# The modulation experiment:
Searching for trends in radioactive decays

Thomas Mons

July 2019

Supervisor: Dr. A.P. Colijn   Second reader: Dr. T. Peitzmann

University Utrecht

Abstract

This master thesis investigates the decay of three radioactive isotopes: ^{137}Cs,^{60}Co and ^{44}Ti. The direct motivation for this experiment are claims of modulation in radioactive isotopes of yearly, monthly and daily frequencies. The stability of the modulation experiment is first investigated, and estimations of the effect of environmental parameters are made. Then the claim of yearly modulation is investigated for all sources and excluded at 95% confidence for amplitudes above 2.57×10^{-4}. Via power spectrum analysis other frequencies are also investigated, and no evidence of modulation is found.

Contents

Chapter 1

Introduction

Radioactivity was discovered by Henri Becquerel in 1896, for which he got the Nobel prize in 1903 The Nobel Prize in Physics 1903 together with Pierre and Marie Curie [1]. At first it was thought to be similar to the x-rays discovered by Röntgen, but it was quickly found out that the radiation came from within the atoms and consisted of matter. The discovery of radioactivity was a great scientific breakthrough, and lead to other discoveries like nuclear fission.

Radioactive decays have always been believed to be a purely random Poissonian process, with exponentially decaying activity. Recently however, claims have been made concerning the stability of radioactive decays. These papers publish results where a modulation is visible in the rate of decay. The reported frequency of this modulation varies from 1 day to 1 year with an amplitude of approximately 0.1% of the activity of the source.

In order to investigate these claims a collaboration called the modulation experiment was set up. It consists of four identical setups at Amsterdam, Zürich, Purdue and Rio de Janeiro measuring the radioactive rate of four sources. The modulation experiment aims to have a better controlled environment than any activity measuring experiment, and the goal is to either confirm or reject the claims of modulation.

In this thesis I will analyse the data from the modulation experiment taken at Nikhef in Amsterdam over almost three years. I will analyse the stability of the setup, and will try to confirm or reject modulation of any frequency.

Chapter 2

Theory

2.1 Radioactivity

Radioactivity concerns the decay of atomic nuclei, during which a particle is emitted. Generally speaking there are three types of radio-active decay:

- α-decay, in which a nucleus emits a $^4_2$He nucleus

- β^--decay, in which a neutron becomes a proton emitting an electron

- β^+-decay, in which a proton becomes a neutron emitting an positron

- γ-decay, in which an excited nucleus emits a photon and returns to its ground state.

There are other types like neutron emission and spontaneous fission, but these are the most common. Although no α-sources are used in our experiment, there have been claims of radioactive modulation with α-sources. Inside buildings ^{222}Rn is present in the air, as it is part of the Uranium decay chain, which can be found in concrete.

2.1.1 The sources of the modulation experiment

The sources used in the modulation experiment are ^{54}Mn, ^{137}Cs, ^{44}Ti and ^{60}Co. ^{54}Mn, ^{137}Cs and ^{60}Co are β^- emitters that decay because a downquark in a neutron emits a W$^-$-boson to become an upquark, changing the neutron to a proton. The W$^-$-boson then decays into an electron and an antielectron neutrino. This process is shown in figure 2.1a. Afterwards the new nucleus often is in an excited state. In the case of ^{54}Mn and ^{137}Cs there is just one mode of excitation, which quickly deexcites. Producing a 835 keV and 661.7 keV γ-ray respectively. ^{60}Co decays into ^{60}Ni which has two excited states meaning that two separate γ-rays will be emitted from a single decay. The decays of ^{54}Mn, ^{60}Co and ^{137}Cs can be seen in figure 2.2 together with the decay product and the possible excitation levels of the nucleus which produce the γ-rays.

The decay of ^{44}Ti to ^{44}Sc happens via electron capture. As is displayed in figure 2.1b an electron from one of the orbitals is captured by a proton, with a W$^-$-boson as force mediator. The result is a neutron and an electron neutrino. The γ-rays from the deexcitation of ^{44}Sc are below 100 keV, as they are caused by other electrons falling back into the shell of the captured electron. This low energy makes them hard to measure. Instead we measure the decay of ^{44}Sc to ^{44}Ca via β^+ decay, which has a half life of 4 hours. In β^+ decay the opposite of the usual decay happens, and a proton becomes a neutron as shown in figure 2.1c. The produced positron annihilates and produces two 511 keV γ-rays which we measure, and we measure the deexcitation of ^{44}Ca. The complete decay of ^{44}Ti to ^{44}Ca is shown in figure 2.2d. In the modulation experiment, we do not measure the decay of ^{44}Ti directly, instead we only measure ^{44}Sc. The rate of ^{44}Sc is dependent on the rate of ^{44}Ti , where the four hours previous to the measurement have the biggest contribution to the amount of ^{44}Sc decays.

(a) A neutron decays to an electron

(b) A proton decays to a neutron via electron capture

(c) A proton decays to a neutron

Figure 2.1: Feynman diagrams of the decays occurring in the modulation experiment, made using [2]

.

2.2 Overview of modulation claims

The stability of radioactive decays over time has been analysed since its discovery. Potential processes affecting the half life have been investigated. Negligible effects were found for parameters like magnetic field, temperature or pressure [7]. Half lives can however be affected by introducing different chemical elements in the vicinity, or forming molecules with the radioactive atoms [8].
In 2009 Jenkins, Fischbach et al. [9] claimed a yearly modulation in radioactive decays with an amplitude in the rate after subtracting an exponential fit of $\approx 0.1\%$. They used data taken between 1983 and 1998, taken by Brookhaven National Laboratory and Physikalish-Technische-Bundesandstalt and correlated this to the earth-sun distance. This lead to the plots as shown in figure 2.3. With a Pearson correlation coefficient they then find a probability of $6 \cdot 10^{-18}$ that this correlation is not due to a physical effect but due to statistical fluctu-

(a) The decay of ^{137}Cs from [3]

(b) The decay of ^{54}Mn from [4]

(c) The decay of ^{60}Co from [5]

(d) The decay of ^{44}Ti from [6]

Figure 2.2: The decay diagrams for all sources in the modulation experiment.

ations.

There are many objectionable choices in this research, and I will mention some. First of all, the standard deviation is not shown in figure 2.3, creating a very suggestive picture. Second of all, no environmental parameters were measured.

Figure 2.3: The residual rate after subtracting an exponential decay of $^{32}Si/^{36}Cl$ in blue. In red the earth-sun distance to the power -2. From [9]

Instead they correct for environmental influences by dividing the rate of ^{32}Si by the rate of ^{36}Cl (half life of 301000 years), which they claim should largely cancel environmental effects. Why the rate of ^{36}Cl would not also show modulation which would then be divided out is not explained. Third of all, they do not explain why they choose to investigate this particular isotope. They claim the modulation is caused by a change in neutrino flux. If their is a measurable influence from neutrinos in β-decays, this should be measurable for all β-emitting isotopes, not just ^{32}Si. By cherry-picking a single isotope, the probability of the correlation coefficient should reflect all isotopes that did not show modulation, leading to a higher probability that this measurement was due to statistical fluctuations. This was unfortunately not done.

Another set of claims was made by Parkhomov. In this paper [10] he shows the residual rate for multiple sources, and claims that ^{60}Co shows a clear yearly modulation. No actual analysis is made of this data. For ^{90}Sr-^{90}Y a frequency analysis is done via a Fourier transform which I have shown in figure 2.4. Here a strong peak becomes visible for a yearly frequency. The article does look into the temperature of the setup, and shows the daily average of each day of the year. They do not correlate this to the rate however, and so the influence of temperature via thermal expansion cannot be excluded. When studying the graphs of temperature and rate, both graphs show a dip in September, this is however not discussed in the article. Via averaging residuals they try to show evidence of a modulation with the period of a lunar month. Strangely enough, no evidence of this monthly modulation can be seen in the Fourier analysis.

These two papers were just an example of the many papers on this subject,

Figure 2.4: A Fourier transformation of the residual rate of ^{90}Sr-^{90}Y, with a clear peak at yearly frequency. Peaks at a frequency of 2,3,4 yr^{-1} are also visible. From [10]

and the objections that can be made against them. In table 2.1 I have shown an overview of modulation claims per isotope, it was based on the table one can find in this article [11] by Jenkins and Fischbach.

2.2.1 Correlation between astronomical events and radioactive decay

Aside from modulation of radioactive decays, there are also claims of solar flares or neutron star mergers affecting the nuclear decay rate [12]. An improvement is that Jenkins and Fischbach take their own data, instead of relying on others. Once again they do seem to be cherry-picking. Each year roughly ten solar flares occur, but not all are investigated in this paper. It remains unexplained why these specific solar flares are investigated, and the significance of the correlation should be corrected for that. The period leading up to the solar flare and neutron star merger is also investigated, which implies effects faster than the speed of light .

2.3 Photon-matter interactions

The modulation experiment described in this thesis aims to measure radioactive rates very precisely. Instead of a simple Geiger-counter, the energy spectrum is measured with scintillation detectors and photomultiplier tubes. To extract the rate from those spectra, an understanding of the components of the spectrum is necessary. Here I will explain what processes affect the γ-rays as they pass through the scintillation detector. We do not measure the electrons from the decay, as a 1.1 meV electron passes through a maximum of 4 mm of plastic (7 half-value ranges) [13]. The γ-rays can interact with matter in a number of

Table 2.1: An overview of isotopes that have claims of modulation attached to them, based on [11]

Isotope	Observed periods
^{3}H	1 yr, 1 day, 12.1 yr
^{22}Na	1yr
^{36}Cl	1 yr, 11.7 yr, 2.1 yr
^{44}Ti	1 yr
^{54}Mn	1 yr
^{56}Mn	1 yr
^{60}Co	1 yr, 12.1 yr, 1 day
^{85}Kr	1 yr
^{90}Sr-^{90}Y	1 yr, 11.7 yr
^{108}Ag	1 yr
^{133}Ba	1 yr
^{137}Cs	12.1 yr, 1 day
^{152}Eu	1 yr
^{222}Rn	1 yr, 2.1 yr, 11.7 yr
^{226}Ra	1 yr, 2.1 yr, 11.7 yr
^{239}Pu	1 day, 1yr, 13.5 yr

ways, mainly via:

- the photo-electric effect,

- Rayleigh scattering,

- Compton scattering,

- and pair production.

