

Bachelorarbeit

aus der Physik

Search for periodic neutrino emission from pulsars with
KM3NeT

Robert Petri

Betreuer: PD Dr. Thomas Eberl, Dr. Rodrigo Gracia-Ruiz

Erlangen Centre for Astroparticle Physics

Abgabedatum: 20-02-2026

Abstract

KM3NeT has two neutrino detectors, ARCA (reconstruction threshold of 100GeV – 1TeV) and ORCA (reconstruction threshold of 3GeV – 5TeV), submerged in the Mediterranean Sea. These are Cherenkov detectors designed to detect high-energy neutrinos.

This thesis presents the building of a pipeline for a search for a possible periodic low-energy neutrino emission, with the idea that a possible source of such emissions is pulsars. Instead of relying on reconstructed neutrino events, the analysis is based directly on the PMT counting rates. This approach removes the energy threshold imposed by track reconstruction and enables sensitivity to lower-energy neutrino fluxes, in principle as low as cross-section and flux allow it.

To identify periodic signals, the counting-rate time series is transformed into the frequency domain using the Fast Fourier Transformation (FFT). The Fourier power at the known spin frequencies of the pulsars is used as the test statistic.

A sensitivity study is performed to determine the minimum neutrino flux required to produce a detectable periodic signal within the analysed data set. The obtained sensitivity is a function of the neutrino energy, but it roughly corresponds to a flux of $10^{17} \text{ m}^{-2} \text{ s}^{-1} \text{ GeV}^{-1}$ at 2.5 GeV at Earth taking the considered 9-hour exposure as basis. For a catalogue of five pulsars, the necessary neutrino Luminosity is calculated. Depending on the distance and duty cycle, one gets values in the range of $L_\nu \sim 10^{54}$ to $10^{60} \text{ GeV}^{-1} \text{ s}^{-1}$.

Contents

1	Introduction	1
1.1	Pulsars	1
1.2	Neutrinos associated with pulsars	2
1.3	KM3NeT	4
1.3.1	Detector: Overview and Layout	4
1.3.2	Energy threshold for the detector and detection principle	5
1.4	Goal of this thesis	7
2	Analysis Prerequisite	8
2.1	Mathematical Background	8
2.1.1	Definition and Properties of Fast Fourier Transform	8
2.1.1.1	Discrete Fourier Transform	8
2.1.1.2	Fast Fourier Transform (FFT)	9
2.1.1.3	Properties of the FFT	10
2.1.1.4	Zero padding	11
2.1.1.5	Scalloping effect	11
2.1.2	Definition and Properties of the Test Statistic used	12
2.1.2.1	Fourier coefficients and choice of observable	12
2.1.2.2	Define hypotheses	13
2.1.2.3	Define a test statistic	13
2.1.2.4	Null distribution	14
2.1.2.5	Observed test statistic	14
2.1.2.6	Definition of the p-value	14
2.1.2.7	Decision rule	14
2.2	Ingredients for Analysis and the pipeline	15
2.2.1	Detector rate	15
2.2.1.1	Justification of total detector rate	15
2.2.1.2	Extracting total detector rate and combining different runs	17
2.2.2	Correct time series	17
2.2.2.1	Barycentric correction	17
2.2.3	Selected Pulsars	19
3	Analysis	20
3.1	FFT analysis	20
3.1.1	Red noise	20
3.1.2	Harmonic summation	20

3.1.3	Statistical significance	21
3.2	Verifying the pipeline	21
3.3	Sensitivity Study	22
3.3.1	Signal Injection and Significance Estimation	22
3.3.2	Cross section and effective volume	24
3.3.2.1	Cross section	24
3.3.2.2	Definition of the Effective Volume	25
3.3.3	Neutrino propagation and flavour ratio at Earth	26
3.3.3.1	Case A:	26
3.3.3.2	Case B:	26
3.3.4	Necessary Rate	27
3.3.4.1	Detection Threshold	28
3.3.4.2	Flux Sensitivity	28
3.3.4.3	Result	28
3.3.4.4	Comparison with ANTARES	28
4	Future prospects	31
A	Appendix A: Implementation Details of the Analysis Pipeline	33
A.1	Extract total detector rate	33
A.2	Combining time series	34
A.3	Correct time series	34
A.4	Analysis	36
B	Appendix B: Implementation Details of the Sensitivity Pipeline	37
B.1	Inject Signal and Significance Estimation	37
B.2	Cross-section	38
B.3	Effective volume	38
B.4	Necessary Rate	38
C	Appendix C: Plots	39
	Bibliography	45

Chapter 1

Introduction

For millennia, humans have looked up into the sky at night and wondered what is out there. A new area was unlocked when telescopes became available. And then, when the digital revolution allowed us to also record events in the sky, not only on paper but digitally as well, a new area was unlocked again. The rise of global communication networks made it possible to send data around the world in seconds. Ultimately, the deployment of space-based telescopes expanded access to additional regions of the electromagnetic spectrum and reduced the atmospheric interference that affects ground-based observations.

At some point, we figured out that light is not the only messenger reaching us from the cosmos; other particles also carry information. The first event detected through multiple messengers was Supernova 1987A, during which neutrinos from the core collapse were observed several hours before the visible light of the stellar explosion. This was the first supernova observed via both neutrinos and photons [1]. Supernovae are among the most powerful explosions in the universe, capable of outshining entire galaxies, as they have peak luminosities of order $10^9 - 10^{10} L_{\odot}$, corresponding to absolute magnitudes around -17 to -19 [2].

We discovered holes in the fabric of spacetime from which nothing can escape once entered. And when such objects merge, they produce gravitational waves that can be observed by detectors such as LIGO [3].

1.1 Pulsars

We also discovered objects that behave like extremely precise clocks. These objects are known as pulsars. Pulsars are neutron stars with masses comparable to the solar mass and radii on the order of 10 km. They rotate very rapidly, with periods ranging from a few milliseconds to hundreds of seconds. The fastest known pulsar is PSR J1748-2446ad [4], which has a rotational frequency of 716.35 Hz. This corresponds to a rotational velocity at its equator of approximately $0.24c$.

Pulsars also have strong magnetic fields, typically on the order of 10^{11} G to 10^{15} G [5]. For comparison, contemporary laboratory technology is capable of generating magnetic fields of only short duration and on the order of 1.65×10^5 G [6].

These magnetic fields make pulsars highly efficient natural particle accelerators, particularly near their magnetic poles. As a consequence, strong electromagnetic emission comes from these regions.




Figure 1.1: Schematic of a radio pulsar with the misaligned spin and magnetic axes, magnetic field lines (blue) and beamed radiation cones (redish). Credit: NASA/Goddard Space Flight Centre Conceptual Image Lab.

Pulsars differ from ordinary neutron stars in that the axis of rotation is not aligned with the magnetic axis, a schematic of which can be seen in Figure 1.1. This misalignment causes the emission beams to sweep through space as the pulsar rotates, producing the appearance of discrete pulses for an observer.

1.2 Neutrinos associated with pulsars

There are models that attempt to describe those objects. Some of them [7, 8] also propose the idea that neutrinos are produced alongside the electromagnetic radiation. The model described in [7], for example, proposes a model that produces low-energy neutrinos in the energy range from approximately 100 keV to 3 MeV via pair-annihilation processes in the hot stellar plasma. In this scenario, accretion columns become optically thick across magnetic field lines and are dominated by advection. Radiation generated higher up in the column diffuses downward rather than escaping, leading to an accumulation of energy at the base of the column and a runaway increase in temperature. A schematic of this effect can be seen in Figure 1.2, where the orange shaded region is the proposed region of neutrino production. As a result, electron-positron pairs become abundant. Photon cooling is inefficient due to radiation trapping, and the dominant cooling channel becomes

$$e^- + e^+ \rightarrow \nu\bar{\nu}.$$

Since the neutrino emissivity scales as $Q_\nu \propto T^9$, the neutrino luminosity could, in this case, exceed the photon luminosity.

Another model, as described in [8], proposes that high-energy (TeV–PeV) neutrinos can be produced in pulsar wind nebulae via interactions of relativistic protons in the pulsar wind with ambient matter and radiation fields, leading to charged pion production and




Figure 1.2: Schematic view of an accreting pulsar and its accretion columns. The orange shaded area shows the region of possible neutrino production [7]

subsequent decay primarily through inelastic proton–proton collisions within the nebula:

$$p + p \rightarrow \pi^{\pm}, \pi^0 + X$$

$$\pi^+ \rightarrow \mu^+ + \nu_{\mu}$$

$$\pi^- \rightarrow \mu^- + \bar{\nu}_{\mu}$$

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_{\mu}$$

$$\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_{\mu}$$

Neutrinos may also be produced through photomeson interactions $p + \gamma$, although proton–proton interactions are expected to dominate [8].

But as of yet, there is no neutrino emitted from a pulsar that has been identified as such.




Figure 1.3: This shows the frequency distributions of pulsars in the ATNF Pulsar Catalogue, as well as in the HMXRB and LMXRB catalogues. Taken from [5]

Nevertheless, there is no shortage of pulsars that may serve as potential targets. For example, an extensive compilation of known pulsars and their properties, derived from multiple surveys, is provided by the Australia Telescope National Facility (ATNF)

Pulsar Catalogue [9]. The distribution of their period and, hence, their frequency is depicted in Figure 1.3.

Recently, large-volume neutrino telescopes have been built or are planned to be built to look for astrophysical neutrinos. Prominent examples include ANTARES, IceCube, and P-ONE in the Pacific Ocean, as well as KM3NeT. The present work concentrates on the detectors of the KM3NeT collaboration.

But given the large catalogue of known pulsars and the growing neutrino telescope data sets, a future detection of a neutrino associated with a pulsar appears plausible.

1.3 KM3NeT

The KM3NeT Collaboration [10] is responsible for the construction and oversight of the ARCA and ORCA detectors. The designation KM3NeT derives from KM3, indicating a cubic-kilometre scale, and NeT, representing a neutrino telescope. ARCA stands for Astroparticle Research with Cosmics in the Abyss, while ORCA refers to Oscillation Research with Cosmics in the Abyss.

1.3.1 Detector: Overview and Layout

The ORCA detector is located in the Mediterranean Sea at approximately $42^{\circ}48\text{N}$ latitude and $6^{\circ}02\text{E}$ longitude, at a depth of about 2450 m below the sea surface, with an additional 30 m offset. The site corresponds to the French KM3NeT location near Toulon [11, 12]. And the ARCA detector is also in the Mediterranean Sea at approximately $36^{\circ}16\text{N}$ latitude and $16^{\circ}06\text{E}$ longitude, at a depth of about 3500 m below the sea surface, including a 72 m offset. This corresponds to the Italian KM3NeT site near Capo Passero [13, 11, 14].

The KM3NeT/ARCA detector is specifically designed to detect high-energy neutrinos from astrophysical sources in the TeV–PeV energy range.

In contrast, the KM3NeT/ORCA detector is optimised for neutrino physics, oscillation parameters and neutrino mass ordering [15] and also allows for probing physics beyond the Standard Model [16, 17] in the neutrino sector using atmospheric neutrinos in the GeV–TeV energy range.

Although both configurations share the same geometric concept, their dimensions differ due to the distinct physics objectives. A schematic of the layout is depicted in Figure 1.4. The detection system is based on three-dimensional assemblies of light-sensitive elements enclosed in pressure-resistant glass housings known as Digital Optical Modules (DOMs) [19]. Each DOM contains 31 photomultiplier tubes (PMTs), together with the required readout electronics and calibration devices. The PMTs are oriented to cover nearly the entire 4π steradian solid angle, with a higher density in the lower hemisphere to improve sensitivity to particles arriving from below.

These DOMs are mounted on Detection Units (DUs), which are flexible vertical structures arranged in a hexagonal grid on the seafloor. Their upright orientation is maintained by the combined buoyant force of the DOMs and an additional buoy located at the top. Every DU supports 18 DOMs suspended along two parallel ropes, while a central cable extends along the full length of the unit for the transfer of the data.

The KM3NeT/ORCA configuration will consist of 115 such DUs, spaced horizontally by roughly 20m. Along each unit, adjacent DOMs are separated vertically by about




Figure 1.4: Schematic layout of the KM3NeT/ARCA detector, composed of two building blocks each with 115 detection units, instrumented with 18 DOMs spaced vertically by 36m, and horizontal detector unit spacing 95m [18]

9m. Once fully deployed, the detector will instrument a total volume of approximately $7 \times 10^6 \text{m}^3$ [20].

