of neutrino masses, mixings and phases: entering the era of leptonic CP violation searches*, Phys. Rev. **D86** (2012) 013012, arXiv:1205.5254 [hep-ph].

[132] T. Benson, *CHIPS-M Structure Design*, Internal presentation, CHIPS-doc-137, 2014.

[133] F. Halzen and S. R. Klein, *IceCube: An Instrument for Neutrino Astronomy*, Rev. Sci. Instrum. **81** (2010) 081101, arXiv:1007.1247.

[134] L. Whitehead, *The CHIPS Experiment*, in *Seminar given at the EPP Group, University of Warwick, 10 December 2015.* 2015. http://www.hep.ucl.ac.uk/~lwhitehead/chipsSeminarWarwick2015.pdf.

[135] KM3NeT Collaboration, P. Bagley et al., *KM3NeT: Technical Design Report for a Deep-Sea Research Infrastructure in the Mediterranean Sea Incorporating a Very Large Volume Neutrino Telescope*, .

[136] KM3Net Collaboration, S. Adrian-Martinez et al., *Letter of intent for KM3NeT 2.0*, J. Phys. **G43** (2016) no. 8, 084001, arXiv:1601.07459.

[137] KM3NeT website. https://www.km3net.org/artists-impression/. accessed 14 April 2017.

[138] KM3NeT Collaboration, D. Vivolo, *The KM3NeT Digital Optical Module*, in *Proceedings, 7th Very Large Volume Neutrino Telescope Workshop (VLVnT 2015): Rome, Italy, September 14-16, 2015*, vol. 116, p. 01002. 2016.

[139] *Photomultiplier Tubes – Basics and Applications.* Hamamatsu Photonics,

3a ed., 2007.

[140] KM3NeT Collaboration, S. Aiello, L. Classen, V. Giordano, O. Kalekin, E. Leonora, H. Peek, J. Reubelt, D. Samtleben, and E. Visser, *Characterization of the 80-mm diameter Hamamatsu PMTs for the KM3NeT project*, AIP Conf. Proc. **1630** (2014) 118–121.

[141] D. Gajanana, V. Gromov, and P. Timmer, *ASIC design in the KM3NeT detector*, Journal of Instrumentation **8** (2013) no. 02, C02030. http://stacks.iop.org/1748-0221/8/i=02/a=C02030.

[142] H. Band. Private communication, 2015.

[143] G. Martin-Chassard, S. Conforti, F. Dulucq, M. El Berni, C. de La Taille, and W. Wei, *PARISROC, a photomultiplier array readout chip (PMm2 collaboration)*, Nuclear Instruments and Methods in Physics Research A **623** (Nov., 2010) 492–494, arXiv:0912.2915.

[144] R. M. Pope and E. S. Fry, *Absorption spectrum (380–700 nm) of pure water. II. Integrating cavity measurements*, Appl. Opt. **36** (Nov, 1997) 8710–8723. http://ao.osa.org/abstract.cfm?URI=ao-36-33-8710.

[145] L. Whitehead. Private communication, 2014.

[146] R. Bruijn and D. Samtleben. Private communication, 2017.

[147] KM3NeT Collaboration, S. Adrián-Martínez et al., *Deep sea tests of a prototype of the KM3NeT digital optical module*, Eur. Phys. J. **C74** (2014) no. 9, 3056, arXiv:1405.0839.

[148] F. Kaether and C. Langbrandtner, *Transit time and charge correlations of single photoelectron events in R7081 photomultiplier tubes*, Journal of Instrumentation **7** (2012) no. 09, P09002. http://stacks.iop.org/1748-0221/7/i=09/a=P09002.

[149] R. Abbasi et al., *Calibration and characterization of the IceCube photomultiplier tube*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **618** (2010) no. 1–3, 139 – 152. http://www.sciencedirect.com/science/article/pii/S0168900210006662.

[150] L. N. Bogdanova, M. G. Gavrilov, V. N. Kornoukhov, and A. S. Starostin, *Cosmic muon flux at shallow depths underground*, Phys. Atom. Nucl. **69** (2006) 1293–1298, arXiv:nucl-ex/0601019.

[151] E. V. Bugaev, A. Misaki, V. A. Naumov, T. S. Sinegovskaya, S. I. Sinegovsky, and N. Takahashi, *Atmospheric muon flux at sea level, underground and underwater*, Phys. Rev. **D58** (1998) 054001, arXiv:hep-ph/9803488.

[152] A. Buckley, *A class for typesetting academic theses*, 2010.