Each process has a photon energy region in which it is most prominent. This is displayed in figure 2.5. As we start out with roughly 1 MeV γ-rays, pair production is unlikely to happen a lot. And Rayleigh scattering is only relevant for much lower energies, as the γ-rays between 0.5 and 2.5 meV. The main processes in our scintillation crystals are Compton scattering and photo-electric absorption. The information in this section was based on [14].

2.3.1 The photo-electric effect

The photo-electric effect consist of a photon being absorbed by an electron, and that electron being ejected from its atomic shell. This only happens when the photon energy is close to the binding energy of the electron, which is lower than the initial energy of the γ-rays. The electron then deexcites and emits a new photon.

Figure 2.5: A graph of the cross section between photons and atoms depending on the photon energy for carbon ($A = 12$). The coherent scattering is Rayleigh scattering, whereas the incoherent scattering is Compton scattering. In low energy the photo-electric effect dominates, at high energy pair production. From [15]

2.3.2 Rayleigh scattering

In Rayleigh scattering a photon scatters elastically on electrons in the shell of the atoms, losing no energy in the process. That means that it is largely irrelevant for our detectors. It also happens at much lower energies (e.g. visible light, causing blue skies), while this experiment only measures γ-rays.

2.3.3 Compton scattering

In Compton scattering the incident photon collides inelastically with a particle of the matter it is passing through. In most cases this will be an electron in the outer shells of the atom, as the nucleus of the atom is very small, and is unlikely to be hit. During the collision the photon transfers some of its energy to the electron and changes direction. The new wavelength of the photon (and

Figure 2.6: The Feynman diagram for pair production in which a photon interacts with a (virtual) photon, producing an electron and a positron. Made with [2].

thus its energy) is given by:

$$\lambda' - \lambda = \frac{h}{m_e c}(1 - \cos\theta), \tag{2.1}$$

where λ' is the new wavelength, λ the old wavelength, m_e the electron mass and θ the angle of scattering. The electron hit by the photon is ejected from the atom, resulting in an ionized atom and a free electron. This free electron will quickly fall back to a bound state in the atom producing a photon.

2.3.4 Pair production

In pair production a virtual photon from atoms in the material collides with the incident photon to produce an electron and positron, which can scatter and will produce new photons. This process is shown in figure 2.6 The positron will eventually annihilate with an electron from a nearby atom, producing two 511 keV photons, whereas the electron will fall back into an atom orbit and produce a photon.

Chapter 3

The modulation experiment

In the previous chapter I discussed radioactivity and claims of radioactive modulation. The data analysed in this thesis were taken by the modulation experiment, an experiment specifically designed to confirm or reject claims of modulation in β-emitters. The experiment is designed to control or measure all possible influences on decay rates and setup, for example temperature and background rate. The setup is run at 4 different locations in the world:

- Amsterdam

- Zurich

- Brazil

- Purdue

The benefit is that radioactive isotopes are monitored in different time zones in both hemispheres, so that any seasonal or daily effects of weather and climate can be excluded.

3.1   The setup

The setup for each of the four modulation experiments is identical, except for the sources. I will discuss the setup briefly, but extensive information can be found in this thesis [16]. The full setup including surrounding computers is shown in figure 3.1.

Each setup monitors three sources and background. At Nikhef detectors (called channels) 0 and 1 measure background, 2 and 3 measure ^{44}Ti, 4 and 5 measure ^{60}Co, and 6 and 7 measure ^{137}Cs. All sources had an activity of approximately 1 kBq when the experiment started. The distribution of the sources over the setups is shown in table 3.1. We have two thallium doped sodium iodine detectors with photomultipliers attached to them for each. These detectors were all build by Scionix, and specifications can be found here [17].

Figure 3.1: The complete modulation setup.(A) contains the sources, detectors, heaters and environmental measuring instruments. (B) is the insulation of about 10 cm polyurethane foam. (C) contains all electronics such as the high voltage supply, heater supply, radon monitor and data acquisition computer. Figure from [16]

Table 3.1: All sources measured in the modulation experiment, their half lives and their locations

| Source	Q-value	γ-energy (keV)	Half life (yr)	Locations
Background (^{40}K)	1504.9 keV	1460 keV	N.A.	all
^{44}Ti	511,1157	267.5 keV	59.1 ± 0.3 [6]	all
^{60}Co	1173,1333	2823.07 keV	5.2711 ± 0.0004 [5]	all
^{54}Mn	835	1377.2 keV	0.85534 ± 0.00055 [4]	Purdue, UZH
^{137}Cs	661.7	1175.63 keV	30.08 ± 0.09 [3]	CBPF, Nikhef

Figure 3.2: A pair of NaI (Tl) detectors in their holders, with a container with a source in the middle. The whole setup is fixed to a aluminium frame. Taken from [18]

All detectors are fixed to a frame, facing inwards towards a source to get as close as possible to a 4π coverage. A single source with two detectors is shown in figure 3.2. One detector pair is operated without a radioactive source in order to measure the natural background level in the setup over time. These four arrays of detectors are fastened to an aluminium frame and placed in a box. The detectors are then encased in approximately 5 cm of lead. Between the arrays two heater bands are placed. An Arduino computer was also placed in a box with a thermometer, hygrometer, magnetometer and barometer connected to it. Specifications and resolution for these instruments can be found in table 3.2. The box is then sealed, except for a feedthrough for a small (max 10 L per hour) N_2 flush to lower the ^{222}Rn concentration in the air. The goal of the modulation experiment is to measure radioactive isotopes in a stable environment, this is why the box is heated above room temperature (29 °C), and placed in an insulated container. Note that the specific temperature is irrelevant, as long as it is stable.

Outside this insulated box is our data acquisition system (DAQ). The output signal from the PMT's is fed into an analog to digital converter (ADC), which then feeds into our DAQ computer (PZIe-1071). All details on electronics compononents can be found in the thesis of F. Gjaltema [19]. The DAQ computer runs a Labview program to control the temperature, collect the data, and show some details on a monitor. The data is then transferred to a normal PC via an optical fiber and is then saved on the Nikhef /data storage.

Table 3.2: A table of all measuring instruments included in the modulation experiment. All except the high voltage and Radon activity measurements are connected to an Arduino inside the setup. Table from [18]

parameter	Manufacturer	Model	Resolution	Unit
Temperature	Bosch	BMP085	0.1	$^{\circ}$C
Pressure	Bosch	BMP085	3	Pa
Magnetic field	Honeywell	HMC5883L	12	mG
Relative Humidity	Hoe RF	HH10D	0.1	%
High voltage	CAEN	DT5533P	0.1	V
Radon activity	Durridge	RAD7	0.8	Bq/m^3

3.2 Scintillation detectors

Scintillation detectors are any detector where a particle goes through a detector material and deposits energy mainly via Compton scattering and the photo-electric effect, details on these processes can be found in section 2.3. The produced light is subsequently measured, usually with a photomultiplier (PMT). There are many scintillation materials available, but sodium iodide (NaI) is one of the most frequently used. It has a high light yield compared to other sources [20], and thus a high energy resolution as this is proportional to the square root of the number of photons produced. It also has a relatively quick excitation decay time of 0.23 μs [20]. It is also one of the cheaper scintillation materials. The biggest downside of NaI crystals is that they dissolve in water, and thus need to be stored in a moisture free environment.

All 32 detectors used in the modulation experiment are of the same type made by Scionix [17]. They consist of a thallium doped NaI crystal of 3×3 inch connected to a photomultiplier tube.

3.2.1 Photomultipliers

To measure the photons produced in the NaI crystal a photomultiplier (PMT) is used. A schematic drawing of a PMT is shown in figure 3.3. When incoming photons hit the photocathode of the PMT, they produce an electron via the photo-electric effect. The PMT consists of multiple electrical fields produced by dynodes which accelerate electrons. These accelerated electrons smash into a dynode and release more electrons, which are then accelerated to the next dynode. This process happens multiple times and produces an exponential growth in the number of electrons. At the end of the PMT the final bunch of electrons lands on the anode, and results in a pulse that is large enough to be measured. The number of electrons produced is strongly dependent on acceleration of the electrons, and thus of the voltage over the PMT. This makes a stable voltage crucial to the energy measurements.

Figure 3.3: An image of a scintillation crystal with PMT on it. Note that there are several dynodes between which the electrons get accelerated. A sketch of a PMT pulse is also shown. Image from [14]

Chapter 4

Processing

In this chapter I will discuss the processing procedure for the modulation experiment. This starts at the waveforms coming from the PMT's into the analog to digital converter (ADC), to the extraction of the rates of decay.

4.1 Waveforms

The PMT's measure the photons produced by the crystal with a sampling rate of roughly 10 ns. Since the excitation of thallium has a decay time of 230 ns [20], we expect the waveform to be an exponential decay with a half life of the same size. A measured waveform is shown in figure 4.1. The maximum height is reached at approximately 0.4 μs, at approximately 0.65 μs the pulse is halved, meaning the expected half life is measured. Each PMT has a baseline output, which is subtracted from the waveform. The integral of the remaining waveform is then proportional to the energy deposited in the scintillator.