On the other hand, the KM3NeT/ARCA has an average vertical distance between DOMs along a detection unit of 36m, and the DUs are placed with an average spacing of 95m. The total volume is planned to be approximately 1km^3 [21].

1.3.2 Energy threshold for the detector and detection principle

Due to the differing geometrical configurations of ORCA and ARCA, their effective energy thresholds are also different. For ORCA, the threshold lies in the range 3–5 GeV, whereas for ARCA it is approximately 100 GeV–1 TeV.

The general idea of detecting a neutrino via KM3NeT is as follows:

The neutrino, itself not visible, interacts with matter as shown in Figure 1.5 via charged (CC) or neutral (NC) current. The CC current produces a charged lepton. When this propagates through water, it travels faster than the local speed of light in the medium. This produces detectable Cherenkov radiation.

Each recorded signal in a PMT is referred to as a hit. The rate is then the number of hits in a given time slice. This time slice for KM3NeT is 100ms. In the latter text, it might also be referred to as a time series. Those data-taking periods are called runs and are 3 hours long. A normal run has about 10^6 bins as can be seen in Figure 1.6.

By correlating the arrival times of hits with the known three-dimensional positions of the PMTs, coincident signals across multiple PMTs can be identified. Such coincidences are likely to originate from the same secondary particle and enable the reconstruction of both the direction and the energy of the incident neutrino.

The so-called golden channel corresponds to charged-current interactions producing a muon, which causes a characteristic track-like event. In contrast, other interaction channels typically produce shower-like (cascade) events. While cascade events generally provide better energy resolution, track-like events are better for direction reconstruction. The amount of energy transferred to the charged lepton or hadronic shower and thus



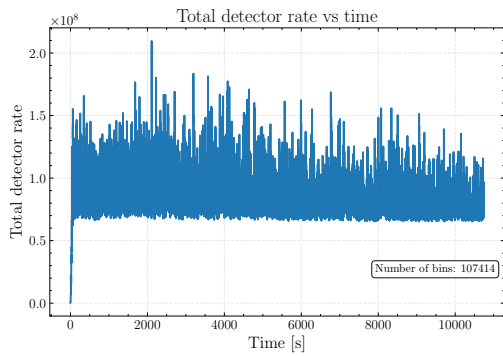



Figure 1.5: Interaction channels of neutrinos via charged-current (CC) and neutral-current (NC) processes. Graph adapted from [22].



(a) ORCA-18 run 00018992, with 107414 bins. (b) ORCA-18 run 00018994, with 107592 bins.



(c) ORCA-18 run 00018993, with 107449 bins.

Figure 1.6: This graph shows the total detector rate for the ORCA site for three runs (data taking time of 3 hours) on an effective DOM.

the intensity of Cherenkov light detected by the PMTs, is constrained by the initial neutrino energy. At low neutrino energies, the light yield is insufficient either to generate

enough coincident hits for reliable reconstruction or to propagate to multiple DUs. This limitation ultimately determines the effective energy thresholds of ARCA and ORCA.

1.4 Goal of this thesis

A single photon hit on a DOM is insufficient to associate the signal with a track and thus reconstruct the properties of the neutrino. Nevertheless, it still contains limited information about the interaction. A sufficiently large flux of low-energy neutrinos, particularly if modulated periodically, would produce a statistically detectable signature in the data.

Essentially, that is the idea for this work. The motivation is that pulsars are expected not only to be luminous in photons but also to emit neutrinos by the processes described in section 1.2. In fact, theoretical models suggest that pulsars may be even brighter in neutrinos than in photons [7]. However, these neutrinos are predicted to be produced at low energies, up to a few GeV.

At such low energies, it is hard to trigger events from neutrino interactions due to the effect described in 1.3.2 and thus infer the arrival direction of the neutrino. For this reason, this analysis focuses on the KM3NeT photomultiplier tube (PMT) counting rate rather than on reconstructed events. Using a small catalogue of known pulsars, a barycentric correction is applied to the PMT rate time series. That is, in order to perform a consistent timing analysis, the effects of the Earth's orbital and rotational motion must be accounted for. This is done by transforming the timestamps to the Solar System Barycentre (SSB), thereby correcting for the detector's motion along the line of sight to the source. Further is explained in 2.2.2.1. And the power spectrum obtained from a fast Fourier transform (FFT), which quantifies the strength of periodic components in the time series as a function of frequency, is analysed at the known rotational (spin) frequencies of the pulsars.

Noteworthy is that this assumes that the periodicity of the neutrino emission corresponds to the observed photon periodicity. It is also worth mentioning that the rate is dominated by photons originating from potassium-40 decay and bioluminescence. However, since an FFT-based analysis is used and the time distribution of the background is approximately constant, it is not expected to have a significant impact on the resulting power spectrum. Although a big luminescent fish in the immediate vicinity of a DOM would push the DOM into the high veto rate, large-scale biological activity in the detector environment may still induce slow variations in the PMT counting rate. Since such variations occur on timescales of minutes to hours or months, they contribute only to low-frequency components of the power spectrum and are not expected to affect the analysis at pulsar spin frequencies.

In addition, a sensitivity study is performed to determine what periodic neutrino flux, and at which energies, would be required at Earth to achieve a detection significance of 5σ . This study will assume a sinusoidal neutrino signal. Since any periodic signal can be expressed as a sum of sinusoidal components, this assumption should be sufficient for an initial sensitivity estimate until theoretical models predict a more specific emission profile for neutrinos.

Chapter 2

Analysis Prerequisite

2.1 Mathematical Background

2.1.1 Definition and Properties of Fast Fourier Transform

The Fast Fourier Transform (FFT) is not a new concept, but in order to understand the FFT, one must first understand the Discrete Fourier Transform (DFT). The general idea is that, after transforming a series in real space, one obtains a representation in k -space, also known as frequency space. For our goal of finding the frequency content of a signal in the PMT rate over time of KM3NeT, this is an important aspect.

Generally, if there are questions on the Fourier transformation, it is advisable to read [23], as it is comprehensive and, at the time of writing, accessible online free of charge. The definitions and properties of the DFT/FFT can therefore be found there, with some of them also discussed in [5].

2.1.1.1 Discrete Fourier Transform

Given a sequence $x_0, x_1, x_2, \dots, x_{n-1}$ where $x_m \in \mathbb{C}$. In the present analysis, however, the time series is real-valued, i.e. $x_m \in \mathbb{R}$, the DFT is defined as

$$X_k = \sum_{m=0}^{n-1} x_m \exp\left(-\frac{2\pi i m k}{n}\right), \quad k = 0, \dots, n-1.$$

This is a linear transformation on \mathbb{C}^n , and its matrix representation is an $n \times n$ DFT matrix F with entries

$$F_{km} = \exp\left(-\frac{2\pi i m k}{n}\right).$$

When introducing $\omega_n = e^{-2\pi i/n}$, one can write the matrix F as

$$F_{km} = (\omega_n)^{km},$$

where ω_n is a primitive n -th root of unity.

The DFT can be understood as expressing the input vector in the basis of complex exponentials

$$\left\{ \exp\left(-\frac{2\pi i m k}{n}\right) \right\}_{k=0}^{n-1}.$$

As can be easily checked, these exponentials form an orthogonal (up to scaling) basis of \mathbb{C}^n .

The inverse DFT exists and is given by

$$x_m = \frac{1}{n} \sum_{k=0}^{n-1} X_k \exp\left(\frac{2\pi i m k}{n}\right).$$

The prefactor varies depending on the definition one uses. But the important thing is that since the inverse exists, we can from that conclude that F is indeed invertible and unitary up to a scaling factor.

Furthermore, F has a well-defined structure, as it is built from roots of unity in the complex plane. In particular, the cyclic and factorisation properties of these roots enable recursive decompositions of the transform. The following part about the FFT will make use of that.

As is easily seen, computing the DFT requires $O(n^2)$ operations, since it corresponds to multiplying an $n \times n$ matrix by a vector.

However, this is not the optimal scaling for performing a Fourier transformation. The computational complexity can be reduced to $O(n \log_2 n)$. Algorithms that do this improvement are known as Fast Fourier Transforms.

2.1.1.2 Fast Fourier Transform (FFT)

The Fast Fourier Transform (FFT) is not a single algorithm, but a class of algorithms that compute the Discrete Fourier Transform efficiently by exploiting its algebraic structure. In particular, the FFT makes use of the properties of roots of unity and the associated group structure of the cyclic group, which is isomorphic to $\mathbb{Z}/n\mathbb{Z}$. The most widely used realisation of these ideas is the Cooley-Tukey algorithm, which implements a divide-and-conquer strategy.

In Figure 2.1, one can see the difference in scaling between a naive DFT and an optimised FFT implementation.

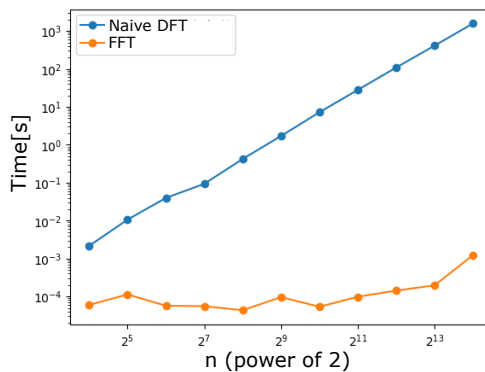


Figure 2.1: Difference in scaling of a naive DFT (blue) and an FFT (orange) algorithm

Cooley-Tukey Divide-and-Conquer Algorithm

The Cooley-Tukey Divide-and-Conquer Algorithm [24] is one such algorithm that speeds up the DFT to become an FFT.

Assume that $n = AB$ with integers $A, B > 1$. The DFT sum

$$X_k = \sum_{m=0}^{n-1} x_m \omega_n^{mk}$$

can then be rewritten by decomposing the index m as

$$m = m_1 + Am_2, \quad m_1 = 0, \dots, A-1, \quad m_2 = 0, \dots, B-1.$$

This gives then

$$X_k = \sum_{m_2=0}^{B-1} \left(\sum_{m_1=0}^{A-1} x_{m_1+Am_2} \omega_n^{m_1k} \right) \omega_n^{Am_2k}.$$

The inner sum is then a DFT of size A . The outer sum corresponds to a DFT of size B with additional multiplicative phase factors

$$\omega_n^{Am_2k}.$$

They are known as twiddle factors. By recursively applying this decomposition, the original DFT of size n is reduced to smaller DFTs, ultimately giving the total computational complexity of $O(n \log_2 n)$.

In the case where $n = 2^r$, this procedure leads to the radix-2 FFT, which is particularly efficient. This is also what will later be used for speeding up the FFT in the analysis.

2.1.1.3 Properties of the FFT

The FFT inherits its properties from the DFT and from the normal Fourier transformation.

Frequency resolution

The FFT relies on a finite time series of length N , sampled at a finite non-zero sampling frequency f_s . This implies a finite frequency resolution (bin spacing) Δf given by

$$\Delta f = \frac{f_s}{N} \tag{2.1}$$

But since the observation time T is related to N and f_s via $T = N/f_s$, one can substitute that into 2.1 to get the more intuitive expression for the frequency resolution:

$$\Delta f = \frac{1}{T}$$

For KM3NeT, and in this case with $N = 2^{18.3}$, corresponding to a total observation time of $T = 9$ h, and a sampling interval of 100 ms (i.e. $f_s = 10$ Hz), the resulting frequency resolution is

$$\Delta f \simeq 3.1 \times 10^{-5} \text{ Hz}. \tag{2.2}$$

This value is the smallest frequency separation that can be resolved in the FFT for my parameters.

Nyquist frequency and aliasing

Another important effect is the Nyquist frequency, which means that only frequencies up to

$$f_{\text{Nyq}} = \frac{f_s}{2}$$

can be unambiguously reconstructed from the sampled data.

Again, for my parameters, that is:

$$f_{\text{Nyq}} = \frac{10\text{Hz}}{2} = 5\text{ Hz} \quad (2.3)$$

Frequency components with $f > f_{\text{Nyq}}$ are subject to aliasing [23] and appear at lower frequencies in the FFT spectrum.