When studying the waveforms, three distinct types of bad waveform can be identified:

- the ADC overflows leading to an incomplete waveform (figure 4.2a),

- the tail of a previous event comes into the baseline measurement for the current event, throwing of the integral (figure 4.2b), classified by setting a maximum value for the root mean squared (RMS) of the baseline,

- a pileup event where there are two waveforms in one event (figure 4.2c).

In figure 4.3 I have shown the rate of error per channel for the channels with sources. As visible the total rate of error per channel stays around 1 Hz, whereas the rate of our sources does not fall under 50 Hz for the low energies. The slight decrease in rate of error over time is due to the gain of the PMT's going down (will later be shown in figure 5.2), leading to less overflow errors. Another factor is the lower rates of the sources leading to less pileup errors. Unfortunately channel 2 and 3 show some strange peaks, the behaviour of these channels will be discussed later
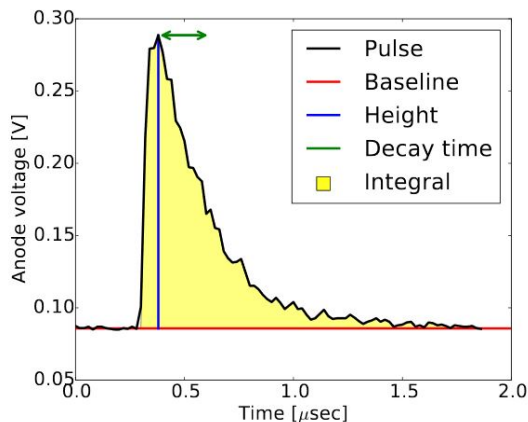
Figure 4.1: An event as recorded by the PMT. In black the waveform. In red the baseline that determines the value for the integral in yellow. In blue the maximum height. In green the decay time of roughly 230 ns. From [21]

.

4.2 Calibration

After cutting out all bad waveforms, each waveform is integrated. This gives us a measure in Volt seconds (Vs) that is proportional to the energy. Filling a histogram of these Vs measurements gives us an energy spectrum. We will use the well know monochromatic γ-ray absorption lines in the integral spectra to calibrate the energy scale of the NaI(Tl) detectors.

To find these absorption peaks in the spectrum we first search for the highest bin in the histogram and fit a Gaussian function to it. If the source has only one absorption peak, like ^{54}Mn or ^{137}Cs, the spectrum can now be calibrated. If however we have three peaks, as for ^{60}Co or ^{44}Ti, we wish to find the other two. As we now have a ratio between the position of the first peak and its energy, we can use that ratio to estimate where the other peaks will be. We then define a fitting region and fit a Gaussian to the two remaining peaks. We calibrate with the following function:

$$E(\text{keV}) = c_1 I(\text{Vs}) + c_0, \tag{4.1}$$

but expect $c_0 \approx 0$ for ^{137}Cs and ^{54}Mn. In figure 4.4 an example of a calibrated ^{60}Co spectrum is shown. The two photopeaks at 1173 and 1333 keV are clearly visible. In the decay of ^{60}Co two γ-rays are emitted by the ^{60}Ni daughter within a ps. So if both gamma rays are absorbed by a single NaI detector these gamma rays cannot be resolved in time, resulting in a so called sumpeak at $1173 + 1333 = 2505$ keV, also visible in the spectrum. The spectrum between 400 and 1000 keV consists of Compton scatters of the photon leaving the crystal. The peak around 700 keV is the backscatter peak of the photon scattering in the opposite crystal and scattering back into this channel.

18

(a) An example of a waveform with an overflow error, the amount of counts is higher than the maximum value in the ADC

(b) An example of a baseline error, the tail from a previous event is visible just before the start of the event displayed here.

(c) An example of a double event, there are two waveforms in this event window

Figure 4.2: Examples of erroneous events recorded by the ADC

4.2.1 Background fitting

In channels were there is no source we calibrate to the ^{40}K background. This is a very stable background with a half life of $1.251(3) * 10^9$ yr, and is primordial in origin [22]. ^{40}K is not continuously produced due to solar influence or as part of a decay chain, but has existed since the formation of earth. The measured rate of ^{40}K relatively low, therefore the exponential background needs to be taken into account. We fit an exponential distribution to the whole spectrum, and subtract that. Then we find the highest peak, and fit to it, its energy should be the ^{40}K 1460 keV peak.

In the case of background measurements for Nikhef channel 6 and 7, which contained the ^{137}Cs source, this was not enough. As the gain for these detectors is so high that the spectrum ends between 1400 and 1500 keV, it was impossible to calibrate to the ^{40}K peak. In figure 4.5 the background spectrum at Nikhef with and without lead are shown. The most prominent peak is the ^{40}K at 1460 keV. At low energy there is a peak visible at 511 keV and around 530 keV. These small peaks were the only two peaks to calibrate on. The 511 keV peak is a positron annihilation peak and the ^{214}Bi+^{208}Tl peak has poorly defined energy, as this is dependent on the ratio between the two isotopes. The choice was made to calibrate to the positron annihilation peak. Because the ^{214}Bi+^{208}Tl peak

Figure 4.3: The total rate of errors over time per channel. Channel 2 and 3 contain ^{44}Ti, channel 4 and 5 ^{60}Co, and chanel 6 and 7 ^{137}Cs.


Figure 4.4: A calibrated spectrum of one hour of measurements of ^{60}Co. All three absorption peaks that we calibrate to are clearly visible.
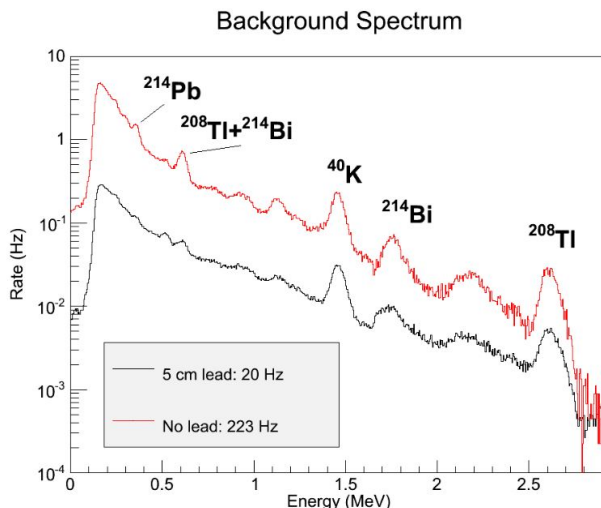
Figure 4.5: The background spectrum in the lab at Nikhef, from [16]. Although both with and without lead are shown, the measurements in this thesis always have a lead encasing.

is very close to this peak, it proved impossible to distinguish between them. In the end the fit range was narrowed to only allow for calibration to the 511 keV peak.

4.2.2 χ^2 values of the Gaussian fit

After we fit a Gaussian distribution to the photopeaks we can calculate the χ^2 value of each fit. As the tails of the Gaussian fit are of course different from the actual Compton background we expect a bad χ^2 value. More important is whether the quality of the fit (which is the same as the χ^2 value) remains stable over time. In figure 4.6 I have shown the χ^2 for all channels with sources for the 3 years analysed in our experiment. Although it seems steady, with a slow decline, due to the decrease in rate leading to a decrease in degrees of freedom, there are some problems. The main problem is that for some channels, namely 3, 4 and 5, there seem to be multiple bands of χ^2 values. This means that most of the time the fitting goes as expected, but a fraction of the time it goes wrong somewhere. A possible explanation for this phenomenon could be that the tails of the fitted Gaussians sometimes get in range of other peaks, completely throwing of the χ^2 value. In the end the χ^2 value is irrelevant as long as the absorption peak is found by the fitting algorithm. Even a slight deviation (± 50 keV) from the peak is corrected by the rate fitting procedure, where the peak energy is a free parameter with a small range around the expected energy.

Figure 4.6: The unreduced χ^2 values for the calibration procedure over time. Problematic are the bands of different values visible for channel 3,4 and 5.

4.3   Rate fit

Now that we have a calibrated energy spectrum per hour we can extract the rates. The rate is calculated by fitting a background template, simulated Compton scattering template and Gaussian photopeaks to the specrum with the Roofit package [23]. The Compton scattering template was made with GEANT, more can be read on that here [24], its rate is coupled to the rate of the sources. The background template was made from the continuous background measurements taken by channel 0 and 1.
After fitting each Gaussian is integrated and divided by the measuring time (usually one hour) to get the rate of the measured peak in Hertz. In figure 4.7 the spectrum with fitted distributions is shown, along with residuals.

4.3.1   Changing the background templates

The rate of background in the experiment is fixed but for statistical fluctuations. As the rate of background events is a few Hertz the statistical fluctuations per hour are around $1/\sqrt{3600} = 1.6\%$, which is small enough to fix the background rate. The sources of the Nikhef setup were removed in September 2018 for a period of four weeks to measure a background template for each channel individually. This way the fit could be done without having the rate of background as a free parameter. In figure 4.8 the resultant background template for two channels is shown. Channel 3 has an unexplained peak at roughly 750 keV, that is not measured in any other channel. Because of this the background templates have not yet been implemented in the rate fitting procedure.

22
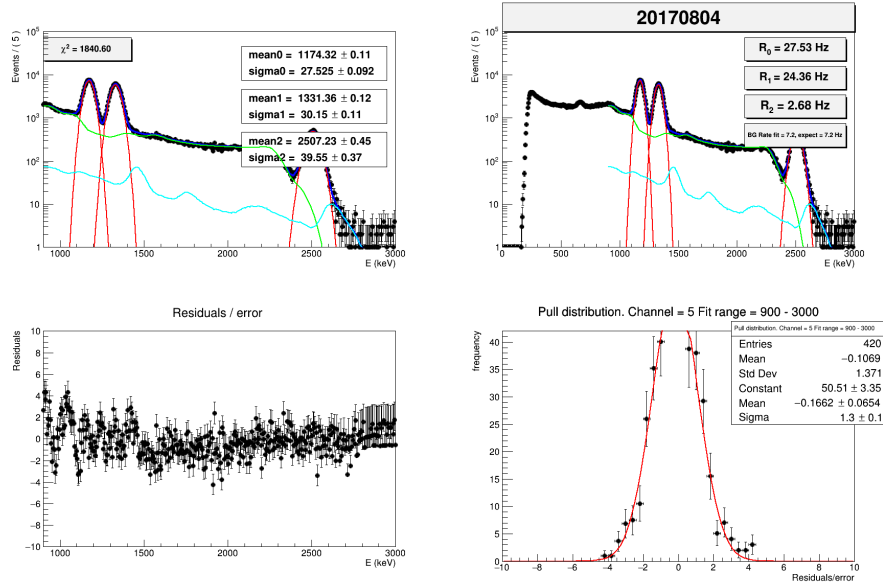
Figure 4.7: The results from the RooFit procedure. In the top two graphs the spectrum with fitted components is shown. The background template made with channel 0 and 1 is visible in light blue, the photopeaks in red and the simulated Compton background in green. This sums up to the blue spectrum that nicely matches the data. In the bottom graph the residuals over the error are shown, which are distributed very close to 0 indicating a good quality of the fit.

Figure 4.8: Two of the measured background spectra in Hz per bin, channel 3 on the left and channel 4 on the right. Clearly visible is the 1460 keV ^{40}K peak which is calibrated to. Channel 3 has an unexplained peak around 750 keV.

Chapter 5

Data Quality and slow control

The data used in this thesis were recorded between November 2015, when the Nikhef setup started running, and September 2018, when the setup was opened to take a background measurement. Before we can analyse the processed data, and determine whether or not we observe modulation of radioactive decay, we need to ensure that the data is trustworthy. This means checking the data, and all parameters generated in processing for anomalies. To be more precise, the following parameters and properties were studied:
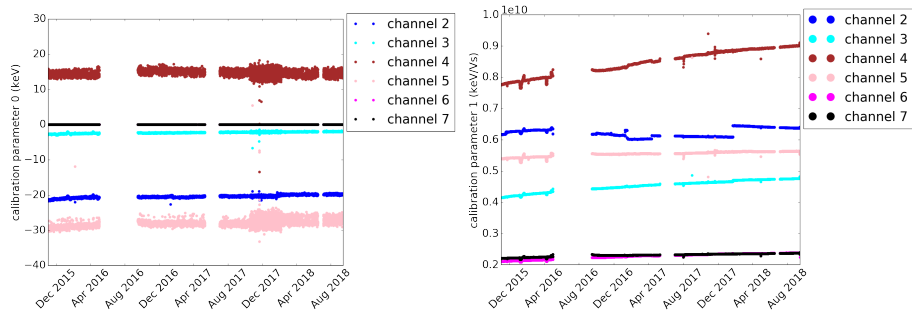
- the calibration parameters over time

- the rate of error over time

- the measured environmental parameters

5.1 The calibration parameters

Each measured waveform is integrated, yielding a result in Volt seconds (Vs). This value is subsequently converted into absorbed energy by calibration. The calibration procedure is explained in section 4.2. Important to remember is that we calibrate to the formula:

$$E = c_1 * I + c_0 \tag{5.1}$$

In figure 5.1 the calibration parameters c_1 and c_0 are plotted per channel over time. The general trend of the linear dependence parameter is upward. This can potentially be explained by the degradation of the PMT's. The gain of the PMT's decreases over time after roughly 1000 operating hours [25]. This leads to less strong responses to a signal, and thus a less strong Vs integral. This needs to be compensated by a higher c_1. In the channels with ^{137}Cs, there is only one peak to calibrate to. This means that the offset parameter c_0 is unnecessary in the fit, and is set to zero. Around January 2017 a jump in channel 2 is visible. At this date the power sockets in the lab had to be changed, and the whole

(a) The calibration parameter c_0 over time (b) The calibration parameter c_1 over time

Figure 5.1: The above figures show the calibration parameters over time. c_0 is the offset parameter in keV, whereas c_1 is the parameter measuring the linear relationship between integral in Vs and energy in keV

Table 5.1: The interpolated dependence for the environmental variables as estimated in the instrumentation paper [18] of the modulation experiment

Parameter	interpolated dependence
Pressure	$(0.7 \pm 1.1) \times 10^{-6} Pa^{-1}$
Magnetic field	$(4 \pm 3) \times 10^{-7} mG^{-1}$
Temperature	$(-4 \pm 2) \times 10^{-5} K^{-1}$
Radon activity	$(-0.1 \pm 2) \times 10^{-6} (Bq/m3)^{-1}$

experiment had to be shut down. This has inexplicably lead to a change in only channel 2. No settings were changed, and no other channels exhibit this behaviour.

In figure 5.2 I have shown the position of the first peak per channel over time as found by the calibration algorithm. This consists of taking the energy of the first peak and 'uncalibrating' that value back to Vs via the calibration parameters. This is both a check on the stability of our detector and on quality of the peak-finding algorithm.

Although the gain of the PMT's slightly decreases over time, this was expected. Other than a strange jump in January 2017 in channel 2, no large short term changes in the calibration parameters or first peak position are visible. This means that the detectors and the calibration procedure are stable enough to reliably produce the energy spectra.

5.2 Environmental effects

As stated in section 3.1 we monitor the conditions in the setup at a frequency of 1 Hz. We also actively keep the temperature at 29 ^{O}C with heaters. In the instrumentation paper for the modulation experiment [18] an estimation was
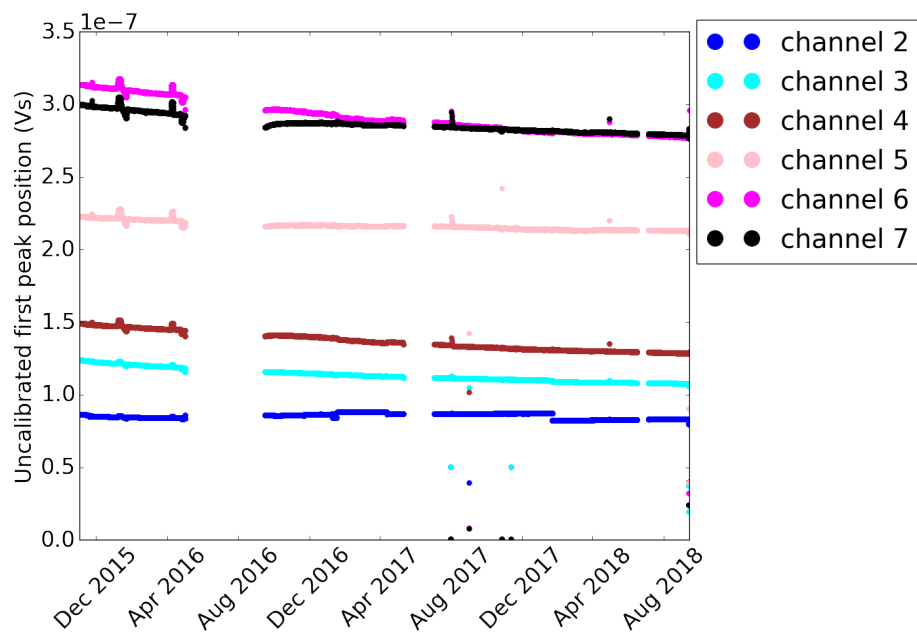
Figure 5.2: The position of the first peak of each channel, found by reversing the calibration of the data. This checks both the fitting algorithm and the stability of the detectors

Figure 5.3: The temperature over time, cut between 25 and 35 ° C. The two strong spikes at the beginning of 2016 are unfortunately unexplained.

made of the possible influence of each environmental parameter. These estimates are shown in table 5.1. Note that humidity is excluded in this, since it can only be relevant if the relative humidity becomes 100%, as that would result in condensation on the electric parts.

5.2.1   Temperature

In figure 5.3 the temperature over time per hour is shown. Fortunately no clear trends are visible, and the temperature is stable, staying mostly within 1 degree from the set temperature of 29 °C. The difference between equilibrium value before and after August 2016 is due to a new temperature control being installed. The two spikes in 2016 are probably due to a temperature calibration measurement.

5.2.2   Pressure

Although the setup is continuously flushed with between 5 and 10 liters of nitrogen per hour, there are pressure fluctuations. The pressure over time is shown in figure 5.4. These fluctuations are of course caused by changes in the weather, and should never go over 1.05 bar or below 0.98 bar. We have some severe outliers (infinite pressure), due to the barometer malfunctioning. We can calculate the variance of the values we trust though, and analyse the distribution in that way.

Figure 5.4: The pressure over time, cut between .96 and 1.05 bar. Notice that this fluctuates wildly. Even though our box is sealed, the pressure is clearly fluctuating with the weather. No further trends are recognisable

5.2.3 Magnetic field

In principle the magnetic field in the Netherlands should be around 0.5 G [26]. However, this is a value for outside, and due to the presence of steel in buildings this value can vary greatly, a change of 0.3 G due to a building being present is no exception [27]. In figure 5.5 the total magnetic field calculated by adding up the squares of all directional components is shown. There are some extreme outliers at 5 G which were cut away. This is probably a malfunction in the magnetometer, as there is no physical explanation for such values suddenly occurring. The changes in equilibrium value that can for example be seen around August 2018 can either be due to some changes in the setup of the lab, or more complications of the measuring instrument.