2.1.1.4 Zero padding

Zero padding is the process of appending zeros to the time-domain signal before computing the FFT. That means that the initial length N is now N' with $N' > N$. Mathematically, zero padding does not change the underlying discrete-time signal; it only changes the sampling of its discrete Fourier transform (DFT) in frequency space. To be noted:

$$\Delta f' = \frac{f_s}{N'} < \Delta f$$

But this is only an interpolation of the spectrum, not an increase in true resolving power.

Advantages:

- Smoother spectra: Zero padding increases the number of FFT bins, making spectral features appear smoother. This makes them easier to visualise.
- Improved peak localisation: The position of spectral peaks can be estimated more accurately by interpolating between bins.
- Convenient FFT sizes: Padding to powers of two can significantly improve computational efficiency.

The convenient size argument is the reason why zero padding will be used later on.

2.1.1.5 Scalloping effect

Let a discrete-time signal contain a single sinusoidal component with frequency f , observed over a finite time interval T .

If the signal frequency satisfies $fT \in \mathbb{Z}$, the sinusoid is exactly aligned with the centre of a DFT bin and all of its power is contained in said single frequency bin. If instead

$$fT \notin \mathbb{Z},$$

the signal frequency lies between two DFT bins, and its power is distributed over all frequency bins. This effect is known as the scalloping effect.

For a rectangular time window, the power contribution to the k -th DFT bin is given by

$$P_k = P_0 \text{sinc}^2[\pi(k - fT)],$$

where P_0 is the Fourier power obtained in the integer-bin case. While the sinc function concentrates most of the power in the bins closest to the true frequency, a fraction of the signal power leaks into neighbouring bins, leading to a reduction of the peak amplitude. In the worst case, when the signal frequency lies exactly halfway between two adjacent DFT bins, nearly 60% of the signal power is lost from the peak bin. Averaged over all possible frequency offsets within a bin, the expected power loss is approximately 23%. For a more detailed explanation, see [25].

This loss of peak power reduces the sensitivity of FFT-based spectral analyses and must be taken into account when estimating signal amplitudes or detection thresholds.

In practice, this effect can lead to unintuitive reductions in recovered signal power.

In this analysis, it was necessary to understand the power spectrum that results from a signal in order to understand the significance. For that reason, the FFT algorithm was performed on a time series where, beforehand, a signal of known strength and phase was injected. A varied phase led to an apparent loss of power and, hence, significance. This can be seen in Figure 2.2. Both upper and lower graphs are a power spectrum for an injected signal with identical amplitude, but there is a phase difference of $\frac{\pi}{2}$. At the injected frequency, the upper spectrum shows a power of 20 a.u., while the lower shows 10 a.u. This difference comes solely from the phase shift. But knowing this effect, it can be attributed directly to discrete sampling and a finite observation window, rather than to an artefact of the analysis itself.




Figure 2.2: Results of applying the later-described analysis to an injected signal with the same amplitude and background, but with a phase shift of 0 in the upper graph and $\frac{\pi}{2}$ in the lower graph. This makes the scalloping effect more apparent.

2.1.2 Definition and Properties of the Test Statistic used

2.1.2.1 Fourier coefficients and choice of observable

The complex coefficients \hat{z}_k resulting from the FFT contain both magnitude and phase. In many time-series analyses, including those used in neutrino astronomy, the absolute phase of the Fourier transform is not physically meaningful because the choice of the observation start time is arbitrary. What matters are the frequencies and relative phases of periodic components.

For the case where the absolute phase does not hold information, there are two commonly used approaches: either using $|\hat{z}_k|$, referred to as the magnitude, or using $|\hat{z}_k|^2$, referred to as the Fourier power. Both approaches contain equivalent information regarding the strength of the periodic components. This analysis, however, makes use of the Fourier power, as in this case the background distribution takes the form of the well-studied χ^2 distribution as can be seen in Figure 2.3.

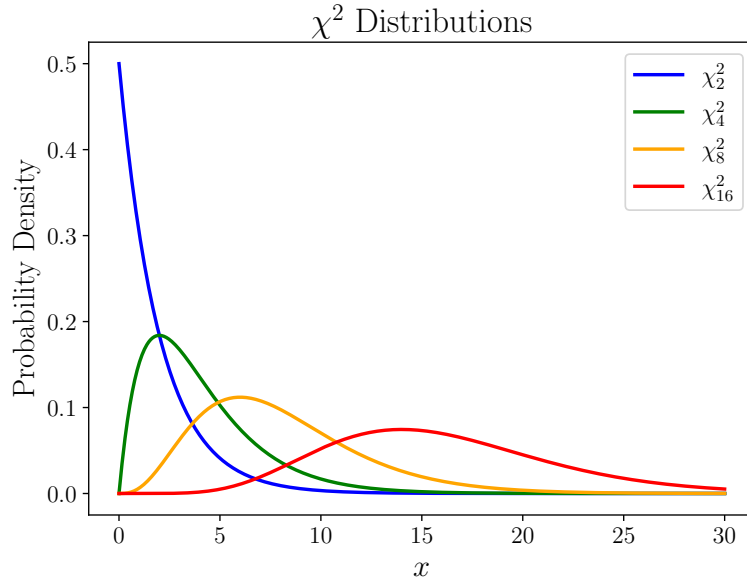


Figure 2.3: χ^2 -distributions for 2 (blue), 4 (green), 8 (orange), and 16 (red) degrees of freedom. The x-axis shows the value of the χ^2 test statistic, while the y-axis represents the probability density.

2.1.2.2 Define hypotheses

The null hypothesis (H_0) is a claim about reality that the proposed effect does not influence the underlying distribution. The alternative hypothesis H_1 represents the presence of the effect under investigation.

In the context of this work,

$$H_0 : \text{background only (no periodic signal),}$$

$$H_1 : \text{presence of a periodic component.}$$

2.1.2.3 Define a test statistic

A test statistic is a real-valued function of observable parameters

$$T = T(\text{data})$$

giving a measure of the incompatibility with H_0 . The choice of T depends on the problem and on the expected structure of the deviations from H_0

In this analysis, the test statistic is the Fourier power at a given frequency,

$$T = |\hat{z}_k|^2$$

This is sensitive to periodic structure at frequency f_k . In [5], Figure 2.2 shows an example of this. If no periodic signal is present, its fluctuations arise purely from statistical background noise.

2.1.2.4 Null distribution

One has to know the reference distribution or null distribution that the test statistic would have given that H_0 is indeed true. That could e.g. be a normal distribution, a polynomial distribution or a χ^2 -distribution.

In this case, under the background-only assumption, the real and imaginary parts of \hat{z}_k are Gaussian distributed. Therefore, the Fourier power follows a χ^2 -distribution with 2 degrees of freedom.

This distribution serves as the reference (null) distribution against which observations are compared.

2.1.2.5 Observed test statistic

The next step is to calculate the test statistic for the data.

$$t_{\text{obs}} = T(\text{observed data}).$$

In our case:

$$t_{\text{obs}} = |\hat{z}_k|^2$$

and large values for t_{obs} indicate a possible periodic contribution

2.1.2.6 Definition of the p-value

The p-value is then the probability of observing data (or a test statistic) at least as extreme as what is observed, assuming H_0 is true.

$$p = P(T \leq t_{\text{obs}} | H_0 \text{ is true})$$

Importantly, the p-value is not the probability that H_0 is true. It is a probability computed under the assumption that H_0 holds.

2.1.2.7 Decision rule

At last a comparison of p to the prior defined significance level α :

If $p \leq \alpha$: H_0 will be rejected, i.e. there is statistically significant evidence for H_0 to be not the correct description of reality.

If $p \geq \alpha$: H_0 cannot be rejected, as there is no statistically significant evidence.

That can lead to H_0 being rejected even though it is true (error of type I). That happens with probability α . Or H_0 not being rejected even though it ought to be (error of type II).

This follows the same approach as is described in [5].

2.2 Ingredients for Analysis and the pipeline

For the analysis, the following inputs are needed:

- The detector rate with corresponding UTC time stamps, i.e. time values expressed in Coordinated Universal Time, the international atomic-based time standard used for precise event timing.
- The geographic coordinates of the detector.
- The target object, specified by its sky coordinates and (if available) an estimate of its spin frequency. In this study, the target objects are pulsars. The details on the selected sources are provided in section 2.2.3.

The flow chart in Figure 2.4 shows the process of getting from the input data to the significance levels: first, the detector rate is extracted and combined with the detector rate from other runs. Secondly, a barycentric correction is applied. Next, the corrected time series is put through an FFT, and the result is filtered to suppress red noise. Optionally, a harmonic summation is performed to enhance sensitivity to signals with power distributed over multiple harmonics. At last, the statistical significance of the periodic component at the spin frequency is calculated. A detailed description of the implementation is provided in Appendix A. The code that was used can be found in the git repository [26].

2.2.1 Detector rate

2.2.1.1 Justification of total detector rate

A neutrino ν with an energy of $E = 1 \text{ GeV}$ propagates with a velocity extremely close to the speed of light. For a relativistic neutrino with mass $m_\nu \ll E/c^2$, the velocity can be approximated by

$$v \simeq c \left(1 - \frac{1}{2} \frac{m_\nu^2 c^4}{E^2} \right).$$

Assuming a neutrino mass of $m_\nu \approx 1 \text{ eV}/c^2$, this yields

$$v \gtrsim c (1 - 5 \times 10^{-19}) \approx c,$$

demonstrating that any deviation from c is entirely negligible for experimental considerations.

Cosmological constraints from the Planck satellite [27] place an upper bound on the sum of the neutrino masses of

$$\sum m_\nu < 0.12 \text{ eV}/c^2,$$

while direct kinematic measurements from the KATRIN [28] experiment give a more conservative limit of

$$m_\nu < 0.8 \text{ eV}/c^2.$$

In the KM3NeT detector, the integration time of a single readout frame is approximately 100 ms. During this time interval, a neutrino travelling at essentially the speed of light would cover a distance of order