5.2.4 Humidity

Unfortunately the hygrometer was not fully calibrated. The humidity is measured relative to an unknown basis point. As visible in figure 5.6 most of the time the humidity stays between -20% and 20% of this value. Measurements with a separate hygrometer show the humidity of the lab as roughly 43%. This means that the humidity in the box is probably at 43% whenever the N_2 flow was off, and below 5% when it was on. This at least guarantees that there will be no condensation forming in the box.

5.2.5 Estimating the systematic error

Our environmental parameters can have some influence on the decay rate, which was estimated in the instrumentation paper [18]. We can calculate the variance

28

Figure 5.5: The total magnetic field over time, cut between 0 and 1 G. Note that the equilibrium value changes in December 2015 and August 2018


Figure 5.6: The relative humidity in % relative to zero point of the hygrometer, cut between -25 and 25 %. The extreme fluctuations to +20% are due to the box being opened, or the N_2 flushing being shut off.

Table 5.2: The measured variance of each environmental parameter and the effect of a variation in that parameter on the rate as estimated in [18]

Environmental parameter	variance	effect (as a fraction of the rate)
Pressure	1050 Pa	$7.3 \times 10^{-4} \mathrm{Pa}^{-1}$
Temperature	0.45 °C	$4.5 \times 10^{-6} {}^{\circ}C^{-1}$
Magnetic field	13.5 mG	$5.4 \times 10^{-6} \mathrm{mG}^{-1}$

of the environmental parameter distribution by calculating the variance of all acceptable measurements. This means we cut values like magnetic fields of 5 G or temperatures of 60 °C as this is clearly a measurement error. From the variance we then get the standard deviation by taking the square root. Multiplying the estimated effect of an environmental parameter by its standard deviation gives an estimate for the systematical uncertainty. These systematical uncertainties are reported in table 5.2.

Note that we have not taken into account what effect any correlation between two environmental parameters might have. As there is for example almost certainly some correlation between temperature and pressure. Any correlation would lead to lower estimates, so the estimates made here are a safe upper bound.

5.3 Bad runs

After analysing environmental parameters and the rates fitted to the spectra, a few problematic runs were identified and had to be cut. In the summer of 2016 the proportional–integral–derivative controller (PID) which controls the temperature in the box broke, and this lead to the temperature falling to 23 °C. In Joran Angevaare's thesis [21] an analysis was made of this failure, and it was concluded that the drop in temperature lead to an increase in rate due to thermal contraction of the frame holding the detectors. Although attempts at correction were made, I decided to ignore these bad months altogether. The reduced statistics due to ignoring two months are compensated by missing the systematic uncertainty introduced by correcting these rates.

In June 2017 the rate in channel 3 (containing ^{44}Ti) dropped significantly and abruptly as can be seen in figure 5.7. It is suspected that this is due to a bad connection either in the high voltage cables (leading to a lower voltage) or in the cables from the PMT to the DAQ computer (leading to loss of signal). The problem disappeared on its own and hasn't reappeared, which is why it was decided to ignore this month of data for all channels. If no cause is given, the run was cut due to bad unphysical rates. In appendix A one can find a table of all runs ignored in this analysis and if possible the reason why.

Figure 5.7: The residual rate for channel 3 during May and June 2017. A sudden drop in rate is visible, although it gets back to normal for a short while around June 8th. It drops again, but restores after this period and no further problems are had in channel 3 up to September 2018. Image from the modulation logbook

Chapter 6

Grid storage for the Raw data

Each modulation setup produces roughly 1 TiB (1024^4 bytes) of raw data per month, and roughly 70 GiB (70×1024^3 bytes) in processed data. Since this works out to roughly 51 TiB of data yearly for the modulation experiment, it was decided to store all data on a central storage at the SURFsara grid facility. The added benefit is that all participants in the experiment can then access all data. The following list of demands must be met by any copying procedure implemented:

- It must be fast enough, as we need to copy 1 TiB per month per setup

- It must be secure

- The data must be copied without errors.

To make copying easier we will make a tarball [28] of each dataset, which consist of data collected during one to seven days. To ensure data quality we will calculate a checksum at both ends and compare them. A checksum does a computation on a file, and returns a string. Since this string should be the same on both ends we can check if the files are the same on both ends the files are the same. After a connection with the storage is established the uberftp protocol [29] is used for access, and globus [30] is used for copying. Uberftp is not used for copying as it has a bug that might destroy data [29]. To get a secure connection we use a proxy and the uberftp protocol [29]. We now have 4 read actions per file, meaning that each file needs to be opened and read four times before the process is done. The Nikhef setup currently has roughly 40 TiB of data, meaning that at 100 MiB/s (roughly the read and internet speed), copying all data would take $4 \times 40 \times 1024^2/(100 \times 3600 \times 24) \approx 20$ days. Even 20 days is an acceptable amount of time, but the internet speed will reduce significantly when connecting from other setups. This makes it important to speed up the copying of data. To reduce this time all actions are done simultaneously through careful use of piping. The great thing is that these processes can work simultaneously, so we can tarball a folder and immediately get a checksum from it. If we introduce
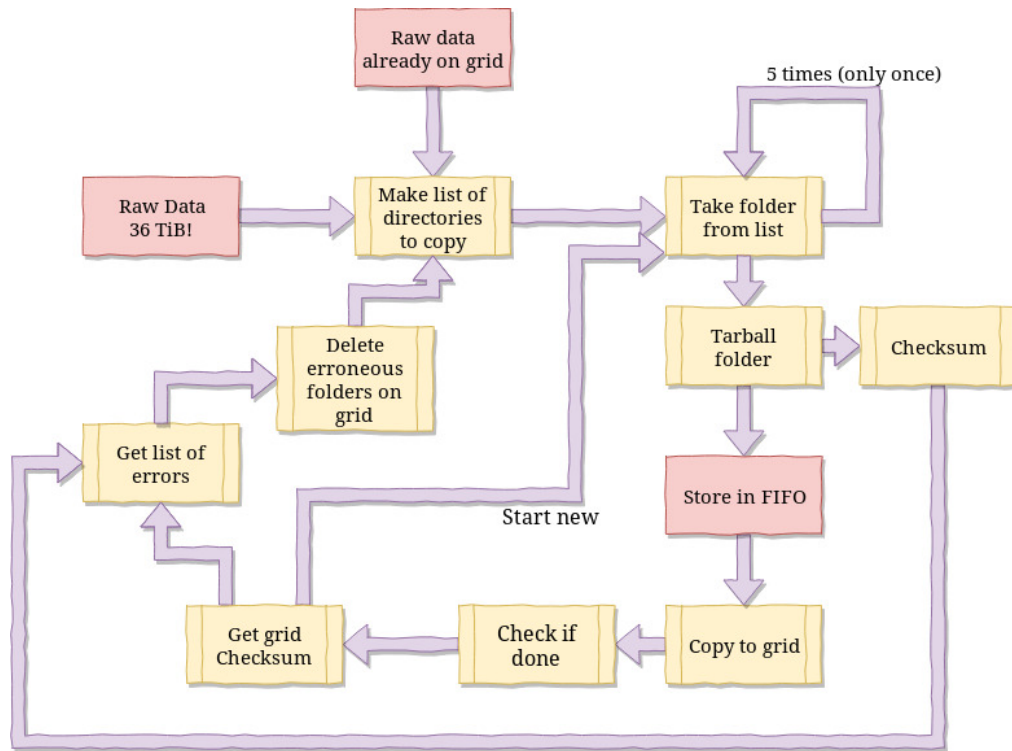
Figure 6.1: A flowchart of the copying procedure, processes are in yellow and storage places in red

a tee (essentially a split) in the pipe we can also copy at the same time. As every copying action is checked immediately, we know when there is a bug in the process and can stop the copying.

In figure 6.1 the process for the copying of data to the Surfsara grid is shown. First a list of local files and files on the grid is made, than a to-do list of files is made. Five copying processes, consisting of tarballing; calculating the checksum; copying; and getting the checksum on the grid, are then started. The amount of processes started can be changed, depending on the system the script is run on. Once one such copying process is finished a new one is started until the to-do list is empty. When this is done, it can be restarted to copy all files that went wrong in the previous copying job.

With the new script described above all raw data was copied from Nikhef to SURFsara. Instead of the approximately 20 days it would have taken when copying the files one by one, it took two days. This improvement in speed makes it possible for all setups to copy their data to SURFsara.

Chapter 7

Modulation analysis

The rates of radioactive decay have been extracted from the spectrum and we have sufficiently established the trustworthiness of our data. The next step is to investigate the claims of modulation in radioactive decay made by others. A maximum likelihood method is used to fit a function of radioactive decay with modulation to the data and measure the most likely amplitude, and set upper limits on the modulation amplitude. To analyse all possible frequencies a power spectrum analysis is then made.