$$d \sim c \times 100 \text{ ms} \approx 3 \times 10^4 \text{ km},$$

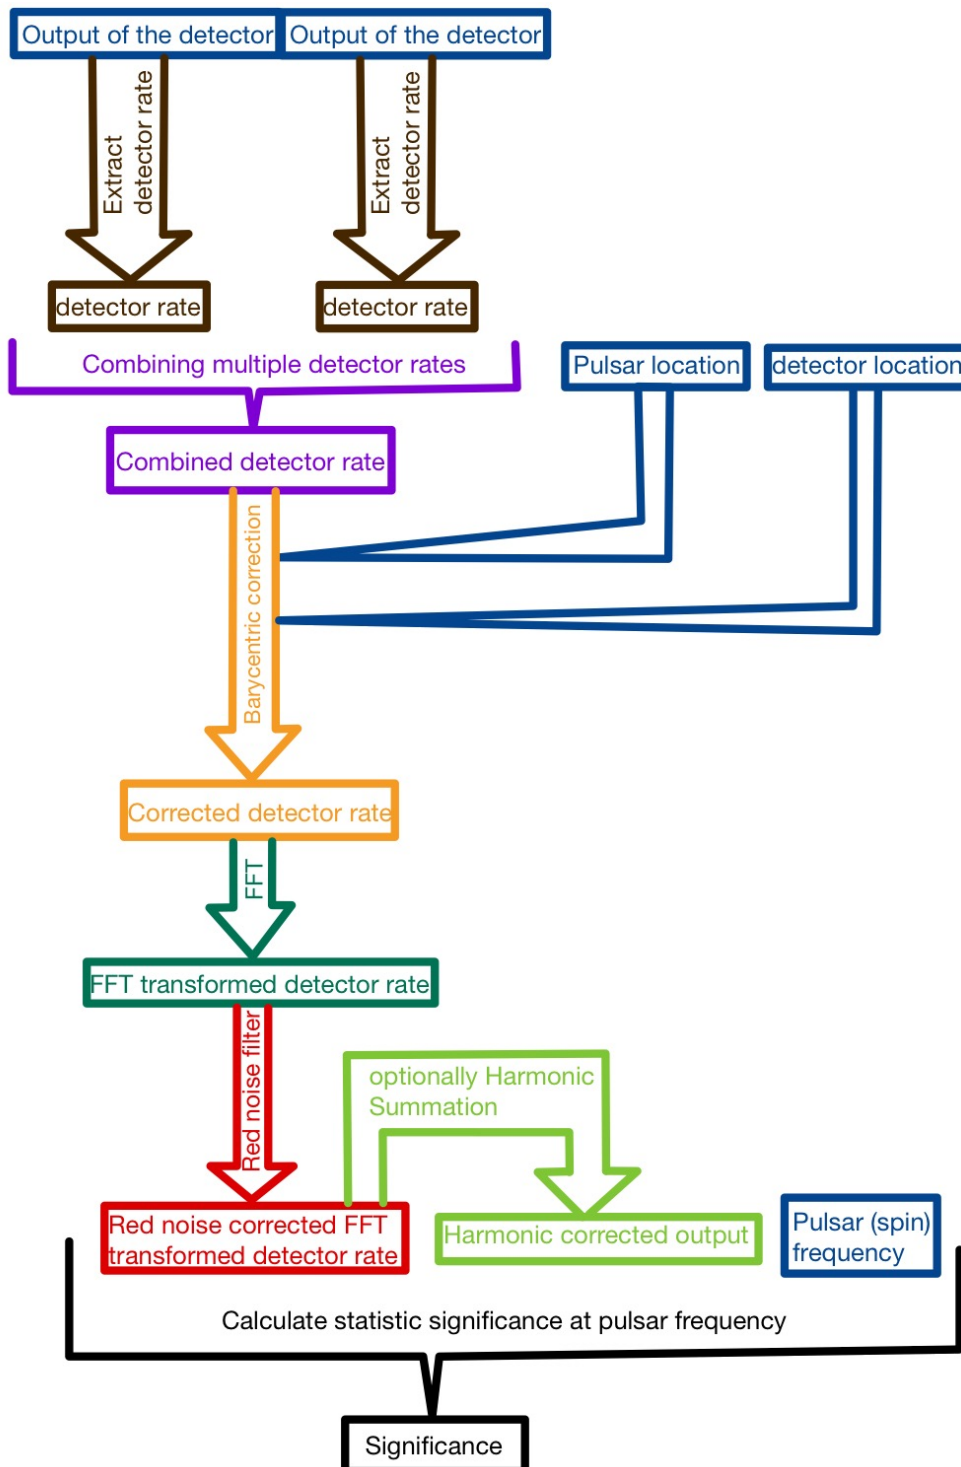


Figure 2.4: Flow chart illustrating the analysis pipeline used to determine the statistical significance of the pulsar signal. Dark blue boxes indicate input data. The data are first converted into the detector rates (brown), after which rates from multiple detectors are combined (violet). A barycentric correction (orange) is applied prior to performing the FFT (dark green). The resulting power spectrum is filtered for red noise (red), and an optional harmonic correction may be applied. Finally, the statistical significance at the pulsar frequency is calculated (black).

which is many orders of magnitude larger than the physical dimensions of the detector. So, if there are periodic neutrino signals with a high flux of low energy, they would virtually be in the entire detector at once. And due to the low energy, the Cherenkov light of the interaction product would mostly interact with the DOM nearby. The idea of which was discussed in section 1.3.2.

Consequently, no additional spatial or temporal information can be extracted by analysing individual photomultiplier tubes (PMTs) or digital optical modules (DOMs) on these timescales.

It is therefore justified to treat the measurement as a global detector rate, rather than attempting a DOM- or PMT-resolved analysis.

2.2.1.2 Extracting total detector rate and combining different runs

This total detector rate is extracted and normalised by dividing by the number of active PMTs (Table 2.2.1.2) divided by 31, corresponding to the effective number of digital optical modules (DOMs) that could be fully populated by the active PMTs. That is done to later easily use the idea of the effective volume for the DOM.

A PMT is considered active if it satisfies all of the following criteria:

- the PMT rate is strictly positive;
- the PMT is not flagged by the high-rate veto (HRV);
- the time-to-digital converter (TDC) FIFO buffer is not almost full. That means:
 - Normally, the PMT signal is stamped by a Time-to-Digital Converter (TDC)
 - the TDC stores these time-stamps temporally in a FIFO (First-In–First-Out) buffer and is read out by the data acquisition (DAQ) system
 - if e.g. due to bioluminescence, this buffer fills faster than it is emptied, these values are not counted, as that might mean that new hits are delayed or lost.

After that, the time series are stored together with the location of the detector, which is taking the data and combined into one file. Without this step, each file would only contain three hours of data-taking time, i.e., one run. Combining the time series is essential because an FFT requires loading the entire dataset into the computer’s RAM, and the location will be necessary in the next section (section 2.2.2.1) for the correction step.

2.2.2 Correct time series

An important step is the barycentric correction. In the context of this study, this step is mandatory, as the analysis targets the pulsars listed in section 2.2.3. The selection of these pulsars, as well as the motivation for their choice, is discussed in more detail in that section.

2.2.2.1 Barycentric correction

The relative motion of the Earth around the Sun (orbital motion) and the rotation of the Earth (diurnal effect) cause a shift in the neutrino arrival times and, hence, in the detected photon hits.

This must be corrected for. To achieve this, a coordinate transformation from the rotating frame of the Earth to a non-rotating (inertial) frame is applied. The frame typically used for this purpose is the barycentre of the Solar System. Although, to be precise, the Solar System Barycentre defines a quasi-inertial reference frame, it is sufficiently inertial for high-precision timing analyses. This limitation arises from the presence of external gravitational fields, such as the Galactic potential and due to nearby stars.

The barycentre refers to the centre of mass of multiple gravitationally interacting objects; in other words, it is the point about which all bodies, including in this case the Earth, orbit.

We therefore perform this transformation to obtain the arrival times of the hits in the detector as they would be at the barycentre of the Solar System (SSB). This takes the projected travel time into account. An overview of the barycentric correction is described in [29] and [5].

On a formal level, the time offset Δt to the SSB is first calculated as follows. Neutrinos are again assumed to propagate at the speed of light, such that barycentric corrections are identical to those applied in photon timing:

$$\Delta t = \frac{\vec{r}_{\text{obs}}(t) \cdot \hat{n}}{c}$$

where

- $\vec{r}_{\text{obs}}(t)$ is the detector position relative to the Solar System Barycentre,
- \hat{n} is the unit vector towards the pulsar,
- c is the speed of light.

The barycentric arrival time is then given by

$$t_{\text{SSB}} = t_{\text{obs}} + \Delta t.$$

A positive Δt corresponds to the detector being closer to the source than the Solar System Barycentre along the line of sight.

Additional timing corrections

If the pulsar is part of a binary system, the orbital motion introduces an additional Doppler shift that must be corrected for (see e.g. [29]). However, since the selected sources do not include binary systems, this effect is not considered here.

Further relativistic timing effects include the Einstein delay, caused by time dilation and gravitational redshift within the Solar System [30, 29], and the Shapiro delay, which arises from signal propagation through curved spacetime. In the case of a signal passing near the solar limb, the latter reaches a maximum of order $120 \mu\text{s}$ [30].

For the present study, with an observation time of $T = 9 \text{ h}$ and the frequency resolution considered (compare with Equation 2.2), these relativistic contributions are subdominant and therefore neglected. Consequently, only the geometric (Römer) delay is applied.

2.2.3 Selected Pulsars

As described earlier in section 1.4, we assume that any potential neutrino emission from pulsars is phase-aligned with the observed photon emission. In addition, we assume that the neutrino signal shares the same duty cycle as the photon signal.

The duty cycle DC is defined as the fraction of the rotational period T_{rot} during which the pulsar emits radiation,

$$DC = \frac{t_{\text{emission}}}{T_{\text{rot}}}.$$

If the pulse width d_{puls} is given as a phase width in radians over a full rotation 2π , the duty cycle can be approximated as

$$DC \approx \frac{d_{\text{puls}}}{2\pi}.$$

It provides a measure of how sharply the emission peaks.

The following criteria are applied for the pulsar selection:

- Nyquist frequency: The pulsar spin frequency must lie below the Nyquist frequency of KM3NeT, as calculated in Equation 2.3, it is $f_{\text{Nyq}} = 5$ Hz.
- Distance: The pulsar should be as close as possible to Earth.
- No binary system: Only isolated pulsars are considered. As discussed in section 2.2.2, binary systems would require additional timing corrections, which are beyond the scope of this analysis.
- Suitable combination of frequency and duty cycle: For later Harmonic Summation, we need a narrow pulse, i.e. low Duty Cycle and also a frequency for which also multiples of that frequency lie below the Nyquist frequency of KM3NeT

Similar constraints were applied in [5]; therefore, the source selection was adopted from that work and is listed in Table 2.1.

PSRJ	RAJD (deg)	DECJD (deg)	F0 (Hz)	DIST (kpc)	Duty Cycle
J1704–6016	256.061 666 67	–60.281 666 67	3.264 528	1.589	0.293 807 5
J0820–4114	125.064 416 67	–41.243 111 11	1.833 364	0.571	0.258 504 3
J0750–6846	117.649 375 00	–68.776 058 33	1.092 638	0.338	0.205 962 4
J2325+6316	351.305 498 33	63.281 211 67	0.696 229	4.855	0.091 345 2
J0828–3417	127.069 272 93	–34.285 291 02	0.540 857	0.354	0.061 116 8

Table 2.1: Selected pulsars and their parameters, taken from [5].

Chapter 3

Analysis

3.1 FFT analysis

3.1.1 Red noise

Initial investigations show that, similarly to the results reported in [5] for ANTARES, the power spectra obtained for both KM3NeT-ORCA and KM3NeT-ARCA have a red-noise behaviour, characterised by an excess of power at low frequencies, with the spectral power decreasing as a function of increasing frequency. In Figure 3.1 panels 1 and 2, one can see this red-noise spectrum.

Such red noise may originate from slow temporal variations in detector conditions, such as environmental effects, or long-timescale modulations in the background event rate. Its presence significantly reduces the sensitivity to weak periodic signals, particularly at low frequencies.

To mitigate this effect, a red-noise filter is applied to the power spectrum.

3.1.2 Harmonic summation

Harmonic summation is optionally applied to increase the sensitivity to periodic signals that do not have a sinusoidal pulse shape. Harmonic summation is reviewed in [5] and discussed in greater detail in [31].

If a pulsar emits a narrow non-sinusoidal pulse at a fundamental frequency f , it does not concentrate all its Fourier power in a single bin. But instead, the power is spread across multiple harmonics at frequencies

$$f, 2f, 3f, 4f, \dots, Nf.$$

Looking only at the fundamental frequency underestimates the total signal power.

To recover the total signal strength, the power at the first H harmonics is summed coherently at the fundamental frequency. If a true periodic signal is present, the harmonics align constructively and enhance the detection statistic. Conversely, if only noise is present, it does not get stronger.

As described in [5], harmonic summation modifies the distribution of the test statistic under the background-only hypothesis. Specifically, the number of degrees of freedom increases from ξ to $\xi \cdot H$. In this case $\xi = 2$, as can be seen in section 2.1.2.4 in the null distribution part.

Finally, since H harmonics are combined into each fundamental frequency bin, the resulting spectrum is reduced in range to

$$(0, \lfloor f_{\text{Nyq}}/H \rfloor].$$

3.1.3 Statistical significance

The statistical significance of peaks in the resulting power spectrum is evaluated using χ^2 statistics. Under the assumption of white Gaussian noise, which the red-noise filter gives a reasonable approximation for, the FFT power in each frequency bin follows a χ^2 distribution with two degrees of freedom. For the case of harmonic summation H , one has $2H$ degrees of freedom. Significance thresholds are derived from that, and frequency bins exceeding the chosen confidence level, in this case 5σ , are identified as candidate periodic signals.

3.2 Verifying the pipeline

This analysis was tested on a 9 h of ORCA 18 running time. This data set is too short for one to reasonably expect to detect a signal.

The purpose was for this to work as a sanity check and to investigate the noise further, and it indeed led to finding some problems that could then be solved.

That data also led to the discovery of the red noise spectrum in the raw data, as can be seen in Figure 3.1. That is a property that one might not have expected purely from simulations.




Figure 3.1: FFT Analysis for the Raw data without barycentric correction or injected signal. Importantly, one sees how there is a red noise background, but no detection of any frequency would cross the detection threshold.

As expected, there was no significant power excess for the pulsars, and p-values are in the expected range. The result of the FFT can be seen in Table 3.1.

Table 3.1: p-value for the selected pulsars at their spin frequency

Pulsar	H	Frequency [Hz]	p-value
J0750-6846	1	1.09264	0.7731
	2	1.09264	0.7138
	4	1.09264	0.8768
J0820-4114	1	1.83336	0.1975
	2	1.83336	0.1108
J0828-3417	1	0.54085	0.7480
	2	0.54085	0.7562
	4	0.54085	0.9244
	8	0.54085	0.9912
J1704-6016	1	3.26452	0.2341
J2325+6316	1	0.69622	0.6842
	2	0.69622	0.9155
	4	0.69622	0.8797

The sensitivity of the analysis scales with \sqrt{T} where T is the total observation time. A longer time series would therefore improve the possibility of a detection. Such an extension would require unblinding the remaining dataset, which was not performed at this stage in order to preserve statistical integrity.

Additionally, the applied red-noise filtering slightly oversubtracts the local noise level, introducing a small downward bias in the power estimator. This effect may further reduce sensitivity to weak signals.

It is noteworthy that the pulsars in the table with higher frequency have a lower p-value, i.e., a higher significance, which could be due to the red noise, as stated in section 3.1.1. This red noise reduces the sensitivity to weak signals, especially at low frequencies. And the other expected effect is that with a higher frequency, there are more data points in the time series.

3.3 Sensitivity Study

3.3.1 Signal Injection and Significance Estimation

An obvious question is how large a neutrino flux would be required from a pulsar in order for it to be observable as a function of frequency, energy and background.

The detailed implementation can again be found in Appendix B.

The general idea is to perform the FFT multiple times on an injected signal with known parameters. Between each iteration, the rate data are randomly scrambled in the time domain, thereby pseudo-simulating different realisations of the background.

The injection of the signal of known frequency and amplitude is as follows:

$$\text{rate}_{\text{new}}(t) = \text{rate}_{\text{old}}(t) + \text{SNR} \cdot \sigma(\text{rate}_{\text{old}}) \cdot \sin(2\pi f t_{\text{bin}}),$$

where f is the injected frequency and t_{bin} denotes the time associated with each bin.

As one can see, the definition of the amplitude is chosen to be the product of the noise $\sigma(\text{rate}_{\text{old}})$ and the SNR signal-to-noise ratio.

For each pseudo-experiment, the FFT is computed, and the resulting spectrum is checked for a peak at the input frequency that exceeds the 5σ detection threshold. A counter records how many of these trials surpass this threshold; this is called the detection probability.

In addition, a median significance is defined as the median of all significance values obtained from the ensemble of pseudo-simulations. This quantity provides a good measure of the typical sensitivity at the chosen SNR.




Figure 3.2: Significance scan for pulsar J1704-6016; The upper panel shows the mean significance σ_{mean} while the lower panel shows the detection probability. Each is plotted against the amplitude of the injected signal in terms of SNR. For this plot, there were 160 pseudo-experiments.

As one can see in Figure 3.2, the σ_{mean} shows the expected behaviour of a linear function. Since the signal amplitude scales with the signal-to-noise ratio,

$$A \propto \text{SNR},$$

and the raw power is proportional to the squared amplitude,

$$P \propto A^2,$$

the corresponding significance, defined as

$$\sigma \sim \sqrt{P},$$

becomes

$$\sigma \propto A \propto \text{SNR}.$$

The detection probability also behaves as expected, as long as $\sigma_{\text{mean}} \ll 5\sigma$ mark, the number of pseudo-experiments that supersede the threshold of 5σ is low and near 0. As soon as $\sigma_{\text{mean}} \approx 5\sigma$, the detection transitions from 0 to 1. And if the $\sigma_{\text{mean}} \gg 5\sigma$, the detection probability reaches 1 and stays constant.

3.3.2 Cross section and effective volume

3.3.2.1 Cross section

As the cross-section, it was decided to take the standard cross-section (G18_02a_00_000.xml) of Genie [32]. Especially, the package of the cross sections can be found at [33]. It was also decided to take the cross-sections that are produced on the interaction with oxygen ^{16}O , as a good proxy for water.




Figure 3.3: Comparative charged-current (CC) neutrino and antineutrino cross sections on ^{16}O for energies below 20 GeV, as obtained from the GENIE cross-section package. The colour coding indicates ν_e (blue), $\bar{\nu}_e$ (orange), ν_μ (green), $\bar{\nu}_\mu$ (red), ν_τ (purple), and $\bar{\nu}_\tau$ (brown).

Figure C.21 shows the cross-sections that were extracted from Genie individually, while Figure 3.3 shows them in the same diagram.

There are some humps for the ν_τ and $\bar{\nu}_\tau$, and also it starts near 2.5 GeV, even though a simple minimum energy threshold for the scattering on a nucleon at rest gives:

$$\nu_\tau + n \rightarrow \tau^- + p^+ \Rightarrow E_\nu^{\text{thr}} = \frac{(m_\tau + m_p)^2 - m_n^2}{2m_n}$$

gives $E_\nu^{\text{thr}} \approx 3.45$ GeV. But the reason is that it is scattering on a bound nuclear target, not free nucleons. So, there is Fermi motion and binding energy corrections. But

also importantly, a high momentum tail due to short-range corrections (SRC) between nucleons. So that it is brought down to around 2.5 GeV.

Further, one can see in Figure 3.3 the influence of the V-A structure of the weak force as the cross-section for the antineutrinos is lower than the cross-section for the neutrinos.

3.3.2.2 Definition of the Effective Volume

The idea of what will be called effective volume V_{eff} in this thesis is as follows: when one looks at all monoenergetic neutrinos that interact within a cylindrical volume V_{original} around the DOM, say the number is N_{interact} . And then we count the number of which produce at least one hit at one of the PMTs of the DOM, say $N_{\text{hit}>0}$, and divide those to get the effective volume V_{eff} .

$$V_{\text{eff}}(E_\nu) = \frac{N_{\text{hit}>0}}{N_{\text{interact}}}\bigg|_{E_\nu=\text{const}} \cdot V_{\text{generation}} \quad (3.1)$$

Since that definition only looks at the neutrinos that have already interacted in the volume, no cross-section is folded into the effective volume. However, because of the idea discussed before with the lower energy neutrino interaction producing less light, this quantity is indeed energy dependent.

To determine the effective volume of a single DOM, a large sample of monoenergetic neutrinos is simulated. The neutrinos are forced to interact within the original generation volume, and the resulting secondary particles are propagated through water to simulate light production and detection.

The generation is done via gSeaGen [32], and the light propagation via jsirene [34].

The effective volume can be seen in Figure 3.4.

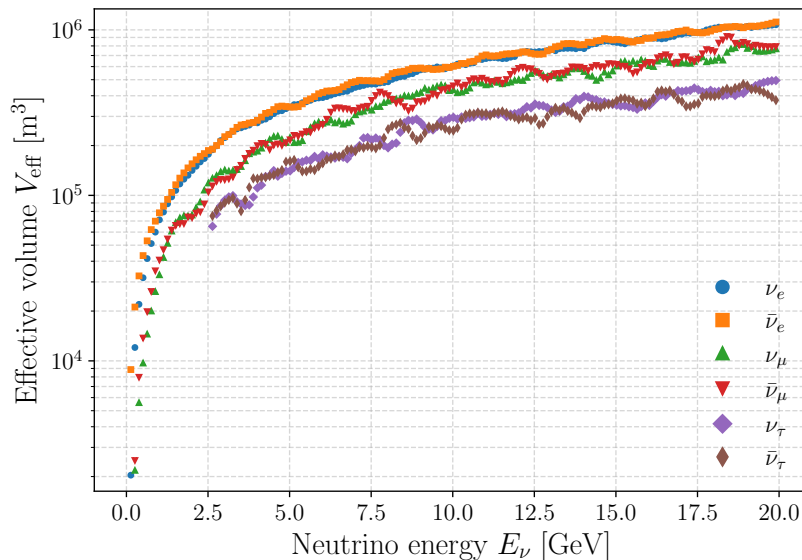


Figure 3.4: Effective volume as a function of neutrino energy for all neutrino flavours and their corresponding antiparticles. The colour coding indicates ν_e (blue), $\bar{\nu}_e$ (orange), ν_μ (green), $\bar{\nu}_\mu$ (red), ν_τ (purple), and $\bar{\nu}_\tau$ (brown).

As one would expect, the effective volume for electron (anti-)neutrinos is larger than that for muon (anti-)neutrinos, while tau (anti-)neutrinos have the smallest effective

volume. This behaviour is consistent with the differing masses of the charged leptons produced in the interaction, which affect the light yield and propagation characteristics. Furthermore, the effective volumes for neutrinos and antineutrinos are nearly identical. This is partially expected, since the definition of the effective volume does not include the interaction cross-section; it depends only on the fraction of interacting events that produce a detectable signal within the considered volume.

However, one would expect that due to the Bjorken y difference, i.e. a difference in how much energy is put into the lepton/hadronic shower, between neutrinos and antineutrinos, that the effective volume for antineutrinos is higher. This is not observed. There are presumably two reasons: especially with the fluctuations visible in the graph, there might not be enough statistics to see that effect. And there could be quenching of the hadronic shower, since protons often do not have enough energy to produce Cherenkov light.

3.3.3 Neutrino propagation and flavour ratio at Earth

Neutrino oscillation, especially over astrophysical distances, modifies the flavour composition of an initially produced neutrino beam. An explanation of that modification through propagation can be found in [35]. Due to the long baseline, the oscillation probabilities average out, resulting in a flavour ratio at Earth that is independent of energy.

In this study, the CP-violating phase is set to $\delta = 0$

The following two cases of initial flavour composition are considered

3.3.3.1 Case A:

For an initial flavour ratio of

$$\{\nu_e, \nu_\mu, \nu_\tau\} = \{1, 2, 0\}$$

the oscillation-averaged composition at Earth becomes then

$$\{1, 1, 1\} \quad \text{i.e. normalised: } \left\{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right\}.$$

3.3.3.2 Case B:

For an initial flavour ratio of

$$\{\nu_e, \nu_\mu, \nu_\tau\} = \{1, 0, 0\}$$

the resulting composition at Earth is approximately

$$\{0.52, 0.24, 0.20\}.$$

Those numbers are taken from [35]. In the case B (3.3.3.2), the corresponding graphic can be found in Figure 3.5.

The new flavour compositions will be used as the weights between the different flavours for the Sensitivity calculation.




Figure 3.5: For initial mono-flavour ν_e , one gets these flavours after averaging over astronomical propagation. The y-axis is the relative amount, the x-axis is the CP-violating phase δ . This study, however, looks at the case of $\delta = 0$, taken from [35].

3.3.4 Necessary Rate

Now, to convert the output of the $\sigma_{\text{mean}}(\text{SNR})$ from section 3.3.1 into a physical rate, we need to do some operations.

The differential neutrino-induced event rate n_ν per unit energy and time is modelled as

$$\frac{dn_\nu}{dt dE} = F_\nu(E) \sum_\alpha \omega_\alpha \left(\frac{d\sigma(E)}{dE} \right)_\alpha \frac{N_A \rho}{M_{\text{water}}} V_{\text{eff},\alpha}(E) \varepsilon. \quad (3.2)$$

In this expression, the variables are defined as follows:

- $F_\nu(E)$ is the neutrino flux at the detector,
- ω_α is the normalised proportion of flavour α in the beam,
- $\frac{d\sigma}{dE}$ is the differential neutrino interaction cross section,
- N_A is Avogadro's constant,
- ρ is the density of seawater at 13.3 to 13.5°C, corresponding to the depth of ORCA and ARCA [36],
- M_{water} is the molar mass of water,
- $V_{\text{eff}}(E)$ is the detector effective volume, and
- ε denotes the overall detection efficiency.

As by the definition here used for the effective volume V_{eff} (Equation 3.1) does not include cross-section effects.

3.3.4.1 Detection Threshold

From the $\sigma_{\text{mean}}(\text{SNR})$, a signal-to-noise ratio corresponding to a median significance of 5σ is obtained.

This threshold is converted into a minimum detectable modulation of the event rate, denoted by

$$\left(\frac{dn_\nu}{dt}\right)_{5\sigma}. \quad (3.3)$$

This quantity represents the smallest periodic signal amplitude that can be expected to be detectable by the analysis pipeline with 5σ .

3.3.4.2 Flux Sensitivity

The neutrino flux sensitivity is obtained by inverting Equation 3.2 for a fixed detection threshold, yielding

$$F_\nu^{(5\sigma)}(E) = \frac{\left(\frac{dn_\nu}{dt}\right)_{5\sigma}}{\sum_\alpha \omega_\alpha \left(\frac{d\sigma(E)}{dE}\right)_\alpha \frac{N_A \rho}{M_{\text{water}}} V_{\text{eff},\alpha}(E) \varepsilon}. \quad (3.4)$$

Equation (3.4) defines the energy-dependent neutrino flux required to produce a detectable periodic modulation at the 5σ level.