7.1 Likelihood fits

From the hourly energy spectra of the radioactive decay the rate of the absorption peak is calculated by fitting and integration a Gaussian distribution. The next goal is to use this data of the rate over time and calculate the half life, and a possible modulation amplitude. This is done via a maximum likelihood fit. The principle of likelihood fits is as follows: calculate the probability of measuring the measured data given a function and try to maximise that probability for different function parameters. The parameters with the maximimal probability give the best possible fit. Of course this can be difficult to implement. Given a time between data points of δt and a function with parameters $f_{a,b,\dots}(t)$ the expected number of events for a datapoint is

$$\bar{N}_i = \int_{t-\frac{\delta t}{2}}^{t+\frac{\delta t}{2}} dt f_{a,b,\dots}(t).$$
(7.1)

This expected number is then the expectation value of a Poissonian distribution (as we are counting events), and we can calculate the probability of measurement i having value n:

$$P_i(n|\bar{N}) = \text{Poisson}(n|\bar{N}) = e^{-\bar{N}} \frac{\bar{N}^n}{n!}.$$
(7.2)

Note that we have now found a probability for a single measurement, and we need to do this for all measurements. This yields:

$$\mathcal{P}(a, b, ...) = \prod_i P_i(a, b, ...),$$
(7.3)

where $\mathcal{P}$ is the probability of the measurement. This product is difficult to numerically calculate. The likelihood often is so small that it is rounded to zero, and multiplying is a more intensive procedure than summation. This is why the negative logarithm of the likelihood is taken to get:

$$-\log(\mathcal{P}(a, b, ...)) = -\log(\prod_i P_i(a, b, ...)) = -\sum_i \log(P_i(a, b, ...)),$$
(7.4)

this gives us a sum of independent terms that we can minimise separately, and minimising the logarithm of the likelihood is the same as minimising the likelihood function.

The migrad algorithm will be used to determine the minimum of functions and the minos algorithm to determine the errors [31].

7.1.1   p-values

After fitting a function we wish to establish the quality of that fit, and the likelihood of our function given the data. For this we always use two hypotheses. H_0 which is the simple model and H_1 which is the model with the new physics. In our case H_0 is an exponential decay without modulation, and H_1 is an exponential decay with modulation.

As we have generally a high amount of events per bin (as our lowest rate is $\approx$ 5 Hz = 18000 counts per hour), the rate is approximately normally distributed, as is true for all Poisson distributions with a mean above 10. For normally distributed bins, a goodness of fit can be measured by calculating the χ^2 value:

$$\chi^2 = \sum_i \frac{(O_i - E_i)^2}{\sigma_i^2},$$
(7.5)

where O_i is the observed rate, E_i the expected rate, and σ_i the standard deviation of the observed rate. This measured χ^2 follows a χ^2-distribution with $k = N - c$ degrees of freedom, where N is the number of measurements and c is the number of fitted parameters. A lower χ^2 indicates a better fit, and we can calculate the probability of a worse fit occurring by integrating from the measured χ^2-value to infinity. This is the p-value, and a high p-value means that our fit is likely. A low p-value indicates an unlikely fit.

7.2   Half-life fit

To determine the half lives of the we fit the following function:

$$R(t) = N_0 e^{-\lambda t}.$$
(7.6)

Figure 7.1: The measured rate of the 667.1 keV ^{137}Cs absorption peak over a span of almost three years for both detectors together. Each data point is an hourly bin and the red dashed line is the half life fit.

Table 7.1: The fitted half lives compared to the literature values.

Source	literature half life (yr)	half life fit (yr)
^{44}Ti	59.1 ± 0.3	57.99 ± 0.067
^{60}Co	5.271 ± 0.004	5.2995 ± 0.00095
^{137}Cs	30.08 ± 0.09	30.242 ± 0.024

Where we give the literature value for the half life as initial guess for λ. In figure 7.1 I have shown the fit for ^{137}Cs. Note that the gaps are due to deleted datasets, as explained in section 5.3. It is possible to fit each energy separately, or fit for the combined rate of all energies together. The half life fits for the other sources are shown in appendix B, and the half lives for the total rate can be seen in table 7.1. Note that ^{137}Cs is the only half life within 2 standard deviations of the literature value. ^{44}Ti is off by roughly 3.5 σ and ^{60}Co is the worst with a deviation of more than 7σ. This means we underestimate the systematic uncertainty.

7.3 Modulation fit

We can now fit a modulation function in the same way as the half life fit from the previous section. We fit the function:

$$R(t) = N_0(1 + a\cos(\frac{f}{2\pi}(t - \phi)))e^{-\lambda t}, \tag{7.7}$$

where a is the modulation amplitude, f the frequency in year^{-1}, and ϕ the phase shift. In figure 7.2 the residual rate of ^{137}Cs is shown and compared with the expected modulation signal, the fit for the other sources is shown in appendix B. It becomes clear that our measurements are in disagreement with the claims of yearly modulation. In table 7.2 the best fit values for amplitude and phase are shown per source. For ^{44}Ti the fit is done per channel, because otherwise iminuit did not succeed in finding a minimum. All fitted amplitudes are within 2 σ of no modulation, meaning that although we do not find zero modulation, we do not find any significant modulation either. The fitted amplitudes are also much smaller than the amplitudes found by other researchers. Note that the phase is very uncertain and in all cases far away from 0, meaning that the explanation of modulating radioactive decays being caused by the sun-earth distance is very unlikely.

7.4 Modulation limits

Now that we have a fitting procedure, we can calculate the p-values of the fits with and without modulation, and hopefully set limits on modulation amplitudes. In table 7.3 the p-values for the half life fit and modulation fit are shown. channel 2, 3, 4, and 5 are separate, as the quality of fit for the channels

Figure 7.2: The modulation fit for ^{137}Cs with weekly bins. The green band signifies the expected rate from modulation claims from others, the red dotted line is our modulation fit. The colour on the datapoints is the temperature of the setup, as shown by the scale on the right.

Table 7.2: The fitted modulation functions for all sources. ^{44}Ti is split per channel because no minimum in the likelihood could otherwise be found.

Source	half life fit (yr)	amplitude fit (%)	phase fit (radians)
^{44}Ti channel 2	57.718 ± 0.096	-0.0027 ± 0.0025	-1.1 ± 0.74
^{44}Ti channel 3	58.272 ± 0.097	0.0036 ± 0.0025	-1.1 ± 0.54
^{60}Co	5.2991 ± 0.0011	-0.0022 ± 0.0023	-0.81 ± 0.87
^{137}Cs	30.253 ± 0.027	0.0022 ± 0.0025	-1.3 ± 0.95

38

Table 7.3: Calculated p-values for the fits with and without modulation. The rightmost column shows the amplitude with 95% CL. It is not available (n.a.) for 2 channels, as the fit without modulation is already below 5% probability.

Source	p-value no modulation	p-value best modulation fit	amplitude at p=0.05
^{44}Ti channel 2	0.1165	0.1099	1.05×10^{-4}
^{44}Ti channel 3	0.016	0.0139	n.a.
^{60}Co channel 4	0.000297	0.000272	n.a.
^{60}Co channel 5	0.093	0.0879	2.14×10^{-4}
^{137}Cs	0.75691	0.7502	2.57×10^{-4}

together was very low. Note that the p-value for no modulation is better than with modulation. This should not happen as it means that the optimal modulation fit would be an even smaller modulation, instead iminuit finds a local minimum. This implies that the optimal modulation fit has an amplitude so close to zero that the algorithm struggles to find it. In the rightmost column of table 7.3 the amplitudes of a 95% certainty limit (CL) are shown. All amplitudes above this are less likely than 5%. All limits are on the order of 10^{-4}, meaning that we can exclude the claims by others of order 10^{-3} with a good certainty.

7.5 Spectral analysis

In several papers [11][10] a power spectrum analysis of nuclear decay data is done to show evidence of several modulation frequencies. Most of these papers claim to show an annual modulation, although frequencies of one period per lunar month are also claimed. In order to study whether the Nikhef data contained this modulation I first fitted and subtracted an exponential decay, and then calculated the residual rate with respect to the expected rate. This yields the residual rate as a fraction. A periodogram is then made using the residual rates with the scipy periodogram function. This consists of first calculating the autocorrelation function:

$$R_{ff}(t) = \int_{-\infty}^{\infty} f(t+\tau)f(t)d\tau, \tag{7.8}$$

and then transforming the function to frequency space via a discrete Fourier transform (DFT):

$$f_k(t) = \sum_{n=0}^{N-1} x_n e^{-\frac{\pi i}{N}kn} \tag{7.9}$$

By autocorrelating the data and transforming it into frequency space any present modulation frequencies show up as peaks. Two different Monte Carlo simulations were made: one with a modulation of 0.1% as claimed by [9], and one without. Both have a Poissonian noise applied to them which is $\frac{1}{\sqrt{3600*\text{rate}}}$

as we use hourly bins. The Monte Carlo simulations are shown in figure 7.3. The rate used is 180 Hz as this is a nice middle ground between the rates in our experiment of ^{60}Co at roughly 100 Hz, ^{137}Cs at roughly 180 Hz and ^{44}Ti at roughly 500 Hz. The periodograms calculated from the data can be seen in figures 7.4,7.5 and 7.6. These are the periodograms for the total rate. It becomes very clear that no yearly modulation is observed. All other frequency peaks seem to have roughly the same maximum, also indicating no modulation of any kind. This steady level is dependent on the noise, which is dependent on the rate. For higher rates we expect less noise. This is visible as the source with the highest rate (^{44}Ti) has a lower level of amplitudes in the power spectrum than the lowest rate source (^{60}Co).

The relation between the top of the measured frequency and the amplitude of the oscillation is $A = \sqrt{2\text{PSD}}$ where A is the amplitude and PSD the power spectrum density. Note that the PSD for the Monte Carlo without signal is just as large as the PSD for the data. This implies that there is almost no signal. The general shapes of the Monte-Carlo periodogram without modulation also matches that of the periodograms of the data. This spectral analysis shows that we do not see any hint of modulation in our data. This does not mean we can fully exclude it, but we can exclude it up to the level of our noise.

(a)   (b)

Figure 7.3: The Power spectrum of the simulated data. Purple is data with modulation, whereas red is data without modulation. On the left we see data with a monthly modulation, on the right a simulation with a yearly modulation. The real data clearly resembles the spectrum without modulation.