3.3.4.3 Result

The resulting flux and luminosity sensitivity are visualised as a contour in the energy–flux and energy–luminosity plane. Fluxes and luminosities above the contour correspond to periodic neutrino signals that would be detected with a global significance of at least 5σ , while fluxes and luminosities below the contour remain undetectable.

For the pulsar J1704-6016, the flux and luminosity sensitivities can be found in Figure 3.10. The other flux sensitivities for case A can be found in Figure C.5, those for case B in Figure C.10. While the other luminosity sensitivities for case A can be seen in Figure C.15 and for case B in Figure C.20.

For clarity, only the 5σ discovery sensitivity is shown, as lower-significance thresholds exhibit the same energy dependence and differ only by a constant normalisation.

As can be seen, cases A and B do not differ significantly, which makes sense, as it only causes a linear prefactor in front of the cross sections. And this linear factor is subdominant for the logarithmic y-axis.

One can see a sharp peak near zero energy, which is also expected, as for the extreme case of exactly 0 energy, one would need an infinite flux

Since the fluxes are high and the luminosities correspondingly high, one would need to increase the sensitivity. This can be done with a longer time window. The sensitivity of the FFT scales with \sqrt{T} with T being the observation time. But because the Fourier power is used, this gets squared.

3.3.4.4 Comparison with ANTARES

In [5], the same pulsars were looked at and produced flux sensitivity estimations for the ANTARES.




Figure 3.6: Flux sensitivity for pulsar J1704-6016 and case A.




Figure 3.7: Luminosity sensitivity for pulsar J1704-6016 and case A.




Figure 3.8: Flux sensitivity for pulsar J1704-6016 and case B.




Figure 3.9: Luminosity sensitivity for pulsar J1704-6016 and case B.

Figure 3.10: For the pulsar J1704-6016 gives for case A (3.3.3.1) and case B (3.3.3.2) the flux and luminosity sensitivity.

There, a detection rate of around $F_{5\sigma}^{\text{ANT}} \approx 10^4 \text{ s}^{-1}$ is necessary at around 2.5 GeV for a 5σ detection.

From the data discussed here, one would need a detection rate of around $\left(\frac{dF}{dE}\right)_{5\sigma}^{\text{ORCA}} \approx 10^9 \text{ s}^{-1} \text{ GeV}^{-1}$ at around 2.5 GeV for the 5σ detection.

The values in [5] are already integrated over energy, so one can approximately do:

$$F_{5\sigma}^{\text{ORCA}}(\Delta E) = \left(\frac{dF}{dE}\right)_{5\sigma}^{\text{ORCA}} \cdot \Delta E$$

However, the sensitivity of the FFT scales with the $\frac{1}{\sqrt{\text{length of the time series}}}$ and goes as a square (due to using the Fourier power) into the 5σ threshold.

In this analysis, there were $N_{\text{ORCA}} \approx 2^{18.3}$ and for the ANTARES analysis, there were $N_{\text{ANT}} \approx 2^{26.1}$, hence

$$F_{5\sigma}^{\text{ORCA}}(N_{\text{ANT}}) = F_{5\sigma}^{\text{ORCA}}(N_{\text{ORCA}}) \cdot \frac{N_{\text{ORCA}}}{N_{\text{ANT}}} \approx 10^6 \text{ s}^{-1}$$

That is in the same order. The remaining differences can be explained by the different assumptions. While this thesis uses a neutrino beam that is not entirely one flavour, the

ANTARES estimate uses a mono-flavour electron neutrino beam, which has amongst the highest cross sections and effective volume as can be seen in Figure 3.3 and Figure 3.4.

Chapter 4

Future prospects

The pipeline is now fully operational and can be applied to a longer time window to enhance the sensitivity.

Understanding the origin of the red noise better might also be useful, as there is this potentially oversubtracting and, hence, reducing the sensitivity of this analysis.

The different flavour composition of the neutrino beam did not significantly change the flux that needs to be detected on Earth, as a factor that more or less goes linearly into the flux. However, it is still an interesting quantity from the theory side and would give bounds to the possible models.

For longer time windows, it might be advisable to also implement corrections for the Einstein delay or the Shapiro delay as is described in section 2.2.2.1. Additionally, there could be interesting candidates in binary systems, so it might be worth looking at an implementation to correct the introduced Doppler shift as well.

Also, the method used here, the FFT, needs to load the entire data set into its RAM. It might be useful to explore other methods. Others have, such as [29], for example, already begun to explore epoch folding.

Also, one could create a framework to not only use KM3NeT data but to combine it with other neutrino experiments, such as IceCube, ANTARES, or P-ONE. That could have not only the benefit of giving more data to perform the analysis on, but also average out the noise caused by the direct environment. In principle, that should work just fine if one were to combine the data after performing the barycentric correction (and others if so decided) and resampling on a common time grid. The only difficulty could be if the time slices were too different, but that seems solvable.

Appendix A

Appendix A: Implementation Details of the Analysis Pipeline

The code that was used can be found in the git repository [26]. This section is intended for someone who may want to understand the pipeline; as such, there is some duplication with the main text.

A.1 Extract total detector rate

For this purpose, the script `ExtractTotalDetectorRate.jl` was written. The script extracts the detector-wide rate from `.root` files and is written in Julia [37], making use of the `KM3io.jl` library [38]. The choice of Julia and this library was motivated by Julia's efficient handling of for loops and the intuitive, easy-to-use functions provided by the library, which allow for rapid and reliable implementation.

Within the script, the individual rates of all active photomultiplier tubes (PMTs) are summed up. This total rate is then normalised by dividing by the number of active PMTs by 31, corresponding to the effective number of digital optical modules (DOMs) that could be fully populated by the active PMTs. The resulting detector rate is then correlated with the corresponding UTC time stamp.

A PMT is considered active if it satisfies all of the following criteria:

- the PMT rate is strictly positive;
- the PMT is not flagged by the high-rate veto (HRV);
- the time-to-digital converter (TDC) FIFO buffer is not almost full. That means:
 - Normally, the PMT signal is stamped by a Time-to-Digital Converter (TDC)
 - the TDC stores these time-stamps temporally in a FIFO (First-In-First-Out) buffer and is read out by the data acquisition (DAQ) system
 - if e.g. due to bioluminescence, this buffer fills faster than it is emptied, these values are not counted, as that might mean that new hits are delayed or lost.

The final output of the script is an HDF5 file containing the time series of the detector rate together with the associated UTC timestamps.

A.2 Combining time series

After running `ExtractTotalDetectorRate.jl` multiple times, one ends up with several HDF5 files containing the time series and the corresponding timestamps. In order to perform analyses over longer periods, these files need to be combined into a single, larger HDF5 file. Without this step, each file would only contain three hours of data-taking time, i.e., one run.

Combining the time series is essential because an FFT requires loading the entire dataset into the computer's RAM.

To address this, the script `KM3NeTCombineTimeSeries.py` was written. At this stage, the pipeline switched to Python [39] as the programming language. This choice was motivated in part by inheriting an initial structure from [5], and also because `astropy` [40] provides convenient tools that proved useful in subsequent analysis steps.

The Python code reads the HDF5 files, stacks the time series, and returns an `astropy`. The `Timeseries` object stores the combined data in a new HDF5 file. Additionally, it assigns the location of the detector site that produced each time series:

The ORCA detector is located in the Mediterranean Sea at approximately $42^{\circ}48\text{N}$ latitude and $6^{\circ}02\text{E}$ longitude, at a depth of about 2450 m below the sea surface, with an additional 30 m offset and the OMF spacing correction applied. The site corresponds to the French KM3NeT location near Toulon [11, 12]. And the ARCA detector is also in the Mediterranean Sea at approximately $36^{\circ}16\text{N}$ latitude and $16^{\circ}06\text{E}$ longitude, at a depth of about 3500 m below the sea surface, including a 72 m offset and the OMF spacing correction. This corresponds to the Italian KM3NeT site near Capo Passero [13, 11, 14].

A.3 Correct time series

After merging the individual time series, is put through a barycentric correction of the time stamps. In the context of this study, this step is mandatory, as the analysis targets the pulsars listed in section 2.2.3. The selection of these pulsars, as well as the motivation for their choice, is discussed in more detail in that section. The resulting dataset can contain gaps and irregular time sampling caused by run transitions, detector downtime, and especially because of the barycentric correction. Since the following Fourier analysis requires evenly spaced data points, the time series is resampled onto a regular time grid. In order to do that, the script `KM3NeTCorrectTimeSeries.py` is applied.

A.3.1 Barycentric correction

The relative motion of the Earth around the Sun (orbital motion) and the rotation of the Earth (diurnal effect) cause a shift in the neutrino arrival times and hence in the detected photon hits.

This must be corrected for. To achieve this, a coordinate transformation from the rotating frame of the Earth to a non-rotating (inertial) frame is applied. The frame typically used for this purpose is the barycentre of the Solar System. This is also implemented in the code used here. Although, to be precise, the Solar System Barycentre defines a quasi-inertial reference frame, it is sufficiently inertial for high-precision timing analyses.

This limitation arises from the presence of external gravitational fields, such as the Galactic potential and due to nearby stars.

The barycentre refers to the centre of mass of multiple gravitationally interacting objects; in other words, it is the point about which all bodies, including in this case the Earth, orbit.

In the case of Astropy, it is computed using JPL Solar System ephemerides, which provide high-accuracy position and velocity data for the planets and the Moon referenced to the Solar System Barycentre. In this work, the DE440/DE441 (Astropy default) ephemerides are employed, as described in [41].

We therefore perform this transformation to obtain the arrival times of the hits in the detector as they would be at the barycentre of the Solar System (SSB). This takes the projected travel time into account. An overview of the barycentric correction is described in [29] and [5].

On a formal level, the time offset Δt to the SSB is first calculated as follows. Neutrinos are again assumed to propagate at the speed of light, such that barycentric corrections are identical to those applied in photon timing:

$$\Delta t = \frac{\vec{r}_{\text{obs}}(t) \cdot \hat{n}}{c}$$

where

- $\vec{r}_{\text{obs}}(t)$ is the detector position relative to the Solar System Barycentre,
- \hat{n} is the unit vector towards the pulsar,
- c is the speed of light.

The barycentric arrival time is then given by

$$t_{\text{SSB}} = t_{\text{obs}} + \Delta t.$$

A positive Δt corresponds to the detector being closer to the source than the Solar System Barycentre along the line of sight.

A.3.2 Resampling

Subsequently, the time series is resampled onto a uniform temporal grid. In this case, on a grid with again 100 ms as grid size. During this step, each data point is assigned to the temporal bin whose centre is closest to its barycentrically corrected timestamp. This effectively converts the irregularly sampled event times into a regularly sampled time series suitable for further analysis.

A.3.3 Mean padding

Finally, with the barycentric correction and the resampling process, there could now be bins that have no assigned value. Those gaps in the time series are filled using mean padding, whereby missing bins are assigned the mean value of the observed time series. This is done to avoid artificial discontinuities. Mean padding allows for a constant baseline and minimises the introduction of high-frequency components that would otherwise arise from abrupt gaps in the data.

A.4 Analysis

The now corrected time series is ready to be transformed into the frequency domain using an FFT.

This procedure is done with the script `KM3NeT_FFT_Analysis.py`.

A.4.1 Zero padding

Before computing the FFT, the time series is zero-padded to the next power of two, as described in 2.1.1.4, point 3. Zero padding improves the computational efficiency of the FFT algorithm without introducing additional information.

A.4.2 Red noise

Initial investigations show that, similarly to the results reported in [5] for ANTARES, the power spectra obtained for both KM3NeT-ORCA and KM3NeT-ARCA have a red-noise behaviour. This is characterised by an excess of power at low frequencies, with the spectral power decreasing as a function of increasing frequency. In Figure 3.1, one can see this red-noise spectrum.

To mitigate this effect, a red-noise filter is applied to the power spectrum. For that, one needs to estimate the underlying red-noise component. This is done by computing smoothed versions of the power spectrum by dividing it into different segments. After having the smoothed version, they represent a local average noise level as a function of frequency. The original power spectrum is then normalised by this estimate, effectively flattening the spectrum and suppressing low-frequency excess power. This transforms the red-noise-dominated spectrum into an approximately white-noise spectrum. And for that, there is a clearer understanding of significance. In this case, the segments are chosen logarithmically based on the size of the time series.