Figure 7.4: The Power spectrum of the ^{44}Ti data. On the left we see no evidence of yearly or monthly modulation, whereas the right plot shows frequencies around daily modulation, also without evidence of modulation

Figure 7.5: The Power spectrum of the ^{60}Co data. On the left we see no evidence of yearly or monthly modulation, whereas the right plot shows frequencies around daily modulation, also without evidence of modulation

(a)   (b)

Figure 7.6: The Power spectrum of the ^{137}Cs data. On the left we see no evidence of yearly or monthly modulation, whereas the right plot shows frequencies around daily modulation, also without evidence of modulation

Chapter 8

Discussion

In the previous chapter limits it was concluded that the Nikhef setup of the modulation experiment measures no evidence of modulation in radioactive decays of any frequency. This chapter will discuss the trustworthiness of those results, and what might be done in the future to improve the results from the previous chapter.

8.1   Half lives

One of the main problems are the measured half lives. ^{60}Co and ^{44}Ti of them are more than three standard deviations away from the literature value. The deviation in the half life of ^{44}Ti is not unexpected, as it has a long half life while we have been measuring for only 3 years. The literature half life of ^{44}Ti is also relatively uncertain, as it is currently reported to be 59.1 years [6], but was reported to be 60.0 years just 10 years before that [32]. ^{60}Co on the other hand has been studied extensively, and has a very well-known half life. It is problematic that even when we include the uncertainty caused by environmental fluctuations in our analysis we do net get within 3 standard deviations of that half life. When investigating potential deviations from radioactive decay as we currently understand it, specifically when trying to disprove claims, we want to be able to report a half life in accordance with literature values.
We hope to improve the half life values in the future by fitting a channel specific background template when extracting the rates from the energy spectrum. It was the goal of this master project to accomplish this, but unfortunately implementing this in RooFit proved more difficult than expected.

8.2   Systematic uncertainties

The uncertainty on each rate measurement was adjusted by adding the standard deviation in the environmental parameters multiplied with the estimated effect on the rate. This is a very general approach, where all data points get the same systematic uncertainty added. A different approach would have tried to find a correlation between an environmental parameter and the rate, and corrected the rate per data point for that influence if it was found.
I decided against this as this possibly also corrects away a modulation signal or induces one. If there is a periodic fluctuation in an environmental parameter correcting the rate takes away sensitivity for modulations of that frequency. Even though the environment could cause this modulation, this is not certain. This is why the same systematic uncertainty was added for each data point.

8.3   Modulation fit procedure

When fitting the modulation function:

$$R(t) = N_0(1 + a\cos(\frac{f}{2\pi}(t - \phi))e^{-\lambda t}, \tag{8.1}$$

iminuit fails to maximise the likelihood. Although the algorithm converges, and produces fitted parameters with uncertainties, it is certain that this is a local maximum. When an exponential decay without modulation is fitted, and the χ^2 is calculated and divided by the degrees of freedom to produce the reduced χ^2, this reduced χ^2 is lower than the reduced χ^2 of the fit with modulation. In theory this is impossible, because if $a = 0$ in the modulation fit the function becomes an exponential decay.
Local maxima are found when maxima are close together in parameter space, and the values of the maxima are similar. This implies that the best fit amplitude lies somewhere between 0 and the amplitude reported in the previous chapter. The problem is to find this best possible fit. A solution is carefully tweaking the input parameters for the migrad algorithm until the actual maximum is found. It is however hard to guarantee that any found fit is the best possible fit. Another method would be downsampling the data by averaging the data over days or weeks. A reduced number of data points will simplify fitting a function, and will increase the likelihood of finding the best fit parameters. The downside is that while averaging data points might work for a yearly modulation search, it is not useful in searches for higher modulation frequencies ($f \geq 12$). Another option would be abandoning the likelihood fit in favour for the power spectrum analysis. A power spectrum analysis uses all data points, and will find all frequencies components in the data. The downsides are that it is harder to estimate the modulation amplitude from the power spectrum and p-values and exclusion limits are harder to calculate.

Chapter 9

Conclusion and outlook

The experiment described in this thesis aims to study radioactive decays more closely than has ever been done before. The goal of the experiment is to look for any variation on the expected exponential decay with Poissonian fluctuations.
The motivation for this experiment comes from multiple claims made by other experience, claiming a modulation on top of the exponential decay.
This thesis extensively analysed the stability of the modulation setup. The environmental parameters were all measured and analysed, adding the uncertainty caused by them to the rate measurements.
Then the half lives were fitted, yielding results that are significantly different from the literature values. This remains an unsolved problem in this thesis, although a better background fit might improve these half lives. A modulation function is then fitted, with amplitudes around 10^{-5}, much lower than the claimed 10^{-3}. A 95% certainty exclusion is then made for all sources at roughly 10^{-4}. Finally a Power spectrum analysis yields no evidence for any frequencies from daily to yearly.
Further analysis on the modulation data will need to improve the measured half lives, hopefully by using a different background template per channel in the rate fitting. The modulation fit can also be improved, as it currently finds a local minimum. In order to find amplitude exclusion limits for all frequencies, it would be better to abandon the likelihood analysis in favour a a power spectrum analysis. A power spectrum analysis does not have the risk of finding local maxima, and will be able to exclude or confirm modulation of any frequency.
Looking further ahead, ^{44}Ti is an isotope that has a uncertain half life. It is growing in importance as it was recently found to be important in supernovae [33]. The modulation experiment might be able to publish the best measured half life of ^{44}Ti.

9.1   Acknowledgements

I would like to thank all members of the Dark matter group at Nikhef for the wonderful year I had during my master thesis. Thank you Auke-Pieter, for being one of the best supervisors anyone could wish for. Even though you were often away, you somehow always found time to help with the pressing questions. Thank you to the PhD students of the dark matter group for your advice, especially Joran and Sander, both modulation experiment alumni. Peter also deserves a special mention for his enthusiasm, especially in the mornings, that seemed to spread to everyone. Without Roel Aaij, I would never have got this far with grid computing, his help and patience were invaluable in learning programming.

Lastly I would like to thank my girlfriend, family and friends for supporting me throughout the year, and always listening when I tried to explain my project, even though they might not have understood it.

Bibliography

[1] The nobel prize in physics 1903. `https://www.nobelprize.org/prizes/physics/1903/summary/`. accessed on 2019-06-25.

[2] Aidan Randle-Conde. Feynman diagram maker. `https://www.aidansean.com/feynman/`, 2019. Accessed: 2019-06-24.

[3] J. Tuli E. Brown. Nuclear data sheets for a=137. *Nuclear data sheets*, 2007.

[4] H. Junde Y. Dong. Nuclear data sheets for a=54. *Nuclear data sheets*, 2014.

[5] J. Tuli. Nuclear data sheets for a=60. *Nuclear data sheets*, 2003.

[6] J. A. Cameron J. Chen, B. Singh. Nuclear data sheets for a=44. *Nuclear data sheets*, 2011.

[7] Katharine Williams Yun Ping Wang. Radioactive decay rates. `https://chem.libretexts.org/Bookshelves/Physical_and_Theoretical_Chemistry_Textbook_Maps/Supplemental_Modules_(Physical_and_Theoretical_Chemistry)/Nuclear_Chemistry/Nuclear_Kinetics/Radioactive_Decay_Rates`, 2019. Accessed: 2019-06-24.

[8] B. Wang, S. Yan, B. Limata, F. Raiola, M. Aliotta, H. W. Becker, J. Cruz, N. De Cesare, A. Donofrio, Z. Fülöp, and et al. Change of the 7be electron capture half-life in metallic environments. *The European Physical Journal A*, 28(3):375–377, Jul 2006.

[9] Jere H. Jenkins, Ephraim Fischbach, John B. Buncher, John T. Gruenwald, Dennis E. Krause, and Joshua J. Mattes. Evidence of correlations between nuclear decay rates and earth–sun distance. *Astroparticle Physics*, 32(1):42–46, 2009.

[10] A.G.Parkhomov. Researches of alpha and beta radioactivity at long-term observations. April 2010.

[11] J.h. Jenkins, E. Fischbach, D. Javorsek, R.h. Lee, and P.a. Sturrock. Concerning the time dependence of the decay rate of 137cs. *Applied Radiation and Isotopes*, 74:50–55, 2013.

[12] Jere H. Jenkins and Ephraim Fischbach. Perturbation of nuclear decay rates during the solar flare of 2006 december 13. *Astroparticle Physics*, 31(6):407–411, 2009.

[13] Wang. *CRC Handbook of radioactive nuclides*. The Chemical Rubber Co., 1969.

[14] Ernest M. Henley and Alejandro Garcia. *Subatomic physics*. World Scientific, 2012.

[15] Stefano Meroli. The interaction of photons with the matter.

[16] P.A. Breur. The performance of nai(tl) scintillation detectors. 2013.

[17] Homepage scionix. `www.scionix.nl`. accessed: 2019-06-24.

[18] J.r. Angevaare, P. Barrow, L. Baudis, P.a. Breur, A. Brown, A.p. Colijn, G. Cox, M. Gienal, F. Gjaltema, A. Helmling-Cornell, and et al. A precision experiment to investigate long-lived radioactive decays. *Journal of Instrumentation*, 13(07), 2018.

[19] F. Gjaltema. Experiment to study modulations inradioactive decay. 2015.

[20] properties and use of scintillation crystals. `https://scionix.nl/scintillation-crystals/#tab-id-4`. accessed: 2019-06-24.

[21] J.R. Angevaare. Limiting modulations. 2017.