A.4.3 Harmonic Summation

Optionally an Harmonic summation is applied, the details of which can be found in section 3.1.2. In the code, it is the option `-H`, and it allows for $H \in \{1, 2, 4, 8\}$ where 1 corresponds to no harmonic summation.

A.4.4 Statistical significance

The statistical significance of peaks in the resulting power spectrum is evaluated using χ^2 statistics. Under the assumption of white Gaussian noise, which the red-noise filter gives a reasonable approximation for, the FFT power in each frequency bin follows a χ^2 distribution with two degrees of freedom (for no harmonic summation, adjust accordingly for harmonic summation). Significance thresholds are derived from that, and frequency bins exceeding the chosen confidence level, in this case 5σ , are identified as candidate periodic signals.

Finally, the script writes the list of detected candidates to a `.txt` file and produces plots of the power spectrum for visual inspection.

Appendix B

Appendix B: Implementation Details of the Sensitivity Pipeline

The code that was used can be found in the git repository [26]. This section is intended for someone who may want to understand the pipeline; as such, there is some duplication with the main text.

First, it is necessary to determine the received flux at a given energy that would be required to achieve a 5σ significance. This calculation is performed using the script `KM3NeT_SensitivityScan.py`.

Subsequently, one must evaluate what this required received flux implies for the intrinsic flux that needs to reach the detector. This step is carried out using `KM3NeT_Sensitivity_plots.py`.

B.1 Inject Signal and Significance Estimation

The code `KM3NeT_SensitivityScan.py` first injects a signal and then performs the FFT analysis to determine the mean σ and detection probability as follows:

Given a chosen frequency and a target signal-to-noise ratio (SNR), an artificial periodic signal is injected into the event-rate data. The standard deviation of the original rate data is used, together with the desired SNR, to define the amplitude of a sinusoidal modulation. The injected signal is then added to the original rate:

$$\text{rate}_{\text{new}}(t) = \text{rate}_{\text{old}}(t) + \text{SNR} \cdot \sigma(\text{rate}_{\text{old}}) \cdot \sin(2\pi f t_{\text{bin}}),$$

where f is the injected frequency and t_{bin} denotes the time associated with each bin. To determine how likely it is to detect a signal with a given SNR, the same FFT-based analysis described in section 3.1 is applied. However, it is not only applied once but multiple times. Between each iteration, the rate data are randomly scrambled in the time domain, thereby pseudo-simulating different realisations of the background.

For each pseudo-experiment, the FFT is computed, and the resulting spectrum is checked for a peak at the input frequency exceeding the detection threshold corresponding to 5σ significance. A counter records how many of these trials surpass this threshold. This is called the detection probability.

In addition, a median significance is defined as the median of all significance values obtained from the ensemble of pseudo-simulations. This quantity provides a robust

measure of the typical sensitivity at the chosen SNR.

The output is then a txt file with metadata, σ_{mean} , and detection probability for each SNR.

Here, the sanity check was performed on the following runs:

- KM3NeT_00000148_00018992.root
- KM3NeT_00000148_00018993.root
- KM3NeT_00000148_00018994.root

B.2 Cross-section

As the cross-section, it was decided to take the standard cross-section of Genie [32]. Especially, the package of the cross-sections can be found at [33]. It was also decided to take the cross-sections that are produced on the interaction with oxygen ^{16}O , as that seemed like a good proxy for water. After opening [33] `genie_xsec-3.06.00-noarch-G1802a00000-k250-e1000.tar.bz2` was downloaded and navigated to `\genie_xsec\v3_06_00\NULL\G1802a00000-k250-e1000\data\xsec_graphs.root`. This Root file was then put through `xs.ipynb` in order to get the cross-section files for the CC interaction for each neutrino flavour on oxygen in the form of a .txt file.

B.3 Effective volume

For that, a Snakemake [42] was written. It first generates, using a singularity file (`gseagen_v7.4.3.sif`), monoenergetic neutrinos in a cylindrical volume around a single DOM. The energy is the same as the energies in the cross-section file, so that they correspond to each other. After that, it uses `jsirene` [34] (`Jpp_v19.3.0.sif`) to propagate the neutrinos. In the end, the number of neutrinos that produced at least one hit is compared to the number produced initially. Both are written down in a .txt file. From the ratio of those, the effective volume is derived.

B.4 Necessary Rate

Now, to convert the output of the sensitivity scan from section B.1 into a physical rate, we need to do some operations.

This is done by the script `KM3NeT_Sensitivity_plots.py`

Firstly, the background is calculated using Equation 3.2 and then secondly using the output .txt file from the sensitivity scan, `KM3NeT_Sensitivity_plots_Cases.py` calculates what the rate for $\sigma_{\text{mean}} = 5\sigma$ would be using Equation 3.4 and draws a contour on the graph. It also calculates the luminosity the source would need to have, then it outputs the graph as a .png and .pdf file.

Appendix C

Appendix C: Plots




Figure C.1: Flux sensitivity for pulsar J2325+6316.




Figure C.2: Flux sensitivity for pulsar J0750-6846.




Figure C.3: Flux sensitivity for pulsar J0820-4114.




Figure C.4: Flux sensitivity for pulsar J0828-3417.

Figure C.5: Neutrino flux sensitivity contours for the analysed pulsars. Each contour corresponds to the minimum flux at Earth required to achieve a periodic signal detection with at least 5σ significance. Fluxes above the contour are detectable, while fluxes below are not. This corresponds to the Case A 3.3.3.1, i.e. the equidistribution of all flavours of neutrino.




Figure C.6: Flux sensitivity for pulsar J2325+6316.

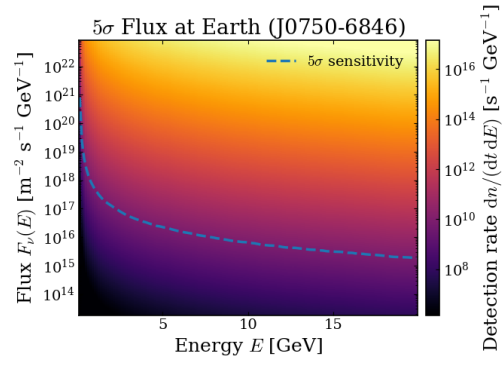


Figure C.7: Flux sensitivity for pulsar J0750-6846.




Figure C.8: Flux sensitivity for pulsar J0820-4114.




Figure C.9: Flux sensitivity for pulsar J0828-3417.

Figure C.10: Neutrino flux sensitivity contours for the analysed pulsars. Each contour corresponds to the minimum flux at Earth required to achieve a periodic signal detection with at least 5σ significance. Fluxes above the contour are detectable, while fluxes below are not. This corresponds to the Case B 3.3.3.2, i.e. an initial electron monoflavour that traverses astronomical distances.




Figure C.11: Luminosity sensitivity for pulsar J2325+6316.




Figure C.12: Luminosity sensitivity for pulsar J0750-6846.

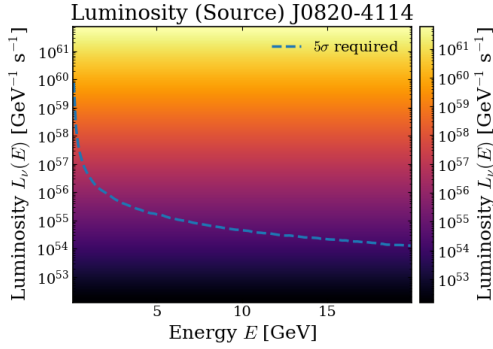


Figure C.13: Luminosity sensitivity for pulsar J0820-4114.




Figure C.14: Luminosity sensitivity for pulsar J0828-3417.

Figure C.15: Neutrino luminosity sensitivity contours for the analysed pulsars. Each contour corresponds to the minimum luminosity of the pulsar required to achieve a periodic signal detection with at least 5σ significance. Luminosities above the contour are detectable, while luminosities below are not. This corresponds to the Case A 3.3.3.1, i.e. the equidistribution of all flavours of neutrino.




Figure C.16: Luminosity sensitivity for pulsar J2325+6316.




Figure C.17: Luminosity sensitivity for pulsar J0750-6846.




Figure C.18: Luminosity sensitivity for pulsar J0820-4114.


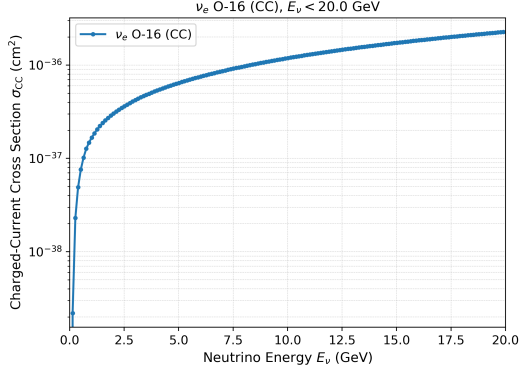



Figure C.19: Luminosity sensitivity for pulsar J0828-3417.


Figure C.20: Neutrino luminosity sensitivity contours for the analysed pulsars. Each contour corresponds to the minimum luminosity of the pulsar required to achieve a periodic signal detection with at least 5σ significance. Luminosities above the contour are detectable, while luminosities below are not. This corresponds to the Case B 3.3.3.2, i.e. an initial electron monoflavour that traverses astronomical distances.



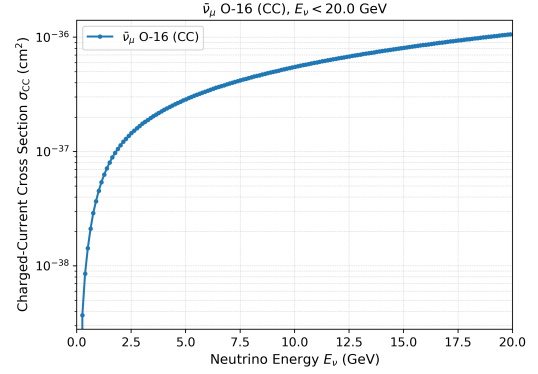
(a) ν_e CC on ^{16}O



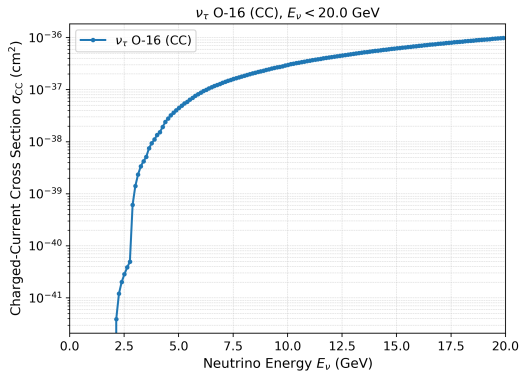
(b) $\bar{\nu}_e$ CC on ^{16}O




(c) ν_μ CC on ^{16}O



(d) $\bar{\nu}_\mu$ CC on ^{16}O



(e) ν_τ CC on ^{16}O



(f) $\bar{\nu}_\tau$ CC on ^{16}O

Figure C.21: Charged-current (CC) neutrino and antineutrino cross sections on ^{16}O for energies below 20 GeV, as obtained from the GENIE cross-section package.