[22] Potassium 40, a curiosity of nature and a very long lived beta emitter. `http://www.radioactivity.eu.com/site/pages/Potassium_40.htm`. Accessed: 2019-06-24.

[23] W. Verkerke and D. Kirkby. The roofit toolkit for data modeling. *Statistical Problems in Particle Physics, Astrophysics and Cosmology*, Jun 2003.

[24] Cassie A. Reuter. *An experiment to search for systematic effects in long-lived radioactive decays*. 2017.

[25] Hamamatsu photonics. *Photomultiplier tubes: Basics and Applications*. 2007.

[26] Find the magnetic declination at your location. `http://www.magnetic-declination.com/`. accessed: 2019-06-24.

[27] E. Čermáková. Magnetization of steel building materials and structures in the natural geomagnetic field. *Acta Polytechnica Vol. 45*, 2005.

[28] Gnu tar. `https://github.com/Distrotech/tar`. Accessed: 2019-06-24.

[29] Grid storage. `http://docs.surfsaralabs.nl/projects/grid/en/latest/Pages/Advanced/grid_storage.html`. Accessed: 2019-06-24.

[30] Ian Foster. Globus toolkit version 4: Software for service-oriented systems. *Journal of Computer Science and Technology*, 21(4):513–520, 2006.

[31] F. James. Function minimization and error analysis version 94.1. *CERN Program Library Long Writeup D506*, 1998.

[32] S. Kumagai Y. Motizuki. Radioactivity of the key isotope 44ti in sn 1987a. *AIP Conference Proceedings*, 2004.

[33] Georgios Magkotsios, Francis X. Timmes, Wiescher Michael, Christopher L. Fryer, Aimee Hungerford, Patrick Young, Michael E. Bennett, Steven Diehl, Falk Herwig, Raphael Hirschi, and et al. Ti44 and ni56 in core-collapse supernovae. *Proceedings of 10th Symposium on Nuclei in the Cosmos — PoS(NIC X)*, 2009.

Appendix A

Appendix A

Table A.1: All runs removed from this analysis starting June 2017. Runs ignored before that time can be found in Joran Angevaare's thesis [21]

| Run	Reason
ANA_mx_n_20170629_0710	unknown
ANA_mx_n_20170629_1140	Unknown
ANA_mx_n_20171017_1144	Network errors
ANA_mx_n_20171027_1201	Network errors
ANA_mx_n_20180115_1258	Error 1
ANA_mx_n_20180116_1142	Weird errors, vangnet implemented
ANA_mx_n_20180117_0835	Weird errors, vangnet implemented
ANA_mx_n_20180119_1215	Unknown
ANA_mx_n_20180122_1503	Unknown, probably corrupted possibly radon
ANA_mx_n_20180125_1018	Unknown
ANA_mx_n_20180126_1218	Unknown
ANA_mx_n_20180219_0914	Unknown
ANA_mx_n_20180220_1148	Unknown
ANA_mx_n_20180221_1042	Unknown
ANA_mx_n_20180222_1031	Unknown
ANA_mx_n_20180913_0747	Sources removed
ANA_mx_n_20180913_1059	Sources removed
ANA_mx_n_20180914_0704	Sources removed

Appendix B

Appendix B

Figure B.1: The half life measurement for ^{44}Ti for all peaks combined

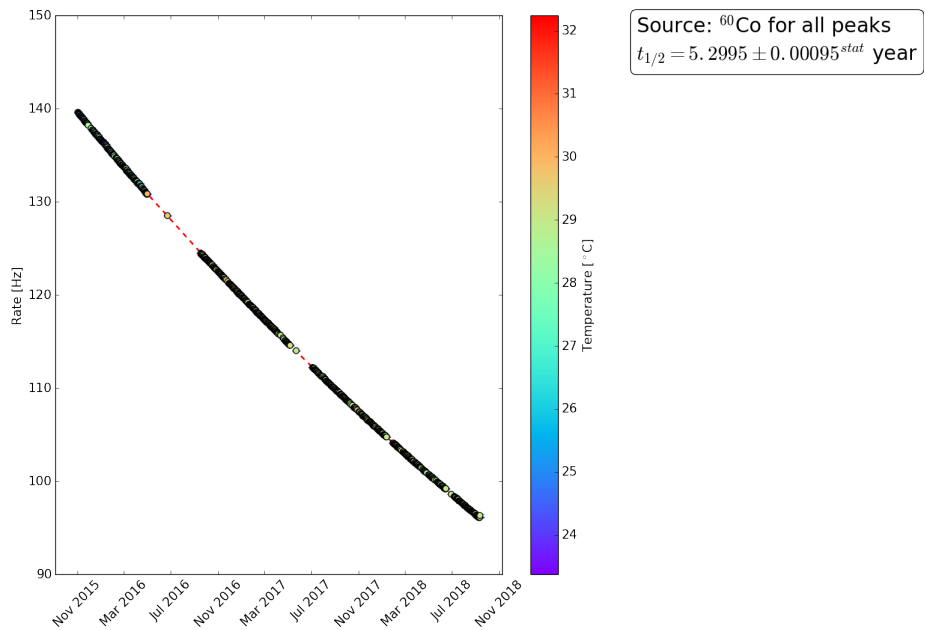

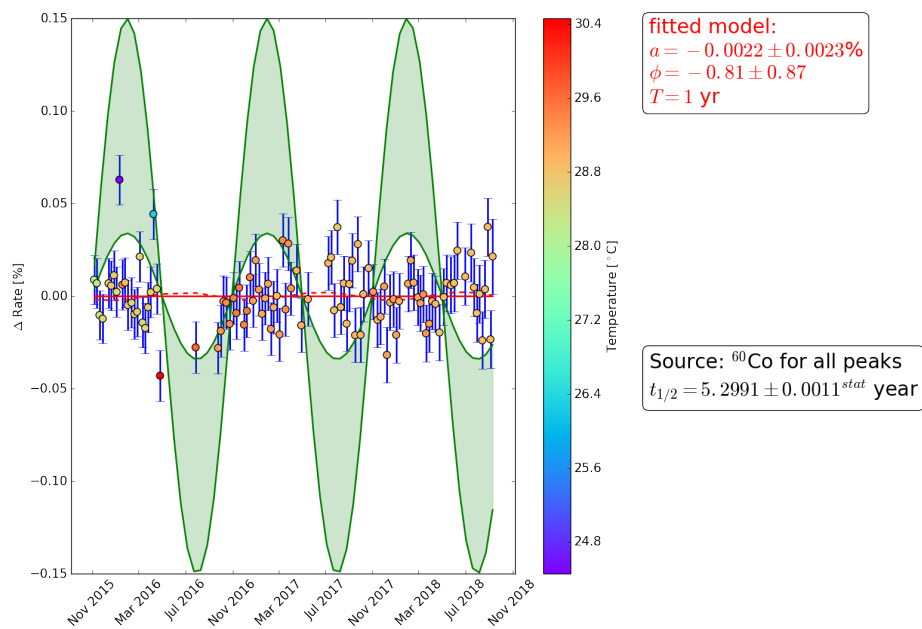Figure B.2: The half life measurement for ^{60}Co for all peaks combined

Figure B.3: The modulation fit for ^{60}Co with weekly bins. The green band signifies the expected rate from modulation claims from others, the red dotted line is our modulation fit. The colour on the data points is the temperature of the measurements, as shown by the scale on the right.
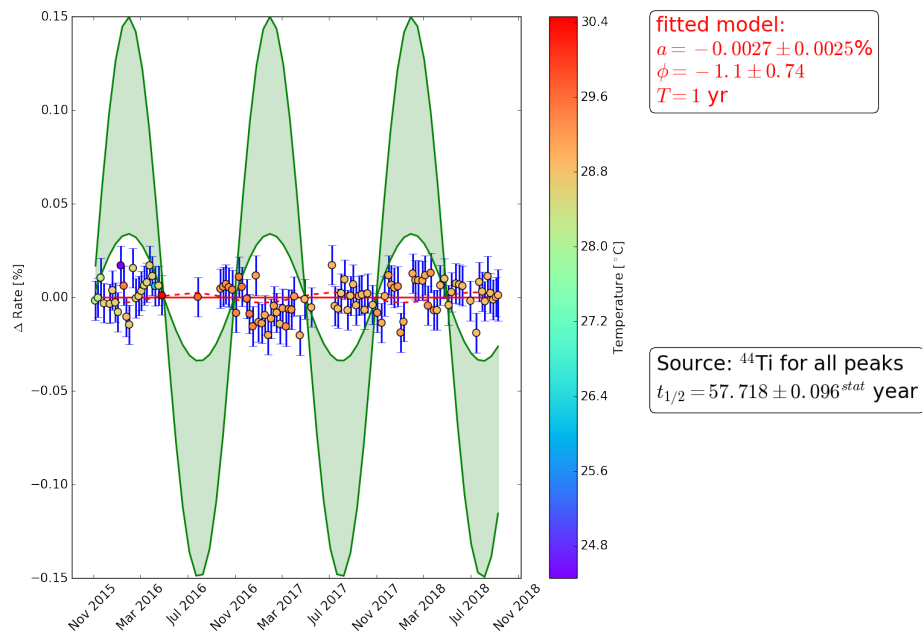
Figure B.4: The modulation fit for channel 2 containing ^{44}Ti with weekly bins. The green band signifies the expected rate from modulation claims from others, the red dotted line is our modulation fit. The colour on the data points is the temperature of the measurements, as shown by the scale on the right.

Figure B.5: The modulation fit for channel 3 containing ^{44}Ti with weekly bins. The green band signifies the expected rate from modulation claims from others, the red dotted line is our modulation fit. The colour on the data points is the temperature of the measurements, as shown by the scale on the right.