Bibliography

- [1] K. Hirata and others (Kamiokande-II Collaboration). “Observation of a Neutrino Burst in Coincidence with Supernova 1987A in the Large Magellanic Cloud”. In: *Physical Review Letters* 58.14 (1987). First reported detection of neutrinos from SN 1987A, pp. 1490–1494. DOI: 10.1103/PhysRevLett.58.1494.
- [2] Alexei V. Filippenko. “Optical Spectra of Supernovae”. In: *Annual Review of Astronomy and Astrophysics* 35 (1997), pp. 309–355. DOI: 10.1146/annurev.astro.35.1.309.
- [3] B. P. Abbott, R. Abbott, T. D. Abbott, et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Phys. Rev. Lett.* 116 (6 Feb. 2016), p. 061102. DOI: 10.1103/PhysRevLett.116.061102. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102>.
- [4] J. W. T. Hessels, S. M. Ransom, I. H. Stairs, et al. “A Radio Pulsar Spinning at 716 Hz”. In: *Science* 311.5769 (2006). Discovery of fastest-spinning pulsar PSR J1748-2446ad, pp. 1901–1904. DOI: 10.1126/science.1123430.
- [5] Maximilian Eff. “Search for low-energy periodic neutrino sources with ANTARES”. Master’s Thesis. Erlangen, Germany: Friedrich-Alexander-Universität Erlangen-Nürnberg, Jan. 2023. URL: https://ecap.nat.fau.de/wp-content/uploads/2023/03/KM3NeT_THES_2023_001_MSc_Periodic_low_energy_neutrino_sources.pdf.
- [6] Corinne Pralavorio. *A demonstrator magnet produces a record magnet field*. Accessed: 2026-02-02. CERN. Mar. 2020. URL: <https://home.cern/news/news/accelerators/demonstrator-magnet-produces-record-magnet-field>.
- [7] Alexander A Mushtukov, Sergey S Tsygankov, Valery F Suleimanov, and Juri Poutanen. “Ultraluminous X-ray sources as neutrino pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 476.3 (Feb. 2018), pp. 2867–2873. ISSN: 0035-8711. DOI: 10.1093/mnras/sty379. eprint: <https://academic.oup.com/mnras/article-pdf/476/3/2867/24422267/sty379.pdf>. URL: <https://doi.org/10.1093/mnras/sty379>.
- [8] Amato, E., Guetta, D., and Blasi, P. “Signatures of high energy protons in pulsar winds”. In: *A & A* 402.3 (2003), pp. 827–836. DOI: 10.1051/0004-6361:20030279. URL: <https://doi.org/10.1051/0004-6361:20030279>.

- [9] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs. *The ATNF Pulsar Catalogue*. <http://www.atnf.csiro.au/research/pulsar/psrcat>. Accessed: 2025-12-21. 2005.
- [10] S Adrián-Martínez, M Ageron, F Aharonian, et al. “Letter of intent for KM3NeT 2.0”. In: *Journal of Physics G: Nuclear and Particle Physics* 43.8 (June 2016), p. 084001. DOI: 10.1088/0954-3899/43/8/084001. URL: <https://doi.org/10.1088/0954-3899/43/8/084001>.
- [11] Uli Katz. *KM3NeT – The next-generation neutrino telescope in the Mediterranean Sea*. Presentation at VHEPA 2016, Honolulu, Hawaii. https://ecap.nat.fau.de/wp-content/uploads/2017/03/UliKatz_vhepa16-1.pdf. Jan. 2016.
- [12] KM3NeT Collaboration. *KM3NeT-Fr neutrino telescope installation site*. <https://www.km3net.org/research/research-infrastructure/km3net-fr-site/>. Accessed on February 20, 2026.
- [13] KM3NeT Collaboration. *KM3NeT-It deep-sea installation site (ARCA) – Capo Passero, Sicily*. <https://www.km3net.org/research/research-infrastructure/km3net-it-site/>. Accessed on February 20, 2026.
- [14] S. Biagi. *The KM3NeT Neutrino Telescope: Status and Perspectives*. Presentation at CERN Indico. https://indico.cern.ch/event/393078/contributions/2195242/attachments/1333594/2005124/Biagi_KM3NeT.pdf. 2017.
- [15] S. Aiello, A. Albert, S. Alves Garre, et al. “Determining the neutrino mass ordering and oscillation parameters with KM3NeT/ORCA”. In: *European Physical Journal C* 82.1 (2022), p. 26. DOI: 10.1140/epjc/s10052-021-09893-0.
- [16] KM3NeT Collaboration, S. Aiello, A. Albert, et al. “Sensitivity to light sterile neutrino mixing parameters with KM3NeT/ORCA”. In: *Journal of High Energy Physics* 2021.10 (2021), p. 180. DOI: 10.1007/JHEP10(2021)180.
- [17] Alba Domi. “Lorentz Invariance Violation with KM3NeT/ORCA115”. In: *PoS ICRC2023* (2023), p. 1086. DOI: 10.22323/1.444.1086.
- [18] S. Aiello et al. (The KM3NeT Collaboration). “Astronomy potential of KM3NeT/ARCA”. In: *European Physical Journal C* (2024). arXiv:2402.08363 [astro-ph.HE].
- [19] S. Aiello, A. Albert, M. Alshamsi, et al. “The KM3NeT multi-PMT optical module”. In: *Journal of Instrumentation* 17.07 (July 2022), P07038. DOI: 10.1088/1748-0221/17/07/P07038. URL: <https://doi.org/10.1088/1748-0221/17/07/P07038>.
- [20] KM3NeT Collaboration. “Measurement of neutrino oscillation parameters with the first six detection units of KM3NeT/ORCA”. In: *J. High Energ. Phys.* 2024.206 (2024). ORCA6 initial configuration; first six detection units; neutrino oscillation parameters, pp. 1–31. DOI: 10.1007/JHEP10(2024)206. arXiv: 2408.07015 [hep-ex].
- [21] S. Aiello, A. Albert, M. Alshamsi, et al. “Astronomy potential of KM3NeT/ARCA”. In: *The European Physical Journal C* 84.9 (Sept. 2024). ISSN: 1434-6052. DOI:

- 10.1140/epjc/s10052-024-13137-2. URL: <http://dx.doi.org/10.1140/epjc/s10052-024-13137-2>.
- [22] Nicole Geisselbrecht. “Event Classification and Energy Reconstruction for ANTARES using Convolutional Neural Networks”. Master’s thesis in Physics. Master’s thesis. Friedrich-Alexander-Universität Erlangen-Nürnberg, 2021. URL: https://ecap.nat.fau.de/wp-content/uploads/2021/05/Masters_Thesis_Nicole_Geisselbrecht.pdf.
- [23] Ronald N. Bracewell. *The Fourier Transform and Its Applications*. 3rd ed. McGraw-Hill Series in Electrical and Computer Engineering. Circuits and Systems. McGraw-Hill, 2000. ISBN: 0-07-303938-1.
- [24] James W. Cooley and John W. Tukey. “An Algorithm for the Machine Calculation of Complex Fourier Series”. In: *Mathematics of Computation* 19.90 (1965), pp. 297–301. ISSN: 00255718, 10886842. URL: <http://www.jstor.org/stable/2003354> (visited on 01/30/2026).
- [25] Scott M. Ransom, Stephen S. Eikenberry, and John Middleditch. “Fourier Techniques for Very Long Astrophysical Time-Series Analysis”. In: *The Astronomical Journal* 124.3 (Sept. 2002), p. 1788. DOI: 10.1086/342285. URL: <https://doi.org/10.1086/342285>.
- [26] R. Petri. *periodic_sources*. https://git.km3net.de/rgracia/periodic_sources. Branch RobertKM3Net2, commit c386f02cdd7927b494e73c8f511eb8a5dff98f45; maintained by Gracia, R. 2026.
- [27] Planck Collaboration. “Planck 2018 results. VI. Cosmological parameters”. In: *Astronomy & Astrophysics* 641 (2020), A6. DOI: 10.1051/0004-6361/201833910.
- [28] KATRIN Collaboration. “Direct neutrino-mass measurement with sub-electronvolt sensitivity”. In: *Nature Physics* 18 (2022), pp. 160–166. DOI: 10.1038/s41567-021-01463-1.
- [29] Katharina Jurk. *Periodic Low Energy Neutrino Search in a High Energy Neutrino Detector*. Research Project. Erlangen Centre for Astroparticle Physics (ECAP), Friedrich-Alexander-Universität Erlangen-Nürnberg, May 2024. URL: https://ecap.nat.fau.de/wp-content/uploads/2024/05/2024-04_KatharinaJurk_RP_PeriodicNeutrinoSources.pdf.
- [30] Ingrid H. Stairs. “Testing General Relativity with Pulsar Timing”. In: *Living Reviews in Relativity* 6.1 (2003), p. 5. ISSN: 1433-8351. DOI: 10.12942/lrr-2003-5. URL: <https://doi.org/10.12942/lrr-2003-5>.
- [31] K. Adánek, J. Roy, and W. Armour. “A novel greedy approach to harmonic summing using GPUs”. In: *Astronomy and Computing* 40 (2022), p. 100621. DOI: 10.1016/j.ascom.2022.100621. URL: <https://www.sciencedirect.com/science/article/pii/S2213133722000464>.
- [32] S. Aiello et al. “gSeaGen: The KM3NeT GENIE-based code for neutrino telescopes”. In: *Computer Physics Communications* 256 (2020), p. 107477. DOI: 10.1016/j.cpc.2020.107477.

- [33] SciSoft Team. *GENIE Cross Section Package v3_06_00*. https://scisoft.fnal.gov/scisoft/packages/genie_xsec/v3_06_00/. Accessed: 2026-02-15. 2025.
- [34] Maarten et al. de Jong. *JSirene: Simulation of the KM3NeT detector response*. Software package. Version used: v19.3.0 in the Jpp framework. Available via Docker: git.km3net.de:4567/common/jpp:latest. Documentation: <https://common.pages.km3net.de/jpp/>. 2023.
- [35] Roberto Aloisio, Eugenio Coccia, and Francesco Vissani, eds. *Multiple Messengers and Challenges in Astroparticle Physics*. Cham: Springer, 2018. ISBN: 978-3-319-65423-2. DOI: 10.1007/978-3-319-65423-2.
- [36] *KM3NeT: Technical Design Report*. Tech. rep. Section 5.2.1: Temperature; shows stable deep-sea water temperatures of about 13.3–13.5°C at detector sites. KM3NeT Collaboration, 2023. URL: https://www.km3net.org/wp-content/uploads/2023/05/KM3NeT_DS_TDR-published-in-2010.pdf.
- [37] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. “Julia: A Fresh Approach to Numerical Computing”. In: *SIAM Review* 59.1 (2017), pp. 65–98. DOI: 10.1137/141000671.
- [38] Tamas Gal, Johannes Schumann, and Santiago Peña Martínez. *KM3io.jl: Julia library to access KM3NeT related dataformats*. 2025. URL: <https://github.com/KM3NeT/KM3io.jl/> (visited on 01/31/2026).
- [39] Python Software Foundation. *Python Language Reference*. 2023. URL: <https://www.python.org/>.
- [40] Astropy Collaboration, T. P. Robitaille, E. J. Tollerud, et al. “Astropy: A community Python package for astronomy”. In: *Astronomy & Astrophysics* 558 (2013), A33. DOI: 10.1051/0004-6361/201322068.
- [41] Ryan S. Park, William M. Folkner, James G. Williams, and Dale H. Boggs. “The JPL Planetary and Lunar Ephemerides DE440 and DE441”. In: *The Astronomical Journal* 161.3 (Feb. 2021), p. 105. DOI: 10.3847/1538-3881/abd414. URL: <https://doi.org/10.3847/1538-3881/abd414>.
- [42] Johannes Köster and Sven Rahmann. “Snakemake—a scalable bioinformatics workflow engine”. In: *Bioinformatics* 28.19 (2012), pp. 2520–2522. DOI: 10.1093/bioinformatics/bts480.

Acknowledgements

A thesis is nothing one does alone, hence I need to thank some people:

- The foremost thanks go to my supervisors, PD Dr Thomas Eberl and Dr Rodrigo Gracia-Ruiz, for their invaluable guidance throughout this thesis. They were always available to answer my questions promptly and provided consistent support in navigating the topic.
- I would also like to thank Frederik Andersen for his help with code-related problems, particularly during the Genie phase of the thesis. His presence in the office was enjoyable and provided great motivation throughout the writing process.
- I would also like to thank the Erlangen National High Performance Computing Centre (NHR@FAU) at the Regionales Rechenzentrum Erlangen, Friedrich-Alexander-Universität Erlangen-Nürnberg, for hosting and providing access to their HPC clusters, which made running my code and simulations possible.

Declaration of Originality

I, Robert Petri, student registration number: 23103917, hereby confirm that I completed the submitted work independently and without the unauthorized assistance of third parties and without the use of undisclosed and, in particular, unauthorized aids. This work has not been previously submitted in its current form or in a similar form to any other examination authorities and has not been accepted as part of an examination by any other examination authority.

Where the wording has been taken from other people's work or ideas, this has been properly acknowledged and referenced. This also applies to drawings, sketches, diagrams and sources from the Internet.

In particular, I am aware that the use of artificial intelligence is forbidden unless its use as an aid has been expressly permitted by the examiner. This applies in particular to chatbots (especially ChatGPT) and such programs in general that can complete the tasks of the examination or parts thereof on my behalf.

Any infringements of the above rules constitute fraud or attempted fraud and shall lead to the examination being graded "fail" ("nicht bestanden").

Place, Date

Signature